Contents

Windows Server 2012 R2 and Windows Server 2012 371

What's New in Windows Server 372

What's New in Windows Server 2012 R2 372

What's New in Windows Server 2012 375

Technical Scenarios for Windows Server 2012 R2 and Windows Server 2012 377
- Windows Server 2012 R2 377
- Windows Server 2012 378

Access and Information Protection 378
- See Also 379

Building Your Cloud Infrastructure: Scenario Overview 380
- Scenario description 380
- In this scenario 380
- Practical applications 380
- Roles and features included in this scenario 381
- Hardware requirements 382
- Software requirements 382
- See also 382

Designing Your Cloud Infrastructure 383
- Cloud Infrastructure Technical Overview 384
  - NIST Definition of Cloud Computing 384
    - Essential Characteristics: 384
    - Service Models: 385
    - Deployment Models: 385
  - Microsoft Private Cloud Overview 385
  - Windows Server 2012 Cloud Infrastructure for Hosting Environments 386
  - Private Cloud Architecture Principles 386
  - The Private Cloud Reference Model 387
  - Conceptual Architecture—Infrastructure 388
    - Scale Units 388
    - Storage 389
    - Networking 389
    - Virtualization Platform 390
- Cloud Infrastructure Design 391
  - Designing the Cloud Storage Infrastructure 391

Storage Options ... 391
Storage Protocols ... 392
Storage Network ... 392
Cluster Shared Volumes ... 393
CSV Requirements ... 394
CSV Volume Sizing ... 395
CSV Design Patterns ... 396
Single CSV per Cluster ... 396
Multiple CSVs per Cluster ... 396
Multiple I/O Optimized CSVs per Cluster ... 396
Storage Design ... 396
Performance ... 397
Drive Types ... 397
Multipathing ... 398
Fibre Channel SAN ... 398
iSCSI SAN ... 399
Storage Spaces ... 401
Data Deduplication ... 402
Thin Provisioning ... 403
Volume Cloning ... 404
Volume Snapshot ... 404
Storage Automation ... 404
Designing the Cloud Network Infrastructure ... 405
Network Infrastructure ... 406
Traffic Flow Isolation ... 407
Security Considerations ... 410
Secure Access Control ... 411
Traffic Control ... 412
Avoid Rogue IP Distribution ... 412
Scalability and Performance Considerations ... 413
Host and Virtualization Design ... 415
Windows Server 2012 Hyper-V Host Design ... 415
Licensing ... 416
Operating System Configuration ... 416
Memory and Hyper-V Dynamic Memory Options ... 416
Storage Adapters ... 417
MPIO Configuration ... 418
Performance Options ... 418
Network Adapter Teaming Configurations ... 419
Hyper-V Host Failover Cluster Design ... 420
Private Cloud Infrastructure without Failover Clustering ... 421
Host Failover Cluster Topology ... 422
Compute Cluster Traffic Profiles ... 422

Hyper-V Guest Virtual Machine Design 425
Virtual Machine Storage 425
Dynamically Expanding Disks 426
Fixed Size Disks 426
Differencing Disks 426
Pass-Through Disks 426
Support for Guest Clustering 427
In-Guest iSCSI Initiator 427
In-Guest Fibre Channel Synthetic HBA 427
Virtual Machine Network Interfaces 428
Virtual Processors 428
Overview of Suggested Cloud Infrastructure Deployment Scenarios 429
The Non-Converged Data Center Configuration 429
Networking 430
Storage 430
Compute 430
Overview of the Non-Converged Data Center Configuration 431
The Converged Data Center with File Server Storage Configuration 433
Design Considerations and Requirements for the Converged Data Center with File Server Storage Pattern 433
Networking 433
Storage 433
Compute 434
Overview of the Converged Data Center with File Server Storage Configuration 434
The Converged Data Center without Dedicated Storage Nodes Configuration 437
Design Considerations and Requirements for the Converged Data Center without Dedicated Storage Node Pattern 437
Networking 437
Storage 437
Compute 438

Building Your Cloud Infrastructure: Non-Converged Data Center Configuration 441
Design Considerations and Requirements for the Non-Converged Data Center Configuration Pattern 441
Networking 441
Storage 442
Compute 442
Overview 442
Non-Converged Data Center Configuration Scenario Overview 443
Install and configure 445
Step 1: Initial node configuration 446
1.1 Enable BIOS settings required for Hyper-V for SR-IOV 447
1.2 Perform a clean operating system installation 447

1.3 Perform post-installation tasks 447
1.4 Install roles and features using the default settings 447
Step 2: Initial network configuration 448
2.1 Disable unused and disconnected interfaces and rename active connections 448
Step 3: Initial storage configuration 449
3.1 Present all shared storage to relevant nodes 449
3.2 Install and configure MPIO as necessary for multipath scenarios 449
Step 4: Failover cluster setup 449
4.1 Run through the Cluster Validation Wizard 450
4.2 Address any indicated warnings and/or errors 450
4.3 Complete the Create Failover Cluster Wizard 450
4.4 Create the witness disk 451
4.5 Create the virtual machine storage disk 451
4.6 Add the virtual machine storage disk and Witness disk to Cluster Shared Volumes .. 451
4.7 Add folders to the cluster shared volume 452
4.8 Configure Quorum Settings 452
4.9 Configure cluster networks to prioritize traffic 452
Step 5: Configure Hyper-V settings 452
5.1 Create the Hyper-V virtual switch 453
5.2 Change default file locations for virtual machine files 453
Step 6: Cloud validation 454
6.1 Create a new virtual machine 454
6.2 Test network connectivity from the virtual machine 455
6.3 Perform a live migration 455
6.4 Perform a quick migration 455

Building Your Cloud Infrastructure: Converged Data Center with File Server Storage 456
Design Considerations and Requirements for the Converged Data Center with File Server Storage Pattern 456
Networking 456
Storage 457
Compute 457
Overview 458
Install and configure the Converged Data Center with File Server Storage cloud infrastructure 460
Step 1: Initial node configuration 464
1.1 Enable BIOS settings required for Hyper-V on the Nodes in the Hyper-V Cluster 464
1.2 Perform a clean operating system installation on all nodes in the Hyper-V and File Server Clusters 464
1.3 Perform post-installation tasks on all nodes in the Hyper-V and File Server Clusters 464
1.4 Install roles and features using the default settings on the Hyper-V Failover Cluster 465
1.5 Install roles and features using the default settings on the File Server Failover Cluster 466

Step 2: Initial network configuration 466
2.1 Disable unused and disconnected interfaces and rename active connections 467
2.2 Create the infrastructure and the tenant networks NIC teams on each member of the Hyper-V cluster 467
2.3 Create the infrastructure network NIC team on each member of the File Server cluster 468
2.4 Configure QoS settings for infrastructure traffic 469
Step 3: Initial storage configuration 470
3.1 Present all shared storage to relevant nodes 471
3.2 Install and configure MPIO as necessary for multipath scenarios 471
3.3 Wipe, bring online, and initialize all shared disks 471
Step 4: File server failover cluster setup 471
4.1 Run through the Cluster Validation Wizard 472
4.2 Address any indicated warnings and/or errors 472
4.3 Complete the Create Failover Cluster Wizard 472
4.4 Create a cluster storage pool 473
4.5 Create the quorum virtual disk 473
4.6 Create the virtual machine storage virtual disk 474
4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes 475
4.8 Add folders to the cluster shared volume 475
4.9 Configure Quorum Settings 476
4.10 Add the Scale-Out File Server for Applications Role 476
Step 5: Hyper-V Failover Cluster Setup 476
5.1 Run through the cluster validation wizard 477
5.2 Address any indicated warnings and/or errors 477
5.3 Complete the create cluster wizard 477
5.4 Verify cluster quorum configuration and modify as necessary 478
5.5 Configure Cluster Networks 478
Step 6: Configure Share and Hyper-V settings using a Script 478
6.1 Create Shares and Configure Hyper-V Settings using a Script 479
6.2 Configure Kerberos Constrained Delegation 479
Step 7: Cloud validation 479
7.1 Create the TenentNetSwitch 480
7.2 Create a new virtual machine 480
7.3 Test network connectivity from the virtual machine 481
7.4 Perform a live migration 481
7.5 Perform a quick migration 482

Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes 482
Design Considerations and Requirements for the Converged Data Center without Dedicated Storage Node Pattern 483
Networking 483
Storage 483

Compute 483
Overview 484
Install and configure the Converged Data Center without Dedicated Storage Server cloud infrastructure 487
Step 1: Initial node configuration 490
1.1 Add appropriate VLANS to the interface ports on the physical switch 491
1.2 Enable BIOS settings required for Hyper-V 491
1.3 Perform a clean operating system installation 491
1.4 Perform post-installation tasks 491
1.5 Install roles and features using the default settings 492
Step 2: Initial network configuration 493
2.1 Disable unused and disconnected interfaces and rename active connections 493
2.2 Create a converged network adapter team 493
2.3 Create the Hyper-V virtual switch and management virtual network adapter 494
2.4 Rename the management virtual network adapter (optional) 495
2.5 Create additional virtual network adapters and assign VLAN IDs 495
2.6 Rename virtual network adapters (optional) 495
2.7 Assign static IP addresses to the virtual network adapters 496
2.8 Configure QoS for different traffic types and configure the default minimum bandwidth for the switch 496
Step 3: Initial storage configuration 496
3.1 Present all shared storage to relevant nodes 497
3.2 Install and configure MPIO as necessary for multipath scenarios 497
3.3 Wipe, bring online, and initialize all shared disks 497
Step 4: Failover cluster setup 497
4.1 Run through the Cluster Validation Wizard 498
4.2 Address any indicated warnings and/or errors 498
4.3 Complete the Create Failover Cluster Wizard 498
4.4 Create a clustered storage pool 499
4.5 Create the quorum virtual disk 499
4.6 Create the virtual machine storage virtual disk 500
4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes 501
4.8 Add folders to the cluster shared volume 501
4.9 Configure Quorum Settings 502
4.10 Configure cluster networks to prioritize traffic 502
Step 5: Configure Hyper-V settings 502
5.1 Change default file locations for virtual machine files 502
Step 6: Cloud validation 503
6.1 Create a new virtual machine 503
6.2 Test network connectivity from the virtual machine 504
6.3 Perform a live migration 504
6.4 Perform a quick migration 505

Dynamic Access Control: Scenario Overview 505
In this scenario 505
Dynamic Access Control Content Roadmap 506
See also 514

Scenario: Central Access Policy 514
In this scenario 517
Roles and features included in this scenario 517

Plan for a Central Access Policy Deployment 518
Process to map a business request to a central access policy 519
Understand and translate business intent 519
Express access policy in Windows Server 2012 constructs 519
Determine the user groups, resource properties and claim types 519
Determine the servers where this policy should be applied to 520
Planning Guidelines for Deploying Central Access Policies 520
Using Security Groups for Dynamic Access Control 520
Using security groups to limit access to data 520
Using conditional expressions to reduce complexity of security groups 521
Using User Claims 521
Operations to enable user claims 522
Enable the domain controllers to provide claims and compound authentication on request 522
Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies 523
Using Device Claims and Device Security Groups 523
Considerations for using static device claims 523
Operations to enable device claims 523
Enable the Windows 8 devices in domain to request claims and compound authentication 523
Enable the Windows 8 devices to request claims and compound authentication using custom policy 524
Enable the Windows 8 device to receive compound authentication 524
Configuring central access policies with different options 525
Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs 525
Configuring forest root DCs 525
Configuring domains which provide claims and compound authentication 525
Configuring devices to request claims and compound authentication 526
Configuring resources to receive compound authentication 526
Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites 526
Configuring forest root DCs 526

Configuring domains which provides claims and compound authentication 527
Configuring file servers to request claims on the behalf of users 527
Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded 527
Considerations for using smartcards for Central Access Polices 527
Best Practices for Deploying Central Access Policies 528
Delegating of administration for Dynamic Access Control 528
Exception Mechanisms for Planning Central Access Policies 529
Tools for Deployment 530
Appendix: Deployment Configurations for Central Access Policies 531

Deploy a Central Access Policy (Demonstration Steps) 532
Set up a test environment 532
Plan: Identify the need for policy and the configuration required for deployment 533
Implement: Configure the components and policy 534
Deploy the central access policy 540
Maintain: Change and stage the policy 542
Next Steps 545

Deploy Claims Across Forests 545
Claim transformation rules 545
Linking claim transformation policies to forests 545
In this scenario 546
Roles and features included in this scenario 546

Deploy Claims Across Forests (Demonstration Steps) 547
Scenario overview 547
Set up the prerequisites and the test environment 547
Set up claims transformation on trusted forest (Adatum) 548
Create a claims transformation policy in Adatum 549
Set a claims transformation link on Adatum’s trust domain object 549
Set up claims transformation in the trusting forest (Contoso) 549
Create a claims transformation policy in Contoso 550
Set a claims transformation link on Contoso’s trust domain object 550
Validate the scenario 551
Additional scenarios for claims transformation policies 551
See also 552

Claims Transformation Rules Language 553
Tools for authoring claims transformation policies 553
Active Directory claims transformation rules language 554
Syntax overview 554
Runtime operation 554
Special rules semantics 556

Security considerations 556
Other language considerations 557
Sample transformation rules 558
Examples of rules parser errors 558
Language terminals 560
Language syntax 561

Scenario: File Access Auditing 563
In this scenario 564
Roles and features included in this scenario 565

Plan for File Access Auditing 565

Deploy Security Auditing with Central Audit Policies (Demonstration Steps) 567
Configure global object access policy 567
Update Group Policy settings 568
Verify that the global object access policy has been applied 569
See also 569

Scenario: Access-Denied Assistance 569
Scenario description 570
In this scenario 570
Practical applications 571
Features included in this scenario 571

Plan for Access-Denied Assistance 571
1.1 Determine the access-denied assistance model 572
1.2. Determine who should handle access requests 572
1.3. Customize the access-denied assistance message 572
1.4. Plan for exceptions 573
1.5. Determine how access-denied assistance is deployed 573
See also 573

Deploy Access-Denied Assistance (Demonstration Steps) 574
Step 1: Configure access-denied assistance 574
Step 2: Configure the email notification settings 578
Step 3: Verify that access-denied assistance is configured correctly 578
See also 579

Scenario: Classification-Based Encryption for Office Documents 579
Scenario description 579
In this scenario 580
Roles and features included in this scenario 580

Planning Considerations for Encryption of Office Documents 582
Determining files to automatically encrypt 582

Determining the rights policy template to use when encrypting files ... 584
Multi-machine considerations ... 585
Dynamic scope using the FolderUsage property ... 585
Setting management property values ... 586
Moving configurations between computers ... 586

Deploy Encryption of Office Files (Demonstration Steps) ... 587
Step 1: Enable resource properties ... 587
Step 2: Create classification rules ... 588
Step 3: Use file management tasks to automatically protect documents with AD RMS ... 591
Step 4: View the results ... 592
Step 5: Verify protection with AD RMS ... 593

Scenario: Get Insight into Your Data by Using Classification ... 593
Scenario description ... 593
In this scenario ... 594
Practical applications ... 594
Features included in this scenario ... 594

Plan for Automatic File Classification ... 594
1.1. Identify what information to classify in your environment ... 595
1.2. Identify how to classify files ... 595
1.3. Considerations for multiple computers ... 596
See also ... 596

Deploy Automatic File Classification (Demonstration Steps) ... 596
Step 1: Create resource property definitions ... 597
Step 2: Create a string content classification rule ... 597
Step 3: Create a regular expression content classification rule ... 599
Step 4: Verify that the files are classified correctly ... 600
See also ... 600

Set up Manual File Classification ... 600
Create Resource Properties ... 601
Set Group Policy Settings for Manual File Classification ... 601
Classify files and folders manually ... 602
Classification properties lists ... 602

Scenario: Implement Retention of Information on File Servers ... 603
Scenario description ... 603
In this scenario ... 603
Features included in this scenario ... 603

Plan for Retention of Information on File Servers ... 604
1.1. Determine the retention schedule ... 604

1.2. Identify files to be retained 605
1.3 Considerations for multiple computers 605
See also 605

Deploy Implementing Retention of Information on File Servers (Demonstration Steps) 605
Prerequisites 606
Step 1: Create resource property definitions 606
Step 2: Configure notifications 606
Step 3: Create a file management task 607
Step 4: Classify a file manually 609
See also 609

Appendix A: Dynamic Access Control Glossary 609
See Also 612

Appendix B: Setting Up the Test Environment 612
Prerequisites 612
Build the test lab virtual machines 613
Install the Hyper-V role 613
Create an internal virtual network 613
Build the domain controller 613
Build the file server and AD RMS server (FILE1) 616
Install File Services Resource Manager 616
Install the Microsoft Office Filter Packs on the file server 616
Configure email notifications on FILE1 617
Create groups on FILE1 617
Create files and folders on FILE1 617
Install Active Directory Rights Management Services 618
Build the mail server (SRV1) 623
Build the client virtual machine (CLIENT1) 623
Lab setup for deploying claims across forests scenario 624
Build a virtual machine for DC2 624
Set up a new forest called adatum.com 624
Set contoso.com as a trusting forest to adatum.com 625
Create additional users in the Adatum forest 626
Create the Company claim type on adataum.com 626
Enable the Company resource property on contoso.com 627
Enable Dynamic Access Control on adatum.com 627
Create the Company claim type on contoso.com 628
Create the central access rule 628
Create the central access policy 629
Publish the new policy through Group Policy 629
Create the Earnings folder on the file server 630
Set classification and apply the central access policy on the Earnings folder 630

Hosting-Friendly Web Server Platform (IIS): Scenario Overview 631
Scenario Description 631
Web Server Scenarios 631
Practical Applications 632
See Also 633

Build a Static Website on IIS 633
Prerequisites 634
Step 1: Install the IIS Web Server 634
Step 2: Add a Website 635
Step 3: Configure Anonymous Authentication 637
Step 4: Configure the Default Documents 638
Step 5: Configure Static Content Compression 638
Next Steps 639
See also 639

Configure Request Filtering in IIS 640
Prerequisites 640
General Request Filter Settings 640
File Name Extensions 642
Filtering Rules 643
Hidden Segments 644
URL Filtering 644
HTTP Verbs 645
Header Size Limits 646
Query Strings 646
Request Filter Logging 646
See Also 647

Configure Logging in IIS 647
Prerequisites 648
Configure Logging at the Site Level 648
Configure Per-site Logging at the Server Level 650
Configure Per-server Logging at the Server Level 650
Select W3C Fields to Log 650
Configure Log File Rollover Options 651
See Also 652

Build a Classic ASP Website on IIS 652
Prerequisites 653
Step 1: Install the IIS Web Server 653
Step 2: Add a Classic ASP Website 655
Step 3: Edit ASP Application Settings 656
Next Steps 663

See also 663

Build an ASP.NET Website on IIS 663
Scenario Description 663
In This Scenario 664
Practical Applications 664
Software Requirements 664
See Also 664

Plan an ASP.NET Website on IIS 665

Step 1: Plan IIS Web Server and ASP.NET Modules Installation 665
1.1. Plan to Install IIS and ASP.NET Modules 666
1.2. Plan to Add the ASP.NET Application 666
See Also 666

Step 2: Plan ASP.NET Settings 667
2.1. Session State Settings 667
Store session state in process 667
Store session state by using state server 668
Store session state by using SQL server 669
Cookie mode for session state 669
2.2. Pages and Controls Settings 671
2.3. Application Settings 671
2.4. .NET Compilation Settings 672
2.5. .NET Globalization Settings 672

Step 3: Plan Data Source Settings 673
3.1. Data source connection strings 673
3.2. ASP.NET providers 673
3.3. .NET profiles 674
3.4. .NET roles 675
3.5. .NET users 675

Step 4: Plan Application Security 675
4.1. Isolate Web Applications 676
4.2. .NET Trust Levels 676
4.3. .NET Authentication 677
ASP.NET Forms Authentication 677
Forms authentication basics 678
Authentication cookies 678
ASP.NET Impersonation Authentication 680
4.4. Machine Key Settings 681
4.5. TLS/SSL Communication 681
Server Certificates 682

SSL Binding 682
Require SSL for Your Site 683
Client Certificates 683

Configure an ASP.NET Website on IIS 683

Step 1: Install IIS and ASP.NET Modules 684
Installing IIS and ASP.NET Modules 684
Adding the ASP.NET Application 686
See Also 687

Step 2: Configure ASP.NET Settings 687
2.1. Session State Settings 688
Store Session State in Process 688
Store Session State by using State Server 689
Store Session State by using SQL Server 690
Cookie Mode for Session State 692
2.2. Pages and Controls Settings 694
Edit Pages and Controls 694
Add a Custom Control 694
2.3. Application Settings 695
2.4. .NET Compilation Settings 696
2.5. .NET Globalization Settings 698

Step 3: Configure Data Source Settings 700
3.1. Data Source Connection Strings 700
3.2. ASP.NET Providers 701
3.3. .NET Profiles 704
3.4. .NET Roles 706
3.5. .NET Users 707

Step 4: Configure Application Security 707
4.1. Isolate Web Applications 708
4.1. .NET Trust Levels 709
4.2. .NET Authentication 709
ASP.NET Forms Authentication 710
ASP.NET Impersonation Authentication 712
4.3. Machine Key Settings 714
4.4. TLS/SSL Communication 714
SSL Binding 714
Require SSL for Your Site 715
Client Certificates 716

Build an FTP Site on IIS 717
Prerequisites 717

Step 1: Install FTP on an Existing IIS Web Server 718
Step 2: Add an FTP Site 718
Step 3: Configure FTP Site Defaults 720
Step 4: Configure Firewall Support 721
Step 5: Configure User Isolation 722
Step 6: Configure Directory Browsing Options 723
Step 7: Configure Logon Attempt Restrictions 724
Step 8: Configure Request Filtering 724
Step 9: Configure FTP Logging 725
Step 10: Configure FTP Messages 725
See Also 726

Build a PHP Website on IIS 727
Scenario Description 727
In This Scenario 727
Practical Applications 727
Software Requirements 728
See Also 728

Plan a PHP Website on IIS 728

Step 1: Plan IIS Web Server and PHP Installation 729
1.1. Plan to Install IIS 729
1.2. Plan to Download and Install PHP 729
1.3. Plan to Add a PHP Application 730
See Also 730

Step 2: Plan PHP Settings 730
2.1. Plan WinCache Configuration 731
2.2. Plan Other PHP Settings 731
Required Settings 731
Optional Settings 732
2.3 Plan PHP Extensions 733
See Also 734

Step 3: Plan PHP Application Security 734
3.1. PHP Configuration Settings for Security 734
Disable File Handling for Remote URLs 734
Disable Register_Globals 735
Restrict File System Read/Write 735
Disable Safe Mode 735
Limit Script Execution Time 735
Limit Memory Usage and File Size 735
Configure Error Logging 736
Enable FastCGI Impersonation 736

Disable FastCGI Logging ... 736
Hide PHP Presence ... 736
3.2. Web Server and PHP Application Security ... 736
Isolate Web Applications ... 736
Enable Per-site PHP Configuration ... 737
Use Request Filtering ... 737
See Also ... 737

Configure a PHP Website on IIS ... 737

Step 1: Install IIS and PHP ... 738
1.1. Install IIS ... 738
1.2 Install PHP by using Web PI ... 739
1.3. Download and Install PHP Manually ... 740
1.4. Add Your PHP Application ... 741
See Also ... 742

Step 2: Configure PHP Settings ... 742
2.1. Configure WinCache ... 743
2.2. Configure Other PHP Settings ... 743
2.3 Configure PHP Extensions ... 744
See Also ... 744

Step 3: Configure PHP Application Security ... 744
3.1. Configure PHP Settings for Security ... 744
3.2. Configure Web Server and PHP Application Security ... 745
Isolate Web Applications ... 745
Enable Per-site PHP Configuration ... 746
Per-site PHP Process Pools ... 746
Specifying Php.ini Location ... 747
Use Request Filtering ... 748
See Also ... 748

Build a Web Farm with IIS Servers ... 749
Scenario Description ... 749
In This Scenario ... 749
Practical Applications ... 750
Software Requirements ... 750
See Also ... 750

Plan a Web Farm with IIS Servers ... 750

Step 1: Plan IIS Web Farm Infrastructure ... 751
1.1. Decide on Web Farm Infrastructure ... 751
Local Content Infrastructure ... 751

Shared Network Content Infrastructure 752
The Infrastructure Chosen for This Scenario 752
1.2. Use ARR for Load Balancing 753
1.3. Start with a Functional Website 753
See Also 753

Step 2: Plan IIS Web Farm Configuration 753
2.1. Plan for Shared Content 754
2.2. Plan for Shared Configuration 754
2.3 Plan to Add Web Servers 754
See Also 754

Step 3: Plan IIS Web Farm Load Balancing 755
3.1. Plan Load Balancing with ARR 755
3.2. Survey Other Features of ARR 755
See Also 757

Step 4: Plan SSL Central Certificate Store 758
4.1. Introduction to Centralized Certificates 758
4.2 Plan a Central Certificate Store 758
See Also 758

Step 5: Plan Application Deployment 759
5.1. Deploy Websites with FTP 759
5.2 Deploy Web Applications with Web Deploy 759
See Also 760

Configure a Web Farm with IIS Servers 760

Step 1: Install IIS Web Farm Infrastructure 760
1.1. Install IIS with Appropriate Modules 761
1.2 Install ARR for Load Balancing 762
1.3 Set up Your Website on One Web Server 763
See Also 763

Step 2: Configure IIS Web Farm Servers 763
2.1. Prepare Your Back-end File Server 763
2.2. Configure Shared Content 764
2.3. Set up Shared Configuration 765
2.4. Add Web Servers to Your Farm 765
See Also 766

Step 3: Configure IIS Web Farm Load Balancing 766
3.1. Create a Server Farm with ARR 767
3.2. Configure Load Balancing with ARR 767
3.3. Change Application Pool Settings 767

See Also 768

Step 4: Configure SSL Central Certificate Store 768
4.1. Configure a Central Certificate Store 768
See Also 769

Step 5: Configure Application Deployment 769
5.1. Install and Configure FTP for Your Web Farm 769
5.2. Install and Test Web Deploy for Your Web Farm 771
5.3. Where do I go from here? 772
See Also 773

Increasing Server, Storage, and Network Availability: Scenario Overview 773
In this scenario 774
See also 774

Deploying Fast and Efficient File Servers for Server Applications 775
Overview 775
Requirements and recommendations 775
SMB Direct 775
SMB Multichannel 776
Install the required roles, role services, and features 776
Step-by-step instructions 776
Step 1: Verify the basic network configuration 777
Step 2: Configure a failover cluster 777
Step 3: Configure the networks for the failover cluster 777
Step 4: Configure a Scale-out File Server 778
Step 5: Verify each file server name has two addresses 778
Step 6: Configure a Hyper-V or Microsoft SQL Server client 779
Step 7: Verify servers are using SMB Multichannel and SMB Direct 779
Step 8: Monitor file shares using Performance Counters 779
See also 780

Scale-Out File Server for Application Data Overview 780
Scenario description 780
In this scenario 781
When to use Scale-Out File Server 782
Practical applications 784
Features included in this scenario 784

Plan for Scale-Out File Server 784
See also 785

Step 1: Plan for Storage in Scale-Out File Server 785
Review Failover Clustering requirements 786

Review server application storage requirements ... 786
Review existing storage in your organization ... 786
See also ... 787

Step 2: Plan for Networking in Scale-Out File Server ... 787
2.1. Review your network adapter configurations ... 788
2.2. Review the CSV redirection traffic network configuration ... 788
2.3. Review the DNS configuration for the cluster nodes ... 788
See also ... 788

Deploy Scale-Out File Server ... 788

Step 1: Install Prerequisites for Scale-Out File Server ... 789
1.1. Install role services and features ... 789
1.2. Validate hardware and create a cluster ... 790
1.3. Add storage to a cluster shared volume ... 791
See also ... 791

Step 2: Configure Scale-Out File Server ... 792
Configure Scale-Out File Server ... 792
Create a continuous availability file share on the cluster shared volume ... 793
See also ... 794

Step 3: Configure Hyper-V to Use Scale-Out File Server ... 794
3.1. Verify permissions ... 794
3.2. Create a new virtual machine ... 795
See also ... 795

Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server ... 796
4.1. Verify permissions ... 796
4.2. Create the database files ... 796
See also ... 797

Network Performance and Availability ... 797
Scenario description ... 797
Practical applications ... 797
Features included in this experience ... 797
Hardware requirements ... 798
Software requirements ... 798

Deploy Hyper-V over SMB ... 798
Prerequisites ... 799
Step 1: Configuring file server clusters ... 800
Step 2: Install Hyper-V ... 802
Step 3: Create an SMB file share ... 802
Step 4: Create a virtual machine and virtual hard disk file on the file share ... 804

Step 5: Migrate virtual machine storage to an SMB file share 805
Step 6: Initiate a live migration of a virtual machine to another cluster node 805
Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage 806
Troubleshooting 807
See also 808

Protect Data on Remote SMB File Shares using VSS 809
VSS for SMB File Shares: overview 809
Requirements and supported configurations 809
Deployment scenarios 810
Step 1: Install File Server VSS Agent Service 812
Step 2: Add a user to the Backup Operators local group on the file server 813
Step 3: Perform a shadow copy 814
See also 817

Install and Deploy Windows Server 2012 R2 and Windows Server 2012 818
Windows Server 2012 R2 818
Windows Server 2012 818

System Requirements and Installation Information for Windows Server 2012 R2 819
Clean installation 820
Review system requirements 820
Processor 820
RAM 821
Disk space requirements 821
Other requirements 821
Review preinstallation documentation 822
Obtain the server product 822
Perform preinstallation tasks 822
Evaluation versions of Windows Server 2012 823
Limits of evaluation versions 823
Installing versions distributed as VHDs 823
Copyright 824

Upgrade Options for Windows Server 2012 R2 825
Upgrading previous retail versions of Windows Server to Windows Server 2012 R2 825
Per-server-role considerations for upgrading 826
Converting a current evaluation version to a current retail version 827
Converting a current retail version to a different current retail version 828
Converting a current volume-licensed version to a current retail version 828

Release Notes: Important Issues in Windows Server 2012 R2 828
Setup on virtual machines 829
Add/Remove Features wizard 829
Internet Explorer 11 829

Storage Spaces ... 829
Trusts ... 829
Windows Server Essentials Experience ... 830
Work Folders ... 830
Copyright ... 830

Features Removed or Deprecated in Windows Server 2012 R2 ... 831
Quick reference table ... 831
Features removed from Windows Server 2012 R2 ... 834
Backup and file recovery ... 834
Drivers ... 835
Recovery disk creation ... 835
Slmgr.vbs options ... 835
Subsystem for UNIX-based Applications ... 835
Windows Authorization Manager (AzMan) ... 835
WMI root\virtualization namespace v1 (used in Hyper-V) ... 835
Features deprecated starting with Windows Server 2012 R2 ... 835
Active Directory ... 836
Application Server ... 836
COM and Inetinfo interfaces of the Web Server role ... 836
DNS ... 836
File and storage services ... 836
IIS Manager 6.0 ... 837
Networking ... 837
Network Information Service (NIS) and Tools (in RSAT) ... 837
RSAT: Identity management for Unix/NIS ... 837
SMB ... 837
Telnet server ... 837
Windows Identity Foundation ... 838
SQL Lite ... 838
WMI providers and methods ... 838
Copyright ... 838

Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012 ... 838
Open the Start screen ... 839
Shut down or restart the computer ... 839
Lock the computer or sign out ... 840
Close a Windows app ... 840
Access Settings for the current screen ... 840
Access Control Panel ... 840
Access Administrative Tools ... 841
Create shortcuts ... 842
Open the Run dialog box ... 843

Run a program as administrator or as another user ... 843
Open Server Manager ... 844
Start Windows PowerShell ... 844
Open Remote Desktop Connection ... 844
Open Command Prompt ... 845
Open Microsoft Management Console (MMC) and snap-ins ... 845
Keyboard shortcuts ... 846
Use keyboard shortcuts in a Remote Desktop session ... 849
Use keyboard shortcuts in Hyper-V virtual machines ... 850

Installing Windows Server 2012 ... 851
Preinstallation information ... 851
System requirements ... 851
Processor ... 851
RAM ... 851
Disk space requirements ... 851
Other requirements ... 852
Important information for x64-based operating systems ... 852
Before you start Setup ... 852
Supported upgrade paths ... 853
Copyright ... 854

Release Notes: Important Issues in Windows Server 2012 ... 854
Upgrade ... 854
Copyright ... 855

Evaluation Versions and Upgrade Options for Windows Server 2012 ... 855
Evaluation versions of Windows Server 2012 ... 855
Where to find evaluation versions ... 856
Limits of evaluation versions ... 856
Converting evaluation versions of Windows Server 2012 to full retail versions ... 856
Upgrading previous retail versions of Windows Server to Windows Server 2012 ... 857
Per-server-role considerations for upgrading ... 858
Converting existing Windows Server 2012 versions ... 861

Windows Server Installation Options ... 862
Installation options description ... 862
If you choose the Server Core Installation option ... 863
If you choose the Server with a GUI option ... 864
Minimal Server Interface ... 864
Features on Demand ... 865
Practical applications ... 866
Reference table ... 867
See also ... 868

Server Core and Full Server Integration Overview ... 868
Requirements ... 869
Technical overview ... 869

Configure and Manage Server Core Installations ... 870

Deploy a Server Core Server ... 870
Initial installation ... 871
Using an unattend file to install the server directly in Server Core mode ... 871
See also ... 871

Configure a Server Core Server ... 872
1.1. Set the administrative password ... 872
1.2. Set a static IP address ... 873
1.3 Join a domain ... 874
1.4 Rename the server ... 874
1.5 Activate the server ... 874
1.6 Configure Windows Firewall ... 874
1.7. Enable Windows PowerShell remoting ... 875
See also ... 875

Configure a Server Core Server with Sconfig.cmd ... 875
Domain/Workgroup settings ... 876
Computer name settings ... 876
Local administrator settings ... 877
Network settings ... 877
Windows Update settings ... 877
Remote Desktop settings ... 877
Date and time settings ... 878
To enable remote management ... 878
To log off, restart, or shut down the server ... 878
To exit to the command line ... 878

Install Server Roles and Features on a Server Core Server ... 878
Installing and uninstalling server roles and features ... 880
Working with Features on Demand ... 880
See also ... 881

Manage a Server Core Server ... 881
1.1. Manage with Windows PowerShell ... 883
1.2 Manage with Server Manager ... 883
1.3 Manage with Microsoft Management Console ... 883
1.4 Manage with Remote Desktop Services ... 885
1.5 Switch to Server with a GUI mode ... 886
1.6 Add hardware and manage drivers locally ... 886

See also 887

Service Updates on a Server Core Server 888
1.1. Manage updates automatically with Windows Update 888
1.2. Manage updates with WSUS 889
1.3. Manage updates manually 889
See also 890

Quick Reference for Server Core Tasks 890
1.1. Configuration and installation 891
1.2. Networking and firewall 893
1.3. Updates and error reporting 894
1.4. Services, processes, and performance 895
1.5. Event logs 895
1.6. Disk and file system 895
1.7. Hardware 896
See also 896

Features Removed or Deprecated in Windows Server 2012 896
Features removed from Windows Server 2012 897
Active Directory Federation Services 897
Server Core components 897
Clustering 897
Graphics 897
Hyper-V 898
Networking 898
Server roles 898
Server Message Block 898
SQL Server 898
Storage 898
Visual Studio 898
Windows Help 898
Features deprecated starting with Windows Server 2012 899
Active Directory 899
Database management systems 899
Networking 899
Hyper-V 899
Printing 900
Remote Data Service 900
SMTP 900
Subsystem for UNIX-based Applications 900
Transport protocols 900
SNMP 900
SQL Server 900

Windows System Resource Manager ... 901
WMI providers ... 901
XML ... 901
Copyright ... 901

Copyright attributions ... 901
Third Party Notices ... 902
Copyright ... 902

Migrate Roles and Features to Windows Server ... 932
Migration guides ... 933
Windows Server roles, role services, and features ... 933
Windows Server Migration Tools ... 933
See Also ... 933

Migrate Roles and Features to Windows Server 2012 R2 ... 933
In this section ... 934
See also ... 934

Active Directory Certificate Services Migration Guide for Windows Server 2012 R2 ... 934
About this guide ... 934
Target audience ... 934
Supported migration scenarios ... 935
Supported operating systems ... 935
What this guide does not provide ... 937
CA migration overview ... 937
Preparing to migrate ... 938
Migrating the certification authority ... 938
Verifying the migration ... 938
Post-migration tasks ... 938
Impact of migration ... 938
Impact of migration on the source server ... 938
Impact of migration on other computers in the enterprise ... 939
Permissions required to complete the migration ... 939
Estimated duration ... 939
See also ... 939

Prepare to Migrate ... 939
Preparing your destination server ... 940
Hardware requirements for the destination server ... 940
Hardware requirements for AD CS ... 940
Software requirements for the destination server ... 940
Installing the Operating System ... 941
Backing up your source server ... 942
Preparing your source server ... 942

Backing up a CA templates list ... 942
Recording a CA's signature algorithm and CSP ... 943
Publishing a CRL with an extended validity period ... 943
Next steps ... 944
See also ... 944

Migrating the Certification Authority ... 944
Backing up a CA database and private key ... 945
Backing up a CA database and private key by using the Certification Authority snap-in ... 945
Backing up a CA database and private key by using Windows PowerShell ... 946
Backing up a CA database and private key by using Certutil.exe ... 947
Backing up CA registry settings ... 948
Backing up CAPolicy.inf ... 949
Removing the CA role service from the source server ... 949
Removing the source server from the domain ... 950
Joining the destination server to the domain ... 951
Adding the CA role service to the destination server ... 952
Special instructions for migrating to a failover cluster ... 952
Importing the CA certificate ... 952
Adding the CA role service by using Server Manager ... 953
Adding the CA role service by using Windows PowerShell ... 954
Restoring the CA database and configuration on the destination server ... 956
Restoring the source CA database on the destination server ... 956
Restoring the source CA registry settings on the destination server ... 958
Verifying certificate extensions on the destination CA ... 962
Restoring the certificate templates list ... 962
Granting permissions on AIA and CDP containers ... 963
Additional procedures for failover clustering ... 964
Configuring failover clustering for the destination CA ... 964
Granting permissions on public key containers ... 965
Editing the DNS name for a clustered CA in AD DS ... 966
Configuring CRL distribution points for failover clusters ... 967
Next steps ... 968
See also ... 968

Verifying the Certification Authority Migration ... 968
Verifying certificate enrollment ... 968
Verifying CRL publishing ... 970
Next steps ... 970
See also ... 971

Post-Migration Tasks ... 971
Upgrading certificate templates in Active Directory Domain Services (AD DS) ... 971
Retrieving certificates after a host name change ... 972

Restoring Active Directory Certificate Services (AD CS) to the source server in the event of migration failure ... 973
Troubleshooting migration ... 973
See also ... 973

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2 ... 974
About this guide ... 974
Target audience ... 974
Supported migration scenarios ... 974
Supported operating systems ... 975
Supported AD FS role services and features ... 975
See Also ... 976

Preparing to Migrate the AD FS Federation Server ... 976
Migration Process Outline ... 976
New AD FS functionality in Windows Server 2012 R2 ... 977
AD FS Requirements in Windows Server 2012 R2 ... 978
SQL Server support for AD FS in Windows Server 2012 R2 ... 979
Increasing your Windows PowerShell limits ... 979
Other migration tasks and considerations ... 979
See Also ... 980

Migrating the AD FS Federation Server ... 980
Export and backup the AD FS configuration data ... 980
Create a Windows Server 2012 R2 federation server farm ... 983
Import the original configuration data into the Windows Server 2012 R2 AD FS farm ... 984
See Also ... 987

Migrating the AD FS Federation Server Proxy ... 987
See Also ... 988

Verifying the AD FS Migration to Windows Server 2012 R2 ... 988
See Also ... 989

Migrate DHCP Server to Windows Server 2012 R2 ... 989
About this guide ... 989
Target audience ... 989
What this guide does not provide ... 989
Supported migration scenarios ... 990
Supported operating systems ... 990
Supported role configurations ... 992
DHCP Server migration overview ... 992
DHCP Server migration process ... 993
Impact of migration on other computers in the enterprise ... 993
Permissions required to complete migration ... 994

Estimated duration ... 994
See also ... 994

DHCP Server Migration: Preparing to Migrate ... 994
Migration planning ... 994
Install migration tools ... 995
Working with Windows PowerShell cmdlets ... 995
Prepare the destination server ... 996
Prepare the source server ... 998
See also ... 998

DHCP Server Migration: Migrating the DHCP Server Role ... 999
Migrating DHCP Server to the destination server ... 999
Migrating DHCP Server from the source server ... 999
Destination server final migration steps ... 1001
See also ... 1003

DHCP Server Migration: Verifying the Migration ... 1003
Verifying destination server configuration ... 1003
See also ... 1004

DHCP Server Migration: Post-Migration Tasks ... 1004
Completing migration ... 1004
Retiring DHCP on your source server ... 1005
Retiring your source server ... 1005
Restoring DHCP in the event of migration failure ... 1005
Estimated time to complete a rollback ... 1005
Troubleshooting cmdlet-based migration ... 1006
Viewing the content of Windows Server Migration Tools result objects ... 1007
Result object descriptions ... 1007
Examples ... 1009
More information about querying results ... 1010
See also ... 1010

DHCP Server Migration: Appendix A ... 1011
Migration tools ... 1011
Installing and using Windows PowerShell with migration cmdlets ... 1011
Known issues ... 1011
See also ... 1012

Migrate Hyper-V to Windows Server 2012 R2 from Windows Server 2012 ... 1012
About this guide ... 1012
Target audience ... 1012
What this guide does not provide ... 1012
Supported migration scenarios ... 1013

Migration dependencies ... 1013
Migration scenarios that are not supported ... 1013
Overview of migration process for this role ... 1014
Estimated duration ... 1014
Additional references ... 1014

Hyper-V: Migration Options ... 1015
Hyper-V migration options ... 1015
Cross-version live migration ... 1017
Hyper-V Replica ... 1018
See also ... 1019

Hyper-V: Stand-alone Migration ... 1019
Migration options ... 1019
In-place upgrade ... 1019
Perform an in-place upgrade ... 1020
Cross-version live migration ... 1020
Move a virtual machine from Hyper-V in Windows Server 2012 to Windows Server 2012 R2 ... 1021
Modify the Hyper-V Replica settings ... 1021
Verify that the virtual machine runs correctly ... 1022
See also ... 1023

Hyper-V: Hyper-V Cluster Migration ... 1023
Hyper-V Cluster Migrations ... 1023

Hyper-V Cluster Using Separate Scale-Out File Server Migration ... 1023
Cross-version live migration ... 1024
Cross-version live migration scenario ... 1025
Migrate the old cluster node to the new cluster ... 1028
To move the remaining virtual machines ... 1028
Copy Cluster Roles Wizard ... 1028
See also ... 1031

Hyper-V Cluster Using Cluster Shared Volumes (CSV) Migration ... 1031
Copy Cluster Roles Wizard ... 1031
See also ... 1034

Migrate File and Storage Services to Windows Server 2012 R2 ... 1034
About this guide ... 1035
Target audience ... 1035
What this guide does not provide ... 1036
Supported migration scenarios ... 1036
Supported operating systems ... 1037
File services migration overview ... 1038

Impact of migration on other computers in the enterprise ... 1039
Impact of data migration by copying data and shared folders ... 1039
Impact of data migration by physically moving data drives ... 1039
Impact on DFS Namespaces ... 1039
Impact on DFS Replication ... 1039
Permissions required to complete migration ... 1039
Permissions required for data and shared folder migration ... 1040
Permissions required to complete migration on the destination server ... 1040
Permissions required to migrate DFS Namespaces ... 1040
Permissions required to complete migration on the source server ... 1040
Permissions required to migrate DFS Namespaces ... 1040
Permissions required for DFS Replication ... 1041
See also ... 1041

File and Storage Services: Prepare to Migrate ... 1041
Install migration tools ... 1041
Prepare for migration ... 1042
Prepare the destination server ... 1042
Hardware requirements for the destination server ... 1042
Software requirements for the destination server ... 1042
Prepare for local user and group migration on the destination server ... 1043
Prepare for File and Storage Services on destination server ... 1043
Prepare File Server Resource Manager on destination server ... 1044
Data and file share preparation on destination server ... 1044
Data integrity and security considerations on destination server ... 1044
Prepare DFS Namespaces on destination server ... 1044
Back up the source server ... 1044
Prepare the source server ... 1045
Prepare all file services on source server ... 1045
Data and file share preparation on the source server ... 1045
Prepare DFS on the source server ... 1046
Prepare DFS Namespaces on source server ... 1046
Prepare other computers in the enterprise ... 1046
For copy data migration scenarios ... 1046
For physical data migration scenarios ... 1047
See also ... 1047

File and Storage Services: Migrate the File and Storage Services Role ... 1047
Migrate File Services ... 1047
Freeze administration configuration ... 1047
Install the Windows Server Migration Tools ... 1048
Export settings ... 1048
BranchCache for Network Files server key ... 1049
Group Policy setting or local policy setting specific to SMB and Offline Files ... 1050

Server message block ... 1050
Offline Files ... 1051
DFS Namespace configuration ... 1052
Considerations for namespaces ... 1054
Inventory advanced registry keys ... 1055
DFS Replication configuration ... 1055
File Server Resource Manager configuration on the source server ... 1056
Shadow Copies of Shared Folders ... 1057
Migrate local users and groups to the destination server ... 1058
Export local users and groups from the source server ... 1059
Import local users and groups to the destination server ... 1059
Migrate data ... 1059
Data copy migration ... 1060
Physical data migration ... 1062
Using disk drives or LUNs to migrate data from the source server to the destination server ... 1062
Migrate shared folders ... 1064
DFS Replication migration ... 1066
Migrate the source server identity ... 1067
Rename the source server ... 1067
Migrate IP address ... 1067
Rename destination server ... 1067
Export Remote VSS settings ... 1067
If you migrated the data by copying it ... 1068
If you migrated the data by physically moving it ... 1068
Import settings to the destination server ... 1069
Group Policy or local policy specific to server message block and Offline Files ... 1070
DFS Namespace configuration ... 1071
Stand-alone namespaces ... 1071
Domain-based namespaces with more than one namespace server ... 1072
Domain-based namespaces with one namespace server ... 1072
File Server Resource Manager configuration on the destination server ... 1073
Shadow Copies of Shared Folders ... 1075
Deduplication ... 1076
Migrating Deduplication from Windows Server 2012 to Windows Server 2012 ... 1076
Migrating SIS from Windows Storage Server 2008 to Windows Server 2012 ... 1076
Migrating SIS volumes ... 1077
Import Remote VSS settings ... 1077
See also ... 1078

File and Storage Services: Verify the Migration ... 1078
Verify the File Services migration ... 1078
Verify migration of BranchCache for Network File Services server key ... 1079

Verify migration of local users and groups 1079
Verify data and shared folder migration 1079
Verify the migration of DFS Namespaces 1079
Verify the configuration on other computers 1080
Verify the File Server Resource Manager migration 1081
See Also 1081

File and Storage Services: Migrate an iSCSI Software Target 1081
Supported migration scenarios 1082
Supported operating systems 1082
Supported role configurations 1083
Supported role services and features 1084
Migrating multiple roles 1084
Migration scenarios that are not supported 1084
Migration overview 1085
Migration process 1085
Impact of migration 1087
Permissions required for migration 1087
Estimated time duration 1087
See Also 1088

Prepare to Migrate iSCSI Software Target 1089
Prepare the destination server 1089
Back up the source server 1089
Prepare the source server 1089
Cluster resource group configuration 1089
iSCSI Target portal configuration 1091
iSNS configuration 1092
CHAP and Reverse CHAP configuration 1092
Snapshot storage configuration 1092
Disconnect the iSCSI initiators 1092
Capture the existing settings: stand-alone configuration 1093
Capture the existing settings: clustered configuration 1093
Remove the network identity of the iSCSI Software Target computer 1094
Prepare the iSCSI initiator computers 1095
Capture the session information 1095
Disconnect the session 1095
See Also 1095

Migrate iSCSI Software Target 1096
Migrating iSCSI Software Target in a standalone configuration 1096
Establish network identity of the iSCSI Target Server computer 1096
Configure the iSCSI Target Server portal 1096
Configure iSNS settings 1097

Configure storage ... 1097
Configure the Volume Shadow Copy Service ... 1097
Transfer the virtual disk ... 1098
Import the iSCSI Software Target settings in a stand-alone configuration ... 1098
Configure shadow storage for the virtual disks ... 1099
Configure CHAP and Reverse CHAP ... 1099
Migrating iSCSI Software Target in a failover cluster ... 1099
Migrate resource groups ... 1100
Import the iSCSI Software Target settings in a failover cluster ... 1100
Migrate iSCSI Target Server Providers ... 1101
See Also ... 1101

Verify the iSCSI Software Target Migration ... 1102
Verifying the destination server configuration ... 1102
Verify the listening endpoints ... 1102
Verify the basic connectivity ... 1102
Perform a Best Practices Analyzer scan ... 1103
Verifying the configuration of iSCSI initiator computers ... 1103
Verify that the iSCSI initiators can discover iSCSI Target Server ... 1103
Verify that the iSCSI initiators can log on ... 1103
See Also ... 1104

Troubleshoot the iSCSI Software Target Migration ... 1104
Understanding the messages from the iSCSI Target Migration tool ... 1104
See Also ... 1106

Roll Back a Failed iSCSI Software Target Migration ... 1106
Restoring the role if the migration failed ... 1107
Rollback requirements ... 1107
Roll back iSCSI initiators on other computers ... 1107
Roll back iSCSI Software Target on a stand-alone source server ... 1107
Roll back iSCSI Software Target on a clustered source server ... 1108
Roll back iSCSI Target Server on a stand-alone destination server ... 1108
Roll back iSCSI Target Server on a clustered destination server ... 1109
Retiring iSCSI Software Target on a source server ... 1109
Retiring a source server ... 1109
See Also ... 1109

File and Storage Services: Migrate Network File System ... 1109
Network File System Migration overview ... 1110
Migrating NFS Server from Windows Server°2012 to Windows Server°2012°R2 ... 1110
Export the server configuration ... 1110
Export NFS shares ... 1110
Export NFS share permissions ... 1111

Copy local mapping data ... 1111
Export identity mapping ... 1111
Export netgroups and client groups ... 1111
Importing NFS shares and settings from Windows Server°2012 to Windows Server°2012°R2 ... 1112
Import the server configuration ... 1112
Import NFS shares ... 1112
Import NFS share permissions ... 1113
Import local mapping data ... 1113
Import non-local identity mapping ... 1113
Import netgroups and client groups ... 1113
Migrating NFS Server from Windows Server°2008°R2, Windows Server°2008, or Windows Server°2003°R2 to Windows Server°2012°R2 ... 1114
Get server configuration ... 1114
Collect NFS shares information ... 1115
Collect identity mapping and group identifier information ... 1116
Reconfiguring NFS shares and settings from Windows Server°2008°R2, Windows Server°2008, or Windows Server°2003°R2 to Windows Server°2012°R2 ... 1116
Set up the NFS server configuration ... 1117
Configure NFS shares ... 1118
Configure identity mapping and group identifier information ... 1119
See Also ... 1120

File and Storage Services: Post-Migration Tasks ... 1120
Completing the migration ... 1120
Retire File and Storage Services on the source server ... 1120
Remove DFS Namespaces from the source server ... 1120
Restoring File and Storage Services in the event of migration failure ... 1121
Roll back DFS Namespaces ... 1121
Roll back data and shared folders ... 1122
Roll back migration on the other computers in the enterprise ... 1123
Troubleshooting migration issues ... 1123
Troubleshoot data migration that does not complete ... 1123
Troubleshoot data migration connectivity ... 1124
Troubleshoot unexpected Windows PowerShell session closure ... 1125
Locate the deployment log file ... 1126
View the content of Windows Server Migration Tools result objects ... 1126
Result object descriptions ... 1126
Examples ... 1128
More information about querying results ... 1130
See Also ... 1130

File and Storage Services: Appendix A: Optional Procedures ... 1130
Opening ports in Windows Firewall ... 1130

Closing ports in Windows Firewall 1131
Detect reparse points and hard links 1132
Migrated and nonmigrated attributes for local users and groups 1133
See Also 1133

File and Storage Services: Appendix B: Migration Data Collection Worksheets 1133
SMB data collection worksheet 1133
BranchCache data collection worksheet 1134
See Also 1135

Migrate Remote Desktop Services to Windows Server 2012 R2 1135
About this guide 1135
Target audience 1135
What this guide does not provide 1136
Supported migration scenarios 1136
Supported operating systems 1137
Policy and configuration settings 1137
Supported role services and features 1137
Migration scenarios that are not supported 1137
Order of migration for multiple role services 1138
Impact of migration on Remote Desktop Services 1138
Additional references 1140

Remote Desktop Services: Prepare to Migrate 1141
Assign permissions required to migrate Remote Desktop Services 1141
Migration dependencies 1141
Prerequisite features to migrate separately 1141
Prerequisite features already installed 1142
Prepare your source server 1142
Back up your source server 1142
Gather data from your source server 1143
Prepare your destination server 1143
Hardware requirements for the destination server 1143
Software requirements for the destination server 1143
Other servers and client computers in the enterprise 1143
Additional references 1143

Remote Desktop Services: Migrate Remote Desktop Services Role Services 1144
Migrate the RD Connection Broker server 1144
Migrate session collections 1145
Migrate virtual desktop collections 1145
Migrate RD Web Access servers 1146
Migrate RD Gateway servers 1146
Migrate RD Licensing servers 1146

Migrate standalone Remote Desktop Services servers ... 1147
Migrate certificates ... 1147
Remote Desktop Services features that use certificates ... 1147
Preparing certificates for migration ... 1147
Additional references ... 1148

Remote Desktop Services: Verify the Migration ... 1148
Run a pilot program ... 1148
Additional references ... 1148

Remote Desktop Services: Post-Migration Tasks ... 1149
Retire the source servers ... 1149

Migrate Cluster Roles to Windows Server 2012 R2 ... 1149
Operating system requirements for clustered roles and feature migrations ... 1150
Target audience ... 1150
What this guide does not provide ... 1151
Planning considerations for migrations between failover clusters ... 1151
Migration scenarios that use the Copy Cluster Roles Wizard ... 1152
In this guide ... 1152
Related references ... 1152

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2 ... 1153
Migration paths for specific migrations ... 1153
Cluster roles that cannot be migrated ... 1155
Roles restricted to a single instance per cluster ... 1155
Migrations for which the Copy Cluster Roles Wizard performs most or all steps ... 1155
Migration within mixed environments ... 1156
Additional steps for a wizard-based migration ... 1157
Failover Cluster Copy Roles reports ... 1157
Clustered role and feature migrations that require extra steps ... 1158
Clustered DFS Replication migrations ... 1158
Clustered DHCP migrations ... 1158
Clustered DTC migrations ... 1159
Clustered File Server and Scale-out File Server migrations ... 1159
Choosing the best migration method for your file server ... 1160
Virtual machine storage migration ... 1160
Copy Cluster Roles Wizard - Migrate to a new multi-node cluster ... 1160
Copy Cluster Roles Wizard – In-place migration ... 1161
Storage pool migration ... 1162
Additional tasks for file server migration using the Copy Cluster Roles Wizard ... 1163
Clustered FSRM migrations ... 1163
Clustered Message Queuing (MSMQ) migrations ... 1163
Other Server migrations involving resource types not built into failover clusters ... 1164

Migration of highly available virtual machines 1164
Alternate methods for migrating HAVMs to a Windows Server 2012 R2 failover cluster 1165
Additional tasks for using the Copy Cluster Roles Wizard to migrate HAVMs 1166
Additional references 1166

Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2 1167
Overview of migration of cluster roles between two multi-node failover clusters 1167
Impact of a migration between two multi-node clusters 1168
Access rights required to complete migration 1168
Additional references 1168
Cluster roles: Prepare to migrate between two multi-node clusters 1169
Cluster roles: Migrate the cluster roles 1170
Cluster roles: Post-migration tasks for a migration between two multi-node clusters 1172
Cluster roles: Verify the migration 1173

In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2 1174
Overview of an in-place migration for a two-node cluster 1175
Impact of the migration 1175
Access rights required to complete migration 1176
Additional references 1176
Create a new cluster from a node in the old cluster 1176
Copy the cluster roles to the new cluster 1178
Perform post-migration tasks 1180
Add the second node to the new cluster 1181
Verify failover for the migrated cluster roles 1184

Cluster Migrations Involving New Storage: Mount Points 1185
Additional references 1186

Additional References 1186

Migrate Network Policy Server to Windows Server 2012 R2 1187
About this guide 1187
Target audience 1188
What this guide does not provide 1188
Supported migration scenarios 1188
Supported operating systems 1188
Supported NPS role configurations 1189
IP address and host name configuration 1190
Migration scenarios that are not supported 1190
Overview of migration process for this role 1190
Impact of migration 1191
Impact of migration on the source server 1191
Impact of migration on other computers in the enterprise 1191
Permissions required to complete migration 1191

Estimated duration ... 1192

Prepare to Migrate ... 1192
Choose a migration file storage location ... 1192
Prepare your source server ... 1192
Prepare your destination server ... 1193

Migrating the NPS Server ... 1193
Known issues ... 1194
Exporting settings from the source server ... 1194
Exporting settings from Windows Server 2003 ... 1194
Exporting settings from Windows Server 2008 ... 1196
Exporting settings from Windows Server 2008 R2 ... 1197
Exporting settings from Windows Server 2012 or Windows Server 2012 R2 ... 1198
Importing settings to the destination server ... 1201
Importing settings from Windows Server 2003 ... 1201
Importing settings from Windows Server 2008 or Windows Server 2008 R2 ... 1203
Importing settings from Windows Server 2012 or Windows Server 2012 R2 ... 1204
Using the NPS console to migrate NPS settings ... 1205

Verifying the NPS Server Migration ... 1206
Verifying NPS Migration ... 1206

Post-Migration Tasks ... 1208
Post migration tasks ... 1208
Restoring the role in the event of migration failure ... 1209

Appendix A - Data Collection Worksheet ... 1209
Migration data collection worksheet ... 1209

Migrate Roles and Features to Windows Server 2012 ... 1211
In this section ... 1211
See Also ... 1212

Install, Use, and Remove Windows Server Migration Tools ... 1212
In this guide ... 1212
Supported operating systems ... 1213
Permission requirements ... 1214
Prepare for installation ... 1215
Windows Server 2012 source server ... 1215
Windows Server 2008 R2 source server ... 1215
Windows Server 2008 source server ... 1215
Windows Server 2003 or Windows Server 2003 R2 source server ... 1215
Other computers in your enterprise ... 1216
Install Windows Server Migration Tools ... 1216

Full installation option of Windows Server 2012 R2 or Windows Server 2012 1216
Server Core installation option of Windows Server 2012 R2 or Windows Server 2012 1217
Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 source computers 1218
Creating a deployment folder on destination computers 1218
Registering Windows Server Migration Tools on source computers 1219
Use Windows Server Migration Tools 1221
Full installation option of Windows Server 2012 R2 1221
Server Core installation option of Windows Server 2012 R2 1221
Full installation option of Windows Server 2012 1222
Server Core installation option of Windows Server 2012 1222
Source computer running full installation option of Windows Server 2008 R2 1222
Source computer running Server Core installation option of Windows Server 2008 R2 1223
Windows Server 2003 or Windows Server 2008 source computers 1223
Additional resources and next steps for using Windows Server Migration Tools 1224
Remove Windows Server Migration Tools 1225
Full installation option of Windows Server 2012 R2 or Windows Server 2012 1225
Server Core installation option of Windows Server 2012 R2 or Windows Server 2012 1226
Source computers running full and Server Core installation options of Windows Server 2012 1226
Source computers running full and Server Core installation options of Windows Server 2008 R2 1227
Windows Server 2003 or Windows Server 2008 source computers 1227
See Also 1228

Migrate Active Directory Federation Services Role Services to Windows Server 2012 1228
About this guide 1228
Target audience 1228
Supported migration scenarios 1228
Supported operating systems 1229
Supported AD FS role services and features 1230
See Also 1231

Prepare to Migrate the AD FS 2.0 Federation Server 1231
Prepare to migrate a stand-alone AD FS federation server or a single-node AD FS farm 1232
Step 1: Export service settings 1232
Step 2: - Export claims provider trusts 1234
Step 3: - Export relying party trusts 1234
Step 4: - Back up custom attribute stores 1235
Step 5: Back up webpage customizations 1235
Prepare to migrate a WID farm 1235
Step 1: - Export service settings 1235
Step 2: Back up custom attribute stores 1236
Step 3: Back up webpage customizations 1236

Prepare to migrate a SQL Server farm ... 1236
Step 1: Export service settings ... 1236
Step 2: Back up custom attribute stores ... 1238
Step 3: Back up webpage customizations ... 1238
See Also ... 1238

Prepare to Migrate the AD FS 2.0 Federation Server Proxy ... 1238
Step 1: Export proxy service settings ... 1238
Step 2: Back up webpage customizations ... 1239
See Also ... 1239

Migrate the AD FS 2.0 Federation Server ... 1239
Migrate a stand-alone AD FS federation server or a single-node AD FS farm ... 1240
Migrate a WID farm ... 1242
Migrate a SQL Server farm ... 1244
Restoring the Remaining AD FS Farm Configuration ... 1245
See Also ... 1246

Migrate the AD FS 2.0 Federation Server Proxy ... 1246
See Also ... 1247

Migrate the AD FS 1.1 Web Agents ... 1247
See Also ... 1248

Migrate File and Storage Services to Windows Server 2012 ... 1248
About this guide ... 1248
Target audience ... 1249
What this guide does not provide ... 1249
Supported migration scenarios ... 1250
Supported operating systems ... 1250
File services migration overview ... 1252
Impact of migration on other computers in the enterprise ... 1252
Impact of data migration by copying data and shared folders ... 1252
Impact of data migration by physically moving data drives ... 1252
Impact on DFS Namespaces ... 1253
Impact on DFS Replication ... 1253
Permissions required to complete migration ... 1253
Permissions required for data and shared folder migration ... 1253
Permissions required to complete migration on the destination server ... 1253
Permissions required to migrate DFS Namespaces ... 1253
Permissions required to complete migration on the source server ... 1254
Permissions required to migrate DFS Namespaces ... 1254
Permissions required for DFS Replication ... 1254
See Also ... 1254

File and Storage Services: Prepare to Migrate 1255
Install migration tools 1255
Prepare for migration 1255
Prepare the destination server 1256
Hardware requirements for the destination server 1256
Software requirements for the destination server 1256
Prepare for local user and group migration on the destination server 1256
Prepare for File and Storage Services on destination server 1256
Prepare File Server Resource Manager on destination server 1257
Data and shared folder preparation on destination server 1257
Data integrity and security considerations on destination server 1257
Prepare DFS Namespaces on destination server 1258
Back up the source server 1258
Prepare the source server 1258
Prepare all file services on source server 1258
Data and shared folder preparation on the source server 1259
Prepare DFS on the source server 1259
Prepare DFS Namespaces on source server 1259
Prepare other computers in the enterprise 1260
For copy data migration scenarios 1260
For physical data migration scenarios 1260
See Also 1260

File and Storage Services: Migrate the File and Storage Services Role 1261
Migrate File Services 1261
Freeze administration configuration 1261
Install the Windows Server Migration Tools 1261
Export settings 1262
BranchCache for Network Files server key 1262
Group or local policy specific to SMB and Offline Files 1263
Server message block 1263
Offline Files 1264
DFS Namespace configuration 1266
Considerations for namespaces 1267
Inventory advanced registry keys 1269
DFS Replication configuration 1269
File Server Resource Manager configuration on the source server 1269
Shadow Copies of Shared Folders 1271
Migrate local users and groups to the destination server 1272
Export local users and groups from the source server 1272
Import local users and groups to the destination server 1272
Migrate data 1273
Data copy migration 1273

Physical data migration ... 1275
Using disk drives or LUNs to migrate data from the source server to the destination server ... 1275
Migrate shared folders ... 1278
DFS Replication migration ... 1279
Migrate the source server identity ... 1280
Rename the source server ... 1280
Migrate IP address ... 1280
Rename destination server ... 1280
Configure DFS Replication on the destination server ... 1280
If you migrated the data by copying it ... 1280
If you migrated the data by physically moving it ... 1281
Import settings to the destination server ... 1282
Group Policy or local policy specific to server message block and Offline Files ... 1283
DFS Namespace configuration ... 1284
Stand-alone namespaces ... 1284
Domain-based namespaces with more than one namespace server ... 1285
Domain-based namespaces with one namespace server ... 1285
File Server Resource Manager configuration on the destination server ... 1286
Shadow Copies of Shared Folders ... 1288
Deduplication ... 1289
Migrating Deduplication from Windows Server 2012 to Windows Server 2012 ... 1289
Migrating SIS from Windows Storage Server 2008 to Windows Server 2012 ... 1289
Migrating SIS volumes ... 1289
See Also ... 1290

File and Storage Services: Verify the Migration ... 1290
Verify the File Services migration ... 1291
Verify migration of BranchCache for Network File Services server key ... 1291
Verify migration of local users and groups ... 1291
Verify data and shared folder migration ... 1291
Verify the migration of DFS Namespaces ... 1292
Verify the configuration on other computers ... 1293
Verify the File Server Resource Manager migration ... 1293
See Also ... 1293

File and Storage Services: Post-Migration Tasks ... 1294
Completing the migration ... 1294
Retire File and Storage Services on the source server ... 1294
Remove DFS Namespaces from the source server ... 1294
Restoring File and Storage Services in the event of migration failure ... 1295
Roll back DFS Namespaces ... 1295
Roll back data and shared folders ... 1296
Roll back migration on the other computers in the enterprise ... 1296

Troubleshooting migration issues ... 1296
Troubleshoot data migration that does not complete ... 1297
Troubleshoot data migration connectivity ... 1298
Troubleshoot unexpected Windows PowerShell session closure ... 1299
Locate the deployment log file ... 1299
View the content of Windows Server Migration Tools result objects ... 1300
Result object descriptions ... 1300
Examples ... 1302
More information about querying results ... 1303
See Also ... 1304

File and Storage Services: Appendix A: Optional Procedures ... 1304
Opening ports in Windows Firewall ... 1304
Closing ports in Windows Firewall ... 1305
Detect reparse points and hard links ... 1305
Migrated and non-migrated attributes for local users and groups ... 1306
See Also ... 1307

File and Storage Services: Appendix B: Migration Data Collection Worksheets ... 1307
SMB data collection worksheet ... 1307
BranchCache data collection worksheet ... 1308
See Also ... 1308

File and Storage Services: Appendix C: Migrate iSCSI Software Target ... 1309
See Also ... 1309

iSCSI SoftwareTarget Migration Overview ... 1309
Migration overview ... 1309
Migration process ... 1310
Impact of migration ... 1311
Permissions required for migration ... 1312
Estimated time duration ... 1312
Supported migration scenarios ... 1313
Supported operating systems ... 1313
Supported role configurations ... 1314
Supported role services and features ... 1315
Migrating multiple roles ... 1315
Migration scenarios that are not supported ... 1315

Prepare to Migrate iSCSI Software Target ... 1316
Prepare the destination server ... 1316
Backup the source server ... 1316
Prepare the source server ... 1317
Cluster resource group configuration ... 1317
iSCSI Target portal configuration ... 1318

iSNS configuration ... 1319
CHAP and Reverse CHAP configuration ... 1319
Snapshot storage configuration ... 1319
Disconnect the iSCSI initiators ... 1320
Capture the existing settings: standalone configuration ... 1320
Capture the existing settings: clustered configuration ... 1321
Remove the network identity of the iSCSI Software Target computer ... 1322
Prepare the iSCSI initiator computers ... 1322
Capture the session information ... 1322
Disconnect the session ... 1323

Migrate iSCSI Software Target ... 1323
Migrating ISCSI Software Target in a standalone configuration ... 1323
Establish network identity of the iSCSI Target Server computer ... 1323
Configure the iSCSI Target Server portal ... 1323
Configure iSNS settings ... 1324
Configure storage ... 1324
Configure the Volume Shadow Copy Service ... 1324
Transfer the virtual disk ... 1325
Import the iSCSI Software Target settings in a standalone configuration ... 1325
Configure shadow storage for the virtual disks ... 1326
Configure CHAP and Reverse CHAP ... 1326
Migrating iSCSI Software Target in a failover cluster ... 1326
Migrate resource groups ... 1327
Import the iSCSI Software Target settings in a failover cluster ... 1327

Verify the iSCSI Software Target Migration ... 1328
Verifying the destination server configuration ... 1328
Verify the listening endpoints ... 1328
Verify the basic connectivity ... 1329
Perform a Best Practices Analyzer scan ... 1329
Verifying the configuration of iSCSI initiator computers ... 1329
Verify that the iSCSI initiators can discover iSCSI Target Server ... 1329
Verify that the iSCSI initiators can log on ... 1329

Troubleshoot the iSCSI Software Target Migration ... 1330
Understanding the messages from the iSCSI Target Migration tool ... 1330

Roll Back a Failed iSCI Software Target Migration ... 1332
Restoring the role if the migration failed ... 1332
Rollback requirements ... 1332
Roll back iSCSI initiators on other computers ... 1333
Roll back iSCSI Software Target on a standalone source server ... 1333
Roll back iSCSI Software Target on a clustered source server ... 1333

Roll back ISCSI Target Server on a standalone destination server 1334
Roll back ISCSI Target Server on a clustered destination server 1334
Retiring ISCSI Software Target on a source server 1334
Retiring a source server 1335

Migrate Health Registration Authority to Windows Server 2012 1335
About this guide 1335
Target audience 1335
What this guide does not provide 1335
Supported migration scenarios 1336
Supported operating systems 1336
Supported role configurations 1337
Migrating prerequisite roles 1337
Migration scenarios that are not covered 1338
Overview of migration process for this role 1338
Impact of migration 1339
Impact of migration on the source server 1339
Impact of migration on other computers in the enterprise 1339
Permissions required to complete migration 1339
Estimated duration 1340
See Also 1340

HRA Server Migration: Preparing to Migrate 1340
Choose a migration file storage location 1340
Prepare your source server 1340
Prepare your destination server 1341
See Also 1341

HRA Server Migration: Migrating the HRA Server 1341
Migrating settings from the source server 1342
Configuring the destination server 1342
Migrating settings to the destination server 1344
Configuring the Certification Authority 1345
Configuration tips for migrating the Certification Authority 1346
See Also 1346

HRA Server Migration: Verifying the Migration 1346
Verifying HRA Functionality 1347
Adding a new trusted server group for testing 1347
Testing the HRA with a NAP client 1347
See Also 1348

HRA Server Migration: Post-migration Tasks 1348
Deploying final client settings 1348
Restoring the role in the event of migration failure 1349

Retiring the Source Server 1349
Troubleshooting migration 1350
See Also 1350

Migrate Hyper-V to Windows Server 2012 from Windows 2008 R2 1350
About this guide 1350
Target audience 1351
What this guide does not provide 1351
Supported migration scenarios 1352
Supported operating systems 1352
Supported role configurations and settings 1353
Migration dependencies 1354
Migration scenarios that are not supported 1354
Hyper-V migration overview 1355
Impact of migration 1355
Impact of migration on the source server 1355
Impact of migration on other computers in the enterprise 1355
Access rights required to complete migration 1356
Estimated duration 1356
Additional references 1356

Hyper-V: Prepare to Migrate 1356
Select and prepare your destination server 1356
Hardware requirements for the destination server 1356
Software requirements for the destination server 1357
Back up your source server 1357
Install migration tools 1357
Collect configuration details from your source server 1358
Prepare other computers in the enterprise 1359
Additional references 1359

Hyper-V: Migrate the Hyper-V Role 1359
Migrate the Hyper-V Role 1359
Perform migration steps on the source server 1360
Migrate virtual machine data 1361
Perform migration steps on the destination server 1363

Hyper-V: Verify the Migration 1365
Verify the Hyper-V security policy 1365
Verify the networking configuration 1365
Verify the configuration and availability of the virtual machines 1365

Hyper-V: Post-migration Tasks 1367
Retiring your source server 1367
Restoring the role in the event of migration failure 1367

Roll back migration of Hyper-V on the source server 1367
Roll back migration of Hyper-V on the destination server running Windows Server 2012 1367
Roll back migration changes on other computers in the enterprise 1368
Troubleshooting cmdlet-based migration 1368
Viewing the content of Windows Server Migration Tools result objects 1369
Result object descriptions 1369
Examples 1371
More information about querying results 1372

Migrate IP Configuration to Windows Server 2012 1373
Supported operating systems 1373
Supported scenarios and features 1374
Scenarios and features that are not supported 1377
See Also 1377

IP Configuration: Prepare to Migrate 1378
Impact on the source server 1378
Impact on the destination server 1378
Impact on other servers in your enterprise 1378
Impact on other client computers in your enterprise 1378
Expected downtime during IP configuration migration 1379
User rights required to perform migration on both source and destination servers 1379
Preparing the destination server 1379
Preparing the source server 1379
Preparing other computers in the enterprise 1380
See Also 1380

IP Configuration: Migrate IP Configuration Data 1380
Migrating Global and NIC IP configuration 1380
IP configuration migration tools 1380
Migrating IP configuration by using Windows Server Migration Tools 1381
Export IP configuration settings from the source server 1381
Import IP configuration settings to the destination server 1382
See Also 1383

IP Configuration: Post-migration Tasks 1383
Verifying the migration 1383
Rolling back migration 1384
Troubleshooting cmdlet-based migration 1384
Viewing the content of Windows Server Migration Tools result objects 1385
Result object descriptions 1385
Examples 1387
See Also 1389

IP Configuration: Appendix 1389

Migrating manually-configured IPv6 interface metrics from Windows Server 2003 1389
Additional resources 1390
See Also 1391

Migrate Network Policy Server to Windows Server 2012 1391
About this guide 1391
Target audience 1392
What this guide does not provide 1392
Supported migration scenarios 1392
Supported operating systems 1392
Supported NPS role configurations 1393
IP address and host name configuration 1394
Migration scenarios that are not supported 1394
Overview of migration process for this role 1394
Process diagram 1395
Impact of migration 1395
Impact of migration on the source server 1395
Impact of migration on other computers in the enterprise 1396
Permissions required to complete migration 1396
Estimated duration 1396
See Also 1396

NPS Server Migration: Preparing to Migrate 1396
Choose a migration file storage location 1397
Prepare your source server 1397
Prepare your destination server 1397
See Also 1398

NPS Server Migration: Migrating the NPS Server 1398
Known issues 1398
Exporting settings from the source server 1399
Exporting settings from Windows Server 2003 1399
Exporting settings from Windows Server 2008 1400
Exporting settings from Windows Server 2008 R2 1402
Exporting settings from Windows Server 2012 1403
Importing settings to the destination server 1406
Importing settings from Windows Server 2003 1406
Importing settings from Windows Server 2008 or Windows Server 2008 R2 1408
Importing settings from Windows Server 2012 1409
Using the NPS console to migrate NPS settings 1410
See Also 1410

NPS Server Migration: Verifying the Migration 1411
Verifying NPS Migration 1411

See Also ... 1413

NPS Server Migration: Post-migration Tasks ... 1413
Post migration tasks ... 1413
Restoring the role in the event of migration failure ... 1414
See Also ... 1414

NPS Server Migration: Appendix A - Data Collection Worksheet ... 1415
Migration data collection worksheet ... 1415
See Also ... 1417

Migrate Print and Document Services to Windows Server 2012 ... 1417
Overview ... 1417
About this guide ... 1419
Target audience ... 1419
What this guide does not provide ... 1419
Supported migration scenarios ... 1419
Supported operating systems ... 1419
Supported role configurations ... 1421
Supported role services and features ... 1421
Migrating from x86-based to x64-based v3 printer drivers ... 1421
Unsupported scenarios ... 1422
Print and Document Services migration overview ... 1422
Migrate print servers (overview) ... 1423
Impact of migration ... 1423
Impact of migration on the source server ... 1423
Impact of migration on other computers in the enterprise ... 1423
Permissions required to complete migration ... 1424
Permissions required to complete migration on other computers in the enterprise ... 1424
Estimated duration ... 1424
See Also ... 1424

Preparing to Migrate ... 1425
Access the migration tools ... 1425
To access the Printer Migration Wizard ... 1425
To access the Printbrm.exe command-line tool ... 1426
Prepare the destination server ... 1426
Hardware requirements for the destination server ... 1426
Software requirements for the destination server ... 1426
Installing the Print and Document Services role on the destination server ... 1427
Preparing for cross-architecture migrations ... 1427
Preparing for additional scenarios ... 1427
Prepare the source server ... 1428
See Also ... 1429

Migrating the Print and Document Services Role ... 1429
Back up the source server ... 1429
Cross-architecture migrations ... 1431
Restoration ... 1431
See Also ... 1432

Verifying the Migration ... 1433
Verify the migration ... 1433
To verify destination server configuration ... 1433
Rename the destination server to the name of the source server ... 1434
To verify configuration of other computers in the enterprise ... 1434
Print a test job from a client with an existing connection ... 1435
See Also ... 1435

Post-Migration Tasks ... 1435
Post-migration ... 1435
Success ... 1435
Retire the source server ... 1435
Failure ... 1436
Restoring the role in the event of migration failure ... 1436
Rollback requirements ... 1436
Estimated time to complete rollback ... 1436
Roll back migration on the source server ... 1437
Roll back migration on the destination server ... 1437
Troubleshooting ... 1437
Log file locations ... 1437
Migrating cross-platform driver language monitors ... 1437
Mitigating a failure in the Print Spooler service ... 1437
Additional references ... 1438
See Also ... 1438

Appendix A - Printbrm.exe Command-Line Tool Details ... 1438
Printbrm.exe command-line tool syntax ... 1438
Printbrm enhancements ... 1439
Printbrm usage scenarios ... 1440
Using the configuration file ... 1440
Selectively restoring your printers ... 1441
Moving printers to a different domain ... 1441
See Also ... 1442

Appendix B - Additional Destination Server Scenarios ... 1442
If your server hosts Line Printer Remote (LPR) printers ... 1442
If your server offers Internet Printing Protocol (IPP) printer connections ... 1443
If your server hosts Web Services on Devices (WSD) printers ... 1443

If your print server is a highly available virtual machine ... 1443
If your server hosts local bus printers (LPT and USB) ... 1443
If your server hosts plug and play printers ... 1443
See Also ... 1444

Appendix C - Printbrm Event IDs ... 1444
Printbrm Event IDs ... 1444
See Also ... 1457

Migrate Remote Access to Windows Server 2012 ... 1458
About this guide ... 1458
Target audience ... 1458
What this guide does not provide ... 1458
Supported migration scenarios ... 1459
Supported operating systems ... 1459
Supported role configurations ... 1460
Migration dependencies ... 1460
Migration components that are not supported in all operating system versions ... 1461
Migration components that are not automatically migrated ... 1464
Overview of the Routing and Remote Access service migration process ... 1465
Impact of migration ... 1466
Permissions required to complete migration ... 1466
Estimated duration ... 1467
See Also ... 1467

Remote Access: Prepare to Migrate ... 1467
Prepare your destination server ... 1467
Hardware requirements for the destination server ... 1467
Prepare the destination server for migration ... 1468
Prepare your source server ... 1469
Back up your source server ... 1469
Install the migration tools ... 1470
See Also ... 1470

Remote Access: Migrate Remote Access ... 1470
Migrating Remote Access from the source server ... 1471
Migrating Remote Access to the destination server ... 1474
Completing the required manual migration steps ... 1475
DirectAccess ... 1475
Dial-up demand-dial connections ... 1475
Certificates for IKEv2, SSTP, and L2TP/IPsec connections ... 1476
Routing and Remote Access service policies and accounting settings ... 1476
PEAP, smart card, and other certificate settings on Network Policy Server ... 1476
Weak encryption settings ... 1476

Connection Manager profiles ... 1477
Group forwarded fragments ... 1477
RAS administration and security DLLs ... 1477
See Also ... 1478

Remote Access: Verify the Migration ... 1478
Verifying the destination server configuration ... 1478
Installation state of Remote Access ... 1478
Status of Remote Access Service ... 1478
Remote access Operations Status ... 1479
DirectAccess configuration ... 1479
VPN configuration ... 1479
Dial-up configuration ... 1480
Demand-dial VPN configuration ... 1480
Router settings ... 1480
User and Group accounts ... 1482
Final checks ... 1482
See Also ... 1482

Remote Access: Post-migration Tasks ... 1482
Completing the migration ... 1483
Configuring firewall rules for VPN ... 1483
Configuring firewall rules for DirectAccess ... 1484
Restoring Remote Access in the event of migration failure ... 1484
Estimated time to complete a rollback ... 1485
Retiring Remote Access on your source server ... 1485
Troubleshooting cmdlet-based migration ... 1485
Viewing the content of Windows Server Migration Tools result objects ... 1486
Result object descriptions ... 1486
Examples ... 1488
More information about querying results ... 1490
See Also ... 1490

Migrate Windows Server Update Services to Windows Server 2012 ... 1490

Step 1: Plan for WSUS Migration ... 1491
1.1. Know supported operating systems ... 1491
1.2. Review supported migration scenarios ... 1491
1.3. Review migration scenarios that are not supported ... 1492
See also ... 1492

Step 2: Prepare to Migrate WSUS ... 1492
2.1. Prepare before you start the migration ... 1493
2.2. Prepare the destination server ... 1494
2.3. Prepare the source server ... 1494

See also ... 1494

Step 3: Migrate WSUS ... 1495
3.1. Migrate WSUS update binaries ... 1495
3.2. Migrate WSUS security groups ... 1496
3.3. Back up the WSUS database ... 1497
3.4. Change the WSUS server identity ... 1501
3.5. Apply security settings ... 1501
Point the downstream servers to the new WSUS server ... 1502
Point the WSUS clients to the new WSUS server ... 1502
3.6. Review additional considerations ... 1503
See also ... 1503

Step 4: Verify the WSUS Migration ... 1504
4.1. Verify the destination server configuration ... 1504
4.2. Verify client computer functionality ... 1504
See also ... 1504

Migrating Clustered Services and Applications to Windows Server 2012 ... 1505
Operating system requirements for clustered roles and feature migrations ... 1505
Target audience ... 1505
What this guide does not provide ... 1506
Planning considerations for migrations between failover clusters ... 1506
Migration scenarios that use the Migrate a Cluster Wizard ... 1507
In this guide ... 1507
Related references ... 1508

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 ... 1508
Migration paths for specific migrations ... 1508
Cluster roles that cannot be migrated ... 1510
Roles restricted to a single instance per cluster ... 1510
Migrations for which the Migrate a Cluster Wizard performs most or all steps ... 1510
Migration within mixed environments ... 1511
Additional steps for a wizard-based migration ... 1512
Migration reports ... 1512
Clustered role and feature migrations that require extra steps ... 1512
Clustered DFS Replication migrations ... 1513
Clustered DHCP migrations ... 1513
Clustered DTC migrations ... 1514
Clustered File Server and Scale-out File Server migrations ... 1514
Clustered file server migrations ... 1514
Scale-out File Server migrations ... 1515
Clustered FSRM migrations ... 1515
Clustered Message Queuing (MSMQ) migrations ... 1515

Other Server migrations involving resource types not built into failover clusters 1516
Clustered virtual machine migrations 1516
Additional references 1517

Migration Between Two Multi-Node Clusters 1517
Overview of migration between two multi-node clusters 1518
Steps for creating a failover cluster 1518
Preparation 1519
After you create the failover cluster 1519
Steps for migrating clustered services and applications to a failover cluster running Windows Server 2012 1520
Steps for completing the transition from the old cluster to the new cluster 1521
Related references 1523

In-Place Migration for a Two-Node Cluster 1523
Overview of an in-place migration for a two-node cluster 1524
Steps for evicting a node and creating a new single-node Windows Server 2012 failover cluster 1525
Step 1: Evict one node from the old cluster, and perform a clean installation of Windows Server 2012 1525
Step 2: Create a single-node cluster and install other needed software 1525
Preparation 1525
After you create the failover cluster 1526
Steps for migrating clustered services and applications to the new cluster 1527
Steps for making existing data available to the new cluster and bringing it online 1528
Steps for adding the second node to the new cluster 1529
Related references 1531

Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard 1531
Supported operating systems 1532
Overview of the migration process 1532
Impact of the migration 1533
Required permissions 1533
Prepare to migrate 1533
Migrate the highly available virtual machines to the new failover cluster 1534
Verify a successful migration 1536
Related references 1537

Cluster Migrations Involving New Storage: Mount Points 1537
Additional references 1539

Additional References 1539

Secure Windows Server 2012 R2 and Windows Server 2012 1539
Assessment tools 1539

Security technologies 1540
Administration tools 1540
Support and enhance critical security needs 1540
Additional resources 1542

Security Tools to Administer Windows Server 2012 1542
Manage user accounts, groups, and credentials 1543
Modify or create new security principals 1544
Manage certificates and encryption 1545
Manage a CA and other Active Directory Certificate Services tasks 1547
Manage access to network resources 1547
Take ownership or securely delete files 1548
Manage security auditing and audit logs 1548
Analyze and manage security policies 1550
Analyze and manage computer processes and performance 1552
Diagnose, plan and remediate overall system security 1554
See also 1555

Manage Privacy 1556
Related Resources 1556

Managing Internet Communication and Privacy 1557
What this document includes 1557
Standard computer information sent by Internet-enabled features 1557
Types of features covered in this document 1557
Types of features not covered in this document 1558
Security and privacy basics that are beyond the scope of this document 1558
Resources about security basics 1559

Manage Privacy: Activation and Resulting Internet Communication 1559
Purposes of activation 1560
Overview: Activation in a managed environment 1560
Activation options with volume licensing 1560
How a computer communicates with sites on the Internet during activation 1562
Additional references 1563

Managing Privacy: Dynamic Update and Resulting Internet Communication 1563
Benefits and purposes of Dynamic Update 1564
Overview: Using Dynamic Update in a managed environment 1565
How Dynamic Update communicates with sites on the Internet 1565
Controlling Dynamic Update to limit the flow of information to and from the Internet 1567

Manage Privacy: Internet Explorer 10 and Resulting Internet Communication 1568
Benefits and purposes 1569
Enhanced Security Configuration 1569

Security-related features 1569
Resources for learning about security in Internet Explorer 10 1572
Learn about security and privacy settings 1572
Learn about mitigating the risks inherent in web-based applications and scripts 1572
Learn about Group Policy Objects that control configuration settings 1573
Learn about the Internet Explorer Administration Kit 1573
Procedures for controlling Internet Explorer 1574
Procedures for controlling web browsers 1574
Procedure to turn Internet Explorer Enhanced Security Configuration on or off 1575
Procedures for setting the security level to High for specific websites 1576

Manage Privacy: SmartScreen Filter and Resulting Internet Communication 1577
Benefits and purposes of SmartScreen 1577
Overview: Using SmartScreen Filter in a managed environment 1578
How SmartScreen Filter communicates with a web service on the Internet 1578
Controlling SmartScreen Filter to limit the flow of information to and from the Internet 1579
Additional references 1580

Managing Privacy: User Access Logging and Resulting Internet Communication 1580
Benefits and purposes of User Access Logging 1581
User and device-related data recorded with User Access Logging 1581
Viewing or changing settings that affect User Access Logging 1583
Additional references 1583

Manage Privacy: Windows Customer Experience Improvement Program and Resulting Internet Communication 1583
Purpose of the Windows Customer Experience Improvement Program 1584
Overview: Using the Windows Customer Experience Improvement Program in a managed environment 1584
How the Windows Customer Experience Improvement Program communicates with a site on the Internet 1584
Procedures for controlling the Windows Customer Experience Improvement Program 1585

Manage Privacy: Windows Defender and Resulting Internet Communication 1587
Benefits and purposes of Windows Defender and the online Microsoft Active Protection Service community 1588
Windows Defender 1588
The online Microsoft Active Protection Service community 1589
Overview: Using Windows Defender and information from the MAPS community in a managed environment 1589
How Windows Defender communicates with Internet sites without a MAPS membership 1589
How Windows Defender communicates with Internet sites when combined with MAPS 1591
Procedures for configuring Windows Defender 1592
Additional references 1593

Manage Privacy: Windows Error Reporting and Resulting Internet Communication 1594
Benefits and purposes of Windows Error Reporting and the Problem Reports and Solutions feature 1594
Consent levels in Windows Error Reporting 1595
Options for controlling Windows Error Reporting 1595
Overview: Using Windows Error Reporting and the Problem Reports and Solutions feature in a managed environment 1595
How Windows Error Reporting communicates with an Internet site 1596
Types of data collected 1597
Overview of the data that Windows Error Reporting collects 1597
Data collected from application errors 1598
Data collected from handwriting recognition errors 1598
Data collected from the Japanese Input Method Editor 1598
Data collected from Windows kernel failures 1599
Controlling Windows Error Reporting to prevent the flow of information to and from the Internet 1600
Using an answer file with an unattended installation 1600
Selected Group Policy settings for Windows Error Reporting 1600
Setting to redirect Windows Error Reporting to a server on your intranet 1600
Setting to control the degree of prompting that occurs before data is sent 1600
Setting to disable reporting handwriting recognition errors 1601
Setting for disabling Windows Error Reporting 1601
Procedures to configure Windows Error Reporting 1602
Additional references 1604

Managing Privacy: Windows Store and Resulting Internet Communication 1604
Benefits and Purpose of Accessing the Windows Store 1605
Windows Store App Feature Disclosure Requirements to Users 1606
Windows Store Access on Windows Server 1607
Controlling Windows Store Access Using Group Policy 1607
Additional references 1608

Managing Privacy: Using a Microsoft Account to Logon and Resulting Internet Communication 1608
Overview of Using a Microsoft Account to Logon to Windows 1608
Benefits and purpose of using a Microsoft Account to Login 1609
Microsoft Account and User Information Synchronized 1610
Safeguarding Microsoft Account Information Stored in the Cloud 1610
Additional references 1611

Appendix A: Resources for Learning About Automated Installation and Deployment 1611

Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category 1612

Overview of Group Policy settings listed under Internet Communication Management 1612
Controlling multiple Group Policy settings through the Restrict Internet Communications setting 1612
Group Policy settings that affect computer configuration 1613
Individual Group Policy settings that affect computer configuration 1613
Group Policy settings that affect user configuration 1616
Group Policy settings that affect user configuration 1616

Support Windows Server 2012 R2 and Windows Server 2012 1618
Best Practices Analyzer 1618
Windows Server 2012 Understand and Troubleshoot Guides 1619

Best Practices Analyzer 1619
What is Best Practices Analyzer? 1619
About content in this section 1620

Best Practices Analyzer for Remote Access 1620
More information about Remote Access 1620

Best Practices Analyzer for Remote Access: Prerequisites 1621
Best Practices Analyzer and prerequisite rules 1621
Topics in this section 1621

RRAS: The Remote Access server role should be configured in Multitenant mode 1621
Issue 1622
Impact 1622
Resolution 1622

RRAS: The Remote Access gateway should be configured with Multitenancy support 1622
Issue 1623
Impact 1623
Resolution 1623

RRAS: All Routing Domains should be enabled 1624
Issue 1624
Impact 1624
Resolution 1624

RRAS: All enabled Routing Domains should be available 1625
Issue 1625
Impact 1625
Resolution 1625

Best Practices Analyzer for Remote Access: Configuration (Section 1) 1626
Best Practices Analyzer and configuration rules 1626
Topics in this section 1626

RRAS: The inbound Certification Authority (CA) should be configured 1627
Issue 1627
Impact 1627
Resolution 1627

RRAS: At least one valid IKEv2 certificate should be present on the RRAS server 1628
Issue 1628
Impact 1628
Resolution 1628

RRAS: A valid CA certificate for the Remote Access Server certificate must be present in the TRCA certificate store 1629
Issue 1629
Impact 1629
Resolution 1630

RRAS: A valid CA certificate for the Site-to-Site VPN interface certificate must be present in the TRCA certificate store 1630
Issue 1630
Impact 1630
Resolution 1631

RRAS: The server certificate expires within 7 days 1631
Issue 1631
Impact 1631
Resolution 1631

RRAS: The CA certificate of the Remote Access Server certificate in the TRCA certificate store expires within 7 days 1632
Issue 1632
Impact 1633
Resolution 1633

RRAS: The certificate for the Site-to-Site VPN interface expires within 7 days 1633
Issue 1633
Impact 1634
Resolution 1634

RRAS: The CA certificate for the Site-to-Site VPN interface certificate expires within 7 days . 1634
Issue 1634
Impact 1635
Resolution 1635

RRAS: The CA certificate for the destination server of Site-to-Site VPN interface expires within 7 days 1635
Issue 1635

Impact 1636
Resolution 1636

RRAS: The Remote Access server certificate must have a public IP address for Alternate Subject Name 1636
Issue 1637
Impact 1637
Resolution 1637

RRAS: The Site-to-Site VPN interface name must match the Username 1637
Issue 1638
Impact 1638
Resolution 1638

RRAS: No two Site-to-Site VPN interfaces with PSK based authentication should have the same destination 1639
Issue 1639
Impact 1639
Resolution 1639

RRAS: For PSK authentication, the destination cannot be configured as a Fully Qualified Domain Name (FQDN) 1640
Issue 1641
Impact 1641
Resolution 1641

RRAS: The Site-to-Site VPN interface should be configured with a Source IP address 1641
Issue 1642
Impact 1642
Resolution 1642

RRAS: Custom policies configured for the Site-to-Site VPN interface should be a subset of Remote Access server global policies 1642
Issue 1643
Impact 1643
Resolution 1643

Best Practices Analyzer for Remote Access: Configuration (Section 2) 1643
Best Practices Analyzer and configuration rules 1643
Topics in this section 1644

RRAS: The number of ports available for the Routing Domain should not be less than the number of VPN and Site-to-Site VPN interfaces 1644
Issue 1645
Impact 1645
Resolution 1645

RRAS: A static pool should be configured for IPv4 address assignment to the VPN client 1645
Issue 1646
Impact 1646
Resolution 1646

RRAS: The static pool IPv4 addresses must be valid unicast IPv4 addresses 1647
Issue 1647
Impact 1647
Resolution 1647

RRAS: The VPN Tenant Name must be specified 1648
Issue 1648
Impact 1648
Resolution 1648

RRAS: The VPN Tenant Name for a Routing Domain must not be a subset of another Routing Domain's Tenant Name 1649
Issue 1649
Impact 1649
Resolution 1649

RRAS: The default route (IPv4 or IPv6) should not be advertised to the peers 1650
Issue 1650
Impact 1650
Resolution 1650

RRAS: The default route (IPv4 or IPv6) should not be accepted from the peers 1651
Issue 1651
Impact 1651
Resolution 1651

RRAS: The BGP peer IP address should not be assigned to a local network interface 1652
Issue 1652
Impact 1652
Resolution 1652

RRAS: A local global IPv6 address must be configured on the BGP Router 1653
Issue 1653
Impact 1653
Resolution 1653

RRAS: Multiple routes with different Next-Hop values and the same Destination prefix are configured 1654
Issue 1654
Impact 1654
Resolution 1654

RRAS: The BGP peer's Hold-Timer should not be set to the value 0 1655
Issue 1655
Impact 1655
Resolution 1655

RRAS: BGP peer's Hold-Timer should not be set to a very low value 1656
Issue 1656
Impact 1656
Resolution 1656

RRAS: All the ingress route advertisements should not be dropped because of a routing policy 1657
Issue 1657
Impact 1657
Resolution 1657

RRAS: All the egress route advertisements should not be dropped because of a routing policy 1658
Issue 1658
Impact 1658
Resolution 1658

RRAS: BGP peers should not be configured for manual (passive) peering mode 1659
Issue 1659
Impact 1659
Resolution 1659

Best Practices Analyzer for Remote Access: Configuration (Section 3) 1660
Best Practices Analyzer and configuration rules 1660
Topics in this section 1660

RRAS: For BGP Peering over IPv6 addresses, IPv4 routes should not be configured for advertisement 1661
Issue 1661
Impact 1661
Resolution 1661

RRAS: IdleHoldTimer should not be set to a high value (> 10 sec) 1662
Issue 1662
Impact 1663
Resolution 1663

RRAS: Max Prefix policy should be configured for all BGP Peers 1663
Issue 1663
Impact 1663
Resolution 1664

RRAS: The total number of prefixes learned is in proximity of the Maximum Allowed Prefixes 1664
Issue 1664
Impact 1664
Resolution 1665

RRAS: A triggering route must be configured on the Site-to-Site VPN interface for the BGP peers 1665
Issue 1665
Impact 1666
Resolution 1666

RRAS: The Site-to-Site VPN triggering route should be a specific address (/32 for IPv4 address, /128 for IPv6 address) 1666
Issue 1667
Impact 1667
Resolution 1667

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking BGP traffic 1668
Issue 1668
Impact 1668
Resolution 1668

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VSID interface traffic 1669
Issue 1669
Impact 1670
Resolution 1670

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VPN client traffic 1670
Issue 1671
Impact 1671
Resolution 1671

RRAS: The VPN static address pool should be configured as custom networks on the BGP router 1672
Issue 1672
Impact 1672
Resolution 1672

Best Practices Analyzer for Remote Access: Performance and Operation 1673
Best Practices Analyzer performance and operation rules 1673
Topics in this section 1673

RRAS: Rate-limiting (Tx/Rx BandwidthKbps) should be set to a value as per the network requirements 1673
Issue 1674

Impact ... 1674
Resolution ... 1674

RRAS: Rate-limiting parameters (Tx/Rx BandwidthKbps) should not have a significant difference in the values ... 1674
Issue ... 1675
Impact ... 1675
Resolution ... 1675

RRAS: The CapacityKbps parameter should be configured with a value as per the network requirements ... 1676
Issue ... 1676
Impact ... 1676
Resolution ... 1676

RRAS: The routes being advertised to the peers must be locally resolvable ... 1677
Issue ... 1677
Impact ... 1677
Resolution ... 1677

Work Folders Best Practices Analyzer ... 1678

The Windows Sync Share service should be set to start automatically ... 1678
Issue ... 1678
Impact ... 1678
Resolution ... 1678
See also ... 1679

Work Folders should be installed on all nodes of the failover cluster ... 1679
Issue ... 1679
Impact ... 1679
Resolution ... 1679
See also ... 1680

All nodes in a failover cluster should be reachable to Work Folders ... 1680
Issue ... 1680
Impact ... 1680
Resolution ... 1680
See also ... 1681

The Work Folders server should be domain joined ... 1681
Issue ... 1681
Impact ... 1681
Resolution ... 1681
See also ... 1682

The sync share should be located in a valid folder 1682
Issue 1682
Impact 1682
Resolution 1682
See also 1682

A current SSL certificate should be configured for this Work Folders server 1683
Issue 1683
Impact 1683
Resolution 1683
See also 1683

Windows Firewall should open port 80 and 443 for Work Folders 1683
Issue 1684
Impact 1684
Resolution 1684
See also 1684

A staging area should exist for the sync share 1684
Issue 1685
Impact 1685
Resolution 1685
See also 1685

Best Practices Analyzer for Web Application Proxy 1685
More information about Web Application Proxy 1686
Topics in this section 1686

Web Application Proxy must be configured before it is used 1686
Issue 1687
Impact 1687
Resolution 1687

Web Application Proxy: The external and backend server URLs are different and URL translation is disabled 1688
Issue 1688
Impact 1688
Resolution 1688

Web Application Proxy: The service is not configured to run automatically 1690
Issue 1691
Impact 1691
Resolution 1691

Web Application Proxy: The AD FS Proxy service is not configured to run automatically 1691
Issue 1692

Impact 1692
Resolution 1692

Web Application Proxy: This server is not included in the ConnectedServersName list 1693
Issue 1693
Impact 1693
Resolution 1693

Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high 1694
Issue 1694
Impact 1694
Resolution 1694

Web Application Proxy: Application is using an external certificate that is not yet valid 1695
Issue 1696
Impact 1696
Resolution 1696

Web Application Proxy: Application is using an external certificate that is about to expire 1698
Issue 1698
Impact 1699
Resolution 1699

Web Application Proxy: Application is using an external certificate that has no private key 1701
Issue 1701
Impact 1701
Resolution 1702

Web Application Proxy: Application is using an external certificate that has expired 1704
Issue 1704
Impact 1704
Resolution 1704

Web Application Proxy: Application is configured to use an external certificate that is not present on this server 1707
Issue 1707
Impact 1707
Resolution 1707

Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows authentication but the server is not joined to a domain 1708
Issue 1709
Impact 1709
Resolution 1709

Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed 1710
Issue 1711
Impact 1711
Resolution 1711

Server Roles and Technologies in Windows Server 2012 R2 and Windows Server 2012 1712

Active Directory 1715

What's New in Active Directory in Windows Server 2012 R2 1715

Active Directory Certificate Services Overview 1716
Role description 1717
Practical applications 1717
New and changed functionality 1717
Server Manager information 1717
See also 1718

What's New in Certificate Services in Windows Server 2012 R2 1720
Role/Feature description 1720
New and changed functionality 1720
Policy Module support for the Network Device Enrollment Service 1720
TPM key attestation 1721
Windows PowerShell for Certificate Services 1722

What's New in Certificate Services in Windows Server 2012 1722
Role description 1723
New and changed functionality 1723
Integration with Server Manager 1723
Deployment and management capabilities from Windows PowerShell® 1724
All AD CS role services run on any version 1724
All AD CS role services can be run on Server Core 1725
Support for key-based renewal 1725
Certificate Template Compatibility 1725
Support for certificate renewal with same key 1726
Support for Internationalized Domain Names 1727
Increased security enabled by default on the CA role service 1728
AD DS Site Awareness for AD CS and PKI Clients 1729
Group-protected PFX format 1729
Certificate lifecycle notifications 1730
CA private keys are included in the System State Backup image 1731
See also 1731

Protecting Against Weak Cryptographic Algorithms 1732

What does this software update do? ... 1733
How to configure policies for blocking cryptographic algorithms ... 1734
Updating client registry settings through Group Policy ... 1737
Examples ... 1744
Setting logging directory and enabling logging ... 1745
Logging with Audit only Mode ... 1745

Certification Authority Guidance ... 1746
Plan for PKI ... 1746
Use an HSM ... 1746
Consider a CAPolicy.inf file ... 1747
Select CA configuration settings ... 1748
Select setup type ... 1748
Choose CA type ... 1749
Designate a root CA ... 1749
Subordinate CAs ... 1750
Store a private key ... 1750
Locate an existing key ... 1750
Locate an existing certificate ... 1751
Select cryptographic options ... 1751
Establish a CA name ... 1753
Obtain a certificate request ... 1754
Verify the validity period ... 1755
Choose a CA database ... 1755
Configure the CA ... 1755
Publish the AIA extension ... 1757
Use the interface to publish the AIA extension ... 1758
Use Windows PowerShell to publish the AIA extension ... 1760
Use certutil to publish the AIA extension ... 1760
Publish the CDP extension ... 1760
Use the interface to publish the CDP extension ... 1762
Use Windows PowerShell to publish the CDP extension ... 1764
Use certutil to publish the CDP extension ... 1764
Verify the configuration ... 1765
Related content ... 1765

Configure Trusted Roots and Disallowed Certificates ... 1765
Certificates and trust ... 1766
Software update description ... 1767
Configuration options ... 1768
Configure a file or web server to download the CTL files ... 1769
Redirect the Microsoft Automatic Update URL for a disconnected environment ... 1770
Redirect the Microsoft Automatic Update URL for untrusted CTLs only ... 1772
Use a subset of the trusted CTLs ... 1774

Registry settings modified 1776
New Certutil Options 1777
Potential errors with Certutil -SyncWithWU 1778
Related content 1779

Certification Authority Web Enrollment Guidance 1779
CA for Web Enrollment 1780
Web Enrollment Configuration 1781
Use the CA Web Enrollment pages 1781
Request a basic certificate 1782
Request a certificate with advanced options 1782
Check a pending certificate request 1783
Retrieve the CA certificate 1783
Retrieve the current base and delta CRLs 1784
Submit a certificate request by using a PKCS #10 file or a PKCS #7 file 1784
Related content 1785

Certificate Enrollment Web Service Guidance 1785
Installation requirements 1786
Configure a CA for the Certificate Enrollment Web Service 1787
Set the authentication type for Certificate Enrollment Web Service 1788
Allow key-based renewal for Certificate Enrollment Web Service 1788
Configure a Service Account 1788
Select a Server Certificate 1791
Complete Certificate Enrollment Web Services Configuration 1792
Related content 1793

Certificate Enrollment Policy Web Service Guidance 1793
Set the authentication type for Certificate Enrollment Policy Web Service 1793
Determine whether to enable key-based renewal for Certificate Enrollment Policy Web Service 1794
Server Certificate 1794
Certificate Enrollment Policy Web Service Configuration 1794
Related content 1797

Network Device Enrollment Service Guidance 1797
NDES configuration settings 1798
Configure a service account for NDES 1798
Select a CA for NDES 1800
Set RA information 1800
Configure cryptography for NDES 1800
Complete NDES configuration 1801
Related content 1801

Using a Policy Module with the Network Device Enrollment Service 1801

How the policy module works ... 1802
Deployment options for the Network Device Enrollment Service and a policy module ... 1804
Deployment in a separate forest ... 1804
Deployment in an isolated network ... 1806
Deployment on an internal domain ... 1807
How to install and uninstall the policy module ... 1807
Related content ... 1809

Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy ... 1809
In this guide ... 1809
Test lab overview ... 1810
Hardware and software requirements ... 1810
Step 1: Complete the Base TLG Configuration ... 1811
Step 2: Configure ORCA1 ... 1811
Step 3: Configure APP1 to distribute certificates and CRLs ... 1820
Step 4: Configure APP1 as an Enterprise Subordinate CA ... 1821
Step 5: Configure computer certificate autoenrollment ... 1827
Step 6: Configuring SSL for APP1 ... 1828
See Also ... 1830

Test Lab Guide: Demonstrating Certificate Key-Based Renewal ... 1830
In this guide ... 1830
Test lab overview ... 1830
Hardware and software requirements ... 1831
Step 1: Complete the Base Configuration Test Lab ... 1832
Step 2: Complete the Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy ... 1832
Step 3: Configure the CEP1 server ... 1832
Step 4: Configure the CES1 server ... 1837
Step 5: Prepare an appropriate certificate template ... 1843
Step 6: Configure WEB1 ... 1844
Step 7: Obtain a certificate and test automatic renewal ... 1846
See Also ... 1849

Active Directory Domain Services Overview ... 1849
Requirements for running Active Directory Domain Services ... 1850
Running Active Directory Domain Services ... 1852
How do I deploy and configure this role by using Windows PowerShell? ... 1852
How do I deploy and configure this role in a multi-server environment? ... 1852
How can I run this role on virtual machines? ... 1852
Security considerations for running this role ... 1852
Special considerations for managing this role remotely ... 1852
Special considerations for managing the role on the Server Core installation option ... 1852
Role services for Active Directory Domain Services ... 1853
Additional references ... 1853

What's New in Active Directory Domain Services (AD DS) ... 1854
Virtualization that just works ... 1855
Rapid deployment with cloning ... 1855
Safer virtualization of domain controllers ... 1855
Simplified deployment and upgrade preparation ... 1856
Simplified management ... 1857
Dynamic Access Control ... 1857
DirectAccess Offline Domain Join ... 1858
Active Directory Federation Services (AD FS) ... 1859
Windows PowerShell History Viewer ... 1859
Active Directory Recycle Bin User Interface ... 1859
Fine-Grained Password Policy User Interface ... 1860
Active Directory Replication and Topology Windows PowerShell cmdlets ... 1860
Active Directory Based Activation (AD BA) ... 1861
Group Managed Service Accounts (gMSA) ... 1862
AD DS Platform Changes ... 1863
AD DS Claims in AD FS ... 1863
Relative ID (RID) Improvements ... 1863
Deferred Index Creation ... 1864
Kerberos Enhancements ... 1864
Kerberos Constrained Delegation across domains ... 1865
Flexible Authentication Secure Tunneling (FAST) ... 1866

Active Directory Domain Services (AD DS) Virtualization ... 1866

Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100) ... 1867
Safe virtualization of domain controllers ... 1867
How do these virtualization safeguards work? ... 1869
Virtualized domain controller cloning ... 1871
Scenarios that benefit from virtual domain controller cloning ... 1871
Clear separation of responsibilities ... 1872
How does virtual domain controller cloning work? ... 1872
Cloning components ... 1873
Deployment scenarios ... 1875
Steps for deploying a clone virtualized domain controller ... 1876
Prerequisites ... 1876
Step 1: Grant the source virtualized domain controller the permission to be cloned ... 1878
Step 2: Run Get-ADDCCloningExcludedApplicationList cmdlet ... 1878
Step 3: Run New-ADDCCloneConfigFile ... 1879
Running New-ADDCCloneConfigFile in offline mode ... 1881
Step 4: Export and then import the virtual machine of the source domain controller ... 1882
Troubleshooting ... 1884

Virtualized Domain Controller Technical Reference (Level 300) ... 1885

Virtualized Domain Controller Architecture 1885
Virtualized domain controller cloning architecture 1885
Overview 1885
Cloning Detailed Processing 1886
Virtualized domain controller safe restore architecture 1891
Overview 1891
Safe Restore Detailed Processing 1892

Virtualized Domain Controller Deployment and Configuration 1894
Installation Considerations 1895
Platform Requirements 1895
Critical Caveats 1896
Virtualized Domain Controller Cloning 1896
Step 1 - Validate the Hypervisor 1898
Step 2 - Verify the PDCE FSMO role 1899
Active Directory Users and Computers Method 1900
Windows PowerShell Method 1900
Step 3 - Authorize a Source DC 1901
Active Directory Administrative Center Method 1901
Windows PowerShell Method 1901
Rebuilding Default Permissions 1901
Active Directory Administrative Center Method 1902
Windows PowerShell Method 1902
Step 4 - Remove Incompatible applications or services (if not using CustomDCCloneAllowList.xml) 1902
Step 5 - Create DCCloneConfig.xml 1903
Generating with New-ADDCCloneConfigFile 1903
Step 6 - Take the Source Domain Controller Offline 1906
Graphical Method 1906
Windows PowerShell Method 1906
Step 7 - Copy Disks 1907
Manually Copying Disks 1907
Hyper-V Manager Method 1907
Windows PowerShell Method 1908
Exporting the VM 1910
Hyper-V Manager Method 1910
Windows PowerShell Method 1910
Exporting merged disks, using Hyper-V 1911
Hyper-V Manager Method 1911
Windows PowerShell Method 1911
Adding XML to the Offline System Disk 1912
Windows Explorer Method 1913
Windows PowerShell Method 1915

Step 8 - Create the New Virtual Machine 1916
Associating a New VM with Copied Disks 1916
Hyper-V Manager Method 1916
Windows PowerShell Method 1917
Import VM 1917
Hyper-V Manager Method 1917
Windows PowerShell Method 1920
Step 9 - Clone the New Virtual Machine 1920
Virtualization safeguards 1921
Validate the Hypervisor 1921
Validate the Replication Topology 1922
Writable Domain Controller Availability 1922
Simultaneous Restore 1922
Post-Snapshot Replication 1923
Windows PowerShell Snapshot Cmdlets 1925

Virtualized Domain Controller Troubleshooting 1925
Introduction 1925
Troubleshooting virtualized domain controller cloning 1926
Tools for Troubleshooting 1928
Logging Options 1928
Tools and Commands for Troubleshooting Domain Controller Configuration 1928
General Methodology for Troubleshooting Domain Controller Cloning 1928
Disabling DSRM Boot 1929
Removing DSRM with Msconfig.exe 1929
Removing DSRM with Bcdedit.exe 1930
Server Core and the Event Log 1930
Troubleshooting Specific Problems 1931
Events 1931
Directory Services Event Log 1931
Error Messages 1961
Known Issues and Support Scenarios 1962
Advanced Troubleshooting 1968
Cloning a Domain Controller 1968
Directory Services Event Log 1968
System Event Log 1975
DCPROMO.LOG 1978
Active Directory Web Services Event Log 1992
DNS Server Event Log 1993
File Replication Service Event Log 1995
DFS Replication Event Log 1998
Troubleshooting virtualized domain controller safe restore 2000
Tools for Troubleshooting 2000

Logging Options 2000
Tools and Commands for Troubleshooting Domain Controller Configuration 2001
General Methodology for Troubleshooting Domain Controller Safe Restore 2001
Troubleshooting Specific Problems 2002
Events 2002
Error Messages 2011
Known Issues and Support Scenarios 2011
Advanced Troubleshooting 2012
Restoring a Domain Controller that Replicates SYSVOL Using DFSR 2013
Directory Services Event Log 2013
System Event Log 2019
Application Event Log 2020
DFS Replication Event Log 2021
Restoring a Domain Controller that Replicates SYSVOL Using FRS 2023
File Replication Service Event Log 2023
Application Event Log 2025

Virtualized Domain Controller Technical Reference Appendix 2028
Terminology 2028
FixVDCPermissions.ps1 2028

Virtualized Domain Controller Additional Resources 2029

Virtualized Domain Controller Cloning Test Guidance for Application Vendors 2030
Overview of virtualized DC cloning 2030
What is interesting for Application Vendors? 2031
CustomDCCloneAllowList.xml 2031
Distributed System Interactions 2031
Additional scenarios suggested for testing 2032
Cloning Failure 2032
PDC emulator cloning 2032
Writable versus read-only domain controllers 2032

Support for using Hyper-V Replica for virtualized domain controllers 2032
Windows Server 2012 domain controllers required 2033
Supported and unsupported scenarios 2034
Windows Server 2012 2034
Windows Server 2008 R2 and earlier versions 2035

Deploy Active Directory Domain Services (AD DS) in Your Enterprise 2036

What's New in Active Directory Domain Services Installation and Removal 2036
The Active Directory Domain Services Configuration Wizard 2037
Adprep.exe integration 2038
Group membership check against Windows Server 2003 operations master roles 2039

Syntax for Adprep in Windows Server 2012 2040
Running adprep using smartcard 2040
Adprep /domainprep /gpprep command is not run automatically 2040
AD DS installation prerequisite validation 2041
System requirements 2041
Known issues 2041
Active Directory Domain Services installation hangs if an error prevents critical replication 2043
Active Directory Domain Services Configuration Wizard warns when a non-normalized DNS name is specified 2045

Upgrade Domain Controllers to Windows Server 2012 2046
What's new in Windows Server 2012? 2046
Automatic Maintenance and changes to restart behavior after updates are applied by Windows Update 2047
What’s new in AD DS in Windows Server 2012? 2047
Automatic Maintenance and changes to restart behavior after updates are applied by Windows Update 2049
AD DS server role installation changes 2050
Deprecated features and behavior changes related to AD DS in Windows Server 2012 2051
Operating system requirements 2052
Disk space requirements for upgrading domain controllers to Windows Server 2012 2052
Available SKUs 2053
Windows client and Windows Server operating systems that are supported to join Windows Server 2012 domains 2053
Supported in-place upgrade paths 2054
Functional level features and requirements 2054
AD DS interoperability with other server roles and Windows operating systems 2056
Operations master roles 2056
Virtualizing domain controllers that run Windows Server 2012 2057
Administration of Windows Server 2012 servers 2057
Application compatibility 2057
Known issues 2061
Domain controller upgrade workflow 2066
See Also 2067

AD DS Simplified Administration 2067
Technical Overview 2069
ADPREP Integration 2069
Server Manager AD DS Integration 2069
Active Directory Administrative Center Recycle Bin 2070
Active Directory Administrative Center Fine-Grained Password Policy 2071
Active Directory Administrative Center Windows PowerShell History Viewer 2072
AD Replication Windows PowerShell 2072

RID Management and Issuance Improvements 2073
New AD DS Deployment Architecture 2073
AD DS Role Deployment and Management Architecture 2073
ADPrep and Prerequisite Checking Architecture 2074
ADPrep Executables, DLLs, LDFs, files 2075
Prerequisite Checking 2076
Prerequisite Windows PowerShell 2076
Prerequisite Tests 2076

Simplified Administration Appendix 2078
Server Manager Add Servers Dialog (Active Directory) 2079
Server Manager Remote Server Status 2079
Windows PowerShell Module Loading 2079
RID Issuance Hotfixes for Previous Operating Systems 2080
Ntdsutil.exe Install from Media Changes 2081

Install Active Directory Domain Services (Level 100) 2081
Credential requirements to run Adprep.exe and install Active Directory Domain Services 2082
Installing AD DS by Using Windows PowerShell 2082
ADDSDeployment cmdlet arguments 2084
Specifying Windows PowerShell Credentials 2093
Using test cmdlets 2093
Installing a new forest root domain using Windows PowerShell 2094
Installing a new child or tree domain using Windows PowerShell 2095
Installing an additional (replica) domain controller using Windows PowerShell 2095
Performing a staged RODC installation using Windows PowerShell 2096
Installing AD DS by using Server Manager 2097
Creating server pools 2097
Installing AD DS 2098
Performing a Staged RODC Installation using the Graphical User Interface 2102
See Also 2106

Install a New Windows Server 2012 Active Directory Forest (Level 200) 2106
Active Directory Domain Services Simplified Administration 2106
What Is AD DS Simplified Administration? 2106
Purpose and Benefits 2107
Technical Overview 2107
What You Should Know Before You Begin 2107
Functional Descriptions 2108
AD DS Role Installation 2108
AD DS Role Configuration 2109
Prerequisite Checking 2110
Deploying a Forest with Server Manager 2110
Server Manager AD DS Role Installation Process 2110

Server Pool and Add Roles 2111
Installation Type 2115
Server Selection 2116
Server Roles and Features 2117
Active Directory Domain Services 2119
Confirmation 2120
Results 2121
Promote to Domain Controller 2123
Uninstalling/Disabling 2123
Create an AD DS Forest Root Domain with Server Manager 2124
Deployment Configuration 2125
Domain Controller Options 2126
DNS Options and DNS Delegation Credentials 2127
Additional Options 2128
Paths 2129
Review Options and View Script 2130
Prerequisites Check 2132
Installation 2133
Results 2134
Deploying a Forest with Windows PowerShell 2134
Windows PowerShell AD DS Role Installation Process 2134
Create an AD DS Forest Root Domain with Windows PowerShell 2138
See Also 2142

Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200) 2143
Upgrade and Replica Workflow 2143
Upgrade and Replica Windows PowerShell 2144
Deployment 2146
Deployment Configuration 2146
Domain Controller Options 2149
DNS Options and DNS Delegation Credentials 2152
Additional Options 2153
Paths 2155
Preparation Options 2156
Review Options and View Script 2158
Prerequisites Check 2160
Installation 2161
Results 2165

Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200) 2165
Child and Tree Domain Workflow 2166
Child and Tree Domain Windows PowerShell 2166
Deployment 2167
Deployment Configuration 2167

Domain Controller Options 2170
DNS Options and DNS Delegation Credentials 2173
Additional Options 2174
Paths 2175
Review Options and View Script 2176
Prerequisites Check 2179
Installation 2180
Results 2182

Install a Windows Server 2012 Active Directory Read-Only Domain Controller (RODC) (Level 200) 2182
Stage RODC Workflow 2182
Stage RODC Windows PowerShell 2183
Attach RODC Workflow 2184
Attach RODC Windows PowerShell 2185
Staging 2186
Welcome 2187
Network Credentials 2188
Specify the Computer Name 2189
Select a Site 2190
Additional Domain Controller Options 2191
Specify the Password Replication Policy 2192
Delegation of RODC Installation and Administration 2193
Summary 2194
Creation 2195
Attaching 2197
Deployment Configuration 2197
Domain Controller Options 2198
Additional Options 2201
Paths 2203
Review Options and View Script 2204
Prerequisites Check 2206
Installation 2207
Results 2209
RODC without Staging Workflow 2209
RODC without Staging Windows PowerShell 2210
RODC without Staging Deployment 2212
Deployment Configuration 2212
Domain Controller Options 2213
RODC Options 2215
Additional Options 2217
Paths 2218
Preparation Options 2219

Review Options and View Script ... 2220
Prerequisites Check ... 2223
Installation ... 2224
Results ... 2226

Remove Active Directory Domain Services (Level 100) ... 2226
Remove AD DS using Windows PowerShell ... 2227
Using Test-ADDSDomainControllerUninstallation cmdlet ... 2228
Remove AD DS using the Remove Roles Wizard in Server Manager ... 2229
See Also ... 2231

Demoting Domain Controllers and Domains (Level 200) ... 2231
AD DS Removal Workflow ... 2232
Demotion and Role Removal Windows PowerShell ... 2232
Demote ... 2233
Remove Roles and Features ... 2233
Server Selection ... 2235
Server Roles and Features ... 2236
Credentials ... 2238
Warnings ... 2240
Removal Options ... 2241
New Administrator Password ... 2242
Confirmation ... 2243
Demotion ... 2244
Results ... 2247

AD DS Installation and Removal Wizard Page Descriptions ... 2247
Deployment Configuration ... 2248
Domain Controller Options ... 2251
DNS Options ... 2254
RODC Options ... 2255
Additional Options ... 2257
Paths ... 2258
Preparation Options ... 2260
Review Options ... 2261
Prerequisites Check ... 2262
Results ... 2263
Role Removal credentials ... 2264
AD DS Removal Options and Warnings ... 2265
New Administrator Password ... 2267
Review Options ... 2268

Windows Server 2012: Changes Made by Adprep.exe ... 2269
See Also ... 2269

Windows Server 2012: Schema Updates ... 2269
Schema Updates ... 2270
Sch48.ldf ... 2270
Sch49.ldf ... 2290
Sch50.ldf ... 2314
Sch51.ldf ... 2328
Sch52.ldf ... 2340
Sch53.ldf ... 2355
Sch54.ldf ... 2356
Sch55.ldf ... 2356
Sch56.ldf ... 2357

Windows Server 2012: Read-Only Domain Controller Updates ... 2358

Windows Server 2012: Domain-Wide Updates ... 2358

Windows Server 2012: Forest-Wide Updates ... 2359

Troubleshooting Domain Controller Deployment ... 2384
Introduction to Troubleshooting ... 2385
Troubleshooting Options ... 2385
Logging Options ... 2385
Tools and Commands for Troubleshooting Domain Controller Configuration ... 2386
General Methodology for Troubleshooting Domain Controller Configuration ... 2386
Troubleshooting Specific Problems ... 2388
Events and Error Messages ... 2388
Known/Likely Issues and Support Scenarios ... 2397

Active Directory Administrative Center Enhancements ... 2409

Introduction to Active Directory Administrative Center Enhancements (Level 100) ... 2409
Active Directory Recycle Bin ... 2409
Active Directory Recycle Bin step-by-step ... 2410
Step 1: Raise the forest functional level ... 2411
Step 2: Enable Recycle Bin ... 2411
Step 3: Create test users, group and organizational unit ... 2412
Step 4: Restore deleted objects ... 2414
Fine-Grained Password Policy ... 2416
Fine-Grained Password Policy step-by-step ... 2416
Step 1: Raise the domain functional level ... 2417
Step 2: Create test users, group, and organizational unit ... 2417
Step 3: Create a new fine-grained password policy ... 2417
Step 4: View a resultant set of policies for a user ... 2419
Step 5: Edit a fine-grained password policy ... 2419
Step 6: Delete a fine-grained password policy ... 2420

Windows PowerShell History Viewer 2420
Windows PowerShell History Viewer step-by-step 2421
See Also 2421

Advanced AD DS Management Using Active Directory Administrative Center (Level 200) 2422
Active Directory Administrative Center Architecture 2422
ADPrep Executables, DLLs 2422
Enabling and Managing the Active Directory Recycle Bin Using Active Directory Administrative Center 2423
Capabilities 2423
Limitations 2423
Enabling Active Directory Recycle Bin using Active Directory Administrative Center 2424
Managing Active Directory Recycle Bin using Active Directory Administrative Center 2426
Storage and Filtering 2427
Restoration 2429
Filtering 2429
Single Object 2431
Multiple Peer Objects 2432
Multiple Parent and Child Objects 2433
Server-side Filtering 2436
Configuring and Managing Fine-Grained Password Policies Using Active Directory Administrative Center 2437
Configuring Fine-Grained Password Policies 2437
Managing Fine-Grained Password Policies 2437
Using the Active Directory Administrative Center Windows PowerShell History Viewer 2441
Troubleshooting AD DS Management 2444
Introduction to Troubleshooting 2444
Troubleshooting Options 2445
Logging Options 2445
Known/Likely Issues and Support Scenarios 2450
See Also 2450

Active Directory Replication and Topology Management Using Windows PowerShell 2451

Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100) 2451
Installing the Active Directory Module for Windows PowerShell 2451
Scenarios for testing Windows PowerShell for Active Directory replication and topology management cmdlets 2452
Lab Requirements 2452
View domain controllers and their sites 2452
Manage replication topology 2453
Verification 2454
View replication status information 2454

See Also 2455

Advanced Active Directory Replication and Topology Management Using Windows PowerShell (Level 200) 2455
Introduction 2456
Replication and Metadata 2457
Get-ADReplicationAttributeMetadata 2458
Get-ADReplicationPartnerMetadata 2460
Get-ADReplicationFailure 2462
Get-ADReplicationQueueOperation and Get-ADReplicationUpToDatenessVectorTable 2463
Sync-ADObject 2463
Topology 2463
See Also 2465

Managing RID Issuance 2465
Managing RID Issuance 2465
Periodic Consumption Warnings 2466
RID Pool Invalidation Events 2467
RID Block Size Limit 2467
Global RID Space Size Unlock 2467
Important Caveats 2468
Implementing Unlocked Global RID space 2468
RID Ceiling Enforcement 2469
Removing the Ceiling Block 2470
Other RID Fixes 2472
Unfixed RID Issues 2472
RID Fixes for earlier versions of Windows Server 2472
Troubleshooting RID Issuance 2472
Introduction to Troubleshooting 2472
Troubleshooting Options 2473
Logging Options 2473
Utilities and Commands for Troubleshooting 2473
General Methodology for Troubleshooting Domain Controller Configuration 2474
Troubleshooting Specific Problems 2474
See Also 2478

Dynamic Access Control: Scenario Overview 2478
In this scenario 2478
Dynamic Access Control Content Roadmap 2478
See also 2487

Scenario: Central Access Policy 2487
In this scenario 2490
Roles and features included in this scenario 2490

Plan for a Central Access Policy Deployment 2491
Process to map a business request to a central access policy 2492
Understand and translate business intent 2492
Express access policy in Windows Server 2012 constructs 2492
Determine the user groups, resource properties and claim types 2492
Determine the servers where this policy should be applied to 2493
Planning Guidelines for Deploying Central Access Policies 2493
Using Security Groups for Dynamic Access Control 2493
Using security groups to limit access to data 2493
Using conditional expressions to reduce complexity of security groups 2494
Using User Claims 2494
Operations to enable user claims 2495
Enable the domain controllers to provide claims and compound authentication on request 2495
Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies 2496
Using Device Claims and Device Security Groups 2496
Considerations for using static device claims 2496
Operations to enable device claims 2496
Enable the Windows 8 devices in domain to request claims and compound authentication 2496
Enable the Windows 8 devices to request claims and compound authentication using custom policy 2497
Enable the Windows 8 device to receive compound authentication 2497
Configuring central access policies with different options 2498
Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs 2498
Configuring forest root DCs 2498
Configuring domains which provide claims and compound authentication 2498
Configuring devices to request claims and compound authentication 2499
Configuring resources to receive compound authentication 2499
Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites 2499
Configuring forest root DCs 2499
Configuring domains which provides claims and compound authentication 2500
Configuring file servers to request claims on the behalf of users 2500
Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded 2500
Considerations for using smartcards for Central Access Polices 2500
Best Practices for Deploying Central Access Policies 2501
Delegating of administration for Dynamic Access Control 2501
Exception Mechanisms for Planning Central Access Policies 2502

Tools for Deployment 2503
Appendix: Deployment Configurations for Central Access Policies 2504

Deploy a Central Access Policy (Demonstration Steps) 2505
Set up a test environment 2505
Plan: Identify the need for policy and the configuration required for deployment 2506
Implement: Configure the components and policy 2507
Deploy the central access policy 2513
Maintain: Change and stage the policy 2515
Next Steps 2518

Scenario: File Access Auditing 2518
In this scenario 2519
Roles and features included in this scenario 2520

Plan for File Access Auditing 2520

Deploy Security Auditing with Central Audit Policies (Demonstration Steps) 2522
Configure global object access policy 2523
Update Group Policy settings 2523
Verify that the global object access policy has been applied 2524
See also 2524

Scenario: Access-Denied Assistance 2524
Scenario description 2525
In this scenario 2525
Practical applications 2526
Features included in this scenario 2526

Deploy Access-Denied Assistance (Demonstration Steps) 2526
Step 1: Configure access-denied assistance 2526
Step 2: Configure the email notification settings 2530
Step 3: Verify that access-denied assistance is configured correctly 2531
See also 2531

Scenario: Classification-Based Encryption for Office Documents 2531
Scenario description 2531
In this scenario 2532
Roles and features included in this scenario 2532

Deploy Encryption of Office Files (Demonstration Steps) 2534
Step 1: Enable resource properties 2535
Step 2: Create classification rules 2535
Step 3: Use file management tasks to automatically protect documents with AD RMS 2538
Step 4: View the results 2539
Step 5: Verify protection with AD RMS 2540

Scenario: Get Insight into Your Data by Using Classification 2540
Scenario description 2540
In this scenario 2541
Practical applications 2541
Features included in this scenario 2541

Deploy Automatic File Classification (Demonstration Steps) 2542
Step 1: Create resource property definitions 2542
Step 2: Create a string content classification rule 2543
Step 3: Create a regular expression content classification rule 2544
Step 4: Verify that the files are classified correctly 2545
See also 2546

Scenario: Implement Retention of Information on File Servers 2546
Scenario description 2546
In this scenario 2546
Features included in this scenario 2547

Deploy Implementing Retention of Information on File Servers (Demonstration Steps) 2547
Prerequisites 2547
Step 1: Create resource property definitions 2548
Step 2: Configure notifications 2548
Step 3: Create a file management task 2549
Step 4: Classify a file manually 2550
See also 2551

Deploy Claims Across Forests 2551
Claim transformation rules 2552
Linking claim transformation policies to forests 2552
In this scenario 2552
Roles and features included in this scenario 2552

Deploy Claims Across Forests (Demonstration Steps) 2553
Scenario overview 2553
Set up the prerequisites and the test environment 2553
Set up claims transformation on trusted forest (Adatum) 2555
Create a claims transformation policy in Adatum 2555
Set a claims transformation link on Adatum’s trust domain object 2555
Set up claims transformation in the trusting forest (Contoso) 2556
Create a claims transformation policy in Contoso 2556
Set a claims transformation link on Contoso’s trust domain object 2556
Validate the scenario 2557
Additional scenarios for claims transformation policies 2557
See also 2559

Claims Transformation Rules Language 2559
Tools for authoring claims transformation policies 2559
Active Directory claims transformation rules language 2560
Syntax overview 2560
Runtime operation 2561
Special rules semantics 2562
Security considerations 2563
Other language considerations 2563
Sample transformation rules 2564
Examples of rules parser errors 2564
Language terminals 2566
Language syntax 2567

Appendix A: Dynamic Access Control Glossary 2569
See Also 2571

Appendix B: Setting Up the Test Environment 2571
Prerequisites 2572
Build the test lab virtual machines 2572
Install the Hyper-V role 2572
Create an internal virtual network 2573
Build the domain controller 2573
Build the file server and AD RMS server (FILE1) 2575
Install File Services Resource Manager 2576
Install the Microsoft Office Filter Packs on the file server 2576
Configure email notifications on FILE1 2576
Create groups on FILE1 2577
Create files and folders on FILE1 2577
Install Active Directory Rights Management Services 2578
Build the mail server (SRV1) 2582
Build the client virtual machine (CLIENT1) 2582
Lab setup for deploying claims across forests scenario 2583
Build a virtual machine for DC2 2583
Set up a new forest called adatum.com 2584
Set contoso.com as a trusting forest to adatum.com 2585
Create additional users in the Adatum forest 2585
Create the Company claim type on adataum.com 2586
Enable the Company resource property on contoso.com 2587
Enable Dynamic Access Control on adatum.com 2587
Create the Company claim type on contoso.com 2587
Create the central access rule 2588
Create the central access policy 2588
Publish the new policy through Group Policy 2589
Create the Earnings folder on the file server 2590

Set classification and apply the central access policy on the Earnings folder 2590

Best Practices for Securing Active Directory 2590

Foreword 2592

Acknowledgements 2593

Executive Summary 2593
Avenues to Compromise 2594
Reducing the Active Directory Attack Surface 2595
Monitoring Active Directory for Signs of Compromise 2596
Planning for Compromise 2596
Summary of Best Practices for Securing Active Directory Domain Services 2597

Introduction 2600
Account and Group Naming Conventions 2600
About This Document 2601
Microsoft IT and ISRM 2601
Active Directory Security Assessments 2602
Content Origin and Organization 2602
Executive Summary 2602
Introduction 2602
Avenues to Compromise 2602
Reducing the Active Directory Attack Surface 2603
Monitoring Active Directory for Signs of Compromise 2603
Planning for Compromise 2603
Summary of Best Practice Recommendations 2604
Appendices 2604

Avenues to Compromise 2605
Initial Breach Targets 2606
Gaps in Antivirus and Antimalware Deployments 2607
Incomplete Patching 2608
Patch and Vulnerability Management Software 2608
Outdated Applications and Operating Systems 2609
Misconfiguration 2609
In Active Directory 2610
On Domain Controllers 2610
Protecting Domain Controllers 2611
Within the Operating System 2611
Disabling Security Features 2611
Granting Excessive Privilege 2612
Standardizing Local Administrator Credentials 2612
Permitting Installation of Unauthorized Applications 2613

Applications 2614
Lack of Secure Application Development Practices 2614

Attractive Accounts for Credential Theft 2616
Activities that Increase the Likelihood of Compromise 2617
Logging on to Unsecured Computers with Privileged Accounts 2617
Not Maintaining Separate Administrative Credentials 2617
Logons to Compromised Workstations or Member Servers with Privileged Accounts 2618
Unsecured Administrative Workstations 2618
Browsing the Internet with a Highly Privileged Account 2619
Configuring Local Privileged Accounts with the Same Credentials across Systems 2619
Overpopulation and Overuse of Privileged Domain Groups 2619
Poorly Secured Domain Controllers 2619
Privilege Elevation and Propagation 2620
Permanent Privileged Accounts 2620
VIP Accounts 2620
“Privilege-Attached” Active Directory Accounts 2620

Reducing the Active Directory Attack Surface 2621
Privileged Accounts and Groups in Active Directory 2621
Built-in Privileged Accounts and Groups 2621
Highest Privilege Groups in Active Directory 2622
Enterprise Admins 2622
Domain Admins 2622
Administrators 2622
Schema Admins 2622
Protected Accounts and Groups in Active Directory 2623
AdminSDHolder and SDProp 2624

Implementing Least-Privilege Administrative Models 2625
The Privilege Problem 2626
In Active Directory 2627
On Member Servers 2627
On Workstations 2627
In Applications 2628
In Data Repositories 2628
Reducing Privilege 2628
Securing Local Administrator Accounts on Workstations and Member Servers 2628
Securing Local Administrator Accounts 2628
Controls for Local Administrator Accounts 2629
Configuring GPOs to Restrict Administrator Accounts on Domain-Joined Systems 2629
Securing Local Privileged Accounts and Groups in Active Directory 2630
Securing Built-in Administrator Accounts in Active Directory 2630
Controls for Built-in Administrator Accounts 2630

Enable the “Account is sensitive and cannot be delegated” flag on the account 2631
Enable the “Smart card is required for interactive logon” flag on the account 2631
Disable the Account 2631
Configuring GPOs to Restrict Domains’ Administrator Accounts on Domain-Joined Systems 2631
Configuring GPOs to Restrict Administrator Accounts on Domain Controllers 2632
Configure Auditing of Built-in Administrator Accounts 2633
Securing Administrators, Domain Admins and Enterprise Admins Groups 2633
Securing Enterprise Admin Groups 2633
Securing Domain Admins Groups 2634
Securing Administrators Groups in Active Directory 2635
Role-Based Access Controls (RBAC) for Active Directory 2637
Native Approaches to RBAC for Active Directory 2637
Privileged Identity Management 2639
Creating Unprivileged Accounts to Manage Privileged Accounts 2639
Implementing Robust Authentication Controls 2639
General Authentication Controls 2640
Additional Controls for VIP Accounts 2640
Configuring Privileged Account Authentication 2641
UPN Hijacking for Certificate Spoofing 2641

Implementing Secure Administrative Hosts 2642
Principles for Creating Secure Administrative Hosts 2643
Account Configuration 2643
Physical Security 2644
Operating System Versions and Configuration 2645
Microsoft Security Configuration Wizard 2645
Microsoft Security Compliance Manager 2645
AppLocker 2645
RDP Restrictions 2646
Patch and Configuration Management 2646
Blocking Internet Access 2646
Virtualization 2646
Sample Approaches to Implementing Secure Administrative Hosts 2647
Implementing Separate Physical Workstations 2647
Pros 2648
Cons 2648
Implementing a Secure Physical Workstation with a Virtualized Productivity Workstation 2648
Pros 2648
Cons 2648
Implementing a Single Secure Workstation with Connections to Separate “Productivity” and “Administrative” Virtual Machines 2648
Pros 2649

Cons ... 2649
Implementing Secure Administrative Workstations and Jump Servers ... 2649
Pros ... 2650
Cons ... 2650

Securing Domain Controllers Against Attack ... 2650
Physical Security for Domain Controllers ... 2651
Datacenter Domain Controllers ... 2651
Physical Domain Controllers ... 2651
Virtual Domain Controllers ... 2651
Branch Locations ... 2652
Physical Domain Controllers ... 2652
Virtual Domain Controllers ... 2652
Remote Locations with Limited Space and Security ... 2652
Domain Controller Operating Systems ... 2652
Secure Configuration of Domain Controllers ... 2653
Security Configuration Wizard ... 2653
Microsoft Security Compliance Manager ... 2653
AppLocker ... 2653
RDP Restrictions ... 2654
Patch and Configuration Management for Domain Controllers ... 2654
Blocking Internet Access for Domain Controllers ... 2654
Perimeter Firewall Restrictions ... 2655
DC Firewall Configurations ... 2655
Preventing Web Browsing from Domain Controllers ... 2655

Monitoring Active Directory for Signs of Compromise ... 2655
Windows Audit Policy ... 2656
Windows Audit Categories ... 2656
Audit Policy Category Descriptions ... 2657
Audit Account Logon Events ... 2657
Audit Account Management ... 2657
Audit Directory Service Access ... 2657
Audit Logon Events ... 2657
Audit Object Access ... 2657
Auditing Policy Change ... 2658
Audit Privilege Use ... 2658
Audit Process Tracking ... 2658
System Events Audit ... 2658
Advanced Audit Policies ... 2658
Auditing Subcategories Descriptions ... 2659
Account Logon ... 2659
Credential Validation ... 2659
Kerberos Service Ticket Operations ... 2660

Kerberos Authentication Service 2660
Other Account Logon Events 2660
Account Management 2660
User Account Management 2660
Computer Account Management 2661
Security Group Management 2661
Distribution Group Management 2661
Application Group Management 2661
Other Account Management Events 2661
Detailed Process Tracking 2661
Process Creation 2661
Process Termination 2661
DPAPI Activity 2661
RPC Events 2661
Directory Service Access 2662
Directory Service Access 2662
Directory Service Changes 2662
Directory Service Replication 2662
Detailed Directory Service Replication 2662
Logon/Logoff 2662
Logon 2662
Network Policy Server 2662
IPsec Main Mode 2662
IPsec Extended Mode 2663
Other Logon/Logoff Events 2663
Logoff 2663
Account Lockout 2663
IPsec Quick Mode 2663
Special Logon 2663
Policy Change 2663
Audit Policy Change 2663
Authentication Policy Change 2663
Authorization Policy Change 2663
MPSSVC Rule-Level Policy Change 2663
Filtering Platform Policy Change 2664
Other Policy Change Events 2664
Privilege Use 2664
Sensitive Privilege Use 2664
Nonsensitive Privilege Use 2664
Other Privilege Use Events 2664
Object Access 2664
File System 2664
Registry 2665

Kernel Object 2665
SAM 2665
Certification Services 2665
Application Generated 2665
Handle Manipulation 2665
File Share 2665
Filtering Platform Packet Drop 2665
Filtering Platform Connection 2666
Other Object Access Events 2666
System 2666
Security State Change 2666
Security System Extension 2666
System Integrity 2666
Other System Events 2666
Configuring Windows Audit Policy 2666
Setting Windows Audit Policy by Using Group Policy 2667
Setting Windows Audit Policy Using Auditpol.exe 2668
Scripting Auditpol 2669
Other Auditpol Commands 2669
Enforcing Traditional Auditing or Advanced Auditing 2670

Audit Policy Recommendations 2670
Recommended Audit Policies by Operating System 2671
Set Audit Policy on Workstations and Servers 2683
Events to Monitor 2683
Active Directory Objects and Attributes to Monitor 2684
Additional Information for Monitoring Active Directory Domain Services 2684
General List of Security Event ID Recommendation Criticalities 2685

Planning for Compromise 2685
Rethinking the Approach 2686
Identifying Principles for Segregating and Securing Critical Assets 2688
Defining a Limited, Risk-Based Migration Plan 2689
Leveraging “Nonmigratory” Migrations 2689
User Accounts 2689
Servers and Workstations 2690
Applications 2690
Implementing Creative Destruction 2690
Isolating Legacy Systems and Applications 2691
Simplifying Security for End Users 2692

Maintaining a More Secure Environment 2692
Creating Business-Centric Security Practices for Active Directory 2693
Assign a Business Owner to Active Directory Data 2694

Implement Business-Driven Lifecycle Management ... 2694
Classify all Active Directory Data ... 2695
Systems ... 2695
Applications ... 2695
Users ... 2696

Summary of Best Practices ... 2696

Appendices ... 2700

Appendix A: Patch and Vulnerability Management Software ... 2702
Appendix A: Patch and Vulnerability Management Software ... 2702

Appendix B: Privileged Accounts and Groups in Active Directory ... 2702
Appendix B: Privileged Accounts and Groups in Active Directory ... 2702
Rights, Privileges, and Permissions in Active Directory ... 2703
Rights and Privileges ... 2703
Table B-1: User Rights and Privileges ... 2704
Permissions ... 2706
Built-in Privileged Accounts and Groups ... 2706
Enterprise Admins ... 2707
Domain Admins ... 2707
Administrators ... 2708
Schema Admins ... 2708
Additional Built-in and Default Groups in Active Directory ... 2709
Table B-1: Built-in and Default Accounts and Groups in Active Directory ... 2709

Appendix C: Protected Accounts and Groups in Active Directory ... 2729
Appendix C: Protected Accounts and Groups in Active Directory ... 2729
Protected Groups ... 2729
AdminSDHolder ... 2730
SDProp ... 2730
Changing SDProp Interval ... 2730
Running SDProp Manually ... 2731

Appendix D: Securing Built-In Administrator Accounts in Active Directory ... 2738
Appendix D: Securing Built-In Administrator Accounts in Active Directory ... 2738
Controls for Built-in Administrator Accounts ... 2739
Step-by-Step Instructions to Secure Built-in Administrator Accounts in Active Directory 2740
Configuring GPOs to Restrict Administrator Accounts at the Domain-Level ... 2743
Verification Steps ... 2747
Verify “Smart card is required for interactive logon” Account Option ... 2747
Verify “Account is disabled” Account Option ... 2748
Verify “Deny access to this computer from the network” GPO Settings ... 2748
Verify “Deny log on as a batch job” GPO Settings ... 2749

Verify “Deny log on as a service” GPO Settings 2750
Revert Changes to the Printer Spooler Service 2751
Verify “Deny log on through Remote Desktop Services” GPO Settings 2751

Appendix E: Securing Enterprise Admins Groups in Active Directory 2752
Appendix E: Securing Enterprise Admins Groups in Active Directory 2752
Step-by-Step Instructions for Removing All Members from the Enterprise Admins Group 2753
Step-by-Step Instructions to Secure Enterprise Admins in Active Directory 2754
Verification Steps 2759
Verify “Deny access to this computer from the network” GPO Settings 2759
Verify “Deny log on as a batch job” GPO Settings 2760
Create a Batch File 2760
Schedule a Task 2760
Verify “Deny log on as a service” GPO Settings 2761
Revert Changes to the Printer Spooler Service 2762
Verify “Deny log on locally” GPO Settings 2762
Verify “Deny log on through Remote Desktop Services” GPO Settings 2763

Appendix F: Securing Domain Admins Groups in Active Directory..... 2763
Appendix F: Securing Domain Admins Groups in Active Directory 2763
Step-by-Step Instructions for Removing all Members from the Domain Admins Group . 2764
Step-by-Step Instructions to Secure Domain Admins in Active Directory..... 2765
Verification Steps..... 2769
Verify “Deny access to this computer from the network” GPO Settings 2769
Verify “Deny log on as a batch job” GPO Settings..... 2770
Verify “Deny log on as a service” GPO Settings..... 2771
Revert Changes to the Printer Spooler Service..... 2772
Verify “Deny log on locally” GPO Settings 2772
Verify “Deny log on through Remote Desktop Services” GPO Settings 2773

Appendix G: Securing Administrators Groups in Active Directory 2774
Appendix G: Securing Administrators Groups in Active Directory 2774
Step-by-Step Instructions for Removing All Members from the Administrators Group ... 2775
Step-by-Step Instructions to Secure Administrators Groups in Active Directory 2775
Step-by-Step Instructions to Grant User Rights to the Administrators Group 2780
Verification Steps..... 2783
Verify “Deny access to this computer from the network” GPO Settings 2783
Verify “Deny log on as a batch job” GPO Settings..... 2784
Verify “Deny log on as a service” GPO Settings..... 2785
Revert Changes to the Printer Spooler Service..... 2786

Appendix H: Securing Local Administrator Accounts and Groups 2787
Appendix H: Securing Local Administrator Accounts and Groups..... 2787
Controls for Local Administrator Accounts 2787

Step-by-Step Instructions to Secure Local Administrators Groups 2787
Verification Steps 2792
Verify “Deny access to this computer from the network” GPO Settings 2792
Verify “Deny log on as a batch job” GPO Settings 2793

Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory 2796
Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory 2796
Creating Management Accounts for Protected Accounts and Groups in Active Directory 2797
Step-by-Step Instructions for Creating Management Accounts for Protected Groups 2798
Creating a Group to Enable and Disable Management Accounts 2798
Creating the Management Accounts 2802
Enabling Management Accounts to Modify the Membership of Protected Groups 2813
Verifying Group and Account Configuration Settings 2816

Appendix J: Third-Party RBAC Vendors 2823
Appendix J: Third-Party RBAC Vendors 2823
The Dot Net Factory 2823
Key Features and Benefits: 2824
IBM 2825
Domain RBAC 2825
Oracle 2825
Oracle Solaris RBAC Elements and Basic Concepts 2825
Centrify 2826
IT Security & Access Control 2826
The Challenge 2826
The Centrify Solution 2826

Appendix K: Third-Party PIM Vendors 2827
Appendix K: Third-Party PIM Vendors 2827
Cyber-Ark 2827
What is the PIM Suite? 2828
The PIM Suite: features and components 2828
Quest 2829
Privileged Account Management 2829
Controlling and Auditing Superuser Access 2829
Lieberman Software 2829
Privileged Identity Management 2829
Risks of Unsecured Privileged Identities 2829
Controlling Privileged Account Access 2830
Business Value of Privileged Identity Management 2830
Novell 2831
NetIQ Privileged User Manager 2831

Secure access to UNIX, Linux and Windows systems 2831
CA 2831
CA IdentityMinder 2831
Business Challenges 2831
Solution Overview 2832
Key features 2832

Appendix L: Events to Monitor 2833
Appendix L: Events to Monitor 2833

Appendix M: Document Links and Recommended Reading 2858
Appendix M: Document Links and Recommended Reading 2858
Document Links 2858
Recommended Reading 2875
Copyright Information 2876

Active Directory Domain Services Component Updates for Windows Server 2012 R2 2877
What You Will Learn 2877

Identity component updates 2878
Lesson 1: Identity component updates 2878
What You Will Learn 2878

SPN and UPN uniqueness 2879
Overview 2879
Background 2879
Symptoms 2879
New user creation fails if UPN is not unique 2880
DSA.msc 2880
Active Directory Administrative Center (DSAC.exe) 2881
Event 2974 Source: ActiveDirectory_DomainService 2881
SetSPN: 2882
ADSIEDIT: 2883
Windows PowerShell 2883
Restore of an object that would result in a duplicate UPN fails: 2885
Identify the conflicting UPN on the deleted objectUsing repadmin.exe 2886
To identify all objects with the same UPN:Using Repadmin.exe 2886
Using Global Search 2886
Using Windows PowerShell 2887
Duplicate SPN 2889
Workflow 2890

8648: 2890

8647: 2891

8648: ... 2891

8647: ... 2891
Try This: Exploring SPN and UPN uniqueness ... 2892

Winlogon Automatic Restart Sign-On (ARSO) ... 2893
Overview ... 2893
What's changed? ... 2893
Group Policy: Sign-in last interactive user automatically after a system-initiated restart ... 2896
Troubleshooting ... 2897
Reasons why autologon might fail ... 2898
User Must Change Password at Next Login ... 2898
User Account Disabled ... 2898
Logon Hours and Parental Controls ... 2898
Additional Resources ... 2898

TPM Key Attestation ... 2899
Overview ... 2899
Terminology ... 2899
Background ... 2900
TPM key attestation ... 2900
Why is TPM key attestation important? ... 2900
How does TPM key attestation work? ... 2900
Deployment overview ... 2901
Deployment details ... 2902
Configure a certificate template ... 2902
CA configuration ... 2907
Troubleshooting ... 2909
Key attestation fields are unavailable on a certificate template ... 2909
Verification of TPM device for attestation ... 2910
See Also ... 2911

CA Backup and Restore Windows PowerShell cmdlets ... 2911
Overview ... 2911
Backup-CARoleService ... 2911
-Password <Secure String> ... 2913
Restore-CARoleService ... 2913
Issues ... 2915
Additional Resources ... 2916
Try This: Backup the CA in your lab using Windows PowerShell ... 2916

Command line process auditing ... 2916
Overview ... 2917
Configuration ... 2918
You must have Audit Process Creation auditing enabled to see event ID 4688. ... 2918

In order to see the additions to event ID 4688, you must enable the new policy setting:
Include command line in process creation events 2918
To ensure that Advanced Audit Policy Configuration settings are not overwritten 2921
Additional Resources 2922
Try This: Explore command line process auditing 2922

Directory Services component updates 2922
What You Will Learn 2923
Domain and Forest Functional Levels 2923
Overview 2923
New DFL and FFL 2923
The Windows Server 2012 R2 Domain Functional Level enables support for the following:
..... 2923
Minimum DFL enforced on new domain creation 2924
Lowering the forest and domain functional levels 2924
ADPREP 2925
Deprecation of NTFRS 2926
Overview 2926
LDAP Query Optimizer changes 2927
Overview 2927
Background 2927
Details of change 2927
Comparison between old and new algorithm 2928
Sample results using the new algorithm 2929
To enable the Stats control in LDP 2930
Try This: Use LDP to return query statistics 2931
Additional Resources 2932
1644 Event improvements 2932
Overview 2932
Background 2932
Additional search statistics added to event 1644 2933
New time-based threshold registry value for event 1644 logging 2933
Comparison of the old and new event ID 1644 2933
Try This: Use the event log to return query statistics 2935
Active Directory Replication throughput improvement 2936
Overview 2936
Additional Resources 2936

How to Configure Protected Accounts 2936
Protected Users 2937
Requirements for using protected accounts 2938
Troubleshoot events related to Protected Users 2938
New logs for Protected Users 2938
Troubleshoot TGT expiration 2939

Troubleshoot delegation issues ... 2939
Audit authentication attempts ... 2940
Provide DC-side protections for services and computers ... 2940
Authentication policies ... 2941
Quick Kerberos refresher ... 2941
Overview ... 2942
Requirements for using authentication policies ... 2942
Restrict a user account to specific devices and hosts ... 2943
Configure domain controller support ... 2943
Create a user account audit for authentication policy with ADAC ... 2944
Add computer account or group conditions ... 2948
Add computer claim conditions ... 2952
Troubleshoot missing computer claims ... 2953
Provision a user account with an authentication policy with ADAC ... 2954
Configure Dynamic Access Control support on devices and hosts ... 2955
Troubleshoot Authentication Policies ... 2956
Determine the accounts that are directly assigned an Authentication Policy ... 2956
Use the Authentication Policy Failures – Domain Controller administrative log ... 2957
Manage authentication policies by using Windows PowerShell ... 2957
Authentication policy silos ... 2957
Manage authentication policy silos by using Windows PowerShell ... 2960

Active Directory Federation Services Overview ... 2961
Role description ... 2961
Practical applications ... 2962
AD FS in Windows Server 2012 ... 2963
AD FS in Windows Server 2012 R2 ... 2963
Enable users to access resources on their personal devices from anywhere ... 2963
Enhanced access control risk management tools ... 2963
Simplified deployment experience ... 2964
Enhanced sign-in with AD FS experience ... 2964
Enable developers to build modern applications ... 2965
Other improvements ... 2965
See Also ... 2965

Getting Started with AD FS ... 2965
See Also ... 2966

Set up the lab environment for AD FS in Windows Server 2012 R2 ... 2966
Step 1: Configure the domain controller (DC1) ... 2966
Create test Active Directory accounts ... 2967
Create a GMSA account ... 2967
Step 2: Configure the federation server (ADFS1) by using Device Registration Service ... 2967
Install a server SSL certificate ... 2968

Install the AD FS server role ... 2968
Configure the federation server ... 2968
Configure Device Registration Service ... 2969
Add Host (A) and Alias (CNAME) Resource Records to DNS ... 2970
Step 3: Configure the web server (WebServ1) and a sample claims-based application ... 2971
Step 4: Configure the client computer (Client1) ... 2974
See Also ... 2974

Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications ... 2975
Key concepts introduced in the solution ... 2975
Workplace Join ... 2975
Device Registration Service ... 2975
Workplace Join as a seamless second factor authentication ... 2976
Single Sign-On ... 2976
Solution Overview ... 2976
See Also ... 2976

Walkthrough Guide: Workplace Join with an iOS Device ... 2976
Join an iOS device with Workplace Join ... 2977
See Also ... 2977

Walkthrough Guide: Workplace Join with a Windows Device ... 2978
Access the web application before device registration ... 2978
Join your device with Workplace Join ... 2978
Access the web application after joining the workplace ... 2979
See Also ... 2979

Overview: Connect to Applications and Services from Anywhere with Web Application Proxy 2980
Providing Access to Applications and Services ... 2981
Using Active Directory Federation Services ... 2981
AD FS Proxy ... 2982
Roles and Features Included in this Scenario ... 2982
Scenario Steps ... 2982
See Also ... 2983

Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy ... 2983
Step 1: Attempt to access the web application from an Internet client ... 2985
Step 2: Configure the Web Application Proxy server and publish the application ... 2986
Step 3: Configure and test accessing a website using Integrated Windows authentication ... 2989
Install Windows Authentication on WebServ1 ... 2989
Create a new website using IIS ... 2989
Create a non-claims-aware relying party trust ... 2990
Configure Kerberos constrained delegation ... 2991

Test accessing the application internally 2992
Publish the application 2992
Test accessing the application 2993
Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control 2993
See also 2995

Overview: Manage Risk with Multi-Factor Access Control 2995
Key concepts - Multi-factor access control in AD FS 2995
Managing Risk with Multi-factor Access Control 2999
Common Scenarios 2999
Advanced Scenarios 2999
See Also 3001

Walkthrough Guide: Manage Risk with Multi-Factor Access Control 3001
About This Guide 3001
Step 1: Setting up the lab environment 3002
Step 2: Verify the default AD FS access control mechanism 3002
Step 3: Configure multi-factor access control policy based on user data 3003
Step 4: Verify multi-factor access control mechanism 3004
See Also 3005

Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications 3005
In this guide 3005
Key Concepts - Authentication mechanisms in AD FS 3005
Benefits of authentication mechanisms in AD FS 3005
Authentication scope 3006
Primary and additional authentication methods 3007
Configuring MFA 3008
Scenario Overview 3009
See Also 3010

Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications 3010
About This Guide 3010
Step 1: Setting up the lab environment 3011
Step 2: Verify the default AD FS authentication mechanism 3011
Step 3: Configure MFA on your federation server 3012
Select an additional authentication method 3012
Certificate authentication 3012
Windows Azure Multi-Factor Authentication 3013
Set up MFA policy 3015
Step 4: Verify MFA mechanism 3017
See Also 3017

AD FS Design ... 3017
See Also ... 3017

AD FS Design Guide in Windows Server 2012 R2 ... 3018
About this guide ... 3018
In this guide ... 3018
Acknowledgments ... 3018
See Also ... 3018

Identify Your AD FS Deployment Goals ... 3019
Enable your users to access resources on their personal devices from anywhere ... 3020
Enhance your access control risk management tools ... 3020
Use AD FS to enhance the sign-in experience ... 3020
See Also ... 3021

Plan Your AD FS Deployment Topology ... 3021
Determining which type of AD FS configuration database to use ... 3021
SQL Server considerations ... 3023
How the configuration database type you select may impact hardware resources ... 3023
Where to place a federation server ... 3024
Supported deployment topologies ... 3024
See Also ... 3024

Federation Server Farm Using WID ... 3024
Deployment considerations ... 3025
Who should use this topology? ... 3025
What are the benefits of using this topology? ... 3025
What are the limitations of using this topology? ... 3026
Server placement and network layout recommendations ... 3026
See Also ... 3027

Federation Server Farm Using WID and Proxies ... 3027
Deployment considerations ... 3027
Who should use this topology? ... 3027
What are the benefits of using this topology? ... 3028
What are the limitations of using this topology? ... 3028
Server placement and network layout recommendations ... 3028
See Also ... 3029

Federation Server Farm Using SQL Server ... 3029
Deployment considerations ... 3029
Who should use this topology? ... 3030
What are the benefits of using this topology? ... 3030
What are the limitations of using this topology? ... 3030
Supported SQL Server Versions ... 3030

Server placement and network layout recommendations 3030
See Also 3031

AD FS Requirements 3031
Certificate requirements 3032
Hardware requirements 3037
Software requirements 3037
AD DS requirements 3037
Configuration database requirements 3038
Browser requirements 3038
Network requirements 3039
Attribute store requirements 3040
Application requirements 3041
Authentication requirements 3041
Workplace join requirements 3042
Cryptography requirements 3042
See Also 3044

AD FS Design Guide in Windows Server 2012 3044
About this guide 3045
In this guide 3045
Acknowledgments 3045

Identifying Your AD FS Deployment Goals 3046

Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services 3046

Provide Your Active Directory Users Access to the Applications and Services of Other Organizations 3048

Provide Users in Another Organization Access to Your Claims-Aware Applications and Services 3050

Mapping Your Deployment Goals to an AD FS Design 3051

Web SSO Design 3052

Federated Web SSO Design 3053

Determine Your AD FS Deployment Topology 3055

AD FS Deployment Topology Considerations 3056
Determining which type of AD FS configuration database to use 3056
SQL Server considerations 3059
How the configuration database type you select may impact hardware resources 3059
Verifying that your production environment can support an AD FS deployment 3059

Stand-Alone Federation Server Using WID ... 3060
Test lab considerations ... 3060
Who should use this topology? ... 3060
What are the benefits of using this topology? ... 3060
What are the limitations of using this topology? ... 3060

Federation Server Farm Using WID ... 3061
Deployment considerations ... 3061
Who should use this topology? ... 3061
What are the benefits of using this topology? ... 3062
What are the limitations of using this topology? ... 3062
Server placement and network layout recommendations ... 3062

Federation Server Farm Using WID and Proxies ... 3063
Deployment considerations ... 3063
Who should use this topology? ... 3063
What are the benefits of using this topology? ... 3063
What are the limitations of using this topology? ... 3063
Server placement and network layout recommendations ... 3064

Federation Server Farm Using SQL Server ... 3064
Deployment considerations ... 3065
Who should use this topology? ... 3065
What are the benefits of using this topology? ... 3065
What are the limitations of using this topology? ... 3065
Supported SQL Server Versions ... 3065
Server placement and network layout recommendations ... 3066

Planning Your Deployment ... 3066
In this section ... 3067

Using AD DS Claims with AD FS ... 3067
About Dynamic Access Control ... 3067
What's New in AD FS? ... 3068
Benefits of Using AD DS Claims with AD FS ... 3068
Differences Between AD DS and AD FS Issued Claims ... 3069
How AD DS Issued Claims Work with AD FS ... 3069

Best Practices for Secure Planning and Deployment of AD FS ... 3070
Core security best practices for AD FS ... 3070
SQL Server–specific security best practices for AD FS ... 3073

Planning for Interoperability with AD FS 1.x ... 3075
Understanding the Name ID claim type ... 3076

When to Use Identity Delegation ... 3077

What is identity delegation? ... 3077
How does identity delegation work? ... 3078
Configuring AD FS for identity delegation ... 3079
Configuring the front-end Web application for identity delegation ... 3079

Deploying AD FS in the Account Partner Organization ... 3079
In this section ... 3080

Review the Role of the Federation Server in the Account Partner ... 3080

Review the Role of the Federation Server Proxy in the Account Partner ... 3081

Prepare Client Computers in the Account Partner ... 3081

Deploying AD FS in the Resource Partner Organization ... 3082
In this section ... 3082

Review the Role of the Federation Server in the Resource Partner ... 3082

Review the Role of the Federation Server Proxy in the Resource Partner ... 3083

Determine Your Federated Application Strategy in the Resource Partner ... 3084

Planning Federation Server Placement ... 3084

When to Create a Federation Server ... 3085
Determine the organizational role for the federation server ... 3085
Determine which AD FS design to deploy ... 3086
Differences between a federation server and a federation server proxy ... 3086
How to create a federation server ... 3087

Where to Place a Federation Server ... 3087
Configuring your firewall servers for a federation server ... 3088

When to Create a Federation Server Farm ... 3088
Best practices for deploying a federation server farm ... 3089
Configuring federation servers for a farm ... 3089

Certificate Requirements for Federation Servers ... 3091
Determining your CA strategy ... 3093
Certificate revocation lists ... 3094

Token-Signing Certificates ... 3094
Token-signing certificate requirements ... 3095
How token-signing certificates are used across partners ... 3095
Deployment considerations for token-signing certificates ... 3096

Service Communications Certificates ... 3097

Service communication certificate requirements 3097
Deployment considerations for service communication certificates 3098

Name Resolution Requirements for Federation Servers 3098
Configure corporate DNS 3099

Planning Federation Server Proxy Placement 3101

When to Create a Federation Server Proxy 3101
How to create a federation server proxy 3102

Where to Place a Federation Server Proxy 3102
Configuring your firewall servers for a federation server proxy 3103
Configuring ISA Server to allow SSL 3103

When to Create a Federation Server Proxy Farm 3104
Configuring federation server proxies for a farm 3104

Certificate Requirements for Federation Server Proxies 3105

Name Resolution Requirements for Federation Server Proxies 3106
DNS zone serving only the perimeter network 3106
1. Configure the hosts file on the federation server proxy 3108
2. Configure perimeter DNS 3109
DNS zone serving both the perimeter network and Internet clients 3109
1. Configure perimeter DNS 3110
2. Configure Internet DNS 3110

Planning for AD FS Server Capacity 3111
AD FS capacity planning terms 3113
Configuration environment used during AD FS testing 3114
Measure AD FS server capacity 3114
Continue reading more about AD FS capacity planning 3115

Planning for Federation Server Capacity 3115
AD FS configuration database size and growth 3116
Memory, CPU and disk space requirements 3116
Estimate the number of federation servers for your organization 3116
Using the AD FS Capacity Planning Sizing Spreadsheet 3117
How to use this spreadsheet 3117

Planning for Federation Server Proxy Capacity 3118
Estimate the number of federation server proxies required for your organization 3119

Appendix A: Reviewing AD FS Requirements 3119
Hardware requirements 3120
Software requirements 3120

Certificate requirements ... 3120
Federation server certificates ... 3120
Federation server proxy certificates ... 3122
Browser requirements ... 3123
Cookies ... 3124
Network requirements ... 3124
TCP/IP network connectivity ... 3124
DNS ... 3124
Attribute store requirements ... 3125
AD DS ... 3125
Schema requirements ... 3125
Functional-level requirements ... 3126
Service account requirements ... 3126
LDAP ... 3126
SQL Server ... 3126
Custom attribute stores ... 3126
Application requirements ... 3127
Authentication requirements ... 3127
Smart card logon ... 3127
Smart card authentication ... 3127

AD FS Deployment ... 3128
See Also ... 3128

Windows Server 2012 R2 AD FS Deployment Guide ... 3128
See Also ... 3128

Deploying a Federation Server Farm ... 3128
See Also ... 3129

Join a Computer to a Domain ... 3130
See Also ... 3131

Enroll an SSL Certificate for AD FS ... 3131
See Also ... 3132

Install the AD FS Role Service ... 3132
See Also ... 3133

Configure a Federation Server ... 3133
Configure the first federation server in a new federation server farm ... 3133
To configure the first federation server in a new federation server farm by using the Active Directory Federation Service Configuration Wizard ... 3133
To configure the first federation server in a new federation server farm via Windows PowerShell ... 3135

Add a federation server to an existing federation server farm 3137
To add a federation server to an existing federation server farm via the Active Directory Federation Service Configuration Wizard 3137
To add a federation server to an existing federation server farm via Windows PowerShell 3138
See Also 3140

Configure a federation server with Device Registration Service 3140
Prepare your Active Directory forest to support devices 3141
Enable Device Registration Service on a federation server farm node 3141
Enable seamless second factor authentication 3141
Update the Web Application Proxy configuration 3142
See Also 3142

Configure Corporate DNS for the Federation Service and DRS 3142
Step 6: Add a Host (A) and Alias (CNAME) Resource Record to Corporate DNS for the Federation Service and DRS 3142
See Also 3143

Verify That a Federation Server Is Operational 3143
See Also 3144

Deploying Federation Server Proxies 3144
See Also 3145

Windows Server 2012 AD FS Deployment Guide 3145
About this guide 3145
What this guide does not provide 3146
In this guide 3146
Acknowledgments 3146

Planning to Deploy AD FS 3147
Reviewing your AD FS design 3147

Implementing Your AD FS Design Plan 3147
How to implement your AD FS design using this guide 3148

Checklist: Implementing a Web SSO Design 3148

Checklist: Implementing a Federated Web SSO Design 3150

Configuring Partner Organizations 3152
About account partner organizations 3152
About resource partner organizations 3152

Checklist: Configuring the Account Partner Organization 3153

Checklist: Configuring the Resource Partner Organization 3157

Add an Attribute Store 3160
Additional references 3161

Create a Claims Provider Trust Using Federation Metadata 3161
Additional references 3162

Create a Claims Provider Trust Manually 3163
Additional references 3164

Create a Relying Party Trust Using Federation Metadata 3164
Additional references 3165

Create a Relying Party Trust Manually 3165
Additional references 3166

Add a Claim Description 3167
Additional references 3167

Configure Client Computers to Trust the Account Federation Server 3168
Configuring Internet Explorer settings manually 3168
Configuring Internet Explorer settings by using Group Policy 3168

Distribute Certificates to Client Computers by Using Group Policy 3169

Configuring Claim Rules 3170
About Claim Rules 3170

Checklist: Creating Claim Rules for a Claims Provider Trust 3170

Checklist: Creating Claim Rules for a Relying Party Trust 3172

Create a Rule to Pass Through or Filter an Incoming Claim 3174
Additional references 3175

Create a Rule to Permit All Users 3176
Additional references 3176

Create a Rule to Permit or Deny Users Based on an Incoming Claim 3176
Additional references 3177

Create a Rule to Send LDAP Attributes as Claims 3178
Additional references 3179

Create a Rule to Send Group Membership as a Claim 3179
Additional references 3180

Create a Rule to Transform an Incoming Claim 3180

Additional references 3181

Create a Rule to Send an Authentication Method Claim 3182

Create a Rule to Send an AD FS 1.x Compatible Claim 3185
Creating a rule to issue an AD FS 1 3186
Creating a rule to issue an AD FS 1 3186

Create a Rule to Send Claims Using a Custom Rule 3187
Additional references 3188

Deploying Federation Servers 3189
About federation servers 3189

Checklist: Setting Up a Federation Server 3190

Join a Computer to a Domain 3195
Additional references 3195

Add a Host (A) Resource Record to Corporate DNS for a Federation Server 3195
Additional references 3196

Export the Private Key Portion of a Server Authentication Certificate 3196
Additional references 3197

Import a Server Authentication Certificate to the Default Web Site 3197
Additional references 3198

Manually Configure a Service Account for a Federation Server Farm 3198

Install the Federation Service Role Service 3200
Prerequisites 3200

Create the First Federation Server in a Federation Server Farm 3201
Additional references 3203

Create a Stand-Alone Federation Server 3203
Additional references 3205

Add a Federation Server to a Federation Server Farm 3205
Additional references 3206

Add a Token-Signing Certificate 3206
Additional references 3207

Add a Token-Decrypting Certificate 3207
Additional references 3208

Set a Service Communications Certificate 3208

Additional references 3209

Verify That a Federation Server Is Operational 3209
Additional references 3210

Configure Performance Monitoring 3210
AD FS performance counters 3211

Deploying Federation Server Proxies 3212
About federation server proxies 3213

Checklist: Setting Up a Federation Server Proxy 3213

Join a Computer to a Domain 3217
Additional references 3218

Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Only the Perimeter Network 3218
Add the IP address of a federation server to the hosts file 3218
Add a host (A) resource record to perimeter DNS for a federation server proxy 3219
Additional references 3220

Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Both the Perimeter Network and Internet Clients 3220
Add a host (A) resource record to the Internet DNS zone for a federation server proxy 3220
Add a host (A) resource record to the perimeter DNS zone for a federation server proxy 3221
Additional references 3222

Export the Private Key Portion of a Server Authentication Certificate 3222
Additional references 3223

Import a Server Authentication Certificate to the Default Web Site 3223
Additional references 3224

Install the Federation Service Proxy Role Service 3224
Additional references 3225

Configure a Computer for the Federation Server Proxy Role 3225
Configuring an Alternate TCP/IP Port for Proxy Operations 3226
Additional references 3228

Verify That a Federation Server Proxy Is Operational 3228
Additional references 3229

Configure Performance Monitoring 3229
AD FS performance counters 3230

Interoperating with AD FS 1.x 3232

Differences between Federation Service settings ... 3232
See Also ... 3233

Checklist: Configuring AD FS to Consume Claims from AD FS 1.x ... 3233
Checklist: Configuring AD FS to consume claims from AD FS 1.x ... 3233

Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Federation Service ... 3235
Checklist: Configuring AD FS to send claims to an AD FS 1.x Federation Service ... 3235

Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Claims-Aware Web Agent ... 3238
Checklist: Configuring AD FS to send claims to an AD FS 1.x claims-aware Web agent ... 3238

Create a Relying Party Trust Manually ... 3240
Additional references ... 3241

Create a Claims Provider Trust Manually ... 3241
Additional references ... 3242

Create a Rule to Send an AD FS 1.x Compatible Claim ... 3243
Creating a rule to issue an AD FS 1 ... 3243
Creating a rule to issue an AD FS 1 ... 3244

AD FS Operations ... 3245
See Also ... 3245

Customizing the AD FS Sign-in Pages ... 3246
Overview ... 3246
Customizing the AD FS sign-in experience ... 3247
General sign-in page ... 3247
Change company name ... 3248
Change company logo ... 3248
Change Illustration ... 3249
Add sign-in page description ... 3249
Add Help Desk Link ... 3249
Add Home Link ... 3250
Add Privacy Link ... 3250
Custom Web Themes ... 3250
Custom Error Messages ... 3251
Customize the error page description ... 3252
Customize a generic error message ... 3252
Customize an authorization error message ... 3252
Customize a device authentication error message ... 3252
Customize a support email error message ... 3252
Customize a relying party authorization message ... 3253
Home Realm Discovery ... 3253
Configure Identity Provider to use certain email suffixes ... 3254

Configure an identity provider list per relying party ... 3254
Bypass Home Realm Discovery for the intranet ... 3254
Update password ... 3255
Customize the Update Password page description ... 3256
Multi-factor authentication and external authentication providers ... 3256
Localization ... 3258
Removing the Microsoft copyright ... 3259
See Also ... 3259

Add an Attribute Store ... 3259
See Also ... 3260

Create a Relying Party Trust ... 3260
See Also ... 3262

Create a Non-Claims-Aware Relying Party Trust ... 3262
See Also ... 3263

Create a Claims Provider Trust ... 3263
Additional references ... 3264
See Also ... 3265

Configuring Claim Rules ... 3265
See Also ... 3265

Create a Rule to Pass Through or Filter an Incoming Claim ... 3265
See Also ... 3266

Create a Rule to Permit All Users ... 3266
See Also ... 3267

Create a Rule to Permit or Deny Users Based on an Incoming Claim ... 3267
See Also ... 3268

Create a Rule to Send LDAP Attributes as Claims ... 3268
See Also ... 3269

Create a Rule to Send Group Membership as a Claim ... 3270
See Also ... 3271

Create a Rule to Transform an Incoming Claim ... 3271
See Also ... 3272

Create a Rule to Send an Authentication Method Claim ... 3272
See Also ... 3275

Create a Rule to Send an AD FS 1.x Compatible Claim ... 3276
Create a rule to issue an AD FS 1. ... 3276

Create a rule to issue an AD FS 1. ... 3277
See Also ... 3278

Create a Rule to Send Claims Using a Custom Rule ... 3278
See Also ... 3279

Configuring Authentication Policies ... 3279
Configure authentication policies via the AD FS Management snap-in ... 3280
Configure authentication policies via Windows PowerShell ... 3282
See Also ... 3288

Configuring AD FS Extranet Lockout ... 3288
See Also ... 3289

Configuring Client Access Policies ... 3289
Client Access Policy Claim Types ... 3290
Client Access Policy Scenarios ... 3290
Enabling Client Access Policy ... 3290
Scenario 1: Block all external access to Office 365 ... 3291
Scenario 2: Block all external access to Office 365 except Exchange ActiveSync ... 3292
Scenario 3: Block all external access to Office 365 except browser-based applications ... 3294
Scenario 4: Block all external access to Office 365 except for designated Active Directory groups ... 3296
Building the IP address range expression ... 3297
Regular Expressions ... 3298
Testing the Expression ... 3299
Claim Types ... 3299
X-MS-Forwarded-Client-IP ... 3299
X-MS-Client-Application ... 3300
X-MS-Client-User-Agent ... 3300
X-MS-Proxy ... 3301
InsideCorporateNetwork ... 3301
X-MS-Endpoint-Absolute-Path (Active vs Passive) ... 3301
See Also ... 3302

AD FS Technical Reference ... 3302
See Also ... 3302

Understanding Key AD FS Concepts ... 3302
AD FS terminology used in this guide ... 3302
Overview of AD FS ... 3305

The Role of Attribute Stores ... 3306
How attribute stores fit in with your AD FS deployment goals ... 3307
Attribute stores that are supported by AD FS ... 3307

The Role of the AD FS Configuration Database 3307
Using WID to store the AD FS configuration database 3308
How a WID federation server farm works 3309
Primary federation server 3309
Secondary federation servers 3309
How the AD FS configuration database is synchronized 3309
Using SQL Server to store the AD FS configuration database 3310
SAML artifact resolution 3311
SAML/WS-Federation token replay detection 3311

The Role of Claims 3312
What are claims? 3312
How claims are sourced 3312
How claims flow 3313
How claims are issued 3313
What are claim types? 3313
What are claim descriptions? 3315
When generating Federation Metadata 3316
When claims rules are processed 3316

The Role of Claim Rules 3316
What are claim rules? 3316
How claim rules are processed 3317
What are claim rule templates? 3317
How claim rule templates work 3317
How to create a claim rule 3318
Using claim rule templates 3318
Using the claim rule language 3318
Using Windows PowerShell 3319
What is a claim rule set? 3319
What are claim rule set types? 3320

The Role of the Claims Engine 3322
Claim rules execution process 3323
Step 1 – Initialization 3323
Adding a claim to the input claim set for a rule set 3323
Step 2 – Execution 3324
Adding a claim to the output claim set for a rule set 3325
Adding a claim to both claim sets for a rule set 3325
Step 3 – Execution Result 3326
Sending the execution output to the claims pipeline 3326
Processing authorization rules 3327

The Role of the Claims Pipeline 3328

Claims pipeline process 3328

The Role of the Claim Rule Language 3329
Creating custom claim rules using the claim rule language 3330
Using claim rule templates to learn about the claim rule language syntax 3330
Understanding the components of the claim rule language 3330
Conditions 3330
Single-condition examples 3331
Multiple-condition examples 3332
Regular-condition examples 3332
Issuance statements 3332
Claim issuance actions 3333
Expressions 3335
Exists functions 3336
Rule body 3336
Additional references 3336

How URIs Are Used in AD FS 3337
URIs as partner network addresses 3337
URIs as object identifiers 3337
URI prefix matching for relying party identifiers 3339

Determine the Type of Claim Rule Template to Use 3341
Claim rule template types 3342

When to Use an Authorization Claim Rule 3344
About claim rules 3345
Permit All Users 3345
Permit access to users with this incoming claim 3345
Deny access to users with this incoming claim 3346
Authorizing users 3346
Authorization rule sets 3346
Supported claim types 3346
How to create this rule 3347
Using the claim rule language 3347
Example of how to create an authorization rule based on multiple claims 3347
Example of how to create authorization rules that will delegate who can create or remove federation server proxy trusts 3347

When to Use a Pass Through or Filter Claim Rule 3350
About claim rules 3350
Pass through all claim values 3351
Filtering a claim 3351
Configuring this rule on a claims provider trust 3352
Configuring this rule on a relying party trust 3352

How to create this rule 3352
Using the claim rule language 3353
Examples of how to construct a pass through or filter rule syntax 3353
Best practices for creating custom rules 3353

When to Use a Transform Claim Rule 3355
About claim rules 3355
Pass through all claim values 3356
Transforming a claim 3356
Configuring this rule on a claims provider trust 3357
Configuring this rule on a relying party trust 3357
How to create this rule 3357
Using the claim rule language 3357
Examples of how to construct a transform rule syntax 3358
Best practices for creating custom rules 3358

When to Use a Send LDAP Attributes as Claims Rule 3359
About claim rules 3360
Mapping of LDAP attributes to outgoing claim types 3360
How to create this rule 3360
Using the claim rule language 3361
Example: How to query an AD LDS attribute store and return a specified value 3361
Example: How to query an Active Directory attribute store and return a specified value 3362
Example: How to query an Active Directory attribute store based on the value of an incoming claim 3362
Example: How to use two custom rules to extract the manager e-mail from an attribute in Active Directory 3363
Additional references 3364

When to Use a Send Group Membership as a Claim Rule 3364
About claim rules 3364
Outgoing claim value 3365
Configuring this rule on a claims provider trust 3365
How to create this rule 3365
Using the claim rule language 3366
Example: How to issue group claims based on the user’s group membership 3366
Additional references 3366

When to Use a Custom Claim Rule 3366
About claim rules 3367
How to create this rule 3367
Using the claim rule language 3368
Example: How to combine first and last names based on a user’s name attribute values 3368
Example: How to issue a manager claim based on whether users have direct reports 3368

Example: How to issue a PPID claim based on an LDAP attribute 3368

Active Directory Lightweight Directory Services Overview 3369

Active Directory Rights Management Services Overview 3370
Role description 3370
Practical applications 3370
New and changed functionality 3371
Server Manager information 3371
See also 3371

What's New in Active Directory Rights Management Services (AD RMS)? 3372
Role/Feature description 3373
Changes in AD RMS and SQL Server requirements 3373
Changes in deployment of AD RMS for Server Manager and Windows PowerShell 3374
Using Server Manager to deploy AD RMS 3374
Using Windows PowerShell to deploy AD RMS 3375
Using Windows PowerShell to uninstall AD RMS 3376
Changes to Windows PowerShell for deploying AD RMS 3376
AD RMS fails to install if multiple installations are active simultaneously in Server Manager 3380
Deployment considerations for virtualized AD RMS servers 3380
Server Core Support for AD RMS 3381
Recent feature updates for AD RMS 3382
Simple Delegation for AD RMS 3382
Strong cryptography for AD RMS 3383
See also 3384

Test Lab Guide: Deploying an AD RMS Cluster 3385
In this guide 3385
What this guide does not provide 3386
Test lab overview 3386
Hardware and software requirements 3387
Step 1: Complete the Base TLG Configuration and add users and groups 3388
Step 2: Prepare SQL1 as a SQL Server database server for supporting AD RMS 3392
Step 3: Install and configure AD RMS on APP1 3399
Step 4: Verify AD RMS configuration after you complete the configuration 3402

How AD RMS Works 3405
How AD RMS clients work 3405
How AD RMS servers work 3406
Working with Templates in AD RMS 3407
A simplified view of AD RMS in action 3407
How AD RMS works within an organization 3408
How AD RMS works between organizations 3409

AD RMS Infrastructure Deployment Tips 3410
Centralize servers as much as possible 3410
Use a single cluster if possible 3411
Use licensing-only cluster only if required 3411
Centralize licensing from multiple forests in one central cluster 3412

Using AD RMS with Hardware Security Modules 3412
In this guide 3413
How AD RMS Benefits From HSMs 3413
How AD RMS Does Encryption and Uses Encryption Keys 3413
How HSMs Help Improve Server Performance 3413
How HSMs Help to Ensure the Security of Private Keys 3414
Recommendations and best practices for using HSMs with AD RMS 3415
Implementation guidance 3415
Deployment 3415
Configuration 3416
Trust 3416
Decommissioning 3417
Operations guidance 3417
Management 3417
Maintenance 3417
Disaster recovery 3418

AD RMS Best Practices Guide 3418
AD RMS Installation Best Practices 3418
AD RMS Clustering Best Practices 3421
AD RMS Database Server Best Practices 3423
AD RMS Client Deployment Best Practices 3424
AD RMS Rights Policy Templates Best Practices 3425
AD RMS Performance and Logging Best Practices 3425

Application Server Overview 3426

Failover Clustering Overview 3426
Feature description 3426
Practical applications 3427
New and changed functionality 3427
Hardware requirements 3427
Software requirements 3428
Server Manager information 3428
See also 3428

What's New in Failover Clustering in Windows Server 2012 R2 3429
Feature description 3429
New and changed functionality 3430

Highly available virtual machine improvements ... 3432
Shared virtual hard disk (for guest clusters) ... 3432
Virtual machine drain on shutdown ... 3432
Virtual machine network health detection ... 3433
Cluster Shared Volume (CSV) improvements ... 3434
Optimized CSV placement policies ... 3434
Increased CSV resiliency ... 3435
CSV cache allocation ... 3435
CSV diagnosibility ... 3436
CSV interoperability ... 3436
Deploy an Active Directory-detached cluster ... 3436
Quorum improvements ... 3437
Dynamic witness ... 3437
Quorum user interface improvements ... 3438
Force quorum resiliency ... 3439
Tie breaker for 50% node split ... 3440
Configure the Global Update Manager mode ... 3441
Cluster node health detection ... 3443
Turn off IPsec encryption for inter-node cluster communication ... 3444
Cluster dashboard ... 3444
See also ... 3445

What's New in Failover Clustering in Windows Server 2012 ... 3445
Feature description ... 3446
New and changed functionality ... 3446
Cluster scalability ... 3447
Management of large-scale clusters by using Server Manager and Failover Cluster Manager ... 3448
Management and mobility of clustered virtual machines and other clustered roles ... 3448
Cluster Shared Volumes ... 3451
Support for Scale-Out File Servers ... 3452
Cluster-Aware Updating ... 3452
Virtual machine application monitoring and management ... 3452
Cluster validation tests ... 3453
Active Directory Domain Services integration ... 3453
Quorum configuration and dynamic quorum ... 3454
Cluster upgrade and migration ... 3454
Task Scheduler integration ... 3454
Windows PowerShell support ... 3455
Removed or deprecated functionality ... 3455
See also ... 3456

Failover Clustering Hardware Requirements and Storage Options ... 3456
Hardware requirements for Hyper-V ... 3458

Deploying storage area networks with failover clusters ... 3458
See also ... 3459

Validate Hardware for a Failover Cluster ... 3459
Step 1: Prepare to validate hardware for a failover cluster ... 3460
What is cluster validation? ... 3460
Supported cluster configurations ... 3460
Common validation scenarios ... 3461
Categories of validation tests ... 3462
Step 2: Validate a new or existing failover cluster ... 3463
Step 3: Analyze validation results ... 3465
What to do if a validation test fails ... 3466
Provide a validation report when you request support from Microsoft ... 3466
Updates to validation tests ... 3466
Advanced validation scenarios ... 3467
Including storage tests ... 3470
Considerations when you include storage tests ... 3470
Considerations when you do not include storage tests ... 3471
Frequently asked questions ... 3472
If a cluster passes all tests in the Validate a Configuration Wizard, is it supported? ... 3472
Will failover cluster solutions be listed in the Windows Server Catalog? ... 3472
How does Microsoft customer support check if the solution has been validated? ... 3472
What if I make a change to the cluster configuration, like add a node? Do I have to run the Validate a Configuration Wizard again? ... 3472
See also ... 3472

Create a Failover Cluster ... 3473
Verify the prerequisites ... 3473
Install the Failover Clustering feature ... 3474
Validate the configuration ... 3475
Create the failover cluster ... 3476
Create clustered roles ... 3478
Create a failover cluster by using Windows PowerShell ... 3479
See also ... 3480

Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster ... 3480
Overview of the quorum in a failover cluster ... 3481
Why configure the quorum? ... 3481
Overview of the quorum configuration options ... 3482
Witness configuration ... 3483
Node vote assignment ... 3485
Dynamic quorum management ... 3486
General recommendations for quorum configuration ... 3486
Configure the cluster quorum ... 3487

Configure the cluster quorum settings ... 3487
Recover a cluster by starting without quorum ... 3490
Force start cluster nodes ... 3490
Prevent quorum on remaining cluster nodes ... 3491
Quorum considerations for disaster recovery configurations ... 3492
Automatic failover ... 3492
Manual failover ... 3493
See Also ... 3493

Use Cluster Shared Volumes in a Failover Cluster ... 3494
Review requirements and considerations for using CSV in a failover cluster ... 3495
Network configuration considerations ... 3495
About I/O synchronization and I/O redirection in CSV communication ... 3496
Storage and disk configuration requirements ... 3497
Node requirements ... 3498
Plan to use CSV in a failover cluster ... 3498
Arrangement of LUNs, volumes, and VHD files ... 3499
Number and size of LUNs and volumes ... 3499
Add a disk to CSV on a failover cluster ... 3500
Enable the CSV cache for read-intensive workloads (optional) ... 3501
Back up CSV ... 3503
See also ... 3504

Prestage Cluster Computer Objects in Active Directory Domain Services ... 3504
Step 1: Prestage the CNO in AD DS ... 3505
Step 2: Grant the user permissions to create the cluster ... 3506
Step 3: Grant the CNO permissions to the OU or prestage VCOs for clustered roles ... 3508
See also ... 3510

Deploy an Active Directory-Detached Cluster ... 3510
Deployment considerations ... 3510
How to deploy an Active Directory-detached cluster ... 3512

Deploy a Hyper-V Cluster ... 3512
Prerequisites ... 3513
Hardware requirements ... 3513
Software requirements ... 3514
Network infrastructure and domain account requirements ... 3514
Limitations ... 3515
Step 1: Connect both physical computers to the networks and storage ... 3515
Step 2: Install Hyper-V and Failover Clustering on both physical computers ... 3516
Step 3: Create a virtual switch ... 3516
Step 4: Validate the cluster configuration ... 3517
Step 5: Create the cluster ... 3517

Step 6: Add a disk as CSV to store virtual machine data ... 3517
Step 7: Create a highly available virtual machine ... 3518
Considerations for creating a virtual machine that will be highly available ... 3518
Step 8: Install the guest operating system on the virtual machine ... 3519
Step 9: Test a planned failover ... 3520
Step 10: Test an unplanned failover ... 3520
Step 11: Modify the settings of a virtual machine ... 3521
Step 12: Remove a virtual machine from a cluster ... 3523
See also ... 3524

Network Recommendations for a Hyper-V Cluster in Windows Server 2012 ... 3524
Overview of different network traffic types ... 3525
Management traffic ... 3525
Cluster traffic ... 3525
Live migration traffic ... 3526
Storage traffic ... 3526
Replica traffic ... 3527
Virtual machine access traffic ... 3527
How to isolate the network traffic on a Hyper-V cluster ... 3528
Isolate traffic on the management network ... 3528
Isolate traffic on the cluster network ... 3528
Isolate traffic on the live migration network ... 3530
Isolate traffic on the storage network ... 3531
Isolate traffic for replication ... 3532
NIC Teaming (LBFO) recommendations ... 3532
Quality of Service (QoS) recommendations ... 3533
Virtual machine queue (VMQ) recommendations ... 3534
Example of converged networking: routing traffic through one Hyper-V virtual switch ... 3535
Appendix: Encryption ... 3539
Cluster traffic ... 3539
Live migration traffic ... 3539
SMB traffic ... 3539
Replica traffic ... 3539
See also ... 3539

Using Guest Clustering for High Availability ... 3540
Physical host considerations ... 3540
Deploy a guest cluster on a physical host cluster ... 3540
Place the virtual machines on different physical hosts ... 3541
Guest cluster considerations for deployment and management ... 3542
Support requirements ... 3542
Node limit ... 3542
Active Directory domain requirements ... 3543
Storage options ... 3543

Network considerations ... 3544
See also ... 3546

Deploy a Guest Cluster Using a Shared Virtual Hard Disk ... 3547
Deployment scenarios ... 3547
Step 1: Configure the physical servers ... 3548
Step 2: Create the virtual machines ... 3549
Step 3: Create and enable a shared virtual hard disk ... 3550
Step 4: Create the guest cluster ... 3552

Cluster-Aware Updating Overview ... 3553
Feature description ... 3553
Practical applications ... 3553
Important functionality ... 3554
Hardware and software requirements ... 3555
Additional requirements and best practices ... 3556
Starting CAU ... 3556
See also ... 3557

Cluster-Aware Updating: Frequently Asked Questions ... 3557
What is CAU? ... 3558
Does CAU work with Windows Server 2008 R2 or Windows 7? ... 3558
Is CAU functionality limited to only specific clustered applications? ... 3558
Does CAU support updates from Microsoft Update and Windows Update? ... 3558
Does CAU support WSUS updates? ... 3558
Can CAU apply limited distribution release updates? ... 3558
Can I use CAU to apply cumulative updates? ... 3559
Can I schedule CAU to apply updates? ... 3559
Can CAU updates be scheduled when a backup is in progress? ... 3559
Can CAU work well with System Center Configuration Manager? ... 3559
Do I need administrative credentials to run CAU? ... 3560
Can I script CAU functionality to automate it further? ... 3560
When CAU coordinates updates, what happens to the clustered roles that are active on each cluster node? ... 3560
How does CAU select the target nodes that each clustered role should fail over to? ... 3560
Does CAU load balance the clustered roles at the end of the updating process? ... 3560
How does CAU select the order of nodes to update? ... 3561
When a CAU Updating Run is initiated on the cluster, what happens if all the cluster nodes are not online? ... 3561
Can I use CAU to select and update only a single node? ... 3561
Can CAU report cluster node updates that are initiated from outside CAU? ... 3561
Is CAU flexible enough to support my unique IT process needs? ... 3561
How can I export the CAU preview and update results? ... 3562
How do I install CAU? ... 3562

Does CAU need any service or component running on the cluster nodes that are being updated? 3563
What is the difference between using CAU and using System Center 2012 Virtual Machine Manager to update Hyper-V clusters? 3563
Can I use remote-updating on a cluster that is configured for self-updating? 3563
Can I reuse my cluster update settings across clusters? 3563
Where can I get more information about the public CAU plug-in specification? 3563
See also 3564

Requirements and Best Practices for Cluster-Aware Updating 3564
Install the Failover Clustering feature and the Failover Clustering Tools 3564
Obtain an administrator account 3565
Verify the cluster configuration 3565
Configure the nodes for remote management 3566
Enable Windows Management Instrumentation 3567
Enable Windows PowerShell and Windows PowerShell remoting 3567
Install .NET Framework 4.5 3568
Enable a firewall rule to allow automatic restarts 3568
Best practices recommendations for using CAU 3568
Recommendations for applying Microsoft updates 3568
Apply Microsoft updates in branch office scenarios 3569
Recommendations for using the Microsoft.HotfixPlugin 3570
Additional recommendations 3570
Test cluster updating readiness 3571
Tests for cluster updating readiness 3571
See also 3577

How CAU Plug-ins Work 3577
Install a plug-in 3577
Specify a plug-in and plug-in arguments 3578
Specify a CAU plug-in 3578
Specify CAU plug-in arguments 3579
Specify optional plug-in arguments 3579
Manage plug-ins using Windows PowerShell cmdlets 3579
Use the Microsoft.WindowsUpdatePlugin 3580
Configure the Windows Update Agent query string 3582
Use the Microsoft.HotfixPlugin 3583
Configure a hotfix root folder structure 3585
Customize the hotfix configuration file 3587
Restrict access to the hotfix root folder 3590
See also 3592

Advanced Options and Updating Run Profiles for CAU 3592
Options that can be set in an Updating Run Profile 3593

Options that you specify when you request an Updating Run 3597
Use Updating Run Profiles 3598
See also 3599

File and Storage Services Overview 3599
Role description 3599
Practical applications 3599
New and changed functionality 3600
Data Deduplication 3602
iSCSI Target Server 3603
Storage Spaces and storage pools 3604
Unified remote management of File and Storage Services in Server Manager 3605
Windows PowerShell cmdlets for File and Storage Services 3606
Removed or deprecated functionality 3607
Requirements for running File and Storage Services 3607
How do I deploy and configure this role in a multiserver environment? 3608
Can I run this role on virtual machines? 3608
Can I run this role in a clustered environment? 3608
Are there considerations for managing this role remotely? 3608
Are there considerations for managing the role on the Server Core installation option? 3608
Server Manager information 3609
See also 3611

Data Deduplication Overview 3612
Feature description 3612
Practical applications 3613
New and changed functionality 3614
Requirements 3615
Interoperability with Windows Azure virtual machines 3615
Architecture overview 3616
See also 3616

What's New in Data Deduplication in Windows Server 2012 R2 3617
Feature description 3617
New and changed functionality 3617
Data deduplication for remote storage of VDIworkloads 3618
Expand an optimized file on its original path 3619
See also 3619

Plan to Deploy Data Deduplication 3620
Step 1: Target deployments 3620
Step 2: Determine which volumes are candidates for deduplication 3620
Server and volume requirements for deduplication 3622
Unsupported configurations 3623

Deduplication considerations 3624
Step 3: Evaluate savings with the Deduplication Evaluation Tool 3625
Determine potential space reclamation with the Measure-DedupFileMetadata cmdlet..... 3626
Step 4: Plan the rollout, scalability, and deduplication policies..... 3627

Backup and Restore Considerations for Deduplicated Volumes 3628
Windows Server Backup 3628
Optimized backup considerations 3629
Incremental optimized backup considerations 3629
Full volume restore from optimized backup 3629
Selective restore from an optimized backup 3630
Selective partial volume restore to a newly formatted volume 3630
Selective restore from a backup store that allows data deduplication..... 3630
Non-optimized backup and restore 3631
Incremental non-optimized backup 3631

Data Deduplication Interoperability..... 3631
BranchCache..... 3632
Failover clusters 3632
DFS Replication 3632
File Server Resource Manager quotas 3632
Single Instance Storage 3633
Migrating SIS volumes..... 3633

Install and Configure Data Deduplication 3634
Prerequisites 3634
Step 1: Set up the server 3634
Step 2: Enable data deduplication 3635
Step 3: Set data deduplication jobs..... 3636
Optimization jobs 3636
Data Scrubbing jobs 3639
Garbage Collection jobs 3640
Step 4: Set data deduplication schedules 3640

Monitor and Report for Data Deduplication 3644
How do I tell if deduplication is keeping up with the rate of incoming data?..... 3648

DFS Namespaces and DFS Replication Overview 3649
Role service descriptions 3649
Practical applications 3650
New and changed functionality 3650
Removed or deprecated functionality 3652
Requirements for running DFS 3652
Server Manager information..... 3653
Installing DFS Namespaces, DFS Replication, and DFS Management..... 3653

Interoperability with Windows Azure virtual machines ... 3654
See also ... 3655

What's New in DFS Replication in Windows Server 2012 R2 ... 3656
Role service description ... 3656
New and changed functionality ... 3656
Windows PowerShell module for DFS Replication ... 3657
DFS Replication Windows Management Infrastructure provider ... 3658
Database cloning for initial sync ... 3658
Database corruption recovery ... 3659
Cross-file RDC disable ... 3659
File staging tuning ... 3660
Preserved file restoration ... 3660
Unexpected shutdown database recovery improvements ... 3661
Membership disabling improvements ... 3661
Removed or deprecated functionality ... 3662
See also ... 3662

What's New in DFS Namespaces and DFS Replication in Windows Server 2012 ... 3663
Role service descriptions ... 3663
New and changed functionality ... 3663
Windows PowerShell module for DFS Namespaces ... 3663
DFS Namespaces: Site awareness for DirectAccess clients ... 3664
DFS Namespaces: Windows Management Infrastructure provider ... 3665
DFS Replication: Support for Data Deduplication volumes ... 3665
Removed or deprecated functionality ... 3665
See also ... 3665

Export a Clone of the DFS Replication Database ... 3666
Prerequisites ... 3667
Step 1: Create an example replication environment ... 3668
Step 2: Export a clone of a DFS Replication database ... 3669
Step 3: Import a clone of a DFS Replication database ... 3671
Use cloning to replace a corrupted DFS Replication database ... 3673
See also ... 3674

DFS Replication: Copying Files to Preseed or Stage Initial Synchronization ... 3675
Overview of preseeding files for DFS Replication ... 3675
Comparison of tools ... 3676
In this guide ... 3678
Additional resources ... 3678

Step 1: Preseed Files for DFS Replication ... 3679

Use Robocopy to Preseed Files for DFS Replication ... 3679

Prerequisites ... 3680
Step 1: Download and install the latest version of Robocopy ... 3680
Step 2: Stabilize files that will be replicated ... 3681
Step 3: Copy the replicated files to the destination server ... 3682
Next step ... 3684

Use Windows Server Backup to Preseed Files for DFS Replication ... 3685
Prerequisites ... 3685
Step 1: Install the Windows Server Backup feature ... 3686
Step 2: Back up the replicated files from the source server ... 3686
Step 3: Restore the replicated files to the destination server ... 3688
Next step ... 3690

Use Windows NTBackup to Preseed Files for DFS Replication ... 3690
Prerequisites ... 3691
Step 1: Install NTBackup on the destination server ... 3691
Step 2: Back up the files from the source server ... 3692
Step 3: Restore the backup to the destination server ... 3692
Next step ... 3693

Step 2: Validate Preseeded Files for DFS Replication ... 3693
Best practices for validating preseeded files ... 3694
Two methods for validating preseeded files ... 3694
Get-DfsrFileHash method ... 3694
Compare a file hash for a single file ... 3695
Compare file hashes for a folder's immediate contents ... 3695
Compare file hashes between two servers by using Compare-Object ... 3696
Dfsrdiag method ... 3697
Compare a file hash for a single file ... 3698
Related resources ... 3698

Troubleshoot Preseeding for Initial Synchronization in DFS Replication ... 3698
Troubleshoot mismatched hashes ... 3699
Troubleshoot missing files and folders ... 3699
Troubleshoot poor copy or restore performance ... 3699
Related resources ... 3700

File Server Resource Manager Overview ... 3700
Feature descriptions ... 3700
Practical applications ... 3701
See also ... 3701

What's New in File Server Resource Manager in Windows Server 2012 R2 ... 3702
Feature description ... 3702
New and changed functionality ... 3702

Clear classification property values that no longer apply to an updated file 3703
Set maximum values for storage reports 3703
See also 3704

What's New in File Server Resource Manager in Windows Server 2012 3704
Feature description 3704
New and changed functionality 3705
Dynamic Access Control 3706
Automatic classification 3706
Manual classification 3707
File management tasks 3707
Access-denied assistance 3707
Removed or deprecated functionality 3708
See also 3708

Folder Redirection, Offline Files, and Roaming User Profiles overview 3708
Did you mean… 3708
Technology description 3709
Practical applications 3709
New and changed functionality 3710
Always Offline mode 3710
Cost-aware synchronization 3711
Primary computers for Folder Redirection and Roaming User Profiles 3711
Hardware requirements 3712
Software requirements 3713
See also 3713

Deploy Folder Redirection, Offline Files, and Roaming User Profiles 3714
See Also 3714

Deploy Folder Redirection with Offline Files 3714
Prerequisites 3715
Hardware requirements 3715
Software requirements 3715
Step 1: Create a folder redirection security group 3715
Step 2: Create a file share for redirected folders 3716
Step 3: Create a GPO for Folder Redirection 3718
Step 4: Configure folder redirection with Offline Files 3718
Step 5: Enable the Folder Redirection GPO 3719
Step 6: Test Folder Redirection 3719
Appendix A: Checklist for deploying Folder Redirection 3720
See Also 3721

Deploy Roaming User Profiles 3721
Prerequisites 3722

Hardware requirements 3722
Software requirements 3722
Considerations when using Roaming User Profiles on multiple versions of Windows 3723
Step 1: Enable the use of separate profile versions 3723
Step 2: Create a Roaming User Profiles security group 3724
Step 3: Create a file share for roaming user profiles 3725
Step 4: Optionally create a GPO for Roaming User Profiles 3727
Step 5: Optionally set up Roaming User Profiles on user accounts 3727
Step 6: Optionally set up Roaming User Profiles on computers 3729
Step 7: Enable the Roaming User Profiles GPO 3730
Step 8: Test Roaming User Profiles 3730
Appendix A: Checklist for deploying Roaming User Profiles 3731
Appendix B: Profile version reference information 3732
Change History 3732
See Also 3733

Deploy Primary Computers for Folder Redirection and Roaming User Profiles 3733
Prerequisites 3733
Software requirements 3733
Step 1: Designate primary computers for users 3734
Step 2: Optionally enable primary computers for Folder Redirection in Group Policy 3735
Step 3: Optionally enable primary computers for Roaming User Profiles in Group Policy 3735
Step 4: Enable the GPO 3736
Step 5: Test primary computer function 3736
See Also 3737

Enable Advanced Offline Files Functionality 3737
See Also 3737

Enable Optimized Moving of Redirected Folders 3737
Prerequisites 3738
Step 1: Enable optimized move in Group Policy 3738
Step 2: Relocate the file share for redirected folders 3738
See Also 3739

Enable the Always Offline Mode to Provide Faster Access to Files 3740
Prerequisites 3740
Enabling the Always Offline mode 3740
See Also 3741

Enable Background File Synchronization on Metered Networks 3741
Prerequisites 3742
Enable background file synchronization of Offline Files on metered networks 3742
See Also 3743

Disable Offline Files on Individual Redirected Folders ... 3743
Prerequisites ... 3743
Disabling Offline Files on individual redirected folders ... 3743
See Also ... 3745

Troubleshoot User Profiles with Events ... 3745
Step 1: Checking events in the Application log ... 3746
Step 2: Viewing the Operational log for the User Profile Service ... 3746
Step 3: Enabling and viewing analytic and debug logs ... 3747
Step 4: Creating and decoding a trace ... 3747
See Also ... 3748

iSCSI Target Block Storage Overview ... 3748
Block storage requirements ... 3749
See Also ... 3749

What's New in iSCSI Target Server in Windows Server 2012 R2 ... 3749
Feature description ... 3749
New and changed functionality ... 3750
See also ... 3751

iSCSI High-Availability Block Storage Technical Preview ... 3752
Requirements ... 3752
Technical overview ... 3752

iSCSI Target Block Storage, How To ... 3753
Prerequisites ... 3753
Step 1: Install iSCSI Target Server and failover clustering ... 3753
Step 2: Configure the iSCSI initiator ... 3754
Step 3: Create the high availability iSCSI Target Server role service ... 3754
Step 4: Configure iSCSI storage provider identity credentials ... 3755
See also ... 3756

iSCSI Target Implementation Notes ... 3756
Known Issues ... 3757
SCSI Response Excess Immediate Data ... 3757
SCSI Response to Incorrect Amount of Data Condition ... 3757
SCSI Response to Unexpected Unsolicited Data ... 3757
Handling Text Response Text Field in Discovery Session ... 3757
Handling Text Request During Normal Session ... 3758
Handling Text Request During Discovery Session ... 3758
Handling Text Request with SendTargets= ... 3759
Handling Text Request with SendTargets=All during Normal Session ... 3759
Text Response Negotiation Failure ... 3759
Handling Text Request with C bit ... 3759

Standard Login Key Negotiation ... 3760
Duplicate or Out of Order CmdSN ... 3760
Task Management Response for Non-Existent Task ... 3760
Incorrect Logout Reason Code Being Accepted During Discovery ... 3761
Logout Response Non-existent Connection ... 3761
Handling Not Understood for Required Keys ... 3761
Handling Re-negotiation During Login ... 3761
Reject SNACK ... 3762
Text mode negotiation ... 3762

iSCSI Target Boot Overview ... 3762
Feature description ... 3763
Practical applications ... 3763
Hardware requirements ... 3763
Software requirements ... 3764

Datacenter Diskless Boot Technical Preview ... 3764
Requirements ... 3765
Technical overview ... 3765

Network File System Overview ... 3765
Feature description ... 3766
Practical applications ... 3766
New and changed functionality ... 3766
NFS version 4.1 ... 3767
NFS infrastructure ... 3767
NFS version 3 continuous availability ... 3768
Deployment and manageability improvements ... 3768
Server Manager information ... 3768
See also ... 3769

Deploy Network File System ... 3769
What’s new in Network File System ... 3769
Scenarios for using Network File System ... 3770
System requirements ... 3771
Deploy NFS infrastructure ... 3771
Install Network File System ... 3771
Configure NFS authentication ... 3772
Create an NFS file share ... 3772

NTFS ... 3773
Feature description ... 3773
Practical applications ... 3773
See also ... 3774

What's New in NTFS for Windows Server 2012 R2 3774
Feature description 3774
New and changed functionality 3775
See also 3775

NTFS Health and Chkdsk 3776
Requirements 3776
Technical overview 3776

Resilient File System Overview 3777
Feature description 3778
Deployment scenarios 3778
Important functionality 3778
New and updated functionality 3779
Corruptions on parity spaces 3780
Subfolder recovery from ReFS metadata corruption 3780
ReFS is available on client operating systems 3780
ReFS registry entry 3780
Storage cmdlets in Windows 3780
See also 3780

Server for NFS Data Store 3781
Requirements 3781
Technical overview 3782

Server Message Block overview 3783
Feature description 3783
Practical applications 3783
New and changed functionality 3783
Hardware requirements 3786
See also 3787

Deploy Hyper-V over SMB 3787
Prerequisites 3788
Step 1: Configuring file server clusters 3789
Step 2: Install Hyper-V 3791
Step 3: Create an SMB file share 3791
Step 4: Create a virtual machine and virtual hard disk file on the file share 3793
Step 5: Migrate virtual machine storage to an SMB file share 3793
Step 6: Initiate a live migration of a virtual machine to another cluster node 3794
Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage 3795
Troubleshooting 3796
See also 3797

Improve Performance of a File Server with SMB Direct ... 3798
Requirements ... 3799
Enabling and disabling SMB Direct ... 3799
Test performance of SMB Direct ... 3800
See also ... 3801

Deploy SMB Direct with InfiniBand Network Adapters ... 3801
Overview ... 3801
Configure a subnet manager ... 3802
Use a managed switch with a built-in subnet manager. ... 3802
Use OpenSM with an unmanaged switch (recommended only for test environments) ... 3802
Configure the IP address ... 3803
Verify the configuration ... 3803
Review the performance counters ... 3803
Review the event logs for RDMA network adapters ... 3803
See also ... 3803

Deploy SMB Direct with RoCE Network Adapters ... 3804
Overview ... 3804
Configuring Priority Flow Control ... 3805
Configure PFC on cluster nodes ... 3805
Configure PFC on a switch ... 3806
Configure the IP address ... 3807
Verify the configuration ... 3807
Review the performance counters ... 3807
Review the event logs for RDMA network adapters ... 3807
See also ... 3807

Deploy SMB Direct with Ethernet (iWARP) Network Adapters ... 3808
Overview ... 3808
Configure the IP addresses ... 3808
Configure Windows Firewall ... 3809
Allow access across multiple subnets ... 3809
Verify the configuration ... 3810
Verify the network adapter configuration ... 3810
Verify the SMB configuration ... 3810
Verify the SMB connection ... 3811
Review the performance counters ... 3811
Review the event logs for RDMA network adapters ... 3811
See also ... 3812

What's New in SMB in Windows Server 2012 R2 ... 3812
Feature description ... 3812
New and changed functionality ... 3812

See also ... 3815

SMB Security Enhancements ... 3815
SMB Encryption ... 3815
Secure dialect negotiation ... 3817
New signing algorithm ... 3818
Disabling SMB 1.0 ... 3818
See also ... 3819

Storage Management Overview ... 3819
Technology description ... 3819
Requirements ... 3819
Technical overview ... 3820
See also ... 3821

Storage Spaces Overview ... 3822
Technology description ... 3822
Important functionality ... 3822
New and changed functionality ... 3824
Requirements ... 3824
See also ... 3825

Deploy Storage Spaces on a Stand-Alone Server ... 3826
Prerequisites ... 3828
Step 1: Create a storage pool ... 3830
Step 2: Create a virtual disk ... 3832
Step 3: Create a volume ... 3834
See also ... 3835

Deploy Clustered Storage Spaces ... 3835
Prerequisites ... 3837
Step 1: Enable MPIO on each server ... 3839
Step 2: Verify that all servers can see the shared disks ... 3840
Step 3: (optional) Create storage spaces through File and Storage Services ... 3840
Step 4: Create a failover cluster ... 3841
Step 5: (optional) Create clustered storage spaces in Failover Cluster Manager ... 3843
Step 6: (optional) Add a cluster disk to a CSV ... 3847
See also ... 3848

How Storage Spaces Makes Use of Hot Spares ... 3848
How to designate hot spares ... 3848
When creating a storage pool ... 3848
When adding a physical disk ... 3848
How Storage Spaces monitors disk health ... 3849
Write errors ... 3849

Read errors ... 3849
Disconnected disks ... 3850

What's New in Storage Spaces in Windows Server 2012 R2 ... 3850
Technology description ... 3850
New and changed functionality ... 3851
Storage tiers ... 3851
Write-back cache ... 3852
Parity space support for failover clusters ... 3853
Dual parity ... 3853
Automatically rebuild storage spaces from storage pool free space ... 3854
See also ... 3854

Supporting Information Workers with Reliable File Services and Storage ... 3854
Availability and Performance ... 3855
Always Offline mode ... 3855
Cluster-Aware Updating for minimal downtime ... 3855
SMB and NFS v3 Transparent Failover ... 3855
Chkdsk improvements ... 3856
Storage Spaces and Storage Pools ... 3856
SMB Directory Leasing ... 3856
Cost Effectiveness ... 3856
Thin provisioning of storage ... 3857
Data Deduplication ... 3857
BranchCache for Network Files version 2 ... 3857
Enhanced support for low cost SATA disks ... 3857
Access-denied assistance ... 3858
Security and Manageability ... 3858
Primary Computers support for Folder Redirection and Roaming User Profiles ... 3858
SMB Encryption ... 3858
File Classification Infrastructure and Dynamic Access Control ... 3859
Unified remote management of File and Storage Services in Server Manager ... 3859
Improved management of Folder Redirection, Offline Files, and Roaming User Profiles ... 3859
See Also ... 3859

Thin Provisioning and Trim Storage Overview ... 3860
Technology description ... 3860
Requirements ... 3860
Technical overview ... 3860
See also ... 3861

Plan and Deploy Thin Provisioning ... 3861
Determine if thin provisioning is appropriate for your environment ... 3862
Perform capacity forecasting ... 3862

Determine threshold notification and resource exhaustion settings 3863
Monitor and respond to events 3864
Consider space reclamation and potential performance impact 3865
See also 3867

Windows Offloaded Data Transfers Overview 3867
Feature description 3867
Practical applications 3867
Important functionality 3868
Hardware requirements 3870
Software requirements 3870
Hyper-V Requirements 3871
See also 3871

Deploy Windows Offloaded Data Transfers 3872
Prerequisites 3872
Hardware requirements 3872
Software requirements 3873
Hyper-V Requirements 3873
Step 1: Gather storage array information 3874
Step 2: Validate file system filter drivers 3874
Step 3: Establish a performance baseline 3874
Disable ODX 3875
Create a System Performance Report during a data transfer 3875
Step 4: Test ODX performance 3875
Enable ODX 3876
Verify ODX performance 3876
Appendix: Deployment checklist 3876
See Also 3877

Work Folders Overview 3877
Role description 3877
Practical applications 3877
Important functionality 3878
Software requirements 3879
Work Folders compared to other sync technologies 3880
Server Manager information 3881
Interoperability with Windows Azure virtual machines 3881
See also 3881

Designing a Work Folders Implementation 3882
Software requirements 3883
Deployment scenarios 3884
Single-Site Deployment 3884

Multiple-Site Deployment ... 3884
Hosted Deployment ... 3884
Deployment technologies ... 3884
Active Directory Domain Services ... 3885
File Servers ... 3885
Group Policy ... 3885
Windows Intune ... 3885
Web Application Proxy/Device Registration Service ... 3886
Additional design considerations ... 3886
Number of Sync Servers ... 3886
Number of Sync Shares ... 3887
Access to Sync Shares ... 3887
Design checklist ... 3888
Next steps ... 3889
See also ... 3889

Deploying Work Folders ... 3889
Step 1: Obtain SSL certificates ... 3890
Step 2: Create DNS records ... 3890
Step 3: Install Work Folders on file servers ... 3890
Step 4: Binding the SSL certificate on the sync servers ... 3891
Step 5: Create security groups for Work Folders ... 3891
Step 6: Optionally delegate user attribute control to Work Folders administrators ... 3892
Step 7: Create sync shares for user data ... 3893
Step 8: Optionally specify a tech support email address and Active Directory Federation Services authentication ... 3895
Step 9: Optionally set up server automatic discovery ... 3896
Step 10: Configure Web Application Proxy or another reverse proxy ... 3897
Step 11: Optionally use Group Policy to configure domain-joined PCs ... 3897
See also ... 3897

Group Policy Overview ... 3898
Feature description ... 3898
Practical applications ... 3899
New and changed functionality ... 3899
See also ... 3899

What's New in Group Policy in Windows Server 2012 R2 ... 3900
Feature description ... 3900
New and changed functionality ... 3900
IPv6 support ... 3901
Policy caching ... 3901
Event logging ... 3901
See also ... 3902

What's New in Group Policy in Windows Server 2012 ... 3902
Feature description ... 3902
New and changed functionality ... 3902
Remote Group Policy Update ... 3903
Group Policy Results report improvements ... 3903
Group Policy infrastructure status ... 3904
Local Group Policy support for Windows RT ... 3905
Sign-in optimizations ... 3905
Fast Startup ... 3906
New Group Policy starter GPOs ... 3906
Group Policy cmdlet changes ... 3906
Registry.pol changes ... 3907
Group Policy Client service idle state ... 3907
Group Policy settings in Internet Explorer 10 ... 3907
Group Policy Preferences for Internet Explorer 10 ... 3908
Removed or deprecated functionality ... 3908
See also ... 3908

Configure Firewall Port Requirements for Group Policy ... 3909
Remote Resultant Set of Policy (RSoP) Group Policy results: ports that require firewall rules ... 3910
Configure firewall rules by creating a GPO from the Group Policy Reporting Firewall Ports Starter GPO and linking to the domain ... 3910
Remote Group Policy refresh: ports that require firewall rules ... 3911
Configure firewall rules by creating a GPO from the Group Policy Remote Update Firewall Ports Starter GPO and linking to the domain ... 3912

Group Policy Management Console ... 3913
Install Group Policy Management Console (GPMC) ... 3913
Open GPMC ... 3913

Resultant Set of Policy ... 3914
RSoP overview ... 3914

Local Group Policy Editor ... 3915
Use the Local Group Policy Editor ... 3915

Group Policy Preferences ... 3916
Prerequisite Fundamentals ... 3917
Group Policy ... 3917
Client-side Extensions ... 3918
Group Policy Processing ... 3918
Scope ... 3919
Group Policy Preferences ... 3928
Client-side Extensions ... 3929

Common Configurations ... 3930
Item-level Targeting ... 3931
Processing ... 3936

Drive Map ... 3938
Overview ... 3938
Capabilities ... 3938
Configurable Options ... 3939
Action (1) ... 3939
Location (2) ... 3940
Reconnect (3) ... 3940
Label as (4) ... 3940
Drive Letter (5) ... 3940
Connect as (6) ... 3941
Hide/Show this drive (7) ... 3941
Hide/Show all drives (8) ... 3941
How does it work? ... 3942
Client-side and Tool identifiers ... 3942
XML File Configuration ... 3943
Elements ... 3944
Outer Element ... 3944
Attributes ... 3944
Inner Element ... 3945
Attributes ... 3945
Properties Element ... 3948
Attributes ... 3948
Xml to User Interface mapping ... 3955
CSE Processing ... 3957
Startup ... 3958
Prerequisite information ... 3958
Out-of-Scope Group Policy objects ... 3959
Group Policy Preference History ... 3959
Configuration File ... 3961
New and Changed Group Policy objects ... 3961
Configuration File ... 3961
Outer Element Processing ... 3962
Inner Element Processing ... 3962
Action Processing ... 3962
Create ... 3962
Delete ... 3968
Replace ... 3969
Update ... 3975
Dependencies ... 3981

Interaction with other components ... 3982
Logging ... 3982

Printers ... 3984
Overview ... 3984
Capabilities ... 3984
Pushed Printers ... 3985
Processing Complete ... 3985
Local Printer ... 3985
Action (1) ... 3986
Name (2) ... 3987
Port (3) ... 3987
Printer Path (4) ... 3987
Set this printer as the default printer (5) ... 3987
Location (6) ... 3987
Comment (7) ... 3987
Shared Printer ... 3988
Action (1) ... 3988
Action (1) ... 3989
Share Path (3) ... 3989
Set this printer as the default printer (4) ... 3989
Only if local printer is not present (5) ... 3990
Local Port (6) ... 3990
Reconnect (7) ... 3990
Unmap all local ports (8) ... 3990
TCP/IP Printer (Port printer) ... 3990
General Tab ... 3991
Port Settings Tab ... 3993
How does it work? ... 3994
Client-side and Tool Identifiers ... 3994
XML Configuration ... 3995
XML Declaration ... 3996
Elements ... 3996
Outer Elements ... 3996
Attributes ... 3996
Inner Elements ... 3997
Local Printer Element ... 3997
Port Printer Element ... 4004
Shared Printer Element ... 4015
Evaluation of XML ... 4023
CSE Processing ... 4025
Startup ... 4025
Prerequisite information ... 4025

Out-of-Scope Group Policy objects 4027
Group Policy Preference History 4027
Configuration File 4028
New and Changed Group Policy objects 4028
Configuration File 4029
Outer Element Processing 4029
Inner Element Processing 4029
Action Processing 4030
Local Printer 4030
Create 4030
Delete 4032
Replace 4033
Update 4035
Port Printer 4038
Shared Printer 4066
Dependencies 4085
Interaction with other components 4086
Logging 4086

Group Policy Analysis and Troubleshooting Overview 4088
Group Policy troubleshooting description 4088
Practical considerations 4088
See also 4089

Check Group Policy Infrastructure Status 4090
Prerequisites 4090
Step 1: Check Group Policy infrastructure health 4092
Step 2: Check the results of the Group Policy infrastructure status report 4092
Step 3: Check Active Directory replication issues 4093
Step 4: Check file services replication issues 4094
See also 4094

Force a Remote Group Policy Refresh (GPUpdate) 4094
Prerequisites 4095
Step 1: Configure firewall rules on each client that will be managed with remote Group Policy refresh 4096
Step 2: Schedule a remote Group Policy refresh 4097
See also 4099

Understand the Effect of Fast Logon Optimization and Fast Startup on Group Policy 4099
Prerequisites 4099
Group Policy settings and CSEs 4099
Asynchronous and synchronous processing 4099
Foreground and background processing 4100

CSE processing requirements 4100
Fast Logon Optimization and Group Policy processing 4106
How Fast Logon Optimization affects CSE processing 4107
Fast Logon Optimization and required synchronous policy application 4107
Fast Logon Optimization, required synchronous processing, and required foreground processing for a CSE 4109
Fast startup and Group Policy processing 4109
Related resources 4110

Hyper-V Overview 4110
Role and technology description 4110
Practical applications 4111
New and changed functionality 4112
Hardware requirements 4112
Software requirements (for supported guest operating systems) 4112
See also 4116

Feature Overviews 4116

Hyper-V Automation Overview 4117
Key benefits 4118
Requirements 4118
Technical overview 4118
Designed for IT pros 4118

Hyper-V Dynamic Memory Overview 4120
Key benefits 4120
Technical overview 4120
Minimum memory configuration with reliable restart operation 4120
Run-time Dynamic Memory configuration changes 4122
About the Dynamic Memory Settings 4122
Guest operating systems that support Dynamic Memory 4123
Dynamic Memory / Hyper-V Memory Management FAQs 4125
Do I need to manually configure the size of the page file after installing the Hyper-V role? 4125
How is the memory divided up between the host operating system and running virtual machines on Windows Server 2012 R2? 4125
My virtual machine failed to start with the following error message “Not enough memory in the system to start the virtual machine” or “Could not initialize memory: Ran out of memory (0x8007000E)”. How can I tell how much memory is available for virtual machines? 4125
Should I configure the MemoryReserve registry setting? 4126

Hyper-V Offloaded Data Transfer Overview 4126
Key benefits 4126
Requirements 4126

Technical overview ... 4127

Hyper-V Replica Feature Overview ... 4127
Key benefits ... 4128
Key features ... 4128

Hyper-V Resource Metering Overview ... 4128
Key benefits ... 4128
Use of network metering port ACLs ... 4129
Virtual machine metrics ... 4129

Hyper-V Support for Large Sector Disks Overview ... 4129
Key benefits ... 4130
Requirements ... 4130
Technical overview ... 4130
Support for improved performance of VHDs on 512e disks ... 4130
Support for hosting VHDs on native 4 KB disks ... 4131

Hyper-V Virtual Fibre Channel Overview ... 4131
Key benefits ... 4131
Requirements ... 4131
Technical overview ... 4132
NPIV support ... 4132
Virtual SAN support ... 4132
Tape library support ... 4133
Live migration ... 4133
MPIO functionality ... 4133
See also ... 4134

Implement Hyper-V Virtual Fibre Channel ... 4134
Prerequisites ... 4134
Step 1: Review considerations for live migration ... 4136
Step 2: Review considerations for MPIO connectivity to virtual Fibre Channel storage ... 4136
Step 3: Configure a virtual Fibre Channel adapter ... 4137
Terms and definitions ... 4138
See also ... 4139

Hyper-V Virtual Hard Disk Format Overview ... 4139
Key benefits ... 4139
Technical overview ... 4139

Hyper-V Support for Scaling Up and Scaling Out Overview ... 4140
Key benefits ... 4141
Requirements ... 4141
Technical overview ... 4141

Virtual Machine Storage Migration Overview 4142
Key benefits 4143
Requirements 4143
Technical overview 4143

Virtual Machine Live Migration Overview 4144
Key benefits 4145
Requirements 4145
Performance options 4146
Technical overview 4147
See Also 4150

Simplified Import Overview 4150
Key benefits 4150
Requirements 4151
Technical overview 4151

Hyper-V Network Virtualization Overview 4151
Feature description 4152
Practical applications 4153
Important functionality 4154
Software requirements 4156
See also 4156

What's New in Hyper-V Network Virtualization in Windows Server 2012 R2 4157
Feature description 4157
New and changed functionality 4157
Inbox HNV Gateway 4158
HNV Architecture 4158
HNV interoperability with Hyper-V Virtual Switch Extensions 4159
HNV VM Network Diagnostics 4159
Dynamic IP Address Learning 4159
HNV + Windows NIC Teaming 4160
NVGRE Encapsulated Task Offload 4160
See also 4161

Hyper-V Network Virtualization technical details 4161
Hyper-V Network Virtualization Concepts 4162
Routing in Hyper-V Network Virtualization 4164
Routing Between Virtual Subnets 4165
Routing Outside a Virtual Network 4165
Private Cloud (Routing) 4165
Hybrid Cloud (Site to site VPN) 4166
Packet Encapsulation 4167
Network virtualization through address virtualization 4169

Multitenant deployment example ... 4170
Hyper-V Network Virtualization architecture ... 4172
Hyper-V Network Virtualization Policy Management ... 4174
Summary ... 4174
See also ... 4174

Hyper-V Network Virtualization Gateway Architectural Guide ... 4175
Gateway Scenarios ... 4175
Private Cloud (routing) ... 4176
Hybrid Cloud (S2S VPN) ... 4176
Load Balancer ... 4177
External Router Gateway Configuration ... 4181
Management ... 4181
Gateway Console ... 4181
Multiple Gateways ... 4182
System Center Virtual Machine Manager ... 4182
Windows-Based Gateway Appliances ... 4183
Private Cloud Router Architecture ... 4183
Cross Premise Gateway ... 4184
Hardware Considerations ... 4186
Additional Resources ... 4187

Virtual Hard Disk Sharing Overview ... 4188
Key benefits ... 4188
Guest Failover Clusters ... 4188
Shared virtual hard disk format ... 4188
Supported virtual machines ... 4188
Supported guest operating systems ... 4189
Deploying a shared virtual hard disk in a guest failover cluster ... 4189
Using Cluster Shared Volume ... 4189
Using Scale-Out File Server with SMB 3.0 ... 4190
Requirements ... 4191

Hyper-V Virtual NUMA Overview ... 4192
Introduction to NUMA ... 4192
NUMA topology in Hyper-V virtual machines ... 4192
Virtual NUMA ... 4193
Configuring Virtual NUMA ... 4193
Host settings – NUMA Spanning ... 4193
Virtual machine settings - NUMA ... 4195
Virtual NUMA and Dynamic Memory ... 4196
Virtual NUMA FAQs ... 4196
Why is NUMA virtualized in Windows Server 2012? ... 4197
Do all virtual machines take advantage of virtual NUMA? ... 4197

How much of a benefit does virtual NUMA provide? ... 4197
Can virtual NUMA and Dynamic Memory be used together? ... 4197
How can I tell if my system in NUMA-based? ... 4197
How can I determine if a virtual machine's performance is being impacted by NUMA spanning? ... 4198
What can I do if a performance counters indicate that a virtual machine's virtual NUMA nodes do not align well with the Hyper-V host's physical NUMA topology? ... 4199
How can I see which physical NUMA nodes a virtual machine is using? ... 4200
Why do I see that a virtual machine is assigned memory from different physical NUMA nodes than expected? ... 4205

Automatic Virtual Machine Activation ... 4205
Practical Applications ... 4205
System Requirements ... 4206
How to implement AVMA ... 4206
AVMA keys ... 4206
Reporting and Tracking ... 4206

Online Virtual Hard Disk Resizing Overview ... 4207
Key benefits ... 4207
Requirements ... 4208
Understanding online resizing functionality ... 4208
Expanding a virtual hard disk ... 4208
Shrinking a virtual hard disk ... 4208
Performing online resizing operations ... 4209

Storage Quality of Service for Hyper-V ... 4209
Key benefits ... 4209
Key features ... 4209
Requirements ... 4210
Technical overview ... 4210
Virtual hard disk maximum IOPS ... 4210
Virtual hard disk minimum IOPS threshold notifications ... 4210

Generation 2 Virtual Machine Overview ... 4211
Requirements ... 4211
Generation 2 virtual machine features ... 4212
PXE boot by using a standard network adapter ... 4212
Boot from SCSI controller ... 4212
Secure Boot ... 4212
Device support comparison ... 4212
Generation 2 FAQ ... 4214
Can I run generation 1 and generation 2 virtual machines together? ... 4214
Is there a performance gain to generation 2 virtual machines? ... 4214

What are the benefits of using a generation 2 virtual machine? 4214
How can I enable kernel debugging by using a COM port on a generation 2 virtual machine? 4214
Is RemoteFX supported by generation 2 virtual machines? 4215
Can I attach a physical CD or DVD to a generation 2 virtual machine? 4215
Can I attach a virtual hard disk in VHD format to a generation 2 virtual machine? 4215
Can a VHDX file that was converted from a VHD file be used to boot a generation 2 virtual machine? 4215
Can I resize a VHDX file that contains the boot volume for a generation 2 virtual machine while the virtual machine is running? 4215
What is the maximum size of a VHDX file that is supported by a generation 2 virtual machine? 4215
Can I create a VHDX file that can be used to boot both generation 1 and generation 2 virtual machines? 4215
Can I mount a VHDX file that is used by a generation 2 virtual machine? 4216
Can I tell whether a VHDX file was created by a generation 1 or a generation 2 virtual machine? 4216
Can a generation 2 VHDX file be attached to a generation 1 virtual machine? 4216
Can I import a generation 2 virtual machine to Hyper-V in Windows Server 2012? 4216
Is Secure Boot or UEFI firmware required on the physical host? 4216
Does the UEFI firmware in a generation 2 virtual machine support setup mode for Secure Boot? 4216
Can I run a UEFI shell or other UEFI application in a generation 2 virtual machine? 4216
Where are the boot entries for generation 2 virtual machines stored? 4217
What is the default generation of a new virtual machine? 4217
Can I change the generation of a virtual machine after it has been created? 4217
Why are 64-bit versions of Windows Server 2008 R2 and Windows 7 not supported as a guest operating system for generation 2 virtual machines? 4217
Why are 32-bit versions of Windows 8.1 and Windows 8 not supported as guest operation systems for generation 2 virtual machines? 4217
What is the maximum number of network adapters that can be supported by a generation 2 virtual machine? 4217
What is the maximum number of storage devices (VHDX files or DVD drives), that can be supported by generation 2 virtual machines? 4217
When I create a new virtual machine with the New-VM Windows PowerShell cmdlet, why does the generation 2 virtual machine not have a DVD drive? 4218
Does a generation 2 virtual machine support iSCSI or Fibre Channel boot? 4218
Can I perform a network installation that uses IPv6? 4218
Can I use a virtual floppy disk (.vfd) as an answer file for an unattended installation of a generation 2 virtual machine? 4218

Virtual Machine Specifications for Hyper-V in Windows Server 2012 R2 4218
Generation 1 virtual machine components 4218

Generation 2 virtual machine components ... 4222
Hypervisor top-level functional specification (TLFS) ... 4223

Virtual Machine Connection - Enhanced Session Mode Overview ... 4224
Enhanced session mode ... 4224
Practical applications ... 4224
Configuring enhanced session mode ... 4225
Server settings ... 4225
User settings ... 4225
Guest operating system ... 4226
FAQs ... 4226
Does enhanced session mode require Terminal Services CALs? ... 4226
Can I save the settings for a connection to a virtual machine? ... 4226

What’s New in Hyper-V for Windows Server 2012 R2 ... 4227
Role description ... 4227
New and changed functionality ... 4227
Shared virtual hard disk ... 4228
Resize virtual hard disk ... 4228
Storage Quality of Service ... 4228
Live migrations ... 4229
Improved performance ... 4229
Cross-version live migrations ... 4230
Virtual machine generation ... 4230
Integration services ... 4231
Export ... 4231
Failover Clustering and Hyper-V ... 4232
Enhanced session mode ... 4232
Hyper-V Replica ... 4233
Linux support ... 4233
Management ... 4234
Automatic Virtual Machine Activation ... 4234
Hyper-V Networking ... 4235

What's New in Hyper-V for Windows Server 2012 ... 4235
Role description ... 4235
New and changed functionality ... 4236
Client Hyper-V ... 4236
Dynamic Memory ... 4237
Hyper-V module for Windows PowerShell ... 4237
Hyper-V Replica ... 4238
Importing of virtual machines ... 4238
Live migration ... 4238
Resource metering ... 4239

Significantly increased scale and improved resiliency ... 4239
Simplified authorization ... 4239
SR-IOV ... 4240
Storage migration ... 4240
Storage on SMB 3.0 file shares ... 4240
Virtual Fibre Channel ... 4240
Virtual hard disk format ... 4241
Virtual machine snapshots ... 4241
Virtual NUMA ... 4241
Virtual switch ... 4242
Sleep support ... 4242
Sleep support ... 4242
Removed or deprecated functionality ... 4243
See also ... 4243

Hyper-V Scalability in Windows Server 2012 ... 4243
Virtual machines ... 4244
Server running Hyper-V ... 4246
Failover Clusters and Hyper-V ... 4247

Hyper-V Component Architecture Posters ... 4248

Install Hyper-V and Configure a Virtual Machine ... 4249
Prerequisites ... 4249
Step 1: Install Hyper-V ... 4249
To add the Hyper-V role in Windows Server ... 4250
To enable Client Hyper-V ... 4250
To install the Hyper-V role using the Install-WindowsFeature cmdlet on Windows Server 4251
To install Client Hyper-V using the Get-WindowsOptionalFeature cmdlet ... 4252
Step 2: Create a virtual machine ... 4252
Step 3: Install the guest operating system ... 4253
Step 4: Install or upgrade integration services ... 4254
See also ... 4254

Hyper-V Replica Overview ... 4255
Role/Feature description ... 4255
Key concepts ... 4255
New features starting with Windows Server 2012 R2 ... 4256
Practical applications ... 4256
Hardware requirements ... 4257
Software requirements ... 4257
See also ... 4258

Demonstrate Planned Failover in Hyper-V Replica ... 4258
Prerequisites ... 4259

Step 1: Enable the Hyper-V Role and create a virtual machine 4259
Step 2: Configure the Replica server to accept replication 4259
Step 3: Enable and begin replication of virtual machines 4261
Step 4: Monitor replication 4263
Step 5: Demonstrate test failover 4263
Step 6: Demonstrate planned failover 4264
See also 4265

Deploy Hyper-V Replica 4265

Step 1: Prepare to Deploy Hyper-V Replica 4265
1.1. Make basic planning decisions 4266
1.2. Install the Hyper-V server role 4268
1.3. Configure the firewall 4268
1.4. Configure Hyper-V Replica Broker 4269
1.5. Provide a certificate for encrypted data 4270
See also 4271

Step 2: Enable Replication 4271
2.1 Configure the Replica server 4272
2.2. Configure a Replica server that is part of a failover cluster (optional) 4273
2.3 Enable replication for virtual machines 4274
2.4 Configure primary server to receive replication 4276
See also 4276

Step 3: Test the Replication Deployment 4276
3.1. Perform test failover 4277
See also 4278

Step 4: Perform a Planned Failover 4278
4.1. Perform a planned failover 4278
See also 4279

Step 5: Respond to an Unplanned Failover 4280
5.1. Start the Replica virtual machines 4280
5.2. Start reverse replication 4281
See also 4282

Step 6: Configure and Use Extended Replication 4282
Enable extended replication 4283
Monitor extended replication 4284
Pause and resume extended replication 4284
Resynchronization 4285
Test failover with an extended Replica server 4286
New options for planned failover 4286

Unplanned failover from primary to Replica with an extended Replica server 4288
Unplanned failover when both primary and Replica servers go offline 4289
Moving primary or Replica virtual machines to different servers or clusters 4291
Removing replication or extended replication 4291
See also 4291

Hyper-V Module for Windows PowerShell 4292
Get updated Help 4292
See Also 4292

Hyper-V Configuration 4292

Configure Live Migration and Migrating Virtual Machines without Failover Clustering 4293
Prerequisites 4293
Step 1: [Optional] Configure constrained delegation 4295
Step 2: Configure the source and destination computers for live migration 4297
Step3: [optional] Configure performance options for live migration 4298
Step 4: Move a running virtual machine 4299
Step 5: [Optional] Move a running virtual machine again (back to the original host or to another host) 4299
See also 4300

Configure Storage Quality of Service 4300
To configure a virtual hard disk for Storage Quality of Service 4300

Configure a Shared Virtual Hard Disk 4301
To create and enable a shared virtual hard disk on a virtual machine 4301

Configure Online Virtual Hard Disk Resize 4302
Prerequisites 4302
To create a virtual hard disk for a virtual machine 4302
Adding a virtual hard disk to the virtual machine 4303
To expand the size of a virtual hard disk 4303
Expand the volume of a virtual hard disk connected to a virtual machine 4304
To shrink the size of a virtual hard disk 4304

Linux Virtual Machines on Hyper-V 4305
In this section: 4305

CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map ... 4306

SUSE virtual machines feature distribution map 4308

Ubuntu virtual machines feature distribution map 4310

Feature Descriptions 4313
Core 4313

Networking 4314
Storage 4314
Memory 4315
Video 4317
Miscellaneous 4317

Networking Overview 4318
Networking in Windows Server 2012 and Windows Server 2012 R2 4318

What's New in Networking in Windows Server 2012 R2 4324
New and improved technologies for Windows Server Networking 4324
802.1X Authenticated Wired Access 4324
802.1X Authenticated Wireless Access 4324
Domain Name System 4324
Dynamic Host Configuration Protocol 4324
Hyper-V Network Virtualization 4325
Hyper-V Virtual Switch 4325
Internet Protocol Address Management 4325
Remote Access (DirectAccess, Routing and Remote Access) 4325
Virtual Receive-side Scaling 4325
Windows Server Gateway 4325

What's New in Networking in Windows Server 2012 4326
New and improved technologies for Windows Server Networking 4326
802.1X Authenticated Wired and Wireless Access 4326
BranchCache 4326
Data Center Bridging (DCB) 4327
Domain Name System (DNS) 4327
Dynamic Host Configuration Protocol (DHCP) 4327
Hyper-V Network Virtualization 4327
Hyper-V Virtual Switch 4328
IP Address Management (IPAM) 4328
Low Latency Workloads Technologies 4328
Network Load Balancing (NLB) 4328
Network Policy and Access Services 4329
NIC Teaming 4329
Quality of Service (QoS) 4329
Remote Access 4330
DirectAccess and RRAS unified server role 4330
Windows Firewall with Advanced Security 4330

Core Network Guide and Companion Guides Overview 4331
Introduction to the Windows Server Core Network 4331
Windows Server 2012 Core Network Guide 4332

Windows Server 2012 Core Network Companion Guides 4332
Core Network Companion Guide: Server Certificate Deployment 4333
Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access 4333
Core Network Companion Guide: Computer and User Certificate Deployment 4334
Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode 4334
Core Network Companion Guide: Group Policy Deployment 4335

Core Network Guide 4335
About this guide 4335
Network hardware requirements 4336
What this guide does not provide 4337
Technology Overviews 4337
Active Directory Domain Services 4337
DNS 4337
DHCP 4337
TCP/IP 4338
Core Network Overview 4338
Core Network Components 4339
Router 4339
Static TCP/IP configurations 4339
Active Directory Domain Services global catalog and DNS server DC1 4340
DHCP server DHCP1 4340
Client computers 4340
Core Network Planning 4340
Planning subnets 4340
Planning basic configuration of all servers 4341
Planning naming conventions for computers and devices 4341
Planning static IP addresses 4342
Planning the deployment of DC1 4342
Planning the name of the forest root domain 4342
Planning the forest functional level 4343
Planning DNS zones 4344
Planning domain access 4345
Planning the deployment of DHCP1 4345
Planning DHCP servers and DHCP forwarding 4345
Planning IP address ranges 4346
Planning subnet masks 4346
Planning exclusion ranges 4347
Planning TCP/IP static configuration 4348
Core Network Deployment 4349
Configuring All Servers 4349
Rename the computer 4349

Configure a static IP address 4350
Deploying DC1 4351
Install AD DS and DNS for a New Forest 4352
Create a User Account in Active Directory Users and Computers 4354
Add a Group 4355
Assign Group Membership 4355
Configure a DNS Reverse Lookup Zone 4356
Joining Server Computers to the Domain and Logging On 4357
Deploying DHCP1 4358
Install Dynamic Host Configuration Protocol (DHCP) 4359
Create and Activate a New DHCP Scope 4360
Joining Client Computers to the Domain and Logging On 4361
Deploying optional features for network access authentication and Web services 4362
Deploying NPS1 4363
Planning the deployment of NPS1 4364
Install Network Policy Server (NPS) 4364
Register the NPS Server in the Default Domain 4365
Deploying WEB1 4366
Install the Web Server (IIS) server role 4366
Additional Technical Resources 4367
Appendices A through E 4368
Appendix A - Renaming computers 4368
Windows Server 2008 R2 and Windows 7 4368
Windows Server 2008 and Windows Vista 4369
Windows Server 2003 and Windows XP 4369
Appendix B - Configuring static IP addresses 4369
Windows Server 2008 R2 4370
Windows Server 2008 4370
Windows Server 2003 4371
Appendix C – Joining computers to the domain 4372
Windows Server 2008 R2 and Windows 7 4372
Windows Server 2008 and Windows Vista 4373
Windows Server 2003 and Windows XP 4373
Appendix D – Log on to the domain 4374
Windows Server 2008 R2 and Windows 7 4374
Windows Server 2008 and Windows Vista 4374
Windows Server 2003 and Windows XP 4375
Appendix E - Core Network Planning Preparation Sheet 4375
Installing Active Directory Domain Services and DNS 4376
Pre-installation configuration items for AD DS and DNS 4376
AD DS and DNS installation configuration items 4376
Configuring a DNS Reverse Lookup Zone 4377
Installing DHCP 4377

Pre-installation configuration items for DHCP 4378
DHCP installation configuration items 4378
Creating an exclusion range in DHCP 4379
Creating a new DHCP scope 4379
Installing Network Policy Server (optional) 4380
Pre-installation configuration items 4380
Network Policy Server installation configuration items 4380

Core Network Companion Guide: Server Certificate Deployment 4381
Prerequisites for using this guide 4381
About this guide 4382
Requirements for deploying server certificates 4382
What this guide does not provide 4383
Technology overviews 4383
EAP 4383
EAP in Windows Server 2012 4384
PEAP 4384
Features of PEAP 4385
Active Directory Certificate Services 4385

Server Certificate Deployment Overview 4385
Server certificate deployment components 4386
CA1 running the AD CS server role 4386
CAPolicy.inf 4387
Copy of the RAS and IAS servers certificate template 4387
Additional CA1 configuration 4387
WEB1 running the Web Services (IIS) server role 4387
Virtual directory for the CRL and AIA 4387
DC1 running the AD DS and DNS server roles 4388
Group Policy default domain policy 4388
DNS alias (CNAME) resource record 4388
NPS1 running the Network Policy Server role service of the Network Policy and Access Services server role 4388
Group Policy applied and certificate enrolled to NPS1 4388
Server certificate deployment process overview 4388

Server Certificate Deployment Planning 4389
Plan basic server configuration 4390
Plan domain access 4390
Plan the location and name of the virtual directory on your Web server 4390
Plan a DNS alias (CNAME) record for your Web server 4390
Plan configuration of CAPolicy.inf 4391
Plan configuration of the CDP and AIA extensions on CA1 4392
Plan the copy operation between the CA and the Web server 4393

Plan the configuration of the server certificate template on the CA 4393
Server Certificate Deployment 4393
Create an Alias (CNAME) Record in DNS for WEB1 4394
Configure WEB1 to Distribute Certificate Revocation Lists (CRLs) 4395
Prepare the CAPolicy.inf File 4396
Install the Certification Authority 4398
Configure the CDP and AIA Extensions on CA1 4400
Copy the CA Certificate and CRL to the Virtual Directory 4402
Configure the Server Certificate Template 4403
Configure Server Certificate Autoenrollment 4404
Refresh Group Policy 4405
Verify NPS Server Enrollment of a Server Certificate 4405
Additional Resources 4407
Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access 4408
About this guide 4408
Requirements 4409
What this guide does not provide 4410
Technology overviews 4411
IEEE 802.1X 4411
802.1X-capable wireless access points (APs) 4411
Wireless clients 4412
Support for IEEE 802.11 Standards 4412
Wireless network security methods 4413
Wireless authentication 4413
Wireless security encryption 4414
Enhanced encryption 4415
Wireless authentication and encryption pairs 4415
Active Directory Doman Services (AD DS) 4416
Active Directory Users and Computers 4416
Group Policy Management 4416
Server certificates 4417
EAP, PEAP, and PEAP-MS-CHAP v2 4417
Network Policy Server 4418
Bootstrap profiles 4419

Wireless Access Deployment Overview 4420
Wireless access deployment components 4421
802.1X-capable Wireless access points 4422
Active Directory Domain Services 4422
NPS 4424
Wireless client computers 4424
Wireless access deployment process 4424
Wireless Access Deployment Planning 4425
Planning wireless AP installations 4426
Verify wireless AP support for standards 4426
Identify areas of coverage for wireless users 4426
Determine where to install wireless APs 4427
Plan the configuration of wireless APs in NPS 4427
Plan the use of PEAP Fast Reconnect 4428
Wireless AP configuration 4428
Planning wireless client configuration and access 4430
Planning support for multiple standards 4431
Planning restricted access to the wireless network 4431
Planning methods for adding new wireless computers 4432
Wireless Access Deployment 4433
Deploying and Configuring Wireless APs 4433
Specify Wireless AP Channel Frequencies 4433
Configure Wireless APs 4433
Creating Security Groups for Wireless Users 4435
Create a Wireless Users Security Group 4435
Add Users to the Wireless Users Security Group 4435
Configuring Wireless Network (IEEE 802.11) Policies 4436
Open or Add and Open a Group Policy Object 4436
Activate Default Wireless Network (IEEE 802.11) Policies 4437
Open Wireless Network (IEEE 802.11) Policies for Editing 4438
Configure the New Wireless Network Policy 4439
Configure a Wireless Connection Profile for PEAP-MS-CHAP v2 4439
Set the Preference Order for Wireless Connection Profiles 4442
Define Network Permissions 4443
Configure the New Windows XP Wireless Network Policy 4444
To configure a Windows XP wireless connection profile for PEAP-MS-CHAP v2 4445
Configuring your NPS Server 4446
Register NPS in Active Directory Domain Services 4446
Configure a Wireless AP as an NPS RADIUS Client 4447
Create NPS Policies for 802.1X Wireless Using a Wizard 4448
Joining New Wireless Computers to the Domain 4451
Join the Domain and Log On by using Wireless Method 1 4452
Join the Domain and Log On by using Wireless Method 2 4453

Join the Domain and Log On by using Wireless Method 3 4455
Additional Resources 4458

Core Network Companion Guide: Computer and User Certificates Deployment 4459
Prerequisites for using this guide 4459
About this guide 4460
Requirements for deploying computer and user certificates 4461
What this guide does not provide 4462
Technology overviews 4462
AD CS 4462
Client computer and user certificates 4463
Certificate store 4464
EAP 4464
EAP in Windows Server 2012 4464
PEAP 4465
Features of PEAP 4465
EAP-TLS and PEAP-TLS deployment overview 4466
Group Policy 4466

Certificate Deployment Overview 4466
Computer and user certificate deployment components 4467
CA1 running the AD CS server role 4467
Copy of the User certificate template 4467
Copy of the Workstation Authentication certificate template 4468
DC1 running the AD DS and DNS server roles 4468
Client1 4468
Computer and user certificate deployment process 4468

Certificate Deployment Planning 4469
Planning certificate configuration 4469
Client computer display of certificates 4470
Planning User certificate enrollment 4470

Certificate Deployment 4471

Certificate Deployment for Computers 4471

Configure the Workstation Authentication Certificate Template 4471

Configure Computer Certificate Autoenrollment 4473

Refresh Group Policy 4474

Certificate Deployment for Users 4474

Configure the User Certificate Template 4474

Configure User Certificate Autoenrollment 4476

Refresh Group Policy 4477

Verify Computer or User Enrollment of a Certificate 4477

Additional Resources 4478

Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode 4479
- Prerequisites for using this guide 4480
- About this guide 4481
- What this guide does not provide 4481
- Technology overviews 4482
  - BranchCache 4482
  - Group Policy 4482

BranchCache Hosted Cache Mode Deployment Overview 4482
- Hosted Cache Server deployment components 4482
  - HCS1 in the branch office 4483
  - WEB1 in the cloud data center 4483
  - FILE1 in the cloud data center 4483
  - DC1 in the main office 4484
  - Client computers in the branch office 4484
- Hosted Cache Server deployment process overview 4484

BranchCache Hosted Cache Mode Deployment Planning 4485
- Plan basic server configuration 4485
- Plan domain access 4485
- Plan the location and size of the hosted cache 4486
- Plan the share to which the content server packages are to be copied 4486
- Plan prehashing and data package creation on content servers 4486

BranchCache Hosted Cache Mode Deployment 4486

Install the BranchCache Feature and Configure the Hosted Cache Server by Service Connection Point 4487

Move and Resize the Hosted Cache (Optional) 4488

Prehash and Preload Content on the Hosted Cache Server (Optional) 4489

Create Content Server Data Packages for Web and File Content (Optional) 4490

Import Data Packages on the Hosted Cache Server (Optional) 4491

Configure Client Automatic Hosted Cache Discovery by Service Connection Point 4492

Additional Resources 4493

Core Network Companion Guide: Group Policy Deployment 4494
About this guide 4494
Requirements 4495
What this guide does not provide 4496
Technology overview of Group Policy 4496

Group Policy Deployment Overview 4497
Group Policy deployment components 4497
Domain controller: AD-DNS-01 4497
Group Policy: GPO_Membership 4497
WMI Filters 4497
Membership group: GRP_Membership 4497
Exception group: GRP_Exception 4498
Group Policy deployment process 4498

Group Policy Deployment Planning 4499
Planning the membership and exception groups 4499
Planning domain access 4499

Group Policy Deployment 4500

Create the Membership Group 4500

Create the Exception Group 4501

Create the WMI Filters 4501
Create WMI Filters 4502

Create the Group Policy Objects for Each Version of Windows 4504

Apply Security Group Filters and WMI Filters to the GPOs 4504

Link the GPOs to the Domain 4505

Add Computers to the Exception Group 4506

Add Test Computers to the Membership Group 4507

Add Production Computers to the Membership Group 4508

Additional Resources 4509

802.1X Authenticated Wired Access Overview 4509
Feature description 4509
Important terminology and technology overviews 4510
IEEE 802.1X 4510
IEEE 802.1X-capable wired Ethernet switches 4510
IEEE 802.3 Ethernet 4510

Network Policy Server 4510
Server certificates 4511
EAP 4511
New and changed functionality 4511
See also 4511

What's New in 802.1X Authenticated Wired Access in Windows Server 2012 R2 4512
Extending the use of passwords for Enterprise wired Ethernet access 4512

Managing the New Wired Network (IEEE 802.3) Policies Settings 4513
Configuring Wired Network (IEEE 802.3) Policies 4514
Configure a wired connection profile for PEAP-MS-CHAP v2 4514
Configure a wired connection profile for PEAP-TLS 4516
Configure a wired connection profile for EAP-TLS 4518
Per-setting details 4519
General - settings 4519
Security - settings 4521
See Also 4523

802.1X Authenticated Wireless Access Overview 4523
Feature description 4524
Important terminology and technology overviews 4524
IEEE 802.1X 4524
IEEE 802.1X-capable wired Ethernet switches 4524
IEEE 802.11 wireless 4524
Network Policy Server 4525
Server certificates 4525
EAP 4525
New and changed functionality 4525
See also 4526

What's New in 802.1X Authenticated Wireless Access in Windows Server 2012 R2 4527
Role and technology description 4527
New and changed functionality 4527
Support for 802.11ac 4528
Wireless Display 4529
Extending the use of passwords for Enterprise wireless access 4529

Improvements to Certificate-based Authentication 4530
Certificate filtering 4531
Certificate weight as a filtering mechanism 4533
Private Key Protected Registry Certificates 4534
Certificate filtering compatibility information 4534

Wireless LAN Service Overview 4534

Enabling the Wireless LAN Service 4534

Managing the Wireless Network (IEEE 802.11) Policies 4535

Managing the New Wireless Network (IEEE 802.11) Policies Settings 4536
Configuring Wireless Network (IEEE 802.11) Policies 4537
Configure a wireless connection profile for PEAP-MS-CHAP v2 4537
Configure a wireless connection profile for PEAP-TLS 4539
Configure a wireless connection profile for EAP-TLS 4541
Wireless Network (IEEE 802.11) Policies settings 4543
Wireless Policies for computers running Windows Vista and subsequent releases of Windows 4543
General - settings 4544
Import and Export Wireless Network Profiles 4545
Open for import a profile (Import Profiles) 4545
Save export profile as (Export Profiles) 4545
Connection - settings 4546
Network Permissions - settings 4547
New Permissions Entry - settings 4548
Security – configuration items 4549
Select the security methods for this network - settings 4549
Select a network authentication method - settings 4550
Additional resources 4552

Advanced Security Settings for Wired and Wireless Network Policies 4552
Advanced security settings for the Wireless Network Policies and Wired Network Policies 4553
IEEE 802.1X - settings 4553
Single Sign On - settings 4554
Fast Roaming settings 4555
Additional resources 4556

Managing the Windows XP-based Wireless Network Policies 4556
General - settings 4557
802.1X - settings 4558
Network Properties - settings 4561
Preferred Networks - settings 4563

New Wireless Connection Processes 4564

Connections to Organization Networks that Require RADIUS Server Authentication by Wireless Clients 4565

Connections to Organization Networks with Multiple RADIUS Servers 4566
Deployment requirements for multiple RADIUS servers 4567

Connections to Networks Whose RADIUS Server Certificate is Issued by a Public Certification Authority 4568

Connections to Public Wi-Fi Hotspots 4568
Connections to public Wi-Fi hotspots 4569
Selecting the Authentication Protocol 4571
Caveats for third party authentication methods 4573

BranchCache Overview 4573
What is BranchCache? 4574
BranchCache modes 4575
Increased cache availability 4576
Centralized caching for multiple-subnet branch offices 4576
BranchCache-enabled content servers 4577
Web servers 4577
File servers 4577
Application servers 4578
BranchCache and the cloud 4578
Content information versions 4578
BranchCache installation guide 4579
Operating system versions for BranchCache 4580
Operating systems for BranchCache client computer functionality 4581
Operating systems for BranchCache content server functionality 4581
Operating systems for BranchCache hosted cache server functionality 4581
BranchCache Security 4581
How BranchCache generates content information 4582
Content information details 4582
Content flow and processes 4583
BranchCache processes: Request content 4583
BranchCache processes: Locate content 4584
BranchCache processes: Retrieve content 4585
Security threats 4586
BranchCache processes: Cache content 4586
Hosted cache mode cache population 4587
Cache Security 4588
Client computer cache security 4588
Hosted cache server cache security 4588

BranchCache Learning Roadmap 4589
Prerequisite information 4590
Level 100 4590
Level 200 4591
Level 300 4592
Optional information 4593

What's New in BranchCache 4594
New BranchCache Group Policy settings 4595
Additional BranchCache Group Policy settings 4597
BranchCache Network Shell and Windows PowerShell Commands 4598
BranchCache Deployment Guide 4599
What this guide provides 4600
BranchCache deployment requirements 4600
Choosing a BranchCache Design 4601
Deploy BranchCache 4602
Install and Configure Content Servers 4603
Install Content Servers that Use the BranchCache Feature 4603
Install the BranchCache Feature 4603
Configure Windows Server Update Services (WSUS) Content Servers 4604
Install File Services Content Servers 4605
Configure the File Services server role 4605
Install a New File Server as a Content Server 4605
Configure an Existing File Server as a Content Server 4606
Enable Hash Publication for File Servers 4607
Enable Hash Publication for Non-Domain Member File Servers 4608
Enable Hash Publication for Domain Member File Servers 4609
Create the BranchCache File Servers Organizational Unit 4609
Move File Servers to the BranchCache File Servers Organizational Unit 4610
Create the BranchCache Hash Publication Group Policy Object 4611
Configure the BranchCache Hash Publication Group Policy Object 4611
Enable BranchCache on a File Share (Optional) 4613
Deploy Hosted Cache Servers (Optional) 4613
Prehashing and Preloading Content on Hosted Cache Servers (Optional) 4614
Configure BranchCache Client Computers 4615

Use Group Policy to Configure Domain Member Client Computers 4615

Use Windows PowerShell to Configure Non-Domain Member Client Computers 4618

Configure Firewall Rules for Non-Domain Members to Allow BranchCache Traffic 4619
[MS-PCCRD]: Peer Content Caching and Retrieval Discovery Protocol 4619
[MS-PCCRR]: Peer Content Caching and Retrieval: Retrieval Protocol 4620

Verify Client Computer Settings 4620

Additional Resources 4621

Data Center Bridging (DCB) Overview 4621
Feature description 4622
Practical applications 4622
Important functionality 4622
Hardware requirements 4623
See also 4623

DCB Windows PowerShell User Scripting Guide 4624
DCB Overview 4624
Prerequisites 4624
Windows Server 2012 4624
Enable Data Center Bridging 4625
Import DCB PowerShell modules 4626
Windows PowerShell tips 4627
Find Windows PowerShell commands 4627
Learn command syntax 4628
Learn more about command parameters 4629
DCB configurations 4631
Traffic Class management 4632
Create Traffic Class 4632
Display traffic class 4633
Modify Traffic Class 4633
Remove Traffic Class 4634
Priority Flow Control settings: 4634
Enable Priority Flow Control 4634
Display Priority Flow Control… 4635
Disable Priority Flow Control 4635
Application Priority assignment 4635
Create QoS Policy 4636
Display QoS Policy 4637
Modify QoS Policy 4638
Remove QoS Policy 4639
DCBX settings 4639

Set the Willing bit ... 4640
Display the Willing bit ... 4640
DCB configuration on network adapters ... 4640
Display DCB setting on network adapters ... 4640
Disable DCB on network adapters ... 4641
Enable DCB on network adapters ... 4642
More Information ... 4642

Domain Name System (DNS) Overview ... 4642
Role description ... 4643
Practical applications ... 4643
New and changed functionality ... 4643
Server Manager information ... 4643

What's New in DNS ... 4644
Role description ... 4644
See Also ... 4645

What's New in DNS Server in Windows Server 2012 ... 4645
New and changed functionality ... 4645
Enhanced support for DNSSEC ... 4645
Windows PowerShell support ... 4646

What's New in DNS Client in Windows 8 ... 4646
New and changed functionality ... 4646

What's New in DNS Server in Windows Server 2012 R2 ... 4647
New and changed functionality ... 4647
Enhanced zone level statistics ... 4648
Enhanced DNSSEC support ... 4649
Enhanced Windows PowerShell support ... 4649

What's New in DNS Client in Windows 8.1 ... 4650
New and changed functionality ... 4650

DNSSEC in Windows Server 2012 ... 4652
In this guide ... 4652
See Also ... 4653

Overview of DNSSEC ... 4653
In this section ... 4653
How DNSSEC works ... 4654
DNSSEC-related resource records ... 4655
Validation of DNS responses ... 4657
Example DNS queries ... 4660
DNSSEC scenarios ... 4663

See also ... 4665

DNSSEC in Windows ... 4666
See also ... 4666

DNS Servers ... 4667
In this section ... 4667
DNSSEC support in Windows Server ... 4668
DNSSEC mixed-mode deployment ... 4669
WINS lookup for signed zones ... 4670
DNSSEC in Windows Server 2012 ... 4672
Online signing and dynamic updates ... 4673
Read-only domain controllers (RODCs) ... 4674
DNSSEC in DNS Manager ... 4675
The Key Master ... 4676
Transferring the Key Master role ... 4677
Seizing the Key Master role ... 4679
DNSSEC standards ... 4682
Zone transfers ... 4683
See also ... 4683

DNS Clients ... 4683
In this section ... 4684
Security-aware client ... 4684
DNS queries and responses ... 4685
Managing validation ... 4686
DNS client resolver behavior ... 4687
See also ... 4688

DNS Zones ... 4688
In this section ... 4689
Signing a zone ... 4689
DNS Manager ... 4690
Zone Signing Wizard ... 4690
Signing options ... 4691
Custom parameters ... 4693
Existing parameters ... 4704
Default parameters ... 4705
Parameter values ... 4705
Unsign zone wizard ... 4709
DNSSEC properties page ... 4710
Modifying DNSSEC properties ... 4712
Windows PowerShell ... 4713
Sign a zone ... 4713

Unsign a zone ... 4714
Review DNSSEC properties ... 4714
Review the status of signing keys ... 4715
Modify DNSSEC properties ... 4718
See also ... 4718

Trust Anchors ... 4718
In this section ... 4718
Types of trust anchors ... 4719
Delegations and the chain of trust ... 4721
The delegation signer (DS) record ... 4724
Working with trust anchors ... 4726
Trust anchor status ... 4729
See also ... 4730

The NRPT ... 4730
In this section ... 4730
Introduction to the NRPT ... 4731
NRPT rule processing ... 4732
Filter NRPT policy ... 4733
View NRPT policy ... 4734
See also ... 4735

DNSSEC Deployment Planning ... 4736
See also ... 4736

Why DNSSEC ... 4736
In this section ... 4736
Why to deploy DNSSEC ... 4736
DNS spoofing ... 4737
Signing domain zones ... 4738
Signed zones ... 4738
See also ... 4739

Stage a DNSSEC Deployment ... 4739
In this section ... 4739
Identify goals ... 4740
DNSSEC pilot ... 4740
DNSSEC staging ... 4741
See also ... 4742

DNSSEC Performance Considerations ... 4743
In this section ... 4743
DNS performance ... 4743
Cryptographic algorithms ... 4744

See also 4745

DNSSEC Requirements 4745
See also 4749

Deploy DNSSEC with Windows Server 2012 4749
Deploying DNSSEC 4749
DNSSEC concepts 4749
DNSSEC deployment checklists 4749
See also 4750

Checklist: Deploy DNSSEC 4750
See also 4751

Procedure: Review Name Resolution Policy Settings 4752
Review NRPT configuration 4752
Review effective name resolution policy 4754
See also 4755

Checklist: Sign a Zone 4755
See also 4755

Procedure: Customize Zone Signing Parameters 4756
Customize zone signing parameters in DNS Manager 4756
Customize zone signing parameters in Windows PowerShell 4758
See also 4760

Procedure: Sign the Zone with Parameters of an Existing Zone 4760
Sign a zone with existing zone signing parameters in DNS Manager 4760
Sign a zone with existing zone signing parameters in Windows PowerShell 4762
See also 4764

Procedure: Use Default Settings to Sign the Zone 4764
Sign a zone using default parameters in DNS Manager 4764
Sign a zone using default parameters in Windows PowerShell 4765
See also 4765

Procedure: Verify Zone Signing 4765
Verify zone signing in DNS Manager 4766
Verify zone signing in Windows PowerShell 4768
See also 4770

Checklist: Distribute Trust Anchors 4770
See also 4771

Procedure: Enable Automatic Update of Trust Anchors on Key Rollover 4771
Enable automatic update of trust anchors on key rollover using DNS Manager 4772

Enable automatic update of trust anchors on key rollover using Windows PowerShell 4772
See also 4773

Procedure: Distribute Trust Anchors in Active Directory 4774
Enable trust anchor distribution in Active Directory using DNS Manager 4774
Enable trust anchor distribution in Active Directory using Windows PowerShell 4775
See also 4775

Procedure: Export a Trust Point 4776
Export a trust point using File Explorer 4776
Export a trust point using Windows PowerShell 4777
See also 4777

Procedure: Import a Trust Point 4777
Import a trust point using DNS Manager 4778
Import a trust point using Windows PowerShell 4778
See also 4779

Procedure: Add a Trust Point 4779
Add a trust point in DNS Manager 4780
Add a trust point in Windows PowerShell 4783
See also 4784

Procedure: Deploy a Root Trust Point 4784
Deploy a root trust point using Windows PowerShell 4785
See also 4786

Checklist: Deploy DNSSEC Policies to DNS Clients 4786
See also 4787

Procedure: Configure the NRPT 4787
Configure the NRPT with the Group Policy Management console 4787
Configure the NRPT with Windows PowerShell 4788
See also 4790

Procedure: Verify Name Resolution Policy 4790
Review NRPT configuration 4791
Review effective name resolution policy 4792
Verify DNSSEC validation 4793
See also 4798

Checklist: Review and Manage a Signed Zone 4798
See also 4799

Procedure: Review DNSSEC Parameters and Settings 4799
Review DNSSEC parameters and settings in DNS Manager 4800

Review DNSSEC parameters and settings in Windows PowerShell 4800
See also 4804

Procedure: Enable DNS Diagnostic Events 4804
In this section 4804
View and modify the event logging status 4804
See also 4808

Procedure: Review Key Rollover Status 4808
Review key rollover status in DNS Manager 4808
Review key rollover status in Windows PowerShell 4808
See also 4811

Checklist: Revert to an Unsigned Zone 4811
See also 4812

Procedure: Configure the NRPT 4812
Configure the NRPT with the Group Policy Management console 4813
Configure the NRPT with Windows PowerShell 4813
See also 4815

Procedure: Verify Name Resolution Policy 4815
Review NRPT configuration 4816
Review effective name resolution policy 4817
Verify DNSSEC validation 4818
See also 4823

Procedure: Remove a Trust Point 4823
Delete a trust point using DNS Manager 4824
Delete a trust point using Windows PowerShell 4824
See also 4825

Procedure: Unsign a Zone 4825
Unsign a zone in DNS Manager 4825
Unsign a zone in Windows PowerShell 4825
See also 4826

Procedure: Verify DNS Resolution 4826
Verify DNSSEC validation 4826
See also 4831

Checklist: Manage Signing Keys 4831
See also 4832

Procedure: Review Signing Keys 4832
Review DNSSEC signing keys in DNS Manager 4832

Review DNSSEC signing keys in Windows PowerShell 4833
See also 4835

Procedure: Retire a Signing Key 4836
Revoke a signing key using DNS Manager 4836
Revoke a signing key using Windows PowerShell 4837
See also 4838

Checklist: Move the Key Master Role 4838
See also 4838

Procedure: Move the Key Master Role 4839
Move the Key Master role using DNS Manager 4839
Move the Key Master role using Windows PowerShell 4840
See also 4841

Procedure: Seize the Key Master Role 4841
Seize the Key Master role using DNS Manager 4842
Seize the Key Master role using Windows PowerShell 4842
See also 4844

Checklist: Reconfigure Zone Signing Parameters on a Signed Zone 4844
See also 4845

Procedure: Modify Zone Signing Parameters 4845
Modify DNSSEC parameters and settings in DNS Manager 4845
Modify DNSSEC parameters and settings in Windows PowerShell 4846
See also 4848

Procedure: Review DNSSEC Parameters and Settings 4848
Review DNSSEC parameters and settings in DNS Manager 4849
Review DNSSEC parameters and settings in Windows PowerShell 4849
See also 4853

Checklist: Perform an Emergency Key Revocation 4853
See also 4853

Procedure: Import a Trust Point 4854
Import a trust point using DNS Manager 4854
Import a trust point using Windows PowerShell 4855
See also 4855

Procedure: Replace Signing Keys 4856
Replace a signing key using DNS Manager 4856
Replace a signing key using Windows PowerShell 4857
See also 4858

Checklist: Perform a Manual Key Rollover 4858
See also 4859

Procedure: Perform a Manual Signing Key Rollover 4859
Perform a manual signing key rollover using DNS Manager 4860
Perform a manual signing key rollover using Windows PowerShell 4860
See also 4862

Procedure: Review Key Rollover Status 4862
Review key rollover status in DNS Manager 4862
Review key rollover status in Windows PowerShell 4863
See also 4865

Checklist: Secure Zone Transfers 4866
See also 4866

Procedure: Deploy Certificates for DNS Server Authentication 4866
Configure certificate templates 4867
Publish certificate templates 4868
Enable certificate auto-enrollment 4869
See also 4869

Procedure: Deploy IPsec Policy to DNS Servers 4870
Configure IPsec policy 4870
Certificate Selection 4872
See also 4873

Appendix A: DNSSEC Terminology 4873

Appendix B: Windows PowerShell for DNS Server 4878
See also 4890

Step-by-Step: Demonstrate DNSSEC in a Test Lab 4890
In this guide 4890
DNSSEC overview 4891
DNS threats 4891
How DNSSEC works 4891
Digital signatures 4892
Zone signing 4892
Authenticated denial of existence 4892
Trust anchors 4892
DNSSEC key management 4892
DNSSEC-aware clients 4893
NRPT 4893
Scenario overview 4893
Hardware and software requirements 4894

Steps for configuring the test lab ... 4894
Configure DC1 ... 4894
Install the operating system and configure TCP/IP on DC1 ... 4895
Install Active Directory and DNS on DC1 ... 4895
Configure the sec.contoso.com DNS zone ... 4897
Enable Remote Desktop on DC1 ... 4898
Configure DNS1 ... 4898
Install the operating system and configure TCP/IP on DNS1 ... 4898
Install and configure DNS on DNS1 ... 4899
Optional: Install Network Monitor on DNS1 ... 4900
Configure DC2 ... 4900
Install the operating system and configure TCP/IP on DC2 ... 4901
Install Active Directory and DNS on DC2 ... 4901
Configure Client1 ... 4902
Install the operating system and configure TCP/IP on Client1 ... 4902
Join Client1 to the contoso.com domain ... 4903
Pin Windows PowerShell to the taskbar ... 4903
DNSSEC demonstration ... 4904
Query an unsigned zone without DNSSEC validation required ... 4904
Sign a zone on DC1 and distribute trust anchors ... 4905
Query a signed zone without DNSSEC validation required ... 4910
Query a signed zone with DNSSEC validation required ... 4911
Query DNSSEC records in the sec.contoso.com zone ... 4912
Unsign the zone and then re-sign the zone with custom parameters ... 4913
Demonstrate failed validation ... 4916
Demonstrate Remote Desktop failure ... 4917
Demonstrate Active Directory replication of DNSSEC signed resource records ... 4918
Transfer the Key Master role for sec.contoso.com to DC2 ... 4919
Appendix: Network Monitor results ... 4920
Packet capture results ... 4920

Dynamic Host Configuration Protocol (DHCP) Overview ... 4921
Role description ... 4921
Practical applications ... 4922
New and changed functionality ... 4922
Server Manager information ... 4922
See Also ... 4922

What's New in DHCP in Windows Server 2012 R2 ... 4923
Role/Feature description ... 4923
New and changed functionality ... 4923
DNS registration enhancements ... 4923
DNS PTR registration options ... 4924
Windows PowerShell for DHCP server ... 4924

What's New in DHCP in Windows Server 2012 4927
Role/Feature description 4927
New and changed functionality 4927
See Also 4927

Understand and Deploy DHCP Failover 4928
DHCP high availability options 4928
In this guide 4929
See Also 4930

What is DHCP Failover? 4930
Introduction to DHCP failover 4931
DHCP failover specifications 4932
DHCP failover and IPv6 4933
DHCP failover and Windows Failover Clustering 4933
DHCP failover and DNS dynamic updates 4934
See also 4935

DHCP Failover Relationships 4935
Understand failover relationships 4935
Create failover relationships 4937
View failover relationships 4938
Edit failover relationships 4940
Delete failover relationships 4941
See also 4942

DHCP Failover Modes 4943
Hot standby mode 4943
Load balance mode 4944
See also 4945

DHCP Failover Communications 4945
DHCP failover messages 4946
See also 4949

DHCP Failover Settings 4949
Name 4951
Partner server 4951
Mode 4951
Load balance percentage 4951
Server role 4951
Reserve percentage 4951
Maximum client lead time (MCLT) 4952
State switchover interval 4952
State 4952

Scope IDs ... 4955
Automatic state transition ... 4955
Enable authentication ... 4955
See also ... 4955

DHCP Failover Examples ... 4956
Hot standby example ... 4956
Load balancing example ... 4958
Partner down example ... 4960
See also ... 4962

DHCP Failover Events and Performance ... 4962
DHCP server event channels ... 4962
Administrative event logging ... 4963
Operational event logging ... 4965
Performance counters ... 4968
See also ... 4968

Deploy DHCP Failover ... 4969
Overview: Configure DHCP failover using the DHCP console ... 4969
Overview: Configure DHCP failover using Windows PowerShell ... 4970
See also ... 4970

DHCP Failover Requirements ... 4970
Prerequisite checks ... 4973
See also ... 4974

DHCP Failover Architecture ... 4974
DHCP failover design ... 4974
Hot standby design ... 4974
Load balance design ... 4975
Load balance in a single site with a single subnet ... 4975
Load balance in a single site with multiple subnets ... 4976
Deployment considerations ... 4977
Time synchronization ... 4977
BOOTP support ... 4977
Policy based assignment ... 4977
Windows Firewall ... 4978
Relay agents ... 4978
Duplicate relay agents ... 4978
See also ... 4979

Checklist: Deploy DHCP Failover ... 4979
See also ... 4980

Configure DHCP Failover using the Command Line ... 4980
Configure DHCP failover in load balance mode using the command line ... 4981
Configure DHCP failover in hot standby mode using the command line ... 4982
See also ... 4983

Configure DHCP Failover using the DHCP Console ... 4983
Configure DHCP failover in load balance mode using the DHCP console ... 4984
Configure DHCP failover in hot standby mode using the DHCP console ... 4987
See also ... 4988

Replicate DHCP Failover Using the Command Line ... 4989
Replicate failover settings at the server level using the command line ... 4989
Replicate failover settings at the relationship level using the command line ... 4990
Replicate failover settings at the scope level using the command line ... 4991
See also ... 4991

Replicate DHCP Failover Using the DHCP Console ... 4991
Replicate failover settings at the server level using the DHCP console ... 4992
Replicate failover settings at the relationship level using the DHCP console ... 4992
Replicate failover settings at the scope level using the DHCP console ... 4993
See also ... 4993

Migrate to DHCP Failover ... 4993
Migrate to DHCP failover ... 4994
Install the DHCP Server role on DHCP2012-1 and DHCP2012-2 ... 4994
Export DHCP settings and leases from DHCP2008-1 and DHCP2008-2 ... 4995
Import DHCP settings and leases to DHCP2012-1 ... 4998
Import server level settings to DHCP2012-2 ... 4999
Configure DHCP failover on DHCP2012-1 ... 4999
Update DHCP relay agents and complete migration ... 4999
See also ... 4999

DHCP Policy Based Assignment (PBA) Scenario Guide ... 4999
In this guide ... 5000
See also ... 5000

Introduction to DHCP Policies ... 5000
Why DHCP PBA? ... 5000
How DHCP PBA works ... 5002
DHCP policy conditions and settings ... 5005
Conditions ... 5005
Settings ... 5006
Policy processing ... 5007
Deploying DHCP policies ... 5008
See also ... 5009

Scenario: Manage the network configuration of virtual machines 5009
Problem description 5009
Virtual machines 5010
MAC addressing in Hyper-V 5011
Configure a MAC address based DHCP policy 5012
Conclusion 5013
See also 5013

Scenario: Secure a subnet to a specific set of clients 5013
Problem description 5013
Creating a scope-based allow-filter list using DHCP policies 5014
Creating a scope-based deny-filter list using DHCP policies 5016
Windows PowerShell 5017
See also 5018

Scenario: Customize lease duration based on device type 5018
Problem description 5018
Grouping devices on the network 5019
Configure DHCP policies for different lease durations 5019
Use Windows PowerShell to configure policies with different lease duration 5024
See also 5024

Step-by-Step: Configure DHCP Using Policy-based Assignment 5025
In this guide 5025
DHCP policy based assignment overview 5025
Policy settings and evaluation 5026
Address assignment 5026
Option assignment 5027
Scenario overview 5027
Hardware and software requirements 5027
Steps for configuring the test lab 5028
Configure DHCP1 5028
Install the operating system and configure TCP/IP on DHCP1 5028
Install AD DS, DNS Server, and DHCP Server 5029
Create a domain administrator account 5030
Create a DHCP scope on DHCP1 5031
Configure Client1 5032
Install the operating system on Client1 5032
Pin Windows PowerShell to the taskbar 5032
Configure Client2 5033
DHCP policy based assignment demonstration 5033
Determine MAC addresses 5033
Create policies 5034
Demonstrate policies 5036

Step-by-Step: Configure DHCP for Failover ... 5038
In this guide ... 5038
DHCP failover overview ... 5038
DHCP failover architecture ... 5039
Hot standby mode ... 5039
Load sharing mode ... 5040
Load sharing in a single site with a single subnet ... 5040
Load sharing in a single site with multiple subnets ... 5041
Scenario overview ... 5041
Hardware and software requirements ... 5042
Steps for configuring the test lab ... 5042
Configure DHCP1 ... 5043
Install the operating system and configure TCP/IP on DHCP1 ... 5043
Install AD DS, DNS Server, and DHCP Server ... 5044
Create a domain administrator account ... 5045
Authorize DHCP1 in Active Directory ... 5046
Configure DHCP2 ... 5046
Install the operating system and configure TCP/IP on DHCP2 ... 5046
Install and configure DHCP on DHCP2 ... 5047
Configure Client1 ... 5049
Install the operating system on Client1 ... 5049
Pin Windows PowerShell to the taskbar ... 5049
DHCP failover demonstration ... 5050
Configure a failover relationship ... 5050
View or edit properties of the failover configuration ... 5052
Edit properties of the failover relationship and demonstrate load balancing ... 5054
Edit properties of the failover relationship and demonstrate hot standby mode ... 5055
Conclusion ... 5056
See also ... 5056

Extensible Authentication Protocol (EAP) for Network Access Overview ... 5057

Extensible Authentication Protocol (EAP) Settings for Network Access ... 5057
Authentication methods ... 5057
EAP-TLS, PEAP, and EAP-TTLS ... 5058
Protected EAP Properties configuration items ... 5059
Secure password (EAP-MSCHAP v2) Properties configuration items ... 5065
Smart Card or other Certificate Properties configuration items ... 5065
Configure New Certificate Selection configuration items ... 5069
Select EKUs ... 5071
Add or Edit EKU ... 5072
TTLS configuration items ... 5072
EAP-SIM, EAP-AKA, and EAP-AKA' ... 5076
EAP-SIM ... 5076

EAP-AKA ... 5077
EAP-AKA' ... 5078
Additional resources ... 5079

High-Speed Networking ... 5079

Hyper-V Network Virtualization Overview ... 5079
Feature description ... 5080
Practical applications ... 5081
Important functionality ... 5082
Software requirements ... 5084
See also ... 5084

What's New in Hyper-V Network Virtualization in Windows Server 2012 R2 ... 5085
Feature description ... 5085
New and changed functionality ... 5086
Inbox HNV Gateway ... 5086
HNV Architecture ... 5086
HNV interoperability with Hyper-V Virtual Switch Extensions ... 5087
HNV VM Network Diagnostics ... 5087
Dynamic IP Address Learning ... 5087
HNV + Windows NIC Teaming ... 5088
NVGRE Encapsulated Task Offload ... 5088
See also ... 5089

Hyper-V Network Virtualization technical details ... 5090
Hyper-V Network Virtualization Concepts ... 5090
Routing in Hyper-V Network Virtualization ... 5093
Routing Between Virtual Subnets ... 5093
Routing Outside a Virtual Network ... 5093
Private Cloud (Routing) ... 5093
Hybrid Cloud (Site to site VPN) ... 5094
Packet Encapsulation ... 5096
Network virtualization through address virtualization ... 5097
Multitenant deployment example ... 5098
Hyper-V Network Virtualization architecture ... 5100
Hyper-V Network Virtualization Policy Management ... 5102
Summary ... 5102
See also ... 5102

Hyper-V Network Virtualization Gateway Architectural Guide ... 5103
Gateway Scenarios ... 5103
Private Cloud (routing) ... 5104
Hybrid Cloud (S2S VPN) ... 5104
Load Balancer ... 5105

External Router Gateway Configuration 5109
Management 5109
Gateway Console 5109
Multiple Gateways 5110
System Center Virtual Machine Manager 5110
Windows-Based Gateway Appliances 5111
Private Cloud Router Architecture 5111
Cross Premise Gateway 5112
Hardware Considerations 5114
Additional Resources 5115

Hyper-V Virtual Switch Overview 5116
Hyper-V Virtual Switch 5116
Practical applications 5117
Important functionality 5118
Hyper-V virtual switch 5118
Hardware requirements 5119
See also 5119

What's New in Hyper-V Virtual Switch in Windows Server 2012 R2 5120
Hyper-V Virtual Switch Extended Port ACLs 5120
Dynamic Load Balancing of Network Traffic 5121
Hyper-V Network Virtualization coexists with third party forwarding extensions for the Hyper-V Virtual Switch 5121
Traffic bottlenecks to VMs are reduced with vRSS 5122
Network tracing is streamlined and provides more detail 5122

What's New in Hyper-V Virtual Switch in Windows Server 2012 5122
Windows PowerShell commands for Hyper-V Virtual Switch 5123
Multiple Virtual NICs 5124
New security features for Hyper-V Virtual Switch 5126
Port Access Control Lists (ACLs) 5127
MacAddressSpoofing 5128
RouterGuard 5128
DHCPGuard 5129
IPsec Task Offload (IpsecTO) 5129
Port Virtual Local Area Network (PVLAN) and Trunk Mode 5130
Port Mirroring 5131
Receive Side Scaling (RSS) and Dynamic Virtual Machine Queue (dVMQ) 5131

Create Security Policies with Extended Port Access Control Lists for Windows Server 2012 R2 5132
Detailed ACL rules 5132
Configuring ACL rules with Windows PowerShell 5132

Detailed ACL rule examples ... 5134
Enforce application-level security ... 5134
Enforce both user-level and application-level security ... 5135
Provide security support to a non-TCP/UDP application ... 5136
Stateful ACL rules ... 5136
Allow inbound remote server traffic only after it is contacted by the local server ... 5137
See Also ... 5138

Unified Tracing Overview ... 5138
See also ... 5139

IP Address Management (IPAM) Overview ... 5139
Feature description ... 5140
IPAM architecture ... 5140
IPAM specifications ... 5142
Practical applications ... 5143
New and changed functionality ... 5143
Server Manager information ... 5143
See also ... 5144

What's New in IPAM in Windows Server 2012 R2 ... 5144
Role/Feature description ... 5144
New and changed functionality ... 5145
Role based access control ... 5146
Virtual address space management ... 5147
Enhanced DHCP server management ... 5147
External database support ... 5147
Upgrade and migration support ... 5148
Enhanced Windows PowerShell support ... 5148
See also ... 5148

Walkthrough: Demonstrate IPAM in Windows Server 2012 R2 ... 5148
Prerequisites ... 5148
Lab setup ... 5149
Objectives ... 5151
Objective 1: Demonstrate role based access control and delegated administration ... 5151
Concepts ... 5151
Procedures ... 5151
Objective 2: Manage DHCP policy based assignment with IPAM ... 5153
Concepts ... 5153
Procedures ... 5154
Objective 3: Automate IP address lifecycle management ... 5155
Concepts ... 5155
Procedures ... 5155

Objective 4: Administer DHCP failover with IPAM ... 5157
Concepts ... 5157
Procedures ... 5157
Objective 5: Manage DHCP MAC address filters with IPAM ... 5160
Concepts ... 5160
Procedures ... 5160
Objective 6: Manage DHCP superscopes with IPAM ... 5161
Concepts ... 5161
Procedures ... 5161

What's New in IPAM in Windows Server 2012 ... 5162
Role/Feature description ... 5162
New and changed functionality ... 5163
See also ... 5163

Walkthrough: Demonstrate IPAM in Windows Server 2012 ... 5163
In this guide ... 5163
IPAM overview ... 5164
IPAM discovery ... 5164
IP address space management ... 5164
Multi-server management and monitoring ... 5165
Operational auditing and IP address tracking ... 5165
IPAM architecture ... 5165
IPAM security groups ... 5166
IPAM tasks ... 5166
Privacy ... 5167
Information collected, processed, or transmitted ... 5168
Audit control ... 5168
IPAM requirements ... 5168
Scenario overview ... 5169
Hardware and software requirements ... 5169
Configuring the test lab ... 5169
Configure DC1 ... 5170
Install the operating system and configure TCP/IP on DC1 ... 5170
Install Active Directory and DNS on DC1 ... 5171
Create a domain administrator account ... 5172
Configure DHCP1 ... 5173
Install the operating system and configure TCP/IP on DHCP1 ... 5173
Install and configure DHCP on DHCP1 ... 5174
Configure Client1 ... 5175
Install the operating system and configure TCP/IP on Client1 ... 5175
Join Client1 to the contoso.com domain ... 5176
Configure IPAM1 ... 5176
Install the operating system and configure TCP/IP on IPAM1 ... 5176

Install and configure IPAM on IPAM1 ... 5177
IPAM demonstration ... 5182
Address space management ... 5182
Create, delete, import and export IP addresses ... 5183
Find available IP addresses and create reservations ... 5189
Create custom logical groups ... 5194
Infrastructure monitoring and management ... 5196
Review audit logs and events ... 5200
See also ... 5200

IPAM Planning and Design Guide ... 5201
In this guide ... 5201
Benefits of IPAM ... 5201
Server administration and management ... 5202
Planning and organization of address space ... 5202
Flexibility and automation ... 5202
Forensics ... 5203
See also ... 5203

What is IPAM? ... 5203
Address Space Management ... 5204
Multi-Server Management and Monitoring ... 5205
Network Audit ... 5206
See also ... 5206

IPAM Terminology ... 5207
See also ... 5210

Getting Started with IPAM ... 5210
The IPAM Client ... 5212
Provisioning IPAM ... 5213
Choosing a provisioning method ... 5213
Configuring managed server access settings ... 5214
Server discovery ... 5216
Scope of discovery ... 5216
Discovering servers ... 5217
Domain controllers ... 5218
DHCP servers ... 5218
DNS servers ... 5218
NPS servers ... 5219
Select managed servers ... 5219
Verify managed server access ... 5221
Verifying managed DHCP servers ... 5222
Verifying managed DNS servers ... 5222

Retrieve data ... 5223
See also ... 5223

IPAM Architecture ... 5223
IPAM Client ... 5224
Client communications ... 5225
IPAM Server ... 5225
Server communications ... 5225
Scheduled tasks ... 5226
The IPAM database ... 5227
Role-based access control ... 5227
Managed servers ... 5229
See also ... 5230

IPAM Deployment Planning ... 5230
In this topic ... 5230
Deployment topologies ... 5231
Hardware and software requirements ... 5231
IPAM specifications ... 5233
Capacity planning ... 5234
Planning disk capacity ... 5234
See also ... 5236

IPAM Deployment Guide ... 5237
About this guide ... 5237
In this guide ... 5237
Terminology used in this guide ... 5237
See Also ... 5237

Planning to Deploy IPAM ... 5238
Reviewing IPAM concepts ... 5238

Implementing Your IPAM Design Plan ... 5239
See Also ... 5240

Checklist: Deploy IPAM ... 5240

Deploying IPAM Server ... 5241
Configuring an IPAM server ... 5241
See Also ... 5242

Checklist: Deploy IPAM Server ... 5242
See Also ... 5243

Install IPAM Server ... 5243
Installing IPAM Server ... 5244

See Also 5244

Choose an IPAM Provisioning Method 5244
Choosing a provisioning method 5245
See Also 5246

Configure Server Discovery 5246
Configuring server discovery 5247
See Also 5248

Discover Servers on the Network 5249
Discovering servers on the network 5249
See Also 5250

Manually Add a Server to Server Inventory 5250
Manually adding a server 5250
See Also 5251

Create IPAM Provisioning GPOs 5251
Creating IPAM provisioning GPOs 5251
Sample Usage 5254
Configuring IPAM GPO security filtering 5254
See Also 5255

Manually Configuring Managed Server Access Settings 5255
Manually configuring managed server access settings 5255
See Also 5255

Manually Configure DHCP Access Settings 5256
Manually configure managed DHCP server access settings 5256
Configure Windows Firewall on a managed DHCP server 5257
Configure security groups on a managed DHCP server 5258
Configure a DHCP audit share on a managed DHCP server 5260
Restart the DHCP Server service 5261
See Also 5261

Manually Configure DNS Access Settings 5262
Manually configure managed DNS server access settings 5262
Configure Windows Firewall on a managed DNS server 5263
Configure security groups on a managed DNS server 5264
Enable event log monitoring on a managed DNS server 5266
Configure the DNS DACL on a managed DNS server 5267
See Also 5268

Manually Configure DC and NPS Access Settings 5268
Manually configure managed DC and NPS server access settings 5269

Configure Windows Firewall on a managed DC or NPS server 5269
Configure security groups on a managed DC or NPS server 5270
See Also 5272

Choose Managed Servers 5272
Choosing managed servers 5272
See Also 5273

Verify Managed Server Access 5273
Verifying managed server access 5274
See Also 5275

Retrieve Data from Managed Servers 5275
Retrieving data from managed servers 5275

Deploying IPAM Client 5276
Installing IPAM Client 5276
See Also 5276

Checklist: Deploy IPAM Client 5276
See Also 5277

Install RSAT 5278
Installing RSAT 5278
See Also 5279

Install IPAM Client 5279
The server pool and IPAM Client 5280
See Also 5282

Connect to an IPAM Server 5282
Connecting to an IPAM server 5282
See Also 5283

Assigning IPAM Server and Administrator Roles 5284
Configuring IPAM server roles 5284
Configuring IPAM security groups 5284
See Also 5284

Checklist: Assign Roles 5285
See Also 5285

Configure IPAM Server Roles 5285
Configuring IPAM server roles 5286
See Also 5289

Assign Administrator Roles 5289

Assigning IPAM administrator roles ... 5289
See Also ... 5292

IPAM Operations Guide ... 5292
In this guide ... 5292
See also ... 5293

Using the IPAM Client Console ... 5293
Upper navigation pane ... 5294
OVERVIEW ... 5294
SERVER INVENTORY ... 5295
IP ADDRESS SPACE ... 5296
Context-sensitive menus ... 5297
IP Address Blocks ... 5298
IP Address Inventory ... 5299
IP Address Range Groups ... 5301
MONITOR AND MANAGE ... 5303
DNS and DHCP Servers ... 5303
DHCP Scopes ... 5305
DNS Zone Monitoring ... 5306
Server Groups ... 5307
EVENT CATALOG ... 5309
IPAM Configuration Events ... 5310
DHCP Configuration Events ... 5310
IP Address Tracking ... 5311
By IP Address ... 5311
By Client ID ... 5312
By Host Name ... 5312
By User Name ... 5313
Lower navigation pane ... 5313
Lower navigation list view ... 5314
Lower navigation tree view ... 5314
Display pane search view ... 5315
Display pane details view ... 5315
Display pane refresh ... 5315
Notifications area ... 5315
Server Manager menu bar ... 5316
TASKS menu ... 5317
See also ... 5317

Managing Server Inventory ... 5318
Server inventory ... 5318
Adding servers to the inventory ... 5319
Manageability status ... 5319

IPAM access status ... 5320
See also ... 5320

Managing IP Address Space ... 5321
IP address space mapping ... 5321
IP address blocks ... 5322
IP address ranges ... 5322
IP addresses ... 5322
IP address operations ... 5323
Add or edit IP address space ... 5323
IPv4 blocks ... 5323
IPv6 blocks ... 5324
IPv4 ranges ... 5325
IPv6 ranges ... 5325
IPv4 addresses ... 5326
IPv6 addresses ... 5326
Import ... 5326
Import and update ... 5327
Export ... 5328
Find and allocate available IP address ... 5328
Reclaim IP addresses ... 5330
IP address expiry log settings ... 5332
IP address inventory ... 5333
Adding IP addresses to the inventory ... 5333
Basic configurations ... 5334
Managing DHCP reservations ... 5334
Managing DNS records ... 5336
Custom configurations ... 5337
Importing DHCP lease data from the network ... 5339
IP address range groups ... 5340
See also ... 5341

Multi-server Management ... 5342
DNS and DHCP server monitoring ... 5342
DNS zone monitoring ... 5343
DNS and DHCP record synchronization ... 5344
DHCP server management ... 5344
DHCP scope management ... 5345
See also ... 5347

IP Address Tracking ... 5347
Tracking by IP address ... 5347
Tracking by client ID ... 5348
Tracking by host name ... 5349

Tracking by user name ... 5349
See also ... 5350

Operational Event Tracking ... 5350
DHCP configuration events ... 5350
IPAM configuration events ... 5352
See also ... 5352

Using Windows PowerShell with IPAM ... 5352
IpamGpoProvisioning ... 5352
IpamCustomValue ... 5352
IpamAddress ... 5353
IpamCustomField ... 5353
IpamRange ... 5354
IpamConfiguration ... 5354
See also ... 5354

Best Practices ... 5355
IPAM best practices ... 5355
General ... 5355
Provisioning ... 5355
Discovery ... 5356
IP address management ... 5356
Monitor and manage ... 5357
Event catalog ... 5357
See also ... 5358

Troubleshooting IPAM ... 5358
Event logs ... 5358
Common problems and solutions ... 5358
Connecting to the IPAM server ... 5359
Server access ... 5359
Server discovery ... 5359
IP address space management ... 5360
Monitoring and management ... 5360
See also ... 5360

IPAM Backup and Restore ... 5361
Requirements for using Windows Server Backup ... 5361
Backup and restore IPAM ... 5361
See also ... 5364

Low Latency Workloads Technologies ... 5365

Data Center Bridging ... 5366

Data Center Bridging 5366

Data Center Transmission Control Protocol (DCTCP) 5366
Data Center Transmission Control Protocol (DCTCP) 5366

Kernel Mode Remote Direct Memory Access (kRDMA) 5367
Kernel Mode Remote Direct Memory Access (kRDMA) 5367

Network Interface Card (NIC) Teaming 5368
NIC teaming 5368
NIC teaming in virtual machines 5369

NetworkDirect 5370
NetworkDirect 5370

Receive Segment Coalescing (RSC) 5370
Receive Segment Coalescing (RSC) 5370

Receive Side Scaling (RSS) 5372
Receive Side Scaling (RSS) 5372

Registered Input/Output (RIO) API Extensions 5374
Registered Input/Output (RIO) API Extensions 5374

Transmission Control Protocol (TCP) Loopback Optimization 5376
Transmission Control Protocol (TCP) loopback optimization 5376

Low Latency Workloads Management and Operations 5377
New Performance Counters 5377
UDP Datagram Loss 5377
Rejected Connections 5377
NetworkDirect Activity 5377
Per Processor Network interface card Activity 5377
NUMA systems synchronization 5377
Windows PowerShell management of networking features 5378

Network Load Balancing Overview 5378
Feature description 5378
Practical applications 5379
High availability 5379
Scalability 5379
Manageability 5379
Important functionality 5380
Hardware requirements 5382
Software requirements 5382
Server Manager information 5382

See also 5382

Network Policy and Access Services Overview 5383
Role description 5383
New and changed functionality 5384
Support for Windows PowerShell 5385
Removed functionality 5385
Deprecated functionality 5385
Server Manager information 5386
Running Network Policy and Access Services 5386
How do I deploy and configure Network Policy and Access Services using Windows PowerShell? 5387
How do I deploy and configure this role in a multi-server environment? 5387
Can I run this role on virtual machines? 5387
Can I run this role in a clustered environment? 5387
Special considerations for managing this role remotely 5387
Special considerations for managing the role on the Server Core installation option 5387
See also 5387

Network Shell (Netsh) Overview 5388
Network Shell (Netsh) Technical Reference 5389
Content availability 5389

Formatting Legend 5389
Formatting legend 5389

Netsh Commands for Trace 5390
Learn More 5390

Netsh Commands for Network Trace 5390
In this topic 5390
Using Netsh trace commands 5391
Identifying scenarios 5391
Obtaining trace provider details 5393
Using filters to limit ETL trace file details 5394
Example Netsh trace filter parameters and usage 5396
Starting and stopping trace 5396
Using the files rendered by trace 5396
Commands in the Netsh trace context 5397
convert 5397
correlate 5398
diagnose 5399
show CaptureFilterHelp 5401
show globalKeywordsAndLevels 5401
show helperclass 5403

show interfaces 5403
show provider 5403
show providers 5404
show scenario 5404
show scenarios 5404
show status 5405
start 5405
stop 5408

Network Subsystem Performance Tuning 5408

Choosing a Network Adapter 5409
Offload Capabilities 5410
Receive Side Scaling (RSS) 5411
Understanding RSS Performance 5413
RSS and Virtualization 5414
NIC Teaming and RSS 5414
Receive Segment Coalescing (RSC) 5414
Understanding RSC Diagnostics 5414
RSC and Virtualization 5415
Network Adapter Resources 5415
Message-Signaled Interrupts (MSI/MSI-X) 5416
Interrupt Moderation 5416
Suggested Network Adapter Features for Server Workloads 5416

Performance Tuning Network Adapters 5417
Enabling Offload Features 5417
Enabling Receive Side Scaling (RSS) for Web Servers 5418
RSS Profiles and RSS Queues 5418
Increasing Network Adapter Resources 5418
Enabling Interrupt Moderation 5419
Performance Tuning for Low Latency Packet Processing 5419
System Management Interrupts 5420
Performance Tuning TCP 5420
TCP Receive Window Auto-Tuning 5420
Windows Filtering Platform 5421
TCP Parameters 5421

Network-Related Performance Counters 5421
Resource Utilization 5421
Potential Network Problems 5422
Receive Side Coalescing (RSC) performance 5423

Performance Tools for Network Workloads 5423
Tuning for NTttcp 5423

TCP/IP Window Size 5424
Microsoft Server Performance Advisor 3.0 5424

NIC Teaming Overview 5425
Requirements 5425
Technical overview 5425
NIC Teaming architecture 5425
NIC Teaming configurations 5426
Traffic distribution algorithms 5427
NIC Teaming in virtual machines 5428
Incompatibilities 5428
Requirements 5428
Configuring NIC Teaming 5429
NIC Teaming PowerShell commands 5430

Quality of Service (QoS) Overview 5430
QoS features 5430
Bandwidth management 5431
Classification and tagging 5432
Priority based flow control 5433
Policy-based QoS and Hyper-V QoS 5433
Policy-based QoS 5433
Hyper-V QoS 5434

QoS Common Configurations 5435
2 NICs without NIC Teaming 5435
2 NICs with NIC Teaming 5437
4 NICs in two NIC teams 5440
4 NICs with a standard NIC team and two RDMA NICs 5441
Alternate configuration of 4 NICs with a standard NIC team and two RDMA NICs 5443

QoS Minimum Bandwidth Best Practices 5444
Guidelines for using Minimum Bandwidth 5444
Minimum Bandwidth Modes 5445
Minimum Bandwidth and Data Center Bridging (DCB) 5446
Configuring QoS and NIC Teaming 5447
Related topics 5450

Policy-based Quality of Service (QoS) 5450
Advantages of Policy-based QoS 5451
Scenario 1: Prioritizing a line-of-business application 5452
Prerequisites for prioritizing a line-of-business application 5454
Administrative credentials 5454
Setting up the test environment for prioritizing a line-of-business application 5454
Steps for prioritizing a line-of-business application 5454

Scenario 2: Prioritizing an HTTP server application ... 5454
Precedence rules for URL-based policies ... 5455
1. URL scheme ... 5455
2. URL host ... 5455
3. URL port ... 5456
4. URL path ... 5456
URL-based policy vs. Quintuple policy ... 5456
Configuring Policy-based QoS ... 5456
Create a QoS policy ... 5457
Prioritizing traffic with DSCP ... 5457
Throttling traffic ... 5457
Wizard page 1—Policy Profile ... 5458
Wizard page 2—Application Name ... 5458
Wizard page 3—IP Addresses ... 5459
Wizard page 4—Protocols and Ports ... 5460
View, edit, or delete a QoS policy ... 5461
QoS Policy GPMC Reporting ... 5461
Advanced settings for roaming and remote users ... 5462
Advanced QoS settings ... 5462
Advanced QoS settings: inbound TCP traffic ... 5463
Advanced QoS settings: DSCP Marking Override ... 5463
Wireless Multimedia and DSCP values ... 5464
QoS Policy Precedence Rules ... 5464
Reference ... 5466
Error and Event Messages ... 5466
Informational Messages ... 5466
Warning Messages ... 5469
Error Messages ... 5472
Additional resources ... 5475
FAQ ... 5476
What operating system does my domain controller need to be running? ... 5476
Do I need Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to deploy QoS policies? ... 5476
Do QoS policies apply to the sender or receiver of traffic? ... 5476
What happens if conflicting policies are deployed to the same computer? ... 5476
Additional resources ... 5477

Configuring Policy-based Quality of Service (QoS) ... 5477
QoS policies overview ... 5477
Prioritizing traffic by using DSCP ... 5478
Throttling traffic ... 5478
Create a QoS policy ... 5478
Policy Profile - Wizard page 1 ... 5478

Application Name - Wizard page 2 5479
IP Addresses - Wizard page 3 5479
Protocols and Ports - Wizard page 4 5480
View, edit, or delete a QoS policy 5480
Additional resources 5481

Quality Windows Audio Video Experience (qWave) 5481
qWave Overview 5482
See Also 5482

QoS Frequently Asked Questions 5482
Frequently Asked Questions 5482

Remote Access (DirectAccess, Routing and Remote Access) Overview 5483
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set 5484
Deploy Basic Remote Access 5484
Deploy Advanced Remote Access 5484
Deploy Remote Access in an Enterprise 5484
Manage Remote Access 5485
Migrate Remote Access 5485
Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites 5485
Overview of the Remote Access Technologies 5485
DirectAccess and RRAS Unified Server Role 5485
DirectAccess in Windows Server 2012 Essentials 5486
What works differently? 5486
DirectAccess and RRAS coexistence 5487
Simplified DirectAccess Deployment 5487
Removal of PKI Deployment as a DirectAccess Prerequisite 5487
NAT64 and DNS64 Support for Accessing IPv4-only Resources 5488
Support for DirectAccess Server behind a NAT Device 5489
Simplified Network Security Policy 5490
Load Balancing Support 5490
Support for Multiple Domains 5491
Support for OTP (Token Based Authentication) 5491
Automated Support for Force Tunneling 5491
IP-HTTPS Interoperability and Performance Improvements 5492
DirectAccess Manage-out to Clients Support 5493
Multisite Support 5493
Support for Server Core 5494
PowerShell Support 5494
User Monitoring 5495
Server Operational Status 5497
Diagnostics 5498

Accounting and Reporting ... 5498
Site-to-site IKEv2 IPsec tunnel mode VPN ... 5499
See also ... 5499

What's New in Remote Access in Windows Server 2012 R2 ... 5500
New Features in Windows Server 2012 R2 ... 5500
Multi-tenant Site-to-site VPN Gateway ... 5500
Multi-tenant Remote Access VPN Gateway ... 5501
Border Gateway Protocol (BGP) ... 5501
Web Application Proxy ... 5501
Auto-triggered VPN ... 5501
SSL VPN plugin from 3rd Party VPN vendors ... 5501
Advanced VPN Client PowerShell configuration ... 5502
Enhanced VPN IPsec ... 5502
Create and Edit VPN profiles in PC settings ... 5502

DirectAccess Offline Domain Join ... 5503
Offline domain join overview ... 5503
Offline domain join with DirectAccess policies scenario overview ... 5503
Prepare for offline domain join ... 5503
Operating system requirements ... 5504
Credential requirements ... 5504
Granting user rights to join workstations to the domain ... 5504
Offline domain join process ... 5505
Steps for performing a DirectAccess offline domain join ... 5505
See Also ... 5506

Remote Access (DirectAccess) Known Issues ... 5507
Recommended hotfixes and updates for Windows Server 2012 DirectAccess ... 5507

Remote Access (DirectAccess) Prerequisites ... 5507

Remote Access (DirectAccess) Unsupported Configurations ... 5511
Network Access Protection for DirectAccess clients ... 5511
Multisite support for Windows 7 clients ... 5511
User-based access control ... 5511
Customizing DirectAccess policy ... 5512
KerbProxy authentication ... 5512
Using ISATAP ... 5512

Remote Access (DirectAccess) Troubleshooting ... 5513

DirectAccess Capacity Planning ... 5515
High-end hardware test environment ... 5516
Test Environment ... 5517

Testing results for low-end hardware: 5518
Testing results for high-end hardware: 5520

Deploy a Single Remote Access Server using the Getting Started Wizard 5524
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set 5525
Deploy Basic Remote Access 5525
Deploy Advanced Remote Access 5525
Deploy Remote Access in an Enterprise 5525
Manage Remote Access 5525
Migrate Remote Access 5525
Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites 5526
Scenario description 5526
In this scenario 5526
Prerequisites 5526
Planning steps 5527
Deployment steps 5527
Practical applications 5527
Roles and features included in this scenario 5528
Hardware requirements 5529
Software requirements 5530
See also 5530

Plan a Basic Remote Access Deployment 5531
Next steps 5531

Step 1: Plan the Remote Access Infrastructure 5531
Plan network topology and settings 5532
Plan network adapters and IP addressing 5532
Plan firewall requirements 5537
Plan certificate requirements 5537
Plan certificates for IP-HTTPS and network location server 5538
Plan DNS requirements 5538
DNS server requirements 5540
Plan the network location server 5540
Plan Active Directory 5540
Plan Group Policy Objects 5541
Automatically-created GPOs 5542
Manually-created GPOs 5542
Recovering from a deleted GPO 5543
See also 5543

Step 2: Plan the Remote Access Deployment 5543
Planning for client deployment 5543

Planning for Remote Access server deployment 5544
See also 5545

Install and Configure Basic Remote Access 5545

Step 1: Configure the Remote Access Infrastructure 5545
Configure server network settings 5546
Configure routing in the corporate network 5547
Configure firewalls 5547
Configure the DNS server 5548
Configure Active Directory 5549
Configure GPOs 5550
Configure security groups 5551

Step 2: Configure the Remote Access Server 5551
Install the Remote Access role 5552
Configure DirectAccess with the Getting Started Wizard 5552
Update clients with the DirectAccess configuration 5553
See also 5554

Step 3: Verify the Deployment 5554

Deploy a Single Remote Access Server with Advanced Settings 5554
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set 5555
Deploy Basic Remote Access 5555
Deploy Advanced Remote Access 5556
Deploy Remote Access in an Enterprise 5556
Manage Remote Access 5556
Migrate Remote Access 5556
Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites 5556
Scenario description 5557
In this scenario 5557
Prerequisites 5557
Planning steps 5557
Deployment steps 5558
Practical applications 5558
Roles and features required for this scenario 5558
Hardware requirements 5560
Software requirements 5561
See also 5562

Plan an Advanced Remote Access Deployment 5562
See also 5563

Step 1: Plan the Remote Access Infrastructure ... 5563
1.1 Plan network topology and settings ... 5564
1.1.1 Plan network adapters and IP addressing ... 5564
1.1.2 Plan IPv6 intranet connectivity ... 5568
1.1.3 Plan for force tunneling ... 5573
1.2 Plan firewall requirements ... 5574
1.3 Plan certificate requirements ... 5575
1.3.1 Plan computer certificates for IPsec authentication ... 5576
1.3.2 Plan certificates for IP-HTTPS ... 5576
1.3.3 Plan website certificates for the network location server ... 5578
1.4 Plan DNS requirements ... 5578
1.4.1 Plan for DNS server requirements ... 5581
1.4.2 Plan for local name resolution ... 5581
1.5 Plan the network location server ... 5583
1.5.1 Plan certificates for the network location server ... 5584
1.5.2 Plan DNS for the network location server ... 5584
1.6 Plan management servers ... 5584
1.7 Plan Active Directory ... 5585
1.7.1 Plan client authentication ... 5586
1.7.2 Plan multiple domains ... 5587
1.8 Plan Group Policy Objects ... 5587
1.8.1 Configure automatically created GPOs ... 5588
1.8.2 Configure manually created GPOs ... 5589
1.8.3 Manage GPOs in a multi-domain controller environment ... 5589
1.8.4 Manage Remote Access GPOs with limited permissions ... 5590
1.8.5 Recover from a deleted GPO ... 5591
See also ... 5592

Step 2: Plan the Remote Access Deployment ... 5592
2.1 Plan for client deployment ... 5593
2.2 Plan for Remote Access server deployment ... 5594
2.3 Plan infrastructure servers ... 5595
2.4 Plan application servers ... 5596
2.5 Plan DirectAccess and third-party VPN clients ... 5596
See also ... 5597

Install and Configure Advanced Remote Access ... 5597
See also ... 5597

Step 1: Configure the Remote Access Infrastructure ... 5597
1.1 Configure server network settings ... 5598
1.2 Configure force tunneling ... 5599
1.3 Configure routing in the corporate network ... 5600
1.4 Configure firewalls ... 5600

1.5 Configure CAs and certificates ... 5601
1.5.1 Configure IPsec authentication ... 5601
1.5.2 Configure certificate templates ... 5602
1.5.3 Configure the IP-HTTPS certificate ... 5602
1.6 Configure the DNS server ... 5603
1.7 Configure Active Directory ... 5605
1.8 Configure GPOs ... 5606
1.8.1 Configure Remote Access GPOs with limited permissions ... 5607
1.9 Configure security groups ... 5609
1.10 Configure the network location server ... 5610
See also ... 5612

Step 2: Configure the Remote Access Server ... 5612
2.1. Install the Remote Access role ... 5613
2.2. Configure the deployment type ... 5613
2.3. Configure DirectAccess clients ... 5614
2.4. Configure the Remote Access server ... 5615
2.5. Configure the infrastructure servers ... 5616
2.6. Configure application servers ... 5617
2.7. Configuration summary and alternate GPOs ... 5617
2.8. How to configure the Remote Access server by using Windows PowerShell ... 5618
See also ... 5618

Step 3: Verify the Deployment ... 5619

Deploy Remote Access in an Enterprise ... 5619
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set ... 5620
Deploy Basic Remote Access ... 5620
Deploy Advanced Remote Access ... 5620
Deploy Remote Access in an Enterprise ... 5621
Manage Remote Access ... 5621
Migrate Remote Access ... 5621
Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites ... 5621
Scenario description ... 5621
In this scenario ... 5622
Practical applications ... 5622
Roles and features included in this scenario ... 5622
See also ... 5624

Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB ... 5625
About this guide ... 5625
Known issues ... 5625

Overview of the Test Lab Scenario ... 5626

Configuration Requirements 5627
Steps for Configuring the Test Lab 5628
STEP 1: Complete the DirectAccess Configuration 5628
STEP 2: Configure EDGE1 5628
STEP 3: Install and Configure EDGE2 5629
STEP 4: Create the Network Load Balanced Remote Access Cluster 5632
Prerequisites 5632
Install the Network Load Balancing feature on EDGE1 and EDGE2 5633
Enable load balancing on EDGE1 5633
Add EDGE2 to the load balanced cluster 5634
STEP 5: Test DirectAccess Connectivity from the Internet and Through the Cluster 5635
STEP 6: Test DirectAccess Client Connectivity from Behind a NAT Device 5637
Prerequisites 5638
STEP 7: Test Connectivity When Returning to the Corpnet 5640
STEP 8: Snapshot the Configuration 5640
Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment 5641
About this guide 5641
Overview of the Test Lab Scenario 5641
Configuration Requirements 5643
Steps for Configuring the Test Lab 5644
STEP 1: Complete the DirectAccess Configuration 5645
STEP 2: Install and Configure ROUTER1 5645
Install the operating system on ROUTER1 5645
Configure TCP/IP properties and rename the computer 5645
Turn off the firewall 5646
Configure routing and forwarding 5647
STEP 3: Install and Configure CLIENT2 5648
User account control 5648
STEP 4: Configure APP1 5649
STEP 5: Configure DC1 5649
Create security groups for Windows 7 DirectAccess clients on DC1 5651

STEP 6: Install and Configure 2-DC1 5651
Install the operating system on 2-DC1 5652
Configure TCP/IP properties 5652
Configure 2-DC1 as a domain controller and DNS server 5653
Provide Group Policy permissions to CORP\User1 5654
Allow CORP2 computers to obtain computer certificates 5654
Force replication between DC1 and 2-DC1 5655

STEP 7: Install and Configure 2-APP1 5656
Install the operating system on 2-APP1 5656
Configure TCP/IP properties 5656
Join 2-APP1 to the CORP2 domain 5657
Install the Web Server (IIS) role on 2-APP1 5657
Create a shared folder on 2-APP1 5658

STEP 8: Configure INET1 5658

STEP 9: Configure EDGE1 5659

STEP 10: Install and Configure 2-EDGE1 5659

STEP 11: Configure the Multisite Deployment 5663

STEP 12: Test DirectAccess Connectivity 5665
Prerequisites 5665

STEP 13: Test DirectAccess Connectivity from Behind a NAT Device 5669
Prerequisites 5670

STEP 14: Snapshot the Configuration 5672

Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID 5672
About this guide 5672

Overview of the Test Lab Scenario 5673

Configuration Requirements 5674

Steps for Configuring the Test Lab 5675

STEP 1: Complete the DirectAccess Configuration 5675

STEP 2: Configure APP1 5675

STEP 3: Configure DC1 5678

STEP 4: Install and Configure RSA and EDGE1 5678
Install RSA SecurID software token on CLIENT1 5682
Configure EDGE1 as an RSA Authentication Agent 5683

Configure EDGE1 to support OTP authentication 5683
STEP 5: Verify OTP Health on EDGE1 5684
STEP 6: Test DirectAccess Connectivity from the Homenet Subnet 5685
STEP 7: Test DirectAccess Connectivity from the Internet 5686
STEP 8: Snapshot the Configuration 5687
Deploy Remote Access in a Cluster 5687
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set 5687
Deploy Basic Remote Access 5687
Deploy Advanced Remote Access 5687
Deploy Remote Access in an Enterprise 5688
Manage Remote Access 5688
Migrate Remote Access 5688
Scenario description 5688
Prerequisites 5688
In this scenario 5689
Practical applications 5689
Roles and features included in this scenario 5689
Hardware requirements 5691
Software requirements 5691
Known issues 5692
Plan a Remote Access Cluster Deployment 5693
Step 1: Plan an Advanced Single-Server Deployment 5693
Plan a single server deployment 5693
Step 2: Plan Cluster Servers 5693
2.1 Installing roles and features 5694
2.2 Configure server settings 5694
See also 5694
Step 3: Plan a Load-Balanced Cluster Deployment 5695
3.1 Plan load balancing 5695
3.1.1 External Load Balancer configuration example 5696
3.2 Plan IP-HTTPS 5698
3.3 Plan for VPN client connections 5698
3.4 Plan the network location server 5698
See also 5699
Configure a Remote Access Cluster 5699
Step 1: Implement a Single-Server Remote Access Deployment 5699

Implement a single server deployment ... 5699
See also ... 5699

Step 2: Prepare Cluster Servers ... 5700
2.1 Configure the Remote Access infrastructure ... 5700
2.2 Install the Remote Access role ... 5701
2.3 Install NLB ... 5701
See also ... 5701

Step 3: Configure a Load-Balanced Cluster ... 5701
3.1 Configure the IPv6 prefix ... 5703
3.2 Enable load balancing ... 5703
3.3 Install the IP-HTTPS certificate ... 5705
3.4 Install the network location server certificate ... 5706
3.5 Add servers to the cluster ... 5707
3.6 Remove a server from the cluster ... 5708
3.7 Disable load balancing ... 5709
3.8 Change from unicast to multicast ... 5710
See also ... 5710

Step 4: Verify the Cluster ... 5710

Deploy Multiple Remote Access Servers in a Multisite Deployment ... 5711
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set ... 5711
Deploy Basic Remote Access ... 5711
Deploy Advanced Remote Access ... 5711
Deploy Remote Access in an Enterprise ... 5711
Manage Remote Access ... 5712
Migrate Remote Access ... 5712
Scenario description ... 5712
Prerequisites ... 5712
In this scenario ... 5713
Practical applications ... 5713
Roles and features included in this scenario ... 5713
Hardware requirements ... 5715
Software requirements ... 5715
Known issues ... 5716

Plan a Multisite Deployment ... 5718

Step 1: Plan an Advanced Single Server Deployment ... 5718
Plan a single server deployment ... 5719
See also ... 5719

Step 2: Plan the Multisite Infrastructure ... 5719

2.1 Plan Active Directory ... 5719
Active Directory best practices and recommendations ... 5721
2.2 Plan security groups ... 5722
2.3 Plan Group Policy Objects ... 5723
2.3.1 Automatically-created GPOs ... 5723
2.3.2 Manually created GPOs ... 5724
2.3.3 Managing GPOs in a multi-domain controller environment ... 5724
Modifying domain controller association ... 5725
2.4 Plan DNS ... 5726
See also ... 5726

Step 3: Plan the Multisite Deployment ... 5727
3.1 Plan IP-HTTPS certificates ... 5727
3.2 Plan the network location server ... 5728
3.2.1 Certificate requirements for the network location server ... 5728
3.2.2DNS for the network location server ... 5729
3.3 Plan the IPsec root certificate for all Remote Access servers ... 5729
3.4 Plan global server load balancing ... 5729
3.5 Plan DirectAccess client entry point selection ... 5730
3.6 Plan prefixes and routing ... 5731
Internal IPv6 prefix ... 5731
IPv6 prefix for DirectAccess client computers (IP-HTTPS prefix) ... 5731
IPv6 prefix for VPN clients ... 5731
Routing ... 5731
Active Directory site-specific IPv6 prefixes ... 5733
3.7 Plan the transition to IPv6 when multisite Remote Access is deployed ... 5733
Transition from an IPv4-only to an IPv6+IPv4 corporate network ... 5734
Transition from an IPv6+IPv4 to an IPv6-only corporate network ... 5734
Transition from an IPv4-only to an IPv6-only deployment using dual DirectAccess deployments ... 5735
See also ... 5736

Configure a Multisite Deployment ... 5736

Step 1: Implement a Single Server Remote Access Deployment ... 5736
Implement a single server deployment ... 5737
See also ... 5737

Step 2: Configure the Multisite Infrastructure ... 5737
2.1. Configure additional Active Directory sites ... 5737
2.2. Configure additional domain controllers ... 5739
2.3. Configure security groups ... 5740
2.4. Configure GPOs ... 5741
Domain controller maintenance and downtime ... 5742

Change two or more domain controllers that manage server GPOs 5743
Optimization of configuration distribution 5746
See also 5747

Step 3: Configure the Multisite Deployment 5747
3.1. Configure Remote Access servers 5748
3.2. Grant administrator access 5749
3.3. Configure IP-HTTPS for a multisite deployment 5749
3.4. Configure the network location server for a multisite deployment 5750
3.5. Configure DirectAccess clients for a multisite deployment 5752
3.6. Enable the multisite deployment 5753
3.7. Add entry points to the multisite deployment 5755
See also 5757

Step 4: Verify the Multisite Deployment 5757

Troubleshoot a Multisite Deployment 5758

Troubleshooting Enabling Multisite 5759
User connectivity issues 5759
Kerberos proxy authentication 5759
IP-HTTPS certificates 5760
Network location server 5760
Windows 7 client computers 5761
Active Directory site 5761
Saving server GPO settings 5762
Internal error occurred 5762

Troubleshooting Adding Entry Points 5762
Internal error occurred 5763
Missing RemoteAccessServer parameter 5763
Remote Access is not configured 5763
Multisite not enabled 5763
IPv6 prefix issues 5764
ConnectTo address 5764
DirectAccess or VPN already installed 5765
IPsec root certificate 5765
Mixing IPv6 and IPv4 entry points 5766
Domain issues with the ServerGpoName 5767
Split-brain DNS 5768
Saving server GPO settings 5768
GPO updates cannot be applied 5768

Troubleshooting Setting the Entry Point Domain Controller 5769
Saving server GPO settings 5769

Remote Access is not configured ... 5769
Multisite is not enabled ... 5769
Entry point and domain controller not provided in cmdlet ... 5770
Could not locate domain controller ... 5770
Could not connect to domain controller ... 5771
Read-only domain controller ... 5772
Cannot retrieve GPO ... 5772
Entry point not part of multisite deployment ... 5773
Remote Access server settings ... 5773
Problem resolving FQDN ... 5774
No entry points to update ... 5774

Troubleshooting Web Probe URLs ... 5775
Adding and removing web probe URLs ... 5775
IP address of management server associated with web probe URL ... 5775
Cannot verify IP addresses ... 5776
DNS entry cannot be updated ... 5776

Troubleshooting General Issues ... 5776
GPO retrieval error ... 5776
Windows 7 to Windows 8 client upgrade ... 5777
General cmdlet errors ... 5778

Deploy Remote Access with OTP Authentication ... 5778
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set ... 5779
Deploy Basic Remote Access ... 5779
Deploy Advanced Remote Access ... 5779
Deploy Remote Access in an Enterprise ... 5779
Manage Remote Access ... 5779
Migrate Remote Access ... 5779
Scenario description ... 5780
Prerequisites ... 5780
In this scenario ... 5780
Practical applications ... 5781
Roles and features included in this scenario ... 5781
Hardware requirements ... 5782
Software requirements ... 5783
Known Issues ... 5784

Plan Remote Access with OTP Authentication ... 5784

Step 1: Plan an Advanced Single Server Deployment ... 5785
Plan a single server deployment ... 5785

Step 2: Plan the RADIUS Server Deployment ... 5785

2.1 Plan the RADIUS server ... 5785
See also ... 5786

Step 3: Plan OTP Certificate Deployment ... 5786
3.1 Plan the OTP CA ... 5786
3.2 Plan the OTP certificate template ... 5787
3.3 Plan the registration authority certificate ... 5787
See also ... 5788

Step 4: Plan for OTP on the Remote Access Server ... 5788
4.1 Plan for OTP client exemptions ... 5788
4.2 Plan for Windows 7 clients ... 5788
4.3 Plan for smart cards ... 5789
See also ... 5789

Configure DirectAccess with OTP Authentication ... 5789

Step 1: Implement a Single Server Remote Access Deployment ... 5789
Implement a single server deployment ... 5789
See also ... 5790

Step 2: Configure the RADIUS Server ... 5790
2.1 Configure the RADIUS software distribution tokens ... 5790
2.2 Configure the RADIUS security information ... 5790
2.3 Adding user account for OTP probing ... 5791
2.4 Synchronize with Active Directory ... 5791
2.5 Configure the RADIUS authentication agent ... 5791

Step 3: Configure the Remote Access Server for OTP ... 5791
3.1 Exempt users from OTP authentication (optional) ... 5792
3.2 Configure the Remote Access server to support OTP ... 5792
3.3 Smart cards for additional authorization ... 5794
Allowing access for users with unusable smart cards ... 5795
Under the covers: Smart card authorization ... 5795

Step 4: Verify DirectAccess with OTP ... 5795

Troubleshoot an OTP Deployment ... 5796

Troubleshooting Authentication Issues ... 5797
Failed to access the CA that issues OTP certificates ... 5797
DirectAccess server connectivity issues ... 5797
Failed to enroll for the DirectAccess OTP logon certificate ... 5798
Missing or invalid computer account certificate ... 5798
Missing CA that issues OTP certificates ... 5799
Misconfigured DirectAccess server address ... 5799

Failed to generate the OTP logon certificate request 5800
No connection to the domain controller 5800
OTP provider requires challenge/response 5801
Incorrect OTP logon template used 5801
Missing OTP signing certificate 5801
Missing or incorrect UPN/DN for the user 5802
OTP certificate is not trusted for login 5802
Windows could not verify user credentials 5803

Troubleshooting Enabling OTP 5803
Failed to enroll the OTP signing certificate 5804
Failed to enable DirectAccess OTP when WebDAV is installed 5804
No templates available in the Remote Access Management console 5805
Cannot set renewal period of OTP template to 1 hour 5805

Deploy Remote Access in a Multi-Forest Environment 5806

Plan a Multi-Forest Deployment 5806
Prerequisites 5806
Plan trust between forests 5806
Plan Remote Access administrator permissions 5806
Plan client security groups 5807
Plan certification authorities 5807
Plan OTP exemptions 5808

Configure a Multi-Forest Deployment 5808
Access resources from Forest2 5808
Enable clients from Forest2 to connect via DirectAccess 5808
Add entry points from Forest2 5809
Configure OTP in a multi-forest deployment 5809
Configuration procedures 5811
Add NRPT rules and DNS suffixes 5811
Add internal IPv6 prefix 5812
Add client security groups 5813
Refresh the management servers list 5813

Deploy Remote Access in the Cloud 5814
Scenario description 5814
In this scenario 5814
Practical applications 5815
Cloud Bursting 5815
Disaster Recovery 5815
Roles and features included in this scenario 5816
Hardware requirements 5817
Software requirements 5817

See also 5817

Test Lab Guide : Corp to Cloud: Configure an IKEv2-based Site-to-site VPN Connection 5818
About this guide 5818

Overview of the test lab scenario 5819

Configuration component requirements 5821

Steps for configuring the test lab 5821

STEP 1: Complete the base configuration 5821
STEP 1: Complete the Base Configuration 5821

STEP 2: Install and Configure 3-EDGE1 5822
STEP 2: Install and Configure 3-EDGE1 5822

STEP 3: Configure routing and forwarding between EDGE-1 and 3-EDGE1 using pre-shared key authentication 5824
STEP 3: Configure routing and forwarding between EDGE1 and 3-EDGE1 using EAP and (optionally) preshared key authentication 5824

STEP 4: Install a basic PKI certificate infrastructure 5829
Step 4: Install a basic PKI certificate infrastructure for site-to-site certificate authentication method. 5829

STEP 5: Configure EDGE1 and 3-EDGE1 to use EAP for site-to-site authentication 5839

STEP 6: Configure EDGE1 and 3-EDGE1 to use certificates for site-to-site authentication 5840
Configure computer certificate authentication on EDGE1 5840
Configure name resolution 5841
Export the root certificate (corp-DC1-CA) from EDGE1 5841
Configure computer certificate authentication on 3-EDGE1 5842
Configure name resolution 5842
Export the root certificate (corp-3-DC1-CA) from 3-EDGE1 5843
Import the root certificate (corp-DC1-CA) to 3-EDGE1 5843
Change the authentication method to certificates on EDGE1 5844
Change the authentication method to certificates on 3-EDGE1 5844
Reinitiate the site-to-site connection 5845
Test the site-to-site connection 5845

Snapshot the configuration 5845

Manage Remote Access 5846
Scenario description 5846
In this scenario 5846
Plan the deployment 5846

Configure the deployment ... 5847
Practical applications ... 5847
Roles and features included in this scenario ... 5848
Hardware requirements ... 5849
Server requirements ... 5849
Client requirements ... 5849
Infrastructure and management server requirements ... 5849
Software requirements ... 5850
Server requirements ... 5850
Remote access client requirements ... 5850
See also ... 5850

Use Remote Access Monitoring and Accounting ... 5851
In this guide ... 5851
Understand monitoringand accounting ... 5851

Monitor the existing load on the Remote Access server ... 5852

Monitor the configuration distribution status of the Remote Access server ... 5853

Monitor the operations status of the Remote Access server and its components ... 5856

Identify and resolve Remote Access server operations problems ... 5857
Simulate an operations issue ... 5858
Identify the operations issue and take corrective action ... 5858
Restore the IP Helper service ... 5859

Monitor connected remote clients for activity and status ... 5859

Generate a usage report for remote clients using historical data ... 5861

Manage DirectAccess Clients Remotely ... 5863
Scenario description ... 5863
In this scenario ... 5863
Planning steps ... 5863
Prerequisites ... 5864
Deployment steps ... 5864
Practical applications ... 5864
Roles and features included in this scenario ... 5865
Hardware requirements ... 5866
Software requirements ... 5867
See also ... 5867

Plan Deployment for Remote Management of DirectAccess Clients ... 5868
Next steps ... 5869

Step 1: Plan the Remote Access Infrastructure ... 5869
Plan network topology and settings ... 5870
Plan network adapters and IP addressing ... 5870
Plan ISATAP requirements ... 5874
Plan firewall requirements ... 5878
Plan certificate requirements ... 5879
Plan computer certificates for IPsec authentication ... 5880
Plan certificates for IP-HTTPS ... 5880
Plan website certificates for the network location server ... 5881
Plan DNS requirements ... 5881
DirectAccess client requests ... 5881
Infrastructure servers ... 5882
Connectivity verifiers ... 5882
DNS server requirements ... 5883
Plan for local name resolution ... 5883
NRPT ... 5883
Single label names ... 5884
Split-brain DNS ... 5884
Plan local name resolution behavior for DirectAccess clients ... 5885
Plan the network location server configuration ... 5885
Plan certificates for the network location server ... 5886
Plan DNS for the network location server ... 5886
Plan management servers' configuration ... 5886
Plan Active Directory requirements ... 5887
Plan client authentication ... 5888
Plan multiple domains ... 5888
Plan Group Policy Object creation ... 5889
Automatically created GPOs ... 5889
Manually created GPOs ... 5890
Recovering from a deleted GPO ... 5890
See also ... 5890

Step 2: Plan the Remote Access Deployment ... 5891
Plan a client deployment strategy ... 5891
Plan a Remote Access server deployment strategy ... 5892
Plan the infrastructure servers' configurations ... 5894
See also ... 5894

Install and Configure Deployment for Remote Management of DirectAccess Clients ... 5894
See also ... 5895

Step 1: Configure the Remote Access Infrastructure ... 5895
Configure server network settings ... 5896
Configure routing in the corporate network ... 5897

Configure firewalls ... 5897
Remote Access server on IPv4 Internet ... 5897
Remote Access server on IPv6 Internet ... 5897
Remote Access traffic ... 5898
Configure CAs and certificates ... 5898
Configure IPsec authentication ... 5898
Configure certificate templates ... 5898
Configure the IP-HTTPS certificate ... 5899
Configure the DNS server ... 5900
Configure Active Directory ... 5901
Configure GPOs ... 5903
Configure security groups ... 5903
Configure the network location server ... 5904
See also ... 5906

Step 2: Configure the Remote Access Server ... 5906
Install the Remote Access role ... 5907
Configure the deployment type ... 5907
Configure DirectAccess clients ... 5908
Configure the Remote Access server ... 5909
Configure the infrastructure servers ... 5910
Configure application servers ... 5910
Configuration summary and alternate GPOs ... 5910
See also ... 5911

Step 3: Verify the Deployment ... 5911

Deploy VPN When Connecting Remotely with Windows 8 ... 5912
Deploy Windows 8 VPN client ... 5912
Use self-service portal ... 5912
Use Windows PowerShell scripts ... 5913
Advanced scenarios ... 5914
Use Connection Manager Administration Kit ... 5915
Manage profiles ... 5915
Connect VPN by using Windows 8 user interface ... 5917
Launch Windows 8 user interface for VPN connections ... 5918
Connect to VPN ... 5918
Deploy the VPN ... 5919
Password-based deployment ... 5919
Certificate-based deployments ... 5919
Enable split tunneling ... 5920
Understand route implications ... 5920
Deploy web proxy and intranet settings ... 5921
Deploy multisite VPN ... 5922

Interoperate with third-party VPN servers 5922
Use Connected Standby state 5922
Appendix A: Glossary 5922
Experience terminology 5923

Manage VPN Connections in Windows 8 by Using Windows PowerShell 5924
Use VPN deployment and management scripts: For IT pros 5925
Create VPN connections 5925
Configure and edit VPN connections 5926
New-EapConfiguration 5926
VPN connection lookup 5927
VPN connection removal 5927
Write configuration scripts for VPN connection management 5927
Use VPN deployment and management scripts: For knowledge workers 5931
Use Windows PowerShell cmdlets for third-party EAP methods 5932

Set Up and Edit VPN Connections in Windows 8 5932
Create a new VPN connection 5933
Discovery Profile 5933
Connect to a corporate network by use an existing VPN connection 5934
Edit connection properties 5938
Delete a connection 5939

VPN Compatibility and Interoperability in Windows 8 and Windows Server 2012 5940

Windows Server 2012 R2 Test Lab Guide: Demonstrate VPN Auto trigger 5940
In this Guide 5941
Test Lab Overview 5941
Hardware and software requirements 5942
Steps for Configuring the Remote Access VPN Auto-trigger Test Lab 5942
Step 1: Set up the Base Configuration Test Lab 5943
Step 2: Configure EDGE1 5943
Deploy VPN access using the Remote Access Configuration wizard 5943
Step 3: Configure INET1 5944
Create a new DNS host record on INET1 5945
Step 4: Add and Configure CLIENT2 5945
Install the operating system on CLIENT2 5945
Connect CLIENT2 to the Internet subnet and configure remote access with VPN 5946
Configure and test VPN auto-trigger functionality on CLIENT2 5946
Connect CLIENT2 to the Corpnet subnet and verify that auto-trigger functionality is disabled 5948
Snapshot the Configuration 5949

Deploy and Configure a VPN for Devices Running Windows RT 5949
Work with third-party VPN servers 5949

Verify third-party VPN interoperability ... 5950
VPN deployment options ... 5951
Multi-site VPN deployment ... 5951
Windows RT VPN client provisioning ... 5951
Get Connected Wizard (GCW) ... 5951
Single-Click (PowerShell) scripts ... 5952
Provisioning via System Center ... 5952
Profile management ... 5953
Enabling split-tunneling ... 5954
Implications on routes ... 5954
Provisioning web proxy and Intranet settings for VPN connection ... 5954
Connect to the VPN ... 5955
Use one-time password for authentication ... 5955

Migrate Remote Access to Windows Server 2012 ... 5956
About this guide ... 5957
Target audience ... 5957
What this guide does not provide ... 5957
Supported migration scenarios ... 5958
Supported operating systems ... 5958
Supported role configurations ... 5959
Migration dependencies ... 5959
Migration components that are not supported in all operating system versions ... 5959
Migration components that are not automatically migrated ... 5962
Overview of the Routing and Remote Access service migration process ... 5964
Impact of migration ... 5964
Permissions required to complete migration ... 5965
Estimated duration ... 5965
See Also ... 5966

Remote Access: Prepare to Migrate ... 5966
Prepare your destination server ... 5966
Hardware requirements for the destination server ... 5966
Prepare the destination server for migration ... 5967
Prepare your source server ... 5968
Back up your source server ... 5968
Install the migration tools ... 5968
See Also ... 5969

Remote Access: Migrate Remote Access ... 5969
Migrating Remote Access from the source server ... 5969
Migrating Remote Access to the destination server ... 5973
Completing the required manual migration steps ... 5974
DirectAccess ... 5974

Dial-up demand-dial connections ... 5974
Certificates for IKEv2, SSTP, and L2TP/IPsec connections ... 5975
Routing and Remote Access service policies and accounting settings ... 5975
PEAP, smart card, and other certificate settings on Network Policy Server ... 5975
Weak encryption settings ... 5975
Connection Manager profiles ... 5976
Group forwarded fragments ... 5976
RAS administration and security DLLs ... 5976
See Also ... 5976

Remote Access: Verify the Migration ... 5977
Verifying the destination server configuration ... 5977
Installation state of Remote Access ... 5977
Status of Remote Access Service ... 5977
Remote access Operations Status ... 5978
DirectAccess configuration ... 5978
VPN configuration ... 5978
Dial-up configuration ... 5979
Demand-dial VPN configuration ... 5979
Router settings ... 5979
User and Group accounts ... 5981
Final checks ... 5981
See Also ... 5981

Remote Access: Post-migration Tasks ... 5981
Completing the migration ... 5982
Configuring firewall rules for VPN ... 5982
Configuring firewall rules for DirectAccess ... 5982
Restoring Remote Access in the event of migration failure ... 5983
Estimated time to complete a rollback ... 5984
Retiring Remote Access on your source server ... 5984
Troubleshooting cmdlet-based migration ... 5984
Viewing the content of Windows Server Migration Tools result objects ... 5985
Result object descriptions ... 5985
Examples ... 5987
More information about querying results ... 5989
See Also ... 5989

Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012 ... 5989
Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set ... 5989
Deploy Basic Remote Access ... 5990
Deploy Advanced Remote Access ... 5990
Deploy Remote Access in an Enterprise ... 5990
Manage Remote Access ... 5990

Migrate Remote Access ... 5990
Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites ... 5990
Scenario description ... 5991
In this scenario ... 5991
Prerequisites ... 5991
Practical applications ... 5991
Hardware requirements ... 5992
Software requirements ... 5992
Using ISATAP ... 5993
Using NAP ... 5993

Side-by-Side Migration of Forefront UAG DirectAccess ... 5993

Side-by-Side Migration Steps ... 5994
Step 1: Configure ISATAP ... 5995
Step 2: Export the Forefront UAG DirectAccess settings ... 5999
Step 3: Collect Forefront UAG DirectAccess GPO settings ... 5999
Step 4: Install the Remote Access role ... 5999
Step 5: Configure server and infrastructure settings ... 6000
Step 6: Obtain an IP-HTTPS certificate ... 6000
Step 7: Create a security group for DirectAccess clients ... 6000
Step 8: Prepare GPOs ... 6000
Step 9: Configure DirectAccess ... 6001
See also ... 6004

Offline Migration of Forefront UAG DirectAccess ... 6004

Offline Migration Steps ... 6005
Step 1: Install the Remote Access role ... 6006
Step 2: Configure IP addresses ... 6006
Step 3: Obtain an IP-HTTPS certificate ... 6007
Step 4: Prepare GPOs ... 6007
Step 5: Configure DirectAccess ... 6007
See also ... 6007

Connection Manager Overview ... 6007

Connection Manager Administration Kit Overview ... 6008

Add DirectAccess to an Existing Remote Access (VPN) Deployment ... 6008
Deployment documentation set for Windows Server 2012 Remote Access (DirectAccess) ... 6008
Basic Remote Access deployment ... 6009
Advanced Remote Access deployment ... 6009
Remote Access deployment in an enterprise ... 6009

Manage Remote Access ... 6009
Migrate Remote Access ... 6009
Scenario description ... 6010
In this scenario ... 6010
Planning steps ... 6010
Deployment steps ... 6010
Practical applications ... 6010
Roles and features required for this scenario ... 6011
Hardware requirements ... 6012
Software requirements ... 6013
See also ... 6014

Plan to Enable DirectAccess ... 6014
See also ... 6015

Step 1: Plan the DirectAccess Infrastructure ... 6015
Plan network topology and settings ... 6016
Plan network adapters and IP addressing ... 6016
Plan firewall requirements ... 6020
Plan certificate requirements ... 6020
Plan certificates for IP-HTTPS ... 6021
Plan website certificates for the network location server ... 6022
Plan DNS requirements ... 6022
DNS server requirements ... 6023
Plan Active Directory ... 6023
Plan Group Policy Objects ... 6024
Automatically-created GPOs ... 6025
Manually-created GPOs ... 6026
Recovering from a deleted GPO ... 6026
See also ... 6026

Step 2: Plan the DirectAccess Deployment ... 6026
Planning for client deployment ... 6027
Planning for Remote Access server deployment ... 6027
See also ... 6028

Enable DirectAccess ... 6028

Step 1: Configure the DirectAccess Infrastructure ... 6029
Configure server network settings ... 6030
Configure routing in the corporate network ... 6030
Configure firewalls ... 6031
Configure CAs and certificates ... 6031
Configure certificate templates ... 6031
Configure the IP-HTTPS certificate ... 6031

Configure the DNS server ... 6033
Configure Active Directory ... 6034
Configure GPOs ... 6035
Configure security groups ... 6035
Configure the network location server ... 6036
See also ... 6037

Step 2: Configure the DirectAccess Server ... 6037
Configure DirectAccess clients ... 6038
Configure the Network Topology ... 6038
Configure the DNS Suffix Search List ... 6039
GPO Configuration ... 6040
Summary ... 6040

Step 3: Verify the Deployment ... 6040

Working with Web Application Proxy ... 6041

Web Application Proxy Overview ... 6041

Getting Started with Web Application Proxy ... 6041
Providing Access to Applications ... 6042
Publishing Applications ... 6042
Accessing Applications ... 6042
Protecting Applications from External Threats ... 6043
Defense in Depth ... 6043
Authentication and Authorization ... 6043
Authenticating Users and Devices ... 6044
Authentication Capabilities ... 6045
Web Application Proxy Technical Overview ... 6045
Web Application Proxy Configuration Storage ... 6046
AD FS Proxy Functionality ... 6046
Managing Web Application Proxy ... 6047
Interoperability with Other Remote Access Products ... 6047

Overview: Connect to Applications and Services from Anywhere with Web Application Proxy 6048
Providing Access to Applications and Services ... 6049
Using Active Directory Federation Services ... 6049
AD FS Proxy ... 6050
Roles and Features Included in this Scenario ... 6050
Scenario Steps ... 6050
See Also ... 6051

Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy ... 6051

Step 1: Attempt to access the web application from an Internet client 6053
Step 2: Configure the Web Application Proxy server and publish the application 6054
Step 3: Configure and test accessing a website using Integrated Windows authentication .. 6057
Install Windows Authentication on WebServ1 6057
Create a new website using IIS 6057
Create a non-claims-aware relying party trust 6058
Configure Kerberos constrained delegation 6059
Test accessing the application internally 6060
Publish the application 6060
Test accessing the application 6061
Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control 6061
See also 6063

Installing and Configuring Web Application Proxy for Publishing Internal Applications 6063
Providing Access to Applications and Services 6064
Roles and Features Included in this Scenario 6064
Scenario Steps 6064
Hardware requirements 6065
Software requirements 6065
Known issues 6065

Plan to Publish Applications through Web Application Proxy 6066

Step 1: Plan the Web Application Proxy Infrastructure 6067
1.1. Plan Network Location 6067
Plan Firewalls 6068
1.2. Plan DNS 6068
1.3. Plan Load Balancing 6068
1.4. Plan Active Directory 6069
Plan Integrated Windows authentication and Kerberos constrained delegation 6069
1.5. Plan Active Directory Federation Services 6070
AD FS Requirements 6070
1.6. Plan Preauthentication 6070
See also 6071

Step 2: Plan the Web Application Proxy Server 6071
2.1. Plan the Web Application Proxy Role Service Installation 6071
2.2. Plan Multiple Servers 6072
2.3. Plan Certificates 6072
2.4. Plan NTP 6074
See also 6074

Step 3: Plan to Publish Applications using AD FS Preauthentication 6074
3.1. Plan Claims-Based Applications 6076

3.2. Plan Integrated Windows Authentication-Based Applications 6076
3.3. Plan Applications for MSOFBA Clients 6077
3.4. Plan Applications for Clients that use Windows Store Apps 6078
See also 6079

Step 4: Plan to Publish Applications using Client Certificate Preauthentication 6080
4.1. Plan the External Servers 6080
4.2. Plan Applications for Client Certificate Preauthentication 6081
See also 6081

Step 5: Plan to Publish Applications using Pass-through Preauthentication 6081
5.1. Plan Applications for Pass-Through Preauthentication 6082
See also 6082

Step 6: Plan to Publish SharePoint Server and Exchange Server 6082
6.1. Plan to publish SharePoint Server 6083
6.2. Plan to publish Exchange Server 6083
See also 6084

Configure Publishing Applications through Web Application Proxy 6084

Step 1: Configure the Web Application Proxy Infrastructure 6085
1.1. Configure server network settings 6085
1.2. Configure firewalls 6086
1.3. Configure DNS 6086
1.4. Configure Active Directory 6086
Configure Kerberos constrained delegation 6087
1.5. Configure Active Directory Federation Services 6088
See also 6088

Step 2: Install and Configure the Web Application Proxy Server 6088
2.1. Configure CAs and certificates 6089
Configure certificate templates 6089
Configure web application certificates 6089
2.2. Install the Remote Access role 6090
2.3. Configure Web Application Proxy 6091
See also 6092

Step 3: Publish Applications using AD FS Preauthentication 6092
3.1. Publish a Claims-based Application for Web Browser Clients 6093
3.2. Publish an Integrated Windows authenticated-based Application for Web Browser Clients 6095
3.3. Publish an Application that uses MSOFBA 6096
3.4. Publish an Application for a Windows Store App 6096
See also 6097

Step 4: Publish Applications using Client Certificate Preauthentication 6097
4.1. Publish an Application using Client Certificate Preauthentication 6098
See also 6098

Step 5: Publish Applications using Pass-through Preauthentication 6099
5.1. Publish an Application using Pass-Through Preauthentication 6099
See also 6100

Best Practices Analyzer for Web Application Proxy 6101
More information about Web Application Proxy 6101
Topics in this section 6101

Web Application Proxy must be configured before it is used 6102
Issue 6102
Impact 6102
Resolution 6102

Web Application Proxy: The external and backend server URLs are different and URL translation is disabled 6103
Issue 6104
Impact 6104
Resolution 6104

Web Application Proxy: The service is not configured to run automatically 6106
Issue 6106
Impact 6106
Resolution 6106

Web Application Proxy: The AD FS Proxy service is not configured to run automatically 6107
Issue 6107
Impact 6107
Resolution 6107

Web Application Proxy: This server is not included in the ConnectedServersName list 6108
Issue 6108
Impact 6108
Resolution 6109

Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high 6109
Issue 6110
Impact 6110
Resolution 6110

Web Application Proxy: Application is using an external certificate that is not yet valid 6111
Issue 6111
Impact 6111

Resolution 6111

Web Application Proxy: Application is using an external certificate that is about to expire 6113
Issue 6114
Impact 6114
Resolution 6114

Web Application Proxy: Application is using an external certificate that has no private key 6116
Issue 6117
Impact 6117
Resolution 6117

Web Application Proxy: Application is using an external certificate that has expired 6119
Issue 6120
Impact 6120
Resolution 6120

Web Application Proxy: Application is configured to use an external certificate that is not present on this server 6122
Issue 6122
Impact 6123
Resolution 6123

Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows authentication but the server is not joined to a domain 6124
Issue 6124
Impact 6124
Resolution 6125

Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed 6126
Issue 6126
Impact 6126
Resolution 6126

Web Application Proxy: Could not perform Integrated Windows authentication to the backend servers 6127
Issue 6128
Impact 6128
Resolution 6128
Additional references 6129

Web Application Proxy could not publish an application 6129
Issue 6129
Impact 6129
Resolution 6129

Web Application Proxy could not publish an application due to certificate problems 6130
Issue 6130
Impact 6130
Resolution 6130

Web Application Proxy was not able to check for configuration updates 6132
Issue 6132
Impact 6133
Resolution 6133

Web Application Proxy stopped working because it could not retrieve its configuration 6133
Issue 6134
Impact 6134
Resolution 6134

Web Application Proxy: The AD FS Proxy service is stopped 6135
Issue 6135
Impact 6135
Resolution 6135

Web Application Proxy service is stopped 6136
Issue 6136
Impact 6136
Resolution 6137

TCP/IP 6137

Telephony Server and the Telephony API Overview 6138

Virtual Receive-side Scaling in Windows Server 2012 R2 6138
Feature description 6138
Practical applications 6139
Processor load balancing 6139
Software load balancing 6139
Software requirements 6139
Hardware requirements 6139
Enable and Configure Virtual Receive-side scaling 6140
Before enabling Virtual Receive-side scaling 6140
Enabling Virtual Receive-side scaling 6140
Configuring Virtual Receive-side scaling 6141
Inside the Host: 6142
Disabling Virtual Receive-side scaling 6142
Managing Virtual Receive-side scaling FAQ 6143

Windows Firewall with Advanced Security Overview 6144
Feature description 6144

Practical applications 6144
New and changed functionality 6145
IKEv2 for IPsec transport mode 6145
Windows Store app network isolation 6145
Windows PowerShell cmdlets for Windows Firewall 6146
See also 6146

Isolating Windows Store Apps on Your Network 6146
Prerequisites 6147
Step 1: Define your network 6148
Step 2: Create custom firewall rules 6149
See also 6154

Securing End-to-End IPsec Connections by Using IKEv2 in Windows Server 2012 6154
Prerequisites 6155
Computers joined to a domain 6155
Computers not joined to a domain 6157
Troubleshooting 6158
See also 6159

Windows Firewall with Advanced Security Administration with Windows PowerShell 6159
Scope 6159
Audience and user requirements 6160
System requirements 6160
In this guide 6160
Set profile global defaults 6161
Enable Windows Firewall 6161
Control firewall behavior 6161
Deploy basic firewall rules 6162
Create firewall rules 6162
GPO Caching 6163
Modify an existing firewall rule 6163
Delete a firewall rule 6165
Manage remotely 6166
Deploy basic IPsec rule settings 6166
Create IPsec rules 6167
Add custom authentication methods to an IPsec rule 6167
IKEv2 IPsec transport rules 6168
Copy an IPsec rule from one policy to another 6169
Handling Windows PowerShell errors 6169
Monitor 6170
Find the source GPO of a rule 6170
Deploy a basic domain isolation policy 6171
Configure IPsec tunnel mode 6171

Deploy secure firewall rules with IPsec ... 6172
Create a secure firewall rule (allow if secure) ... 6172
Isolate a server by requiring encryption and group membership ... 6173
Create a firewall rule that requires group membership and encryption ... 6174
Endpoint security enforcement ... 6175
Create firewall rules that allow IPsec-protected network traffic (authenticated bypass) ... 6175
Additional resources ... 6176

Windows Firewall with Advanced Security Design Guide ... 6176
About this guide ... 6177
Terminology used in this guide ... 6178

Understanding the Windows Firewall with Advanced Security Design Process ... 6181

Identifying Your Windows Firewall with Advanced Security Deployment Goals ... 6182

Protect Computers from Unwanted Network Traffic ... 6183

Restrict Access to Only Trusted Computers ... 6184

Require Encryption When Accessing Sensitive Network Resources ... 6186

Restrict Access to Only Specified Users or Computers ... 6188

Mapping Your Deployment Goals to a Windows Firewall with Advanced Security Design ... 6190

Basic Firewall Policy Design ... 6191

Domain Isolation Policy Design ... 6193

Server Isolation Policy Design ... 6195

Certificate-based Isolation Policy Design ... 6197

Evaluating Windows Firewall with Advanced Security Design Examples ... 6198

Firewall Policy Design Example ... 6199
Design requirements ... 6199
Design details ... 6201

Domain Isolation Policy Design Example ... 6203
Design Requirements ... 6203
Design Details ... 6204

Server Isolation Policy Design Example ... 6206
Server isolation without domain isolation ... 6206
Design requirements ... 6207
Design details ... 6208

Certificate-based Isolation Policy Design Example 6209
Design requirements 6209
Design details 6210

Designing a Windows Firewall with Advanced Security Strategy 6211
If you already have firewall or IPsec rules deployed 6212

Gathering the Information You Need 6212

Gathering Information about Your Current Network Infrastructure 6212
Network segmentation 6213
Network address translation (NAT) 6214
Network infrastructure devices 6214
Current network traffic model 6216

Gathering Information about Your Active Directory Deployment 6217

Gathering Information about Your Computers 6218
Automated Discovery 6219
Manual Discovery 6219

Gathering Other Relevant Information 6219
Capacity considerations 6219
Group Policy deployment groups and WMI filters 6220
Different Active Directory trust environments 6220
Creating firewall rules to permit IKE, AH, and ESP traffic 6221
Network load balancing and server clusters 6221
Network inspection technologies 6222

Determining the Trusted State of Your Computers 6223
Trust states 6223
Trusted state 6223
Trustworthy state 6224
Known, untrusted state 6225
Unknown, untrusted state 6226
Capturing upgrade costs for current computers 6226

Planning Your Windows Firewall with Advanced Security Design 6228
Basic firewall design 6228
Algorithm and method support and selection 6228
IPsec performance considerations 6228
Domain isolation design 6229
Server isolation design 6229
Certificate-based authentication design 6229
Documenting your design 6230
Designing groups and GPOs 6230

Planning Settings for a Basic Firewall Policy 6230

Planning Domain Isolation Zones 6232

Exemption List 6232

Isolated Domain 6233
GPO settings for isolated domain members running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 6234

Boundary Zone 6235
GPO settings for boundary zone servers running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2 6237

Encryption Zone 6238
GPO settings for encryption zone servers running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2 6238

Planning Server Isolation Zones 6239
Isolated domains and isolated servers 6240
Creating multiple isolated server zones 6240
Creating the GPOs 6240
GPO settings for isolated servers running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 6241

Planning Certificate-based Authentication 6242
Deploying certificates 6242
Using Active Directory Certificate Services 6243
Using a commercially purchased certificate for computers running Windows 6243
Using a commercially purchased certificate for computers running a non-Windows operating system 6243
Configuring IPsec to use the certificates 6243

Documenting the Zones 6244

Planning Group Policy Deployment for Your Isolation Zones 6245

Planning Isolation Groups for the Zones 6245

Planning Network Access Groups 6247

Planning the GPOs 6248
General considerations 6248
Woodgrove Bank example GPOs 6249

Firewall GPOs 6250

GPO_DOMISO_Firewall 6250

Firewall settings 6250
Firewall rules 6251

Isolated Domain GPOs 6251

GPO_DOMISO_IsolatedDomain_Clients 6252
General settings 6252
Connection Security Rules 6253

GPO_DOMISO_IsolatedDomain_Servers 6254

Boundary Zone GPOs 6254

GPO_DOMISO_Boundary_WS2008 6255
IPsec settings 6255
Connection security rules 6255
Registry settings 6256
Firewall rules 6256

Encryption Zone GPOs 6256

GPO_DOMISO_Encryption_WS2008 6256
IPsec settings 6257
Connection security rules 6257
Registry settings 6257
Firewall rules 6257

Server Isolation GPOs 6257
GPO_SRVISO_WS2008 6258

Planning GPO Deployment 6258
General considerations 6259
Test your deployed groups and GPOs 6259
Do not enable require mode until deployment is complete 6260
Example Woodgrove Bank deployment plans 6260
GPO_DOMISO_Firewall_2008_Win7-Vista 6260
GPO_DOMISO_IsolatedDomain_Clients_Win7Vista 6261
GPO_DOMISO_IsolatedDomain_Servers_WS2008 6261
GPO_DOMISO_Boundary_WS2008 6261
GPO_DOMISO_Encryption_WS2008 6261

Appendix A: Sample GPO Template Files for Settings Used in this Guide 6262

Additional Resources 6264
Windows Firewall with Advanced Security 6264
IPsec 6265
Server and Domain Isolation 6265

Group Policy ... 6265
Active Directory Domain Services ... 6265

Windows Firewall with Advanced Security Deployment Guide ... 6266
About this guide ... 6266
What this guide does not provide ... 6267
Overview of Windows Firewall with Advanced Security ... 6267

Planning to Deploy Windows Firewall with Advanced Security ... 6268
Reviewing your Windows Firewall with Advanced Security Design ... 6268

Implementing Your Windows Firewall with Advanced Security Design Plan ... 6269
How to implement your Windows Firewall with Advanced Security design using this guide ... 6270

Checklist: Creating Group Policy Objects ... 6271
About membership groups ... 6271
About exclusion groups ... 6272

Checklist: Implementing a Basic Firewall Policy Design ... 6273

Checklist: Configuring Basic Firewall Settings ... 6275

Checklist: Creating Inbound Firewall Rules ... 6276

Checklist: Creating Outbound Firewall Rules ... 6277

Checklist: Implementing a Domain Isolation Policy Design ... 6278

Checklist: Configuring Rules for the Isolated Domain ... 6280

Checklist: Configuring Rules for the Boundary Zone ... 6282

Checklist: Configuring Rules for the Encryption Zone ... 6283

Checklist: Configuring Rules for an Isolated Server Zone ... 6284

Checklist: Implementing a Standalone Server Isolation Policy Design ... 6288

Checklist: Configuring Rules for Servers in a Standalone Isolated Server Zone ... 6289

Checklist: Creating Rules for Clients of a Standalone Isolated Server Zone ... 6293

Checklist: Implementing a Certificate-based Isolation Policy Design ... 6295

Procedures Used in This Guide ... 6296

Add Production Computers to the Membership Group for a Zone ... 6298

Add Test Computers to the Membership Group for a Zone ... 6299

Assign Security Group Filters to the GPO 6300

Change Rules from Request to Require Mode 6302

Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6303

Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6306

Configure Group Policy to Autoenroll and Deploy Certificates 6308

Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6308

Configure the Rules to Require Encryption on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6310

Configure the Windows Firewall Log 6311

Configure the Workstation Authentication Certificate Template 6312

Configure Windows Firewall to Suppress Notifications When a Program Is Blocked 6313

Confirm That Certificates Are Deployed Correctly 6315

Copy a GPO to Create a New GPO 6315

Create a Group Account in Active Directory 6316

Create a Group Policy Object 6317

Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6318

Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6319

Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6322

Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6324

Create an Inbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6326

Create an Outbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2 6328

Create an Outbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2 6329

Create Inbound Rules to Support RPC on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6331

Create WMI Filters for the GPO 6333

Enable Predefined Inbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6335

Enable Predefined Outbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6336

Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 6337

Install Active Directory Certificate Services 6338

Link the GPO to the Domain 6339

Modify GPO Filters to Apply to a Different Zone or Version of Windows 6340

Open the Group Policy Management Console to IP Security Policies 6342

Open the Group Policy Management Console to Windows Firewall 6342

Open the Group Policy Management Console to Windows Firewall with Advanced Security 6342

Open Windows Firewall with Advanced Security 6343
- Opening Windows Firewall with Advanced Security 6343

Restrict Server Access to Members of a Group Only 6344

Start a Command Prompt as an Administrator 6345

Turn on Windows Firewall and Configure Default Behavior 6346

Verify That Network Traffic Is Authenticated 6346
- For computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 6347

Additional Resources 6348
- Windows Firewall with Advanced Security 6348
- IPsec 6349
- Group Policy 6349
- Active Directory Domain Services 6349

Windows Internet Name Service (WINS) Overview 6349

Windows PowerShell Cmdlets for Networking 6350

Windows Server Gateway 6351
What is Windows Server Gateway? 6352
Windows Server Gateway integration with Hyper-V Network Virtualization 6353
Clustering Windows Server Gateway for high availability 6353
Windows Server Gateway as a forwarding gateway for private cloud environments 6354
Windows Server Gateway as a site-to-site VPN gateway for hybrid cloud environments . 6355
Multitenant Network Address Translation (NAT) for VM Internet access 6355
Multitenant remote access VPN connections 6356

Windows Server Gateway Hardware and Configuration Requirements 6357
Hyper-V hardware recommendations 6357
Failover Clustering and Load Balancing 6358
Hyper-V host configuration 6359
Windows Server Gateway VM configuration 6361
See Also 6363

Print and Document Services Overview 6363
Role description 6363
Practical applications 6364
New and changed functionality 6364
Server Manager information 6364
See also 6365

What's New in Print and Document Services in Windows Server 2012 R2 6366
Role and technology description 6366
New and changed functionality 6366
Event Logging for Branch Office Direct Printing 6367
Printer Migration for Web Services for Devices (WSD) print devices 6368
Roaming Settings include Printer Connections 6369
Easier Printing in Windows RT 6369
Near Field Communication (NFC) Connections to Printers 6369
Common framework for PIN-protected printing support by IHVs 6370
Print and Fax services now include user access logging 6370
See Also 6370

Install Print and Document Services 6371
Deployment Scenarios 6371
Print Server 6371
Internet Printing 6371
LPD Service 6371
Installation Process 6372
Installation UI/Wizard 6372
Installing Print and Document Services 6373

Verifying Installation ... 6374
Installing on Server Core ... 6375
Installing with Minimal Server Interface ... 6375
Uninstalling Role Services ... 6376
See Also ... 6376

Install Distributed Scan Server ... 6376
Step 1: Install software ... 6377
Step 2: Configure the server ... 6378
Step 3: Add and share scan devices ... 6380
Additional resources ... 6381

Distributed Scan Management (DSM) Scanners ... 6382

Fax Server Step-by-Step Guide ... 6382
What is a fax server? ... 6382
Who should use a fax server? ... 6383
Benefits of a fax server ... 6383
Requirements for using the Fax service ... 6383
Deploying a fax server ... 6383
Step 1: Install and open Fax Service Manager ... 6384
Step 2: Install and share a fax printer ... 6385
Step 3: Add and remove fax devices ... 6386
Step 4: Configure fax devices to send and receive ... 6386
Step 5: Configure fax settings ... 6387
Configure incoming fax settings ... 6388
Start or stop fax reception ... 6388
Configure incoming fax routing ... 6388
Configure SMTP settings for e-mail routing ... 6389
Configuring outgoing fax settings ... 6390
Start or stop outgoing faxes ... 6390
Configure the outbox and outgoing transmissions ... 6390
Configure fax delivery receipts ... 6392
Configure outgoing routing ... 6393
Create a fax device group ... 6393
Add or remove a fax device in a group ... 6393
Add or delete a fax routing rule ... 6394
Step 6: Set up user accounts ... 6395
User accounts ... 6395
Additional Resources ... 6395

Configure Print and Document Services ... 6396
Step 1: Install v4 drivers ... 6396
Step 2: Install v3 drivers (if necessary) ... 6398

Step 3: Create a shared print queue 6399
Step 4: Connect to the print queue 6399
Step 5: Print from Windows 6400
Step 6: Print from a Windows app 6403
See also 6405

Create Custom Separator Pages in Windows Server 2012 6405
Create a custom separator page file 6406
See also 6408

Assign Delegated Print Administrator and Printer Permission Settings in Windows Server 2012 6408
Configuring security settings 6409
The print server security user interface 6409
Setting permissions in Print Server Properties 6410
Creating a delegated print administrator 6412
Print-related permissions and the tasks they enable 6413
Designing and creating print security groups 6415

Branch Office Direct Printing Overview 6419
Feature description 6419
Practical applications 6419
Important functionality 6419
Software requirements 6420
See also 6420

Configure Branch Office Direct Printing 6420
Prerequisites 6421
Step 1: Configure Branch Office Direct Printing 6421
See also 6421

Branch Office Direct Printing Technical Details 6422
Overview 6422
Benefits of Branch Office Direct Printing 6422
Limitations of Branch Office Direct Printing 6422
Architecture 6423
Implementation Details 6424
Configuration 6424
Setup 6425
Synchronization 6425
Branch Office Status 6426
Synchronization Before Printing 6426
Fallback 6427
Branch Office Direct Printing Event IDs 6427

Branch Office Direct Print (BODP) Logging in Windows Server 2012 R2 6428
Branch Office Direct Print (BODP) Logging Overview 6428
Branch Office Direct Print (BODP) Logging Configuration 6429
Branch Office Direct Print (BODP) Logging Limitations 6429
Branch Office Direct Print (BODP) Logging Diagram 6430

Print and Document Services Architecture 6430
Printer driver overview 6430
V4 printer drivers 6431
V3 printer drivers 6431
Changes to the printer driver model 6431
Reducing driver size 6432
Driver isolation 6434
Application isolation 6436
V4 driver model design and architecture 6437
V4 driver design 6438
V4 driver architecture 6438
Printer sharing overview 6440
Architecture 6441
Appendix A: Terms and Definitions 6444

Printer Sharing Technical Details 6445
Types of printer sharing in Windows Server 2012 6446
Enhanced Point and Print 6447
Package Aware Point and Print 6449
Legacy Point and Print 6450
Internet Print Protocol (IPP) 6451
Technologies to Augment Printer Sharing 6452
Branch Office Direct Printing 6452
WSD Secure Printing 6452
AD Printer Publishing 6453
Push Printer Connections 6454
Deployment Scenarios 6455
Legacy and Package Aware Point and Print 6455
IPP 6455
Using v3 print drivers for IPP 6455
Using v4 print drivers on the Server 6455
Manually Creating an IPP Connection in the User Interface 6456
Enhanced Point and Print 6458
Basic Deployment of enhanced Point and Print shares 6459
Custom User Interface for enhanced Point and Print Shares 6460
Deploying Desktop Printer Extension Apps as an Application 6460
Deploying Desktop Printer Extension Apps in a v4 Driver 6460
Deploying Desktop Printer Extension Apps in Print Preferences 6460

Blocking Desktop Printer Extensions ... 6461
Deploying Windows Store Device Apps ... 6461
Enabling CSR for Client Computers Running Windows 8 ... 6461
Deploying v4 drivers with WSUS ... 6462
Deploying v4 drivers with DevicePath ... 6463
Deploying v4 drivers with standard deployment tools ... 6463
Windows on ARM in Enterprises ... 6463
Connecting Directly to the Printer ... 6463
Connecting to v4 Shares ... 6464
Connecting to v3 Shares ... 6464
Connecting to IPP Shares ... 6464
Branch Office Printing ... 6465
WSD Secure Printing ... 6465
Requirements for deploying a WSD Secure Printer ... 6465
Steps to deploy a WSD Secure Printer ... 6465
Notes on WSD Secure Print scenarios ... 6466

Migrate Print and Document Services to Windows Server 2012 ... 6466
Overview ... 6466
About this guide ... 6468
Target audience ... 6468
What this guide does not provide ... 6468
Supported migration scenarios ... 6468
Supported operating systems ... 6468
Supported role configurations ... 6470
Supported role services and features ... 6470
Migrating from x86-based to x64-based v3 printer drivers ... 6470
Unsupported scenarios ... 6471
Print and Document Services migration overview ... 6471
Migrate print servers (overview) ... 6472
Impact of migration ... 6472
Impact of migration on the source server ... 6472
Impact of migration on other computers in the enterprise ... 6472
Permissions required to complete migration ... 6473
Permissions required to complete migration on other computers in the enterprise ... 6473
Estimated duration ... 6473
See Also ... 6473

Preparing to Migrate ... 6474
Access the migration tools ... 6474
To access the Printer Migration Wizard ... 6474
To access the Printbrm.exe command-line tool ... 6475
Prepare the destination server ... 6475
Hardware requirements for the destination server ... 6475

Software requirements for the destination server ... 6475
Installing the Print and Document Services role on the destination server ... 6476
Preparing for cross-architecture migrations ... 6476
Preparing for additional scenarios ... 6476
Prepare the source server ... 6477
See Also ... 6478

Migrating the Print and Document Services Role ... 6478
Back up the source server ... 6478
Cross-architecture migrations ... 6480
Restoration ... 6480
See Also ... 6481

Verifying the Migration ... 6482
Verify the migration ... 6482
To verify destination server configuration ... 6482
Rename the destination server to the name of the source server ... 6483
To verify configuration of other computers in the enterprise ... 6483
Print a test job from a client with an existing connection ... 6484
See Also ... 6484

Post-Migration Tasks ... 6484
Post-migration ... 6484
Success ... 6484
Retire the source server ... 6484
Failure ... 6485
Restoring the role in the event of migration failure ... 6485
Rollback requirements ... 6485
Estimated time to complete rollback ... 6485
Roll back migration on the source server ... 6486
Roll back migration on the destination server ... 6486
Troubleshooting ... 6486
Log file locations ... 6486
Migrating cross-platform driver language monitors ... 6486
Mitigating a failure in the Print Spooler service ... 6486
Additional references ... 6487
See Also ... 6487

Appendix A - Printbrm.exe Command-Line Tool Details ... 6487
Printbrm.exe command-line tool syntax ... 6487
Printbrm enhancements ... 6488
Printbrm usage scenarios ... 6489
Using the configuration file ... 6489
Selectively restoring your printers ... 6490

Moving printers to a different domain ... 6490
See Also ... 6491

Appendix B - Additional Destination Server Scenarios ... 6491
If your server hosts Line Printer Remote (LPR) printers ... 6491
If your server offers Internet Printing Protocol (IPP) printer connections ... 6492
If your server hosts Web Services on Devices (WSD) printers ... 6492
If your print server is a highly available virtual machine ... 6492
If your server hosts local bus printers (LPT and USB) ... 6492
If your server hosts plug and play printers ... 6492
See Also ... 6493

Appendix C - Printbrm Event IDs ... 6493
Printbrm Event IDs ... 6493
See Also ... 6506

Print Server Scalability and Capacity Planning ... 6507
What this guide includes ... 6507
Factors that influence Print Server performance ... 6507
Diagnostic and performance analysis ... 6511
Reference systems case study ... 6517
Recommendations for deployment ... 6518
Print Server performance data ... 6520
See also ... 6520

High Availability Printing Overview ... 6521
Feature description ... 6521
Practical applications ... 6522
Important functionality ... 6522
Hardware requirements ... 6523
Software requirements ... 6523
See also ... 6524

Install and Configure High Availability Printing ... 6524
Configure and manage high availability printing ... 6524
Step 1: Configure the firewall ... 6525
Step 2: Configure Virtual Machine Monitoring ... 6525
Step 3: Verify the Virtual Machine Monitoring configuration ... 6529
Step 4: Testing Virtual Machine Monitoring ... 6530
See Also ... 6530

Client Name Instrumentation in Windows Server 2012 R2 ... 6530
Client Name Instrumentation Overview ... 6530
Client Name Instrumentation Limitations ... 6530
Client Name Instrumentation UAL Cmdlets ... 6531

Remote Desktop Services Overview ... 6531
Role description ... 6531
Practical applications ... 6531
New and changed functionality for Windows Server 2012 ... 6532
New and changed functionality for Windows Server 2012 R2 ... 6532
New Microsoft Remote Desktop Clients ... 6533
Removed or deprecated functionality ... 6533
Hardware requirements ... 6533
Server Manager information ... 6534
See also ... 6535

What's New in Remote Desktop Services in Windows Server 2012 R2 ... 6537
Session Shadowing ... 6538
Online Data Deduplication ... 6538
Improved RemoteApp behavior ... 6538
Quick reconnect for remote desktop clients ... 6538
Improved compression and bandwidth usage ... 6538
Dynamic display handling ... 6539
RemoteFX virtualized GPU supports DX11.1 ... 6539
RestrictedAdmin Mode Remote Desktop ... 6539
Remote Desktop Services role services description ... 6540
Removed or deprecated functionality ... 6541
See also ... 6541

What's New in Remote Desktop Services in Windows Server 2012 ... 6543
Virtual Desktop Infrastructure (VDI) deployment ... 6543
Session Virtualization deployment ... 6543
Centralized resource publishing ... 6545
Rich user experience with Remote Desktop Protocol (RDP) ... 6545
Rich graphics experience with RemoteFX vGPU ... 6546
Remote Desktop Services role services description ... 6546
Removed or deprecated functionality ... 6548
See also ... 6548

Test Lab Guide: Virtual Desktop Infrastructure Quick Start ... 6549
In this guide ... 6549
Test lab overview ... 6550
Hardware and software requirements ... 6551
Steps for configuring the test lab ... 6551
Step 1: Complete the base configuration test lab ... 6552
Step 2: Configure APP1 ... 6552
Step 3: Test Remote Desktop Services connectivity ... 6552
Step 4: Take a snapshot of the configuration ... 6553
Additional Resources ... 6553

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment 6554
In this guide 6554
Test lab overview 6554
Hardware and software requirements 6555
Steps for configuring the test lab 6556
Step 1: Complete the base configuration test lab 6556
Step 2: Install and Configure RDVH1 6556
Install the operating system on RDVH1 6557
Configure TCP/IP properties 6557
Join RDVH1 to the CORP domain 6558
Step 3: Install and Configure RDWA1 6558
Install the operating system on RDWA1 6558
Configure TCP/IP properties 6559
Join RDWA1 to the CORP domain 6560
Step 4: Install and Configure RDCB1 6560
Install the operating system on RDCB1 6560
Configure TCP/IP properties 6561
Join RDCB1 to the CORP domain 6562
Step 5: Deploy the Virtual Desktop Infrastructure standard deployment 6562
Step 6: Test the VDI standard deployment connectivity 6563
Step 7: Take a snapshot of the configuration 6564
Additional Resources 6564

Test Lab Guide: Managed Pooled Virtual Desktop Collections 6564
In this guide 6565
Test lab overview 6565
Hardware and software requirements 6566
Steps for configuring the test lab 6567
Step 1: Complete the VDI standard deployment test lab 6567
Step 2: Configure RDCB1 6567
Step 3: Configure RDVH1 6568
Step 4: Create the managed pooled virtual desktop collection 6569
Step 5: Test Remote Desktop Services connectivity 6570
Step 6: Take a snapshot of the configuration 6570
Additional Resources 6571

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections 6571
In this guide 6571
Test lab overview 6572
Hardware and software requirements 6573
Steps for configuring the test lab 6573
Step 1: Complete the VDI standard deployment test lab 6574
Step 2: Configure RDCB1 6574
Step 3: Create the unmanaged pooled virtual desktop collection 6575

Step 4: Test Remote Desktop Services connectivity 6576
Step 5: Take a snapshot of the configuration 6576
Additional Resources 6577

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start 6577
Introduction 6577
In this guide 6578
Test lab overview 6578
Hardware and software requirements 6579
Steps for configuring the test lab 6579
Step 1: Complete the Base Configuration Test Lab 6580
Step 2: Configure APP1 6580
Deploy the Session Virtualization Quick Start 6580
Step 3: Test Remote Desktop Service Connectivity 6581
Connect to the Session Collection 6581
Step 4: Snapshot the Configuration 6581
Additional Resources 6582

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment 6582
In this guide 6583
Test lab overview 6583
Hardware and software requirements 6584
Steps for configuring the test lab 6584
Step 1: Complete the base configuration test lab 6585
Step 2: Install and Configure RDSH1 6585
Install the operating system on RDSH1 6585
Configure TCP/IP properties 6586
Join RDSH1 to the CORP domain 6587
Step 3: Install and Configure RDWA1 6587
Install the operating system on RDWA1 6587
Configure TCP/IP properties 6588
Join RDWA1 to the CORP domain 6589
Step 4: Install and Configure RDCB1 6589
Install the operating system on RDCB1 6589
Configure TCP/IP properties 6590
Join RDCB1 to the CORP domain 6591
Step 5: Deploy Session Virtualization Standard 6591
Add RDSH1, RDWA1, and RDCB1 to the server pool 6591
Deploy the Session Virtualization Standard deployment 6592
Step 6: Create Session Collection 6592
Create a session collection 6593
Step 7: Test Remote Desktop Services Connectivity 6593
Connect to the session collection 6593
Step 8: Snapshot the Configuration 6594

Additional Resources 6594

Test Lab Guide: Remote Desktop Services Publishing 6595
Introduction 6595
In this guide 6595
Test lab overview 6595
Hardware and software requirements 6596
Steps for configuring the test lab 6597
Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab 6597
Step 2: Publish a RemoteApp Program 6598
Publish a RemoteApp program 6598
Step 3: Enable File Type Association on a RemoteApp Program 6598
Enable the default connection URL using Group Policy 6598
Configure the File Type Association 6599
Step 4: Configure DNS Feed Lookup for RemoteApp and Desktop Connections 6599
Configure DNS Feed Lookup for RemoteApp and Desktop Connections. 6599
Step 5: Test Published RemoteApp Program using RD Web Access 6600
Test Published RemoteApp Program using RD Web Access 6600
Step 6: Snapshot the Configuration 6600
Additional Resources 6601

Test Lab Guide: Remote Desktop Services Licensing 6601
Introduction 6601
In this guide 6601
Test lab overview 6602
Hardware and software requirements 6603
Steps for configuring the test lab 6603
Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab 6604
Step 2: Install the RD Licensing role service 6604
Install the RD Licensing role service 6605
Step 3: Activate the Remote Desktop License Server 6605
Activate the Remote Desktop license server 6605
Step 4: Request RDS Beta CALs 6606
Request RDS Beta CALs 6606
Step 5: Install RDS CALs on the Remote Desktop License Server 6606
Install RDS CALs on the Remote Desktop license server 6606
Step 6: Verify Licensing for the Remote Desktop Services 6607
Establish a client connection in a session-based desktop deployment using a valid license 6607
Establish a client connection in a virtual machine-based desktop deployment using a valid license 6607
Verify the RDS CAL by using Remote Desktop Licensing Manager 6608

Step 7: Snapshot the Configuration ... 6608
Additional Resources ... 6608

Microsoft Remote Desktop Clients ... 6608
Remote Desktop Client for Android ... 6609
Remote Desktop Client for iOS ... 6609
Remote Desktop Client for Mac ... 6609
Remote Desktop Client for Windows ... 6609
See also ... 6609

Getting Started with Remote Desktop Client on Android ... 6611
Help contents for the Microsoft Remote Desktop Client on Android ... 6611
Here's how to get started with the Remote Desktop Client: ... 6612
Getting to know the Connection Center ... 6612
Create / Edit Remote Desktops ... 6612
Manage Favorites ... 6614
Search Remote Desktops ... 6614
Manage Settings ... 6614
Remote Resources ... 6615
Remote Desktop Gateway ... 6616
Certificates ... 6616
Trusted Hosts ... 6617
Navigating the Remote Desktop Session ... 6617
Add a Remote Desktop Widget to the Home Screen ... 6619
See also ... 6619

Getting Started with Remote Desktop Client on iOS ... 6619
Help contents for the Microsoft Remote Desktop Client on iOS ... 6620
Here's how to get started with the Remote Desktop Client: ... 6620
Getting to know the Connection Center ... 6620
Create / Edit Remote Desktops ... 6621
Search Remote Desktops ... 6622
Manage Configuration ... 6622
Certificates ... 6623
Gateway ... 6623
Remote Desktop licensing ... 6624
Remote Resources ... 6624
Navigating the Remote Desktop Session ... 6625
See also ... 6627

Getting Started with Remote Desktop Client on Mac ... 6627
Help contents for the Microsoft Remote Desktop Client on Mac ... 6627
Here's how to get started with the Remote Desktop Client: ... 6628
Getting to know the Connection Center ... 6628

Create / Edit Remote Desktops ... 6628
Export Remote Desktops ... 6630
Import Remote Desktops ... 6630
Search Remote Desktops ... 6630
Rearrange Remote Desktops ... 6631
Manage Preferences ... 6631
Remote Desktop Gateway ... 6631
Resolution ... 6632
Remote Resources ... 6632
Navigating the Remote Desktop Session ... 6633
See also ... 6633

Remote Desktop Client on Android: FAQ ... 6634
Remote Desktop Client: FAQ ... 6634
How do I set up a PC for Remote Desktop Connection? ... 6634
Which PCs can I connect to using Remote Desktop Connection? ... 6635
Why can’t I connect using Remote Desktop Connection? ... 6636
Why can’t I sign in to a remote PC? ... 6637
Why are the characters mixed up in the session? ... 6637
I cannot find or connect to my computer ... 6637
VPN does not work ... 6638
Which connection methods are supported for company networks? ... 6638
Is bi-directional sound supported? ... 6638
Is the Remote Desktop Client compatible with RDP 7.1? ... 6638
How to configure L2TP or PPTP VPN connections ... 6638
Error: Insufficient privileges ... 6639
Error: Failed to parse NTLM challenge ... 6639
What can I do if the sound does not play? ... 6639
Unsupported Windows Versions and Editions ... 6639
RD Gateway messaging is not supported ... 6639
ERROR: TS_RAP "You are not allowed to connect to the given host ..." ... 6639
Error message that there is no CAL ... 6640
Error: Access Denied ... 6640
How can I test if VPN is working properly? ... 6640
Error: RPC Error 23014 or Error 0x59e6 ... 6640
See also ... 6640

Remote Desktop Client on iOS: FAQ ... 6641
Remote Desktop Client: FAQ ... 6641
How do I set up a PC for Remote Desktop Connection? ... 6641
Which PCs can I connect to using Remote Desktop Connection? ... 6642
Why can’t I connect using Remote Desktop Connection? ... 6643
Why can’t I sign in to a remote PC? ... 6644
Error: Insufficient privileges ... 6644

How to configure L2TP or PPTP VPN connections ... 6644
What do I need to connect to an office computer using Remote Desktop Gateway? ... 6644
I cannot find or connect to my computer ... 6644
VPN does not work ... 6645
What can I do when the app does not start anymore? ... 6645
Which connection methods are supported for company networks? ... 6645
Error: STOP error in Windows XP ... 6645
Is AirPrint supported? ... 6645
Is bi-directional sound supported? ... 6646
Why are the characters mixed up in the session? ... 6646
Why don’t special keys work on a Bluetooth keyboard? ... 6646
Is the Remote Desktop Client compatible with RDP 7.1? ... 6646
Error: Failed to parse NTLM challenge ... 6646
What can I do if the sound does not play? ... 6646
Unsupported Windows Versions and Editions ... 6647
RD Gateway messaging is not supported ... 6647
ERROR: TS_RAP "You are not allowed to connect to the given host ..." ... 6647
Error message that there is no CAL ... 6647
Error: Access Denied ... 6648
How can I test if VPN is working properly? ... 6648
Error: RPC Error 23014 or Error 0x59e6 ... 6648
See also ... 6648

Remote Desktop Client on Mac: FAQ ... 6648
Remote Desktop Client: FAQ ... 6649
How do I set up a PC for Remote Desktop Connection? ... 6649
Which PCs can I connect to using Remote Desktop Connection? ... 6650
Why can’t I connect using Remote Desktop Connection? ... 6651
Why can’t I sign in to a remote PC? ... 6652
Is retina resolution supported? ... 6652
How do I enable secondary right-click on your Mac? ... 6652
How do I make use of all of my monitors? ... 6652
How do I type various symbols on a Mac keyboard? ... 6652
I cannot find or connect to my computer ... 6652
VPN does not work ... 6653
Which connection methods are supported for company networks? ... 6653
Is bi-directional sound supported? ... 6653
Why are the characters mixed up in the session? ... 6653
Is the Remote Desktop Client compatible with RDP 7.1? ... 6654
How to configure L2TP or PPTP VPN connections ... 6654
Error: Insufficient privileges ... 6654
Error: Failed to parse NTLM challenge ... 6654
What can I do if the sound does not play? ... 6654

Unsupported Windows Versions and Editions ... 6654
RD Gateway messaging is not supported ... 6655
ERROR: TS_RAP "You are not allowed to connect to the given host ..." ... 6655
Error message that there is no CAL ... 6655
Error: Access Denied ... 6655
How can I test if VPN is working properly? ... 6655
Error: RPC Error 23014 or Error 0x59e6 ... 6656
See also ... 6656

Security and Protection ... 6656
See also ... 6664

Access Control Overview ... 6665
Feature description ... 6665
Practical applications ... 6665
Permissions ... 6666
Ownership of objects ... 6667
Inheritance of permissions ... 6667
User rights ... 6667
Object auditing ... 6667
Additional resources ... 6667

Dynamic Access Control Overview ... 6668
Central access rules ... 6668
Central access policies ... 6668
Claims ... 6669
Expressions ... 6669
Proposed permissions ... 6669
Additional changes ... 6670
Support in the Kerberos authentication protocol to reliably provide user claims, device claims, and device groups. ... 6670
Support for using the Key Distribution Center (KDC) Group Policy setting to enable Dynamic Access Control for a domain. ... 6670
Support for using the Key Distribution Center (KDC) Group Policy setting to enable Dynamic Access Control for a domain. ... 6670
Support in Active Directory to store user and device claims, resource properties, and central access policy objects. ... 6670
Support for using Group Policy to deploy central access policy objects. ... 6670
Support for claims-based file authorization and auditing for file systems by using Group Policy and Global Object Access Auditing ... 6671
Support for transforming or filtering claim policy objects that traverse Active Directory forest trusts ... 6671
Software requirements ... 6671
Additional resource ... 6672

Security Principals Technical Overview ... 6672
What are security principals? ... 6672
How security principals work ... 6673
Authorization and access control components ... 6673
Security identifiers ... 6674
Access tokens ... 6674
Security descriptors and access control lists ... 6675
Permissions ... 6675
Security context in authentication ... 6676
Accounts and security groups ... 6676
User accounts ... 6676
Security groups ... 6677

Service Accounts ... 6678
Overview ... 6678
Standalone managed service accounts ... 6678
Software requirements ... 6679
Group managed service accounts ... 6679
Practical applications ... 6679
Software requirements ... 6680
Virtual accounts ... 6680
Software requirements ... 6680
See also ... 6681

Special Identities ... 6681
Anonymous Logon ... 6682
Authenticated Users ... 6683
Batch ... 6683
Creator Group ... 6684
Creator Owner ... 6684
Dialup ... 6685
Digest Authentication ... 6685
Enterprise Domain Controllers ... 6685
Everyone ... 6686
Interactive ... 6686
Local Service ... 6687
LocalSystem ... 6688
Network ... 6688
Network Service ... 6688
NTLM Authentication ... 6689
Other Organization ... 6689
Principal Self ... 6690
Remote Interactive Logon ... 6690
Restricted ... 6691

SChannel Authentication ... 6691
Service ... 6691
Terminal Server User ... 6692
This Organization ... 6692
Window Manager\Window Manager Group ... 6692
See also ... 6693

Microsoft Accounts ... 6693
How a Microsoft account works ... 6693
How Microsoft accounts are created ... 6694
How the Microsoft account information is safeguarded ... 6694
The Microsoft account in the enterprise ... 6695
Managing the Microsoft account in the domain ... 6696
Restrict the use of the Microsoft account ... 6696
Configure connected accounts ... 6696
Provision Microsoft accounts in the enterprise ... 6697
Audit account activity ... 6697
Perform password resets ... 6697
Restrict app installation and usage ... 6697
See also ... 6697

Active Directory Security Groups ... 6697
About Active Directory groups ... 6698
Distribution groups ... 6698
Security groups ... 6698
Group scope ... 6699
Special identity groups ... 6701
Default security groups ... 6701
Active Directory default security groups by operating system version ... 6702
Access Control Assistance Operators ... 6704
Account Operators ... 6705
Administrators ... 6706
Allowed RODC Password Replication Group ... 6708
Backup Operators ... 6709
Certificate Service DCOM Access ... 6709
Cert Publishers ... 6710
Cloneable Domain Controllers ... 6711
Cryptographic Operators ... 6711
Denied RODC Password Replication Group ... 6712
Distributed COM Users ... 6713
DnsUpdateProxy ... 6714
DnsAdmins ... 6714
Domain Admins ... 6715
Domain Computers ... 6716

Domain Controllers ... 6717
Domain Guests ... 6717
Domain Users ... 6718
Enterprise Admins ... 6719
Enterprise Read-Only Domain Controllers ... 6720
Event Log Readers ... 6721
Group Policy Creators Owners ... 6721
Guests ... 6722
Hyper-V Administrators ... 6723
IIS_IUSRS ... 6724
Incoming Forest Trust Builders ... 6724
Network Configuration Operators ... 6725
Performance Log Users ... 6726
Performance Monitor Users ... 6727
Pre–Windows 2000 Compatible Access ... 6728
Print Operators ... 6729
Protected Users ... 6730
RAS and IAS Servers ... 6732
RDS Endpoint Servers ... 6732
RDS Management Servers ... 6733
RDS Remote Access Servers ... 6734
Remote Desktop Users ... 6734
Read-Only Domain Controllers ... 6735
Remote Management Users ... 6736
Replicator ... 6737
Schema Admins ... 6738
Server Operators ... 6739
Terminal Server License Servers ... 6740
Users ... 6741
Windows Authorization Access Group ... 6741
WinRMRemoteWMIUsers_ ... 6742
See also ... 6743

AppLocker Overview ... 6744
What does AppLocker do? ... 6744
When to use AppLocker ... 6745
Versions, interoperability, and differences in functionality ... 6746
System requirements ... 6747
Installing AppLocker ... 6748
Using AppLocker on Server Core ... 6748
Virtualization considerations ... 6748
Security considerations ... 6748
Maintaining AppLocker policies ... 6749

See also 6749

Administer AppLocker 6750
Using the MMC snap-ins to administer AppLocker 6752
Administer Applocker using Group Policy 6752
Administer AppLocker on the local computer 6752
Using Windows PowerShell to administer AppLocker 6752

Maintain AppLocker Policies 6752
Maintaining AppLocker policies by using Group Policy 6753
Step 1: Understand the current behavior of the policy 6753
Step 2: Export the AppLocker policy from the GPO 6754
Step 3: Update the AppLocker policy by editing the appropriate AppLocker rule 6754
Step 4: Test the AppLocker policy 6754
Step 5: Import the AppLocker policy into the GPO 6754
Step 6: Monitor the resulting policy behavior 6754
Maintaining AppLocker policies by using the Local Security Policy snap-in 6755
Step 1: Understand the current behavior of the policy 6755
Step 2: Update the AppLocker policy by modifying the appropriate AppLocker rule 6755
Step 3: Test the AppLocker policy 6755
Step 4: Deploy the policy with the modified rule 6755
Step 5: Monitor the resulting policy behavior 6755
Additional resources 6756
See Also 6756

Edit an AppLocker Policy 6756
Editing an AppLocker policy by using Group Policy 6756
Step 1: Use Group Policy management software to export the AppLocker policy from the GPO 6756
Step 2: Import the AppLocker policy into the AppLocker reference computer or the computer you use for policy maintenance 6757
Step 3: Use AppLocker to modify and test the rule 6757
Step 4: Use AppLocker and Group Policy to import the AppLocker policy back into the GPO 6757
Editing an AppLocker policy by using the Local Security Policy snap-in 6758
Step 1: Import the AppLocker policy 6758
Step 2: Identify and modify the rule to change, delete, or add 6758
Step 3: Test the effect of the policy 6758
Step 4: Export the policy to an XML file and propagate it to all targeted computers 6759
Additional resources 6759
See Also 6759

Test and Update an AppLocker Policy 6759
Step 1: Enable the Audit only enforcement setting 6759

Step 2: Configure the Application Identity service to start automatically 6759
Step 3: Test the policy 6760
Step 4: Analyze AppLocker events 6760
Step 5: Modify the AppLocker policy 6760
Step 6: Repeat policy testing, analysis, and policy modification 6761
Additional resources 6761
See Also 6761

Deploy AppLocker Policies by Using the Enforce Rules Setting 6761
Background and prerequisites 6761
Step 1: Retrieve the AppLocker policy 6761
Step 2: Alter the enforcement setting 6762
Step 3: Update the policy 6762
Step 4: Monitor the effect of the policy 6762
Additional resources 6762
See Also 6762

Use the AppLocker Windows PowerShell Cmdlets 6763
AppLocker Windows PowerShell cmdlets 6763
Import the AppLocker PowerShell cmdlet module 6763
Retrieve application information 6763
Set AppLocker policy 6763
Retrieve an AppLocker policy 6763
Generate rules for a given user or group 6764
Test the AppLocker Policy against a file set 6764
Additional resources 6764
See Also 6764

Use AppLocker and Software Restriction Policies in the Same Domain 6764
Using AppLocker and Software Restriction Policies in the same domain 6764
See Also 6768

Optimize AppLocker Performance 6768
Optimization of Group Policy 6768
AppLocker rule limitations 6768
Using the DLL rule collection 6769
See Also 6769

Monitor Application Usage with AppLocker 6769
Discover the effect of an AppLocker policy 6769
Review AppLocker events with Get-AppLockerFileInformation 6770
View the AppLocker Log in Event Viewer 6770
See Also 6771

Manage Packaged Apps with AppLocker 6771

Understanding Packaged apps and Packaged app installers for AppLocker 6771
Comparing classic apps and Packaged apps 6772
Design and deployment decisions 6772
Using AppLocker to manage Packaged apps 6773
See Also 6774

Working with AppLocker Rules 6774
Procedures 6774
Rule collections 6775
Rule conditions 6776
Publisher 6776
Path 6778
File hash 6779
AppLocker default rules 6779
AppLocker rule behavior 6779
Rule exceptions 6780
DLL rule collection 6780
AppLocker wizards 6781
Additional considerations 6781
See Also 6782

Create a Rule That Uses a File Hash Condition 6782

Create a Rule That Uses a Path Condition 6783

Create a Rule That Uses a Publisher Condition 6784

Create AppLocker Default Rules 6785

Configure Exceptions for an AppLocker Rule 6786

Create a Rule for Packaged Apps 6787

Delete an AppLocker Rule 6790

Edit AppLocker Rules 6790

Enable the DLL Rule Collection 6792

Enforce AppLocker Rules 6792

Run the Automatically Generate Rules Wizard 6793

Working with AppLocker Policies 6794
What does AppLocker do? 6795
See Also 6795

Configure the Application Identity Service 6795

Configure an AppLocker Policy for Audit Only 6796
Configure an AppLocker Policy for Enforce Rules 6797
Display a Custom URL Message When Users Try to Run a Blocked Application 6798
Export an AppLocker Policy from a GPO 6798
Export an AppLocker Policy to an XML File 6799
Import an AppLocker Policy from Another Computer 6800
Import an AppLocker Policy into a GPO 6800
Add Rules for Packaged Apps to Existing AppLocker Rule-set 6801
  See Also 6801
Merge AppLocker Policies by Using Set-ApplockerPolicy 6801
  Example 6802
    Description 6802
    Code 6802
Merge AppLocker Policies Manually 6802
Refresh an AppLocker Policy 6804
Test an AppLocker Policy by Using Test-AppLockerPolicy 6805
Packaged Apps and Packaged App Installer Rules in AppLocker 6806
  See also 6807
AppLocker Policies Design Guide 6808
  Purpose of this guide 6808
  Contents of this guide 6808
Understand the AppLocker Policy Deployment Process 6809
  Resources to support the deployment process 6811
Understand AppLocker Policy Design Decisions 6811
    Which applications do you need to control in your organization? 6812
    Comparing classic desktop applications and Windows Store apps for AppLocker policy design decisions 6814
    How do you currently control application usage in your organization? 6815
    Which Windows desktop and server operating systems are running in your organization? 6816
    Are there specific groups in your organization that need customized application control policies? 6817

Does your IT department have resources to analyze application usage, and to design and manage the policies? 6817
Does your organization have Help Desk support? 6818
Do you know what applications require restrictive policies? 6818
How do you deploy or sanction applications (upgraded or new) in your organization? 6819
Does your organization already have SRP deployed? 6819
What are your organization's priorities when implementing application control policies? .. 6820
How are applications currently accessed in your organization? 6821
Is the structure in Active Directory Domain Services based on the organization's hierarchy? 6822
Record your findings 6822

Determine Your Application Control Objectives 6823

Create List of Applications Deployed to Each Business Group 6827
Determining application usage 6827
How to perform the application usage assessment 6827
Prerequisites to completing the inventory 6828
Next steps 6828

Document Your Application List 6829
Record your findings 6829
Next steps 6830

Select Types of Rules to Create 6830
Select the rule collection 6831
Determine the rule condition 6831
Determine how to allow system files to run 6832
Next steps 6833

Document Your AppLocker Rules 6833
Record your findings 6833
Next steps 6835

Determine Group Policy Structure and Rule Enforcement 6835

Understand AppLocker Enforcement Settings 6836

Understand AppLocker Rules and Enforcement Setting Inheritance in Group Policy 6837

Document Group Policy Structure and AppLocker Rule Enforcement 6839
Record your findings 6839
Next steps 6840

Plan for AppLocker Policy Management 6841
Policy management 6841

Application and user support policy 6841
Policy maintenance 6843
Next steps 6844

Document Your Application Control Management Processes 6844
Record your findings 6844
Next steps 6847

Create Your AppLocker Planning Document 6847
The AppLocker deployment design 6848
AppLocker planning document contents 6848
Sample template for an AppLocker planning document 6848
Example of an AppLocker planning document 6849
Additional resources 6851

AppLocker Policies Deployment Guide 6852
Purpose of this guide 6852
Prerequisites to deploying AppLocker policies 6852
Contents of this guide 6853
Additional resources 6853

Understand the AppLocker Policy Deployment Process 6853
Resources to support the deployment process 6856

Requirements for Deploying AppLocker Policies 6856
Deployment plan 6856
Supported operating systems 6858
Policy distribution mechanism 6859
Event collection and analysis system 6859
See also 6859

Using Software Restriction Policies and AppLocker Policies 6859
Understand the difference between SRP and AppLocker 6859
Use SRP and AppLocker in the same domain 6859
Test and validate SRPs and AppLocker policies that are deployed in the same environment 6861
Step 1: Test the effect of SRPs 6861
Step 2: Test the effect of AppLocker policies 6861
See also 6861

Create Your AppLocker Policies 6861
Step 1: Use your plan 6862
Step 2: Create your rules and rule collections 6862
Step 3: Configure the enforcement setting 6862
Step 4: Update the GPO 6862

Step 5: Test the effect of the policy ... 6863
Step 6: Implement the policy ... 6863
Step 7: Test the effect of the policy and adjust ... 6863
Next steps ... 6863
See also ... 6863

Create Your AppLocker Rules ... 6863
Creating AppLocker rules ... 6863
Automatically generate your rules ... 6864
Create your rules individually ... 6864
About selecting rules ... 6864
Next steps ... 6865
See Also ... 6865

Test and Update an AppLocker Policy ... 6865
Step 1: Enable the Audit only enforcement setting ... 6865
Step 2: Configure the Application Identity service to start automatically ... 6866
Step 3: Test the policy ... 6866
Step 4: Analyze AppLocker events ... 6866
Step 5: Modify the AppLocker policy ... 6867
Step 6: Repeat policy testing, analysis, and policy modification ... 6867
Additional resources ... 6867
See Also ... 6867

Deploy the AppLocker Policy into Production ... 6867
Understand your design decisions ... 6867
AppLocker deployment methods ... 6868
See also ... 6868

Use a Reference Computer to Create and Maintain AppLocker Policies ... 6868
Background and prerequisites ... 6868
Step 1: Automatically generate rules on the reference computer ... 6869
Step 2: Create the default rules on the reference computer ... 6869
Step 3: Modify rules and the rule collection on the reference computer ... 6870
Step 4: Test and update AppLocker policy on the reference computer ... 6870
Step 5: Export and import the policy into production ... 6871
Step 6: Monitor the effect of the policy in production ... 6871
See also ... 6871

Determine Which Applications Are Digitally Signed on a Reference Computer ... 6871
See Also ... 6872

Configure the AppLocker Reference Computer ... 6872
See also ... 6873

Maintain AppLocker Policies 6873
Maintaining AppLocker policies by using Group Policy 6874
Step 1: Understand the current behavior of the policy 6874
Step 2: Export the AppLocker policy from the GPO 6875
Step 3: Update the AppLocker policy by editing the appropriate AppLocker rule 6875
Step 4: Test the AppLocker policy 6875
Step 5: Import the AppLocker policy into the GPO 6875
Step 6: Monitor the resulting policy behavior 6875
Maintaining AppLocker policies by using the Local Security Policy snap-in 6875
Step 1: Understand the current behavior of the policy 6876
Step 2: Update the AppLocker policy by modifying the appropriate AppLocker rule 6876
Step 3: Test the AppLocker policy 6876
Step 4: Deploy the policy with the modified rule 6876
Step 5: Monitor the resulting policy behavior 6876
Additional resources 6876
See Also 6877

AppLocker Technical Reference 6877

What Is AppLocker? 6878
What features are different between Software Restriction Policies and AppLocker? 6879
See Also 6881

Requirements to Use AppLocker 6881
General requirements 6881
Operating system requirements 6882
See also 6884

AppLocker Policy Use Scenarios 6884
Use scenarios 6885
See also 6885
See Also 6885

How AppLocker Works 6886
Additional resources 6886
See Also 6886

Understanding AppLocker Rule Behavior 6886
See Also 6887

Understanding AppLocker Rule Exceptions 6887
See Also 6887

Understanding AppLocker Rule Collections 6888
See Also 6888

Understanding AppLocker Allow and Deny Actions on Rules ... 6888
Allow action versus deny action on rules ... 6888
Deny rule considerations ... 6889
See Also ... 6889

Understanding AppLocker Rule Condition Types ... 6889
Considerations ... 6890
See Also ... 6891

Understanding the Publisher Rule Condition in AppLocker ... 6891
See Also ... 6893

Understanding the Path Rule Condition in AppLocker ... 6893
See Also ... 6894

Understanding the File Hash Rule Condition in AppLocker ... 6895
See Also ... 6895

Understanding AppLocker Default Rules ... 6895
See Also ... 6896

Executable Rules in AppLocker ... 6896
See Also ... 6897

Windows Installer Rules in AppLocker ... 6897
See Also ... 6898

Script Rules in AppLocker ... 6898
See Also ... 6899

DLL Rules in AppLocker ... 6899
See Also ... 6899

Packaged Apps and Packaged App Installer Rules in AppLocker ... 6900
See also ... 6900

AppLocker Architecture and Components ... 6901
See Also ... 6902

AppLocker Processes and Interactions ... 6902
How policies are implemented by AppLocker ... 6902
Understanding AppLocker rules ... 6903
Understanding AppLocker policies ... 6904
Understanding AppLocker and Group Policy ... 6904
See Also ... 6904

AppLocker Functions ... 6904
Functions ... 6905

Security level ID ... 6905
See Also ... 6905

Security Considerations for AppLocker ... 6905
See Also ... 6907

Tools to Use with AppLocker ... 6907
See Also ... 6908

Using Event Viewer with AppLocker ... 6909
See Also ... 6911

AppLocker Settings ... 6911
See Also ... 6912

BitLocker Overview ... 6912
Feature description ... 6912
Practical applications ... 6913
New and changed functionality ... 6914
Removed or deprecated functionality ... 6915
System requirements ... 6916
See also ... 6916

What's New in BitLocker for Windows 8 and Windows Server 2012 ... 6917
BitLocker provisioning ... 6917
Used Disk Space Only encryption ... 6918
Standard User PIN and password change ... 6919
Network Unlock ... 6919
Support for Encrypted Hard Drives for Windows ... 6920
Additional information ... 6921

BitLocker Frequently Asked Questions (FAQ) ... 6921
Overview and requirements ... 6924
What is BitLocker? How does it work? ... 6924
Does BitLocker support multifactor authentication? ... 6925
What are the BitLocker hardware and software requirements? ... 6925
Why are two partitions required? Why does the system drive have to be so large? ... 6927
Which Trusted Platform Modules (TPMs) does BitLocker support? ... 6928
How can I tell if a TPM is on my computer? ... 6928
Can I use BitLocker on an operating system drive without a TPM? ... 6928
How do I obtain BIOS support for the TPM on my computer? ... 6929
What credentials are required to use BitLocker? ... 6929
What is the recommended boot order for computers that are going to be BitLocker-protected? ... 6929
Upgrading ... 6930

Can I upgrade my Windows 7–based computer to Windows 8 with BitLocker enabled? .. 6930
What is the difference between suspending and decrypting BitLocker? 6930
Do I have to decrypt my BitLocker-protected drive to download and install system updates and upgrades? 6930
Deployment and administration 6931
Can BitLocker deployment be automated in an enterprise environment? 6931
Can BitLocker encrypt more than just the operating system drive? 6931
Is there a noticeable performance impact when BitLocker is enabled on a Windows 8–based computer? 6931
Approximately how long will initial encryption take when BitLocker is turned on? 6932
What happens if the computer is turned off during encryption or decryption? 6932
Does BitLocker encrypt and decrypt the entire drive all at once when reading and writing data? 6932
How can I prevent users on a network from storing data on an unencrypted drive? 6932
What system changes would cause the integrity check on my operating system drive to fail? 6933
What causes BitLocker to start into recovery mode when attempting to start the operating system drive? 6933
Can I swap hard disks on the same computer if BitLocker is enabled on the operating system drive? 6935
Can I access my BitLocker-protected drive if I insert the hard disk into a different computer? 6935
Why is "Turn BitLocker on" not available when I right-click a drive? 6936
What type of disk configurations are supported by BitLocker? 6936
What if my disk configuration is not listed? 6937
Key management 6937
What is the difference between a TPM owner password, recovery password, recovery key, password, PIN, enhanced PIN, and startup key? 6937
How can the recovery password and recovery key be stored? 6940
Is it possible to add an additional method of authentication without decrypting the drive if I only have the TPM authentication method enabled? 6940
If I lose my recovery information, will the BitLocker-protected data be unrecoverable? 6940
Can the USB flash drive that is used as the startup key also be used to store the recovery key? 6940
Can I save the startup key on multiple USB flash drives? 6941
Can I save multiple (different) startup keys on the same USB flash drive? 6941
Can I generate multiple (different) startup keys for the same computer? 6941
Can I generate multiple PIN combinations? 6941
What encryption keys are used in BitLocker? How do they work together? 6941
Where are the encryption keys stored? 6941
Why do I have to use the function keys to enter the PIN or the 48-character recovery password? 6942

How does BitLocker help prevent an attacker from discovering the PIN that unlocks my operating system drive? ... 6942
How can I determine the manufacturer of my TPM? ... 6942
How can I evaluate a TPM's dictionary attack mitigation mechanism? ... 6943
Can PIN length and complexity be managed with Group Policy? ... 6943
BitLocker To Go ... 6943
What is BitLocker To Go? ... 6943
How can I authenticate or unlock my removable data drive? ... 6943
What happens if I try to open a BitLocker-protected, NTFS-formatted removable drive by using a computer running Windows XP or Windows Vista? ... 6943
Is there a way to ensure the BitLocker To Go Reader is not installed on FAT-formatted drives? ... 6944
Can I save files to my BitLocker-protected removable drive when I am using Windows XP or Windows Vista? ... 6944
Can I download a copy of the BitLocker To Go Reader? ... 6944
Why am I unable to access my removable drive on computers running Windows XP or Windows Vista when using the BitLocker To Go Reader? ... 6944
Active Directory Domain Services (AD DS) ... 6945
Does BitLocker require a schema extension to store recovery information in AD DS? ... 6945
What type of information is stored in AD DS? ... 6945
Does BitLocker encrypt recovery information as it is sent to AD DS? ... 6946
Is the BitLocker recovery information stored in plaintext in AD DS? ... 6946
What if BitLocker is enabled on a computer before the computer has joined the domain? 6947
Is there an event log entry recorded on the client computer to indicate the success or failure of the Active Directory backup? ... 6947
If I change the BitLocker recovery password on my computer and store the new password in AD DS, will AD DS overwrite the old password? ... 6947
What happens if the backup initially fails? Will BitLocker retry the backup? ... 6948
Security ... 6948
What form of encryption does BitLocker use? Is it configurable? ... 6948
What is the best practice for using BitLocker on an operating system drive? ... 6948
What are the implications of using the sleep or hibernate power management options? ... 6949
What are the advantages of a TPM? ... 6949
Is Microsoft pursuing any security certification for BitLocker? ... 6949
BitLocker Network Unlock ... 6949
What is BitLocker Network Unlock? ... 6949
What is required to use BitLocker Network Unlock? ... 6950
How is BitLocker Network Unlock different from automatic unlock? ... 6950
How do I unlock my computer when not connected to the network? ... 6950
Other questions ... 6950
Can I use EFS with BitLocker? ... 6950
Can I run a kernel debugger with BitLocker? ... 6951
How does BitLocker handle memory dumps? ... 6951

Can BitLocker support smart cards for pre-boot authentication? 6951
Can I use a non-Microsoft TPM driver? 6951
Can other tools that manage or modify the master boot record work with BitLocker? 6951
Why is the system check failing when I am encrypting my operating system drive? 6951
What can I do if the recovery key on my USB flash drive cannot be read? 6952
Why am I unable to save my recovery key to my USB flash drive? 6952
Why am I unable to automatically unlock my drive? 6952
Can I use BitLocker in Safe Mode? 6952
How do I "lock" a data drive? 6953
Can I use BitLocker with the Volume Shadow Copy Service? 6953
Does BitLocker support virtual hard disks (VHDs)? 6953
More information 6953

BitLocker Basic Deployment 6954
Using BitLocker to encrypt volumes 6954
Encrypting volumes using the BitLocker control panel 6954
Operating system volume 6955
Data volume 6957
Windows SkyDrive option 6958
Using BitLocker within Windows Explorer 6958
Down-level compatibility 6958
Encrypting volumes using the manage-bde command line interface 6959
Operating system volume 6959
Data volume 6960
Using manage-bde to encrypt volumes with BitLocker 6961
Encrypting volumes using the BitLocker Windows PowerShell cmdlets 6961
Operating system volume 6964
Data volume 6964
Using a SID based protector in Windows PowerShell 6965
Using PowerShell to encrypt volumes with BitLocker 6966
Checking BitLocker Status 6966
Checking BitLocker status with the control panel 6966
Checking BitLocker status with manage-bde 6966
Checking BitLocker status with Windows PowerShell 6967
Provisioning BitLocker during operating system deployment 6967
Decrypting BitLocker volumes 6967
Decrypting volumes using the BitLocker control panel applet 6967
Decrypting volumes using the manage-bde command line interface 6968
Decrypting volumes using the BitLocker Windows PowerShell cmdlets 6968
More information 6968

Prepare your organization for BitLocker: Planning and Policies 6969
Audit your environment 6969
Encryption keys and authentication 6969

TPM hardware configurations 6972
TPM states of existence 6972
Endorsement keys 6973
Non-TPM hardware configurations 6974
Disk configuration considerations 6974
BitLocker provisioning 6975
Used Disk Space Only encryption 6975
Active Directory Domain Services considerations 6976
FIPS support for recovery password protector 6978
More information 6979

BitLocker Group Policy Settings 6979
Overview 6980
BitLocker Group Policy settings 6980
Allow network unlock at startup 6982
Require additional authentication at startup 6983
Allow enhanced PINs for startup 6984
Configure minimum PIN length for startup 6985
Disallow standard users from changing the PIN or password 6986
Configure use of passwords for operating system drives 6987
Require additional authentication at startup (Windows Server 2008 and Windows Vista). 6988
Configure use of smart cards on fixed data drives 6990
Configure use of passwords on fixed data drives 6991
Configure use of smart cards on removable data drives 6992
Configure use of passwords on removable data drives 6993
Validate smart card certificate usage rule compliance 6995
Enable use of BitLocker authentication requiring preboot keyboard input on slates 6996
Deny write access to fixed drives not protected by BitLocker 6997
Deny write access to removable drives not protected by BitLocker 6998
Control use of BitLocker on removable drives 6999
Choose drive encryption method and cipher strength 7000
Configure use of hardware-based encryption for fixed data drives 7001
Configure use of hardware-based encryption for operating system drives 7002
Configure use of hardware-based encryption for removable data drives 7003
Enforce drive encryption type on fixed data drives 7005
Enforce drive encryption type on operating system drives 7006
Enforce drive encryption type on removable data drives 7007
Choose how BitLocker-protected operating system drives can be recovered 7008
Choose how users can recover BitLocker-protected drives (Windows Server 2008 and Windows Vista) 7009
Store BitLocker recovery information in Active Directory Domain Services (Windows Server 2008 and Windows Vista) 7010
Choose default folder for recovery password 7012

Choose how BitLocker-protected fixed drives can be recovered 7013
Choose how BitLocker-protected removable drives can be recovered 7014
Allow Secure Boot for integrity validation 7016
Provide the unique identifiers for your organization 7017
Prevent memory overwrite on restart 7018
Configure TPM platform validation profile for BIOS-based firmware configurations 7019
Configure TPM platform validation profile (Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2) 7021
Configure TPM platform validation profile for native UEFI firmware configurations 7022
Reset platform validation data after BitLocker recovery 7024
Use enhanced Boot Configuration Data validation profile 7025
Allow access to BitLocker-protected fixed data drives from earlier versions of Windows 7026
Allow access to BitLocker-protected removable data drives from earlier versions of Windows 7027
FIPS setting 7029
Power management Group Policy settings: Sleep and Hibernate 7029
About the Platform Configuration Register (PCR) 7030
See also 7031

BitLocker: How to deploy on Windows Server 2012 7031
Installing BitLocker 7031
To install BitLocker using Server Manager 7031
To install BitLocker using Windows PowerShell 7032
Using the servermanager module to install BitLocker 7032
Using the dism module to install BitLocker 7033
More information 7034

Protecting cluster shared volumes and storage area networks with BitLocker 7034
Configuring BitLocker on Cluster Shared Volumes 7034
Using BitLocker with Clustered Volumes 7034
Active Directory based Protector 7035
Turning on BitLocker before adding disks to a cluster using Windows PowerShell 7035
Turning on BitLocker for a clustered disk using Windows PowerShell 7036
Adding BitLocker encrypted volumes to a cluster using manage-bde 7037
Physical Disk Resources 7037
Restrictions on BitLocker actions with cluster volumes 7038
Other considerations when using BitLocker on CSV 2.0 7039

BitLocker: How to enable Network Unlock 7039
Network Unlock Core requirements 7040
Network Unlock sequence 7041
Configuring Network Unlock 7041
Step One: Install the WDS Server role 7041
Step Two: Confirm the WDS Service is running 7041

Step Three: Install the Network Unlock feature ... 7042
Step Four: Create the Network Unlock certificate ... 7042
Step Five: Deploy the private key and certificate to the WDS server ... 7043
Step Six: Configure Group Policy settings for Network Unlock ... 7043
Step Seven: Require TPM+PIN protectors at startup (recommended) ... 7044
Creating the certificate template for Network Unlock ... 7044
Subnet policy configuration files on WDS Server (Optional) ... 7045
Turning off Network Unlock ... 7047
Updating Network Unlock certificates ... 7047
Troubleshooting Network Unlock ... 7047
More information ... 7048

BitLocker Recovery Guide ... 7049
What is BitLocker recovery? ... 7049
What causes BitLocker recovery? ... 7049
Testing recovery ... 7052
Planning your recovery process ... 7053
Self-recovery ... 7053
Recovery password retrieval ... 7054
Record the name of the user's computer ... 7054
Verify the user's identity ... 7055
Locate the recovery password in AD DS ... 7055
Multiple recovery passwords ... 7055
Gather information to determine why recovery occurred ... 7055
Give the user the recovery password ... 7055
Post-recovery analysis ... 7056
Determine the root cause of the recovery ... 7056
Resolve the root cause ... 7057
Unknown PIN ... 7057
Lost startup key ... 7057
Changes to boot files ... 7058
Windows RE and BitLocker ... 7058
Using additional recovery information ... 7058
BitLocker key package ... 7058
Resetting recovery passwords ... 7059
Retrieving the BitLocker key package ... 7063
See also ... 7076

BCD Settings and BitLocker ... 7077
BitLocker and BCD Settings ... 7077
When secure boot is enabled ... 7077
Customizing BCD validation settings ... 7077
Default BCD validation profile ... 7078
Full list of friendly names for ignored BCD settings ... 7079

BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker ... 7085
Manage-bde ... 7085
Using manage-bde with operating system volumes ... 7086
Using manage-bde with data volumes ... 7087
Repair-bde ... 7087
BitLocker cmdlets for Windows PowerShell ... 7088
Using the BitLocker Windows PowerShell cmdlets with operating system volumes ... 7091
Using the BitLocker Windows PowerShell cmdlets with data volumes ... 7092
Using an AD Account or Group protector in Windows PowerShell ... 7092
More information ... 7093

BitLocker: Use BitLocker Recovery Password Viewer ... 7093
Before you start ... 7093
More information ... 7094

Credentials Protection and Management ... 7094
Restricted Admin mode for Remote Desktop Connection ... 7095
LSA protection ... 7095
Protected Users security group ... 7095
Authentication Policy and Authentication Policy Silos ... 7096
See also ... 7096

Configuring Additional LSA Protection ... 7096
Protected process requirements for plug-ins or drivers ... 7097
Recommended practices ... 7097
How to identify LSA plug-ins and drivers that fail to run as a protected process ... 7097
Before opting in: How to identify plug-ins and drivers loaded by the lsass.exe ... 7098
After opting in: How to identify plug-ins and drivers loaded by the lsass.exe ... 7099
How to configure additional LSA protection of credentials ... 7100
On x86-based or x64-based devices using Secure Boot and UEFI or not ... 7100
Verifying LSA protection ... 7101
Additional resources ... 7101

Cached and Stored Credentials Technical Overview ... 7101
Introduction ... 7102
Credential authenticators ... 7102
Plaintext credentials ... 7102
NT hash ... 7103
LM hash ... 7103
Windows logon cached password verifiers ... 7103
Credentials storage ... 7103
Security Accounts Manager database ... 7104
LSASS process memory ... 7104
LSA secrets on the hard disk drive ... 7105

AD DS database (NTDS.DIT) 7105
Credential Manager store 7105
Related resources 7106

Protected Users Security Group 7106
How the Protected Users group works 7107
Event log information 7108
Deployment requirements 7110
Additional resources 7110

Authentication Policies and Authentication Policy Silos 7110
About authentication policy silos 7111
About authentication policies 7112
How it works 7115
How the Kerberos protocol is used with authentication policy silos and policies 7116
How restricting a user sign-in works 7117
How restricting service ticket issuance works 7118
Associated error and informational event messages 7119
See also 7120

Encrypted Hard Drive 7121
System Requirements 7122
Technical overview 7122
Configuring Encrypted Hard Drives as Startup drives 7122
Encrypted Hard Drive Architecture 7123
Re-configuring Encrypted Hard Drives 7124

Exchange ActiveSync Policy Engine Overview 7124
Feature description 7124
Practical applications 7124
How it works 7125
Policies supported by a mail app and EAS policy engine 7125
About password policies and accounts 7131
Password length and complexity supported by account types 7131
Policy application on administrator and standard user accounts 7132
Policies specified by different sources 7133
Multiuser support 7133
Policy reset 7133
EAS policies and provisioning refresh 7133
Device lock 7134
Autologon behavior 7134
New and changed functionality 7134
Software requirements 7135
Additional resources 7135

Use Exchange ActiveSync Policies for Device Management ... 7136
Addressing policy compliance ... 7136
Supported EAS policy engine policies details ... 7137
AllowSimpleDevicePassword ... 7137
DisallowConvenienceLogon ... 7138
MaxInactivityTimeDeviceLock ... 7138
MaxDevicePasswordFailedAttempts ... 7139
MinDevicePasswordComplexCharacters ... 7140
MinDevicePasswordLength ... 7140
DevicePasswordExpiration ... 7141
DevicePasswordHistory ... 7141
RequireDeviceEncryption ... 7142
Related resources ... 7143

Group Managed Service Accounts Overview ... 7143
Feature description ... 7143
Practical applications ... 7143
New and changed functionality ... 7144
Deprecated functionality ... 7144
Software requirements ... 7144
Server Manager information ... 7145
See also ... 7145

What's New for Managed Service Accounts ... 7146
Managed Service Accounts description ... 7146
New and changed functionality ... 7146
Group Managed Service Accounts ... 7146
Removed or deprecated functionality ... 7147
See also ... 7147

Getting Started with Group Managed Service Accounts ... 7147
Prerequisites ... 7148
Introduction ... 7148
Requirements for group Managed Service Accounts ... 7149
Deploying a new server farm ... 7151
Step 1: Provisioning group Managed Service Accounts ... 7151
Step 2: Configuring service identity application service ... 7154
Adding member hosts to an existing server farm ... 7155
Updating the group Managed Service Account properties ... 7156
Decommissioning member hosts from an existing server farm ... 7156
Step 1: Remove member host from gMSA ... 7156
Step 2: Removing a group Managed Service Account from the system ... 7158
Decommissioning an existing server farm ... 7158
Step 1: Deleting Active Directory Objects ... 7159

Step 2: Removing a group Managed Service Account from the system 7159
See also 7159

Create the Key Distribution Services KDS Root Key 7160
See Also 7161

Kerberos Authentication Overview 7161
Feature description 7161
Practical applications 7161
New and changed functionality 7162
See also 7162
See Also 7163

What's New in Kerberos Authentication 7163
Feature description 7163
Improvements to reduce authentication failures due to large service tickets 7164
KDC resource group compression 7164
Increase in the Kerberos SSPI context token buffer size 7165
Group Policy to set a maximum for the Kerberos SSPI context token buffer size 7165
KDC with warning events for large Kerberos tickets 7165
New and changed functionality for developers 7166
Extended protection for untrusted targets 7166
Kernel support for EncryptMessage and DecryptMessage 7166
Device sign-in with certificates 7166
New and changed functionality for IT professionals 7167
Branch office support for authentication to resources outside the branch office 7167
Support for claims, compound authentication, and Kerberos armoring 7167
Claims 7170
Compound authentication 7171
Kerberos Armoring (Flexible Authentication Secure Tunneling (FAST)) 7172
Resource-based constrained delegation across domains and forest 7173
Enforcement for forest boundary for Kerberos full delegation 7173
KDC proxy service 7174
Strict KDC Validation default changes 7174
Configuration and maintenance improvements 7175
KDC service servicing without a reboot 7175
New KDC system events 7175
KDC operational log 7176
Kerberos operational log 7176
Performance monitor counters 7177
Changes to SetSPN 7178
See also 7179

Kerberos Constrained Delegation Overview 7179

Feature description 7179
Practical applications 7180
New and changed functionality 7180
Software requirements 7181
See also 7181

NTLM Overview 7182
Feature description 7182
Current applications 7183
New and changed functionality 7183
Removed or deprecated functionality 7183
Server Manager information 7183
See also 7183

Passwords Overview 7185
Feature description 7185
Practical applications 7185
New and changed functionality 7185
Deprecated functionality 7186
Software requirements 7186
See also 7186

Security Auditing Overview 7187
Feature description 7187
Practical applications 7187
New and changed functionality 7188
Managing security auditing 7189
Related resources 7189

Advanced Security Auditing FAQ 7189
What is Windows security auditing and why might I want to use it? 7190
What is the difference between audit policies located in Local Policies\Audit Policy and audit policies located in Advanced Audit Policy Configuration? 7191
What is the interaction between basic audit policy settings and advanced audit policy settings? 7192
How are audit settings merged by Group Policy? 7193
What is the difference between an object DACL and an object SACL? 7194
Why are audit policies applied on a per-computer basis rather than per user? 7194
What are the differences in auditing functionality between versions of Windows? 7195
Can I use advanced audit policies from a domain controller running Windows Server 2003 or Windows 2000 Server? 7195
What is the difference between success and failure events? Is something wrong if I get a failure audit? 7195
How can I set an audit policy that affects all objects on a computer? 7195

How do I figure out why someone was able to access a resource? 7196
How do I know when changes are made to access control settings, by whom, and what the changes were? 7196
How can I roll back security audit policies from the advanced audit policy to the basic audit policy? 7197
How can I monitor if changes are made to audit policy settings? 7197
How can I minimize the number of events that are generated? 7197
What are the best tools to model and manage audit policies? 7198
Where can I find information about all the possible events that I might receive? 7198
Where can I find more detailed information? 7198

Advanced Security Auditing Walkthrough 7199
About this guide 7199
Deploying advanced audit policy settings in a test environment 7200
Steps for deploying advanced audit policies in a test environment 7202
Step 1: Set up the infrastructure 7202
Configure the domain controller (CONTOSO-DC) 7203
Configure the member server (CONTOSO-SRV) running Windows Server 2008 R2 7204
Configure the client computer (CONTOSO-CLNT) 7206
Step 2: Create and verify an advanced audit policy 7208
Step 3: Create and verify an audit policy that provides the reason for object access 7210
Step 4: Create and verify a global object access policy 7212
Step 5: Create and verify additional advanced audit policies 7213
Managing per-user auditing 7214
Optional section: Roll back security audit policy from advanced audit policy to basic audit policy 7215

Which Editions of Windows Support Advanced Audit Policy Configuration 7215
Are there any special considerations? 7216

Planning and Deploying Advanced Security Audit Policies 7217
Overview 7217
About this guide 7218
Supported versions of Windows 7218
Terminology used in this guide 7218
Understanding the security audit policy design process 7219
Identifying your Windows security audit policy deployment goals 7220
Network environment 7221
Data and resources 7221
Users 7222
Computers 7223
Regulatory requirements 7224
Mapping the security audit policy to groups of users, computers, and resources in your organization 7225

Mapping your security auditing goals to a security audit policy configuration 7226
Exploring audit policy options 7227
Choosing audit settings to use 7227
Data and resource activity 7228
User activity 7229
Network activity 7231
Confirm operating system version compatibility 7232
Success, failure, or both 7233
Planning for security audit monitoring and management 7233
Deploying the security audit policy 7235

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects 7235
In this guide 7236
Related resources 7236

Monitor the Central Access Policies that Apply on a File Server 7236
Related resource 7238

Monitor the Use of Removable Storage Devices 7238
Related resource 7240

Monitor Resource Attribute Definitions 7240
Related resource 7241

Monitor Central Access Policy and Rule Definitions 7241
Related resource 7243

Monitor User and Device Claims During Sign-in 7243
Related resource 7244

Monitor the Resource Attributes on Files and Folders 7244
Related resource 7246

Monitor the Central Access Policies Associated with Files and Folders 7246
Related resource 7248

Monitor Claim Types 7248
Related resource 7249

Advanced Security Audit Policy Settings 7249

Audit Credential Validation 7253
Related resource 7254

Audit Kerberos Authentication Service 7254
Related resource 7255

Audit Kerberos Service Ticket Operations 7255

Related resource 7255

Audit Other Account Logon Events 7256
Related resource 7257

Audit Application Group Management 7257
Related resource 7258

Audit Computer Account Management 7258
Related resource 7259

Audit Distribution Group Management 7259
Related resource 7260

Audit Other Account Management Events 7260
Related resource 7261

Audit Security Group Management 7261
Related resource 7262

Audit User Account Management 7262
Related resource 7264

Audit DPAPI Activity 7264
Related resource 7264

Audit Process Creation 7265
Related resource 7265

Audit Process Termination 7265
Related resource 7266

Audit RPC Events 7266
Related resource 7266

Audit Detailed Directory Service Replication 7266
Related resource 7267

Audit Directory Service Access 7267
Related resource 7268

Audit Directory Service Changes 7268
Related resource 7269

Audit Directory Service Replication 7269
Related resource 7269

Audit Account Lockout 7270
Related resource 7270

Audit IPsec Extended Mode 7270
Related resource 7273
Audit IPsec Main Mode 7273
Related resource 7274
Audit IPsec Quick Mode 7275
Related resource 7276
Audit Logoff 7276
Related resource 7276
Audit Logon 7276
Related resource 7277
Audit Network Policy Server 7277
Related resource 7278
Audit Other Logon/Logoff Events 7279
Related resource 7280
Audit Special Logon 7280
Related resource 7280
Audit Application Generated 7281
Related resource 7281
Audit Certification Services 7281
Related resource 7284
Audit Detailed File Share 7284
See Also 7284
Audit File Share 7284
Related resource 7285
Audit File System 7285
Related resource 7286
Audit Filtering Platform Connection 7286
Related resource 7288
Audit Filtering Platform Packet Drop 7288
Related resource 7288
Audit Handle Manipulation 7288
Related resource 7289
Audit Kernel Object 7289

Related resource 7290

Audit Other Object Access Events 7290
Related resource 7291

Audit Registry 7292
Related resource 7292

Audit SAM 7292
Related resource 7293

Audit Audit Policy Change 7294
Related resource 7295

Audit Authentication Policy Change 7295
Related resource 7296

Audit Authorization Policy Change 7296
Related resource 7297

Audit Filtering Platform Policy Change 7297
Related resource 7301

Audit MPSSVC Rule-Level Policy Change 7301
Related resource 7303

Audit Other Policy Change Events 7303
Related resource 7304

Audit Sensitive Privilege Use 7304
Related resource 7305

Audit Non-Sensitive Privilege Use 7306
Related resource 7307

Audit Other Privilege Use Events 7307
Related resource 7307

Audit IPsec Driver 7307
Related resource 7310

Audit Other System Events 7310
Related resource 7313

Audit Security State Change 7313
Related resource 7314

Audit Security System Extension 7314
Related resource 7315

Audit System Integrity 7315
Related resource 7316

Registry (Global Object Access Auditing) 7317
Related resource 7317

File System (Global Object Access Auditing) 7317
Related resource 7317

Security Policy Settings Overview 7318
Security policy settings description 7318
Practical applications 7318
New and changed functionality 7318
Deprecated functionality 7319
See also 7320

Security Policy Settings Technical Overview 7321
Introduction 7321
Policy-based Security Settings management 7322
Common scenarios for using Security Settings policies 7323
Dependencies on other operating system technologies 7323
Security Settings policies and Group Policy 7325
Security Settings extension architecture 7326
Security Settings policy processes and interactions 7329
Group Policy processing 7329
Group Policy Objects storage 7330
Group Policy processing order 7330
Security Settings policy processing 7331
Merging of security policies on domain controllers 7332
Special considerations for domain controllers 7333
When security settings are applied 7333
Persistence of Security Settings policy 7333
Permissions required for policy to apply 7334
Filtering security policy 7334
Migration of GPOs containing Security Settings 7334
Additional resources 7335
See Also 7335

Security Policy Settings Reference 7335
Overview 7336
Contents of this guide 7336
Additional resources 7337

How to Configure Security Policy Settings 7338
To configure a setting for your local computer 7338

To configure a setting for computer that is joined to a domain 7339
To configure a setting for a domain controller 7339
See also 7340

Account Policies 7340

Password Policy 7341

Enforce password history 7342
Reference 7342
Possible values 7343
Best practices 7343
Location 7343
Default values 7343
Operating system version differences 7344
Policy management 7344
Restart requirement 7344
Security considerations 7344
Vulnerability 7344
Countermeasure 7344
Potential impact 7345
See Also 7345

Maximum password age 7345
Reference 7345
Possible values 7345
Best practices 7346
Location 7346
Default values 7346
Operating system version differences 7346
Policy management 7346
Restart requirement 7346
Security considerations 7347
Vulnerability 7347
Countermeasure 7347
Potential impact 7347
See Also 7347

Minimum password age 7347
Reference 7347
Possible values 7348
Best practices 7348
Location 7348
Default values 7348
Operating system version differences 7349

Policy management ... 7349
Restart requirement ... 7349
Security considerations ... 7349
Vulnerability ... 7349
Countermeasure ... 7349
Potential impact ... 7350
See Also ... 7350

Minimum password length ... 7350
Reference ... 7350
Possible values ... 7350
Best practices ... 7350
Location ... 7351
Default values ... 7351
Operating system version differences ... 7351
Policy management ... 7351
Restart requirement ... 7351
Security considerations ... 7351
Vulnerability ... 7352
Countermeasure ... 7352
Potential impact ... 7352
See Also ... 7352

Password must meet complexity requirements ... 7353
Reference ... 7353
Possible values ... 7354
Best practices ... 7354
Location ... 7354
Default values ... 7354
Operating system version differences ... 7355
Security considerations ... 7355
Vulnerability ... 7355
Countermeasure ... 7355
Potential impact ... 7355
See Also ... 7356

Store passwords using reversible encryption ... 7356
Reference ... 7356
Possible values ... 7356
Best practices ... 7357
Location ... 7357
Default values ... 7357
Operating system version differences ... 7357
Security considerations ... 7357

Vulnerability 7358
Countermeasure 7358
Potential impact 7358
See Also 7358

Account Lockout Policy 7358
Overview of Account Lockout Policy settings 7358

Account lockout duration 7359
Reference 7359
Possible values 7359
Location 7359
Default values 7360
Operating system version differences 7360
Security considerations 7360
Vulnerability 7360
Countermeasure 7360
Potential impact 7361
See Also 7361

Account lockout threshold 7361
Reference 7361
Possible values 7362
Best practices 7362
Location 7362
Default values 7362
Operating system version differences 7363
Security considerations 7363
Vulnerability 7363
Countermeasure 7363
Potential impact 7364
See Also 7364

Reset account lockout counter after 7364
Reference 7365
Possible values 7365
Best practices 7365
Location 7365
Default values 7365
Operating system version differences 7366
Security considerations 7366
Vulnerability 7366
Countermeasure 7366
Potential impact 7366

See Also 7366

Kerberos Policy 7366

Enforce user logon restrictions 7367
Reference 7367
Best practices 7368
Location 7368
Default Values 7368
Operating system version differences 7368
Policy management 7368
Group Policy 7369
Security considerations 7369
Vulnerability 7369
Countermeasure 7369
Potential impact 7369
See Also 7369

Maximum lifetime for service ticket 7369
Reference 7370
Best practices 7370
Location 7370
Default Values 7370
Operating system version differences 7371
Policy management 7371
Group Policy 7371
Security considerations 7371
Vulnerability 7371
Countermeasure 7372
Potential impact 7372
See Also 7372

Maximum lifetime for user ticket 7372
Reference 7372
Best practices 7372
Location 7372
Default Values 7373
Operating system version differences 7373
Policy management 7373
Group Policy 7373
Security considerations 7374
Vulnerability 7374
Countermeasure 7374
Potential impact 7374

See Also 7374

Maximum lifetime for user ticket renewal 7374
Reference 7374
Best practices 7375
Location 7375
Default values 7375
Operating system version differences 7375
Policy management 7375
Group Policy 7375
Security considerations 7376
Vulnerability 7376
Countermeasure 7376
Potential impact 7376
See Also 7376

Maximum tolerance for computer clock synchronization 7376
Reference 7376
Best practices 7377
Location 7377
Default values 7377
Operating system version differences 7378
Policy management 7378
Group Policy 7378
Security considerations 7378
Vulnerability 7379
Countermeasure 7379
Potential impact 7379
Additional references 7379
See Also 7379

Audit Policy 7379

Security Options 7380

Accounts: Administrator account status 7386
Reference 7386
Possible values 7386
Best practices 7386
Location 7387
Default values 7387
Operating system version differences 7387
Policy management 7387
Restart requirement 7388
Safe mode considerations 7388

How to access a disabled Administrator account 7388
Security considerations 7389
Vulnerability 7389
Countermeasure 7389
Potential impact 7389

Accounts: Block Microsoft accounts 7389
Reference 7390
Possible values 7390
Best practices 7390
Location 7390
Default values 7390
Operating system version differences 7391
Policy management 7391
Restart requirement 7391
Group Policy 7391
Security considerations 7391
Vulnerability 7391
Countermeasure 7392
Potential impact 7392
See Also 7392

Accounts: Guest account status 7392
Reference 7392
Possible values 7392
Best practices 7392
Location 7393
Default values 7393
Operating system version differences 7393
Security considerations 7393
Vulnerability 7393
Countermeasure 7394
Potential impact 7394

Accounts: Limit local account use of blank passwords to console logon only 7394
Reference 7394
Possible values 7395
Best practices 7395
Location 7395
Default values 7395
Operating system version differences 7395
Policy management 7395
Restart requirement 7396
Policy conflict considerations 7396

Group Policy 7396
Security considerations 7396
Vulnerability 7396
Countermeasure 7396
Potential impact 7396

Accounts: Rename administrator account 7397
Reference 7397
Possible values 7397
Best practices 7397
Location 7397
Default values 7397
Operating system version differences 7398
Policy management 7398
Restart requirement 7398
Policy conflict considerations 7398
Group Policy 7398
Security considerations 7398
Vulnerability 7399
Countermeasure 7399
Potential impact 7399

Accounts: Rename guest account 7399
Reference 7399
Possible values 7400
Best practices 7400
Location 7400
Default values 7400
Operating system version differences 7400
Policy management 7401
Restart requirement 7401
Policy conflict considerations 7401
Group Policy 7401
Security considerations 7401
Vulnerability 7401
Countermeasure 7401
Potential impact 7401

Audit: Audit the access of global system objects 7402
Reference 7402
Possible values 7402
Best practices 7403
Location 7403
Default values 7403

Operating system version differences 7403
Policy management 7403
Restart requirement 7403
Group Policy 7403
Auditing 7404
Security considerations 7405
Vulnerability 7405
Countermeasure 7405
Potential impact 7405
Additional resources 7406

Audit: Audit the use of Backup and Restore privilege 7406
Reference 7406
Possible values 7406
Best practices 7406
Location 7406
Default values 7406
Operating system version differences 7407
Policy management 7407
Restart requirement 7407
Auditing 7407
Security considerations 7408
Vulnerability 7408
Countermeasure 7408
Potential impact 7408

Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings 7408
Reference 7408
Possible values 7409
Best practices 7409
Location 7409
Default values 7409
Operating system version differences 7409
Policy management 7409
Restart requirement 7410
Group Policy 7410
Auditing 7410
Command-line tools 7410
Security considerations 7410
Vulnerability 7410
Countermeasure 7410
Potential impacts 7410

Audit: Shut down system immediately if unable to log security audits 7411
Reference 7411
Possible values 7412
Best practices 7412
Location 7412
Default values 7412
Operating system version differences 7412
Policy management 7413
Restart requirement 7413
Group Policy 7413
Security considerations 7413
Vulnerability 7413
Countermeasure 7413
Potential impact 7413

DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax 7414
Reference 7414
Possible values 7414
Location 7415
Default values 7415
Operating system version differences 7415
Policy management 7415
Restart requirement 7415
Group Policy 7416
Security considerations 7416
Vulnerability 7416
Countermeasure 7417
Potential impact 7417

DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax 7417
Reference 7417
Possible values 7418
Location 7418
Default values 7418
Operating system version differences 7419
Policy management 7419
Restart requirement 7419
Group Policy 7419
Security considerations 7419
Vulnerability 7419
Countermeasure 7420
Potential impact 7420

Devices: Allow undock without having to log on ... 7420
Reference ... 7420
Possible values ... 7421
Best practices ... 7421
Location ... 7421
Default values ... 7421
Operating system version differences ... 7421
Policy management ... 7422
Restart requirement ... 7422
Security considerations ... 7422
Vulnerability ... 7422
Countermeasure ... 7422
Potential impact ... 7422

Devices: Allowed to format and eject removable media ... 7422
Reference ... 7422
Possible values ... 7423
Best practices ... 7423
Location ... 7423
Default values ... 7423
Operating system version differences ... 7423
Policy management ... 7423
Restart requirement ... 7424
Security considerations ... 7424
Vulnerability ... 7424
Countermeasure ... 7424
Potential impact ... 7424

Devices: Prevent users from installing printer drivers ... 7424
Reference ... 7424
Possible values ... 7425
Best practices ... 7425
Location ... 7425
Default values ... 7425
Operating system version differences ... 7426
Policy management ... 7426
Restart requirement ... 7426
Security considerations ... 7426
Vulnerability ... 7426
Countermeasure ... 7426
Potential impact ... 7426

Devices: Restrict CD-ROM access to locally logged-on user only ... 7427
Reference ... 7427

Possible values 7427
Best practices 7427
Location 7427
Default values 7427
Operating system version differences 7428
Policy management 7428
Restart requirement 7428
Security considerations 7428
Vulnerability 7428
Countermeasure 7428
Potential impact 7429

Devices: Restrict floppy access to locally logged-on user only 7429
Reference 7429
Possible values 7429
Best practices 7430
Location 7430
Default values 7430
Operating system version differences 7430
Policy management 7430
Restart requirement 7430
Security considerations 7431
Vulnerability 7431
Countermeasure 7431
Potential impact 7431

Domain controller: Allow server operators to schedule tasks 7431
Reference 7431
Possible values 7432
Best practices 7432
Location 7432
Default values 7432
Operating system version differences 7433
Policy management 7433
Restart requirement 7433
Command-line tools 7433
Security considerations 7433
Vulnerability 7433
Countermeasure 7433
Potential impact 7433

Domain controller: LDAP server signing requirements 7434
Reference 7434
Possible values 7434

Best practices 7434
Location 7435
Default values 7435
Operating system version differences 7435
Policy management 7435
Restart requirement 7435
Security considerations 7435
Vulnerability 7436
Countermeasure 7436
Potential impact 7436

Domain controller: Refuse machine account password changes 7436
Reference 7436
Possible values 7436
Best practices 7437
Location 7437
Default values 7437
Operating system version differences 7437
Policy management 7437
Restart requirement 7437
Security considerations 7438
Vulnerability 7438
Countermeasure 7438
Potential impact 7438

Domain member: Digitally encrypt or sign secure channel data (always) 7438
Reference 7438
Possible values 7439
Best practices 7439
Location 7440
Default values 7440
Operating system version differences 7440
Policy management 7440
Restart requirement 7441
Group Policy 7441
Security considerations 7441
Vulnerability 7441
Countermeasure 7441
Potential impact 7441

Domain member: Digitally encrypt secure channel data (when possible) 7442
Reference 7442
Possible values 7443
Best practices 7443

Location ... 7444
Default values ... 7444
Operating system version differences ... 7444
Policy management ... 7444
Restart requirement ... 7444
Group Policy ... 7445
Security considerations ... 7445
Vulnerability ... 7445
Countermeasure ... 7445
Potential impact ... 7445

Domain member: Digitally sign secure channel data (when possible) ... 7446
Reference ... 7446
Possible values ... 7447
Best practices ... 7447
Location ... 7447
Default values ... 7447
Operating system version differences ... 7448
Policy management ... 7448
Restart requirement ... 7448
Group Policy ... 7448
Security considerations ... 7448
Vulnerability ... 7448
Countermeasure ... 7449
Potential impact ... 7449

Domain member: Disable machine account password changes ... 7449
Reference ... 7449
Possible values ... 7450
Best practices ... 7450
Location ... 7450
Default values ... 7450
Operating system version differences ... 7450
Policy management ... 7451
Restart requirement ... 7451
Security considerations ... 7451
Vulnerability ... 7451
Countermeasure ... 7451
Potential impact ... 7451

Domain member: Maximum machine account password age ... 7451
Reference ... 7451
Possible values ... 7452
Best practices ... 7452

Location 7452
Default values 7452
Operating system version differences 7453
Policy management 7453
Restart requirement 7453
Security considerations 7453
Vulnerability 7453
Countermeasure 7453
Potential impact 7453

Domain member: Require strong (Windows 2000 or later) session key 7453
Reference 7454
Possible values 7454
Best practices 7454
Location 7454
Default values 7454
Operating system version differences 7455
Policy management 7455
Restart requirement 7455
Group Policy 7455
Security considerations 7455
Vulnerability 7455
Countermeasure 7456
Potential impact 7456

Interactive logon: Display user information when the session is locked 7456
Reference 7456
Possible values 7456
Best practices 7457
Location 7457
Default values 7457
Operating system version differences 7458
Policy management 7458
Restart requirement 7458
Policy conflict considerations 7458
Group Policy 7458
Security considerations 7458
Vulnerability 7458
Countermeasure 7458
Potential impact 7459

Interactive logon: Do not display last user name 7459
Reference 7459
Possible values 7459

Best practices 7459
Location 7460
Default values 7460
Operating system version differences 7460
Policy management 7460
Restart requirement 7460
Policy conflict considerations 7460
Group Policy 7461
Security considerations 7461
Vulnerability 7461
Countermeasure 7461
Potential impact 7461

Interactive logon: Do not require CTRL+ALT+DEL 7461
Reference 7461
Possible values 7462
Best practices 7462
Location 7462
Default values 7462
Operating system version differences 7463
Policy management 7463
Restart requirement 7463
Policy conflict considerations 7463
Group Policy 7463
Security considerations 7463
Vulnerability 7464
Countermeasure 7464
Potential impact 7464

Interactive logon: Machine inactivity limit 7464
Reference 7464
Possible values 7464
Best practices 7465
Location 7465
Default values 7465
Operating system version differences 7465
Policy management 7465
Restart requirement 7465
Group Policy 7466
Security considerations 7466
Vulnerability 7466
Countermeasure 7466
Potential impact 7466
See Also 7466

Interactive logon: Machine account lockout threshold ... 7466
Reference ... 7467
Possible values ... 7467
Best practices ... 7467
Location ... 7467
Default values ... 7467
Operating system version differences ... 7468
Policy management ... 7468
Restart requirement ... 7468
Group Policy ... 7468
Security considerations ... 7468
Vulnerability ... 7468
Countermeasure ... 7469
Potential impact ... 7469
See Also ... 7469

Interactive logon: Message text for users attempting to log on ... 7469
Reference ... 7469
Possible values ... 7470
Best practices ... 7470
Location ... 7470
Default values ... 7470
Operating system version differences ... 7470
Policy management ... 7471
Restart requirement ... 7471
Effect on Windows 2000 operating systems ... 7471
Security considerations ... 7471
Vulnerability ... 7471
Countermeasure ... 7472
Potential impact ... 7472

Interactive logon: Message title for users attempting to log on ... 7472
Reference ... 7472
Possible values ... 7473
Best practices ... 7473
Location ... 7473
Default values ... 7473
Operating system version differences ... 7474
Policy management ... 7474
Restart requirement ... 7474
Effect on Windows 2000 operating systems ... 7474
Security considerations ... 7474
Vulnerability ... 7475
Countermeasure ... 7475

Potential impact 7475

Interactive logon: Number of previous logons to cache (in case domain controller is not available) 7476
Reference 7476
Possible values 7476
Best practices 7476
Location 7477
Default values 7477
Operating system version differences 7477
Policy management 7477
Restart requirement 7477
Policy conflict considerations 7478
Group Policy 7478
Security considerations 7478
Vulnerability 7478
Countermeasure 7478
Potential impact 7478

Interactive logon: Prompt user to change password before expiration 7479
Reference 7479
Possible values 7479
Best practices 7479
Location 7479
Default values 7479
Operating system version differences 7480
Policy management 7480
Restart requirement 7480
Policy conflict considerations 7480
Group Policy 7480
Security considerations 7480
Vulnerability 7480
Countermeasure 7481
Potential impact 7481

Interactive logon: Require Domain Controller authentication to unlock workstation 7481
Reference 7481
Possible values 7482
Best practices 7482
Location 7482
Default values 7482
Operating system version differences 7483
Policy management 7483
Restart requirement 7483

Policy conflict considerations 7483
Group Policy 7483
Security considerations 7483
Vulnerability 7483
Countermeasure 7483
Potential impact 7484

Interactive logon: Require smart card 7484
Reference 7484
Possible values 7484
Best practices 7484
Location 7485
Default values 7485
Operating system version differences 7485
Policy management 7485
Restart requirement 7485
Policy conflict considerations 7485
Group Policy 7486
Security considerations 7486
Vulnerability 7486
Countermeasure 7486
Potential impact 7486

Interactive logon: Smart card removal behavior 7486
Reference 7487
Possible values 7487
Best practices 7487
Location 7487
Default values 7488
Operating system version differences 7488
Policy management 7488
Restart requirement 7488
Policy conflict considerations 7488
Group Policy 7488
Security considerations 7489
Vulnerability 7489
Countermeasure 7489
Potential impact 7489

Microsoft network client: Digitally sign communications (always) 7489
Reference 7489
Possible values 7490
Best practices 7490
Location 7491

Default values 7491
Operating system version differences 7491
Policy management 7491
Restart requirement 7491
Security considerations 7491
Vulnerability 7492
Countermeasure 7492
Potential impact 7492

Microsoft network client: Digitally sign communications (if server agrees) 7493
Reference 7493
Possible values 7494
Best practices 7494
Location 7494
Default values 7494
Operating system version differences 7495
Policy management 7495
Restart requirement 7495
Security considerations 7495
Vulnerability 7495
Countermeasure 7495
Potential impact 7496

Microsoft network client: Send unencrypted password to third-party SMB servers 7496
Reference 7496
Possible values 7496
Best practices 7497
Location 7497
Default values 7497
Operating system version differences 7497
Policy management 7497
Restart requirement 7498
Security considerations 7498
Vulnerability 7498
Countermeasure 7498
Potential impact 7498

Microsoft network server: Amount of idle time required before suspending session 7498
Reference 7498
Possible values 7499
Best practices 7499
Location 7499
Default values 7499
Operating system version differences 7499

Policy management ... 7500
Restart requirement ... 7500
Security considerations ... 7500
Vulnerability ... 7500
Countermeasure ... 7500
Potential impact ... 7500

Microsoft network server: Attempt S4U2Self to obtain claim information ... 7500
Reference ... 7500
Possible values ... 7501
Best practices ... 7501
Location ... 7501
Default values ... 7501
Operating system version differences ... 7502
Policy management ... 7502
Restart requirement ... 7502
Group Policy ... 7502
Security considerations ... 7502
Vulnerability ... 7502
Countermeasure ... 7503
Potential impact ... 7503
See Also ... 7503

Microsoft network server: Digitally sign communications (always) ... 7503
Reference ... 7503
Possible values ... 7504
Best practices ... 7504
Location ... 7504
Default values ... 7504
Operating system version differences ... 7505
Policy management ... 7505
Restart requirement ... 7505
Security considerations ... 7505
Vulnerability ... 7505
Countermeasure ... 7506
Potential impact ... 7506

Microsoft network server: Digitally sign communications (if client agrees) ... 7506
Reference ... 7507
Possible values ... 7507
Best practices ... 7508
Location ... 7508
Default values ... 7508
Operating system version differences ... 7508

Policy management ... 7508
Restart requirement ... 7509
Security considerations ... 7509
Vulnerability ... 7509
Countermeasure ... 7509
Potential impact ... 7510

Microsoft network server: Disconnect clients when logon hours expire ... 7510
Reference ... 7510
Possible values ... 7510
Best practices ... 7511
Location ... 7511
Default values ... 7511
Operating system version differences ... 7511
Policy management ... 7511
Restart requirement ... 7511
Group Policy ... 7511
Security considerations ... 7512
Vulnerability ... 7512
Countermeasure ... 7512
Potential impact ... 7512

Microsoft network server: Server SPN target name validation level ... 7512
Reference ... 7512
Possible values ... 7512
Best practices ... 7513
Location ... 7513
Default values ... 7513
Operating system version differences ... 7514
Policy management ... 7514
Restart requirement ... 7514
Policy conflict considerations ... 7514
Group Policy ... 7514
Security considerations ... 7514
Vulnerability ... 7514
Countermeasure ... 7514
Potential impact ... 7515

Network access: Allow anonymous SID/Name translation ... 7515
Reference ... 7515
Possible values ... 7515
Best practices ... 7515
Location ... 7516
Default values ... 7516

Operating system version differences 7516
Policy management 7516
Restart requirement 7516
Group Policy 7516
Security considerations 7517
Vulnerability 7517
Countermeasure 7517
Potential impact 7517

Network access: Do not allow anonymous enumeration of SAM accounts 7517
Reference 7517
Possible values 7518
Best practices 7518
Location 7518
Default values 7518
Operating system version differences 7519
Policy management 7519
Restart requirement 7519
Policy conflicts 7519
Group Policy 7519
Security considerations 7519
Vulnerability 7519
Countermeasure 7519
Potential impact 7520

Network access: Do not allow anonymous enumeration of SAM accounts and shares 7520
Reference 7520
Possible values 7520
Best practices 7520
Location 7521
Default values 7521
Operating system version differences 7521
Policy management 7521
Restart requirement 7521
Policy conflicts 7522
Group Policy 7522
Security considerations 7522
Vulnerability 7522
Countermeasure 7522
Potential impact 7522

Network access: Do not allow storage of passwords and credentials for network authentication 7522
Reference 7523

Possible values 7523
Best practices 7523
Location 7523
Default values 7523
Operating system version differences 7524
Policy management 7524
Restart requirement 7524
Group Policy 7524
Security considerations 7524
Vulnerability 7524
Countermeasure 7525
Potential impact 7525
Additional resources 7525

Network access: Let Everyone permissions apply to anonymous users 7525
Reference 7526
Possible values 7526
Best practices 7526
Location 7526
Default values 7526
Operating system version differences 7527
Policy management 7527
Restart requirement 7527
Security considerations 7527
Vulnerability 7527
Countermeasure 7527
Potential impact 7527

Network access: Named Pipes that can be accessed anonymously 7528
Reference 7528
Possible values 7528
Best practices 7528
Location 7528
Default values 7528
Operating system version differences 7529
Policy management 7529
Restart requirement 7529
Group Policy 7529
Security considerations 7529
Vulnerability 7529
Countermeasure 7530
Potential impact 7530

Network access: Remotely accessible registry paths 7530

Reference ... 7530
Possible values ... 7531
Best practices ... 7531
Location ... 7531
Default values ... 7531
Operating system version differences ... 7531
Policy management ... 7532
Restart requirement ... 7532
Security considerations ... 7532
Vulnerability ... 7532
Countermeasure ... 7532
Potential impact ... 7532

Network access: Remotely accessible registry paths and subpaths ... 7533
Reference ... 7533
Possible values ... 7533
Best practices ... 7533
Location ... 7533
Default values ... 7533
Operating system version differences ... 7534
Policy management ... 7534
Restart requirement ... 7534
Security considerations ... 7534
Vulnerability ... 7535
Countermeasure ... 7535
Potential impact ... 7535

Network access: Restrict anonymous access to Named Pipes and Shares ... 7535
Reference ... 7535
Possible values ... 7535
Best practices ... 7536
Location ... 7536
Default values ... 7536
Operating system version differences ... 7536
Policy management ... 7536
Restart requirement ... 7536
Security considerations ... 7537
Vulnerability ... 7537
Countermeasure ... 7537
Potential impact ... 7537

Network access: Shares that can be accessed anonymously ... 7538
Reference ... 7538
Possible values ... 7538

Best practices 7538
Location 7538
Default values 7538
Operating system version differences 7539
Policy management 7539
Restart requirement 7539
Security considerations 7539
Vulnerability 7539
Countermeasure 7539
Potential impact 7539

Network access: Sharing and security model for local accounts 7539
Reference 7540
Possible values 7540
Best practices 7540
Location 7541
Default values 7541
Operating system version differences 7541
Policy management 7541
Restart requirement 7541
Group Policy 7541
Security considerations 7542
Vulnerability 7542
Countermeasure 7542
Potential impact 7542

Network security: Allow Local System to use computer identity for NTLM 7542
Reference 7542
Possible values 7543
Location 7543
Default values 7543
Operating system version differences 7544
Policy management 7544
Restart requirement 7544
Policy conflict considerations 7544
Group Policy 7544
Security considerations 7545
Vulnerability 7545
Countermeasure 7545
Potential impact 7545

Network security: Allow LocalSystem NULL session fallback 7545
Reference 7546
Possible values 7546

Best practices 7546
Location 7547
Default values 7547
Operating system version differences 7547
Security considerations 7547
Vulnerability 7547
Countermeasure 7547
Potential impact 7548

Network Security: Allow PKU2U authentication requests to this computer to use online identities 7548
Reference 7548
Possible values 7548
Best practices 7549
Location 7549
Default values 7549
Operating system version differences 7549
Security considerations 7549
Vulnerability 7550
Countermeasure 7550
Potential impact 7550

Network security: Configure encryption types allowed for Kerberos 7550
Reference 7550
Possible values 7552
Best practices 7552
Location 7552
Default values 7552
Operating system version differences 7552
Security considerations 7553
Vulnerability 7553
Countermeasure 7553
Potential impact 7553

Network security: Do not store LAN Manager hash value on next password change 7553
Reference 7553
Possible values 7554
Best practices 7554
Location 7554
Default values 7554
Operating system version differences 7555
Policy management 7555
Restart requirement 7555
Security considerations 7555

Vulnerability 7555
Countermeasure 7555
Potential impact 7555

Network security: Force logoff when logon hours expire 7556
Reference 7556
Possible values 7556
Best practices 7556
Location 7556
Default values 7557
Operating system version differences 7557
Policy management 7557
Restart requirement 7557
Security considerations 7557
Vulnerability 7557
Countermeasure 7558
Potential impact 7558

Network security: LAN Manager authentication level 7558
Reference 7558
Possible values 7558
Best practices 7560
Location 7560
Default values 7560
Operating system version differences 7560
Policy management 7561
Restart requirement 7561
Group Policy 7561
Security considerations 7561
Vulnerability 7561
Countermeasure 7562
Potential impact 7562

Network security: LDAP client signing requirements 7562
Reference 7562
Possible values 7563
Best practices 7563
Location 7563
Default values 7563
Operating system version differences 7564
Policy management 7564
Restart requirement 7564
Group Policy 7564
Security considerations 7564

Vulnerability 7564
Countermeasure 7564
Potential impact 7564

Network security: Minimum session security for NTLM SSP based (including secure RPC) clients 7565
Reference 7565
Possible values 7565
Best practices 7565
Location 7565
Default values 7565
Operating system version differences 7566
Policy management 7566
Restart requirement 7566
Policy conflicts 7566
Security considerations 7566
Vulnerability 7566
Countermeasure 7566
Potential impact 7567

Network security: Minimum session security for NTLM SSP based (including secure RPC) servers 7567
Reference 7567
Possible values 7567
Best practices 7567
Location 7567
Default values 7568
Operating system version differences 7568
Policy management 7568
Restart requirement 7568
Policy dependencies 7568
Security considerations 7568
Vulnerability 7569
Countermeasure 7569
Potential impact 7569

Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication 7569
Reference 7569
Possible values 7569
Best practices 7570
Location 7570
Default values 7570
Operating system version differences 7570
Policy management 7570

Restart requirement 7571
Group Policy 7571
Auditing 7571
Security considerations 7571
Vulnerability 7571
Countermeasure 7571
Potential impact 7572

Network security: Restrict NTLM: Add server exceptions in this domain 7572
Reference 7572
Possible values 7572
Best practices 7572
Location 7573
Default values 7573
Operating system version differences 7573
Policy management 7573
Restart requirement 7573
Group Policy 7573
Auditing 7574
Security considerations 7574
Vulnerability 7574
Countermeasure 7574
Potential impact 7574

Network Security: Restrict NTLM: Incoming NTLM Traffic 7575
Reference 7575
Possible values 7575
Best practices 7575
Location 7575
Default values 7576
Operating system version differences 7576
Policy management 7576
Restart requirement 7576
Group Policy 7576
Auditing 7576
Security considerations 7577
Vulnerability 7577
Countermeasure 7577
Potential impact 7577

Network Security: Restrict NTLM: NTLM authentication in this domain 7577
Reference 7577
Possible values 7578
Best practices 7578

Location 7579
Default values 7579
Operating system version differences 7579
Policy management 7579
Restart requirement 7579
Group Policy 7579
Auditing 7579
Security considerations 7580
Vulnerability 7580
Countermeasure 7580
Potential impact 7580

Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers 7580
Reference 7581
Possible values 7581
Best practices 7581
Location 7581
Default values 7581
Operating system version differences 7582
Policy management 7582
Restart requirement 7582
Group Policy 7582
Auditing 7582
Security considerations 7582
Vulnerability 7583
Countermeasure 7583
Potential impact 7583

Network Security: Restrict NTLM: Audit Incoming NTLM Traffic 7583
Reference 7583
Possible values 7584
Best practices 7584
Location 7584
Default values 7584
Operating system version differences 7585
Policy management 7585
Restart requirement 7585
Group Policy 7585
Auditing 7585
Security considerations 7585
Vulnerability 7586
Countermeasure 7586
Potential impact 7586

Network Security: Restrict NTLM: Audit NTLM authentication in this domain 7586
Reference 7586
Possible values 7586
Best practices 7587
Location 7587
Default values 7587
Operating system version differences 7587
Policy management 7588
Restart requirement 7588
Group Policy 7588
Auditing 7588
Security considerations 7588
Vulnerability 7588
Countermeasure 7588
Potential impact 7589

Recovery console: Allow automatic administrative logon 7589
Reference 7589
Possible values 7589
Best practices 7589
Location 7589
Default values 7590
Operating system version differences 7590
Policy management 7590
Restart requirement 7590
Group Policy 7590
Policy conflicts 7590
Security considerations 7590
Vulnerability 7591
Countermeasure 7591
Potential impact 7591

Recovery console: Allow floppy copy and access to all drives and folders 7591
Reference 7591
Possible values 7591
Best practices 7592
Location 7592
Default values 7592
Operating system version differences 7592
Policy management 7592
Restart requirement 7592
Group Policy 7592
Policy conflicts 7593
Command-line tools 7593

Security considerations 7593
Vulnerability 7593
Countermeasure 7593
Potential impact 7593

Shutdown: Allow system to be shut down without having to log on 7593
Reference 7594
Possible values 7594
Best practices 7594
Location 7594
Default values 7594
Operating system version differences 7595
Policy management 7595
Restart requirement 7595
Group Policy 7595
Security considerations 7595
Vulnerability 7595
Countermeasure 7596
Potential impact 7596

Shutdown: Clear virtual memory pagefile 7596
Reference 7596
Possible values 7596
Best practices 7597
Location 7597
Default values 7597
Operating system version differences 7597
Policy management 7597
Restart requirement 7597
Security considerations 7598
Vulnerability 7598
Countermeasure 7598
Potential impact 7598

System cryptography: Force strong key protection for user keys stored on the computer - Explain text 7598
Settings page: Options 7598
Explain text 7599

System cryptography: Force strong key protection for user keys stored on the computer 7599
Reference 7599
Possible values 7599
Best practices 7600
Location 7600

Default values 7600
Operating system version differences 7600
Policy management 7600
Restart requirement 7600
Security considerations 7601
Vulnerability 7601
Countermeasure 7601
Potential impact 7601

System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing 7601
Reference 7601
Possible values 7602
Best practices 7602
Location 7602
Default values 7603
Operating system version differences 7603
Policy management 7604
Restart requirement 7604
Group Policy 7604
Security considerations 7604
Vulnerability 7604
Countermeasure 7604
Potential impact 7604

System objects: Require case insensitivity for non-Windows subsystems 7605
Reference 7605
Possible values 7605
Best practices 7605
Location 7606
Default values 7606
Operating system version differences 7606
Policy management 7606
Restart requirement 7606
Security considerations 7606
Vulnerability 7607
Countermeasure 7607
Potential impact 7607

System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links) 7607
Reference 7607
Possible values 7607
Best practices 7608
Location 7608

Default values 7608
Operating system version differences 7608
Policy management 7608
Restart requirement 7608
Security considerations 7608
Vulnerability 7609
Countermeasure 7609
Potential impact 7609

System settings: Optional subsystems 7609
Reference 7609
Possible values 7610
Best practices 7610
Location 7610
Default values 7610
Operating system version differences 7610
Policy management 7610
Restart requirement 7611
Security considerations 7611
Vulnerability 7611
Countermeasure 7611
Potential impact 7611

System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies 7611
Reference 7612
Possible values 7612
Best practices 7612
Location 7612
Default values 7612
Operating system version differences 7613
Policy management 7613
Restart requirement 7613
Security considerations 7613
Vulnerability 7613
Countermeasure 7613
Potential impact 7613

User Account Control: Admin Approval Mode for the Built-in Administrator account 7614
Reference 7614
Possible values 7614
Best practices 7614
Location 7614
Default values 7615

Operating system version differences 7615
Policy management 7615
Restart requirement 7615
Security considerations 7615
Vulnerability 7615
Countermeasure 7616
Potential impact 7616

User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop 7616
Reference 7616
Possible values 7617
Best practices 7618
Location 7618
Default values 7618
Operating system version differences 7618
Policy management 7618
Restart requirement 7619
Group Policy 7619
Policy interactions 7619
Security considerations 7619
Vulnerability 7619
Countermeasure 7620
Potential impact 7620

User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode 7620
Reference 7620
Possible values 7620
Best practices 7621
Location 7621
Default values 7621
Operating system version differences 7622
Policy management 7622
Restart requirement 7622
Group Policy 7622
Security considerations 7622
Vulnerability 7622
Countermeasure 7623
Potential impact 7623

User Account Control: Behavior of the elevation prompt for standard users 7623
Reference 7623
Possible values 7623

Best practices 7623
Location 7624
Default values 7624
Operating system version differences 7624
Policy management 7624
Restart requirement 7624
Group Policy 7624
Security considerations 7625
Vulnerability 7625
Countermeasure 7625
Potential impact 7625

User Account Control: Detect application installations and prompt for elevation 7625
Reference 7626
Possible values 7626
Best practices 7626
Location 7626
Default values 7626
Operating system version differences 7627
Policy management 7627
Restart requirement 7627
Security considerations 7627
Vulnerability 7627
Countermeasure 7627
Potential impact 7627

User Account Control: Only elevate executables that are signed and validated 7628
Reference 7628
Possible values 7628
Best practices 7628
Location 7629
Default values 7629
Operating system version differences 7629
Policy management 7629
Restart requirement 7629
Group Policy 7629
Security considerations 7630
Vulnerability 7630
Countermeasure 7630
Potential impact 7630

User Account Control: Only elevate UIAccess applications that are installed in secure locations 7630
Reference 7630

Possible values 7631
Best practices 7632
Location 7632
Default values 7632
Operating system version differences 7632
Policy management 7632
Restart requirement 7633
Group Policy 7633
Security considerations 7633
Vulnerability 7633
Countermeasure 7633
Potential impact 7633

User Account Control: Run all administrators in Admin Approval Mode 7634
Reference 7634
Possible values 7634
Best practices 7634
Location 7634
Default values 7634
Operating system version differences 7635
Policy management 7635
Restart requirement 7635
Group Policy 7635
Security considerations 7635
Vulnerability 7635
Countermeasure 7635
Potential impact 7636

User Account Control: Switch to the secure desktop when prompting for elevation 7636
Reference 7636
Possible values 7636
Best practices 7636
Location 7636
Default values 7637
Operating system version differences 7637
Policy management 7637
Restart requirement 7637
Group Policy 7637
Security considerations 7637
Vulnerability 7638
Countermeasure 7638
Potential impact 7638

User Account Control: Virtualize file and registry write failures to per-user locations 7638

Reference ... 7638
Possible values ... 7638
Best practices ... 7639
Location ... 7639
Default values ... 7639
Operating system version differences ... 7639
Policy management ... 7639
Restart requirement ... 7639
Group Policy ... 7639
Security considerations ... 7640
Vulnerability ... 7640
Countermeasure ... 7640
Potential impact ... 7640

Advanced Audit Policy Configuration ... 7640
See Also ... 7644

User Rights Assignment ... 7644

Access Credential Manager as a trusted caller ... 7646
Reference ... 7646
Possible values ... 7646
Best practices ... 7646
Location ... 7647
Default values ... 7647
Operating system version differences ... 7647
Policy management ... 7647
Group Policy ... 7647
Security considerations ... 7648
Vulnerability ... 7648
Countermeasure ... 7648
Potential impact ... 7648
See Also ... 7648

Access this computer from the network ... 7648
Reference ... 7648
Possible values ... 7649
Best practices ... 7649
Location ... 7649
Default values ... 7649
Operating system version differences ... 7650
Policy management ... 7650
Group Policy ... 7650
Security considerations ... 7650

Vulnerability 7651
Countermeasure 7651
Potential impact 7651
See Also 7651

Act as part of the operating system 7651
Reference 7652
Possible values 7652
Best practices 7652
Location 7652
Default values 7652
Operating system version differences 7653
Policy management 7653
Group Policy 7653
Security considerations 7653
Vulnerability 7653
Countermeasure 7653
Potential impact 7654
See Also 7654

Add workstations to domain 7654
Reference 7654
Possible values 7654
Best practices 7654
Location 7654
Default values 7655
Operating system version differences 7655
Policy management 7655
Group Policy 7655
Security considerations 7656
Vulnerability 7656
Countermeasure 7656
Potential impact 7656
See Also 7656

Adjust memory quotas for a process 7656
Reference 7657
Possible values 7657
Best practices 7657
Location 7657
Default values 7657
Operating system version differences 7658
Policy management 7658
Group Policy 7658

Security considerations 7659
Vulnerability 7659
Countermeasure 7659
Potential impact 7659
See Also 7659

Allow log on locally 7659
Reference 7659
Possible values 7660
Best practices 7660
Location 7660
Default values 7660
Operating system version differences 7661
Policy management 7661
Group Policy 7662
Security considerations 7662
Vulnerability 7662
Countermeasure 7662
Potential impact 7662
See Also 7663

Allow log on through Remote Desktop Services 7663
Reference 7663
Possible values 7663
Best practices 7663
Location 7663
Default values 7664
Operating system version differences 7664
Policy management 7664
Group Policy 7664
Security considerations 7665
Vulnerability 7665
Countermeasure 7665
Potential impact 7666
See Also 7666

Back up files and directories 7666
Reference 7666
Possible values 7667
Best practices 7667
Location 7667
Default values 7667
Operating system version differences 7668
Policy management 7668

Group Policy 7668
Security considerations 7668
Vulnerability 7668
Countermeasure 7668
Potential impact 7669
See Also 7669

Bypass traverse checking 7669
Reference 7669
Possible values 7669
Best practices 7669
Location 7669
Default values 7670
Operating system version differences 7671
Policy management 7671
Group Policy 7671
Security considerations 7671
Vulnerability 7671
Countermeasure 7672
Potential impact 7672
See Also 7672

Change the system time 7672
Reference 7672
Possible values 7673
Best practices 7673
Location 7673
Default values 7673
Operating system version differences 7674
Policy management 7674
Group Policy 7674
Security considerations 7674
Vulnerability 7674
Countermeasure 7675
Potential impact 7675
See Also 7675

Change the time zone 7675
Reference 7675
Possible values 7676
Best practices 7676
Location 7676
Default values 7676
Operating system version differences 7676

Policy management ... 7676
Group Policy ... 7677
Security considerations ... 7677
Vulnerability ... 7677
Countermeasure ... 7677
Potential impact ... 7677
See Also ... 7677

Create a pagefile ... 7677
Reference ... 7678
Possible values ... 7678
Best practices ... 7678
Location ... 7678
Default values ... 7678
Operating system version differences ... 7679
Policy management ... 7679
Group Policy ... 7679
Security considerations ... 7679
Vulnerability ... 7679
Countermeasure ... 7679
Potential impact ... 7679
See Also ... 7680

Create a token object ... 7680
Reference ... 7680
Possible values ... 7680
Best practices ... 7680
Location ... 7680
Default values ... 7680
Operating system version differences ... 7681
Policy management ... 7681
Group Policy ... 7681
Security considerations ... 7681
Vulnerability ... 7682
Countermeasure ... 7682
Potential impact ... 7682
See Also ... 7682

Create global objects ... 7682
Reference ... 7682
Possible values ... 7683
Best practices ... 7683
Location ... 7683
Default values ... 7683

Operating system version differences 7684
Policy management 7684
Group Policy 7684
Security considerations 7684
Vulnerability 7685
Countermeasure 7685
Potential impact 7685
See Also 7685

Create permanent shared objects 7685
Reference 7685
Possible values 7686
Best practices 7686
Location 7686
Default values 7686
Operating system version differences 7686
Policy management 7687
Group Policy 7687
Security considerations 7687
Vulnerability 7687
Countermeasure 7687
Potential impact 7687
See Also 7687

Create symbolic links 7688
Reference 7688
Possible values 7688
Best practices 7688
Location 7688
Default values 7688
Operating system version differences 7689
Policy management 7689
Group Policy 7689
Command-line tools 7689
Security considerations 7690
Vulnerability 7690
Countermeasure 7690
Potential impact 7690
See Also 7690

Debug programs 7690
Reference 7690
Possible values 7691
Best practices 7691

Location 7691
Default values 7691
Operating system version differences 7691
Policy management 7691
Group Policy 7692
Security considerations 7692
Vulnerability 7692
Countermeasure 7692
Potential impact 7692
See Also 7692

Deny access to this computer from the network 7693
Reference 7693
Possible values 7693
Best practices 7693
Location 7693
Default values 7693
Operating system version differences 7694
Policy management 7694
Group Policy 7694
Security considerations 7694
Vulnerability 7694
Countermeasure 7695
Potential impact 7695
See Also 7695

Deny log on as a batch job 7695
Reference 7695
Possible values 7696
Best practices 7696
Location 7696
Default values 7696
Operating system version differences 7696
Policy management 7697
Group Policy 7697
Security considerations 7697
Vulnerability 7697
Countermeasure 7697
Potential impact 7698
See Also 7698

Deny log on as a service 7698
Reference 7698
Possible values 7698

Best practices 7698
Location 7698
Default values 7699
Operating system version differences 7699
Policy management 7699
Group Policy 7699
Security considerations 7700
Vulnerability 7700
Countermeasure 7700
Potential impact 7700
See Also 7700

Deny log on locally 7700
Reference 7700
Possible values 7701
Best practices 7701
Location 7701
Default values 7701
Operating system version differences 7701
Policy management 7701
Group Policy 7702
Security considerations 7702
Vulnerability 7702
Countermeasure 7702
Potential impact 7702
See Also 7703

Deny log on through Remote Desktop Services 7703
Reference 7703
Possible values 7703
Best practices 7703
Location 7703
Default values 7703
Operating system version differences 7704
Policy management 7704
Group Policy 7704
Security considerations 7705
Vulnerability 7705
Countermeasure 7705
Potential impact 7705
See Also 7705

Enable computer and user accounts to be trusted for delegation 7705
Reference 7705

Possible values 7706
Best practices 7706
Location 7706
Default values 7706
Operating system version differences 7707
Policy management 7707
Group Policy 7707
Security considerations 7707
Vulnerability 7708
Countermeasure 7708
Potential impact 7708
See Also 7708

Force shutdown from a remote system 7708
Reference 7708
Possible values 7708
Best practices 7709
Location 7709
Default values 7709
Operating system version differences 7709
Policy management 7709
Group Policy 7710
Security considerations 7710
Vulnerability 7710
Countermeasure 7710
Potential impact 7710
See Also 7710

Generate security audits 7711
Reference 7711
Possible values 7711
Best practices 7711
Location 7711
Default values 7711
Operating system version differences 7712
Policy management 7712
Group Policy 7712
Security considerations 7713
Vulnerability 7713
Countermeasure 7713
Potential impact 7713
See Also 7713

Impersonate a client after authentication 7713

Reference ... 7713
Possible values ... 7714
Best practices ... 7714
Location ... 7714
Default values ... 7714
Operating system version differences ... 7715
Policy management ... 7715
Group Policy ... 7716
Security considerations ... 7716
Vulnerability ... 7716
Countermeasure ... 7716
Potential impact ... 7716
See Also ... 7716

Increase a process working set ... 7716
Reference ... 7717
Possible values ... 7717
Best practices ... 7717
Location ... 7717
Default values ... 7717
Operating system version differences ... 7718
Policy management ... 7718
Group Policy ... 7718
Security considerations ... 7718
Vulnerability ... 7718
Countermeasure ... 7718
Potential impact ... 7718
See Also ... 7719

Increase scheduling priority ... 7719
Reference ... 7719
Possible values ... 7719
Best practices ... 7719
Location ... 7719
Default values ... 7719
Operating system version differences ... 7720
Policy management ... 7720
Group Policy ... 7720
Security considerations ... 7720
Vulnerability ... 7721
Countermeasure ... 7721
Potential impact ... 7721
See Also ... 7721

Load and unload device drivers ... 7721
Reference ... 7721
Possible values ... 7722
Best practices ... 7722
Location ... 7722
Default values ... 7722
Operating system version differences ... 7722
Policy management ... 7722
Group Policy ... 7723
Security considerations ... 7723
Vulnerability ... 7723
Countermeasure ... 7723
Potential impact ... 7723
See Also ... 7724

Lock pages in memory ... 7724
Reference ... 7724
Possible values ... 7724
Best practices ... 7724
Location ... 7725
Default values ... 7725
Operating system version differences ... 7725
Policy management ... 7725
Group Policy ... 7725
Security considerations ... 7726
Vulnerability ... 7726
Countermeasure ... 7726
Potential impact ... 7726
See Also ... 7726

Log on as a batch job ... 7726
Reference ... 7726
Possible values ... 7727
Best practices ... 7727
Location ... 7727
Default values ... 7727
Operating system version differences ... 7728
Policy management ... 7728
Group Policy ... 7728
Security considerations ... 7728
Vulnerability ... 7728
Countermeasure ... 7728
Potential impact ... 7729
See Also ... 7729

Log on as a service 7729
Reference 7729
Possible values 7729
Best practices 7729
Location 7730
Default values 7730
Operating system version differences 7730
Policy management 7730
Group Policy 7730
Security considerations 7731
Vulnerability 7731
Countermeasure 7731
Potential impact 7731
See Also 7731

Manage auditing and security log 7732
Reference 7732
Possible values 7732
Best practices 7732
Location 7732
Default values 7732
Operating system version differences 7733
Policy management 7733
Group Policy 7733
Security considerations 7734
Vulnerability 7734
Countermeasure 7734
Potential impact 7734
See Also 7734

Modify an object label 7734
Reference 7734
Possible values 7735
Best practices 7735
Location 7735
Default values 7735
Operating system version differences 7736
Policy management 7736
Group Policy 7736
Security considerations 7736
Vulnerability 7736
Countermeasure 7737
Potential impact 7737
See Also 7737

Modify firmware environment values 7737
Reference 7737
Possible values 7738
Best practices 7738
Location 7738
Default values 7738
Operating system version differences 7738
Policy management 7739
Group Policy 7739
Security considerations 7739
Vulnerability 7739
Countermeasure 7739
Potential impact 7739
See Also 7739

Perform volume maintenance tasks 7740
Reference 7740
Possible values 7740
Best practices 7740
Location 7740
Default values 7740
Operating system version differences 7741
Policy management 7741
Group Policy 7741
Security considerations 7741
Vulnerability 7741
Countermeasure 7742
Potential impact 7742
See Also 7742

Profile single process 7742
Reference 7742
Possible values 7742
Best practices 7742
Location 7742
Default values 7743
Operating system version differences 7743
Policy management 7743
Group Policy 7743
Security considerations 7744
Vulnerability 7744
Countermeasure 7744
Potential impact 7744
See Also 7744

Profile system performance ... 7744
Reference ... 7744
Possible values ... 7745
Best practices ... 7745
Location ... 7745
Default values ... 7745
Operating system version differences ... 7745
Policy management ... 7745
Group Policy ... 7746
Security considerations ... 7746
Vulnerability ... 7746
Countermeasure ... 7746
Potential impact ... 7746
See Also ... 7746

Remove computer from docking station ... 7747
Reference ... 7747
Possible values ... 7747
Best practices ... 7747
Location ... 7747
Default values ... 7747
Operating system version differences ... 7748
Policy management ... 7748
Group Policy ... 7748
Security considerations ... 7748
Vulnerability ... 7749
Countermeasure ... 7749
Potential impact ... 7749
See Also ... 7749

Replace a process level token ... 7749
Reference ... 7749
Possible values ... 7750
Best practices ... 7750
Location ... 7750
Default values ... 7750
Operating system version differences ... 7751
Policy management ... 7751
Group Policy ... 7751
Security considerations ... 7751
Vulnerability ... 7751
Countermeasure ... 7752
Potential impact ... 7752
See Also ... 7752

Restore files and directories 7752
Reference 7752
Possible values 7752
Best practices 7753
Location 7753
Default values 7753
Operating system version differences 7753
Policy management 7754
Group Policy 7754
Security considerations 7754
Vulnerability 7754
Countermeasure 7754
Potential impact 7755
See Also 7755

Shut down the system 7755
Reference 7755
Possible values 7755
Best practices 7756
Location 7756
Default values 7756
Operating system version differences 7757
Policy management 7757
Group Policy 7757
Security considerations 7757
Vulnerability 7758
Countermeasure 7758
Potential impact 7758
See Also 7758

Synchronize directory service data 7758
Reference 7758
Possible values 7759
Best practices 7759
Location 7759
Default values 7759
Operating system version differences 7759
Policy management 7759
Group Policy 7760
Security considerations 7760
Vulnerability 7760
Countermeasure 7760
Potential impact 7760
See Also 7760

Take ownership of files or other objects 7761
Reference 7761
Possible values 7761
Best practices 7761
Location 7761
Default values 7761
Operating system version differences 7762
Policy management 7762
Group Policy 7762
Security considerations 7763
Vulnerability 7763
Countermeasure 7763
Potential impact 7763
See Also 7763

Administer Security Policy Settings 7763
Introduction 7763
What's changed in how settings are administered? 7764
Using the Local Security Policy snap-in 7765
Using the Secedit command-line tool 7766
Using the Security Compliance Manager 7766
Using the Security Configuration Wizard 7767
Working with the Security Configuration Manager 7768
Security Configuration and Analysis 7768
Security analysis 7768
Security configuration 7769
Security templates 7769
Security Settings extension to Group Policy 7770
Local Security Policy 7770
Using the Security Configuration Manager 7771
Applying security settings 7771
Importing and exporting security templates 7772
Analyzing security and viewing results 7772
Resolving security discrepancies 7773
Automating security configuration tasks 7774
Working with Group Policy tools 7774

Network List Manager Policies 7774
Policy settings for Network List Manager Policies 7775
Unidentified Networks 7775
Identifying Networks 7776
All Networks 7776

Smart Card Overview 7777

Feature description ... 7777
Practical applications ... 7777
See also ... 7777

What's New in Smart Cards ... 7778
Feature description ... 7778
New and changed functionality ... 7778
Virtual smart cards ... 7779
Changes to the smart card sign-in experience ... 7779
Smart Card Service start and stop behavior ... 7779
Smart card transactions ... 7779
Smart card support on Windows RT ... 7780
Smart card support in Windows 8 applications ... 7780
See also ... 7780

Software Restriction Policies ... 7780
Software Restriction Policies description ... 7781
Practical applications ... 7781
New and changed functionality ... 7781
Removed or deprecated functionality ... 7781
Software requirements ... 7782
Server Manager information ... 7782
See also ... 7782

Software Restriction Policies Technical Overview ... 7783
Introduction ... 7783
Procedures ... 7784
Software restriction policy usage scenarios ... 7784
Differences and changes in functionality ... 7785
System requirements ... 7788
Software restriction policies components and architecture ... 7788
Best practices ... 7789
Do not modify the default domain policy. ... 7789
Create a separate Group Policy Object for software restriction policies. ... 7789
If you experience problems with applied policy settings, restart Windows in Safe Mode. ... 7789
Use caution when defining a default setting of Disallowed. ... 7790
For best security, use access control lists in conjunction with software restriction policies. ... 7790
Test new policy settings thoroughly in test environments before applying the policy settings to your domain. ... 7790
Filter user policy settings based on membership in security groups. ... 7790
Do not link to a GPO in another domain or site. ... 7790
Additional resources ... 7791

Administer Software Restriction Policies 7791
Introduction 7791
To open Software Restriction Policies 7792
For your local computer 7792
For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain 7792
For a domain or organizational unit, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed 7793
For a site, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed 7793
To create new software restriction policies 7794
To add or delete a designated file type 7794
To prevent software restriction policies from applying to local administrators 7795
To change the default security level of software restriction policies 7795
To apply software restriction policies to DLLs 7796

Determine Allow/Deny List and Application Inventory for Software Restriction Policies 7796
Introduction 7797
What default rule to choose: Allow or Deny 7797
Create an inventory of your applications for the Allow list 7797

Work with Software Restriction Policies Rules 7798
Introduction 7798
Working with certificate rules 7799
Enabling certificate rules 7800
Set trusted publisher options 7802
Working with hash rules 7803
Working with Internet Zone rules 7804
Working with path rules 7805

Use Software Restriction Policies to Help Protect Your Computer Against an Email Virus 7806
Introduction 7806

Troubleshoot Software Restriction Policies 7807
Introduction 7807
Windows cannot open a program 7808
Modified software restriction policies are not taking effect 7808
After adding a rule through SRP, you cannot log on to your computer 7809
A new policy setting is not applying to a specific file name extension 7809
A default rule is not restricting as expected 7809
Unable to discover which restrictions are applied 7809

TLS/SSL (Schannel SSP) Overview 7810
TLS/SSL (Schannel) description 7810
Practical applications 7810

New and changed functionality 7811
Deprecated functionality 7811
Software requirements 7811
Server Manager information 7811
See also 7812

What's New in TLS/SSL (Schannel SSP) 7812
TLS/SSL (Schannel SSP) description 7812
New and changed functionality in Windows Server 2012 R2 7813
TLS session resumption 7813
Application Protocol Negotiation 7813
New and changed functionality in Windows Server 2012 7814
Management of trusted issuers for client authentication 7814
TLS support for Server Name Indicator (SNI) Extensions 7816
Datagram Transport Layer Security (DTLS) 7816
Deprecated functionality 7817
See also 7817

Trusted Platform Module Technology Overview 7817
Feature description 7818
Practical applications 7818
New and changed functionality 7819
Automated provisioning and management of the TPM 7819
Measured Boot with support for attestation 7820
TPM-based Virtual Smart Card 7820
TPM-based certificate storage 7820
TPM Owner Authorization Value 7820
TPM Cmdlets 7821
Supported versions 7821
Additional Resources 7822

Schema Extensions for Windows Server 2008 R2 to support AD DS backup of TPM information from Windows 8 clients 7822
Why a schema extension is needed 7822
TpmSchemaExtension.ldf 7823
TpmSchemaExtensionACLChanges.ldf 7829

User Account Control Overview 7831
Feature description 7832
Practical applications 7832
New and changed functionality 7832

How User Account Control Works 7832
UAC Process and Interactions 7833
Windows Server® 2012 Logon Process 7833

The UAC User Experience ... 7834
UAC Architecture ... 7838
Virtualization ... 7843
Request Execution Levels ... 7844
Installer Detection Technology ... 7844

Windows Authentication Overview ... 7845
Feature description ... 7845
Practical applications ... 7846
New and changed functionality ... 7849
Removed or deprecated functionality ... 7850
Software requirements ... 7850
Server Manager information ... 7850

Windows Server 2012 R2 Software Inventory Logging Overview ... 7850
Feature description ... 7851
Practical applications ... 7851
Software requirements ... 7852
See Also ... 7852

Manage Software Inventory Logging in Windows Server 2012 R2 ... 7852
Starting and Stopping Software Inventory Logging ... 7853
Displaying Software Inventory Logging data ... 7854
Forwarding Software Inventory Logging data ... 7856
Deleting data logged by Software Inventory Logging ... 7857
Backing up and restoring data logged by Software Inventory Logging ... 7857
Configuring certificates for Software Inventory Logging ... 7858
Enabling Software Inventory using unattended installation ... 7859
Creating a master VHD image to deploy Software Inventory Logging to servers in an enterprise ... 7859
See also ... 7860

Telemetry Overview ... 7860
Feature description ... 7860
See also ... 7860

Deploy Windows Feedback Forwarder ... 7861
Participate in the Customer Experience Improvement Program ... 7861
Participate in Windows Error Reporting ... 7862
Configure all servers to participate in CEIP and WER ... 7864
Install Windows Feedback Forwarder ... 7864
Install Windows Feedback Forwarder by using Add Roles and Features Wizard ... 7865
Install Windows Feedback Forwarder on a Server Core installation ... 7867
Configure Windows Feedback Forwarder ... 7867

What's New in Telemetry ... 7868
Feature description ... 7868
New and changed functionality ... 7869
Windows Feedback Forwarder ... 7870
Windows Automatic Feedback ... 7870
See also ... 7870

User Access Logging Overview ... 7870
Feature description ... 7871
Practical applications ... 7871
Important functionality ... 7872
Data logged with UAL ... 7873
Software requirements ... 7874
See also ... 7874

Manage User Access Logging ... 7875
Disabling and enabling the UAL service ... 7875
Collecting UAL data ... 7877
Deleting data logged by UAL ... 7879
Managing UAL in high volume environments ... 7879
Recovering from a corrupt state ... 7880
Enable Work Folders usage license tracking ... 7880
See also ... 7880

Volume Activation Overview ... 7881
Volume activation description ... 7881
How does volume activation work? ... 7882
Practical applications ... 7884
What is activation? ... 7884
What are my activation options? ... 7884
What happens if systems are not activated? ... 7885
When should I use volume activation? ... 7885
System requirements ... 7886
Additional information ... 7886

Test Lab Guide: Demonstrate Volume Activation Services ... 7887
Setting up Volume Activation Services in a test environment ... 7887
Step 1: Set up the base test lab infrastructure ... 7887
Step 2: Install the Volume Activation Services server role ... 7888
Step 3: Configure Active Directory-based Activation ... 7889
Step 4: Verify that Active Directory-based Activation works ... 7890
Step 5: Configure Key Management Services ... 7890
Step 6: Verify that KMS volume activation works ... 7891
Additional information ... 7892

Plan for Volume Activation 7892

Step 1: Review and Select Activation Methods 7892
1.1 Plan for Active Directory-based Activation 7893
Using Active Directory-based Activation in mixed Windows activation environments 7894
Utilizing Active Directory-based Activation failover capability 7894
1.2 Plan for Key Management Services activation 7895
Addressing KMS operational requirements 7896
Addressing KMS functional requirements 7897
1.3 Plan for Multiple Activation Key activation 7897
See also 7898

Step 2: Evaluate Client Connectivity 7898
Activation scenarios 7899
2.1. Identify scenario requirements for the central network 7900
2.2. Identify scenario requirements for the branch office network 7901
2.3. Identify high-security zone scenario requirements 7902
2.4. Identify disconnected computer scenario requirements 7903
2.5 Identify test and development lab requirements 7903
See also 7904

Step 3: Determine Activation Method and Product License Requirements 7904
See also 7905

Step 4: Determine Monitoring and Reporting Needs 7906
4.1 Managing Active Directory-based Activation objects 7906
4.2 Managing Activation with System Center Configuration Manager 7907
4.3 Managing Activation by Using Event Logs 7907
4.4 Managing Activation by Using the Volume Activation Management Tool 7908
See also 7908

Appendix A: KMS Client Setup Keys 7908
Windows Server 2012 R2 and Windows 8.1 Client Setup Keys 7908
Windows Server 2012 and Windows 8 Client Setup Keys 7909
Windows 7 and Windows Server 2008 R2 7909
Windows Vista and Windows Server 2008 7910

Deploy Volume Activation 7911
Methods of volume activation 7911

Deploy Active-Directory-based Activation 7911
Using Active Directory-based activation 7912
Managing the activation object 7912
Appendix A: Administration delegation script 7913

Deploy KMS Activation 7916

Configuring KMS hosts 7916
Configuring DNS 7916
Change the Default DNS Permissions for SRV Records 7916
Publish to Multiple DNS Domains 7917
Manually Create SRV Records in DNS 7918
Installing KMS hosts 7919
Configuring KMS clients 7920

Deploy MAK Activation 7921
Installing a MAK During Operating System Installation 7921

Volume Activation Technical Reference 7923
Technical Reference for Volume Activation 7923
See also 7924

Slmgr.vbs Options for Volume Activation 7924
See also 7932

Registry Settings for Volume Activation 7932
Registry settings 7932
See also 7935

KMS Activation Timing and Discovery for Volume Activation 7935
On this page 7935
KMS activation timing 7935
KMS discovery search order 7936
See also 7936

KMS and MAK Activation Scenarios for Volume Activation 7936
On this page 7937
KMS scenarios 7937
Default KMS implementation for a single-site network 7937
KMS implementation in a complex, global network 7938
MAK scenarios 7939
MAK independent activation 7940
MAK proxy activation 7942
VAMT support for KMS activation 7947
Converting KMS to MAK activation 7948
Installing a MAK during operating system installation 7948
Installing a MAK after operating system installation 7948
See also 7949

KMS Host Keys to Products Activated for Volume Activation 7949
See also 7951

Activation Policy Values for Volume Activation 7951

See also 7952

KMS Client Setup Keys for Volume Activation 7952
See also 7952

Activation Error Codes for Volume Activation 7952
See also 7958

WMI Properties and Methods for Volume Activation 7958
See also 7967

Web Server (IIS) Overview 7967
Role Description 7967
Practical Applications 7968
New and Changed Functionality 7968
Deprecated Functionality 7969
See Also 7970

Hosting-Friendly Web Server Platform (IIS): Scenario Overview 7970
Scenario Description 7970
Web Server Scenarios 7971
Practical Applications 7972
See Also 7972

Build a Static Website on IIS 7973
Prerequisites 7973
Step 1: Install the IIS Web Server 7973
Step 2: Add a Website 7975
Step 3: Configure Anonymous Authentication 7976
Step 4: Configure the Default Documents 7977
Step 5: Configure Static Content Compression 7978
Next Steps 7979
See also 7979

Configure Logging in IIS 7979
Prerequisites 7979
Configure Logging at the Site Level 7980
Configure Per-site Logging at the Server Level 7981
Configure Per-server Logging at the Server Level 7982
Select W3C Fields to Log 7982
Configure Log File Rollover Options 7983
See Also 7984

Configure Request Filtering in IIS 7984
Prerequisites 7985
General Request Filter Settings 7985

File Name Extensions 7987
Filtering Rules 7988
Hidden Segments 7988
URL Filtering 7989
HTTP Verbs 7989
Header Size Limits 7990
Query Strings 7991
Request Filter Logging 7991
See Also 7992

Build a Classic ASP Website on IIS 7992
Prerequisites 7992
Step 1: Install the IIS Web Server 7993
Step 2: Add a Classic ASP Website 7994
Step 3: Edit ASP Application Settings 7996
Next Steps 8002
See also 8002

Build an ASP.NET Website on IIS 8003
Scenario Description 8003
In This Scenario 8003
Practical Applications 8004
Software Requirements 8004
See Also 8004

Plan an ASP.NET Website on IIS 8004

Step 1: Plan IIS Web Server and ASP.NET Modules Installation 8005
1.1. Plan to Install IIS and ASP.NET Modules 8005
1.2. Plan to Add the ASP.NET Application 8005
See Also 8006

Step 2: Plan ASP.NET Settings 8006
2.1. Session State Settings 8006
Store session state in process 8007
Store session state by using state server 8008
Store session state by using SQL server 8008
Cookie mode for session state 8009
2.2. Pages and Controls Settings 8010
2.3. Application Settings 8011
2.4. .NET Compilation Settings 8011
2.5. .NET Globalization Settings 8012

Step 3: Plan Data Source Settings 8012
3.1. Data source connection strings 8012

3.2. ASP.NET providers ... 8013
3.3. .NET profiles ... 8014
3.4. .NET roles ... 8014
3.5. .NET users ... 8014

Step 4: Plan Application Security ... 8015
4.1. Isolate Web Applications ... 8015
4.2. .NET Trust Levels ... 8016
4.3. .NET Authentication ... 8016
ASP.NET Forms Authentication ... 8017
Forms authentication basics ... 8017
Authentication cookies ... 8018
ASP.NET Impersonation Authentication ... 8020
4.4. Machine Key Settings ... 8020
4.5. TLS/SSL Communication ... 8021
Server Certificates ... 8021
SSL Binding ... 8022
Require SSL for Your Site ... 8022
Client Certificates ... 8022

Configure an ASP.NET Website on IIS ... 8023

Step 1: Install IIS and ASP.NET Modules ... 8023
Installing IIS and ASP.NET Modules ... 8023
Adding the ASP.NET Application ... 8026
See Also ... 8027

Step 2: Configure ASP.NET Settings ... 8027
2.1. Session State Settings ... 8027
Store Session State in Process ... 8028
Store Session State by using State Server ... 8029
Store Session State by using SQL Server ... 8030
Cookie Mode for Session State ... 8031
2.2. Pages and Controls Settings ... 8033
Edit Pages and Controls ... 8033
Add a Custom Control ... 8034
2.3. Application Settings ... 8035
2.4. .NET Compilation Settings ... 8035
2.5. .NET Globalization Settings ... 8038

Step 3: Configure Data Source Settings ... 8039
3.1. Data Source Connection Strings ... 8039
3.2. ASP.NET Providers ... 8041
3.3. .NET Profiles ... 8044
3.4. .NET Roles ... 8046

3.5. .NET Users ... 8046

Step 4: Configure Application Security ... 8047
4.1. Isolate Web Applications ... 8047
4.1. .NET Trust Levels ... 8048
4.2. .NET Authentication ... 8049
ASP.NET Forms Authentication ... 8049
ASP.NET Impersonation Authentication ... 8052
4.3. Machine Key Settings ... 8053
4.4. TLS/SSL Communication ... 8053
SSL Binding ... 8053
Require SSL for Your Site ... 8054
Client Certificates ... 8055

Build an FTP Site on IIS ... 8056
Prerequisites ... 8057
Step 1: Install FTP on an Existing IIS Web Server ... 8057
Step 2: Add an FTP Site ... 8058
Step 3: Configure FTP Site Defaults ... 8059
Step 4: Configure Firewall Support ... 8061
Step 5: Configure User Isolation ... 8061
Step 6: Configure Directory Browsing Options ... 8063
Step 7: Configure Logon Attempt Restrictions ... 8063
Step 8: Configure Request Filtering ... 8064
Step 9: Configure FTP Logging ... 8064
Step 10: Configure FTP Messages ... 8065
See Also ... 8066

IIS Manager Feature Overview ... 8066
IIS manager features ... 8066

.NET Authorization Rules ... 8071
UI Elements for .NET Authorization Rules ... 8071
Feature Page Elements ... 8071
Actions Pane Elements ... 8072
Allow or Deny Authorization Rule Dialog Boxes ... 8073

.NET Compilation ... 8074
UI Elements for .NET Compilation ... 8074
Feature Page Elements ... 8074
Actions Pane Elements ... 8076

.NET Error Pages ... 8076
UI Elements for .NET Error Pages ... 8077
Feature Page Elements ... 8077

Actions Pane Elements ... 8077
Add or Edit Custom Error Page Dialog Box ... 8078
Edit ASP.NET Error Pages Settings Dialog Box ... 8078

.NET Globalization ... 8079
UI Elements for .NET Globalization ... 8079
Feature Page Elements ... 8080
Actions Pane Elements ... 8081

.NET Profile ... 8081
UI Elements for .NET Profile ... 8082
Feature Page Elements ... 8082
Actions Pane Elements ... 8082
Add or Edit .NET Profile Property Dialog Box ... 8083
Add or Edit Group Dialog Box ... 8084
Edit Profile Settings Dialog Box ... 8084

.NET Roles ... 8085
UI Elements for .NET Roles ... 8085
Feature Page Elements ... 8085
Actions Pane Elements ... 8085
Add .NET Role Dialog Box ... 8086
Edit .NET Roles Settings Dialog Box ... 8086

.NET Trust Levels ... 8087
UI Elements for .NET Trust Levels ... 8087
Feature Page Elements ... 8087
Actions Pane Elements ... 8088

.NET Users ... 8088
UI Elements for .NET Users ... 8089
Feature Page Elements ... 8089
Actions Pane Elements ... 8089
Edit .NET User Dialog Box ... 8090
Add .NET User Roles ... 8090
Add .NET User Account Details ... 8090
Edit .NET Users Settings Dialog Box ... 8091

Application Pools ... 8091
UI Elements for Application Pools ... 8092
Feature Page Elements ... 8092
Actions Pane Elements ... 8093
Add or Edit Application Pool Dialog Box ... 8094
Application Pool Identity Dialog Box ... 8096
Recycling Conditions Wizard ... 8097

Recycling Events to Log Wizard 8099

Application Settings 8100
UI Elements for Application Settings 8101
Feature Page Elements 8101
Actions Pane Elements 8101
Add or Edit Application Setting Dialog Box 8102

Applications 8102
UI Elements for Applications 8102
Feature Page Elements 8103
Actions Pane Elements 8103
Add or Edit Application Dialog Box 8104

ASP 8106
UI Elements for ASP 8106
Feature Page Elements 8106
Actions Pane Elements 8110

Authentication 8111
UI Elements for Authentication 8111
Active Directory Client Certificate Authentication 8113
Feature Page Elements 8114
Actions Pane Elements 8114
Anonymous Authentication 8114
Feature Page Elements 8114
Actions Pane Elements 8115
Edit Anonymous Authentication Credentials Dialog Box 8115
Set Credentials Dialog Box (Anonymous) 8116
Basic Authentication 8117
Feature Page Elements 8117
Actions Pane Elements 8117
Edit Basic Authentication Settings Dialog Box 8117
Digest Authentication 8118
Feature Page Elements 8118
Actions Pane Elements 8118
Edit Digest Authentication Settings Dialog Box 8119
Forms Authentication 8119
Feature Page Elements 8119
Actions Pane Elements 8119
Edit Forms Authentication Settings Dialog Box 8120
Windows Authentication 8122
Feature Page Elements 8122
Actions Pane Elements 8122

Advanced Settings Dialog Box (Windows) ... 8123
ASP.NET Impersonation ... 8123
Feature Page Elements ... 8123
Actions Pane Elements ... 8123
ASP.NET Impersonation Settings Dialog Box ... 8124
Set Credentials Dialog Box (Impersonation) ... 8124

Authorization Rules ... 8124
UI Elements for Authorization Rules ... 8125
Feature Page Elements ... 8125
Actions Pane Elements ... 8126
Allow Authorization Rule Dialog Boxes ... 8126

Centralized Certificates ... 8127
UI Elements for Centralized Certificates ... 8127
Feature Page Elements ... 8127
Actions Pane Elements ... 8128
Edit Centralized Certificates Settings Dialog Box ... 8128

CGI ... 8129
UI Elements for CGI ... 8129
Feature Page Elements ... 8129
Actions Pane Elements ... 8129

Compression ... 8130
UI Elements for Compression ... 8130
Feature Page Elements ... 8130
Actions Pane Elements ... 8132

Configuration Editor ... 8132
UI Elements for Configuration Editor ... 8132
Feature Page Elements ... 8133
Actions Pane Elements ... 8133
Configuration Search Dialog Box ... 8135
Collection Editor Dialog Box ... 8135

Connection Strings ... 8136
UI Elements for Connection Strings ... 8136
Feature Page Elements ... 8136
Actions Pane Elements ... 8136
Add or Edit Connection String Dialog Box ... 8137
Set Credentials Dialog Box ... 8138

Default Document ... 8138
UI Elements Default Documents ... 8139

Feature Page Elements ... 8139
Actions Pane Elements ... 8139
Add Default Document Dialog Box ... 8141

Directory Browsing ... 8141
UI Elements for Directory Browsing ... 8141
Feature Page Elements ... 8141
Actions Pane Elements ... 8142

Error Pages ... 8142
UI Elements for Error Pages ... 8143
Feature Page Elements ... 8143
Actions Pane Elements ... 8143
Add or Edit Custom Error Page Dialog Box ... 8144
Edit Error Pages Settings Dialog Box ... 8145
Set Localized Custom Error Path Dialog Box ... 8146

Failed Request Tracing Rules ... 8147
UI Elements for Failed Request Tracing Rules ... 8147
Feature Page Elements ... 8147
Actions Pane Elements ... 8148
Specify Content to Trace Wizard ... 8150
Define Trace Conditions Wizard ... 8150
Select Trace Providers Wizard ... 8151

FastCGI Settings ... 8154
UI Elements for FastCGI Settings ... 8155
Feature Page Elements ... 8155
Actions Pane Elements ... 8155
Add or Edit FastCGI Application Dialog Box ... 8156
Environment Variables Collection Editor ... 8159

Feature Delegation ... 8160
Feature Delegation Overview ... 8160
UI Elements for Feature Delegation ... 8161
Feature Page Elements ... 8161
Actions Pane Elements ... 8162
Custom Site or Application Delegation ... 8165
Feature Page Elements ... 8166
Actions Pane Elements ... 8166
Copy Delegation Dialog Box ... 8168

FTP Authentication ... 8168
UI Elements for FTP Authentication ... 8169
Feature Page Elements ... 8169

Actions Pane Elements ... 8170
Note ... 8171
Custom Providers Dialog Box ... 8171
Add or Edit Custom Authentication Provider Dialog Box ... 8172
Edit Anonymous Authentication Credentials Dialog Box ... 8173
Edit Basic Authentication Settings ... 8173

FTP Authorization Rules ... 8173
UI Elements for FTP Authorization Rules ... 8174
Feature Page Elements ... 8174
Actions Pane Elements ... 8174
Note ... 8175
Allow and Deny Authorization Rule Dialog Boxes ... 8175

FTP Directory Browsing ... 8176
UI Elements for FTP Directory Browsing ... 8176
Feature Page Elements ... 8176
Actions Pane Elements ... 8177

FTP Firewall Support ... 8178
UI Elements for FTP Firewall Support ... 8178
Feature Page Elements ... 8178
Actions Pane Elements ... 8179

FTP IP Address and Domain Restrictions ... 8179
UI Elements for IP Address and Domain Restrictions ... 8179
Feature Page Elements ... 8179
Actions Pane Elements ... 8180
Add Allow or Deny Restriction Rule Dialog Box ... 8182
Edit IP and Domain Restrictions Dialog Box ... 8183

FTP Logging ... 8184
UI Elements for FTP Logging ... 8184
Feature Page Elements ... 8184
Actions Pane Elements ... 8186
Information To Log Dialog Box ... 8186

FTP Logon Attempt Restrictions ... 8189
UI Elements for FTP Logon Attempt Restrictions ... 8189
Feature Page Elements ... 8189
Actions Pane Elements ... 8189

FTP Messages ... 8190
UI Elements for FTP Messages ... 8190
Feature Page Elements ... 8190

Actions Pane Elements ... 8192
Note ... 8192

FTP Network Security ... 8192
UI Elements for FTP Network Security ... 8192
Feature Page Elements ... 8192
Action Pane Elements ... 8193

FTP Request Filtering ... 8193
UI Element List ... 8194
Feature Page Elements ... 8194
Actions Pane Elements ... 8194
File Name Extensions ... 8195
Feature Page Elements ... 8196
Actions Pane Elements ... 8196
Allow or Deny File Name Extension Dialog Box ... 8196
Hidden Segments ... 8197
Feature Page Elements ... 8197
Actions Pane Elements ... 8197
Add Hidden Segment Dialog Box ... 8197
Denied URL Sequences ... 8198
Feature Page Elements ... 8198
Actions Pane Elements ... 8198
Add Deny Sequence Dialog Box ... 8198
Edit FTP Request Filtering Settings Dialog Box ... 8199
FTP Commands ... 8199
Feature Page Elements ... 8200
Actions Pane Elements ... 8200
Allow or Deny Command Dialog Box ... 8201

FTP Sites ... 8201
Add FTP Site Wizard ... 8201
Site Information Page ... 8202
Binding and SSL Settings Page ... 8202
Authentication and Authorization Information Page ... 8203
FTP Site Defaults Dialog Box ... 8203
FTP Site Advanced Settings Dialog Box ... 8205
Add FTP Site Publishing Wizard ... 8207
Binding and SSL Settings Page ... 8207
Authentication and Authorization Information Page ... 8208
FTP Current Sessions Page ... 8209
Feature Page Elements ... 8209
Actions Pane Elements ... 8209

FTP SSL Settings 8209
UI Elements for FTP SSL Settings 8210
Feature Page Elements 8210
Actions Pane Elements 8211
Advanced SSL Policy Dialog Box 8211

FTP User Isolation 8212
UI Elements for FTP User Isolation 8213
Feature Page Elements 8213
Actions Pane Elements 8216
Set Credentials Dialog Box 8217

Handler Mappings 8217
UI Elements for Handler Mappings 8218
Feature Page Elements 8218
Actions Pane Elements 8219
Add or Edit Managed Handler Dialog Box 8221
Add or Edit Script Map Dialog Box 8222
Add or Edit Wildcard Script Map Dialog Box 8223
Add or Edit Module Mapping Dialog Box 8224
Request Restrictions Dialog Box 8225
Edit Feature Permissions Dialog Box 8227

HTTP Redirect 8229
UI Elements for HTTP Redirect 8229
Feature Page Elements 8229
Actions Pane Elements 8230

HTTP Response Headers 8231
UI Elements for HTTP Response Headers 8231
Feature Page Elements 8231
Actions Pane Elements 8232
Add or Edit Custom HTTP Response Header Dialog Box 8232
Set Common HTTP Response Headers Dialog Box 8233

IIS Manager 8234
Start Page 8234
Home Page 8235
Content View Page 8236
Feature Page Elements 8236
Actions Pane Elements 8236
Defaults Dialog Boxes 8238
Advanced Settings Dialog Boxes 8239
New Feature Available Dialog Boxes 8239
Change .NET Framework Version Dialog Box 8240

Create New Connection ... 8240
Specify Server Connection Details Wizard Page ... 8241
Specify Site Connection Details Wizard Page ... 8241
Specify Application Connection Details Wizard Page ... 8242
Provider Credentials Wizard Page ... 8243
Specify a Connection Name Wizard Page ... 8243
Certificate Name Mismatch Help Identifier Dialog Box ... 8244
Certificate Verification Help Identifier Dialog Box ... 8244

IIS Manager Permissions ... 8244
UI Elements for IIS Manager Permissions ... 8245
Feature Page Elements ... 8245
Actions Pane Elements ... 8246
Allow User Dialog Box ... 8246
Users Dialog Box ... 8247

IIS Manager Users ... 8247
UI Elements for IIS Manager Users ... 8248
Feature Page Elements ... 8248
Actions Pane Elements ... 8248
Add User and Change Password Dialog Boxes ... 8249

IP Address and Domain Restrictions ... 8249
UI Elements for IP Address and Domain Restrictions ... 8250
Feature Page Elements ... 8250
Actions Pane Elements ... 8251
Add Allow or Add Deny Restriction Rule Dialog Boxes ... 8252
Edit IP and Domain Restrictions Dialog Box ... 8254
Dynamic IP Restriction Settings Dialog Box ... 8255

ISAPI and CGI Restrictions ... 8255
UI Elements for ISAPI and CGI Restrictions ... 8256
Feature Page Elements ... 8256
Actions Pane Elements ... 8256
Add or Edit ISAPI or CGI Restriction Dialog Box ... 8257
Edit ISAPI and GCI Restrictions Settings Dialog Box ... 8257

ISAPI Filters ... 8258
UI Elements for ISAPI Filters ... 8258
Feature Page Elements ... 8258
Actions Pane Elements ... 8259
Add or Edit ISAPI Filter Dialog Box ... 8260

Logging ... 8260
UI Elements for Logging ... 8261

Feature Page Elements ... 8261
Actions Pane Elements ... 8266
W3C Logging Fields Dialog Box ... 8267
Enhanced Logging with Custom Fields ... 8268

Machine Key ... 8269
UI Elements for Machine Key ... 8269
Feature Page Elements ... 8269
Actions Pane Elements ... 8272

Management Service ... 8272
UI Elements for Management Service ... 8273
Feature Page Elements ... 8273
Actions Pane Elements ... 8274
Add Allow or Deny Connection Rule Dialog Box ... 8275

MIME Types ... 8276
UI Elements for MIME Types ... 8276
Feature Page Elements ... 8276
Actions Pane Elements ... 8276
Add or Edit MIME Type Dialog Box ... 8277

Modules ... 8277
UI Elements for Modules ... 8278
Feature Page Elements ... 8278
Actions Pane Elements ... 8278
Add or Edit Managed Module Dialog Box ... 8280
Configure Native Modules Dialog Box ... 8281
Register Native Module or Edit Native Module Registration Dialog Box ... 8282

Output Caching ... 8283
UI Elements for Output Caching ... 8283
Feature Page Elements ... 8283
Actions Pane Elements ... 8284
Add or Edit Output Cache Rule Dialog Box ... 8285
Edit Output Cache Settings Dialog Box ... 8288
Advanced Output Caching Settings Dialog Box ... 8288

Pages and Controls ... 8289
UI Elements for Pages and Controls ... 8289
Feature Page Elements ... 8290
Actions Pane Elements ... 8291
Controls Page ... 8291
Feature Page Elements ... 8291
Actions Pane Elements ... 8292

Add or Edit Custom Control and Add or Edit User Control Dialog Boxes 8292

Providers 8293
UI Elements for Providers 8293
Feature Page Elements 8293
Actions Pane Elements 8294
Add or Edit Provider Dialog Box 8294

Request Filtering 8296
UI Elements for Request Filtering 8297
Feature Page Elements 8297
Actions Pane Elements 8298
Edit Request Filtering Settings Dialog Box 8298
File Name Extensions Tab 8299
Feature Page Elements 8299
Actions Pane Elements 8299
Allow or Deny File Name Extension Dialog Boxes 8300
Rules Tab 8300
Feature Page Elements 8300
Actions Pane Elements 8300
Add Filtering Rule Dialog Box 8301
Hidden Segments Tab 8301
Feature Page Elements 8301
Actions Pane Elements 8302
Add Hidden Segment Dialog Box 8302
URL Tab 8302
Feature Page Elements 8302
Actions Pane Elements 8303
Allow URL Dialog Box 8303
Add Deny Sequence Dialog Box 8303
HTTP Verbs Tab 8303
Feature Page Elements 8303
Actions Pane Elements 8304
Allow or Deny Verb Dialog Boxes 8304
Headers Tab 8304
Feature Page Elements 8305
Actions Pane Elements 8305
Add Header Dialog Box 8305
Query Strings Tab 8305
Feature Page Elements 8305
Actions Pane Elements 8306
Allow or Deny Query String Dialog Boxes 8306

Server Certificates 8306

UI Elements for Server Certificates ... 8307
Feature Page Elements ... 8307
Actions Pane Elements ... 8308
Import Certificate Dialog Box ... 8309
Request Certificate Wizard ... 8309
Distinguished Name Properties Wizard Page ... 8309
Cryptographic Service Provider Properties Wizard Page ... 8310
File Name Wizard Page ... 8311
Complete Certificate Request Dialog Box ... 8311
Create Certificate Wizard ... 8312
Distinguished Name Properties Wizard Page ... 8312
Online Certification Authority Wizard Page ... 8312
Select Certification Authority Dialog Box ... 8313
Create Self-Signed Certificate Dialog Box ... 8313
Export Certificate Dialog Box ... 8314
Renew an Existing Certificate Wizard ... 8314

Session State ... 8315
UI Elements for Session State ... 8315
Feature Page Elements ... 8315
Actions Pane Elements ... 8318

Shared Configuration ... 8318
UI Elements for Shared Configuration ... 8319
Feature Page Elements ... 8319
Actions Pane Elements ... 8319
Export Configuration Dialog Box ... 8320
Set Credentials Dialog Box ... 8320
Encryption Keys Password Dialog Box ... 8321

Sites ... 8321
UI Elements for Sites ... 8322
Feature Page Elements ... 8322
Actions Pane Elements ... 8322
Add Website Dialog Box ... 8324
Edit Site Dialog Box ... 8328
Select Application Pool Dialog Box ... 8328
Connect As Dialog Box ... 8329
Set Credentials Dialog Box ... 8330
Site Bindings Dialog Box ... 8330
Add or Edit Site Binding Dialog Box ... 8331
Advanced Settings ... 8334
Edit Website Limits Dialog Box ... 8334
Edit Website Failed Request Tracing Settings Dialog Box ... 8335

SMTP E-mail ... 8336
UI Elements for SMTP E-mail ... 8336
Feature Page Elements ... 8336
Actions Pane Elements ... 8337
Set Credentials Dialog Box ... 8338

SSL Settings ... 8338
UI Elements for SSL Settings ... 8338
Feature Page Elements ... 8338
Actions Pane Elements ... 8339

Virtual Directories ... 8340
UI Elements for Virtual Directories ... 8340
Feature Page Elements ... 8340
Actions Pane Elements ... 8341
Add or Edit Virtual Directory Dialog Box ... 8341

WebDAV Authoring Rules ... 8343
UI Elements for WebDAV Authoring Rules ... 8343
Feature Page Elements ... 8343
Actions Pane Elements ... 8344
Add Authoring Rule Dialog Box ... 8344
WebDAV Settings Page ... 8345
Feature Page Elements ... 8346
Actions Pane Elements ... 8348

Worker Processes ... 8348
UI Elements for Worker Processes ... 8348
Feature Page Elements ... 8348
Actions Pane Elements ... 8349
Requests Page ... 8349

IIS Manager UI (IIS 8) ... 8350
Navigation Toolbar ... 8350
Connection Manager and Tree ... 8351
Workspace ... 8351
Actions Pane ... 8351
Procedures ... 8352

Open IIS Manager (IIS 8) ... 8352
See Also ... 8353

Manage Connections in IIS Manager (IIS 8) ... 8353
Connect to a web server by specifying connection details manually ... 8353
Connect to a Site by Using IIS Manager ... 8354

Connect to an Application by Using IIS Manager ... 8355
Save a List of Current Connections in IIS Manager ... 8355
See Also ... 8356

Navigation in IIS Manager (IIS 8) ... 8356
Navigate to a web server in IIS Manager ... 8356
Navigate to a site in IIS Manager ... 8357
Navigate to an Application in IIS Manager ... 8357
See Also ... 8358

Start or Stop the Web Server (IIS 8) ... 8358
To start or stop a web server ... 8358
See Also ... 8359

View the Contents of a Site, Application, or Directory (IIS 8) ... 8359
See Also ... 8360

Security Best Practices for IIS 8 ... 8360
Installation and Configuration ... 8360
Web Application Isolation ... 8361
Authentication ... 8361
Request Filtering ... 8362
Application Pool Identities ... 8362
More Security Practices ... 8363

Appcmd.exe (IIS 8) ... 8363
Start Appcmd.exe ... 8364
Get help about Appcmd.exe ... 8365
Working with objects in Appcmd.exe ... 8365
Working with commands in Appcmd.exe ... 8365
Working with attributes and values in Appcmd.exe ... 8366
Errors and the AppHostConfig attribute ... 8366
Configure server-level settings by using Appcmd.exe ... 8366
Configure settings for a site, application, virtual directory, or URL by using Appcmd.exe ... 8367
Use Appcmd.exe to change configuration in a Web.config file ... 8367
Use Appcmd.exe to change configuration in a parent-level configuration file ... 8368

Configuration Store ... 8369
Working with configuration files in IIS 8.0 ... 8369
Delegating configuration in IIS 8.0 ... 8369
Configuration levels ... 8369
Configuration files ... 8370
Inheritance ... 8370
Copy configuration files to a remote server ... 8372
Requirements ... 8372

To copy and deploy a configuration file ... 8372
Copy configuration to another location on the same computer ... 8372
Copy configuration from one computer to another computer ... 8373
See Also ... 8373

Windows Management Instrumentation (WMI) ... 8374
See Also ... 8374

Windows Deployment Services Overview ... 8374
Benefits of Windows Deployment Services ... 8374
Prerequisites for installing Windows Deployment Services ... 8375
Tools for managing Windows Deployment Services ... 8376
Deploying and configuring WDS in a multi-computer environment ... 8376
Installing the WDS role on virtual machines ... 8377
Support for the WDS role in a clustered environment ... 8377
Considerations for managing this role remotely ... 8377
Managing the WDS role using the Server Core installation option ... 8377
Configuring the WDS role for high availability ... 8377
Known issues ... 8377
Common usage scenarios ... 8378
Scenario one: The small business ... 8378
Scenario two: The medium-sized business ... 8379
Scenario three: The large enterprise ... 8380
Scenario Four: A Custom Deployment Using Transport Server ... 8381
Additional references ... 8382

What's New in Windows Deployment Services in Windows Server 2012 R2 ... 8382
Role description ... 8383
New and changed functionality ... 8383
Windows PowerShell cmdlets for WDS ... 8388
Additional references ... 8388

What's New in Windows Deployment Services in Windows Server 2012 ... 8389
Role description ... 8389
New and changed functionality ... 8389
ARM Architecture and support ... 8394
WDS infrastructure for custom deployments ... 8394
WDS Management Console ... 8395
WDSclient.exe ... 8395
Standalone server mode ... 8396
DDP enhancements ... 8396
Expected Deployment Results Wizard ... 8396
TFTP enhancements ... 8396
Troubleshooting enhancements ... 8397

Boot Image and Install Image Priorities 8397
Additional references 8397

Windows Deployment Services Getting Started Guide for Windows Server 2012 8398
In this guide 8398
Quick start checklist 8398
What is Windows Deployment Services? 8399
What's new in Windows Deployment Services? 8399
Who should use this guide? 8399
Benefits of Windows Deployment Services 8399
Installing Windows Deployment Services 8400
Prerequisites for installing Windows Deployment Services 8400
Installation Methods 8401
Install Windows Deployment Services 8403
Configuring Windows Deployment Services 8404
Prerequisites for configuring Windows Deployment Services 8404
Steps for configuring Windows Deployment Services for Standalone Server 8404
Installing Windows Deployment Services integrated with Active Directory 8405
Prerequisites for installing Windows Deployment Services 8405
Steps for configuring Windows Deployment Services integrated with Active Directory 8406
Steps for adding images 8407
Installing an install image 8408
Prerequisites for installing an install image 8408
Steps for installing an install image 8408
Boot Image and Install Image Priorities 8408
Steps for configuring the boot menu 8409
Creating custom install images 8409
Prerequisites for creating custom install images 8409
Known issues when creating custom install images 8410
Steps for creating a capture image 8410
Steps for creating an install image 8411
Creating discover images 8412
Prerequisites for creating discover images 8412
Steps for creating discover images 8412
Performing an unattended installation 8413
Prerequisites for performing an unattended installation 8414
Steps for performing an unattended installation 8414
Creating a multicast transmission 8415
Prerequisites for creating a multicast transmission 8415
Steps for creating a multicast transmission 8415
Steps for configuring transmissions 8415
Steps for configuring clients in a transmission 8416
Additional References 8416

Understanding the Windows Assessment and Deployment Kit (ADK) for Windows 8.1 8416

Windows System Resource Manager Overview 8418
Role/Feature description 8418
Practical applications 8419
Methods of resource management 8419
Built-in resource management policies 8419
Custom resource management 8421
Removed or deprecated functionality 8422

Windows Server Backup Feature Overview 8422
Feature description 8422
Practical applications 8422
New and changed functionality 8423
Virtual Machine Support 8425
More information 8425

Windows Server Backup and Storage Pools 8425
What is a storage pool? 8426
Backup and recovery process and storage pools 8427
Windows Server Backup Considerations 8427
Bare-metal recovery considerations 8427
Example 1: Same Server Recovery Process with disk recreation 8428
Example 2: Same server recovery without disk recreation 8430
Example 3: Alternate server recovery 8431
Critical volumes and storage spaces 8433
Special Considerations when using RAID and Storage Spaces 8435
Sample Scripts 8435

Windows Server Essentials Experience Overview 8436
See also 8437

Windows Server Update Services Overview 8437
WSUS server role description 8437
Practical applications 8438
New and changed functionality 8438
Using Windows PowerShell to manage WSUS 8439
Removed or deprecated functionality 8439
Hardware requirements 8439
See also 8440

Migrate Windows Server Update Services to Windows Server 2012 8440

Step 1: Plan for WSUS Migration 8440
1.1. Know supported operating systems 8441

1.2. Review supported migration scenarios 8441
1.3. Review migration scenarios that are not supported 8442
See also 8442

Step 2: Prepare to Migrate WSUS 8442
2.1. Prepare before you start the migration 8443
2.2. Prepare the destination server 8443
2.3. Prepare the source server 8444
See also 8444

Step 3: Migrate WSUS 8445
3.1. Migrate WSUS update binaries 8445
3.2. Migrate WSUS security groups 8446
3.3. Back up the WSUS database 8447
3.4. Change the WSUS server identity 8451
3.5. Apply security settings 8451
Point the downstream servers to the new WSUS server 8452
Point the WSUS clients to the new WSUS server 8452
3.6. Review additional considerations 8453
See also 8453

Step 4: Verify the WSUS Migration 8454
4.1. Verify the destination server configuration 8454
4.2. Verify client computer functionality 8454
See also 8454

Deploy Windows Server Update Services in Your Organization 8455

Step 1: Prepare for Your WSUS Deployment 8455
1.1. Review considerations and system requirements 8456
WSUS database requirements 8457
1.2. Choose a WSUS deployment scenario 8458
Simple WSUS deployment 8458
Multiple WSUS servers 8459
Disconnected WSUS server 8460
WSUS server hierarchies 8461
Autonomous mode 8461
Replica mode 8462
Branch offices 8463
Network Load Balancing 8464
WSUS deployment with roaming client computers 8464
1.3. Choose a WSUS storage strategy 8465
WSUS database 8465
WSUS with Windows Internal Database 8466
WSUS with SQL Server 8466

WSUS update storage ... 8467
Local WSUS server storage ... 8467
Remote storage on Microsoft Update servers ... 8467
1.4. Choose WSUS update languages ... 8468
1.5. Plan WSUS computer groups ... 8470
Conflict Resolution ... 8471
Priority ... 8471
Priority of Install and Uninstall ... 8472
Priority of Deadlines ... 8472
1.6. Plan WSUS performance considerations ... 8472
Network setup ... 8472
Deferred download ... 8473
Filters ... 8473
Installation ... 8473
Large update deployment ... 8475
Background Intelligent Transfer Service ... 8475
1.7. Plan Automatic Updates settings ... 8475
See also ... 8476

Step 2: Install the WSUS Server Role ... 8476
See also ... 8477

Step 3: Configure WSUS ... 8478
3.1. Configure network connections ... 8478
3.1.1. Connection from the WSUS server to the Internet ... 8479
3.1.2. Connection between WSUS servers ... 8479
3.1.3. Connection between clients (Windows Update Agent) and WSUS servers ... 8480
Configure the proxy server ... 8480
3.2. Configure WSUS by using the WSUS Configuration Wizard ... 8481
3.3. Configure computer groups ... 8483
3.4. Configure client updates ... 8484
Configure Automatic Updates in Group Policy ... 8485
3.5. Secure WSUS with the Secure Sockets Layer Protocol ... 8487
Limitations of WSUS SSL deployments ... 8487
Configure SSL on the WSUS server ... 8487
Configure SSL on client computers ... 8488
Configure SSL for downstream WSUS servers ... 8489
Additional SSL resources ... 8489
3.6. Complete IIS Configuration ... 8490
3.6. Configure a Signing Certificate ... 8490
See also ... 8490

Step 4: Approve and Deploy WSUS Updates ... 8490
4.1. Approve and deploy WSUS updates ... 8491

4.2. Configure auto-approval rules ... 8491
4.3. Review installed updates with WSUS Reports ... 8492
See also ... 8492

Management and Tools for Windows Server 2012 R2 and Windows Server 2012 ... 8493

Windows PowerShell Support for Windows Server ... 8494

Manage Multiple, Remote Servers with Server Manager ... 8494
Review initial considerations and system requirements ... 8495
Hardware requirements ... 8495
Software and configuration requirements ... 8495
Manage remote computers from a client computer ... 8496
Configure remote management on servers that you want to manage ... 8497
Tasks that you can perform in Server Manager ... 8499
Start Server Manager ... 8502
Restart remote servers ... 8502
Export Server Manager settings to other computers ... 8503

Manage the Local Server and the Server Manager Console ... 8504
Shut down the local server ... 8505
Configure Server Manager properties ... 8505
Manage the Server Manager console ... 8508
Add servers to Server Manager ... 8508
Refresh data that is displayed in Server Manager ... 8508
Refresh limitations ... 8509
Add or remove roles or features ... 8509
Create server groups ... 8509
Prevent Server Manager from opening automatically at logon ... 8509
Zoom in or out ... 8510
Customize tools that are displayed in the Tools menu ... 8510
Manage roles on role home pages ... 8511
See Also ... 8512

Configure Remote Management in Server Manager ... 8513
Enabling or disabling remote management ... 8514
Windows Remote Management (WinRM) listener settings ... 8518
See Also ... 8518

Install or Uninstall Roles, Role Services, or Features ... 8518
Install roles, role services, and features by using the Add Roles and Features Wizard ... 8519
Install roles, role services, and features by using Windows PowerShell cmdlets ... 8521
Remove roles, role services, and features by using the Remove Roles and Features Wizard ... 8523
Remove roles, role services, and features by using Windows PowerShell cmdlets ... 8525

Install roles and features on multiple servers by running a Windows PowerShell script 8527
Install .NET Framework 3.5 and other features on-demand 8529
Configure alternate sources for feature files in Group Policy 8531
See Also 8532

Configure Features on Demand in Windows Server 8533
Create a feature file or side-by-side store 8533
Methods of removing feature files 8534
Remove feature files by using Uninstall-WindowsFeature 8535
See Also 8536

Add Servers to Server Manager 8536
Provide credentials with the Manage As command 8537
Add servers to manage 8538
Add and manage servers in workgroups 8539
Add and manage servers in clusters 8541
See Also 8541

View and Configure Performance, Event, and Service Data 8541
What are thumbnails? 8542
View and configure events 8544
View and configure performance log data 8546
Analyze performance data and solve problems 8548
Manage services and configure service alerts 8548
View and copy event, service, or performance entries 8549
See Also 8549

View Task Details and Notifications 8549
The Notifications area 8550
Viewing and troubleshooting tasks by using Task Details 8550
See Also 8551

Run Best Practices Analyzer Scans and Manage Scan Results 8551
Find BPA 8552
How BPA works 8552
Rule categories 8552
Performing Best Practices Analyzer scans on roles 8554
Scanning roles by using the BPA GUI 8554
Scanning roles by using Windows PowerShell cmdlets 8555
Manage scan results 8556
Exclude and include BPA results 8557
Exclude scan results 8557
Include scan results 8558
View and export BPA scan results in Windows PowerShell 8559
See Also 8560

Create and Manage Server Groups 8560
Server groups 8561
See Also 8562

Filter, Sort, and Query Data in Server Manager Tiles 8562
Filter list entries in tiles 8562
Sort list entries in tiles 8563
Create and run custom queries on tile data 8563
See Also 8564

Keyboard Shortcuts for Server Manager 8564
Access keys 8564

Deploy Remote Server Administration Tools 8566
Remote Server Administration Tools for Windows 8.1 8567
Tools available in this release 8567
System requirements 8567
Install, uninstall, or disable Remote Server Administration Tools for Windows 8.1 8568
Run Remote Server Administration Tools 8570
Remote Server Administration Tools for Windows 8 8571
Tools available in this release 8571
System requirements 8571
Install, uninstall, or disable Remote Server Administration Tools for Windows 8 8572
Run Remote Server Administration Tools 8574
See Also 8575

Install and Use Windows PowerShell Web Access 8575
Requirements for running Windows PowerShell Web Access 8577
Browser and client device support 8577
Supported desktop computer browsers 8577
Minimally-tested mobile devices or browsers 8577
Browser requirements 8578
Recommended (quick) deployment 8578
Step 1: Install Windows PowerShell Web Access 8578
Step 2: Configure the gateway 8579
Step 3: Configure a restrictive authorization rule 8582
Custom deployment 8583
Step 1: Install Windows PowerShell Web Access 8583
Step 2: Configure the gateway 8584
Step 3: Configure a restrictive authorization rule 8587
Configure a genuine certificate 8588
Using the web-based Windows PowerShell console 8589
See Also 8589

Authorization Rules and Security Features of Windows PowerShell Web Access 8589

Configuring authorization rules and site security ... 8589
Security ... 8591
Configuring authorization rules ... 8595
Other authorization rule scenario examples ... 8597
Using a single set of authorization rules for multiple sites ... 8599
Session management ... 8599
Setting default parameters on the sign-in page ... 8599
Time-outs and unplanned disconnections ... 8600
See Also ... 8600

Use the Web-based Windows PowerShell Console ... 8600
Supported browsers and client devices ... 8601
Supported desktop computer browsers ... 8601
Minimally-tested mobile devices or browsers ... 8601
Browser requirements ... 8601
Signing in to Windows PowerShell Web Access ... 8602
Signing out and timing out ... 8603
Differences in the web-based Windows PowerShell console ... 8604
Functional disparity with PowerShell.exe ... 8604
Limitations of the web-based console ... 8606
See Also ... 8607

Troubleshooting Access Problems in Windows PowerShell Web Access ... 8608
Troubleshooting access problems ... 8608
See Also ... 8610

Uninstall Windows PowerShell Web Access ... 8611
Uninstalling Windows PowerShell Web Access ... 8611
Recommended (quick) uninstallation ... 8611
Step 1: Delete the web application ... 8611
Step 2: Uninstall Windows PowerShell Web Access ... 8612
Custom uninstallation ... 8612
Step 1: Delete the web application ... 8612
Step 2: Uninstall Windows PowerShell Web Access ... 8613
See Also ... 8613

Getting Started with Windows PowerShell Workflow ... 8614
Overview of Windows PowerShell Workflow ... 8614
Activities ... 8615
Benefits of Windows PowerShell Workflow ... 8615
How Windows PowerShell Workflow and Windows PowerShell scripts differ ... 8616
Creating and importing workflows by using the Visual Studio Workflow Designer ... 8616
Windows PowerShell Workflow activities in the Visual Studio Workflow Designer ... 8616
See Also ... 8617

Writing a Script Workflow 8617
Planning a Workflow 8617
The Workflow Keyword 8618
Naming Workflows and Workflow Elements 8618
How to Add Parameters to a Workflow 8619
How to Add Activities to a Workflow 8620
How to Use the Parameters of Activities 8622
How to Use Activity Common Parameters 8622
How to Get Common Parameter Values 8623
How to Get Runtime Variables and Common Parameter Values 8624
How to Run a Script in a Workflow 8625
How to Run Commands in Parallel 8626
Parallel 8626
ForEach -Parallel 8627
Sequence 8628

Adding Nested Functions and Nested Workflows 8629
Adding Nested Functions and Nested Workflows 8629
Using Parameters and Variables in Nested Functions and Workflows 8629
Calling Nested Workflows and Functions 8631

Running Windows PowerShell Commands in a Workflow 8634
The InlineScript activity 8634
Rules for Using the InlineScript Activity 8634
Variables in InlineScript 8635
Running InlineScript in the Workflow Process 8636
Examples 8637

Making a Workflow Suspend Itself 8638
The Suspend-Workflow Activity 8638
Suspending a Workflow 8638
Resuming a Workflow Job 8639
Getting the Output of a Workflow Job 8640

Restarting the Computer in a Workflow 8640
Restart-Computer 8641
Restarting the Workflow Computer 8641
Manually Resuming the Workflow 8642
Automate Resuming the Workflow 8643
See Also 8644

Adding Custom Activities to a Script Workflow 8644

Adding Checkpoints to a Script Workflow 8645
What is a checkpoint? 8645

Where to place checkpoints ... 8646
How to add a checkpoint ... 8646
PSPersist workflow common parameter ... 8647
PSPersist activity common parameter ... 8647
Checkpoint-Workflow Activity ... 8648
$PSPersistPreference Preference Variable ... 8648
Checkpoints in Pipeline and Parallel Script Blocks ... 8649

Writing Help for a Script Workflow ... 8649
Supporting the Get-Help Cmdlet ... 8649
Supporting Online Help ... 8650
Supporting Updatable Help ... 8651
See Also ... 8651

Saving Your Workflow in a Module ... 8652
Prepare your workflow by saving it in a module ... 8652
See Also ... 8652

Workflow Authoring Reference Topics ... 8652

Syntactic Differences Between Script Workflows and Scripts ... 8653
Scripts and Script Workflows ... 8653
The Workflow Keyword ... 8653
Workflow Names ... 8654
CmdletBinding Attribute ... 8654
Workflow Parameters ... 8654
Workflow Common Parameters ... 8655
Activity Parameters ... 8655
Activity Common Parameters ... 8655
Commands Run in Independent Sessions ... 8656
No Method Calls ... 8656
Statements in a Workflow ... 8657
Control Statements in a Workflow ... 8657
Reserved Words in Workflows ... 8658
Using Provider Drives in Workflows ... 8658

Using Activities in Script Workflows ... 8659
Script Workflow Activities ... 8659
Cmdlet Activities ... 8660
Activities that Differ from Cmdlets ... 8660
Excluded Cmdlets ... 8661
Activity Common Parameters ... 8662

Using Variables in Script Workflows ... 8668
Automatic Variables in Workflows ... 8668

Windows PowerShell Preference Variables ... 8668
Workflow Preference Variables ... 8668
$PSParentActivityId <string> ... 8668
$PSPersistPreference <Boolean> ... 8669
$PSRunInProcessPreference ... 8670
Workflow Runtime Variables and Common Parameter Variables ... 8670
Rules for Variables in Workflows ... 8673
Rules for Variables in Loops, Inline Scripts, and Parallel statements ... 8674
Variables in Loops and Control Statements ... 8675
Variables in Parallel and Sequence Statements ... 8675
Variables in an InlineScript Activity ... 8677

Configuring Your Workflow Environment ... 8679
Supported operating systems and architectures ... 8679
Workflow configurations ... 8680
Download locations for .NET Framework and Windows PowerShell ... 8682
Prepare computers to run workflows ... 8682
Workflow session configurations ... 8683
See Also ... 8684

Running a Windows PowerShell Workflow ... 8684
Create a session to the computer that runs your workflow ... 8685
Connect or reconnect to a running workflow ... 8686
Run the workflow ... 8686
Run workflows as jobs ... 8687
View data about workflows ... 8689
Add checkpoints to a workflow ... 8689
Get help for your workflow ... 8690
Find error logs for workflows ... 8690
See Also ... 8690

Command-Line Reference for Windows Server 2012 ... 8691

Windows Remote Management (WinRM) Overview ... 8691
WinRM description ... 8691
Practical applications ... 8691
See also ... 8692

Windows Management Instrumentation (WMI) Overview ... 8693
WMI description ... 8693
Practical applications ... 8694
See also ... 8694

Datacenter Abstraction Layer (DAL) Overview ... 8695
DAL description ... 8695

Practical applications 8697
Important functionality 8698
See also 8699

Physical Computer System View (PCSV) Property Mappings 8700
SMASH/DASH property mappings 8700
IPMI property mappings 8704
Unimplemented properties 8705
See also 8707

Background Intelligent Transfer Service (BITS) Overview 8707
BITS description 8707
Practical applications 8708
Server Manager information 8708
See also 8708

Windows Server 2012 R2 and Windows Server 2012

Welcome to the Windows Server 2012 R2 and Windows Server 2012 Technical Library.

This library provides the core content that IT pros need to evaluate, plan, deploy, manage, troubleshoot, and support servers running the Windows Server 2012 R2 and Windows Server 2012 operating systems. To find the information you need, browse the content listed below, or use our search feature. If you need help with using Windows at home, see Windows Help and How-To.

- What's New in Windows Server 2012 R2
  The content in this section describes what's new and changed in Windows Server 2012 R2. This content focuses on changes that will potentially have the greatest impact on your use of this release.
- What's New in Windows Server 2012
  This section describes what's new and changed in Windows Server 2012. This content focuses on changes that will potentially have the greatest impact on your use of this release.
- Technical Scenarios for Windows Server 2012 R2 and Windows Server 2012
  These technical scenarios provide guidance that addresses an area of interest or particularly compelling capability enabled by Windows Server 2012 R2 and Windows Server 2012, such as virtualization and Dynamic Access Control.
- Install and Deploy Windows Server 2012 R2 and Windows Server 2012
  This section provides information about how to migrate to, install, and deploy Windows Server 2012 R2 and Windows Server 2012.
- Migrate Roles and Features to Windows Server
  Migration documentation and tools ease the process of migrating server roles, features, operating system settings, and data from an existing server that is running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 to a computer that is running Windows Server 2012.
- Secure Windows Server 2012 R2 and Windows Server 2012
  This section provides the security products and assessment tools needed to help secure servers running Windows Server 2012 R2 or Windows Server 2012.
- Manage Privacy
  This section offers not only guidance about managing privacy-related settings in Windows Server 2012, but also links that can be useful to administrators and others concerned about privacy.
- Support Windows Server 2012 R2 and Windows Server 2012
  This section contains information to help IT pros find workarounds for known issues in

Windows Server 2012 R2 and Windows Server 2012, and troubleshoot and resolve specific Windows Server 2012 R2 and Windows Server 2012 system errors and events.

- Server Roles and Technologies in Windows Server 2012 R2 and Windows Server 2012
  This section contains information to design, deploy, manage, and troubleshoot technologies in Windows Server 2012 R2 and Windows Server 2012.
- Management and Tools for Windows Server 2012 R2 and Windows Server 2012
  This section contains tool and reference information for IT pros using Windows Server 2012 R2 and Windows Server 2012. Learn about Windows PowerShell, Server Manager, Remote Server Administration Tools, and Command-Line Reference.

What's New in Windows Server

The content in this section describes what's new and changed in the latest releases of Windows Server.

- What's New in Windows Server 2012 R2
- What's New in Windows Server 2012

What's New in Windows Server 2012 R2

The content in this section describes what's new and changed in Windows Server 2012 R2. This content focuses on changes that will potentially have the greatest impact on your use of this release.

- What's New in 802.1X Authenticated Wired Access in Windows Server 2012 R2
  This topic provides information about the new features for 802.1X Authenticated Wired Access in Windows Server 2012 R2 and Windows 8.1.
- What's New in 802.1X Authenticated Wireless Access in Windows Server 2012 R2
  This topic provides information about the new features for 802.1X Authenticated Wireless Access in Windows Server 2012 R2 and Windows 8.1, including Miracast Wireless Display and faster Wi-Fi with 802.11ac.
- What's New in Active Directory in Windows Server 2012 R2
  You can leverage new features in Active Directory to enable employees and partners to access protected corporate data from their personal devices and at the same time manage risk and govern the use of corporate resources.
- **What's New in BitLocker in Windows 8.1 and Windows Server 2012 R2**
  BitLocker now provides support for device encryption on x86-based and x64-based computers with a Trusted Platform Module that supports connected standby. This topic describes the new functionality.

- What's New in Certificate Services in Windows Server 2012 R2
  Active Directory Certificate Services in Windows Server 2012 R2 supports a policy module for the Network Device Enrollment Service, TPM key attestation, and new Windows PowerShell cmdlets for backup and restore.
- What's New in Data Deduplication in Windows Server 2012 R2
  Data Deduplication can now be installed on a scale-out file share and used to optimize live virtual hard disks (VHDs) for Virtual Desktop Infrastructure (VDI) workloads. This topic describes this and other new functionality.
- What's New in DFS Replication in Windows Server 2012 R2
  This topic describes the features that were added to DFS Replication (DFSR or DFS-R) in Windows Server 2012 R2.
- What's New in DHCP in Windows Server 2012 R2
  Dynamic Host Configuration Protocol (DHCP) in Windows Server 2012 R2 provides new features and capabilities over previous versions. This document describes new deployment, manageability, and capabilities added to the DHCP Server role in Windows Server 2012 R2.
- What's New in DNS Server in Windows Server 2012 R2
  This topic provides information about new and changed functionality in the DNS Server service in Windows Server 2012 R2.
- What's New in Failover Clustering in Windows Server 2012 R2
  This topic describes the Failover Clustering functionality that is new or changed in Windows Server 2012 R2.
- New and Changed Functionality in File and Storage Services
  File and Storage Services provides a number of new management, scalability, and functionality improvements in Windows Server 2012 R2.
- What's New in File Server Resource Manager in Windows Server 2012 R2
  This topic summarizes the File Server Resource Manager functionality in Windows Server 2012 R2 that is new or changed since Windows Server 2012.
- What's New in Group Policy in Windows Server 2012 R2
  This topic describes the new and changed functionality of the Group Policy feature in Windows Server 2012 R2.
- What's New in Hyper-V for Windows Server 2012 R2
  This topic describes the new and changed functionality of the Hyper-V role in Windows Server 2012 R2.
- What's New in Hyper-V Network Virtualization in Windows Server 2012 R2
  This topic describes the new or changed features and functionality in Hyper-V Network Virtualization in Windows Server 2012 R2.
- What's New in Hyper-V Virtual Switch in Windows Server 2012 R2
  This topic provides information about the new features in Hyper-V Virtual Switch in Windows Server 2012 R2.
- What's New in IPAM in Windows Server 2012 R2
  IP Address Management (IPAM) is a feature that was first introduced in Windows Server 2012 that provides highly customizable administrative and monitoring capabilities for the IP address infrastructure on a corporate network. IPAM in Windows Server 2012 R2 includes many enhancements.

- What's New in iSCSI Target Server in Windows Server 2012 R2
  This topic describes the new and changed functionality of iSCSI Target Server in Windows Server 2012 R2.
- What's New in Networking in Windows Server 2012 R2
  This topic describes the new and changed functionality of networking in Windows Server 2012 R2.
- What's New in Print and Document Services in Windows Server 2012 R2
  This topic describes the new and changed functionality of Print and Document Services in Windows Server 2012 R2.
- What's New in Remote Access in Windows Server 2012 R2
  A number of new Remote Access server and client features are included in Windows Server 2012 R2 and Windows 8.1.
- What's New in Remote Desktop Services in Windows Server 2012 R2
  This topic describes the Remote Desktop Services functionality that is new or changed in Windows Server 2012 R2.
- New and Changed Functionality in Security and Protection in Windows Server 2012 R2
  This topic describes the significant changes to security technologies in Windows Server 2012 R2 and Windows Server 2012 and how those changes impact Windows 8.1.
- What's New in SMB in Windows Server 2012 R2
  This topic introduces the new features and functionality for Server Message Block (SMB) in Windows Server 2012 R2.
- What's New in Storage Spaces in Windows Server 2012 R2
  This topic describes the features that were added to Storage Spaces in Windows Server 2012 R2, including storage tiers, write-back cache, and dual parity.
- What's New in Windows Deployment Services in Windows Server 2012 R2
  A Windows Deployment Services (WDS) server running Windows Server 2012 R2 can be managed using the Windows PowerShell cmdlets for WDS. Using Windows PowerShell cmdlets, you can add driver packages, add client images, enable and disable boot and install images, and do many other common WDS tasks. For a full reference, see Windows PowerShell Support for Windows Server.
- What's New in Windows PowerShell
  Windows PowerShell includes several significant features that extend its use, improve its usability, and allow you to control and manage Windows-based environments more easily and comprehensively.
- **What's New in Windows Server 2012 R2 Essentials**
  The content in this section describes what's new and changed in Windows Server 2012 R2 Essentials. This content focuses on changes that will potentially have the greatest impact on your use of this release.

What's New in Windows Server 2012

The content in this section describes what's new and changed in Windows Server 2012. This content focuses on changes that will potentially have the greatest impact on your use of this release.

- What's New in Certificate Services in Windows Server 2012
  Active Directory Certificate Services (AD CS) in Windows Server 2012 provides multiple new features and capabilities over previous versions. This document describes new deployment, manageability, and capabilities added to AD CS in Windows Server 2012.
- What's New in Active Directory Domain Services (AD DS)
  Active Directory Domain Services (AD DS) in Windows Server 2012 includes new features that make it simpler and faster to deploy domain controllers (both on-premises and in the cloud), more flexible and easier to both audit and authorize access to files with Dynamic Access Control, and easier to perform administrative tasks at scale, either locally or remotely, through consistent graphical and scripted management experiences.
- What's New in Active Directory Rights Management Services (AD RMS)?
  Active Directory Rights Management Services (AD RMS) is the server role that provides you with management and development tools that work with industry security technologies—including encryption, certificates, and authentication—to help organizations create reliable information protection solutions.
- What's New in BitLocker for Windows 8 and Windows Server 2012
  BitLocker encrypts the hard drives on your computer to provide enhanced protection against data theft or exposure on computers and removable drives that are lost or stolen.
- What's New in BranchCache
  BranchCache in Windows Server 2012 and Windows 8 provides substantial performance, manageability, scalability, and availability improvements.
- What's New in DFS Namespaces and DFS Replication in Windows Server 2012
  DFS Namespaces and DFS Replication in Windows Server 2012 provide new management functionality as well as interoperability with DirectAccess and Data Deduplication.
- What's New in DHCP in Windows Server 2012
  Dynamic Host Configuration Protocol (DHCP) is an Internet Engineering Task Force (IETF) standard designed to reduce the administration burden and complexity of configuring hosts on a TCP/IP-based network, such as a private intranet.
- What's New in DNS
  Domain Name System (DNS) services in Windows Server 2012 and Windows 8 are used in TCP/IP networks for naming computers and network services. DNS naming locates computers and services through user-friendly names.
- New and changed functionality in File and Storage Services
  File and Storage Services provides a number of new management, scalability, and functionality improvements in Windows Server 2012.

- What's New in Failover Clustering in Windows Server 2012
  Failover clusters provide high availability and scalability to many server workloads. These include file share storage for server applications such as Hyper-V and Microsoft SQL Server, and server applications that run on physical servers or virtual machines.
- What's New in File Server Resource Manager in Windows Server 2012
  File Server Resource Manager provides a set of features that allow you to manage and classify data that is stored on file servers.
- What's New in Group Policy in Windows Server 2012
  Group Policy is an infrastructure that enables you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences
- What's New in Hyper-V for Windows Server 2012
  The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built in to Windows Server 2012. Hyper-V virtualizes hardware to provide an environment in which you can run multiple operating systems at the same time on one physical computer, by running each operating system in its own virtual machine.
- What's New in IPAM in Windows Server 2012
  IP Address Management (IPAM) is an entirely new feature in Windows Server 2012 that provides highly customizable administrative and monitoring capabilities for the IP address infrastructure on a corporate network.
- What's New in Kerberos Authentication
  The Microsoft Windows Server operating systems implement the Kerberos version 5 authentication protocol and extensions for public key and password-based authentication. The Kerberos authentication client is implemented as a security support provider (SSP) and can be accessed through the Security Support Provider Interface (SSPI).
- What's New for Managed Service Accounts
  Standalone Managed Service Accounts, which were introduced in Windows Server 2008 R2 and Windows 7, are managed domain accounts that provide automatic password management and simplified SPN management, including delegation of management to other administrators.
- What's New in Networking in Windows Server 2012
  Discover new networking technologies and new features for existing technologies in Windows Server 2012. Technologies covered include BranchCache, Data Center Bridging, NIC Teaming, and more.
- What's New in Remote Desktop Services in Windows Server 2012
  The Remote Desktop Services server role in Windows Server 2012 provides technologies that enable users to connect to virtual desktops, RemoteApp programs, and session-based desktops. With Remote Desktop Services, users can access remote connections from within a corporate network or from the Internet.
- **What's new in Security Auditing**
  Security auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is to verify regulatory compliance.
- What's new in Server Manager
  In this blog post, senior Server Manager program manager Wale Martins describes the

innovations and value of the new Server Manager. Server Manager in Windows Server 2012 lets administrators manage multiple, remote servers that are running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.

- What's New in Smart Cards
  Smart cards and their associated personal identification numbers (PINs) are an increasingly popular, reliable, and cost-effective form of two-factor authentication. With the right controls in place, a user must have the smart card and know the PIN to gain access to network resources.
- What's New in TLS/SSL (Schannel SSP)
  Schannel is a Security Support Provider (SSP) that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols. The Security Support Provider Interface (SSPI) is an API used by Windows systems to perform security-related functions including authentication.
- What's New in Windows Deployment Services in Windows Server 2012
  Windows Deployment Services is a server role that enables you to remotely deploy Windows operating systems. You can use it to set up new computers by using a network-based installation.
- What's new in Windows PowerShell 3.0
  Windows PowerShell 3.0 includes many new features and improvements in the scripting and automation experience, such as Windows PowerShell Workflow, multiple new features in Windows PowerShell ISE to help make scripting and debugging faster and easier, updatable Help, Windows PowerShell Web Access, and over 2,200 new cmdlets and functions.

Technical Scenarios for Windows Server 2012 R2 and Windows Server 2012

The Technical Scenarios provide guidance that addresses an area of interest or particularly compelling capability enabled by Windows Server 2012 R2 or Windows Server 2012. Typically, multiple features or product capabilities must come together to solve a problem or address a need. The guidance in this area will grow over time as we address more scenarios and get feedback from our customers.

Windows Server 2012 R2

- Access and Information Protection
  You can leverage new features in Active Directory to enable employees and partners to access protected corporate data from their personal devices and at the same time manage risk and govern the use of corporate resources.
- **Hybrid Cloud Multi-Tenant Networking Solution Guide**
  You can connect multiple tenant networks to your network and offer secure network isolation to each tenant in a way that is simple to set up and efficient to operate.

Windows Server 2012

- Building Your Cloud Infrastructure: Scenario Overview
  You can leverage new features around network and storage virtualization that, when combined with improved server virtualization, enable the building of your cloud infrastructure based on Windows Server 2012. This will help with your strategy in delivering Infrastructure as a Service (IaaS) or building hosted services.
- Dynamic Access Control: Scenario Overview
  You can apply data governance across your file servers to control who can access information and to audit who has accessed information.
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
  Rapid and efficient scaling of your web applications makes for a cloud-ready web platform. Enhanced security, application initialization, NUMA-aware scalability, and the sharing of resources across sites allows for this rapid scaling with minimal management overhead.
- Increasing Server, Storage, and Network Availability: Scenario Overview
  New experiences in Windows Server 2012 work together to improve availability, performance, and reliability at the single-server and multiple-server (scale-up and scale-out) levels.

Access and Information Protection

One of the most prevalent IT industry trends at the moment is the proliferation of consumer devices in the workplace. Employees and partners want to access protected corporate data from their personal devices, from checking email to the consumption of advanced business applications. IT administrators in organizations, while wanting to enable this level of productivity, would like to ensure that they can manage risk and govern the use of corporate resources.

Following are some of the cross-technology Access and Information Protection solutions you can implement in your organization:

Solution	Area	Description
Secure Access to Company Resources from any Location, on any Device	Device and Access Management	Addresses how you can implement an end-to-end solution for managing corporate-owned and personal devices, and at the same time providing secure and seamless access to corporate applications and resources.
Hybrid Identity Solution Guide – DirSync with Password Sync	Hybrid Identity	Describes a hybrid identity solution that uses DirSync with Password Sync to synchronize

Solution	Area	Description
		on-premises Active Directory with Windows Azure Active Directory.
Hybrid Identity Solution Guide – DirSync with Federation (SSO)	Hybrid Identity	Describes a hybrid identity solution that uses DirSync with AD FS and Federation (SSO) to synchronize on-premises Active Directory with Windows Azure Active Directory and provide a seamless single sign-on experience.
Dynamic Access Control: Scenario Overview	Information Protection	Provides a solution to apply data governance across corporate resources and set up policies to exercise control to who can access the data. You can also set up audit access to files for compliance reporting, and apply rights-managed protection for sensitive documents.

In addition to security and access, IT also needs to have a good strategy in place to manage PCs and personal devices from a single administrator console. Managing devices includes setting security and compliance settings, gathering software and hardware inventory, or deploying software. IT also must have a solution in place to protect the company by wiping corporate data stored on the mobile device when the device is lost, stolen, or retired from use.The solution, Manage mobile devices and PCs by migrating to Configuration Manager with Windows Intune explains in detail the Unified Device Management solution.

See Also

Technical Scenarios for Windows Server 2012 R2 and Windows Server 2012

Manage mobile devices and PCs by migrating to Configuration Manager with Windows Intune

Building Your Cloud Infrastructure: Scenario Overview

An overview of the process for building your cloud infrastructure in Windows Server® 2012 for building an infrastructure for public and private clouds, including a scenario description, its practical applications, the roles and services that enable it, and links to topics that describe how to deploy it.

Did you mean…

1. Private cloud
2. Public cloud Software as a Service
3. Public Cloud Platform as a Service

Scenario description

The process for building your cloud infrastructure uses a combination of Hyper-V, failover clustering, storage, and networking technologies to more easily create a Microsoft cloud infrastructure. Windows Server 2012 introduces a significant number of new features that provide all of the required capabilities for building an effective cloud infrastructure in an open platform. By using automation, having an open platform, and being standards based, a Windows Server 2012-based cloud infrastructure decreases the total cost of ownership and reduces susceptibility to failures due to interoperability issues. The Windows Server 2012 open platform allows partners to extend the functionality beyond what is in the platform.

In this scenario

The following topics for the process for building your cloud infrastructure allow you to design a Hyper-V cluster for hosting the virtual machines of an IaaS cloud, configure the cloud, and then perform operations such as the on boarding of existing virtual machines and virtual machine maintenance:

- Building Your Cloud Infrastructure: Non-Converged Data Center Configuration
- Building Your Cloud Infrastructure: Converged Data Center with File Server Storage
- Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes

Practical applications

Building a cloud infrastructure with Windows Server 2012 is much easier than previous versions of Windows Server because of built-in support for the following:

- Multi-tenancy

  Designing a data center for dynamic and automatic placement of virtual machines is not enough, especially when you are serving more than one customer. Multi-tenancy is the ability of a cloud infrastructure to support the virtual machine workloads of multiple tenants, but

isolate them from each other, but all of the workloads run on the same infrastructure. The multiple workloads of an individual tenant can interconnect and be managed remotely, but these systems do not interconnect with the workloads of other tenants, nor can other tenants remotely manage them.

- Highly Scalable, Low-Cost Data Center

  You can deploy different scales of clouds, such as:

  - A mid-market business with a need for a small number of servers
  - An enterprise with hundreds or thousands of servers
  - An IaaS hosting provider with thousands of servers for multiple customers

  In all of these cases, Windows Server 2012 supports features that use low-cost alternatives to traditional data center resources that allow small clouds to be easily built, but also supports features that enable high scale operations.

- Managing and Extending the Data Center

  Windows Server 2012 simplifies data center initial configuration and ongoing management with the support of PowerShell 3.0 and PowerShell workflows. With the Hyper-V virtual switch in Windows Server 2012, vendors and partners can extend the capabilities of the virtual switch to develop additional functionality, such as network monitoring, security, or routing add-ons.

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Hyper-V Overview	The Hyper-V server role hosts the virtual machines that make up the IT workloads running in the cloud. The process for building your cloud infrastructure takes advantage of many new features of Hyper-V in Windows Server 2012, including the Hyper-V virtual switch, live migration and storage migration improvements, Hyper-V Replica, and resource metering.
Failover Clustering Overview	The Failover Clustering feature allows a set of computers to act as a single computer, providing scale-out and failover to clustered services and resources, including storage and virtual machines. Windows Server 2012 is tightly integrated with Hyper-V and enables much of the fabric management capabilities

Role/feature	How it supports this scenario
	required for a cloud infrastructure.
File and Storage Services Overview	File and Storage Services allows a pool of storage to be made available to the cluster to assign to virtual machines and to store virtual hard disk files.
Network Adapter Teaming	The network adapter teaming feature allows you to group multiple network adapters into a team that appears as a single network adapter for bandwidth aggregation and fault tolerance.
Hyper-V Virtual Switch Overview	The Hyper-V virtual switch platform allows network partners to easily hook into the Hyper-V virtual switch network flows and build monitoring, security, and forwarding extensions.

Hardware requirements

The exact hardware requirements depend on the types of workloads you are planning to run on the Hyper-V servers in the cloud infrastructure. However, some features included in the cloud infrastructure scenarios require specialized BIOS capabilities, such as those needed to support Single Root I/O Virtualization (SR-IOV). Any technologies specific hardware requirements are discussed within the pertinent scenario..

Software requirements

Clouds based on this process for building your cloud infrastructure require Windows Server 2012 and its platform capabilities.

See also

See the following table for links to additional resources about building your cloud infrastructure.

Content type	References
Product evaluation	Building an Infrastructure as a Service (IaaS) Cloud Using Windows Server 2012
Deployment	Building Your Cloud Infrastructure: Non-Converged Data Center Configuration | Building Your Cloud Infrastructure: Converged Data

Content type	References
	Center with File Server Storage | Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes
Community resources	Private Cloud Architecture Blog
Related technologies	Hyper-V Overview | Failover Clustering Overview | File and Storage Services Overview | Networking Overview

Designing Your Cloud Infrastructure

The Build Your Private Cloud Infrastructure Design Guide provides you with the information that you need to make informed decisions about the storage, network and compute design decision points to build a cloud infrastructure that meets the requirements of your organization. This guide provides information that you can use to assess Windows Server® 2012 platform technologies that enable a cloud-based infrastructure.

This Design Guide is focused on the design of the cloud infrastructure and the components that make up a cloud infrastructure. It does not provide information on how build a complete private cloud, public cloud, or hosted cloud infrastructure as a service (IaaS), platform as a service (PaaS), or software as a service (SaaS) solution. The cloud infrastructure contains the building blocks on which any Windows Server 2012 cloud service or delivery model is built.

This document is comprised of the following sections:

- **Cloud Infrastructure Technical Overview**. This section provides a short overview of cloud computing and the requirements of a cloud infrastructure.
- **Cloud Infrastructure Design**. This section provides an introduction to the cloud infrastructure design process.
- **Designing the Cloud Storage Infrastructure**. This section provides information related to design considerations for building the cloud storage infrastructure using Windows Server 2012 platform features and capabilities.
- **Designing the Cloud Network Infrastructure**. This section provides information related to design considerations for building the cloud network infrastructure by using Windows Server 2012 platform features and capabilities.
- **Designing the Cloud Compute (Virtualization) Infrastructure**. This section provides information related to design considerations for building the cloud compute (virtualization) infrastructure using Windows Server 2012 platform features and capabilities.
- **Overview of Suggested Cloud Infrastructure Deployment Scenarios.** This section provides information on three suggested cloud infrastructure deployment scenarios and the design decisions that drive selecting one over the other.

It should be noted from the beginning that while the terms "cloud" and "private cloud" are used throughout this document, they are used interchangeably. Private cloud, public cloud, and hybrid cloud refer to deployment models for cloud computing and it is assumed that all cloud computing infrastructures share common design considerations. This includes cloud hosting environments, which can host private or public clouds. Therefore, this documentation is topical for corporate IT organizations, private cloud hosting providers, and even public cloud providers who are interested in designing a Windows Server 2012 cloud infrastructure.

The primary focus of this Design Guide is to provide design information for the cloud infrastructure that will support the IaaS service model. Subsequent design guides will provide details on how to design, plan and implement a complete IaaS solution.

Cloud Infrastructure Technical Overview

NIST Definition of Cloud Computing

Cloud computing is a model for enabling ubiquitous, on-demand network access to a shared pool of computing resources (which include networks, servers, storage, applications, and services) that can be dynamically and rapidly provisioned and released with minimal human interaction. The United States National Institute of Standards and Technology (NIST) definition of cloud computing is one of the more popular and well-received models, and therefore is used as a basis for understanding and creating a common vocabulary and definition of cloud computing. The NIST definition includes five essential characteristics, three service models, and four deployment models.

Essential Characteristics:

- **On-demand self-service**. A consumer of the cloud service can unilaterally and automatically provision computing capabilities, such as server time and network storage, as needed without requiring human interaction with each service provider.
- **Broad network access**. The ability of a consumer of a cloud service to obtain access to cloud resources from any location, using a wide array of devices.
- **Resource pooling.** The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- **Rapid elasticity.** Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand.
- **Measured service.** Cloud systems automatically control and optimize resource use by taking advantage of a metering capability at some level of abstraction appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts).

It should be noted that in the case of public clouds, broad network access to the cloud is crucial. Broad network access enables the cloud to be accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (for example, mobile phones, tablets,

laptops, and workstations). Private clouds may or may not require support for broad network access.

Service Models:

- **Software as a service (SaaS).** The capability provided to the consumer is to use the cloud service provider's applications running on a cloud infrastructure.
- **Platform as a service (PaaS).** The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider.
- **Infrastructure as a service (IaaS).** The capability provided to the consumer is to provision processing (compute), storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications.

Deployment Models:

- **Private cloud.** The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (for example, business units).
- **Community cloud.** The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (for example, mission, security requirements, policy, and compliance considerations).
- **Public cloud.** The cloud infrastructure is provisioned for open use by the general public. It might be owned, managed, and operated by a business, academic, or government organization, or some combination of those.
- **Hybrid cloud.** The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (for example, cloud bursting for load balancing between clouds).

It is important to note that there is no specific alignment between cloud service and delivery models. Any service model can be instantiated using any of the delivery models. For example, while Software as a Service (SaaS) is often considered a public cloud service model, there is no reason why SaaS cannot be delivered in an on-premises or hosted private cloud.

For more information and a detailed discussion of the NIST definition of private cloud, please see the NIST Definition of Cloud Computing.

Microsoft Private Cloud Overview

Private cloud is a computing model that uses infrastructure dedicated to your organization. A private cloud shares many of the characteristics of public cloud computing including resource pooling, self-service, elasticity, and metered services delivered in a standardized manner with the additional control and customization available from dedicated resources.

The only fundamental difference between a private cloud and a public cloud is that a public cloud provides cloud resources to multiple organizations while the private cloud hosts resources for a

single organization. However, a single organization may have multiple business units and divisions which can lend itself to being multi-tenant in nature. In these circumstances, private cloud shares many of the security and isolation requirements of public cloud.

While virtualization is an important technological enabler of private cloud, the key differentiator is the continued abstraction of computing resources from the infrastructure and the machines (virtual or otherwise) used to deliver those resources. Only by delivering this abstraction can you achieve the benefits of cloud computing—including improved agility and responsiveness, reduced total cost of ownership (TCO), and increased business alignment and focus. Most importantly, a private cloud promises to exceed the cost effectiveness of a virtualized infrastructure through higher workload density and greater resource utilization.

The Microsoft Private Cloud is a unique and comprehensive offering, built on four key "pillars."

- **All About the App:** Application-centric cloud platform that helps you focus on business value.
- **Cross-Platform from the Metal Up:** Cross-platform support for multi-hypervisor environments, operating systems, and application frameworks.
- **Foundation for the Future:** Microsoft Private Cloud lets you go beyond virtualization to a true cloud platform.
- **Cloud on Your Terms:** Ability to consume cloud on your terms, providing you the choice and flexibility of a hybrid cloud model through common management, virtualization, identity, and developer tools.

For further reading on the Microsoft Private Cloud, please see Microsoft Private Cloud Overview

Windows Server 2012 Cloud Infrastructure for Hosting Environments

A Microsoft cloud infrastructure forms the core compute, storage and networking capabilities that power any of the cloud service or delivery models. This means that, in addition to enterprises, cloud service providers (CSPs) or private cloud hosters can also take advantage of the core Windows Server 2012 platform cloud technologies to build hosted cloud solutions to support hosted private and public cloud solutions. The many innovations included with Windows Server 2012 enable you to deliver the security, isolation, performance, availability and scalability hosters require for providing commercial cloud services.

Private Cloud Architecture Principles

Private cloud principles provide general rules and guidelines to support the evolution of a cloud infrastructure, whether they be on-premises or in a hosted cloud environment. They are enduring, seldom amended, and inform and support the way a cloud fulfills its mission and goals. They also strive to be compelling and aspirational in some respects, since there needs to be a connection with business drivers for change. These principles are often interdependent and together form the basis on which a cloud infrastructure is planned, designed, and created:

- Resource Pooling
- Elasticity and Perception of Infinite Capacity

- Perception of Continuous Availability
- Drive Predictability
- Metering/Chargeback (Service Provider's Approach to Delivering IT)
- Multi-tenancy
- Security and Identity

All decisions regarding how the cloud infrastructure is designed should keep the principles in mind. For a detailed discussion of the private cloud principles, please see Private Cloud Principles, Concepts, and Patterns.

The Private Cloud Reference Model

IaaS is an application of the private cloud architecture principles to deliver a cloud infrastructure. You can use the reference model displayed in the following figure to ensure that you are delivering a holistic solution that spans all the layers required for a mature IaaS solution. The model acts as the "guide-rails" to assist cloud infrastructure designers' and architects' efforts to holistically address the development of a private cloud architecture. This model is used as a reference only, and some pieces are emphasized more than others in the technical reference architecture based on experience operating private clouds in real-world environments.

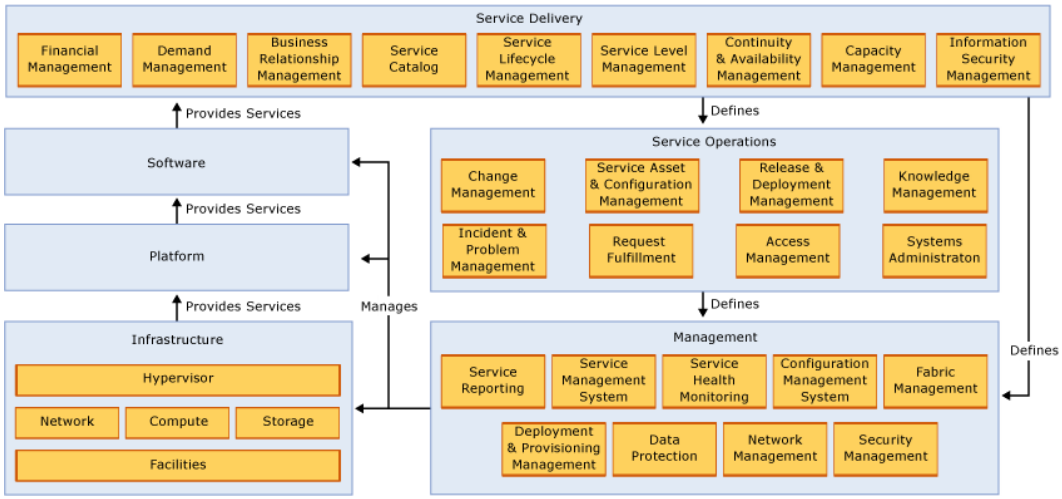

Figure 1. The private cloud reference model—IaaS view

The primary focus of this document is on the Infrastructure section of the Private Cloud Reference Model, where the focus is on the core network, compute, and storage components that let you build a cloud-ready infrastructure. However, we do not cover issues that apply to facilities considerations.

For more information and comprehensive coverage of the private cloud reference model, please see Private Cloud Reference Model.

Conceptual Architecture—Infrastructure

One of the key drivers of the layered approach to the infrastructure architecture presented here is to enable complex workflow and automation to be developed over time by creating a collection of simple automation tasks, assembling them into procedures that are managed by the management layer, and then creating workflows and process automation that are controlled by the orchestration layer. However, before you can implement these key features of a cloud solution, you need to understand key definitions and design points relevant to a well architected cloud infrastructure. The cloud infrastructure should address the following components:

1. Scale Units
2. Storage
3. Networking
4. Virtualization platform

Scale Units

In a modular architecture, the concept of a scale unit refers to the point where a module in the architecture can scale to before another module is required. For example, an individual server is a scale unit, it can be expanded to a certain point in terms CPU and RAM but beyond its maximums, an additional server is required to continue scaling. Each scale unit also has an associated amount of physical installation labor, configuration labor, and other labor overhead. With large scale units such as a pre-configured full rack of servers, the labor overhead can be minimized. When designing a cloud infrastructure, care should be taken to define the size of a scale unit.

Server scale limits are well published and include number and speed of CPU cores, maximum amount and speed of RAM, number and type of expansion slots, and so on. Particularly important are the number and type of onboard input/output (I/O) ports and the number and type of supported I/O cards. Infiniband, Ethernet, SAS (Serial Attached SCSI) and Fibre Channel expansion cards often provide multi-port options where a single card can have 4 ports. Additionally, in blade server architectures, there are often limitations in the amount of I/O card and/or supported combinations. It is important to be aware of these limitations and the oversubscription ratio between blade I/O ports and any blade chassis switch modules.

A single server is not a good scale unit for a private cloud solution due to the amount of overhead required to install and configure an individual server. Instead, scale units should be designed with considerations regarding historical rates of growth and how often your organization wants to add capacity to the infrastructure while providing adequate lead time so that overall service levels are not negatively impacted at any point in time. Instead, scale units should be designed with considerations regarding historical rates of growth and how often your organization wants to add capacity to the infrastructure while providing adequate lead time so that overall service levels are not negatively impacted at any point in time. Scale units should also take advantage of the private cloud principle of homogeneity.

For more information about scale units, please the Private Cloud Principles, Concepts, and Patterns.

Storage

Storage architecture is a critical design consideration for private cloud solutions. Storage and supporting storage networking are critical to the overall performance of the environment in addition to the overall cost, because storage tends to be one of the more costly items.

Storage architectures today have several layers including the storage array(s), the storage network, the storage protocol, and for virtualization, the file system utilizing the physical storage.

The traditional approach to storage arrays and storage protocols was to install a Storage Area Network (SAN) solution directly connected to the cloud compute (virtualization) cluster using either iSCSI or Fibre Channel. This storage design pattern provided for the key performance, reliability, availability and scalability requirements for enterprise storage, all of which are critical factors when considering storage design for a cloud infrastructure.

Windows Server 2012 provides a new option that can satisfy these cloud infrastructure storage requirements: Storage Spaces. Storage Spaces in Windows Server 2012 enables cost-effective, optimally used, highly available, scalable, and flexible storage solutions for business-critical (virtual or physical) deployments. You can use cost-effective SAS drives in a JBOD (Just a Bunch of Disks) enclosure with Storage Spaces and benefit from many of the same capabilities of expensive SAN-based storage arrays.

Historically, iSCSI and Fibre Channel have been the storage protocols of choice. However, Windows Server 2012 offers another option – file based storage over SMB 3.0. Using this approach, you can create separate compute and storage clusters that enable you to scale compute and storage independently. This also has the potential to greatly simplify your cloud infrastructure as management of file shared based storage is typically simpler than SAN storage. The SMB 3.0 protocol makes it possible for the compute cluster to access the storage cluster at near Direct Attached Storage (DAS) speeds over a 10 Gbps network fabric.

In addition to highly available storage over SMB 3.0, Windows Server 2012 cloud infrastructure administrators have the option to use another file based storage option – NFS high-availability transparent failover. This makes it possible to deploy Server for NFS in a clustered configuration and take advantage of better resilience to hardware and software outages that may afflict individual cluster nodes.

We will cover detailed cloud storage infrastructure design decisions later in this document.

Networking

Many network architectures include a tiered design with three or more tiers such as core, distribution, and access. Designs are driven by the port bandwidth and quantity required at the edge, in addition to the ability of the distribution and core tiers to provide higher speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, spanning tree or other loop-avoidance technologies.

A dedicated management network is a frequent feature of advanced data center virtualization solutions. Most virtualization vendors recommend that hosts be managed via a dedicated network so that there is no competition with tenant traffic and to provide a degree of separation for security and ease of management purposes. This historically implied dedicating a network

adapter per host and port per network device to the management network. However, with new technologies included with Windows Server 2012, dedicating network adapters to each traffic class is no longer required or recommended. Windows Server 2012 includes new capabilities that enable you to converge your network traffic classes to a single NIC or NIC team. .

With advanced data center virtualization, a frequent use case is to provide isolated networks where different owners such as particular departments or applications are provided their own dedicated networks. In the past, multi-tenant networking referred to using technologies such as virtual local area networks (VLANs) or Internet Protocol Security (IPsec) isolation techniques to provide dedicated networks that utilize a single network infrastructure or wire. However, VLANs are complex to configure and manage, and are not scalable past 4094 VLANs. Windows Server 2012 Hyper-V includes a number of improvements in the Hyper-V virtual switch that enable you to get the isolation you need without using VLANs by providing new features such as Private VLANs (PVLANs), Port ACLs and virtual networking, sometimes referred to as software defined networking (SDN).

Managing the network environment in a private cloud can present challenges that must be addressed. Ideally, network settings and policies are defined centrally and applied universally by the management solution, such as the Virtual Machine Manager component of Microsoft System Center 2012 with Service Pack 1 (SP1). In the case of IPsec-based isolation, this can be accomplished using Active Directory Domain Services (AD DS) and Active Directory Group Policy to control firewall settings across the hosts and guests, in addition to the IPsec policies controlling network communication. In addition, you can optimize VM density when using IPsec because Windows Server 2012 now includes support for IPsec Talk Offload (IPsecTO) for guest virtual machines.

For VLAN-based network segmentation, several components including the host servers, host clusters, Virtual Machine Manager, and the network switches must be configured correctly to enable both rapid provisioning and network segmentation. With Hyper-V and host clusters, identical virtual switches should be defined on all nodes in order for a virtual machine to be able to failover to any node and maintain its connection to the network. At large scale, this can be accomplished via Windows PowerShell scripting.

In a Windows Server 2012–based cloud infrastructure, VLAN assignments are typically limited to the major traffic classes; these are the host traffic and tenant traffic classes. Bandwidth guarantees can be assigned using Windows Server 2012 Quality of Service (QoS) so that subclasses of traffic, such as cluster traffic, management traffic, live migration traffic, and storage traffic, have the bandwidth they require. In addition, it's important to minimize or avoid the use of VLANs for tenant traffic, as VLANs exclude the option of using Windows Server 2012 Network Virtualization. We will discuss cloud infrastructure network design decisions in detail later in this document.

Virtualization Platform

The virtualization layer is one of the primary enablers for private cloud. The decoupling of hardware, operating systems, data, applications, and user state opens a wide range of options for optimal management and distribution of workloads across the physical infrastructure. Windows

Server 2012 Hyper-V provides the ability to migrate running virtual machines from one server to another with zero downtime and many other Hyper-V features that provide a rich set of capabilities. These capabilities can be utilized by the automation, management, and orchestration layers to maintain desired states (such as load distribution) or to proactively address decaying hardware or other issues that would otherwise cause faults or service disruptions.

As with the hardware layer, the virtualization layer must be able to be managed by the automation, management, and orchestration layers. The abstraction of software from hardware that virtualization provides moves the majority of management and automation into the software space, instead of requiring you to perform manual operations on physical hardware.

Cloud Infrastructure Design

In the next three sections we will discuss design considerations and options for designing a cloud infrastructure using Windows Server 2012 platform capabilities. This discussion will include design discussions around:

1. Storage architecture
2. Network architecture
3. Compute (virtualization and cluster) architecture

After a review of the design decisions you will need to make around the storage, network and compute architectures, there will be a discussion on three design patterns that based on key design decisions made in each of these domains.

Designing the Cloud Storage Infrastructure

The storage design for any virtualization-based solution is a critical element that is typically responsible for a large percentage of the solution's overall cost, performance, and agility.

The storage solution should provide transport-independent, seamless data access using block-level and/or file-level protocols from the same platform. A storage solution can provide block-level data access over a Fibre Channel SAN fabric using Fibre Channel Protocol (FCP) and over an IP-based Ethernet network using iSCSI or Fibre Channel over Ethernet (FCoE). Infiniband is another option for high performance storage networking. File-access protocols such as NFS, SMB 3.0, HTTP, or FTP provide file-level access over an IP-based Ethernet network.

Storage Options

This section discusses a number of storage options available to you that will inform some key design decisions you need to make when considering alternatives. Specific options discussed include:

1. Storage protocols
2. Storage network
3. Cluster Shared Volumes (CSVs)

Storage Protocols

Although many storage options exist, organizations should choose their storage devices based on their specific data-management needs. Storage devices are typically modular and flexible midrange and high-end SANs. Modular midrange SANs are procured independently and can be chained together to provide greater capacity. They are efficient, can grow with the environment as needed, and require less up-front investment than high-end SANs. Large enterprises and hosters might have larger storage demands and might need to serve a larger set of customers and workloads. In this case, high-end SANs can provide the highest performance and capacity. High-end SANs typically include more advanced features such as continuous data availability through technologies like replication and clustering. However, the price of these high-end SANs can be prohibitive, and you should factor in the costs with the value provided over other options.

Fibre Channel has historically been the storage protocol of choice for enterprise data centers for a variety of reasons, including good performance and low latency. Over the past several years, however, the advancing performance of Ethernet from 1 Gbps to 10 Gbps and beyond has led to great interest in storage protocols that make use of Ethernet transport—such as iSCSI, and more recently FCoE. SMB 3.0 can also be used for file based storage access and provide the same or superior performance when compared to iSCSI of FCoE.

A key advantage of the protocols that use the Ethernet transport is the ability to use a "converged" network architecture in which a single Ethernet infrastructure serves as the transport for both LAN and storage traffic. FCoE is an emerging technology that brings the benefits of using an Ethernet transport while retaining the advantages of the Fibre Channel protocol and the ability to use Fibre Channel storage arrays. SMB 3.0 can be used to connect to file servers that host virtual machine files. Windows Server 2012 is designed to fully support a converged network infrastructure.

In the past a common practice in large-scale Hyper-V deployments was to use both Fibre Channel and iSCSI. Fibre Channel and iSCSI can provide the host storage connectivity. Similarly, both Fibre Channel and iSCSI can be used directly by guests—for example, for the shared disks in a guest cluster. This is made possible by the new Windows Server 2012 feature that surfaces Fibre Channel HBAs to guest virtual machines. In Windows Server 2012 SMB 3.0 can provide the same advantages. In addition, with Windows Server 2012 you can take advantage of a new storage option called Storage Spaces where you can attach SAS JBODs to a file server cluster and access the virtual machine files from the compute cluster from over the Ethernet network.

Note

Windows Server 2012 now supports Offload Data Transfer (ODX) which is a new high performance storage protocol that significantly reduces the processing and time it takes to copy files from one storage unit to another. You can learn more about ODX at Windows Offloaded Data Transfers overview.

Storage Network

iSCSI, SMB 3.0 and FCoE use an Ethernet transport for storage networking. This provides another architecture choice in terms of whether to use a dedicated Ethernet network with

separate switches, cables, paths, and other infrastructure or, instead, to use a converged network in which multiple traffic types are run over the same cabling and infrastructure.

The storage solution must provide logical or physical isolation between storage and Ethernet I/O. If it's a converged network, QoS must be provided to guarantee storage performance. The storage solution must provide iSCSI or Fibre Channel connectivity for guest clustering and fully redundant, independent paths for storage I/O.

Standards-based converged network adapters, switches, and Fibre Channel storage arrays should be used for FCoE. If iSCSI, FCoE or SMB 3.0 is used, you can assign a dedicated network adapter for storage traffic or use Windows Server 2012 new Hyper-V switch capabilities to create a converged network infrastructure.

Cluster Shared Volumes

CSVs (Cluster Shared Volumes) provide a distributed file-access solution so that multiple nodes in the cluster can simultaneously access the same NTFS file system. For example, virtual machines that are distributed across multiple cluster nodes can access their virtual hard disk (VHD) files, even if the VHD files are on a single disk (logical unit number [LUN]) in storage. This also enables virtual machines to move to any node in the cluster, since each cluster node has access to files contained in the CSV. The clustered virtual machines can also all fail over independently of one another and no longer present themselves as a single resource group.

The following is a partial list of new CSV functionality in Windows Server 2012 that you can use to when designing the shared storage infrastructure for your private cloud:

- Storage system for scale-out file servers, which can provide continuously available and scalable file-based SMB server application storage. The scale-out file server feature is used when you choose to host the virtual machine files in a storage cluster. For more information, see Support for scale-out file servers.
- A single consistent file name space, with CSV volumes that now appear as CSV File System (CSVFS). The underlying technology is still the NTFS file system, and volumes are still formatted with NTFS. In a clustered environment, the CSV must be formatted as NTFS as ReFS is not supported on CSVs in Windows Server 2012.
- Direct I/O for file data access, including sparse files, which enhances virtual machine creation and copy performance. Redirected I/O is used when Storage Spaces is deployed using either mirrored or parity based resiliency.
- No external authentication dependencies, which provides improved performance and resiliency. The cluster can now start and the CSV will be mounted even if an Active Directory domain controller is not available when the cluster reboots.
- Support for SMB 3.0 storage for Hyper-V and applications such as SQL Server. For more information on SMB 3.0 functionality, see Server Message Block overview.
- Integration with SMB Multichannel and SMB Direct, which allow CSV traffic to stream across multiple networks in the cluster and to take advantage of network adapters that support remote direct memory access (RDMA). SMB also provides for transparent failover so that in the event that a member of the file server cluster becomes unavailable, connections to the virtual machine files continue to be available and with no downtime.

- Support for BitLocker Drive Encryption for CSVs in addition to traditional cluster disks. This is particularly useful when clusters are deployed in branch offices and other low security environments.
- The capability to make storage visible to only a subset of nodes, which enables scenarios for a single cluster that contains application and data nodes. This is an example of an "asymmetric storage cluster". In this design pattern, some of the nodes in the cluster are connected to storage and some of the nodes are dedicated to the compute role. The compute nodes then connect to file based storage over the SMB 3.0 protocol.
- Integration with the Storage Spaces feature in Windows Server 2012, which can provide virtualized storage on clusters of inexpensive SAS disks. For more information about Storage Spaces, see File and Storage Services overview.
- Ability to scan and repair volumes with zero offline time due to new capabilities deliver by the Windows Server 2012 CHKDSK application. For more information on advances in CHKDSK, please see Multi-terabyte volumes.

CSVs provide not only shared access to the disk, but also storage path I/O fault tolerance (dynamic I/O redirection). In the event the storage path on one node becomes unavailable, the I/O for that node will be rerouted via SMB 3.0 through another node. This feature can use any Cluster Communications Network and further increases the need for high-speed 10 GB Ethernet networks.

CSVs maintain metadata information about the volume access and require that some I/O operations take place over the cluster communications network. One node in the cluster is designated as the coordinator node and is responsible for these disk operations. Virtual machines, however, have direct I/O access to the volumes, and only use the dedicated storage paths for disk I/O, unless a failure scenario occurs as described above. The exception to this is with Storage Spaces when resiliency is added to the CSV.

CSV Requirements

Requirements for implementing a Windows Server 2012 private cloud that uses CSVs include:

- All cluster nodes must use Windows Server 2012.
- All cluster nodes must use the same drive letter for the system disk.
- All cluster nodes must be on the same logical network subnet.
- SMB must be enabled for each network on each node that will carry CSV cluster communications.
- "Client for Microsoft Networks" and "File and Printer Sharing for Microsoft Networks" must be enabled in the network adapter's properties to enable all nodes in the cluster to communicate with the CSV.
- The Hyper-V role must be installed on any cluster node that might host a virtual machine.

Note that when you configure the failover cluster using the failover cluster wizard, most of these configuration requirements will be handled for you.

CSV Volume Sizing

Because all cluster nodes can access all CSV volumes simultaneously, you can now use standard LUN allocation methodologies based on performance and capacity requirements of the workloads running within the virtual machines themselves. In general, it's considered a good idea to isolate the virtual machine operating system I/O from the application data I/O, in addition to application-specific I/O considerations such as segregating databases and transaction logs and creating SAN volumes and/or Storage Spaces–based storage pools that factor in the I/O profile itself (for example, random read and write operations versus sequential write operations).

CSV architecture differs from traditional clustered file systems, which frees it from common scalability limitations. As a result, there is no special guidance for scaling the number of Hyper-V nodes or virtual machines on a CSV volume other than ensuring that the overall I/O requirements of the expected virtual machines running on the CSV are met by the underlying storage system and storage network.

While rare, disks and volumes can enter a state where a chkdsk is required, which with large disks can take a long time to complete, causing downtime of the volume during this process somewhat proportional to the volume's size. However, significant improvements have been made to the chkdsk feature in Windows Server 2012 so that for even very large volumes with tens of millions of files, it should not take more than a minute to complete the check.

Each enterprise application you plan to run within a virtual machine can have unique storage recommendations and even perhaps virtualization-specific storage guidance. That guidance applies to use with CSV volumes as well. The important thing to keep in mind is that all virtual machine virtual disks running on a particular CSV will contend for storage I/O.

Also worth noting is that individual SAN LUNs do not necessarily equate to dedicated disk spindles. A SAN Storage Pool or redundant array of independent disks (RAID) array can contain many LUNs. A LUN is simply a logical representation of a disk provisioned from a pool of disks. Therefore, if an enterprise application requires specific storage I/O per second (IOPS) or disk-response times you must consider all the LUNs in use on that Storage Pool. An application which would require dedicated physical disks were it not virtualized might require dedicated Storage Pools and CSV volumes running within a VM.

Note

When considering the use of Storage Spaces to host your CSVs, you might want to consider the fact that when enabling resiliency for the CSV, only a single node can write to it directly; the other nodes will use I/O redirection to the coordinator node in this scenario. For this reason, you may want to create multiple CSVs (perhaps one per cluster member) and distribute your virtual machine files across the CSVs so that a single node does not receive a disproportionate amount of network traffic due to I/O redirection.

CSV Design Patterns

Single CSV per Cluster

In the Single CSV per Cluster design pattern, storage is configured to present a single large LUN to all the nodes in the cluster. The LUN is configured as a CSV in Failover Clustering. All virtual-machine related files (such as VHDs and configuration files) belonging to the virtual machines hosted on the cluster are stored on the CSV. Optionally, data deduplication functionality provided by the SAN or the Windows Server 2012 deduplication feature can be utilized. Note that when you configure a file server cluster using the scale-out file server capability, Windows Server 2012 deduplication will not be available. However, you can still take advantage of Windows Server 2012 data deduplication in your virtual machine libraries.

Multiple CSVs per Cluster

In the Multiple CSVs per Cluster design pattern, storage is configured to present two or more large LUNs to all the nodes in the host cluster. The LUNs or Storage Spaces virtual disks are configured as a CSV in Failover Clustering. All virtual-machine related files (such as VHDs and configuration files) belonging to the virtual machines hosted on the cluster are stored on the CSVs. Optionally, data deduplication functionality provided by the SAN or Windows Server 2012 deduplication can be utilized, with some limitations as previously mentioned.

For both the Single and Multiple CSV patterns, each CSV has the same I/O characteristics, so each individual virtual machine has all its associated VHDs stored on one of the CSVs. Alternatively, you can use storage tiering and assign CSVs to different tiers. This is useful if you want to present storage options as part of your private cloud service catalog.

Multiple I/O Optimized CSVs per Cluster

In the Multiple I/O Optimized CSVs per Cluster design pattern, the SAN is configured to present multiple LUNs or Storage Spaces virtual disks to all the nodes in the cluster but they are optimized for particular I/O patterns such as fast sequential read performance, or fast random write performance. The volumes are configured as CSVs in Failover Clustering. All VHDs belonging to the virtual machines hosted on the cluster are stored on the CSVs but targeted to the most appropriate CSV for the given I/O needs.

In the Multiple I/O Optimized CSVs per Cluster design pattern, each individual virtual machine has all its associated VHDs stored on the appropriate CSV following required I/O requirements.

Note that a single virtual machine can have multiple VHDs, and each VHD can be stored on a different CSV (provided all CSVs are available to the host cluster the virtual machine is created on).

In addition, you can further optimize the CSV I/O by taking advantage of CSV read caching with write through. Learn more about this capability at How to enable CSV Cache.

Storage Design

A highly available storage design, which could use either a SAN or SAS-based JBOD solution and Storage Spaces, should have no single points of failure, including:

- Redundant power from independent power distribution units (PDUs)
- Redundant storage controllers
- Redundant storage paths (supported, for example, by redundant target ports of network adapters per controller, redundant Fibre Channel or IP network switches, and redundant cabling)
- Data storage redundancy similar to what occurs with volume mirroring or parity, or synchronous or asynchronous replication

You will need to address the following elements when designing or modifying your storage solution as the basis of your Microsoft Private Cloud storage infrastructure:

- Performance
- Drive types
- Multipathing
- Fibre Channel SAN
- iSCSI SAN
- Storage Spaces
- Data deduplication
- Thin provisioning
- Volume cloning
- Volume snapshot

Performance

Storage performance is a mix of drive, interface, controller, cache, protocol, SAN, host bus adapter (HBA), driver, and operating system considerations. The overall performance of the storage architecture is typically measured in terms of maximum throughput and/or maximum IOPS for a given latency or response time. Although each of these performance measurements is important, IOPS for a given latency is the most relevant to server virtualization in private cloud environments.

Another important consideration is what type of cloud infrastructure do you want to provide? In most cases the private cloud will present a general purpose infrastructure that is not designed around specific workloads. Workloads that require specific disk performance characteristics that lie outside the capabilities of a general purpose private cloud infrastructure may be better served by creating a dedicated infrastructure for those workloads.

Drive Types

The type of hard drive used in the host server or the storage array has the most significant impact on the overall storage architecture performance. As with storage connectivity, high IOPS and low latency are more critical than maximum sustained throughput when it comes to host server sizing and guest performance. When selecting drives, this translates into selecting those with the highest rotational speed and lowest latency possible. Using 15K revolution per minute (RPM) drives over 10K RPM drives can result in up to 35 percent more IOPS per drive.

You will also want to consider storage tiering, so that you can provide options to consumers of the cloud service. A storage tiering scheme might include a high-performance tier, a medium-performance tier, and a high-capacity tier. These tiers can be included in the service catalog. Windows Server 2012 includes support for newer 4K high-capacity drives, which can be part of one or more storage tiers. In addition, SSD drives will become the norm in the future as prices continue to come down on this ultra-fast storage option.

Multipathing

In order to support a highly resilient private cloud storage infrastructure, multipathing should be used. Among the features enabling high availability for connecting servers based on Windows to SANs is integrated Multipath I/O (MPIO) support. Microsoft MPIO architecture supports iSCSI, Fibre Channel, and SAS SAN connectivity by establishing multiple sessions or connections to the storage array. Multipathing solutions use redundant physical path components—adapters, cables, and switches—to create logical paths between the server and the storage device. In the event that one or more of these components fails, causing the path to fail, multipathing logic uses an alternate path for I/O so that applications can still access their data. Each Ethernet interface or HBA should be connected by using redundant switch infrastructures to provide continued access to storage in the event of a failure in a storage fabric component.

Failover times vary by storage vendor, and can be configured by using timers in the Microsoft iSCSI Software Initiator driver, or modifying the Fibre Channel host bus adapter driver parameter settings.

Fibre Channel SAN

Fibre Channel is a robust and high-speed storage protocol that supports multipathing through Microsoft Windows Server 2012 MPIO. Windows Server 2012 provides Fibre Channel ports for both host and guest operating systems. Support for guest operating system connectivity through Fibre Channel is new in Windows Server 2012 and allows you to connect to Fibre Channel directly from within virtual machines. This feature enables you to virtualize workloads that use direct access to Fibre Channel storage and allows you to cluster guest operating systems over Fibre Channel.

Mid-range and high-end storage arrays are capable of advanced storage functionality that helps offload certain management tasks from the hosts to the SANs. Virtual Fibre Channel presents an alternate hardware-based I/O path to the Windows software virtual hard disk stack. This allows you to use the advanced functionality offered by your Fibre Channel SANs directly from Hyper-V virtual machines. For example, you can use Hyper-V to offload storage functionality (for example, taking a snapshot of a LUN) on the SAN hardware by using a hardware volume shadow copy service (VSS) provider from within a Hyper-V virtual machine.

Windows Server 2012 Hyper-V allows you to define virtual SANs on the host to accommodate scenarios where a single Hyper-V host is connected to different SANs through multiple Fibre Channel ports. A virtual SAN defines a named group of physical Fibre Channel ports that are connected to the same physical SAN. For example, assume that a Hyper-V host is connected to two SANs—a production SAN and a test SAN. The host is connected to each SAN through two

physical Fibre Channel ports. In this example, you might configure two virtual SANs—one named "Production SAN" that has the two physical Fibre Channel ports connected to the production SAN, and one named "Test SAN" that has two physical Fibre Channel ports connected to the test SAN. You can use the same technique to name two separate paths to a single storage target.

You can configure as many as four virtual Fibre Channel adapters on a virtual machine and associate each one with a virtual SAN. Each virtual Fibre Channel adapter connects with one WWN address or two WWN addresses to support live migration. You can set each WWN address automatically or manually.

Hyper-V in Windows Server 2012 can use the multipath I/O (MPIO) functionality to ensure continuous connectivity to Fibre Channel storage from within a virtual machine.

You can use MPIO functionality with Fibre Channel in the following ways:

- Virtualize workloads that use MPIO. Install multiple Fibre Channel ports on the host, and use MPIO to provide highly available connectivity to the LUNs accessible by the host.
- Configure multiple virtual Fibre Channel adapters inside a virtual machine, and use a separate copy of MPIO within the guest operating system of the virtual machine to connect to the LUNs that the virtual machine can access. This configuration can coexist with a host MPIO setup.
- Use different DSMs for the host or each virtual machine. This approach allows live migration of the virtual machine configuration, including the configuration of DSM and connectivity between hosts and compatibility with existing server configurations and DSMs.

iSCSI SAN

Unlike a Fibre Channel–connected SAN, which is on its own Fibre Channel network, the iSCSI SAN can be on its own isolated Ethernet or Infiniband network or it can benefit from Windows Server 2012 capabilities that enable converged networking to support isolation that is so critical in a multi-tenant architecture. Any networking standard practice method for achieving this goal is acceptable, including a physically separate, dedicated storage network and a physically shared network with the iSCSI SAN running on a private VLAN. Alternatively, you can use new QoS and port access control list (ACL) features in Windows Server 2012 to converge the storage network with other non-tenant and tenant network flows.

In Windows Server 2012, the iSCSI Software Target is available as a built-in feature under the File and Storage Service role. As an inbox feature, the management experience is integrated with Server Manager and the failover cluster management console.

When used in a cloud infrastructure, the iSCSI Software Target is ideal for:

- **Network/diskless boot:** By using boot-capable network adapters or a software loader, you can deploy hundreds of diskless host servers. This is ideal for large deployments of identical operating system images, such as a Hyper-V compute cluster.
- **Server application storage:** Some applications require block storage. The iSCSI Software Target can provide these applications with continuously available block storage. Since the storage is remotely accessible it can also consolidate block storage for central or branch office locations.

- **Heterogeneous storage:** iSCSI Software Target supports non-Windows iSCSI initiators, making it easy to share storage on Windows Servers in a mixed environment.
- **Development/test/demo/lab clouds:** When the iSCSI Software Target feature is enabled, it turns any server running Windows Server into a network-accessible block storage device. This is ideal for testing applications prior to deployment on SAN storage.

Enabling Microsoft iSCSI Software Target to provide block storage takes advantage of your existing Ethernet network. No additional hardware is needed. To obtain high availability, consider setting up a highly available storage cluster. With a highly available storage cluster, you will need shared storage for the cluster—such as Fibre Channel storage or an SAS storage array, which can be configured to use Storage Spaces.

If you enable guest clustering, you need to provide block storage. Any servers running Windows Server software with Microsoft iSCSI Software Target can provide block storage. An iSCSI Target Server enables you to network boot multiple computers from a single operating system image that is stored in a centralized location. iSCSI Software Target in Windows Server 2012 can boot hundreds of computers by using a single operating system image, and it provides the several key benefits.

For example, by using differencing virtual disks, you can use a single operating system image (the "golden image") to boot up to 256 computers. In a deployment of Windows Server 2008 R2 HPC Edition, the operating system image is approximately 20 GB. A common deployment is to have two mirrored disk drives that act as the boot volume. Rounding the operating system storage to 40 GB per instance, you would need approximately 10 TB of storage—for only the operating system image—to boot 256 computers. With iSCSI Software Target boot, however, you will use 40 GB for the operating system base image, and 2 GB for differencing virtual hard disks (VHDs) per server instance, totaling 552 GB for the operating system images. This provides a savings of over 90 percent on storage for the operating system images alone.

Other advantages to using Windows Server 2012 iSCSI Target Boot in your private cloud storage infrastructure include:

- **Controlled operating system images make it more secure and easy to manage.** Some enterprises require that data be secured by physically locking storage in a centralized location. In this scenario, servers access the data remotely, including the operating system image. With iSCSI Software Target boot, you can centrally manage the operating system boot images, and control which applications to put in the golden image.
- **Rapid deployment.** Because the golden image is a sysprepped operating system image, when the computers boot from the golden image, they skip the file copying and installation phase that occurs during Windows Setup, and they go straight to the customization phase. In our testing, we deployed 256 computers in 34 minutes.
- **Fast recovery.** Because the operating system images are hosted on the iSCSI Target Server, if the diskless client needs to be replaced, the new computer can point to the operating system image, and boot up immediately.
- **Centralized updating of the golden image.** You can update all host systems in your private cloud by updating the golden image. When cluster members are restarted they reboot from the new golden image. Post-boot tasks can be integrated into the cloud fabric controller so that the updated server is joined back into the cluster.

SAN boot is a solution that has been offered from various vendors. Now with Windows Server 2012, the new iSCSI Software Target feature provides this network boot capability on commodity hardware which can lead to significant cost savings since there are no special hardware requirements. In data centers with large-scale deployments, the design should be validated against specific hardware.

Note

For reference, Microsoft internal testing indicated that for a 256-iSCSI boot deployment, 24 x 15k-RPM disks in a RAID 10 configuration were required for storage. A network bandwidth of 10 GB is optimal. A general estimate is 60 iSCSI boot servers per 1 GB network adapter. However, an iSCSI boot-capable network adapter is not required for this scenario. If the network adapter does not support it, a software boot loader can be used (such as iPXE open source boot firmware).

Storage Spaces

Storage Spaces in Windows Server 2012 enables cost-effective, optimally used, highly available, scalable, and flexible storage solutions for business-critical (virtual or physical) deployments. Windows Server 2012 delivers sophisticated storage virtualization capabilities, which enables you to use cost-effective industry-standard storage, for single-node and scalable multi-node deployments

Storage Spaces delivers storage virtualization capabilities within Windows Server 2012. The storage stack has been fundamentally enhanced to incorporate two new abstractions:

- Storage pools are administrative units of physical disks. Pools permit storage aggregation, elastic capacity expansion, and delegated administration.
- Storage spaces are virtual disks with associated attributes such as a desired level of resiliency, thin or fixed provisioning, automatic or controlled allocation on heterogeneous storage media, and precise administrative control.

Storage Spaces is completely integrated with failover clustering for high availability, and it is integrated with CSV for scale-out deployments. Storage Spaces includes the following features:

1. **Storage pools**. Storage pools are the fundamental building blocks for Storage Spaces. You can flexibly create storage pools based on the needs of the deployment. For example, given a set of physical disks, you can create one pool (by using all the available physical disks) or multiple pools (by dividing the physical disks as required). Furthermore, to maximize the value from storage hardware, you can map a storage pool to combinations of hard disks as well as solid-state drives (SSDs). Pools can be expanded dynamically by simply adding drives.
2. **Multitenancy**. Administration of storage pools can be controlled through access control lists (ACLs) and delegated on a per-pool basis, supporting hosting scenarios that require tenant isolation. Storage Spaces follows the familiar Windows security model; therefore, it can be fully integrated with Active Directory Domain Services.
3. **Resilient storage**. Storage Spaces support two optional resiliency modes: mirroring and parity. Per-pool support for disks that are reserved for replacing failed disks (hot spares), background scrubbing, and intelligent error correction allow continuous service availability

despite storage component failures. Note that only the mirroring option is available when deploying resiliency to cluster shared volumes.

4. **Continuous availability**. Storage Spaces is fully integrated with failover clustering, which allows it to deliver continuously available service deployments. One or more pools can be clustered across multiple nodes within a single cluster. Storage spaces can then be instantiated on individual nodes, and the storage will seamlessly fail over to a different node when necessary (in response to failure conditions or due to load balancing). Integration with CSVs permits scale-out access to data.
5. Optimal storage use. Server consolidation often results in multiple data sets sharing the same storage hardware. Storage Spaces supports thin provisioning to allow businesses to easily share storage capacity among multiple unrelated data sets and thereby maximize capacity use. Trim support permits capacity reclamation when possible.

   Storage Spaces is the ideal cost-effective solution for cloud infrastructures that are designed to support a generic population of virtual workloads that do not require specialized storage requirements or exceptional IOPS.

Data Deduplication

Data deduplication can yield significant storage cost savings in virtualized environments. Some common considerations are performance implications during the deduplication cycle and achieving maximum efficiency by locating similar data types on the same volume or LUN. Data deduplication can be carried out at the SAN level, or if you are using SAS storage with Storage Spaces, you can use the new Windows Server 2012 built-in deduplication feature.

Windows Server 2012 includes the following data duplication features that you can take advantage of in your cloud storage design:

- **Capacity optimization**. Data deduplication in Windows Server 2012 stores more data in less physical space. It achieves greater storage efficiency than was possible in previous releases with single-instance storage (SIS) or NTFS compression: data deduplication that can deliver optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data.
- **Scale and performance**. Windows Server 2012 data deduplication can run on dozens of large volumes of primary data simultaneously without affecting other workloads on the server. Low impact on the server workloads is maintained by throttling of CPU and memory resources consumed. You also have the option to set times when data deduplication should run, specify the resources available to it, and establish policies on file selection.
- **Reliability and data integrity**. When data deduplication is applied, it is essential to maintain data integrity. Windows Server 2012 takes advantage of checksum, consistency, and identity validation to ensure data integrity. Windows Server 2012 data deduplication also maintains redundancy to ensure that the data is recoverable in the event of data corruption.

Note that data deduplication is a feature that potentially will process all of the data on a selected volume, so careful planning should be done to determine if a server and attached volumes are suitable candidates for deduplication prior to enabling the feature. We strongly advise that during deployment of deduplication, you make regular backups of important data.

Note

Windows Server 2012 data deduplication is not available on CSVs running on a highly available scale-out file server cluster.

Thin Provisioning

In private cloud virtualization environments, thin provisioning is a common practice. This allows efficient use of the available storage capacity. The LUN or Storage Spaces virtual disk and corresponding CSV can grow as needed, typically in an automated fashion. However, storage can become overprovisioned in this scenario, so careful management and capacity planning are critical. Because physical allocation of data within a thin-provisioned volume is done on demand, theoretically the volume size can be set to a very high value that can easily keep all application data and snapshot copies.

The unallocated space in the volume is not exclusively reserved for the volume itself; therefore, all other applications can benefit from the shared pool of unallocated storage. Additionally, the volume size limits when using deduplication should be taken into account because the maximum sizes depend on the storage controllers when using SAN-based storage solutions.

Sophisticated and costly storage solutions offer just-in-time allocations (also known as thin provisioning) and the ability to reclaim storage that is no longer needed (also known as trim). Windows Server 2012 integrates with these sophisticated storage solutions to enable organizations to get the most out of their storage infrastructures at no additional cost. You can maximize the benefits of sophisticated storage infrastructure that is accessed through Windows Server 2012.

Requirements for thin provisioning by Windows Server 2012 include:

- Storage infrastructure that complies with the certification that is required for Windows Server 2012
- Standards-compliant hardware for identification.

To confidently deploy sophisticated storage solutions that support just-in-time allocation, you need to know that you can provision additional capacity as needed. Windows Server 2012 identifies thinly provisioned virtual disks, provides standardized notifications when use thresholds are crossed, and provides a platform that enables applications to release storage when it is no longer needed.

When designing for thin provisioning in your cloud infrastructure, consider the following Windows Server 2012 capabilities:

- **Identification**. Windows Server 2012 uses a standardized method to detect and identify thinly provisioned virtual disks, thus enabling additional capabilities delivered by the storage stack that is provided in the operating system and through storage-management applications.
- **Notification**. When configured physical storage use thresholds are reached, Windows Server 2012 notifies the administrator through events, which helps you to take appropriate action as soon as possible. These events can also be used for automated actions by management applications, such as System Center.
- **Optimization**. Windows Server 2012 provides a new application programming interface (API) that lets applications return storage when it is no longer needed. NTFS issues trim

notifications in real time when appropriate. Additionally, trim notifications are issued as part of storage consolidation (optimization), which is performed regularly on a scheduled basis.

Note
Server consolidation often results in multiple data sets sharing the same storage hardware. Storage Spaces supports thin provisioning to allow you to easily share storage capacity among multiple unrelated data sets and thereby maximize capacity use. Trim support permits capacity reclamation when possible.

Volume Cloning

Volume cloning is another common practice in virtualization environments. This can be used for both host and virtual machine volumes to dramatically decrease host installation times and virtual machine provisioning times.

Rapid provisioning is a common feature for private cloud implementations. In these environments, the expectation is that end users or departmental administrators will deploy virtual machines. Because of this, the system needs to respond rapidly to provisioning requests, and must scale those requests to accept large numbers of simultaneous requests. Clone-based provisioning has several key advantages over traditional copy-based provisioning.

Volume Snapshot

SAN volume snapshot copies are a common method of providing a point-in-time, instantaneous backup of a SAN volume or LUN. These snapshot copies are typically block-level and only use storage capacity as blocks change on the originating volume. Some SANs provide tight integration with Hyper-V, integrating both the Hyper-V VSS Writer on hosts and volume snapshot copies on the SAN. This integration provides a comprehensive and high-performing backup and recovery solution.

Windows Server 2012 has the ability to create application-consistent snapshots of the server application data. In Windows Server 2012, this is accomplished using the VSS infrastructure. VSS for SMB file shares extends the VSS infrastructure to perform application-consistent shadow copies of data stored on remote SMB file shares for backup and restore purposes. In addition, VSS for SMB file shares enable backup applications to read the backup data directly from a shadow copy file share rather than involving the server application computer during the data transfer. Because this feature takes advantage of the existing VSS infrastructure, it is easy to integrate with existing VSS-aware backup software and VSS-aware applications, such as Hyper-V.

Storage Automation

One of the objectives of the Microsoft private cloud solution is to enable rapid provisioning and deprovisioning of virtual machines. Doing so on a large scale requires tight integration with the storage architecture, in addition to robust automation. Provisioning a new virtual machine on an already existing LUN is a simple operation. However, provisioning a new CSV LUN and adding it to a host cluster are relatively complicated tasks that should be automated.

Historically, many storage vendors have designed and implemented their own storage management systems, APIs, and command-line utilities. This has made it a challenge to use a common set of tools and scripts across heterogeneous storage solutions.

Windows Server 2012 enables storage management that is comprehensive and fully scriptable, and administrators can manage it remotely. A WMI-based interface provides a single mechanism through which to manage all storage, including non-Microsoft intelligent storage subsystems and virtualized local storage (known as Storage Spaces). Additionally, management applications can use a single Windows API to manage different storage types by using standards-based protocols such as Storage Management Initiative Specification (SMI-S).

Designing the Cloud Network Infrastructure

The demand for hosted cloud requires integrated protection, the ability to add virtual appliances, and the capability to adapt virtual networking to the data center infrastructure. One of the most important issues to consider while building your private cloud is to ensure that the networking infrastructure is able to provide security, high availability, high performance, predictability, and resiliency, together with the ability to adapt for future cloud requirements. These prerequisites should be provided not only on the virtual layer by taking advantage of the Hyper-V networking features, but also in the physical layer where the switches and routers are connected. Figure 2 shows a common private cloud infrastructure and different layers that must be covered during the design phase.

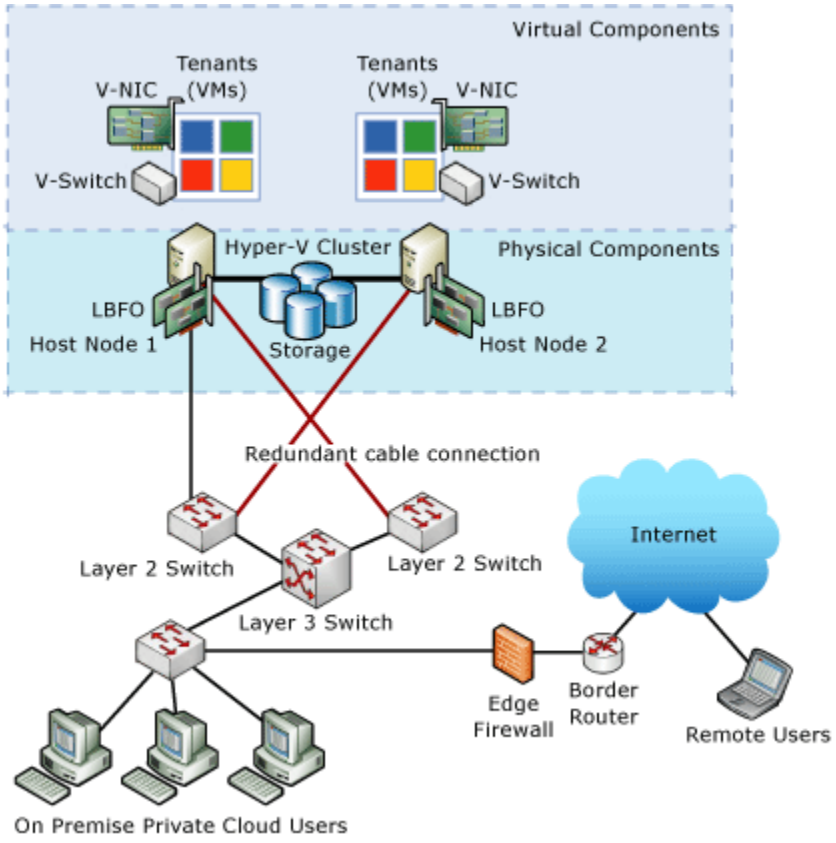

Figure 2. Network infrastructure overview

This section will cover the networking design considerations to assist you with building your cloud infrastructure.

Network Infrastructure

You must access your current network infrastructure and verify the network is capable of receiving the new set of servers that will be hosting your cloud infrastructure. As shown in Figure 2, there are virtual components and physical components. Because the physical components will be in place before the initial setup of the cloud infrastructure, it is important to review the following elements of your current network infrastructure:

Requirements	Description
Switches where the Hyper-V clusters will be connected	• Layer 2 or layer 3 switches that support major Windows Server 2012 networking features. • Ethernet 1 GB or 10 GB, Fibre Channel, or InfiBand.
Cabling system	• Use the cabling system that matches with your layer 1 technology (Ethernet, Fibre Channel, or InfiBand).
Router	• Border router to route traffic from the private cloud to the Internet.
Edge Firewall/Gateway	• A firewall/gateway that allows secure web access from private cloud to Internet and from remote users to the private cloud.

Note

For more information about InfiniBand technology see InfiniBand® Roadmap.

When designing network connectivity for a well-managed cloud infrastructure, the virtualization hosts should have the following specific networking requirements:

Requirements	Description
Support for 802.1q VLAN tagging	To provide network segmentation for the virtualization hosts, supporting management infrastructure, and workloads. This is one method you can use to help secure and isolate data traffic for a private cloud.
Remote out-of-band management capability	To monitor and manage servers remotely over

	the network regardless of whether the server is turned on or off.
Support for PXE version 2 or later	To facilitate automated physical server provisioning.
Support for ECN switches (RFC 3168 compliant)	To support Windows Server 2012 Datacenter TCP (DCTPC).
Support for 802.1az and 802.1Qbb (priority flow control)	To support Datacenter Bridging (DCB).
Network interfaces and virtual networking gateway that support GRE offload (NVGRE)	To support Windows Server 2012 Network Virtualization.

Before deploying a cloud infrastructure, you should also address the following requirements:

1. Address any network bottleneck prior to building the cloud infrastructure.

   Ensure that redundant components are working. For example, if a switch has two (or more) power supplies, perform a failover test to validate if it is properly working.
2. Make note of the results.
3. Document (or update the documentation) for the current network infrastructure.
4. Additionally, a cloud infrastructure should meet the following requirements to make sure that it is highly available and well-managed:
   - Multiple paths to the switches for redundancy. Should a switch fail this can provide resiliency in the provisioned switch. Consult the switch vendor for specific recommendations.
   - Redundant power supplies and cooling fans, to increase the number of faults the switches and routers can survive.

Traffic Flow Isolation

In a cloud infrastructure it is important that the traffic generated by each tenant is securely isolated from other tenants. The existing Windows Server 2008 Hyper-V Switch architecture does not allow a component outside the switch to add to the security and isolation that is already provided by the switch. The new Windows Server 2012 Hyper-V Extensible Switch has support for isolation policies as well as allowing for extensibility and lets third parties add in filters to provide their own forwarding rules.

Note

For more information about Hyper-V Virtual Switch, see Hyper-V Virtual Switch Overview

Another important feature that Windows Server 2012 Hyper-V introduces is called Network Virtualization. This feature removes the constraints of VLAN and hierarchical IP address assignment for virtual machine provisioning. In addition, network virtualization maintains the necessary multi-tenant isolation and security requirements. This feature introduces the concept of location-independent addressing by virtualizing the IP address. It creates virtual Layer-2/Layer-3

topologies over any physical network that supports bidirectional IP connectivity. This means that physical networks can be hierarchical three-tier networks, with full directional bandwidth, a Clos network, or a large Layer-2 infrastructure. The advantage is that virtual networks can dynamically be configured by software as virtual machines are added to the cloud, and it can span over multiple physical subnets and across multiple sites.

Windows Server 2012 Network Virtualization can support a larger number of network isolation groups than the 4,094 maximum provided by VLANs (per IEEE 802.1Q standard). With VLAN-based network isolation, all of the workloads in an isolation group must be on a single subnet. Network Virtualization also allows network isolation groups to be defined across network boundaries such as subnets. This improves flexibility by removing physical network configuration constraints in determining where workloads can be placed.

Note

For more information about Hyper-V Network Virtualization see Hyper-V Network Virtualization Overview.

Enterprises that already have a large Fibre Channel SAN installation for storage service but want to migrate away from additional investment in the Fibre Channel technology will be able to use data center bridging (DCB) to build an Ethernet-based converged fabric for both storage and data networking. Having a unified Ethernet fabric can reduce the future TCO and simplify the management. For enterprises who have already adopted, or who plan to adopt iSCSI as their storage solution, DCB can also provide hardware-assisted bandwidth reservation for iSCSI traffic to ensure performance isolation. DCB-capable Ethernet network adapters must be installed in computers that are providing Windows Server 2012 DCB, and DCB-capable hardware switches must be deployed on the network that will host the cloud infrastructure.

The following table summarizes the design requirements for traffic isolation and the Windows Server 2012 feature that should be used to accomplish this:

Design requirements	Windows Server 2012 feature	How does it help?
Separate workloads from internal IP addresses per tenant	Hyper-V Network Virtualization	You can separate workloads from internal IP addresses in order to minimize configuration changes that are necessary for IP addresses, Domain Name System (DNS), and other virtual machine configurations.
Separate server and network administrator duties to allow operational isolation		You can simplify management since migration and placement of workloads are independent from the underlying physical network configuration.
Provide isolation beyond		Using software based policy in

VLAN configuration		conjunction with policy-based data center networks makes it possible to implement isolation without using VLANs.
Migrate servers across data centers in a flexible way while traffic is isolated from other tenants		You can use flexible workload placement while keeping the current IP address schema as it is not limited to a physical IP subnet or VLAN configuration.
Optimize server and network utilization to not affect other tenants		By removing the dependency of virtual machine placement based on physical network infrastructure you will increase the flexibility for workload placement and improve the overall server and network utilization.
Automate management and deployment		By using the Network Virtualization PowerShell cmdlet it is possible to accomplish this design requirement.
Allow lower-level traffic inspection with rules that can be applied per tenant	Hyper-V Switch Extensibility	Filtering extensions can inspect, drop, modify, and insert packets.
Allow different tenants to implement different filters according to their traffic-isolation needs		Extensions plug in to each instance of a vSwitch on a machine, and those extensions are either filters or Windows Filtering Platform (WFP) providers.
Monitor tenant traffic without modifying the existing vSwitch traffic		Capture extensions can inspect traffic and generate new traffic for reporting purposes. Capture extensions do not modify existing vSwitch traffic.

Another mechanism that can be used to assist traffic isolation between virtual machines in Windows Server 2012 is new support for Hyper-V switch port ACLs. This feature will be covered in more detail in the Security Considerations next section of this document.

Security Considerations

Considering what was mentioned previously in the section Cloud Architecture Considerations, security is considered to be a wrapper around the entirety of the solution. When designing network connectivity for a secure cloud infrastructure, there are a series of security considerations for the current physical network infrastructure that must be verified before deploying the cloud infrastructure.

The first recommended step is to perform a network vulnerability assessment (NVA) in the current infrastructure. Before securing the assets you must evaluate what's in place via an analytical assessment of every control point on the network (such as switches and routers). When performing an NVA it is important to follow a core framework that must include at least the following items:

Task	Example
Detect and document the vulnerabilities	• Description: Switches located in the data center do not have the latest firmware and are vulnerable to Internet Control Message Protocol (ICMP) Redirect Attack.
Assign risk levels to discovered vulnerabilities	• Criticality: Severe • Reason: If an attack successfully exploits the vulnerabilities on the switches located in the data center, they could redirect the traffic and disrupt traffic flow, causing major downtime for the hosts connected to that switch.
Identify vulnerabilities that have not been remediated and take actions to remediate	• Description: Two months ago the NVA Report 00001 detected that switches located in the data center were vulnerable to ICMP Redirect Attack. The problem was not fixed. • Action: Escalate the problem to upper management to address this issue immediately.

It is important to perform the NVA on the current network infrastructure because an attacker usually looks for poorly configured network devices to exploit. Besides the vulnerabilities that are resolved by applying updates the come from the device vendor, there are other vulnerabilities that need to be addressed, such as those caused by insecure administration practices, such as weak default installation settings (including weak password) and wide-open access controls (no ACLs). The following include some of the main network threats:

- Sniffing
- Spoofing

- Session hijacking
- Denial of service

Cloud operators and private cloud network administrators must have knowledge of the threats that can affect the network infrastructure in order to apply effective countermeasures.

Note

For more information about Network Security, see Advanced Network Security.

Once the current physical network infrastructure is secured, the next phase is to plan how to extend this security to the virtual network while allowing the traffic between physical and virtual to flow in a secure manner. Security considerations while designing and planning the cloud infrastructure should also include:

- Secure access control between virtual machines (on the same host or in different hosts)
- Traffic control to avoid tenants overloading the network
- Avoiding rogue IP distribution
- Ensure network availability and performance

Secure Access Control

Windows Server 2012 supports Hyper-V port ACLs that are enforced by the Hyper-V virtual switch. These ACLs are fine-grained rules that can either allow or deny traffic destined to or from a virtual machine's virtual network adapter. ACL filters are based on the IP address, IP address prefix, or the MAC address of the incoming or outgoing packets.

Since ACLs are a port property of the Hyper-V virtual switch, when a virtual machine live migrates to another host, the ACLs move with the virtual machine. Although it is technically possible to provide multi-tenancy isolation by using only ACLs, the challenge is managing and keeping all ACLs updated. For this reason, ACLs are intended to ensure that virtual machines do not spoof their IP or MAC addresses or to control specific network traffic for particular address ranges.

Consider developing an access control plan that describes a method to establish a secure and usable environment. A typical access control plan might include the following sections:

- Security goals: define the resources and processes that you are protecting.
- Security risks: enumerate the types of security hazards that affect your enterprise, including what poses the threats and how significant these threats are.
- Security strategies: describe the general security strategies necessary to meet the threats and mitigate the risks.
- Security policy: policy to define and enforce your security strategy on any number of virtual machines.
- Information security strategies: define how you plan to implement information security solutions.
- Administrative policies: document policies for delegation of administrative tasks and monitoring of audit logs to detect suspicious activity.

Note

For more information about Access Control List on Hyper-V see Hyper-V Virtual Switch Overview.

Traffic Control

As documented in the Cloud Security Challenges Secondary to Cloud Essential Characteristics document, one of the concerns as a designer of a private cloud solution is that: "A rogue application, client, or DoS attack might destabilize the data center by requesting a large amount of resources. How do I balance the requirement that individual consumers/tenants have the perception of infinite capacity with the reality of limited shared resources?"

One way to address this concern is by controlling the network traffic. QoS in Windows Server 2012 is designed to help manage network traffic on the physical network and on the virtual network. Policy-based QoS is designed to manage network on the physical network. In addition, a new functionality in QoS, referred to as Hyper-V QoS, is designed to manage traffic at the virtual switch level.

The use of policy-based QoS allows cloud operators to specify network bandwidth control based on application type, users, and computers. Policy-based QoS can also be used to help control bandwidth costs and negotiate service levels with bandwidth providers or departments (tenants). Hyper-V QoS enables cloud operators to guarantee specific performance levels based on service level agreements (SLAs). Hyper-V QoS helps ensure that no tenant is impacted or compromised by other tenants on their shared infrastructure, which includes computing, storage, and network resources.

Consider developing a QoS plan that describes how to establish a secure a usable environment. A typical QoS plan might include the following sections:

- SLA: plan QoS policy based on the tenants' SLA.
- Network utilization: gauge network utilization to understand traffic flow and how QoS can be used to optimize performance.
- Policy enforcement: enforce QoS policies on Single Root I/O Virtualization (SR-IOV) network adapters that support bandwidth reservation per virtual port.

These sections should be considered per tenant because it is possible to create multiple virtual network adapters in Hyper-V and specify QoS on each virtual network adapter individually.

Note

For more information about Policy Based QoS see Quality of Service (QoS) Overview.

Avoid Rogue IP Distribution

The same core concern that a rogue application, client, or denial of service (DoS) attack might destabilize the data center, mentioned previously, also applies in this case. If a rogue Dynamic Host Configuration Protocol (DHCP) server is capable of providing IP addresses to computers in the tenant network, these computers might lose access to resources and disrupt the whole tenant network infrastructure.

The DHCP server role in Windows Server 2012 introduces a new policy-based IP address and option assignment feature. With the Policy Based Assignment (PBA) feature the cloud operator will be able to group DHCP clients by specific attributes based on fields contained in the DHCP client request packet. This feature allows a greater control of configuration parameters delivered to network devices.

A typical DHCP PBA plan might include the following sections:

- Define conditions: DHCP PBA can be defined according to fields in the DHCP client request. Define which condition better fits on your cloud infrastructure network.
- IP scope: define the DHCP IP scope that will be used for the tenants.
- Lease duration: define the DHCP lease duration for each tenant network.
- Additional options: define the DHCP scope or server options that will be provided to the tenants. If policies exist at the server and scope levels, the server applies both sets of policies and evaluates the scope policies before the server policies. The processing order for a scope-level policy defines the order of evaluation within the scope.

Note

For more information about DHCP PBA see Step-by-Step: Configure DHCP Using Policy-based Assignment.

Scalability and Performance Considerations

Many network architectures include a tiered design with three or more tiers such as core, distribution, and access. When designing a cloud infrastructure the current network infrastructure should be evaluated to understand the port bandwidth and quantity required at all layers, in addition to the ability of the distribution and core tiers to provide higher-speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, and spanning tree and/or other loop avoidance technologies should be considered.

Physical separation of the networks should also be considered in order to provide another layer of isolation and security and increase the overall performance per segment. When designing a cloud infrastructure ensure that the servers that will be used to host the virtual machines support different traffic profiles by either using separate physical adapters for each traffic type, or by using VLAN tagging or virtual NICs. The table below describes a cloud infrastructure model with five physical network segments:

Traffic	Description
Management	Network that handles all the management traffic. This traffic will be from the host to the management workstation. Cloud operators will connect to the host system to manage the cloud infrastructure.
Cluster	Network dedicated to cluster communications.

Storage	Handle all the storage traffic in scenarios where you use Ethernet-based storage protocols.
Live migration	Responsible for handling the live migration traffic that will take place from one host to another.
Tenant	Network dedicated to tenant traffic. An additional layer of isolation within the tenant network can be provided by using VLAN IDs in the virtual switch.

Note

It is also possible to route all these traffic profiles through a single or teamed 10 GB network adapter. This converged networking model reduces the number of network adapters and cabling required for your private cloud infrastructure.

If business requirements dictate that your cloud infrastructure should have five physical network segments, the same layout can be applied to the virtual network infrastructure. When implementing those requirements it is also important to create a standard nomenclature for the network adapters names in order to reflect which network they are connected to. This not only helps to easily identify the network but also to create a standard for future automations via Windows PowerShell.

Although physical separation of the traffic can be a business requirement, there will be scenarios that this is not the case and the main business requirement is to reduce implementation cost while maintaining security isolation between tenants. When the business requirements lead you to this design decision you can converge all the data center networks into basically two networks, physically isolated: a data center network can carry all the storage, live migration, clustering, and management traffic flows, and a second network can carry all of the virtual machine tenant-generated traffic. You can still apply QoS policies to guarantee minimum traffic for each flow and use VLAN ID in the virtual switch to isolate the traffic within the tenant network.

When choosing the networking hardware for your cloud infrastructure, consider the following options:

Network technology	Advantages	Disadvantages
10 GB Ethernet	• Great performance • Offers QoS (DCB) and flexible bandwidth allocation • New hardware offloads available	• RDMA optional (for SMB 3.0 file access) • Physical switch ports more expensive
InfiniBand (32 GB and 56	• Very high performance,	• Network management

GB)	low latency • RDMA included (for SMB 3.0 file access)	different than Ethernet • More expensive to implement
1 GB Ethernet	• Adequate performance for many workloads	• Not very scalable

Remote interfaces and management of the network infrastructure via Secure Shell (SSH) or similar protocol is important to both automation and resiliency of the data center network. Remote access and administration protocols can be used by management systems to automate complex or error prone configuration activities. For example, adding a VLAN to a distributed set of access-tier switches can be automated to avoid the potential for human error.

Host and Virtualization Design

There are a number of critical design decisions you need to make that apply to your virtualization infrastructure. Hardware virtualization uses software to create a virtual machine that emulates a physical computer. This creates a separate operating system environment that is logically isolated from the host server. By providing multiple virtual machines at once, this approach allows several operating systems to run simultaneously on a single physical machine.

Windows Server 2012 Hyper-V technology is based on a 64-bit hypervisor-based microkernel architecture that enables standard services and resources to create, manage, and execute virtual machines with the cloud infrastructure. The Windows Hypervisor runs directly above the hardware and ensures strong isolation between the partitions by enforcing access policies for critical system resources such as memory and processors. The Windows Hypervisor does not contain any third-party device drivers or code, which minimizes its attack surface and provides a more secure architecture.

In addition to the Windows Hypervisor, there are two other major elements to consider with Hyper-V: a parent partition and child partition. The parent partition is a special virtual machine that runs Windows Server 2012, controls the creation and management of child partitions, and maintains direct access to hardware resources. In this model, device drivers for physical devices are installed in the parent partition. In contrast, the role of a child partition is to provide a virtual machine environment for the installation and execution of guest operating systems and applications.

Windows Server 2012 Hyper-V Host Design

There are many important considerations you need to make when designing the Windows Server 2012 host and Hyper-V components as part of your cloud infrastructure. In this section we will discuss decision points regarding:

- Licensing
- Operating system configuration
- Memory and Hyper-V Dynamic Memory

- Storage adapters
- Hyper-V host failover design
- Hyper-V guest virtual machine design

Licensing

Different versions of Windows Server 2012 include different virtualization rights, which is the right and license to run a specified number of virtual machines based on Windows. For a private cloud environment, you will want to use Windows Server 2012 Datacenter Edition.

Operating System Configuration

The following outlines the general considerations for the Hyper-V Host Operating system. Note that these are not meant to be installation instructions but rather the process requirements and order.

To install and use the Hyper-V role, you will need the following:

- An x64-based processor
- Hardware-assisted virtualization support in the processor and BIOS
- Hardware-enforced data execution prevention (DEP) support in the processor and BIOS
- Use the latest hardware device drivers and system BIOS updates
- Hyper-V parent partition operating system must be domain-joined. This is required to support failover clustering and other management capabilities.
- Hyper-V server roles and failover clustering features. Failover clustering will provide high availability so that virtual machines remain available even if a member of the cluster is disabled.
- Apply relevant Windows updates, including out-of-band (OOB) updates not offered on Microsoft Update
- All nodes, networks, and storage must pass the Cluster Validation Wizard

Please see Hyper-V Overview for more information on hardware and software requirements.

Memory and Hyper-V Dynamic Memory Options

Dynamic Memory, introduced in Windows Server 2008 R2 Service Pack 1 (SP1), defined startup memory as the minimum amount of memory that a virtual machine can have. However, Windows requires more memory during startup than the steady state. As a result, administrators sometimes assign extra memory to a virtual machine because Hyper-V cannot reclaim memory from these virtual machines after startup. In Windows Server 2012, Dynamic Memory introduces a minimum memory setting, which allows Hyper-V to reclaim the unused memory from the virtual machines. This is reflected as increased virtual machine consolidation numbers.

Windows Server 2012 also introduces Smart Paging for reliable virtual machine restart operations. Although minimum memory increases virtual machine consolidation numbers, it also brings a challenge. If a virtual machine has a smaller amount of memory than its startup memory and it is restarted,

Hyper-V needs additional memory to restart the virtual machine. Due to host memory pressure or virtual machine states, Hyper-V might not always have additional memory available. This can cause sporadic virtual machine restart failures. Smart Paging is used to bridge the memory gap between minimum memory and startup memory, and allow virtual machines to restart reliably.

For more information regarding Dynamic Memory, please see Hyper-V Dynamic Memory Overview.

In addition to the general guidance above, specific applications or workloads, particularly those with built-in memory management capabilities, such as SQL Server or Exchange Server, might provide workload-specific guidance. An example of such guidance is Running SQL Server with Hyper-V Dynamic Memory.

When designing for memory requirements for host servers in the Hyper-V cluster, you should consider how many virtual machines you want to support on each server. Windows Server 2012 supports up to 4 TB of memory for the host operating system and up to 1 TB for each virtual machine.

One approach you can take is to consider how many virtual machines you wish to support of each service class you plan to offer in your private cloud. For example, you may want to support up to 10 virtual machines on each host system in the cluster. You also want to offer the following service classes in terms of amount of memory offered:

- **Bronze** – 1 GB of RAM for the virtual machine
- **Silver** – 4 GB of RAM for the virtual machine
- **Gold** – 16 GB of RAM for the virtual machine

You estimate that consumers of your private cloud will acquire these service offerings in the following percentages:

- **Bronze** – 30%
- **Silver** – 60%
- **Gold** – 10%

You can compute the amount of RAM required in each host in this example using the following:

(1GBx3) + (4GBx6) + (16GBx1) + 2GB (for the host OS) = 45 GB

Storage Adapters

Design decisions regarding storage adapters on the Hyper-V private cloud host are important. The type of storage adapter you choose will determine the speed, latency, reliability, performance, and cost of your cloud infrastructure storage components. In this section we will discuss:

- MPIO configuration
- Performance settings
- Network adapter teaming configurations

MPIO Configuration

As discussed earlier, Microsoft MPIO architecture supports iSCSI, Fibre Channel, and SAS SAN connectivity by establishing multiple sessions or connections to the storage array.

Performance Options

The following Hyper-V network performance improvements should be tested and considered for production use (note that many of these technologies require hardware support):

- **Transmission Control Protocol (TCP) checksum offload** benefits both CPU and overall network throughput performance, and it is fully supported by live migration.
- **Jumbo frames** capability is extended to virtual machines with Windows Server 2012 Hyper-V. Just as in physical network scenarios, jumbo frames add the same basic performance enhancements to virtual networking. That includes up to six times larger payloads per packet, which improves overall throughput and also reduces CPU utilization for large file transfers. For more information on jumbo frames, please see Hyper-V Networking Options – Jumbo Frames.
- **Dynamic Virtual Machine Queue (DVMQ)** enables you to span processing of virtual machine network traffic across all processors in the host operating system. This feature enables you to scale up and scale down the CPU utilization based on demand.
- **Scale and resiliency.** The computer that runs Hyper-V can be configured with up to 320 logical processors and 4 TB of memory. Virtual machines can be configured with 32 virtual processors and 1 TB of memory. Improved handling of hardware errors increases the resiliency and stability of the virtualization environment. For more information on scale and resiliency features, please see Hyper-V Support for Scaling Up and Scaling Out Overview.
- **Single root I/O virtualization (SR-IOV).** Use of SR-IOV maximizes network throughput while minimizing network latency, in addition to the CPU overhead required for processing network traffic. SR-IOV bypasses the virtual switch stack and enables the virtual machine direct access to the hardware. Network adapter and BIOS support is required. For more information on SR-IOV please see Hyper-V Support for Scaling Up and Scaling Out Technical Preview.
- **New .VDHX disk format.** The new format increases the maximum storage size per virtual hard disk, and improves the stability and efficiency of those disks. The new VDHX disk format supports up to 64 TB of storage and supports newer 4K disk architecture. It also provides built-in protection from corruption stemming from power failures and prevents performance degradation on some large-sector physical disks. For more information on the new VHDX disk format, please see Hyper-V Virtual Hard Disk Format Technical Preview.
- **Virtual non-uniform memory access (NUMA).** Newer operating systems and high-performance applications such as SQL Server include optimizations that recognize a computer's NUMA topology to increase performance by considering NUMA when scheduling threads or allocating memory. The virtual NUMA feature makes it possible for the guest operating system and NUMA-aware applications running in the virtual machine (such as SQL Server) to take advantage of these performance optimizations. A virtual NUMA topology is projected to the guest operating system. The default virtual NUMA topology is optimized to match the NUMA topology of the physical computer. Note that if a virtual machine uses dynamic memory, then a flat NUMA topology is reflected to that VM. In addition, in Windows Server 2012 failover clustering, NUMA topology of members of the cluster is taken into consideration when automated decisions are made to move virtual machines.

- **Hyper-V Offloaded Data Transfer.** To take advantage of innovations in storage hardware that provide near-instantaneous copying of large amounts of data, Hyper-V in Windows Server 2012 introduces Offloaded Data Transfer (ODX). With this new feature, Hyper-V workloads use the offload semantics of the host hardware, in addition to the virtual storage stack, to perform certain internal operations on virtual hard disks that require large amounts of data to be copied. Hyper-V performs these operations faster than was previously possible. For more information on Windows Server 2012 ODX, please see Hyper-V Offloaded Data Transfer Overview.

Network Adapter Teaming Configurations

Network adapter teaming can be utilized to enable multiple redundant network adapters and connections between servers and access-tier network switches. Teaming can be enabled via hardware or software-based approaches. Windows Server 2012 includes built-in support for network adapter teaming via the new load balancing and failover (LBFO) feature, where any two network adapters can be teamed, regardless of make, model, or speed. Teaming can enable multiple scenarios including path redundancy, failover, and load-balancing.

Windows Server 2012 LBFO (Load Balancing and Failover, also known as NIC Teaming), allows multiple network adapters on a computer to be placed into a team for the following purposes:

- Bandwidth aggregation
- Traffic failover to prevent connectivity loss in the event of a network component failure

Note

Windows Server 2012 NIC Teaming enables you to configure the team in ways that can have all adapters active or reserve some adapters for standby. We recommend that you use the all active configuration for your private cloud NIC teams.

This feature has been a requirement for independent hardware vendors (IHVs) to enter the server network adapter market, but until now network adapter teaming has not been included in Windows Server operating systems. NIC Teaming enables you to team network adapters of different speeds and from different manufacturers.

Note

It is recommended that you do not team network adapters of different speeds. The reason for this is that while it is a supported configuration, if the higher speed adapter fails, performance on the slower adapter will be severely impacted.

It is recommended that you do not team network adapters of different speeds. The reason for this is that while it is a supported configuration, if the higher speed adapter fails, performance on the slower adapter will be severely impacted.

Network adapter teaming in Windows Server 2012 also works within a virtual machine. This allows a virtual machine to have virtual network adapters that are connected to more than one Hyper-V switch and still have connectivity even if the network adapter under that switch gets disconnected. This is particularly important when working with features such as SR-IOV because SR-IOV traffic does not go through the Hyper-V switch. Thus, it cannot be protected by a team that is under a Hyper-V switch. With the virtual machine teaming option, an administrator can set

up two Hyper-V switches, each connected to its own SR-IOV-capable network adapter. At that point:

- Each virtual machine can then install a virtual function from one or both SR-IOV network adapters. Then, in the event of a network adapter disconnect, the virtual machine can fail over from the primary virtual function to the backup virtual function.
- Alternately, the virtual machine might have a virtual function from one network adapter and a non-virtual function network adapter from the other switch. If the network adapter that is associated with the virtual function gets disconnected, the traffic can fail over to the other switch without loss of connectivity.

Note
NIC Teaming is compatible with all networking capabilities in Windows Server 2012 with three exceptions:

- SR-IOV
- RDMA
- TCP Chimney (TCP Chimney is disabled by default in Windows Server 2012)

For SR-IOV and RDMA, data is delivered directly to the network adapter without passing through the virtual networking stack. Therefore, it is not possible for the network adapter team to look at or redirect the data to another path in the team. In addition, TCP Chimney is not supported with network adapter teaming in Windows Server 2012. Network adapter teaming requires the presence of at least one Ethernet network adapter, which can be used for separation of traffic using VLANs. All modes that provide fault protection through failover require at least two Ethernet network adapters. The Windows Server 2012 implementation supports up to 32 network adapters in a team.

For more information on Windows Server 2012 load balancing and failover, please see Load Balancing and Failover Overview.

Dedicated Compute and Storage Clusters Using SMB 3.0 and Windows Server 2012 File Server

In this design pattern, the Hyper-V cluster providing the compute component for the cloud infrastructure is separate from the storage cluster. When you separate the compute from the storage cluster, you have the opportunity to scale compute capacity and storage capacity separately. This provides you more flexibility when designing your scale units for compute and storage. In this pattern, two 10 GbE adapters on the compute cluster are teamed which support management, cluster, live migration, and storage traffic. The Windows Server 2012 storage cluster is also configured with teamed 10 GbE adapters and hosts VHDX and configuration files that are accessed through SMB 3.0 file shares over the network. The storage configuration on the storage cluster can connect to any type of block storage solution.

Hyper-V Host Failover Cluster Design

A Hyper-V host failover cluster is a group of independent servers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to

provide service (a process known as failover). In case of a planned migration (called live migration), users experience no perceptible service interruption.

The host servers are one of the critical components of a dynamic, virtual infrastructure. Consolidation of multiple workloads onto the host servers requires that those servers be highly available. Windows Server 2012 provides advances in failover clustering that enable high availability and live migration of virtual machines between physical nodes. Some of the new features which you will need to include in your private cloud design decisions for host cluster members include:

- Cluster scalability features
- CSVs
- Support for scale-out file servers
- Cluster-aware updating
- Virtual machine application monitoring
- Cluster validation tests
- Active Directory Domain Services integration
- Multisite support
- Cluster upgrade and migration
- iSCSI software target integration
- Windows PowerShell support

For more information on each of these features and how they can fit into your cloud infrastructure design, please see What's New in Failover Clustering.

Private Cloud Infrastructure without Failover Clustering

It is important to note that a private cloud infrastructure does not require failover clustering. Failover clustering provides high availability for stateful applications that are not specifically designed to work on and support cloud capabilities. Specific cloud capabilities are targeted at the stateless applications of the future.

Prior to Windows Server 2012, Live Migration of virtual machines from one host to another required failover clustering; thus a failover cluster defined the scope of virtual machine and workload mobility. However, Windows Server 2012 introduces what is known as the "shared nothing" Live Migration feature.

Shared nothing Live Migration enables the cloud service provider to move virtual machines from one Hyper-V host to another without requiring failover clustering or shared storage. With only a network cable (or even wireless connection) the virtual machine and its virtual disk and configuration files can be moved from one machine to another. Private cloud service providers could take advantage of this capability to run virtual machines hosting stateless workloads and place virtual machines on specific hosts based on their fabric controller of choice.

This design guide focuses on implementing failover clustering in the private cloud infrastructure because the Windows Server 2012 failover clustering feature adds many vital capabilities for the management, monitoring and control of the overall cloud solution and is tightly integrated with Windows Server 2012 Hyper-V technologies. In addition, we believe that, for at least the near

term that stateful applications will represent the most common workload running in a private cloud.

Host Failover Cluster Topology

In a Microsoft Private Cloud infrastructure, we recommend two standard design patterns. The server topology should consist of at least two Hyper-V host clusters. The first should have at least two nodes, and will be referred to as the management cluster. The second, and any additional clusters, will be referred to as fabric host clusters. The second cluster might be the compute cluster, and the third cluster might be the storage cluster. Alternatively, the second cluster can be a combined computer and storage cluster that is either symmetric or asymmetric.

Note

In the current context, a symmetric cluster is a cluster where each member of the cluster is directly attached to storage. Conversely, an asymmetric cluster has some members attached to storage and others unattached to storage. Typically, the unattached cluster members are performing the compute function (actually running the virtual machines) and the attached members of the cluster are acting as scale-out file servers that host the virtual machine files for the compute nodes.

In some cases such as smaller-scale scenarios or specialized solutions, the management and fabric clusters can be consolidated onto the fabric host cluster. Special care has to be taken in that case to ensure resource availability for the virtual machines that host the various parts of the management stack. There are significant security issues with this design, so assiduous attention to traditional and cloud-based security measures is mandatory.

Each host cluster can contain up to 64 nodes. However, for file server clusters running the Windows Server 2012 scale-out file server role that host VHDX files used by the computer cluster, there is an informal 8-node limitation. The compute cluster can host its own storage, or can use SMB 3.0 file based storage to access storage on the file server cluster.

Note

The 8-node informal limitation is based on what has been tested as of the writing of this document. You can exceed the 8-node limitation and customer support services will work with you in identifying and resolving issues should they arise.

Compute Cluster Traffic Profiles

A variety of host cluster traffic profiles or types are used in a Hyper-V failover cluster. The network requirements to support these traffic types enable high availability and high performance.

The Microsoft cloud infrastructure configurations support the following Ethernet traffic profiles:

- **Management network traffic.** The management network is required so that hosts can be managed to avoid competition with guest and other infrastructure traffic needs. The management network provides a degree of separation for security and ease of management purposes. This network is used for remote administration of the host, communication to management systems (System Center agents), and other administrative tasks.

- **iSCSI, Fibre Channel over Ethernet, or SMB 3.0 traffic (Storage Traffic)**. If using iSCSI, FCoE, or SMB 3.0, the iSCSI, FCoE, or SMB 3.0 network is configured so that storage traffic is not in contention with any other infrastructure or tenant traffic. For all these storage connections, an MPIO configuration with two independent physical ports is required. In the case of SMB 3.0 access to file storage, multiple network adapters can be used in conjunction with SMB multichannel and transparent failover to provide benefits similar to MPIO.
- **CSV/cluster communication traffic.** Usually, when the cluster coordinator node that "owns" a virtual hard disk (VHD) file in CSV performs disk I/O, the node communicates directly with the storage devices; for example, through a SAN. However, storage connectivity failures and other non-failure scenarios sometimes prevent a given node from communicating directly with the storage device. To maintain functionality until the failure is corrected, the node redirects the disk I/O through a cluster network (the preferred network for CSV) to the node where the disk is currently mounted. This is called CSV redirected I/O mode.
- **Live migration traffic.** During live migration, the contents of the memory of the virtual machine running on the source node needs to be transferred to the destination node over a LAN connection. Large virtual machines can consume many gigabytes of memory that need to be transferred over the network. To provide a high-speed transfer, a dedicated, redundant, 10 Gbps live migration network is required. This significantly reduces the time required to evacuate the virtual machines off a host with zero downtime during maintenance or Windows updates. The time it takes to complete an evacuation of a cluster member depends on the total amount of memory consumed by running virtual machines on that system and the amount of bandwidth available on the Live Migration network.
- **Tenant traffic.** Tenant traffic is dedicated to virtual machine LAN traffic. The tenant traffic can use two or more 1 GB or 10 GB networks using network adapter teaming, or virtual networks created from shared network adapters. You can implement one or more dedicated virtual machine networks. The amount of bandwidth required on the tenant network may be less than what you require on any of the infrastructure networks, depending on the type of workloads you expect to support on your private cloud. One way to determine the bandwidth you want to make available on your tenant network is to define your network service classes and then determine what is the total amount of bandwidth required for you to meet the SLA for your network bandwidth defined service classes for all VMs running on your host.

When you look at each of the traffic profiles presented here, you can categorize each of them into one of two categories:

- Private cloud infrastructure traffic
- Private cloud tenant traffic

Private cloud infrastructure traffic includes management, storage, CSV/cluster and Live Migration traffic. Tenant traffic is defined by traffic to and from virtual machines running within the private cloud infrastructure.

Note

There are some security considerations that you need to address which extend past the isolation methods we've talked about until now. While you can use Hyper-V virtual switch port ACLs and VLAN tagging and even private VLANs to obtain a good degree of separation between host and guest traffic, you may want to introduce additional security

for your Live Migration traffic. During the process of a Live Migration, tenant data, by default, is moving unencrypted over the Live Migration network. There is a reasonable chance that this data contains private information that might be of interest to intruders. If the physical or logical infrastructure of the Live Migration network were to be compromised, the Live Migration traffic could be accessible in an unencrypted format to the intruder. Because of this, we recommend that you use IPsec to secure the connections between hosts on the Live Migration network.

There are several options available regarding how you manage these traffic profiles so that they receive the bandwidth they require while enforcing separation between the tenant traffic and the cloud infrastructure traffic (management, storage, CSV/cluster, and live migration). These are:

- **Dedicated network adapters for each traffic profile.** This is a traditional approach that was recommended for Windows Server 2008 R2 Hyper-V clusters. With new features included in Windows Server 2012, this is no longer considered the preferred design. However, it simplifies upgrading the current Windows Server 2008 R2 to a Windows Server 2012 infrastructure.
- **Dedicated network adapters for cloud infrastructure traffic and tenant traffic.** In this design pattern, separate adapters and networks are used for the cloud infrastructure traffic and the tenant traffic. This provides the required isolation between infrastructure and tenant traffic, and lessens the impact of tenant traffic on overall bandwidth availability. Windows Server 2012 QoS policies can be used to provide minimum and maximum bandwidth availability for each cloud infrastructure traffic profile.
- **No dedicated adapters for any traffic profile.** In this design pattern, all traffic moves through the same network adapters over the same physical network. Infrastructure traffic and tenant traffic share the same physical adapter or adapter team. Hyper-V QoS policies are applied so that each traffic profile has a guaranteed amount of bandwidth. VLAN tagging or Hyper-V virtual switch port ACLs can be used to provide the necessary isolation between cloud infrastructure and tenant traffic. This pattern requires that infrastructure traffic flow through virtual network adapters that are created on the same Hyper-V virtual switch through which tenant traffic also flows. The advantage of this converged networking pattern is that it is simpler to manage, is more cost effective and enables you to take advantage of security and performance capabilities included with the Hyper-V virtual switch.

Note that the conversation to this point has been focused on traffic patterns and network design for compute node cluster traffic. If you have chosen to separate your compute and file server clusters so that you can scale compute and storage separately, then you will need to consider what traffic profiles need to be defined on the file server cluster.

The file server cluster will have the management, cluster/CSV, and perhaps a storage network, depending on how you decide to present block storage to the file server cluster. There will be no Live Migration traffic profile on the storage cluster. But you may want to define a file server network traffic profile, which delineates the path that connects the compute node to the file server node using the SMB 3.0 protocol.

Hyper-V Guest Virtual Machine Design

Standardization is a key tenet of private cloud architectures. This also applies to virtual machines. A standardized collection of virtual machine templates can both drive predictable performance and greatly improve capacity planning capabilities. These templates also provide the foundation for your private cloud's service catalog. The service catalog can also be used to help you determine how you size the hosts in your private cloud as well as the scale units you might wish to define.

As an example, the table below illustrates what a basic virtual machine template library might look like.

Template	Specs	Network	Operating System	Unit Cost
Template 1: Small	1 vCPU, 2 GB memory, 50 GB disk, 100Mbps	VLAN 20	Windows Server 2008 R2	1
Template 2: Medium	4 vCPU, 4 GB memory, 250 GB disk, 400Mbps	VLAN 20	Windows Server 2008 R2	2
Template 3: Extra-Large	16 vCPU, 16 GB memory, 500 GB disk, 800Mbps	VLAN 20	Windows Server 2008 R2	4
Template 4: Small	1 vCPU, 2 GB memory, 50 GB disk, 100Mbps	VLAN 10	Windows Server 2012	1
Template 5: Medium	4 vCPU, 4 GB memory, 250 GB disk, 400Mbps	VLAN 10	Windows Server 2012	2
Template 6: Extra-Large	16 vCPU, 16 GB memory, 500 GB disk, 800Mbps	VLAN 10	Windows Server 2012	4

Virtual Machine Storage

This section discusses the different types of Hyper-V disks. Note that while in the past Microsoft recommended using only fixed VHDs for production, significant improvements to the virtual disk format (.VHDX) have made using dynamically expanding disks a viable format for production use. Therefore, for performance reasons, you can use either fixed or dynamically expanding disks in your private cloud infrastructure.

Dynamically Expanding Disks

Dynamically expanding virtual hard disks provide storage capacity as needed to store data. The size of the VHDX file is small when the disk is created and grows as data is added to the disk. In Windows Server 2012 the size of the VHDX file will shrink automatically when data is deleted from the virtual hard disk. Dynamically expanding disks can be provisioned very quickly and can be used as part of your thin provisioning scheme.

Fixed Size Disks

Fixed virtual hard disks provide storage capacity by using a VHDX file that is in the size specified for the virtual hard disk when the disk is created. The size of the VHDX file remains 'fixed' regardless of the amount of data stored, similar to a physical hard disk. However, you can use the Edit Virtual Hard Disk Wizard to increase the size of the virtual hard disk, which increases the size of the VHDX file. By allocating the full capacity at the time of creation, fragmentation at the host level is not an issue (fragmentation inside the VHD itself must be managed within the guest). A disadvantage of fixed sized disks is that they can take a long time to be provisioned, the time depending on the size of the fixed size disk. Fixed disks can provide incremental performance improvement. You should weigh the advantages and disadvantages of using dynamically expanding disks versus fixed sized disks by considering whether the disk space advantages conferred by the dynamically expanding disks outweigh the incremental performance improvement in fixed sized disks.

Differencing Disks

Differencing virtual hard disks provide storage to enable you to make changes to a parent virtual hard disk without altering that disk. The size of the VHD file for a differencing disk grows as changes are stored to the disk. Differencing disks are useful in lab/test environments but can also be used in private cloud infrastructure deployments.

For example, you might want to configure your compute node host infrastructure to use a combination iSCSI and VHDX boot. Given that there is minimal state required on Hyper-V compute cluster nodes, each host can have its own differencing disk. When updates to the operating system are required, the parent disk is serviced and the compute nodes are rebooted after virtual machines are migrated away from the compute node host. A new differencing disk is created at that point. This enables centralized management of a single golden image that can be used for each compute cluster host and significant reduces the amount of storage required to host the cluster node virtual disk files.

Pass-Through Disks

Hyper-V enables virtual machine guests to directly access local disks or SAN LUNs that are attached to the physical server without requiring the volume to be presented to the host server. The virtual machine guest accesses the disk directly (utilizing the disk's GUID) without having to utilize the host's file system. Given that the performance difference between fixed-disk and pass-through disks is negligible in Windows Server 2012, the decision is now based on manageability.

For instance, if the data on the volume will be very large (hundreds of gigabytes or terabytes), a VHDX is hardly portable at that size given the extreme amounts of time it takes to copy. Also, bear in mind the backup scheme. With pass-through disks, the data can only be backed up from within the guest. In addition, virtual machine portability will be limited.

Also note that when utilizing pass-through disks, there is no VHDX file created; the LUN is used directly by the guest. Since there is no VHDX file, there is no dynamic sizing capability or snapshot capability.

Support for Guest Clustering

Creating a guest cluster gives you the ability to failover or migrate your applications independently of the guest operating system and provides you the ability to decouple the application from the guest operating system in a way similar the abstraction that the guest operating system gets from virtualization. In the case of the guest operating system itself, you are decoupling the guest operating system from the hardware. With a guest cluster you decouple the application from the virtual machine. This gives you much greater flexibility and increased uptime in the event that the virtual machine has a failure in which a migration or restart of the VM will not solve.

Windows Server 2012 provides support for guest clustering and access to shared storage required for guest clustering with:

- In-Guest iSCSI Initiator
- In-Guest Fibre Channel synthetic HBA (new in Windows Server 2012)

In-Guest iSCSI Initiator

Hyper-V can also utilize iSCSI storage by directly connecting to iSCSI LUNs utilizing the guest's virtual network adapters. This is mainly used for access to large volumes, volumes on SANs that the Hyper-V host itself is not connected to, or for guest-clustering. Guests cannot boot from iSCSI LUNs accessed through the virtual network adapters without utilizing a third-party iSCSI initiator. However, given that guests can boot from SMB 3.0 file shares, there is no advantage to using iSCSI boot for virtual machines.

In-Guest Fibre Channel Synthetic HBA

A new feature in Windows Server 2012 allows you to connect directly to Fibre Channel storage from within the guest operating system that runs in a virtual machine. This feature makes it possible to virtualize workloads and applications that require direct access to Fibre Channel-based storage. It also makes it possible to configure clustering directly within the guest operating system. This feature makes HBA ports available within the guest operating system.

Hyper-V can utilize Fibre Channel storage by directly connecting to Fibre Channel LUNs utilizing the guest's virtual HBAs. This is mainly used for access to large volumes, volumes on SANs which the Hyper-V host itself is not connected to, or for guest-clustering. Guests cannot boot from Fibre Channel LUNs accessed through the virtual network adapters without utilizing a third-party

Fibre Channel HBA. However, similar to the situation with iSCSI boot, given that guests can boot from SMB 3.0 file shares, there is no advantage to using Fibre Channel boot for virtual machines.

For more information about the in-guest Fibre Channel HBA, see Hyper-V Virtual Fibre Channel Overview.

Virtual Machine Network Interfaces

Hyper-V guests support two types of virtual network adapters: synthetic and emulated. Synthetic makes use of the Hyper-V VMBUS architecture and is the high-performance, native device. Synthetic devices require the Hyper-V integration services be installed within the guest. Emulated adapters are available to all guests even if the integration services are not available. They are much slower performing, and should only be used if synthetic is unavailable.

Note
When you configure a virtual machine to perform a PXE boot for installation, it will initially configure itself to use an emulated adapter. If you wish to take advantage of the synthetic adapter, you will need to change the adapter type after installation completes.

You can create many virtual networks on the server running Hyper-V to provide a variety of communications channels. For example, you can create networks to provide the following:

- Communications between virtual machines only. This type of virtual network is called a private network.
- Communications between the host server and virtual machines. This type of virtual network is called an internal network.
- Communications between a virtual machine and a physical network by creating an association to a physical network adapter on the host server. This type of virtual network is called an external network.

Virtual Processors

In Windows Server® 2008 R2, Hyper-V supported a maximum ratio of 8 virtual processors per 1 logical processor for server workloads, and 12 virtual processors per 1 logical processor for VDI workloads. A logical processor is defined as a processing core seen by the host operating system or parent partition. In the case of Intel Hyper-Threading Technology, each thread is considered a logical processor.

Therefore a 16 logical-processor server supports a maximum of 128 virtual processors. That would in turn equate to 128 single-processor virtual machines, 64 dual-processor virtual machines, or 32 quad-processor virtual machines. The 8:1 or 12:1 virtual to logical processor ratios are maximum supported limits; it is recommended that lower limits be utilized than the maximum.

In Windows Server 2012, there are no hard or soft coded virtual to logical processor ratios. The recommendation is use dual socket servers with the highest core density available. When planning on scale units and how many virtual machines the scale unit will support, it's more effective to consider the amount of RAM you will require, as discussed earlier. If more nodes are

required to provide required processing power, you can add more nodes to your scale unit. Historical trend analysis will be useful in this context. Trends and thresholds are more important than specific ratios.

Overview of Suggested Cloud Infrastructure Deployment Scenarios

Throughout this document you have been presented with hundreds of potential options for how to design a private cloud infrastructure from the storage, networking and compute perspectives. The possible permutations are virtually limitless and you might find it a bit challenging to determine which options would work best for you, and which options work best together. You may even be interested in whether there are tested private cloud infrastructure design patterns that you can use to get your infrastructure started.

To help you in your testing and evaluation of a Microsoft private cloud infrastructure, we have worked on developing and testing three private cloud infrastructure design patterns that you may want to adopt in your own private cloud environment. These design patterns:

- Represent a hardware and software configuration that has been demonstrated to work in the Microsoft Enterprise Engineering Center (EEC)
- Provide three options that you may want to choose from, where many of the design decisions have been made for you
- Help you to better understand how you can take advantage of the many improvements in Windows Server 2012 to get the most out of your private cloud infrastructure investments.

The three design patterns that we've identified and documented are:

- The Non-Converged Data Center Configuration
- The Converged Data Center with File Server Storage Configuration
- The Converged Data Center without Dedicated Storage Configuration

It is important to note that within the context of these three configurations, the term "converged" refers to the networking configuration. A converged networking design consolidates multiple traffic profiles onto a single network adapter (or team) and network and then uses a variety of software constructs to provide the required isolation and Quality of Service for each profile.

Finally, be aware that while these three cloud infrastructure design patterns have made a number of decisions regarding storage, networking and compute functions for you, they do not span the entire gamut of options made available to you through Windows Server 2012. You may want to begin with these patterns and build on top of them by taking advantage of other Windows Server 2012 platform capabilities.

The Non-Converged Data Center Configuration

The Non-Converged Data Center Configuration cloud infrastructure design pattern is aimed at allowing for easy upgrade of an existing cloud infrastructure based on networking design decisions and hardware configuration recommendations for a Windows Server 2008 R2

infrastructure. The Non-Converged Data Center Configuration focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You have an existing investment in separate networks based on the recommended configuration of Hyper-V in Windows Server 2008 R2 and you require that physical network traffic segmentation be kept in place to avoid re-architecting your network. Each type of infrastructure traffic (management, cluster/CSV, Live Migration and storage) and tenant traffic are carried over physically separate networks and network adapters. This requirement is met by installing physically separate NICs for each traffic type and assigning VLAN 802.1q tags to each adapter. Alternatively, you can avoid using VLAN tagging and use port based VLANs on your network switches.
- You require that each traffic type is dedicated to a specific adapter. This requirement is met by configuring each of the traffic flows to use the correct subnet/IP address on the dedicated NIC.
- You require that virtual machine workloads have access to the highest network performance possible. This requirement is met by using the new Windows Server 2012 Single Root I/O Virtualization (SR-IOV) feature, which enables the virtual machines running on the Hyper-V server to have direct access to the network adapter hardware, thus bypassing the virtual networking stack.

Storage

- You have an existing investment in Fibre Channel or iSCSI SANs and require well-connected storage so that all members of the Hyper-V failover cluster are connected to block storage. This requirement is met by configuring your Hyper-V hosts to keep using the traditional SAN storage. Each member of the Hyper-V failover cluster can connect to iSCSI, Fibre Channel or Fibre Channel over Ethernet.
- You require that each member of the Hyper-V failover cluster be able to access the shared storage. This requirement is met by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

1. You require that you are able to repurpose previous Hyper-V hardware that ran Windows Server 2008 R2 servers. This requirement is met by reusing previous hardware and making a single change to the hardware configuration, which is to add a 10 GB network adapter that supports SR-IOV. In addition, you will need to update the BIOS if it doesn't current support SR-IOV. Alternatively, you could use a 1 GbE adapter and not deploy SR-IOV. In this document we will demonstrate how to use SR-IOV.
2. You require that access to virtual machines is resilient to hardware failure. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.

Overview of the Non-Converged Data Center Configuration

The Non-Converged Data Center cloud configuration consists of the following:

- Multiple computers in a Hyper-V failover cluster serving both the compute and storage roles.A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility and management which is critical in a cloud environment.
- A non-converged networking infrastructure that supports multiple cloud traffic profiles.Each computer in the Hyper-V failover cluster must have enough network adapters installed to support each traffic type that needs to be isolated from other traffic types. In most cases, this will include network adapters that are assigned to infrastructure traffic such as management, cluster/CSV, and Live Migration. If you decide to use an Ethernet-based storage protocol (such as iSCSI or FCoE), then you will also need a network adapter for the storage traffic. The infrastructure traffic network adapters in this scenario are non-teamed 1 GB adapters, although you can use 10 GB adapters or mix 1 GB and 10 GB adapters for the infrastructure traffic types you need to support. In addition, you also have the option of teaming any adapters you like, with the caveat being that RDMA is not supported on teamed adapters. The tenant traffic adapter is a 10 GB adapter that supports SR-IOV and is not teamed because SR-IOV and NIC Teaming are not compatible.
- SAN-based storage that can be Ethernet or non-Ethernet based. Storage options include iSCSI, Fibre Channel, Fibre Channel over Ethernet or Infiniband (non-Ethernet). You can also choose to use SAS storage together with Windows Server 2012 Storage Spaces.
- The appropriate networking hardware to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

The specific design pattern documented in the Non-Converged Data Center Configuration scenario includes the following:

- Multiple computers in a Hyper-V failover cluster
- Separate 1 GB network adapters over which live migration, cluster, management, and tenant traffic traverse.
- One 10 GB adapter that is dedicated to tenant traffic that supports SR-IOV. Note that in the SR-IOV scenario that all traffic bypasses the Hyper-V virtual switch, so all traffic isolation depends on the capabilities the unique SR-IOV NIC may provide in this area.
- SAN storage based on iSCSI, Fibre Channel, Fibre Channel over Ethernet or Infiniband.

Figure 3 depicts the Non-Converged Enterprise configuration.

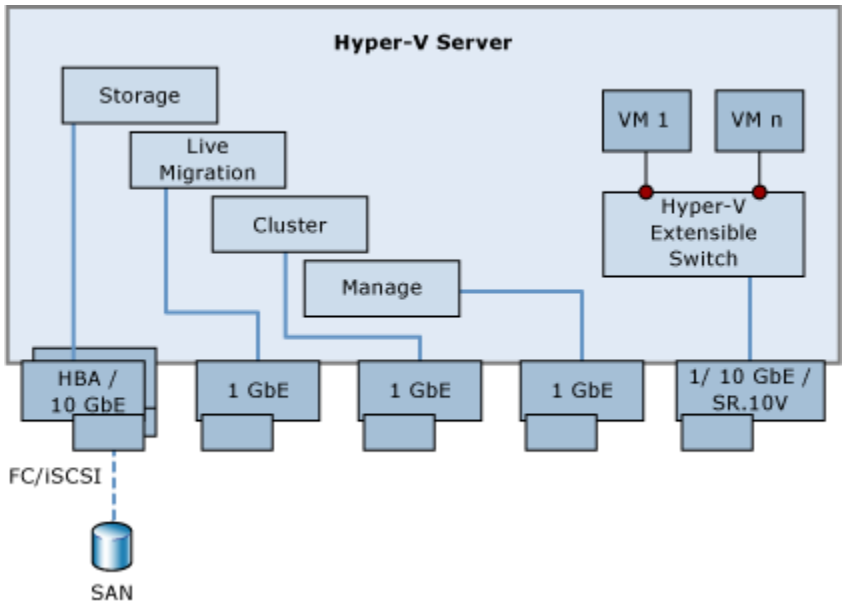

Figure 3: High level overview of cluster member networking configuration

This configuration highlights the following technologies and features of Windows Server 2012:

- Failover Clustering: A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. With the Failover Clustering feature, users experience a minimum of disruptions in service.
- Single Root I/O Virtualization: Single Root I/O Virtualization (SR-IOV) is a standard introduced by the PCI-SIG. SR-IOV works in conjunction with system chipset support for virtualization technologies. This provides remapping of interrupts and DMA and allows SR-IOV capable devices to be assigned directly to a virtual machine. Hyper-V in Windows Server 2012 enables support for SR-IOV–capable network devices and allows an SR-IOV virtual function of a physical network adapter to be assigned directly to a virtual machine. This increases network throughput and reduces network latency, while also reducing the host CPU overhead required for processing network traffic.

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

For detailed setup and configuration information that you can use to build the Non-Converged Data Center Configuration in your own datacenter, please see Building Your Cloud Infrastructure: Non-Converged Data Center Configuration.

The Converged Data Center with File Server Storage Configuration

The Converged Datacenter with File Server Storage cloud design pattern is essentially defined by the fact that the Hyper-V computer cluster is separate from the storage cluster. In the Non-Converged Datacenter Configuration cloud design pattern, the Hyper-V cluster was attached to block storage. In contrast, the Converged Data Center with File Server Storage Configuration uses a highly available file server cluster to host the VHDX and virtual machine configuration files. The nodes on the Hyper-V compute cluster access the virtual machine files over the network using the SMB 3.0 protocol. In addition, the Converged Data Center with File Server Storage Configuration uses a converged network pattern on the Hyper-V compute cluster where the infrastructure traffic profiles (management, cluster/CSV, Live Migration and storage) are hosted on one subnet/NIC and the tenant traffic is hosted on a separate subnet/NIC.

The Converged Data Center with File Server Storage Configuration uses:

- Two subnets - one for cloud infrastructure traffic (live migration, cluster, storage and management) one for tenant traffic.
- NIC Teaming for network bandwidth aggregation and failover for the infrastructure and tenant subnet.
- A dedicated file server storage cluster that hosts virtual machine VHDX and configuration files.
- A dedicated Hyper-V compute cluster that runs the virtual machine workloads.

Design Considerations and Requirements for the Converged Data Center with File Server Storage Pattern

The Converged Data Center with File Server Storage cloud infrastructure design pattern focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You require that cloud infrastructure traffic be physically separated from cloud tenant traffic. The requirement is met by creating separate NICs (or teams) for infrastructure and tenant traffic and connecting them to different subnets/segments. The tenant subnet may use NIC Teaming if you like in this scenario.
- You require that infrastructure traffic (i.e. live migration, cluster, storage, management) all receive guaranteed levels of bandwidth. The requirement is met by using Windows QoS policies on the parent Hyper-V partition.
- You require that tenant traffic from different tenants receive guaranteed levels of bandwidth. The requirement is met by using Hyper-V virtual switch QoS policies.

Storage

- You require the ability to scale and manage storage separately from the compute infrastructure. This requirement can be met by creating a dedicated storage cluster that will host the VHDX and virtual machine configuration files. The Hyper-V compute cluster will connect to the file server cluster using SMB 3.0 to connect to file shares on the storage

cluster. The file server cluster will be configured to use the new Windows Server 2012 Scale-out File Server feature so as to provide continuous availability of the virtual machine files.

- You require a low cost storage option. This requirement is met by using Serial Attached SCSI (SAS) disks in shared JBOD enclosures managed through Storage Spaces. Alternatively, if you have existing investments in other storage technologies, each member of the file server failover cluster can connect to iSCSI, Fibre Channel or Fibre Channel over Ethernet.
- You require a resilient storage solution. This requirement is met by using at least two file servers configured as a failover cluster, with well-connected (shared JBODs) storage so that all members of the file server failover cluster are directly connected to storage with Storage Spaces configured as a mirrored space to guarantee against data loss in the case of disk failures. Each cluster member has redundant connections to the JBOD enclosure and Windows Server 2012 MPIO is enabled. In addition, each member of the file server failover cluster is able to access the shared storage by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

1. You require the ability to scale your compute infrastructure separately from your storage infrastructure. This requirement is met by creating a dedicated Hyper-V compute cluster that connects to remote file server storage for virtual machines and virtual machine configuration files. Local disks on the compute nodes are only used for the boot partition but not for the virtual machines.
2. You require that virtual machines will be continuously available and resilient to hardware failures. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.
3. You require the highest number of virtual machines possible per host server (i.e. increased VM density). This requirement is met by using processor offload technologies, such as Remote Direct Memory Access (RDMA), Receive Side Scaling (RSS), Receive Side Coalescing (RSC), and Datacenter Bridging (DCB).

Overview of the Converged Data Center with File Server Storage Configuration

Based on the above requirements, the Converged Data Center with File Server Storage Configuration has the following characteristics:

- Multiple computers in a dedicated Hyper-V compute cluster.A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility and manageability which is critical in a cloud environment. The Hyper-V cluster is a dedicated compute cluster and does not host storage for virtual machines and virtual machine configuration files.
- Multiple computers in a dedicated Scale-out File Server storage cluster.A File Server failover cluster is created using the Windows Server 2012 Failover Cluster feature. Windows Server 2012 includes a new file server capability known as "Scale-out File Server for applications" that enables you to store virtual machine and virtual machine configuration files in a file share

and make these files continuously available. When you separate the file server cluster from the compute cluster you are able to scale compute and storage resources independently.

- A converged networking infrastructure that supports physical segmentation of infrastructure and tenant traffic.Each computer in the Hyper-V failover cluster must have at least two network adapters so that one adapter can host the cloud infrastructure traffic and one adapter can support tenant traffic. If resiliency against NIC failures is required, then you can add additional network adapters on each of the networks, and team them using Windows Server 2012 Load Balancing and Failover (LBFO) NIC Teaming. The NICs can be 10 GbE or 1 GbE network adapters. These NICs will be used for live migration, cluster, storage, management (together referred to as "infrastructure" traffic) and tenant traffic. However, keep in mind that NIC Teaming is not compatible with RDMA and is the reason why RDMA is not enabled on the infrastructure network.
- The appropriate networking hardware (e.g. Ethernet switches, cables, etc.) to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

Figure 4 provides a high-level view of the scenario layout. Key elements of the configuration include:

- A File Server cluster that hosts the virtual hard disks and virtual machine configuration files
- The File Server cluster is connected only to the infrastructure network
- The File Server cluster is connected to block storage either through HBAs or over a 10 GB Ethernet network (as in the case of iSCSI or Fibre Channel of Ethernet [FCoE])
- A Hyper-V compute cluster that hosts the virtual machine workloads.
- The Hyper-V compute cluster is connected to the datacenter (infrastructure) network and the tenant network using teamed network adapters.
- Cloud infrastructure traffic flows to and from the host-based 10GbE network adapter team
- Tenant network traffic to and from the virtual machines flows through the Hyper-V virtual switch which is bound to the Tenant Network NIC team.

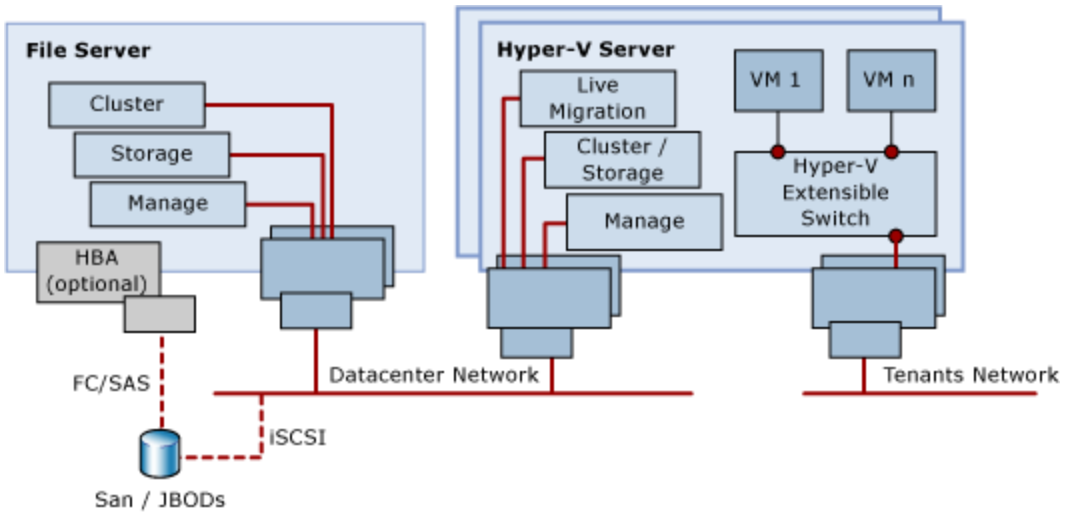

Figure 4: High level overview of Converged Datacenter with File Server storage infrastructure design pattern

Note

Although this configuration uses SAS storage on the file server cluster, you can easily choose to use other types of storage, such as iSCSI or Fibre Channel-based SAN storage. You can find more information about storage configuration for a non-SAS scenario in the document Building Your Cloud Infrastructure: Non-Converged Data Center Configuration, which describes how to configure the SAN storage.

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

This configuration highlights the following technologies and features of Windows Server 2012:

- **Load Balancing and Failover (LBFO)**: Load Balancing and Failover (NIC Teaming) logically combines multiple network adapters to provide bandwidth aggregation and traffic failover to prevent connectivity loss in the event of a network component failure. Load Balancing with Failover is also known as NIC Teaming in Windows Server 2012.
- **Windows Server Quality of Service (QoS)**: Windows Server 2012 includes improved QoS features that lets you manage bandwidth and provide predictable network performance to traffic to and from the host operating system.
- **Hyper-V Extensible Switch Quality of Service (QoS)**: In Windows Server 2012, The Hyper-V extensible switch includes QoS features that can be applied on virtual switch ports, and that lets you manage bandwidth and provide predictable network performance to virtual machines running on that server.
- **Data Center Bridging (DCB)**: DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control. Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support Internet Small Computer System Interface (iSCSI), Remote Direct Memory Access (RDMA) over Converged Ethernet, or Fiber Channel over Ethernet (FCoE).
- **Storage Spaces**: Storage Spaces makes it possible for you to create cost-effective disk pools that present themselves as a single mass storage location on which virtual disks or volumes can created and formatted.

For detailed setup and configuration information that you can use to build the Non-Converged Data Center Configuration in your own datacenter, please see Building Your Cloud Infrastructure: Converged Data Center with File Server Storage.

The Converged Data Center without Dedicated Storage Nodes Configuration

The Converged Data Center without Dedicated Storage Nodes cloud infrastructure design pattern takes advantage of new capabilities in Windows Server 2012 that enable you to converge all traffic profiles (infrastructure and tenant) to run through a single network adapter or NIC Team. This differs from the previous two design patterns, where in the first one each traffic profile was dedicated to a physical network adapter and where in the second design pattern infrastructure traffic and tenant traffic were carried over different network adapters and subnets. The Converged Data Center without Dedicated Storage Nodes cloud design pattern can greatly simplify your cloud infrastructure by reducing the number of cables and ports you need to deal with.

The Converged Datacenter without Dedicated Storage Nodes Configuration includes:

- A converged network infrastructure for live migration, cluster, storage, management, and tenant traffic
- All network traffic moves through the Hyper-V virtual switch, this includes the host operating system traffic that moves over the infrastructure network.
- Hyper-V Virtual Switch Quality of Service (QoS)
- Hyper-V Virtual Switch port ACLs and 802.1q VLAN tagging
- NIC Teaming for network bandwidth aggregation and failover
- Well-connected storage using SAS JBOD enclosures

Design Considerations and Requirements for the Converged Data Center without Dedicated Storage Node Pattern

The Converged Data Center without Dedicated Storage Nodes cloud infrastructure design pattern focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You prefer that network traffic to and from both the host operating system and the guest operating systems running on the host move through a single network adapter team. This requirement is met by using Windows Server 2012 NIC Teaming (LBFO) and passing all traffic through the Hyper-V virtual switch.
- You require that live migration, cluster, storage, management and tenant traffic all receive guaranteed levels of bandwidth. The requirement is met by using Hyper-V virtual switch QoS policies.
- You require that infrastructure traffic (which includes Live Migration, cluster, storage and management traffic) and tenant traffic be isolated from each other. This requirement is met by using Hyper-V virtual switch port ACLs and 802.1q VLAN tagging.

Storage

- You prefer to scale your cloud infrastructure by adding scale units consisting of compute and storage capacity together. This requirement is met by connecting the Hyper-V servers directly to SAS storage, without having dedicated file servers. Note that you can also use SAN based storage connected to the compute cluster in this design pattern.

- You require cost-effective storage. This requirement is met by using SAS disks in shared JBOD enclosures managed through Storage Spaces.
- You require a resilient storage solution. This requirement is met by having multiple Hyper-V servers configured as a failover cluster, and having a well-connected (shared JBODs) storage so that all members of the failover cluster are directly connected to storage, and by having Storage Spaces configured as a mirrored space to guarantee against data loss in the case of disk failures
- You require that each member of the Hyper-V failover cluster be able to access the shared storage where the VHDs are located. This requirement is met by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

- You require that the virtual machines will be continuously available and resilient to hardware failures. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.
- You require the highest number of virtual machines possible per host server (i.e. increased density). This requirement is met by using processor offload technologies, such as Remote Direct Memory Access (RDMA), Receive Side Scaling, Receive Side Coalescing (RSC), and Datacenter Bridging (DCB). Please note that in the default configuration presented here (without a dedicated storage access NIC), RDMA and DCB cannot be used because these technologies require direct access to the hardware and must bypass much of the virtual networking stack and this compatible with NIC Teaming. This is similar to the situation with Single Root I/O Virtualization (SR-IOV). For optimal performance, especially in the context of network access to storage, a separate NIC team would be required to support these hardware offload acceleration technologies.

The Converged Datacenter without Dedicated Storage nodes cloud infrastructure design pattern consists of the following:

- Multiple computers in a Hyper-V failover cluster. A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility which is critical in a cloud environment. For example, Live Migration is enhanced when performed in a failover cluster deployment because the cluster can automatically evaluate which node in the cluster is optimal for migrated virtual machine placement.
- A converged networking infrastructure that supports multiple cloud traffic profiles. Each computer in the Hyper-V failover cluster should have at least two network adapters that will be used for the converged network. This converged network will host all traffic to and from the server, which includes both host system traffic and guest/tenant traffic. The network adapters will be teamed by using Windows Server 2012 Load Balancing and Failover (LBFO) NIC Teaming. The NICs can be either two or more 10 GbE or 1 GbE network adapters. These NICs will be used for live migration, cluster, storage, management (together referred to as "infrastructure" traffic) and tenant traffic.

- The appropriate networking hardware to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

The following figure provides a high-level view of the scenario architecture. The teamed network adapters on each member of the failover cluster are connected to what will be referred to as a converged subnet in this document. We use the term converged subnet to make it clear that all traffic to and from the Hyper-V cluster members and the tenant virtual machines on each cluster member must flow through the teamed converged subnet network adapter. Both the host operating system and the tenants connect to the network through the Hyper-V virtual switch. The figure also shows an optional network adapter that is RDMA-capable that can be used for storage traffic, such as when storage is being hosted on a share on a remote file server.

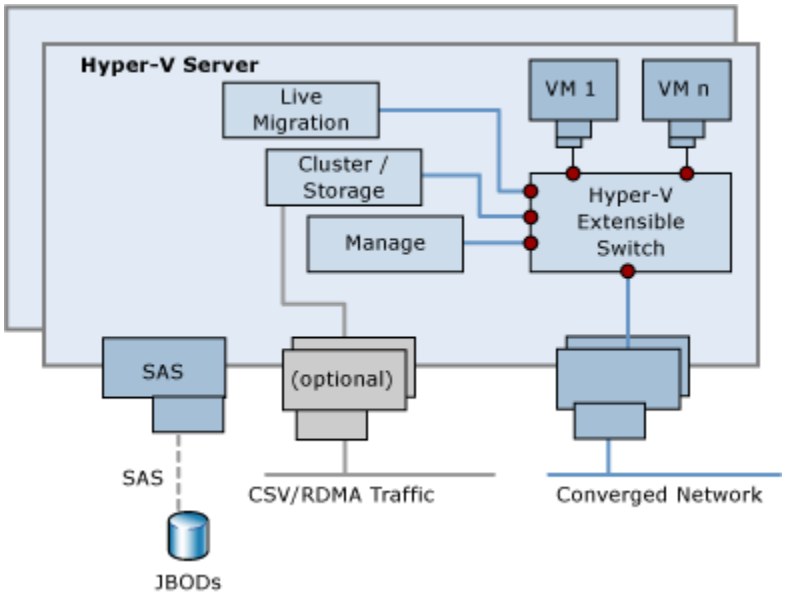

Figure 5 High level overview of cluster member networking configuration

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

Figure 6 provides an overview of traffic flows on each member of the Hyper-V cluster. The figure calls out the following significant issues in the configuration:

- Each cluster node member uses a virtual network adapter to connect to the Hyper-V Extensible Switch, which connects it to the physical network.
- Each tenant virtual machine is also connected to the Hyper-V Extensible Switch using a virtual network adapter.
- Network adapters named ConvergedNet1 and ConvergedNet2 participate in a teamed physical network adapter configuration using the Windows Server 2012 Failover and Load Balancing feature.

- Windows Server 2012 Hyper-V virtual switch QoS is used to assure that each traffic type (such as live migration, cluster, management and tenant) has a predictable amount of bandwidth available.
- Traffic isolation is enabled by 802.1q VLAN tagging so that host traffic is not visible to the tenants.
- Windows Server 2012 Hyper-V virtual switch port ACLs can also be used for more granular access control at the network level.

It is important to note that Remote Direct Memory Access (RDMA) cannot be used on the converged network because it does not work together with the Hyper-V virtual switch and NIC Teaming. This will be an issue if you prefer to use high performance SMB 3 connectivity to file server based storage for virtual machine disk and configuration files.

Note

Virtual local area networks (VLANs) are not assigned to each tenant because VLAN-based network isolation is not a scalable solution and is not compatible with Windows Server 2012 network virtualization. VLANs are used to isolate infrastructure traffic from tenant traffic in this scenario.

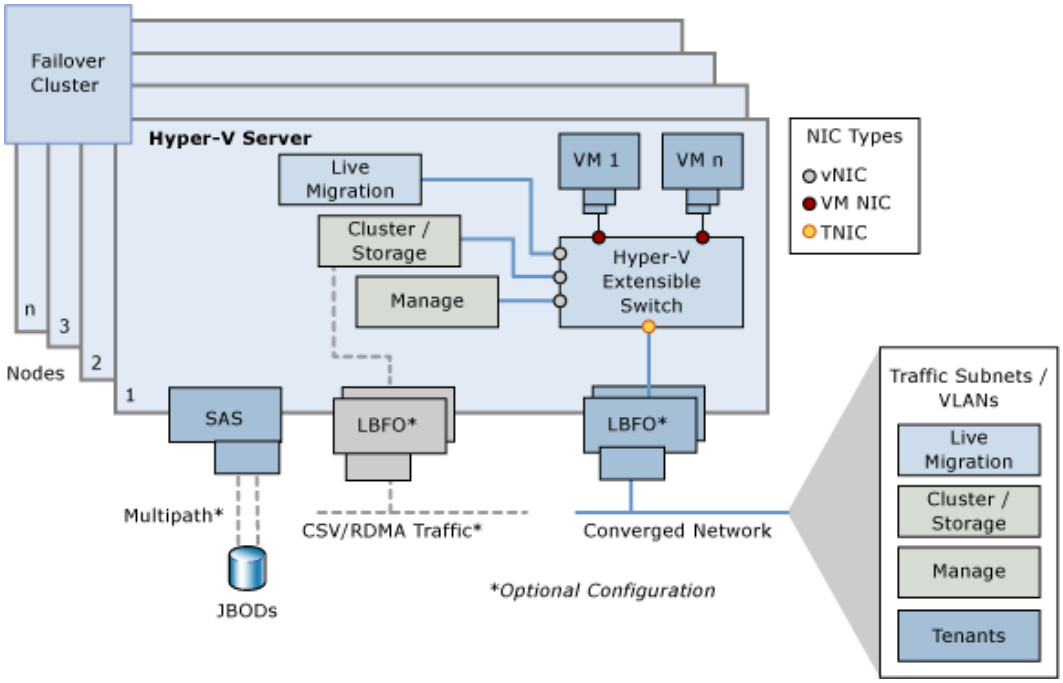

Figure 6 Overview of cluster member traffic flows

This configuration highlights the following technologies and features of Windows Server 2012:

- **Load Balancing and Failover (LBFO)**: Load Balancing and Failover logically combines multiple network adapters to provide bandwidth aggregation and traffic failover to prevent connectivity loss in the event of a network component failure. Load Balancing with Failover is also known as NIC Teaming in Windows Server 2012.

- **Hyper-V Virtual Switch Quality of Service (QoS)**: In Windows Server 2012, QoS includes new bandwidth management features that let you provide predictable network performance to virtual machines on a server running Hyper-V.
- **Hyper-V Virtual Switch Quality of Service (QoS)**: In Windows Server 2012 the Hyper-V virtual switch includes new capabilities that enhance the security of the cloud infrastructure. You can now use Port Access Control Lists (Port ACLs) and VLAN support to get network isolation similar to what you find when using physical network isolation.
- **Storage Spaces**: Storage Spaces makes it possible for you to create cost-effective disk pools that present themselves as a single mass storage location on which virtual disks or volumes can created and formatted.

For detailed setup and configuration information that you can use to build the Non-Converged Data Center Configuration in your own datacenter, please see Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes.

Building Your Cloud Infrastructure: Non-Converged Data Center Configuration

This document contains the instructions that you need to follow to create a private or public cloud configuration that uses:

- Separate network adapters for live migration, cluster, management, and tenant traffic
- Traditional SAN storage
- Optionally use Single Root I/O Virtualization (SR-IOV)

The design pattern discussed in this document is one of three design patterns we suggest for building the core cloud compute and storage infrastructure. For information about the other two cloud infrastructure design patterns, please see:

- Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes
- Building Your Cloud Infrastructure: Converged Data Center with File Server Storage

Design Considerations and Requirements for the Non-Converged Data Center Configuration Pattern

The Non-Converged Data Center Configuration cloud infrastructure design pattern focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You have an existing investment in separate networks based on the recommended configuration of Hyper-V in Windows Server® 2008 R2 and you require that this physical network traffic segmentation be kept in place to avoid re-architecting your network. Each type of infrastructure traffic (management, cluster/CSV, Live Migration and storage) and tenant traffic are carried over physically separate networks and network adapters. This requirement

is met by installing physically separate NICs for each traffic type and assigning VLAN 802.1q tags to each adapter.

- You require that each traffic type is dedicated to a specific adapter. This requirement is met by configuring each of the traffic flows to use the correct subnet/IP on the dedicated NIC.
- You require that the virtual machine workload has access to the highest network performance possible. This requirement is met by using the new Windows Server® 2012 Single Root I/O Virtualization (SR-IOV) feature, which enables the virtual machines running on the Hyper-V server to have direct access to the network adapter hardware, thus bypassing the virtual networking stack.

Storage

- You have an existing investment in FC or iSCSI SANs and require well-connected storage so that all members of the Hyper-V failover cluster are connected to block storage. This requirement is met by configuring your Hyper-V hosts to keep using the traditional SAN storage. Each member of the Hyper-V failover cluster can connect to iSCSI, Fibre Channel or Fibre Channel over Ethernet.
- You require that each member of the Hyper-V failover cluster be able to access the shared storage. This requirement is met by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

- You require that you are able to reuse previous Hyper-V hardware that ran Windows Server 2008 R2 server. This requirement is met by reusing previous hardware and making a single change to the hardware configuration, which is to add a 10GbE network adapter that supports SR-IOV. Alternatively, you could use a 1 GbE adapter and not deploy SR-IOV. In this document we will demonstrate how to use SR-IOV.
- You require that access to virtual machines is resilient to hardware failure. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.

Overview

A Windows Server 2012 cloud infrastructure is a high-performing and highly available Hyper-V cluster that hosts virtual machines that can be managed to create private or public clouds using the Non-Converged Data Center Configuration infrastructure design pattern. This document explains how to configure the basic building blocks for such a cloud. It does not cover the System Center or other management software aspects of deployments; the focus is on configuring the core Windows Server hosts that are used to build cloud infrastructure.

This cloud configuration consists of the following:

- Multiple computers in a Hyper-V failover cluster.

A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility which is critical in a cloud environment.

- A non-converged networking infrastructure that supports multiple cloud traffic profiles.

  Each computer in the Hyper-V failover cluster must have enough network adapters installed to support each traffic type that needs to be isolated from other traffic types. In most cases, this will include network adapters that are assigned to infrastructure traffic such as management, cluster/CSV, and Live Migration. If you decide to use an Ethernet-based storage protocol (such as iSCSI or FCoE), then you will also need a network adapter for the storage traffic. The infrastructure traffic network adapters in this scenario are non-teamed 1 GB adapters, although you can use 10 GB adapters or mix 1 GB and 10 GB adapters for the infrastructure traffic types you need to support. The tenant traffic adapter is a 10 GB adapter that supports SR-IOV.
- SAN-based storage what can be Ethernet or non-Ethernet based. Storage options include iSCSI, Fibre Channel, Fibre Channel over Ethernet or Infiniband.
- The appropriate networking hardware to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

This configuration highlights the following technologies and features of Windows Server 2012:

- **Failover Clustering**: A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. With the Failover Clustering feature, users experience a minimum of disruptions in service.
- **Single Root I/O Virtualization**: Single Root I/O Virtualization (SR-IOV) is a standard introduced by the PCI-SIG. SR-IOV works in conjunction with system chipset support for virtualization technologies. This provides remapping of interrupts and DMA and allows SR-IOV capable devices to be assigned directly to a virtual machine. Hyper-V in Windows Server 2012 enables support for SR-IOV–capable network devices and allows an SR-IOV virtual function of a physical network adapter to be assigned directly to a virtual machine. This increases network throughput and reduces network latency, while also reducing the host CPU overhead required for processing network traffic.

Non-Converged Data Center Configuration Scenario Overview

The Non-Converged Data Center Configuration scenario includes the following:

- Multiple computers in a Hyper-V failover cluster
- Separate 1 GB network adapters over which live migration, cluster, management, and tenant traffic traverse.
- One 10 GB adapter that is dedicated to tenant traffic that supports SR-IOV. Note that in this scenario that all traffic bypasses the Hyper-V virtual switch, so all traffic isolation depends on the capabilities of the unique SR-IOV may provide in this area.
- SAN storage based on iSCSI, Fibre Channel, Fibre Channel over Ethernet or Infiniband.

Figure 1 shows the Non-Converged Data Center configuration. The following are the salient aspects of this cloud infrastructure design pattern:

- SAN connectivity is enabled either via an HBA or 10 GB Ethernet connection (depending on the type of SAN deployed)
- Live Migration, Cluster/CSV and management traffic is assigned to separate 1 GB network adapters
- Virtual machine workloads connect to the tenant network using a 10 GB SR-IOV capable adapter. The Hyper-V extensible switch is configured to support SR-IOV to enable the virtual machine to communicate directly with the adapter hardware.

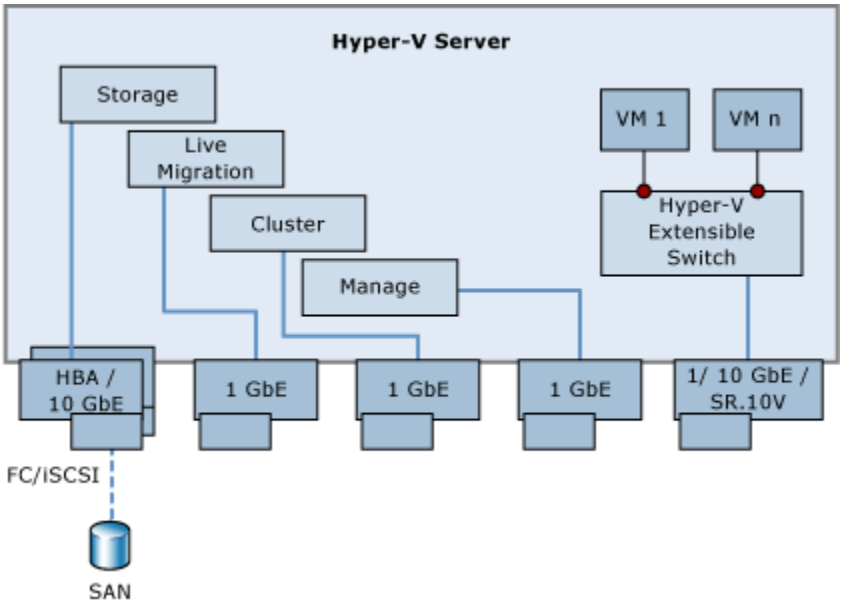

Figure 1 High level overview of cluster member networking configuration

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

The following sections describe how to set up this cloud configuration using UI-based tools and Windows PowerShell.

- After the cloud is built, you can validate the configuration by doing the following:
- Install and configure virtual machines
- Migrate running virtual machines between servers in the Hyper-V cluster (live migration)
- Failover storage from one owner to another

Install and configure

Creating this cloud infrastructure configuration consists of the following steps:

- Step 1: Initial node configuration
- Step 2: Initial network configuration
- Step 3: Initial storage configuration
- Step 4: Failover cluster setup
- Step 5: Configure Hyper-V settings
- Step 6: Cloud validation

The following table summarizes the steps that this document describes:

Step	Task	Target	Tasks

1	Initial Node Configuration	All Nodes	• 1.1-Upgrade the machine to the latest BIOS and enable BIOS settings required for Hyper-V and SR-IOV
			• 1.2-Perform a clean operating system installation
			• 1.3-Perform post installation tasks:
			  • Set Windows PowerShell execution policy
			  • Enable Windows PowerShell remoting
			  • Enable Remote Desktop Protocol and Firewall rule
			  • Join the domain
			• 1.4-Install roles and features using default settings, rebooting as needed
			  • Hyper-V (plus management tools)
			  • Failover clustering (plus management tools)

2	Initial Network Configuration	All Nodes	• 2.1-Disable unused and disconnected interfaces and rename active connections (static IP addresses on all interfaces)
3	Initial Storage Configuration	Single Node	• 3.1-Present all shared storage to relevant nodes • 3.2-For multipath scenarios, install and configure multipath I/O (MPIO) as necessary

4	Failover Cluster Setup	Single Node	• 4.1-Run through the Cluster Validation Wizard
			• 4.2-Address any indicated warnings and/or errors
			• 4.3-Complete the Create Cluster Wizard (setting name and IP address but do not add eligible storage)
			• 4.4-Create a witness disk
			• 4.5-Create a virtual machine storage disk
			• 4.6-Add the virtual machine storage disk to cluster shared volumes
			• 4.7-Add folders to the cluster shared volume:
			  • C:\ClusterStorage\Volume1\VHDdisks\
			  • C:\ClusterStorage\Volume1\VHDsettings\
			• 4.8-Configure quorum settings
			• 4.9-Configure cluster networks to prioritize traffic

5	Hyper-V Configuration	All Nodes	• 5.1-Create the virtual switch
			• 5.2-Change default file locations, mapping to CSV volumes:
			  • C:\ClusteredStroage\Volume1\VHDfiles\
			  • C:\ClusterStorage\Volume1\VHDsettings\

6	Cloud Validation	Single Node	• 6.1-Create a virtual machine, attaching an existing operating system VHD and tagging to the appropriate VLAN • 6.2-Test network connectivity from the virtual machine • 6.3-Perform a Live Migration • 6.4-Perfomr a quick migration

Step 1: Initial node configuration

In step 1, you will perform the following steps on all nodes of the Hyper-V cluster:

- 1.1 Upgrade the machine to the latest BIOS release and enable BIOS settings for Hyper-V and SR-IOV.
- 1.2 Perform a clean operating system installation.
- 1.3 Perform post-installation tasks.
- 1.4 Install roles and features using the default settings.

1.1 Enable BIOS settings required for Hyper-V for SR-IOV

You will need to enable virtualization support in the BIOS of each cluster member prior to installing the Hyper-V server role. The procedure for enabling processor virtualization support will vary with your processors' make and model and the system BIOS. In addition, if you choose to enable SR-IOV, this capability will need to be enabled in the system BIOS. Please refer to your hardware documentation for the appropriate procedures. In addition, make sure to upgrade your servers with their latest BIOS revision.

1.2 Perform a clean operating system installation

Install Windows Server 2012 using the Full Installation option.

1.3 Perform post-installation tasks

There are several tasks you need to complete on each node after the operating system installation is complete. These include:

- Join each node to the domain
- Enable remote access to each node via the Remote Desktop Protocol.
- Set the Windows PowerShell execution policy.
- Enable Windows PowerShell remoting.

Perform the following steps to join each node to the domain:

1. Press the **Windows Key** on the keyboard and then press **R**. Type **Control Panel** and then click **OK**.
2. In the **Control Panel** window, click **System and Security**, and then click **System**.
3. In the **System** window under **Computer name, domain, and workgroup settings**, click **Change settings**.
4. In the **System Properties** dialog box, click **Change**.
5. Under **Member of**, click **Domain**, type the name of the domain, and then click **OK**.

Run the following Windows PowerShell commands on each node to enable remote access using the Remote Desktop Protocol, to enable PowerShell execution policy and enable PowerShell Remoting:

```
(Get-WmiObject Win32_TerminalServiceSetting -Namespace
root\cimv2\terminalservices).SetAllowTsConnections(1,1)

Set-ExecutionPolicy Unrestricted –Force

Enable-PSRemoting –Force
```

1.4 Install roles and features using the default settings

The following roles and features will be installed on each node of the cluster:

- Hyper-V and Hyper-V management Tools
- Failover cluster and failover cluster management tools

Perform the following steps on each node in the cluster to install the required roles and features:

1. In **Server Manager**, click **Dashboard** in the console tree.
2. In **Welcome to Server Manager**, click **2 Add roles and features**, and then click **Next**.
3. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
4. On the **Installation Type** page, click **Next**.
5. On the **Server Selection** page, click **Next**.
6. On the **Server Roles** page, select **Hyper-V** from the **Roles** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Click **Next**.
7. On the **Features** page, select **Failover Clustering** from the **Features** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Click Next.

   Note
   If you plan to use Multipath I/O for your storage solution, select the Multipath I/O feature while performing step 7.
8. On the **Hyper-V** page, click **Next**.
9. On the **Virtual Switches** page, click **Next**.
10. On the **Migration** page, click **Next**.
11. On the **Default Stores** page, click **Next**.
12. On the **Confirmation** page, put a checkmark in the **Restart the destination server automatically if required** checkbox and then in the **Add Roles and Features** dialog box click **Yes**, then click **Install**.
13. On the **Installation progress** page, click **Close** after the installation has succeeded.
14. Restart the computer. This process might require restarting the computer twice. If so, the installer will trigger the multiple restarts automatically.

After you restart the server, open **Server Manager** and confirm that the installation completed successfully. Click **Close** on the **Installation Progress** page.

Step 2: Initial network configuration

The network configuration on each node in the cluster needs to be configured to support the non-converged networking scenario where all traffic types move through their own dedicated network adapters. You will perform the following procedures on each of the nodes in the cluster to complete the initial network configuration:

- 2.1 Disable unused and disconnected interfaces and rename active connections.

2.1 Disable unused and disconnected interfaces and rename active connections

You can simplify the configuration and avoid errors when running the wizards and running PowerShell commands by disabling all network interfaces that are either unused or disconnected. You can disable these network interfaces in the **Network Connections** window.

For the remaining network adapters, do the following:

1. Connect them to the network switch ports.

2. Configure appropriate IP addressing information for each network adapter based on your network addressing scheme for each traffic type. Make sure that each traffic type is assigned to a different network ID (subnet ID).
3. To help you more easily recognize the active network adapters, rename them with names that indicate their use or their connection to the intranet or Internet (for example, **ManagementNet** and **LiveMigrationNet** and **ClusterNet** and **TenantNet**). You can do this in the **Network Connections** window.

Note
Configure a DNS server and default gateway address only on the management interface adapter.

Step 3: Initial storage configuration

With the initial cluster node configuration complete, you are ready to perform initial storage configuration tasks on all nodes of the cluster. Initial storage configuration tasks include:

- 3.1 Present all shared storage to relevant nodes.
- 3.2 Install and configure MPIO as necessary for multipath scenarios.

3.1 Present all shared storage to relevant nodes

Connect the host bus adapter (HBA) fiber optic cables to a FC hub or switch that is connected to the storage device. Each cluster node should have two HBAs in them if high availability of storage access is required.

Configure the storage device with the appropriate RAID settings and volumes (LUNs) so that the storage on the device is available from the Disk Management console of the Computer Management snap-in of each cluster member computer.

For the first computer that is attached to the storage, create the following from this storage space:

- A 1 GB NTFS volume that may be used for the cluster witness disk.
- An NTFS volume with an appropriate amount of space that will be used to store the virtual machine folders and files of the Hyper-V cluster.

3.2 Install and configure MPIO as necessary for multipath scenarios

If you have multiple data paths to storage (for example, two SAS cards) make sure to install the Microsoft® Multipath I/O (MPIO) on each node. This step might require you to restart the system. For more information about MPIO, see What's New in Microsoft Multipath I/O.

Step 4: Failover cluster setup

You are now ready to complete the failover cluster settings. Failover cluster setup includes the following steps:

- 4.1 Run through the Cluster Validation Wizard.
- 4.2 Address any indicated warnings and/or errors.

- 4.3 Complete the Create Failover Cluster Wizard.
- 4.4 Create the witness virtual disk.
- 4.5 Create the virtual machine storage virtual disk.
- 4.6 Add the virtual machine storage virtual disk and Witness disk to Cluster Shared Volumes.
- 4.7 Add folders to the cluster shared volume.
- 4.8 Configure Quorum Settings.
- 4.9 Configure cluster networks to prioritize traffic.

4.1 Run through the Cluster Validation Wizard

The Cluster Validation Wizard will query multiple components in the intended cluster hosts and confirm that the hardware and software is ready to support failover clustering. On one of the nodes in the server cluster, perform the following steps to run the Cluster Validation Wizard:

1. In the **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** console, in the **Management** section, click **Validate Configuration**.
3. On the **Before You Begin** page of the **Validate a Configuration Wizard**, click **Next**.
4. On the **Select Servers or a Cluster** page, type the name of the local server, and then click **Add**. After the name appears in the **Selected servers** list, type the name of another Hyper-V cluster member computer, and then click **Add**. Repeat this step for all computers in the Hyper-V cluster. When all of the servers of the Hyper-V cluster appear in the **Selected servers** list, click **Next**.
5. On the **Testing Options** page, click **Next**.
6. On the **Confirmation** page, click **Next**. The time to complete the validation process will vary with the number of nodes in the cluster and can take some time to complete.
7. On the **Summary** page, the summary text will indicate that the configuration is suitable for clustering. Confirm that there is a checkmark in the **Create the cluster now using the validated nodes...** checkbox.

4.2 Address any indicated warnings and/or errors

Click the **Reports** button to see the results of the Cluster Validation. Address any issues that have led to cluster validation failure. After correcting the problems, run the Cluster Validation Wizard again. After the cluster passes validation, then proceed to the next step. Note that you may see errors regarding disk storage. You may see this if you haven't yet initialized the disks. Click **Finish**.

4.3 Complete the Create Failover Cluster Wizard

After passing cluster validation, you are ready to complete the cluster configuration.

Perform the following steps to complete the cluster configuration:

1. On the **Before You Begin** page of the **Create Cluster Wizard**, click **Next**.

2. On the **Access Point for Administering the Cluster** page, enter a valid NetBIOS name for the cluster, and then select the network you want the cluster on and then type in a static IP address for the cluster object, and then click **Next**. In this example, the network you would select is the Management Network. Unselect all other networks that appear here.
3. On the **Confirmation** page, clear **Add all eligible storage to the cluster** checkbox and then click **Next**.
4. On the **Creating New Cluster** page you will see a progress bar as the cluster is created.
5. On the **Summary** page, click **Finish**.
6. In the console tree of the **Failover Cluster Manager** snap-in, open the **Networks** node under the cluster name.
7. Right-click the cluster network that corresponds to the management network adapter network ID (subnet), and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is *not* selected and that **Allow clients to connect through this network** is enabled. In the **Name** text box, enter a friendly name for this network (for example, ManagmentNet), and then click **OK**.
8. Right-click the cluster network that corresponds to the Cluster network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is not enabled. In the **Name** text box, enter a friendly name for this network (for example, ClusterNet), and then click **OK**.
9. Right-click the cluster network that corresponds to the live migration network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is not enabled. In the **Name** text box, enter a friendly name for this network (for example, LiveMigrationNet), and then click **OK**.

4.4 Create the witness disk

You will need a disk of at least 1 GB and a volume on it formatted as NTFS to use it as a witness disk. Later you will use this disk when you configure the quorum model.

4.5 Create the virtual machine storage disk

You will next need to a disk that you will place the virtual machine disk and configuration files on. This disk can be of any size but should be large enough to support the number of virtual machines you want to run on the cluster. You can use multiples disks if you like. This disk must have volumes on it that are formatted as NTFS so that they can be placed into a cluster shared volume.

4.6 Add the virtual machine storage disk and Witness disk to Cluster Shared Volumes

The disk you created for virtual machine storage is now ready to be added to a Cluster Shared Volume. Perform the following steps to add the virtual disk to a Cluster Shared Volume.

1. In the **Failover Cluster Manager**, in the left pane of the console, expand the **Storage** node and click **Disks**. In the middle pane of the console, in the **Disks** section, right click the disk you created in the previous step and then click **Add to Cluster Shared Volumes**.
2. Repeat step 1 to add the witness disk to a cluster shared volume.

4.7 Add folders to the cluster shared volume

Now you need to create the folders on the disk located on the Cluster Shared Volume to store the virtual machine files and the virtual machine data files.

Perform the following steps to create the folders that will store the running VMs of the Hyper-V cluster:

1. Open **Windows Explorer** and navigate to the C: drive and then double-click **Cluster Storage** and then double-click **Volume 1**.
2. Create two folders in **Volume 1**. One of the folders will contain the .vhd files for the virtual machines (for example, VHDdisks) and one folder will contain the virtual machine configuration files (for example, VHDsettings)

4.8 Configure Quorum Settings

Perform the following steps to configure quorum settings for the cluster:

1. In the left pane of the **Failover Cluster Manager** console, right click on the name of the cluster and click **More Actions** and click **Configure Cluster Quorum Settings**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Quorum Configuration Option** page, select **Use typical settings (recommended)** and click **Next**.
4. On the **Confirmation** page, click **Next**.

4.9 Configure cluster networks to prioritize traffic

The cluster will use the network with the lowest metric for CSV traffic and the second lowest metric for live migration. Windows PowerShell® is the only method available to prescriptively specify the CSV network.

Run the following Windows PowerShell commands on one node of the failover cluster to set the metric for the cluster network traffic, set the metric for the live migration network traffic and set the metric for the management network traffic:

```
(Get-ClusterNetwork "ClusterNet" ).Metric = 100

(Get-ClusterNetwork "LiveMigrationNet" ).Metric = 500

(Get-ClusterNetwork "ManagementNet" ).Metric = 1000.
```

Step 5: Configure Hyper-V settings

To finalize the Hyper-V configuration, you will need to:

- 5.1 Create the Hyper-V virtual switch

- 5.2 Change default file locations for virtual machine files.

5.1 Create the Hyper-V virtual switch

On *each Hyper-V cluster member*, perform the following steps to create the Hyper-V virtual switch that will be used by the tenant virtual machines:

1. Right-click the network icon in the notification area of the desktop, and then click **Open Network and Sharing Center**.
2. From the **Network and Sharing Center** window, click **Change adapter settings**.
3. In the **Network Connections** window, double-click the adapter with the name of the network adapter that is attached to the tenant subnet (for example, TenantNet), and then click **Details**.
4. In the **Network Connections Details** window, note the value of the **Description** field, and then click **Close**. You will need this information to determine which physical network adapter to bind to the virtual switch.
5. In the console tree of the **Hyper-V Manager**, right-click the name of the server, and then click **Virtual Switch Manager**.
6. In **Create virtual switch**, click **External**, and then click **Create virtual switch**.
7. In **Name**, enter a name for the virtual network to indicate that it is connected to the tenant subnet (for example, **TenantNetSwitch**).
8. On **External**, select the description of the network adapter attached to the Tenant subnet, as noted in step 7, select **Enable Single-Root I/O Virtualization (SR-IOV)**, and then click **OK**.

   Note
   If your hardware does not support SR-IOV or if SR-IOV is not enabled in the BIOS, the option will not be available. You can use the Get-NetAdapterSriov command to help make this assessment http://technet.microsoft.com/en-us/library/jj130915.aspx
9. When prompted with the **Apply Networking Changes** window, click **Yes**.

   Note
   If you are performing these steps over a remote desktop connection, you will temporarily lose connectivity.

5.2 Change default file locations for virtual machine files

On *each Hyper-V cluster member*, perform the following steps to change the default file locations for virtual machine files:

1. In **Server Manager**, click **Tools**, then click **Hyper-V Manager**.
2. From the console tree of the **Hyper-V Manager**, right-click the name of the Hyper-V server, and then click **Hyper-V Settings**.
3. In the **Hyper-V Settings** dialog box, click **Virtual Hard Disks** under **Server**, type the path to the folder where the virtual hard disk files are stored or use the **Browse** button in **Specify the default folder to store virtual hard disk files**, and then click **Apply**.

4. Click **Virtual Machines under Server**, type the file folder location or use the **Browse** button in **Specify the default folder to store virtual machine configuration files**, and then click **OK**

Step 6: Cloud validation

To verify the configuration of your cloud environment, perform the following operations.

- 6.1 Create a new virtual machine.
- 6.2 Test network connectivity from the virtual machine.
- 6.3 Perform a live migration.
- 6.4 Perform a quick migration.

6.1 Create a new virtual machine

To create a new virtual machine in the cluster environment, perform the following steps.

1. **Open Failover Cluster Manager**, click **Roles** under the cluster name, click **Virtual Machines** under the **Actions** pane, and then click **New Virtual Machine**.
2. On the **New Virtual Machine** page, select the cluster node where you want to create the virtual machine, and then click **OK**.
3. On the **Before you Begin** page of the **New Virtual Machine Wizard**, click **Next**.
4. On the **Specify Name and Location** page, enter a friendly name for this virtual machine and then click **Next**.
5. On the **Assign Memory** page, enter the amount of memory that will be used for this virtual machine (minimum for this lab is 1024 MB RAM) and then click **Next**.
6. On the **Configuring Networking** page, select the **TenantNetSwitch** and then click **Next**.
7. On the **Connect Virtual Hard Disk** page, leave the default options selected and click **Next**.
8. On the **Installation Options** page, select **Install an operating system from a boot CD/DVD-ROM** and then select the location where the CD/DVD is located. If you are installing the new operating system based on an ISO file, make sure to select the option Image file (.iso) and browse for the file location. After you select the appropriate option for your scenario, click **Next**.
9. On the **Completing the New Virtual Machine Wizard** page, review the options, and then click **Finish**.
10. The virtual machine creation process starts. After it is finished, you will see the **Summary** page, where you can access the report created by the wizard. If the virtual machine was created successfully, click **Finish**.

At this point your virtual machine is created and you should use the Failover Cluster Manager to start the virtual machine and perform the operating system installation according to the operating system that you choose. For the purpose of this validation, the guest operating system can be any Windows Server version.

6.2 Test network connectivity from the virtual machine

Once you finish installing the operating system in the virtual machine you should log on and verify if this virtual machine was able to obtain IP address from the enterprise network. Assuming that in this network you have a DHCP server, this virtual machine should be able to obtain the IP address. To perform the basic network connectivity test use the following approach.

- Use `ping` command for a reachable IP address in the same subnet.
- Use `ping` command for the same destination but now using the full qualified domain name for the destination host. The goal here is to test basic name resolution.

Note

You may need to open Windows Firewall with Advanced Security and create a new rule to allow Internet Control Message Protocol (ICMP) before performing the previous tests. This may be true for other hosts you want to ping − confirm that the host-based firewall on the target allows for ICMP Echo Requests.

After you confirm that this basic test is working properly, leave a command prompt window open and enter the command ping <Destination_IP_Address_or_FQDN> -t. The goal here is to have a continuous test while you perform the live migration to the second node.

Note

If you prefer to work with PowerShell, instead of the ping command you can use the **Test-Connection** command. This cmdlet provides you a number of connectivity testing options that exceed what is available with the simple ping command.

6.3 Perform a live migration

To perform a live migration of this virtual machine from the current cluster node to the other node in the cluster, perform the following steps.

1. In the **Failover Cluster Manager**, click **Roles** under the cluster name. On the Roles pane, right click the virtual machine that you created, click **Move**, click **Live Migration**, and then click **Select Node**.
2. On the **Move Virtual Machine** page, select the node that you want to move the virtual machine to and click **OK**.

You will notice in the **Status** column when the live migration starts, it will take some time for the **Information** column to update the current state of the migration. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

6.4 Perform a quick migration

To perform the quick migration of this virtual machine from the current node to the other one, perform the following steps.

1. On the **Failover Cluster Manager**, click **Roles** under the cluster name. In the **Roles** pane, right-click the virtual machine that you created, click **Move**, click **Quick Migration** and then click **Select Node**.

2. On the **Move Virtual Machine** window, select the node that you want to move the virtual machine to, and then click **OK**.

You will notice in the status that the quick migration will start faster than the live migration did. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

Building Your Cloud Infrastructure: Converged Data Center with File Server Storage

This document contains the instructions to create a private or public cloud configuration that uses:

- Two subnets - one for cloud infrastructure traffic (live migration, cluster, storage and management) one for tenant traffic.
- NIC Teaming for network bandwidth aggregation and failover for the infrastructure and tenant subnet.
- A dedicated file server storage cluster that hosts virtual machine VHDX and configuration files.
- A dedicated Hyper-V compute cluster that runs the virtual machine workloads.

The design pattern discussed in this document is one of three design patterns we suggest for building the core cloud network, compute and storage infrastructure. For information about the other two cloud infrastructure design patterns, please see:

- Building Your Cloud Infrastructure: Non-Converged Data Center Configuration
- Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes

Design Considerations and Requirements for the Converged Data Center with File Server Storage Pattern

The Converged Data Center with File Server Storage cloud infrastructure design pattern focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You require that cloud infrastructure traffic be physically separated from cloud tenant traffic. The requirement is met by creating separate NIC teams for infrastructure and tenant traffic and connecting them to different subnets/segments.
- You require that infrastructure traffic (i.e. live migration, cluster, storage, management) all receive guaranteed levels of bandwidth. The requirement is met by using Windows QoS policies on the parent Hyper-V partition..

- You require that tenant traffic from different tenants receive guaranteed levels of bandwidth. The requirement is met by using Hyper-V virtual switch QoS policies.
- You require the highest networking performance possible to support hosting Virtual Machines and VHD/X files on a file server using the SMB 3.0 protocol. The requirement is met by installing and enabling Remote Direct Memory Access (Remote DMA or RDMA) network adapters on both the Hyper-V Server cluster and the File server cluster.

Storage

- You require the ability to scale and manage storage separately from the compute infrastructure. This requirement can be met by creating a dedicated storage cluster that will host the VHDX and virtual machine configuration files. The Hyper-V compute cluster will connect to the file server cluster using SMB 3.0 to connect to file shares on the storage cluster. The file server cluster will be configured to use the new Windows Server 2012 Scale-out File Server feature.
- You require a low cost storage option. This requirement is met by using Serial Attached SCSI (SAS) disks in shared JBOD enclosures managed through Storage Spaces. Alternatively, each member of the file server failover cluster can connect to iSCSI, Fibre Channel or Fibre Channel over Ethernet. Only the SAS scenario is described in this document.
- You require a resilient storage solution. This requirement is met by using two file servers configured as a failover cluster, with well-connected (shared JBODs) storage so that all members of the file server failover cluster are directly connected to storage with Storage Spaces configured as a mirrored space to guarantee against data loss in the case of disk failures. In addition, each member of the file server failover cluster is able to access the shared storage by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

- You require the ability to scale your compute infrastructure separately from your storage infrastructure. This requirement is met by creating a dedicated Hyper-V compute cluster that connects to remote file server storage for virtual machines and virtual machine configuration files. Local disks on the compute nodes are only used for the boot partition but not for the virtual machines.
- You require that virtual machines will be continuously available and resilient to hardware failures. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.
- You require the highest number of virtual machines possible per host server (i.e. increased density). This requirement is met by using processor offload technologies, such as Remote Direct Memory Access (RDMA), Receive Side Scaling, Receive Side Coalescing (RSC), and Datacenter Bridging (DCB).

Overview

A Windows Server® 2012 cloud infrastructure described in this document is a high-performing and highly available Hyper-V cluster that hosts virtual machines that can be managed to create private or public clouds connected to a converged 10Gb Ethernet network, and using dedicated file servers as the storage nodes. This document explains how to configure the basic building blocks for such a cloud. It does not cover the System Center or other management software aspects of deployments; the focus is on configuring the core Windows Server computers that are used to build cloud infrastructure.

For background information on creating clouds using Windows Server 2012, see Building Infrastructure as a Service Clouds using Windows Server "8".

This cloud configuration consists of the following:

- Multiple computers in a dedicated Hyper-V compute cluster.

  A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility which is critical in a cloud environment. The Hyper-V cluster is a dedicated compute cluster and does not host storage for virtual machines and virtual machine configuration files.

- Multiple computers in a dedicated Scale-out File Server storage cluster.

  A File Server cluster is created using the Windows Server 2012 Failover Cluster feature. Windows Server 2012 includes a new file server capability known as "Scale-out File Server for applications" that enables you to store virtual machine and virtual machine configuration files in a file share and make these files continuously available. When you separate the file server cluster from the compute cluster you are able to scale compute and storage resources independently.

- A converged networking infrastructure that supports physical segmentation of infrastructure and tenant traffic.

  Each computer in the Hyper-V failover cluster must have at least two network adapters so that one adapter can host the cloud infrastructure traffic and one adapter can support tenant traffic. If resiliency against NIC failures is required, then you can add additional network adapters on each of the networks, and team them using Windows Server 2012 Load Balancing and Failover (LBFO) NIC Teaming. The NICs can be 10 GbE or 1 GbE network adapters. These NICs will be used for live migration, cluster, storage, management (together referred to as "infrastructure" traffic) and tenant traffic.

- The appropriate networking hardware (e.g. Ethernet switches, cables, etc.) to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

Figure 1 provides a high-level view of the scenario layout. Key elements of the configuration include:

- A File Server cluster that hosts the virtual hard disks and virtual machine configuration files
- The File Server cluster is connected only to the infrastructure network

- The File Server cluster is connected to block storage either through HBAs or over a 10 GB Ethernet network (as in the case of iSCSI or Fibre Channel of Ethernet [FCoE])
- A Hyper-V compute cluster that hosts the virtual machine workloads.
- The Hyper-V compute cluster is connected to the datacenter network and the tenant network using teamed network adapters.
- Cloud infrastructure traffic flows to and from the host-based 10GbE network adapter
- Tenant network traffic to and from the virtual machines flows through the Hyper-V virtual switch which is bound to the Tenant Network NIC team.

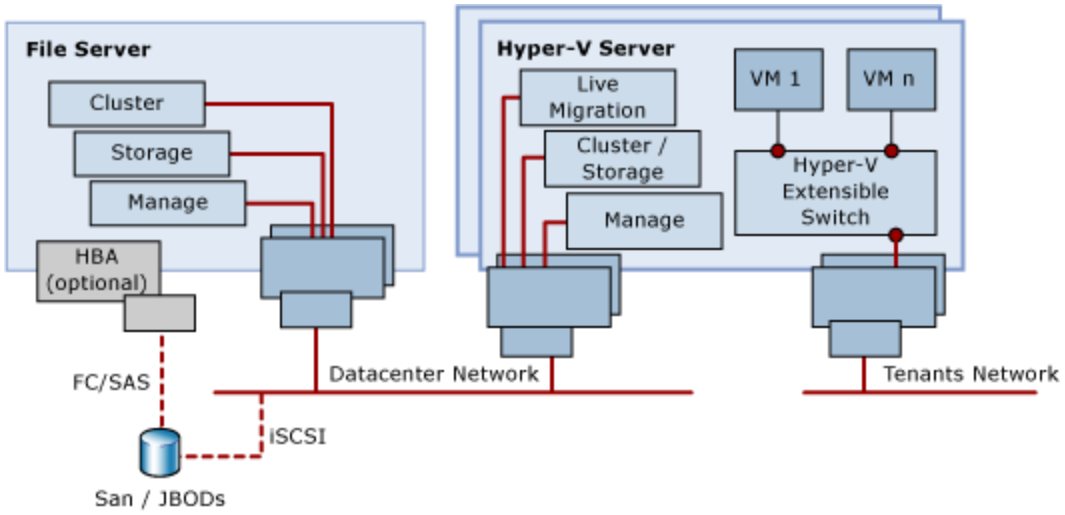

Note

Although this configuration uses SAS storage on the file server cluster, you can easily choose to use other types of storage, such as iSCSI or Fibre Channel-based SAN storage. You can find more information about storage configuration for a non-SAS scenario in the document Building Your Cloud Infrastructure: Non-Converged Enterprise Configuration, which describes how to configure the SAN storage.

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

This configuration highlights the following technologies and features of Windows Server 2012:

- **Load Balancing and Failover (LBFO)**: Load Balancing and Failover logically combines multiple network adapters to provide bandwidth aggregation and traffic failover to prevent connectivity loss in the event of a network component failure. Load Balancing with Failover is also known as NIC Teaming in Windows Server 2012.
- **Windows Server Quality of Service (QoS)**: Windows Server 2012 includes improved QoS features that lets you manage bandwidth and provide predictable network performance to traffic to and from the host operating system.

- **Data Center Bridging (DCB)**: DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control. Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support Internet Small Computer System Interface (iSCSI), Remote Direct Memory Access (RDMA) over Converged Ethernet, or Fiber Channel over Ethernet (FCoE).
- **Storage Spaces**: Storage Spaces makes it possible for you to create cost-effective disk pools that present themselves as a single mass storage location on which virtual disks or volumes can created and formatted.

The following sections describe how to set up this cloud configuration using UI-based tools and Windows PowerShell.

After the cloud is built, you can validate the configuration by doing the following:

- Install and configure virtual machines
- Migrate running virtual machines between servers in the Hyper-V cluster (live migration)

Install and configure the Converged Data Center with File Server Storage cloud infrastructure

In this section we will cover step by step how to configure the cloud infrastructure compute and storage scale units described in this document.

Creating this cloud infrastructure configuration consists of the following steps:

- Step 1: Initial node configuration
- Step 2: Initial network configuration
- Step 3: Initial storage configuration
- Step 4: File Server failover cluster setup
- Step 5: Hyper-V failover cluster setup
- Step 6: Configure Hyper-V settings
- Step 7: Cloud validation

The following table summarizes the steps that this document describes:

Step	Task	Target	Tasks
1	Initial Node Configuration for Compute and Storage Clusters	All Nodes	• 1.1-Enable BIOS settings required for Hyper-V on the nodes in the Hyper-V cluster • 1.2-Perform a clean operating system installation on all nodes in the Hyper-V and File Server clusters • 1.3-Perform post

			installation tasks on all nodes in the Hyper-V and File Server clusters: • Set Windows PowerShell execution policy • Enable Windows PowerShell remoting • Enable Remote Desktop Protocol and Firewall rules • Join the domain • 1.4-Install roles and features using default settings, rebooting as needed on the Hyper-V failover cluster: • Hyper-V (plus management tools) • Data Center Bridging (DCB) • Failover clustering (plus management tools) • 1.5-Install roles and features using default settings, rebooting as needed on the File Server failover cluster: • Failover clustering (plus management tools) • File Server Role • File Sharing and storage management tools
2	Initial Network Configuration	All Nodes	• 2.1-Disable unused and disconnected interfaces and rename active connections • 2.2-Create the infrastructure network

			NIC team and the tenant NIC team on each member of the Hyper-V Cluster and assign IP addressing information. • 2.3-Create the infrastructure NIC team on each member of the file server cluster and assign IP addressing information • 2.4-Configure QoS settings for infrastructure traffic
3	Initial Storage Configuration	Single Node	• 3.1-Present all shared storage to relevant nodes • 3.2-For multipath scenarios, install and configure multipath I/O (MPIO) as necessary • 3.3-All shared disks: Wipe, bring online and initialize
4	File Server Failover Cluster Setup	Single Node	• 4.1-Run through the Cluster Validation Wizard • 4.2-Address any indicated warnings and/or errors • 4.3-Complete the Create Cluster Wizard (setting name and IP but do not add eligible storage) • 4.4-Create the cluster storage pool • 4.5-Create the quorum virtual disk • 4.6-Create the virtual machine storage virtual disk

			• 4.7-Add the virtual machine storage virtual disk to Cluster Shared Volumes • 4.8-Add folders to the cluster shared volume • 4.9-Configure Quorum Settings • 4.10-Add the Scale-Out File Server for Applications Role
5	Hyper-V Failover Cluster Setup	Single Node	• 5.1-Run through the Cluster Validation Wizard • 5.2-Address any indicated warnings and/or errors • 5.3-Complete the Create Cluster Wizard • 5.4-Verify cluster quorum configuration and modify as necessary • 5.5-Configure cluster network metrics
6	File Shares Setup	Single Node	• 6.1-Create Shares and Configure Hyper-V Settings using a Script • 6.2-Configure Kerberos Constrained Delegation
7	Cloud Validation	Single Node	• 7.1-Create the TenantNetSwitch • 7.2-Create a New Virtual machine • 7.3-Test network connectivity from the virtual machine • 7.4-Perform a Live Migration • 7.5-Perform a quick migration

Step 1: Initial node configuration

In step 1, you will perform the following steps on all nodes of the Hyper-V and File Server clusters:

- 1.1 Enable BIOS settings required for Hyper-V on the nodes in the Hyper-V cluster.
- 1.2 Perform a clean operating system installation on all nodes in the Hyper-V and File Server clusters.
- 1.3 Perform post-installation tasks on all nodes in the Hyper-V and File Server clusters.
- 1.4 Install roles and features using the default settings on the Hyper-V failover cluster.
- 1.5 Install roles and features using the default settings on the filer server failover cluster

1.1 Enable BIOS settings required for Hyper-V on the Nodes in the Hyper-V Cluster

You will need to enable virtualization support in the BIOS of each cluster member prior to installing the Hyper-V server role. The procedure for enabling processor virtualization support will vary with your processors' make and model and the system BIOS. Please refer to your hardware documentation for the appropriate procedures. In addition, confirm that all systems have the latest BIOS updates installed.

1.2 Perform a clean operating system installation on all nodes in the Hyper-V and File Server Clusters

Install Windows Server 2012 using the Full Installation option.

1.3 Perform post-installation tasks on all nodes in the Hyper-V and File Server Clusters

There are several tasks you need to complete on each node of the compute and file server clusters after the operating system installation is complete. These include:

- Join each node to the domain
- Enable remote access to each node via the Remote Desktop Protocol.
- Set the Windows PowerShell execution policy.
- Enable Windows PowerShell remoting.

Perform the following steps to join each node to the domain:

1. Press the Windows Key on the keyboard and then press R. Type Control Panel and then click OK.
2. In the Control Panel window, click System and Security, and then click System.
3. In the System window under Computer name, domain, and workgroup settings, click Change settings.
4. In the System Properties dialog box, click Change.
5. Under Member of, click Domain, type the name of the domain, and then click OK.

Run the following Windows PowerShell commands on each node of the compute and file server clusters to enable remote access using the Remote Desktop Protocol (RDP), to enable PowerShell execution policy and enable PowerShell Remoting:

```
(Get-WmiObject Win32_TerminalServiceSetting -Namespace
root\cimv2\terminalservices).SetAllowTsConnections(1,1)

Set-ExecutionPolicy Unrestricted –Force

Enable-PSRemoting –Force
```

1.4 Install roles and features using the default settings on the Hyper-V Failover Cluster

The following roles and features will be installed on *each node* of the Hyper-V compute cluster:

- Hyper-V and Hyper-V management Tools
- Data Center Bridging (DCB)
- Failover cluster and failover cluster management tools

Perform the following steps on *each node* in the cluster to install the required roles and features:

1. In **Server Manager**, click **Dashboard** in the console tree.
2. In **Welcome to Server Manager**, click **2 Add roles and features**, and then click **Next**.
3. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
4. On the **Installation Type** page, click **Next**.
5. On the **Server Selection** page, click **Next**.
6. On the **Server Roles** page, select **Hyper-V** from the **Roles** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Click **Next**.
7. On the **Features** page, select **Data Center Bridging and Failover Clustering** from the **Features** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Click **Next**.

   Note
   If you plan to use Multipath I/O for your storage solution, select the **Multipath I/O** feature while performing step 7.
8. On the **Hyper-V** page, click **Next**.
9. On the **Virtual Switches** page, click **Next**.
10. On the **Migration** page, click **Next**.
11. On the **Default Stores** page, click **Next**.
12. On the **Confirmation** page, put a checkmark in the **Restart the destination server automatically if required** checkbox and then in the **Add Roles and Features** dialog box click **Yes**, then click **Install**.
13. On the **Installation progress** page, click **Close** after the installation has succeeded.
14. Restart the computer. This process might require restarting the computer twice. If so, the installer will trigger the multiple restarts automatically.

After you restart the server, open **Server Manager** and confirm that the installation completed successfully. Click **Close** on the **Installation Progress** page.

1.5 Install roles and features using the default settings on the File Server Failover Cluster

The following roles and features will be installed on each node of the file server failover cluster:

- Failover cluster and failover cluster management tools
- Datacenter Bridging (DCB)
- Storage management tools

Perform the following steps on each node in the file server failover cluster to install the required roles and features:

1. In **Server Manager**, click **Dashboard** in the console tree.
2. In **Welcome to Server Manager**, click **2 Add roles and features**, and then click **Next**.
3. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
4. On the **Installation Type** page, click **Next**.
5. On the **Server Selection** page, click **Next**.
6. On the **Server Roles** page, expand the **File and Storage Services** node, then expand the **File and iSCSI Services** node and select **File Server**. Click **Next**.
7. On the **Features** page, select **Data Center Bridging and Failover Clustering** from the **Features** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Expand **Remote Server Administrator Tools** and then expand **Role Administration Tools. Expand File Services Tools**. Select **Share and Storage Management Tool**. Click **Next**.

   Note
   If you plan to use Multipath I/O for your storage solution, select the **Multipath I/O** feature while performing step 7.
8. On the **Confirmation** page, put a checkmark in the **Restart the destination server automatically if required** checkbox and then in the **Add Roles and Features** dialog box click **Yes**, then click **Install**.
9. On the **Installation progress** page, click **Close** after the installation has succeeded.
10. Restart the computer. This process might require restarting the computer twice. If so, the installer will trigger the multiple restarts automatically.

After you restart the server, open **Server Manager** and confirm that the installation completed successfully. Click **Close** on the **Installation Progress** page.

Step 2: Initial network configuration

The network configuration on each node in both the compute and file server clusters needs to be configured. The networking configuration on the compute cluster needs to be configured to support the converged networking scenario where all traffic, including infrastructure and tenant traffic, moves through the Hyper-V virtual switch. The network configuration on the file server

cluster will connect it to the infrastructure network. You will perform the following procedures to complete the initial networking configuration for the compute and file server clusters:

- 2.1 Disable unused and disconnected interfaces and rename active connections.
- 2.2 Create the infrastructure network NIC team and the tenant network NIC team on each member of the Hyper-V cluster.
- 2.3 Create the infrastructure network NIC team on each member of the file server cluster.
- 2.4 Configure QoS settings for infrastructure traffic

2.1 Disable unused and disconnected interfaces and rename active connections

You can simplify the configuration and avoid errors when running the wizards and running PowerShell commands by disabling all network interfaces that are either unused or disconnected. You can disable these network interfaces in the **Network Connections** window.

For the remaining network adapters for all servers in the compute and storage failover clusters, do the following:

1. Connect them to the appropriate network switch ports.
2. To help you more easily recognize the active network adapters, rename them with names that indicate their use or their connection to the intranet or Internet (for example, **HosterNet1** and **HosterNet2** for the infrastructure network NICs and **TenantNet1** and **TenantNet2** for the tenant network NICs). You can do this in the **Network Connections** window.

2.2 Create the infrastructure and the tenant networks NIC teams on each member of the Hyper-V cluster

Network Load Balancing and Failover (LBFO) enables bandwidth aggregation and network adapter failover to prevent connectivity loss in the event of a network card or port failure. This feature is commonly referred to as "NIC Teaming". In this scenario you will create one team that will be connected to the HosterNet subnet (which is the cloud infrastructure network).

To configure the network adapter teams by using **Server Manager**, do the following on each computer in the Hyper-V compute cluster:

Note

Some steps in the following procedure will temporarily interrupt network connectivity. We recommend that all servers be accessible over a keyboard, video, and mouse (KVM) switch so that you can check on the status of these machines if network connectivity is unavailable for more than five minutes.

1. From **Server Manager**, click **Local Server** in the console tree.
2. In **Properties**, click **Disabled**, which you'll find next to **Network adapter teaming**.
3. In the **NIC Teaming window**, click the name of the server computer in **Servers**.
4. In **Teams**, click **Tasks**, and then click **New Team**.
5. In the **New Team** window, in the **Team Name** text box, enter the name of the network adapter team for the infrastructure traffic subnet (example: **HosterNetTeam**).

6. In the **Member adapters** list select the two network adapters connected to the converged traffic subnet (in this example, **HosterNet1** and **HosterNet2**), and then click **OK**. Note that there may be a delay of several minutes before connectivity is restored after making this change. To ensure that you see the latest state of the configuration, right click your server name in the **Servers** section in the **NIC Teaming** window and click **Refresh Now**. There may be a delay before the connection displays as **Active**. You may need to refresh several times before seeing the status change.
7. Repeat the procedure to create a NIC Team for the tenant network NICs and give it an informative name, such as **TenantNetTeam**.
8. Close the **NIC Teaming** window.
9. Assign static IP addresses to your NIC teams.

Configure static IPv4 addressing information for the new network adapter team connected to the infrastructure and tenant traffic subnets (example: **HosterNetTeam** and **TenantNetTeam**). The IP addresses will be the ones that you will use when connecting to the host system for management purposes. You can configure the IP addressing information in the **Properties** of the team in the **Network Connections** window. You will see new adapters where the names of the teamed network adapters are the names you assigned in step 5. You will lose connectivity for a few moments after assigning the new IP addressing information.

Note

You might need to manually refresh the display of the **NIC Teaming** window to show the new team and there may be a delay in connectivity as the network adapter team is created. If you are managing this server remotely, you might temporarily lose connectivity to the server.

2.3 Create the infrastructure network NIC team on each member of the File Server cluster

You now are ready to team the network adapters on the servers in the file server failover cluster. The file server failover cluster is connected only to the infrastructure network, so you will create a single team on each server in the failover cluster..

To configure the network adapter teams by using **Server Manager**, do the following on *each computer in the file server failover cluster*:

Note

Several steps in the following procedure will temporarily interrupt network connectivity. We recommend that all servers be accessible over a keyboard, video, and mouse (KVM) switch so that you can check on the status of these machines if network connectivity is unavailable for more than five minutes.

1. From **Server Manager**, click **Local Server** in the console tree.
2. In **Properties**, click **Disabled**, which you'll find next to **Network adapter teaming**.
3. In the **NIC Teaming** window, click the name of the server computer in **Servers**.
4. In **Teams**, click **Tasks**, and then click **New Team**.

5. In the **New Team** window, in the **Team Name** text box, enter the name of the network adapter team for the converged traffic subnet (example: **HosterNetTeam**).
6. In the **Member** adapters list select the two network adapters connected to the converged traffic subnet (in this example, **HosterNet1** and **HosterNet2**), and then click **OK**. Note that there may be a delay of several minutes before connectivity is restored after making this change. To ensure that you see the latest state of the configuration, right click your server name in the **Servers** section in the **NIC Teaming** window and click **Refresh Now**. There may be a delay before the connection displays as **Active**. You may need to refresh several times before seeing the status change.
7. Close the **NIC Teaming** window.

Configure a static IPv4 addressing configuration for the new network adapter team connected to the infrastructure traffic subnet (example: **HosterNet Team**). This IP address is the one that you will use when connecting to the host system for management purposes. You can do this in the **Properties** of the team in the **Network Connections** window. You will see a new adapter where the name of the teamed network adapter is the name you assigned in step 5. You will lose connectivity for a few moments after assigning the new IP addressing information.

Note

You might need to manually refresh the display of the **NIC Teaming** window to show the new team and there may be a delay in connectivity as the network adapter team is created. If you are managing this server remotely, you might temporarily lose connectivity to the server.

2.4 Configure QoS settings for infrastructure traffic

If both your 10 GbE network adapter and 10 GbE-capable switch support Data Center Bridging and you want to use DCB for QoS offload you can take advantage of Windows Server 2012 support for Data Center Bridging (DCB).

Whether or not you use DCB, you must create QoS policies to classify and tag each network traffic type. You must use Windows PowerShell commands (**New-NetQosPolicy -Name**) to create new QoS policies for each type of traffic on each computer in the cluster. Here are some example commands (make sure to open all PowerShell windows as administrator):

```
New-NetQosPolicy –Name "Live Migration policy" –LiveMigration –MinBandwidthWeightAction
20

New-NetQosPolicy –Name "SMB policy" –SMB –MinBandwidthWeightAction 50

New-NetQosPolicy –Name "Cluster policy" -IPProtocolMatchCondition UDP -
IPDstPortMatchCondition 3343 –MinBandwidthWeightAction 20

New-NetQosPolicy –Name "Management policy" –DestinationAddress 10.7.124.0/24 –
MinBandwidthWeightAction 10
```

These commands use the **MinBandwidthWeightAction** parameter, which specifies a minimum bandwidth as a relative weighting of the total. The **-LiveMigration** and **-SMB** filters are built in Windows Server 2012. They match packets sent to TCP port 6600 (live migration) and TCP port

445 (SMB protocol used for file storage), respectively. The cluster service traffic uses UDP port 3343. The example for management traffic is the address range 10.7.124.0/24, which corresponds to the Hoster subnet in this example.

As a result, live migration, SMB, cluster, and management traffic will have roughly 20 percent, 50 percent, 20 percent, and 10 percent of the total bandwidth, respectively. To display the resulting traffic classes, run the **Get-NetQosTrafficClass** Windows PowerShell command.

If you are not using DCB, Windows will enforce performance isolation through these QoS policies.

If your network adapters and switch support DCB, enable DCB on the network adapters attached to the converged network subnet using Windows PowerShell. The following are examples:

```
Enable-NetAdapterQos "HosterNet1"

Enable-NetAdapterQos "HosterNet2"
```

Note that **HosterNet1** and **HosterNet2** are the names assigned to the individual network adapters and not the NIC Team names.

To verify the settings on a network adapter, use the **get-netadapterqos** command. The following is example output after running the Windows PowerShell command (note that if QoS is not supported on the adapter you will see an error in the output):

```
Network Adapter Name : HosterNet 1

QOS Enabled : True

MACsec Bypass Supported : False

Pre-IEEE DCBX Supported : True

IEEE DCBX Supported : False

Traffic Classes (TCs) : 8

ETS-Capable TCs : 8

PFC-Enabled TCs : 8

Operational TC Mappings : TC TSA Bandwidth Priorities

                  -- --- --------- ----------

                         0  ETS 20%       0-3,5,7

                         1  ETS 50%       4

                         2  ETS 30%       6
```

For an example of a Windows PowerShell script that configures LBFO and performance isolation settings, see ConfigureNetworking.ps1 which can be downloaded at: http://gallery.technet.microsoft.com/scriptcenter/Windows-Server-2012-Cloud-e0b7753a

Step 3: Initial storage configuration

With the initial cluster node configuration complete, you are ready to perform initial storage configuration tasks on all nodes of the cluster. Initial storage configuration tasks include:

- 3.1 Present all shared storage to relevant nodes.

- 3.2 Install and configure MPIO as necessary for multipath scenarios.
- 3.3 Wipe, bring online, and initialize all shared disks.

3.1 Present all shared storage to relevant nodes

In a SAS scenario, connect the SAS adapters to each storage device. Each cluster node should have two adapters in them if high availability to storage access is required.

3.2 Install and configure MPIO as necessary for multipath scenarios

If you have multiple data paths to storage (for example, two SAS cards) make sure to install the Microsoft® Multipath I/O (MPIO) on each node. This step might require you to restart the system. For more information about MPIO, see What's New in Microsoft Multipath I/O.

3.3 Wipe, bring online, and initialize all shared disks

To prevent issues with the storage configuration procedures that are detailed later is this document; confirm that the disks in your storage solution have not been previously provisioned. The disks should have no partitions or volumes. They should also be initialized so that there is a master book record (MBR) or GUID partition table (GPT) on the disks and then brought online. You can use the **Disk Management** console or Windows PowerShell to accomplish this task. This task must be completed on each node in the cluster.

Note

If you have previously configured these disks with Windows Server 2012 Storage Spaces pools, you will need to delete these storage pools prior to proceeding with the storage configuration described in this document. Please refer to the TechNet Wiki Article "How to Delete Storage Pools and Virtual Disks Using PowerShell".

Step 4: File server failover cluster setup

You are now ready to complete the failover cluster settings for the file server cluster. Failover cluster setup includes the following steps:

- 4.1 Run through the Cluster Validation Wizard.
- 4.2 Address any indicated warnings and/or errors.
- 4.3 Complete the Create Failover Cluster Wizard.
- 4.4 Create the clustered storage pool.
- 4.5 Create the quorum virtual disk.
- 4.6 Create the virtual machine storage virtual disk.
- 4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes.
- 4.8 Add folders to the cluster shared volume.
- 4.9 Configure Quorum Settings.
- 4.10 Add the Scale-Out File Server for Applications role.

4.1 Run through the Cluster Validation Wizard

The Cluster Validation Wizard will query multiple components in the intended cluster hosts and confirm that the hardware and software is ready to support failover clustering. On one of the nodes in the server cluster, perform the following steps to run the Cluster Validation Wizard:

1. In the **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** console, in the **Management** section, click **Validate Configuration**.
3. On the **Before You Begin** page of the **Validate a Configuration Wizard**, click **Next**.
4. On the **Select Servers or a Cluster** page, type the name of the local server, and then click **Add**. After the name appears in the **Selected servers** list, type the name of another Hyper-V cluster member computer, and then click **Add**. Repeat this step for all computers in the Hyper-V cluster. When all of the servers of the Hyper-V cluster appear in the **Selected servers** list, click **Next**.
5. On the **Testing Options** page, click **Next**.
6. On the **Confirmation** page, click **Next**. The time to complete the validation process will vary with the number of nodes in the cluster and can take some time to complete.
7. On the **Summary** page, the summary text will indicate that the configuration is suitable for clustering. Confirm that there is a checkmark in the **Create the cluster now using the validated nodes...** checkbox.

4.2 Address any indicated warnings and/or errors

Click the **Reports** button to see the results of the Cluster Validation. Address any issues that have led to cluster validation failure. After correcting the problems, run the Cluster Validation Wizard again. After the cluster passes validation, then proceed to the next step. Note that you may see errors regarding disk storage. You may see this if you haven't yet initialized the disks. Click **Finish**.

4.3 Complete the Create Failover Cluster Wizard

After passing cluster validation, you are ready to complete the cluster configuration.

Perform the following steps to complete the cluster configuration:

1. On the **Before You Begin** page of the **Create Cluster Wizard**, click **Next**.
2. On the **Access Point for Administering the Cluster** page, enter a valid NetBIOS name for the cluster, and then select the network you want the cluster on and then type in a static IP address for the cluster, and then click **Next**. In this example, the network you would select is the Management Network. Unselect all other networks that appear here.
3. On the **Confirmation** page, clear **Add all eligible storage to the cluster** checkbox and then click **Next**.
4. On the **Creating New Cluster** page you will see a progress bar as the cluster is created.
5. On the **Summary** page, click **Finish**.
6. In the console tree of the **Failover Cluster Manager** snap-in, open the **Networks** node under the cluster name.

7. Right-click the cluster network that corresponds to the management network adapter network ID (subnet), and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is not selected and that **Allow clients to connect through this network** is enabled. In the **Name** text box, enter a friendly name for this network (for example, **ManagmentNet**), and then click **OK**.
8. Right-click the cluster network that corresponds to the Cluster network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is not enabled. In the **Name** text box, enter a friendly name for this network (for example, **ClusterNet**), and then click **OK**.
9. Right-click the cluster network that corresponds to the live migration network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is not enabled. In the **Name** text box, enter a friendly name for this network (for example, **LiveMigrationNet**), and then click **OK**.

4.4 Create a cluster storage pool

Perform the following steps on one of the members of the cluster to create the storage pool:

1. In the left pane of the **Failover Cluster Manager**, expand the server name and then expand the **Storage** node. Click **Storage Pools**.
2. In the **Actions** pane, click **New Storage Pool**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Storage Pool Name** page, enter a name for the storage pool in the **Name** text box. Enter an optional description for the storage pool in the **Description** text box. In the **Select the group of available disks (also known as a primordial pool) that you want to use** list, select the name you assigned to the cluster (this is the NetBIOS name you assigned to the cluster when you created the cluster). Click **Next**.
5. On the **Physical Drives** page, select the drives that you want to participate in the storage pool. Then click **Next**.
6. On the **Confirmation** page, confirm the settings and click **Create**.
7. On the **Results** page, you should receive the message **You have successfully completed the New Storage Pool Wizard**. Remove the checkmark from the **Create a virtual disk when the wizard closes** checkbox. Then click **Close**.

4.5 Create the quorum virtual disk

Now that you have created the storage pool, you can create virtual disks within that storage pool. A virtual disk is sometimes called a logical unit number or LUN and it represents a collection of one or more physical disks from the previously created storage pool. The layout of data across the physical discs can increase the reliability and performance of the physical disk.

You will need to create at least two virtual disks:

- A virtual disk that can be used as a quorum witness disks. This disk can be configured as a 1 GB virtual disk.

- A virtual disk that will be assigned to a cluster shared volume.

Perform the following steps to create the quorum virtual disk:

1. In the **Failover Cluster Manager** console, expand the **Storage** node in the left pane of the console. Right click **Pools** and click **Add Disk**.
2. In the **New Virtual Disk Wizard** on the **Before You Begin** page, click **Next**.
3. On the **Storage Pool** page, select your server name in the **Server** section and then select the storage pool you created earlier in the **Storage pool** section. Click **Next**.
4. On the **Virtual Disk Name** page, enter a name for the virtual disk in the **Name** text box. You can also enter an optional description in the **Description** text box. Click **Next**.
5. On the **Storage Layout** page, in the **Layout** section, select **Mirror**. Click **Next**.
6. On the **Resiliency Settings** select **Two-way mirror** and click **Next**.
7. On the **Size** page, in the **Virtual disk size** text box, enter a size for the new virtual disk, which in this example will be 1 GB. Use the drop down box to select GB. Also, you can put a checkmark in the **Create the largest virtual disk possible, up to the specified size** checkbox, but this is not required or desired when creating a witness disk. When this option is selected it allows the wizard to calculate the largest size virtual disk you can create given the disks you have assigned to the pool, regardless of the number you put in the **Virtual disk size** text box. Click **Next**.
8. On the **Confirmation** page, review your settings and click **Create**.
9. On the **Results** page, put a checkmark in the **Create a volume when this wizard closes** checkbox. Click **Close**.
10. On the **Before You Begin** page of the **New Volume Wizard**, click **Next**.
11. On the **Server and Disk** page, select the name of the cluster from the **Server** list. In the **Disk** section, select the virtual disk you just created. You can identify this disk by looking in the **Virtual Disk** column, where you will see the name of the virtual disk you created. Click **Next**.
12. On the **Size** page, accept the default volume size, and click **Next**.
13. On the **Drive Letter or Folder** page, select the **Drive letter** and select a drive letter. Click **Next**.
14. On the **File System Settings** page, from the **File system** drop down list, select **NTFS**. Use the default setting in the **Allocation unit size** list. Click **Next**.
15. On the **Confirmation** page, click **Create**.
16. On the **Results** page, click **Close**.

4.6 Create the virtual machine storage virtual disk

Perform the following steps to create the **virtual machine storage** virtual disk:

1. In the **Failover Cluster Manager** console, expand the **Storage** node in the left pane of the console. Right click **Pools** and click **Add Disk**.
2. In the **New Virtual Disk Wizard** on the **Before You Begin** page, click **Next**.
3. On the **Storage Pool** page, select your server name in the **Server** section and then select the storage pool you created earlier in the **Storage pool** section. Click **Next**.

4. On the **Virtual Disk Name** page, enter a name for the virtual disk in the **Name** text box. You can also enter an optional description in the **Description** text box. Click **Next**.
5. On the **Storage Layout** page, in the **Layout** section, select **Mirror**. Click **Next**.
6. On the **Resiliency Settings** select **Two-way mirror** and click **Next**.
7. On the **Size** page, in the **Virtual disk size** text box, enter a size for the new virtual disk. Use the drop down box to select MB, GB or TB. Also, you can put a checkmark in the **Create the largest virtual disk possible, up to the specified size** checkbox. When this option is selected it allows the wizard to calculate the largest size virtual disk you can create given the disks you have assigned to the pool, regardless of the number you put in the **Virtual disk size** text box. Click **Next**.
8. On the **Confirmation** page, review your settings and click **Create**.
9. On the **Results** page, put a checkmark in the **Create a volume when this wizard closes** checkbox. Click **Close**.
10. On the **Before You Begin** page of the **New Volume Wizard**, click **Next**.
11. On the **Server and Disk** page, select the name of the cluster from the **Server** list. In the **Disk** section, select the virtual disk you just created. You can identify this disk by looking in the **Virtual Disk** column, where you will see the name of the virtual disk you created. Click **Next**.
12. On the **Size** page, accept the default volume size, and click **Next**.
13. On the **Drive Letter or Folder** page, select the **Don't Assign to a drive letter or folder** and select a drive letter. Click **Next**.
14. On the **File System Settings** page, from **the File system** drop down list, select **NTFS**. Use the default setting in the Allocation unit size list. Note that **ReFS** is not supported in a Cluster Shared Volume configuration. Click **Next**.
15. On the **Confirmation** page, click **Create**.
16. On the **Results** page, click **Close**.

4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes

The virtual disk you created for virtual machine storage is now ready to be added to a Cluster Shared Volume. Perform the following steps to add the virtual disk to a Cluster Shared Volume.

1. In the **Failover Cluster Manager**, in the left pane of the console, expand the **Storage** node and click **Disks**. In the middle pane of the console, in the **Disks** section, right click the virtual disk you created in the previous step and then click **Add to Cluster Shared Volumes**.
2. Proceed to the next step.

4.8 Add folders to the cluster shared volume

Now you need to create the folders on the virtual disk located on the Cluster Shared Volume to store the virtual machine files and the virtual machine data files.

Perform the following steps to create a file share to store the running VMs of the Hyper-V cluster:

1. Open **Windows Explorer** and navigate to the C: drive and then double-click **Cluster Storage** and then double-click **Volume 1**.

2. Create two folders in **Volume 1**. One of the folders will contain the .vhd files for the virtual machines (for example, VHDdisks) and one folder will contain the virtual machine configuration files (for example, VHDsettings)

4.9 Configure Quorum Settings

Perform the following steps to configure quorum settings for the cluster:

1. In the left pane of the **Failover Cluster Manager** console, right click on the name of the cluster and click **More Actions** and click **Configure Cluster Quorum Settings**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Quorum Configuration Option** page, select **Use typical settings (recommended)** and click **Next**.
4. On the **Confirmation** page, click **Next**.

4.10 Add the Scale-Out File Server for Applications Role

The file server failover cluster will provide continuous availability of the virtual machine files to the nodes in the compute cluster. Windows Server 2012 includes the Scale Out File Server for Applications feature the enables the level of continuous availability required for hosting virtual machine files.

Perform the following steps to configure the file server failover cluster as a Scale-Out File Server for Applications:

1. In the **Failover Cluster Management** console, right click on **Roles** and click **Configure Role**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Select Role** page, select **File Server** and click **Next**.
4. On the **File Server Type** page, select **Scale-Out File Server for application data** and click **Next**.
5. On the **Client Access Point** page, enter a NetBIOS name for the client access point and then click **Next**.
6. On the **Confirmation** page, click **Next**.
7. On the **Summary** page, click **Finish**.

Step 5: Hyper-V Failover Cluster Setup

To setup the Hyper-V compute failover cluster you will need to:

- 5.1 Run through the cluster validation wizard
- 5.2 Address any indicated warnings and/or errors
- 5.3 Complete the Create Cluster Wizard
- 5.4 Verify cluster quorum configuration and modify as necessary
- 5.5 Configure Cluster Networks

5.1 Run through the cluster validation wizard

The Cluster Validation Wizard will query multiple components in the intended cluster hosts and confirm that the hardware and software is ready to support failover clustering. On one of the nodes in the server cluster, perform the following steps on one of the members of the Hyper-V failover cluster to run the Cluster Validation Wizard:

1. In the **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** console, in the **Management** section, click **Validate Configuration**.
3. On the **Before You Begin** page of the **Validate a Configuration Wizard**, click **Next**.
4. On the **Select Servers or a Cluster** page, type the name of the local server, and then click **Add**. After the name appears in the **Selected servers** list, type the name of another Hyper-V cluster member computer, and then click **Add**. Repeat this step for all computers in the Hyper-V cluster. When all of the servers of the Hyper-V cluster appear in the **Selected servers** list, click **Next**.
5. On the **Testing Options** page, click **Next**.
6. On the **Confirmation** page, click **Next**. The time to complete the validation process will vary with the number of nodes in the cluster and can take some time to complete.

On the **Summary** page, the summary text will indicate that the configuration is suitable for clustering. Confirm that there is a checkmark in the **Create the cluster now using the validated nodes...** checkbox.

5.2 Address any indicated warnings and/or errors

Click the **Reports** button to see the results of the Cluster Validation. Address any issues that have led to cluster validation failure. After correcting the problems, run the Cluster Validation Wizard again. After the cluster passes validation, then proceed to the next step. Note that you may see errors regarding disk storage. You may see this if you haven't yet initialized the disks. Click **Finish**.

5.3 Complete the create cluster wizard

After passing cluster validation, you are ready to complete the cluster configuration.

Perform the following steps to complete the cluster configuration:

1. On the **Before You Begin** page of the **Create Cluster Wizard**, click **Next**.
2. On the **Access Point for Administering the Cluster** page, enter a valid NetBIOS name for the cluster, and then select the network you want the cluster on and then type in a static IP address for the cluster, and then click **Next**. The network you will use in this scenario is the infrastructure network.
3. On the **Confirmation** page, clear **Add all eligible storage to the cluster** checkbox and then click **Next**.
4. On the **Creating New Cluster** page you will see a progress bar as the cluster is created.
5. On the **Summary** page, click **Finish**.

5.4 Verify cluster quorum configuration and modify as necessary

In most situations, use the quorum configuration that the cluster software identifies as appropriate for your cluster. Change the quorum configuration only if you have determined that the change is appropriate for your cluster. For more information about quorum configuration see Understanding Quorum Configurations in a Failover Cluster.

In most cases the Hyper-V compute cluster will not have shared storage. Because of this, you may want to set up a file share witness on the file server cluster or another location.

Perform the following steps to configure the quorum model to use the file share witness:

1. In the **Failover Cluster Manager**, in the left pane of the console, right click the cluster name, point to **More Actions** and click **Configure Cluster Quorum Settings**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Select Quorum Configuration Option** page, select **Add** or change the quorum witness. Click **Next**.
4. On the **Select Quorum Witness** page, select **Configure a file share witness (recommended for special configurations)** and click **Next**.
5. On the **Configure File Share Witness** page, enter the path to the file share witness and click **Next**.
6. On the **Confirmation** page, click **Next**.
7. On the **Configure Cluster Quorum Settings** page, click **Next**.
8. On the **Summary** page, click **Finish**.

5.5 Configure Cluster Networks

Perform the following steps to complete the compute cluster by configuring the cluster network settings:

1. In the console tree of the **Failover Cluster Manager** snap-in, open the **Networks** node under the cluster name.
2. Right-click the cluster network that corresponds to the TenantNet network adapter network ID (subnet), and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is *not* selected and that *Do not allow cluster network communication on this network* is enabled. In the **Name** text box, enter a friendly name for this network (for example, TenantNet), and then click **OK**.
3. Right-click the cluster network that corresponds to the infrastructure network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is enabled. In the **Name** text box, enter a friendly name for this network (for example, **HosterNet**), and then click **OK**.

Step 6: Configure Share and Hyper-V settings using a Script

To finalize the Hyper-V configuration, you will need to:

- 6.1 Create Shares and Configure Hyper-V Settings using a Script
- 6.2 Configure Kerberos Constrained Delegation

6.1 Create Shares and Configure Hyper-V Settings using a Script

You will need shares available that will store the virtual machine disk files and the virtual machine configuration files. There are a number of ways that this can be done. In the example used in this document, you will use a script that will:

1. Create shares that will contain the virtual machine disk and configuration files on the CSV in the file server cluster
2. Configure the correct share and NTFS permissions on those folders
3. Configure the servers in the Hyper-V compute cluster to use the correct share locations as default virtual machine disk and configuration file location.

The following script will perform these actions for you.

```
Insert script here (when we get it)
```

6.2 Configure Kerberos Constrained Delegation

In order to fully manage the storage on the file server failover cluster from a machine in the Hyper-V cluster, you will need to configure Kerberos Constrained Delegation. For details on why this is required please see the article, Using Constrained Delegation to remotely manage a server running Hyper-V that uses CIFS/SMB file shares.

Perform the following steps to configure the members of the compute cluster to be trusted for Kerberos Constrained Delegation:

1. On a domain controller responsible for your compute and file server cluster environment, open **Control Panel** and then open **Administrative Tools**. Open **Active Directory Users and Computers**.
2. In the **Active Directory Users and Computers** console, expand the domain name and then click on **Computers**.
3. In the right pane of the console, right click on the name of one of the computers in the Hyper-V compute failover cluster and then click **Properties**.
4. In the **Properties** dialog box for the member of the Hyper-V compute failover cluster member, click on the **Delegation** tab.
5. On the **Delegation** tab, select the **Trust this computer for delegation to the specified services only**. Then select **Use Kerberos only**. Click **Add**.
6. In the **Add Services** dialog box, in the **Service Type** column, find the **cifs** entries. Click on the entry that corresponds to the name of the file server failover cluster and click **OK**.
7. Click **OK** in the server's **Properties** dialog box.

Step 7: Cloud validation

To verify the configuration of your cloud environment, perform the following operations.

- 7.1 Create the TenantNetSwitch.
- 7.2 Create a new virtual machine.
- 7.3 Test network connectivity from the virtual machine.
- 7.4 Perform a live migration.

- 7.5 Perform a quick migration

7.1 Create the TenentNetSwitch

Before you create a virtual machine, you will need to create a virtual switch that is connected to the TenantNet so that the virtual machine can connect to the network. Perform the following steps to create the TenantNet virtual switch:

1. Open the **Hyper-V Manager** console. In the **Hyper-V Manager** console, in the **Actions** pane, click **Virtual Switch Manager**.
2. In the right pane of the **Virtual Switch Manager** dialog box, select **External** and then click **Create Virtual Switch**.
3. In the **Virtual Switch Properties** section of the dialog box, enter a name for the virtual switch in the **Name** text box (in this example it will be **TenantNetSwitch**). In the **Connection type** section, select the **External network** option. Then select the NIC team representing the TenantNet from the drop down box. Click **OK**.

7.2 Create a new virtual machine

To create a new virtual machine in the cluster environment, perform the following steps.

1. Open **Failover Cluster Manager**, click **Roles** under the cluster name, click **Virtual Machines** under the **Actions** pane, and then click **New Virtual Machine**.
2. On the **New Virtual Machine** page, select the cluster node where you want to create the virtual machine, and then click **OK**.
3. On the **Before you Begin** page of the **New Virtual Machine Wizard**, click **Next**.
4. On the **Specify Name and Location** page, enter a friendly name for this virtual machine and then click **Next**.
5. On the **Assign Memory** page, enter the amount of memory that will be used for this virtual machine (minimum for this lab is 1024 MB RAM) and then click **Next**.
6. On the **Configuring Networking** page, click **Next**.
7. On the **Connect Virtual Hard Disk** page, leave the default options selected and click **Next**.
8. On the **Installation Options** page, select **Install an operating system from a boot CD/DVD-ROM** and then select the location where the CD/DVD is located. If you are installing the new operating system based on an ISO file, make sure to select the option Image file (.iso) and browse for the file location. If you prefer to PXE boot, that option will be described in later steps. After you select the appropriate option for your scenario, click **Next**.
9. On the **Completing the New Virtual Machine Wizard** page, review the options, and then click **Finish**.
10. The virtual machine creation process starts. After it is finished, you will see the **Summary** page, where you can access the report created by the wizard. If the virtual machine was created successfully, click **Finish**.
11. If you want to PXE boot the virtual machine, you will need to create a **Legacy Network Adapter**. Right click the new virtual machine and click **Settings**.
12. In the **Settings** dialog box, select the **Legacy Network Adapter** option and click **Add**.

13. In the **Legacy Network Adapter** dialog box, connect it to the virtual switch (such as TenantNetSwitch) and enable virtual LAN identification and assign the appropriate network identifier.

Note that if the virtual machine continues to use the legacy network adapter it will not be able to leverage many of the features available in the Hyper-V virtual switch. You may want to replace the legacy network adapter after the operating system is installed.

At this point your virtual machine is created and you should use the Failover Cluster Manager to start the virtual machine and perform the operating system installation according to the operating system that you choose. For the purpose of this validation, the guest operating system can be any Windows Server version.

7.3 Test network connectivity from the virtual machine

Once you finish installing the operating system in the virtual machine you should log on and verify if this virtual machine was able to obtain IP address from the enterprise network. Assuming that in this network you have a DHCP server, this virtual machine should be able to obtain the IP address. To perform the basic network connectivity test use the following approach.

- Use `ping` command for a reachable IP address in the same subnet.
- Use `ping` command for the same destination but now using the full qualified domain name for the destination host. The goal here is to test basic name resolution.

Note

If you installed Windows 8 Developer Preview in this virtual machine you need to open Windows Firewall with Advanced Security and create a new rule to allow Internet Control Message Protocol (ICMP) before performing the previous tests. This may be true for other hosts you want to ping − confirm that the host-based firewall on the target allows for ICMP Echo Requests.

After you confirm that this basic test is working properly, leave a command prompt window open and enter the command ping <Destination_IP_Address_or_FQDN> -t. The goal here is to have a continuous test while you perform the live migration to the second node.

Note

If you prefer to work with PowerShell, instead of the ping command you can use the Test-Connection command. This cmdlet provides you a number of connectivity testing options that exceed what is available with the simple ping command.

7.4 Perform a live migration

To perform a live migration of this virtual machine from the current cluster node to the other node in the cluster, perform the following steps.

1. In the **Failover Cluster Manager**, click **Roles** under the cluster name. On the **Roles** pane, right click the virtual machine that you created, click **Move**, click **Live Migration**, and then click **Select Node**.

2. On the **Move Virtual Machine** page, select the node that you want to move the virtual machine to and click **OK**.

You will notice in the **Status** column when the live migration starts, it will take some time for the **Information** column to update the current state of the migration. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

7.5 Perform a quick migration

To perform the quick migration of this virtual machine from the current node to the other one, perform the following steps.

1. On the **Failover Cluster Manager**, click **Roles** under the cluster name. In the **Roles** pane, right-click the virtual machine that you created, click **Move**, click **Quick Migration** and then click **Select Node**.
2. On the **Move Virtual Machine** window, select the node that you want to move the virtual machine to, and then click **OK**.

You will notice in the status that the quick migration will start faster than the live migration did. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

Building Your Cloud Infrastructure: Converged Data Center without Dedicated Storage Nodes

This document contains the instructions that you need to follow to create a private or public cloud configuration that uses:

- A converged network infrastructure for live migration, cluster, storage, management, and tenant traffic
- All network traffic moves through the Hyper-V virtual switch
- Hyper-V Virtual Switch Quality of Service (QoS)
- Hyper-V Virtual Switch port ACLs and 802.1q VLAN tagging
- NIC Teaming for network bandwidth aggregation and failover
- Well-connected storage using SAS JBOD enclosures

The design pattern discussed in this document is one of three design patterns we suggest for building the core cloud network, compute and storage infrastructure. For information about the other two cloud infrastructure design patterns, please see:

- Building Your Cloud Infrastructure: Non-Converged Data Center Configuration
- Building Your Cloud Infrastructure: Converged Data Center with File Server Storage

Design Considerations and Requirements for the Converged Data Center without Dedicated Storage Node Pattern

The Converged Data Center without Dedicated Storage Nodes cloud infrastructure design patterns focuses on the following key requirements in the areas of networking, compute and storage:

Networking

- You prefer that network traffic to and from both the host operating system and the guest operating systems running on the host move through a single network adapter team. This requirement is met by using Windows Server 2012 NIC Teaming (LBFO) and passing all traffic through the Hyper-V virtual switch.
- You require that live migration, cluster, storage, management and tenant traffic all receive guaranteed levels of bandwidth. The requirement is met by using Hyper-V virtual switch QoS policies.
- You require that infrastructure traffic (which includes Live Migration, cluster, storage and management traffic) and tenant traffic be isolated from each other. This requirement is met by using Hyper-V virtual switch port ACLs and 802.1q VLAN tagging.

Storage

- You prefer to scale your cloud infrastructure by adding scale units consisting of compute and storage capacity together. This requirement is met by connecting the Hyper-V servers directly to SAS storage, without having dedicated file servers.
- You require cost-effective storage. This requirement is met by using SAS disks in shared JBOD enclosures managed through Storage Spaces.
- You require a resilient storage solution. This requirement is met by having multiple Hyper-V servers configured as a failover cluster, and having a well-connected (shared JBODs) storage so that all members of the failover cluster are directly connected to storage, and by having Storage Spaces configured in as a mirrored space to guarantee against data loss in the case of disk failures
- You require that each member of the Hyper-V failover cluster be able to access the shared storage where the VHDs are located. This requirement is met by using Windows Server 2012 Failover Clustering and Cluster Shared Volumes Version 2 (CSV v2) volumes to store virtual machine files and metadata.

Compute

- You require that the virtual machines will be continuously available and resilient to hardware failures. This requirement can be met by using Windows Server 2012 Failover Clustering together with the Hyper-V Server Role.

- You require the highest number of virtual machines possible per host server (i.e. increased density). This requirement is met by using processor offload technologies, such as Remote Direct Memory Access (RDMA), Receive Side Scaling, Receive Side Coalescing (RSC), and Datacenter Bridging (DCB). Please note that in the default configuration presented here (without a dedicated storage access NIC), RDMA and DCB cannot be used because these technologies require direct access to the hardware and must bypass much of the virtual networking stack. This is similar to the situation with Single Root I/O Virtualization (SR-IOV). For optimal performance, especially in the context of network access to storage, a separate NIC team would be required to support these hardware offload acceleration technologies.

Overview

A Windows Server® 2012 cloud infrastructure is a high-performing and highly available Hyper-V cluster that hosts virtual machines that can be managed to create private or public clouds using the Converged Data Center without Dedicated Storage Nodes infrastructure design pattern. This document explains how to configure the basic building blocks for such a cloud. It does not cover the System Center or other management software aspects of deployments; the focus is on configuring the core Windows Server hosts that are used to build cloud infrastructure.

For background information on creating clouds using Windows Server 2012, see Building Infrastructure as a Service Clouds using Windows Server "8".

This cloud configuration consists of the following:

- Multiple computers in a Hyper-V failover cluster.

  A Hyper-V cluster is created using the Windows Server 2012 Failover Cluster feature. The Windows Server 2012 Failover Clustering feature set is tightly integrated with the Hyper-V server role and enables a high level of availability from a compute and networking perspective. In addition, Windows Server 2012 Failover Clustering enhances virtual machine mobility which is critical in a cloud environment. For example, Live Migration is enhanced when performed in a failover cluster deployment because the cluster can automatically evaluate which node in the cluster is optimal for migrated virtual machine placement.
- A converged networking infrastructure that supports multiple cloud traffic profiles.

  Each computer in the Hyper-V failover cluster should have at least two network adapters that will be used for the converged network. This converged network will host all traffic to and from the server, which includes both host system traffic and guest/tenant traffic. The network adapters will be teamed by using Windows Server 2012 Load Balancing and Failover (LBFO) NIC Teaming. The NICs can be either two or more 10 GbE or 1 GbE network adapters. These NICs will be used for live migration, cluster, storage, management (together referred to as "infrastructure" traffic) and tenant traffic.
- The appropriate networking hardware to connect all of the computers in the Hyper-V cluster to each other and to a larger network from which the hosted virtual machines are available.

The following figure provides a high-level view of the scenario architecture. The teamed network adapters on each member of the failover cluster are connected to what will be referred to as a converged subnet in this document. We use the term converged subnet to make it clear that all traffic to and from the Hyper-V cluster members and the tenant virtual machines on each cluster

member must flow through the teamed converged subnet network adapter. Both the host operating system and the tenants connect to the network through the Hyper-V virtual switch. The figure also shows an optional network adapter that is RDMA-capable that can be used for storage traffic, such as when storage is being hosted on a share on a remote file server. This document does not discuss this optional configuration option. For more information about this storage option, please see the document on the Converged Data Center with File Server Storage design pattern at http://technet.microsoft.com/en-us/library/hh831738.

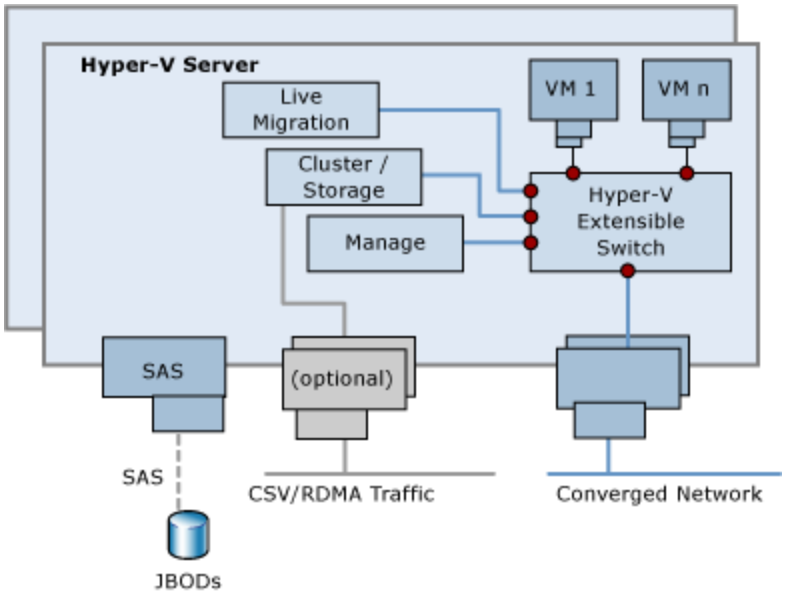

Figure 1 High level overview of cluster member networking configuration

Note

At least one Active Directory Domain Services (AD DS) domain controller is needed for centralized security and management of the cluster member computers (not shown). It must be reachable by all of the cluster member computers, including the members of the shared storage cluster. DNS services are also required and are not depicted.

Figure 2 provides an overview of traffic flows on each member of the Hyper-V cluster. The figure calls out the following significant issues in the configuration:

- Each cluster node member uses a virtual network adapter to connect to the Hyper-V Extensible Switch, which connects it to the physical network.
- Each tenant virtual machine is also connected to the Hyper-V Extensible Switch using a virtual network adapter.
- Network adapters named **ConvergedNet1** and **ConvergedNet2** participate in a teamed physical network adapter configuration using the Windows Server 2012 Failover and Load Balancing feature.
- Windows Server 2012 Hyper-V virtual switch QoS is used to assure that each traffic type (such as live migration, cluster, management and tenant) has a predictable amount of bandwidth available.

- Traffic isolation is enabled by 802.1q VLAN tagging so that host traffic is not visible to the tenants.
- Windows Server 2012 Hyper-V virtual switch port ACLs can also be used for more granular access control at the network level.

It is important to note that Remote Direct Memory Access (RDMA) cannot be used on the converged network because it does not work together with the Hyper-V virtual switch. This will be an issue if you prefer to use high performance SMB 3 connectivity to file server based storage for virtual machine disk and configuration files. In the file server storage scenario, you can introduce addition RDMA capable adapters to connect to storage.

Note

Virtual local area networks (VLANs) are not assigned to each tenant because VLAN-based network isolation is not a scalable solution and is not compatible with Windows Server 2012 network virtualization. VLANs are used to isolate infrastructure traffic from tenant traffic in this scenario.

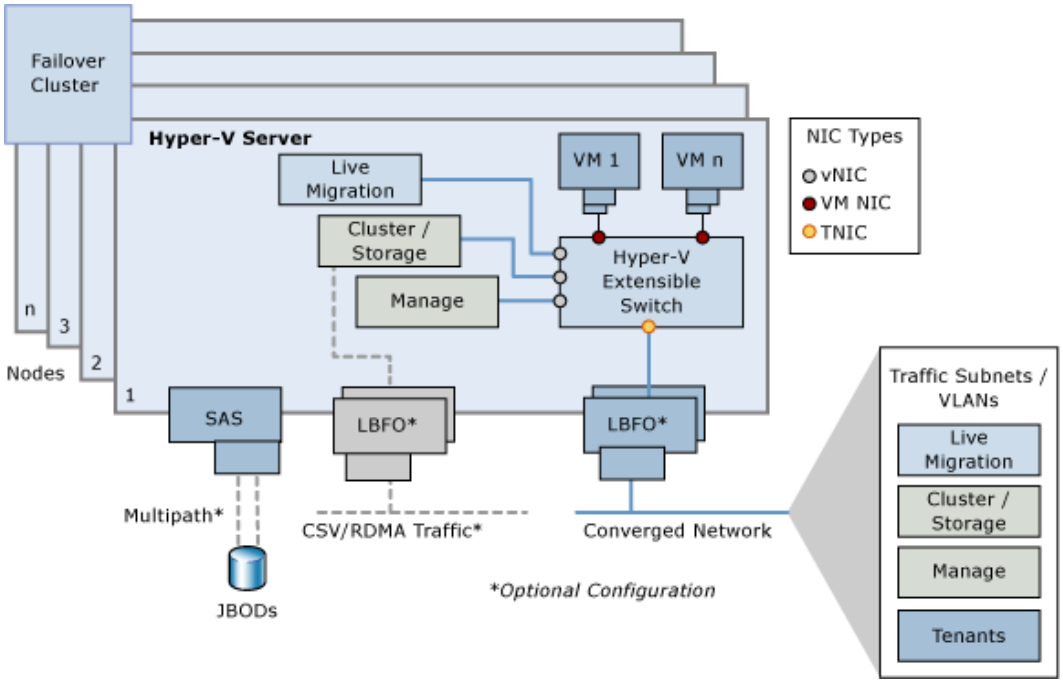

Figure 2 Overview of cluster member traffic flows

This configuration highlights the following technologies and features of Windows Server 2012:

- **Load Balancing and Failover (LBFO)**: Load Balancing and Failover logically combines multiple network adapters to provide bandwidth aggregation and traffic failover to prevent connectivity loss in the event of a network component failure. Load Balancing with Failover is also known as NIC Teaming in Windows Server 2012.
- **Hyper-V Virtual Switch Quality of Service (QoS)**: In Windows Server 2012, QoS includes new bandwidth management features that let you provide predictable network performance to virtual machines on a server running Hyper-V.

- **Hyper-V Virtual Switch Quality of Service (QoS)**: In Windows Server 2012 the Hyper-V virtual switch includes new capabilities that enhance the security of the cloud infrastructure. You can now use Port Access Control Lists (Port ACLs) and VLAN support to get network isolation similar to what you find when using physical network isolation.
- **Storage Spaces**: Storage Spaces makes it possible for you to create cost-effective disk pools that present themselves as a single mass storage location on which virtual disks or volumes can created and formatted.

Note

Although this configuration uses local SAS storage to meet the cost-effective storage requirement, you can easily choose to use other types of storage, such as SAN storage. You can find more information about storage configuration for a non-SAS scenario in the document Building Your Cloud Infrastructure: Non-Converged Enterprise Configuration, which describes how to configure the SAN storage.

The following sections describe how to set up this cloud configuration using UI-based tools and Windows PowerShell.

After the cloud is built, you can validate the configuration by doing the following:

- Install and configure virtual machines
- Migrate running virtual machines between servers in the Hyper-V cluster (live migration)

Install and configure the Converged Data Center without Dedicated Storage Server cloud infrastructure

In this section, we will cover the step by step of how to configure the cloud infrastructure scale unit described in this document.

Creating this cloud infrastructure configuration consists of the following steps:

- Step 1: Initial node configuration
- Step 2: Initial network configuration
- Step 3: Initial storage configuration
- Step 4: Failover cluster setup
- Step 5: Configure Hyper-V settings
- Step 6: Cloud validation

The following table summarizes the steps that this document describes:

Step	Task	Target	Tasks
1	Initial Node Configuration	All Nodes	• 1.1-Add appropriate VLANs to interface ports on the physical switch for each traffic type:

			• Management (untagged, default) • Tenants (tagged) • Live migration (tagged) • Cluster/cluster shared volumes (CSV) (tagged) • 1.2-Enable BIOS settings required for Hyper-V • 1.3-Perform a clean operating system installation • 1.4-Perform post installation tasks: • Set Windows PowerShell execution policy • Enable Windows PowerShell remoting • Enable Remote Desktop Protocol and Firewall rule • Join the domain • 1.5-Install roles and features using default settings, rebooting as needed • Hyper-V (plus management tools) • Storage Services • Failover clustering (plus management tools) • File Sharing and storage management tools
2	Initial Network Configuration	All Nodes	• 2.1-Disable unused and disconnected interfaces and rename active connections • 2.2-Create the converged network adapter team (rename as necessary) and

			assign IP addresses or configure DHCP as appropriate.
			• 2.3-Create the Hyper-V vSwitch and Management virtual network adapter (PowerShell)
			• 2.4-Rename Management virtual network adapter (optional)
			• 2.5-Create additional virtual network adapters and assign VLAN IDs (PowerShell)
			  • Live migration
			  • Cluster
			• 2.6-Rename the virtual network adapters
			• 2.7-Assign static IPs as necessary
			• 2.8-Configure QoS for different traffic types and configure the default minimum bandwidth for the switch

3	Initial Storage Configuration	Single Node	• 3.1-Present all shared storage to relevant nodes • 3.2-For multipath scenarios, install and configure multipath I/O (MPIO) as necessary • 3.3-All shared disks: Wipe, bring online and initialize
4	Failover Cluster Setup	Single Node	• 4.1-Run through the Cluster Validation Wizard • 4.2-Address any indicated warnings and/or errors • 4.3-Complete the Create Cluster Wizard (setting name and IP but do not add eligible storage) • 4.4-Create the clustered

			storage pool. • 4.5-Create the quorum virtual disk • 4.6-Create the virtual machine storage virtual disk. • 4.7-Add the virtual machine storage virtual disk to cluster shared volumes. • 4.8-Add folders to the cluster shared volume. • 4.9-Configure quorum settings • 4.10-Configure cluster networks to prioritize traffic.
5	Hyper-V Configuration	All Nodes	• 5.1-Change default file locations, mapping to CSV volumes
6	Cloud Validation	Single Node	• 6.1-Create a virtual machine, attaching an existing operating system VHD and tagging to the appropriate VLAN • 6.2-Test network connectivity from the virtual machine. • 6.3-Perform a Live Migration • 6.4-Perform a quick migration

Step 1: Initial node configuration

In step 1, you will perform the following steps on all nodes of the Hyper-V cluster:

- 1.1 Add appropriate VLANS to the interface ports on the physical switch.
- 1.2 Enable BIOS settings required for Hyper-V.
- 1.3 Perform a clean operating system installation.
- 1.4 Perform post-installation tasks.
- 1.5 Install roles and features using the default settings.

1.1 Add appropriate VLANS to the interface ports on the physical switch

Cluster nodes will be configured to use different VLAN tags for the following traffic types:

- **Management traffic** – untagged/default
- **Tenant traffic** – tagged
- **Live migration traffic** – tagged
- **Cluster and CSV traffic** – tagged

VLANs are configured to enable traffic isolation and quality of service policies. Define the VLAN tag numbers for each traffic type and then configure your switch with the appropriate VLAN port numbers. The procedures for doing this will vary with the switch make and model. Please refer to your switch documentation for more information. Note that management traffic is typically not tagged because it can interfere with a number of core host system activities. While you can tag the management traffic, you may run into problem with issues such as PXE boot. Therefore, we recommend that you do not tag the management traffic.

1.2 Enable BIOS settings required for Hyper-V

You will need to enable virtualization support in the BIOS of each cluster member prior to installing the Hyper-V server role. The procedure for enabling processor virtualization support will vary with your processors' make and model and the system BIOS. Please refer to your hardware documentation for the appropriate procedures.

1.3 Perform a clean operating system installation

Install Windows Server 2012 using the Full Installation option.

1.4 Perform post-installation tasks

There are several tasks you need to complete on each node after the operating system installation is complete. These include:

- Join each node to the domain
- Enable remote access to each node via the Remote Desktop Protocol.
- Set the Windows PowerShell execution policy.
- Enable Windows PowerShell remoting.

Perform the following steps to join each node to the domain:

1. Press the **Windows Key** on the keyboard and then press **R**. Type **Control Panel** and then click **OK**.
2. In the **Control Panel** window, click **System and Security**, and then click **System**.
3. In the **System** window under **Computer name, domain, and workgroup settings**, click **Change settings**.
4. In the **System Properties** dialog box, click **Change**.
5. Under **Member of**, click **Domain**, type the name of the domain, and then click **OK**.

Run the following Windows PowerShell commands on each node to enable remote access using the Remote Desktop Protocol, to enable PowerShell execution policy and enable PowerShell Remoting:

```
(Get-WmiObject Win32_TerminalServiceSetting -Namespace
root\cimv2\terminalservices).SetAllowTsConnections(1,1)

Set-ExecutionPolicy Unrestricted –Force

Enable-PSRemoting –Force
```

1.5 Install roles and features using the default settings

The following roles and features will be installed on each node of the cluster:

- Hyper-V and Hyper-V management Tools
- Failover cluster and failover cluster management tools
- Storage management tools

Perform the following steps on each node in the cluster to install the required roles and features:

1. In **Server Manager**, click **Dashboard** in the console tree.
2. In **Welcome to Server Manager**, click **2 Add roles and features**, and then click **Next**.
3. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
4. On the **Installation Type** page, click **Next**.
5. On the **Server Selection** page, click **Next**.
6. On the **Server Roles** page, select **Hyper-V** from the **Roles** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Click **Next**.
7. On the **Features** page, select **Failover Clustering** from the **Features** list. In the **Add Roles and Features Wizard** dialog box, click **Add Features**. Expand **Remote Server Administrator Tools** and then expand **Role Administration Tools**. Expand **File Services Tools**. Select **Share and Storage Management Tool**. Click **Next**.

   Note
   If you plan to use Multipath I/O for your storage solution, select the Multipath I/O feature while performing step 7.
8. On the **Hyper-V** page, click **Next**.
9. On the **Virtual Switches** page, click **Next**.
10. On the **Migration** page, click **Next**.
11. On the **Default Stores** page, click **Next**.
12. On the **Confirmation** page, put a checkmark in the **Restart the destination server automatically if required** checkbox and then in the **Add Roles and Features** dialog box click **Yes**, then click **Install**.
13. On the **Installation progress** page, click **Close** after the installation has succeeded.
14. Restart the computer. This process might require restarting the computer twice. If so, the installer will trigger the multiple restarts automatically.

After you restart the server, open **Server Manager** and confirm that the installation completed successfully. Click Close on the Installation Progress page.

Step 2: Initial network configuration

The network configuration on each node in the cluster needs to be configured to support the converged networking scenario where all traffic, including infrastructure and tenant traffic, moves through the Hyper-V virtual switch. You will perform the following procedures on each of the nodes in the cluster to complete the initial network configuration:

- 2.1 Disable unused and disconnected interfaces and rename active connections.
- 2.2 Create a converged network adapter team and configure IP addressing information.
- 2.3 Create the Hyper-V virtual switch and management virtual network adapter.
- 2.4 Rename the management virtual network adapter (optional).
- 2.5 Create additional virtual network adapters and assign VLAN IDs.
- 2.6 Rename virtual network adapters (optional).
- 2.7 Assign static IP addresses to the virtual network adapters.
- 2.8 Configure QoS for different traffic types and configure the default minimum bandwidth for the virtual switch.

2.1 Disable unused and disconnected interfaces and rename active connections

You can simplify the configuration and avoid errors when running the wizards and running PowerShell commands by disabling all network interfaces that are either unused or disconnected. You can disable these network interfaces in the **Network Connections** window.

For the remaining network adapters, do the following:

1. Connect them to the converged network switch ports.
2. To help you more easily recognize the active network adapters, rename them with names that indicate their use or their connection to the intranet or Internet (for example, **ConvergedNet1** and **ConvergedNet2**). You can do this in the **Network Connections** window.

2.2 Create a converged network adapter team

Network Load Balancing and Failover (LBFO) enables bandwidth aggregation and network adapter failover to prevent connectivity loss in the event of a network card or port failure. This feature is often referred to as "NIC Teaming". In this scenario you will create one team that will be connected to the ConvergedNet subnet.

To configure the network adapter teams by using **Server Manager**, do the following on *each computer in the cluster*:

Note

Several steps in the following procedure will temporarily interrupt network connectivity. We recommend that all servers be accessible over a keyboard, video, and mouse (KVM) switch so that you can check on the status of these machines if network connectivity is unavailable for more than five minutes.

1. From **Server Manager**, click **Local Server** in the console tree.
2. In **Properties**, click **Disabled**, which you'll find next to **Network adapter teaming**.
3. In the **NIC Teaming** window, click the name of the server computer in **Servers**.
4. In **Teams**, click **Tasks**, and then click **New Team**.
5. In the **New Team** window, in the **Team Name** text box, enter the name of the network adapter team for the converged traffic subnet (example: **ConvergedNet Team**).
6. In the **Member adapters** list select the two network adapters connected to the converged traffic subnet (in this example, **ConvergedNet1** and **ConvergedNet2**), and then click **OK**. Note that there may be a delay of several minutes before connectivity is restored after making this change. To ensure that you see the latest state of the configuration, right click your server name in the **Servers** section in the **NIC Teaming** window and click **Refresh Now**. There may be a delay before the connection displays as **Active**. You may need to refresh several times before seeing the status change.
7. Close the **NIC Teaming** window.

Configure a static IPv4 addressing configuration for the new network adapter team connected to the converged traffic subnet (example: **ConvergedNet Team**). This IP address is the one that you will use when connecting to the host system for management purposes. You can do this in the **Properties** of the team in the **Network Connections** window. You will see a new adapter where the name of the teamed network adapter is the name you assigned in step 5. You will lose connectivity for a few moments after assigning the new IP addressing information.

Note

You might need to manually refresh the display of the NIC Teaming window to show the new team and there may be a delay in connectivity as the network adapter team is created. If you are managing this server remotely, you might temporarily lose connectivity to the server.

2.3 Create the Hyper-V virtual switch and management virtual network adapter

In this scenario, all traffic will flow through the Hyper-V virtual switch. This includes the host operating system traffic (cluster/CSV, management and live migration) and guest/tenant traffic. You will create the virtual switch in Windows PowerShell instead of using the Hyper-V console. The reason for this is that when you create the virtual switch in the Hyper-V console, you are unable to specify the Minimum Bandwidth Mode, which defaults to Absolute (requiring a bits per second value) as opposed to Weighted mode (configurable via Windows PowerShell) which allows for a relative range from 1 to 100. For more information on the New-VMSwitch cmdlet, please see New-VMSwitch.

Run the following Windows PowerShell command on each member of the cluster to create the Hyper-V virtual switch and the management traffic virtual network adapter:

```
New-VMSwitch "ConvergedNetSwitch" -MinimumBandwidthMode weight -NetAdapterName
"ConvergedNetTeam" -AllowManagementOS 1
```

Please note that during this step you created both the Hyper-V virtual switch and the management virtual NIC.

Note that if you are performing this action of an RDP connection, the connection may drop for a few moments.

2.4 Rename the management virtual network adapter (optional)

The management virtual network adapter that was created when you created the virtual switch now appears in the **Network Connections** window and it was assigned a generic name such as **vEthernet** (**ConvergedNetSwitch**). You should rename this virtual network adapter to make it easier to identify in subsequent operations. Right-click the new virtual network adapter and click **Rename** and assign the virtual network adapter a new name (for example, Management).

2.5 Create additional virtual network adapters and assign VLAN IDs

The Hyper-V virtual switch now has a single virtual network adapter that will be used for hosting operating system management traffic. You now will create two additional virtual network adapters: one for live migration traffic and one for cluster traffic.

Run the following Windows PowerShell commands to create the live migration traffic virtual network adapter, the cluster traffic virtual network adapter, assign the live migration virtual network adapter a VLAN ID and assign the cluster virtual network adapter a VLAN ID:

```
Add-VMNetworkAdapter -ManagementOS -Name "LiveMigration" -SwitchName "ConvergedNetSwitch"

Add-VMNetworkAdapter -ManagementOS -Name "Cluster" -SwitchName "ConvergedNetSwitch"

Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName LiveMigration -Access -
VlanId 2160

Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Cluster -Access -VlanId 2161
```

In the preceding example, the **–VMNetworkAdapterName** represents the name of the virtual network adapter.

2.6 Rename virtual network adapters (optional)

The live migration and cluster virtual network adapters you created in the previous step now appear in the **Network Connections** window and they were assigned default names, such as **vEthernet (Cluster)**. You should rename these virtual network adapters to make them easier to identify in subsequent operations. Right-click each of these virtual network adapters, click **Rename** and assign the virtual network adapter a new name (for example, live migration and cluster)

2.7 Assign static IP addresses to the virtual network adapters

You now need to assign IP addresses to your virtual network adapters. This can be done through DHCP or you can assign static addresses. Make sure that each of the virtual network adapters is assigned an IP address on a different network ID – this will become important later when you configure your cluster networking configuration. You can use the Networking Control Panel applet or Windows PowerShell to assign IP addressing information to the virtual network adapters. For example:

```
Set-NetIPInterface -InterfaceAlias "LiveMigration" -dhcp Disabled; new-NetIPAddress -
PrefixLength 8 -InterfaceAlias "LiveMigration" -IPv4Address 11.0.0.x
```

2.8 Configure QoS for different traffic types and configure the default minimum bandwidth for the switch

In this step you will configure QoS weightings that define the minimal share of bandwidth assigned to each of the virtual network adapters. You can determine the percentage of bandwidth that can be allocated to a particular virtual network adapter by adding all the weight values together and then dividing the individual weight assigned to a virtual network adapter by the total.

Perform the following steps to add weight values to the cluster, management and live migration virtual network adapters:

Run the following Windows PowerShell commands to assign a weight value to the cluster virtual network adapter, assign a weight value to the management virtual network adapter, assign a weight value to the live migration virtual network adapter and assign a default weight value for any future virtual network adapters you create:

```
Set-VMNetworkAdapter -ManagementOS -Name "Cluster" -MinimumBandwidthWeight 40

Set-VMNetworkAdapter -ManagementOS -Name ConvergedNetSwitch -MinimumBandwidthWeight 5

Set-VMNetworkAdapter -ManagementOS -Name "LiveMigration" -MinimumBandwidthWeight 20

Set-VMSwitch "ConvergedNetSwitch" -DefaultFlowMinimumBandwidthWeight 10
```

Note

VMNetworkAdapter name is listed under 'Device Name' in the **Network Connections** user interface.

Step 3: Initial storage configuration

With the initial cluster node configuration complete, you are ready to perform initial storage configuration tasks on all nodes of the cluster. Initial storage configuration tasks include:

- 3.1 Present all shared storage to relevant nodes.
- 3.2 Install and configure MPIO as necessary for multipath scenarios.
- 3.3 Wipe, bring online, and initialize all shared disks.

3.1 Present all shared storage to relevant nodes

In a SAS scenario, connect the SAS adapters to each storage device. Each cluster node should have two adapters in them if high availability to storage access is required.

3.2 Install and configure MPIO as necessary for multipath scenarios

If you have multiple data paths to storage (for example, two SAS cards) make sure to install the Microsoft® Multipath I/O (MPIO) on each node. This step might require you to restart the system. For more information about MPIO, see What's New in Microsoft Multipath I/O.

3.3 Wipe, bring online, and initialize all shared disks

To prevent issues with the storage configuration procedures that are detailed later is this document; confirm that the disks in your storage solution have not been previously provisioned. The disks should have no partitions or volumes. They should also be initialized so that there is a master book record (MBR) or GUID partition table (GPT) on the disks, and then brought online. You can use the **Disk Management** console or Windows PowerShell to accomplish this task. This task must be completed on each node in the cluster.

To discover disks that can participate in a pool you can use the PowerShell command Get-PhysicalDisk | ? BusType -Eq "SAS". All disks in the **CanPool** column that read TRUE are eligible for pooling.

Note

If you have previously configured these disks with Windows Server 2012 Storage Spaces pools, you will need to delete these storage pools prior to proceeding with the storage configuration described in this document. Please refer to the TechNet Wiki Article, How to Delete Storage Pools and Virtual Disks Using PowerShell.

Step 4: Failover cluster setup

You are now ready to complete the failover cluster settings. Failover cluster setup includes the following steps:

- 4.1 Run through the Cluster Validation Wizard.
- 4.2 Address any indicated warnings and/or errors.
- 4.3 Complete the Create Failover Cluster Wizard.
- 4.4 Create the clustered storage pool.
- 4.5 Create the quorum virtual disk.
- 4.6 Create the virtual machine storage virtual disk.
- 4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes.
- 4.8 Add folders to the cluster shared volume.
- 4.9 Configure Quorum Settings.
- 4.10 Configure cluster networks to prioritize traffic.

4.1 Run through the Cluster Validation Wizard

The Cluster Validation Wizard will query multiple components in the intended cluster hosts and confirm that the hardware and software is ready to support failover clustering. On one of the nodes in the server cluster, perform the following steps to run the Cluster Validation Wizard:

1. In the **Server Manager**, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** console, in the **Management** section, click **Validate Configuration**.
3. On the **Before You Begin** page of the **Validate a Configuration Wizard**, click **Next**.
4. On the **Select Servers or a Cluster** page, type the name of the local server, and then click **Add**. After the name appears in the **Selected servers** list, type the name of another Hyper-V cluster member computer, and then click **Add**. Repeat this step for all computers in the Hyper-V cluster. When all of the servers of the Hyper-V cluster appear in the **Selected servers** list, click **Next**.
5. On the **Testing Options** page, click **Next**.
6. On the **Confirmation** page, click **Next**. The time to complete the validation process will vary with the number of nodes in the cluster and can take some time to complete.
7. On the **Summary** page, the summary text will indicate that the configuration is suitable for clustering. Confirm that there is a checkmark in the **Create the cluster now using the validated nodes...** checkbox.

4.2 Address any indicated warnings and/or errors

Click the **Reports** button to see the results of the Cluster Validation. Address any issues that have led to cluster validation failure. After correcting the problems, run the Cluster Validation Wizard again. After the cluster passes validation, then proceed to the next step. Note that you may see errors regarding disk storage. You may see this if you haven't yet initialized the disks. Click **Finish**.

4.3 Complete the Create Failover Cluster Wizard

After passing cluster validation, you are ready to complete the cluster configuration.

Perform the following steps to complete the cluster configuration:

1. On the **Before You Begin** page of the **Create Cluster Wizard**, click **Next**.
2. On the **Access Point for Administering the Cluster** page, enter a valid NetBIOS name for the cluster, and then select the network you want the cluster on and then type in a static IP address for the cluster, and then click **Next**. In this example, the network you would select is the Management Network. Unselect all other networks that appear here.
3. On the **Confirmation** page, clear **Add all eligible storage to the cluster** checkbox and then click **Next**.
4. On the **Creating New Cluster** page you will see a progress bar as the cluster is created.
5. On the **Summary** page, click **Finish**.
6. In the console tree of the **Failover Cluster Manager** snap-in, open the **Networks** node under the cluster name.

7. Right-click the cluster network that corresponds to the management network adapter network ID (subnet), and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is enabled. In the **Name** text box, enter a friendly name for this network (for example, **ManagmentNet**), and then click **OK**.
8. Right-click the cluster network that corresponds to the Cluster network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network is selected** and that **Allow clients to connect through this network** is *not* enabled. In the **Name** text box, enter a friendly name for this network (for example, **ClusterNet**), and then click **OK**.
9. Right-click the cluster network that corresponds to the live migration network adapter network ID (subnet) and then click **Properties**. On the **General** tab, confirm that **Allow cluster communications on this network** is selected and that **Allow clients to connect through this network** is *not* enabled. In the **Name** text box, enter a friendly name for this network (for example, **LiveMigrationNet**), and then click **OK**.

4.4 Create a clustered storage pool

Perform the following steps on one of the members of the cluster to create the storage pool:

1. In the left pane of the **Failover Cluster Manager**, expand the server name and then expand the **Storage** node. Click **Storage Pools**.
2. In the **Actions** pane, click **New Storage Pool**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Storage Pool Name** page, enter a name for the storage pool in the **Name** text box. Enter an optional description for the storage pool in the Description text box. In the **Select the group of available disks (also known as a primordial pool) that you want to use** list, select the name you assigned to the cluster (this is the NetBIOS name you assigned to the cluster when you created the cluster). Click **Next**.
5. On the **Physical Drives** page, select the drives that you want to participate in the storage pool. Then click **Next**.
6. On the **Confirmation** page, confirm the settings and click **Create**.
7. On the **Results** page, you should receive the message **You have successfully completed the New Storage Pool Wizard**. Remove the checkmark from the **Create a virtual disk when the wizard closes** checkbox. Then click **Close**.

4.5 Create the quorum virtual disk

Now that you have created the storage pool, you can create virtual disks within that storage pool. A virtual disk is sometimes called a logical unit number or LUN and it represents a collection of one or more physical disks from the previously created storage pool. The layout of data across the physical discs can increase the reliability and performance of the physical disk.

You will need to create at least two virtual disks:

- A virtual disk that can be used as a quorum witness disks. This disk can be configured as a 1 GB virtual disk.

- A virtual disk that will be assigned to a cluster shared volume.

Perform the following steps to create the **quorum** virtual disk:

1. In the **Failover Cluster Manager** console, expand the **Storage** node in the left pane of the console. Right click **Pools** and click **Add Disk**.
2. In the **New Virtual Disk Wizard** on the **Before You Begin** Page, click **Next**.
3. On the **Storage Pool** page, select your server name in the **Server** section and then select the storage pool you created earlier in the **Storage pool** section. Click **Next**.
4. On the **Virtual Disk Name** page, enter a name for the virtual disk in the **Name** text box. You can also enter an optional description in the **Description** text box. Click **Next**.
5. On the **Storage Layout** page, in the **Layout** section, select **Mirror**. Click **Next**.
6. On the **Resiliency Settings** select **Two-way mirror** and click **Next**.
7. On the **Size** page, in the **Virtual disk size** text box, enter a size for the new virtual disk, which in this example will be 1 GB. Use the drop down box to select GB. Also, you can put a checkmark in the **Create the largest virtual disk possible, up to the specified size** checkbox, but this is not required or desired when creating a witness disk. When this option is selected, it allows the wizard to calculate the largest size virtual disk you can create given the disks you have assigned to the pool, regardless of the number you put in the **Virtual disk size** text box. Click **Next**.
8. On the **Confirmation** page, review your settings and click **Create**.
9. On the **Results** page, put a checkmark in the **Create a volume when this wizard closes** checkbox. Click **Close**.
10. On the **Before You Begin** page of the **New Volume Wizard**, click **Next**.
11. On the **Server and Disk** page, select the name of the cluster from the **Server** list. In the **Disk** section, select the virtual disk you just created. You can identify this disk by looking in the **Virtual Disk** column, where you will see the name of the virtual disk you created. Click **Next**.
12. On the **Size** page, accept the default volume size, and click **Next**.
13. On the **Drive Letter or Folder** page, select **Drive letter** and select a drive letter. Click **Next**.
14. On the **File System Settings** page, from the **File system** drop down list, select **NTFS**. Use the default setting in the **Allocation unit size** list. Click **Next**.
15. On the **Confirmation** page, click **Create**.
16. On the **Results** page, click **Close**.

4.6 Create the virtual machine storage virtual disk

Perform the following steps to create the **virtual machine storage** virtual disk:

1. In the **Failover Cluster Manager** console, expand the **Storage** node in the left pane of the console. Right click **Pools** and click **Add Disk**.
2. In the **New Virtual Disk Wizard** on the **Before You Begin** page, click **Next**.
3. On the **Storage Pool** page, select your server name in the **Server** section and then select the storage pool you created earlier in the **Storage pool** section. Click **Next**.
4. On the **Virtual Disk Name** page, enter a name for the virtual disk in the **Name** text box. You can also enter an optional description in the **Description** text box. Click **Next**.

5. On the **Storage Layout** page, in the **Layout** section, select **Mirror**. Click **Next**.
6. On the **Resiliency Settings** select **Two-way mirror** and click **Next**.
7. On the **Size** page, in the **Virtual disk size** text box, enter a size for the new virtual disk. Use the drop down box to select MB, GB or TB. Also, you can put a checkmark in the **Create the largest virtual disk possible, up to the specified size** checkbox. When this option is selected, it allows the wizard to calculate the largest size virtual disk you can create given the disks you have assigned to the pool, regardless of the number you put in the **Virtual disk size** text box. Click Next.
8. On the **Confirmation** page, review your settings and click **Create**.
9. On the **Results** page, put a checkmark in the **Create a volume when this wizard closes** checkbox. Click **Close**.
10. On the **Before You Begin** page of the **New Volume Wizard**, click **Next**.
11. On the **Server and Disk** page, select the name of the cluster from the **Server** list. In the **Disk** section, select the virtual disk you just created. You can identify this disk by looking in the **Virtual Disk** column, where you will see the name of the virtual disk you created. Click **Next**.
12. On the **Size** page, accept the default volume size, and click **Next**.
13. On the **Drive Letter or Folder** page, select the **Don't Assign to a drive letter or folder** and select a drive letter. Click **Next**.
14. On the **File System Settings** page, from the **File system** drop down list, select **NTFS**. Use the default setting in the **Allocation unit size** list. Note that **ReFS** is not supported in a **Cluster Shared Volume** configuration. Click **Next**.
15. On the **Confirmation** page, click **Create**.
16. On the **Results** page, click **Close**.

4.7 Add the virtual machine storage virtual disk to Cluster Shared Volumes

The virtual disk you created for virtual machine storage is now ready to be added to a Cluster Shared Volume. Perform the following steps to add the virtual disk to a Cluster Shared Volume.

1. In the **Failover Cluster Manager**, in the left pane of the console, expand the **Storage** node and click **Disks**. In the middle pane of the console, in the **Disks** section, right click the virtual disk you created in the previous step and then click **Add to Cluster Shared Volumes**.
2. Proceed to the next step.

4.8 Add folders to the cluster shared volume

Now you need to create the folders on the virtual disk located on the Cluster Shared Volume to store the virtual machine files and the virtual machine data files.

Perform the following steps to create a file share to store the running VMs of the Hyper-V cluster:

1. Open **Windows Explorer** and navigate to the C: drive and then double-click **Cluster Storage** and then double-click **Volume 1**.
2. Create two folders in **Volume 1**. One of the folders will contain the .vhd files for the virtual machines (for example, VHDdisks) and one folder will contain the virtual machine configuration files (for example, VHDsettings)

4.9 Configure Quorum Settings

Perform the following steps to configure quorum settings for the cluster:

1. In the left pane of the **Failover Cluster Manager** console, right click on the name of the cluster and click **More Actions** and click **Configure Cluster Quorum Settings**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Quorum Configuration Option** page, select **Use typical settings (recommended)** and click **Next**.
4. On the **Confirmation** page, click **Next**.

4.10 Configure cluster networks to prioritize traffic

The cluster will use the network with the lowest metric for CSV traffic and the second lowest metric for live migration. Windows PowerShell® is the only method available to prescriptively specify the CSV network. You can set the live migration network via the Hyper-V management console, which you will do in Step 5: Configure Hyper-V settings.

Run the following Windows PowerShell commands on one node of the failover cluster to set the metric for the cluster network traffic, set the metric for the live migration network traffic and set the metric for the management network traffic:

```
(Get-ClusterNetwork "ClusterNet" ).Metric = 100

(Get-ClusterNetwork "LiveMigrationNet" ).Metric = 500

(Get-ClusterNetwork "ManagementNet" ).Metric = 1000
```

Step 5: Configure Hyper-V settings

To finalize the Hyper-V configuration, you will need to take the following step:

- 5.1 Change default file locations for virtual machine files.

5.1 Change default file locations for virtual machine files

On each Hyper-V cluster member, perform the following steps on to change the default file locations for virtual machine files:

1. In **Server Manager**, click **Tools**, then click **Hyper-V Manager**.
2. From the console tree of the **Hyper-V Manager**, right-click the name of the Hyper-V server, and then click **Hyper-V Settings**.
3. In the **Hyper-V Settings** dialog box, click **Virtual Hard Disks** under **Server**, type the file share location in **Specify the default folder to store virtual hard disk files**, and then click **Apply**. For example, `c:\clusterstorage\volume1\VHDdisks`.
4. Click **Virtual Machines** under **Server**, type the file folder location in **Specify the default folder to store virtual machine configuration files**, and then click **OK** For example, `c:\clusterstorage\volume1\VHDsettings`.

Step 6: Cloud validation

To verify the configuration of your cloud environment, perform the following operations.

- 6.1 Create a new virtual machine.
- 6.2 Test network connectivity from the virtual machine.
- 6.3 Perform a live migration.
- 6.4 Perform a quick migration.

6.1 Create a new virtual machine

To create a new virtual machine in the cluster environment, perform the following steps.

1. Open **Failover Cluster Manager**, click **Roles** under the cluster name, click **Virtual Machines** under the **Actions** pane, and then click **New Virtual Machine**.
2. On the **New Virtual Machine** page, select the cluster node where you want to create the virtual machine, and then click **OK**.
3. On the **Before you Begin** page of the **New Virtual Machine Wizard**, click **Next**.
4. On the **Specify Name and Location** page, enter a friendly name for this virtual machine and then click **Next**.
5. On **the Assign Memory** page, enter the amount of memory that will be used for this virtual machine (minimum for this lab is 1024 MB RAM) and then click **Next**.
6. On the **Configuring Networking** page, click **Next**.
7. On the **Connect Virtual Hard Disk** page, leave the default options selected and click **Next**.
8. On the **Installation Options** page, select **Install an operating system from a boot CD/DVD-ROM** and then select the location where the CD/DVD is located. If you are installing the new operating system based on an ISO file, make sure to select the option **Image file** (.iso) and browse for the file location. If you prefer to PXE boot, that option will be described in later steps. After you select the appropriate option for your scenario, click **Next**.
9. On the **Completing the New Virtual Machine Wizard** page, review the options, and then click **Finish**.
10. The virtual machine creation process starts. After it is finished, you will see the **Summary** page, where you can access the report created by the wizard. If the virtual machine was created successfully, click **Finish**.
11. If you want to PXE boot the virtual machine, you will need to create a **Legacy Network Adapter**. Right click the new virtual machine and click settings.
12. In the **Settings** dialog box, select the **Legacy Network Adapter** option and click **Add**.
13. In the **Legacy Network Adapter** dialog box, connect it to the virtual switch (such as ConvergedNetSwitch) and enable virtual LAN identification and assign the appropriate network identifier.

Note

If the virtual machine continues to use the legacy network adapter it will not be able to leverage many of the features available in the Hyper-V virtual switch. You may want to replace the legacy network adapter after the operating system is installed.

At this point your virtual machine is created and you should use the Failover Cluster Manager to start the virtual machine and perform the operating system installation according to the operating system that you choose. For the purpose of this validation, the guest operating system can be any Windows Server version.

6.2 Test network connectivity from the virtual machine

Once you finish installing the operating system in the virtual machine you should log on and verify if this virtual machine was able to obtain IP address from the enterprise network. Assuming that in this network you have a DHCP server, this virtual machine should be able to obtain the IP address. To perform the basic network connectivity test use the following approach.

- Use `ping` command for a reachable IP address in the same subnet.
- Use `ping` command for the same destination but now using the full qualified domain name for the destination host. The goal here is to test basic name resolution.

Note

If you installed Windows "8" Developer Preview in this virtual machine you need to open Windows Firewall with Advanced Security and create a new rule to allow Internet Control Message Protocol (ICMP) before performing the previous tests. This may be true for other hosts you want to ping − confirm that the host-based firewall on the target allows for ICMP Echo Requests.

After you confirm that this basic test is working properly, leave a command prompt window open and enter the command, `ping <Destination_IP_Address_or_FQDN> -t`. The goal here is to have a continuous test while you perform the live migration to the second node.

Note

If you prefer to work with PowerShell, instead of the ping command you can use the Test-Connection command. This cmdlet provides you a number of connectivity testing options that exceed what is available with the simple ping command.

6.3 Perform a live migration

To perform a live migration of this virtual machine from the current cluster node to the other node in the cluster, perform the following steps.

1. In the **Failover Cluster Manager**, click **Roles** under the cluster name. On the **Roles** pane, right click the virtual machine that you created, click **Move**, click **Live Migration**, and then click **Select Node**.
2. On the **Move Virtual Machine** page, select the node that you want to move the virtual machine to and click **OK**.

You will notice in the **Status** column when the live migration starts, it will take some time for the **Information** column to update the current state of the migration. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

6.4 Perform a quick migration

To perform the quick migration of this virtual machine from the current node to the other one, perform the following steps.

1. On the **Failover Cluster Manager**, click **Roles** under the cluster name. In the **Roles** pane, right-click the virtual machine that you created, click **Move**, click **Quick Migration** and then click **Select Node**.
2. On the **Move Virtual Machine** window, select the node that you want to move the virtual machine to, and then click **OK**.

You will notice in the status that the quick migration will start faster than the live migration did. While the migration is taking place you can go back to the virtual machine that has the ping running and observe if there is any packet loss.

Dynamic Access Control: Scenario Overview

In Windows Server 2012, you can apply data governance across your file servers to control who can access information and to audit who has accessed information. Dynamic Access Control lets you:

- Identify data by using automatic and manual classification of files. For example, you could tag data in file servers across the organization.
- Control access to files by applying safety-net policies that use central access policies. For example, you could define who can access health information within the organization.
- Audit access to files by using central audit policies for compliance reporting and forensic analysis. For example, you could identify who accessed highly sensitive information.
- Apply Rights Management Services (RMS) protection by using automatic RMS encryption for sensitive Microsoft Office documents. For example, you could configure RMS to encrypt all documents that contain Health Insurance Portability and Accountability Act (HIPAA) information.

The Dynamic Access Control feature set is based on infrastructure investments that can be used further by partners and line-of-business applications, and the features can provide great value for organizations that use Active Directory. This infrastructure includes:

- A new authorization and audit engine for Windows that can process conditional expressions and central policies.
- Kerberos authentication support for user claims and device claims.
- Improvements to the File Classification Infrastructure (FCI).
- RMS extensibility support so partners can provide solutions that encrypt non-Microsoft files.

In this scenario

The following scenarios and guidance are included as part of this content set:

Dynamic Access Control Content Roadmap

Scenario	Evaluate	Plan	Deploy	Operate
Scenario: Central Access Policy Creating Central access policies for files allow organizations to centrally deploy and manage authorization policies that include conditional expressions using user claims, device claims, and resource properties. These polices are based on compliance and business regulatory requirements. These policies are created and hosted in Active Directory, therefore making it easier to manage and deploy. **Deploying Claims Across Forests** In Windows Server 2012, the AD DS maintains a 'claims dictionary' in each forest and all claim types in use within the forest are defined at the Active Directory forest level. There	Dynamic Access Control: Scenario Overview Deploy Claims Across Forests	Plan for a Central Access Policy Deployment • Process to map a business request to a central access policy • Delegating of administration for Dynamic Access Control • Exception Mechanisms for Planning Central Access Policies Best Practices for Using User Claims • Choosing the right configuration to enable claims in your user domain • Operations to enable user claims • Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies Using Device Claims and Device Security Groups	Deploy a Central Access Policy (Demonstration Steps) Deploy Claims Across Forests (Demonstration Steps)	• Modeling a central access policy

Scenario	Evaluate	Plan	Deploy	Operate
are many scenarios where a principal may need to traverse a trust boundary. This scenario describes how a claim traverses a trust boundary.		• Considerations for using static device claims • Operations to enable device claims Tools for Deployment • Data Classification Toolkit		
Scenario: File Access Auditing Security auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is regulatory compliance. For example, industry standards such as Sarbanes Oxley, HIPAA, and Payment Card Industry (PCI) require enterprises to follow a strict set of rules related to data security and privacy. Security audits help establish the presence or absence of such policies; thereby, they prove compliance or	Scenario: File Access Auditing	Plan for File Access Auditing	Deploy Security Auditing with Central Audit Policies (Demonstration Steps)	• Monitor the Central Access Policies that Apply on a File Server • Monitor the Central Access Policies Associated with Files and Folders • Monitor the Resource Attributes on Files and Folders • Monitor Claim Types • Monitor User and Device Claims During

Scenario	Evaluate	Plan	Deploy	Operate
noncompliance with these standards. Additionally, security audits help detect anomalous behavior, identify and mitigate gaps in security policy, and deter irresponsible behavior by creating a record of user activity that can be used for forensic analysis.				Sign-in • Monitor Central Access Policy and Rule Definitions • Monitor Resource Attribute Definitions • Monitor the Use of Removable Storage Devices.
Scenario: Access-Denied Assistance Today, when users try to access a remote file on the file server, the only indication that they would get is that access is denied. This generates requests to helpdesk or IT administrators that need to figure out what the issue is and often the administrators have a hard time getting the appropriate context from users which makes it harder to resolve the issue. In Windows Server 2012, the goal is to	Scenario: Access-Denied Assistance	Plan for Access-Denied Assistance • Determine the access-denied assistance model • Determine who should handle access requests • Customize the access-denied assistance message • Plan for exceptions • Determine how access-denied assistance is deployed	Deploy Access-Denied Assistance (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
try and help the information worker and business owner of the data to deal with the access denied issue before IT gets involved and when IT gets involved, provide all the right information for a quick resolution. One of the challenges in achieving this goal is that there is no central way to deal with access denied and every application deals with it differently and thus in Windows Server 2012, one of the goals is to improve the access-denied experience for Windows Explorer.				
Scenario: Classification-Based Encryption for Office Documents Protection of sensitive information is mainly about mitigating risk for the organization. Various compliance regulations, such as HIPAA or Payment	Scenario: Classification-Based Encryption for Office Documents	Planning Considerations for Encryption of Office Documents	Deploy Encryption of Office Files (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
Card Industry Data Security Standard (PCI-DSS), dictate encryption of information, and there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it might impair business productivity. Thus, organizations tend to have different approaches and priorities for encrypting their information. To support this scenario, Windows Server 2012 provides the ability to automatically encrypt sensitive Windows Office files based on their classification. This is done through file management tasks that invoke Active Directory Rights Management Server (AD RMS) protection for sensitive documents a few seconds after the				

Scenario	Evaluate	Plan	Deploy	Operate
file is identified as being a sensitive file on the file server.				
Scenario: Get Insight into Your Data by Using Classification Reliance on data and storage resources has continued to grow in importance for most organizations. IT administrators face the growing challenge of overseeing larger and more complex storage infrastructures while simultaneously being tasked with the responsibility to ensure total cost of ownership is maintained at reasonable levels. Managing storage resources is not just about the volume or availability of data anymore, but also about the enforcement of company policies and knowing how storage is consumed to enable efficient utilization and	Scenario: Get Insight into Your Data by Using Classification	Plan for Automatic File Classification	Deploy Automatic File Classification (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
compliance to mitigate risk. File Classification Infrastructure provides insight into your data by automating classification processes so that you can manage your data more effectively. The following classification methods are available with File Classification Infrastructure: manual, programmatically, and automatic. This scenario focuses on the automatic file classification method.				
Scenario: Implement Retention of Information on File Servers A retention period is the amount of time that a document should be kept before it is expired. Depending on the organization, the retention period can be different. You can classify files in a folder as having a	Scenario: Implement Retention of Information on File Servers	Plan for Retention of Information on File Servers	Deploy Implementing Retention of Information on File Servers (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
short, medium, or long-term retention period and then assign the timeframe for each period. You may want to keep a file indefinitely by putting it on legal hold. File Classification Infrastructure and File Server Resource Manager uses file management tasks and file classification to apply retention periods for a set of files. You can assign a retention period on a folder and then use a file management task to configure how long an assigned retention period is to last. When the files in the folder are about to expire, the owner of the file gets a notification email. You can also classify a file as being on legal hold so that the file management task will not expire the file.				

Note
Dynamic Access Control is not supported on ReFS (Resilient File System).

See also

Content type	References
Product evaluation	• Dynamic Access Control Reviewers Guide • Dynamic Access Control Developer Guidance
Planning	• Plan for a Central Access Policy Deployment • Plan for File Access Auditing
Deployment	• Active Directory Deployment • File and Storage Services Deployment
Operations	Dynamic Access Control PowerShell Reference
Tools and settings	Data Classification Toolkit
Community resources	Directory Services Forum

Scenario: Central Access Policy

Central access policies for files enable organizations to centrally deploy and manage authorization policies that include conditional expressions that use user groups, user claims, device claims, and resource properties. (Claims are assertions about the attributes of the object with which they are associated). For example, to access high-business-impact (HBI) data, a user must be a full-time employee, obtain access from a managed device, and log on with a smart card. These policies are defined and hosted in Active Directory Domain Services (AD DS).

Organizational access policies are driven by compliance and business regulatory requirements. For example, if an organization has a business requirement to restrict access to personally identifiable information (PII) in files to only the file owner and members of the human resources (HR) department who are allowed to view PII information, this policy applies to PII files wherever they are located on file servers across the organization. In this example, you need to be able to:

- Identify and mark the files that contain PII.
- Identify the group of HR members who are allowed to view PII information.
- Create a central access policy that applies to all files that contain PII wherever they are located on file servers across the organization.

The initiative to deploy and enforce an authorization policy can come for many reasons and apply to multiple levels of the organization. The following are some example policy types:

- **Organization-wide authorization policy.** Most commonly initiated from the information security office, this authorization policy is driven by compliance or a high-level organization requirements, and it is relevant across the organization. For example, HBI files are accessible to only full-time employees.
- **Departmental authorization policy.** Each department in an organization has some special data-handling requirements that they want to enforce. For example, the finance department might want to limit access to finance servers to the finance employees.
- **Specific data-management policy.** This policy usually relates to compliance and business requirements, and it is targeted at protecting the correct access to the information that is being managed. For example, financial institutions might implement information walls so that analysts do not access brokerage information and brokers do not access analysis information.
- **Need-to-know policy.** This authorization policy type is typically used in conjunction with the previous policy types. For example, vendors should be able to access and edit only files that pertain to a project they are working on.

Real-life environments also teach us that every authorization policy needs to have exceptions so that organizations can quickly react when important business needs arise. For example, executives who cannot find their smart cards and need quick access to HBI information can call the Help Desk to get a temporary exception to access that information.

Central access policies act as security umbrellas that an organization applies across its servers. These policies enhance (but do not replace) the local access policies or discretionary access control lists (DACL) that are applied to files and folders. For example, if a DACL on a file allows access to a specific user, but a central policy that is applied to the file restricts access to the same user, the user cannot obtain access to the file. If the central access policy allows access, but the DACL does not allow access, the user cannot obtain access to the file.

A central access policy rule has the following logical parts:

- **Applicability.** A condition that defines which data the policy applies to, such as Resource.BusinessImpact=High.
- **Access conditions.** A list of one or more access control entries (ACEs) that define who can access the data, such as Allow | Full Control | User.EmployeeType=FTE.
- **Exceptions.** An additional list of one or more ACEs that define an exception for the policy, such as MemberOf(HBIExceptionGroup).

The following two figures show the workflow in central access and audit policies.

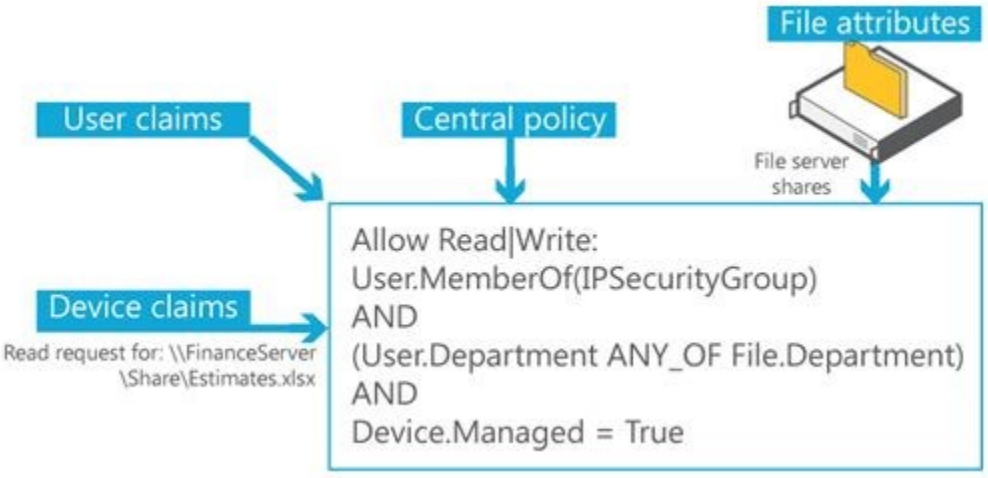

Figure 1 Central access and audit policy concepts

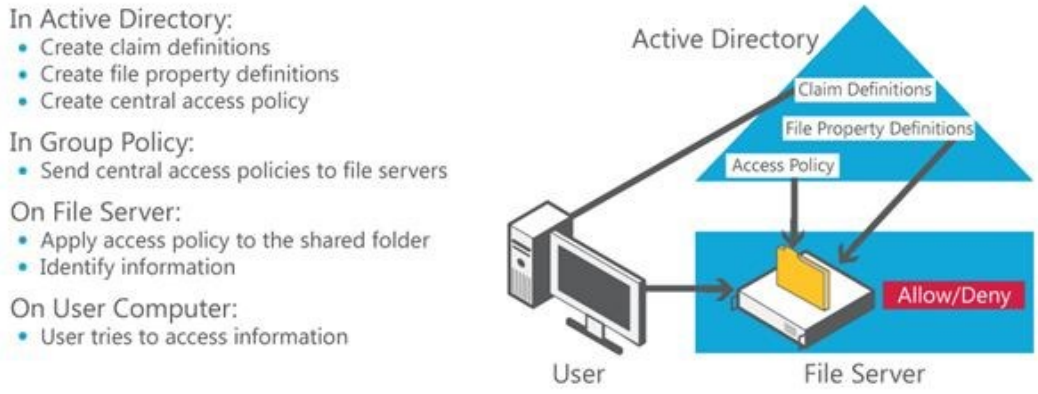

Figure 2 Central access policy workflow

The central authorization policy combines the following components:

- A list of centrally defined access rules that target specific types of information, such as HBI or PII.
- A centrally defined policy that contains a list of rules.
- A policy identifier that is assigned to each file on the file servers to point to a specific central access policy that should be applied during the access authorization.

The following figure demonstrates how you can combine policies into policy lists to centrally control access to files.

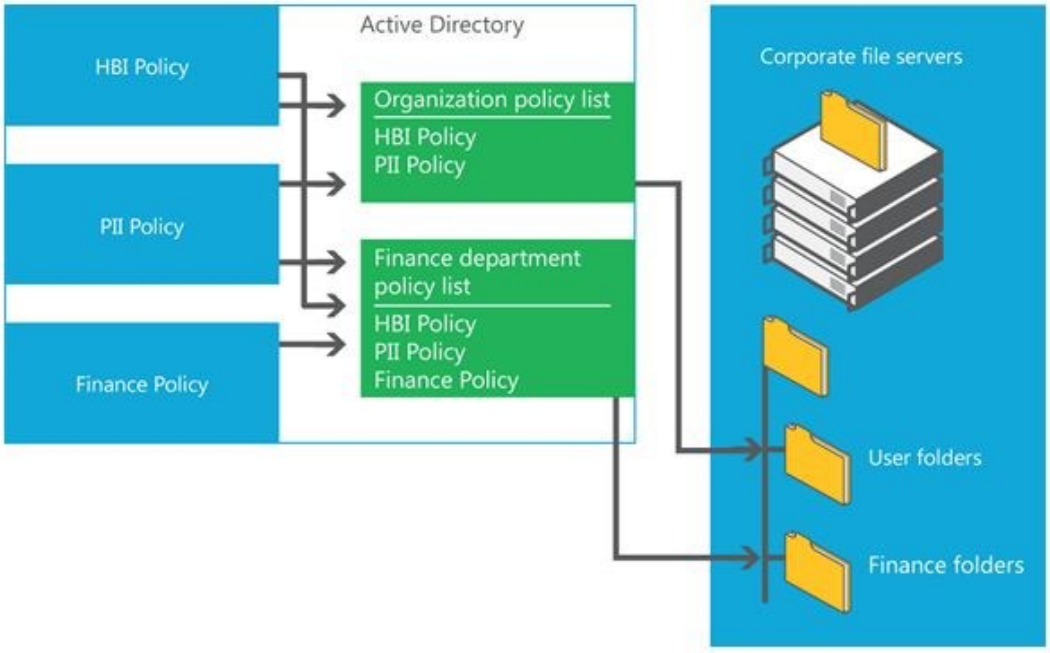

Figure 3 Combining policies

In this scenario

The following guidance is available to you for central access policies:

- Plan for a Central Access Policy Deployment
- Deploy a Central Access Policy (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services role	AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables the creation of user claims and device claims, compound identity, (user plus device claims), new central access policy (CAP) models, and the use of file-classification information in authorization decisions.

Role/feature	How it supports this scenario
File and Storage Services Server role	File and Storage Services provides technologies that help you set up and manage one or more file servers that provide central locations on your network where you can store files and share them with users. If your network users need access to the same files and applications, or if centralized backup and file management are important to your organization, you should set up one or more computers as a file server by adding the File and Storage Services role and the appropriate role services to the computers.
Windows client computer	Users can access files and folders on the network through the client computer.

Plan for a Central Access Policy Deployment

The need to control the information in enterprise-level organizations for compliance and business regulations is one of the drivers in the consolidation trend where large amounts of information from users' desktops and departmental file shares are moved into centrally managed file servers. The initiative to deploy and enforce an authorization policy may come for different reasons and from multiple levels of the organization:

- **Organization-wide authorization policy**: most commonly initiated from the Information Security office, this authorization policy is driven from compliance or very high level organization requirement and would be relevant across the organization. For example: High Business Impact files should be accessible by full time employees only
- **Departmental authorization policy**: each department in an organization has some special data handling requirements that they would like to enforce. This is very common in distributed organization. For example: the finance department might want to limit all access to finance servers to the finance employees
- **Specific data management policy**: this policy usually relates to compliance and business requirements and is targeted at protecting the right access to information that is being managed. For example: Preventing modification or deletion of files that are under retention or files that are under eDiscovery
- **Need to know policy**: This is a catch all authorization policy type and most probably used in conjunction with the policy types mentioned above. Examples include: Vendors should be able to access and edit only files that pertain to a project that they are working on.
  In financial institutions, information walls are important so that analysts do not access brokerage information and brokers do not access analysis information

Process to map a business request to a central access policy

1. Understand and translate business intent
2. Express access policy in Windows Server 2012 constructs
3. Determine the user groups, resource properties and claim types
4. Determine the servers where this policy should be applied to

Understand and translate business intent

Business determines that a central access policy is needed. The next step is to decide on the content (resources) that you want to apply policies. Then create a list of all the policies you want to apply to your content. For example some of the common policies that apply to the finance department of an organization would be:

- Archived finance documents should only be read by members of the Finance department.
- Members of the Finance department should only be able to access documents in their own country.
- Only Finance Administrators should have write access. An exception will be allowed for members of the FinanceException group. This group will have read access.

Express access policy in Windows Server 2012 constructs

The next step in the planning process is to translate the policies you require into expressions. A central access policy is targeted at providing an easy interpretation from a business requirement language to an authorization language.
A central policy in Windows Server 2012has the following distinct parts:

- Applicability: A condition that defines which data the policy applies to
  (Example: Resource.BusinessImpact=High)
- Access conditions: A list of one or more Access Control Entries (ACE) that defines who can access the data
  (Example: Allow | Full Control | User.EmployeeType=FTE)
- Exception: An additional list of one or more access control entries that define an exception for the policy
  (Example: MemberOf(HBIExceptionGroup)

Determine the user groups, resource properties and claim types

In this step, the expressions that were created for the policies are broken down and analyzed to figure out what resource properties, security groups as well as potential user claims need to be created to deploy the specific policies. Note that before using user claims, you should review the relevant section on user claims in this topic. Remember that you can always use security groups in case that you do not have the appropriate user claim attribute available in your environment.

Determine the servers where this policy should be applied to

The next step is to determine what file servers you want to deploy the access policies you have decided on. For example you can have a finance access policy that you want to roll out only to the finance file servers.

Planning Guidelines for Deploying Central Access Policies

There are multiple ways to deploy central access policies, based on different configurations in your environment. You can choose to have a simple deployment with security groups or have more advanced central access policies using user claims and device claims. The following sections in this topic discuss the different deployment options available to you base on the configuration you choose. The design and deployment guidance for each of these is discussed in further detail:

- Using Security Groups for Dynamic Access Control
- Using User Claims
- Using Device Claims and Device Security Groups

For a list of all configuration options and guidance on which configuration to choose, see the section, Configuring central access policies with different options

For a high-level view of the different deployment options, requirements and configurations, see the Appendix: Deployment Configurations for Central Access Policies

Using Security Groups for Dynamic Access Control

You don't have to use user claims in order to implement Dynamic Access Control in your organization. In fact, you can use security groups with very minimal upgrade requirements to your current environment. You can use security groups in conjunction with central access policies and conditional expressions in Windows Server 2012. This way you can use Dynamic Access Control to limit access to specific data using the existing security group mechanism.

To use Dynamic Access Control with security groups you need the following in your environment:

- A Windows Server 2012 File Server
- A domain with a Windows Server 2012 schema (so that you can define central access policies)

Using security groups to limit access to data

Dynamic Access Control can help you limit access to data to the specific groups of people. The steps to achieve this are:

1. Tagging the data by marking the folders that contain confidential data

2. Configuring a Central Access Rule that specifies that only specific security groups can access data that are tagged in a specific way
3. Applying a Central Access Policy to the appropriate Windows Server 2012 File Servers in your organization

For a detailed walkthrough of using security groups, see the Dynamic Access Control Blog.

Using conditional expressions to reduce complexity of security groups

You can use Dynamic Access Control to considerably reduce the complexity of a combinatorial number of security groups (for example, from 2,000 to less than a 100) so that you can have a clear understanding of who can access what data and you are able to easily adjust access when people move between different roles in the company.

In a typical enterprise environment there are a large number of users, departments, security groups, and thus, a large number of access control lists (ACLs). Also, if a person moves from one department to the other, that involves updating a large number of security groups. This results in a lot of IT overhead and Dynamic Access Control can help with reducing the workload of IT admins for managing security group. The steps required for this are:

1. Tag all the folders with the appropriate values, for example, department, country, and sensitive.
2. Decide on the combinations of expressions to be used in in Windows ACL. For example, you would use **MemberOf (Spain_Security_Group) AND MemberOf (Finance_Security_Group) AND MemberOf(Sensitive_Security_Group)** to limit access to Spain's finance department sensitive information.
3. Create specific central access rules with these expressions that target certain security groups and specific folders on the files servers.

For a detailed walkthrough of using Dynamic Access Control to reduce the complexity of security groups, see the Dynamic Access Control Blog.

Using User Claims

As a rule of thumb, you should use user claims (vs. security group) when:

- You want to be able to use conditions such as User.Project = File.Project in your policy so that you can compress thousands of conditions to a simple expression (avoiding conditions such as File.Project=Cosmos AND User.Memberof(Cosmos_security_group)).
- The user attribute in Active Directory that you are sourcing the user claim from has the appropriate security setting on who or what can set that attribute.
- High integrity of the attribute value in Active Directory and the system that sets this value has operational procedures that take into consideration the use of that value for authorization decisions.
- No foreseeable changes to the values in the attribute. For example if the attribute is a department name and these often change due to re-organization, then it is not fit to be used as a user claim.

When using user claims, you should make sure that you have the appropriate Windows Server 2012 domain controller environment to support these user claims. The two main domains of interest are the User domain and the File Server domain.

- Both domains cannot include Windows Server 2003 (or earlier version) domain controllers.
- Have adequate number of Windows Server 2012 domain controllers in the user domain to support the number of Windows 8 clients deployed in that domain or alternatively set the Windows 8 clients to not require claims in which case having at least one Windows Server 2012 domain controller in the user domain will be sufficient (see more details below
- If the User domain and File Server domain are in a different forest, you need to:
  - Have two-way trust relations between the two forests.
  - Have all forest root domain controllers for both domains be Windows Server 2012

Operations to enable user claims

- Enable the domain controllers to provide claims and compound authentication on request

Enable the domain controllers to provide claims and compound authentication on request

To enable the domain controllers to provide claims and compound authentication on request

1. Open Group Policy Management and navigate to **Domain Controllers** OU in the domain.
2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **KDC**.
4. Select **KDC support for claims, compound authentication, and Kerberos armoring**.
5. Under **Claims, compound authentication for Dynamic Access Control and Kerberos armoring** options select **Supported**.
6. Click **OK**. Close Group Policy Management.

Important

Until the Kerberos Group Policy **Kerberos client support for claims, compound authentication and Kerberos armoring** is enabled on all Windows 8 devices, claims and compound authentication will not be requested and claims-based and device-based access will fail.

Warning

Windows 8 devices enabled to support requesting claims and compound authentication will fail authentication when a Windows Server 2012 DC cannot be found in domains configured to support claims and compound authentication.

Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies

Windows Server 2012 File servers have a group policy setting (On/Off/Automatic) that specifies whether it needs to get user claims for user tokens that do not carry claims. This setting is by default set to “automatic” which results in this group policy setting to be turned **On** if there is a central policy that contains user and/or device claims for that file server. If the file server contains local access policy (Discretionary ACLs) that uses user claims, you need to set this group policy to “On” so that the server knows to request claims on behalf of users that do not provide claims when accessing the server (For example, non-Windows 8 clients.

Using Device Claims and Device Security Groups

Windows 8 introduces the concept of compound identity that allows for the user token to include both the user details (ID, security groups and claims) as well as the device details (ID, security groups and claims) so that you can use these details in the authorization decision such as: Only specific users using a specific device can access highly sensitive data.

Much like user claims, device claims are sourced from the device object attribute in Active Directory (claim source) that contains the value of the claim.

Important
Device claims are supported in Windows 8clients only.

Considerations for using static device claims

To configure static device claims the administrator needs to:

- Apply the appropriate security on the device attribute
- Enable the claim definition – This operation will result in the claim to be available in the Kerberos ticket with the value sourced from the device attribute.
- Enable domain controller support for Dynamic Access Control – This operation will result in the directory service returning claims to the KDC and the KDC creating TGTs containing claims
- Apply the appropriate mechanisms to populate the device attribute in Active Directory.
- Set Windows 8 clients and Windows Server 2012 to request compound tokens.

Operations to enable device claims

- Enable the Windows 8 devices in domain to request claims and compound authentication
- Enable the Windows 8 devices to request claims and compound authentication using custom policy
- Enable the Windows 8 device to receive compound authentication

Enable the Windows 8 devices in domain to request claims and compound authentication

To enable the Windows 8 devices to request claims and compound authentication

1.

Note

Apply the group policy on the domain.

Open Group Policy Management and navigate to **Domain**.

2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **Kerberos**.
4. Select **Enable KDC support for claims, compound authentication, and Kerberos armoring**.
5. Click **OK**. Close Group Policy Management.

Enable the Windows 8 devices to request claims and compound authentication using custom policy

To enable Windows 8 devices to request claims and compound authentication using custom policy

1.

Note

Apply the group policy on the domain, site, OU or object scope.

Open Group Policy Management and navigate the appropriate policy.

2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **Kerberos**.
4. Select **Enable KDC support for claims, compound authentication, and Kerberos armoring**.
5. Click **OK**. Close Group Policy Management.

Enable the Windows 8 device to receive compound authentication

This group policy setting must be applied on the local group policy.

To enable the Windows 8 device to receive compound authentication

1. Open Local Group Policy Editor.
2. In the Local Group Policy window, expand to **Computer Configuration**, expand **Administrative Templates**, expand **System** and click **Kerberos**. Select **Enable** and **Support compound authentication**.
3. Click **OK** and close the Group Policy Management Editor.

Configuring central access policies with different options

As mentioned before, if you are using security groups, you do not have any specific domain consideration other than the need to upgrade your schema to the Windows Server 2012 schema.

If you decide to use user claims or you would like to use user/device compound identity, your deployment design for Windows Server 2012 can be any of the following three configurations:

- **Configuration 1**: Domains providing claims and compound authentication have all Windows Server 2012domain controllers.
- **Configuration 2**: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 DCs in all the file server sites.
- **Configuration 3**: Device-based access control needed, but cannot wait until all domain controllers can be upgraded.

Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs

This configuration provides complete infrastructure for claims and compound authentication for Dynamic Access Control and is the easiest to support. It requires that:

- If a cross-forest trusts exist, then root domain have all Windows Server 2012 domain controllers.
- Domains that provide claims and compound authentication have all Windows Server 2012 domain controllers
- All Windows 8 devices must be enabled to support requesting claims and compound authentication

Configuring forest root DCs

Simply upgrade all the domain controllers in the forest root to Windows Server 2012 and configure egress and ingress claims transformation.

Note

This will help ensure that claims are not lost from trusted forests. If pre-Windows Server 2012 domain controllers exist, those domain controllers will discard claims which will result in intermittent access control failures. Additionally pre-Windows Server 2012 domain controllers do not transform claims so the claims data is disclosed to all trusting forests

Configuring domains which provide claims and compound authentication

First identify which domains will be providing claims and compound authentication. These will be the account and resource domains. If all the domains in the environment have Windows Server

2012DCs then configure all.If claims-based access is required then claims will need to be provisioned.

Configure each domain which provides compound authentication and claims to DFL and Enable the DCs to support compound authentication and always provide claims.

Configuring devices to request claims and compound authentication

Update the default domain policy to Enable the Windows 8 devices in domain to request claims and compound authentication or Enable the Windows 8 devices to request claims and compound authentication using custom policy which applies to all the Windows 8 computer objects.

Note

This configuration that you set up to enable Windows 8 devices to request claims and compound authentication is ignored by versions of Windows which do not support it.

Configuring resources to receive compound authentication

If the resources are using CAP, then simply apply the CAP. Otherwise, Enable the Windows 8 device to receive compound authentication.

Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites

This configuration provides limited infrastructure for only user claims for Dynamic Access Control. It requires:

- If a cross-forest trusts exist, then root domain has all Windows Server 2012 DCs
- For each domain which provides user claims on request has Windows Server 2012 DCs in the sites with file servers and no Windows Server 2003 domain controllers.
- Enable all Windows 8 file servers to support requesting claims on the behalf of users.

Important

Authentication mechanism assurance-based universal groups and device-based access control are incompatible with file server retrieval of user claims.

Configuring forest root DCs

Simply upgrade all the domain controllers in the forest root to Windows Server 2012 and configure egress and ingress claims transformation.

Note

This will help ensure that claims are not lost from trusted forests. If pre-Windows Server 2012 domain controllers exist, those domain controllers will discard claims which will result in intermittent access control failures. Additionally pre-Windows Server 2012

domain controllers do not transform claims so the claims data is disclosed to all trusting forests

Configuring domains which provides claims and compound authentication

1. First identify which domains will be providing claims and compound authentication. These will be the account and resource domains.
2. Configure and provision claims as explained in Deploy a Central Access Policy (Demonstration Steps)
3. Configure each domain to Enable the domain controllers to provide claims and compound authentication on request.

Configuring file servers to request claims on the behalf of users

If the resources are using a central access policy, then simply apply the policy. Otherwise, Enable the Windows 8 devices to request claims and compound authentication using custom policy which applies to Windows 8 file server computer objects.

Note
This setting is ignored by versions of Windows which do not support it.

Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded

This configuration will be unique to your environment and can be difficult to support when Windows 8 devices have different configurations.

General requirements for all environments:

- If across-forest trusts exist, then root domain must have all Windows Server 2012domain controllers
- For each domain which provides claims and compound authentication on request, there cannot be Windows Server 2003 domain controllers
- For resources using device-based access control, receiving compound authentication must be enabled unless a central access policy is being used.

Considerations for using smartcards for Central Access Polices

In Windows 7 and Windows 8 clients, smart card logon is mapped through security groups. In order to create this mapping, the domain administrator needs to:

- Create a security group that will represent the smart card logon.
- Create a mapping from the smart card certificate OID to and smart card security group.
- Domain controllers need to have Windows Server 2008 R2 Domain Functional Level

Note

This group membership will be lost if the server performs S4U2Self on behalf of the user

Best Practices for Deploying Central Access Policies

The following sections provide additional guidance for best practices for delegation of administration, setting up exception mechanisms, and so on.

Delegating of administration for Dynamic Access Control

Ideally, you should delegate permissions for all Dynamic Access Control containers in the forest root of your domain controller. You would need to grant read-write permissions for to specific security groups that have access to certain objects. For example, you can create universal groups like:

- DAC Claim Admins
- DAC Resource Property Admins
- DAC Central Access Rule Admins
- DAC Central Access Policy Admins

You can then delegate the corresponding rights to these groups.

Following are the built-in Dynamic Access Control containers in Windows Server 2012. To access these containers, browse to configuration partition-> system->services->claims configuration.

- Claim Types
- Resource Properties
- Central Access Rules
- Central Access Policies

You can set up delegation of policies from Active Directory Administrative Center (ADAC).

To set up delegation of permissions from ADAC

1. From Server Manager, on the **Tools** menu, select **Active Directory Administrative Center**.
2. Select Dynamic Access Control on the left pane, and select the container that has the object that you want to delegate permissions to. For example, select **Central Access Policies** and select a policy from the list.
3. Right-click and select **Properties**. In the Policy Properties window, select the **Extensions** tab.
4. Select the **Security** tab. Click **Advanced**. In the Advanced Security Settings dialog box, click **Add**. In the Permission entry dialog box, click **Select a Principal** and type the name of the security group to which you want to grant access to. Click OK.
5. In the Permission entry dialog box, select the permissions you want to grant for the group. Click **OK** three times.

Exception Mechanisms for Planning Central Access Policies

Exceptions to the common access rules are a key component of every access policy and in particular a central access policy. For example, if a central policy for the organization protects access to high business impact data (HBI) so that only full-time employees can access this data, what happens if a vendor in the finance department requires access to financial information that is considered to be HBI so that he can do his job?

There are several valid answers ranging from using one security group in the central access policy to grant access to HBI data (which would then allow the vendor in this example to access both financial HBI data and other HBI data such as HR or engineering) to having multiple security groups for exceptions to having per file- or user-specific exception mechanisms.

Exception Mechanism	Central Access Policy scenario	Description	Pros	Cons
Security group	Any access rule where we are OK with providing access to all the information covered by the central access policy when granting the exception Example is a central access rule for department data that implements a safety net over ALL of the departments data (e.g.: Resource.Department=Finance)	One security group managed by the department IT or Information Security personal with a delegation to allow department senior personal to add / remove (manage) users from the group. Content owners from the department will be able to manage the exception group so that they control	Simple to implement Provides a good delegation model with existing tools to manage groups	Exception is wide and grants access to information that might not be relevant for this specific user

Exception Mechanism	Central Access Policy scenario	Description	Pros	Cons
		who can access information to the department data		
Multiple security groups	Central access rule for access to a range of countries with a finite (e.g.: <20) number of permutations where we would like to have each one controlled by a different content owner Example is a central access rule for access to country data Resource.Country = User.Country	A security group for each of the countries. Each security group is managed by the content owners of that country	Simple to implement Provides a good delegation model with existing tools to manage groups	The more groups there are, the more complex the access conditions becomes. Each time a new instance is added (for example, a new country), the administrator needs to create a new group to manage the exception for it

Tools for Deployment

Data Classification Toolkithttp://go.microsoft.com/fwlink/?LinkId=217654: The Data Classification Toolkit for Windows Server 2012 is designed to help organizations identify, classify, and protect data on their file servers. The out-of-the-box classification and rule examples help organizations build and deploy their policies to protect critical information in a cost-effective manner.The toolkit supports both Windows Server 2012 file servers and Windows Server 2008 R2 file servers. In addition to configuring file classification infrastructure, the latest version of the toolkit allows you to manage a Central Access Policy across the file servers in your organization. It provides tools for you to provision user and device claim values, and manage Central Access Policy across the forest to help simplify the configuration process of Dynamic Access Controls in Windows Server 2012. The toolkit also provides a new report template that you can use to review existing Central Access Policy on file shares.

Appendix: Deployment Configurations for Central Access Policies

Deployment Option	Requirements and Configuration Options	Guidance
Using Security Groups	• Windows Server 2012 File Server • A domain with a Windows Server 2012 schema	Using Security Groups for Dynamic Access Control
Using User Claims	• Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites For cross-forest claims deployment, you need a two-way trust been the user domain and a resource domain. For more information, see Deploy Claims Across Forests • Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs	Using User Claims
Using Device Claims	• Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites • Configuration 1: Domains providing claims and compound authentication have all Windows Server	Using Device Claims and Device Security Groups

Deployment Option	Requirements and Configuration Options	Guidance
	2012 DCs • Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded	

Deploy a Central Access Policy (Demonstration Steps)

In this scenario, the finance department security operations is working with central information security to specify the need for a central access policy so that they can protect archived finance information stored on file servers. The archived finance information from each country can be accessed as read-only by finance employees from the same country. A central finance admin group can access the finance information from all countries.

Deploying a central access policy includes the following phases:

Phase	Description
Plan: Identify the need for policy and the configuration required for deployment	Identify the need for a policy and the configuration required for deployment.
Implement: Configure the components and policy	Configure the components and policy.
Deploy the central access policy	Deploy the policy.
Maintain: Change and stage the policy	Policy changes and staging

Set up a test environment

Before you begin, you need to set up lab to test this scenario. The steps for setting up the lab are explained in detail in Appendix B: Setting Up the Test Environment.

Plan: Identify the need for policy and the configuration required for deployment

This section provides the high-level series of steps that aid in the planning phase of your deployment.

<table>
<tr><th></th><th>Step</th><th>Example</th></tr>
<tr><td>1.1</td><td>Business determines that a central access policy is needed</td><td>To protect finance information that is stored on file servers, the finance department security operations is working with central information security to specify the need for a central access policy.</td></tr>
<tr><td>1.2</td><td>Express the access policy</td><td>Finance documents should only be read by members of the Finance department. Members of the Finance department should only access documents in their own country. Only Finance Administrators should have write-access. An exception will be allowed for members of the FinanceException group. This group will have Read access.</td></tr>
<tr><td>1.3</td><td>Express the access policy in Windows Server 2012 constructs</td><td><table>
<tr><td>Targeting:</td><td>• Resource.Department Contains Finance</td></tr>
<tr><td>Access rules:</td><td>• Allow read User.Country=Resource.Country AND User.department = Resource.Department
• Allow Full control User.MemberOf(FinanceAdmin)</td></tr>
<tr><td>Exception:</td><td>Allow read memberOf(FinanceException)</td></tr>
</table></td></tr>
<tr><td>1.4</td><td>Determine the file properties required for the policy</td><td>Tag files with:
• Department
• Country</td></tr>
<tr><td>1.5</td><td>Determine the claim types and groups required for the policy</td><td>Claim types:
• Country
• Department</td></tr>
</table>

	Step	Example
		User groups: • FinanceAdmin • FinanceException
1.6	Determine the servers on which to apply this policy	Apply the policy on all finance file servers.

Implement: Configure the components and policy

This section presents an example that deploys a central access policy for finance documents.

No	Step	Example
2.1	Create claim types	Create the following claim types: • Department • Country
2.2	Create resource properties	Create and enable the following resource properties: • Department • Country
2.3	Configure a central access rule	Create a Finance Documents rule that includes the policy determined in the previous section.
2.4	Configure a central access policy (CAP)	Create a CAP called Finance Policy and add the Finance Documents rule to that CAP.
2.5	Target central access policy to the file servers	Publish the Finance Policy CAP to the file servers.
2.6	Enable KDC Support for claims, compound authentication and Kerberos armoring.	Enable KDC Support for claims, compound authentication and Kerberos armoring for contoso.com.

In the following procedure, you create two claim types: Country and Department.

To create claim types

1. Open Server DC1 in Hyper-V Manager and log on as contoso\administrator, with the

password **pass@word1**.

2. Open Active Directory Administrative Center.
3. Click the **Tree View icon**, expand **Dynamic Access Control**, and then select **Claim Types**.

   Right-click **Claim Types**, click **New**, and then click **Claim Type**.

   Tip

   You can also open a **Create Claim Type:** window from the **Tasks** pane. On the **Tasks** pane, click **New**, and then click **Claim Type**.
4. In the **Source Attribute** list, scroll down the list of attributes, and click **department**. This should populate the **Display name** field with **department**. Click **OK**.
5. In **Tasks** pane, click **New**, and then click **Claim Type**.
6. In the **Source Attribute** list, scroll down the list of attributes, and then click the **c** attribute (Country-Name). In the **Display name** field, type **country**.
7. In the **Suggested Values** section, select **The following values are suggested:**, and then click **Add**.
8. In the **Value** and **Display name** fields, type **US**, and then click **OK**.
9. Repeat the above step. In the **Add a suggest value** dialog box, type **JP** in the **Value** and **Display name** fields, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADClaimType country -SourceAttribute c -SuggestedValues:@((New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("US","US","")), (New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("JP","JP","")))
New-ADClaimType department -SourceAttribute department
```

Tip

You can use the Windows PowerShell History Viewer in Active Directory Administrative Center to look up the Windows PowerShell cmdlets for each procedure you perform in Active Directory Administrative Center. For more information, see Windows PowerShell History Viewer

The next step is to create resource properties. In the following procedure you create a resource property that is automatically added to the Global Resource Properties list on the domain controller, so that it is available to the file server.

To create and enable pre-created resource properties

1. In the left pane of Active Directory Administrative Center, click **Tree View**. Expand **Dynamic Access Control**, and then select **Resource Properties**.
2. Right-click **Resource Properties**, click **New**, and then click **Reference Resource Property**.

   Tip
   You can also choose a resource property from the **Tasks** pane. Click **New** and then click **Reference Resource Property**.
3. In **Select a claim type to share its suggested values list**, click **country**.
4. In the **Display name** field, type **country**, and then click **OK**.
5. Double-click the **Resource Properties** list, scroll down to the **Department** resource property. Right-click, and then click **Enable**. This will enable the built-in **Department** resource property.
6. In the **Resource Properties** list on the navigation pane of the Active Directory Administrative Center, you will now have two enabled resource properties:
   - Country
   - Department

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADResourceProperty Country -IsSecured $true -ResourcePropertyValueType MS-DS-MultivaluedChoice -SharesValuesWith country

Set-ADResourceProperty Department_MS -Enabled $true

Add-ADResourcePropertyListMember "Global Resource Property List" -Members Country

Add-ADResourcePropertyListMember "Global Resource Property List" -Members Department_MS
```

The next step is to create central access rules that define who can access resources. In this scenario the business rules are:

- Finance documents can be read only by members of the Finance department.
- Members of the Finance department can access only documents in their own country.
- Only Finance Administrators can have Write access.
- We will allow an exception for members of the FinanceException group. This group will have Read access.
- The administrator and document owner will still have full access.

Or to express the rules with Windows Server 2012 constructs:

Targeting: Resource.Department Contains Finance

Access Rules:

Allow Read User.Country=Resource.Country AND User.department = Resource.Department

Allow Full control User.MemberOf(FinanceAdmin)

Allow Read User.MemberOf(FinanceException)

To create a central access rule

1. In the left pane of the Active Directory Administrative Center, click **Tree View**, select **Dynamic Access Control**, and then click **Central Access Rules**.
2. Right-click **Central Access Rules**, click **New**, and then click **Central Access Rule**.
3. In the **Name** field, type **Finance Documents Rule**.
4. In the **Target Resources** section, click **Edit**, and in the **Central Access Rule** dialog box, click **Add a condition**. Add the following condition:
   [**Resource**] [**Department**] [**Equals**] [**Value**] [**Finance**], and then click **OK**.
5. In the **Permissions** section, select **Use following permissions as current permissions**, click **Edit**, and in the **Advanced Security Settings for Permissions** dialog box click **Add**.

   Note
   Use the following permissions as proposed permissions option lets you create the policy in staging. For more information on how to do this refer to the Maintain: Change and stage the policy section in this topic.
6. In the **Permission entry for Permissions** dialog box, click **Select a principal**, type **Authenticated Users**, and then click **OK**.
7. In the **Permission Entry for Permissions** dialog box, click **Add a condition**, and add the following conditions:
   [**User**] [**country**] [**Any of**] [**Resource**] [**country**]
   Click **Add a condition**.
   [**And**]
   Click [**User**] [**Department**] [**Any of**] [**Resource**] [**Department**]. Set the **Permissions** to **Read**.
8. Click **OK**, and then click **Add**. Click **Select a principal**, type **FinanceAdmin**, and then click **OK**.
9. Select the **Modify, Read and Execute, Read, Write** permissions, and then click **OK**.
10. Click **Add**, click **Select a principal**, type **FinanceException**, and then click **OK**. Select the permissions to be **Read** and **Read and Execute**.
11. Click **OK** three times to finish and return to Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$countryClaimType = Get-ADClaimType country
$departmentClaimType = Get-ADClaimType department
$countryResourceProperty = Get-ADResourceProperty Country
```

```
$departmentResourceProperty = Get-ADResourceProperty
Department
$currentAcl =
"O:SYG:SYD:AR(A;;FA;;;OW)(A;;FA;;;BA)(A;;0x1200a9;;;S-1-5-21-
1787166779-1215870801-2157059049-1113)(A;;0x1301bf;;;S-1-5-
21-1787166779-1215870801-2157059049-
1112)(A;;FA;;;SY)(XA;;0x1200a9;;;AU;((@USER." +
$countryClaimType.Name + " Any_of @RESOURCE." +
$countryResourceProperty.Name + ") && (@USER." +
$departmentClaimType.Name + " Any_of @RESOURCE." +
$departmentResourceProperty.Name + ")))"
$resourceCondition = "(@RESOURCE." +
$departmentResourceProperty.Name + " Contains {`"Finance`"})"
New-ADCentralAccessRule "Finance Documents Rule" -CurrentAcl
$currentAcl -ResourceCondition $resourceCondition
```

Important

In the above cmdlet example, the security identifiers (SIDs) for the group FinanceAdmin and users is determined at creation time and will be different in your example. For example, the provided SID value (S-1-5-21-1787166779-1215870801-2157059049-1113) for the FinanceAdmins needs to be replaced with the actual SID for the FinanceAdmin group that you would need to create in your deployment. You can use Windows PowerShell to look up the SID value of this group, assign that value to a variable, and then use the variable here. For more information, see Windows PowerShell Tip: Working with SIDs.

You should now have a central access rule that allows people to access documents from the same country and the same department. The rule allows the FinanceAdmin group to edit the documents, and it allows the FinanceException group to read the documents. This rule targets only documents classified as Finance.

To add a central access rule to a central access policy

1. In the left pane of the Active Directory Administrative Center, click **Dynamic Access Control**, and then click **Central Access Policies**.
2. In the **Tasks** pane, click **New**, and then click **Central Access Policy**.
3. In **Create Central Access Policy:**, type **Finance Policy** in the **Name** box.
4. In **Member central access rules**, click **Add**.
5. Double-click the **Finance Documents Rule** to the add it to the **Add the following central access rules** list , and then click **OK**.

6. Click **OK** to finish. You should now have a central access policy named Finance Policy.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADCentralAccessPolicy "Finance Policy" Add-
ADCentralAccessPolicyMember
-Identity "Finance Policy"
-Member "Finance Documents Rule"
```

To apply the central access policy across file servers by using Group Policy

1. On the **Start** screen, in the **Search** box, type **Group Policy Management**. Double-click **Group Policy Management**.

   Tip
   If the **Show Administrative tools** setting is disabled, the **Administrative Tools** folder and its contents will not appear in the **Settings** results.

   Tip
   In your production environment, you should create a File Server Organization Unit (OU) and add all your file servers to this OU, to which you want to apply this policy. You can then create a group policy and add this OU to that policy..

2. In this step, you edit the group policy object you created in Build the domain controller section in the Test Environment to include the central access policy that you created. In the Group Policy Management Editor, navigate to and select the organizational unit in the domain (contoso.com in this example): **Group Policy Management**, **Forest: contoso.com**, **Domains**, **contoso.com**, **Contoso**, **FileServerOU**.
3. Right-click **FlexibleAccessGPO**, and then click **Edit**.
4. In the Group Policy Management Editor window, navigate to **Computer Configuration**, expand **Policies**, expand **Windows Settings**, and click **Security Settings**.
5. Expand **File System**, right-click **Central Access Policy**, and then click **Manage Central access policies**.
6. In the **Central Access Policies Configuration** dialog box, add **Finance Policy**, and then click **OK**.
7. Scroll down to **Advanced Audit Policy Configuration**, and expand it.
8. Expand **Audit Policies**, and select **Object Access**.
9. Double-click **Audit Central Access Policy Staging**. Select all three check boxes and then click **OK**. This step allows the system to receive audit events related to Central Access Staging Policies.
10. Double-click **Audit File System Properties**. Select all three check boxes then click **OK**.

11. Close the Group Policy Management Editor. You have now included the central access policy to the Group Policy.

For a domain's domain controllers to provide claims or device authorization data, the domain controllers need to be configured to support dynamic access control.

To enable support for claims and compound authentication for contoso.com

1. Open Group Policy Management, click **contoso.com**, and then click **Domain Controllers**.
2. Right-click **Default Domain Controllers Policy**, and then click **Edit**.
3. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.
4. Double-click **KDC Support for claims, compound authentication and Kerberos armoring**. In the **KDC Support for claims, compound authentication and Kerberos armoring** dialog box, click **Enabled** and select **Supported** from the **Options** drop-down list. (You need to enable this setting to use user claims in central access policies.)
5. Close **Group Policy Management**.
6. Open a command prompt and type `gpupdate /force`.

Deploy the central access policy

	Step	Example
3.1	Assign the CAP to the appropriate shared folders on the file server.	Assign the central access policy to the appropriate shared folder on the file server.
3.2	Verify that access is appropriately configured.	Check the access for users from different countries and departments.

In this step you will assign the central access policy to a file server. You will log onto a file server that is receiving the central access policy that you created the previous steps and assign the policy to a shared folder.

To assign a central access policy to a file server

1. In Hyper-V Manager, connect to server FILE1. Log on to the server by using contoso\administrator with the password: **pass@word1**.
2. Open an elevated command prompt and type: **gpupdate /force**. This ensures that your Group Policy changes take effect on your server.
3. You also need to refresh the Global Resource Properties from Active Directory. Open an

elevated Windows PowerShell window and type `Update-FSRMClassificationpropertyDefinition`. Click ENTER, and then close Windows PowerShell.

Tip

You can also refresh the Global Resource Properties by logging on to the file server. To refresh the Global Resource Properties from the file server, do the following

a. Logon to File Server FILE1 as contoso\administrator, using the password **pass@word1**.

b. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.

c. In the File Server Resource Manager, click **File Classification Management** , right-click **Classification Properties** and then click **Refresh**.

4. Open Windows Explorer, and in the left pane, click drive D. Right-click the **Finance Documents** folder, and click **Properties**.
5. Click the **Classification** tab, click **Country**, and then select **US** in the **Value** field.
6. Click **Department**, then select **Finance** in the **Value** field and then click **Apply**.

Note

Remember that the central access policy was configured to target files for the Department of Finance. The previous steps mark all documents in the folder with the Country and Department attributes.

7. Click the **Security** tab, and then click **Advanced**. Click the **Central Policy** tab.

8. Click **Change**, select **Finance Policy** from the drop-down menu, and then click **Apply**. You can see the **Finance Documents Rule** listed in the policy. Expand the item to view all of the permissions that you set when you created the rule in Active Directory.
9. Click **OK** to return to Windows Explorer.

In the next step, you ensure that access is appropriately configured. User accounts need to have the appropriate Department attribute set (set this using Active Directory Administrative Center). The simplest way to view the effective results of the new policy is to use the **Effective Access** tab in Windows Explorer. The **Effective Access** tab shows the access rights for a given user account.

To examine the access for various users

1. In Hyper-V Manager, connect to server FILE1. Log on to the server by using contoso\administrator. Navigate to D:\ in Windows Explorer. Right-click the **Finance Documents** folder, and then click **Properties**.
2. Click the **Security** tab, click **Advanced**, and then click the **Effective Access** tab.
3. To examine the permissions for a user, click **Select a user**, type the user's name, and

then click **View effective access** to see the effective access rights. For example:

- Myriam Delesalle (MDelesalle) is in the Finance department and should have Read access to the folder.
- Miles Reid (MReid) is a member of the FinanceAdmin group and should have Modify access to the folder.
- Esther Valle (EValle) is not in the Finance department; however, she is a member of the FinanceException group and should have Read access.
- Maira Wenzel (MWenzel) is not in the Finance department and is not a member of either the FinanceAdmin or FinanceException group. She should not have any access to the folder.

Notice that the last column named **Access limited by** in the effective access window. This column tells you which gates are effecting the person's permissions. In this case, the Share and NTFS permissions allow all users full control. However, the central access policy restricts access based on the rules you configured earlier.

Maintain: Change and stage the policy

Number	Step	Example
4.1	Configure Device Claims for Clients	Set the group policy setting to enable device claims
4.2	Enable a claim for devices.	Enable the country claim type for devices.
4.3	Add a staging policy to the existing central access rule that you would like to modify.	Modify the Finance Documents Rule to add a staging policy.
4.4	View the results of the staging policy.	Check for Ester Velle's permissions.

To set up group policy setting to enable claims for devices

1. Log on to DC1, open Group Policy Management, click **contoso.com**, click **Default Domain Policy**, right-click and select **Edit**.
2. In the Group Policy Management Editor window, navigate to **Computer Configuration**, **Policies**, **Administrative Templates**, **System**, **Kerberos**.
3. Select **Kerberos client support for claims, compound authentication and Kerberos armoring** and click **Enable**.

To enable a claim for devices

1. Open Server DC1 in Hyper-V Manager and log on as contoso\Administrator, with the password **pass@word1**.
2. From the **Tools** menu, open Active Directory Administrative Center.
3. Click **Tree View**, expand **Dynamic Access Control**, double-click **Claim Types**, and double-click the **country** claim.
4. In **Claims of this type can be issued for the following classes**, select the **Computer** check box. Click **OK**.
   Both the **User** and **Computer** check boxes should now be selected. The country claim can now be used with devices in addition to users.

The next step is to create a staging policy rule. Staging policies can be used to monitor the effects of a new policy entry before you enable it. In the following step, you will create a staging policy entry and monitor the effect on your shared folder.

To create a staging policy rule and add it to the central access policy

1. Open Server DC1 in Hyper-V Manager and log on as contoso\Administrator, with the password **pass@word1**.
2. Open Active Directory Administrative Center.
3. Click **Tree View**, expand **Dynamic Access Control**, and select **Central Access Rules**.
4. Right-click **Finance Documents Rule**, and then click **Properties**.
5. In the **Proposed Permissions** section, select the **Enable permission staging configuration** check box, click **Edit**, and then click **Add**. In the **Permission Entry for Proposed Permissions** window, click the **Select a Principal** link, type **Authenticated Users**, and then click **OK**.
6. Click the **Add a condition** link and add the following condition:
   [**User**] [**country**] [**Any of**] [**Resource**] [**Country**].
7. Click **Add a condition** again, and add the following condition: [**And**]
   [**Device**] [**country**] [**Any of**] [**Resource**] [**Country**]
8. Click **Add a condition** again, and add the following condition. [And]
   [**User**] [**Group**] [**Member of any**] [**Value**](**FinanceException**)
9. To set the FinanceException, group, click **Add items** and in the **Select User, Computer, Service Account, or Group** window, type **FinanceException**.
10. Click **Permissions**, select **Full Control**, and click **OK**.
11. In the Advance Security Settings for Proposed Permissions window, select **FinanceException** and click **Remove**.
12. Click **OK** two times to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADCentralAccessRule
-Identity:
"CN=FinanceDocumentsRule,CN=CentralAccessRules,CN=ClaimsConfiguration,CN=Configuration,DC
=Contoso.com"
-ProposedAcl: "O:SYG:SYD:AR(A;;FA;;;BA)(A;;FA;;;SY)(A;;0x1301bf;;;S-1-21=1426421603-
1057776020-1604)"
-Server: "WIN-2R92NN8VKFP.Contoso.com"
```

Note

In the above cmdlet example, the Server value reflects the Server in the test lab environment. You can use the Windows PowerShell History Viewer to look up the Windows PowerShell cmdlets for each procedure you perform in Active Directory Administrative Center. For more information, see Windows PowerShell History Viewer

In this proposed permissions set, members of the FinanceException group will have Full Access to files from their own country when they access them through a device from the same country as the document. Audit entries are available in the File Servers security log when someone from the Finance department attempts to access files. However, security settings are not enforced until the policy is promoted from staging.

In the next procedure, you verify the results of the staging policy. You access the shared folder with a user name that has permissions based on the current rule. Esther Valle (EValle) is a member of FinanceException, and she currently has Read rights. According to our staging policy, EValle should not have any rights.

To verify the results of the staging policy

1. Connect to the File Server FILE1 in Hyper-V Manager and log on as contoso\administrator, with the password **pass@word1**.
2. Open a Command Prompt window and type **gpupdate /force**. This ensures that your Group Policy changes will take effect on your server.
3. In Hyper-V Manager, connect to server CLIENT1. Log off the user who is currently logged on. Restart the virtual machine, CLIENT1. Then log on to the computer by using contoso\EValle pass@word1.
4. Double-click the desktop shortcut to \\FILE1\Finance Documents. EValle should still have access to the files. Switch back to FILE1.
5. Open **Event Viewer** from the shortcut on the desktop. Expand **Windows Logs**, and then select **Security**. Open the entries with **Event ID 4818** under the **Central Access Policy Staging** task category. You will see that EValle was allowed access; however, according to the staging policy, the user would have been denied access.

Next Steps

If you have a central server management system such as System Center Operations Manager, you can also configuring monitoring for events. This allows Administrators to monitor the effects of central access policies before enforcing them.

Deploy Claims Across Forests

In Windows Server 2012, a claim type is an assertion about the object with which it's associated. Claim types are defined per forest in Active Directory. There are many scenarios where a security principal may need to traverse a trust boundary to access resources in a trusted forest. Cross-forest claims transformation in Windows Server 2012 enables you to transform egress and ingress claims that traverse forests so that the claims are recognized and accepted in the trusting and trusted forests. Some of the real-world scenarios for transformation of claims are:

- Trusting forests can use claim transformation as a guard against elevation of privilege by filtering the incoming claims with specific values.

  Trusting forests can also issue claims for principals coming over a trust boundary if the trusted forest does not support or issue any claims.
- Trusted forests can use claim transformation to prevent certain claim types and claims with certain values from going out to the trusting forest.
- You can also use claim transformation to map different claim types between trusting and trusted forests. This can be used to generalize the claim-type, the claim value, or both. Without this, you need to standardize the data between the forests before you can use the claims. Generalizing claims between the trusting and trusted forests reduces the IT costs.

Claim transformation rules

The transformation rule language syntax divides a single rule into two main parts: a series of condition statements and the issue statement. Each condition statement has two subcomponents: the claim identifier and the condition. The issue statement contains keywords, delimiters, and an issue expression. The condition statement optionally begins with a claim identifier variable, which represents the matched input claim. The condition checks for the expression. If the input claim does not match the condition, then the transformation engine ignores the issue statement and evaluates the next input claim against the transformation rule. If all conditions match the input claim, it processes the issue statement.

For detailed information on claim rules language, see Claims Transformation Rules Language.

Linking claim transformation policies to forests

There are two components involved in setting up claim transformation policies: claim transformation policy objects and the transformation link. The policy objects live in the

configuration naming context in a forest, and they contain mapping information for the claims. The link specifies which trusting and trusted forests the mapping applies to.

It is important to understand if the forest is the trusting or trusted forest because this is basis for linking transformation policy objects. For example, the trusted forest is the forest that contains user accounts that require access. The trusting forest is the forest that contains resources that you want to give users access to. Claims travel in the same direction as the security principal that requires access. For example, if there is a one-way trust from the contoso.com forest to the adatum.com forest, the claims will flow from adatum.com to contoso.com, which allows users from adatum.com to access resources in contoso.com.

By default, a trusted forest allows all outgoing claims to pass, and a trusting forest drops all incoming claims that it receives.

In this scenario

The following guidance is available for this scenario:

- Deploy Claims Across Forests (Demonstration Steps)
- Claims Transformation Rules Language

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services	In this scenario, you are required to set up two Active Directory forests with a two-way trust. You have claims in both forests. You also set central access policies on the trusting forest where the resources reside.
File and Storage Services role	In this scenario, the data classification is applied to the resources on the file servers. The central access policy is applied to the folder where you want to grant user access. After transformation, the claim grants user access to resources based on the central access policy that is applied to the folder on the file server.

Deploy Claims Across Forests (Demonstration Steps)

In this topic, we'll cover a basic scenario that explains how to configure claims transformations between trusting and trusted forests. You will learn how claims transformation policy objects can be created and linked to the trust on the trusting forest and the trusted forest. You will then validate the scenario.

Scenario overview

Adatum Corporation provides financial services to Contoso, Ltd. Each quarter, Adatum accountants copy their account spreadsheets to a folder on a file server located at Contoso, Ltd. There is a two-way trust set up from Contoso to Adatum. Contoso, Ltd. wants to protect the share so that only Adatum employees can access the remote share.

In this scenario:

1. Set up the prerequisites and the test environment
2. Set up claims transformation on trusted forest (Adatum)
3. Set up claims transformation in the trusting forest (Contoso)
4. Validate the scenario

Set up the prerequisites and the test environment

The test configuration involves setting up two forests: Adatum Corporation and Contoso, Ltd, and having a two-way trust between Contoso and Adatum. "adatum.com" is the trusted forest and "contoso.com" is the trusting forest.

The claims transformation scenario demonstrates transformation of a claim in the trusted forest to a claim in the trusting forest. To do this, you need to set up a new forest called adatum.com and populate the forest with a test user with a company value of "Adatum". You then have to set up a two-way trust between contoso.com and adatum.com.

Important

When setting up the Contoso and Adatum forests, you must ensure that both the root domains are at the Windows Server 2012 Domain Functional Level for claims transformation to work.

You need to set up the following for the lab. These procedures are explained in detail in Appendix B: Setting Up the Test Environment

You need to implement the following procedures to set up the lab for this scenario:

1. **Set Adatum as trusted forest to Contoso**
2. Create the "Company" claim type on Contoso

3. Enable the “Company” resource property on Contoso
4. Create the central access rule
5. Create the central access policy
6. Publish the new policy through Group Policy
7. Create the Earnings folder on the file server
8. Set classification and apply the central access policy on the new folder

Use the following information to complete this scenario:

Objects	Details
Users	Jeff Low, Contoso
User claims on Adatum and Contoso	ID: ad://ext/Company:ContosoAdatum, Source attribute: company Suggested values: Contoso, Adatum **Important** You must set the ID on the “Company” claim type on both Contoso and Adatum to be the same for the claims transformation to work.
Central access rule on Contoso	AdatumEmployeeAccessRule
Central access policy on Contoso	Adatum Only Access Policy
Claims Transformation policies on Adatum and Contoso	DenyAllExcept Company
File folder on Contoso	D:\EARNINGS

Set up claims transformation on trusted forest (Adatum)

In this step you create a transformation policy in Adatum to deny all claims except “Company” to pass to Contoso.

The Active Directory module for Windows PowerShell provides the **DenyAllExcept** argument, which drops everything except the specified claims in the transformation policy.

To set up a claims transformation, you need to create a claims transformation policy and link it between the trusted and trusting forests.

Create a claims transformation policy in Adatum

To create a transformation policy Adatum to deny all claims except “Company”

1. Sign in to the domain controller, adatum.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell, and type the following:

```
New-ADClaimTransformPolicy `
-Description:"Claims transformation policy to deny all claims
except Company"`
-Name:"DenyAllClaimsExceptCompanyPolicy" `
-DenyAllExcept:company `
-Server:"adatum.com" `
```

Set a claims transformation link on Adatum’s trust domain object

In this step, you apply the newly created claims transformation policy on Adatum’s trust domain object for Contoso.

To apply the claims transformation policy

1. Sign in to the domain controller, adatum.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell, and type the following:

```
Set-ADClaimTransformLink `
-Identity:"contoso.com" `
-Policy:"DenyAllClaimsExceptCompanyPolicy" `
-TrustRole:Trusted `
```

Set up claims transformation in the trusting forest (Contoso)

In this step you create a claims transformation policy in Contoso (the trusting forest) to deny all claims except “Company.” You need to create a claims transformation policy and link it to the forest trust.

Create a claims transformation policy in Contoso

To create a transformation policy Adatum to deny all except “Company”

1. Sign in to the domain controller, contoso.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell and type the following:

```
New-ADClaimTransformPolicy `
-Description:"Claims transformation policy to deny all claims
except company" `
-Name:"DenyAllClaimsExceptCompanyPolicy" `
-DenyAllExcept:company `
-Server:"contoso.com" `
```

Set a claims transformation link on Contoso’s trust domain object

In this step, you apply the newly created claims transformation policy on the contoso.com trust domain object for Adatum to allow “Company” be passed through to contoso.com. The trust domain object is named adatum.com.

To set the claims transformation policy

1. Sign in to the domain controller, contoso.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell and type the following:

```
Set-ADClaimTransformLink
-Identity:"adatum.com" `
-Policy:"DenyAllClaimsExceptCompanyPolicy" `
-TrustRole:Trusting `
```

Validate the scenario

In this step you try to access the D:\EARNINGS folder that was set up on the file server FILE1 to validate that the user has access to the shared folder.

To ensure that the Adatum user can access the shared folder

1. Sign in to the Client machine, CLIENT1 as Jeff Low with the password **pass@word1**.
2. Browse to the folder \\FILE1.contoso.com\Earnings.
3. Jeff Low should be able to access the folder.

Additional scenarios for claims transformation policies

Following is a list of additional common cases in claims transformation.

Scenario	Policy

Allow all claims that come from Adatum to go through to Contoso Adatum	`New-ADClaimTransformPolicy ` `-Description:"Claims transformation policy to allow all claims" ` `-Name:"AllowAllClaimsPolicy" ` `-AllowAll ` `-Server:"contoso.com" `
	
	`Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"AllowAllClaimsPolicy" ` `-TrustRole:Trusting ` `-Server:"contoso.com" `

Deny all claims that come from Adatum to go through to Contoso Adatum	`New-ADClaimTransformPolicy ` `-Description:"Claims transformation policy to deny all claims" ` `-Name:"DenyAllClaimsPolicy" ` `-DenyAll ` `-Server:"contoso.com" `

Scenario	Policy
	`Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"DenyAllClaimsPolicy" ` `-TrustRole:Trusting ` `-Server:"contoso.com"`
Allow all claims that come from Adatum except “Company” and “Department” to go through to Contoso Adatum	`New-ADClaimTransformationPolicy ` `-Description:"Claims transformation policy to allow all claims except company and department" ` `-Name:"AllowAllClaimsExceptCompanyAndDepartmentPolicy" ` `-AllowAllExcept:company,department ` `-Server:"contoso.com" ` `Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"AllowAllClaimsExceptCompanyAndDepartmentPolicy" ` `-TrustRole:Trusting ` `-Server:"contoso.com" `

See also

- For a list of all Windows PowerShell cmdlets that are available for claims transformation, see Active Directory PowerShell Cmdlet Reference.
- For advanced tasks that involve export and import of DAC configuration information between two forests, use the Dynamic Access Control PowerShell Reference
- Deploy Claims Across Forests
- Claims Transformation Rules Language
- Dynamic Access Control: Scenario Overview

Claims Transformation Rules Language

The across-forest claims transformation feature enables you to bridge claims for Dynamic Access Control across forest boundaries by setting claims transformation policies on across-forest trusts. The primary component of all policies is rules that are written in claims transformation rules language. This topic provides details about this language and provides guidance about authoring claims transformation rules.

The Windows PowerShell cmdlets for transformation policies on across-forest trusts have options to set simple policies that are required in common scenarios. These cmdlets translate the user input into policies and rules in the claims transformation rules language, and then store them in Active Directory in the prescribed format. For more information about cmdlets for claims transformation, see the AD DS Cmdlets for Dynamic Access Control.

Depending on the claims configuration and the requirements placed on the across-forest trust in your Active Directory forests, your claims transformation policies may have to be more complex than the policies supported by the Windows PowerShell cmdlets for Active Directory. To effectively author such policies, it is essential to understand the claims transformation rules language syntax and semantics. This claims transformation rules language ("the language") in Active Directory is a subset of the language that is used by Active Directory Federation Services for similar purposes, and it has a very similar syntax and semantics. However, there are fewer operations allowed, and additional syntax restrictions are placed in the Active Directory version of the language.

This topic briefly explains the syntax and semantics of the claims transformation rules language in Active Directory and considerations to be made when authoring policies. It provides several sets of example rules to get you started, and examples of incorrect syntax and the messages they generate, to help you decipher error messages when you author the rules.

Tools for authoring claims transformation policies

Windows PowerShell cmdlets for Active Directory: This is the preferred and recommended way to author and set claims transformation policies. These cmdlets provide switches for simple policies and verify rules that are set for more complex policies.

LDAP: Claims transformation policies can be edited in Active Directory through Lightweight Directory Access Protocol (LDAP). However, this is not recommended because the policies have several complex components, and the tools you use may not validate the policy before writing it to Active Directory. This may subsequently require a considerable amount of time to diagnose problems.

Active Directory claims transformation rules language

Syntax overview

Here is a brief overview of the syntax and semantics of the language:

- The claims transformation rule set consists of zero or more rules. Each rule has two active parts: **Select Condition List** and **Rule Action**. If the **Select Condition List** evaluates to TRUE, the corresponding rule action is executed.
- **Select Condition List** has zero or more **Select Conditions**. All of the **Select Conditions** must evaluate to TRUE for the **Select Condition List** to evaluate to TRUE.
- Each **Select Condition** has a set of zero or more **Matching Conditions**. All the **Matching Conditions** must evaluate to TRUE for the Select Condition to evaluate to TRUE. All of these conditions are evaluated against a single claim. A claim that matches a **Select Condition** can be tagged by an **Identifier** and referred to in the **Rule Action**.
- Each **Matching Condition** specifies the condition to match the **Type** or **Value** or **ValueType** of a claim by using different **Condition Operators** and **String Literals**.
    - When you specify a **Matching Condition** for a **Value**, you must also specify a **Matching Condition** for a specific **ValueType** and vice versa. These conditions must be next to each other in the syntax.
    - **ValueType** matching conditions must use specific **ValueType** literals only.
- A **Rule Action** can copy one claim that is tagged with an **Identifier** or issue one claim based on a claim that is tagged with an Identifier and/or given String Literals.

Example rule

This example shows a rule that can be used to translate the claims Type between two forests, provided that they use the same claims ValueTypes and have the same interpretations for claims Values for this type. The rule has one matching condition and an Issue statement that uses String Literals and a matching claims reference.

```
C1: [TYPE=="EmployeeType"]
                  => ISSUE (TYPE= “EmpType”, VALUE = C1.VALUE, VALUETYPE = C1.VALUETYPE);
```

`[TYPE=="EmployeeType"] == Select Condition List with one Matching Condition for claims Type.`

`ISSUE (TYPE= “EmpType”, VALUE = C1.VALUE, VALUETYPE = C1.VALUETYPE) == Rule Action that issues a claims using string literal and matching claim referred with the Identifier.`

Runtime operation

It is important to understand the runtime operation of claims transformations to author the rules effectively. The runtime operation uses three sets of claims:

1. **Input claims set**: The input set of claims that are given to the claims transformation operation.
2. **Working claims set**: Intermediate claims that are read from and written to during the claims transformation.
3. **Output claims set**: Output of the claims transformation operation.

Here is a brief overview of the runtime claims transformation operation:

1. Input claims for claims transformation are used to initialize the working claims set.
   a. When processing each rule, the working claims set is used for the input claims.
   b. The Selection Condition List in a rule is matched against all possible sets of claims from the working claims set.
   c. Each set of matching claims is used to run the action in that rule.
   d. Running a rule action results in one claim, which is appended to the output claims set and the working claims set. Thus, the output from a rule is used as input for subsequent rules in the rule set.
2. The rules in the rule set are processed in sequential order starting with the first rule.
3. When the entire rule set is processed, the output claims set is processed to remove duplicate claims and for other security issues. The resulting claims are the output of the claims transformation process.

It is possible to write complex claims transformations based on the previous runtime behavior.

Example: Runtime operation

This example shows the runtime operation of a claims transformation that uses two rules.

```
     C1:[Type==”EmpType”, Value==”FullTime”,ValueType==”string”] =>
                Issue(Type==”EmployeeType”, Value==”FullTime”,ValueType==”string”);
     [Type==”EmployeeType”] =>
               Issue(Type==”AccessType”, Value==”Privileged”, ValueType==”string”);
Input claims and Initial Evaluation Context:
  {(Type= “EmpType”),(Value=”FullTime”),(ValueType=”String”)}
{(Type= “Organization”),(Value=”Marketing”),(ValueType=”String”)}
After Processing Rule 1:
 Evaluation Context:
  {(Type= “EmpType”),(Value=”FullTime”),(ValueType=”String”)}
{(Type= “Organization”), (Value=”Marketing”),(ValueType=”String”)}
  {(Type= “EmployeeType”),(Value=”FullTime”),(ValueType=”String”)}
Output Context:
  {(Type= “EmployeeType”),(Value=”FullTime”),(ValueType=”String”)}
```

```
After Processing Rule 2:

Evaluation Context:

  {(Type= “EmpType”),(Value=”FullTime”),(ValueType=”String”)}

{(Type= “Organization”),(Value=”Marketing”),(ValueType=”String”)}

  {(Type= “EmployeeType”),(Value=”FullTime”),(ValueType=”String”)}

  {(Type= “AccessType”),(Value=”Privileged”),(ValueType=”String”)}

Output Context:

  {(Type= “EmployeeType”),(Value=”FullTime”),(ValueType=”String”)}

  {(Type= “AccessType”),(Value=”Privileged”),(ValueType=”String”)}


Final Output:

  {(Type= “EmployeeType”),(Value=”FullTime”),(ValueType=”String”)}

  {(Type= “AccessType”),(Value=”Privileged”),(ValueType=”String”)}
```

Special rules semantics

The following are special syntax for rules:

1. Empty Rule Set == No Output Claims
2. Empty Select Condition List == Every Claim matches the Select Condition List

   Example: Empty Select Condition List

   The following rule matches every claim in the working set.

   ```
   => Issue (Type = “UserType”, Value = “External”, ValueType =
   “string”)
   ```
3. Empty Select Matching List == Every claim matches the Select Condition List

   Example: Empty Matching Conditions

   The following rule matches every claim in the working set. This is the basic “Allow-all” rule if it is used alone.

   ```
   C1:[] => Issule (claim = C1);
   ```

Security considerations

Claims that enter a forest

The claims presented by principals that are incoming to a forest need to be inspected thoroughly to ensure that we allow or issue only the correct claims. Improper claims can compromise the forest security, and this should be a top consideration when authoring transformation policies for claims that enter a forest.

Active Directory has the following features to prevent misconfiguration of claims that enter a forest:

- If a forest trust has no claims transformation policy set for the claims that enter a forest, for security purposes, Active Directory drops all the principal claims that enter the forest.
- If running the rule set on claims that enters a forest results in claims that are not defined in the forest, the undefined claims are dropped from the output claims.

Claims that leave a forest

Claims that leave a forest present a lesser security concern for the forest than the claims that enter the forest. Claims are allowed to leave the forest as-is even when there is no corresponding claims transformation policy in place. It is also possible to issue claims that are not defined in the forest as part of transforming claims that leave the forest. This is to easily set up across-forest trusts with claims. An administrator can determine if claims that enter the forest need to be transformed, and set up the appropriate policy. For example, an administrator could set a policy if there is a need to hide a claim to prevent information disclosure.

Syntax errors in claims transformation rules

If a given claims transformation policy has a rules set that is syntactically incorrect or if there are other syntax or storage issues, the policy is considered invalid. This is treated differently than the default conditions mentioned earlier.

Active Directory is unable to determine the intent in this case and goes into a fail-safe mode, where no output claims are generated on that trust+direction of traversal. Administrator intervention is required to correct the issue. This could happen if LDAP is used to edit the claims transformation policy. Windows PowerShell cmdlets for Active Directory have validation in place to prevent writing a policy with syntax issues.

Other language considerations

1. There are several key words or characters that are special in this language (referred to as terminals). These are presented in the Language terminals table later in this topic. The error messages use the tags for these terminals for disambiguation.
2. Terminals can sometimes be used as string literals. However, such usage may conflict with the language definition or have unintended consequences. This kind of usage is not recommended.
3. The rule action cannot perform any type conversions on claim Values, and a rule set that contains such a rule action is considered invalid. This would cause a runtime error, and no output claims are produced.
4. If a rule action refers to an Identifier that was not used in the Select Condition List portion of the rule, it is an invalid usage. This would cause a syntax error.

Example: Incorrect Identifier reference
The following rule illustrates an incorrect Identifier used in rule action.

```
C1:[] => Issue (claim = C2);
```

Sample transformation rules

- **Allow all claims of a certain type**

  Exact type

  ```
  C1:[type==”XYZ”] => Issue (claim = C1);
  ```

  Using Regex

  ```
  C1: [type =~ “XYZ*”] => Issue (claim = C1);
  ```

- **Disallow a certain claim type**

  Exact type

  ```
  C1:[type != “XYZ”] => Issue (claim=C1);
  ```

  Using Regex

  ```
  C1:[Type !~ “XYZ?”] => Issue (claim=C1);
  ```

Examples of rules parser errors

Claims transformation rules are parsed by a custom parser to check for syntax errors. This parser is run by related Windows PowerShell cmdlets before storing rules in Active Directory. Any errors in parsing the rules, including syntax errors, are printed on the console. Domain controllers also run the parser before using the rules for transforming claims, and they log errors in the event log (add event log numbers).

This section illustrates some examples of rules that are written with incorrect syntax and the corresponding syntax errors that are generated by the parser.

1. Example:

   ```
   c1;[]=>Issue(claim=c1);
   ```

   This example has an incorrectly used semicolon in place of a colon.
   Error message:
   POLICY0002: Could not parse policy data.
   Line number: 1, Column number: 2, Error token: ;. Line: 'c1;[]=>Issue(claim=c1);'.
   Parser error: 'POLICY0030: Syntax error, unexpected ';', expecting one of the following: ':' .'

2. Example:

   ```
   c1:[]=>Issue(claim=c2);
   ```

In this example, the Identifier tag in the copy issuance statement is undefined.
Error message:
POLICY0011: No conditions in the claim rule match the condition tag specified in the CopyIssuanceStatement: 'c2'.

3. Example:

```
c1:[type=="x1", value=="1", valuetype=="bool"]=>Issue(claim=c1)
```

"bool" is not a Terminal in the language, and it is not a valid ValueType. Valid terminals are listed in the following error message.
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 39, Error token: "bool". Line: 'c1:[type=="x1", value=="1",valuetype=="bool"]=>Issue(claim=c1);'.
Parser error: 'POLICY0030: Syntax error, unexpected 'STRING', expecting one of the following: 'INT64_TYPE' 'UINT64_TYPE' 'STRING_TYPE' 'BOOLEAN_TYPE' 'IDENTIFIER'

4. Example:

```
c1:[type=="x1", value==1,
valuetype=="boolean"]=>Issue(claim=c1);
```

The numeral **1** in this example is not a valid token in the language, and such usage is not allowed in a matching condition. It has to be enclosed in double quotes to make it a string.
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 23, Error token: 1. Line: 'c1:[type=="x1", value==1, valuetype=="bool"]=>Issue(claim=c1);'.Parser error: 'POLICY0029: Unexpected input.

5. Example:

```
c1:[type == "x1", value == "1", valuetype == "boolean"] =>

      Issue(type = c1.type, value="0", valuetype == "boolean");
```

This example used a double equal sign (==) instead of a single equal sign (=).
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 91, Error token: ==. Line: 'c1:[type=="x1", value=="1", valuetype=="boolean"]=>Issue(type=c1.type, value="0", valuetype=="boolean");'.
Parser error: 'POLICY0030: Syntax error, unexpected '==', expecting one of the following: '='

6. Example:

```
c1:[type=="x1", value=="boolean", valuetype=="string"] =>
```

```
Issue(type=c1.type, value=c1.value, valuetype = "string");
```

This example is syntactically and semantically correct. However, using “boolean” as a string value is bound to cause confusion, and it should be avoided. As previously mentioned, using language terminals as claims values should be avoided where possible.

Language terminals

The following table lists the complete set of terminal strings and the associated language terminals that are used in the claims transformation rules language. These definitions use case-insensitive UTF-16 strings.

String	Terminal
"=>"	IMPLY
";"	SEMICOLON
":"	COLON
","	COMMA
"."	DOT
"["	O_SQ_BRACKET
"]"	C_SQ_BRACKET
"("	O_BRACKET
")"	C_BRACKET
"=="	EQ
"!="	NEQ
"=~"	REGEXP_MATCH
"!~"	REGEXP_NOT_MATCH
"="	ASSIGN
"&&"	AND
"issue"	ISSUE
"type"	TYPE
"value"	VALUE
"valuetype"	VALUE_TYPE
"claim"	CLAIM

String	Terminal
"[_A-Za-z][_A-Za-z0-9]*"	IDENTIFIER
"\"[^\"\n]*\""	STRING
"uint64"	UINT64_TYPE
"int64"	INT64_TYPE
"string"	STRING_TYPE
"boolean"	BOOLEAN_TYPE

Language syntax

The following claims transformation rules language is specified in ABNF form. This definition uses the terminals that are specified in the previous table in addition to the ABNF productions defined here. The rules must be encoded in UTF-16, and the string comparisons must be treated as case insensitive.

```
Rule_set        = ;/*Empty*/
             / Rules
Rules         = Rule
             / Rule Rules
Rule          = Rule_body
Rule_body       = (Conditions IMPLY Rule_action SEMICOLON)
Conditions       = ;/*Empty*/
             / Sel_condition_list
Sel_condition_list   = Sel_condition
             / (Sel_condition_list AND Sel_condition)
Sel_condition     = Sel_condition_body
             / (IDENTIFIER COLON Sel_condition_body)
Sel_condition_body   = O_SQ_BRACKET Opt_cond_list C_SQ_BRACKET
Opt_cond_list     = /*Empty*/
             / Cond_list
Cond_list       = Cond
             / (Cond_list COMMA Cond)
Cond          = Value_cond
             / Type_cond
```

```
Type_cond       = TYPE Cond_oper Literal_expr
Value_cond       = (Val_cond COMMA Val_type_cond)
             /(Val_type_cond COMMA Val_cond)
Val_cond        = VALUE Cond_oper Literal_expr
Val_type_cond     = VALUE_TYPE Cond_oper Value_type_literal
claim_prop       = TYPE
             / VALUE
Cond_oper       = EQ
             / NEQ
             / REGEXP_MATCH
             / REGEXP_NOT_MATCH
Literal_expr      = Literal
             / Value_type_literal

Expr          = Literal
             / Value_type_expr
             / (IDENTIFIER DOT claim_prop)
Value_type_expr    = Value_type_literal
             /(IDENTIFIER DOT VALUE_TYPE)
Value_type_literal   = INT64_TYPE
             / UINT64_TYPE
             / STRING_TYPE
             / BOOLEAN_TYPE
Literal        = STRING
Rule_action      = ISSUE O_BRACKET Issue_params C_BRACKET
Issue_params      = claim_copy
             / claim_new
claim_copy       = CLAIM ASSIGN IDENTIFIER
claim_new       = claim_prop_assign_list
claim_prop_assign_list = (claim_value_assign COMMA claim_type_assign)
             /(claim_type_assign COMMA claim_value_assign)
claim_value_assign   = (claim_val_assign COMMA claim_val_type_assign)
             /(claim_val_type_assign COMMA claim_val_assign)
claim_val_assign    = VALUE ASSIGN Expr
```

```
claim_val_type_assign = VALUE_TYPE ASSIGN Value_type_expr

Claim_type_assign   = TYPE ASSIGN Expr
```

Scenario: File Access Auditing

Security Auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is regulatory compliance. Industry standards such as Sarbanes Oxley, Health Insurance Portability and Accountability Act (HIPAA), and Payment Card Industry (PCI) require enterprises to follow a strict set of rules related to data security and privacy. Security audits help establish the presence of such policies and prove compliance with these standards. Additionally, security audits help detect anomalous behavior, identify and mitigate gaps in security policies, and deter irresponsible behavior by creating a trail of user activity that can be used for forensic analysis.

Audit policy requirements are typically driven at the following levels:

- **Information security.** File access audit trails are often used for forensic analysis and intrusion detection. Being able to get targeted events about access to high-value information lets organizations considerably improve their response time and investigation accuracy.
- **Organizational policy.** For example, organizations regulated by PCI standards could have a central policy to monitor access to all files that are marked as containing credit card information and personally identifiable information (PII).
- **Departmental policy.** For example, the finance department may require that the ability to modify certain finance documents (such as a quarterly earnings report) be restricted to the finance department, and thus the department would want to monitor all other attempts to change these documents.
- **Business policy.** For example, business owners may want to monitor all unauthorized attempts to view data that belongs to their projects.

Additionally, the compliance department may want to monitor all changes to central authorization policies and policy constructs such as user, computer, and resource attributes.

One of the biggest considerations of security audits is the cost of collecting, storing, and analyzing audit events. If the audit policies are too broad, the volume of audit events collected rises, and this increases costs. If the audit policies are too narrow, you risk missing important events.

With Windows Server 2012, you can author audit policies by using claims and resource properties. This leads to richer, more targeted, and easier-to-manage audit policies. It enables scenarios that, until now, were impossible or too difficult to perform. The following are examples of audit policies that administrators can author:

- Audit everyone who does not have a high-security clearance and tries to access an HBI document. For example, Audit | Everyone | All-Access | Resource.BusinessImpact=HBI AND User.SecurityClearance!=High.

- Audit all vendors when they try to access documents that are related to projects that they are not working on. For example, Audit | Everyone | All-Access | User.EmploymentStatus=Vendor AND User.Project Not_AnyOf Resource.Project.

These policies help regulate the volume of audit events and limit them to only the most relevant data or users.

After administrators have created and applied the audit policies, the next consideration for them is gleaning meaningful information from the audit events that they collected. Expression-based audit events help reduce the volume of audits. However, users need a way to query these events for meaningful information and ask questions such as, “Who is accessing my HBI data?” or “Was there an unauthorized attempt to access sensitive data?”

Windows Server 2012 enhances existing data access events with user, computer, and resource claims. These events are generated on a per-server basis. To provide a full view of events across the organization, Microsoft is working with partners to provide event collection and analysis tools, such as the Audit Collection Services in System Center Operation Manager .

Figure 4 shows an overview of a central audit policy.

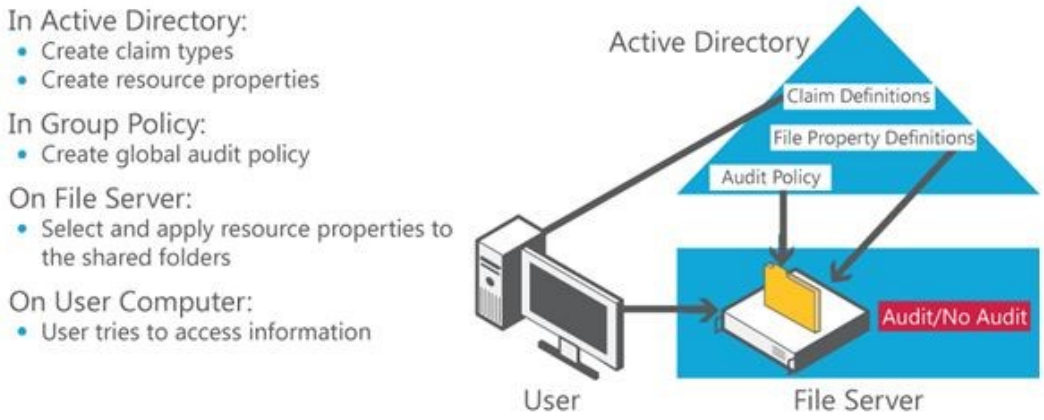

Figure 4 Central auditing experiences

Setting up and consuming security audits typically involves the following general steps:

1. Identify the correct set of data and users to monitor
2. Create and apply appropriate audit policies
3. Collect and analyze audit events
4. Manage and monitor the policies that were created

In this scenario

The following topics provide additional guidance for this scenario:

- Plan for File Access Auditing
- Deploy Security Auditing with Central Audit Policies (Demonstration Steps)

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Doman Services role	AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables creating user claims and device claims, compound identity, (user plus device claims), new central access policies (CAP) model, and the use of file classification information in authorization decisions.
File and Storage Services role	File servers in Windows Server 2012 provide a user interface where administrators can view the effective permissions for users for a file or folder and troubleshoot access issues and grant access as required.

Plan for File Access Auditing

The information in this topic explains the security auditing enhancements that are introduced in Windows Server 2012 and new audit settings that you should consider as you deploy Dynamic Access Control in your enterprise. The actual audit policy settings that you deploy will depend on your goals, which can include regulatory compliance, monitoring, forensic analysis, and troubleshooting.

Note
Detailed information about how to plan and deploy an overall security auditing strategy for your enterprise is explained in Planning and Deploying Advanced Security Audit Policies. For more information about configuring and deploying a security audit policy, see the Advanced Security Audit Policy Step-by-Step Guide.

The following security auditing capabilities in Windows Server 2012 can be used with Dynamic Access Control to extend your overall security auditing strategy.

- **Expression-based audit policies**. Dynamic Access Control enables you to create targeted audit policies by using expressions based on user, computer, and resource claims. For example, you could create an audit policy to track all Read and Write operations on files classified as high-business impact by employees who do not have a high-security clearance. Expression-based audit policies can be authored directly for a file or folder or centrally through Group Policy. For more information, see Group Policy using Global Object Access Auditing.

- **Additional information from object access auditing**. File access auditing is not new to Windows Server 2012. With the right audit policy in place, the Windows and Windows Server operating systems generate an audit event each time a user accesses a file. Existing File Access events (4656, 4663) contain information about the attributes of the file that was accessed. This information can be used by event log filtering tools to help you identify the most relevant audit events. For more information, see Audit Handle Manipulation and Audit Security Accounts Manager.
- **More information from user logon events**. With the right audit policy in place, Windows operating systems generate an audit event every time a user signs in to a computer locally or remotely. In Windows Server 2012 or Windows 8, you can also monitor user and device claims associated with a user's security token. Examples can include Department, Company, Project, and Security clearances.Event 4626 contains information about these user claims and device claims, which can be leveraged by audit log management tools to correlate user logon events with object access events to enable event filtering based on file attributes and user attributes. For more information about user logon auditing, see Audit Logon.
- **Change tracking for new types of securable objects**. Tracking changes to securable objects can be important in the following scenarios:
    - **Change tracking for central access policies and central access rules**. Central access policies and central access rules define the central policy that can be used to control access to critical resources. Any change to these can directly impact the file access permissions that are granted to users on multiple computers. Therefore, tracking changes to central access policies and central access rules can be important for your organization. Because central access policies and central access rules are stored in Active Directory Domain Services (AD DS), you can audit attempts to modify them, like auditing changes to any other securable object in AD DS. For more information, see Audit Directory Service Access.
    - **Change tracking for definitions in the claim dictionary**. Claim definitions include the claim name, description, and possible values. Any change to the claim definition can impact the access permissions on critical resources. Therefore, tracking changes to claim definitions can be important to your organization. Like central access policies and central access rules, claim definitions are stored in AD DS; therefore, they can be audited like any another securable object in AD DS. For more information, see Audit Directory Service Access.
    - **Change tracking for file attributes**. File attributes determine which central access rule applies to the file. A change to the file attributes can potentially impact the access restrictions on the file. Therefore, it can be important to track changes to file attributes. You can track changes to file attributes on any computer by configuring the authorization policy change auditing policy. For more information, see Authorization Policy Change auditing and Object Access auditing for File Systems. In Windows Server 2012, Event 4911 differentiates file attribute policy changes from other authorization policy change events.
    - **Chang tracking for the central access policy associated with a file.** Event 4913 displays the security identifiers (SIDs) of the old and new central access policies. Each central access policy also has a user friendly name that can be looked up using this security identifier. For more information, see Authorization Policy Change auditing.

- **Change tracking for user and computer attributes**. Like files, user and computer objects can have attributes, and changes to these attributes can impact the user's ability to access files. Therefore, it can be valuable to track changes to user or computer attributes. User and computer objects are stored in AD DS; therefore, changes to their attributes can be audited. For more information, see DS Access .

- **Policy change staging**. Changes to central access policies can impact the access control decisions on all computers where the policies are enforced. A loose policy could grant more access than desired, and an overly restrictive policy could generate an excessive number of Help Desk calls. As a result, it can be extremely valuable to verify changes to a central access policy before enforcing the change. For that purpose, Windows Server 2012 introduces the concept of "staging." Staging enables users to verify their proposed policy changes before enforcing them. To use policy staging, proposed policies are deployed with the enforced policies, but staged policies do not actually grant or deny permissions. Instead, Windows Server 2012 logs an audit event (4818) any time the result of the access check that uses the staged policy is different from the result of an access check that uses the enforced policy.

Deploy Security Auditing with Central Audit Policies (Demonstration Steps)

In this scenario, you will audit access to files in the Finance Documents folder by using the Finance Policy that you created in Deploy a Central Access Policy (Demonstration Steps). If a user who is not authorized to access the folder attempts to access it, the activity is captured in the event viewer.

The following steps are required to test this scenario.

Task	Description
Configure global object access policy	In this step, you configure the global object access policy on the domain controller.
Update Group Policy settings	Sign in to the file server and apply the Group Policy update.
Verify that the global object access policy has been applied	View the relevant events in the event viewer. The events should include metadata for the country and document type.

Configure global object access policy

In this step, you configure the global object access policy in the domain controller.

To configure a global object access policy

1. Sign in to the domain controller DC1 as contoso\administrator with the password **pass@word1**.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Domains**, double-click **contoso.com**, click **Contoso**, and then double-click **File Servers**.
4. Right-click **FlexibleAccessGPO**, and click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **Audit Policies**.
7. Double-click **Object Access**, and then double-click **Audit File System**.
8. Select the **Configure the following events** check box, select the **Success** and **Failure** check boxes, and then click **OK**.
9. In the navigation pane, double-click **Global Object Access Auditing**, and then double-click **File system**.
10. Select the **Define this policy setting** check box, and click **Configure**.
11. In the **Advanced Security Settings for Global File SACL** box, click **Add**, then click **Select a principal**, type **Everyone**, and then click **OK**.
12. In the **Auditing Entry for Global File SACL** box, select **Full control** in the **Permissions** box.
13. In the **Add a condition:** section, click **Add a condition** and in the drop-down lists select [**Resource**] [**Department**] [**Any of**] [**Value**] [**Finance**].
14. Click **OK** three times to complete the configuration of the global object access audit policy setting.
15. In the navigation pane, click **Object Access**, and in the results pane, double-click **Audit Handle Manipulation**. Click **Configure the following audit events**, **Success**, and **Failure**, click **OK**, and then close the flexible access GPO.

Update Group Policy settings

In this step, you update the Group Policy settings after you have created the audit policy.

To update Group Policy settings

1. Sign in to the file server, FILE1 as contoso\Administrator, with the password **pass@word1**.
2. Press the Windows key+R, then type **cmd** to open a Command Prompt window.

   Note
   If the **User Account Control** dialog box appears, confirm that the action it

displays is what you want, and then click **Yes**.

3. Type **gpupdate /force** and then press ENTER.

Verify that the global object access policy has been applied

After the Group Policy settings have been applied, you can verify that the audit policy settings were applied correctly.

To verify that the global object access policy has been applied

1. Sign in to client computer, CLIENT1 as Contoso\MReid. Browse to the folder \\FILE1\Finance Documents, and modify Word Document 2.
2. Sign in to the file server, FILE1 as contoso\administrator. Open Event Viewer, browse to **Windows Logs**, select **Security**, and confirm that your activities resulted in audit events **4656** and **4663** (even though you did not set explicit auditing SACLs on the files or folders that you created, modified, and deleted).

Important

A new logon event is generated on the computer where the resource is located, on behalf of the user for whom effective access is being checked. When analyzing security audit logs for user sign-in activity, to differentiate between logon events that are generated because of effective access and those generated because of an interactive network user sign in, the Impersonation Level information is included. When the logon event is generated because of effective access, the Impersonation Level will be Identity. A network interactive user sign in typically generates a logon event with the Impersonation Level = Impersonation or Delegation.

See also

- Scenario: File Access Auditing
- Plan for File Access Auditing
- Dynamic Access Control: Scenario Overview

Scenario: Access-Denied Assistance

Users will get an access-denied message when they try to access shared files and folders on a file server for which they do not have permissions. Administrators often do not have the appropriate context to troubleshoot the access issue, which makes it hard to resolve the issue.

Scenario description

Access-denied assistance is a new feature in Windows Server 2012, which provides the following ways to troubleshoot issues that are related to access to files and folders:

- **Self-assistance.** If a user can determine the issue and remediate the problem so that they can get the requested access, the impact to the business is low, and no special exceptions are needed in the central access policy. Access-denied assistance provides an access-denied message that file server administrators can customize with information specific to their organizations. For example, an administrator could set the message so that users can request access from a data owner without involving the file server administrator.
- **Assistance by the data owner.** You can define a distribution list for shared folders, and configure it so that the folder owner receives an email notification when a user needs access. If the data owner does not know how to help the user get access, the owner can forward this information to the file server administrator.
- **Assistance by the file server administrator.** This type of assistance is available when the user cannot fix an issue and the data owner cannot help. Windows Server 2012 provides a user interface where file server administrators can view the effective permissions for a user on a file or folder so that it is easier to troubleshoot access issues.

Access-denied assistance in Windows Server 2012 provides file server administrators the relevant access details so that they can determine the issue and appropriate tools so that they can make configuration changes to satisfy the access request. For example, a user might follow this process to access a file that they currently do not have access to:

- The user attempts to read a file in the \\financeshares shared folder, but the server displays an access-denied message.
- Windows Server 2012 displays the access-denied assistance information to the user with an option to request assistance.
- If the user requests access to the resource, the server sends an email with the access request information to the folder owner.

You can find planning information for configuring access-denied assistance in Plan for Access-Denied Assistance.

You can find steps about configuring access-denied assistance in Deploy Access-Denied Assistance (Demonstration Steps).

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Practical applications

Access-denied assistance in Windows Server 2012 contributes to Dynamic Access Control by giving users the ability to request access to shared files and folders directly from an access-denied message.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	Access-denied assistance can be configured by using the File Server Resource Manager console on the file server.
File and Storage Services Overview	File Server Resource Manager is a File and Storage Services role service, and it is comprised of a set of features that can be used to administer the file servers on your network.

Plan for Access-Denied Assistance

There are a few considerations and decisions that should be made before you deploy access-denied assistance.

Use the following table to plan your access-denied assistance deployment in your organization.

Task	Description
1.1 Determine the access-denied assistance model	Determine whether your organization should use an email model or a Web services model for access-denied assistance.
1.2. Determine who should handle access requests	You can assign each file share an owner distribution list that will receive access requests.
1.3. Customize the access-denied assistance message	The access-denied assistance message should be customized for your organization. The included message is only a sample.
1.4. Plan for exceptions	Exceptions happen when a user account needs access to a specific file share but they do not

Task	Description
	need access to everything that the security group has.
1.5. Determine how access-denied assistance is deployed	Access-denied assistance can be configured at the file server level or at the file share level.

1.1 Determine the access-denied assistance model

There are two ways that you can configure access-denied assistance in your organization:

- **Email model** In an email model, if a user is denied access, a customized error message is shown with a button to request assistance. When the user clicks the **Request Assistance** button, an email is sent to the folder owner with the specified information.
- **Web services model** The Web services model is similar to the email model. Instead of the request assistance button being shown, a link is included in the access-denied assistance message that directs the user to request access through a self-service portal, such as Forefront Identity Manager.

The model that you choose is dependent on your organization.

1.2. Determine who should handle access requests

When using the email model each of the file shares, you can determine whether access requests to each file share will be received by the administrator, a distribution list that represents the file share owners, or both.

The owner distribution list is configured by using the **SMB Share – Advanced** file share profile in the New Share Wizard in Server Manager.

You can also use the File Server Resource Manager console to configure the owner distribution list by editing the management properties of the classification properties.

1.3. Customize the access-denied assistance message

Windows Server 2012 contains a sample message when you enable access-denied assistance. You should customize this message to meet the needs of your organization. One thing to consider including in the message is a link to your Intranet or help desk location.

You can also provide a specific access-denied assistance message per file share. This message will replace the global message when a user tries to access files within file share. For more

information on how to configure a separate access-denied assistance message for a file share, see Deploy Access-Denied Assistance (Demonstration Steps).

1.4. Plan for exceptions

How to deal with exceptions is an important consideration to plan for before you deploy access-denied assistance. Exceptions can happen if a user account needs access to a file share.

We recommend that the user account should not be added to any of the security groups that have access as part of the user's role. Instead, you should create another security group for exceptions that contains those user accounts and grant access to the appropriate file share so that you can monitor this security group separately and enforce rules, such as expiring membership.

1.5. Determine how access-denied assistance is deployed

Access-denied assistance can be configured on a per-file server or per-share basis. By configuring access-denied assistance at the file share level, you can customize the message to include specific information about the file share itself. For example, you could specify the exception security group for that share in the message so the user would know which group in which they should request access. Another example is to specify a specific owner distribution list for the folder that represents the file share.

Note

Some settings are global for the server. Enabling or disabling access-denied assistance is done on a server basis and the **Request Assistance** button is configured the file server level, so if you configure the file server to show the **Request Assistance** button, you cannot disable it on specific file shares.

For more information on configuring access-denied assistance at both a file share level and at a server level, see Deploy Access-Denied Assistance (Demonstration Steps).

See also

- Scenario: Access-Denied Assistance
- Deploy Access-Denied Assistance (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Deploy Access-Denied Assistance (Demonstration Steps)

This topic explains how to configure access-denied assistance, and verify that it is working properly.

In this document

- Step 1: Configure access-denied assistance
- Step 2: Configure the email notification settings
- Step 3: Verify that access-denied assistance is configured correctly

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Configure access-denied assistance

You can configure access-denied assistance within a domain by using Group Policy, or you can configure the assistance individually on each file server by using the File Server Resource Manager console. You can also change the access-denied message for a specific shared folder on a file server.

You can configure access-denied assistance for the domain by using Group Policy as follows:

Do this step using Windows PowerShell

To configure access-denied assistance by using Group Policy

1. Open Group Policy Management. In Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Right-click the appropriate Group Policy, and then click **Edit**.
3. Click **Computer Configuration**, click **Policies**, click **Administrative Templates**, click **System**, and then click **Access-Denied Assistance**.
4. Right-click **Customize message for Access Denied errors**, and then click **Edit**.
5. Select the **Enabled** option.
6. Configure the following options:
   a. In the **Display the following message to users who are denied access** box, type a message that users will see when they are denied access to a file or folder.

      You can add macros to the message that will insert customized text. The macros include:

      - **[Original File Path]** The original file path that was accessed by the user.
      - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.

- **[Admin Email]** The administrator email recipient list.
- **[Data Owner Email]** The data owner email recipient list.

b. Select the **Enable users to request assistance** check box.

c. Leave the remaining default settings.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName AllowEmailRequests
–Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName GenerateLog –Type
DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName
IncludeDeviceClaims –Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName IncludeUserClaims
–Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName PutAdminOnTo –Type
DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName PutDataOwnerOnTo –
Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName ErrorMessage –Type
MultiString –value “Type the text that the user will see in the error message dialog
box.”

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName Enabled –Type
DWORD –value 1
```

Alternatively, you can configure access-denied assistance individually on each file server by using the File Server Resource Manager console.

Do this step using Windows PowerShell

To configure access-denied assistance by using File Server Resource Manager

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
3. Click the **Access-Denied Assistance** tab.
4. Select the **Enable access-denied assistance** check box.
5. In the **Display the following message to users who are denied access to a folder or file** box, type a message that users will see when they are denied access to a file or folder.

   You can add macros to the message that will insert customized text. The macros include:

   - **[Original File Path]** The original file path that was accessed by the user.
   - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
   - **[Admin Email]** The administrator email recipient list.
   - **[Data Owner Email]** The data owner email recipient list.
6. Click **Configure email requests**, select the **Enable users to request assistance** check box, and then click **OK**.
7. Click **Preview** if you want to see how the error message will look to the user.
8. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FSRMAdrSetting -Event "AccessDenied" –DisplayMessage “Type the text that the user will see in the error message dialog box.” -Enabled:$true -AllowRequests:$true
```

After you configure the access-denied assistance, you must enable it for all file types by using Group Policy.

Do this step using Windows PowerShell

To configure access-denied assistance for all file types by using Group Policy

1. Open Group Policy Management. In Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Right-click the appropriate Group Policy, and then click **Edit**.
3. Click **Computer Configuration**, click **Policies**, click **Administrative Templates**, click **System**, and then click **Access-Denied Assistance**.
4. Right-click **Enable access-denied assistance on client for all file types**, and then click **Edit**.
5. Click **Enabled**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\SOFTWARE\Policies\Microsoft\Windows\Explorer” –ValueName
EnableShellExecuteFileStreamCheck –Type DWORD –value 1
```

You can also specify a separate access-denied message for each shared folder on a file server by using the File Server Resource Manager console.

Do this step using Windows PowerShell

To specify a separate access-denied message for a shared folder by using File Server Resource Manager

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Expand **File Server Resource Manager (Local)**, and then click **Classification Management**.
3. Right-click **Classification Properties**, and then click **Set Folder Management Properties**.
4. In the **Property** box, click **Access-Denied Assistance Message**, and then click **Add**.
5. Click **Browse**, and then choose the folder that should have the custom access-denied message.
6. In the **Value** box, type the message that should be presented to the users when they cannot access a resource within that folder.

   You can add macros to the message that will insert customized text. The macros include:

   - **[Original File Path]** The original file path that was accessed by the user.
   - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
   - **[Admin Email]** The administrator email recipient list.
   - **[Data Owner Email]** The data owner email recipient list.
7. Click **OK**, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FSRMMgmtProperty -Namespace "folder path" -Name "AccessDeniedMessage_MS" -Value “Type
the text that the user will see in the error message dialog box.”
```

Step 2: Configure the email notification settings

You must configure the email notification settings on each file server that will send the access-denied assistance messages.

Do this step using Windows PowerShell

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
3. Click the **Email Notifications** tab.
4. Configure the following settings:
   - In the **SMTP server name or IP address** box, type the name of IP address of the SMTP server in your organization.
   - In the **Default administrator recipients** and **Default "From" e-mail address** boxes, type the email address of the file server administrator.
5. Click **Send Test E-mail** to ensure that the email notifications are configured correctly.
6. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
set-FSRMSetting -SMTPServer "server1" -AdminEmailAddress "fileadmin@contoso.com" -
FromEmailAddress "fileadmin@contoso.com"
```

Step 3: Verify that access-denied assistance is configured correctly

You can verify that the access-denied assistance is configured correctly by having a user who is running Windows 8 try to access a share or a file in that share that they do not have access to. When the access-denied message appears, the user should see a **Request Assistance** button. After clicking the Request Assistance button, the user can specify a reason for access and then send an email to the folder owner or file server administrator. The folder owner or file server administrator can verify for you that the email arrived and contains the appropriate details.

Important

If you want to verify access-denied assistance by having a user who is running Windows Server 2012, you must install the Desktop Experience before connecting to the file share.

See also

- Scenario: Access-Denied Assistance
- Plan for Access-Denied Assistance
- Dynamic Access Control: Scenario Overview

Scenario: Classification-Based Encryption for Office Documents

Protection of sensitive information is mainly about mitigating risk for the organization. Various compliance regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and Payment Card Industry Data Security Standard (PCI-DSS), dictate encryption of information, and there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it might impair business productivity. Thus, organizations tend to have different approaches and priorities for encrypting their information.

Scenario description

Windows Server 2012 provides the ability to automatically encrypt sensitive Microsoft Office files, based on their classification. This is done through file management tasks that invoke Active Directory Rights Management Services (AD RMS) protection for sensitive documents a few seconds after the file is identified as being a sensitive file on the file server. This is facilitated by continuous file management tasks on the file server.

AD RMS encryption provides another layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the file is protected by the AD RMS encryption. Users who want to access the file must first authenticate themselves to an AD RMS server to receive the decryption key. The following figure shows this process.

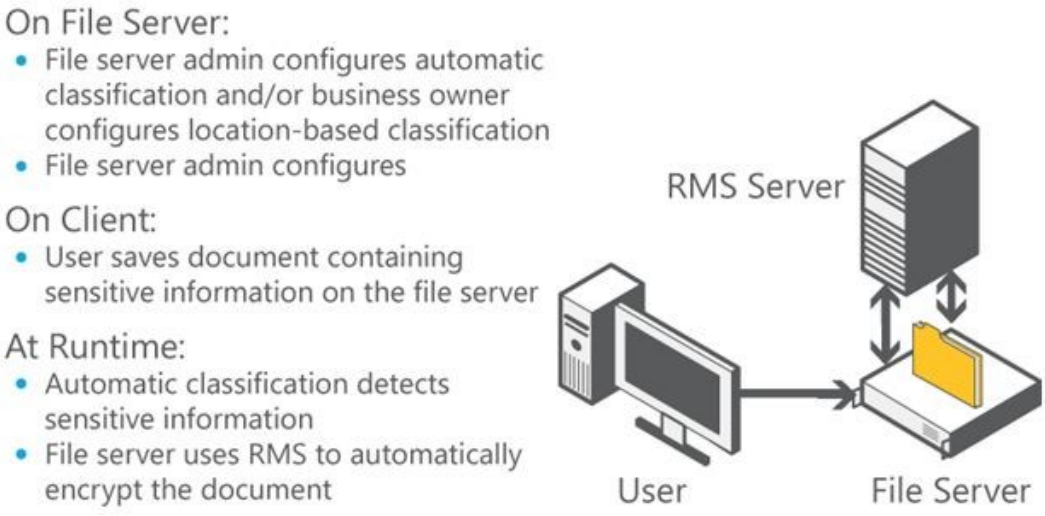

Figure 6 Classification-based RMS protection

Support for non-Microsoft file formats is available through non-Microsoft vendors. After a file has been protected by AD RMS encryption, data management features such as search- or content-based classification are no longer available for that file.

In this scenario

Following is the guidance that is available for this scenario:

- Planning Considerations for Encryption of Office Documents
- Deploy Encryption of Office Files (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services role (AD DS)	AD DS provides a distributed database that stores and manages information about network resources and application-specific data from directory-enabled applications. In this scenario, AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables the creation of user claims and device

Role/feature	How it supports this scenario
	claims, compound identity (user plus device claims), a new central access policies model, and the use of file-classification information in authorization decisions.
File and Storage Services role File Server Resource Manager	File and Storage Services provides technologies to help you set up and manage one or more file servers that provide central locations on your network where you can store files and share them with users. If your network users need access to the same files and applications, or if centralized backup and file management are important to your organization, you should set up one or more computers as a file server by adding the File and Storage Services role and the appropriate role services to the computers. In this scenario, file server administrators can configure file management tasks that invoke AD RMS protection for sensitive documents a few seconds after the file is identified as being a sensitive file on the file server (continuous file management tasks on the file server).
Active Directory Rights Management Services (AD RMS) role	AD RMS enables individuals and administrators (through Information Rights Management (IRM) policies) to specify access permissions to documents, workbooks, and presentations. This helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. After permission for a file has been restricted by using IRM, the access and usage restrictions are enforced no matter where the information is, because the permission to a file is stored in the document file itself. In this scenario, AD RMS encryption provides another layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the file is protected by the AD RMS encryption. Users who want to access the file must first authenticate themselves to an AD RMS server

Role/feature	How it supports this scenario
	to receive the decryption key.

Planning Considerations for Encryption of Office Documents

Active Directory Rights Management Services (AD RMS) is an information protection technology that can be used with dynamic access control to protect files as they are flexibly exchanged within an organization or across organizational boundaries. AD RMS works with Microsoft Office to help safeguard documents by enabling information rights management (IRM) features when you create documents using enabled applications such as Microsoft Word and Microsoft Outlook.

By using AD RMS to support the IRM feature capabilities within Microsoft Office, you can create documents or email messages and protect them from unauthorized use as they are flexibly shared in digital form. When assigning rights, content owners can define exactly how a recipient can use the information; including such rights as the ability to open, modify, print, or forward the information.

Note
AD RMS can be implemented to provide a comprehensive information-protection platform in complex business environments. Additionally, it can be effectively implemented to protect information shared among multiple organizations. For more on planning deployment scenarios around AD RMS, see AD RMS Architecture Design and Secure Collaboration Scenarios(http://go.microsoft.com/fwlink/?LinkId=256399).

This topic provides guidance around the following additional considerations that should be looked at when you are using AD RMS to protect and encrypt Office documents that you will provide limited access to by sharing them using dynamic access control. (For more about how dynamic access control works, see Dynamic Access Control: Scenario Overview.)

- Determining files to automatically encrypt
- Determining the rights policy template to use when encrypting files
- Multi-machine considerations

Determining files to automatically encrypt

In determining which files or folders to automatically encrypt using AD RMS technology, there are two aspects to consider. First, you should determine what types of files need to be encrypted. For example, do you have files containing confidential information or personal information?

Next, you should determine which resource properties are useful ways to describe those types of files. To get started, try to use the predefined resource properties whenever possible. Typically, this will mean selecting the most appropriate resource property from the list of resource

properties that are either built-in or user-defined using Active Directory Administrative Center. For example, the following table provides suggestions for some of the exact built-in resource properties you might choose to use when identifying files to be marked for encryption and rights protection.

Resource Property	Description
PII	The Personal Identifying Information (PII) resource property can be used to identify files that have PII within them to ensure those files are protected at the appropriate assigned level of sensitivity. For example,
Compliancy	The Compliancy property for confidentiality purposes according HIPAA guidelines in an organization that provides health care services.
Project	This resource property can be used to set a list of name values for which the name of a special project in your organization could be applied to all of its related files. This could identify project files for the purposes of them having encryption and protection applied to them, ensuring that only project team members have access to open and view these files.

In addition to the above resource properties, you might choose to have other custom resource properties defined in your deployment that apply the following criterion:

- **Based on file location**. You might create a new resource property definition called "Location" and then set list of supported values for it to the names of specific sites or locations within your enterprise or organization. You could then assign the values to this resource property for each of the files by their location. This might be a good strategy to implement if you have a secure location that requires all documents within it to be encrypted and protected.
- **Based on file contents**. Another way to create a resource property to help protect content based on identifying the type of content it is. For example, if you only need to protect files or documents that have sensitive financial or payroll information, you might choose to use "Financial" or "Payroll" as supported values for a "Content" resource property and use that property to identify and tag files that need to be encrypted and protected.
- **Based on manual classification**. Another possible way to identify files for classification is to manually classify them. Manual classification gives users and content owners the ability to classify their files and folders by using the properties sheet of that file or folder. For example, you can classify folders so that any file added to the folder will inherit the classifications of the parent folder. For more information, see Working with File Classification.

Determining the rights policy template to use when encrypting files

Rights policy templates are used in AD RMS to control the rights that a user or group has on a particular piece of rights-protected content. When you are deciding upon which rights policy template to use when encrypting files in your organization, it's helpful to consider the purpose and scope of how you are planning to protect content.

For example, you might consider using a rights policy template for the purpose of protecting PII content that would be generated by the finance or payroll department within your organization. In this case you could name the template "Finance" and select the following configuration user and rights settings when creating the template in the AD RMS console.

- Configure an Active Directory group to be assigned use of the template, one that contains only the regular staff employees who work within the finance or payroll department. Note that this group will also need to be configured for an email name for AD RMS to use in identifying it such as "finance@contoso.com" or "payroll@contoso.com".
- Select **View** and **View Rights** to limit the department employees in their ability to modify payroll or other financial details for documents.

Similar template configurations might also be done to protect content considered to be High Business Impact (HBI) information by having a "Company Confidential" template that enforces similar view-only rights and limited to all full-time employees only.

Another way of determining a rights policy template to use would be to apply scope as the determining factor. One way to do this would be to reserve the use of separate network mapped drives or volumes for each department so they can be used by workers within the department to store their documents. In this scenario, you might have templates that are specific to each department and ensuring that they are only available to be applied to files on that department's mapped drive or volume.

For example, the finance department might have their work files managed using a template for volume (drive) F and the engineering department might use volume G for storing their protected work files. File management has been configured so that workers in each department do not have access to the other department's volume. This will be useful if you if you need to create more than one file management task to manage protection for multiple departments and rights policy templates.

Note

There should be only a 1:1 relationship between templates and file management tasks. In general, you should try to use only a single file management task and rights policy template throughout your organization to schedule and manage template usage. If you do use more than one template and file management task as suggested here, be sure to take the suggested precautions and manage separate scope well so that tasks and volumes where files are managed are not able to be assigned protection using more than a single template to avoid rights conflicts or collisions.

For more information on working with AD RMS rights policy templates, see the following additional resources:

- Configuring Rights Policy Templates
- Working with File Management Tasks
- Creating a New Rights Policy Template (Using Windows PowerShell)
- AD RMS Rights Policy Templates Deployment Step-by-Step Guide
- AD RMS Rights Policy Template Considerations
- AD RMS Rights Policy Templates Best Practices

Multi-machine considerations

In working to plan encryption for the files you use in your organization, there are a number of aspects to consider in preparing for multi-machine deployment. To begin with, there is no easy way to push file management tasks and classification rules out to multiple computers. Therefore, it is common practice for those working with dynamic access control to do new configuration by transferring settings from one computer to another manually, although there are some tools that can assist in this process.

Dynamic scope using the FolderUsage property

For each file management task and classification rule, you can use a management property in a way that allows for dynamic scoping in contrast to a statically specified scope. In static scoping, a file management task, a classification rule or a report is defined using specific folder shares, paths or volumes that are specific to the computer where the rule was made. Since each computer has a different share / folder / volume structure, a statically scoped rule is likely not work if copied to other computers.

Dynamic scoping allows a task, rule or report to be defined in such a way that it can be correct calculated on any computer where the task, rule or report is placed. A dynamically scoped rule might use a management property to help calculate and apply the rule to folders by looking at how they are marked for usage. For example, you might have two servers, Server1 and Server2. On Server1, D:\share might contain files with user data and on Server2, E:\share is the share used for the same purpose. To apply dynamic scope, you could then specify a file management task that applies to all folders marked as containing FolderUsage="User Data". By doing so, the rule would work correctly on both servers with no modification.

Note

For dynamic scope, always use the FolderUsage property definition, one of three special management property definitions that are built in to the File Services infrastructure.

Setting management property values

While "User Data" (as was described in the previous section) is one possible value for the FolderUsage management property that could be applied, others include "Collaboration Data" or "Application Data" and these classifications can be modified, extended or deleted as needed. If you do want to modify a management property such as FolderUsage, the following are all possible ways to do so:

1. You can set the management property in the New Share Wizard.
2. In the File Services Resource Manager (FSRM) console, in the treeview on the left, select **Classification Properties**, and then select **Set Management Properties** in the task pane on the right.
3. Using Windows PowerShell, you can use the **Get-FSRMMgmtProperty** and **Set-FSRMMgmtProperty** cmdlets to view status on management properties or configure them.

Moving configurations between computers

Once you have set or updated management properties on a single computer you will want to be able to move your configuration changes to others computers. The following are three possible options for doing so:

- **Write a script that configures File Classification Infrastructure (FCI).**

  To alter the FCI configuration on multiple computers, you can write a COM-based script that uses the FCI API to create new file classifications and to store new properties. You can then create applications or scripts that can be initiated by file management tasks in File Services Resource Manager (FSRM) or you can also use Group Policy to launch completed scripts. For more information on the FCI API, see IFsrmPipelineModuleImplementation Interface
- **Use Windows PowerShell remoting to use the File Services Resource Manager (FSRM) cmdlets to configure computers.**

  Windows PowerShell provides another scripting platform you can leverage to perform configuration updates. For more information on using Windows PowerShell to create FCI scripts, see Using Windows PowerShell Script for File Classification.
- **Use the Data Classification Toolkit to export and then import configurations.**

  The Data Classification Toolkit provides another option for exporting and importing FCI configurations. It was designed to help enable an organization to identify, classify, and protect data on its file servers and it provides out-of-the-box classification and rule examples that can help organizations build and deploy their policies to protect critical information in a cost-effective manner. For more information, see Data Classification Toolkit.

Deploy Encryption of Office Files (Demonstration Steps)

Contoso's Finance Department has a number of file servers that store their documents. These documents can be general documentation or they can have a high-business impact (HBI). For example, any document that contains confidential information is deemed, by Contoso, to have a high-business impact. Contoso wants to ensure that all their documentation has a minimum amount of protection and that their HBI documentation is restricted to the appropriate people. To accomplish this, Contoso is exploring using the File Classification Infrastructure (FCI) and AD RMS that is available in Windows Server 2012. By using FCI, Contoso will classify all of the documents on their file server, based on the content, and then use AD RMS to apply the appropriate rights policy.

In this scenario, you'll perform the following steps:

Task	Description
Step 1: Enable resource properties	Enable the **Impact** and **Personally Identifiable Information** resource properties.
Step 2: Create classification rules	Create the following classification rules: **HBI Classification Rule** and **PII Classification Rule**.
Step 3: Use file management tasks to automatically protect documents with AD RMS	Create a file management task that automatically used AD RMS to protect documents with high personally identifiable information (PII). Only members of the FinanceAdmin group will have access to documents that contain high PII.
Step 4: View the results	Examine the classification of documents and observe how they change as you change the content in the document. Also verify how the document gets protected by AD RMS.
Step 5: Verify protection with AD RMS	Verify that the document is protected with AD RMS.

Step 1: Enable resource properties

To enable resource properties

1. In Hyper-V Manager, connect to server ID_AD_DC1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. Open Active Directory Administrative Center, and click **Tree View**.
3. Expand **DYNAMIC ACCESS CONTROL**, and select **Resource Properties**.
4. Scroll down to the **Impact** property in the **Display name** column. Right-click **Impact**, and then click **Enable**.
5. Scroll down to the **Personally Identifiable Information** property in the **Display name** column. Right-click **Personally Identifiable Information**, and then click **Enable**.
6. To publish the resource properties in the **Global Resource List**, in the left pane, click **Resource Property Lists**, and then double-click **Global Resource Property List**.
7. Click **Add**, and then scroll down to and click **Impact** to add it to the list. Do the same for **Personally Identifiable Information**. Click **OK** twice to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Impact_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=PII_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Create classification rules

This step explains how to create the **High Impact** classification rule. This rule will search the content of documents and if the string “Contoso Confidential” is found, it will classify this document as having high-business impact. This classification will override any previously assigned classification of low-business impact.

You will also create a **High PII** rule. This rule searches the content of documents, and if a Social Security number is found, it classifies the document as having high PII.

To create the high-impact classification rule

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. You need to refresh the Global Resource Properties from Active Directory. Open Windows PowerShell and type: `Update-FSRMClassificationPropertyDefinition`, and then press ENTER. Close Windows PowerShell.
3. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
4. In the left pane of File Server Resource Manager, expand **Classification Management**,

and then select **Classification Rules**.

5. In the **Actions** pane, click **Configure Classification Schedule**. On the **Automatic Classification** tab, select **Enable fixed schedule**, select a **Day of the week**, and then select the **Allow continuous classification for new files** check box. Click **OK**.
6. In the **Actions** pane, click **Create Classification Rule**. This opens the **Create Classification Rule** dialog box.
7. In the **Rule name** box, type **High Business Impact**.
8. In the **Description** box, type **Determines if the document has a high business impact based on the presence of the string "Contoso Confidential"**
9. On the **Scope** tab, click **Set Folder Management Properties**, select **Folder Usage**, click **Add**, then click **Browse**, browse to D:\Finance Documents as the path, click **OK**, and then choose a property value named **Group Files** and click **Close**. Once management properties are set, on the **Rule Scope** tab select **Group Files**.
10. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list.
11. Under **Choose a property to assign to files**, select **Impact** from the drop-down list.
12. Under **Specify a value**, select **High** from the drop-down list.
13. Click **Configure** under **Parameters**. In the **Classification Parameters** dialog box, in the **Expression Type** list, select **String**. In the **Expression** box, type: **Contoso Confidential**, and then click **OK**.
14. Click the **Evaluation Type** tab. Click **Re-evaluate existing property values**, click **Overwrite** the existing value, and then click **OK** to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Update-FSRMClassificationPropertyDefinition

$date = Get-Date

$AutomaticClassificationScheduledTask = New-FsrmScheduledTask -Time $date -Weekly @(3, 2, 4, 5,1,6,0) -RunDuration 0;

Set-FsrmClassification -Continuous -schedule $AutomaticClassificationScheduledTask

New-FSRMClassificationRule -Name "High Business Impact" -Property "Impact_MS" -Description "Determines if the document has a high business impact based on the presence of the string 'Contoso Confidential'" -PropertyValue "3000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("StringEx=Min=1;Expr=Contoso Confidential") -ReevaluateProperty Overwrite
```

To create the high-PII classification rule

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.

2. On the desktop, open the folder named **Regular Expressions**, and then open the text document named **RegEx-SSN**. Highlight and copy the following regular expression string: **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**. This string will be used later in this step so keep it on your clipboard.
3. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
4. In the left pane of File Server Resource Manager, expand **Classification Management**, and then select **Classification Rules**.
5. In the **Actions** pane, click **Configure Classification Schedule**. On the **Automatic Classification** tab, select **Enable fixed schedule**, select a **Day of the week**, and then select the **Allow continuous classification for new files** check box. Click OK.
6. In the **Rule name** box, type **High PII**. In the **Description** box, type **Determines if the document has a high PII based on the presence of a Social Security Number.**
7. Click the **Scope** tab, select the **Group Files** check box.
8. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list.
9. Under **Choose a property to assign to files**, select **Personally Identifiable Information** from the drop-down list.
10. Under **Specify a value**, select **High** from the drop-down list.
11. Click **Configure** under **Parameters**.
    In the **Classification Parameters** window, in the **Expression Type** list, select **Regular Expression**. In the **Expression** box, paste the text from your clipboard: **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**, and then click **OK**.

    Note
    This expression will allow invalid Social Security numbers. This allows us to use fictitious Social Security numbers in the demonstration.
12. Click the **Evaluation Type** tab. Select **Re-evaluate existing property values**, **Overwrite** the existing value, and then click **OK** to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-FSRMClassificationRule -Name "High PII" -Description "Determines if the document has a high PII based on the presence of a Social Security Number." -Property "PII_MS" -PropertyValue "5000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("RegularExpressionEx=Min=1;Expr=^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -]?)(?!00)\d\d\3(?!0000)\d{4}$") -ReevaluateProperty Overwrite
```

You should now have two classification rules:

- High Business Impact

- High PII

Step 3: Use file management tasks to automatically protect documents with AD RMS

Now that you've created rules to automatically classify documents based on content, the next step is to create a file management task that uses AD RMS to automatically protect certain documents based on their classification. In this step, you will create a file management task that automatically protects any documents with a high PII. Only members of the FinanceAdmin group will have access to documents that contain high PII.

To protect documents with AD RMS

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
3. In the left pane, select **File Management Tasks**. In the **Actions** pane, select **Create File Management Task**.
4. In the **Task name:** field, type **High PII**. In the **Description** field, type **Automatic RMS protection for high PII documents**.
5. Click the **Scope** tab, select the **Group Files** check box.
6. Click the **Action** tab. Under **Type**, select **RMS Encryption**. Click **Browse** to select a template, and then select the **Contoso Finance Admin Only** template.
7. Click the **Condition** tab, and then click **Add**. Under **Property**, select **Personally Identifiable Information**. Under **Operator**, select **Equal**. Under **Value**, select **High**. Click **OK**.
8. Click the **Schedule** tab. In the **Schedule** section, click **Weekly**, and then select **Sunday**. Running the task once-a-week will ensure that you catch any documents that may have been missed due to a service outage or other disruptive event.
9. In the **Continuous operation** section, select **Run task continuously on new files**, and then click **OK**. You should now have a file management task named High PII.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$fmjRmsEncryption = New-FSRMFmjAction -Type 'Rms' -RmsTemplate 'Contoso Finance Admin Only'

$fmjCondition1 = New-FSRMFmjCondition -Property 'PII_MS' -Condition 'Equal' -Value '5000'

$date = get-date
```

```
$schedule = New-FsrmScheduledTask -Time $date -Weekly @('Sunday')
$fmj1=New-FSRMFileManagementJob -Name "High PII" -Description "Automatic RMS protection
for high PII documents" -Namespace @('D:\Finance Documents') -Action $fmjRmsEncryption -
Schedule $schedule -Continuous -Condition @($fmjCondition1)
```

Step 4: View the results

It's time to take a look at your new automatic classification and AD RMS protection rules in action. In this step you will examine the classification of documents and observe how they change as you change the content in the document.

To view the results

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. In Windows Explorer, navigate to D:\Finance Documents.
3. Right-click the Finance Memo document and click **Properties**.Click the **Classification** tab, and notice that the Impact property currently has no value. Click **Cancel**.
4. Right-click the **Request for Approval to Hire document**, and then select **Properties**.
5. Click the **Classification** tab, and notice that **the Personally Identifiable Information property** currently has no value. Click **Cancel**.
6. Switch to CLIENT1. Sign off any user who is signed in, and then sign in as Contoso\MReid with the password **pass@word1**.
7. From the Desktop, open the **Finance Documents** shared folder.
8. Open the **Finance Memo** document. Near the bottom of the document, you will see the word **Confidential**. Modify it to read: **Contoso Confidential**. Save the document and close it.
9. Open the **Request for Approval to Hire** document. In the **Social Security#:** section, type: 777-77-7777. Save the document and close it.

   Note
   You may need to wait 30 seconds for the classification to occur.
10. Switch back to ID_AD_FILE1. In Windows Explorer, navigate to D:\Finance Documents.
11. Right-click the Finance Memo document, and click **Properties**. Click the **Classification** tab. Notice that the **Impact** property is now set to **High**. Click **Cancel**.
12. Right-click the Request for Approval to Hire document and click **Properties**.
13. . Click the **Classification** tab. Notice that the **Personally Identifiable Information** property is now set to **High**. Click **Cancel**.

Step 5: Verify protection with AD RMS

To verify that the document is protected

1. Switch back to ID_AD_CLIENT1.
2. Open the **Request for approval to Hire** document.
3. Click **OK** to allow the document to connect to your AD RMS server.
4. You can now see that the document has been protected by AD RMS because it contains a Social Security number.

Scenario: Get Insight into Your Data by Using Classification

Reliance on data and storage resources has continued to grow in importance for most organizations. IT administrators face the growing challenge of overseeing larger and more complex storage infrastructures, while simultaneously being tasked with the responsibility to ensure that total cost-of-ownership is maintained at reasonable levels. Managing storage resources is not only about the volume or availability of data; it is also about enforcing company policies and knowing how storage is consumed to enable efficient utilization and compliance to mitigate risk. File Classification Infrastructure provides insight into your data by automating classification processes so that you can manage your data more effectively. The following classification methods are available with File Classification Infrastructure: manual, programmatic, and automatic. This topic focuses on the automatic file classification method.

Scenario description

File Classification Infrastructure uses classification rules to automatically scan files and classify them according to the contents of the file. Classification properties are defined centrally in Active Directory so that these definitions can be shared across file servers in the organization. You can create classification rules that scan files for a standard string or for a string that matches a pattern (regular expression). When a configured classification parameter is found in a file, that file is classified as configured in the classification rule. Some examples of classification rules include:

- Classify any file that contains the string “Contoso Confidential” as having high business impact
- Classify any file that contains at least 10 social security numbers as having personally identifiable information

When a file is classified, you can use a file management task to take action on any files that are classified a specific way. The actions in a file management task include protecting the rights

associated with the file, expiring the file, and running a custom action (such as posting information to a web service).

You can find planning information for configuring automatic file classification in Plan for Automatic File Classification.

You can find steps for how to automatically classify files in Deploy Automatic File Classification (Demonstration Steps).

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Practical applications

File Classification Infrastructure in Windows Server 2012 contributes to Dynamic Access Control by enabling business data owners to easily classify and label data. The classification information that is stored in the central access policy allows you to define access policies for data classes that are critical to business.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	File Classification Infrastructure is a feature that is included in File Server Resource Manager.
File and Storage Services Overview	File Server Resource Manager is a feature that is included with the File Services server role.

Plan for Automatic File Classification

Before you deploy file classification across the file servers in your organization, you should identify the information that should be classified and identify the appropriate classification method.

Use the following table to plan your automatic file classification deployment in your organization.

Task	Description
1.1. Identify what information to classify in your environment	Inventory your existing data and decide what information needs to be classified before creating rules.
1.2. Identify how to classify files	Choose the appropriate classification method for the classification rule: manual, location-based, or content-based.
1.3. Considerations for multiple computers	Use the Data Classification Toolkit when possible to export the configuration from a baseline computer and import it on the file servers.

1.1. Identify what information to classify in your environment

You should start by doing an inventory on the existing data on all files in your organization. From there, you should list the classification requirements and use that list to figure out which file and folder should be classified. Be sure to ensure that you consider any regulations, such as HIPAA, while you are listing the classification requirements.

Before you create any new classification properties, ensure that a matching resource property does not already exist in Active Directory Domain Services. If one already exists, you should enable that resource property before creating a new one.

1.2. Identify how to classify files

When you identify the files that should be classified, you should include both new and existing files in your organization. There are three ways to classify files:

- **Manual** You can classify files manually by using the **Classification** tab of the properties sheet of the file.
- **Location-based** Location-based classification can be used by classifying folders manually using the **Classification** tab or by using the folder classifier in a classification rule.
- **Content-based** Content-based classification can be deployed by using the content classifier in a classification rule. The Data Classification Toolkit includes some built-in rules for determining personally identifiable information. More information on the Data Classification Toolkit can be found on the TechNet website.

If the folder and content classifiers do not fit the requirements of your organization, you can use the Windows PowerShell classifier or purchase a non-Microsoft classifier. The Windows PowerShell classifier allows you to write a Windows PowerShell script that returns true or false. If true, the file is classified according to the classification rule.

1.3. Considerations for multiple computers

There are several things to consider when you have more than one file server in your organization:

- The Data Classification Toolkit uses Windows PowerShell cmdlets to import and export classification rules. You should be this to export the configuration from a baseline computer and import to another computer to ensure that the configuration is the same.
- You should use dynamic name spaces when the source and destination servers use the same drive letters for the storage on the server. When you create a new file share by using Server Manager, you can specify the name space. For more information about dynamic name spaces, see What's New in File Server Resource Manager in Windows Server 2012.

See also

- Scenario: Get Insight into Your Data by Using Classification
- Deploy Automatic File Classification (Demonstration Steps)
- Dynamic Access Control: Scenario Overview
- TechNet website

Deploy Automatic File Classification (Demonstration Steps)

This topic explains how to enable resource properties in Active Directory, create classification rules on the file server, and then assign values to the resource properties for files on the file server. For this example, the following classification rules are created:

- A content classification rule that searches a set of files for the string “Contoso Confidential.” If the string is found in a file, the Impact resource property is set to High on the file.
- A content classification rule that searches a set of files for a regular expression that matches a social security number at least 10 times in one file. If the pattern is found, the file is classified as having personally identifiable information and the Personally Identifiable Information resource property is set to High.

In this document

- **Step 1: Create resource property definitions**
- **Step 2: Create a string content classification rule**
- **Step 3: Create a regular expression content classification rule**
- Step 4: Verify that the files are classified correctly

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Create resource property definitions

The Impact and Personally Identifiable Information resource properties are enabled so that File Classification Infrastructure can use these resource properties to tag the files that are scanned on a network shared folder.

Do this step using Windows PowerShell

To create resource property definitions

1. On the domain controller, sign in to the server as a member of the Domain Admins security group.
2. Open Active Directory Administrative Center. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Expand **Dynamic Access Control**, and then click **Resource Properties**.
4. Right-click **Impact**, and then click **Enable**.
5. Right-click **Personally Identifiable Information**, and then click **Enable**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Impact_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=PII_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Create a string content classification rule

A string content classification rule scans a file for a specific string. If the string is found, the value of a resource property can be configured. In this example, we will scan each file on a network shared folder and look for the string "Contoso Confidential." If the string is found, the associated file is classified as having high business impact.

Do this step using Windows PowerShell

To create a string content classification rule

1. Log on to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition** and then press ENTER. This will synchronize the property definitions created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Expand **Classification Management**, right-click **Classification Rules**, and then click

Configure Classification Schedule.

5. Select the **Enable fixed schedule** check box, select the **Allow continuous classification for new files** check box, choose a day of the week to run the classification, and then click **OK**.
6. Right-click **Classification Rules**, and then click **Create Classification Rule**.
7. On the **General** tab, in the **Rule name** box, type a rule name such as **Contoso Confidential**.
8. On the **Scope** tab, click **Add**, and choose the folders that should be included in this rule, such as D:\Finance Documents.

   Note
   You can also choose a dynamic name space for the scope. For more information about dynamic name spaces for classification rules, see What's New in File Server Resource Manager in Windows Server 2012.
9. On the **Classification** tab, configure the following:
   - In the **Choose a method to assign a property to files** box, ensure that **Content Classifier** is selected.
   - In the **Choose a property to assign to files** box, click **Impact**.
   - In the **Specify a value** box, click **High**.
10. Under the **Parameters** heading, click **Configure**.
11. In the **Expression Type** column, select **String**.
12. In the **Expression** column, type **Contoso Confidential**, and then click **OK**.
13. On the **Evaluation Type** tab, select the **Re-evaluate existing property values** check box, click **Overwrite the existing value**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$date = Get-Date

$AutomaticClassificationScheduledTask = New-FsrmScheduledTask -Time $date -Weekly @(3, 2, 4, 5,1,6,0) -RunDuration 0;$AutomaticClassificationScheduledTask

Set-FsrmClassification -Continuous -schedule $AutomaticClassificationScheduledTask

New-FSRMClassificationRule -Name "Contoso Confidential" -Property "Impact_MS" -PropertyValue "3000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("StringEx=Min=1;Expr=Contoso Confidential") -ReevaluateProperty Overwrite
```

Step 3: Create a regular expression content classification rule

A regular expression classification rule scans a file for a pattern that matches the regular expression. If a string that matches the regular expression is found, the value of a resource property can be configured. In this example, we will scan each file on a network shared folder and look for a string that matches the pattern of a social security number (XXX-XX-XXXX). If the pattern is found, the associated file is classified as having personally identifiable information.

Do this step using Windows PowerShell

To create a regular expression content classification rule

1. Sign in to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition**, and then press ENTER. This will synchronize the property definitions that are created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Right-click **Classification Rules**, and then click **Create Classification Rule**.
5. On the **General** tab, in the **Rule name** box, type a name for the classification rule, such as PII Rule.
6. On the **Scope** tab, click **Add**, and then choose the folders that should be included in this rule, such as D:\Finance Documents.
7. On the **Classification** tab, configure the following:
   - In the **Choose a method to assign a property to files** box, ensure that **Content Classifier** is selected.
   - In the **Choose a property to assign to files** box, click **Personally Identifiable Information**.
   - In the **Specify a value** box, click **High**.
8. Under the **Parameters** heading, click **Configure**.
9. In the **Expression Type** column, select **Regular expression**.
10. In the **Expression** column, type **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**
11. In the **Minimum Occurrences** column, type **10**, and then click **OK**.
12. On the **Evaluation Type** tab, select the **Re-evaluate existing property values** check box, click **Overwrite the existing value**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-FSRMClassificationRule -Name "PII Rule" -Property "PII_MS" -PropertyValue "5000" -
Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -
Parameters @("RegularExpressionEx=Min=10;Expr=^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -
]?)(?!00)\d\d\3(?!0000)\d{4}$") -ReevaluateProperty Overwrite
```

Step 4: Verify that the files are classified correctly

You can verify that the files are properly classified by viewing the properties of a file that was created in the folder specified in the classification rules.

To verify that the files are classified correctly

1. On the file server, run the classification rules by using File Server Resource Manager.
   a. Click **Classification Management**, right-click **Classification Rules**, and then click **Run Classification With All Rules Now**.
   b. Click the **Wait for classification to complete** option, and then click **OK**.
   c. Close the Automatic Classification Report.
   d. You can do this by using Windows PowerShell with the following command: **Start-FSRMClassification –RunDuration 0 -Confirm:$false**
2. Navigate to the folder that was specified in the classification rules, such as D:\Finance Documents.
3. Right-click a file in that folder, and then click **Properties**.
4. Click the **Classification** tab, and verify that the file is classified correctly.

See also

- Scenario: Get Insight into Your Data by Using Classification
- Plan for Automatic File Classification
- Dynamic Access Control: Scenario Overview

Set up Manual File Classification

Users can manually classify files and folders using the File Classification Infrastructure in Windows Server 2012. Users can tag their data by the resource properties defined or enabled in the domain controller.

The following procedures provide manual file classification in your organization:

- Create/enable resource properties
- Set group policy settings to enable the manual classification UI on Windows 8 devices
- Classify the files and folders based on your resource properties

Create Resource Properties

The following procedure demonstrates how to create two resource properties: Country and Department.

To create and enable pre-created resource properties

1. In the left pane of Active Directory Administrative Center, click **Tree View**. Expand **Dynamic Access Control**, and then select **Resource Properties**.
2. Right-click **Resource Properties**, click **New**, and then click **Reference Resource Property**.

   Tip

   You can also choose a resource property from the **Tasks** pane. Click **New** and then click **Reference Resource Property**.
3. In **Select a claim type to share its suggested values list**, click **country**.
4. In the **Display name** field, type **country**, and then click **OK**.
5. Double-click the **Resource Properties** list, scroll down to the **Department** resource property. Right-click, and then click **Enable**. This will enable the built-in **Department** resource property.
6. In the **Resource Properties** list on the navigation pane of the Active Directory Administrative Center, you will now have two enabled resource properties:
   - Country
   - Department

Set Group Policy Settings for Manual File Classification

You need to set the following two group policy settings in order to enable the manual classification UI on Windows 8 devices:

- **File Classification Infrastructure: Display Classification tab in File Explorer**: This setting controls the display of the classification tab in the Properties tab of File Explorer.
- **File Classification Infrastructure: Specify classification properties list**: This setting controls which set of resource properties are displayed in the classification tab in your resource location. (see resource properties list section below)

To edit group policy settings for manual classification:

1. Open Group Policy Management, select your domain, and then click **Domain Controllers**.
2. Right-click **Default Domain Controllers Policy**, and then click **Edit**.
3. **To enable group policy settings to display Classification tab in File Explorer**:

   In the Group Policy Management Editor window, double-click **Computer Configuration**,

double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **File Classification Infrastructure**.

4. Double-click **File Classification Infrastructure: Display Classification tab in File Explorer**. In the File Classification Infrastructure: Display Classification tab in File Explorer dialog box, select **Enabled**
5. **To enable group policy setting to specify classification properties list**:

   Double-click **File Classification Infrastructure: Specify classification properties list**. In the File Classification Infrastructure: Specify classification properties list dialog box, select **Global Properties list** and then select **Enabled**.

   Note
   Specify the resource property list you want displayed in the **Classification Properties list** box. By default, the Global Resource Properties List is used if you do not specify anything in the **Classification Properties List** box.
6. Close Group Policy Management.
7. On the Windows 8 client, open a command prompt and type **gpupdate /force**.

Classify files and folders manually

To manually classify files and folder

1. Navigate to the folder on your file server that contains files that you want to classify.
2. Right-click a file in that folder, and then click Properties.
3. Click the Classification tab, select the resource property you want to tag the folder and click the value, and click OK.

Classification properties lists

To allow you to expose different sets of classification properties to different users, you can create classification properties lists that contain specific resource properties. For example, you might want all the HR employees to set a resource property called “HR file type” and the finance department employees to set a resource property called “Finance file type”

In order to expose HR and Finance with different resource properties, you can create two classification properties lists in the Dynamic Access Control node in Active Directory and populate them with the appropriate resource properties (note that the built-in Global properties list is populated by default)Then in the group policy for the HR department, use the “HR properties list” and for the Finance department, use the “Finance properties list”

Scenario: Implement Retention of Information on File Servers

A retention period is the amount of time that a document should be kept before it is expired. Depending on the organization, the retention period can be different. You can classify files in a folder as having a short, medium, or long-term retention period, and then assign a timeframe for each period. You may want to keep a file indefinitely by putting it on legal hold.

Scenario description

File Classification Infrastructure and File Server Resource Manager uses file management tasks and file classification to apply retention periods for a set of files. You can assign a retention period on a folder and then use a file management task to configure how long an assigned retention period is to last. When the files in the folder are about to expire, the owner of the file receives a notification email. You can also classify a file as being on legal hold so that the file management task will not expire the file.

You can find planning information for configuring retention in Plan for Retention of Information on File Servers.

You can find steps for classifying files for legal hold and configuring a retention period in Deploy Implementing Retention of Information on File Servers (Demonstration Steps).

Note
That scenario only discusses how to manually classify a document for legal hold. However, it is possible in Windows Server 2012 to automatically classify documents for legal hold. One way to do this is to create a Windows PowerShell classifier that compares the file owner to a list of user accounts that are under legal hold. If the file owner is a part of the user account list, the file is classified for legal hold.

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	File Management Tasks is a feature that is

Feature	How it supports this scenario
	included in File Server Resource Manager.
File and Storage Services Overview	File Server Resource Manager is a feature that is included with the File Services server role.

Plan for Retention of Information on File Servers

A data retention policy is important for any organization.

Use the following table to plan how you retain information in your organization.

Task	Description
1.1. Determine the retention schedule	Determining how long to retain data reduces your storage requirements, reduces liability, and satisfies regulations.
1.2. Identify files to be retained	Before you implement a data retention policy, you must first know the type of data that is stored on each file server in your organization.
1.3 Considerations for multiple computers	Use the Data Classification Toolkit when possible to export the configuration from a baseline computer and import it on the file servers.

1.1. Determine the retention schedule

Determining how long data is stored on file servers in your network and developing a data retention schedule offers the following advantages:

- Limits the amount of data to store which lowers the overall cost of storage in your organization
- Reduces liability
- Satisfies regulations

When you are determining your retention schedule, you should consult an attorney to ensure that the retention schedule meets the regulatory compliance of the industry your organization is in. For example, if your company is a healthcare provider, you should understand the HIPAA regulations.

1.2. Identify files to be retained

Before you can implement your data retention policy, you must first identify the data that is stored on each file server. Once the data has been identified, you should mark the folders with a retention period and retention start date. Also, you should remove the Delete Child permission from any folders that could contain files that are being retained. This will ensure that the files are not accidentally deleted.

Important
Classification properties should not be specified as a date. You should use the retention period to classify the file. If the retention period should change, you can update the retention period interval without having to classify every file again.

1.3 Considerations for multiple computers

There are several things to consider when you have more than one file server in your organization:

- Data retention policies are usually the same across the organization so you can reuse the same set of rules on multiple computers.
- The Data Classification Toolkit uses Windows PowerShell cmdlets to import and export classification rules. You should be this to export the configuration from a baseline computer and import to another computer to ensure that the configuration is the same.
- You should use dynamic name spaces when the source and destination servers use the same drive letters for the storage on the server. When you create a new file share by using Server Manager, you can specify the name space. For more information about dynamic name spaces, see What's New in File Server Resource Manager in Windows Server 2012.

See also

- Scenario: Implement Retention of Information on File Servers
- Deploy Implementing Retention of Information on File Servers (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Deploy Implementing Retention of Information on File Servers (Demonstration Steps)

You can set retention periods for folders and put files on legal hold by using File Classification Infrastructure and File Server Resource Manager.

In this document

- **Prerequisites**

- **Step 1: Create resource property definitions**
- Step 2: Configure notifications
- **Step 3: Create a file management task**
- Step 4: Classify a file manually

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

The steps in this topic assume you have a SMTP server configured for file expiration notifications.

Step 1: Create resource property definitions

In this step, we enable the Retention Period and Discoverability resource properties so that File Classification Infrastructure can use these resource properties to tag the files that are scanned in a network shared folder.

Do this step using Windows PowerShell

To create resource property definitions

1. On the domain controller, sign in to the server as a member of the Domain Admins security group.
2. Open Active Directory Administrative Center. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Expand **Dynamic Access Control**, and then click **Resource Properties**.
4. Right-click **Retention Period**, and then click **Enable**.
5. Right-click **Discoverability**, and then click **Enable**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=RetentionPeriod_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Discoverability_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Configure notifications

In this step, we use the File Server Resource Manager console to configure the SMTP server, the default administrator email address, and the default email address that the reports are sent from.

Do this step using Windows PowerShell

To configure notifications

1. Sign in to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition**, and then press ENTER. This will synchronize the property definitions that are created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Right-click **File Server Resource Manager (local)**, and then click **Configure Options**.
5. On the **Email Notifications** tab, configure the following:
   - In the **SMTP server name or IP address** box, type the name of the SMTP server on your network.
   - In the **Default administrator recipients** box, type the email address of the administrator who should get the notification.
   - In the **Default "From" e-mail address** box, type the email address that should be used to send the notifications.
6. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FsrmSetting -SmtpServer IP address of SMTP server -FromEmailAddress
"FromEmailAddress" -AdminEmailAddress "AdministratorEmailAddress"
```

Step 3: Create a file management task

In this step, we use the File Server Resource Manager console to create a file management task that will run on the last day of the month and expire any files with the following criteria:

- The file is not classified as being on legal hold.
- The file is classified as having a long-term retention period.
- The file has not been modified in the last 10 years.

Do this step using Windows PowerShell

To create a file management task

1. Sign in to the file server as a member of the Administrators security group.
2. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
3. Right-click **File Management Tasks**, and then click **Create File Management Task**.
4. On the **General** tab, in the **Task name** box, type a name for the file management task,

such as Retention Task.

5. On the **Scope** tab, click **Add**, and choose the folders that should be included in this rule, such as D:\Finance Documents.
6. On the **Action** tab, in the **Type** box, click **File expiration**. In the **Expiration directory** box, type a path to a folder on the local file server where the expired files will be moved. This folder should have an access control list that grants only file server administrators access.
7. On the **Notification** tab, click **Add**.
   - Select the **Send e-mail to the following administrators** check box.
   - Select the **Send an email to users with affected files** check box, and then click **OK**.
8. On the **Condition** tab, click **Add**, and add the following properties:
   - In the **Property** list, click **Discoverability**. In the **Operator** list, click **Not equal**. In the **Value** list, click **Hold**.
   - In the **Property** list, click **Retention Period**. In the **Operator** list, click **Equal**. In the **Value** list, click **Long-Term**.
9. On the **Condition** tab, select the **Days since file was last modified** check box, and then set the value to **3650**.
10. On the **Schedule** tab, click the **Monthly** option, and then select the **Last** check box.
11. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$fmjexpiration = New-FSRMFmjAction -Type 'Expiration' -ExpirationFolder folder

$fmjNotificationAction = New-FsrmFmjNotificationAction -Type Email -MailTo
"[FileOwner],[AdminEmail]"

$fmjNotification = New-FsrmFMJNotification -Days 10 -Action @($fmjNotificationAction)

$fmjCondition1 = New-FSRMFmjCondition -Property 'Discoverability_MS' -Condition
'NotEqual' -Value "Hold"

$fmjCondition2 = New-FSRMFmjCondition -Property 'RetentionPeriod_MS' -Condition 'Equal' -
Value "Long-term"

$fmjCondition3 = New-FSRMFmjCondition -Property 'File.DateLastAccessed' -Condition
'Equal' -Value 3650

$date = get-date
$schedule = New-FsrmScheduledTask -Time $date -Monthly @(-1)

$fmj1=New-FSRMFileManagementJob -Name "Retention Task" -Namespace @('D:\Finance
Documents') -Action $fmjexpiration -Schedule $schedule -Notification @($fmjNotification)
-Condition @( $fmjCondition1, $fmjCondition2, $fmjCondition3)
```

Step 4: Classify a file manually

In this step, we manually classify a file to be on legal hold. The parent folder of this file will be classified with a long-term retention period.

To manually classify a file

1. Sign in to the file server as a member of the Administrators security group.
2. Navigate to the folder that was configured in the scope of the file management task created in Step 3.
3. Right-click the folder, and then click **Properties**.
4. On the **Classification** tab, click **Retention Period**, click **Long-Term**, and then click **OK**.
5. Right-click a file within that folder, and then click **Properties**.
6. On the **Classification** tab, click **Discoverability**, click **Hold**, click **Apply**, and then click **OK**.
7. On the file server, run the file management task by using the File Server Resource Manager console. After the file management task completes, check the folder and ensure the file was not moved to the expiration directory.
8. Right-click the same file within that folder, and then click **Properties**.
9. On the **Classification** tab, click **Discoverability**, click **Not Applicable**, click **Apply**, and then click **OK**.
10. On the file server, run the file management task again by using the File Server Resource Manager console. After the file management task completes, check the folder and ensure that file was moved to the expiration directory.

See also

- Scenario: Implement Retention of Information on File Servers
- Plan for Retention of Information on File Servers
- Dynamic Access Control: Scenario Overview

Appendix A: Dynamic Access Control Glossary

Following are the list of terms and definitions that are included in the Dynamic Access Control scenario.

Term	Definition
Automatic classification	Classification that occurs based on classification properties that are determined by

Term	Definition
	classification rules configured by an administrator.
CAPID	Central access policy ID. This ID references a specific central access policy, and it is used to reference the policy from the security descriptor of files and folders.
Central access rule	A rule that includes a condition and an access expression.
Central access policy	Policies that are authored and hosted in Active Directory.
Claims-based access control	A paradigm that utilizes claims to make access control decisions to resources.
Classification	The process of determining the classification properties of resources and assigning these properties to the metadata that is associated with the resources. See also Automatic classification, Inherited classification, and Manual classification.
Device claim	A claim that is associated with the system. With user claims, it is included in the token of a user attempting to access a resource.
Discretionary access control list (DACL)	An access control list that identifies trustees who are allowed or denied access to a securable resource. It can be modified at the discretion of the resource owner.
Resource property	Properties (such as labels) that describe a file and are assigned to files by using automatic classification or manual classification. Examples include: Sensitivity, Project, and Retention period.
File Server Resource Manager	A feature in the Windows Server operating system that offers management of folder quotas, file screening, storage reports, file classification, and file management jobs on a file server.
Folder properties and labels	Properties and labels that describe a folder and are assigned manually by administrators and

Term	Definition
	folder owners. These properties assign default property values to the files within these folders, for example, Secrecy or Department.
Group Policy	A set of rules and policies that controls the working environment of users and computers in an Active Directory environment.
Near real time classification	Automatic classification that is performed shortly after a file is created or modified.
Near real-time file management tasks	File management tasks that are performed shortly after (a file is created or modified. These tasks are triggered by the Near real-time classification.
Organizational Unit (OU)	An Active Directory container that represents hierarchical, logical structures within an organization. It is the smallest scope to which Group Policy settings are applied.
Secure property	A classification property that the authorization runtime can trust to be a valid assertion about the resource at a certain point-in-time. In claims-based access control, a secure property that is assigned to a resource is treated as a resource claim.
Security descriptor	A data structure that contains security information associated with a securable resource, such as access control lists.
Security descriptor definition language	A specification that describes the information in a security descriptor as a text string.
Staging policy	A central access policy that is not yet in effect.
System access control list (SACL)	An access control list that specifies the types of access attempts by particular trustees for which audit records need to be generated.
User claim	Attributes of a user that are provided within the user security token. Examples include: Department, Company, Project, and Security clearance. Information in the user token from systems prior to Windows Server 2012, such as the security groups that the user is part of, can

Term	Definition
	also be considered user claims. Some user claims are provided through Active Directory and others are calculated dynamically, such as whether the user logged in with a smart card.
User token	A data object that identifies a user and the user claims and device claims that are associated with that user. It is used to authorize the user's access to resources.

See Also

Dynamic Access Control: Scenario Overview

Appendix B: Setting Up the Test Environment

This topic outlines the steps to build a hands-on lab to test Dynamic Access Control. The instructions are meant to be followed sequentially because there are many components that have dependencies.

Prerequisites

Hardware and software requirements

Requirements for setting up the test lab:

- A host server running Windows Server 2008 R2 with SP1 and Hyper-V
- A copy of the Windows Server 2012 ISO
- A copy of the Windows 8 ISO
- Microsoft Office 2010
- A server running Microsoft Exchange Server 2003 or later

You need to build the following virtual machines to test the Dynamic Access Control scenarios:

- DC1 (domain controller)
- DC2 (domain controller)
- FILE1 (file server and Active Directory Rights Management Services)
- SRV1 (POP3 and SMTP server)
- CLIENT1 (client computer with Microsoft Outlook)

The passwords for the virtual machines should be as follows:

- BUILTIN\Administrator: pass@word1
- Contoso\Administrator: pass@word1

- All other accounts: pass@word1

Build the test lab virtual machines

Install the Hyper-V role

You need to install the Hyper-V role on a computer running Windows Server 2008 R2 with SP1.

To install the Hyper-V Role

1. Click **Start**, and then click Server Manager.
2. In the Roles Summary area of the Server Manager main window, click **Add Roles**.
3. On the **Select Server Roles** page, click **Hyper-V**.
4. On the **Create Virtual Networks** page, click one or more network adapters if you want to make their network connection available to virtual machines.
5. On the **Confirm Installation Selections** page, click **Install**.
6. The computer must be restarted to complete the installation. Click **Close** to finish the wizard, and then click **Yes** to restart the computer.
7. After you restart the computer, sign in with the same account you used to install the role. After the Resume Configuration Wizard completes the installation, click **Close** to finish the wizard.

Create an internal virtual network

Now you will create an internal virtual network called ID_AD_Network.

To create a virtual network

1. Open Hyper-V Manager.
2. From the **Actions** menu, click **Virtual Network Manager**.
3. Under **Create virtual network**, select the **Internal**.
4. Click **Add**. The **New Virtual Network** page appears.
5. Type **ID_AD_Network** as the name for the new network. Review the other properties and modify them if necessary.
6. Click **OK** to create the virtual network and close Virtual Network Manager, or click **Apply** to create the virtual network and continue using Virtual Network Manager.

Build the domain controller

Build a virtual machine to be used as the domain controller (DC1). Install the virtual machine using Windows Server 2012 ISO, and name it DC1.

To install Active Directory Domain Services

1. Connect the virtual machine to the ID_AD_Network. Sign in to the DC1 as Administrator with the password **pass@word1**.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or Feature-based Install**, and then click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, click **Active Directory Domain Services**. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Domain Services** page, review the information, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**. The Feature installation progress bar on the Results page indicates that the role is being installed.
10. On the **Results** page, verify that the installation succeeded, and click **Close**. In Server Manager, click the warning icon with an exclamation mark on top right corner of the screen, next to **Manage**. In the Tasks list, click the **Promote this server to a domain controller** link.
11. On the **Deployment Configuration** page, click **Add a new forest**, type the name of the root domain, **contoso.com**, and then click **Next**.
12. On the **Domain Controller Options** page, select the domain and forest functional levels as Windows Server 2012, specify the DSRM password **pass@word1**, and then click **Next**.
13. On the **DNS Options** page, click **Next**.
14. On the **Additional Options** page, click **Next**.
15. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder (or accept default locations), and then click **Next**.
16. On the **Review Options** page, confirm your selections, and then click **Next**.
17. On the **Prerequisites Check** page, confirm that the prerequisites validation is completed, and then click **Install**.
18. On the **Results** page, verify that the server was successfully configured as a domain controller, and then click **Close**.
19. Restart the server to complete the AD DS installation. (By default, this happens automatically.)

Create the following users by using Active Directory Administrative Center.

Create users and groups on DC1

1. Sign in to contoso.com as Administrator. Launch Active Directory Administrative Center.
2. Create the following security groups:

Group Name	Email Address
FinanceAdmin	financeadmin@contoso.com
FinanceException	financeexception@contoso.com

3. Create the following organizational unit (OU):

OU Name	Computers
FileServerOU	FILE1

4. Create the following users with the attributes indicated:

User	Username	Email address	Department	Group	Country/Region
Myriam Delesalle	MDelesalle	MDelesalle@contoso.com	Finance		US
Miles Reid	MReid	MReid@contoso.com	Finance	FinanceAdmin	US
Esther Valle	EValle	EValle@contoso.com	Operations	FinanceException	US
Maira Wenzel	MWenzel	MWenzel@contoso.com	HR		US
Jeff Low	JLow	JLow@contoso.com	HR		US
RMS Server	rms	rms@contoso.com			

For more information about creating security groups, see Create a New Group on the Windows Server website.

To create a Group Policy Object

1. Hover the cursor on the upper right corner of screen and click the search icon. In the Search box, type **group policy management**, and click **Group Policy Management**.
2. Expand **Forest: contoso.com**, and then expand **Domains**, navigate to **contoso.com**, expand **(contoso.com)**, and then select **FileServerOU**. Right-click **Create a GPO in this domain and Link it here…**

3. Type a descriptive name for the GPO, such as **FlexibleAccessGPO**, and then click **OK**.

To enable Dynamic Access Control for contoso.com

1. Open the Group Policy Management Console, click **contoso.com**, and then double-click **Domain Controllers**.
2. Right-click **Default Domain Controllers Policy**, and select **Edit**.
3. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.
4. Double-click **KDC support for claims, compound authentication, and Kerberos armoring** and select the option next to **Enabled**. You need to enable this setting to use Central Access Policies.
5. Open an elevated command prompt, and run the following command:

```
gpupdate /force
```

Build the file server and AD RMS server (FILE1)

1. Build a virtual machine with the name FILE1 from the Windows Server 2012 ISO.
2. Connect the virtual machine to the ID_AD_Network.
3. Join the virtual machine to the contoso.com domain, and then sign in to FILE1 as contoso\administrator using the password **pass@word1**.

Install File Services Resource Manager

To install the File Services role and the File Server Resource Manager

1. In Server Manager, click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select Server Roles** page, expand **File and Storage Services**, select the check-box next to **File and iSCSI Services**, expand, and select **File Server Resource Manager**.

   In the Add Roles and Features Wizard, click **Add Features**, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. On the **Installation progress** page, click **Close**.

Install the Microsoft Office Filter Packs on the file server

You should install the Microsoft Office Filter Packs on Windows Server 2012 to enable IFilters for a wider array of Office files than are provided by default. Windows Server 2012 does not have

any IFilters for Microsoft Office Files installed by default, and the file classification infrastructure uses IFilters to perform content analysis.

To download and install the IFilters, see Microsoft Office 2010 Filter Packs.

Configure email notifications on FILE1

When you create quotas and file screens, you have the option of sending email notifications to users when their quota limit is approaching or after they have attempted to save files that have been blocked. If you want to routinely notify certain administrators of quota and file screening events, you can configure one or more default recipients. To send these notifications, you must specify the SMTP server to be used for forwarding the email messages.

To configure email options in File Server Resource Manager

1. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
2. In the File Server Resource Manager interface, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.
3. On the **E-mail Notifications** tab, under SMTP server name or IP address, type the host name or the IP address of the SMTP server that will forward email notifications.
4. If you want to routinely notify certain administrators of quota or file screening events, under **Default administrator recipients**, type each email address such as fileadmin@contoso.com. Use the format account@domain, and use semicolons to separate multiple accounts.

Create groups on FILE1

To create security groups on FILE1

1. Sign in to FILE1 as contoso\administrator, with the password: **pass@word1**.
2. Add NT AUTHORITY\Authenticated Users to the **WinRMRemoteWMIUsers__** group.

Create files and folders on FILE1

1. Create a new NTFS volume on FILE1 and then create the following folder: D:\Finance Documents.
2. Create the following files with the details specified:
   - **Finance Memo.docx**: Add some finance related text in the document. For example, “The business rules about who can access finance documents have changed. Finance documents are now only accessed by members of the FinanceExpert group. No other departments or groups have access.” You need to evaluate the impact of this change before implementing it in the environment. Ensure that this document has CONTOSO CONFIDENTIAL as the footer on every page.

- **Request for Approval to Hire.docx**: Create a form in this document that collects applicant information. You must have the following fields in the document: **Applicant Name, Social Security number, Job Title, Proposed Salary, Starting Date, Supervisor name, Department**. Add an additional section in the document that has a form for **Supervisor Signature, Approved Salary, Conformation of Offer**, and **Status of Offer**.
  Make the document rights-management enabled.
- **Word Document1.docx**: Add some test content to this document.
- **Word Document2.docx**: Add test content to this document.
- **Workbook1.xlsx**
- **Workbook2.xlsx**
- Create a folder on the desktop called Regular Expressions. Create a text document under the folder called **RegEx-SSN**. Type the following content in the file, and then save and close the file:
  ^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$

3. Share the folder D:\Finance Documents as Finance Documents and allow everyone to have Read and Write access to the share.

Note
Central access policies are not enabled by default on the system or boot volume C:.

Install Active Directory Rights Management Services

Add the Active Directory Rights Management Services (AD RMS) and all required features through Server Manager. Choose all the defaults.

To install Active Directory Rights Management Services

1. Sign in to the FILE1 as CONTOSO\Administrator or as a member of the Domain Admins group.

   Important
   In order to install the AD RMS server role the installer account (in this case, CONTOSO\Administrator) will have to be given membership in both the local Administrators group on the server computer where AD RMS is to be installed as well as membership in the Enterprise Admins group in Active Directory.
2. In Server Manager, click **Add Roles and Features**. The Add Roles and Features Wizard appears.
3. On the **Before you Begin** screen, click **Next**.
4. On the **Select Installation Type** screen, click **Role/Feature Based Install**, and then click **Next**.
5. On the **Select Server Targets** screen, click **Next**.
6. On the **Select Server Roles** screen, select the box next to **Active Directory Rights Management Services**, and then click **Next**.

7. In the **Add features that are required for Active Directory Rights Management Services?** dialog box, click **Add Features**.
8. On the **Select Server Roles** screen, click **Next**.
9. On the **Select Features to Install** screen, click **Next**.
10. On the **Active Directory Rights Management Services** screen, click Next.
11. On the **Select Role Services** screen, click **Next**.
12. On the **Web Server Role (IIS)** screen, click **Next**.
13. On the **Select Role Services** screen, click **Next**.
14. On the **Confirm Installation Selections** screen, click **Install**.
15. After the installation has completed, on the **Installation Progress** screen, click **Perform additional configuration**. The AD RMS Configuration Wizard appears.
16. On the **AD RMS** screen, click **Next**.
17. On the **AD RMS Cluster** screen, select **Create a new AD RMS root cluster** and then click **Next**.
18. On the **Configuration Database** screen, click **Use Windows Internal Database on this server**, and then click **Next**.

    Note

    Using the Windows Internal Database is recommended for test environments only because it does not support more than one server in the AD RMS cluster. Production deployments should use a separate database server.
19. On the **Service Account** screen, in **Domain User Account**, click **Specify** and then specify the user name (**contoso\rms**), and Password (**pass@word1**) and click **OK**, and then click **Next**.
20. On the **Cryptographic Mode** screen, click **Cryptographic Mode 2**.
21. On the **Cluster Key Storage** screen, click **Next**.
22. On the **Cluster Key Password** screen, in the **Password** and **Confirm password** boxes, type **pass@word1**, and then click **Next**.
23. On the **Cluster Web Site** screen, make sure that **Default Web Site** is selected, and then click **Next**.
24. On the **Cluster Address** screen, select the **Use an unencrypted connection** option, in the **Fully Qualified Domain Name** box, type **FILE1.contoso.com**, and then click **Next**.
25. On the **Licensor Certificate Name** screen, accept the default name (**FILE1**) in the text box and click **Next**.
26. On the **SCP Registration** screen, select **Register SCP now**, and then click **Next**.
27. On the **Confirmation** screen, click **Install**.
28. On the **Results** screen, click **Close**, and then click **Close** on **Installation Progress** screen. When complete, log off and log on as contoso\rms using the password provided (**pass@word1**).
29. Launch the AD RMS console and navigate to **Rights Policy Templates**.

    To open the AD RMS console, in Server Manager, click **Local Server** in the console tree,

then click **Tools**, and then click **Active Directory Rights Management Services**.

30. Click the **Create Distributed Rights Policy** template located on the right panel, click **Add**, and select the following information:
    - Language: US English
    - Name: Contoso Finance Admin Only
    - Description: Contoso Finance Admin Only

    Click **Add**, and then click **Next**.
31. Under the Users and Rights section, click **Users and rights**, click **Add**, type **financeadmin@contoso.com**, and click **OK**.
32. Select **Full Control**, and leave **Grant owner (author) full control right with no expiration** selected.
33. Click though the remaining tabs with no changes, and then click **Finish**. Sign in as CONTOSO\Administrator.
34. Browse to the folder, C:\inetpub\wwwroot_wmcs\certification, select the ServerCertification.asmx file, and add Authenticated Users to have Read and Write permissions to the file.
35. Open Windows PowerShell and run `Get-FsrmRmsTemplate`. Verify that you are able to see the RMS template you created in the previous steps in this procedure with this command.

Important

If you want your file servers to immediately change so you can test them, you need to do the following:

1. On the file server, FILE1, open an elevated command prompt, and run the following commands:
   - gpupdate /force.
   - NLTEST /SC_RESET:contoso.com
2. On the domain controller (DC1), replicate Active Directory.

   For more information about steps to force the replication of Active Directory, see Active Directory Replication

Optionally, instead of using the Add Roles and Features Wizard in Server Manager, you can use Windows PowerShell to install and configure the AD RMS server role as show in the following procedure.

To install and configure an AD RMS cluster in Windows Server 2012 using Windows PowerShell

1. Logon on as CONTOSO\Administrator with the password: **pass@word1**.

   Important

   In order to install the AD RMS server role the installer account (in this case, CONTOSO\Administrator) will have to be given membership in both the local Administrators group on the server computer where AD RMS is to be installed as

well as membership in the Enterprise Admins group in Active Directory.

2. On the Server desktop, right-click the Windows PowerShell icon on the taskbar and select **Run as Administrator** to open a Windows PowerShell prompt with administrative privileges.
3. To use Server Manager cmdlets to install the AD RMS server role, type:

```
Add-WindowsFeature ADRMS –IncludeAllSubFeature –
IncludeManagementTools
```

4. Create the Windows PowerShell drive to represent the AD RMS server you are installing.

   For example, to create a Windows PowerShell drive named RC to install and configure the first server in an AD RMS root cluster, type:

```
Import-Module ADRMS
New-PSDrive -PSProvider ADRMSInstall -Name RC -Root
RootCluster
```

5. Set properties on objects in the drive namespace that represent required configuration settings.

   For example, to set the AD RMS service account, at the Windows PowerShell command prompt, type:

```
$svcacct = Get-Credential
```

   When the Windows security dialog box appears, type the AD RMS service account domain user name CONTOSO\RMS and the assigned password.

   Next, to assign the AD RMS service account to the AD RMS cluster settings, type the following:

```
Set-ItemProperty –Path RC:\ –Name ServiceAccount –Value
$svcacct
```

   Next, to set the AD RMS server to use the Windows Internal Database, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty –Path RC:\ClusterDatabase –Name
UseWindowsInternalDatabase –Value $true
```

   Next, to securely store the cluster key password in a variable, at the Windows PowerShell command prompt, type:

```
$password = Read-Host –AsSecureString –Prompt "Password:"
```

   Type the cluster key password, and then press the ENTER key.

   Next, to assign the password to your AD RMS installation, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty –Path RC:\ClusterKey –Name
CentrallyManagedPassword –Value $password
```

   Next, to set the AD RMS cluster address, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name ClusterURL -Value
"http://file1.contoso.com:80"
```

Next, to assign the SLC name for your AD RMS installation, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name SLCName -Value "FILE1"
```

Next, to set the service connection point (SCP) for the AD RMS cluster, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name RegisterSCP -Value $true
```

6. Run the **Install-ADRMS** cmdlet. In addition to installing the AD RMS server role and configuring the server, this cmdlet also installs other features required by AD RMS if necessary.

   For example, to change to the Windows PowerShell drive named RC and install and configure AD RMS, type:

   ```
   Set-Location RC:\
   Install-ADRMS –Path.
   ```

   Type "Y" when the cmdlet prompts you to confirm you want to start the installation.

7. Log out as CONTOSO\Administrator and log on as CONTOSO\RMS using the provided password ("pass@word1").

   Important
   In order to manage the AD RMS server the account you are logged on to and using to manage the server (in this case, CONTOSO\RMS) will have to be given membership in both the local Administrators group on the AD RMS server computer as well as membership in the Enterprise Admins group in Active Directory.

8. On the Server desktop, right-click the Windows PowerShell icon on the taskbar and select **Run as Administrator** to open a Windows PowerShell prompt with administrative privileges.
9. Create the Windows PowerShell drive to represent the AD RMS server you are configuring.

   For example, to create a Windows PowerShell drive named RC to configure the AD RMS root cluster, type:

   ```
   Import-Module ADRMSAdmin `
   New-PSDrive -PSProvider ADRMSAdmin -Name RC -Root
   http://localhost -Force -Scope Global
   ```

10. To create new rights template for the Contoso finance administrator and assign it user rights with full control in your AD RMS installation, at the Windows PowerShell command prompt, type:

    ```
    New-Item -Path RC:\RightsPolicyTemplate –LocaleName en-us -
    DisplayName "Contoso Finance Admin Only" -Description
    ```

```
"Contoso Finance Admin Only" -UserGroup
financeadmin@contoso.com  -Right ('FullControl')
```

11. To verify that you can see the new rights template for the Contoso finance administrator, at the Windows PowerShell command prompt:

    ```
    Get-FsrmRmsTemplate
    ```

    Review the output of this cmdlet to confirm the RMS template you created in the previous step is present.

Build the mail server (SRV1)

SRV1 is the SMTP/POP3 mail server. You need to set it up so that you can send email notifications as part of the Access-Denied assistance scenario.

Configure Microsoft Exchange Server on this computer. For more information, see How to Install Exchange Server.

Build the client virtual machine (CLIENT1)

To build the client virtual machine

1. Connect the CLIENT1 to the ID_AD_Network.
2. Install Microsoft Office 2010.
3. Sign in as Contoso\Administrator, and use the following information to configure Microsoft Outlook.
   - Your name: File Administrator
   - Email address: fileadmin@contoso.com
   - Account type: POP3
   - Incoming mail server: Static IP address of SRV1
   - Outgoing mail server: Static IP address of SRV1
   - User name: fileadmin@contoso.com
   - Remember password: Select
4. Create a shortcut to Outlook on the contoso\administrator desktop.
5. Open Outlook and address all the “first time launched” messages.
6. Delete any test messages that were generated.
7. Create a new short cut on desktop for all users on the client virtual machine that points to \\FILE1\Finance Documents.
8. Reboot as needed.

Enable Access-Denied assistance on the client virtual machine

1. Open Registry Editor, and navigate to

HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\Explorer.

- Set **EnableShellExecuteFileStreamCheck** to **1**.
- Value: DWORD

Lab setup for deploying claims across forests scenario

Build a virtual machine for DC2

- Build a virtual machine from the Windows Server 2012 ISO.

- Create the virtual machine name as DC2.
- Connect the virtual machine to the ID_AD_Network.

Important

Joining virtual machines to a domain and deploying claim types across forests require that the virtual machines be able to resolve the FQDNs of the relevant domains. You may have to manually configure the DNS settings on the virtual machines to accomplish this. For more information, see Configuring a virtual network.

All the virtual machine images (servers and clients) must be reconfigured to use a static IP version 4 (IPv4) address and Domain Name System (DNS) client settings. For more information, see Configure a DNS Client for Static IP Address.

Set up a new forest called adatum.com

To install Active Directory Domain Services

1. Connect the virtual machine to the ID_AD_Network. Sign in to the DC2 as Administrator with the password **Pass@word1**.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select Installation Type** page, click **Role-based or Feature-based Install**, and then click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, click the names of the server where you want to install Active Directory Domain Services (AD DS), and then click **Next**.
6. On the **Select Server Roles** page, click **Active Directory Domain Services**. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
7. On the **Select Features** page, click **Next**.
8. On the **AD DS** page, review the information, and then click **Next**.
9. On the **Confirmation** page, click **Install**. The Feature installation progress bar on the

Results page indicates that the role is being installed.

10. On the **Results** page, verify that the installation succeeded, and then click the warning icon with an exclamation mark on top right corner of the screen, next to **Manage**. In the Tasks list, click the **Promote this server to a domain controller** link.

    Important
    If you close the installation wizard at this point rather than click **Promote this server to a domain controller**, you can continue the AD DS installation by clicking **Tasks** in Server Manager.

11. On the **Deployment Configuration** page, click **Add a new forest**, type the name of the root domain, **adatum.com**, and then click **Next**.
12. On the **Domain Controller Options** page, select the domain and forest functional levels as Windows Server 2012, specify the DSRM password **pass@word1**, and then click **Next**.
13. On the **DNS Options** page, click **Next**.
14. On the **Additional Options** page, click **Next**.
15. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder (or accept default locations), and then click **Next**.
16. On the **Review Options** page, confirm your selections, and then click **Next**.
17. On the **Prerequisites Check** page, confirm that the prerequisites validation is completed, and then click **Install**.
18. On the **Results** page, verify that the server was successfully configured as a domain controller, and then click **Close**.
19. Restart the server to complete the AD DS installation. (By default, this happens automatically.)

Important

To ensure that the network is configured properly, after you have set up both the forests, you must do the following:

- Sign in to adatum.com as adatum\administrator. Open a Command Prompt window, type **nslookup contoso.com**, and then press ENTER.
- Sign in to contoso.com as contoso\administrator. Open a Command Prompt window, type **nslookup adatum.com**, and then press ENTER.

  If these commands execute without errors, the forests can communicate with each other. For more information on nslookup errors, see the troubleshooting section in the topic Using NSlookup.exe

Set contoso.com as a trusting forest to adatum.com

In this step, you create a trust relationship between the Adatum Corporation site and the Contoso, Ltd. site.

To set Contoso as a trusting forest to Adatum

1. Sign in to DC2 as administrator. On the **Start** screen, type domain.msc.
2. In the console tree, right-click adatum.com, and then click Properties.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type **contoso.com**, in the Domain Name System (DNS) name field, and then click **Next**.
5. On the **Trust Type** page, click **Forest Trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.
8. Continue to follow the instructions in the wizard.

Create additional users in the Adatum forest

Create the user Jeff Low with the password **pass@word1**, and assign the company attribute with the value **Adatum**.

To create a user with the Company attribute

1. Open an elevated command prompt in Windows PowerShell, and paste the following code:

```
New-ADUser `
-SamAccountName jlow `
-Name "Jeff Low" `
-UserPrincipalName jlow@adatum.com `
-AccountPassword (ConvertTo-SecureString `
-AsPlainText "pass@word1" -Force) `
-Enabled $true `
-PasswordNeverExpires $true `
-Path 'CN=Users,DC=adatum,DC=com' `
-Company Adatum`
```

Create the Company claim type on adataum.com

To create a claim type by using Windows PowerShell

1. Sign in to adatum.com as an administrator.
2. Open an elevated command prompt in Windows PowerShell, and type the following code:

```
New-ADClaimType `
-AppliesToClasses:@('user') `
-Description:"Company" `
-DisplayName:"Company" `
-ID:"ad://ext/Company:ContosoAdatum" `
-IsSingleValued:$true `
-Server:"adatum.com" `
-SourceAttribute:Company `
-SuggestedValues:@((New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("C
ontoso", "Contoso", "")), (New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("A
datum", "Adatum", ""))) `
```

Enable the Company resource property on contoso.com

To enable the Company resource property on contoso.com

1. Sign in to contoso.com as an administrator.
2. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. In the left pane of Active Directory Administrative Center, click **Tree View**. In the left pane, click **Dynamic Access Control**, and then double-click **Resource Properties**.
4. Select **Company** from the **Resource Properties** list, right-click and select **Properties**. In the **Suggested Values** section, click **Add** to add the suggested values: Contoso and Adatum, and then click **OK** twice.
5. Select **Company** from the **Resource Properties** list, right-click and select **Enable**.

Enable Dynamic Access Control on adatum.com

To enable Dynamic Access Control for adatum.com

1. Sign in to adatum.com as an administrator.
2. Open the Group Policy Management Console, click **adatum.com**, and then double-click **Domain Controllers**.
3. Right-click **Default Domain Controllers Policy**, and select **Edit**.
4. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.

5. Double-click **KDC support for claims, compound authentication, and Kerberos armoring** and select the option next to **Enabled**. You need to enable this setting to use Central Access Policies.
6. Open an elevated command prompt, and run the following command:

```
gpupdate /force
```

Create the Company claim type on contoso.com

To create a claim type by using Windows PowerShell

1. Sign in to contoso.com as an administrator.
2. Open an elevated command prompt in Windows PowerShell, then type the following code:

```
New-ADClaimType -SourceTransformPolicy `
-DisplayName "Company" `
-ID "ad://ext/Company:ContosoAdatum" `
-IsSingleValued $true `
-ValueType "string" `
```

Create the central access rule

To create a central access rule

1. In the left pane of Active Directory Administrative Center, click **Tree View**. In the left pane, click **Dynamic Access Control**, and then click **Central Access Rules**.
2. Right-click **Central Access Rules**, click **New**, and then **Central Access Rule**.
3. In the **Name** field, type **AdatumEmployeeAccessRule**.
4. In the **Permissions** section, select the **Use following permissions as current permissions** option, click **Edit**, and then click **Add**. Click the **Select a principal** link, type **Authenticated Users**, and then click **OK**.
5. In the **Permission Entry for Permissions** dialog box, click **Add a condition**, and enter the following conditions: [**User**] [**Company**] [**Equals**] [**Value**] [**Adatum**]. Permissions should be **Modify, Read and Execute, Read, Write**.
6. Click **OK**.
7. Click **OK** three times to finish and return to Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear

word-wrapped across several lines here because of formatting constraints.

```
New-ADCentralAccessRule `
-
CurrentAcl:"O:SYG:SYD:AR(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)(
XA;;0x1301bf;;;AU;(@USER.ad://ext/Company:ContosoAdatum ==
`"Adatum`"))" `
-Name:"AdatumEmployeeAccessRule" `
-ProposedAcl:$null `
-ProtectedFromAccidentalDeletion:$true `
-Server:"contoso.com" `
```

Create the central access policy

To create a central access policy

1. Sign in to contoso.com as an administrator.
2. Open an elevated command prompt in Windows PowerShell, and then paste the following code:

```
New-ADCentralAccessPolicy "Adatum Only Access Policy"
Add-ADCentralAccessPolicyMember "Adatum Only Access Policy" `
-Member "AdatumEmployeeAccessRule" `
```

Publish the new policy through Group Policy

To apply the central access policy across file servers through Group Policy

1. On the **Start** screen, type **Administrative Tools**, and in the **Search** bar, click **Settings**. In the **Settings** results, click **Administrative Tools**. Open the Group Policy Management Console from the **Administrative Tools** folder.

   Tip
   If the **Show Administrative tools** setting is disabled, the Administrative Tools folder and its contents will not appear in the **Settings** results.
2. Right-click the contoso.com domain, click **Create a GPO in this domain and Link it here…**
3. Type a descriptive name for the GPO, such as **AdatumAccessGPO**, and then click **OK**.

To apply the central access policy to the file server through Group Policy

1. On the **Start** screen, type **Group Policy Management**, in the **Search** box. Open **Group Policy Management** from the Administrative Tools folder.

Tip

If the **Show Administrative tools** setting is disabled, the Administrative Tools folder and its contents will not appear in the Settings results.

2. Navigate to and select Contoso as follows: Group Policy Management\Forest: contoso.com\Domains\contoso.com.
3. Right-click the **AdatumAccessGPO** policy, and select **Edit**.
4. In Group Policy Management Editor, click **Computer Configuration**, expand **Policies**, expand **Windows Settings**, and then click **Security Settings**.
5. Expand **File System**, right-click **Central Access Policy**, and then click **Manage Central access policies**.
6. In the **Central Access Policies Configuration** dialog box, click **Add**, select **Adatum Only Access Policy**, and then click **OK**.
7. Close the Group Policy Management Editor. You have now added the central access policy to Group Policy.

Create the Earnings folder on the file server

Create a new NTFS volume on FILE1, and create the following folder: D:\Earnings.

Note

Central access policies are not enabled by default on the system or boot volume C:.

Set classification and apply the central access policy on the Earnings folder

To assign the central access policy on the file server

1. In Hyper-V Manager, connect to server FILE1. Sign in to the server by using Contoso\Administrator, with the password **pass@word1**.
2. Open an elevated command prompt and type: **gpupdate /force**. This will ensure that your Group Policy changes will take effect on your server.
3. You also need to refresh the Global Resource Properties from Active Directory. Open Windows PowerShell, type `Update-FSRMClassificationpropertyDefinition`, and then press ENTER. Close Windows PowerShell.
4. Open Windows Explorer, and navigate to D:\EARNINGS. Right-click the **Earnings** folder, and click **Properties**.
5. Click the **Classification** tab. Select **Company**, and then select **Adatum** in the **Value** field.
6. Click **Change**, select **Adatum Only Access Policy** from the drop-down menu, and then click **Apply**.
7. Click the **Security** tab, click **Advanced**, and then click the **Central Policy** tab. You

should see the **AdatumEmployeeAccessRule** listed. You can expand the item to view all of the permissions that you set when you created the rule in Active Directory.

8. Click **OK** to return to Windows Explorer.

Hosting-Friendly Web Server Platform (IIS): Scenario Overview

This document contains an overview for the hosting-friendly web server, IIS 8, by presenting a collection of scenarios for configuring IIS 8 to support various types of web applications.

Did you mean…

- Building Your Cloud Infrastructure: Scenario Overview
- Web Server (IIS) Overview

Scenario Description

The IIS 8 web server provides a secure, easy-to-manage, modular and extensible platform for reliably hosting websites, services, and applications. The web server scenarios provide end-to-end instructions on how to plan, install, and configure IIS to host various types of services and applications.

Each scenario is self-contained. You can read them in any order and concentrate on those scenarios that most closely meet your needs.

Web Server Scenarios

Scenario	Description
Build a Static Website on IIS	Shows how to install IIS with the minimum module configuration to server static websites. Describes how to add a website, configure anonymous authentication, set up a default document, set up static content compression, use IIS logging, and filter requests.
Build a Classic ASP Website on IIS	Shows how to install IIS with the modules required to server classic ASP websites. Explains how to add a websites and edit ASP application settings.
Build an ASP.NET Website on IIS	Guides you through planning, installing, and configuring an ASP.NET website on IIS.

Scenario	Description
	Describes what IIS modules to install. Shows how to configure ASP.NET settings, such as session state, pages and controls, application settings, .NET compilation and globalization settings. Explains how to configure data source settings. And shows how to improve security with application isolation, .NET trust levels, .NET authentication, machine key settings, and SSL communication.
Build a PHP Website on IIS	Guides you through planning, installing, and configuring a PHP website on IIS. Describes what IIS modules to install. Describes how to configure essential PHP settings. Explains how to configure WinCache and other PHP extensions. Describes how to improve security of a PHP application on IIS.
Build an FTP Site on IIS	Shows how to install FTP on an existing IIS web server. Describes how to add an FTP site, configure FTP site defaults, fire wall support, user isolation, and directory browsing options. Shows how to configure logon attempt restrictions, request filtering, FTP logging, and FTP messages.
Build a Web Farm with IIS Servers	Guides you through planning, installing, and configuring a simple server farm with IIS servers. Helps you plan and configure the farm infrastructure. Shows how set up load balancing by using Application Request Routing (ARR). Describes how to set up shared content and shared configuration. Shows how to set up a central certificate store for SSL certificates. Explains how to set up FTP and Web Deploy on the server farm.

Practical Applications

Whether you are an IT professional, a web developer, or you want to set up your own web server, this scenario can help you install IIS and configure it to serve your web application.

See Also

The following table contains links to resources related to IIS 8.

Content Type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center | IIS Media Services | What's New in IIS 8.0 for Windows 8?
Troubleshooting	IIS Troubleshooting
Security	Secure Windows Server 2012 R2 and Windows Server 2012 | Security and Protection
Tools and Settings	Web Server (IIS) Administration Cmdlets in Windows PowerShell
Community Resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related Technologies	ASP.NET | ASP.NET Web Projects

Build a Static Website on IIS

This document guides you through the process of installing an IIS web server and configuring it to serve static content. Static content is a web page (HTML) that is delivered to the user exactly as stored. By contrast, dynamic content is generated by a web application, such as an ASP.NET, classic ASP, or PHP application. Static content displays the same information for all users; dynamic content can display user-specific information, such as the user name.

A static-content web server is the most basic configuration of IIS for supporting HTML websites. You can use a static-content web server to host internal or external (public) websites. When you install IIS 8, the default installation provides all the IIS modules required to support a static-content web server. The default installation includes the ability to serve static HTML files, documents, and images. IIS 8 supports default documents, directory browsing, logging, and anonymous authentication for the static content server.

In this document

- Prerequisites
- Step 1: Install the IIS Web Server
- Step 2: Add a Website

- Step 3: Configure Anonymous Authentication
- Step 4: Configure the Default Documents
- Step 5: Configure Static Content Compression
- Next Steps

Prerequisites

To get the most from this tutorial, you must have access to a computer that is running one of the following operating systems:

- Windows Server® 2012
- Windows® 8

Step 1: Install the IIS Web Server

You can use the Web Platform Installer (Web PI) to install IIS and applications that run on IIS. The Web PI installs the latest versions of web platform offerings, including updates. To learn more about the Web PI, see Learn more and install the Web PI. If you use the Web PI to install IIS, you can skip to step 2.

You can also perform this procedure by using the Windows user interface (UI) or from a command line.

To install IIS on Windows Server 2012 by using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select the name of your computer, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, notice the preselected features, and then click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select Role Services** page, note the preselected role services, and then click **Next**.

   Note
   Install the IIS 8 default role services for a static-content web server.
10. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
11. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.

12. To verify that IIS installed successfully, enter the following into a web browser:

    http://localhost

    You see the default IIS Welcome page.

To install IIS on Windows 8 by using the UI

1. On the **Start** page, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services**, and then click **OK**.

   This action installs the IIS 8 default features. Install only the default features for a static-content web server.
4. To verify that IIS installed successfully, enter the following into a web browser:

   http://localhost

   You see the default IIS Welcome page.

To Install IIS by using the command line

- Enter the following command at an elevated command prompt or into a script:

```
Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI
```

Step 2: Add a Website

This step tells you how to add a website to IIS by using the IIS Manager UI, or by running Appcmd.exe commands in a command-line window.

To add a website by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In the **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, right-click the **Sites** node in the tree, and then click **Add**

Website.

3. In the **Add Website** dialog box, enter a friendly name for your website in the **Site name** box.
4. If you want to select a different application pool than the one listed in the **Application Pool** box, click **Select**. In the **Select Application Pool** dialog box, select an application pool from the **Application Pool** list and then click **OK**.
5. In the **Physical path** box, enter the physical path of the website's folder, or click the browse button (**...**) to navigate the file system to find the folder.
6. If the physical path that you entered in the previous step is to a remote share, click **Connect as** to specify credentials that have permission to access the path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option in the **Connect As** dialog box.
7. Select the protocol for the website from the **Type** list.
8. . If you must specify a static IP address for the website (by default, this is set to **All Unassigned**), enter the IP address in the **IP address** box.
9. Enter a port number in the **Port** text box.
10. Optionally, enter a host header name for the website in the **Host Header** box.
11. If you do not have to make any changes to the site, and you want the website to be immediately available, select the **Start Web site immediately** check box.
12. Click **OK**.

To add a website by using the command line

- Use the following syntax at an elevated command prompt or in a script:

  Note
  For this syntax to work, you either must be in the following directory, or have the directory in your path: `%windir%\system32\inetsrv`

  `appcmd add site /name:string /id:`*`uint`* `/physicalPath:`*`string`* `/bindings:`*`string`*

  The variable `name` is the name, and the variable `id` is a positive integer that you want to assign to the site. The variables `name` and `id` are the only variables that are required to add a site by using appcmd. However, if you add a site without specifying the values for the `bindings` and `physicalPath` attributes, the site will not be able to start.

  The variable `physicalPath` is the absolute path of the site content in the file system.

  The variable `bindings` contains information that is used to access the site. It must be in the form *protocol/IP_Address:port:hostheader*. For example, for a website, the binding `http/*:85:` configures the site to listen for HTTP requests on port 85 for all IP addresses and domain names (also known as host headers or host names). On the other hand, a binding of `http/*:85:`**marketing.contoso.com** configures a website to listen for HTTP requests on port 85 for all IP addresses and for the domain name **marketing.contoso.com**.

  To add a website named **contoso** with an ID of 2 that has content in **c:\contoso**, and

that listens for HTTP requests on port 85 for all IP addresses and a domain name of **marketing.contoso.com**, enter the following at the command prompt:

```
appcmd add site /name:contoso /id:2 /physicalPath:c:\contoso
/bindings:http/*:85:marketing.contoso.com
```

Step 3: Configure Anonymous Authentication

Anonymous authentication gives users access to the public areas of your website without prompting them for a user name or password. You can configure anonymous authentication by using the default anonymous user account (IUSR), or you can set up a local user account for anonymous users.

To configure anonymous authentication by using the UI

1. In **Features View** of IIS Manager, double-click **Authentication**.
2. On the **Authentication** page, select **Anonymous Authentication**.
3. In the **Actions** pane, click **Edit** to set the security principal (user credentials) under which anonymous users will connect to the site.
4. In the **Edit Anonymous Authentication Credentials** dialog box, select one of the following options:
   - If you want to configure a specific user account that IIS uses to access your site or application, select **Specific user**. Then click **Set** to open the **Set Credentials** dialog box, and enter a user name and password for the identity. Then click **OK**.
   - If you want IIS processes to run by using the account that is currently specified on the property page for the application pool, select **Application pool identity**. By default, this identity is the IUSR account.

     Important
     If you use the IUSR account, you grant anonymous users all the internal network access associated with that account.
5. Click **OK** to close the **Edit Anonymous Authentication Credentials** dialog box.

To configure anonymous authentication by using the command line

- Use the following syntax to change the default account for anonymous access:

  `appcmd set config /section:anonymousAuthentication /userName:`*`string`* `/password:`*`string`*

  The variable `username` is the account that IIS uses for anonymous authentication, and the variable `password` is the password, which is encrypted in the configuration file by default. For example, to use an account named **Moe** and a password of **pssword1** for anonymous access, enter the following at the command prompt:

  ```
  appcmd set config /section:anonymousAuthentication /userName:Moe
  /password:pssword1
  ```

Step 4: Configure the Default Documents

When a client request to your website doesn't include a document name, IIS looks for a file whose name is defined as a default document. Typically, the default document name is Default.htm. You can define a list of default document names in order of precedence.

To configure the default document by using the UI

1. In **Features View** of IIS Manager, double-click **Default Document**.
2. In the **Actions** pane, click **Add**.
3. In the **Name** box, enter the file name that you want to add to the list of default documents and then click **OK**. This file name is added to the top of the default document list.
4. Optionally, select a default document in the list and in the **Actions** pane, click **Move Up** or **Move Down** to change the file's precedence.
5. Optionally, select a default document in the list, and in the **Actions** pane, click **Remove** to remove any file names that you do not want to use as default documents.

To configure the default document by using the command line

- To add a file name to the list of default documents, use the following syntax:

  ```
  appcmd set config /section:defaultDocument /+files.[value='string']
  ```

  The variable `string` is the file name that you want to add to the list. For example, to add a file named **home.html** to the default document list, enter the following at the command prompt:

  ```
  appcmd set config /section:defaultDocument /+files.[value='home.html']
  ```

  To remove a file named **home.html** from the default document list, enter the following at the command prompt, and then press ENTER:

  ```
  appcmd set config /section:defaultDocument /-files.[value='home.html']
  ```

Step 5: Configure Static Content Compression

You can optionally configure your web server to compress static content to user bandwidth more efficiently and to enhance the performance of your website.

To configure static content compression by using the UI

1. In **Features View** of IIS Manager, double-click **Compression**.
2. Select **Enable static content compression** to configure IIS to compress static content.
3. In the **Static Compression** box, configure the following settings:
   a. Optionally, select **Only compress files larger than (in bytes)** and enter the minimum file size that you want IIS to compress. The default size is 256 bytes.
   b. In the **Cache directory** text box, enter the path of a local directory or click the browse button (**…**) to locate a directory. After a static file is compressed, it is cached in this

temporary directory until it expires, or until the content changes. The temporary directory must be on a local drive on an NTFS-formatted partition. The directory cannot be compressed, and should not be shared.

c. Optionally, select **Per application pool disk space limit (in MB)** and enter the maximum amount of space per application pool, in megabytes, you want IIS to use when it compresses static content. For example, if there are 20 application pools on the server and the **Disk space limit** is set to 100, the maximum disk space will be 2GB. If you click the **Per application pool disk space limit (in MB)** option and enter a number in the text box under it, IIS automatically cleans up the temporary directory according to a least recently used rule when the set limit is reached. The default is 100 MB per application pool.

4. Click **Apply** in the **Actions** pane.

To configure static content compression by using the command line

- To enable HTTP compression of static content, at the command prompt, enter the following command, and then press Enter:

  ```
  appcmd set config /section:urlCompression /doStaticCompression:True
  ```

- To configure static-content compression settings, use the following syntax:

  ```
  appcmd set config /section:urlCompression /minFileSizeforComp:int /directory:string
  /maxDiskSpace:int
  ```

  The variable `minFileSizeforComp` sets the minimum number of bytes that a file must contain for it to be compressed. The default value is 256. The variable `directory` specifies the directory where compressed versions of static files are temporarily stored and cached. The default is the following folder:

  `%SystemDrive%\inetpub\temp\IIS Temporary Compressed Files`

  The variable `maxDiskSpace` sets the maximum amount of space per application pool, in megabytes, that you want IIS to use when it compresses static content. The default is 100 MB per application pool.

Next Steps

Test your website thoroughly to ensure that it functions as expected. Then consider configuring the following features.

- To help you troubleshoot or optimize the performance of your web server, set up IIS logging. For instructions see, Configure Logging in IIS.
- To improve the security of your web server, configure request filtering. For instructions see, Configure Request Filtering in IIS.

See also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build an ASP.NET Website on IIS

- Build a Classic ASP Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Configure Request Filtering in IIS

This document shows you how to use common request-filter settings to improve the security of your IIS 8 web server.

Request filters restrict the types of HTTP requests that IIS 8 processes. By blocking specific HTTP requests, request filters help prevent potentially harmful requests from reaching the server. The request filter module scans incoming requests and rejects requests that are unwanted based upon the rules that you set up.

By default, IIS rejects requests to browse critical code segments. It also rejects requests for some file name extensions.

You can configure a request filter at the server-wide level and then override the configuration at a website level.

- Prerequisites
- General Request Filter Settings
- File Name Extensions
- Filtering Rules
- Hidden Segments
- URL Filtering
- HTTP Verbs
- Header Size Limits
- Query Strings
- Request Filter Logging

Prerequisites

This guide was written for, and tested on, the following operating systems:

- Windows Server® 2012
- Windows® 8

General Request Filter Settings

The general settings include such settings as the following:

1. Whether to allow access to a file with an extension that is not listed for request filter.
2. Whether to allow requests that use HTTP verbs that are not listed.

3. Whether to allow requests that contain high-bit characters (non-ASCII).
4. Whether to allow requests that are double encoded.
5. Maximum length of the content requested.
6. Maximum length of the URL
7. Maximum size of a query string

To configure general request-filter options by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. In the **Actions** pane, click **Edit Feature Settings**.
4. In the **Edit Request Filtering Settings** dialog, edit the settings as desired, and then click **OK**.

To configure general request-filter options by using the command line

- To configure high-bit characters, use the following syntax:

  appcmd set config /section:requestfiltering /allowhighbitcharacters:true | false

  For example, to allow high-bit characters, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /allowhighbitcharacters:true
- To configure double escaping, use the following syntax:

  appcmd set config /section:requestfiltering /allowdoubleescaping:true | false

  For example, to enable double escaping, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /allowdoubleescaping:true
- To configure a value for the maximum allowed length of content, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxallowedcontentlength:*unit*

  The variable **requestlimits.maxallowedcontentlength** *unit* specifies the maximum length of content.

  For example, to specify 30000000 as the maximum length of content, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /requestlimits.maxallowedcontentlength:30000000
- To configure a value for the maximum allowed length of an incoming URL, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxurl:*unit*

  The variable **requestlimits.maxurl** *unit* specifies the maximum length of an incoming URL.

For example, to specify 4096 as the maximum incoming URL length, type the following at the command prompt, and then press ENTER:

appcmd set config /section:requestfiltering /requestlimits.maxurl:4096

- To configure a value for the maximum allowed length of an incoming query string, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxquerystring:*unit*

  The variable **requestlimits.maxquertystring** *unit* specifies the maximum length of an incoming query string.

  For example, to specify 2048 as the maximum incoming query string, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /requestlimits.maxquerystring:2048

- To configure a size limit for a specific HTTP header, use the following syntax:

  appcmd set config /section:requestfiltering /+requestlimits.headerLimits.[header='*string***',sizelimit='***unit***']**

  The variable **header** *string* specifies the header this restriction applies to. The variable **sizelimit** *unit* specifies the maximum size of this header.

  For example, to specify a maximum size of 2048 for headers that include a value of contoso.com, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+requestlimits.headerLimits.[header='contoso.com',sizelimit='2048']

File Name Extensions

For each file name extension you add, you can indicate whether to allow or reject requests for that type of file.

To configure file name extensions by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **File Name Extensions** tab.
4. In the **Actions** pane, click either **Allow File Name Extension** or **Deny File Name Extension**.
5. Type the file name extension in the box, and then click **OK**.

To configure file name extensions by using the command line

- To configure how IIS deals with unlisted file name extensions, use the following syntax:

  appcmd set config /section:requestfiltering /fileExtensions.allowunlisted:true | false

  For example, to deny unlisted file name extensions, type the following at the command

prompt, and then press ENTER:

appcmd set config /section:requestfiltering /fileExtensions.allowunlisted:false

- To configure whether file name extensions apply to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /fileExtensions.applyToWebDAV:true | false

  For example, to configure IIS so that file name extensions do not apply to WebDAV requests, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /fileExtensions.applyToWebDAV:false

- To add a file name extension, use the following syntax:

  appcmd set config /section:requestfiltering /+fileExtensions.[fileextension='*.string***',allowed='true | false']**

  The variable **fileextension** *string* is the file name extension you want to allow or deny.

  For example, to add an allow rule for the file name extension .xxx, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+fileExtensions.[fileextension='*.xxx***',allowed='true']**

- To remove a rule for the file name extension .xxx, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /-fileExtensions.[fileextension='*.xxx***']**

Filtering Rules

IIS 8 permits you to define custom filter rules that apply to incoming requests. Using this feature, you can define filters that can do the following:

1. Scan the request URL.
2. Scan for query strings contained in the URL.
3. Scan for specific header fields.
4. Define what file name extensions the filter applies to.
5. Define strings you want to deny.

To configure a filtering rule by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Rules** tab.
4. In the **Actions** pane, double-click **Add Filtering Rule**.
5. In the **Name** box, type a name for the filtering rule.
6. If you want the URL scanned, select the **Scan Url** check box.
7. If you want the query string scanned, select the **Scan query string** check box.

8. Under **Scan Headers**, type one or more headers to scan for.
9. Under **Applies To**, type one or more file name extensions that the rule applies to.
10. Under the **Deny Strings**, type one or more strings you want to deny.
11. Click **OK**.

Hidden Segments

This feature allows you to reject requests that contain a URL segment (for example, a folder name).

To configure hidden segments by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Hidden Segments** tab.
4. In the **Actions** pane, click **Add Hidden Segment**.
5. Type the URL segment in the box, and then click **OK**.

To configure hidden segments by using the command line

- To configure whether hidden segments apply to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /hiddensegments.applyToWebDAV:true | false

  For example, to configure IIS so that hidden segments do not apply to WebDAV requests, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /hiddensegments.applyToWebDAV:false
- To configure a hidden segment, use the following syntax:

  appcmd set config /section:requestfiltering /+hiddensegments.[segment='*string***']**

  The variable **segment** *string* specifies a URL segment that is hidden.

  For example, to specify that /bin is a hidden segment, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+hiddensegments.[segment='/bin']

URL Filtering

You can configure IIS to accept a specified URL. In addition, you can configure it to deny a specified URL sequence.

To configure URL filtering by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **URL** tab.
4. In the **Actions** pane, select either **Allow URL** or **Deny Sequence**.
5. Type the URL or the URL sequence in the box, and click **OK**.

To configure URL filtering by using the command line

- To deny a URL sequence, use the following syntax:

  appcmd set config /section:requestfiltering /+denyurlsequences.[sequence='*string***']**

  The variable **sequence** *string* specifies a sequence of characters in a URL that IIS is never allowed to parse.

  For example, to specify that IIS never parse URLs that contain two periods, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+denyurlsequences.[sequence='..']

HTTP Verbs

You can define a list of verbs that IIS 8 accepts as part of a request. When IIS rejects a request based on this feature, the error code logged is 404.6.

To configure unlisted HTTP verbs by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **HTTP Verbs** tab.
4. In the **Actions** pane, click either **Allow Verb** or **Deny Verb**.
5. Enter the verb in the box, and then click **OK**.

To configure unlisted HTTP verbs by using the command line

- To configure how IIS deals with unlisted verbs, use the following syntax:

  appcmd set config /section:requestfiltering /verbs.allowunlisted:true | false

  For example, to deny unlisted verbs, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /verbs.allowunlisted:false
- To configure whether verb filtering applies to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /verbs.applyToWebDAV:true | false

  For example, to configure IIS so that verb filtering does not apply to WebDAV requests, type the following at the command prompt and then press ENTER:

appcmd set config /section:requestfiltering /verbs.applyToWebDAV:false

- To configure a verb to filter, use the following syntax:

  appcmd set config /section:requestfiltering /+verbs.[verb='*string***',allowed='true | false']**

  The variable **verb** *string* specifies the verb this restriction applies to.

  For example, to specify GET is allowed, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+verbs.[verb='GET',allowed='true']

Header Size Limits

You can limit the size of HTTP request headers to improve performance and security. Headers are name/value pairs that define the operating parameters of an HTTP transaction.

To configure header size limits by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Headers** tab, and click **Add Header**.
4. In the **Header** box, type the header field name.
5. In the **Size Limit** box, type a positive integer that represents the header size limit in bytes.
6. Click **OK**.

Query Strings

You can configure IIS 8 to allow or deny specific query strings contained in the requested URL. For example, if a denied query string is found in a request URL, the request is denied.

To configure query strings by using the UI

1. Open IIS Manager and select the level you want to configure request filter for.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Query Strings** tab, and click either **Allow Query String** or **Deny Query String**.
4. In the **Query string** box, type the query string.
5. Click **OK**.

Request Filter Logging

You can use IIS logging to evaluate and optimize your request filter configuration.

The following table shows the request filter error codes that you see in the log:

Error Description	Status Code
Request Filtering: URL Sequence denied	404.5
Request Filtering: Verb denied	404.6
Request Filtering: File name extension denied	404.7
Request Filtering: Denied by hidden segment	404.8
Request Filtering: Denied because request header is too long.	404.10
Request Filtering: Denied because URL doubled escaping	404.11
Request Filtering: Denied because of high bit characters	404.12
Request Filtering: Denied because content length too large	404.13
Request Filtering: Denied because URL too long	404.14
Request Filtering: Denied because query string too long	404.15

See Also

- Build an ASP.NET Website on IIS
- Build a Classic ASP Website on IIS
- Build a Static Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS ServersBuild a web farm using IIS servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Configure Logging in IIS

You can configure logging on your web server or website that records information about HTTP requests and errors. The information in your log can help you troubleshoot or optimize your website.

Choose one or more of the following procedures to set up logging appropriately for your needs:

- Prerequisites
- Configure Logging at the Site Level
- Configure Per-site Logging at the Server Level
- Configure Per-server Logging at the Server Level
- Select W3C Fields to Log
- Configure Log File Rollover Options

Prerequisites

This guide was written for, and tested on, the following operating systems:

- Windows Server® 2012 R2
- Windows Server® 2012
- Windows® 8

Configure Logging at the Site Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure logging at the site level by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** tree view, select your website.
3. In **Features View**, double-click **Logging**.
4. On the **Logging** page, in the **Log file** section under **Format**, select one of the following log file formats:
   - **IIS**: to use the Microsoft IIS log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by commas, and time is recorded as local time. For more information about the IIS log file format, see IIS Log File Format (IIS 6.0).
   - **NCSA**: to use the National Center for Supercomputing Applications (NCSA) Common log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by spaces, and time is recorded as local

time with the Coordinated Universal Time (UTC) offset. For more information about the NCSA log file format, see NCSA Common Log File Format (IIS 6.0).

- **W3C**: to use the centralized W3C log file format to log information about all sites on the server. This format is handled by HTTP.sys, and is a *customizable* ASCII text-based format, which means that you specify the fields that are logged. Specify the fields that are logged on the **W3C Logging Fields** dialog box by clicking **Select Fields** on the **Logging** page. Fields are separated by spaces, and time is recorded in Coordinated Universal Time (UTC). For more information about the W3C log file format, see W3C Extended Log File Format (IIS 6.0).

  Note
  In Windows Server 2012 R2 and later, you can write log information to a log file, as Event Tracing for Windows (ETW) events, or both. For more information on ETW, see Event Tracing.

- **Custom**: to use a custom format for a custom logging module. When you select this option, the **Logging** page becomes disabled, because custom logging cannot be configured in IIS Manager. For more information about how to use custom log file formats, see Custom Logging Modules (IIS 6.0).

5. Under **Directory**, specify the path where the log file should be stored. The default is *%SystemDrive%*\inetpub\logs\LogFiles

   Note
   As a best practice, store log files, such as failed request trace logs, in a directory other than *systemroot*.

6. In the **Log File Rollover** section, select one of the following options:
   - **Schedule**: to create new log file that is based on one of the following values:
   - **Hourly**: a new log file is created each hour.
   - **Daily**: a new log file is created each day.
   - **Weekly**: a new log file is created each week.
   - **Monthly**: a new log file is created each month.
   - **Maximum file size (in bytes)**: to create a log file when the file reaches a certain size (in bytes).The minimum file size is 1048576 bytes. If this attribute is set to a value less than 1048576 bytes, the default value is implicitly assumed as 1048576 bytes.
   - **Do not create a new log file**: there is a single log file that continues to grow as information is logged.
7. Select **Use local time for file naming and rollover** to specify that log file naming and time for log file rollover uses the local server time. When this option is not selected, Coordinated Universal Time (UTC) is used.

   Note
   Regardless of this setting, timestamps in the actual log file will use the time format for the log format that you select from the **Format** list. For example, NCSA and W3C log file formats use UTC time format for timestamps.

8. Click **Apply** in the **Actions** pane.

Configure Per-site Logging at the Server Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure per-site logging at the server level by using the UI

1. In the **Connections** tree view of IIS Manager, select your web server.
2. In **Features View**, double-click **Logging**.
3. On the **Logging** page under **One log file per site**, select **Site** from the drop-down list. By default, **Site** is selected.
4. Complete this procedure by following the site-level procedure by starting with step 4.

Configure Per-server Logging at the Server Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure per-server logging at the server level by using the UI

1. In the **Connections** tree view of IIS Manager, select your web server.
2. In **Features View**, double-click **Logging**.
3. On the **Logging** page, under **One log file per site**, select **Server** from the drop-down list. By default, **Site** is selected.
4. Complete this procedure by following the site-level procedure by starting with step 4.

Select W3C Fields to Log

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To select W3C fields to log by using the UI

1. In **Features View** of IIS Manager, double-click **Logging**.
2. On the **Logging** page, in the **Log file** section under **Format**, click **Select Fields**.
3. In the **W3C Logging Fields** dialog box, select one or more of the following options:
   - **Date (date)**: the date on which the request occurred.
   - **Time (time)**: the time, in Coordinated Universal Time (UTC), at which the request occurred.
   - **Client IP Address (c-ip)**: the IP address of the client that made the request.
   - **User Name (cs-username)**: the name of the authenticated user who accessed your

server. Anonymous users are indicated by a hyphen.

- **Service Name (s-sitename)**: the site instance number that fulfilled the request.
- **Server Name (s-computername)**: the name of the server on which the log file entry was generated.
- **Server IP Address (s-ip)**: the IP address of the server on which the log file entry was generated.
- **Server Port (s-port)**: the server port number that is configured for the service.
- **Method (cs-method)**: the requested action, for example, a GET method.
- **URI Stem (cs-uri-stem)**: the Universal Resource Identifier, or target, of the action.
- **URI Query (cs-uri-query)**: the query, if any, that the client was trying to perform. A Universal Resource Identifier (URI) query is necessary only for dynamic pages.
- **Protocol Status (sc-status)**: the HTTP or FTP status code.
- **Protocol Sub-status (sc-substatus)**: the HTTP or FTP substatus code.
- **Win32 Status (sc-win32-status)**: the Windows status code.
- **Bytes Sent (sc-bytes)**: the number of bytes that the server sent.
- **Bytes Received (cs-bytes)**: the number of bytes that the server received.
- **Time Taken (time-taken)**: the length of time that the action took in milliseconds.
- **Protocol Version (cs-version)**: the protocol version, HTTP or FTP, that the client used.
- **Host (cs-host)**: the host name, if any.
- **User Agent (cs(UserAgent))**: the browser type that the client used.
- **Cookie (cs(Cookie))**: the content of the cookie sent or received, if any.
- **Referer (cs(Referer))**: the site that the user last visited. This site provided a link to the current site.

Note
In Windows Server 2012 R2 and later, you can specify additional custom fields to log from the HTTP request and response headers, and from server variables. To add custom fields, select **Add Field** in the **W3C Logging Fields** dialog box. Enhanced logging is available only for site-level logging - if you selected server-level logging, then **Add Field** is disabled. Note that if the total size of the custom fields that you define is greater than 64K bytes, the logged content is truncated to 64K bytes.

4. Click **Apply** in the **Actions** pane.

Configure Log File Rollover Options

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure log file rollover options by using the UI

1. In **Features View** of IIS Manager, double-click **Logging**.
2. On the **Logging** page, in the **Log File Rollover** section, select one of the following options:
   - **Schedule**: to create new log file that is based on one of the following values:
   - **Hourly**: a new log file is created each hour.
   - **Daily**: a new log file is created each day.
   - **Weekly**: a new log file is created each week.
   - **Monthly**: a new log file is created each month.
   - **Maximum file size (in bytes)**: to create a log file when the file reaches a certain size (in bytes).The minimum file size is 1048576 bytes. If this attribute is set to a value less than 1048576 bytes, the default value is implicitly assumed as 1048576 bytes.
   - **Do not create a new log file**: This option means that there is a single log file that continues to grow as information is logged. If you use a single log file for your site, it is helpful when you use log parsing utilities, but it also creates larger log files that could affect the overall performance of the server.
3. Select **Use local time for file naming and rollover** to specify that log file naming and time for log file rollover uses the local server time. When this option is not selected, Coordinated Universal Time (UTC) is used.

   Note
   Regardless of this setting, timestamps in the actual log file will use the time format for the log format that you select from the **Format** list. For example, NCSA and W3C log file formats use UTC time format for timestamps.
4. Click **Apply** in the **Actions** pane.

See Also

- Build an ASP.NET Website on IIS
- Build a Classic ASP Website on IIS
- Build a Static Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Build a Classic ASP Website on IIS

This document guides you through installing IIS and configuring a classic ASP website. Classic ASP is a server-side scripting environment that you can use to create and run dynamic web applications. With ASP, you can combine HTML pages, script commands, and COM components

to create interactive web pages that are easy to develop and modify. Classic ASP is the predecessor to ASP.NET, but it is still in wide use today.

The Classic ASP server configuration adds IIS modules for ASP and ISAPI extensions to the default IIS installation.

In this document

- Prerequisites
- Step 1: Install the IIS Web Server
- Step 2: Add a Classic ASP Website
- Step 3: Edit ASP Application Settings
- Next Steps

Prerequisites

This guide was written for, and tested on, the following operating systems:

1. Windows Server® 2012
2. Windows® 8

Step 1: Install the IIS Web Server

You can use the Web Platform Installer (Web PI) to install IIS, and applications that run on IIS. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

You can also perform this procedure by using the Windows user interface (UI) or from a command line.

To install IIS on Windows Server 2012 by using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select your server, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, note the preselected features that are installed by default, and then select the following additional role services:
   - ASP
   - ISAPI Extensions
8. Click **Next**.

9. On the **Web Server Role (IIS)** page, click **Next**.
10. On the **Select Role Services** page, note the preselected role services that are installed by default, and then click **Next**.

    Note
    You only have to install the IIS 8 default role services for a static-content web server.
11. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
12. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
13. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    You should see the default IIS Welcome page.

To install IIS on Windows 8 by using the UI

1. On the **Start** page, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services**, note the preselected features that are installed by default, and then select the following additional role services:
   - ASP
   - ISAPI Extensions
4. Click **OK**.
5. To verify that IIS installed successfully, type the following into a web browser:

   http://localhost

   You see the default IIS Welcome page.

To Install IIS by using the command line

- Type the following command at a command prompt or into a script:

  Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-ApplicationDevelopment;IIS-ASP;IIS-ISAPIExtensions;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI

Step 2: Add a Classic ASP Website

You can perform this procedure by using the user interface (UI), by running Appcmd.exe commands in a command-line window, by editing configuration files directly, or by writing WMI scripts.

To add a website by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. On the **Server Manager Dashboard**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, right-click the **Sites** node in the tree, and then click **Add Website**.
3. In the **Add Website** dialog box, type a friendly name for your website in the **Site name** box.
4. If you want to select a different application pool than the one listed in the **Application Pool** box, click **Select**. In the **Select Application Pool** dialog box, select an application pool from the **Application Pool** list and then click **OK**.
5. In the **Physical path** box, type the physical path of the Web site's folder, or click the browse button (**...**) to navigate the file system to find the folder.
6. If the physical path that you entered in step 5 is to a remote share, click **Connect as** to specify credentials that have permission to access the path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option in the **Connect As** dialog box.
7. Select the protocol for the Web site from the **Type** list.
8. The default value in the **IP address** box is **All Unassigned**. If you must specify a static IP address for the Web site, type the IP address in the **IP address** box.
9. Type a port number in the **Port** text box.
10. Optionally, type a host header name for the Web site in the **Host Header** box.
11. If you do not have to make any changes to the site, and you want the Web site to be immediately available, select the **Start Web site immediately** check box.
12. Click **OK**.

To add a website by using the command line

- Use the following syntax at the command prompt or in a script:

  Note
  For this syntax to work, you either must be in the following directory, or have the

directory in your path: `%windir%\system32\inetsrv`

appcmd add site /name:*string* **/id:***uint* **/physicalPath:***string* **/bindings:***string*

The variable **name** *string* is the name, and the variable **id** *uint* is the unsigned integer that you want to assign to the site. The variables **name** *string* and **id** *uint* are the only variables that are required when you add a site in Appcmd.exe.

Note

When you add a site without specifying the values for the **bindings** and **physicalPath** attributes, the site will not be able to start.

The variable **physicalPath** *string* is the path of the site content in the file system.

The variable **bindings** *string* contains information that is used to access the site, and it should be in the form of *protocol*/*IP_address*:*port*:*host_header*. For example, a web site binding is the combination of protocol, IP address, port, and host header. A binding of **http/*:85:** enables a web site to listen for HTTP requests on port 85 for all IP addresses and domain names (also known as host headers or host names). On the other hand, a binding of **http/*:85:marketing.contoso.com** enables a web site to listen for HTTP requests on port 85 for all IP addresses and the domain name **marketing.contoso.com**.

To add a web site named **contoso** with an ID of 2 that has content in **c:\contoso**, and that listens for HTTP requests on port 85 for all IP addresses and a domain name of **marketing.contoso.com**, type the following at the command prompt, and then press ENTER:

appcmd add site /name:contoso /id:2 /physicalPath:c:\contoso /bindings:http/*:85:marketing.contoso.com

Step 3: Edit ASP Application Settings

IIS 8 provides default settings for ASP applications, but you can change those settings as needed. For example, you can enable client-side debugging on a test server to aide in troubleshooting issues during a test pass.

To edit ASP application settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **ASP**.
3. On the **ASP** page, edit settings as desired.
4. When finished, click **Apply** in the **Actions** pane.

To edit ASP application settings by using the command line

- **Specify default character set**

  To specify the default character set for an application, use the following syntax:

  appcmd set config /section:asp /codePage:*integerRange*

The variable *integerRange* is the default character set. For example, to set the code page to a Latin character set used in American English and many European alphabets, type the following at the command prompt, and then press Enter:

appcmd set config /section:asp /codePage:*1252*

- **Enable or disable buffering**

  To enable or disable buffering of ASP application output, use the following syntax:

  appcmd set config /section:asp /bufferingOn:True|False

  True enables buffering whereas **False** disables buffering. The default value is **True**.

- **Enable or disable HTTP 1.1 chunked transfer encoding**

  To enable HTTP 1.1 chunked transfer encoding for the World Wide Web publishing service, use the following syntax:

  appcmd set config /section:asp /enableChunkedEncoding:True|False

  True enables HTTP 1.1 chunked transfer encoding whereas **False** disables HTTP 1.1 chunked transfer encoding. The default value is **True**.

- **Enable or disable HTML fallback**

  To enable or disable HTML fallback, use the following syntax:

  appcmd set config /section:asp /enableASPHTMLFallback:True|False

  True causes an .htm file that has the same name as the requested .asp file, if it exists, to be sent instead of the .asp file if the request is rejected because of a full request queue. The default value is **True**.

- **Enable or disable parent paths**

  To enable or disable paths relative to the current directory or above the current directory, use the following syntax:

  appcmd set config /section:asp /enableParentPaths:True|False

  True sets ASP pages to allow paths relative to the current directory or above the current directory. The default value is **True**.

- **Set client connection test interval**

  To set a time interval after which ASP will check to see whether the client is still connected before executing a request, use the following syntax:

  appcmd set config /section:asp /queueConnectionTestTime:*timeSpan*

  The variable *timeSpan* sets the time interval (hh:mm:ss) after which ASP checks to see whether the client is still connected before executing a request. The default value is 00:00:03.

- **Set maximum requesting entity body limit**

  To specify the maximum number of bytes allowed in the entity body of an ASP request, use the following syntax:

  appcmd set config /section:asp /maxRequestEntityAllowed:*int*

  The variable *int* represents the maximum number of bytes allowed in the body of an ASP

request. The default value is 200000 bytes.

- **Set request queue length**

  To specify the maximum number of concurrent ASP requests allowed into the queue, use the following syntax:

  appcmd set config /section:asp /requestQueueMax:*int*

  The variable *int* represents the maximum number of concurrent ASP requests that are allowed into the request queue. The default value is 3000.

- **Set request queue time-out**

  To specify the period that an ASP request can wait in the request queue, use the following syntax:

  appcmd set config /section:asp /queueTimeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that an ASP request can wait in the request queue. The default value is 00:00:00.

- **Specify response buffering limit**

  To control the maximum number of bytes that an ASP page can write to the response buffer before a flush occurs, use the following syntax:

  appcmd set config /section:asp /bufferingLimit:*int*

  The variable *int* represents the maximum size, in bytes, of the ASP buffer. The default value is 4194304 bytes.

- **Set script time-out**

  To specify the default length of time that ASP pages let a script run before terminating the script and writing an event to the Windows Event Log, use the following syntax:

  appcmd set config /section:asp /scriptTimeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that an ASP request can run before an event is written to the Windows Event Log. The default value is 00:01:30.

- **Specify threads per processor limit**

  To specify the maximum number of worker threads per processor that ASP can create, use the following syntax:

  appcmd set config /section:asp /processorThreadMax:*int*

  The variable *int* represents the maximum number of worker threads per processor that ASP can create. The default value is 25.

- **Specify default locale identifier**

  To define how dates, times, and currencies are formatted for an ASP application, use the following syntax:

  appcmd set config /section:asp /lcid:*int*

  The variable *int* represents the default locale identifier for an ASP application. The default value is 0.

- **Enable or disable automatic application restart**

  To enable or disable automatic restart of ASP applications whenever a configuration setting is changed, use the following syntax:

  appcmd set config /section:asp /enableApplicationRestart:True|False

  True enables ASP applications to be automatically restarted whenever a configuration setting is changed. The default value is **True**.
- **Enable or disable line number calculation**

  To enable or disable ASP to calculate and store the line number of each executed line of code to provide the number in an error report, use the following syntax:

  appcmd set config /section:asp /calLineNumber:True|False

  True enables line number calculation and storage. The default value is **True**.
- **Enable or disable COM component exception trapping**

  To enable or disable ASP pages to catch exceptions thrown by COM components, use the following syntax:

  appcmd set config /section:asp /exceptionCatchEnable:True|False

  True enables COM component exception trapping. If set to **False**, the Microsoft Script Debugger tool does not catch exceptions sent by the component that you are debugging. The default value is **True**.
- **Enable or disable client-side debugging**

  To enable or disable client-side debugging, use the following syntax:

  appcmd set config /section:asp /appAllowClientDebug:True|False

  True enables client-side debugging. The default value is **False**.
- **Enable or disable log error requests**

  To enable or disable the writing of ASP errors to the application section of the Windows event log, use the following syntax:

  appcmd set config /section:asp /logErrorRequests:True|False

  True enables log error requests. By default, ASP errors are written to the client browser and the IIS logs. The default value is **True**.
- **Enable or disable server-side debugging**

- **Enable or disable Windows event logging of ASP errors**

  To enable or disable ASP debugging on the server, use the following syntax:

  appcmd set config /section:asp /appAllowDebugging:True|False

  True enables server-side debugging for ASP applications. The default value is **False**.
- **Run On End Functions Anonymously**

  To enable or disable **SessionOnEnd** and **ApplicationOnEnd** global ASP functions to run as the anonymous user, use the following syntax:

  appcmd set config /section:asp /runOnEndAnonymously:True|False

True enables **SessionOnEnd** and **ApplicationOnEnd** global ASP functions to run as the anonymous user. The default value is **True**.

- **Specify script error message**

  To specify the error message to send to the browser if specific debugging errors are not sent to the client, use the following syntax:

  appcmd set config /section:asp /scriptErrorMessage:*string*

  The variable *string* represents the error message that is sent to the browser when specific debugging errors are not sent to the client. The default value is "An error occurred on the server when processing the URL. Please contact the system administrator".

- **Enable or disable sending errors to the browser**

  To enable or disable the writing of debugging specifics (file name, error, line number, and description) to the client browser in addition to logging them to the Windows event log, use the following syntax:

  appcmd set config /section:asp /scriptErrorSentToBrowser:True|False

  True enables the writing of debugging specifics to the client browser. The default value is **False**.

- **Specify default script language**

  To specify the default script language for all ASP applications that are running on the web server, use the following syntax:

  appcmd set config /section:asp /scriptLanguage:*string*

  The variable *string* represents the default script language. The default value is VBScript.

- **Specify cache directory path**

  To specify the name of the directory where ASP stores compiled ASP templates when the in-memory cache overflows, use the following syntax:

  appcmd set config /section:asp /diskTemplateCacheDirectory:*string*

  The variable *string* represents the cache directory path. The default value is *%windir%*\system32\inetsrv\ASP Compiled Templates.

- **Enable or disable type library caching**

  To enable or disable the caching of type libraries, use the following syntax:

  appcmd set config /section:asp /enableTypelibCache:True|False

  True enables the caching of type libraries. The default value is **True**.

- **Set maximum number of compiled ASP templates to store**

  To set the maximum number of compiled ASP templates that can be stored, use the following syntax:

  appcmd set config /section:asp /maxDiskTemplateCacheFiles:*int*

  The variable *int* represents the maximum number of compiled ASP templates to store. The default value is 2000.

- **Set maximum number of compiled ASP templates to store**

  To set the maximum number of precompiled script files to cache, use the following syntax:

  appcmd set config /section:asp /scriptFileCacheSize:*int*

  The variable *int* represents the number of precompiled script files to cache. If set to 0, no script files are cached. If set to 4294967295, all requested script files are cached. The default value is 500.

- **Set maximum number of scripting engines to cache**

  To set the maximum number of scripting engines that ASP pages keep cached in memory, use the following syntax:

  appcmd set config /section:asp /scriptEngineCacheMax:*int*

  The variable *int* represents the maximum number of scripting engines that are cached. The default value is 250.

- **Enable or disable COM+ side-by-side assemblies**

  To enable or disable side-by-side COM+ assemblies, which allow ASP applications to specify which version of a system DLL or classic COM component to use, use the following syntax:

  appcmd set config /section:asp /appServiceFlags:True|False

  True enables COM+ side-by-side assemblies. The default value is **False**.

- **Enable or disable COM+ tracker**

  To enable or disable COM+ tracker, use the following syntax:

  appcmd set config /section:asp /enableTypelibCache:True|False

  True enables COM+ tracker, which lets administrators or developers debug ASP applications. The default value is **False**.

- **Enable or disable multithreaded environments**

  To enable or disable ASP to run in a multithreaded environment, use the following syntax:

  appcmd set config /section:asp /executeInMta:True|False

  True enables ASP to run in a multithreaded environment. The default value is **False**.

- **Enable or disable thread model checking**

  To enable or disable whether IIS checks the threading model of any component that your application creates, use the following syntax:

  appcmd set config /section:asp /trackThreadingModel:True|False

  True enables thread model checking. The default value is **False**.

- **Specify COM+ partition ID**

  To specify the Globally Unique Identifier (GUID) of the COM+ partition, use the following syntax:

  appcmd set config /section:asp /partitionID:*string*

The variable *string* represents the GUID of the COM+ partition. The default value is 00000000-0000-0000-0000-000000000000.

Note
You must also set the **appServiceFlags** flag to **True**.

- **Specify COM+ application**

  To specify the name of the COM+ application, use the following syntax:

  appcmd set config /section:asp /sxsName:*string*

  The variable *string* represents name of the COM+ application.

  Note
  You must also set the **appSeviceFlags** flag to **True**.

- **Enable or disable COM+ partitioning**

  To enable or disable COM+ partitioning, use the following syntax:

  appcmd set config /section:asp /appServiceFlags:True|False

  True enables COM+ partitioning, which can be used to isolate applications in their own COM+ partition. The default value is **False**.

  Note
  If set to **True**, you must also set a value for the **partitionID** property.

- **Enable or disable session state**

  To enable or disable session state persistence for an ASP application, use the following syntax:

  appcmd set config /section:asp /allowSessionState:True|False

  True enables session state persistence. The default value is **True**.

- **Set maximum number of concurrent sessions**

  To set the maximum number of concurrent sessions that ASP allows, use the following syntax:

  appcmd set config /section:asp /max:*int*

  The variable *int* represents the maximum number of concurrent sessions. The default value is -1.

- **Enable or disable secure session ID**

  To enable or disable the sending of a session ID as a secure cookie if assigned over a secure session channel, use the following syntax:

  appcmd set config /section:asp /keepSessionIdSecure:True|False

  True enables secure session ID. The default value is **True**.

- **Set session time-out**

  To specify the default time that a session object is maintained after the last request associated with the object is made, use the following syntax:

appcmd set config /section:asp /timeout:*timeSpan*

The variable *timeSpan* represents the maximum time (hh:mm:ss) that a session object is maintained after the last request associated with the object is made. The default value is 00:20:00.

Next Steps

Test your website thoroughly to ensure that it functions as expected. Then consider configuring the following features.

- To help you troubleshoot or optimize the performance of your web server, set up IIS logging. For instructions see, Configure Logging in IIS.
- To improve the security of your web server, configure request filtering. For instructions see, Configure Request Filtering in IIS.

See also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build a Static Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Build an ASP.NET Website on IIS

This document contains an overview of the **Build an ASP.NET Website on IIS** scenario. It also contains links to additional information and community resources related to the scenario.

Did you mean…

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an FTP Site on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Scenario Description

This scenario shows how to plan and configure an ASP.NET website on an IIS 8 webserver. It is divided into two phases: a plan and design phase, and an install and configure phase. In the plan and design phase, you are provided the information needed to make informed decisions about

web server installation, ASP.NET settings, data source settings, and basic application security. In the install and configure phase, you are guided through the procedures required to install IIS, add an ASP.NET application, and configure IIS.

This scenario does not cover how to write an ASP.NET application.

In This Scenario

- Plan an ASP.NET Website on IIS
  - Step 1: Plan IIS Web Server and ASP.NET Modules Installation
  - Step 2: Plan ASP.NET Settings
  - Step 3: Plan Data Source Settings
  - Step 4: Plan Application Security
- Configure an ASP.NET Website on IIS
  - Step 1: Install IIS and ASP.NET Modules
  - Step 2: Configure ASP.NET Settings
  - Step 3: Configure Data Source Settings
  - Step 4: Configure Application Security

Practical Applications

Whether you are an IT professional, a web developer, or you just want to set up your own webserver, this scenario can help you install IIS and configure it to serve your ASP.NET web application.

Software Requirements

To get the most from this scenario, you must have access to a computer running one of the following operating systems:

- Windows Server® 2012
- Windows® 8

See Also

The following table contains links to resources related to this scenario.

Content type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center
Tools and Settings	Web Server (IIS) Administration Cmdlets in

Content type	References
	Windows PowerShell
Community Resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related Technologies	ASP.NET | ASP.NET Web Projects

Plan an ASP.NET Website on IIS

To develop a plan for installing an IIS web server and configuring it for ASP.NET web applications, follow the steps listed.

- Step 1: Plan IIS Web Server and ASP.NET Modules Installation
- Step 2: Plan ASP.NET Settings
- Step 3: Plan Data Source Settings
- Step 4: Plan Application Security

After you have reviewed these planning steps, see Configure an ASP.NET Website on IIS. For more information, see Build an ASP.NET Website on IIS.

Step 1: Plan IIS Web Server and ASP.NET Modules Installation

The first step in planning to build an ASP.NET website on IIS 8 is to install IIS along with the ASP.NET modules and to add your application files to IIS.

The following list shows the tasks required to complete this step:

1.1. Plan to Install IIS and ASP.NET Modules

1.2. Plan to Add the ASP.NET Application

When you are done with these tasks, record your design decisions before going on to Step 2: Plan ASP.NET Settings.

1.1. Plan to Install IIS and ASP.NET Modules

An ASP.NET Web server is an extended IIS server that has the ASP.NET runtime extensibility model integrated into the core server. The ASP.NET server configuration adds the following IIS modules to the default IIS web server installation:

- ASP.NET 4.5
- .NET Extensibility 4.5
- ISAPI Extensions
- ISAPI Filters

Important

To install IIS 8 on either Windows Server® 2012 or Windows® 8, you must sign on as a member of the Administrators group.

By default, Windows Server 2012 and Windows 8 installs .NET 4.5 only. If you want to run .NET 2 applications, install .NET 3.5, which supports web applications written using ASP.NET 2 to 3.5.

1.2. Plan to Add the ASP.NET Application

An application is a grouping of content at the root level of a website or a grouping of content in a separate folder under the website root directory. When you add an application in IIS 8, you designate a directory as the *application root*, or starting point, for the application. Then specify properties specific to that particular application, such as the application pool that the application runs in.

To configure your ASP.NET application on an IIS website, provide the following information:

- **Alias**: The alias is used as part of the application root URL and should be short and descriptive. For example, the alias *marketing* added to Default Web Site on the local host computer, would produce the following URL: `//localhost/marketing`.
- **Application pool**: An application pool enables an application or a group of applications to run in isolation from one or more applications in another application pool.
- **Physical path**: The local path to the application files on the server.

See Also

- Step 2: Plan ASP.NET Settings
- Plan an ASP.NET Website on IIS
- Build an ASP.NET Website on IIS
- Step 1: Install IIS and ASP.NET Modules

Step 2: Plan ASP.NET Settings

In this phase of building your website, consider the following IIS server and website settings that support ASP.NET:

2.1. Session State Settings

2.2. Pages and Controls Settings

2.3. Application Settings

2.4. .NET Compilation Settings

2.5. .NET Globalization Settings

2.1. Session State Settings

When clients visit a site, they generally navigate from one page to another and frequently change some of the pages they visit. If you want to track pages that users browse and the changes they make during a visit to your website, configure session state.

When session state is enabled for your application, a user receives a unique session ID on their first request to a web page from your ASP.NET application. Session-state data is stored on the server in one of the following ways:

- **In-process**: Session state is stored in the worker process where the ASP.NET application runs.
- **State Server**: Session state is stored outside the worker process where the ASP.NET application runs.
- **SQL Server**: Session state is stored in a SQL Server database.

Note

You can also configure custom out-of-state storage for session-state data. However, it is beyond the scope of this tutorial. Visual Studio 11 project uses custom storage to support SQL Server, SQL Compact, and SQL Azure.

Session-state data can also be stored on the client in a cookie. A cookie is a text file that contains data used for storing information about a user, such as user preferences and user authentication information.

The following sections describe your session-state storage options in more detail.

- Store session state in process
- Store session state by using state server
- Store session state by using SQL server
- Cookie mode for session state

Store session state in process

The in-process session state mode stores session-state data for an ASP.NET application in the worker process where the application runs. This mode is the default for IIS 8.

The advantage of using in-process session state is that it provides the fastest access to session-state data. However, as you store more data in a session, you consume more memory, which can slow server performance.

Before you configure in-process session state, consider the effect of worker process recycling on session state data. If the worker process recycles, all session state data is lost. If your ASP.NET applications must preserve session-state data, and if speed of access to the data is not your primary objective, consider using an out-of-process session state mode for storing this data.

Important
The Windows state service (Aspnet_state.exe) must be running for in-process session state to take effect. By default, this service is installed when Windows Server® 2012 is installed and is configured for manual start. It is recommended that you change the start behavior to Automatic.

By default, the session expires when the user has not requested or refreshed a page in the ASP.NET application for 20 minutes. Because Session objects consume memory on the web server, consider decreasing the time-out value to conserve resources.

Important
Use caution when you adjust the session time-out value, because information stored in a user's Session object is lost when the user is not active on the website for the length of the time-out period.

If you decide to use in-process session-state storage, decide also if you also want to use cookies. For more information about cookies, see Cookie mode for session state.

Store session state by using state server

A *state server* maintains session-state date in memory that is outside the worker process. The advantage of this configuration is that session state is preserved when the application worker process recycles. Using a state server is recommended for medium-sized applications.

A state server depends on the Windows state service (Aspnet_state.exe) and requires a machine key to verify session state across the connection.

When a state server runs on the same web server that contains the applications for which it maintains state, a web garden configuration is supported. For increased protection of session state data, consider using a web farm configuration with a separate server that stores session state and shares it among all servers in the farm. Another approach is to use SQL server to maintain out-of-process session state.

Important
The Windows state service (Aspnet_state.exe) must be running for in-process session state to take effect. By default, this service is installed when Windows Server 2012 is installed and is configured for manual start. Change the start behavior to Automatic.

If you decide to store session state by using a state server, make the following design decisions:

- Define a connection string for the state server.

- Specify the number of seconds to wait before the connection time out.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie. For more information about cookies, see Cookie mode for session state.

Store session state by using SQL server

One type of out-of-process session state uses a SQL server to store session state data. The advantage of this configuration is that session state is preserved despite recycling of the application worker process, or if either the Windows state service or the web server goes down.

Note

This setting does not support SQL Azure.

When a SQL server runs on the same web server that has the applications for which it maintains state, it supports a web garden configuration, which increases web server scalability. When the SQL server runs on another server, it supports a web farm configuration, which greatly increases scalability across a group of servers.

Important

The Windows state service (Aspnet_state.exe) must be running for out-of-process session state to take effect. By default, this service is installed when Windows Server 2012 is installed and is configured for manual start. Change the start behavior to Automatic.

Important

Before you configure a SQL server for session state, run the InstallSqlState.sql script on the server. By default, this script is stored in %systemroot%\Microsoft.NET\Framework\V4.0.30319.

If you decide to store session state in a SQL Server database, make the following design decisions:

- Define a connection string for the database.
- Specify the number of seconds to wait before the connection time-out.
- Specify the number of seconds to wait before trying to reconnect.
- Decide whether to enable a custom database.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie. For more information about cookies, see Cookie mode for session state.

Cookie mode for session state

One way to track session state for clients that connect to a web server is to use cookies. You can configure a web server to use cookies, not to use cookies, or to select cookie behavior that depends on the browser that is used for the connection.

A session cookie associates session information with client information for the session. A session is the duration of a user's connection to a site. The cookie is passed together with all requests between a client and a web server in an HTTP header.

Using cookies to track session state is more efficient than any other method that does not use cookies, because ookies do not require any redirection. In addition, they allow users to bookmark web pages, and they retain state if a user leaves one site to visit another and then returns to the original site. The one drawback of user cookies is that users can disable cookies in their browser.

The *Use Device Profile* cookie mode causes the brower to use cookies if the it cookies; otherwise, no cookies are used. If the device profile indicates support for cookies, they are used regardless of whether the user has disabled cookie support.

Important

When you use the Use Device Profile cookie mode, set expired session IDs be regenerated. Doing so allows a web server to expire and regenerate tokens, which gives a potential attacker less time to capture a cookie and gain access to web server content.

The *Auto-Detect* cookie mode causes the mobile device to use cookies if its profile supports cookies; otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET tries to use cookies when cookie support is enabled in the browser. If cookie support is disabled, session state is stored in the URL.

Important

When you use the Auto-Detect cookie mode, set expired session IDs be regenerated. Doing so enables a web server to expire and regenerate tokens, which gives a potential attacker less time to capture a cookie and gain access to web server content. Consider changing the time-out value to less than the 20-minute default.

You can configure session state without using cookies. When you use a Uniform Resource Identifier (URI) to handle session state, the session ID is embedded as a query string in the URI request, and then the URI is redirected to the originally requested URL. The changed URI request is used for the duration of the session, so that no cookie is necessary.

Important

When you use a URI, set expired session IDs be regenerated. Doing so enables a web server to expire and regenerate tokens, which gives a potential attacker less time to capture a cookie and gain access to web server content.

Using a URI to track session state can help you avoid the disadvantages of cookies, including browser support problems and the possibility that users disable cookies. However, using a URI has the following disadvantages:

- Cannot use absolute URLs without losing session state, which means that if a user goes to another application and then returns to the previous one, the user's input no longer exists on the page.
- Does not allow users to bookmark web pages, because session state is lost.

If you decide to use cookies to store session state, make the following design decisions:

- Select a cookie mode: auto detect, use cookies, use device profile, or use URI.

- Unless you selected use URI, specify the name of the cookie.
- Unless you selected use URI, specify the number of minutes before the cookie times out.
- Unless you selected use cookies, decide whether to regenerate an expired session ID.

2.2. Pages and Controls Settings

ASP.NET pages include extra elements that ASP.NET recognizes and processes when the page runs. ASP.NET pages can also contain custom, reusable controls. These custom controls are processed on the server. This lets you use server code to set ASP.NET web page properties.

Note

These settings apply only the ASP.NET Web Forms. They do not apply to ASP.NET MVC or ASP.NET Web Pages.

IIS 8 lets you configure the following ASP.NET page and user controls settings:

- Behavior settings: For example, whether the web page maintains its view state and the view state of any server controls it contains when the current page request ends.
- General settings: For example, namespaces that are included for all pages.
- Compilation settings: For example, whether pages are compiled or interpreted.
- Services: For example, whether session state is enabled.

IIS 8 provides default settings for ASP.NET pages and controls, but you can change those settings as needed. For example, you can set the master page file for a site or enable view state.

Web custom controls are compiled components that run on the server and that encapsulate user-interface and other related functionality into reusable packages. In IIS 8, you can specify the tag prefix and namespace mapping for a custom control that can be used in multiple pages in an application.

Add a custom control when you want to specify the tag prefix/namespace mapping for a custom control that is used in multiple pages in an application.

Note

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

If you decide to configure ASP.NET custom controls, you need the following information for each control you want to configure:

- I must specify the tag prefix of the control.
- I must specify the .NET namespace of the control.
- I must specify the assembly the control is in.

2.3. Application Settings

Configure application settings when you want to store key/value pairs as part of your configuration in the Web.config file. Application settings provide quick and easy access to stored configuration data for your application.

To manage custom controls, you can view a list that contains all custom controls for a particular configuration level. You can sort this list by tag prefix, source or assembly, or scope (local or inherited). You can also group controls by scope to see which custom controls apply at the current configuration level and which custom controls are inherited from a parent level.

Add a custom control when you want to specify the tag prefix/namespace mapping for a custom control that is used in multiple pages in an application.

Note

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

If you decide to configure application settings, you need the following information for each setting you want to configure:

1. Specify a name for the setting.
2. Specify a value for the setting.

2.4. .NET Compilation Settings

For application code to service requests by users, ASP.NET must first compile the code into one or more assemblies. Assemblies are files that have the file name extension .dll. Configure .NET compilation settings in IIS 8 when you want to control how ASP.NET code is compiled.

IIS lets you configure the following .NET compilation settings:

- Batch settings, such as the maximum file size that you can batch and the maximum number of pages that you can have per batched compilation.
- Behavior settings, such as the number of times resources are dynamically compiled before the application is restarted.
- General settings, such as the default programming language that is used in dynamic compilation files.

2.5. .NET Globalization Settings

Globalization is the process of internationalizing application code, then localizing the application to other languages and cultures. The internationalization process makes it possible to translate, store, retrieve, and present application content for any locale by using the same application code base whenever possible. Locale is the combination of both language and cultural environment. This includes date formats, times, currencies, telephone numbers, and so on. Localization means adapting your application to other locales by translating and formatting content according to culture, preferably without touching the code.

You can change globalization settings for ASP.NET applications at the web server level when you want them to apply to all ASP.NET applications on the server. You can also edit ASP.NET globalization settings for sites, applications, directories, and files.

IIS lets you configure the following globalization settings:

- Culture settings, such as the UI culture or UI language.

- Encoding settings, such as encoding for response headers.

Note
Editing a configuration setting changes the setting at the local level and for any child levels that inherit the setting.

Step 3: Plan Data Source Settings

In this phase of building your website, consider the data storage needs of your ASP.NET application. The following sections describe various data-source settings available in IIS:

3.1. Data source connection strings

3.2. ASP.NET providers

3.3. .NET profiles

3.4. .NET roles

3.5. .NET users

3.1. Data source connection strings

A connection string provides the information that an application or provider must have to communicate with a particular database. A connection string usually supplies the server or location of the database server, the particular database to use, and the authentication information. A connection string enables you to connect to databases from managed code applications in a centralized manner.

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

To add a connection string to IIS, provide the following information:

- Specify a name for the connection string. This name must be the same name you reference in your application code to retrieve data from the database.
- Specify the server the database is on.
- Specify the name of the database.
- Provide the credentials, unless using Windows integrate security.

3.2. ASP.NET providers

ASP.NET 2.0 includes several services that store state in a database or other data store. A provider is a software module that implements a uniform interface between one of these services and a data source. In IIS 7, you can set the default provider for your application. You can also configure the provider properties. For example, **Users** is a provider-based feature where one provider stores the user data in SQL whereas another provider stores the user data in a text file.

Add a provider in IIS when you have an application that uses a provider-based service to store data in a database or other data store. For example, the session state service in ASP.NET is a provider-based service that manages per-user session state by storing it in process (in memory in the application domain of the host application), in memory in an external process (the "state server process"), or in a Microsoft SQL Server database.

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

To add a provider for your application, provide the following configuration information:

- Select the IIS feature for the provider to provide: .NET Profiles, .NET Roles, or .NET Users.
- Select the provider type.
- Enter a name for the provider.
- If the feature selected is .NET Users, select any of the following provider behaviors that you want:
  - Enable password reset
  - Enable password retrieval
  - Requires question and answer
  - Requires unique email
  - Store password in secure format
- Provide the name of the connection string to the database.
- Enter the name of the application.

3.3. .NET profiles

The .NET profile feature associates information with an individual user and stores the information in a persistent format. .NET profiles let you manage user information without requiring you to create and maintain your own database.

In .NET profile, you can add properties or groups. A property stores information that is unique to a user, such as the user name. You can then use the information you have stored to present the user with a personalized version of your application. A group organizes related properties together. For example, the different properties of a user’s address information can be grouped together in an Address group.

If you decide to add profile properties or groups for your ASP.NET application, provide the following information:

- For each profile property, provide the property name, data type (such as string or Boolean), a default value, a serialization option (string, XML, binary, or provider specific), whether it is read-only, and whether it is available for anonymous users.
- For each profile group, provide the group name.

3.4. .NET roles

Roles give you an easy way to manage access rules for groups of users. You create users and then assign the users to roles (in Windows, you assign users to groups). For example, you can create a set of pages that you want to restrict to certain users and store those pages in a folder by themselves. Then you can use IIS 8 to define rules that grant and deny access to restricted folders. If an unauthorized user tries to view a restricted page, the user either sees an error or is redirected to a page that you specify.

Roles do not work with anonymous users who access a site, application, or file.

Add roles when you have security settings that you want to apply to a group of users who differ from all existing groups.

If you decide to define roles for your application, provide a name for each role.

Important

To configure roles for your application, you must have previously configured a .NET roles provider.

3.5. .NET users

Associating user identities with an application helps you manage authentication, authorization, and other security-related operations for that application.

Add a user when you want to use IIS to define a user name, e-mail address, password, and security question for initiating an automatic account reset if a user loses or forgets their password. If the **.NET Roles** feature is enabled, you can add users to roles as you create them.

If you decide to configure users for your application, provide the following information for each user:

- User Name (must be unique).
- E-mail (must use standard format: name@domain.com).
- Password (must be a strong password).
- Question (enter a custom question or select from the list).
- Answer to the foregoing question.
- Roles the user is assigned to.

Important

To configure users for your application, you must have previously configured a .NET users provider.

Step 4: Plan Application Security

In this phase of building your website, consider the security needs of your ASP.NET application. The following sections describe application security settings available in IIS 8:

4.1. Isolate Web Applications
4.2. .NET Trust Levels
4.3. .NET Authentication
4.4. Machine Key Settings
4.5. TLS/SSL Communication

4.1. Isolate Web Applications

One of the most effective ways to improve security for your web application is to isolate it from other applications on your web server. An application pool has its own worker process, which processes requests and runs application code. The worker process has a security identifier (SID). And each application pool has a unique application-pool identity. By default, when you create a web application, a new application pool is also created with the same name as the application. If you keep web applications in separate application pools, you can isolate them from one another.

Web application isolation entails the following:

- Site isolation: Separate different applications into different sites with different application pools.
- Least privilege: Run your worker process as a low privileged identity (virtual application pool identity) that is unique per site.
- Temp isolation: Set up a separate ASP.NET temp folder per site and only give access to appropriate process identity.
- Content isolation: Make sure to set an ACL (access control list) on each site root to allow only access to the appropriate process identity.

Tip

It is a good idea to host your website and web application content on a drive other than your system drive (C:).

4.2. .NET Trust Levels

An application *trust level* determines the permissions that the ASP.NET code access security (CAS) policy grants. CAS defines two trust categories: full trust and partial trust. An application that has full trust permissions can access all resource types on a server and perform privileged operations. Applications with full trust are affected only by the security settings of the operating system.

Partial-trust web applications are applications that do not have full trust and have a restricted set of code access permissions. As a result, partial-trust applications are limited in their ability to access secured resources and perform other privileged operations. Certain permissions are denied to partial-trust applications, so resources that require those permissions cannot be directly accessed. Other permissions are granted in a restricted way, so resources that require those permissions might be accessible, but in a limited way. For example, restricted file IO permission

give the application can access to the file system, but only in directories beneath the application's virtual directory root.

By configuring a web application or web service for partial trust, you can restrict the application ability to access crucial system resources or resources that belong to other web applications. By granting only the permissions that the application requires and no more, you can build least privileged web applications and limit potential damage if the web application is compromised by a code injection attack.

The following list shows the restrictions associated with each trust level:

- Full trust applications have unrestricted access to all resource types and can perform privileged operations.
- High, medium, low, or minimal trust applications are unable to call unmanaged code or serviced components, write to the event log, access Message Queuing queues, or access OLE DB data sources.
- High trust applications have unrestricted access to the file system.
- Medium trust applications have restricted file system access and can only access files in their own application-directory hierarchy.
- Low or minimal trust applications cannot access SQL Server databases.
- Minimal trust applications cannot access any resources.

4.3. .NET Authentication

Authentication helps you confirm the identity of clients who request access to your sites and applications. When authentication is enabled, IIS 8 uses the account credentials supplied by the user to determine what permissions the user has been granted and what resources the user can access.

This section describes the authentication modes that are specific to ASP.NET applications.

1. ASP.NET Forms Authentication
2. ASP.NET Impersonation Authentication

ASP.NET Forms Authentication

Forms authentication uses client-side redirection to forward unauthenticated users to an HTML form where they can enter their credentials, which are usually a user name and password. After the credentials are validated, users are redirected to the page they originally requested. Forms authentication often employs cookies to pass user credentials between the server and the client browser.

The following sections describe what you need to know to plan adding forms authentication to your site:

1. Forms authentication basics
2. Authentication cookies

Forms authentication basics

ASP.NET Forms-based authentication works well for sites or applications on public web servers that receive many requests. This authentication mode lets you manage client registration and authentication at the application level, instead of relying on the authentication mechanisms the operating system provides.

Important
Because Forms authentication sends the user name and password to the web server as plaintext, use Secure Sockets Layer (SSL) encryption for the logon page and for all other pages in your application except the home page. For information about SSL, see 4.5. TLS/SSL Communication.

Forms authentication lets users log on by using identities from an ASP.NET membership database. This authentication method uses redirection to an HTML logon page to confirm the identity of the user. You can configure Forms authentication at the site or application levels.

Forms authentication is convenient for the following reasons:

- It allows either a custom data store, such as a SQL server database, or Active Directory to be used for authentication.
- It integrates easily with a web user interface.
- Clients can use any browser.

If you want to use membership roles for authorization, use Forms authentication or a similar custom authentication method.

Important
If you select Forms authentication, you cannot use any of the challenge-based authentication methods at the same time.

By default, the login URL for Forms authentication is Login.aspx. You can create a unique login page for clients who visit a site or application. For example, you might want to collect specific information from visitors, or offer membership to selected pages on the site or selected applications.

The default time-out value for Forms authentication is 30 minutes. Consider changing the time-out value to a shorter period, to shorten the session lifetime and to reduce the chance of cookie replay attacks.

Authentication cookies

Authentication cookies are used as a token to verify that a client has access to some or all pages of an application. By contrast, personalization cookies contain user-specific settings that determine user experience on a specific site or application.

Important
Because authentication cookies are passed between client and server together with every request, always secure authentication cookies using Secure Sockets Layer (SSL). For information about SSL, see 4.5. TLS/SSL Communication.

Cookies are a more efficient way to track visitors to a site than query strings, because they do not require redirection. However, they are browser-dependent, and some browsers do not support their use. In addition, the use of cookie-based authentication is not always effective because some users disable cookie support in their browsers.

By default, the cookie name for ASP.NET applications is .ASPXAUTH. However, you can instead use a unique cookie name and path for each application. Doing so can prevent users who are authenticated for one application from being authenticated for other applications on a web server.

You can choose one of the following cookie modes for your site or application.

Mode	Description
Use cookies	Cookies are always used regardless of device.
Do not use cookies	Cookies are not used.
Auto Detect	Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET checks to determine whether cookies are enabled. This setting is the default.
Use device profile	Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. ASP.NET does not check to determine whether cookies are enabled on devices that support cookies. This setting is the default for IIS 8.

The cookie protection mode defines the function a Forms authentication cookie performs for a specific application. The following table shows the cookie protection modes that you can define:

Mode	Description
Encryption and validation	Specifies that the application use both data validation and encryption to help protect the cookie. This option uses the configured data validation algorithm (based on the machine key). If triple-DES (3DES) is available and if the key is long enough (48 bytes or more 3DES is used for encryption. This setting is the default (and recommended) value.
None	Specifies that both encryption and validation are disabled for sites that are using cookies only for personalization and have weaker

Mode	Description
	security requirements. We do not recommend that you use cookies in this manner; however, it is the least resource-intensive way to enable personalization by using the .NET Framework.
Encryption	Specifies that the cookie is encrypted by using Triple-DES or DES, but data validation is not performed on the cookie. Cookies used in this manner might be subject to plaintext attacks.
Validation	Specifies that a validation scheme verifies that the contents of an encrypted cookie have not been changed in transit.

Important
For security reasons, consider keeping Encryption and Validation cookies separate from each other. The theft of encryption cookies would be a greater security exposure than the theft of validation cookies.

If an application contains objects that clients request frequently, improve application performance by caching those objects. If the user accesses the cached object before the authentication cookie times out, IIS 8 allows the cached object to remain in the cache, and the timer is reset. However, if the user does not access the cached object during that time, IIS 8 removes the cached object from the cache.

Consider enabling this setting under the following circumstances:

- You have a limited amount of memory available for caching.
- You have many objects to cache, because this setting allows only the most frequently requested objects to remain in the cache.

Note
You specify the number of minutes before an authentication cookie times out with **Authentication cookie time-out (in minutes)**.

ASP.NET Impersonation Authentication

Use ASP.NET impersonation when you want to run your ASP.NET application under a security context different from the default security context for ASP.NET applications.

If you enable impersonation for an ASP.NET application, that application can run in one of two different contexts: either as the user authenticated by IIS 8 or as an arbitrary account that you set up. For example, if you use Anonymous authentication and choose to run the ASP.NET application as the authenticated user, the application would run under an account that is set up for anonymous users (typically, IUSR). Likewise, if you chose to run the application under an arbitrary account, it would run under whatever security context was set up for that account.

By default, ASP.NET impersonation is disabled. If you enable impersonation, your ASP.NET application runs under the security context of the user authenticated by IIS 8.

4.4. Machine Key Settings

Machine keys help protect Forms authentication cookie data and page-level view state data. They also verify out-of-process session state identification. ASP.NET uses the following types of machine keys:

- A *validation key* computes a Message Authentication Code (MAC) to confirm the integrity of the data. This key is appended to either the Forms authentication cookie or the view state for a specific page.
- A *decryption key* is used to encrypt and decrypt Forms authentication tickets and view state.

IIS 8 enables you to configure validation and encryption settings and generate machine keys for use with ASP.NET application services, such as view state, forms authentication, membership, roles, and anonymous identification.

Before you generate machine keys for your application, make the following design decisions:

- Decide what validation method to use: AES, MD5, SHA1 (default), TripleDES, HMACSHA256, HMACSHA384, or HMACSHA512.
- Decide what encryption method to use: Auto (default), AES, TripleDES, or DES.
- Decide whether to generate the validation key at runtime automatically.
- Decide whether to generate a unique validation key for each application.
- Decide whether to generate the decryption key at runtime automatically.
- Decide whether to generate a unique decryption key for each application.

4.5. TLS/SSL Communication

Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL) are protocols that provide communication security your website. You can use TLS/SSL to authenticate servers and clients and then use it to encrypt messages between the authenticated parties.

In the authentication process, a TLS/SSL client sends a message to a TLS/SSL server, and the server responds with the information that the server needs to authenticate itself. The client and server perform an additional exchange of session keys, and the authentication dialog ends. When authentication is completed, SSL-secured communication can begin between the server and the client by using the symmetric encryption keys that are established during the authentication process.

To configure TSL/SSL for your website, do the following:

1. Obtain a server certificate from a certification authority (CA). See Server Certificates.
2. Add SSL binding to the site. See SSL Binding.
3. Set IIS to require SSL on the site. See Require SSL for Your Site.
4. Consider using client certificates for your site. See Client Certificates.

Server Certificates

You can obtain a server certificate from a certification authority (CA). Obtaining a server certificate from a certification authority is one step in configuring Secure Sockets Layer (SSL) or Transport Layer Security (TLS). You can obtain server certificates from a third-party CA. A third-party CA might require you to provide proof of identity before a certificate is issued. You can also issue your own server certificates by using an online CA, such as Microsoft Certificate Services.

Digital certificates are electronic files that work like an online password to verify the identity of a user or a computer. They are used to create the SSL encrypted channel that is used for client communications. A certificate is a digital statement that is issued by a certification authority (CA) that vouches for the identity of the certificate holder and enables the parties to communicate in a secure manner by using encryption.

Digital certificates do the following:

- They authenticate that their holders—people, web sites, and even network resources such as routers—are truly who or what they claim to be.
- They protect data that is exchanged online from theft or tampering.

Digital certificates are issued by a trusted third-party CA or a Microsoft Windows public key infrastructure (PKI) using Certificate Services, or they can be self-signed. Each type of certificate has advantages and disadvantages. Each type of digital certificate is tamper-proof and can't be forged.

Certificates can be issued for several uses. These uses include web user authentication, web server authentication, Secure/Multipurpose Internet Mail Extensions (S/MIME), Internet Protocol security (IPsec), Transport Layer Security (TLS), and code signing.

A certificate contains a public key and attaches that public key to the identity of a person, computer, or service that holds the corresponding private key. The public and private keys are used by the client and the server to encrypt the data before it is transmitted. For Windows-based users, computers, and services, trust in a CA is established when there is a copy of the root certificate in the trusted root certificate store and the certificate contains a valid certification path. For the certificate to be valid, the certificate must not have been revoked and the validity period must not have expired.

SSL Binding

You can assign multiple bindings to a site when you have site content that serves different purposes or for which you must use a different protocol. For example, a commerce site might have an application that requires that users log on to an account to purchase merchandise. The company hosts the site over HTTP, but users must log on to their account on an HTTPS page. In this example, the site would have two bindings: one for the HTTP portion and one for the HTTPS portion.

Out of the box, you cannot add bindings for protocols other than HTTP and HTTPS by using IIS Manager. If you want to add a binding for a different protocol, such as a protocol supported by Windows Communication Foundation (WCF), use one of the other administration tools. However, if you install the IIS File Transfer Protocol (FTP) server, you can add FTP bindings by using IIS

Manager. There might also be other modules or third-party functionality available for download that extend the UI.

Require SSL for Your Site

Secure Sockets Layer (SSL) encryption protects confidential or personal information sent between a client and a server. When SSL is enabled, remote clients access your site by using URLs that start with https://.

First configure a server certificate and create an HTTPS binding to enable any SSL settings. Then require Secure Sockets Layer (SSL) encryption in one or more of the following circumstances:

1. When confidential or personal content on your server must be protected by an encrypted channel.
2. When you want users to be able to confirm the identity of your server before they transmit personal information.
3. When you want to use client certificates to authenticate clients that access your server.

Client Certificates

When you want clients to verify their identity before they access content on your web server, configure client certificates. By default, client certificates are ignored.

Before you can configure client certificate on your website, configure a server certificate and create an HTTPS binding to enable any Secure Sockets Layer (SSL) settings.

If you want all clients to verify their identity, specify that client certificates are required. If some clients can access content without first verifying their identity, specify that client certificates are accepted.

Configure an ASP.NET Website on IIS

To install an IIS web server and configure it for ASP.NET web applications, follow the steps listed.

- Step 1: Install IIS and ASP.NET Modules
- Step 2: Configure ASP.NET Settings
- Step 3: Configure Data Source Settings
- Step 4: Configure Application Security

For planning information to review before deployment, see Plan an ASP.NET Website on IIS. For more information, see Build an ASP.NET Website on IIS.

Step 1: Install IIS and ASP.NET Modules

The first step in building an ASP.NET website on IIS 8 is to install IIS along with the ASP.NET modules. Then add your ASP.NET application files to IIS.

The following list shows the tasks required to complete this step:

Installing IIS and ASP.NET Modules

Adding the ASP.NET Application

When you are done, make sure that IIS and the ASP.NET modules are installed, and your ASP.NET application has been added to your website. Then go on to Step 2: Configure ASP.NET Settings.

Installing IIS and ASP.NET Modules

This section describes how to install IIS and the modules that support ASP.NET applications in any of the following ways:

- Install IIS on Windows Server® 2012 using the IIS Manager UI.
- Install IIS on Windows® 8 using the IIS Manager UI.
- Install IIS using the command line.

Important
To install IIS on either Windows Server 2012 or Windows 8, you must sign on as a member of the Administrators group.

If you prefer, you could use the Web Platform Installer (Web PI) to install IIS and third-party applications that run on IIS. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

If you support web applications written using ASP.NET 2 to 3.5, install .NET 3.5 before installing IIS and ASP.NET modules.

To install .NET 3.5 on Windows Server 2012 or Windows 8

1. On the **Start** screen, right-click the **Command Prompt** tile, and then click **Run as administrator**.
2. At the command prompt, type the following: **dism /online /enable-feature /featurename:netfx3**
3. Wait for the command to complete. It could take several minutes.
4. Close the command prompt window.

To install IIS and ASP.NET modules on Windows Server 2012 using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, select Role-based or feature-based installation, and click **Next**.
5. On the **Select destination server** page, select **Select a server from the server pool**, select your server, and click **Next**.
6. On the **Select server roles** page, select **Web Server (IIS)**, and click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select role services** page, note the preselected role services that are installed by default, expand the **Application Development** node, and then select **ASP.NET 4.5**. (If you installed **.NET 3.5**, select **ASP.NET 3.5** also.)
10. On the **Summary of Features to Install** page, confirm your selections, and then click **Install**.
11. In the **Add features that are required for ASP.NET 4.5?** box, click **Add Features**.

    The following additional features are added:

    - **.NET Extensibility 4.5**
    - **ISAPI Extensions**
    - **ISAPI Filters**
    - **.NET Extensibility 3.5** (If **ASP.NET 3.5** was selected)
12. Click **Next**.
13. On the **Confirm installation selections** page, click **Install**.
14. On the **Installation progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
15. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    The default IIS Welcome page is displayed.

To install IIS and ASP.NET modules on Windows 8 by using the UI

1. On the **Start** page, click the **Control Panel** tile.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services** to install the default features.
4. Expand the **Application Development Features** node and click **ASP.NET 4.5** to add the features that support ASP.NET. (If you installed **.NET 3.5**, select **ASP.NET 3.5** also.)

   The following additional features are automatically selected:

- **.NET Extensibility 4.5**
- **ISAPI Extensions**
- **ISAPI Filters**
- **.NET Extensibility 3.5** (If **ASP.NET 3.5** was selected)

5. Click **OK** to close the **Windows Features** dialog box.
6. To verify that IIS installed successfully, type the following into a web browser:

   http://localhost

   The default IIS Welcome page is displayed.

To install IIS and ASP.NET modules by using the command line

- Type the following command at a Command Prompt or into a script:

  Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-ApplicationDevelopment;IIS-ASPNET;IIS-NetFxExtensibility;IIS-ISAPIExtensions;IIS-ISAPIFilter;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI

Adding the ASP.NET Application

This section describes how to add your ASP.NET application to your website in the following ways:

- Adding an ASP.NET Application with the IIS Manager UI
- Adding an ASP.NET application with the command line

In Plan an ASP.NET Website on IIS, you made the following decisions:

- Alias string that becomes part of the application root URL.
- Application pool to run the application in.
- Physical path to the application files on the server.

To add an ASP.NET application by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services**

(IIS) Manager.

2. In the **Connections** pane, expand the **Sites** node.
3. Right-click the site for which you want to create an application, and click **Add Application**.
4. In the **Alias** text box, type a value for the application URL, such as **marketing**. This value is used to access the application in a URL.
5. Click **Select** if you want to select a different application pool than the one listed in the **Application pool** box. In the **Select Application Pool** dialog box, select an application pool from the **Application pool** list and then click **OK**.
6. In the **Physical path** text box, type the physical path of the application's folder, or click the browse button (**...**) to navigate the file system to find the folder.
7. Optionally, click **Connect as** to specify credentials that have permission to access the physical path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option on the **Connect As** dialog box.
8. Optionally, click **Test Settings** to verify the settings that you specified for the application.
9. Click **OK**.

To add an ASP.NET application by using the command line

- To add an application to a site, use the following syntax:

  appcmd add app /site.name:*string* **/path:***string* **/physicalPath:***string*

  The variable **site.name** *string* is the name of the website to which you want to add the application. The variable **path** *string* is the virtual path of the application, such as **/application**, and **physicalPath** *string* is the physical path of the application content in the file system.

  For example, to add an application named **marketing** to a site named **contoso**, with content at **c:\application**, type the following at the command prompt, and then press ENTER:

  appcmd add app /site.name:contoso /path:/marketing /physicalPath:c:\application

See Also

- Step 2: Configure ASP.NET Settings
- Configure an ASP.NET Website on IIS
- Build an ASP.NET Website on IIS
- Step 1: Plan IIS Web Server and ASP.NET Modules Installation

Step 2: Configure ASP.NET Settings

In this phase of building your website, you configure the following IIS server and website settings that support ASP.NET:

2.1. Session State Settings
2.2. Pages and Controls Settings
2.3. Application Settings
2.4. .NET Compilation Settings
2.5. .NET Globalization Settings

2.1. Session State Settings

In Plan an ASP.NET Website on IIS, you decided to use one of the following session-state-storage options:

- **In-process**: Session state is stored in the worker process where the ASP.NET application runs.
- **State Server**: Session state is stored outside the worker process where the ASP.NET application runs.
- **SQL Server**: Session state is stored in a SQL Server database.

You might have also chosen to store session data on the client computer in a cookie. You can even combine the use of cookies with one of the other storage methods. For example, you could store authentication data in a cookie while storing other session-state data in a SQL Server database.

The following sections describe how to configure session state based on the planning decisions you have made:

- Store Session State in Process
- Store Session State by using State Server
- Store Session State by using SQL Server
- Cookie Mode for Session State

Store Session State in Process

This section describes how to configure in-process session state by using either the IIS Manager UI or the command line.

To enable in-process session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **In process**.
4. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
5. In the **Time-out** field, enter a time-out value in the format hh:mm:ss. For example, enter

00:15:00 for 15 minutes.

6. (Optional) Check the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for remote connections.
7. Click **Apply** in the **Actions** pane.

To enable in-process session state by using the command line

- To enable in-process session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /mode:InProc

  The variable **mode:InProc** enables in-process session state. By default, in-process session state is enabled.

  Note
  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To specify the default time that a session object is maintained after the last request associated with the object is made, use the following syntax:

  appcmd set config /section:asp /timeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that a session object is maintained after the last request associated with the object is made. The default value is 00:20:00.

Store Session State by using State Server

This section describes how to configure a state server for session-state storage by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Define a connection string for the state server.
- Specify the number of seconds to wait before the connection time-out.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie.

To configure a state server to maintain session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **State Server.**
4. Type a connection string in the **Connection string** text box, or click **Create** to create a connection string.

5. Type a time-out value in the **Time-out (in seconds)** text box. The default time-out value is 10 seconds.
6. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
7. (Optional) Select the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for connections to the database.
8. Click **Apply** in the **Actions** pane.

To configure a state server to maintain session state by using the command line

- To configure a state server to maintain session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /mode:StateServer /stateConnectionString:*string* **/stateNetworkTimeout:** *timeSpan* **/useHostingIdentity:True|False**

  The variable **mode:StateServer** sets the session state mode to store session data in a state server. By default, the variable *string* defines the connection string that the state server uses. The default setting is tcpip=loopback:42424. The variable *timeSpan* sets the time, in seconds, that the connection to the state server is maintained. The default is 10 seconds. The variable **useHostingIdentity:True|False** enables or disables the use of Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for connections to the database. The default value is **True**.

  Note
  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

Store Session State by using SQL Server

This section describes how to configure a SQL Server database for session-state storage using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Define a connection string for the database.
- Specify the number of seconds to wait before the connection time-out.
- Specify the number of seconds to wait before trying to reconnect.
- Decide whether to enable a custom database.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie.

To configure a SQL server to maintain session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **SQL Server**.
4. Type a connection string in the **Connection string** text box, or click **Create** to create a connection string. If the SQL server is on the server you are managing, select **LocalSqlServer** in the **Connection string** text box.
5. Type a time-out value in the **Time-out** text box.
6. (Optional) Check the **Enable custom database** check box to use a custom database for storing session state data.
7. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
8. (Optional) Check the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for remote connections to the SQL database.
9. Click **Apply** in the **Actions** pane.

To configure a SQL server to maintain session state by using the command line

1. To configure a state server to maintain session state, use the following syntax:

   appcmd set config /commit:WEBROOT /section:sessionState /mode:SqlServer /sqlConnectionString:*string* **/sqlCommandTimeout:** *timeSpan* **/useHostingIdentity:True|False**

   The variable **mode:SqlServer** sets the session state mode to store session data in a SQL server database. By default, the variable *string* defines the connection string that the SQL server uses. The default setting is LocalSqlServer. The variable *timeSpan* sets the time, in seconds, that the connection to the SQL server is maintained. The default is 30 seconds. The variable **useHostingIdentity:True|False** enables or disables the use of Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for connections to the database. The default value is **True**.

   Note

   When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

Cookie Mode for Session State

This section describes how to configure a SQL Server database for session-state storage by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Select a cookie mode: auto detect, use cookies, use device profile, or use URI.
- Unless you selected use URI, specify the name of the cookie.
- Unless you selected use URI, specify the number of minutes before the cookie times out.
- Unless you selected use cookies, decide whether to regenerate an expired session ID.

To configure Use Cookies mode for session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Cookie Settings** area, select **Use Cookies** from the **Mode** drop-down list.
4. Type a cookie name in the **Name** text box, or use the default cookie name, ASP.NET_SessionId.
5. Type a time-out value in the **Time-out** text box, or use 20 minutes, the default time-out value.
6. Click **Apply** in the **Actions** pane.

To configure Use Cookies mode for session state by using the command line

- To configure use cookies mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseCookies /cookieName:string **/timeout:**timeSpan

  The variable **cookieless:UseCookies** configures IIS 8 to use cookies mode for session state. This is the default value. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. For example, to set a cookie-mode cookie for session state (called MyCookie and that expires after 40 minutes), type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseCookies /cookieName:MyCookie /timeout:40

  Note

  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure use device profile cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState

/cookieless:UseDeviceProfile /cookieName:string **/timeout:**timeSpan **/regenerateExpiredSessionId:True|False**

The variable **cookieless:UseDeviceProfile** configures IIS 8 to use device-profile cookie mode for session state. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

Note

When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure auto-detect cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:AutoDetect /cookieName:string **/timeout:**timeSpan **/regenerateExpiredSessionId:True|False**

  The variable **cookieless:AutoDetect** configures IIS 8 to use auto-detect cookie mode for session state. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

  Note

  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure URI cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseUri /regenerateExpiredSessionId:True|False

  The variable **cookieless:UseUri** configures IIS 8 to use URI cookie mode for session state. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

  Note

  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.2. Pages and Controls Settings

This section describes how to configure ASP.NET pages and control settings by using either the IIS Manager UI or the command line.

Note
These settings apply only to ASP.NET Web Forms. They do not apply to ASP.NET MVC or ASP.NET Web Pages.

Edit Pages and Controls

You can edit the settings of existing pages and controls, and the changes are reflected into the Web.config file.

To edit settings for pages and controls by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Pages and Controls**.
3. On the **Pages and Controls** page, edit settings as necessary.
4. When finished, click **Apply** in the **Actions** pane.

To edit settings for pages and controls by using the command line

- To enable or disable page output buffering, use the following syntax:
  appcmd set config /commit:WEBROOT /section:pages /buffer:True|False
  The variable **buffer:True** enables buffering of page output. The default is **True**.

Add a Custom Control

In Plan an ASP.NET Website on IIS, you made the following design decisions for each custom control you want to add:

- Specify the tag prefix of the control.
- Specify the .NET namespace of the control.
- Specify the assembly the control is in.

To add a custom control by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Pages and Controls**.
3. In the **Actions** pane, click **Register Controls**.
4. In the **Actions** pane, click **Add Custom Control**.
5. In the **Add Custom Control** dialog box, type a tag prefix in the **Tag prefix** text box.
6. In the **Namespace** text box, type the namespace to which the custom control belongs. This is the same namespace that was specified in the application code.

7. In the **Assembly** text box, type the name of the source or assembly for the custom control, and then click **OK**.

To add a custom control by using the command line

- To add a custom control, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/pages /+"controls.[assembly='*string***',namespace='***string***',src='***string***', tagName='***string***', tagPrefix='***string***']"**

  The variable **assembly** *string* is the name of the assembly that contains the COM control implementation and requires that the **namespace** attribute is also set. The variable **namespace** *string* is the namespace that is associated with the tag prefix and requires that the **assembly** attribute is also set, if the COM control is not located in the application code directory. The variable **src** *string* is the name of the file that contains the user control and requires that the **tagName** attribute is also set. The variable **tagName** *string* is the name of the control to use in the page and requires that the **src** attribute is also set. The variable **tagPrefix** *string* is the tag prefix that is being mapped to a source file or namespace and assembly. This attribute requires one of the following combinations of other attributes:

  - **namespace** to define a COM control, if the control is in the application code directory.
  - **namespace** and **assembly** to define a COM control.
  - **tagName** and **src** to define a user control.

  For example, to add a custom control that uses the specified source for a user control, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:system.web/pages /+"controls.[tagPrefix='MyTag', tagName='MyControl',src='controls/MyControl.asx']"

  Note
  When you use Appcmd.exe to configure the **<controls>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.3. Application Settings

This section describes how to configure application settings by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions for each application setting you want to configure:

1. Specify a name for the setting.
2. Specify a value for the setting.

To create an application setting by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Application Settings**.
3. In the **Actions** pane, click **Add**.
4. In the **Add Application Setting** dialog box, type a name for the application setting in the **Name** text box.
5. In the **Add Application Setting** dialog box, type a value for the application setting in the **Value** text box, and then click **OK**.

To create an application setting by using the command line

- To add an application setting, use the following syntax:

  appcmd set config /commit:MACHINE /section:appSettings /+"[key='*string***',value='***string***']"**

  The variable **key** *string* is the name of the application setting and the variable **value** *string* is the value for the application setting. For example, to add an application setting for the name of your application, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:appSettings /+"[key='Application Name',value='MyApplication']"

2.4. .NET Compilation Settings

Edit .NET compilation settings in IIS 8 when you want to control how ASP.NET code is compiled. This section describes how to edit the compilation settings for you ASP.NET application.

To configure .NET compilation settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Compilation**.
3. On the **.NET Compilation** page, edit settings as necessary.
4. When finished, click **Apply** in the **Actions** pane.

To configure .NET compilation settings by using the command line

- To enable or disable batch compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /batch:True|False

  The variable **batch:True** enables batch compilation, and the variable **batch:False** disables batch compilation. The default value is **True**.
- To change the maximum file size (in KB) of the batch compilation, use the following syntax:

appcmd set config /commit:WEBROOT /section:compilation /maxBatchGeneratedFileSize:*int*

The variable *int* is the maximum file size (in KB). The default value is 1000.

- To specify the maximum number of pages per batched compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /maxBatchSize:*int*

  The variable *int* is the maximum number of pages per batched compilation.

- To define the time-out period, in seconds, for batch compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /batchTimeout:*timeSpan*

  The variable *int* is the maximum number of pages per batched compilation.

- To specify compilation of retail or debug binaries for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /debug:True|**False**

  The variable **debug:True** enables compilation of debug binaries and the variable **debug:False** enables compilation of retail binaries. The default value is **False**.

- To specify the number of times resources can be dynamically recompiled before an application restarts for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /numRecompilesBeforeAppRestart:*int*

  The variable *int* is the number of times resources are dynamically recompiled before the application is restarted. The default is 15.

- To enable or disable URL line pragma for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /urlLinePragmas:True|**False**

  The variable **urlLinePragmas:True** enables URL line pragma and the variable **urlLinePragmas:False** disables URL line pragma. The default value is **False**.

- To enable or disable Visual Basic explicit compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /explicit:True|**False**

  The variable **explicit:True** enables Visual Basic explicit compilation and the variable **explicit:False** disables explicit compilation. The default value is **False**.

- To enable or disable Visual Basic strict compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /strict:True|**False**

  The variable **strict:True** enables Visual Basic strict compilation and the variable **strict:False** disables strict compilation. The default value is **False**.

- To add an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /+"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.

- To remove an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /-"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.
- To remove an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /-"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.
- To change the default programming language that is used in dynamic compilation files for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /defaultLanguage:*string*

  The variable *string* is the default programming language. The default value is vb. For example, to change the default programming language that is used in dynamic compilation files from Visual Basic to C#, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:compilation /defaultLanguage:c#
- To specify the directory to use for temporary file storage during compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /tempDirectory:*string*

  The variable *string* is the directory path.

  Note
  When you use Appcmd.exe to configure the **<compilation>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.5. .NET Globalization Settings

This section describes how to edit the globalization settings for you ASP.NET application by using either the IIS Manager UI or the command line.

To edit .NET globalization settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Globalization**.
3. On the **.NET Globalization** page, in the property sheet, click to select the global setting you want to edit, and select a value from the drop-down list.

4. In the **Actions** pane, click **Apply**.

To edit .NET globalization settings by using the command line

- To edit the default culture for processing Web requests, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /culture:*string*

  The variable *string* is the default culture for processing web requests. For example, to change the default culture for processing Web requests to US English, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:globalization /culture:en-us

  Note
  When you use Appcmd.exe to configure the **<globalization>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To enable or disable client-based culture, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /enableClientBasedCulture:True|False

  The variable **enableClientBasedCulture:True** enables client-based culture and the variable **enableClientBasedCulture:False** disables client-based culture. The default is **False**.

- To edit the default culture for processing locale-dependent resource searches, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /uiCulture:*string*

  The variable *string* is the default culture for processing locale-dependent resource searches.

- To edit the default culture for processing locale-dependent resource searches, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /uiCulture:*string*

  The variable *string* is the default culture for processing locale-dependent resource searches.

- To change the default encoding when parsing .aspx, .asmx, and .asax files, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /fileEncoding:*string*

  The variable *string* is the default encoding to be used when parsing .aspx, .asmx, and .asax files.

- To change the header encoding for responses, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /responseHeaderEncoding:*string*

  The variable *string* is the header encoding used for responses. The default is utf-8.

- To edit the content encoding for responses, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /responseEncoding:*string*

  The variable *string* is the content encoding used for responses. The default is utf-8.

Step 3: Configure Data Source Settings

In this phase of building your ASP.NET website, you configure the data source settings that are available in IIS.

3.1. Data Source Connection Strings

3.2. ASP.NET Providers

3.3. .NET Profiles

3.4. .NET Roles

3.5. .NET Users

3.1. Data Source Connection Strings

This section describes how to create a database connection string in IIS using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions about the need to add a connection string:

1. Specify a name for the connection string.
2. Specify the server the database is on.
3. Specify the name of the database.
4. Provide the credentials, unless using Windows integrate security.

To create a database connection string by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Connection Strings**.
3. On the **Connection Strings** page, click **Add** in the **Actions** pane.
4. In the **Add Connection String** dialog box, type a name for the connection string, such as MyConnection, in the **Name** text box.

   Note

   The name that you enter in IIS Manager is the same name that you reference in your application code to retrieve data by using this connection string.
5. With the **SQL Server** option selected, type the name of the server that hosts the database in the **Server** text box and type the name of the database in the **Database** text

box.

6. Under **Credentials**:
   - Select **Use Windows Integrated Security**.

     -Or-
   - Select **Specify credentials** and click **Set**. Type a user name and password for an account that can connect to the server and database in the **User name** and **Password** text boxes. Then type the same password in the **Confirm password** text box, and click **OK**.
7. Click **OK**.

To create a database connection string by using the command line

- To create a database connection string, use the following syntax:

  appcmd set config /commit:MACHINE /section:connectionStrings /+"[connectionString='*string***',name='***string***',providerName='***string***']"**

  The variable **connectionString** *string* is the connection string value, the variable **name** *string* is the key to use to access the connection string value, and the variable **providerName** *string* is the name of the ADO.NET provider to use to access the underlying data store. For example, to create a connection string for an application that uses the Northwind database, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:connectionStrings /+"[connectionString='Data Source=localhost;Integrated Security=SSPI;Initial Catalog=Northwind;',name='Northwind',providerName='System.Data.SqlClient ']"

3.2. ASP.NET Providers

This section describes how to add a .NET provider to IIS using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions about adding a .NET provider to IIS:

- Select the IIS feature for the provider to provide: .NET Profiles, .NET Roles, or .NET Users.
- Select the provider type.
- Enter a name for the provider.
- If the feature selected is .NET Users, select any of the following provider behaviors that you want:
  - Enable password reset
  - Enable password retrieval
  - Requires question and answer
  - Requires unique email
  - Store password in secure format

- Provide the name of the connection string to the database.
- Enter the name of the application.

To add a provider by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Providers**.
3. On the **Providers** page, under **Feature**, select one of the following features:
   - **.NET Roles**: to configure the provider to provide an interface between the ASP.NET role management service (the "role manager") and role data sources.
   - **.NET Users**: to configure the provider to provide an interface between the ASP.NET membership service and membership data sources.
   - **.NET Profile**: to configure the provider to provide an interface between the ASP.NET profile service and profile data sources.
4. In the **Actions** pane, click **Add**.
5. In the **Add Provider** dialog box, select a provider type from the **Type** drop-down list.
6. In the **Name** text box, type a name for the provider.
7. If the **.NET Users** feature was selected in step #3, in the **Profile properties** section, under **Behavior**, set the value of one or more of the following behaviors to **True** to enable the behavior:
   - **EnablePasswordReset**: indicates whether passwords can be reset by using the provider *ResetPassword* method. The default setting is **False**.
   - **EnablePasswordRetrieval**: indicates whether passwords can be retrieved by using the provider *GetPassword* method. The default setting is **False**.

     Important
     Some providers, such as the Active Directory provider, do not support the retrieval of passwords. For these providers, the value of the **enablePasswordRetrieval** attribute is always **False** and cannot be changed in configuration.
   - **RequiresQuestionAndAnswer**: indicates whether a password answer is supplied when the program calls the provider *GetPassword* and *ResetPassword* methods. The default setting is **False**.
   - **RequiresUniqueEmail**: indicates whether each registered user must have a unique e-mail address. The default setting is **False**.
   - **StorePasswordInSecureFormat**: indicates whether passwords are hashed. The default setting is **False**.
8. Under **Data**, type the name of the connection string that is used to connect to the database in the **ConnectionStringName** text box.

   Note
   This is the same connection string that was configured under the **Connection Strings** feature. For more information about how to configure connection strings,

see **Configuring database connection strings**. If the *WindowsTokenRoleProvider* was selected in step #5 of this procedure, a connection string name is optional.

9. Optionally, in the **Profile properties** section, under **General**, type the virtual path of the application in the **ApplicationName** text box. If you do not specify a value in the **ApplicationName** text box, the membership API defaults to HttpContext.Current.Request.ApplicationPath. This path can be determined only at runtime.
10. Optionally, in the **Profile properties** section, under **General**, type a description of the provider in the **Description** text box.
11. Click **OK**.

To add a provider by using the command line

1. To add a .NET role provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:roleManager /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. For example, to configure an ASP.NET application to use the **SqlRoleProvider** class to store and retrieve role information, type the following at the command prompt, and then press Enter:

   appcmd set config /commit:MACHINE /section:roleManager /+"providers.[name='SqlProvider',type='System.Web.Security.SqlRoleProvider',connectionStringName='SqlServices',applicationName='SampleApplication']"
2. To add a .NET user provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:membership /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***',enablePasswordReset='True|False',enablePasswordRetrieval='True|False',requiresQuestionAndAnswer='True|False',requiresUniqueEmail='True|False',storePasswordInSecureFormat='True|False']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. The variable **enablePasswordReset True** enables password reset by using the provider *ResetPassword* method. The default setting is **False**. The variable **enablePasswordRetrieval True** enables password retrieval by using the provider *GetPassword* method. The default setting is **False**. The variable **requiresQuestionAndAnswer True** requires that a password answer must be supplied when the program calls the provider *GetPassword* and *ResetPassword* methods. The

default setting is **False**. The variable **requiresUniqueEmail True** requires that each registered user has a unique e-mail address. The default setting is **False.** The variable **storePasswordInSecureFormat True** requires that all passwords must be hashed. The default setting is **False**. For example, to configure an ASP.NET application to use the **SqlMembershipProvider** class to store and retrieve user information, type the following at the command prompt, and then press Enter:

appcmd set config /commit:MACHINE /section:membership /+"providers.[name='SqlProvider',type='System.Web.Security.SqlMembershipProvider',connectionStringName='SqlServices',applicationName='/',enablePasswordRetrieval='False',enablePasswordReset='True',requiresQuestionAndAnswer='True']"

3. To add a .NET profile provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:profile /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. For example, to configure an ASP.NET application to use the **SqlProfileProvider** class to store and retrieve profile information, type the following at the command prompt, and then press Enter:

   appcmd set config /commit:MACHINE /section:profile /+"providers.[name='SqlProvider',type='System.Web.Profile.SqlProfileProvider',connectionStringName='SqlServices',applicationName='SampleApplication']"

3.3. .NET Profiles

This section describes how to add a profile property and a profile group.

In Plan an ASP.NET Website on IIS, you made the following design decisions about profile properties and groups to IIS:

- For each profile property, provide the property name, data type (such as string or Boolean), a default value, a serialization option (string, XML, binary, or provider specific), whether it is read-only, and whether it is available for anonymous users.
- For each profile group, provide the group name.

To add a .NET profile property by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Profile**.
3. In the **Actions** pane, click **Add Property** to add a .NET profile property or, to add a .NET profile property to a group, select the group to which you want to add the .NET profile property, and then click **Add Property To Group**.
4. In the **Add .NET Profile Property** dialog box, type a name for the .NET profile property

in the **Name** text box.

5. Under **Data Type**, select one of the following data types:
   - **System.Boolean**: to configure the .NET profile property to have a value of either **True** or **False**.
   - **System.Char**: to configure the .NET profile property to contain Unicode Characters.
   - **System.DateTime**: To configure the .NET profile property to contain dates and times with values that range from 12:00:00 midnight, January 1, 0001 Anno Domini (A.D.) or Common Era (C.E.) through 11:59:59 P.M., December 31, 9999 A.D. (C.E.)
   - **System.Decimal**: to configure .NET profile property to contain decimal numbers that range from positive 79,228,162,514,264,337,593,543,950,335 to negative 79,228,162,514,264,337,593,543,950,335. The **Decimal** value type is appropriate for financial calculations that require many significant integral and fractional digits and no round-off errors.
   - **System.Double**: To configure the .NET profile property to contain a double-precision 64-bit number with values that range from negative 1.79769313486232e308 to positive 1.79769313486232e308, as well as positive or negative zero, **PositiveInfinity**, **NegativeInfinity**, and Not-a-Number (NaN). For more information about the **Double** value type, see Double Structure.
   - **System.Int32**: to configure the .NET profile property to contain a signed integer with a value that ranges from negative 2,147,483,648 through positive 2,147,483,647.
   - **System.Int64**: to configure the .NET profile property to contain an integer with a value that ranges from negative 9,223,372,036,854,775,808 through positive 9,223,372,036,854,775,807.
   - **System.Single**: To configure the .NET profile property to contain a single-precision 32-bit number with values that range from negative 3.402823e38 to positive 3.402823e38, as well as positive or negative zero, **PositiveInfinity**, **NegativeInfinity**, and Not-a-Number (NaN). For more information about the **Single** value type, see Single Structure.
   - **System.String**: To configure the .NET profile property to be a sequential collection of **Char** structures that represents a string of text. This setting is the default.
6. In the **Default value** text box, type a value that the property will be initialized with.
7. Under **Serialization option**, select one of the following serialization formatters:
   - **String**: Select this option when the settings property is serialized as plaintext. This setting is the default.
   - **XML**: Select this option to serialize only public properties and fields. XML serialization does not preserve type fidelity. This option is useful when you want to provide or consume data without restricting the application that uses the data. Because XML is an open standard, it is an attractive choice for sharing data on the Web.
   - **Binary**: Select this option to preserve type fidelity. Binary serialization is useful for preserving the state of an object between different invocations of an application. For example, you can share an object between different applications by serializing it to the Clipboard. You can serialize an object to a stream, to a disk, to memory, over the network, and so on.

- **Provider Specific**: Select this option when the settings provider has implicit knowledge of the property or its type and can pick an appropriate serialization mechanism. This option is often used for custom serialization.

8. Check the **Read only** box to configure the .NET profile property so that it cannot be modified.
9. Check the **Available for anonymous users** box to make the .NET profile property available for unauthenticated users.
10. Click **OK**.

To add a .NET profile property by using the command line

- To add a .NET Profile property, use the following syntax:

  appcmd set config /commit:MACHINE /section:profile /+"properties.[name='*string*',type='*string*',defaultValue='*string*',serializeAs='String|Xml|Binary|ProviderSpecific',readOnly='True|False',allowAnonymous='True|False']"

  The variable **name** *string* is the name of the profile property. The variable **type** *string* is the data type. The variable **defaultValue** *string* is the value with which the profile property is initialized. The variable **serializeAs 'String|Xml|Binary|ProviderSpecific** sets the serialization formatter. The variable **readOnly True** configures the .NET profile property so that it cannot be modified. The default value is **False**. The variable **allowAnonymous True** makes the .NET profile property available for unauthenticated users. The default value is **False**. For example, to specify a profile property that holds a recent search list collection, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:profile /+"properties.[name='RecentSearchList',type='System.Collections.Specialized.StringCollection',serializeAs='Xml',allowAnonymous='true']"

To add a .NET profile group by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Profile**.
3. In the **Actions** pane, click **Add Group**.
4. In the **Add Group** dialog box, type a name for the .NET profile group in the **Name** text box, and then click **OK**.

3.4. .NET Roles

This section describes how to add a .NET role by using the IIS Manager UI.

To add a .NET role by using the UI

1. Open IIS Manager and navigate to the level you want to manage.

2. In **Features View**, double-click **.NET Roles**.
3. In the **Actions** pane, click **Add**.
4. In the **Add .NET Role** dialog box, type the name of the role in the **Name** text box, and then click **OK**.

3.5. .NET Users

This section describes how to configure .NET users by using the IIS Manager UI.

In Plan an ASP.NET Website on IIS, you made the following design decisions for each user you want to add to IIS:

- User Name
- E-mail
- Password
- Question
- Answer
- Role or roles

To add a .NET user by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Users**.
3. In the **Actions** pane, click **Add**.
4. In the **.NET User Account Details** dialog box, enter the following information:
   - User Name (must be unique).
   - E-mail (must use standard format: name@domain.com).
   - Password (must be a strong password).
   - Confirm Password (must match the password).
   - Question (enter a custom question or select from the list).
   - Answer to the question.
5. Click **Next** to select roles for this user. If you have not enabled **.NET Roles**, click **Finish**.
6. In the **.NET User Roles** dialog box, optionally select available roles from the **Roles** box, and then click **Finish**.

Step 4: Configure Application Security

In this phase of building your ASP.NET website, you configure the security settings that are available in IIS. The following sections discuss common security settings for ASP.NET applications:

4.1. Isolate Web Applications
4.1. .NET Trust Levels
4.2. .NET Authentication
4.3. Machine Key Settings
4.4. TLS/SSL Communication

4.1. Isolate Web Applications

Implement the following recommendations to isolate websites and web applications on your server.

- Use one application pool per website or web application.
- Limit access to site folders and files to the application pool identity.
- Set up a separate ASP.NET temp folder per site and only give access to the application pool identity.
- Make sure to set an ACL (access control list) on each site root to allow only access to the application pool identity.

If you have more than one application per application pool, consider creating enough application pools and moving some of the applications to the new pools.

To create an application pool

1. Open IIS Manager.
2. In the **Connections** pane, click **Application Pools**.
3. In the **Actions** pane, click **Add Application Pool**.
4. In the **Name** box, type a unique name for the application pool.
5. Select the **.NET Framework version** and **Managed pipeline mode**.
6. Click **OK**.

To move an application to another application pool

1. Open IIS Manager.
2. In the **Connections** page, select the website or web application you want to move.
3. In the **Actions** pane, click **Basic Settings**.
4. On the **Edit Site** dialog, click **Select** to open the **Select Application Pool** dialog, and then select the application pool from the **Application pool** menu.
5. Click **OK** to close the **Select Application Pool** dialog, and click **OK** to close the **Edit Site** menu.

To add an application pool identity to a folder or file ACL

1. Open Windows Explorer and navigate to the folder or file.
2. Right click the folder or file, and then click **Properties**.

3. Select the **Security** tab, and then click **Edit**.
4. Click **Add**, click **Locations**, and select your server as the location to search.
5. In the **Enter the object names to select** box, type **IIS APPPOOL\applicationPoolName**, where *applicationPoolName* is the application pool identity.
6. Click **OK**, click **OK**, and click **OK** again to close the dialogs.

4.1. .NET Trust Levels

This section describes how to set application trust level by using either IIS Manager UI or the command line.

To set a trust level by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Trust Levels**.
3. On the **.NET Trust Levels** page, select a trust level from the **Trust level** drop-down list, and then click **Apply** in the **Actions** pane.

To set a trust level by using the command line

1. To set a trust level, use the following syntax:

   appcmd set config /commit:WEBROOT /section:trust /level: Full | High | Medium | Low | Minimal

   The level attribute uses one of five values that correspond to preconfigured CAS policy files. For example, to set a trust level of **Full**, type the following at the command prompt, and then press ENTER:

   appcmd set config /commit:WEBROOT /section:trust /level:Full

   Note
   When you use Appcmd.exe to configure the trust element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

4.2. .NET Authentication

In Plan an ASP.NET Website on IIS, you made design decisions about what authentication mode was right for your application. The following sections decide how to configure authentication for your ASP.NET application:

1. ASP.NET Forms Authentication
2. ASP.NET Impersonation Authentication

ASP.NET Forms Authentication

This section describes how to configure ASP.NET forms authentication by using either the IIS Manager UI or the command line.

To configure forms authentication by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Authentication**.
3. On the **Authentication** page, select **Forms Authentication**.
4. In the **Actions** pane, click **Enable** to use Forms authentication with the default settings.
5. In the **Actions** pane, click **Edit**.
6. In the **Edit Forms Authentication Settings** dialog box, in the **Login URL** text box, type the name of the page where clients log in.
7. In the **Authentication cookie time-out (in minutes)** text box, type the number of minutes you want to use for the time-out value.
8. From the **Mode** list, select the cookie mode you want to use.
9. In the **Name** text box, type the name of the cookie.
10. From the **Protection mode** list, select the protection mode you want to use.
11. Select the **Requires SSL** check box.
12. Select the **Extend cookie expiration on every request** check box, and then click **OK**.

To configure forms authentication by using the command line

- To enable forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /mode: None | Windows | Windows Live ID | Forms

  By default, IIS 8 sets the **mode** attribute to Windows, which disables Forms authentication. If you set the attribute to Forms, you enable Forms authentication. For example, to enable Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /mode:Forms

  Note
  When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To specify the login URL for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.loginURL:*string*

The variable **forms.loginURL** *string* is the name of the page where clients login. The default value is Login.aspx. For example, to specify the login URL for Forms authentication, type the following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.loginURL:login.aspx

Note

When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To specify the authentication time-out for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.timeout:*TimeSpan*

  The variable **forms.timeout** *TimeSpan* is the time in minutes when the cookie used for authentication expires. The default value is 30 minutes. For example, to specify the authentication time-out for Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.timeout:30

  Note

  When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, you must specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure the cookie name for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.name:*string*

  The variable **forms.name** *name* is the name of the cookie used for Forms authentication. The default value is .ASPXAUTH. For example, to configure the cookie name for Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.name:.ASPXUTH

- To configure the cookie mode for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.cookieless: UseUri | UseCookies | AutoDetect | UseDeviceProfile

  The default value for **forms.cookieless** is **UseDeviceProfile**. For example, to configure the cookie mode for Forms authentication to use the setting Use Device Profile, type the

following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.cookieless:UseDeviceProfile

- To configure the cookie protection mode for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.protection: All | None | Encryption | Validation

  The default value for **forms.protection** is **All**. For example, to configure the cookie protection mode for Forms authentication to use the setting Encryption and Validation, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.protection:All

- To require SSL for an authentication cookie, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.requireSSL: True | False

  The default value for **forms.requireSSL** is False. If you set this attribute to **True**, you require SSL. For example, to require SSL for an authentication cookie, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.requireSSL:True

- To cache frequently requested content, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.slidingExpiration: True | False

  The default value for **forms.slidingExpiration** is **True**. For example, to cache frequently requested content, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.slidingExpiration:True

ASP.NET Impersonation Authentication

To configure Impersonation Authentication by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Authentication**.
3. On the **Authentication** page, select **ASP.NET Impersonation**.
4. In the **Actions** pane, click **Enable** to use ASP.NET Impersonation authentication with the

default settings.

5. Optionally, in the **Actions** pane, click **Edit** to set the security principal.
6. In the **Edit ASP.NET Impersonation Settings** dialog box, select either **Specific user** or **Authenticated user**. Whichever you decide, IIS uses this identity for the security context of the ASP.NET application. By default, IIS 8 is set to impersonate the authenticated user.
7. Click **OK** to finish or proceed to the next optional steps to change the identity to impersonate.
8. Optionally, click **Set** to change the **Specific user** identity.
9. In the **Set Credentials** dialog box, enter the name of an existing user account in **User name,** the password associated with that user account in **Password**, and then the exact same value in **Confirm password** for a new account IIS should use for anonymous access.
10. Click **OK** to close the **Set Credentials** dialog box.
11. Click **OK** to close the **Edit ASP.NET Impersonation Settings** dialog box.

To configure Impersonation Authentication by using the command line

- To enable or disable ASP.NET Impersonation, use the following syntax:

  appcmd set config /commit:WEBROOT /section:identity /impersonate:true | **false**

  By default, IIS sets the **impersonate** attribute to false, which disables ASP.NET Impersonation authentication. If you set the attribute to true, you enable ASP.NET Impersonation authentication. For example, to enable ASP.NET Impersonation authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:identity /impersonate:true

- Optionally, you can set the account for IIS to impersonate, using the following syntax:

  appcmd set config /commit:WEBROOT /section:identity /userName:*string* **/password:***string*

  The variable **userName** *string* is the account IIS uses to impersonate and the variable **password** *string* is the password. For example, to use an account named **Moe** for IIS to impersonate, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:identity /userName:Moe /password:pass@word1

  Note

  When you use Appcmd.exe to configure the identity element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

4.3. Machine Key Settings

This section describes how to generate machine keys for your ASP.NET application by using the IIS Manager UI.

To generate machine keys by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Machine Key**.
3. On the **Machine Key** page, select a validation method from the **Validation method** list. The default validation method is **SHA1**.
4. Choose an encryption method from the **Encryption method** list. The default encryption method is **Auto**.
5. Optionally, configure settings for validation and decryption keys.
6. In the **Actions** pane, click **Generate Keys**, and then click **Apply**.

4.4. TLS/SSL Communication

This section describes how to configure TLS/SSL security for your application.

After obtaining a server certificate from a certification authority (CA), work through the procedures in the following sections:

1. SSL Binding
2. Require SSL for Your Site
3. Client Certificates

SSL Binding

This section describes how to add an SSL binding to your site by using either the IIS Manager UI or the command line.

To add an SSL binding to a site

1. Open IIS Manager.
2. In the **Connections** pane, expand the **Sites** node in the tree, and then click to select the site for which you want to add a binding.
3. In the **Actions** pane, click **Bindings**.
4. In the **Site Bindings** dialog box, click **Add**.
5. In the **Add Site Binding** dialog box, in the **Type** list, select **https**.
6. From the **IP address** list, select **All Unassigned** (unless there is a specific IP address you want to use).
7. In the **Port** box, type the number of the port (the default is 443).
8. In the **Host name** box, type the name of the host computer.
9. If you want multiple secure websites to be served using the same IP address, select the

Require Server Name Indication check box.

10. From the **SSL certificate** list, select the certificate for your website. If your certificate doesn't appear in the list, click **Select** and search for the certificate using the **Select Certificate** dialog box.
11. Click **OK**.

1. To add a binding to a site, use the following syntax:

   appcmd set site /site.name:*string* **/+bindings.[protocol='***string***',bindingInformation='***string***']**

   The variable **site.name** *string* is name of the site to which you want to add a binding. The variable **protocol** *string* is the protocol that you want to use, and the variable **bindingInformation** *string* is the combination of IP address, port, and host header.

   For example, to configure a site named **contoso** to have an HTTPS binding for all IP addresses, on port 443, without a host header, type the following at the command prompt, and then press ENTER:

   appcmd set site /site.name:contoso /+bindings.[protocol='https',bindingInformation='*:443:']

Require SSL for Your Site

This section describes how to require SSL for your website by using the IIS Manager UI or the command line.

To require SSL using the UI

1. Open IIS Manager and navigate to the level you want to manage.

   Make sure that you are at the site, application, or directory level; **SSL Settings** are not available at the server level.

   Note
   If you want to configure SSL at the file level, navigate to the file in **Content View** and then click **Switch to Features View** in the **Actions** pane.
2. In **Features View**, double-click **SSL Settings**.
3. On the **SSL Settings** page, select **Require SSL**.
4. In the **Actions** pane, click **Apply**.

To require SSL using the command line

- To require SSL, use the following syntax:

  appcmd set config "*site* | *URL***" /section:access /sslFlags:Ssl /commit:APPHOST**

  The variable *site* | *URL* is the site, application, virtual directory, or file where you want IIS 8 to require SSL. For example, to require SSL for the Default Web Site, type the following

at the command prompt, and then press ENTER:

appcmd set config "Default Web Site" /section:access /sslFlags:Ssl /commit:APPHOST

- To require SSL for the file iisstart.htm on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "http://localhost/iisstart.htm" /section:access /sslFlags:Ssl /commit:APPHOST

Client Certificates

This section describes how to

To specify client certificates by using the UI

1. Open IIS Manager and navigate to the level you want to manage.

   Make sure that you are at the site, application, or directory level; **SSL Settings** are not available at the server level.

   Note
   If you want to configure SSL at the file level, navigate to the file in **Content View** and then click **Switch to Features View** in the **Actions** pane.
2. In **Features View**, double-click **SSL Settings**.
3. On the **SSL Settings** page, optionally select **Require SSL**. You do not need SSL to **Ignore** or **Accept** client certificates.
4. On the **SSL Settings** page, in the **Client certificates** area, use one of the following procedures:
   - Select **Ignore** if you do not want to accept a client certificate even if a client presents one.
   - Select **Accept** to accept client certificates.
   - Select **Require** to require client certificates. To use **Require Client Certificates**, you must enable **Require SSL**.
5. In the **Actions** pane, click **Apply**.

To specify client certificates by using the command line

- To specify whether to use client certificates, use the following syntax:

  appcmd set config "*site* | *URL***"/section:access /sslFlags: Ssl | SslNegotiateCert | SslRequireCert /commit:APPHOST**

  The variable *site* | *URL* is the site, application, virtual directory, or file where you want IIS to enable client certificates. For example, to accept client certificates for the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "Default Web Site"/section:access /sslFlags:SslNegotiateCert /commit:APPHOST

- To accept client certificates for the file iisstart.htm on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "http://localhost/iisstart.htm"/section:access /sslFlags:SslNegotiateCert /commit:APPHOST

  You can specify one or more of the values for the sslFlags attribute. If you want more than one value, separate each value with a comma (,). For example, to specify a requirement for both SSL and client certificates on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "Default Web Site"/section:access /sslFlags:Ssl,SslRequireCert /commit:APPHOST

Build an FTP Site on IIS

File Transfer Protocol (FTP) simple is a protocol for transferring files between computer systems. IIS 8 includes an FTP server that is easy to configure.

The document shows how to install and configure the FTP server on an existing IIS 8 web server. The first two steps are required. All other steps are optional but recommended.

In this document

- Prerequisites
- Step 1: Install FTP on an Existing IIS Web Server
- Step 2: Add an FTP Site
- Step 3: Configure FTP Site Defaults
- Step 4: Configure Firewall Support
- Step 5: Configure User Isolation
- Step 6: Configure Directory Browsing Options
- Step 7: Configure Logon Attempt Restrictions
- Step 8: Configure Request Filtering
- Step 9: Configure FTP Logging
- Step 10: Configure FTP Messages

Prerequisites

To get the most from this tutorial, you must have access to a computer that runs one of the following operating systems:

- Windows Server® 2012
- Windows® 8

Step 1: Install FTP on an Existing IIS Web Server

This step shows you how to install the FTP service on an existing IIS web server that runs on either Windows Server 2012 or Windows 8.

To install FTP on Windows Server 2012

1. On the **Start** screen, click the **Server Manager** tile, then click **Yes**.
2. In the **Server Manager Dashboard**, click **Add roles and features**.
3. If the **Before you begin** page of **Add Roles and Features Wizard** is displayed, click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation**, and click **Next**.
5. On the **Select destination server** page, select **Select a server from the server pool**, select your server from the **Server Pool** list, and then click **Next**.
6. On the **Select server roles** page, expand the **Web Server (IIS)** node, and then expand the **FTP Server** node.
7. Select the **FTP Server** check box and the **FTP Service** check box, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.

To install FTP on Windows 8

1. On the **Start** screen, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. Click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, expand the **Internet Information Services** node.
4. Expand the **FTP Server** node.
5. Select the **FTP Server** check box and the **FTP Service** check box, and then click **OK**.

Step 2: Add an FTP Site

Once the FTP service is installed on your IIS web server, you can add one or more FTP sites. Add an FTP site when you want to enable clients to transfer files to and from a site by using the File Transfer Protocol (FTP).

Note

Because FTP settings are contained in the sites section, changing any FTP setting also forces website application recycling. If you want to avoid this side effect, add a site that is configured exclusively for FTP, instead of for both HTTP and FTP.

To add an FTP site

1. Open IIS Manager.

2. In the **Connections** pane, expand the server node and click the **Sites** node.
3. In the **Actions** pane, click **Add FTP Site** to open the **Add FTP Site** wizard.
4. On the **Site Information** page, in the **FTP site name** box, type a unique friendly name for the FTP site.
5. In the **Physical path** box, type the physical path or click the browse button (**...**) to locate the physical path of the content directory.
6. Click **Next** to open the **Binding and SSL Settings** page.
7. Under **Binding**, in the **IP Address** list, select or type an IP address if you do not want the IP address to remain **All Unassigned**.
8. In the **Port** box, type the port number.
9. Optionally, in the **Virtual Host** box, type a host name if you want to host multiple FTP sites on a single IP address. For example, type **www.contoso.com**.
10. Clear the **Start FTP site automatically** box if you want to start the site manually.
11. Under **SSL**, from the **SSL Certificate** list, select a certificate. Optionally, click **View** to open the **Certificates** dialog box and verify information about the selected certificate.
12. Select one of the following options:
    - **Allow SSL**: Allows the FTP server to support both non-SSL and SSL connections with a client.
    - **Require SSL**: Requires SSL encryption for communication between the FTP server and a client.
13. Click **Next** to open the **Authentication and Authorization Information** page.
14. Under **Authentication**, select the authentication method or methods that you want to use:
    - **Anonymous**: Allows any user to access content providing only the user name **anonymous** or **ftp**. (Most, but not all, FTP clients enter user name for you automatically.)
    - **Basic**: Requires users to provide a valid user name and password to access content. Because Basic authentication transmits unencrypted passwords across the network, use this authentication method only when you know that the connection between the client and FTP server is secure, such as by using Secure Sockets Layer (SSL).
15. Under **Authorization**, from the **Allow access to** list, select one of the following options:
    - **All Users**: All users, whether they are anonymous or identified, can access the content.
    - **Anonymous Users**: Anonymous users can access the content.
    - **Specified Roles or User Groups**: Only members of certain roles or user groups can access the content. Type the role or user group in the corresponding box.
    - **Specified Users**: Only specified users can access the content. Type the user name in the corresponding box.
16. If you selected an option from the **Allow access to** list, select one or both of the following permissions:
    - **Read**: Permits authorized users to read content from the directory.

- **Write**: Permits authorized users to write to the directory.

17. Click **Finish**.

Step 3: Configure FTP Site Defaults

Change an FTP site default value when you want new FTP sites to use a different default value.

Note

When you change a default value, existing sites are not overridden with the new value. Change the value for any existing sites manually.

To configure FTP site defaults

1. Open IIS Manager.
2. In the **Connections** pane, click the server node and the **Sites** node.
3. In the **Actions** pane, click **FTP Site Defaults**.
4. Edit the settings, and then click **OK**.

The following table lists the settings available for configuration in the **FTP Site Defaults** dialog box.

Setting Category	Setting Name	Description
General	Allow UTF-8	Specifies whether to use UTF8 encoding. Default is true.
General	Start Automatically	If true, the FTP site is started upon creation or when the FTP service is started. Default is true.
Connections	Control Channel Time-out	Specifies the time-out (in seconds) when a connection times out due to inactivity.
Connections	Data Channel Time-out	Specifies the time-out (in seconds) when the data channel times out due to inactivity.
Connections	Disable Socket Pooling	Specifies whether socket pooling is used for sites distinguished by IP address rather than port number or host name.
Connections	Max Connections	Specifies the maximum

Setting Category	Setting Name	Description
		simultaneous connections to a server.
Connections	Reset On Max Connections	Specifies whether to disconnect FTP session when sending max connections response.
Connections	Server Listen Backlog	Specifies the number of outstanding sockets that can be queued.
Connections	Unauthenticated Time-out	Specifies the timeout (in seconds) between when a new connection is made and authentication succeeds.
Credential Caching	Enabled	Specifies whether credential caching is enabled for the FTP service.
Credential Caching	Flush Interval	Specifies the cache lifetime, in seconds, for the credentials that are stored in the cache.
File Handling	Allow Reading Files While Uploading	Specifies whether files can be read while being transferred to the server.
File Handling	Allow Replace on Rename	Specifies whether files can overwrite other files when renamed.
File Handling	Keep Partial Uploads	Specifies whether to keep files that have been partially uploaded.

Step 4: Configure Firewall Support

Use the **FTP Firewall Support** feature to configure the following settings that enable the FTP server to accept passive data connections from a firewall:

- **Data Channel Port Range**: Specify a range of ports for passive data connections. You must also open that same range of ports on your firewall. This option can only be configured at the server-level, and the port range of “0-0” means to use the ephemeral port range of the server.

An ephemeral port is a short-lived transport protocol port that TCP/IP allocates from a predefined range.

- **External IP Address of Firewall**: Specify the external IP address of your firewall so that clients know which IP address to use when they communicate with the FTP server through the firewall.

To configure firewall support

1. Open IIS Manager.
2. In the **Connections** pane, select the server node.
3. In **Features View**, double-click **FTP Firewall Support**.
4. In the **Data Channel Port Range** box, type a range of port numbers (separated by a hyphen). For example, type **5000-6000**. Or type **0-0** to use the default port range specified in Windows TCP/IP settings.

   Note
   Do not use ports 0-1024 because these ports are reserved ports.
5. In the **External IP Address of Firewall** box, type the IP address of your firewall.
6. In the **Actions** pane, click **Apply**.

Step 5: Configure User Isolation

The user isolation feature allows you to configure your FTP server to isolate users, which prevents users from accessing the directories of other users on the same FTP site. If you choose not to isolate users, they share a common directory structure.

For example, you can choose not to isolate users on a site that offers only download capabilities for shared content or for a site that does not require the protection of data between users.

If you want to isolate users on your site, you can choose one of the following isolation options:

1. **User name directory**: Isolates user sessions to the physical or virtual directory with the same name of the FTP user account. The user sees only their FTP root location and is, therefore, restricted from navigating higher up the physical or virtual directory tree. Any global virtual directories that are created are ignored.
2. **User name physical directory**: Isolates user sessions to the physical directory with the same name of the FTP user account. The user sees only their FTP root location and is, therefore, restricted from navigating higher up the physical directory tree. Any global virtual directories that are created apply to all users.
3. **FTP home directory configured in Active Directory**: Isolates user sessions to the home directory that is configured in the Active Directory account settings for each FTP user.

To configure user isolation

1. Open IIS Manager.
2. In **Features View**, double-click **FTP User Isolation**.

3. If you don't want to isolate users, under **Do not isolate users. Start users in**, select one of the following options:
   - **FTP root directory**: specifies that all FTP sessions start in the root directory for the FTP site. This option disables all user isolation and starting folder logic.
   - **User name directory**: specifies that all FTP sessions start in the physical or virtual directory with the same name of the currently logged-on user if the folder exists; otherwise, the FTP session starts in the root directory for the FTP site.
4. If you want to isolate users, under **Isolate users. Restrict users to the following directory**, select one of the following options:

   User name directory (disable global virtual directories): isolates user sessions to the physical or virtual directory with the same name of the FTP user account.

   User name physical directory (enable global virtual directories): isolates user sessions to the physical directory with the same name of the FTP user account.

   FTP home directory configured in Active Directory: isolates user sessions to the home directory that is configured in the Active Directory account settings for each FTP user.

   Custom: This option is an advanced feature, and enables developers to create custom providers that provide home directory lookups based on their unique business needs.
5. If you selected **FTP home directory configured in Active Directory** in the previous step, click the **Set** button, and then type a user name and password in the **User name** and **Password** boxes of the **Set Credentials** dialog box that has access to your Active Directory server. Enter the password again in the **Confirm Password** box, then click **OK**.
6. In the **Actions** pane, click **Apply**.

Step 6: Configure Directory Browsing Options

The directory browsing feature gives you control over what is displayed when users browse FTP directories.

To configure directory browsing options

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Features View**, double-click **FTP Directory Browsing**.
4. On the **FTP Directory Browsing** page, under **Directory Listing Style**, select one of the following options:
   - **MS-DOS**: Displays directory content in a manner consistent with MS-DOS.
   - **UNIX**: Displays directory content in a manner consistent with UNIX.
5. Under **Directory Listing Options**, select the information you want to display in directory listings. You can select any of the following options:
   a. **Virtual directories**: Shows virtual directories.
   b. **Available bytes**: Displays file size in bytes.

c. **Four-digit years**: Displays years by using four digits rather than two.

6. In the **Actions** pane, click **Apply**.

Step 7: Configure Logon Attempt Restrictions

This feature enables you to configure the maximum number of failed sign-on attempts that you allow within a specified time before the IP address is denied.

To configure logon attempt restrictions

1. Open IIS Manager.
2. In the **Connections** pane, select the server level.
3. In **Features View**, double click **FTP Logon Attempt Restrictions**.
4. On the **FTP Logon Attempt Restrictions** page, select the **Enable FTP Logon Attempt Restrictions** check box.
5. In the **Maximum number of failed login attempts** box, type a positive integer. The default value is 4.
6. In the **Time period (in seconds)** box, type a positive integer. The default is 30.
7. Select either **Deny IP addresses based on the number of failed login attempts** or **Write to the log only**. If you choose **Write to the log only**, IIS will not restrict clients even if the maximum number of failed attempts is exceeded.
8. In the **Actions** page, click **Apply**.

Step 8: Configure Request Filtering

Use the **FTP Request Filtering** feature page to define the request filtering settings for your FTP site. FTP request filtering is a security feature that allows internet service providers (ISPs) and application service providers to restrict protocol and content behavior.

To configure request filtering

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Feature View**, double click the **FTP Request Filtering** icon.
4. To change the general request filtering settings, click **Edit Feature Settings** in the **Actions** pane. Then edit the general settings as required.
5. To add a filter based on file name extensions, select the **File Name Extension** tab. In the **Actions** pane, click either **Allow File Name Extension** or **Deny File Name Extensions**. Then type the extension in the **File name extension** box, and click **OK**.
6. To add a filter based on a URL segment (for example, a folder name), select the **Hidden Segments** tab. In the Actions pane, click **Add Hidden Segment**. Then type the segment in the **Hidden segment** box, and click **OK**.
7. To add a filter based on a URL sequence, select the **Denied URL Sequences** tab. In the

Actions pane, click **Add URL Sequence**. Then type the sequence in the **URL sequence** box, and click **OK**.

8. To add a filter based on a command, select the **Commands** tab. In the **Actions** pane, click either **Allow Command** or **Deny Command**. Then type the command in the **Command** box, and click **OK**.

   Warning
   The ability to add a command-based filter is an advanced feature. If you use this feature incorrectly, you could deny access to all FTP clients on your server.

Step 9: Configure FTP Logging

You can use the **FTP Logging** feature to configure logging features at the server or site level, and to configure logging settings.

To configure FTP logging

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Feature View**, double click the **FTP Logging** icon.
4. In the **One log file per** menu, select either **Site** or **Server**.
5. Under **Log File**, click **Select W3C Fields**, and then select the information you want go log.
6. Under **Directory**, either type the path to the base folder where you want the FTP log files stored, or click **Browse** to browse to the base folder.
7. Under **Encoding**, select either **UTF8** (single-byte and multi-byte characters) or **ANSI** (single-byte characters only).
8. Under **Log File Rollover**, select how you want FTP to create new log files from the following list:
   a. **Schedule**: select **Hourly**, **Daily**, **Weekly**, or **Monthly** to create new log files on a fixed time interval.
   b. **Maximum file size (in bytes)**: enter a positive integer to create new log files when the file size exceeds that number of bytes.
   c. **Do not create new log files**
9. Select the **Use local time for file naming and rollover** check box when you want log file naming and rollover to be based on the local time zone instead of Coordinated Universal Time (UTC).
10. In the **Actions** pane, click **Apply**.

Step 10: Configure FTP Messages

Use the **FTP Messages** feature to modify the settings for messages sent when a user connects to your FTP site.

To configure FTP messages

1. In the **Connections** pane, select either the server level or the site level.
2. In **Features View**, double-click **FTP Messages**.
3. On the **FTP Messages** page, under **Message Behavior**, select how you want your FTP messages to behave. You can select any of the following options:
   a. **Suppress default banner**: Specifies whether to display the default identification banner for the FTP server.
   b. **Support user variables in messages**: Specifies whether to display a specific set of user variables in FTP messages. The following user variables are supported:
      - %BytesReceived% - The number of bytes sent from the server to the client for the current session.
      - %BytesSent% - The number of bytes sent from the client to the server for the current session.
      - %SessionID% - The unique identifier for the current session.
      - %SiteName% - The name of the FTP site that is hosting the current session.
      - %UserName% - The account name of the currently logged-on user.
   c. **Show detailed messages for local requests**: Specifies whether to display detailed error messages when the FTP client is connecting to the FTP server on the server itself (local host).
4. Under Message Text, type messages in the following boxes:
   a. **Banner**: Specifies the message the FTP server displays when FTP clients first connect to the FTP server.
   b. **Welcome**: Specifies the message the FTP server displays when FTP clients have logged in to the FTP server.
   c. **Exit**: Specifies the message the FTP server displays when FTP clients log off the FTP server.
   d. **Maximum Connections**: Specifies the message the FTP server displays when clients try to connect and cannot because the FTP service has reached the maximum number of client connections allowed.
5. In the **Actions** pane, click **Apply**.

See Also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build a Static Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Build a PHP Website on IIS

This document contains an overview of the **Build a PHP website on IIS** scenario. It also contains links to additional information and community resources related to the scenario.

Did you mean…

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Scenario Description

This scenario shows how to plan and configure a PHP website on an IIS 8 webserver. It is divided into two phases: a plan and design phase, and an install and configure phase. In the plan and design phase, you are provided the information needed to make informed decisions about web server installation, PHP settings, and basic application security. In the install and configure phase, you are guided through the procedures required to install IIS, add a PHP application, and configure IIS and PHP.

This scenario does not cover how to write a PHP application.

In This Scenario

- Plan a PHP Website on IIS
  - Step 1: Plan IIS Web Server and PHP Installation
  - Step 2: Plan PHP Settings
  - Step 3: Plan PHP Application Security
- Configure a PHP Website on IIS
  - Step 1: Install IIS and PHP
  - Step 2: Configure PHP Settings
  - Step 3: Configure PHP Application Security

Practical Applications

Whether you are an IT professional, a web developer, or you just want to set up your own webserver, this scenario can help you install IIS and configure it to serve your PHP web application.

Software Requirements

To get the most from this scenario, you must have access to a computer running one of the following operating systems:

- Windows Server® 2012
- Windows® 8

See Also

The following table contains links to resources related to this scenario.

Content type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center
Tools and Settings	**N:Microsoft.Iis.PowerShell.Provider** namespace
Community Resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related Technologies	ASP.NET | ASP.NET Web Projects

Plan a PHP Website on IIS

To develop a plan for installing an IIS web server and configuring it for PHP web applications, follow the steps listed.

- Step 1: Plan IIS Web Server and PHP Installation
- Step 2: Plan PHP Settings
- Step 3: Plan PHP Application Security

After you have reviewed these planning steps, see Configure a PHP Website on IIS. For more information, see Build a PHP Website on IIS.

Step 1: Plan IIS Web Server and PHP Installation

The first step in planning to build a PHP website on IIS 8 is to determine how to install IIS along with the CGI module, where to download and install the appropriate PHP version, and how to add your application files to IIS.

The following list shows the tasks required to complete this step:

1.1. Plan to Install IIS

1.2. Plan to Download and Install PHP

1.3. Plan to Add a PHP Application

When you are done with these tasks, record your design decisions before going on to Step 2: Plan PHP Settings.

1.1. Plan to Install IIS

To support PHP applications on an IIS web server, install the default configuration of IIS and the CGI module. The CGI module contains both CGI (Common Gateway Interface protocol) and FastCGI. FastCGI is a high-performance extension to CGI and is recommended for PHP applications. Once your machine has this IIS configuration installed, you can download and install PHP.

Important

To install IIS 8 on either Windows Server® 2012 or Windows® 8, you must sign on as a member of the Administrators group.

1.2. Plan to Download and Install PHP

The simplest way to install PHP is to use Web Platform Installer (Web PI). Web PI installs the latest version of PHP along with any dependencies that version requires. To learn more about the Web PI, see Learn More and Install the Web PI.

To add PHP support manually to your IIS web server, download and install the PHP runtime version that you require. You can download Windows versions of PHP from the PHP for Windows site.

There are two builds for each PHP version: one is thread-safe, and the other is non-thread-safe. The thread-safe version is designed for environments where the web server keeps the PHP engine in memory, running multiple treads of execution simultaneously. The architecture of IIS and the FastCGI extension provide an isolation model that keeps requests separate, removing the need for a thread-safe version of PHP. The non-thread-safe version of PHP gives better performance on IIS with FastCGI install than does the tread-safe version.

PHP offers several extensions that extend its base functionality. One extension is important for PHP web applications that run on IIS—the Windows Cache Extension for PHP (WinCache).

WinCache is a PHP accelerator that is used to increase the speed of PHP applications that run on Windows and Windows Server. Once the WinCache is enabled and loaded by the PHP engine, PHP applications can take advantage of the functionality without any code modifications.

You can download WinCache extension from the PHP for Windows site.

Warning
The current version of WinCache (php_winchache-svn20110402-5.2-nts-vc6-x86.zip) does not work with PHP 5.4.0. Therefore, we recommend that you install the non-thread-safe version of PHP 5.3.6 (or earlier).

1.3. Plan to Add a PHP Application

An application is a grouping of content at the root level of a website or a grouping of content in a separate folder under the website root directory. When you add an application in IIS 8, you designate a directory as the *application root*, or starting point, for the application. Then specify properties specific to that particular application, such as the application pool that the application runs in.

To configure your PHP application on an IIS website, provide the following information:

- **Alias**: The alias is used as part of the application root URL and should be short and descriptive. For example, the alias *marketing* added to the default website on the local host computer, would produce the following URL: `//localhost/marketing`.
- **Application pool**: An application pool enables an application or a group of applications to run in isolation from one or more applications in another application pool.
- **Physical path**: The local path to the application files on the server.

See Also

- Plan a PHP Website on IIS
- Step 1: Install IIS and PHP
- Step 2: Plan PHP Settings
- Step 3: Plan PHP Application Security
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 2: Plan PHP Settings

In this phase of building your PHP website, plan to set up the WinCache PHP extension, determine what configuration settings you require for PHP, and what PHP extensions your application requires.

The following list shows the tasks required to complete this step:

2.1. Plan WinCache Configuration

2.2. Plan Other PHP Settings

2.3 Plan PHP Extensions

When you are done with these tasks, record your design decisions before going on to Step 3: Plan PHP Application Security.

2.1. Plan WinCache Configuration

WinCache is a PHP extension that accelerates PHP applications, which run on Windows and Windows Server. For the most installations, the WinCache extension must be added to the php.ini file and no other configuration is required. However, if you want more control, WinCache offers many configuration settings. You can see a full list of setting in the WinCache Runtime Configuration page.

2.2. Plan Other PHP Settings

Configuration and environmental settings for PHP are contained within the Php.ini, which is located in the PHP folder that you create during installation, for example C:\PHP. The following sections contain a number of settings for the Php.ini file that help PHP work better with Windows.

Required Settings

extension_dir = <PATH TO EXTENSIONS>

The extension_dir points to the directory where the PHP extensions are stored. The path can be fully qualified (for example, C:\PHP\ext) or relative (for example, .\ext). Extensions that are specified lower in the Php.ini file are located in the extension_dir. If the extensions specified are not in the extension_dir, then PHP displays a warning message at the start of script execution, and the application may show errors because of the missing functionality.

extension = xxxxxx.dll

For each extension enabled, a corresponding extension= directive that tells PHP which extensions in the extension_dir to load at startup time is necessary.

log_errors = On

PHP errors can also go through the PHP error logging facility. This can be used to send errors to a file or to a service (for example, syslog) and works with the error_log directive described below. When running under IIS, log_errors must be enabled with a valid error_log. Otherwise, FastCGI considers any startup messages (that may be benign) as an error condition, which generates an HTTP 500 return error code to the browser.

error_log = <path_to_error_log_file>

The error_log must specify the fully qualified, or relative, path to the file where the PHP error log is stored. This file must be writable for the IIS service. The most common places for this file are in various temporary directories (for example, C:\inetpub\temp\php-errors.log). That puts the log in a place that IIS can use and keeps the log close to where PHP applications are running.

cgi.force_redirect = 0

This directive is required for running under IIS. It is a directory security facility that is required by many other web servers. However, enabling it under IIS causes the PHP engine to fail on Windows.

cgi.fix_pathinfo = 1

This setting lets PHP access real path information per the CGI specification. The IIS FastCGI implementation needs this extension set.

fastcgi.impersonate = 1

FastCGI under IIS supports the ability to impersonate security tokens of the calling client. This setting allows IIS to define the security context that the request runs under.

fastcgi.logging = 0

FastCGI logging should be disabled when you use PHP with IIS. If it is left enabled, then any messages of any class are treated by FastCGI as error conditions, which causes IIS to generate an HTTP 500 exception.

Optional Settings

max_execution_time = ##

This directive sets the maximum amount of time that can be taken executing any given script. The default is 30 seconds. Some applications need more time to process batch operations (for example, Gallery2 loading multiple images from a remote location). However, setting the execution time higher than 300 seconds is not advised because there are often other parts of the connection that cannot support such a long execution time.

memory_limit = ###M

The amount of memory available for the PHP process (in MB). The default is 128 MB,

which is sufficient for most PHP applications.

display_errors = Off

This directive determines whether to include any error messages in the stream that it returns to the web server. If turned on, PHP sends the classes of errors that are defined with the error_reporting directive back to IIS as part of the error stream. Many of the open-source applications bypass error reporting by executing commands prefaced with @. This setting allows the applications to control error handling.

Mail functions

PHP is configured by default to send outbound mail through an SMTP server that is located on the same system as the web server. Most Windows installations usually have the web and mail servers on separate systems.

2.3 Plan PHP Extensions

It is important to determine which extensions are required by the applications you intend to run, and then to limit the installed extensions to those extensions only. For example, a typical open-source PHP application might require the following types of extensions:

- **Database Extensions**—Open-source applications that use MySQL for a database engine use either the php_mysql or the php_mysqli extension. For new development work, either of these extensions work well. Consider using the PDO versions of the MySQL driver. PDO is a PHP extension that provides a data-access abstraction layer that can be used with various databases. Microsoft provides PHP drivers for SQL server that you can download. To learn more about PHP drivers for SQL Server, see Microsoft Drivers for PHP for SQL Server.
- **Image Handling Extensions**—Open-source applications that enable work with images use the GD2 extension – php_gd2, which has several application programming interfaces (APIs) for basic image manipulation. Some applications use the ImageMagick application and libraries. There is also a php_exif library for working with the extended information that modern digital cameras store within the images.
- **Internationalization and Localization Extensions**—The two most commonly used extensions for i18n and l10n are php_mbstring (Multi-Byte String) and php_gettext (Native Language Support). Many of the open-source applications use one or both of these extensions.
- **Web Services Extensions**—Choose the web services extensions based on the services desired. For PHP, the SOAP extension is widely used. The XML-RPC extension is often used with SOAP and other services.

To view the entire list of PHP extensions available for applications that run on Windows platforms, see the complete list of Windows extensions for PHP.

See Also

- Plan a PHP Website on IIS
- Step 2: Configure PHP Settings
- Step 1: Plan IIS Web Server and PHP Installation
- Step 3: Plan PHP Application Security
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 3: Plan PHP Application Security

In this phase of Building a PHP website on IIS, consider what PHP configuration settings, web server settings, and PHP application settings you need to strengthen security.

3.1. PHP Configuration Settings for Security

3.2. Web Server and PHP Application Security

When you are done with these tasks, record your design decisions before going on to Step 1: Install IIS and PHP.

3.1. PHP Configuration Settings for Security

The dynamic capabilities of PHP also make it a potential security risk because data is actively fetched, received, and processed from anywhere on the Internet. Attackers may attempt to send in malicious data and scripts and trick your server into fetching malicious scripts and running them. Attackers may also attempt to read and write files on your server to take control of the web site and use it for their own purposes.

You can configure PHP settings to tighten the security of a PHP installation and help protect the website from malicious attacks. The Php.ini file specifies the configuration settings PHP uses when it is running on your website. The Php.ini file determines what PHP scripts are allowed to do and prohibited from doing.

This section describes the configuration settings that help to protect your PHP applications.

Disable File Handling for Remote URLs

```
allow_url_fopen = Off
 allow_url_include = Off
```

This setting is very important because it prevents URLs from being used in statements such as include(). Setting **allow_url_fopen** to “Off” means that only files that reside within your website can be included. You can’t include a file from a different server, but neither can other people through Remote File Inclusion (RFI) attacks. In an RFI attack, someone embeds a URL in an HTTP request hoping that your script is tricked into running theirs. A command such as `include("http://website.com/page.php")`, for example, is not allowed to execute.

Include a file from your own site by specifying its path and filename. For example, if you have a URL **include** line, convert it to: `include($_SERVER['DOCUMENT_ROOT'] . '/page.php');`

The code segment `$_SERVER['DOCUMENT_ROOT']` is a superglobal variable set to the root folder of your site. (Note that there is no trailing "/"; you must provide a leading "/" in '/page.php'.)

If you want to include static content from another one of your websites, such as `include('http://myothersite.com/includes/footer.php')`, make a copy of that content in the current site and then include it locally.

Disable Register_Globals

```
register_globals = Off
```

This setting makes it difficult for an attacker to inject code into your script. For example, in `URL http://site.com/index.php?variable=value`, the variable passes into your script with its value set to **value** when **register_globals** is **On**. When **register_globals** is **Off**, however, variables do not automatically pass into your script's variable list.

Restrict File System Read/Write

```
open_basedir = "c:\inetpub\"
```

This setting restricts PHP scripts from accessing files outside the specified base directory.

Disable Safe Mode

```
safe_mode = Off
 safe_mode_gid = Off
```

This setting restricts the permissions with which PHP scripts run. Some third-party scripts do not run properly when safe_mode is set to "On." Note that beginning with PHP 6 safe_mode does not exist.

Limit Script Execution Time

```
max_execution_time = 30
 max_input_time = 60
```

These settings control the number of seconds a script is allowed to run and parse user input. These settings help prevent poorly written scripts from typing up the server.

Limit Memory Usage and File Size

```
memory_limit = 16M
 upload_max_filesize = 2M
 post_max_size = 8M
 max_input_nesting_levels = 64
```

These settings help manage memory and input/output efficiently. In addition, they prevent poorly written scripts from hogging memory and handling huge files.

Configure Error Logging

```
display_errors = Off
display_startup_errors = Off log_errors = On
error_log = "C:\path\of\your\choice"
```

These settings specify that all errors and warnings get logged to your error log text file and specify that none of the errors or warnings get displayed on any web page that is sent out from your server. Errors should not be displayed publicly because they can help someone figure out how to attack your server. Always check your error log when you are testing new code.

Enable FastCGI Impersonation

```
fastcgi.impersonate = 1
```

This setting enables IIS to impersonate the security tokens of calling client and define the security context that is used to run the request.

Disable FastCGI Logging

```
fastcgi.logging = 0
```

The FastCGI module will fail the request when PHP sends any data on **stderr** by using FastCGI protocol. Disabling FastCGI logging will prevent PHP from sending error information over **stderr**, and generating 500 response codes for the client.

Hide PHP Presence

```
expose_php = Off
```

With this setting, the headers that accompany outgoing pages do not reveal that PHP is running or its version.

3.2. Web Server and PHP Application Security

Isolate Web Applications

One of the most effective ways to improve security for your web application is to isolate it from other applications on your web server. An application pool has its own worker process, which processes requests and runs application code. The worker process has a security identifier (SID). And each application pool has a unique application-pool identity. By default, when you create a web application, a new application pool is also created with the same name as the application. If you keep web applications in separate application pools, you can isolate them from one another.

Web application isolation entails the following:

- Site isolation: Separate different applications into different sites with different application pools.

- Least privilege: Run your worker process as a low privileged identity (virtual application pool identity) that is unique per site.
- Temp isolation: Set up a separate PHP temp folder per site and only give access to appropriate process identity.
- Content isolation: Make sure to set an ACL (access control list) on each site root to allow only access to the appropriate process identity.

Tip
It is a good idea to host your website and web application content on a drive other than your system drive (C:).

Enable Per-site PHP Configuration

The FastCGI handler makes it possible to use a different Php.ini file for every application mapping. You can customize your PHP configuration around the specific requirements of your users or your PHP applications, letting you tighten configuration.

Use Request Filtering

The request filtering module inspects known malicious patterns in the requests and prevents such requests from being serviced if the module determines that the requests may be harmful. For example, this module lets you filter requests that are double escaped, filter requests that use certain HTTP verbs, or block requests to specific folders. You can enforce tighter security policies on your web servers with the request filtering module. For more information about request filtering, see Configure Request Filtering in IIS.

See Also

- Plan a PHP Website on IIS
- Step 1: Install IIS and PHP
- Step 1: Plan IIS Web Server and PHP Installation
- Step 2: Plan PHP Settings
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Configure a PHP Website on IIS

To install an IIS web server and configure it for PHP web applications, follow the steps listed.

- Step 1: Install IIS and PHP
- Step 2: Configure PHP Settings
- Step 3: Configure PHP Application Security

For planning information to review before deployment, see Plan a PHP Website on IIS. For more information, see Build a PHP Website on IIS.

Step 1: Install IIS and PHP

In this step of building a PHP website, you install IIS and FastCGI, download and install PHP and the WinCache extension, and upload your PHP application.

The following list shows the tasks required to complete this step:

1.1. Install IIS

1.2 Install PHP by using Web PI

1.3. Download and Install PHP Manually

1.4. Add Your PHP Application

When you are done, make sure that IIS and the PHP are installed, and your PHP application has been added to your website. Then go on to Step 2: Configure PHP Settings.

1.1. Install IIS

You can use the Web Platform Installer (Web PI) to install IIS, and applications that run on IIS. Web PI installs the latest versions of available Web Platform offerings with just a few simple clicks. Using Web PI, you can download and install any new tools or updates, including PHP. To learn more about the Web PI, see Learn more and install the Web PI.

If you do not use Web PI to install IIS, you can install IIS manually. To install IIS manually, use the following steps:

To install IIS on Windows Server 2012

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select your server, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, note the preselected features that are installed by default, and then select **CGI**. This selection also installs FastCGI, which is recommended for PHP

applications.

8. Click **Next**.
9. On the **Web Server Role (IIS)** page, click **Next**.
10. On the **Select Role Services** page, note the preselected role services that are installed by default, and then click **Next**.

    Note
    You only have to install the IIS 8 default role services for a static-content web server.

11. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
12. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
13. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    You should see the default IIS Welcome page.

To install IIS on Windows 8

1. On the **Start** page, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services**, note the preselected features that are installed by default, and then select **CGI**. This selection also installs FastCGI, which is recommended for PHP applications.
4. Click **OK**.
5. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    You see the default IIS Welcome page.

1.2 Install PHP by using Web PI

The preferred method to install PHP on a Windows or Windows Server computer is to use Web Platform Installer (Web PI).

To install PHP manually

1. Open a browser to the following website: Microsoft Web Platform Installer 3.0.
2. Click **Download It Now**, and then click **Run**.
3. At the top of the **Web Platform Installer** window, click **Products**.
4. Click **Frameworks**, and then select the current version of PHP. (At this writing, the current version is PHP 5.3.13.)

5. Click **Install**.

   The **Web Platform Installation** page displays the version of PHP and its dependencies that will be installed.
6. Click **I Accept**.

   Web PI installs the PHP packages.
7. Click **Finish**.

1.3. Download and Install PHP Manually

If you decide to download PHP and install it manually, the procedures in this section guide you the following tasks:

- Download PHP and the WinCache extension.
- Install PHP and WinCache.
- Add the PHP installation folder to the Path environment variable.
- Set up a handler mapping for PHP.
- Add default document entries for PHP.
- Test your PHP installation.

To keep this procedure simple, install the WinCache extension but do not configure it. You will configure and test WinCache in Step 2: Configure PHP Settings.

To download and install PHP and WinCache

1. Open your browser to Windows for PHP Download Page and download the PHP 5.3 non-thread-safe zip package.

   Caution
   The PHP 5.4 version does not work with the WinCache extension version 1.1. Use PHP 5.3 until this problem is resolved.
2. Download the WinCache extension (Php_wincache-svn20110402-5.2-nts-vc6-x86.zip) from the List of Windows Extensions for PHP.
3. Extract all files in the PHP .zip package to a folder of your choice, for example C:\PHP\.
4. Extract the WinCache .zip package to the PHP extensions folder (\ext), for example C:\PHP\ext.

   The WinCache .zip package contains one file (Php_wincache.dll).
5. Open **Control Panel**, click **System and Security**, click **System**, and then click **Advanced system settings**.
6. In the **System Properties** window, select the **Advanced** tab, and then click **Environment Variables**.
7. Under **System variables**, select **Path**, and then click **Edit**.
8. Add the path to your PHP installation folder to the end of the **Variable value**, for example **;C:\PHP**. Click **OK**.

9. Open IIS Manager, select the hostname of your computer in the **Connections** panel, and then double-click **Handler Mappings**.
10. In the **Action** panel, click **Add Module Mapping**.
11. In **Request path**, type ***.php**.
12. From the **Module** menu, select FastCgiModule.
13. In the **Executable** box, type the full path to Php-cgi.exe, for example **C:\PHP\Php-cgi.exe**.
14. In **Name**, type a name for the module mapping, for example **FastCGI**.
15. Click **OK**.
16. Select the hostname of your computer in the **Connections** panel, and double-click **Default Document**.
17. In the **Action** panel, click **Add**. Type **Index.php** in the **Name** box, and then click **OK**.
18. Click **Add** again. Type **Default.php** in the **Name** box, and then click **OK**.

To test your PHP installation

1. Open a text editor, for example Notepad, as Administrator.
2. In a new file, type the following text: `<?php phpinfo(); ?>`
3. Save the file as **C:\inetpub\wwwroot\Phpinfo.php**.
4. Open a browser and enter the following URL: **http://localhost/phpinfo.php**

   A nicely formatted webpage is displayed showing the current PHP settings.

1.4. Add Your PHP Application

Once you have IIS and PHP installed, you can add a PHP application to your web server. This section describes how to set up your PHP application on an IIS web server with PHP installed. It does not explain how to develop a PHP application.

To add a PHP web application

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. On the **Server Manager Dashboard**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, right-click the **Sites** node in the tree, and then click **Add Website**.
3. In the **Add Website** dialog box, type a friendly name for your website in the **Site name** box.

4. If you want to select a different application pool than the one listed in the **Application Pool** box, click **Select**. In the **Select Application Pool** dialog box, select an application pool from the **Application Pool** list and then click **OK**.
5. In the **Physical path** box, type the physical path of the website's folder, or click the browse button (**...**) to navigate the file system to find the folder.
6. If the physical path that you entered in step 5 is to a remote share, click **Connect as** to specify credentials that have permission to access the path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option in the **Connect As** dialog box.
7. Select the protocol for the website from the **Type** list.
8. The default value in the **IP address** box is **All Unassigned**. If you must specify a static IP address for the website, type the IP address in the **IP address** box.
9. Type a port number in the **Port** text box.
10. Optionally, type a host header name for the website in the **Host Header** box.
11. If you do not have to make any changes to the site, and you want the website to be immediately available, select the **Start Web site immediately** check box.
12. Click **OK**.

See Also

- Configure a PHP Website on IIS
- Step 1: Plan IIS Web Server and PHP Installation
- Step 2: Configure PHP Settings
- Step 3: Configure PHP Application Security
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 2: Configure PHP Settings

In this step of building a PHP website on IIS, you configure the WinCache PHP extension, configure PHP settings, and download and configure any other PHP extensions that your application requires.

2.1. Configure WinCache

2.2. Configure Other PHP Settings

2.3 Configure PHP Extensions

When you complete these tasks, continue to Step 3: Configure PHP Application Security.

2.1. Configure WinCache

Before you perform this procedure, you should download in install PHP and WinCache as described in 1.2 Download and Install PHP.. For more information about WinCache, see 2.1. Plan WinCache Configuration.

To configure the WinCache PHP extension

1. In Windows Explorer, open your PHP installation folder, for example **C:\PHP**.
2. Choose either the **php.ini - development** or **php.ini - production** file, and rename it **php.ini**.
3. In a text editor, open the php.ini file and added the following line at the end of the file:
   `extension = php_wincache.dll.`
4. Save and close the php.ini file.
5. Recycle the IIS Application Pools for PHP to pick up the configuration changes.

To view WinCache configuration and other PHP settings

1. Open a text editor.
2. In a new file, type the following text: `<?php phpinfo(); ?>`
3. Save the file as **c:\inetpub\wwwroot\phpinfo.php**.
4. Open a browser and enter the following URL: **http://localhost/phpinfo.php**

   A nicely formatted web page is displayed showing the current PHP settings. The WinCache settings appear in a section called **wincache**.

   Warning
   Delete the phpinfo.php file when it's no longer needed.

2.2. Configure Other PHP Settings

The following procedure explains how to configure PHP settings in the php.ini file. For more information about PHP settings, see 2.2. Plan other PHP settings.

To configure a PHP setting

1. In Windows Explorer, open your PHP installation folder, for example **C:\PHP**.
2. In a text editor, open the php.ini file.
3. Search the file for the setting you want to change.

   If the setting is commented out (line begins with a semicolon [;]), delete the semicolon and set the value. If you can't find the setting, add the line to the end of the file.
4. Save and close the php.ini file.
5. Recycle the IIS Application Pools for PHP to pick up the configuration changes.

2.3 Configure PHP Extensions

The following procedure shows how to download and install a PHP extension. For information about PHP extensions, see 2.3 Plan PHP Extensions.

To configure a PHP extension

1. Download the PHP extension you want from the list of Windows extensions for PHP.
2. Extract the extension zip package to the PHP extensions folder (\ext), for example C:\PHP\ext.
3. In Windows Explorer, open your PHP installation folder, for example C:\PHP.
4. In a text editor, open the php.ini file.
5. Search the file for the extension you want to configure.

   If the extension is commented out (line begins with a semicolon [;]), delete the semicolon. If you can't find the extension, add it to the end of the file. A line that adds an extension is in the form: `extension = extension_name.dll`. For example: `extension = php_soap.dll`.
6. Save and close the php.ini file.
7. Recycle the IIS Application Pools for PHP to pick up the configuration changes.

See Also

- Configure a PHP Website on IIS
- Step 2: Plan PHP Settings
- Step 1: Install IIS and PHP
- Step 3: Configure PHP Application Security
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 3: Configure PHP Application Security

In this last step in building a PHP website on IIS, you configure PHP settings, web server settings, and PHP application settings that improve your website security.

3.1. Configure PHP Settings for Security

3.2. Configure Web Server and PHP Application Security

3.1. Configure PHP Settings for Security

The following procedure shows you how to configure PHP settings in the php.ini file. For information about security-related PHP settings, see 3.1. PHP Configuration Settings for Security.

To configure a PHP setting for security

1. In Windows Explorer, open your PHP installation folder, for example **C:\PHP**.
2. In a text editor, open the php.ini file.
3. Search the file for the setting you want to change.

   If the setting is commented out (line begins with a semicolon [;]), delete the semicolon and set the value. If you can't find the setting, add the line to the end of the file.
4. Save and close the php.ini file.
5. Recycle the IIS Application Pools for PHP to pick up the configuration changes.

3.2. Configure Web Server and PHP Application Security

This section shows how to configure several web server and application settings for IIS. These settings include isolating web applications, enabling per-site PHP configurations, and using request filtering. For more information about web server and PHP application security settings, see 3.2. Web Server and PHP Application Security.

Isolate Web Applications

Implement the following recommendations to isolate websites and web applications on your server.

- Use one application pool per website or web application.
- Limit access to site folders and files to the application pool identity.
- Set up a separate PHP temp folder per site and only give access to the application pool identity.
- Make sure to set an ACL (access control list) on each site root to allow only access to the application pool identity.

If you have more than one application per application pool, consider creating enough application pools and moving some of the applications to the new pools.

To create an application pool

1. Open IIS Manager.
2. In the **Connections** pane, click **Application Pools**.
3. In the **Actions** pane, click **Add Application Pool**.
4. In the **Name** box, type a unique name for the application pool.
5. Under **.NET Framework version**, select **No Managed Code**.
6. Select the **Managed pipeline mode**. The **Integrated** mode is recommended.
7. Click **OK**.

To move an application to another application pool

1. Open IIS Manager.
2. In the **Connections** page, select the website or web application you want to move.
3. In the **Actions** pane, click **Basic Settings**.
4. On the **Edit Site** dialog, click **Select** to open the **Select Application Pool** dialog, and then select the application pool from the **Application pool** menu.
5. Click **OK** to close the **Select Application Pool** dialog, and click **OK** to close the **Edit Site** menu.

To add an application pool identity to a folder or file ACL

1. Open Windows Explorer and navigate to the folder or file.
2. Right click the folder or file, and then click **Properties**.
3. Select the **Security** tab, and then click **Edit**.
4. Click **Add**, click **Locations**, and select your server as the location to search.
5. In the **Enter the object names to select** box, type **IIS APPPOOL\applicationPoolName**, where *applicationPoolName* is the application pool identity.
6. Click **OK**, click **OK**, and click **OK** again to close the dialogs.

Enable Per-site PHP Configuration

When you have multiple PHP applications on an IIS web server, you can improve security by configuring a PHP process pool and a php.ini file for each application. This section explains how to configure process pools and multiple pnp.ini files by using an **applicationHost.config** file.

Per-site PHP Process Pools

When each website has its own application pool, which is a recommended practice on IIS, it is possible to associate a dedicated FastCGI process pool with each website. This association is done in the **fastCgi** section of the **applicationHost.config** file. A FastCGI process pool is uniquely identified by the combination of **fullPath** and **arguments** attributes of the **application** element. To create several FastCGI process pools for the same process executable, such as php-cgi.exe, use the **arguments** attribute to distinguish the process pool definitions. With php-cgi.exe processes, use the command line switch "-d" to define an INI entry for a PHP process. And use this switch to set a PHP setting that makes the arguments string unique.

For example, if there are two Web sites "website1" and "website2" that must have their own set of PHP settings, define the FastCGI process pools as follows:

```
<fastCgi>
  <application fullPath="C:\PHP\php-cgi.exe"
    arguments="-d open_basedir=C:\Websites\Website1" />
  <application fullPath="C:\PHP\php-cgi.exe"
    arguments="-d open_basedir=C:\Websites\Website2" />
```

```
</fastCgi>
```

In this example the PHP setting `open_basedir` is used to distinguish between the process pool definitions. The setting also enforces that the PHP executable for each process pool can perform file operations only within the root folder of the corresponding website.

Therefore, the PHP handler mapping for website1 is as follows:

```
<system.webServer>
  <handlers accessPolicy="Read, Script">
    <add name="PHP via FastCGI"
      path="*.php" verb="*"
      modules="FastCgiModule"
      scriptProcessor="C:\PHP\php-cgi.exe|-d open_basedir=C:\Websites\Website1"
      resourceType="Unspecified"
      requireAccess="Script" />
  </handlers>
</system.webServer>
```

And the handler mapping for website2 is as follows:

```
<system.webServer>
  <handlers accessPolicy="Read, Script">
    <add name="PHP via FastCGI"
      path="*.php" verb="*"
      modules="FastCgiModule"
      scriptProcessor="C:\PHP\php-cgi.exe|-d open_basedir=C:\Websites\Website2"
      resourceType="Unspecified" requireAccess="Script" />
  </handlers>
</system.webServer>
```

Specifying Php.ini Location

When the PHP process starts, it determines the location of the configuration php.ini file by using various settings. The PHP documentation provides a detailed description of the PHP startup process. One of the places where the PHP process searches for the php.ini location is the PHPRC environment variable. If the PHP process finds a php.ini file in the path that is specified in this environment variable, it will use it. Otherwise, the PHP process will revert to using the default location of the php.ini file. This environment variable can be used to allow hosting customers to use their own versions of php.ini files.

For example, if there are two websites, "website1" and "website2," that are located at the following file paths: `C:\WebSites\website1` and `C:\WebSites\website2`, you can configure the php-cgi.exe process pools in the **fastCgi** section of the **applicationHost.config** file as follows:

```
<fastCgi>
  <application fullPath="C:\PHP\php-cgi.exe" arguments="-d
open_basedir=C:\Websites\Website1">
    <environmentVariables>
      <environmentVariable name="PHPRC" value="C:\WebSites\website1" />
    </environmentVariables>
  </application>
  <application fullPath="C:\PHP\php-cgi.exe" arguments="-d
open_basedir=C:\WebSites\Website2">
    <environmentVariables>
      <environmentVariable name="PHPRC" value="C:\WebSites\website2" />
    </environmentVariables>
  </application>
</fastCgi>
```

This way website1 can have its own version of the php.ini file that is located in the `C:\WebSites\website1`, while website2 can have its own version of the php.ini file that is located in `C:\WebSites\website2`. This configuration also ensures that if a php.ini file cannot be found in the location that is specified by the PHPRC environment variable, then PHP will use the default php.ini file that is located in the same folder where the php-cgi.exe is located.

Use Request Filtering

For information about how to configure request filtering, see Configure Request Filtering in IIS.

See Also

- Configure a PHP Website on IIS
- Step 3: Plan PHP Application Security
- Step 1: Install IIS and PHP
- Step 2: Configure PHP Settings
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Build a Web Farm with IIS Servers

This document contains an overview of the **Build a web farm with IIS servers** scenario. It also contains links to additional information and community resources related to the scenario.

Did you mean…

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Web Farm Framework

Scenario Description

This scenario shows how to plan and configure a web farm with servers that run the Windows Server® 2012 operating system.

The scenario is divided into two phases: a plan and design phase, and an install and configure phase. In the plan and design phase, you learn what the concepts needed to configure a web farm that keeps its server configuration and content synchronized, balances server load, and stores SSL certificate centrally. In the install and configure phase, you are guided through the procedures required to install IIS web server, configure a back-end content server, configure load balancing, and configure a central certificate store for SLL.

In This Scenario

- Plan a Web Farm with IIS Servers
  - Step 1: Plan IIS Web Farm Infrastructure
  - Step 2: Plan IIS Web Farm Configuration
  - Step 3: Plan IIS Web Farm Load Balancing
  - Step 4: Plan SSL Central Certificate Store
  - Step 5: Plan Application Deployment
- Configure a Web Farm with IIS Servers
  - Step 1: Install IIS Web Farm Infrastructure
  - Step 2: Configure IIS Web Farm Servers
  - Step 3: Configure IIS Web Farm Load Balancing
  - Step 4: Configure SSL Central Certificate Store
  - Step 5: Configure Application Deployment

Practical Applications

Whether you are an IT professional, a web developer, or you just want to set up your own web farm, this scenario can help you install and configure a fully functional web farm with IIS web servers.

Software Requirements

To get the most from this scenario, you must have administrator access to at least five computers (or virtual machines) that run the Windows Server 2012 operating system.

See Also

The following table contains links to resources related to this scenario.

Content type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center
Tools and settings	**N:Microsoft.Iis.PowerShell.Provider** namespace
Community resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related technologies	ASP.NET | ASP.NET Web Projects

Plan a Web Farm with IIS Servers

To develop a plan for installing and configuring an IIS web farm, follow the steps listed.

- Step 1: Plan IIS Web Farm Infrastructure
- Step 2: Plan IIS Web Farm Configuration
- Step 3: Plan IIS Web Farm Load Balancing
- Step 4: Plan SSL Central Certificate Store
- Step 5: Plan Application Deployment

For planning information you should review before deployment, see Plan a Web Farm with IIS Servers. For more information, see Build a Web Farm with IIS Servers.

Step 1: Plan IIS Web Farm Infrastructure

The first step in planning your web farm is to decide its infrastructure. What server configuration do you want? What technologies do you employ? How do you perform load balancing? Does your site need SSL security?

This topic presents the information you need to create a high-level plan for your web farm. It selects a simple web farm configuration and shows the types of servers that are required. In addition, this topic explains how Application Request Routing (ARR) provides load balancing for the web farm.

The following list shows the tasks required to complete this step:

1.1. Decide on Web Farm Infrastructure

1.2. Use ARR for Load Balancing

1.3. Start with a Functional Website

When you are done with these tasks, record your design decisions before going on to Step 2: Plan IIS Web Farm Configuration.

1.1. Decide on Web Farm Infrastructure

There are many ways to design the infrastructure of a web farm. This discussion looks briefly at two common infrastructures—local content and shared network content. It compares advantages and disadvantages of both designs, and selects one for this scenario.

Local Content Infrastructure

The local content infrastructure requires that each web server keeps the content locally. A way must be provided to push the content to all nodes of the web farm.

Advantages

- Local content provides isolation between servers. If one server goes down, other servers are not affected.
- Local content requires fewer computers, because it does not need a back-end content file server.
- You can easily take a web server off line for testing or troubleshooting.

Disadvantages

- Content must be replicated between servers. A common way to handle content replication is by using Distributed File System Replication (DFSR). However, DFSR requires access to a domain server.
- If the website writes to a disk, the data is not available to other servers until it is replicated.
- Each server must have a complete copy of all content. For large websites, the cost of content storage can be high.

Shared Network Content Infrastructure

Shared network content configures all web servers point to a central location (for example, a back-end file server) that contains the website content.

Advantages

- Shared network content is relatively simple to configure and is a good starting point for individuals who are new to web farms or to Microsoft web farm technologies.
- Content written to disk is immediately available on all servers.
- Adding additional servers to the farms is relatively easy. You simply point to the content UNC path.
- Only a few copies of the website files need to be kept. Hard disks on the web servers need contain only the operating system.

Disadvantages

- The back-end file server is a single point of failure. This problem can be minimized by mirroring the file server and providing a means of failover control.
- Network bandwidth can become a limiting factor for a busy web farm.
- File locking issues might arise as multiple servers use the same files.

The Infrastructure Chosen for This Scenario

This scenario uses the shared network content infrastructure. It describes and demonstrates shared content, shared configuration, software load balancing, centralized SSL certificates, and application deployment techniques.

This web farm infrastructure requires you to configure the following servers:

- One Application Request Routing (ARR) server for load balancing. This server requires a default installation of IIS web server. For more information about ARR, see Step 3: Plan IIS Web Farm Load Balancing.
- Two or more web servers. These servers require you to install IIS web server and the IIS modules required to support your website.
- One file server for content and configuration sharing. Set up this server with shared folders for configuration file and website content.
- One file server for the SSL central certificate store. For more information about central certificate store, see Step 4: Plan SSL Central Certificate Store.

1.2. Use ARR for Load Balancing

Microsoft Application Request Routing (ARR) is a proxy-based routing module that forwards HTTP requests to content servers based on HTTP headers, server variables, and load balance algorithms. For more information about ARR, see Step 3: Plan IIS Web Farm Load Balancing.

The easiest way to install ARR and its dependencies is with Web Platform Installer.

To configure ARR load balancing, you need the IP address for all of your servers. To find an IP address, use the Ipconfig.exe command-line tool.

1.3. Start with a Functional Website

The easiest way to configure your web farm is to start by getting your website working on one server. If you need help configuring your website, refer to one or more of the following IIS scenarios:

1. Build a Static Website on IIS
2. Build a Classic ASP Website on IIS
3. Build an ASP.NET Website on IIS
4. Build a PHP Website on IIS
5. Build an FTP Site on IIS

Note
To keep this scenario simple, it is assumed that your web farm servers only one site. However, you can set up multiple sites with little difficulty.

See Also

- Step 2: Plan IIS Web Farm Configuration
- Step 1: Install IIS Web Farm Infrastructure
- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 2: Plan IIS Web Farm Configuration

In the second phase of planning a web farm, determine what is needed to configure shared content and shared configuration. In addition, learn how to add web servers to your farm.

The following list shows the tasks required to complete this step:

2.1. Plan for Shared Content

2.2. Plan for Shared Configuration

2.3 Plan to Add Web Servers

When you are done with these tasks, record your design decisions before going on to Step 3: Plan IIS Web Farm Load Balancing.

2.1. Plan for Shared Content

Network shared content uses a back-end file server to manage website content. All web servers point to a shared folder on the file server over a UNC path. To reduce the risk of failure, the file server is often mirrored to another server with some method of failover provided. For information about failover clusters, see Failover Clustering.

To configure shared content, you first set up the shared folder on the file server. Create a custom user for each application pool and assign that user to the shared folder. You can create local users and groups as long as the same username and password is assigned to each web server.

If your web farm supports ASP.NET applications, use the Code Access Security Policy tool (Caspol.exe) to grant ASP.NET approval for the UNC path. With the Caspol tool, you can grant full trust to the UNC path on your web server. The tool resides on your server at one of the following locations:

- 64-bit systems: %windir%\Microsoft.NET\Framework64\v4.030319
- 32-bit systems: %windir%\Microsoft.NET\Framework\v4.030319

Run the tool by using either a system command prompt, a PowerShell command prompt, or a Visual Studio command prompt. For more information, see Step 2: Configure IIS Web Farm Servers.

2.2. Plan for Shared Configuration

Shared configuration is an IIS feature that helps support homogeneous web farms where all web servers share the same configuration. By using a UNC share, any changes to a master configuration file propagate across different servers without extra tools or programmatic support.

You enable shared configuration in two steps by using the IIS Manager.

1. Export the configuration files to a shared folder on the back-end file server.
2. Point IIS to the UNC path for that shared folder.

2.3 Plan to Add Web Servers

To add web servers to you farm, first use IIS Manager to set up shared configuration by pointing to the configuration UNC path. When you restart IIS Manager, the shared site or sites will be listed. If your web farm supports ASP.NET applications, run the Caspol tool on the new server.

See Also

- Step 3: Plan IIS Web Farm Load Balancing
- Step 1: Plan IIS Web Farm Infrastructure
- Step 2: Configure IIS Web Farm Servers

- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 3: Plan IIS Web Farm Load Balancing

In this phase of planning a web farm with IIS servers, you learn about load balancing. You also learn how to configure load balancing by using Application Request Routing (ARR). ARR is a proxy-based request routing and load balancing module for IIS.

The following list shows the tasks required to complete this step:

3.1. Plan Load Balancing with ARR

3.2. Survey Other Features of ARR

When you are done with these tasks, record your design decisions before going on to Step 4: Plan SSL Central Certificate Store.

3.1. Plan Load Balancing with ARR

Load balancing is a way to distribute workload across multiple web servers. The purpose is to attain optimal resource utilization, maximize request throughput, minimize response time, and avoid server overload.

Load balancers use various algorithms to accomplish this task. One of the simplest algorithms is round robin, which sends each new request to a different web server in an attempt to load all servers equally. Other algorithms send specific types of requests to specific servers in an attempt to reduce response time.

Application Request Routing (ARR) is a proxy-based routing module that uses HTTP headers, server variables, and load balance algorithms to determine how to forward requests to content servers. ARR leverages the URL Rewrite module to inspect incoming requests.

A key feature of ARR is called *host name affinity*. Host name affinity creates an affinity (close connection) between requests and a given server (or set of servers). This ensures that a given site is consuming resources only on a defined number of servers.

Prepare a list of IP addresses for all servers in your farm. You will need the list to configure ARR load balancing.

3.2. Survey Other Features of ARR

In addition to load balancing, ARR provides many features that are beyond the scope of this scenario. The following table summarizes the features that are available in ARR.

Feature	Description
Browsing cached contents with the UI	Cached contents across multiple drives, including cached directories and files, are viewable in IIS Manager.
Byte-range support	ARR segments byte-range requests into smaller pieces to help increase the cache hit/miss ratio without impacting the response time.
Cache hierarchy management	ARR lets you define and manage relationships between the cache nodes. It supports the Cache Array Routing Protocol (CARP), so that you can configure ARR as an edge cache node.
Cache proxy node in CDN/ECN environment	ARR can be used as a cache proxy in a content delivery network (CDN) or edge cache network (ECN) deployment. Because ARR is a proxy-based HTTP request-routing module with support for disk cache and CARP.
Caching compressed objects	ARR compresses and stores objects in a disk-based cache so that these objects remain uncompressed in real time for every request.
Caching while serving responses	ARR can cache large files while serving the response in real time.
Client affinity	ARR can use cookies to affinitize all requests from a client to a content server.
Disk-based caching	ARR supports cached content on user-specified primary cache drives and a secondary network-based cache drive.
Failed request tracing rules	ARR uses tracing rules to troubleshoot and diagnose failed requests.
Health monitoring	ARR provides configuration parameters to use in live traffic and specific URL tests to determine the health of content servers.
Host name affinity	Host name affinity—a feature specific to shared hosters—changes the deployment topology by providing two providers, round robin and memory, to determine which server the host

Feature	Description
	affinitizes.
HTTP-based routing decisions	Application Request Routing works with the URL Rewrite module to write routing rules that are based on HTTP headers and server variables, and to make routing decisions at the application level.
Live request support	ARR supports live content requests by consolidating requests and checking the cache-miss requests before forwarding requests to the origin server.
Load balance algorithms	ARR provides six load balance algorithms that you can use to identify which content server to use to service HTTP requests.
Management and monitoring using the UI	ARR configuration settings and runtime statistics are viewable in IIS Manager.
Multiple server groups	ARR can manage multiple groups of content servers in pilot management and A/B testing scenarios.
Overriding cache-control directives	ARR lets you manually override cache-control directives that control caching behavior, such as the ability to cache and cache duration.
Removing cached contents	ARR lets you delete cached contents by matching URL patterns.
Warming up cache nodes	ARR lets you pre-cache content, anticipating demand before the content is requested. You can pre-cache the parent cache tier only or both the parent and child cache tiers.

See Also

- Step 4: Plan SSL Central Certificate Store
- Step 2: Plan IIS Web Farm Configuration
- Step 3: Configure IIS Web Farm Load Balancing
- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 4: Plan SSL Central Certificate Store

In this phase of planning for a web farm, you add support for SSL-secured websites by configuring a central certificate store. Centralized SSL certificate support is a new feature on IIS 8.

The following list shows the tasks required to complete this step:

4.1. Introduction to Centralized Certificates

4.2 Plan a Central Certificate Store

When you are done with these tasks, record your design decisions before going on to Step 5: Plan Application Deployment.

4.1. Introduction to Centralized Certificates

On Windows Server® 2012, the **Centralized SSL Certificate Support** feature allows server administrators to store and access the certificates centrally on a file share. You can configure a centralized certificate store on you web farm to load the certificates from the file share.

Using centralized certificates simplifies management SSL bindings. SSL requires the DNS name and CN name of a certificate to match. Similar contract can be extended to the file names of the certificates. For example, *www.contoso.com* would use the certificate with a file name *www.contoso.com.pfx*. This contract enables Windows Server 2012 to have just one SSL binding, regardless of the number of secure sites that are using this feature. IIS 8 infers what certificate to use by the SNI value or hostname of the requested website, and by matching it to the file name of the certificate.

4.2 Plan a Central Certificate Store

Similar to shared configuration, the centralized certificates uses a shared folder on a dedicated back-end file server to store the certificates for the web farm. Do not put the certificate shared folder on the content file server.

The central certificate store requires certificates to use the following naming conventions:

- Certificate names must the following form: CN_name.pfx (for example, www.contoso.com.pfx).
- If the certificate is a wild character certificate, use an underscore (_) as the wild character (for example, _.contoso.com.pfx).
- If the certificate has multiple CN names, they must be named as individual files (for example, www.contoso1.com.pfx, www.contoso2.com.pfx, and so forth).

See Also

- Step 5: Plan Application Deployment
- Step 3: Plan IIS Web Farm Load Balancing

- Step 4: Configure SSL Central Certificate Store
- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 5: Plan Application Deployment

When your web farm is up and working, you will want to update your website or add new websites to your server. File Transfer Protocol (FTP) has existed since the early days of the internet for this purpose. Microsoft's Web Deployment Tool (Web Deploy) simplifies deployment of the complex web applications of today.

In this step of building a web farm, you plan the installation and configuration of FTP and Web Deploy for you web farm.

The following list shows the tasks required to complete this step:

5.1. Deploy Websites with FTP

5.2 Deploy Web Applications with Web Deploy

When you are done with these tasks, record your design decisions before going on to Step 1: Install IIS Web Farm Infrastructure.

5.1. Deploy Websites with FTP

File Transfer Protocol (FTP) simple is a protocol for transferring files between computer systems. IIS 8 includes an FTP server that is easy to configure. For the web farm in this scenario, install FTP on the load balancing computer (ARR server) and configure it to point to the back-end shared content. For more information about FTP configuration, see Build an FTP Site on IIS.

5.2 Deploy Web Applications with Web Deploy

The Web Deployment Tool (Web Deploy) enables IIS to synchronize, package, and deploy web applications, websites, or web server content and configuration. It has numerous features that let you include those components that you want to process and exclude those components that you do not. For you to be able to use Web Deploy, first install IIS on the source and destination computers.

The easiest way to install Web Deploy is by using the Web Platform Installer. You can install Web Deploy by running the Web Deploy Windows® Installer or from the command line by using Msiexec.exe. The Web Platform Installer requires fewer steps, while the other methods allow you to customize your installation. In all cases, log on as an administrator.

For more information about Web Deploy, see Web Deployment Tool Overview and Introduction to Web Deploy.

See Also

- Step 1: Install IIS Web Farm Infrastructure
- Step 4: Plan SSL Central Certificate Store
- Configure a Web Farm with IIS Servers
- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Configure a Web Farm with IIS Servers

To install and configure an IIS web farm, follow the steps listed.

- Step 1: Install IIS Web Farm Infrastructure
- Step 2: Configure IIS Web Farm Servers
- Step 3: Configure IIS Web Farm Load Balancing
- Step 4: Configure SSL Central Certificate Store
- Step 5: Configure Application Deployment

After you have completed these planning steps, see Configure a Web Farm with IIS Servers. For more information, see Build a Web Farm with IIS Servers.

Step 1: Install IIS Web Farm Infrastructure

The first step in installing and configuring an IIS 8 web farm is to install IIS on the web servers and load balancing server. Then install Application Request Routing (ARR) on the load balancing server. Finally, set up your website on one of the web servers. The goal of this step is to get the software you need installed on your servers. You configure your web farm later in the process.

The following list shows the tasks required to complete this step:

1.1. Install IIS with Appropriate Modules

1.2 Install ARR for Load Balancing

1.3 Set up Your Website on One Web Server

When you are done with these tasks, record your design decisions before going on to Step 2: Configure IIS Web Farm Servers.

1.1. Install IIS with Appropriate Modules

Install IIS on all servers in your web farm except your back-end file servers (content server and certificate store server). The ARR server requires only a default install of IIS with **Centralized SLL Certificate Support**. All of the web servers require you to install IIS and all IIS modules that support the types of web applications you intend to run.

To install IIS on an ARR server

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select the name of your computer, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, notice the preselected features, and then click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select Features** page, note the preselected features that are installed by default, and then select **Centralized SSL Certificate Support** under the **Security** node.
10. Click **Next**.
11. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
12. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
13. To verify that IIS installed successfully, enter the following into a web browser:

    http://localhost

    You see the default IIS Welcome page.

To install IIS and IIS modules on web servers

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select your server, and click **Next**.

6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, note the preselected features that are installed by default, and then select additional role services for the type of web applications you plan to support.

   Caution
   Install only the role services your websites require. This practice increases your website security.
8. Click **Next**.
9. On the **Web Server Role (IIS)** page, click **Next**.
10. On the **Select Role Services** page, note the preselected role services that are installed by default, and then click **Next**.

    Note
    You only have to install the IIS 8 default role services for a static-content web server.
11. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
12. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
13. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    You should see the default IIS Welcome page.

1.2 Install ARR for Load Balancing

Once you have a default installation of IIS on a server, you can use Web Platform Installer to install the current version of Application Request Routing (ARR).

To install ARR load balancing and its dependencies

1. Open a browser to the Application Request Routing website, and click the **Install** button.
2. Click **Allow**, and then click **Yes**.
3. In the **Web Platform Installer** window, click **Install**.
4. Look over the components to be installed, and then click **I Accept**.

   When the installer completes the installation, it displays and summary showing what was installed.
5. Click **Finish**, and then click **Exit**.

1.3 Set up Your Website on One Web Server

Before you proceed to the next step, set up and test your web application on one of your web servers. This will make configuring shared configuration and shared content much easier.

If you need help configuring your website on your web server, refer to one of the following IIS scenarios:

1. Build a Static Website on IIS
2. Build a Classic ASP Website on IIS
3. Build an ASP.NET Website on IIS
4. Build a PHP Website on IIS
5. Build an FTP Site on IIS

See Also

- Step 2: Configure IIS Web Farm Servers
- Step 4: Plan SSL Central Certificate Store
- Configure a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 2: Configure IIS Web Farm Servers

In the second phase of building your web farm, you configure shared content and shared configuration on your first web server. Then you add more web servers to your farm.

The following list shows the tasks required to complete this step:

2.1. Prepare Your Back-end File Server

2.2. Configure Shared Content

2.3. Set up Shared Configuration

2.4. Add Web Servers to Your Farm

When you are done with these tasks, record your design decisions before going on to Step 3: Configure IIS Web Farm Load Balancing.

2.1. Prepare Your Back-end File Server

In this procedure, you set up a back-end file server with two shared folders. One folder contains the shared content files for the website. The other folder contains the shared configuration files.

Once the shared content folder is ready, copy the files for the website you set up on the primary web server to the content folder.

To set up your back-end file server

1. On the file server, open Windows Explorer.
2. Create a folder for the shared content, for example `C:\Content`.
3. Create a folder for the shared configuration, fox example `C:\Config`.
4. Share the content folder by using the custom user credentials.
5. Share the configuration folder by using the custom user credentials.
6. Copy the files from the working website on your primary web server to the content folder on the file server, for example `C:\Content\MySite`.

2.2. Configure Shared Content

In this procedure, you set up shared content of your primary web server. When finished, your web server retrieves content from the back-end file server.

To configure shared content

1. On your primary, open IIS Manager, select you website, and click **Advanced Settings** in the **Actions** pane.
2. On the **Advancing Settings** page, change the value of **Physical Path** to the configuration UNC path, for example `\\FileServer\Content\MySite`.
3. Select **Physical Path Credentials**, and click the ellipsis (…) button.
4. In the **Connect As** dialog, select **Specific user**, and click **Set**.
5. Enter the credentials of the custom shared user account, and click **OK**.
6. Click **OK** two more times to exit **Advanced Settings**.
7. If your web farm supports ASP.NET, open command prompt, and enter one the following commands depending on your system:
   - 64-bit systems: `cd %windir%\Microsoft.NET\Framework64\v4.030319`
   - 32-bit systems: `cd %windir%\Microsoft.NET\Framework\v4.030319`
8. Enter the following command:

   ```
   caspol –m –ag 1. –url file://<file server>/<content>/<website>/* FullTrust
   ```

   Replace `<file server>` with the name of your file server. Replace `<content>` with the name of your shared content folder. Replace `<website>` with the folder name for your website. Note **FullTrust** is case-sensitive.
9. Close the command prompt, select **Application Pools** in the **Connections** page.
10. Select the application pool for your site, and click **Recycle** in the **Actions** pane.
11. Open a browser, view your site with **localhost**. The site is displayed as usual, but the content is coming from the back-end file server.

2.3. Set up Shared Configuration

To configure shared configuration, export the configuration files to the back-end file server. Then point the web server to the shared configuration location.

To set up shared configuration

1. Open IIS Manager, select your server name, and double-click **Shared Configuration**.
2. In the **Actions** pane, click **Export Configuration**.
3. In the **Physical path** box, type the UNC path for the shared configuration folder on the file server, for example `\\FileServer\Config`.
4. Click **Connect As**.
5. Enter the credentials of the custom shared user account, and click **OK**.
6. Under **Encryption Keys**, type a password for encryption the configuration files, and confirm the password.
7. Click **OK**.
8. On the Shared Configuration page, select the **Enable shared configuration** check box.
9. In the **Physical path** box, type the UNC path to the configuration folder on the file server.
10. In the **User name** box, type the custom shared user name.
11. In the **Password** and **Confirm password** boxes, type the password for the custom share user account.
12. In the **Actions** pane, click **Apply**.
13. In the **Encryption Keys Password** dialog, type the encryption keys password and click **OK**.
14. Close IIS Manager, and reopen it.

    IIS is now using the configuration files that are located of the file server.

2.4. Add Web Servers to Your Farm

A web server added to your farm by using this procedure retrieves both configuration information and website content from the back-end file server.

To add a web server to your farm

1. Open IIS Manager, select your server name, and double-click **Shared Configuration**.
2. On the Shared Configuration page, select the **Enable shared configuration** check box.
3. In the **Physical path** box, type the UNC path to the configuration folder on the file server.
4. In the **User name** box, type the custom shared user name.
5. In the **Password** and **Confirm password** boxes, type the password for the custom share user account.
6. In the **Actions** pane, click **Apply**.
7. In the **Encryption Keys Password** dialog, type the encryption keys password and click

OK.

8. Close and reopen IIS Manager.
9. If your web farm supports ASP.NET, open command prompt, and enter one the following commands depending on your system:
   - 64-bit systems: `cd %windir%\Microsoft.NET\Framework64\v4.030319`
   - 32-bit systems: `cd %windir%\Microsoft.NET\Framework\v4.030319`
10. Enter the following command:

    ```
    caspol –m –ag 1. –url file://<file server>/<content>/<website>/* FullTrust
    ```

    Replace `<file server>` with the name of your file server. Replace `<content>` with the name of your shared content folder. Replace `<website>` with the folder name for your website. Note **FullTrust** is case-sensitive.
11. Close the command prompt, select **Application Pools** in the **Connections** page.
12. Select the application pool for your site, and click **Recycle** in the **Actions** pane.
13. Open a browser, view your site with **localhost**. The site displays like it did on the primary server.

See Also

- Step 3: Configure IIS Web Farm Load Balancing
- Step 1: Install IIS Web Farm Infrastructure
- Step 2: Plan IIS Web Farm Configuration
- Configure a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 3: Configure IIS Web Farm Load Balancing

Now that IIS is installed all web servers. All web servers share both contend and configuration. And IIS and ARR are installed on your load balancing server. It is time to configure ARR and turn your servers into a functioning web farm.

The following list shows the tasks required to complete this step:

3.1. Create a Server Farm with ARR

3.2. Configure Load Balancing with ARR

3.3. Change Application Pool Settings

When you are done with these tasks, record your design decisions before going on to Step 4: Configure SSL Central Certificate Store.

3.1. Create a Server Farm with ARR

To create your server farm, you must enter the IP address of every web server in your farm.

To create a server farm with ARR

1. Open IIS Manager.
2. In the **Connections** pane, expand the server node, and select **Server Farms**.
3. In the **Actions** pane, click **Create Server Farm**.

   The **Create Server Farm** wizard opens to the **Specify Server Farm Name** page.
4. In the **Server farm name** box, type a name for your server farm, and click **Next**.

   The **Add Server** page is displayed.
5. In the **Server address** box, type the IP address of the first server, and click **Add**.
6. Continue typing IP addresses and clicking **Add** until all web servers are entered.
7. Click **Finish**.

3.2. Configure Load Balancing with ARR

This procedure configures ARR load balancing with an algorithm that distributes incoming request evenly among the web servers. It then sets up the server-farm health test feature.

To configure load balancing with ARR

1. Open IIS Manager.
2. In the Connections pane, expend the server
3. Under the server node, expand **Server Farms**, and then select the server farm that you created.
4. In the **Server Farm** pane, double-click **Load Balance**.
5. On the **Load Balance** page, select **Weighted round robin** from the **Load balance algorithm** list, and then click **Apply**.
6. In the **Connections** pane, select the server farm that you created.
7. In the **Server Farm** pane, double-click **Health Test**.
8. On the **Health Test** page, in the **URL** box, enter a URL that you want ARR to test by sending a GET request to determine the health of the servers. Click **Apply**.
9. To verify the health of your farm, click **Verify URL Test**.

3.3. Change Application Pool Settings

Because all HTTP requests and responses go through ARR, delays or errors occur if the application pool times out or recycles unexpectedly. Set the idle time-out and application pool recycle settings to zero to avoid problems.

To change application pool settings

1. On one of the web servers in your farm, open IIS Manager.
2. In the **Connections** pane, select **Application Pools**.
3. In the **Application Pools** pane, select the application pool for you website, and then click **Advanced Settings** in the **Actions** pane.
4. In the **Advanced Settings** dialog box, expand **Process Model**, and change the **Idle Time-out (minutes)** value to 0.
5. Expand **Recycling**, and change the **Regular time intervals (in minutes)** value to 0.
6. Click **OK**.

See Also

- Step 4: Configure SSL Central Certificate Store
- Step 2: Configure IIS Web Farm Servers
- Step 3: Plan IIS Web Farm Load Balancing
- Configure a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 4: Configure SSL Central Certificate Store

In the final phase of building a web farm with IIS servers, you add a central certificate store to your farm. This store enables you to place your SSL certificates on one file server and share them with all web servers in farm. You no longer have to have copies of every certificate on every server.

This article guides you through configuring a central certificate store and testing it on your web farm.

The following list shows the tasks required to complete this step:

4.1. Configure a Central Certificate Store

When you are done with these tasks, record your design decisions before going on to Step 5: Configure Application Deployment.

4.1. Configure a Central Certificate Store

Before attempting to configure a central certificate store, create a shared folder on a dedicated back-end file server. You must provide a UNC path to the shared folder (for example, \\CertServer\Certs).

To configure a central certificate store

1. On your ARR server, open IIS Manager.
2. In the **Connections** pane, select the server node.
3. Double-click **Centralized Certificates**.
4. In the **Actions** pane, click **Edit Feature Settings**.
5. In the **Edit Centralized Certificates Settings** dialog, select the **Enable Centralized Certificates** check box.
6. In the **Physical path** box, type the UNC path to the shared certificate folder.
7. In the **User name** box, type the name of the custom user account on the file server.
8. In the **Password** and **Confirm password** boxes, type the password for the account.
9. If your certificates use a private key password, type the password in the **Password** and **Confirm password** boxes.
10. Click **OK**.

See Also

- Step 3: Configure IIS Web Farm Load Balancing
- Step 5: Configure Application Deployment
- Step 4: Plan SSL Central Certificate Store
- Configure a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Step 5: Configure Application Deployment

In this step of building a web farm, you install and configure FTP for deploying website updates and new websites. You also install and configure Web Deploy for deploying web applications.

The following list shows the tasks required to complete this step:

5.1. Install and Configure FTP for Your Web Farm

5.2. Install and Test Web Deploy for Your Web Farm

5.3. Where do I go from here?

5.1. Install and Configure FTP for Your Web Farm

The procedures in this section guide you through installing and configuring FTP for your web farm. Install FTP on your load balancing computer (ARR server) only. The web servers do not need FTP installed. For more information about FTP configuration, see Build an FTP Site on IIS.

To install FTP for your farm

1. On the **Start** screen, and click the **Server Manager** tile.
2. In the **Server Manager Dashboard**, click **Add roles and features**.
3. If the **Before you begin** page of **Add Roles and Features Wizard** is displayed, click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation**, and click **Next**.
5. On the **Select destination server** page, select **Select a server from the server pool**, select your server from the **Server Pool** list, and then click **Next**.
6. On the **Select server roles** page, expand the **Web Server (IIS)** node, and then expand the **FTP Server** node.
7. Select the **FTP Server** check box and the **FTP Service** check box, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.

To add an FTP site

1. Open IIS Manager.
2. In the **Connections** pane, expand the server node and click the **Sites** node.
3. In the **Actions** pane, click **Add FTP Site** to open the **Add FTP Site** wizard.
4. On the **Site Information** page, in the **FTP site name** box, type a unique friendly name for the FTP site.
5. In the **Physical path** box, type the UNC path to the shared content folder on your back-end file server.
6. Click **Next** to open the **Binding and SSL Settings** page.
7. Under **Binding**, in the **IP Address** list, select or type an IP address if you do not want the IP address to remain **All Unassigned**.
8. In the **Port** box, type the port number.
9. Optionally, in the **Virtual Host** box, type a host name if you want to host multiple FTP sites on a single IP address. For example, type **www.contoso.com**.
10. Clear the **Start FTP site automatically** box if you want to start the site manually.
11. Under **SSL**, from the **SSL Certificate** list, select a certificate. Optionally, click **View** to open the **Certificates** dialog box and verify information about the selected certificate.
12. Select one of the following options:
    - **Allow SSL**: Allows the FTP server to support both non-SSL and SSL connections with a client.
    - **Require SSL**: Requires SSL encryption for communication between the FTP server and a client.
13. Click **Next** to open the **Authentication and Authorization Information** page.
14. Under **Authentication**, select the authentication method or methods that you want to

use:

- **Anonymous**: Allows any user to access content providing only the user name **anonymous** or **ftp**. (Most, but not all, FTP clients enter user name for you automatically.)
- **Basic**: Requires users to provide a valid user name and password to access content. Because Basic authentication transmits unencrypted passwords across the network, use this authentication method only when you know that the connection between the client and FTP server is secure, such as by using Secure Sockets Layer (SSL).

15. Under **Authorization**, from the **Allow access to** list, select one of the following options:
    - **All Users**: All users, whether they are anonymous or identified, can access the content.
    - **Anonymous Users**: Anonymous users can access the content.
    - **Specified Roles or User Groups**: Only members of certain roles or user groups can access the content. Type the role or user group in the corresponding box.
    - **Specified Users**: Only specified users can access the content. Type the user name in the corresponding box.
16. If you selected an option from the **Allow access to** list, select one or both of the following permissions:
    - **Read**: Permits authorized users to read content from the directory.
    - **Write**: Permits authorized users to write to the directory.
17. Click **Finish**.

5.2. Install and Test Web Deploy for Your Web Farm

Install Web Deploy on your source computer (your development computer with IIS installed). Install Web Deploy on each of your web application servers. Use Web Platform Installer to install the current version of Web Deploy and its dependencies.

To test your Web Deploy installation, use it to update your website.

To install Web Deploy for you farm

1. Open a browser to the Web Deployment Tool website, and click the **Install** button.
2. Click **Allow**, and then click **Yes**.
3. In the **Web Platform Installer** window, click **Install**.
4. Look over the components to be installed, and then click **I Accept**.

   When the installer completes the installation, it displays and summary showing what was installed.
5. Click **Finish**, and then click **Exit**.

To update your website with Web Deploy

1. On your development computer, open IIS Manager.
2. In the **Connections** pane, select the updated website.
3. In the **Actions** pane, click **Export Application**.

   The **Export Application Package** wizard appears.
4. On the **Select the Contents of the Package** page, make sure all check boxes are selected, and click **Next**.
5. On the **Select Parameters** page, add parameters if needed for your application, and click **Next**.
6. On the **Save Package** page, type the path (including the package name) to the location where you want the package saved.

   This location could to a shared folder on the development computer or a shared folder on the destination computer (the ARR server).
7. Click **Next**.

   The package is exported and saved.
8. Click **Finish**.
9. On one of your web application servers, open IIS Manager.
10. In the **Actions** page, click **Import Application**.

    The **Import Application Package** appears.
11. On the **Select the Package** page, type the path (including the package name) to the location where you exported the package.
12. Click **Next**.
13. On the **Select the Contents of the Package** page, make sure all check boxes are selected, and click **Next**.
14. On the **Enter Application Package Information** page, type the path to your website, and click **Next**.
15. On the **Overwrite Existing Files** page, choose whether to delete all files that are in the new package, and click **Next**.
16. Click **Finish**.

    Web Deploy has updated your website.

5.3. Where do I go from here?

You now have a working web farm. It uses ARR for load balancing. It employs shared content and shared configuration. You have a central certificate store for use with SSL-secured websites. And you can deploy websites and web applications with either FTP or Web Deploy.

There are many ways to improve your web farm, here are a few suggestions:

- Allow users who are not administrators to deploy web applications with Web Deploy. For more information, see Allowing non-admin users to deploy web applications.
- Use failover clustering to mirror your back-end file servers, eliminate signal points of failure on the back-end, and improve throughput. For more information, see Failover Clustering.

- You should also to eliminate signal points on failure on the front end. You can use Network Load Balancing (NLB) with multiple ARR servers achieve high availability, scalability, and stability on the front end of your farm. To learn more about using ARR and NLB together, see Achieving High Availability and Scalability—ARR and NLB.

See Also

- Step 4: Configure SSL Central Certificate Store
- Step 5: Plan Application Deployment
- Configure a Web Farm with IIS Servers
- Plan a Web Farm with IIS Servers
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Increasing Server, Storage, and Network Availability: Scenario Overview

Server Availability in Windows Server 2012 provides a collection of experiences that are designed to improve availability, performance, and reliability at the single-server and multiple-server (scale-up and scale-out) levels. The experiences in Server Availability include the following:

- **Improved NTFS availability** Install and configure large volumes that are resilient against downtime caused by potential file corruption.
- **File server for scale-out application data** Deploy server applications, such as Microsoft SQL Server and Hyper-V, on continuously available file servers.
- **Continuously available DHCP service** Improve network availability by using Dynamic Host Configuration Protocol (DHCP) failover services.
- **Network failure resilience** Use Load Balancing and Failover to aggregate bandwidth and maximize network reliability.
- **Continuously available block storage for server applications** Create continuously available block storage for server applications by using iSCSI Transparent Failover.
- **Fast and efficient file servers for server applications** Improve network and file performance by using Remote Direct Memory Access (RDMA) and SMB Direct.
- **Network performance and availability for Windows file servers** Take advantage of all available network bandwidth by using SMB Multichannel.
- **Storage spaces** Enable cost-effective and flexible storage solutions for business critical deployments.
- **Offloaded data transfers** Move data quickly between storage devices by using Offloaded Data Transfer.

- **Storage mobility for virtual machines** Move virtual hard disks to other locations with no downtime.
- **Direct connections for virtualized workloads to Fibre Channel-based storage** Provide server applications with access to Fibre Channel-based storage directly from a guest operating system.
- **Continuously available file servers for heterogeneous environments** Create continuously available Network File System (NFS) file shares by using NFS transparent failover.
- **Data protection** Protect your application data by using SMB Volume Shadow Copy for remote file shares.
- **Hyper-V over SMB** Enable virtual machines to reside on an SMB file shares, giving you the power to design new flexible storage solutions for your virtual or cloud infrastructure.

In this scenario

Use the following links to help understand the Server Availability experiences:

- Deploying Fast and Efficient File Servers for Server Applications
- Scale-Out File Server for Application Data Overview
- Step-by-Step: Configure DHCP for Failover
- NIC Teaming Overview
- iSCSI Target Block Storage Overview
- What's New in Failover Clustering in Windows Server 2012
- Network Performance and Availability
- Install SQL Server with SMB fileshare as a storage option
- Application Compatibility and API Support for SMB 3.0, CSVFS, and ReFS
- Hyper-V Virtual Fibre Channel Overview
- Configure Live Migration and Migrating Virtual Machines without Failover Clustering
- Protect Data on Remote SMB File Shares using VSS
- Deploy Hyper-V over SMB
- Windows Offloaded Data Transfers Overview

See also

Use the links in the following table to learn more about Server Availability.

Content type	References
Product evaluation	Server Message Block overview| What's New in Failover Clustering in Windows Server 2012| What's New in Hyper-V for Windows Server 2012| File and Storage Services Overview

Deploying Fast and Efficient File Servers for Server Applications

This document outlines the planning and deployment steps for setting up fast and efficient file servers for server applications, such as Hyper-V™ and Microsoft SQL Server.

Overview

In Windows Server® 2012, the SMB protocol for remote storage has been enhanced to allow for faster and more efficient file servers for server applications, such as Hyper-V and SQL Server. As part of the SMB protocol, two new features, SMB Direct and SMB Multichannel, enable customers to deploy storage for server applications on cost efficient, continuously available, high performance file servers.

SMB Direct supports the use of network adapters that have remote direct memory access (RDMA) capability. SMB Direct (SMB over RDMA) is a new storage protocol in Windows Server 2012 that includes:

- **Increased throughput**: Leverages the full throughput of high speed networks where the network adapters coordinate the transfer of large amounts of data at line speed.
- **Low latency**: Provides extremely fast responses to network requests, and, as a result, makes remote file storage feel as if it is directly attached block storage.
- **Low CPU utilization**: Uses fewer CPU cycles when transferring data over the network, which leaves more power available to server applications.

SMB Direct is automatically configured by Windows Server 2012.

SMB Multichannel allows file servers to use multiple network connections simultaneously and includes the following capabilities:

- **Fault tolerance**. When using multiple network connections at the same time, the file server continues functioning despite the loss of a network connection.
- **Increased throughput**. The file server can simultaneously transmit more data using multiple connections for high speed network adapters or multiple network adapters.

SMB Multichannel is automatically configured by Windows Server 2012.

Requirements and recommendations

Before deploying the file server, you should review the requirements and recommendations outlined below.

SMB Direct

SMB Direct requires the following:

- At least two computers running Windows Server 2012

- A network adapter with RDMA capability. Currently, these network adapters are available in three different types: iWARP, Infiniband, or RoCE (RDMA over Converged Ethernet).

SMB Multichannel

SMB Multichannel requires the following:

- At least two computers running Windows Server 2012 or Windows RT.
- At least one of the configurations below:

  Multiple network adapters

  One or more network adapters that support RSS (Receive Side Scaling)

  One or more network adapters that support RDMA

The following are sample network configurations that can be used for SMB Multichannel:

- **Single 10 GbE network adapter**. Each computer is configured with a 10 GbE network adapter, which is RSS-capable or RDMA-capable.
- **Dual 1 GbE network adapters configured in a team**. Each computer is configured with two 1 gigabit Ethernet network adapters in a load balancing and failover environment, also known as a network adapter team.
- **Dual 1 GbE network adapters**. Each computer is configured with dual 1 gigabit Ethernet network adapters.
- **Dual 10 GbE network adapters**. Each computer is configured with dual 10 GbE network adapters. These adapters could also be RSS-capable and/or RDMA-capable.
- **Dual Infiniband network adapters**. Each computer is configured with dual Infiniband network adapters. These adapters are typically RDMA-capable.

Install the required roles, role services, and features

By default, both SMB Multichannel and SMB Direct are enabled on Windows Server 2012. You do not need to add any roles, role services, or features to utilize the functionality of SMB Multichannel and SMB Direct. For information on installing roles, role services, and features, see “Step 1: Install Prerequisites for Scale-out File Servers” in Deploy Scale-Out File Server.

Step-by-step instructions

Use the following steps to validate a configuration that leverages SMB Multichannel or SMB Direct. Both SMB Multichannel and SMB Direct can be used with different file server configurations, including standalone file server clusters or Scale-Out file server clusters. In this document, the focus is on the Scale-out File Server cluster configuration.

- Step 1: Verify the basic network configuration
- Step 2: Configure a failover cluster
- Step 3: Configure the networks for the failover cluster

- Step 4: Configure a Scale-out File Server
- Step 5: Verify each file server name has two addresses
- Step 6: Configure a Hyper-V or Microsoft SQL Server client
- Step 7: Verify servers are using SMB Multichannel and SMB Direct
- Step 8: Monitor file shares using Performance Counters

Step 1: Verify the basic network configuration

After installing Windows Server 2012, you can use the following Windows PowerShell commands to verify the network adapters that are configured. If you are using RSS-capable and/or RDMA-capable network adapters, you can verify that these capabilities are being properly detected.

To verify the basic network configuration

1. Open **Windows PowerShell**.
2. On each server, type the following to view a list of network adapters:

   ```
   Get-NetAdapter
   ```

   Review the list of network adapters installed on the system, along with their basic characteristics. You should have at least two adapters configured.
3. On each server, type the following to view a list of network adapters available to SMB:

   ```
   Get-SmbServerNetworkInterface
   ```

   Review the list of network adapters, and note whether they are RSS-capable and/or RDMA-capable. You should have at least two adapters configured.
4. You can also verify the network configuration using Server Manager. In Server Manager, click **Local Server** and verify that you have at least two network adapters configured for the server.

Step 2: Configure a failover cluster

For instructions on configuring a failover cluster in Windows Server 2012, see Deploy Scale-Out File Server.

Step 3: Configure the networks for the failover cluster

In this example, use two networks in your failover cluster, which allows you to use SMB Multichannel in a two network adapter configuration. You should enable two of the networks on the failover cluster for client access.

To configure the networks using Failover Cluster Manager

1. From one of the cluster nodes, open Server Manager.
2. In the Server Manager menu bar, click **Tools** to access the contents of the **Administrative Tools** folder. Select **Failover Cluster Manager** from the list of tools.

3. On the left pane in Failover Cluster Manager, click to expand the failover cluster you want to use, and then click **Networks**.
4. Right-click the cluster network you want to configure, and click **Properties**.
5. Make sure both **Allow cluster network communication on this network** and the checkbox next to **Allow clients to connect through this network** are selected.
6. Repeat, as required, so that two of the cluster networks are configured with both of these properties.

To configure the networks using Windows PowerShell

1. Open **Windows PowerShell**.
2. From one of the cluster nodes, type the following to configure the networks:

   ```
   Get-ClusterNetwork
   ```

   The list of networks available to the failover cluster is displayed, along with their role.
3. From the cluster node, type the following to specify client access:

   ```
   (Get-ClusterNetwork <NetworkName>).Role = 3
   ```

   The specified cluster network is configured for client access. You should repeat this step for two of the networks that are listed.

Step 4: Configure a Scale-out File Server

For instructions on configuring a Scale-out File Server in Windows Server 2012, see Deploy Scale-Out File Server.

Step 5: Verify each file server name has two addresses

In this example, confirm that the Scale-out File Server is properly configured for SMB Multichannel by verifying that you have two network adapters configured for the Scale-out File Server name. If you are using RDMA network adapters, you can also verify if these are properly configured.

To configure the networks using Windows PowerShell

1. Open **Windows PowerShell**.
2. From one of the file servers, type the following:

   ```
   Get-ClusterNetwork
   ```

   The list of networks available to the failover cluster is displayed, along with their role.
3. From the cluster node, type the following to specify client access:

   ```
   Get-SmbServerNetworkInterface
   ```

   In addition to the list of network interfaces you viewed in Step 1, you should see additional entries for the file server name you have created. You should have at least two network adapters configured for the file server name, and you can also verify if they are

RSS-capable and/or RDMA-capable.

Step 6: Configure a Hyper-V or Microsoft SQL Server client

To configure Hyper-V or Microsoft SQL Server as a client for your Scale-out File Server, see Deploy Scale-Out File Server.

Step 7: Verify servers are using SMB Multichannel and SMB Direct

In this example, confirm that the application server (running either Hyper-V or Microsoft SQL Server) is properly leveraging the two network adapters using SMB Multichannel. If you are using RDMA adapters, you can also verify if you are using SMB Direct.

To verify that the servers are using SMB Multichannel and SMB Direct

1. Open **Windows PowerShell**.
2. On the application server using SMB, type the following:

   ```
   Get-SmbClientNetworkInterface
   ```

   The list of network adapters available to the SMB Client is displayed, along with the indication on whether they are RSS-capable and/or RDMA-capable.
3. On the application server using SMB, type the following:

   ```
   Get-SmbMultichannelConnection
   ```

   You can view the connections that are actively used for the currently established sessions, and you can also verify if they are RSS-capable and/or RDMA-capable. You should have two paths listed for the Scale-out File Server name.

   Note
   When you use this cmdlet, it will not show any information unless the application server is actively accessing the file server.

Step 8: Monitor file shares using Performance Counters

In Windows Server 2012, new SMB performance counters provide detailed information about I/O size, I/O latency, and IOPS, allowing administrators to analyze the performance of SMB file shares where their data is stored. These counters are specifically designed for server applications, such as Hyper-V and SQL Server, which store files on remote Windows file shares.

To monitor application server activity on the SMB share

1. In the Server Manager menu bar, click **Tools** to access the contents of the **Administrative Tools** folder. Select **Performance Monitor** from the list of tools.
2. Under **Performance**, and then under **Monitoring Tools**, click **Performance Monitor**.
3. Right-click in the Performance Monitor window, and click **Add Counters**. You can also

do this by pressing CTRL+N.

4. To view activity for virtual machines, under **Available counters**, expand **SMB2 Client Shares** and select all of the counters.
5. Select the share used by the application server, and click **Add**. Click **OK** to continue.
6. On the Taskbar, click **Change graph type** to change the type to **Report**. You can also do this by pressing CTRL+G twice.
7. You can now view the activity of the counters in the Performance Monitor window.

See also

- File and Storage Services Overview
- Server Message Block overview

Scale-Out File Server for Application Data Overview

Scale-Out File Server is a feature that is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provides the ability to share the same folder from multiple nodes of the same cluster. This scenario focuses on how to plan for and deploy Scale-Out File Server.

You can deploy and configure a clustered file server by using either of the following methods:

- **Scale-Out File Server for application data** This clustered file server feature was introduced in Windows Server 2012, and it lets you store server application data, such as Hyper-V virtual machine files, on file shares, and obtain a similar level of reliability, availability, manageability, and high performance that you would expect from a storage area network. All file shares are simultaneously online on all nodes. File shares associated with this type of clustered file server are called scale-out file shares. This is sometimes referred to as active-active. This is the recommended file server type when deploying either Hyper-V over Server Message Block (SMB) or Microsoft SQL Server over SMB.
- **File Server for general use** This is the continuation of the clustered file server that has been supported in Windows Server since the introduction of Failover Clustering. This type of clustered file server, and therefore all the shares associated with the clustered file server, is online on one node at a time. This is sometimes referred to as active-passive or dual-active. File shares associated with this type of clustered file server are called clustered file shares. This is the recommended file server type when deploying information worker scenarios.

Scenario description

With scale-out file shares, you can share the same folder from multiple nodes of a cluster. For instance, if you have a four-node file server cluster that is using Server Message Block (SMB) Scale-Out, a computer running Windows Server 2012 R2 or Windows Server 2012 can access

file shares from any of the four nodes. This is achieved by leveraging new Windows Server Failover Clustering features and the capabilities of the Windows file server protocol, SMB 3.0. File server administrators can provide scale-out file shares and continuously available file services to server applications and respond to increased demands quickly by simply bringing more servers online. All of this can be done in a production environment, and it is completely transparent to the server application.

Key benefits provided by Scale-Out File Server in include:

- **Active-Active file shares** All cluster nodes can accept and serve SMB client requests. By making the file share content accessible through all cluster nodes simultaneously, SMB 3.0 clusters and clients cooperate to provide transparent failover to alternative cluster nodes during planned maintenance and unplanned failures with service interruption.
- **Increased bandwidth** The maximum share bandwidth is the total bandwidth of all file server cluster nodes. Unlike previous versions of Windows Server, the total bandwidth is no longer constrained to the bandwidth of a single cluster node; but rather, the capability of the backing storage system defines the constraints. You can increase the total bandwidth by adding nodes.
- **CHKDSK with zero downtime** CHKDSK in Windows Server 2012 is significantly enhanced to dramatically shorten the time a file system is offline for repair. Clustered shared volumes (CSVs) take this one step further by eliminating the offline phase. A CSV File System (CSVFS) can use CHKDSK without impacting applications with open handles on the file system.
- **Clustered Shared Volume cache** CSVs in Windows Server 2012 introduces support for a Read cache, which can significantly improve performance in certain scenarios, such as in Virtual Desktop Infrastructure (VDI).
- **Simpler management** With Scale-Out File Server, you create the scale-out file servers, and then add the necessary CSVs and file shares. It is no longer necessary to create multiple clustered file servers, each with separate cluster disks, and then develop placement policies to ensure activity on each cluster node.
- **Automatic rebalancing of Scale-Out File Server clients** In Windows Server 2012 R2, automatic rebalancing improves scalability and manageability for scale-out file servers. SMB client connections are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share. This improves efficiency by reducing redirection traffic between file server nodes. Clients are redirected following an initial connection and when cluster storage is reconfigured.

In this scenario

The following topics are available to help you deploy a Scale-Out File Server:

- Plan for Scale-Out File Server

  Step 1: Plan for Storage in Scale-Out File Server

  Step 2: Plan for Networking in Scale-Out File Server
- Deploy Scale-Out File Server

  Step 1: Install Prerequisites for Scale-Out File Server

Step 2: Configure Scale-Out File Server

Step 3: Configure Hyper-V to Use Scale-Out File Server

Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server

When to use Scale-Out File Server

You should not use Scale-Out File Server if your workload generates a high number of metadata operations, such as opening files, closing files, creating new files, or renaming existing files. A typical information worker would generate a lot of metadata operations. You should use a Scale-Out File Server if you are interested in the scalability and simplicity that it offers and if you only require technologies that are supported with Scale-Out File Server.

The following table lists the capabilities in SMB 3.0, the common Windows file systems, file server data management technologies, and common workloads. You can see whether the technology is supported with Scale-Out File Server, or if it requires a traditional clustered file server (also known as a file server for general use).

Technology	File Server for General Use	Scale-Out File Server
SMB capability: SMB Transparent Failover	Yes	Yes
SMB capability: SMB Scale Out	No	Yes
SMB capability: SMB Multichannel	Yes	Yes
SMB capability: SMB Direct	Yes	Yes
SMB capability: SMB Encryption	Yes	Yes
File system: NTFS file system	Yes	No
File system: Resilient File System (ReFS)	Yes	No
File system: CSV File System (CSVFS)	No	Yes
Data management: BranchCache	Yes	No
Data management: Data Deduplication	Yes	Yes **Warning** In Windows Server 2012 R2, Data Deduplication is

Technology	File Server for General Use	Scale-Out File Server
		only supported in a scale-out file server deployment for Virtual Desktop Infrastructure (VDI) workloads with separate storage and compute nodes. The storage must be remote.
Data management: DFS Namespaces: Namespace Server	Yes	No
Data management: DFS Namespaces: Folder Target	Yes	Yes
Data management: DFS Replication	Yes	No
Data management: File Server Resource Manager	Yes	No
Data management: File Classification Infrastructure	Yes	No
Data management: File Server Volume Shadow Copy Service (VSS) Agent	Yes	Yes
Data management: Folder Redirection	Yes	Yes
Data management: Client-Side Caching	Yes	Yes
Workload: Information worker	Yes	Not recommended
Workload: Hyper-V	Yes	Yes
Workload: Microsoft SQL Server	Yes	Yes

Practical applications

Scale-Out File Servers are ideal for server applications that keep files open for a long amount of time, doing mostly data operations with infrequent metadata operations on the file system. Hyper-V virtual hard disks and SQL Server database files can be stored on a scale-out file share.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
Failover Clustering Overview	Failover clusters added the following features in Windows Server 2012 to support scale-Out file server: Distributed Network Name, the Scale-Out File Server resource type, Cluster Shared Volumes (CSV) 2, and the Scale-Out File Server High Availability role. For more information about these features, see What's New in Failover Clustering in Windows Server 2012 on Microsoft TechNet.
Server Message Block overview	SMB 3.0 added the following features in Windows Server 2012 to support scale-Out File Server: SMB Transparent Failover, SMB Multichannel, and SMB Direct. For more information on new and changed functionality for SMB in Windows Server 2012 R2, see What's New in SMB in Windows Server 2012 R2.

Plan for Scale-Out File Server

Before you deploy Scale-Out File Server, you should review the requirements and plan for the deployment. The information you should review includes:

Failover Clustering requirements Because a Scale-Out File Server is built on Failover Clustering, any requirements for Failover Clustering also apply to Scale-Out File Server.

Server application storage requirements Microsoft SQL Server and Hyper-V are the two server applications that are supported by Scale-Out File Server for storage. You should review the storage requirements for the server application you plan to use.

Existing storage in your organization You can use storage that may already be deployed within your organization.

Network configuration You should review your network adapter configuration, the bandwidth of the CSV redirection traffic, and the DNS configuration for the cluster nodes.

Use the following steps to plan a Scale-Out File Server deployment in your organization:

- Step 1: Plan for Storage in Scale-Out File Server
- Step 2: Plan for Networking in Scale-Out File Server

See also

After you have completed these planning steps, see Deploy Scale-Out File Server.

Scale-Out File Server for Application Data Overview.

Step 1: Plan for Storage in Scale-Out File Server

When you plan the storage that will support your scale-out file server, you can leverage the existing storage in your organization.

Note

Before you deploy Scale-Out File Server, you should review the storage requirements for the server applications that will use Scale-Out File Server.

After you complete this step, see Step 2: Plan for Networking in Scale-Out File Server.

This topic includes the sections that are outlined in the following table, and it can help you through the process of identifying requirements for the server applications in your organization and identifying existing storage that you can leverage.

Task	Description
Review Failover Clustering requirements	Scale-Out File Server uses the features and functionality that are included with Failover Clustering. You should review the requirements for Failover Clustering when planning for a scale-out file server.
Review server application storage requirements	Microsoft SQL Server and Hyper-V are server applications that are supported with Scale-Out File Server. If you plan to store your database files, database transaction logs, or virtual hard disks on scale-out file shares, the same storage recommendations for their respective server

Task	Description
	applications apply.
Review existing storage in your organization	You can use the storage that has already been deployed within your organization.

Review Failover Clustering requirements

Scale-Out File Server is built on Failover Clustering, so any requirements for Failover Clustering also apply to Scale-Out File Server. You should have an understanding of Failover Clustering before deploying Scale-Out File Server. For more information about Failover Clustering, see:

- Use Cluster Shared Volumes in a Failover Cluster
- Installing the Failover Cluster Feature and Tools in Windows Server 2012
- Creating a Windows Server 2012 Failover Cluster
- How to Troubleshoot Create Cluster failures in Windows Server 2012

Some important considerations for Failover Clustering and Scale-Out File Server are as follows:

- The storage configuration must be supported by Failover Clustering before you deploy Scale-Out File Server. You must successfully run the Cluster Validation Wizard before you add Scale-Out File Server.
- Scale-Out File Server requires the use of Clustered Shared Volumes (CSVs).
- Scale-Out File Server is not supported for use with Resilient File System.
- Accessing a continuously available file share as a loopback share is not supported. For example, if Microsoft SQL Server or Hyper-V store data files on SMB file shares, they must run on computers that are not a member of the file server cluster for the SMB file shares.

Review server application storage requirements

Microsoft SQL Server and Hyper-V are the two server applications that are supported by Scale-Out File Server. If you choose to use Microsoft SQL Server or Hyper-V with Scale-Out File Server, the storage recommendations for direct-attached storage apply. For more information about storage requirements, see:

- SQL Server Best Practices Article
- Analyzing I/O Characteristics and Sizing Storage Systems for SQL Server Database Applications

Review existing storage in your organization

It is not necessary to deploy new storage in your organization to support Scale-Out File Server. You can use existing storage that may already be deployed within your organization. Some supported storage configurations that can be used as the storage subsystem for Scale-Out File Server include:

- **Storage Spaces** Storage Spaces was introduced in Windows Server 2012.

  For more information, see Storage Spaces Overview.
- **iSCSI Storage Area Network**

  For more information about iSCSI in Windows Server 2012, see iSCSI Target Block Storage Overview.
- **Fibre Channel Storage Area Network**

  For more information about using a Fibre Channel SAN, see Failover Clustering Hardware Requirements and Storage Options.
- **Clustered RAID controller**

  For more information, see Enable Support for Clustered Windows Servers using clustered RAID controllers.

See also

- Step 2: Plan for Networking in Scale-Out File Server
- Plan for Scale-Out File Server
- Scale-Out File Server for Application Data Overview
- Increasing Server, Storage, and Network Availability: Scenario Overview

Step 2: Plan for Networking in Scale-Out File Server

Networking planning for Scale-Out File Server includes reviewing the network adapter configuration, ensuring that the CSV redirection traffic network bandwidth is sufficient, and that DNS dynamic update protocol is being used for Scale-Out File Server.

Use the following table to help you review the configuration of your network to ensure that it will work when deploying Scale-Out File Server.

Task	Description
2.1. Review your network adapter configurations	Ensure that the network adapter configurations are consistent across all of your nodes in Scale-Out File Server.
2.2. Review the CSV redirection traffic network configuration	Ensure that the network that includes the CSV redirection traffic has sufficient bandwidth.
2.3. Review the DNS configuration for the cluster nodes	Use DNS dynamic update protocol for the cluster node name and all of the cluster nodes.

2.1. Review your network adapter configurations

The network adapters in each of the cluster nodes for the Scale-Out File Server should be the same. For more information on the different options that are available, see Deploying Fast and Efficient File Servers for Server Applications.

2.2. Review the CSV redirection traffic network configuration

The network containing the CSV redirection traffic must have enough bandwidth to service all requests. If the network bandwidth is not adequate, you could create an internal network for just CSV redirection traffic and use an external network for the rest.

2.3. Review the DNS configuration for the cluster nodes

You should make sure that the cluster node name is registered by using DNS dynamic update protocol. This should include the name of the Scale-Out File Server and the IP addresses of all of the network adapters in every cluster node on the client network.

See also

- Step 1: Plan for Storage in Scale-Out File Server
- Plan for Scale-Out File Server
- Scale-Out File Server for Application Data Overview
- Increasing Server, Storage, and Network Availability: Scenario Overview

Deploy Scale-Out File Server

Use the following steps to deploy Scale-Out File Server in your organization:

- Step 1: Install Prerequisites for Scale-Out File Server
- Step 2: Configure Scale-Out File Server
- Step 3: Configure Hyper-V to Use Scale-Out File Server
- Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server

Important

To take full advantage of Scale-Out File Server, all servers running the server applications that are using scale-out file shares should be running Windows Server 2012 R2 or Windows Server 2012. If the server application is running on Windows Server 2008 or Windows Server 2008 R2, the servers can connect to the scale-out file shares, but they cannot take advantage of any of the new features. If the server application is running

Windows Server 2003, the server receives access-denied error when connecting to the scale-out file share.

Step 1: Install Prerequisites for Scale-Out File Server

Scale-Out File Server leverages features that are included in the File and Storage Services role and in the Failover Clustering feature. You must install the prerequisite features and role services, create a failover cluster, add then add storage to a cluster shared volume.

Task	Description
1.1. Install role services and features	Install the File Server role service and Failover Clustering by using Server Manager or by using Windows PowerShell.
1.2. Validate hardware and create a cluster	Validate the hardware configuration and create a cluster.
1.3. Add storage to a cluster shared volume	A cluster shared volume is required for a scale-out file share.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Install role services and features

The File Server role service, which is part of the File and Storage Services server role, and the Failover Clustering feature are required to configure Scale-Out File Server. You configure these by using Server Manager or by using Windows PowerShell.

Do this step using Windows PowerShell

To install the roles and features

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Before you begin** page of the **Add Roles and Features Wizard**, click **Next**.
5. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.

6. On the **Select destination server** page, select the appropriate server, and then click **Next**. The local server is selected by default.
7. On the **Select server roles** page, expand **File and Storage Services**, expand **File Services**, and then select the **File Server** check box. Click **Next**.
8. On the **Select features** page, select the **Failover Clustering** check box, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. Repeat the steps in this procedure for each server that will be added to the cluster.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WindowsFeature –name File-Services,Failover-Clustering –IncludeManagementTools
```

1.2. Validate hardware and create a cluster

You can ensure that you are using the proper hardware by validating the configuration. After the validation is complete, you can create the cluster.

Do this step using Windows PowerShell

To validate hardware and create a cluster

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, and then click **Server Manager**.
3. Click **Tools**, and then click **Failover Cluster Manager**.
4. Under the **Management** heading, click **Validate Configuration**.
5. On the **Before You Begin** page, click **Next**.
6. On the **Select Servers or a Cluster** page, in the **Enter name** box, type the FQDN of one of the servers that will be part of the cluster, and then click **Add**. Repeat this step for each server that will be in the cluster.

   Note
   If you already have a cluster configured, you can enter the name of the cluster or one of its nodes instead of entering the name of each server.
7. Click **Next**.
8. On the **Testing Options** page, ensure that the **Run all tests (recommended)** option is selected, and then click **Next**.
9. On the **Confirmation** page, click **Next**.
10. On the **Summary** page, ensure that the **Create the cluster now using the validated nodes** check box is selected, and then click **Finish**. The Create Cluster Wizard appears.

11. On the **Before You Begin** page, click **Next**.
12. On the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type a name for the cluster, and then click **Next**.
13. On the **Confirmation** page, click **Next**.
14. On the **Summary** page, click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Test-Cluster –Node server1, server2

New-Cluster –Name ClusterName –Node server1, server2
```

1.3. Add storage to a cluster shared volume

A cluster shared volume is used by a failover cluster, and it can be accessed by more than one node at a time. You can add to a cluster shared volume by using Failover Cluster Manager.

Do this step using Windows PowerShell

To add storage to a cluster shared volume

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, and then click **Server Manager**.
3. Click **Tools**, and then click **Failover Cluster Manager**.
4. Click **Storage**, right-click the disk that you want to add to the cluster shared volume, and then click **Add to Cluster Shared Volumes**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-ClusterSharedVolume "ClusterDiskName"
```

See also

- Step 2: Configure Scale-Out File Server
- Step 3: Configure Hyper-V to Use Scale-Out File Server
- Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server
- Deploy Scale-Out File Server
- Scale-Out File Server for Application Data Overview

Step 2: Configure Scale-Out File Server

In this step, you configure Scale-Out File Server for high availability and create a continuously available file share on the cluster shared volume.

Task	Description
Configure Scale-Out File Server	Configure Scale-Out File Server to provide continuous availability for server applications.
Create a continuous availability file share on the cluster shared volume	Create a file share on the cluster shared volume.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Configure Scale-Out File Server

Scale-Out File Server must be configured for continuous availability.

To configure the File Server role

1. Sign in to the server as a member of the local Administrators group.
2. To open Failover Cluster Manager in Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
3. Right-click the name of the cluster, and then click **Configure Role**.
4. On the **Before You Begin** page, click **Next**.
5. On the **Select Role** page, click **File Server**, and then click **Next**.
6. On the **File Server Type** page, select the **Scale-Out File Server for application data** option, and then click **Next**.
7. On the **Client Access Point** page, in the **Name** box, type a NETBIOS name that will be used to access Scale-Out File Server, and then click **Next**.
8. On the **Confirmation** page, confirm your settings, and then click **Next**.
9. On the **Summary** page, click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-ClusterScaleOutFileServerRole -Name DistributedNetworkName -Cluster ClusterName
```

Create a continuous availability file share on the cluster shared volume

You must create a file share for continuous availability on the cluster shared volume by using Failover Cluster Manager.

Do this step using Windows PowerShell

To create a continuous availability file share on the cluster shared volume

1. Sign in to the server as a member of the local Administrators group.
2. To open Failover Cluster Manager in Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
3. Expand the cluster, and then click **Roles**.
4. Right-click the file server role, and then click **Add File Share**.
5. On **the Select the profile for this share** page, click **SMB Share – Applications**, and then click **Next**.
6. On the **Select the server and path for this share** page, click the name of the cluster shared volume, and then click **Next**.
7. On the **Specify share name** page, in the **Share name** box, type a name for the file share, and then click **Next**.
8. On the **Configure share settings** page, ensure that the **Enable continuous availability** check box is selected, and then click **Next**.

   Note
   You should not use access-based enumeration on file shares for Scale-Out File Server because of the increased metadata traffic that is generated on the coordinator node.
9. On the **Specify permissions to control access** page, click **Customize permissions**, grant the following permissions, and then click **Next**:
   - For Hyper-V: All Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators must be granted full control on the share and the file system.
   - For Microsoft SQL Server: The SQL Server service account must be granted full control on the share and the file system.
10. On the **Confirm selections** page, click **Create**.
11. On the **View results** page, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-Item -Name "folder" -ItemType Directory

New-SmbShare -Name ShareName -Path Path -FullAccess Domain\Account
```

```
Set-SmbPathAcl -ShareName ShareName
```

See also

- Step 1: Install Prerequisites for Scale-Out File Server
- Step 3: Configure Hyper-V to Use Scale-Out File Server
- Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server
- Deploy Scale-Out File Server
- Scale-Out File Server for Application Data Overview

Step 3: Configure Hyper-V to Use Scale-Out File Server

Hyper-V is one of two Microsoft applications that are supported when using Scale-Out File Server for data storage. When the failover cluster is configured, virtual machines can be configured by using Hyper-V Manager.

Task	Description
3.1. Verify permissions	Ensure that the permissions on the continuously available file share are configured correctly to host virtual machines.
3.2. Create a new virtual machine	Create a new virtual machine by using Hyper-V Manager.

3.1. Verify permissions

You should ensure that the proper permissions are granted on the continuously available file share before you create the virtual machines.

To verify permissions

1. Log on to the server as a member of the local Administrators group.
2. Open Windows Explorer and navigate to the shared folder.
3. Right-click the folder, and then click **Properties**.
4. Click the **Sharing** tab, click **Advanced Sharing**, and then click **Permissions**.
5. Ensure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions.

6. Click **OK** twice.
7. Click the **Security** tab.
8. Ensure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions.

3.2. Create a new virtual machine

You create a virtual machine by using Hyper-V Manager. When you are prompted for the virtual machine location, enter the name of the client access point that is configured in the failover cluster for Scale-Out File Server.

To create a virtual machine

1. Log on to the server running Hyper-V as a member of the local Administrators group.
2. Open Hyper-V Manager. Click **Start**, and then click **Hyper-V Manager**.
3. Right-click the server running Hyper-V name, point to **New**, and then click **Virtual Machine**.
4. On the **Before You Begin** page, click **Next**.
5. On the **Specify Name and Location** page, in the **Name** box, type a name for the virtual machine.
6. Select the **Store the virtual machine in a different location** check box. In the **Location** box, type the client access point that is configured in the failover cluster for Scale-Out File Server, and then click **Next**.
7. On the **Assign Memory** page, type the desired amount of memory in the **Startup memory** box, and then click **Next**.
8. On the **Configure Networking** page, in the **Connection** box, choose the appropriate network, and then click **Next**.
9. On the **Connect Virtual Hard Disk** page, create a virtual hard disk and set the location to the client access point that is configured in the failover cluster for Scale-Out File Server, and then click **Next**.
10. On the **Completing the New Virtual Machine Wizard** page, click **Finish**.

See also

- Step 1: Install Prerequisites for Scale-Out File Server
- Step 2: Configure Scale-Out File Server
- Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server
- Deploy Scale-Out File Server
- Scale-Out File Server for Application Data Overview

Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server

SQL Server is one of two Microsoft applications that are supported when using File Server for scale-out application data storage. SQL Server 2008 R2 and SQL Server 2012 are supported in this scenario. SQL Server 2008 R2 is supported in a stand-alone configuration, and SQL Server 2012 adds support for clustered servers.

Task	Description
4.1. Verify permissions	Ensure that the permissions on the continuously available file share are configured correctly to host the database files.
4.2. Create the database files	Create the database files by using SQL Server Management Studio or by using a query.

4.1. Verify permissions

You should ensure that the proper permissions are granted on the continuously available file share before you create SQL Server database files.

To verify permissions

1. Log on to the file server as a member of the local Administrators group.
2. Open Windows Explorer and navigate to the scale-out file share.
3. Right-click the folder, and then click **Properties**.
4. Click the **Sharing** tab, click **Advanced Sharing**, and then click **Permissions**.
5. Ensure that the SQL Server service account has full-control permissions.
6. Click **OK** twice.
7. Click the **Security** tab.
8. Ensure that the SQL Server service account has full-control permissions.

4.2. Create the database files

When creating the database files, you can store the database files in a scale-out file share by using SQL Server Management Studio or by using a query. In SQL Server 2012, you can choose to store all database files in a scale-out file share during installation.

See also

- Step 1: Install Prerequisites for Scale-Out File Server
- Step 2: Configure Scale-Out File Server
- Step 3: Configure Hyper-V to Use Scale-Out File Server
- Deploy Scale-Out File Server
- Scale-Out File Server for Application Data Overview

Network Performance and Availability

This topic describes how Server Message Block (SMB) Multichannel provides a network performance and availability experience in Server Availability.

Scenario description

In Windows Server 2012, the SMB Multichannel feature provides multiple connections within a single SMB session to enhance the network performance and availability experience for Server Availability. SMB Multichannel provides the following advantages:

- **Failover** Uses a different network adapter if a network adapter fails in a computer.
- **Throughput** Aggregates bandwidth across multiple network adapters, and uses multiple processors to process network interruptions on network adapters that support Receive Side Scaling.
- **Configuration** Automatically detects and uses multiple network paths.

Practical applications

Some practical applications of the network performance and availability experience include:

- **SQL over SMB** Install SQL database files on an SMB network share. For more information about SQL over SMB, see Install SQL Server with SMB fileshare as a storage option.
- **Hyper-V over SMB** Enable virtual machines to reside on an SMB network shares, giving you the power to design new flexible storage solutions for your virtual or cloud infrastructure.

Features included in this experience

The following table lists the features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Server Message Block overview	SMB Multichannel is a new feature in the SMB version 3.0 protocol.

Role/feature	How it supports this scenario
NIC Teaming Overview	Load balancing and failover (also known as network adapter teaming), allows multiple network adapters on a computer to be placed into a team.

Hardware requirements

SMB Multichannel requires at least one of the following configurations:

- Multiple network adapters
- Network adapters that support Receive Side Scaling
- Network adapters that support Remote Direct Memory Access (RDMA)

The following are a few sample network configurations that can be used for SMB Multichannel:

- **Dual 1 gigabit Ethernet network adapters** Each computer running Windows Server 2012 is configured with a dual 1 gigabit Ethernet network adapter.
- **Dual 1 gigabit Ethernet network adapters configured in a team** Each computer is configured with two 1 gigabit Ethernet network adapters that are configured in a load balancing and failover environment, also known as a network adapter team.
- **Single 10 gigabit Ethernet network adapters** Each computer running Windows Server 2012 is configured with a 10 gigabit Ethernet network adapter that support Receive Side Scaling.
- **Dual 10 gigabit Ethernet network adapters** Each computer is configured with a dual 10 gigabit Ethernet network adapter.
- **Dual Infiniband network adapters** Each computer is configured with dual Infiniband network adapters that support RDMA.

Software requirements

SMB Multichannel is enabled by default on computers running Windows Server 2012.

Deploy Hyper-V over SMB

SMB 3.0 file shares can be used as shared storage for Hyper-V in Windows Server 2012 R2 and Windows Server 2012. With this capability, Hyper-V can store virtual machine files, which includes configuration, virtual hard disk (VHD) files, and snapshots, on SMB file shares. Listed below are the main advantages of storing application data for Hyper-V on SMB file shares:

- **Ease of provisioning and management.** You can manage file shares instead of storage fabric and logical unit numbers (LUNs).

- **Increased flexibility.** You can dynamically migrate virtual machines or databases in the data center.
- **Ability to take advantage of existing investment in a converged network.** You can use your existing converged network with no specialized storage networking hardware.
- **Reduced capital expenditures.** Capital expenses (acquisition costs) are reduced.
- **Reduced operating expenditures.** You can reduce operating costs because there is no need for specialized storage expertise.

Note
Some Hyper-V features in Windows Server 2012 R2 and Windows Server 2012 now rely on using SMB file shares, such as some types of live migration.

In this guide

- Prerequisites
- Step 1: Configuring file server clusters
- Step 2: Install Hyper-V
- Step 3: Create an SMB file share
- Step 4: Create a virtual machine and virtual hard disk file on the file share
- Step 5: Migrate virtual machine storage to an SMB file share
- Step 6: Initiate a live migration of a virtual machine to another cluster node
- Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage
- Troubleshooting

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

Using Hyper-V with SMB has the following requirements:

- One or more computers running Windows Server 2012 R2 or Windows Server 2012 with the Hyper-V role installed. You can also use non-Microsoft file servers that implement the SMB 3.0 protocol.
- One or more computers running Windows Server 2012 R2 or Windows Server 2012 with the File and Storage Services role installed.
- A common Active Directory infrastructure. The servers running Active Directory Domain Services (AD DS) do not need to run Windows Server 2012 R2 or Windows Server 2012.

The three most common file server configurations for Hyper-V over SMB are a single-node file server, a dual-node file server, and a multi-node file server, as shown in the following figure.

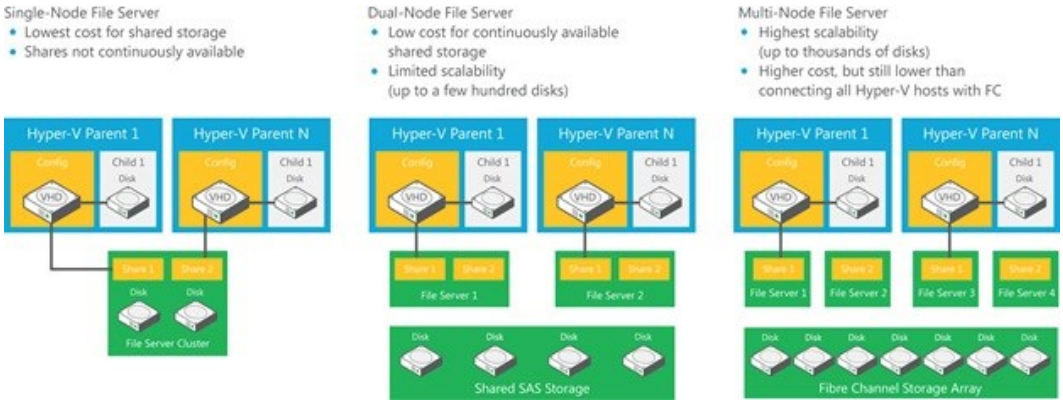

Figure 1 Common configurations for Hyper-V over SMB2

The two supported Hyper-V configurations for Hyper-V over SMB are:

- Standalone Hyper-V servers (not a high availability solution)
- Hyper-V servers configured in a failover cluster

Note

The Hyper-V host must have Windows Server 2012 R2 or Windows Server 2012 installed.

Considerations when using Hyper-V with SMB

- An Active Directory infrastructure is required, so you can grant permissions to the computer account of the Hyper-V hosts.
- The file server must have Windows Server 2012 R2 or Windows Server 2012 installed, so the new SMB 3.0 protocol is available. You can also use non-Microsoft file servers that implement the SMB 3.0 protocol. Hyper-V does not block older versions of SMB, however, the Hyper-V Best Practice Analyzer issues an alert when an older version of SMB is detected.
- Loopback configurations (where the computer that is running Hyper-V is used as the file server for virtual machine storage) are not supported.
- You must have separate failover clusters for Hyper-V and for the file server.

Step 1: Configuring file server clusters

To deploy Hyper-V over SMB, use one of the following procedures for your file server configuration. All servers in your file server configuration must have Windows Server 2012 R2 or Windows Server 2012 installed.

To configure a standalone file server

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type

servermanager.exe, and then click **Server Manager**.

3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, select the appropriate server, and then click **Next**. The local server is selected by default.
6. On the **Select server roles** page, click **File and Storage Services**, and then click **Next**.
7. On the **Confirm installation selections** page, click **Install**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To add the File and Storage Services role, type:

```
Install-WindowsFeature File-Services, FS-FileServer
```

If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available. To view the list of network adapters, type:

```
Get-NetAdapter

Get-SmbServerNetworkInterface
```

To configure a clustered file server

1. With Windows Server 2012 R2 or Windows Server 2012 installed on two servers, add the File and Storage Services role and the Failover Clustering feature on each server by typing:

   ```
   Install-WindowsFeature File-Services, FS-FileServer,
   Failover-Clustering

   Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature
   ```

2. If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available and that they are configured on different subnets. To view the list of network adapters, type:

   ```
   Get-NetAdapter

   Get-SmbServerNetworkInterface
   ```

3. To create a failover cluster using the two servers, type:

   ```
   New-Cluster –Name ClusterName -Node FileServer1, FileServer2
   ```

4. To create a file server cluster to host continuously available SMB file shares, where `FST` is the name of the file server cluster and `Cluster Disk 1` is the storage, type:

   ```
   Add-ClusterFileServerRole -Name FST -Storage "Cluster Disk 1"
   –StaticAddress 192.168.101.22/24, 192.168.102.22/24
   ```

The above example assumes you are using two networks with addresses, 192.168.101.22/24 and 192.168.102.22/24, for SMB network traffic. Two networks are recommended for network fault tolerance.

To configure a failover cluster with a scale-out file server

1. Follow steps 1-3 in the previous procedure, To configure a clustered file server.
2. To create a scale-out file server on the failover cluster to host continuously available SMB file shares, where `FSO` is the name of the scale-out file server cluster and `Cluster Disk 2` is the name of the storage, type:

```
Add-ClusterSharedVolume "Cluster Disk 2"
Add-ClusterScaleOutFileServerRole -Name FSO
```

Step 2: Install Hyper-V

To continue deploying Hyper-V over SMB, install the Hyper-V role on a separate server.

To install the Hyper-V role

1. To install the Hyper-V role on a separate server, follow steps 1-5 from To configure a standalone file server.
2. On the **Select server roles** page, click **Hyper-V**, and then click **Next**.
3. On the **Confirm installation selections** page, click **Install**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To install the Hyper-V role and Hyper-V Windows PowerShell cmdlets and tools, type:

```
Install-WindowsFeature Hyper-V, Hyper-V-PowerShell, Hyper-V-Tools
```

If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available, and they are not connected to the virtual switch. To view the list of network adapters, type:

```
Get-NetAdapter

Get-SmbClientNetworkInterface
```

Step 3: Create an SMB file share

The folder used by Hyper-V to store virtual machine data requires specific permissions to access the SMB file share. You need to make sure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions.

To create an SMB file share by using Server Manager

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. On the left, click **File and Storage Services**.
4. Click **Tasks**, and then click **New Share** to open the **New Share Wizard**.
5. On the **Select Profile** page, select **SMB Share – Applications**, and click **Next**.
6. On the **Share Location** page, select a server and a volume, and click **Next**.
7. On the **Share Name** page, specify a name for the new share, and click **Next**.
8. On the **Permissions** page, click **Customize Permissions**.
9. Click **Add**, click **Select a Principal**, and then click **Object Types**.
10. In **Object Types**, click to select **Computers**, and click **OK**.
11. Enter the name of the computer, and click **OK**.
12. In **Permissions Entry**, select **Full Control**, and click **OK**.
13. Repeat the previous three steps for the second Hyper-V server. Click OK when finished.
14. On the **Permissions** page, click **Next**.
15. Click **Create** to create the SMB file share.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

For a standalone file server or a clustered file server, type the following to configure an SMB file share (where `HV1` and `HV2` are servers running Hyper-V, `HVC` is the Hyper-V cluster account, and `HVadmin` is the Hyper-V administrator account):

```
# Create folder

MD X:\VMS

# Create file share

New-SmbShare -Name VMS1 -Path X:\VMS -FullAccess Domain\HVAdmin, Domain\HV1$,
Domain\HV2$, Domain\HVC$

# Set NTFS permissions from the file share permissions

Set-SmbPathAcl -Name VMS1
```

For a scale-out file server cluster, type the following to configure an SMB file share (where `HV1` and `HV2` are servers running Hyper-V and `HVadmin` is the Hyper-V administrator account):

```
# Create folder

MD X:\VMS

# Create file share
```

```
New-SmbShare -Name VMS1 -Path X:\VMS -FullAccess Domain\HVAdmin, Domain\HV1$,
Domain\HV2$, Domain\HVC$
apply permissions
# Set NTFS permissions from the file share permissions
Set-SmbPathAcl -Name VMS1
```

Step 4: Create a virtual machine and virtual hard disk file on the file share

To create a virtual hard disk (VHD) and virtual machine on an SMB file share, use Hyper-V Manager or Hyper-V Windows PowerShell cmdlets. You will also need to specify a Universal Naming Convention (UNC) path (for example, \\servername\sharename).

To create a virtual machine

1. Open Hyper-V Manager. From the Server Manager Tools menu, click **Hyper-V Manager**.
2. From the navigation pane of Hyper-V Manager, select the computer running Hyper-V.
3. From the Actions pane, click **New** and then click **Virtual Machine**.
4. The New Virtual Machine wizard opens. Click **Next**.
5. On the **Specify Name and Location** page, type an appropriate name and location (using a UNC path). To configure Hyper-V over SMB, you must use a UNC path.
6. On the **Assign Memory** page, specify enough memory to start the guest operating system.
7. On the **Configure Networking** page, connect the virtual machine to the switch you created when you installed Hyper-V.
8. On the **Connect Virtual Hard Disk** and **Installation Options** pages, choose **Create a virtual hard disk**. Click **Next**, and then click the option that describes the type of media you will use. For example, to use an .iso file, click **Install an operating system from a boot CD/DVD** and then specify the path to the .iso file.
9. On the Summary page, verify your selections and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

For a file server named `FS1` and a file share named `VMS`, to create a virtual machine named `VM1`, type:

```
New-VHD -Path \\FS1\VMS\VM1.VHDX -VHDType Dynamic -SizeBytes 127GB
New-VM -Name VM1 -Path \\FS1\VMS -Memory 1GB -VHDPath \\FS1\VMS\VM1.VHDX
```

Step 5: Migrate virtual machine storage to an SMB file share

You can migrate virtual machine storage from both direct-attached storage (DAS) to an SMB file share and from a SMB file share to another SMB file share.

To migrate virtual machine storage from local storage to an SMB file share

1. To confirm that there is a virtual machine using local storage, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

   To migrate virtual machine storage from an SMB file share to another SMB file share, specify the SMB file share in this step.
2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To migrate the virtual machine storage to an SMB file share, type:

   ```
   Move-VMStorage –VMName VM1 –DestinationStoragePath \\FST\VMS
   ```
4. To confirm that a virtual machine is using an SMB file share and the workload is not interrupted, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

Step 6: Initiate a live migration of a virtual machine to another cluster node

You can transparently move running virtual machines from one cluster node to another node in the same cluster without a dropped network connection or perceived downtime.

To initiate a live migration of a virtual machine to another cluster node

1. To confirm that a virtual machine is running in a cluster node and is using an SMB file share, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   Get-ClusterGroup VM1 | FT Name, OwnerNode, State
   ```
2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To perform a live migration of a virtual machine to another cluster node, type:

   ```
   Move-ClusterVirtualMachineRole -Name VM1 -Node HV2 –
   ```

```
VmMigrationType Live
```

4. To confirm that the virtual machine moved to another cluster node and the workload is not interrupted, type:

```
Get-ClusterGroup VM1 | FT Name, OwnerNode, State
```

Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage

To move running virtual machines from one cluster node to another node in the same cluster and migrate virtual machine storage, use one of the following procedures:

To move a virtual machine to another Hyper-V host

1. To confirm that a virtual machine is running on the Hyper-V host, type:

```
Get-VM VM1 | FT Name, Path, State
```

2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To perform a live migration of a virtual machine to another Hyper-V host, type:

```
Move-VM –Name VM1 –DestinationHost HV2
```

4. To confirm that the virtual machine moved to another Hyper-V host and the workload is not interrupted, type:

```
Get-VM VM1 | FT Name, Path, State
```

To move a virtual machine with direct-attached storage and migrate its virtual machine storage to an SMB file share

1. To confirm that a virtual machine with local storage is running on the Hyper-V host, type:

```
Get-VM VM1 | FT Name, Path, State

Get-VMHardDiskDrive VM1 | FT VMName, Path
```

   To migrate virtual machine storage from an SMB file share to another SMB file share, specify the SMB file share in this step.

2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To move the virtual machine to another Hyper-V host and the storage to an SMB file share, type:

```
Move-VM –Name VM1 -DestinationHost HV2 –
DestinationStoragePath \\FST\VMS
```

4. To confirm that the virtual machine moved to another Hyper-V host using the SMB file share, and the workload is not interrupted, type:

```
Get-VM VM1 | FT Name, Path, State
Get-VMHardDiskDrive VM1 | FT VMName, Path
```

Troubleshooting

This section covers some common issues that you might encounter when using Hyper-V over SMB.

Check and fix issues with permissions

You may experience issues with permissions that are related to accessing an SMB file share or the NTFS folder on which the share was created. To check permissions on an SMB file share (where `VMS1` is the share and `X:\VMS` is the NTFS folder), type the following:

```
Get-SmbShareAccess -Name VMS1 -Path X:\VMS
```

If you determine that a specific computer account is missing from permissions, you can add the account to both the file share and the folder to fix the issue. To fix the permissions, type the following:

```
Grant-SmbShareAccess –Name VMS1 –AccountName Domain\HV3$ -AccessRight Full
Set-SmbPathAcl –Name VMS1
```

Using constrained delegation

When using Hyper-V Manager from a computer running Windows Server 2012 R2 or Windows Server 2012 to manage virtual machines on another computer running Windows Server 2012 R2 or Windows Server 2012, you may experience an error that says access to an SMB file share is denied. Typically, this is because you need delegation rights to use your credentials to access the remote share on another computer. This is a security feature that prevents a user from gaining access to a computer in your network for the purpose of performing actions on other computers in your network. To address this issue, you have two choices:

Option 1: Use Remote Desktop. Use Remote Desktop to access the computer and run Hyper-V Manager directly on that computer.

Option 2: Configure constrained delegation. You can change the properties of the computer account in Active Directory Users and Computers to allow delegation. When enabled, constrained delegation gives you the ability to use a specific SMB remote file share without requiring you to perform an action on any computer. Constrained delegation tells Active Directory Users and Computers that between two computers, (in this case, the Hyper-V server and the SMB file server), and for specific services, (in this case, SMB), it is allowed to re-issue access to the resources.

To configure constrained delegation, for each server running Hyper-V, perform one of the following procedures.

Note

In Windows Server 2012 R2, you can perform this procedure using a new set of Windows PowerShell SMB cmdlets that simplifies the configuration of constrained delegation.

To configure constrained delegation

1. In Active Directory Users and Computers, click to open **Properties** for the computer account, and then click to open the **Delegation** tab.
2. Select both **Trust this computer for delegation to the specified services only** and **Use Kerberos only**.
3. Click **Add**, and provide the name of the SMB file server (or the **Cluster Access Point** for a Scale-Out File Server).
4. Select the **CIFS** service. Note that Common Internet File System (CIFS) is the previous name for SMB.
5. On the SMB file share created for virtual machines, add Full Control permissions for the Hyper-V Administrators.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Before you perform this procedure using the Windows PowerShell SMB cmdlets, the Active Directory module for Windows PowerShell must be available. To install the Active Directory cmdlets, type:

```
Install-WindowsFeature RSAT-AD-PowerShell
```

To configure constrained delegation (where `FileServer1` and `FileServer2` are the servers running Hyper-V and `HV1` and `HV2` are the SMB clients), type:

```
Enable-SmbDelegation –SmbServer FileServer1 –SmbClient HV1

Enable-SmbDelegation –SmbServer FileServer1 –SmbClient HV2

Enable-SmbDelegation –SmbServer FileServer2 –SmbClient HV1

Enable-SmbDelegation –SmbServer FileServer2 –SmbClient HV2
```

Note

This procedure works only with the resource-based constrained delegation available starting in Windows Server 2012, therefore, the Active Directory forest must be at the Windows Server 2012 functional level.

See also

- SMB Share Cmdlets in Windows PowerShell
- File and Storage Services Overview
- Increasing Server, Storage, and Network Availability: Scenario Overview

Protect Data on Remote SMB File Shares using VSS

In Windows Server 2012, a new feature is introduced called VSS for SMB File Shares, which allows VSS-aware backup applications to create a volume shadow copy of VSS-aware server applications that store data on remote SMB 3.0 file shares. This guide provides information on how to perform backups of server applications using this feature, and therefore, protect your data.

In this document

- VSS for SMB File Shares: overview
- Requirements and supported configurations
- Deployment scenarios
- Step 1: Install File Server VSS Agent Service
- Step 2: Add a user to the Backup Operators local group on the file server
- Step 3: Perform a shadow copy

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

VSS for SMB File Shares: overview

VSS for SMB File Shares is an extension to the existing VSS infrastructure and consists of four parts:

- **File Share Shadow Copy Provider (fssprov.dll)**. This is a new VSS provider on the server running the VSS-aware application. It manages shadow copies on remote Universal Naming Convention (UNC) paths where the application stores its data files, and then relays the shadow copy request to the File Share Shadow Copy Agents. Enhancements to the VSS infrastructure to support the File Share Shadow Copy Provider include updates to the API.
- **File Share Shadow Copy Agent (fssagent.dll)**. This is a new VSS requester on the file server hosting the SMB 3.0 file shares (using an UNC path) that store the application's data files. It manages mappings between file shares and volumes, and also interacts with the file server's VSS infrastructure to perform shadow copies of the volumes that backup the SMB 3.0 file shares (where the VSS-aware applications stores their data).
- **File Server Remote VSS Protocol (MSFSRVP)**. The File Share Shadow Copy Provider and the File Share Shadow Copy Agent use this new RPC-based protocol to coordinate shadow copy requests of data stored on SMB file shares.

Requirements and supported configurations

The requirements for using VSS for SMB File Shares are listed below.

- The application server and file server must be running Windows Server 2012.

- The application server and file server must be joined to the same Active Directory domain.
- The File Server VSS Agent Service role service must be installed on the file server.
- The File Share Shadow Copy agent must run in a security environment that has backup operators or administrator privileges on both the application server and file server.
- The File Share Shadow Copy agent and the application must run in a security environment that has no less than read-only permission on the file share data that will be backed up.

Note
VSS for SMB File Shares also works with third-party Network Attached Storage (NAS) appliances or other similar solutions. These appliances or solutions must support SMB 3.0 and File Server Remote VSS Protocols.

VSS for SMB File Shares supports the following configurations:

- An application server configured as a single server or in a failover cluster.
- A file server configured as a single server or in a failover cluster with continuously available or scale-out file shares.
- File shares with a single DFS Namespaces link target.

VSS for SMB File Shares has the following limitations:

- Unsupported VSS capabilities, such as: hardware transferrable shadow copies, writable shadow copies, VSS fast recovery where a volume can quickly revert to a shadow copy, and client-accessible shadow copies (Shadow Copy of Shared Folders).
- Loopback configurations where an application server accesses its data on SMB file shares that are hosted on the same application server.
- Shadow copies of Hyper-V virtual machines, where the application for which the virtual machine stores its data on SMB file shares is not supported.
- Data on mounted drives that are below the root of the file share are not included in the shadow copy.
- Shadow copies of file shares that do not support failover clustering.

Deployment scenarios

VSS for SMB File Shares is most commonly deployed with Hyper-V, where a server running Hyper-V stores the virtual machine files on remote SMB file share. Some example deployments are described below.

Single server running Hyper-V and single file server

In this scenario, there is a single, non-clustered server running Hyper-V and a single, non-clustered file server. As shown in the diagram below, the file server has two volumes attached to it, with a file share on each volume. The virtual machine files for `VM A` are stored on `\\fileserv\share1`, which is backed by `Volume 1`. For `VM B`, some of the virtual machine files are stored on `\\fileserv\share1`, and some are stored on `\\fileserv\share2`, which is backed by `Volume 2`. The virtual machine files for `VM C` are stored on `\\fileserv\share2`.

Diagram of a single server running Hyper-V and a single file server

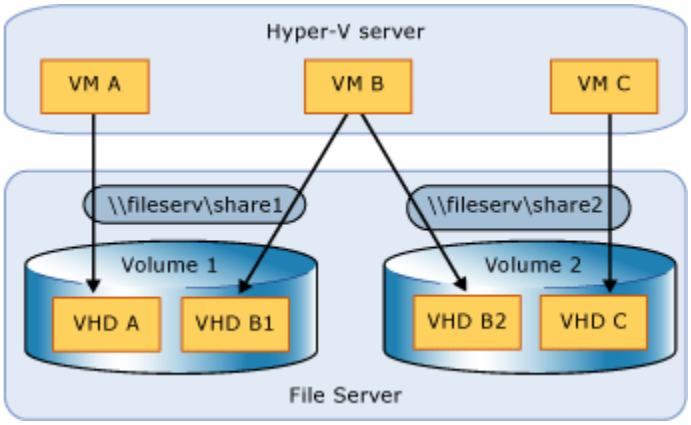

When a backup operator performs a shadow copy of `VM A`, the Hyper-V VSS writer adds `\\fileserv\share1` to the shadow copy set. When ready, the File Share Shadow Copy Provider sends the shadow copy request to `\\fileserv`. On the file server, the File Share Shadow Copy Agent invokes the local VSS service to perform a shadow copy of `Volume 1`. `Volume 2` is not part of the shadow copy set since only `\\fileserv\share1` was reported by the VSS writer. When the shadow copy sequence is complete, a shadow copy share `\\fileserv\share1@{GUID}` is available for the backup application to move the backup data. When the backup is complete, the backup application releases the shadow copy set, and the associated shadow copies and shadow copy shares are removed.

If the backup operator performs a shadow copy of `VM B`, the Hyper-V VSS writer reports both `\\fileserv\share1` and `\\fileserv\share2` in the shadow copy set. On the file server, this results in a shadow copy of both `Volume 1` and `Volume 2`, and two shadow copy shares `\\fileserv\share1@{GUID}` and `\\fileserv\share2@{GUID}` are created.

If the backup operator performs a shadow copy of `VM A` and `VM B`, the Hyper-V VSS writer reports both `\\fileserv\share1` and `\\fileserv\share2` in the shadow copy set. On the file server, a shadow copy of both volumes is generated and two shadow copy shares are created.

Two servers running Hyper-V and a file server cluster

In this scenario, there are two servers running Hyper-V and a file server cluster. As shown in the diagram below, the file server cluster has two cluster nodes, `node1` and `node2`. The file server cluster, `\\fs1`, is currently online on `node1`, with a single share, `\\fs1\share`, on `Volume 1`. To use both cluster nodes, a second file server cluster, `\\fs2`, is configured and is currently online on `node2`, with a single share, `\\fs2\share`, on `Volume 2`.

Diagram of two servers running Hyper-V and a file server cluster

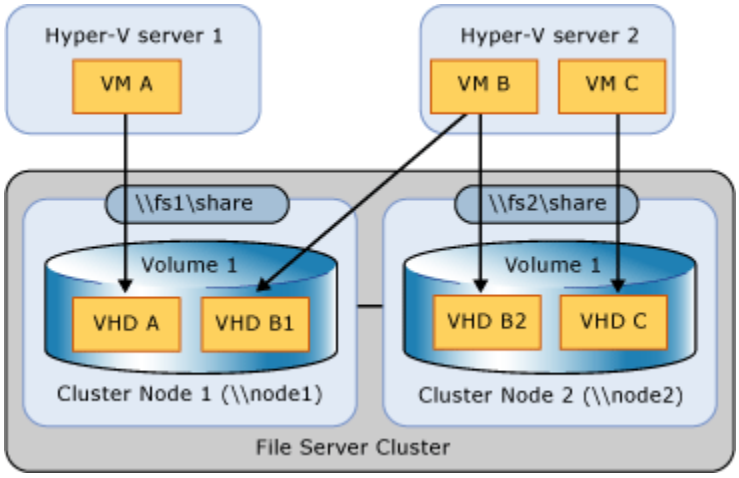

When the backup operator performs a shadow copy of `VM A`, the Hyper-V VSS writer reports `\\fs1\share` in the shadow copy set. When ready, the File Share Shadow Copy Provider sends a shadow copy request to `\\fs1`. As part of the exchange between the File Share Shadow Copy Provider and the File Share Shadow Copy Agent, the agent notifies the provider of the physical computer name, `node1`, which is actually performing the shadow copy.

On `node1`, the File Share Shadow Copy Agent calls the local VSS service to perform a shadow copy of the volume that backs the file share. When the shadow copy sequence is complete, a shadow copy share `\\node1\share@{GUID}` is available for the backup application to stream the backup data. Note the shadow copy share, `\\node1\share@{GUID}`, is targeted to the cluster node, `node1`, and not the virtual computer name, `\\fs1`.

Once the backup is complete, the backup application releases the shadow copy set and the associated shadow copies and shadow copy shares are removed. If the file server cluster moves to, or fails over to, `node2` before the backup sequence is complete, the shadow copy share and the shadow copy become invalid. If the file server cluster moves back to `node1`, the shadow copy and the corresponding shadow copy share will become valid again.

Step 1: Install File Server VSS Agent Service

Use one of the following procedures to install File Server VSS Agent Service.

To install File Server VSS Agent Service using Server Manager

1. From the Add Roles and Features Wizard, under Server Roles, select **File and Storage Services** if it has not already been installed.
2. Under **File and iSCSI Services**, select **File Server** and **File Server VSS Agent Service**.
3. On the **Select Features** page, click **Next**.
4. On the **Confirmation** page, verify that **File Server** and **File Server VSS Agent Service**

are listed, and then click **Install**.

The next part of Step 1 is where you include the Windows PowerShell equivalent to the GUI procedure. Delete this part if there is no PowerShell equivalent.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Provide the Windows PowerShell code for the procedure. For example:

```
PS C:\> Import-Module ServerManager

PS C:\> Add-WindowsFeature -name File-Services,FS-VSS-Agent
```

Step 2: Add a user to the Backup Operators local group on the file server

The user who is performing the shadow copy must have the backup administrative privileges on the remote file servers that are included in the shadow copy set. Usually, this is accomplished by adding the user to the Backup Operators group on the file servers. Use one of the following procedures to do this.

To add a user to the Backup Operators local group on the file server using Server Manager

1. From the dashboard in Server Manager, click **Tools**, and then click **Computer Management**.
2. In Computer Management, expand **Local Users and Groups**, and then expand **Groups**.
3. Double-click **Backup Operators**.
4. On the **Backup Operators Properties** page, click **Add**.
5. Type of the name of the user you want to add to the Backup Operators group, and click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

In the example below, modify the user account and the name of the file server to apply to your environment.

```
$objUser = [ADSI]("WinNT://domain/user")

$objGroup = [ADSI]("WinNT://fileserv/Backup Operators")

$objGroup.PSBase.Invoke("Add",$objUser.PSBase.Path)
```

Step 3: Perform a shadow copy

You must have a VSS-aware backup application that supports VSS for SMB File Shares to perform a shadow copy of an application's data that is stored on a remote file share.

Note
Windows Server Backup in Windows Server 2012 does not support VSS for SMB File Shares.

This section contains examples of performing a shadow copy of a virtual machine that has its data files stored on an SMB file share using DiskShadow.exe (a tool that exposes the functionality offered by VSS), and Microsoft System Center Data Protection Manager 2012 Service Pack 1. For information on how to backup and protect data using Microsoft System Center Data Protection Manager 2012 SP1, see Managing Hyper-V computers.

To perform a shadow copy of virtual machine on the Hyper-V host (HV1) using DiskShadow

1. Start Windows PowerShell. Right-click the PowerShell icon on the taskbar, and select **Run as Administrator**.

   Run the following Windows PowerShell command:

   ```
   PS C:\Users\administrator> DISKSHADOW
   ```

2. Type the following:

   ```
   DISKSHADOW> Writer Verify {66841cd4-6ded-4f4b-8f17-fd23f8ddc3de}
   DISKSHADOW> Set Context Persistent
   DISKSHADOW> Set MetaData vm1backup.cab
   DISKSHADOW> Begin Backup
   DISKSHADOW> Add Volume \\smbsofs\vm\vm1
   DISKSHADOW> Create
   ```

   The following output appears:

   ```
   Alias VSS_SHADOW_1 for shadow ID {7b53b887-76e5-4db8-821d-
   6828e4cbe044} set as environment variable.
   Alias VSS_SHADOW_SET for shadow set ID {2bef895d-5d3f-4799-8368-
   f4bfc684e95b} set as environment variable.
   Querying all shadow copies with the shadow copy set ID
   {2bef895d-5d3f-4799-8368-f4bfc684e95b}
   * Shadow copy ID = {7b53b887-76e5-4db8-821d-6828e4cbe044}
   %VSS_SHADOW_1%
   ```

```
- Shadow copy set: {2bef895d-5d3f-4799-8368-f4bfc684e95b}
%VSS_SHADOW_SET%
- Original count of shadow copies = 1
- Original volume name: \\SMBSOFS\VM\ [volume not on this
machine]
- Creation time: 5/30/2012 5:35:52 PM
- Shadow copy device name: \\FSF-260403-09\VM@{F1C5E17A-4168-
4611-9CD4-8366F9F935C3}
- Originating machine: FSF-260403-09
- Service machine: CONTOSO.SMBTEST.stbtest.microsoft.com
- Not exposed
- Provider ID: {89300202-3cec-4981-9171-19f59559e0f2}
- Attributes: No_Auto_Release Persistent FileShare
Number of shadow copies listed: 1
```

Type the following to end the backup operation:

```
DISKSHADOW> End Backup
```

In this example, the following command parameters are used:

Writer Verify. This parameter specifies that the backup or restore operation must fail if the writer or the component is not included. For more information, see Invoking diskshadow to back up a Virtual Machine from a Hyper-V Host.

Set Context Persistent. This parameter sets the shadow copy to be persistent, which means that the user or the application controls when to delete the shadow copy.

Set MetaData. This parameter stores the metadata information for the shadow copy (which is needed for restore) in the specified file.

Add Volume. This parameter adds the UNC path to the shadow copy set. You can specify multiple paths by repeating the Add Volume parameter.

Create. This parameter initiates the shadow copy. When the shadow copy is created, DiskShadow generates the properties of the shadow copy. The shadow Copy device name is the path for the shadow copy data, which can be copied to the backup store using XCOPY or other tools.

Note

During the backup session, you can see the virtual machine status reporting **Backing up..** in Hyper-V Manager. The backup session starts with the Create parameter, and then ends with the End Backup command in the DiskShadow sequence above.

To copy the backup data to an alternate location

After the shadow copy is complete, you can view the shadow copy share (which is the shadow copy device name from above), and then copy the data you want to back up to an alternate location.

1. Start Windows PowerShell. Right-click the PowerShell icon on the taskbar, and select **Run as Administrator**.

   Run the following Windows PowerShell command:

```
PS C:\Users\administrator> Get-ChildItem -Recurse -Path "\\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}"

Directory: \\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}

Mode LastWriteTime Length Name

---- ------------- ------ ----

d---- vm1

Directory: \\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}\vm1

Mode LastWriteTime Length Name

---- ------------- ------ ----

d---- vm1

-a--- vm1.vhd

Directory: \\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}\vm1\vm1

Mode LastWriteTime Length Name

---- ------------- ------ ----

d---- Virtual Machines

Directory: \\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}\vm1\vm1\Virtual Machines

Mode LastWriteTime Length Name

---- ------------- ------ ----

d---- 87B27972-46C2-406B-87A4-C3FFA1FB6822

-a--- 28800 87B27972-46C2-406B-87A4-C3FFA1FB6822.xml

Directory: \\FSF-260403-09\VM@{F1C5E17A-4168-4611-9CD4-8366F9F935C3}\vm1\vm1\Virtual

Machines\87B27972-46C2-406B-87A4-C3FFA1FB6822

Mode LastWriteTime Length Name

---- ------------- ------ ----

-a--- 2147602688 87B27972-46C2-406B-87A4-C3FFA1FB6822.bin
```

```
-a--- 20971520 87B27972-46C2-406B-87A4-C3FFA1FB6822.vsv
```

2.

To delete the shadow copy

When the backup data is copied, you can delete the shadow copy.

1. Start Windows PowerShell. Right-click the PowerShell icon on the taskbar, and select **Run as Administrator**.

   Run the following Windows PowerShell command:

   ```
   PS C:\Users\administrator> DISKSHADOW
   ```

2. Type the following:

   ```
   DISKSHADOW> Delete Shadows Volume \\smbsofs\vm
   Deleting shadow copy {7b53b887-76e5-4db8-821d-6828e4cbe044} on
   volume \\SMBSOFS\VM\ from provider {89300202-3cec-4981-91
   71-19f59559e0f2} [Attributes: 0x04400009]...
   Number of shadow copies deleted: 1
   ```

To restore data from a shadow copy

Start Windows PowerShell. Right-click the PowerShell icon on the taskbar, and select **Run as Administrator**.

Run the following Windows PowerShell command:

```
PS C:\Users\administrator> DISKSHADOW
DISKSHADOW> Set Context Persistent
DISKSHADOW> Load MetaData vm1backup.cab
DISKSHADOW> Begin Restore
DISKSHADOW> //xcopy files from backup store to the original
location
DISKSHADOW> End Restore
```

The Load MetaData parameter loads the metadata information for the shadow copy to restore data from the specified file. After running the Begin Restore parameter, you can copy the virtual machine files from the backup store to the original location (\\smbsofs\vm\vm1). For more information, see this Microsoft TechNet blog.

See also

- File and Storage Services Overview
- Server Message Block overview

Install and Deploy Windows Server 2012 R2 and Windows Server 2012

Use the articles in this section to learn how to install and deploy Windows Server 2012 R2 and Windows Server 2012.

Windows Server 2012 R2

- System Requirements and Installation Information for Windows Server 2012 R2
  This document provides information about installing the Windows Server 2012 R2 operating system, including important steps to take prior to installation, supported upgrade paths for the release, and information about installing versions distributed as virtual hard disk (VHD) files.
- **Release Notes: Important Issues in Windows Server 2012 R2**
  These release notes address the most critical issues and information about the Windows Server 2012 R2 operating system.
- Upgrade Options for Windows Server 2012 R2
  This document summarizes key information about the supported upgrade paths from previously licensed retail versions of Windows Server to Windows Server 2012 R2.
- Features Removed or Deprecated in Windows Server 2012 R2
  These features and functionalities in Windows Server 2012 R2 have either been removed from the product in the current release or are planned for potential removal in subsequent releases (“deprecated”).
- Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012
  Windows Server 2012 R2 and Windows Server 2012 feature the new Modern user interface. This topic helps you find and open common management tools, create shortcuts to frequently used programs, run programs with elevated privileges, and perform common tasks such signing in and out, restarting, and shutting down computers that are running Windows Server 2012 R2 and Windows Server 2012.

Windows Server 2012

- Installing Windows Server 2012
  This document provides information about installing the Windows Server 2012 operating system, including any known issues that you might need to work around before starting an installation. It also provides information that you can use to troubleshoot problems that may occur during the installation.
- Release Notes: Important Issues in Windows Server 2012
  These release notes address the most critical issues and information about the Windows Server 2012 operating system. For information about by-design changes, new features, and fixes in this release, see documentation and announcements from the specific feature teams.
- Evaluation Versions and Upgrade Options for Windows Server 2012
  This document summarizes key information about evaluation versions of Windows Server

2012, including where to obtain them, the limits on their use, and how to convert them to full retail versions. It also summarizes the supported upgrade paths from previously licensed retail versions of Windows Server to Windows Server 2012.

- Windows Server Installation Options
  This document summarizes the differences between the installation options available for Windows Server 2012, including the features that are installed with each option, the management options available after installation, and how to switch between the installation options during use.
- Server Core and Full Server Integration Overview
  In Windows Server 2012, the Server Core installation option is no longer an irrevocable selection that is made during setup. An administrator now has the flexibility to change to a Server Core installation or a full, GUI-based installation as needed, after operating system installation is finished.
- Configure and Manage Server Core Installations
  This collection of topics provides the information needed to install and deploy Server Core servers; install, manage, and uninstall server roles and features; and manage the server locally or remotely. It also includes a quick reference table of common tasks and the commands for accomplishing them locally on a Server Core server.
- Features Removed or Deprecated in Windows Server 2012
  This is a list of features and functionalities in Windows Server 2012 that have either been removed from the product in the current release or are planned for potential removal in subsequent releases. The list is intended for IT pros who are updating operating systems in a commercial environment.
- Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012
  Windows 8 and Windows Server 2012 feature the new Windows user interface. This topic helps you find and open common management tools, create shortcuts to frequently used programs, run programs with elevated user rights, and perform common tasks such as signing in and out, restarting, and shutting down computers that are running Windows Server 2012 and Windows 8.

System Requirements and Installation Information for Windows Server 2012 R2

This topic addresses the information you need to install Windows Server® 2012 R2. The process of moving to Windows Server 2012 R2 might vary greatly depending on which operating system you are starting with and the pathway you take. We use the following terms to distinguish among different actions, any of which could be involved in a new Windows Server 2012 R2 deployment.

- **Installation** is the basic concept of getting the new operating system on your hardware. Specifically, a **clean installation** requires deleting the previous operating system.
- **Upgrade** means moving from your existing operating system to Windows Server 2012 R2, while staying on the same hardware. You can upgrade from an evaluation version of

Windows Server 2012 R2, a current retail version of certain editions of Windows Server 2012 R2 (an action known as “license conversion”), a previous retail version of certain editions of Windows Server, or from a volume-licensed edition of Windows Server 2012 R2. Even while staying on the same hardware, some server roles might require specific steps to ensure a smooth upgrade. For detailed information about upgrading by any of these methods to Windows Server 2012 R2, see upgrade options

- **Migration** means moving from your existing operating system to Windows Server 2012 R2 by transferring to a different set of hardware. Migration, which might vary considerably depending on the server roles you have installed, is discussed in detail at http://technet.microsoft.com/en-us/windowsserver/dn458795.

Clean installation

Before you start a clean installation of Windows Server 2012 R2, you should perform the following steps.

1. Review the system requirements
2. Review preinstallation documentation
3. Obtain the product
4. Perform preinstallation tasks

Review system requirements

The following are estimated system requirements for the Windows Server 2012 R2. If your computer has less than the "minimum" requirements, you will not be able to install this product correctly. Actual requirements will vary based on your system configuration and the applications and features you install.

Important
The highly diverse scope of potential deployments makes it unrealistic to state “recommended” system requirements that would be generally applicable. Consult documentation for each of the server roles you intend to deploy for more details about the resource needs of particular server roles. For the best results, conduct test deployments to determine appropriate system requirements for your particular deployment scenarios.

Processor

Processor performance depends not only on the clock frequency of the processor, but also on the number of processor cores and the size of the processor cache. The following are the processor requirements for this product:

- **Minimum**: 1.4 GHz 64-bit processor

RAM

The following are the estimated RAM requirements for this product:

- **Minimum**: 512 MB

Important

If you create a virtual machine with the minimum supported hardware parameters (1 processor core and 512 MB RAM) and then attempt to install this release on the virtual machine, Setup will fail.

To avoid this, do one of the following:

- Allocate more than 800 MB RAM to the virtual machine you intend to install this release on. Once Setup has completed, you can change the allocation to as little as 512 MB RAM, depending on the actual server configuration.
- Interrupt the boot process of this release on the virtual machine with SHIFT+F10. In the command prompt that opens, use Diskpart.exe to create and format an installation partition. Run **Wpeutil createpagefile /path=C:\pf.sys** (assuming the installation partition you created was C:). Close the command prompt and proceed with Setup.

Disk space requirements

The following are the estimated **minimum** disk space requirements for the system partition.

- **Minimum**: 32 GB

Notes

Be aware that 32 GB should be considered an absolute minimum value for successful installation. This minimum should allow you to install Windows Server 2012 R2 in Server Core mode, with the Web Services (IIS) server role. A server in Server Core mode is about 4 GB smaller than the same server in Server with a GUI mode. For the smallest possible installation footprint, start with a Server Core installation and then completely remove any server roles or features you do not need by using Features on Demand. For more information about Server Core and Minimal Server Interface modes, see Windows Server Installation Options.

The system partition will need extra space for any of the following circumstances:

- If you install the system over a network.
- Computers with more than 16 GB of RAM will require more disk space for paging, hibernation, and dump files.

Other requirements

You also must have the following:

- Gigabit (10/100/1000baseT) Ethernet adapter
- DVD drive (if you intend to install the operating system from DVD media)

The following items are not strictly required, but are necessary for certain features:

- Super VGA (1024 x 768) or higher-resolution monitor
- Keyboard and Microsoft® mouse (or other compatible pointing device)
- Internet access (fees may apply)

Review preinstallation documentation

To ensure that you are aware of any issues that might require workarounds, features that have been removed from the product, and features that have been added to the product, review these topics:

- Features Removed or Deprecated in Windows Server 2012 R2
- Release Notes: Important Issues in Windows Server 2012 R2
- What's New in Windows Server 2012 R2

Obtain the server product

You can obtain evaluation versions of Windows Server 2012 R2 (convertible to retail versions) in ISO or VHD format from the TechNet Evaluation Center.

You can obtain evaluation versions (convertible to retail versions) in ISO or VHD format from the TechNet Evaluation Center. You can also buy retail versions from the Microsoft Store.

Perform preinstallation tasks

Before you install Windows Server 2012 R2, follow the steps in this section to prepare for the installation.

- **Disconnect UPS devices.** If you have an uninterruptible power supply (UPS) connected to your destination computer, disconnect the serial cable before running Setup. Setup automatically attempts to detect devices that are connected to serial ports, and UPS equipment can cause issues with the detection process.
- **Back up your servers.** Your backup should include all data and configuration information that is necessary for the computer to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP) servers. When you perform the backup, be sure to include the boot and system partitions and the system state data. Another way to back up configuration information is to create a backup set for Automated System Recovery.
- **Disable your virus protection software.** Virus protection software can interfere with installation. For example, it can make installation much slower by scanning every file that is copied locally to your computer.
- **Provide mass storage drivers.** If your manufacturer has supplied a separate driver file, save the file to a floppy disk, CD, DVD, or Universal Serial Bus (USB) flash drive in either the root directory of the media or the **amd64** folder. To provide the driver during Setup, on the disk selection page, click **Load Driver** (or press F6). You can browse to locate the driver or have Setup search the media.

- **Be aware that Windows Firewall is on by default.** Server applications that must receive unsolicited inbound connections will fail until you create inbound firewall rules to allow them. Check with your application vendor to determine which ports and protocols are necessary for the application to run correctly.

  For more information about Windows Firewall, see http://go.microsoft.com/fwlink/?LinkID=84639.

Evaluation versions of Windows Server 2012

Evaluation versions are 64-bit only and can be installed with the Server Core option or the Server with a GUI option. For more information about these installation options, how to convert between them, and how to use the Minimal Server Interface and Features on Demand, see http://technet.microsoft.com//library/hh831786.

For all editions, you have 10 days to complete online activation, at which point the evaluation period begins and runs for 180 days. During the evaluation period, a notification on the Desktop displays the days remaining the evaluation period (except in Windows Server 2012 Essentials). You can also run **slmgr.vbs /dlv** from an elevated command prompt to see the time remaining.

Limits of evaluation versions

All evaluation versions are fully functional during the evaluation period, although booting to Safe mode is not available. The Windows Server 2012 Standard and Windows Server 2012 Datacenter editions come with the activation key pre-installed. After the 180-day evaluation period elapses, the server warns you in various ways depending on the edition:

Windows Server 2012 Standard; Windows Server 2012 Datacenter:

- The following warning appears on the Desktop: **Windows License is expired**
- When you log on to Windows, you are prompted with the following options:
  - Activate now
  - Ask me later
- The system shuts down every hour.
- The only updates that can be installed are security updates.
- Event ID 100 from source WLMS “The license period for this installation of Windows has expired. The operating system will shut down every hour.” appears in the Application log.

Windows Server 2012 Essentials: you receive warnings on the Desktop and on the dashboard, but the server does not shut down.

Installing versions distributed as VHDs

In addition to the other distribution channels, Windows Server 2012 R2 is also available as a pre-configured virtual hard disk (VHD) file, which you can obtain from the TechNet Evaluation Center (http://technet.microsoft.com/en-us/evalcenter/dn205286.aspx). VHD files are available with Windows Server 2012 R2 pre-configured in either Server Core mode or Server with a GUI mode

(though you can switch modes after installation at will). The resulting virtual machine has the following characteristics:

- English only
- 2 GB RAM
- 1 CPU
- 80 GB hard drive

To use the VHD distribution, you must have a computer running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2. The Hyper-V server role must be installed.

To install the VHD

1. Download the VHD file.
2. Start Hyper-V Manager. On the **Action** menu, select **Import Virtual Machine**.
3. Navigate to the directory that the virtual machine file was extracted to and select the directory (not the directory where the VHD file is located).
4. Select the **Copy the virtual machine option**.
5. Confirm that the import was successful by checking Hyper-V Manager.
6. Configure the network adapter for the resulting virtual machine: right-click the virtual machine and select **Settings**. In the left pane, click **Network Adapter**. In the menu that appears, select one of the network adapters of the virtualization server, and then click **OK**.
7. Start the virtual machine.

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2013 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

1.0

Upgrade Options for Windows Server 2012 R2

This topic includes information about upgrading to Windows Server® 2012 R2 from a variety of previous operating systems using a variety of methods.

The process of moving to Windows Server 2012 R2 might vary greatly depending on which operating system you are starting with and the pathway you take. We use the following terms to distinguish among different actions, any of which could be involved in a new Windows Server 2012 R2 deployment.

- **Installation** is the basic concept of getting the new operating system on your hardware. Specifically, a **clean installation** requires deleting the previous operating system. For information about installing Windows Server 2012 R2, see System Requirements and Installation Information for Windows Server 2012 R2. For information about installing other versions of Windows Server, see Windows Server Installation and Upgrade.
- **Upgrade** means moving from your existing operating system release to a more recent release, while staying on the same hardware. For example, if your server is running Windows Server 2012, you can upgrade it to Windows Server 2012 R2. You can upgrade from an evaluation version of the operating system to a retail version, from an older retail version to a newer version, or, in some cases, from a volume-licensed edition of the operating system to an ordinary retail edition.
- **License conversion** in some operating system releases, you can convert a particular edition of the release to another edition of the same release in a single step with a simple command and the appropriate license key. We call this “license conversion.” For example, if you are running Windows Server 2012 R2 Standard, you can convert it to Windows Server 2012 R2 Datacenter.
- **Migration** means moving from your existing operating system to Windows Server 2012 R2 by transferring to a different set of hardware. Migration, which might vary considerably depending on the server roles you have installed, is discussed in detail at http://technet.microsoft.com/en-us/windowsserver/dn458795.

Depending on your scenario, you might encounter a variety of different upgrade pathways.

Upgrading previous retail versions of Windows Server to Windows Server 2012 R2

The table below briefly summarizes which already licensed (that is, not evaluation) Windows operating systems can be upgraded to which editions of Windows Server 2012 R2.

Note the following general guidelines for supported paths:

- In-place upgrades from 32-bit to 64-bit architectures are not supported. All editions of Windows Server 2012 R2 are 64-bit only.
- In-place upgrades from one language to another are not supported.
- In-place upgrades from one build type (fre to chk, for example) are not supported.

- If the server is a domain controller, see http://technet.microsoft.com/library/hh994618.aspx for important information.
- Upgrades from pre-release versions of Windows Server 2012 R2 are not supported. Perform a clean installation to Windows Server 2012 R2.
- Upgrades that switch from a Server Core installation to the Server with a GUI mode of Windows Server 2012 R2 in one step (and vice versa) are not supported. However, after upgrade is complete, Windows Server 2012 R2 allows you to switch freely between Server Core and Server with a GUI modes. For more information about these installation options, how to convert between them, and how to use the new Minimal Server Interface and Features on Demand, see http://technet.microsoft.com/en-us/library/hh831786.

If you do not see your current version in the left column, upgrading to this release of Windows Server 2012 R2 is not supported.

If you see more than one edition in the right column, upgrade to either edition from the same starting version is supported.

If you are running:	You can upgrade to these editions:
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2012 R2 Datacenter
Windows Server 2008 R2 Enterprise with SP1	Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter
Windows Server 2008 R2 Standard with SP1	Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter
Windows Web Server 2008 R2 with SP1	Windows Server 2012 R2 Standard
Windows Server 2012 Datacenter	Windows Server 2012 R2 Datacenter
Windows Server 2012 Standard	Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter
Hyper-V Server 2012	Hyper-V Server 2012 R2
Windows Storage Server 2012 Standard	Windows Storage Server 2012 R2 Standard
Windows Storage Server 2012 Workgroup	Windows Storage Server 2012 R2 Workgroup

Per-server-role considerations for upgrading

Even in supported upgrade paths from previous retail versions to Windows Server 2012 R2, certain server roles that are already installed might require additional preparation or actions for the role to continue functioning after the upgrade. Consult the specific TechNet Library topics for each server role you intend to install for details of additional steps that might be required.

Converting a current evaluation version to a current retail version

You can convert the evaluation version of Windows Server 2012 R2 Standard to either Windows Server 2012 R2 Standard (retail) or Datacenter (retail). Similarly, you can convert the evaluation version of Windows Server 2012 R2 Datacenter to the retail version.

Before you attempt to convert from evaluation to retail, verify that your server is actually running an evaluation version. To do this, do either of the following:

- From an elevated command prompt, run **slmgr.vbs /dlv**; evaluation versions will include “EVAL” in the output.
- From the **Start** screen, open **Control Panel**. Open **System and Security**, and then **System**. View Windows activation status in the Windows activation area of the System page. Click **View details in Windows activation** for more information about your Windows activation status.

If you have already activated Windows, the Desktop shows the time remaining in the evaluation period.

If the server is running a retail version instead of an evaluation version, see the “Upgrading previous retail versions of Windows Server to Windows Server 2012 R2” section of this topic for instructions to upgrade to Windows Server 2012.

For Windows Server 2012 Essentials: You can convert to the full retail version by entering a retail, volume license, or OEM key in the command slmgr.vbs.

If the server is running an evaluation version of Windows Server 2012 Standard or Windows Server 2012 Datacenter, you can convert it to a retail version as follows:

1. If the server is a domain controller, you cannot convert it to a retail version. In this case, install an additional domain controller on a server that runs a retail version and remove AD DS from the domain controller that runs on the evaluation version. For more information, see http://technet.microsoft.com/en-us/library/hh994618.aspx.
2. Read the license terms.
3. From an elevated command prompt, determine the current edition name with the command **DISM /online /Get-CurrentEdition**. Make note of the edition ID, an abbreviated form of the edition name. Then run **DISM /online /Set-Edition:<edition ID> /ProductKey:XXXXX-XXXXX-XXXXX-XXXXX-XXXXX /AcceptEula**, providing the edition ID and a retail product key. The server will restart twice.

For the evaluation version of Windows Server 2012 Standard, you can also convert to the retail version of Windows Server 2012 Datacenter in one step using this same command and the appropriate product key.

Tip

For more information about Dism.exe, see http://go.microsoft.com/fwlink/?LinkId=192466.

Converting a current retail version to a different current retail version

At any time after installing Windows Server 2012, you can run Setup to repair the installation (sometimes called “repair in place”) or, in certain cases, to convert to a different edition.

You can run Setup to perform a “repair in place” on any edition of Windows Server 2012; the result will be the same edition you started with.

For Windows Server 2012 Standard, you can convert the system to Windows Server 2012 Datacenter as follows: From an elevated command prompt, determine the current edition name with the command **DISM /online /Get-CurrentEdition**. Make note of the edition ID, an abbreviated form of the edition name. Then run **DISM /online /Set-Edition:<edition ID> /ProductKey:XXXXX-XXXXX-XXXXX-XXXXX-XXXXX /AcceptEula**, providing the edition ID and a retail product key. The server will restart twice.

Converting a current volume-licensed version to a current retail version

At any time after installing Windows Server 2012, you can freely convert it between a volume-licensed version, a retail version, or an OEM version. The edition remains the same during this conversion.

To do this, from an elevated command prompt, run:

slmgr /ipk <key>

Where <key> is the appropriate volume-license, retail, or OEM product key.

Release Notes: Important Issues in Windows Server 2012 R2

These release notes address the most critical issues and information about the Windows Server 2012 R2 operating system. For information about by-design changes, new features, and fixes in this release, see documentation and announcements from the specific feature teams. For information about important steps to take before installing this release, including issues that you may need to work around, see “System Requirements and Installation Options for Windows Server 2012 R2.” Unless otherwise specified, these notes apply to all editions and installation options of Windows Server 2012 R2.

This document is continuously updated; as critical issues requiring a workaround are discovered, they are added.

Setup on virtual machines

If you create a virtual machine with the minimum supported hardware parameters (1 processor core and 512 MB RAM) and then attempt to install this release on the virtual machine, Setup will fail.

To avoid this, do one of the following:

- Allocate more than 800 MB RAM to the virtual machine you intend to install this release on. Once Setup has completed, you can change the allocation to as little as 512 MB RAM, depending on the actual server configuration.
- Interrupt the boot process of this release on the virtual machine with SHIFT+F10. In the command prompt that opens, use Diskpart.exe to create and format an installation partition. Run **Wpeutil createpagefile /path=C:\pf.sys** (assuming the installation partition you created was C:). Close the command prompt and proceed with Setup.

Add/Remove Features wizard

In two screens, the wizard refers to Windows Server 2012; the text should read “Windows Server 2012 and Windows Server 2012 R2.” There is no impact to functionality.

Internet Explorer 11

- Input Method Editor (IME) input does not accept non-Latin characters with Adobe Flash in Internet Explorer. Only Latin characters can be inserted with an IME. To avoid this, use Internet Explorer for desktop.
- If you visit a website that contains Adobe Flash Player-based content, no PlayTo endpoints are listed in the Devices charm. There is no workaround at this time.
- When you pause Adobe Flash Player-based content in Internet Explorer and switch to Internet Explorer for desktop, the content might not be visible when you return. To avoid this, do not pause content when switching to the desktop. If this has already occurred, try refreshing the web page.

Storage Spaces

If you host a shared VHDX resource on Storage Spaces and then enable deduplication, virtual machines will not boot while the shared VHDX is attached. Do not enable deduplication on volumes that host shared VHDX resources.

Trusts

The Selective Authentication feature of selective trusts is not functional. Access to resources enabled by “Allowed to Authenticate” will fail. There is no workaround at this time.

Windows Server Essentials Experience

These issues affect both Windows Server 2012 R2 Essentials and other editions of Windows Server 2012 R2 when the Essentials Experience is installed with Server Manager.

- The Windows Server Essentials Experience only works in a single-domain environment that does not include a read-only domain controller. There is no workaround at this time.
- If you install Windows Server Essentials as a virtual machine, and if your server is not connected to the network (or the DHCP service is not available) while running the Configure Windows Server Essentials wizard, Anywhere Access functionalities (such as Remote Web Access, virtual private networking, and DirectAccess) are blocked.

  To avoid this, ensure that your server has a network connection when you install Windows Server Essentials as a virtual machine and run the Configure Windows Server Essentials Wizard. If this has already occurred, manually configure the DNS forwarder settings.
- If you enable NIC teaming and then run the Anywhere Access setup wizard, auto-configuration of the router is skipped and you will receive a false “No internet connection” warning on the dashboard.

  To correct this, manually configure the router to forward traffic on port 80/443 to the address of the teamed network interface adapter. If the server is a domain controller and had the DNS server role installed, you should also ensure that DNS on the teamed NiC is set to (‘127.0.0.1’,’::1’) and that DNS forwarders are properly configured.

Work Folders

The client and server must be running the same milestone release for Work Folders to function properly. For example, if the server is running this milestone release of Windows Server 2012 R2, the client must be running the same milestone release of Windows 8.1.

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2013 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

3.0

Features Removed or Deprecated in Windows Server 2012 R2

The following is a list of features and functionalities in Windows Server® 2012 R2 that have either been removed from the product in the current release or are planned for potential removal in subsequent releases (“deprecated”). It is intended for IT professionals who are updating operating systems in a commercial environment. This list is subject to change in subsequent releases and may not include every deprecated feature or functionality. For more details about a particular feature or functionality and its replacement, see the documentation for that feature.

For your quick reference, following table briefly summarizes the status of features that have been removed or deprecated in either Windows Server® 2012 or Windows Server 2012 R2.

Note
This table is necessarily abbreviated; if you see a feature marked for deprecation or removal, please consult the detailed information in this topic or in Features Removed or Deprecated in Windows Server 2012.

Quick reference table

	Windows Server 2012		**Windows Server 2012 R2**	
Feature	Removed	Deprecated	Removed	Deprecated
AD FS v1 Web Agent			x	
AD FS in-place upgrade from AD FS 1.0 or “out of box” AD FS 2.0	x			
AD FS support for “Resource Group”	x			
AD FS support for NT Token mode	x			
AD FS support for using AD LDS as an authentication store	x			
AD RMS license revocation				x
AD RMS SDK				x
Application Server role				x
Built-in drivers for tape drives				x

	Windows Server 2012		Windows Server 2012 R2	
Cluster Automation Server COM API	x (made optional)		x (optional)	
Cluster.exe command-line interface	x (made optional)		x (optional)	
CertObj COM and InetInfo interfaces of the Web Server role				x
Dcpromo.exe		x		
Dfscmd.exe				x
Drivers for Jet Red RDBMS and ODBC		x		x
File Replication Service				x
GAA_FLAG_INCLUDE_TUNNEL_BINDINGORDER flat in GetAdaptersAddresses				x
Internet Information Service (IIS) 6.0 Manager				x
Layered Service Providers		x		x
IscsiVirtualDiskSnapshot and associated WMI methods				x
LPR/LPD protocol		x		x
Namespace for version 1.0 of WMI; WMIC (in WMI)		x	x	
NDIS version 5.0, 5.1, and 5.2 APIs		x	x	
Net DMA	x			
Network Access Protection (NAP)				x
Network Information Service (NIS) and Tools (in RSAT)				x
Nfsshare.exe				x
NFSv2 support				x
Oclist.exe	x			

	Windows Server 2012		Windows Server 2012 R2	
ODBC support for 16- and 32-bit applications and drivers		x		x
ODBC/OLEDB support for Microsoft Oracle		x		
ODBC/OLEDB support for SQL beyond SQL Server 7 and SQL 2000		x		x
Providers for SNMP, Win32_ServerFeature API, Active Directory, MSClus WMI1.0 (in WMI)		x		x
Recovery disk creation			x	
Remote Data Service		x		
Role Collector (Ceiprole.exe) and associated API	x			
SCSIport host-bus adapter	x			
Servermanagercmd.exe		x	x	
SIS Limited API				x
Slmgr.vbs options			x	
SMB 1.0				x
SMB.sys	x			
SMTP and associated management tools		x		x
SQLXMLX		x		x
Storage Explorer snap-in for MMC	x			
Storage Manager for SANs snap-in for MMC	x			
Subsystem for UNIX-based Applications		x	x	
Support for 32-bit cluster resource DLLs	x			
Support for hardware drivers for XDDM	x			
Support for Microsoft SQL Server prior to 7.0	x			
Support for native VGA via the PC/AT BIOS or UEFI CSM	x			
Support for Static VMQ	x			
Support for Token Rings	x			

	Windows Server 2012		Windows Server 2012 R2	
Support for Visual Studio Analyzer 2003 over ODBC, OLEDB, and ADO	x			
System Image Backup (“Windows 7 File Recovery”)			x	
Telnet server				x
VM Chimney (also called TCP Offload) (in Hyper-V)	x			
Windows Server 2003 domain and functional levels of Active Directory				x
Windows Authorization Manager (AzMan)		x	x	
Windows Help executable (WinHlp32.exe)	x			
Windows Identity Foundation 3.5				x
Windows Server Resource Manager		x	x	
Winsock Direct		x	x	
WMI root\virtualization namespace v1 (in Hyper-V)		x	x	
XDR schema elements, XSl pattern feature of MSXML3 (in XML)		x		x

Features removed from Windows Server 2012 R2

The following features and functionalities have been removed from this release of Windows Server 2012 R2. Applications, code, or usage that depend on these features will not function in this release unless you employ an alternate method.

Note
If you are moving to Windows Server 2012 R2 from a server release prior to Windows Server 2012, you should also review Features Removed or Deprecated in Windows Server 2012.

Backup and file recovery

- The File Backup and Restore feature has been removed. Use the File History feature instead.

- System Image Backup (the “Windows 7 File Recovery” feature) has been removed. Instead, use “Reset your PC.”

Drivers

Drivers for tape drives have been removed from the operating system. Instead, use the drivers provided by the manufacturer of your tape drive.

Recovery disk creation

The ability to create a recovery disk on CD or DVD has been removed. Use the Recovery Disk to USB feature.

Slmgr.vbs options

The **/stao** (sets token activation only) and **/ctao** (clears activation tokens only) options of Slmgr.vbs have been removed. These options have been replaced by more flexible activation command options.

Subsystem for UNIX-based Applications

The Subsystem for UNIX-based Applications (SUA) has been removed. If you use the SUA POSIX subsystem with this release, use Hyper-V to virtualize the server. If you use the tools provided by SUA, switch to Cygwin's POSIX emulation, or use either mingw-w64 (available from Sourceforge.net) or MinGW (available from MinGW.org) for doing a native port.

Windows Authorization Manager (AzMan)

Windows Authorization Manager (AzMan) has been removed. You might need to use new management tools for virtual machines or redesign the authorization model.

WMI root\virtualization namespace v1 (used in Hyper-V)

The WMI root\virtualization\v1 namespace has been removed. Instead use the root\virtualization\v2 namespace.

Features deprecated starting with Windows Server 2012 R2

The following features and functionalities are deprecated starting with this release. Eventually, they will be completely removed from the product, but they are still available in this release, sometimes with certain functionality removed. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features.

Active Directory

- The Active Directory Rights Management Services (AD RMS) SDK has been deprecated. To build applications for AD RMS, migrate to AD RMS SDK 2.0, which leverages functionality exposed by the client in Msipc.dll.
- The license revocation functionality in AD RMS is deprecated. Use the protection policy to control the document lifecycle. To remove access to a particular document, set the validity time to “0” in the template, or select **Require a connection to verify a user’s permission** in Microsoft Office. Be aware that both of these options require a connection to a Rights Management Server in order to open the files.
- The File Replication Service (FRS; part of the Active Directory Domain Services role) is deprecated. You should migrate any FRS-based SYSVOLs to use Distributed File System Replication.
- The Windows Server 2003 domain and forest functional levels are deprecated. When you create a new domain or forest, you should consider using a functional level from Windows Server 2008 or newer. When you deploy Windows Server 2012 R2 into an existing Windows Server 2003 environment, you will be notified to move to a newer functional level.

Application Server

The Application Server role is deprecated and will eventually no longer be available as an installable server role. Instead, install individual features and roles separately.

COM and Inetinfo interfaces of the Web Server role

- The IIS CertObj COM interface is deprecated. Use alternate methods for managing certificates.
- The Inetinfo interface is deprecated.

DNS

The GAA_FLAG_INCLUDE_TUNNEL_BINDINGORDER flag in GetAdaptersAddresses is deprecated. There is no specific replacement.

File and storage services

- The SIS Limited API set is deprecated. Once it is removed, backups of SIS-enabled volumes on servers running Windows Server 2008 R2 or earlier operating systems will not be able to be restored to newer servers. If you use this API, you should limit the target operating system to older versions that still support the API.
- Dfscmd.exe is deprecated. Instead, use Windows PowerShell cmdlets for Distributed File System Namespaces or the Dfsutil.exe command set.
- Support for the Network File System version 2 protocol is deprecated. This means that UNMP protocol support, version 2 of the MOUNT protocol, and versions 1, 2, and 3 of the NLM protocol are also deprecated. Migrate to using NFS version 3 or 4.1.

- Nfsshare.exe is deprecated. Instead, use Windows PowerShell cmdlets in scripts for share provisioning.
- The Local Mount driver, related cmdlets (Mount-IscsiVirtualDiskSnapshot, Dismount-IscsiVirtualDiskSnapshot) and WMI methods (see “WMI” section of this document), as well as conversion of VHD files from prior to Windows Server 2012 R2 using Convert-IscsiVirtualDisk are deprecated. Change any scripts that use these cmdlets or WMI methods to use the Export-IscsiVirtualDiskSnapshot cmdlet. Then assign it to an iSCSI target and use it like any other iSCSI virtual disk. Alternately, you can access virtual disks using a local loopback iSCSI initiator.

IIS Manager 6.0

Internet Information Services (IIS) Manager 6.0 is deprecated. It has been replaced by a newer management console.

Networking

Network Access Protection (NAP) is deprecated. Other options for keeping client computers up to date and secure for remote access include DirectAccess, Windows Web Application Proxy, and various non-Microsoft solutions.

Network Information Service (NIS) and Tools (in RSAT)

The Server for Network Information Service (NIS) is deprecated. This includes the associated administration tools in Remote Server Administration Tools (RSAT). Use native LDAP, Samba Client, Kerberos, or non-Microsoft options.

RSAT: Identity management for Unix/NIS

The Server for Network Information Service (NIS) Tools option of Remote Server Administration Tools (RSAT) is deprecated. Use native LDAP, Samba Client, Kerberos, or non-Microsoft options.

SMB

SMB 1.0 is deprecated. Once this is removed, systems running Windows XP or Windows Server 2003 (or older) operating systems will not be able to access file shares. SMB 1.0 has been replaced by SMB 2.0 and newer versions.

Telnet server

Telnet server is deprecated. Instead, use Remote Desktop.

Windows Identity Foundation

Windows Identity Foundation (WIF) 3.5 is deprecated and has been replaced by WIF 4.5. You should start migrating applications that use WIF to WIF 4.5 or Windows .NET Framework 4.5.

SQL Lite

SQL Lite is deprecated. Migrate to alternatives such as SQL LocalDb.

WMI providers and methods

- The WMI version 1 provider for MSCluster is deprecated; it is being replaced by a WMI version 2 provider.
- The WMI methods related to the file server Local Mount driver and related cmdlets (WT_Snapshot.DVMount, WT_SnapShot.DVDismount, WT_Disk.GetDVMountPoints, and the WT_DVMountedPath class) are deprecated. See the “File and storage services” section of this document for more information.

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2013 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

4.1

Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012

Windows Server® 2012 R2, Windows® 8.1, Windows Server® 2012, and Windows® 8 feature the new Modern user interface. This topic helps you find and open common management tools, create shortcuts to frequently-used programs, run programs with elevated privileges, and perform

common tasks like signing in and out, restarting, and shutting down computers that are running Windows Server 2012 R2, Windows 8.1, Windows Server 2012, or Windows 8.

You can install Windows Server 2012 R2 or Windows Server 2012 with a minimal user interface that is well-suited to remote management. For more information, see Windows Server Installation Options.

In this topic:

Open the Start screen

Shut down or restart the computer

Lock the computer or sign out

Close a Windows app

Access Settings for the current screen

Access Control Panel

Access Administrative Tools

Create shortcuts

Open the Run dialog box

Run a program as administrator or as another user

Use common IT Pro tools:

- Open Server Manager
- Start Windows PowerShell
- Open Remote Desktop Connection
- Open Command Prompt
- Open Microsoft Management Console (MMC) and snap-ins

Keyboard shortcuts

Use keyboard shortcuts in a Remote Desktop session

Use keyboard shortcuts in Hyper-V virtual machines

Open the Start screen

The **Start** screen is the home of Windows apps. To open the **Start** screen, use one of these methods:

- Press the Windows logo key. In a virtual machine, you can press **Ctrl**+**Esc**.
- Hover the mouse cursor in the upper right corner of the screen, and then click **Start**.
- On the desktop, hover the mouse cursor in the lower left corner of the screen, and click when the thumbnail of the **Start** screen appears.

Shut down or restart the computer

To shut down the computer

1. Hover the mouse cursor in the upper right corner of the screen, and then click **Settings**.
2. Click **Power**, and then click **Shut down**.

To restart the computer

1. Hover the mouse cursor in the upper right corner of the screen, and then click **Settings**.
2. Click **Power**, and then click **Restart**.

Lock the computer or sign out

To lock the computer

- On the **Start** screen, click your user name in the upper right corner, and then click **Lock**. Alternatively, on the **Start** screen, press the Windows logo key+L.

To sign out from the computer

- On the **Start** screen, click your user name in the upper right corner, and then click **Sign out**.

Close a Windows app

When Windows apps like Internet Explorer are open but not in use, they are automatically minimized and become inactive, freeing resources for other apps. This functionality is similar to a cell phone, tablet PC, or other mobile computer. It is not necessary to close an app, but you can follow steps in this section if you want to do so.

To close a Windows app

- With the app active, hover at the top edge of the screen until the mouse cursor becomes a hand. Click and drag the app to the bottom of the screen, and then release.

Access Settings for the current screen

To open the Settings bar

- Press the Windows logo key+i to open the **Settings** bar for the current screen (for example, **Start**, the desktop, or a Windows app).
- Alternatively, hover the mouse cursor in the upper right corner of the screen, and then click **Settings**.

Access Control Panel

You can access Control Panel from the **Start** screen and from the desktop.

To open Control Panel from the Start screen

- On the **Start** screen, click **Control Panel**.

To open Control Panel from the desktop

1. On the desktop, hover the mouse cursor in the upper right corner of the screen, and then click **Settings**.
2. Click **Control Panel**.

To open common Control Panel tools by using typed commands

1. On the **Start** screen, type one of the following, and then press **Enter** to open common Control Panel management tools.
   - **ncpa.cpl** to open **Control Panel\Network and Internet\Network Connections**.
   - **sysdm.cpl** to open the **System Properties** dialog box that is available at **Control Panel\System and Security\System\Advanced System Settings**.
   - **appwiz.cpl** to open **Control Panel\Programs\Programs and Features\Uninstall or change a program**.
   - **inetcpl.cpl** to open the **Internet Properties** dialog box that is available at **Control Panel\Network and Internet\Internet Options**.

   Tip
   For more Control Panel tool commands that you can type directly on the **Start** screen, see How to run Control Panel tools by typing a command.

To add Control Panel to the desktop

1. Open Control Panel.
2. In the Control Panel **Search** box, type **desktop**.
3. In the Control Panel **Search** results, in **Display**, click **Show or hide common icons on the desktop**.
4. In **Desktop Icon Settings**, select **Control Panel**, and then click **OK**.

Access Administrative Tools

The **Administrative Tools** folder contains links to many common Microsoft Management Console (MMC) snap-ins, including **Computer Management**, **Event Viewer**, and the management tools for installed roles or features.

Tip
The **Start** screen settings (accessed by pressing the Windows logo key+i on the **Start** screen) include an option to allow or prevent tools in the **Administrative Tools** folder

from being displayed in search results and on the **Start** screen. This setting is enabled by default.

To access Administrative Tools in Server Manager

- In the Server Manager menu bar, click **Tools** to access the contents of the **Administrative Tools** folder.

To open the Administrative Tools folder from the Start screen

1. On the **Start** screen, click **Administrative Tools**. You can also type **Administrative Tools** on the **Start** screen, and then click **Administrative Tools** in the list of results.

Tip

If the **Show Administrative tools** setting is turned off, the Administrative Tools folder and its contents will not appear in the **Settings** results. The **Show Administrative tools** setting is available when you hover the mouse cursor over the upper or lower right edge of the **Start** screen, click **Settings**, and then click **Tiles**.

To open the Administrative Tools folder from the Control Panel

- Open **Control Panel**, click **System and Security**, and then click **Administrative Tools**.

Create shortcuts

To create a shortcut on the desktop

1. In File Explorer, navigate to the location of the program for which you want to create a shortcut.
2. Right-click the executable file in the program folder, and then click **Create shortcut**. Because of restricted user access rights, Windows does not allow shortcuts to be created in some folders, including the **Program Files** folder.
3. Drag the shortcut to a folder of your choice, to the desktop, or to the desktop taskbar. If the shortcut cannot be created in the location of the executable file (typically because the access rights for the **Program Files** folder are restricted), and you are prompted to select a location for the shortcut, browse to the **Desktop** folder.

To pin a program to the desktop taskbar from the Start screen

1. On the **Start** screen, search for or navigate to the app that you want to pin to the desktop taskbar.
2. Right-click the app tile, and in the app bar, click **Pin to taskbar**.

To pin a program to the desktop taskbar from File Explorer

1. Open File Explorer.

2. Navigate to the folder where the program that you want to pin to the desktop taskbar is located.
3. Right-click the executable file in the program folder, and then click **Pin to Taskbar**.

To pin a program or folder to the Start screen

1. Open File Explorer.
2. Navigate to the program or folder that you want to pin to the **Start** screen.
3. Right-click the executable file or folder, and then click **Pin to Start**.

Open the Run dialog box

To open the Run dialog box

- On the desktop, press the **Windows logo key+R** to open the **Run** dialog box.

  Alternatively, on the **Start** screen, type **Run**, and then press **Enter**.

Run a program as administrator or as another user

To run a program as administrator from the Start screen

1. On the **Start** screen, navigate to the app that you want to run as Administrator.
2. Right-click the app tile, and in the app bar, click **Run as administrator**.

To run a program as administrator from the desktop

1. Right-click the executable file in File Explorer, or right-click the program shortcut on the desktop.
2. Click **Run as administrator**.

To run a program as another user

- Do one of the following.
  - Add the **Run as a different user** command to the app bar by enabling the following Group Policy setting: `User Configuration/Administrative Templates/Start Menu and Taskbar/Show "Run as different user" command on Start`. To start Local Group Policy Editor, on the **Start** screen, type **gpedit.msc**, and then press **Enter**.
  - Use the **runas** command from a command prompt. For more information about how to use the **runas** command, at a command prompt, type **runas /?**, and then press **Enter**.

Open Server Manager

By default, Server Manager starts when a member of the **Administrators** group signs in to a computer that is running Windows Server 2012 R2 or Windows Server 2012. If Server Manager is not already open, if you are a standard (non-Administrator) user on the server, or if administrators have changed Server Manager default settings so that it does not open automatically at sign-in, open Server Manager by using procedures in this section.

To open Server Manager from the Start screen

- On the **Start** screen, click **Server Manager**.

Note

If the **Show Administrative tools** setting is disabled, the Server Manager tile does not appear on the **Start** screen.

To open Server Manager from the desktop

- On the taskbar, click **Server Manager**.

Start Windows PowerShell

To start Windows PowerShell™ from the Start screen

- On the **Start** screen, click **Windows PowerShell**.

To start Windows PowerShell from the desktop

- On the taskbar, click **Windows PowerShell**.

To start Windows PowerShell with elevated user rights (Run as administrator)

- To run Windows PowerShell as an administrator from the **Start** screen, right-click the **Windows PowerShell** tile, and in the app bar, click **Run as administrator**.
- To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.

Open Remote Desktop Connection

To open Remote Desktop Connection from the Start screen

1. On the **Start** screen, type **mstsc**.
2. In the **Search Apps** results, click **mstsc**.

To open Remote Desktop Connection from the desktop

1. On the desktop, press the **Windows logo key+R** to open the **Run** dialog box.
2. In the **Run** dialog box, type **mstsc**, and then press **Enter**.

Open Command Prompt

To open Command Prompt from the Start screen

1. On the **Start** screen, type **cmd**.
2. In the **Apps** results, click **cmd**.

To open Command Prompt from the desktop

1. On the desktop, press the **Windows logo key+R** to open the **Run** dialog box.
2. In the **Run** dialog box, type **cmd**, and then press **Enter**.

To open Command Prompt with elevated user rights (Run as administrator)

- Do one of the following.
  - To run Command Prompt as an administrator from the **Start** screen on Windows Server 2012 R2, type **cmd**, right-click **Command Prompt** in the **Search** results, and on the shortcut menu, click **Run as administrator**.
  - To run Command Prompt as an administrator from the **Start** screen on Windows Server 2012, type **cmd**, right-click **Command Prompt** in the **Apps** results, and then in the taskbar, click **Run as Administrator**.

Open Microsoft Management Console (MMC) and snap-ins

To open MMC from the Start screen

1. On the **Start** screen, type **mmc**.
2. In the **Apps** or **Search** results, click **mmc**.

To open MMC from the desktop

1. On the desktop, press the **Windows logo key+R** to open the **Run** dialog box.
2. In the **Run** dialog box, type **mmc**, and then press **Enter**.

To open an MMC snap-in from the Start screen

1. On the **Start** screen, type the executable file name of a snap-in.

   Example: Type **gpedit.msc**.
2. When the snap-in is displayed in the **Apps** or **Search** results, click the tile.

To open an MMC snap-in from the desktop

1. On the desktop, press the **Windows logo key+R** to open the **Run** dialog box.
2. Type the executable file name of a snap-in, and then press **Enter**.

   Example: Type **gpedit.msc**.

To open Event Viewer

- Do one of the following.
  a. To open Event Viewer from the **Start** screen, type **eventvwr.msc**, and then press **Enter**.
  b. To open Event Viewer from the desktop, press the **Windows logo key+R** to open the **Run** dialog box, type **eventvwr.msc**, and then press **Enter**.

Tip

Alternatively, you can open a technology-specific snap-in from the **Administrative Tools** folder. For information about how to access the **Administrative Tools** folder, see Access Administrative Tools in this topic.

Keyboard shortcuts

This section provides keyboard shortcuts that are unchanged from Windows 7 and Windows Server 2008 R2, and a table of shortcuts that are new for the Modern user interface in Windows 8 and Windows Server 2012.

Note

Keyboard shortcuts require certain settings and environments to work in a Remote Desktop or virtual machine session. For more information, see Use keyboard shortcuts in a Remote Desktop session and Use keyboard shortcuts in Hyper-V virtual machines in this topic.

Keyboard shortcuts that are unchanged from Windows 7 or Windows Server 2008 R2

Key	Windows 7 or Windows Server 2008 R2 Functionality
Windows logo key	Display or hide the **Start** screen
Windows logo key+left arrow	Dock active window of a desktop app to left half of screen (no effect on Windows Store apps)
Windows logo key+right arrow	Dock active window of a desktop app to right half of screen (no effect on Windows Store apps)
Windows logo key+up arrow	Maximize active desktop app window (no effect

Key	Windows 7 or Windows Server 2008 R2 Functionality
	on Windows Store apps)
Windows logo key+down arrow	Restore or minimize active desktop app window (no effect on Windows Store apps)
Windows logo key+**Shift**+up arrow	Maximize active desktop app window vertically, maintaining width (no effect on Windows Store apps)
Windows logo key+**Shift**+down arrow	Restore or minimize active desktop app window vertically (no effect on Windows Store apps)
Windows logo key+**Shift**+left arrow	Move active desktop app window to monitor on the left (no effect on Windows Store apps)
Windows logo key+**Shift**+right arrow	Move active desktop app window to monitor on the right (no effect on Windows Store apps)
Windows logo key+**P**	Display projection options
Windows logo key+**Home**	Minimize all non-active windows, restore on second keystroke (no effect on Windows apps)
Windows logo key+<number>	Open or switch to the program located at the specified position on the taskbar (Example: Windows logo key+1 to open first program.)
Windows logo key+**Shift**+<number>	Open a new or additional session in the program located at the specified position on the taskbar
Windows logo key+**Ctrl**+**Shift**+<number>	Open a new or additional session of the program located at the specified position on the taskbar, running as Administrator
Windows logo key+**B**	Set focus in the notification area.
Windows logo key+**Break**	Display the System Properties dialog box.
Windows logo key+**D**	Show the desktop, restore on second keystroke (no effect on Windows apps)
Windows logo key+**E**	Open File Explorer to display the **Computer** page
Windows logo key+**Ctrl**+**F**	Search for computers (if you are on a network)
Windows logo key+**G**	Cycle through installed Windows Desktop Gadgets

Key	Windows 7 or Windows Server 2008 R2 Functionality
Windows logo key+**L**	Lock your computer (if you are connected to a network domain), or switch users (if you are not connected to a network domain)
Windows logo key+**M**	Minimize all windows
Windows logo key+**Shift**+**M**	Restore minimized windows to the desktop (no effect on Windows apps)
Windows logo key+**R**	Open the **Run** dialog box.
Windows logo key+**T**	Set focus on the taskbar and cycle through programs
Windows logo key+**Alt**+**Enter**	Open Windows Media Center
Windows logo key+**U**	Open Ease of Access Center
Windows logo key+**X**	Open Windows Mobility Center
Windows logo key+**F1**	Open Windows Help and Support
Windows logo key+**N**	Creates a new note (OneNote)
Windows logo key+**S**	Opens screen clipper (OneNote)
Windows logo key+**Q**	Opens Lync (Lync)
Windows logo key+**A**	Accepts incoming call (Lync)
Windows logo key+**X**	Rejects incoming call (Lync)
Windows logo key+Minus (**-**)	Zoom out (Magnifier)
Windows logo key+Plus (**+**)	Zoom in (Magnifer)
Windows logo key+**Esc**	Close Magnifier

Keyboard shortcuts that are new for Windows Server 2012 R2, Windows 8.1, Windows Server 2012, and Windows 8

Key	Windows Server 2012 R2, Windows 8.1, Windows 8, or Windows Server 2012 Functionality
Windows logo key+**spacebar**	Switch input language and keyboard layout
Windows logo key+**O**	Locks device orientation
Windows logo key+**Y**	Temporarily displays the desktop
Windows logo key+**V**	Cycles through notifications

Key	Windows Server 2012 R2, Windows 8.1, Windows 8, or Windows Server 2012 Functionality
Windows logo key+**Shift**+**V**	Cycles through notifications in reverse order
Windows logo key+**Enter**	Opens Narrator
Windows logo key+**PgUp**	Moves Windows apps to the monitor on the left
Windows logo key+**PgDown**	Moves Windows apps to the monitor on the right
Windows logo key+**Shift**+period (**.**)	Moves the gutter to the left (snaps an application)
Windows logo key+period (**.**)	Moves the gutter to the right (snaps an application)
Windows logo key+**C**	Opens charms bar
Windows logo key+**I**	Opens Settings pane
Windows logo key+**K**	Opens Devices pane
Windows logo key+**H**	Opens Share pane
Windows logo key+**Q**	Opens Search pane
Windows logo key+**W**	Opens Settings Search app
Windows logo key+**F**	Opens File Search app
Windows logo key+**Tab**	Cycles through Windows apps
Windows logo key+**Shift**+**Tab**	Cycles through Windows apps in reverse order
Windows logo key+**Ctrl**+**Tab**	Cycles through Windows apps and snaps them as they are cycled
Windows logo key+**Z**	Opens app bar

Use keyboard shortcuts in a Remote Desktop session

Before you connect to a Remote Desktop (also known as RDP) session, you can configure the session to accept Windows key combinations, whether the session is contained within a window, or occupies the full screen.

To apply keyboard shortcuts to a Remote Desktop session

1. If the Remote Desktop Connection dialog box is not already open, open it by typing **mstsc** on the **Start** screen, and then pressing **Enter**.
2. On the Remote Desktop Connection dialog box, click **Show Options** to display

connection setting tabs.

3. In the **Keyboard** area of the **Local Resources** tab, select one of the following from the **Apply Windows key combinations** drop-down list.
   - To apply keyboard shortcuts to a full-screen Remote Desktop session, select **Only when using the full screen**.
   - To apply keyboard shortcuts to a Remote Desktop session that is contained within a window, select **On the remote computer**.
4. When you are finished configuring other settings for your Remote Desktop session, click **Connect** to connect to the session and start working, or click **Save** on the **General** tab to save your connection settings as an RDP file that you can use for future connections.

Use keyboard shortcuts in Hyper-V virtual machines

Before you start a virtual machine connection, you can apply Windows key combinations to virtual machine connections on a physical host computer by editing the Hyper-V settings for the physical computer in the Hyper-V Manager console.

Note
The setting in this procedure is selected by default if the Hyper-V host computer is running Windows Server 2012 R2 or Windows Server 2012. If the host computer is running Windows 8.1, Windows 8, Windows Server 2008 R2, or Windows Server 2008, you must change the setting to apply Windows key combinations to virtual machine connections.

To apply keyboard shortcuts to new virtual machine connections

1. Open the Hyper-V Manager snap-in if it is not already open.
   - If you are running Remote Server Administration Tools for Windows 8.1 or Windows 8, or you are running Windows Server 2012 R2 or Windows Server 2012, open Server Manager, and then open **Hyper-V Manager** from the **Tools** menu in Server Manager.
   - On the **Start** screen, click **Hyper-V Manager**.
   - If the Hyper-V Manager tile is not on the **Start** screen, type all or part of the name, **Hyper-V Manager** until the Hyper-V Manager tile appears on the **Start** screen.
2. In the tree pane, right-click the physical host computer, and then click **Hyper-V Settings**.
3. In the **User** area of the navigation pane, click **Keyboard** to display keyboard shortcut settings.
4. Select **Use on the virtual machine** to allow new virtual machine connections to accept Windows key combinations from the physical computer. Click **OK** to save your changes and close the Hyper-V Settings dialog box.

   Note
   This setting does not apply to virtual machine connections that are already open.

Installing Windows Server 2012

This document provides information about installing the Windows Server 2012 operating system, including any known issues that you may need to work around before starting an installation. It also provides information that you can use to troubleshoot problems that may occur during the installation. For information about serious known issues that you may need to work around after installation is complete, see the release notes, available at the same location as this document.

Setup works in several stages. You will be prompted for some basic information, and then Setup will copy files and restart the computer. Setup concludes by presenting a menu for Initial Configuration Tasks, which you can use to configure your server for your specific needs.

Preinstallation information

System requirements

The following are estimated system requirements for the Windows Server 2012. If your computer has less than the "minimum" requirements, you will not be able to install this product correctly. Actual requirements will vary based on your system configuration and the applications and features you install.

Processor

Processor performance depends not only on the clock frequency of the processor, but also on the number of processor cores and the size of the processor cache. The following are the processor requirements for this product:

- **Minimum**: 1.4 GHz 64-bit processor

RAM

The following are the estimated RAM requirements for this product:

- **Minimum**: 512 MB

Disk space requirements

The following are the estimated **minimum** disk space requirements for the system partition.

- **Minimum**: 32 GB

  Notes
  Be aware that 32 GB should be considered an absolute minimum value for successful installation. The system partition will need extra space for any of the following circumstances:

  - If you install the system over a network.

- Computers with more than 16 GB of RAM will require more disk space for paging, hibernation, and dump files.

Other requirements

You also need to have the following:

- DVD drive
- Super VGA (800 x 600) or higher-resolution monitor
- Keyboard and Microsoft® mouse (or other compatible pointing device)
- Internet access (fees may apply)

Important information for x64-based operating systems

Ensure that you have updated and digitally signed kernel-mode drivers for Windows Server 2012

If you install a Plug and Play device, you may receive a warning if the driver is not digitally signed. If you install an application that contains a driver that is not digitally signed, you will not receive an error during Setup. In both cases, Windows Server 2012 will not load the unsigned driver.

If you are not sure whether the driver is digitally signed, or if you are unable to boot into your computer after the installation, use the following procedure to disable the driver signature requirement. This procedure enables your computer to start correctly, and the unsigned driver will load successfully.

To disable the signature requirement for the current boot process:

1. Restart the computer and during startup, press F8.
2. Select **Advanced Boot Options**.
3. Select **Disable Driver Signature Enforcement**.
4. Boot into Windows® and uninstall the unsigned driver.

For more information, see http://go.microsoft.com/fwlink/?LinkId=66577.

Before you start Setup

Before you install Windows Server 2012, follow the steps in this section to prepare for the installation.

- **Disconnect UPS devices.** If you have an uninterruptible power supply (UPS) connected to your destination computer, disconnect the serial cable before running Setup. Setup automatically attempts to detect devices that are connected to serial ports, and UPS equipment can cause issues with the detection process.
- **Back up your servers.** Your backup should include all data and configuration information that is necessary for the computer to function. It is important to perform a backup of configuration information for servers, especially those that provide network infrastructure, such as Dynamic Host Configuration Protocol (DHCP) servers. When you perform the backup, be sure to include the boot and system partitions and the system state data. Another

way to back up configuration information is to create a backup set for Automated System Recovery.

- **Disable your virus protection software.** Virus protection software can interfere with installation. For example, it can make installation much slower by scanning every file that is copied locally to your computer.
- **Provide mass storage drivers.** If your manufacturer has supplied a separate driver file, save the file to a floppy disk, CD, DVD, or Universal Serial Bus (USB) flash drive in either the root directory of the media or the **amd64** folder. To provide the driver during Setup, on the disk selection page, click **Load Driver** (or press F6). You can browse to locate the driver or have Setup search the media.
- **Be aware that Windows Firewall is on by default.** Server applications that must receive unsolicited inbound connections will fail until you create inbound firewall rules to allow them. Check with your application vendor to determine which ports and protocols are necessary for the application to run correctly.

  For more information about Windows Firewall, see http://go.microsoft.com/fwlink/?LinkID=84639.

Supported upgrade paths

The table below briefly summarizes supported upgrade paths to Windows Server 2012. For more details on upgrade paths, important caveats to upgrade, and additional information about evaluation versions, see http://go.microsoft.com/fwlink/?LinkId=260917.

If you do not see your current version in the left column, upgrading to this release of Windows Server 2012 is not supported.

If you see more than one edition in the right column, upgrade to either edition from the same starting version is supported.

If you are running:	You can upgrade to these editions:
Windows Server 2008 Standard with SP2 or Windows Server 2008 Enterprise with SP2	Windows Server 2012 Standard, Windows Server 2012 Datacenter
Windows Server 2008 Datacenter with SP2	Windows Server 2012 Datacenter
Windows Web Server 2008	Windows Server 2012 Standard
Windows Server 2008 R2 Standard with SP1 or Windows Server 2008 R2 Enterprise with SP1	Windows Server 2012 Standard, Windows Server 2012 Datacenter
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2012 Datacenter
Windows Web Server 2008 R2	Windows Server 2012 Standard

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2012 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

2.0

Release Notes: Important Issues in Windows Server 2012

These release notes address the most critical issues and information about the Windows Server® 2012 operating system. For information about by-design changes, new features, and fixes in this release, see documentation and announcements from the specific feature teams. For information about important steps to take before installing this release, including issues that you may need to work around, see Installing Windows Server 2012, a document available at the same location as this document. Unless otherwise specified, these notes apply to all editions and installation options of the Windows Server 2012.

This document is continuously updated; as critical issues requiring a workaround are discovered, they are added.

Upgrade

If you upgrade from a Full installation of Windows Server® 2008 or Windows Server® 2008 R2 to Windows Server 2012 in Server with a GUI mode, and then switch Windows Server 2012 to Server Core mode, conversion back to Server with a GUI mode will fail.

To avoid this, delete these registry keys with the following commands:

reg delete HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\WINEVT\Publishers\{bc2eeeec-b77a-4a52-b6a4-dffb1b1370cb}

reg delete HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\WINEVT\Publishers\{57e0b31d-de8c-4181-bcd1-f70e880b49fc}

reg delete HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\WINEVT\Publishers\{8c9dd1ad-e6e5-4b07-b455-684a9d879900}

After running these commands, restart the upgrade.

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2012 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

1.0

Evaluation Versions and Upgrade Options for Windows Server 2012

This document summarizes key information about evaluation versions of Windows Server® 2012, including where to obtain them, the limits on their use, and how to convert them to full retail versions. It also summarizes the supported upgrade paths from previously licensed retail versions of Windows Server to Windows Server 2012.

Evaluation versions of Windows Server 2012

Fully functional time-limited evaluation versions of Windows Server 2012 are available for the following editions:

- Windows Server 2012 Standard
- Windows Server 2012 Datacenter
- Windows Server® 2012 Essentials

These evaluation versions are 64-bit only and can be installed with the Server Core option or the Server with a GUI option. For more information about these installation options, how to convert between them, and how to use the new Minimal Server Interface and Features on Demand, see http://technet.microsoft.com//library/hh831786.

For all editions, you have 10 days to complete online activation, at which point the evaluation period begins and runs for 180 days. During the evaluation period, a notification on the Desktop displays the days remaining the evaluation period (except in Windows Server 2012 Essentials). You can also run **slmgr.vbs /dlv** from an elevated command prompt to see the time remaining.

Where to find evaluation versions

You can get evaluation versions by the following methods:

- Windows Server 2012 Standard; Windows Server 2012 Datacenter: from the TechNet Evaluation Center (http://technet.microsoft.com/en-US/evalcenter/hh670538.aspx). These editions are also available as pre-created VHDs for use in virtualized environments.
- Windows Server 2012 Essentials: from http://go.microsoft.com/fwlink/p/?LinkId=266340

Limits of evaluation versions

All evaluation versions are fully functional during the evaluation period, although booting to Safe mode is not available. The Windows Server 2012 Standard and Windows Server 2012 Datacenter editions come with the activation key pre-installed. After the 180-day evaluation period elapses, the server warns you in various ways depending on the edition:

Windows Server 2012 Standard; Windows Server 2012 Datacenter:

- The following warning appears on the Desktop: **Windows License is expired**
- When you log on to Windows, you are prompted with the following options:
  - Activate now
  - Ask me later
- The system shuts down every hour.
- The only updates that can be installed are security updates.
- Event ID 100 from source WLMS “The license period for this installation of Windows has expired. The operating system will shut down every hour.” appears in the Application log.

Windows Server 2012 Essentials: you receive warnings on the Desktop and on the dashboard, but the server does not shut down.

Converting evaluation versions of Windows Server 2012 to full retail versions

Most evaluation versions can be converted to full retail versions, but the method varies slightly depending on the edition. Before you attempt to convert the version, verify that your server is actually running an evaluation version. To do this, do either of the following:

1. From an elevated command prompt, run **slmgr.vbs /dlv**; evaluation versions will include “EVAL” in the output.
2. From the **Start** screen, open **Control Panel**. Open **System and Security**, and then **System**. View Windows activation status in the Windows activation area of the System page. Click **View details in Windows activation** for more information about your Windows activation status.

If you have already activated Windows, the Desktop shows the time remaining in the evaluation period.

If the server is running a retail version instead of an evaluation version, see the “Upgrading previous licensed versions” section of this document for instructions to upgrade to Windows Server 2012.

For Windows Server 2012 Essentials: You can convert to the full retail version by entering a retail, volume license, or OEM key in the command slmgr.vbs.

If the server is running an evaluation version of Windows Server 2012 Standard or Windows Server 2012 Datacenter, you can convert it to a retail version as follows:

1. If the server is a domain controller, you cannot convert it to a retail version. In this case, install an additional domain controller on a server that runs a retail version and remove AD DS from the domain controller that runs on the evaluation version. For more information, see http://technet.microsoft.com/en-us/library/hh994618.aspx.
2. Read the license terms.
3. From an elevated command prompt, determine the current edition name with the command **DISM /online /Get-CurrentEdition**. Make note of the edition ID, an abbreviated form of the edition name. Then run **DISM /online /Set-Edition:<edition ID> /ProductKey:XXXXX-XXXXX-XXXXX-XXXXX-XXXXX /AcceptEula**, providing the edition ID and a retail product key. The server will restart twice.

For the evaluation version of Windows Server 2012 Standard, you can also convert to the retail version of Windows Server 2012 Datacenter in one step using this same command and the appropriate product key.

Tip

For more information about Dism.exe, see http://go.microsoft.com/fwlink/?LinkId=192466.

Upgrading previous retail versions of Windows Server to Windows Server 2012

The table below briefly summarizes which already licensed (that is, not evaluation) Windows operating systems can be upgraded to which editions of Windows Server 2012.

Note the following general guidelines for supported paths:

- In-place upgrades from 32-bit to 64-bit architectures are not supported. All editions of Windows Server 2012 are 64-bit only.
- In-place upgrades from one language to another are not supported.
- In-place upgrades from one build type (fre to chk, for example) are not supported.
- If the server is a domain controller, see http://technet.microsoft.com/library/hh994618.aspx for important information.
- Upgrades from pre-release versions of Windows Server 2012 (such as the Release Candidate) are not supported. Perform a clean installation to Windows Server 2012.

- Upgrades that switch from a Server Core installation to the Server with a GUI mode of Windows Server 2012 in one step (and vice versa) are not supported. However, after upgrade is complete, Windows Server 2012 allows you to switch freely between Server Core and Server with a GUI modes. For more information about these installation options, how to convert between them, and how to use the new Minimal Server Interface and Features on Demand, see http://technet.microsoft.com/en-us/library/hh831786.
- If you do not see your current version in the left column, upgrading to this release of Windows Server 2012 is not supported.

  If you see more than one edition in the right column, upgrade to either edition from the same starting version is supported.

If you are running:	You can upgrade to these editions:
Windows Server 2008 Standard with SP2 or Windows Server 2008 Enterprise with SP2	Windows Server 2012 Standard, Windows Server 2012 Datacenter
Windows Server 2008 Datacenter with SP2	Windows Server 2012 Datacenter
Windows Web Server 2008	Windows Server 2012 Standard
Windows Server 2008 R2 Standard with SP1 or Windows Server 2008 R2 Enterprise with SP1	Windows Server 2012 Standard, Windows Server 2012 Datacenter
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2012 Datacenter
Windows Web Server 2008 R2	Windows Server 2012 Standard

Per-server-role considerations for upgrading

Even in supported upgrade paths from previous retail versions to Windows Server 2012, certain server roles that are already installed might require additional preparation or actions for the role to continue functioning after the upgrade. The following table summarizes these considerations.

Server role	Upgrade information
Active Directory	**Active Directory Domain Services (AD DS):** Active Directory domains can be upgraded by upgrading the operating system of existing domain controllers or by replacing existing domain controllers with domain controllers that run Windows Server 2012. For more information, see Deploy Active Directory Domain Services (AD DS) in Your Enterprise (http://go.microsoft.com/fwlink/?LinkId=262195). **Active Directory Lightweight Directory**

Server role	Upgrade information
	Services (AD LDS): For more information, see Upgrading from ADAM to AD LDS (http://go.microsoft.com/fwlink/?LinkId=186351).
Active Directory Federation Services (AD FS)	For more information about upgrading AD FS from Windows Server 2008 R2 to Windows Server 2012, see http://technet.microsoft.com/library/jj647765.aspx
Active Directory Rights Management Services (AD RMS)	You can perform an in-place upgrade from either Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012. After completing upgrade of your operating system for any servers running the AD RMS server role, you will need to run the AD RMS Upgrade wizard to upgrade your AD RMS cluster and ensure consistency. Otherwise, your AD RMS cluster might not be in a consistent state. The only other consideration when upgrading that you should be aware of when upgrading from these versions of the Windows Server operating system is that if the Windows Internal Database (WID) was previously selected for use to support your AD RMS database needs, that configuration will block your upgrade to Windows Server 2012 and require you to take additional steps. To unblock your upgrade to Windows Server 2012 you must first uninstall the AD RMS server role and migrate the existing WID instance to a SQL Server instance. For more information, see http://go.microsoft.com/fwlink/?LinkId=229299.
Fax Server	See http://technet.microsoft.com/library/jj134193.aspx.
File and Storage Services	After you upgrade a Windows Server 2008 R2-based server that has DFS Management installed, you must reinstall the DFS Management Tools on the server. To install the DFS Management Tools, run the following Windows PowerShell cmdlet as an administrator: `Install-WindowsFeature RSAT-DFS-Mgmt-Con`

Server role	Upgrade information
	You can also use the Add Roles and Features Wizard in Server Manager. On the **Features** page of the wizard, expand **Remote Server Administration Tools**, expand **Role Administration Tools**, expand **File Services Tools**, and then select the **DFS Management Tools** check box.
Hyper-V	Shut down and save all virtual machines prior to starting upgrade. For additional information, see http://technet.microsoft.com /library/hh831531
Print and Document Services	See http://technet.microsoft.com/library/jj134150.
Remote Access	Routing and Remote Access Service (RRAS) was a role service in Windows Server operating systems prior to Windows Server 2012 that enabled you to use a computer as an IPv4 or IPv6 router, as an IPv4 network address translation (NAT) router, or as a remote access server that hosted dial-up or virtual private network (VPN) connections from remote clients. Now, that feature has been combined with DirectAccess to make up the Remote Access server role in Windows Server 2012. For information about migration from Windows Server 2008 R2 and other versions prior to Windows Server 2012, see http://technet.microsoft.com/library/hh831423.aspx
Remote Desktop Services	Windows Server® 2008 R2 Remote Desktop Services role services cannot be migrated to Windows Server 2012, however an existing Windows Server 2008 R2 RD Session Host server deployment can be integrated into a Windows Server 2012 RDS deployment. The Windows Server 2012 RD Web Access can be configured to point to an existing Windows Server 2008 R2 RD Session Host server farm. Desktops and RemoteApp programs published on the Windows Server® 2008 R2 RD Session Host server farm can be accessed from a Windows Server 2012 RD Web Access server.The Following steps need to be completed in order to add an existing Windows Server® 2008 R2 RD Session Host server farm to

Server role	Upgrade information
	a Windows Server 2012 Remote Desktop Services deployment: • Populate the TS Web Access Computers Security Group. • Configure Windows Server 2008 R2 Session Host servers with the right certificate to sign RDP files. • Pointing a Windows Server 2012 RD Web Access server to a Windows Server® 2008 R2 RD Session Host server farm.
Volume Activation Services	With Active Directory-based Activation, you do not need an additional host server; your existing domain controllers can support activation clients, with the following limitations: • Active Directory-based Activation cannot be configured on read-only domain controllers. • Active Directory-based Activation cannot be used with non-Microsoft directory services. • AD DS must be at the Windows Server 2012 schema level to store activation objects. Domain controllers running earlier versions of Windows Server can activate clients after their schemas have been updated using the Windows Server 2012 version of Adprep.exe. For more information, see What's New in Active Directory Domain Services Installation and Removal.
Web Server (IIS)	No functionality has been removed or changed. Web applications that work in IIS 7.0 run normally in IIS 8.0.

Converting existing Windows Server 2012 versions

At any time after installing Windows Server 2012, you can run Setup to repair the installation (sometimes called “repair in place”) or, in certain cases, to convert to a different edition.

You can run Setup to perform a “repair in place” on any edition of Windows Server 2012; the result will be the same edition you started with.

For Windows Server 2012 Standard, you can convert the system to Windows Server 2012 Datacenter as follows: From an elevated command prompt, determine the current edition name with the command **DISM /online /Get-CurrentEdition**. Make note of the edition ID, an

abbreviated form of the edition name. Then run **DISM /online /Set-Edition:<edition ID> /ProductKey:XXXXX-XXXXX-XXXXX-XXXXX-XXXXX /AcceptEula**, providing the edition ID and a retail product key. The server will restart twice.

For Windows Server 2012 Essentials, you can run Setup and convert it to Windows Server 2012 Standard by providing the appropriate retail license key.

Windows Server Installation Options

This document summarizes the differences between the installation options available for Windows Server® 2012, including the features that are installed with each option, the management options available after installation, and how to switch between the installation options during use. It also explains the differences between the Server Graphical Shell and the Minimal Server Interface and how to switch between them. In addition, it discusses how to use Features on Demand to further reduce the disk footprint by including the binary files for only the server roles you actually use.

Installation options description

When you install Windows Server 2012, you can choose between **Server Core Installation** and **Server with a GUI**. The “Server with a GUI” option is the Windows Server 2012 equivalent of the Full installation option available in Windows Server 2008 R2. The “Server Core Installation” option reduces the space required on disk, the potential attack surface, and especially the servicing requirements, so we recommend that you choose the Server Core installation unless you have a particular need for the additional user interface elements and graphical management tools that are included in the “Server with a GUI” option. For this reason, the Server Core installation is now the default. Because you can freely switch between these options at any time later, one approach might be to initially install the Server with a GUI option, use the graphical tools to configure the server, and then later switch to the Server Core Installation option.

An intermediate state is possible where you start with a Server with a GUI installation and then remove Server Graphical Shell, resulting in a server that comprises the “Minimal Server Interface,” Microsoft Management Console (MMC), Server Manager, and a subset of Control Panel. See the “Minimal Server Interface” section of this document for more information.

In addition, after installation of either option is complete, you can completely remove the binary files for server roles and features that you do not need, thereby conserving disk space and reducing the attack surface still further. See the “Features on Demand” section of this document for more information.

A server in Minimal Server Interface mode is about 300 MB smaller than the same server in Server with a GUI mode. A server in Server Core mode is about 4 GB smaller than the same server in Server with a GUI mode. For the smallest possible installation footprint, start with a Server Core installation and then completely remove any server roles or features you do not need by using Features on Demand.

If you choose the Server Core Installation option

With this option, the standard user interface (the “Server Graphical Shell”) is not installed; you manage the server using the command line, Windows PowerShell, or by remote methods.

- **User interface:** command prompt (Server Graphical Shell is not installed)
- **Install, configure, uninstall server roles locally:** at a command prompt with Windows PowerShell.
- **Install, configure, uninstall server roles remotely:** with Server Manager, Remote Server Administration Tools (RSAT), or Windows PowerShell.

  Note
  For RSAT, you must use the Windows 8 version.
- **Microsoft Management Console:** not available locally.
- **Desktop Experience:** not available.
- **Server roles available:**
  - Active Directory Certificate Services
  - Active Directory Domain Services
  - DHCP Server
  - DNS Server
  - File Services (including File Server Resource Manager)
  - Active Directory Lightweight Directory Services (AD LDS)
  - Hyper-V
  - Print and Document Services
  - Streaming Media Services
  - Web Server (including a subset of ASP.NET)
  - Windows Server Update Server
  - Active Directory Rights Management Server
  - Routing and Remote Access Server and the following sub-roles:
    - Remote Desktop Services Connection Broker
    - Licensing
    - Virtualization
- **To convert to a Server with GUI installation with Windows PowerShell:** follow the steps in the procedure below.

  Note
  All cmdlets must be run from an elevated Windows PowerShell prompt.

To use Windows PowerShell to convert from a Server Core installation to a Server with a GUI installation

1. Determine the index number for a Server with a GUI image (for example, **SERVERDATACENTER**, not **SERVERDATACENTERCORE**) with `Get-WindowsImage -ImagePath <path to wim>\install.wim`.
2. Run `Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell –Restart –Source c:\mountdir\windows\winsxs`
3. Alternatively, if you want to use Windows Update as the source instead of a WIM file, use this Windows PowerShell cmdlet:

   ```
   Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell –Restart
   ```

If you choose the Server with a GUI option

With this option, the standard user interface and all tools are installed. Server roles and features are installed with Server Manager or by other methods.

- **User interface:** standard graphical user interface (“Server Graphical Shell”). The Server Graphical Shell includes the new Windows 8 shell, but does not include the Windows Store or support for Windows Store apps. To enable support for the Windows Store and Windows Store apps, install the Desktop Experience feature.
- **Install, configure, uninstall server roles locally:** with Server Manager or with Windows PowerShell

  Install, configure, uninstall server roles remotely: with Server Manager, Remote Server, RSAT, or Windows PowerShell
- **Microsoft Management Console:** installed
- **Desktop Experience:** installable with Server Manager or Windows PowerShell
- **To convert to a Server Core installation with Windows PowerShell:** run the following cmdlet:

  ```
  Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -Restart
  ```

If you initially install with the Server with a GUI option and then use the above command to convert to a Server Core installation, you can later revert to a Server with a GUI installation without specifying a source. This is because the necessary files remain stored on the disk, even though they are no longer installed. For more information, and for instructions to completely remove the Server with a GUI files from disk, see the “Features on Demand” section of this document.

If you convert to a Server Core installation, Windows features, server roles, and GUI management tools that require a Server with a GUI installation will be uninstalled automatically. You can specify the `-WhatIf` option in Windows PowerShell to see exactly which features will be affected by the conversion.

Minimal Server Interface

In Windows Server 2012, you can remove the Server Graphical Shell, resulting in the “Minimal Server Interface.” This is similar to a Server with a GUI installation, but Internet Explorer 10,

Windows Explorer, the desktop, and the Start screen are not installed. Microsoft Management Console (MMC), Server Manager, and a subset of Control Panel are still present.

Starting with a Server with a GUI installation, you can convert to the Minimal Server Interface at any time using Server Manager.

Note

When you change any of these options, you will have to restart the server for the change to take effect.

See the table below for a summary of which selections to make in Server Manager (or cmdlets to use in Windows PowerShell) in order to get a given installation state:

To reach the installation state in each column…	Server Core installation option	Minimal Server Interface	Server with a GUI installation option	Desktop Experience feature installed
Select these features in Server Manager:	none	Graphical Management Tools and Infrastructure	Graphical Management Tools and Infrastructure, Server Graphical Shell	Graphical Management Tools and Infrastructure, Server Graphical Shell, Desktop Experience
Or, run the Windows PowerShell install/uninstall commands with these values for the Name parameter:	none	`Server-Gui-Mgmt-Infra`	`Server-Gui-Mgmt-Infra, Server-Gui-Shell`	`Server-Gui-Mgmt-Infra, Server-Gui-Shell, Desktop-Experience`

Features on Demand

In previous versions of Windows, even if a server role or feature was disabled, the binary files for it were still present on the disk, consuming space. In Windows Server 2012, not only can you disable a role or feature, but you can also completely remove its files, a state shown as “removed” in Server Manager or “disabled with payload removed” in Dism.exe. To reinstall a role or feature that been completely removed, you must have access to an installation source.

To completely remove a role or feature, use `-Remove` with the `Uninstall-WindowsFeature` cmdlet of Windows PowerShell. For example, to completely remove Windows Explorer, Internet Explorer, and dependent components, run the following Windows PowerShell command:

```
Uninstall-WindowsFeature Server-Gui-Shell -Remove
```

To install a role or feature that has been completely removed, use the Windows PowerShell `-Source` option of the `Install-WindowsFeature` Server Manager cmdlet. The `-Source` option specifies a path to a WIM image and the index number of the image. If you do not specify a `-Source` option, Windows will use Windows Update by default. Offline VHDs cannot be used as a source for installing roles or features which have been completely removed.

Only component sources from the exact same version of Windows are supported. For example, a component source derived from the Windows Server Developer Preview is not a valid installation source for a server running Windows Server 2012.

To install a removed role or feature using a WIM image, use the steps and Windows PowerShell cmdlets:

1. Run `Get-windowsimage –imagepath <path to wim>\install.wim`, and make note of the index of the Windows Server 2012 image.
2. `Install-WindowsFeature <featurename> -Source wim:<path>:<index>`, where:

   Featurename is the name of the role or feature from Get-WindowsFeature

   Path is the path to the WIM mount point

   Index is the index of the server image from Step 1.

For example: `Install-WindowsFeature <featurename> -Source wim:d:\sources\install.wim:4`

You can also specify a source for servers that are domain members using Group Policy. Navigate to **Computer Configuration > Administrative Templates > System > "Specify settings for optional component installation and component repair"**

Note

Permissions might affect the system's ability to access Windows features for installation over a network. The Trusted Installer process runs within the machine account. If you encounter network access issues, try issuing a net use command (for example, **net use * \\path\to\network** or the cmdlet `New-PSDrive -Name Z -PSProvider Filesystem -Root "\\path\to\network`) to mount the network source and then copy the source locally. Then use the local copy as the installation source.

Practical applications

These examples give you an idea of how you can choose the installation option that might be most appropriate for your deployment needs:

- Server Core installations require approximately 4 GB less space than a Server with a GUI installation. By using Server Core installations on virtual machines, you can achieve a significant space savings by not having to store the GUI files on the virtual machine's disk.
- Servers often have comparatively large amounts of memory and complex disk arrays, both of which can take a significant amount of time to initialize at startup. Because Server Core installations minimize the number of restarts required for updates, the frequency at which disk arrays and memory must be re-initialized is reduced.
- Certain server applications have dependencies on certain Windows services, libraries, applications, and files that are not available in Server Core installations, but the administrator

wants to take advantage of the reduced need for updating typical of Server Core installations. The Minimal Server Interface offers additional compatibility while still maintaining a reduced system footprint (though to a lesser extent than a Server Core installation).

- Features on Demand can be used to reduce the footprint for your virtual machine deployments by removing roles and features that will never be deployed in your virtual machines. Depending on the roles and features used in your virtual machines, it is possible to reduce the size by over 1 GB.

Reference table

This table summarizes which features are available locally depending on which installation option you choose.

	Server Core installation option	**Minimal Server Interface**	**Server with a GUI installation option**	**Desktop Experience feature installed**
Command prompt	available	available	available	available
Windows PowerShell/Microsoft .NET	available	available	available	available
Server Manager	not available	available	available	available
Microsoft Management Console	not available	available	available	available
Control Panel	not available	not available	available	available
Control Panel applets	not available	some available	available	available
Windows Explorer	not available	not available	available	available
Taskbar	not available	not available	available	available
Notification area	not available	not available	available	available
Internet Explorer	not available	not available	available	available
Built-in help system	not available	not available	available	available
Themes	not available	not available	not available	available
Windows 8 Shell	not available	not available	not available	available
Windows Store and support for Windows Store apps	not available	not available	not available	available

	Server Core installation option	Minimal Server Interface	Server with a GUI installation option	Desktop Experience feature installed
Windows Media Player	not available	not available	not available	available

See also

For detailed information about working with a server in Server Core mode, see Configure and Manage Server Core Installations.

1.4

Server Core and Full Server Integration Overview

In Windows Server 2012, the Server Core installation option is no longer an irrevocable selection that is made during setup. In Windows Server 2008 R2 and Windows Server 2008, if your requirements changed, there was no way to convert to a full installation or a Server Core installation without completely reinstalling the operating system. An administrator now has the ability to convert between a Server Core installation and a full installation as needed.

There are several scenarios in which this capability is especially useful:

- An administrator installed and is running a full installation option of Windows Server 2012, but exclusively using the roles that run on a Server Core installation. The administrator can convert the servers to Server Core installations to reduce the image size and increase servicing advantages without having to reprovision all of their servers.
- An administrator installs a Server Core installation and now needs to make a change or troubleshoot something that is not possible with the remote GUI. The administrator may not know how to make the change from the command line or cannot find a command-line equivalent. The administrator can convert the server to a full installation, perform the changes, and then convert it back to a Server Core installation to reduce the image size and maintain servicing advantages.
- An administrator wants to use the GUI for all of the initial configuration steps to make the initial configuration experience as easy as possible, yet wants to reduce the image size and maintain the servicing advantages that a Server Core installation provides. The administrator can install a full installation, configure the server as needed, and then convert it to a Server Core installation.
- An enterprise mandates a single server operating system image, so it cannot use a Server Core installation because it requires two images. Windows Server 2012 integrates the Server

Core installation and the full server installation options. Now the enterprise can use a single server operating system image to deploy full installations of Windows Server 2012 and then convert them to Server Core installations to reduce the image size and provide the servicing advantages that it offers.

Requirements

You need a Server Core installation or a full installation of Windows Server 2012. No special hardware is required.

Technical overview

In Windows Server 2008 R2 and Windows Server 2008, the Server Core installation and the full installation options were selections that an administrator made at the time of installation.

Figure 1 Earlier Windows Server installations

In Windows Server 2012, the installation options are integrated, and three large optional features are provided. An administrator can install or uninstall these options to move between Server Core and full server installations.

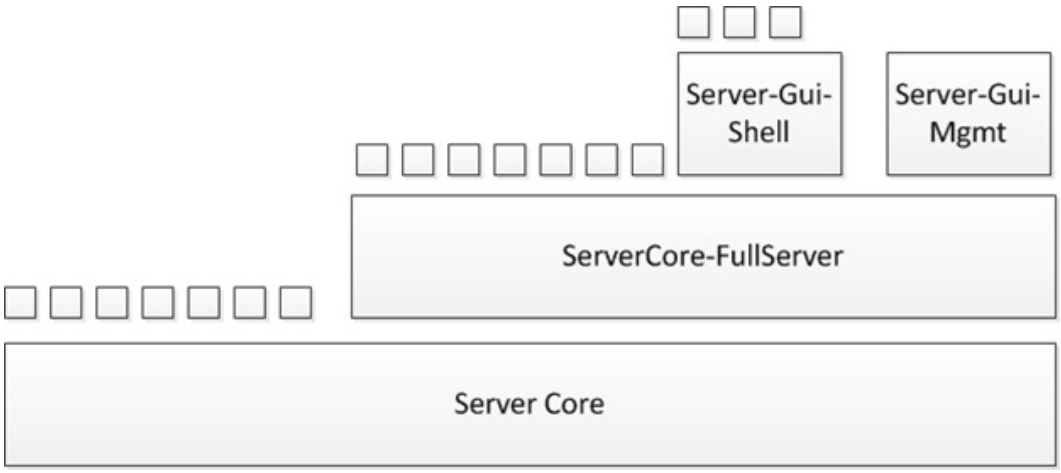

Figure 2 Windows Server 2012 installation options

Configure and Manage Server Core Installations

Because when Windows Server® 2012 is in Server Core mode, the help system and browser are not available, this collection of topics provides the information needed to install and deploy Server Core servers; install, manage, and uninstall server roles and features; configure settings such as server activation and naming, domain membership, and Windows Firewall; installing updates; and managing the server locally or remotely. It also includes a quick reference table of common tasks and the commands for accomplishing them locally on a Server Core server.

If you choose the Server Core option, the standard user interface (the “Server Graphical Shell”) is not installed; you manage the server using remote user interface tools, Windows PowerShell, or, if necessary, locally using the command line (or Windows PowerShell). For more information about the differences between Server Core and Server with a GUI, the features included in each mode, switching between the modes, the Minimal Server Interface mode, and Features on Demand, see http://technet.microsoft.com/library/hh831786.

- Deploy a Server Core Server
- Configure a Server Core Server
- Install Server Roles and Features on a Server Core Server
- Manage a Server Core Server
- Service Updates on a Server Core Server
- Quick Reference for Server Core Tasks

Deploy a Server Core Server

This topic addresses basic information concerning installation and deployment of Windows Server® 2012 in Server Core mode.

When you install Windows Server 2012, you can choose between **Server Core Installation** and **Server with a GUI**. The “Server with a GUI” option is the Windows Server 2012 equivalent of the Full installation option available in Windows Server 2008 R2. The “Server Core Installation” option reduces the space required on disk, the potential attack surface, and especially the requirements for servicing and restarting the server, so we recommend that you choose the Server Core installation unless you have a particular need for the additional user interface elements and graphical management tools that are included in the “Server with a GUI” option. For this reason, the Server Core installation is now the default. Because you can freely switch between these options at any time later, one approach might be to initially install the Server with a GUI option, use the graphical tools to configure the server, and then later switch to the Server Core Installation option.

If you choose the Server Core option, the standard user interface (the “Server Graphical Shell”) is not installed; you manage the server using remote user interface tools, Windows PowerShell, or,

if necessary, locally using the command line (or Windows PowerShell). For more information about the differences between Server Core and Server with a GUI, the features included in each mode, switching between the modes, the Minimal Server Interface mode, and Features on Demand, see http://technet.microsoft.com/library/hh831786.

Initial installation

At the time you run Setup, you have the option to install the server in Server Core mode or Server with a GUI mode.

Using Setup to install the server directly in Server Core mode

1. Insert the appropriate Windows Server 2012 installation media into your DVD drive.
2. When the auto-run dialog box appears, click **Install Now**.
3. Follow the instructions on the screen to complete Setup.
4. You will be prompted to set a password for the Administrator account.

Using an unattend file to install the server directly in Server Core mode

Using an unattend file to install directly in Server Core mode enables you to perform most of the initial configuration tasks during Setup.

To install directly in Server Core mode with an unattend file

1. Create a .xml file titled Unattend.xml by using a text editor or Windows System Image Manager.
2. Copy the Unattend.xml file to a local drive or shared network resource.
3. Boot your computer to Windows Preinstallation Environment (Windows PE).
4. Insert the media disk with the Server Core installation of Windows Server 2012 into your disk drive. If the auto-run Setup window appears, click **Cancel**.
5. Change to the drive that contains the installation media.
6. Run **setup /unattend:<path>\unattend.xml**, where *path* is the path to the Unattend.xml file.
7. Allow Setup to complete.

See also

- Configure a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012

Configure a Server Core Server

This topic explains how to accomplish common server configuration tasks, such as setting passwords, configuring Windows Firewall, joining a domain, and activating the server, all while the server is in Server Core mode.

This topic assumes that you have installed the server in Server Core mode initially and are proceeding directly to these configuration tasks. If you have installed the server in Server with a GUI mode and have since switched to Server Core mode, these steps still apply.

Note
If you close all command prompts, you will have no way to manage the server in Server Core mode. To recover, you can press CTRL+ALT+DELETE, click Start Task Manager, click File, click Start New Task, and then type cmd.exe. Alternatively, you can log off and log back on again.

Note
Because there is no Web browser in Server Core mode, you cannot access the Internet or activate the product through a proxy server that requires users to log on. If you have a proxy server that requires users to log on, you can switch the server to Server with a GUI mode, activate the server, and then switch back to Server Core mode.

Task	Description
1.1. Set the administrative password	Set the administrative password
1.2. Set a static IP address	Set a static IP address
1.3 Join a domain	Join a domain
1.4 Rename the server	Rename the server
1.5 Activate the server	Activate the server
1.6 Configure Windows Firewall	Configure Windows Firewall
1.7. Enable Windows PowerShell remoting	Enable remote use of Windows PowerShell

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Set the administrative password

To set the administrative password

1. When your computer starts for the first time after the installation completes, you will be

prompted to enter a new password.

2. Type an appropriate administrative password.

You can later change the administrative password. To do this, log in and press CTRL+ALT+DELETE, and then choose **Change Password** from the Windows Security menu.

1.2. Set a static IP address

Note

A DHCP address is provided by default. You should perform this procedure only if you need a static IP address.

To view your current network configuration use the `Get-NetIPConfiguration` Windows PowerShell cmdlet.

To view the IP addresses you are already using, use the `Get-NetIPAddress` Windows PowerShell cmdlet.

For a complete reference on Windows PowerShell cmdlets for Net TCP/IP, see http://technet.microsoft.com/en-us/library/hh826150.aspx.

To set a static IP address

1. In Windows PowerShell, run `Get-NetIPInterface`.
2. Make a note of the number shown in the **IfIndex** column of the output for your IP interface or the **InterfaceDescription** string. If your computer has more than one network adapter, make a note of the number or string corresponding to the interface for which you wish to set a static IP address.
3. In Windows PowerShell, run `New-NetIPAddress –InterfaceIndex 12 –IPAddress -192.0.2.2 –PrefixLength 24 –DefaultGateway -192.0.2.1`

   Where:

   `InterfaceIndex` is the value of **IfIndex** from Step 2 (in this example, 12)

   `IPAddress` is the static IP address you intend to set (in this example, 192.0.2.2)

   `PrefixLength` is the prefix length (another form of subnet mask) for the IP address you intend to set (in this example, 24)

   `DefaultGateway` is the default gateway (in this example, 192.0.2.1)
4. In Windows PowerShell, run `Set-DNSClientServerAddress –InterfaceIndex 12 -ServerAddresses 192.0.2.4`

   Where:

   `InterfaceIndex` is the value of **IfIndex** from Step 2

   `ServerAddresses` is the IP address of your DNS server
5. To add multiple DNS servers, run `Set-DNSClientServerAddress –InterfaceIndex 12 -ServerAddresses 192.0.2.4,192.0.2.5`

   Where in this example, 192.0.2.4, 192.0.2.5 are both IP addresses of DNS servers

If you need to switch to using DHCP, use the Windows PowerShell command `Set-DnsClientServerAddress –InterfaceIndex 12 –ResetServerAddresses`.

1.3 Join a domain

To join a domain

1. In Windows PowerShell, run `Add-Computer`. You will be prompted for both credentials to join the domain and the domain name.
2. If you need to add a domain user account to the local Administrators group, either use the Windows PowerShell cmdlets documented at http://technet.microsoft.com/en-us/library/hh826150.aspx, or at a command prompt, run the following command:

   net localgroup administrators /add <DomainName>\<UserName>
3. Restart the computer. You can do this in Windows PowerShell with the command `Restart-Computer`.

1.4 Rename the server

To rename the server

1. Determine the current name of the server with the **hostname** or **ipconfig** command.
2. In Windows PowerShell, run `Rename-Computer`.
3. Restart the computer.

1.5 Activate the server

In Windows PowerShell, run `slmgr.vbs –ipk<productkey>`. Then run `slmgr.vbs –ato`. If activation is successful, no message will return.

Note

You can also activate the server by phone, using a Key Management Service (KMS) server, or remotely. To activate remotely, from a computer that is running Windows Vista, Windows 7, Windows Server 2008, Windows Server 2008 R2, Windows® 8, or Windows Server 2012, use Windows PowerShell to run `cscript windows\system32\slmgr.vbs <ServerName> <UserName> <password>:-ato`.

1.6 Configure Windows Firewall

You can configure Windows Firewall locally on the Server Core computer using Windows PowerShell cmdlets and scripts. See http://technet.microsoft.com/library/hh831755.aspx for documentation of basic Windows Firewall tasks using Windows PowerShell.

1.7. Enable Windows PowerShell remoting

You can enable Windows PowerShell Remoting, in which commands typed in Windows PowerShell on one computer run on another computer. Enable Windows PowerShell Remoting with `Enable-PSRemoting`.

For more information, see http://go.microsoft.com/fwlink/?LinkID=135183

See also

- Install Server Roles and Features on a Server Core Server
- Deploy a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012

Configure a Server Core Server with Sconfig.cmd

In Windows Server 2012, you can use the Server Configuration tool (Sconfig.cmd) to configure and manage several common aspects of Server Core installations. You must be a member of the Administrators group to use the tool.

Sconfig.cmd is available in the Minimal Server Interface and in Server with a GUI mode.

To start the Server Configuration Tool

1. Change to the system drive.
2. Type **Sconfig.cmd**, and then press ENTER. The Server Configuration tool interface opens:

Screenshot of Sconfig.cmd user interface

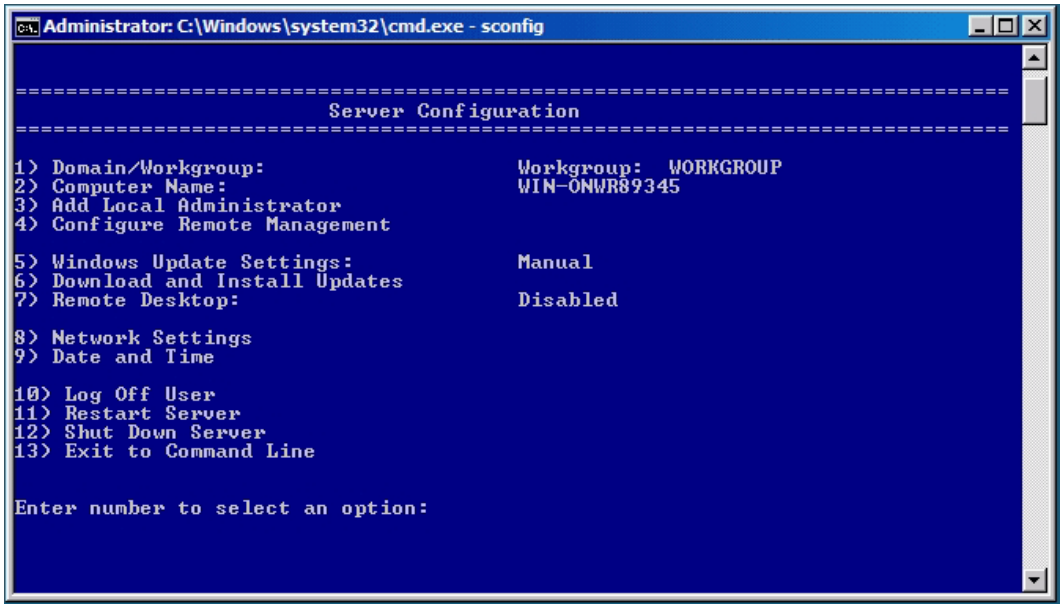

Domain/Workgroup settings

The current Domain/Workgroup settings are displayed in the default Server Configuration tool screen. You can join a domain or a workgroup by accessing the **Domain/Workgroup** settings page from the main menu and following the instructions on the following pages, supplying any required information.

If a domain user has not been added to the local administrators group, you will not be able to make system changes, such as changing the computer name, by using the domain user. To add a domain user to the local administrators group, allow the computer to restart. Next, log on to the computer as the local administrator and follow the steps in the **Local administrator settings** section later in this document.

Note

You are required to restart the server to apply any changes to domain or workgroup membership. However, you can make additional changes and restart the server after all of the changes to avoid the need to restart the server multiple times. By default, running virtual machines are automatically saved prior to restarting the Hyper-V Server.

Computer name settings

The current computer name is displayed in the default Server Configuration Tool screen. You can change the computer name by accessing the “Computer Name” settings page from the main menu and following the instructions.

Note

You are required to restart the server to apply any changes to domain or workgroup membership. However, you can make additional changes and restart the server after all of the changes to avoid the need to restart the server multiple times. By default, running virtual machines are automatically saved prior to restarting the Hyper-V Server.

Local administrator settings

To add additional users to the local administrators group, use the **Add Local Administrator** option on the main menu. On a domain joined machine, enter the user in the following format: domain\username. On a non-domain joined machine (workgroup machine), enter only the user name. The changes take effect immediately.

Network settings

You can configure the IP address to be assigned automatically by a DHCP Server or you can assign a static IP address manually. This option allows you to configure DNS Server settings for the server as well.

Note
These options and many more are now available using the Networking Windows PowerShell cmdlets. For more information, see Network Adapter Cmdlets in the Windows Server Library.

Windows Update settings

The current Windows Update settings are displayed in the default Server Configuration Tool screen. You can configure the server to use Automatic or Manual updates on the **Windows Update Settings** configuration option on the main menu.

When **Automatic Updates** are selected, the system will check for and install updates every day at 3:00 AM. The settings take effect immediately. When **Manual** updates are selected, the system will not check for updates automatically.

At any time, you can download and install applicable updates from the **Download and Install Updates** option on the main menu.

Remote Desktop settings

The current status of remote desktop settings is displayed in the default Server Configuration Tool screen. You can configure the following Remote Desktop settings by accessing the **Remote Desktop** main menu option and following the instructions on screen.

- enable Remote Desktop for Clients running Remote Desktop with Network Level Authentication
- enable Remote Desktop for clients running any version of Remote Desktop
- disable Remote

Date and time settings

You can access and change date and time settings by accessing the **Date and Time** main menu option

To enable remote management

You can enable various remote management scenarios from the **Configure Remote Management** main menu option:

- Microsoft Management Console remote management
- Windows PowerShell
- Server Manager

To log off, restart, or shut down the server

To log off, restart, or shut down the server, access the corresponding menu item from the main menu. These options are also available from the Windows Security menu which can be accessed from any application at any time by pressing CTRL+ALT+DEL.

To exit to the command line

Select the **Exit to the Command Line** option and press ENTER to exit to the command line. To return to the Server Configuration Tool, type **Sconfig.cmd**, and then press ENTER

Install Server Roles and Features on a Server Core Server

When Windows Server 2012 is in Server Core mode, the following server roles are supported:

- Active Directory Certificate Services
- Active Directory Domain Services
- DHCP Server
- DNS Server
- File Services (including File Server Resource Manager)
- Active Directory Lightweight Directory Services (AD LDS)
- Hyper-V
- Print and Document Services
- Streaming Media Services
- Web Server (including a subset of ASP.NET)
- Windows Server Update Server
- Active Directory Rights Management Server

- Routing and Remote Access Server, including the following sub-roles:
  - Remote Desktop Services Connection Broker
  - Licensing
  - Virtualization

When Windows Server 2012 is in Server Core mode, the following server features are supported:

- Microsoft .NET Framework 3.5
- Microsoft .NET Framework 4.5
- Windows PowerShell
- Background Intelligent Transfer Service (BITS)
- BitLocker Drive Encryption
- BitLocker Network Unlock
- BranchCache
- Data Center Bridging
- Enhanced Storage
- Failover Clustering
- Multipath I/O
- Network Load Balancing
- Peer Name Resolution Protocol
- Quality Windows Audio Video Experience
- Remote Differential Compression
- Simple TCP/IP Services
- RPC over HTTP Proxy
- SMTP Server
- SNMP Service
- Telnet client
- Telnet server
- TFTP client
- Windows Internal Database
- Windows PowerShell Web Access
- Windows Process Activation Service
- Windows Standards-based Storage Management
- WinRM IIS extension
- WINS server
- WoW64 support

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Installing and uninstalling server roles and features

To discover the server roles and features on your computer, run the Windows PowerShell cmdlet `Get-WindowsFeature`.

To install a server role or feature, make a note of the role or feature name in the output from `Get-WindowsFeature` and then run `Install-WindowsFeature <featurename>`.

You can install (or uninstall) more than one feature at a time by separating feature names with commas. For example, to install both Server Graphical Shell and Windows PowerShell ISE (and its dependencies), if you run `Install-WindowsFeature Server-Gui-Shell,PowerShell-ISE`, the server is converted to Server with a GUI mode.

Appending the `-WhatIf` parameter to either `Install-WindowsFeature <featurename>` or `Uninstall-WindowsFeature <featurename>` displays any dependent features that will be installed or uninstalled along with the feature you have specified.

To uninstall a server role or feature, run the Windows PowerShell cmdlet `Uninstall-WindowsFeature <featurename>`

Note

When you uninstall a role or feature with this command, the binary files for it remain on the disk. This allows you to add the role or feature later without having to access an installation source (such as an installation DVD or WIM image). To completely remove the files for the role or feature from the disk or to reinstall a role or feature that has been completely removed, see “Working with Features on Demand” in this topic.

Working with Features on Demand

In previous versions of Windows, even if a server role or feature was disabled, the binary files for it were still present on the disk, consuming space. In Windows Server 2012, not only can you disable a role or feature, but you can also completely remove its files, a state shown as “removed” in Server Manager or “disabled with payload removed” in Dism.exe. To reinstall a role or feature that has been completely removed, you must have access to an installation source.

To completely remove a role or feature, use `-Remove` with the `Uninstall-WindowsFeature` cmdlet of Windows PowerShell. For example, to completely remove Windows Explorer, Internet Explorer, and dependent components, run the following Windows PowerShell command:

```
Uninstall-WindowsFeature Server-Gui-Shell -remove
```

To install a role or feature that has been completely removed, use the Windows PowerShell `-Source` option of the `Install-WindowsFeature` Server Manager cmdlet. The `-Source` option specifies a path to a WIM image and the index number of the image. If you do not specify a `-Source` option, Windows will use Windows Update by default.

To install a removed role or feature using a WIM image, use these steps and Windows PowerShell cmdlets:

1. Run `Get-windowsimage –imagepath <path to wim>\install.wim`, and make note of the index of the image for the Server with a GUI version of Windows Server 2012.
2. `Install-WindowsFeature <featurename> -Source wim:<path>:<index>`, where:

   Featurename is the name of the role or feature from Get-WindowsFeature

   Path is the path to the WIM mount point

   Index is the index of the server image from Step 1.

For example, if the Server with a GUI image is in D:\sources: `Install-WindowsFeature <featurename> -Source wim:d:\sources\install.wim:4`

You can also specify a source for servers that are domain members using Group Policy. Navigate to **Computer Configuration > Administrative Templates > System > "Specify settings for optional component installation and component repair"**

Note

Permissions might affect the system's ability to access Windows features for installation over a network. The Trusted Installer process runs within the machine account. If you encounter network access issues, try issuing a net use command (for example, **net use * \\path\to\network**) to connect to the network source and then copy the source locally. Then use the local copy as the installation source.

See also

- Manage a Server Core Server
- Configure a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012

Manage a Server Core Server

This topic explains the various methods available for managing a Windows Server 2012 server in Server Core mode, as well as how to add hardware and manage drivers. You can manage the server in the following ways:

- Locally and remotely using Windows PowerShell
- Remotely using Server Manager
- Remotely using an MMC snap-in
- Remotely with Remote Desktop Services
- Converting the server to Server with a GUI mode

Task	Description
1.1. Manage with Windows PowerShell	You can manage servers in Server Core mode using Windows PowerShell either locally or remotely.
1.2 Manage with Server Manager	By running Server Manager on a remote computer, you can connect to a server that is in Server Core mode.
1.3 Manage with Microsoft Management Console	By using a Microsoft Management Console (MMC) snap-in, you can connect to a server that is in Server Core mode.
1.4 Manage with Remote Desktop Services	By using another computer running Windows, you can use Remote Desktop Services to run scripts and tools on a server that is in Server Core mode.
1.5 Switch to Server with a GUI mode	You can switch the computer to Server with a GUI mode, use the usual user interface tools to accomplish your tasks, and then switch back to Server Core mode.
1.6 Add hardware and manage drivers locally	You can add hardware and manage drivers locally while in Server Core mode.

There are some important limitations and tips to keep in mind when you work with a server in Server Core mode:

- If you close all command prompt windows and want to open a new Command Prompt window, press CTRL+ALT+DELETE, click **Start Task Manager**, click **More Details**, click **File**, click **Run**, and then type **cmd.exe**. Alternatively, you can log off and log back on.
- If you close all Windows PowerShell windows and want to open a new one, press CTRL+ALT+DELETE, open Task Manager, click **More Details**. The File menu opens. In the **File** menu, click **Run new task**, and then start either Cmd.exe or Powershell.exe.
- Any command or tool that attempts to start Windows Explorer will not work. For example, **start .** used from a command prompt will not work.
- There is no support for HTML rendering or HTML help in Server Core mode.
- When in Server Core mode, there are no notifications for activation, new updates, or password expiration because these notifications require the Windows Explorer shell, which is not present in Server Core mode.
- If you need to write a script for managing a server in Server Core mode that requires the secure inclusion of an administrative password, see the scripting column on Microsoft TechNet (http://go.microsoft.com/fwlink/?LinkID=56421).
- Server Core mode supports Windows Installer in quiet mode so that you can install tools and utilities from Windows Installer files.

- When installing Windows Installer packages on a server in Server Core mode, use the **/qb** option to display the basic user interface.
- To change the time zone while in Server Core mode, run `Set-Date`.
- To change international settings while in Server Core mode, run **control intl.cpl**.
- Control.exe will not run on its own. You must run it with either Timedate.cpl or Intl.cpl.
- Winver.exe is not available in Server Core mode. To obtain version information use Systeminfo.exe.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Manage with Windows PowerShell

You can accomplish most management tasks locally or from a remote computer by using Windows PowerShell cmdlets and scripts. To start Windows PowerShell, run **powershell** in a command prompt.

To run Windows PowerShell remotely, see 1.7. Enable Windows PowerShell remoting

For more information about working with Windows PowerShell, see http://technet.microsoft.com/library/hh857343.aspx.

1.2 Manage with Server Manager

Server Manager is a management console in Windows Server® 2012 that helps IT professionals provision and manage both local and remote Windows-based servers from their desktops, without requiring either physical access to servers, or the need to enable Remote Desktop protocol (RDP) connections to each server. Although Server Manager is available in Windows Server 2008 R2 and Windows Server 2008, Server Manager has been completely redesigned for Windows Server 2012, to support remote, multi-server management, and help increase the number of servers an administrator can manage.

Note
You must use the version of Server Manager in Windows Server 2012 or that is available as part of the Remote Server Administration Tools for Windows 8 download package. Older versions of Server Manager will not work.

To enable your local server to be managed by Server Manager running on a remote server, run the Windows PowerShell cmdlet `Configure-SMRemoting.exe -Enable`.

1.3 Manage with Microsoft Management Console

Many snap-ins for Microsoft Management Console (MMC) can be use remotely to manage your Server Core server.

To manage a server in Server Core mode that is a domain member with an MMC snap-in

1. Start an MMC snap-in, such as **Computer Management**.
2. In the left pane, right-click the top of the tree and click **Connect to another computer**. (In the Computer Management example, you would right-click **Computer Management (Local)**.)
3. In **Another computer**, type the computer name of the server that is in Server Core mode and click **OK**. You can now use the MMC snap-in to manage the Server Core server as you would any other computer running a Windows Server operating system.

To manage a server in Server Core mode that is not a domain member with an MMC snap-in

1. If the Server Core server is not a member of a domain, establish alternate credentials to use to connect to the Server Core computer by typing the following command at a command prompt on the remote computer:

   cmdkey /add:<ServerName> /user:<UserName> /pass:<password>

   Where:

   ServerName is the name of the Server Core server

   UserName is the name of an administrator account

   To be prompted for a password, omit the **/pass** option.
2. When prompted, type the password for the user name that is specified in the previous step.
3. If the firewall on the Server Core server is not already configured to allow MMC snap-ins to connect, follow the steps in "To configure Windows Firewall to allow MMC snap-in(s) to connect." Then return to this procedure.
4. On a different computer, start an MMC snap-in, such as Computer Management.
5. In the left pane, right-click the top of the tree and click **Connect to another computer**. (In the Computer Management example, you would right-click **Computer Management (Local)**.)
6. In **Another computer**, type the computer name of the server that is in Server Core mode and click **OK**. You can now use the MMC snap-in to manage the Server Core server as you would any other computer running a Windows Server operating system.

To configure Windows Firewall to allow MMC snap-in(s) to connect

- To allow all MMC snap-ins to connect, run

  ```
  Enable-NetFirewallRule -DisplayGroup "Remote Administration"
  ```
- To allow only specific MMC snap-ins to connect, run:

  ```
  Enable-NetFirewallRule -DisplayGroup "<rulegroup>"
  ```

  Where:

  Rulegroup is one of the values from the table below, depending on which snap-in you

want to connect.

MMC snap-in	Rule group
Event Viewer	Remote Event Log Management
Services	Remote Service Management
Shared Folders	File and Printer Sharing
Task Scheduler	• Performance Logs and Alerts • File and Printer Sharing
Disk Management	Remote Volume Management
Windows Firewall with Advanced Security	Windows Firewall Remote Management

Notes

Some MMC snap-ins do not have a corresponding rule group that allows them to connect through the firewall. However, enabling the rule groups for Event Viewer, Services, or Shared Folders will allow most other snap-ins to connect. Additionally, certain snap-ins require further configuration before they can connect through Windows Firewall:

- **Disk Management.** You must first start the Virtual Disk Service (VDS) on the Server Core computer. You must also configure the Disk Management rules appropriately on the computer that is running the MMC snap-in.
- **IP Security Monitor.** You must first enable remote management of this snap-in. To do this, at a command prompt, type:

  Cscript \windows\system32\scregedit.wsf /im 1
- **Reliability and Performance.** The snap-in does not require any further configuration, but when you use it to monitor a Server Core computer, you can only monitor performance data. Reliability data is not available.

1.4 Manage with Remote Desktop Services

You can use Remote Desktop to manage a server in Server Core mode from remote computers by using Remote Desktop Services.

To manage a Server Core server using Remote Desktop

1. On the Server Core server, run:

   cscript C:\Windows\System32\Scregedit.wsf /ar 0

   This enables the Remote Desktop for Administration mode to accept connections.
2. On another computer, click **Start**, click **Run**, type **mstsc**, and then click **OK**.
3. In **Computer**, enter the name of the server that is in Server Core mode, and click **Connect**.

4. Log on using an administrator account.
5. When the command prompt appears, you can manage the computer using the Windows command-line tools.
6. When you have finished remotely managing the Server Core computer, type **logoff** in the command prompt to end your Remote Desktop session

Note

If you are running the Remote Desktop Services client on a previous version of Windows, you must turn off the higher security level that is set by default in Windows Server 2012. To do this, after step 1, type the following command at the command prompt:**cscript C:\Windows\System32\Scregedit.wsf /cs 0**

1.5 Switch to Server with a GUI mode

There may be situations in which you need to use the graphical user interfaces available in Server with a GUI mode. You can switch the system to Server with a GUI mode by following the steps below, although a restart is required. For more information about the differences between Server Core mode and Server with a GUI mode, as well as information about the Minimal Server Interface and Features on demand, see Windows Server Installation Options.

To convert from Server Core mode to Server with a GUI mode when the server was initially installed in Server with a GUI mode

- Run `Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart`

To convert from Server Core mode to Server with a GUI mode when the server was initially installed in Server Core mode

1. Determine the index number for a Server with a GUI image (for example, **SERVERDATACENTER**, not **SERVERDATACENTERCORE**) using this cmdlet: `Get-WindowsImage -ImagePath <path to wim>\install.wim`
2. Run `Install-WindowsFeature Server-Gui-Mgmt-Infra, Server-Gui-Shell -Restart -Source wim:<path to wim>\install.wim:<Index # from step 1>`
3. Alternatively, if you want to use Windows Update as the source instead of a WIM file, use this Windows PowerShell cmdlet:

   ```
   Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart
   ```

After you have accomplished the management tasks, you can switch the server back to Server Core mode whenever it is convenient (a restart is required) with the Windows PowerShell cmdlet

```
Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -restart
```

1.6 Add hardware and manage drivers locally

To add hardware to a server in Server Core mode

1. Follow the instructions provided by the hardware vendor for installing new hardware:
   - If the driver for the hardware is included in Windows Server 2012, Plug and Play will start and install the driver.
   - If the driver for the hardware is not included, proceed with the steps 2 and 3.
2. Copy the driver files to a temporary folder on the server running a Server Core installation.
3. At a command prompt, open the folder where the driver files are located, and then run the following command:

   pnputil –i –a <driverinf>

   Where:

   driverinf is the file name of the .inf file for the driver.
4. If prompted, restart the computer.

To obtain a list of drivers that are installed

- At a command prompt, run

  sc query type= driver

Note

You must include the space after the equal sign for the command to complete successfully.

To disable a device driver

- At a command prompt, run

  sc delete <service_name>

  Where:

  service_name is the name of the service that you obtain by running **sc query type= driver**

See also

- Service Updates on a Server Core Server
- Install Server Roles and Features on a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012

Service Updates on a Server Core Server

This topic discusses the various ways you can keep a Server Core server up to date with updates. You can service updates in the following ways:

- **Using Windows Update automatically or with WSUS.** By using Windows Update, either automatically or with command-line tools, or Windows Server Update Services (WSUS), you can service servers running a Server Core installation.
- **Manually.** Even in organizations that do not use Windows update or WSUS, you can apply updates manually.

Task	Description
1.1. Manage updates automatically with Windows Update	Managing updates automatically with Windows Update
1.2. Manage updates with WSUS	Managing updates with Windows Server Update Services
1.3. Manage updates manually	Managing updates manually

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Manage updates automatically with Windows Update

To verify the current Windows Update setting, at a command prompt, run the following command:

Cscript scregedit.wsf /AU /v

To enable automatic updates, run the following commands:

Net stop wsuaserv

cscript scregedit.wst /AU 4

net start wsuaserv

To disable automatic updates, run the following commands:

Net stop wsuaserv

cscript scregedit.wsf /AU 1

net start wsuaserv

If the server is a member of a domain, you can also configure Windows Update using Group Policy. For more information, see http://go.microsoft.com/fwlink/?LinkId=192470. However, when you use this method, only option 4 ("Auto download and schedule the install") is relevant to Server Core installations because of the lack of a graphical interface. For more control over which updates are installed and when, you can use a script which provides a command-line equivalent of most of the Windows Update graphical interface. For information about the script, see http://go.microsoft.com/fwlink/?LinkId=192471.

To force Windows Update to immediately detect and install any available updates, run the following command:

wuauclt /detectnow

Depending on the updates that are installed, you might need to restart the computer, although the system will not notify you of this. To determine if the installation process has completed, use Task Manager to verify that the **Wuauclt** or **Trusted Installer** processes are not actively running. You can also use the methods in the "Viewing installed updates" section to check the list of installed updates.

1.2. Manage updates with WSUS

If the Server Core server is a member of a domain, you can configure it to use a WSUS server with Group Policy. For more information, see http://go.microsoft.com/fwlink/?LinkId=192472.

If the server is not a member of a domain, edit the Registry to configure it to use a WSUS server. For more information, see http://go.microsoft.com/fwlink/?LinkId=192473.

Whenever you configure WSUS settings, be sure to select options that are valid for Server Core installations. For example, since there is no graphical interface, there is no way to receive WSUS notifications. For more control over which updates are installed and when, you can use a script which provides a command-line equivalent of most of the Windows Update graphical interface. For information about the script, see http://go.microsoft.com/fwlink/?LinkId=192471.

1.3. Manage updates manually

To install an update manually, download the update and make it available to the Server Core computer, and then run the following command:

Wusa <update>.msu /quiet

Note

Depending on the updates that are installed, you may need to restart the computer, although the system will not notify you of this.

To uninstall an update manually, download the update and make it available to the Server Core computer, and then run the following command:

Wusa /uninstall <update>.msu /quiet

Note
Depending on the updates that are installed, you may need to restart the computer, although the system will not notify you of this.

To view installed updates, run either of these commands:

systeminfo

wmic qfe list

See also

- Quick Reference for Server Core Tasks
- Manage a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012

Quick Reference for Server Core Tasks

This topic is a collection of common administrative tasks for Server Core servers that you can perform locally or remotely from a command prompt.

Whenever possible, Windows PowerShell commands are used and are styled thus: `Get-NetIPConfiguration`. When no Windows PowerShell equivalent is available, run the indicated command at an ordinary command prompt or a Windows PowerShell prompt. These command are styled thus: **Netdom remove**

Task	Description
1.1. Configuration and installation	Configuration and installation
1.2. Networking and firewall	Networking and firewall
1.3. Updates and error reporting	Updates and error reporting
1.4. Services, processes, and performance	Services, processes, and performance
1.5. Event logs	Event logs
1.6. Disk and file system	Disk and file system
1.7. Hardware	Hardware

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Configuration and installation

Task	Steps
Set the local administrative password	Run: **Net user administrator ***
Join a computer to a domain	Run `Add-Computer` For more information see Configure a Server Core Server
Confirm that the domain has changed	Run: **set**
Remove a computer from a domain	Run: `Remove-Computer`
Add a user to the local Administrators group.	Run: **net localgroup Administrators /add <domain>\<username>**
Remove a user from the local Administrators group	Run: **net localgroup Administrators /delete <domain\username>**
Add a user to the local computer	Run: **net user <domain\user name> /add ***
Add a group to the local computer	Run: **net localgroup <group name> /add**
Change the name of a domain-joined computer	Run `Rename-Computer`
Confirm the new computer name	Run: **set c**
Change the name of a computer in a work group	Run: `Rename-Computer`

Task	Steps
Disable paging file management	Run: **wmic computersystem where name="<computername>" set AutomaticManagedPagefile=False**
Configure the paging file	Run: **wmic pagefileset where name="<path/filename>" set InitialSize=<initialsize>,MaximumSize=<maxsize>** Where: *path/filename* is the path to and name of the paging file *initialsize* is the starting size of the paging file in bytes *maxsize* is the maximum size of the paging file in bytes
Change to a static IP address	See http://technet.microsoft.com/library/hh826123.
Set a static DNS address	See http://technet.microsoft.com/library/hh826123.
Change to a DHCP-provided IP address from a static IP address	See http://technet.microsoft.com/library/hh826123.
Enter a product key	Run `slmgr.vbs -ipk <product key>`, and then run `slmgr.vbs -ato`. For more information, see Configure a Server Core Server.
Activate the server locally	Run **slmgr.vbs -ato**
Activate the server remotely	1. Run **cscript slmgr.vbs –ipk <product key><server name><username><password>** 2. Run **cscript slmgr.vbs -ato <servername> <username> <password>** 3. Retrieve the GUID of the computer with cscript **slmgr.vbs -did** 4. Run **cscript slmgr.vbs -dli <GUID>** 5. Verify that **License status** is set to **Licensed (activated)**.

Note

If Slmgr.vbs fails to activate the computer, the resulting error message advises you to run Slui.exe, along with the error code for more information. Slui.exe is not included in Server Core installations—it is available in <versions>.

1.2. Networking and firewall

Task	Steps
Configure your server to use a proxy server.	Run **netsh Winhttp set proxy <servername>:<port number>** **Note** Servers in Server Core mode cannot access the Internet through a proxy that requires a password to allow connections.
Configure your server to bypass the proxy for Internet addresses.	Run **netsh winttp set proxy <servername>:<port number>bypass-list="<local>"**
Display or modify IPSEC configuration.	Run **netsh ipsec**
Display or modify NAP configuration.	Run **netsh nap**
Display or modify IP-to-physical address translation.	Run **arp**
Display or configure the local routing table.	Run **route**
View or configure DNS server settings.	Run **nslookup**
Display protocol statistics and current TCP/IP network connections	Run **netstat**
Display protocol statistics and current TCP/IP connections using NetBIOS over TCP/IP (NBT).	Run **nbtstat**
Display hops for network connections.	Run **pathping**
Trace hops for network connections.	Run **tracert**
Display the configuration of the multicast router	Run **mrinfo**
Enable remote administration of the firewall.	Run **netsh advfirewall firewall set rule group=”Windows Firewall Remote Management” new enable=yes**

Note

See additional Windows PowerShell cmdlets for networking at http://technet.microsoft.com/library/hh826123, http://technet.microsoft.com/library/jj134956, and http://technet.microsoft.com/library/hh848620.

1.3. Updates and error reporting

Task	Steps
Install an update	Run **wusa <update>.msu /quiet**
List installed updates	Run **systeminfo**
Remove an update	1. Run **expand /f:* <update>.msu c:\test** 2. Navigate to c:\test\ and open <update>.xml in a text editor. 3. In <update>.xml, replace **Install** with **Remove** and save the file. 4. Run **pkgmgr /n:<update>.xml**
Configure automatic updates	• To verify the current setting, run **cscript scregedit.wsf /AU /v** • To enable automatic updates, run **cscript scregedit.wsf /AU 4** • To disable automatic updates, run **cscript scregedit.wsf /AU 1**
Enable error reporting	• To verify the current setting, run **serverWerOptin /query** • To automatically send detailed reports, run **serverWerOptin /detailed** • To automatically send summary reports, run **serverWerOptin /summary** • To disable error reporting, run **serverWerOptin /disable**
Participate in the Customer Experience Improvement Program (CEIP)	• To verify the current setting, run **serverCEIPOptin /query** • To enable CEIP, run **serverCEIPOptin /enable** • To disable CEIP, run **serverCEIPOptin /disable**

1.4. Services, processes, and performance

Task	Steps
List running services	Run `Get-Service`
Start a service	Run `Start-Service`
Stop a service	Run `Stop-Service`
Retrieve a list of running applications and associated processes	See http://technet.microsoft.com/library/dd347650.aspx.
Stop a process forcibly	See http://technet.microsoft.com/library/dd347650.aspx.
Start Task Manager	Run **taskmgr**
Manage the performance counters and logging with commands such as **typeperf**, **logman**, **relog**, **tracerprt**.	See http://go.microsoft.com/fwlink/?LinkId=84872.

1.5. Event logs

Use `Get-EventLog` and related cmdlets; see http://technet.microsoft.com/library/ee176846.

1.6. Disk and file system

Task	Steps
Manage disk partitions	For a complete list of commands, run **diskpart /?**
Manage software RAID	For a complete list of commands, run **diskraid /?**
Manage volume mount points	For a complete list of commands, run **mountvol /?**
Defragment a volume	For a complete list of commands, run **defrag /?**
Convert a volume to the NTFS file system	Run **convert <volume letter> /FS:NTFS**
Compact a file	For a complete list of commands, run **compact /?**
Administer open files	For a complete list of commands, run **openfiles /?**

Task	Steps
Administer VSS folders	For a complete list of commands, run **vssadmin /?**
Administer the file system	For a complete list of commands, run **fsutil /?**
Verify a file signature	For a complete list of commands, run **sigverif /?**
Take ownership of a file or folder	For a complete list of commands, run **icacls /?**

1.7. Hardware

Task	Steps
Add a driver for a new hardware device	1. Copy the driver to a folder at %homedrive%\<driver folder>. 2. Run **pnputil -i -a %homedrive%\<driver folder>\<driver>.inf**
Remove a driver for a hardware device	1. Obtain a list of loaded drivers by running **sc query type= driver** 2. Run **sc delete <service_name>**

See also

- Service Updates on a Server Core Server
- Configure and Manage Server Core Installations
- Windows Server Installation Options
- Evaluation Versions and Upgrade Options for Windows Server 2012
- Windows PowerShell Support for Windows Server 2012

Features Removed or Deprecated in Windows Server 2012

The following is a list of features and functionalities in Windows Server® 2012 that have either been removed from the product in the current release or are planned for potential removal in subsequent releases. It is intended for IT professionals who are updating operating systems in a commercial environment. This list is subject to change in subsequent releases and may not include every deprecated feature or functionality.

Features removed from Windows Server 2012

The following features and functionalities have been removed from this release of Windows Server 2012. Applications, code, or usage that depend on these features will not function in this release unless you employ an alternate method.

Active Directory Federation Services

- Support for applications that use the “NT Token mode” configuration of the web agent is removed. Applications are expected to move to Windows Identity Foundation and use the “Claims to Windows Token Service” to convert a UPN from a SAML token to a Windows token for consumption in the application.
- Support for “Resource Group” is removed. (Resource groups are explained at http://technet.microsoft.com/library/cc753670(WS.10).aspx)
- Support for using Active Directory Lightweight Directory Services (AD LDS) as an authentication store is removed
- You are required to migrate to the AD FS version in Windows Server 2012. In-place upgrade from AD FS 1.0 or “out of box” AD FS 2.0 is not supported.

Server Core components

Oclist.exe has been removed. Instead, use Dism.exe. For documentation on using Dism.exe, see http://technet.microsoft.com/library/dd772580(WS.10).aspx.

Clustering

- The Cluster Automation Server (MSClus) COM application programming interface (API) has been made an optional component called FailoverCluster-AutomationServer which is not installed by default. Cluster programmatic functionality is now provided by the Failover Cluster API and the Failover Cluster WMI provider.
- The Cluster.exe command-line interface has been made an optional component called FailoverCluster-CmdInterface which is not installed by default. Cluster command-line functionality is provided by the Failover Cluster PowerShell cmdlets.
- Support for 32-bit cluster resource DLLs has been deprecated. Use 64-bit versions instead.

Graphics

- Support for hardwire drivers for XDDM has been removed. As a result, XDDM graphics drivers will not load in Windows Server 2012. Instead, you can do any of the following:
  - Use the WDDM “basic display-only driver” included in the operating system.
  - Use a WDDM “display-only driver” provided by a hardware vendor.
  - Use a full WDDM driver provided by a hardware vendor.
- Support for native VGA via the PC/AT BIOS is removed. The WDDM “basic display-only driver” included in the operating system will replace this functionality. In UEFI systems, you may see fewer high-resolution (VESA) modes, but there is no other impact.

Hyper-V

- VM Chimney (also called TCP Offload) has been removed. The TCP chimney will not be available to guest operating systems.
- Support for Static VMQ has been removed. Drivers using NDIS 6.3 will automatically access Dynamic VMQ capabilities that are new in Windows Server 2012.

Networking

- NetDMA has been removed.
- Support for Token Rings has been removed.

Server roles

The Role Collector (Ceiprole.exe) and the associated API (Ceiprole.dll) have been removed. To collect telemetry data on server roles, use Server Manager.

Server Message Block

SMB.sys has been removed. The operating system now uses the Winsock Kernel (WSK) to provide the same functionality.

SQL Server

Versions of Microsoft SQL Server prior to 7.0 are no longer supported. Computers running Windows Server 2012 that connect to computers running SQL Server 6.5 (or earlier) will receive an error message.

Storage

- The Storage Manager for SANs snap-in for MMC has been removed. Instead, manage storage with PowerShell cmdlets and Server Manager.
- The Storage Explorer snap-in for MMC has been removed.
- The SCSIport host-bus adapter driver has been removed. Instead, either use a Storport driver or a different host-bus adapter.

Visual Studio

Support for Visual Studio Analyzer 2003 over ODBC, OLEDB, and ADO has been removed.

Windows Help

The Windows Help program (specifically, WinHlp32.exe, the executable file that opens *.hlp help files) has been removed from Windows since Windows Server 2008. Previously, downloadable packages that provide this functionality were made available (see http://support.microsoft.com/kb/917607). Starting with this release, no download will be provided

to enable you to view *.hlp files on Windows Server 2012. For this milestone release, Windows Help is also currently unsupported in Windows® 8.

Features deprecated starting with Windows Server 2012

The following features and functionalities are deprecated starting with this release. Eventually, they will be completely removed from the product, but they are still available in this release, sometimes with certain functionality removed. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features.

Active Directory

Dcpromo.exe has been deprecated. In Windows Server 2012, if you run dcpromo.exe (without any parameters) from a command prompt, you receive a message directing you to Server Manager, where you can install Active Directory Domain Services using the Add Roles wizard. If you run **dcpromo /unattend** from a command prompt, you can still perform unattended installations that use Dcpromo.exe. This allows organizations to continue to use automated Active Directory Domain Services (AD DS) installation routines based on dcpromo.exe until they can rewrite those routines using Windows PowerShell. For more information, see http://technet.microsoft.com/library/hh472160.aspx.

Database management systems

- ODBC support for 16- and 32-bit applications and drivers is deprecated. Use 64-bit versions instead.
- ODBC/OLEDB support for Microsoft Oracle is deprecated. Migrate to drivers and providers supplied by Oracle.
- Jet Red RDBMS and ODBC drivers are deprecated.

Networking

The Network Driver Interface Specification (NDIS) version 5.0, 5.1, and 5.2 APIs are deprecated. New drivers for NDIS 6.0 must be written.

Hyper-V

- The WMI root\virtualization namespace is deprecated. The new namespace is root\virtualization\v2.
- Windows Authorization Manager (AzMan) is deprecated. You may need to use new management tools for virtual machines or redesign the authorization model.

Printing

The line printer daemon protocol (LPR/LPD) is deprecated. When this feature is eventually removed, clients that print to a server using this protocol, such as UNIX clients, will not be able to connect or print. Instead, UNIX clients should use IPP. Windows clients can connect to UNIX shared printers using the Windows Standard Port Monitor (see http://support.microsoft.com/kb/246868for more information).

Remote Data Service

The Remote Data service is deprecated. Migrate to the Windows Web Services API.

SMTP

SMTP and the associated management tools are deprecated. Though the functionality is still available in Windows Server 2012, you should begin using System.Net.Smtp. With this API, you will not be able to insert a message into a file for pickup; instead configure Web applications to connect on port 25 to another server using SMTP.

Subsystem for UNIX-based Applications

The Subsystem for UNIX-based Applications (SUA) is deprecated. If you use the SUA POSIX subsystem with this release, use Hyper-V to virtualize the server. If you use the tools provided by SUA, switch to Cygwin's POSIX emulation, or use either mingw-w64 (available from Sourceforge.net) or MinGW (available from MinGW.org) for doing a native port.

Transport protocols

- The Transport Driver Interface (TDI) is deprecated. Use Windows Filtering Platform instead.
- Layered Service Providers (LSP) are deprecated. Use Windows Filtering Platform instead.
- Winsock Direct is deprecated. Use Network Direct instead.

SNMP

SNMP is deprecated. Instead, use the Common Information Model (CIM), which is supported by the WS-Management web services protocol and implemented as Windows Remote Management.

SQL Server

- ODBC/OLEDB support for SQL is deprecated for versions beyond SQL Server 7 and SQL 2000. Migrate to SQL Native Client (SNAC) to use features provided by SQL Server 2005, SQL Server 2008, and later versions.
- SQLXMLX is deprecated. Migrate code to use SQLXML.

Windows System Resource Manager

Windows System Resource Manager (WSRM) is deprecated. Similar functionality is provided by Hyper-V.

WMI providers

- The WMI provider for Simple Network Management Protocol (SNMP) is deprecated because the SNMP service is being deprecated.
- The WMI provider for the Win32_ServerFeature API is deprecated.
- The WMI provider for Active Directory is deprecated. Manage Active Directory with PowerShell cmdlets.
- The WMI command-line tool (Wmic) is deprecated. Use PowerShell cmdlets instead.
- The namespace for version 1.0 of WMI is deprecated. Prepare to adapt scripts for a revised namespace.

XML

- XML-Data Reduced (XDR) schema elements are deprecated. Migrate Web applications that use this schema to the W3C Standards-compliant XML schema.
- The XSL pattern feature of MSXML3 is deprecated. Migrate Web applications that use this feature to the W3C Standards-compliant XML Path Language (XPath) feature set.

Copyright

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

©2012 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Hyper-V, MS-DOS, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

All other trademarks are property of their respective owners.

4.1

Copyright attributions

This document provides copyright and credit information for Windows®.

Third Party Notices

This file is based on or incorporates material from the projects listed below (collectively, “Third Party Code”). Microsoft is not the original author of the Third Party Code. The original copyright notice and license, under which Microsoft received such Third Party Code, are set out below. Such licenses and notices are provided for informational purposes only. Microsoft, not the third party, licenses the Third Party Code to you under the terms set forth in the license terms for the Microsoft product. Microsoft reserves all other rights not expressly granted under this agreement, whether by implication, estoppel or otherwise

Copyright

Information in this document, including URL and other Internet website references, is subject to change without notice. Unless otherwise noted, the companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted in examples herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license terms from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2013 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Visual Basic, Visual Studio, Windows, Windows NT, Windows Server, Active Directory, Aero, BitLocker, HotStart, ReadyBoost, ReadyDrive, SideShow, and SuperFetch are either registered trademarks or trademarks of Microsoft group of companies.

Portions of this software are based in part on the work of Spider Systems ® Limited. Because Microsoft has included the Spider Systems Limited software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright 1987 Spider Systems Limited

Copyright 1988 Spider Systems Limited

Copyright 1990 Spider Systems Limited

Portions of this software are based in part on the work of Seagate Software.

Portions of this software are based in part on the work of ACE*COMM Corp. Because Microsoft has included the ACE*COMM Corp. software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright 1995-1997 ACE*COMM Corp

Portions of this software are based in part on the work of Sam Leffler and Silicon Graphics, Inc. Because Microsoft has included the Sam Leffler and Silicon Graphics software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1988-1997 Sam Leffler

Copyright © 1991-1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, providedthat (i) the above copyright notices and this permission notice appear inall copies of the software and related documentation, and (ii) the names ofSam Leffler and Silicon Graphics may not be used in any advertising orpublicity relating to the software without the specific, prior writtenpermission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FORANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND,OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA, OR PROFITS,WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Portions of this software are based in part on the work of Highground Systems. Because Microsoft has included the Highground Systems software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1996-1999 Highground Systems

This product incorporates compression code from the Info-ZIP group. There are no extra charges or costs due to the use of this code, and the original compression sources are freely available from http://www.info-zip.org/ or ftp://ftp.info-zip.org/pub/infozip/src/ on the Internet.

Portions Copyright © 2000 SRS Labs, Inc

This product includes software from the 'zlib' general purpose compression library.Portions of this software are based in part on the work of ScanSoft, Inc. Because Microsoft has included the ScanSoft, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:TextBridge® OCR © by ScanSoft, Inc.

Portions of this software are based in part on the work of University of Southern California. Because Microsoft has included the University of Southern California software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1996 by the University of Southern California

All rights reserved.

Permission to use, copy, modify, and distribute this software and its documentation in source and binary forms for any purpose and without fee is hereby granted, provided that both the above copyright notice and this permission notice appear in all copies, and that any documentation,

advertising materials, and other materials related to such distribution and use acknowledge that the software was developed in part by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.THE UNIVERSITY OF SOUTHERN CALIFORNIA makes no representations about the suitability of this software for any purpose. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Other copyrights might apply to parts of this software and are so noted when applicable.

Portions of this software are based in part on the work of James Kanze. Because Microsoft has included the James Kanze software in this product, Microsoft is required to include the following text that accompanied such software:

COPYRIGHT AND PERMISSION NOTICE

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation. Permission is also given to modify the software to any extend, under the condition that, in the modified software, the prefix "GB_" is changed to something else, and the name directories for includes files ("gb" in this distribution) is also changed.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use, or other dealings in this Software without prior written authorization of the copyright holder.

This product contains software from Cisco ISAKMP Services.

Portions of this software are based in part on the work of RSA Data Security, Inc. Because Microsoft has included the RSA Data Security, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1990, RSA Data Security, Inc. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Portions of this software are based in part on the work of Regents of The University of Michigan. Because Microsoft has included the Regents of The University of Michigan software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 1995, 1996 Regents of The University of Michigan.

All Rights Reserved.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of The University of Michigan not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. This software is supplied as is without expressed or implied warranties of any kind.

Copyright © 1993, 1994 Regents of the University of Michigan.

All rights reserved.

Redistribution and use in source and binary forms are permitted provided that this notice is preserved and that due credit is given to the University of Michigan at Ann Arbor. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. This software is provided ``as is'' without express or implied warranty.

Portions of this software are based in part on the work of Massachusetts Institute of Technology. Because Microsoft has included the Massachusetts Institute of Technology software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright 1989, 1990 by the Massachusetts Institute of Technology. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Under U.S. law, this software may not be exported outside the US without license from the U.S. Commerce department.

Copyright 1994 by the Massachusetts Institute of Technology. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

This product includes software developed by the University of California, Berkeley and its contributors.

Portions of this software are based in part on the work of Autodesk, Inc. Because Microsoft has included the Autodesk, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:

© Copyright 1995 by Autodesk, Inc.

This product contains graphics filter software; this software is based in part on the work of the Independent JPEG Group.

This product includes software from the 'libpng' PNG reference library.

This product includes “True Verb” technology from KS Waves Ltd.

Contains Runtime Modules of Lotus C++ API Toolkit for Notes/Domino. (c) Copyright IBM Corporation 2003. All rights reserved.

Portions of this software are based in part on the work of SGS-Thomson Microelectronics, Inc. Because Microsoft has included the SGS-Thomson Microelectronics, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright 1996 SGS-Thomson Microelectronics, Inc. All Rights Reserved

Portions of this software are based in part on the work of Unicode, Inc. Because Microsoft has included the Unicode, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:

COPYRIGHT AND PERMISSION NOTICE

Copyright © 1991-2005 Unicode, Inc. All rights reserved. Distributed under the Terms of Use in http://www.unicode.org/copyright.html.

Permission is hereby granted, free of charge, to any person obtaining a copy of the Unicode data files and any associated documentation (the "Data Files") or Unicode software and any associated documentation (the "Software") to deal in the Data Files or Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Data Files or Software, and to permit persons to whom the Data Files or

Software are furnished to do so, provided that (a) the above copyright notice(s) and this permission notice appear with all copies of the Data Files or Software, (b) both the above copyright notice(s) and this permission notice appear in associated documentation, and (c) there is clear notice in each modified Data File or in the Software as well as in the documentation associated with the Data File(s) or Software that the data or software has been modified.

THE DATA FILES AND SOFTWARE ARE PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THE DATA FILES OR SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in these Data Files or Software without prior written authorization of the copyright holder.

The Combined PostScript Driver was the result of a cooperative development process by Adobe Systems Incorporated and Microsoft Corporation.

Portions of this software are based in part on the work of Media Cybernetics. Because Microsoft has included the Media Cybernetics software in this product, Microsoft is required to include the following text that accompanied such software:

HALO Image File Format Library © 1991-1992 Media Cybernetics, Inc.

Portions of this software are based in part on the work of Luigi Rizzo. Because Microsoft has included the Luigi Rizzo software in this product, Microsoft is required to include the following text that accompanied such software:

© 1997-98 Luigi Rizzo (luigi@iet.unipi.it)

Portions derived from code by Phil Karn (karn@ka9q.ampr.org), Robert Morelos-Zaragoza (robert@spectra.eng.hawaii.edu) and Hari Thirumoorthy (harit@spectra.eng.hawaii.edu), Aug 1995

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS “AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS BE LIABLE FOR ANY DIRECT, INDIRECT,

INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions of this software are based in part on the work of W3C. Because Microsoft has included the W3C software in this product, Microsoft is required to include the following text that accompanied such software:

W3C ® SOFTWARE NOTICE AND LICENSE

http://www.w3.org/Consortium/Legal/2002/copyright-software-20021231

This work (and included software, documentation such as READMEs, or other related items) is being provided by the copyright holders under the following license. By obtaining, using and/or copying this work, you (the licensee) agree that you have read, understood, and will comply with the following terms and conditions.Permission to copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications:

1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.
2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, the W3C Software Short Notice should be included (hypertext is preferred, text is permitted) within the body of any redistributed or derivative code.
3. Notice of any changes or modifications to the files, including the date changes were made. (We recommend you provide URLs to the location from which the code is derived.)

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

Portions of this software are based in part on the work of Sun Microsystems, Inc. Because Microsoft has included the Sun Microsystems, Inc. software in this product, Microsoft is required to include the following text that accompanied such software:

Sun RPC is a product of Sun Microsystems, Inc. and is provided for unrestricted use provided that this legend is included on all tape media and as a part of the software program in whole or

part. Users may copy or modify Sun RPC without charge, but are not authorized to license or distribute it to anyone else except as part of a product or program developed by the user.

SUN RPC IS PROVIDED AS IS WITH NO WARRANTIES OF ANY KIND INCLUDING THE WARRANTIES OF DESIGN, MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE.

Sun RPC is provided with no support and without any obligation on the part of Sun Microsystems, Inc. to assist in its use, correction, modification or enhancement.

SUN MICROSYSTEMS, INC. SHALL HAVE NO LIABILITY WITH RESPECT TO THE INFRINGEMENT OF COPYRIGHTS, TRADE SECRETS OR ANY PATENTS BY SUN RPC OR ANY PART THEREOF.

In no event will Sun Microsystems, Inc. be liable for any lost revenue or profits or other special, indirect and consequential damages, even if Sun has been advised of the possibility of such damages.

Sun Microsystems, Inc.

2550 Garcia Avenue

Mountain View, California 94043

Manufactured under license from Dolby Laboratories. “Dolby” and the double-D symbol are trademarks of Dolby Laboratories. Confidential unpublished works. Copyright 1992-1997 Dolby Laboratories. All rights reserved.

Contains Adobe® [Flash® and/or Shockwave®] Player technology by Adobe Systems Incorporated. Copyright © 1996 – 2013 Adobe Systems Incorporated. All rights reserved. Adobe, Flash, and Shockwave are trademarks or registered trademarks of Adobe Systems Incorporated in the United States and internationally.

Portions of this software are based in part on the work of Andrei Alexandrescu. Because Microsoft has included the Andrei Alexandrescu software in this product, Microsoft is required to include the following text that accompanied such software:

The Loki Library

Copyright © 2001 by Andrei Alexandrescu

This code accompanies the book: Alexandrescu, Andrei. "Modern C++ Design: Generic Programming and Design Patterns Applied." Copyright © 2001. Addison-Wesley.

Permission to use, copy, modify, distribute and sell this software for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

The author or Addison-Welsey Longman make no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Portions of this software are based in part on the work of the Distributed Management Task Force, Inc. (DMTF). Because Microsoft has included software based on DMTF specifications in this product, Microsoft is required to include the following text:

Copyright © 2007 Distributed Management Task Force, Inc. (DMTF). All rights reserved.

Portions of this work are derived from "The Draft Standard C++ Library" Copyright © 1995 by P.J. Plauger published by Prentice-Hall and are used with permission.

Portions Copyright 2002-2007 Charlie Poole or Copyright 2002-2004 James W. Newkirk, Michael C. Two, Alexei A. Vorontsov or Copyright 2000-2002 Philip A. Craig

Portions of this software are based in part on the work of Hewlett-Packard Company. Because Microsoft has included the Hewlett-Packard Company software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright © 2002, 2003 Hewlett-Packard Company.

About Notice:

This software is based on software available from http://mpvtools.sourceforge.net.

This software processes a format called MPV. MPV is an open specification for managing collections and multimedia playlists of photo, video, and music content and associated metadata and is available at no cost from the Optical Storage Technology Association. More information about the MPV specification can be found at http://www.osta.org/mpv.

Permission Notice:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice, this permission notice, and the above About Notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Portions of International CorrectSpell™ spelling correction system © 1993 by Lernout & Hauspie Speech Products N.V. All rights reserved. The American Heritage® Dictionary of the English Language, Third Edition Copyright © 1992 Houghton Mifflin Company. Electronic version licensed from Lernout & Hauspie Speech Products N.V. All rights reserved.

Portions Copyright © 2005 Gregory Wild-Smith

Authors Website: http://www.twilightuniverse.com/

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s), authors website url, and this permission notice appear in all copies of the Software and that the above copyright notice(s), authors url, and this permission notice appear in supporting documentation.

This product includes software from the Box2D engine developed by Erin Catto (www.box2d.org).

This product uses materials from Bigelow and Holmes

Lucida® typeface(s)

Copyright© 1985, 1986, 1987, 1988 and 1990 by Bigelow & Holmes

U.S. Patents Des. 289,420; Des. 289,421; Des. 289,422; Des. 289,773

Portions of this software are based in part on flex. Because Microsoft has included flex software in this product, Microsoft is required to include the following text that accompanied such software:

Copyright (c) 2001, 2002, 2003, 2004, 2005, 2006, 2007 The Flex Project.

Copyright (c) 1990, 1997 The Regents of the University of California.

All rights reserved.

This code is derived from software contributed to Berkeley by Vern Paxson.

The United States Government has rights in this work pursuant to contract no. DE-AC03-76SF00098 between the United States Department of Energy and the University of California.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

This product uses materials from FastCGI

fastcgi.h

Defines for the FastCGI protocol.

Copyright (c) 1995-1996 Open Market, Inc.

This FastCGI application library source and object code (the "Software") and its documentation (the "Documentation") are copyrighted by Open Market, Inc ("Open Market"). The following terms apply to all files associated with the Software and Documentation unless explicitly disclaimed in individual files.

Open Market permits you to use, copy, modify, distribute, and license this Software and the Documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this Software and Documentation may be copyrighted by their authors and need not follow the licensing terms described here. If modifications to this Software and Documentation have new licensing terms, the new terms must be clearly indicated on the first page of each file where they apply.

OPEN MARKET MAKES NO EXPRESS OR IMPLIED WARRANTY WITH RESPECT TO THE SOFTWARE OR THE DOCUMENTATION, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL OPEN MARKET BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY DAMAGES ARISING FROM OR RELATING TO THIS SOFTWARE OR THE DOCUMENTATION, INCLUDING, WITHOUT LIMITATION, ANY INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES OR SIMILAR DAMAGES, INCLUDING LOST PROFITS OR LOST DATA, EVEN IF OPEN MARKET HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS". OPEN MARKET HAS NO LIABILITY IN CONTRACT, TORT, NEGLIGENCE OR OTHERWISE ARISING OUT OF THIS SOFTWARE OR THE DOCUMENTATION.

This product uses materials from PJ Naughter’s Base64 Implementation.

Copyright (c) 1999 - 2004 by PJ Naughter.

All rights reserved.

Used by permission.

This product uses materials from RFC 3174.

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This product uses materials from RFC 3550.

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This product uses materials copyright Bob Jenkins.

(c) Bob Jenkins, 1996. 74512.261@compuserve.com.

You may use this code any way you wish, private, educational, or commercial, as long as this whole comment accompanies it.

See http://ourworld.compuserve.com/homepages/bob_jenkins/evahash.htm

Use to detect changes between revisions of documents, assuming nobody is trying to cause collisions. Do NOT use for cryptography.

This product uses materials copyright OpenVision Technologies, Inc. Copyright 1993 by OpenVision Technologies, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting

documentation, and that the name of OpenVision not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. OpenVision makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

OPENVISION DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL OPENVISION BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright (c) 2006 by BEA, BMC, CA, Cisco, Dell, EMC, HP, IBM, Intel, Microsoft, and Sun. All rights reserved.

Permission to copy, display, and distribute this Service Modeling Language (SML) Schema Document, in any medium without fee or royalty is hereby granted, provided that you include the following on ALL copies of the SML Schema

Document, or portions thereof, that you make:

1. A link or URL to the SML Schema Document at this location:

   http://schemas.serviceml.org/sml/2007/02/sml.xsd
2. The copyright notice as shown in the SML Schema Document.

BEA, BMC, CA, Cisco, Dell, EMC, HP, IBM, Intel, Microsoft, and Sun (collectively, the “Authors”) each agree to grant you a royalty-free license, under reasonable, non-discriminatory terms and conditions to their respective patents that they deem necessary to implement the Service Modeling Language Schema Document.

THE SML SCHEMA DOCUMENT IS PROVIDED "AS IS," AND THE AUTHORS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE SML SCHEMA DOCUMENT ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.THE AUTHORS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO ANY USE OR DISTRIBUTION OF THE SML SCHEMA DOCUMENT.

The name and trademarks of the Authors may NOT be used in any manner, including advertising or publicity pertaining to the SML Schema Document or its contents without specific, written prior permission. Title to copyright in the SML Schema Document will at all times remain with the Authors.

(c) 2001-2006 BEA Systems, Inc., BMC Software, CA, Inc.,

International Business Machines Corporation, Layer 7 Technologies, Microsoft Corporation, Inc., Novell, Inc. and VeriSign, Inc. All rights reserved.

Permission to copy, display, perform, modify and distribute WS-AUTHORIZATION.XSD (the "Document"), and to authorize others to do the foregoing, in any medium without fee or royalty is hereby granted for the purpose of developing and evaluating the Document.

BEA Systems, BMC Software, CA Inc., IBM, Layer 7 Technologies, Microsoft, Novell and VeriSign (collectively, the "Authors") each agree to grant a license to third parties, under royalty-free and otherwise reasonable, non-discriminatory terms and conditions, to their respective essential patent claims that they deem necessary to implement the Document.

THE DOCUMENT IS PROVIDED "AS IS," AND THE AUTHORS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE DOCUMENT IS SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY, PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

THE AUTHORS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE DOCUMENT OR THE PERFORMANCE OR IMPLEMENTATION OF THE CONTENTS THEREOF.

You may remove these disclaimers from your modified versions of the Document provided that you effectively disclaim all warranties and liabilities on behalf of all copyright holders in the copies of any such modified versions you distribute.

The name and trademarks of the Authors may NOT be used in any manner, including advertising or publicity pertaining to the Document or its contents without specific, written prior permission. Title to copyright in the Document will at all times remain with the Authors.

(c) 2001-2006 BEA Systems Inc., International Business Machines Corporation, Microsoft Corporation, Inc., SAP AG, Sonic Software, and VeriSign, Inc. All rights reserved.

Permission to copy and display the WS-Policy Specification (the "Specification", which includes WSDL and schema documents), in any medium without fee or royalty is hereby granted, provided that you include the following on ALL copies of the WS-Policy Specification, that you make:

1. A link or URL to the WS-Policy Specification at one of the Authors' websites
2. The copyright notice as shown in the WS-Policy Specification.

BEA Systems, IBM, Microsoft, SAP, Sonic Software, and VeriSign (collectively, the "Authors") each agree to grant you a license, under royalty-free and otherwise reasonable, non-discriminatory terms and conditions, to their respective essential patent claims that they deem necessary to implement the WS-Policy Specification.

THE WS-POLICY SPECIFICATION IS PROVIDED "AS IS," AND THE AUTHORS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE WS-POLICY SPECIFICATION ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

THE AUTHORS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO ANY USE OR DISTRIBUTION OF THE WS-POLICY SPECIFICATION.

The name and trademarks of the Authors may NOT be used in any manner, including advertising or publicity pertaining to the WS-Policy Specification or its contents without specific, written prior permission. Title to copyright in the WS-Policy Specification will at all times remain with the Authors.

No other rights are granted by implication, estoppel or otherwise.

This product uses materials authored by J. Zobel.Quickly computes a hash value from a sequence of bytes.This type of hash is NOT CRYPTOGRAPHIC; it is designedonly for use in hash tables. The author has made itpublic domain as long as it has this comment:

Author: J. Zobel, April 2001. Permission to use this codeis freely granted, provided that this statement is retained.

This product uses materials copyright Regents of the University of California.

Copyright (c) 1982, 1986, 1990, 1993

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

   This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)

ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses materials copyright OASIS.

Copyright OASIS Open 2002-2006. All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself does not be modified in any way, such as by removing the copyright notice or references to OASIS, except as needed for the purpose of developing OASIS specifications, in which case the procedures for copyrights defined in the OASIS Intellectual Property Rights document must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an AS IS basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This product uses materials copyright David Gottner.

Copyright 1994, David Gottner

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice, this permission notice and the following disclaimer notice appear unmodified in all copies.

I DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL I BE LIABLE FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA, OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

This product uses materials copyright William E. Kempf, Inc.

(c) Copyright William E. Kempf 2001

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all

copies and that both that copyright notice and this permission notice appear in supporting documentation. William E. Kempf makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

This file uses materials from LibTIFF.

Copyright (c) 1988-1997 Sam Leffler

Copyright (c) 1991-1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the names of Sam Leffler and Silicon Graphics may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

This product uses materials from the TWAIN toolkit

The TWAIN Toolkit is distributed as is. The developer and distributors of the TWAIN Toolkit expressly disclaim all implied, express or statutory warranties including, without limitation, the implied warranties of merchantability, noninfringement of third party rights and fitness for a particular purpose. Neither the developers nor the distributors will be liable for damages, whether direct, indirect, special, incidental, or consequential, as a result of the reproduction, modification, distribution or other use of the TWAIN Toolkit.

This product uses materials copyright Intel Corporation.

Copyright (c) 2004 - 2007, Intel Corporation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses jQuery UI.

Copyright 2010, AUTHORS.txt (http://jqueryui.com/about)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product uses xHCI Specification Information.

Copyright © 2008-2010 Intel Corporation. All rights reserved.

http://www.intel.com/technology/usb/download/xHCI_Adopters_Agreement.pdf

3.2 Copyright License. Subject to Adopter’s compliance with the terms of this Agreement, Intel, on behalf of itself and the Contributors, also hereby grants to Adopter a non-exclusive, royalty-free, non- transferable, non-sublicenseable, worldwide copyright license to the Final Specification to reproduce the Final Specification as necessary in order to exercise the patent rights granted in Section 3.1(a), provided that all reproductions thereof shall include any copyright notices and disclaimers contained in the Final Specification.

INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE,

MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

This product uses materials from RFC 4034.

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

This this software is based in part on the work of the Independent JPEG Group.

Copyright (C) 1991-2013, Thomas G. Lane, Guido Vollbeding.

All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

1. If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
2. If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".

3. Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

This product uses materials copyright Stephen Satchell.

Copyright (C) 1986 Stephen Satchell.

Programmers may incorporate any or all code into their programs, giving proper credit within the source. Publication of the source routines is permitted so long as proper credit is given to Stephen Satchell, Satchell Evaluations and Chuck Forsberg, Omen Technology.

This product uses materials copyright Jon Zeeff.

Copyright 1988 Jon Zeeff (zeeff@b-tech.ann-arbor.mi.us)

You can use this code in any manner, as long as you leave my name on it and don't hold me responsible for any problems with it.

This product uses materials copyright Lexmark International

Copyright (c) 1993-2003 Lexmark International Inc. All Rights Reserved.

Permission is granted for redistribution of this file as long as this copyright notice is intact and the content of the file is not altered in any way from its original form.

This product uses zlib.

Copyright (C) 1995-2002 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler

jloup@gzip.org madler@alumni.caltech.edu

This product uses SafeInt.

SafeInt.hpp

Version 3.0.14p

This software is licensed under the Microsoft Public License (Ms-PL).

For more information about Microsoft open source licenses, refer to

http://www.microsoft.com/opensource/licenses.mspx

This license governs use of the accompanying software. If you use the software, you accept this license. If you do not accept the license, do not use the software.

Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under U.S. copyright law. A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees, or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

This product uses materials from “Graphic Gems,” Academic Press, 1990.

This product uses materials copyright Gary P. Mussar

Copyright (c) 1990 Gary P. Mussar.

/***

* CRC utility routines for general 16 and 32 bit CRCs.

* 1990 Gary P. Mussar

* This code is released to the public domain. There are no restrictions, however, acknowledging the author by keeping this comment around* would be appreciated.

***/

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)

Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (http://www.openssl.org/)"

The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

==========

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT

LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License.]

Portions of this software are based in part on FreeBSD. Because Microsoft has included the FreeBSD software in this product, Microsoft is required to include the following text that accompanied such software:

All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by The Regents of the University of California.

Copyright 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The Institute of Electrical and Electronics Engineers and the American National Standards Committee X3, on Information Processing Systems have given us permission to reprint portions of their documentation.

In the following statement, the phrase ``this text'' refers to portions of the system documentation. Portions of this text are reprinted and reproduced in electronic form in the second BSD Networking Software Release, from IEEE Std 1003.1-1988, IEEE Standard Portable Operating System Interface for Computer Environments (POSIX), copyright C 1988 by the Institute of Electrical and Electronics Engineers, Inc. In the event of any discrepancy between these versions and the original IEEE Standard, the original IEEE Standard is the referee document.

In the following statement, the phrase ``This material'' refers to portions of the system documentation.

This material is reproduced with permission from American National Standards Committee X3, on Information Processing Systems. Computer and Business Equipment Manufacturers Association (CBEMA), 311 First St., NW, Suite 500, Washington, DC 20001-2178. The developmental work of Programming Language C was completed by the X3J11 Technical Committee.

The views and conclusions contained in the software and documentation are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the Regents of the University of California.

NOTE: The copyright of UC Berkeley's Berkeley Software Distribution ("BSD") source has been updated. The copyright addendum may be found at ftp://ftp.cs.berkeley.edu/pub/4bsd/README.Impt.License.Change and is includedbelow.

July 22, 1999

To All Licensees, Distributors of Any Version of BSD:

As you know, certain of the Berkeley Software Distribution ("BSD") source code files require that further distributions of products containing all or portions of the software, acknowledge within their advertising materials that such products contain software developed by UC Berkeley and its contributors.

Specifically, the provision reads: “All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors."

Effective immediately, licensees and distributors are no longer required to include the acknowledgement within advertising materials. Accordingly, the foregoing paragraph of those BSD Unix files containing it is hereby deleted in its entirety.

William Hoskins

Director, Office of Technology Licensing----------

Portions of this software are based in part on LZ4.

Copyright (C) 2011-2013, Yann Collet.

BSD 2-Clause License (http://www.opensource.org/licenses/bsd-license.php)

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses materials Copyright Carnegie Mellon University.

Copyright (c) 1989-2005 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

This work was supported in part by funding from the Defense Advanced Research Projects Agency and the National Science Foundation of the United States of America, and the CMU Sphinx Speech Consortium.

THIS SOFTWARE IS PROVIDED BY CARNEGIE MELLON UNIVERSITY ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY NOR ITS EMPLOYEES BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses materials copyright Silicon Graphics, Inc.

Copyright (C) Silicon Graphics, Inc. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice including the dates of first publication and either this permission notice or a reference to http://oss.sgi.com/projects/FreeB/ shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIESOF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL SILICON GRAPHICS, INC. BELIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN ANACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF ORIN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Silicon Graphics, Inc. shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from Silicon Graphics, Inc.

This product uses materials copyright The NetBSD Foundation, Inc.

Copyright (c) 2008 The NetBSD Foundation, Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses materials copyright John Dyer.

Portions derived from MediaElement.js

Copyright 2010-2013, John Dyer (http://j.hn)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product uses materials copyright the World Wide Web Consortium.

Copyright (c) 1994-2000 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University).

All Rights Reserved.

This program is distributed under the W3C's Software Intellectual Property License. This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See W3C License http://www.w3.org/Consortium/Legal/ for more details.

Copyright © 1995 CERN. "This product includes computer software created and made available by CERN. This acknowledgment shall be mentioned in full in any product which includes the CERN computer software included herein or parts thereof."

License

By obtaining, using and/or copying this work, you (the licensee) agree that you have read, understood, and will comply with the following terms and conditions.

Permission to copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, the W3C Software Short Notice should be included (hypertext is preferred, text is permitted) within the body of any redistributed or derivative code.

Notice of any changes or modifications to the files, including the date changes were made. (We recommend you provide URIs to the location from which the code is derived.)

Disclaimers

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

This product uses materials copyright W3C.

Copyright 1998-2004 W3C (MIT, ERCIM, Keio),

All Rights Reserved. Permission to use, copy, modify and distribute the SSML core schema and its accompanying documentation for any purpose and without fee is hereby granted in perpetuity, provided that the above copyright notice and this paragraph appear in all copies. The copyright holders make no representation about the suitability of the schema for any purpose. It is provided "as is" without expressed or implied warranty.

This product uses materials copyright Ralph Hancock, and John Hudson.

Hebrew OpenType Layout logic copyright (c) 2003 & 2007, Ralph Hancock & John Hudson. This layout logic for Biblical Hebrew is open source software under the MIT License.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE,

ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product uses materials copyright the CodePlex Foundation.

Copyright (c) 2009 CodePlex Foundation.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the CodePlex Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL <COPYRIGHT HOLDER> BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product uses materials copyright the University of Illinois.

University of Illinois/NCSA

Open Source License

Copyright (c) 2003-2010 University of Illinois at Urbana-Champaign.

All rights reserved.

Developed by:

LLVM Team

University of Illinois at Urbana-Champaign

http://llvm.org

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal with the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense,

and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimers.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimers in the documentation and/or other materials provided with the distribution.

* Neither the names of the LLVM Team, University of Illinois at Urbana-Champaign, nor the names of its contributors may be used to endorse or promote products derived from this Software without specific prior written permission.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE CONTRIBUTORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS WITH THESOFTWARE.

All other trademarks are property of their respective owners.

Migrate Roles and Features to Windows Server

Migration documentation and tools ease the process of migrating server roles, features, operating system settings, and data from an existing server that is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 to a computer that is running Windows Server 2012 R2. By using migration guides linked to on this page (and where appropriate, Windows Server Migration Tools) to migrate roles, role services, and features, you can simplify deployment of new servers (including those that are running the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012, and virtual servers), reduce migration downtime, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

Most of the migration documentation and tools featured in this section support cross-architecture migrations (x86-based to x64-based computing platforms), migrations between physical and virtual environments, and migrations between both the full and Server Core installation options of the Windows Server operating system, where available.

In Windows Server 2012 and later releases of Windows Server, Windows Server Migration Tools supports cross-subnet migrations.

Migration guides

The following are available resources for migrating roles to Windows Server 2012 or Windows Server 2012 R2.

Windows Server roles, role services, and features

Windows Server Migration guides provide you with instructions for migrating a single role, role service, or feature to a server that is running Windows Server 2012 or Windows Server 2012 R2. Guides do not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure specific to your server environment, based on the information provided in other migration guides.

- Migrate Roles and Features to Windows Server 2012 R2
- Migrate Roles and Features to Windows Server 2012

Windows Server Migration Tools

Windows Server Migration Tools, available as a feature in Windows Server 2012 R2 and Windows Server 2012, allows an administrator to migrate some server roles, features, operating system settings, shares, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 to computers that are running Windows Server 2012 or Windows Server 2012 R2.

Not all migrations require or use Windows Server Migration Tools. Guides for migrations that require Windows Server Migration Tools clearly state that Windows Server Migration Tools setup is part of the migration process, and provide specific instructions for how to use Windows Server Migration Tools.

To use Windows Server Migration Tools, the feature must be installed on both source and destination computers as described in the following guide.

Install, use, and remove Windows Server Migration Tools

See Also

Migrating Roles and Features to Windows Server

Migrate Roles and Features to Windows Server 2012 R2

Migration documentation and tools ease the process of migrating server roles, features, operating system settings, and data from an existing server that is running Windows Server 2003, Windows

Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 to a computer that is running Windows Server 2012 R2. By using migration guides linked to on this page (and where appropriate, Windows Server Migration Tools) to migrate roles, role services, and features, you can simplify deployment of new servers (including those that are running the Server Core installation option of Windows Server 2012 or Windows Server 2012 R2, and virtual servers), reduce migration downtime, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

In this section

- Active Directory Certificate Services Migration Guide for Windows Server 2012 R2
- Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2
- Migrate DHCP Server to Windows Server 2012 R2
- Migrate Hyper-V to Windows Server 2012 R2 from Windows Server 2012
- Migrate File and Storage Services to Windows Server 2012 R2
- Migrate Remote Desktop Services to Windows Server 2012 R2
- Migrate Cluster Roles to Windows Server 2012 R2
- Migrate Network Policy Server to Windows Server 2012 R2

See also

- **Migrate from Previous Versions to Windows Server 2012 R2 Essentials**
- **Transition from Windows Server 2012 R2 Essentials to Windows Server 2012 R2 Standard**
- Migrating Roles and Features to Windows Server
- Migrate Roles and Features to Windows Server

Active Directory Certificate Services Migration Guide for Windows Server 2012 R2

About this guide

This document provides guidance for migrating a certification authority (CA) to a server that is running Windows Server 2012 R2 from a server that is running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 R2, or Windows Server 2003.

Target audience

- Administrators or IT operations engineers responsible for planning and performing CA migration.

- Administrators or IT operations engineers responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications.
- IT operations managers accountable for network and server management.
- IT architects responsible for computer management and security throughout an organization.

Supported migration scenarios

This guide provides you with instructions for migrating an existing server that is running Active Directory® Certificate Services (AD CS) to a server that is running Windows Server 2008 R2 or Windows Server 2012 R2. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, you should design a custom migration procedure that is specific to your server environment, based on the information provided in other role migration guides. To view migration guides for other server roles, see Migrating Roles and Features in Windows Server (http://go.microsoft.com/fwlink/?LinkID=128554).

Note

This guide can be used to migrate a CA from a source server that is also a domain controller to a destination server with a different name. However, migration of a domain controller is not covered by this guide. For information about Active Directory Domain Services (AD DS) migration, see Active Directory Domain Services and DNS Server Migration Guide (http://go.microsoft.com/fwlink/?LinkId=179357).

Supported operating systems

This guide supports migrations from source servers running the operating system versions and service packs listed in the following table. All migrations described in this document assume that the destination server is running Windows Server 2012 R2 as specified in the following table.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Server 2012 R2	Windows Server 2012 R2, Server with a GUI only (not Server Core or Minimal Server Interface)	x64-based
x64-based	Windows Server 2012	Windows Server 2012 R2 or Windows Server 2012, Server with a GUI only (not Server Core or Minimal Server Interface)	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Server 2008 R2	Windows Server 2012 R2or Windows Server 2012, Server with a GUI only (not Server Core or Minimal Server Interface) or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86-based or x64-based	Windows Server 2008	Windows Server 2012 R2or Windows Server 2012, Server with a GUI only (not Server Core or Minimal Server Interface) or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86-based or x64-based	Windows Server 2003 R2	Windows Server 2012 R2or Windows Server 2012, Server with a GUI only (not Server Core or Minimal Server Interface) or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86-based or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012 R2or Windows Server 2012, Server with a GUI only (not Server Core or Minimal Server Interface) or Windows Server 2008 R2, both full and Server Core installation options	x64-based

Note

In-place upgrades directly from Windows Server 2003 with Service Pack 2 or Windows Server 2003 R2 to Windows Server 2012 R2 are not supported. If you are running an x64-based computer, you can upgrade the CA role service from Windows Server 2003 with Service Pack 2 or Windows Server 2003 R2 to Windows Server 2008 or Windows Server 2008 R2 first and then upgrade to Windows Server 2012 R2 or Windows Server 2012.

What this guide does not provide

- Procedures to upgrade to Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2
- Procedures to migrate additional server roles
- Procedures to migrate additional AD CS role services

In general, migration is not required for the following AD CS role services. Instead, you can install and configure these role services on computers running Windows Server 2008 R2 or Windows Server 2012 by completing the role service installation procedures. For information about the impact of CA migration on other AD CS role services, see Impact of migration on other computers in the enterprise.

- CA Web Enrollment (http://go.microsoft.com/fwlink/?LinkId=179360)
- Online Responder (http://go.microsoft.com/fwlink/?LinkId=143098)
- Network Device Enrollment (http://go.microsoft.com/fwlink/?LinkId=179362)
- Certificate Enrollment Web Services (http://go.microsoft.com/fwlink/?LinkId=179363)

CA migration overview

Certification authority (CA) migration involves several procedures, which are covered in the following sections.

Warning

During the migration procedure, you are asked to turn off your existing CA (either the computer or at least the CA service). You are asked to name the destination CA with the same name that you used for the original CA. The computer name, (hostname or NetBIOS name), does not have to match that of the original CA. However, the destination CA name must match that of the source CA. Further, the destination CA name must not be identical to the destination computer name.

Note

It is possible to install a new PKI hierarchy while still leveraging an existing PKI hierarchy. However, doing so requires designing a new PKI, which is not covered in this guide. For an informal overview of how a dual PKI could work for an organization, see the following Ask DS blog post: Moving Your Organization from a Single Microsoft CA to a Microsoft Recommended PKI.

Preparing to migrate

- Preparing your destination server
- Backing up your source server
- Preparing your source server

Migrating the certification authority

- Backing up a CA database and private key
- Backing up CA registry settings
- Backing up CAPolicy.inf
- Removing the CA role service from the source server
- Removing the source server from the domain
- Joining the destination server to the domain
- Adding the CA role service to the destination server
- Restoring the CA database and configuration on the destination server
- Granting permissions on AIA and CDP containers
- Additional procedures for failover clustering (optional)

Verifying the migration

- Verifying certificate enrollment
- Verifying CRL publishing

Post-migration tasks

- Upgrading certificate templates in Active Directory Domain Services (AD DS)
- Retrieving certificates after a host name change
- Restoring Active Directory Certificate Services (AD CS) to the source server in the event of migration failure
- Troubleshooting migration

Impact of migration

Impact of migration on the source server

The CA migration procedures described in this guide include decommissioning the source server after migration is completed and CA functionality on the destination server has been verified. If the source server is not decommissioned, then the source server and destination server must have different names. Additional steps are required to update the CA configuration on the destination server if the name of the destination server is different from the name of the source server.

Impact of migration on other computers in the enterprise

During migration, the CA cannot issue certificates or publish CRLs.

To ensure that revocation status checking can be performed by domain members during CA migration, it is important to publish a CRL that is valid beyond the planned duration of the migration.

Because the authority identification access and CRL distribution point extensions of previously issued certificates may reference the name of the source CA, it is important to either continue to publish CA certificates and CRLs to the same location or provide a redirection solution. For an example of configuring IIS redirection, see Redirecting Web Sites in IIS 6.0.

Permissions required to complete the migration

To install an enterprise CA or a standalone CA on a domain member computer, you must be a member of the Enterprise Admins group or Domain Admins group in the domain. To install a standalone CA on a server that is not a domain member, you must be a member of the local Administrators group. Removal of the CA role service from the source server has the same group membership requirements as installation.

Estimated duration

The simplest CA migration can typically be completed within one to two hours. The actual duration of CA migration depends on the number of CAs and the sizes of CA databases.

See also

- Prepare to Migrate
- Migrating the Certification Authority
- Verifying the Certification Authority Migration
- Post-Migration Tasks
- Migrating Roles and Features in Windows Server

Prepare to Migrate

To reduce the duration of the migration process, you can complete the procedures detailed in this topic before beginning the migration process and taking the certification authority (CA) offline.

- Preparing your destination server
- Backing up your source server
- Preparing your source server

Preparing your destination server

Hardware requirements for the destination server

The hardware requirements to install any of the Active Directory Certificate Services (AD CS) role services are the same as the minimum and recommended configurations for installation of Windows Server 2012 R2. This section includes the general hardware recommendations for Windows Server 2012 R2. For detailed requirements, see System Requirements and Installation Information for Windows Server 2012 R2.

Hardware requirements for AD CS

In addition to the hardware requirements for the operating system, consider these storage and performance requirements for optimal CA performance and availability:

- The disk space requirements for a CA database depend on the number of certificates that the CA issues. Because a CA stores certificate requests, the issued certificates, and optionally, archived key material, 64 KB of database space per certificate is recommended.
- The operating system, the CA database, and the CA log files should be stored on separate physical disk drives in a multidisk configuration. For optimal CA performance and reliability, consider a redundant array of independent disks (RAID) system, such as RAID 5 for the CA database and log files and RAID 1 or RAID 0+1 for the operating system. A recommended minimum hard disk speed is 10,000 RPM.
- Processor power is generally more important to CA performance than system memory capacity.
- Failover clusters have additional hardware, software, and networking requirements. For more information, see Failover Cluster Requirements (http://go.microsoft.com/fwlink/?LinkId=179369).
- If a hardware security module (HSM) is used by the CA, consult with your HSM vendor to verify compatibility with Windows Server 2012 R2.

Software requirements for the destination server

Enterprise CAs can be installed on computers running any version of Windows Server 2012 R2.

When AD CS in Windows Server 2012 R2 is installed in an Active Directory Domain Services (AD DS) domain, the AD DS schema version must be at least 30 and all domain controllers in the domain must be running one of the following operating systems:

- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 with Service Pack 1 (SP1)
- Windows Server 2008
- Windows Server 2003 R2
- Windows Server 2003 with Service Pack 2 (SP2)
- Windows Server 2003 with SP1

- Windows Server 2003

Note
Domain controllers running Windows 2000 Server with Service Pack 4 (SP4) or Windows 2000 Server with Service Pack 3 (SP3) are technically compatible with AD CS deployments. However, the use of Windows 2000 Server is not recommended because Mainstream Support is no longer available for this operating system. For more information, see Microsoft Support Lifecycle (http://go.microsoft.com/fwlink/?LinkId=117347).

If an HSM is used by the CA, consult your HSM vendor to verify cryptographic service provider (CSP) and key service provider (KSP) compatibility with Windows Server 2012 R2 depending on the operating system to be used.

Installing the Operating System

To reduce the duration of the migration process, you can prepare the destination server by completing the following procedures before beginning the migration process and taking the source CA offline.

- Review the hardware and software requirements in the previous sections.
- Install Windows Server 2012 R2. For more information, see System Requirements and Installation Information for Windows Server 2012 R2.
- Install updates by using Windows Update.
- (Optional) Install failover clustering by reviewing the Active Directory Certificate Services (AD CS) Clustering documentation.

If you are migrating to a Server Core installation you should configure the server for remote management, which is disabled by default.

Configure remote management on Server Core

1. Log on as an administrator.
2. Type **sconfig.cmd** and press ENTER.
3. Perform the following tasks by completing the procedures described in Configuring a Server Core installation with Sconfig.cmd:
   a. Configure network settings as required for your environment.
   b. Join the server to your domain. This step is required if you are setting up an enterprise CA and optional if you are setting up a standalone CA.
   c. Configure Remote Management to enable **MMC Remote Management** or **Server Manager Remote Management**.
   d. Enable **Remote Desktop** (optional).
4. Type **13** and press ENTER to close sconfig.cmd.

Note

For information on configuring remote management in see Configure Remote Management in Server Manager.

Backing up your source server

Back up your source server to prepare for recovery of the source CA in the event of migration failure.

For information about backing up Windows Server 2012 R2 or Windows Server 2012, see Windows Server Backup.

For more information about creating backups in Windows Server 2008, see the Windows Server Backup Step-by-Step Guide for Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=119141).

For more information about creating system state backups in Windows Server 2003, see article 326216 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=117369).

Detailed procedures for backing up the source CA database, private key, and registry settings are provided in the topic Migrating the Certification Authority.

Preparing your source server

To reduce the duration and impact of CA migration, the following procedures should be completed before you begin migration:

- Back up the CA templates list (required only for enterprise CAs).
- Record the CA's CSP and signature algorithm.
- Publish a CRL with an extended validity period.

Backing up a CA templates list

An enterprise CA can have certificate templates assigned to it. You should record the assigned certificate templates before beginning the CA migration. The information is not backed up with the CA database or registry settings backup. This is because certificate templates and their association with enterprise CAs are stored in AD DS. You will need to add the same list of templates to the destination server to complete CA migration.

Note

It is important that the certificate templates assigned to the source CA are not changed after this procedure is completed.

You can determine the certificate templates assigned to a CA by using the Certification Authority snap-in or the **Certutil.exe –catemplates** command.

To record a CA templates list by using the Certification Authority snap-in

1. Log on with local administrative credentials to the CA computer.
2. Open the Certification Authority snap-in.

3. In the console tree, expand **Certification Authority**, and click **Certificate Templates**.
4. Record the list of certificate templates by taking a screen shot or by typing the list into a text file.

To record a CA templates list by using Certutil.exe

1. Log on with local administrative credentials to the CA computer.
2. Open a Command Prompt window.
3. Type the following command and press ENTER.

```
certutil.exe -catemplates > catemplates.txt
```

4. Verify that the catemplates.txt file contains the templates list.

Note
If no certificate templates are assigned to the CA, the file contains an error message: 0x80070490 (Element not found).

Recording a CA's signature algorithm and CSP

During CA installation on the destination server, you can specify the signature algorithm and CSP used by the CA, or accept the default configuration. If your source CA is not using the default configuration, then you should complete the following procedure to record the CSP and signature algorithm.

Note
If an HSM is used by the source CA, follow procedures provided by the HSM vendor to determine the HSM CSP.

To record a CA's CSP by using Certutil.exe

1. Log on with local administrative credentials to the CA computer.
2. Open a Command Prompt window.
3. Type the following command and press ENTER.

```
certutil.exe -getreg ca\csp\* > csp.txt
```

4. Verify that the csp.txt file contains the CSP details.

Publishing a CRL with an extended validity period

Before beginning CA migration, it is a good practice to publish a CRL with a validity period that extends beyond the planned migration period. The validity period of the CRL should be at least the length of time that is planned for the migration. This is necessary to enable certificate validation processes on client computers to continue during the migration period.

You should publish a CRL with an extended validity period for each CA being migrated. This procedure is particularly important in the case of a root CA because of the potentially large number of certificates that would be affected by the unavailability of a CRL.

By default, the CRL validity period is equal to the CRL publishing period plus 10 percent. After determining an appropriate CRL validity period, set the CRL publishing interval and manually publish the CRL by completing the following procedures:

Important
Record the value of the CRL publishing period before changing it. After migration is complete, the CRL publishing period should be reset to its previous value.

- Schedule the publication of the certificate revocation list
- Manually publish the certificate revocation list

Caution
Client computers download a new CRL only after the validity period of a locally cached CRL expires. Therefore, you should not use a CRL validity period that is excessively long.

Next steps

After completing the procedures to prepare the source and destination servers, you should review the topic Migrating the Certification Authority and complete the procedures appropriate for your specific migration scenario.

See also

- Active Directory Certificate Services Migration Guide for Windows Server 2012 R2
- Migrating the Certification Authority
- Verifying the Certification Authority Migration
- Post-Migration Tasks
- Migrating Roles and Features in Windows Server

Migrating the Certification Authority

Review all procedures in this topic and complete only the procedures that are required for your migration scenario.

- Backing up a CA database and private key
- Backing up CA registry settings
- Backing up CAPolicy.inf
- Removing the CA role service from the source server
- Removing the source server from the domain
- Joining the destination server to the domain
- Adding the CA role service to the destination server
- Restoring the CA database and configuration on the destination server

- Granting permissions on AIA and CDP containers
- Additional procedures for failover clustering

  This is an optional set of steps if you are migrating to a failover cluster.

Backing up a CA database and private key

You can back up the CA database and private key by using the Certification Authority snap-in or by using Certutil.exe at a command prompt. Complete either one of the backup procedures described in this section.

Note

If a hardware security module (HSM) is used by the CA, back up the private keys by following procedures provided by the HSM vendor.

After completing backup steps, the Active Directory Certificate Services service (Certsvc) should be stopped to prevent issuance of additional certificates. Before adding the CA role service to the destination server, the CA role service should be removed from the source server.

The backup files created during these procedures should be stored in the same location to simplify the migration. The location should be accessible from the destination server; for example, removable media or a shared folder on the destination server or another domain member.

Backing up a CA database and private key by using the Certification Authority snap-in

The following procedure describes the steps to back up the CA database and private key by using the Certification Authority snap-in while logged on to the source CA.

Note

If you prefer, you can use the certutil application to back up the CA database and private key. Using certutil for CA backup is covered in the next section.

You must use an account that is a CA administrator. On an enterprise CA, the default configuration for CA administrators includes the local Administrators group, the Enterprise Admins group, and the Domain Admins group. On a standalone CA, the default configuration for CA administrators includes the local Administrators group.

To back up a CA database and private key by using the Certification Authority snap-in

1. Choose a backup location and attach media, if necessary.
2. Log on to the source CA.
3. Open the **Certification Authority** snap-in.
4. Right-click the node with the CA name, point to **All Tasks**, and then click **Back Up CA**.
5. On the **Welcome** page of the CA Backup wizard, click **Next**.
6. On the **Items to Back Up** page, select the **Private key and CA certificate** and **Certificate database and certificate database log** check boxes, specify the backup

location, and then click **Next**.

7. On the **Select a Password** page, type a password to protect the CA private key, and click **Next**.

   Security

   Use a strong password; for example, at least eight characters long with a combination of uppercase and lowercase characters, numbers, and punctuation characters.

8. On the **Completing the Backup Wizard** page, click **Finish**.
9. After the backup completes, verify the following files in the location you specified:
   - *CAName*.p12 containing the CA certificate and private key
   - Database folder containing files certbkxp.dat, edb#####.log, and *CAName*.edb
10. Open a Command Prompt window, and type **net stop certsvc** to stop the Active Directory Certificate Services service.

    Important

    The service should be stopped to prevent issuance of additional certificates. If certificates are issued by the source CA after a database backup is completed, repeat the CA database backup procedure to ensure the database backup contains all issued certificates.

11. Copy all backup files to a location that is accessible from the destination server; for example, a network share or removable media.

    Security

    The private key must be protected against compromise. Protect a shared folder by limiting its access control list to authorized CA administrators. Protect removable media against unauthorized access and damage.

Backing up a CA database and private key by using Windows PowerShell

The following procedure describes the steps to back up the CA database and private key by using the Backup-CARoleService cmdlet while logged on to the source CA.

Important

You must use an account that is a CA administrator. On an enterprise CA, the default configuration for CA administrators includes the local Administrators group, the Enterprise Admins group, and the Domain Admins group. On a standalone CA, the default configuration for CA administrators includes the local Administrators group.

To back up a CA database and private key by using Windows PowerShell

1. Log on with local administrative credentials to the CA computer.
2. Right-click **Windows PowerShell** and click **Run as Administrator**.

3. Type the following command and press ENTER:

```
Backup-CARoleService –path <BackupDirectory>
```

Note

BackupDirectory specifies the directory in which the backup files are created. The specified value can be a relative or absolute path. If the specified directory does not exist, it is created. The backup files are created in a subdirectory named Database.

4. The service must be stopped to prevent issuance of additional certificates. Type the following command and press ENTER:

```
Stop-service certsvc
```

5. After the backup completes, verify the following files in the location you specified:
   - *CAName*.p12 containing the CA certificate and private key
   - Database folder containing files certbkxp.dat, edb#####.log, and *CAName*.edb
6. Copy all backup files to a location that is accessible from the destination server; for example, a network share or removable media.

Security

The private key must be protected against compromise. Protect a shared folder by granting permission to only authorized CA administrators. Protect removable media against unauthorized access and damage.

Backing up a CA database and private key by using Certutil.exe

The following procedure describes the steps to back up the CA database and private key by using Certutil.exe while logged on to the source CA.

Important

You must use an account that is a CA administrator. On an enterprise CA, the default configuration for CA administrators includes the local Administrators group, the Enterprise Admins group, and the Domain Admins group. On a standalone CA, the default configuration for CA administrators includes the local Administrators group.

To back up a CA database and private key by using Certutil.exe

1. Log on with local administrative credentials to the CA computer.
2. Open a Command Prompt window.
3. Type **Certutil.exe –backupdb** <*BackupDirectory*> and press ENTER.
4. Type **Certutil.exe –backupkey** <*BackupDirectory*> and press ENTER.

Note

BackupDirectory specifies the directory in which the backup files are created. The specified value can be a relative or absolute path. If the specified directory does not exist, it is created. The backup files are created in a subdirectory named

Database.

5. Type a password at the prompt, and press ENTER. You must retain a copy of the password to access the key during CA installation on the destination server.

   Security

   Use a strong password; for example, at least eight characters with a combination of uppercase and lowercase characters, numbers, and symbols.

6. Type **net stop certsvc** and press ENTER to stop the Active Directory Certificate Services service. The service must be stopped to prevent issuance of additional certificates.
7. After the backup completes, verify the following files in the location you specified:
   - *CAName*.p12 containing the CA certificate and private key
   - Database folder containing files certbkxp.dat, edb#####.log, and *CAName*.edb
8. Copy all backup files to a location that is accessible from the destination server; for example, a network share or removable media.

   Security

   The private key must be protected against compromise. Protect a shared folder by granting permission to only authorized CA administrators. Protect removable media against unauthorized access and damage.

Backing up CA registry settings

Complete one of the following procedures to back up the CA registry settings.

The files created during the backup procedure should be stored in the same location as the database and private key backup files to simplify the migration. The location should be accessible from the destination server; for example, removable media or a shared folder on the destination server or another domain member.

You must be logged on to the source CA using an account that is a member of the local Administrators group.

To back up CA registry settings by using Regedit.exe

1. Click **Start**, point to **Run**, and type **regedit** to open the Registry Editor.
2. In `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc`, right-click **Configuration**, and then click **Export**.
3. Specify a location and file name, and then click **Save**. This creates a registry file containing CA configuration data from the source CA.
4. Copy the registry file to a location that is accessible from the destination server; for example, a shared folder or removable media.

To back up CA registry settings by using Reg.exe

1. Open a Command Prompt window.
2. Type **reg export HKLM\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration** <*output file*>.reg and press ENTER.
3. Copy the registry file to a location that is accessible from the destination server; for example, a shared folder or removable media.

Backing up CAPolicy.inf

If your source CA is using a custom CAPolicy.inf file, you should copy the file to the same location as the source CA backup files.

The CAPolicy.inf file is located in the %SystemRoot% directory, which is usually C:\Windows.

Removing the CA role service from the source server

It is important to remove the CA role service from the source server after completing backup procedures and before installing the CA role service on the destination server. Enterprise CAs and standalone CAs that are domain members store in Active Directory Domain Services (AD DS) configuration data that is associated with the common name of the CA. Removing the CA role service also removes the CA's configuration data from AD DS. Because the source CA and destination CA share the same common name, removing the CA role service from the source server after installing the CA role service on the destination server removes configuration data that is required by destination CA and interferes with its operation.

The CA database, private key, and certificate are not removed from the source server by removing the CA role service. Therefore, reinstalling the CA role service on the source server restores the source CA if migration fails and performing a rollback is required. See Restoring Active Directory Certificate Services (AD CS) to the source server in the event of migration failure.

Warning

Although it is not recommended, some administrators may choose to leave the CA role service installed on the source server to enable the source CA to be brought online quickly in the case of migration failure. If you choose not to remove the CA role service from the source server before installing the CA role service on the destination server, it is important that you disable the Active Directory Certificate Services service (Certsvc) and shut down the source server before installing the CA role service on the destination server. Do not remove the CA role service from the source server after completing the migration to the destination server. Removing the CA role service from the source server after migrating to the destination server interferes with the operation of the destination CA.

- To remove the CA on a computer running Windows Server 2003, use the Add/Remove Windows Components wizard.
- To remove the CA on a computer running Windows Server 2008 or later, use the Remove Roles and Features Wizard in Server Manager.

Removing the source server from the domain

Because computer names must be unique within an Active Directory domain, it is necessary to remove the source server from its domain and delete the associated computer account from Active Directory before joining the destination server to the domain.

If you have access to a domain member computer running Windows Server 2008 or later , complete the following procedure to remove the source server from the domain by using Netdom.exe.

If you do not have access to a computer running Windows Server 2008 or later, then complete the procedure Join a Workgroup (http://go.microsoft.com/fwlink/?LinkId=207683). Joining a workgroup also removes a domain member computer from its domain.

To remove the source server from the domain by using Netdom.exe

1. On a domain member computer running Windows Server 2008 or later, open an elevated Command Prompt window.
2. Type **netdom remove** <*source server name*> /d:<*domain name*> /ud:<*domain user account*> /pd:* and press ENTER. For additional command-line options, see Netdom remove syntax (http://go.microsoft.com/fwlink/?LinkID=207681).

   Tip
   Using Windows PowerShell, you can run the command: **remove-computer** <*computer name*>

   For more information, see Remove-Computer (http://technet.microsoft.com/en-us/library/dd347703.aspx).
3. Shut down the source server.

After removing the source server from its domain, delete the source server's computer account from AD DS by completing the procedure Delete a Computer Account (http://go.microsoft.com/fwlink/?LinkID=138386).

Tip
You can also use Windows PowerShell to remove the computer account from AD DS. For more information, see Remove-ADComputer (http://technet.microsoft.com/library/hh852313).

Joining the destination server to the domain

Before joining the destination server to the domain, change the computer name to the same name as the source server. Then complete the procedure to join the destination server to the domain.

If your destination server is running on the Server Core installation option, you must use the command-line procedure.

To rename the destination server, you must be a member of the local Administrators group. To join the server to the domain, you must be a member of the Domain Admins or Enterprise Admins groups, or have delegated permissions to join the destination server to an organizational unit (OU) in the domain.

Important
If you are migrating a standalone CA that is not a domain member, complete only the steps to rename the destination server and do not join the destination server to the domain.

To join the destination server to the domain by using Netdom.exe

1. On the destination server, open an elevated Command Prompt window.
2. Type **netdom renamecomputer** <*computer name*> /newname:<*new computer name*>

   Tip
   Using Windows PowerShell, you can run the command: **rename-computer** <*new computer name*>
3. Restart the destination server.
4. After the destination server restarts, log on by using an account that has permission to join computers to the domain.
5. Open an elevated Command Prompt window, type **netdom join** <*computer name*> /d:<*domain name*> /ud:<*domain user account*> /pd:* [/ou:<*OU name*>] and press ENTER. For additional command-line options, see Netdom join syntax (http://go.microsoft.com/fwlink/?LinkID=207680).

   Tip
   Using Windows PowerShell, you can run the command: **add-computer -DomainName** <*domain name*>

   For more information, see Add-Computer (http://technet.microsoft.com/en-us/library/dd347556.aspx).
6. Restart the destination server.

Adding the CA role service to the destination server

This section describes two different procedures for adding the CA role service to the destination server, including special instructions for using failover clustering.

Review the following statements to determine which procedures to complete.

- If your destination server is running the Server Core installation option, you can use Windows PowerShell to install the CA. See Install-AdcsCertificationAuthority for more information.
- If you are migrating to a CA that uses failover clustering, you must review the section "Special instructions for migrating to a failover cluster" and complete the procedures Importing the CA certificate and Adding the CA role service by using Server Manager.
- If you are migrating to a CA that uses an HSM, you must complete the procedures Importing the CA certificate and Adding the CA role service by using Server Manager.
- If none of the above statements describes your migration scenario, you can use the following procedure to add the CA role service: Adding the CA role service by using Server Manager. If you use Server Manager, you must also complete the procedure Importing the CA certificate.

Special instructions for migrating to a failover cluster

If you are migrating to a failover cluster, the procedures to import the CA certificate and add the CA role service must be completed on each cluster node. After the CA role service is added to each node, you should stop the Active Directory Certificate Services service (Certsvc).

Additionally, it is important to ensure that the shared storage used by the CA is online and assigned to the node you are adding the CA role service to.

The CA database and log files must be located on shared storage. Specify the shared storage location during step 12 of the CA installation procedure.

To verify shared storage is online

1. Log on to the destination server.
2. Start Server Manager.
3. In the console tree, double-click **Storage**, and click **Disk Management**.
4. Ensure that the shared storage is online and assigned to the node you are logged on to.

Importing the CA certificate

If you are adding the CA role service by using Server Manager, you must complete the following procedure to import the CA certificate.

To import the CA certificate

1. Start the Certificates snap-in for the local computer account.
2. In the console tree, double-click **Certificates (Local Computer)**, and click **Personal**.

3. On the **Action** menu, click **All Tasks**, and then click **Import** to open the Certificate Import Wizard. Click **Next**.
4. Locate the <*CAName*>.p12 file created by the CA certificate and private key backup on the source CA, and click **Open**.
5. Type the password, and click **OK**.
6. Click **Place all certificates in the following store**.
7. Verify **Personal** is displayed in **Certificate store**. If it is not, click **Browse**, click **Personal**, click **OK**.

   Note
   If you are using a network HSM, complete steps 8 through 10 to repair the association between the imported CA certificate and the private key that is stored in the HSM. Otherwise, click **Finish** to complete the wizard and click **OK** to confirm that the certificate was imported successfully.
8. In the console tree, double-click **Personal Certificates**, and click the imported CA certificate.
9. On the **Action** menu, click **Open**. Click the **Details** tab, copy the serial number to the Clipboard, and then click **OK**.
10. Open a Command Prompt window, type **certutil –repairstore My "{Serialnumber}"** and then press ENTER.

Adding the CA role service by using Server Manager

If your destination server is a domain member, you must use an account that is a member of the Domain Admins or Enterprise Admins group in order for the installation wizard to access objects in AD DS.

Important
If you made a backup CAPolicy.inf file from the source CA, review the settings and make adjustments, if necessary. Copy the CAPolicy.inf file to the %windir% folder (C:\Windows by default) of the destination CA before adding the CA role service.

To add the CA role service by using Server Manager

1. Log on to the destination server, and start Server Manager.
2. In the console tree, click **Roles**.
3. On the **Action** menu, click **Add Roles**.
4. If the **Before you Begin** page appears, click **Next**.
5. On the **Select Server Roles** page, select the **Active Directory Certificate Services** check box, and click **Next**.
6. On the **Introduction to AD CS** page, click **Next**.
7. On the **Role Services** page, click the **Certification Authority** check box, and click **Next**.

   Note

If you plan to install other role services on the destination server, you should complete the CA installation first, and then install other role services separately. Installation procedures for other AD CS role services are not described in this guide.

8. On the **Specify Setup Type** page, specify either **Enterprise** or **Standalone**, to match the source CA, and click **Next**.
9. On the **Specify CA Type** page, specify either **Root CA** or **Subordinate CA**, to match the source CA, and click **Next**.
10. On the **Set Up Private Key** page, select **Use existing private key** and **Select a certificate and use its associated private key**.

    Note
    If an HSM is used by the CA, select the private key by following procedures provided by the HSM vendor.

11. In the **Certificates** list, click the imported CA certificate, and then click **Next**.

    Note
    If you are using a custom CSP that requires strong private key protection, click **Allow administrator interaction when the private key is accessed by the CA**. The CSPs included with Windows Server do not require this setting to be enabled.

12. On the **CA Database** page, specify the locations for the CA database and log files.

    Note
    If you are migrating the CA to a failover cluster, the specified locations for database and log files must be on shared storage that is attached to all nodes. Because the location is common to cluster nodes, click **Yes** to overwrite the existing CA database as you add the CA role service to other nodes.

    Important
    If you specify locations that are different from the locations used on the source CA, then you must also edit the registry settings backup file before the CA is restored. If the locations specified during setup are different from the locations specified in the registry settings, the CA cannot start.

13. On the **Confirmation** page, review the messages, and then click **Configure**.
14. If you are migrating to a failover cluster, stop the Active Directory Certificate Services service (Certsvc) and HSM service if your CA uses an HSM. Then repeat the procedures to import the CA certificate and add the CA role service on other cluster nodes.

Adding the CA role service by using Windows PowerShell

Use the following procedure to add the CA role service by using the Install-ADCSCertificateAuthority cmdlet with the ExistingCertificateParameterSet:

```
Install-AdcsCertificationAuthority [-AllowAdministratorInteraction [<SwitchParameter>]]
[-CAType <CAType>]
[-CertFile <String>] [-CertFilePassword <SecureString>] [-CertificateID <String>] [-
Credential <PSCredential>]
[-DatabaseDirectory <String>] [-Force [<SwitchParameter>]] [-LogDirectory <String>] [-
OverwriteExistingDatabase
[<SwitchParameter>]] [-OverwriteExistingKey [<SwitchParameter>]] [-Confirm
[<SwitchParameter>]] [-WhatIf
[<SwitchParameter>]] [<CommonParameters>]
```

The ExistingCertificateParameterSet is the preferred for migration because you can use the `-CertificateID` parameter to identify the CA certificate from the Importing the CA certificate section in order to configure for the CA. The value for `-CertificateID` can be either the thumbprint or the serial number of the imported certificate.

To add the CA role service by using Windows PowerShell

1. Right-click **Windows PowerShell** and click **Run as Administrator**.
2. To install the CA role service binaries with the Certification Authority and Certificate Templates MMC tools, type the following command and press ENTER:

   ```
   Add-WindowsFeature ADCS-Cert-Authority -
   IncludeManagementTools
   ```

3. Type the `Install-AdcsCertificationAuthority` cmdlet with the appropriate parameters and press ENTER. For example, to restore an Enterprise Subordinate CA by using the certificate you imported in the Importing the CA certificate section, type the following command and press ENTER:

   ```
   Install-AdcsCertificationAuthority –CAType
   EnterpriseSubordinateCA –CertificateID "YourCertSerialNumber
   or YourCertThumbprint" –credential (get-credential
   domain\administrator)
   ```

   Type the password for a member of the Enterprise Admins group or Domain Admins group as needed.

For more information about using Windows PowerShell to install other AD CS role services, see Deploying AD CS Using Windows PowerShell. For more general information about Windows PowerShell cmdlets for AD CS, see AD CS Deployment Cmdlets in Windows PowerShell.

Restoring the CA database and configuration on the destination server

The procedures in this section should be completed only after the CA role service has been installed on the destination server.

If you are migrating to a failover cluster, add the CA role service to all cluster nodes before restoring the CA database. The CA database should be restored on only one cluster node and must be located on shared storage.

Restoring the source CA backup includes the following tasks:

- Restoring the source CA database on the destination server
- Restoring the source CA registry settings on the destination server
- Verifying certificate extensions on the destination CA
- Restoring the certificate templates list (required only for enterprise CAs)

Restoring the source CA database on the destination server

This section describes two different procedures for restoring the source CA database backup on the destination server.

If you are migrating to a Server Core installation, you must use the procedure "To restore the CA database by using Certutil.exe" or "To restore the CA database by using Windows PowerShell." In general, it is possible to remotely manage a CA running on a Server Core installation by using the Certification Authority snap-in and Server Manager; however, it is only possible to restore a CA database remotely by using Windows PowerShell.

If you are migrating to a failover cluster, ensure that shared storage is online and restore the CA database on only one cluster node.

To restore the CA database by using the Certification Authority snap-in

1. Log on to the destination server by using an account that is a CA administrator.
2. Start the Certification Authority snap-in.
3. Right-click the node with the CA name, point to **All Tasks**, and then click **Restore CA**.
4. On the **Welcome** page, click **Next**.
5. On the **Items to Restore** page, select **Certificate database and certificate database log**.
6. Click **Browse**. Navigate to the parent folder that holds the **Database** folder (the folder that contains the CA database files created during the CA database backup).

   Caution
   Do not select the **Database** folder. Select its parent folder.
7. Click **Next** and then click **Finish**.
8. Click **Yes** to start the CA service (certsvc).

To restore only the CA database by using Windows PowerShell

1. Log on to the destination server by using an account that is a CA administrator.
2. Right-click **Windows PowerShell** and click **Run as Administrator**.
3. Type the following command to stop the CA service and press ENTER:

```
Stop-service certsvc
```

4. Type the following command and press ENTER:

```
Restore-CARoleService –path < CA Database Backup Directory> -
DatabaseOnly - Force
```

> **Note**
> The value of <*CA Database Backup Directory*> is the parent directory of the **Database** directory. For example, if the CA database backup files are located in C:\Temp\Database, then the value of <*CA Database Backup Directory*> is C:\Temp. Include the force flag because an empty CA database will already be present after you perform the steps in Adding the CA role service by using Server Manager.

5. Type the following command to restart the CA service and press ENTER:

```
Start-service certsvc
```

To restore the only CA database by using Certutil.exe

1. Log on to the destination server by using an account that is a CA administrator.
2. Open a Command Prompt window.
3. Type the following command to stop the CA service and press ENTER:

```
Net stop certsvc
```

4. Type **certutil.exe -f -restoredb** <*CA Database Backup Directory*> and press ENTER.

> **Note**
> The value of <*CA Database Backup Directory*> is the parent directory of the **Database** directory. For example, if the CA database backup files are located in C:\Temp\Database, then the value of <*CA Database Backup Directory*> is C:\Temp.

5. Type the following command to restart the CA service and press ENTER:

```
Net start certsvc
```

Restoring the source CA registry settings on the destination server

The CA configuration information is stored in the registry in:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc

Before importing the registry settings from the source CA to the target CA, create a backup of the default target CA registry configuration by using the procedure Backing up CA registry settings. Be sure to perform these steps on the target CA and to name the registry file a name such as "DefaultRegCfgBackup.reg" to avoid confusion.

Important
Some registry parameters should be migrated without changes from the source CA computer, and some should not be migrated. If they are migrated, they should be updated in the target system after migration because some values are associated with the CA itself, whereas others are associated with the domain environment, the physical host, the Windows version, or other factors that may be different in the target system.

A suggested way of performing the registry configuration import is first to open the registry file you exported from the source CA in a text editor and analyze it for settings that may need to be changed or removed. The following table shows the configuration parameters that should be transferred from the source CA to the target CA.

Registry location	Configuration parameter
HKEY_LOCAL_MACHINE\system\currentcontrolset\services\certsvc\Configuration	LDAPFlags
HKEY_LOCAL_MACHINE\system\currentcontrolset\services\certsvc\Configuration*CAname*	DSConfigDN ForceTeletex CRLEditFlags CRLFlags InterfaceFlags (required only if has been changed manually) EnforceX500NameLengths SubjectTemplate ValidityPeriod ValidityPeriodUnits KRACertHash

Registry location	Configuration parameter
	KRACertCount KRAFlags CRLPublicationURLs CRLPeriod CRLPeriodUnits CRLOverlapPeriod CRLOverlapUnits CRLDeltaPeriod CRLDeltaPeriodUnits CRLDeltaOverlapPeriod CRLDeltaOverlapUnits CACertPublicationURLs (check for custom entries with hard-coded host names or other data specific to the source CA) CACertHash
HKEY_LOCAL_MACHINE\system\currentcontrolset\services\certsvc\Configuration*CAname*\ExitModules\CertificateAuthority_MicrosoftDefault.Exit	PublishCertFlags
HKEY_LOCAL_MACHINE\system\currentcontrolset\services\certsvc\Configuration*CAname*\PolicyModules\CertificateAuthority_MicrosoftDefault.Policy	EnableRequestExtensionList EnableEnrolleeRequestExtensionList DisableExtensionList SubjectAltName SubjectAltName2 RequestDispositio

Registry location	Configuration parameter
	n EditFlags

To analyze the registry file

1. Right-click the .reg text file created by exporting the settings from the source CA.
2. Click **Edit** to open the file in a text editor.
3. If the target CA's computer name is different from the source CA's computer name, search the file for the host name of the source CA computer. For each instance of the host name found, ensure that it is the appropriate value for the target environment. Change the host name, if necessary. Update the **CAServerName** value.

   Important
   If the host name is located in the .reg file as part of the CA name, such as in the **Active** value within the **Configuration** key or the **CommonName** value within the **CAName** key, do not change the setting. The CA name must not be changed as part of the migration. This means the new target CA must have the old CA's name, even if part of that name is the old CA's host name.
4. Check any registry values that indicate local file paths, such as the following, to ensure drive letter names and paths are correct for the target CA. If there is a mismatch between the source and the target CA, either update the values in the file or remove them from the file so that the default settings are preserved on the target CA.

   These storage location settings are elected during CA setup. They exist under the Configuration registry key:

   - DBDirectory
   - DBLogDirectory
   - DBSystemDirectory
   - DBTempDirectory

   The following settings under the Configuration\{*CA Name*} registry key contain, in their default values, a local path. (Alternatively, you can update these values after importing them by using the Certification Authority snap-in. The values are located on the CA properties **Extensions** tab.)

   - CACertPublicationURLs
   - CRLPublicationURLs

Warning

Some registry values are associated with the CA, while others are associated with the domain environment, the physical host computer, the Windows version, or even other role services. Consequently, some registry parameters should be migrated without

changes from the source CA computer and others should not. Any value that is not listed in the .reg text file that is restored on the target CA retains its existing setting or default value. An issue that can occur, if the registry values are not properly verified, is explained in the following TechNet Wiki article: AD: Certification Authority Web Enrollment Configuration Failed 0x80070057 (WIN32: 87).

Remove any registry values that you do not want to import into the target CA. Once the .reg text file is edited, it can be imported into the target CA. By importing the source server registry settings backup into the destination server, the source CA configuration is migrated to the destination server.

To import the source CA registry backup on the destination CA

1. Log on to the destination server as a member of the local Administrators group.
2. Open a Command Prompt window.
3. Type **net stop certsvc** and press ENTER.
4. Type **reg import** *<Registry Settings Backup.reg>* and press ENTER.

To edit the CA registry settings

1. Click **Start**, type **regedit.exe** in the **Search programs and files** box, and press ENTER to open the Registry Editor.
2. In the console tree, locate the key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\CertSvc\Configuration`, and click **Configuration**.
3. In the details pane, double-click **DBSessionCount**.
4. Click **Hexadecimal**. In **Value data**, type **64**, and then click **OK**.
5. Verify the locations specified in the following settings are correct for your destination server, and change them as needed to indicate the location of the CA database and log files.
   - DBDirectory
   - DBLogDirectory
   - DBSystemDirectory
   - DBTempDirectory

   Important
   Complete steps 6 through 8 only if the name of your destination server is different from the name of your source server.
6. In the console tree of the registry editor, expand **Configuration**, and click your CA name.
7. Modify the values of the following registry settings by replacing the source server name with the destination server name.

   Note
   In the following list, CACertFileName and ConfigurationDirectory values are

created only when certain CA installation options are specified. If these two settings are not displayed, you can proceed to the next step.

- CAServerName
- CACertFileName
- ConfigurationDirectory – This value should appear in Windows Registry under the following location:
  `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\CertSvc\Configuration`.

Verifying certificate extensions on the destination CA

The steps described for importing the source CA registry settings and editing the registry in case of a server name change are intended to retain the network locations that were used by the source CA to publish CRLs and CA certificates. If the source CA was published to default Active Directory locations, after completing the previous procedure, there should be an extension with publishing options enabled and an LDAP URL that references the source server's NetBIOS name; for example,

`ldap:///CN=<CATruncatedName><CRLNameSuffix>,CN=<ServerShortName>,CN=CDP,CN=Public Key Services,CN=Services,<ConfigurationContainer><CDPObjectClass>`.

Because many administrators configure extensions that are customized for their network environment, it is not possible to provide exact instructions for configuring CRL distribution point and authority information access extensions.

Carefully review the configured locations and publishing options, and ensure that the extensions are correct according to your organization's requirements.

To verify extensions by using the Certification Authority snap-in

1. Review and modify the CRL distribution point and authority information access extensions and publishing options by following example procedures described in Specify CRL Distribution Points (http://go.microsoft.com/fwlink/?LinkID=145848).
2. If the destination server name is different from the source server name, add an LDAP URL specifying a location that references the destination server's NetBIOS name with the substitution variable <*ServerShortName*>; for example
   `ldap:///CN=<CATruncatedName><CRLNameSuffix>,CN=<ServerShortName>,CN=CDP,CN=Public Key Services,CN=Services,<ConfigurationContainer><CDPObjectClass>`.
3. Ensure that the CDP options are set so that the former CDP location is not included in the CDP extension of newly issued certificates or in the Freshest CRL extension of CRLs.

Restoring the certificate templates list

The following procedure is required only for an enterprise CA. A standalone CA does not have certificate templates.

To assign certificate templates to the destination CA

1. Log on with administrative credentials to the destination CA.
2. Open a command prompt window.
3. Type **certutil -setcatemplates +** <*templatelist*> and press ENTER.

   Note
   Replace <*templatelist*> with a comma-separated list of the template names that are listed in the catemplates.txt file created during the procedure "To record a CA templates list by using Certutil.exe." For example, **certutil -setcatemplates +Administrator,User,DomainController**. Review the list of templates created during Backing up a CA templates list.

Granting permissions on AIA and CDP containers

If the name of the destination server is different from the source server, the destination server must be granted permissions on the source server's CDP and AIA containers in AD DS to publish CRLs and CA certificates. Complete the following procedure in the case of a server name change.

To grant permissions on the AIA and CDP containers

1. Log on as a member of the Enterprise Admins group to a computer on which the Active Directory Sites and Services snap-in is installed. Open Active Directory Sites and Services (dssite.msc).
2. In the console tree, click the top node.
3. On the **View** menu, click **Show services node**.
4. In the console tree, expand **Services**, expand **Public Key Services**, and then click **AIA**.
5. In the details pane, right-click the name of the CA, and then click **Properties**.
6. Click the **Security** tab, and then click **Add**.
7. Click **Object Types**, click **Computers**, and then click **OK**.
8. Type the name of the CA, and click **OK**.
9. In the **Allow** column, click **Full Control**, and click **Apply**.
10. The previous CA computer object is displayed (as **Account Unknown** with a security identifier following it) in **Group or user names**. You can remove that account. To do so, select it and then click **Remove**. Click **OK**.
11. In the console tree, expand **CDP**, and then click the folder with the same name as the CA.
12. In the details pane, right-click the **cRLDistributionPoint** item at the top of the list, and then click **Properties**.
13. Click the **Security** tab, and then click **Add**.
14. Click **Object Types**, click **Computers**, and then click **OK**.
15. Type the name of the destination server, and click **OK**.
16. In the **Allow** column, click **Full Control**, and click **Apply**.

17. The previous CA computer object is displayed (as **Account Unknown** with a security identifier following it) in **Group or user names**. You can remove that account. To do so, select it and then click **Remove**. Click **OK**.
18. Repeat steps 13 through 18 for each **cRLDistributionPoint** item.

Notes

If you are using file//\\computer\share syntax in the CDP Extensions for publishing the CRL to a shared folder location, then you may need to adjust the permissions to that shared folder so that the destination CA has the ability to write to that location.

If you are hosting the CDP on the destination server and using a AIA or CDP path that includes an alias name (for example, pki.contoso.com) for the destination, you may need to adjust the DNS record so that it points to the correct destination IP address.

Additional procedures for failover clustering

If you are migrating to a failover cluster, complete the following procedures after the CA database and registry settings have been migrated to the destination server.

- Configuring failover clustering for the destination CA
- Granting permissions on public key containers
- Editing the DNS name for a clustered CA in AD DS
- Configuring CRL distribution points for failover clusters

Note

Migration of a CA to a failover cluster running on the Server Core installation option of Windows Server 2008 R2 is not described in this guide.

Configuring failover clustering for the destination CA

If you are migrating to a failover cluster, complete the following procedures to configure failover clustering for AD CS.

To configure AD CS as a cluster resource

1. Click **Start**, point to **Run**, type **Cluadmin.msc**, and then click **OK**.
2. In the console tree of the Failover Cluster Management snap-in, click **Services and Applications**.
3. On the **Action** menu, click **Configure a service or Application**. If the **Before you begin** page appears, click **Next**.
4. In the list of services and applications, select **Generic Service**, and click **Next**.
5. In the list of services, select **Active Directory Certificate Services**, and click **Next**.
6. Specify a service name, and click **Next**.
7. Select the disk storage that is still mounted to the node, and click **Next**.
8. To configure a shared registry hive, click **Add**, type

SYSTEM\CurrentControlSet\Services\CertSvc, and then click **OK**. Click **Next** twice.

9. Click **Finish** to complete the failover configuration for AD CS.
10. In the console tree, double-click **Services and Applications**, and select the newly created clustered service.
11. In the details pane, click **Generic Service**. On the **Action** menu, click **Properties**.
12. Change **Resource Name** to **Certification Authority**, and click **OK**.

If you use a hardware security module (HSM) for your CA, complete the following procedure.

To create a dependency between a CA and the network HSM service

1. Open the Failover Cluster Management snap-in. In the console tree, click **Services and Applications**.
2. In the details pane, select the previously created name of the clustered service.
3. On the **Action** menu, click **Add a resource**, and then click **Generic Service**.
4. In the list of available services displayed by the **New Resource** wizard, click the name of the service that was installed to connect to your network HSM. Click **Next** twice, and then click **Finish**.
5. Under **Services and Applications** in the console tree, click the name of the clustered services.
6. In the details pane, select the newly created **Generic Service**. On the **Action** menu, click **Properties**.
7. On the **General** tab, change the service name if desired, and click **OK**. Verify that the service is online.
8. In the details pane, select the service previously named **Certification Authority**. On the **Action** menu, click **Properties**.
9. On the **Dependencies** tab, click **Insert**, select the network HSM service from the list, and click **OK**.

Granting permissions on public key containers

If you are migrating to a failover cluster, complete the following procedures to grant all cluster nodes permissions to on the following AD DS containers:

- The AIA container
- The Enrollment container
- The KRA container

To grant permissions on public key containers in AD DS

1. Log on to a domain member computer as a member of the Domain Admins group or Enterprise Admins group.
2. Click **Start**, point to **Run**, type **dssite.msc**, and then click **OK**.

3. In the console tree, click the top node.
4. On the **View** menu, click **Show services node**.
5. In the console tree, expand **Services**, then **Public Key Services**, and then click **AIA**.
6. In the details pane, right-click the name of the source CA, and then click **Properties**.
7. Click the **Security** tab, and then click **Add**.
8. Click **Object Types**, click **Computers**, and then click **OK**.
9. Type the computer account names of all cluster nodes, and click **OK**.
10. In the **Allow** column, select the **Full Control** check box next to each cluster node, and click **OK**.
11. In the console tree, click **Enrollment Services**.
12. In the details pane, right-click the name of the source CA, and then click **Properties**.
13. Click the **Security** tab, and then click **Add**.
14. Click **Object Types**, click **Computers**, and then click **OK**.
15. Type the computer account names of all cluster nodes, and click **OK**.
16. In the **Allow** column, select the **Full Control** check box next to each cluster node, and click **OK**.
17. In the console tree, click **KRA**.
18. In the details pane, right-click the name of the source CA, then click **Properties**.
19. Click the **Security** tab, and then click **Add**.
20. Click **Object Types**, click **Computers**, and then click **OK**.
21. Type the names of all cluster nodes, and click **OK**.
22. In the **Allow** column, select the **Full Control** check box next to each cluster node, and click **OK**.

Editing the DNS name for a clustered CA in AD DS

When the CA service was installed on the first cluster node, the Enrollment Services object was created and the DNS name of that cluster node was added to the dNSHostName attribute of the Enrollment Services object. Because the CA must operate on all cluster nodes, the value of the dNSHostName attribute of the Enrollment Services object must be the service name specified in step 6 of the procedure "To configure AD CS as a cluster resource."

If you are migrating to a clustered CA, complete the following procedure on the active cluster node. It is necessary to complete the procedure on only one cluster node.

To edit the DNS name for a clustered CA in AD DS

1. Log on to the active cluster node as a member of the Enterprise Admins group.
2. Click **Start**, point to **Run**, type **adsiedit.msc**, and then click **OK**.
3. In the console tree, click **ADSI Edit**.
4. On the **Action** menu, click **Connect to**.
5. In the list of well-known naming contexts, click **Configuration**, and click **OK**.

6. In the console tree, expand **Configuration**, **Services**, and **Public Key Services**, and click **Enrollment Services**.
7. In the details pane, right-click the name of the cluster CA, and click **Properties**.
8. Click **dNSHostName**, and click **Edit**.
9. Type the service name of the CA as displayed under **Failover Cluster Management** in the Failover Cluster Manager snap-in, and click **OK**.
10. Click **OK** to save changes.

Configuring CRL distribution points for failover clusters

In a CA's default configuration, the server's short name is used as part of the CRL distribution point and authority information access locations.

When a CA is running on a failover cluster, the server's short name must be replaced with the cluster's short name in the CRL distribution point and authority information access locations. To publish the CRL in AD DS, the CRL distribution point container must be added manually.

Important
The following procedures must be performed on the active cluster node.

To change the configured CRL distribution points

1. Log on to the active cluster node as a member of the local Administrators group.
2. Click **Start**, click **Run**, type **regedit**, and then click **OK**.
3. Locate the registry key **\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration**.
4. Click the name of the CA.
5. In the right pane, double-click **CRLPublicationURLs**.
6. In the second line, replace **%2** with the service name specified in step 6 of the procedure "To configure AD CS as a cluster resource."

   Tip
   The service name also appears in the Failover Cluster Management snap-in under **Services and Applications**.
7. Restart the CA service.
8. Open a command prompt, type **certutil -CRL**, and press ENTER.

   Note
   If a "Directory object not found" error message is displayed, complete the following procedure to create the CRL distribution point container in AD DS.

To create the CRL distribution point container in AD DS

1. At a command prompt, type **cd %windir%\System32\CertSrv\CertEnroll**, and press

ENTER. The CRL file created by the certutil –CRL command should be located in this directory.

2. To publish the CRL in AD DS, type **certutil -f -dspublish** *"CRLFile.crl"* and press ENTER.

Next steps

After completing the procedures to migrate the CA, you should complete the procedures described in Verifying the Certification Authority Migration.

See also

- Active Directory Certificate Services Migration Guide for Windows Server 2012 R2
- Prepare to Migrate
- Verifying the Certification Authority Migration
- Post-Migration Tasks
- Migrating Roles and Features in Windows Server

Verifying the Certification Authority Migration

Complete the following procedures to verify the operation of the destination certification authority (CA).

- Verifying certificate enrollment
- Verifying CRL publishing

Verifying certificate enrollment

To verify migration to an enterprise CA, complete the procedure Request a Certificate (http://go.microsoft.com/fwlink/?LinkId=179367).

You can start autoenrollment for user certificates by completing the following procedure or by running the following command: **certutil.exe -pulse**.

To verify autoenrollment

1. Log on to a domain member computer by using an account that has Autoenroll, Enroll, and Read permissions for the certificate templates that are assigned to the destination CA.
2. Click **Start**, and then click **Run**.
3. Type **certmgr.msc**, and then click **OK** to open the Certificates snap-in.
4. In the console tree, right-click **Certificates – Current User**, click **All Tasks**, and then click **Automatically Enroll and Retrieve Certificates** to start the Certificate Enrollment

wizard.

5. On the **Before You Begin** page, click **Next**.
6. On the **Request Certificates** page, a list of one or more certificate templates should be displayed. Select the check box next to each certificate template that you want to request, and then click **Enroll**.

   Note
   If the correct certificate templates are not displayed, click **Show all templates** to display all certificate templates that are assigned to the issuing CA. A status of **Unavailable** indicates the user account does not have permission to autoenroll for a certificate. Follow the steps in the "To configure certificate templates for autoenrollment" procedure earlier in this topic.
7. Click **Finish** to complete the enrollment process.
8. In the console tree, double-click **Personal**, and then click **Certificates** to display a list of installed user certificates and to verify that the certificate that you requested is displayed.

To verify migration to a standalone CA, complete the following procedure.

To verify manual enrollment by using Certreq.exe

1. Create a certificate request, and save it to a file by completing the procedure Create a Custom Certificate Request (http://go.microsoft.com/fwlink/?LinkId=179368).
2. Open a Command Prompt window.
3. Type **certreq -submit -config "<DestinationServerName\CAName>" "<CertificateRequestInput>" "<CertificateResponseOutput>"** and press ENTER.

   Note
   If a message is displayed indicating that the certificate request is pending, the certificate must be issued by a certificate manager or CA administrator by using the Certification Authority snap-in. After the certificate is issued, it must be retrieved by using the command in step 4. If the certificate is issued immediately by the CA, the file specified in <CertificateResponseOutput> contains the certificate. Use the command in step 5 to install the certificate into the certificate store.
4. Type **certreq –retrieve -config "<DestinationServerName\CAName>" <RequestID> <CertificateResponseOutput>** and press ENTER.
5. Type **certreq –accept -config "<DestinationServerName\CAName>" <CertificateResponseOutput>** and press ENTER.

Option	Description	Example
-config	The –config option is followed by a string specifying a host name and CA name in the	Certreq.exe –submit –config Server1\CA1 C:\RequestFile.txt

Option	Description	Example
	format HostName\CAName.	C:\ResponseFile.cer
DestinationServerName	The host name of the destination server.	
CAName	The CA name being migrated.	
CertificateRequestInput	The path and name of the file containing the certificate request that was created by using the procedure "Create a Custom Certificate Request."	
CertificateResponseOutput	The path and name of the file receiving the issued certificate from the CA. If the certificate request is pending, the file contains a message from the CA indicating the status of the request and the request ID. The request ID is used to retrieve the certificate after it is issued by a certificate manager or CA administrator.	

Verifying CRL publishing

If you published a certificate revocation list (CRL) with an extended validity period before beginning migration, you should change the CRL publishing period back to its pre-migration value by completing the procedure Schedule the publication of the certificate revocation list.

Manually publish a CRL by completing one of the procedures described in Manually Publish a CRL.

Next steps

After completing verification steps, you should review the topic Post-Migration Tasks and complete the procedures appropriate for your environment.

See also

- Active Directory Certificate Services Migration Guide for Windows Server 2012 R2
- Prepare to Migrate
- Migrating the Certification Authority
- Post-Migration Tasks
- Migrating Roles and Features in Windows Server

Post-Migration Tasks

Post-migration steps can be performed after migration has been completed and the operation of the destination CA has been verified.

If verification steps have failed, review the Troubleshooting section in this topic.

- Upgrading certificate templates in Active Directory Domain Services (AD DS)
- Retrieving certificates after a host name change
- Restoring Active Directory Certificate Services (AD CS) to the source server in the event of migration failure
- Troubleshooting migration

Upgrading certificate templates in Active Directory Domain Services (AD DS)

Review the post-migration steps below and perform only those that are appropriate for your environment and migration scenario.

The following additional default certificate templates are included in enterprise certification authorities (CAs) running on Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2 and Windows Server 2008 but are not included in Windows Server 2003:

- OCSP Response Signing
- Kerberos Authentication

These certificate templates are not required for CA operation. OCSP Response Signing certificates are required if you are deploying the Online Responder role service.

If you require these additional certificate templates, complete the following procedure.

To upgrade certificate templates in AD DS by using the Certificate Templates snap-in

1. Log on to the destination server as a member of the Enterprise Admins group.
2. Open the Certificate Templates snap-in. The snap-in automatically adds the default certificate templates to AD DS.

Retrieving certificates after a host name change

If the destination server name is different from the source server name, you might need to manually retrieve any certificates that were issued by the source CA and had not been retrieved before migration.

Complete this procedure on the computer that was used to submit the certificate request to the source CA.

To retrieve a certificate by using Certreq.exe

1. Open a Command Prompt window.
2. Type **certreq –retrieve -config "<DestinationServerName\CAName>" <RequestID> <CertificateResponseOutput>** and press ENTER.
3. Type **certreq –accept <CertificateResponseOutput>** and press ENTER.

Option	Description	Example
-config	The –config option is followed by a string specifying a host name and CA name in the format HostName\CAName.	Certreq.exe –submit –config Server1\CA1 C:\RequestFile.txt C:\ResponseFile.cer
DestinationServerName	The host name of the destination server.	
CAName	The CA name being migrated.	
CertificateRequestInput	The path and name of the file containing the certificate request that was created by using the procedure "Create a Custom Certificate Request."	
CertificateResponseOutput	The path and name of the file receiving the issued certificate from the CA. If the certificate request is pending, the file contains a message from the CA indicating the status of the request and the request ID. The request ID is used to retrieve the certificate after it is issued by a certificate manager or CA administrator.	
RequestID	The Request ID value returned	

Option	Description	Example
	by a CA in response to a certificate request. The Request ID value is displayed in command output and written to the CertificateResponseOutput file.	

Restoring Active Directory Certificate Services (AD CS) to the source server in the event of migration failure

If you removed the CA role service from the source server as described in the procedure Removing the CA role service from the source server, you can restore the source CA by reinstalling the CA role service on the source server. It is important to remove the CA role service from the destination server before reinstalling the CA role service on the source server.

If you did not remove the CA role service from the source server, you should not remove the CA role service from the destination server. Simply shut down the destination CA and start the source CA.

Rollback procedures can be completed in less than one hour.

To remove the CA role service from the destination server, use the Remove Roles Wizard in Server Manager.

To add the CA role service to a source server running Windows Server 2003, use the Add/Remove Windows Components wizard.

To add the CA role service to a source server running Windows Server 2008 or later, use the Add Roles Wizard in Server Manager.

Troubleshooting migration

If you encounter errors during verification procedures, use Event Viewer to review the Application log on the destination CA. View an Error event in the preview pane or event properties, and click **Event Log Online Help** to open a Web page with troubleshooting procedures for that event.

For the full collection of documented AD CS events, see AD CS Events and Errors.

See also

- Active Directory Certificate Services Migration Guide for Windows Server 2012 R2
- Prepare to Migrate
- Migrating the Certification Authority

- Verifying the Certification Authority Migration
- Migrating Roles and Features in Windows Server

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2

About this guide

This guide provides instructions to migrate the following role services to Active Directory Federation Services (AD FS) that is installed with Windows Server® 2012 R2:

- AD FS 2.0 federation server installed on Windows Server 2008 or Windows Server 2008 R2
- AD FS federation server installed on Windows Server 2012

Target audience

- IT architects who are responsible for computer management and security throughout an organization
- IT operations engineers who are responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications
- IT operations managers who are accountable for network and server management

Supported migration scenarios

The migration instructions in this guide consist of the following tasks:

- Exporting the AD FS 2.0 configuration data from your server that is running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012
- Performing an in-place upgrade of the operating system of this server from Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012 to Windows Server® 2012 R2
- Recreating the original AD FS configuration and restoring the remaining AD FS service settings on this server, which is now running the AD FS server role that is installed with Windows Server® 2012 R2.

This guide does not include instructions to migrate a server that is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration process specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Supported operating systems

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2008, both full and Server Core installation options	Windows Server® 2012 R2 (Server Core and full installation options)	x64-based
x64-based	Windows Server 2008 R2		
x64-based	Server Core installation option of Windows Server 2008 R2		
x64-based	Server Core and full installation options of Windows Server 2012		

Notes

- The versions of operating systems that are listed in the preceding table are the oldest combinations of operating systems and service packs that are supported.
- The Foundation, Standard, Enterprise, and Datacenter editions of the Windows Server operating system are supported as the source or the destination server.
- Migrations between physical operating systems and virtual operating systems are supported.

Supported AD FS role services and features

The following table describes the migration scenarios of the AD FS role services and their respective settings that are described in this guide.

From	To AD FS installed with Windows Server® 2012 R2
AD FS 2.0 federation server installed on Windows Server 2008 or Windows Server 2008 R2	Migration on the same server is supported. For more information, see: • Preparing to Migrate the AD FS Federation Server • Migrating the AD FS Federation Server
AD FS federation server installed on Windows Server 2012	Migration on the same server is supported. For more information see: • Preparing to Migrate the AD FS Federation Server • Migrating the AD FS Federation Server

See Also

Preparing to Migrate the AD FS Federation Server

Migrating the AD FS Federation Server

Migrating the AD FS Federation Server Proxy

Verifying the AD FS Migration to Windows Server 2012 R2

Preparing to Migrate the AD FS Federation Server

To perform the same server migration of your AD FS federation server farm for Windows Server 2012 R2, you must review the following information:

Note
The information below applies to migrating a one-node federation server, as well as a WID or a SQL Server federation server farm. It applies to the migration of a federation server running AD FS 2.0 running on Windows Server 2008 or Windows Server 2008 R2 or AD FS installed with Windows Server 2012.

- Migration Process Outline
- New AD FS functionality in Windows Server 2012 R2
- AD FS Requirements in Windows Server 2012 R2
- Increasing your Windows PowerShell limits
- Other migration tasks and considerations

Migration Process Outline

To complete the migration of your AD FS federation server farm to Windows Server 2012 R2, you must complete the following tasks:

1. Export, record, and backup the following configuration data in your existing AD FS farm. For detailed instructions on how to complete these tasks, see Migrating the AD FS Federation Server.

   The following settings are migrated with the scripts located in the \support\adfs folder on the Windows Server 2012 R2 installation CD:

   - Claims provider trusts, with the exception of custom claim rules on the Active Directory Claims provider trust. For more information, see Migrating the AD FS Federation Server.
   - Relying party trusts.
   - AD FS internally generated, self-signed token signing and token decryption certificates.

   Any of the following custom settings must be migrated manually:

- Service settings:
  - Non-default token signing and token decryption certificates that were issued by an enterprise or public certification authority.
  - The SSL server authentication certificate used by AD FS.
  - The service communications certificate used by AD FS (by default, this is the same certificate as the SSL certificate.
  - Non-default values for any federation service properties, such as AutoCertificateRollover or SSO lifetime.
  - Non-default AD FS endpoint settings and claim descriptions.
- Custom claim rules on the Active Directory claims provider trust.
- AD FS sign-in page customizations

For more information, see Migrating the AD FS Federation Server.

2. Create a Windows Server 2012 R2 federation server farm.
3. Import the original configuration data into this new Windows Server 2012 R2 AD FS farm.
4. Configure and customize the AD FS sign-in pages.

New AD FS functionality in Windows Server 2012 R2

The following AD FS functionality changes in Windows Server 2012 R2 impact a migration from AD FS 2.0 or AD FS in Windows Server 2012:

- **IIS dependency**

  AD FS in Windows Server 2012 R2 is self-hosted and does not require IIS installation. Make sure you note the following as a result of this change:

  - SSL certificate management for both federation servers and proxy computers in your AD FS farm must now be performed via Windows PowerShell.
- **Changes to AD FS sign-in pages' settings and customizations**

  In AD FS in Windows Server 2012 R2, there are several changes intended to improve the sign-in experience for both administrators and users. The IIS-hosted web pages that existed in the previous version of AD FS are now removed. The look and feel of the AD FS sign-in web pages are self-hosted in AD FS and can now be customized to tailor the user experience. The changes include:

  - Customizing the AD FS sign-in experience, including the customization of the company name, logo, illustration, and sign-in description.
  - Customizing the error messages.
  - Customizing the ADFS Home Realm Discovery experience, which includes the following:
    - Configuring your identity provider to use certain email suffixes.
    - Configuring an identity provider list per relying party.
    - Bypassing Home Realm Discovery for intranet.
    - Creating custom web themes.

For detailed instructions on configuring the look and feel of the AD FS sign-in pages, see Customizing the AD FS Sign-in Pages.

If you have web page customization in your existing AD FS farm that you want to migrate to Windows Server 2012 R2, you can recreate them as part of the migration process using the new customization features in Windows Server 2012 R2.

- **Other changes**
  - AD FS in Windows Server 2012 R2 is based on Windows Identity Foundation (WIF) 3.5, not WIF 4.5. Therefore, some specific features of WIF 4.5 (for example, Kerberos claims and dynamic access control) are not supported in AD FS in Windows Server 2012 R2.
  - Device Registration Service (DRS) in Windows Server 2012 R2 operates on port 443; ClientTLS for user certificate authentication operates on port 49443
    - For active, non-browser clients using certificate transport mode authentication that are specifically hard-coded to point to port 443, a code change is required to continue to use user certificate authentication on port 49443.
    - For passive applications no change is required because AD FS redirects to the correct port for user certificate authentication.
    - Firewall ports between the client and the proxy must enable port 49443 traffic to pass through for user certificate authentication.

AD FS Requirements in Windows Server 2012 R2

In order to successfully migrate your AD FS farm to Windows Server 2012 R2, you must meet the following requirements:

For AD FS to function, each computer that you want to be a federation server must be joined to a domain.

For AD FS running on Windows Server 2012 R2 to function, your Active Directory domain must run either of the following:

- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2
- Windows Server 2008

If you plan to use a group Managed Service Account (gMSA) as the service account for AD FS, you must have at least one domain controller in your environment that is running on Windows Server 2012 or Windows Server 2012 R2 operating system.

If you plan to deploy Device Registration Service (DRS) for AD Workplace Join as a part of your AD FS deployment, the AD DS schema needs to be updated to the Windows Server 2012 R2level. There are three ways to update the schema:

1. In an existing Active Directory forest, run adprep /forestprep from the \support\adprep folder of the Windows Server 2012 R2 operating system DVD on any 64-bit server that runs Windows Server 2008 or later. In this case, no additional domain controller needs to be installed, and no existing domain controllers need to be upgraded.

To run adprep/forestprep, you must be a member of the Schema Admins group, the Enterprise Admins group, and the Domain Admins group of the domain that hosts the schema master.

2. In an existing Active Directory forest, install a domain controller that runs Windows Server 2012 R2. In this case, adprep /forestprep runs automatically as part of the domain controller installation.

   During the domain controller installation, you may need to specify additional credentials in order to run adprep /forestprep.

3. Create a new Active Directory forest by installing AD DS on a server that runs Windows Server 2012 R2. In this case, adprep /forestprep does not need to be run because the schema will be initially created with all the necessary containers and objects to support DRS.

SQL Server support for AD FS in Windows Server 2012 R2

If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

Increasing your Windows PowerShell limits

If you have more than 1000 claims provider trusts and relying party trusts in your AD FS farm, or if you see the following error while trying to run the AD FS migration export/import tool, you must increase your Windows PowerShell limits:

```
'Exception of type 'System.OutOfMemoryException' was thrown. At
E:\dev\ds\security\ADFSv2\Product\Migration\Export-FederationConfiguration.ps1:176
char:21 + $configData = Invoke-Command -ScriptBlock $GetConfig -Argume ...
```

This error is thrown because the Windows PowerShell session default memory limit is too low. In Windows PowerShell 2.0, the session default memory is 150MB. In Windows PowerShell 3.0, the session default memory is 1024MB. You can verify Windows PowerShell remote session memory limit using the following command: `Get-Item wsman:localhost\Shell\MaxMemoryPerShellMB`. You can increase the limit by running the following command: `Set-Item wsman:localhost\Shell\MaxMemoryPerShellMB 512`.

Other migration tasks and considerations

In order to successfully migrate your AD FS farm to Windows Server 2012 R2, make sure you are aware of the following:

- The migration scripts located in the \support\adfs folder on the Windows Server 2012 R2 installation CD require that you retain the same federation server farm name and service account identity name that you used in your legacy AD FS farm when you migrate it to Windows Server 2012 R2.
- If you want to migrate a SQL Server AD FS farm, note that the migration process involves creating a new SQL database instance into which you must import the original configuration data.

See Also

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2

Migrating the AD FS Federation Server

To migrate an AD FS federation server that belongs to a single-node AD FS farm, a WIF farm, or a SQL Server farm to Windows Server 2012 R2, you must perform the following tasks:

1. Export and backup the AD FS configuration data
2. Create a Windows Server 2012 R2 federation server farm
3. Import the original configuration data into the Windows Server 2012 R2 AD FS farm
4. **Configure and customize the AD FS sign-in pages in the migrated AD FS farm**
5. **Other migration tasks**

Export and backup the AD FS configuration data

To export the AD FS configuration settings, perform the following procedures:

- To export service settings
- To export claims provider trusts and relying party trusts
- **To export relying party trusts**
- To back up custom attribute stores
- **To back up AD FS sign-in pages' settings and customizations**

To export service settings

1. Make sure that you have access to the following certificates and their private keys in a .pfx file:
   - The SSL certificate that is used by the federation server farm that you want to migrate
   - The service communication certificate (if it is different from the SSL certificate) that is used by the federation server farm that you want to migrate
   - All third-party party token-signing or token-encryption/decryption certificates that are used by the federation server farm that you want to migrate

   To find the SSL certificate, open the Internet Information Services (IIS) management console, Select **Default Web Site** in the left pane, click **Bindings…** in the **Action** pane, find and select the https binding, click **Edit**, and then click **View**.

   You must export the SSL certificate used by the federation service and its private key to a .pfx file. For more information, see Export the Private Key Portion of a Server Authentication Certificate.

   Note

If you plan to deploy the Device Registration Service as part of running your AD FS in Windows Server 2012 R2, you must obtain a new SSL cert. For more information, see Enroll an SSL Certificate for AD FS and Configure a federation server with Device Registration Service.

To view the token signing, token decryption and service communication certificates that are used, run the following Windows PowerShell command to create a list of all certificates in use in a file:

```
Get-ADFSCertificate | Out-File ".\certificates.txt"
```

2. Export AD FS federation service properties, such as the federation service name, federation service display name, and federation server identifier to a file.

   To export federation service properties, open Windows PowerShell and run the following command: `PSH:> Get-ADFSProperties | Out-File ".\properties.txt".`

   The output file will contain the following important configuration values:

Federation Service Property name as reported by Get-ADFSProperties	Federation Service Property name in AD FS management console
HostName	Federation Service name
Identifier	Federation Service identifier
DisplayName	Federation Service display name

3. Back up the application configuration file. Among other settings, this file contains the policy database connection string.

   To back up the application configuration file, you must manually copy the `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config` file to a secure location on a backup server.

   Notes

   Make note of the database connection string in this file, located immediately after "policystore connectionstring="). If the connection string specifies a SQL Server database, the value is needed when restoring the original AD FS configuration on the federation server.

   The following is an example of a WID connection string: `"Data Source=\\.\pipe\mssql$microsoft##ssee\sql\query;Initial Catalog=AdfsConfiguration;Integrated Security=True"`. The following is an example of a SQL Server connection string: `"Data Source=databasehostname;Integrated Security=True"`.

4. Record the identity of the AD FS federation service account and the password of this account.

   To find the identity value, examine the **Log On As** column of **AD FS 2.0 Windows**

Service in the **Services** console and manually record this value.

Note

For a stand-alone federation service, the built-in NETWORK SERVICE account is used. In this case, you do not need to have a password.

5. Export the list of enabled AD FS endpoints to a file.

   To do this, open Windows PowerShell and run the following command: `PSH:> Get-ADFSEndpoint | Out-File ".\endpoints.txt".`

6. Export any custom claim descriptions to a file.

   To do this, open Windows PowerShell and run the following command: `Get-ADFSClaimDescription | Out-File ".\claimtypes.txt".`

7. If you have custom settings such as useRelayStateForIdpInitiatedSignOn configured in the web.config file, ensure you back up the web.config file for reference. You can copy the file from the directory that is mapped to the virtual path **"/adfs/ls"** in IIS. By default, it is in the **%systemdrive%\inetpub\adfs\ls** directory.

To export claims provider trusts and relying party trusts

1. To export AD FS claims provider trusts and relying party trusts, you must log in as Administrator (however, not as the Domain Administrator) onto your federation server and run the following Windows PowerShell script that is located in the **media/server_en-us/support/adfs** folder of the Windows Server 2012 R2 installation CD: `export-federationconfiguration.ps1`.

   Important

   The export script takes the following parameters:

   - Export-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-CertificatePassword <securestring>]
   - Export-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-CertificatePassword <securestring>] [-RelyingPartyTrustIdentifier <string[]>] [-ClaimsProviderTrustIdentifier <string[]>]
   - Export-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-CertificatePassword <securestring>] [-RelyingPartyTrustName <string[]>] [-ClaimsProviderTrustName <string[]>]

     -RelyingPartyTrustIdentifier <string[]> - the cmdlet only exports relying party trusts whose identifiers are specified in the string array. The default is to export NONE of the relying party trusts. If none of RelyingPartyTrustIdentifier, ClaimsProviderTrustIdentifier, RelyingPartyTrustName, and ClaimsProviderTrustName is specified, the script will export all relying party trusts and claims provider trusts.

     -ClaimsProviderTrustIdentifier <string[]> - the cmdlet only exports claims provider trusts whose identifiers are specified in the string array. The default is to export NONE of the claims provider trusts.

-RelyingPartyTrustName <string[]> - the cmdlet only exports relying party trusts whose names are specified in the string array. The default is to export NONE of the relying party trusts.

-ClaimsProviderTrustName <string[]> - the cmdlet only exports claims provider trusts whose names are specified in the string array. The default is to export NONE of the claims provider trusts.

-Path <string> - the path to a folder that will contain the exported files.

-ComputerName <string> - specifies the STS server host name. The default is the local computer. If you are migrating AD FS 2.0 or AD FS in Windows Server 2012 to AD FS in Windows Server 2012 R2, this is the host name of the legacy AD FS server.

-Credential <PSCredential> - specifies a user account that has permission to perform this action. The default is the current user.

-Force – specifies to not prompt for user confirmation.

-CertificatePassword <SecureString> - specifies a password for exporting AD FS certificates' private keys. If not specified, the script will prompt for a password if an AD FS certificate with private key needs to be exported.

Inputs: None

Outputs: string - this cmdlet returns the export folder path. You can pipe the returned object to Import-FederationConfiguration.

To back up custom attribute stores

1. You must manually export all custom attribute stores that you want to keep in your new AD FS farm in Windows Server 2012 R2.

   Note
   In Windows Server 2012 R2, AD FS requires custom attribute stores that are based on .NET Framework 4.0 or above. Follow the instructions in Microsoft .NET Framework 4.5 to install and setup .Net Framework 4.5.

   You can find information about custom attribute stores in use by AD FS by running the following Windows PowerShell command: `PSH:>Get-ADFSAttributeStore`. The steps to upgrade or migrate custom attribute stores vary.
2. You must also manually export all .dll files of the custom attribute stores that you want to keep in your new AD FS farm in Windows Server 2012 R2. The steps to upgrade or migrate .dll files of custom attribute stores vary.

Create a Windows Server 2012 R2 federation server farm

1. Install the Windows Server 2012 R2 operating system on a computer that you want to function as a federation server and then add the AD FS server role. For more information, see Install the AD FS Role Service. Then configure your new federation service either through the Active Directory Federation Service Configuration Wizard or via Windows PowerShell. For more information, see “Configure the first federation server in a new federation server farm” in Configure a Federation Server.

   While completing this step, you must follow these instructions:

   - You must have Domain Administrator privileges in order to configure your federation service.
   - You must use the same federation service name (farm name) as was used in the AD FS 2.0 or AD FS in Windows Server 2012. If you do not use the same federation service name, the certificates that you backed up will not function in the Windows Server 2012 R2 federation service that you are trying to configure.
   - Specify whether this is a WID or SQL Server federation server farm. If it is a SQL farm, specify the SQL Server database location and instance name.
   - You must provide a pfx file containing the SSL server authentication certificate that you backed up as part of preparing for the AD FS migration process.
   - You must specify the same service account identity that was used in the AD FS 2.0 or AD FS in Windows Server 2012 farm.
2. Once the initial node is configured, you can add additional nodes to your new farm. For more information, see “Add a federation server to an existing federation server farm” in Configure a Federation Server.

Import the original configuration data into the Windows Server 2012 R2 AD FS farm

Now that you have an AD FS federation server farm running in Windows Server 2012 R2, you can import the original AD FS configuration data into it.

1. Import and configure other custom AD FS certificates, including externally enrolled token-signing and token- decryption/encryption certificates, and the service communication certificate if it is different from the SSL certificate.

   In the AD FS management console, select **Certificates**. Verify the service communications, token-encryption/decryption, and token-signing certificates by checking each against the values you exported into the certificates.txt file while preparing for the migration.

   To change the token-decrypting or token-signing certificates from the default self-signed certificates to external certificates, you must first disable the automatic certificate rollover feature that is enabled by default. To do this, you can use the following Windows PowerShell command:

```
Set-ADFSProperties –AutoCertificateRollover $false.
```

2. Configure any custom AD FS service settings such as AutoCertificateRollover or SSO lifetime using the Set-AdfsProperties cmdlet.
3. To import AD FS relying party trusts and claims provider trusts, you must be logged in as Administrator (however, not as the Domain Administrator) onto your federation server and run the following Windows PowerShell script that is located in the \support\adfs folder of the Windows Server 2012 R2 installation CD:

```
import-federationconfiguration.ps1
```

Important

The import script takes the following parameters:

- Import-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-LogPath <string>] [-CertificatePassword <securestring>]
- Import-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-LogPath <string>] [-CertificatePassword <securestring>] [-RelyingPartyTrustIdentifier <string[]>] [-ClaimsProviderTrustIdentifier <string[]>
- Import-FederationConfiguration.ps1 -Path <string> [-ComputerName <string>] [-Credential <pscredential>] [-Force] [-LogPath <string>] [-CertificatePassword <securestring>] [-RelyingPartyTrustName <string[]>] [-ClaimsProviderTrustName <string[]>]

  -RelyingPartyTrustIdentifier <string[]> - the cmdlet only imports relying party trusts whose identifiers are specified in the string array. The default is to import NONE of the relying party trusts. If none of RelyingPartyTrustIdentifier, ClaimsProviderTrustIdentifier, RelyingPartyTrustName, and ClaimsProviderTrustName is specified, the script will import all relying party trusts and claims provider trusts.

  -ClaimsProviderTrustIdentifier <string[]> - the cmdlet only imports claims provider trusts whose identifiers are specified in the string array. The default is to import NONE of the claims provider trusts.

  -RelyingPartyTrustName <string[]> - the cmdlet only imports relying party trusts whose names are specified in the string array. The default is to import NONE of the relying party trusts.

  -ClaimsProviderTrustName <string[]> - the cmdlet only imports claims provider trusts whose names are specified in the string array. The default is to import NONE of the claims provider trusts.

  -Path <string> - the path to a folder that contains the configuration files to be imported.

  -LogPath <string> - the path to a folder that will contain the import log file. A log file named “import.log” will be created in this folder.

  -ComputerName <string> - specifies host name of the STS server. The default

is the local computer. If you are migrating AD FS 2.0 or AD FS in Windows Server 2012 to AD FS in Windows Server 2012 R2, this parameter should be set to the hostname of the legacy AD FS server.

-Credential <PSCredential>- specifies a user account that has permission to perform this action. The default is the current user.

-Force – specifies to not prompt for user confirmation.

-CertificatePassword <SecureString> - specifies a password for importing AD FS certificates' private keys. If not specified, the script will prompt for a password if an AD FS certificate with private key needs to be imported.

Inputs: string - this command takes the import folder path as input. You can pipe Export-FederationConfiguration to this command.

Outputs: None.

Any trailing spaces in the WSFedEndpoint property of a relying party trust may cause the import script to error. In this case, manually remove the spaces from the file prior to import. For example, these entries cause errors:

```
<URI N="WSFedEndpoint">https://127.0.0.1:444 /</URI>
<URI N="WSFedEndpoint">https://myapp.cloudapp.net:83 /</URI>
```

They must be edited to:

```
<URI N="WSFedEndpoint">https://127.0.0.1:444/</URI>
<URI N="WSFedEndpoint">https://myapp.cloudapp.net:83/</URI>
```

Important

If you have any custom claim rules (rules other than the AD FS default rules) on the Active Directory claims provider trust in the source system, these will not be migrated by the scripts. This is because Windows Server 2012 R2 has new defaults. Any custom rules must be merged by adding them manually to the Active Directory claims provider trust in the new Windows Server 2012 R2 farm.

4. Configure all custom AD FS endpoint settings. In the AD FS Management console, select **Endpoints**. Check the enabled AD FS endpoints against the list of enabled AD FS endpoints that you exported to a file while preparing for the AD FS migration.

   - And -

   Configure any custom claim descriptions. In the AD FS Management console, select **Claim Descriptions**. Check the list of AD FS claim descriptions against the list of claim descriptions that you exported to a file while preparing for the AD FS migration. Add any custom claim descriptions included in your file but not included in the default list in AD FS. Note that Claim identifier in the management console maps to the ClaimType in the file.

5. Install and configure all backed up custom attribute stores. As an administrator, ensure any custom attribute store binaries are upgrade to .NET Framework 4.0 or higher before updating the AD FS configuration to point to them.

6. Configure service properties that map to the legacy web.config file parameters.
    - If **useRelayStateForIdpInitiatedSignOn** was added to the **web.config** file in your AD FS 2.0 or AD FS in Windows Sever 2012 farm, then you must configure the following service properties in your AD FS in Windows Server 2012 R2 farm:
        - AD FS in Windows Server 2012 R2 includes a **%systemroot%\ADFS\Microsoft.IdentityServer.Servicehost.exe.config** file. Create an element with the same syntax as the **web.config** file element: `<useRelayStateForIdpInitiatedSignOn enabled="true" />`. Include this element as part of **<microsoft.identityserver.web>** section of the **Microsoft.IdentityServer.Servicehost.exe.config** file.
    - If **<persistIdentityProviderInformation enabled="true|false" lifetimeInDays="90" enablewhrPersistence="true|false" />** was added to the **web.config** file in your AD FS 2.0 or AD FS in Windows Sever 2012 farm, then you must configure the following service properties in your AD FS in Windows Server 2012 R2 farm:

        i. In AD FS in Windows Server 2012 R2, run the following Windows PowerShell command: `Set-AdfsWebConfig -HRDCookieEnabled -HRDCookieLifetime`.
    - If **<singleSignOn enabled="true|false" />** was added to the **web.config** file in your AD FS 2.0 or AD FS in Windows Sever 2012 farm, you do not need to set any additional service properties in your AD FS in Windows Server 2012 R2 farm. Single sign-on is enabled by default in AD FS in Windows Server 2012 R2 farm.
    - If localAuthenticationTypes settings were added to the **web.config** file in your AD FS 2.0 or AD FS in Windows Sever 2012 farm, then you must configure the following service properties in your AD FS in Windows Server 2012 R2 farm:
        - Integrated, Forms, TlsClient, Basic Transform list into equivalent AD FS in Windows Server 2012 R2 has global authentication policy settings to support both federation service and proxy authentication types. These settings can be configured in the AD FS in Management snap-in under the **Authentication Policies**.

After you import the original configuration data, you can customize the AD FS sign in pages as needed. For more information, see Customizing the AD FS Sign-in Pages.

See Also

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2

Migrating the AD FS Federation Server Proxy

In Active Directory Federation Services (AD FS) in Windows Server 2012 R2, the role of a federation server proxy is handled by a new Remote Access role service called Web Application Proxy. In Windows Server 2012 R2, to enable your AD FS for accessibility from outside of the corporate network, you can deploy one or more Web Application Proxies. However, you cannot

migrate a federation server proxy running on Windows Server 2008 R2 or Windows Server 2012 to a Web Application Proxy running on Windows Server 2012 R2.

Important
The migration of a federation server proxy running on Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to a Web Application Proxy running on Windows Server 2012 R2 is NOT supported.

If you want to configure AD FS in a Windows Server 2012 R2 migrated farm for extranet access, you must perform a fresh deployment of one or more Web Application Proxy computers as part of your AD FS infrastructure.

To plan Web Application Proxy deployment, you can review the information in the following topics:

- Step 1: Plan the Web Application Proxy Infrastructure
- Step 2: Plan the Web Application Proxy Server

To deploy Web Application proxy, you can follow the procedures in the following topics:

- Step 1: Configure the Web Application Proxy Infrastructure
- Step 2: Install and Configure the Web Application Proxy Server

See Also

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2

Verifying the AD FS Migration to Windows Server 2012 R2

Once you complete the same server migration of your Active Directory Federation Service (AD FS) farm to Windows Server 2012 R2, you can use the following procedure to verify that federation servers in your farm are operational; that is, that any client on the same network can reach your federation servers.

Membership in **Users**, **Backup Operators**, **Power Users**, **Administrators** or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To verify that a federation server is operational

1. Open a browser window and in the address bar, type the federation server name, and then append it with **federationmetadata/2007-06/federationmetadata.xml** to browse to the federation service metadata endpoint. For example, **https://fs.contoso.com/federationmetadata/2007-06/federationmetadata.xml** .

   If in your browser window you can see the federation server metadata without any SSL

errors or warnings, your federation server is operational.

2. You can also browse to the AD FS sign-in page (your federation service name appended with **adfs/ls/idpinitiatedsignon.htm**, for example, **https://fs.contoso.com/adfs/ls/idpinitiatedsignon.htm**). This displays the AD FS sign-in page where you can sign in with domain administrator credentials.

   Important
   Make sure to configure your browser settings to trust the federation server role by adding your federation service name (for example, **https://fs.contoso.com**) to the browser's local intranet zone.

See Also

Migrating Active Directory Federation Services Role Service to Windows Server 2012 R2

Migrate DHCP Server to Windows Server 2012 R2

DHCP server role migration involves moving the settings for your existing DHCP server to a new DHCP server on the network. The goal of this server migration is to install the DHCP server role on the Windows Server® 2012 R2 operating system so that it provides DHCP leases on a network without any perceptible change to DHCP client computers.

About this guide

This guide describes the steps for migrating existing DHCP server settings to a server that is running Windows Server 2012 R2. Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012 R2. By using the tools that are described in this guide to migrate a DHCP server, you can simplify migration, reduce migration time, increase the accuracy of the migration process, and help eliminate possible conflicts that might otherwise occur during DHCP migration. For more information about the migration tools, see DHCP Server Migration: Appendix A.

Target audience

This guide is intended for information technology (IT) administrators, IT professionals, and other knowledge workers who are responsible for the operation and deployment of DHCP servers in a managed environment.

What this guide does not provide

The following scenarios are not supported or are beyond the scope of this guide.

- Clustering scenarios are not supported by this migration process. For more information about migrating DHCP Server in a cluster environment, see Migrating DHCP to a Cluster Running Windows Server 2008 R2 Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=140512) on the Windows Server TechCenter.

  Also see Migrate to DHCP Failover. DHCP failover is a new option for DHCP high availability, introduced in Windows Server 2012.
- Upgrading roles on the same computer is out of scope for this guide.
- Scenarios in which the new operating system is installed on existing server hardware by using the **Upgrade** option during setup (in-place upgrades) are not covered in this guide.
- Migrating more than one server role is not covered in this guide.

Supported migration scenarios

This guide gives you the instructions to migrate an existing DHCP server to a server that is running Windows Server 2012 R2. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration procedure specific to your server environment based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server 2012 TechCenter (http://technet.microsoft.com/library/jj134039.aspx).

Caution

If the source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

This guide provides instructions only for migrating DHCP data and settings from a server that is being replaced by an x64-based server running Windows Server 2012 R2.

Supported operating systems

This guide provides instructions for migration of a DHCP server from a server that is running Windows Server 2003 or a later operating system to a server running Windows Server 2012 R2. Supported operating systems are listed in the following table.

Supported operating systems for migration

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012 R2, both full and	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
		Server Core installation options	
x86- or x64-based	Windows Server 2008	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2012	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2012	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2012 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2012 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based

The versions of operating systems shown in the previous table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported for the migration of DHCP server settings.

Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers.

Migrations between physical operating systems and virtual operating systems are supported.

Migration from a source server to a destination server that is running an operating system in a different system user interface (UI) language than the source server is not supported. The system UI language is the language of the localized installation package that was used to set up the Windows operating system. For example, you cannot use Windows Server migration tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 R2 in the French system UI language to a computer that is running Windows Server 2012 R2 in the German system UI language.

Both x86-based and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 R2 are x64-based.

Roles that are running on Server Core installations of Windows Server 2008 cannot be migrated, because there is no .NET Framework available on Server Core installations of Windows Server 2008.

We recommend migration rather than an upgrade even when the hardware is native x64-based. For example, with a server role split, a scenario in which the source server has more than one server role, because of increased use of this server you might decide to separate the roles onto several additional x64-based servers. In this case, migrating (not upgrading) individual server roles to other servers may be the best solution.

The server administrator can choose which components of an existing installation to migrate; together with the server role, these components usually include configuration, data, system identity, and operating system settings.

Supported role configurations

You can migrate all DHCP Server settings by using this guide, including registry and database settings.

Notes

If you are migrating a DHCP server in a cluster configuration, see Migrating DHCP to a Cluster Running Windows Server 2008 R2 Step-by-Step Guide (http://go.microsoft.com/fwlink/?LinkId=140512) on the Windows Server TechCenter.

Also see Migrate to DHCP Failover. DHCP failover is a new option for DHCP high availability, introduced in Windows Server 2012.

DHCP Server migration overview

DHCP Server migration is divided into the following major sections:

- DHCP Server Migration: Preparing to Migrate
- DHCP Server Migration: Migrating the DHCP Server Role
- DHCP Server Migration: Verifying the Migration
- DHCP Server Migration: Post-Migration Tasks

DHCP Server migration process

As shown in the following illustration, the pre-migration process involves the manual collection of data, followed by procedures on the destination and source servers. The migration process includes source and destination server procedures that use the **Export** and **Import** cmdlets to automatically collect, store, and then migrate server role settings. Post-migration procedures include verifying that the destination server successfully replaced the source server and then retiring or repurposing the source server. If the verification procedure indicates that the migration failed, troubleshooting begins. If troubleshooting fails, rollback instructions are provided to return to the use of the original source server.

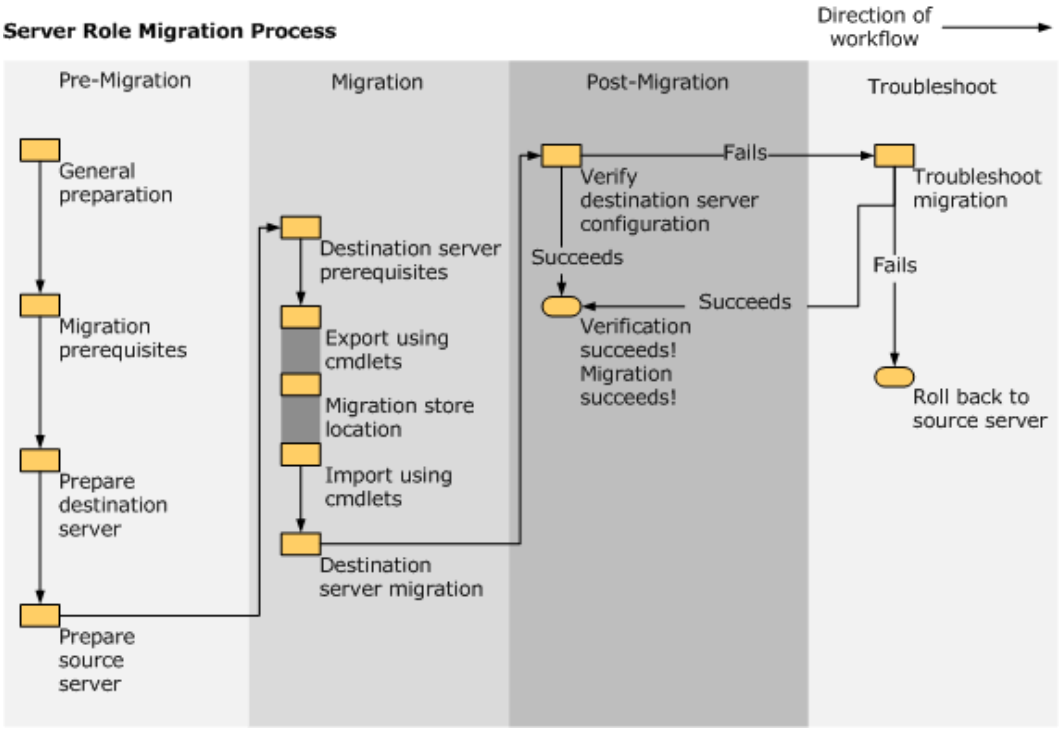

Impact of migration on other computers in the enterprise

During migration, the source DHCP server might not be available. Therefore, client computers will not be able to obtain IP addresses from this DHCP server. We recommend that you maintain or create an auxiliary DHCP server so that client computers can obtain IP addresses while you migrate the primary DHCP server.

Be aware that if you choose to perform the migration without any auxiliary DHCP servers, all clients with valid leases must keep using those leases. If a lease for an existing client expires, that client will not be able to obtain an IP address. In addition, any new client that connects to the network will not be able to obtain an IP address when the single-source DHCP server is not available.

Permissions required to complete migration

The following permissions are required on the source server and the destination server:

- Domain administrative rights that are required to authorize DHCP Server.
- Local administrative rights are required to install or manage DHCP Server.
- Write permissions are required to the migration store location. For more information, see DHCP Server Migration: Preparing to Migrate.

Estimated duration

The migration can take two to three hours, including testing.

See also

DHCP Server Migration: Preparing to Migrate

DHCP Server Migration: Migrating the DHCP Server Role

DHCP Server Migration: Verifying the Migration

DHCP Server Migration: Post-Migration Tasks

DHCP Server Migration: Appendix A

DHCP Server Migration: Preparing to Migrate

Complete the following procedures before you migrate a DHCP Server from an x86-based or x64-based server to an x64-based server running Windows Server 2012 R2.

Migration planning

Install migration tools

Prepare the destination server

Prepare the source server

Migration planning

Membership in **Domain Administrators**, or equivalent, is the minimum required to complete these procedures. Review details about how to use the appropriate accounts and group memberships at Run a program with administrative credentials (http://go.microsoft.com/fwlink/?LinkId=131210).

To prepare for migration

- Identify your DHCP Server source and destination servers.
- Determine the domain, server name, and passwords on the source server. To identify the domain of the original server, click **Start**, right-click **Computer**, and then click

Properties.

- If you have not already done so, install Windows Server Migration Tools on the destination and source servers as instructed in Install migration tools.
- Before migration, install all critical updates and service packs on the source server that were released before Windows Server 2012 R2. It is a recommended best practice that all current critical updates and service packs are installed on both the source and the destination servers.
- Count the number of network adapters in the source and destination servers and make sure that they are equal in number. If the source server that is running DHCP Server has multiple network adapters and the DHCP Server service is bound to all and serving IP addresses on different subnets, the destination server that is running DHCP Server must also have multiple network adapters so that it can serve the same subnets as on the source server.
- Prepare a migration store file location. The store location must be accessible from the source server during the export and from the destination server during the import. Use a common drive that can contain all DHCP Server–related information from the source server. The storage location should be similar to the following: \\fileserver\users\username\.

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who own files or shares that are being migrated.

Install migration tools

Install Windows Server Migration Tools on the destination and source servers. For more information, see Install, Use, and Remove Windows Server Migration Tools (http://technet.microsoft.com/library/jj134202.aspx).

Working with Windows PowerShell cmdlets

Cmdlets (pronounced *command-lets*) are built-in commands, installed by default when you install role services and features in Windows Server 2012 R2. Throughout this guide, there are several PowerShell cmdlets that you will have to run to carry out some of the migration steps. For more information about Windows PowerShell, see Windows PowerShell Support for Windows Server on the Microsoft Web site (http://technet.microsoft.com/en-us/library/hh801904.aspx).

Except where specifically noted, cmdlets are not case-sensitive.

You can obtain detailed Help about specific syntax, parameters, and usage guidelines for any installed cmdlet by typing **Get-Help** <cmdlet name> **-full** in a wps session, in which *cmdlet name* represents the name of the cmdlet for which you want help. Add the **-Verbose** parameter to a cmdlet to display detailed information about the operation in the Windows PowerShell session.

Although most commands for DHCP Server migration are cmdlets, you can run executable files in a session by adding an ampersand (&) before the executable file name. The ampersand is the call operator.

If the executable file is not in the current directory, add the fully qualified path, as shown in the following examples. If an executable file name contains spaces enclose the file name in quotation marks. If you are running the executable file from the current directory, precede the file name with **.**.

- Executable file that is not in the current directory: `PS C:\> & C:\Windows\System32\notepad.exe`
- Executable file that is in the current directory: `PS C:\Windows\System32> & .\notepad.exe`
- Executable file name that contains a space and is in the current directory: `PS C:\Windows\System32> & ".\executable test.exe"`

The commands in this document are provided in Windows PowerShell format. For more information, see DHCP Server Migration: Appendix A. You can run Command Prompt commands in a Windows PowerShell session by adding **cmd /C** before the command, as shown in the following example. The example shows the use of the **dir** command in wps.

```
cmd /C dir c:\*
```

Prepare the destination server

To install DHCP Server on the destination server, complete the menu-driven installation process. Complete the following procedure to prepare the destination server.

To prepare the destination server

1. Install Windows Server 2008 R2 and configure the destination server.
2. Make sure that there is sufficient disk space to store the DHCP Server database. The disk space needed varies with each installation and should be equal to or greater than the space for the DHCP Server database.
3. Add the destination server as a member server in the domain of the source server that is being replaced.
4. Verify that the destination server can resolve the names of domain users who are members of the local group during the import operation. If source and destination servers are in different domains, the destination server must be able to contact a global catalog server for the forest in which the source domain user accounts are located.
5. On a computer that is running Windows Server 2008 R2, open a wps session with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, open the **Windows PowerShell** folder, right-click **Windows PowerShell**, and then click **Run as administrator**.

6. Load the Server Manager module into your wps session. To load the Server Manager module, type the following, and then press **Enter**.

```
Import-Module ServerManager
```

Note

It is not mandatory that DHCP Server is installed on the destination server before you import the settings. If the role is not installed on the destination server, it will be installed automatically during the import process. However, because installation of the role during import might extend downtime, we recommend that you install DHCP Server by using the Server Manager console on the destination server as part of your preparation for the migration.

7. On the destination server, run the following command to install DHCP Server:

```
Add-WindowsFeature DHCP
```

You can also install DHCP Server manually by using Server Manager. For more information, see Install Dynamic Host Configuration Protocol (DHCP) (http://go.microsoft.com/fwlink/?LinkId=128465).

Note

If you use the Add Roles Wizard in Server Manager to install DHCP Server on the destination server, you do not have to answer every question in the wizard. You can leave settings empty (the default) and then click **Next** through each wizard page. If you do not want to use the wizard, you can install DHCP Server by using the **Add-WindowsFeature** cmdlet, as described in this step.

8. By the end of the migration process, the destination server should have a static IP address. Although you will not change the destination server IP address now, consider the following scenarios in preparation for changing it when migration is complete.
   - If your migration scenario requires that you decommission and disconnect the source server from the network, only then can you make the IP address on the destination server the same as the IP address on the source server. The source server must be disconnected from the network or shut down so that there is no IP address conflict between the source server and destination server. However, the destination server can still serve clients that are searching for the legacy (source) server that was running DHCP Server.
   - If your migration scenario calls for continuing to run the source server on the network for other, non-DHCP purposes, you have to assign the destination server an unallocated IP address in the same subnet as the source server to avoid IP conflicts.
   - DHCP Server clients that attempt to renew an IP address lease send the renew request to the previous IP address of the DHCP server. If the source server has been decommissioned and then disconnected from the network and the new DHCP destination server is operating with a different IP address, this request initially fails because of the changed IP address. However, clients try to rediscover the IP address of the DHCP server on the network and therefore recover from this transient failure.

     Warning

If the source server is running multiple roles, renaming the source server or changing its IP address can cause other roles that are running on the source server to fail.

9. If the DHCP Server database path does not match the default path, you must ensure that the destination server has a disk with the same drive letter as seen in source server's DHCP Server database path. For more information, see the "Known issues" section of DHCP Server Migration: Appendix A.

The destination server is now prepared for migration.

Prepare the source server

Follow these steps to prepare the source server for migration.

To prepare the source server

1. Back up the source server. The backup should be a DHCP Server-specific backup, not a Windows backup. (A Windows backup backs up the complete operating system.) You can create the DHCP Server-specific backup by using the **Netsh** command-line tool or Microsoft Management Console (MMC).
   - In the DHCP MMC tree, right-click the server node to open DHCP backup options.
   - Create the backup by using the **Netsh** command-line tool. For more information, see Netsh Commands for Dynamic Host Configuration Protocol server (http://go.microsoft.com/fwlink/?LinkId=128496).

   Note
   The Windows Server 2003 operating system does not support **Netsh**-based backup.
2. If it is running, stop the DHCP Server service. In a session that was opened as described in step 5 of To prepare the destination server, type the following, and then press **Enter**.

   ```
   Stop-Service DHCPserver
   ```
3. If the DHCP Server database path does not match the default path, make sure that the destination server has a disk with the same drive letter as in source server's DHCP Server database path. For more information, see the "Known issues" section of DHCP Server Migration: Appendix A.

You are now ready to begin DHCP Server migration, as described in DHCP Server Migration: Migrating the DHCP Server Role.

See also

Migrate DHCP Server to Windows Server 2012 R2

DHCP Server Migration: Migrating the DHCP Server Role

DHCP Server Migration: Verifying the Migration

DHCP Server Migration: Post-Migration Tasks

DHCP Server Migration: Migrating the DHCP Server Role

Complete the following procedures to migrate a DHCP Server.

Migrating DHCP Server to the destination server

Migrating DHCP Server from the source server

Destination server final migration steps

Migrating DHCP Server to the destination server

Membership in **Domain Administrators** or equivalent is the minimum required to complete these procedures. Review details about how to use the appropriate accounts and group memberships at Run a program with administrative credentials (http://go.microsoft.com/fwlink/?LinkId=131210).

To migrate DHCP Server to the destination server

1. If it is not already installed, install DHCP Server on the destination server, as previously described in the “Prepare the destination server” section in DHCP Server Migration: Preparing to Migrate.
2. If it is running, stop the DHCP Server service by running the following command:

   ```
   Stop-Service DHCPserver
   ```

   If you are unsure whether the service is running, you can check its state by running the following command:

   ```
   Get-Service DHCPServer
   ```

Migrating DHCP Server from the source server

Follow these steps to migrate DHCP Server from the source server.

To migrate DHCP Server from the source server

1. Open a Windows PowerShell session with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, open the **Windows PowerShell** folder, right-click **Windows PowerShell,** and then click **Run as administrator**.
2. Load Windows Server Migration Tools into your session.

   If you opened the current session by using the Windows Server Migration Tools shortcut on the **Start** menu, skip this step, and go to step 3. Only load the Windows Server Migration Tools snap-in in a session that was opened by using some other method, and into which the snap-in has not already been loaded. To load Windows Server Migration

Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

3. Collect data from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The **Export-SmigServerSetting** cmdlet parameters can collect all source DHCP server data in a single file (Svrmig.mig). Or, the **Export-SmigServerSetting** cmdlet can be run multiple times, with each iteration using one or more parameters to collect and store data in multiple Svrmig.mig files. For more information, see DHCP Server Migration: Preparing to Migrate. Before you run this command, review the following:

   - When you run the command in step 4, you are prompted to provide a password to encrypt the migration store data. You must provide this same password to import from the migration store.
   - The **path** parameter can be an empty or nonempty directory. The actual data file in the directory (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Therefore, the user does not have to specify a file name.
   - If the path is not a shared location that the destination server can read from, you must manually copy the migration store to the destination server or a location that the destination server can access.
   - If a migration store location already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must move the Svrmig.mig file from that location and store it elsewhere, rename or first delete the migration store.
   - You can perform both IP and DHCP Server migration at the same time from a Windows PowerShell prompt by using the **Export-SmigServerSetting** cmdlet combined with the **IPConfig** switch, on a single command line.
   - Additional command line parameter information:
     - **-Users** and **-Group** parameters

     The **-Users** parameter must be specified only if the DHCP Administrators group includes local users. Otherwise, you can use the **-Group** parameter and all members of DHCP administrators will be migrated. Administrator group members can include domain users.

     Important
     If the source server is a domain controller, but the destination server is not, Domain Local groups are migrated as local groups, and domain users are migrated as local users.

     - The **-IPConfig** parameter collects IP information when it is used with the **Export-SmigServerSetting** cmdlet on the source server; the **-IPConfig** parameter applies settings when the **Import-SmigServerSetting** cmdlet is used on the destination server.

     If the source DHCP Server has multiple network adapters and the DHCP server service is bound to more than one network adapter and serving IP addresses on different subnets, the destination DHCP Server must also have multiple network

adapters so that it can serve the same subnets as the source DHCP Server. For more information, see Migrate IP Configuration to Windows Server 2012. Because IP configuration details will be used later when importing IP configuration settings to the destination server, it is a best practice to save the IP configuration settings by using the following command:

```
IPConfig /all > IPSettings.txt
```

The **Import-SmigServerSetting** cmdlet requires you to map the source physical address to the destination physical address.

Note

The destination server can be assigned the same static IP address as the source server, unless other roles on the source server must continue to run on it. In that case, the static IP address of the destination server can be any unallocated static IP address in the same subnet as the source server.

4. On the source server, run the **Export-SmigServerSetting** cmdlet, where <*storepath*> is the path that will contain the Svrmig.mig file after this step is completed. An example of the path is \\fileserver\users\username\dhcpstore.

   ```
   Export-SmigServerSetting -featureID DHCP -User All -Group -
   IPConfig -path <storepath> -Verbose
   ```

   For more information about how to export IP configuration settings, see Migrate IP Configuration to Windows Server 2012.

5. On the source server, delete the DHCP authorization for the source DHCP server by running the following command, where *Server FQDN* is the fully qualified domain name (FQDN) of the DHCP server and *Server IPAddress* is the IP address of the server. The command parameters are case-sensitive and must appear exactly as shown.

   ```
   Netsh DHCP delete server <Server FQDN> <Server IPAddress>
   ```

Destination server final migration steps

Return to the destination server and follow these steps to complete the migration.

1. Before you use the **Import-SmigServerSetting** cmdlet to import the DHCP server settings, be aware of the following conditions:
   - You can either use a single command line with all the parameters to import DHCP settings (as when you export data from the source server) or you can use the **Import** cmdlet multiple times to import data one parameter at a time.
   - If you decide to run the **Import-SmigServerSetting** cmdlet separately to import the IP settings, see Migrate IP Configuration to Windows Server 2012. Use the source IPSettings.txt file, referred to in step 3 of the previous procedure. You will map the source physical addresses to the destination physical addresses in step 3 of this procedure.

     Important

If you will be importing role and IP settings separately, you should import IP settings first to avoid any IP conflicts. You can then import the DHCP role.

- If the DHCP Administrators group includes local users, then use the **-Users** parameter combined with the **-Group** parameter to import local users into the DHCP Administrators group. If it only contains domain users, then use only the **-Group** parameter.

  Security

  If the source server is a domain member server, but the destination server is a domain controller, imported local users are elevated to domain users, and imported local groups become Domain Local groups on the destination server.

- If the DHCP Server role that you are migrating has not yet been installed on the destination server, the **Import-SmigServerSetting** cmdlet will install that DHCP Server role and its dependencies, described in the next step. you might have to restart the destination computer to complete the installation after the DHCP Server role is installed by the cmdlet. Then, to complete the import operation after you restart the computer you must run the **Import-SmigServerSetting** cmdlet again along with the **-Force** parameter..

2. On the destination server, run the following command, where *<storepath>* is the available path that contains the Svrmig.mig file, *<SourcePhysicalAddress-1>* and *<SourcePhysicalAddress-2>* are comma-separated lists of the physical addresses of the source network adapter, and *<TargetPhysicalAddress-1>* and *<TargetPhysicalAddress-2>* are comma-separated lists of the physical addresses of the destination network adapter:

   ```
   Import-SmigServerSetting -featureid DHCP -User All -Group -
   IPConfig <All | Global | NIC>
   -SourcePhysicalAddress <SourcePhysicalAddress-
   1>,<SourcePhysicalAddress-2>
   -TargetPhysicalAddress <TargetPhysicalAddress-
   1>,<TargetPhysicalAddress-2>
   -Force -path <storepath> -Verbose
   ```

   The **-IPConfig** switch should be used with the value **All** in case the user wants to import all source settings. For more information, see Migrate IP Configuration to Windows Server 2012.

   Important

   If you import the source server IP address to the target server together with the DHCP role without disconnecting or changing the IP address of the source server, an IP address conflict will occur.

3. Run the following command to start the DHCP service:

   ```
   Start-Service DHCPServer
   ```

4. Authorize the destination server. Command parameters are case-sensitive and must appear exactly as shown. On the destination server, run the following command where *Server FQDN* is the FQDN of the DHCP Server and *Server IPAddress* is the IP address of the server:

   ```
   netsh DHCP add server <Server FQDN> <Server IPAddress>
   ```

Note

After authorization, the Server Manager event log might show event ID 1046. This is a known issue and is expected to occur only once. The event can be safely ignored.

When this migration is finished, client computers on the network server are served by the new x64-based destination server running Windows Server 2012 R2. The migration is complete when the destination server is ready to serve IP addresses to the network.

See also

Migrate DHCP Server to Windows Server 2012 R2

DHCP Server Migration: Preparing to Migrate

DHCP Server Migration: Verifying the Migration

DHCP Server Migration: Post-Migration Tasks

DHCP Server Migration: Appendix A

DHCP Server Migration: Verifying the Migration

When all the migration steps are complete, you can use the following procedure to verify that the DHCP server role migration was successful. If the migration failed, you can return to the previous valid configuration by following the steps in DHCP Server Migration: Post-Migration Tasks.

Verifying destination server configuration

Follow these steps to confirm that the DHCP destination server is now serving the domain.

Membership in **Domain Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about how to use the appropriate accounts and group memberships at Understanding Groups: Default groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To verify the configuration of the destination server

1. Make sure that the destination server is authorized by running the following command in a Windows PowerShell window:

   ```
   netsh DHCP show server
   ```

   The output of this command must contain the name of the DHCP destination server.
2. Check whether DHCP server is running on the destination server. In Task Manager, on the **Services** tab, the DHCP server status should be **Started**. You also use Task Manager to confirm that the status of the source server is **Stopped**.
3. Verify that the client computers are correctly receiving IP addresses on request by

running the following commands at a command prompt on a client computer:

```
ipconfig /release
ipconfig /renew
```

If the IP address of the DHCP server has not changed, you do not have to run the **ipconfig /release** command. Running **ipconfig /renew** should be sufficient.

The output of these commands should show that the client computer was issued an IP address.

4. Use the DHCP console to verify that the scopes and other settings were migrated. To connect to the destination server, click **Action**, click **Add Server**, and then type the IP address or host name of the DHCP server. In the console tree, expand the server node, and then expand the IPv4 and IPv6 nodes to confirm that the scopes have been migrated. Then locate the folders for the scopes and view the address range, reservations, scope options, and active leases to verify the same. You can also go to the Server Options folder and verify the migrated server options.

See also

Migrate DHCP Server to Windows Server 2012 R2

DHCP Server Migration: Preparing to Migrate

DHCP Server Migration: Migrating the DHCP Server Role

DHCP Server Migration: Post-Migration Tasks

DHCP Server Migration: Appendix A

DHCP Server Migration: Post-Migration Tasks

The post-migration tasks for the source server are optional, depending on your migration scenario.

Completing migration

Restoring DHCP in the event of migration failure

Troubleshooting cmdlet-based migration

Completing migration

Migration is complete after you have verified that the destination server, not the source server, is now serving the network. If your verification efforts demonstrate that the migration failed, see “Restoring DHCP in the event of migration failure” later in this topic.

Retiring DHCP on your source server

After you have verified the migration, you can disconnect, repurpose, or retire the source server. If the source server is running other server roles, it should be left on the network. If you do not have to use this computer, you can store it as a backup in case you ever have to revert to your previous DHCP configuration.

- If your migration scenario includes a standalone DHCP Server, then this source server was taken offline after the export file was created, as described in DHCP Server Migration: Preparing to Migrate. In this scenario, the DHCP service was interrupted from the time that it was stopped until the migration was complete on the new server, as described in DHCP Server Migration: Migrating the DHCP Server Role.
- If your migration scenario includes more than one DHCP Server in a domain, a backup or other DHCP server continues to serve IP addresses during the migration so that services to clients are never interrupted. The migration is complete on the new server when the IP address of the source server is migrated to the destination server.

Retiring your source server

After you have confirmed that the destination server is performing the functions previously handled by the source server, you can retire or repurpose the source server, depending on your needs. Follow your organization's policy regarding server decommissioning. For information about decommissioning a domain controller, see Decommissioning a Domain Controller (http://go.microsoft.com/fwlink/?LinkID=128290).

Warning
After the source server is repurposed as a member server, otherwise repurposed or retired from service, you cannot roll that server back to its previous working state.

Restoring DHCP in the event of migration failure

If the migration of DHCP Server fails, you have these options:

- If the source server has not been repurposed, an administrator can reassign the IP configuration settings, reauthorize the server, and restart the DHCP service on the original server.
- Use the backup files that were created on the source server, as described in DHCP Server Migration: Preparing to Migrate, to restore DHCP server on the original DHCP server.

Estimated time to complete a rollback

You should be able to complete a rollback in one to two hours.

Troubleshooting cmdlet-based migration

The Windows Server Migration Tools deployment log file is located at %*windir*%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations.

- %*windir*%\Logs\ServerMigration.log
- On Windows Server 2008 and Windows Server 2008 R2: %*localappdata*%\SvrMig\Log
- On Windows Server 2003: %*userprofile*%\Local Settings\Application Data\SvrMig\Log

If migration log files cannot be created in the previous locations, **ServerMigration.log** and **SmigDeploy.log** are created in %*temp*%, and other logs are created in %*windir*%\System32.

For DHCP-specific troubleshooting tips, see Troubleshooting DHCP servers on the Windows Server TechCenter (http://go.microsoft.com/fwlink/?LinkId=128533). Although these tips are written for Windows Server 2003, they also address common issues that apply to later versions of the operating system.

If a migration cmdlet fails, and the wps session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at 0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050). Minidump attached (317793 bytes).

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Viewing the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects, and query them for more information about settings and data that were migrated. You can also use result objects as input for other wps commands and scripts.

Result object descriptions

The Windows Server Migration Tools **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or share that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.
SourceLocation	String	The source location of the item,

Property name	Type	Definition
		shown as a path.
DestinationLocation	String	The destination location of the item, shown as a path.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to both **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.
Title	String	Title of the result object. Examples of values include the network adapter physical

Property name	Type	Definition
		address for IP configuration, or user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   ```
   $VariableName = $(Cmdlet)
   ```

   The following is an example.

   ```
   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -
   User all -Group -Path D:\rmt\DemoStore -force -Verbose)
   ```

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.

2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information that is contained in the result object by typing a command in the following format, and then pressing **Enter**.

   ```
   $VariableName
   ```

   In the following example, the variable is named **ImportResult**.

   ```
   $ImportResult
   ```

   This command returns information that was contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable.

   ```
               ItemType  ID                                Success
   DetailsList
               --------  --                                -------
   -----------
             OSSetting  Local User                          True
   {Local User, Loc...
             OSSetting  Local Group                         True
   {Local Group, Lo...
        WindowsFeature  DHCP                                True
   {}
   ```

Each line of the previous sample is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return more detail about result objects, as shown by examples in step 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   ```
   $<VariableName>| Select-Object -ExpandProperty <PropertyName>
   ```

   The following is an example.

   ```
   $importResult | Select-Object -ExpandProperty DetailsList
   ```

4. You can run more advanced queries to analyze result objects by using wps cmdlets. The following are examples.

   - The following command returns only those details of result objects that use the ID **Local User**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local User" } |
     Select-Object -ExpandProperty DetailsList
     ```

   - The following command returns only those details of result objects that use an ID of **Local User** that have a message severity equal to **Warning**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local User" } |
     Select-Object -ExpandProperty DetailsList | ForEach-Object {
     if ($_.Severity -eq "Warning") {$_} }
     ```

   - The following command returns only the details of result objects that use an ID of **Local Group** that also have the title **Remote Desktop Users**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local Group" } |
     Select-Object -ExpandProperty DetailsList | ForEach-Object {
     if ($_.Title -eq "Remote DesktopUsers") {$_} }
     ```

More information about querying results

For more information about the cmdlets that are used in the previous examples, see the following additional resources.

- Where-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134853).
- Select-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134858).
- ForEach-Object on the Microsoft Script Center Web site (http://www.microsoft.com/technet/scriptcenter/topics/msh/cmdlets/foreach-object.mspx)

See also

Migrate DHCP Server to Windows Server 2012 R2

DHCP Server Migration: Preparing to Migrate
DHCP Server Migration: Migrating the DHCP Server Role
DHCP Server Migration: Verifying the Migration
DHCP Server Migration: Appendix A

DHCP Server Migration: Appendix A

Migration tools

Migration tools are provided in Windows Server® 2012 R2. The tools for earlier versions of the Windows operating system are also available in Windows Server 2012 R2.

Follow these steps to access the tools on the destination server:

1. Open Server Manager.
2. Click **Action**, and then select **Add Features**. The Add Features Wizard opens.
3. On the **Select Features** page, from the **Features** list, select **Windows Server Migration Tools**, and then click **Next**.
4. Complete the steps in the wizard, and then click **Close**.

The previous steps do not work for Server Core installations. To install the migration tools on a Server Core installation, see the Install Windows Server Migration Tools topic in Install, Use, and Remove Windows Server Migration Tools.

Installing and using Windows PowerShell with migration cmdlets

To access, download, and install migration tools (the migration toolkit), any role-specific tools, and Windows PowerShell, see Install, Use, and Remove Windows Server Migration Tools.

Known issues

If the DHCP installation on the source server has a database path that varies from the default, before you perform the import, provide the destination server with a volume with the same drive letter on which the DHCP server database exists on the source server. For example, if the DHCP server database on the source server is located on D:\, then the destination server should have a volume with the driver letter D.

If you cannot match the volume on the destination server that has the same driver letter as that shown for the source DHCP server database, then the DHCP database path on the source server must be changed back to the default path (%systemroot%\system32\dhcp) before you start the migration.

See also

Migrate DHCP Server to Windows Server 2012 R2

DHCP Server Migration: Preparing to Migrate

DHCP Server Migration: Migrating the DHCP Server Role

DHCP Server Migration: Verifying the Migration

DHCP Server Migration: Post-Migration Tasks

Migrate Hyper-V to Windows Server 2012 R2 from Windows Server 2012

With Hyper-V, you can create a virtualized server computing environment by using a technology that is part of Windows. This guide provides information and instructions about migrating the Hyper-V role that include virtual machines, data, and operating system settings from the source server running Hyper-V in Windows Server 2012 to the destination server that is running the Windows Server 2012 R2 operating system.

About this guide

This guide describes how to migrate the Hyper-V role by providing preparation, migration, and verification steps.

Migration documentation and tools facilitate the migration of server role settings and data from an existing source server to a destination server that is running Windows Server 2012 R2. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help eliminate possible conflicts that might otherwise occur during the migration process.

In addition to the migration options that are described in this topic, Virtual Machine Manager in Microsoft System Center 2012 R2 can facilitate and automate a considerable part of the migration process. For more information about Virtual Machine Manager, see Virtual Machine Manager.

Target audience

This document is intended for information technology (IT) professionals who are responsible for operating and deploying Hyper-V in a managed environment.

What this guide does not provide

- Migration of Hyper-V from one server that runs Windows Server 2008 R2 to another server that runs Windows Server 2012 R2.
- Instructions for migrating more than one server role at one time.

- Migration of Hyper-V from one server that runs Windows Server 2012 R2 to another server that runs Windows Server 2012 R2. Instead, this process is supported by Hyper-V management tools and features. The general process is as follows:
    a. Determine whether to use export and import or live migration to move the virtual machines:
        - Export and import can be used in either a workgroup or a domain environment. In Hyper-V running on Windows Server 2012 R2, you can now export a running virtual machine.
        - Live migration requires a domain environment and some additional configuration, but the virtual machine is running throughout the move process.
    b. Add the Hyper-V role to the destination server. You can configure the default storage locations and live migration when you add the role.
    c. Configure virtual switches, and optionally, other networking features on the destination server. Management tools include the Windows PowerShell cmdlets New-VMSwitch and Set-VMSwitch, and the Hyper-V Virtual Switch Manager in the Hyper-V Manager Console.
    d. Move the virtual machines by using export and import or live migration. Management tools include the Windows PowerShell cmdlets Export-VM, Import-VM, and Move-VM and the **Export**, **Import**, and **Move** menu commands in the Hyper-V Manager Console.

Supported migration scenarios

This guide provides you with instructions to migrate a server that is running Hyper-V in Windows Server 2012 to a server that is running Windows Server 2012 R2. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration procedure that is specific to your server environment and is based on the information in other role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Migration dependencies

The Hyper-V role does not depend on any other roles. As a best practice, we recommend that no other roles are installed on a server running Hyper-V.

Migration scenarios that are not supported

The following migration scenarios are not supported:

- Virtual machine configuration under one of the following conditions:
    - When the number of virtual processors that are configured for the virtual machine is greater than the number of logical processors on the destination server.
    - When the memory that is configured for a virtual machine is greater than the available memory on the destination server.

Overview of migration process for this role

Hyper-V role migration involves moving the virtual machines, virtual networks, and all the associated settings from one physical computer to another physical computer in the enterprise. The process supports moving from a server running Hyper-V in Windows Server 2012 to a server running Hyper-V in Windows Server 2012 R2. The Hyper-V role does not depend on any other roles.

The migration tools include Windows PowerShell cmdlets that you can use to perform some of the tasks that are required to migrate the Hyper-V role and script or to automate the migration process.

In previous versions of Hyper-V, you were required to shut down a virtual machine before you moved it to a new server. If the move was performed correctly, downtime was limited, but still, there was downtime. A new feature in Windows Server 2012 R2, cross-version live migration, supports moving a running virtual machine from Windows Server 2012 to Windows Server 2012 R2. The Windows PowerShell Export-VM cmdlet captures the majority of the Hyper-V settings that are required to perform a successful migration, which includes the virtual machine configurations, virtual networks, and virtual hard disks. Now you can decide how to move virtual machines to Windows Server 2012 R2, where in the past, your options where limited.

The following options are available to move a virtual machine:

- In-place upgrade
- Cross-version live migration
- Export and Import
- Copy Cluster Role Wizard

For additional information about each option, see Hyper-V: Migration Options.

This guide explains the migration process for the following three main scenarios:

- Hyper-V: Stand-alone Migration
- Hyper-V Cluster Using Separate Scale-Out File Server Migration
- Hyper-V Cluster Using Cluster Shared Volumes (CSV) Migration

Estimated duration

The length of time it takes to migrate the Hyper-V role depends on the size of the data to be transferred and on the tools that are used. Of the various types of files to be transferred, the virtual hard disk (VHD), .vhd and .vhdx files, have the largest file sizes from a few gigabytes to many gigabytes in size. The length of time that is required for migration is largely affected by the size of the VHD files and by the network bandwidth.

Additional references

Windows Server Migration forum

Hyper-V: Migration Options

When you migrate virtual machines from the Windows Server 2012 operating system to the Windows Server 2012 R2 operating system, you have various options on how to migrate your virtual machines. You now can select the migration options that meet the requirements of your environment.

Hyper-V migration options

Depending on your requirements and service level agreements that must be maintained, you can use one migration option or a combination of migration options. For example, if you have virtual machines that either must be running all the time or that only have a short maintenance period for you to shut them down, you might select to use cross-version live migration to move the virtual machines from Windows Server 2012 to Windows Server 2012 R2. For other virtual machines that are not as critical, have a long maintenance period, or might take too long to move by using live migration, you can use the Copy Cluster Roles Wizard or **Export / Import**, which depends on your environment.

The following table shows the benefits and disadvantages of the various migration options.

Migration option	Benefits	Disadvantages
In-place upgrade	• No new hardware required.	• Virtual machines must be shut off during the upgrade.
Cross-version live migration	• Virtual machines continue running during migration. • If the virtual hard disk is stored on a Scale-out File Server share that is accessible from both servers, the virtual hard disks do not have to be copied. • Moves virtual machines from one Hyper-V cluster to another cluster without any downtime. • Migrates individual virtual machines that are part of the Hyper-V cluster.	• Requires additional hardware or extra capacity in the existing cluster to create the destination cluster. • The amount of time it takes to migrate a virtual machine depends on various factors, for example, the size of memory that is configured for the virtual machine and the network configuration. If the virtual hard disks are not stored on a Scale-out File Server, additional time is required to move the virtual hard disk. • The virtual machine must be removed from the existing cluster before it is

Migration option	Benefits	Disadvantages
		moved to the new cluster. After the virtual machine has successfully moved to new cluster node, high availability is added to the virtual machine. During the move process, the virtual machine is not protected by the cluster services.
Copy Cluster Roles Wizard	• Easily migrates a Hyper-V cluster from Windows Server 2012 to Windows Server 2012 R2. • Tests the Copy Cluster Roles process without affecting production. • Reverses the process if you encounter any issues. • Copies roles on test clusters to production clusters.	• All virtual machines on the same Shared Clustered Volume are migrated at the same time. • Virtual machines must be shut down for a short period of time. • The Copy Cluster Wizard does not copy the Hyper-V replication settings when it copies a virtual machine to a new failover cluster. Hyper-V replication must be re-enabled on the virtual machine after it is copied. For the **Initial Replication Method**, select **Use an existing virtual machine on the Replica server as the initial copy**. • Requires additional hardware.
Export / Import	• Migrates one virtual machine at a time. • Controls the method of copying the virtual machine to the new server.	• The virtual machine is shut down during the Export / Import process. • Requires additional hardware. • The Import of a virtual machine removes any Hyper-V Replica configuration settings for a virtual machine. Hyper-V

Migration option	Benefits	Disadvantages
		replication must be re-enabled on the virtual machine after it is imported. For the **Initial Replication Method**, select **Use an existing virtual machine on the Replica server as the initial copy**.

The following table shows the available options to use in different deployments of Hyper-V.

Scenario / Migration option	In-place upgrade	Export / Import	Cross-version live migration	Copy Cluster Roles Wizard
Standalone host	Yes	Yes	Yes	No
Hyper-V Cluster with Cluster Shared Volumes (CSV)	No	Yes	Yes, the virtual machine must be removed from the cluster first, and the virtual hard disks must be copied as part of the live migration.	Yes
Hyper-V Cluster with a separate Scale-out File Server for storage	No	Yes	Yes, the virtual machine must be removed from the cluster first.	Yes

Important
When Hyper-V Replica is enabled, we recommend that you migrate the virtual machines on the Replica site first, and then migrate the primary site.

Cross-version live migration

The upgrade to a new version of Windows Server no longer requires downtime of the virtual machines. In Windows Server 2012 R2, Hyper-V live migration has been updated to support the migration of virtual machines in Windows Server 2012 to Windows Server 2012 R2. If the virtual hard disk files are stored on an SMB 3.0 share that is accessible from both the source and destination server, you only must move the virtual machine configuration and memory files, but not the virtual hard disk files. If the virtual hard disk files are not stored on an SMB 3.0 share, or if the share is not accessible to the destination server, you can use the Shared Nothing Live

Migration to migrate the virtual hard disk files, virtual machine configuration files, and the running virtual machine with no downtime.

Hyper-V Replica

Hyper-V Replica was introduced in Windows Server 2012 and provides asynchronous replication of Hyper-V virtual machines between two hosting servers. It is simple to configure and does not require either shared storage or any particular storage hardware. Any server workload that can be virtualized in Hyper-V can be replicated. Replication works over any ordinary IP-based network, and the replicated data can be encrypted during transmission. Hyper-V Replica works with standalone servers, failover clusters, or a mixture of both. The servers can be physically co-located or widely separated geographically. The physical servers do not have to be in the same domain or even be joined to any domain at all.

Consider the following factors for the move from Windows Server 2012 to Windows Server 2012 R2 when you use Hyper-V Replica:

- You must upgrade the Replica server first. A Windows Server 2012 R2 Replica server can accept replication from a primary server that runs Windows Server 2012. However, a Windows Server 2012 Replica server cannot accept replication from a primary server that runs Windows Server 2012 R2.
- When you upgrade the Replica server, consider the following factors:
  - When you perform an in-place upgrade on the Replica server, the post-upgrade of the Replica server to Windows Server 2012 R2 replication continues from the primary server that runs Windows Server 2012 at the default replication frequency of 5 minutes.
  - When you move the virtual machines to a new server that runs Windows Server 2012 R2, you must update the virtual machine replication settings on the primary server with the name of the new Replica server or Hyper-V Replica Broker. Until the Replica server name is updated, replication does not resume.
  - You can start to use new Hyper-V Replica features, such as extended replication from the Replica server.
  - You can add new virtual machines to the primary server that runs Windows Server 2012 and start replication to a Replica server that runs Windows Server 2012 R2.
  - In case of emergency, you can fail over your virtual machines from the primary server to the Replica server. You cannot start reverse replication because replication is not supported from Windows Server 2012 R2 to Windows Server 2012.

    Note

    At this point, the virtual machine is no longer protected by Hyper-V Replica. You can configure extended replication by using another server running Hyper-V in Windows Server 2012 R2. After the primary server has been upgraded to Windows Server 2012 R2, you can reverse replication back to the primary server. When you reverse replication, you can select to use an existing virtual machine to limit the amount of replication that must be transmitted over the network.

    - Migration cancels a test failover that currently runs for a Replica virtual machine and deletes the test virtual machine.
- When you upgrade the primary server, consider the following factors:
    - The Replica server has already been upgraded to Windows Server 2012 R2. If the Replica server has not been upgraded to Windows Server 2012 R2, replication fails until the Replica server is upgraded to Windows Server 2012 R2.
    - Replication continues at the default frequency of 5 minutes, which can be modified if it is required.
    - When you use certificate-based authentication for Hyper-V Replica, after you move the primary virtual machine to a new server, you must update the certificate thumbprint for the virtual machine.

      You can update the certificate thumbprint in the Hyper-V Manager Console by editing the Replication settings of the virtual machine, or you can use the following Windows PowerShell cmdlet, **Set-VMReplication**.

      ```
      Set-VMReplication –VMName <virtual machine name. –
      CertificateThumbprint <thumbprint>
      ```

See also

Hyper-V Replica Overview

Virtual Machine Live Migration Overview

Deploy Hyper-V over SMB

Failover Clustering Overview

Migrating Clustered Services and Applications to Windows Server 2012

Hyper-V: Stand-alone Migration

This scenario describes how to migrate a single server running the Hyper-V role in Windows Server 2012 to Windows Server 2012 R2.

Migration options

When you migrate a single server, you have the following migration options available:

- In-place upgrade
- **Cross-version live migration**
- Export and Import (not covered in this guide)

In-place upgrade

This scenario describes how to use the existing hardware that runs the Windows Server 2012 operating system and to perform an in-place upgrade of the operating system to Windows Server

2012 R2. This scenario does not require any additional hardware; however, during the upgrade process, all of the virtual machines must be turned off or be in a saved state.

Notes

- We recommend that you shut down or turn off all of the virtual machines before you upgrade. Virtual machines can be in a saved state during the upgrade, but we do not recommend it. You receive a warning during the upgrade process if any of the virtual machines are in a saved state.
- Before you run an in-place upgrade, we recommend that you back up the management operating system and the virtual machines.

Perform an in-place upgrade

Use the following steps to perform an in-place upgrade.

Note

If Hyper-V Replica has been enabled on any of the virtual machines, we recommend that you upgrade the Replica server first. During the upgrade of the Replica server, the primary server continues to send updates to the Replica server, and you might see warning messages about the health of the replication. After the Replica server has successfully upgraded, the replication should continue normally.

- Log on to the server by using a user account with local Administrator rights.
- Insert media for Windows Server 2012 R2 and run **Setup.exe** if the installation program did not start automatically.
- Review the upgrade report and fix any blocking warning messages.
- After the server running Hyper-V has restarted, confirm that the server running Hyper-V was successfully upgraded.
- Install the latest updates.
- Start each of the virtual machines that were running before the upgrade.
- Confirm that each virtual machine operates as expected.
- Upgrade the integration services for each virtual machine. A restart might be required to complete the integration services update.

Cross-version live migration

The upgrade to a new version of the Windows Server operating system no longer requires downtime of the virtual machines. In Windows Server 2012 R2, live migration has been updated to support the migration of Hyper-V virtual machines in Windows Server 2012 to Hyper-V in Windows Server 2012 R2. If the virtual hard disk (VHD) files are stored on an SMB 3.0 file share, you must only move the virtual machine, but not the storage.

This scenario requires additional hardware for a destination server. Ensure that the destination server has the capacity to run the virtual machines that you are currently running and has room for future expansion.

Use the following steps to move a virtual machine from Windows Server 2012 to Windows Server 2012 R2.

Prepare the new server hardware

1. Install Windows Server 2012 R2 on the new server hardware.
2. Install the Hyper-V role on the server.
3. Configure the following Hyper-V settings, for example:
   - The default location for virtual hard disks and virtual machine configuration files.
   - NUMA settings.
   - Live migration settings. Even if live migration was not previously configured, you must enable and configure live migration on both servers.
   - Replication settings if Hyper-V Replica is used. If certificate-based authentication is configured, an appropriate certificate must be installed on the new server.
   - Virtual switches.
   - **Hyper-V Administrators** local group membership.
4. Install the latest updates.

Move a virtual machine from Hyper-V in Windows Server 2012 to Windows Server 2012 R2

In this section, you move a virtual machine from Hyper-V in Windows Server 2012 to Hyper-V in Windows Server 2012 R2.

Perform this procedure on the source server running Hyper-V in Windows Server 2012.

To move the virtual machine to Hyper-V in Windows Server 2012 R2

1. On the **source server** running Hyper-V in Windows Server 2012, open the Hyper-V Manager Console, and then select the virtual machine that you want to move.
2. From the **Actions** pane, click **Move**. This action opens the **Move Wizard**.
   a. On the **Choose Move Type** page, select **Move the virtual machine**.
   b. On the **Specify Destination Computer**, specify the name or server that is running Windows Server 2012 R2.
   c. On the **Choose Move Options** page, select **Move only the virtual machine**.

   You can also use the Windows PowerShell cmdlet **Move-VM**. The following example shows a virtual machine *test VM* that is moved to a remote computer *NewServer* where the virtual machine is stored on an SMB share.

   ```
   PS C:\> Move-VM –Name "Test VM" -DestinationHost NewServer
   ```

Modify the Hyper-V Replica settings

Note

Perform the following procedure on the primary server after moving a virtual machine on the Replica server.

[Optional] To modify Hyper-V Replica settings

1. On the primary server, open the Hyper-V Manager Console, and then select the virtual machine whose Replica virtual machine was just moved.
2. Right-click the virtual machine to select **Settings**.
3. Select **Replication** and update the value for **Replica server** with the name of the destination Replica server.
4. Confirm that replication has successfully started.

   You can also use the **Set-VMReplication** cmdlet to update the name of the Replica server.

Verify that the virtual machine runs correctly

This procedure describes how to confirm that the virtual machine that was moved runs correctly on the destination server running Hyper-V in Windows Server 2012 R2.

Note

Skip this step when you move a virtual machine on a Replica server. Replica virtual machines are in an off state until the virtual machine is failed over by the administrator.

To verify that virtual machine runs correctly

1. Open the Hyper-V Manager Console on the destination server.
2. Verify that the virtual machine is running. If the virtual machine is not running, attempt to start it. If the virtual machine does not start, check the event log to see why it failed to start.
3. [Optional] - Run some basic operations that change the state of the virtual machine.
4. Run the necessary application-specific tests to ensure that the application on the virtual machine can provide the same service levels as it provided before the virtual machine was migrated. Although the virtual machine was moved while it was running the services that the virtual machine provides, the services should not have been interrupted.
5. Verify that you can connect to the virtual machine by using **Remote Desktop** or **Virtual Machine Connection**.
6. Upgrade the integration services on the virtual machine. Because the virtual machine was never shut down during the migration, you can update the integration services silently without a restart. The update occurs the next time that the virtual machine is restarted during its scheduled maintenance period.
   a. Modify the settings of the virtual machine and specify the following media to be used for the CD/DVD drive, *%Systemroot%\System32\Vmguest.iso*.
   b. Run the following command from an elevated command prompt in the virtual machine:

i. For 64-bit Windows Server operating systems, *drive:\Support\Amd64\Setup.exe /quiet /norestart*

ii. For 32-bit Windows Server operating systems, *drive:\Support\X86\Setup.exe /quiet /norestart*

See also

Virtual Machine Live Migration Overview

Configure Live Migration and Migrating Virtual Machines without Failover Clustering

Install and Deploy Windows Server 2012 R2 and Windows Server 2012

Hyper-V Replica Overview

Hyper-V: Hyper-V Cluster Migration

Hyper-V Cluster Migrations

The following sections describe how to migrate a Hyper-V cluster running in Windows Server 2012 to a Hyper-V cluster running in Windows Server 2012 R2. Depending on the configuration of the storage that the cluster uses, the following migration options are available:

- Hyper-V Cluster Using Separate Scale-Out File Server Migration
- Hyper-V Cluster Using Cluster Shared Volumes (CSV) Migration

Hyper-V Cluster Using Separate Scale-Out File Server Migration

This scenario describes how to migrate virtual machines from a Hyper-V cluster by using a separate Scale-out File Server that runs on the Windows Server 2012 operating system to the Windows Server 2012 R2 operating system. In this scenario, you move the virtual machines that run on a Hyper-V cluster from Windows Server 2012 to a Hyper-V cluster that runs on Windows Server 2012 R2.

Depending on your requirements, you have two main options to move your virtual machines from a Hyper-V cluster that runs on Windows Server 2012 to a Hyper-V cluster that runs on Windows Server 2012 R2. For information about the advantages or disadvantages for each option, see Hyper-V: Migration Options.

- Cross-version live migration
- Copy Cluster Roles Wizard

Note
Because the Hyper-V cluster and Scale-Out File Server run on separate clusters, you can upgrade each cluster independently of the other. This topic only describes how to move the virtual machines to a new Hyper-V cluster, while the storage remains on the same Scale-out File Server.

Cross-version live migration

With cross-version live migration, you can move a running virtual machine from a server running Hyper-V in Windows Server 2012 to a server running Hyper-V in Windows Server 2012 R2. Cross-version live migration does not support moving a virtual machine to a down-level version of Hyper-V.

Important
To use cross-version live migration, the virtual machine must be removed from the cluster. The virtual machine is then moved to one of the servers in the new cluster. After the virtual machine has successfully moved to the new server, the virtual machine is configured for high availability on the new cluster. During this process, the virtual machine is not highly available.

In this option, you must create a new Windows Server 2012 R2 Hyper-V cluster. You have various options to create the new Hyper-V cluster:

- Evict two nodes from the existing cluster and create a new two-node cluster.
- Evict one node from the existing cluster and use new hardware to create a new two-node cluster.
- Use two new servers to create the new cluster.
- Evict one node from the existing cluster and create the new one-node cluster. Until a second node is added, the virtual machines that are moved to the new cluster are not highly available.
- Use one new server to create a new cluster with one node. Until a second node is added, the virtual machines that are moved to the new cluster are not highly available.

Now that the Windows Server 2012 R2 Hyper-V cluster is running, you can move the virtual machines that are currently running on one of the nodes to the new cluster.

Important
The folder that Hyper-V uses to store virtual machine data requires specific permissions to access the Server Message Block (SMB) file share. You must ensure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions. For more information about deploying Hyper-V over SMB, see Deploy Hyper-V over SMB.

Cross-version live migration scenario

The following migration scenario is based on the following factors:

The following table lists the servers at the start of the migration.

For the old cluster, there are three servers running Hyper-V with eight highly available virtual machines.

A two-node Windows Server 2012 R2 cluster has been prepared and is ready to receive the virtual machines from the server running Hyper-V in Windows Server 2012.

Note
You must enable and configure Hyper-V live migration on all of the servers running Hyper-V.

Name	Windows Server operating system	Cluster name
HVSRV1	Windows Server 2012	HVHA2012
HVSRV2	Windows Server 2012	HVHA2012
HVSRV3	Windows Server 2012	HVHA2012
HVR2_1	Windows Server 2012 R2	HVHAR2
HVR2_2	Windows Server 2012 R2	HVHAR2

The following table lists the virtual machines that are currently running on each of the nodes at the start of the migration.

Virtual machine name	Server running Hyper-V
Testvm_1	HVSRV1
Testvm_2	HVSRV1
Testvm_3	HVSRV2
Testvm_4	HVSRV2
Testvm_5	HVSRV3
Testvm_6	HVSRV3

The following are the general steps to move the virtual machines from the HVHA2012 cluster to the new HVHAR2 cluster.

1. Create a new Hyper-V cluster by using a separate Scale-Out File Server.
2. Move all of the virtual machines from HVSRV1 to HVR2_1.
    a. On the HVHA2012 cluster, remove one virtual machine that runs on HVSRV1. The virtual machine still runs on Hyper-V, but it is no longer highly available.
    b. From the Hyper-V Manager, move the virtual machine that was removed in step 2a to the HVR2_1 server. Because there is shared storage, you must move only the virtual machine, not the storage.
    c. On the HVHAR2 cluster, add the virtual machine that was moved to HVR2_1 in step 2b to the new cluster. The virtual machine is now highly available.
    d. Repeat steps 2a-2c until all of the virtual machines from HVSRV1 node have been moved to the new R2Cluster cluster.
3. Evict HVSRV1 from the HVHA2012 cluster.
4. Install Windows Server 2012 R2 on HVSRV1, and then join the server to the HVHAR2 cluster.
5. Repeat steps 2-4 for HVSRV2 and HVSRV3.

Note

When you get down to the last two servers in the Windows Server 2012 cluster, if you evict another node, you have a single-node cluster, and the remaining virtual machines are no longer highly available. If there is enough capacity on the new cluster to run the remaining virtual machines, move all of the remaining virtual machines to the new cluster, and then evict the last two servers at the same time.

To create a Windows Server 2012 R2 Hyper-V cluster

1. Create a new Hyper-V cluster that is connected to the same Scale-out File Server, to which Windows Server 2012 is connected.

    Configure live migration on the new servers running Hyper-V and the old servers running Hyper-V.

    If Hyper-V replication is enabled, configure the Hyper-V Replica Broker on the new cluster, HVHAR2.

    For more information about creating a Hyper-V cluster, see Deploy a Hyper-V Cluster.

    Important

    The folder that Hyper-V uses to store virtual machine data requires specific permissions to access the SMB file share. You must ensure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions. For more information about deploying Hyper-V over SMB, see Deploy Hyper-V over SMB.

To move the virtual machines to the new cluster

1. On HVSRV1, open the Failover Cluster Manager and select **Nodes**.

2. Right-click the HVSRV1 node, and then select **Pause** and **Do Not Drain Roles**.
3. In the information pane for the HVSRV1, select **Roles** to see the virtual machines on the node.
4. Right-click the **Testvm_1** virtual machine, and then select **Remove**. The virtual machine is still running but is no longer highly available.
5. On HVSRV1, open the Hyper-V Manager.
6. Right-click the **Testvm_1** virtual machine, and then select **Move**.
   a. On the **Choose Move Type** page, select **Move the virtual machine**.
   b. On the **Specify Destination Computer** tab, specify the name or server running Windows Server 2012 R2, HVR2_1. Do not enter the name of the new cluster, HVHAR2.
   c. On the **Choose Move Options** page, select **Move only the virtual machine**.

   You can also use the Windows PowerShell cmdlet **Move-VM**. In the following example, a virtual machine *test VM* was moved to a remote computer *NewServer* where the virtual machine is stored on an SMB share.

   ```
   PS C:\> Move-VM –Name "Test VM" –DestinationServer NewServer
   ```
7. After the move finishes successfully, on HVR2_1, open Hyper-V Manager and confirm that the virtual machine runs correctly.
8. On HVR2_1, open the Failover Cluster Manager, and then select **Roles**.
9. In the **Actions** pane, select **Configure Roles** to open the **High Availability Wizard**.

   On the **Select Role** page, select **Virtual Machine**.

   On the **Select Virtual Machine** page, select **Testvm_1**.
10. The virtual machine is now highly available.
11. Update the integration services on **Testvm_1**. Because the virtual machine was never shut down during the migration, you can update the integration services silently without a restart. The update occurs the next time that the virtual machine is restarted during its scheduled maintenance period.
    a. Modify the settings of the virtual machine and specify the following media to be used for the CD/DVD drive, *%Systemroot%\System32\Vmguest.iso*.
    b. Run the following command from an elevated command prompt in the virtual machine:
       i. For 64-bit Windows Server operating systems, *drive:\\Support\Amd64\Setup.exe /quiet /norestart*.
       ii. For 32-bit Windows Server operating systems, *drive:\\Support\X86\Setup.exe /quiet /norestart*.
12. Repeat steps 1 – 11 for all of the virtual machines on HVSRV1.
13. **[Optional]** For virtual machines that are moved from a Hyper-V Replica server, you must update the virtual machine replication settings on the Hyper-V primary server to reestablish replication.
    a. Open Failover Cluster Manager on the cluster where the primary virtual machine is running and select **Roles**.

b. Select the virtual machine, and then select **Settings** from the **Actions** pane.

c. Select **Replication** and update the value for the **Replica server** with the name of the Hyper-V Replica Broker that runs on the new cluster, HVHAR2.

Migrate the old cluster node to the new cluster

After all of the virtual machines from HVSRV1 have been moved to the HVHAR2 cluster, you can evict HVSRV1 from the HVHA201 cluster, install Windows Server 2012 R2, and join the HVHAR2 cluster.

To evict the node from the old cluster

1. On HVSRV2, open the Failover Cluster Manager and select **Nodes**.
2. Select **HVSRV1** and confirm that there are no virtual machines that have moved to **HVSRV1**.
3. Right-click **HVSRV1** to select **More Actions**, and then select **Evict**.

To install Windows Server 2012 R2 and join the new Hyper-V cluster

1. Install Windows Server 2012 R2 on HVSRV1.
2. Install the Hyper-V role and Failover Clustering feature if a clean installation was performed.
3. On HVR2_1, open the Failover Cluster Manager, and then select **Nodes**.
4. In the **Actions** pane, select **Add Node** to open the **Add Node Wizard**.

   Enter the name of the server to be added to the cluster, **HVSRV1**.

   Review the report.

To move the remaining virtual machines

Repeat the following procedures for the HVSRV2 and HVSRV3 servers to complete the migration.

- To move the virtual machines to the new cluster
- Migrate the old cluster node to the new cluster

Copy Cluster Roles Wizard

The Copy Cluster Role Wizard helps you copy clustered roles from clusters running Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2 to a new cluster running Windows Server 2012 R2. After the virtual machine has been created on the new cluster, to complete the migration, you must shut down each virtual machine on the source cluster before you start the virtual machine on the destination cluster.

You can use the Copy Cluster Roles Wizard to do the following:

- Test the Copy Cluster Roles process without affecting production.
- Reverse the process if you encounter any issues.

- Copy roles on test clusters to production clusters.

The Copy Cluster Role Wizard assumes that storage is reused between the old cluster and the new cluster. The cluster role settings are the only data that is copied.

Before migration of the virtual machines from the old cluster, perform the following actions:

- Before running the wizard, you must ensure that the new Windows Server 2012 R2 cluster is configured and is connected to the same storage as the old Hyper-V cluster. For more information about installing a Hyper-V cluster, see Deploy a Hyper-V Cluster.
- Before you work with shadow copies, you should back up all volumes that are attached to the virtual machines.
- Merge or discard all shadow copies for the volumes that store the virtual machines.
- Install the latest updates on all cluster nodes on both clusters.

Important

When you run the Copy Cluster Roles Wizard, the virtual machines are created on the new cluster, but they are not turned on. The virtual machines on the old cluster are still running. After the wizard has finished, you must turn off the virtual machines on the old cluster, and then, on the new cluster, you must start the virtual machines. There is some downtime but its duration should be limited, and you control when the downtime occurs.

To run the Copy Cluster Roles Wizard

1. You must have local Administrator rights on the new and old cluster to run the Copy Cluster Roles Wizard.
2. On the new cluster, open Failover Cluster Manager.
3. Select the top node for the cluster, and click **Copy Cluster Roles** from the **Configure** window.
4. On the **Specify Old Cluster** page, enter the name of the old cluster.
5. On the **Select Roles** page, select the role that you must copy.
6. On the **Customize Virtual Machine Networks** page, specify which virtual network switch the virtual machines are to use on the new cluster.
7. Review the settings and complete the wizard.
8. Review the Failover Cluster Copy Roles Report for any issues.
9. The virtual machines are still running on the old cluster, and the virtual machines that are created on the new cluster are shut off.

To run the virtual machine on new cluster

1. On the old cluster, open Failover Cluster Manager.
2. Turn off the virtual machines that have been copied over to the new cluster.

   Warning

   At no time should a virtual machine be running on both the old cluster and the new cluster. A virtual machine that runs on both the old cluster and the new

cluster at the same time might become corrupted. You can run a virtual machine on the old cluster while you migrate it to a new cluster with no problems; the virtual machine on the new cluster is created in a Stopped state. However, to avoid corruption, it is important that you do not turn on the virtual machine on the new cluster until after you stop the virtual machine on the old cluster.

3. On the new cluster, open Failover Cluster Manager.
4. Start the virtual machines and verify that the virtual machine runs correctly.

   Note
   If the migrated cluster is a Hyper-V Replica server, do not start the virtual machines and go to step 6.

   - Run the necessary application-specific tests to ensure that the application on the virtual machine can provide the same service levels as it provided before the virtual machine was migrated.
   - Verify that you can connect to the virtual machine by using **Remote Desktop** or **Virtual Machine Connection**.
5. Update integration services on each virtual machine.
6. **[Optional]** For virtual machines that were copied from a Hyper-V Replica server, you must remove replication and re-enable replication of the virtual machine on the Hyper-V primary server to reestablish replication.
   a. Open Failover Cluster Manager on the cluster where the primary virtual machine is running and select **Roles**.
   b. Select the virtual machine whose Replica virtual machine was copied, and in the **Actions** pane, select **Replication**, and then select **Remove Replication**.
   c. Select the virtual machine, and in the **Actions** pane, select **Replication**, and then select **Enable Replication**. This process opens the **Enable Replication Wizard**.
      - On the **Specify Replica Server** page, specify the name of the Hyper-V Replica Broker in the **Replica server** box.
      - On the **Choose Initial Replication Method** page, select **Use an existing virtual machine on the Replica server as the initial copy**.
7. **[Optional]** For virtual machines that are copied from a Hyper-V primary server, you must remove replication from the Replica virtual machine and enable replication on the virtual machine on the Hyper-V primary server to re-establish replication.

   Perform the following steps on the Replica virtual machine:

   a. Open Failover Cluster Manager on the cluster where the Replica virtual machine is running and select **Roles**.
   b. Select the virtual machine whose primary virtual machine was copied, and in the **Actions** pane, select **Replication**, and then select **Remove Replication**.

   Perform the following steps on the primary virtual machine:

   a. Open Failover Cluster Manager on the new cluster where the primary virtual machine is running and select **Roles**.
   b. Select the virtual machine that was just copied, and in the **Actions** pane, select

Replication, and then select **Enable Replication**. This process opens the **Enable Replication Wizard**.

- On the **Specify Replica Server** page, specify the name of the Hyper-V Replica Broker in the **Replica server** box.
- On the **Choose Initial Replication Method** page, select **Use an existing virtual machine on the Replica server as the initial copy**.

See also

Migrating Clustered Services and Applications to Windows Server 2012

Configuring Hyper-V Replica Broker in a Failover Cluster

Hyper-V Replica Feature Overview

Hyper-V Cluster Using Cluster Shared Volumes (CSV) Migration

This scenario describes how to migrate virtual machines from a Hyper-V cluster by using Cluster Shared Volumes (CSV) that run on the Windows Server 2012 operating system to the Windows Server 2012 R2 operating system. This scenario reuses the existing CSVs. Migrating the storage to a Scale-out File Server is beyond the scope of the scenario.

The Copy Cluster Roles Wizard is used to move the virtual machine roles to the new cluster.

Copy Cluster Roles Wizard

The Copy Cluster Role Wizard helps you copy cluster roles from clusters that are running Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2 to a new cluster that is running Windows Server 2012 R2.

The Copy Cluster Role Wizard assumes that storage is to be reused between the old cluster and the new cluster. The only data that is copied is the cluster role settings.

You can use the Copy Cluster Roles Wizard to do the following:

- Test the Copy Cluster Roles process without affecting production.
- Reverse the process if you encounter any issues.
- Copy roles on test clusters to production clusters.

Before you migrate the virtual machines from the old cluster, perform the following actions:

- Before you run the wizard, you must ensure that the new Windows Server 2012 R2 cluster is configured and is connected to the same logical unit numbers (LUNs) storage as the old

Hyper-V cluster. For more information about installing a Hyper-V cluster, see Deploy a Hyper-V Cluster.

- Before you work with shadow copies, you should back up all volumes that are attached to the virtual machines.
- Merge or discard all shadow copies for the volumes that store the virtual machines.
- Ensure that no virtual machines that you do not want to migrate share a CSV volume with virtual machines that you plan to migrate.
- Install the latest updates on all cluster nodes on both clusters.

Important

When you run the Copy Cluster Roles Wizard, the virtual machines are created on the new cluster, but they are not turned on. The virtual machines on the old cluster are still running. After the wizard has finished, you must turn off the virtual machines and take the disk offline on the old cluster. Then, on the new cluster, you must enable the disk and start the virtual machines. There is some downtime, but its duration should be limited, and you control when the downtime occurs.

Caution

The Copy Cluster Roles Wizard does not copy the replication settings for a virtual machine. After a virtual machine that has Hyper-V replication enabled is moved by using the Copy Cluster Roles Wizard, Hyper-V replication must be removed and be re-enabled.

To run the Copy Cluster Roles Wizard

1. You must be a local administrator on the new and old clusters to run the Copy Cluster Roles Wizard.
2. On the new cluster, open Failover Cluster Manager.
3. Select the top node for the cluster, and then click **Copy Cluster Roles** in the **Configure** window.
4. On the **Specify Old Cluster** page, enter the name of the old cluster.
5. On the **Select Roles** page, select the role that you want to copy.

   Note

   All the virtual machines that are running on a CSV must be migrated at the same time. When you select one virtual machine on a CSV, it automatically selects all of the virtual machines on that CSV.
6. On the **Customize Virtual Machine Networks** page, specify which virtual network switch is to be used by the virtual machines on the new cluster.

   Click **View Report**, to view the Pre-migration report.
7. Review the settings and complete the wizard.
8. Review the Failover Cluster Copy Roles Report to verify that the virtual machines were migrated.
9. The virtual machines are still running on the old cluster, and the virtual machines that are

created on the new cluster are off. Additionally, the CSV disk on the new cluster is offline.

To run the virtual machine on new cluster

1. On the old cluster, open Failover Cluster Manager.
2. Turn off the virtual machines that have been migrated over to the new cluster.
3. Take the CSV disk offline.
4. [Optional] In the storage, unmask the CSV disk so that the old cluster can no longer use it.

   Note
   Depending on storage topology, LUN masking and LUN unmasking might be necessary at this stage to ensure that the old cluster does not have write permission to the disks or LUNS that are used by the new cluster.

   Warning
   At no time should a virtual machine run on both the old cluster and the new cluster. A virtual machine that runs on both the old cluster and the new cluster at the same time might become corrupted. You can run a virtual machine on the old cluster while you migrate it to a new cluster with no problems; the virtual machine on the new cluster is created in a Stopped state. However, to avoid corruption, it is important that you do not turn on the virtual machine on the new cluster until after you stop the virtual machine on the old cluster.
5. On the new cluster, open Failover Cluster Manager.
6. Bring the CSV disk online.
7. Start the virtual machines and verify that the virtual machine runs correctly.

   Note
   If the cluster that is migrated is a Hyper-V Replica server, do not start the virtual machines and go to step 9.

   - Run the necessary application-specific tests to ensure that the application on the virtual machine can provide the same service levels as it provided before the virtual machine was migrated.
   - Verify that you can connect to the virtual machine by using **Remote Desktop** or **Virtual Machine Connection**.
8. Update integration services on each virtual machine.
9. **[Optional]** For virtual machines that are copied from a Hyper-V Replica server, you must remove replication and re-enable replication for the virtual machine on the Hyper-V primary server to re-establish replication.
   a. Open Failover Cluster Manager on the cluster where the primary virtual machine is running and select **Roles**.
   b. Select the virtual machine whose Replica virtual machine was copied, and in the **Actions** pane, select **Replication**, and then select **Remove Replication**.

c. Select the virtual machine, and in the **Actions** pane, select **Replication**, and then select **Enable Replication**. This action opens the **Enable Replication Wizard**.
    - On the **Specify Replica Server** page, specify the name of the Hyper-V Replica Broker in the **Replica server** box.
    - On the **Choose Initial Replication Method** page, select **Use an existing virtual machine on the Replica server as the initial copy**.

10. **[Optional]** For virtual machines that are copied from a Hyper-V primary server, you must remove replication from the Replica virtual machine and enable replication on the virtual machine on the Hyper-V primary server to re-establish replication.

    Perform the following steps on the Replica virtual machine:

    a. Open Failover Cluster Manager on the cluster where the Replica virtual machine is running and select **Roles**.
    b. Select the virtual machine whose primary virtual machine was copied, and in the **Actions** pane, select **Replication**, and then select **Remove Replication**.

    Perform the following steps on the primary virtual machine:

    a. Open Failover Cluster Manager on the new cluster where the primary virtual machine is running and select **Roles**.
    b. Select the virtual machine that was just copied, and in the **Actions** pane, select **Replication**, and then select **Enable Replication**. This action opens the **Enable Replication Wizard**.
        - On the **Specify Replica Server** page, specify the name of the Hyper-V Replica Broker in the **Replica server** box.
        - On the **Choose Initial Replication Method** page, select **Use an existing virtual machine on the Replica server as the initial copy**.

See also

Migrating Clustered Services and Applications to Windows Server 2012

Cluster Migrations Involving New Storage: Mount Points

Deploy Hyper-V Replica

Migrate File and Storage Services to Windows Server 2012 R2

The File and Storage Services Migration Guide provides step-by-step instructions for how to migrate the File and Storage Services role, including data, shared folders, and operating system settings from a source server to a destination server that is running Windows Server 2012 R2.

About this guide

Note
Your detailed feedback is very important and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Click **Rate this topic** at the top of the page and describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server 2012 forum.

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012 R2. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process. For more information about installing and using the migration tools on both source and destination servers, see Migrating Roles and Features in Windows Server.

Specifically, this guide includes information about migrating the following:

- Information about the server's identity
- Local users and groups
- Data and shared folders
- Shadow Copies of Shared Folders
- Data Deduplication
- DFS Namespaces
- DFS Replication
- File Server Resource Manager (FSRM)
- Group Policy settings that are specific to server message block (SMB)
- Group Policy settings for Offline Files (also known as client-side caching or CSC)
- iSCSI Software Target

  Note
  iSCSI Software Target was previously an optional Windows Server and Windows Storage Server component download. Because of the amount of content, all iSCSI-specific migration information is located in File and Storage Services: Appendix C: Migrate iSCSI Software Target.
- Network File System (NFS) file shares
- Remote Volume Shadow Copy Service (RVSS)

Target audience

This document is intended for information technology (IT) professionals and knowledge workers who are responsible for operating and deploying file servers in a managed environment.

What this guide does not provide

This guide does not provide information or support for the following migration scenarios:

- Migrating Roaming User Profiles (for additional information see **Upgrading or Migrating a Roaming User Profiles Environment to Windows 8.1 or Windows Server 2012 R2**).
- Upgrading roles on the same computer
- Migrating more than one server role
- Migrating data across subnets
- Migrating file servers by using File Server Resource Manager
- Migrating encrypted files from Encrypting File System (EFS)
- Migrating file allocation tables (FAT) and FAT32 file systems
- Migrating hardware and software installation for storage resources

In addition to these unsupported scenarios, you should understand the following migration limitations:

- If the hard disk drive that contains your data is physically moved from the source server to the destination server, file and folder permissions for local users are not preserved.
- Reparse points, hard links, and mounted volumes are not migrated when data is copied, and they need to be migrated manually.
- To facilitate migrating file and shared folder permissions, you must migrate local users and groups as part of the migration procedure. However, not all user and group attributes are migrated.

  For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide.

Supported migration scenarios

This guide provides instructions for migrating an existing server that is running File and Storage Services to a server that is running Windows Server 2012 R2 or Windows Server 2012. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure for your server environment, based on the information that is provided in other server role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other server roles that are running on the source server to fail.

Supported migration scenarios include the following configurations or features:

- File server is joined to a domain
- File server is in a workgroup

- File server data and file shares are located in a storage area network (SAN) or other external storage location that preserves data and file share permissions (except data for local users and groups)
- File server data and file shares are located on the server disk (direct-attached storage) that is preserving data and files shares permissions
- DFS Namespaces
- File Server Resource Manager
- iSCSI Software Target
- Network File System (NFS) file shares
- Shadow Copies of Shared Folders

Important
The file migration portion of the Windows Server Migration Tools is designed for smaller data sets (less than 100 GB of data). It copies files one at a time over HTTPS. For larger datasets, we recommend using the version of Robocopy.exe included with Windows Server 2012 R2 or Windows Server 2012.

Supported operating systems

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86-based or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86-based or x64-based	Windows Server 2003 R2	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86-based or x64-based	Windows Server 2008, full installation option	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
		installation options	
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core installation options	x64-based
x64-based	Server Core and full installation options of Windows Server 2012	Windows Server 2012 R2 or Windows Server 2008 R2, both full and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported. Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers.

Migrations between physical operating systems and virtual operating systems are supported. Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use to migrate roles, operating system settings, data, or shares from a computer that is running in the French system UI language to a computer that is running in the German system UI language.Windows Server 2012Windows Server 2008Windows Server Migration Tools

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Both x86-based and x64-based migrations are supported for Windows Server 2008 R2 and Windows Server 2003. All editions of Windows Server 2008 R2 are x64-based.

File services migration overview

The following topics contain step-by-step information about how to migrate File and Storage Services from a computer that is running Windows Server 2003 or later to a computer that is running Windows Server 2012 R2:

- File and Storage Services: Prepare to Migrate
- File and Storage Services: Migrate the File and Storage Services Role
- File and Storage Services: Verify the Migration

- File and Storage Services: Post-Migration Tasks

Impact of migration on other computers in the enterprise

The content in this section describes the impact to the computers in your enterprise during migration.

Impact of data migration by copying data and shared folders

- The performance of your source server can be affected during the data migration. This can result in slower access to files that are stored on the server.
- At the beginning of the second phase of the data migration, all open files are closed, which can lead to data loss.
- During the second phase of data migration, clients are unable to access the file server.

Impact of data migration by physically moving data drives

Clients cannot access the file server from the moment the storage device is disconnected from the source server until it is attached to the destination server.

Impact on DFS Namespaces

The DFS Namespaces are unavailable at several times during the migration process. You should plan your migration when you can take the namespace offline that is hosted on the source server.

Impact on DFS Replication

The impact of migration activity on other servers in the enterprise depends largely on the configuration of the replication topology. Typically, DFS Replication is configured in a hub and spoke replication topology with multiple branch office servers (spokes) replicating with a single hub server. Depending on which server in the replication topology is migrated and how the data is migrated, the remaining servers in the enterprise can be affected. Client workstations that are accessing data that is contained in the replicated folder on the server can be affected during the migration process.

Client computers may be accessing data in the folder that is being replicated by using DFS Replication. The replicated folder is often exposed to client computers as an SMB shared folder.

For more information about the impact of the migration process on client computers, see Impact of data migration by copying data and shared folders earlier in this document.

Permissions required to complete migration

This section describes permissions that are required to perform the migration.

Permissions required for data and shared folder migration

For data and shared folder migration, local Administrator permissions are required on the source server and destination server.

Permissions required to complete migration on the destination server

This section describes permissions that are required to perform the migration on the destination server.

Permissions required to migrate DFS Namespaces

For a stand-alone namespace, the user must be a member of the local Administrators group on the destination server.

There are three permissions options for a domain-based namespace:

- Option 1: Membership in the Domain Admins group
- Option 2 (if there are more than one namespace servers):
  - Permission to administer all namespaces that are hosted on the source server
  - Member of the local Administrators group on the destination server
- Option 3 (if there is a single namespace server):
  - Permission to delete and create domain-based namespaces in the domain
  - Member of the local Administrators group on the destination server

Permissions required to complete migration on the source server

This section describes permissions that are required to perform the migration on the source server.

Permissions required to migrate DFS Namespaces

For a stand-alone namespace, the user must have membership in the local Administrators group on the source server.

There are three permissions options for a domain-based namespace:

- Option 1: Membership in the Domain Admins group
- Option 2 (if there are more than one namespace servers):
  - Permission to administer the all namespaces that are hosted on the source server
  - Member of the local Administrators group on the source server
- Option 3 (if there is a single namespace server):
  - Permission to delete and create domain-based namespaces in the domain
  - Member of the local Administrators group on the destination server

Permissions required for DFS Replication

For DFS Replication, the user who starts the migration must be a member of the Domain Admins group or delegated permissions to the replication groups and replication members. This permission is required to remove the source server from replication groups to which it belongs. If the permissions to administer a replication group have been delegated to a particular user through the DFS Management snap-in, that user can use the DFS Management snap-in to perform tasks such as removing the source server from a replication group. The user must also be a member of the local Administrators group on the source server and the destination server.

See also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Migrate an iSCSI Software Target

File and Storage Services: Migrate Network File System

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

Migrating Roles and Features in Windows Server

File and Storage Services: Prepare to Migrate

This guide provides you with instructions for migrating the File and Storage Services role to a server that is running Windows Server 2012 R2.

Install migration tools

Windows Server Migration Tools in Windows Server 2012 R2 allows an administrator to migrate some server roles, features, operating system settings, file shares, and other data from computers that are running certain editions of Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 to computers that are running Windows Server 2012 R2.

For complete installation, configuration, and removal instructions for Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.

Migration documentation and tools ease the process of migrating server role settings and data from an existing server that is running a Windows server operating system to another computer. For a complete list of supported operating systems, see Migrate File and Storage Services to Windows Server 2012 R2.

By using these tools to migrate roles, you can simplify migration, reduce migration time, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

Prepare for migration

The following list outlines the major steps for preparing to migrate the File and Storage Services role.

- Prepare the destination server
- Back up the source server
- Prepare the source server
- Prepare other computers in the enterprise

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

Prepare the destination server

Use the following information to prepare the destination server for migration.

Hardware requirements for the destination server

Verify that the data locations for the destination server have sufficient free space to migrate the data. Ensure that the destination server hard disk drives are the same size or larger than the source server hard disk drives.

Software requirements for the destination server

There are several software requirements that must be met to ensure a successful migration.

- Consult the migration matrix to determine if you can migrate the version of Windows Server that you are running on the source server to Windows Server 2012 R2. For a complete list of supported operating systems, see Migrate File and Storage Services to Windows Server 2012 R2.
- Before migration, install all critical updates and service packs on the source server that were released before Windows Server 2012 R2. It is a recommended best practice that you install all current critical updates and service packs on the source server and the destination server.

Prepare for local user and group migration on the destination server

Verify that the destination server can resolve the names of domain users who are members of the local group during the import operation. If the source server and destination server are in different domains, the destination server must be able to contact a global catalog server for the forest in which the source domain user accounts are located.

Prepare for File and Storage Services on destination server

1. Install Windows Server 2012 R2 on the destination server.
2. Ensure that the time and date are set correctly on the destination server, and that they are in sync with the source server.
3. Determine the File Services role services that have been installed on the source server and then install the same File and Storage Services role services on the destination server.
4. Install Windows Server Migration Tools on the destination server.

   For more information about how to install Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.
5. Open UDP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to establish a data transfer connection.

   Note
   If you have changed the default behavior of Windows Firewall to block outbound traffic on computers that are running Windows Server 2012 R2, you must explicitly allow outbound traffic on UDP port 7000.
6. Open TCP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to perform the data transfer.

   For more information about how to open UDP port 7000 and TCP port 7000, see File and Storage Services: Appendix A: Optional Procedures.

   For more information about how to determine if a port is in use, see How To Determine Which Program Uses or Blocks Specific Transmission Control Protocol Ports in Windows Server 2003.
7. Verify that the destination path has sufficient disk space to migrate the data. If NTFS or folder quota management (in File Server Resource Manager) is enabled on the destination server disk drive, verify that the quota limit allows for sufficient free disk space to migrate data. For more information about quota management in File Server Resource Manager, see one of the following:

   - Quota Management for Windows Server 2008 R2 and Windows Server 2008
   - Quota Management for Windows Server 2003 R2

   For more information about NTFS quota management, see one of the following.

   - Setting Disk Quotas for Windows Server 2008 R2 and Windows Server 2008
   - Enable disk quotas for Windows Server 2003 R2 and Windows Server 2003

Prepare File Server Resource Manager on destination server

If you are using File Classification Infrastructure plug-ins from a non-Microsoft vendor, you should register the non-Microsoft plug-ins on the destination server and refer to additional instructions for migration from the non-Microsoft plug-in vendor. You should register the plug-in after File Server Resource Manager (FSRM) has been installed and started on the destination server.

Use the same drive letters for the destination server volumes as for the source server. This is required, because FSRM migration requires the drive letter to remain the same.

Data and file share preparation on destination server

Do not allow users to access the destination server until migration is fully completed. This ensures data integrity and prevents failure when an open file on the destination server cannot be overwritten during migration.

Data integrity and security considerations on destination server

Server migration tools preserve file and folder permissions during data migration. When you are planning the migration, keep in mind that if the migrated files and folders inherit permissions from their parents, during migration it is the inheritance setting that is migrated, not the inherited permissions. Therefore, it is important to make sure that the parent folders on the source server and the destination server have the same permissions to maintain the permissions on migrated data that has inherited permissions.

For example:

1. Migrate folder **c:\A\C** from the source server to folder **c:\B\D** on the destination server.
2. Verify that on the source server, only Mary has access to folder **c:\A** and folder **c:\A\C** is specified to inherit permission from its parent.
3. Verify that on the destination server, only John has access to folder **c:\B**. After **c:\A\C** is migrated to **c:\B\D**, John will have access to folder **D**, but Mary will not.

If you use permissions inheritance for the migrated data, ensure that the parent folder for the migrated data on the destination server has the required permission set.

Prepare DFS Namespaces on destination server

The DFS Namespaces role service must be installed, and the DFS Namespace service must be running before migration. If the namespaces that you are migrating are domain-based, both source and destination servers must be in the same Active Directory Domain Services (AD DS) domain. If the namespaces are stand-alone namespaces, AD DS membership does not matter.

Back up the source server

If DFS Namespaces are being migrated, back up the source server by using a full server backup or system state backup. If the DFS Namespaces are part of an AD DS domain, you need to back up the AD DS domain to save the Active Directory configuration information for DFS Namespaces.

For each domain-based DFS namespace, you should also back up the configuration information for the namespace. Repeat the following command for each namespace and save the output file name to a safe location:

```
DFSUtil.exe root export <\\<DomainName>\Namespace> <Filename>
```

Note

DFSUtil.exe is available on computers that are running Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. It is available to download for use on Windows Server 2003 R2 and Windows Server 2003 as part of the Windows Server 2003 Service Pack 1 32-bit Support Tools.

Prepare the source server

The following sections describe how to prepare the source server for the migration.

Prepare all file services on source server

- Install Windows Server Migration Tools on the source server.

  For more information about how to install Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.
- Verify that the time and date are set correctly on the destination server and that they are synchronized with the source server.
- Open UDP port 7000 and make sure that is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to establish a data transfer connection.

  Note

  If you have changed the default behavior of Windows Firewall to block outbound traffic on computers that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, you must explicitly allow outbound traffic on UDP port 7000.
- Open TCP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to perform the data transfer.

For more information about how to open UDP port 7000 and TCP port 7000, see File and Storage Services: Appendix A: Optional Procedures.

For more information about how to determine if a port is in use, see How To Determine Which Program Uses or Blocks Specific Transmission Control Protocol Ports in Windows Server 2003.

Data and file share preparation on the source server

To minimize downtime and reduce impact to users, plan your data migration to occur during off-peak hours. Use the `net share` command to list all file shares on the source server.

You can use this list during the verification step to verify that all the required file shares have migrated. Reparse points and hard links will not migrate when data is copied (versus a physical

migration), and the reparse points need to be migrated manually. When you migrate hard links, a separate file is created on the destination server for each link. If your migration involves copying the data to the destination server, follow the instructions for how to detect the reparse points and hard links in File and Storage Services: Appendix A: Optional Procedures. Afterward, you can remap and recreate them during migration, as instructed in the For copy data migration scenarios section.

Prepare DFS on the source server

DFS Namespaces role services must be installed, and the DFS Namespace service must be running before migration.

For information about DFS Namespaces preparation, see Prepare DFS Namespaces on source server.

Prepare DFS Namespaces on source server

For domain-based namespaces with one namespace server, determine if you will add a temporary server to the namespace or if you will perform a manual inventory of the namespace permissions.

- Option 1 (recommended):

  Add a temporary server as a namespace server to each domain-based namespace on the source server when the source server is the only namespace server.
- Option 2:

  Inventory the permissions for managing each of the namespaces that are hosted on the source server when the source server is the only namespace server. This process can be completed by using the DFS Management MMC Snap-in.

Prepare other computers in the enterprise

Data and file share migration requires preparing other computers in the enterprise. Perform the following steps for copy data migration scenarios, and for physical data scenarios.

For copy data migration scenarios

- Notify the users that the server performance may be reduced during the first phase of data migration.
- Ask users to stop writing to the server before the second phase of data migration begins (to prevent possible data loss). We recommend that you prevent access to the file shares so that users do not ignore this advice. For example, you could temporarily set all file shares to be read-only by setting the share permissions to Everyone = Read Only.
- Notify users that they cannot access their files on the server when the second phase of the migration begins until the file server migration is fully completed.

For physical data migration scenarios

Notify the users that they cannot access the file server from the moment the storage is disconnected from the source server until the server migration is fully completed.

See also

- Migrate File and Storage Services to Windows Server 2012 R2
- File and Storage Services: Migrate the File and Storage Services Role
- File and Storage Services: Verify the Migration
- File and Storage Services: Migrate an iSCSI Software Target
- File and Storage Services: Migrate Network File System
- File and Storage Services: Post-Migration Tasks
- File and Storage Services: Appendix A: Optional Procedures
- File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Migrate the File and Storage Services Role

Migrate File Services

Perform the following tasks to migrate the File and Storage Services server role.

- Freeze administration configuration
- Export settings
- Migrate local users and groups to the destination server
- Migrate data
- Migrate the source server identity
- Export Remote VSS settings
- Import settings to the destination server

Freeze administration configuration

Administrators must stop all configuration changes to the File and Storage Services role services on the source server before starting migration. When the migration begins, you must not make any configuration changes to the source server other than those that are required for the migration. For example, no links can be added to a DFS namespace after migration starts until the migration is successfully verified.

Install the Windows Server Migration Tools

Before you can use any of the following Windows PowerShell cmdlets for migration on the source server or destination server, ensure that the Windows Server Migration Tools are added. You can do this by using Server Manager or by using Windows PowerShell.

To install the Windows Server Migration Tools

1. Log on to the computer as a member of the local Administrators security group.
2. In Server Manager, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, select **the Role-base or feature-based installation** option, and then click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, accept the default selections, and then click **Next**.
7. On the **Select features** page, click **Windows Server Migration Tools**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. After the installation is complete, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature Migration
```

The following is a list of Windows Server Migration Tools cmdlets:

- **Export-SmigServerSetting**
- **Import-SmigServerSetting**
- **Get-SmigServerFeature**
- **Send-SmigServerData**
- **Receive-SmigServerData**

For more information about how to work with the Windows Server Migration Tools see Install, Use, and Remove Windows Server Migration Tools.

Export settings

Export the following settings from the source server to the destination server:

- Server Message Block (SMB)
- Offline Files (also known as called client-side caching or CSC)
- DFS Namespaces
- File Server Resource Manager (FSRM)
- Shadow Copies of Shared Folders

BranchCache for Network Files server key

The following procedure applies only if the source server is running Windows Server 2012 or Windows Server 2008 R2.

Notes

This procedure, which is used to migrate the seed value that is used by the BranchCache for Network Files component, enables data that was stored in branch office locations by using BranchCache to be used after the file server is migrated from the source server to the destination server.

For information about how to migrate a BranchCache host server, see the BranchCache Migration Guide.

To migrate BranchCache for Network Files settings from the source server

1. In your Windows PowerShell session, collect data from the source server by running the **Export-SmigServerSetting** cmdlet as a member of the Administrators security group. This step runs the **Export-SmigServerSetting** cmdlet with all parameters from a single command line. The **Export-SmigServerSetting** cmdlet parameters can collect all source BranchCache feature data in a single file (Svrmig.mig), or you can run the **Export-SmigServerSetting** cmdlet multiple times by using one or more parameters to collect and store data in multiple Svrmig.mig files.

   For more information, see the section "Prepare for migration" in File and Storage Services: Prepare to Migrate.

   Review the following dependencies before you run the command.

   - When you run the **Export-SmigServerSetting** cmdlet, you are prompted to provide a password to encrypt the migration store data. You must provide this same password to import data from the migration store.
   - The *path* parameter can be to a folder that is empty or one that contains data. The actual data file in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Therefore, the user does not have to specify a file name.
   - If the path is not a shared location that the destination server can read, you must manually copy the migration store to the destination server or a location that the destination server can access.
   - If a migration store location already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must move the Svrmig.mig file from the migration store location and then store it elsewhere, or rename or delete the Svrmig.mig file first.

2. On the source server, type the following, and then press Enter, where *<storepath>* is the path that will contain the Svrmig.mig file after this step is completed. An example of the path is *\\fileserver\users\username\branchcachestore*.

```
Export-SmigServerSetting -featureID BranchCache -Path
```

```
<storepath\BranchCache> -Verbose
```

Group Policy setting or local policy setting specific to SMB and Offline Files

Most SMB and Offline Files settings are migrated as part of shared folders migration. The remaining settings that affect the server are set through Group Policy settings or local policy settings. This section describes how to inventory SMB and Offline Files settings that are controlled by Group Policy.

Server message block

Determine the policy settings that affect the SMB server. The SMB settings are controlled by Group Policy settings or local policy settings. If a Group Policy Object (GPO) is applied, these policies override the local settings. First, determine if the policy settings are controlled by a GPO, and then determine local settings for anything that is not controlled by the GPO.

To determine if a GPO is applied to the source server

1. Open the Resultant Set of Policy snap-in. To open the Resultant Set of Policy snap-in, open a command prompt, type **rsop.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
3. Note in the SMB data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets any Group Policy setting that affects the following **Microsoft network server** settings:
   - Microsoft network server: Amount of idle time required before suspending session
   - Microsoft network server: Attempt S4USelf to obtain claim information
   - Microsoft network server: Digitally sign communications (always)
   - Microsoft network server: Digitally sign communications (if client agrees)
   - Microsoft network server: Disconnect clients when logon hours expire

   On source servers that are running the Server Core installation option of the Windows Server 2012 or Windows Server 2008 R2 operating system, run the **gpresult** command to review Group Policy settings (for more information about **gpresult**, type **gpresult /?** at a command prompt.)

Notes

For any setting that is controlled by Group Policy, you must apply the same GPO to the destination server, or you can set the local policy of the destination server for the same behavior.

For any setting that is not controlled by Group Policy, use the following procedure to determine the local policy setting. Note the local policy setting in the SMB data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets.

To determine local policy settings

1. Open the Local Group Policy Editor. To open the Local Group Policy Editor, open a command prompt, type **gpedit.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
3. Note the following settings for **Microsoft network server**:
   - Microsoft network server: Amount of idle time required before suspending session
   - Microsoft network server: Attempt S4USelf to obtain claim information
   - Microsoft network server: Digitally sign communications (always)
   - Microsoft network server: Digitally sign communications (if client agrees)
   - Microsoft network server: Disconnect clients when logon hours expire

   On source servers that are running the Server Core installation, run the **secedit** command to export and review local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt.)

Offline Files

Note
This section applies only to source servers that are running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or. Previous operating system releases do not have Offline Files settings that affect the server.

Determine the policy settings that affect shared folders on the server for which client computers use Offline Files. The Offline Files settings are controlled through Group Policy or local policy. If Group Policy is applied, these policies override local settings. First, determine if the settings are controlled through Group Policy, and then determine the local settings for anything that is not controlled by using Group Policy.

To determine if Group Policy is applied to the source server

1. Open the Resultant Set of Policy snap-in. To open the Resultant Set of Policy snap-in, open a command prompt, type **rsop.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Administrative Templates**, click **Network**, and then click **Lanman Server**.

   Note
   If no policies are set, the preceding path will not exist. If the path does not exist, skip to the procedure To determine local policy settings. If the path exists and policies are found, proceed to the next step.
3. Note in the BranchCache data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets any Group Policy settings that control the **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

On source servers that are running the Server Core installation option, run the **gpresult** command to review Group Policy settings (for more information about **gpresult**, type **gpresult /?** at a command prompt).

For any setting controlled by Group Policy, have the same Group Policy setting apply to the destination server, or you can choose to set the local policy setting of the destination server to get the same behavior.

For any setting not controlled by Group Policy, use the following instructions to determine the local policy setting.

To determine local policy settings

1. Open the Local Group Policy Editor. To open the Local Group Policy Editor, open a command prompt, type **gpedit.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Administrative Templates**, click **Network**, and then click **Lanman Server**.
3. Note in the BranchCache data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets the value of the **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

   On source servers that are running the Server Core installation option, run the **secedit** command to export and review local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

DFS Namespace configuration

Procedures in this section describe how to migrate DFS Namespaces from the source server to the destination server.

Before the migration of the namespace begins, you can inventory the namespaces that are hosted on the source server for tracking purposes. You can do this by using DFS Management or DFSUtil.exe.

The following procedure (To inventory DFS Namespaces by using DFS Management) applies only to computers that are running at least the Windows Server 2003 R2 version of the Windows Server operating system. For computers that are running Windows Server 2003, you can perform a DFS Namespace inventory by using **DFSUtil.exe** as described in To inventory DFS Namespaces by using DFSutil.exe.

You can also perform a DFS Namespace inventory from a client computer that is running Windows 8, Windows 7, or Windows Vista, by using DFS Management that is part of Remote Server Administration Tools.

- To download Remote Server Administration Tools for Windows 8, see Deploy Remote Server Administration Tools.
- To download Remote Server Administration Tools for Windows 7, see Remote Server Administration Tools for Windows 7.

- To download Remote Server Administration Tools for Windows Vista, see Microsoft Remote Server Administration Tools for Windows Vista.

To inventory DFS Namespaces by using DFS Management

1. Under DFS Management in the left pane, right-click **Namespaces**.
2. Click **Add Namespaces to Display**.
3. In the dialog box that is displayed, select **Server** from the Scope options.
4. Type the name of source server and click **Show Namespaces**.
5. Select all namespaces listed in the list box and click **OK**.
6. Right-click the first namespace listed in the left pane.
7. Click **Properties**.
8. On the **General** tab, check the **Type** field. The type of namespace that is hosted on the server is described here. Possible values are stand-alone, domain-based (Windows Server 2000 mode), and domain-based (Windows Server 2008 mode).
9. In the case of a domain-based namespace, click the **Namespace Servers** tab to identify the number of servers that host the namespace.
10. Repeat steps 7 through 10 for the remaining namespaces listed in the left pane.

To inventory DFS Namespaces by using DFSutil.exe

1. You can inventory your DFS Namespaces using DFSUtil.exe by using the command prompt. From a command prompt, type **DFSUtil.exe server SourceServer** where *SourceServer* represents the name of the source server.
2. Identify the namespaces (DFS roots) listed for the source server.
3. Type the following command, which lists the namespace properties for the first namespace identified in step 2:

   ```
   DFSUtil.exe root <\\SourceServer\Namespace>
   ```

4. Identify the namespace type; possible values are stand-alone root, domain root (domain-based namespace in Windows 2000 Server mode), domainV2 root (domain-based namespace in Windows 2008 mode).
5. Identify the DFS folders present in the namespace in each of the **Link Name** items displayed.
6. In the case of domain-based namespaces, identify all the namespace servers by typing the following command:

   ```
   DFSUtil.exe root <\\Domain\Namespace>
   ```

7. Identify the namespace servers that host the namespace in each of the **Target** items displayed under **Root Name**.
8. Repeat steps 3 through 7 for the remaining namespaces on the source server.

Considerations for namespaces

Is the namespace stand-alone or domain-based? If the namespace is stand-alone, see the following section in this document: Stand-alone namespaces.

If the namespace is domain-based, consider the number of namespace servers for each namespace. For more information, see the following sections in this document:

- Domain-based namespaces with more than one namespace server
- Domain-based namespaces with one namespace server

Stand-alone namespaces

Complete the following procedure to export a stand-alone namespace configuration.

To export the namespace configuration to an export file

1. On the destination server, open a Command Prompt window.
2. Type **DFSUtil.exe root export \\SourceServer\Namespace FileName** to export the stand-alone namespace to a file (where *FileName* represents the exported file), and then press Enter.

Domain-based namespaces with more than one namespace server

For more than one namespace server, remove the namespace server from the namespace by using DFSUtil.exe.

To remove the namespace server from the namespace

1. On the destination server, open a Command Prompt window.
2. Type **DFSUtil.exe target remove <\\SourceServer\Namespace>**, and then press Enter.

Domain-based namespaces with one namespace server

There are two options that you can use in this scenario: export the namespace settings, or add a temporary server to the namespace.

To export namespace settings

1. On the destination server, open a Command Prompt window.
2. Type **DFSUtil.exe root export \\Domain\Namespace FileName** where *FileName* represents the file containing settings for export, and then press Enter.

   Note
   For each namespace, there must be a different file name to export settings.
3. Repeat step 2 for each namespace for which the source server is a namespace server.

You can use either of the following two options if a temporary server can be added to the namespace. This provides the ability to maintain the namespace online while the migration progresses. If this is not possible, follow the steps in To remove the namespace server from the namespace instead.

To add a temporary server to the namespace by using DFS Management

1. In the left pane, select the namespace to be migrated.
2. Click the **Namespace servers** tab.
3. Select **Add Namespace Server**.
4. In the **Namespace server** box, type the name of the temporary server, and then click **OK**.

The temporary server will be added to the namespace.

To add a temporary server to the namespace by using DFSUtil.exe

1. Create a shared folder for the namespace on the temporary server with the same permissions as on the source server.
2. On the destination server, open a Command Prompt window.
3. Type **DFSUtil.exe target add \\TemporaryServer\Namespace** and then press Enter.

```
DFSUtil.exe target add <\\TemporaryServer\Namespace>
```

The temporary server will be added to the namespace.

After the namespace settings are exported or a temporary server is added to the namespace, the namespace source server can be removed from the namespace as described in To remove the namespace server from the namespace.

Inventory advanced registry keys

This section describes the process for determining if there are any settings that have been applied to the DFS Namespace component on the source server. These settings are stored in the registry and set or viewed through the DFSUtil.exe tool. To inventory these settings, run the following commands from the destination server:

```
DFSUtil.exe server registry DfsDnsConfig <SourceServer>

DFSUtil.exe server registry LdapTimeoutValue <SourceServer>

DFSUtil.exe server registry SyncInterval <SourceServer>
```

Note the setting for any registry modification. Registry keys that have not been modified display a value similar to the following:

```
<KeyName> does not exist in the Registry.
```

DFS Replication configuration

To migrate DFS Replication settings, use the following Microsoft Enterprise Support blog series: Replacing DFSR Member Hardware or OS.

File Server Resource Manager configuration on the source server

When you migrate File Server Resource Manager, remember to use the same drive letters for the destination server volumes as for the source server. This is required because the File Server Resource Manager migration requires that the drive letter remains the same.

1. Stop the File Server Resource Manager and File Server Storage Reports Manager services. You can stop these services in Windows PowerShell by using the following command: **Stop-Service –Name “srmsvc”, “srmreports”**.
2. Export the File Server Resource Manager configuration. You can export the File Server Resource Manager configuration in Windows PowerShell by using the following command: **Export-SmigServerSetting -FeatureID FS-Resource-Manager -Path <storepath\FSRM> -Verbose**.
3. For each volume, get the configuration files by running the following commands in the Windows PowerShell session.
   a. Stop the file screen driver. Type **fltmc detach datascrn <VolumeLetter>:** and then press Enter.
   b. Stop the quota driver. Type **fltmc detach quota <VolumeLetter>:** and then press Enter.
   c. Grant Read permissions to the Administrator account for the "<*VolumeLetter*>:\System Volume information\SRM" folder and the following child objects:
      - takeown /F "<VolumeLetter>:\System Volume Information" /A /R /D Y
      - cacls "<VolumeLetter>:\System Volume Information" /T /E /G Administrators:F
      - attrib -S -H "<VolumeLetter>:\System Volume Information*" /S /D
   d. Copy the following files from the **SRM** folder to an external storage device:
      - Quota.xml
      - Quota.md
      - Datascrn.md
      - DataScreenDatabase.xml
   e. Start the file screen driver. Type **fltmc attach datascrn <VolumeLetter>:** and then press Enter.
   f. Start the quota driver. Type **fltmc attach quota <VolumeLetter>:** and then press Enter.
4. Restart the File Server Resource Manager and File Server Storage Reports Manage services. Type **Start-Service -name "srmsvc","srmreports"**, and then press Enter.
5. Configure scheduled reports.

   File Server Resource Manager reports and classification rule configurations are dependent on the drive letters and mount points. Any drives or mount points on the source server that are used by report or classification rule configurations must be available on the destination server, or the configurations must be updated during import.

   To configure scheduled reports, follow step (a). However, if you are migrating from Windows Server 2003, follow step (b).

   - To configure scheduled reports on all servers except Windows Server 2003, run the following commands in a Windows PowerShell session on the source server that was opened with elevated user rights.

- To get a list of all the task names associated with storage reports: `storrept r l`
- For each task name listed, run the following command on the source server: `schtasks /query /tn:"TASKNAME" /xml > "TASKNAME.xml"`

- To configure scheduled reports when you migrate from Windows Server 2003:
  - On the source server, do the following:
    - Open File Server Resource Manager.
    - In **Storage Report Management**, for each report task, note the report task, target, and schedule.
  - On the destination server, after you import the file server resource manager configuration, do the following:
    - Open File Server Resource Manager.
    - In **Storage Report Management**, for each report task, edit the report task properties.
    - On the **Schedule** tab, manually add the appropriate schedule for the report.

6. Configure scheduled file management tasks. This step applies only to source servers that are running Windows Server 2012 or Windows Server 2008 R2.
   a. To display a list of all task names associated with file management tasks, type the following command on the source server in a Windows PowerShell session opened with elevated user rights:

      ```
      (new-object -com
      Fsrm.FsrmFileManagementJobManager).EnumFileManagementJobs()
      ```

   b. For each entry listed, locate the task element, and then type the following command:

      ```
      Schtasks /query /tn:"TASK" /xml > "TASK.xml"
      ```

7. Export the classification schedule. This is only applicable to servers running Windows Server 2012 or Windows Server 2008 R2 that already have a classification schedule configured. From an elevated command prompt, type the following command:

   ```
   Schtasks /query /tn:”FsrmAutoClassification{c94c42c4-08d5-473d-
   8b2d-98ea77d55acd}” /xml > “classification.xml”
   ```

Shadow Copies of Shared Folders

The following procedures describe how to migrate shadow copy settings.

To migrate shadow copy settings

1. Open Windows Explorer on the source server to view shadow copy storage locations and the creation schedule.

   Important
   This procedure applies to shadow copies for a server running the full installation option of Windows Server. If your source server is running the Server Core installation option of Windows Server, skip this procedure and follow the

instructions in the following section: To migrate shadow copies in a Server Core installation.

2. For each volume on the source server, right-click the volume, and then click **Configure Shadow Copies**.

   On source servers that are running Windows Server 2003, right-click the volume, click **Properties**, and then click the **Shadow Copies** tab.
3. Click **Settings**, and note the location and size of the shadow copy storage.
4. Click **Schedule** and note the details for the snapshot creation task.

To migrate shadow copies in a Server Core installation

1. Log on to the computer that is running a Server Core installation remotely as follows:
   a. In Server Manager, click **Tools**, and then click **Computer Management**.
   b. In the **Computer Management** snap-in pane, right-click the top node, and then click **Connect to another computer**.
2. Type the computer name, and then click **OK**.
3. Expand **System Tools**, right-click **Shared Folders**, click the **All Tasks** tab, and then click **Configure Shadow Copies**.
4. For each volume on the source server, right-click the volume, click **Configure Shadow Copies**, click **Settings**, and note the location and size of the shadow copy storage.
5. Click **Schedule**, and then note details for the snapshot creation task.

Migrate local users and groups to the destination server

Before migrating data and shared folders or completing your migration of the FSRM configuration, you must migrate local users and groups. Export local users and groups from the source server, and then import local users and groups to the destination server.

Important
If the source server is a domain member server, but the destination server is a domain controller, imported local users are elevated to domain users, and imported local groups become Domain Local groups on the destination server.

If the source server is a domain controller, but the destination server is not, Domain Local groups are migrated as local groups, and domain users are migrated as local users.

Export local users and groups from the source server

On the source server, export local users and groups to a migration store (as shown in the following example) in a Windows PowerShell session that has been opened with elevated user rights.

```
Export-SmigServerSetting -User All -Group -Path <storepath\UsersGroups> -Verbose
```

You can use one of the following values with the *-user* parameter:

- Enabled: Specify to export only enabled local users.
- Disabled: Specify to export only disabled local users.
- All: Specify to export enabled and disabled local users.

For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide.

You are prompted to provide a password to encrypt the migration store. Remember this password, because you must provide the same password to import from the migration store.

If the path is not a shared location that is accessible to the destination server, you must manually copy the contents of the migration store folder to the destination server or a location that is accessible to the destination server.

Import local users and groups to the destination server

On the destination server, import local users and groups from the migration store to which you exported the users and groups in Export local users and groups from the source server, as illustrated by the following example. Use a Windows PowerShell session that has been opened with elevated user rights.

```
Import-SmigServerSetting -User All -Group -Path <storepath\UsersGroups> -Verbose
```

You can use one of the following values with the *-user* parameter:

- Enabled: Specify to import only enabled local users.
- Disabled: Specify to import only disabled local users.
- All: Specify to import enabled and disabled local users.

For the list of user attributes that are supported for migration, see the Local User and Group Migration Guide.

You are prompted to provide the same password that you provided during export to decrypt the migration store.

Migrate data

To migrate data, you can copy file data or physically move it, for example, by attaching the storage drive from the source server to the destination server. If you copy the data, follow the copy data migration steps in the following section.

If you physically move the data, follow the steps described in the Physical data migration section later in this document.

Data copy migration

If you are planning a two-phase data copy migration as described in the previous section, note that if files have been deleted on the source server between the start of the first copy and the start of the final copy, copies of the deleted files might have already transferred to the destination server. So if a file is deleted between the two copy processes, the file might still be available on the destination server after the migration is complete. If this is unacceptable in your environment, perform data and shared folder migration in a single phase, and disconnect all users before starting migration.

Important

The file migration portion of the Windows Server Migration Tools is designed for smaller data sets (less than 100 GB of data). It copies files one at a time over HTTPS. For larger datasets, we recommend using the version of robocopy.exe included with Windows Server 2012 R2 or Windows Server 2012.

To copy data and shared folders and migrate all data without disconnecting users

1. Verify that the destination path has sufficient disk space to migrate the data. If NTFS or folder quota management is enabled on the destination server disk drive, verify that the NTFS or File Server Resource Manager quota limit allows for sufficient free disk space to migrate data. For more information about quota management in File Server Resource Manager, see one of the following:
   - Quota Management for Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008
   - Quota Management for Windows Server 2003 R2

   For more information about NTFS quota management, see one of the following:
   - Setting Disk Quotas for Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008
   - Enable disk quotas for Windows Server 2003 R2 and Windows Server 2003
2. Ensure that you have completed the migration of local users and groups.

   The **Send-SmigServerData** and **Receive-SmigServerData** cmdlets must be run on the source and destination server within five minutes of each other. By default, **Send-SmigServerData** and **Receive-SmigServerData** time out if a connection cannot be established within 300 seconds (five minutes). This maximum connection time-out for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets is stored in the following registry subkey, which is user-defined.

   Subkey: HKEY_LOCAL_MACHINE\Software\Microsoft\ServerMigration

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents connection time-out, in seconds)

   If a value larger than 3600 is specified, 3600 seconds (1 hour) is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key on

the Microsoft Web site.

3. Use the following command to run the **Receive-SmigServerData** cmdlet on the destination server. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Receive-SmigServerData
   ```

   Note
   All output for the Send and Receive operations occurs on the source server only. The destination server will appear to be done before the operation has completed.

4. Use the following command to run the **Send-SmigServerData** cmdlet on the source server to migrate data and shared folders. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Send-SmigServerData -ComputerName <DestinationServer> -
   SourcePath d:\users -DestinationPath d:\shares\users -Recurse
   -Include All -Force
   ```

The destination data location does not have to be the same as the source location, and you can change it, if desired.

Notes
The Server service startup type must be set to Automatic on the destination server for shared folder migration to complete.

Data that is transferred is encrypted automatically. You are prompted to enter a password to encrypt the transferred data on the source server, and the same password to decrypt the received data on the destination server.

After the first data copy is finished, you must freeze the source server and all data changes.

To disconnect users and migrate new or updated files

1. Make sure that users are notified that they should stop using the source server at this time to prevent any possible data loss. You can run the following command to list all the currently open files to determine the potential impact of performing this step.

   ```
   net file
   ```

2. Disconnect all users from the source server by stopping the LanMan Server service.

   ```
   Stop-Service LanmanServer -force
   ```

   Stopping the LanMan Server service invalidates all open remote files to the shared folders on the source server, which can lead to potential data loss. It is best to perform this step when the fewest users are expected to access files on this server.

3. Use the following command to run the **Receive-SmigServerData** cmdlet on the destination server. Use a Windows PowerShell session that is running with elevated user

rights.

```
Receive-SmigServerData
```

4. Use the following command to run the **Send-SmigServerData** cmdlet on the source server to migrate data and shared folders. Use a Windows PowerShell session that is running with elevated user rights.

```
Send-SmigServerData -ComputerName <DestinationServer> -
SourcePath d:\users -DestinationPath d:\shares\users -Recurse
-Include All -Force
```

5. If your scenario requires migrating reparse points, hard links, and mount points, recreate them on the destination server by using the **mklink** command for reparse points and hard links, and using the **mountvol** command for mounted volumes. For more information about these commands, enter `mklink /?` or `mountvol /?` in a Command Prompt window.

It is important to maintain the same destination path that you used during the first copy of data and shared folders. The cmdlets transfer files, folders, and shared folders only if they do not exist on the destination server, or if there is a new version on the source server.

Physical data migration

The next sections describe data migration by physically moving external drives or logical unit numbers (LUNs).

Using disk drives or LUNs to migrate data from the source server to the destination server

You can migrate data from the source server by moving the disk drives. Or, if your data resides on a LUN storage device, you have the option of moving the file server data by masking the LUNs from the source server and unmasking them on the destination server.

For the ideal migration, make sure that you maintain the same mapping of the drive letters (for example, drive D) and the volume IDs (see the following explanation) so that relevant data and application information remains as consistent as possible during the move.

Caution
You should not move a disk drive or LUN if it contains both data and the operating system.

Benefits of physical migration:

- For large amounts of data, this is a faster operation.
- You maintain all data on the disk drive, such as hard links and mount points.
- Shadow copies are preserved if the shadow copies are on the migrated disk drive.

Potential issues to be aware of:

- Permissions for local users that are not default computer accounts (such as local administrators) will be lost even if the same user name is used when creating the user account on the destination server.

- Encrypted files (EFS) cannot be migrated.
- Encrypted volumes with BitLocker cannot be migrated without first decrypting the volumes.
- Remote Storage cannot be migrated.
- When you are physically migrating disk drives that have File Server Resource Manager quotas enabled on them, it is a best practice to dismount the drive gracefully to avoid marking the quotas as dirty. Otherwise, unnecessary scans may occur later.

To migrate data by physically moving the disk drive or by masking and unmasking the LUNs

1. Collect information from the source server.

   Tip
   You can use Server Manager or Windows PowerShell on a computer running Windows Server 2012 or Windows 8 to collect information from source computers running Windows Server 2012.

   a. Record the drive letter and volume label for each data volume on the source server that you would like to move to the destination server.

   b. On the source server, export the volume GUID paths by exporting the following registry subkey to a file: **\HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices**. To do this, open the Registry Editor (regedit.exe), browse to the registry subkey, right-click the registry subkey, and clicking **Export**.

   Alternatively, to export the volume GUID paths from a server running Windows Server 2012 or Windows Server 2008 R2, open a Windows PowerShell session, and then type the following commands, where *<SourceServer>* is the name of the source server, *<Domain\User>* is a user account with administrative permissions on the source server and *<LocalPath>\<Filename>* is a local path and filename of the exported registry keys:

   ```
   Enter-PSSession <SourceServer> -Credential <Domain\User>
   Regedit.exe /E <LocalPath>\<Filename>.reg
   "HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices"
   ```

   Note
   To use Server Manager or Windows PowerShell to remotely collect information from earlier versions of Windows Server, you must first setup the source server for remote management. For more information, see Managing Downlevel Windows-based Servers from Server Manager in Windows Server 2012.

   c. Open Notepad, and copy the exported .reg file. Remove all entries that are in the following form: **\DosDevices\D:**. Save the.reg file (all remaining entries should be in the following form: **\??\Volume{ef93fe94-5dd7-11dd-961a-001e4cdb4059}**).

2. Prepare the destination server.

   a. In the Server Manager navigation pane, click **File and Storage Services**, and then

click **Volumes** to display the Volumes page. Use the Volumes tile to make sure that the drive letters for the data volumes are available. If there is a drive letter that is currently assigned to an existing volume on the destination server, change the drive letter for that volume.

Alternatively, use the Windows PowerShell **Get-Volume** and **Set-Partition** cmdlets. For example, to get any volumes with the drive letters of F, G, or H, type `Get-Volume F,G,H`. To change the drive letter of a partition with the F drive letter, type `Set-Partition -DriveLetter F -NewDriveLetter Z`

b. To import the volume GUID paths into the destination server, copy the.reg file that you created previously to the destination server, and then double-click that file to update the destination server.

3. Move the disk drives or LUNs from the source server to the destination server.
   a. On the source server, remove the disk drives or unassign the LUNs by using Storage Manager for SANs. (To open Storage Manager for SANs, click **Start**, click **Administrative Tools**, and then click **Storage Manager for SANs**.) If the source server is running Windows Server 2012, use the File and Storage Services role in Server Manager instead to view the disks or virtual disks (when using storage pools) that you want to move. If the disk is part of a storage pool, on the **Storage Pools** page of the File and Storage Services role right-click the virtual disk, and then click **Detach Virtual Disk**. For other types of disks, on the **Disks** page, right-click the disk that you want to move and then click **Take Offline**.
   b. On the destination server, attach each disk drive or assign the LUNs, and then assign the appropriate drive letter by using the **Disks** and **Storage Pools** pages of the File and Storage Services role in Server Manager.
4. If any files or folders on the migrated drive use local users or local groups permissions (except default users and groups), re-create these permission. Note that all domain users and groups permissions will remain intact, assuming that the source server and the destination server are members of the same domain.

Notes

You can use the `icacls` command to modify file and folder permissions (type `icacls /?` in a Command Prompt window for details). Type this same command in a Windows PowerShell session or a command prompt that has been opened with elevated user rights.

The list of the default users and groups is available in the topic Default User Accounts and Groups.

Migrate shared folders

If any of the folders on the migrated drive were shared on the source server and must be shared on the destination server, the following steps explain how to migrate shared folders.

1. If any of the migrated shared folders use local users and group permissions, ensure that you have completed the migration of local users and groups.

   The **Send-SmigServerData** and **Receive-SmigServerData** cmdlets must be started on the source server and the destination server within five minutes of each other. By default, **Send-SmigServerData** and **Receive-SmigServerData** operations terminate if a connection cannot be established within 300 seconds (five minutes). The maximum connection time-out for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets is stored in the following registry subkey, which is user-defined.

   Subkey: **\HKLM\Software\Microsoft\ServerMigration**

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents connection time-out, in seconds). If a value larger than 3600 is specified, 3600 seconds (one hour) is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key.
2. Open port 7000 on the source server and destination server (if this has not already been done).

   For information about how to open a port in Windows Firewall, see File and Storage Services: Appendix A: Optional Procedures.
3. On the destination server:
   a. Open a Windows PowerShell session with elevated user rights and enter the following command: **Receive-SmigServerData**.
4. On the source server:
   a. Open a Windows PowerShell session in Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. On computers that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, the Windows PowerShell session must be opened with elevated user rights. Enter the following command: **Send-SmigServerData -ComputerName <DestinationServerName> -SourcePath <SourcePath> -DestinationPath <DestinationPath> -Recurse -Include Share -Force**

Notes

The <*SourcePath*> value specifies the local path on the source server that contained the shared folder before the drive was migrated. Shared folder information is not stored on the data drive, so do not be concerned that the drive no longer resides on the source server.

The <*DestinationPath*> value specifies the local path on the destination server that contains folders that were previously shared on the source server. Unless the root drive letter or the folder structure has been changed on the migrated drive, the <*SourcePath*> and <*DestinationPath*> values should be the same.

During shared folder migration, permissions for local users and groups and domain users and groups are migrated—no manual remapping is required.

LanMan Server service automatically restarts on the destination server, and the shared folders migrate.

DFS Replication migration

If you physically migrated data, clean-up the DFS Replication configuration state, which is stored on the migrated volume.

1. To clean up volumes (for each physically migrated volume)
   a. Navigate to the path <volume>\System Volume Information.

      Note
      This is a hidden system folder. To view this folder: in File Explorer, click **View**, and then select the **Hidden Items** check box. Also ensure that local administrators are granted **Full Control** of the folder.

   b. Delete the **DFSR** folder and all content in the folder.
   c. Revert any security permissions modifications that you made to perform the migration process.
   d. Repeat this process for all physically migrated volumes.
2. To clean up replicated folders (for replicated folders on physically migrated volumes)
   a. Navigate to the root of a replicated folder.
   b. Delete the **DfsrPrivate** folder and all subfolders.
   c. If the staging folder for the replicated folder is not located in the default location, remove the staging folder and all content in the staging folder.

      Note
      The default location for the staging folder is in the DfsrPrivate folder, and this step is not required if the path is at the default location.

   d. If the **Conflict and Deleted** folder for the replicated folder is not located in the default location, remove the **Conflict and Deleted** folder and all content in the **Conflict and Deleted** folder.

      Note
      The default location for the **Conflict and Deleted** folder is in the **DfsrPrivate** folder, and this step is not required if the path is at the default location.

Use the inventoried information that you collected for the source server to detect all replication groups to which the source server belongs. Add the destination server as a member server to all these replication groups.

Migrate the source server identity

You need to rename the source server and migrate its previous identity to the destination server. You might also need to migrate the source server IP address to the destination server.

Rename the source server

Rename the source server to a temporary name.

Migrate IP address

When a static IP address is used on the source server, it is recommended that the IP address be migrated from the source server to the destination server. This is because client computers locally cache the IP address that is associated with a server name. Client computers will still attempt to access the source server even if it has been renamed.

When the server IP address is not migrated, you must stop the LanMan Server service on the source server. This is done to prevent users from accessing shared folders on the source server after they have been migrated to the destination server. Open a Windows PowerShell session with elevated user rights, and then run the following cmdlet:

```
Stop-Service LanmanServer -Force
```

For more information about IP address migration, see IP Configuration Migration Guide.

Rename destination server

Rename the destination server to the name that was originally used for the source server.

Export Remote VSS settings

Follow the procedure in this section to migrate Remote VSS settings from Windows Server 2012 R2 or Windows Server 2012.

To migrate Remote VSS from Windows Server 2012 R2 or Windows Server 2012, you must first export the remote VSS settings using the configuration information that is included in the registry and in Group Policy. There are two configuration Group Policy settings for Remote VSS:

- Computer Policy->Administrator Templates->System->File Share Shadow Copy Provider
- Computer Policy->Administrator Templates->System->File Share Shadow Copy Agent

You can configure these settings using either local or domain-based Group Policy. It is recommended that you use a domain-based policy setting because it does not require migration steps—you simply ensure that the policy setting applies to the new destination computer. If you are using a local policy setting, you must document the current settings for these two policy settings by running gpedit.msc. For the remaining policy settings, export the following registry key (using reg.exe from a command prompt with Administrative privileges), and then copy the rvss.reg file to the destination server:

```
Reg.exe export "HKLM\SYSTEM\CurrentControlSet\Services\fssagent\Settings" rvss.reg
```

If you migrated the data by copying it

Follow this procedure to add a replication connection between the source server and the destination server for each replication group on the source server:

1. In Server Manager, click **Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, select **Add Replication Groups to Display**, enter the name of the source, and then click **Show Replication Groups**. Select all of the replication groups that are displayed, and then click **OK**.
3. For each replication group, do the following:
    a. Click the replication group, and then click **New Member**. The **New Member Wizard** appears. Follow the instructions in the wizard to add the destination server to the replication group by using the information from row #2 in the DFS Replication data collection worksheet (File and Storage Services: Appendix B: Migration Data Collection Worksheets).
    b. In the console tree, under the **Replication** node, right-click the replication group to which you just added the destination server, and then click **New Connection**.
    c. Specify the source server and destination server as sending and receiving members, and specify a schedule so that the connection is always enabled. At this point, the replication is one-way.
    d. Select **Create a second connection in the opposite direction** to create a second connection for two-way replication between the sending and receiving members.

If you migrated the data by physically moving it

Follow this procedure to add a replication connection between the destination server and the closest server to the destination server other than the source server:

1. In Server Manager, click **Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, select **Add Replication Groups to Display**, enter the name of the source, and then click **Show Replication Groups**. Select all of the replication groups that are displayed, and then click **OK**.
3. For each replication group:
    a. Click the replication group, and then click **New Member**. The **New Member Wizard** appears. Follow the instructions in the wizard to add the destination server to the replication group by using the information from row #2 in the DFS Replication data collection worksheet (File and Storage Services: Appendix B: Migration Data Collection Worksheets).
    b. In the console tree, under the **Replication** node, right-click the replication group to which you just added the destination server, and then click **New Connection**.
    c. Specify the destination server as the sending member, and then specify any other server except the source server as the receiving member. Specify the schedule to use for the connection. It is recommended that you select a server that has a good network connection to the destination server as the receiving member.
    d. Select **Create a second connection in the opposite direction** to create a connection for two-way replication between the sending and receiving members.

Notes

The folder does not begin to replicate immediately. The new DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory Domain Services (AD DS) replication latency and the polling interval (60 minutes) on each member. The **dfsrdiag /pollad** command can be used to force DFS Replication on the source server and destination server to poll AD DS and retrieve the latest configuration information instead of waiting for the next normal polling interval which could be up to 60 minutes.

After DFS Replication on the destination server polls AD DS, it begins to replicate the folders that it configured, and it performs an initial synchronization. Event ID 4102 (MSG_EVENT_DFSR_CS_INITIAL_SYNC_NEEDED) is registered in the event log on the destination server for each replicated folder.

During initial sync, DFS Replication downloads all files in the replicated folders from the source server and builds up a local copy of the database per volume. This process can be time consuming. It is possible to speed up the initial sync by seeding the data from the source server onto the destination server (from the backup that was taken prior to commencing migration).

When the initial sync completes, event ID 4104 (MSG_EVENT_DFSR_CS_INITIAL_SYNC_COMPLETED) is registered for each replicated folder on the destination server. Monitor each replicated folder on the destination server and check to ensure that all of them have completed the initial sync.

Import settings to the destination server

Follow the procedures in this section to import settings to the destination server.

Note

If the source server is not running Windows Server 2012 or Windows Server 2008 R2, the first procedure in this section does not apply. (This procedure is used to migrate the seed value that is used by BranchCache for the Network Files component, and it enables data that is stored in BranchCache on the source server to be used after it is migrated to the destination server. For information about how to migrate a BranchCache host server, see the BranchCache Migration Guide.

To set up BranchCache for Network Files migration on the destination server

1. On the destination server, open a Windows PowerShell session with elevated user rights.
2. Type the following command, where *storepath* is the available path that contains the Svrmig.mig file, and then press Enter.

```
Import-SmigServerSetting -featureid BranchCache -Path
<storepath\BranchCache> -Force -Verbose
```

Group Policy or local policy specific to server message block and Offline Files

Use a Group Policy Object or a local policy setting on the destination server to change the settings to the same values as the source server. These settings are recorded in the SMB and BranchCache data collection worksheets in File and Storage Services: Appendix B: Migration Data Collection Worksheets.

To import SMB settings

1. Do one of the following:
   - If the policy settings are set by using Group Policy Objects, use the Group Policy editing tools to apply appropriate policy settings to the destination server.
   - If the policy settings are set by using a local policy setting, do the following:
     i. On the destination server, open the Local Group Policy Editor snap-in.
     ii. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
2. Use a Group Policy Object or a local policy setting to set the following settings to the same values as noted in Export settings. Set the destination server settings to the same values as were noted on the source server for the following **Microsoft network server** settings:
   - Microsoft network server: Amount of idle time required before suspending session
   - Microsoft network server: Attempt S4USelf to obtain claim information
   - Microsoft network server: Digitally sign communications (always)
   - Microsoft network server: Digitally sign communications (if client agrees)
   - Microsoft network server: Disconnect clients when logon hours expire

   Note
   For any setting that is controlled by Group Policy, you must have the same Group Policy Object apply to the destination server, or you can set the local policy of the destination server to get the same behavior.

   On destination servers that are running the Server Core installation, run the **secedit** command to change local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

Note

The following procedure applies only if the source server is Windows Server 2012 or Windows Server 2008 R2.

To import Offline Files settings

1. Do one of the following:
   - If the policy settings are set by using Group Policy, use the Group Policy editing tools to apply appropriate policy settings to the destination server.
   - If the policy settings are set by using local policy, do the following:
     i. On the destination server, open the Local Group Policy Editor snap-in.
     ii. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Administrative Templates**, click **Network**, and then click **LanMan Server**.
2. Use a Group Policy Object or a local policy setting to set the destination server policy settings to the same values as the source server policy settings for **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

   On destination servers that are running the Server Core installation, run the **secedit** command to change local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

DFS Namespace configuration

Complete the configuration of namespaces on the destination server. The procedure you use depends on whether you want a stand-alone or a domain-based namespace.

- Stand-alone namespaces
- Domain-based namespaces with more than one namespace server
- Domain-based namespaces with one namespace server

Stand-alone namespaces

If you want a stand-alone namespace, you must first create the namespace on the destination server. You can do this by using DFS Management, or the **DFSUtil.exe** command-line utility.

To create the namespace on the destination server

1. Do one of the following:
   - On the destination server, open DFS Management, and create the namespace by using the same name as on the source server.
   - On the destination server, in a Command Prompt window opened with elevated user rights, type the following, and then press Enter.

     ```
     Dfsutil.exe root addstd <\\DestinationServer\Namespace>
     ```

To import a namespace configuration from the export file

1. On the destination server, in a Command Prompt window opened with elevated user rights, type the following (in which *filename* represents the file name into which you exported namespace settings from the source server in To export the namespace configuration to an export file), and then press Enter.

```
Dfsutil.exe root import set <filename>
<\\DestinationServer\Namespace>
```

Domain-based namespaces with more than one namespace server

If you have more than one domain-based namespace server, you can add namespace servers to your destination server by using DFS Management or the **DFSUtil.exe** command-line utility.

To use DFS Management

1. Select the namespace being migrated in the left pane.
2. Click the **Namespace servers** tab in the right pane.
3. Select **Add Namespace Server**.
4. In the dialog box that opens, type the name of the destination server, and then click **OK**.

The destination server is added to the namespace.

To use DFSUtil.exe

1. On the destination server, open a Command Prompt window.
2. Type the following command, and then press Enter.

```
DFSUtil.exe target add <\\DestinationServer\Namespace>
```

Domain-based namespaces with one namespace server

This section applies only if a temporary server was not added to the namespace. If you added a temporary server to the namespace as part of your export process, see Domain-based namespaces with more than one namespace server.

To create the namespace on the destination server

1. Do one of the following:
   a. In DFS Management on the destination server, create the namespace with the same name that was used on the source server.
   b. Type the following command at a command prompt, and then press Enter.

```
Dfsutil.exe root adddom <\\DestinationServer\Namespace>
```

To import a namespace configuration from the export file

1. On the destination server, open a Command Prompt window.
2. Type the following command (in which *Filename* represents the export file names you

created in To export namespace settings). Run this command for each of the namespaces for which the source server was a namespace server.

```
DFSUtil.exe root import set <Filename>
\\DestinationServer\Namespace
```

Note
For each namespace, there must be a file name from which settings are imported.

To manually reset delegation permissions on the namespace

1. On the destination server, open DFS Management.
2. Set the permissions that you inventoried in DFS Namespace configuration. When complete, close DFS Management.

If any advanced registry keys were configured on *SourceServer*, use **DFSUtil.exe** to configure *DestinationServer* to have the same registry key settings. Run the following commands on the destination server to set the advanced registry keys.

To set advanced registry keys

1. On the destination server, open a Command Prompt window.
2. Run the following commands to set the advanced registry keys by using **DFSUtil.exe**.

```
DFSUtil.exe server registry DfsDnsConfig set
<DestinationServer>
DFSUtil.exe server registry LdapTimeoutValue set <Value>
<DestinationServer>
DFSUtil.exe server registry SyncInterval set <Value>
<DestinationServer>
```

File Server Resource Manager configuration on the destination server

When you are migrating File Server Resource Manager, remember to use the same drive letters for the destination server volumes as for the source server. This is required because File Server Resource Manager migration requires that the drive letter remains the same.

1. Stop the File Server Resource Manager and File Server Storage Reports Manager services.

   Open a Windows PowerShell session with elevated user rights, and then run the following command:

   ```
   Stop-Service -name "srmsvc","srmreports"
   ```

2. Type the following in the Windows PowerShell session, and then press Enter.

   ```
   Import-SmigServerSetting -FeatureID FS-Resource-Manager -Path
   <storepath\FSRM> -Force
   ```

Notes

If the Windows features that you are migrating have not been installed on the destination server, the **Import-SmigServerSetting** cmdlet installs them as part of the import process, along with any Windows features that they depend on. Some Windows features might require that you restart the destination server to complete the installation. After restarting the computer, you must run the cmdlet again with the **-Force** parameter to complete the import operation.

Importing FSRM settings to the destination server replaces any global FSRM configuration information that is already on the destination server.

3. Set the configuration files for each volume.

   Type the following commands in a Windows PowerShell session, and then press Enter.

   Note

   Running the following commands on a clean computer returns an error message. It is safe to ignore this error message.

   a. Type the following command to stop the file screen driver:

      ```
      fltmc detach datascrn <VolumeLetter>:
      ```

   b. Type the following command to stop the quota driver:

      ```
      fltmc detach quota <VolumeLetter>:
      ```

   c. Add administrator Write permissions to the "<VolumeLetter>:\System Volume information\SRM" folder and the following subfolders:

      - takeown /F "<VolumeLetter>:\System Volume Information" /A /R /D Y
      - cacls "<VolumeLetter>:\System Volume Information" /T /E /G Administrators:F
      - attrib -S -H "<VolumeLetter>:\System Volume Information*" /S /D

   d. Copy the following files from the external storage to the **SRM** folder:

      - Quota.xml
      - Quota.md
      - Datascrn.md
      - DataScreenDatabase.xml

   e. Type the following command to start the file screen driver:

      ```
      fltmc attach datascrn <VolumeLetter>:
      ```

   f. Type the following command to start the quota driver:

      ```
      fltmc attach quota <VolumeLetter>:
      ```

4. Restart the File Server Resource Manager and File Server Storage Reports Manager services.

   Type the following command in a Windows PowerShell session, and then press Enter.

   ```
   Start-Service -name "srmsvc","srmreports"
   ```

5. Configure scheduled reports and file management tasks.

   For each scheduled report, you need to create a scheduled task on the destination server.

   Note
   File Server Resource ManagerReports and classification rule configurations are dependent on the drive letters and mount points. Any drives or mount points on the source server that are used by report or classification rule configurations must be available on the destination server or the configurations must be updated during import.

   After you have an XML file for each task, copy them to the destination server and run the following command in a Windows PowerShell session on the destination server for each task:

   ```
   schtasks /create /xml:"TASKNAME.xml" /tn:"TASKNAME"
   ```

6. Import the classification schedule. The classification schedule requires a scheduled task on the destination server.

   ```
   schtasks /create /xml:"classification.xml"
   /tn:"FsrmAutoClassification{c94c42c4-08d5-473d-8b2d-
   98ea77d55acd}"
   ```

   Note that *classification.xml* is the name of the XML file that was exported from the target server.

Shadow Copies of Shared Folders

Apply the same settings from the source server to the corresponding volumes on the destination server.

To migrate shadow copy settings for Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003

1. To configure shadow copies, right-click each volume on the destination server that had shadow copies configured on the source server, right-click the volume, and then click **Configure Shadow Copies**.
2. Click **Settings** and verify that the location and size of shadow copy storage matches the settings from the source server.
3. Click **Schedule** and verify that the details for the snapshot creation task match the settings from the source server.

To migrate shadow copy settings for a Server Core installation

1. Log on to the destination server that is remotely running the Server Core installation by doing the following:
   a. In Server Manager, click **Tools**, and then click **Computer Management**.
   b. In the **Computer Management** snap-in tree pane, right-click the top node, and then click **Connect to another computer**.

2. Enter the computer name, and then click **OK**.
3. Expand **System Tools**, right-click **Shared Folders**, click the **All Tasks** tab, and then click **Configure Shadow Copies**.
4. For each volume on the destination server that had shadow copies configured on the source server, right-click the volume, click **Configure Shadow Copies**, click **Settings**, and verify that the location and size of shadow copy storage match the settings from the source server.
5. Click **Schedule**, and verify that these details for the snapshot creation task match the settings from the source server.

Deduplication

Use the following section to migrate Deduplication.

Migrating Deduplication from Windows Server 2012 to Windows Server 2012

All configuration information needed for migration is included on the deduplicated volume.

If a disk is physically moved, or if a deduplicated volume is restored from a backup onto a different Windows Server 2012 computer, install the Deduplication role service using Server Manager on the new computer. If the Deduplication role service is not installed on the new server, only normal nondeduplicated files will be accessible. After a volume has been mounted, the deduplication filter will detect that the volume is deduplicated and will redirect input/output requests appropriately.

Note

Any previous custom deduplication job schedules that were created using Task Scheduler must be created again on the new computer using Task Scheduler.

Migrating SIS from Windows Storage Server 2008 to Windows Server 2012

Volumes that have been created and optimized using the down-level Windows Storage Server version of deduplication, Single Instance Storage (SIS), should not be enabled for data deduplication. Microsoft recommends that SIS be removed from the volume by using SISAdmin.exe within Windows Storage Server to remove SIS or by copying the data to a different volume that is not running SIS prior to migrating the volume.

Windows Server 2012 supports reading and writing to SIS-controlled volumes, but you cannot continue to SIS files using Windows Server 2012. You can install the SIS filter driver on Windows Server 2012 by installing the SIS-Limited feature using the following command syntax:

```
dism /online /enable-feature:SIS-Limited
```

The SIS filter driver can be loaded so that you can read SIS-controlled volumes and the data can be copied to a non-SIS controlled volume so that data deduplication can be installed on the volume. Note that Windows Server 2012 does not support sisadmin.exe and cannot be used to remove SIS from a volume.

1. You should remove SIS from your volumes before installing the Windows Server 2012 data deduplication feature. (This process is also known as un-SIS.)
2. Do not restore SIS links from a backup to a Windows Server 2012 deduplicated volume.
3. Restoring SIS volumes to Windows Server 2012 is supported if you load the SIS-Limited filter.

Migrating SIS volumes

You have several options when it comes to migrating Windows Storage Server 2008 volumes to Windows Server 2012 to take advantage of the new Data Deduplication feature.

You can migrate your existing SIS-installed Windows Storage Server 2008 volumes to Windows Server 2012; however, migration is not automatic. SIS and data deduplication are mutually-exclusive technologies.

Caution

You will need to open the volumes in Windows Storage Server 2008 first, un-SIS them, and then uninstall SIS before migrating to Windows Server 2012 as described in the procedures below.

To unSIS a Windows Storage Server 2008 R2 or 2008 SIS volume, type **sisadmin.exe [/m <server>] [/u <volumes>]** where:

1. **/m <server>** shifts the focus of the command line to a remote server. If the */m* option is not specified, the command line is applied to the local server. <*server*> can be expressed as a host name, fully qualified domain name (FQDN), or IP address.
2. **/u <volumes>** is used to un-SIS a volume (that is, to restore all file copies and remove reparse points).

For each command option that uses <*volumes*> as a parameter, <*volumes*> represents a space-delimited list of volume names (for example, d:, e:, f:, and g:).

To unSIS or remove SIS entirely from the F: volume of a remote server using the IP address of the server, you might use the following command: **sisadmin.exe /m 192.168.1.50 /u F:**

Import Remote VSS settings

Follow the procedure in this section to migrate Remote VSS settings from Windows Server 2012 R2 or Windows Server 2012.

To finish migrating Remote VSS from Windows Server 2012 R2 or Windows Server 2012, import the remote VSS settings using the configuration information that is included in the registry and in Group Policy. There are two configuration Group Policy settings for Remote VSS:

- Computer Policy->Administrator Templates->System->File Share Shadow Copy Provider
- Computer Policy->Administrator Templates->System->File Share Shadow Copy Agent

You can configure these policy settings using either local or domain-based Group Policy. It is recommended that you use a domain-based policy setting because it does not require migration steps—you simply ensure that the policy setting applies to the new destination computer. If you

are using a local policy setting, open gpedit.msc and recreate the policy settings that you documented in Export Remote VSS settings.

For the remaining policy settings, export the registry key that you previously exported by using reg.exe from a command prompt with Administrative privileges:

For the remaining policy settings, import the registry key by using reg.exe from a command prompt with Administrative privileges:

```
Reg.exe import rvss.reg
```

See also

- Migrate File and Storage Services to Windows Server 2012 R2
- File and Storage Services: Prepare to Migrate
- File and Storage Services: Verify the Migration
- File and Storage Services: Migrate an iSCSI Software Target
- File and Storage Services: Migrate Network File System
- File and Storage Services: Post-Migration Tasks
- File and Storage Services: Appendix A: Optional Procedures
- File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Verify the Migration

To verify that the migration was successful, follow the appropriate verification steps based on the File and Storage Services role services that have been migrated.

The following overview describes the steps to verify the migration.

Verify the File Services migration

Perform the following tasks to verify the File and Storage Services role migration.

- Verify the File Services migration (only if running Windows Server 2012 or Windows Server 2008 R2)
- Verify migration of local users and groups
- Verify data and shared folder migration
- Verify the migration of DFS Namespaces
- Verify the configuration on other computers
- Verify the File Server Resource Manager migration

Verify migration of BranchCache for Network File Services server key

Perform this step only if the source server is running Windows Server 2012 or Windows Server 2008 R2:

Verify that the server key was migrated correctly by checking the key value, and ensure that the key values are identical on the source server and destination server, as shown in the following example:

```
Key: HKLM\Software\Microsoft\WindowsNT\CurrentVersion\PeerDist\SecurityManager\Restricted
Value: Seed
```

Verify migration of local users and groups

Check that all the local users and groups you expected to migrate are present on the destination server by comparing the list of users and groups on the Local Users and Groups snap-in on the source server with the list on the destination server.

To open the Local Users and Groups

1. In Server Manager, click **Tools**, and then click **Computer Management**.

Alternatively, you can compare the list of users and groups on the source server and destination server by typing **net** commands in a Command Prompt window.

- To get the list of all local users and save it in a text file, type the following command:

  ```
  net user > localusers.txt
  ```

- To get the list of all local groups and save it in a text file, type the following command:

  ```
  net localgroup > localgroups.txt
  ```

Verify data and shared folder migration

1. Check that all the data you expected to migrate is present at the correct location on the destination server and that the data has the correct permissions associated with it.

   To list files and folders with their permissions, type the following command in a Command Prompt window or in a Windows PowerShell session opened with elevated user rights:

   ```
   icacls <path>
   ```

2. Verify that all the expected shared folders have migrated and that they have the correct permissions associated with them. To list all shared folders and their permissions, type the following command in a Windows PowerShell session opened with elevated user rights:

   ```
   gwmi win32_share | %{net share $_.name}
   ```

Verify the migration of DFS Namespaces

The procedure that you use to verify the migration of DFS Namespaces depends on whether your namespaces are stand-alone or domain-based.

To verify the migration of a stand-alone namespace

1. Open DFS Management on the destination server.
2. Right-click **Namespaces**, or click the **Action** menu.
3. Click **Add Namespaces to Display**.
4. Type the name of the destination server, and then click the **Show Namespaces** button. Select the namespace that you migrated, and then click **OK**.
5. In the namespaces tree, click the namespace that you migrated.
6. Click the **Namespace** tab, and check that all the namespace links are present.
7. Click the **Namespace server** tab, and check that the destination server is listed.
8. Right-click the destination server name, and then click **Open in Windows Explorer**. All namespace links should be visible in the new window.
9. Using DFSUtil.exe on the destination server, type the following command for each stand-alone namespace:

   ```
   Dfsutil.exe root \\DestinationServer\Namespace
   ```

   The information displayed should contain the destination server and all the namespace links.

To verify the migration of a domain-based namespace

1. Open DFS Management, and then right-click **Namespaces** or click the **Action** menu.
2. Click **Add Namespaces to Display**.
3. Type the name of the domain where the namespace is located, and then click the **Show Namespaces** button. Select the namespace that you migrated, and click **OK**.
4. In the namespaces tree, click the namespace that you migrated.
5. Click the **Namespace** tab, and check that all the namespace links are present.
6. Click the **Namespace** server tab, and check that all the namespace servers are listed.
7. Right-click the destination server name, and then click **Open in Windows Explorer**. All namespace links should be visible in the new window.
8. Using DFSUtil.exe on the destination server, type the following command in a Command Prompt window, where *\\Domain\Namespace* is the name of the appropriate domain and namespace that you migrated.

   ```
   Dfsutil.exe root <\\Domain\Namespace>
   ```

   The information displayed should contain all namespace servers and namespace links.

Verify the configuration on other computers

To verify that File and Storage Services migration completed successfully on other computers, you must test the configuration on the client computers in your enterprise.

To verify DFS Namespaces on a client computer

1. Log on to a client computer with the credentials of a user who has access to the migrated namespace.
2. Verify that you can access the namespace by using File Explorer, a command prompt window, or another application, by entering the same name that you used before the migration.

Verify the File Server Resource Manager migration

Follow these steps to ensure that File Server Resource Manager migrated:

1. If any custom actions are configured for quota notification or file management tasks, the user should ensure that the folders that contain the executable files configured for the actions and the working folders have the correct access control lists.
2. As a best practice, ensure that all email message text for notifications, reports, and other purposes migrated correctly.
3. Administrators should send a test email message through the File Server Resource Manager console to verify that the Simple Mail Transfer Protocol (SMTP) server is configured correctly for the destination server.
4. Ensure that expiration folders that are used by File Management Tasks are reachable on the destination server.
5. Ensure that executable files that are used by custom actions (such as quota notifications and file management tasks) are accessible or executable on the destination server.

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Migrate an iSCSI Software Target

This section describes how to migrate Microsoft iSCSI Software Target 3.2 or 3.3 settings and data from an existing Windows Storage Server 2008 R2 or Windows Storage Server 2008 computer to a destination server that is running the iSCSI Target Server role service that is included with Windows Server 2012 R2 or Windows Server 2012 and Windows Storage Server 2012.

The naming for iSCSI Software Target has changed. To reduce the potential for confusion, in the context of this document, any naming that refers to “iSCSI Software Target”, refers to prior product versions installed on Windows Storage Server 2008 R2 and Windows Storage Server 2008, which are source servers. By contrast, any naming that refers to “iSCSI Target Server” refers to the new role service included with Windows Server 2012 R2, Windows Server 2012, and Windows Storage Server 2012, which are destination servers.

Note

This section contains only iSCSI-specific migration information. For generic information, such as the use of Windows Server Migration Tools, refer to File and Storage Services: Prepare to Migrate.

Supported migration scenarios

This section details both supported and unsupported migration scenarios.

Supported operating systems

The versions of operating systems that are listed are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

Migrations between physical operating systems and virtual operating systems are supported.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shared resources from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2012 R2 or Windows Server 2012 in the German system UI language.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Server 2008 R2	Windows Server 2012 R2 and Windows Storage Server 2012	x64-based
x64-based	Windows Storage Server 2008 R2	Windows Server 2012 R2 and Windows Storage Server 2012	x64-based
x64-based	Windows Server 2012	Windows Server 2012 R2 and Windows Storage Server 2012	x64-based
x64-based	Windows Storage Server 2012	Windows Server 2012 R2 and Windows	x64-based

		Storage Server 2012	

x64-based migrations are supported for Windows Storage Server 2012 R2 and Windows Server 2012 R2. All editions of Windows Storage Server 2008 R2 and Windows Server 2008 R2 are x64-based.

x86-based migrations are not supported because Windows Storage Server 2012 R2 is not offered in the x86 platform.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Supported role configurations

This migration guide is applicable to stand-alone and clustered configurations, with certain limitations.

The following general restrictions are applicable to all the supported configurations:

- Authentication settings for iSCSI initiators that use CHAP and Reverse CHAP settings are not automatically migrated.
- Snapshot storage settings for each virtual disk in the configuration are not automatically migrated.
- Configuration settings for virtual disks that are derived from snapshots are not automatically migrated.
- For clustered configurations, the migration process includes iSCSI target settings that are scoped to the virtual computer object, to a cluster node, or to the cluster node that owns the code cluster group.
- For clustered configurations, the migration of resource groups, network name resources, IP addresses, and cluster disks that are associated with resource groups is outside of the scope of this guide, and the migration needs to be performed independently as a preliminary step.
- iSCSI Naming Services (iSNS) settings for iSCSI Software Target are not automatically migrated.
- iSCSI target portal settings (such as IP addresses that are used by the iSCSI target service to listen for incoming network connections) are not automatically migrated
- The schedule for snapshots of virtual disks is not migrated.

The following configurations are supported:

- Migration from a stand-alone configuration to stand-alone configuration
- Migration from a clustered configuration to a stand-alone configuration (with the restrictions listed previously regarding the scope of the settings).
- Migration from a clustered configuration to a clustered configuration (with the restrictions listed previously regarding the scope of the settings).

Supported role services and features

iSCSI Target Server (as included with Windows Storage Server 2012 and Windows Server 2012 R2) does not have role dependencies or feature dependencies.

It is possible to install iSCSI Target Server with failover clustering, and this configuration is supported with the migration limitations listed previously.

Migrating multiple roles

If you are migrating one clustered configuration to a different clustered configuration, the Failover Clustering feature needs to be migrated or set up prior to migrating iSCSI target settings.

Migration scenarios that are not supported

The following migration scenarios are not supported:

- Migration from Windows Unified Storage Server 2003 R2.
- Migration from a stand-alone configuration to a clustered configuration. This migration is not supported because there is no default mechanism to associate target and virtual disk settings to resource groups without knowing how the file paths are mapped to the cluster disk and how IP addresses are mapped to resource groups.
- Snapshots of virtual disks are not automatically migrated. Snapshots are based on a snapshot of the volume that contains the virtual hard disk (VHD) file at the time the snapshot was taken. Their existence and implementation depends on the volume of the computer from which the migration process happens, and it cannot be replicated or exported.
- Snapshot storage settings for virtual disks are not automatically migrated. The snapshot storage settings (such as volume and maximum size per volume) are dependent on the hardware and software configuration of the computer to which the settings are being migrate, and they cannot automatically be migrated. For detailed information about how to manually migrate the snapshot storage settings, see Import the iSCSI Software Target settings in a stand-alone configuration.
- The configuration settings of the iSCSI target portal are not automatically migrated. This configuration is based on the IP addresses of the destination server, and those settings cannot be migrated outside the knowledge of the network configuration of the computer to which the settings are being migrate. For detailed information about how to manually configure the portal settings, see Configure the iSCSI Target Server portal.
- iSNS settings are not automatically migrated. The iSNS settings are based on the network infrastructure and configuration of the destination server, and those settings cannot be migrated outside the knowledge of the network configuration of the computer to which the settings are being migrated. For detailed information about how to manually configure iSNS settings, see Configure iSNS settings.
- Settings for virtual disks that are attached as local disks on the source server are not automatically migrated. The ability to attach a disk locally is expected to be a temporary operation that can be replicated if. For detailed information about how to configure settings for virtual disks that are to be attached as local disks, see Configure storage.

- The schedule for snapshots of virtual disks is not migrated. Those settings must be manually discovered and replicated from the source to the destination server.

Migration overview

This section describes the high-level migration process, which involves harvesting configuration settings from the source, moving the virtual disks from the source server to the destination server, and restoring the configuration settings.

Migration process

This section describes the high-level migration process.

Migration planning

The migration planning phase involves gathering information based on the following questions:

- Are the source server and destination server configured in a cluster?
- If the servers are configured in a cluster, what are the virtual computer objects or client access points that contain the iSCSI target resources?
- Is the storage system of the destination server capable and configured appropriately to host the virtual disks of the source server, and does it have appropriate space to store the volume snapshots?
- Are there any iSCSI initiators that have a critical dependency on iSCSI targets for the duration of the migration process (such as a computer that uses iSCSI boot nodes, or clusters that use shared storage)?
- Are there any IP address or portal settings that are unique to the source server that need to be accounted for (such as IP addresses that are known to the firmware of devices)?
- Are there any iSNS settings that need to be manually recorded and migrated?
- Are there any virtual disks attached as local disks that might need to be exposed?

Preparing to migrate

The preparation to migrate data from the source server to the destination server involves the following steps:

1. If the destination server will have a clustered configuration, install the Failover Clustering feature and form a cluster before performing the migration.
2. If the destination server will have a clustered configuration, create a number of cluster resource groups with client access points and cluster disks as appropriate to replicate the existing configuration. If possible, use the same resource group names for the source clusters and the destination clusters.
3. Install the iSCSI Target Server role service on the destination server.
4. Disconnect all the iSCSI initiators. This step is required to maintain consistent data on the virtual disks while they are being moved.
5. Run the Windows PowerShell script, iSCSITargetSettings.ps1, to capture the existing settings on the source server in an XML file. For a cluster, run the script on each node in the

cluster or on each virtual computer object, as appropriate for the scope of the planned migration.

The Windows PowerShell script displays the virtual disks that are eligible for migration and those that are not (for the snapshot-based reasons discussed previously).

Migration

The actual migration process includes the following steps:

1. Move the files for all the virtual disks that are eligible for migration from the source server to the destination server. If there are any file path changes, note the source to destination mapping.
2. In a cluster configuration, ensure that the destination path of the file copy is on a cluster disk and that the cluster disk has been assigned to a resource group. Note the resource group that owns the path.
3. If the file paths have changed between the source and the destination servers, open the settings .xml file in a text editor, and identify the **<MigrationDevicePath>** tags that need to be changed to reflect the new path.
4. In a cluster configuration, if the file path or the resource group name have changed between the source server and the destination server, open the settings .xml file in a text editor, and identify the **<MigrationResourceGroup>** tags that need to be changed to reflect the new resource group.
5. Run the Windows PowerShell script, iSCSITargetSettings.ps1, to import the settings to the destination server. In a cluster configuration, the destination server can be specified as a cluster node or as a virtual computer object. The cluster node or virtual computer object must be the owner of the resource group that is indicated in the settings .xml file.
6. If there are snapshot storage settings relevant to the new configuration, apply those settings manually.
7. If there are virtual disks that need to be attached as local disks, perform those actions.
8. If there are any iSNS settings that are relevant to the new configuration, apply those settings manually.
9. If there are any iSCSI target portal settings that are relevant to the new configuration, apply those settings manually.
10. If there are any iSCSI initiators that are configured to authenticate by using CHAP and Reverse CHAP, manually restore those settings.

Verification

The verification process for the migration involves the following steps:

1. Validate the iSCSI target portal settings by opening a Command Prompt window and typing **netstat.exe –nao | findstr 3260**. (This assumes that the default TCP port for the iSCSI protocol 3260 is used.) Alternatively, type **Get-WmiObject –Namespace root\wmi –Class WT_Portal** to cross-check the results.
2. Inspect the iSCSI Target Server configuration by using the Windows PowerShell cmdlet, **Get-IScsiServerTarget**
3. Inspect the iSCSI virtual disk configuration by using the Windows PowerShell cmdlet, **Get-IScsiVirtualDisk**

4. Validate the configuration for each iSCSI initiator that you expect to use with iSCSI Target Server by using the iscsicpl.exe UI tool or the iscsicli.exe command-line tool.

Impact of migration

The migration process does not impact or affect the source server. There are no resources or configuration settings that are altered or deleted as part of the migration process.

No servers in the enterprise, other than the destination servers, will be affected by the migration.

Client computers that are running as iSCSI initiators are expected to be explicitly disconnected during the migration to ensure data integrity. During the migration, the source server will be unavailable. When the migration process is complete, it is expected that the iSCSI initiators will log on to the destination server without any issues.

The downtime for the iSCSI initiators is expected to be proportionate to the time it takes to move the virtual disk files from the source server to the destination server, plus the time needed to restore the configuration settings and to establish the network identity of the destination server.

Permissions required for migration

Local Administrator permissions are required on the source and the destination server.

If the iSNS server has additional access control policies, permission to alter the iSNS settings are required as appropriate for the iSNS server.

To perform the migration process for the iSCSI initiators, permissions to log on and log off iSCSI sessions are required. For the iSCSI initiator, Local Administrator permissions are required.

For iSCSI initiators that are firmware based, such as a network interface with the option to boot from iSCSI, being at the actual console may be required to configure logon credentials or the network identity of the destination server if the authentication settings (CHAP and Reverse CHAP) have changed.

Estimated time duration

This section detail the various factors that impact how long a migration may take to complete.

Planning

The planning phase is expected to be influenced by the following factors:

- Stand-alone versus a cluster configuration. A cluster setup may require one to two hours to configure if all the validations are performed.
- Storage configuration. Understanding and configuring a storage array to host potentially huge files requires that you plan the spindle and volume configurations so that they use the tools that are provided by the storage subsystem vendor.
- Network identity. This planning involves understanding if the source server has specially or purposely configured IP addresses, if configuring Level-2 components (such as switches) is required, and if specific DNS or NetBIOS names need to be known to and cached by the iSCSI initiators.

Preparation

The preparation process involves understanding which settings (that are specific to the source server) cannot be automatically migrated, and gathering those settings. For each step in the preparation phase, the mechanism that is used to retrieve the settings depends on which step is applicable and which tool is used to recover those settings.

- Cluster resource group names and configuration. These settings can be gathered from the cluster administration tools and the user interfaces.
- iSCSI target portal configuration. These settings can be gathered by typing the following code at a command prompt: `PS > Get-WmiObject –Namespace root\wmi –Class WT_Portal`
- iSNS Server settings. These settings can be gathered by typing the following code at a command prompt: `PS > Get-WmiObject –Namespace root\wmi –Class WT_ISnsServer`
- CHAP and Reverse CHAP authentication settings. These settings cannot be automatically retrieved because the iSCSI target server does not offer a mechanism to retrieve passwords. These settings have been stored elsewhere in the enterprise, and they need to be retrieved independently.
- Locally mounted virtual disk settings.

Migration

The estimated time for the actual migration process is largely dominated by the time that it takes to move the virtual disk files from the source server to the destination server.

A network-based file copy, using a 1 GB link used at 50% for 1 TB of data, is estimated to take over five hours. Techniques that use a file transfer process involving external media, such as an External Serial Advanced Technology Attachment (eSATA) device, may take less time.

The execution of the Windows PowerShell import script is estimated to take few minutes for approximately 100 resources (with a combination of iSCSI target settings and virtual disk settings).

Verification

The estimated time for the verification is proportionate to the time it takes to reconnect or log on to the iSCSI initiators.

For each iSCSI initiator, the target portal needs to be reconfigured, credentials related to authentication settings must be entered (if required), and the sessions have to be logged on.

The estimated time is 5 to 15 minutes to verify each iSCSI initiator, depending on the process that is being used. iSCSI initiators can be verified through the iscsicpl.exe UI, through the iscsicli.exe command-line tool, or through ad hoc Windows Management Instrumentation (WMI)-based scripts).

See Also

Migrate File and Storage Services to Windows Server 2012 R2

Prepare to Migrate iSCSI Software Target

Migrate iSCSI Software Target

Verify the iSCSI Software Target Migration

Troubleshoot the iSCSI Software Target Migration

Roll Back a Failed iSCSI Software Target Migration

Prepare to Migrate iSCSI Software Target

This topic discusses the tasks that are necessary before you start the migration process. The first step is to install the Windows Server Migration Tools. For more information, see File and Storage Services: Prepare to Migrate.

Prepare the destination server

The destination server is a computer that is configured and shipped by an OEM with Windows Storage Server 2012 pre-installed, or that is running Windows Server 2012 R2.

iSCSI Target Server hardware requirements for the destination server are the following:

- The amount of free disk space on the destination server must be sufficient to host the iSCSI virtual disk from the source server with adequate room for the snapshot storage.
- For clustered configurations, the resource groups that are created in the destination server must have associated cluster disks with adequate free space to host the iSCSI virtual disk from the source server.
- The destination server must have one or more network interfaces to be utilized for the iSCSI network traffic.

Installing the Failover Cluster feature in Windows Server 2012 R2 or Windows Storage Server 2012 or is required if the source server was configured with failover clusters. For more information, see the Failover Clustering Overview.

Back up the source server

Before you start migration, as a best practice, it is recommended that you back up the source server. For more information, see Windows Server Backup.

Prepare the source server

The following tasks are performed on the source server.

Cluster resource group configuration

Use the following steps to obtain the cluster resource groups:

1. Gather the resource groups that have iSCSI Software Target resources by using the following Windows PowerShell commands:

```
PS > Import-Module FailoverClusters
```

```
PS > $iSCSITargetResources = Get-ClusterResource | Where-Object
{ ( $_.ResourceType.Name -eq "Host" ) –or ($_.ResourceType.Name
-eq "WTDisk") }
PS > $iSCSITargetResources
```

2. From the cluster resources obtained in the previous step, gather the cluster disk dependencies by using the following Windows PowerShell commands:

```
PS > $Dependencies = &{ $iSCSITargetResources | Get-
ClusterResourceDependency  }
PS > $Dependencies
```

If the source server is running Windows Storage Server 2008, the following steps can be followed to gather the equivalent information:

1. Gather the iSCSI Software Target resources, and then gather the groups by using the following Windows PowerShell commands:

```
PS > $iSCSITargetResources = Get-WmiObject -NameSpace
root\mscluster -Authentication PacketPrivacy -Class
MsCluster_Resource -Filter "Type = `"WTDisk`" or Type =
`"Host`""
PS > $iSCSITargetResources
PS > $Groups = &{foreach($res in $iSCSITargetResources) { Get-
WmiObject -NameSpace root\mscluster -Authentication
PacketPrivacy -Query "associators of {$($res.__RELPATH)}  WHERE
ResultClass = MSCluster_ResourceGroup" }}
PS > $Groups
```

2. From the cluster resources obtained in the previous step, gather the cluster disk dependencies by using the following Windows PowerShell commands:

```
PS > $Dependencies = &{foreach($res in $iSCSITargetResources) {
Get-WmiObject -NameSpace root\mscluster -Authentication
PacketPrivacy -Query "associators of {$($res.__RELPATH)}  WHERE
ResultClass = MSCluster_Resource ResultRole = Dependent" }}
PS > $Dependencies
```

The resource groups obtained in step 1 have network name resources and IP addresses that need to be migrated to the destination server.

For information about how to migrate these settings, see Migrate Roles and Features to Windows Server.

The cluster disk that you obtained in step 2 is the physical disk where the volumes reside that are hosting the iSCSI Software Target virtual disks.

To obtain the volumes from the cluster disk, use the following steps:

1. Obtain the disk signature of the cluster disk by using the following Windows PowerShell command:

```
PS > & cluster.exe res "<cluster resource name>" /priv
```

2. Obtain the **Win32_DiskDrive** object from the disk signature by using the following Windows PowerShell command:

```
PS > $DiskObj = Get-WmiObject -Namespace root\cimv2 -Class
Win32_DiskDrive -Filter "Signature = <disk signature>"
PS > $DiskObj
```

3. Obtain the **Win32_DiskDriveToDiskPartition** association by using the following Windows PowerShell command:

```
PS > $DiskToDiskPartition = Get-WmiObject -Namespace root\cimv2
-Class Win32_DiskDriveToDiskPartition | Where-Object {
$_.Antecedent -eq $DiskObj.__PATH }
PS > $DiskToDiskPartition
```

4. Obtain the **Win32_LogicalDiskToDiskPartition** association that points to the volume association by using the following Windows PowerShell command:

```
PS > Get-WmiObject -Namespace root\cimv2 -Class
Win32_LogicalDiskToPartition | Where-Object { $_.Antecedent -eq
$ DiskToDiskPartition.Dependent }
```

Steps 2–4 must be applied on the source server cluster node that currently owns the physical disk cluster resource.

iSCSI Target portal configuration

Use the following steps to obtain the portal associations:

1. Gather the configured portals association for the iSCSI target portal by using the following Windows PowerShell command:

```
PS> Get-WmiObject -Namespace root\wmi -Class WT_portal | Format-
List -Property Address,Listen,Port
```

2. The IP addresses that have the **Listen** state set to **True** are the IP addresses that an iSCSI initiator can use to reach the server. For more information about migrating the IP addresses, see Migrate Roles and Features to Windows Server.

iSNS configuration

Gather the configured iSCSI Naming Services (iSNS) association for the server by using the following Windows PowerShell command:

```
PS> Get-WmiObject -Namespace root\wmi -Class WT_ISnsServer | Format-List -Property
ServerName
```

The server names that are listed need to be added to the list of iSNS servers that can be used to retrieve information about the iSCSI initiators in the enterprise.

CHAP and Reverse CHAP configuration

Gather the CHAPUserName and ReverseCHAPUserName association for the servers that are configured with CHAP and Reverse CHAP by using the following Windows PowerShell command:

```
PS > Get-WmiObject -Namespace root\wmi -Class WT_Host | Where-Object { ( $_.EnableCHAP )
-or ( $_.EnableReverseCHAP )  } | Format-List -Property
Hostname,CHAPUserName,ReverseCHAPUserName
```

The passwords that are used in conjunction with the credentials listed previously cannot be retrieved, and they must be known through other mechanisms.

Snapshot storage configuration

The snapshot storage configuration can be obtained by using the following Windows PowerShell command:

```
PS > & vssadmin.exe list shadowstorage
```

This command shows the volume snapshot shadow storage configuration for the entire source server. Not all the volumes listed may be relevant to the current iSCSI Software Target server configuration.

For the volumes that are relevant (that is, the volumes that host iSCSI virtual disks), the associated shadow storage volume is listed, in addition to the amount of disk space used with the maximum amount of configured space.

Disconnect the iSCSI initiators

Follow the instruction in the following section to disconnect the iSCSI initiators: Prepare other computers in the enterprise.

Capture the existing settings: stand-alone configuration

All of the settings on the iSCSI Software Target source server that are not hardware configuration specific and are not dependent on an IP address and the network identity of the server can be captured with the following Windows PowerShell commands:

Windows Server 2008 R2 and Windows Server 2008 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> Export-IscsiTargetServerConfiguration –FileName <settings XML file>
```

Windows Server 2012 R2 file path:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerShell\V1.0\Modules\IscsiTarget"

PS> Export-IscsiTargetServerConfiguration –FileName <settings XML file>
```

If the procedure is performed on a source server that is running iSCSI Target 3.3 from a destination server that is prepared as illustrated in the previous sections, the settings can be captured using the following Windows PowerShell commands:

Windows Server 2012 R2 file path:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget"

PS> Export-IscsiTargetServerConfiguration –FileName <settings XML file> -ComputerName
<source server computer name>
```

Windows Server 2008 R2 and Windows Server 2008 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> Export-IscsiTargetServerConfiguration –FileName <settings XML file> -ComputerName
<source server computer name>
```

At the end of the settings capture process, the Windows PowerShell script will display the set of VHD files that are eligible for migration. This list is needed for the destination server during migration.

Capture the existing settings: clustered configuration

Before capturing the iSCSI Software Target source server settings that are not hardware configuration specific, we recommend that all the resource groups with iSCSI target resources are moved to a single node in the cluster.

This can be accomplished by using the following Windows PowerShell commands. These commands assume that you previously followed the steps in Cluster resource group configuration.

```
PS > $iSCSITargetResources | Format-List -Property OwnerGroup

PS > foreach($Res in $iSCSITargetResources) { & cluster group  $Res.OwnerGroup
/moveto:$ENV:COMPUTERNAME }
```

After all the resource groups have been moved to a single node, the settings can be gathered by using the following Windows PowerShell commands:

Windows Server 2012 R2 file path:

```
PS > cd "$ENV:Programfiles\ISCSI Target"

PS> .\ Export-IscsiTargetServerConfiguration –FileName <settings XML file>
```

Windows Server 2008 R2 and Windows Server 2008 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> .\ Export-IscsiTargetServerConfiguration –FileName <settings XML file>
```

If the procedure is performed on a source server that is running iSCSI Target 3.2, the resources can be moved to a single node by using the following Windows PowerShell commands:

```
PS >  $Groups = &{foreach($res in $iSCSITargetResources) { Get-WmiObject -NameSpace
root\mscluster -Authentication PacketPrivacy -Query "associators of {$($res.__RELPATH)}
WHERE ResultClass = MSCluster_ResourceGroup" }}

PS > foreach($Group in $Groups) { & cluster group $Group.Name /moveto:<node name source
server> }
```

The iSCSI Target Server settings need to be gathered from a destination server that is prepared as illustrated in the previous sections. Run the script from a source server that is running iSCSI Target 3.3 by using the following Windows PowerShell command:

Windows Server 2012 R2 file path:

```
PS > cd "$ENV:Programfiles\ISCSI Target"

PS> .\ Export-IscsiTargetServerConfiguration –FileName <settings XML file> -ComputerName
<source server computer name>
```

Windows Server 2008 R2 and Windows Server 2008 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> .\ Export-IscsiTargetServerConfiguration –Export –FileName <settings XML file> -
ComputerName <source server computer name>
```

In the previous example, the source server computer name is the name of the node. At the end of the settings capture process, the Windows PowerShell script will display the set of VHD files that are eligible for migration. This list is needed for the destination server during migration.

Remove the network identity of the iSCSI Software Target computer

In a network with an iSCSI Software Target source computer, the identity of the server is known to iSCSI initiators in the form of NetBIOS names, fully qualified domain names (FQDN), or IP

addresses. When a server is being replaced, as part of planning, a strategy to replace the server network identity must be devised. Possible scenarios include:

- Transfer the NetBIOS and FQDNs to the destination server, and then assign new IP addresses to the destination server.
- Create new NetBIOS and FQDNs for the destination server, and then assign the existing IP addresses to the destination server.
- Create new NetBIOS and FQDNs for the destination server, and then assign new IP addresses to the destination server.

Each scenario requires potentially updating information in the DNS server, Active Directory, or DHCP server, according to the methodology that is used to assign IP addresses and network names to the servers in the enterprise.

The intent of this step is to ensure that upon completion of the migration steps, the iSCSI initiators are able to locate the destination server (either through explicit reconfiguration, or implicitly through the computer name or IP address re-assignment).

For more information, see Migrate Roles and Features to Windows Server.

Prepare the iSCSI initiator computers

The other computers in the enterprise that are affected by migration are the iSCSI initiators. The users of the computers that are acting as iSCSI initiators should be sent an outage notification. If the iSCSI Software Target is being used as a boot node for the iSCSI initiator computers, the computers may be completely unusable for the duration of the migration.

Capture the session information

The information regarding the active session for an iSCSI Software Target source server can be obtained by using the following Windows PowerShell command:

```
PS > & iscsicli.exe sessionlist
```

This information is needed to disconnect the session.

Disconnect the session

The session can be disconnected by using the following Windows PowerShell command:

```
PS > & iscsicli.exe LogoutTarget <session id>
```

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Migrate an iSCSI Software Target

Migrate iSCSI Software Target

Verify the iSCSI Software Target Migration

Troubleshoot the iSCSI Software Target Migration

Roll Back a Failed iSCSI Software Target Migration

Migrate iSCSI Software Target

This topic discusses the actual migration steps for iSCSI Software Target 3.2 or iSCSI Software Target 3.3 for both the stand-alone configuration and the clustered configuration:

Migrating iSCSI Software Target in a standalone configuration

The migration of iSCSI Target Server 3.3 or iSCSI Target Server 6.2 has equivalent steps, whether you are migrating from Windows Storage Server 2008 R2 or Windows Server 2012 or Windows Storage Server 2012 to Windows Server 2012 R2.

Establish network identity of the iSCSI Target Server computer

As part of the planning process, a strategy should have been devised regarding how iSCSI Target Server will be accessed from the network, based on key questions including but not limited to:

- Which computer name will be used?
- Which IP addresses on which subnet or set of network interfaces will be used?
- What relationship should be maintained between the IP addresses and computer name of the source server and the destination server? Will you keep the same addresses and names or create new ones?

Based on the desired final configuration, configuration changes are potentially needed in the following areas:

- The DHCP Server that might assign IP addresses to the destination iSCSI Target servers
- The DHCP Server that might assign IP addresses to the iSCSI initiators
- The DNS Server or Active Directory domain controller that might perform naming resolution services for the computers in the enterprise

Configure the iSCSI Target Server portal

After you have configured IP addresses for the network interfaces of the iSCSI Target Server computer, it is possible to verify the existing configuration by using the following Windows PowerShell command:

```
PS > $Portals = Get-WmiObject -Namespace root\wmi -Class WT_Portal | Where-Object {
$_.Listen }

PS > $Portals
```

The configuration of the access surface for iSCSI Target Server from the network can be restricted by disabling certain portals. For example, you can disable the fourth portal in the array returned in the previous step by using the following Windows PowerShell commands:

```
PS > $Portals[3].Listen = 0

PS > $Portals[3].Put()
```

The default port can also be changed from 3260 to any available TCP port on the destination server.

Configure iSNS settings

The iSNS servers that were configured for the source server can be configured for the destination server by using the following Windows PowerShell commands:

```
PS > $WT_ISnsServerClass =  Get-WmiObject -namespace root\wmi -class meta_class -filter
"__CLASS = 'WT_ISnsServer'"

PS > $WtiSNSInstanace = $WT_ISnsServerClass.CreateInstance()

PS > $WtiSNSInstanace.ServerName = "<iSNS computer name or IP>"

PS > $WtIsnsInstanace.Put()
```

Note
The set of iSNS servers that are configured for iSCSI Target Server was obtained during the preparation of the source server.

Configure storage

The destination server is expected to have sufficient storage space to host all of the virtual disks that are present on the source server.

The space does not need to be contiguous or in a single volume, and it does not need to replicate the same file system structure or volume mount point structure of the source server. The storage that is prepared to host the virtual disks must not be a nested volume, and it must be formatted with the NTFS file system.

Configure the Volume Shadow Copy Service

For the storage that was prepared in the previous step, it is appropriate to configure the Volume Shadow Copy Service, in case the default per-volume settings are not adequate. To inspect that current configuration, use the following Windows PowerShell command:

```
PS > & vssadmin.exe list shadowstorage
```

To modify the current configuration, use the following Windows PowerShell commands:

```
PS > & vssadmin.exe add ShadowStorage

PS > & vssadmin.exe delete ShadowStorage

PS > & vssadmin.exe resize ShadowStorage
```

Transfer the virtual disk

For all the files in the list of files that was captured in the source server preparation step, copy the files from the source server to the destination server. For more information, see Capture the existing settings: stand-alone configuration.

You will need the destination paths in the following steps. So if the absolute file path is different between the source server and the destination server, create a table with the mapping; for example:

Source path	Destination path
G:\WS08R2_OpsMgr2007_R2.vhd	H:\VHDS\WS08R2_OpsMgr2007_R2.vhd
F:\Dynamic_Spanned_GPT_2.vhd	D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd

Import the iSCSI Software Target settings in a stand-alone configuration

To import the iSCSI Software Target settings in a stand-alone configuration, you need the .xml file that you previously created. For more information, see Capture the existing settings: stand-alone configuration.

If there is no change in the absolute path of the virtual disk files, the import process can be performed by using the following Windows PowerShell commands:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget"

PS> .\ iSCSITargetSettings.PS1 -Import -FileName <settings XML file>
```

If the absolute path is different between the source server and the destination server, before you import the settings, the settings .xml file needs to be altered to reflect the new path.

Locate the records for the virtual disk, and alter the path in the <MigrationDevicePath> tag to reflect the absolute file path in the destination server:

```
<iSCSIVirtualDisk>
 <DevicePath>F:\Dynamic_Spanned_GPT_2.vhd</DevicePath>
  <MigrationDevicePath>D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd</MigrationDevicePath>
</iSCSIVirtualDisk>
```

After the XML has been altered to reflect the path in the destination server, you can import the settings by using the Windows PowerShell commands previously presented.

Configure shadow storage for the virtual disks

If certain virtual disks have shadow storage requirements that are different than the ones configured in the section Configure the Volume Shadow Copy Service, it is possible to alter the default or previously configured settings by using the following Windows PowerShell commands:

```
PS > $VirtDisk = Get-WmiObject -Namespace root\wmi -Class WT_Disk | Where-Object {
$_.DevicePath -eq '<full path of virtual disk>' }

PS > $VirtDisk.SnapshotStorageSizeInMB = <new size>

PS > $VirtDisk.Put()
```

Configure CHAP and Reverse CHAP

The authentication settings for iSCSI Target Server that are configured with CHAP and Reverse CHAP need to be manually configured. The list of targets that require CHAP and Reverse CHAP configuration is listed at the end of the import script, as described in the section Import the iSCSI Software Target settings in a standalone configuration.

To configure the CHAP and Reverse CHAP settings, use the following Windows PowerShell commands:

```
PS > $Target = Get-WmiObject -Namespace root\wmi -Class WT_Host | Where-Object {
$_.HostName -eq '<name of the target>' }

PS > $Target.EnableCHAP = 1

PS > $Target.CHAPUserName = "<user name>"

PS > $Target.CHAPSecret = "<CHAP Secret>"

PS  $Target.Put()
```

Migrating iSCSI Software Target in a failover cluster

The migration process for the failover cluster configuration is largely identical to migrating a stand-alone configuration, with the following differences:

- You will migrate resource groups instead of merely establishing the network identity of the server.
- You will use different Windows PowerShell commands to import the resource groups.

Migrate resource groups

This step replaces the “Establishing the network identity of iSCSI Target Server” step when you migrate a stand-alone configuration because the network identity of an iSCSI Target server in a cluster is given by the union of the client access point. (A client access point in the cluster is the logical union of a network name resource and one or more IP addresses that are assigned to the network name resource.)

Assuming the initial cluster resource groups and network names were configured in the default state, those can be recreated by using the following Windows PowerShell command:

```
PS > Add-ClusterServerRole – Name <resource group name>
```

Use this command for each of the resource groups that were in the original configuration. If the default client access point configuration does not match the initial configuration (for example, because the network name is bound to the incorrect cluster network, or the configuration required statically assigned IP addresses), modifications can be made. For more information, see Migrate Roles and Features to Windows Server.

After the resource groups have been created, clustered disks must be assigned to the network resources to match the configuration that you captured. For more information, see the Cluster resource group configuration section.

Import the iSCSI Software Target settings in a failover cluster

Follow these instructions to import settings in a failover cluster configuration. (This information differs from the how you would import settings in a stand-alone configuration.)

A prerequisite for the import phase is to have all of the resource groups that will host iSCSI Target Server resources owned by the same cluster node. Use the following Windows PowerShell command to validate the current ownership:

```
PS > Get-ClusterGroup
```

If there is no change in the absolute path of the virtual disk files, the import process can be performed by using the following commands:

```
PS > cd “$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget”

PS> .\iSCSITargetSettings.PS1 –Import –FileName <settings XML file>
```

If the absolute path is different between the source server and the destination server, before you import the settings, the settings .xml file needs to be altered to reflect the new path.Locate the records for the virtual disk, and alter the path in the **<MigrationDevicePath>** tag to reflect the absolute file path in the destination server:

```
<iSCSIVirtualDisk>

      <DevicePath>F:\Dynamic_Spanned_GPT_2.vhd</DevicePath>
```

```
    <MigrationDevicePath>D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd</MigrationDevicePath>
</iSCSIVirtualDisk>
```

After the XML has been altered to reflect the path in the destination server, you can import the settings by using the Windows PowerShell commands.

Migrate iSCSI Target Server Providers

This section provides details about migrating iSCSI Target Server Virtual Disk Service (VDS), Volume Shadow Copy Service (VSS), and SMI-S providers.

Migrate VDS and VSS hardware providers

If you are upgrading from Windows Server 2012 to Windows Server 2012 R2, the previous storage provider is automatically upgraded to Windows Server 2012 R2, and no additional action is required.

If you are upgrading from Windows Server 2008 R2 to Windows Server 2012 R2, you must first manually uninstall the currently installed iSCSI Target storage provider. Because iSCSI Target storage provider was installed separately from Windows Server 2008 R2, the provider cannot be automatically upgraded. When the iSCSI Target storage provider is uninstalled, do the following:

a. Upgrade the server to Windows Server 2012 R2.
b. Install the **iSCSI Target Storage Provider (VDS and VSS hardware providers)** role service on the upgraded server. You can do this using **Server Manager dashboard**.
c. The iSCSI VDS and VSS storage providers must be configured to run under the administrative credentials of the iSCSI Target Server. For more information, see iSCSI Target Block Storage, How To.

Migrate SMI-S providers

You must first manually uninstall the currently installed SMI-S provider for Windows Server 2012. Because the SMI-S provider was installed separately from Windows Server 2012, the provider cannot be automatically upgraded. When the SMI-S provider is uninstalled, do the following:

1. Upgrade the server to Windows Server 2012 R2. The SMI-S provider is automatically installed along with the iSCSI Target Server role service.
2. From any System Center Virtual Machine Manager (VMM) or SMI-S management client, unregister and reregister using the appropriate credentials. For information on configuring the SMI-S provider using VMM, see Configuring an SMI-S Provider for iSCSI Target Server. For information about configuring the SMI-S provider using the SMI-S management client, see Register-SmisProvider.

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Migrate an iSCSI Software Target

Prepare to Migrate iSCSI Software Target
Verify the iSCSI Software Target Migration
Troubleshoot the iSCSI Software Target Migration
Roll Back a Failed iSCSI Software Target Migration

Verify the iSCSI Software Target Migration

This topic discusses the steps you can use to verify that the migration successfully completed.

Verifying the destination server configuration

To verify that the destination server has been properly configured after migration, you can verify the listening endpoints and connectivity and run a scan with the Best Practices Analyzer.

Verify the listening endpoints

On the iSCSI Target destination server, you can validate that the target portals have been configured as desired by using the following Windows PowerShell command:

```
PS > & netstat.exe -nao | findstr 3260 | findstr LISTENING

  TCP    10.121.26.107:3260     0.0.0.0:0              LISTENING       1560

  TCP    10.121.26.126:3260     0.0.0.0:0              LISTENING       1560

  TCP    [2001:4898:0:fff:0:5efe:10.121.26.126]:3260  [::]:0                 LISTENING
1560

  TCP    [2001:4898:f0:1001:f063:8fc5:52e6:2310]:3260  [::]:0                 LISTENING
1560
```

The list of IP addresses and port pairs in the listening state needs to match the desired set of target portals.

Note
If ports other than the default 3260 are being used, the command needs to be altered to reflect the alternate IP ports.

Verify the basic connectivity

To validate that the iSCSI Target Server computer is reachable from other computers on the network, from a computer that has the Telnet Client feature installed, use the following Windows PowerShell command:

```
PS > telnet.exe <iSCSI Software Target machine name or IP> 3260
```

If there is a successful connection, Telnet Client will show a blinking cursor at the top of the window. Press any key to close Telnet Client.

Perform a Best Practices Analyzer scan

To verify that iSCSI Target Server is optimally configured on Windows Server 2012 or Windows Storage Server 2012 after migration, we recommend that you run a Best Practices Analyzer (BPA) scan on the role.

BPA is a server management tool that is available in Windows Server 2012. After the migration of iSCSI Target 3.3 is complete, BPA can help you ensure that your server is configured according to best practices. You can use the Server Manager console UI or Windows PowerShell to perform BPA scans and view results. For detailed information about how to scan your role and view results, see Run Best Practices Analyzer Scans and Manage Scan Results.

Verifying the configuration of iSCSI initiator computers

Validating the migration of iSCSI Software Target to the destination server includes ensuring that the iSCSI initiators can discover and fully access all features of the iSCSI protocol.

Verify that the iSCSI initiators can discover iSCSI Target Server

To verify that the iSCSI initiators can discover iSCSI Target Server, use the following Windows PowerShell commands:

```
PS > & iscsicli AddTargetPortal <ip-address> 3260
```

```
PS > & iscsicli.exe ListTargets
```

If the commands execute without errors, the initiator is capable of discovering the targets that are offered by the server

Verify that the iSCSI initiators can log on

The second step is to verify that the iSCSI initiators are able to log on to the iSCSI targets that are exposed by iSCSI Target Server. This can be accomplished by using the following Windows PowerShell command:

```
PS > & iscsicli.exe  LoginTarget <target IQN> T <ip address> 3260 Root\ISCSIPRT\0000_0 *
* * * * * * * * * * * * *
```

Note

If you are using CHAP and Reverse CHAP authentication, you may need to specify more parameters. For more information, consult the documentation in the iscsicli.exe.

If the command executes without errors, the iSCSI initiator has successfully logged on to the target, and the disks are exposed to iSCSI Target Server.

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Migrate an iSCSI Software Target

Prepare to Migrate iSCSI Software Target

Migrate iSCSI Software Target

Troubleshoot the iSCSI Software Target Migration

Roll Back a Failed iSCSI Software Target Migration

Troubleshoot the iSCSI Software Target Migration

Troubleshooting iSCSI Software Target migration issues involves first viewing the contents of the Windows Server Migration Tools deployment log and the result objects. For more information, see Locate the deployment log file and View the content of Windows Server Migration Tools result objects.

Understanding the messages from the iSCSI Target Migration tool

The iSCSI migration tool (iSCSITargetSettings.ps1) does not produce a log file, but it prints diagnostics messages on the console. These messages show the outcome of the operations that are being attempted and performed.

For example, the following message shows information about saved settings:

```
PS C:\Windows\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget>

.\iSCSITargetSettings.PS1 -Export -FileName $env:temp\test00000000.xml


Number of Target(s) saved in the Export settings: 4.

Target Names:

   test000

   test001

   test002

   test1111
```

```
Number of Virtual Disk(s) saved in the Export settings: 3.
Virtual Disk DevicePaths:
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd

Number of Virtual Disk(s) NOT saved in the Export settings: 0.
Virtual Disk DevicePaths:
```

The following message shows that not all the virtual disks are eligible for migration:

```
PS D:\Program Files\ISCSI Target> .\iSCSITargetSettings.PS1 -Export -FileName
$env:temp\test00000001.xml

Number of Target(s) saved in the Export settings: 4.
Target Names:
   test000
   test001
   test002
   test1111

Number of Virtual Disk(s) saved in the Export settings: 3.
Virtual Disk DevicePaths:
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd

Number of Virtual Disk(s) NOT saved in the Export settings: 1.
Virtual Disk DevicePaths:
    \\?\GLOBALROOT\Device\HarddiskVolumeShadowCopy{B6B3C77C-93CC-11DF-B3FE-
001CC0C60A6E}\test000.vhd
```

The following message shows information about the settings restore phase:

```
PS C:\Program Files\ISCSI Target> .\iSCSITargetSettings.PS1 -Import -file
$env:temp\test00000000.xml
```

```
Importing settings from file
'E:\Users\administrator\AppData\Local\Temp\test00000001.xml'.
The operation may take a long time.


Number of Target(s) imported from the Import settings: 4.
Targets:
    test000
    test001
    test002
    test1111


Number of Virtual Disk(s) imported from the Import settings: 3.
Virtual Disk:
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd
```

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Migrate an iSCSI Software Target

Prepare to Migrate iSCSI Software Target

Migrate iSCSI Software Target

Verify the iSCSI Software Target Migration

Roll Back a Failed iSCSI Software Target Migration

Roll Back a Failed iSCSI Software Target Migration

If iSCSI initiators have successfully reconnected to the iSCSI Target Server computer, the migration is successful and complete. This topic discusses the tasks that should be performed in the event of a failed migration.

Restoring the role if the migration failed

If migration does not complete successfully, a rollback procedure is required to undo any changes to the source server, other servers, and client computers, and then restore the source server back into service.

Rollback requirements

The rollback procedure requires that the source server is available in the same state as it was after removing the network identity of the iSCSI Software Target server. For more information, see Remove the network identity of the iSCSI Software Target computer.

During the source server preparation steps, none of the steps performed permanently changed the existing configuration of the server because all of the operations were substantially read operations.

The estimated time to complete the rollback is equivalent to the time that it takes to re-establish the network identity of the source server. This operation may require rolling back changes to the DHCP servers, DNS server, or Active Directory Domain controllers.

Roll back iSCSI initiators on other computers

The other computers in the enterprise that are affected by migrating ISCSI Software Target are the iSCSI initiators.

In the case of a rollback, the iSCSI initiators that were configured to log on to the destination server need to be rolled back to the source server. Use the following Windows PowerShell commands:

1. To log out of an existing iSCSI session:

   ```
   PS > & iscsicli.exe sessionlistPS > & iscsicli.exe LogoutTarget <session id>
   ```

2. To discover the iSCSI Software Target source server:

   ```
   PS > & iscsicli AddTargetPortal <source server ip address> 3260PS > iscsicli.exe
   ListTargets
   ```

3. To log on to the targets on the iSCSI Software Target source server:

   ```
   PS > & iscsicli.exe  LoginTarget <target IQN> T < source server ip address> 3260
   Root\ISCSIPRT\0000_0 * * * * * * * * * * * * *
   ```

Roll back iSCSI Software Target on a stand-alone source server

This step will undo the network identity removal that is described in Remove the network identity of the iSCSI Software Target computer.

Possible scenarios include:

- Restore the NetBIOS fully qualified domain name to the source server, and assign the required IP addresses to the source server.
- Restore any DNS assignments (for example, reverse lookup and DHCP assignment).

- Restore any identities that were previously assigned in Active Directory.

Each scenario requires potentially updating information in the DNS server, Active Directory, or DHCP server, according to the methodology that is used to assign IP addresses and network names to the servers in the enterprise.

The intent of this step is to ensure that upon completion of the migration steps, the iSCSI initiators are able to locate the source server (either through explicit reconfiguration, or implicitly through the computer name or IP address reassignment).

Roll back iSCSI Software Target on a clustered source server

Rolling back iSCSI Software Target on a clustered source server requires two steps:

Step 1: Roll back cluster network name changes

This step will undo the network identity removal described in Remove the network identity of the iSCSI Software Target computer.

In a clustered configuration, network names are established by the server principal name that is assigned in Active Directory to the cluster when the cluster was formed.

To re-establish network names that were possibly deleted or retired, the cluster administration utilities must be used. For more information, see Migrating Clustered Services and Applications to Windows Server 2012.

Step 2: Move resource groups to the preferred owner node

After the client access points have been re-established, the resource groups need to be moved back to their preferred owner node.

The resource groups were moved to a single node as part of the steps performed in Capture the existing settings: clustered configuration.

To move the resource groups back to their preferred owner node, use the following Windows PowerShell command:

```
PS > & cluster.exe /cluster:<cluster name> GROUP <group name> /moveto:<node name>
```

Note

The group name and the node names were obtained during the previous preparation tasks.

Roll back iSCSI Target Server on a stand-alone destination server

To roll back iSCSI Target Server on a stand-alone destination server that is running Windows Server 2012 or Windows Storage Server 2012, uninstall the **iSCSI Target Server** role service using Server Manager.

Roll back iSCSI Target Server on a clustered destination server

To roll back iSCSI Target Server on a destination server that is running Windows Server 2012 or Windows Storage Server 2012 in a clustered configuration, first remove any client access point that was created for iSCSI Target Server and then uninstall the **iSCSI Target Server** role service using Server Manager.

Retiring iSCSI Software Target on a source server

Retiring iSCSI Software Target 3.2 or iSCSI Software Target 3.3 on your source server requires using the following Windows PowerShell commands:

Retire iSCSI Software Target

1. Find the package GUID:

   ```
   PS > Get-WmiObject -Class Win32_product | Where-Object { $_.packageName -match 'iscsitarget'}
   ```

2. Uninstall the package:

   ```
   PS > & msiexec /uninstall <package GUID> /qr
   ```

Retiring a source server

In a stand-alone configuration, there are no particular procedures for retiring the source server. In a clustered configuration, the client access points that are devoted to iSCSI Software Target access can be removed by using the following Windows PowerShell command:

```
PS > Remove-ClusterGroup -Name <resource group name> -RemoveResources -Force
```

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Migrate an iSCSI Software Target

Prepare to Migrate iSCSI Software Target

Migrate iSCSI Software Target

Verify the iSCSI Software Target Migration

Troubleshoot the iSCSI Software Target Migration

File and Storage Services: Migrate Network File System

This topic describes how to migrate Network File System shares and settings from previous versions of Windows Server to Windows Server 2012 R2.

Network File System Migration overview

You can migrate Network File System (NFS) from a server running Windows Server 2012, Windows Server 2008 R2, Windows Server°2008, or Windows Server°2003°R2 to a server running Windows Server 2012 R2 using the procedures described in this topic. Some of the methods that are available for migrating NFS are the following:

- You can gather output from NFS servers running on previous versions of Windows Server, and then modify and use this information to input into the new NFS server running Windows Server 2012 R2. This can be done using the NFS cmdlets in Windows PowerShell, or using command tools such as **nfsshare** and **nfsadmin**.
- You can gather output from the NFS servers running on previous versions of Windows Server, and then use this information as a reference when manually configuring NFS settings for the new NFS server running Windows Server 2012 R2. This can be done using the NFS cmdlets in Windows PowerShell or the File Server Administration Console in Server Manager.

Migrating NFS Server from Windows Server°2012 to Windows Server°2012°R2

This section explains how to migrate NFS shares and permissions from Windows Server 2012 to Windows Server 2012 R2. Introduced in Windows Server 2012, the NFS cmdlets in Windows PowerShell allow you to manage NFS shares and settings, export shares and configuration metadata to .xml files, and then import the files into Windows Server 2012 R2. During this process, UNIX or Linux-based style group and password files are copied to Windows Server 2012 R2. If you used Active Directory Lightweight Directory Services (AD°LDS) to configure name mapping, see Active Directory Lightweight Directory Services Overview.

Export the server configuration

Before starting the export process, you must first create a directory (for example C:\tmp) where all the files will be exported.

Open Windows PowerShell, and to export the NFS server configuration information, type:

```
PS C:\tmp> Get-NfsServerConfiguration | Export-Clixml NfsServerConfig.xml
```

Export NFS shares

To export the NFS share settings information, open Windows PowerShell, and type the following. Note that this procedure does not include exporting the NFS share permissions.

```
PS C:\tmp> Get-NfsShare | Export-Clixml NfsShares.xml
```

Next, update the host configuration information using the following steps. You can ignore these steps if the net name and host names are going to remain the same.

1. Open the file where the exported shares are located (for example, c:\tmp\NfsShares.xml).
2. Find the network name and host name, and then rename them as appropriate.
3. If necessary, update the location of the directory path.
4. Save the file that contains the exported shares (such as c:\tmp\NfsShares.xml).

Export NFS share permissions

To export the NFS share permissions for all the NFS shares, type:

```
PS C:\tmp> Get-NfsShare | Get-NfsSharePermission | Export-Clixml NfsSharePermission.xml
```

Next, update the host configuration information using the following steps. You can ignore these steps if the net name and host names are going to remain the same.

1. Open the file where the exported permissions are located (for example, c:\tmp\NfsSharePermission.xml).
2. Find the network name and host name, and then rename them as appropriate.
3. If necessary, update the location of the directory path.
4. Save the file that contains the exported permissions (such as c:\tmp\NfsShares.xml).

Copy local mapping data

If you are using the UNIX or Linux-based local password and group files to map between UNIX and Linux-based users and Windows, copy the following files from Windows Server 2012. You can ignore this step if you are not using UNIX or Linux-based local password and group files.

```
PS C:\tmp> COPY %SystemRoot%\system32\drivers\etc\passwd C:\tmp
```

```
PS C:\tmp> COPY %SystemRoot%\system32\drivers\etc\group C:\tmp
```

Export identity mapping

In Windows PowerShell, type the following to display identity mapping information (such as Lightweight Directory Access Protocol or AD LDS) used by the NFS Server. This information must be manually recreated in Windows Server 2012 R2. If no identity mapping stores are configured, you can ignore this step.

```
PS C:\tmp> Get-NfsMappingStore | Export-Clixml nfsmappingstore.xml
```

Note

The group and user identity mapping are expected to remain the same after the migration.

Export netgroups and client groups

Configuring netgroups and client groups makes it easier to manage computer and user authentication. In Windows PowerShell, type the following to display information about netgroups and client groups, which can then be exported to Windows Server 2012 R2.

```
PS C:\tmp> Get-NfsNetgroup | Export-Clixml nfsnetgroup.xml

PS C:\tmp> Get-NfsNetgroupStore | Export-Clixml nfsnetgroupstore.xml

PS C:\tmp> Get-NfsClientGroup | Export-Clixml nfsclientgroup.xml
```

Importing NFS shares and settings from Windows Server°2012 to Windows Server°2012°R2

This section describes how to import NFS shares and settings that you exported from Windows Server 2012 to Windows Server 2012 R2. First, create a directory (for example C:\tmp) on the computer running Windows Server 2012 R2 and copy all the files exported from Windows Server 2012.

Note

The settings for NFS shares are metadata used over existing volumes and directories. Therefore, you should make sure the data and directory structure are correct before NFS share settings are applied. After the directory structure is in place, you can proceed to the following procedure. For more information about data migration, see Impact of data migration by copying data and shared folders.

Import the server configuration

Before importing the server configuration, make sure that you have installed the **Server for NFS** role service in Server Manager. To import the server configuration, open Windows PowerShell, and type:

```
PS C:\tmp> Import-Clixml NfsServerConfig.xml | Set-NfsServerConfiguration
```

Restart **Server for NFS** by using either **Control Panel** or by typing `Restart-Service NfsService` at a command prompt.

Import NFS shares

Before performing this step, make sure that the directory structure is already in place and that the Nfsshares.xml file is updated with the appropriate location, server names, and any additional important information.

To import NFS share settings, open Windows PowerShell, and type:

```
PS C:\tmp> Import-Clixml NfsShares.xml | %{New-NfsShare -Name $_.Name -Path $_.Path -
NetworkName $_.NetworkName -EnableAnonymousAccess
$_.AnonymousAccess -AnonymousUid $_.AnonymousUid -AnonymousGid $_.AnonymousGid -
EnableUnmappedAccess
$_.UnmappedUserAccess -Authentication $_.Authentication}
```

You should resolve any errors before proceeding to the next step.

Import NFS share permissions

Before performing this step, make sure that the Nfssharepermission.xml file is updated with the correct server names. To import NFS share permissions, open Windows PowerShell, and type:

```
PS C:\tmp> Import-Clixml NfsSharePermission.xml | foreach { $_ |Grant-NfsSharePermission}
```

Import local mapping data

If UNIX and Linux-based local password and group files are used for mapping between UNIX and Linux users and Windows, copy the following files (which were exported from Windows Server 2012). You can ignore this step if you do not use UNIX and Linux-based password and group files.

```
PS C:\tmp> COPY  C:\tmp\passwd %SystemRoot%\system32\drivers\etc\passwd

PS C:\tmp> COPY C:\tmp\group %SystemRoot%\system32\drivers\etc\group
```

Import non-local identity mapping

If you are using methods, such as LDAP or AD LDS, to configure identity mapping, use the following Windows PowerShell script to import the .xml file:

```
PS C:\tmp> Import-Clixml nfsmappingstore.xml | Set-NfsMappingStore
```

Import netgroups and client groups

In Windows PowerShell, type the following to export netgroups and client groups to Windows Server 2012 R2:

```
PS C:\tmp> Import-Clixml nfsnetgroup.xml | Set-NfsNetgroup

PS C:\tmp> Import-Clixml nfsnetgroupstore.xml | Set-NfsNetgroupStore

PS C:\tmp> Import-Clixml nfsclientgroup.xml | Set-NfsClientGroup
```

After the netgroups and client groups are defined, permission to access shares that an NFS server exports can be configured using the Grant-NfsSharePermission Windows PowerShell cmdlet. Some examples for granting share permissions are shown in the following generated information in Windows PowerShell:

```
PS C:\> New-NfsClientgroup -ClientGroupName MIGRATION -AddMember 'MACHINE1','MACHINE2'

PS C:\> Get-NfsClientgroup MIGRATION

ClientGroupName                                          ClientGroupMembers

---------------                                          ------------------

MIGRATION                                                {MACHINE1, MACHINE2}

PS C:\> Grant-NfsSharePermission -Name NfsTestShare1 -ClientName MIGRATION -ClientType
clientgroup -Permission readonly

PS C:\> Get-NfsSharePermission NfsTestShare1
```

```
Name                 ClientName             Permission        AllowRootAccess

----                 ----------             ----------      ---------------

NFSTestShare1         MIGRATION               READ              False
```

If you are using Unmapped UNIX User Access (UUUA), see NFS Identity Mapping in Windows Server 2012, which provides information about the various methods of identity mapping. You should note that both Windows Server 2012 R2 and Windows Server 2012 support UNIX and Linux-based password and group files.

NFS server and share settings migration from Windows Server 2012 to Windows Server 2012 R2 is complete.

Migrating NFS Server from Windows Server°2008°R2, Windows Server°2008, or Windows Server°2003°R2 to Windows Server°2012°R2

This section describes how to migrate NFS shares and permissions from Windows Server 2008 R2 and earlier versions of the Windows Server operating system to Windows Server 2012 R2. Using the command-line tools, **nfsshare** and **nfsadmin**, you can export NFS shares and settings, and then import the files into Windows Server 2012 R2.

Get server configuration

To retrieve information from the NFS server configuration, type the following at a command prompt:

```
C:\tmp> nfsadmin server
```

After running the command, create a copy of the information that is generated. An example of output from Windows Server 2008 R2 follows:

```
The following are the settings on localhost


Locking Daemon Grace Period : 45 seconds

Activity logging Settings   :

Protocol for Portmap        : TCP+UDP

Protocol for Mount          : TCP+UDP

Protocol for NFS            : TCP+UDP

Protocol for NLM            : TCP+UDP

Protocol for NSM            : TCP+UDP

Protocol for Mapping Server : TCP+UDP
```

```
Protocol for NIS            : TCP+UDP

Enable NFS V3 Support       : Enabled

Renew Authentication        : Enabled

Renew Authentication Interval : 600 seconds

Directory Cache             : 128 KB

Translation File Name       :

Dot Files Hidden            : Disabled

Case Sensitive Lookups      : Enabled

NTFS Case                   : Preserve Case

NetGroup Source             : none

NIS Server                  :

NIS Domain                  :

LDAP Server or AD Domain    :

LDAP naming context (DN)    :
```

Collect NFS shares information

You can display NFS share settings by using the following commands. You should note that running this command does not display NFS share permissions.

To retrieve the list of NFS shares configured in the server, type the following at a command prompt:

```
C:\tmp> nfsshare
```

To retrieve detailed information for each NFS share listed after using the preceding command, type:

```
C:\tmp> nfsshare <share-name>
```

Example output:

```
C:\tmp> nfsshare

        share1 = C:\shares\share1

        share2 = C:\shares\share2

C:\tmp> nfsshare share1

   Alias = share1

   Path = C:\shares\share1

   Supported security flavors are SYS:KRB5:KRB5I

   Encoding = ansi

   UNMAPPED UNIX USER access allowed
```

```
ANONYMOUS access disallowed

Anonymous UID = -2

Anonymous GID = -2

HOST ACCESS :

   ALL MACHINES            read-write     Root Access Allowed      ansi
```

Note
Kerberos authentication was introduced in Windows Server 2008 R2 for use with NFS, and therefore, so it is not available in earlier versions of the operating system.

Collect identity mapping and group identifier information

To display identity mapping settings (such as the Network Information Service [NIS] server, NIS domain, and LDAP or AD LDS information), type the following at a command prompt:

```
C:\tmp> nfsadmin server
```

To display the identity mapping methods that are used, type:

```
C:\tmp> nfsadmin mapping
```

Example output:

```
C:\tmp> nfsadmin mapping

The following are the settings on localhost

Mapping Server Lookup       : Disabled

Mapping Server              :

AD Lookup                   : Disabled

AD Domain                   :
```

To display the names of all client groups, type:

```
C:\tmp\nfsadmin server listgroups

C:\tmp\nfsadmin server listmembers <client group name>
```

Reconfiguring NFS shares and settings from Windows Server°2008°R2, Windows Server°2008, or Windows Server°2003°R2 to Windows Server°2012°R2

This section explains the process of reconfiguring the NFS shares and settings that you exported from Windows Server 2008 R2 or previous versions of Windows Server to Windows Server 2012 R2. You can reconfigure NFS shares and settings using the **nfsshare** or **nfsadmin** command tools, or, if you are migrating from Windows Server 2008 R2, you can use the NFS cmdlets in Windows PowerShell.

Before you import NFS shares and settings, make sure that you have installed the **Server for NFS** role service in **Server Manager** on the computer running Windows Server 2012 R2.

The settings for NFS shares are metadata used over existing volumes and directories. Therefore, you should make sure the data and directory structure are correct before NFS share settings are applied. After the directory structure is in place, you can proceed to the following procedure. For more information about data migration, see Impact of data migration by copying data and shared folders.

Set up the NFS server configuration

To configure the NFS server, type the following in Windows PowerShell. Instructions for setting up the NFS server using the **nfsadmin** command are provided later in this section.

In Windows PowerShell, type:

```
PS C:\tmp> Set-NfsServerConfiguration  -[parameters as displayed below]
```

Windows Server°2008°R2 output	Equivalent Windows PowerShell cmdlet parameters in Windows Server°2012°R2
Locking daemon grace period	GracePeriodSec
Protocol for Portmap	PortmapProtocol
Protocol for Mount	MountProtocol
Protocol for NFS	NfsProtocol
Protocol for NLM	NlmProtocol
Protocol for NSM	NsmProtocol
Protocol for Mapping Server	MapServerProtocol
Protocol for NIS	NisProtocol
Enable NFS V3 support	EnableNFSV3
Renew Authentication	EnableAuthenticationRenewal
Renew Authentication Interval	AuthenticationRenewalIntervalSec
Directory Cache	DirectoryCacheSize
Translation File Name	CharacterTranslationFile
Dot Files Hidden	HideFilesBeginningInDot
Activity Logging Setting	LogActivity

Notes

In Windows Server 2012 R2, there is a new parameter *LeasePeriodSec* for the Set-NfsServerConfiguration Windows PowerShell cmdlet. When setting the *GracePeriodSec* value, make sure that the *LeasePeriodSec* value is set to 50 percent of *GracePeriodSec*.

Case-sensitive file lookups and case-sensitive preservation can no longer be configured in Windows Server 2012 R2 because they are now system-wide Windows settings.

If you prefer to use the **nfsadmin** command tool, type the following at a command prompt:

```
C:\tmp> nfsadmin server config config_options
```

For a detailed list of configuration options for **nfsadmin**, type **nfsadmin server /?**.

Configure NFS shares

To configure NFS shares using the information you previously gathered on the shares, type the following in Windows PowerShell. Instructions for configuring NFS shares using the **nfsshare** command are provided later in this section.

```
PS C:\tmp> New-NfsShare <parameters>
```

Windows Server output	Equivalent NfsShare cmdlet parameters in Windows Server°2012°R2
Alias	-Name
Path	-Path
Encoding	-LanguageEncoding
Anonymous access	-EnableAnonymousAccess
Anonymous UID	-AnonymousUID
-Anonymous GID	-AnonymousGID
Host access	-Permission, -AllowRootAccess, …

An example of configuring an NFS share follows:

```
PS C:\> New-NfsShare –Name roshare – Path C:\shares\roshare =AnonymousUid -2 –
AnonymousGid -2 –LanguageEncoding ANSI –EnableAnonymousAccess $false –
EnableUnmappedAccess $false –AllowRootAccess $false
```

If you prefer to use the **nfsshare** command tool, type the following at a command prompt:

```
C:\tmp> nfsshare sharename=drive:path [ -o options ]
```

For a detailed list of configuration options for **nfsshare**, type **nfsshare server /?**.

Configure identity mapping and group identifier information

Using the information you gathered earlier for identity mapping, type the following in Windows PowerShell to configure ID mapping:

```
PS C:\tmp> Set-NfsMappingStore <Parameters>

PS C:\tmp> Set-NfsMappedIdentity <Parameters>
```

An example of configuring identity mapping follows:

```
PS c:\tmp> Set-NfsMappingStore -EnableADLookup $true -ADDomainName "Contoso.com"
```

If you prefer to use the **nfsadmin** command tool, type the following at a command prompt:

```
C:\tmp> nfsadmin server  <parameter for NIS server or LDAP server information>

C:\tmp> nfsadmin mapping <parameters>
```

For a detailed list of configuration options for **nfsadmin**, type **nfsadmin server /?**.

An example of configuring identity mapping using **nfsadmin** follows:

```
c:\tmp> nfsadmin mapping

The following are the settings on localhost

Mapping Server Lookup       : Disabled

Mapping Server              :

AD Lookup                   : Enabled

AD Domain                   : Contoso.com
```

Using the information you gathered earlier for group identifiers, type the following in Windows PowerShell to configure groups:

```
Set-NfsgroupStore

Set-NfsClientGroup  (or) New-NfsClientGroup

Set-NfsNetGroup (or) New-NfsNetGroup
```

After you have configured the netgroup and client group, you can set the NFS share permissions using the Grant-NfsSharePermission Windows PowerShell cmdlet . Some examples of configuring share permissions follow:

```
PS C:\> New-NfsClientgroup -ClientGroupName MIGRATION -AddMember 'MACHINE1','MACHINE2'

PS C:\> Get-NfsClientgroup MIGRATION

ClientGroupName                                          ClientGroupMembers

---------------                                          ------------------

MIGRATION                                                {MACHINE1, MACHINE2}

PS C:\> Grant-NfsSharePermission -Name NfsTestShare1 -ClientName MIGRATION -ClientType
clientgroup -Permission readonly

PS C:\> Get-NfsSharePermission NfsTestShare1
```

```
Name                 ClientName           Permission       AllowRootAccess
----                 ----------           ----------     ---------------
NFSTestShare1         MIGRATION              READ             False
```

If you are using Unmapped UNIX User Access (UUUA), see NFS Identity Mapping in Windows Server 2012, which provides information about the various methods of identity mapping. You should note that both Windows Server 2012 R2 and Windows Server 2012 support UNIX and Linux-based password and group files.

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Migrate an iSCSI Software Target

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Post-Migration Tasks

This topic explains how to complete the migration if it was successful, and how to roll back or troubleshoot the migration if it failed.

Completing the migration

After you verify the migration, you can retire the source server.

Retire File and Storage Services on the source server

After you complete and verify the migration, the source server can be shut down or disconnected from the network.

Remove DFS Namespaces from the source server

The procedure you use to remove DFS Namespaces from the source server depends on whether the namespaces are stand-alone or domain-based. If you want to remove the namespace from the source server, you must use **DFSUtil.exe**.

Note

By default, clients cache the list of namespace servers for 300 seconds (five minutes), so we recommend that you do not run the **DFSUtil.exe remove** command within five minutes of completing verification of the DFS namespace migration. During migration, clients have only the temporary server in the cache of namespace servers. Waiting five minutes after you add the destination server to the namespace allows clients to list the destination server in their cache.

To remove stand-alone namespaces

1. Open a Command Prompt window on the destination server.
2. Type the following code, and then press Enter.

```
Dfsutil.exe root remove <\\SourceServer\Namespace>
```

To remove domain-based namespaces with one namespace server

1. On the destination server, open a Command Prompt window.
2. Type the following, and then press **Enter**.

```
DFSUtil.exe target remove <\\TemporaryServer\Namespace>
```

Note

This procedure applies only if a temporary server was added to the namespace for migration purposes. For domain-based namespaces with more than one namespace server, no additional actions are required.

Restoring File and Storage Services in the event of migration failure

The following sections describe how to restore the File and Storage Services server role in the event of migration failure.

Roll back DFS Namespaces

The steps that you perform to roll back DFS Namespaces depend on whether the namespaces are stand-alone or domain-based, and whether a temporary namespace was created during the migration process.

To roll back DFS Namespaces (do one of the following)

1. For stand-alone namespaces, no action is required other than migrating the identity back to the source server.
2. For domain-based namespaces with more than one namespace server, or if a temporary server was added to a namespace that initially had only one namespace server, do the

following:

a. Remove the destination server from the namespace.
b. Migrate identity and shared folder information to the source server.
c. Add the source server to the namespace.

3. For domain-based namespaces with only one namespace server, where no temporary namespace server was added during migration, do the following:
   a. Migrate identity and shared folder information to the source server.
   b. Verify that the export file for the namespace that was created during migration is still available.
   c. Delete the namespace.
   d. Create the namespace on the source server.
   e. Import the namespace configuration from the export file created during the migration.
   f. Manually reset delegation permissions to the namespace.

Note

Another option to migrate domain-based namespaces with one namespace server is to temporarily add a second namespace server before the migration, and then remove the temporary server after the migration.

Roll back data and shared folders

If no changes have been made to migrated files, folders, and shared folders on the destination server and this data has not been deleted from the source server, no additional steps to roll back data and shared folders are required.

If the migrated files, folders, or shared folders may have been modified on the destination server by the administrators or users, perform the following steps to synchronize the changes from the destination server back to the source server:

1. Type the following command in a Command Prompt window to copy the updated migrated data (files and folders) from the destination server back to the source server:

   ```
   robocopy <copy from path> <copy to path> /E
   ```

   This command can be executed on the source server or on the destination server, and it will recursively copy updated data. Type `robocopy /?` in a Command Prompt window for additional copy options, including options to copy file and folder permissions.

   Caution

   Permissions that you set for nondefault local users and groups will not copy properly and need to be created manually.

2. Compare the lists of shared folders and their permissions on the source server and destination server and manually synchronize any changes.

   To list all shared folders and their permissions, type the following command in a Windows PowerShell session that has been opened with elevated user rights:

```
gwmi win32_share | %{net share $_.name}
```

Roll back migration on the other computers in the enterprise

If the migration failed, verify that the other computers in the enterprise can access the source server after you roll back the migration data.

Troubleshooting migration issues

Troubleshooting tips include the following:

- For physical migration issues:

  When some files are migrated physically and others are copied, there is a chance that the File Server Resource Manager configuration is not synchronized. To remedy this, delete and create new copies of the Quota.md and Datascrn.md files.

- For domain-joined computers:

  If a custom action (quota notification or file management task) fails to execute with an access-denied failure and a corresponding event log, you should remove the custom action and create it on the destination server.

Troubleshoot data migration that does not complete

If the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets run indefinitely without completing, your destination server might not have sufficient disk space or a large enough File Server Resource Manager or NTFS quota limit to allow for data migration to finish. To determine whether insufficient disk space is preventing the data send-receive process from completing, do the following on the destination server.

1. Open %*localappdata*%/Svrmig/Log/SetupAct.log.
2. Review the most recent log entries. If the following exception occurs, your destination server has insufficient disk space or File Server Resource Manager or NTFS quota limits to complete data migration.

   ```
   Win32Exception: unable to write to FileStream: There is not enough space on the disk.
   ```

To resolve this issue, do the following

1. Press Ctrl+C to cancel **Send-SmigServerData** and **Receive-SmigServerData** on both source and destination servers.
2. Check for sufficient disk space on the destination server's hard disk drive. If the destination server's hard disk drive has insufficient space, do one of the following:
   - Clear additional space.
   - Identify a different hard disk drive that has sufficient space.

3. If the destination server's hard disk drive, the destination path, or any folders that contain the destination path have a File Server Resource Manager or NTFS quota enabled, and the quota limit does not allow for sufficient disk space to migrate data, do one of the following:
   - Increase the quota limit to set sufficient disk space to migrate the data. For more information about FSRM quota management, see one of the following.
     - Quota Management for Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008
     - Quota Management for Windows Server 2003 R2

     For more information about NTFS quota management, see one of the following.
     - Setting Disk Quotas for Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008
     - Enable disk quotas for Windows Server 2003 R2 and Windows Server 2003
   - Identify a different hard disk drive that already has sufficient space and File Server Resource Manager or NTFS quota limits.
4. Run the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets again, specifying a destination path that has sufficient disk space, and large enough File Server Resource Manager or NTFS quota limits, if applicable.

Troubleshoot data migration connectivity

If the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets cannot establish connectivity, verify the following conditions and then try again:

1. In the **Send-SmigServerData** command on the source server, the *ComputerName* parameter correctly specifies the name of the destination server.
2. The **Receive-SmigServerData** and **Send-SmigServerData** commands are entered on the destination server and the source server respectively within five minutes of one another. This is the default maximum connection timeout for **Send-SmigServerData** and **Receive-SmigServerData**. You can change the maximum connection timeout for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets by modifying the following user-defined registry subkey on the source server and destination server.

   Subkey: **\HKEY_Local_Machine\Software\Microsoft\ServerMigration**

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents the connection time-out in seconds). If a value larger than 3600 is specified, 3600 seconds is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key.
3. The same password is entered on the source server and destination server.
4. The source server and destination server are available on the same subnet:
   a. On the destination server, in a Command Prompt window, type `ipconfig` and note the subnet mask value.
   b. On the source server, in a Command Prompt window, type `ipconfig` and note the subnet mask value.

c. Ensure that the subnet mask values are the same on the source server and destination server.

5. Port 7000 is open on the source and destination server and is not in use by another application.
   a. To check if port 7000 is open, in a Command Prompt window, enter the command:

      ```
      netsh firewall show portopening
      ```

      If port 7000 is not in the list, follow the instructions in File and Storage Services: Appendix A: Optional Procedures to open port 7000.

   b. If port 7000 is open, type the following command to check if port 7000 is being used by another application:

      ```
      netstat
      ```

      - In the Local Address column, you will see <*IP Address*>: <*port number*>.
      - If port 7000 is in the list, it is being used by another application.

Troubleshoot unexpected Windows PowerShell session closure

If a migration cmdlet fails and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at 0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050). Minidump attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting**, or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Locate the deployment log file

The Windows Server Migration Tools deployment log file is located at %windir%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations:

- `%windir%\Logs\ServerMigration.log`
- On Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008: `%localappdata%\SvrMig\Log`
- On Windows Server 2003: `%userprofile%\Local Settings\Application Data\SvrMig\Log`

If migration log files are not created in the preceding locations, ServerMigration.log and SmigDeploy.log are created in `%temp%`, and other logs are created in `%windir%\System32`.

View the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects and query them for more information about the settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets in Windows Server Migration Tools return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property Name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if the migration was successful; otherwise, **False** is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or shared folder that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property Name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item shown as a path name.
Success	Boolean	The value **True** is displayed if the migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event

Property name	Type	Definition
		messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.
Title	String	Title of the result object. Examples of values include the physical address of the network adapter for IP configuration, or the user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after the migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in a variable, type a command in the following format, and then press **Enter**.

   $*VariableName* **=** $(*Cmdlet*)

   The following is an example.

   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -User all -Group -Path D:\rmt\DemoStore -force -Verbose)

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.

2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

   $*VariableName*

   In the following example, the variable is named **ImportResult**.

   $*ImportResult*

   This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable:

```
           ItemType  ID                                  Success
DetailsList
           --------  --                                  -------
-----------
          OSSetting  Local User                            True
{Local User, Loc...
          OSSetting  Local Group                           True
{Local Group, Lo...
     WindowsFeature  DHCP                                  True
{}
```

Each line of the preceding example is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by the examples that follow in steps 3 and 4.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   $<VariableName>| Select-Object -ExpandProperty <PropertyName>

   The following is an example.

   $importResult | Select-Object -ExpandProperty DetailsList

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples:

   - The following command returns only those details of result objects that have the ID **Local User**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Severity -eq "Warning") {$_} }

   - The following command returns only the details of result objects with an ID of **Local User** that also have the title **Remote Desktop Users**.

     $ImportResult | Where-Object { $_.ID -eq "Local Group" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Title -eq "Remote DesktopUsers") {$_} }

More information about querying results

For more information about the cmdlets that are used in the preceding examples, see the following additional resources:

- Using the Where-Object Cmdlet
- Using the Select-Object Cmdlet
- Using the ForEach-Object Cmdlet

For more information about Windows PowerShell scripting techniques, see What Can I Do With Windows PowerShell?

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Migrate an iSCSI Software Target

File and Storage Services: Migrate Network File System

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix A: Optional Procedures

Opening ports in Windows Firewall

The following instructions are for opening ports in Windows Firewall. If you have a non-Microsoft firewall installed, consult the guide for that firewall about how to open ports. Opening ports in Windows Firewall can be done through the command line.

Important
Opening ports in your firewall can leave your server exposed to malicious attacks. Make sure that you understand firewall systems before you open ports.

To open Windows Firewall ports by using the command line (do one of the following):

1. Open a Command Prompt window with elevated user rights, type the following, and then press Enter.
   - On computers that are running Windows Server 2003, type:

     ```
     netsh firewall add allowedprogram
     ```

```
program=%windir%\System32\WindowsPowerShell\v1.0\powershell.exe name="ServerMigration" mode=ENABLE
```

- On computers that are running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, type the following commands in order, and press Enter after each command.

  i.

  ii.

2. If you have changed the default behavior of Windows Firewall to block all outbound traffic on computers that are running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, you must explicitly allow outbound traffic on UDP port 7000. To do this, open a Command Prompt window with elevated user rights, type the following, and then press Enter.

```
netsh advfirewall firewall add rule name=ServerMigration(UDP-Out) dir=out
program=%windir%\System32\WindowsPowerShell\v1.0\powershell.exe action=allow protocol=UDP localport=7000
```

Closing ports in Windows Firewall

As a best practice, we recommend that you close Windows Firewall ports after the data transfer operation is completed.

To close Windows Firewall ports by using the command line

- Do one of the following:
  - On computers that are running Windows Server 2003, open a Command Prompt window, type the following, and then press Enter.

```
netsh firewall delete allowedprogram
program=%windir%\System32\WindowsPowerShell\v1.0\powershell.exe
```

  - On computers that are running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, open a Command Prompt window with elevated user rights, type the following two commands, and press Enter after each command.

```
netsh advfirewall firewall delete rule
name=ServerMigration(TCP-In)
```

```
netsh advfirewall firewall delete rule
name=ServerMigration(UDP-Out)
```

Detect reparse points and hard links

The following commands can be used to detect reparse points and mounted volumes in any folder and its subfolders. Open a Command Prompt window, type the following commands to detect reparse points, in which **D:\Test** represents the hard disk drive and folder that you want to search, and then press Enter.

```
dir D:\Test\* /S /A:L
```

The option **/A:L** specifies that only reparse points need to be enumerated. The output is similar to the following:

```
Volume in drive D has no label.

Volume Serial Number is 3AE4-E412


Directory of D:\Test\Links


10/07/2008  03:44 PM    <JUNCTION>     JunctionMSIT [d:\test\targets\msit]

10/07/2008  03:42 PM    <SYMLINK>      LinkMSIT [d:\test\targets\msit]

10/07/2008  03:41 PM    <SYMLINKD>     SymLinkMSIT [d:\test\targets\msit]

1 File(s)              0 bytes


Directory of D:\Test\Targets


10/07/2008  05:35 PM    <JUNCTION>     Volume [\??\Volume{0674413f-760d-11dd-beb3-
806e6f6e6963}\]

0 File(s)              0 bytes


Total Files Listed:

1 File(s)              0 bytes

3 Dir(s)  17,918,840,832 bytes free
```

To enumerate hard links on a file on Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2, open a Command Prompt window with elevated user rights, type the following command, and then press Enter.

```
fsutil hardlink list D:\Test\File.txt
```

To enumerate hard links on all files in a folder on Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2, run the following command in a Windows PowerShell session that has been opened with elevated user rights:

```
Get-ChildItem D:\* | %{'Links for: ' + $_.FullName; fsutil hardlink list $_.FullName; ""}
```

For more information about enumerating hard links on computers that are running Windows Server 2008 or Windows Server 2003, see FindFirstFileNameW function on MSDN.

Migrated and nonmigrated attributes for local users and groups

For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide.

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Migrate an iSCSI Software Target

File and Storage Services: Migrate Network File System

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix B: Migration Data Collection Worksheets

SMB data collection worksheet

Use this server message block (SMB) data collection worksheet to record data for SMB policies that are set on the source server.

#	Source Server Essential Settings	Setting Identification
01	Idle time The setting name is: **Microsoft network server: Amount of idle time required before suspending a session**.	Idle time (in minutes): _____ Group or Local Policy: _____
02	S4USelf The setting name is: **Microsoft network server: Attempt**	Claim information: __ Default __ Enabled or __ Disabled Group or Local Policy:

#	Source Server Essential Settings	Setting Identification
	S4USelf to obtain claim information.	_____
03	Sign (always) The setting name is: **Microsoft network server: Digitally sign communications (always)**.	Sign always: __ Enabled or __ Disabled Group or Local Policy: _____
04	Sign (if client agrees) The setting name is: **Microsoft network server: Digitally sign communications (if client agrees)**.	Sign if client agrees: __ Enabled or __ Disabled Group or Local Policy: _____
05	Disconnect when logon hours expire The setting name is: **Microsoft network server: Disconnect clients when logon hours expire**.	Disconnect: __ Enabled or __ Disabled Group or Local Policy: _____

BranchCache data collection worksheet

Use this BranchCache data collection worksheet to record data for the BranchCache policies that are set on the source server.

#	Source Server Essential Settings	Setting Identification
01	BranchCache The setting name is: **Hash Publication for BranchCache**.	BranchCache: __ Not configured, __ Enabled, or __ Disabled Group or Local Policy: _____
	BranchCache The setting name is: **Hash Version support for BranchCache**.	BranchCache: __ Not configured, __ Enabled, or __ Disabled Group or Local Policy: _____

See Also

Migrate File and Storage Services to Windows Server 2012 R2

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Migrate an iSCSI Software Target

File and Storage Services: Migrate Network File System

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

Migrate Remote Desktop Services to Windows Server 2012 R2

Remote Desktop Services is a role in the Windows Server operating system that provides multi-user access to applications and desktops for non-administrative purposes. This guide describes how to migrate Remote Desktop Services, what Remote Desktop Services role services will be migrated, and tasks that apply to migrating the role services.

About this guide

This guide describes how to migrate the Remote Desktop Services role by providing preparation, migration and verification steps.

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012 R2. By using the process described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help eliminate possible conflicts that might otherwise occur during the migration process.

Note

Your detailed feedback is very important, and it helps us make Windows Server Migration Guides as reliable, complete, and easy-to-use as possible. Please take a moment to rate this topic and add comments that support your rating. Click **Rate this topic** at the top of the page. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Windows Server migration guides or tools, please write a post on the Windows Server Migration forum.

Target audience

This guide is intended for the following audiences:

- IT architects who are responsible for computer management and security throughout an organization
- IT operations engineers who are responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications
- IT operations managers who are accountable for network and server management

What this guide does not provide

This guide does not cover migration of the following:

- Customizations made to any Remote Desktop Services role service. In particular, this may apply to the RD Session Host, RD Virtualization Host, RD Web Access, or RD Connection Broker role services.
- Third-party application settings, programs, or plug-ins
- More than one server role at the same time
- More than one role service at a time
- Group Policy settings
- User profiles, including roaming profiles
- Event history
- Microsoft applications or application settings
- RD Connection Broker servers that are configured in a clustered or load-balanced environment (except High-Availability mode)

This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure that is specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional server roles are available on the Windows Server Migration Portal.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

Supported migration scenarios

This guide provides you with instructions for the following:

- Migrating a server that is running Remote Desktop Services on Windows Server 2012 to a server that is running Remote Desktop Services on Windows Server 2012 R2
- Migrating between two servers running Remote Desktop Services on Windows Server 2012 R2

For the migration scenarios that are described in this guide, each of the Remote Desktop Services role services is migrated separately. You can migrate one, some, or all role services by

following the steps in this guide. For information about the order of migration, see Order of migration for multiple role serviceslater in this topic.

Supported operating systems

The Remote Desktop Services role services are available in Windows Server 2012 R2 as follows:

- All Remote Desktop Services role services are available in Windows Server 2012 R2 Standard, Windows Server 2012 R2 Enterprise, and Windows Server 2012 R2 Datacenter.
- RD Web Access is available in Windows Web Server 2012 R2.
- RD Session Host, RD Licensing, RD Web Access, and RD Gateway are available in Windows Server 2012 R2 Foundation.

Physical to virtual machine migration

Migration between physical operating systems and virtual operating systems are supported for the RD Connection Broker, RD Session Host, RD Web Access, RD Licensing, and RD Gateway role services. However, the RD Virtualization Host role services and the Hyper-V role do not run on virtual machines.

Backward compatibility

You can migrate the Remote Desktop Services role services from computers running Windows Server 2012 or Windows Server 2012 R2 to a computer running Windows Server 2012 R2.

Policy and configuration settings

Some Remote Desktop Services settings can be configured by using Group Policy. Information about migrating Group Policy settings is not included in this migration guide.

Supported role services and features

This migration guide describes how to migrate the Remote Desktop Services role services from a source server running Windows Server 2012 or Windows Server 2012 R2 to a destination server running Windows Server 2012 R2.

Following are the Remote Desktop Services role services that can be migrated to a computer running Windows Server 2012 R2:

1. RD Connection Broker
2. RD Virtualization Host
3. RD Session Host
4. RD Web Access
5. RD Licensing
6. RD Gateway

Migration scenarios that are not supported

The following scenarios are not supported:

- Upgrading Windows Server 2008 Terminal Services or Windows Server 2012 R2 Remote Desktop Services server role or role services
- Migrating or upgrading from Windows Server 2003 or Windows Server 2003 R2
- Migrating from a source server to a destination server that is running an operating system with a different system UI language installed
- Migrating the RD Virtualization Host or RD Session Host role services from physical computers to virtual machines
- Migrating any applications or application settings from the RD Session Host server

Order of migration for multiple role services

The steps in this guide are based on migrating the role services in the following order when you are migrating more than one role service:

1. RD Connection Broker
2. RD Session Host
3. RD Virtualization Host
4. RD Web Access

The following role services can be migrated at any time during the migration:

- RD Licensing
- RD Gateway

In a Remote Desktop Services deployment, RD Connection Broker servers must be migrated first. All other services can be migrated independently. If you do not have RD Connection Broker servers, you can migrate other role services by following the steps provided in this document.

The Remote Desktop license server can be migrated at any time, but if the new license server does not have the same name as the source server, the Remote Desktop deployments and standalone RD Session Host servers that use that license server must be configured after migration to use the new license server.

The RD Gateway server migration is not dependent on the other role services for migration. It can be migrated at any time.

Impact of migration on Remote Desktop Services

A Remote Desktop Services role service will not be available during migration. This is also the case for any role services that are dependent on it. In addition, applications and add-ons on the affected servers will not be available.

Migration times will be affected by the dependencies between role services. For example, RD Session Host servers, RD Virtualization Host servers, and RD Web Access servers are dependent on RD Connection Broker servers. These dependencies should be considered when you are estimating downtime.

Plan your data migration to occur during off-peak hours to minimize downtime and reduce impact to users. Notify users that the resources will be unavailable during that time.

In some deployments, replication may extend the length of time that the services are unavailable.

If there is more than one role service on the source server, after you remove the source server from the domain, you will not have access to role services that you didn't migrate.

The following table details the expected impacts during the migration process.

Role service	Dependent role services	Impact of migration	Downtime estimates
RD Connection Broker	RD Virtualization Host, RD Session Host, RD Web Access	Users will not have access to any resources that are managed by the RD Connection Broker or TS Session Broker server that is being migrated. These resources include RemoteApp programs, virtual desktop pools, and personal virtual desktops.	Three hours
RD Session Host	RD Web Access may be dependent on RD Session Host in your deployment.	Session collections will not be available until migration of all destination servers in the virtual desktop collection is complete. RemoteApp programs will not be available until they are installed on the destination servers.	One hour
RD Virtualization Host	RD Virtualization Host is dependent on RD Connection Broker.	Virtual desktop collections will not be available until migration of all destination servers in the virtual desktop collection is complete.	Three hours or more depending on the number of virtual machines being migrated
RD Web Access	RD Web Access cannot serve connections to session	Resources that are accessed by RD Web Access and managed	One hour

Role service	Dependent role services	Impact of migration	Downtime estimates
	collections or virtual desktop collections while they are being migrated.	by the associated RD Connection Broker server will not be available. These resources include session collections and virtual desktop collections.	
RD Licensing	Remote Desktop deployments and standalone RD Session Host servers must be configured with at least one Remote Desktop license server that is available to serve licenses. If not, users cannot connect to the RD Session Host servers while they are being migrated.	Remote Desktop deployments and standalone RD Session Host servers that are configured to use the license server may not be able to receive licenses during the migration.	One hour
RD Gateway	RD Gateway	Users cannot access the network with the RD Gateway server that is being migrated. The Remote Desktop Gateway service may be slow or not available.	One hour

Additional references

- **You are here in this migration process document:** Migrate Remote Desktop Services to Windows Server 2012
- Remote Desktop Services: Prepare to Migrate
- Remote Desktop Services: Migrate Remote Desktop Services Role Services
- Remote Desktop Services: Verify the Migration

- Remote Desktop Services: Post-Migration Tasks

- Windows Server Migration forum
- Windows Server Migration Portal

Remote Desktop Services: Prepare to Migrate

This topic explains how to prepare to migrate the Remote Desktop Services role services from a source server running Windows Server 2012 or Windows Server 2012 R2 to a destination server running Windows Server 2012 R2. It assumes that you are migrating some or all of the role services, including dependencies, from a functional deployment of Remote Desktop Services.

The general preparation instructions provided in this topic apply to the following role services in Remote Desktop Services.

- RD Connection Broker
- RD Session Host
- RD Virtualization Host
- RD Web Access
- RD Licensing
- RD Gateway

Assign permissions required to migrate Remote Desktop Services

At a minimum, you must be a member of the **Administrators** group on the source server and the destination server to install, remove, or set up Remote Desktop Services.

Migration dependencies

Remote Desktop Services role services have dependencies or prerequisites for migration, as described in this section.

Prerequisite features to migrate separately

The following features in Remote Desktop Services must be migrated separately:

- DNS Server
- Active Directory Domain Services

  Remote Desktop User Profiles are stored in Active Directory.:

    - To migrate Active Directory Domain Services, see Active Directory Domain Services and DNS Server Migration Guide
    - To deploy user profiles, see User Profiles on Windows Server 2008 R2 Remote Desktop Services
- Active Directory Certificates Services
    - If you are migrating an enterprise certification authority (CA) within the same domain, before you migrate Remote Desktop Services, follow the instructions in AD CS Migration: Migrating the Certification Authority.
    - If you are migrating an enterprise CA within the same domain, before you migrate Remote Desktop Services, follow the instructions in AD CS Migration: Migrating the Certification Authority.
    - Group Policy

      You can migrate Group Policy objects (GPOs) by using the Import Settings Wizard in the Group Policy Management Console (GPMC). For more information, see Import Settings from a Group Policy Object.
-

Prerequisite features already installed

Remote Desktop Services role services require the following roles and features in Windows Server 2012 R2. With the exception of Network Policy and Access Services (NPAS), these roles and features are installed automatically when the role service is installed, if they are not already installed on the server.

- RD Web Access requires Web Server (IIS)
- RD Virtualization Host requires Hyper-V
- RD Gateway requires Web Server (IIS), RPC over HTTP Proxy, and Network Policy and Access Services (NPAS)

Prepare your source server

To prepare your source server for migration, you need to back it up and gather data.

Back up your source server

Migrating some Remote Desktop Services role services require import or export of registry settings. You should back up the computer before working with the registry.

You can find information about backing up Windows Server 2012 and Windows Server 2012 R2 in the following topics:

- Backup and Recovery.
- Registry Editor.

Gather data from your source server

Settings for applications on the Remote Desktop Session Host server will not be gathered or recorded during this migration. Before you retire the RD Session Host server, gather and record the data that you will migrate from the source server into a data collection worksheet for each role service.

Prepare your destination server

The following steps are necessary to prepare all destination servers for the migration of Remote Desktop Services role services.

Hardware requirements for the destination server

Verify that the computer meets the hardware requirements for the role service and its prerequisites. Minimally, you should migrate to servers with comparable memory, disk space, processors, and GPUs.

The RD Virtualization Host server must meet the hardware requirements for the Hyper-V server role. For more information about Hyper-V hardware requirements, see Hardware Considerations.

RD Session Host, RD Web Access, and RD Virtualization Host cannot run on virtual machines.

Software requirements for the destination server

Remote Desktop Services is a server role in Windows Server 2012 R2. Windows Server 2012 R2 must be installed on the destination server before you migrate any of the Remote Desktop Services role services.

RD Session Host, RD Virtualization Host, RD Connection Broker, and RD Web Access require that the name of the destination server is the same as the name of the source server.

Other servers and client computers in the enterprise

Within the domain, if the destination server has the same name as the source server, no preparations are needed on other computers in the deployment.

To migrate Remote Desktop Services role services across domains, RD Session Host, RD Virtualization Host, RD Connection Broker, and RD Web Access must have accounts with permissions to join the new domain.

When you migrate RD Gateway and Remote Desktop license servers across domains, domain trust relationships are required.

Additional references

- Migrate Remote Desktop Services to Windows Server 2012 R2
- **You are here in this migration process document ->** Remote Desktop Services: Prepare to Migrate

- Remote Desktop Services: Migrate Remote Desktop Services Role Services
- Remote Desktop Services: Verify the Migration
- Remote Desktop Services: Post-Migration Tasks

- Windows Server Migration forum
- Windows Server Migration Portal

Remote Desktop Services: Migrate Remote Desktop Services Role Services

Migration for a Remote Desktop Services deployment is supported from source servers running Windows Server 2012 or Windows Server 2012 R2 to destination servers running Windows Server 2012 R2. Migration from any other major or minor releases to Windows Server 2012 R2 is not supported.

Following are the steps for migrating a Remote Desktop Services deployment:

1. Migrate the RD Connection Broker server
2. Migrate session collections
3. Migrate virtual desktop collections
4. Migrate RD Web Access servers
5. Migrate RD Gateway servers
6. Migrate RD Licensing servers
7. Migrate standalone Remote Desktop Services servers
8. Migrate certificates

Migrate the RD Connection Broker server

This is the first and most important step for migrating to a destination server running Windows Server 2012 R2.

- The Remote Desktop Connection Broker (RD Connection Broker) destination server must be configured for high availability to support migration.

  For more information, see RD Connection Broker High Availability in Windows Server 2012.
- If you have more than one RD Connection Broker server in the high availability setup, remove all the RD Connection Broker servers except the one that is currently active.
- Upgrade the remaining RD Connection Broker server to Windows Server 2012 R2.
- After the server is upgraded, add it to the high availability deployment.

Notes

A mixed high availability configuration with Windows Server 2012 and Windows Server 2012 R2 is not supported for RD Connection Broker servers.

An RD Connection Broker running Windows Server 2012 R2 can serve session collections with RD Session Host servers running Windows Server 2012, and it can serve virtual desktop collections with RD Virtualization Host servers running Windows Server 2012.

Migrate session collections

Follow these steps to migrate a session collection in Windows Server 2012 to a session collection in Windows Server 2012 R2.

Important
Migrate session collections only after successfully completing the previous step, Migrate the RD Connection Broker server.

1. Upgrade the session collection from Windows Server 2012 to Windows Server 2012 R2.
2. Add the new RD Session Host server running Windows Server 2012 R2 to the session collection.

   Tip
   Use drain mode when you are setting the RD Session Host servers. For more information about drain mode, see Introducing Terminal Services Server Drain Mode.
3.
4. Sign out of all sessions in the RD Session Host servers, and remove the servers that require migration from the session collection.

Notes
If the UVHD template (UVHD-template.vhdx) is enabled in the session collection and the file server has been migrated to a new server, update the **User Profile Disks: Location** collection property with the new path. The User Profile Disks must be available at the same relative path in the new location as they were on the source server.

A session collection of RD Session Host servers with a mix of servers running Windows Server 2012 and Windows Server 2012 R2 is not supported.

Migrate virtual desktop collections

Follow these steps to migrate a virtual desktop collection from a source server running Windows Server 2012 to a destination server running Windows Server 2012 R2.

Important
Migrate virtual desktop collections only after successfully completing the previous step, Migrate the RD Connection Broker server.

1. Upgrade the virtual desktop collection from the server running Windows Server 2012 to Windows Server 2012 R2.
2. Add the new Windows Server 2012 R2 RD Virtualization Host servers to the session collection.

Tip

Use drain mode when you set the RD Session Host servers that need to be migrated.

3.
4. Migrate all virtual machines in the current virtual desktop collection that are running on RD Virtualization Host servers to the new servers.
5. Remove all RD Virtualization Host servers that required migration from the virtual desktop collection in the source server.

Notes

If the UVHD template (UVHD-template.vhdx) is enabled in the session collection and the file server has been migrated to a new server, update the **User Profile Disks: Location** collection property with the new path. The User Profile Disks must be available at the same relative path in the new location as they were on the source server.

A session collection of RD Session Host servers with a mix of servers running Windows Server 2012 and Windows Server 2012 R2 is not supported.

Migrate RD Web Access servers

To migrate the RD Web Access servers, see Remote Desktop Web Access Role Service Migration.

Migrate RD Gateway servers

To migrate the RD Gateway Servers, see Remote Desktop Gateway Role Service Migration.

Migrate RD Licensing servers

Follow these steps to migrate an RD Licensing server from a source server running Windows Server 2012 or Windows Server 2012 R2 to a destination server running Windows Server 2012 R2.

1. Migrate the Remote Desktop Services client access licenses (RDS CALs) from the source server to the destination server.

   For more information, see Migrate Remote Desktop Services Client Access Licenses (RDS CALs).
2. Use the Deployment Properties Wizard to list the new RD Licensing servers on the server running Windows Server 2012 R2.
3. Remove the RDS CALs from the source RD Licensing server.

   For more information, see Remove Remote Desktop Services Client Access Licenses.
4. Remove the source RD Licensing servers from the deployment.

Migrate standalone Remote Desktop Services servers

The following list contains the complete migration guides for each role service. Each guide include information about preparing to migrate, verifying the migration, and post-migration tasks:

- Remote Desktop Session Host Role Service Migration
- Remote Desktop Virtualization Host Role Service Migration
- Remote Desktop Web Access Role Service Migration
- Remote Desktop Licensing Role Service Migration
- Remote Desktop Gateway Role Service Migration

Migrate certificates

Migrating certificates simply requires updating certificate information in **Deployment Properties: Manage certificates**

Remote Desktop Services features that use certificates

Although this guide does not describe how to migrate Remote Desktop Services features, the following list of features that use certificates is included for reference. Each of the following features uses certificates for at least one role service:

- Single sign-on (SSO) for RemoteApp and Desktop Connection
- Web Single Sign-On (Web SSO)
- HTTPS connections to RD Web Access
- Digital signing of Remote Desktop Protocol (.rdp) files for personal virtual desktops and virtual desktop pools
- Digital signing of Remote Desktop Protocol files for Remote App programs
- RD Gateway connections to Remote Desktop Services
- RD Session Host server connections in a farm configuration

Preparing certificates for migration

In most cases, the migration of certificates for Remote Desktop Services requires you to export the certificate with the private key. After export, you should store the certificate in a safe location.

A certificate with a private key can be migrated by using the following steps:

1. To export the certificate to a PFX file, see Export a certificate with the private key.
2. To import the certificate from a PFX file, see Import a certificate.

After you have imported the certificate to the certificate store on the destination server, follow the instructions for configuring the certificate for the specific role service.

Additional references

- Migrate Remote Desktop Services to Windows Server 2012 R2
- Remote Desktop Services: Prepare to Migrate
- **You are here in this migration process document ->**Remote Desktop Services: Migrate Remote Desktop Services Role Services
- Remote Desktop Services: Verify the Migration
- Remote Desktop Services: Post-Migration Tasks

- Windows Server Migration forum
- Windows Server Migration Portal

Remote Desktop Services: Verify the Migration

Verifying the destination server configuration is best done by running a pilot program. Use an Administrator account and an account for a valid user.

Run a pilot program

We recommend that you create a pilot program in the production environment to ensure that the migration of all role services was successful. Run the program on the servers before you put the migrated role services into production to verify that your deployment works as you expect. Depending on the role service that you migrated, you should limit connections at first, and slowly increase the number of users or connections.

Additional references

- Migrate Remote Desktop Services to Windows Server 2012 R2
- Remote Desktop Services: Prepare to Migrate
- Remote Desktop Services: Migrate Remote Desktop Services Role Services
- **You are here in this migration process document ->**Remote Desktop Services: Verify the Migration
- Remote Desktop Services: Post-Migration Tasks

- Windows Server Migration forum
- Windows Server Migration Portal

Remote Desktop Services: Post-Migration Tasks

This topic contains information about general post-migration tasks that you can perform after you migrate Remote Desktop Services role services from a source server running Windows Server 2012 or Windows Server 2012 R2 to a destination server running Windows Server 2012 R2.

The post-migration tasks include:

1. Retire the source servers

Retire the source servers

In each case, the source server is retired by removing it from the domain. After you complete and verify the migration, the source server can be shut down or disconnected from the network.

Caution
If there is more than one role service on the server, after removing the source server from the domain, you will not have access to the other role services on the computer.

- Migrate Remote Desktop Services to Windows Server 2012 R2
- Remote Desktop Services: Prepare to Migrate
- Remote Desktop Services: Migrate Remote Desktop Services Role Services
- Remote Desktop Services: Verify the Migration
- **You are here in this migration process document ->**Remote Desktop Services: Post-migration Tasks

- Windows Server Migration forum
- Windows Server Migration Portal

Migrate Cluster Roles to Windows Server 2012 R2

This guide provides step-by-step instructions for migrating clustered services and applications to a failover cluster running Windows Server 2012 R2 by using the Copy Cluster Roles Wizard. Not all clustered services and applications can be migrated using this method. This guide describes supported migration paths and provides instructions for migrating between two multi-node clusters or performing an in-place migration with only two servers. Instructions for migrating a highly available virtual machine to a new failover cluster, and for updating mount points after a clustered service migration, also are provided.

Operating system requirements for clustered roles and feature migrations

The Copy Cluster Roles Wizard supports migration to a cluster running Windows Server 2012 R2 from a cluster running any of the following operating systems:

- Windows Server 2008 R2 with Service Pack 1 (SP1)
- Windows Server 2012
- Windows Server 2012 R2

Migrations are supported between different editions of the operating system (for example, from Windows Server Enterprise to Windows Server Datacenter), between x86 and x64 processor architectures, and from a cluster running Windows Server Core or the Microsoft Hyper-V Server R2 operating system to a cluster running a full version of Windows Server.

The following migrations scenarios are not supported:

- Migrations from Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008 to Windows Server 2012 R2 are not supported. You should first upgrade to Windows Server 2008 R2 SP1 or Windows Server 2012, and then migrate the resources to Windows Server 2012 R2 using the steps in this guide. For information about migrating to a Windows Server 2012 failover cluster, see Migrating Clustered Services and Applications to Windows Server 2012. For information about migrating to a Windows Server 2008 R2 failover cluster, see Migrating Clustered Services and Applications to Windows Server 2008 R2 Step-by-Step Guide.
- The Copy Cluster Roles Wizard does not support migrations from a Windows Server 2012 R2 failover cluster to a cluster with an earlier version of Windows Server.

Important

Before you perform a migration, you should install the latest updates for the operating systems on both the old failover cluster and the new failover cluster.

Target audience

This migration guide is designed for cluster administrators who want to migrate their existing clustered roles, on a failover cluster running an earlier version of Windows Server, to a Windows Server 2012 R2 failover cluster. The focus of the guide is the steps required to successfully migrate the clustered roles and resources from one cluster to another by using the Copy Cluster Roles Wizard in Failover Cluster Manager.

General knowledge of how to create a failover cluster, configure storage and networking, and deploy and manage the clustered roles and features is assumed.

It is also assumed that customers who will use the Copy Cluster Roles Wizard to migrate highly available virtual machines have a basic knowledge of how to create, configure, and manage highly available Hyper-V virtual machines.

What this guide does not provide

This guide does not provide instructions for migrating clustered roles by methods other than using the Copy Cluster Roles Wizard.

This guide identifies clustered roles that require special handling before and after a wizard-based migration, but it does not provide detailed instructions for migrating any specific role or feature. To find out requirements and dependencies for migrating a specific Windows Server role or feature, see Migrate Roles and Features to Windows Server 2012 R2.

This guide does not provide detailed instructions for migrating a highly available virtual machine (HAVM) by using the Copy Cluster Roles Wizard. For a full discussion of migration options and requirements for migrating HAVMs to a Windows Server 2012 R2 failover cluster, and step-by-step instructions for performing a migration by using the Copy Cluster Roles Wizard, see Hyper-V: Hyper-V Cluster Migration.

Planning considerations for migrations between failover clusters

As you plan a migration to a failover cluster running Windows Server 2012 R2, consider the following:

- For your cluster to be supported by Microsoft, the cluster configuration must pass cluster validation. All hardware used by the cluster should be Windows logo certified. If any of your hardware does not appear in the Windows Server Catalog in hardware certified for Windows Server 2012 R2, contact your hardware vendor to find out their certification timeline.

  In addition, the complete configuration (servers, network, and storage) must pass all tests in the Validate a Configuration Wizard, which is included in the Failover Cluster Manager snap-in. For more information, see Validate Hardware for a Failover Cluster.
- Hardware requirements are especially important if you plan to continue to use the same servers or storage for the new cluster that the old cluster used. When you plan the migration, you should check with your hardware vendor to ensure that the existing storage meets certification requirements for use with Windows Server 2012 R2. For more information about hardware requirements, see Failover Clustering Hardware Requirements and Storage Options.
- The Copy Cluster Roles Wizard assumes that the migrated role or feature will use the same storage that it used on the old cluster. If you plan to migrate to new storage, you must copy or move of data or folders (including shared folder settings) manually. The wizard also does not copy any mount point information used in the old cluster. For information about handling mount points during a migration, see Cluster Migrations Involving New Storage: Mount Points.
- Not all clustered services and features can be migrated to a Windows Server 2012 R2 failover cluster by using the Copy Cluster Roles Wizard. To find out which clustered services and applications can be migrated by using the Copy Cluster Roles Wizard, and operating system requirements for the source failover cluster, see Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2.

Migration scenarios that use the Copy Cluster Roles Wizard

When you use the Copy Cluster Roles Wizard for your migration, you can choose from a variety of methods to perform the overall migration. This guide provides step-by-step instructions for the following two methods:

- **Create a separate failover cluster running Windows Server 2012 and then migrate to that cluster**. In this scenario, you migrate from a multi-node cluster running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2. For more information, see Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2.
- **Perform an in-place migration involving only two servers**. In this scenario, you start with a two-node cluster that is running Windows Server 2008 R2 SP1 or Windows Server 2012, remove a server from the cluster, and perform a clean installation (not an upgrade) of Windows Server 2012 R2 on that server. You use that server to create a new one-node failover cluster running Windows Server 2012 R2. Then you migrate the clustered services and applications from the old cluster node to the new cluster. Finally, you evict the remaining node from the old cluster, perform a clean installation of Windows Server 2012 R2 and add the Failover Clustering feature to that server, and then add the server to the new failover cluster. For more information, see In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2.

Note

We recommend that you test your migration in a test lab environment before you migrate a clustered service or application in your production environment. To perform a successful migration, you need to understand the requirements and dependencies of the service or application and the supporting roles and features in Windows Server in addition to the processes that this migration guide describes.

In this guide

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2

Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2

In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2

Cluster Migrations Involving New Storage: Mount Points

Additional References

Related references

What's New in Failover Clustering in Windows Server 2012 R2

Failover Clustering Overview

Failover Clustering Hardware Requirements and Storage Options

Create a Failover Cluster

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2

This topic provides guidance for migrating specific cluster roles to a failover cluster running the Windows Server 2012 R2 operating system by using the Copy Cluster Roles Wizard in Failover Cluster Manager. The topic covers supported migration paths, provides an overview of wizard-based migration, and notes which cluster roles require special handling during migration.

Migration paths for specific migrations

The following table lists the operating system versions on a source failover cluster that can be migrated to a failover cluster running Windows Server 2012 R2 for each clustered service or application. Migrations between failover clusters created with physical computers and failover clusters that are created from virtual machines (also known as *guest clusters*) are supported.

Supported migrations for clustered roles and resources to a Windows Server 2012 R2 failover cluster

Clustered role or resource	From Windows Server 2008 R2 SP1	From Windows Server 2012	From Windows Server 2012 R2
Cluster Registry settings	Yes	Yes	Yes
Cluster Shared Volume (CSV) volumes	Yes	Yes	Yes
DFS Namespace (DFS-N)	Yes	Yes	Yes
DFS Replication (DFS-R)	Yes	Yes	Yes
DHCP Server	Yes	Yes	Yes
Distributed Network Name (DNN)	No	Yes	Yes
File Server	Yes	Yes	Yes
Scale-Out File Server for application data	No	Yes	Yes
Generic Application	Yes	Yes	Yes
Generic Script	Yes	Yes	Yes
Generic Service	Yes	Yes	Yes

Clustered role or resource	From Windows Server 2008 R2 SP1	From Windows Server 2012	From Windows Server 2012 R2
Virtual Machine	Yes	Yes	Yes
Hyper-V Replica Broker	No	Yes	Yes
IP addresses (IPV4, IPV6, IPv6 tunnel addresses)	Yes	Yes	Yes
iSCSI Target Server	Yes	Yes	Yes
Internet Storage Name Service (iSNS)	Yes	Yes	Yes
Message Queuing (MSMQ), MSMQ triggers	Yes	Yes	Yes
Network Name resources	Yes	Yes	Yes
NFS shares	Yes	Yes	Yes
Other Server	Yes	Yes	Yes
Physical Disk resource	Yes	Yes	Yes
WINS Server	Yes	Yes	Yes

Note

In Windows Server 2012 R2, you can designate a virtual hard disk (.vhdx file) as shared storage for multiple virtual machines that are configured as a guest failover cluster. This new type of guest cluster, known as a *shared VHDX guest cluster*, enables scenarios such as Microsoft SQL Server Failover Cluster Instance (FCI) guest clusters. The Copy Cluster Roles Wizard supports migration of the roles in the table above (except for the Virtual Machines role, which cannot exist in a guest cluster) between shared-VHDX guest clusters running the released version of Windows Server 2012 R2. However, if you created a shared-VHDX guest cluster in Windows Server 2012 R2 Preview, you cannot use the wizard to copy the cluster roles to a shared VDX guest cluster running the released version of Windows Server 2012 R2.

Cluster roles that cannot be migrated

Some services and applications that can run in a failover cluster on Windows Server 2012 R2 cannot be migrated by using the Copy Cluster Roles Wizard—in some cases because they were not supported on earlier versions of clustering. The Copy Cluster Roles Wizard in Windows Server 2012 R2 cannot be used to migrate the following clustered roles:

- Microsoft SQL Server - For upgrade guidance for SQL Server, see the whitepaper SQL Server 2012 Upgrade Technical Guide.
- Microsoft Exchange Server - For upgrade guidance for Exchange Server, see Understanding Upgrade to Exchange 2010.
- Print Spooler from Windows Server 2008 R2 – In Windows Server 2012 R2 and Windows Server 2012, the print spooler is no longer a clustered resource. Instead, high availability is defined as a highly available virtual machine running on a single cluster node. The Print Server role is installed on a single virtual machine, which can be migrated to other nodes automatically or manually. For more information, see High Availability Printing Overview.
- Remote Desktop Connection Broker from Windows Server 2008 R2 - In Windows Server 2012 R2 and Windows Server 2012, the active/passive clustering model for the RD Connection Broker role service, used in earlier versions of Windows Server, is replaced by the Active/Active Broker feature, which eliminates the need for clustering and provides a fully active/active model. For more information, see the blog entry RD Connection Broker High Availability in Windows Server 2012.
- Volume Shadow Copy Service tasks
- Task Scheduler tasks (Windows Server 2012 R2 and Windows Server 2012 only)
- Cluster Aware Updating (CAU) settings (Windows Server 2012 R2 and Windows Server 2012 only)

Roles restricted to a single instance per cluster

For the following roles, only one instance per failover cluster is supported:

- DHCP Server
- WINS Server
- iSCSI Target Server
- Hyper-V Replica Broker (Windows Server 2012 R2 and Windows Server 2012 only)

For those roles, the Copy Cluster Roles Wizard will not attempt to create a second role instance if one instance already exists on the target cluster.

Migrations for which the Copy Cluster Roles Wizard performs most or all steps

For the following clustered services or applications, The Copy Cluster Roles Wizard performs most or all steps for a migration to a Windows Server 2012 R2 failover cluster:

- Distributed File System (DFS) Namespace

- Generic Application
- Generic Script
- Generic Service
- IPv4 Address, when migrating within the same subnet
- IPv6 Address or IPv6 Tunnel Address
- Internet Storage Name Service (iSNS)
- Network Name (other than the cluster name)

  If Kerberos authentication is enabled for the Network Name resource, the migration wizard prompts you for the password for the Cluster service account that is used by the old cluster.
- NFS
- Physical Disk (resource settings only; does not copy data to new storage)
- Windows Internet Name Service (WINS) (Extra steps might be required if you migrate to new storage, and you use a different drive letter on the path to the new database.)

For more information about the Copy Cluster Roles Wizard, see Create a Failover Cluster. For step-by-step instructions for performing a migration between two multimode failover clusters, see Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2. For step-by-step instructions for performing a stand-alone migration while upgrading a single failover cluster, see In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2.

Migration within mixed environments

The Copy Cluster Roles Wizard can migrate clustered resources within mixed environments. For example, the wizard accommodates the following differences in the source and destination environments:

- Migrate static IP addresses to a cluster using DHCP.
- Migrate IPv4 resources into an IPv6 environment.
- Migrate across routed subnets.
- Migrate a physical cluster to a guest (virtual) cluster (with the exception of Hyper-V clusters, which must run on physical computers).
- Migrate between different editions of the operating system (for example, from Windows Server Enterprise to Windows Server Datacenter), between x86 and x64 processor architectures, and from a cluster running Windows Server Core or Microsoft Hyper-V Server to a cluster running a full version of Windows Server.

During migration, the wizard allows you to address name conflicts between resource groups, resources, and share names and to address drive letter collisions. The wizard resolves the conflicts as part of the post-migration repair process.

Important

The Copy Cluster Roles Wizard moves resources, not data. If you plan to migrate to new storage, you must move the data and folders yourself.

Additional steps for a wizard-based migration

Some additional steps typically are needed before or after you run the wizard, including the following:

- Install server roles and features that are needed in the new cluster. In most cases, you must install the role or feature on all nodes of the cluster. For example, before you migrate a highly available virtual machine, you must install the Hyper-V server role on each cluster node.
- Copy or install any associated applications, services, or scripts on the new cluster (all nodes).
- If a migrated role or feature uses the same storage, take the services and storage offline on the old cluster and then make the storage available to the new cluster.
- If a migrated role or feature uses new storage, ensure that any data and folders are copied to new storage. Verify permissions on any shared subfolders that were migrated.
- If the new cluster is on a different subnet, provide static IP addresses.
- If the new cluster uses a different volume letter, update drive path locations for applications.
- Configure Task Manager tasks on the new cluster. (Windows Server 2012 R2 or Windows Server 2012 only)
- For a virtual machine, install the latest integration services on the virtual machine. Configure Volume Shadow Copy Service (VSS) backups. For a migration from Windows Server 2012 R2 or Windows Server 2012, configure Hyper-V Replica settings.
- Configure Cluster Aware Updating (CAU). (Windows Server 2012 R2 and Windows Server 2012 only)

Failover Cluster Copy Roles reports

The wizard provides a Failover Cluster Copy Roles Pre-Copy Report (formerly the Pre-Migration Report) and a Failover Cluster Copy Roles Post-Copy Report (formerly the Post-Migration Report), which provide important information. We recommend that you review both reports while performing a migration:

- The Pre-Copy Roles Report explains whether each resource that you plan to migrate is eligible for migration.
- The Post-Copy Roles Report contains information about the success of the migration, and describes additional steps that might be needed before you bring the migrated resources online.

Note

Two resource groups are never migrated: **Cluster Core Resources Group** and **Available Storage Group**. You can ignore these resource groups in the Failover Cluster Copy Roles reports.

Clustered role and feature migrations that require extra steps

This section provides guidance for migrating clustered roles and features that require additional steps before or after you run the Copy Cluster Roles Wizard to perform a migration between clusters.

- Clustered DFS Replication migrations
- Clustered DHCP migrations
- Clustered DTC migrations
- Clustered File Server and Scale-out File Server migrations
- Clustered FSRM migrations
- Clustered Message Queuing (MSMQ) migrations
- Other Server migrations involving resource types not built into failover clusters
- Migration of highly available virtual machines

Clustered DFS Replication migrations

Before you migrate clustered Distributed File System (DFS) Replication (also known as DFS-R or DFSR) to a cluster running Windows Server 2012 R2, you must add the new cluster to the DFS replication group to which the old cluster belongs, and then wait until DFS Replication synchronizes the data to the new cluster. After data synchronization is complete, you can decommission the old cluster. For step-by-step guidance, see Migrate File and Storage Services to Windows Server 2012 R2 and File and Storage Services: Post-migration Tasks.

To migrate clustered instances of DFS Replication to a cluster running Windows Server 2012 R2

1. Obtain the name of the cluster to which you will migrate. In Active Directory, this is the name that is used for the computer account of the cluster itself (also called the cluster name object or CNO). Add this name to the replication group that you will migrate. For more information, see Add a member to a replication group.
2. Wait until DFS Replication finishes synchronizing the replicated data to the cluster to which you will migrate.
3. If you plan to decommission the cluster from which you migrated, remove its network name from the replication group. If necessary, destroy the cluster.

For more information about DFS Replication in Windows Server 2012 R2, see DFS Namespaces and DFS Replication Overview. For step-by-step instructions for migrating DEF Replication, see Migrate File and Storage Services to Windows Server 2012 R2.

Clustered DHCP migrations

When migrating clustered Dynamic Host Configuration Protocol (DHCP) to a cluster running Windows Server 2012 R2, the Copy Cluster Roles Wizard migrates resources and settings, but

not the DHCP database. For information about how to migrate the DHCP database, see DHCP Server Migration: Migrating the DHCP Server Role. The information in the topic also applies to migrations from Windows Server 2008 R2 or Windows Server 2012 to Windows Server 2012 R2. The topic includes information about migrating from a cluster.

Note

Although the migration of the clustered DHCP role is supported, in Windows Server 2012 R2 there is the option to use DHCP failover. DHCP failover provides redundancy and load balancing without clustered DHCP. For more information, see Migrate to DHCP Failover and Understand and Deploy DHCP Failover.

Clustered DTC migrations

Before you begin the migration of clustered Distributed Transaction Coordinator (DTC) to a cluster running Windows Server 2012 R2, you must make sure the list of transactions stored by DTC is empty. This is referred to as *draining the transaction logs*. If you do not drain the logs, the information in the logs (the transaction state information for unresolved transactions) will be lost during the migration. Unresolved transactions include **Active**, **In Doubt**, and **Cannot Notify** transactions.

To drain DTC transaction logs of unresolved transactions

1. Stop the application that creates transactions on the clustered instance of DTC that is being migrated.
2. On a node of the cluster that you are migrating from, click **Start**, point to **Administrative Tools**, and then click **Component Services**. (In Windows Server 2012 R2, open **Component Services** directly from the Start screen.)
3. Expand **Component Services**, expand **Computers**, expand **My Computer**, expand **Distributed Transaction Coordinator**, and then expand **Clustered DTCs**.
4. Expand the clustered instance of DTC that you are migrating, and then click **Transaction List**.
5. View the transaction list to see if it is empty. If there are transactions listed, then either wait for them to be completed or right-click each transaction, click **Resolve**, and then select **Forget**, **Commit**, or **Abort**.

   For information about the effect of each of these options, see Transaction State Resolution After System Failure.

For additional information, see View Transaction Information.

Clustered File Server and Scale-out File Server migrations

Several methods are available for migrating a scale-out file server or traditional clustered file server to Windows Server 2012 R2. For all methods, there are trade-offs among required downtime, migration duration, resource usage, and required hardware. The best method for your

environment depends on hardware and resources you have available, the volume of data to be moved, the number of clustered file servers that are affected, and service requirements.

Choosing the best migration method for your file server

When you plan your clustered file server migration, consider these methods:

- Virtual machine storage migration
- Copy Cluster Roles Wizard - Migrate to a new multi-node cluster
- Copy Cluster Roles Wizard – In-place migration
- Migrate storage pools

Note
For a fuller discussion of storage upgrade options as an integral part of upgrading your private cloud infrastructure, view the presentation Upgrading Your Private Cloud with Windows Server 2012 R2, presented at TechEd 2013.

Virtual machine storage migration

Introduced in Windows Server 2012, virtual machine storage migration enables you to the virtual hard disks used by one clustered file server to another clustered file server while the virtual machine remains running. This is known as *storage migration*. After you migrate storage for each virtual machine, you migrate the virtual machines to the new Windows Server 2012 R2 failover cluster. For more information, see Virtual Machine Storage Migration Overview.

This method is useful for moving to new storage if you have the resources available to maintain required service levels on all of the virtual machines during migration.

Migration method: Virtual machine storage migration

Advantages	Disadvantages
Live-migrate storage without any downtime for the virtual machines.	The process moves lots of data over the network, using lots of resources. If you are migrating a large number of virtual machines, and don't have the network capacity to gracefully handle the large loads, this can have a large impact on performance. You must move to new storage.

Copy Cluster Roles Wizard - Migrate to a new multi-node cluster

With this method, you set up a new Windows Server 2012 R2 failover cluster, migrate the File Server role to the new cluster, and then take the file server offline while you redirect storage to the new cluster. The wizard does not move data; if you migrate to new storage, the wizard updates the storage settings for the role, but you must move the data and files manually during the migration. For step-by-step instructions, see Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2.

Use this method if you have too much data to move over the network without unacceptable impact on the performance of your clustered file servers.

Migration method: Copy Cluster Roles Wizard – Migrate to a new multi-node cluster

Advantages	Disadvantages
This method is much faster than storage migration. In a large enterprise with hundreds of clustered file servers, the migration can take hours rather than days.	Downtime is required. You must take the File Server roles offline on the old cluster while you redirect the storage to the new cluster, However, this method is faster than moving VHDs over the network, and you can schedule the downtime for a maintenance window, when you will experience a limited interruption in service but will not risk degrading service for running virtual machines over long periods. Additional hardware is required to create the new failover cluster.

Copy Cluster Roles Wizard – In-place migration

If you do not have the hardware available to create a new multi-node Windows Server 2012 R2 failover cluster before you migrate the cluster roles, you can perform an in-place migration. In an in-place migration, you use hardware from an existing cluster to create the new cluster, evicting one node to use as the first node in the new cluster.

For a two-node cluster, you would evict one node, perform a clean installation of Windows Server 2012 R2 on that node, create a new single-node failover cluster with that node, and then migrate the File Server role from the old cluster to the new cluster. At that point, you must take the File Server roles offline on the old cluster while you redirect storage to the new cluster. When the migration is complete, you then destroy the old cluster, install Windows Server 2012 R2 on the other cluster node, and add that node to the new cluster. For step-by-step instructions, see In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2.

Migration method: Copy Cluster Roles Wizard – In-place migration

Advantages	Disadvantages
No new hardware required. Data is not migrated over the network.	Downtime is required: you must take the File Server roles offline on the old cluster before you can redirect storage to the new cluster and then bring the roles and the storage online on the new cluster. While migrating a two-node cluster in place, you take on the added risk of losing high availability for your file servers from the time

Advantages	Disadvantages
	when you remove the first node from the old cluster until you add the second node to the new cluster. Service can be degraded on the nodes that remain online during the migration, particularly if you are migrating large numbers of clustered file servers.

Storage pool migration

If you are migrating from a Windows Server 2012 failover cluster that uses storage pools, you can minimize the impact of migration by migrating one storage pool at a time, from the old cluster to the new cluster. With storage pools, instead of managing each disk individually, you add physical disks to one or more pools and then create virtual disks from the available capacity. You then create volumes on the virtual disks, as if they were physical disks. When you run low on the available capacity in the pool, add physical disks to the pool to create bigger pools with more capacity for more virtual disks.

Storage Spaces uses commodity drives that are attached via Serial-Attached SCSI (SAS), Serial ATA (SATA), or USB. When it is time to change from the old cluster using the storage to the new cluster using the storage, you might need to change the cabling. If you are reusing hardware (that is, you are performing an in-place migration), when you evict a node from the old cluster, you need to disconnect that server's connection to the disks. When it is time to change the storage from the old cluster to the new cluster, disconnect the storage from the old cluster before you connect the storage to the new cluster, so that only one cluster is connected to the disks at one time. When you connect the storage to the new cluster, the Storage Spaces and associated storage pools becomes available to the new cluster so that the migration can complete.

For more information about using Storage Spaces and storage pools, see Storage Spaces Overview, What's New in Storage Spaces in Windows Server 2012 R2, and Deploy Clustered Storage Spaces. For a video presentation from TechEd 2013 that demonstrates Storage Spaces basics and new features, see Storage Spaces: What's New in Windows Server 2012 R2.

Migration method: Storage pool migration

Advantages	Disadvantages
No downtime is required. High availability is maintained throughout migration.	Data moves over the network. A four-node cluster is required to enable you to maintain two nodes on both the old and new clusters during migration.

Additional tasks for file server migration using the Copy Cluster Roles Wizard

If you choose to use the Copy Cluster Roles Wizard to migrate your file server, be aware of the following requirements:

- If you plan to migrate to new storage, keep in mind that if the migrated files and folder inherit permissions from their parents, during migration it is the inheritance setting that is migrated, not the inherited permissions. Therefore it is important to make sure that the parent folders on the source server and the destination server have the same permissions to maintain the permissions on migrated data that has inherited permissions. After the file server migration, it's important to verify the folder permissions after the migration. Sometimes folder permissions reset to Read-only during a file server migration.
- You do not need to migrate the quorum resource. When you run the Create a Cluster Wizard in Windows Server 2012 R2 or win8_server_2 , the cluster software automatically chooses the quorum configuration that provides the highest availability for your new failover cluster, and it dynamically updates the quorum configuration if you add or evict nodes. You can change the quorum configuration on the new cluster if necessary for your specific environment. However, Dynamic Quorum is not in effect on a Windows Server 2008 R2 failover cluster. If you evict a node to perform an in-place migration, you will need to update the quorum configuration.

Clustered FSRM migrations

To migrate the File Server Resource Manager (FSRM) classification, storage reporting, and file management task configuration on a clustered file server running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 to a failover cluster running Windows Server 2012 R2, you must export the configuration from one FSRM server node in the cluster and then import the configuration to another FSRM server. These steps must be performed locally on one node of the cluster. You then fail over the other nodes until this process is complete. For step-by-step instructions, see Migrate File and Storage Services to Windows Server 2012 R2.

Important
When you migrate the configuration, FSRM requires that you use the same drive letters on both the source and destination servers.

Clustered Message Queuing (MSMQ) migrations

When you migrate a clustered instance of Message Queuing (also known as MSMQ) to a cluster running Windows Server 2012 R2, it's important to take the following precautions to ensure that the data is preserved and you can bring the service online on the new cluster:

- Before you migrate, you should back up the data that is associated with clustered instances of Message Queuing. This ensures that you can restore service-specific Message Queuing data if it is accidentally deleted during migration. For more information about Message Queuing backup and restore, see Backing up and restoring messages.
- During the migration, it's important to make sure that the migration is complete before you delete either clustered instance of Message Queuing (old or new). Otherwise, service-specific

data for Message Queuing might be deleted from the shared storage, which prevents the remaining Message Queuing resource from coming online. After the migration is complete and you are ready to delete a clustered instance of Message Queuing (old or new), first remove the disk resource from that clustered instance and take the disk offline. Then delete the clustered instance of Message Queuing.

Other Server migrations involving resource types not built into failover clusters

Before you use the Copy Cluster Roles Wizard to migrate an application that uses a clustered resource type that is not built into failover clustering, be sure to add the resource type to the new cluster. You can then use the Copy Cluster Roles Wizard to migrate your clustered application. In this situation, the Copy Cluster Roles Wizard attempts a "best effort" migration.

To add a resource type to a failover cluster running Windows Server 2012 R2

1. Open Failover Cluster Manager from the Start screen of any node in the cluster running Windows Server 2012 R2.
2. If the cluster to which you want to migrate is not displayed, in the console tree, right-click **Failover Cluster Manager**, click **Connect to Cluster**, select the cluster that you want to migrate to, and then click **OK**.
3. In the console tree, right-click the cluster, and then click **Properties**.
4. Click the **Resource Types** tab, and then click **Add**.
5. Specify the following information for the resource type:
   - **Resource DLL path and file name**: The path and file name of the resource dynamic-link library (DLL) that the Cluster service should use when it communicates with your service or application.
   - **Resource type name**: The name that the Cluster service uses for the resource type. This name stays the same regardless of the regional and language options that are currently selected.
   - **Resource type display name**: The name that is displayed for the resource type. This name might vary when you make changes to regional and language options.

Migration of highly available virtual machines

You can use the Copy Cluster Roles Wizard to migrate highly available virtual machines created in Hyper-V from a Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 failover cluster to a cluster running Windows Server 2012 R2. Using the wizard, you migrate the Virtual Machine clustered role, select highly available virtual machines to migrate, and update virtual network settings for the virtual machines on the new cluster.

Migrating HAVMs by using the Copy Cluster Roles Wizard has the advantage of not copying VHDs over the network, so migration completes fairly quickly and downtime is limited. However, the wizard cannot migrate virtual machines to new storage. Also, if you migrate one virtual machine on a Cluster Shared Volume (CSV) volume, all virtual machines on that volume are

migrated. And downtime is required: after you copy the Virtual Machine roles to the cluster, you must take the virtual machines on the old cluster offline, mask the storage to the old cluster, unmask the storage to the new cluster, then bring the storage online on the new cluster, and then start the virtual machines on the new cluster.

Caution

It is very important that you not turn on the migrated virtual machines on the new cluster before you take the virtual machines offline on the old cluster. Running a virtual machine on both clusters at the same time might corrupt the virtual machine.

For step-by-step instructions for migrating highly available virtual machines from a Windows Server 2012 failover cluster to a Windows Server 2012 R2 failover cluster by using the Copy Cluster Roles Wizard, see Copy Cluster Roles Wizard in Hyper-V Cluster Using Separate Scale-out File Server Migration, or, if your virtual machines are stored on Cluster Shared Volume (CSV) volumes, see Hyper-V Cluster Using Cluster Shared Volumes (CSV) Migration. You can use the same procedures to migrate virtual machines from CSV volumes on a Windows Server 2008 R2 cluster to a Windows Server 2012 R2 cluster.

Alternate methods for migrating HAVMs to a Windows Server 2012 R2 failover cluster

Depending upon your environment and the service requirements for the migrated workloads, you should consider two alternate methods for migrating highly available virtual machines:

- **Cross version live migration** – Windows Server 2012 R2 introduces a new method for migrating a highly available virtual machine from a Windows Server 2012 cluster to a Windows Server 2012 R2 cluster. Using cross version live migration, you can migrate vritual machines to the new failover cluster without any downtime. If the virtual hard disks (VHDs) are stored on a Scale-out File Server share that is accessible to both clusters, you don't have to copy files over the network. However, depending on factors such as the amount of memory configured for the virtual machine, migration can be slow, and resource consumption during the live migrations can be high.
- **Export/Import method** - You also can migrate individual virtual machines by using the **Export** and **Import** actions in Hyper-V Manager (also available in Windows PowerShell). The Export/Import method lets you migrate virtual machines one at a time and control the method by which they the VHDs are copied to the new cluster. The virtual machine must be taken offline during the export and import, and you must re-enable Hyper-V replication on the virtual machine after migration.

For a comparison of migration methods for migrating HAVMs to a Windows Server 2012 R2 failover cluster, see Hyper-V: Migration Options.

Note

You must use the Copy Cluster Roles Wizard or the **Export** and **Import** actions to migrate from a Windows Server 2008 R2 cluster to a Windows Server 2012 R2 cluster. Cross version live migration is only available when you migrate from Windows Server 2012.

Additional tasks for using the Copy Cluster Roles Wizard to migrate HAVMs

When you migrate HAVMs by using the Copy Cluster Roles Wizard, a few extra steps are required:

- You must merge or discard all shadow copies before you migrate the volumes that are attached to the virtual machines. Before you begin working with shadow copies, you should back up volumes.
- After you migrate the virtual machines to the new cluster, install the latest Hyper-V integration services on the new virtual machines.
- After you migrate, The Copy Cluster Roles Wizard does not migrate the following settings. You will need to configure the settings on the new cluster after migration.
  - Hyper-V Replica settings

    Important
    If you using Hyper-V Replica with the workload that you are migrating, see the "Hyper-V Replica" section of Hyper-V: Migration Options for special considerations when migrating from Windows Server 2012 to Windows Server 2012 R2.
  - Volume Shadow-Copy Service (VSS) tasks
  - Cluster-Aware Updating (CAU) settings

Additional references

- Migrate Cluster Roles to Windows Server 2012 R2
- Windows Server Migration forum
- Failover cluster basics:
  - What's New in Failover Clustering in Windows Server 2012 R2
  - Failover Clustering Overview
- Instructions for migrations that use the Copy Cluster Role Wizard:
  - Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2
  - In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2
  - Hyper-V: Hyper-V Cluster Migration (Migrating highly available virtual machines from Windows Server 2012 or Windows Server 2012 R2)
  - Cluster Migrations Involving New Storage: Mount Points
- Migrating individual roles and features:
  - Migrating Roles and Features in Windows Server
  - Migrate File and Storage Services to Windows Server 2012 R2
- High availability for Microsoft Exchange Server 2013: Deploying High Availability and Site Resilience
- High availability for Microsoft SQL Server 2014: High Availability Solutions (SQL Server)
- High availability for Microsoft SQL Server 2012: Microsoft SQL Server AlwaysOn Solutions Guide for High Availability and Disaster Recovery (whitepaper)

Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2

This topic provides step-by-step instructions for migrating cluster roles from a multi-node failover cluster running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 to a multimode cluster running Windows Server 2012 R2. (Alternatively, you can perform an in-place migration using a single two-node cluster. For more information, see In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2.) If you plan to migrate highly available Hyper-V virtual machines (by migrating the Virtual Machine cluster role), see Hyper-V: Hyper-V Cluster Migration for step-by-step instructions.

Important
Before you begin your migration, review Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2 to confirm that the clustered service or application can be migrated by using the Copy Cluster Roles Wizard.

Overview of migration of cluster roles between two multi-node failover clusters

The procedures in this topic describe the following process for migrating cluster roles from an existing multi-node cluster to a new multi-node Windows Server 2012 R2 failover cluster.

1. Cluster roles: Prepare to migrate between two multi-node clusters
   - To prepare servers for the new cluster – Install the Windows Server 2012 R2 operating system, required server roles and features, and any services or applications that will run on the new failover cluster. Pre-test services and applications to make sure they are compatible with Windows Server 2012 R2. Verify that your storage is certified for use with Windows Server 2012 R2.
   - To prepare storage for the new cluster – Storage preparation differs slightly depending on whether you will use the same storage for the new cluster that the old cluster is using or you plan to migrate to new storage. If you migrate highly available virtual machines, some additional storage preparation is required.
   - To create the new failover cluster and configure your firewall
2. Cluster roles: Migrate the cluster roles – Use the Copy Cluster Roles Wizard in Failover Cluster Manager to migrate the cluster roles to the new cluster.
3. Cluster roles: Post-migration tasks for a migration between two multi-node clusters – Before you can bring a cluster role online on the new cluster, take the role offline on the new server, take the storage offline on the old cluster, and bring the storage on the new cluster online on the new cluster online. If you migrated highly available virtual machines, install the latest integration services on the virtual machines.
4. Cluster roles: Verify the migration:
   - To verify that the migrated cluster roles are performing as expected on the new cluster - Verify that the workload is available on the new cluster, and that service is provided at the

required service-level agreement (SLA). For virtual machines, verify the status of the virtual machines in Hyper-V Manager, and confirm that you can connect to the virtual machines by using Remote Desktop or Virtual Machine Connection.

- To verify that the migrated cluster roles can fail over successfully
- To troubleshoot issues with failover for the migrated cluster roles

Impact of a migration between two multi-node clusters

When you migrate a cluster role between two multi-node failover clusters, you can prepare the destination failover cluster, configure storage on the new cluster, and copy the cluster role to the new cluster while maintaining service availability on the old cluster. However, customers will experience a brief downtime after the cluster role is migrated and before you bring the role online on the new cluster.

If the new cluster will use the same storage that the old cluster is using, before you bring the role online on the new cluster, you must take the clustered role offline on the old cluster, mask the storage to the old cluster, unmask the storage to the new cluster, and then bring the volumes or disks online on the new cluster.

If you are migrating to new storage, before you can bring the role online on the new cluster, you must take the role offline on the old cluster, copy data and files for the clustered role to the new storage (the Copy Cluster Roles Wizard does not move data), and then bring the new storage online on the new cluster.

Warning

If you plan to use the Copy Cluster Roles Wizard to migrate a highly available Hyper-V virtual machine from a Windows Server 2008 R2 failover cluster to a Windows Server 2012 R2 failover cluster, be aware that live migration is not supported for that scenario. However, you can live migrate an HAVM from a cluster running Windows Server 2012 or Windows Server 2012 R2 to a Windows Server 2012 R2 failover cluster.

Access rights required to complete migration

To migrate a cluster role by using the Copy Cluster Roles Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

Additional references

In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2

Hyper-V: Migration Options

Hyper-V: Hyper-V Cluster Migration

Cluster Migrations Involving New Storage: Mount Points

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2

Windows Server Migration Forum

Clustering Forum for Windows Server 2012

Cluster roles: Prepare to migrate between two multi-node clusters

To prepare to migrate cluster roles to a new failover cluster, perform the following tasks:

- Prepare the servers for the new cluster
- Prepare storage for the new cluster
- Create a new failover cluster, and configure the firewall

To prepare servers for the new cluster

1. Perform a clean installation of Windows Server 2012 R2 on each server that you will add to the new failover cluster.
2. Install the Failover Clustering feature on each server.
3. If you plan to migrate highly available virtual machines, add the Hyper-V role to each server.
4. Install any needed services, applications, and server roles. For example, if you plan to migrate clustered Windows Internet Name Service (WINS) to the new cluster, install the WINS Server feature by using Server Manager.
5. If you are migrating a Generic Application, Generic Script, or Generic Service resource, confirm that any associated application is compatible with Windows Server 2012 R2. You also must confirm that any associated service exists in Windows Server 2012 R2 and has the same name that it had in the old cluster. Test the application or service (separately, not as part of a cluster) to confirm that it runs as expected.

To prepare storage for the new cluster

1. If you plan to migrate to existing storage, verify that your existing storage is certified for use with Windows Server 2012 R2.
2. Make an appropriate number of LUNs or disks accessible to the servers, and do not make those LUNs or disks accessible to any other servers. If the new cluster will use old storage, for testing purposes, you can limit the number of LUNs or disks to one or two. If the new cluster will use new storage, make as many disks or LUNs accessible to the new server as you think the cluster will need.

   Note
   We recommend that you keep a small disk or LUN available (unused by clustered services and applications) throughout the life of the cluster, so that you can always run storage validation tests without taking your services and applications offline.
3. Confirm that the intended cluster disks are visible and are formatted appropriately:
   a. On one of the servers that you plan to include in the cluster, open **Computer Management** from the Start screen, and then click **Disk Management** in the console

tree.

b. In Disk Management, confirm that the intended cluster disks are visible.

c. Check the format of any exposed volume or LUN. We recommend that you use NTFS for the format. (For a disk witness, you must use NTFS.)

4. To prepare to migrate a highly available virtual machine, you must merge or discard all shadow copies that have been created for the virtual machine:

   a. Back up the volumes that store the virtual machines.

   b. Merge or discard shadow copies for each virtual hard disk (VHD).

   c. If you are migrating virtual machines stored on a Cluster Shared Volume (CSV) volume, make sure that you want to migrate all of the virtual machines on any volume that you plan to migrate. If you migrate one virtual machine that is stored on Cluster Shared Volume (CVS) volume, the Copy Cluster Roles Wizard migrates all virtual machines on that volume. This restriction does not apply when you migrate a Scale-out File Server cluster, which does not use CSV volumes.

5. If you are using new storage, and your disk configuration uses mount points, review Cluster Migrations Involving New Storage: Mount Points to identify any additional steps that you need to perform.

To create the new failover cluster and configure your firewall

1. Create the new failover cluster. For information about how to create a Windows Server 2012 R2 failover cluster, see Create a Failover Cluster.
2. After you create the cluster, ensure that your firewall is configured appropriately. For example, if you are using Windows Firewall, and you will be sharing folders and files, use your preferred Windows Firewall interface to allow the exception for **Remote Volume Management**.

Cluster roles: Migrate the cluster roles

Use the following instructions to migrate clustered services and applications from your old cluster to your new cluster. The Copy Cluster Roles Wizard leaves most of the migrated resources offline so that you can perform additional steps before you bring them online.

Note

To migrate a cluster role by using the Copy Cluster Roles Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

Before you copy cluster roles to a new failover cluster

- If you plan to use new storage with the migrated clustered roles, before you run the Copy Cluster Roles Wizard, ensure that the storage is available to the new cluster – that is, ensure that the volumes have been added to the new cluster and that the volumes are online. This enables the wizard to update storage settings during migration.

To copy cluster roles from an existing cluster to a new cluster

1. From the **Start** screen or from **Server Manager** (**Tools**), open **Failover Cluster Manager**.
2. In the console tree, if the cluster that you created is not displayed, right-click **Failover Cluster Manager**, click **Connect to Cluster**, and then select the cluster that you want to configure.
3. In the console tree, expand the cluster that you created to see the items underneath it.
4. If the clustered servers are connected to a network that is not to be used for cluster communications (for example, a network intended only for iSCSI), then under **Networks**, right-click that network, click **Properties**, and then click **Do not allow cluster network communication on this network**. Click **OK**.
5. In the console tree, select the cluster.
6. Under **Configure**, click **Copy Cluster Roles**.

   The Copy Cluster Roles Wizard opens.
7. Read the Welcome page, and then click **Next**.
8. Specify the name or IP address of the cluster or cluster node from which you want to migrate services and applications, and then click **Next**.
9. The **Select Roles** page lists the clustered roles that can be migrated from the old cluster. The list does not contain any role that is not eligible for migration. Click **View Report** to view details in the Failover Cluster Pre-Copy Report. Then select each cluster role that you want to copy to the new cluster, and click **Next**.

   Important
   We recommend that you read the report, which explains whether each resource is eligible for migration. (The wizard also provides a report after it finishes, which describes any additional steps that might be needed before you bring the migrated resource groups online.)

   If storage is available on the new cluster, the **Specify Storage for Migration** page appears, giving you the option to migrate to new storage. If storage is not available on the new cluster, the wizard retains existing storage settings and does not display the page.

   Note
   Not all clustered roles can be migrated to new storage. For example, the wizard cannot be used to migrate highly available virtual machines (the Virtual Machine role) to new storage. For an overview of options for migrating highly available virtual machines to a Windows Server 2012 R2 failover cluster and step-by-step instructions for each migration option, see Hyper-V: Hyper-V Cluster Migration.
10. If you want to use new storage for a service or application:
    a. On the **Specify Storage for Migration** page, select the cluster disk that you want to migrate to new storage, and then click **Select Storage**.
    b. In the **Select Storage for Resource Group** dialog box, under **Available Storage in**

New Cluster, select the cluster disk that you want the service or application to use in the new cluster, and then click **OK**.

c. Repeat these steps for each cluster disk that you want to migrate to new storage. Then click **Next**.

Important
The Copy Cluster Roles Wizard does not move existing data and folders to the new storage. You must copy the folders and data manually.

11. Follow the instructions in the wizard to perform the migration. From the **Summary** page, we recommend that you read the Failover Cluster Post-Copy Roles Report, which describes any additional steps that you might need to complete before you bring the roles online. For example, if you have not already installed needed applications on the new cluster node, you might need to install them.

After the wizard completes, most migrated resources will be offline. Leave them offline at this stage.

Caution
At no time should a virtual machine be running on both the old cluster and the new cluster. A virtual machine that runs on both the old cluster and the new cluster at the same time might become corrupted. You can run a virtual machine on the old cluster while you migrate it to a new cluster with no problems; the virtual machine on the new cluster is created in a Stopped state. However, to avoid corruption, it is important that you do not turn on the virtual machine on the new cluster until after you stop the virtual machine on the old cluster.

Cluster roles: Post-migration tasks for a migration between two multi-node clusters

To complete the transition to the new cluster running Windows Server 2012 R2, perform the following steps. After you complete the transition, verify that the migrated workloads are available and are performing at the expected service levels, and verify that the cluster roles can successfully fail over within the new cluster.

To complete the transition from the old cluster to the new cluster

1. Prepare for clients to experience downtime, probably briefly.
2. On the old cluster, take the role and resource that were copied to the new cluster offline.
3. Complete the transition for the storage:
   - If the new cluster will use old storage, follow your plan for making LUNs or disks inaccessible to the old cluster and accessible to the new cluster.
   - If the new cluster will use new storage, copy the appropriate folders and data to the storage. As needed for disk access on the old cluster, bring individual disk resources online on that cluster. (Keep other resources offline, to ensure that clients cannot

change data on the disks in storage.) On the new cluster, use Disk Management to confirm that the appropriate LUNs or disks are visible to the new cluster and not visible to any other servers.

4. If the new cluster uses mount points, adjust the mount points as needed, and make each disk resource that uses a mount point dependent on the resource of the disk that hosts the mount point. For more information about mount points, see Cluster Migrations Involving New Storage: Mount Points.
5. Bring the services and resources that were copied to the new cluster online.
6. If you migrated virtual machines, install the latest integration services on each virtual machine.

Note

The Copy Cluster Roles Wizard does not migrate Volume Shadow Copy Service (VSS) tasks, Hyper-V Replica Broker settings, Task Scheduler tasks, and Cluster-Aware Updating (CAU) settings. If you were using any of these features on the old cluster, you will need to configure them on the new cluster.

Cluster roles: Verify the migration

After you complete the transition to the new failover cluster, verify that the migrated workloads are available and are performing at the expected service levels, and verify that the cluster roles can successfully fail over within the new cluster.

- Verify that the migrated cluster roles are performing as expected on the new cluster
- Verify that the migrated cluster roles can fail over successfully
- Troubleshoot issues with failover for the migrated cluster roles

To verify that the migrated cluster roles are performing as expected on the new cluster

1. Verify that you can access the workload that was migrated. For example, can you connect to a highly available file server after it is migrated? Can you see the data that the server stores?
2. Run the necessary application-specific tests to ensure that the new cluster can provide the same service levels for the migrated workload that was provided before the clustered role was migrated.
3. If you migrated virtual machines, verify the status of the virtual machines in Hyper-V Manager, and confirm that you can connect to the virtual machines by using Remote Desktop or Virtual Machine Connection.

To verify that the migrated cluster roles can fail over successfully

1. In the console tree of **Failover Cluster Manager**, click the failover cluster on which the role is running.
2. Expand **Roles**, and then click a migrated role that you want to test.
3. On the **Actions** pane, expand **Roles**, and then click the cluster role that you want to test.

To perform a basic test of failover for the copied cluster role, on the **Actions** pane, click **Move**, and then either select a node to move the role to (**Select Node** option) or move the role to the best possible node. When prompted, confirm your choice.

You can observe the status changes in the center pane of Failover Cluster Manager as the cluster role is moved.

If there are any issues with failover, use the following procedure to troubleshoot those issues.

To troubleshoot issues with failover for the migrated cluster roles

1. View events in Failover Cluster Manager. To do this, in the console tree, right-click **Cluster Events**, and then click **Query**. In the **Cluster Events Filter** dialog box, select the criteria for the events that you want to display, or to return to the default criteria, click the **Reset** button. Click **OK**. To sort events, click a heading, for example, **Level** or **Date and Time**.
2. Confirm that necessary services, applications, or server roles are installed on all nodes. Confirm that services or applications are compatible with Windows Server 2012 R2 and run as expected.
3. If you used old storage for the new cluster, use the **Validate Cluster** action in Failover Cluster Manager to rerun the Validate a Configuration Wizard and confirm the validation results for all LUNs or disks in the storage.
4. Review migrated resource settings and dependencies. If you are using new storage that includes disks that use mount points, see Cluster Migrations Involving New Storage: Mount Points.
5. If you migrated one or more Network Name resources with the Kerberos protocol enabled, confirm that the computer account for the failover cluster has **Full Control** permission for the computer accounts (computer objects) of your Kerberos protocol-enabled Network Name resources. On a domain controller, open **Active Directory Users and Computers**, and then verify the permissions for the appropriate computer accounts (computer objects).

In-Place Migration for a Two-Node Cluster: Migration to Windows Server 2012 R2

This topic provides an overview and steps for upgrading an existing failover cluster to Windows Server 2012 R2 when you have only two servers - that is, for performing an *in-place migration*.

Important

Before you begin the migration, confirm that the cluster role that you want to migrate can be migrated by using the Copy Cluster Roles Wizard, as described in Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2, and note any preparation or follow-up steps that are required for the role that is being migrated.

Note

For an alternative approach to failover cluster migration, see Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2. If you plan to migrate highly available Hyper-V virtual machines (by migrating the Virtual Machine cluster role), see Hyper-V: Hyper-V Cluster Migration for step-by-step instructions that use the Copy Cluster Roles Wizard to migrate virtual machines.

Overview of an in-place migration for a two-node cluster

The procedures in this topic describe the following process for upgrading an existing failover cluster to Windows Server 2012 R2 when only two servers are available.

1. Create a new cluster from a node in the old cluster – You will evict a node from the old cluster; upgrade that server to Windows Server 2012 R2 and install roles, features, and any needed software; prepare storage for the new cluster; and then create the Windows Server 2012 R2 failover cluster.
2. Copy the cluster roles to the new cluster – Use the Copy Cluster Roles Wizard to copy the clustered roles and features from the old cluster to the new cluster.
3. Perform post-migration tasks – Make existing data and files available to the new cluster; if you migrated to new storage, you will need to copy the data and files to the new storage location. Then bring the new cluster online, and verify that the migrated cluster roles and resources are available and are performing as expected.
4. Add the second node to the new cluster – First, you will destroy the old cluster. Then you will prepare the remaining node for the new cluster as you did the first node. Perform a complete set of cluster validation tests to validate the configurations of both nodes. Add the second node to the new cluster. Then configure quorum settings on the new cluster.
5. Verify failover for the migrated cluster roles – After you add the second node to the new cluster, you can verify that the migrated cluster roles fail over successfully, and you can troubleshoot any issues with failover.

Impact of the migration

When you perform an in-place migration on a two-node failover cluster, you can prepare the destination failover cluster, configure storage, and copy the cluster role to the new cluster while maintaining service availability. High availability is lost when you evict the first node from the old cluster to use for the new cluster, and it is not restored until you have repurposed the remaining cluster node and added it to the new cluster.

Customers will experience a brief downtime after the cluster role is migrated and before you bring the role online on the new cluster.

If the new cluster will use the same storage that the old cluster is using, before you bring the role online on the new cluster, you must take the clustered role offline on the old cluster, mask the

storage to the old cluster, unmask the storage to the new cluster, and then bring the volumes or disks online on the new cluster.

If you are migrating to new storage, before you can bring the role online on the new cluster, you must take the role offline on the old cluster, copy data and files for the clustered role to the new storage (the Copy Cluster Roles Wizard does not move data), and then bring the new storage online on the new cluster.

Access rights required to complete migration

To migrate a clustered role by using the Copy Cluster Roles Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

Additional references

Migrate Between Two Multi-Node Clusters: Migration to Windows Server 2012 R2

Hyper-V: Migration Options

Hyper-V: Hyper-V Cluster Migration

Cluster Migrations Involving New Storage: Mount Points

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 R2

Windows Server Migration forum

High Availability (Clustering) forum for Windows Server 2012

Create a new cluster from a node in the old cluster

For this phase of the migration, allow one existing server to continue running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 and the Cluster service while you prepare to migrate the cluster roles.

In this phase, you will perform the following tasks:

- Evict a node from the old cluster
- Prepare the node for the new cluster
- Prepare storage for the new cluster
- Create a new Windows Server 2012 R2 failover cluster

To evict a node from the old cluster

1. Before you evict a node from a failover cluster, take the following precautions:
    - For each clustered role that you plan to migrate, verify that there are no special requirements or procedures for removing or evicting a node from the cluster. You can evict a node from a clustered file server or a cluster with the Hyper-V role with no special preparation. However, you might need to uncluster some services or

applications before you evict a node.

- If you are migrating from Windows Server 2008 R2, you should migrate all roles, including cluster core roles, to the remaining node before you evict a node.
- To prevent any loss of application data when the node is evicted, shut down all services and applications on the cluster before you evict the node.

2. From the Start screen of either node in the cluster, open **Failover Cluster Manager**.
3. In the console tree, expand the cluster, expand Nodes, and then click the node that you want to evict to select it.
4. Right-click the node, click **More Actions**, and then click **Evict**.

To prepare the node for the new cluster

1. Perform a clean installation of Windows Server 2012 R2 on the server that you removed from the old cluster.
2. Add the Failover Clustering feature to the server.
3. If you plan to migrate highly available virtual machines, add the Hyper-V role to the server.
4. Install any other needed services, applications, and server roles. For example, if you plan to migrate clustered Windows Internet Name Service (WINS) to the new cluster, install the WINS Server feature by using Server Manager.
5. If you are migrating a Generic Application, Generic Script, or Generic Service resource, confirm that any associated application is compatible with Windows Server 2012 R2. You also must confirm that any associated service exists in Windows Server 2012 R2 and has the same name that it had in the old cluster. Test the application or service (separately, not as part of a cluster) to confirm that it runs as expected.

To prepare storage for the new cluster

1. If you plan to migrate to existing storage, verify that your existing storage is certified for use with Windows Server 2012 R2.
2. Make an appropriate number of LUNs or disks accessible to the servers, and do not make those LUNs or disks accessible to any other servers. If the new cluster will use old storage, for testing purposes, you can limit the number of LUNs or disks to one or two. If the new cluster will use new storage, make as many disks or LUNs accessible to the new server as you think the cluster will need.

   Note
   We recommend that you keep a small disk or LUN available (unused by clustered services and applications) throughout the life of the cluster, so that you can always run storage validation tests without taking your services and applications offline.

3. Confirm that the intended cluster disks are visible and are formatted appropriately:
   a. On one of the servers that you plan to include in the cluster, open **Computer Management** from the Start screen, and then click **Disk Management** in the console

tree.

b. In Disk Management, confirm that the intended cluster disks are visible.

c. Check the format of any exposed volume or LUN. We recommend that you use NTFS for the format. (For a disk witness, you must use NTFS.)

4. To prepare to migrate a highly available virtual machine, you must merge or discard all shadow copies that have been created for the virtual machine:

   a. Back up the volumes that store the virtual machines.

   b. Merge or discard shadow copies for each virtual hard disk (VHD).

   c. If you are migrating virtual machines stored on a Cluster Shared Volume (CSV) volume, make sure that you want to migrate all of the virtual machines on any volume that you plan to migrate. If you migrate one virtual machine that is stored on Cluster Shared Volume (CVS) volume, the Copy Cluster Roles Wizard migrates all virtual machines on that volume. This restriction does not apply when you migrate a Scale-out File Server cluster, which does not use CSV volumes.

5. If you are using new storage and your disk configuration uses mount points, review Cluster Migrations Involving New Storage: Mount Points to identify any additional steps that you need to perform.

To create the new failover cluster and configure your firewall

1. Create the new failover cluster. For information about how to create a Windows Server 2012 R2 failover cluster, see Create a Failover Cluster.
2. After you create the cluster, ensure that your firewall is configured appropriately. For example, if you are using Windows Firewall, and you will be sharing folders and files, use your preferred Windows Firewall interface to allow the exception for **Remote Volume Management**.

Copy the cluster roles to the new cluster

Use the following instructions to copy cluster roles from your old one-node cluster to your new one-node cluster. The Copy Cluster Roles Wizard leaves most of the migrated resources offline so that you can perform additional steps before you bring them online.

Before you copy cluster roles to a new failover cluster

- If you plan to use new storage with the migrated clustered roles, before you run the Copy Cluster Roles Wizard, ensure that the storage is available to the new cluster – that is, ensure that the volumes have been added to the new cluster and that the volumes are online. This enables the wizard to update storage settings during migration.

To copy cluster roles from an existing cluster to a new cluster

1. From the **Start** screen or from **Server Manager** (**Tools**), open **Failover Cluster Manager**.
2. In the console tree, if the cluster that you created is not displayed, right-click **Failover**

Cluster Manager, click **Connect to Cluster**, and then select the cluster that you want to configure.

3. In the console tree, expand the cluster that you created to see the items underneath it.
4. If the clustered servers are connected to a network that is not to be used for cluster communications (for example, a network intended only for iSCSI), then under **Networks**, right-click that network, click **Properties**, and then click **Do not allow cluster network communication on this network**. Click **OK**.
5. In the console tree, select the cluster.
6. Under **Configure**, click **Copy Cluster Roles**.

   The Copy Cluster Roles Wizard opens.
7. Read the Welcome page, and then click **Next**.
8. Specify the name or IP address of the cluster or cluster node from which you want to migrate services and applications, and then click **Next**.
9. The **Select Roles** page lists the clustered roles that can be migrated from the old cluster. The list does not contain any role that is not eligible for migration. Click **View Report** to view details in the Failover Cluster Pre-Copy Report. Then select each cluster role that you want to copy to the new cluster, and click **Next**.

   Important
   We recommend that you read the report, which explains whether each resource is eligible for migration. (The wizard also provides a report after it finishes, which describes any additional steps that might be needed before you bring the migrated resource groups online.)

   If storage is available on the new cluster, the **Specify Storage for Migration** page appears, giving you the option to migrate to new storage. If storage is not available on the new cluster, the wizard retains existing storage settings and does not display the page.

   Note
   Not all clustered roles can be migrated to new storage. For example, the wizard cannot be used to migrate highly available virtual machines (the Virtual Machine role) to new storage. For an overview of options for migrating highly available virtual machines to a Windows Server 2012 R2 failover cluster and step-by-step instructions for each migration option, see Hyper-V: Hyper-V Cluster Migration.
10. If you want to use new storage for a service or application:
    a. On the **Specify Storage for Migration** page, select the cluster disk that you want to migrate to new storage, and then click **Select Storage**.
    b. In the **Select Storage for Resource Group** dialog box, under **Available Storage in New Cluster**, select the cluster disk that you want the service or application to use in the new cluster, and then click **OK**.
    c. Repeat these steps for each cluster disk that you want to migrate to new storage. Then click **Next**.

    Important

The Copy Cluster Roles Wizard does not move existing data and folders to the new storage. You must copy the folders and data manually.

11. Follow the instructions in the wizard to perform the migration. From the **Summary** page, we recommend that you read the Failover Cluster Post-Copy Roles Report, which describes any additional steps that you might need to complete before you bring the roles online. For example, if you have not already installed needed applications on the new cluster node, you might need to install them.

After the wizard completes, most migrated resources will be offline. Leave them offline at this stage.

Caution

At no time should a virtual machine be running on both the old cluster and the new cluster. A virtual machine that runs on both the old cluster and the new cluster at the same time might become corrupted. You can run a virtual machine on the old cluster while you migrate it to a new cluster with no problems; the virtual machine on the new cluster is created in a Stopped state. However, to avoid corruption, it is important that you do not turn on the virtual machine on the new cluster until after you stop the virtual machine on the old cluster.

Perform post-migration tasks

During this phase of migration, you will perform the following tasks:

- Make existing data available to the new cluster, and bring the cluster online
- Verify that the migrated cluster roles and resources are available and are performing as expected

To make existing data available to the new cluster and bring the cluster online

1. Prepare for clients to experience downtime, probably briefly.
2. On the old cluster, take the roles and resources that were copied to the new cluster offline.
3. Complete the transition of storage to the new cluster:
   - If the new cluster will use old storage, follow your plan for making LUNs or disks inaccessible to the old cluster and accessible to the new cluster.
   - If the new cluster will use new storage, copy the appropriate folders and data to the storage. As needed for disk access on the old cluster, bring individual disk resources online on that cluster. (Keep other resources offline to ensure that clients cannot change data on the disks in storage.) Then, on the new cluster node, use Disk Management to confirm that the appropriate LUNs or disks are visible to the new cluster and not visible to any other servers.
4. If the new cluster uses mount points, adjust the mount points as needed, and make each disk resource that uses a mount point dependent on the resource of the disk that hosts the mount point. For more information about mount points, see Cluster Migrations

Involving New Storage: Mount Points.

5. Bring the cluster roles and resources that were copied to the new cluster online.
6. If you migrated virtual machines, install the latest integration services on each virtual machine.

Note

The Copy Cluster Roles Wizard does not migrate Volume Shadow Copy Service (VSS) tasks, Hyper-V Replica Broker settings, Task Scheduler tasks, and Cluster-Aware Updating (CAU) settings. If you were using any of these features on the old cluster, you will need to configure them on the new cluster.

To verify that the migrated cluster roles are performing as expected on the new cluster

1. Verify that you can access the workload that was migrated. For example, can you connect to a highly available file server after it is migrated? Can you see the data that the server stores?
2. Run the necessary application-specific tests to ensure that the new cluster can provide the same service levels for the migrated workload that was provided before the clustered role was migrated.
3. If you migrated virtual machines, verify the status of the virtual machines in Hyper-V Manager, and confirm that you can connect to the virtual machines by using Remote Desktop or Virtual Machine Connection.

Add the second node to the new cluster

Use the following procedures to prepare the second node and then add it to the new cluster. As part of this process, you will run validation tests that include both servers.

- Delete migrated cluster roles, and destroy the old cluster
- Prepare the second node for the new cluster
- Validate the configuration of both cluster nodes
- Add the node to the cluster
- Verify storage for the new cluster
- Configure quorum settings on the new cluster

To delete the copied cluster roles and destroy the old cluster

1. From the Start screen, open **Failover Cluster Manager**.
2. Remove cluster roles that were copied to the new cluster:
   a. Expand the old cluster in the console tree, and then expand **Roles**.
   b. To delete a role, right-click the role, and click **Delete**.
3. To destroy the cluster, right-click the cluster in the console tree, click **More Actions**, and then click **Destroy Cluster**.

To prepare the second node for the new cluster

1. Perform a clean installation of Windows Server 2012 R2.
2. Add the Failover Clustering feature in the same way that you added it to the other server.
3. If the new cluster hosts virtual machines, add the Hyper-V role to the server.
4. Connect the newly installed server to the same networks and storage that the existing failover cluster node is connected to.
5. Install any other needed server roles, services, and applications.
6. Identify the disks or LUNs that are exposed to the new one-node failover cluster, and expose them to the newly installed server also.

   We recommend that you keep a small disk or LUN accessible to both nodes, and unused by clustered services and applications, throughout the life of the cluster. With this LUN, you can always run storage validation tests without taking your services and applications offline.

To perform a full validation of both cluster nodes

1. On either server running Windows Server 2012 R2, open **Failover Cluster Manager** from the Start screen.
2. In the console tree, confirm that **Failover Cluster Manager** is selected, and then, in the center pane, under **Management**, click **Validate Cluster**.
3. Follow the instructions in the Validate a Configuration Wizard, but this time, be sure to specify both servers (not just the existing cluster node) and specify that you want to run all tests. Then, run the tests. Because two nodes are now being tested, a more complete set of tests runs, which takes longer than testing one node.

   Important
   If any cluster role is using a disk when you start the wizard, the wizard asks whether to take that cluster role offline for testing. If you choose to take a cluster role offline, the role remains offline until the tests are complete.
4. On the **Summary** page, which appears after the tests run, review the test results:
   - Click **View Report** and view the full set of test results in the Failover Cluster Validation Report.

     Notes
     To view the results of the tests after you close the wizard, open the report on the following path:

     <*SystemRoot*>\Cluster\Reports\Validation Report <*date and time*>.mht

     where <*SystemRoot*> is the folder in which the operating system is installed (for example, C:\Windows\).
   - To view Help topics to help you interpret the results, click **More about cluster validation tests**.

5. As necessary, make changes in the configuration, and then rerun the tests.

Note

For more information about failover cluster validation tests, see Validate Hardware for a Failover Cluster.

To add the node to the cluster

1. If the new cluster is not displayed, in the console tree, right-click **Failover Cluster Manager**, click **Connect to Cluster**, and then select the new cluster.
2. Select the new cluster in the console tree. Then, on the **Actions** pane, click **Add Node**.
3. Follow the instructions in the wizard to specify the server that you want to add to the cluster. On the **Summary** page, click **View Report** to review the tasks that the wizard performed.

   After the wizard closes, you can view the report in the <*SystemRoot*>\Cluster\Reports\ folder.

   Note

   After you close the wizard, in the center pane, you might see a warning about "Node Majority." You will correct this issue when you configure quorum settings for the new cluster.

To verify storage for the new cluster

1. In the console tree of Failover Cluster Manager, select the new cluster.
2. Expand **Storage**. Then check to see if all the disks that you want to make available to the new cluster are shown, either in one of the clustered services or applications or in **Available Storage**.
3. In most cases, you need at least one disk in **Available Storage** for your next task (specifying a witness disk). If you need to add a disk, on the **Actions** pane, click **Add Disk**, and follow the steps in the wizard.

   Before you can add a disk to storage, the disk must be accessible from both nodes in the cluster. To be used for a witness disk, a disk can be a relatively small, but must be at least 512 MB.

To configure quorum settings for the new cluster

1. In the console tree of Failover Cluster Manager, right-click the new cluster, click **More Actions**, and then click **Configure Cluster Quorum Settings**.
2. Follow the instructions in the wizard to select the most appropriate quorum setting for your needs. In most cases, this is the Node Majority quorum configuration, which requires that you specify an appropriate disk (from **Available Storage**) for the witness disk. For more information about quorum settings in Windows Server 2012 R2, see Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster.

Verify failover for the migrated cluster roles

Earlier, after you copied the cluster roles to the new single-node cluster, you verified that workloads were accessible on the new cluster and that the services and applications performed as expected. Now that you have a multi-node cluster, and have configured quorum settings, you can verify failover for the migrated cluster roles.

To verify that the migrated cluster roles can fail over successfully

1. In the console tree of **Failover Cluster Manager**, click the failover cluster on which the role is running.
2. Expand **Roles**, and then click a migrated role that you want to test.
3. On the **Actions** pane, expand **Roles**, and then click the cluster role that you want to test.

   To perform a basic test of failover for the copied cluster role, on the **Actions** pane, click **Move**, and then either select a node to move the role to (**Select Node** option) or move the role to the best possible node. When prompted, confirm your choice.

   You can observe the status changes in the center pane of Failover Cluster Manager as the cluster role is moved.

If there are any issues with failover, use the following procedure to troubleshoot those issues.

To troubleshoot issues with failover for the migrated cluster roles

1. View events in Failover Cluster Manager. To do this, in the console tree, right-click **Cluster Events**, and then click **Query**. In the **Cluster Events Filter** dialog box, select the criteria for the events that you want to display, or to return to the default criteria, click the **Reset** button. Click **OK**. To sort events, click a heading, for example, **Level** or **Date and Time**.
2. Confirm that necessary services, applications, or server roles are installed on all nodes. Confirm that services or applications are compatible with Windows Server 2012 R2 and run as expected.
3. If you used old storage for the new cluster, use the **Validate Cluster** action in Failover Cluster Manager to rerun the Validate a Configuration Wizard and confirm the validation results for all LUNs or disks in the storage.
4. Review migrated resource settings and dependencies. If you are using new storage that includes disks that use mount points, see Cluster Migrations Involving New Storage: Mount Points.
5. If you migrated one or more Network Name resources with the Kerberos protocol enabled, confirm that the computer account for the failover cluster has **Full Control** permission for the computer accounts (computer objects) of your Kerberos protocol-enabled Network Name resources. On a domain controller, open **Active Directory Users and Computers**, and then verify the permissions for the appropriate computer accounts (computer objects).

Cluster Migrations Involving New Storage: Mount Points

This topic describes considerations for configuring mount points during a migration to a failover cluster running Windows Server 2012 R2 or Windows Server 2012 when the destination cluster will use new storage after the migration.

Caution

If you want to use new storage, you must copy or move the data or folders (including shared folder settings) during a migration. The wizard for migrating clustered resources does not copy data from one shared storage location to another.

The Migrate a Cluster Wizard does not migrate mount point information (that is, information about hard disk drives that do not use drive letters, but are mounted instead in a folder on another hard disk drive). However, the wizard can migrate Physical Disk Resource settings to and from disks that use mount points. The wizard also does not configure the necessary dependency between the resources for mounted disks and the resource for a host disk (the disk on which the other disks are mounted). You must configure those dependencies after the wizard completes.

When you work with new storage for your cluster migration, you have some flexibility in the order in which you complete the tasks. You must create the mount points, run the Migrate a Cluster Wizard, copy the data to the new storage, and confirm the disk letters and mount points for the new storage. After completing those tasks, configure the disk resource dependencies in Failover Cluster Manager.

A useful way to keep track of disks in the new storage is to give them labels that indicate your intended mount point configuration. For example, in the new storage, when you are mounting a new disk in a folder called **Mount1-1** on another disk, you can also label the mounted disk as **Mount1-1**. (This assumes that the label **Mount1-1** is not already in use in the old storage.) When you run the Migrate a Cluster Wizard, and you need to specify that disk for a particular migrated resource, you can select the disk labeled **Mount1-1** from the list. After the wizard completes, you can return to Failover Cluster Manager to configure the disk resource for **Mount1-1** so that it is dependent on the appropriate resource - for example, the resource for disk **F**. Similarly, you would configure the disk resources for all other disks mounted on disk F so that they depended on the disk resource for disk F.

After you run the wizard and fully configure the mounted disk, your last task is to configure the disk dependencies in Failover Cluster Manager. For each disk resource for a mounted hard disk drive, open the Properties sheet and, on the **Dependencies** tab, specify a dependency on the disk resource for the host drive (where the mounted drives reside). This ensures that the Cluster service brings the host drive online first, followed by the drives that are dependent on it.

After you configure the dependencies, you can view a dependency report. To view a dependency report, click the service or application in Failover Cluster Manager, and then, under **Actions**, click **Show Dependency Report**. The following illustration shows four mount points that are configured with the correct dependencies on the disk on which they are mounted:

Four mount points with dependencies configured

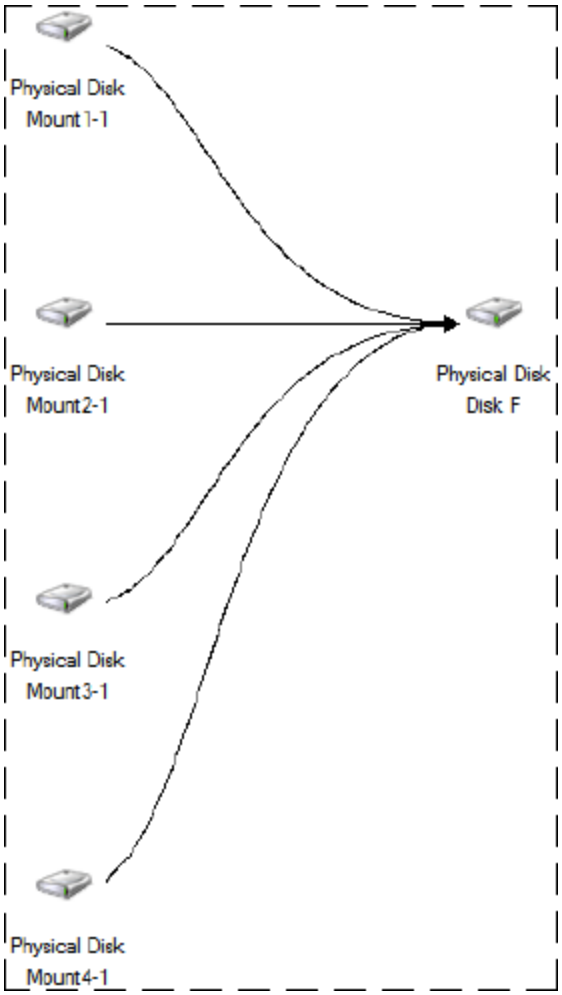

Additional references

Migrate Cluster Roles to Windows Server 2012 R2

Migrating Clustered Services and Applications to Windows Server 2012

Additional References

- Overview of failover clusters:
  - What's New in Failover Clustering in Windows Server 2012 R2
  - Failover Clustering Overview
  - Failover Clustering Hardware Requirements and Storage Options
  - Validate Hardware for a Failover Cluster

- Community resources:
  - Windows Server Migration forum
  - High Availability (Clustering) forum for Windows Server 2012
- Deploying failover clusters:
  - Create a Failover Cluster
  - Deploy an Active Directory-Detached Cluster
  - Deploy a Hyper-V Cluster
- Cluster configuration:
  - Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster
  - Use Cluster Shared Volumes in a Failover Cluster
- Migrating highly available virtual machines:
  - Migrate Hyper-V to Windows Server 2012 R2 from Windows Server 2012
  - Hyper-V: Migration Options
  - Hyper-V: Hyper-V Cluster Migration

Migrate Network Policy Server to Windows Server 2012 R2

This document provides guidance for migrating the Network Policy Server (NPS) or Internet Authentication Server (IAS) role service from an x86-based or x64-based server running Windows Server 2003, Windows Server® 2008, Windows Server® 2008 R2, or Windows Server® 2012 to a new Windows Server® 2012 server.

About this guide

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic by clicking the stars in the upper-right corner of the page (1=poor, 5=excellent), and then add comments that support your rating. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

NPS migration documentation and tools ease the migration of NPS role service settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the IAS/NPS migration process, reduce migration time, increase the accuracy of the IAS/NPS migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process.

Target audience

This guide is intended for the following IT professionals:

- IT architects responsible for computer management and security throughout an organization.
- IT operations engineers who are responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications.
- IT operations managers who are accountable for network and server management.

What this guide does not provide

This guide does not provide detailed steps to migrate the configuration of other services that might be running on the source server.

Guidance is not provided for scenarios in which the new operating system is installed on existing server hardware by using the upgrade option during setup.

Supported migration scenarios

This guide provides the instructions for migrating an existing server that is running NPS or IAS to a server that is running Windows Server 2012. This guide does not contain instructions for Network Policy Server migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available on Migrating Roles and Features in Windows Server.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

Supported operating systems

The following table displays the minimum operating system requirements that are supported by this guide.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 SP2	Windows Server 2012 R2	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012 R2	x64-based
x86- or x64-based	Windows Server® 2008	Windows Server 2012 R2	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Server 2008 R2	Windows Server 2012 R2	x64-based
x64-based	Windows Server 2012	Windows Server 2012 R2	x64-based
x64-based	Windows Server 2012 R2	Windows Server 2012 R2	x64-based

- The NPS role service is not available in Server Core editions. Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers. Windows Server Foundation edition is not available for Windows Server 2003.
- Migration from a source server to a destination server that is running an operating system with a different installed language is not supported. For example, migration of server roles from a computer that is running Windows Server 2008 with a system language of French to a computer that is running Windows Server 2012 R2 with a system language of German is not supported. The system language is the language of the localized installation package that was used to set up the Windows operating system.
- Both x86-based and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 R2 are x64-based.

Supported NPS role configurations

Migration of the following NPS settings are supported by this guide:

1. **Policies**. Migration of NPS policy configuration, including connection request policies, network policies, and health policies is supported by using this guide.
2. **Authentication methods**. All supported authentication method settings can be migrated using this guide. For more information about authentication methods, see NPS Authentication Methods (http://go.microsoft.com/fwlink/?LinkId=169629).
3. **System Health Validators (SHVs)**. Migration of SHV configuration settings implemented using Microsoft published SDK are supported.
4. **NPS templates**. Template settings are migrated using NPS UI export and import functionality. You cannot migrate template settings using the command line.
5. **RADIUS clients and remote RADIUS servers**. RADIUS clients and remote RADIUS server configuration settings, including shared secrets can be migrated using this guide.
6. **SQL accounting**. The configuration of SQL parameters, including connection, description, accounting, authentication, periodic accounting status, periodic authentication status, and max sessions settings can be migrated using this guide. It is recommended to manually configure SQL connection string settings. For more information, see Configure SQL Server Logging in NPS (http://go.microsoft.com/fwlink/?LinkId=169631).

IP address and host name configuration

This guide supports the following scenarios:

1. The destination server is configured with the same host name or IP address as source server.
2. The destination server is configured with a different host name or IP address than the source server.

Migration scenarios that are not supported

The following migration scenarios are not covered in this document:

- **Upgrade**. Guidance is not provided for scenarios in which the new operating system is installed on existing server hardware by using the **Upgrade** option during setup.
- **Extension DLLs**. This guide does not support migration of extension DLL registry key settings. For more information about extension DLL registry key migration, see Setting Up the Extension DLLs (http://go.microsoft.com/fwlink/?LinkId=169632).
- **Non-Microsoft authentication methods**. The migration of settings for non-Microsoft authentication methods is not supported. To migrate these settings, refer to your vendor documentation.
- **Non-Microsoft SHVs**. The migration of settings for non-Microsoft SHVs is supported only if the SHV is developed using guidance from the NAP SHA/SHV SDK. To migrate these settings, refer to your vendor documentation.

Overview of migration process for this role

Network Policy Server (NPS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy in Windows Server 2012 R2. NPS is the replacement for Internet Authentication Service (IAS) in Windows Server 2003.

The current topic provides an overview of the NPS migration process. The NPS migration guide also includes the following major sections:

- Prepare to Migrate
- Migrating the NPS Server
- Verifying the NPS Server Migration
- Post-Migration Tasks
- Appendix A - Data Collection Worksheet

The pre-migration process involves establishing a storage location for migration data, collection of information that will be used to perform the server migration, and operating system installation on the destination server. The NPS migration process includes using the **iasmigreader** tool if the source server is running Windows Server 2003. If the source server is running Windows Server 2008 or Windows Server 2008 R2, the Network Shell (netsh) utility is used to obtain NPS settings. When migrating a source server running Windows Server 2012 or Windows Server 2012 R2, you can use netsh or Windows PowerShell®. Procedures are then performed on the destination server to install the required roles and migrate NPS settings. Verification procedures

include testing the destination server to ensure it works correctly. Post-migration procedures include retiring or repurposing the source server.

Impact of migration

In its recommended configuration, the destination server has the same host name and IP address as the source server. In this scenario, the source server will be unavailable to process network access requests for the duration of the migration process (estimated 1-2 hours).

This guide also includes procedures for migration of the NPS server configuration from the source server to a destination server with a different host name or IP address. This allows the source and destination NPS servers to run simultaneously until all testing and verification is complete, and reduces service disruption. If you change the name or IP address of the server running NPS, RADIUS clients must also be updated with the new NPS server name and IP address.

Impact of migration on the source server

- When deploying the destination server with the same host name and IP address as the source server, the source server must be decommissioned and taken offline prior to renaming the destination server from *tempNPS* to the host name of the source server.
- When deploying the destination server with a different host name and IP address, there is no impact to the source server.

Impact of migration on other computers in the enterprise

- When deploying the destination server with the same host name and IP address, network access requests cannot be evaluated by NPS while the source server is offline and before the destination server brought online with the same name and IP address. During this time, client computers requesting access to the network cannot authenticate and are denied network access.
- When deploying the destination server with a different host name and IP address, RADIUS client settings for all network access servers that are configured to use the source server must be updated.

Permissions required to complete migration

The following permissions are required on the source server and the destination server:

- Membership in the **Administrators** group, or the equivalent, is the minimum required to install and configure server running NPS.
- Membership in the SQL database rights are required for SQL settings migration.
- If the destination server is a domain member, membership in the **Domain Admins** group, or the equivalent, is the minimum required to authorize the NPS server.

Estimated duration

The work required to migrate NPS settings from the source to destination server, including testing, can require 1 to 2 hours. Additional time may be required for migration of non-Microsoft authentication methods, SHVs or extension DLLs.

Prepare to Migrate

Migration of Network Policy Server (NPS) includes the following tasks:

- Choose a migration file storage location
- Prepare your source server
- Prepare your destination server

Complete the steps or procedures in these sections to prepare your environment for migration.

If the server running NPS will be joined to a domain, membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. If the server running NPS is not domain joined, membership in the **Administrators** group, or equivalent, is required. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Choose a migration file storage location

First, choose a location where migration files will be kept.

To choose a storage location

1. Select a file storage location where migration files will be kept. The storage location can be a network share that is accessible by both the source and destination server, or portable media that can be transferred from one server to another.

Prepare your source server

Follow these steps to prepare an x86-based or x64-based server running Windows Server 2003, Windows Server® 2008, Windows Server® 2008 R2, Windows Server® 2012, or Windows Server 2012 R2 for NPS migration.

To prepare the source server

1. Determine the domain, server name, IP address, and passwords on the source server.
2. If the source server is domain joined, determine the group membership of the source server in Active Directory Domain Services (AD DS), including security group and OU membership. This can be done using the Active Directory Users and Computers console (dsa.msc) or Server Manager on a domain controller.

Prepare your destination server

Follow these steps to prepare an x64-based destination server running Windows Server 2012 R2 for NPS migration.

Scenario 1: Prepare the destination server using the same host name and IP address

1. Install Windows Server 2012 R2 on the destination server.
2. If the source server host name is used by RADIUS clients or remote RADIUS server groups, name the destination server with a temporary server name, for example: *TempNPS*.
3. If the source server IP address is used by RADIUS clients or remote RADIUS server groups, assign a different temporary static IP address to the destination server.
4. If the source server is domain joined, add the destination server to the domain of the source server. Configure AD DS group membership settings on the destination server that are identical to the source server, including security group and OU membership.
5. Install the NPS role service using the steps provided in Install Network Policy Server (NPS) (http://go.microsoft.com/fwlink/?LinkId=169633).
6. If the source server has non-Microsoft authentication methods installed, then install same authentication methods on the destination server using your vendor documentation before importing the source server configuration.
7. If the source server has extension DLLs installed, install the same extension DLLs on the destination server before importing the source server configuration. For more information, see Setting Up the Extension DLLs (http://go.microsoft.com/fwlink/?LinkId=169632).
8. If the source server has non-Microsoft SHVs installed, then install same SHVs on the destination server using your vendor documentation before importing the source server configuration.

Scenario 2: Prepare the destination server using a different host name and IP address

1. Follow the same steps as provided for scenario 1, replacing the temporary server name with the new destination server host name, and assigning a permanent static IP address.

The destination server is now prepared for migration.

Migrating the NPS Server

This topic contains steps and procedures for migrating the Network Policy Server (NPS) role service from a legacy source server to a new x64-based destination server running Windows Server 2012 R2.

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Known issues

If you previously created conditional attributes for your remote access policy using **Called Station ID** and **Calling Station ID**, the comparison of these attributes in Windows Server 2012 R2 uses a regular expression instead of matching the exact string. For a description of these attributes, see Remote Access Policy Conditions in the **IAS Authorization** section.

Exporting settings from the source server

Use the following procedures to export the NPS settings from your x86-based or x64-based source server prior to migrating to an x64-based server running Windows Server 2012. Follow the steps in the appropriate section based on the version of Windows Server that is running on the source server:

- Exporting settings from Windows Server 2003
- Exporting settings from Windows Server 2008
- Exporting settings from Windows Server 2008 R2
- Exporting settings from Windows Server 2012 or Windows Server 2012 R2

Warning

When you use the following procedures to export configuration settings, apply appropriate precautions when moving these files from the source server to destination servers. NPS server configurations are not encrypted in the exported XML file, and contain shared secrets for RADIUS clients and members of remote RADIUS server groups. Therefore, sending these files over a network connection might pose a security risk. You can add the file to an encrypted, password protected archive file before moving the file to provide greater security. In addition, store the file in a secure location to prevent access by unauthorized users.

Exporting settings from Windows Server 2003

Configuration settings for Internet Authentication Service (IAS) in Windows Server 2003 are stored in **.MDB** files. Configuration settings for Network Policy Server (NPS) in Windows Server 2012 are stored in **.XML** files. **Iasmigreader.exe** is a command-line tool that exports the configuration settings of IAS on a computer running Windows Server 2003 to a text file. You can obtain the **iasmigreader.exe** command line migration tool for migrating Windows Server 2003 IAS settings to Windows Server 2012 from the following locations:

1. Windows Server 2012 installation media provides a copy of the migration tool in the **\sources\dlmanifests\microsoft-windows-iasserver-migplugin** directory.
2. The migration tool is available in the **%windir%\syswow64** directory on a server running Windows Server 2012.

To export settings from a source server running Windows Server 2003

1. Copy **iasmigreader.exe** to the source server into a directory configured in the **%path%**

environment variable.

Tip

To review the source server's **%path%** configuration, type **echo %path%** at a command prompt and press Enter.

2. At an elevated command prompt, type **iasmigreader.exe**, and then press Enter. The migration tool will automatically export settings to a text file.

Important

Configuration changes made to IAS will take at least one minute to be available for export.

3. IAS settings are stored in the file **ias.txt** located in the **%windir%\system32\ias** directory on the source server. If you are running a 64-bit version of Windows Server 2003, the **ias.txt** file is located in the **%windir%\syswow64\ias** directory.
4. You must manually copy SQL log configuration settings on the source server to a file (example: sql.txt).

   To record these settings:

   a. At an elevated command prompt, type **ias.msc**, and then press Enter.
   b. In the IAS console tree, click **Remote Access Logging**, right-click **SQL Server**, and then click **Properties**.
   c. Record the configuration settings on the **Settings** tab, and then click **Configure**.
   d. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.
5. Copy the ias.txt and sql.txt files to the migration store file location.

Warning

Store the ias.txt and sql.txt files in a secure location. These files contain shared secret information and SQL connection strings.

Important

When you migrate the configuration settings of the IAS role service that is running on a 32-bit or a 64-bit Windows Server 2003–based source server to the NPS role service that is running on a Windows Server 2012 R2–based destination server, the import procedure seems to complete successfully. However, the Extensible Authentication Protocol (EAP) method is misconfigured. This occurs because the migration tool generates a faulty parameter that is stored in the configuration text file (ias.txt). For more information about this issue and for a workaround, see The EAP method is configured incorrectly during the migration process from Windows Server 2003 32-bit or a 64-bit to Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=181982).

Exporting settings from Windows Server 2008

Configuration settings for NPS in Windows Server 2008 are stored in **.XML** files that can be directly imported to the destination server. The Network Shell (NetSh) command line utility can be used to export and import these settings. You can also use the Windows interface to import and export these settings.

Warning

You cannot use the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

To export settings from a source server running Windows Server 2008 using a command line

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.
2. Confirm that a message appears indicating that the export to file was successful.
3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.
4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2008 using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, open **Roles\Network Policy and Access Services\NPS**.
3. Right click **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

To record these settings:

a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

7. Copy the ias.txt and sql.txt files to the migration store file location.

Exporting settings from Windows Server 2008 R2

Configuration settings for NPS in Windows Server 2008 R2 are stored in **.XML** files that can be directly imported to the destination server. The Network Shell (NetSh) command line utility can be used to export and import these settings. You can also use the Windows interface to import and export settings.

Warning

You cannot use the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

Important

The netsh utility does not support migration of template configuration settings. To migrate these settings, you must use the Windows interface.

To export settings from a source server running Windows Server 2008 R2 using a command line

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.

2. Confirm that a message appears indicating that the export to file was successful.

3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2008 R2 using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, open **Roles\Network Policy and Access Services\NPS**.
3. Right click **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. In the console tree, right-click **Templates Management** and then click **Export Templates to a file**.
7. Next to **File name**, type **iastemplates.xml**, navigate to the migration store file location, and then click **Save**.
8. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

   a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

   b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

   c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.
9. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

Exporting settings from Windows Server 2012 or Windows Server 2012 R2

Configuration settings for NPS in Windows Server 2012 R2 are stored in **.XML** files that can be directly imported to the destination server. You can use the following methods to export and import these settings:

1. The Network Shell (NetSh) command line utility
2. The Windows interface
3. Windows PowerShell cmdlets

Warning

You cannot use Windows PowerShell, the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

Important

The netsh utility and Windows PowerShell do not support migration of template configuration settings. To migrate these settings, you must use the Windows interface.

To export settings from a source server using Windows PowerShell

1. On the source server, create a new folder for your settings (for example: C:\ConfigSettings).
2. Export your configuration settings to an .xml file in that folder, by following these steps.
   a. On the **Start** screen, type **PowerShell**, and then click Enter.
   b. To switch to the NPS context enter the following Windows PowerShell command and then press Enter:

      Import-Module NPS

   c. To export the configuration file to an .xml file, enter the following Windows PowerShell command, using the -path parameter to identify the name of the .xml file to be created and the folder into which it should be placed:

      Export-NpsConfiguration [-Path] <String>

      Tip

      For example:

      Export-NpsConfiguration –Path C:\ConfigSettings -Path nps01.xml

      Caution

      The exported file contains unencrypted shared secrets for RADIUS clients and members of remote RADIUS server groups. Because of this, you should ensure that the file is stored in a secure location to prevent malicious users from accessing the file.
3. Confirm that no errors were reported by Windows PowerShell.
4. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

   a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.
   b. Record the configuration settings on the **Settings** tab, and then click **Configure**.
   c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text

logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

5. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server using the Netsh utility

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.
2. Confirm that a message appears indicating that the export to file was successful.
3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.
4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
3. Right click the root node **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. In the console tree, right-click **Templates Management** and then click **Export Templates to a file**.
7. Next to **File name**, type **iastemplates.xml**, navigate to the migration store file location, and then click **Save**.
8. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

   a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see Appendix A - Data Collection Worksheet.

9. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

Importing settings to the destination server

Use the following procedures to import the NPS settings from your x86-based or x64-based source server to an x64-based destination server running Windows Server 2012 R2.

- Importing settings from Windows Server 2003
- Importing settings from Windows Server 2008 or Windows Server 2008 R2
- Importing settings from Windows Server 2012 or Windows Server 2012 R2

Importing settings from Windows Server 2003

The configuration file **ias.txt** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012 or Windows Server 2012 R2. If SQL accounting settings were saved, these settings are recorded manually in the **sql.txt** file.

Important

When you migrate the configuration settings of the IAS role service that is running on a 32-bit or a 64-bit Windows Server 2003–based source server to the NPS role service that is running on a Windows Server 2012 R2–based destination server, the import procedure seems to complete successfully. However, the Extensible Authentication Protocol (EAP) method is misconfigured. This occurs because the migration tool generates a faulty parameter that is stored in the configuration text file (ias.txt). For more information about this issue and for a workaround, see The EAP method is configured incorrectly during the migration process from Windows Server 2003 32-bit or a 64-bit to Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=181982).

To import settings from a source server running Windows Server 2003

1. Copy the configuration file **ias.txt** that was exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, use either netsh or Windows PowerShell to import the configuration.
   - To use netsh, do the following:

a. Open an elevated command prompt, type the following command and then press Enter:

```
netsh nps import filename="path\ias.txt"
```

Replace *path* with the directory where the **ias.txt** file is located. Verify that a message appears indicating that the import process was successful.

Tip

If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\ias.txt"**.

- To use Windows PowerShell, do the following:

a. On the **Start** screen, type **PowerShell**, and then click Enter.

b. Switch to the NPS context, enter the following Windows PowerShell command:

Import-Module NPS

c. To import the configuration, enter the following:

Import-NpsConfiguration [-Path] <String>

Replace *String* with the directory where the **ias.txt** file is located. Verify that a message appears indicating that the import process was successful.

Tip

For example:

Import-NpsConfiguration –Path c:\temp\ias.txt

3. If required, configure SQL accounting. To configure SQL accounting:
   a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
   b. Click **Accounting** and then click **Change SQL Server Logging Properties**.
   c. Manually enter SQL settings from the **sql.txt** file that you created.

Importing settings from Windows Server 2008 or Windows Server 2008 R2

The configuration file **file.xml** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012. SQL accounting settings are saved in the **sql.txt** file.

Note

For source servers running Windows Server 2008 R2: If you saved a templates configuration file, **iastemplates.xml**, you must use the Windows interface to import these settings.

To import settings from a source server running Windows Server 2008 or Windows Server 2008 R2

1. Copy the configuration files **file.xml** and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, use either netsh or Windows PowerShell to import the configuration.
   - To use netsh, do the following:

     a. Open an elevated command prompt, type the following command and then press Enter:

     ```
     netsh nps import filename="path\file.xml"
     ```

     Replace *path* with the directory where the **file.xml** file is located. Verify that a message appears indicating that the import process was successful.

     Tip

     If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\file.xml"**.

   - To use Windows PowerShell, do the following:

     a. On the **Start** screen, type **PowerShell**, and then click Enter.

     b. Switch to the NPS context, enter the following Windows PowerShell command:

     Import-Module NPS

c. To import the configuration, enter the following:

Import-NpsConfiguration [-Path] <String>

Replace <*String*> with the directory where the **file.xml** file is located.

Tip

For example:

Import-NpsConfiguration –Path c:\temp\file.xml

d. Confirm that no errors were reported by Windows PowerShell.

3. If required, configure SQL accounting. To configure SQL accounting:
   a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
   b. Click **Accounting** and then click **Change SQL Server Logging Properties**.
   c. Manually enter SQL settings from the **sql.txt** file.

Importing settings from Windows Server 2012 or Windows Server 2012 R2

The configuration file **file.xml** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012 or Windows Server 2012 R2. SQL accounting settings are saved in the **sql.txt** file. If you saved a templates configuration file, **iastemplates.xml**, you must use the Windows interface to import these settings.

To import settings from a source server

1. Copy the configuration files **file.xml** and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps import filename="path\file.xml"
   ```

   Replace *path* with the directory where the **file.xml** file is located. Verify that a message appears indicating that the import process was successful.

   Tip

   If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\file.xml"**.

   The following Windows PowerShell command performs the same function:

   ```
   Import-NpsConfiguration –Path c:\temp\file.xml
   ```

3. If required, configure SQL accounting. To configure SQL accounting:

a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
b. Click **Accounting** and then click **Change SQL Server Logging Properties**.
c. Manually enter SQL settings from the **sql.txt** file.

Using the NPS console to migrate NPS settings

You can also use the Windows interface on the destination server to import configuration settings.

To import settings from a source server using the Windows interface

1. Copy the configuration files **file.xml**, **iastemplates.xml**, and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command. If you have custom settings that were recorded using the Appendix A - Data Collection Worksheet, these must be configured manually on the destination server.
2. On the destination server, open Server Manager.
3. In the Server Manager console tree, click **ALL SERVERS**, and then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
4. To import template configuration settings, follow steps 5 to 13. If you do not have template settings, skip to step 7.
5. In the console tree, right-click **Templates Management** and then click **Import Templates from a file**.
6. Select the template configuration file **iastemplates.xml** that you copied from the source server and then click **Open**.
7. In the console tree, right-click **NPS** and then click **Import Configuration**.
8. Select the configuration file **file.xml** or **ias.txt** that you copied from the source server and then click **Open**.
9. Verify that a message appears indicating the import was successful.
10. Configure SQL accounting if required using the **sql.txt** file and the data collection worksheet. To configure SQL accounting, follow steps 11 to 13.
11. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties** in the details pane.
12. Modify the properties on the **Settings** tab if required, and then click **Configure** to enter detailed settings.
13. Using information recorded in the **sql.txt** file, enter the required settings on the **Connection** and **Advanced** tabs, and then click **OK**.

Verifying the NPS Server Migration

After the migration of your Network Policy Server (NPS) server is complete, you can perform some tasks to verify that the migration was successful.

Verifying NPS Migration

To verify the functionality of NPS on the destination server, confirm that the service is running, that the correct configuration was migrated, and that client computers can authenticate successfully.

To verify NPS migration

1. To verify that the NPS service is running on the destination server, type the following command at an elevated command prompt on the destination server and then press ENTER.

   ```
   sc query ias
   ```

   In the command output, verify that **RUNNING** is displayed next to **STATE**.
2. To verify that the source NPS configuration has been migrated to the destination server, type the following command at an elevated command prompt on the destination server and then press ENTER:

   ```
   netsh nps show config
   ```

   Verify that the destination server is not using default NPS settings. For example, default settings display a single policy under **Connection request policy configuration** with the name **Use Windows authentication for all users**.
3. To verify that the NPS console on the destination server displays the correct settings, type the following command at an elevated command prompt on the destination server and then press ENTER:

   ```
   nps.msc
   ```

   a. The NPS console will open. In the console tree, click **Accounting**, click **Change SQL Server Logging Properties**, click **Configure**, and verify that the correct settings are displayed on the **Connection** and **Advanced** tabs.

   b. In the NPS console tree, click **Policies** and then click **Connection Request Policies**, **Network Policies**, and **Health Policies**. For each type of policy, verify that the correct policies are displayed.

   c. In the NPS console tree, click **RADIUS Clients and Servers** and then click **RADIUS Clients and Remote RADIUS Server Groups**. Verify that the correct RADIUS clients and remote RADIUS server groups are displayed.

   d. In the NPS console tree, click **Network Access Protection**, and then click **System Health Validators** and **Remediation Server Groups**. Verify that the correct Network Access Protection (NAP) related settings are displayed.

e. In the NPS console tree, click **Templates Management**. If the source server was running Windows Server 2008 R2, verify that the correct templates settings are displayed.

f. In the NPS console tree, right-click **NPS**, click **Properties**, and then click the **Ports** tab. Verify that the correct **Authentication** and **Accounting** ports are displayed.

4. To verify the configuration of authentication methods, you must manually review settings in connection request policy and network policy. Certificate based EAP methods require that the proper certificate is chosen, and might require that you provision a computer certificate on the destination server.

Verifying authentication methods

a. If you use certificate based EAP methods, your destination server might already be provisioned with a suitable certificate through autoenrollment. You might also be required to manually enroll the destination server with a computer certificate. For an overview of certificate requirements for network authentication, see Network access authentication and certificates (http://go.microsoft.com/fwlink/?LinkId=169625).

b. To view certificates associated with EAP methods, click **Start**, click **Run**, type **nps.msc**, and press ENTER.

c. In the NPS console tree, open **Policies** and then open the type of policy you are using to perform authentication. For example, if the option to **Override network policy authentication settings** is enabled on the **Settings** tab in a connection request policy, then authentication is performed in connection request policy. Otherwise, authentication is performed in network policy. Authentication can be configured in both types of policies.

d. For connection request policy, double-click the policy name and then click the **Settings** tab. For network policy, double-click the policy name and then click the **Constraints** tab.

e. Click **Authentication Methods**, and then under **EAP Types** click the name of the certificate-based authentication method. For example: **Microsoft: Protected EAP (PEAP)** or Microsoft: Smart Card or other certificate.

f. Click **Edit**, verify that the correct certificate is chosen next to **Certificate issued** or **Certificate issued to**, and then click **OK**.

Note
Client computers using certificate based authentication methods must trust the certification path for this certificate.

5. To verify that client computers can authenticate using the destination server, attempt to connect to the network using client VPN connection, an 802.1X connection, or another connection that requires successful RADIUS authentication for network access.

Verifying client connections

a. To verify that client computers are successfully connecting to the network, click **Start**, click **Run**, type **eventvwr.msc**, and then press ENTER.
b. In the event viewer console tree, open **Custom Views\Server Roles\Network Policy and Access Services**.
c. In the details pane, verify under **Event ID** that event number 6272 is displayed.
d. Events 6273 or 6274 indicate that client authentication attempts are unsuccessful.
e. If no events are displayed, client connection requests are unable to reach the destination server, or the server is not logging authentication attempts.

Post-Migration Tasks

After all migration steps are complete and you have verified the migration of the Network Policy Server (NPS) role service, perform the following post-migration tasks.

Post migration tasks

After verifying NPS configuration is working on the destination server, the following steps need to be performed:

To decommission a source server using the same host and IP address

1. Remove the source server from your Active Directory domain.
2. Shut down the source server.
3. Rename the destination server from *tempNPS* to the name of the source server and configure the same static IP address as that used by the source server.
4. Perform verification steps in Verifying the NPS Server Migration with the updated host name and IP address configured on the destination server.

To decommission a source server using a different host and IP address

1. NPS server name/ IP address should be updated on Remote RADIUS servers and RADIUS clients. It requires manual update of the configurations on RADIUS clients and Network Access Servers (NAS). Please refer to your RADIUS client configuration guide for more information.
2. Perform verification steps in Verifying the NPS Server Migration.
3. When the destination server has been configured, tested, and verified, then the NPS role on the source server may be retired.

Restoring the role in the event of migration failure

If the destination server is deployed simultaneously with the source server using a different host name and IP address, then the migration can be reversed by changing RADIUS clients, remote RADIUS server groups, and network access device settings to use the source NPS server name and IP address. If the destination server is replacing the source server using the same host name and IP address, then the destination server will need to be renamed, the IP address must be updated, and the destination server must be removed from the domain to reverse the migration and bring the source server back online.

Appendix A - Data Collection Worksheet

Migration data collection worksheet

You can use this migration data collection worksheet to collect data about your source server and help ensure that the destination server functions properly after the migration.

NPS data worksheet

#	Source server essential settings	Setting values
1	**Server name** At a command prompt, type the following command, and then press ENTER. `ipconfig /all` The host name of a server is the first part of the fully qualified domain name (FQDN). The FQDN is the full computer name, including both the host name and the primary DNS suffix, separated by dots (.). For example, the FQDN of a computer named *host* with a primary DNS suffix of *example.microsoft.com* is *host.example.microsoft.com*.	Computer host name: _____ FQDN: _____
2	**Authentication, authorization, and accounting (AAA) roles** Determine what types of network access requests are validated using the RADIUS protocol on the source	Check all that apply (√) ☐ Network Access Protection (NAP) ☐ RADIUS server for dial-Up or VPN connections ☐ RADIUS server for 802.1X wireless

#	Source server essential settings	Setting values
	server.	or wired connections
3	**Text logging** Record the path and settings used for text logging. By default, local file accounting logs are stored in **%windir%\system32\LogFiles**.	Local file logging directory: _____ Format: _____ Create a new log file: _____
4	**SQL settings** Manually record any customized SQL data link properties.	Application Name: _____ Auto Translate: _____ Connect Timeout: _____ Current Language: _____ Data Source: _____ Extended Properties: _____ General Timeout: _____ Initial Catalog: _____ Initial File Name: _____ Integrated Security: _____ Locale Identifier: _____ Network Address: _____ Network Library: _____ Packet Size: _____ Password: _____ Persistent Security Info: _____

#	Source server essential settings	Setting values
		Replication server name connect option: _____ Tag with column collation when possible: _____ Use Encryption for Data: _____ Use Procedure for Prepare: _____ User ID: _____ Workstation ID: _____

Migrate Roles and Features to Windows Server 2012

Migration documentation and tools ease the process of migrating server roles, features, operating system settings, and data from an existing server that is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to a computer that is running Windows Server 2012. By using migration guides linked to on this page (and where appropriate, Windows Server Migration Tools) to migrate roles, role services, and features, you can simplify deployment of new servers (including those that are running the Server Core installation option of Windows Server 2012, and virtual servers), reduce migration downtime, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

In this section

- Install, Use, and Remove Windows Server Migration Tools
- Migrate Active Directory Federation Services Role Services to Windows Server 2012
- Migrate File and Storage Services to Windows Server 2012
- Migrate Health Registration Authority to Windows Server 2012
- Migrate Hyper-V to Windows Server 2012 from Windows 2008 R2
- Migrate IP Configuration to Windows Server 2012
- Migrate Network Policy Server to Windows Server 2012
- Migrate Print and Document Services to Windows Server 2012
- Migrate Remote Access to Windows Server 2012
- Migrate Windows Server Update Services to Windows Server 2012

- Migrating Clustered Services and Applications to Windows Server 2012

See Also

Migrating Roles and Features to Windows Server

Install, Use, and Remove Windows Server Migration Tools

Windows Server Migration Tools Installation, Access, and Removal describes how to locate, install, use, and remove Windows Server Migration Tools. Administrators can use Windows Server Migration Tools to migrate server roles, features, operating system settings, and other data and shares to computers that are running Windows Server® 2012 R2 or Windows Server® 2012.

This topic supports migrations in which the migration destination servers are running Windows Server 2012 R2 or Windows Server 2012. For information about how to prepare to use Windows Server Migration Tools for migrations to servers that are running Windows Server 2008 R2, see Windows Server Migration Tools Installation, Access, and Removal.

Windows Server Migration Tools installation and preparation can be divided into the following stages.

1. Installing Windows Server Migration Tools on destination servers that run Windows Server 2012 R2 or Windows Server 2012.
2. Creating deployment folders on migration destination servers, for copying to source servers.
3. Copying deployment folders from destination servers to source servers.
4. Registering Windows Server Migration Tools on source servers.

In this guide

Supported operating systems

Permission requirements

Prepare for installation

Install Windows Server Migration Tools

Use Windows Server Migration Tools

Remove Windows Server Migration Tools

Supported operating systems

The following table indicates the Windows Server operating systems that Windows Server Migration Tools supports.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2008, full installation option	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Windows Server 2012	Windows Server 2012 R2 or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2012	Windows Server 2012 R2 or Windows Server 2012, both full	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
		and Server Core installation options	
x64-based	Windows Server 2012 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2012 R2	Windows Server 2012 R2, both full and Server Core installation options	x64-based

The versions of operating systems shown in the previous table are the oldest combinations of operating systems and service packs that are supported. If available, newer service packs are supported.

Migrations between physical operating systems and virtual operating systems are supported. Migrations that use Windows Server Migration Tools to migrate to Windows Server 2012 or Windows Server 2012 R2 support cross-subnet migrations.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Both x86- and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2 are x64-based.

Roles that are running on the Server Core installation option of Windows Server 2008 cannot be migrated, because the Microsoft .NET Framework is not available in the Server Core installation option of Windows Server 2008.

Permission requirements

At minimum, you must be a member of the **Administrators** group on both source and destination servers to install, remove, or set up Windows Server Migration Tools.

Prepare for installation

Follow the steps in this section if you are registering Windows Server Migration Tools on migration source servers that are running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, and if the source server is running an older release of Windows Server than the migration destination server. For example, if the source server is running Windows Server 2012, but the destination server is running Windows Server 2012 R2. Otherwise, see Install Windows Server Migration Tools.

Note

All commands in this guide are case-insensitive unless specifically noted.

Windows Server 2012 source server

Complete the following tasks to prepare a source server that is running Windows Server 2012 for migration in which the destination server is running Windows Server 2012 R2.

- Verify that the source server has sufficient disk space (at least 23 MB) to store the Windows Server Migration Tools deployment folder.

Windows Server 2008 R2 source server

Complete the following tasks to prepare a source server that is running Windows Server 2008 R2 for Windows Server Migration Tools.

- Verify that the source server has sufficient disk space (at least 23 MB) to store the Windows Server Migration Tools deployment folder.

Windows Server 2008 source server

Complete the following tasks to prepare a source server that is running Windows Server 2008 for Windows Server Migration Tools.

- Verify that the source server has sufficient disk space (at least 23 MB) to store the Windows Server Migration Tools deployment folder.
- Install Windows PowerShell by using Server Manager or by running the Server Manager command prompt tool, **ServerManagerCmd.exe**. For more information about how to add features to the server by using **ServerManagerCmd.exe**, see Overview of Server Manager Commands in the Windows Server 2008 Server Manager Help.

Windows Server 2003 or Windows Server 2003 R2 source server

Complete the following tasks to prepare a source server that is running Windows Server 2003 or Windows Server 2003 R2 for Windows Server Migration Tools.

- Verify that the source server has sufficient disk space (at least 25 MB) to store the Windows Server Migration Tools deployment folder.
- Download and install Microsoft .NET Framework 2.0. Microsoft .NET Framework 2.0 is available for download from the Microsoft Web site.

- Download and install Windows PowerShell 1.0, or a later version. Windows PowerShell 1.0 is available for download from the Microsoft Web site.

Note
Windows PowerShell 2.0 and 3.0 are available in a graphically-oriented version, Windows PowerShell ISE. For more information about Windows PowerShell ISE, see Windows PowerShell 3.0 Integrated Scripting Environment (ISE).

Other computers in your enterprise

Because you might have to restart the server after you install Windows Server Migration Tools, notify users in advance that they might experience downtime while the server operating system loads. To minimize downtime, and reduce its effect on users in your enterprise, install Windows Server Migration Tools during off-peak hours.

Install Windows Server Migration Tools

This section describes how to install Windows Server Migration Tools on both source and destination servers. If both source and destination computers are running the same operating system on which Windows Server Migration Tools is available for installation (if both servers are running Windows Server 2012 R2, or both servers are running Windows Server 2012), install Windows Server Migration Tools on both computers by following installation steps in either Full installation option of Windows Server 2012 R2 or Windows Server 2012 or Server Core installation option of Windows Server 2012 R2 or Windows Server 2012.

If you plan to migrate roles, features, or other data from computers that are running older releases of Windows Server than your destination server—that is, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003—you must complete the following additional tasks after you install Windows Server Migration Tools on destination servers.

1. Create a Windows Server Migration Tools deployment folder on destination servers. For more information, see Creating a deployment folder on destination computers.
2. Register Windows Server Migration Tools on source computers that are running older releases of Windows Server than your destination server; that is, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. For more information, see Registering Windows Server Migration Tools on source computers.

For more detailed information, see Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 source computers.

Full installation option of Windows Server 2012 R2 or Windows Server 2012

To install Windows Server Migration Tools

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note
   If you are installing Windows Server Migration Tools from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
2. Type the following, and then press **Enter**. If you are installing the feature on the local server, omit the `ComputerName` parameter.

   ```
   Install-WindowsFeature Migration -ComputerName
   <computer_name>
   ```

   Note
   You can also install Windows Server Migration Tools on a full installation of Windows Server 2012 R2 or Windows Server 2012 by using the Add Roles and Features Wizard in Server Manager. For more information about how to use the Add Roles and Features Wizard, see Install or uninstall roles, role services, or features.

Server Core installation option of Windows Server 2012 R2 or Windows Server 2012

Windows PowerShell is installed by default on the Server Core installation option of Windows Server 2012 R2 and Windows Server 2012. By default, programs on the Server Core installation option run as Administrator, so there is no need to start Windows PowerShell with elevated user rights.

To install Windows Server Migration Tools on a Server Core installation of Windows Server 2012

1. Open a Windows PowerShell session by typing the following in the current command prompt session, and then press **Enter**.

   ```
   powershell.exe
   ```
2. In the Windows PowerShell session, install Windows Server Migration Tools by using the Windows PowerShell **Install-WindowsFeature** cmdlet for Server Manager. In the Windows PowerShell session, type the following, and then press **Enter**. Omit the `ComputerName` parameter if you are installing Windows Server Migration Tools on the local server.

   ```
   Install-WindowsFeature Migration -ComputerName
   <computer_name>
   ```

Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 source computers

Complete the following two tasks to install Windows Server Migration Tools.

1. Create deployment folders for source computers by running the **smigdeploy.exe** tool (included with Windows Server Migration Tools) on your destination server. For more information, see Creating a deployment folder on destination computers.
2. Register Windows Server Migration Tools on source computers that are running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 by using **SmigDeploy.exe**. For more information, see Registering Windows Server Migration Tools on source computers.

Creating a deployment folder on destination computers

This procedure describes how to create the deployment folder on your destination server that is running Windows Server Migration Tools. After you create the deployment folder, copy it to the local drive of a migration source server that is running an older release of Windows Server; that is, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.

To create a deployment folder on destination computers

1. If you have not already installed Windows Server Migration Tools on the destination server, see Install Windows Server Migration Tools in this topic.
2. Open a Command Prompt window with elevated user rights. On the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012, an elevated command prompt is already opened by default. On the full installation option, type **cmd** on the **Start** screen, right-click the **Command Prompt** tile, and then click **Run as administrator**.
3. At the command prompt, change to the directory in which the **smigdeploy.exe** tool is stored. Type the following, and then press **Enter**.

   ```
   cd %Windir%\System32\ServerMigrationTools\
   ```

4. Do one of the following to create a Windows Server Migration Tools deployment folder.
   - To create a folder to copy to an x64-based computer that is running Windows Server 2012, where Windows Server 2012 R2 is running on the destination server, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

     ```
     SmigDeploy.exe /package /architecture amd64 /os WS12 /path
     <deployment folder path>
     ```

   - To create a folder to copy to an x64-based computer that is running Windows Server 2008 R2, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

     ```
     SmigDeploy.exe /package /architecture amd64 /os WS08R2 /path
     <deployment folder path>
     ```

- To create a folder to copy to an x64-based source computer that is running Windows Server 2008, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

  ```
  SmigDeploy.exe /package /architecture amd64 /os WS08 /path <deployment folder path>
  ```

- To create a folder to copy to an x64-based source computer that is running Windows Server 2003, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

  ```
  SmigDeploy.exe /package /architecture amd64 /os WS03 /path <deployment folder path>
  ```

- To create a folder to copy to an x86-based source computer that is running Windows Server 2008, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

  ```
  SmigDeploy.exe /package /architecture X86 /os WS08 /path <deployment folder path>
  ```

- To create a folder to copy to an x86-based source computer that is running Windows Server 2003, type the following, in which *deployment folder path* represents the path of the deployment folder on the source computer, and then press **Enter**.

  ```
  SmigDeploy.exe /package /architecture X86 /os WS03 /path <deployment folder path>
  ```

Note

Each of these commands creates a deployment folder named in the format SMT_<*Operating System*>_<*Architecture*> and stores it in the specified deployment folder path.

You can also specify a network path as the path for the deployment folder. Verify that you have access to the network path before you create the deployment folder.

For more information about **SmigDeploy.exe**, at a command prompt, type **SmigDeploy.exe /?**, and then press **Enter**.

Registering Windows Server Migration Tools on source computers

Before you can run the Windows Server Migration Tools snap-in for the first time on a source server that is running an older release of Windows Server than your destination server, it must be registered with Windows PowerShell. Use **SmigDeploy.exe** to register the Windows Server Migration Tools snap-in on a migration source server that is running an older release of Windows Server than your destination server (that is, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008 or Windows Server 2003).

Before you start the procedure in this section, verify the following.

- Microsoft .NET Framework 2.0 is installed on computers that are running Windows Server 2003.

- Windows PowerShell 1.0 or a later version is installed on source computers that are running either Windows Server 2008 or Windows Server 2003. (Windows PowerShell is already installed on computers that are running Windows Server 2008 R2 and Windows Server 2012.)

To register Windows Server Migration Tools

1. Copy the Windows Server Migration Tools deployment folder that was created by using the procedure in Creating a deployment folder on destination computers to a local drive on the source computer that is running an older release of Windows Server than your destination server. Be sure that the operating system architecture of the deployment folder matches that of the source computer to which you are copying the folder.

   For example, the **SMT_WS08_amd64** folder should only be copied to the local drive of an AMD64 source computer that is running Windows Server 2008.
2. On the source computer, open a Command Prompt window.
   - On computers that are running Windows Server 2003 or the Server Core installation option of Windows Server 2008 R2, you do not have to run a Command Prompt window with elevated user rights. Click **Start**, click **Run**, type **cmd**, and then click **OK**.
   - On computers that are running the full installation options of Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008, you must open a Command Prompt window with elevated user rights. To do this, right-click the shortcut for Command Prompt, and then click **Run as Administrator**.
3. At the command prompt, change to the directory to which you copied the Windows Server Migration Tools deployment folder in step 1.

   Note

   You can register and run Windows Server Migration Tools cmdlets from a removable drive, CD-ROM, or DVD-ROM. However, to increase the reliability of registering the cmdlets, we recommend that you copy the deployment folder to a local drive of the source computer. You cannot register or run Windows Server Migration Tools cmdlets from a network location.
4. In the deployment folder directory, type the following command to register Windows Server Migration Tools cmdlets, and then press **Enter**.

   ```
   .\Smigdeploy.exe
   ```

Note

When registration is finished, a status message is displayed that indicates that the registration finished successfully, and a Windows PowerShell session opens. You can run Windows Server Migration Tools cmdlets in this Windows PowerShell session. If you close the Windows PowerShell session, see Windows Server 2003 or Windows Server 2008 source computers for information about how to access and use Windows Server Migration Tools cmdlets.

Use Windows Server Migration Tools

This section describes how to run Windows Server Migration Tools cmdlets.

- Full installation option of Windows Server 2012 R2
- Server Core installation option of Windows Server 2012 R2
- Full installation option of Windows Server 2012
- Server Core installation option of Windows Server 2012
- Source computer running full installation option of Windows Server 2008 R2
- Source computer running Server Core installation option of Windows Server 2008 R2
- Windows Server 2003 or Windows Server 2008 source computers

Full installation option of Windows Server 2012 R2

Start Windows PowerShell and run Windows Server Migration Tools cmdlets by using either of the following procedures. These can apply to either source or destination servers.

To run Windows Server Migration Tools from the Start screen

- To open a Windows Server Migration Tools custom Windows PowerShell session, right-click the **Windows Server Migration Tools** tile, and then on the app bar, click **Run as administrator**.

To run Windows Server Migration Tools in a new Windows PowerShell session

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
2. Load Windows Server Migration Tools into your Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Server Core installation option of Windows Server 2012 R2

This procedure applies to either source or destination servers.

To run Windows Server Migration Tools cmdlets

1. Type **powershell** into a command prompt, and then press **Enter**.
2. Load Windows Server Migration Tools into your Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Full installation option of Windows Server 2012

Start Windows PowerShell and run Windows Server Migration Tools cmdlets by using either of the following procedures. These can apply to either source or destination servers.

To run Windows Server Migration Tools from the Start screen

- To open a Windows Server Migration Tools custom Windows PowerShell session, right-click the **Windows Server Migration Tools** tile, and then on the app bar, click **Run as administrator**.

To run Windows Server Migration Tools in a new Windows PowerShell session

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
2. Load Windows Server Migration Tools into your Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Server Core installation option of Windows Server 2012

This procedure applies to either source or destination servers.

To run Windows Server Migration Tools cmdlets

1. Type **powershell** into a command prompt, and then press **Enter**.
2. Load Windows Server Migration Tools into your Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Source computer running full installation option of Windows Server 2008 R2

If you close the Windows PowerShell session that is opened automatically when **SmigDeploy.exe** finishes registering the Windows Server Migration Tools cmdlets, you can run Windows Server Migration Tools cmdlets by using any of the following procedures.

To run Windows Server Migration Tools from the Start menu

- To open a Windows Server Migration Tools custom Windows PowerShell session, click **Start**, point to **Administrative Tools**, open the **Windows Server Migration Tools** folder, right-click **Windows Server Migration Tools**, and then click **Run as administrator**.

To run Windows Server Migration Tools in a new Windows PowerShell session

1. Open a Windows PowerShell session with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, click **Windows PowerShell**, right-click the Windows PowerShell shortcut, and then click **Run as administrator**.
2. Load Windows Server Migration Tools into your Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Source computer running Server Core installation option of Windows Server 2008 R2

Start Windows PowerShell and use Windows Server Migration Tools cmdlets by using any of the following procedures.

To open Windows PowerShell together with Windows Server Migration Tools

- At a command prompt on a computer that is running the Server Core installation option of Windows Server 2008 R2, type the following, and then press **Enter**.

```
powershell.exe -PSConsoleFile
%SystemRoot%\system32\ServerMigrationTools\ServerMigration.ps
c1
```

To open Windows PowerShell and load Windows Server Migration Tools separately

1. At a command prompt, type the following, and then press **Enter**.

```
powershell
```

2. Load Windows Server Migration Tools into the Windows PowerShell session. To load Windows Server Migration Tools, type the following, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

Windows Server 2003 or Windows Server 2008 source computers

If you close the Windows PowerShell session that is opened automatically when **SmigDeploy.exe** finishes registering the Windows Server Migration Tools cmdlets, you can run Windows Server Migration Tools cmdlets by using any of the following procedures.

To open Windows Server Migration Tools from the Start menu

- Do one of the following:
  - On computers that are running Windows Server 2003, click **Start**, point to **Administrative Tools**, open the **Windows Server Migration Tools** folder, and then

click **Windows Server Migration Tools**.

- On computers that are running Windows Server 2008, click **Start**, point to **Administrative Tools**, open the **Windows Server Migration Tools** folder, right-click **Windows Server Migration Tools**, and then click **Run as administrator**.

To open Windows PowerShell and load Windows Server Migration Tools separately

1. Do one of the following:
   - On computers that are running Windows Server 2003, open a Windows PowerShell session by clicking **Start**, clicking **All Programs**, opening the **Windows PowerShell** folder, and clicking the **Windows PowerShell** shortcut.
   - On computers that are running Windows Server 2008, open a Windows PowerShell session with elevated user rights. To do this, click **Start**, click **All Programs**, open the **Windows PowerShell** folder, right-click the **Windows PowerShell** shortcut, and then click **Run as administrator**.
2. In the Windows PowerShell session, type the following to load the Windows Server Migration Tools snap-in, and then press **Enter**.

```
Add-PSSnapin Microsoft.Windows.ServerManager.Migration
```

To open Windows PowerShell together with Windows Server Migration Tools from a Command Prompt window

1. Do one of the following.
   - On computers that are running Windows Server 2003, open a Command Prompt window by clicking **Start**, clicking **Run**, typing **cmd**, and then pressing **Enter**.
   - On computers that are running Windows Server 2008, open a Command Prompt window with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, right-click the **Command Prompt** shortcut, and then click **Run as administrator**.
2. At the command prompt, change directories to the location of the Windows Server Migration Tools deployment folder.
3. In the deployment directory, type the following to open a Windows PowerShell session with preloaded Windows Server Migration Tools cmdlets, and then press **Enter**.

```
PowerShell.exe -PSConsoleFile ServerMigration.psc1
```

Additional resources and next steps for using Windows Server Migration Tools

For more information about Windows Server Migration Tools and Windows PowerShell, see the following resources.

- For detailed, step-by-step information about how to migrate specific roles or data, see the Windows Server Migration Portal on the Windows Server TechCenter.
- In a Windows PowerShell session, type the following, and then press **Enter** to view detailed information about how to use a specific Windows Server Migration Tools cmdlet.

```
Get-Help <cmdlet_name> -full
```

- See the Windows PowerShell page on the Microsoft Web site.

Remove Windows Server Migration Tools

Follow the procedures in this section to remove Windows Server Migration Tools from computers.

Full installation option of Windows Server 2012 R2 or Windows Server 2012

You can use either Server Manager deployment cmdlets, or the Add Roles and Features Wizard in Server Manager to remove Windows Server Migration Tools. If Windows Server 2012 was a source computer for a migration to a server running Windows Server 2012 R2, unregister Windows Server Migration Tools on the source computer instead of uninstalling Windows Server Migration Tools. For more information, see Source computers running full and Server Core installation options of Windows Server 2012.

To uninstall Windows Server Migration Tools from the full installation option

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note
   If you are uninstalling Windows Server Migration Tools from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
2. Type the following, and then press **Enter**. If you are uninstalling the feature from the local server, omit the `ComputerName` parameter.

   ```
   Uninstall-WindowsFeature Migration –ComputerName
   <computer_name>
   ```

   Note
   You can also uninstall Windows Server Migration Tools from a full installation of Windows Server 2012 R2 or Windows Server 2012 by using the Add Roles and Features Wizard in Server Manager. For more information about how to use the Add Roles and Features Wizard, see Install or uninstall roles, role services, or features.

Server Core installation option of Windows Server 2012 R2 or Windows Server 2012

Windows PowerShell is installed by default on the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012. By default, programs on the Server Core installation option run as Administrator, so there is no need to start Windows PowerShell with elevated user rights.

To uninstall Windows Server Migration Tools from the Server Core installation option

1. Open a Windows PowerShell session by typing the following in the current command prompt session, and then press **Enter**.

   ```
   powershell.exe
   ```

2. In the Windows PowerShell session, uninstall Windows Server Migration Tools by using the Windows PowerShell **Uninstall-WindowsFeature** cmdlet for Server Manager. In the Windows PowerShell session, type the following, and then press **Enter**. Omit the `ComputerName` parameter if you are uninstalling Windows Server Migration Tools from the local server.

   ```
   Uninstall-WindowsFeature Migration –ComputerName
   <computer_name>
   ```

Source computers running full and Server Core installation options of Windows Server 2012

To remove Windows Server Migration Tools from a source computer that is running Windows Server 2012, and on which Windows Server Migration Tools was registered for migrating to a destination server running Windows Server 2012 R2, you must first reverse the registration of Windows Server Migration Tools cmdlets, and then remove the deployment folder.

To remove Windows Server Migration Tools from Windows Server 2012

1. Do one of the following.
   - On computers that are running the full installation option of Windows Server 2012, open a Command Prompt window with elevated user rights. To do this, on the **Start** screen, type **cmd**. Right-click the **Command Prompt** tile, and then click **Run as Administrator**.
   - On computers that are running the Server Core installation option of Windows Server 2012, select the Command Prompt window to bring it in focus. You do not need to open a command prompt with elevated user rights on Server Core installations.
2. Change directories to the location of the Windows Server Migration Tools deployment folder.
3. Type the following to reverse the registration of Windows Server Migration Tools cmdlets, and then press **Enter**.

   ```
   SmigDeploy.exe /unregister
   ```

4. When **SmigDeploy.exe** has finished, delete the Windows Server Migration Tools deployment folder and its contents.

Source computers running full and Server Core installation options of Windows Server 2008 R2

To remove Windows Server Migration Tools, you must first reverse the registration of Windows Server Migration Tools cmdlets, and then remove the deployment folder.

To remove Windows Server Migration Tools from Windows Server 2008 R2

1. Do one of the following.
   - On computers that are running the full installation option of Windows Server 2008 R2, open a Command Prompt window with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.
   - On computers that are running the Server Core installation option of Windows Server 2008 R2, select the Command Prompt window to bring it in focus. You do not need to open a command prompt with elevated user rights on Server Core installations.
2. Change directories to the location of the Windows Server Migration Tools deployment folder.
3. Type the following to reverse the registration of Windows Server Migration Tools cmdlets, and then press **Enter**.

   ```
   SmigDeploy.exe /unregister
   ```
4. When **SmigDeploy.exe** has finished, delete the Windows Server Migration Tools deployment folder and its contents.

Windows Server 2003 or Windows Server 2008 source computers

To remove Windows Server Migration Tools, you must first reverse the registration of Windows Server Migration Tools cmdlets, and then remove the deployment folder.

To remove Windows Server Migration Tools from Windows Server 2003 or Windows Server 2008

1. Do one of the following.
   - On computers that are running Windows Server 2003, open a Command Prompt window by clicking **Start**, clicking **Run**, typing **cmd**, and then pressing **Enter**.
   - On computers that are running Windows Server 2008, open a Command Prompt window with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.
2. At a command prompt, change directories to the location of the Windows Server

Migration Tools deployment folder.

3. Type the following to reverse the registration of Windows Server Migration Tools cmdlets, and then press **Enter**.

```
SmigDeploy.exe /unregister
```

4. When **SmigDeploy.exe** has finished, delete the Windows Server Migration Tools deployment folder and its contents.

See Also

Windows Server Migration Portal

Windows PowerShell

Install or uninstall roles, role services, or features

Adding Server Roles and Features

Migrate Active Directory Federation Services Role Services to Windows Server 2012

About this guide

This guide provides instructions to migrate the following role services to Active Directory Federation Services (AD FS) that is installed with Windows Server 2012:

- AD FS 1.1 Windows token-based agent and AD FS 1.1 claims-aware agent installed with Windows Server 2008 or Windows Server 2008 R2
- AD FS 2.0 federation server and AD FS 2.0 federation server proxy installed on Windows Server 2008 or Windows Server 2008 R2

Target audience

- IT architects who are responsible for computer management and security throughout an organization
- IT operations engineers who are responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications
- IT operations managers who are accountable for network and server management

Supported migration scenarios

The migration instructions in this guide consist of the following tasks:

- Exporting the AD FS 2.0 configuration data from your server that is running Windows Server 2008 or Windows Server 2008 R2

- Performing an in-place upgrade of the operating system of this server from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012
- Recreating the original AD FS configuration and restoring the remaining AD FS service settings on this server, which is now running the AD FS server role that is installed with Windows Server 2012.

This guide does not include instructions to migrate a server that is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration process specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Supported operating systems

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012 or Windows Server 2008 R2 (Server Core and full installation options)	x64-based
x86- or x64-based	Windows Server 2003 R2		
x86- or x64-based	Windows Server 2008, both full and Server Core installation options		
x64-based	Windows Server 2008 R2		
x64-based	Server Core installation option of Windows Server 2008 R2		
x64-based	Server Core and full installation options of Windows Server 2012		

Notes

- The versions of operating systems that are listed in the preceding table are the oldest combinations of operating systems and service packs that are supported.
- The Foundation, Standard, Enterprise, and Datacenter editions of the Windows Server operating system are supported as the source or the destination server.
- Migrations between physical operating systems and virtual operating systems are supported.

Supported AD FS role services and features

The following table describes the migration scenarios of the AD FS role services and their respective settings that are described in this guide.

From	To AD FS installed with Windows Server 2012
AD FS 1.0 federation server installed with Windows Server 2003 R2	Migration is not supported
AD FS 1.0 federation server proxy installed with Windows Server 2003 R2	Migration is not supported
AD FS 1.0 Windows token-based agent installed with Windows Server 2003 R2	Migration is not supported
AD FS 1.0 claims-aware agent installed with Windows Server 2003 R2)	Migration is not supported
AD FS 1.1 federation server installed with Windows Server 2008 or Windows Server 2008 R2	Migration is not supported
AD FS 1.1 federation server proxy installed with Windows Server 2008 or Windows Server 2008 R2	Migration is not supported
AD FS 1.1 Windows token-based agent installed with Windows Server 2008 or Windows Server 2008 R2	Migration on the same server is supported, but the migrated AD FS Windows token-based agent will function only with an AD FS 1.1 federation service installed with Windows Server 2008 or Windows Server 2008 R2. For more information, see: • Migrate the AD FS 1.1 Web Agents • Interoperating with AD FS 1.x
AD FS 1.1 claims-aware agent installed with Windows Server 2008 or Windows Server 2008 R2)	Migration on the same server is supported. The migrated AD FS 1.1 claims-aware web agent will function with the following: • AD FS 1.1 federation service installed with Windows Server 2008 or Windows Server 2008 R2 • AD FS 2.0 federation service installed on Windows Server 2008 or Windows Server 2008 R2 • AD FS federation service installed with Windows Server 2012

From	To AD FS installed with Windows Server 2012
	For more information, see: • Migrate the AD FS 1.1 Web Agents • Interoperating with AD FS 1.x
AD FS 2.0 federation server installed on Windows Server 2008 or Windows Server 2008 R2	Migration on the same server is supported. For more information, see: • Prepare to Migrate the AD FS 2.0 Federation Server • Migrate the AD FS 2.0 Federation Server
AD FS 2.0 federation server proxy installed on Windows Server 2008 or Windows Server 2008 R2	Migration on the same server is supported. For more information see: • Prepare to Migrate the AD FS 2.0 Federation Server Proxy • Migrate the AD FS 2.0 Federation Server Proxy

See Also

Prepare to Migrate the AD FS 2.0 Federation Server

Prepare to Migrate the AD FS 2.0 Federation Server Proxy

Migrate the AD FS 2.0 Federation Server

Migrate the AD FS 2.0 Federation Server Proxy

Migrate the AD FS 1.1 Web Agents

Prepare to Migrate the AD FS 2.0 Federation Server

This topic includes the following information:

- Prepare to migrate a stand-alone AD FS federation server or a single-node AD FS farm
- Prepare to migrate a WID farm
- Prepare to migrate a SQL Server farm

Prepare to migrate a stand-alone AD FS federation server or a single-node AD FS farm

To prepare to migrate (same server migration) a stand-alone AD FS 2.0 federation server or a single-node AD FS farm to Windows Server 2012, you must export and back up the AD FS configuration data from this server.

To export the AD FS configuration data, perform the following tasks:

- Step 1: Export service settings
- Step 2: - Export claims provider trusts
- Step 3: - Export relying party trusts
- Step 4: - Back up custom attribute stores
- Step 5: Back up webpage customizations

Step 1: Export service settings

To export service settings, perform the following procedure:

To export service settings

1. Record the certificate subject name and thumbprint value of the SSL certificate used by the federation service. To find the SSL certificate, open the Internet Information Services (IIS) management console, Select **Default Web Site** in the left pane, click **Bindings…** in the **Action** pane, find and select the https binding, click **Edit**, and then click **View**.

   Notes

   Optionally, you can also export the SSL certificate used by the federation service and its private key to a .pfx file. For more information, see Export the Private Key Portion of a Server Authentication Certificate.

   Exporting the SSL certificate is optional because this certificate is stored in the local computer Personal certificates store and is preserved in the operating system upgrade.

2. Record the configuration of the AD FS Service communications, token-decrypting and token-signing certificates. To view all the certificates that are used, open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to create a list of all certificates in use in a file `PSH:>Get-ADFSCertificate | Out-File ".\certificates.txt"`

   Notes

   Optionally, you can also export any token-signing, token-encryption, or service-communications certificates and keys that are not internally generated, in addition to all self-signed certificates. You can view all the certificates that are in use on your server by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows

PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell`. Then run the following command to view all certificates that are in use on your server `PSH:>Get-ADFSCertificate`. The output of this command includes StoreLocation and StoreName values that specify the store location of each certificate. You can then use the guidance in Export the Private Key Portion of a Server Authentication Certificate to export each certificate and its private key to a .pfx file.

Exporting these certificates is optional because all external certificates are preserved during the operating system upgrade.

3. Export AD FS 2.0 federation service properties, such as the federation service name, federation service display name, and federation server identifier to a file.

   To export federation service properties, open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to export federation service properties: `PSH:> Get-ADFSProperties | Out-File ".\properties.txt"`.

   The output file will contain the following important configuration values:

Federation Service Property name as reported by Get-ADFSProperties	Federation Service Property name in AD FS management console
HostName	Federation Service name
Identifier	Federation Service identifier
DisplayName	Federation Service display name

4. Back up the application configuration file. Among other settings, this file contains the policy database connection string.

   To back up the application configuration file, you must manually copy the `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config` file to a secure location on a backup server.

   Notes

   Make note of the database connection string in this file, located immediately after "policystore connectionstring="). If the connection string specifies a SQL Server database, the value is needed when restoring the original AD FS configuration on the federation server.

   The following is an example of a WID connection string: `"Data Source=\\.\pipe\mssql$microsoft##ssee\sql\query;Initial Catalog=AdfsConfiguration;Integrated Security=True"`. The following is an example of a SQL Server connection string: `"Data Source=databasehostname;Integrated Security=True"`.

5. Record the identity of the AD FS 2.0 federation service account and the password of this

account.

To find the identity value, examine the **Log On As** column of **AD FS 2.0 Windows Service** in the **Services** console and manually record this value.

Note
For a stand-alone federation service, the built-in NETWORK SERVICE account is used. In this case, you do not need to have a password.

6. Export the list of enabled AD FS endpoints to a file.

   To do this, open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to export the list of enabled AD FS endpoints to a file: `PSH:> Get-ADFSEndpoint | Out-File ".\endpoints.txt"`.

7. Export any custom claim descriptions to a file.

   To do this, open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to export any custom claim descriptions to a file: `Get-ADFSClaimDescription | Out-File ".\claimtypes.txt"`.

Step 2: - Export claims provider trusts

To export the claims provider trusts, perform the following procedure:

To export claims provider trusts

1. You can use Windows PowerShell to export all claims provider trusts. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to export all claims provider trusts: `PSH:>Get-ADFSClaimsProviderTrust | Out-File ".\cptrusts.txt"`.

Step 3: - Export relying party trusts

To export relying party trusts, perform the following procedure:

To export relying party trusts

1. To export all relying party trusts, open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to export all relying party trusts:`PSH:>Get-ADFSRelyingPartyTrust | Out-File ".\rptrusts.txt"`.

Step 4: - Back up custom attribute stores

You can find information about custom attribute stores in use by AD FS by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to find information about the custom attribute stores: `PSH:>Get-ADFSAttributeStore`. The steps to upgrade or migrate custom attribute stores vary.

Step 5: Back up webpage customizations

To back up any webpage customizations, copy the AD FS webpages and the **web.config** file from the directory that is mapped to the virtual path **"/adfs/ls"** in IIS. By default, it is in the **%systemdrive%\inetpub\adfs\ls** directory.

Prepare to migrate a WID farm

To prepare to migrate AD FS 2.0 federation servers that belong to a Windows Internal Database (WID) farm to Windows Server 2012, you must export and back up the AD FS configuration data from these servers.

To export the AD FS configuration data, perform the following tasks:

- Step 1: - Export service settings
- Step 2: Back up custom attribute stores
- Step 3: Back up webpage customizations

Step 1: - Export service settings

To export service settings, perform the following procedure:

To export service settings

1. Record the certificate subject name and thumbprint value of the SSL certificate used by the federation service. To find the SSL certificate, open the Internet Information Services (IIS) management console, select **Default Web Site** in the left pane, click **Bindings…** in the **Action** pane, find and select the https binding, click **Edit**, then click **View**.

   Notes

   Optionally, you can also export the SSL certificate and its private key to a .pfx file. For more information, see Export the Private Key Portion of a Server Authentication Certificate.

   This step is optional because this certificate is stored in the local computer Personal certificates store and will be preserved in the operating system upgrade.

2. Export any token-signing, token-encryption, or service-communications certificates and keys that are not internally generated, in addition to self-signed certificates.

   You can view all the certificates that are in use on your server by using Windows

PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to view all certificates that are in use on your server `PSH:>Get-ADFSCertificate`. The output of this command includes StoreLocation and StoreName values that specify the store location of each certificate. You can then use the guidance in Export the Private Key Portion of a Server Authentication Certificate to export each certificate and its private key to a .pfx file.

> **Note**
> This step is optional, because all external certificates are preserved during the operating system upgrade.

3. Record the identity of the AD FS 2.0 federation service account and the password of this account.

   To find the identity value, examine the **Log On As** column of **AD FS 2.0 Windows Service** in the **Services** console and manually record the value.

Step 2: Back up custom attribute stores

You can find information about custom attribute stores in use by AD FS by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to find information about the custom attribute stores: `PSH:>Get-ADFSAttributeStore`. The steps to upgrade or migrate custom attribute stores vary.

Step 3: Back up webpage customizations

To back up any webpage customizations, copy the AD FS webpages and the **web.config** file from the directory that is mapped to the virtual path **"/adfs/ls"** in IIS. By default, it is in the **%systemdrive%\inetpub\adfs\ls** directory.

Prepare to migrate a SQL Server farm

To prepare to migrate AD FS 2.0 federation servers that belong to a SQL Server farm to Windows Server 2012, you must export and back up the AD FS configuration data from these servers.

To export the AD FS configuration data, perform the following tasks:

- Step 1: Export service settings
- Step 2: Back up custom attribute stores
- Step 3: Back up webpage customizations

Step 1: Export service settings

To export service settings, perform the following procedure:

To export service settings

1. Record the certificate subject name and thumbprint value of the SSL certificate used by the federation service. To find the SSL certificate, open the Internet Information Services (IIS) management console, select **Default Web Site** in the left pane, click **Bindings…** in the **Action** pane, find and select the https binding, click **Edit**, and then click **View**.

   Notes

   Optionally, you can also export the SSL) certificate and its private key to a .pfx file. For more information, see Export the Private Key Portion of a Server Authentication Certificate.

   This step is optional because this certificate is stored in the local computer Personal certificates store and will be preserved in the operating system upgrade.

2. Export any other token-signing, token-encryption, or service-communications certificates and keys that are not internally generated by AD FS.

   You can view all certificates that are in use by AD FS on your server by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to view all certificates that are in use on your server `PSH:>Get-ADFSCertificate`. The output of this command includes StoreLocation and StoreName values that specify the store location of each certificate.

   Note

   Optionally, you can then use the guidance in Export the Private Key Portion of a Server Authentication Certificate to export each certificate and its private key to a .pfx file. This step is optional, because all external certificates are preserved during the operating system upgrade.

3. Back up the application configuration file. Among other settings, this file contains the policy database connection string.

   To back up the application configuration file, you must manually copy the `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config` file to a secure location on a backup server.

   Note

   Record the SQL Server connection string after “policystore connectionstring=” in the following file: `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config`. You need this string when you restore the original AD FS configuration on the federation server.

4. Record the identity of the AD FS 2.0 federation service account and the password of this account.

To find the identity value, examine the **Log On As** column of **AD FS 2.0 Windows Service** in the **Services** console and manually record the value.

Step 2: Back up custom attribute stores

You can find information about custom attribute stores in use by AD FS by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to find information about the custom attribute stores: `PSH:>Get-ADFSAttributeStore`. The steps to upgrade or migrate custom attribute stores vary.

Step 3: Back up webpage customizations

To back up any webpage customizations, copy the AD FS webpages and the **web.config** file from the directory that is mapped to the virtual path **"/adfs/ls"** in IIS. By default, it is in the **%systemdrive%\inetpub\adfs\ls** directory.

See Also

Migrate Active Directory Federation Services Role Services to Windows Server 2012

Prepare to Migrate the AD FS 2.0 Federation Server Proxy

To prepare to migrate an AD FS 2.0 federation server proxy to Windows Server 2012, you must export and back up the AD FS configuration data from this server proxy. The steps in this topic apply to a scenario with one proxy federation server or multiple proxy federation servers.

To export the AD FS configuration data, perform the following tasks:

- Step 1: Export proxy service settings
- Step 2: Back up webpage customizations

Step 1: Export proxy service settings

To export federation server proxy service settings, perform the following procedure:

To export proxy service settings

1. Export the Secure Sockets Layer (SSL) certificate and its private key to a .pfx file. For more information, see Export the Private Key Portion of a Server Authentication

Certificate.

Note
This step is optional because this certificate is preserved during the operating system upgrade.

2. Export AD FS 2.0 federation proxy properties to a file. You can do that by using Windows PowerShell.

   Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell".` Then run the following command to export federation proxy properties to a file: `PSH:> Get-ADFSProxyProperties | out-file ".\proxyproperties.txt".`

3. Ensure you know the credentials of an account that is either an administrator of the AD FS federation server or the service account under which the AD FS federation service runs. This information is required for the proxy trust setup.

Completing this step results in gathering the following information that is required to configure your AD FS federation server proxy:

- AD FS federation service name
- Name of the domain account that is required for the proxy trust setup
- The address and the port of the HTTP proxy (if there is an HTTP proxy between the AD FS federation server proxy and the AD FS federation servers)

Step 2: Back up webpage customizations

To back up webpage customizations, copy the AD FS proxy webpages and the **web.config** file from the directory that is mapped to the virtual path **"/adfs/ls"** in IIS. By default, it is in the **%systemdrive%\inetpub\adfs\ls** directory.

See Also

Migrate Active Directory Federation Services Role Services to Windows Server 2012

Migrate the AD FS 2.0 Federation Server

This topic provides instructions for the following migration scenarios:

- Migrate a stand-alone AD FS federation server or a single-node AD FS farm
- Migrate a WID farm
- Migrate a SQL Server farm

Migrate a stand-alone AD FS federation server or a single-node AD FS farm

To migrate a stand-alone AD FS federation server or a single-node AD FS farm to Windows Server 2012, perform the following procedure:

1. Review and perform the procedures in the “Prepare to migrate a stand-alone AD FS federation server or a single-node AD FS farm” section of Prepare to Migrate the AD FS 2.0 Federation Server.
2. Perform an in-place upgrade of the operating system on your server from Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012. For more information, see Installing Windows Server 2012.

   Important
   As the result of the operating system upgrade, the AD FS configuration on this server is lost and the AD FS 2.0 server role is removed. The Windows Server 2012 AD FS server role is installed instead, but it is not configured. You must manually create the original AD FS configuration and restore the remaining AD FS settings to complete the federation server migration.
3. Create the original AD FS configuration. You can create the original AD FS configuration by using either of the following methods:
   - Use the **AD FS Federation Server Configuration Wizard** to create a new federation server. For more information, see Create the First Federation Server in a Federation Server Farm.

     As you go through the wizard, use the information you gathered while preparing to migrate your AD FS federation server as follows:

Federation Server Configuration Wizard input option	**Use the following value**
SSL Certificate on the **Specify the Federation Service Name** page	Select the SSL certificate whose subject name and thumbprint you recorded while preparing for the AD FS federation server migration.
Service account and **Password** on the **Specify a Service Account** page	Enter the service account information that you recorded while preparing for the AD FS federation server migration. **Note** If you select stand-alone federation server on the second page of the

	wizard, NETWORK SERVICE is used automatically as the service account.

Important
You can employ this method only if you are using Windows Internal Database (WID) to store the AD FS configuration database for your stand-alone federation server or a single-node AD FS farm.

If you are using SQL Server to store the AD FS configuration database for your single-node AD FS farm, you must use Windows PowerShell to create the original AD FS configuration on your federation server.

- Use Windows PowerShell

  Important
  You must use Windows PowerShell if you are using SQL Server to store the AD FS configuration database for your stand-alone federation server or a single-node AD FS farm.

  The following is an example of how to use Windows PowerShell to create the original AD FS configuration on a federation server in a single-node SQL Server farm. Open the Windows PowerShell module and run the following command: `$fscredential = Get-Credential`. Enter the name and the password of the service account that you recorded while preparing your SQL server farm for migration. Then run the following command: `C:\PS> Add-AdfsFarmNode -ServiceAccountCredential $fscredential -SQLConnectionString "Data Source=<Data Source>;Integrated Security=True"` where `Data Source` is the data source value in the policy store connection string value in the following file: `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config`.

4. Restore the remaining AD FS service settings and trust relationships. This is a manual step during which you can use the files that you exported and the values that you collected while preparing for the AD FS migration. For detailed instructions, see Restoring the Remaining AD FS Farm Configuration.

   Note
   This step is only required if you are migrating a stand-alone federation server or a single node WID farm. If the federation server uses a SQL Server database as the configuration store, the service settings and trust relationships are preserved in the database.

5. Update your AD FS webpages. This is a manual step. If you backed up your customized AD FS webpages while preparing for the migration, use your backup data to overwrite the default AD FS webpages that were created by default in the

%systemdrive%\inetpub\adfs\ls directory as a result of the AD FS configuration on Windows Server 2012.

6. Restore any remaining AD FS customizations, such as custom attribute stores.

Migrate a WID farm

To migrate a Windows Internal Database (WID) farm to Windows Server 2012, perform the following procedure:

1. For every node (server) in the WID farm, review and perform the procedures in the "Prepare to migrate a WID farm" section of Prepare to Migrate the AD FS 2.0 Federation Server.
2. Remove any non-primary nodes from the load balancer.
3. Upgrade of the operating system on this server from Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012. For more information, see Installing Windows Server 2012.

   Important
   As the result of the operating system upgrade, the AD FS configuration on this server is lost and the AD FS 2.0 server role is removed. The Windows Server 2012 AD FS server role is installed instead, but it is not configured. You must create the original AD FS configuration and restore the remaining AD FS settings to complete the federation server migration.
4. Create the original AD FS configuration on this server.

   You can create the original AD FS configuration by using the **AD FS Federation Server Configuration Wizard** to add a federation server to a WID farm. For more information, see Add a Federation Server to a Federation Server Farm.

   Notes
   - When you reach the **Specify the Primary Federation Server and a Service Account** page in the **AD FS Federation Server Configuration Wizard**, enter the name of the primary federation server of the WID farm and be sure to enter the service account information that you recorded while preparing for the AD FS migration. For more information, see the "Prepare to migrate a WID farm" section in Prepare to Migrate the AD FS 2.0 Federation Server.
   - When you reach the **Specify the Federation Service Name** page, be sure to select the same SSL certificate you recorded in the "Prepare to migrate a WID farm" section in Prepare to Migrate the AD FS 2.0 Federation Server.

5. Update your AD FS webpages on this server. If you backed up your customized AD FS webpages while preparing for the migration, you need to use your backup data to overwrite the default AD FS webpages that were created by default in the

%systemdrive%\inetpub\adfs\ls directory as a result of the AD FS configuration on Windows Server 2012.

6. Add the server that you just upgraded to Windows Server 2012 to the load balancer.
7. Repeat steps 1 through 6 for the remaining secondary servers in your WID farm.
8. Promote one of the upgraded secondary servers to be the primary server in your WID farm. To do this, open Windows PowerShell and run the following command: `PSH:> Set-AdfsSyncProperties –Role PrimaryComputer`.
9. Remove the original primary server of your WID farm from the load balancer.
10. Demote the original primary server in your WID farm to be a secondary server by using Windows PowerShell. Open Windows PowerShell and run the following command to add the AD FS cmdlets to your Windows PowerShell session: `PSH:>add-pssnapin "Microsoft.adfs.powershell"`. Then run the following command to demote the original primary server to be a secondary server: `PSH:> Set-AdfsSyncProperties – Role SecondaryComputer –PrimaryComputerName <FQDN of the Primary Federation Server>`.
11. Upgrade of the operating system on this last node (server) in your WID farm from Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012. For more information, see Installing Windows Server 2012.

    Important
    As the result of upgrading the operating system, the AD FS configuration on this server is lost and the AD FS 2.0 server role is removed. The Windows Server 2012 AD FS server role is installed instead, but it is not configured. You must manually create the original AD FS configuration and restore the remaining AD FS settings to complete the federation server migration.

12. Create the original AD FS configuration on this last node (server) in your WID farm.

    You can create the original AD FS configuration by using the **AD FS Federation Server Configuration Wizard** to add a federation server to a WID farm. For more information, see Add a Federation Server to a Federation Server Farm.

    Notes
    - When you reach the **Specify the Primary Federation server and a Service Account** page in the **AD FS Federation Server Configuration Wizard**, enter the service account information that you recorded while preparing for the AD FS migration. For more information, see the “Prepare to migrate a WID farm” section in Prepare to Migrate the AD FS 2.0 Federation Server.
    - When you reach the **Specify the Federation Service Name** page, be sure to select the same SSL certificate you recorded in the “Prepare to migrate a WID farm” section in Prepare to Migrate the AD FS 2.0 Federation Server.

13. Update your AD FS webpages on this last server in your WID farm. If you backed up your customized AD FS webpages while preparing for the migration, use your backup data to overwrite the default AD FS webpages that were created by default in the **%systemdrive%\inetpub\adfs\ls** directory as a result of the AD FS configuration on Windows Server 2012.
14. Add this last server of your WID farm that you just upgraded to Windows Server 2012 to

the load balancer.

15. Restore any remaining AD FS customizations, such as custom attribute stores.

Migrate a SQL Server farm

To migrate a SQL Server farm to Windows Server 2012, perform the following procedure:

1. For each server in your SQL Server farm, review and perform the procedures in the “Prepare to migrate a SQL Server farm” section of Prepare to Migrate the AD FS 2.0 Federation Server.
2. Remove any server in your SQL Server farm from the load balancer.
3. Upgrade the operating system on this server in your SQL Server farm from Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012. For more information, see Installing Windows Server 2012.

   Important
   As the result of the operating system upgrade, the AD FS configuration on this server is lost and the AD FS 2.0 server role is removed. The Windows Server 2012 AD FS server role is installed instead, but it is not configured. You must manually create the original AD FS configuration and restore the remaining AD FS settings to complete the federation server migration.
4. Create the original AD FS configuration on this server in your SQL Server farm by using AD FS Windows PowerShell cmdlets to add a server to an existing farm.

   Important
   You must use Windows PowerShell to create the original AD FS configuration if you are using SQL Server to store your AD FS configuration database.

   a. Open Windows PowerShell and run the following command: `$fscredential = Get-Credential`.
   b. Enter the name and the password of the service account that you recorded while preparing your SQL Server farm for migration.
   c. Run the following command: `Add-AdfsFarmNode -ServiceAccountCredential $fscredential -SQLConnectionString "Data Source=<Data Source>;Integrated Security=True"`, where `Data Source` is the data source value in the policy store connection string value in the following file: `%programfiles%\Active Directory Federation Services 2.0\Microsoft.IdentityServer.Servicehost.exe.config`.
5. Add the server that you just upgraded to Windows Server 2012 to the load balancer.
6. Repeat steps 2 through 6 for the remaining nodes in your SQL Server farm.
7. When all of the servers in your SQL Server farm are upgraded to Windows Server 2012, restore any remaining AD FS customizations, such as custom attribute stores.

Restoring the Remaining AD FS Farm Configuration

- Restore the following AD FS service settings to a single node WID farm or stand-alone federation service as follows:
  - In the AD FS management console, select **Service** and click **Edit Federation Service…**. Verify the federation service settings by checking each of the values against the values you exported into the properties.txt file while preparing for the migration:

Federation Service Property name as reported by Get-ADFSProperties	Federation Service Property name in AD FS Management console
DisplayName	Federation Service display name
HostName	Federation Service name
Identifier	Federation Service identifier

  - In the AD FS management console, select **Certificates**. Verify the service communications, token-decrypting, and token-signing certificates by checking each against the values you exported into the certificates.txt file while preparing for the migration.

    To change the token-decrypting or token-signing certificates from the default self-signed certificates to external certificates, you must first disable the automatic certificate rollover feature that is enabled by default. To do this, you can use the following Windows PowerShell command: `PSH: Set-ADFSProperties –AutoCertificateRollover $false`.
  - In the AD FS Management console, select **Endpoints**. Check the enabled AD FS endpoints against the list of enabled AD FS endpoints that you exported to a file while preparing for the AD FS migration.
  - In the AD FS Management console, select **Claim Descriptions**. Check the list of AD FS claim descriptions against the list of claim descriptions that you exported to a file while preparing for the AD FS migration. Add any custom claim descriptions included in your file but not included in the default list in AD FS. Note that Claim identifier in the management console maps to the ClaimType in the file. For more information on adding claim descriptions, see Add a Claim Description. For more information, see the “Step 1 - Export Service Settings” section in Prepare to Migrate the AD FS 2.0 Federation Server.
- In the AD FS Management console, select **Claims Provider Trusts**. You must recreate each Claims Provider trust manually using the **Add Claims Provider Trust Wizard**. Use the list of claims provider trusts that you exported and recorded while preparing for the AD FS migration. You can disregard the claims provider trust with Identifier “AD AUTHORITY” in the file because this is the “Active Directory” claims provider trust that is part of the default AD FS configuration. However, check for any custom claim rules you may have added to the Active Directory trust prior to the migration. For more information on creating claims provider trusts, see Create a Claims Provider Trust Using Federation Metadata or Create a Claims Provider Trust Manually.

- In the AD FS Management console, **select Relying Party Trusts**. You must recreate each Relying Party trust manually using the **Add Relying Party Trust Wizard**. Use the list of relying party trusts that you exported and recorded while preparing for the AD FS migration. For more information on creating relying party trusts, see Create a Relying Party Trust Using Federation Metadata or Create a Relying Party Trust Manually.

See Also

Migrate Active Directory Federation Services Role Services to Windows Server 2012

Migrate the AD FS 2.0 Federation Server Proxy

To migrate an AD FS 2.0 federation server proxy to Windows Server 2012, perform the following procedure:

1. For every federation server proxy that you plan to migrate to Windows Server 2012, review and perform the procedures in Prepare to Migrate the AD FS 2.0 Federation Server Proxy.
2. Remove a federation server proxy from the load balancer.
3. Perform an in-place upgrade of the operating system on this server from Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012. For more information, see Installing Windows Server 2012.

   Important
   As the result of the operating system upgrade, the AD FS proxy configuration on this server is lost and the AD FS 2.0 server role is removed. The Windows Server 2012 AD FS server role is installed instead, but it is not configured. You must manually create the original AD FS proxy configuration and restore the remaining AD FS proxy settings to complete the federation server proxy migration.
4. Create the original AD FS proxy configuration by using the **AD FS Federation Server Proxy Configuration Wizard**. For more information, see Configure a Computer for the Federation Server Proxy Role. As you execute the wizard, use the information you gathered in Prepare to Migrate the AD FS 2.0 Federation Server Proxy as follows:

Federation Server Proxy Wizard input option	Use the following value

Federation Service Name	Enter the BaseHostName value from proxyproperties.txt file
Use an HTTP proxy server when sending requests to this Federation Service check box	Check this box if your proxyproperties.txt file contains a value for the ForwardProxyUrl property
HTTP proxy server address	Enter the ForwardProxyUrl value from proxyproperties.txt file
Credential prompt	Enter the credentials of an account that is either an administrator of the AD FS federation server or the service account under which the AD FS federation service runs.

5. Update your AD FS webpages on this server. If you backed up your customized AD FS proxy webpages while preparing your federation server proxy for the migration, use your backup data to overwrite the default AD FS webpages that were created by default in the **%systemdrive%\inetpub\adfs\ls** directory as a result of the AD FS proxy configuration in Windows Server 2012.
6. Add this server back to the load balancer.
7. If you have other AD FS 2.0 federation server proxies to migrate, repeat steps 2 through 6 for the remaining federation server proxy computers.

See Also

Migrate Active Directory Federation Services Role Services to Windows Server 2012

Migrate the AD FS 1.1 Web Agents

To migrate the AD FS 1.1 Windows token-based agent or the AD FS 1.1 claims-aware agent that is installed with Windows Server 2008 R2 or Windows Server 2008 to Windows Server 2012, perform an in-place upgrade of the operating system of the computer that hosts either agent to Windows Server 2012. For more information, see Installing Windows Server 2012. No further configuration is required.

Important

The migrated AD FS 1.1 Windows token-based agent functions only with an AD FS 1.1 federation service that is installed with Windows Server 2008 R2 or Windows Server 2008. For more information, see Interoperating with AD FS 1.x.

The migrated AD FS 1.1 claims-aware web agent functions with the following:

- AD FS 1.1 federation service installed with Windows Server 2008 R2 or Windows Server 2008
- AD FS 2.0 federation service installed on Windows Server 2008 R2 or Windows Server 2008
- AD FS federation service installed with Windows Server 2012

  For more information, see Interoperating with AD FS 1.x.

See Also

Migrate Active Directory Federation Services Role Services to Windows Server 2012

Migrate File and Storage Services to Windows Server 2012

The File and Storage Services Migration Guide provides step-by-step instructions for how to migrate the File and Storage Services role, including data, shared folders, and operating system settings from a source server to a destination server that is running Windows Server 2012.

About this guide

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic, and then add comments that support your rating. Click **Rate this topic** at the top of the page, and describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process. For more information about installing and using the migration tools on both source and destination servers, see the Windows Server Migration Tools Installation, Access, and Removal Guide.

Specifically, this guide includes information about migrating the following:

- Information about the server's identity
- Local users and groups
- Data and shared folders
- Shadow Copies of Shared Folders
- Data Deduplication

- DFS Namespaces
- DFS Replication
- File Server Resource Manager (FSRM)
- Group Policy settings that are specific to server message block (SMB)
- Group Policy settings for Offline Files (also known as client-side caching or CSC)
- ISCSI Software Target

> **Note**
> ISCSI Software Target was previously an optional Windows Server and Windows Storage Server component download. Due to the amount of content, all iSCSI-specific migration information is located in File and Storage Services: Appendix C: Migrate iSCSI Software Target.

Target audience

This document is intended for information technology (IT) professionals and knowledge workers who are responsible for operating and deploying file servers in a managed environment.

What this guide does not provide

This guide does not provide information or support for the following migration scenarios:

- Clustering migration for clustered server configurations
- Migrating Roaming User Profiles (for additional information see **Upgrading or Migrating a Roaming User Profiles Environment to Windows 8.1 or Windows Server 2012 R2**).
- Upgrading roles on the same computer
- Migrating more than one server role
- Migrating data across subnets
- Migrating Network File System (NFS) shared folders
- Migrating file servers by using File Server Resource Manager
- Migrating encrypted files from Encrypting File System (EFS)
- Migrating file allocation tables (FAT) and FAT32 file systems
- Migrating hardware and software installation for storage resources

In addition to these unsupported scenarios, you should understand the following migration limitations:

- If the hard disk drive that contains your data is physically moved from the source server to the destination server, file and folder permissions for local users are not preserved.
- Reparse points, hard links, and mounted volumes are not migrated when data is copied, and they need to be migrated manually.
- To facilitate migrating file and shared folder permissions, you must migrate local users and groups as part of the migration procedure. However, not all user and group attributes are migrated.

For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?LinkId=258341) on the Microsoft Web site.

Supported migration scenarios

This guide provides instructions for migrating an existing server that is running File and Storage Services to a server that is running Windows Server 2012 R2 or Windows Server 2012. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure for your server environment, based on the information that is provided in other server role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other server roles that are running on the source server to fail.

Supported migration scenarios include the following configurations or features:

- File server is joined to a domain
- File server is in a workgroup
- File server data and file shares are located in a storage area network (SAN) or other external storage location that preserves data and file share permissions (except data for local users and groups).
- File server data and file shares are located on the server disk (direct-attached storage) that is preserving data and files shares permissions.
- DFS Namespaces
- File Server Resource Manager
- Shadow Copies of Shared Folders

Important

The file migration portion of the Windows Server Migration Tools is designed for smaller data sets (under 100 GB of data). It copies files one at a time over HTTPS. For larger datasets, we recommend using the version of Robocopy.exe included with Windows Server 2012 R2 or Windows Server 2012.

Supported operating systems

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with	Windows Server 2008 R2 or	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
	Service Pack 2	Windows Server 2012, both full and Server Core installation options	
x86- or x64-based	Windows Server 2003 R2	Windows Server 2008 R2or Windows Server 2012, both full and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2008, full installation option	Windows Server 2008 R2or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2008 R2or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2008 R2or Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Server Core and full installation options of Windows Server 2012	Windows Server 2008 R2or Windows Server 2012, both full and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers.

Migrations between physical operating systems and virtual operating systems are supported.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Both x86- and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008 R2. All editions of Windows Server 2008 R2 are x64-based.

File services migration overview

The following topics contain step-by-step information about how to migrate File and Storage Services from a computer that is running Windows Server 2003 or later to a computer that is running Windows Server 2012:

- File and Storage Services: Prepare to Migrate
- File and Storage Services: Migrate the File and Storage Services Role
- File and Storage Services: Verify the Migration
- File and Storage Services: Post-Migration Tasks

Impact of migration on other computers in the enterprise

The content in this section describes the impact to the computers in your enterprise during migration.

Impact of data migration by copying data and shared folders

- The performance of your source server can be affected during the data migration. This can result in slower access to files that are stored on the server.
- At the beginning of the second phase of the data migration, all open files are closed, which can lead to data loss.
- During the second phase of data migration, clients are unable to access the file server.

Impact of data migration by physically moving data drives

Clients cannot access the file server from the moment the storage device is disconnected from the source server until it is attached to the destination server.

Impact on DFS Namespaces

The DFS Namespaces are unavailable at several times during the migration process. You should plan your migration when you can take the namespace that is hosted on the source server offline.

Impact on DFS Replication

The impact of migration activity on other servers in the enterprise depends largely on the configuration of the replication topology. Typically, DFS Replication is configured in a hub and spoke replication topology with multiple branch office servers (spokes) replicating with a single hub server. Depending on which server in the replication topology is migrated, and how the data is migrated, the remaining servers in the enterprise can be affected. Client workstations that are accessing data that is contained in the replicated folder on the server can be affected during the migration process.

Client computers may be accessing data in the folder that is being replicated by using DFS Replication. The replicated folder is often exposed to client computers as an SMB shared folder.

For more information about the impact of the migration process on client computers, see Impact of data migration by copying data and shared folders earlier in this document.

Permissions required to complete migration

This section describes permissions that are required to perform the migration.

Permissions required for data and shared folder migration

For data and shared folder migration, local Administrator permissions are required on the source server and destination server.

Permissions required to complete migration on the destination server

This section describes permissions that are required to perform the migration on the destination server.

Permissions required to migrate DFS Namespaces

For a stand-alone namespace, the user must be a member of the local Administrators group on the destination server.

There are three permissions options for a domain-based namespace:

- Option 1: Membership in the Domain Admins group
- Option 2 (if there are more than one namespace server):
  - Permission to administer all namespaces that are hosted on the source server
  - Member of the local Administrators group on the destination server
- Option 3 (if there is a single namespace server):

- Permission to delete and create domain-based namespaces in the domain
- Member of the local Administrators group on the destination server

Permissions required to complete migration on the source server

This section describes permissions that are required to perform the migration on the source server.

Permissions required to migrate DFS Namespaces

For a stand-alone namespace, the user must have membership in the local Administrators group on the source server.

There are three permissions options for a domain-based namespace:

- Option 1: Membership in the Domain Admins group
- Option 2 (if there are more than one namespace servers):
  - Permission to administer the all namespaces that are hosted on the source server
  - Member of the local Administrators group on the source server
- Option 3 (if there is a single namespace server):
  - Permission to delete and create domain-based namespaces in the domain
  - Member of the local Administrators group on the destination server

Permissions required for DFS Replication

For DFS Replication, the user who starts the migration must be a member of the Domain Admins group or delegated permissions to the replication groups and replication members. This permission is required to remove the source server from replication groups to which it belongs. If the permissions to administer a replication group have been delegated to a particular user through the DFS Management snap-in, that user can use the DFS Management snap-in to perform tasks such as removing the source server from a replication group. The user must also be a member of the local Administrators group on the source server and the destination server.

See Also

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Prepare to Migrate

This guide provides you with instructions for migrating the File and Storage Services role to a server that is running Windows Server® 2008 R2.

Install migration tools

Windows Server Migration Tools in Windows Server® 2012 allows an administrator to migrate some server roles, features, operating system settings, shared folders, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to computers that are running Windows Server 2012.

For complete installation, configuration, and removal instructions for Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.

Migration documentation and tools ease the process of migrating server role settings and data from an existing server that is running a Windows server operating system to another computer. For a complete list of supported operating systems, see Migrate File and Storage Services to Windows Server 2012.

By using these tools to migrate roles, you can simplify migration, reduce migration time, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

Prepare for migration

The following list outlines the major steps for preparing to migrate the File and Storage Services role.

- Prepare the destination server
- Back up the source server
- Prepare the source server
- Prepare other computers in the enterprise

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

Prepare the destination server

The following steps are necessary to prepare the destination server for migration.

Hardware requirements for the destination server

Verify that the data locations for the destination server have sufficient free space to migrate the data. Ensure that the destination server hard disk drives are the same size or larger than the source server hard disk drives.

Software requirements for the destination server

There are several software requirements that must be met to ensure a successful migration.

- Consult the migration matrix to determine if you can migrate the version of Windows Server that you are running on the source server to Windows Server 2012. For a complete list of supported operating systems, see Migrate File and Storage Services to Windows Server 2012.
- Before migration, install all critical updates and service packs on the source server that were released before Windows Server 2012. It is a recommended best practice that you install all current critical updates and service packs on the source server and the destination server.

Prepare for local user and group migration on the destination server

Verify that the destination server can resolve the names of domain users who are members of the local group during the import operation. If source server and destination server are in different domains, the destination server must be able to contact a global catalog server for the forest in which the source domain user accounts are located.

Prepare for File and Storage Services on destination server

1. Install Windows Server 2012 on the destination server.
2. Ensure that the time and date are set correctly on the destination server, and that they are in sync with the source server.
3. Determine the File Services role services that have been installed on the source server and then install the same File and Storage Services role services on the destination server.
4. Install Windows Server Migration Tools on the destination server.

   For more information about how to install Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.
5. Open UDP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to establish a data transfer connection.

   Note
   If you have changed the default behavior of Windows Firewall to block outbound traffic on computers that are running Windows Server 2012, you must explicitly allow outbound traffic on UDP port 7000.

6. Open TCP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to perform the data transfer.

   For more information about how to open UDP port 7000 and TCP port 7000, see File and Storage Services: Appendix A: Optional Procedures.

   For more information about how to determine if a port is in use, see the following article on the Microsoft Web site: How To Determine Which Program Uses or Blocks Specific Transmission Control Protocol Ports in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=149887).

7. Verify that the destination path has sufficient disk space to migrate the data. If NTFS or folder quota management (in File Server Resource Manager) is enabled on the destination server disk drive, verify that the quota limit allows for sufficient free disk space to migrate data. For more information about quota management in File Server Resource Manager, see one of the following.

   - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154277) for Windows Server 2008 and Windows Server 2008 R2
   - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154241) for Windows Server 2003 R2

   For more information about NTFS quota management, see one of the following.

   - Setting Disk Quotas (http://go.microsoft.com/fwlink/?LinkId=154243) for Windows Server 2008 and Windows Server 2008 R2
   - Enable disk quotas (http://go.microsoft.com/fwlink/?LinkId=154245) for Windows Server 2003 and Windows Server 2003 R2

Prepare File Server Resource Manager on destination server

If you are using File Classification Infrastructure plug-ins from a non-Microsoft vendor, you should register the non-Microsoft plug-ins on the destination server and refer to additional instructions for migration from the non-Microsoft plug-in vendor. You should register the plug-in after File Server Resource Manager (FSRM) has been installed and started on the destination server.

Use the same drive letters for the destination server volumes as for the source server. This is required, because FSRM migration requires the drive letter to remain the same.

Data and shared folder preparation on destination server

Do not allow users to access the destination server until migration is fully completed. This ensures data integrity and prevents failure when an open file on the destination server cannot be overwritten during migration.

Data integrity and security considerations on destination server

Server migration tools preserve file and folder permissions during data migration. When you are planning the migration, keep in mind that if the migrated files and folder inherit permissions from their parents, during migration it is the inheritance setting that is migrated, not the inherited permissions. Therefore it is important to make sure that the parent folders on the source server

and the destination server have the same permissions to maintain the permissions on migrated data that has inherited permissions.

For example:

1. Migrate folder c:\A\C from the source server to folder c:\B\D on the destination server.
2. Verify that on the source server, only Mary has access to folder c:\A and folder c:\A\C is specified to inherit permission from its parent.
3. Verify that on the destination server, only John has access to folder c:\B. After c:\A\C is migrated to c:\B\D, John will have access to folder D, but Mary will not.

If you use permissions inheritance for the migrated data, ensure that the parent folder for the migrated data on the destination server has the required permission set.

Prepare DFS Namespaces on destination server

The DFS Namespaces role service must be installed, and the DFS Namespace service must be running before migration. If the namespaces that you are migrating are domain-based, both source and destination servers must be in the same Active Directory® Domain Services (AD DS) domain. If the namespaces are stand-alone namespaces, AD DS membership does not matter.

Back up the source server

If DFS Namespaces are being migrated, back up the source server by using a full server backup or system state backup. If the DFS Namespaces are part of an AD DS domain, you need to back up the AD DS domain to save the Active Directory configuration information for DFS Namespaces.

For each domain-based DFS namespace, you should also back up the configuration information for the namespace. Repeat the following command for each namespace and save the output filename to a safe location:

```
DFSUtil.exe root export <\\<DomainName>\Namespace> <Filename>
```

Note

DFSUtil.exe is available on computers that are running Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012. It is available to download for use on Windows Server 2003 and Windows Server 2003 R2 as part of the Windows Server 2003 Service Pack 1 32-bit Support Tools (http://go.microsoft.com/fwlink/?LinkId=147453).

Prepare the source server

The following sections describe how to prepare the source server for the migration.

Prepare all file services on source server

- Install Windows Server Migration Tools on the source server.

  For more information about how to install Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools.

- Verify that the time and date are set correctly on the destination server and that they are synchronized with the source server.
- Open UDP port 7000 and make sure that is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to establish a data transfer connection.

  Note
  If you have changed the default behavior of Windows Firewall to block outbound traffic on computers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, you must explicitly allow outbound traffic on UDP port 7000.
- Open TCP port 7000 and make sure that it is not in use by other applications. This port is used by **Send-SmigServerData** and **Receive-SmigServerData** to perform the data transfer.

For more information about how to open UDP port 7000 and TCP port 7000, see **File Services Migration: Appendix A: Optional Procedures**.

For more information about how to determine if a port is in use, see the following article on the Microsoft Web site: How To Determine Which Program Uses or Blocks Specific Transmission Control Protocol Ports in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=149887).

Data and shared folder preparation on the source server

To minimize downtime and reduce impact to users, plan your data migration to occur during off-peak hours. Use the `net share` command to list all shared folders on the source server.

You can use this list during the verification step to verify that all the required shared folders have migrated. Reparse points and hard links will not migrate when data is copied (versus a physical migration), and the reparse points need to be migrated manually. When you migrate hard links, a separate file is created on the destination server for each link. If your migration involves copying the data to the destination server, follow the instructions for how to detect the reparse points and hard links in File and Storage Services: Appendix A: Optional Procedures. Then you can remap and recreate them during migration, as instructed in the For copy data migration scenarios section.

Prepare DFS on the source server

DFS Namespaces role services must be installed, and DFS Namespace service must be running before migration.

For information about DFS Namespaces preparation, see Prepare DFS Namespaces on source server.

Prepare DFS Namespaces on source server

For domain-based namespaces with one namespace server, determine if you will add a temporary server to the namespace or if you will perform a manual inventory of the namespace permissions.

- Option 1 (recommended):

  Add a temporary server as a namespace server to each domain-based namespace on the source server when the source server is the only namespace server.
- Option 2:

  Inventory the permissions for managing each of the namespaces that are hosted on the source server when the source server is the only namespace server. This process can be completed by using the DFS Management MMC Snap-in.

Prepare other computers in the enterprise

Data and shared folder migration requires preparing other computers in the enterprise. Following are the steps that you should perform for copy data migration scenarios, and for physical data scenarios.

For copy data migration scenarios

- Notify the users that the server performance may be reduced during the first phase of data migration.
- Ask users to stop writing to the server before the second phase of data migration begins (to prevent possible data loss). We recommend that you prevent access to the file shares so that users don't ignore this advice. For example, you could temporarily set all file shares to be read-only by setting the share permissions to Everyone = Read Only.
- Notify users that they cannot access their files on the server when the second phase of the migration begins until the file server migration is fully completed.

For physical data migration scenarios

Notify the users that they cannot access the file server from the moment the storage is disconnected from the source server until the server migration is fully completed.

See Also

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Migrate the File and Storage Services Role

Migrate File Services

Perform the following tasks to migrate the File and Storage Services server role.

- Freeze administration configuration
- Export settings
- Migrate local users and groups to the destination server
- Migrate data
- Migrate the source server identity
- Configure DFS Replication on the destination server
- Import settings to the destination server

Freeze administration configuration

Administrators must stop all configuration changes to the File and Storage Services role services on the source server before starting migration. When the migration begins, you must not make any configuration changes to the source server other than those that are required for the migration (for example, no links can be added to a (DFS namespace after migration starts until the migration is verified successfully).

Install the Windows Server Migration Tools

Before you can use any of the following Windows PowerShell cmdlets for migration on the source server or destination server, ensure that the Windows Server Migration Tools is added. You can do this by using Server Manager or by using Windows PowerShell.

Do this step using Windows PowerShell

To install the Windows Server Migration Tools

1. Log on to the computer as a member of the local Administrators security group.
2. In Server Manager, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, select **the Role-base or feature-based installation** option, and then click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, accept the default selections, and then click **Next**.
7. On the **Select features** page, click **Windows Server Migration Tools**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.

9. After the installation is complete, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature Migration
```

The following is a list of Windows Server Migration Tools cmdlets:

- **Export-SmigServerSetting**
- **Import-SmigServerSetting**
- **Get-SmigServerFeature**
- **Send-SmigServerData**
- **Receive-SmigServerData**

For more information on how to work with the Windows Server Migration Tools see Install, Use, and Remove Windows Server Migration Tools.

Export settings

Export the following settings from the source server to the destination server. Settings include Server Message Block (SMB), Offline Files (also known as called client-side caching or CSC), DFS Namespaces, File Server Resource Manager (FSRM), and Shadow Copies of Shared Folders.

BranchCache for Network Files server key

The following procedure applies only if the source server is running Windows Server 2008 R2 or Windows Server 2012.

Notes

This procedure, which is used to migrate the seed value that is used by the BranchCache™ for Network Files component, enables data that was stored in branch office locations by using BranchCache to be used after the file server is migrated from the source server to the destination server.

For information about how to migrate a BranchCache host server, see the BranchCache Migration Guide (http://go.microsoft.com/fwlink/?LinkID=139091).

To migrate BranchCache for network files settings from the source server

1. In your Windows PowerShell session, collect data from the source server by running the **Export-SmigServerSetting** cmdlet as an member of the Administrators security group. This step runs the **Export-SmigServerSetting** cmdlet with all parameters from a single command line. The **Export-SmigServerSetting** cmdlet parameters can collect all source BranchCache feature data in a single file (Svrmig.mig), or you can run the **Export-SmigServerSetting** cmdlet multiple times by using one or more parameters to collect

and store data in multiple Svrmig.mig files.

For more information, see the section "Prepare for Migration" in File and Storage Services: Prepare to Migrate.

Review the following dependencies before you run the command.

- When you run the **Export-SmigServerSetting** cmdlet, you are prompted to provide a password to encrypt the migration store data. You must provide this same password to import data from the migration store.
- The *path* parameter can be to a folder that is empty or one that contains data. The actual data file in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Therefore, the user does not have to specify a file name.
- If the path is not a shared location that the destination server can read, you must manually copy the migration store to the destination server or a location that the destination server can access.
- If a migration store location already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must move the Svrmig.mig file from the migration store location and then store it elsewhere, or rename or delete the Svrmig.mig file first.

2. On the source server, type the following, and then press ENTER, where <storepath> is the path that will contain the Svrmig.mig file after this step is completed. An example of the path is \\fileserver\users\username\branchcachestore.

```
Export-SmigServerSetting -featureID BranchCache -Path
<storepath\BranchCache> -Verbose
```

Group or local policy specific to SMB and Offline Files

Most SMB and Offline Files settings are migrated as part of shared folder migration. The remaining settings that affect the server are set through group or local policies. This section describes how to inventory SMB and Offline Files settings that are controlled by Group Policy.

Server message block

Determine the policy settings that affect the SMB server. The SMB settings are controlled by Group Policy settings or local policy settings. If a Group Policy object (GPO) is applied, these policies override the local settings. First, determine if the settings are controlled by a GPO, and then determine local settings for anything that is not controlled by the GPO.

To determine if a GPO is applied to the source server

1. Open the Resultant Set of Policy snap-in. To open the Resultant Set of Policy snap-in, open a command prompt, type **rsop.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
3. Note in the SMB data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets any Group Policy setting that affects the following

Microsoft® network server settings:

- Microsoft network server: Amount of idle time required before suspending session
- Microsoft network server: Attempt S4USelf to obtain claim information
- Microsoft network server: Digitally sign communications (always)
- Microsoft network server: Digitally sign communications (if client agrees)
- Microsoft network server: Disconnect clients when logon hours expire

On source servers that are running the Server Core installation option of the Windows Server 2008 R2 or Windows Server 2012 operating system, run the **gpresult** command to review Group Policy settings (for more information about **gpresult**, type **gpresult /?** at a command prompt.)

Notes

For any setting that is controlled by Group Policy, you must apply the same GPO to the destination server, or you can set the local policy of the destination server for the same behavior.

For any setting that is not controlled by Group Policy, use the following procedure to determine the local policy setting. Note the local policy setting in the SMB data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets.

To determine local policy settings

1. Open the Local Group Policy Editor. To open the Local Group Policy Editor, open a command prompt, type **gpedit.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
3. Note the following settings for **Microsoft network server**:
   - Microsoft network server: Amount of idle time required before suspending a session
   - Microsoft network server: Attempt S4USelf to obtain claim information
   - Microsoft network server: Digitally sign communications (always)
   - Microsoft network server: Digitally sign communications (if client agrees)
   - Microsoft network server: Disconnect clients when logon hours expire

   On source servers that are running the Server Core installation, run the **secedit** command to export and review local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt.)

Offline Files

Note

This section only applies to source servers that are running Windows Server 2012 R2 Windows Server 2012, Windows Server 2008 R2, or. Previous operating system releases do not have Offline Files settings that affect the server.

Determine the policy settings that affect file shares on the server for which client computers use Offline Files. The Offline Files settings are controlled through Group Policy or local policy. If Group Policy is applied, then these policies override local settings. First, determine if the settings are controlled through Group Policy, then determine the local settings for anything that is not controlled by using Group Policy.

To determine if Group Policy is applied to the source server

1. Open the Resultant Set of Policy snap-in. To open the Resultant Set of Policy snap-in, open a command prompt, type **rsop.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Administrative Templates**, click **Network**, and then click **Lanman Server**.

   Note
   If no policies are set, the preceding path won't exist. If the path does not exist, skip to the procedure To determine local policy settings. If the path exists and policies are found, go on to the next step.
3. Note in the BranchCache data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets any Group Policy settings that control the **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

   On source servers that are running the Server Core installation option, run the **gpresult** command to review Group Policy settings (for more information about **gpresult**, type **gpresult /?** at a command prompt).

For any setting controlled by Group Policy, have the same Group Policy apply to the destination server, or you can choose to set the local policy of the destination server to get the same behavior.

For any setting not controlled by Group Policy, use the following instructions to determine the local policy setting.

To determine local policy settings

1. Open the Local Group Policy Editor. To open the Local Group Policy Editor, open a command prompt, type **gpedit.msc**, and then press Enter.
2. In the snap-in tree pane, click **Computer Configuration**, click **Administrative Templates**, click **Network**, and then click **Lanman Server**.
3. Note in the BranchCache data collection worksheet in File and Storage Services: Appendix B: Migration Data Collection Worksheets the value of the **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

   On source servers that are running the Server Core installation option, run the **secedit** command to export and review local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

DFS Namespace configuration

Procedures in this section describe how to migrate DFS Namespaces from the source server to the destination server.

Before the migration of the namespace begins, you can inventory the namespaces that are hosted on the source server for tracking purposes. You can do this by using DFS Management or DFSUtil.exe.

The following procedure (To inventory DFS Namespaces by usingDFS Management) applies only to computers that are running at least the Windows Server 2003 R2 version of the Windows Server operating system. For computers that are running Windows Server 2003, you can perform a DFS Namespace inventory by using **DFSUtil.exe** as described in To inventory DFS Namespaces by using DFSutil.exe.

You can also perform a DFS Namespace inventory from a client computer that is running Windows Vista®, Windows® 7, or Windows® 8 by using DFS Management that is part of Remote Server Administration Tools.

- To download Remote Server Administration Tools for Windows Vista, see Microsoft Remote Server Administration Tools for Windows Vista (http://go.microsoft.com/fwlink/?LinkID=113192).
- To download Remote Server Administration Tools for Windows 7, see Remote Server Administration Tools for Windows 7 (http://go.microsoft.com/fwlink/?LinkID=131280).
- To download Remote Server Administration Tools for Windows® 8, see Remote Server Administration Tools for Windows 7 (http://go.microsoft.com/fwlink/?LinkID=131280).

To inventory DFS Namespaces by using DFS Management

1. Under DFS Management in the left pane, right-click **Namespaces**.
2. Select **Add Namespaces to Display**.
3. In the dialog box that is displayed, select **Server** from the Scope options.
4. Type the name of source server and click **Show Namespaces**.
5. Select all namespaces listed in the list box and click **OK**.
6. Right-click the first namespace listed in the left pane.
7. Select **Properties**.
8. On the **General** tab, check the **Type** field. The type of namespace that is hosted on the server is described here. Possible values are stand-alone, domain-based (Windows Server 2000 mode), and domain-based (Windows Server 2008 mode).
9. In the case of a domain-based namespace, click the **Namespace Servers** tab to identify the number of servers that host the namespace.
10. Repeat steps 7 through 10 for the remaining namespaces listed in the left pane.

To inventory DFS Namespaces by using DFSutil.exe

1. You can inventory your DFS Namespaces using DFSUtil.exe by using the command prompt. From a command prompt, type **DFSUtil.exe server SourceServer** where

SourceServer represents the name of the source server.

2. Identify the namespaces (DFS roots) listed for the source server.
3. Type the following command, list the namespace properties for the first namespace identified in step 2:

```
DFSUtil.exe root <\\SourceServer\Namespace>
```

4. Identify the namespace type; possible values are stand-alone root, domain root (domain-based namespace in Windows 2000 Server mode), domainV2 root (domain-based namespace in Windows 2008 mode).
5. Identify the DFS folders present in the namespace in each of the **Link Name** items displayed.
6. In the case of domain-based namespaces, identify all the namespace servers by typing the following command:

```
DFSUtil.exe root <\\Domain\Namespace>
```

7. Identify the namespace servers that host the namespace in each of the **Target** items displayed under **Root Name**.
8. Repeat steps 3 through 7 for the remaining namespaces on the source server.

Considerations for namespaces

Is the namespace stand-alone or domain-based? If the namespace is stand-alone, see the following section in this document:

Stand-alone namespaces.

If the namespace is domain-based, consider the number of namespace servers for each namespace. For more information, see the following sections in this document:

Domain-based namespaces with more than one namespace server

Domain-based namespaces with one namespace server

Stand-alone namespaces

Complete the following procedure to export a stand-alone namespace configuration.

To export the namespace configuration to an export file

1. On the destination server, open a Command Prompt window.
2. Type **DFSUtil.exe root export \\SourceServer\Namespace FileName** the following to export the standalone namespace to a file (where *FileName* represents the exported file), and then press ENTER.

Domain-based namespaces with more than one namespace server

For more than one namespace server, remove the namespace server from the namespace by using DFSUtil.exe.

To remove the namespace server from the namespace

1. On the destination server, open a Command Prompt window.

2. Type **DFSUtil.exe target remove <\\SourceServer\Namespace>**, and then press ENTER.

Domain-based namespaces with one namespace server

There are two options that you can use in this scenario: Export the namespace settings or add a temporary server to the namespace.

To export namespace settings

1. On the destination server, open a Command Prompt window.
2. Type **DFSUtil.exe root export \\Domain\Namespace FileName** where *FileName* represents the file containing settings for export, and then press ENTER.

   Note
   For each namespace, there must be a different file name to export settings.
3. Repeat step 2 for each namespace for which the source server is a namespace server.

You can use either of the following two options if a temporary server can be added to the namespace. This provides the ability to maintain the namespace online while the migration progresses. If this is not possible, follow the steps in To remove the namespace server from the namespace instead.

To add a temporary server to the namespace by using DFS Management

1. In the left pane, select the namespace to be migrated.
2. Click the **Namespace servers** tab.
3. Select **Add Namespace Server**.
4. In the **Namespace server** box, type the name of the temporary server, and then click **OK**.

The temporary server will be added to the namespace.

To add a temporary server to the namespace by using DFSUtil.exe

1. Create a shared folder for the namespace on the temporary server with the same permissions as on the source server.
2. On the destination server, open a Command Prompt window.
3. Type **DFSUtil.exe target add \\TemporaryServer\Namespace**and then press ENTER.

```
DFSUtil.exe target add <\\TemporaryServer\Namespace>
```

The temporary server will be added to the namespace.

After the namespace settings are exported or a temporary server is added to the namespace, the namespace source server can be removed from the namespace as described in To remove the namespace server from the namespace.

Inventory advanced registry keys

This section describes the process for determining if there are any settings that have been applied to the DFS Namespace component on the source server. These settings are stored in the registry and set or viewed through the dfsutil.exe tool. To inventory these settings, run the following commands from the destination server:

```
DFSUtil.exe server registry DfsDnsConfig <SourceServer>

DFSUtil.exe server registry LdapTimeoutValue <SourceServer>

DFSUtil.exe server registry SyncInterval <SourceServer>
```

Note the setting for any registry modification. Registry keys that have not been modified display a value similar to the following:

```
<KeyName> does not exist in the Registry.
```

DFS Replication configuration

To migrate DFS Replication settings, use the following Microsoft Enterprise Support blog series:Replacing DFSR Member Hardware or OS.

File Server Resource Manager configuration on the source server

When you migrate File Server Resource Manager, remember to use the same drive letters for the destination server volumes as for the source server. This is required because the File Server Resource Manager migration requires that the drive letter remains the same.

1. Stop the File Server Resource Manager and File Server Storage Reports Manager services. You can stop these services in Windows PowerShell by using the following command: **Stop-Service –Name “srmsvc”, “srmreports”**.
2. Export the File Server Resource Manager configuration. You can export the File Server Resource Manager configuration in Windows PowerShell by using the following command: **Export-SmigServerSetting -FeatureID FS-Resource-Manager -Path <storepath\FSRM> -Verbose**.
3. For each volume, get the configuration files by running the following commands in the Windows PowerShell session.
   a. Stop the file screen driver. Type **fltmc detach datascrn <VolumeLetter>:**, and then press ENTER.
   b. Stop the quota driver. Type **fltmc detach quota <VolumeLetter>:**, and then press ENTER.
   c. Grant Read permissions to the Administrator account for the "<*VolumeLetter*>:\System Volume information\SRM" folder and the following child objects:
      - takeown /F "<VolumeLetter>:\System Volume Information" /A /R /D Y
      - cacls "<VolumeLetter>:\System Volume Information" /T /E /G Administrators:F
      - attrib -S -H "<VolumeLetter>:\System Volume Information*" /S /D
   d. Copy the following files from the SRM folder to an external storage device:
      - Quota.xml

- Quota.md
- Datascrn.md
- DataScreenDatabase.xml

e. Start the file screen driver. Type **fltmc attach datascrn <VolumeLetter>:**, and then press ENTER.

f. Start the quota driver. Type **fltmc attach quota <VolumeLetter>:**, and then press ENTER.

4. Restart the File Server Resource Manager and File Server Storage Reports Manage services. Type **Start-Service -name "srmsvc","srmreports"**, and then press ENTER.

5. Configure scheduled reports.

   File Server Resource Manager reports and classification rule configurations are dependent on the drive letters and mount points. Any drives or mount points on the source server that are used by report or classification rule configurations must be available on the destination server or the configurations must be updated during import.

   To configure scheduled reports, follow step (a). However, if you are migrating from Windows Server 2003, follow step (b).

   - To configure scheduled reports on all servers except Windows Server 2003, run the following commands in a Windows PowerShell session on the source server that was opened with elevated user rights.
     - To get a list of all the task names associated with storage reports: `storrept r l`
     - For each task name listed run the following command on the source server: `schtasks /query /tn:"TASKNAME" /xml > "TASKNAME.xml"`
   - To configure scheduled reports when you migrate from Windows Server 2003:
     - On the source server, do the following:
       - Open File Server Resource Manager.
       - In storage report management, for each report task, note the report task, target, and schedule.
     - On the destination server, after you import the file server resource manager configuration, do the following:
       - Open File Server Resource Manager.
       - In **Storage Report Management**, for each report task, edit the report task properties.
       - On the **Schedule** tab, manually add the appropriate schedule for the report.

6. Configure scheduled file management tasks. This step applies only to source servers that are running Windows Server 2008 R2 or Windows Server 2012.

   a. To display a list of all task names associated with file management tasks, type the following command on the source server in a Windows PowerShell session opened with elevated user rights:

   ```
   (new-object -com
   Fsrm.FsrmFileManagementJobManager).EnumFileManagementJobs()
   ```

b. For each entry listed, locate the task element, and then type the following command:

```
Schtasks /query /tn:"TASK" /xml > "TASK.xml"
```

7. Export the classification schedule. This is only applicable to servers running Windows Server 2008 R2 or Windows Server 2012 that already have a classification schedule configured. From an elevated command prompt, type the following command:

```
Schtasks /query /tn:”FsrmAutoClassification{c94c42c4-08d5-473d-8b2d-98ea77d55acd}” /xml > “classification.xml”
```

Shadow Copies of Shared Folders

The following procedures describe how to migrate shadow copy settings.

To migrate shadow copy settings

1. Open Windows Explorer on the source server to view shadow copy storage locations and the creation schedule.

   Important
   This procedure applies to shadow copies for a server running the full installation option of Windows Server. If your source server is running the Server Core installation option of Windows Server, skip this procedure and follow the instructions in the following section: To migrate shadow copies in a Server Core installation.

2. For each volume on the source server, right-click the volume, select **Configure Shadow Copies**.

   On source servers that are running Windows Server 2003, right-click the volume, click **Properties**, and then click the **Shadow Copies** tab.

3. Click **Settings**, and note the location and size of the shadow copy storage.
4. Click **Schedule** and note the details for the snapshot creation task.

To migrate shadow copies in a Server Core installation

1. Log on to the computer that is running a Server Core installation remotely as follows:
   a. In Server Manager, click **Tools**, and then click **Computer Management**.
   b. In the **Computer Management** snap-in pane, right-click the top node, and then click **Connect to another computer**.
2. Type the computer name, and then click **OK**.
3. Expand **System Tools**, right-click **Shared Folders**, click the **All Tasks** tab, and then click **Configure Shadow Copies**.
4. For each volume on the source server, right-click the volume, select **Configure Shadow Copies**, click **Settings**, and note the location and size of the shadow copy storage.
5. Click **Schedule**, and then note details for the snapshot creation task.

Migrate local users and groups to the destination server

Before migrating data and shared folders, or completing your migration of the FSRM configuration, you must migrate local users and groups. Export local users and groups from the source server, and then import local users and groups to the destination server.

Important

If the source server is a domain member server, but the destination server is a domain controller, imported local users are elevated to domain users, and imported local groups become Domain Local groups on the destination server.

If the source server is a domain controller, but the destination server is not, Domain Local groups are migrated as local groups, and domain users are migrated as local users.

Export local users and groups from the source server

On the source server, export local users and groups to a migration store (as shown in the following example) in a Windows PowerShell session that has been opened with elevated user rights.

```
Export-SmigServerSetting -User All -Group -Path <storepath\UsersGroups> -Verbose
```

You can use one of the following values with the *-user* parameter:

- Enabled: Specify to export only enabled local users.
- Disabled: Specify to export only disabled local users.
- All: Specify to export enabled and disabled local users.

For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?LinkID=258341) on the Microsoft Web site.

You are prompted to provide a password to encrypt the migration store. Remember this password, because you must provide the same password to import from the migration store.

If the path is not a shared location that is accessible to the destination server, you must manually copy the contents of the migration store folder to the destination server or a location that is accessible to the destination server.

Import local users and groups to the destination server

On the destination server, import local users and groups from the migration store to which you exported the users and groups in Export local users and groups from the source server, as illustrated by the following example. Use a Windows PowerShell session that has been opened with elevated user rights.

```
Import-SmigServerSetting -User All -Group -Path <storepath\UsersGroups> -Verbose
```

You can use one of the following values with the *-user* parameter:

- Enabled: Specify to import only enabled local users.

- Disabled: Specify to import only disabled local users.
- All: Specify to import enabled and disabled local users.
- For the list of user attributes that are supported for migration, see the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?LinkID=258341).

You are prompted to provide the same password that you provided during export to decrypt the migration store.

Migrate data

To migrate data, you can copy file data or physically move it, for example, by attaching the storage drive from the source server to the destination server. If you copy the data, follow the copy data migration steps in the following section.

If you physically move the data, follow the steps described in the Physical data migration section later in this document.

Data copy migration

If you are planning a two-phase data copy migration as described in the previous section, note that if files have been deleted on the source server between the start of the first copy and the start of the final copy, copies of the deleted files might have already transferred to the destination server. So if a file is deleted between the two copy processes, the file might still be available on the destination server after the migration is complete. If this is unacceptable in your environment, perform data and shared folder migration in a single phase, and disconnect all users before starting migration.

Important
The file migration portion of the Windows Server Migration Tools is designed for smaller data sets (under 100GB of data). It copies files one at a time over HTTPS. For larger datasets, we recommend using the version of robocopy.exe included with Windows Server 2012 R2 or Windows Server 2012.

To copy data and shared folders and migrate all data without disconnecting users

1. Verify that the destination path has sufficient disk space to migrate the data. If NTFS or folder quota management is enabled on the destination server disk drive, verify that the NTFS or File Server Resource Manager quota limit allows for sufficient free disk space to migrate data. For more information about quota management in File Server Resource Manager, see one of the following.
   - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154277) for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
   - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154241) for Windows Server 2003 R2

   For more information about NTFS quota management, see one of the following.
   - Setting Disk Quotas (http://go.microsoft.com/fwlink/?LinkId=154243) for Windows

Server 2008, Windows Server 2008 R2, and Windows Server 2012

- Enable disk quotas (http://go.microsoft.com/fwlink/?LinkId=154245) for Windows Server 2003 and Windows Server 2003 R2

2. Ensure that you have completed the migration of local users and groups.

   Send-SmigServerData and **Receive-SmigServerData** cmdlets must be run on the source and destination server within five minutes of each other. By default, **Send-SmigServerData** and **Receive-SmigServerData** time out if a connection cannot be established within 300 seconds (five minutes). This maximum connection time-out for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets is stored in the following registry key, which is user-defined.

   Key: HKEY_LOCAL_MACHINE\Software\Microsoft\ServerMigration

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents connection time-out, in seconds)

   If a value larger than 3600 is specified, 3600 seconds (1 hour) is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key (http://go.microsoft.com/fwlink/?LinkId=147298) on the Microsoft Web site.

3. Use the following command to run the **Receive-SmigServerData** cmdlet on the destination server. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Receive-SmigServerData
   ```

   Note
   All output for the Send and Receive operations occurs on the source server only. The destination server will appear to be done before the operation has completed.

4. Use the following command to run the **Send-SmigServerData** cmdlet on the source server to migrate data and shared folders. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Send-SmigServerData -ComputerName <DestinationServer> -
   SourcePath d:\users -DestinationPath d:\shares\users -Recurse
   -Include All -Force
   ```

The destination data location does not have to be the same as the source location, and you can change it, if desired.

Notes

The Server service startup type must be set to Automatic on the destination server for shared folder migration to complete.

Data that is transferred is encrypted automatically. You are prompted to enter a password to encrypt the transferred data on the source server, and the same password to decrypt the received data on the destination server.

After the first data copy is finished, you must freeze the source server and all data changes.

To disconnect users and migrate new or updated files

1. Make sure that users are notified that they should stop using the source server at this time to prevent any possible data loss. You can run the following command to list all the currently open files to determine the potential impact of performing this step.

   ```
   net file
   ```

2. Disconnect all users from the source server by stopping the LanMan server service.

   ```
   Stop-Service LanmanServer -force
   ```

   Stopping the LanMan Server service invalidates all open remote files to the shared folders on the source server, which can lead to potential data loss. It is best to perform this step when the fewest users are expected to access files on this server.

3. Use the following command to run the **Receive-SmigServerData** cmdlet on the destination server. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Receive-SmigServerData
   ```

4. Use the following command to run the **Send-SmigServerData** cmdlet on the source server to migrate data and shared folders. Use a Windows PowerShell session that is running with elevated user rights.

   ```
   Send-SmigServerData -ComputerName <DestinationServer> -
   SourcePath d:\users -DestinationPath d:\shares\users -Recurse
   -Include All -Force
   ```

5. If your scenario requires migrating reparse points, hard links, and mount points, recreate them on the destination server by using the **mklink** command for reparse points and hard links, and using the **mountvol** command for mounted volumes. For more information about these commands, enter `mklink /?` or `mountvol /?` in a Windows Command Prompt.

It is important to maintain the same destination path that you used during the first copy of data and shared folders. The cmdlets transfer files, folders, and shared folders only if they do not exist on the destination server, or if there is a new version on the source server.

Physical data migration

The next sections describe data migration by physically moving external drives or logical unit numbers (LUNs).

Using disk drives or LUNs to migrate data from the source server to the destination server

You can migrate data from the source server by moving the disk drives. Or, if your data resides on a LUN storage device, you have the option of moving the file server data by masking the LUNs from the source server and unmasking them on the destination server.

For the ideal migration, make sure that you maintain the same mapping of the drive letters (for example, drive D) and the volume IDs (see the following explanation) so that relevant data and application information remains as consistent as possible during the move.

Caution

You should not move a disk drive or LUN if it contains both data and the operating system.

Benefits of physical migration:

- For large amounts of data, this is a faster operation.
- You maintain all data on the disk drive, such as hard links and mount points.
- Shadow copies are preserved if the shadow copies are on the migrated disk drive.

Potential issues to be aware of:

- Permissions for local users that are not default computer accounts (such as local administrators) will be lost even if the same user name is used when creating the user account on the destination server.
- Encrypted files (EFS) cannot be migrated.
- Encrypted volumes with BitLocker cannot be migrated without first decrypting the volumes.
- Remote Storage cannot be migrated.
- When you are physically migrating disk drives that have File Server Resource Manager quotas enabled on them, it is a best practice to dismount the drive gracefully to avoid marking the quotas as dirty. Otherwise, unnecessary scans may occur later.

To migrate data by physically moving the disk drive or by masking and unmasking the LUNs

1. Collect information on the source server.

   Tip

   You can use Server Manager or Windows PowerShell on a computer running Windows Server 2012 or Windows 8 to collect information from source computers running Windows Server 2012.

   a. Record the drive letter and volume label for each data volume on the source server that you would like to move to the destination server.

   b. On the source server, export the volume GUID paths by exporting the following registry key to a file: HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices. To do this, open the Registry Editor (regedit.exe), browse to the registry key, right-click the registry key, and clicking **Export**.

      Alternatively, to export the volume GUID paths from a server running Windows Server 2012 or Windows Server 2008 R2, open a Windows PowerShell session, and then type the following commands, where *<SourceServer>* is the name of the source server, *<Domain\User>* is a user account with administrative permissions on the source server and *<LocalPath>\<Filename>* is a local path and filename of the exported registry keys:

```
Enter-PSSession <SourceServer> -Credential <Domain\User>
Regedit.exe /E <LocalPath>\<Filename>.reg
"HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices"
```

Note

To user Server Manager or Windows PowerShell to remotely collect information from earlier versions of Windows Server you must first setup the source server for remote management. For more information, see Managing Downlevel Windows-based Servers from Server Manager in Windows Server 2012.

c. Open Notepad and copy the exported .reg file. Remove all entries that are in the following form: \DosDevices\D:. Save the.reg file (all remaining entries should be in the following form: \??\Volume{ef93fe94-5dd7-11dd-961a-001e4cdb4059}).

2. Prepare the destination server.

a. In the Server Manager navigation pane, click **File and Storage Services**, and then click **Volumes** to display the Volumes page. Use the Volumes tile to make sure that the drive letters for the data volumes are available. If there is a drive letter that is currently assigned to an existing volume on the destination server, change the drive letter for that volume.

Alternatively, use the Windows PowerShell **Get-Volume** and **Set-Partition** cmdlets. For example, to get any volumes with the drive letters of F, G, or H, type `Get-Volume F,G,H`. To change the drive letter of a partition with the F drive letter, type `Set-Partition -DriveLetter F -NewDriveLetter Z`

b. To import the volume GUID paths into the destination server, copy the.reg file that you created previously to the destination server, and then double-click that file to update the destination server.

3. Move the disk drives or LUNs from the source server to the destination server.

a. On the source server, remove the disk drives or unassign the LUNs by using Storage Manager for SANs. (To open Storage Manager for SANs, click **Start**, click **Administrative Tools**, and then click **Storage Manager for SANs**.) If the source server is running Windows Server 2012, instead use the File and Storage Services role in Server Manager to view the disks or virtual disks (when using storage pools) that you want to move. If the disk is part of a storage pool, on the **Storage Pools** page of the File and Storage Services role right-click the virtual disk, and then click **Detach Virtual Disk**. For other types of disks, on the **Disks** page, right-click the disk that you want to move and then click **Take Offline**.

b. On the destination server, attach each disk drive or assign the LUNs, and then assign the appropriate drive letter by using the **Disks** and **Storage Pools** pages of the File and Storage Services role in Server Manager.

4. If any files or folders on the migrated drive use local users or local groups permissions (except default users and groups), re-create these permission. Note that all domain users and groups permissions will remain intact, assuming that the source server and the destination server are members of the same domain.

Notes

You can use the `icacls` command to modify file and folder permissions (type `icacls /?` in a Command Prompt window for details). Type this command in a Windows PowerShell session or a command prompt that has been opened with elevated user rights.

The list of the default users and groups is available in the topic Default User Accounts and Groups (http://go.microsoft.com/fwlink/?LinkId=149889) on the Microsoft Web site.

Migrate shared folders

If any of the folders on the migrated drive were shared on the source server, and they must be shared on the destination server, the following steps explain how to migrate shared folders.

1. If any of the migrated shared folders use local users and group permissions, ensure that you have completed the migration of local users and groups.

   Send-SmigServerData and **Receive-SmigServerData** cmdlets must be started on the source server and the destination server within five minutes of each other. By default, **Send-SmigServerData** and **Receive-SmigServerData** operations terminate if a connection cannot be established within 300 seconds (five minutes). The maximum connection time-out for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets is stored in the following registry key, which is user-defined.

   Key: HKLM\Software\Microsoft\ServerMigration

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents connection time-out, in seconds). If a value larger than 3600 is specified, 3600 seconds (one hour) is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key (http://go.microsoft.com/fwlink/?LinkId=147298) on the Microsoft Web site.

2. Open port 7000 on the source server and destination server (if this has not already been done).

   For information about how to open a port in Windows Firewall, see File and Storage Services: Appendix A: Optional Procedures.

3. On the destination server:
   a. Open a Windows PowerShell session with elevated user rights and enter the following command: **Receive-SmigServerData**.

4. On the source server:
   a. Open a Windows PowerShell session in Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012. On computers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, the Windows PowerShell session must be opened with elevated user rights. Enter the following command: **Send-SmigServerData -ComputerName <DestinationServerName> -SourcePath <SourcePath> -DestinationPath <DestinationPath> -Recurse -Include Share -Force**

Notes

The <*SourcePath*> value specifies the local path on the source server that contained the shared folder before the drive was migrated. Shared folder information is not stored on the data drive, so do not be concerned that the drive no longer resides on the source server.

The <*DestinationPath*> value specifies the local path on the destination server that contains folders that were previously shared on the source server. Unless the root drive letter or the folder structure has been changed on the migrated drive, <*SourcePath*> and <*DestinationPath*> values should be the same.

During shared folder migration, permissions for local users and groups and domain users and groups are migrated—no manual remapping is required.

LanMan Server service automatically restarts on the destination server, and the shared folders migrate.

DFS Replication migration

If you physically migrated data, clean-up the DFS Replication configuration state, which is stored on the migrated volume:

1. To clean up volumes (for each physically migrated volume)
   a. Navigate to the path <volume>\System Volume Information.

      Note

      This is a hidden system folder. To view this folder: in File Explorer, click **View**, and then select the **Hidden Items** check box.. Also ensure that local administrators are granted **Full Control** of the folder.

   b. Delete the **DFSR** folder and all content in the folder.
   c. Revert any security permissions modifications that you made to perform the migration process.
   d. Repeat this process for all physically migrated volumes.
2. To clean up replicated folders (for replicated folders on physically migrated volumes)
   a. Navigate to the root of a replicated folder.
   b. Delete the **DfsrPrivate** folder and all subfolders.
   c. If the staging folder for the replicated folder is not located in the default location, remove the staging folder and all content in the staging folder.

      Note

      The default location for the staging folder is in the DfsrPrivate folder, and this step is not required if the path is at the default location.

   d. If the Conflict and Deleted folder for the replicated folder is not located in the default location, remove the Conflict and Deleted folder and all content in the Conflict and Deleted folder.

Note

The default location for the Conflict and Deleted folder is in the DfsrPrivate folder, and this step is not required if the path is at the default location.

Use the inventoried information that you collected for the source server to detect all replication groups to which the source server belongs. Add the destination server as a member server to all these replication groups.

Migrate the source server identity

You need to rename the source server and migrate its previous identity to the destination server. You might also need to migrate the source server IP address to the destination server.

Rename the source server

Rename the source server to a temporary name.

Migrate IP address

When a static IP address is used on the source server, it is recommended that the IP address be migrated from the source server to the destination server. This is because client computers locally cache the IP address that is associated with a server name. Client computers will still attempt to access the source server even if it has been renamed.

When the server IP address is not migrated, you must stop the LanMan Server service on the source server. This is done to prevent users from accessing shared folders on the source server after they have been migrated to the destination server. Open a Windows PowerShell session with elevated user rights, and then run the following cmdlet:

```
Stop-Service LanmanServer -Force
```

For more information on IP address migration, see http://go.microsoft.com/fwlink/?LinkId=128513

Rename destination server

Rename the destination server to the name that was originally used for the source server.

Configure DFS Replication on the destination server

Configuration of DFS Replication on the destination server is determined by whether you migrated the data by copying or physically moving it

If you migrated the data by copying it

Follow this procedure to add a replication connection between the source server and the destination server for each replication group on the source server:

1. In Server Manager, click **Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, select **Add Replication Groups to Display**, enter the name of the source, and then click **Show Replication Groups**. Select all of the replication groups that are displayed, and then click **OK**.
3. For each replication group, do the following:
    a. Click the replication group, and then click **New Member**. The **New Member Wizard** appears. Follow the instructions in the wizard to add the destination server to the replication group by using the information from question #2 in the DFS Replication Data Collection Worksheet (File and Storage Services: Appendix B: Migration Data Collection Worksheets).
    b. In the console tree, under the **Replication** node, right-click the replication group that you just added the destination server to, and then click **New Connection**.
    c. Specify the source server and destination server as sending and receiving members, and specify a schedule so that the connection is always enabled. At this point, the replication is one-way.
    d. Select **Create a second connection in the opposite direction** to create a second connection for two-way replication between the sending and receiving members.

If you migrated the data by physically moving it

Follow this procedure to add a replication connection between the destination server and the closest server to the destination server other than the source server:

1. In Server Manager, click **Tools**, and then click **DFS Management**.
2. In the console tree, under the **Replication** node, select **Add Replication Groups to Display**, enter the name of the source, and then click **Show Replication Groups**. Select all of the replication groups that are displayed, and then click **OK**.
3. For each replication group:
    a. Click the replication group, and then click **New Member**. The **New Member Wizard** appears. Follow the instructions in the wizard to add the destination server to the replication group by using the information from question #2 in the DFS Replication Data Collection Worksheet (File and Storage Services: Appendix B: Migration Data Collection Worksheets).
    b. In the console tree, under the **Replication** node, right-click the replication group that you just added the destination server to, and then click **New Connection**.
    c. Specify the destination server as the sending member, and then specify any other server except the source server as the receiving member. Specify the schedule to use for the connection. It is recommended that you select a server that has a good network connection to the destination server as the receiving member.
    d. Select **Create a second connection in the opposite direction** to create a connection for two-way replication between the sending and receiving members.

Notes

The folder does not begin to replicate immediately. The new DFS Replication settings must be replicated to all domain controllers, and each member in the replication group

must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory Domain Services (AD DS) replication latency and the polling interval (60 minutes) on each member.

The **dfsrdiag /pollad** command can be used to force DFS Replication on the source server and destination server to poll AD DS and retrieve the latest configuration information instead of waiting for the next normal polling interval which could be up to 60 minutes.

After DFS Replication on the destination server polls AD DS, it begins to replicate the folders that it configured, and it performs an initial synchronization. Event ID 4102 (MSG_EVENT_DFSR_CS_INITIAL_SYNC_NEEDED) is registered in the event log on the destination server for each replicated folder.

During initial sync, DFS Replication downloads all files in the replicated folders from the source server and builds up a local copy of the database per volume. This process can be time consuming. It is possible to speed up the initial sync by seeding the data from the source server onto the destination server (from the backup that was taken prior to commencing migration).

When the initial sync completes, event ID 4104 (MSG_EVENT_DFSR_CS_INITIAL_SYNC_COMPLETED) is registered for each replicated folder on the destination server. Monitor each replicated folder on the destination server, and check to ensure that all of them have completed the initial sync.

Import settings to the destination server

Follow the procedures in this section to import settings to the destination server.

Note

If the source server is not running Windows Server 2008 R2 or Windows Server 2012, the first procedure in this section does not apply. (This procedure is used to migrate the seed value that is used by BranchCache for the Network Files component, and it enables data that is stored in BranchCache on the source server to be used after it is migrated to the destination server. For information about how to migrate a BranchCache host server, see the BranchCache Migration Guide (http://go.microsoft.com/fwlink/?LinkID=139091).

To set up BranchCache for Network Files migration on the destination server

1. On the destination server, open a Windows PowerShell session with elevated user rights.
2. Type the following command, where *storepath* is the available path that contains the Svrmig.mig file, and then press ENTER.

```
Import-SmigServerSetting -featureid BranchCache -Path
<storepath\BranchCache> -Force -Verbose
```

Group Policy or local policy specific to server message block and Offline Files

Use a Group Policy object or a local policy on the destination server to change the settings to the same values as the source server. These settings are recorded in the SMB and BranchCache data collection worksheets in File and Storage Services: Appendix B: Migration Data Collection Worksheets.

To import server message block settings

1. Do one of the following:
   - If the policies are set by using Group Policy objects, use the Group Policy editing tools to apply appropriate policies to the destination server.
   - If the policies are set by using a local policy, do the following:
     i. On the destination server, open the Local Group Policy Editor snap-in.
     ii. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Security Settings**, click **Local Policies**, and then click **Security Options**.
2. Use a Group Policy object or a local policy to set the following settings to the same values as noted in Export settings. Set the destination server settings to the same values as were noted on the source server for the following settings:
   - Microsoft network server: Amount of idle time required before suspending a session
   - Microsoft network server: Attempt S4USelf to obtain claim information
   - Microsoft network server: Digitally sign communications (always)
   - Microsoft network server: Digitally sign communications (if client agrees)
   - Microsoft network server: Disconnect clients when logon hours expire

   Note
   For any setting that is controlled by Group Policy, you must have the same Group Policy object apply to the destination server, or you can set the local policy of the destination server to get the same behavior.

   On destination servers that are running the Server Core installation, run the **secedit** command to change local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

Note

The following procedure applies only if the source server is Windows Server 2008 R2 or Windows Server 2012.

To import Offline Files settings

1. Do one of the following:
   - If the policies are set by using Group Policy, use the Group Policy editing tools to apply appropriate policies to the destination server.

- If the policies are set by using local policy, do the following:
  i. On the destination server, open the Local Group Policy Editor snap-in.
  ii. In the snap-in tree pane, click **Computer Configuration**, click **Windows Settings**, click **Administrative Templates**, click **Network**, and then click **LanMan Server**.

2. Use a Group Policy object or a local policy to set the destination server policy settings to the same values as the source server settings for **Hash Publication for BranchCache** and **Hash Version support for BranchCache** settings.

   On destination servers that are running the Server Core installation, run the **secedit** command to change local policy settings (for more information about **secedit**, type **secedit /?** at a command prompt).

DFS Namespace configuration

Complete the configuration of namespaces on the destination server. The procedure you use depends on whether you want a stand-alone or a domain-based namespace.

Stand-alone namespaces

Domain-based namespaces with more than one namespace server

Domain-based namespaces with one namespace server

Stand-alone namespaces

If you want a stand-alone namespace, you must first create the namespace on the destination server. You can do this by using DFS Management, or the **DFSUtil.exe** command-line utility.

To create the namespace on the destination server

1. Do one of the following:
   - On the destination server, open DFS Management, and create the namespace by using the same name as on the source server.
   - On the destination server, in a Command Prompt window opened with elevated user rights, type the following, and then press ENTER.

     ```
     Dfsutil.exe root addstd <\\DestinationServer\Namespace>
     ```

To import a namespace configuration from the export file

1. On the destination server, in a Command Prompt window opened with elevated user rights, type the following (in which *filename* represents the file name into which you exported namespace settings from the source server in To export the namespace configuration to an export file), and then press ENTER.

   ```
   Dfsutil.exe root import set <filename>
   <\\DestinationServer\Namespace>
   ```

Domain-based namespaces with more than one namespace server

If you have more than one domain-based namespace server, you can add namespace servers to your destination server by using DFS Management or the **DFSUtil.exe** command-line utility.

To use DFS Management

1. Select the namespace being migrated in the left pane.
2. Click the **Namespace servers** tab in the right pane.
3. Select **Add Namespace Server**.
4. In the dialog box that opens, type the name of the destination server, and then click **OK**.

The destination server is added to the namespace.

To use DFSUtil.exe

1. On the destination server, open a Command Prompt window.
2. Type the following command, and then press ENTER.

```
DFSUtil.exe target add <\\DestinationServer\Namespace>
```

Domain-based namespaces with one namespace server

This section applies only if a temporary server was not added to the namespace. If you added a temporary server to the namespace as part of your export process, see Domain-based namespaces with more than one namespace server.

To create the namespace on the destination server

1. Do one of the following:
   a. In DFS Management on the destination server, create the namespace with the same name that was used on the source server.
   b. Type the following command at a command prompt, and then press ENTER.

```
Dfsutil.exe root adddom <\\DestinationServer\Namespace>
```

To import a namespace configuration from the export file

1. On the destination server, open a Command Prompt window.
2. Type the following command (in which *filename* represents the export file names you created in To export namespace settings). Run this command for each of the namespaces for which the source server was a namespace server.

```
DFSUtil.exe root import set <Filename>
\\DestinationServer\Namespace
```

Note
For each namespace, there must be a file name from which settings are imported.

To manually reset delegation permissions on the namespace

1. On the destination server, open DFS Management.
2. Set the permissions that you inventoried in DFS Namespace configuration. When complete, close DFS Management.

If any advanced registry keys were configured on *SourceServer*, use **DFSUtil.exe** to configure *DestinationServer* to have the same registry key settings. Run the following commands on the destination server to set the advanced registry keys.

To set advanced registry keys

1. On the destination server, open a Command Prompt window.
2. Run the following commands to set the advanced registry keys by using **DFSUtil.exe**.

```
DFSUtil.exe server registry DfsDnsConfig set
<DestinationServer>
DFSUtil.exe server registry LdapTimeoutValue set <Value>
<DestinationServer>
DFSUtil.exe server registry SyncInterval set <Value>
<DestinationServer>
```

File Server Resource Manager configuration on the destination server

When you are migrating File Server Resource Manager, remember to use the same drive letters for the destination server volumes as for the source server. This is required because File Server Resource Managermigration requires that the drive letter remains the same.

1. Stop the File Server Resource Manager and File Server Storage Reports Manager services.

   Open a Windows PowerShell session with elevated user rights, and then run the following command:

   ```
   Stop-Service -name "srmsvc","srmreports"
   ```

2. Type the following in the Windows PowerShell session, and then press ENTER.

   ```
   Import-SmigServerSetting -FeatureID FS-Resource-Manager -Path
   <storepath\FSRM> -Force
   ```

   Notes

   If the Windows features that you are migrating have not been installed on the destination server, the **Import-SmigServerSetting** cmdlet installs them as part of the import process, along with any Windows features that they depend on. Some Windows features might require that you restart the destination server to complete the installation. After restarting the computer, you must run the cmdlet again with the **-Force** parameter to complete the import operation.

Importing FSRM settings to the destination server replaces any global FSRM configuration information that is already on the destination server.

3. Set the configuration files for each volume.

   Type the following commands in a Windows PowerShell session, and then press ENTER.

   Note
   Running the following commands on a clean computer returns an error message. It is safe to ignore this error message.

   a. Type the following code to stop the file screen driver:

      ```
      fltmc detach datascrn <VolumeLetter>:
      ```

   b. Type the following code to stop the quota driver:

      ```
      fltmc detach quota <VolumeLetter>:
      ```

   c. Add administrator Write permissions to the "<VolumeLetter>:\System Volume information\SRM" folder and the following subfolders:

      - takeown /F "<VolumeLetter>:\System Volume Information" /A /R /D Y
      - cacls "<VolumeLetter>:\System Volume Information" /T /E /G Administrators:F
      - attrib -S -H "<VolumeLetter>:\System Volume Information*" /S /D

   d. Copy the following files from the external storage to the SRM folder:

      - Quota.xml
      - Quota.md
      - Datascrn.md
      - DataScreenDatabase.xml

   e. Type the following code to start the file screen driver:

      ```
      fltmc attach datascrn <VolumeLetter>:
      ```

   f. Type the following code to start the quota driver:

      ```
      fltmc attach quota <VolumeLetter>:
      ```

4. Restart the File Server Resource Manager and File Server Storage Reports Manager services.

   Type the following in a Windows PowerShell session, and then press ENTER.

   ```
   Start-Service -name "srmsvc","srmreports"
   ```

5. Configure scheduled reports and file management tasks.

   For each scheduled report, you need to create a scheduled task on the destination server.

   Note
   File Server Resource ManagerReports and classification rule configurations are dependent on the drive letters and mount points. Any drives or mount points on the source server that are used by report or classification rule configurations must be available on the destination server or the configurations must be updated during import.

After you have an eXtensible Markup Language (XML) file for each task, copy them to the destination server and run the following command in a Windows PowerShell session on the destination server for each task:

```
schtasks /create /xml:"TASKNAME.xml" /tn:"TASKNAME"
```

6. Import the classification schedule. The classification schedule requires a scheduled task on the destination server.

```
schtasks /create /xml:"classification.xml"
/tn:"FsrmAutoClassification{c94c42c4-08d5-473d-8b2d-
98ea77d55acd}"
```

classification.xml is the name of the XML file that was exported from the target server.

Shadow Copies of Shared Folders

Apply the same settings from the source server to the corresponding volumes on the destination server.

To migrate shadow copy settings for Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012

1. To configure shadow copies, right-click each volume on the destination server that had shadow copies configured on the source server, right-click the volume and select **Configure Shadow Copies**.
2. Click **Settings** and verify that the location and size of shadow copy storage matches the settings from the source server.
3. Click **Schedule** and verify that the details for the snapshot creation task match the settings from the source server.

To migrate shadow copy settings for a Server Core installation

1. Log on to the destination server that is remotely running the Server Core installation by doing the following:
   a. In Server Manager, click **Tools**, and then click **Computer Management**.
   b. In the **Computer Management** snap-in tree pane, right-click the top node, and then click **Connect to another computer**.
2. Enter the computer name, and then click **OK**.
3. Expand **System Tools**, right-click **Shared Folders**, click the **All Tasks** tab, and then click **Configure Shadow Copies**.
4. For each volume on the destination server that had shadow copies configured on the source server, right-click the volume, select **Configure Shadow Copies**, click **Settings**, and verify that the location and size of shadow copy storage match the settings from the source server.
5. Click **Schedule**, and verify that these details for the snapshot creation task match the settings from the source server.

Deduplication

Use the following section to migrate Deduplication.

Migrating Deduplication from Windows Server 2012 to Windows Server 2012

All configuration information needed for migration is included on the deduplicated volume.

If a disk is physically moved, or if a deduplicated volume is restored from a backup onto a different Windows Server 2012 computer, install the Deduplication role service using Server Manager on the new computer. If the Deduplication role service is not installed on the new server, only normal non-deduplicated files will be accessible. Once a volume has been mounted, the deduplication filter will detect that the volume is deduplicated and will redirect input/output requests appropriately.

Note

Any previous custom deduplication job schedules that were created using Task Scheduler must be created again on the new computer using Task Scheduler.

Migrating SIS from Windows Storage Server 2008 to Windows Server 2012

Volumes that have been created and optimized using the down-level Windows Storage Server version of deduplication, Single Instance Storage (SIS), should not be enabled for data deduplication. Microsoft recommends that SIS be removed from the volume by using SISAdmin.exe within Windows Storage Server to remove SIS or by copying the data to a different volume that is not running SIS prior to migrating the volume.

Windows Server 2012 supports reading and writing to SIS-controlled volumes, but you cannot continue to SIS files using Windows Server 2012. You can install the SIS filter driver on Windows Server 2012 by installing the SIS-Limited feature using the following command syntax:

```
dism /online /enable-feature:SIS-Limited
```

The SIS filter driver can be loaded so that you can read SIS-controlled volumes and the data can be copied to a non-SIS controlled volume so that data deduplication can be installed on the volume. Note that Windows Server 2012 does not support sisadmin.exe and cannot be used to remove SIS from a volume.

1. You should remove SIS from your volumes before installing the Windows Server 2012 data deduplication feature. (This process is also known as un-SIS.)
2. Do not restore SIS links from a backup to a Windows Server 2012 deduplicated volume.
3. Restoring SIS volumes to Windows Server 2012 is supported if you load the SIS-Limited filter.

Migrating SIS volumes

You have several options when it comes to migrating Windows Storage Server 2008 volumes to Windows Server 2012 to take advantage of the new Data Deduplication feature.

You can migrate your existing SIS-installed Windows Storage Server 2008 volumes to Windows Server 2012, however, migration is not automatic. Single Instance Storage (SIS) and data deduplication are mutually-exclusive technologies.

Caution
You will need to open the volumes in Windows Storage Server 2008 first, un-SIS them, and then uninstall SIS before migrating to Windows Server 2012 as described in the procedures below.

To unSIS a Windows Storage Server 2008 or 2008 R2 SIS volume type **sisadmin.exe [/m <server>] [/u <volumes>]** where:

/m <server> - Shifts the focus of the command line to a remote server. If the /m option is not specified, the command line is applied to the local server. <server> can be expressed as a host name, fully qualified domain name (FQDN), or as an IP address.

/u <volumes> - Is used to un-SIS a volume (that is, to restore all file copies, and remove reparse points).

For each command option that uses <volumes> as a parameter, <volumes> represents a space-delimited list of volume names (for example: d: e: f: g:). For example:

To unSIS or remove SIS entirely from the F: volume of a remote server using the IP address of the server, you might use the following command: **sisadmin.exe /m 192.168.1.50 /u F:**

See Also

Migrate File and Storage Services to Windows Server 2012

File and Storage Services: Prepare to Migrate

File and Storage Services: Verify the Migration

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Verify the Migration

Verify that the migration was successful. Follow the appropriate verification steps based on the File and Storage Services role services that have been migrated.

The following overview describes the steps to verify the migration.

Verify the File Services migration

Perform the following tasks to verify the File and Storage Services role migration.

- Verify the File Services migration (only if running Windows Server 2008 R2 or Windows Server 2012)
- Verify migration of local users and groups
- Verify data and shared folder migration
- Verify the migration of DFS Namespaces
- Verify the configuration on other computers
- **Verify the File Server Resource Manager migration**

Verify migration of BranchCache for Network File Services server key

Perform this step only if the source server is running Windows Server 2008 R2 or Windows Server 2012:

Verify that the server key was migrated correctly by checking the key value, and ensure that the key values are identical on source server and the destination server, as shown in the following example:

```
Key: HKLM\Software\Microsoft\WindowsNT\CurrentVersion\PeerDist\SecurityManager\Restricted
Value: Seed
```

Verify migration of local users and groups

Check that all the local users and groups you expected to migrate are present on the destination server by comparing the list of users and groups on the Local Users and Groups snap-in on the source server with the list on the destination server.

To open the Local Users and Groups

1. In Server Manager, click **Tools**, and then click **Computer Management**.

Alternatively, you can compare the list of users and groups on the source server and destination server by typing **net** commands in a Command Prompt window.

- To get the list of all local users and save it in a text file, type the following command:

  ```
  net user > localusers.txt
  ```

- To get the list of all local groups and save it in a text file, type the following command:

  ```
  net localgroup > localgroups.txt
  ```

Verify data and shared folder migration

1. Check that all the data you expected to migrate are present at the correct location on the destination server and that they have the correct permissions associated with them.

To list files and folders with their permissions, type the following command in a Command Prompt window or in a Windows PowerShell session opened with elevated user rights:

```
icacls <path>
```

2. Verify that all the expected shared folders have migrated and that they have the correct permissions associated with them. To list all shared folders and their permissions, type the following command in a Windows PowerShell session opened with elevated user rights:

```
gwmi win32_share | %{net share $_.name}
```

Verify the migration of DFS Namespaces

The procedure that you use to verify the migration of DFS Namespaces depends on whether your namespaces are stand-alone or domain-based.

To verify the migration of a stand-alone namespace

1. Open DFS Management on the destination server.
2. Right-click **Namespaces**, or click the **Action** menu.
3. Click **Add Namespaces to Display**.
4. Type the name of destination server, and then click the **Show Namespaces** button.
5. Select the namespace that you migrated, and then click **OK**.
6. In the namespaces tree, click the namespace that you migrated.
7. Click the **Namespace** tab, and check that all the namespace links are present.
8. Click the **Namespace server** tab, and check that the destination server is listed.
9. Right-click the destination server name, and then click **Open in Windows Explorer**. All namespace links should be visible in the new window.
10. Using DFSUtil.exe on the destination server, type the following command for each stand-alone namespace:

    ```
    Dfsutil.exe root \\DestinationServer\Namespace
    ```

    The information displayed should contain the destination server and all the namespace links.

To verify the migration of a domain-based namespace

1. Open DFS Management, and then right-click **Namespaces** or click the **Action** menu.
2. Click **Add Namespaces to Display**.
3. Type the name of the domain where the namespace is located, and then click the **Show Namespaces** button.

   Select the namespace that you migrated, and click **OK**.
4. In the namespaces tree, click the namespace that you migrated.
5. Click the **Namespace** tab, and check that all the namespace links are present.
6. Click the **Namespace** server tab, and check that all the namespace servers are listed.
7. Right-click the destination server name, and then click **Open in Windows Explorer**. All

namespace links should be visible in the new window.

8. Using DFSUtil.exe on the destination server, type the following command in a command Prompt window, where *domain**namespace* is the name of the appropriate domain and namespace that you migrated.

   ```
   Dfsutil.exe root <\\Domain\Namespace>
   ```

   The information displayed should contain all namespace servers and namespace links.

Verify the configuration on other computers

To verify that File and Storage Services migration completed successfully on other computers, you must test the configuration on the client computers in your enterprise.

To verify DFS Namespaces on a client computer

1. Log on to a client computer with the credentials of a user who has access to the migrated namespace.
2. Verify that you can access the namespace by using File Explorer, a command prompt window, or another application, by entering the same name that you used before the migration.

Verify the File Server Resource Manager migration

Follow these steps to ensure that File Server Resource Manager migrated:

1. If any custom actions are configured for quota notification or file management tasks, the user should ensure that the folders that contain the executable files configured for the actions and the working folders have the correct access control lists.
2. As a best practice, ensure that all e-mail message text for notifications, reports, and so on migrated correctly.
3. Administrators should send a test e-mail message through the File Server Resource Manager console to verify that the Simple Mail Transfer Protocol (SMTP) server is configured correctly for the destination server.
4. Ensure that expiration folders that are used by File Management Tasks are reachable on the destination server.
5. Ensure that executable files that are used by custom actions (such as quota notifications and file management tasks) are accessible or executable on the destination server.

See Also

Migrate File and Storage Services to Windows Server 2012

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets
File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Post-Migration Tasks

This topic explains how to complete the migration if it was successful, and how to roll back or troubleshoot the migration if it failed.

Completing the migration

After you verify the migration, retire the source server.

Retire File and Storage Services on the source server

After you complete and verify the migration, the source server can be shut down or disconnected from the network.

Remove DFS Namespaces from the source server

The procedure you use to remove DFS Namespaces from the source server depends on whether the namespaces are stand-alone or domain-based. If you want to remove the namespace from the source server, you must use **DFSUtil.exe**.

Note
By default, clients cache the list of namespace servers for 300 seconds (five minutes), so we recommend that you do not run the **DFSUtil.exe remove** command within five minutes of completing verification of the DFS namespace migration. During migration, clients have only the temporary server in the cache of namespace servers. Waiting five minutes after you add the destination server to the namespace allows clients to list the destination server in their cache.

To remove stand-alone namespaces

1. Open a Command Prompt window on the destination server.
2. Type the following code, and then press **Enter**.

```
Dfsutil.exe root remove <\\SourceServer\Namespace>
```

To remove domain-based namespaces with one namespace server

1. On the destination server, open a Command Prompt window.
2. Type the following, and then press **Enter**.

```
DFSUtil.exe target remove <\\TemporaryServer\Namespace>
```

Notes

This procedure applies only if a temporary server was added to the namespace for migration purposes.

For domain-based namespaces with more than one namespace server, no additional actions are required.

Restoring File and Storage Services in the event of migration failure

The following sections describe how to restore the File and Storage Services server role in the event of migration failure.

Roll back DFS Namespaces

The steps that you perform to roll back DFS Namespaces depend on whether the namespaces are stand-alone or domain-based, and whether a temporary namespace was created during the migration process.

To roll back DFS Namespaces (do one of the following)

1. For stand-alone namespaces, no action is required other than migrating the identity back to the source server.
2. For domain-based namespaces with greater than one namespace server, or if a temporary server was added to a namespace that initially had only one namespace server, do the following:
   a. Remove destination server from the namespace.
   b. Migrate identity and shared folder information to the source server.
   c. Add the source server to namespace.
3. For domain-based namespaces with only one namespace server, where no temporary namespace server was added during migration, do the following:
   a. Migrate identity and shared folder information to source server.
   b. Verify the export file for the namespace that was created during migration is still available.
   c. Delete the namespace.
   d. Create the namespace on the source server.
   e. Import the namespace configuration from the export file created during the migration.
   f. Manually reset delegation permissions to the namespace.

Note

Another option to migrate domain-based namespaces with one namespace server is to temporarily add a second namespace server before the migration, and then remove the temporary server after the migration.

Roll back data and shared folders

If no changes have been made to migrated files, folders, and shared folders on the destination server and this data has not been deleted from the source server, no additional steps to roll back data and shared folders are required.

If the migrated files, folders, or shared folders may have been modified on the destination server by the administrators or users, perform the following steps to synchronize the changes from the destination server back to the source server:

1. Type the following code in a Command Prompt window to copy the updated migrated data (files and folders) from the destination server back to the source server:

   ```
   robocopy <copy from path> <copy to path> /E
   ```

   This command can be executed on the source server or on the destination server, and it will recursively copy updated data. Type `robocopy /?` in a Command Prompt window for additional copy options, including options to copy file and folder permissions.

   Caution
   Permissions that you set for non-default local users and groups will not copy properly and need to be created manually.

2. Compare the lists of shared folders and their permissions on the source server and destination server and manually synchronize any changes.

   To list all shared folders and their permissions, type the following command in a Windows PowerShell session that has been opened with elevated user rights:

   ```
   gwmi win32_share | %{net share $_.name}
   ```

Roll back migration on the other computers in the enterprise

If the migration failed, verify that the other computers in the enterprise can access the source server after you roll back the migration data.

Troubleshooting migration issues

Troubleshooting tips include the following:

- For physical migration issues:

  When some files are migrated physically and others are copied, there is a chance that the File Server Resource Manager configuration is not synchronized. To remedy this, delete and create new copies of the Quota.md and Datascrn.md files.

- For domain-joined machines:

If a custom action (quota notification or file management task) fails to execute with an access-denied failure and a corresponding event log, you should remove the custom action and create it on the destination server.

Troubleshoot data migration that does not complete

If the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets run indefinitely without completing, your destination server might not have sufficient disk space or a large enough File Server Resource Manager or NTFS quota limit to allow for data migration to finish. To determine whether insufficient disk space is preventing the data send-receive process from completing, do the following on the destination server.

1. Open %*localappdata*%/Svrmig/Log/SetupAct.log.
2. Review the most recent log entries. If the following exception occurs, your destination server has insufficient disk space or File Server Resource Manager or NTFS quota limits to complete data migration.

   ```
   Win32Exception: unable to write to FileStream: There is not enough space on the disk.
   ```

To resolve this issue, do the following:

1. Press **Ctrl+C** to cancel **Send-SmigServerData** and **Receive-SmigServerData** on both source and destination servers.
2. Check for sufficient disk space on the destination server's hard disk drive. If the destination server's hard disk drive has insufficient space, do one of the following.
   - Clear additional space.
   - Identify a different hard disk drive that has sufficient space.
3. If the destination server's hard disk drive, the destination path, or any folders that contain the destination path have an File Server Resource Manager or NTFS quota enabled, and the quota limit does not allow for sufficient disk space to migrate data, do one of the following.
   - Increase the quota limit to set sufficient disk space to migrate the data. For more information about FSRM quota management, see one of the following.
     - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154277) for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
     - Quota Management (http://go.microsoft.com/fwlink/?LinkId=154241) for Windows Server 2003 R2

     For more information about NTFS quota management, see one of the following.
     - Setting Disk Quotas (http://go.microsoft.com/fwlink/?LinkId=154243) for Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
     - Enable disk quotas (http://go.microsoft.com/fwlink/?LinkId=154245) for Windows Server 2003 and Windows Server 2003 R2

- Identify a different hard disk drive that already has sufficient space and File Server Resource Manager or NTFS quota limits.

4. Run the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets again, specifying a destination path that has sufficient disk space, and large enough File Server Resource Manager or NTFS quota limits, if applicable.

Troubleshoot data migration connectivity

If the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets cannot establish connectivity, verify the following conditions and then try again:

1. In the **Send-SmigServerData** command on the source server, the *ComputerName* parameter correctly specifies the name of the destination server.
2. The **Receive-SmigServerData** and **Send-SmigServerData** commands are entered on the destination server and the source server respectively within five minutes of one another. This is the default maximum connection time-out for **Send-SmigServerData** and **Receive-SmigServerData**. You can change the maximum connection time-out for the **Send-SmigServerData** and **Receive-SmigServerData** cmdlets by modifying the following user-defined registry key on the source server and destination server.

   Key: HKEY_Local_Machine\Software\Microsoft\ServerMigration

   Value: MaxConnectionTime (REG_DWORD)

   Data: Between 1 and 3600 (represents the connection time-out in seconds). If a value larger than 3600 is specified, 3600 seconds is used as the maximum connection time-out.

   For information about how to create a Windows Registry key, see Add a Registry Key (http://go.microsoft.com/fwlink/?LinkId=147298) on the Microsoft Web site.
3. The same password is entered on the source server and destination server.
4. The source server and destination server are available on the same subnet:
   a. On the destination server, in a command prompt window, type `ipconfig` and note the subnet mask value.
   b. On the source server, in a command prompt window, type `ipconfig` and note the subnet mask value.
   c. Ensure that the subnet mask values are the same on the source server and destination server.
5. Port 7000 is open on the source and destination server, and they are not in use by another application.
   a. To check if port 7000 is open, in a Command Prompt window, enter the command:

      ```
      netsh firewall show portopening
      ```

      If port 7000 is not in the list, follow the instructions in File and Storage Services: Appendix A: Optional Procedures to open port 7000.
   b. If port 7000 is open, type the following command to check if port 7000 is being used by another application:

      ```
      netstat
      ```

- In the Local Address column, you will see <IP Address>:<port number>.
- If port 7000 is in the list, it is being used by another application.

Troubleshoot unexpected Windows PowerShell session closure

If a migration cmdlet fails, and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at
0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050). Minidump
attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Locate the deployment log file

The Windows Server Migration Tools deployment log file is located at %windir%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations:

- `%windir%\Logs\ServerMigration.log`
- On Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012: `%localappdata%\SvrMig\Log`
- On Windows Server 2003: `%userprofile%\Local Settings\Application Data\SvrMig\Log`

If migration log files are not created in the preceding locations, ServerMigration.log and SmigDeploy.log are created in `%temp%`, and other logs are created in `%windir%\System32`.

View the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects, and query them for more information about the settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets in Windows Server Migration Tools return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property Name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if the migration was successful; otherwise, **False** is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or shared folder that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property Name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.

Property Name	Type	Definition
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item shown as a path name.
Success	Boolean	The value **True** is displayed if the migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.

Property name	Type	Definition
Title	String	Title of the result object. Examples of values include the physical address of the network adapter for IP configuration, or the user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after the migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   $*VariableName* **=** $(*Cmdlet*)

   The following is an example.

   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -User all -Group -Path D:\rmt\DemoStore -force -Verbose)

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.

2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

   $*VariableName*

   In the following example, the variable is named **ImportResult**.

   $*ImportResult*

   This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable:

```
            ItemType  ID                                  Success
DetailsList
            --------  --                                  -------
-----------
          OSSetting  Local User                            True
{Local User, Loc...
          OSSetting  Local Group                           True
```

```
{Local Group, Lo...
     WindowsFeature  DHCP                                    True
{}
```

Each line of the preceding example is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by the examples that follow in steps 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   $<VariableName>| Select-Object -ExpandProperty <PropertyName>

   The following is an example.

   $importResult | Select-Object -ExpandProperty DetailsList

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples:

   - The following command returns only those details of result objects that have the ID **Local User**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Severity -eq "Warning") {$_} }

   - The following command returns only the details of result objects with an ID of **Local User** that also have the title **Remote Desktop Users**.

     $ImportResult | Where-Object { $_.ID -eq "Local Group" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Title -eq "Remote DesktopUsers") {$_} }

More information about querying results

For more information about the cmdlets that are used in the preceding examples, see the following additional resources.

- Where-Object (http://go.microsoft.com/fwlink/?LinkId=134853).
- Select-Object (http://go.microsoft.com/fwlink/?LinkId=134858).
- ForEach-Object (http://go.microsoft.com/fwlink/?LinkId=134860)

For more information about Windows PowerShell scripting techniques, see What Can I Do With Windows PowerShell? - Scripting Techniques on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134862).

See Also

Migrate File and Storage Services to Windows Server 2012

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Appendix A: Optional Procedures

Opening ports in Windows Firewall

The following instructions are for opening ports in Windows® Firewall. If you have a non-Microsoft firewall installed, consult the guide for that firewall about how to open ports. Opening ports in Windows Firewall can be done through the command interface.

Important
Opening ports in your firewall can leave your server exposed to malicious attacks. Make sure that you understand firewall systems before you open ports.

To open Windows Firewall ports by using the command line (do one of the following):

1. Open a Command Prompt window with elevated user rights, type the following, and then press ENTER.
   - On computers that are running Windows Server 2003, type:

     ```
     netsh firewall add allowedprogram
     program=%windir%\System32\WindowsPowerShell\v1.0\powershell.ex
     e name="ServerMigration" mode=ENABLE
     ```
   - On computers that are running Windows Server® 2008, Windows Server 2008 R2, or Windows Server 2012, type the following commands, in order, pressing ENTER after each command.

     i.

ii.

2. If you have changed the default behavior of Windows Firewall to block all outbound traffic on computers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, you must explicitly allow outbound traffic on UDP port 7000. To do this, open a Command Prompt window with elevated user rights, type the following, and then press ENTER.

```
netsh advfirewall firewall add rule name=ServerMigration(UDP-Out) dir=out program=%windir%\System32\WindowsPowerShell\v1.0\powershell.exe action=allow protocol=UDP localport=7000
```

Closing ports in Windows Firewall

As a best practice, we recommend that you close Windows Firewall ports after the data transfer operation is completed.

To close Windows Firewall ports by using the command line

- Do one of the following:
  - On computers that are running Windows® Server 2003, open a Command Prompt window, type the following, and then press ENTER.

    ```
    netsh firewall delete allowedprogram program=%windir%\System32\WindowsPowerShell\v1.0\powershell.exe
    ```
  - On computers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, open a command prompt window with elevated user rights, type the following two commands. Press ENTER after each.

    ```
    netsh advfirewall firewall delete rule name=ServerMigration(TCP-In)
    ```

    ```
    netsh advfirewall firewall delete rule name=ServerMigration(UDP-Out)
    ```

Detect reparse points and hard links

The following commands can be used to detect reparse points and mounted volumes in any folder and its subfolders. Open a Command Prompt window, type the following commands to detect reparse points, in which **D:\Test** represents the hard disk drive and folder that you want to search, and then press ENTER.

```
dir D:\Test\* /S /A:L
```

The option **/A:L** specifies that only reparse points need to be enumerated. The output is similar to the following:

```
Volume in drive D has no label.
Volume Serial Number is 3AE4-E412

Directory of D:\Test\Links

10/07/2008  03:44 PM    <JUNCTION>     JunctionMSIT [d:\test\targets\msit]
10/07/2008  03:42 PM    <SYMLINK>      LinkMSIT [d:\test\targets\msit]
10/07/2008  03:41 PM    <SYMLINKD>     SymLinkMSIT [d:\test\targets\msit]
1 File(s)              0 bytes

Directory of D:\Test\Targets

10/07/2008  05:35 PM    <JUNCTION>     Volume [\??\Volume{0674413f-760d-11dd-beb3-
806e6f6e6963}\]
0 File(s)              0 bytes

Total Files Listed:
1 File(s)              0 bytes
3 Dir(s)  17,918,840,832 bytes free
```

To enumerate hard links on a file on Windows Server 2008 R2, or Windows Server 2012 open a command prompt window with elevated user rights, type the following, and then press ENTER.

```
fsutil hardlink list D:\Test\File.txt
```

To enumerate hard links on all files in a folder on Windows Server 2008 R2 or Windows Server 2012, run the following command in a Windows PowerShell session that has been opened with elevated user rights:

```
Get-ChildItem D:\* | %{'Links for: ' + $_.FullName; fsutil hardlink list $_.FullName; ""}
```

For more information about enumerating hard links on computers that are running Windows Server 2003 or Windows Server 2008, see FindFirstFileNameW Function (http://go.microsoft.com/fwlink/?LinkId=147392) on MSDN.

Migrated and non-migrated attributes for local users and groups

For more information about the attributes of local users and groups that can be migrated, see the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?LinkID=128751) on the Microsoft Web site.

See Also

Migrate File and Storage Services to Windows Server 2012

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix B: Migration Data Collection Worksheets

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Appendix B: Migration Data Collection Worksheets

SMB data collection worksheet

Use this server message block (SMB) data collection worksheet to record data for SMB policies that are set on the source server.

#	Source Server Essential Settings	Setting Identification
01	Idle time The setting name is: **Microsoft network server: Amount of idle time required before suspending a session**.	Idle time (in minutes): _____ Group or Local Policy: _____
02	S4USelf The setting name is: **Microsoft network server: Attempt S4USelf to obtain claim information**.	Claim information: __ Default __ Enabled or __ Disabled Group or Local Policy: _____
03	Sign (always) The setting name is: **Microsoft network server: Digitally sign communications (always)**.	Sign always: __ Enabled or __ Disabled Group or Local Policy: _____
04	Sign (if client agrees) The setting name is: **Microsoft network server: Digitally sign communications (if client**	Sign if client agrees: __ Enabled or __ Disabled Group or Local Policy:

#	Source Server Essential Settings	Setting Identification
	agrees).	_____
05	Disconnect when logon hours expire The setting name is: **Microsoft network server: Disconnect clients when logon hours expire**.	Disconnect: __ Enabled or __ Disabled Group or Local Policy: _____

BranchCache data collection worksheet

Use this BranchCache data collection worksheet to record data for the BranchCache policies that are set on the source server.

#	Source Server Essential Settings	Setting Identification
01	BranchCache The setting name is: **Hash Publication for BranchCache**.	BranchCache: __ Not configured, __ Enabled, or __ Disabled Group or Local Policy: _____
	BranchCache The setting name is: **Hash Version support for BranchCache**.	BranchCache: __ Not configured, __ Enabled, or __ Disabled Group or Local Policy: _____

See Also

Migrate File and Storage Services to Windows Server 2012

File and Storage Services: Prepare to Migrate

File and Storage Services: Migrate the File and Storage Services Role

File and Storage Services: Verify the Migration

File and Storage Services: Post-Migration Tasks

File and Storage Services: Appendix A: Optional Procedures

File and Storage Services: Appendix C: Migrate iSCSI Software Target

File and Storage Services: Appendix C: Migrate iSCSI Software Target

This document describes how to migrate Microsoft® iSCSI Software Target 3.2 or 3.3 settings and data from an existing Windows Storage Server 2008 or Windows Storage Server 2008 R2 computer to a destination server that is running the ISCSI Target Server role service that is included with Windows Server® 2012 and Windows Storage Server® 2012.

The naming for iSCSI Software Target has changed. To reduce the potential for confusion, in the context of this document, any naming that refers to “iSCSI Software Target”, refers to prior product versions installed on Windows Storage Server 2008 and Windows Storage Server 2008 R2, which are source servers. By contrast, any naming that refers to “iSCSI Target Server” refers to the new role service included with Windows Server 2012 and Windows Storage Server 2012, which are destination servers.

Note

This document only contains iSCSI-specific migration information. For generic information, such as the use of Windows Server Migration Tools, refer to the application section in the main File and Storage Services Migration Guide.

See Also

iSCSI SoftwareTarget Migration Overview

Prepare to Migrate iSCSI Software Target

Migrate iSCSI Software Target

Verify the iSCSI Software Target Migration

Troubleshoot the iSCSI Software Target Migration

Roll Back a Failed iSCI Software Target Migration

iSCSI SoftwareTarget Migration Overview

Insert introduction here.

Migration overview

This section describes the high-level migration process, which involves harvesting configuration settings from the source, moving the virtual disks from the source server to the destination server, and restoring the configuration settings.

Migration process

This section describes the high-level migration process.

Migration planning

The migration planning phase involves gathering the following information:

- Are the source server and destination server are configured in a cluster?
- If the servers are configured in a cluster, what are the virtual computer objects or client access points that contain the iSCSI target resources?
- Is the storage system of the destination server capable and configured appropriately to host the virtual disks of the source server, and does it have appropriate space to store the volume snapshots?
- Are there any iSCSI initiators that have a critical dependency on iSCSI targets for the duration of the migration process (such as a computer that uses iSCSI boot nodes, or clusters that use shared storage)?
- Are there any IP address or portal settings that are unique to the source server that need to be accounted for (such as IP addresses that are known to the firmware of devices)?
- Are there any iSNS settings that need to be manually recorded and migrated?
- Are there any virtual disks surfaced as local disks that might need to be exposed?

Preparing to migrate

The preparation to migrate data from the source server to the destination server involves the following steps:

1. If the destination server will have a clustered configuration, install the Failover Clustering feature and form a cluster before performing the migration.
2. If the destination server will have a clustered configuration, create a number of cluster resource groups with client access points and cluster disks as appropriate to replicate the existing configuration. If possible, use the same resource group names for the source clusters and the destination clusters.
3. Install the iSCSI Target Server role service on the destination server.
4. Disconnect all the iSCSI initiators. This step is required to maintain consistent data on the virtual disks while they are being moved.
5. Run the Windows® PowerShell™ script, iSCSITargetSettings.ps1, to capture the existing settings on the source server in an XML file. For a cluster, run the script on each node in the cluster or on each virtual computer object, as appropriate for the scope of the planned migration.

   The Windows PowerShell script displays the virtual disks that are eligible for migration and those that are not (for the snapshot-based reasons discussed previously).

Migration

The actual migration process consists in the following steps:

1. Move the files for all the virtual disks that are eligible for migration from the source server to the destination server. If there are any file path changes, note the source to destination mapping.

2. In a cluster configuration, ensure that the destination path of the file copy is on a cluster disk and that the cluster disk has been assigned to a resource group. Note of the resource group that owns the path.
3. If the file paths have changed between the source and the destination servers, open the settings .xml file in a text editor, and identify the <MigrationDevicePath> tags that need to be changed to reflect the new path.
4. In a cluster configuration, if the file path or the resource group name have changed between the source server and the destination server, open the settings .xml file in a text editor, and identify the <MigrationResourceGroup> tags that need to be changed to reflect the new resource group.
5. Run the Windows PowerShell script, iSCSITargetSettings.ps1, to import the settings to the destination server. In a cluster configuration, the destination server can be specified as a cluster node or as a virtual computer object. The cluster node or virtual computer object must be the owner of the resource group that is indicated in the settings .xml file.
6. If there are snapshot storage settings relevant to the new configuration, apply those settings manually.
7. If there are virtual disks that need to be surfaced as local disks, perform those actions.
8. If there are any iSNS settings that are relevant to the new configuration, apply those settings manually.
9. If there are any iSCSI target portal settings that are relevant to the new configuration, apply those settings manually.
10. If there are any iSCSI initiators that are configured to authenticate by using CHAP and Reverse CHAP, manually restore those settings.

Verification

The verification process for the migration involves the following steps:

1. Validate the iSCSI target portal settings by opening a Command Prompt window and typing **netstat.exe –nao | findstr 3260**. (This assumes that the default TCP port for the iSCSI protocol 3260 is used). Alternatively, type **Get-WmiObject –Namespace root\wmi –Class WT_Portal** to cross-check the results.
2. Inspect the iSCSI Target Server configuration by using the Windows PowerShell cmdlet, **Get-IScsiServerTarget**
3. Inspect the iSCSI virtual disk configuration by using the Windows PowerShell cmdlet, **Get-IScsiVirtualDisk**
4. Validate the configuration for each iSCSI initiator that you expect to use with iSCSI Target Server by using the iscsicpl.exe UI tool or the iscsicli.exe command line tool.

Impact of migration

The migration process does not impact or affect the source server. There are no resources or configuration settings that are altered or deleted as part of the migration process.

No servers in the enterprise, other than the destination servers, will be affected by the migration.

Client computers that are running as iSCSI initiators are expected to be explicitly disconnected during the migration to ensure data integrity. During the migration, the source server will be

unavailable. When the migration process is complete, it is expected that the iSCSI initiators will log on to the destination server without any issues.

The downtime for the iSCSI initiators is expected to be proportionate to the time it takes to move the virtual disk files from the source server to the destination server, plus the time needed to restore the configuration settings and to establish the network identity of the destination server.

Permissions required for migration

Local Administrator permissions are required on the source and the destination server.

If the iSNS server has additional access control policies, permission to alter the iSNS settings are required as appropriate for the iSNS server.

To perform the migration process for the iSCSI initiators, permissions to log on and log off iSCSI sessions are required. For the iSCSI initiator, Local Administrator permissions are required.

For iSCSI initiators that are firmware based, such as a network interface with the option to boot from iSCSI, being at the actual console may be required to configure logon credentials or the network identity of the destination server if the authentication settings (CHAP and Reverse CHAP) have changed.

Estimated time duration

This section detail the various factors that impact how long a migration may take to complete.

Planning

The planning phase is expected to be influenced by the following factors:

- Standalone versus a cluster configuration. A cluster setup may require one to two hours to configure if all the validations are performed.
- Storage configuration. Understanding and configuring a storage array to host potentially huge files requires that you plan the spindle and volume configurations so that they use the tools that are provided by the storage subsystem vendor.
- Network identity. This planning involves understanding if the source server has specially or purposely configured IP addresses, if configuring Level-2 components (such as switches) is required, and if specific DNS or NetBIOS names need to be known to and cached by the iSCSI initiators.

Preparation

The preparation process involves understanding which settings (that are specific to the source server) cannot be automatically migrated, and gathering those settings. For each step in the preparation phase, the mechanism that is used to retrieve the settings depends on which step is applicable and which tool is used to recover those settings.

- Cluster resource group names and configuration. These settings can be gathered from the cluster administration tools and the user interfaces.
- iSCSI target portal configuration. These settings can be gathered by typing the following code at a command prompt: `PS > Get-WmiObject –Namespace root\wmi –Class WT_Portal`

- iSNS Server settings. These settings can be gathered by typing the following code at a command prompt: `PS > Get-WmiObject –Namespace root\wmi –Class WT_ISnsServer`
- CHAP and Reverse CHAP authentication settings. These settings cannot be automatically retrieved because the iSCSI target server does not offer a mechanism to retrieve passwords. These settings have been stored elsewhere in the enterprise, and they need to be retrieved independently.
- Locally mounted virtual disk settings.

Migration

The estimated time for the actual migration process is largely dominated by the time that it takes to move the virtual disk files from the source server to the destination server.

A network-based file copy, using a 1 GB link used at 50% for 1 TB of data, is estimated to take over five hours. Techniques that use a file transfer process involving external media, such as an External Serial Advanced Technology Attachment (eSATA) device, may take less time.

The execution of the Windows PowerShell import script is estimated to take few minutes for approximately 100 resources (with a combination of iSCSI target settings and virtual disk settings).

Verification

The estimated time for the verification is proportionate to the time it takes to reconnect or log on to the iSCSI initiators.

For each iSCSI initiator, the target portal needs to be reconfigured, credentials related to authentication settings must be entered (if required), and the sessions have to be logged on.

The estimated time is 5 to 15 minutes to verify each iSCSI initiator, depending on the process that is being used. iSCSI initiators can be verified through the iscsicpl.exe UI, through the iscsicli.exe command line tool, or through ad hoc Windows Management Instrumentation (WMI)-based scripts).

Supported migration scenarios

This section details both supported and unsupported migration scenarios.

Supported operating systems

The versions of operating systems that are listed are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

Migrations between physical operating systems and virtual operating systems are supported.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shared resources from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Storage Server 2008, full installation options	Windows Server 2012 and Windows Storage Server 2012	x64-based
x64-based	Windows Storage Server 2008 R2	Windows Server 2012 and Windows Storage Server 2012	x64-based
x64-based	Windows Storage Server 2008 R2	Windows Server 2012 and Windows Storage Server 2012	x64-based

x64-based migrations are supported for Windows Storage Server 2012 and Windows Server 2012. All editions of Windows Storage Server 2008 R2 and Windows Server 2008 R2 are x64-based.

x86-based migrations are not supported because Windows Storage Server 2012 is not offered in the x86 platform.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Supported role configurations

This migration guide is applicable to standalone and clustered configurations, with certain limitations.

The following general restrictions are applicable to all the supported configurations:

- Authentication settings for iSCSI initiators that use CHAP and Reverse CHAP settings are not automatically migrated.
- Snapshot storage settings for each virtual disk in the configuration are not automatically migrated.
- Configuration settings for virtual disks that are derived from snapshots are not automatically migrated.
- For clustered configurations, the migration process includes iSCSI target settings that are scoped to the virtual computer bject, to a cluster node, or to the cluster node that owns the code cluster group.
- For clustered configurations, the migration of resource groups, network name resources, IP addresses, and cluster disks that are associated with resource groups is outside of scope for this guide, and it needs to be performed independently as a preliminary step.
- iSCSI Naming Services (iSNS) settings for ISCSI Software Target are not automatically migrated.

- iSCSI target portal settings (such as IP addresses that are used by the iSCSI target service to listen for incoming network connections) are not automatically migrated
- The schedule for snapshots of virtual disks is not migrated.

The following configurations are supported:

- Migration from a standalone configuration to standalone configuration
- Migration from a clustered configuration to a standalone configuration (with the restrictions listed previously regarding the scope of the settings).
- Migration from a clustered configuration to a clustered configuration (with the restrictions listed previously regarding the scope of the settings).

Supported role services and features

ISCSI Target Server (as included with Windows Storage Server 2012 and Windows Server 2012) does not have role dependencies or feature dependencies.

It is possible to install ISCSI Target Server with failover clustering, and this configuration is supported with the migration limitations listed previously.

Migrating multiple roles

If you are migrating one clustered configuration to a different clustered configuration, the Failover Clustering feature needs to be migrated or set up prior to migrating iSCSI target settings.

Migration scenarios that are not supported

The following migration scenarios are not supported:

- Migration from Windows Unified Storage Server 2003 R2.
- Migration from a standalone configuration to a clustered configuration. This migration is not supported because there is no default mechanism to associate target and virtual disk settings to resource groups without knowing how the file paths are mapped to the cluster disk and how IP Addresses are mapped to resource groups.
- Snapshots of virtual disks are not automatically migrated. Snapshots are based on a snapshot of the volume that contains the virtual hard disk (VHD) file at the time the snapshot was taken. Their existence and implementation depends on the volume of the computer from which the migration process happens, and it cannot be replicated or exported.
- Snapshot storage settings for virtual disks are not automatically migrated. The snapshot storage settings (such as volume and maximum size per volume) are dependent on the hardware and software configuration of the computer that the settings are being migrate to, and they cannot automatically be migrated. For detailed information about how to manually migrate the snapshot storage settings, see Migrating ISCSI Target.
- The configuration settings of the iSCSI target portal are not automatically migrated. This configuration is based on the IP addresses of the destination server, and those settings cannot be migrated outside the knowledge of the network configuration of the computer that the settings are being migrate to. For detailed information about how to manually configure the portal settings, see Migrating ISCSI Target.

- iSNS settings are not automatically migrated. The iSNS settings are based on the network infrastructure and configuration of the destination server, and those settings cannot be migrated outside the knowledge of the network configuration of the computer that the settings are being migrated to. For detailed information about how to manually configure iSNS settings, see Migrating ISCSI Target.
- Settings for virtual disks that are surfaced as local disks on the source server are not automatically migrated. The ability to surface a disk locally is expected to be a temporary operation that can be replicated if. For detailed information about how to configure settings for virtual disks that are to be surfaced as local disks, see Migrating ISCSI Target.
- The schedule for snapshots of virtual disks is not migrated. Those settings must be manually discovered and replicated from the source to the destination server.

Prepare to Migrate iSCSI Software Target

This topic discusses the tasks that are necessary before you start the migration process. The first step is to install the Windows Server Migration Tools. For more information, see File and Storage Services: Prepare to Migrate.

Prepare the destination server

The destination server is a computer that is configured and shipped by an OEM with Windows Storage Server 2012 pre-installed, or that is running Windows Server 2012.

ISCSI Target Server hardware requirements for the destination server are as follows:

- The amount of free disk space on the destination server must be sufficient to host the iSCSI virtual disk from the source server with adequate room for the snapshot storage.
- For clustered configurations, the resource groups that are created in the destination server must have associated cluster disks with adequate free space to host the iSCSI virtual disk from the source server.
- The destination server must have one or more network interfaces to be utilized for the iSCSI network traffic.

Installing the Failover Cluster feature in Windows Storage Server 2012 or Windows Server 2012 is required if the source server was configured with failover clusters. For more information, see the Failover Clustering Failover Clustering (http://technet.microsoft.com/en-us/library/hh831579.aspx).

Backup the source server

Before you start migration, as a best practice, it is recommended that you back up the source server. For more information, see Windows Server Backup (http://technet.microsoft.com/en-us/library/cc770757.aspx).

Prepare the source server

The following are tasks that are performed on the source server.

Cluster resource group configuration

Use the following steps to obtain the cluster resource groups:

1. Gather the resource groups that have iSCSI Software Target resources by using the following Windows PowerShell command:

```
PS > Import-Module FailoverClusters
PS > $iSCSITargetResources = Get-ClusterResource | Where-Object
{ ( $_.ResourceType.Name -eq "Host" ) -or ($_.ResourceType.Name
-eq "WTDisk") }
PS > $iSCSITargetResources
```

2. From the cluster resources obtained in the previous step, gather the cluster disk dependencies by using the following Windows PowerShell command:

```
PS > $Dependencies = &{ $iSCSITargetResources | Get-
ClusterResourceDependency  }
PS > $Dependencies
```

If the source server is running Windows Storage Server 2008, the following steps can be followed to gather the equivalent information.

1. Gather the iSCSI Software Target resources, and then gather the groups by using the following Windows PowerShell command:

```
PS > $iSCSITargetResources = Get-WmiObject -NameSpace
root\mscluster -Authentication PacketPrivacy -Class
MsCluster_Resource -Filter "Type = `"WTDisk`" or Type =
`"Host`""
PS > $iSCSITargetResources
PS > $Groups = &{foreach($res in $iSCSITargetResources) { Get-
WmiObject -NameSpace root\mscluster -Authentication
PacketPrivacy -Query "associators of {$($res.__RELPATH)}  WHERE
ResultClass = MSCluster_ResourceGroup" }}
PS > $Groups
```

2. From the cluster resources obtained in the previous steps, gather the cluster disk dependencies by using the following Windows PowerShell command:

```
PS > $Dependencies = &{foreach($res in $iSCSITargetResources) {
Get-WmiObject -NameSpace root\mscluster -Authentication
PacketPrivacy -Query "associators of {$($res.__RELPATH)}  WHERE
ResultClass = MSCluster_Resource ResultRole = Dependent" }}
PS > $Dependencies
```

The resource groups obtained in step 1 have network name resources and IP addresses that need to be migrated to the destination server.

For information about how to migrate these settings, see Migrate Roles and Features to Windows Server 2012 (http://technet.microsoft.com/en-us/library/jj134039.aspx).

The cluster disk that you obtained in step 2 is the physical disk where the volumes that are hosting the iSCSI Software Target virtual disks reside.

To obtain the volumes from the cluster disk, use the following steps:

1. Obtain the disk signature of the cluster disk by using the following Windows PowerShell command:

   ```
   PS > & cluster.exe res "<cluster resource name>" /priv
   ```

2. Obtain the Win32_DiskDrive object from the disk signature by using the following Windows PowerShell command:

   ```
   PS > $DiskObj = Get-WmiObject -Namespace root\cimv2 -Class
   Win32_DiskDrive -Filter "Signature = <disk signature>"
   PS > $DiskObj
   ```

3. Obtain the Win32_DiskDriveToDiskPartition association by using the following Windows PowerShell command:

   ```
   PS > $DiskToDiskPartition = Get-WmiObject -Namespace root\cimv2
   -Class Win32_DiskDriveToDiskPartition | Where-Object {
   $_.Antecedent -eq $DiskObj.__PATH }
   PS > $DiskToDiskPartition
   ```

4. Obtain the Win32_LogicalDiskToDiskPartition association that points to the volume association by using the following Windows PowerShell command:

   ```
   PS > Get-WmiObject -Namespace root\cimv2 -Class
   Win32_LogicalDiskToPartition | Where-Object { $_.Antecedent -eq
   $ DiskToDiskPartition.Dependent }
   ```

Steps 2 through 4 need to be applied on the source server cluster node that currently owns the physical disk cluster resource.

iSCSI Target portal configuration

Use the following steps to obtain the portal associations:

1. Gather the configured portals association for the iSCSI target portal by using the following Windows PowerShell command:

   ```
   PS> Get-WmiObject -Namespace root\wmi -Class WT_portal | Format-List -Property Address,Listen,Port
   ```

2. The IP addresses that have the Listen state set to True are the IP addresses that an iSCSI initiator can use to reach the server. For more information about migrating the IP addresses, see Migrate Roles and Features to Windows Server 2012 (http://technet.microsoft.com/en-us/library/jj134039.aspx).

iSNS configuration

Gather the configured iSCSI Naming Services (iSNS) association for the server by using the following Windows PowerShell command:

```
PS> Get-WmiObject -Namespace root\wmi -Class WT_ISnsServer | Format-List -Property ServerName
```

The server names that are listed need to be added to the list of iSNS servers that can be used to retrieve information about the iSCSI initiators in the enterprise.

CHAP and Reverse CHAP configuration

Gather the UserName and ReverseCHAPUserName association for the servers that are configured with CHAP and Reverse CHAP by using the following Windows PowerShell command:

```
PS > Get-WmiObject -Namespace root\wmi -Class WT_Host | Where-Object { ( $_.EnableCHAP ) -or ( $_.EnableReverseCHAP )  } | Format-List -Property Hostname,CHAPUserName,ReverseCHAPUserName
```

The passwords that are used in conjunction with the credentials listed previously cannot be retrieved, and they must be known through other mechanisms.

Snapshot storage configuration

The snapshot storage configuration can be obtained by using the following Windows PowerShell command:

```
PS > & vssadmin.exe list shadowstorage
```

This command shows the volume snapshot shadow storage configuration for the entire source server. Not all the volumes listed may be relevant to the current iSCSI Software Target server configuration.

For the volumes that are relevant (that is, the volumes that host iSCSI virtual disks), the associated shadow storage volume is listed, in addition to the amount of disk space used with the maximum amount of configured space.

Disconnect the iSCSI initiators

Follow the instruction in the following section to disconnect the iSCSI initiators: Prepare other computers in the enterprise.

Capture the existing settings: standalone configuration

All of the settings on the iSCSI Software Target source server that are not hardware configuration specific and are not dependent on an IP address and the network identity of the server can be captured with the following Windows PowerShell command:

Windows Server 2008 and Windows Server 2008 R2 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> .\ iSCSITargetSettings.PS1 –Export –FileName <settings XML file>
```

Windows Server 2012 file path:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerShell\V1.0\Modules\IscsiTarget"

PS> .\ iSCSITargetSettings.PS1 –Export –FileName <settings XML file>
```

If the procedure is performed on a source server that is running ISCSI Target 3.3 from a destination server that is prepared as illustrated in the previous sections, the settings can be captured using the following Windows PowerShell command:

Windows Server 2012 file path:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget"

PS> .\ iSCSITargetSettings.PS1 –Export –FileName <settings XML file> -ComputerName
<source server computer name>
```

Windows Server 2008 and Windows Server 2008 R2 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> .\ iSCSITargetSettings.PS1 –Export –FileName <settings XML file> -ComputerName
<source server computer name>
```

At the end of the settings capture process, the Windows PowerShell script will display the set of VHD files that are eligible for migration. This list is needed for the destination server during migration.

Capture the existing settings: clustered configuration

Before capturing the iSCSI Software Target source server settings that are not hardware configuration specific, we recommend that all the resource groups with iSCSI target resources are moved to a single node in the cluster.

This can be accomplished by using the following Windows PowerShell commands. These commands assume that you previously followed the steps in the following section: Cluster resource group configuration.

```
PS > $iSCSITargetResources | Format-List -Property OwnerGroup

PS > foreach($Res in $iSCSITargetResources) { & cluster group  $Res.OwnerGroup
/moveto:$ENV:COMPUTERNAME }
```

After all the resource groups have been moved to a single node, the settings can be gathered by using the following Windows PowerShell commands:

Windows Server 2012 file path:

```
PS > cd "$ENV:Programfiles\ISCSI Target"

PS> .\ iSCSITargetSettings.PS1 -Export -FileName <settings XML file>
```

Windows Server 2008 and Windows Server 2008 R2 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"

PS> .\ iSCSITargetSettings.PS1 --Export -FileName <settings XML file>
```

If the procedure is performed on a source server that is running ISCSI Target 3.2, the resources can be moved to a single node by using the following Windows PowerShell command:

```
PS >  $Groups = &{foreach($res in $iSCSITargetResources) { Get-WmiObject -NameSpace
root\mscluster -Authentication PacketPrivacy -Query "associators of {$($res.__RELPATH)}
WHERE ResultClass = MSCluster_ResourceGroup" }}

PS > foreach($Group in $Groups) { & cluster group $Group.Name /moveto:<node name source
server> }
```

The ISCSI Target Server settings need to be gathered from a destination server that is prepared as illustrated in the previous sections. Run the script for a source server that is running ISCSI Target 3.3 by using the following Windows PowerShell command:

Windows Server 2012 file path:

```
PS > cd "$ENV:Programfiles\ISCSI Target"

PS> .\ iSCSITargetSettings.PS1 -Export -FileName <settings XML file> -ComputerName
<source server computer name>
```

Windows Server 2008 and Windows Server 2008 R2 file path

```
PS > cd "$ENV:SystemRoot\Program Files\Microsoft iSCSI Software Target"
```

```
PS> .\ iSCSITargetSettings.PS1 –Export –FileName <settings XML file> -ComputerName
<source server computer name>
```

In this command, the source server computer name is the name of the node. At the end of the settings capture process, the Windows PowerShell script will display the set of VHD files that are eligible for migration. This list is needed for the destination server during migration.

Remove the network identity of the iSCSI Software Target computer

In a network with an iSCSI Software Target source computer, the identity of the server is known to iSCSI initiators in the form of NetBIOS names, fully qualified domain names (FQDN), or IP addresses. When a server is being replaced, as part of planning, a strategy to replace the server network identity must be devised. Possible scenarios include:

- Transfer the NetBIOS and fully qualified domain names to the destination server, and then assign new IP addresses to the destination server.
- Create new NetBIOS and fully qualified domain names for the destination server, and then assign the existing IP addresses to the destination server.
- Create new NetBIOS and fully qualified domain names for the destination server, and then assign new IP addresses to the destination server.

Each scenario requires potentially updating information in the DNS server, Active Directory, or DHCP server, according to the methodology that is used to assign IP addresses and network names to the servers in the enterprise.

The intent of this step is to ensure that upon completion of the migration steps, the iSCSI initiators are able to locate the destination server (either through explicit reconfiguration, or implicitly through the computer name or IP address re-assignment).

For more information, see Migrate Roles and Features to Windows Server 2012 (http://technet.microsoft.com/en-us/library/jj134039.aspx).

Prepare the iSCSI initiator computers

The other computers in the enterprise that are affected by migration are the iSCSI initiators. The users of the computers that are acting as iSCSI initiators should be sent an outage notification. If the iSCSI Software Target is being used as a boot node for the iSCSI initiator computers, the computers may be completely unusable for the duration of the migration.

Capture the session information

The information regarding the active session for an iSCSI Software Target source server can be obtained by using the following Windows PowerShell command:

```
PS > & iscsicli.exe sessionlist
```

This information is needed to disconnect the session in the following step.

Disconnect the session

The session can be disconnected by using the following Windows PowerShell command:

```
PS > & iscsicli.exe LogoutTarget <session id>
```

Migrate iSCSI Software Target

This topic discusses the actual migration steps for ISCSI Software Target 3.2 or iSCSI Software Target 3.3 for both the standalone configuration and the clustered configuration:

Migrating ISCSI Software Target in a standalone configuration

The migration of ISCSI Software Target 3.2 or iSCSI Software Target 3.3 has equivalent steps, whether you are migrating from Windows Storage Server 2008 or Windows Storage Server 2008 R2 to Windows Server 2012 or Windows Storage Server 2012.

Establish network identity of the iSCSI Target Server computer

As part of the planning process, a strategy should have been devised regarding how iSCSI Target Server will be accessed from the network, including but not limited to:

- Which computer name will be used?
- Which IP addresses on which subnet or set of network interfaces will be used?
- What relationship should be maintained between the IP addresses and computer name of the source server and the destination server—will you keep the same addresses and names or create new ones?

Based on the desired final configuration, configuration changes are potentially needed in the following areas:

- The DHCP Server that might assign IP addresses to the destination iSCSI Target servers
- The DHCP Server that might assign IP addresses to the iSCSI initiators
- The DNS Server or Active Directory domain controller that might perform naming resolution services for the computers in the enterprise

Configure the iSCSI Target Server portal

After you have configured IP addresses for the network interfaces of the iSCSI Target Server computer, it is possible to verify the existing configuration by using the following Windows PowerShell command:

```
PS > $Portals = Get-WmiObject -Namespace root\wmi -Class WT_Portal | Where-Object {
$_.Listen }

PS > $Portals
```

The configuration of the access surface for iSCSI Target Server from the network can be restricted by disabling certain portals. For example, you can disable the fourth portal in the array returned in the previous step by using the following Windows PowerShell commands:

```
PS > $Portals[3].Listen = 0

PS > $Portals[3].Put()
```

The default port can also be changed from 3260 to any available TCP port on the destination server.

Configure iSNS settings

The iSNS servers that were configured for the source server can be configured for the destination server by using the following Windows PowerShell commands:

```
PS > $WT_ISnsServerClass =  Get-WmiObject -namespace root\wmi -class meta_class -filter
"__CLASS = 'WT_ISnsServer'"

PS > $WtiSNSInstanace = $WT_ISnsServerClass.CreateInstance()

PS > $WtiSNSInstanace.ServerName = "<iSNS computer name or IP>"

PS > $WtIsnsInstanace.Put()
```

Note
The set of iSNS servers that are configured for iSCSI Target Server was obtained during the preparation of the source server.

Configure storage

The destination server is expected to have sufficient storage space to host all of the virtual disks that are present on the source server.

The space does not need to be contiguous or in a single volume, and it does not need to replicate the same file system structure or volume mount point structure of the source server.The storage that is prepared to host the virtual disks must not be a nested volume, and it must be formatted with the NTFS file system.

Configure the Volume Shadow Copy Service

For the storage that was prepared in the previous step, it is appropriate to configure the Volume Shadow Copy Service, in case the default per-volume settings are not adequate.To inspect that current configuration, use the following Windows PowerShell command:

```
PS > & vssadmin.exe list shadowstorage
```

To modify the current configuration, use the following Windows PowerShell commands:

```
PS > & vssadmin.exe add ShadowStorage

PS > & vssadmin.exe delete ShadowStorage

PS > & vssadmin.exe resize ShadowStorage
```

Transfer the virtual disk

For all the files in the list of files that was captured in the source server preparation step, copy the files from the source server to the destination server. For more information, see Capture the existing settings section.

You will need the destination paths in the following steps. So if the absolute file path is different between the source server and the destination server, create a table with the mapping, for example

Source path	Destination path
G:\WS08R2_OpsMgr2007_R2.vhd	H:\VHDS\WS08R2_OpsMgr2007_R2.vhd
F:\Dynamic_Spanned_GPT_2.vhd	D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd

Import the iSCSI Software Target settings in a standalone configuration

To import the iSCSI Software Target settings in a standalone configuration, you need the .xml file that you previously created. For more information, see Capture the existing settings section.

If there is no change in the absolute path of the virtual disk files, the import process can be performed by using the following Windows PowerShell commands:

```
PS > cd "$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget"

PS> .\ iSCSITargetSettings.PS1 -Import -FileName <settings XML file>
```

If the absolute path is different between the source server and the destination server, before you import the settings, the settings .xml file needs to be altered to reflect the new path.

Locate the records for the virtual disk, and alter the path in the <MigrationDevicePath> tag to reflect the absolute file path in the destination server, for example:

```
    <iSCSIVirtualDisk>

     <DevicePath>F:\Dynamic_Spanned_GPT_2.vhd</DevicePath>

      <MigrationDevicePath>D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd</MigrationDevicePath>

    </iSCSIVirtualDisk>
```

After the XML has been altered to reflect the path in the destination server, you can import the settings by using the Windows PowerShell commands previously presented.

Configure shadow storage for the virtual disks

If certain virtual disks have shadow storage requirements that are different than the ones configured in the section Configure the Volume Shadow Copy Service , it is possible to alter the default or previously configured settings by using the following Windows PowerShell commands:

```
PS > $VirtDisk = Get-WmiObject -Namespace root\wmi -Class WT_Disk | Where-Object {
$_.DevicePath -eq '<full path of virtual disk>' }

PS > $VirtDisk.SnapshotStorageSizeInMB = <new size>

PS > $VirtDisk.Put()
```

Configure CHAP and Reverse CHAP

The authentication settings for iSCSI Target Server that are configured with CHAP and Reverse CHAP need to be manually configured. The list of targets that require CHAP and Reverse CHAP configuration is listed at the end of the import script, as described in the section Import the iSCSI Software Target settings in a standalone configuration.

To configure the CHAP and Reverse CHAP settings, use the following Windows PowerShell commands:

```
PS > $Target = Get-WmiObject -Namespace root\wmi -Class WT_Host | Where-Object {
$_.HostName -eq '<name of the target>' }

PS > $Target.EnableCHAP = 1

PS > $Target.CHAPUserName = "<user name>"

PS > $Target.CHAPSecret = "<CHAP Secret>"

PS  $Target.Put()
```

Migrating iSCSI Software Target in a failover cluster

The migration process for the failover cluster configuration is largely identical to migrating a standalone configuration, with the following differences:

- You will migrate resource groups instead of merely establishing the network identity of the server.
- You will use different Windows PowerShell commands to import the resource groups.

Migrate resource groups

This step replaces the “Establishing the network identity of iSCSI Target Server” step when you migrate a standalone configuration. The reason is that the network identity of an iSCSI Target server in a cluster is given by the union of the client access point. (A client access point in the cluster is the logical union of a network name resource and one or more IP addresses that are assigned to the network name resource.)

Assuming the initial cluster resource groups and network names were configured in the default state, those can be recreated by using the following Windows PowerShell command:

```
PS > Add-ClusterServerRole – Name <resource group name>
```

Use this command for each of the resource groups that were in the original configuration.If the default client access point configuration does not match the initial configuration (for example, because the network name is bound to the incorrect cluster network, or the configuration required statically assigned IP addresses), modifications can be made. For more information, see Migrate Roles and Features to Windows Server 2012 (http://technet.microsoft.com/en-us/library/jj134039.aspx).

After the resource groups have been created, clustered disks must be assigned to the network resources to match the configuration that you captured. For more information, see the Cluster resource group configuration section.

Import the iSCSI Software Target settings in a failover cluster

Follow these instructions to import settings in a failover cluster configuration. (This information differs from the how you would import settings in a standalone configuration.)

A prerequisite for the import phase is to have all of the resource groups that will host iSCSI Target Server resources owned by the same cluster node. Use the following Windows PowerShell command to validate the current ownership:

```
PS > Get-ClusterGroup
```

If there is no change in the absolute path of the virtual disk files, the import process can be performed by using the following commands:

```
PS > cd “$ENV:SystemRoot\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget”

PS> .\iSCSITargetSettings.PS1 –Import –FileName <settings XML file>
```

If the absolute path is different between the source server and the destination server, before you import the settings, the settings .xml file needs to be altered to reflect the new path.Locate the records for the virtual disk, and alter the path in the <MigrationDevicePath> tag to reflect the absolute file path in the destination server, for example:

```
<iSCSIVirtualDisk>
```

```
      <DevicePath>F:\Dynamic_Spanned_GPT_2.vhd</DevicePath>
      <MigrationDevicePath>D:\DYNVHDS\Dynamic_Spanned_GPT_2.vhd</MigrationDevicePath>
    </iSCSIVirtualDisk>
```

After the XML has been altered to reflect the path in the destination server, you can import the settings by using the Windows PowerShell commands previously presented.

Verify the iSCSI Software Target Migration

This topic discusses the steps that you can use to verify that the migration successfully completed.

Verifying the destination server configuration

To verify that the destination server has been properly configured after migration, you can verify the listening endpoints and connectivity and run a scan with the Best Practices Analyzer.

Verify the listening endpoints

On the iSCSI Target destination server, you can validate that the target portals have been configured as desired by using the following Windows PowerShell command:

```
PS > & netstat.exe -nao | findstr 3260 | findstr LISTENING
  TCP    10.121.26.107:3260     0.0.0.0:0              LISTENING       1560
  TCP    10.121.26.126:3260     0.0.0.0:0              LISTENING       1560
  TCP    [2001:4898:0:fff:0:5efe:10.121.26.126]:3260  [::]:0                 LISTENING
1560
  TCP    [2001:4898:f0:1001:f063:8fc5:52e6:2310]:3260  [::]:0                 LISTENING
1560
```

The list of IP addresses and port pairs in the listening state needs to match the desired set of target portals.

Note

If ports other than the default 3260 are being used, the command needs to be altered to reflect the alternate IP ports.

Verify the basic connectivity

To validate that the iSCSI Target Server computer is reachable from other computers on the network, from a computer that has the Telnet Client feature installed, use the following Windows PowerShell command:

```
PS > telnet.exe <iSCSI Software Target machine name or IP> 3260
```

If there is a successful connection, Telnet Client will show a blinking cursor at the top of the window. Press any key to close Telnet Client.

Perform a Best Practices Analyzer scan

To verify that ISCSI Target Server is optimally configured on Windows Server 2012 or Windows Storage Server 2012 after migration, we recommend that you run a Best Practices Analyzer (BPA) scan on the role.

BPA is a server management tool that is available in Windows Server 2012. After the migration of ISCSI Target 3.3 is complete, BPA can help you ensure that your server is configured according to best practices. You can use the Server Manager console UI or Windows PowerShell to perform BPA scans and view results. For detailed information about how to scan your role and view results, see the Best Practices Analyzer Help (http://go.microsoft.com/fwlink/?LinkId=122786).

Verifying the configuration of iSCSI initiator computers

Validating the migration of ISCSI Software Target to the destination server includes ensuring that the iSCSI initiators can discover and fully access all features of the iSCSI protocol.

Verify that the iSCSI initiators can discover iSCSI Target Server

To verify that the iSCSI initiators can discover iSCSI Target Server, use the following Windows PowerShell commands:

```
PS > & iscsicli AddTargetPortal <ip-address> 3260
```

```
PS > & iscsicli.exe ListTargets
```

If the commands execute without errors, the initiator is capable of discovering the targets that are offered by the server

Verify that the iSCSI initiators can log on

The second step is to verify that the iSCSI initiators are able to log on to the iSCSI targets that are exposed by iSCSI Target Server. This can be accomplished by using the following Windows PowerShell command:

```
PS > & iscsicli.exe  LoginTarget <target IQN> T <ip address> 3260 Root\ISCSIPRT\0000_0 *
* * * * * * * * * * * * *
```

Note

If you are using CHAP and Reverse CHAP authentication, you may need to specify more parameters. For more information, consult the documentation in the iscsicli.exe.

If the command executes without errors, the iSCSI initiator has successfully logged on to the target, and the disks are exposed to iSCSI Target Server.

Troubleshoot the iSCSI Software Target Migration

Troubleshooting iSCSI Software Target migration issues involves first viewing the contents of the Windows Server Migration Tools deployment log and the result objects. For more information, see Locate the deployment log file and View the content of Windows Server Migration Tools result objects.

Understanding the messages from the iSCSI Target Migration tool

The iSCSI migration tool (iSCSITargetSettings.PS1) does not produce a log file, but it prints diagnostics messages on the console. These messages show the outcome of the operations that are being attempted and performed.

For example, the following message shows information about saved settings:

```
PS C:\Windows\System32\WindowsPowerSehll\V1.0\Modules\IscsiTarget>

.\iSCSITargetSettings.PS1 -Export -FileName $env:temp\test00000000.xml


Number of Target(s) saved in the Export settings: 4.

Target Names:

   test000

   test001

   test002

   test1111


Number of Virtual Disk(s) saved in the Export settings: 3.

Virtual Disk DevicePaths:
```

```
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd

Number of Virtual Disk(s) NOT saved in the Export settings: 0.
Virtual Disk DevicePaths:
```

The following message shows that not all the virtual disks are eligible for migration:

```
PS D:\Program Files\ISCSI Target> .\iSCSITargetSettings.PS1 -Export -FileName
$env:temp\test00000001.xml

Number of Target(s) saved in the Export settings: 4.
Target Names:
   test000
   test001
   test002
   test1111

Number of Virtual Disk(s) saved in the Export settings: 3.
Virtual Disk DevicePaths:
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd

Number of Virtual Disk(s) NOT saved in the Export settings: 1.
Virtual Disk DevicePaths:
    \\?\GLOBALROOT\Device\HarddiskVolumeShadowCopy{B6B3C77C-93CC-11DF-B3FE-
001CC0C60A6E}\test000.vhd
```

The following message shows information about the settings restore phase:

```
PS C:\Program Files\ISCSI Target> .\iSCSITargetSettings.PS1 -Import -file
$env:temp\test00000000.xml

Importing settings from file
'E:\Users\administrator\AppData\Local\Temp\test00000001.xml'.
```

```
The operation may take a long time.


Number of Target(s) imported from the Import settings: 4.
Targets:
    test000
    test001
    test002
    test1111


Number of Virtual Disk(s) imported from the Import settings: 3.
Virtual Disk:
   s:\test000.vhd
   t:\test000.vhd
   z:\test000.vhd
```

Roll Back a Failed iSCI Software Target Migration

If iSCSI initiators have successfully reconnected to the iSCSI Target Server computer, the migration is successful and complete. This topic discusses the tasks that should be performed in the event of a failed migration.

Restoring the role if the migration failed

If migration does not complete successfully, a rollback procedure is required to undo any changes to the source server, other servers, and client computers, and then restore the source server back into service.

Rollback requirements

The rollback procedure requires that the source server is available in the same state as it was after the “Remove the network identity of the iSCSI Software Target server” step in the “Prepare your source server” section. For more information, see Remove the network identity of the iSCSI Software Target server.

During the source server preparation steps, none of the steps performed permanently changed the existing configuration of the server because all of the operations were substantially read operations.

The estimated time to complete the rollback is equivalent to the time that it takes to re-establish the network identity of the source server. This operation may require rolling back changes to the DHCP servers, DNS server, or Active Directory Domain controllers.

Roll back iSCSI initiators on other computers

The other computers in the enterprise that are affected by migrating ISCSI Software Target are the iSCSI initiators.

In the case of a rollback, the iSCSI initiators that were configured to log on to the destination server need to be rolled back to the source server. Use the following Windows PowerShell commands:

1. To log out of an existing iSCSI session:

   ```
   PS > & iscsicli.exe sessionlistPS > & iscsicli.exe LogoutTarget <session id>
   ```

2. To discover the iSCSI Software Target source server:

   ```
   PS > & iscsicli AddTargetPortal <source server ip address> 3260PS > iscsicli.exe
   ListTargets
   ```

3. To log on to the targets on the iSCSI Software Target source server:

   ```
   PS > & iscsicli.exe  LoginTarget <target IQN> T < source server ip address> 3260
   Root\ISCSIPRT\0000_0 * * * * * * * * * * * * *
   ```

Roll back iSCSI Software Target on a standalone source server

This step will undo the network identity removal that is described in “Remove the network identity of the iSCSI Software Target server’.

Possible scenarios include:

- Restore the NetBIOS fully qualified domain name to the source server, and assign the required IP addresses to the source server.
- Restore any DNS assignments (for example, reverse lookup and DHCP assignment).
- Restore any identities that were previously assigned in Active Directory.

Each scenario requires potentially updating information in the DNS server, Active Directory, or DHCP server, according to the methodology that is used to assign IP addresses and network names to the servers in the enterprise.

The intent of this step is to ensure that upon completion of the migration steps, the iSCSI initiators are able to locate the source server (either through explicit reconfiguration, or implicitly through the computer name or IP address re-assignment).

Roll back iSCSI Software Target on a clustered source server

Rolling back iSCSI Software Target on a clustered source server requires two steps:

Step 1: Roll back cluster network name changes

This step will undo the network identity removal described in “Remove the network identity of the iSCSI Software Target server”.

In a clustered configuration, network names are established by the Server Principal Name that is assigned in Active Directory to the cluster when the cluster was formed.

To re-establish network names that were possibly deleted or retired, the cluster administration utilities must be used. For more information, see Migrating Settings to a Failover Cluster Running Windows Server 2008 R2.

Step 2: Move resource groups to the preferred owner node

After the client access points have been re-established, the resource groups need to be moved back to their preferred owner node.

The resource groups were moved to a single node as part of the steps performed in “Capture the existing settings: clustered configuration”.

To move the resource groups back to their preferred owner node, use the following Windows PowerShell command:

```
PS > & cluster.exe /cluster:<cluster name> GROUP <group name> /moveto:<node name>
```

Note

The group name and the node names were obtained during the previous preparation tasks.

Roll back ISCSI Target Server on a standalone destination server

To roll back iSCSI Target Server on a standalone destination server that is running Windows Server 2012 or Windows Storage Server 2012, uninstall the **iSCSI Target Server** role service using Server Manager.

Roll back ISCSI Target Server on a clustered destination server

To roll back iSCSI Target Server on a destination server that is running Windows Server 2012 or Windows Storage Server 2012 in a clustered configuration, first remove any client access point that was created for iSCSI Target Server and then uninstall the **iSCSI Target Server** role service using Server Manager.

Retiring ISCSI Software Target on a source server

Retiring ISCSI Software Target 3.2 or iSCSI Software Target 3.3 on your source server requires using the following Windows PowerShell commands:

Retire iSCSI Software Target

1. Find the package GUID:

   ```
   PS > Get-WmiObject -Class Win32_product | Where-Object { $_.packageName -match
   'iscsitarget'}
   ```

2. Uninstall the package:

   ```
   PS > & msiexec /uninstall <package GUID> /qr
   ```

Retiring a source server

In a standalone configuration, there are no particular procedures for retiring the source server. In a clustered configuration, the client access points that are devoted to iSCSI Software Target access can be removed by using the following Windows PowerShell command:

```
PS > Remove-ClusterGroup -Name <resource group name> -RemoveResources -Force
```

Migrate Health Registration Authority to Windows Server 2012

This document provides guidance for migrating the Health Registration Authority (HRA) role service from an x86-based or x64-based server running Windows Server® 2008, Windows Server® 2008 R2, or Windows Server® 2012 to a new Windows Server 2012 server.

About this guide

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic by clicking the stars in the upper-right corner of the page (1=poor, 5=excellent), and then add comments that support your rating. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

This guide describes the steps for migrating existing HRA server settings to a server that is running Windows Server 2012. By using this documentation, you can simplify migration, reduce or eliminate server downtime, and help eliminate possible conflicts that might otherwise occur during HRA migration.

Target audience

This guide is intended for information technology (ITOS) administrators, IT professionals, and other knowledge workers who are responsible for the operation and deployment of HRA servers in a managed environment.

What this guide does not provide

This guide does not provide detailed steps to migrate the configuration of other services used with NAP, such as Network Policy Server (NPS) or Active Directory Certificate Services (AD CS). These procedures are found in the Migrate Network Policy Server to Windows Server 2012 and the Active Directory Certificate Services Migration Guide

(http://go.microsoft.com/fwlink/p/?LinkID=156771). Instructions to perform specific procedures in these other guides are provided as necessary to complete migration of the HRA server.

Supported migration scenarios

This guide provides you with instructions for migrating an existing server that is running the HRA role service to a server that is running Windows Server 2012. This includes guidance for installing the prerequisite IIS server role and NPS role service. If your server is running additional services, it is recommended that you design a custom migration procedure specific to your server environment based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server 2008 R2 TechCenter (http://go.microsoft.com/fwlink/p/?LinkID=128554).

Caution

If your source server provides other roles and services in addition to HRA, migrating the computer name and IP configuration can cause these services to fail. You must verify the impact of these procedures before performing them during HRA migration.

Supported operating systems

The following table displays the minimum operating system requirements.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2008	Windows Server 2012	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012	x64-based
x64-based	Windows Server 2012	Windows Server 2012	x64-based

- The NPS and HRA roles services are not available in Server Core editions. Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers. However, If you have configured AD CS on the source server as an enterprise certification authority (CA), the destination CA server must be running Enterprise or Datacenter editions of Windows Server 2012.
- Migration from a source server to a destination server that is running an operating system with a different installed language is not supported. For example, migration of server roles from a computer that is running Windows Server 2008 with a system language of French to a computer that is running Windows Server 2012 with a system language of German is not supported. The system language is the language of the localized installation package that was used to set up the Windows operating system.
- Both x86- and x64-based migrations are supported for Windows Server 2008. All editions of Windows Server 2008 R2 and Windows Server 2012 are x64-based.

Supported role configurations

This guide provides procedures to migrate all HRA server settings, including any custom CA and request policy settings. This guide also provides instructions for configuring minimum IIS role requirements on the destination server.

Migrating prerequisite roles

HRA is a role service under the Network Policy and Access Services (NPAS) server role. To install HRA, you must also install NPS and IIS on the same computer. If these services are not already installed, they will be added automatically by the Add Roles and Features Wizard when you choose to install HRA.

HRA also requires a connection to one or more servers running AD CS that are configured to provide NAP health certificates. AD CS can be installed on the same computer with HRA, or it can be installed on another computer. If any HRA severs in your organization are configured to use AD CS on the source server for health certificate requests, you must install AD CS on the destination HRA server and configure it to provide health certificates, or you can change the CA configuration of your HRA servers.

Consider the following information about prerequisite roles and required services on the destination HRA server.

1. **NPS**. The NPS role service must be migrated before you can test and verify the functionality of HRA on the destination server. If NPS on the source server is only used with HRA, either as a standalone NAP IPsec health policy server or as a RADIUS proxy for another health policy server, this guide provides references to specific procedures in the Migrate Network Policy Server to Windows Server 2012 that are required to migrate required NPS policies and settings. If the NPS role on the source server is used for purposes other than IPsec NAP, or if the source server is a member of RADIUS clients or remote RADIUS server groups on other servers in your organization, consult the Migrate Network Policy Server to Windows Server 2012 for detailed migration instructions prior to migrating HRA.
2. **AD CS**. During installation of HRA, you can choose to install AD CS on the same computer, to use an existing NAP CA on a different computer, or to select a CA later. You can also choose to install AD CS as an enterprise CA or a standalone CA.

   Warning
   After you install AD CS on the HRA server, you cannot change the name of the HRA server.

   - If you install AD CS on the same computer with HRA, you must configure AD CS on the destination HRA server to provide NAP health certificates.
     - If AD CS is installed as an enterprise CA, use procedures in this guide to configure permission settings for the NAP CA. See the Active Directory Certificate Services Migration Guide (http://go.microsoft.com/fwlink/p/?LinkID=156771) for procedures to migrate health certificate templates to the destination server.
     - If AD CS is installed as a standalone CA, this guide provides all permission setting procedures that are required to configure a NAP CA on the destination server. If you use the local CA for other purposes than issuing NAP health certificates, or you have

a custom configuration, see the Active Directory Certificate Services Migration Guide (http://go.microsoft.com/fwlink/p/?LinkID=156771) for detailed instructions to migrate CA settings.

- If you use an existing NAP CA on a different computer, you do not need to configure AD CS on the destination server.
- If you choose to select a CA later, you do not need to configure AD CS on the destination server. If you choose to install AD CS on the destination HRA server later, see Deploying NAP Certification Authorities.

If AD CS on the source server is also used to issue certificates that are not health certificates, see the Active Directory Certificate Services Migration Guide (http://go.microsoft.com/fwlink/?LinkID=156771) for procedures to migrate AD CS.

3. **IIS**. If the prerequisite IIS server role is used for any purposes other than the HRA, or is run with customized settings beyond adding an SSL certificate, follow procedures in the Internet Information Services Migration Guide prior migrating the HRA. If the IIS server role is only used with HRA, use the procedures in this guide to migrate IIS.

   Important
   To maintain HRA performance, the default IIS connection settings must be modified to increase the maximum number of concurrent connections. To perform this procedure, see the Configure IIS connection settings section in Configure an HRA server for NAP.

Migration scenarios that are not covered

The following migration scenarios are not covered in this document:

- **Upgrade**. Guidance is not provided for scenarios in which the new operating system is installed on existing server hardware by using the **Upgrade** option during setup.
- **Workgroup**. Guidance is not provided for migration of HRA settings to or from a non-domain-joined server.

Overview of migration process for this role

HRA server migration is divided into the following major sections:

- HRA Server Migration: Preparing to Migrate
- HRA Server Migration: Migrating the HRA Server
- HRA Server Migration: Verifying the Migration
- HRA Server Migration: Post-migration Tasks

The pre-migration process involves establishing a storage location for migration data, collection of information that will be used to perform the server migration, and operating system installation on the destination server. The HRA migration process includes using the Network Shell (netsh) utility from a command line on the source server to obtain the required HRA settings, and procedures on the destination server to install the required roles and migrate the HRA settings. Verification

procedures include testing the destination server to ensure it works correctly. Post-migration procedures include retiring or repurposing the source server.

Impact of migration

If your migration plan involves configuring the destination server with a different host name from the source server, the trusted server group settings on NAP client computers that use the source HRA server must be updated to use the destination HRA server. This approach has the advantage that it allows the source and destination HRA servers to run simultaneously until testing and verification is complete.

If your migration plan involves configuring the destination server with the same name as the source server, then the source server must be decommissioned and taken offline prior to joining the destination server to the same domain with the same host name. To eliminate downtime in this scenario, NAP client computers must have access to a secondary HRA server in addition to the source and destination servers. To eliminate short term name resolution issues, use the same IP address configuration on the source and destination server.

If the NPS role on the source server is used for purposes other than IPsec NAP, client computers might fail to access the network during the server migration process. For example, if the source server is used for VPN client authentication, consult the Migrate Network Policy Server to Windows Server 2012 for detailed migration instructions prior to migrating HRA.

Impact of migration on the source server

- When deploying the destination server with a different host name, there is no impact to the source server.
- When deploying the destination server with the same host name, the source server must be decommissioned and taken offline prior to joining the destination server to the domain.

Impact of migration on other computers in the enterprise

- When deploying the destination server with a different host name, the NAP client settings for all machines configured to use the HRA must be updated. There is little to no downtime in this scenario if the procedures in this guide are followed.
- When deploying the destination server with the same host name, clients will not be able to obtain a health certificate shortly after the source server is decommissioned, unless a secondary HRA server is deployed.

Permissions required to complete migration

The following permissions are required on the source server and the destination server:

- Domain administrative rights are required to configure and authorize the HRA server, and configure group policy settings for NAP clients.
- Local administrative rights are required to install or manage the server running HRA.

Estimated duration

The migration can take two to three hours, including testing.

See Also

HRA Server Migration: Preparing to Migrate

HRA Server Migration: Migrating the HRA Server

HRA Server Migration: Verifying the Migration

HRA Server Migration: Post-migration Tasks

Network Access Protection Design Guide

Network Access Protection Deployment Guide

HRA Server Migration: Preparing to Migrate

Migration of Health Registration Authority (HRA) Server includes the following tasks:

- Choose a migration file storage location
- Prepare your source server
- Prepare your destination server

Complete the steps or procedures in these sections to prepare your environment for migration.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Choose a migration file storage location

First, choose a location where migration files will be kept.

To choose a storage location

1. Select a file storage location where migration files will be kept. The storage location can be a network share that is accessible by both the source and destination server, or portable media that can be transferred from one server to another.

Prepare your source server

Follow these steps to prepare an x64 or x86-based server running Windows Server® 2008, Windows Server® 2008 R2, or Windows Server® 2012 for HRA migration.

To prepare the source server

1. Determine the domain, server name, IP address, and passwords on the source server.
2. Determine the group membership of the source server in Active Directory Domain Services (AD DS), including security group and OU membership. This can be done using the Active Directory Users and Computers console (dsa.msc) or Server Manager on a domain controller.

Prepare your destination server

Follow these steps to prepare an x64-based destination server running Windows Server 2012 for HRA migration.

To prepare the destination server

1. Install Windows Server 2012 on the destination server.
2. Configure the host name of the computer, and configure network settings as desired. Do not join the computer to the domain yet.
3. Install all critical updates and service packs on the destination server.
4. Verify the server has access to the migration file storage location.

See Also

Migrate Health Registration Authority to Windows Server 2012

HRA Server Migration: Migrating the HRA Server

HRA Server Migration: Verifying the Migration

HRA Server Migration: Post-migration Tasks

Network Access Protection Design Guide

Network Access Protection Deployment Guide

HRA Server Migration: Migrating the HRA Server

This topic contains steps and procedures for migrating the Health Registration Authority (HRA) role service from a legacy source server to a new x64-based destination server running Windows Server® 2012.

Important

The NPS role service must be installed before HRA can be configured on the destination server. If NPS on the destination server will only be used with HRA, you can use the Add Roles and Features Wizard in Server Manager to install both HRA and NPS role services together. Following service installation, see the Migrate Network Policy Server to

Windows Server 2012 for procedures to migrate NPS settings to the destination server. When you have completed migration of NPS, continue performing the procedures in this guide to complete HRA migration.

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Migrating settings from the source server

Use the following procedures to export the HRA settings from your x86-based or x64-based source HRA server prior to migrating to an x64-based server running Windows Server 2012.

Important

If your migration plan involves configuring the destination server with the same host name as the source server, then the source server must be decommissioned and taken offline prior to joining the destination server to the domain. To eliminate downtime in this scenario, a secondary HRA server should already be deployed before proceeding. For information about deploying a new HRA server, see Install HRA using the Add Roles and Features Wizard.

To export settings from the source server

1. On the source HRA server, type the following command at an elevated command prompt, and then press ENTER:

   ```
   netsh nap hra export filename=c:\hra_export.xml
   ```

2. Copy the *hra_export.xml* file from the *c:*\ directory to the migration file storage location you have chosen.
3. Configuration settings for the NPS role service must also be exported from the source server. Use the procedures provided in the **Migrating settings from the source server** section of the NPS Server Migration: Migrating the NPS Server topic to export these settings.
4. Copy the exported HRA configuration file to the migration file storage location you have chosen.

Configuring the destination server

Use the following procedures to configure the destination with the required identity, certificates, and services. If the destination server will have a different host name and IP address from the source server, then the source server can remain online and in service until testing and verification of the destination server is complete. When you have completed configuring the destination server's identity, certificates, and services, you can begin migrating HRA settings from the source to destination server.

Important

Some services and settings on the destination server might already be migrated due to the migration of prerequisite roles. Before you configure the destination HRA server, consult the **Migrating prerequisite roles** topic in this guide to determine the configuration settings for NPS, AD CS, and IIS that must be migrated first.

To configure the destination server

1. Add the destination server to the domain of the source server. If the destination server will use the same name as the source server, you must ensure the source server is decommissioned as described in the **Impact of migration** topic.
2. Add the destination server to all security groups and organizational units (OUs) of which the source HRA server is a member. In most cases, the HRA server is a member of the IPsec boundary OU. Members of the boundary OU typically have IPsec policies applied that allow communication with both compliant and noncompliant computers. For more information on OUs and required IPsec policy settings, see Checklist: Deploy IPsec Policies for NAP (http://go.microsoft.com/fwlink/p/?linkid=229649).
3. To update Group Policy settings on the destination server, run the following command at an elevated command prompt:

   ```
   gpupdate /force
   ```

   Note
   To apply new security group membership settings, you must restart the destination server.
4. If client computers will use SSL to request health certificates from HRA, you must provision the destination server with an SSL certificate. For more information, see Configure an SSL Certificate for HRA (http://go.microsoft.com/fwlink/p/?LinkId=229650), or use the process defined within your organization for provisioning an SSL certificate.
5. Install the HRA role service on the destination server.

Install HRA using the Add Roles and Features Wizard

a. In Server Manager, click **Manage** and click **Add Roles and Features**.
b. On the **Before you begin** page, click **Next**.
c. On the **Select Installation Type** page, click **Role/Feature Based Install** and then click **Next**.
d. On the **Select destination server** page, click **Select a server from the server pool**, click the names of the servers where you want to install HRA and then click **Next**.
e. On the **Select server roles** page, click **Network Policy and Access Services**, and then click **Next** three times.

   Note
   If the Network Policy Server role service is already installed, expand the NPAS node and select **Health Registration Authority**. Click

Next five times and continue with step below.

f. On the **Select Role Services** page, click **Health Registration Authority**, and in the **Add Roles and Features Wizard** dialog box, verify that **Include management tools (if applicable)** is selected, click **Add Features**, and then click **Next** five times.

g. On the **Certification Authority** page, choose **Select a CA later using the HRA console**, and then click **Next**.

> **Note**
> Certification Authority settings for HRA will be configured when you migrate settings from the source server.

h. On the **Authentication Requirements** page, choose **No, allow anonymous requests for health certificates**, if the destination HRA will provide health certificates to workgroup computers. If health certificates will be issued to domain-joined clients only, choose **Yes, require requestors to be authenticated as members of a domain (recommended)**. Click **Next** to continue.

i. On the **Server Authentication Certificate** page, click **Choose an existing certificate for SSL encryption (recommended)**, click the certificate displayed under this option, and then click **Next**. If multiple certificates are displayed, or you are not sure if the certificate displayed can be used for SSL encryption, see Install the HRA Role Service for more information.

j. Click **Next**, and then click **Install**.

k. On the **Installation Results** page, verify that installation was successful and then click **Close**.

The following Windows PowerShell command performs the same function:

```
Add-WindowsFeature NPAS-Health
```

Migrating settings to the destination server

Follow the procedure below to migrate HRA settings from the source to destination server.

To migrate the settings to the destination server

1. On the destination server, type the following command at an elevated command prompt, and then press ENTER:

   ```
   netsh nap hra import filename = c:\hra_export.xml
   ```

   Replace *c:\hra_export.html* with the path and file name of the HRA configuration file that you exported in the previous procedure: Migrating settings from the source server.

   > **Note**
   > If you receive the error message “Cannot create a file when that file already

exists," reset the HRA configuration and then perform this procedure again. To reset the HRA configuration, type the following command at an elevated command prompt and then press ENTER: **reg delete HKLM\Software\Microsoft\HCS\CAServers**.

2. Verify that the settings have been imported successfully. To review HRA settings, type the following command at a command prompt and then press ENTER:

```
netsh nap hra show configuration
```

3. If the name of the certification authority will change as a result of the migration, type the following commands at an elevated command prompt to add the name of the correct CA and delete the name of the old CA. Replace *\\srv1.woodgrovebank.com\woodgrovebank-srv1-CA* and *1* with the name and processing order of the CA you wish to use.

```
netsh nap hra delete caserver name =
"\\srv1.woodgrovebank.com\woodgrovebank-srv1-CA"
netsh nap hra add caserver name =
"\\srv2.woodgrovebank.com\woodgrovebank-srv2-CA"
processingorder = "1"
```

You can use the output of the **netsh nap hra show configuration** command to view the name and processing order format for the previous CA. For more information, see HRA Certification Authority Commands.

Configuring the Certification Authority

The destination HRA server name must be given security permissions to request, issue, and manage certificates. It must also be granted permission to manage the CA so that it can periodically clear expired certificates from the certificate store.

If the host name of the destination server is different from the source server, then the certification authority for the NAP deployment must be configured with permissions settings for the new HRA. If the destination HRA server is already a member of an OU or group that has permissions to manage the NAP CA, then this procedure is not required.

To configure the Certification Authority with permissions for the destination HRA

1. On the **Start** screen, type **certsrv.msc**, and then press ENTER on the CA server.
2. In the Certification Authority console tree, right-click the CA name, and then click **Properties**.
3. Click the **Security** tab, and then click **Add**.
4. Click **Object Types**, click the **Computers** check box, and then click **OK**.
5. If the CA is located on a different computer than the destination HRA server, type the name of the destination HRA server under **Enter the object names to select**, and then click **OK**.

Note

If the CA is installed on the same computer as the destination HRA server, type **NETWORK SERVICE** under **Enter the object names to select**, and then click OK.

6. Click the name of the destination server, or click **NETWORK SERVICE**, select **Allow** for the **Issue and Manage Certificates**, **Manage CA**, and **Request Certificates** check boxes, and then click **OK**.
7. Close the Certification Authority console.

Configuration tips for migrating the Certification Authority

If the HRA uses a CA that was recently migrated in parallel using the Active Directory Certificate Services Migration Guide (http://go.microsoft.com/fwlink/?LinkID=156771), consider the following:

1. If the HRA uses an Enterprise CA that was recently migrated, the template for the System Health Authentication certificate used by the HRA must be re-issued in Active Directory before it can be used. This procedure is described in the Restoring the certificate templates list section of the AD CS Migration: Migrating the Certification Authority topic and in the Backing up a CA templates list procedure of the AD CS Migration: Preparing to Migrate topic in the Active Directory Certificate Services Migration Guide (http://go.microsoft.com/fwlink/?LinkID=156771).
2. If the HRA uses a Root CA that was recently migrated, then all NAP IPsec policies configured in Group Policy need to be edited to use the correct Root CA. For more information, see Configure IPsec GPOs.

See Also

Migrate Health Registration Authority to Windows Server 2012

HRA Server Migration: Preparing to Migrate

HRA Server Migration: Verifying the Migration

HRA Server Migration: Post-migration Tasks

Network Access Protection Design Guide

Network Access Protection Deployment Guide

HRA Server Migration: Verifying the Migration

After the migration of your Health Registration Authority (HRA) server is complete, you can perform some tasks to verify that the migration was successful.

Verifying HRA Functionality

In order to verify the HRA functionality, the URL of the destination server must be configured in the NAP client trusted server group settings. This is typically done using Group Policy.

To test the destination server with minimal impact to your current NAP deployment, you can add a secondary trusted server group to NAP client settings. The new trusted server group can contain the URL of the newly migrated destination server. When a secondary trusted server group is configured, compliant client computers will receive a health certificate from both the source HRA and the destination HRA. Once you have verified that client computers are successfully receiving health certificates from the destination server, the new trusted server group can be removed, and the original trusted server group can be updated to use the destination server instead of the source server.

Adding a new trusted server group for testing

To add a new trusted server group in group policy that will be used to test the destination HRA, see Configure Trusted Server Groups in Group Policy.

The new trusted server group should be ordered below any other groups configured, and only the URL of the destination server (for example: *https://destination.contoso.com/domainhra/hcsrvext.dll*) should be added.

Note

If there are multiple GPOs for NAP clients in your organization, you can make these changes to one GPO that applies to a group of clients you wish to test.

Testing the HRA with a NAP client

Use the following procedure to test the functionality of the destination server using a domain-joined NAP client in your deployment.

To test the HRA functionality using a NAP client

1. On the client computer, On the **Start** screen, type **gpupdate /force**, and then press ENTER. This updates the Group Policy configuration for the client.
2. On the **Start** screen, type **cmd**, type **netsh nap client show grouppolicy**, and then press ENTER.
3. In the command output, under **Enforcement clients**, verify that the **Admin** status of the **IPSec Relying Party** is **Enabled**.
4. In the command output, under **Trusted server group configuration**, verify that the trusted server group and destination server URL you configured previously are displayed.
5. Next, the NAP Agent service will be restarted to verify that the client computer successfully receives a health certificate from the new destination HRA.
6. To restart the NAP Agent service, at the command prompt, type `net stop napagent && net start napagent`, and then press ENTER. Verify that the commands completed

successfully.

7. At the command prompt, type `eventvwr.msc`, and then press ENTER. This launches the Event Viewer.
8. In Event Viewer, browse to **Windows Logs /Application and Services Logs/Microsoft/Windows/Network Access Protection/Operational**.
9. In the details pane, under **Event ID**, locate the most recent occurrences of event 22. Event 22 is displayed each time a client computer acquires a health certificate from HRA. Double-click these events to review detailed information about the certificate acquisition. Verify that the URL of the destination server is displayed in at least one event as the source of the certificate.
10. Close Event Viewer.

See Also

Migrate Health Registration Authority to Windows Server 2012

HRA Server Migration: Preparing to Migrate

HRA Server Migration: Migrating the HRA Server

HRA Server Migration: Post-migration Tasks

Network Access Protection Design Guide

Network Access Protection Deployment Guide

HRA Server Migration: Post-migration Tasks

After all migration steps are complete and you have verified the migration of the Health Registration Authority (HRA) role service, perform the following post-migration tasks.

Deploying final client settings

To finish deploying the destination server, all NAP clients must be updated to obtain a health certificate from the destination server URL instead of the source server URL (if different). These settings are typically configured using Group Policy. If the source and destination URLs are different, each GPO in your NAP deployment that uses the new trusted server group settings must be modified. If your organization uses other methods to push NAP client settings to clients, then you might also need to modify those procedures.

Warning

If you have configured HRA automatic discovery on your network and the name of your source and destination HRA servers are different, you must modify DNS service (SRV) records to deploy the new trusted server group setting to client computers. For more information, see Configure HRA Automatic Discovery.

To configure final NAP client settings in group policy

1. On a domain controller or member server with the Group Policy Management feature installed, click **Start**, click **Run**, type **gpmc.msc**, and then press ENTER.
2. In the Group Policy Management console tree, open **Group Policy Objects**, right-click the name of the GPO you want to edit, and then click **Edit**. The Group Policy Management Editor opens.
3. In the console tree, open **Computer Configuration/Policies/Windows Settings/Security Settings/Network Access Protection/NAP Client Configuration/Health Registration Settings/Trusted Server Groups**.
4. Delete the secondary trusted server group that was added for testing purposes. To delete this group, right-click the name of the trusted server group, and click **Delete**.
5. Double-click the name of the primary trusted server group you wish to edit.
6. Click the URL of the source server in the list, and then click **Edit**.
7. Replace the source server URL with the destination server URL.
8. Click **OK**.
9. In the console tree, right-click **NAP Client Configuration**, and then click **Apply**.
10. Close the Group Policy Management Editor window.
11. If you are prompted to apply settings, click **Yes**.
12. Repeat the testing procedure as described in HRA Server Migration: Verifying the Migration to verify that deployment of the destination server is successful.

Restoring the role in the event of migration failure

If the destination server is deployed simultaneously with the source server using a different host name, then the configuration prior to migration can be restored by changing the NAP client settings to use the URL of the source HRA server. To restore the previous configuration, perform the steps described in the **Deploying final client settings** section of the HRA Server Migration: Verifying the Migration topic, replacing the destination server URL with the source server URL.

If the destination server replaced the source server using the same host name, then the destination server will need to be renamed, unjoined from the domain, or otherwise decommissioned in order to bring the source server back online.

Retiring the Source Server

Once the destination HRA has been configured, tested, and verified, and the URL of the source HRA has been removed from group policy, then the HRA role on the source server may be retired.

- The source server can be taken offline and physically retired or repurposed. Follow your organization's policy regarding server decommissioning.

- To retire only the HRA role on the source server, in the Server Manager console tree, click **Network Policy and Access Services**. In the details pane, click **Remove Role Services**, and then use the Remove Role Services wizard to select and remove the HRA role service.

Note
If the source server was configured to use a certification authority on a different machine, consider removing permissions for the source server from the certification authority.

Troubleshooting migration

If you encounter problems while verifying the HRA migration, see Fixing Health Certificate Problems in the NAP Troubleshooting Guide for help troubleshooting these problems.

See Also

Migrate Health Registration Authority to Windows Server 2012

HRA Server Migration: Preparing to Migrate

HRA Server Migration: Migrating the HRA Server

HRA Server Migration: Verifying the Migration

Network Access Protection Design Guide

Network Access Protection Deployment Guide

Migrate Hyper-V to Windows Server 2012 from Windows 2008 R2

Hyper-V enables you to create a virtualized server computing environment using a technology that is part of Windows. This guide provides information and instructions about migrating the Hyper-V role—including virtual machines, data, and operating system settings—from the source server running Hyper-V in an earlier version of Windows to the destination server that is running the Windows Server® 2012 operating system.

About this guide

Note
Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic, and then add comments that support your rating. If you are viewing this topic in Lightweight View, click **Rate this topic** at the top of the page. In Classic View, click the stars in the upper-right corner of the page (1=poor, 5=excellent). Describe what you liked, did not like, or want to see in future versions of the topic. To

submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

This guide describes how to migrate the Hyper-V role by providing preparation, migration, and verification steps.

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process. For more information about installing and using the migration tools on both source and destination servers, see the Windows Server Migration Tools Installation, Access, and Removal Guide.

Target audience

This document is intended for information technology (IT) professionals who are responsible for operating and deploying Hyper-V in a managed environment.

What this guide does not provide

The following items are not covered in this guide because they are not supported by the migration tools:

- Clustering scenarios are not supported by this migration process. For information about how to perform a migration in a clustered environment, see the Migrating Clustered Services and Applications to Windows Server 2012 Step-by-Step Guide Migrating Clustered Services and Applications to Windows Server 2012.
- Upgrading roles on the same computer is out of scope for this guide.
- Migrating more than one server role at one time.
- Migrating Hyper-V from one server running Windows Server 2012 to another server running 2012. Instead, this process is supported by several of the new Hyper-V management tools and features. The general process is as follows:
    - Determine whether you will use export and import or live migration to move the virtual machines. Export and import can be used in either a workgroup or a domain environment but requires that the virtual machine is turned off. Live migration requires a domain environment as well as some configuration, but allows you to move running virtual machines.
    - Add the Hyper-V role to the destination server. You can configure default storage locations and live migration when you add the role. For instructions, see Install Hyper-V and Configure a Virtual Machine.
    - Configure virtual switches and, optionally, other networking features on the destination server. Management tools include the cmdlets New-VMSwitch and Set-VMSwitch in the Hyper-V module, and the Virtual Switch Manager in the Hyper-V Manager snap-in.
    - Move the virtual machines by exporting and importing them, or performing live migrations. Management tools include the cmdlets Export-VM and Import-VM, and the

Export, **Import**, and **Move** menu commands in Hyper-V Manager. For more information about using live migration to move a virtual machine, see Configure Live Migration and Migrating Virtual Machines without Failover Clustering.

- For a list of the cmdlets included in the Hyper-V module, see http://technet.microsoft.com/library/hh848559.

Supported migration scenarios

This guide provides you with instructions for migrating an existing server that is running the Hyper-V role on an earlier version of Windows Server to a server that is running Windows Server 2012. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server Migration Portal.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

Supported operating systems

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x64-based	Windows Server 2008 with Service Pack 2, full installation option only	Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012, both full and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012, both full and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are

supported. If an operating system is not listed, then it is not supported. The stand-alone product Microsoft Hyper-V Server is not supported.

Standard, Enterprise, and Datacenter editions of Windows Server running Hyper-V are supported as either source or destination servers.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Supported role configurations and settings

This section identifies the configurations and settings that can be migrated by using the migration tools, and the configurations and settings that must be migrated manually. The following table provides a summary.

Configurations and settings	Type of migration
Virtual machine (configuration and data)	Automated, except as noted below
Hyper-V settings	Automated
Virtual network adapter settings in the management operating system	Automated
External virtual networks	Partially automated, as described below
Virtual machine queue (VMQ) networking settings	Automated
Customized remote administration settings	Manual

The following configurations and settings can be migrated automatically:

- **Most virtual machine configurations**. Virtual machines and their data are moved as part of the migration, but some configurations require manual intervention, as described below.
- **Hyper-V settings**. These include the system-wide settings and the authorization store.

  Note
  If you are migrating from a source server running Windows Server 2008 R2 and have set a MAC address range, that value also is automatically migrated to the destination server.

- **Internal and private virtual networks**.
- **Virtual network adapter settings in the management operating system**. When Hyper-V is configured to use a physical network adapter as a bridge that virtual machines can use to access a physical network, a virtual network adapter is created in the management operating system (which runs the Hyper-V role). For this virtual network adapter, the migration process automatically migrates the IP settings, bindings, and MAC address of this virtual network adapter. However, the connection between the virtual network adapter and the physical network adapter must be re-established manually, as described in the migration steps.
- **Virtual machine queue (VMQ) settings for networking**.

The following configurations and settings require manual intervention after the migration tools are used:

- **Firewall settings**. Firewall settings are recreated on the destination server using the default values that Hyper-V is installed with. If you have modified any of the firewall settings from these default values, you will need to make the same modifications on the destination server.
- **External virtual networks**. The migration tool recreates the virtual networks on the destination server, but recreates external virtual networks as internal virtual networks. You will need to modify each of these networks to connect it to the appropriate physical network adapter on the destination server, as described in the migration steps.
- **VFD and ISO files**. These files are not migrated because they are not required for the virtual machine to operate and are not supported by the Import and Export cmdlets. To make them available to a migrated virtual machine, manually copy these files to the destination server and then reattach them to the virtual machine after it is migrated.
- **Connections to physical disks directly attached to virtual machines**. These connections (sometimes referred to as “pass-through disks”) are not migrated because the disk references might not be valid on the destination server. To make a physical disk available to a migrated virtual machine, connect the disk to the destination server and then to the virtual machine after it is migrated, as described in the migration steps.
- **Customized remote administration settings**. If you have customized Hyper-V for remote access, you will need to perform some additional procedures to recreate the DCOM and WMI Namespace settings. The migration steps identify the point at which you should take perform these procedures, as well as provide a recommended tool or script to complete the procedure.

Migration dependencies

The Hyper-V role is not dependent on any other roles. As a best practice, we recommend that no other roles are installed on a server running Hyper-V.

Migration scenarios that are not supported

The following migration scenarios are not supported:

- The saved state of a virtual machine.
- Virtual machine configuration under one of the following conditions:

    - When the number of virtual processors configured for the virtual machine is more than the number of logical processors on the destination server.
    - When the memory configured for a virtual machine is greater than the available memory on the destination server.
- Consolidation of physical servers to virtual machines, or consolidation of multiple instances of Hyper-V to one instance.

Hyper-V migration overview

Hyper-V role migration involves moving the virtual machines, virtual networks, and all the associated settings from one physical computer to another physical computer in the enterprise. The process supports moving from a server running Hyper-V in Windows Server® 2008 R2 to a server running Hyper-V in Windows Server 2012. The Hyper-V role is not dependent on any other roles.

The migration tools include cmdlets that you use to perform some of the tasks required to migrate the Hyper-V role. The Export cmdlet captures the majority of the Hyper-V settings that are required to perform a successful migration, including the virtual machine configurations, virtual networks, and virtual hard disks. The DCOM and WMI namespace security settings must be migrated separately. The instructions for this are provided later in the guide. On the destination server, the import cmdlets will recreate the virtual machines.

Impact of migration

The following section describes the impact of migration on the source server and on other computers in the enterprise.

Impact of migration on the source server

The source server should be turned off or removed from the network before you run the import cmdlets on the destination server so that there are no conflicts between the virtual machines running on the source server and the virtual machines that will be recreated on the destination server. The point at which you should perform this task is identified in the migration steps, later in this guide.

Impact of migration on other computers in the enterprise

This migration may impact any computer (either virtual or physical) that relies on the applications or workloads running in the virtual machines to be migrated as part of the Hyper-V role migration, because the virtual machines will be offline for the duration of the migration. For example, if a virtual machine hosts a database, any applications in the enterprise that require access to that database will be impacted. As a result, you will need to plan for this downtime by either scheduling a planned outage or by redirecting traffic to other servers to provide the services.

Access rights required to complete migration

The user account that runs the cmdlets and tools must be a member of the local Administrators group on the source server and the destination server.

Estimated duration

The length of time it takes to migrate the Hyper-V role depends on the size of the data to be transferred. Of the various types of files to be transferred, the .vhd files have the largest file sizes (from a few gigabytes to many gigabytes in size). The length of time is affected by the size of the .vhd files and by the network bandwidth.

Additional references

- Hyper-V: Prepare to Migrate
- Hyper-V: Migrate the Hyper-V Role
- Hyper-V: Verify the Migration
- Hyper-V: Post-migration Tasks
- Hyper-V Overview
- Hyper-V Migration Guide (for migration of Hyper-V running in Windows Server 2008 to Windows Server 2008 R2.)
- Windows Server Migration Portal

Hyper-V: Prepare to Migrate

Follow these steps to prepare for migration.

Select and prepare your destination server

To select and prepare the destination server for migration, perform the steps in the order they are given.

Hardware requirements for the destination server

The computer you select as the destination server must meet the following hardware requirements:

- **Storage**. The destination server requires enough storage to hold the virtual hard disks from the source server.
- **Network**. The destination server requires at least as many physical network adapters as the number of physical network adapters in use as external virtual networks on the source server. To determine this, open Hyper-V Manager on the source server, and then open Virtual

Network Manager. Under **Virtual Networks** (in the left pane) note the number of the networks designated as “External”.

- **Memory**. The destination server requires enough memory to run all the virtual machines you plan to run at the same time, as well as run the Hyper-V role. For example, if you run all the virtual machines configured on the source server at the same time, the destination server must have at least as much memory as the sum of memory configured for all virtual machines, plus memory to run the Hyper-V role in Windows Server® 2012.
- **Processor**. The destination server requires at least as many logical processors as the largest number of processors configured on a virtual machine on the source server. Note that if you want to migrate virtual machines with saved states, the processor on the destination server must be compatible with the processor on the source server. The processors must be from the same manufacturer and have compatible steppings.

Software requirements for the destination server

After you select the destination server, prepare the software by doing the following:

Important

If you install Windows Server Migration Tools before you install the Hyper-V role, you must remove the tools and then install the Hyper-V role before you install the tools. For removal instructions, see Install, Use, and Remove Windows Server Migration Tools.

1. Install Windows Server 2012. For more information, see Installing Windows Server 2012.
2. Add the Hyper-V role. For instructions, see Install Hyper-V and Configure a Virtual Machine.

Back up your source server

Before you start migration, back up the source server. If the migration fails, you can use this backup to restore the source server. For information about the different types of backups, see Planning for Backup (http://go.microsoft.com/fwlink/?LinkId=178128).

Install migration tools

Windows Server Migration Tools in Windows Server® 2008 R2 allows an administrator to migrate some server roles, features, operating system settings, shares, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2 to computers that are running Windows Server 2008 R2 or Windows Server 2012.

Complete installation, configuration, and removal instructions for Windows Server Migration Tools are available on the World Wide Web, in Install, Use, and Remove Windows Server Migration Tools. Windows Server Migration Tools must be installed on both the destination server and the source server, in that order.

Migration documentation and tools ease the process of migrating server role settings and data from an existing server that is running Windows Server 2003 and later releases of the Windows operating system to another computer. By using these tools to migrate roles, you can simplify

migration, reduce migration time, increase accuracy of the migration process, and help eliminate conflicts that could otherwise occur during the migration process.

Windows Server Migration Tools is a set of Windows PowerShell® cmdlets. For more information about Windows PowerShell and working with cmdlets, see Windows PowerShell Core on Microsoft TechNet.

Collect configuration details from your source server

Collect the following configuration details about the source server. You will use this information as part of the verification process after you perform the migration.

- Gather identifying information about the set of virtual machines on the source server. If there is a relatively small number of virtual machines on the server (for example, less than 20), you could take a screenshot of the list of virtual machines displayed in the Hyper-V Manager snap-in. Additionally, record the following configuration information for each virtual machine:
  - Amount of memory
  - Number of virtual processors
  - Virtual hard disks (.vhd files) connected to the virtual machine
- For each virtual machine that has snapshots, gather information about the number of snapshots and the structure of the snapshot tree.
- Record information about the external virtual networks. Include information such as the name of each external virtual network. If you plan to migrate the IP settings of the physical network adapters (for example, if the network adapters use static IP addresses that you want to retain), save the IP configuration settings by using the following command:

  ```
  IPConfig /all > IPSettings.txt
  ```

- If you have made any customizations to the Hyper-V security policy, gather all the information about scopes and roles from Authorization Manager (see Using Authorization Manager for Hyper-V Security (http://go.microsoft.com/fwlink/?LinkId=183469)).
- If you have turned off any exceptions in Windows Firewall, record that information.
- If you have granted remote access to the server for Hyper-V management to any user account that is not a member of the Administrators group, record that information. The Hyper-V Remote Management Configuration Utility is a tool is that you can use for this task. However, this tool is not published or supported by Microsoft. For more information, see Hyper-V Remote Management Configuration Utility (http://go.microsoft.com/fwlink/?LinkId=178138).
- If you have granted remote access to the server for Hyper-V management to any user account that is not a member of the Administrators group or Distributed COM Users group, open the registry and navigate to the HKLM\Software\Microsoft\OLE\ key. Find the MachineLaunchRestriction value and record all the information about that value.

  Caution

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

Prepare other computers in the enterprise

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

Depending on the workloads that you have deployed in your virtual machines, you need to take the necessary actions to ensure that the users and clients that obtain services from the virtual machines are not negatively impacted during the migration. In other words, you could either have a planned downtime or redirect the clients and users to alternate, redundant virtual machines while the migration is in progress. The specific actions you take depend on the best practices you have in place for the workloads deployed in the virtual machines.

Additional references

1. Migrate Hyper-V to Windows Server 2012 from Windows 2008 R2
2. Hyper-V: Migrate the Hyper-V Role
3. Hyper-V: Verify the Migration
4. Hyper-V: Post-migration Tasks
5. Hyper-V Overview
6. Hyper-V Migration Guide (for migration of Hyper-V running in Windows Server 2008 to Windows Server 2008 R2.)
7. Windows Server Migration Portal

Hyper-V: Migrate the Hyper-V Role

Migrate the Hyper-V Role

The steps to migrate the Hyper-V role are the same for all of the scenarios defined in “Supported migration scenarios,” earlier in this guide.

Perform migration steps on the source server

To perform migration steps on the source server

1. Prepare the virtual machines for migration by shutting down the virtual machines.
2. Open a Windows PowerShell session with elevated user rights by doing one of the following:
   - To run Windows PowerShell as an administrator from the **Start** screen, right-click the **Windows PowerShell** tile, and in the app bar, click **Run as administrator**.
   - To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.
3. Load Windows Server Migration Tools into your Windows PowerShell session.

   Only load the Windows Server Migration Tools snap-in into a Windows PowerShell session that was opened by using some other method (and into a session where it has not already been loaded). To load Windows Server Migration Tools, type the following, and then press ENTER.

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```
4. From Windows PowerShell, collect data from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The **Export-SmigServerSetting** cmdlet creates an XML file, StoragePathMappings.xml, that contains information about where the virtual machine storage (.vhd and .avhd) files are stored, in the form of folder paths on the source server. The import process uses the StoragePathMappings.xml file to associate the storage files to the appropriate virtual machines on the destination server. If the destination server will use the same drive mapping and folder structure for virtual machine storage as the source server, you do not need to edit the file after the **Export-SmigServerSetting** cmdlet creates it. Otherwise, you must edit this file before you import it to the destination server. Before you run this command, review the following information to determine the information to include in the command:
   - Determine where to store StoragePathMappings.xml. The **-Path** parameter is a required parameter that specifies the location. You can either choose a location that can be accessed by both the source server and the destination server, such as a network location, or copy it from the source server to the destination server.
   - To export user groups, which are used for access control by the Hyper-V security policy and the Hyper-V remote management tools, include the **-User** and **-Group** parameters:

     ```
     -User <Enabled | Disabled | All> -Group
     ```
   - To recreate the same IP settings on the physical network adapters on the destination server as are configured on the source server, include the **-IPConfig** parameter. The **-IPConfig** parameter collects IP information when it is used with the **Export-SmigServerSetting** cmdlet on the source server. The **-IPConfig** parameter applies settings when the **Import-SmigServerSetting** cmdlet is used on the destination server.

     ```
     -IPConfig
     ```

5. After you determine the parameters, run the **Export-SmigServerSetting** cmdlet, where <storepath> specifies the path to the folder where the configuration data file (Svrmig.mig) will be stored. (For example, C:\Migration.) The StoragePathMappings.xml file is created in a subfolder of the <storepath> folder, named VirtualMachines. (For example, C:\Migration\VirtualMachines.)

```
Export-SmigServerSetting -FeatureId Hyper-V -IPConfig -User
All -Group -path <storepath> -Verbose
```

Migrate virtual machine data

The following steps show you how to use Robocopy and a Windows PowerShell script to copy the data from the source server to the destination server. The script parses the folder paths specified in the StoragePathMappings.xml to migrate the data. You can use the StoragePathMappings.xml file stored under <storepath> as a reference to determine which folders need to be transferred.

Important
If you want to use a different drive mapping and/or folder structure on the destination server, edit the StoragePathMappings.xml file before you attempt to migrate the data to the destination server.

To migrate virtual machine data

1. Copy the data from the source server to the destination server. The recommended way to do this is to use the Robocopy command. You can run the command for each file and specify the source and destination locations. (You can use the StoragePathMappings.xml file to determine the source paths.) Or, you can automate this process by using a Windows PowerShell script. Before you run the script, update the StoragePathMappings.xml file with the locations where you want to paste the files on the destination server. The script can parse the StoragePathMappings.xml file and then call the Robocopy command to copy and paste the files.

   The following is an example of such a script. To use this sample, copy the code and paste it into a text editor, then save the file with a .ps1 file name extension in a directory where you want to run the script from. For example, save CopyData.ps1 to C:\migration\.

```
param(
    [string]$xmlFilePath = $(throw "Must pass the fully
qualified file name of Storage Path XML in the command
string"),
    [string]$destinationHost = $(throw "Must pass the
Destination Host Name (NetBiosName), where the files will be
copied to")
)
```

```
Write-Host "XML File Path: " $xmlFilePath
Write-Host "Destination Host Name: " $destinationHost

# Get the content of the XML file
[xml]$xmlFile = Get-content $xmlFilePath

# For each storage path, if the "Copy" attribute is true copy
the files to the destination Host.
foreach ($storagePath in $xmlFile.StoragePaths.storagePath)
{
         if($storagePath.Copy -eq "true")
         {
                # Get the Source directory
                $sourceDirectory = $storagePath.Source

                # Get the Destination directory
                $destinationDirectory =
$storagePath.Destination
                $destinationDirectory = $destinationDirectory
-replace ":","$"
                $destinationDirectory = "\\" +
$destinationHost + "\" + $destinationDirectory

                # Copy the files to the destination host
                robocopy $sourceDirectory
$destinationDirectory /E /XF *.xml /R:5 /W:60 /V
    }
}
```

To run the script, type the full path to the script at the command prompt and pass the fully qualified file name of the StoragePathMappings.xml file (full path and file name) and the name of the destination server as parameters. The file name extension of the script is optional. For more information, see Support for Scripting (http://go.microsoft.com/fwlink/?LinkID=178144).

For example, if you used the folder and file name example shown above, type:

```
c:\migration\copyData.ps1 <XMLPathName>
<DestinationServerName>
```

2. Disconnect the source server from the network so that you avoid any potential MAC address conflicts between the virtual machines on the source and destination servers. MAC address conflicts may impact the availability of the workloads that run on the virtual machines.

Perform migration steps on the destination server

1. If the <storepath> is located anywhere other than locally on the destination server, edit the permissions of the shared folder to grant **Full Control** to the following accounts:
   - The user account that will run the import and export commands. If the same account is used, only one entry is required.
   - The computer account of the source server.
   - The computer account of the destination server.
2. If you used another method instead of the Robocopy command to copy data to the destination server, check the destination folder and delete any .xml files that were copied to that folder.
3. Open a Windows PowerShell session with elevated user rights by doing one of the following:
   - To run Windows PowerShell as an administrator from the **Start** screen, right-click the **Windows PowerShell** tile, and in the app bar, click **Run as administrator**.
   - To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.
4. Load Windows Server Migration Tools into your Windows PowerShell session.

   Only load the Windows Server Migration Tools snap-in into a Windows PowerShell session that was opened by using some other method (and into a session where it has not already been loaded). To load Windows Server Migration Tools, type the following, and then press ENTER.

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```
5. To import the Hyper-V settings to the destination server, run the **Import-SmigServerSetting** cmdlet and all additional parameters that you used with the **Export-SmigServerSetting**.

   ```
   Import-SmigServerSetting -FeatureId Hyper-V
   <additionalparameters> -path <storepath> -Verbose -Force
   ```

   Additional parameters:

   - To import the same IP settings on the destination server that were on the source server, where <*SourcePhysicalAddress-1*> and <*SourcePhysicalAddress-2*> are comma-separated lists of the physical addresses of the source network adapter, and <*TargetPhysicalAddress-1*> and <*TargetPhysicalAddress-2*> are comma-separated lists of the physical addresses of the destination network adapter, include:

     ```
     -IPConfig All -SourcePhysicalAddress
     ```

```
"<SourcePhysicalAddress1>","<SourcePhysicalAddress2>" -
TargetPhysicalAddress
"<TargetPhysicalAddress1>","<TargetPhysicalAddress2>"
```

- To import the user groups that are used by the Hyper-V security policy and remote administration, include:

```
-User <Enabled | Disabled | All> -Group
```

6. If a failure occurred while running the **Import-SmigServerSetting** cmdlet, review the Setupact.log and Setuperr.log under %localappdata%\SvrMig\Log.
7. Use the information you gathered about physical-to-virtual network connections to establish the connections between the physical network adapters and external virtual switches on the destination server. (Virtual networks are now referred to as virtual switches in Windows Server 2012.) All external virtual networks are migrated to the destination server as internal virtual switches because the import process cannot map virtual switches to physical networks. To establish the connections:
   a. .a. Open Hyper-V Manager. (From the **Start** screen, click the **Hyper-V Manager** tile.)
   b. In the **Action** pane, click **Virtual Switch Manager**.
   c. In the left pane, under **Virtual Switches**, click the name of the first internal switch that you want to convert to an external switch.
   d. In the right pane, under **Connection type**, select **External**. From the drop-down list, select the physical network adapter to use for access to the physical network.
   e. Click **OK** to save the changes and close Virtual Switch Manager, or click **Apply** to save the changes and modify another virtual switch.
8. For each virtual machine that used a physical disk connected directly to the physical computer, establish this connection on the destination server:
   f. .a. Open the Disk Management snap-in and verify that the disk is in an Offline state. If the disk is not in an Offline state, it will not be available when configuring storage for a virtual machine.
   b. Open Hyper-V Manager. (From the **Start** screen, click the **Hyper-V Manager** tile.)
   c. Under **Virtual Machines**, select the virtual machine that you want to connect to the physical disk.
   d. In the **Action** pane, under the virtual machine name, click **Settings**.
   e. In the navigation pane (left pane), click the controller that you want to attach the disk to. If you plan to use the disk as a startup disk, make sure you attach it to an IDE controller. Click **Add**.
   f. On the **Hard Drive** page, select the location on the controller to attach the disk.
   g. Under **Media**, specify the physical hard disk. If the disk does not appear in the drop-down list under **Physical hard disks**, make sure the disk is in an Offline state in Disk Management.
   h. Select **Physical hard disk**, and then click **OK**.
9. Restore any customizations you made to the WMI namespace security settings. For more information, see [WMI namespace security customizations are missing after upgrading to

Windows Server 2008 R2(http://go.microsoft.com/fwlink/?LinkId=178143).

10. If you turned off any exceptions in Windows Firewall on the source server, turn off those same exceptions on the destination server.
11. If you have granted remote access to the server for Hyper-V management to any user account that is not a member of the Administrators group or Distributed COM Users group, open the registry and navigate to the HKLM\Software\Microsoft\OLE\ key. Add the MachineLaunchRestriction value that you recorded from the source server.

    Caution
    Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

Hyper-V: Verify the Migration

After you have completed the migration, perform the following verification steps to ensure that the migration succeeded.

Verify the Hyper-V security policy

Use the Hyper-V security policy information (that you gathered when you prepared for migration) to verify that the roles and scopes are the same on the destination server as the corresponding roles and scopes on the source server.

Verify the networking configuration

Use the information you gathered about the virtual and physical networks to verify the following:

- The virtual networks are the same on the destination server as they were on the source server.
- If you migrated IP settings for the physical network adapters, such as static IP addresses, they are applied to the corresponding network adapter on the destination server.

Verify the configuration and availability of the virtual machines

Perform the following steps to determine whether the migrated virtual machines will operate as expected.

To verify the virtual machines

1. Use the virtual machine information you gathered when you prepared for migration to verify the following:
    a. Check to see that the set of virtual machines on the destination server has all the virtual machines that were on the source server.
    b. For each virtual machine, verify that the state of the virtual machine on the destination server is the same as it was on the source server before the migration.
    c. For each virtual machine, verify that the snapshots it has are identical in number and structure to the snapshots of the corresponding virtual machine on the source server.
    d. Verify that the memory and number of virtual processors are the same as they were on the source server.
    e. Verify that the storage configuration (virtual hard disks and/or physical disks attached directly to the virtual machine) is identical to that on the source server.
2. Start each migrated virtual machine. If a virtual machine does not start, check the event log under Applications and Service Logs\Microsoft\Windows\Hyper-V virtual machineMS\Admin to see why it failed to start. Common reasons for failure include:
    - The virtual machine is not in the correct scope in the authorization policy.
    - The storage is misconfigured. For example, one or more virtual hard disks might not be in the specified location. Check the hard disk settings for the virtual machine to make sure that the path to the .vhd file is valid. If the virtual machine is configured to use a directly attached physical disk, make sure it is attached to the destination server and shows as offline in Disk Management on the server.
    - One or more virtual hard disks do not have the required security permissions in the file system where the .vhd files are stored.
3. Run some basic operations that change the state of each virtual machine to verify that the operations work as expected on the migrated virtual machine. For example, saving and restoring, pausing and resuming, starting and stopping, or taking and applying or deleting snapshots.
4. Delete any of the snapshots you have taken as part of the previous step, turn off the virtual machine to merge the snapshot disks, and then turn on the virtual machine.
5. After the virtual machine has booted into the operating system, run the necessary application-specific tests to ensure that the application on the virtual machine can provide the same service levels as it provided before the virtual machine was migrated.
6. Verify that you can access the desktop of each virtual machine using Remote Desktop or Virtual Machine Connection, if you had access to the desktop on the source server.
7. If the virtual machine passes all of the above tests, it is ready to be put into production.

Hyper-V: Post-migration Tasks

After you have performed the verification steps, you are ready to complete the migration. Completing the migration for Hyper-V consists of either retiring the source server if the migration succeeded, or rolling back the source server to its pre-migration state if the migration failed.

Retiring your source server

If the migration succeeded, you can repurpose the server for another use or retain it as a backup.

Important
We recommend that you remove the Hyper-V role as soon as you verify that the migration succeeded, to avoid unintentionally placing the source server back online, which could result in running duplicate virtual machines on the same network.

Restoring the role in the event of migration failure

If verification was not successful, follow these steps to roll back the migration.

Roll back migration of Hyper-V on the source server

To roll back migration of Hyper-V on the source server

1. Disconnect the destination server from the network.
2. If you removed the Hyper-V role from the source server, add the Hyper-V role.
3. Reconnect the source server to the network.
4. Restart all the virtual machines.

Roll back migration of Hyper-V on the destination server running Windows Server 2012

To roll back migration of Hyper-V on the destination server

1. Delete the migrated virtual machines.

   Important
   Do not delete the migrated data if you plan to retry the migration to the destination server. This will allow you to save time by not having to copy the data files again.

2. Remove the Hyper-V role.

Roll back migration changes on other computers in the enterprise

For each client that depends on workloads running on virtual machines on the source server, verify that the clients can communicate with the virtual machines.

Troubleshooting cmdlet-based migration

The Windows Server Migration Tools deployment log file is located at %*windir*%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations.

- %*windir*%\Logs\ServerMigration.log
- On Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012: %*localappdata*%\SvrMig\Log
- On Windows Server 2003: %*userprofile*%\Local Settings\Application Data\SvrMig\Log

If migration log files cannot be created in the preceding locations, **ServerMigration.log** and **SmigDeploy.log** are created in %*temp*%, and other logs are created in %*windir*%\System32.

If a migration cmdlet fails, and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at
0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050).  Minidump
attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on

the migration source server.

Viewing the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects, and query them for more information about settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The Windows Server Migration Tools **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or share that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**,

Property name	Type	Definition
		Share, and **Encrypted File**.
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item, shown as a path name.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to both **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.

Property name	Type	Definition
Title	String	Title of the result object. Examples of values include NIC physical address for IP configuration, or user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   $*VariableName* **=** $(*Cmdlet*)

   The following is an example.

   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -User all -Group -Path D:\rmt\DemoStore -force -Verbose)

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.

2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

   $*VariableName*

   In the following example, the variable is named **ImportResult**.

   $*ImportResult*

   This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable.

```
            ItemType  ID                                Success
DetailsList
            --------  --                                -------
-----------
          OSSetting  Local User                          True
{Local User, Loc...
          OSSetting  Local Group                         True
{Local Group, Lo...
```

```
        WindowsFeature  DHCP                                    True
{}
```

Each line of the preceding sample is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by examples in step 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   $<VariableName>| Select-Object -ExpandProperty <PropertyName>

   The following is an example.

   $importResult | Select-Object -ExpandProperty DetailsList

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples.

   - The following command returns only those details of result objects that have the ID **Local User**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Severity -eq "Warning") {$_} }

   - The following command returns only the details of result objects with an ID of **Local User** that also have the title **Remote Desktop Users**.

     $ImportResult | Where-Object { $_.ID -eq "Local Group" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Title -eq "Remote DesktopUsers") {$_} }

More information about querying results

For more information about the cmdlets that are used in the preceding examples, see the following additional resources.

- Where-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134853).
- Select-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134858).
- ForEach-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134860)

For more information about Windows PowerShell scripting techniques, see What Can I Do With Windows PowerShell? - Scripting Techniques on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134862).

Migrate IP Configuration to Windows Server 2012

Migration of IP configuration data is a necessity for the migration of some server roles to Windows Server 2012, including DHCP Server, Domain Name System (DNS) Server, and Active Directory Domain Services (AD DS). This guide describes how to migrate core IPv4 and IPv6 configuration settings and data.

Supported operating systems

The following table indicates the Windows Server operating systems that are supported by Windows Server Migration Tools.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2008 R2, both full and Server Core installation options	x64-based
x86- or x64-based	Full installation option of Windows Server 2008	Windows Server 2008 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2, both full and Server Core installation options	Windows Server 2008 R2, both full and Server Core installation options	x64-based
x64-based	Windows Server 2012, both full and Server Core installation options	Windows Server 2012, both full and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers.

Migrations between physical operating systems and virtual operating systems are supported.

Windows Server Migration Tools does not support migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server® 2012 in the German system UI language.

Note

The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Both x86- and x64-based migrations are supported for Windows Server 2003, Windows Server 2008 R2 and Windows Server 2012. All editions of Windows Server 2008 R2, and Windows Server 2012 are x64-based.

Roles that are running on Server Core installations of Windows Server 2008 cannot be migrated, because there is no .NET Framework available on Server Core installations of Windows Server 2008.

Supported scenarios and features

Windows Server Migration Tools supports migration of the following frequently used IP configuration settings and data. Settings are migrated in the order in which they are listed in the Windows interface. For example, DNS server settings are migrated in the order in which they are used.

Setting type	Supported settings and notes
Manually-configured IP settings for all enabled network adapters (also known as network interface cards, or NICs) that are connected to the network	IPv4 addresses
	IPv4 subnet mask
	IPv4 DHCP status
	IPv4 default gateway addresses (but not gateway metrics)
	IPv4 interface metric
	IPv4 Windows Internet Name Service (WINS) server settings

Setting type	Supported settings and notes
	• WINS server addresses • NetBIOS setting (Default, Enable NetBIOS over TCP/IP, or Disable NetBIOS over TCP/IP)
	IPv6 addresses and corresponding subnet prefix lengths **Note** Migration of IPv6 subnet prefix lengths is supported only on destination servers that are running Windows Server 2008, Windows Server 2008 R2 or Windows Server® 2012. If an IPv6 address is imported from a Windows Server 2003 migration store, the subnet prefix length is set to the default value of 64 on the destination server.
	IPv6 router discovery setting
	IPv6 **Managed address configuration** flag, and **Other stateful configuration** flag **Notes** These settings are supported only on Windows Server 2008, Windows Server 2008 R2 and Windows Server® 2012. Windows Server 2003 does not support DHCPv6-based IPv6 address configuration. If **Router Discovery** is enabled on the source server, you must ensure that the **Advertising** setting is the same on both source and destination servers to make sure that **Managed address configuration** and **Other stateful configuration** values are configured the same on both source and destination servers. For example, if **Managed address configuration** is configured automatically on the source server,

Setting type	Supported settings and notes
	then the value for the **Advertising** setting must be the same on both source and destination servers for **Managed address configuration** to be configured automatically on the destination server. **Managed address configuration** and **Other stateful configuration** settings are not imported to the destination server if they are configured as **Automatic**. In other words, they are not imported if the **Router Discovery** setting is enabled, yet the **Advertising** setting is disabled.
	IPv6 default gateway addresses (but not gateway metrics)
	IPv6 interface metric **Note** IPv6 interface metric is supported only on Windows Server 2008, Windows Server 2008 R2 and Windows Server® 2012. If this setting is manually configured on Windows Server 2003 or Windows Server 2003 R2, it will not be migrated. For more information about how to migrate this setting manually from Windows Server 2003 or Windows Server 2003 R2, see **IP Configuration: Appendix**.
	DNS settings • IPv4 DNS server addresses • IPv6 DNS server addresses • **DNS suffix for this connection** • **Register this connection's addresses in DNS** • **Use this connection's DNS suffix in DNS registration**
For global (Windows-based) IP configuration	• For resolution of unqualified names:

Setting type	Supported settings and notes
	• **Append primary and connection-specific DNS suffixes** • **Append parent suffixes of the primary DNS suffixes** • **Append these DNS suffixes and the list of DNS suffixes** (also known as the DNS search list)
	Enable LMHOSTS lookup (but not LMHOSTS file)
	IPv6 **DisabledComponents** property **Note** This setting is supported only on Windows Server 2008, Windows Server 2008 R2 and Windows Server® 2012. If this setting is configured in Windows Server 2003 or Windows Server 2003 R2, it will not be migrated.

Scenarios and features that are not supported

Group Policy settings or other autoconfigured settings related to IP configuration are not supported.

If the source server uses additional settings for advanced IP configuration that are not in the previous list, define a custom migration procedure based on the configuration of your organization's network environment.

For more information about IP configuration settings that are not supported, see IP Configuration: Appendix.

See Also

IP Configuration: Prepare to Migrate

IP Configuration: Migrate IP Configuration Data

IP Configuration: Post-migration Tasks

IP Configuration: Appendix

IP Configuration: Prepare to Migrate

This topic helps you prepare to migrate IP configuration settings and data.

Impact on the source server

To prevent IP address conflicts when the source server has static IP addresses that you want to migrate, you must do one of the following after you export configuration from the source server, but before you import IP configuration data from the migration store to the destination server.

- Disconnect the source server from the network.
- Change static IP addresses on the source server.

Important
Changing the source server's IP address can cause roles that are running on the source server to fail.

Impact on the destination server

The destination server has only an intermittent connection to the network from the start of the migration data importation process until importation is complete. During migration, IP configuration settings that are migrated from the source server overwrite IP configuration settings on the destination server.

Important
If you migrate the static IP address from the source server to the destination server, changing the IP address can cause roles that are running on the destination server to fail.

Impact on other servers in your enterprise

Other servers in the enterprise might be affected during IP configuration migration if they depend on server roles or features that are running on the source server.

Impact on other client computers in your enterprise

Client computers cannot access either the source or destination servers during the import process. If the servers from which you are migrating IP configuration data are configured as routers, computers that are configured to use the router will not be able to connect to some networks.

Expected downtime during IP configuration migration

The following conditions can be expected during IP configuration migration.

- Users cannot access the source server while it is disconnected from the network.
- After the import operation starts, users cannot access the destination server until the import operation is fully completed.
- The destination server must be restarted for changes to the IPv6 **DisabledComponents** property to take effect.

User rights required to perform migration on both source and destination servers

Local Administrator rights are required on both the source and destination servers to perform IP configuration migration.

Preparing the destination server

Prepare the destination server for IP configuration migration by using the following steps.

To prepare the destination server

1. Install Windows Server Migration Tools on the destination server. For more information, see Install, Use, and Remove Windows Server Migration Tools.
2. Verify that all network adapters that you want to configure are enabled and connected to the network.
3. If you choose to import global IP configuration settings, verify that you can restart the server after the import operation is completed.

Preparing the source server

Perform the following steps to prepare the source server for IP configuration migration.

To prepare the source server

1. Install Windows Server Migration Tools on the destination server. For more information, see Install, Use, and Remove Windows Server Migration Tools.
2. Verify that all network adapters that have configurations that you want to migrate are enabled and connected to the network.

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination

servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

Preparing other computers in the enterprise

Perform the following steps to prepare the source server for IP configuration migration.

To prepare other computers

- Notify users that they will be unable to access either source or destination servers after the import operation has started, and will not have access until IP configuration migration is complete.

See Also

Migrate IP Configuration to Windows Server 2012

IP Configuration: Migrate IP Configuration Data

IP Configuration: Post-migration Tasks

IP Configuration: Appendix

IP Configuration: Migrate IP Configuration Data

After you have prepared for IP configuration migration by performing steps in IP Configuration: Prepare to Migrate, migrate IP configuration settings and data by using procedures in this section.

Migrating Global and NIC IP configuration

Perform steps in this section to migrate IP configuration data.

IP configuration migration tools

Windows PowerShell cmdlets that are used for data and share migration include **Export-SmigServerSetting** (used on the source server), and **Import-SmigServerSetting** (used on the destination server).

The **Export-SmigServerSetting** cmdlet lets the user copy all supported IP configuration settings from the source server to migration store at a specified location.

On the destination server, the **Import-SmigServerSetting** cmdlet applies the IP configuration settings specified in the migration store to the destination computer. To import the IP configuration settings of network adapters, you must provide the mapping between the source and destination network adapters by listing the physical addresses (also called MAC addresses) of all network adapters.

Note
To import and export IP configuration settings for a network adapter, it must be enabled and connected to a network.

For more information, see Help for Windows PowerShell cmdlets. To view Help for a cmdlet, in a Windows PowerShell session, type **Get-Help** <*cmdlet_name*>, and then press **Enter**.

Note
Windows Server Migration Tools must be installed on the computer on which you want to view Help for the migration cmdlets.

Migrating IP configuration by using Windows Server Migration Tools

Use the two procedures in this section to migrate global and network adapter-specific configuration settings by using Windows Server Migration Tools.

Export IP configuration settings from the source server

If you have already exported IP configuration settings from your source server as part of another migration guide, go to the procedure Import IP configuration settings to the destination server.

To export IP configuration settings from the source server

1. Do one of the following.
   - To open a Windows Server Migration Tools custom Windows PowerShell session on computers that are running Windows Server® 2012 go to **Start** and then click **Windows Server Migration Tools**.
   - To open a Windows Server Migration Tools custom Windows PowerShell session on computers that are running Windows Server 2008 R2 or Windows Server 2008 click **Start**, point to **Administrative Tools**, open the **Windows Server Migration Tools** folder, right-click **Windows Server Migration Tools**, and then click **Run as administrator**.
   - To open a Windows Server Migration Tools custom Windows PowerShell session on computers that are running Windows Server® 2012 go to **Start** and then click **Windows Server Migration Tools**.
   - To open a Windows Server Migration Tools custom Windows PowerShell session on computers that are running Windows Server 2003, click **Start**, point to **Administrative Tools**, open the **Windows Server Migration Tools** folder, and then click **Windows Server Migration Tools**.

2. In the same Windows PowerShell session, run the **Export-SmigServerSetting** cmdlet on the source server by typing the following command, in which *MigrationStorePath* represents the path of your migration store location, and then pressing **Enter**.

```
Export-SmigServerSetting -IPConfig -Path <MigrationStorePath>
-Verbose
```

Notes

Because network connectivity might be interrupted during the import operation, be sure to verify that the migration store is created on the destination computer.

You are prompted to provide a password to encrypt the migration store. Remember this password, because you must provide the same password to import settings from the migration store.

3. Because it contains information that you must have to perform the import operation, save the output of the **Ipconfig -all** command. Type the following, and then press **Enter**, in which *FileName* represents the path of the location in which you want to save the output text file, and the file name.

```
IPConfig -all > <FileName>
```

4. If the source server has a static IP address, disconnect the source server, or change the static IP address.

Import IP configuration settings to the destination server

To import IP configuration settings to the destination server

1. For network adapter IP configuration migration, map physical (MAC) addresses for both source and destination network adapters. View the IPConfig output you generated by using the **IPConfig -all > <FilePath>** command in To export IP configuration settings from the source server to determine network adapter physical address mapping.
2. If the migration store is not already on the destination server, copy the migration store to a local drive on the destination server by typing the following and then pressing **Enter**, in which *NetworkPath* is the path of the location of the migration store, and *LocalPath* is the path of a location on the destination server.

```
Copy <NetworkPath> <LocalPath>
```

3. Log on to the destination server as a member of the Administrators group, if you have not already done so.
4. On the destination server, migrate all IP configuration by using the **Import-SmigServerSetting** cmdlet as shown in the following example, in which each *SourcePhysicalAddress* and *TargetPhysicalAddress* in quotation marks represents the physical address of a network adapter that you want to migrate, and *MigrationStorePath* represents the path of the location of your migration store. Specify each network adapter physical address in the format *AA-AA-AA-AA-AA-AA*, and separate the physical addresses of multiple network adapters by using commas.

```
Import-SmigServerSetting -IPConfig All -SourcePhysicalAddress
```

```
"<SourcePhysicalAddress1>","<SourcePhysicalAddress2>" -
TargetPhysicalAddress
"<TargetPhysicalAddress1>","<TargetPhysicalAddress2>" -Path
<MigrationStorePath> -Verbose
```

You can use one of the following values with the **-IPConfig** parameter. For **All** or **Global** IP configuration migration, the destination server must be restarted for modifications to the IPv6 **DisabledComponents** property to take effect. You cannot use any of the Windows Server Migration Tools cmdlets until the server has restarted.

- **Global:** only import global Windows IP configuration settings.
- **NIC:** only import specific IP configuration settings for certain network adapters. You must specify the physical address mapping by using the **-SourcePhysicalAddress** and **-TargetPhysicalAddress** parameters.
- **All:** import both global and network adapter IP configuration-specific settings. You must also specify the physical address mapping by using the **-SourcePhysicalAddress** and **-TargetPhysicalAddress** parameters.

For the list of supported settings for network adapters, see IP Configuration: Prepare to Migrate.

5. You are prompted to provide the same password that was provided during the export process to decrypt the migration store. Type the password, and then press **Enter**.

See Also

Migrate IP Configuration to Windows Server 2012

IP Configuration: Prepare to Migrate

IP Configuration: Post-migration Tasks

IP Configuration: Appendix

IP Configuration: Post-migration Tasks

After you have migrated IP configuration settings and data as directed in IP Configuration: Migrate IP Configuration Data, verify the migration and, if necessary, roll back IP configuration migration by using the procedure in this section.

Verifying the migration

Perform the following steps to verify your IP configuration migration.

To verify the IP configuration migration

1. Open a Command Prompt window on the destination server. To do this, click **Start**, click **Run**, type **cmd**, and then press **Enter**.

2. Verify that all IP configurations that you wanted to migrate exist on the correct network adapters on the destination server. To do this, type the following, , and then press **Enter**.

   ```
   IPConfig -all
   ```

3. Compare the results of the **IPConfig -all** command with the IPConfig output you generated on the source server in the procedure “To export IP configuration settings from the source server” in IP Configuration: Migrate IP Configuration Data.
4. For static IP address migration, verify that you can access the destination server by using the same IP address as the source server had before the migration. You can verify this by using the **ping** command in a Windows Command Prompt session.

Rolling back migration

If necessary, perform the following steps to roll back IP configuration migration.

To roll back migration

1. If you obtained a different static IP address for your source server, and migrated the statically-configured IP address to the destination server, either disconnect the destination server from the network or obtain a new static IP address for the destination server.
2. Set the IP address of the source server back to the pre-migration static IP address.
3. Connect the source server back to the network if you disconnected it in step 1.

Troubleshooting cmdlet-based migration

The Windows Server Migration Tools deployment log file is located at %*windir*%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations.

- %*windir*%\Logs\ServerMigration.log
- On Windows Server® 2012, Windows Server 2008 and Windows Server 2008 R2: %*localappdata*%\SvrMig\Log
- On Windows Server 2003: %*userprofile*%\Local Settings\Application Data\SvrMig\Log

If migration log files cannot be created in the preceding locations, **ServerMigration.log** and **SmigDeploy.log** are created in %*temp*%, and other logs are created in %*windir*%\System32.

If a migration cmdlet fails, and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at
0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050).  Minidump
attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares

that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Viewing the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects and query them for more information about settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The Windows Server Migration Tools **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise,

Property name	Type	Definition
		False is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or share that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item, shown as a path name.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to both **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values

Property name	Type	Definition
		include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.
Title	String	Title of the result object. Examples of values include NIC physical address for IP configuration, or user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and after migration is complete, use the variable in a query to return the content of result objects.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   ```
   $VariableName = $(Cmdlet)
   ```

   The following is an example.

   ```
   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -
   User all -Group -Path D:\rmt\DemoStore -force -Verbose)
   ```

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.

2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

   ```
   $VariableName
   ```

   In the following example, the variable is named **ImportResult**.

```
$ImportResult
```

This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable.

```
            ItemType  ID                                  Success
DetailsList
            --------  --                                  -------
-----------
          OSSetting  Local User                            True
{Local User, Loc...
          OSSetting  Local Group                           True
{Local Group, Lo...
     WindowsFeature  DHCP                                  True
{}
```

Each line of the preceding sample is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by examples in step 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   ```
   $<VariableName>| Select-Object -ExpandProperty <PropertyName>
   ```

   The following is an example.

   ```
   $importResult | Select-Object -ExpandProperty DetailsList
   ```

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples.

   - The following command returns only those details of result objects that have the ID **Local User**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local User" } |
     Select-Object -ExpandProperty DetailsList
     ```

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local User" } |
     Select-Object -ExpandProperty DetailsList | ForEach-Object {
     if ($_.Severity -eq "Warning") {$_} }
     ```

   - The following command returns only the details of result objects with an ID of **Local Group** that also have the title **Remote Desktop Users**.

     ```
     $ImportResult | Where-Object { $_.ID -eq "Local Group" } |
     ```

```
Select-Object -ExpandProperty DetailsList | ForEach-Object {
if ($_.Title -eq "Remote DesktopUsers") {$_} }
```

See Also

Migrate IP Configuration to Windows Server 2012

IP Configuration: Prepare to Migrate

IP Configuration: Migrate IP Configuration Data

IP Configuration: Appendix

IP Configuration: Appendix

This Appendix contains additional information that might help you prepare for a custom IP configuration migration when your migration is not supported according to the list of supported migration scenarios and features in IP Configuration: Prepare to Migrate.

Migrating manually-configured IPv6 interface metrics from Windows Server 2003

If you have manually-configured IPV6 interface metrics on a computer that is running Windows Server 2003, manually migrate the interface metrics to Windows Server® 2012 by using **netsh** commands.

To obtain index numbers for source and destination network adapters

1. On source servers that are running either Windows Server 2003 or Windows Server 2003 R2, open a Command Prompt session by clicking **Start**, clicking **Run**, typing **cmd** in the **Open** box, and then either clicking **OK** or pressing **Enter**.
2. On the destination server, open a Command Prompt session with elevated user rights. To do this, click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.
3. On both source and destination servers, obtain the index numbers of source and destination network adapters. Type the following in each Command Prompt session, and then press **Enter**.

   ```
   netsh interface ipv6 show interface
   ```
4. Record the numbers in the **Index** column that correspond to the names of the interfaces that you want to migrate.

To obtain the manually-configured IPv6 metric from the source server

1. If a Command Prompt window is not already open on the source server, open one as directed in Sstep 1 of To obtain index numbers for source and destination network

adapters.

2. Type the following, in which *Index* represents the index numbers that you obtained in To obtain index numbers for source and destination network adapters, and then press **Enter**.

   ```
   netsh interface ipv6 show interface <Index>
   ```

   For example, if your interface has an index number of 11, use the following command.

   ```
   netsh interface ipv6 show interface 11
   ```

3. Record the IPv6 metric that you want to migrate to the destination server in the **Metric** field.

To migrate the manually-configured IPv6 metric to the destination server

- If a Command Prompt window is not already open on the destination server, open one as directed in step 1 of To obtain index numbers for source and destination network adapters.
- Type the following, in which *Index* represents the number that you obtained in To obtain index numbers for source and destination network adapters, and *Integer* represents the number that you obtained in To obtain the manually-configured IPv6 metric from the source server, and then press **Enter**.

  ```
  netsh interface ipv6 set interface <Index> metric=<Integer>
  ```

  For example, for an interface with an index of 22, use the following command to set the metric to **2**.

  ```
  netsh interface ipv6 set interface 22 metric=2
  ```

Additional resources

Some advanced IPv4 and IPv6 configuration settings for a network adapter are not displayed in the Windows interface. Depending upon the configuration of the destination network, you might need to migrate these settings manually to a destination server.

- If you have manually configured nondefault routes for an interface, use the following command to view these settings on the source server. The value for *ipvx* can be either IP (for IPv4) or IPv6.

  ```
  netsh interface <ipvx> show route
  ```

  To set these settings on the destination server, see **netsh** help by entering the following command, in which the value for *ipvx* can be either IPv4 or IPv6.

  ```
  netsh interface <ipvx> add route
  ```

- If you have manually configured general interface settings for a network adapter, use the following **netsh** command to view these settings on the source server. The value for *ipvx* can be either IP (for Ipv4) or IPv6.

  ```
  netsh interface <ipvx> show interface <InterfaceIndex>
  ```

  To set these configuration settings on the destination server, view **netsh** help by entering the following command, in which the value for *ipvx* can be either IPv4 or IPv6.

```
netsh interface <ipvx> set interface
```

For the complete list of all settings that can be viewed and configured by using the **netsh** command, see the following articles on the Microsoft Web site.

- Netsh commands for Interface IP (http://technet.microsoft.com/en-us/library/cc738592.aspx)
- Netsh commands for Interface IPv6 (http://technet.microsoft.com/en-us/library/cc740203.aspx)

Additional general TCP/IP configuration parameters are stored in registry keys. For the complete list of general TCP/IP configuration settings that are stored in registry keys, see TCP/IP Configuration Parameters on the Microsoft Web site (http://technet.microsoft.com/en-us/library/cc739819.aspx).

For additional information about IP configuration, the following resources are recommended.

- TCP/IP Fundamentals for Microsoft Windows (http://www.microsoft.com/downloads/details.aspx?FamilyID=c76296fd-61c9-4079-a0bb-582bca4a846f&displaylang=en)
- Understanding IPv6, Second Edition (http://www.microsoft.com/MSPress/books/11607.aspx)

See Also

Migrate IP Configuration to Windows Server 2012

IP Configuration: Prepare to Migrate

IP Configuration: Migrate IP Configuration Data

IP Configuration: Post-migration Tasks

Migrate Network Policy Server to Windows Server 2012

This document provides guidance for migrating the Network Policy Server (NPS) or Internet Authentication Server (IAS) role service from an x86-based or x64-based server running Windows Server 2003, Windows Server® 2008, Windows Server® 2008 R2, or Windows Server® 2012 to a new Windows Server® 2012 server.

About this guide

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic by clicking the stars in the upper-right corner of the page (1=poor, 5=excellent), and then add comments that support your rating. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional

suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

NPS migration documentation and tools ease the migration of NPS role service settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the IAS/NPS migration process, reduce migration time, increase the accuracy of the IAS/NPS migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process.

Target audience

This guide is intended for the following IT professionals:

- IT architects responsible for computer management and security throughout an organization.
- IT operations engineers who are responsible for the day-to-day management and troubleshooting of networks, servers, client computers, operating systems, or applications.
- IT operations managers who are accountable for network and server management.

What this guide does not provide

This guide does not provide detailed steps to migrate the configuration of other services that might be running on the source server.

Guidance is not provided for scenarios in which the new operating system is installed on existing server hardware by using the upgrade option during setup.

Supported migration scenarios

This guide provides the instructions for migrating an existing server that is running NPS or IAS to a server that is running Windows Server 2012. This guide does not contain instructions for Network Policy Server migration when the source server is running multiple roles. If your server is running multiple roles, it is recommended that you design a custom migration procedure specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available on the Windows Server 2008 R2 TechCenter (http://go.microsoft.com/fwlink/?LinkID=128554).

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

Supported operating systems

The following table displays the minimum operating system requirements that are supported by this guide.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 SP2	Windows Server 2012	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012	x64-based
x86- or x64-based	Windows Server® 2008	Windows Server 2012	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012	x64-based
x64-based	Windows Server 2012	Windows Server 2012	x64-based

- The NPS role service is not available in Server Core editions. Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported as either source or destination servers. Windows Server Foundation edition is not available for Windows Server 2003.
- Migration from a source server to a destination server that is running an operating system with a different installed language is not supported. For example, migration of server roles from a computer that is running Windows Server 2008 with a system language of French to a computer that is running Windows Server 2012 with a system language of German is not supported. The system language is the language of the localized installation package that was used to set up the Windows operating system.
- Both x86-based and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 are x64-based.

Supported NPS role configurations

Migration of the following NPS settings are supported by this guide:

1. **Policies**. Migration of NPS policy configuration, including connection request policies, network policies, and health policies is supported by using this guide.
2. **Authentication methods**. All supported authentication method settings can be migrated using this guide. For more information about authentication methods, see NPS Authentication Methods (http://go.microsoft.com/fwlink/?LinkId=169629).
3. **System Health Validators (SHVs)**. Migration of SHV configuration settings implemented using Microsoft published SDK are supported.
4. **NPS templates**. Template settings are migrated using NPS UI export and import functionality. You cannot migrate template settings using the command line.
5. **RADIUS clients and remote RADIUS servers**. RADIUS clients and remote RADIUS server configuration settings, including shared secrets can be migrated using this guide.

6. **SQL accounting**. The configuration of SQL parameters, including connection, description, accounting, authentication, periodic accounting status, periodic authentication status, and max sessions settings can be migrated using this guide. It is recommended to manually configure SQL connection string settings. For more information, see Configure SQL Server Logging in NPS (http://go.microsoft.com/fwlink/?LinkId=169631).

IP address and host name configuration

This guide supports the following scenarios:

1. The destination server is configured with the same host name or IP address as source server.
2. The destination server is configured with a different host name or IP address than the source server.

Migration scenarios that are not supported

The following migration scenarios are not covered in this document:

- **Upgrade**. Guidance is not provided for scenarios in which the new operating system is installed on existing server hardware by using the **Upgrade** option during setup.
- **Extension DLLs**. This guide does not support migration of extension DLL registry key settings. For more information about extension DLL registry key migration, see Setting Up the Extension DLLs (http://go.microsoft.com/fwlink/?LinkId=169632).
- **Non-Microsoft authentication methods**. The migration of settings for non-Microsoft authentication methods is not supported. To migrate these settings, refer to your vendor documentation.
- **Non-Microsoft SHVs**. The migration of settings for non-Microsoft SHVs is supported only if the SHV is developed using guidance from the NAP SHA/SHV SDK. To migrate these settings, refer to your vendor documentation.

Overview of migration process for this role

Network Policy Server (NPS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy in Windows Server 2012. NPS is the replacement for Internet Authentication Service (IAS) in Windows Server 2003.

The current topic provides an overview of the NPS migration process. The NPS migration guide also includes the following major sections:

- NPS Server Migration: Preparing to Migrate
- NPS Server Migration: Migrating the NPS Server
- NPS Server Migration: Verifying the Migration
- NPS Server Migration: Post-migration Tasks

The pre-migration process involves establishing a storage location for migration data, collection of information that will be used to perform the server migration, and operating system installation on the destination server. The NPS migration process includes using the **iasmigreader** tool if the source server is running Windows Server 2003. If the source server is running Windows

Server 2008 or Windows Server 2008 R2, the Network Shell (netsh) utility is used to obtain NPS settings. When migrating a source server running Windows Server 2012, you can use netsh or Windows PowerShell®. Procedures are then performed on the destination server to install the required roles and migrate NPS settings. Verification procedures include testing the destination server to ensure it works correctly. Post-migration procedures include retiring or repurposing the source server.

Process diagram

The following diagram provides an overview of the migration process.

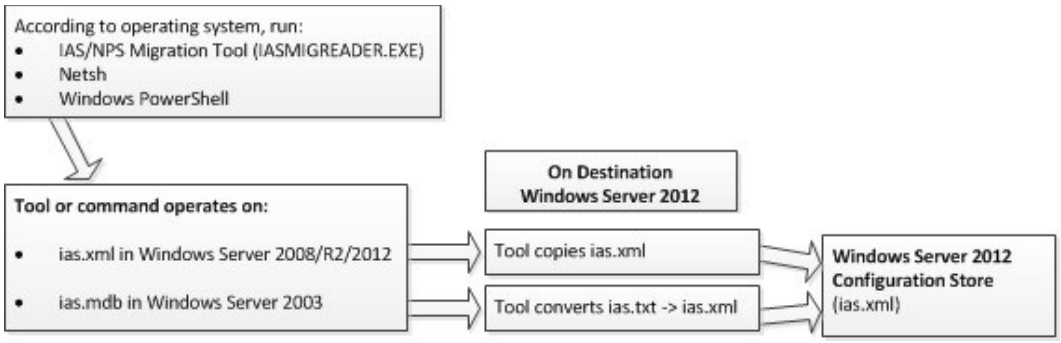

Figure 1. NPS server migration overview

Impact of migration

In its recommended configuration, the destination server has the same host name and IP address as the source server. In this scenario, the source server will be unavailable to process network access requests for the duration of the migration process (estimated 1-2 hours).

This guide also includes procedures for migration of the NPS server configuration from the source server to a destination server with a different host name or IP address. This allows the source and destination NPS servers to run simultaneously until all testing and verification is complete, and reduces service disruption. If you change the name or IP address of the server running NPS, RADIUS clients must also be updated with the new NPS server name and IP address.

Impact of migration on the source server

- When deploying the destination server with the same host name and IP address as the source server, the source server must be decommissioned and taken offline prior to renaming the destination server from *tempNPS* to the host name of the source server.
- When deploying the destination server with a different host name and IP address, there is no impact to the source server.

Impact of migration on other computers in the enterprise

- When deploying the destination server with the same host name and IP address, network access requests cannot be evaluated by NPS while the source server is offline and before the destination server brought online with the same name and IP address. During this time, client computers requesting access to the network cannot authenticate and are denied network access.
- When deploying the destination server with a different host name and IP address, RADIUS client settings for all network access servers that are configured to use the source server must be updated.

Permissions required to complete migration

The following permissions are required on the source server and the destination server:

- Membership in the **Administrators** group, or the equivalent, is the minimum required to install and configure server running NPS.
- Membership in the SQL database rights are required for SQL settings migration.
- If the destination server is a domain member, membership in the **Domain Admins** group, or the equivalent, is the minimum required to authorize the NPS server.

Estimated duration

The work required to migrate NPS settings from the source to destination server, including testing, can require 1 to 2 hours. Additional time may be required for migration of non-Microsoft authentication methods, SHVs or extension DLLs.

See Also

NPS Server Migration: Preparing to Migrate

NPS Server Migration: Migrating the NPS Server

NPS Server Migration: Verifying the Migration

NPS Server Migration: Post-migration Tasks

NPS Server Migration: Appendix A - Data Collection Worksheet

NPS Server Migration: Preparing to Migrate

Migration of Network Policy Server (NPS) includes the following tasks:

- Choose a migration file storage location
- Prepare your source server
- Prepare your destination server

Complete the steps or procedures in these sections to prepare your environment for migration.

If the server running NPS will be joined to a domain, membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. If the server running NPS is not domain joined, membership in the **Administrators** group, or equivalent, is required. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Choose a migration file storage location

First, choose a location where migration files will be kept.

To choose a storage location

1. Select a file storage location where migration files will be kept. The storage location can be a network share that is accessible by both the source and destination server, or portable media that can be transferred from one server to another.

Prepare your source server

Follow these steps to prepare an x86-based or x64-based server running Windows Server 2003, Windows Server® 2008, Windows Server® 2008 R2 or Windows Server® 2012 for NPS migration.

To prepare the source server

1. Determine the domain, server name, IP address, and passwords on the source server.
2. If the source server is domain joined, determine the group membership of the source server in Active Directory Domain Services (AD DS), including security group and OU membership. This can be done using the Active Directory Users and Computers console (dsa.msc) or Server Manager on a domain controller.

Prepare your destination server

Follow these steps to prepare an x64-based destination server running Windows Server 2012 for NPS migration.

Scenario 1: Prepare the destination server using the same host name and IP address

1. Install Windows Server 2012 on the destination server.
2. If the source server host name is used by RADIUS clients or remote RADIUS server groups, name the destination server with a temporary server name, for example: *TempNPS*.
3. If the source server IP address is used by RADIUS clients or remote RADIUS server groups, assign a different temporary static IP address to the destination server.
4. If the source server is domain joined, add the destination server to the domain of the source server. Configure AD DS group membership settings on the destination server

that are identical to the source server, including security group and OU membership.

5. Install the NPS role service using the steps provided in Install Network Policy Server (NPS) (http://go.microsoft.com/fwlink/?LinkId=169633).
6. If the source server has non-Microsoft authentication methods installed, then install same authentication methods on the destination server using your vendor documentation before importing the source server configuration.
7. If the source server has extension DLLs installed, install the same extension DLLs on the destination server before importing the source server configuration. For more information, see Setting Up the Extension DLLs (http://go.microsoft.com/fwlink/?LinkId=169632).
8. If the source server has non-Microsoft SHVs installed, then install same SHVs on the destination server using your vendor documentation before importing the source server configuration.

Scenario 2: Prepare the destination server using a different host name and IP address

1. Follow the same steps as provided for scenario 1, replacing the temporary server name with the new destination server host name, and assigning a permanent static IP address.

The destination server is now prepared for migration.

See Also

Migrate Network Policy Server to Windows Server 2012

NPS Server Migration: Migrating the NPS Server

NPS Server Migration: Verifying the Migration

NPS Server Migration: Post-migration Tasks

NPS Server Migration: Appendix A - Data Collection Worksheet

NPS Server Migration: Migrating the NPS Server

This topic contains steps and procedures for migrating the Network Policy Server (NPS) role service from a legacy source server to a new x64-based destination server running Windows Server 2012.

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Known issues

If you previously created conditional attributes for your remote access policy using **Called Station ID** and **Calling Station ID**, the comparison of these attributes in Windows Server 2012 now uses

a regular expression instead of matching the exact string. For a description of these attributes, see Remote Access Policy Conditions in the **IAS Authorization** section.

Exporting settings from the source server

Use the following procedures to export the NPS settings from your x86-based or x64-based source server prior to migrating to an x64-based server running Windows Server 2012. Follow the steps in the appropriate section based on the version of Windows Server that is running on the source server:

- Exporting settings from Windows Server 2003
- Exporting settings from Windows Server 2008
- Exporting settings from Windows Server 2008 R2
- Exporting settings from Windows Server 2012

Warning

When you use the following procedures to export configuration settings, apply appropriate precautions when moving these files from the source server to destination servers. NPS server configurations are not encrypted in the exported XML file, and contain shared secrets for RADIUS clients and members of remote RADIUS server groups. Therefore, sending these files over a network connection might pose a security risk. You can add the file to an encrypted, password protected archive file before moving the file to provide greater security. In addition, store the file in a secure location to prevent access by unauthorized users.

Exporting settings from Windows Server 2003

Configuration settings for Internet Authentication Service (IAS) in Windows Server 2003 are stored in **.MDB** files. Configuration settings for Network Policy Server (NPS) in Windows Server 2012 are stored in **.XML** files. **Iasmigreader.exe** is a command-line tool that exports the configuration settings of IAS on a computer running Windows Server 2003 to a text file. You can obtain the **iasmigreader.exe** command line migration tool for migrating Windows Server 2003 IAS settings to Windows Server 2012 from the following locations:

1. Windows Server 2012 installation media provides a copy of the migration tool in the **\sources\dlmanifests\microsoft-windows-iasserver-migplugin** directory.
2. The migration tool is available in the **%windir%\syswow64** directory on a server running Windows Server 2012.

To export settings from a source server running Windows Server 2003

1. Copy **iasmigreader.exe** to the source server into a directory configured in the **%path%** environment variable.

   Tip

   To review the source server's **%path%** configuration, type **echo %path%** at a

command prompt and press Enter.

2. At an elevated command prompt, type **iasmigreader.exe**, and then press Enter. The migration tool will automatically export settings to a text file.

   Important
   Configuration changes made to IAS will take at least one minute to be available for export.

3. IAS settings are stored in the file **ias.txt** located in the **%windir%\system32\ias** directory on the source server. If you are running a 64-bit version of Windows Server 2003, the **ias.txt** file is located in the **%windir%\syswow64\ias** directory.
4. You must manually copy SQL log configuration settings on the source server to a file (example: sql.txt).

   To record these settings:

   a. At an elevated command prompt, type **ias.msc**, and then press Enter.
   b. In the IAS console tree, click **Remote Access Logging**, right-click **SQL Server**, and then click **Properties**.
   c. Record the configuration settings on the **Settings** tab, and then click **Configure**.
   d. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

5. Copy the ias.txt and sql.txt files to the migration store file location.

Warning

Store the ias.txt and sql.txt files in a secure location. These files contain shared secret information and SQL connection strings.

Important

When you migrate the configuration settings of the IAS role service that is running on a 32-bit or a 64-bit Windows Server 2003–based source server to the NPS role service that is running on a Windows Server 2012–based destination server, the import procedure seems to complete successfully. However, the Extensible Authentication Protocol (EAP) method is misconfigured. This occurs because the migration tool generates a faulty parameter that is stored in the configuration text file (ias.txt). For more information about this issue and for a workaround, see The EAP method is configured incorrectly during the migration process from Windows Server 2003 32-bit or a 64-bit to Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=181982).

Exporting settings from Windows Server 2008

Configuration settings for NPS in Windows Server 2008 are stored in **.XML** files that can be directly imported to the destination server. The Network Shell (NetSh) command line utility can be

used to export and import these settings. You can also use the Windows interface to import and export these settings.

Warning

You cannot use the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

To export settings from a source server running Windows Server 2008 using a command line

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.
2. Confirm that a message appears indicating that the export to file was successful.
3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.
4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2008 using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, open **Roles\Network Policy and Access Services\NPS**.
3. Right click **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

7. Copy the ias.txt and sql.txt files to the migration store file location.

Exporting settings from Windows Server 2008 R2

Configuration settings for NPS in Windows Server 2008 R2 are stored in **.XML** files that can be directly imported to the destination server. The Network Shell (NetSh) command line utility can be used to export and import these settings. You can also use the Windows interface to import and export settings.

Warning

You cannot use the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

Important

The netsh utility does not support migration of template configuration settings. To migrate these settings, you must use the Windows interface.

To export settings from a source server running Windows Server 2008 R2 using a command line

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.

2. Confirm that a message appears indicating that the export to file was successful.

3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data

Collection Worksheet.

4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2008 R2 using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, open **Roles\Network Policy and Access Services\NPS**.
3. Right click **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. In the console tree, right-click **Templates Management** and then click **Export Templates to a file**.
7. Next to **File name**, type **iastemplates.xml**, navigate to the migration store file location, and then click **Save**.
8. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

   a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

   b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

   c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.
9. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

Exporting settings from Windows Server 2012

Configuration settings for NPS in Windows Server 2012 are stored in **.XML** files that can be directly imported to the destination server. You can use the following methods to export and import these settings:

1. The Network Shell (NetSh) command line utility
2. The Windows interface
3. Windows PowerShell cmdlets

Warning

You cannot use Windows PowerShell, the Windows interface or a command line to export or import detailed SQL configuration settings. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

Important
The netsh utility and Windows PowerShell do not support migration of template configuration settings. To migrate these settings, you must use the Windows interface.

To export settings from a source server running Windows Server 2012 using Windows PowerShell

1. On the source server, create a new folder for your settings (for example: C:\ConfigSettings).
2. Export your configuration settings to an .xml file in that folder, by following these steps.
   a. On the **Start** screen, type **PowerShell**, and then click Enter.
   b. To switch to the NPS context enter the following Windows PowerShell command and then press Enter:

      Import-Module NPS

   c. To export the configuration file to an .xml file, enter the following Windows PowerShell command, using the -path parameter to identify the name of the .xml file to be created and the folder into which it should be placed:

      Export-NpsConfiguration [-Path] <String>

      Tip
      For example:

      Export-NpsConfiguration –Path C:\ConfigSettings -Path nps01.xml

      Caution
      The exported file contains unencrypted shared secrets for RADIUS clients and members of remote RADIUS server groups. Because of this, you should ensure that the file is stored in a secure location to prevent malicious users from accessing the file.
3. Confirm that no errors were reported by Windows PowerShell.
4. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

   a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.
   b. Record the configuration settings on the **Settings** tab, and then click **Configure**.
   c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text

logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

5. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2012 using the Netsh utility

1. On the source NPS server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps export filename="path\file.xml" exportPSK=YES
   ```

   Replace *path* with the directory location where you want to save the source server configuration file, and replace *file* with the name of the .XML file that you want to save.
2. Confirm that a message appears indicating that the export to file was successful.
3. On the source server, type the following command and then press Enter:

   ```
   netsh nps show sqllog > path\sql.txt
   ```

   Replace *path* with the directory location where you want to save the source server SQL configuration file, and replace *sql* with the name of the .TXT file that you want to save. This file contains the basic configuration for SQL logging that is found on the **Settings** tab in SQL logging properties. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.
4. Copy the **file.xml** and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

To export settings from a source server running Windows Server 2012 using the Windows interface

1. On the source server, open Server Manager.
2. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
3. Right click the root node **NPS**, and then click **Export Configuration**.
4. In the dialog box that appears, select the check box next to **I am aware that I am exporting all shared secrets**, and then click **OK**.
5. Next to **File name**, type **file.xml**, navigate to the migration store file location, and then click **Save**.
6. In the console tree, right-click **Templates Management** and then click **Export Templates to a file**.
7. Next to **File name**, type **iastemplates.xml**, navigate to the migration store file location, and then click **Save**.
8. If you have configured SQL logging, you must manually record detailed SQL configuration settings.

   To record these settings:

a. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties**.

b. Record the configuration settings on the **Settings** tab, and then click **Configure**.

c. Manually record all configuration settings from the Connection and Advanced tabs by copying them into the **sql.txt** file. Alternatively, you can click the **All** tab and enter **Name** and **Value** settings displayed on each line into the **sql.txt** file. For a list of text logging and SQL configuration settings that you need to record manually, see NPS Server Migration: Appendix A - Data Collection Worksheet.

9. Copy the **file.xml**, **iastemplates.xml**, and **sql.txt** files to the migration store file location. This information will be required for configuration of the destination server.

Importing settings to the destination server

Use the following procedures to import the NPS settings from your x86-based or x64-based source server to an x64-based destination server running Windows Server 2012.

- Importing settings from Windows Server 2003
- Importing settings from Windows Server 2008 or Windows Server 2008 R2
- Importing settings from Windows Server 2012

Importing settings from Windows Server 2003

The configuration file **ias.txt** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012. If SQL accounting settings were saved, these settings are recorded manually in the **sql.txt** file.

Important

When you migrate the configuration settings of the IAS role service that is running on a 32-bit or a 64-bit Windows Server 2003–based source server to the NPS role service that is running on a Windows Server 2012–based destination server, the import procedure seems to complete successfully. However, the Extensible Authentication Protocol (EAP) method is misconfigured. This occurs because the migration tool generates a faulty parameter that is stored in the configuration text file (ias.txt). For more information about this issue and for a workaround, see The EAP method is configured incorrectly during the migration process from Windows Server 2003 32-bit or a 64-bit to Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=181982).

To import settings from a source server running Windows Server 2003

1. Copy the configuration file **ias.txt** that was exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, use either netsh or Windows PowerShell to import the configuration.

- To use netsh, do the following:

a. Open an elevated command prompt, type the following command and then press Enter:

```
netsh nps import filename="path\ias.txt"
```

Replace *path* with the directory where the **ias.txt** file is located. Verify that a message appears indicating that the import process was successful.

Tip

If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\ias.txt"**.

- To use Windows PowerShell, do the following:

a. On the **Start** screen, type **PowerShell**, and then click Enter.
b. Switch to the NPS context, enter the following Windows PowerShell command:

Import-Module NPS

c. To import the configuration, enter the following:

Import-NpsConfiguration [-Path] <String>

Replace *String* with the directory where the **ias.txt** file is located. Verify that a message appears indicating that the import process was successful.

Tip

For example:

Import-NpsConfiguration –Path c:\temp\ias.txt

3. If required, configure SQL accounting. To configure SQL accounting:
   a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
   b. Click **Accounting** and then click **Change SQL Server Logging Properties**.
   c. Manually enter SQL settings from the **sql.txt** file that you created.

Importing settings from Windows Server 2008 or Windows Server 2008 R2

The configuration file **file.xml** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012. SQL accounting settings are saved in the **sql.txt** file.

Note

For source servers running Windows Server 2008 R2: If you saved a templates configuration file, **iastemplates.xml**, you must use the Windows interface to import these settings.

To import settings from a source server running Windows Server 2008 or Windows Server 2008 R2

1. Copy the configuration files **file.xml** and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, use either netsh or Windows PowerShell to import the configuration.
   - To use netsh, do the following:

     a. Open an elevated command prompt, type the following command and then press Enter:

     ```
     netsh nps import filename="path\file.xml"
     ```

     Replace *path* with the directory where the **file.xml** file is located. Verify that a message appears indicating that the import process was successful.

     Tip

     If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\file.xml"**.

   - To use Windows PowerShell, do the following:

     a. On the **Start** screen, type **PowerShell**, and then click Enter.

     b. Switch to the NPS context, enter the following Windows PowerShell command:

     Import-Module NPS

c. To import the configuration, enter the following:

Import-NpsConfiguration [-Path] <String>

Replace <*String*> with the directory where the **file.xml** file is located.

Tip

For example:

Import-NpsConfiguration –Path c:\temp\file.xml

d. Confirm that no errors were reported by Windows PowerShell.

3. If required, configure SQL accounting. To configure SQL accounting:
   a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
   b. Click **Accounting** and then click **Change SQL Server Logging Properties**.
   c. Manually enter SQL settings from the **sql.txt** file.

Importing settings from Windows Server 2012

The configuration file **file.xml** that was exported from the source server is in a format that can be imported to a destination server running Windows Server 2012. SQL accounting settings are saved in the **sql.txt** file. If you saved a templates configuration file, **iastemplates.xml**, you must use the Windows interface to import these settings.

To import settings from a source server running Windows Server 2012

1. Copy the configuration files **file.xml** and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command.
2. On the destination server, open an elevated command prompt, type the following command and then press Enter:

   ```
   netsh nps import filename="path\file.xml"
   ```

   Replace *path* with the directory where the **file.xml** file is located. Verify that a message appears indicating that the import process was successful.

   Tip

   If the configuration file is located on a network share, provide full path to the file. For example: **netsh nps import filename = "\\fileserver1\Data\file.xml"**.

   The following Windows PowerShell command performs the same function:

   ```
   Import-NpsConfiguration –Path c:\temp\file.xml
   ```

3. If required, configure SQL accounting. To configure SQL accounting:
   a. In the Server Manager console tree, click **ALL SERVERS**, then from the list of

servers in the right pane, right-click the relevant server and select **Network Policy Server**.

b. Click **Accounting** and then click **Change SQL Server Logging Properties**.

c. Manually enter SQL settings from the **sql.txt** file.

Using the NPS console to migrate NPS settings

You can also use the Windows interface on the destination server to import configuration settings.

To import settings from a source server using the Windows interface

1. Copy the configuration files **file.xml**, **iastemplates.xml**, and **sql.txt** that were exported to the migration store file location to the destination NPS server. Alternatively you can import configuration settings directly from the migration store file location by supplying the appropriate path to the file in the import command. If you have custom settings that were recorded using the NPS Server Migration: Appendix A - Data Collection Worksheet, these must be configured manually on the destination server.
2. On the destination server, open Server Manager.
3. In the Server Manager console tree, click **ALL SERVERS**, and then from the list of servers in the right pane, right-click the relevant server and select **Network Policy Server**.
4. To import template configuration settings, follow steps 5 to 13. If you do not have template settings, skip to step 7.
5. In the console tree, right-click **Templates Management** and then click **Import Templates from a file**.
6. Select the template configuration file **iastemplates.xml** that you copied from the source server and then click **Open**.
7. In the console tree, right-click **NPS** and then click **Import Configuration**.
8. Select the configuration file **file.xml** or **ias.txt** that you copied from the source server and then click **Open**.
9. Verify that a message appears indicating the import was successful.
10. Configure SQL accounting if required using the **sql.txt** file and the data collection worksheet. To configure SQL accounting, follow steps 11 to 13.
11. In the NPS console tree, click **Accounting** and then click **Change SQL Server Logging Properties** in the details pane.
12. Modify the properties on the **Settings** tab if required, and then click **Configure** to enter detailed settings.
13. Using information recorded in the **sql.txt** file, enter the required settings on the **Connection** and **Advanced** tabs, and then click **OK**.

See Also

Migrate Network Policy Server to Windows Server 2012

NPS Server Migration: Preparing to Migrate
NPS Server Migration: Verifying the Migration
NPS Server Migration: Post-migration Tasks
NPS Server Migration: Appendix A - Data Collection Worksheet

NPS Server Migration: Verifying the Migration

After the migration of your Network Policy Server (NPS) server is complete, you can perform some tasks to verify that the migration was successful.

Verifying NPS Migration

To verify the functionality of NPS on the destination server, confirm that the service is running, that the correct configuration was migrated, and that client computers can authenticate successfully.

To verify NPS migration

1. To verify that the NPS service is running on the destination server, type the following command at an elevated command prompt on the destination server and then press ENTER.

   ```
   sc query ias
   ```

   In the command output, verify that **RUNNING** is displayed next to **STATE**.
2. To verify that the source NPS configuration has been migrated to the destination server, type the following command at an elevated command prompt on the destination server and then press ENTER:

   ```
   netsh nps show config
   ```

   Verify that the destination server is not using default NPS settings. For example, default settings display a single policy under **Connection request policy configuration** with the name **Use Windows authentication for all users**.
3. To verify that the NPS console on the destination server displays the correct settings, type the following command at an elevated command prompt on the destination server and then press ENTER:

   ```
   nps.msc
   ```

   a. The NPS console will open. In the console tree, click **Accounting**, click **Change SQL Server Logging Properties**, click **Configure**, and verify that the correct settings are displayed on the **Connection** and **Advanced** tabs.

b. In the NPS console tree, click **Policies** and then click **Connection Request Policies**, **Network Policies**, and **Health Policies**. For each type of policy, verify that the correct policies are displayed.

c. In the NPS console tree, click **RADIUS Clients and Servers** and then click **RADIUS Clients and Remote RADIUS Server Groups**. Verify that the correct RADIUS clients and remote RADIUS server groups are displayed.

d. In the NPS console tree, click **Network Access Protection**, and then click **System Health Validators** and **Remediation Server Groups**. Verify that the correct Network Access Protection (NAP) related settings are displayed.

e. In the NPS console tree, click **Templates Management**. If the source server was running Windows Server 2008 R2, verify that the correct templates settings are displayed.

f. In the NPS console tree, right-click **NPS**, click **Properties**, and then click the **Ports** tab. Verify that the correct **Authentication** and **Accounting** ports are displayed.

4. To verify the configuration of authentication methods, you must manually review settings in connection request policy and network policy. Certificate based EAP methods require that the proper certificate is chosen, and might require that you provision a computer certificate on the destination server.

Verifying authentication methods

a. If you use certificate based EAP methods, your destination server might already be provisioned with a suitable certificate through autoenrollment. You might also be required to manually enroll the destination server with a computer certificate. For an overview of certificate requirements for network authentication, see Network access authentication and certificates (http://go.microsoft.com/fwlink/?LinkId=169625).

b. To view certificates associated with EAP methods, click **Start**, click **Run**, type **nps.msc**, and press ENTER.

c. In the NPS console tree, open **Policies** and then open the type of policy you are using to perform authentication. For example, if the option to **Override network policy authentication settings** is enabled on the **Settings** tab in a connection request policy, then authentication is performed in connection request policy. Otherwise, authentication is performed in network policy. Authentication can be configured in both types of policies.

d. For connection request policy, double-click the policy name and then click the **Settings** tab. For network policy, double-click the policy name and then click the **Constraints** tab.

e. Click **Authentication Methods**, and then under **EAP Types** click the name of the certificate-based authentication method. For example: **Microsoft: Protected EAP (PEAP)** or Microsoft: Smart Card or other certificate.

f. Click **Edit**, verify that the correct certificate is chosen next to **Certificate issued** or **Certificate issued to**, and then click **OK**.

Note

Client computers using certificate based authentication methods must trust the certification path for this certificate.

5. To verify that client computers can authenticate using the destination server, attempt to connect to the network using client VPN connection, an 802.1X connection, or another connection that requires successful RADIUS authentication for network access.

Verifying client connections

a. To verify that client computers are successfully connecting to the network, click **Start**, click **Run**, type **eventvwr.msc**, and then press ENTER.
b. In the event viewer console tree, open **Custom Views\Server Roles\Network Policy and Access Services**.
c. In the details pane, verify under **Event ID** that event number 6272 is displayed.
d. Events 6273 or 6274 indicate that client authentication attempts are unsuccessful.
e. If no events are displayed, client connection requests are unable to reach the destination server, or the server is not logging authentication attempts.

See Also

Migrate Network Policy Server to Windows Server 2012

NPS Server Migration: Preparing to Migrate

NPS Server Migration: Migrating the NPS Server

NPS Server Migration: Post-migration Tasks

NPS Server Migration: Appendix A - Data Collection Worksheet

NPS Server Migration: Post-migration Tasks

After all migration steps are complete and you have verified the migration of the Network Policy Server (NPS) role service, perform the following post-migration tasks.

Post migration tasks

After verifying NPS configuration is working on the destination server, the following steps need to be performed:

To decommission a source server using the same host and IP address

1. Remove the source server from your Active Directory domain.

2. Shut down the source server.
3. Rename the destination server from *tempNPS* to the name of the source server and configure the same static IP address as that used by the source server.
4. Perform verification steps in NPS Server Migration: Verifying the Migration with the updated host name and IP address configured on the destination server.

To decommission a source server using a different host and IP address

1. NPS server name/ IP address should be updated on Remote RADIUS servers and RADIUS clients. It requires manual update of the configurations on RADIUS clients and Network Access Servers (NAS). Please refer to your RADIUS client configuration guide for more information.
2. Perform verification steps in NPS Server Migration: Verifying the Migration.
3. When the destination server has been configured, tested, and verified, then the NPS role on the source server may be retired.

Restoring the role in the event of migration failure

If the destination server is deployed simultaneously with the source server using a different host name and IP address, then the migration can be reversed by changing RADIUS clients, remote RADIUS server groups, and network access device settings to use the source NPS server name and IP address. If the destination server is replacing the source server using the same host name and IP address, then the destination server will need to be renamed, the IP address must be updated, and the destination server must be removed from the domain to reverse the migration and bring the source server back online.

See Also

Migrate Network Policy Server to Windows Server 2012

NPS Server Migration: Preparing to Migrate

NPS Server Migration: Migrating the NPS Server

NPS Server Migration: Verifying the Migration

NPS Server Migration: Appendix A - Data Collection Worksheet

NPS Server Migration: Appendix A - Data Collection Worksheet

Migration data collection worksheet

You can use this migration data collection worksheet to collect data about your source server and help ensure that the destination server functions properly after the migration.

NPS data worksheet

#	Source server essential settings	Setting values
1	**Server name** At a command prompt, type the following command, and then press ENTER. `ipconfig /all` The host name of a server is the first part of the fully qualified domain name (FQDN). The FQDN is the full computer name, including both the host name and the primary DNS suffix, separated by dots (.). For example, the FQDN of a computer named *host* with a primary DNS suffix of *example.microsoft.com* is *host.example.microsoft.com*.	Computer host name: _____ FQDN: _____
2	**Authentication, authorization, and accounting (AAA) roles** Determine what types of network access requests are validated using the RADIUS protocol on the source server.	Check all that apply (√) ☐ Network Access Protection (NAP) ☐ RADIUS server for dial-Up or VPN connections ☐ RADIUS server for 802.1X wireless or wired connections
3	**Text logging** Record the path and settings used for text logging. By default, local file accounting logs are stored in **%windir%\system32\LogFiles**.	Local file logging directory: _____ Format: _____ Create a new log file: _____
4	**SQL settings**	Application Name:

#	Source server essential settings	Setting values
	Manually record any customized SQL data link properties.	_____ Auto Translate: _____ Connect Timeout: _____ Current Language: _____ Data Source: _____ Extended Properties: _____ General Timeout: _____ Initial Catalog: _____ Initial File Name: _____ Integrated Security: _____ Locale Identifier: _____ Network Address: _____ Network Library: _____ Packet Size: _____ Password: _____ Persistent Security Info: _____ Replication server name connect option: _____ Tag with column collation when possible: _____ Use Encryption for Data: _____ Use Procedure for Prepare:

#	Source server essential settings	Setting values
		_____ User ID: _____ Workstation ID: _____

See Also

Migrate Network Policy Server to Windows Server 2012

NPS Server Migration: Preparing to Migrate

NPS Server Migration: Migrating the NPS Server

NPS Server Migration: Verifying the Migration

NPS Server Migration: Post-migration Tasks

Migrate Print and Document Services to Windows Server 2012

Overview

This article provides guidance to migrate a print server running Windows Server® 2003, Windows Server 2003 R2, Windows Server 2008, or Windows Server 2008 R2 operating systems to a server running the Windows Server 2012 operating system with the Print and Document Services role installed. This includes cross-architecture and stand-alone migrations, as well as configurations for a server in a cluster. This document provides step-by-step instructions for migrating from old hardware to new hardware and consolidating print servers.

Print and Document Services enables print server tools and configures the server to act as a print server. Print and Document Services is not dependent on any other features or roles. However, some specific network configurations, clients, and hardware may require you to install additional features or enable certain services.

This guide provides you with instructions for migrating an existing print server to a server that is running Windows Server 2012. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration procedure specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available in the Windows Server TechCenter.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

To manage the migration process, use one of the following:

- The Printer Migration Wizard, which you access through Print Management, a snap-in in Microsoft Management Console (MMC).
- The Printbrm.exe command-line tool.

You can perform the migration locally or remotely, and from either a client computer or server.

Important

As a best practice, run the Printer Migration Wizard or Printbrm.exe from a computer running Windows Server 2012 or Windows 8 to ensure that you are using the newest version of the migration tools that have the latest updates and features. You can run these tools either locally on the server or remotely from any other computer running Windows Server 2012 or Windows 8.

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

For more information about installing and using these tools, see Access the migration tools.

Notes

The Print Management snap-in is not available in Windows Server 2003. However, it is available in Windows Vista Enterprise and Windows Vista Ultimate, which enables you to migrate from Windows Server 2003. It is also available in Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012. For more information about migrating from Windows Server 2003, see Preparing to Migrate.

The Print Management snap-in and the Printbrm.exe command-line tool are not available for the Server Core installation option. To migrate from a print server running on a Server Core installation, use a server running the Printer Migration Wizard, Windows Vista Enterprise, or Windows Vista Ultimate. For more information about migrating from a server running a Server Core installation, see Preparing to Migrate.

There is no equivalent to Print and Document Services for Windows® client operating systems.

You can migrate Print and Document Services from the destination server or from any client with the following:

- The Printer Migration Wizard (provided that the client is running one of the supported operating systems listed in the Supported operating systems matrix).
- Remote access to the destination server.
- Access to the printer settings file created when you back up the source server.

Note

All commands in this guide are case-insensitive unless specifically noted.

About this guide

This guide is designed as a step-by-step tutorial for migrating print servers.

Target audience

This document is intended for IT administrators and other knowledge workers who are responsible for the operation and deployment of print servers in a managed environment.

What this guide does not provide

This document does not provide guidance for the following:

- Upgrading roles on the same computer
- Migrating printer configurations during client installations of Windows
- Migrating settings for a server that is not being used as a print server
- Recovering server information that was not properly saved prior to migration for in-place upgrades
- Instructions for migrating more than Print and Document Services

Supported migration scenarios

You must have access to the Printer Migration Wizard to migrate print servers. For more information about supported scenarios and limitations, see Preparing to Migrate.

Supported operating systems

The following table outlines the supported operating systems for migration covered in this guide.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2008, both full and Server Core	Windows Server 2012, including full,	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
	installation options	MinShell, and Server Core installation options	
x64-based	Windows Server 2008 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
X64-based	Windows Server 2012	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
X64-based	Server Core and MinShell installation options of Windows Server 2012	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

The Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported on full, Server Core, and MinShell installation options, as either source or destination servers.

All versions of Windows Server 2012 are x64-based. Migrating to an x86-based server is not supported.

Migrations between physical operating systems and virtual operating systems are supported.

Note
Both x86-based and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 are x64-based.

You might prefer the migration process, rather than an upgrade, even when the hardware is native x64-based. An example would be a case where there is increased use of the server and

there is a server role split (in which the source server has more than one server role)—and you decide to separate the roles onto several additional x64-based servers. In this case, migration of individual server roles to other servers might be the best solution.

The server administrator can choose which parts of an existing installation to migrate, but along with the server role, this usually includes configuration, data, system identity, and operating system settings.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French language to a computer that is running Windows Server 2012 in the German language.

Note

The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Cross-architecture migration (such as migrating from an x86-based server to an x64-based server) is supported. The source server must have print queue drivers installed for both the source and destination server architectures. If a print queue does not have a driver for the destination server architecture, then it will not be migrated. Similarly, verify that the destination server contains drivers for each supported architecture.

Supported role configurations

The Supported operating systems matrix provides a complete listing of the supported migration scenarios.

Some migration scenarios require additional preparation. For more information about these scenarios, see Appendix B - Additional Destination Server Scenarios.

Supported role services and features

The Printer Migration Wizard migrates:

- Print queues.
- Shared printer settings.
- Printer drivers in use by the print spooler.
- Any security settings specific to the installed printer.

Migrating from x86-based to x64-based v3 printer drivers

There are several things you must consider when managing migrations using v3 print drivers. The first is that a print queue cannot function without the native printer driver for the server architecture (x86 or x64) on which it exists. Since Windows Server 2012 is a 64-bit only operating system, it is important that you have 64-bit drivers installed for all of your printers if you are

migrating from a 32-bit system. The most difficult transition is from 32-bit to 64-bit servers in an organization with 32-bit clients since it is common to have third-party 32-bit printer drivers that do not have 64-bit equivalents. During the print configuration restoration for cross-platform scenarios, if the backup file does not contain driver binaries for the platform of the target server, the drivers will be installed from the target server's driver store, if available.

As a best practice, when migrating from x86-based to x64-based v3 drivers:

1. Verify that x64-based versions of the drivers are available.
2. If you are unable to verify their availability, back up the source server before the migration.
3. Install the x64-based drivers on the source server so that you can determine if there any problems or conflicts before the migration process.
4. If there are conflicts or problems on the destination server after the migration, roll back the migration. For more information, see "Roll back migration on the source server" in Post-Migration Tasks.

Unsupported scenarios

The Printer Migration Wizard does not migrate the following:

- Other services or settings that specific printers may rely on, such as Line Printer Remote (LPR), Internet Printer Protocol (IPP), or Web Services on Devices (WSD). You must enable or install these features before restoring the source print server configuration. For additional information, see "Roll back migration on the source server" in Post-Migration Tasks.
- Local bus printers (LPT and USB), although they are shown during backup. For additional information, see Appendix B - Additional Destination Server Scenarios.
- Plug and play printers. However, plug and play printer drivers will be migrated. For additional information, see Appendix B - Additional Destination Server Scenarios.
- Any print jobs currently in the printer queue.
- Any system or print administrators, or permissions. If you want to retain the same system or print administrators on the destination server as on the source server, you will need to manually add these administrators to the destination server.

Print and Document Services migration overview

While the original server is still running, use the Printer Migration Wizard or the Printbrm.exe command-line tool to export or back up the print information (such as settings, queues, and drivers) in a printer settings file. Then, import or restore this backup image to a destination server running Windows Server 2012 that has been configured to run as a print server.

Some migration scenarios require additional preparation. For more information about these scenarios, see Appendix B - Additional Destination Server Scenarios.

To migrate printers from a server running Windows Server 2003 or a Server Core installation to a server running Windows Server 2008 R2, you must use a computer running the Printer Migration Wizard to remotely manage the server running Windows Server 2003 or a Server Core installation. Using this computer, you can store the printer settings file (containing information

about the printers you want to migrate, such as settings, queues, and drivers) from the server running Windows Server 2003 or a Server Core installation to a file share. You can then use the Printer Migration Wizard to migrate the printers from the file share to the server running Windows Server 2008 R2.

For more information about accessing the Printer Migration Wizard, see Preparing to Migrate.

Note
The **Printing-Server Core** role must be installed on a server running a Server Core installation from which you want to migrate.

Migrate print servers (overview)

The following list provides an overview of the steps to migrating the print servers.

- Access the migration tools
- Prepare the destination server
- Prepare the source server
- Back up the source server
- Restoration
- Verify the migration
- Post-migration

Impact of migration

The objective of the migration process is that the destination server is able to perform the same functions as the source server did, without client computers on the network being aware that the migration has taken place. The following sections describe the impact of migration.

Impact of migration on the source server

The source server is not impacted by print server migration until the destination server takes over as the active server (typically when the name or IP address of the source server is assigned to the destination server). At that point, the source server no longer services print requests that target the print server.

Impact of migration on other computers in the enterprise

If the destination server replaces the source server in the network (same name or IP address), then there should be no impact to other computers in the enterprise.

If the destination server has a different name or IP address than the source server, then all clients with existing print connections must delete and recreate those print connections so that they target the destination server.

Permissions required to complete migration

Administrative permissions are required on both the source print server and the destination print server.

Permissions required to complete migration on other computers in the enterprise

If the destination server replaces the source server in the network, then no permissions are required on other computers in the enterprise. If the destination server has a different name or IP address, then the permissions required on other computers may vary depending on Group Policy settings, Windows Update access, and driver availability.

Estimated duration

The time required to migrate a print server will vary from server to server, depending on the following:

- The number of queues being migrated.
- The number of individual drivers needed for the queue.
- The size of a given driver, in terms of its file size and the number of files.
- The configuration of the server.

Migrating a single printer queue with a typical x86-based and x64-based driver can range from five seconds to several minutes, depending on the factors listed above. Because of this range, a typical migration can take anywhere from less than an hour to several hours.

See Also

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Preparing to Migrate

Access the migration tools

The Printer Migration Wizard and the Printbrm.exe command-line tool support all migrations to Windows Server 2012.

Note

Although the Printer Migration Wizard supports migrations from servers running Windows Server 2003 or a Server Core installation, it cannot run on these servers directly.

To access the Printer Migration Wizard

Open the Print Management snap-in on computers running operating systems previous to Windows 8

1. If necessary, enable the **Administrative Tools** menu, which is hidden by default on Windows-based client operating systems.
   a. Right-click **Start**, and then click **Properties**. The **Start Menu and Taskbar Properties** dialog box opens.
   b. On the **Start Menu** tab, click **Customize**. The **Customize Start Menu** dialog box opens.
   c. Under **System Administrative Tools**, select **Display on the All Programs menu** or **Display on the All Programs menu and the Start menu**.
2. In the **Administrative Tools** menu, click **Print Management**.

Tip

Alternatively, you can click **Start**, and type **printmanagement.msc** in the search text box.

Open the Print Management snap-in on computers running Windows Server 2012

- Open Server Manager, click **Tools**, and then click **Print Management**.

Open the Print Management snap-in on computers running Windows 8

- From the Start screen, type **printmanagement.msc** and click **printmanagement** when it appears in the search results.

Tip

If you run the Print Management Console remotely (for example, from a client computer), simply add the print server to your console (right-click **Print Servers**, click **Add/Remove Servers**) and continue the migration from there.

Start the Printer Migration wizard

- In Print Management, right-click **Print Management** and click **Migrate Printers**.

Important
The Print Management snap-in filter settings will not be migrated and need to be saved independently of the printer migration.

To access the Printbrm.exe command-line tool

1. To open a Command Prompt window, click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.

   To open a command Prompt window on a computer running Windows 8 or Windows Server 2012, right-click the start charm and click Command Prompt (Admin)
2. Type:

   ```
   %WINDIR%\System32\Spool\Tools\Printbrm.exe
   ```

To view the complete syntax for this command, at a command prompt, type:

```
Printbrm.exe /?
```

For a listing of the available syntax for the **Printbrm.exe** command, see Appendix A - Printbrm.exe Command-Line Tool Details.

Prepare the destination server

The second step in the migration process is to prepare the destination server.

Hardware requirements for the destination server

There are no specific hardware requirements for being a print server beyond those for the version of the server operating system you are using.

The amount of disk space needed for a migration is dependent on the number of print drivers to be migrated and the size of the drivers. Because print drivers can vary widely in size, the amount of disk space can range from one megabyte to several gigabytes.

Software requirements for the destination server

Verify that hard drive space is sufficient on the destination server for the backup.

No additional software is needed other than the necessary drivers required for the printers to be hosted. Migrate these drivers from the source server.

For cross-architecture migrations, verify that the destination server contains drivers for each supported architecture.

Installing the Print and Document Services role on the destination server

You must install the Print and Document Services role on the destination server before you begin the migration process. For more information on installing this and other server roles, see the Print and Document Services overview.

Preparing for cross-architecture migrations

If you are migrating from the x86-based architecture of Windows Server 2003 or Windows Server 2008 to the x64-based architecture of Windows Server 2012, you should install x64-based drivers on the source server before creating the backup file. The migration process copies all installed drivers from the source server to the destination server. It recreates the printer queues on the destination server if the printer settings file contains the x64-based drivers.

Verify that each print queue on the source server has a driver installed for the operating system on the destination server before creating the printer settings file. For example, if you are migrating an x86-based source print server to an x64-based destination print server, verify that each print queue has an x64-based driver installed before you create the printer settings file. Any print queue that does not have a cross-architecture driver installed will not be migrated to the destination server.

To install cross-architecture drivers for a printer, you can use:

- The **Add Printer Driver Wizard**, which is available in the Print Management snap-in.
- The **Printer Properties** dialog box, which is available through the Printers folder in the Control Panel.

As a best practice, you need to install a driver with the same name as the native architecture. To add the x86-based driver to the x64-based destination server, use the x86-based client to remotely open the x64-based server using Windows Explorer and navigate to the remote printer folder and add the driver. To install an x64-based driver on the x86-based source server, use the x64-based client to remotely open the x86-based server using Windows Explorer and navigate to the remote printer folder and add the driver.

Tip

In many cases, it can take you a long time to update all the print drivers for all your print queues. To save time, you may want to update just the most used print queues first, and gradually update the others when you have time. To save time, you can set the existing print queues to the **Generic/Text Only** driver for migration and later switch them to the OEM driver at your convenience. Most printers allow basic printing using the Generic Text driver. For more information, see Cross-Architecture print server migrations: Speeding up the migration process at the Microsoft TechNet Blogs web site.

Preparing for additional scenarios

In the following instances, installing a feature on your destination server may require additional preparation before you migrate to it:

- The server hosts Line Printer Remote (LPR) printers.
- The server offers Internet Printing Protocol (IPP) printer connections.
- The server hosts Web Services on Devices (WSD) printers.
- The server is in a server cluster.
- The server hosts plug and play printers.

For more information on these scenarios, see Appendix B - Additional Destination Server Scenarios.

Prepare the source server

Simple system-to-system migrations require no preparation for the source server. However, additional preparation is required for cross-architecture migrations. If performing the migration as quickly as possible is a priority, remove unused drivers, ports, and queues before starting the migration to improve its speed after verifying with users that the items to remove are no longer in use. In general, however, minimize changes to the source server environment to ensure you can roll back to this environment if necessary.

Caution

If your source server is running multiple roles, renaming the source server or changing its IP address can cause other roles that are running on the source server to fail.

Notes

You should delete native print drivers that are not currently associated with a print queue because these drivers increase the size of the printer settings file unnecessarily. The print spooler will not allow a native print driver that is currently associated with a print queue to be deleted.

The Print Spooler service will use non-native drivers. It routes these drivers to the Print Server service when a non-native client connects to a print queue and has to download a driver. You should remove any unused drivers and print queues.

Do not delete a non-native driver with a corresponding native print driver that is associated with a print queue. In this instance, the Print Spooler service will not prevent the non-native driver from being deleted. If the non-native driver's architecture matches the destination server's architecture, then you must block the driver's deletion. Cross-architecture drivers will never appear to be loaded by the Print Spooler service. Administrators should only delete them after confirming the drivers are no longer needed.

To install cross-architecture drivers using the Print Management snap-in on computers running Windows Vista and Windows Server 2008

1. Open the Print Management snap-in. Click **Start**, click **Administrative Tools**, and then click **Print Management**.
2. In the Print Management tree, under **Print Servers**, click the print server you want.
3. Under the print server, right-click **Drivers** and then select **Add Driver** to open the **Add**

Printer Driver Wizard.

4. Follow the steps as indicated by the wizard.

To install cross-architecture drivers by using only the Printer Properties dialog box on computers running Windows XP and Windows Server 2003

1. Click **Start**, click **Control Panel**, and double-click **Printers**.
2. Select **Printer**. Right-click **Sharing**.
3. Click **Additional Drivers** and select **Processor** from the list.
4. Follow the instructions in the dialog boxes to install the correct driver. Only install the driver associated with the printer you are administering.

Note

You can only add a cross-architecture driver if you have already installed a native architecture version of the same driver.

See Also

Migrate Print and Document Services to Windows Server 2012

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Migrating the Print and Document Services Role

Back up the source server

The fourth step in the migration process is to back up your source server data to a printer settings file using either the Printer Migration Wizard or the **Printbrm.exe** command-line tool in preparation for exporting printer queues, print drivers, and printer settings.

Important

As a best practice, run the Printer Migration Wizard or Printbrm.exe from a computer running Windows Server 2012 or Windows 8 to ensure that you are using the newest version of the migration tools that have the latest updates and features. You can run

these tools either locally on the server or remotely from any other computer running Windows Server 2012 or Windows 8.

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

To back up the source server using the Printer Migration Wizard

1. Open the Print Management snap-in.
2. Do one of the following:
   - In the **Print Management** window, right-click **Print Management**, and then click **Migrate Printers** to open the Printer Migration Wizard. Make sure that **Export printer queues and printer drivers to a file** is selected, and then click **Next**. In the **Select a print server** window, select the print server to be migrated, and then click **Next**.
   - In the **Print Management** tree, under **Print Servers**, right-click the print server that contains the printer queues to migrate, and then click **Export printers to a file** to open the Printer Migration Wizard.
3. Review the list of items to be exported, and then click **Next**.
4. In the **Export printer data to** box, enter the path to the printer settings file to use, or browse to the location where you want to store the file. Click **Next** to export the printer-specific information for the server to this file.
5. Verify that the printer settings file is stored on a resource that will be available to the destination server. Optimally, store it on a network share. Click **Finish**.

The backup file that you create by using either the Printer Migration Wizard or the Printbrm.exe tool inherits the permissions allowed by your user credentials. Only you can access the file if you saved the file directly to a share during the backup file creation process. You must either change the file permissions on the **Security** tab of the file's **Properties** dialog box, or you must perform any restorations or migrations by using that file yourself. If you create the backup file on the computer from which you are running the migration and later copy the file to a share, then file access permissions are inherited from the destination folder.

To back up the source server using the Printbrm.exe command-line tool

1. Open an administrator Command Prompt window.
2. Perform a remote print backup. To do this, type the following command in the **%WINDIR%\System32\Spool\Tools** folder at the command prompt, in which *Source Computer1* is the Universal Naming Convention (UNC) name of the source computer, and *Printer1 Settings* is the name of the printer settings file to back up.

   ```
   Printbrm -s \\<Source Computer1> -b -f <Printer1
   Settings>.printerExport
   ```

   Notes

   The Printer Migration Wizard and the Printbrm.exe command-line tool only

support a printer settings file that is created by the migration tool you are using. For example, .cab file backups that were created by using the Printer Migration Wizard are not supported. To view the complete syntax for the **Printbrm.exe** command, type **Printbrm.exe /?** in a Command Prompt session.

Only TCP/IP, WSD, and LPR ports will be migrated. The Printer Migration Wizard will not migrate printers attached through USB, LPT, or other local ports. For more information about these scenarios and migrating Plug and Play printers, see Appendix B - Additional Destination Server Scenarios.

Cross-architecture migrations

For cross-architecture migrations, verify that each print queue has a driver installed on the source server that is compatible with the operating system on the destination server before creating the printer settings file on the source server. For example, if you are migrating an x86-based source print server to an x64-based destination print server, verify that each print queue has an x64-based driver installed before you create the printer settings file. Any print queue that does not have a cross-architecture driver installed will not be migrated to the destination server.

Restoration

The fifth step in the migration process is to restore the printers to the destination server, using the printer settings file you created.

Before beginning the migration process, verify that you installed the Print and Document Services role on the destination server as part of your preparation.

To restore printers to the destination server using the Printer Migration Wizard

1. On the source server, stop the Print Spooler service for all printers so you can preserve all print jobs prior to the migration.
   a. Open Computer Management. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then click **Computer Management**.
   b. In the console tree, expand **Services and Applications**.
   c. In the console tree, under **Services and Applications**, click **Services**.
   d. In the details pane, do one of the following to stop the service:
      i. Right-click **Print Spooler** and select **Stop**.
      ii. Double-click **Print Spooler**. On the **General** tab, under **Service Status**, click **Stop**.
2. From the computer that is running the Printer Migration Wizard, on the **Administrative Tools** menu, click **Print Management**.
3. Do one of the following:
   - Right-click **Print Management**, and then click **Migrate Printers** to open the Printer Migration Wizard. Select **Import printer queues and printer drivers from a file,** and then click **Next**.

- In the **Print Management** tree, under **Print Servers**, right-click the destination print server, and then click **Import printers from a file** to open the Printer Migration Wizard.

4. Specify the printer settings file created in the **Back up the source server** section, and then click **Next**.
5. Review the list of items to be imported, and then click **Next**.
6. In the **Import Mode** list, indicate whether you want to keep or overwrite existing printers.

   If the printer settings file contains a printer already on the destination server, the printer is not restored, and the existing printer on the destination server is not changed.
7. In the **List in the directory** list, indicate which printers to list on the destination server.
8. Optionally, indicate whether you want to convert LPR ports to standard port monitors when you migrate.
9. Click **Next** to import the printers.
10. Click **Finish**.

    Note

    It is recommended that you review the Application events that have a PrintBRM source to determine whether any additional actions are needed. The restored printers are shared in the same manner in which they were shared previously.
11. To view details of the migration, click **Open Event Viewer**. For more information, see "Verify the Migration" in Verifying the Migration. If you identify Error 30 in the Event Viewer, see "Troubleshooting" and "Migrating cross-platform driver language monitors" in Post-Migration Tasks for instructions on resolving the error.

To restore printers to the destination server using the Printbrm.exe command-line tool

1. Open an administrator Command Prompt window.
2. Type the following command in the **%WINDIR%\System32\Spool\Tools** folder at the command prompt, in which *Source Computer1* is the UNC name of the source computer, and *Printer1 Settings* is the name of the printer settings file to restore.

```
Printbrm -s \\<Source Computer1> -r -f <Printer1
Settings>.printerExport
```

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Verifying the Migration

Verify the migration

The sixth step in the migration process is to verify that the migration was successful by testing and validating the new print server. The Printer Migration Wizard provides detailed logging of migration events in the Event Viewer.

To verify destination server configuration

1. View event log messages about the migration.
   - If you are managing Print Services migration from a remote client computer, you can view event messages in **Custom Views/Administrative Events** in Event Viewer on the Windows-based client computer.
   - If you are managing migration from the destination server and Print and Document Services is not yet installed, then migration-related events are logged in **Custom Views/Administrative Events** in Event Viewer on the destination server.
   - If you are managing migration from the destination server and Print and Document Services is installed, events are logged to a different location. See the events located at **Applications and Services Logs/Microsoft/PrintBRM/Admin**

     To view details after closing the Printer Migration Wizard, right-click the **Start** charm, and then click **Event Viewer**. In the Event Viewer pane, under **Custom Views**, click **Server Roles**, and then click **Print and Document Services**. In the center pane, click the printer migration event to view details.
2. To verify that each printer queue was migrated to the new server:
   - Manually check the destination server for each printer migrated from the source server.
   - Verify that the printer associated with each printer queue is online.

   To check each printer queue's online status in the Print Management snap-in, under **Print Servers**, click **Printers**. A list of all migrated printers appears in the center pane, listing the printer queue status for each printer. Clients will be unable to print to printers that were not restored successfully, and any connections to a missing printer queue will be invalidated.
3. Check the printer queue settings.
   - Confirm that a printer queue's special settings, permissions, or drivers were preserved during the migration.
   - Check the properties for each queue on the destination server and verify that any special settings are still applicable.
   - If the driver installs any non-standard settings that have been altered as a result of the migration, verify those as well.

   Note

The migration process only preserves printer queue permissions. Other permissions on the source server, such as system permissions (for example, user accounts) and custom permissions, are not migrated using this process.

4. Make any necessary changes, such as adding a port monitor or a new driver.

Rename the destination server to the name of the source server

Temporarily rename the destination server. For example, you can name the destination server the same as the source server with **_NEW** appended to the source server name. After verifying that printers are restored to the destination server, rename the source server (for example, by appending **_OLD** to the source server name), and then use the source server's pre-migration name as the new name of the destination server.

Important

Validating existing printer connections from client computers can only be completed after the destination server name is the same as the pre-migration source server name. Print connections to the server in the period between renaming the source server and renaming the destination server will fail. All migration steps should be complete on the destination server before renaming to ensure that the downtime occurs only between the renaming of the source server and the final renaming of the destination server.

If you are using Active Directory Domain Services (AD DS) to publish printers, do the following to ensure that AD DS does not contain multiple instances of the same printer.

When you restore printers to the destination server, do not publish printers to AD DS. This prevents duplicate printers from being displayed by AD DS before the destination server configuration is verified.

On the source server, you must unpublish printers before renaming the source server. To do this, select all printers in the Print Management snap-in, right-click the selected printers, and then click **Remove from Directory**. This prevents printers from being published twice to AD DS when the source server is renamed.

After renaming the destination server to the source server's original name, you can publish all printers on the destination server to AD DS. To do this, select all printers in the Print Management snap-in, right-click the selected printers, and then click **List in directory**.

To verify configuration of other computers in the enterprise

In most cases, a new print server will not affect other computers in the enterprise. Existing client connections may be corrupted if you make a change to any of the following print server properties:

- The print server name
- The printer name
- The print share name
- The share permissions
- The printer's availability to the server

Print a test page to each printer queue from a client (or set of clients) that had an existing connection to the source server to verify that other computers are not affected by the new print server. In a cross-architecture environment, test each supported architecture.

Print a test job from a client with an existing connection

From a client computer that is configured to print to the source print server, use the existing print queue to print a page to the new server. If you cannot print a test page:

- Determine if one or more of the print server properties listed above have been changed.
- Check whether the destination server is available to the client on the network. Create a new connection to a printer on the destination server to verify that the client and server are communicating normally.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Post-Migration Tasks

Post-migration

The final step in the migration process is determined by whether the migration was successful or unsuccessful.

No post-migration tasks are necessary beyond the standard migration process. If the server state (new settings, drivers, queues, and so on) changed, create and archive a new backup of the server state for recovery purposes.

Success

The next section provides tasks that should be completed after you have successfully migrated the source printer server.

Retire the source server

After taking the source server offline while backing it up, users are unable to print until the migration to the destination server is complete. To minimize the impact, leave the source server

in service while you complete the migration and testing of the destination server. By leaving it in service, you can also add both the source and destination servers to the Print Management snap-in to simplify verifying the restoration.

Once you have validated the installation, rename the source and destination servers and take the source server offline.

1. On the source server, restart the print spooler for all printers so it can finish spooling delayed print jobs. When it finishes, verify that there are no new print jobs.
2. Rename source and destination servers as directed in Verifying the Migration.
3. Follow your enterprise's normal policy for server decommissioning and retirement until retirement of the source server is complete.

Note

If the destination server has already been published in Active Directory Domain Services (AD DS), then the source server must be unpublished.

Failure

The next section provides tasks that should be completed if your migration of the source printer server did not succeede.

Restoring the role in the event of migration failure

Restoring the source server lets you deploy the settings to a new system or use the source server while determining the cause of the failure.

Rollback requirements

Warning

Rollback can only be completed if retirement of the source server has not been started. After you start retiring a source server—that is, you delete any print queues, close any print connections, reformat any drivers, or remove any hardware from the source server—you cannot roll back migration. After you start retiring the source server, the only method of rolling back migration is to restart the Print Services migration process from the beginning.

Estimated time to complete rollback

Rolling back migration involves renaming the source server to its pre-migration name, and renaming the destination server to either its original name, or another name that is not the same as the pre-migration name of the source server. Renaming the source and destination servers can be completed in a few minutes.

Roll back migration on the source server

Rename the source server to its original, pre-migration name. You might have to rename the destination server to a temporary name first.

Roll back migration on the destination server

Rename the destination server, either to its original name, or to another name that is not the same as the original name of the source server.

Troubleshooting

The following sections can help you troubleshoot any migration issues.

Log file locations

Printer migration events are included in the Application log, which is located at %*SystemRoot*%\System32\Winevt\Logs\Application.evtx and can be viewed using Event Viewer. A custom view for Printer Migration Events is available in Event Viewer.

Note

If the Printer Migration Wizard fails, you are directed to Event Viewer to view error messages. If you cannot find an error that explains the failure in Event Viewer, restore the backup by using the Printbrm.exe command-line tool. Error reporting from Printbrm.exe can often provide more detail than what is available in the event log.

Migrating cross-platform driver language monitors

When a cross-architecture migration includes the migration of printer language monitors, an error will occur during the process of restoring the printers to the destination server using the Backup Restore Migration (Printbrm.exe) command-line tool. The reason for the error is that language monitor driver architecture must always be the same as the source server architecture. Therefore, when migrating from x86-based architecture to x64-based platforms, language monitor migration cannot be successful. An error posted to the event log will state that the source architecture is not the same as that of the destination server.

You can recover from the printer restore error on the destination server by manually installing (or reinstalling) the appropriate standard driver for the migrated printer(s) running on that architecture.

Mitigating a failure in the Print Spooler service

If you encounter a failure in the Print Spooler service during print server migration, you can work around the failure. Using policy settings, you can isolate print drivers in separate processes so that print driver failures will not cause the Print Spooler service to fail—which allows the restoration to continue.

To turn on print driver isolation using Group Policy

1. Open the Group Policy Management Console. Right-click a Group Policy Object with the necessary scope, and then click **Edit**.
2. In the console tree under **Computer Configuration**, expand the **Administrative Templates** folder, and then expand the **Printers** folder.
3. Double-click **Execute print drivers in isolated processes**.
4. Click **Enabled**, and then click **OK**.
5. Double-click **Override print driver execution compatibility setting reported by print driver**.
6. Click **Enabled**, and then click **OK**.
7. At a command prompt, type **gpupdate /force** to reapply Group Policy settings.

Additional references

- Install, Deploy, and Migrate to Windows Server 2012
- Windows PowerShell Blog (http://go.microsoft.com/fwlink/?LinkId=128557)

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Appendix A - Printbrm.exe Command-Line Tool Details

Printbrm.exe command-line tool syntax

The following table lists the available printbrm parameters:

Syntax	Description
-s <server name>	Specifies the destination server.

Syntax	Description
-b	Backs up the server to the specified file.
-r	Restores the configuration in the file to the server.
-q	Queries the server or the backup file.
-f <file name>	Specifies the backup file.
-d <directory name>	Unpacks the backup file to the directory (with -r), or repacks a backup file from the directory (with -b).
-o force	Forces overwriting of existing objects.
-p all:org	Publishes all printers in the directory, or publishes printers that were published originally.
-nobin	Omits binary files from the backup.
-lpr2tcp	Converts LPR ports to standard TCP/IP ports on restore.
-c <file name>	Uses the specified configuration file.
-noacl	Removes ACLs from print queues on restore.
-?	Displays Help.

Printbrm enhancements

Printbrm.exe in Windows Server 2012 has several enhancements and improvements, including the following:

- **Supports both v3 and v4 print drivers**

  Windows Server 2012 supports both driver types, so with Printbrm you have full flexibility to back up, restore and configure the drivers you need to support your environment.

- **Supports backup CAB files greater than 2 GB**

  You should use the latest PrintBRM version when performing a migration or restoration. You will avoid CAB file size issues if you use the latest version of Printbrm to migrate and restore your Windows Server 2008 R2 (or previous) servers.

- **General improvements for reporting and error handling conditions during the backup and restore processes**

  These conditions are primarily logged in the Event Log under **Custom Views\Printer Migration Events**.

Printbrm usage scenarios

There are many ways Printprm can be used to make migrating your printers easier and more flexible.

Using the configuration file

You can use a configuration file to customize your printbrm migration for the following purposes:

- Replace printer drivers during a restore operation.

  For example, you might want to import your printers to a new print server, but use new v4 printer drivers.
- Backup / Restore dependent files from third-party Language Monitors
- Backup / Restore dependent files from third-party party Print Processors

For example, to replace your printer drivers, you can backup your printers using the `-nobin` parameter, and then restore the printers using the `-c <file name>` parameter to specify a configuration file with a DriverMap section.

To use a configuration file to specify updated print drivers

1. Backup you printers using the `-nobin` parameter. For example:

   ```
   Printbrm.exe -b -nobin -s \\myoldprintserver -f printers.printerExport
   ```
2. On the new print server computer, manually install the updated printer drivers.
3. Create a BrmConfig.xml configuration file to map the old drivers to the new drivers. For example:

   ```
   <BrmConfig>
   <PLUGINS>
   </PLUGINS>

   <LanguageMonitors>
   </LanguageMonitors>

   <DriverMap>
   <DRV old="OldDriverName1" new="NewDriverName1"/>
   <DRV old="OldDrverName2" new="NewDriverName2"/>
   </DriverMap>

   </BrmConfig>
   ```
4. Restore the printers specifying your configuration file using your configuration file. For

example: `PrintBrm.exe -r -c BrmConfig.xml -f printers.printerExport -o force`

5. Check your installed printers to verify they are installed with the updated printer drivers.

Important

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

Selectively restoring your printers

After you export the printers from your source server, you can selectively restore the printers and their related objects using the `-d` parameter. You can follow a general procedure to accomplish this:

To selectively restore your printers

1. Export the printer objects from the source server.
2. Restore to a temporary folder using the `-d` parameter.
3. Manually edit the files in the temporary folder.

   Note

   More information about the files created by a printbrm backup is described later.
4. Backup the temporary folder using the `-d` parameter.
5. Import the modified backup file to the target server.

A backup or export operation using PrintBRM produces a compressed file that is used for the restore or import operation. The following XML files are part of the export file in addition to the individual printer driver and configuration files:

- **BrmDrivers** – contains a list of every driver installed on the computer and the driver files for each driver
- **BrmForms** – contains a list of all the installed forms
- **BrmLMons** – will usually contain either Windows NT x86 or Windows x64 as the architecture and a list of port monitors and the port monitor files installed
- **BrmPorts** – contains a list of all the installed printer ports
- **BrmPrinters** – contains a list of all printers that have been installed
- **BrmSpoolerAttrib** – contains information about the spooler directory path and also indicates whether or not the source server was a cluster server

Moving printers to a different domain

If you move printers to a different domain, you will want to prevent the restoration of the print queue's ACLs. Use the `-NoACL` parameter to do this. If you use this parameter, the restored print queues will inherit the permissions of the target print server.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Appendix B - Additional Destination Server Scenarios

In some instances, your destination server may require additional preparation before you migrate to it.

If your server hosts Line Printer Remote (LPR) printers

To enable the hosting of LPR printers, install the LPR Port Monitor feature on the server:

1. Open Server Manager.
2. In the **Server Manager** dashboard, click **Add roles and features**. The **Add Roles and Features Wizard** opens.
3. Click **Next** on the **Before you begin** page.
4. Ensure **Role-based or feature-based installation** is select on the **Select installation type** page, and click **Next**.
5. Ensure your destination server is selected on the **Select destination server** page and click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select Features** page, click **LPR Port Monitor**, click **Next**, and then follow the instructions to complete the installation.

Important

The LPD and LPR Services are deprecated starting with Windows Server 2012. Eventually, they will be completely removed from the product, but they are still available in this release. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features. For more information about features or functionalities that have either been removed from the product in the current

release or are planned for potential removal in subsequent releases, see Features Removed or Deprecated in Windows Server 2012.

If your server offers Internet Printing Protocol (IPP) printer connections

To enable the Internet Printing Protocol (IPP):

- When installing the Print and Document Services role, select **Internet Printing**.

This automatically configures IIS and any other necessary features to support IPP printer hosting.

If your server hosts Web Services on Devices (WSD) printers

To enable WSD printing support:

1. Start the **Network and Sharing Center** from Control Panel and, click **Change advanced sharing settings** and click **Turn on network discovery**.
2. In Computer Management, Services, start the **Function Discovery Provider Host** service.
3. In Computer Management, Services, ensure the **Device Association Service** is started.

The server will be then able to identify and communicate with WSD-enabled printers.

If your print server is a highly available virtual machine

- To create a highly available print environment using Hyper-V and Failover Clustering, see High Availability Printing Overview. For more information about Windows Server 2012 Failover Clustering, see What's New in Failover Clustering in Windows Server 2012 and Clustering and High-Availability.
- Continue with the restoration process on the Printer Server virtual machine on the primary node.

If your server hosts local bus printers (LPT and USB)

The migration of local bus printers (LPT and USB) is not supported, although these printers are shown during backup. After the migration is complete:

1. Share the local bus printers again on the destination server.
2. Verify that each printer's name has not changed.
3. Test the printers to ensure that the shared connections still work.

If your server hosts plug and play printers

The migration of plug-and-play printers is not supported. To migrate plug and play printers:

1. Plug the printer into the destination server. The plug-and-play printer drivers will be installed automatically.
2. Enable printer sharing for the print queues.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix C - Printbrm Event IDs

Appendix C - Printbrm Event IDs

Printbrm Event IDs

The following Printbrm events are logged in the **Applications and Services Logs/Microsoft/PrintBRM/Admin** event log:

Event ID	Description
1	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is beginning a backup of print queues. No user action is required.
2	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is beginning a restore of print queues. No user action is required.
3	Printbrm.exe (the Printer Migration Wizard or the command-line tool) replaced driver map settings %1 with %2 for queue %3. No user action is required.
4	Printer queue %1 will be restored without a separator page. Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create a separator file for this queue. Error: %2.
5	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find a configuration file and is using the default settings. No user action is required.

Event ID	Description
6	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored print queue %1. No user action is required.
7	Printbrm.exe (the Printer Migration Wizard or the command-line tool) restored queue %1, but failed to restore printer driver settings. Error %2. This can occur if the driver on the destination server is newer than the driver in the migration file. Open the printer Properties dialog box on the destination computer and manually specify the appropriate printer settings.
8	While attempting to publish the printer to the Active Directory directory service, Printbrm.exe (the Printer Migration Wizard or the command-line tool) may have failed to adjust publishing settings for %1. This can occur if Printbrm.exe cannot access Active Directory. Manually publish the printer using the printer Properties dialog box.
9	Printer queue %1 already exists on the destination computer. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will update the printer settings to match the settings on the source computer. No user action is required.
10	Printer queue %1 already exists on the destination computer and will not be changed because Printbrm.exe (the Printer Migration Wizard or the command-line tool) was run in 'Keep Existing' mode. No user action is required.
11	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print queue %1 because port %2 is unknown. Printbrm.exe will attempt to restore the print queue on port FILE: instead. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are

Event ID	Description
	incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer and then change the print queue to use the new port.
12	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed printer queue %1 on port FILE:. No user action is required.
13	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore printer queue %1 on port FILE:. Error %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
14	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored %1. No user action is required.
15	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print processor %1 while restoring print queues from a file. Error: %2. Examine the Windows error returned by this event to determine the cause.
16	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not back up a dependent file for a language monitor. Dependent file: %1. Error: %2. This can occur if the file was deleted or moved.
17	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored language monitor %1. No user action is required.
18	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore language monitor %1 while restoring print queues from a file. Check the print queue

Event ID	Description
	backup file and examine the Windows error returned by this event to determine the cause.
19	The language monitors in the backup file are for a different processor architecture than the destination computer. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will not migrate any language monitors. Source architecture: %1. Destination architecture: %2.
20	Printbrm.exe (the Printer Migration Wizard or the command-line tool) restored a driver for a different processor architecture than that of the destination computer. Printbrm.exe will attempt to locate and install a native version of driver %1 on destination %2. Try to print to the print queue, and if necessary, manually install a native version of the driver.
21	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 for the processor architecture of the destination computer. No user action is required.
22	The driver in the backup file is for a different processor architecture than the destination computer, and Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not locate and install a native version of the driver. Driver: %1. Destination architecture: %2. Error: %3. Install a native version of the driver on the destination computer and then retry importing the print queues.
23	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from a cabinet (.cab) file. No user action is required.
24	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) because it was backed up on the source computer without its binary files. Printbrm.exe could not install the driver from the

Event ID	Description
	local driver cabinet (.cab) file on the destination computer. Error: %3. This can occur if the user did not save the driver binary files while backing up the print queue, or when restoring a print queue on a destination computer that uses a different processor architecture than the source computer. Install the driver manually on the destination computer or back up the print queue with its binary files. Then, retry importing the print queues.
25	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from files. No user action is required.
26	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) from files. Error reported: %3. This can occur if the driver requires a file that Printbrm.exe did not back up or if the user does not have permission to install drivers on the destination computer.
27	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to remove driver temporary folder %1. Error %2. Manually delete the temporary files and folder.
28	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to copy %1 to %2. Error %3 occurred while it was restoring print queues from a file. This can occur if the user does not have proper permissions on the destination computer, if the backup file is corrupted, or if the system is unstable. Retry exporting the printer queues on the source computer.
29	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not convert Line Printer Remote (LPR) port %1 to a standard TCP/IP printer port. Cannot get device settings. Error %2. This can occur if the LPR port was

Event ID	Description
	incorrectly configured. Manually create a standard TCP/IP printer port for the printer.
30	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) printer port %1 because the length of the port name is too long. len >= MAX_PORTNAME_LEN: Error %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
31	Port %1 already exists. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will skip restoring this port. No user action is required.
32	Printbrm.exe (the Printer Migration Wizard or the command-line tool) port install status: %1. No user action is required.
33	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is backing up printer forms. No user action is required.
34	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is restoring printer forms. No user action is required.
35	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully saved %1 user forms. No user action is required.
36	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find any user forms to restore. No user action is required.
37	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore one or more printer forms from the backup file. This can occur if the forms already exist on the destination computer or if the user does not have permissions to create forms on the destination computer.

Event ID	Description
38	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not load .xml file %1 while restoring print queues from a file. Error: %2. This can occur if the user does not have proper permissions on the destination computer, if the backup file is corrupted, or if the system is unstable. Retry exporting the printer queues on the source computer.
39	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find the printer driver migration plug-in DLL %1 for print processor %2. Printbrm.exe will skip this file and attempt to migrate the driver, but the printer might not work. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
40	Printbrm.exe (the Printer Migration Wizard or the command-line tool) changed the spooler folder to: %1. No user action is required.
41	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change spooler folder to %1. This path does not exist, so Printbrm.exe used the default spooler folder location (%WINDIR%\System32\Spool\Printers). This can occur if the configuration of the source computer is different from the destination computer. No user action is required unless you want to use a custom spooler folder location.
42	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change the location of the spooler folder. Error %1. Printbrm.exe used the default spooler folder location (%WINDIR%\System32\Spool\Printers). This

Event ID	Description
	can occur if the configuration of the source computer is different from the destination computer. No user action is required unless you want to use a custom spooler folder location.
43	Printbrm.exe (the Printer Migration Wizard or the command-line tool) changed the spooler log level to: %1. No user action is required.
44	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change the spooler log level. Error reported: %1. No user action is required unless you want to change the default spooler log level.
45	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is backing up printer objects on server: %1. No user action is required.
46	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to place a file in cabinet (CAB) file %1 while backing up print queues. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
47	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create destination file %1 while backing up print queues. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
48	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not load printer driver migration plug-in DLL %1. Error: %2. Printbrm.exe will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer

Event ID	Description
	and install a newer version of the driver, if necessary.
49	The printer driver migration plug-in %1 is incompatible with this version of Windows. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
50	The printer driver migration plug-in %1 ran with error: %2. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
51	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully ran printer driver migration plug-in %1. No user action is required.
52	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore print processor %1 because it is already installed on the destination computer. No user action is required.
53	While Printbrm.exe (the Printer Migration Wizard or the command-line tool) was restoring a print queue from backup, the following service control function %1 on computer %2 returned error %3. This can occur if Printbrm.exe cannot locate all files for a print processor, cannot stop the print spooler or the Cluster service, or if either of these services become unresponsive during installation of the print processor. If

Event ID	Description
	printing fails, manually install the print processor on the destination computer and examine the Print Spooler service and/or the Cluster service.
54	The service %1 is now running on server %2. No user action is required.
55	The service %1 is now stopped on server %2. No user action is required.
56	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to stop service %1 on computer %2. Error %3. Try printing to the affected printer to determine if there is a problem.
57	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to start service %1 on computer %2. Error %3. Try printing to the affected printer to determine if there is a problem.
58	Printbrm.exe (the Printer Migration Wizard or the command-line tool) did not overwrite print processor file %1 because the file on the destination system has an identical (or newer) time stamp. No user action is required.
59	Printbrm.exe (the Printer Migration Wizard or the command-line tool) copied print processor file %1 to %2. No user action is required.
60	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to copy print processor file %1 to %2. Error %3. This can occur if Printbrm.exe cannot locate all files for a print queue, or if the Print Spooler service cannot be stopped during installation of the print processor. If printing fails, manually install the print processor on the destination computer.
61	The device %1 %2 supports more than one

Event ID	Description
	port. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will convert the first port on the device from a Line Printer Remote (LPR) port to a standard TCP/IP port. The port on the device might need to be reconfigured.
62	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) port %1 because the LPR port monitor is not installed on the destination computer. Install the LPR port monitor on the destination computer.
63	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) port %1 while restoring print queues from a previously created backup file. Error code: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
64	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the standard TCP/IP printer port %1. Error: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
65	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore printer driver settings (%1, %2). Error %3. This can occur if the driver on the destination server is newer than the driver in the migration file. Open the printer Properties dialog box on the destination computer and manually specify the appropriate printer settings.
66	Printbrm.exe (the Printer Migration Wizard or

Event ID	Description
	the command-line tool) cannot apply new settings for printer queue %1 because port %2 that is saved in the cabinet (.cab) file does not exist on the destination server. This can occur if Printbrm.exe failed to migrate the port. Manually install the appropriate port, and then specify the appropriate printer settings or retry importing the print queue.
67	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not open registry key %1. Error code %2. This can occur if an important Windows resource (such as the registry) is unavailable, if the Component Object Model (COM) cannot be initialized, or if Printbrm.exe cannot allocate memory. Examine the Windows error returned by this event to determine the cause.
68	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not create a cabinet (CAB) file at this location: %1. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
69	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the standard TCP/IP printer port %2 because the length of the host IP address %1 is too long to be restored on operating systems older than Windows Vista.
70	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to save separator file %1 because the file or path does not exist. This can occur when a separator page cannot be found during the backup process on the source computer.
71	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not access the print$ share on %1. Error reported: %2. This

Event ID	Description
	can occur if the destination computer does not have any shared printers.
72	Printbrm.exe (the Printer Migration Wizard or the command-line tool) was unable to access the Remote Registry service on %1. Error reported: %2. This can occur if the Remote Registry service is not started or if the computer is behind a firewall.
73	Printbrm.exe (the Printer Migration Wizard or the command-line tool) reset the availability information for the print queue because the StartTime or UntilTime values were invalid. Error %1. The printer will be always available until you use the Advanced tab of the printer's Properties dialog box to specify the correct availability.
74	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from a driver package. No user action is required.
75	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) from a driver package. Error: %3.
76	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to backup driver %1 (%2). The backup process will continue, skipping this driver. Error: %3.
77	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to backup port %1. The backup process will continue, skipping this port. Error: %2.
78	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print processor %1 (%2). Error: %3.
79	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print processor files for architecture %1. Error: %2.

Event ID	Description
80	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print queue %1. The backup process will continue, skipping this queue. Error: %2.
81	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print queue %1. The restore process will continue, skipping this queue. Error: %2.
82	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create a new separator file folder while restoring print queue %1. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources. The print queue was not restored successfully. Error: %2.
83	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the WSD printer port described by the remote URL or global identifier %1. Error: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.

See Also

Migrate Print and Document Services to Windows Server 2012

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Migrate Remote Access to Windows Server 2012

Routing and Remote Access Service (RRAS) was a role service in Windows Server operating systems prior to Windows Server 2012 that enabled you to use a computer as an IPv4 or IPv6 router, as an IPv4 network address translation (NAT) router, or as a remote access server that hosted dial-up or virtual private network (VPN) connections from remote clients. Now, that feature has been combined with DirectAccess to make up the Remote Access server role in Windows Server 2012 This guide describes how to migrate a server that is hosting the Routing and Remote Access service (in Windows Server 2008 R2 and other down-level versions) to a computer that is running Windows Server 2012.

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic, and then add comments that support your rating. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

About this guide

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process. For more information about installing and using the migration tools on the source and destination servers, see the Windows Server Migration Tools Installation, Access, and Removal Guide (http://go.microsoft.com/fwlink/?LinkId=247607).

Target audience

This document is intended for information technology (IT) administrators, IT professionals, and other knowledge workers who are responsible for the operation and deployment of the Remote Access servers in a managed environment. Some scripting knowledge may be required to perform some of the migration steps that are contained in this guide.

What this guide does not provide

This guide does not describe the architecture or detailed functionality of the Remote Access role. The following scenarios are not supported in this migration guide:

- Any process for an in-place upgrade, in which the new operating system is installed on the existing server hardware by using the **Upgrade** option during setup

- Clustering and multisite scenarios
- Migrating more than one server role

  If your server is running multiple roles, it is recommended that you design a custom migration procedure that is specific to your server environment and based on the information that is provided in this and other role migration guides.

Supported migration scenarios

This guide provides instructions for migrating an existing server to a server that is running Windows Server 2012.

Caution

This guide does not contain instructions for migration when the source server is running multiple roles. If your source server is running multiple roles, some migration steps in this guide, such as those for migrating user accounts and network interface names, can cause other roles that are running on the source server to fail.

Supported operating systems

This guide provides instructions for migrating data and settings from an existing server that is being replaced by a new physical or virtual 64-bit server with a clean-installed operating system, as described in the following table.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	win8_Server_2	x64-based
x86- or x64-based	Windows Server 2003 R2	win8_Server_2	x64-based
x86- or x64-based	Windows Server 2008, full installation option only	win8_Server_2	x64-based
x64-based	Windows Server 2008 R2, full installation option only	win8_Server_2	x64-based
x64-based	Win8_Server_2	win8_Server_2	x64-based

- The versions of operating systems shown in the preceding table are the earliest combinations of operating systems and service packs that are supported. Newer service packs are supported.
- Migration with the destination server already having DirectAccess configured is not supported.
- The Foundation, Standard, Enterprise, and Datacenter editions of the Windows Server operating system are supported as either source or destination servers. This includes

migrating across editions. For example, you can migrate from a server running Windows Server 2003 Standard to a server running Windows Server 2012.

- Migrations between physical operating systems and virtual operating systems are supported.
- Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shared resources from a computer that is running Windows Server 2008 R2 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

  Note

  The system UI language is the language of the localized installation package that was used to set up the Windows operating system.
- Both x86- and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2008 R2 and win8_server_2 are x64-based.

Supported role configurations

The following is a broad list of the migration scenarios that are supported for Remote Access. All settings under these scenarios are migrated.

- **DirectAccess (supported in Windows Server 2012 to Windows Server 2012 migration only)**
- **VPN server**
- **Dial-up server**
- **Network address translation (NAT)**
- **Routing**, with the following optional components:
  - DHCP Relay Agent
  - Routing Information Protocol (RIP)
  - Internet Group Management Protocol (IGMP)

In addition to the above scenarios, migration also automatically adjusts configuration of the destination server to account for features that are no longer supported and to support features that are new in Windows Server 2012 and not supported on earlier versions of Windows.

Migration dependencies

If a local or remote NPS server that is used for authentication, accounting, or policy management must also be migrated, then migrate the NPS service before migrating Remote Access. For more information, see Migrate Network Policy Server to Windows Server 2012.

If you are upgrading from Windows 2008 R2 DirectAccess to Windows Server 2012, ensure that all the DirectAccess configuration settings have been applied on the Windows 2008 R2 server. It is possible to save settings through the console but not apply them. Before upgrading, ensure that the saved settings have also been applied.

Migration components that are not supported in all operating system versions

The following Remote Access components are not supported by all operating systems:

Component	UI Dialog/Settings	Action
New Components, not available on Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2		
Specifying adapter to obtain DNS/WINS addresses	RAS Properties – IPv4 Tab: Adapter	This component is not supported on Windows Server 2003. The default value should be used for this setting on the target computer.
SSTP	SSTP ports	SSTP is not supported on Windows Server 2003. SSTP ports should be enabled on the target computer. The number depends on the default value for the SKU on the target computer
IPv6	1. RAS Properties – General Tab: IPv6 router check-box and corresponding radio-buttons, IPv6 Remote access server 1. RAS Properties – IPv6 Tab: All settings 1. Demand-Dial (VPN/PPPoE) properties – Networking tab – TCP/IPv6 1. Demand-dial (VPN/PPPoE) – IPv6 filters for connection initiation 1. IPv6 - Router	• IPv6 is not supported on Windows Server 2003. It should be disabled on the target computer if DirectAccess is not deployed (legacy VPN). Under the General tab of RRAS properties, IPv6 Router and IPv6 Remote access server should not be selected. • The adapter setting under the IPv6 tab of RAS properties was introduced in Windows Server 2008 R2 for IKEv2. It was not present in Windows Server 2008. During migration this setting should be set to the

		default value on the target computer for Windows Server 2008 and be “as is” for Windows Server 2008 R2 and Windows Server 2012
IP Filters under RA Logging and Policies	Remote Access Logging & Policies – IP Filters	Windows Server 2003 and Windows Server 2008 do not have an option to create IP filters under Remote Access Logging and Policies. Hence there would be no filters to migrate.
Automatically obtaining IPv6 address	Demand-Dial (VPN/PPPoE) properties - Networking tab - IPv6 Properties – ‘Obtain IP address automatically’ radio-button	This setting is not present in Windows Server 2008. The value of this setting should be set to default on the target computer.
IKEv2	1. IKEv2 ports 1. RAS Properties – Security Tab: computer certificate authentication for IKEv2 1. RAS Properties – IKEv2 Tab: All settings	• IKEv2 is not supported on Windows Server 2003 and Windows Server 2008. IKEv2 ports should be enabled on the target computer. The number would depend on the default value for the SKU on the target computer. • Default values should be used for all IKEv2 settings on the target computer.
SSTP Cert. Selection	RAS Properties – Security Tab: ‘Use HTTP’ check-box, drop-down to select certificate, crypto binding settings	This component is not supported on Windows Server 2003 and Windows Server 2008. Default values should be used for all these settings on the target computer.
VPN Accounting	Remote Access Logging & Policies – Accounting: Logging failure action settings under ‘SQL Server Logging Properties’ and ‘Log File	These are not present in Windows Server 2008. The default value should be used on the target computer.

	Properties'	
Deprecated Features: Not available on Windows Server 2008, Windows Server 2008 R2, or Win8_Server_2		
SPAP, MS-CHAP, EAP-MD5 protocols and related settings	VPN/PPPoE Demand-dial interface properties - Security tab	SPAP, EAP-MD5 and MS-CHAP settings are not supported on Windows Server 2008 R2, or win8_Server_2, and will not be migrated.
Local Area Connection interface configuration under Routing	Routing – General – Local Area Connection properties – Configuration tab	This tab provides settings to configure how an IP address should be obtained for this interface. It is only present on Windows Server 2003 and will not be migrated.
RAS Firewall (integrated with NAT)	1. NAT – Interface – NAT/Basic Firewall Tab 1. NAT – Interface – ICMP Tab	Windows Server 2003 supported RAS firewall functionality which was removed in Windows Server 2008. These settings will not be migrated.
Weak Encryption settings		Weak encryption is supported on Windows Server 2003 but on Windows Server 2008 and Windows Server 2008 R2, it can only be enabled through the registry. During migration from Windows Server 2003 the registry settings will not be created automatically. For Windows Server 2008 and higher versions, if these registry settings happen to be present already, they will be migrated.

Migration components that are not automatically migrated

The following Remote Access elements and settings are not migrated by the Windows PowerShell cmdlets that are supplied with the Windows Server Migration Tools. Instead, you must manually configure the element or setting on the new RRAS server as described in **Completing the required manual migration steps** in this guide.

Important
Perform the manual configuration of these elements only when directed later in this guide.

- **SSL certificate bindings**. SSL Certificate binding and crypto-binding settings for SSTP are migrated as follows:
  a. The migration Wizard looks for a source certificate on the destination computer. If one is found, SSTP uses that certificate.
  b. If a source certificate is not found, the migration wizard will look for a valid certificate with the same trusted root as the source certificate.
  c. If still no certificate is found, then the SSTP configuration on the destination computer is 'Default'.
  d. If self-signed certificates are being used (valid for win8_Server_2), they will be automatically created on the destination computer.
- **User accounts on the local RRAS server**. If you use domain-based user and group accounts, and both the old and new RRAS servers are part of the same domain, no migration of the accounts is required. Local user accounts can be used if **Windows Authentication** is configured on the RRAS source server.
- **Only routing/VPN/DirectAccess when all are installed**. If your Remote Access server configuration includes all of the available services, then all services must be migrated together. Migrating only one of the services to the destination server is not supported.
- **A local or remote server that is running Network Policy Server (NPS)** that provides authentication, accounting, and policy management. This guide does not include the steps that are required to migrate a server that is running NPS. To migrate a server that is running NPS, use Migrate Network Policy Server to Windows Server 2012. NPS migration should be performed when directed later in this guide.

  Note
  If you are not using a server that is running NPS, the default Remote Access policies and accounting settings that are automatically created while configuring RRAS are not migrated.
- **Dial-up based demand-dial connections**. The destination server might have different modems, and there are many demand-dial settings that are specific to the modem or ISDN device that is selected.
- **Certificates** used for authenticating IKEv2, SSTP, and L2TP/IPsec connections.
- **SSL Certificate Binding** for SSTP when the **Use HTTP** check box is not selected.

- **IKEv2 VPN connections that use IPv6 network adapters**. IKEv2 is supported on RRAS servers that are running Windows Server 2008 R2 only. In the RRAS Microsoft Management Console (MMC) on Windows Server 2008 R2, you can specify the network interface used to acquire IPV6 DHCP and DNS addresses that are used for IKEv2 VPN clients. If you migrate RRAS from Windows Server 2008 R2 to another server running Windows Server 2008 R2, the setting is migrated. However, if you migrate from a previous version of Windows, there is no setting to migrate, and the default value of **Allow RAS to select the adapter** is used.
- **Weak encryption**. In Windows Server 2003, weak encryption is enabled, but on later versions of Windows it is disabled by default. You can enable weak encryption only by modifying the registry. During migration from Windows Server 2003 the required registry settings are not created on the new server by the migration process, and they must manually be configured. For later versions of Windows, if these registry settings are present, they are migrated.
- **Admin DLLs and Security DLLs** and their corresponding registry keys. These DLLs are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration.
- **Custom DLLs used for dialing a demand-dial connection**. These DLLs are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration. Any corresponding registry settings also are not migrated.
- **Connection Manager profiles**. The Connection Manager Administration Kit is used to create VPN and dial-up remote access profiles. Profiles created are stored under specific folders on the RRAS server. Profiles that are created on a 32-bit version of Windows do not work on computers that are running a 64-bit version of Windows, and vice versa. For more information about connection profiles, see Connection Manager Administration Kit (http://go.microsoft.com/fwlink/?linkid=55986).
- **The Group Forwarded Fragments setting on NAT**. This setting is enabled if the RRAS server is deployed behind a NAT router running on the Windows operating system. This is required for L2TP/IPsec connections that are using computer certificate authentication to succeed. We recommend that you enable this value to work around a known issue in RRAS.
- The **Log additional Routing and Remote Access information (used for debugging)** setting in the **Routing and Remote Access Properties** dialog box on the **Logging** tab.

Overview of the Routing and Remote Access service migration process

The pre-migration process involves the manual collection of data, followed by running procedures on the destination and source servers. The migration process includes source and destination server procedures that use the **Export** and **Import** cmdlets to automatically collect, store, and migrate server role settings. Post-migration procedures include verifying that the destination server successfully replaced the source server and then retiring or repurposing the source server. If the verification procedure indicates that the migration failed, troubleshooting begins. If troubleshooting fails, rollback instructions are provided to return the network to the use of the original source server.

10afc962-3596-46ec-b83c-7886d9af025b

Impact of migration

During migration, the Remote Access server is not available to accept incoming connections or to route traffic.

- New remote clients cannot connect to the server by using dial-up, VPN or DirectAccess connections. Existing connections on the server are disconnected. If you have multiple remote access servers, then the loss of availability of this server results in a reduction in capacity until the new server is operational again. For demand-dial connections, you must provide alternate connectivity between offices or reconfigure the connections to point to an alternate server.
- Routing and NAT functionalities are not available. If the functionality is required during the migration, an alternate router can be deployed until the new destination server is available.

Post-migration impacts include the following:

- If you plan to reuse the name of the source server as the name of the destination server, the name can be configured on the destination server after the source server is disconnected from the network. Otherwise, there is a name conflict that can affect availability. If you plan to run both servers, then the destination server must be given a unique name.
- A VPN server can be directly connected to the Internet, or it can be placed on a perimeter network that is behind a firewall or NAT router. If the IP address or DNS name of the destination server changes as part of the migration, or after the migration is completed, then the mappings in the firewall or NAT device must be reconfigured to point to the correct address or name. You must also update any intranet or Internet DNS servers with the new name and IP address. Also, remember to provide information about any server name or IP address changes to your users so that they can connect to the correct server. If you use connection profiles that are created by using the Connection Manager Administration Kit, then deploy a new profile with the updated server address information.

Note

For DirectAccess, the source computer and the destination computer must have the same IP addresses and interface names.

We recommend that you advertise the expected date and time of the migration so that users can plan accordingly, and make other arrangements as needed.

Permissions required to complete migration

The following permissions are required on the source server and the destination Remote Access servers:

- Domain user rights are required join the new server to the domain.
- Local administrative rights are required to install and manage the Remote Access role.
- Equivalent administrator permissions for DirectAccess GPOs as were configured on the source computer.

- Write permissions are required to the migration store location. For more information, see Remote Access: Prepare to Migrate in this guide.

Estimated duration

The migration can take two to three hours, including testing.

See Also

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Prepare to Migrate

Perform the following steps before you begin migrating Remote Access from your x86-based or x64-based source server to an x64-based destination server that is running Windows Server® 2012.

- Prepare your destination server
- Prepare your source server
- Install the migration tools

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, you might have to run the following steps by using the **Run as administrator** option.

Prepare your destination server

Complete the following procedures to prepare the destination server for the migration of Remote Access.

Hardware requirements for the destination server

Your destination server should have the same number or more network adapters as your source server. You can have more network adapters on the destination server than the source server, but the migration fails if you have fewer.

Important

The names of the network adapters on the destination server must be the same as those on the source server, and they must have the same intention (for example, Internet facing versus intranet facing). Most Remote Access server components have interface-specific settings and configuration. Having the same number of interfaces, with the same names

and intent, helps ensure that the settings are migrated to the right interface. This is critical to a successful migration. If there are more adapters on the destination server than on the source server, you must still have a one-to-one match between the names and intention of the network adapters on the source server and those on the destination server.

Note

DirectAccess configuration can only be migrated from a computer running Windows Server 2012 to another computer running Windows 2012. Migration from Windows Server 2008 R2 DirectAccess to Windows Server 2012 DirectAccess is not supported.

Prepare the destination server for migration

To prepare the destination server

1. Install Windows Server 2012 on the destination server.
2. Whether or not you intend to migrate the source server name to the destination server, give the destination server a temporary computer name at this time.
3. If you store the user accounts for remote access users locally on the Remote Access server instead of in Active Directory, and if you use the Challenge Handshake Authentication Protocol (CHAP) for authentication, then you must perform the following additional steps before migrating Remote Access:
   a. To enable the use of CHAP authentication, you must manually configure a local security policy setting that enables passwords to be stored by using a reversible encryption algorithm.

      Security

      We recommend that you do not use CHAP for authentication, and that you do not enable the setting to store passwords with reversible encryption. These options are not considered secure, and they are provided only for backwards compatibility. Use them only if your environment requires the use of CHAP.

      i. On the destination server, in the **Start** screen, click **Administrative tools**, and then click **Local Security Policy**.
      ii. In the navigation tree, expand **Account Policies**, and then select **Password Policy**.
      iii. In the details pane, double-click **Store passwords using reversible encryption**, click **Enabled**, and then click **OK**.

   b. Migrate the local users and groups from the source server to the destination server. Do this separately and before you begin migrating Remote Access.
4. If the source server that is being replaced is joined to a domain, join the destination server to the same domain.
5. In the dashboard of the Server Manager console click **Add roles and Features**.

6. Click **Next** until you reach the **Select Server Roles** screen.
7. On the **Select Server Roles** screen, select **Remote Access**. Click **Add Required Features**, and then click **Next**.
8. On the **Select features** screen, click **Windows Server Migration Tools**. Click **Next** until you reach the **Select Role Services** screen.
9. In the **Select Role Services** screen, select **Routing** and then click **Next**.
10. On the **Confirm Installation selections** screen, click **Install**.
11. On the **Installation progress** screen, verify that the installation was successful, and then click **Close**.
12. Server roles were introduced in Windows Server 2008, and they are also used in Windows Server 2012. Remote Access is a role service that consists of the Routing service and the DirectAccess and VPN service. In Windows Server 2003, the Routing and VPN services were not separate. If the source server is running Windows Server 2003, ensure that both the Routing service and the DirectAccess and VPN service are installed on the destination server. If the source server is running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012, ensure that the destination server has the same Remote Access services installed as has the source server. If the source server has the routing service and the DirectAccess and VPN service, then you must install all these components on the destination server.
13. If DirectAccess is being migrated, the IP-HTTPS and Network Location certificate must be imported to the destination server. Note that if self-signed certificates were being used on the source server, this step is not required.
14. If DirectAccess is being migrated, run the Windows PowerShell cmdlet **install-remoteaccess –prerequisite** to ensure that the destination server meets all the requirements for DirectAccess.

The destination server is now prepared for migration.

Prepare your source server

Important
Before you begin migration, as a best practice, we recommend that you perform a backup of the source server. If the migration fails, and the recovery steps to restore the source server also fail, this backup can be critical for the quick restoration of service.

Back up your source server

- For information about backing up Windows Server 2003, see Backing up and restoring data in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=163718).
- For information about backing up Windows Server 2008 or Windows Server 2008 R2, see Backup and Recovery in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=163719).

Install the migration tools

Windows Server Migration Tools in Windows Server 2012 allows an administrator to migrate some server roles, features, operating system settings, shares, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2 to computers that are running Windows Server 2012.

Install Windows Server Migration Tools on the source and the destination servers. Complete installation, configuration, and removal instructions for Windows Server Migration Tools are available in Install, Use, and Remove Windows Server Migration Tools.

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Migrate Remote Access

Complete the following procedures to migrate the Routing and Remote Access service from a source server to a destination server.

- Migrating Remote Access from the source server
- Migrating Remote Access to the destination server
- Completing the required manual migration steps

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, you might have to run the following steps by using the **Run as administrator** option. For more information, see Run a program with administrative credentials in the Windows Server TechCenter (http://go.microsoft.com/fwlink/?LinkId=131210).

Migrating Remote Access from the source server

Follow these steps to capture the configuration of Remote Access on the source server.

To capture Remote Access configuration: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2

1. On the source server, open a Windows PowerShell session with elevated user rights.
2. Load Windows Server Migration Tools into your Windows PowerShell session.

   Note
   If you opened the current Windows PowerShell session by using the Windows Server Migration Tools shortcut on the **Start** menu, skip this step and go to the next step. You should only load the Windows Server Migration Tools snap-in in a Windows PowerShell session that was opened by using another method, and into which the snap-in has not already been loaded.

   To load Windows Server Migration Tools, run the following cmdlet:

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```

3. Remote Access can be running on the source server while you are capturing its configuration. However, if you made configuration changes to Remote Access that require a service restart, then you must stop Remote Access before starting the migration. Use the following PowerShell command to stop the service:

   ```
   stop-service remoteaccess -force
   ```

   Note
   You must use the `-force` parameter because Remote Access has dependent services.

   To verify that the service is stopped, run the following command:

   ```
   get-service remoteaccess
   ```

4. On the source server, from Windows PowerShell, collect the settings from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The following is the syntax for the cmdlet:

   ```
   Export-SmigServerSetting -featureID NPAS-RRAS -User All -
   Group -path StorePath -verbose
   ```

   The **Export-SmigServerSetting** cmdlet parameters can collect all Routing and Remote Access service settings on the source server in a single file (Svrmig.mig). Before you run this command, review the following:

   - When you run the **Export-SmigServerSetting** command, you are prompted to provide a password to encrypt the migration store data. You must provide this same password when you later import from the migration store. Make sure you provide a strong password to encrypt the migration data and that the location of the migration data file is secure.
   - The *StorePath* variable that is provided as the value of the **path** parameter can be an

empty or nonempty folder. The actual data file that is placed in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Do not specify a file name. If a migration data file already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must first move the Svrmig.mig file from that location and store it elsewhere, rename it, or delete it.

- If the path is not a shared location that the destination server can access, you must manually copy the migration store to the destination server or to a location that the destination server can access.
- Migrating users and groups can be combined with the cmdlets that are used to migrate Remote Access. The **-Users** and **-Group** parameters can be used in the **Export-SmigServerSetting** command to migrate the user and group accounts that are present locally on the Remote Access source server. If you are using an Active Directory domain or RADIUS for authentication, then these parameters are not needed.

  The **-Users** command supports the following parameters:

  - **All**. All user accounts on the source server are included in the migration output file.
  - **Enabled**. Only enabled user accounts on the source server are included in the migration output file.
  - **Disabled**. Only disabled user accounts on the source server are included in the migration output file.

  To prevent migrating any user accounts, do not include the **-Users** parameter in the command.

  The **-Group** command takes no additional parameters. If it is present, then all groups defined locally on the source server are included in the migration output file.

Note

The process described in the Local User and Group Migration Guide does not migrate some settings, including those under the **Dial-in** tab. We recommend that you thoroughly review the Local User and Group Migration Guide to understand which settings are migrated and which are not.

To capture the Routing and Remote Access service configuration: Windows Server 2012

1. From the Start screen, click **Windows Server Migration Tools**.
2.

Note

This step is only required if Routing/VPN is configured on the source computer.

Remote Access can be running on the source server while you are capturing its configuration. However, if you made configuration changes to Remote Access that require a service restart, then you must stop Remote Access before starting the migration. Use the following PowerShell command to stop the service:

```
stop-service remoteaccess -force
```

Note

You must use the `-force` parameter because Remote Access has dependent services.

To verify that the service is stopped, run the following command:

```
get-service remoteaccess
```

3. Before you start to capture the configuration, you must stop the Remote Access Management service. Use the following PowerShell command to stop the service:

   ```
   stop-service Ramgmtsvc
   ```

   To verify that the service is stopped, run the following command:

   ```
   get-service Ramgmtsvc
   ```

   Once the export and migration are complete, you can restart the Remote Access Management service:

   ```
   start-service Ramgmtsvc
   ```

4. On the source server, from Windows PowerShell, collect the settings from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The following is the syntax for the cmdlet:

   ```
   Export-SmigServerSetting -featureID DirectAccess-VPN [-User
   All] -Group -path StorePath
   ```

   The **Export-SmigServerSetting** cmdlet parameters can collect all Remote Access settings on the source server in a single file (Svrmig.mig). Before you run this command, review the following:

   - When you run the **Export-SmigServerSetting** command, you are prompted to provide a password to encrypt the migration store data. You must provide this same password when you later import from the migration store. Make sure you provide a strong password to encrypt the migration data and that the location of the migration data file is secure.
   - The *StorePath* variable that is provided as the value of the **path** parameter can be an empty or nonempty folder. The actual data file that is placed in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Do not specify a file name. If a migration data file already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must first move the Svrmig.mig file from that location and store it elsewhere, rename it, or delete it.
   - If the path is not a shared location that the destination server can access, you must manually copy the migration store to the destination server or to a location that the destination server can access.
   - Migrating users and groups can be combined with the cmdlets that are used to migrate Remote Access. The **-Users** and **-Group** parameters can be used in the **Export-SmigServerSetting** command to migrate the user and group accounts that

are present locally on the Remote Access source server. If you are using an Active Directory domain or RADIUS for authentication, then these parameters are not needed.

The **-Users** command supports the following parameters:

- **All**. All user accounts on the source server are included in the migration output file.
- **Enabled**. Only enabled user accounts on the source server are included in the migration output file.
- **Disabled**. Only disabled user accounts on the source server are included in the migration output file.

To prevent migrating any user accounts, do not include the **-Users** parameter in the command.

The **-Group** command takes no additional parameters. If it is present, then all groups defined locally on the source server are included in the migration output file.

Note

The process described in the Local User and Group Migration Guide does not migrate some settings, including those under the **Dial-in** tab. We recommend that you thoroughly review the Local User and Group Migration Guide to understand which settings are migrated and which are not.

Migrating Remote Access to the destination server

Return to the destination server, and use the following procedure to complete the migration:

To import the Routing and Remote Access service configuration to the destination server

1. Before you use the **Import-SmigServerSetting** cmdlet to import the Routing and Remote Access service settings, be aware of the following condition:
   - If you chose to migrate the users and groups on the source computer, you need to specify the **-User** and **-Group** parameters in the **Import-SmigServerSetting** cmdlet on the destination server.
2. On the destination Server that is running Windows Server 2012, from the Start screen, click **Windows Server Migration Tools**.
3. On the destination server, from Windows PowerShell, run the following command, where *StorePath* is the folder that contains the Svrmig.mig file that you exported from the source server. Do not include the name of the file in the path.

```
Import-SmigServerSetting -featureID DirectAccess-VPN [-User
All] -Group -path StorePath –Force
```

For more information about running the **Import-SmigServerSetting** cmdlet, see the "Using Windows Server Migration Tools" section in the Windows Server Migration Tools Install, Use, and Remove Windows Server Migration Tools guide.

4.

Note

This step is only required if Routing/VPN is configured on the source computer.

Before starting the Remote Access service, you must manually stop the RASMAN service. Run the following command in the Windows PowerShell Command Prompt window:

```
Stop-service -force rasman
```

5.

Note

This step is only required if Routing/VPN is configured on the source computer.

Then run the following command in the Windows PowerShell Command Prompt window to start the Routing and Remote Access service:

```
Start-Service RemoteAccess
```

If a failure occurs while running the **Import-SmigServerSetting** cmdlet, review the Setupact.log, Setuperr.log, and ServerMigration.log files under %localappdata%\SvrMig\Log. Information about how the Remote Access components migrated is included in the Servermigration.log file.

After the script completes, review the following section and adjust any remaining settings that require manual configuration.

Completing the required manual migration steps

Certain settings cannot be migrated by the Windows PowerShell scripts, and they must be configured manually on the destination server. Review the following configuration options, and apply those that are relevant to your environment.

DirectAccess

To ensure that the destination server meets all DirectAccess requirements, run the following Windows PowerShell cmdlet: **Install-remoteaccess –prerequisite**.

Dial-up demand-dial connections

Because of the differences in modem hardware that might exist between the source and destination servers, dial-up connections are not migrated. Use the Demand-Dial Interface Wizard in the Remote Access MMC snap-in.

To create a dial-up demand-dial connection

1. If you are using Server Manager, in Tools click Routing and Remote Access.
2. Right-click the server in the tree, and then click **Configure and Enable Routing and Remote Access**.
3. Follow the steps in the wizard to configure the connection.

Certificates for IKEv2, SSTP, and L2TP/IPsec connections

Certificates can be exported from the source server and imported to the destination server by using the Certificates MMC snap-in.

Routing and Remote Access service policies and accounting settings

If you are not using a local or remote server to run NPS, then default remote access policies and accounting settings are automatically created on the destination server when Remote Access is configured.

To migrate NPS settings, refer to Migrate Network Policy Server to Windows Server 2012.

PEAP, smart card, and other certificate settings on Network Policy Server

If you also migrated a local or remote server running NPS to support the Remote Access server that you are migrating, we recommend that you verify that the server that is running NPS has the correct certificate configuration. Specifically, confirm that any certificates that are associated with Protected Extensible Authentication Protocol (PEAP) and the **Smart card or other certificate** authentication settings are set properly. You can find these settings on the server that is running NPS, in the NPS MMC snap-in under **Connection Request Policies** or **Network Policies** (depending on where the authentication protocols are configured). If no certificates are present, or if the certificates are not configured correctly, perform the following steps:

To reconfigure PEAP or smart card certificates

1. Remove the **PEAP** or **SmartCard or other certificate** methods from the list of authentication methods.
2. Add the method back to the list.
3. Reconfigure the certificate for the specified method.

Weak encryption settings

In Windows Server 2003 weak encryption is enabled, but on later versions of Windows it is disabled by default. You can enable weak encryption only by modifying the registry. During migration from Windows Server 2003, the required registry settings are not created on the new server by the migration process, and they must manually be configured. For later versions of Windows, if these registry settings are present, they are migrated. For more information about the

registry entries that Remote Access adds, see “Registry entries that Routing and Remote Access adds in Windows Server 2008”, article 947054 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=159112). The description of the settings for the weak encryption settings are at the end of the article, and they are named **AllowPPTPWeakCrypto** and **AllowL2TPWeakCrypto**.

Security

Weak encryption includes the use of 40-bit or 56-bit encryption in PPTP, and the use of MD5 or DES for L2TP/IPsec. By default, these weak algorithms are disabled, and we recommend that you do not use them unless they are required.

Connection Manager profiles

Profiles that are created by the Connection Manager Administration Kit (CMAK) can only be created on a computer with the same 32-bit or 64-bit architecture as the client computer on which they are to be run. If your source server is 64-bit, and you have created 64-bit profiles on that source server, you can copy them from the **%PROGRAMFILES%\CMAK\Profiles** folder to the appropriate folder on the destination server.

If the source server is 32-bit, you must use a computer running a 32-bit version of Windows to create and manage the profiles. You can set up a computer running a 32-bit version of Windows 7 or Windows 8, and then install CMAK on it to manage the profiles for your 32-bit client computers. For more information, see Connection Manager Administration Kit in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=136440).

Group forwarded fragments

The **Group Forwarded Fragments** setting on NAT is enabled if the Remote Access server is deployed behind a NAT device that runs Windows. This is required for L2TP/IPsec connections that are using computer certificate authentication to succeed. We recommend that you enable this setting. Group Forwarded Fragments can be enabled for IPv4 on the Windows NAT computer by running the following command at the command prompt:

```
netsh int ipv4 set global groupforwardfragments=enabled
```

RAS administration and security DLLs

Administration DLLs and security DLLs and their corresponding registry keys are not migrated. This is because they are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration. If the source and destination computers are 64-bit, the administration and security DLLs can be reused. For more information, refer to the following topics:

RAS Administration DLL (http://go.microsoft.com/fwlink/?linkid=163778)

RAS Security DLL (http://go.microsoft.com/fwlink/?linkid=163779)

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Verify the Migration

After all the migration steps are completed, you can use the following procedure to verify that the migration of Remote Access was successful. If the migration failed, you can return to the previous valid configuration by following the roll-back steps in Remote Access: Post-migration Tasks.

Verifying the destination server configuration

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, then you might have to run the following steps by using the **Run as administrator** option.

We recommend that you check the configuration of the destination Remote Access server, from the service start-up to the detailed configuration of individual components. The following sections provide a list of items to check. Depending on which Remote Access components are enabled on your server, only some of these checks might be necessary.

Installation state of Remote Access

The first verification step is to confirm that the Remote Access feature installed successfully.

To verify that Remote Access installed on the destination server

1. Click **Windows Server Migration Tools** on the **Start** screen.
2. View the installation status of the Routing and Remote Access service by running the following command:

   ```
   Get-WindowsFeature RemoteAccess
   ```

   The check box on the left of the **Remote Access** feature name is selected if the service is installed on the destination server. If it is not installed, the check box is clear.

Status of Remote Access Service

Verify that the Remote Access service is running.

To verify that the Routing and Remote Access service is running on the destination server

1. Click **Windows Server Migration Tools** on the **Start** screen.
2. View the service status of the Routing and Remote Access service by running the following command:

   ```
   Get-service RemoteAccess
   ```
3. Examine the **Status** column. It should read **Running**.

Remote access Operations Status

Verify the operations status of the deployment.

To verify the Remote Access operations status

1. In **Server Manager** click **Tools** and then click **Remote Access Management**.
2. Click **OPERATIONS STATUS** to navigate to **Operations Status** in the **Remote Access Management Console**. **Operations Status** lists the server operational status and that of all its components.

DirectAccess configuration

Verify the operations status of the deployment.

To verify the DirectAccess configuration settings

1. In **Server Manager** click **Tools** and then click **Remote Access Management**.
2. Click **CONFIGURATION** to navigate to Configuration tab in **Remote Access Management console**. Step through each of the wizards to ensure that the configuration has been migrated successfully.

VPN configuration

Confirm the configuration settings for the Remote Access server and ports.

To verify the Remote Access configuration settings

1. Start Server Manager.
2. In **Tools**, click **Routing and Remote Access**.
3. Right-click the Remote Access server node, and then click **Properties**.

   On each tab, confirm that the destination server is configured the same as the source server, and then click **OK**.
4. In the navigation pane, select **Ports**.

   Confirm that any modem or ISDN devices that are attached to the computer are included in the list.
5. In the navigation pane, right-click **Remote Access Logging and Policies**, and then click **Launch NPS**. In the Network Policy Server navigation pane, select **Network Policies**.

Confirm that the NPS policies that are currently configured are those required for your environment. If you migrated them from an NPS source server to an NPS destination server, confirm that you are connected to the destination server and that the policies migrated successfully.

Dial-up configuration

You must confirm that the correct phone lines are attached to the modems or ISDN ports on the destination server.

Demand-dial VPN configuration

Examine all of your demand-dial VPN connections to ensure that they migrated with the correct settings.

To verify the settings for a demand-dial VPN connection

1. Start Server Manager.
2. Click **Routing and Remote Access**, and then select **Network Interfaces**.
3. In the details pane, right-click a demand-dial interface, and then click **Properties**.

   On each tab, confirm that the connection is configured the same as the source server, and then click **OK**.

Router settings

Confirm that the router components installed, and verify that each is configured correctly. The available routing components include:

- **IPv4:** Static Routes, DHCP Relay Agent, IGMP, NAT, and RIPv2
- **IPv6:** Static Routes and DHCPv6 Relay Agent

To verify the routing components

1. Start Server Manager.
2. Click **Routing and Remote Access**.
3. Expand **IPv4**. Examine the list of installed routing components, and ensure that the components required for your deployment are installed.
4. Expand **IPv6** and follow the same process as the previous step.
5. In the navigation pane, under **IPv4**, click **General**.

   The details pane identifies the interfaces that are configured to route packets for each version of IP. Confirm that the list contains the expected interfaces, including configured demand-dial interfaces.
6. In the navigation pane, under **IPv6**, click **General** and follow the same process as the previous step.

7. In the details pane for **General** under **IPv4**, right-click each interface and select **Properties**.

   On each tab confirm that the interface is configured as required for its routing role on the server.

8. Follow the same process as described in the previous step for the interfaces listed on the under **IPv6** / **General**.
9. Under **IPv4** select **Static Routes** and confirm that the routes to destination networks are correctly configured with the associated interface and destination gateway address.
10. Follow the same process as described in the previous step for the **Static Routes** under **IPv6**.
11. Under **IPv4**, select **NAT**. The details pane shows the interfaces that NAT is configured to use. Right-click each interface and click **Properties**.
    - Confirm that each interface is configured correctly for NAT. There should be at least two interfaces enabled for NAT, one configured as the **Private** interface, and one configured as the **Public** interface.
    - If NAT is responsible for providing IPv4 addresses to clients on the private network, then on **NAT Properties** page, on the **Address Assignment** tab, select the **Automatically assign IP addresses by using the DHCP allocator** check box and enter the address information to be used.
    - If your ISP has provided a pool of addresses to be used by the NAT public interface, ensure that they are configured correctly. The addresses are under **NAT**, on the **Properties** page for the interface, on the **Address Pool** tab. If the addresses that were migrated are not applicable to the target computer, modify the list to use the correct addresses.
    - For each interface under **NAT**, on the interface's **Properties** page on the **Services and Ports** tab, examine the port mappings for services that must be routed to a specific server IP address. Confirm that each service that is to be mapped has the correct address pool entry, private IP address, and port settings configured.
12. Under **IPv4**, select each enabled routing protocol. The details pane shows the interfaces on which the selected routing protocol is enabled. Right-click each interface, and then click **Properties**.

    Confirm that each interface is configured correctly for the selected routing protocol. For example, under **IPv4**/**NAT**, there should be at least two interfaces, one configured as the **Private** interface, and one configured as the **Public** interface.

13. Under **IPv6**, select each enabled routing protocol and follow the same process described in the previous step.
14. Under **IPv4**, right-click each routing protocol, and then select **Properties** to examine the global configuration for that routing protocol.

    Confirm that each protocol is configured correctly for your environment. For example, ensure that the **DHCP Relay Agent** has a list of DHCP server addresses to which it can forward DHCP requests from clients.

15. Under **IPv6**, select each enabled routing protocol and follow the same process described

in the previous step.

User and Group accounts

If you migrated the user and group accounts by using the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?linkid=163774), follow the procedures in its verification section to confirm that the required users and group were migrated successfully.

If you instead used the **-user** and **-group** parameters on the **Import-SmigServerSetting** command, you can manually verify the accounts by using the **Local Users and Groups** MMC snap-in to examine the user and group accounts and confirm that the properties for the accounts are set properly.

Final checks

- If your computer is configured to host VPN/DirectAccess connections, test each type of supported connection to confirm that users can connect.
- If your server is configured to host dial-up connections, verify that client computers can successfully dial-in and connect to the server by using the modems that are installed.
- If your server is configured as an IPv4 or IPv6 router, verify that clients on each attached network can connect through the router to computers on all of the other attached networks. If you use the **ping** command for this test, ensure that Windows Firewall on the router and the client computers is configured to allow ICMP Echo Request and ICMP Echo Reply messages.

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Post-migration Tasks

Remote Access: Post-migration Tasks

Perform the following post-migration tasks to complete your migration:

- Completing the migration
- Configuring firewall rules for VPN
- Configuring firewall rules for DirectAccess
- Restoring Remote Access in the event of migration failure
- Retiring Remote Access on your source server
- Troubleshooting cmdlet-based migration

Note

The post-migration tasks for the source server are optional, depending on your migration scenario.

Completing the migration

Migration is complete as soon as verification efforts demonstrate that the destination server has replaced the source server in serving the network.

If your verification efforts indicate that the migration failed, follow the steps in Restoring Remote Access in the event of migration failure to return to the previous valid configuration. The following list identifies some settings that must be manually configured after the migration is complete.

- **Server name**: If the source server is to be decommissioned, the same name is available to be configured on the destination server. If the IP addresses of the intranet and internet interfaces have changed, it is important to ensure that the intranet and internet DNS servers are updated with the new name and IP addresses. These configuration details also need to be deployed to your users by updating their connection profiles so that they can connect to the correct server.
- **Perimeter network firewall or NAT**: If the destination server is in a perimeter network, the firewall or NAT through which the server is accessed from the Internet must be configured with the new IP address of the RRAS server. Refer to the documentation for your Firewall or NAT router for the relevant configuration instructions.

Configuring firewall rules for VPN

Firewall rules that permit VPN network traffic are included with the Windows Firewall with Advanced Security, but they are disabled by default. The rules that enable the required inbound network traffic must be enabled on the destination server. If there were any other rules explicitly configured on the source server to support RRAS or its roles, they should be configured on the destination server also. If you use non-Microsoft firewall software, refer to the documentation that is provided by the vendor for instructions about how to configure the appropriate rules.

When you migrate RRAS from the source server to the destination server, the firewall rules are not automatically enabled. You must use Windows Firewall with Advanced Security or a Group Policy setting to enable the rules that correspond to the types of RRAS network traffic that must enter your server. The following is a list of rules that you should enable, depending on which RRAS protocols you use:

- Routing and Remote Access (GRE-In)
- Routing and Remote Access (L2TP-In)
- Routing and Remote Access (PPTP-In)
- Secure Socket Tunneling Protocol (SSTP-In)

If you change the default firewall behavior to blocking all traffic that does not match an allow rule then you must enable the outbound allow rules in addition to the inbound rules.

Configuring firewall rules for DirectAccess

When using additional firewalls in your deployment, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- Teredo traffic—User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound.
- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required.

  Note
  This exemption must be configured on the Remote Access server, while all the other exemptions have to be configured on the edge firewall.

Note
For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server. For IP-HTTPS the exceptions need only be applied to the address where the public name of the server resolves.

When using additional firewalls, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.
- Internet Control Message Protocol for IPv6 (ICMPv6) traffic inbound and outbound – for Teredo implementations only.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic
- ICMP for all IPv4/IPv6 traffic

Restoring Remote Access in the event of migration failure

If migration of Remote Access to the destination server fails, you can put the source server back into operation by following these steps:

- If the source server has not been repurposed and the computer name and IP address have not been migrated from the source to the destination server, simply connect the source server to the network and start the Routing and Remote Access service to allow users to connect again.

- If the computer name and IP address have been migrated from the source server to the destination server, rename the destination server to a temporary name and change its IP address to a different IP address. Set the source server computer name and IP address to the values that were used before the migration, and restart the Routing and Remote Access service on the source server.
- If the previous steps are not valid options, such as when the source server has been repurposed or is otherwise unavailable, use the backup files that were created from the source server, as described in Remote Access: Prepare to Migrate. You can use the backup files any time after migration to restore the original Routing and Remote Access source server if a catastrophic failure occurs.

Estimated time to complete a rollback

You should be able to complete a rollback in one to two hours.

Retiring Remote Access on your source server

After you verify that the migration is complete, the source server can be disconnected from the network and removed from service. Stop RRAS before you remove the computer from the network and turn it off. You can keep this computer as a backup server in the event that you want to revert to your previous Routing and Remote Access configuration.

Troubleshooting cmdlet-based migration

The Windows Server Migration Tools deployment log file is located at %*windir*%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations.

- %*windir*%\Logs\ServerMigration.log
- On Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012: %*localappdata*%\SvrMig\Log
- On Windows Server 2003: %*userprofile*%\Local Settings\Application Data\SvrMig\Log

If migration log files cannot be created in the preceding locations, **ServerMigration.log** and **SmigDeploy.log** are created in %*temp*%, and other logs are created in %*windir*%\System32.

If you are migrating Remote Access and a failure occurs while running the **Import-SmigServerSetting** cmdlet, review the Setupact.log, Setuperr.log, and ServerMigration.log files under %localappdata%\SvrMig\Log. Information about how the Remote Access components migrated is included in the Servermigration.log file.

If a migration cmdlet fails, and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at
0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050).  Minidump
attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Viewing the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects, and query them for more information about settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The Windows Server Migration Tools **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.

Property name	Type	Definition
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or share that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item, shown as a path name.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to both **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.
Title	String	Title of the result object. Examples of values include NIC physical address for IP configuration, or user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   $*VariableName* **=** $(*Cmdlet*)

   The following is an example.

   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -User all -Group -Path D:\rmt\DemoStore -force -Verbose)

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.
2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

$*VariableName*

In the following example, the variable is named **ImportResult**.

$*ImportResult*

This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable.

```
            ItemType  ID                                Success
DetailsList
            --------  --                                -------
-----------
          OSSetting  Local User                          True
{Local User, Loc...
          OSSetting  Local Group                         True
{Local Group, Lo...
     WindowsFeature  DHCP                                True
{}
```

Each line of the preceding sample is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by examples in step 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   $<VariableName>| Select-Object -ExpandProperty <PropertyName>

   The following is an example.

   $importResult | Select-Object -ExpandProperty DetailsList

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples.

   - The following command returns only those details of result objects that have the ID **Local User**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Severity -eq "Warning") {$_} }

   - The following command returns only the details of result objects with an ID of **Local User** that also have the title **Remote Desktop Users**.

```
$ImportResult | Where-Object { $_.ID -eq "Local Group" } | Select-Object -
ExpandProperty DetailsList | ForEach-Object { if ($_.Title -eq "Remote
DesktopUsers") {$_} }
```

More information about querying results

For more information about the cmdlets that are used in the preceding examples, see the following additional resources.

- Where-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134853).
- Select-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134858).
- ForEach-Object on the Microsoft Script Center Web site (http://www.microsoft.com/technet/scriptcenter/topics/msh/cmdlets/foreach-object.mspx)

For more information about Windows PowerShell scripting techniques, see What Can I Do With Windows PowerShell? - Scripting Techniques on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134862).

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Migrate Windows Server Update Services to Windows Server 2012

This document describes the process to migrate an existing Windows Server Update Services (WSUS) 3.0 SP2 server role to a destination server that is running Windows Server 2012 or Windows Server 2012 R2. This document includes instructions for moving the updates, settings, target groups, and computers to the new server. By using the tools that are described in this document, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help eliminate possible conflicts that might otherwise occur during the migration process.

- Step 1: Plan for WSUS Migration
- Step 2: Prepare to Migrate WSUS
- Step 3: Migrate WSUS
- Step 4: Verify the WSUS Migration

Step 1: Plan for WSUS Migration

The first step in the migration of your Windows Server Update Services (WSUS) to Windows Server 2012 or Windows Server 2012 R2 is to understand the supported and unsupported scenarios and the supported operating systems for this migration. The following checklist describes the steps involved in planning for your WSUS migration.

Task	Description
1.1. Know supported operating systems	Review the list of supported source operating systems and WSUS versions.
1.2. Review supported migration scenarios	Review the list of supported migration scenarios.
1.3. Review migration scenarios that are not supported	Review the list of unsupported migration scenarios.

1.1. Know supported operating systems

Migration from the following operating systems is supported on Windows Server 2012 and Windows Server 2012 R2:

- Windows Server 2008 R2 running WSUS 3.0 SP2
- Windows Server 2008 (full installation option) running WSUS 3.0 SP2
- Windows Server 2003 SP2 running WSUS 3.0 SP2

1.2. Review supported migration scenarios

The following WSUS migration scenarios are supported:

- Windows Server 2012 Standard and Datacenter editions, and servers running Windows Server 2012 R2 can be used as source or destination servers.
- Windows Server 2012 Enterprise edition can be used as a source server.
- Migration between physical operating systems and virtual operating systems.
- Migration from a source server that is running SQL Server 2005 to a destination server that is running SQL Server 2008 R2 SP1.
- Migration from a source server that is running Windows Internal Database to a destination server that is running SQL Server 2008 R2 SP1.
- Migration from a domain to a workgroup or from a workgroup to a domain. However, if the source server is running SQL Server from a remote location, migration from the domain to a workgroup is not supported.
- The destination server must meet the Windows Server 2012 or Windows Server 2012 R2 WSUS role minimum system requirements for hardware and software.

Important

For more information about minimum system requirements and hardware capacity requirements for the WSUS server, see the Deploy Windows Server Update Services in Your Organization.

1.3. Review migration scenarios that are not supported

The following WSUS migration scenarios are not supported:

- Migration from an unsupported version of WSUS (prior to WSUS 3.0 SP2). Upgrade the existing WSUS server to a supported version before you migrate the WSUS server role to Windows Server 2012 or Windows Server 2012 R2.
- Migration from a Server Core installation option (WSUS 3.0 SP2 does not support a Server Core installation).
- Migration from a domain that is using SQL Server from a remote location to a workgroup.
- Migration from a source server that is running SQL Server to a destination server that is running Windows Internal Database.
- Migration from a source server that stores updates on Microsoft Update to a destination server that stores updates on a local WSUS server, and vice versa. Changing the configuration during the migration process is not supported.
- Migration from a source server to a destination server that is running an operating system in a different system UI language. The system UI language is the language of the localized installation package that was used to set up the Windows operating system. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2008 R2 in the German system UI language.

See also

- Step 2: Prepare to Migrate WSUS
- Migrate Windows Server Update Services to Windows Server 2012
- WSUS server role description

Step 2: Prepare to Migrate WSUS

The second step in the migration of your Windows Server Update Services (WSUS) server role involves preparing the destination and source servers. The following checklist describes the steps involved to prepare for your WSUS migration.

Task	Description
2.1. Prepare before you start the migration	Review the recommended guidelines before starting the migration process.
2.2. Prepare the destination server	Understand the steps that must be completed on the WSUS destination server before the migration.
2.3. Prepare the source server	Understand the steps that must be completed on the WSUS source server before the migration.

2.1. Prepare before you start the migration

This migration procedure assumes a working knowledge of deployment basics for Windows Server Update Services (WSUS) 3.0 SP2. For more information about WSUS deployment, see Deploy Windows Server Update Services in Your Organization. We recommend that you make the following decisions and preparations before you start the migration process:

Warning

Upgrade from any version of Windows Server that supports WSUS 3.2 to Windows Server® 2012 R2 requires that you first uninstall WSUS 3.2.

In Windows Server 2012, upgrading from any version of Windows Server with WSUS 3.2 installed is blocked during the installation process if WSUS 3.2 is detected, and you are prompted to first uninstall Windows Server Update Services prior to upgrading Windows Server 2012.

However, because of changes in Windows Server 2012 R2, when upgrading from any version of Windows Server and WSUS 3.2 to Windows Server 2012 R2, the installation is not blocked. Failure to uninstall WSUS 3.2 prior to performing a Windows Server 2012 R2 upgrade will cause the post installation tasks for WSUS in Windows Server 2012 R2 to fail. In this case, the only known corrective measure is to format the hard drive and reinstall Windows Server 2012 R2.

- Configure a location to store updates on the source server. Changing the content storage configuration as part of the migration process is not supported. You can store updates on the local WSUS server or on Microsoft Update. If you want the destination server to store updates in a different location than the source server, the new location must be configured on the source server before migration.
- Confirm that the destination server meets the minimum WSUS hardware requirements and database requirements. For more information about those requirements see Deploy Windows Server Update Services in Your Organization on Microsoft TechNet.

2.2. Prepare the destination server

Before migrating WSUS, set up a new Windows Server 2012 in your organization as the WSUS destination server and install WSUS server role on the destination server. After you have successfully installed the WSUS server role, the Configuration Wizard starts automatically. Close the Configuration Wizard. Do not try to sync the updates at this point, because you will copy the update binary files later in the migration process. The WSUS installation procedure assumes that updates for the new server come from Windows Update.

After this is complete, follow these guidelines:

- If you have decided to use the full installation of SQL Server as the WSUS database, install SQL Server 2008 R2 Standard or SQL Server 2008 R2 Enterprise.
- Download a graphical tool to manage your database on the destination server from Microsoft SQL Server Management Studio Express or Microsoft SQL Server 2008 R2 Management Studio Express.
- Open TCP port 7000 and make sure that it is not being used by other applications. This port is used by Send-SmigServerData and Receive-SmigServerData to perform the data transfer.
- If the destination server is not joined to the source server's domain, visually verify that the time, date, and time zone on the destination server are synchronized with the source server. Use the Windows Control Panel to update the date and time if it is necessary.

Important
For more information about minimum system requirements and hardware capacity requirements for the WSUS server, see the Deploy Windows Server Update Services in Your Organization.

2.3. Prepare the source server

Review and take action based on the following guidelines:

- Refer to Appendix A: Migration Data Collection Worksheet to collect data about the source server.
- Open TCP port 7000 and make sure that it is not being used by other applications. This port is used by Send-SmigServerData and Receive-SmigServerData to perform the data transfer.
- If you have changed the default behavior of Windows Firewall (or another firewall program) to block outgoing traffic on computers that are running Windows Server 2012, you must enable outgoing traffic on UDP port 7000.
- Download a graphical tool for managing your database on the source server at Microsoft SQL Server Management Studio Express.

See also

- Step 3: Migrate WSUS
- Step 1: Plan for WSUS Migration
- WSUS server role description

Step 3: Migrate WSUS

During the third step in the migration of your Windows Server Update Services (WSUS) server role, you will migrate binaries and security groups, back up the database, change the server identity, and apply security settings The following checklist describes the steps involved.

Task	Description
3.1. Migrate WSUS update binaries	Move WSUS update binaries from the source server to the destination server.
3.2. Migrate WSUS security groups	Migrate local users and groups manually or by using Windows Server Migration Tools.
3.3. Back up the WSUS database	Use SQL Server Management Studio to back up and restore the WSUS database, computer groups, update approvals, and WSUS settings.
3.4. Change the WSUS server identity	Change the WSUS server identity on the destination server. Performing this step ensures that there is no effect on clients that are managed by WSUS during the migration process.
3.5. Apply security settings	Configure security settings on the new server. This includes configuring security settings on the destination server that you were using on the source server.
3.6. Review additional considerations	Review some additional actions that you should take after the migration is complete.

3.1. Migrate WSUS update binaries

Use your preferred method to copy WSUS update binaries in the WSUS folder from the source server to the destination server (for example, Windows Server Migration Tools, Windows Explorer, Xcopy, or Robocopy). If you decide to use Windows Server Migration Tools to migrate WSUS update binaries to a destination server that is running Windows Server 2012, see Migrate WSUS Update Binaries from the Source Server to the Destination Server Using Windows Server Migration Tools.

Important
Migrating WSUS update binaries is unnecessary if update files are stored on Microsoft Update.

3.2. Migrate WSUS security groups

You have the option of manually migrate only the WSUS Administrators and WSUS Reporters local security groups. Or, you can use Windows Server Migration Tools to migrate all local users and groups (including the WSUS Administrators and WSUS Reporters local security groups) from the source server to the destination server.

Warning
The WSUS Server Migration Tools can be installed on the server using the Server Manager Add Features option.

Before you perform this procedure, verify that the destination server can resolve the names of domain users who are members of the local group during the import operation. If the source and destination servers are in different domains, the destination server must be able to contact a global catalog server for the forest in which the source domain user accounts are located. Use the following guidelines:

- **If the source server is a member of the domain and the destination server is a domain controller**: Imported local users are elevated to domain users, and imported local groups become Domain Local groups on the destination server.
- **If the source server is a domain controller, and the destination server is not**: Domain Local groups are migrated as local groups, and domain users are migrated as local users.

Use the following procedure to manually migrate users to the WSUS Administrators and WSUS Reporters local security groups.

To manually migrate local users and groups

1. Right-click in the Taskbar, click **Properties**, highlight **Toolbars**, and then click **Address**.
2. Type **lusrmgr.msc**, and then press ENTER.
3. In the console tree of the Local Users and Groups MMC snap-in, double-click **Users**.
4. Manually create a list of the local users.
5. In the console tree of the Local Users and Groups MMC snap-in, double-click **Groups**.
6. Manually add the users from the source server to the WSUS Administrators and WSUS Reporters groups.

Use the following procedure to use Windows Server Migration Tools to migrate users to the WSUS Administrators and WSUS Reporters local security groups.

To use Windows Server Migration Tools to migrate users

1. Open a Windows PowerShell session on the source server and on the destination server.
2. Type the command below and press ENTER:

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```

3. In the Windows PowerShell session on the source server, type the following command to export local users and groups to a migration store:

   ```
   Export-SmigServerSetting -User <Enabled | Disabled | All> -Group -Path
   ```

```
<MigrationStorePath> -Verbose
```

MigrationStorePath represents the path of the location where you want to store migrated data. You can also use one of the following values with the **-User** parameter:

- **Enabled**: Export only enabled local users
- **Disabled**: Export only disabled local users
- **All**: Export enabled and disabled local users

4. Press ENTER.

Important

You are prompted to provide a password to encrypt the migration store. Remember this password, because you must provide the same password to import data from the migration store on the destination server. If the path is not a shared location to which the destination server has access, you must copy the migration store to the destination server manually, or to a location that this destination server can access as it runs the **Import-SmigServerSetting** cmdlet.

5. In the Windows PowerShell session on the destination server, type the following command to import local users and groups from the migration store that you created in Step 2:

```
Import-SmigServerSetting -User <Enabled | Disabled | All> -Group -Path
<MigrationStorePath> -Verbose
```

MigrationStorePath represents the path of the location from which you want to import migrated data. You can also use one of the following values with the **-User** parameter:

- **Enabled**: Export only enabled local users
- **Disabled**: Export only disabled local users
- **All**: Export enabled and disabled local users

6. Press ENTER.

Warning

After you enter the **Import-SmigServerSetting** cmdlet, you are prompted to provide the same password to decrypt the migration store that you created during the export process.

3.3. Back up the WSUS database

WSUS servers can be configured to use Windows Internal Database, the database software that is included with WSUS, or the full version of SQL Server. Regardless of which database option the source server is running, perform the following procedures to back up the WSUS database on the source server and restore the database to the destination server.

For an overview of backup and command-line syntax, see the following topics in SQL Server TechCenter:

- Backup Overview

- BACKUP (Transact-SQL)

For an overview of restore and command-line syntax, see the following topics in SQL Server TechCenter:

- Restore and Recovery Overview
- RESTORE (Transact-SQL)

Important
SQL Server Management Studio must be run with elevated administrator permissions throughout this procedure.

To back up the WSUS database on the source server

1. After you connect to the appropriate instance of the database in Object Explorer, click the server name to expand the server tree.

   Note
   If the source server is using Windows Internal Database, the query changes depending on which version of WSUS you are currently running. For WSUS 3.2, the query is: **\\.\pipe\mssql$microsoft##ssee\sql\query**, and for WSUS on Windows Server 2012, the query is: **\\.\pipe\Microsoft##WID\tsql\query**.
2. Expand **Databases**, and select the **SUSDB** database.
3. Right-click the database, point to **Tasks**, and then click **Back Up**. The **Back Up Database** dialog box appears.
4. In the **Database** list, verify the database name.
5. In the **Backup type** list, select **Full**.
6. Select **Copy Only Backup**. A copy-only backup is a SQL Server backup that is independent of the sequence of conventional SQL Server backups.
7. For **Backup component**, click **Database**.
8. Accept the default backup set name that is suggested in the **Name** text box, or enter a different name for the backup set.
9. Optionally, in the **Description** text box, enter a description of the backup set.
10. Specify when the backup set will expire and can be overwritten without explicitly skipping verification of the expiration data.
11. Choose the backup destination by clicking Disk.

    Important
    To remove a backup destination, select it and then click **Remove**. To view the contents of a backup destination, select it and then click **Contents**.
12. In the **Select a page** pane, click **Options** to view or select the advanced options.
13. On the **Overwrite Media** option, click one of the following:
    - **Back up to the existing media set** For this option, click **Append to the existing backup set** or **Overwrite all existing backup sets**. Optionally:
      - Click **Check media set name and backup set expiration** to cause the backup

operation to verify the date and time at which the media set and backup set expire.

- Enter a name in the **Media set name** text box. If no name is specified, a media set with a blank name is created. If you specify a media set name, the media (tape or disk) is checked to see whether the name matches the name that you enter here.

- **Back up to a new media set, and erase all existing backup sets** For this option, enter a name in the **New media set name** text box, and, optionally, describe the media set in the **New media set description** text box.

14. In the **Reliability** section, optionally select:
    - **Verify backup when finished.**
    - **Perform checksum before writing to media.**
    - **Continue on checksum error.**
15. SQL Server 2008 Enterprise support backup compression. By default, whether a backup is compressed depends on the value of the **Backup-compression default** server configuration option. Regardless of the current server-level default, you can compress the backup or prevent compression at this time. To compress the backup, select **Compress backup**, or to prevent compression, select **Do not compress backup**.

To restore the WSUS database backup on the destination server by using SQL Server Management Studio

1. After you connect to the appropriate instance of the database in Object Explorer, click the server name to expand the server tree.

   Important
   If the source server is using Windows Internal Database, the database name is: **\\.\pipe\Microsoft##WID\tsql\query**.
2. Click **New Query** and copy the following SQL command to drop the WSUS database

   USE masterGOALTER DATABASE SUSDB SET SINGLE_USER WITH ROLLBACK IMMEDIATEGODROP DATABASE SUSDBGO
3. Click **Execute**, to run the query
4. Run the following query:

   RESTORE DATABASE [SUSDB] FROM DISK = N'C:\SUSDB.bak' WITH FILE = 1, MOVE N'SUSDB' TO N'c:\Windows\WID\Data\susdb.mdf', MOVE N'SUSDB_log' TO N'c:\Windows\WID\Data\SUSDB_log.ldf', NOUNLOAD, STATS = 10

   Important
   Drive C: that is used in this example will vary according to the actual storage location for the files.
5. In the **Backup type** list, select **Full**.

   Warning
   Running the previous query will result in the following error message:

Msg 3605, Level 16, State 1, Line 5 Schema verification failed for database 'SUSDB'.Msg 3013, Level 16, State 1, Line 5RESTORE DATABASE is terminating abnormally

Disregard the error message and continue.

6. Open an elevated command prompt in Windows Server 2012, and run the following command:

   %programfiles%\update services\tools\wsusutil postinstall [sql parameter] [content parameter]

   Important
   For WID, do not specify the SQL parameter.

To restore the WSUS database backup on the destination server by using SQL Server Management Studio

1. After you connect to the appropriate instance of the database in **Object Explorer**, click the server name to expand the server tree.

   Important
   If the source server is using Windows Internal Database, the database name is *\\.\pipe\Microsoft##WID\tsql\query*.

2. Click **New Query** and copy the following SQL command to drop the WSUS database (SUSDB):

```
USE master

GO

ALTER DATABASE SUSDB SET SINGLE_USER WITH ROLLBACK IMMEDIATE

GO

DROP DATABASE SUSDB

GO
```

3. Click **Execute** to run the query.
4. Run the following query:

```
RESTORE DATABASE [SUSDB] FROM  DISK = N'C:\SUSDB.bak' WITH  FILE = 1,  MOVE
N'SUSDB' TO N'c:\Windows\WID\Data\susdb.mdf',  MOVE N'SUSDB_log' TO
N'c:\Windows\WID\Data\SUSDB_log.ldf',  NOUNLOAD,  STATS = 10
```

   Important
   Drive C: that is mentioned in this example will vary according to the actual storage location for the files.

5. This will result in the following error message:

```
Msg 3605, Level 16, State 1, Line 5

Schema verification failed for database 'SUSDB'.

Msg 3013, Level 16, State 1, Line 5
```

```
RESTORE DATABASE is terminating abnormally
```

Disregard the error message and continue.

6. Open an elevated command prompt in Windows Server 2012, and run the following command:

```
%programfiles%\update services\tools\wsusutil postinstall [sql parameter] [content parameter]
```

Important
For WID, do not specify the SQL parameter.

When a database is restored to a different server, it contains a set of users and permissions, although there may be no corresponding user log on information, or the log on information may not be associated with the same users. This condition is known as having "orphaned users." See article 168644 in the Microsoft Knowledge Base for instructions about how to resolve orphaned users.

After you restore a SQL Server 2005 database to SQL Server 2008, the database becomes available immediately, and it is then automatically upgraded.

3.4. Change the WSUS server identity

The WSUS server identity on the destination server must be changed. Performing this step guarantees that WSUS-managed clients are not affected during the migration process. If the source server and the destination server run with the same identity, and a change is made to one of the servers, the communication between the client and server will fail.

To change the WSUS server identity

1. On the destination server, open a Windows PowerShell session with elevated user rights and run the following script:

```
$updateServer = get-wsusserver

$config = $updateServer.GetConfiguration()

$config.ServerId = [System.Guid]::NewGuid()

$config.Save()
```

2. As soon as the server identity is changed, run the following command to generate a new encryption key:

```
%ProgramFiles%\Update Services\Tools\wsusutil.exe postinstall
```

3.5. Apply security settings

Refer to the settings that you recorded in the Migration Data Collection Worksheet, and then complete the following tasks to apply the security settings that you were using on the source server to the destination server.

- **SMTP server settings**: If you are using an authenticating proxy or the email notification feature to an SMTP server that requires a password (or both), you must manually configure the proxy and email notification to the SMTP server, and enter the SMTP password on the new destination server if you are using email notification.
- **Code signing certificate**: If you are using an advanced management tool that exposes local update publishing (such as Microsoft System Center Essentials 2007 or Microsoft System Center Configuration Manager 2007), copy the code-signing certificate.

To initialize a trust relationship between the update server and its clients, use the following procedures to point the downstream servers to the new WSUS server, and point the WSUS clients to the new WSUS server.

Point the downstream servers to the new WSUS server

If you have downstream servers in your WSUS configuration, and if the server identity on the destination server was changed, perform the following procedure to point them to the new WSUS server.

To connect a downstream server to an upstream server

1. In the navigation pane of the downstream WSUS Server Administration console, click **Options**.
2. Click **Update Source and Proxy Server**, and then click the **Update Source** tab.
3. Select the **Synchronize from another Windows Server Update Services server** check box, and then type the server name and port number in the corresponding text boxes.
4. Repeat step 1 through step 3 for additional downstream servers, if applicable. The synchronization can take several minutes to several hours to finish.
5. To confirm that the downstream servers are synchronizing with the upstream server, in the WSUS Administration Console on the upstream WSUS server, click **Downstream Servers**. In the **Status** pane, confirm that the server's **Last Synchronization** date is after the date that the previous steps were completed.

Point the WSUS clients to the new WSUS server

If the server identity on the destination server was changed, use the following procedure to point the WSUS clients to the new WSUS destination server.

To point the WSUS clients to the new destination server

1. Open the Local Group Policy Editor, and in **Specify intranet Microsoft update service policy**, change the URL to reflect the new WSUS server.
2. Update the Group Policy settings that are used to point WSUS clients to the WSUS server by entering the FQDN of the new WSUS server. After you have updated the Group Policy settings, WSUS clients will synchronize with the new WSUS server.
3. To force the clients to detect the new destination server, open a command prompt, and run `wuauclt.exe /resetauthorization /detectnow`.

Note

To make sure that WSUS clients point to the new WSUS server immediately, you must force detection, which causes WSUS to update computer group membership. If you do not force a detection, it can take up to four hours for clients to point to the new WSUS server.

4. Depending on the number of clients, the initial synchronization can take several minutes to several hours to finish. To confirm that the synchronization is complete, in the WSUS Administration Console, expand **Computers**, and then click **All Computers**. In the **Status** pane, click **Any**, and then click **Refresh**. Confirm that the computers that you expect to see synchronizing to this WSUS server are listed. The **Last Contact Date** has to be refreshed with a post-migration time stamp.

   Tip

   To force a report that the **Last Contact Date** was updated, run `wuauclt.exe /resetauthorization /detectnow`, and then run `wuauclt.exe /reportnow`.

5. After the clients have synchronized, confirm that clients are installing approved updates based on your WSUS configuration settings. In the WSUS Administration Console, click **Reports**, and then click **Computer Tabular Status**. Select the **Report Options** that are applicable to the clients, and then click **Run Report**.
6. To make sure that no WSUS clients are still pointing at the old WSUS server, wait a week and then open the WSUS Administration Console on the old WSUS server. Expand **Computers**, and then click **All Computers**. In the **Status** pane, click **Any**, and then click **Refresh**. Sort on **Last Status Report**. There should be no clients that have a **Last Status Report** date after the date that the synchronization completed.

3.6. Review additional considerations

After the migration is complete, consider the following:

- It is important to have a backup plan for restoring the WSUS server role if there is a migration failure. You do not need to roll back the migration on the source server because the migration process makes no changes to it. You do not need to roll back the migration on the destination server because it is a new server.
- After you have confirmed that no WSUS clients are contacting the old WSUS server, you can uninstall WSUS from the source server. To perform this operation, see the section titled **Retire the WSUS role on the source server (optional)** in the Windows Server Update Services 3.0 SP2 Migration Guide topic: Post-migration Tasks for WSUS.

See also

- Step 4: Verify the WSUS Migration
- Step 2: Prepare to Migrate WSUS
- WSUS server role description

Step 4: Verify the WSUS Migration

The final step in the migration of your Windows Server Update Services (WSUS) server role is to verify that the migration was performed correctly and if the clients can obtain updates from the new WSUS server.

Task	Description
4.1. Verify the destination server configuration	Verify if the destination server is synchronized.
4.2. Verify client computer functionality	Verify if the clients are correctly obtaining updates from the new WSUS server.

4.1. Verify the destination server configuration

Perform the following procedure on the new WSUS destination server to verify that it is configured properly and functioning correctly before you point the WSUS clients and any downstream servers to the new WSUS server.

1. In Server Manager, click **Tools**, and then click **Windows Server Update Services**.
2. In the WSUS Administration Console, expand **Computers**, and verify that all the Computer Groups that existed on the source server are displayed.
3. Expand **Synchronizations**. In the **Actions** pane, click **Synchronize now**. After the synchronization is complete, (this may take several minutes), confirm that **Succeeded** is displayed in the **Results** column.

If the synchronization fails, click **Options**. Confirm that the **Update Source and Proxy Server** settings and password are correct. Confirm that the firewall access is configured correctly for the new server's environment. Make the necessary changes, and then run the synchronization again.

4.2. Verify client computer functionality

Select a client computer so that you can force a detection to verify that the client and server communication is functioning correctly. Use the ollowing procedure to perform this verification:

1. Open a command prompt and type `wuauclt.exe /detectnow` to force the detection.
2. After the detection is finished, open Windows Explorer and check the **%WinDir%\WindowsUpdate.log** to verify that the forced detection was successful.

See also

- Step 3: Migrate WSUS
- WSUS server role description

Migrating Clustered Services and Applications to Windows Server 2012

This guide provides step-by-step instructions for migrating clustered services and applications to a failover cluster running Windows Server 2012 by using the Migrate a Cluster Wizard. Not all clustered services and applications can be migrated using this method. This guide describes supported migration paths and provides instructions for migrating between two multi-node clusters or performing an in-place migration with only two servers. Instructions for migrating a highly available virtual machine to a new failover cluster, and for updating mount points after a clustered service migration, also are provided.

Operating system requirements for clustered roles and feature migrations

The Migrate a Cluster Wizard supports migration to a cluster running Windows Server 2012 from a cluster running any of the following operating systems:

- Windows Server 2008 with Service Pack 2 (SP2)
- Windows Server 2008 R2 with Service Pack 1 (SP1)
- Windows Server 2012

Migrations are supported between different editions of the operating system (for example, from Windows Server Enterprise to Windows Server Datacenter), between x86 and x64 processor architectures, and from a cluster running a core installation of Windows Server or Microsoft Hyper-V Server to a cluster running a full version of Windows Server.

The following migrations scenarios are not supported:

- Migrations from Windows Server 2003 or Windows Server 2003 R2 to Windows Server 2012 are not supported. You should first upgrade to Windows Server 2008 R2 SP1 or Windows Server 2008 SP2, and then migrate the resources to Windows Server 2012 using the steps in this guide.
- The Migrate a Cluster Wizard does not support migrations from a Windows Server 2012 failover cluster to a cluster with an earlier version of Windows Server.

Important

Before you perform a migration, you should install the latest updates for the operating systems on both the old failover cluster and the new failover cluster.

Target audience

This migration guide is designed for cluster administrators who want to migrate their existing clustered services and applications that are running on an existing failover cluster to a Windows Server 2012 failover cluster. The focus of the guide is the steps required to successfully migrate

the resources from one cluster to another by using the Migrate a Cluster Wizard in Failover Cluster Manager.

General knowledge of how to create a failover cluster, configure storage and networking, and deploy and manage the clustered roles and features is assumed.

It is also assumed that customers who will use the Migrate a Cluster Wizard to migrate highly available virtual machines have a basic knowledge of how to create, configure, and manage highly available Hyper-V virtual machines.

What this guide does not provide

The scenarios in this guide provide step-by-step instructions for using the Migrate a Cluster Wizard in Failover Cluster Manager to perform a standard migration of a clustered service or application to a Windows Server 2012 failover cluster. Although this guide identifies services and applications that require special handling during a wizard-based migration, the guide does not provide specific instructions for migrating individual clustered services and applications, including special requirements and dependent server roles and features. For information about migration requirements for specific server roles and features, see Migrate Roles and Features to Windows Server.

This guide does not provide instructions for migrating clustered services and applications by any means other than by using the Copy Cluster Roles Wizard.

Planning considerations for migrations between failover clusters

As you plan a migration to a failover cluster running Windows Server 2012, consider the following:

- Microsoft supports a failover cluster solution for Windows Server 2012 only if all the hardware devices are marked as "Certified for Windows Server 2012." In addition, the complete configuration (servers, network, and storage) must pass all tests in the Validate a Configuration Wizard, which is included in the Failover Cluster Manager snap-in. For more information, see Validate Hardware for a Failover Cluster.
- Hardware requirements are especially important if you plan to continue to use the same servers or storage for the new cluster that the old cluster used. When you plan the migration, you should check with your hardware vendor to ensure that the existing storage is certified for use with Windows Server 2012. For more information about hardware requirements, see Failover Clustering Hardware Requirements and Storage Options.
- The Migrate a Cluster Wizard assumes that the migrated role or feature will use the same storage that it used on the old cluster. If you plan to migrate to new storage, you must copy or move of data or folders (including shared folder settings) manually. The wizard also does not copy any mount point information used in the old cluster. For information about handling mount points during a migration, see Cluster Migrations Involving New Storage: Mount Points.

- Not all clustered services and features can be migrated to a Windows Server 2012 failover cluster by using the Migrate a Cluster Wizard. To find out which clustered services and applications can be migrated by using the Migrate a Cluster Wizard, and operating system requirements for the source failover cluster, see Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012.

Migration scenarios that use the Migrate a Cluster Wizard

When you use the Migrate a Cluster Wizard for your migration, you can choose from a variety of methods to perform the overall migration. This guide provides step-by-step instructions for the following two methods:

- **Create a separate failover cluster running Windows Server 2012 and then migrate to that cluster**. In this scenario, you migrate from a multi-node cluster running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012. For more information, see Migration Between Two Multi-Node Clusters.
- **Perform an in-place migration involving only two servers**. In this scenario, you start with a two-node cluster that is running Windows Server 2008 R2 or Windows Server 2008, remove a server from the cluster, and perform a clean installation (not an upgrade) of Windows Server 2012 on that server. You use that server to create a new one-node failover cluster running Windows Server 2012. Then you migrate the clustered services and applications from the old cluster node to the new cluster. Finally, you evict the remaining node from the old cluster, perform a clean installation of Windows Server 2012 and add the Failover Clustering feature to that server, and then add the server to the new failover cluster. For more information, see In-Place Migration for a Two-Node Cluster.

This guide also provides step-by step instructions that describe how to migrate highly available virtual machines as part of a wizard-based migration. For requirements and process steps, see Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard.

Note

We recommend that you test your migration in a test lab environment before you migrate a clustered service or application in your production environment. To perform a successful migration, you need to understand the requirements and dependencies of the service or application and the supporting roles and features in Windows Server in addition to the processes that this migration guide describes.

In this guide

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012

Migration Between Two Multi-Node Clusters

In-Place Migration for a Two-Node Cluster

Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard

Cluster Migrations Involving New Storage: Mount Points

Additional References

Related references

What's New in Failover Clustering in Windows Server 2012

Failover Clustering Overview

Failover Clustering Hardware Requirements and Storage Options

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012

This topic provides guidance for migrating specific clustered services and applications to a failover cluster running the Windows Server 2012 operating system by using the Migrate a Cluster Wizard in Failover Cluster Manager. The topic covers supported migration paths, provides an overview of wizard-based migration, and notes which clustered services and applications require special handling during migration.

Migration paths for specific migrations

The following table lists the operating system versions on a source failover cluster that can be migrated to a failover cluster running Windows Server 2012 for each clustered service or application. Migrations between failover clusters created with physical computers and failover clusters that are created from virtual machines (also known as a *guest cluster*) are supported.

Supported migrations for clustered roles and resources to a Windows Server 2012 failover cluster

Clustered role or resource	From Windows Server 2008 SP2	From Windows Server 2008 R2 SP1	From Windows Server 2012
Cluster Registry settings	Yes	Yes	Yes
Cluster Shared Volumes (CSV)	No	Yes	Yes
DFS Namespace (DFS-N)	Yes	Yes	Yes
DFS Replication (DFS-R)	No	Yes	Yes
DHCP Server	Yes	Yes	Yes
Distributed Network	No	No	Yes

Clustered role or resource	From Windows Server 2008 SP2	From Windows Server 2008 R2 SP1	From Windows Server 2012
Name (DNN)			
File Server	Yes	Yes	Yes
Scale-out File Server for application data	No	No	Yes
Generic Application	Yes	Yes	Yes
Generic Script	Yes	Yes	Yes
Generic Services	Yes	Yes	Yes
Virtual Machines	Yes	Yes	Yes
Hyper-V Replica Broker	No	No	Yes
IP addresses (IPV4, IPV6, IPv6 tunnel addresses)	Yes	Yes	Yes
iSCSI Target Server	No	Yes	Yes
Internet Storage Name Service (iSNS)	Yes	Yes	Yes
Message Queuing (MSMQ), MSMQ triggers	Yes	Yes	Yes
Microsoft Distributed Transaction Coordinator (MSDTC)	Yes	Yes	Yes
Network Name resources	Yes	Yes	Yes
NFS shares	Yes	Yes	Yes
Other Server	Yes	Yes	Yes
Physical Disk resource	Yes	Yes	Yes
WINS Server	Yes	Yes	Yes

Cluster roles that cannot be migrated

Some services and applications that can run in a failover cluster on Windows Server 2012 cannot be migrated by using the Migrate a Cluster Wizard—in some cases because they were not supported on earlier versions of clustering. The Migrate a Cluster Wizard in Windows Server 2012 cannot be used to migrate the following clustered roles:

- Microsoft SQL Server
- Microsoft Exchange Server
- Print Spooler – In Windows Server 2012, the print spooler is no longer a clustered resource. Instead, high availability is defined as a highly available virtual machine running on a single cluster node. The Print Server role is installed on a single virtual machine, which can be migrated to other nodes automatically or manually. For more information, see High Availability Printing Overview.
- DFS Replication (DFS-R) from Windows Server 2008 - Migration from Windows Server 2008 R2 or Windows Server 2012 is supported, but not from Windows Server 2008.
- Remote Desktop Connection Broker - In Windows Server 2012, the active/passive clustering model for the RD Connection Broker role service, used in earlier versions of Windows Server, is replaced by the Active/Active Broker feature, which eliminates the need for clustering and provides a fully active/active model. For more information, see the blog entry RD Connection Broker High Availability in Windows Server 2012.
- Volume Shadow Copy Service (VSS) tasks
- Task Scheduler tasks (Windows Server 2012 only)
- Cluster Aware Updating (CAU) settings (Windows Server 2012 only)

Roles restricted to a single instance per cluster

For the following roles, only one instance per failover cluster is supported:

- DHCP Server
- WINS Server
- iSCSI Target Server
- Hyper-V Replica Broker (Windows Server 2012 only)

For those roles, the Migrate a Cluster Wizard will not attempt to create a second role instance if one instance already exists on the target cluster.

Migrations for which the Migrate a Cluster Wizard performs most or all steps

For the following clustered services or applications, The Migrate a Cluster Wizard performs most or all steps for a migration to a Windows Server 2012 failover cluster:

- Distributed File System (DFS) Namespace
- Generic Application

- Generic Script
- Generic Service
- IPv4 Address, when migrating within the same subnet
- IPv6 Address or IPv6 Tunnel Address
- Internet Storage Name Service (iSNS)
- Network Name (other than the cluster name)

  If Kerberos authentication is enabled for the Network Name resource, the migration wizard prompts you for the password for the Cluster service account that is used by the old cluster.
- NFS
- Physical Disk (resource settings only; does not copy data to new storage)
- Windows Internet Name Service (WINS) (Extra steps might be required if you migrate to new storage, and you use a different drive letter on the path to the new database.)

For step-by-step instructions for performing a migration between two multimode failover clusters, see Migration Between Two Multi-Node Clusters. For step-by-step instructions for performing a stand-alone migration while upgrading a single failover cluster, see In-Place Migration for a Two-Node Cluster.

Migration within mixed environments

The Migrate a Cluster Wizard can migrate clustered resources within mixed environments. For example, the wizard accommodates the following differences in the source and destination environments:

- Migrate static IP addresses to a cluster using DHCP.
- Migrate IPv4 resources into an IPv6 environment.
- Migrate across routed subnets.
- Migrate a physical cluster to a guest (virtual) cluster (with the exception of Hyper-V clusters, which must run on physical computers).
- Migrate between different editions of the operating system (for example, from Windows Server Enterprise to Windows Server Datacenter), between x86 and x64 processor architectures, and from a cluster running ona Server Core installation or from Microsoft Hyper-V Server to a cluster running a full installation of Windows Server.

During migration, the wizard allows you to address name conflicts between resource groups, resources, and share names and to address drive letter collisions. The wizard resolves the conflicts as part of the post-migration repair process.

Important

The Migrate a Cluster Wizard moves resources, not data. If you plan to migrate to new storage, you must move the data and folders yourself.

Additional steps for a wizard-based migration

Some additional steps typically are needed before or after you run the wizard, including the following:

- Install server roles and features that are needed in the new cluster. In most cases, you must install the role or feature on all nodes of the cluster.
- Copy or install any associated applications, services, or scripts on the new cluster (all nodes).
- If a migrated role or feature uses the same storage, take the services and storage offline on the old cluster and then make the storage available to the new cluster.
- If a migrated role or feature uses new storage, ensure that any data and folders are copied to new storage. Verify permissions on any shared subfolders that were migrated.
- If the new cluster is on a different subnet, provide static IP addresses.
- If the new cluster uses a different volume letter, update drive path locations for applications.
- Configure Task Manager tasks on the new cluster. (Windows Server 2012 only)
- For a virtual machine, install the latest integration services on the new virtual machine. Configure Volume Shadow Copy Service (VSS) backups. For a Windows Server 2012 migration, configure Hyper-V Replica settings.
- Configure Cluster Aware Updating (CAU). (Windows Server 2012 only)

Migration reports

The wizard provides both a Pre-Migration Report and a Post-Migration Report, which provide important information. We recommend that you review both reports while performing a migration:

- The Pre-Migration Report explains whether each resource that you plan to migrate is eligible for migration.
- The Post-Migration Report contains information about the success of the migration, and describes additional steps that might be needed before you bring the migrated resources online.

Note

Two resource groups are never migrated: **Cluster Core Resources Group** and **Available Storage Group**. You can ignore these resource groups in the migration reports.

Clustered role and feature migrations that require extra steps

This section provides guidance for migrating clustered roles and features that require additional steps before or after you run the Migrate a Cluster Wizard to perform a cluster migration.

- Clustered DFS Replication migrations
- Clustered DHCP migrations
- Clustered DTC migrations
- Clustered File Server and Scale-out File Server migrations

- Clustered FSRM migrations
- Clustered Message Queuing (MSMQ) migrations
- Other Server migrations involving resource types not built into failover clusters
- Clustered virtual machine migrations

Clustered DFS Replication migrations

Before you migrate clustered Distributed File System (DFS) Replication (also known as DFS-R or DFSR) to a cluster running Windows Server 2012, you must add the new cluster to the DFS replication group to which the old cluster belongs, and then wait until DFS Replication synchronizes the data to the new cluster. After data synchronization is complete, you can decommission the old cluster. For step-by-step guidance, see File and Storage Services: Prepare to Migrate and File and Storage Services: Post-migration Tasks.

To migrate clustered instances of DFS Replication between clusters running Windows Server 2012

1. Obtain the name of the cluster to which you will migrate. In Active Directory, this is the name that is used for the computer account of the cluster itself (also called the cluster name object or CNO). Add this name to the replication group that you will migrate. For more information, see Add a member to a replication group.
2. Wait until DFS Replication finishes synchronizing the replicated data to the cluster to which you will migrate.
3. If you plan to decommission the cluster from which you migrated, remove its network name from the replication group. If necessary, destroy the cluster.

For more information about DFS Replication in Windows Server 2012, see DFS Namespaces and DFS Replication Overview.

Clustered DHCP migrations

When migrating clustered Dynamic Host Configuration Protocol (DHCP) to a cluster running Windows Server 2012, the Migrate a Cluster Wizard migrates resources and settings, but not the DHCP database. For information about how to migrate the DHCP database, see DHCP Server Migration: Migrating the DHCP Server Role. The information in the topic also applies to migrations from Windows Server 2008 R2 to Windows Server 2012. The topic includes information about migrating from a cluster.

Note

Although the migration of the clustered DHCP role is supported, in Windows Server 2012 there is the option to use DHCP failover. DHCP failover provides redundancy and load balancing without clustered DHCP. For more information, see Migrate to DHCP Failover and Understand and Deploy DHCP Failover.

Clustered DTC migrations

Before you begin the migration of clustered Distributed Transaction Coordinator (DTC) to a cluster running Windows Server 2012, you must make sure the list of transactions stored by DTC is empty. This is referred to as *draining the transaction logs*. If you do not drain the logs, the information in the logs (the transaction state information for unresolved transactions) will be lost during the migration. Unresolved transactions include **Active**, **In Doubt**, and **Cannot Notify** transactions.

To drain DTC transaction logs of unresolved transactions

1. Stop the application that creates transactions on the clustered instance of DTC that is being migrated.
2. On a node of the cluster that you are migrating from, click **Start**, point to **Administrative Tools**, and then click **Component Services**. (In Windows Server 2012, open **Component Services** directly from the Start screen.)
3. Expand **Component Services**, expand **Computers**, expand **My Computer**, expand **Distributed Transaction Coordinator**, and then expand **Clustered DTCs**.
4. Expand the clustered instance of DTC that you are migrating, and then click **Transaction List**.
5. View the transaction list to see if it is empty. If there are transactions listed, then either wait for them to be completed or right-click each transaction, click **Resolve**, and then select **Forget**, **Commit**, or **Abort**.

   For information about the effect of each of these options, see Transaction State Resolution After System Failure.

For additional information, see View Transaction Information.

Clustered File Server and Scale-out File Server migrations

You can migrate a clustered file server from Windows Server 2008 R2 or Windows Server 2008 to a failover cluster running Windows Server 2012 by using the Migrate a Cluster Wizard.

Clustered file server migrations

If you plan to migrate to new storage, keep in mind that if the migrated files and folder inherit permissions from their parents, during migration it is the inheritance setting that is migrated, not the inherited permissions. Therefore it is important to make sure that the parent folders on the source server and the destination server have the same permissions to maintain the permissions on migrated data that has inherited permissions. After the file server migration, it's important to verify the folder permissions after the migration. Sometimes folder permissions reset to Read-only during a file server migration.

You do not need to migrate the quorum resource. When you run the Create a Cluster Wizard in Windows Server 2012, the cluster software automatically chooses the quorum configuration that

provides the highest availability for your new failover cluster. You can change the quorum configuration on the new cluster if necessary for your specific environment.

Scale-out File Server migrations

The Scale-out File Server feature was introduced in Windows Server 2012. You can use the Migrate a Cluster Wizard to migrate a scale-out file server from one Windows Server 2012 failover cluster to another Windows Server 2012 failover cluster.

- The new failover cluster must use the same storage that the old cluster used.
- When you prepare for the migration, after you add the File Server role to each cluster node, you must configure the File Server role as the **Scale-out File Server for application data** role type. For more information, see Deploy Scale-out File Server.

Clustered FSRM migrations

To migrate the File Server Resource Manager (FSRM) classification, storage reporting, and file management task configuration on a clustered file server running Windows Server 2008 R2 or Windows Server 2008 to a failover cluster running Windows Server 2012, you must export the configuration from one FSRM server node in the cluster and then import the configuration to another FSRM server. These steps must be performed locally on one node of the cluster. You then fail over the other nodes until this process is complete. For step-by-step instructions, see Migrate File and Storage Services to Windows Server 2012.

Important

When you migrate the configuration, FSRM requires that you use the same drive letters on both the source and destination servers.

Clustered Message Queuing (MSMQ) migrations

When you migrate a clustered instance of Message Queuing (also known as MSMQ) to a cluster running Windows Server 2012, it's important to take the following precautions to ensure that the data is preserved and you can bring the service online on the new cluster:

- Before you migrate, you should back up the data that is associated with clustered instances of Message Queuing. This ensures that you can restore service-specific Message Queuing data if it is accidentally deleted during migration. For more information about Message Queuing backup and restore, see Backing up and restoring messages.
- During the migration, it's important to make sure that the migration is complete before you delete either clustered instance of Message Queuing (old or new). Otherwise, service-specific data for Message Queuing might be deleted from the shared storage, which prevents the remaining Message Queuing resource from coming online. After the migration is complete and you are ready to delete a clustered instance of Message Queuing (old or new), first remove the disk resource from that clustered instance and take the disk offline. Then delete the clustered instance of Message Queuing.

Other Server migrations involving resource types not built into failover clusters

Before you use the Migrate a Cluster Wizard to migrate an application that uses a clustered resource type that is not built into failover clustering, be sure to add the resource type to the new cluster. You can then use the Migrate a Cluster Wizard to migrate your clustered application. In this situation, the Migrate a Cluster Wizard attempts a "best effort" migration.

To add a resource type to a failover cluster running Windows Server 2012

1. Open Failover Cluster Manager from the Start screen of any node in the cluster running Windows Server 2012.
2. If the cluster to which you want to migrate is not displayed, in the console tree, right-click **Failover Cluster Manager**, click **Connect to Cluster**, select the cluster that you want to migrate to, and then click **OK**.
3. In the console tree, right-click the cluster, and then click **Properties**.
4. Click the **Resource Types** tab, and then click **Add**.
5. Specify the following information for the resource type:
   - **Resource DLL path and file name**: The path and file name of the resource dynamic-link library (DLL) that the Cluster service should use when it communicates with your service or application.
   - **Resource type name**: The name that the Cluster service uses for the resource type. This name stays the same regardless of the regional and language options that are currently selected.
   - **Resource type display name**: The name that is displayed for the resource type. This name might vary when you make changes to regional and language options.

Clustered virtual machine migrations

You can use the Migrate a Cluster Wizard to migrate highly available virtual machines deployed using Hyper-V from a Windows Server 2008 R2 or Windows Server 2008 failover cluster to a cluster running Windows Server 2012. Using the wizard, you can migrate the Virtual Machine clustered role, select highly available virtual machines to migrate, and update virtual network settings for the virtual machines on the new cluster.

Migrating a highly available virtual machine requires some additional steps:

- You must merge or discard all shadow copies before you migrate the volume that contains the virtual machines. You should back up the volumes before you begin merging or discarding shadow copies.
- If you migrate one virtual machine that is stored on Cluster Shared Volume (CVS) volume, the Migrate a Cluster Wizard migrates all virtual machines on that volume. This restriction does not apply if you are migrating a Scale-out File Server cluster between Windows Server 2012 failover clusters. The Scale-out File Server cluster does not use CSV volumes, so you can migrate one virtual machine at a time.

- After you migrate the virtual machines, you must install the latest integration services on the new virtual machines.
- The wizard does not migrate Hyper-V Replica settings or Volume Shadow Copy Service (VSS) tasks. If you are using these with your virtual machines, you must configure them on the new cluster after the migration.

For step-by-step guidance, see Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard.

Additional references

- Migrating Clustered Services and Applications to Windows Server 2012
- Migration forum
- What's New in Failover Clustering in Windows Server 2012
- Failover Clustering Overview
- Migrating Roles and Features in Windows Server
- Migrate File and Storage Services to Windows Server 2012
- Instructions for completing failover cluster migration scenarios:
  - Migration Between Two Multi-Node Clusters
  - In-Place Migration for a Two-Node Cluster
  - Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard
  - Cluster Migrations Involving New Storage: Mount Points
- High availability for Microsoft Exchange Server 2013: Deploying High Availability and Site Resilience
- High availability for Microsoft SQL Server 2014: High Availability Solutions (SQL Server)
- High availability for Microsoft SQL Server 2012: Microsoft SQL Server AlwaysOn Solutions Guide for High Availability and Disaster Recovery (whitepaper)

Migration Between Two Multi-Node Clusters

This topic provides step-by-step instructions for migrating clustered services and applications from a multi-node failover cluster running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to a multimode cluster running Windows Server 2012. (Alternatively, you can perform an in-place migration using a single two-node cluster. For more information, see In-Place Migration for a Two-Node Cluster.) If you plan to migrate highly available Hyper-V virtual machines (by migrating the clustered Virtual Machine role), see the Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard for additional instructions.

Important

Before you begin your migration, review Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012 to confirm that the clustered service or application can be migrated by using the Migrate a Cluster Wizard.

Overview of migration between two multi-node clusters

A migration between two multi-node clusters uses the Migrate a Cluster Wizard, and it has three phases:

1. **Install two or more new servers, run validation, and create a new cluster**. For this phase, while the old cluster continues to run, perform a clean installation of Windows Server 2012 and the Failover Clustering feature on at least two servers. Create the networks that the servers will use, and connect the storage. Make an appropriate number of logical unit numbers (LUNs) or disks accessible to the servers, and do not make those LUNs or disks accessible to any other servers. Next, run the complete set of cluster validation tests to confirm that the hardware and hardware settings can support a failover cluster. Finally, create the new cluster. At this point, you have two clusters.

   For more information, see Steps for creating a failover cluster, later in this topic.

2. **Migrate clustered services and applications to the new cluster, and determine how you will make any existing data available to the new cluster**. When the Migrate a Cluster Wizard completes, all the migrated resources will be offline. Leave them offline at this stage. If the new cluster will reuse old storage, plan how you will make the storage available to the new cluster, but leave the old cluster connected to the storage until you are ready to make the transition.

   For more information, see Steps for migrating clustered services and applications to a failover cluster running Windows Server 2012, later in this topic.

3. **Make the transition from the old cluster to the new cluster**. The first step in the transition is to take the clustered services and applications offline on the old cluster. If the new cluster will use old storage, follow your plan for making LUNs or disks inaccessible to the old cluster and accessible to the new cluster. If the new cluster will use new storage, copy the appropriate folders and data to the storage. Bring the clustered services and applications online on the new cluster. Then verify that failover is working and the clustered services and applications are available.

   For more information, see Steps for completing the transition from the old cluster to the new cluster, later in this topic.

Steps for creating a failover cluster

For information about how to create a Windows Server 2012 failover cluster, see Create a Failover Cluster. To prepare to migrate a clustered service or application to the new failover cluster, make the following preparations.

Preparation

Before you create the failover cluster, prepare storage, and install all required services, applications, and server roles.

1. Prepare storage:
    a. Make an appropriate number of LUNs or disks accessible to the servers, and do not make those LUNs or disks accessible to any other servers. If the new cluster will use old storage, for testing purposes, you can limit the number of LUNs or disks to one or two. If the new cluster will use new storage, make as many disks or LUNs accessible to the new server as you think the cluster will need.

    Note
    We recommend that you keep a small disk or LUN available (unused by clustered services and applications) throughout the life of the cluster, so that you can always run storage validation tests without taking your services and applications offline.

    b. On one of the servers that you plan to include in the cluster, open **Computer Management** from the **Start** screen, and then click **Disk Management** in the console tree. In Disk Management, confirm that the intended cluster disks are visible.

    c. Check the format of any exposed volume or LUN. We recommend that you use NTFS for the format. (For a disk witness, you must use NTFS.)

    d. If you are using new storage and your disk configuration uses mount points, review Cluster Migrations Involving New Storage: Mount Points to identify any additional steps you will need to perform.

2. Install services, applications, and server roles:
    - After you install the Failover Clustering feature on all nodes, install any needed services, applications, and server roles. For example, if you plan to migrate clustered Windows Internet Name Service (WINS) to the new cluster, install the WINS Server feature by using Server Manager.
    - If you plan to migrate highly available virtual machines, add the Hyper-V role to the server. You also must merge or discard all shadow copies on the volumes that contain the virtual machines. For step-by-step instructions for migrating highly available virtual machines, see Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard.
    - If you are migrating a Generic Application, Generic Script, or Generic Service resource, you must confirm that any associated application is compatible with Windows Server 2012. You also must confirm that any associated service exists in Windows Server 2012 and has the same name that it had in the old cluster. Test the application or service (separately, not as part of a cluster) to confirm that it runs as expected.

After you create the failover cluster

After you create the cluster, ensure that your firewall is configured appropriately. For example, if you are using Windows Firewall, and you will be sharing folders and files, use your preferred Windows Firewall interface to allow the exception for **Remote Volume Management**.

Steps for migrating clustered services and applications to a failover cluster running Windows Server 2012

Use the following instructions to migrate clustered services and applications from your old cluster to your new cluster. The Migrate a Cluster Wizard leaves most of the migrated resources offline so that you can perform additional steps before you bring them online.

Note

To migrate a clustered service or application by using the Migrate a Cluster Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

To migrate data and clustered services or applications from an existing cluster to a new cluster

1. If the new cluster uses old storage, plan how you will make LUNs or disks inaccessible to the old cluster and accessible to the new cluster (but do not make changes yet). If you plan to use new storage with the migrated services or applications, before you run the Migrate a Cluster Wizard, make the storage is available to the new cluster – that is, that the volumes have been added to the new cluster and that they are online. This enables the wizard to update storage settings during migration.
2. From the **Start** screen or from **Server Manager** (**Tools**), open **Failover Cluster Manager**.
3. In the console tree, if the cluster that you created is not displayed, right-click **Failover Cluster Manager**, click **Connect to Cluster**, and then select the cluster that you want to configure.
4. In the console tree, expand the cluster that you created to see the items underneath it.
5. If the clustered servers are connected to a network that is not to be used for cluster communications (for example, a network intended only for iSCSI), then under **Networks**, right-click that network, click **Properties**, and then click **Do not allow cluster network communication on this network**. Click **OK**.
6. In the console tree, select the cluster.
7. Under **Configure**, click **Migrate services and applications**.

   The Migrate a Cluster Wizard opens.
8. Read the Welcome page, and then click **Next**.
9. Specify the name or IP address of the cluster or cluster node from which you want to migrate services and applications, and then click **Next**.
10. The **Select Services and Applications** page lists the clustered services and applications that can be migrated from the old cluster. The list does not contain any service or application that is not eligible for migration. Click **View Report** for details. Then select each service and application that you want to migrate to the new cluster, and click **Next**.

Important

We recommend that you read the report, which explains whether each resource is eligible for migration. (The wizard also provides a report after it finishes, which describes any additional steps that might be needed before you bring the migrated resource groups online.)

If storage is available on the new cluster, the **Specify Storage for Migration** page appears, giving you the option to migrate to new storage. If storage is not available on the new cluster, the wizard retains existing storage settings and does not display the page.

Note

Not all clustered roles can be migrated to new storage. For example, the wizard cannot be used to migrate highly available virtual machines (the Virtual Machine role) to new storage. For step-by-step instructions for migrating highly available virtual machines, see Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard.

11. If you want to use new storage for a service or application:
    a. On the **Specify Storage for Migration** page, select the cluster disk that you want to migrate to new storage, and then click **Select Storage**.
    b. In the **Select Storage for Resource Group** dialog box, under **Available Storage in New Cluster**, select the cluster disk that you want the service or application to use in the new cluster, and then click **OK**.
    c. Repeat these steps for each cluster disk that you want to migrate to new storage. Then click **Next**.

    Important

    The Migrate a Cluster Wizard does not move existing data and folders to the new storage. You must copy the folders and data manually.
12. Follow the instructions in the wizard to perform the migration. From the **Summary** page, we recommend that you read the Cluster Migration Report, which contains important information about any additional steps that you might need to complete before you bring the migrated services and applications online. For example, if you have not already installed needed applications on the new cluster node, you might need to install them.

After the wizard completes, most migrated resources will be offline. Leave them offline at this stage.

Steps for completing the transition from the old cluster to the new cluster

You must perform the following steps to complete the transition to the new cluster running Windows Server 2012. After you complete the transition, verify that failover is working correctly for the migrated services and applications and that the services are available.

To complete the transition from the old cluster to the new cluster

1. Prepare for clients to experience downtime, probably briefly.
2. On the old cluster, take each role and resource that was migrated offline.
3. Complete the transition for the storage:
   - If the new cluster will use old storage, follow your plan for making LUNs or disks inaccessible to the old cluster and accessible to the new cluster.
   - If the new cluster will use new storage, copy the appropriate folders and data to the storage. As needed for disk access on the old cluster, bring individual disk resources online on that cluster. (Keep other resources offline, to ensure that clients cannot change data on the disks in storage.) On the new cluster, use Disk Management to confirm that the appropriate LUNs or disks are visible to the new cluster and not visible to any other servers.
4. If the new cluster uses mount points, adjust the mount points as needed, and make each disk resource that uses a mount point dependent on the resource of the disk that hosts the mount point. For more information about mount points, see Cluster Migrations Involving New Storage: Mount Points.
5. Bring the migrated services or applications online on the new cluster.
6. If you migrated highly available virtual machines, install the latest integration services on each virtual machine. You might need to restart the virtual machine to complete the installation.

To verify that the migrated service or application is performing as expected and can fail over successfully

1. Verify that you can access the workload that was migrated. For example, can you connect to a highly available file server after it is migrated? Can you see the data that the server stores?
2. In the console tree of **Failover Cluster Manager**, click the failover cluster on which the service or application is running.
3. Expand **Services and Applications**, and then click a migrated service or application that you want to test.
4. Under **Actions** (on the right), click **Move this service or application to another node**, and then click an available choice of node. When prompted, confirm your choice.

   You can observe the status changes in the center pane of the snap-in as the clustered service or application is moved.
5. If there are any issues with failover, review the following:
   - View events in Failover Cluster Manager. To do this, in the console tree, right-click **Cluster Events**, and then click **Query**. In the **Cluster Events Filter** dialog box, select the criteria for the events that you want to display, or to return to the default criteria, click the **Reset** button. Click **OK**. To sort events, click a heading, for example, **Level** or **Date and Time**.
   - Confirm that necessary services, applications, or server roles are installed on all

nodes. Confirm that services or applications are compatible with Windows Server 2012 and run as expected.

- If you used old storage for the new cluster, rerun the Validate a Cluster Configuration Wizard to confirm the validation results for all LUNs or disks in the storage.
- If you migrated highly available virtual machines, verify the status of the virtual machines in Hyper-V Manager, and ensure that you can connect to the virtual machines by using Remote Desktop or Virtual Machine Connection.
- Review migrated resource settings and dependencies. If you are using new storage that includes disks that use mount points, see Cluster Migrations Involving New Storage: Mount Points.
- If you migrated one or more Network Name resources with the Kerberos protocol enabled, confirm that the following permissions change was made in **Active Directory Users and Computers** on a domain controller. In the computer accounts (computer objects) of your Kerberos protocol-enabled Network Name resources, **Full Control** must be assigned to the computer account for the failover cluster.

Note

The Migrate a Cluster Wizard does not migrate Volume Shadow Copy Service (VSS) tasks, Hyper-V Replica Broker settings, Task Scheduler tasks, and Cluster-Aware Updating (CAU) settings. If you were using any of these features on the old cluster, you will need to configure them on the new cluster.

Related references

In-Place Migration for a Two-Node Cluster

Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard

Cluster Migrations Involving New Storage: Mount Points

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012

In-Place Migration for a Two-Node Cluster

This topic provides an overview and steps for upgrading an existing failover cluster to Windows Server 2012 when you have only two servers - that is, for performing an *in-place migration*.

Important

Before you begin the migration, confirm that the clustered service or application that you want to migrate can be migrated by using the Migrate a Cluster Wizard, as described in Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012, and note any preparation or follow-up steps that are required for the service type that is being migrated.

Note

For an alternative approach to failover cluster migration, see Migration Between Two Multi-Node Clusters.

Overview of an in-place migration for a two-node cluster

This migration uses the Migrate a Cluster Wizard, and it has four phases:

1. **Evict one node, install Windows Server 2012, and create a single-node failover cluster**. For this phase, allow one existing server to continue running Windows Server 2008 R2 or Windows Server 2008 and the Cluster service while you begin the migration process. Evict the other server from the old cluster, and then perform a clean installation of Windows Server 2012 and the Failover Clustering feature on it. On that server, run all tests that the Validate a Configuration Wizard will run. The wizard will recognize that this is a single node without storage and limit the tests that it runs. Tests that require two nodes (for example, tests that compare the nodes or that simulate failover) will not run.

   Note that the tests that you run at this stage do not provide complete information about whether the storage will work in a cluster running Windows Server 2012. As described later in this section, you will run the Validate a Configuration Wizard later with all tests included.

   For steps to complete this phase, see Steps for evicting a node and creating a new single-node Windows Server 2012 failover cluster, later in this topic.

2. **Migrate clustered services and applications to the new single-node cluster**. Run the Migrate a Cluster Wizard, but leave the migrated resources offline on the new cluster.

   For steps to complete this phase, see Steps for migrating clustered services and applications to the new cluster, later in this topic.

3. **Make existing data available to the new cluster, and bring the cluster online**. Confirm that the settings for the migrated services and applications are correct. Next, take the migrated services and applications in the old cluster offline. If the new cluster will use new storage, copy the folders and data to appropriate LUNs or disks in the new storage, and make sure that those LUNs or disks are visible to the new cluster (and not visible to any other servers). If the new cluster will use the old storage, make the appropriate disks or LUNs accessible to the new cluster. Bring the services and applications in the new cluster online, and make sure that the resources are functioning and can access the storage.

   For steps to complete this phase, see Steps for making existing data available to the new cluster and bringing it online, later in this topic.

4. **Add the second node to the new cluster**. Destroy the old cluster and, on that server, install Windows Server 2012 and the Failover Clustering feature. Connect that server to the networks and storage that are used by the new cluster. If the appropriate disks or LUNs are not already accessible to both servers, make them accessible. Run the Validate a Configuration Wizard, specifying both servers, and confirm that all tests pass. Finally, add the second server to the new cluster.

   For information about steps for this phase, see Steps for adding the second node to the new cluster, later in this topic.

Steps for evicting a node and creating a new single-node Windows Server 2012 failover cluster

You must complete the following steps to create a single-node Windows Server 2012 failover cluster:

- Step 1: Evict one node from the old cluster, and perform a clean installation of Windows Server 2012
- Step 2: Create a single-node cluster and install other needed software

Step 1: Evict one node from the old cluster, and perform a clean installation of Windows Server 2012

To begin, you must evict one node from the old cluster, and perform a clean installation of Windows Server 2012 on that node.

Before you evict a node from a cluster

- For each clustered service or application that you plan to migrate, verify that there are no special requirements or procedures for removing or evicting a node from the cluster. You can evict a node from a clustered file server or a cluster with the Hyper-V role with no special preparation. However, you might need to uncluster some services or applications before you evict a node.
- To prevent any loss of application data when the node is evicted, shut down all services and applications on the cluster before you evict the node.

To evict a node from a cluster

1. From the Start screen, open **Failover Cluster Manager**.
2. In the console tree, expand the cluster, expand **Nodes**, and then click the node that you want to evict to select it.
3. Right-click the node, click **More Actions**, and then click **Evict**.

Step 2: Create a single-node cluster and install other needed software

For information about how to create a Windows Server 2012 failover cluster, see Create a Failover Cluster. To prepare to migrate a clustered service or application to the new failover cluster, make the following preparations.

Preparation

Before you create the failover cluster, prepare storage, and install all required services, applications, and server roles.

1. Prepare storage:

a. Make an appropriate number of LUNs or disks accessible to the server, and do not make those LUNs or disks accessible to any other servers. If the new cluster will use old storage, for testing purposes, you can limit the number of LUNs or disks to one or two. If the new cluster will use new storage, make as many disks or LUNs accessible to the new server as you think the cluster will need.

   Note
   We recommend that you keep a small disk or LUN available (unused by clustered services and applications) throughout the life of the cluster, so that you can always run storage validation tests without taking your services and applications offline.

b. On the server, open **Computer Management** from the **Start** screen, and then click **Disk Management** in the console tree. In Disk Management, confirm that the intended cluster disks are visible.

c. Check the format of any exposed volume or LUN. We recommend that you use NTFS for the format. (For a disk witness, you must use NTFS.)

d. If you are using new storage and your disk configuration uses mount points, review Cluster Migrations Involving New Storage: Mount Points to identify any additional steps you will need to perform.

2. Install services, applications, and server roles:
   - After you install the Failover Clustering feature on the server, install any needed services, applications, and server roles. For example, if you plan to migrate clustered Windows Internet Name Service (WINS) to the new cluster, install the WINS Server feature by using Server Manager.
   - If you plan to migrate highly available virtual machines, add the Hyper-V role and install the latest Hyper-V integration components. You also must merge or discard all shadow copies on the volumes that contain the virtual machines. For step-by-step instructions for migrating highly available virtual machines, see Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard.
   - If you are migrating a Generic Application, Generic Script, or Generic Service resource, you must confirm that any associated application is compatible with Windows Server 2012. You also must confirm that any associated service exists in Windows Server 2012 and has the same name that it had in the old cluster. Test the application or service (separately, not as part of a cluster) to confirm that it runs as expected.

After you create the failover cluster

After you create the cluster, ensure that your firewall is configured appropriately. For example, if you are using Windows Firewall, and you will be sharing folders and files, use your preferred Windows Firewall interface to allow the exception for **Remote Volume Management**.

Steps for migrating clustered services and applications to the new cluster

Use the following instructions to migrate clustered services and applications from your old one-node cluster to your new one-node cluster. The Migrate a Cluster Wizard leaves most of the migrated resources offline so that you can perform additional steps before you bring them online.

Note

To migrate a clustered service or application by using the Migrate a Cluster Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

To migrate clustered services and applications from the old cluster to the new cluster

1. If you want to migrate to new storage, before you run the Migrate a Cluster Wizard, ensure that the storage is available to the new cluster – that is, that the volumes have been added to the new cluster and that they are online.
2. From the **Start** screen or **Server Manager** (**Tools**), open **Failover Cluster Manager**.
3. In the console tree, if the cluster that you created is not displayed, right-click **Failover Cluster Manager**, click **Connect to Cluster**, and then select the cluster that you want to configure.
4. In the console tree, expand the cluster that you created to see the items underneath it.
5. If the clustered server is connected to a network that is not to be used for cluster communications (for example, a network intended only for iSCSI), then under **Networks**, right-click that network, click **Properties**, and then click **Do not allow cluster network communication on this network**. Click **OK**.
6. In the console tree, select the cluster.
7. Under **Configure**, click **Migrate services and applications**.
8. Read the first page of the Migrate a Cluster Wizard, and then click **Next**.
9. Specify the name or IP address of the cluster or cluster node from which you want to migrate services and applications, and then click **Next**.
10. The **Select Services and Applications** page lists the clustered services and applications that can be migrated from the old cluster. The list does not contain any service or application that is not eligible for migration. Click **View Report** for details. Then select each service and application that you want to migrate to the new cluster, and click **Next**.

    Important

    We recommend that you read the report, which explains whether each resource is eligible for migration. (The wizard also provides a report after it finishes, which describes any additional steps that might be needed before you bring the migrated resource groups online.)

    If storage is available on the new cluster, the **Specify Storage for Migration** page appears, giving you the option to migrate to new storage. If storage is not available on the

new cluster, the wizard automatically retains existing storage settings and does not display the page.

11. If you want to use new storage for a service or application:
    a. On the **Specify Storage for Migration** page, select the cluster disk that you want to migrate to new storage, and then click **Select Storage**.
    b. In the **Select Storage for Resource Group** dialog box, under **Available Storage in New Cluster**, select the cluster disk that you want the service or application to use in the new cluster, and then click **OK**.
    c. Repeat these steps for each cluster disk that you want to migrate to new storage. Then click **Next**.

    Important
    The Migrate a Cluster Wizard does not move existing folders and data to the new storage. You must copy the folders and data manually.
12. Follow the instructions in the wizard to perform the migration. From the **Summary** page, we recommend that you read the Cluster Migration Report, which contains important information about any additional steps that you might need to complete before you bring the migrated services and applications online. For example, if you have not already installed needed applications on the new cluster node, you might need to install them.

When the wizard completes, most migrated resources will be offline. Leave them offline at this stage.

Steps for making existing data available to the new cluster and bringing it online

Use the following procedure to make existing data available to the new cluster and bring it online.

To make existing data available to the new cluster and bring it online

1. Confirm that the settings for the migrated services and applications appear correct.
2. Prepare for clients to experience downtime, probably briefly.
3. On the old cluster, take each clustered service or application that you migrated offline.
4. Complete the transition for the storage:
    - If the new cluster will use old storage, follow your plan for making LUNs or disks inaccessible to the old cluster and accessible to the new cluster.
    - If the new cluster will use new storage, copy the appropriate folders and data to the storage. As needed for disk access on the old cluster, bring individual disk resources online on that cluster. (Keep other resources offline to ensure that clients cannot change data on the disks in storage.) Then, on the new cluster node, use Disk Management to confirm that the appropriate LUNs or disks are visible to the new cluster and not visible to any other servers.
5. If the new cluster uses mount points, adjust the mount points as needed, and make each disk resource that uses a mount point dependent on the resource of the disk that hosts

the mount point. For more information about mount points, see Cluster Migrations Involving New Storage: Mount Points.

6. Bring the migrated services or applications online on the new cluster.

Steps for adding the second node to the new cluster

Use the following instructions to prepare the second node and then add it to the new cluster. As part of this process, you will run validation tests that include both servers.

To add the second node to the new cluster

1. On the new cluster, confirm that the migrated services or applications are functioning and that clients can connect to them.
2. On the old cluster (the server that is running Windows Server 2008 R2 or Windows Server 2008), delete the migrated services and applications, and then destroy the old cluster:
   a. From the Start screen, open **Failover Cluster Manager**.
   b. Remove services and applications that were migrated. In Failover Cluster Manager, expand the cluster, and expand **Services and applications**. To delete a service or application, right-click the item, and click **Delete**.
   c. To destroy the cluster, right-click the cluster, click **More Actions**, and then click **Destroy Cluster**.
3. On the same server, perform a clean installation of Windows Server 2012.
4. Add the Failover Clustering feature in the same way that you added it to the other server, and install any needed services, applications, and server roles.
5. Connect the newly installed server to the same networks and storage that the existing failover cluster node is connected to.
6. Identify the disks or LUNs that are exposed to the new one-node failover cluster, and expose them to the newly installed server also.

   We recommend that you keep a small disk or LUN accessible to both nodes, and unused by clustered services and applications, throughout the life of the cluster. With this LUN, you can always run storage validation tests without taking your services and applications offline.
7. On either server running Windows Server 2012, open Failover Cluster Manager from the Start screen.
8. Confirm that **Failover Cluster Manager** is selected, and then, in the center pane, under **Management**, click **Validate a Configuration**.

   Follow the instructions in the wizard, but this time, be sure to specify both servers (not just the existing cluster name) and specify that you want to run all tests. Then, run the tests. Because two nodes are now being tested, a more complete set of tests runs, which takes longer than testing one node.

Important

If any clustered service or application is using a disk when you start the wizard, the wizard asks whether to take that clustered service or application offline for testing. If you choose to take a clustered service or application offline, it remains offline until the tests finish.

9. The **Summary** page appears after the tests run. To view Help topics to help you interpret the results, click **More about cluster validation tests**.
10. While still on the **Summary** page, click **View Report** and read the test results.

    To view the results of the tests after you close the wizard, see

    <*SystemRoot*>\Cluster\Reports\Validation Report <*date and time*>.mht

    where <*SystemRoot*> is the folder in which the operating system is installed (for example, C:\Windows\).
11. As necessary, make changes in the configuration and rerun the tests.

    For more information about failover cluster validation tests, see Validate Hardware for a Failover Cluster.
12. If the new cluster is not displayed, in the console tree, right-click **Failover Cluster Manager**, click **Connect to a Cluster**, and then select the new cluster.
13. In the console tree, select the one-node cluster, and then in the **Actions** pane, click **Add Node**.
14. Follow the instructions in the wizard to specify the server that you want to add to the cluster. On the **Summary** page, click **View Report** to review the tasks that the wizard performed.
15. On the **Summary** page, click **View Report** if you want to review the tasks that the wizard performed. Or view the report after the wizard closes in the <*SystemRoot*>\Cluster\Reports\ folder.

    Note

    After you close the wizard, in the center pane, you might see a warning about "Node Majority." You will correct this issue in the next few steps.
16. In the console tree, expand **Storage**. Check to see if all the disks that you want to make available to the new cluster are shown, either in one of the clustered services or applications or in **Available Storage**.

    In most cases, you need at least one disk in **Available Storage** for the next step (specifying a witness disk). If you need to add a disk, in the **Actions** pane, click **Add Disk** and follow the steps in the wizard.

    Before you can add a disk to storage, it must be accessible from both nodes in the cluster. To be used for a witness disk, a disk can be a relatively small, but must be at least 512 MB.
17. In the console tree, right-click the new cluster, click **More Actions**, and then click **Configure Cluster Quorum Settings**.
18. Follow the instructions in the wizard to select the most appropriate quorum setting for

your needs. In most cases, this is the **Node Majority** quorum configuration, which requires that you specify an appropriate disk (from **Available Storage**) for the witness disk. For more information about quorum settings in Windows Server 2012, see Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster.

19. Expand **Services and Applications**, and then click a migrated service or application that you want to test.
20. To perform a basic test of failover for the migrated service or application, under **Actions** (on the right), click **Move this service or application to another node**, and then click an available choice of node. When prompted, confirm your choice.

    You can observe the status changes in the center pane of Failover Cluster Manager as the clustered service or application is moved. If there are any issues with failover, review the following:

    - View events in Failover Cluster Manager. To do this, in the console tree, right-click **Cluster Events**, and then click **Query**. In the **Cluster Events Filter** dialog box, select the criteria for the events that you want to display, or, to return to the default criteria, click the **Reset** button. Click **OK**. To sort events, click a heading, for example, **Level** or **Date and Time**.
    - Confirm that necessary services, applications, or server roles are installed on all nodes. Confirm that services or applications are compatible with Windows Server 2012 and run as expected.
    - Review migrated resource settings and dependencies. If you are using new storage that includes disks that use mount points, see Cluster Migrations Involving New Storage: Mount Points for more information.
    - If you migrated one or more Network Name resources with the Kerberos protocol enabled, confirm that the following permissions change was made in **Active Directory Users and Computers** on a domain controller. In the computer accounts (computer objects) of your Kerberos protocol-enabled Network Name resources, **Full Control** must be assigned to the computer account for the failover cluster.

Related references

Migration Between Two Multi-Node Clusters

Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard

Cluster Migrations Involving New Storage: Mount Points

Migration Paths for Migrating to a Failover Cluster Running Windows Server 2012

Migration of Highly Available Virtual Machines Using the Migrate a Cluster Wizard

This topic provides a process overview and step-by-step instructions for migrating a Hyper-V highly available virtual machine (HAVM) from a failover cluster running Windows Server 2008,

Windows Server 2008 R2, or Windows Server 2012 to a failover cluster running Windows Server 2012 by using the Migrate a Cluster Wizard in Failover Cluster Manager. This is accomplished by migrating the clustered Virtual Machine role from the old cluster to the new cluster. The migrated HAVMs use the same storage that they used in the old cluster. The wizard cannot migrate virtual machines to new storage.

You can use either of the two migration scenarios to migrate an HAVM: migrate between two multi-node clusters or perform an in-place migration.

Note

You can also use this method to migrate HAVMs to a failover cluster running Windows Server 2012 R2 from a failover cluster running Windows Server 2008 R2 with Service Pack 1 (SP1), Windows Server 2012, or Windows Server 2012 R2. In Windows Server 2012 R2, the name of the Migrate a Cluster Wizard was changed to **Copy Cluster Roles**, and the wizard is opened by using the **Copy Roles** action. For consistency with labeling in Windows Server 2012 R2, the items being migrated are referred to as clustered roles instead of clustered services and applications. However, the steps for performing the wizard-based migration are the same.

Supported operating systems

The Migrate a Cluster Wizard in Windows Server 2012 can migrate highly available virtual machines running on any of the following Windows Server operating system versions to Windows Server 2012:

- Windows Server 2008 with Server Pack 2 (SP2)
- Windows Server 2008 R2 with Service Pack 1 (SP1)
- Windows Server 2012

Before you migrate any clustered role or service, you should install the latest operating system updates on all nodes in the old and new clusters.

Overview of the migration process

To migrate a highly available virtual machine from one failover cluster to another, you use the Migrate a Cluster Wizard in Failover Cluster Manager to migrate the clustered Virtual Machine role. After you select the Virtual Machine role, you select the role instances (virtual machines) that you want to migrate.

Note

Be aware that if you migrate one virtual machine that resides on a Cluster Shared Volume (CSV) volume, the wizard migrates all virtual machines on that volume. The wizard allows you to update virtual network settings on the virtual machines to the network settings on the new cluster. You cannot use the wizard to migrate virtual machines to new storage. This restriction does not apply if you are migrating a Scale-out

File Server cluster. A Scale-out File Server cluster does not use CSV volumes, so you can migrate one virtual machine at a time.

To prepare the virtual machines for the migration, you must merge or discard all shadow copies on the volumes that contain the virtual machines. To prepare the new cluster for the migration, you must add the Hyper-V role to the cluster nodes and configure storage and virtual networks on the cluster.

After the migration, you will need to take the virtual machines offline on the old cluster, follow your plans to mask the volumes that contain the virtual machines to the old cluster and unmask the volumes to the new cluster, and then bring the virtual machines online on the new cluster. After you bring the virtual machines on the new cluster online, you must also install the latest integration services on the virtual machines.

To schedule local backups of the virtual machines, you will need to configure Volume Shadow Copy Service (VSS) tasks, which the wizard does not migrate. If you migrated from a Windows Server 2012 failover cluster, you will also need to configure Hyper-V Replica Manager settings and Cluster-Aware Updates (CAU) if you were using them; those settings also do not migrate.

Tip

For a step-by-step walk-through, with screenshots, of migrating a Hyper-V host cluster from Windows Server 2008 R2 to Windows Server 2012 by using the Migrate a Cluster Wizard, see the blog entry How to Move Highly Available (Clustered) VMs to Windows Server 2012 with the Cluster Migration Wizard on MSDN.

Impact of the migration

There will be a brief service interruption during the migration. To minimize the effects on users, schedule the migration during a maintenance window. We also recommend that you pretest and verify the migration before you migrate the virtual machines in your production environment.

Note

You cannot use live migration to migrate a highly available virtual machine to a new failover cluster.

Required permissions

To migrate a clustered service or application by using the Migrate a Cluster Wizard, you must be a local administrator on the destination failover cluster and on the cluster or cluster node from which you are migrating.

Prepare to migrate

While you prepare to migrate the virtual machines, the virtual machines can remain online and continue providing service.

To prepare virtual machine storage for migration

1. Before you begin working with shadow copies, you should back up all volumes that are attached to the virtual machine(s).
2. Merge or discard all shadow copies for the volumes that store the virtual machines.
3. Ensure that no virtual machines that you do not want to migrate share a CSV volume with virtual machines that you plan to migrate. If you migrate one virtual machine on a CSV volume, the Migrate a Cluster Wizard migrates all virtual machines on that volume.

To prepare the old failover cluster for the migration

- Install the latest operating system updates on each cluster node. A Windows Server 2008 failover cluster must be running Windows Server 2008 SP2 or later. A Windows Server 2008 R2 failover cluster must be running Windows Server 2008 R2 SP1 or later.

To prepare the new failover cluster for the migration

1. To create the new failover cluster in Windows Server 2012, use Failover Cluster Manager, or use the **New-Cluster** cmdlet in Windows PowerShell (for information, see New-Cluster). For a detailed description of the steps for preparing a new failover cluster, see Migration Between Two Multi-Node Clusters or In-Place Migration for a Two-Node Cluster.
2. Add the Hyper-V role to each cluster node.
3. Configure virtual switches in Hyper-V.
4. If you are migrating from a Windows Server 2008 or Windows Server 2008 R2 failover cluster, check with your hardware vendor to ensure that the existing storage is supported in Windows Server 2012.

   Note
   You do not need to configure storage for the virtual machines on the new cluster before you run the wizard. The Migrate a Cluster Wizard will migrate existing storage settings to the new cluster. After the wizard completes, you will mask the storage from the old cluster and then unmask the storage on the new cluster.
5. Install the latest operating system updates on each cluster node.

Migrate the highly available virtual machines to the new failover cluster

You will use the Migrate a Cluster Wizard in Failover Cluster Manager to migrate the virtual machines to the new failover cluster. This is done by migrating the clustered Virtual Machine role. After the virtual machines are migrated, you must make the storage available to the new cluster before you bring the virtual machines online.

Before you migrate the highly available virtual machines

- Ensure that the virtual switches are configured on host operating systems in the new cluster.
- Prepare for a brief service interruption on the workloads running on the virtual machines. Live migration is not supported during migration of a virtual machine to a new host cluster.

To migrate the virtual machines to a new failover cluster

1. Log on to any node in the new failover cluster with an Administrator account.
2. From the **Start** screen or **Server Manager (Tools)**, open **Failover Cluster Manager**.
3. In the console tree, expand **Failover Cluster Manager** in the console tree, and select the cluster that you want to migrate the virtual machines to.

   If the new cluster is not displayed, right-click **Failover Cluster Manager**, click **Connect to Cluster**, and then select the cluster to which you want to migrate the virtual machines.
4. With the destination cluster selected, click **Migrate Roles**.

   The Migrate a Cluster Wizard opens.
5. Review the instructions on the **Before You Begin** page, and click **Next**.
6. On the **Specify Old Cluster** page, enter the name or IP address of the source cluster, or use the **Browse** button to find the cluster, and then click **Next**.

   The wizard connects to the cluster and displays the roles and features that can be migrated. For virtual machines, each virtual machine “role” is listed under the cluster shared volume that stores the virtual machine.
7. On the **Select Services and Applications** page, click **View Reports**, and review the resources that can and cannot be migrated. Note that **Available Storage** and **Cluster Group** are never available for migration, and always have a Failed result. For all other resources, review any Warning results to identify any and resolve any issues that might prevent a successful migration. Then close the report.
8. On the **Select Services and Applications** page, select each highly available virtual machine that you want to migrate, and then click **Next**. If a volume stores more than one virtual machine, and you select any virtual machine on that volume, the wizard will migrate all virtual machines on that volume.
9. On the **Customize Virtual Machine Networks** page, optionally select a select virtual switch for the virtual machines to use on the destination host cluster. If you do not select a virtual switch, the wizard retains the default switch that is selected automatically the first time the virtual machine starts on its new host.
10. On the **Configuration** page, review your settings. Then click **Next** to start the migration.
11. After the migration completes, review the Post-Migration Report to verify that the virtual machines were migrated. Then click **Finish**.
12. In Failover Cluster Manager, verify the status of the migrated virtual machines and the related resources:

a. In the console tree, click the name of the new failover cluster, and then click **Roles**. You should see the migrated virtual machine (roles) in the **Roles** pane. The virtual machines will be turned off.

b. Click a virtual machine to display the associated resources at the bottom of the window. For a newly migrated virtual machine, the resources have been registered but are not online.

Before you can start the virtual machines, you must remap the storage to the new cluster and then bring the storage online.

To complete the migration

1. Prepare for clients to experience downtime, probably briefly.
2. Shut down the old cluster to ensure that no one will attempt to start the virtual machine during migration and no connections will be made from storage.

   Caution

   At no time should a virtual machine be running on both the old cluster and the new cluster. A virtual machine that runs on both the old cluster and the new cluster at the same time might become corrupted. You can run a virtual machine on the old cluster while you migrate it to a new cluster with no problems; the virtual machine on the new cluster is created in a Stopped state. However, to avoid corruption, it is important that you do not turn on the virtual machine on the new cluster until after you stop the virtual machine on the old cluster.
3. To complete the transition for the storage:
   a. Make the CSV volume that stores the virtual machines inaccessible to the old cluster, and then make them accessible to the new cluster.
   b. After you move the storage to the new cluster, in Disk Management, bring the CSV volume and Virtual Machine Configuration resource for each virtual machine online.
4. At this point, you should be able to start the virtual machines. To start the virtual machines in Failover Cluster Manager, display and select the virtual machine role, and then click **Start Role**.
5. Install the latest integration services on the new virtual machines. You might need to restart the virtual machine to complete the integration services update.

Note

The Migrate a Cluster Wizard does not migrate Volume Shadow Copy Service (VSS) tasks, Hyper-V Replica Broker settings, Task Scheduler settings, and Cluster Aware Updating (CAU) settings. If you were using any of these features on the old cluster, you will need to configure them on the new cluster.

Verify a successful migration

After you complete the migration, you should bring the virtual machine online, make sure the services that the virtual machine was providing on the old cluster are still available and working

as expected, and test failover for the virtual machine on the new cluster. Verify that you can connect to the virtual machines by using Remote Desktop or Virtual Machine Connection. For detailed steps for verifying a successful role migration and testing failover, see Migration Between Two Multi-Node Clusters.

Related references

How to Move Highly Available (Clustered) VMs to Windows Server 2012 with the Cluster Migration Wizard

Migration Between Two Multi-Node Clusters

In-Place Migration for a Two-Node Cluster

Cluster Migrations Involving New Storage: Mount Points

This topic describes considerations for configuring mount points during a migration to a failover cluster running Windows Server 2012 R2 or Windows Server 2012 when the destination cluster will use new storage after the migration.

Caution
If you want to use new storage, you must copy or move the data or folders (including shared folder settings) during a migration. The wizard for migrating clustered resources does not copy data from one shared storage location to another.

The Migrate a Cluster Wizard does not migrate mount point information (that is, information about hard disk drives that do not use drive letters, but are mounted instead in a folder on another hard disk drive). However, the wizard can migrate Physical Disk Resource settings to and from disks that use mount points. The wizard also does not configure the necessary dependency between the resources for mounted disks and the resource for a host disk (the disk on which the other disks are mounted). You must configure those dependencies after the wizard completes.

When you work with new storage for your cluster migration, you have some flexibility in the order in which you complete the tasks. You must create the mount points, run the Migrate a Cluster Wizard, copy the data to the new storage, and confirm the disk letters and mount points for the new storage. After completing those tasks, configure the disk resource dependencies in Failover Cluster Manager.

A useful way to keep track of disks in the new storage is to give them labels that indicate your intended mount point configuration. For example, in the new storage, when you are mounting a new disk in a folder called **Mount1-1** on another disk, you can also label the mounted disk as **Mount1-1**. (This assumes that the label **Mount1-1** is not already in use in the old storage.) When you run the Migrate a Cluster Wizard, and you need to specify that disk for a particular migrated resource, you can select the disk labeled **Mount1-1** from the list. After the wizard completes, you can return to Failover Cluster Manager to configure the disk resource for **Mount1-1** so that it is

dependent on the appropriate resource - for example, the resource for disk **F**. Similarly, you would configure the disk resources for all other disks mounted on disk F so that they depended on the disk resource for disk F.

After you run the wizard and fully configure the mounted disk, your last task is to configure the disk dependencies in Failover Cluster Manager. For each disk resource for a mounted hard disk drive, open the Properties sheet and, on the **Dependencies** tab, specify a dependency on the disk resource for the host drive (where the mounted drives reside). This ensures that the Cluster service brings the host drive online first, followed by the drives that are dependent on it.

After you configure the dependencies, you can view a dependency report. To view a dependency report, click the service or application in Failover Cluster Manager, and then, under **Actions**, click **Show Dependency Report**. The following illustration shows four mount points that are configured with the correct dependencies on the disk on which they are mounted:

Four mount points with dependencies configured

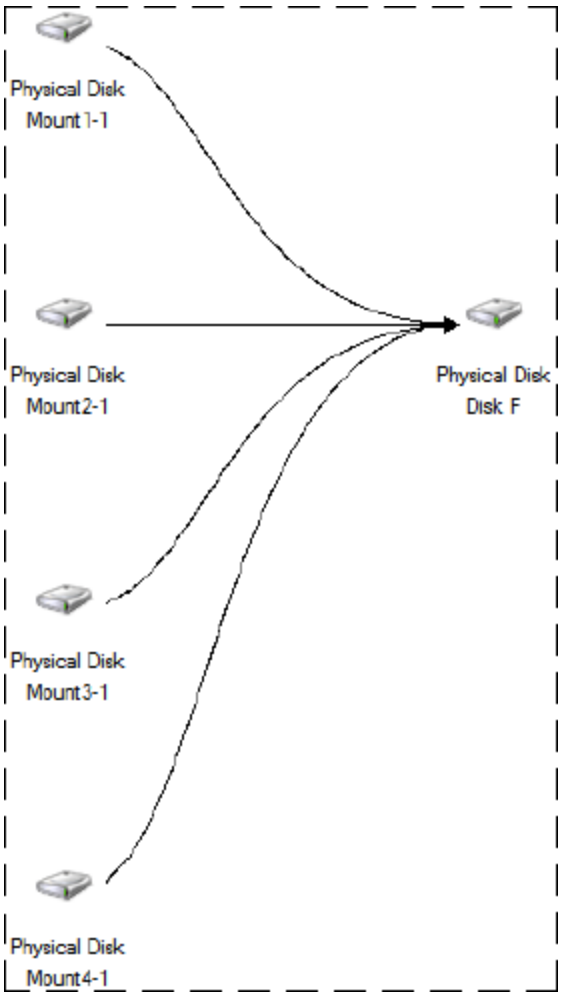

Additional references

Migrate Cluster Roles to Windows Server 2012 R2

Migrating Clustered Services and Applications to Windows Server 2012

Additional References

- Overview of failover clusters:
  - What's New in Failover Clustering in Windows Server 2012
  - Failover Clustering Overview
  - Failover Clustering Hardware Requirements and Storage Options
  - Validate Hardware for a Failover Cluster
- Community resources:
  - Windows Server Migration forum
  - High Availability (Clustering) forum for Windows Server 2012
- Deploying failover clusters:
  - Create a Failover Cluster
  - Deploy a Hyper-V Cluster
- Cluster configuration:
  - Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster
  - Use Cluster Shared Volumes in a Failover Cluster

Secure Windows Server 2012 R2 and Windows Server 2012

To help you locate resources to secure your servers running Windows Server 2012 R2 and Windows Server 2012, the following links have been compiled for essential assessment tools, reference documentation, and product offerings. Get started with the Security and Protection.

Assessment tools

Use these tools to help assess security risks, identify missing security updates, and manage the security and compliance process for the most widely used Microsoft technologies.

- Microsoft Security Assessment Tool 4.0
  The Microsoft Security Assessment Tool is a risk-assessment application designed to provide information and recommendations about best practices for security within an IT infrastructure.

- Microsoft Baseline Security Analyzer 2.2 (for IT pros)
  The Microsoft Baseline Security Analyzer provides a streamlined method to identify missing security updates and common security misconfigurations.
- Microsoft Security Compliance Manager
  This Solution Accelerator provides centralized security baseline management features, a baseline portfolio, customization capabilities, and security baseline export flexibility to accelerate your organization's ability to efficiently manage the security and compliance process for the most widely used Microsoft technologies.

Security technologies

This information is for IT pros to learn about security technologies and features, including important changes, for Windows Server 2012 R2, Windows Server 2012, Windows 8.1 and Windows 8.

- Security and Protection
- **Secure Windows 8 and 8.1**

Administration tools

Use these tools to administer security technologies and address ongoing threats.

- Security Tools to Administer Windows Server 2012
  For the IT pro, this is a list of Microsoft tools that are available for Windows Server 2012 to administer security technologies and address ongoing threats to your computers and network. To help you find the right tool for the job, the security tools are grouped by category and task.

Support and enhance critical security needs

Many technologies, features, and configuration options can be used to enhance the security of computers and networks. Windows Server 2012 R2 and Windows Server 2012 support and enhance critical security needs, including:

- Authentication and identity.
- Authorization and isolation.
- Data protection.
- Secure networking.

The following resources highlight key technologies and features that can be used as building blocks to address these needs (some may be standalone solutions, and others are version interdependent).

Technology or feature	Authentication and identity	Authorization and isolation	Data protection	Secure networking
Access Control	X	X	X	X

Overview				
Active Directory Certificate Services	X	X	X	X
Active Directory Domain Services	X	X	X	X
Active Directory Rights Management Services	X	X	X	
AppLocker		X		
BitLocker	X	X	X	
Remote Access (DirectAccess, Routing, and VPN)	X	X		X
Dynamic Access Control Overview	X	X	X	
Encrypting File System			X	
File and Storage Services	X	X	X	
Group Policy	X	X	X	X
Network Policy and Access Services		X		X
Security Auditing	X	X	X	X
Security Configuration Wizard		X	X	X
Software Restriction Policies		X		
Smart Cards	X			
Web Server (IIS)	X	X		X

Windows Authentication and Logon	X			X
Windows Azure Backup			X	
Windows Firewall with Advanced Security				X
Windows Server Update Services			X	

Additional resources

- Malware Protection Center
  The Microsoft Malware Protection Center provides world class antimalware research and response capabilities that support the Microsoft range of security products and services.
- Wiki: Troubleshooting Portal
  This page is a community-driven list of troubleshooting portals and articles covered on the TechNet Wiki and the top associated articles for each technology.

Security Tools to Administer Windows Server 2012

This topic for the IT professional lists and describes Microsoft tools that are available for Windows Server 2012 to administer security technologies and address ongoing threats to your computers and network.

To help you find the right tool for the job, the following security tools are grouped by category and task:

Category	Task
Access	Manage access to network resources
Auditing	Manage access to network resources
Certificate Services	Manage a CA and other Active Directory Certificate Services tasks
Computer	Analyze and manage computer processes and

Category	Task
	performance
Credentials	Manage user accounts, groups, and credentials
Cryptography	Manage certificates and encryption
Files	Take ownership or securely delete files
Security policies	Analyze and manage security policies
Security principals	Modify or create new security principals
System security	Diagnose, plan and remediate overall system security

The following list provides links to the security cmdlets included in the Windows PowerShell Core Modules, and links to cmdlets for technologies that are sometimes used to manage security in your enterprise.

- Windows PowerShell Security Cmdlets
- **PowerShell Cmdlets for Active Directory**
- **PowerShell Cmdlets for Active Directory Rights Management Services**
- **PowerShell Cmdlets for Applocker**
- **PowerShell Cmdlets for Group Policy**
- **PowerShell Cmdlets for Server Manager**
- **PowerShell Cmdlets for the Best Practice Analyzer**
- Script Center Gallery

Manage user accounts, groups, and credentials

Managing user identities and processes for logon and authentication involve important yet often repetitive tasks. To obtain information about and manage user accounts, groups, and credentials, use one of the following tools.

Tool	Type	Description
Whoami	Windows command-line tool	Displays user, group, and privileges information for the user who is currently logged on to the local computer. If used without parameters, **whoami** displays the current domain and user name.

Tool	Type	Description
cmdkey	Windows command-line tool	Creates, lists, and deletes stored user names and passwords or credentials.
NET LocalGroup	Windows command-line tool	Adds, displays, or modifies local groups.
NET User	Windows command-line tool	Adds or modifies user accounts, or displays user account information.
Get-Credential	Windows PowerShell cmdlet	Gets a credential object based on a user name and password.
Get-Authenticode Signature	Windows PowerShell cmdlet	Gets information about the Authenticode signature in a file.
LogonSessions	Sysinternals utility	Lists active logon sessions.
PsLoggedOn	Sysinternals utility	Lists users logged on to a computer.

Modify or create new security principals

Adding, deleting, and modifying account and group information is one of the most frequent administrator tasks. To modify or create new security principals, use one of the following tools.

Tool	Type	Description
Ktpass	Windows command-line tool	Configures the server principal name for the host or service in Active Directory Domain Services (AD DS) and generates a .keytab file containing the shared secret key of the service. **Note** The .keytab file is based on the Massachusetts Institute of Technology (MIT)

Tool	Type	Description
		implementation of the Kerberos authentication protocol. The Ktpass command-line tool allows UNIX-based services that support Kerberos authentication to use the interoperability features provided by the Key Distribution Center (KDC) service in Windows Server 2008.
cmdkey	Windows command-line tool	Creates, lists, and deletes stored user names and passwords or credentials.
NET LOCALGROUP	Windows command-line tool	Adds, displays, or modifies local groups.
NET USER	Windows command-line tool	Adds or modifies user accounts, or displays user account information.
Dsadd	Windows command-line tool	Allows you to add specific types of objects to the directory.
Add-Computer	Windows PowerShell cmdlet	Adds computers to a workgroup or domain.
Remove-Computer	Windows PowerShell cmdlet	Removes computers from workgroups or domains.
Reset-ComputerMachinePassword	Windows PowerShell cmdlet	Resets the computer account password.

Manage certificates and encryption

Certificate and encryption can significantly strengthen the security of a network and its resources. To manage certificate requests and encrypted files or directories, use the following tools.

Tool	Type	Description
Certreq	Windows command-line tool	Requests certificates from a certification authority (CA), retrieves a response to a previous request from a CA, creates a new request from an .inf file, accepts and installs a response to a request, constructs a cross-certification or qualified subordination request from an existing CA certificate or request, or signs a cross-certification or qualified subordination request.
Cipher	Windows command-line tool	Displays or alters the encryption of directories and files on NTFS volumes. If used without parameters, **cipher** displays the encryption state of the current directory and any files it contains.
Get-PfxCertificate	Windows PowerShell cmdlet	Gets information about .pfx certificate files on the computer.
Certificate Provider	Windows PowerShell provider	Allows you to navigate the certificate namespace and view the certificate stores and certificates. You can also copy, move, and delete certificates and certificate stores, and open the Certificates snap-in for the Microsoft Management Console (MMC).

Manage a CA and other Active Directory Certificate Services tasks

Active Directory Certificate Services (AD CS) role services allow an organization to issue and manage certificates that enable a variety of network infrastructure requirements. To manage a CA and complete a variety of other AD CS tasks, use the following tool.

Tool	Type	Description
Certutil	Windows command-line tool	Collects and displays certification authority (CA) configuration information, configures AD CS, backs up and restores CA components, and verifies certificates, key pairs, and certification paths.

Manage access to network resources

Files, folders, and shares that are protected by using access control lists (ACLs) can be monitored and managed by using the following tools, cmdlets, and utilities. To obtain information about access permissions on resources, use one of the following tools.

Tool	Type	Description
Icacls	Windows command-line tool	Displays or modifies discretionary access control lists (DACLs) on specified files, and applies stored DACLs to files in specified directories. Icacls.exe replaces the Cacls.exe tool for viewing and editing DACLs.
Dsacls	Windows command-line tool	Displays and changes permissions (access control entries) in the ACL of objects in Active Directory Domain Services (AD DS).
Get-Acl	Windows PowerShell cmdlet	Gets the security descriptor for a resource, such as a file or registry key.

Tool	Type	Description
ShareEnum	Sysinternals utility	Allows you to scan file shares on your network and view their security settings.
AccessChk	Sysinternals utility	Displays access permissions to files, registry keys, or Windows services for a specified user or group.
AccessEnum	Sysinternals utility	Displays access permissions to directories, files, and registry keys for all users and groups on computers in your domain.

Take ownership or securely delete files

Administrators might need to modify the ownership of files or ensure that deleted files cannot be accessed. To take ownership or securely delete files, use one of the following tools.

Tool	Type	Description
Takeown	Windows command-line tool	Enables an administrator to recover access to a file that previously was denied, by making the administrator the owner of the file.
SDelete	Sysinternals utility	Allows you to securely overwrite your sensitive files and remove previously deleted files by using this Department of Defense–compliant secure deletion program.

Manage security auditing and audit logs

Security auditing allows you to monitor and analyze a wide variety of computer and network activities. The following utilities can be used to configure event logging and manage event logs and event log entries.

Tool	Type	Description
Auditpol	Windows command-line tool	Displays information about and performs functions to modify audit policy settings.
Logman	Windows command-line tool	Creates and manages Event Trace Session and Performance logs and supports many functions of Performance Monitor from the command line.
Clear-EventLog	Windows PowerShell cmdlet	Deletes all entries from specified event logs on a local or remote computer.
Get-Event	Windows PowerShell cmdlet	Gets the events in the event queue.
Get-EventLog	Windows PowerShell cmdlet	Gets the events in a specified event log or a list of the event logs on a computer.
New-Event	Windows PowerShell cmdlet	Creates a new event.
New-EventLog	Windows PowerShell cmdlet	Creates a new event log and a new event source on a local or remote computer.
Remove-event	Windows PowerShell cmdlet	Deletes events from the event queue.
Remove-EventLog	Windows PowerShell cmdlet	Deletes an event log or unregisters an event source.
Show-EventLog	Windows PowerShell cmdlet	Displays the event logs of the local or a remote computer in Event Viewer.
Write-EventLog	Windows PowerShell cmdlet	Writes an event to an event log.
Limit-EventLog	Windows PowerShell cmdlet	Sets the event log properties that limit the size of the event log and the age of its entries.
PsLogList	Sysinternals utility	Allows you to collect event log records.

Tool	Type	Description
WEvtUtil	Windows command-line tool	Enables you to retrieve information about event logs and publishers. You can also use this command to install and uninstall event manifests, to run queries, and to export, archive, and clear logs.

Analyze and manage security policies

Security policy is the configurable set of rules that the operating system follows when determining the permissions to grant in response to a request for access to resources. You can use the following tools to analyze and manage security policy settings for a single computer or a domain.

Tool	Type	Description
Security Configuration Wizard (SCW)	Windows administrative tool	Determines the minimum functionality required for a server's role or roles and disables functionality that is not required.
Secedit	Windows command-line tool	Configures and analyzes system security by comparing an existing configuration to at least one template.
GPUpdate	Windows command-line tool	Refreshes local and domain Group Policy settings, including security settings. **Note** This command-line tool supersedes the **/refreshpolicy** option for the **secedit** command.
GPResult	Windows command-line tool	Displays Resultant Set of Policy (RSoP) information for a local or domain user and computer.
Local Security Policy	Microsoft Management Console (MMC) snap-in	The Security Policy snap-in (secpol.msc) allows you to adjust settings for Account Policies,

Tool	Type	Description
		Local Policies, Windows Firewall with Advanced Security, Network List Manager Policies, Public Key Policies, Software Restriction Policies, Application Control Policies, IP Security Policies on Local Computer, and Advanced Audit Policy Configuration.
Security templates	Microsoft Management Console (MMC) snap-in	Security templates provide standard security settings to use as a model for your security policies. They help you troubleshoot problems with computers whose security settings are not in compliance with policy or are unknown. Security templates are inactive until imported into a Group Policy object or the Security Configuration and Analysis snap-in to MMC.
AppLocker Overview	Microsoft Management Console (MMC) snap-in	AppLocker helps you control which applications and files users can run. These include executable files, scripts, Windows® Installer files, DLLs, Packaged apps and Packaged app installers. You can also use AppLocker to inventory applications running on your computers.
Software Restriction Policies	Microsoft Management Console (MMC) snap-in	You can use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. Software restriction policies are integrated with Microsoft Active Directory and Group Policy. You can also

Tool	Type	Description
		create software restriction policies on stand-alone computers. Software restriction policies are trust policies, which are regulations set by an administrator to restrict scripts and other code that is not fully trusted from running.

Analyze and manage computer processes and performance

Understanding the configuration and behavior of a computer and the applications and processes running on that computer are important to diagnosing performance issues and system failures but can require detailed investigation. The following tools can assist with many of these tasks.

Tool	Type	Description
Runas	Windows command-line tool	Allows a user to run specific tools and programs with different permissions than the user's current logon provides.
SC	Windows command-line tool	Communicates with the Service Controller and installed services.
Shutdown	Windows command-line tool	Enables you to shut down or restart local or remote computers one at a time.
Tasklist	Windows command-line tool	Displays a list of currently running processes on the local computer or on a remote computer.
Taskkill	Windows command-line tool	Ends one or more tasks or processes. Processes can be ended by process ID or image name.
Bootcfg	Windows command-line tool	Configures, queries, or changes Boot.ini file settings.

Tool	Type	Description
Get-ExecutionPolicy	Windows PowerShell cmdlet	Gets the execution policies in the current session.
Set-ExecutionPolicy	Windows PowerShell cmdlet	Changes the user preference for the execution policy of the shell.
ShellRunAs	Sysinternals utility	Allows you to start programs as a different user via a shell context-menu entry.
PsTools	Sysinternals utility	Includes command-line tools for listing the processes running on local or remote computers, running processes remotely, restarting computers, and obtaining copies of event logs.
Autologon	Sysinternals utility	Allows you to bypass the password screen during logon.
Autoruns	Sysinternals utility	Shows what programs are configured to start automatically when a computer starts and the user logs on. Autoruns also shows the registry and file locations where applications can configure auto-start settings.
Process Explorer	Sysinternals utility	Allows you to find out what files, registry keys, and other objects processes are open, which dynamic link libraries (DLLs) they have loaded, and who owns each process.
PsExec	Sysinternals utility	Allows you to run processes with limited-user rights.

Diagnose, plan and remediate overall system security

Microsoft provides a number of free tools that can be used to diagnose overall system health, plan for improvements and migrations, and security and protect against the risk of infection from malware. The following tools can be used to accomplish these tasks.

Tool	Type	Description
The Security Development Lifecycle Developer Starter Kit	Download	The SDL Developer Starter Kit offers 14 content modules (with speaker notes, presenter guides, and sample comprehension questions) plus eight MSDN virtual labs with lab manuals—all created to help you build a customized SDL training program for your development teams.
Malicious Software Removal Tool	Download	Checks computers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003Windows 8, Windows 7, Windows Vista, and Windows XP for infections by specific, prevalent malicious software and helps remove any infection found.
Microsoft Security Assessment Tool	Download	Provides information and recommendations about best practices to help enhance security within your IT infrastructure.
Enhanced Mitigation Experience Toolkit v4.0	Download	Allows you to design mitigation methods to help prevent malicious users from gaining access to your system.
Microsoft Threat Analysis & Modeling Tool	Download	Allows you to enter information including business requirements and application architecture, which is then used to produce a threat model.
RootkitRevealer	Sysinternals utility	Allows you to scan your computer

Tool	Type	Description
		for rootkit-based malware.
Sigcheck	Sysinternals utility	Allows you to collect file version information and verify that images on your computer are digitally signed.
Attack Surface Analyzer	Download	Allows you to catalogue changes made to the operating system attack surface by the installation of new software.
Microsoft Assessment and Planning Toolkit	Download	The MAP Toolkit is a powerful inventory, assessment and reporting tool that can securely assess IT environments for various platform migrations. Having an inventory of what platforms exist in your environment can enable you to more quickly deploy security updates, react to security incidents, contain any issues that may arise, and recover more quickly from those issues.

See also

The following table provides additional resources for security tools in related technologies.

Group Policy	Group Policy Overview
Active Directory Domain Services	Active Directory Domain Services Overview
Active Directory Certificate Services	Active Directory Certificate Services Overview
Security Troubleshooting	Wiki: Troubleshooting Portal
Windows Server Update Services	Windows Server Update Services Overview
Microsoft System Center	Microsoft System Center 2012

Manage Privacy

There are a variety of technologies that communicate with the Internet to provide increased ease-of-use and functionality. Browser and email technologies are examples, but there are also technologies such as automatic updating that help you obtain the latest software and product information, including bug fixes and software updates. These technologies provide many benefits for users, but they also involve communication with websites, which administrators might want to control.

You can control this communication through a variety of options that are built into individual features, the operating system, and features that are designed for managing configurations across your organization. For example, as an administrator, you can use Group Policy settings to control the way some features communicate. For other features, you can create an environment in which all communication is directed to the organization's internal website instead of to an external website.

This section offers guidance about managing privacy-related settings in Windows Server 2012 and additional links that can be useful to administrators and others concerned about privacy.

- Managing Internet Communication and Privacy
- Manage Privacy: Activation and Resulting Internet Communication
- Managing Privacy: Dynamic Update and Resulting Internet Communication
- Manage Privacy: Internet Explorer 10 and Resulting Internet Communication
- Manage Privacy: SmartScreen Filter and Resulting Internet Communication
- Managing Privacy: User Access Logging and Resulting Internet Communication
- Manage Privacy: Windows Customer Experience Improvement Program and Resulting Internet Communication
- Manage Privacy: Windows Defender and Resulting Internet Communication
- Manage Privacy: Windows Error Reporting and Resulting Internet Communication
- Managing Privacy: Windows Store and Resulting Internet Communication
- Managing Privacy: Using a Microsoft Account to Logon and Resulting Internet Communication
- Appendix A: Resources for Learning About Automated Installation and Deployment
- Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category

Related Resources

- Windows 8 and Windows Server 2012 Privacy Statement
- Windows 8.1 and Windows Server 2012 R2 privacy statement
- Microsoft Safety & Security Central
- Infrastructure Planning and Design
- Microsoft Trustworthy Computing: Privacy

Managing Internet Communication and Privacy

This document provides information about the communication that flows between features in Windows and Internet sites, and it describes steps to take to limit, control, or prevent that communication in an organization with many users. This document is designed to assist you, the administrator, in planning strategies for deploying and maintaining Windows in a way that helps provide an appropriate level of security and privacy for your organization's networked assets.

What this document includes

This document is organized around individual features so that you can find detailed information for any feature you are interested in managing. This information extends the information that is available in the Windows 8 and Windows Server 2012 Privacy Statement.

In this section

Standard computer information sent by Internet-enabled features

Types of features covered in this document

Types of features not covered in this document

Security and privacy basics that are beyond the scope of this document

Standard computer information sent by Internet-enabled features

Software with Internet-enabled features sends information about users' computers ("standard computer information") to the websites that they visit and the online services they use. Microsoft uses standard computer information to provide Internet-enabled services, to help improve our products and services, and for statistical analysis. Standard computer information typically includes information such as the IP address of the computer, operating system version, browser version, and regional and language settings. In some cases, standard computer information may also include a hardware ID, which indicates the device manufacturer, device name, and version. The purpose of this document is not to describe standard computer information that is sent by Internet-enabled features. Instead this document describes the additional information that can be sent or received by these features and how to manage this information.

Types of features covered in this document

This document provides the following:

- Information about features that in the normal course of operation send information to or receive information from Internet sites. An example of this type of feature is Windows Error Reporting. If you choose to use this feature, it sends information to a site on the Internet.

For more information, see Manage Privacy: Windows Error Reporting and Resulting Internet Communication.

- Information about features that routinely display buttons or links that make it easy to initiate communication with Internet sites.
- Brief descriptions of features designed to communicate with the Internet.

It is beyond the scope of this document to describe all aspects of maintaining appropriate levels of security and privacy in an organization running servers that communicate across the Internet. This document does, however, provide basic information about how components such as Internet Information Services work. It provides sources of information about balancing your organization's requirements for Internet communication with requirements for protecting networked assets.

Types of features not covered in this document

This document does not provide the following:

- Information about managing or working with applications, scripts, utilities, web interfaces, Microsoft ActiveX controls, extensible user interfaces, Microsoft .NET Framework, or application programming interfaces (APIs). These are applications or layers that support applications, and they provide extensions that go beyond the operating system.
- Information about Windows Installer—although Windows Installer includes some technology that you can choose to use for installing drivers or other software from the Internet. Windows Installer packages are not described here because they involve scripts or utilities that are created specifically for communicating across the Internet.

  Note
  Web-based and server-based applications such as databases, email, and instant messaging. You must work with your software provider to learn how to mitigate risks that are related to using particular applications (including web-based or server-based applications), scripts, utilities, and other software.
- Information about features that store local logs that could potentially be made available to support personnel or other users. You may want to treat this information like other sensitive information by providing internal guidelines for your support staff about handling logs and other information that you want to protect.

Security and privacy basics that are beyond the scope of this document

This document is designed to assist you, the administrator, in planning strategies for deploying and maintaining Windows in a way that provides an appropriate level of security and privacy for your organization's networked assets. This document does not describe security and privacy basics—that is, strategies and risk-management methods that provide a foundation for security and privacy across your organization. It is assumed that you are actively evaluating and studying these security and privacy basics as a standard part of network administration.

Some security basics that are a standard part of network administration include:

- Monitoring, which includes using a variety of software tools, including tools to assess which ports are open on servers and clients.
- Virus-protection software.
- The principle of least privilege (for example, not signing in as an administrator if signing in as a user is just as effective).
- The principle of running only the services and software that are necessary—that is, stopping unnecessary services and keeping computers (especially servers) free of unnecessary software.
- Strong passwords—that is, requiring all users and administrators to choose passwords that are not easily broken.
- Risk assessment as a basic element for creating and implementing security plans.
- Software deployment and maintenance routines to help ensure that your organization's software is running with the latest security updates and patches.
- Defense-in-depth (also referred to as in-depth defense), which means creating redundancy in security systems. An example is using firewall settings together with Group Policy to control a particular type of communication with the Internet.

Resources about security basics

The following websites are a few of the many sources of information about the security basics described previously:

- Microsoft Safety & Security Central
- Infrastructure Planning and Design
- Secure Windows Server 2012 R2 and Windows Server 2012
- Featured Security Content on the Security Developer Center
- Security TechCenter
- Security Guidance
- Microsoft Trustworthy Computing: Privacy

Manage Privacy: Activation and Resulting Internet Communication

In this section

Purposes of activation

Overview: Activation in a managed environment

How a computer communicates with sites on the Internet during activation

This section discusses the purposes of product activation and how activation-related features communicate over the Internet. It explains steps to take to limit, control, or prevent that communication in an organization with many users.

Purposes of activation

Product activation reduces software piracy, helps ensure that Microsoft customers are receiving genuine Microsoft software, and helps to avoid the risks that are associated with the use of unlicensed software. Genuine Windows provides assurance that the software is reliable, and it helps protect against the security threats and increased cost-of-ownership that can be introduced by counterfeit software. Using genuine Windows products helps ensure that software is reliable as follows:

- Ensures that the software is supported by Microsoft and its partners.
- Assists with license compliance.
- Enhances protection from the risks associated with counterfeit software, such as spyware, malware, and viruses.
- Protects against the potential financial penalties and risks to an organization's reputation due to using non-licensed software.

For Windows client, activation by phone or online is required. OEM installed Windows Server systems are pre-activated by the OEM.

If you acquire licenses through a volume license program, you can perform Windows volume activation and verify that the software is genuine by using the following features:

- Active Directory-Based Activation
- Key Management Service (KMS)
- Multiple Activation Key (MAK)

For more information about volume activation, see Activation options with volume licensing later in this section.

Note
Product activation means that a specific product key becomes associated with the computer hardware that it is installed on. Making significant changes to computer hardware or other significant configuration changes may require that the activation process be completed again.

For more information about product activation, see Microsoft Product Activation.

Overview: Activation in a managed environment

In an environment with many computers you probably want to use an activation option that is designed for use with volume licensing. The following subsection describes these options.

Activation options with volume licensing

Organizations that have a volume license agreement have multiple options for activation:

- **Active Directory-based Activation** Active Directory-based Activation enables you to use Active Directory Domain Services (AD DS) to store activation objects, which can further

simplify the task of maintaining volume activation services for a network. With Active Directory-based Activation, IT pros can complete activations on their local network, which eliminates the need for individual computers to connect to Microsoft for product activation. With Active Directory-based Activation, no additional host server is needed, and activation requests are processed transparently with no user interaction or messages during computer startup.

Any computers with a Generic Volume License Key (GVLK) that are connected to an activated domain activate automatically and transparently. They stay activated as long as they remain members of the domain and maintain periodic contact with a domain controller. Activation takes place after the licensing service starts and renews every seven days. When this service starts, the computer contacts AD DS automatically, receives the activation object, and activates without user intervention.

Note

The AD DS schema must be at the server or client functional level for activation objects to be stored in AD DS.

- **Key Management Service (KMS)** KMS is a role of the Software Licensing Service that allows organizations to activate systems within their network from a server where a KMS host key has been installed. With KMS, IT pros can complete activations on their local network, which eliminates the need for individual computers to connect to Microsoft for product activation. By default, KMS clients (GVLK on Client or Server) contact the KMS host and if that connection doesn't happen within 180 days the system falls into Notifications mode. KMS does not require a dedicated system, and it can be cohosted on a system that provides other services. By default, volume editions connect to a configured KMS host to request activation. No action is required from the user.
- **Multiple Activation Key (MAK)** A MAK is a volume license key that is used for one-time activation with activation services that are hosted by Microsoft. There are two ways to use MAK to activate computers:
  - **MAK independent activation** Each computer must independently connect and be activated by Microsoft over the Internet or by telephone.
  - **MAK proxy activation** A computer that is acting as a MAK proxy gathers activation information from multiple computers on the network, and then sends a centralized activation request to Microsoft on their behalf. MAK proxy activation is configured by using the Volume Activation Management Tool (VAMT).

For more information about the Volume Activation Management Tool, see Volume Activation Management Tool (VAMT) Overview

For information about Automatic Virtual Machine Activation (AVMA) see Automatic Virtual Machine Activation.

How a computer communicates with sites on the Internet during activation

If you are using MAK, OEM or Retail activation, you can activate over the Internet or by phone. The following list describes what is communicated when activation is done directly over the Internet:

- **Specific information sent or received**: During the online activation process, the following information is sent to an activation server that is maintained by Microsoft:
  - Computer make and model
  - Version information for the operating system and software that is using Genuine Advantage
  - Region and language settings
  - A unique number that is assigned to your computer (a globally unique identifier or GUID)
  - Product key (hashed) and product ID
  - BIOS name, revision number, and revision date
  - Hardware ID – non-reversible hash of hardware component IDs.

  Important
  The tools do not collect a user's name, address, email address, or any other information that Microsoft can use to identify or contact a person.

  In addition to the configuration information above, the following status information is also transferred:

  - Whether the installation was successful, if one was performed
  - The result of the validation check, including information about any activation exploits and any related malicious or unauthorized software that is found, disabled, or removed
  - The name and a hash of the contents of the computer's start-up instructions file (commonly called the boot file) to help Microsoft discover activation exploits that modified this file

  Note
  If your system is identified as non-genuine, additional information may be sent to Microsoft to better understand why your system failed validation. This information can include error codes and the names and paths of files that compromise the integrity of your system.

  For activation of an individual computer (where volume licensing is not being used), owners can allow the preceding information to be sent over the Internet to the activation system at Microsoft, or they can present the product key information and hardware hash (combined into one number) by phone.
- **Disabling Activation**: Product activation cannot be disabled, but if you acquire licenses through a volume license program, you can perform Volume Activation through Active Directory-based Activation, the Key Management Service (KMS), or a Multiple Activation Key

(MAK). For more information, see Activation options with volume licensing earlier in this section.

The system must be activated immediately upon installation. Failure to activate the Windows operating systems prevents you from being able to complete system customization.

- **Logging**: Entries that track the progress of activation (for example, return codes and error codes) are logged in Event Viewer. If activation fails, you can use these events to troubleshoot the issue. To locate the events, click **Windows Logs**, click **Application**, click **Source**, and then click **Security-SPP**.
- **Encryption and storage**: Data that is transmitted is encrypted during transmission by using HTTPS (that is, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) with HTTP), and it is stored in Microsoft-controlled facilities. The data is accessible to a restricted number of support personnel who oversee and maintain the activation servers and the product activation program.
- **Privacy**: Customer privacy was a paramount design goal in building the product activation technology. Microsoft uses the information that is sent to confirm that you have a licensed copy of the software. The information is aggregated for statistical analysis. Microsoft does not use the information to identify or contact a person.
- **Transmission protocol and port**: When activating over the Internet, the first transmission uses HTTP through port 80. It communicates with go.microsoft.com to check the HTTP response code. A response code of less than 500 indicates that a product activation server is available. If the product activation server can be reached, any activation data that is sent by Windows Product Activation uses HTTPS through port 443 to sls.microsoft.com. For a complete list of all URLs and ports required to complete activation, see Using MAK Activation.

Additional references

For more information about volume licensing, activation, and Genuine Advantage, see the following pages on the Microsoft website:

- Microsoft Volume Licensing
- Plan for Volume Activation
- Automatic Virtual Machine Activation
- Genuine Windows

Managing Privacy: Dynamic Update and Resulting Internet Communication

In this section

Benefits and purposes of Dynamic Update

Overview: Using Dynamic Update in a managed environment

How Dynamic Update communicates with sites on the Internet

Controlling Dynamic Update to limit the flow of information to and from the Internet

This section explains how Dynamic Update communicates across the Internet, and it explains steps to limit, control, or prevent that communication in an organization with many users.

Benefits and purposes of Dynamic Update

With Dynamic Update, if you start a computer from an existing operating system (for example, Windows 8), and then run Setup from that operating system, Setup can check for new Setup files, including drivers and other files.

Note
If you perform a network boot, for example, from a Pre-Boot Execution Environment (PXE)-enabled computer, and then run Setup, Dynamic Update does not occur. Similarly, if you start a computer with the Windows Preinstallation Environment (Windows PE), even if media is used, Dynamic Update does not occur.

In an interactive installation, the person installing is prompted to choose whether to allow Dynamic Update to occur. In an unattended installation that uses an answer file, an entry in the answer file can control whether Dynamic Update occurs.

Using Dynamic Update reduces the need to apply patches to recently installed systems, and it makes it easier to run Setup with hardware that would otherwise prevent Setup from being completed successfully.

Note
Additional drivers that were recently added or updated that would not prevent Setup from completing successfully are downloaded to the system the first time the user runs Windows Update.

Dynamic Update performs the same type of check for software updates as can be performed through the existing, earlier operating system, However, Dynamic Update runs during Setup, and a limited set of software updates can be downloaded through Dynamic Update. All files that are made available through Dynamic Update are very carefully tested and fall into three categories:

- **Setup software updates**: These updates help Setup run correctly. Dynamic Update handles only limited, important Setup updates.
- **New or changed drivers**: These are drivers that are known to be necessary for success with Setup. They include only network, video, audio, and mass storage drivers. Dynamic Update downloads only the files that are required for a particular computer, which means that the Dynamic Update software briefly examines the computer hardware. The information that is collected is not saved. The only purpose for examining the hardware is to select appropriate drivers for it. This keeps the download time as short as possible, and it ensures that only necessary drivers are downloaded to the hard disk drive.

  Note
  Another alternative for installing drivers during Setup is to use interactive Setup, and press **F6** when prompted. Or you can make use of a deployment technology (such as

unattended setup), which enables you to create operating system images and control the drivers that are included in a specific image.

- **Updates to operating system features**: These are high-priority updates that can help make operating system features more resistant to attack immediately after installation and any blocking issues that prevent Setup from completing. These updates help increase the security of a newly installed operating system when it first connects to a network—before you begin your standard software update process (whether you use the Windows Update Web servers, Windows Server Update Services, or a system management solution).

Dynamic Update checks for the new files in the same location that the existing operating system used for software updates. (This is the same location from which Setup was run.) This location could be any of the following:

- **The Windows Update Web servers**: On a computer that has been receiving software updates from the Internet, Dynamic Update continues to go to the Internet; that is, Windows Update Web servers.
- **A Windows Server Update Services server**: On a computer that previously used Windows Server Update Services (WSUS), Dynamic Update continues to go to a WSUS server.

  For information about WSUS, see Windows Server Update Services:
- **A system management server**: On a computer that previously used system management servers (for example, servers running Systems Center Configuration Manager), Dynamic Update continues to use a management server.

  For more information, see Microsoft Systems Center Configuration Manager

Overview: Using Dynamic Update in a managed environment

In a managed environment where you are Windows on many computers, you might want to prevent Dynamic Update from connecting to the Windows Update Web servers. To do this, you can use Windows Server Update Services or a system management solution, or you can perform an unattended installation with an answer file entry that prevents Dynamic Update. For more information, see Controlling Dynamic Update to limit the flow of information to and from the Internet later in this section.

How Dynamic Update communicates with sites on the Internet

This subsection focuses on the communication that occurs between Dynamic Update and the Windows Update Web servers during an interactive installation or a preinstallation compatibility check when the computer has access to the Internet. This subsection also provides a description of the default behavior of Dynamic Update with an unattended setup.

Note

This subsection describes how Dynamic Update works if a computer runs an earlier operating system, the computer is currently configured to go to the Windows Update Web servers for software updates, and you run Setup from the operating system already running on the computer. You can adjust the description to fit other scenarios, for example, when you are upgrading from Windows 8, or where WSUS is being used.

For a description of how you can control the behavior of Dynamic Update during unattended installations, see Controlling Dynamic Update to limit the flow of information to and from the Internet later in this section.

Specific information sent or received: When Dynamic Update contacts the Windows Update Web servers, it sends only the exact operating system version and the information that is necessary to select appropriate drivers (for example, network, video, audio, or mass storage drivers).

The files that Dynamic Update downloads are only those that are important to:

- Ensure that Setup runs successfully.
- Help protect operating system features immediately after installation (until the normal software-update process can begin).

Files with minor updates that have little impact on the preceding items are not made available through Dynamic Update. Some of the updated files will be replacements (for example, an updated Setup file) and some will be additions (for example, a driver that was not available at the time the Setup CD was created).

Default behavior and triggers: During a conventional interactive installation, Dynamic Update occurs automatically.

During an unattended installation with an answer file, if the answer file does not contain any entries that are related to Dynamic Update, Dynamic Update will occur.

Note

If the computer is not connected to the Internet during installation, Dynamic Update cannot occur during a conventional interactive setup or during an unattended installation with an answer file.

User notification: During an interactive installation, a progress indicator appears that enables the person to track the status of the update process. During an unattended installation, there is no notification. (By definition, an unattended installation means that no user interaction is required.)

Logging: By default, the progress of Setup is logged in *systemroot*\Sources\Panther\setupact.log in the installation folders for the operating system that is being upgraded. After the upgrade is complete, the information about the new installation is stored in *systemroot*\Panther\setupact.log. You can view this log if you have questions about Dynamic Update, for example, if you want to know whether Dynamic Update occurred or which files were successfully downloaded during Dynamic Update.

Encryption: Dynamic Update uses the same encryption methods as Windows Update. This means that the initial data is transferred by using HTTPS (that is, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) with HTTP) and updates are transferred by using HTTP.

Access and privacy: No information about the hardware devices on a particular computer is saved or stored by Dynamic Update, so no one can access this information. The information is used only to select appropriate drivers.

Transmission protocol and port: Dynamic Update uses the same transmission protocols and ports as Windows Update: HTTP with port 80 and HTTPS with port 443.

Ability to disable: During interactive Setup, Dynamic Update cannot be disabled. During an unattended installation with an answer file, Dynamic Update is disabled if the answer file includes the following lines:

```
<DynamicUpdate>
     <Enable>false</Enable>
</DynamicUpdate>
```

Controlling Dynamic Update to limit the flow of information to and from the Internet

If you do not want Dynamic Update to connect to the Windows Update Web servers during the installation, you have several options:

- **Use Windows Server Update Services or a system management solution**: You can use Windows Server Update Services or a system management solution to cause Dynamic Update to use a server that you configure instead of the Windows Update Web servers.

  For more information, see the following websites:

  Windows Server Update Services

  Microsoft Systems Center 2012 Configuration Manager
- **Avoid Dynamic Update**: You can avoid using Dynamic Update, which means that Setup will use only the files and drivers that are provided on the installation media. The method to avoid using Dynamic Update depends on how you are performing the installation. Options include:
  - **Interactive installation**: During an interactive installation, when prompted, you can choose to not use Dynamic Update. As an alternative, you can ensure that the computer does not have Internet access.
  - **Unattended setup**: During an unattended installation with an answer file, Dynamic Update does not occur if the answer file includes the following lines:

    ```
    <DynamicUpdate>
         <Enable>false</Enable>
    </DynamicUpdate>
    ```

    For more information, see the Windows Deployment with the Windows ADK

For additional information about performing automated installations, see Appendix A: Resources for Learning About Automated Installation and Deployment.

Manage Privacy: Internet Explorer 10 and Resulting Internet Communication

In this section

This section provides overview information about using Internet Explorer 10 and suggestions for sources of information about how to balance user requirements for Internet access with your organization's requirements to protect networked assets. It includes the following subsections:

Benefits and purposes Explains the benefits of Internet Explorer 10.

Enhanced Security Configuration Describes Internet Explorer Enhanced Security Configuration.

Security-related features Provides examples of the security-related features that are offered in Internet Explorer 10, including SmartScreen Filter.

Resources for learning about security in Internet Explorer 10 Lists resources for learning about topics that are related to security in Internet Explorer 10. This includes resources to help you learn about:

- Security and privacy settings
- Mitigating the risks inherent in web-based applications and scripts
- Methods for controlling the configuration of Internet Explorer 10 in your organization by using Group Policy settings, the Internet Explorer Administration Kit (IEAK), or both

Procedures for controlling Internet Explorer Details procedures to perform specific actions related to Internet Explorer 10. These actions include:

- Choosing a web browser during unattended installation or by using the Default Programs interface.
- Turning Internet Explorer Enhanced Security Configuration off or on.
- Setting the security level to High for specific websites.

The following information is not included in this document:

- This section of this document describes Internet Explorer 10, but it does not describe related features such as Content Advisor or the wizard for making a connection to the Internet.
- It does not describe error reporting for Internet Explorer. For information about this feature, see **Windows Error Reporting and the Problem Reports and Solutions Feature in Windows 8 and Windows Server 2012**.
- It is beyond the scope of this document to describe all the aspects of maintaining appropriate levels of security in an organization where users perform such actions as connecting to websites, running software from the Internet, or downloading content from the Internet.

For more information about Internet Explorer 10, see the following resources:

- Help for Internet Explorer. (Open Internet Explorer, and press F1.)
- Internet Explorer 10 home page on Microsoft TechNet

- Internet Explorer 10 Privacy Statement

Benefits and purposes

Internet Explorer 10 is designed to make it easy to browse and interact with sites on an intranet or on the Internet. It differs from many of the other features that are described in this document because its main function is to communicate with sites on the Internet or an intranet (which contrasts with features that communicate with the Internet in the process of supporting another activity).

Internet Explorer 10 is designed to be highly configurable, and it has security and privacy settings to help protect your organization's networked assets while providing access to useful information and tools. Internet Explorer Enhanced Security Configuration, which is enabled by default when you install Windows Server 2012, helps make your server more secure by limiting exposure to malicious websites.

Note
Using Internet Explorer 10 allows enterprises to continue using existing line-of-business applications. It also provides a new browsing experience for a corporate workforce that is using Windows touch devices. In addition, there are more than 1,500 Group Policy settings that IT professionals can use to provide management and configuration support in Internet Explorer 10.

Enhanced Security Configuration

Internet Explorer Enhanced Security Configuration is turned on by default when you install Windows Server 2012. This configuration assigns specific levels of security settings to four zones that are defined in Internet Explorer 10: the Internet zone, the Local intranet zone, the Trusted sites zone, and the Restricted sites zone. For example, it assigns High security settings to the Internet zone and the Restricted sites zone.

The configuration also contains a variety of other settings. These include specific settings such as whether the Temporary Internet Files folder is emptied when the browser is closed, and settings that determine which zone standard websites are added to (for example, the Windows Update website is added to the Trusted sites zone).

For more information about Internet Explorer Enhanced Security Configuration, on a server that is running Windows Server 2012, open Internet Explorer, and then click one of the following links:

- If Internet Explorer Enhanced Security Configuration is turned on, click **Effects of Internet Explorer Enhanced Security Configuration**.
- If Internet Explorer Enhanced Security Configuration is turned off, click **Internet Explorer Enhanced Security Configuration**.

Security-related features

Security-related features in Internet Explorer 10 include:

- **SmartScreen Filter** Blocks the download of malicious software and providing enhanced antimalware support. Administrators can use Group Policy to configure the behavior of the SmartScreen Filter, for example, to prevent users from overriding the warning shown when a reported unsafe site or download is detected. The SmartScreen Filter is described in Manage Privacy: SmartScreen Filter and Resulting Internet Communication later in this document.
- **ActiveX Filtering** Provides control over how web pages run on your computers. With ActiveX Filtering, you can turn off ActiveX controls for all websites, and then turn them on selectively. Although ActiveX controls can enable useful web experiences for videos and more, some organizations may want to limit how they run for security and performance.
- **Delete Browsing History** Enables users and organizations to delete browsing history for all websites. Administrators can configure the **Delete Browsing History** options through Group Policy or the Internet Explorer Administration Kit. Administrators can also configure which sites are automatically included in the **Favorites** list. This enables them to create policies that help ensure security by aggressively clearing Internet files, without affecting day-to-day interactions with users' preferred and favorite websites. The **Delete Browser History on Exit** check box (on the **General** tab of the **Internet Options** dialog box) allows users and administrators to automatically delete the browsing history on exit.
- **InPrivate Browsing** Deletes the user's browsing history data that is accumulated on the computer when the Internet Explorer browsing windows for that session are closed. A network administrator can use Group Policy to control how InPrivate Browsing is used in their enterprise.
- **Tracking Protection Lists** Help users stay in control of their privacy as they browse the web. Much of the content, images, ads, and analytics that users see are provided by websites outside your organization. Although this content can provide value to your organization, these websites have the ability to potentially track users' behaviors across multiple sites.

  Tracking Protection Lists contain domains that Internet Explorer will block in addition to domains that Internet Explorer will not block. Tracking Protection stays on until you decide to turn it off. To use this functionality, you simply have to add a Tracking Protection List from one of the Tracking Protection List providers.
- **Enhanced Protected Mode** Extends Protected Mode, which was introduced in Internet Explorer 7 for Windows Vista. Protected Mode helped prevent attackers from installing software or modifying system settings by reducing some of the capabilities that are available to Internet Explorer. Enhanced Protected Mode extends this concept by further restricting capabilities for accessing personal information and for accessing information on corporate intranets as follows:
  - **Protects personal information** Restricts Internet Explorer from locations that contain personal information until you grant permissions to it. This helps prevent unauthorized access to personal information.
  - **Protects corporate assets** Restricts access to valuable information on corporate network resources by controlling access through Internet tab processes as follows:
    - Internet tab processes do not have access to a user's domain credentials.
    - Internet tab processes cannot operate as local web servers.
    - Internet tab processes cannot make connections to intranet servers.

Warning

Internet Explorer always runs with Enhanced Protected Mode enabled. Because Internet Explorer offers free browsing, the compatibility impact of this security feature is minimal. However, some add-ons, such as Adobe Flash and certain toolbars are not yet compatible with Enhanced Protected Mode.

To enable Enhanced Protected Mode in the classic user interface, click **Tools**, click **Internet Options**, click the **Advanced** tab, and then click **Enable Enhanced Protection Mode***.

- **Secure Sockets Layer (SSL)** Provides a security report icon to the right of the address bar when you view a page that uses a Hypertext Transfer Protocol Secure (HTTPS) connection. This makes it easier to see whether web transactions are secured by SSL or Transport Layer Security (TLS). Clicking the icon displays a report that describes the certificate that is used to encrypt the connection and the certification authority (CA) that issued the certificate. The security report also provides links to more detailed information.

  Internet Explorer 10 also supports high assurance certificates, which provide further confidence to users that they are communicating with a verified organization. This verification is granted by existing CAs and shows up in the browser as a clear green fill in the address bar.
- **Microsoft ActiveX Opt-In** Enables users to selectively allow or prevent running the ActiveX control. Internet Explorer 10 disables all ActiveX controls that were not used in Internet Explorer 6 and all ActiveX controls that are not flagged for use on the Internet. When users encounter an ActiveX control for the first time, they are prompted to choose if they want to use the control. By default, the ActiveX opt-in does not apply to **Intranet** and **Trusted Site** zones. Controls for those zones, including preapproved controls, run without prompting.

The following list names some of the security-related features in Internet Explorer 10 that are continued from earlier versions of Internet Explorer.

- **Privacy tab** This tab (click **Tools**, and then click **Internet options**) provides flexibility for blocking or allowing cookies, based on the website that the cookie came from or the type of cookie. Types of cookies include first-party cookies, third-party cookies, and cookies that do not have a compact privacy policy. This tab also includes options to control website requests for physical location data, the ability to block pop-ups, and the ability to run toolbars and extensions when InPrivate browsing is enabled.
- **Security settings that define security zones** For each zone, users can control how Internet Explorer 10 handles higher-risk items such as ActiveX controls, downloads, and scripts.
- **Support for content-restricted inline floating frames (IFrames)** This type of support enables developers to implement IFrames in a way that makes it more difficult for malicious authors to start email-based or content-based attacks.
- **A configurable pop-up blocker** This helps you control pop-ups.
- **An improved interface for managing add-ons** Add-ons are programs that extend the capabilities of the browser.

For more information, see the Internet Explorer 10 home page on Microsoft TechNet.

Resources for learning about security in Internet Explorer 10

This subsection lists resources to help you learn about the following topics that are related to security in Internet Explorer 10:

- Learn about security and privacy settings
- Learn about mitigating the risks inherent in web-based applications and scripts
- Learn about Group Policy Objects that control configuration settings
- Learn about the Internet Explorer Administration Kit

In addition, for information about unattended installation, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment later in this document.

Learn about security and privacy settings

Following are sources of detailed information about the security and privacy settings in Internet Explorer 10:

- Internet Explorer 10 TechCenter
- Windows Internet Explorer 10: Overview for IT Pros

In addition, the privacy statement for Internet Explorer 10 includes information about some of the features in Internet Explorer 10: Windows Internet Explorer 10 Privacy Statement.

Learn about mitigating the risks inherent in web-based applications and scripts

In network-based and Internet-based environments, code can take a variety of forms including scripts within documents, scripts within email messages, or applications or other code objects that are running within web pages. This code can move across the Internet, and it is sometimes referred to as "mobile code." Configuration settings provide ways for you to control how Internet Explorer 10 responds when a user tries to run mobile code.

The following examples explain how you can customize the Internet Explorer configuration that is deployed in your organization.

- You can control the code (in ActiveX controls or in scripts, for instance) that users can run. Do this by customizing Authenticode settings. For example, this can prevent users from running any unsigned code, or enable them to only run code that is signed by specific authors. For more information, see Code-Signing Best Practices.
- If you want to permit the use of ActiveX controls, but you do not want users to download code directly from the Internet, you can specify that when Internet Explorer 10 looks for a requested executable, it looks on your internal website instead of the Internet. You can do this by changing a registry key.

  Caution
  Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer. You

can also use the **Last Known Good Configuration** startup option if you encounter issues after manual changes are applied.

The registry key that you will change specifies an Internet search path for Internet-based code, as follows:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Internet Settings\CodeBaseSearchPath

This registry key usually contains the keyword **CodeBase**. When **CodeBase** is present, calls to **CoGetClassObjectFromURL** check the *szCodeURL* location to download components. After **CodeBase**, the **CodeBaseSearchPath** registry key usually lists additional URLs in the Internet search path, with each URL enclosed in angle brackets and separated by a semicolon.

If you remove **CodeBase** from the registry key, and instead specify a site on your intranet, software will check that site, not an Internet site, for downloadable components. The URL that is specified in **CodeBaseSearchPath** will receive an HTTP POST request with data in the following format, and respond with the object to install and load.

```
CLSID={class id}
Version=a,b,c,d
MIMETYPE=mimetype
```

For more information, search for all instances of **CodeBaseSearchPath** in the following MSDN topic: Implementing Internet Component Download.

Learn about Group Policy Objects that control configuration settings

You can control configuration settings for Internet Explorer 10 by using Group Policy Objects (GPOs). Internet Explorer 10 provides nearly 1,500 Group Policy settings that IT pros can use to manage and control the web browser configuration. For more information, see Group Policy Settings in Internet Explorer 10.

You can also control the configuration of Internet Explorer by using the Internet Explorer Administration Kit. For more information, see Learn about the Internet Explorer Administration Kit later in this section.

To learn about specific Group Policy settings that can be applied to computers running Windows 8 and Windows Server 2012, see the following sources of information:

- Group Policy Overview
- Group Policy Settings Reference for Windows Server 2012 and Windows 8

Learn about the Internet Explorer Administration Kit

You can use the Internet Explorer Administration Kit (IEAK) to create a customized Internet Explorer package for use in your organization. You can then deploy your customized package by using standard means such as network shared folders, intranet sites, or through a system

management solution, such as Microsoft System Center Configuration Manager. (You can also control the configuration of Internet Explorer by using Group Policy.)

A few of the features and resources in the IEAK include:

- **Internet Explorer Customization Wizard.** Step-by-step screens guide you through the process of creating customized browser packages that can be installed on client computers.
- **IEAK Help.** The IEAK Help includes many conceptual and procedural topics that you can view by using the **Contents** and **Search** tabs. You can also print topics from IEAK Help.

For more information about the IEAK, see Internet Explorer Administration Kit (IEAK) Information and Downloads.

Procedures for controlling Internet Explorer

This subsection provides procedures to carry out the following tasks:

- Control the browsers that are available
- Turn Internet Explorer Enhanced Security Configuration on or off
- Set the security level to High for specific websites

Procedures for controlling web browsers

Methods for controlling the browsers that are available include:

- Unattended installation by using an answer file
- The Default Programs interface

To specify a browser during unattended installation by using an answer file

1. Use the methods that you prefer for unattended installation or remote installation to create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment later in this document.
2. Confirm that your answer file includes the following lines. If you already have a <ClientApplications> section in your answer file, the "Internet" line (the line containing information about your browser) should be included in the <ClientApplications> section rather than repeating the section.

   <ClientApplications>

   <Internet>*browser_canonical_name***</Internet>**

   </ClientApplications>

   For *browser_canonical_name*, specify the canonical name that is coded into your web browser.

To remove visible entry points to Internet Explorer during unattended installation by using an answer file

1. Use the methods that you prefer for unattended installation or remote installation to

create an answer file. For more information about unattended and remote installation, see Appendix A: Resources for Learning About Automated Installation and Deployment later in this document.

2. Confirm that your answer file includes the following lines. If you already have a <WindowsFeatures> section in your answer file, the "ShowInternetExplorer" line should be included in the <WindowsFeatures> section rather than repeating the section.

   <WindowsFeatures>

   <ShowInternetExplorer>false</ShowInternetExplorer>

   </WindowsFeatures>

Note

This procedure removes visible entry points to Internet Explorer, but it does not prevent Internet Explorer from running.

To specify a browser through the default programs interface

1. In Control Panel, click **Default Programs**, and then click **Set your default programs**.
2. Under **Programs**, click the browser that you want to select as the default.

   Note

   If the web browser that you want to use does not appear by name, contact the vendor of that program for information about how to configure it as the default.
3. To use the selected program as the default for opening all file types and protocols, click **Set this program as default**.

   As an alternative, you can click **Choose defaults for this program**, and then specify which file types and protocols the selected program should open by default.

Procedure to turn Internet Explorer Enhanced Security Configuration on or off

Before you begin this procedure, confirm that no instances of Internet Explorer are running; otherwise, you will have to close and reopen all instances of Internet Explorer after you complete this procedure.

To turn Internet Explorer Enhanced Security Configuration on or off

1. Open Server Manager and click **Configure this local server** to open the **Local Server** configuration page.
2. In the **Properties** area, next to **IE Enhanced Security Configuration**, click **On** to open the **Internet Explorer Enhanced Security Configuration** dialog box.
3. To allow or prevent members of the Local administrators security group to use Internet Explorer in its default client configuration, under **Administrators**, click **On** or **Off**.
4. To allow or prevent members of all other groups to use Internet Explorer in its default client configuration, under **Users**, click **On** or **Off**.

Procedures for setting the security level to High for specific websites

The procedures that follow provide information about how to set the security level for a particular website to **High**, which prevents actions such as running scripts and downloading files from the site.

For information about planning a configuration for your organization to control whether Internet Explorer allows downloads or if it allows plug-ins, ActiveX controls, or scripts to run, see Security-related features and Learn about security and privacy settings earlier in this section.

To configure a specific computer with a security level of High for specific sites

1. On the computer that you want to configure, open Internet Explorer, click **Tools**, click **Internet Options**, and then click the **Security** tab.
2. Click **Restricted sites**, and under **Security level for this zone**, make sure that the slider for the security level is set to **High**.

   Note
   If the Internet Explorer Enhanced Security Configuration is turned on, the slider will be set to **High**, and it cannot be adjusted.

   If the Internet Explorer Enhanced Security Configuration is turned off, the slider can be adjusted, and the security level can be set to a **Custom level**. If it is set to a **Custom level**, click **Default Level**, and then make sure that the slider for the security level is set to **High**.
3. With **Restricted sites** still selected, click **Sites**.
4. In **Add this website to the zone**, type the website address that you want to add to the list of restricted sites. You can use an asterisk as a wildcard character. For example, for websites at Example.Example.com and www.Example.com, you could type:

   http://*.Example.com
5. Click **Add**.

To use IEAK to set the security level to High for specific sites

1. In Internet Explorer Administration Kit, navigate to the **Security and Privacy Settings** page of the customization wizard.
2. In the **Security Zones and Privacy** section, select **Import the current security zones and privacy settings**. Click **Modify Settings**.
3. In the details pane, double-click **Security Zones and Content Ratings**.
4. Under **Security Zones**, click **Import the current security zones and privacy settings**, and then click **Modify Settings**.
5. Select **Restricted sites**.
6. Under **Security level for this zone**, make sure that the slider for the security level is set to **High**.

7. With **Restricted sites** still selected, click **Sites**.
8. In **Add this website to the zone**, type a website address that you want to restrict. You can use an asterisk as a wildcard character. For example, for websites at Example.Example.com and www.Example.com, you could type:

   http://*.Example.com
9. Click **Add**.

Manage Privacy: SmartScreen Filter and Resulting Internet Communication

In this section

Benefits and purposes of SmartScreen Filter

Overview: Using SmartScreen Filter in a managed environment

How SmartScreen Filter communicates with a web service on the InternetHow SmartScreen Filter communicates with a site on the Internet

Controlling SmartScreen Filter to limit the flow of information to and from the Internet

This section explains how SmartScreen Filter communicates across the Internet, and it explains steps to take to limit, control, or prevent that communication in an organization with many users.

Benefits and purposes of SmartScreen

The SmartScreen Filter provides an early warning system to notify users of suspicious websites that could be engaging in phishing attacks or distributing malware through a socially engineered attack.

Note

SmartScreen Filter is one of the multiple layers of defense in the anti-phishing and malware protection strategies developed by Microsoft. For more information, see What is SmartScreen Filter? on the Microsoft website.

The following list describes the enhancements that SmartScreen Filter provides:

- **Anti-phishing and anti-malware support.** The SmartScreen Filter helps protect users from sites that are reported to host phishing attacks or distribute malicious software through socially engineered attacks. This protection is URL reputation-based, which means that it evaluates the URLs to determine whether they are known to distribute or host unsafe content. SmartScreen Filter also provides application reputation checks, which check the reputation of a downloaded program itself, or the digital signature that is used to sign a file. If the file or certificate has an established reputation, no warnings are shown. If the file does not have an established reputation, the user is at higher risk of malware infection and is shown a more

severe warning. The reputation-based analysis in SmartScreen Filter is an additional layer of protection to help protect against malicious software.

- **Heuristics and enhanced telemetry.** New heuristics combined with enhanced telemetry allow SmartScreen to identify and warn users about malicious sites more quickly.
- **Group Policy support.** A group policy setting can be used to keep the user from managing SmartScreen Filter. If you enable this policy setting, the user is not prompted to turn on SmartScreen Filter. All website addresses that are not on the filter's allow list are sent automatically to Microsoft without prompting the user.If you disable or do not configure this policy setting, the user is prompted to decide whether to turn on SmartScreen Filter during the first-run experience. For more information, see To Control SmartScreen Filter by using Group Policy later in this document.

Overview: Using SmartScreen Filter in a managed environment

In a managed environment, you can use Group Policy to control SmartScreen Filter in a variety of ways, including the following:

- Prevent users from overriding or clicking through SmartScreen Filter warnings.
- Turn on SmartScreen Filter so that it runs automatically.
- Turn off SmartScreen Filter.

For details, see Controlling SmartScreen Filter to limit the flow of information to and from the Internet later in this section.

How SmartScreen Filter communicates with a web service on the Internet

This subsection describes how SmartScreen Filter might communicate with a site on the Internet as it evaluates a website URL that a user is trying to reach.

- **Default settings**: SmartScreen Filter is disabled unless the feature is enabled by the user or through a Group Policy setting.
- **Triggers**: When the user visits an Internet site, the URL of the site is compared against a list of high traffic websites that is built into SmartScreen Filter. If the URL matches a site on the list, no further checks occur for that URL. If the URL does not match a site on the list and SmartScreen Filter is enabled, SmartScreen Filter sends a query to the Microsoft SmartScreen URL Reputation Service (URS). If the URS indicates that a URL is reported to be unsafe, a message is shown to warn the user about entering personal information or downloading malware. Occasionally, a telemetry report containing additional information about the site may be sent to help improve the quality of the SmartScreen services. When a program is downloaded, an application reputation check may be made which requires information about the downloaded file to be sent to the SmartScreen Application Reputation service.
- **Specific information sent**: The following information is sent over an encrypted (HTTPS) connection to the SmartScreen services:

    - **URL**: The full request URL is included only when the site is required to be checked by the URS.
    - **Detailed software version information**: The browser version, the SmartScreen Filter version, and the version of the “high traffic site” list.
    - **Detailed information about the URL**: IP hosting the site, frame URLs, heuristics results, basic network details.
    - **Downloaded file information**: When an application reputation check is made, the download URL, a hash of the full file, information about the digital signature, and some additional data about the downloaded file, such as file size and the hosting IP, is sent.
    - **Operating system version**: The version of Windows that the browser is installed on.
    - **Language and locale setting for the browser**: The language and locale for the browser display, for example, English (United States).
    - **Anonymous statistics about how often SmartScreen Filter is triggered**: SmartScreen Filter tracks basic statistics, such as how often a warning is generated and how often a query is made to the URL Reputation Service. This statistical information is sent to Microsoft periodically, and it is used to analyze the performance and improve the quality of the SmartScreen Filter.
- **User notification**: If SmartScreen Filter is enabled, the user is not notified when SmartScreen Filter performs a check and is notified if SmartScreen Filter detects a URL that is reported as unsafe.
- **Logging**: By default, SmartScreen Filter does not log events. However, if you use the Application Compatibility Toolkit to enable logging for application compatibility events, SmartScreen Filter logs an event when a warning is shown for a website.

  For information, see Application Compatibility.
- **Encryption**: All information that is sent to SmartScreen services is encrypted by using the HTTPS protocol.
- **Access**: The teams that maintain SmartScreen Filter and the URL Reputation Service have access to the data that is sent to the SmartScreen services (including the anonymous statistics described earlier in this list).
- **Privacy**: URLs that are collected may unintentionally contain personal information (depending on the design of the website that is visited). Like the other information that is sent to Microsoft, this information is not used to identify, contact, or target advertising to users. In addition, Microsoft filters address strings to remove personal information where possible.
- **Transmission protocol and port**: The transmission protocol for any information that is transmitted to the URL Reputation Service is over HTTPS using port 443.
- **Ability to disable**: SmartScreen Filter can be disabled through the user interface or through Group Policy.

Controlling SmartScreen Filter to limit the flow of information to and from the Internet

This subsection provides information about how to control settings for SmartScreen Filter.

To Control SmartScreen Filter by using Group Policy

1. Using an account with domain administrator credentials, open the Group Policy Management Console (GPMC) by running **gpmc.msc**, and edit an appropriate Group Policy Object (GPO).
2. If you want the Group Policy setting to apply to all users of a computer and to come into effect when the computer starts or when Group Policy is refreshed, expand **Computer Configuration**. If you want the Group Policy setting to apply to users and to come into effect when users sign in or when Group Policy is refreshed, expand **User Configuration**.
3. Expand **Policies** (if present), expand **Administrative Templates**, expand **Windows Components**, and then click **Internet Explorer**.
4. In the details pane, double-click **Prevent managing SmartScreen Filter**, click **Enabled** (which means that users cannot control SmartScreen Filter settings), and then choose one of the following settings for **Select SmartScreen filter mode**:
   - **On**: Automatic SmartScreen Filter is always turned on in Security Zones for which the feature is Enabled.
   - **Off**: SmartScreen Filter does not automatically perform reputation checks. Users can manually trigger a check by using the **Safety** menu.

     Note
     Disabling the **Turn off Managing SmartScreen filter** Group Policy setting does not disable SmartScreen Filter. Users can control SmartScreen Filter settings on a local computer.

Additional references

- What is SmartScreen Filter?
- Security and Control
- Secure Windows Server 2012

Managing Privacy: User Access Logging and Resulting Internet Communication

In this section

Benefits and purposes of User Access Logging
User and device-related data recorded with User Access Logging
Viewing or changing settings that affect User Access Logging
Additional references

This section provides overview information about User Access Logging and information about some settings that affect User Access Logging. The section also provides suggestions for other sources of information about User Access Logging to help you balance your organization's requirements for communication across the Internet with your organization's requirements for protection of networked assets. It is beyond the scope of this document to describe all aspects of maintaining appropriate levels of privacy and security in an organization running servers that use User Access Logging.

Benefits and purposes of User Access Logging

User Access Logging aggregates unique client device and user request events that are logged into a local database. These records are made available (through a query by a server administrator) to retrieve quantities and instances by server role, by user, by device, by the local server, and by date. In addition, User Access Logging has been extended to enable non-Microsoft software developers to instrument User Access Logging events that are to be aggregated by the server. No data collected by User Access Logging is sent to Microsoft.

This information can be useful to server administrators at all levels. User Access Logging can assist server administrators in performing the following tasks:

- Quantify client user requests for local physical or virtual servers.
- Quantify client user requests for installed software products on a physical computer or virtual machine.
- Retrieve User Access Logging data from multiple remote physical or virtual servers.

Important
User Access Logging is not recommended for use on servers that are connected directly to the Internet, such as web servers on an Internet-accessible address space; and it is not recommended in scenarios where extremely high performance is the primary function of the server (such as in high-performance computing workload environments). User Access Logging is primarily intended for small, medium, and enterprise intranet scenarios where high volume is expected, but not as high as many deployments that serve Internet-facing traffic volume on a regular basis.

User and device-related data recorded with User Access Logging

The following user-related data is logged with User Access Logging.

Data	Description
ActivityCount	Number of times a particular user has accessed the service.
FirstSeen	Date and time when a user first accesses a role

Data	Description
	or service.
LastSeen	Date and time when a user last accessed a role or service.
ProductName	Name of the software parent product (such as Windows) that is providing User Access Logging data.
RoleGUID	GUID that is assigned or registered by User Access Logging, which represents the server role or installed product.
RoleName	Name of the role, component, or subproduct that is providing User Access Logging data. This is also associated with a ProductName and a RoleGUID.
TenantIdentifier	Unique GUID for a tenant client of an installed role or for a product that accompanies the User Access Logging data, if applicable.
UserName	User name on the client that accompanies the User Access Logging entries from installed roles and products, if applicable.
PSComputerName	Name of the target server when you query User Access Logging data from a remote computer.

The following device-related data is logged with User Access Logging.

Data	Description
ActivityCount	Number of times a particular device has been used to access the service.
FirstSeen	Date and time when an IP address is first used to access a role or service.
IPAddress	IP address of a client device that is used to access a role or service.
LastSeen	Date and time when an IP address was last used to access a role or service.

Data	Description
ProductName	Name of the software parent product (such as Windows) that is providing User Access Logging data.
RoleGUID	GUID that is assigned or registered by User Access Logging, which represents the server role or installed product.
RoleName	Name of the role, component, or subproduct that is providing User Access Logging data.
TenantIdentifier	Unique GUID for a tenant client of an installed role or for a product that accompanies the User Access Logging data, if applicable.
PSComputerName	Name of the target server when you query User Access Logging data from a remote computer.

Viewing or changing settings that affect User Access Logging

You can disable or enable User Access Logging, and collect and delete data that is recorded by using User Access Logging. For more information, see Manage User Access Logging.

Additional references

User Access Logging Overview

Manage User Access Logging

Software Inventory Logging

Manage Privacy: Windows Customer Experience Improvement Program and Resulting Internet Communication

In this section

Purpose of the Windows Customer Experience Improvement Program

Overview: Using the Windows Customer Experience Improvement Program in a managed environment

How the Windows Customer Experience Improvement Program communicates with a site on the Internet

Procedures for controlling the Windows Customer Experience Improvement Program

This topic describes how the Windows Customer Experience Improvement Program communicates across the Internet, and it explains steps to take to limit, control, or prevent that communication in an organization with many users.

Purpose of the Windows Customer Experience Improvement Program

The Windows Customer Experience Improvement Program (CEIP) is a voluntary program that collects information about how people use Windows. CEIP collects information about configuration settings, hardware configurations and usage, and user's encounters with the operating system without interrupting their tasks at the computer. The information that is collected helps Microsoft improve the features that are used most often and create solutions to common issues.

Overview: Using the Windows Customer Experience Improvement Program in a managed environment

In a managed environment, the Windows Customer Experience Improvement Program runs only if an administrator chooses to participate. You might decide to disable the program on all computers. You can do this by using Group Policy or by using an answer file with an unattended installation.

If you would like your computers to participate in the CEIP program, you can also use Group Policy to redirect data from the Windows Customer Experience Improvement Program to a Windows Server on your network with the Windows Feedback Forwarder or similar software that is designed to collect data from the Windows Customer Experience Program. For more information, see Procedures for controlling the Windows Customer Experience Improvement Program later in this section.

How the Windows Customer Experience Improvement Program communicates with a site on the Internet

The Windows Customer Experience Improvement Program communicates with a site on the Internet as follows:

- **Specific information sent**: The information that is sent includes details about the computer hardware configuration (such as the number of processors and screen resolution),

performance and reliability (such as how quickly a program responds when you click a button), and information about use of the system (such as how many folders a user typically creates on the desktop). It also includes information about the use of features such as Event Viewer and Remote Assistance. For additional details, see the link to the privacy statement later in this list.

- **Default setting**: By default, the Windows Customer Experience Improvement Program is turned off.
- **Triggers**: Data for the Windows Customer Experience Improvement Program is collected over time and sent periodically. However, data is not collected or sent if the computer is on battery power, and no attempt to send data is made if the computer is not connected to a network.
- **User notification**: After an administrator chooses to participate in the program, there are no notifications. You are not prompted or interrupted in any way when data is collected or sent.
- **Logging**: Events are logged in Event Viewer in **Windows Logs\Application**.
- **Encryption**: Data about software usage is encrypted during transmission by using HTTPS (that is, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) with HTTP).
- **Access and privacy**: Data from the Windows Customer Experience Improvement Program is stored on servers in Microsoft-controlled facilities. Microsoft uses the data to identify trends and usage patterns in Microsoft software and to improve Microsoft products and services. For additional information, see the Windows 8 and Windows Server 2012 Privacy Statement or the Windows 8.1 and Windows Server 2012 R2 privacy statement.
- **Transmission protocol and port**: The transmission protocol is HTTPS and the port is 443.
- **Ability to disable**: You can disable the Windows Customer Experience Improvement Program on an individual computer from the **Customer Experience Improvement Settings** page. You can also disable it by using Group Policy or an answer file with an unattended installation.

Procedures for controlling the Windows Customer Experience Improvement Program

The following procedures explain how to make changes to how the Windows Customer Improvement Program works using Server Manager, Group Policy settings and unattended installation answer file settings.

Note

To change Windows Customer Experience Improvement Program settings, you must be logged on as an administrator.

To view or change the Windows CEIP setting on one or more servers using Server Manager

1. Open **Server Manager**, and click the **All Servers** menu in the navigation pane.
2. In the details pane, select one or more servers.

3. Right-click the selected servers, and select **Configure Windows Automatic Feedback**.
4. In the **Windows Automatic Feedback** dialog, make any necessary changes.
5. Click **OK** to apply the settings and close the dialog box.

To disable the Windows Customer Experience Improvement Program by using Group Policy

1. Using an account with domain administrative credentials, sign in to a computer with the Group Policy Management feature installed. Then open the Group Policy Management Console (GPMC) by running **gpmc.msc**, and edit an appropriate Group Policy Object (GPO).
2. Expand **Computer Configuration**, expand **Policies** (if present), expand **Administrative Templates**, expand **System**, expand **Internet Communication Management**, and then click **Internet Communication settings**.
3. In the details pane, double-click **Turn off Windows Customer Experience Improvement Program**, and then click **Enabled**.

   When you enable this setting, all administrators and users to which the GPO applies are opted out of the Windows Customer Experience Improvement Program.

   You can also restrict Internet access for this and a number of other features by applying the **Restrict Internet communication** policy setting, which is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category.

To disable the Windows Customer Experience Improvement Program by using an answer file with an unattended installation

1. Using the methods you prefer for an unattended installation or a remote installation, create an answer file. For more information about unattended and remote installations, see Appendix A: Resources for Learning About Automated Installation and Deployment.
2. Confirm that your answer file includes the following line:

```
<CEIPEnabled>0</CEIPEnabled>
```

You can also redirect data from multiple computers to a computer on your network. To do this, you need to have a computer running Windows Server with the Windows Feedback Forwarder installed.

To install Windows Feedback Forwarder By using the user interface

1. On the server, open **Server Manager**, click **Manage**, and then click **Add Roles and Features**.
2. In the **Add Roles and Features Wizard**, select the server, and click **Features** in the navigation pane. Scroll down, click **Windows Feedback Forwarder**, and then click **Next** twice.

3. Specify an incoming port number (the default port number is 53533). If the domain has an Internet proxy server, specify the proxy information. Click **Next**, and then click **Install**. When the installation is complete, click **Close**.
4. In Server Manager, click **All Servers** in the navigation pane.
5. In the details pane, right-click the server that you installed Windows Feedback Forwarder on, and then click **Windows Feedback Forwarder Configuration**. Keep the dialog box open, and then continue to the procedure To deploy the Windows Feedback Forwarder Group Policy .

To install Windows Feedback Forwarder by using Windows PowerShell

1. Open Windows PowerShell and run the following command:

   Add-WindowsFeature WFF
2. In Server Manager, click **All Servers** in the navigation pane.
3. In the details pane, right-click the server that you installed Windows Feedback Forwarder on, and then click **Windows Feedback Forwarder Configuration**.
4. Click the **Forwarding Settings** tab, and specify an incoming port number (the default port number is 53533). If the domain has an Internet proxy server, specify the proxy information, and then click **Apply**.

   Keep the dialog box open and continue to the procedure To deploy the Windows Feedback Forwarder Group Policy .

To deploy the Windows Feedback Forwarder Group Policy

1. With the Windows Feedback Forwarder configuration dialog box open, click the **Group Policy** tab.

   Note
   You may have to enter credentials, depending on the settings for the current user.
2. Enter the name of the domain that you want to deploy the Group Policy to, and then click **Find**.
3. After the list of organizational units is populated, select one or more organizational units, and then click **Apply**.

Manage Privacy: Windows Defender and Resulting Internet Communication

In this section

Benefits and purposes of Windows Defender and the online Microsoft Active Protection Service community

Overview: Using Windows Defender and information from the MAPS community in a managed environment

How Windows Defender communicates with Internet sites without a MAPS membership

How Windows Defender communicates with Internet sites when combined with MAPS

Procedures for configuring Windows Defender

Additional references

This section discusses how Windows Defender communicates across the Internet, and it explains steps to take to limit, control, or prevent that communication in an organization with many users.

Benefits and purposes of Windows Defender and the online Microsoft Active Protection Service community

Windows Defender

With Windows Defender, users can be alerted when spyware, malware, adware, random toolbars and such attempts to install or run on their computers. Windows Defender also alerts users when programs attempt to change important settings in the Windows operating system.

With Windows Defender, users can schedule scans on a regular basis, and they can be alerted to harmful software that is detected or removed during the scan.

Windows Defender receives updates to respond to evolving malicious and potentially unwanted software. It is designed to automatically update by using the Windows Update service. It can also be updated from a WSUS server in an environment with Windows Server Update Services (WSUS). The following list briefly describes how Windows Defender obtains updates:

- If Windows Defender is enabled, by default it checks for software updates and updated definitions of spyware and other potentially unwanted software before each scheduled scan. It checks for these updates on the Windows Update service (or in an environment with WSUS, it checks a WSUS server). This check for updates helps ensure that Windows Defender uses the latest available software and definitions when scanning.

  Scheduled scans occur daily by default, so these checks for software updates also occur daily by default.

- Through commands on the **Help** menu, the user can request that Windows Defender check immediately for updated definitions. Users can also view a web-based privacy statement.

For more details about how Windows Defender checks for software updates, see How Windows Defender communicates with Internet sites without a MAPS membership later in this section.

The online Microsoft Active Protection Service community

The online Microsoft Active Protection Service (MAPS) community is designed to help Microsoft continually update and improve definitions of malware, spyware and other potentially unwanted software and to help Microsoft improve Windows Defender and related technologies.

New types and versions of potentially unwanted software are emerging regularly, so MAPS reports help Microsoft researchers discover new threats more rapidly and determine which software to investigate. For example, if many people remove software that has not yet been classified, Microsoft analyzes that software to see if it should be included in future definitions. MAPS uses these definitions to detect and block malware.

Joining the online Microsoft Active Protection Service community is optional, but it is recommended. When the computer is first started, prompts appear that recommend steps that can help protect the computer. These include joining the online Microsoft Active Protection Service community.

Overview: Using Windows Defender and information from the MAPS community in a managed environment

In a managed environment, Windows Defender can help keep potentially unwanted software off of users' computers and help prevent potentially unwanted software from causing issues. Membership in the online MAPS community can provide additional information that might be useful when you are making decisions about questionable software.

If you choose to use other solutions to defend against malware and other potentially unwanted software, you can configure Windows Defender to:

- Use Group Policy settings to prevent users from running Windows Defender.
- Use Group Policy settings to limit access to resources such as the online MAPS community by allowing only designated people to become members.
- Check your WSUS servers for updates. (You must have WSUS set up in your environment for this option.) If the WSUS servers are unavailable, Windows Defender checks the Windows Update website to ensure that it is using the latest definitions when scanning.

  For more information, see Windows Server Update Services.

How Windows Defender communicates with Internet sites without a MAPS membership

The following list describes how Windows Defender communicates with sites on the Internet when users do not have membership in the online MAPS community. Communication that results with Basic or Advanced membership in the online MAPS community is described in the next section.

When enabled by itself, Windows Defender communicates with sites on the Internet as follows:

- **Specific information received**: The following list describes the information that is received in specific situations:
  - **Each time Windows Defender performs a scheduled scan** (if there is a connection to the Internet). By default Windows Defender checks the Windows Update website for software updates and updated definitions. This is the same process that is used to check for updates for other operating system features, which means that the information sent includes the version of the current set of definitions. If updates are available, they are downloaded by Windows Defender.
  - **When the user clicks Help options, and then clicks Check for updates**. Windows Defender performs the same check described in the previous item.
  - **When the user clicks Help options, and then clicks View Privacy Statement Online**. The privacy statement is displayed:

    Windows 8 and Windows Server 2012 Privacy Statement or Windows 8.1 and Windows Server 2012 R2 privacy statement
- **Default settings**: If Windows Defender is enabled, by default it scans the computer daily. (A prompt that recommends Windows Defender be enabled is displayed the first time the computer is started after setup.)
- **Triggers**: When Windows Defender performs a scheduled scan, by default it also searches the Windows Update Web servers for the latest definition file. To cause Windows Defender to check immediately for updates or to display the privacy statement online, the user must click the Help options that are offered.
- **User notification**: When a scan is in progress and the Windows Defender interface is open, status about the scan is displayed. Also when a scan is in progress, the user can click the Windows Defender icon in the notification area to view status.
- **Logging**: Windows Defender logs the following types of information on the local computer:
  - Events are logged in Event Viewer in the System log.
  - Update failures are logged to *systemroot***Temp\Mpsigstub.log**.
  - Actions taken to protect against spyware or potentially unwanted software are logged in the same location as other events for that software.
- **Encryption**: Windows Defender uses the same encryption methods as Windows Update, which means initial data is transferred by using HTTPS, and updates are transferred by using HTTP.
- **Access**: Microsoft staff maintains the functionality of the Windows Update Web servers, and as part of maintaining the servers, they monitor the version information that Windows Defender sends when it checks for updates.
- **Privacy**: To view the privacy statement, see Windows 8 and Windows Server 2012 Privacy Statement or Windows 8.1 and Windows Server 2012 R2 privacy statement.
- **Transmission protocol and port**: Windows Defender uses the same transmission protocols and ports as Windows Update: HTTP with port 80 and HTTPS with port 443.
- **Ability to disable**: You can disable Windows Defender through Control Panel or Group Policy.

How Windows Defender communicates with Internet sites when combined with MAPS

The following list describes communication that results from using Windows Defender with membership in the online MAPS community. When a user has joined the online MAPS community, Windows Defender communicates with sites on the Internet as follows:

- **Specific information sent**: The following list describes the information that is sent depending on the level of membership in MAPS. The information is sent whenever Windows Defender detects software that has not been analyzed for risks or malware:
  - **For Basic members**: The report that is sent by Windows Defender to the MAPS website includes the following information:

    About the computer: A randomly generated, globally unique identifier (GUID) that is used to uniquely identify the computers of MAPS members when they communicate with the MAPS website. (Windows Defender creates the GUID unless the operating system was upgraded from Windows XP, in which case the GUID might have been created previously by the Microsoft Malicious Software Removal Tool running on Windows XP.) This GUID does not contain any personal information.

    Information collected about the computer also includes the operating system name and version (including any service packs that have been applied), the web browser software and version, and identifiers for the country or region and locale. In addition, the report might contain information related to the possible presence of spyware or other potentially unwanted software—for example, information about registry key entries that control actions such as automatically starting an application when the system starts.

    About the software in question: This information includes the file name, size, date stamps, and where applicable, vendor and cryptographic hashes. In addition, full URLs can be collected that indicate the origin of the file. Windows Defender attempts to filter out personal information in the URL and in the **fil** paths for Basic members. The report can also include the action that the user chose to take when the program was detected (Block or Allow).

    Note
    The user's membership in MAPS means that the user might sometimes see a pop-up request for a Sample Submission report. This report requests specific files that Microsoft suspects might be potentially unwanted software on a computer, and these files are used for further analysis. The report is sent only if the user consents.
  - **For Advanced members**: The report that is sent to the MAPS website includes the information that is sent with a Basic membership, plus additional details about the software in question including file paths and partial memory dumps. These file paths and partial memory dumps might unintentionally contain personal information. If any personal information is included in a report, the information is not used to identify or contact a user.

    Note

The user's membership in MAPS means that the user might sometimes see a pop-up request for a Sample Submission report. This report requests specific files that Microsoft suspects might be potentially unwanted software on a computer, and these files are used for further analysis. The report is sent only if the user consents.

- **Default settings**: If a person opts-in to MAPS during the Windows Defender configuration process, the membership is a Basic membership by default.
- **Triggers**: When Windows Defender detects software that has not been analyzed for risks (that is, software not previously categorized in the Windows Defender definition file) and the user is a member of MAPS, Windows Defender sends a report about the software in question.
- **User notification**: For Basic MAPS members, the user notification is the same as for anyone using Windows Defender. For more information, see How Windows Defender communicates with Internet sites without a MAPS membership earlier in this section.

  For Advanced MAPS members, if software is present that has not yet been classified for risk, and it attempts to change computer settings, a prompt asks whether to allow or block the change. (For users who are Basic MAPS members, such software is not blocked.)
- **Logging**: Logging for Windows Defender does not change when the user is a MAPS member. For more information, see How Windows Defender communicates with Internet sites without a MAPS membership earlier in this section.
- **Encryption**: Windows Defender uses Secure Sockets Layer (SSL) to encrypt the information that it sends to MAPS.
- **Access**: MAPS reports are used to improve Microsoft software and services. The reports may also be used for statistical or other testing or analytical purposes, trending, and signature generation. Only Microsoft employees, contractors, and vendors who have a business need to use the reports are provided access to them.
- **Privacy**: To view the privacy statement, which covers MAPS, see Windows 8 and Windows Server 2012 Privacy Statement or Windows 8.1 and Windows Server 2012 R2 privacy statement.
- **Transmission protocol and port**: When Windows Defender sends information to MAPS, it uses HTTPS with port 443.
- **Ability to disable**: A user can decline or end membership in MAPS from an individual computer and an administrator can prevent users from being members by using a Group Policy setting.

Procedures for configuring Windows Defender

This subsection provides procedures for:

- Viewing or changing Windows Defender settings, including MAPS settings.
- Disabling Windows Defender by using Group Policy.
- Preventing MAPS membership by using Group Policy.

To view or change Windows Defender and MAPS settings

1. Open **Control Panel**, and then click **Windows Defender**.
2. Click **Settings**, and then click **MAPS**.
3. View or change the settings, and then click **Cancel** or **Save changes**.

To disable Windows Defender by using Group Policy

1. Using an account with domain administrative credentials, sign, open Group Policy Management Console (GPMC) or Group Policy Object Editor, and then edit an appropriate Group Policy Object (GPO).
2. Expand **Computer Configuration**, expand **Administrative Templates**, expand **Windows Components**, and then click **Windows Defender**.
3. In the details pane, double-click **Turn off Windows Defender**, and then click **Enabled**.

   Note
   If this Group Policy setting is enabled, the user can still click the command to open Windows Defender. However, Windows Defender displays a pop-up window that says it is turned off by Group Policy.

To prevent Windows Active Protection Service membership by using Group Policy

1. Using an account with domain administrative credentials, sign, open the Group Policy Management Console by running **gpmc.msc**, and then edit an appropriate GPO.
2. Expand **Computer Configuration**, expand **Administrative Templates**, expand **Windows Components**, and then click **Windows Defender**.
3. **For Windows 8:** In the details pane, double-click **Configure Microsoft Active Protection Service Reporting**, click **Enabled**, and then click **No Membership**.

   Important
   To prevent Microsoft Active Protection Service reporting, do not disable this setting. You can only block Microsoft Active Protection Service reporting by enabling this setting and then choosing **No Membership**.

   For Windows 8.1: In the navigation pane, expand **MAPS** and then in the details pane, double-click **Join Microsoft MAPS** and choose **Disabled**.

Additional references

For more information, see the following Microsoft websites:

- How Microsoft antimalware products identify potentially unwanted software
- Security and Control
- Microsoft Windows Server Update Services

Manage Privacy: Windows Error Reporting and Resulting Internet Communication

In this section

Benefits and purposes of Windows Error Reporting and the Problem Reports and Solutions feature

Overview: Using Windows Error Reporting and the Problem Reports and Solutions feature in a managed environment

How Windows Error Reporting communicates with an Internet site

Controlling Windows Error Reporting to prevent the flow of information to and from the Internet

Procedures to configure Windows Error Reporting

This section explains how the Windows Error Reporting and the Problem Reports and Solutions feature communicates across the Internet, and it explains steps to limit, control, or prevent that communication in an organization with many users.

Note

The Problem Reports and Solutions feature in Action Center is an interface that displays information from Windows Error Reporting. It communicates with the Internet only through Windows Error Reporting.

Benefits and purposes of Windows Error Reporting and the Problem Reports and Solutions feature

Windows Error Reporting and the Problem Reports and Solutions feature work together to make it easy to find online solutions for computer issues:

- **Windows Error Reporting**: Windows Error Reporting is a feature that allows Microsoft to track and address errors that are relating to the operating system, and Windows features and applications. Windows Error Reporting gives administrators the opportunity to send data about errors to Microsoft and to receive information about solutions.

  Solution information can include instructions for working around an issue, or a link to the Windows Update website or another website for updated drivers, patches, or Microsoft Knowledge Base articles. Microsoft developers can use Windows Error Reporting as a problem-solving tool to address customer issues in a timely manner and to improve the quality of Microsoft products.
- **Problem Reports and Solutions**: The Problem Reports and Solutions feature in Action Center helps you track problem reports and solution information that you have received from Microsoft. Action Center helps you store the solution information, which is displayed by using a web browser. However, all Internet communication that is related to the problem reports and solutions is handled by Windows Error Reporting.

Consent levels in Windows Error Reporting

Windows Error Reporting has the following consent levels to help you control how Windows Error Reporting prompts you before sending data:

- **Automatically check for solutions.** Windows Error Reporting sends the minimum data that is required to check for an existing solution, for example, the application name and version, module name and version, and exception code. After sending this data, Windows Error Reporting prompts you for consent before sending any additional data that is needed to solve the issue.
- **Automatically check for solutions and send additional report data, if needed.** Windows Error Reporting automatically checks for solutions and sends additional information that is needed to solve the issue (typically, the user is not prompted).
- **Each time a problem occurs, ask me before checking for solutions.** Windows Error Reporting always prompts for consent before sending an error report.
- **Never check for solutions (not recommended).** This setting disables Windows Error Reporting.
- **Send all data (Group Policy setting only).** This setting can only be configured through Group Policy, not through the Initial Configuration Tasks interface, Server Manager, or Control Panel. Any data that is requested by Microsoft is sent without prompting the user.

Options for controlling Windows Error Reporting

If a prompt appears for someone who is signed in as an administrator, the person can choose to report application and operating system errors. If a prompt appears for someone who is not signed in as an administrator, the person can choose to report application errors plus errors for operating system software that does not require administrative credentials to run.

Error reporting can be controlled through the Initial Configuration Tasks interface or Server Manager, as outlined in Consent levels in Windows Error Reporting earlier in this section.

Overview: Using Windows Error Reporting and the Problem Reports and Solutions feature in a managed environment

In a managed environment, you can choose to disable Windows Error Reporting or control it as follows:

- You can use Group Policy or an answer file for an unattended installation to control the consent level (described earlier) to determine the amount of prompting that users or administrators see before information about a software issue is sent to Microsoft. For example, you can set the consent level so the person using the computer is always prompted before information is sent.
- You can use Group Policy to disable Windows Error Reporting.

For more information about Microsoft software that is designed for use with the Group Policy setting, **Configure Corporate Windows Error Reporting**, see System Center 2012 Configuration Manager on the Microsoft website.

For more information about the underlying functionality that redirects error reports to a server on your intranet, see WER Settings on the MSDN website.

For more information about the answer-file entries or Group Policy settings that are described in this subsection, see Controlling Windows Error Reporting to prevent the flow of information to and from the Internet later in this section.

How Windows Error Reporting communicates with an Internet site

The data that Microsoft collects through Windows Error Reporting is used strictly for the purpose of tracking and solving issues that users and administrators are experiencing. This subsection describes various aspects of the data that is sent to and from the Internet during error reporting, and how the exchange of information takes place. The next subsection provides additional details.

- **Specific information sent or received**: In most cases, the information that is collected for an error report only includes software parameters, which include such information as the application name and version, module name and version, and exception code. In unusual cases, a more complete crash report might be collected. Rarely, some information that is unique to the person who is using the computer might be collected unintentionally. This information, if present, is not used to identify the person.

  Microsoft may send solution information about an issue to the user or administrator, including links to websites.
- **Default settings**: By default, error reporting is enabled. However, additional configuration steps are needed to configure error reporting, and no reports are sent unless these steps are completed.

  When a computer is started for the first time, the Initial Configuration Tasks interface appears which displays a variety of tasks including **Enable automatic updating and feedback**. In this task, you can choose to enable a default level of automatic updating and feedback (which includes error reporting), or you can manually configure settings. For details about consent levels, see Consent levels in Windows Error Reporting earlier in this section.
- **Triggers**: The opportunity to send an error report is triggered by application or system errors.
- **User notification**: User notification depends on the consent level. See Consent levels in Windows Error Reporting earlier in this section.

  Windows provides reminders (in the form of pop-up notifications) to check for solutions to reports that have not been sent, for example, reports that were generated in the background or while you were offline.
- **Logging**: Descriptions of system and application errors are recorded in the event log. In addition, the Problem Reports and Solutions feature records information about problem reports sent and solution information received on that computer, so that the user or

administrator can investigate solutions later. (New solutions might overwrite old solutions if the number of stored solutions exceeds the allowed maximum.)

- **Encryption**: All report data that could include personal information is encrypted during transmission by using HTTPS (that is, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) with HTTP). The software parameters information, which includes such information as the application name and version, module name and version, and exception code, is also encrypted.
- **Access**: Microsoft employees and contingent staff may access the error reports to maintain Windows Error Reporting or to improve Microsoft products. They may not use the reports for other purposes.

  If the error report indicates that one or more non-Microsoft products were involved in causing the issue, Microsoft may send the report to the respective companies if the companies agree to abide by the terms of the privacy statement. Software or hardware developers (employed by Microsoft or one of its partners) may analyze the fault data and try to identify and correct the problem.
- **Privacy**: For more information, see the Windows 8 and Windows Server 2012 Privacy Statement or the Windows 8.1 and Windows Server 2012 R2 privacy statement.

  Details that are related to privacy of data are presented in Types of data collected later in this section.
- **Transmission protocol and port**: The transmission protocol is MS_SQMCS2. For more information, see [MS-SQMCS2]: Software Quality Metrics (SQM) Client-to-Service Version 2 Protocol Specification.
- **Ability to disable**: The feature can be disabled through Group Policy or on an individual computer. You can also control the feature as described in Overview: Using Windows Error Reporting and the Problem Reports and Solutions feature in a managed environment earlier in this section, and Controlling Windows Error Reporting to prevent the flow of information to and from the Internet later in this section.

Types of data collected

This section provides an overview of the data that Windows Error Reporting collects and information about data that might be collected from four sources:

- Application errors
- Handwriting recognition errors
- Japanese Input Method Editor errors
- Windows kernel failures

Overview of the data that Windows Error Reporting collects

Windows Error Reporting collects information about the computer configuration, what the software was doing when the problem occurred, and other information directly related to the issue. Windows Error Reporting does not intentionally collect anyone's name, address, email address, or computer name. It is possible that such information may be captured in memory or in the data that is collected from open files, but Microsoft does not use it to identify users. Windows

Error Reporting collects Internet Protocol (IP) addresses, but the addresses are not used to identify users, and in many cases, they are the address of a network address translation (NAT) computer or proxy server, not a specific client behind that NAT computer or proxy server.

IP address information is used in aggregate by the operators who maintain the servers that receive error reports.

In rare cases, such as issues that are especially difficult to solve, Microsoft may request additional data, including sections of memory (which may include memory that is shared by any or all applications that were running at the time the issue occurred), some registry settings, and one or more files from your computer. When additional data is requested, users can review the data and choose whether to send the information.

Data collected from application errors

Any application can be written in a way that uses the Error Reporting functionality. If an application error occurs for which Error Reporting is available, and you choose to send the report, the following information is included:

- The digital product ID, which can be used to identify your license.
- Information regarding the condition of the computer and the application at the time the error occurred. This may include data that is stored in memory and stacks, information about files in the application's directory, the operating system version, and the computer hardware in use.

You can use a registry setting to configure Windows Error Reporting so that it collects full user-mode dumps and stores them locally after a user-mode application stops responding. This configuration option in Windows Error Reporting does not involve communication across the Internet.

For more information, see Collecting User-Mode Dumps on the Microsoft website.

Data collected from handwriting recognition errors

If users encounter a handwriting recognition error while using the Tablet PC Input Panel, they can start the error reporting tool and then select recently corrected handwriting samples to send in an error report. The samples are handled according to the consent-level setting, and in most cases, they are sent only with explicit consent. No personal information is intentionally collected; however, the samples that are chosen may include personal information. This information will not be used to personally identify the person.

You can disable the reporting of handwriting recognition errors by using a specific Group Policy setting, as described in Setting for disabling Windows Error Reporting later in this section.

Data collected from the Japanese Input Method Editor

In the Japanese versions of Windows, users can generate a "word registration report" through the Japanese Input Method Editor (IME), and then choose to send the report to Microsoft. The reports are like error reports, but they record a word or word pair to improve the selection of the ideograms that are displayed. Word registration reports can include the information that is

provided in the Add Word dialog box about the words being reported, and the software version number for IME. Each time such a report is generated, the user is asked whether to send the report to Microsoft and can view the information that is contained in the report before sending it.

Microsoft uses the information to help improve IME. Personal information might unintentionally be collected, but Microsoft does not use the information to identify or contact the person. Word registration reports are sent to Microsoft by using HTTPS.

If you configure Windows Error Reporting as described in Procedures to configure Windows Error Reporting later in this section, you can control word registration reports in the same way that you control error reports.

Data collected from Windows kernel failures

When a kernel-mode (system) error occurs a Stop message is displayed and diagnostic information is written to a memory dump file. When someone restarts the computer by using normal mode or Microsoft Windows Safe Mode (with networking), and then signs in as an administrator, Windows Error Reporting responds. As with other errors, Windows Error Reporting uses the consent-level setting to determine when to prompt you before sending a kernel fault report.

Windows kernel fault reports contain information about what the operating system was doing when the problem occurred. These event reports contain the minimum information that can help identify why the operating system stopped unexpectedly. If you choose to send the report, it includes the following:

- Operating system name (for example, Windows Server 2012)
- Operating system version
- Operating system language as represented by the locale identifier (LCID)—(for example, the standard international numeric abbreviation, 1033, for United States English)
- Loaded and recently unloaded drivers. These identify the modules that were in use by the kernel when the Stop error occurred and the modules that were used recently.
- List of drivers in the **Drivers** folder on the hard disk drive (*systemroot*\System32\Drivers)
- File size, date created, version, manufacturer, and full product name for each driver
- Number of available processors
- Amount of random access memory (RAM)
- Time stamp that indicates when the Stop error occurred
- Messages and parameters that describe the Stop error
- Processor context for the process that stopped. This includes the processor, hardware state, performance counters, multiprocessor packet information, deferred procedure call information, and interrupts (that is, requests from software or devices for processor attention)
- Process information and kernel context for the halted process. This includes the offset (location) of the directory table and the database that maintains the information about every physical page (block of memory) in the operating system.

- Process information and kernel context for the thread that stopped. This information identifies registers (data-storage blocks of memory in the processor) and interrupt-request levels, and it includes pointers to data structures for operating system data.
- Kernel-mode call stack for the interrupted thread. This is a data structure that consists of a series of memory locations and one or more pointers.

Controlling Windows Error Reporting to prevent the flow of information to and from the Internet

To control the flow of information to and from the Internet when users or administrators report errors, you can configure Windows Error Reporting by using an answer file with an unattended installation or by using Group Policy. The following subsections provide more details.

Using an answer file with an unattended installation

You can control the consent level for Windows Error Reporting by using an answer file with an unattended installation. To configure a consent level of **Always ask before sending data**, confirm that your answer file includes the following line:

```
<DefaultConsent>1</DefaultConsent>
```

For more information, see **To control the consent level for Windows Error Reporting by using an answer file with an unattended installation** later in this section.

Selected Group Policy settings for Windows Error Reporting

This section provides information about a small set of the Group Policy settings that are available for Windows Error Reporting. For information about viewing these and other Group Policy settings, see **To locate Group Policy settings for configuring Windows Error Reporting** later in this section.

Setting to redirect Windows Error Reporting to a server on your intranet

This setting is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\Windows Components\Windows Error Reporting\Advanced Error Reporting Settings\Configure Corporate Windows Error Reporting**.

Two settings in the **Advanced Error Reporting Settings** refer to the "Report Queue" and the "Report Archive" These refer to information that is stored on the local computer. The Report Queue temporarily stores error reports that are waiting to be sent. The Report Archive stores reports so that the Problem Reports and Solutions interface can display them.

Setting to control the degree of prompting that occurs before data is sent

You can control the degree to which Windows Error Reporting prompts you for consent before data is sent. This setting is located in **Computer Configuration** or in **User Configuration**, under

Policies (if present), in **Administrative Templates\Windows Components\Windows Error Reporting\Consent**.

- **Configure Default consent:** If you enable this setting, you can select one of the following consent levels:
  - **Always ask before sending data**: Windows Error Reporting always prompts for consent before sending an error report.
  - **Send parameters**: Windows Error Reporting automatically sends the minimum data required to check for an existing solution, as well as data which Windows has determined (within a high probability) does not contain personally identifiable data and prompts the user for consent to send any additional data requested by Microsoft.
  - **Send parameters and safe additional data**: Windows Error Reporting sends the minimum data required to check for an existing solution in addition to data that the developer of the program has designated as being highly unlikely to contain personal information. Windows Error Reporting then prompts for consent before sending any additional data that is requested by Microsoft.
  - **Send all data**: Any data requested by Microsoft is sent without prompts. (This setting can only be configured through Group Policy, not through the Initial Configuration Tasks interface, Server Manager, or Control Panel.)

Setting to disable reporting handwriting recognition errors

You can use a Group Policy setting to specifically disable reports for handwriting recognition errors. This setting is located in **Computer Configuration** or in **User Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management\Internet Communication settings**.

- **Turn off handwriting recognition error reporting**: If you enable this setting, you cannot start the error reporting tool for handwriting recognition errors, and corrected handwriting samples will never be sent to Microsoft by Windows Error Reporting.

Setting for disabling Windows Error Reporting

This setting is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management\Internet Communication settings**.

- **Turn off Windows Error Reporting**: If you enable this setting, you can still view settings in the Initial Configuration Tasks interface, Server Manager, or Control Panel, but the display informs you that settings are being managed by a system administrator.

Important

You can also restrict Internet access for Windows Error Reporting and a number of other features by applying the **Restrict Internet communication** Group Policy setting, which is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category.

Procedures to configure Windows Error Reporting

The following procedures explain how to make changes to Windows Error Reporting using Server Manager and Group Policy.

To view or change settings for Windows Error Reporting on one or more Servers using Server Manager

1. Open Server Manager, and in the navigation pane, click **Servers**.
2. In the details pane, select one or more servers.
3. Right-click the selected servers, and then click **Configure Windows Automatic Feedback**.
4. In the **Windows Automatic Feedback** dialog box, make any necessary changes.
5. Click **OK** to apply the settings and close the dialog box.

To use Control Panel to view or change settings for Windows Error Reporting

1. Open **Control Panel**, click **Action Center**, and then click **Maintenance**.
2. Under **Check for solutions to problem reports**, click **Settings**.
3. Under **Choose when to check for solutions to problem reports**, view or change the basic error reporting settings as described earlier in this document.

   You can configure additional error reporting options as follows:

   - Click **Change report settings for all users**. These settings can be used to configure error reporting for all users of the computer or to allow each user to choose their settings (the default).
   - Click **Select programs to exclude from reporting**: This setting allows you to manage the list of programs for which Windows Error Reporting is enabled.

To locate Group Policy settings for configuring Windows Error Reporting

1. Using an account with domain administrative credentials, sign in to a computer with the Group Policy Management feature installed.
2. Open the Group Policy Management Console (GPMC) by running **gpmc.msc**, and then edit an appropriate Group Policy Object (GPO).
3. If you are interested in policy settings that apply to all users of a computer and that come into effect when the computer starts or when Group Policy is refreshed, expand **Computer Configuration**. If you are interested in policy settings that apply to specific users or administrators and that come into effect when a person signs in or when Group Policy is refreshed, expand **User Configuration**.
4. Expand **Policies** (if present), expand **Administrative Templates**, and then expand **Windows Components**.
5. Click **Windows Error Reporting**, and then view the settings that are available.
6. Click **Advanced Error Reporting Settings**, and then view the settings that are available. (What you selected in Step 3 affects what you see in **Advanced Error Reporting**

Settings. If you want to view **Configure Corporate Windows Error Reporting**, you must select **Computer Configuration** in Step 3.)

7. In the left pane, click **Consent**, and then view the settings that are available.

To disable reporting handwriting recognition errors

1. Using an account with domain administrative credentials, sign in to a computer with the Group Policy Management feature installed.
2. Open the Group Policy Management Console (GPMC) by running **gpmc.msc**, and then edit an appropriate Group Policy Object (GPO).
3. If you are interested in policy settings that apply to all users of a computer and that come into effect when the computer starts or when Group Policy is refreshed, expand **Computer Configuration**. If you are interested in policy settings that apply to specific users or administrators and that come into effect when a person signs in or when Group Policy is refreshed, expand **User Configuration**.
4. Expand **Policies** (if present), expand **Administrative Templates**, expand **System**, expand **Internet Communication Management**, and then click **Internet Communication settings**.
5. In the details pane, double-click **Turn off handwriting recognition error reporting**, and then click **Enabled**.

   Note
   You can also restrict Internet access for Windows Error Reporting and a number of other features by applying the **Restrict Internet communication** Group Policy setting, which is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management**. For more information about this Group Policy setting and the policies that it controls, see Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category.

To disable Windows Error Reporting by using Group Policy

1. Using an account with domain administrative credentials, sign in to a computer running with the Group Policy Management feature installed.
2. Open the Group Policy Management Console (GPMC) by running **gpmc.msc**, and then edit an appropriate GPO.
3. Expand **Computer Configuration**, expand **Policies** (if present), expand **Administrative Templates**, expand **Windows Components**, and then expand **Windows Error Reporting**.
4. In the details pane, double-click **Disable Windows Error Reporting**, and then click **Enabled**.

   If you enable this setting, you can still view settings in the Initial Configuration Tasks interface, Server Manager, and Control Panel, but the display informs you that settings are being managed by a system administrator.

Important

You can also restrict Internet access for Windows Error Reporting and a number of other features by applying the **Restrict Internet communication** Group Policy setting, which is located in **Computer Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management**. For more information about this Group Policy and the policies that it controls, see Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category.

To control the consent level for Windows Error Reporting by using an answer file with an unattended installation

1. Use the methods you prefer to create an answer file for an unattended installation. For detailed information about entries to include in the answer file, see Unattend.chm in the Windows Setup Automation Overview.
2. Confirm that your answer file includes one of the following lines:
   - For a consent level of **Always ask before sending data**: `<DefaultConsent>1</DefaultConsent>`
   - For a consent level of **Send parameters**: `<DefaultConsent>2</DefaultConsent>`
   - For a consent level of **Send parameters and safe additional data**: `<DefaultConsent>3</DefaultConsent>`
   - For a consent level of **Send all data**: `<DefaultConsent>4</DefaultConsent>`

   For additional information about an unattended installation, see the resources listed in Appendix A: Resources for Learning About Automated Installation and Deployment.

Additional references

For more information about Windows Error Reporting, see the following resource on the Microsoft website:

- Windows Error Reporting
- Application Recovery and Restart

Managing Privacy: Windows Store and Resulting Internet Communication

In this section

The Windows Store is Microsoft's digital distribution platform which provides access to certified desktop applications and Windows Store apps. A Windows Store app is a new type of app that uses the new Windows user interface (UI) that was introduced in Windows 8. Windows Store apps work with a variety of input sources, including touch, pen, mouse, and keyboard. Windows

Store apps can also connect to a variety of different social network services such as Facebook, while the Photos app as can aggregate photos from services such as Flickr.

Benefits and Purpose of Accessing the Windows Store

Windows Store App Feature Disclosure Requirements to Users

Windows Store Access on Windows Server

Controlling Windows Store Access Using Group Policy

Additional references

Benefits and Purpose of Accessing the Windows Store

- Designed for Discovery: The Windows Store is designed to ensure the discoverability of apps. Discoverability mechanisms such as search, category browse, rankings, and editorial curating help users find apps. Windows Store landing pages are designed to surface compelling apps and categories like new releases, top paid, top free and rising star help organize the catalog. The Windows Store catalog is indexed by search engines, so apps are easy to find.
- App Availability in a Global Marketplace: The Windows Store supports the distribution of free and paid apps in hundreds of markets worldwide, so most customers can find and install the apps they want in the language of their choice. The Windows Store supports market-specific catalogs, tailored for customers in specific locales, as well as market-specific payment providers.
- Support for Enterprise Management: IT administrators can control the method of how apps get onto user's PCs and can control access to the Windows Store by using Group Policy and AppLocker. Specific apps from the Windows Store can be allowed or blocked by using AppLocker..
- Support for Roaming User Settings: Windows Store apps can store user-specific settings so that these settings roam across multiple devices. As with operating-system settings, these user-specific app settings are available whenever the user signs in with the same Microsoft account on any device that is running Windows 8 and is connected to the cloud. After the user signs in, that device automatically downloads the settings from the cloud and applies them when the app is installed.
- Support for Enterprise Deployment: For enterprises looking to take advantage of the rich capabilities of Windows Store apps, the Windows Store offers acquisition options that provide direct control over the app deployment experience. Businesses can load their apps without having to publish their app to the Windows Store. This process, which we call “sideloading”, is available on a Windows 8 PC that is domain-joined or on a Windows 8 PC with an activated sideloading product key. This ensures that an app that is created in an enterprise can stay within the corporate network and be centrally managed, updated, and distributed.

Windows Store App Feature Disclosure Requirements to Users

Apps installed from the Windows Store are designed to take advantage of specific hardware and software features of a PC or tablet. For example, some common consumer scenarios such as a photo app that needs to access your webcam, or a restaurant guide might need to know your location in order to provide dining recommendations near you. You can view what features and functionality an app requires in the Windows Store before installing the app. Windows will ask whether you want to allow or deny access to the most sensitive of these features—location, text messaging, webcam, and microphone—before the first time each app uses them.

Application publishers to the Windows Store must disclose if an app accesses or utilizes sensitive user information such as:

- An Internet connection: Allows the app to connect to the Internet.
- Incoming connections through a firewall: Allows the app to send information to or from your PC through a firewall.
- Usage of a home or work network: Allows the app to send information between your PC and other PCs on the same network.
- App access to your pictures, videos, music, or document libraries: Allows the app to access, change, or delete files in your libraries. This includes access to any additional data embedded in these files, such as location information in photos.
- Removable storage: Allows the app to access, add, change, or delete files on an external hard drive, USB flash drive, or portable device.
- Usage of user windows credentials: Allows the app to use user credentials to authenticate and provide access to a corporate intranet.
- Certificates stored on your PC or a smart card: Allows the app to use certificates to securely connect to organizations like banks, government agencies, or your employer.
- Your location: Allows the app to determine your approximate location based on a GPS sensor or network information.
- Your PC's text messaging feature: Allows the app to send and receive text messages.
- Your PC's near-field communication feature: Allows the app to connect to other nearby devices that the same app is running on.
- Your portable devices: Allows the app to communicate with devices like your mobile phone, digital camera, or portable music player.
- Your information on a portable device: Allows the app to access, add, change, or delete contacts, calendars, tasks, notes, status, or ringtones on your portable device.
- Your mobile broadband account: Allows the app to manage your mobile broadband account.
- Your webcam and microphone: Allows the app to take pictures and record audio and video.

Note

Use of information: Apps which use these features must disclose this in their developers' privacy practices statement. If an app uses sensitive features, a link to its developer's privacy statement must be available on the "App Description page" in the Windows Store.

Windows Store Access on Windows Server

By default, the Windows Store tile is not installed on a graphical user interface (GUI) installation of Windows Server 2012 or Windows Server 2012 R2. The Windows Store tile can be added via the **Server Manager** | **Add Roles and Features Wizard** | **Features** | **User Interfaces and Infrastructure** | **Desktop Experience**. For most customers, accessing the Windows Store from a Windows server is not a common scenario and is included for development purposes only.

Controlling Windows Store Access Using Group Policy

Below are the Group Policy settings which can be used to manage Windows Store access.

- **Allow all trusted apps to install**: This policy setting allows you to manage the installation of app packages that do not originate from the Windows Store. If you enable this policy setting, you can install any trusted app package. A trusted app package is one that is signed with a certificate chain that can be successfully validated by the local computer. This can include line-of-business app packages signed by the enterprise in addition to app packages that originate from the Windows Store. If you disable or do not configure this policy setting, you can only install trusted app packages that come from the Windows Store.
- **Allow Deployment Operations In Special Profiles**: This policy setting allows you to manage the deployment operations of app packages when the user is logged in under special profiles. Deployment operation refers to adding, registering, staging, updating or removing an app package. Special profiles refer to profiles with the following types: mandatory, super-mandatory, temporary or system. Local and roaming profiles are not special profiles. When the user is logged in to a guest account, the profile type is temporary. If you enable this policy setting, the system allows deployment operations when the user is using a special profile. If you disable or do not configure this policy setting, the system blocks deployment operations when the user is using a special profile.
- **Block launching desktop apps associated with a file.**: This policy setting allows you to minimize the risk involved when an app launches the default program for a file. Because desktop programs run at a higher integrity level than apps, there is a risk that an app could compromise the system by launching a file in a desktop program. If you enable this policy setting, Windows prevents apps from launching files that would open in a desktop program. When you enable this policy setting, apps may only launch files that can be opened by another app. If you disable or do not configure this policy setting, apps could launch files that would open in a desktop program.
- **Block launching desktop apps associated with a protocol**: This policy setting allows you to minimize the risk involved when an app launches the default program for a protocol. Because desktop programs run at a higher integrity level than apps, there is a risk that a protocol launched by an app could compromise the system by launching a desktop program. If you enable this policy setting, Windows prevents apps from launching protocols that would be passed to a desktop program. When you enable this policy setting, apps may only launch protocols that can be passed to another app. If you disable or do not configure this policy setting, apps could launch protocols that would be passed to a desktop program. Enabling this policy setting will not block apps from launching http, https, and mailto protocols that

would be passed to a desktop program. The handlers for these protocols are accustomed to handling data from untrusted sources and are therefore hardened against protocol based vulnerabilities. The risk of allowing these protocols to be passed to a desktop program is minimal.

- **Turn off the Store application**: Denies or allows access to the Store application. If you enable this setting, access to the Store application is denied. If you disable or do not configure this setting, access to the Store application is allowed.

Additional references

For more information about the Windows Store see the following topics:

- How to Add and Remove Apps Using DISM
- Packaging, deployment, and query of Windows Store apps
- Managing Packaged Apps with AppLocker
- Managing Client Access to the Windows Store
- Design case study: Enterprise line of business Windows Store app
- Script to Remove Windows Store Apps In Windows 8

Managing Privacy: Using a Microsoft Account to Logon and Resulting Internet Communication

In this section

Overview of Using a Microsoft Account to Logon to Windows

Benefits and purpose of using a Microsoft Account to Login

Microsoft Account and User Information Synchronized

Safeguarding Microsoft Account Information Stored in the Cloud

Additional references

Using a Microsoft Account to login to Windows is designed to extend the capabilities of Windows by enabling cloud services that sync user Windows personalization settings such as the start page, language preferences, Windows Store apps installed, browsing history, and browser favorites.

Overview of Using a Microsoft Account to Logon to Windows

A Microsoft account—an email address and password—is a new way to sign in to any PC running Windows 8 or Windows RT or later. You might already have a Microsoft account. If you use other

Microsoft services like Messenger, Hotmail or Xbox LIVE, the email address and password you use to sign in are a Microsoft account. If you have an existing Windows Live ID, that's the same thing: "Microsoft account" is the new name for what used to be called a "Windows Live ID." When you sign in with a Microsoft account, your PC is connected to the cloud, and many of the settings, preferences, and apps associated with your user account can "follow" a user between different PCs

Signing up for a new Microsoft account for this feature to work is not a requirement. Many online services use a "string" like someone@example.com to represent a user name, even though that string looks like an email address. For example, when you order books at an online bookstore, your user name may look like an email address, even though your online book seller does not manage your email. The someone@example.com address is just a convenient way of identifying you, since most Internet users these days have email addresses. Your email account and password will still be managed by whatever email provider you choose, and the user name and password provided is used to synchronize and manage your settings and state across Windows PCs, even if you haven't signed up for Hotmail or other Microsoft services that use this ID.

During the initial Windows user setup process, users are now prompted to optionally choose to create a new Microsoft account (formerly known as a Windows Live ID) or use an existing ID for login. If you choose to create a new account, you can use any email address you want as your new ID, and then create your unique password. Local Windows account functionality has not been removed and is still an option in managed environments. In order to download apps from the Windows Store a Microsoft account is required.

Benefits and purpose of using a Microsoft Account to Login

- **Download Windows Store Apps**: You can buy and download apps from the Windows Store, and use them on multiple PCs running Windows 8 or Windows RT or later. Windows Store apps leverage the refreshing Microsoft Design principles so content is central to your application experience on Windows 8.
- **Single Sign-On**: Users can use Microsoft account credentials to sign in to devices running Windows 8. When they do this, Windows 8 works with your Windows Store app to enable authenticated experiences for them. On Windows 8, a user can associate a Microsoft account with his or her sign-in credentials for Window Store apps or websites, so that these credentials roam across any devices running Windows 8 or later. After the user signs in with that account to a device running Windows 8 or later and then runs an app or visits a website, if the corresponding stored sign-in credentials are available, Windows attempts to sign the user in automatically. When a user signs in with a Microsoft account to a device running Windows 8, any apps and services running on that device that also use Microsoft accounts for authentication can sign in with that user's Microsoft account and get data that the user has consented to share.
- **Personalized Settings Synchronization**: A user can associate his or her most commonly used operating-system settings with a Microsoft account. These settings are available

whenever the user signs in with that account on any device that is running Windows 8 and connected to the cloud. After the user signs in, that device automatically attempts to get the user's settings from the cloud and apply them to the device.

- **App Synchronization**: Windows Store apps can store user-specific settings so that these settings roam across any devices running Windows 8 or later. As with operating-system settings, these user-specific app settings are available whenever the user signs in with the same Microsoft account on any device that is running Windows 8 and is connected to the cloud. After the user signs in, that device automatically downloads the settings from the cloud and applies them when the app is installed.
- **Integrated Social Media Services**: Your friends' contact info and status automatically stay up to date from places like Hotmail, Outlook, Facebook, Twitter, and LinkedIn. You can get to and share your photos, documents, and other files from places like SkyDrive, Facebook, and Flickr.

Microsoft Account and User Information Synchronized

When you login and connect to a Windows 8 computer with a Microsoft account, users choose which settings to sync. For security purposes, all synced settings are transmitted using SSL/TLS encryption. Some of these settings won't be synced on your PC until you add your PC to your Microsoft account as a trusted PC.

- **Personalize**: Colors, background wallpaper, lock screen and account picture.
- **Desktop personalization**: Themes, taskbar, high contrast, and more.
- **Passwords**: Sign-in info for opt in apps, websites, networks, and HomeGroup.
- **Ease of Access**: Narrator, Magnifier, and more.
- **Language Preferences**: Keyboards, other input methods, display language, and more.
- **App Settings**: Certain app settings and purchases made in an app.
- **Browser**: Settings and info like history and favorites.
- **Other Windows Settings**: Windows Explorer, mouse, and more.
- **Sync Settings over Metered Connections**
- **Sync settings over metered connections even when I'm roaming**: The ability to send information across the Internet or link to a Web site can be prevented through a Group Policy setting.

Safeguarding Microsoft Account Information Stored in the Cloud

Credential information is encrypted once based on your password and then encrypted again as it is sent across the Internet. The data stored is not available to other Microsoft services or third parties.

How Microsoft Account Information is Safeguarded

1. **Strong password is required**. Blank passwords are not allowed. Credential information is encrypted once based on your password and then encrypted again as it is sent across the Internet. The data stored is not available to other Microsoft services or third parties.
2. **Secondary proof of identity is required**. Before user profile information and settings can be accessed on a second Windows 8 computer for the first time, trust must established for that PC by providing secondary proof of your identity. This further proof can be done by providing Windows with a code sent to your mobile phone number or by following the instructions sent to an alternate email address specified in your account settings.
3. **All user profile data is encrypted on the client before transmitted to the cloud.** Profile data is also protected as user data does not roam over WWAN by default. All data and settings that leave your PC are transmitted using SSL (secure socket layer) and TLS (transport layer security).

Additional references

For more information see the following resources on the Microsoft Web site:

- Group Policy Settings Reference for Windows and Windows Server
- Single Sign-on for apps and websites
- How to Switch Between Microsoft Accounts and Local Accounts
- How to Troubleshoot Microsoft Account Issues

Appendix A: Resources for Learning About Automated Installation and Deployment

In a managed environment where one of the goals might be to limit communication with the Internet (as described in other sections of this document), it is often not cost-effective to install using the standard interactive setup on each computer. To greatly lower the total cost-of-ownership and ensure configuration uniformity, you can perform an automated installation on multiple computers. By using an automated installation method, you can ensure that certain features and applications are not available on your organization's servers, or that certain features and applications are preconfigured in a way that helps prevent unwanted communication over the Internet.

The following resources provide information about automated installation:

Windows Server Installation Options

Deployment and Imaging Tools Technical Reference

Appendix B: Group Policy Settings Listed Under the Internet Communication Management Category

In this appendix

Overview of Group Policy settings listed under Internet Communication Management

Controlling multiple Group Policy settings through the Restrict Internet Communications setting

Group Policy settings that affect computer configuration

Group Policy settings that affect user configuration

Overview of Group Policy settings listed under Internet Communication Management

Windows contains a variety of Group Policy settings that can help you control the way that operating system features communicate across the Internet. This appendix describes the Group Policy settings that are presented under **Internet Communication Management**. It also describes how the **Restrict Internet communication** setting controls multiple other policy settings.

Important
The **Restrict Internet communication** policy setting controls the policy settings under **Internet Communication Management** only. In this document, when there is a procedure about a policy setting that is controlled by **Restrict Internet communication**, a note after the procedure describes the control relationship.

You can find **Internet Communication Management** in the following two locations in the Group Policy Management Console:

- In **Computer Configuration**, click **Policies** (if present), and then click **Administrative Templates\System**.
- In **User Configuration**, click **Policies** (if present), and then click **Administrative Templates\System**.

For information about using Group Policy, see Group Policy Overview.

Controlling multiple Group Policy settings through the Restrict Internet Communications setting

There are multiple ways to configure the Group Policy settings under **Internet Communication Management**. You can configure policy settings individually, which means you could configure, for example, **Turn off Event Viewer "Events.asp" links** differently from **Turn off Windows**

Error Reporting. Alternatively, the policy setting called **Restrict Internet communication** allows you to enable or disable the entire collection of policy settings at one time.

If you want to enable or disable **Restrict Internet communication**, and then create exceptions to this master policy setting by configuring individual policy settings in **Internet Communication Management**, you must use two Group Policy Objects (GPOs).

Before you begin, ensure that you understand how processing and precedence works for multiple GPOs. Choose or create a GPO with a lower precedence than another GPO. In the GPO with lower precedence, enable or disable **Restrict Internet communication**. Then, in the GPO that has precedence, apply the individual policy settings that are exceptions to the master policy setting.

If you do not use two GPOs when you set **Restrict Internet communication** and the individual policy settings that are exceptions to the master policy setting, the policy settings might not work as expected. To check the effect of multiple Group Policy settings, you can view Group Policy Results in the GPMC.

For more information, see Group Policy Planning and Deployment Guide on the TechNet website.

Group Policy settings that affect computer configuration

This subsection describes the Group policy settings that are under **Computer Configuration** in **Internet Communication Management** (Computer Configuration\Administrative Templates\System\Internet Communication Management\Internet Communications Settings). These policy settings apply to all users of an affected computer, and they come into effect when the computer starts or when Group Policy is refreshed.

All of the policy settings can also be enabled or disabled in one step by enabling or disabling the master policy setting that controls them, **Restrict Internet communication**. This policy setting is described in Controlling multiple Group Policy settings through the Restrict Internet Communications setting earlier in this section.

Note
This appendix describes only the policy settings that are available under **Internet Communication Management**. For information about all the Group Policy settings that are available, see the Group Policy Settings Reference on the Microsoft website.

Individual Group Policy settings that affect computer configuration

Note
The **Restrict Internet communication** policy setting interacts with all of the policy settings in the following list of **Computer configuration** policy settings.

More details about each policy setting are available in the Explain text for the policy setting. To view Explain text, select the policy setting in Group Policy and click the **Extended** tab, or open the policy setting and click the **Explain** tab.

- **Turn off access to all Windows Update features**: Specifies whether Windows Update can be used to update the operating system on this computer.
- **Turn off access to the Store**: Specifies whether the Store service is used to find an application to open a file with an unhandled file type or protocol association. If you enable this policy setting, **Look for an app in the Store** in the **Open With** dialog is removed. If you disable or do not configure this policy setting, the user is allowed to use the Store service and **Store** is available in the **Open With** dialog.
- **Turn off Automatic Root Certificates Update**: Specifies whether to automatically update root certificates by using the list of trusted certification authorities that Microsoft maintains on the Windows Update website. If you enable this policy setting, when a user is presented with a certificate that is issued by an untrusted root authority, the user's computer will not contact the Windows Update website.
- **Turn off downloading of print drivers over HTTP**: Specifies whether to allow this computer to download print drivers over HTTP when needed.
- **Turn off Event Viewer "Events.asp" links**: Specifies whether the Internet links that are shown within events in Event Viewer are activated. When such a link is activated and the user clicks it, information that identifies the event is sent to a Microsoft website so that explanatory text, if available, can be sent back to the user.
- **Turn off handwriting personalization data sharing**: Turns off data sharing from the handwriting recognition personalization tool. The handwriting recognition personalization tool enables Tablet PC users to adapt handwriting recognition to their own writing style by providing writing samples. The tool can optionally share user writing samples with Microsoft to improve handwriting recognition in future versions of Windows. The tool generates reports and transmits them to Microsoft over a secure connection.

  If you enable this policy, Tablet PC users cannot choose to share writing samples from the handwriting recognition personalization tool with Microsoft. If you disable this policy, Tablet PC user writing samples from the handwriting recognition personalization tool will automatically be shared with Microsoft. If you do not configure this policy, Tablet PC users can choose if they want to share their writing samples from the handwriting recognition personalization tool with Microsoft.
- **Turn off handwriting recognition error reporting**: Specifies whether users can report errors that they encounter in a Tablet PC Input Panel. This policy setting is related to the policy setting **Turn off Windows Error Reporting**. If you turn off Windows Error Reporting, you are also turning off error reporting for handwriting recognition.

  This policy setting is also described in Manage Privacy: Windows Error Reporting and Resulting Internet Communication in this document.
- **Turn off Help and Support Center "Did you know?" content**: This policy setting is deprecated. See **Turn off Windows Online** in **Group Policy settings that affect user configuration** later in this section.

- **Turn off Help and Support Center Microsoft Knowledge Base search**: This policy setting is deprecated. See **Turn off Windows Online** in **Group Policy settings that affect user configuration** later in this section.
- **Turn off Internet Connection Wizard if URL connection is referring to Microsoft.com**: This policy setting is deprecated.
- **Turn off Internet download for web publishing and online ordering wizards**: Specifies whether Windows should download a list of providers for the Order Prints Wizard. By default, Windows displays providers that are downloaded from a Microsoft website in addition to providers that are specified in the registry.

  If you enable this policy setting, Windows will not download providers, and only the service providers that are stored in the local registry are displayed. When Windows 8 is installed, but the Order Prints Wizard has not been used, no service providers are stored in the local registry. If this Group Policy setting is applied at that time, the wizard will not display links to service providers.

  This policy does not affect Windows Server 2012
- **Turn off Internet File Association service**: Specifies whether to use the web-based File Association service or to use only locally stored information about file name extensions and file types. It also specifies the applications or features to use when opening a particular file type. The file association web service is used only when a user tries to open a file and there is no locally stored information about the file name extension.
- **Turn off printing over HTTP**: Specifies whether to allow printing over HTTP from this computer. This policy setting does not control whether the computer can act as an Internet print server.
- **Turn off Registration if URL connection is referring to Microsoft.com**: This policy is deprecated.
- **Turn off Search Companion content file updates**: This policy setting is deprecated.
- **Turn off the "Order Prints" picture task**: Specifies whether the Order Prints Wizard can be run from Windows Photo Gallery.

  This policy does not affect Windows Server 2012.
- **Turn off the "Publish to web" task for files and folders**: This policy setting is deprecated.
- **Turn off the Windows Messenger Customer Experience Improvement Program**: This policy setting is deprecated.
- **Turn off Windows Customer Experience Improvement Program**: Specifies whether to opt out users from the Windows Customer Experience Improvement Program. If you enable this policy setting, all users are opted out of Windows Customer Experience Improvement Program.

  This policy setting is also described in Manage Privacy: Windows Customer Experience Improvement Program and Resulting Internet Communication.
- **Turn off Windows Error Reporting**: Specifies whether error reports from a system or application that has stopped responding are sent to Microsoft. Error reports are used to improve the quality of the product. This policy setting overrides any user setting that is made from the Control Panel for error reporting.

This policy setting and other ways of controlling error reporting through Group Policy are described in Manage Privacy: Windows Error Reporting and Resulting Internet Communication.

- **Turn off Windows Network Connectivity Status Indicator active tests**: Prevents Network Connectivity Status Indicator (NCSI) from performing a network connectivity test that involves attempting to make a connection across the Internet.
- **Turn off Windows Update device driver searching**: Specifies whether Windows searches Windows Update for device drivers when no local drivers for a device are present.

Group Policy settings that affect user configuration

This subsection describes the policy settings under **User Configuration** in **Internet Communication Management** (User Configuration\Administrative Templates\System\Internet Communication Management\Internet Communications Settings). These policy settings apply to the individual user, and they come into effect when the user signs in or when Group Policy is refreshed.

These policy settings are located in **User Configuration** under **Policies** (if present), in **Administrative Templates\System\Internet Communication Management\Internet Communication settings**.

All of the policy settings can be enabled or disabled in one step by enabling or disabling the master policy setting that controls them, **Restrict Internet communication**. This policy setting is described in Controlling multiple Group Policy settings through the Restrict Internet Communications setting earlier in this section.

Note

This appendix describes only the policy settings that are available under **Internet Communication Management**. For information about all the Group Policy settings that are available, see the Group Policy Settings Reference on the Microsoft website.

Group Policy settings that affect user configuration

The **Restrict Internet communication** policy setting interacts with all of the policy settings in the following list of user configuration policy settings under **Internet Communication Management**. You can also select the policy setting in Group Policy and click the **Extended** tab, or open the policy setting and click the **Explain** tab.

- **Turn off access to the Store**: Specifies whether the Store service is used to find an application to open a file that has an unhandled file type or protocol association. When a user opens a file type or protocol that is not associated with any applications on the computer, the user is given the choice to select a local application or use the Store service to find an application. If you enable this policy setting, the **Look for an app in the Store** item in **the Open With** dialog is removed. If you disable or do not configure this policy setting, the user is allowed to use the Store service and the **Store** item is available in the **Open With** dialog.

- **Turn off downloading of print drivers over HTTP**: Specifies whether to allow this computer to download print drivers over HTTP when needed.
- **Turn off handwriting personalization data sharing**: Turns off data sharing from the handwriting recognition personalization tool. The handwriting recognition personalization tool enables Tablet PC users to adapt handwriting recognition to their own writing style by providing writing samples. The tool can optionally share user writing samples with Microsoft to improve handwriting recognition in future versions of Windows. The tool generates reports and transmits them to Microsoft over a secure connection.

  If you enable this policy, Tablet PC users cannot choose to share writing samples from the handwriting recognition personalization tool with Microsoft. If you disable this policy, Tablet PC user writing samples from the handwriting recognition personalization tool will automatically be shared with Microsoft. If you do not configure this policy, Tablet PC users can choose whether they want to share their writing samples from the handwriting recognition personalization tool with Microsoft.
- **Turn off handwriting recognition error reporting**: Specifies whether users can report errors that they encounter in the Tablet PC Input Panel. This policy setting is related to **Turn off Windows Error Reporting**, which is described in Individual Group Policy settings that affect computer configuration. If you turn off Windows Error Reporting, you are also turning off error reporting for handwriting recognition.

  This policy setting is also described in **Windows Error Reporting and the Problem Reports and Solutions Feature in Windows 8 and Windows Server 2012**.
- **Turn off Help Experience Improvement Program**: Specifies whether users can participate in the Help Experience Improvement program. The Help Experience Improvement program collects information about how customers use Windows Help so that Microsoft can improve it. If this setting is enabled, this policy prevents users from participating in the Help Experience Improvement program. If this setting is disabled or not configured, users will be able to turn on the Help Experience Improvement program feature from the Help and Support settings page.
- **Turn off Help Ratings**: Specifies whether, when Online Help is turned on, a user can enter feedback into a form at the bottom of a Help topic, and then send that feedback to Microsoft.
- **Turn off Internet download for web publishing and online ordering wizards**: Specifies whether Windows should download a list of providers for the Order Prints Wizard. By default, Windows displays providers that are downloaded from a Microsoft website in addition to providers that are specified in the registry.

  If you enable this policy setting, Windows will not download providers, and only the service providers that are stored in the local registry are displayed. When Windows is installed, but the Order Prints Wizard has not been used, no service providers are stored in the local registry. If this Group Policy setting is applied at that time, the wizard will not display links to service providers.
- **Turn off Internet File Association service**: Specifies whether to use the web-based File Association service or to use only locally stored information about file name extensions and file types. It also specifies the applications or features to use when opening a particular file type. The File Association service is used only when a user tries to open a file and there is no locally stored information about the file name extension.

- **Turn off printing over HTTP**: Specifies whether to allow printing over HTTP for this user. This policy setting does not control whether the computer can act as an Internet print server.
- **Turn off the "Order Prints" picture task**: Specifies whether the Order Prints Wizard can be run from Windows Photo Gallery.

  This setting was deprecated in Windows Server 2012.
- **Turn off the “Publish to Web” task for files and folders**: Specifies whether the tasks "Publish this file to the Web," "Publish this folder to the Web," and "Publish the selected items to the Web," are available from **File and Folder Tasks** in Windows folders. The Web Publishing Wizard is used to download a list of providers, and it allows users to publish content to the web. If you enable this setting, these tasks are removed from **File and Folder Tasks** in Windows folders. If you disable or do not configure this setting, the tasks will be shown.
- **Turn off Windows Online**: Specifies whether users can see updated Help topics that Microsoft makes available across the Internet. If you turn off Windows Online, you also turn off Help Ratings and the Windows Customer Experience Improvement Program (which are dependent on the Windows Online policy setting).
- **Turn off the Windows Messenger Customer Experience Improvement Program**: This policy setting specifies whether Windows Messenger collects anonymous information about how Windows Messenger software and service is used. If you enable this policy setting, Windows Messenger does not collect usage information, and the user settings to enable the collection of usage information are not shown. If you disable this policy setting, Windows Messenger collects anonymous usage information, and the setting is not shown. If you do not configure this policy setting, users have the choice to opt in and allow information to be collected.

Support Windows Server 2012 R2 and Windows Server 2012

This section contains information to help IT pros find workarounds for known issues in Windows Server 2012, and troubleshoot and resolve specific Windows Server 2012 system errors and events.

Best Practices Analyzer

Best Practices Analyzer

In Windows management, best practices are guidelines that are considered the ideal way, under normal circumstances, to configure a server as defined by experts. Though best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Windows Server 2012 Understand and Troubleshoot Guides

The Windows Server 2012 Understand and Troubleshoot Guides (UTG) help IT administrators and architects develop awareness of key technical concepts, functionality, and troubleshooting techniques. This understanding enables a successful early adoption experience during the product evaluation phase. UTGs cover:

- Technical overview.
- Server role or feature installation, configuration, and management tasks.
- Component architecture and interaction.
- Methodology for troubleshooting.

Note
The UTGs do not provide deployment planning content.

These Understand and Troubleshoot Guides were written for Windows Server “8” Beta, but also apply to Windows Server 2012. The following UTGs are available:

- Understand and Troubleshoot Hyper-V Replica in Windows Server 2012
- Understand and Troubleshoot Printing in Windows Server 2012
- Understand and Troubleshoot Remote Desktop Services Desktop Virtualization in Windows Server 2012
- Understand and Troubleshoot Remote Desktop Services in Windows Server 2012
- Understand and Troubleshoot Storage Spaces in Windows Server 2012

Best Practices Analyzer

In Windows management, *best practices* are guidelines that are considered the ideal way, under normal circumstances, to configure a server as defined by experts. For example, it is considered a best practice for most server technologies to keep open only those ports required for the technologies to communicate with other networked computers, and block unused ports. While best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

What is Best Practices Analyzer?

Best Practices Analyzer (BPA) is a server management tool that is available in Windows Server® 2008 R2, Windows Server® 2012, and Windows Server® 2012 R2. BPA can help administrators reduce best practice violations by scanning one or more roles that are installed on those three Windows Server operating systems, and reporting best practice violations to the administrator. Administrators can filter or exclude results from BPA reports that they don’t

need to see. Administrators can also perform BPA tasks by using either the Server Manager GUI, or Windows PowerShell cmdlets.

For more information about Best Practices Analyzer and scans, see Run Best Practices Analyzer Scans and Manage Scan Results.

About content in this section

Topics in this section can help you bring server roles that are running on Windows Server 2012 R2 and Windows Server 2012 into compliance with best practices. Content in this section is most valuable to administrators who have completed a Best Practices Analyzer scan of one or more roles, and who want information about how to interpret and resolve scan results that identify areas of those roles that are noncompliant with best practices.

Best Practices Analyzer for Remote Access

In Windows management, best practices are guidelines that are considered the ideal way, under normal circumstances, to configure a server as defined by experts. While best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Topics in this section can help you bring Remote Access running on Windows Server® 2012 R2 into compliance with best practices. Content in this section is most valuable when you are preparing to run or have completed a Best Practices Analyzer scan of Remote Access, particularly when you want information about how to interpret and resolve scan results that identify areas of Remote Access that are noncompliant with best practices.

More information about Remote Access

The Remote Access server role includes both DirectAccess and Routing and Remote Access Service (RRAS). Remote Access allows for centralized administration, configuration, and monitoring of both DirectAccess and VPN-based remote access services.

This section contains the following topics.

- Best Practices Analyzer for Remote Access: Prerequisites
- Best Practices Analyzer for Remote Access: Configuration (Section 1)
- Best Practices Analyzer for Remote Access: Configuration (Section 2)
- Best Practices Analyzer for Remote Access: Configuration (Section 3)
- Best Practices Analyzer for Remote Access: Performance and Operation

Best Practices Analyzer for Remote Access: Prerequisites

The topics in this section explain how to prepare Remote Access for running a Best Practices Analyzer scan for the first time. Content in this section is most valuable if you have recently installed Remote Access on Windows Server 2012 R2 and want information about preparations that must be completed before you can run a Best Practices Analyzer scan.

Best Practices Analyzer and prerequisite rules

The Best Practices Analyzer applies prerequisite rules to identify gaps in Remote Access configuration that you can address before you run a Best Practices Analyzer scan.

Topics in this section

This section contains the following topics.

- RRAS: The Remote Access server role should be configured in Multitenant mode
- RRAS: The Remote Access gateway should be configured with Multitenancy support
- RRAS: All Routing Domains should be enabled
- RRAS: All enabled Routing Domains should be available

RRAS: The Remote Access server role should be configured in Multitenant mode

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Prerequisite

Issue

The gateway is configured for Multitenancy, but the Remote Access role is not installed in Multitenant mode.

Impact

The Remote Access gateway cannot act as a multitenant Site-to-Site VPN terminating entity that services multiple customers.

Resolution

Configure the Remote Access role in Multitenant mode.

Before running the Best Practices Analyzer for Remote Access in Multitenant mode, the Remote Access server role must be installed, enabled, and configured in Multitenancy mode. Following are the prerequisites for running Remote Access in Multitenant mode.

- The virtual machine (VM) must be running Windows Server 2012 R2 or later.
- The user that is configuring the Remote Access Service must be a member of the Administrators group on the local computer or VM.
- The computer must have at least two network adapters installed.
- At least one network adapter must be connected to a network.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install Remote Access in Multitenant mode

1. Run Windows PowerShell with Administrative privileges.
2. Type the following commands to install and configure RRAS in Multitenant mode.

```
Add-WindowsFeature -Name RemoteAccess -IncludeAllSubFeature –
IncludeManagementTools

ipmo RemoteAccess

Install-RemoteAccess -Multitenancy
```

RRAS: The Remote Access gateway should be configured with Multitenancy support

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote

Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Prerequisite

Issue

The gateway is not configured for Multitenancy, but the Remote Access role is installed with Multitenancy support.

Impact

The Remote Access gateway cannot be configured until either RRAS is installed in single tenant mode or stack multitenancy is enabled.

Resolution

Determine whether you want to deploy a remote access server with or without multitenant support. After this determination, either configure the host computer with multitenancy support or configure Remote Access in single tenant mode.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

Note

Select and follow one of the procedures below based on your decision to either support multitenancy or to deploy a Remote Access server that does not support multitenancy.

To configure the host computer for Multitenant support

1. Read and follow the steps in the networking blog Multi-tenant Site-to-Site (S2S) VPN Gateway with Windows Server 2012 R2

To configure Remote Access as a VPN or site-to-site VPN server without multitenant support

1. Open Windows PowerShell with Administrator privileges, and then run the following command to uninstall Remote Access in Multitenant mode.

   ```
   Uninstall-RemoteAccess
   ```

2. Run the following Windows PowerShell command to install Remote Access either as a

VPN server (with the parameter **–VPNType** value of **VPN**) or as a site-to-site VPN server (with the parameter **–VPNType** value of **VPNS2S**).

```
Install-RemoteAccess  -VpnType <VPN / VPNS2S>
```

3. Perform the procedures in the topic Enable RRAS as a VPN Server to configure your Remote Access server.

RRAS: All Routing Domains should be enabled

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Information
Category	Prerequisite

Issue

One or more Routing Domains are not in an enabled state.

Impact

The RRAS server cannot accept Site-to-Site VPN connections for Routing Domains that are not enabled for RRAS.

Resolution

Enable the Routing Domains by using the corresponding PowerShell command.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To enable Routing Domains

1. Run Windows PowerShell with Administrative privileges.

2. Type the following command to enable the Routing Domains. Use the **–Name** parameter to identify the target Remote Access Routing Domain to enable. The **–Type** parameter indicates the type of server either as a VPN server using a value of **VPN**, as a site-to-site VPN server using a value of **VPNS2S** or with a value of **All**. **All** enables all VPN, Site-to-Site VPN and routing services for the target Remote Access Routing Domain.

```
Enable-RemoteAccessRoutingDomain –Name <string> -Type <VPN / VPNS2S / All>
```

RRAS: All enabled Routing Domains should be available

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Prerequisite

Issue

Compartments for one or more enabled Routing Domains were deleted from the stack.

Impact

Remote Access will not work for this Routing Domain, and this causes inconsistent behavior.

Resolution

Disable the Routing Domains by using the corresponding Windows PowerShell command.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To disable a Routing Domain

1. Run Windows PowerShell with Administrative privileges.

2. Type the following command to disable a Routing Domain. Use the **–Name** parameter to identify the target Remote Access Routing Domain to disable. The **–Type** parameter indicates the type of server with a value of **All**. **All** enables all VPN, Site-to-Site VPN and routing services for the target Remote Access Routing Domain.

```
Disable-RemoteAccessRoutingDomain -Name -Type All
```

Best Practices Analyzer for Remote Access: Configuration (Section 1)

The topics in this section can help you bring Remote Access running on Windows Server 2012 R2 into compliance with configuration best practices. Content in this section is most valuable when you have completed a Best Practices Analyzer scan of Remote Access and you want information about how to interpret and resolve scan results that identify areas of Remote Access that are noncompliant with configuration best practices.

Best Practices Analyzer and configuration rules

The Best Practices Analyzer applies configuration rules to identify settings that might require modification for Remote Access to perform optimally. Configuration rules can help prevent setting conflicts that can result in error messages or prevent Remote Access from carrying out its prescribed duties in an enterprise.

Topics in this section

This section includes the following topics.

- RRAS: The inbound Certification Authority (CA) should be configured
- RRAS: A valid CA certificate for the Remote Access Server certificate must be present in the TRCA certificate store
- RRAS: A valid CA certificate for the Site-to-Site VPN interface certificate must be present in the TRCA certificate store
- RRAS: The server certificate expires within 7 days
- RRAS: The CA certificate of the Remote Access Server certificate in the TRCA certificate store expires within 7 days
- RRAS: The certificate for the Site-to-Site VPN interface expires within 7 days
- RRAS: The CA certificate for the Site-to-Site VPN interface certificate expires within 7 days
- RRAS: The CA certificate for the destination server of Site-to-Site VPN interface expires within 7 days
- RRAS: The Remote Access server certificate must have a public IP address for Alternate Subject Name

- RRAS: The Site-to-Site VPN interface name must match the Username
- RRAS: No two Site-to-Site VPN interfaces with PSK based authentication should have the same destination
- RRAS: For PSK authentication, the destination cannot be configured as a Fully Qualified Domain Name (FQDN)
- RRAS: The Site-to-Site VPN interface should be configured with a Source IP address
- RRAS: Custom policies configured for the Site-to-Site VPN interface should be a subset of Remote Access server global policies
- RRAS: The number of ports available for the Routing Domain should not be less than the number of VPN and Site-to-Site VPN interfaces

RRAS: The inbound Certification Authority (CA) should be configured

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

No inbound Certification Authority (CA) is configured for the Remote Access server.

Impact

Authentication for the Ikev2 Machine certificate based Site-to-Site VPN interface will fail.

Resolution

Configure an inbound CA for the Site-to-Site VPN interface computer's certificate authentication.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To configure an inbound CA for the Remote Access server

1. Run Windows PowerShell with Administrative privileges.
2. Configure the inbound CA for Site-to-Site VPN interface's computer certificate authentication as follows.

```
Set-VpnAuthProtocol –UserAuthProtocolAccepted "Certificate" –
RootCertificateNameToAccept <root certificate with required
CA>
```

RRAS: At least one valid IKEv2 certificate should be present on the RRAS server

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

A valid IKEv2 Machine Certificate is not installed for the Site-to-Site VPN interface in the Remote Access server's certificate store.

Impact

The IKEv2 Machine Certificate authentication for the Site-to-Site VPN server interface doesn't work without a valid certificate in the Remote Access server's certificate store.

Resolution

Install a valid IKEv2 Machine Certificate for the Site-to-Site VPN server interface in the Remote Access server's certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a IKEv2 Machine Certificate for the Site-to-Site VPN interface

1. Install a valid IKEv2 Machine Certificate for the Site-to-Site VPN interface by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate

RRAS: A valid CA certificate for the Remote Access Server certificate must be present in the TRCA certificate store

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

A valid Certification Authority (CA) certificate for the CA that issued the Remote Access server Certificate is not installed in the Trusted Root Certification Authorities (TRCA) store on the Remote Access server.

Impact

The IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface doesn't work without the valid root CA certificate in the Remote Access server's certificate store.

Resolution

Install a valid CA certificate from the CA that issued the Remote Access server's certificate in the TRCA store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid CA certificate on the Remote Access server

1. Install a valid root CA certificate for the Remote Access server by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate

RRAS: A valid CA certificate for the Site-to-Site VPN interface certificate must be present in the TRCA certificate store

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

A valid Certification Authority (CA) certificate for the CA that issued the Site-to-Site VPN interface certificate is not installed in the Trusted Root Certification Authorities (TRCA) store on the Remote Access server.

Impact

The IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface does not work without a valid CA certificate in the TRCA store on the Remote Access server.

Resolution

Install a valid CA certificate for the Site-to-Site VPN interface in the Remote Access server's certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid root certificate for the Site-to-Site VPN interface

1. Install a valid root certificate for the Site-to-Site VPN interface by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate

RRAS: The server certificate expires within 7 days

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Remote Access server certificate is about to expire.

Impact

IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface does not work without a valid server certificate in the Remote Access server's certificate store.

Resolution

Install a valid server certificate for the Remote Access server in the certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid certificate for the Remote Access server

1. Install a valid root CA certificate from the CA that issued the Remote Access server certificate. To do this, read and follow the steps in the TechNet article Installing a root certificate
2. On the Remote Access server, run Windows PowerShell with Administrative privileges.
3. Type the following command to install a valid server certificate. Use the **–CertificateAdvertised** parameter to set the IKEv2 Machine Certificate type, such as **–X509Certificate2**.

```
Set-VpAuthProtocol –CertificateAdvertised <X509Certificate2>
```

RRAS: The CA certificate of the Remote Access Server certificate in the TRCA certificate store expires within 7 days

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The root CA certificate for the CA that issued the Remote Access server's Server Certificate is about to expire.

Impact

The IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface does not work without a valid root CA certificate in the Trusted Root Certification Authorities (TRCA) store on the Remote Access server.

Resolution

Import a valid Certification Authority (CA) certificate to the Remote Access server certificate store at the location **Certificates (Local Computer)** */* **Trusted Root Certification Authorities** */* **Certificates**.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To import a valid CA certificate to the TRCA certificate store

1. Import a valid root Certification Authority (CA) certificate by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate

RRAS: The certificate for the Site-to-Site VPN interface expires within 7 days

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Remote Access server's Site-to-Site VPN interface certificate is about to expire.

Impact

IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface does not work without a valid Server Certificate in the Remote Access server's certificate store.

Resolution

Install a valid Site-to-Site VPN interface certificate in the Remote Access server's certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid Machine Certificate for the Remote Access server

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to install a valid certificate. Use the **–MachineCertificate** parameter to set the IKEv2 Machine Certificate type, such as **–X509Certificate2**.

```
Set-VpnS2Snterface –MachineCertificate <-X509Certificate2>
```

RRAS: The CA certificate for the Site-to-Site VPN interface certificate expires within 7 days

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The root CA certificate for the CA that issued the Site-to-Site VPN interface's certificate is about to expire.

Impact

IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface does not work without a valid CA certificate in the TRCA certificate store.

Resolution

Install a valid CA certificate in the Remote Access server's TRCA certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid CA certificate for the Site-to-Site VPN interface

1. Import a valid root certificate for the Remote Access server by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate

RRAS: The CA certificate for the destination server of Site-to-Site VPN interface expires within 7 days

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The destination's Certification Authority (CA) certificate is about to expire for the Site-to-Site VPN interface.

Impact

IKEv2 Machine Certificate authentication for the Site-to-Site VPN interface's remote destination does not work without a valid CA certificate in the Remote Access server's certificate store.

Resolution

Install a valid CA certificate for the Site-to-Site VPN interface destination in the Remote Access server's certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid CA certificate for the Site-to-Site VPN interface's remote destination

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to install a valid certificate. Use the **–RootCertificateNameToAccept** parameter to set the IKEv2 Machine Certificate type, such as **–X509Certificate2**.

```
Set-VpnAuthProtocol -RootCertificateNameToAccept
<X509Certificate2>
```

RRAS: The Remote Access server certificate must have a public IP address for Alternate Subject Name

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

There is no valid public IP address configured in the Remote Access server certificate's Subject Name or Alternate Subject Name.

Impact

A remote client cannot connect to the Remote Access server using the IP address.

Resolution

Obtain a valid certificate for the Remote Access server from the Certification Authority (CA). Ensure that the Alternate Subject Name property of the certificate is configured with the Remote Access server's public IP address, and install the certificate in the Remote Access server's certificate store.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To install a valid certificate for the Site-to-Site VPN interface

1. Get a new IKEv2 Machine Certificate from a valid CA with the given subject and alternate subject, and install it in Remote Access Server's Personal certificate store by using the Microsoft Management Console (MMC). Read and follow the steps in the TechNet article Installing a root certificate
2. Run Windows PowerShell with Administrative privileges.
3. Type the following command to install a valid certificate. Use the **–MachineCertificate** parameter to set the IKEv2 Machine Certificate type, such as **–X509Certificate2**.

```
Set-VpnS2Snterface -MachineCertificate <-X509Certificate>
```

RRAS: The Site-to-Site VPN interface name must match the Username

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access

Severity	Warning
Category	Configuration

Issue

The Site-to-Site VPN interface is configured for Extensible Authentication Protocol (EAP) authentication, and the interface name does not match the EAP username.

Impact

The Site-to-Site VPN interface cannot to establish a connection with the remote server.

Resolution

Either add a new Site-to-Site VPN interface with the required name, or add the required EAP User with the same name as the interface.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a new valid Site to Site VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to add the appropriate Site-to-Site VPN interface with required name. Use the **–Name** parameter to identify the Site-to-Site VPN interface name. In the **–Destination** command, type the Destination IP Address for the Site-to-Site interface.

```
Add-VpnS2SInterface –Name <new Interface name> -Destination
<S2S VPN Interface destination IP Address>
```

To add the required EAP User (with the same name as interface)

1. Run Windows PowerShell with Administrative privileges.
2. In an Administrator command prompt, type **lusrmgr.msc** to open the Local Users and Groups console.
3. Right-click **Users** and select **New User**.
4. Type the username (same as Site-to-Site VPN interface), password and other necessary information, and click **Create**.

RRAS: No two Site-to-Site VPN interfaces with PSK based authentication should have the same destination

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

Two or more Site-to-Site VPN interfaces have PSK-based authentication and the same destination.

Impact

Site-to-Site VPN interfaces cannot connect to the same destination, or the Site-to-Site VPN interface cannot serve the incoming connection request from the remote server.

Resolution

Change the authentication method of the Site-to-Site VPN interface, or change the destination IP address of the Site-to-Site VPN interface.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change the Authentication Method of the Site-to-Site VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to change the authentication method. Use the **–Name** parameter to identify the Site-to-Site VPN interface name. In the EAP **–AuthenticationMethod** parameter, type **EAP** for the Extensible Authentication Protocol. The **–EAPMethod** parameter can be defined either as Extensible Authentication Protocol-Transport Layer Security using a value of **EAP-TLS**, as the EAP-Protected

Extensible Authentication Protocol using a value of **EAP-PEAP** or with a value of **EAP-MSCHAPv2**.

```
Set-VpnS2SInterface –Name <VPN S2S Interface name> -
AuthenticationMethod EAP –EapMethod <specify the EAP method>
```

3. Type the following command and use the **–Name** parameter to identify the Site-to-Site VPN interface name. For **–AuthenticationMethod** parameter, type **MachineCertificates**. Use the **–Certificate** parameter to set the certificate type, such as **–X509Certificate2**.

```
Set-VpnS2SInterface –Name <VPN S2S Interface name> -
AuthenticationMethod MachineCertifictes –Certificate
<Certificate object>
```

To change destination IP address of the Site-to-Site VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to change the destination IP address. Use the **–Name** parameter to identify the Site-to-Site VPN interface name. For the **–Destination** parameter, type the new destination IP address.

```
Set-VpnS2SiteInterface –Name <VPN S2S Interface name> -
Destination <Destination IP Address>
```

RRAS: For PSK authentication, the destination cannot be configured as a Fully Qualified Domain Name (FQDN)

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

The destination for Site-to-Site VPN interfaces is configured as a Fully Qualified Domain Name (FQDN).

Impact

The Site-to-Site VPN interface cannot connect to the specified destination.

Resolution

Change the destination for the Site-to-Site VPN interface.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change the destination for the Site-to-Site VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to change the destination IP address. Use the **–Name** parameter to identify the Site-to-Site VPN interface name. For the **–Destination** parameter, type the new destination IP address.

```
Set-VpnS2SInterface –Name <VPN S2S Interface name> -
Destination <Destination IP Address>
```

RRAS: The Site-to-Site VPN interface should be configured with a Source IP address

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Remote Access server has multiple public IP addresses and a source IP address is not configured for the Site-to-Site VPN interface.

Impact

The Site-to-Site VPN interface might not be able to connect to the specified destination.

Resolution

Configure the source IP address of the Site-to-Site VPN interface.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To configure the source IP address of the Site-to-Site VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Type the following command to change the source IP address. Use the **–Name** parameter to identify the Site-to-Site VPN interface. For the **–SourceIPAddress**, type the source IP address of the Site-to-Site VPN interface.

```
Set-VpnS2SInterface –Name <VPN S2S Interface name> -
SourceIPAddress <Source IP Address>
```

RRAS: Custom policies configured for the Site-to-Site VPN interface should be a subset of Remote Access server global policies

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

Custom policies configured for the Site-to-Site VPN interface of the Routing Domain are not a part of the Routing Domain specific or Global (Server level) set of policies. The custom policy is not in the set of policies configured at the Routing Domain level.

Impact

Interoperability with third party VPN devices are affected. The negotiated values might be different for incoming and outgoing connections.

Resolution

Change the set of policies to which the policy belongs.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To set the custom policy at the Routing Domain or Server level

1. Depending upon your deployment, perform one of the following actions.
   - Set the policy as a custom policy in the set of policies at either the Routing Domain level or the Server level.
   - Set the policy as one of the policies from the predefined set of policies defined at the Server level.

Best Practices Analyzer for Remote Access: Configuration (Section 2)

The topics in this section can help you bring Remote Access running on Windows Server 2012 R2 into compliance with configuration best practices. Content in this section is most valuable when you have completed a Best Practices Analyzer scan of Remote Access and you want information about how to interpret and resolve scan results that identify areas of Remote Access that are noncompliant with configuration best practices.

Best Practices Analyzer and configuration rules

The Best Practices Analyzer applies configuration rules to identify settings that might require modification for Remote Access to perform optimally. Configuration rules can help prevent setting conflicts that can result in error messages or prevent Remote Access from carrying out its prescribed duties in an enterprise.

Topics in this section

This section includes the following topics.

- RRAS: A static pool should be configured for IPv4 address assignment to the VPN client
- RRAS: The static pool IPv4 addresses must be valid unicast IPv4 addresses
- RRAS: The VPN Tenant Name must be specified
- RRAS: The VPN Tenant Name for a Routing Domain must not be a subset of another Routing Domain's Tenant Name
- RRAS: The default route (IPv4 or IPv6) should not be advertised to the peers
- RRAS: The default route (IPv4 or IPv6) should not be accepted from the peers
- RRAS: The BGP peer IP address should not be assigned to a local network interface
- RRAS: A local global IPv6 address must be configured on the BGP Router
- RRAS: Multiple routes with different Next-Hop values and the same Destination prefix are configured
- RRAS: The BGP peer's Hold-Timer should not be set to the value 0
- RRAS: BGP peer's Hold-Timer should not be set to a very low value
- RRAS: All the ingress route advertisements should not be dropped because of a routing policy
- RRAS: All the egress route advertisements should not be dropped because of a routing policy
- RRAS: BGP peers should not be configured for manual (passive) peering mode
- RRAS: For BGP Peering over IPv6 addresses, IPv4 routes should not be configured for advertisement

RRAS: The number of ports available for the Routing Domain should not be less than the number of VPN and Site-to-Site VPN interfaces

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning

Category	Configuration

Issue

The total number of ports (SSTP and IKEv2) is less than the number of Site-to-Site VPNs and VPN interfaces configured on the computer.

Impact

VPN clients or Site-to-Site VPN interfaces might not be able to establish a connection if all of the ports are in use.

Resolution

Ensure that the number of available ports is equal to, or greater than, the allowed VPN interfaces and the Site-to-Site VPN interfaces that are configured.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To determine port and interface availability

1. Run Windows PowerShell with Administrative privileges.
2. Check the number of available IKEv2 and SSTP ports using the following command.

   ```
   Get-VpnServerIPsecConfiguration
   ```

3. Run the following command to compare the number of available ports against the total number of configured Site-to-Site VPN interfaces.

   ```
   Get-VpnS2SInterface
   ```

4. And add to it the sum of the maximum allowed number of VPN connections by adding the configured IP Address pools.

RRAS: A static pool should be configured for IPv4 address assignment to the VPN client

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

A static IPv4 address pool is not configured for the VPN interface for the Routing Domain.

Impact

VPN remote clients will not be able to access organization resources.

Resolution

Configure the IP address pool for assigning static addresses for the Routing Domain.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To configure the IP address pool for the Routing Domain

1. Run Windows PowerShell with Administrative privileges.
2. Configure the IP address pool for assigning static addresses by using one of the following commands.
   - To configure an existing IP address pool, use the following command.

     ```
     Set-RemoteAccesRoutingDomain –Name <Routing Domain Name> -
     IPAddressRange <Start IP address, End IP address>
     ```

   - To add a new IP address pool, use the following command.

     ```
     Add-VpnIPAddressRange –RoutingDomain <Routing Domain name> -
     IPAddressRange <Start IP address, End IP address>
     ```

RRAS: The static pool IPv4 addresses must be valid unicast IPv4 addresses

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

The static address pool configured for the VPN interface for the Routing Domain does not contain valid unicast IPv4 addresses.

Impact

Remote VPN clients will not be able to access organization resources.

Resolution

Configure the IP address pool for assigning static addresses for the Routing Domain.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To configure the IP address pool for the Routing Domain

1. Run Windows PowerShell with Administrative privileges.
2. Configure the IP address pool for the Routing Domain by using the following command.

```
Set-RemoteAccesRoutingDomain –Name <Routing Domain Name> -
IPAddressRange <Start IP address, end IP address>
```

RRAS: The VPN Tenant Name must be specified

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

No VPN tenant name is configured for the Routing Domain.

Impact

VPN clients will not be able to connect to the corresponding tenant's Remote Access server.

Resolution

Add a tenant name for the VPN interface for the Routing Domain.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a tenant name for the VPN interface for the Routing Domain

1. Run Windows PowerShell with Administrative privileges.
2. Add a tenant name by using the following command.

```
Set-Remote AccessRoutingDomain -Name <Target Routing Domain Name> -TenantName <Domain Name of the VPN server for the tenant>
```

RRAS: The VPN Tenant Name for a Routing Domain must not be a subset of another Routing Domain's Tenant Name

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Error
Category	Configuration

Issue

The tenant name for the VPN interface for this Routing Domain is a subset of the Tenant Name of the VPN interface for another Routing Domain configured on the same multitenant gateway.

Impact

VPN Clients will not be able to connect to the remote server.

Resolution

Change the Tenant Name for the VPN interface for the Routing Domain.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change the tenant name for the VPN interface

1. Run Windows PowerShell with Administrative privileges.
2. Change the tenant name by running the following command.

```
Set-Remote AccessRoutingDomain -Name <Target RoutingDomain
Name> -TenantName <Domain Name of the VPN server for the
tenant>
```

RRAS: The default route (IPv4 or IPv6) should not be advertised to the peers

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Border Gateway Protocol (BGP) router is advertising a default route to peers.

Impact

Advertising the default route might disrupt the remote BGP peers' default routing.

Resolution

Add a Routing policy to drop the default route from advertisement from the BGP router.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a Routing policy to drop the default route from advertisement

1. Run Windows PowerShell with Administrative privileges.
2. Add a Routing policy to drop the default route from advertisement by running the following commands.

```
Add-BgpRoutingPolicy –Name <PolicyName> -MatchPrefix
@(“0.0.0.0/0”, “::/0”) –PolicyType “Deny”

Add-BgpRoutingPolicyForPeer –PolicyName <PolicyName> -
PeerName <List of target peers> -Direction “Egress”
```

RRAS: The default route (IPv4 or IPv6) should not be accepted from the peers

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Border Gateway Protocol (BGP) router is accepting default route advertisements from peers.

Impact

Accepting default gateway advertisements from the remote peers might disrupt the host router's default routing.

Resolution

Add a Routing policy to drop the default route from ingress route updates from the BGP router.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a Routing policy to drop the default route from advertisement

1. Run Windows PowerShell with Administrative privileges.
2. Add a Routing policy to drop the default route from advertisement by using the following commands.

```
Add-BgpRoutingPolicy –Name <PolicyName> -MatchPrefix
@("0.0.0.0/0", "::/0") –PolicyType "Deny"

Add-BgpRoutingPolicyForPeer –PolicyName <PolicyName> -
PeerName <List of target peers> -Direction "Ingress"
```

RRAS: The BGP peer IP address should not be assigned to a local network interface

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The IP address of the Border Gateway Protocol (BGP) peer is assigned to a local interface.

Impact

This incorrect configuration might lead to unexpected behavior by the BGP router.

Resolution

Check for the correct IP address of the BGP peer and change the peer IP address if required, or change the IP address of the local interface.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change the BGP peer IP address

1. Run Windows PowerShell with Administrative privileges.
2. Modify the peer IP address by using the following command.

```
Set-BgpPeer -Name <BGP Peer name> -PeerIPAddress <a different
non-local IP address>
```

RRAS: A local global IPv6 address must be configured on the BGP Router

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

IPv6 routing is configured for the Border Gateway Protocol (BGP) router, but a local global IPv6 address is not configured.

Impact

The BGP router cannot advertise IPv6 routes to its neighbors.

Resolution

Configure a Local Global IPv6 address for the BGP router.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a Local Global IPv6 address to the BGP router

1. Run Windows PowerShell with Administrative privileges.
2. Add a Local Global IPv6 address for the BGP router by using the following command.

```
Set-BGPRouter -LocalIPv6Address <Local Global IPv6 address to
be used as Next-Hop in IPv6 route advertisements>
```

RRAS: Multiple routes with different Next-Hop values and the same Destination prefix are configured

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Information
Category	Configuration

Issue

Multiple routes with the same destination and different Next-Hop values are configured on the Border Gateway Protocol (BGP) router.

Impact

Routing loops might occur because of multiple misconfigured routes with same destination prefix but different Next-Hop values. Connectivity might fail if routing loops occur.

Resolution

Remove the additional custom routes from the BGP configuration.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To remove additional custom routes

1. Run Windows PowerShell with Administrative privileges.
2. Remove the additional custom routes from the BGP configuration by using the following command.

```
Remove-BgpCustomRoute -Network <IP Prefixes>
```

RRAS: The BGP peer's Hold-Timer should not be set to the value 0

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Hold-Timer value is set to the value 0 for all Border Gateway Protocol (BGP) peers.

Impact

This might cause an infinite wait for a peer when it gets disconnected. When this peer attempts to reconnect, the connection attempt is rejected with collisions.

Resolution

Set the Hold-Timer to a different value (preferably greater than 20 seconds).

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change the Hold-Timer value

1. Run Windows PowerShell with Administrative privileges.
2. Set the Hold-Timer to a value greater than zero by using the following command.

```
Set-BgpPeer –Name <Peer Name> -HoldTimeSec <hold time in seconds>
```

RRAS: BGP peer's Hold-Timer should not be set to a very low value

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Border Gateway Protocol (BGP) router's Hold-Timer is set to a very low value.

Impact

A low Hold-Timer value might cause peering failure with third party routers that do not support very low values.

Resolution

Set the Hold-Timer to a different value (preferably one that is greater than 20 seconds).

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To set the Hold-Timer to a different value

1. Run Windows PowerShell with Administrative privileges.
2. Run the following command to set the Hold-Timer to a different value.

```
Set-BgpPeer –Name <Peer Name> -HoldTimeSec <hold time in seconds>
```

RRAS: All the ingress route advertisements should not be dropped because of a routing policy

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Information
Category	Configuration

Issue

A Routing policy configured on the Border Gateway Protocol (BGP) router is causing all the ingress route advertisements to be dropped.

Impact

None of the destination prefixes that are reachable from those peers is available locally.

Resolution

Remove or modify the misconfigured BGP routing policy or policies.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To remove or modify a BGP routing policy

1. Run Windows PowerShell with Administrative privileges.

   Run one of the following commands, depending on your deployment.

   - Remove the BGP routing policy or policies by running the following command.

     ```
     Remove-BgpRoutingPolicyForPeer -RoutingDomain <Routing Domain name> -PeerName <List of Peers> -PolicyName <List of policies to be removed>
     ```

- Modify the BGP routing policy or policies by running the following command.

```
Set-BgpRoutingPolicy -RoutingDomain <Routing Domain name> -
Name <Policy name to be edited> {Policy parameters}
```

RRAS: All the egress route advertisements should not be dropped because of a routing policy

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Information
Category	Configuration

Issue

A routing policy configured on the Border Gateway Protocol (BGP) peer is causing all the egress route advertisements to be dropped.

Impact

None of the destination prefixes that are reachable from the local router are available to those peers.

Resolution

Remove or modify the misconfigured BGP routing policy or policies.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To remove or modify a BGP routing policy

1. Run Windows PowerShell with Administrative privileges.

2. Run one of the following commands to remove or modify a BGP routing policy or set of policies.
   - You can use the following command to remove policies.

     ```
     Remove-BgpRoutingPolicyForPeer -RoutingDomain <Routing Domain name> -PeerName <List of Peers> -PolicyName <List of policies to be removed>
     ```

   - You can use the following command to modify an existing policy.

     ```
     Set-BgpRoutingPolicy -RoutingDomain <Routing Domain name> -Name <Policy name to be edited> {Policy parameters}
     ```

RRAS: BGP peers should not be configured for manual (passive) peering mode

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

Border Gateway Protocol (BGP) peers are configured for manual peering.

Impact

The manual peering mode is meant for troubleshooting purposes. In this mode, unless you manually start the peer, the BGP peer cannot connect or reconnect.

Resolution

Change the peering mode for the BGP peer from **Manual** *to* **Automatic**.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To change a BGP peering mode

1. Run Windows PowerShell with Administrative privileges.
2. Change the BGP Peer's Peering Mode by running the following command.

```
Set-BgpPeer –Name <Peer Name> -PeeringMode Automatic
```

Best Practices Analyzer for Remote Access: Configuration (Section 3)

The topics in this section can help you bring Remote Access running on Windows Server 2012 R2 into compliance with configuration best practices. Content in this section is most valuable when you have completed a Best Practices Analyzer scan of Remote Access and you want information about how to interpret and resolve scan results that identify areas of Remote Access that are noncompliant with configuration best practices.

Best Practices Analyzer and configuration rules

The Best Practices Analyzer applies configuration rules to identify settings that might require modification for Remote Access to perform optimally. Configuration rules can help prevent setting conflicts that can result in error messages or prevent Remote Access from carrying out its prescribed duties in an enterprise.

Topics in this section

This section includes the following topics.

- RRAS: IdleHoldTimer should not be set to a high value (> 10 sec)
- RRAS: Max Prefix policy should be configured for all BGP Peers
- RRAS: The total number of prefixes learned is in proximity of the Maximum Allowed Prefixes
- RRAS: A triggering route must be configured on the Site-to-Site VPN interface for the BGP peers
- RRAS: The Site-to-Site VPN triggering route should be a specific address (/32 for IPv4 address, /128 for IPv6 address)
- RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking BGP traffic
- RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VSID interface traffic
- RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VPN client traffic

- RRAS: The VPN static address pool should be configured as custom networks on the BGP router

RRAS: For BGP Peering over IPv6 addresses, IPv4 routes should not be configured for advertisement

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

IPv4 routes are configured for advertisement on the Border Gateway Protocol (BGP) router, while peering is established on the IPv6 addresses.

Impact

Advertised IPv4 prefixes are not reachable from the peer routers because of the absence of the IPv4 Next-Hop.

Resolution

Either remove the IPv4 routes from the advertisements, or reconfigure the BGP peering to be on the IPv4 addresses with IPv6 routing enabled and add the local IPv6 address to be used as the IPv6 Next-Hop.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

Perform one of the following procedures based upon your deployment.

To remove the IPv4 routes from advertisements

1. Run Windows PowerShell with Administrative privileges.
2. Remove the IPv4 routes from advertisements by running the following command.

```
Remove-BgpCustomRoute -Network <IPv4 addresses to remove> -
Interface <Interfaces with IPv4 addresses to be removed>
```

To reconfigure the BGP Peering to be on IPv4 addresses with IPv6 Routing enabled

1. Run Windows PowerShell with Administrative privileges.
2. Reconfigure the BGP Peering to be on IPv4 addresses with IPv6 Routing enabled and add the local IPv6 address to be used as the IPv6 Next-HOP by running the following commands.

```
Set-BgpRouter -LocalIPv6Address Enabled –LocalIPv6Address
<Local Global IPv6 address to be used as IPv6 NEXT-HOP>
```

```
Set-BgpPeer -Name <PeerName> -LocalIPAddress <LocalIPv4
Address> -PeerIPAddress <Peer's IPv4 Address>
```

RRAS: IdleHoldTimer should not be set to a high value (> 10 sec)

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Border Gateway Protocol (BGP) router's IdleHoldTimer is set to a very high value for the peers.

Impact

The BGP peer takes a longer time to get connected.

Resolution

Update the IdleHoldTimer for the peers for the Routing Domain to an appropriate value. (A value that is preferably greater than, or equal to, 10 seconds).

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To update the IdleHoldTimer for peers

1. Run Windows PowerShell with Administrative privileges.
2. Update the IdleHoldTimer for the peers that you specify with the following Windows PowerShell command.

```
Set-BgpPeer –Name <Peer Name> -IdleHoldTimer 10
```

RRAS: Max Prefix policy should be configured for all BGP Peers

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Max Prefix policy is not configured for Border Gateway Protocol (BGP) Peers.

Impact

The BGP Router is vulnerable to denial of service attacks from these peers.

Resolution

Configure the MaxAllowedPrefix for the peer.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To configure the MaxAllowedPrefix value for a peer

1. Run Windows PowerShell with Administrative privileges.
2. Provide appropriate values for all parameters, and then run the following command.

```
Set-BgpPeer –Name <PeerName> -MaxAllowedPrefix <new prefix count>
```

RRAS: The total number of prefixes learned is in proximity of the Maximum Allowed Prefixes

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Information
Category	Configuration

Issue

The threshold for the maximum allowed prefixes is about to be reached for the Border Gateway Protocol (BGP) Peer.

Impact

BGP Peering will be restarted and purge all of the previously learned routes.

Resolution

Reset or update the total number of MaxAllowedPrefix for the peer.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the MaxAllowedPrefix value for a peer

1. Run Windows PowerShell with Administrative privileges.
2. You can use one of the following commands to either update the number of prefixes that can be learned from a peer or to remove the set limit on the number of prefixes that are learned.
   - Use this command to update the maximum number of prefixes that can be learned from a peer.

     ```
     Set-BgpPeer –Name <PeerName> -MaxAllowedPrefix <new count>
     ```

   - Use this command to remove the set limit of the number of prefixes that are learned.

     ```
     Set-BgpPeer –Name <PeerName> -ClearPrefixLimit
     ```

RRAS: A triggering route must be configured on the Site-to-Site VPN interface for the BGP peers

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

No triggering route is configured on the Site-to-Site VPN interface.

Impact

The Site-to-Site VPN interface cannot trigger the connection to the remote site as the demand dial. It has to be manually dialed.

Resolution

Add a triggering route on the Site-to-Site VPN interface for corresponding Border Gateway Protocol (BGP) peers.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a specific triggering route to an interface

1. Run Windows PowerShell with Administrative privileges.
2. Add a specific triggering route, where Prefix Length is 32 for IPv4 addresses and 128 for IPv6 addresses, by using one of the following Windows PowerShell commands.

```
Add-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> [configuration options] –Ipv4Subnet <Triggering
IPv4 route/Prefix Length:Metric>


Add-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> [configuration options] –Ipv6Subnet <Triggering
IPv6 route/Prefix Length:Metric>


Set-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> –Ipv4Subnet <Triggering IPv4 route/Prefix
Length:Metric>
Set-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> –Ipv6Subnet <Triggering IPv6 route/Prefix
Length:Metric>
```

RRAS: The Site-to-Site VPN triggering route should be a specific address (/32 for IPv4 address, /128 for IPv6 address)

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote

Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The Site-to-Site VPN triggering address that is configured for the interface is a generic address.

Impact

If the triggering route that is configured on the Site-to-Site VPN interface is a generic address, the Border Gateway Protocol (BGP) router's next-hop resolution that corresponds to those generic addresses might fail.

Resolution

Add a specific, non-generic triggering route, where the Prefix Length is 32 for IPv4 addresses and 128 for IPv6 addresses.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add a specific triggering route to an interface

1. Run Windows PowerShell with Administrative privileges.
2. Add a specific triggering route, where Prefix Length is 32 for IPv4 addresses and 128 for IPv6 addresses, by using one of the following Windows PowerShell commands.

```
Add-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> [configuration options] –Ipv4Subnet <Triggering
IPv4 route/Prefix Length:Metric>


Add-VpnS2SInterface –Name <Name of the Site-to-Site VPN
interface> [configuration options] –Ipv6Subnet <Triggering
IPv6 route/Prefix Length:Metric>


Set-VpnS2SInterface –Name <Name of the Site-to-Site VPN
```

```
interface> -Ipv4Subnet <Triggering IPv4 route/Prefix
Length:Metric>
Set-VpnS2SInterface -Name <Name of the Site-to-Site VPN
interface> -Ipv6Subnet <Triggering IPv6 route/Prefix
Length:Metric>
```

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking BGP traffic

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

Traffic filters that are configured for the Site-to-Site VPN interface are blocking Border Gateway Protocol (BGP) traffic.

Impact

BGP Routing will not work.

Resolution

Reset the Site-to-Site VPN interface's traffic filters to unblock BGP traffic.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Site-to-Site VPN interface's traffic filters to unblock BGP traffic

1. Run Windows PowerShell with Administrative privileges.

2. Reset the Site-to-Site VPN interface's traffic filters by using one of the following Windows PowerShell commands.
   - By using this command, you can specify a list of interface filters to reset.

     ```
     Add-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny> -List <List of filters to be applied in
     the format  ->
     SourceIP/mask:DestinationIP/Mask:Protocol:ProtocolData1:Protoc
     olData2:>
     ```

   - By using this command, all filters for the interface are reset.

     ```
     Set-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny>
     ```

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VSID interface traffic

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The traffic filters that are configured for the Site-to-Site VPN interface are blocking VSID interface traffic.

Impact

The remote VPN server and VPN clients cannot access hosted organization resources.

Resolution

Reset the Site-to-Site VPN interface's traffic filters to unblock VSID interface traffic.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Site-to-Site VPN interface's traffic filters

1. Run Windows PowerShell with Administrative privileges.
2. Reset the Site-to-Site VPN interface's traffic filters by running one of the following Windows PowerShell commands.
   - By using this command, you can specify a list of interface filters to reset.

     ```
     Add-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny> -List <List of filters to be applied in
     the format  ->
     SourceIP/mask:DestinationIP/Mask:Protocol:ProtocolData1:Protoc
     olData2:>
     ```

   - By using this command, all filters for the interface are reset.

     ```
     Set-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny>
     ```

RRAS: The Site-to-Site VPN interface's traffic filters should not be blocking VPN client traffic

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2

Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The traffic filters that are configured for the Site-to-Site VPN interface are blocking VPN client traffic.

Impact

VPN remote clients cannot access hosted organization resources.

Resolution

Reset the Site-to-Site VPN interface's traffic filters to unblock VPN client traffic.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Site-to-Site VPN interface's traffic filters

1. Run Windows PowerShell with Administrative privileges.
2. Reset the Site-to-Site VPN interface's traffic filters by running one of the following Windows PowerShell commands.
   - By using this command, you can specify a list of interface filters to reset.

     ```
     Add-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny> -List <List of filters to be applied in
     the format  ->
     SourceIP/mask:DestinationIP/Mask:Protocol:ProtocolData1:Protoc
     olData2:>
     ```

   - By using this command, all filters for the interface are reset.

     ```
     Set-RemoteAccessIPFilter –InterfaceAlias <Interface Name> -
     AddressFamily <IPv4/ IPv6> -Direction <Inbound / Outbound> -
     Action <Allow / Deny>
     ```

RRAS: The VPN static address pool should be configured as custom networks on the BGP router

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Configuration

Issue

The VPN interface's static address pool is not configured as a custom network on the BGP router.

Impact

VPN clients cannot communicate with the on-premise organization resources using the service provider gateway.

Resolution

Configure the static pool IP addresses to the BGP custom network for the Routing Domain.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To add the static pool IP addresses to the BGP custom network

1. Run Windows PowerShell with Administrative privileges.
2. Add the static pool IP addresses to the BGP custom network for the Routing Domain by running the following Windows PowerShell command.

```
Add-BgpCustomRoute [-RoutingDomain <Routing Domain name> for
MT System] -Network <VPN static pool IP addresses>
```

Best Practices Analyzer for Remote Access: Performance and Operation

The topics in this section can help you bring Remote Access running on Windows Server 2012 R2 into compliance with configuration best practices. Content in this section is most valuable when you have completed a Best Practices Analyzer scan of Hyper-V and you want information about how to interpret and resolve scan results that identify areas of Hyper-V that are noncompliant with configuration best practices.

Best Practices Analyzer performance and operation rules

The Best Practices Analyzer applies performance and operation rules to identify best-practice–related, possible causes of a server role's failure to carry out its prescribed tasks in an enterprise. An example of a violation of operation rules that a Best Practices Analyzer scan might find is a service that is paused or stopped.

Topics in this section

This section contains the following topics.

- RRAS: Rate-limiting (Tx/Rx BandwidthKbps) should be set to a value as per the network requirements
- RRAS: Rate-limiting parameters (Tx/Rx BandwidthKbps) should not have a significant difference in the values
- RRAS: The CapacityKbps parameter should be configured with a value as per the network requirements
- RRAS: The routes being advertised to the peers must be locally resolvable

RRAS: Rate-limiting (Tx/Rx BandwidthKbps) should be set to a value as per the network requirements

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2

Product/Feature	Remote Access
Severity	Warning
Category	Performance

Issue

Rate-limiting parameters are set to the default values for the Site-to-Site VPN interface of the Routing Domain.

Impact

Site-to-Site VPN or VPN interface's throughput is limited by the default value of the rate-limiting parameter.

Resolution

Reset the Site-to-Site VPN or VPN interface's Quality of Service (QoS) parameters to the appropriate values.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Site-to-Site VPN interface's QoS parameters

1. Run Windows PowerShell with Administrative privileges.
2. Set the bandwidth to a specific value by running the following Windows PowerShell command.

```
Set-VpnS2SInterface –Name <Site-to-Site VPN interface Name> -
TxBandwidthKbps <number> -RxBandwidthKbps <number>
```

RRAS: Rate-limiting parameters (Tx/Rx BandwidthKbps) should not have a significant difference in the values

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote

Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Performance

Issue

Rate-limiting parameter values are too far apart for the Site-to-Site VPN interface of the Routing Domain.

Impact

Site-to-Site VPN performance is affected.

Resolution

Set the Quality of Service (QoS) parameters to the appropriate values (where smaller value greater than or equal to 30 percent of the larger value).

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Site-to-Site VPN interface's QoS parameters

1. Run Windows PowerShell with Administrative privileges.
2. Set the QoS parameters to a specific value by running the following Windows PowerShell command.

```
Set-VpnS2SInterface –Name <Site-to-Site VPN interface name> -
TxBandwidthKbps <number> -RxBandwidthKbps <number>
```

RRAS: The CapacityKbps parameter should be configured with a value as per the network requirements

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Performance

Issue

The Capacity Kbps parameter is set to the default value.

Impact

The Remote Access server's performance and throughput is impacted.

Resolution

Set the Remote Access server's CapacityKbps parameter to the appropriate value.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To reset the Remote Access server's CapacityKbps parameter

1. Run Windows PowerShell with Administrative privileges.
2. Reset the Remote Access server's CapacityKbps parameter by running the following Windows PowerShell command.

```
Set-RemoteAccess -CapacityKbps <number>
```

RRAS: The routes being advertised to the peers must be locally resolvable

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Remote Access Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Remote Access
Severity	Warning
Category	Operation

Issue

The next-hops for the custom routes configured on the Border Gateway Protocol (BGP) router for the Routing Domain cannot be resolved locally.

Impact

The remote peers get the routes that are being advertised by the local router, but they cannot perform the routing to the destination prefix.

Resolution

Remove the unresolvable routes from the BGP advertisement.

Membership in Administrators, or equivalent, is the minimum required to perform these procedures.

To remove the unresolvable routes from the BGP advertisement

1. Run Windows PowerShell with Administrative privileges.
2. Remove the unresolvable routes from the BGP advertisement by running the following Windows PowerShell command.

```
Remove-BgpCustomRoute -Network <IP Prefixes>
```

Work Folders Best Practices Analyzer

Topics in this section can help you bring Work Folders running on Windows Server 2012 R2 into compliance with best practices. Content in this section is most valuable to administrators who have completed a Best Practices Analyzer scan of File and Storage Services, and who want information about how to interpret and resolve the scan results that identify areas of Work Folders that are noncompliant with best practices.

For more information about Best Practices Analyzer and scans, see Run Best Practices Analyzer Scans and Manage Scan Results.

The Windows Sync Share service should be set to start automatically

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

The Windows Sync Share service is not set to start automatically.

Impact

Users will not be able to sync with this server using Work Folders.

Resolution

Start the Windows Sync Share service and set it to start automatically.

See also

Work Folders should be installed on all nodes of the failover cluster

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

Work Folders is not installed on a node of a failover cluster that is hosting sync shares.

Impact

Work Folders will not fail over to any nodes that do not have the role service installed, potentially resulting in users not being able to sync with this clustered file server instance.

Resolution

Install the Work Folders role service on the cluster node.

See also

All nodes in a failover cluster should be reachable to Work Folders

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Operation

Issue

The following nodes in a failover cluster cannot be verified because the nodes are not reachable, or because the BPA scan was performed against a remote failover cluster

Impact

Best Practices Analyzer will not be able to check the status of Work Folders on the following nodes

Resolution

Confirm that the nodes are online, log on to the nodes locally (or by using Remote Desktop), and then run the Best Practices Analyzer scan again.

See also

The Work Folders server should be domain joined

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

No groups or users are granted sync access to the sync share.

Impact

Users will not be able to sync with this sync share using Work Folders.

Resolution

Grant sync access to the groups that are allowed to sync with this sync share.

See also

The sync share should be located in a valid folder

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

The path for the sync share does not exist

Impact

Users will not be able to sync with this sync share using Work Folders.

Resolution

Reinitialize the sync share by using the Set-SyncShare cmdlet from a Windows PowerShell session.

See also

T:SyncShare.Set-SyncShare

A current SSL certificate should be configured for this Work Folders server

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

An SSL certificate is either not configured or is expired.

Impact

Users will not be able to sync with this sync share using the encrypted HTTPS protocol. Users might be able to sync with the unencrypted HTTP protocol, if it's configured.

Resolution

Install a new SSL certificate and bind it to the server

See also

Windows Firewall should open port 80 and 443 for Work Folders

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article

can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services
Severity	Error
Category	Configuration

Issue

Work Folders is enabled, but both port 80 and 443 are closed in Windows Firewall.

Impact

Users will not be able to sync with this server using Work Folders.

Resolution

Use Windows Firewall to open port 443 and port 80.

See also

A staging area should exist for the sync share

This article is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this article to computers on which Best Practices Analyzer scans have been run, and that are experiencing the issue addressed by this article. This article can also be useful as general troubleshooting and best practices information to help you configure your server. For more information about best practices and scans, see Run Best Practices Analyzer Scans and Manage Scan Resultshttp://technet.microsoft.com/library/hh831400.aspx.

Operating System	Windows Server 2012 R2
Product/Feature	File and Storage Services

Severity	Error
Category	Configuration

Issue

The staging area for a sync share does not exist.

Impact

Users will not be able to sync with this sync share using Work Folders.

Resolution

Reinitialize the sync share by using the Set-SyncShare cmdlet from a Windows PowerShell session.

See also

T:SyncShare.Set-SyncShare

Best Practices Analyzer for Web Application Proxy

In Windows management, *best practices* are guidelines that are considered the ideal way, under normal circumstances, to configure a server as defined by experts. For example, it is considered a best practice for most server technologies to keep open only those ports required for the technologies to communicate with other networked computers, and block unused ports. While best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Topics in this section can help you bring Web Application Proxy running on Windows Server 2012 R2 into compliance with best practices. Content in this section is most valuable to administrators who have completed a Best Practices Analyzer scan of Web Application Proxy, and who want information about how to interpret and resolve scan results that identify areas of Web Application Proxy that are noncompliant with best practices.

For more information about Best Practices Analyzer and scans, see Run Best Practices Analyzer Scans and Manage Scan Results.

More information about Web Application Proxy

Web Application Proxy enables you to publish selected HTTP- and HTTPS-based applications from your corporate network to client devices outside of the corporate network. It can use Active Directory Federation Services (AD FS) to ensure that users are authenticated before they gain access to published applications. Web Application Proxy also provides proxy functionality for AD FS servers.

Topics in this section

- Web Application Proxy must be configured before it is used
- Web Application Proxy: The external and backend server URLs are different and URL translation is disabled
- Web Application Proxy: The service is not configured to run automatically
- Web Application Proxy: The AD FS Proxy service is not configured to run automatically
- Web Application Proxy: This server is not included in the ConnectedServersName list
- Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high
- Web Application Proxy: Application is using an external certificate that is not yet valid
- Web Application Proxy: Application is using an external certificate that is about to expire
- Web Application Proxy: Application is using an external certificate that has no private key
- Web Application Proxy: Application is using an external certificate that has expired
- Web Application Proxy: Application is configured to use an external certificate that is not present on this server
- Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows authentication but the server is not joined to a domain
- Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed

Web Application Proxy must be configured before it is used

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy

Severity	Error
Category	Configuration

Issue

Web Application Proxy was installed but not configured for initial use.

Impact

If Web Application Proxy is not configured it cannot be used to publish applications.

Resolution

Use 'Remote Access Management' in Server Manager to start the configuration wizard or use the Install-WebApplicationProxy PowerShell command.

After installing the Web Application Proxy role service on a server, you must also configure Web Application Proxy on the server. If you do not configure Web Application Proxy, you will be unable to publish applications and the Web Application Proxy server will not provide AD FS proxy functionality.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To configure Web Application Proxy

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the navigation pane, click **Web Application Proxy**.
3. In the Remote Access Management console, in the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
4. On the **Web Application Proxy Configuration Wizard**, on the **Welcome** dialog, click **Next**.
5. On the **Federation Server** dialog, do the following, and then click **Next**:
   - In the **Federation service name** box, enter the fully qualified domain name (FQDN) of the AD FS server; for example, fs.contoso.com.
   - In the **User name** and **Password** boxes, enter the credentials of a local administrator account on the AD FS servers.
6. On the **AD FS Proxy Certificate** dialog, in the list of certificates currently installed on the Web Application Proxy server, select a certificate to be used by Web Application Proxy for AD FS proxy functionality, and then click **Next**.

7. On the **Confirmation** dialog, review the settings. If required, you can copy the PowerShell cmdlet to automate additional installations. Click **Configure**.
8. On the **Results** dialog, verify that the configuration was successful, and then click **Close**.

Web Application Proxy: The external and backend server URLs are different and URL translation is disabled

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

The external and backend server URLs are different and URL translation is disabled.

Impact

The published application might reject requests from the client.

Resolution

Consider changing the application publishing settings.

By default, Web Application Proxy translates the host portion of requests to a backend server. For example, Web Application Proxy will translate the URLs successfully if the external URL is https://apps.contoso.com/ and the backend server URL is https://appsinternal.contoso.com/. However, URL translation is currently disabled, which might cause client requests to be rejected. You can manually enable the translation of host headers by using the DisableTranslateUrlInRequestHeaders parameter.

You may need to remove the published application and republish it to resolve this issue.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.

     Note
     Web Application Proxy can translate host names in URLs, but cannot translate path names. Therefore, you can enter different host names, but you must enter the same path name. For example, you can enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://app-server/app1/. However, you cannot enter an external URL of

https://apps.contoso.com/app1/ and a backend server URL of https://apps.contoso.com/internal-app1/.

- In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

You can view the settings of a published application and enable URL translation, if required.

To enable URL translation using PowerShell

1. To obtain the application ID of an application, use the following PowerShell command to show the ID, name, and external URL of all applications published by Web Application Proxy, and locate the application ID of the required application in the command output:

```
Get-WebApplicationProxyApplication | Format-Table ID, Name,
ExternalURL
```

2. Use the following PowerShell command to enable URL translation for a specific published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
DisableTranslateUrlInRequestHeaders:$false
```

Web Application Proxy: The service is not configured to run automatically

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Web Application Proxy service (appproxysvc) settings were changed and it is not configured to run automatically.

Impact

After the server restarts, the service will not start and users will not be able to access published applications.

Resolution

Consider changing the Web Application Proxy service startup type to Automatic.

In order for Web Application Proxy to work correctly, the appproxysvc service must be running. It is recommended that this service is set to automatic to ensure that whenever the Web Application Proxy server restarts, the service will restart without any user action.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To set the appproxysvc service startup type to automatic

1. On the Web Application Proxy server, open the Services console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **services.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the details pane, right-click **Web Application Proxy Service**, and then click **Properties**.
3. On the **Web Application Proxy Service** dialog, on the **General** tab, in **Startup type**, click **Automatic**, and then click **OK**.

Web Application Proxy: The AD FS Proxy service is not configured to run automatically

---

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Active Directory Federation Services service (adfssrv) settings were changed and it is not configured to run automatically.

Impact

After the server restarts, the service will not start and users will not be able to authenticate and will not be able to access published applications.

Resolution

Consider changing the Active Directory Federation Services service startup type to Automatic.

In order for Web Application Proxy to work correctly, the adfssrv service must be running. It is recommended that this service is set to automatic to ensure that whenever the Web Application Proxy server restarts, the service will restart without any user action.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To set the adfssrv service startup type to automatic

1. On the Web Application Proxy server, open the Services console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **services.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the details pane, right-click **Active Directory Federation Services**, and then click **Properties**.
3. On the **Active Directory Federation Services** dialog, on the **General** tab, in **Startup type**, click **Automatic**, and then click **OK**.

Web Application Proxy: This server is not included in the ConnectedServersName list

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

This server is not included in the ConnectedServersName list.

Impact

The Remote Access Management console may not work as expected.

Resolution

Use the Set-WebApplicationProxyConfiguration cmdlet to add this server to the ConnectedServersName list.

All Web Application Proxy servers must be included in the ConnectedServersName list, otherwise they will not appear in the Remote Access Management console and attempts to change or configure published applications may result in unexpected behavior. You can manually configure the ConnectedServersName list only using Windows PowerShell.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To add a server to the ConnectedServersName list

1. On the Web Application Proxy server, open a Windows PowerShell window.
2. Use the following command to view the current ConnectedServersName list:

```
Get-WebApplicationProxyConfiguration
```

3. Copy the current ConnectedServersName list to the clipboard.

4. Run the following command to add a server to the ConnectedServersName list:

```
Set-WebApplicationProxyConfiguration –ConnectedServersName
<paste-existing-list-here>,<new_server_name>.<domain_name>
```

Note
You must provide the existing list of connected servers, and then add each new server as a comma-separated list for this command to be effective.

Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

When the ConfigurationChangesPollingIntervalSec value is high, any changes that you make are not propagated in a timely manner, or changes are only partially propagated until the next time the server checks for updates.

Impact

Applications might not be published as expected until the next time the server checks for updates.

Resolution

Consider changing ConfigurationChangesPollingIntervalSec to a lower value using the Set-WebApplicationProxyConfiguration cmdlet.

Web Application Proxy servers must periodically check if there have been any configuration changes made to the Web Application Proxy configuration that is stored on the Active Directory Federation Services (AD FS) servers. The time between these checks is set by the ConfigurationChangesPollingIntervalSec parameter. If the value is low, Web Application Proxy checks for changes more frequently, but does increase internal network traffic. However, if the value is high, Web Application Proxy may operate for a significant period of time before starting to use the new configuration. The default value for this parameter is 30 seconds.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To change how frequently Web Application Proxy checks for configuration changes

1. On the Web Application Proxy server, open a Windows PowerShell window.
2. Use the following command to view the current ConfigurationChangesPollingIntervalSec value:

   ```
   Get-WebApplicationProxyConfiguration
   ```

3. Run the following command to change the ConfigurationChangesPollingIntervalSec value to 15 seconds:

   ```
   Set-WebApplicationProxyConfiguration –
   ConfigurationChangesPollingIntervalSec 15
   ```

Web Application Proxy: Application is using an external certificate that is not yet valid

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that is not yet valid.

Impact

If an application is using a certificate that is not yet valid, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Publish this application again with a valid certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication is not yet valid, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.

10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then

click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that is about to expire

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that is about to expire.

You will receive this message for applications whose certificates will expire in under 60 days.

Impact

If an application is using a certificate that has expired, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Issue a new certificate for this address and publish this application again with the new certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication expires, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell

command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that has no private key

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that has no private key.

Impact

If an application is using a certificate that has no private key, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Publish this application again with a valid certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application has no private key, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web**

Application Proxy.

3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published

application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that has expired

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that has expired.

Impact

If an application is using a certificate that has expired, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Issue a new certificate for this address and publish this application again with the new certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication expires, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new

certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

   ```
   Get-WebApplicationProxyApplication –Name <application_name> |
   Format-Table ID, ExternalURL
   ```

3. Use the following PowerShell command to change the certificate used for the published application:

   ```
   Set-WebApplicationProxyApplication –ID <application_ID> -
   ExternalCertificateThumbprint <New_certificate_thumbprint>
   ```

Web Application Proxy: Application is configured to use an external certificate that is not present on this server

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using an external certificate that is not present on this server.

Impact

If an application uses a certificate that is not present on this server, Web Application Proxy will not publish the application and users will not have access.

Resolution

Obtain the configured external certificate from another server and copy it to this server with its private key.

For every application that you publish through Web Application Proxy, each Web Application Proxy server in a multi-server deployment must have a copy of the certificate used for authentication for the website with the private key.

To export the application certificate

1. On an Web Application Proxy server that has the required certificate, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. In the details pane, right-click the required certificate, click **All Tasks**, and then click **Export**.
8. On the **Certificate Export Wizard**, click **Next**.
9. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
10. On the **Export File Format** page, make sure **Personal Information Exchange** is selected, select the **Include all certificates in the certification path if possible** and **Export all extended properties** check boxes, and then click **Next**.
11. On the **Security** page, decide how to protect this certificate, and then click **Next**.
12. On the **File to Export** page, save the certificate in a location that can be accessed by the Web Application Proxy server that requires this certificate, and then click **Next**, and then click **Finish**.
13. If the export is successful, click **OK**.

After exporting the application certificate, you must import it on the Web Application Proxy server that requires the certificate.

To import the application certificate

1. On the Web Application Proxy server that requires the certificate, open an MMC console as described above.
2. In the Console window, open **Certificates/Personal/Certificates**.
3. Right-click in the details pane, click **All Tasks**, and then click **Import**.
4. On the **Certificate Import Wizard**, click **Next**.
5. On the **File to Import** page, locate the certificate that you exported previously, and then click **Next**.
6. On the **Private key protection** page, enter the password if required, and then click **Next**.
7. On the **Certificate Store** page, click **Next**, and then click **Finish**.
8. If the import is successful, click **OK**.

Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows

authentication but the server is not joined to a domain

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy can perform backend authentication using Integrated Windows authentication only when it is running on a server that is joined to a domain.

Impact

Users will not be able to access this application from the current server.

Resolution

Join this server to a domain.

When publishing applications that use Integrated Windows authentication, the Web Application Proxy server uses Kerberos constrained delegation to authenticate users to the published application.

To use Integrated Windows authentication, the Web Application Proxy server must be joined to an AD DS domain. The following lists the domain and forest requirements for a deployment using Integrated Windows authentication with Kerberos constrained delegation.

- Deployments where users, resources, and Web Application Proxy servers are all in the same forest are supported.
- In deployments with multiple forests where there is a user forest, a resource forest, and a Web Application Proxy forest, the following deployments are supported:
  - Users, resources, and Web Application Proxy servers are all in different forests.
  - Users and Web Application Proxy servers are in the same forest, but resources are in a different forest.

- Resources and Web Application Proxy servers are in the same forest, but users are in a different forest.
- Users and resources are in the same forest, but Web Application Proxy servers are in a different forest.

In multi-forest deployments:

1. The user forest must trust the Web Application Proxy forest, and the Web Application Proxy forest must trust the resource forest.
2. All of the Active Directory domains in a multi-forest deployment must have at least one Windows Server® 2012 or higher domain controller. For more information, see Kerberos Constrained Delegation across Domains

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To join the Web Application Proxy server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. On the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In **Computer Name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
5. When you see a dialog box welcoming you to the domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed on this server.

Impact

Web Application Proxy does not support this configuration and might not work as expected.

Resolution

Consider uninstalling DirectAccess from this server.

When Web Application Proxy is deployed in a multiple server deployment, it cannot be deployed on a server that also has DirectAccess installed. You must uninstall either DirectAccess, or Web Application Proxy. The following table shows the supported deployments:

DirectAccess	VPN	Web Application Proxy
Single server deployment	Single server deployment	Single server deployment
Multisite deployment	Multiple server deployment	Not supported on the same server
Not supported on the same server	Multiple server deployment	Multiple server deployment
Cluster deployment[1]	Multiple server deployment	Multiple server deployment[2]

Notes

1—In a pre-existing DirectAccess cluster deployment, you can install Web Application Proxy only using Windows PowerShell.

2—In a pre-existing multiple server Web Application Proxy deployment, you can install DirectAccess only using Windows PowerShell.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To uninstall DirectAccess

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **DirectAccess and VPN**.
3. In the **Tasks** pane, click **Remove Configuration Settings**.
4. On the **Confirm Remove Configuration** dialog, click **OK**.
5. On the **Removing Configuration Settings** dialog, make sure the configuration was removed, and then click **Close**.

Server Roles and Technologies in Windows Server 2012 R2 and Windows Server 2012

This section contains information to design, deploy, manage, and troubleshoot technologies in Windows Server 2012 R2 and Windows Server 2012.

- Active Directory Certificate Services Overview
  This content provides an overview of Active Directory Certificate Services (AD CS) in Windows Server 2012. AD CS is the server role that allows you to build a public key infrastructure (PKI) and provide public key cryptography, digital certificates, and digital signature capabilities for your organization.
- Active Directory Domain Services Overview
  By using the Active Directory Domain Services (AD DS) server role, you can create a scalable, secure, and manageable infrastructure for user and resource management, and provide support for directory-enabled applications such as Microsoft Exchange Server.
- Active Directory Federation Services Overview
  This topic provides an overview of Active Directory Federation Services (AD FS) in Windows Server 2012.
- Active Directory Lightweight Directory Services Overview
  Active Directory Lightweight Directory Services (AD LDS) is a Lightweight Directory Access Protocol (LDAP) directory service that provides flexible support for directory-enabled applications, without the dependencies and domain-related restrictions of AD DS.
- Active Directory Rights Management Services Overview
  This document provides an overview of Active Directory Rights Management Services (AD RMS) in Windows Server 2012. AD RMS is the server role that provides you with

management and development tools that work with industry security technologies—including encryption, certificates, and authentication—to help organizations create reliable information protection solutions.

- Application Server Overview
  Application Server provides an integrated environment for deploying and running custom, server-based business applications.
- Failover Clustering Overview
  This topic describes the Failover Clustering feature and provides links to additional guidance about creating, configuring, and managing failover clusters on up to 4,000 virtual machines or up to 64 physical nodes.
- File and Storage Services Overview
  This topic discusses the File and Storage Services server role in Windows Server 2012, including what's new, a list of role services, and where to find evaluation and deployment information.
- Group Policy Overview
  This topic describes the Group Policy feature in Windows Server 2012 and Windows 8. Use this topic to find the documentation resources and other technical information you need to accomplish key Group Policy tasks, new or updated functionality in this version compared to previous versions of Group Policy, and ways to automate common Group Policy tasks using Windows PowerShell.
- Hyper-V Overview
  This topic describes the Hyper-V role in Windows Server 2012—practical uses for the role, the most significant new or updated functionality in this version compared to previous versions of Hyper-V, hardware requirements, and a list of operating systems (known as *guest operating systems*) supported for use in a Hyper-V virtual machine.
- Networking Overview
  This section contains detailed information about networking products and features for the IT professional to design, deploy, and maintain Windows Server 2012.
- Network Load Balancing Overview
  By managing two or more servers as a single virtual cluster, Network Load Balancing (NLB) enhances the availability and scalability of Internet server applications such as those used on web, FTP, firewall, proxy, virtual private network (VPN), and other mission-critical servers. This topic describes the NLB feature and provides links to additional guidance about creating, configuring, and managing NLB clusters.
- Network Policy and Access Services Overview
  This topic provides an overview of Network Policy and Access Services in Windows Server 2012, including the specific role services of Network Policy Server (NPS), Health Registration Authority (HRA), and Host Credential Authorization Protocol (HCAP). Use the Network Policy and Access Services server role to deploy and configure Network Access Protection (NAP), secure wired and wireless access points, and RADIUS servers and proxies.
- Print and Document Services Overview
  This is an overview of Print and Document Services, including Print Server, Distributed Scan Server, and Fax Server in Windows Server 2012.

- Remote Desktop Services Overview
  Remote Desktop Services accelerates and extends desktop and application deployments to any device, improving remote worker efficiency, while helping to keep critical intellectual property secure and simplify regulatory compliance. Remote Desktop Services enables both a virtual desktop infrastructure (VDI) and session-based desktops, allowing users to work anywhere.
- Security and Protection
  The table on this page provides links to available information for the IT pro about security technologies and features for Windows Server 2012 and Windows 8.
- Telemetry Overview
  Find out about Windows Feedback Forwarder—a service that enables you to automatically send feedback to Microsoft by deploying a Group Policy setting to one or more organizational units. Windows Feedback Forwarder is available on all editions of Windows Server 2012.
- Volume Activation Overview
  This technical overview for the IT pro describes the volume activation technologies in Windows Server 2012 and how your organization can benefit from using these technologies to deploy and manage volume licenses for a medium to large number of computers.
- Web Server (IIS) Overview
  This document introduces the Web Server (IIS) role of Windows Server 2012, describes new IIS 8 features, and links to additional Microsoft and community information about IIS.
- Windows Deployment Services Overview
  Windows Deployment Services enables you to deploy Windows operating systems over the network, which means that you do not have to install each operating system directly from a CD or DVD.
- Windows Server Backup Feature Overview
  This section provides an overview of the Windows Server Backup feature and lists the new features in Windows Server 2012.
- Windows Server Essentials Experience Overview
  With the Windows Server Essentials Experience role, you can take advantage of Windows Server 2012 R2 Essentials features such as simplified management using the server dashboard, data protection, Remote Web Access, and integration with Microsoft online services—all without enforcement of the Windows Server 2012 R2 Essentials locks and limits.
- Windows Server Update Services Overview
  Windows Server Update Services (WSUS) enables information technology administrators to deploy the latest Microsoft product updates. By using WSUS, administrators can fully manage the distribution of updates that are released through Microsoft Update to computers in their network. In Windows Server 2012, this feature is integrated with the operating system as a server role. This topic provides an overview of this server role and more information about how to deploy and maintain WSUS.
- Windows System Resource Manager Overview
  With Windows System Resource Manager for the Windows Server 2012 operating system, you can manage server processor and memory usage with standard or custom resource policies. Managing your resources can help ensure that all the services provided by a single

server are available on an equal basis or that your resources will always be available to high-priority applications, services, or users.

Active Directory

- What's New in Active Directory in Windows Server 2012 R2
- Active Directory Certificate Services Overview
- Active Directory Domain Services Overview
- Active Directory Federation Services Overview
- Active Directory Lightweight Directory Services Overview
- Active Directory Rights Management Services Overview

What's New in Active Directory in Windows Server 2012 R2

The content in this section describes what's new and changed in Active Directory in Windows Server® 2012 R2. This content focuses on changes that will potentially have the greatest impact on your use of this release.

One of the most prevalent IT industry trends at the moment is the proliferation of consumer devices in the workplace. Employees and partners want to access protected corporate data from their personal devices, from checking email to the consumption of advanced business applications. IT administrators in organizations, while wanting to enable this level of productivity, would like to continue to ensure that they can manage risk and govern the use of corporate resources.

In Windows Server® 2012 R2, Active Directory has been enhanced with the following value propositions to allow IT risk management while also enabling IT to empower their users to be productive from a variety of devices:

- IT administrators can allow devices to be associated with the company's Active Directory and use this association as a seamless second factor authentication.
- Single sign-on (SSO) from devices that are associated with the company's Active Directory
- Enable users to connect to applications and services from anywhere with Web Application Proxy
- Manage the risk of users working from anywhere, accessing protected data from their devices, with Multi-factor Access Control and Multi-Factor Authentication (MFA)

These value propositions are described in detail in the following guides:

Guide Name	Description
Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications	This guide describes the key concepts and provides the step-by-step walkthrough instructions for Workplace Join: configuring the Device Registration Service (DRS), workplace join with a Windows device, workplace join with an iOS device, and how to deal with lost or stolen devices.
Overview: Connect to Applications and Services from Anywhere with Web Application Proxy	This guide describes and provides step-by-step walkthrough instructions for using Web Application Proxy – a new Remote Access role service in Windows Server® 2012 R2 Preview – to provide access to a sample web application using claims-based authentication using AD FS authentication.
Overview: Manage Risk with Multi-Factor Access Control	This guide describes the enhanced IT risk management strategies available in AD FS in Windows Server® 2012 R2 Preview and provides step-by-step walkthrough instructions for configuring and verifying the solution of managing risk with multi-factor access control based on multiple criteria.
Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications	This guide describes the authentication mechanisms available in AD FS in Windows Server® 2012 R2 Preview and provides step-by-step walkthrough instructions for configuring and verifying the solution of using AD FS to enable multi-factor authentication (MFA) based on the user data.

Active Directory Certificate Services Overview

This document provides an overview of Active Directory Certificate Services (AD CS) in Windows Server® 2012. AD CS is the Server Role that allows you to build a public key infrastructure (PKI) and provide public key cryptography, digital certificates, and digital signature capabilities for your organization.

Did you mean…

- Active Directory Domain Services Overview
- Active Directory Rights Management Services Overview
- Active Directory Federation Services Overview
- Active Directory Lightweight Directory Services Overview

Note

To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Role description

AD CS provides customizable services for issuing and managing digital certificates used in software security systems that employ public key technologies.

The digital certificates that AD CS provides can be used to encrypt and digitally sign electronic documents and messages. These digital certificates can be used for authentication of computer, user, or device accounts on a network. Digital certificates are used to provide:

1. Confidentiality through encryption
2. Integrity through digital signatures
3. Authentication by associating certificate keys with computer, user, or device accounts on a computer network

Practical applications

You can use AD CS to enhance security by binding the identity of a person, device, or service to a corresponding private key. AD CS gives you a cost-effective, efficient, and secure way to manage the distribution and use of certificates.

Applications supported by AD CS include Secure/Multipurpose Internet Mail Extensions (S/MIME), secure wireless networks, virtual private network (VPN), Internet Protocol security (IPsec), Encrypting File System (EFS), smart card logon, Secure Socket Layer/Transport Layer Security (SSL/TLS), and digital signatures.

New and changed functionality

There are multiple changes to AD CS in Windows Server 2012 and the What's New in AD CS article (http://go.microsoft.com/fwlink/?LinkID=224385) describes these changes.

Server Manager information

The installation of AD CS role services can be performed through the Server Manager. The following role services can be installed:

Role service	Description
Certification Authority (CA)	Root and subordinate CAs are used to issue certificates to users, computers, and services, and to manage certificate validity.
Web Enrollment	CA Web enrollment allows users to connect to a CA by means of a Web browser in order to request certificates and retrieve certificate revocation lists (CRLs).
Online Responder	The Online Responder service decodes revocation status requests for specific certificates, evaluates the status of these certificates, and sends back a signed response containing the requested certificate status information.
Network Device Enrollment Service	The Network Device Enrollment Service (NDES) allows routers and other network devices that do not have domain accounts to obtain certificates.
Certificate Enrollment Policy Web Service	The Certificate Enrollment Policy Web Service enables users and computers to obtain certificate enrollment policy information.
Certificate Enrollment Web Service	The Certificate Enrollment Web Service is an Active Directory Certificate Services (AD CS) role service that enables users and computers to perform certificate enrollment by using the HTTPS protocol. When used together, the Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service enable policy-based certificate enrollment for • domain member computers not connected to the domain • computers that are not domain members

See also

The following table provides additional resources for evaluating AD CS.

Content type	References
Product evaluation	• Test Lab Guide: Deploying an AD CS Two Tier PKI Hierarchy • Test Lab Guide: Demonstrating Key-Based Renewal • Test Lab Guide Mini-Module: Cross-Forest Certificate Enrollment using Certificate Enrollment Web Services
Community resources	• Community directory for documentation and information: Windows PKI Documentation Reference and Library • Frequently asked questions (FAQs) list Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ) • Support forum: Windows Server Security Forum • Product team blog: Windows PKI Blog • Support Team Blog: Ask the Directory Services team • Script repository: TechNet Script Center Repository search for Certification, Certificate, or PKI. • Community technology overview: Active Directory Certificate Services (AD CS) Overview
Related technologies	Active Directory Domain Services Active Directory Rights Management Services Active Directory Federation Services Active Directory Lightweight Directory Services

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

What's New in Certificate Services in Windows Server 2012 R2

Active Directory Certificate Services (AD CS) in Windows Server 2012 R2 provides new features and capabilities from previous versions. This document describes new deployment, manageability, and capabilities added to the AD CS role in Windows Server 2012 R2.

Role/Feature description

Active Directory Certificate Services (AD CS) provides customizable services for issuing and managing public key infrastructure (PKI) certificates used in software security systems that employ public key technologies.

New and changed functionality

New functionality in AD CS for Windows Server 2012 R2 includes the following.

Feature/functionality	New or improved	Description
Policy Module support for the Network Device Enrollment Service	New	Using a policy module with the Network Device Enrollment Service provides enhanced security so that users and devices can request certificates from the Internet.
TPM key attestation	New	TPM key attestation lets the certification authority (CA) verify that the private key is protected by a hardware-based TPM.
Windows PowerShell for Certificate Services	New	New Windows PowerShell cmdlets are available for backup and restore.

Policy Module support for the Network Device Enrollment Service

The AD CS role service, Network Device Enrollment Service, is designed for secured networks and trusted administrators. Because of this design, enrollment can use a single password, or even no password, to request multiple certificates. In addition, there is no authentication for the

subject name value supplied. However, Windows Server 2012 R2 supports a policy module for the Network Device Enrollment Service, which provides additional authentication that makes it practical to run this role service in a perimeter network. This configuration supports the Bring Your Own Device (BYOD) scenario, where mobile devices such as those that run iOS and Android, and computers that are not domain members, can now use the Network Device Enrollment Service to request user and computer certificates from the Internet. This is sometimes referred to as over-the-air enrollment.

Windows Server 2012 R2 does not come with a policy module. You must install this separately, from a software vendor that provides a policy module, or write your own policy module. If you install a policy module from a software vendor, typically, this will be from a company that provides management for mobile devices. For example, System Center 2012 R2 Configuration Manager provides a policy module that is required for when you deploy certificate profiles.

For more information, see the following resources:

- Using a Policy Module with the Network Device Enrollment Service
- Certificate Profiles in Configuration Manager

TPM key attestation

TPM key attestation lets the certification authority (CA) verify that the private key is protected by a hardware-based TPM and that the TPM is one that the CA trusts. This functionality prevents the certificate from being exported to an unauthorized device, and can bind the user identity to the device.

All TPMs have an endorsement key that is unique to each TPM. In some cases, TPMs have an endorsement key certificate that chains to the manufacturer's issuing CA. Not all TPMs support attestation but when they do, you can optionally choose to validate the key attestation by using the endorsement key, or by using an endorsement key certificate.

To use TPM key attestation, the client operating system must be Windows 8.1 or Windows Server 2012 R2. To configure TPM key attestation, use a version 4 certificate template with an enterprise CA, and configure the settings on the **Key Attestation** tab. Do not select **Do not store certificate and requests in the CA database** on the **Server** tab of the certificate template properties, because this configuration is not supported with TPM key attestation. In addition, standalone CAs and web enrollment do not support TPM key attestation.

When you configure TPM key attestation, you can choose increasing levels of assurance by specifying how to validate the endorsement key that is burned into the TPM by the manufacturer:

- **User credentials**. No additional configuration is required on the CA.
- **Endorsement certificate**. You must add the root and issuing CA certificates for the TPMs to new certificate stores on the CA. The new certificate stores are EKCA for the intermediate store, and EKRROT for the root store.
- **Endorsement key**. You must add each endorsement key for the TPMs to an approved list (EKPUB list).

Tip

If the settings on the **Key Attestation tab** are not available, verify the following settings:

- On the **Compatibility** tab: The **Certification Authority** is set to **Windows Server 2012 R2**, and the **Certificate recipient** is set to **Windows 8.1 / Windows Server 2012 R2**.
- On the **Request Handling** tab: The **Allow private key to be exported** checkbox and the **Archive subject's encryption private key** checkbox must not be selected.
- On the **Cryptography** tab: The **Provider Category** is set to **Key Storage Provider** and the **Algorithm name** is set to **RSA**. In addition, the **Request must use one of the following providers** must be set to **Microsoft Platform Crypto Provider**.

For more information, see the following resources:

- The **Key attestation** section in the **TPM Fundamentals** topic.
- TPM Key Attestation

Windows PowerShell for Certificate Services

New Windows PowerShell cmdlets are available in Windows Server 2012 R2. You can use these cmdlets to back up and restore a certification authority (CA) database.

Cmdlet name	New or improved	Description
Backup-CARoleService	New	Back up the CA database.
Restore-CARoleService	New	Restore the CA database.

For more information about these cmdlets, see Backup-CARoleService and Restore-CARoleservice.

To use these cmdlets in a migration scenario, see the following sections from Active Directory Certificate Services Migration Guide for Windows Server 2012 R2:

- Backing up a CA database and private key
- Restoring the CA database and configuration on the destination server

What's New in Certificate Services in Windows Server 2012

This document describes new public key infrastructure (PKI) features that available in Windows Server® 2012, Windows Server® 2012 R2, and Windows® 8 and Windows® 8.1.

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Role description

Active Directory Certificate Services (AD CS) provides customizable services for issuing and managing public key infrastructure (PKI) certificates used in software security systems that employ public key technologies. The AD CS server role includes six role services:

- Certification Authority (CA)
- Web Enrollment
- Online Responder
- Network Device Enrollment Service
- Certificate Enrollment Policy Web Service
- Certificate Enrollment Web Service

For an overview of AD CS, see the Active Directory Certificate Services (AD CS).

New and changed functionality

The new and changed functionality in AD CS and PKI includes the following.

- Integration with Server Manager
- Deployment and management capabilities from Windows PowerShell®
- All AD CS role services run on any version
- All AD CS role services can be run on Server Core
- Support for key-based renewal
- Certificate Template Compatibility
- Support for certificate renewal with same key
- Support for Internationalized Domain Names
- Increased security enabled by default on the CA role service
- AD DS Site Awareness for AD CS and PKI Clients
- Group-protected PFX format
- Certificate lifecycle notifications
- CA private keys are included in the System State Backup image

Integration with Server Manager

Server Manager provides a centralized graphical user interface for installing and managing the AD CS server role and its six role services.

What value does this change add?

AD CS server role and its role services are integrated into Server Manager, which allows you to install the AD CS role service from the **Manage** menu using **Add Roles and Features**. Once the server role is added, AD CS appears in the Server Manager dashboard as one of the roles that can be managed. This provides you a central location from which you can deploy and manage AD CS and its role services. Further, the new Server Manager allows you to manage multiple servers from one location and you can see the AD CS role services installed on each server,

review related events, and perform management tasks on each server. For more information on how the new Server Manager works, see Manage multiple, remote servers with Server Manager.

What works differently?

To add the AD CS Server Role, you can use the **Add Roles and Features** link on the **Manage** menu in Server Manager. The AD CS installation flow is similar to that in the previous version, except for the division of the binary installation process and the configuration process. Previously the installation and configuration was a single wizard. In the new installation experience, you first install the binary files and then you can launch the AD CS Configuration wizard to configure the role services that have already had their binary files installed. To remove the AD CS Server Role, you can use the **Remove Roles and Features** link on the **Manage** menu.

Deployment and management capabilities from Windows PowerShell®

All AD CS role services can be configured or have their configurations removed by using the AD CS Deployment Windows PowerShell® cmdlets. These new deployment cmdlets are described in the AD CS Deployment cmdlets Overview topic. The AD CS Administration cmdlet allows you to manage the Certification Authority role service. The new administration cmdlets are described in the AD CS Administration cmdlets Overview topic.

What value does this change add?

You can use Windows PowerShell to script deployments of any AD CS role service as well as the ability to manage the CA role service.

What works differently?

You can use either Server Manager or Windows PowerShell cmdlets to deploy the AD CS role services.

All AD CS role services run on any version

All Windows Server 2012 and Windows Server 2012 R2 versions allow you to install all of the AD CS role services.

What value does this change add?

Unlike previous versions, you can install AD CS roles on any version of Windows Server 2012 or Windows Server 2012 R2.

What works differently?

In Windows Server® 2008 R2 operating system the different role services (previously called components) had different operating system version requirements, as described in Active Directory Certificate Services Overview. In Windows Server 2012 or Windows Server 2012 R2, all six of the roles services work as they would on any Windows Server 2012 or Windows Server 2012 R2 version. The only difference is that you will find AD CS with all six role services available for installation on any version of Windows Server 2012 or Windows Server 2012 R2.

All AD CS role services can be run on Server Core

All six of the Windows Server 2012 and Windows Server 2012 R2 AD CS role services can be installed and run using the Server Core or the Minimal Server Interface installation options.

What value does this change add?

Unlike previous versions, you can now run all AD CS role services on Server Core or the Minimal Server Interface installation options in Windows Server 2012 or Windows Server 2012 R2

What works differently?

You can now easily deploy AD CS role services using Server Manager or Windows PowerShell cmdlets working locally at the computer or remotely over the network. In addition, Windows Server 2012 or Windows Server 2012 R2 provides multiple installation options that even allow you to install with a graphical user interface and later switch to a Server Core or Minimal Server Interface installation. For more information on installation options, see Windows Server Installation Options

Support for key-based renewal

Certificate Enrollment Web Services is a feature that was added in Windows® 7 and Windows Server 2008 R2. This feature allows online certificate requests to come from untrusted Active Directory Domain Services (AD DS) domains or even from computers that are not joined to a domain. AD CS in Windows Server 2012 and Windows Server 2012 R2 build on the Certificate Enrollment Web Services by adding the ability to automatically renew certificates for computers that are part of untrusted AD DS domains or not joined to a domain.

What value does this change add?

Administrators no longer need to manually renew certificates for computers that are members of workgroups or possibly joined to a different AD DS domain or forest.

What works differently?

Certificate Enrollment Web Services continues to function as it did before, but now computers that are outside of the domain can renew their certificates using their existing certificate for authentication.

Additional information, see the topic key-based renewal. There are also two Test Lab Guides that demonstrate the use of this key-based renewal:

1. Test Lab Guide: Demonstrating Certificate Key-Based Renewal
2. Test Lab Guide Mini-Module: Cross-Forest Certificate Enrollment using Certificate Enrollment Web Services

Certificate Template Compatibility

AD CS in Windows Server 2012 and Windows Server 2012 R2 include version 4 certificate templates. These templates have several differences from previous template versions. Version 4 certificate templates:

- support both cryptographic service providers (CSPs) and key service providers (KSPs).

- can be set to require renewal with the same key.
- are only available for use by Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2.
- specify the minimum certification authority and certificate client operating systems that can utilize the template.

To help administrators separate what features are supported by which operating system version, the **Compatibility** tab was added to the certificate template properties tab.

What value does this change add?

The new version 4 certificate templates provide additional capabilities, such as enforcing renewal with the same key (available to only Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 certificate clients). The new **Compatibility** tab allows administrators to set different combinations of operating system versions for the certification authority and certificate clients and see only the settings that will work with those client versions.

What works differently?

The **Compatibility** tab appears in the Certificate Template properties user interface. This tab allows you to select the minimum certification authority and minimum certificate client operating system versions. The **Compatibility** tab configuration does a couple of things:

- It marks options as unavailable in the certificate template properties depending upon the selected operating system versions of certificate client and certification authority.
- For version 4 templates, it determines which operating system versions are able to use the template.

Clients prior to Windows 8 and Windows Server 2012 will not be able to take advantage of the new version 4 templates.

Note

There is a statement on the **Compatibility** tab that reads **These settings may not prevent earlier operating systems from using this template**. This statement means that compatibility settings have no restrictive effect on version 1, version 2, or version 3 templates and enrollment may proceed as before. For example, in **Compatibility** tab, if the minimum client operating system version is set to Windows® Vista on a version 2 template, a Windows® XP certificate client may still enroll for a certificate using the version 2 template.

For more information on these changes, see Certificate Template Versions and Options

Support for certificate renewal with same key

AD CS in Windows Server 2012 and Windows Server 2012 allow for a certificate to be configured so that it will be renewed with the same key. This allows the same assurance level of the original key to be maintained throughout its lifecycle. Windows Server 2012 and Windows Server 2012 supports generating Trusted Platform Module (TPM)-protected keys using TPM-based key storage providers (KSPs). The benefit of using TPM-based KSP is true non-exportability of keys backed up by the anti-hammering mechanism of TPMs. Administrators can configure certificate

templates so that Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 to give higher priority to TPM-based KSPs for generating keys. Also, using renewal with the same key, administrators can remain assured that the key still remains on TPM after renewal.

Note

Entering the personal identification number (PIN) incorrectly too many times activates the anti-hammering logic of the TPM. Anti-hammering logic is software or hardware methods that increase the difficulty and cost of a brute force attack on a PIN by not accepting PIN entries until after a certain amount of time has passed.

What value does this change add?

This feature allows an administrator to enforce renewal with the same key, which can reduce administrative costs (when keys are renewed automatically) and increase key security (when keys are stored using TPM-based KSPs).

What works differently?

Clients that receive certificates from templates that are configured for renewal with the same key must renew their certificates using the same key, or renewal will fail. Also, this option is available only for Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 certificate clients.

Notes

- For more information, see Certificate Renewal with the Same Key.
- If **Renew with the same key** is enabled on a certificate template and later key archival (**Archive subject's encryption private key**) is also enabled, some renewed certificates may not be archived. To learn more about this situation and how to mitigate it, see Key Archival and renew with the same key.

  If **Renew with the same key** is enabled on a certificate template and later key archival (**Archive subject's encryption private key**) is also enabled, renewed certificates will not be archived. To learn more about this situation and mitigation for it, see Key Archival and renew with the same key.

Support for Internationalized Domain Names

Internationalized names are names that contain characters that cannot be represented in ASCII. AD CS in Windows Server 2012 and Windows Server 2012 R2 supports Internationalized Domain Names (IDNs) in several scenarios.

What value does this change add?

The following IDN scenarios are now supported

- Certificate enrollment for computers using IDNs
- Generating and submitting a certificate request with an IDN using the certreq.exe command line tool
- Publishing Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) publishing to servers using IDNs

- The **Certificate** user interface supports IDNs
- The Certificate MMC snap-in also allows for IDNs in **Certificate Properties**

What works differently?

There is limited support for IDNs as previously described.

Increased security enabled by default on the CA role service

When a certificate request is received by a certification authority (CA), encryption for the request can be enforced by the CA via the RPC_C_AUTHN_LEVEL_PKT, as described in MSDN article Authentication-Level Constants (http://msdn.microsoft.com/library/aa373553.aspx). On Windows Server 2008 R2 and earlier versions, this setting is not enabled by default on the CA. On a Windows Server 2012 or Windows Server 2012 R2 CA, this enhanced security setting is enabled by default.

What value does this change add?

The CA enforces enhanced security in the requests that are sent to it. This higher security level requires that the packets requesting a certificate are encrypted, so they cannot be intercepted and read. Without this setting enabled, anyone with access to the network can read packets sent to and from the CA using a network analyzer. This means that information could be exposed that might be considered a privacy violation, such as the names of requesting users or machines, the types of certificates for which they are enrolling, the public keys involved, and so on. Within a forest or domain, leaking these data may not be a concern for most organizations. However, if attackers gain access to the network traffic, internal company structure and activity could be gleaned, which could be used for more targeted social engineering or phishing attacks.

The commands to enable the enhanced security level of RPC_C_AUTHN_LEVEL_PKT on Windows Server® 2003, Windows Server® 2003 R2, Windows Server® 2008, or Windows Server 2008 R2 certification authorities are:

certutil -setreg CA\InterfaceFlags +IF_ENFORCEENCRYPTICERTREQUEST
Restart the certification authority
net stop certsvc
net start certsvc

If you still have Windows XP client computers that need to request certificates from a CA that has the setting enabled, you have two options:

1. Upgrade the Windows XP clients to a newer operating system.

To lower CA security for compatibility with Windows XP clients

1. **certutil -setreg CA\InterfaceFlags -IF_ENFORCEENCRYPTICERTREQUEST**
2. **net stop certsvc**
3. **net start certsvc**

What works differently?

Windows XP clients will not be compatible with this higher security setting enabled by default on a Windows Server 2012 or Windows Server 2012 R2 CA. If necessary, you can lower the security setting as previously described.

AD DS Site Awareness for AD CS and PKI Clients

Certificate services in Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 can be configured to utilize Active Directory Domain Services (AD DS) sites to help optimize certificate services client requests. This functionality is not enabled by default on either certification authority (CA) or the public key infrastructure (PKI) client computers.

Note
For information on enabling AD DS Site Awareness, see TechNet Wiki article AD DS Site Awareness for AD CS and PKI Clients.

What value does this change add?

This change enables Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 certificate clients to locate a CA in their local AD DS site.

What works differently?

When enrolling for a template-based certificate, the client queries AD DS for the template and the CA objects. The client then uses a DsGetSiteName function call to get its own site name. For CAs with the msPKI-Site-Name attribute already set, the certificate services client determine the AD DS site link cost from the client site to each target CA site. A DsQuerySitesByCost function call is used to make this determination. The certificate services client uses the returned site costs to prioritize the CAs that allow the client the Enroll permission and support the relevant certificate template. The higher cost CAs are tried to be contacted last (only if former CAs are unavailable).

Note
A CA may return no site cost if the msPKI-Site-Name attribute is not set on the CA. If no site cost is available for an individual CA, then the highest possible cost is assigned to that CA.

Group-protected PFX format

Previously, a PKCS#12 standard (also known as PFX) format was only protected by a password that had the following limitations:

- Difficult to automate
- Not very secure, because usually an administrator used a weak password
- Difficult to share among multiple users

Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 can protect certificates and associated private keys by combining an existing PFX format with a new data protection feature. This allows encrypting the contents of the PFX file with a key that belongs to a group or to an individual, instead of protecting it with a password.

Notes

- To implement this feature, at least one domain controller must be running Windows Server 2012 or Windows Server 2012 R2.
- For more information, see the TechNet Wiki article Certificate PFX Export and Import using AD DS Account Protection.

What value does this change add?

By using this feature, administrators will be able to:

- Deploy, manage, and troubleshoot certificates remotely and across server farms by using Windows PowerShell.
- Share certificates and keys securely across server farms running Windows Server 2012 or Windows Server 2012 R2by using Windows APIs.

Earlier versions of Windows can consume this PFX because internally the operating system assigns a strong random password. The password is included in the PFX, and it is protected by a set of security identifiers (SIDs) with data protection APIs. Any user that has access to the PFX can see that password and share it with previous Windows versions.

What works differently?

A PFX file can now be protected to a security principal instead of just a password. The user interface for certificate export has been updated to allow for the selection of a security principal during export.

Certificate lifecycle notifications

In Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2, certificates provide life cycle notifications in **MY** store from the certificate enrollment API and Windows PowerShell levels. The notifications include expiration, deletion, new, renewal, replacement, close to expiration, archive, and export. Administrators and developers can manage (view, install, copy, request, and delete) certificates and their associated private keys remotely by using Windows PowerShell. This feature allows a script or an executable to launch in response to a certificate lifecycle notification.

Notes

- The expiration notification is supported by stores in addition to **MY** store.
- For more information, see the TechNet Wiki article Certificate Services Lifecycle Notifications.

What value does this change add?

For an application and server-workload developers who use certificates in their product, integrating with the certificate life cycle in Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 is easy and reliable, and it can be done remotely. Developers can develop applications that reconfigure themselves any time a certificate is renewed or replaced with another certificate—by autoenrollment or by a manual or scripted action by an administrator. The investment needed to integrate with the certificate management interfaces is very small.

For an administrator who manages applications that use certificates, Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 certificates are used by those applications

automatically. This occurs because applications integrate with Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 certificate notifications or when the administrator's script is triggered by a certificate event.

What works differently?

Notifications can now be enabled to alert system administrators before certificates expire.

CA private keys are included in the System State Backup image

Windows Server Backup feature can be installed on the certification authority (CA) to create a System State Backup that includes the CA private keys.

What works differently?

In Windows Server 2012 and Windows Server 2012 R2 the System State Backup feature automatically backs up the CAs private key when an administrator or backup operator uses Windows Server Backup feature to perform a System State Backup.

What works differently?

The Windows Server Backup feature now includes the CA private keys.

Tip

- To add this functionality to Windows Server 2008 R2 or Windows Server® 2008, apply the appropriate update listed in document 2603469 in the Microsoft Knowledge Base.
- For details on using this feature, see Windows Server 2012 Active Directory Certificate Services System State Backup and Restore

See also

The following table provides additional resources for evaluating AD CS.

Content type	References
Product evaluation	• Test Lab Guide: Deploying an AD CS Two Tier PKI Hierarchy • Test Lab Guide: Demonstrating Key-Based Renewal • Test Lab Guide Mini-Module: Cross-Forest Certificate Enrollment using Certificate Enrollment Web Services
Community resources	• Community directory for documentation and information: Windows PKI Documentation Reference and Library • Frequently asked questions (FAQs) list Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)

Content type	References
	• Support forum: Windows Server Security Forum • Product team blog: Windows PKI Blog • Support Team Blog: Ask the Directory Services team • Script repository: TechNet Script Center Repository search for Certification, Certificate, or PKI. • Community technology overview: Active Directory Certificate Services (AD CS) Overview
Related technologies	Active Directory Domain Services Active Directory Rights Management Services Active Directory Federation Services Active Directory Lightweight Directory Services

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Protecting Against Weak Cryptographic Algorithms

A software update is available for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 that allows deprecation of weak cryptographic algorithms. In order to use automatic updates from Microsoft as a means to be better protected from weak cryptographic algorithms, this software update must be downloaded and installed on computers that run the aforementioned operating systems.

This software update is built-in to the Windows 8.1 and Windows Server 2012 R2 operating systems.

In this topic:

- What does this software update do?
- How to configure policies for blocking cryptographic algorithms
- Updating client registry settings through Group Policy
- Examples

What does this software update do?

This software update provides an administrator with greater control over the way RSA keys, hash algorithms, and non-RSA asymmetric key algorithms are blocked. This software update allows an administrator to:

- Define policies to selectively block cryptographic algorithms that override settings provided by the operating system.
- Opt-in or opt-out of each policy independently.
- Enable logging per policy (independent of other policies). Logging is off by default.
- Specify a location to which blocked certificates are copied.
- Set policies per algorithm and define hash algorithm policies and asymmetric algorithm policies as described in the following table:

Hash algorithm policies	Asymmetric algorithm policies

• Define the name of the hash algorithm, such as MD5 or SHA1.
• Specify if the policy applies to certificates that chain to third-party root CAs, which excludes the enterprise certificates, or to apply the policy to all certificates.
• Specify a time before which the policy check is disabled (only applicable to code signing certificates used in time-stamped signed binaries).
• Define the type of certificates to which the policy applies, such as:
  • All certificates.
  Note
  If the policy is enabled for all certificates, then an administrator cannot allow weak cryptographic algorithms for a specific Enhanced Key Usage (EKU), such as Server Authentication EKU or Code signing EKU.
  • Certificates that have the Server Authentication EKU.
  • Certificates that have the Code signing EKU.
	• Define the algorithm name and minimum key size, such as RSA, DSA, and ECDSA.
	• Specify if the policy applies to certificates that chain to third-party root CAs, which excludes the enterprise certificates, or if the policy applies to all certificates.
	• Specify a time before which the policy check is disabled (for time-stamped files).
	• Define the type of certificates to which the policy applies, such as:
	  • All certificates.
	  Note
	  If the policy is enabled for all certificates then an administrator cannot allow weak cryptographic algorithms for a specific EKU, such as Server Authentication EKU or Code signing EKU.
	  • Certificates that have the Server Authentication EKU.
	  • Certificates that have the Code signing EKU.
	  • Certificates that have the Time Stamping EKU.

• Certificates that have the Time Stamping EKU. • Specify whether the policy applies only to signing certificates of binaries downloaded from web or to all binaries. • Specify code signing and time stamping certificates that are not blocked (such as those used in legacy code signing or time-stamping situations); the certificates are identified by using their SHA2 thumbprint.	• Specify whether the policy applies only to signing certificates of binaries downloaded from web or to all binaries. • Specify code signing and time stamping certificates that are not blocked (such as those used in legacy code signing or time-stamping situations); the certificates are identified by using their SHA2 thumbprint.

How to configure policies for blocking cryptographic algorithms

The administrator sets cryptographic algorithm blocking policy in the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography\OID\EncodingType 0\CertDllCreateCertificateChainEngine\Config

Use the following commands to display, configure, and remove the cryptographic algorithm blocking policy settings:

- certutil -getreg chain
- certutil -setreg chain
- certutil -delreg chain

The registry entries use the following syntax:

`Weak<CryptoAlg><ConfigType><ValueType>`

The following table lists registry values that can be set to modify default settings for cryptographic algorithm blocking policy. <*CryptoAlg*>, <*ConfigType*> and <*ValueType*> will be present in all names:

Configuration Option	Possible values
CryptoAlg	Md5 Sha1 Rsa Dsa Ecdsa
ConfigType	ThirdParty: Policy is only applied to certificates

	under third-party roots All: Policy is applied to all certs, even under enterprise roots
ValueType	Flags: REG_DWORD data type that can be set to disable the hash algorithm or enable a minimum key length. For more information, see the following table. MinBitLength: REG_DWORD data type that specifies the minimum public key length in bits. Note: MinBitLengh is only applicable to key algorithms policy. AfterTime: REG_BINARY data type that contains an 8 byte FILETIME. The weak crypto algorithm check is disabled for time stamped files before this time. This configuration value isn't applicable to timestamp chains. Sha256Allow: REG_SZ or REG_MULTI_SZ data type that contains the list of certificate SHA256 thumbprints (ASCII_HEX formatted) identifying weak certificates to be explicitly allowed. Non ASCII_HEX characters in the string are skipped, which allows embedded spaces.

The following REG_DWORD values can be set as flags in Weak<*CryptoAlg*><*ConfigType*>Flags:

Flag	Notes
CERT_CHAIN_ENABLE_WEAK_SETTINGS_FLAG (0x80000000)	If this flag is not set, then all other flags and registry values are ignored for this Weak<CryptoAlg><ConfigType>. If the administrator sets this flag for its Weak<CryptoAlg><ConfigType>, then, the corresponding settings provided for the operating system are ignored. If this flag is set in Weak<CryptoAlg>AllFlags: • Resultant "Weak"<CryptoAlg>"ThirdPartyFlags" will or with "Weak"<CryptoAlg>"AllFlags". However, "Weak"<CryptoAlg>"ThirdPartyFlags" logging flags won't be updated. *ThirdPartyFlags* |=

	AllFlags & ~(CERT_CHAIN_ENABLE_WEAK_LOGGING_ FLAG | CERT_CHAIN_ENABLE_ONLY_WEAK_LOG GING_FLAG); • Resultant "Weak<CryptoAlg>"ThirdPartyAfterTime" will be earliest ("Weak"<CryptoAlg>"AllAfterTime", "Weak"<CryptoAlg>"ThirdPartyAfterTime"). Note: Only applicable if "Weak"<CryptoAlg>"AllAfterTime" is defined and nonzero. • Resultant "Weak"<KeyCryptoAlg>"ThirdPartyMinBitLen gth" will be largest ("Weak"<KeyCryptoAlg>"AllMinBitLength", "Weak"<KeyCryptoAlg>"ThirdPartyMinBitLen gth"
CERT_CHAIN_ENABLE_WEAK_LOGGING _FLAG (0x00000004)	This flag is set to enable the logging of weak certificates to the directory identified by CERT_CHAIN_WEAK_SIGNATURE_LOG_DIR_ VALUE_NAME.
CERT_CHAIN_ENABLE_ONLY_WEAK_LO GGING_FLAG (0x00000008)	This flag is set to only log weak certificates to the directory identified by CERT_CHAIN_WEAK_SIGNATURE_LOG_DIR_ VALUE_NAME. Weak signature errors are not returned.

In addition to setting the CERT_CHAIN_ENABLE_WEAK_SETTINGS_FLAG flag described in the preceding table, the following flags corresponding to the EKU must be set to disable weak signature or enable weak hash hygiene checks:

Flag	Notes
CERT_CHAIN_DISABLE_ALL_EKU_WEAK_FLAG (0x00010000)	Disables the algorithm corresponding to that policy for all EKUs.
CERT_CHAIN_DISABLE_SERVER_AUTH_WEAK_FLAG (0x00100000)	Disables the algorithm corresponding to that policy for ServerAuth EKUs.
CERT_CHAIN_DISABLE_CODE_SIGNING_WEAK_FLAG	Disables the algorithm

(0x00400000)	corresponding to that policy for code signing EKUs.
CERT_CHAIN_DISABLE_MOTW_CODE_SIGNING_WEAK_FLAG (0x00800000)	Disables the algorithm corresponding to that policy for code signing EKUs only when the binary is downloaded from the web.
CERT_CHAIN_DISABLE_TIMESTAMP_WEAK_FLAG (0x04000000)	Disables the algorithm corresponding to that policy for timestamp EKUs.
CERT_CHAIN_DISABLE_MOTW_TIMESTAMP_WEAK_FLAG (0x08000000)	Disables the algorithm corresponding to that policy for timestamp EKUs only when the binary is downloaded from the web.

Updating client registry settings through Group Policy

The following procedure shows how to configure the registry settings on all domain-joined machines using GPUpdate. For more details, see Configure a Registry Item.

To update client registry settings

1. On a domain controller, open Group Policy Management Editor.
   a. Open MMC, click **File**, click **Add/Remove Snap ins**, and select **Group Policy Management Editor**.

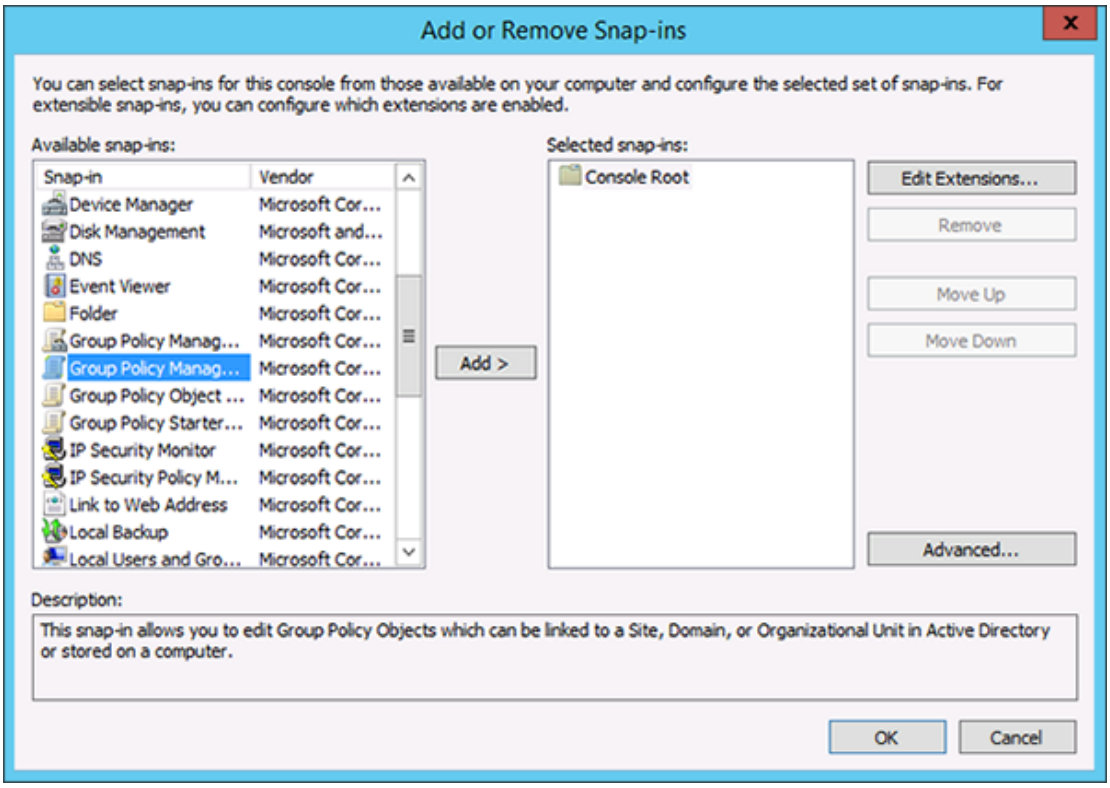

b. Click **Add** to start the Group Policy Wizard.

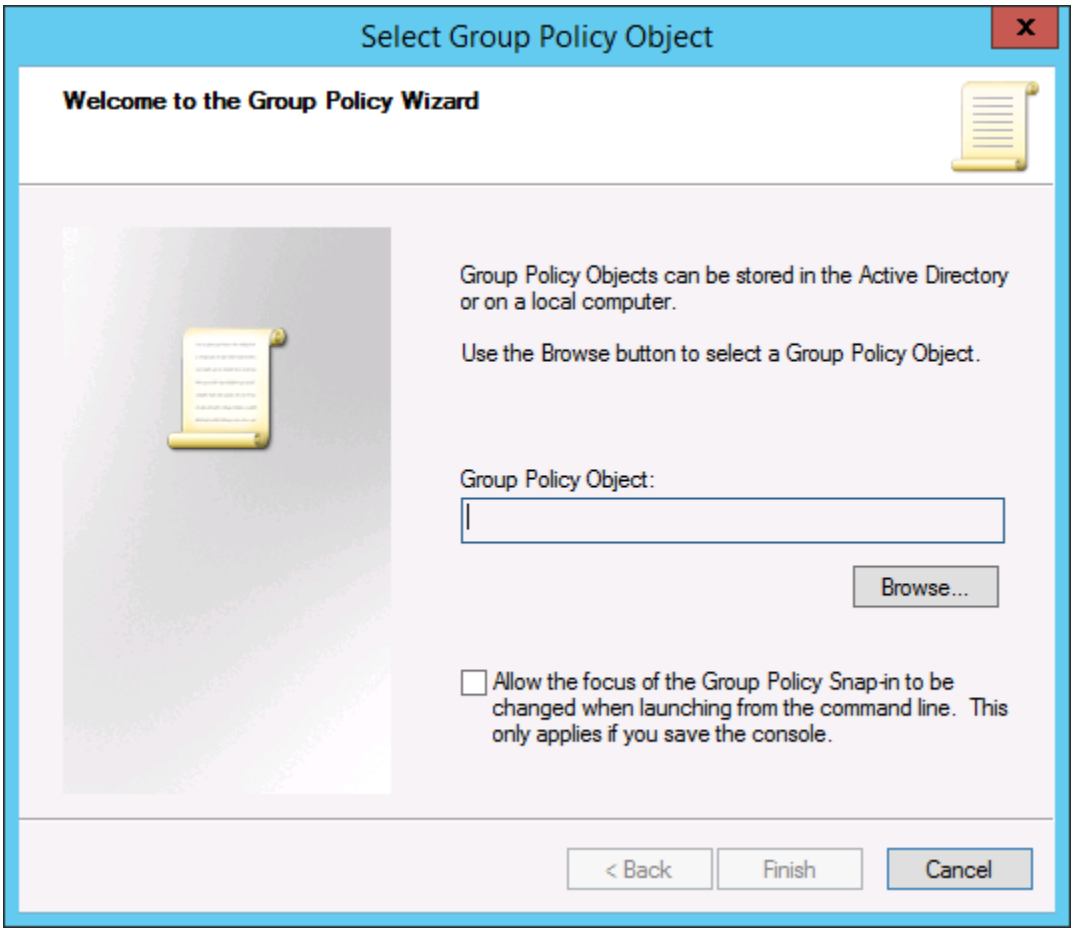

c. Click **Browse**, click **Default Domain Policy**, click **OK**, and click **Finish**.

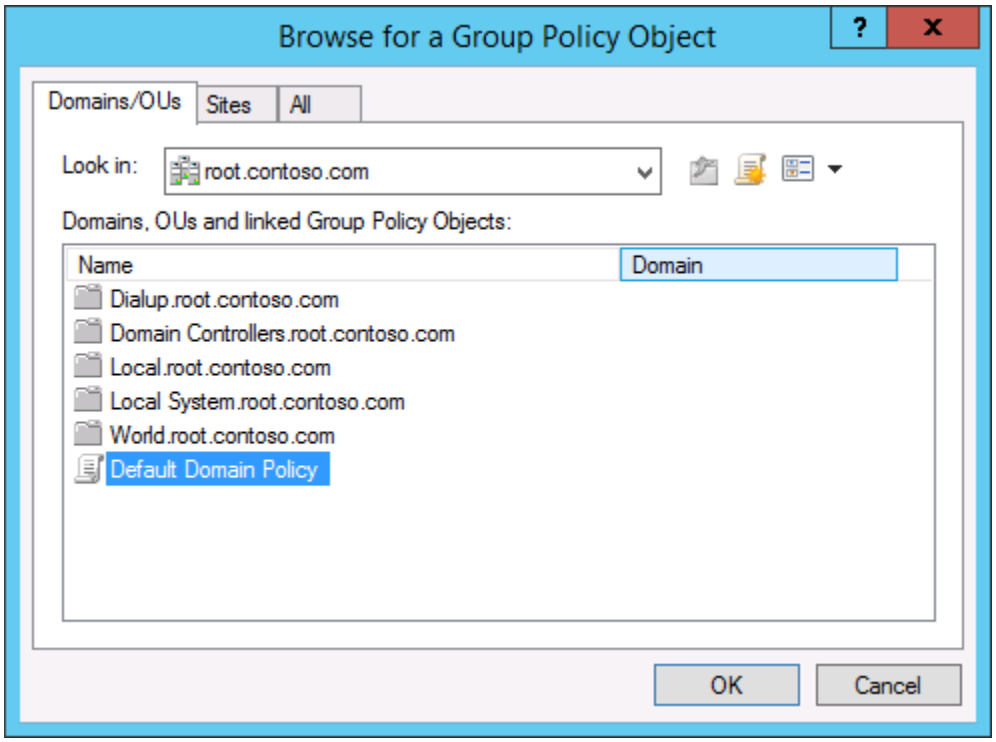

2. Expand Default Domain Policy|Computer Configuration|Preferences|Windows Settings|Registry.

   d. .a. Right-click **Registry**, click **New**, and click **Registry Item**.

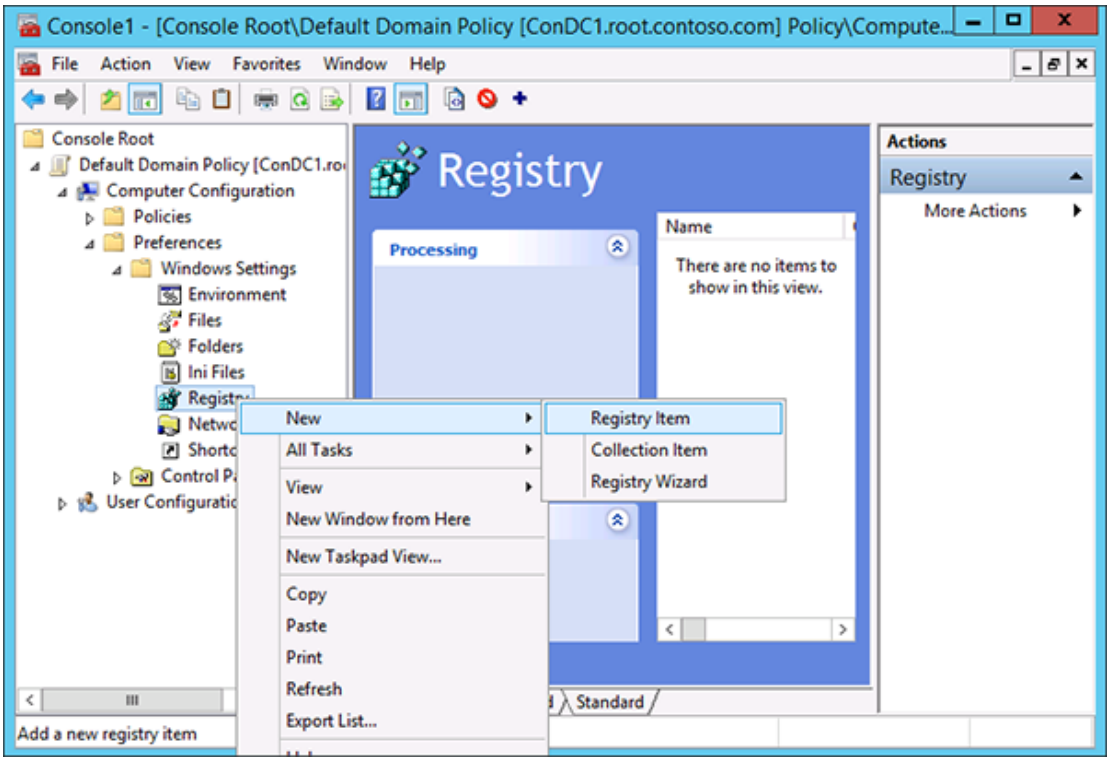

b. In the New Registry Properties, click **Browse** to select the required Key path.

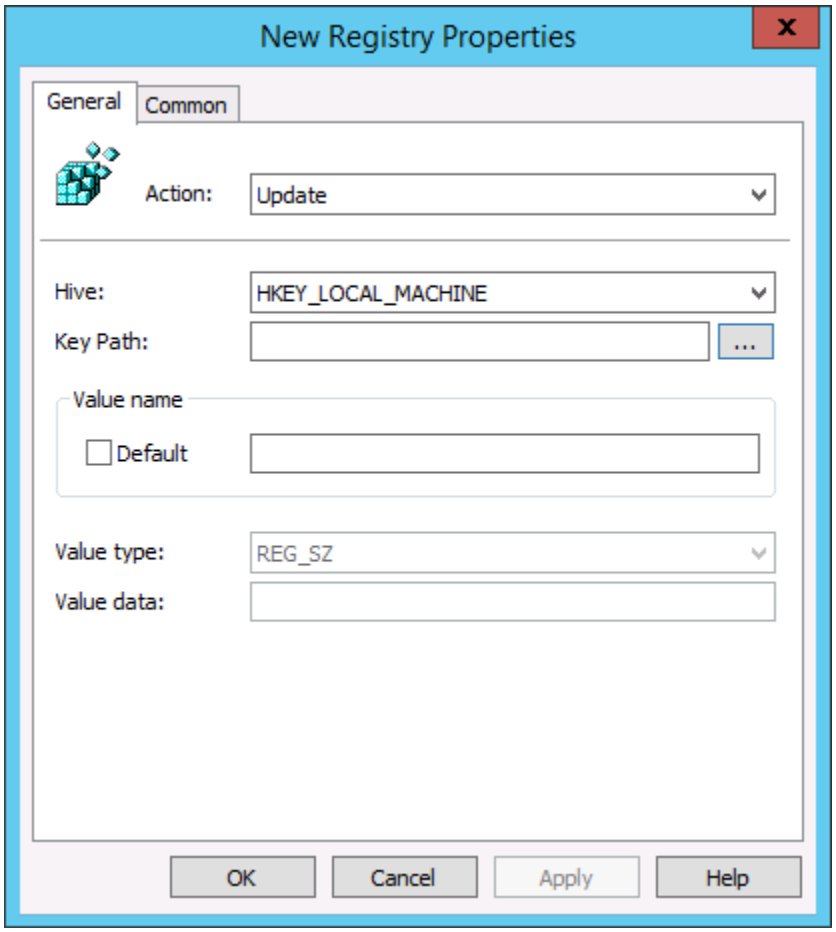

c. Select the registry path. For example:

   SOFTWARE\Microsoft\Cryptography\OID\EncodingType 0\CertDllCreateCertificateChainEngine\Config

d. If the registry value name is already present select the registry value too, for example, **WeakMD5ThirdPartyAfterTime**.

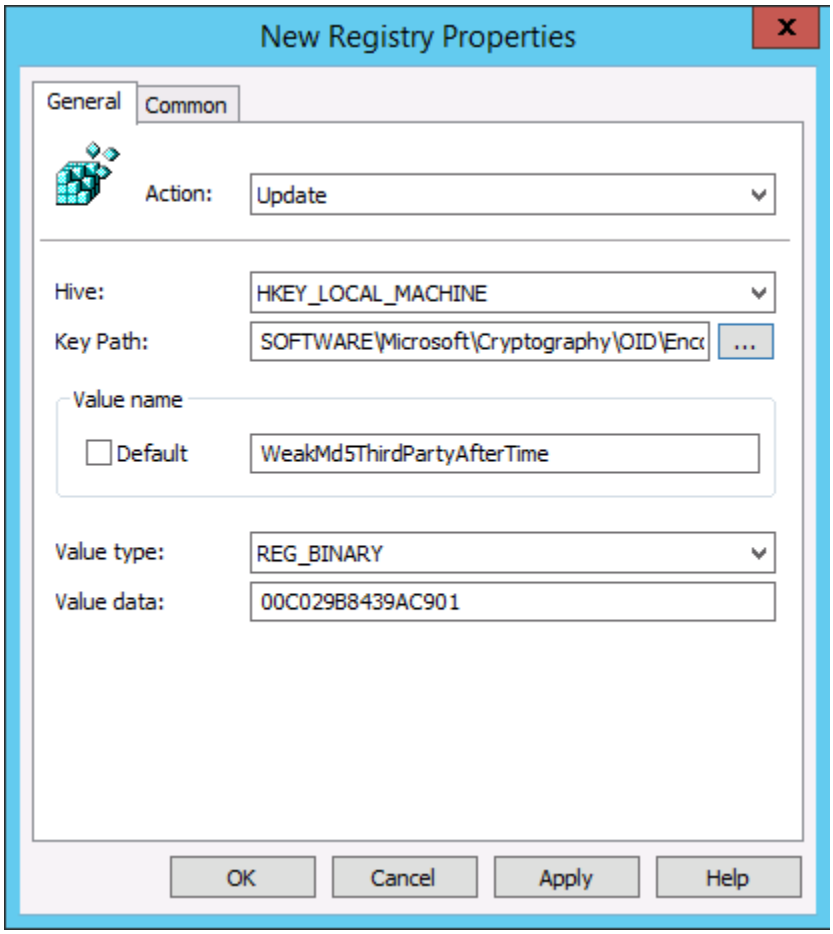

e. If the registry entry is new, select the key location, enter the registry value name, select the appropriate value type, and enter required data.

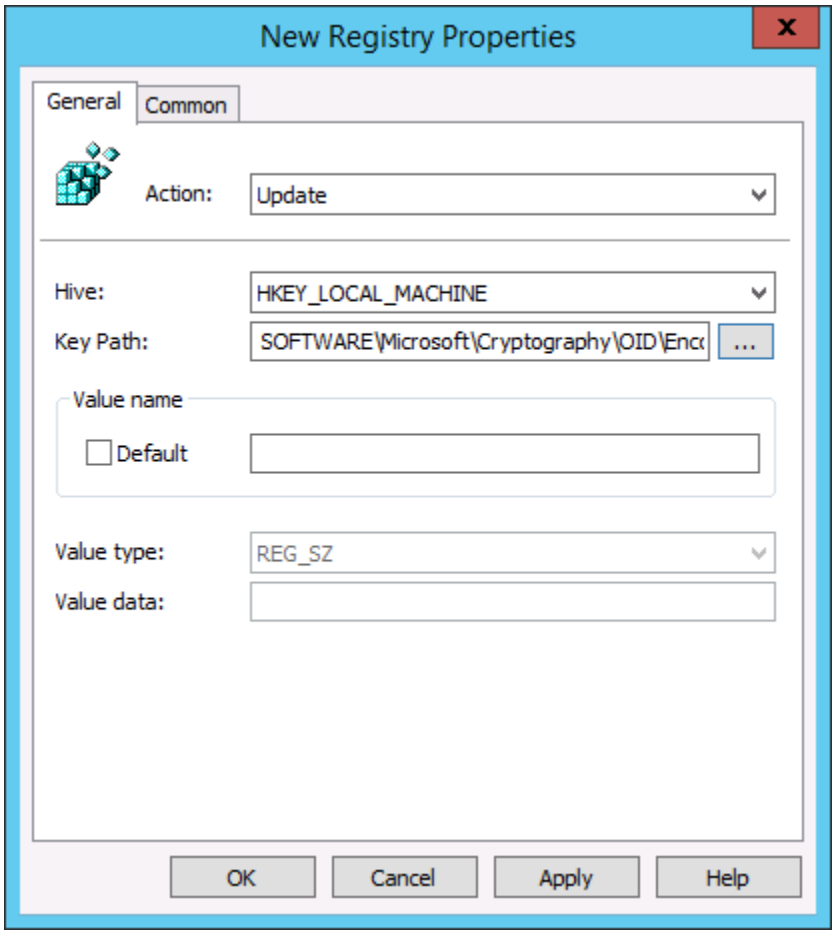

3. Click **Apply** and **OK**. If necessary, run gpupdate /force on domain-joined computers to have the policy setting change be applied immediately.
4. For binary values like **AfterTime**, it is recommended to first apply the value using CertUtil command or in a test computer and then export the values and import them to a domain controller.

   For example, if an administrator needs to apply **WeakMD5ThirdPartyAfterTime**, which is of type REG_BINARY to a date such as 1/1/2010, the administrator can execute the following CertUtil command on a domain controller. The command updates the registry with the correct binary value. After updating the registry, follow the previous steps to apply to the same value to domain-joined computers using Group Policy.

   ```
   Certutil -setreg chain\WeakMD5ThirdPartyAfterTime  @1/1/2010
   ```

Examples

The following example disables MD5 for all SSL server auth certs under third-party root CAs, but allows signed binaries before March 1, 2009 to be accepted. Not applicable to other EKUs. Logging is also enabled, and the message “Setting logging directory section” will appear.

```
Certutil –setreg chain\Default\WeakMd5ThirdPartyFlags 0x80100004

Certutil –setreg chain\Default\WeakMd5ThirdPartyAfterTime @03/01/2009
```

The following example disables RSA 1024 for all timestamp certs under third-party root CAs, but allows signed binaries before March, 2013 to be accepted. Not applicable to other EKUs.

```
Certutil –setreg chain\Default\WeakRSAThirdPartyFlags 0x84000000

Certutil –setreg chain\Default\WeakRSAThirdPartyMinBitLength 1024

Certutil –setreg chain\Default\WeakRSAThirdPartyAfterTime @3/1/2013
```

Setting logging directory and enabling logging

Weak Crypto framework provides a mechanism where administrators can set a log directory for all the certs which are considered weak according to settings.

To enable logging, an administrator can set a logging directory either by adding a registry entry or executing certutil command as follows (c:\Log should be created and with correct permissions) along with weak crypto settings:

```
Certutil –setreg chain\WeakSignatureLogDir c:\log
```

Or updating registry directly:

HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType 0\CertDllCreateCertificateChainEngine\Config\WeakSignatureLogDir

And update the flags to enable logging CERT_CHAIN_ENABLE_WEAK_LOGGING_FLAG (0x00000004), as in the following example, where weak crypto is applied for all EKUs and the weak MD5 third-party certificates are logged to c:\log.

```
Certutil –setreg chain\WeakMD5ThirdPartyFlags  0x80010004
```

Logging with Audit only Mode

Weak crypto framework also provides a feature for an administrator to have logging only without any chain building errors returned. To set this, users need to include CERT_CHAIN_ENABLE_ONLY_WEAK_LOGGING_FLAG (0x00000008). For example, to enable audit only mode for all EKU weak MD5 Third-party certs:

```
Certutil –setreg chain\WeakSignatureLogDir c:\log
```

and

```
Certutil –setreg chain\WeakMD5ThirdPartyFlags  0x80010008
```

Certification Authority Guidance

A certification authority (CA) is responsible for attesting to the identity of users, computers, and organizations. The CA authenticates an entity and vouches for that identity by issuing a digitally signed certificate. The CA can also manage, revoke, and renew certificates.

A certification authority can refer to following:

- An organization that vouches for the identity of an end user
- A server that is used by the organization to issue and manage certificates

By installing the Certification Authority role service of Active Directory Certificate Services (AD CS), you can configure your Windows server to act as a CA.

Before you install the CA role service, you should:

1. Plan a public key infrastructure (PKI) that is appropriate for your organization.
2. Install and configure a Hardware Security Module (HSM) according to the HSM vendor instructions, if you are planning to use one.
3. Create an appropriate CAPolicy.inf, if you want to modify the default installation settings.

Note

To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Plan for PKI

To ensure that your organization can take full advantage of your Active Directory Certificate Services (AD CS) installation, you must plan the PKI deployment appropriately. You should determine how many CAs you will install and in what configuration before you install any CA. Creating an appropriate PKI design can be time consuming, but it is important for the success of your PKI.

For more information and resources, see PKI Design Guidance in Microsoft TechNet.

Use an HSM

Using a hardware security module (HSM) can enhance the security of the CA and the PKI.

An HSM is a dedicated hardware device that is managed separately from the operating system. These modules provide a secure hardware store for CA keys, in addition to a dedicated cryptographic processor to accelerate signing and encrypting operations. The operating system utilizes the HSM through the CryptoAPI interfaces, and the HSM functions as a cryptographic service provider (CSP) device.

HSMs typically are PCI adapters, but they are also available as network-based appliances, serial devices, and USB devices. If an organization plans to implement two or more CAs, you can install a single network-based HSM and share it among multiple CAs.

To set up a CA by using an HSM, the HSM must be installed and configured before you set up any CAs with keys that will be stored on the HSM.

Consider a CAPolicy.inf file

The CAPolicy.inf file is not required to install AD CS, but it can be used to customize the settings of the CA. The CAPolicy.inf file contains various settings that are used when installing a CA or when renewing the CA certificate. The CAPolicy.inf file must be created and stored in the %systemroot% directory (typically C:\Windows) for it to be used.

The settings that you include in the CAPolicy.inf file depend largely on the deployment type that you want to create. For example, a root CA might have a CAPolicy.inf file that looks like this:

```
[Version]

Signature= "$Windows NT$"

[Certsrv_Server]

RenewalKeyLength=4096

RenewalValidityPeriod=Years

RenewalValidityPeriodUnits=20

LoadDefaultTemplates=0
```

Whereas a CAPolicy.inf file for an enterprise that is issuing a CA might look like this:

```
[Version]

Signature= "$Windows NT$"

[PolicyStatementExtension]

Policies = LegalPolicy, LimitedUsePolicy

[LegalPolicy]

OID = 1.1.1.1.1.1.1.1.1

URL = "http://www.contoso.com/pki/Policy/USLegalPolicy.asp"

URL = "ftp://ftp.contoso.com/pki/Policy/USLegalPolicy.txt"

[LimitedUsePolicy]

OID = 2.2.2.2.2.2.2.2.2

URL = "http://www.contoso.com/pki/Policy/USLimitedUsePolicy.asp"

URL = "ftp://ftp.contoso.com/pki/Policy/USLimitedUsePolicy.txt"

LoadDefaultTemplates=0
```

Notes

1. The OIDs shown in the example CAPolicy.inf are examples only. Individual organizations should obtain their own OIDs. For more information about OIDs, see Obtaining a Root OID from an ISO Name Registration Authority.
2. For more information, see CAPolicy.inf Syntax.

Select CA configuration settings

The following sections describe the configuration options that you will select after installing the CA binary installation files.

Select setup type

Enterprise CAs are integrated with Active Directory Domain Services (AD DS). They publish certificates and certificate revocation lists (CRLs) to AD DS. Enterprise CAs use information that is stored in AD DS, including user accounts and security groups, to approve or deny certificate requests. Enterprise CAs use certificate templates. When a certificate is issued, the enterprise CA uses information in the certificate template to generate a certificate with the appropriate attributes for that certificate type.

If you want to enable automated certificate approval and automatic user certificate enrollment, use enterprise CAs to issue certificates. These features are available only when the CA infrastructure is integrated with Active Directory. Additionally, only enterprise CAs can issue certificates that enable smart card sign-in, because this process requires that smart card certificates are mapped automatically to the user accounts in Active Directory.

Notes

Members of the Enterprise Admins group have the appropriate permissions to install an enterprise CA. A local administrator who has been delegated Full Control permissions to the following container in the Configuration directory partition in Active Directory can also install an enterprise CA:

CN=Public Key Services,CN=Services,CN=Configuration,DC=<domainNC>

However, a member of Enterprise Admins or Domain Admins with Write permissions to the CN=Public Key Services location must run `certutil -installdefaulttemplates`. Further, the CA computer account must be added to the **Cert Publishers** and **Pre-Windows 2000 Compatible Access** groups to complete the configuration.

Stand-alone CAs do not require AD DS, and they do not use certificate templates. If you use stand-alone CAs, all information about the requested certificate type must be included in the certificate request. By default, all certificate requests that are submitted to stand-alone CAs are held in a pending queue until a CA administrator approves them. You can configure stand-alone CAs to issue certificates automatically upon request, but this is less secure, and it is usually not recommended because the requests are not authenticated.

From a performance perspective, using stand-alone CAs with automatic issuance enables you to issue certificates at a faster rate than you can by using enterprise CAs. However, unless you are using automatic issuance, using stand-alone CAs to issue large volumes of certificates usually

comes at a high administrative cost because an administrator must manually review and then approve or deny each certificate request. For this reason, stand-alone CAs are best used with public key security applications on extranets and on the Internet, when users do not have user accounts and when the volume of certificates to be issued and managed is relatively low

You must use stand-alone CAs to issue certificates when you are using a non-Microsoft directory service or when AD DS is not available. You can use both enterprise and stand-alone certification authorities in your organization, as explained in the following table.

Option	Enterprise CA	Standalone CA
Publish certificates in Active Directory and use Active Directory to validate certificate requests.	Yes	No
Take the CA offline.	Not recommended	Yes
Configure the CA to issue certificates automatically.	Yes	Not recommended
Allow administrators to approve certificate requests manually.	Yes	Yes
Allow for the use of certificate templates.	Yes	No
Authenticate requests to Active Directory.	Yes	No

Choose CA type

Enterprise and stand-alone CAs can be configured as root CAs or as subordinate CAs. Subordinate CAs can further be configured as intermediate CAs (also referred to as a policy CA) or issuing CAs

Designate a root CA

A root CA is the CA that is at the top of a certification hierarchy. It must be trusted unconditionally by clients in your organization. All certificate chains terminate at a root CA. Whether you use enterprise or stand-alone CAs, you need to designate a root CA.

Since the root CA is the top CA in the certification hierarchy, the Subject field of the certificate that is issued by a root CA has the same value as the Issuer field of the certificate. Likewise, because the certificate chain terminates when it reaches a self-signed CA, all self-signed CAs are root CAs. The decision to designate a CA as a trusted root CA can be made at the enterprise level or locally by the individual IT administrator.

A root CA serves as the foundation upon which you base your certification authority trust model. It guarantees that the subject's public key corresponds to the identity information shown in the subject field of the certificates it issues. Different CAs might also verify this relationship by using different standards; therefore, it is important to understand the policies and procedures of the root certification authority before choosing to trust that authority to verify public keys.

The root CA is the most important CA in your hierarchy. If your root CA is compromised, all CAs in the hierarchy and all certificates issued from it are considered compromised. You can maximize the security of the root CA by keeping it disconnected from the network and by using subordinate CAs to issue certificates to other subordinate CAs or to end users.

Subordinate CAs

CAs that are not root CAs are considered subordinate. The first subordinate CA in a hierarchy obtains its CA certificate from the root CA. This first subordinate CA can use this key to issue certificates that verify the integrity of another subordinate CA. These higher subordinate CAs are referred to as intermediate CAs. An intermediate CA is subordinate to a root CA, but it serves as a higher certifying authority to one or more subordinate CAs.

An intermediate CA is often referred to as a policy CA because it is typically used to separate classes of certificates that can be distinguished by policies. For example, policy separation includes the level of assurance that a CA provides or the geographical location of the CA to distinguish different end-entity populations. A policy CA can be online or offline.

Warning

It is not possible to convert a root CA to a subordinate CA, or vice versa.

Store a private key

The private key is part of the CA identity, and it must be protected from compromise. Many organizations protect CA private keys by using a hardware security module (HSM). If an HSM is not used, the private key is stored on the CA computer. For more information, see Hardware Security Module (HSM) in Microsoft TechNet.

Offline CAs should be stored in secure locations and not connected to the network. Issuing CAs use their private keys when issuing certificates, so the private key must be accessible (online) while the CA is in operation. In all cases, the CA and its private key on the CA should be physically protected.

Locate an existing key

If you already have an existing private key that you want to use during installation, you can use the **Existing Key** screen to locate that key. You can use the **Change** button to modify the cryptographic provider, and optionally, the CA that you want to search for an existing key.

Locate an existing certificate

If you already have a certificate that contains the private key for the CA, you can use the **Existing Certificate** screen to locate it. You can use the **Import** button to open the **Import Existing Certificate** dialog box, and then locate your existing PKCS #12 file.

Select cryptographic options

Selecting cryptographic options for a certification authority (CA) can have significant security, performance, and compatibility implications for that CA. Although the default cryptographic options may be suitable for most CAs, the ability to implement custom options can be useful to administrators and application developers with a more advanced understanding of cryptography and a need for this flexibility. Cryptographic options can be implemented by using cryptographic service providers (CSPs) or key storage providers (KSPs).

Important
When using an RSA certificate for a CA, ensure that the key length is at least 2048 bits. You must not attempt to use an RSA certificate below 1024 bits for the CA. The CA service (certsvc) will not start if an RSA key of less than 1024 bits is installed.

CSPs are hardware and software components in Windows operating systems that provide generic cryptographic functions. CSPs can be written to provide a variety of encryption and signature algorithms.

KSPs can provide strong key protection for computers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, or Windows Vista.

Important
When you select the provider, hash algorithm, and key length, you should carefully consider what cryptographic options the applications and devices that you intend to use will support.

Allow administrator interaction when the private key is accessed by the CA is an option that is typically used with hardware security modules (HSMs). This allows the cryptographic provider to prompt the user for additional authentication when the private key of the CA is accessed. This option can be used to help prevent unapproved use of the CA and its private key by requiring the administrator to enter a password before every cryptographic operation.

The built-in cryptographic providers support specific key lengths and hash algorithms as described in the following table.

Cryptographic provider	Key lengths	Hash algorithm
Microsoft Base Cryptographic Provider v1.0	• 512 • 1024 • 2048 • 4096	• SHA1 • MD2 • MD4 • MD5

Cryptographic provider	Key lengths	Hash algorithm
Microsoft Base DSS Cryptographic Provider	• 512 • 1024	SHA1
Microsoft Base Smart Card Crypto Provider	• 1024 • 2048 • 4096	• SHA1 • MD2 • MD4 • MD5
Microsoft Enhanced Cryptographic Provider v1.0	• 512 • 1024 • 2048 • 4096	• SHA1 • MD2 • MD4 • MD5
Microsoft Strong Cryptographic Provider	• 512 • 1024 • 2048 • 4096	• SHA1 • MD2 • MD4 • MD5
RSA#Microsoft Software Key Storage Provider	• 512 • 1024 • 2048 • 4096	• SHA1 • SHA256 • SHA384 • SHA512 • MD2 • MD4 • MD5
DSA#Microsoft Software Key Storage Provider	• 512 • 1024 • 2048	SHA1
ECDSA_P256#Microsoft Software Key Storage Provider	256	• SHA1 • SHA256 • SHA384 • SHA512
ECDSA_P384#Microsoft Software Key Storage Provider	384	• SHA1 • SHA256 • SHA384 • SHA512
ECDSA_P521#Microsoft Software Key Storage Provider	521	• SHA1 • SHA256

Cryptographic provider	Key lengths	Hash algorithm
		• SHA384 • SHA512
RSA#Microsoft Smart Card Key Storage Provider	• 1024 • 2048 • 4096	• SHA1 • SHA256 • SHA384 • SHA512 • MD2 • MD4 • MD5
ECDSA_P256#Microsoft Smart Card Key Storage Provider	256	• SHA1 • SHA256 • SHA384 • SHA512
ECDSA_P384#Microsoft Smart Card Key Storage Provider	384	• SHA1 • SHA256 • SHA384 • SHA512
ECDSA_P521#Microsoft Smart Card Key Storage Provider	521	• SHA1 • SHA256 • SHA384 • SHA512

Establish a CA name

Before you configure certification authorities (CAs) in your organization, you should establish a CA naming convention.

You can create a name by using any Unicode character, but you might want to use the ANSI character set if interoperability is a concern. For example, certain types of routers will not be able to use the Network Device Enrollment Service to enroll for certificates if the CA name contains special characters such as an underscore.

Important

If you use non-Latin characters (such as Cyrillic, Arabic, or Chinese characters), your CA name must contain fewer than 64 characters. If you use only non-Latin characters, your CA name can be no more than 37 characters in length.

In Active Directory Domain Services (AD DS), the name that you specify when you configure a server as a CA becomes the common name of the CA, and this name is reflected in every

certificate that the CA issues. For this reason, it is important that you do not use the fully qualified domain name for the common name of the CA. This way, malicious users who obtain a copy of a certificate cannot identify and use the fully qualified domain name of the CA to create a potential security vulnerability.

Warning

The CA name should not be identical to the name of the computer (NetBIOS or DNS name). Also, you cannot change the name of a server after Active Directory Certificate Services (AD CS) is installed without invalidating all the certificates that are issued by the CA. For additional considerations regarding CA names, see TechNet Wiki article: Considerations for Certification Authority (CA) Names.

To change the server name after AD CS is installed, you must uninstall the CA, change the name of the server, reinstall the CA using the same keys and modify the registry to use the existing CA keys and database. You do not have to reinstall a CA if you rename a domain; however, you will have to reconfigure the CA to support the name change.

Obtain a certificate request

After a root certification authority (CA) has been installed, many organizations will install one or more subordinate CAs to implement policy restrictions on the public key infrastructure (PKI) and to issue certificates to end clients. Using at least one subordinate CA can help protect the root CA from unnecessary exposure. When you install a subordinate CA, you must obtain a certificate from the parent CA.

If the parent CA is online, you can use the **Send a certificate request to a parent CA** option, and select the parent CA by CA name or computer name.

If the parent CA is offline, you should use the **Save a certificate request to file on the target machine** option. The procedure for this will be unique to the parent CA. At a minimum, the parent CA should provide a file that contains the subordinate CA's newly issued certificate, preferably its full certification path.

If you get a subordinate CA certificate that does not include the full certification path, the new subordinate CA that you install must be able to build a valid CA chain when it starts. Do the following to create a valid certification path:

- Install the parent CA's certificate in the **Intermediate Certification Authorities** certificate store of the computer if the parent CA is not a root CA.
- Install the certificates of any other intermediate CA in the chain.
- Install the certificate of the root CA into the Trusted Root Certification Authorities store.

Note

These certificates should be installed in the certificate store before you install the CA certificate on the subordinate CA you have just set up.

Verify the validity period

Certificate-based cryptography uses public-key cryptography to protect and sign data. Over time, attackers could obtain data that was protected with the public key and attempt to derive the private key from it. Given enough time and resources, this private key could be compromised, effectively rendering all protected data unprotected. Also the names that are guaranteed by a certificate may need to be changed over time. Because a certificate is a binding between a name and a public key, when either of these change, the certificate should be renewed.

Every certificate has a validity period. After the end of the validity period, the certificate is no longer considered an acceptable or usable credential.

CAs cannot issue certificates that are valid beyond their own validity period. A best practice is to renew the CA certificate when half of its validity period is expired. When installing a CA, you should plan this date and ensure that it is recorded as a future task.

Choose a CA database

As in many databases, the certification authority's database is a file on the hard drive. In addition to this file, other files serve as the transaction logs, and they receive all modifications to the database before the changes are made. Because these files may be accessed frequently and simultaneously, it is best to keep the database and transaction logs on separate hard drives or high-performance disk configurations, such as striped volumes.

The location of the certificate database and log files are kept in the following registry location:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration

The registry contains following values:

- DBDirectory
- DBLogDirectory
- DBSystemDirectory
- DBTempDirectory

Note

You can move the certificate database and log files after installation. For information, see article 283193in the Microsoft Knowledge Base.

Configure the CA

After a root or subordinate CA is installed, you must configure the Authority Information Access (AIA) and CRL distribution point (CDP) extensions before the CA issues any certificates. The AIA extension specifies where to find up-to-date certificates for the CA. The CDP extension specifies where to find up-to-date CRLs that are signed by the CA. These extensions apply to all certificates that are issued by that CA.

Configuring these extensions ensures that this information is included in each certificate that the CA issues so that it is available to all clients. This ensures that PKI clients experience the least possible number of failures due to unverified certificate chains or certificate revocations, which

can result in unsuccessful VPN connections, failed smart card sign-ins, or unverified email signatures.

As a CA administrator, you can add, remove, or modify CRL distribution points and the locations for CDP and AIA certificate issuance. Modifying the URL for a CRL distribution point only affects newly issued certificates. Previously issued certificates will continue to reference the original location, which is why you should establish these locations before your CA distributes any certificates.

Consider these guidelines when you configure CDP extension URLs:

- Avoid publishing delta CRLs on offline root CAs. Because you do not revoke many certificates on an offline root CA, a delta CRL is probably not needed.
- Adjust the default LDAP:/// and HTTP:// URL locations on the **Extensions** tab of the certification authority's **Properties Extension** tab according to your needs.
- Publish a CRL on an HTTP Internet or extranet location so that users and applications outside the organization can perform certificate validation. You can publish the LDAP and HTTP URLs for CDP locations to enable clients to retrieve CRL data with HTTP and LDAP.
- Remember that Windows clients always retrieve the list of URLs in sequential order until a valid CRL is retrieved.
- Use HTTP CDP locations to provide accessible CRL locations for clients running non-Windows operating systems.

Note

For more information about CRLs and delta CRLs, see Configuring Certificate Revocation.

Windows PowerShell and `certutil` support variable numbers (preceded by a percent (%) sign) to help in publishing CDP and AIA locations. The CA's **Properties Extension** tab supports bracketed variables. The following table equates the variables between the interfaces and describes their meanings.

Variable	Extensions tab name	Description
%1	<ServerDNSName>	The DNS name for the CA computer. If connected to a DNS domain, it is the fully qualified domain name; otherwise, it is the hostname of the computer.
%2	<ServerShortName>	The NetBIOS name of the CA server
%3	<CaName>	The name of the CA
%4	<CertificateName>	This allows each additional revision of the certificate to

Variable	Extensions tab name	Description
		have a unique suffix.
%4	None	Not used
%6	<ConfigurationContainer>	The location of the configuration container in Active Directory Domain Services (AD DS)
%7	<CATruncatedName>	The name of the CA truncated to 32 characters with a hash at the end
%8	<CRLNameSuffix>	This inserts a suffix on the file name when publishing a CRL to a file or URL location.
%9	<DeltaCRLAllowed>	When a delta CRL is published, this replaces the CRLNameSuffix variable with a separate suffix to distinguish the delta CRL from the CRL.
%10	<CDPObjectClass>	The object class identifier for CRL distribution points, which is used when publishing to an LDAP URL.
%11	<CAObjectClass>	The object class identifier for a CA, which is used when publishing to an LDAP URL.

Publish the AIA extension

The AIA extension tells the client computers where they can find the certificate to be verified. This allows the client to confirm whether the certificate can be trusted.

You can configure the AIA extension by using the Certification Authority interface, Windows PowerShell, or the `certutil` command. The following table describes the options that you can use with the AIA extension by using these methods.

Interface check box name	Windows PowerShell parameter	Certutil value
Include in the AIA extension of issued certificate	-AddToCertificateAia	2

Interface check box name	Windows PowerShell parameter	Certutil value
Include in the online certificate status protocol (OCSP) extension	-AddToCertificateOcsp	32

The examples in this section for publishing the AIA extension represent the following scenario:

- The domain name is corp.contoso.com.
- There is a web server named App1 in the domain.
- App1 has a shared folder named PKI that allows the CA Read and Write permissions.
- App1 has a DNS CNAME of www and a shared virtual directory named PKI.
- The first protocol that client computers should use for the AIA information is HTTP.
- The second protocol that client computers should use for the AIA information is LDAP.
- The CA that is being configured is an online issuing CA.
- OCSP is not in use.

Use the interface to publish the AIA extension

The interface uses the variables and check box names that are described in the previous tables. You can access the interface through the Certification Authority interface. From the contents pane, right-click the CA, click **Properties**, and then click **Extensions**. In **Select extension**, click **Authority Information Access (AIA)**.

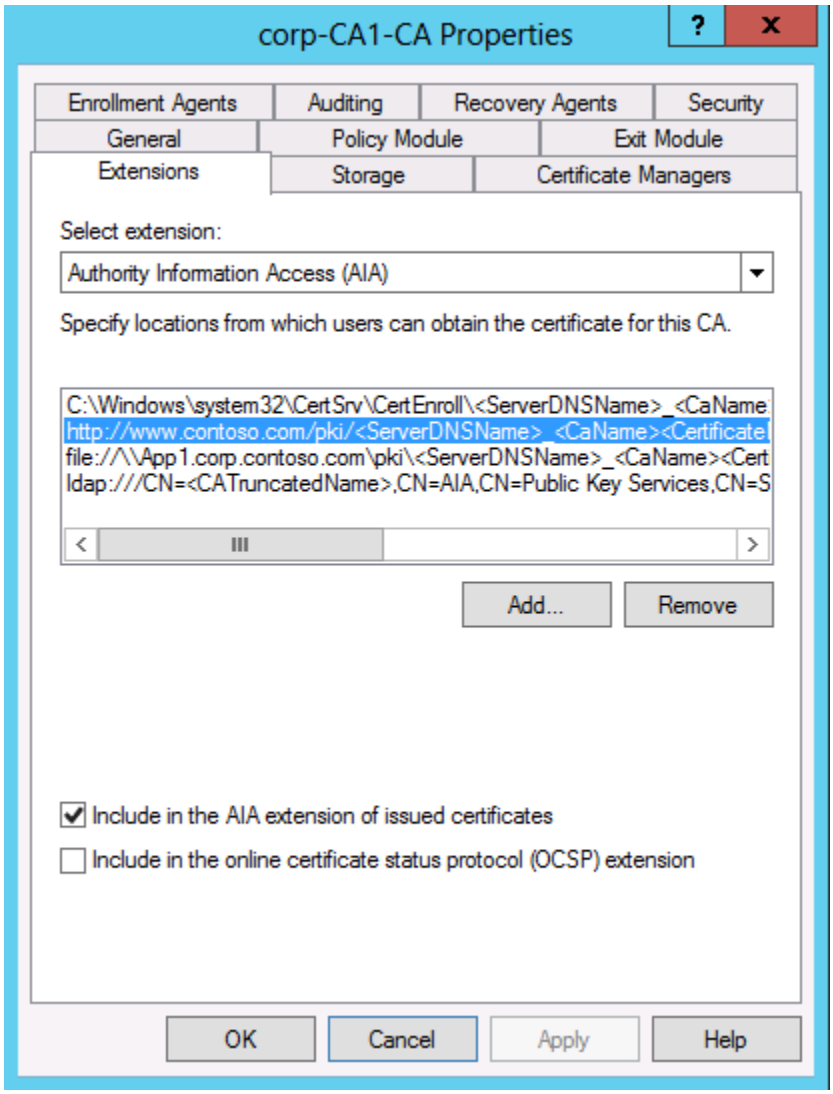

Figure 1 AIA extension menu

The locations and settings configured in the user interface are as follows:

- C:\Windows\system32\CertSrv\CertEnroll\<ServerDNSName>_<CaName><CertificateName>.crt
- http://www.contoso.com/pki/<ServerDNSName>_<CaName><CertificateName>.crt
  - Include in the AIA extension of issued certificates
- file://\\App1.corp.contoso.com\pki\<ServerDNSName>_<CaName><CertificateName>.crt
- ldap:///CN=<CATruncatedName>,CN=AIA,CN=Public Key Services,CN=Services,<ConfigurationContainer><CAObjectClass>
  - Include in the AIA extension of issued certificates

Use Windows PowerShell to publish the AIA extension

The following Windows PowerShell commands can be used to configure the AIA extension for the given scenario:

```
$aialist = Get-CAAuthorityInformationAccess; foreach ($aia in $aialist) {Remove-
CAAuthorityInformationAccess $aia.uri -Force};

Add-CAAuthorityInformationAccess -AddToCertificateAia
http://www.contoso.com/pki/%1_%3%4.crt

Add-CAAuthorityInformationAccess -AddToCertificateAia
file://\\App1.corp.contoso.com\pki\%1_%3%4.crt

Add-CAAuthorityInformationAccess -AddToCertificateAia "ldap:///CN=%7,CN=AIA,CN=Public Key
Services,CN=Services,%6%11"
```

Notes

- If you use Windows PowerShell to add AIA paths, existing paths remain in place. The first Windows PowerShell command in the example removes all the existing paths. For more information about removing AIA paths by using Windows PowerShell, see Remove-CAAuthorityInformationAccess.
- You cannot add a local path by using the **Add-CAAuthorityInformationAccess** Windows PowerShell cmdlet. The CA certificate will automatically be published to the default location of %systemroot%\system32\CertSrv\CertEnroll.

Use certutil to publish the AIA extension

The following `certutil` command can be used to configure the AIA extension for the given scenario:

```
certutil -setreg CA\CACertPublicationURLs
"1:C:\Windows\system32\CertSrv\CertEnroll\%1_%3%4.crt\n2:http://www.contoso.com/pki/%1_%3
%4.crt\n1:file://\\App1.corp.contoso.com\pki\%1_%3%4.crt\n2:ldap:///CN=%7,CN=AIA,CN=Publi
c Key Services,CN=Services,%6%11"
```

Notes

- After you changethese paths, be sure to restart the CertSvc.You can restart the CertSvc by running the following Windows PowerShell command: `restart-service certsvc`
- In a `certutil` command, type all paths as one continuous string enclosed in quotes. Each path is separated by a `\n`.

Publish the CDP extension

The CDP extension tells client computers where they can find the most recent CRL, so the client can confirm that a particular certificate has not been revoked.

You can configure the CDP extension by using the Certification Authority interface, Windows PowerShell, or the `certutil` command. The following table describes the options that you can use with the CDP extension by using these methods.

Interface check box name	Windows PowerShell parameter	Certutil value
Publish CRLs to this location	-PublishToServer	1
Include in all CRLs. (Specifies where to publish in Active Directory when publishing manually.)	-AddToCrlCdp	8
Include in CRLs. (Clients use this to find the delta CRL locations.)	-AddToFreshestCrl	4
Include in the CDP extension of issued certificates.	-AddToCertificateCdp	2
Publish Delta CRLs to this location.	-PublishDeltaToServer	64
Include in the IDP extension of issued CRLs.	-AddToCrlIdp	128

Notes

The Issuing Distribution Point (IDP) extension is used by non-Windows clients to verify certificate revocation. The IDP extension allows partitioned CRLs to be deployed when using third-party CAs. Partitioned CRLs allow a third-party CA to publish CRLs with only specific certificate types within each CRL. For example, you can have separate CRLs for end certificates versus CA certificates. Specifically, the following options can be set in the IDP:

1. onlyContainUserCerts. This option in the IDP allows only certificates that do not have the values cA in the Basic Constraints extension. If the certificate does not contain a Basic Constraints extension, it is assumed it is not a CA.
2. onlyContainsCACerts. This option in the IDP allows only certificates having a Basic Constraints extension with cA set to be included in the CRL.

If you are allowing delta CRL publishing to an Internet Information Services (IIS) web server, you must modify the default IIS configuration by setting **allowDoubleEscaping=true** of the **requestFiltering** element in the **system.web** section of the IIS configuration. For example, if you want to allow double escaping for the PKI virtual directory of the default Web site on IIS, run the following command on the IIS web server: `appcmd set config "Default Web Site/pki" -section:system.webServer/security/requestFiltering -allowDoubleEscaping:true`. For more

information, see AD CS: Web server should allow URI containing the “+” character to enable publishing of delta CRLs.

The examples in this section for publishing the CDP extension represent the following scenario:

- The domain name is corp.contoso.com.
- There is a web server named App1 in the domain.
- App1 has a shared folder named PKI that allows the CA Read and Write permissions.
- App1 has a DNS CNAME of www and a shared virtual directory named PKI.
- The first protocol that client computers should use for the CDP information is HTTP.
- The second protocol that client computers should use for the CDP information is LDAP.
- The CA that is being configured is an online issuing CA.
- IDP is not in use.

Use the interface to publish the CDP extension

The interface uses the variables and check box names that are described in the previous tables. You can access the interface through the Certification Authority interface. From the contents pane, right-click the CA, click **Properties**, and then click **Extensions**. In **Select extension**, click **CRL Distribution Point (CDP)**.

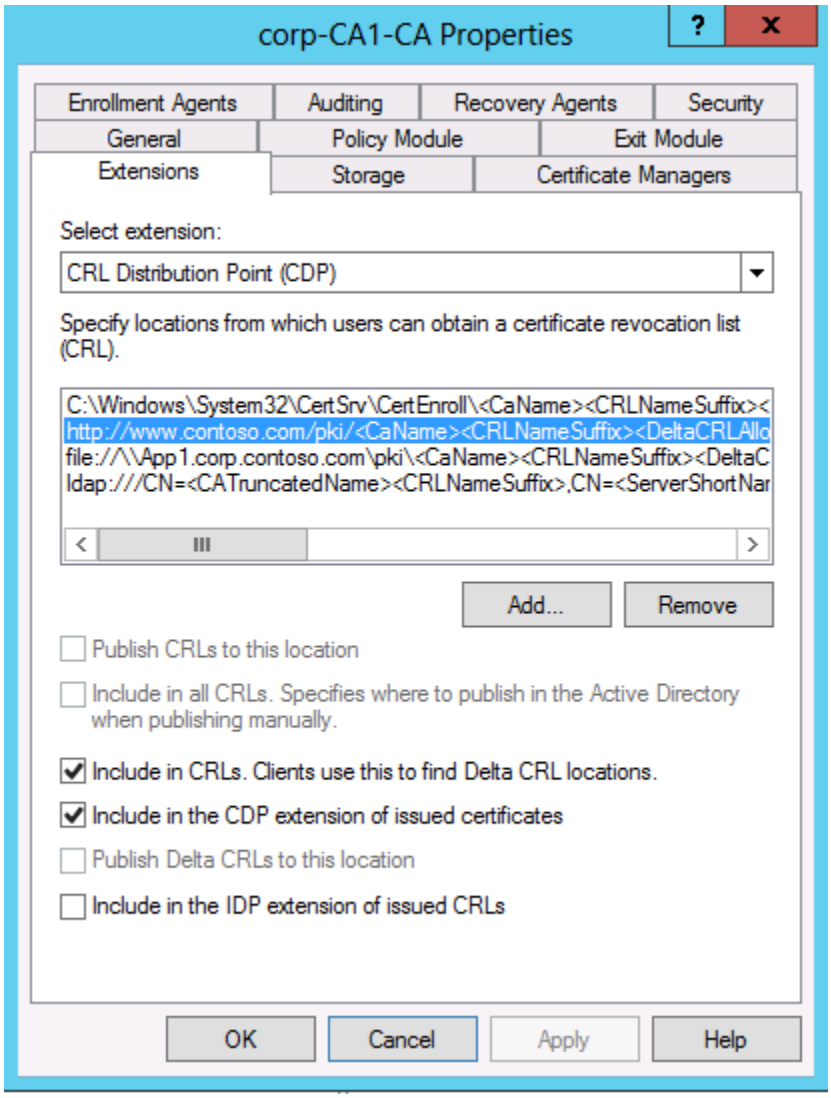

Figure 2 CDP extension menu

The locations and settings configured in the interface are as follows:

- C:\Windows\System32\CertSrv\CertEnroll\<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl
  - Publish CRLs to this location
  - Publish delta CRLs to this location
- http://www.contoso.com/pki/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl
  - Include in CRLs. Clients use this to find delta CRL locations.
  - Include in the CDP extension of issued certificates
- file://\\App1.corp.contoso.com\pki\<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl
  - Publish CRLs to this location

- Publish Delta CRLs to this location

- ldap:///CN=<CATruncatedName><CRLNameSuffix>,CN=<ServerShortName>,CN=CDP, CN=Public Key Services,CN=Services,<ConfigurationContainer><CDPObjectClass>
  - Include in all CRLs. Specifies where to publish in the Active Directory when publishing manually.
  - Include in the CDP extension of certificates

Use Windows PowerShell to publish the CDP extension

The following Windows PowerShell commands are used to configure the CDP extension for the given scenario:

```
$crllist = Get-CACrlDistributionPoint; foreach ($crl in $crllist) {Remove-CACrlDistributionPoint $crl.uri -Force};

Add-CACRLDistributionPoint -Uri C:\Windows\System32\CertSrv\CertEnroll\%3%8%9.crl -PublishToServer -PublishDeltaToServer

Add-CACRLDistributionPoint -Uri http://www.contoso.com/pki/%3%8%9.crl -AddToCertificateCDP -AddToFreshestCrl

Add-CACRLDistributionPoint -Uri file:\\\\App1.corp.contoso.com\pki\%3%8%9.crl -PublishToServer -PublishDeltaToServer

Add-CACRLDistributionPoint -Uri "ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10" -AddToCrlCdp -AddToCertificateCdp
```

Note

If you use Windows PowerShell to add CDP paths, existing paths remain in place. The first Windows PowerShell command in the example removes all the existing paths. For more information about using Windows PowerShell to remove CDP paths, see Remove-CACrlDistributionPoint.

Use certutil to publish the CDP extension

The following `certutil` command configures the CDP extension for the given scenario:

```
certutil -setreg CA\CRLPublicationURLs "65:C:\Windows\system32\CertSrv\CertEnroll\%3%8%9.crl\n6:http://www.contoso.com/pki/%3%8%9.crl\n65:file:\\\\App1.corp.contoso.com\pki\%3%8%9.crl\n10:ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10"
```

Notes

- After you change these paths, be sure to restart the CA service. From Windows PowerShell, you can restart the CertSvc by running the following command: `restart-service certsvc`
- In a `certutil` command, type all paths as one continuous string enclosed in quotes, but separate each path with `\n`.

To publish the CRL, you can run the command `certutil -crl` on the CA from Windows PowerShell or a command prompt run as administrator. For more information about CRL configuration and publishing, see Configuring Certificate Revocation.

Verify the configuration

To verify the CA configuration, you can run the following commands from Windows PowerShell or a from a Command Prompt window:

Command	Description
Certutil -CAInfo	Shows the status of the names, locale, object identifiers (OIDs), and CRLs for the CA.
Certutil -getreg	Displays the CA registry configuration.
Certutil -ADCA	Confirms the configuration of enterprise CAs.

You can use the Enterprise PKI View (PKIView.msc) tool to check your AIA and CDP publication configurations. For more information, see the Enterprise PKI.

You can also use the Online Responder role service to check certificate revocation. For more information about Online Responder, see Online Responder Installation, Configuration, and Troubleshooting Guide.

Related content

1. Windows Server Security Forum
2. Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
3. Windows PKI Documentation Reference and Library
4. Windows PKI Blog

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Configure Trusted Roots and Disallowed Certificates

The Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8 operating systems include an automatic update mechanism that downloads certificate trust lists (CTLs) on a daily basis. In Windows Server 2012 R2 and Windows 8.1, additional capabilities are available to control how the CTLs are updated.

Important

Software updates are available for Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista. To provide the enhancements of the automatic update mechanism that are discussed in this document, apply the following updates:

- For Windows Server 2008 R2, Windows Server 2008, Windows 7, or Windows Vista, apply the appropriate update listed in document 2677070 in the Microsoft Knowledge Base.
- For Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, or Windows Vista, apply the appropriate update listed in document 2813430 in the Microsoft Knowledge Base.

Note

To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Certificates and trust

The Microsoft Root Certificate Program enables distribution of trusted root certificates within Windows operating systems. For more information about the list of members in Windows Root Certificate Program, see Windows Root Certificate Program - Members List (All CAs).

Trusted root certificates are meant to be placed in the Trusted Root Certification Authorities certificate of the Windows operating systems. These certificates are trusted by the operating system and can be used by applications as a reference for which public key infrastructure (PKI) hierarchies and digital certificates that are trustworthy. There are two methods for distributing trusted root certificates:

1. **Automatic**: The list of trusted root certificates is stored in a CTL. Client computers access the Windows Update site by using the automatic update mechanism to update this CTL.

   Note

   The list of trusted root certificates is called the trusted CTL.

2. **Manual**: The list of trusted root certificates is available as a self-extracting IEXPRESS package in the Microsoft Download Center, the Windows catalog, or by using Windows Server Update Services (WSUS). IEXPRESS packages are released at the same time as the trusted CTL.

Note

For more information about these update methods, see document 931125 in the Microsoft Knowledge Base.

Untrusted certificates are certificates that are publicly known to be fraudulent. Similar to the trusted CTL, there are two mechanisms that are used to distribute a list of untrusted certificates:

1. **Automatic**: The list of untrusted certificates is stored in a CTL. Client computers access the Windows Update site by using the automatic update mechanism to update this CTL.

Note
A list of untrusted certificates is called an untrusted CTL. For more information, see Announcing the automated updater of untrustworthy certificates and keys.

2. **Manual**: The list of untrusted certificates comes as a self-extracting IEXPRESS package in a mandatory security Windows Update.

Prior to Windows Server 2012 R2 and Windows 8.1 (or the installation of the software update, as previously discussed), the same registry setting controlled updates for trusted root certificates and untrusted certificates. An administrator could not selectively enable or disable one or the other. This resulting in the following challenges:

- If the organization was in a disconnected environment, the only method for updating CTLs was to use IEXPRESS packages.

  Note
  A computer network where the computers do not have the ability to access the Windows Update site is considered a disconnected environment in this document.

  The IEXPRESS update method is mostly a manual process. Further, the IEXPRESS package may not be immediately available when the CTL is released, so there could be an additional lag for installing these updates when using this method.
- Although disabling automatic updates for trusted CTLs is recommended for administrators who manage their lists of trusted root certificates (in disconnected or connected environments), disabling automatic updates of untrusted CTLs is not recommended.

  For more information, see Controlling the Update Root certificate Certificates Feature to Prevent the Flow of Information to and from the Internet.
- Because there was not a method for network administrators to view and extract only the trusted root certificates in a trusted CTL, managing a customized list of trusted certificates was difficult task.

Software update description

The following improved automatic update mechanisms for a disconnected environment are available in Windows Server 2012 R2 and Windows 8.1 or when the appropriate software update is installed:

- **Registry settings for storing CTLs** New settings enable changing the location for uploading trusted or untrusted CTLs from the Windows Update site to a shared location in an organization. For more information, see the Registry settings modified section.
- **Synchronization options** If the URL for the Windows Update site is moved to a local shared folder, the local shared folder must be synchronized with the Windows Update folder. This software update adds a set of options in the Certutil tool that administrators can use to enable synchronization. For more information, see the New Certutil Options section.
- **Tool to select trusted root certificates** This software update introduces a tool for administrators who manage the set of trusted root certificates in their enterprise environment. Administrators can view and select the set of trusted root certificates, export them to a

serialized certificate store, and distribute them by using Group Policy. For more information, see the New Certutil Options section in this document.

- **Independent configurability** The automatic update mechanism for trusted and untrusted certificates are independently configurable. This enables administrators to use the automatic update mechanism to download only the untrusted CTLs and manage their own list of trusted CTLs. For more information, see the Registry settings modified section in this document.

Configuration options

In Windows Server 2012 R2 and Windows 8.1 (or by installing the previously mentioned software updates on supported operating systems), an administrator can configure a file or web server to download the following files by using the automatic update mechanism:

- authrootstl.cab, which contains a non-Microsoft CTL
- disallowedcertstl.cab, which contains a CTL with untrusted certificates
- disallowedcert.sst, which contains a serialized certificate store, including untrusted certificates
- *thumbprint*.crt, which contains non-Microsoft root certificates

The steps to perform this configuration are described in the Configure a file or web server to download the CTL files section of this document.

By using Windows Server 2012 R2 and Windows 8.1 (or by installing the previously mentioned software updates on supported operating systems), an administrator can:

- Configure Active Directory Domain Services (AD DS) domain member computers to use the automatic update mechanism for trusted and untrusted CTLs, without having access to the Windows Update site. This configuration is described in the Redirect the Microsoft Automatic Update URL for a disconnected environment section of this document.
- Configure AD DS domain member computers to independently opt-in for untrusted and trusted CTL automatic updates. This configuration is described in the Redirect the Microsoft Automatic Update URL for untrusted CTLs only section of this document.
- Examine the set of root certificates in the Windows Root Certificate Program. This enables administrators to select a subset of certificates to distribute by using a Group Policy Object (GPO). This is configuration is described in the Use a subset of the trusted CTLs section of this document.

Important

- All the steps shown in this document require that you use an account that is a member of the local **Administrators** group. For all Active Directory Domain Services (AD DS) configuration steps, you must use an account that is a member of the **Domain Admins** group or that has been delegated the necessary permissions.
- The procedures in this document depend upon having at least one computer that is able to connect to the Internet to download CTLs from Microsoft. The computer requires HTTP (TCP port 80) access and name resolution (TCP and UDP port 53) ability to contact ctldl.windowsupdate.com. This computer can be a domain member or a member of a workgroup. Currently all the downloaded files require approximately 1.5 MB of space.

- The settings described in this document are implemented by using GPOs. These settings are not automatically removed if the GPO is unlinked or removed from the AD DS domain. When implemented, these settings can be changed only by using a GPO or by modifying the registry of the affected computers.
- The concepts discussed in this document are independent of Windows Server Update Services (WSUS).
  - You do not have to use WSUS to implement the configuration discussed in this document.
  - If you do use WSUS, these instructions will not affect its functionality.
  - Implementing WSUS is not a substitute for implementing the configurations discussed in this document.

Configure a file or web server to download the CTL files

To facilitate the distribution of trusted or untrusted certificates for a disconnected environment, you must first configure a file or web server to download the CTL files from the automatic update mechanism.

Tip

The configuration described in this section is not needed for environments where computers are able to connect to the Windows Update site directly. Computers that can connect to the Windows Update site are able to receive updated CTLs on a daily basis (if they are running Windows Server 2012, Windows 8, or the previously mentioned software updates are installed on supported operating systems). For more information, see document 2677070 automatic in the Microsoft Knowledge Base.

To configure a server that has access to the Internet to retrieve the CTL files

1. Create a shared folder on a file or web server that is able to synchronize by using the automatic update mechanism and that you want to use to store the CTL files.

   Tip

   Before you begin, you may have to adjust the shared folder permissions and NTFS folder permissions to allow the appropriate account access, especially if you are using a scheduled task with a service account. For more information on adjusting permissions see Managing Permissions for Shared Folders.

2. From an elevated command prompt, run the following command:

   ```
   Certutil -syncWithWU \\<server>\<share>
   ```

   Substitute the actual server name for <server> and shared folder name for <share>. For example, if you run this command for a server named Server1 with a shared folder named CTL, you would run the command:

   ```
   Certutil -syncWithWU \\Server1\CTL
   ```

3. Download the CTL files on a server that computers on a disconnected environment can access over the network by using a FILE path (for example, FILE://\\Server1\CTL) or an HTTP path (for example, HTTP://Server1/CTL).

Notes

- If the server that synchronizes the CTLs is not accessible from the computers in the disconnected environment, you must provide another method to transfer the information. For example, you can allow one of the domain member computers to connect to the server, then schedule another task on the domain member computer to pull the information into a shared folder on an internal web server. If there is absolutely no network connection, you may have to use a manual process to transfer the files, such as a removable storage device.
- If you plan to use a web server, you should create a new virtual directory for the CTL files. The steps to create a virtual directory by using Internet Information Services (IIS) are nearly the same for all the supported operating systems discussed in this document. For more information, see Create a Virtual Directory (IIS7).
- Be aware that certain system and application folders in Windows have special protection applied to them. For example, the **inetpub** folder requires special access permissions, which makes it difficult to create a shared folder for use with a scheduled task to transfer files. As an administrator, you are typically able to create a folder location at the root of a logical drive system to use for file transfers.

Redirect the Microsoft Automatic Update URL for a disconnected environment

If the computers in your network are configured in a domain environment and they are unable to use the automatic update mechanism or download CTLs, you can implement a GPO in AD DS to configure those computers to obtain the CTL updates from an alternate location.

Note

The configuration in this section requires that you have already completed the steps in Configure a file or web server to download the CTL files.

To configure a custom administrative template for a GPO

1. On a domain controller, create a new administrative template. You can start this as a text file and then change the file name extension to .adm. The contents of the file should be as follows:

```
CLASS MACHINE

CATEGORY !!SystemCertificates

    KEYNAME
"Software\Microsoft\SystemCertificates\AuthRoot\AutoUpdate"

    POLICY !!RootDirURL

       EXPLAIN !!RootDirURL_help
```

```
        PART !!RootDirURL EDITTEXT
              VALUENAME "RootDirURL"
        END PART
    END POLICY
END CATEGORY
[strings]
RootDirURL="URL address to be used instead of default
ctldl.windowsupdate.com"
RootDirURL_help="Enter a FILE or HTTP URL to use as the
download location of the CTL files."
SystemCertificates="Windows AutoUpdate Settings"
```

2. Use a descriptive name to save the file, such as **RootDirURL.adm**.

   Tip
   - Ensure that the file name extension is .adm and not .txt.
   - If you have not already enabled file name extension viewing, see How To: View File Name Extensions.
   - If you save the file to the %windir%\inf folder, it will be easier to locate in the following steps.

3. Open the Group Policy Management Editor.
   - If you are using Windows Server 2008 R2 or Windows Server 2008, click **Start**, and then click **Run**.
   - If you are using Windows Server 2012 R2 or Windows Server 2012, press the **Windows** key plus the **R** key simultaneously.

   Type **GPMC.msc**, and then press ENTER.

   Caution
   You can link a new GPO to the domain or to any organizational unit (OU). The GPO modifications implemented in this document alter the registry settings of the affected computers. You cannot undo these settings by deleting or unlinking the GPO. The settings can only be undone by reversing them in the GPO settings or by modifying the registry using another technique.

4. In the Group Policy Management console, expand the **Forest** object, expand the **Domains** object, and then expand the specific domain that contains the computer accounts that you want to change. If you have a specific OU that you want to modify, then navigate to that location. Click an existing GPO or right-click and then click **Create a GPO in this domain, and Link it here** to create a new GPO. Right-click the GPO you want to modify and then click **Edit**.
5. In the navigation pane, under **Computer Configuration**, expand **Policies**.

6. Right-click **Administrative Templates**, and then click **Add/Remove Templates**.
7. In **Add/Remove Templates**, click **Add**. In the **Policy Templates** dialog box, select the .adm template that you previously saved. Click **Open**, and then click **Close**.
8. In the navigation pane, expand **Administrative Templates**, and then expand **Classic Administrative Templates (ADM)**.
9. Click **Windows AutoUpdate Settings**, and in the details pane, double-click **URL address to be used instead of default ctldl.windowsupdate.com**.
10. Select **Enabled**. In the **Options** section, enter the URL to the file server or web server that contains the CTL files. For example, **http://server1/CTL** or **file://\\server1\CTL**. Click **OK**. Close the Group Policy Management Editor.

The policy is effective immediately, but the client computers must be restarted to receive the new settings, or you can type **gpupdate /force** from an elevated command prompt or from Windows PowerShell.

Important
The trusted and untrusted CTLs can be updated on a daily basis, so ensure that you keep the files synchronized by using a scheduled task or another method (such as a script that handles error conditions) to update the shared folder or web virtual directory. For additional details about creating a scheduled task, see Schedule a Task. If you plan to write a script to make daily updates, see the New Certutil Options and Potential errors with Certutil -SyncWithWU sections of this document. These sections provide more information about command options and the error conditions.

Redirect the Microsoft Automatic Update URL for untrusted CTLs only

Some organizations may want only the untrusted CTLs (not the trusted CTLs) to be automatically updated. To accomplish this, you can create two .adm templates to add to Group Policy.

Important
1. In a disconnected environment, you can use the following procedure with the previous procedure (redirect the Microsoft Automatic Update URL for trusted CTLs and untrusted CTLs). This procedure explains how to selectively disable the automatic update of trusted CTLs.
2. You can also use this procedure in a connected environment in isolation to selectively disable the automatic update of trusted CTLs.

To selectively redirect only untrusted CTLs

1. On a domain controller, create the first new administrative template by starting with a text file and then changing the file name extension to .adm. The contents of the file should be as follows:

```
CLASS MACHINE
```

```
CATEGORY !!SystemCertificates
    POLICY !!DisableRootAutoUpdate
       EXPLAIN !!Certificates_config
       VALUENAME "DisableRootAutoUpdate"
       VALUEON NUMERIC 0
          VALUEOFF NUMERIC 1
       KEYNAME
"Software\Policies\Microsoft\SystemCertificates\AuthRoot"
    END POLICY
END CATEGORY
[strings]
DisableRootAutoUpdate="Auto Root Update"
Certificates_config="By default automatic updating of the
trusted CTL is enabled. To disable the automatic updating
trusted CTLe, select Disabled."
SystemCertificates="Windows AutoUpdate Settings"
```

2. Use a descriptive name to save the file, such as **DisableAllowedCTLUpdate.adm**.
3. Create a second new administrative template. The contents of the file should be as follows:

```
CLASS MACHINE
CATEGORY !!SystemCertificates
    POLICY !!EnableDisallowedCertAutoUpdate
       EXPLAIN !!Certificates_config
       VALUENAME "EnableDisallowedCertAutoUpdate"
       VALUEON NUMERIC 1
          VALUEOFF NUMERIC 0
       KEYNAME
"Software\Policies\Microsoft\SystemCertificates\AuthRoot"
    END POLICY
END CATEGORY
[strings]
EnableDisallowedCertAutoUpdate="Untrusted CTL Automatic
Update"
Certificates_config="By default untrusted CTL automatic
```

```
update is enabled. To disable trusted CTL update, select
Disabled."
SystemCertificates="Windows AutoUpdate Settings"
```

4. Use a descriptive file name to save the file, such as **EnableUntrustedCTLUpdate.adm**.

   Tip
   - Ensure that the file name extensions of these files are .adm and not .txt.
   - If you have not already enabled file name extension viewing, see How To: View File Name Extensions.
   - If you save the file to the %windir%\inf folder, it will be easier to locate in the following steps.
5. Open the Group Policy Management Editor.
6. In the Group Policy Management console, expand the **Forest**, **Domains**, and specific domain object that you want to modify. Right-click the **Default Domain Policy GPO**, and then click **Edit**.
7. In the navigation pane, under **Computer Configuration**, expand **Policies**.
8. Right-click **Administrative Templates**, and then click **Add/Remove Templates**.
9. In **Add/Remove Templates**, click **Add**. Use the **Policy Templates** dialog box to select the .adm templates that you previously saved. (You can hold the CTRL key, and click each file to select both.) Click **Open**, and then click **Close**.
10. In the navigation pane, expand **Administrative Templates** and then expand **Classic Administrative Templates (ADM)**.
11. Click **Windows AutoUpdate Settings** and then in the details pane, double-click **Auto Root Update**.
12. Select **Disabled**. This setting prevents the automatic update of the trusted CTLs. Click **OK**.
13. In the details pane, double-click **Untrusted CTL Automatic Update**. Select **Enabled**. Click **OK**.

The policy is effective immediately, but the client computers must be restarted to receive the new settings, or you can type **gpupdate /force** from an elevated command prompt or from Windows PowerShell.

Important
The trusted and untrusted CTLs can be updated on a daily basis, so ensure that you keep the files synchronized by using a scheduled task or another method to update the shared folder or virtual directory.

Use a subset of the trusted CTLs

This section describes how you can produce, review, and filter the trusted CTLs that you want computers in your organization to use. You must implement the GPOs described in the previous

procedures to make use of this resolution. This resolution is available for disconnected and connected environments.

There are two procedures to complete to customize the list of trusted CTLs.

1. Create a subset of trusted certificates
2. Distribute the trusted certificates by using Group Policy

To create a subset of trusted certificates

1. From a computer that is connected to the Internet, open Windows PowerShell as an Administrator or open an elevated command prompt, and type the following command:

   ```
   Certutil -generateSSTFromWU WURoots.sst
   ```

2. You can run the following command in Windows Explorer to open the WURoots.sst::

   ```
   start explorer.exe wuroots.sst
   ```

   Tip
   You can also use Internet Explorer to navigate to the file and double-click it to open it. Depending on where you stored the file, you may also be able to open it by typing `wuroots.sst`.

3. In the navigation pane of Certificate Manager, expand the file path under **Certificates - Current User** until you see **Certificates**, and then click **Certificates**.
4. In the details pane, you can see the trusted certificates. Hold down the CTRL key and click each of the certificates that you want to allow. When you have finished selecting the certificates you want to allow, right-click one of the selected certificates, click **All Tasks**, and then click **Export**.

   Important
   You must select a minimum of two certificates to export the .sst file type. If you select only one certificate, the .sst file type is not available and the .cer file type is selected instead.

5. In the Certificate Export Wizard, click **Next**.
6. On the **Export File Format** page, select **Microsoft Serialized Certificate Store (.SST)**, and then click **Next**.
7. On the **File to Export** page, enter a file path and an appropriate name for the file, such as **C:\AllowedCerts.sst**, and then click **Next**. Click **Finish**. When you are notified that the export was successful, click **OK**.
8. Copy the .sst file that you created to a domain controller.

To distribute the list of trusted certificates by using Group Policy

1. On the domain controller that has the customized .sst file, open the Group Policy Management Editor.
2. In the Group Policy Management console, expand the **Forest**, **Domains**, and specific domain object that you want to modify. Right-click **Default Domain Policy GPO**, and

then click **Edit**.

3. In the navigation pane, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then expand **Public Key Policies**.
4. Right-click **Trusted Root Certification Authorities**, and then click **Import**.
5. In the Certificate Import Wizard, click **Next**.
6. Enter the path and file name of the file that you copied to the domain controller, or use the **Browse** button to locate the file. Click **Next**.
7. Confirm that you want to place these certificates in the **Trusted Root Certification Authorities** certificate store by clicking **Next**. Click **Finish**. When you are notified that the certificates imported successfully, click **OK**.
8. Close the Group Policy Management Editor.

The policy is effective immediately, but the client computers must be restarted to receive the new settings, or you can type **gpupdate /force** from an elevated command prompt or from Windows PowerShell.

Registry settings modified

The settings described in this document configure the following registry keys on the client computers. These settings are not automatically removed if the GPO is unlinked or removed from the domain. These settings must be specifically reconfigured, if you want to change them.

Registry keys	**Value and Description**
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\AuthRoot\DisableRootAutoUpdate	A value of 1 disables the Windows AutoUpdate of the trusted CTL.
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\SystemCertificates\AuthRoot\EnableDisallowedCertAutoUpdate	A value of 1 enables the Window

Registry keys	Value and Description
	s AutoUpdate of the untrusted CTL.
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SystemCertificates\AuthRoot\AutoUpdate\RootDirUrl	Configures the shared location (the HTTP or the FILE path).

New Certutil Options

The following options were added to Certutil:

Syntax	Description	Example

CertUtil [Options] -syncWithWU DestinationDir	Sync with Windows Update.
	• DestinationDir is the folder that receives the files by using the automatic update mechanism.
	• The following files are downloaded by using the automatic update mechanism:
	  • The authrootstl.cab contains the CTLs of non-Microsoft root certificates.
	  • The disallowedcertstl.cab contains the CTLs of
		CertUtil -syncWithWU \\server1\PKI\CTLs

Syntax	Description	Example
	untrusted certificates. • The disallowedcert.sst contains the serialized certificate store, including the untrusted certificates. • <thumbprint>.crt contains the non-Microsoft root certificates.	
CertUtil [Options] -generateSSTFromWU SSTFile	Generate SST by using the automatic update mechanism. SSTFile: .sst file to be created. The generated .sst file contains the non_Microsoft root certificates that were downloaded by using the automatic update mechanism.	CertUtil –generateSSTFromWU TRoots.sst

Tip

`Certutil -SyncWithWU -f <folder>` updates existing files in the target folder.

`Certutil -syncWithWU -f -f <folder>` removes and replaces files in the target folder.

Potential errors with Certutil -SyncWithWU

You may encounter the following errors and warnings when running the `Certutil -syncWithWU` command:

- If you use a non-existent local path or folder as the destination folder, you will see the error:

  The system cannot find the file specified. 0x80070002 (WIN32: 2 ERROR_FILE_NOT_FOUND)
- If you use a non-existent or unavailable network location as the destination folder, you will see the error:

  The network name cannot be found. 0x80070043 (WIN32: 67 ERROR_BAD_NET_NAME)
- If your server cannot connect over TCP port 80 to Microsoft Automatic Update servers, you will receive the following error:

  A connection with the server could not be established 0x80072efd (INet: 12029 ERROR_INTERNET_CANNOT_CONNECT)
- If your server is unable to reach the Microsoft Automatic Update servers with the DNS name ctldl.windowsupdate.com, you will receive the following error:

  The server name or address could not be resolved 0x80072ee7 (INet: 12007 ERROR_INTERNET_NAME_NOT_RESOLVED).

- If you do not use the **-f** switch, and any of the CTL files already exist in the directory, you will receive a *file* exists error:

  CertUtil: -syncWithWU command FAILED: 0x800700b7 (WIN32/HTTP: 183 ERROR_ALREADY_EXISTS) Certutil: Cannot create a file when that file already exists.
- If there is a change in the trusted root certificates, you will see: "Warning! Encountered the following no longer trusted roots: <folder path>\<thumbprint>.crt. Use "-f -f" options to force the delete of the above ".crt" files. Was "authrootstl.cab" updated? If yes, consider deferring the delete until all clients have been updated."

Related content

- How to Write a Simple .Adm File for Registry-based Group Policy
- Writing Custom ADM Files for System Policy Editor
- Managing Group Policy ADMX Files Step-by-Step Guide
- Windows Root certificate Certificate Program - Members List (All CAs)
- Controlling the Update Root certificate Certificates Feature to Prevent the Flow of Information to and from the Internet
- Windows Server Security Forum
- Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
- Windows PKI Documentation Reference and Library
- Windows PKI Blog

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Certification Authority Web Enrollment Guidance

The Certification Authority (CA) Web Enrollment role service provides a set of web pages that allow interaction with the Certification Authority role service. These web pages are located at https://<servername>/certsrv, where <servername> is the name of the server that hosts the hosts the CA Web Enrollment pages. The certsrv portion of the URL should always be in lowercase letters; otherwise, users may have trouble checking and retrieving pending certificates.

Note
The CA Web Enrollment role service pages require that you secure them with secure sockets layer (SSL) / transport layer security (TLS)> If you do not, you will see an error: "In order to complete the certificate enrollment, the Web site for the CA must be configured to use HTTPS authentication." To resolve this issue, you must configure

HTTPS authentication, which is discussed in the TechNet Wiki article: Active Directory Certificate Services (AD CS): Error: "In order to complete certificate enrollment, the Web site for the CA must be configured to use HTTPS authentication".

The CA Web Enrollment role service pages allow you to connect to the CA by using a web browser and performing common tasks, such as:

- Requesting certificates from the CA.
- Requesting the CA's certificate.
- Submitting a certificate request by using a PKCS #10 file.
- Retrieving the CA's certificate revocation list (CRL).

CA Web Enrollment is useful when you interact with a stand-alone CA because the Certificates Microsoft Management Console (MMC) snap-in cannot be used to interact with a stand-alone CA. Enterprise CAs can accept certificate requests through the Certificates snap-in or the CA Web Enrollment role service pages.

Starting in Windows Server® 2008, the CA Web Enrollment role service includes updated sample web pages for web-based certificate enrollment operations. These web pages are updated to work together with the CertEnroll component (available starting with Windows Vista). These web pages also work together with Xenroll.

The certificate enrollment Web pages starting in Windows Server 2008 detect the client operating system and then select the appropriate control.

- If a client computer is running Windows Server 2003 or Windows XP, the certificate enrollment web pages use Xenroll.
- If the client computer is running at least Windows Vista® or Windows Server 2008, the CA Web Enrollment role service uses CertEnroll.

Important

In Windows® 8, CA Web Enrollment pages will work only with Internet Explorer 10 for the desktop.

Starting in Windows Server 2012 R2, client computers that run Windows XP are not supported for web enrollment.

For more information about CertEnroll and Xenroll, see the following:

- How to use Certificate Services Web enrollment pages together with Windows Vista or Windows Server 2008
- Certificate-Related Changes for Windows Vista
- Certificate Enrollment API

CA for Web Enrollment

You can install CA Web Enrollment on a server that is not a CA to separate web traffic from the CA. Installing CA Web Enrollment configures the computer as an enrollment registration authority. You must select a CA to be used with the CA Web Enrollment pages. The CA that CA Web Enrollment uses is called the **Target CA** in the user interface. You can select the target CA

by using the CA name or the computer name that is associated with the CA. Click the **Select** button to locate the CA that you want to use.

Web Enrollment Configuration

If you install the CA Web Enrollment pages on a computer that is not the target CA, the computer account where the CA Web Enrollment pages are installed must be trusted for delegation. See the following resources for more information:

- How to configure the Windows Server 2008 CA Web Enrollment Proxy
- Install Web Enrollment Support on Another Computer (Optional)

Tip

If CA Web Enrollment pages installation fails on a migrated CA, it could be that the setup status in the registry is incorrectly set. For more information, see Certification Authority Web Enrollment Configuration Failed 0x80070057 (WIN32: 87)

Use the CA Web Enrollment pages

If you have been granted access permissions, you can perform the following tasks from the CA Web Enrollment pages:

- Request a basic certificate.
- Request a certificate with advanced options.

  This gives you greater control over the certificate request. Some of the user-selectable options that are available in an advanced certificate request include:

  - **Cryptographic service provider (CSP) options**. The name of the cryptographic service provider, the key size (1024, 2048, and so on), the hash algorithm (such as SHA/RSA, SHA/DSA, MD2, or MD5) and the key specification (exchange or signature).
  - **Key generation options**. Create a new key set or use an existing key set, mark the keys as exportable, enable strong key protection, and use the local computer store to generate the key.
  - **Additional options**. Save the request to a PKCS #10 file or add specific attributes to the certificate.
- Check a pending certificate request. If you have submitted a certificate request to a stand-alone certification authority, you need to check the status of the pending request to see if the certification authority has issued the certificate. If the certificate has been issued, it will be available for you to install it.
- Retrieve the certification authority's certificate to place in your trusted root store or install the entire certificate chain in your certificate store.
- Retrieve the current base and delta CRLs.
- Submit a certificate request by using a PKCS #10 file or a PKCS #7 file.

  Note

In general, you use a PKCS #10 file to submit a request for a new certificate and a PKCS #7 file to submit a request to renew an existing certificate. Submitting requests with files is useful when the certificate requester is unable to submit a request online to the certification authority.

Notes

- You might need to make https://servername a trusted site for Internet Explorer to browse for a file on the computer's hard disk drive. To make https://servername a trusted site, in Internet Explorer, click **Tools**, then point to **Internet Options**, point to **Security**, point to **Trusted Sites**, and click **Sites**. Type **https://***<servername>*, and click **OK**. Replace <servername> with the actual host name of the server to which you want to connect. If you typically use the fully qualified domain name (FQDN) to connect to the server, create your entry by using that instead or in addition to the host name.
- If you submit the request, and you immediately get a message that asks if you want to submit the request even though it does not contain a BEGIN or END tag, click **OK**.

Request a basic certificate

To use Internet Explorer to request a basic certificate

1. In Internet Explorer, connect to https://<servername>/certsrv, where <servername> is the host name of the computer running the CA Web Enrollment role service.
2. Click **Request a certificate**.
3. On **Request a Certificate**, click **User Certificate**.
4. On the **User Certificate Identifying Information** page, do one of the following:
   - Comply to the message "No further identifying information is required. To complete your certificate, press Submit."
   - Enter your identifying information for the certificate request.
5. (Optional) Click **More Options** to specify the cryptographic service provider (CSP) and choose if you want to enable strong private key protection. (You receive a prompt every time you use the private key that is associated with the certificate.)
6. Click **Submit**.
7. Do one of the following:
   - If you see the **Certificate Pending** page, the CA administrator will have to approve the request before you can retrieve and install the certificate.
   - If you see the **Certificate Issued** page, click **Install this certificate**.

Request a certificate with advanced options

To use Internet Explorer to create an advanced certificate request

1. In Internet Explorer, connect to https://<servername>/certsrv, where <servername> is the host name of the computer running the CA Web Enrollment role service.

2. Click **Request a certificate**.
3. Click **Advanced certificate request**.
4. Click **Create and submit a certificate request to this CA**.
5. Fill in the requested identifying information and other options that you require.
6. Click **Submit**.
7. Do one of the following:
   - If you see the **Certificate Pending** page, the CA administrator will have to approve the request before you can retrieve and install the certificate.
   - If you see the **Certificate Issued** page, click **Install this certificate**.

Check a pending certificate request

To check a pending certificate request using Internet Explorer

1. In Internet Explorer, open https://<servername>/certsrv, where <servername> is the hostname of the computer running the CA Web Enrollment role service.
2. Click **View the status of a pending certificate request**.
3. If there are no pending certificate requests, you will see a message to that effect. Otherwise, select the certificate request that you want to check, and click **Next**.
4. Check the following pending certificate requests:
   - **Still pending**. You must wait for the administrator of the certification authority to issue the certificate. To remove the certificate request, click **Remove**.
   - **Issued**. To install the certificate, click **Install this certificate**.
   - **Denied**. Contact the administrator of the certification authority for further information.

Retrieve the CA certificate

To retrieve a CA certificate by using Internet Explorer

1. In Internet Explorer, connect to https://<servername>/certsrv, where <servername> is the name of the computer running the CA Web Enrollment role service.
2. Click **Download a CA certificate, certificate chain, or CRL**.
3. Do one of the following:
   - If you want to trust all the certificates that are issued by this CA, click **Install this CA certificate chain**.
   - If the CA has been renewed, you have the choice of which version of the CA certificate you want to download.
4. Select the encoding method that you want to use for the CRL: **DER** or **Base 64**.
5. Under **CA Certificate**, click the CA certificate that you want to download, and then click **Download CA certificate** or click **Download CA certificate chain**.
6. In **File Download**, click **Open this file from its current location**, and then click **OK**.

7. When the **Certificate** dialog box appears, click **Install this certificate**.
8. In the Certificate Import Wizard, click **Automatically select the certificate store based on the type of certificate**.

Retrieve the current base and delta CRLs

To retrieve a certificate revocation list by using Internet Explorer

1. In Internet Explorer, connect to https://<servername>/certsrv, where <servername> is the name of the computer running the CA Web Enrollment role service.
2. Click **Download a CA certificate, certificate chain, or CRL**.
3. Click the encoding method that you want to use for the CRL, **DER** or **Base 64**.
4. Do one of the following:
   - Click **Download CA certificate**.
   - Click **Download CA certificate chain**.
   - Click **Download latest base CRL**.
   - Click **Download latest delta CRL**.

     Note
     The latest base CRL must already be installed for the delta CRL to function.
5. When the **File Download** dialog box appears, click **Save**. Select a folder on your computer to store the .crl file, and then click **Save**.
6. Open Windows Explorer and locate the .crl file you just saved.
7. Right-click the .cer or .crl file and click **Install Certificate** or **Install CRL**, and then click **Next**.
8. When the Certificate Import Wizard opens, click **Automatically select the certificate store based on the type of certificate**.

Submit a certificate request by using a PKCS #10 file or a PKCS #7 file

To submit a certificate request by using a PKCS #10 or PKCS #7 file by using Internet Explorer

1. In Internet Explorer, connect to https://<servername>/certsrv, where <servername> is the name of the computer running the CA Web Enrollment role service.
2. Click **Request a certificate**, and then click **Advanced certificate request**.
3. Click **Submit a certificate request using a base-64-encoded CMC or PKCS #10 file, or submit a renewal request by using a base-64-encoded PKCS #7 file**.
4. In Notepad, click **File**, click **Open**, select the PKCS #10 or PKCS #7 file, click **Edit**, click **Select all**, click **Edit**, and then click **Copy**. On the Web page, click the **Saved request**

scroll box. Click **Edit**, and then click **Paste** to paste the contents of certificate request into the scroll box.

5. If you are connected to an enterprise CA, choose the certificate template that you want to use. By default, the appropriate template is named **Subordinate Certification Authority**.
6. If you have any attributes to add to the certificate request, enter them into **Additional Attributes**.
7. Click **Submit**.
8. Do one of the following:
   a. If you see the **Certificate Pending** web page, see Check a pending certificate request earlier in this document.
   b. If you see the **Certificate Issued** web page, click **Download certificate chain**. Choose to save the file to your hard disk drive, and then import the certificate into your certificate store.

Related content

- You cannot download CA certificate from web enrollment pages
- AD CS: Web Enrollment
- Windows Server Security Forum
- Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
- Windows PKI Documentation Reference and Library
- Windows PKI Blog

Certificate Enrollment Web Service Guidance

The Certificate Enrollment Web Service is an Active Directory Certificate Services (AD CS) role service that enables users and computers to perform certificate enrollment by using the HTTPS protocol. Together with the Certificate Enrollment Policy Web Service, this enables policy-based certificate enrollment when the client computer is not a member of a domain or when a domain member is not connected to the domain.

The Certificate Enrollment Web Service uses the HTTPS protocol to accept certificate requests from and return issued certificates to network client computers. The Certificate Enrollment Web Service uses the DCOM protocol to connect to the certification authority (CA) and complete certificate enrollment on behalf of the requester. In versions of AD CS prior to Windows Server 2008 R2, policy-based certificate enrollment can be completed only by domain member client computers that are using the DCOM protocol. This limits certificate issuance to the trust boundaries that are established by Active Directory domains and forests.

Certificate enrollment over HTTPS enables the following new deployment scenarios:

- Certificate enrollment across forest boundaries to reduce the number of CAs in an enterprise

- Extranet deployment to issue certificates to mobile workers and business partners

For more information about the Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service, see Certificate Enrollment Web Services. The remaining sections of this document provide the installation requirements for Certificate Enrollment Web Service and information about the configuration options that are presented when you use Server Manager to install the role service.

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Installation requirements

The requirements for installing the Certificate Enrollment Web Service are:

- The administrator who performs the installation must be a member of the Enterprise Admins group.
- The administrator who installs the Certificate Enrollment Web Service must have Request Certificates permissions on the target certification authority (CA).
- The computer on which the Certificate Enrollment Web Service is to be installed must be a member of the domain and must be running Windows Server 2008 R2 or Windows Server 2012.
- An AD DS forest with at least a Windows Server 2008 R2 schema. For more information, see Prepare a Windows 2000 or Windows Server 2003 Forest Schema for a Domain Controller That Runs Windows Server 2008 or Windows Server 2008 R2.
- An enterprise certification authority (CA) on a computer running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.
    - If the Certificate Enrollment Web Service is configured for client certificate authentication, the CA must be running at least Windows Server 2008.
    - For enrollment across AD DS forests, the CA must be installed on a computer running Windows Server 2008 R2 Enterprise, Windows Server 2008 R2 Datacenter, or Windows Server 2012.
    - For automatic renewal of certificates across AD DS forests or from computers that are not part of an AD DS forest or domain, the CA and Certificate Enrollment Web Services servers must be running Windows Server 2012.
- Client computers must be running at least Windows 7 or Windows Server 2008 R2. For automatic renewal of certificates across AD DS forests or from computers that are not part of an AD DS forest or domain, the CA and Certificate Enrollment Web Services clients must be running at least Windows 8 or Windows Server 2012.
- A Server Authentication certificate installed for HTTPS.

Note

The Web Server (IIS) role service with the Microsoft .NET Framework are automatically added during the Certificate Enrollment Web Services installation, if they are not already installed.

Configure a CA for the Certificate Enrollment Web Service

If the CA role service is installed on the local computer, then the local computer is automatically selected as the CA. However, the Certificate Enrollment Web Service and CA role service cannot be installed at the same time. If you intend to install both the Certificate Enrollment Web Service and CA role service, complete the CA role service installation first. The following requirements apply to configuring the CA for the Certificate Enrollment Web Service:

- The Certificate Enrollment Web Service can be configured to work with an enterprise CA on the same or on a different computer. The CA must be on a computer running at least Windows Server 2003.
- The Certificate Enrollment Web Service cannot be configured to work with a stand-alone CA—an enterprise CA is required.
- If client certificate authentication is used, the CA must be on a computer running at least Windows Server 2008. A CA on a computer running Windows Server 2003 will not work as the targeted CA of an enrollment service that is configured for client certificate authentication.
- Running the enrollment service in renewal-only mode requires a CA on a computer running at least Windows Server 2008 R2.

Notes

- If you want to avoid having to trust the Certificate Enrollment Web Service account for delegation, and you need to process only certificate renewal requests, you can enable renewal-only mode. To do so, select **Configure the Certificate Enrollment Web Service for renewal-only mode**.
- If you want the Certificate Enrollment Web Service to process new certificate enrollment requests and certificate renewals, do not select **Configure the Certificate Enrollment Web Service for renewal-only mode**. Also, ensure the Certificate Enrollment Web Service account is trusted for delegation, as explained in Configure a Service Account.
- You can install multiple instances of the Certificate Enrollment Web Service on a single computer. However, you can only install one instance by using Server Manager. To install a second instance, you must use Windows PowerShell as described in Install-AdcsEnrollmentWebService.

Warning

If the certification authority that the Certificate Enrollment Web Service will be using has spaces in the name, such as Margies Travel Issuing CA, instead of Margies-Travel-Issuing-CA, then additional configuration steps are required after installation of the service. The additional steps required are documented in the article Implementing Certificate Enrollment Web Services in Windows Server® 2012 that uses an Issuing CA with spaces in the name.

Set the authentication type for Certificate Enrollment Web Service

Clients communicating with the Certificate Enrollment Web Service must use one of the following authentication types:

- Windows integrated authentication, also known as Kerberos authentication
- Client certificate authentication, also known as X.509 certificate authentication
- User name and password authentication

Notes

- If you want to enable key-based renewal, you must enable client certificate authentication for the Certificate Enrollment Web Service.
- Anonymous authentication to the web services is not supported.

Allow key-based renewal for Certificate Enrollment Web Service

Key-based renewal mode is a feature introduced in Windows Server 2012 that allows an existing valid certificate to be used to authenticate a certificate renewal request. This enables computers that are not connected directly to the internal network the ability to automatically renewal an existing certificate. To take advantage of this feature, the certificate client computers must be running at least Windows 8 or Windows Server 2012.

Configure a Service Account

During Certificate Enrollment Web Service configuration, you have the option to specify one of the following types of accounts as the service account:

- User account that acts as the service account (recommended)
- Built-in application pool identity of the Internet Information Services (IIS) installation on the local computer

Using a specific user account as the service account is the recommended configuration. However, there are more steps required to configure a user account as a service account. The user account that is to be configured as a service account must be:

- A domain account in the domain in which the Certificate Enrollment Web Services computer is a member.
- A member of the local IIS_IUSRS group on which the Certificate Enrollment Web Service is installed.
- Configured with a service principal name (SPN), if Kerberos authentication is selected or delegation required because the Certificate Enrollment Web Service is hosted on a different computer than the one running the CA role service.

- Trusted for delegation for the host service and the Remote Procedure Call system service (RPCSS) , if Certificate Enrollment Web Service is installed on a different computer than the CA, and if new certificates are to be issued by the Certificate Enrollment Web Service.

To create a domain user account to act as the service account

1. Sign in to the domain controller or administrative computer with Active Directory Domain Services Remote Server Administration Tools installed. Open **Active Directory Users and Computers** by using an account that has permissions to add users to the domain.
2. In the console tree, locate the container where you want to create the user account. For example, some organizations have a Services OU or similar account. Right-click the container, click **New**, and then click **User**.
3. In the **New Object - User** text boxes, enter appropriate names for all the fields so that it is clear that you are creating a user account. Be sure to follow your organization's policy for creating a service account, if such a policy exists. As an example, you could enter the following, and then click **Next**.
   a. **First name**: **CES**
   b. **Last name**: **Service**
   c. **User logon name**: **CES**
4. Set a complex password for the account and confirm the password. Configure the password options to correspond to your organization's security policies regarding service accounts.
5. Click **Next**, and then click **Finished**.

Tip

You can also use the New-ADUser Windows PowerShell cmdlet to add a domain user account.

To add the service account to the local IIS_IUSERS group

1. On the server that is hosting Certificate Enrollment Web Service, open **Computer Management** (compmgmt.msc).
2. In the Computer Management console tree, under **System Tools**, expand **Local User and Groups**, and then click **Groups**.
3. In the details pane, double-click **IIS_IUSRS**.
4. On the **General** tab, click **Add**.
5. In the **Select Users, Computers, Service Accounts, or Groups** text box, type the user sign-in name for the account that you configured to be the service account.
6. Click **Check Names**, click **OK** twice, and then close **Computer Management**.

Tip

You can also type `net localgroup IIS_IUSRS <domain>\<username> /Add` to add the service account for the Certificate Enrollment Web Service to the local IIS_IUSRS group.

The command prompt or Windows PowerShell must be run as an administrator. For more information, see Add a member to a local group.

To set a service principal name for the service account

1. Ensure that you are using an account that is a member of the Domain Admins group. Open Windows PowerShell or a Command Prompt window as an administrator.
2. Use the following command syntax to register the server principal name (SPN) for the service account: `setspn -s http/<computername> <domainname>\<accountname>`. For example, to register a service account with the sign-in name CES in the cpandl.com domain for a computer named CES1, you would run the following command: `setspn -s http/CES1.cpandl.com cpandl\CES`.

To configure the Certificate Enrollment Web Service user account for constrained delegation

1. Sign in to the domain controller or administrative computer with Active Directory Domain Services Remote Server Administration Tools installed. Open **Active Directory Users and Computers** by using an account that has permissions to add users to the domain.
2. In the console tree, expand the structure until you see the container that holds the user account. Click that container.
3. In the details pane, right-click the user account that is the service account for the Certificate Enrollment Web Service, and then click **Properties**.
4. On the **Delegation** tab, select **Trust this user for delegation to specified services only**.

   Note
   The **Delegation** tab is only available in the user account properties after you have created an SPN for the user account.
5. Ensure that **Use Kerberos only** is selected (if the authentication type was set to Windows integrated authentication during installation) or **Use any authentication protocol** (if the authentication type was set to Client certificate authentication during installation), and then click **Add**.
6. In the **Add Services** dialog box, click **Users or Computers**.
7. In the **Select Users or Computers** dialog box, type the name of the computer that is hosting the CA. Click **Check Names**, and then click **OK**.
8. In the **Add Services** dialog box, press the CRTL key, and then click both of the following services:
   - **HOST**
   - **rpcss**

   Click **OK** twice.
9. Close Active Directory Users and Computers.

If you specified the default application pool instead of a user account to act as the service account for Certificate Enrollment Web Service, you must trust the computer account on which the Certificate Enrollment Web Service is installed for delegation.

Important

You need to perform the following procedure only if you selected **Use the built-in application pool identity** when you specified the service account for the Certificate Enrollment Web Service.

To configure the Certificate Enrollment Web Service computer account for constrained delegation

1. Sign in to the domain controller or administrative computer with Active Directory Domain Services Remote Server Administration Tools installed. Open **Active Directory Users and Computers** by using an account that has permissions to add users to the domain.
2. In the console tree, expand the structure until you see the container that holds the computer account that is hosting the Certificate Enrollment Web Service. Click that container.
3. In the details pane, right-click the computer account that is hosting the Certificate Enrollment Web Service, and then click **Properties**.
4. On the **Delegation** tab, select **Trust this computer for delegation to the specified services only**.
5. Ensure that **Use Kerberos only** is selected, and then click **Add**.
6. In the **Add Services** dialog box, click **Users or Computers**.
7. In the **Select Users or Computers** dialog box, type the name of the computer that is hosting the Certificate Enrollment Web Service. Click **Check Names**, and then click **OK**.
8. In the **Add Services** dialog box, press the CRTL key, and then click both of the following services
   - **HOST**
   - **rpcss**

   Click **OK** twice.
9. Close Active Directory Users and Computers.

Select a Server Certificate

The Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service must use Secure Sockets Layer (SSL) for communication with clients (by using HTTPS). Each service must have a valid certificate that has an enhanced key usage (EKU) policy of Server Authentication in the local computer certificate store.

Note

If you have not yet provided an SSL certificate to the server that is hosting the Certificate Enrollment Web Service, you can do so by following the instructions in the article Configure SSL/TLS on a Web site in the domain with an Enterprise CA.

Complete Certificate Enrollment Web Services Configuration

If you need to enable delegation, see Configure a Service Account.

If you enabled renewal-only mode, you need to complete the following additional configuration steps.

- Update CA permissions
- Set the CA policy module flag

To Update CA permissions

1. On the CA, sign in as a CA administrator.
2. Open the Certification Authority console.
3. Right-click the CA, and then click **Properties**.
4. On the **Security** tab, click **Add**.
5. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, type the name of service account for the Certificate Enrollment Service. Click **Check Names**, and then click **OK**.

   Note
   If you assigned the built-in application pool as the service account, you will enter the computer account name that is hosting the Certificate Enrollment Web Service.
6. In the **CA Properties** dialog box, under **Group or user names**, click the service account name for the Certificate Enrollment Web Service. Ensure that the **Allow** check box is selected for **Read** permission. Clear the **Allow** check box for **Request Certificates** permission (which is selected by default), and then click **OK**.
7. Close the Certification Authority console.

To Set the CA policy module flag

1. On the CA, sign in as a CA administrator.
2. Open a Command Prompt window or Windows PowerShell as an administrator.
3. Run `certutil | findstr "Config"` and note the output of the CA Configuration.
4. Run `certutil -config "<CAConfig>" -setreg policy\EditFlags +EDITF_ENABLERENEWONBEHALFOF`, where *<CAConfig>* is the actual CA Config information that was returned in the previous step.
5. Restart the CA service. To do so from a command prompt, run `net stop certsvc && net start certsvc`. To restart from Windows PowerShell, run `restart-service certsvc`.

Related content

1. Certificate Enrollment Web Services
2. Windows Server Security Forum
3. Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
4. Windows PKI Documentation Reference and Library
5. Windows PKI Blog

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Certificate Enrollment Policy Web Service Guidance

This document provides additional information for the Server Manager configuration pages for the Certificate Enrollment Policy Web Service. For an overview of the service and its installation requirements, see Certificate Enrollment Web Service Guidance. For more information about the Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service, see Certificate Enrollment Web Services.

The remaining sections of this document provide more information for the configuration options that are presented when you use Server Manager to install the Certificate Enrollment Policy Web Service.

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Set the authentication type for Certificate Enrollment Policy Web Service

Clients that communicate with the Certificate Enrollment Policy Web Service must use one of the following authentication types:

- Windows integrated authentication, also known as Kerberos authentication
- Client certificate authentication, also known as X.509 certificate authentication
- User name and password authentication

Notes
- If you want to configure key-based renewal, you must enable user name and password authentication or client certificate authentication.

- Anonymous authentication to the web services is not supported.

Determine whether to enable key-based renewal for Certificate Enrollment Policy Web Service

Key-based renewal mode is a feature introduced in Windows Server 2012 that allows an existing valid certificate to be used to authenticate a certificate renewal request. This enables computers that are not connected directly to the internal network the ability to automatically renew an existing certificate. To take advantage of this feature, the certificate client computers must be running at least Windows 8 or Windows Server 2012.

Note

When key-based renewal mode is enabled for the Certificate Enrollment Policy Web Service, it will not accept requests for new certificates. You can install multiple instances of the Certificate Enrollment Policy Web Service on Windows Server 2012, but you must use the Windows PowerShell Install-AdcsEnrollmentPolicyWebService to install additional instances.

Server Certificate

The Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service must use Secure Sockets Layer (SSL) for communication with clients (by using HTTPS). Each service must have a valid certificate that has an enhanced key usage (EKU) policy of Server Authentication in the local computer certificate store.

Note

If you have not yet provided an SSL certificate to the server that is hosting the Certificate Enrollment Web Service, you can do so by following the instructions in the article Configure SSL/TLS on a Web site in the domain with an Enterprise CA.

Certificate Enrollment Policy Web Service Configuration

After you install the Certificate Enrollment Policy Web Service, there are two additional configuration steps to complete.

1. Configure a friendly name value for the Certificate Enrollment Policy Web Service.
2. Configure Group Policy to enable use of the Certificate Enrollment Policy Web Service.

To configure a friendly name value for the Certificate Enrollment Policy Web Service

1. Open the Internet Information Services (IIS) Manager console.
2. In the **Connections** pane, expand the web server that is hosting the Certificate Enrollment Policy Web Service.

Note

If you are asked to get started with the Microsoft Web Platform, click **No**.

3. Expand **Sites**, expand **Default Web Site**, and then click the appropriate installation virtual application name. The name of the virtual application name varies with the type of installation that you performed. The variation is as follows:

   *KeyBasedRenewal*_**ADPolicyProvider_CEP**_*AuthenticationType*

   For example:

   - **KeyBasedRenewal_ADPolicyProvider_CEP_Certificate** is the virtual application name if you enabled key-based renewal and configured client certificate authentication.
   - **ADPolicyProvider_CEP_UsernamePassword** is the virtual application name if you did not enable key-based renewal and you configured user name and password authentication.
   - **ADPolicyProvider_CEP_Kerberos** is the virtual application name if you did not enable key-based renewal and you configured Windows integrated authentication.
4. In the *virtual application name* Home pane, double-click **Application Settings**, and then double-click **FriendlyName**.
5. In the **Edit Application Setting** dialog box, under **Value**, type the name that you want to configure as a friendly name for the service. For example, you might type **Client Certificate Enrollment** as the friendly name for the service. Click **OK**.
6. In the **Application Settings** pane, double-click **URI**. The value that is shown for URI is significant because that is the path that clients will use to connect to the service. Copy this value, because you will use it when you configure Group Policy. Click **Cancel**.
7. Close the Internet Information Services (IIS) Manager console.

To provide domain client users or their computers with the ability to obtain certificates using Certificate Enrollment Policy Web Services, you can set the URI that you obtained by using the previous procedure. This will allow domain clients to request certificates by using the Certificates console, without the clients having to know the URI to the Certificate Enrollment Policy Web Services virtual application name.

Note

Domain users could input the URI by configuring a custom certificate request, but this is typically not a practical solution because the URI is long and the procedure is complex. However, administrators can perform custom certificate requests to validate the configuration of the Certificate Enrollment Policy Web Service. For more information, see Certificate Enrollment Web Services

To configure Group Policy to enable use of the Certificate Enrollment Policy Web Service

1. Open the Group Policy Management console. To do so, from Server Manager, click **Tools**, and then click **Group Policy Management**.

Note

Ensure that you sign in by using an account with membership in Domain Admins or Enterprise Admins so that you can configure Group Policy settings. You can configure a Group Policy setting for the entire domain, an OU, or (if the account you are using is a member of Enterprise Admins), an entire site. The following instructions assume that you want to set a new Group Policy for the domain.

2. Expand the forest that you want to target for the new Group Policy. Expand **Domains**. Right-click the domain, and then click **Create a GPO in this domain, and link it here**.
3. In the **New GPO** dialog box, under **Name**, type a name that is appropriate for the new Group Policy Object (GPO), for example, **Certificate Enrollment Policy Web Service Certificates**. Click **OK**.
4. Click the linked GPO that you just created. If you see a warning message about Group Policy Management Console, review the message, and then click **OK**.
5. Right-click the linked GPO that you just created, and then click **Edit**.
6. There are two types of certificates that you can distribute by using a GPO: computer certificates or user certificates. The following instructions describe setting the URI for both the **Computer Configuration** and **User Configuration** parts of the GPO. You can set either separately or set them both.
7. To distribute certificates for computers, in the console pane, under **Computer Configuration**, click **Policies**, click **Windows Settings**, click **Security Settings**, and then click **Public Key Policies**.
   a. In the details pane, double-click **Certificate Services Client - Certificate Enrollment Policy**.
   b. Set **Configuration Model** to **Enabled**, and then click **Add**.
   c. In the **Certificate Enrollment Policy Server** dialog box, under **Enter enrollment policy server URI**, enter the URI that you copied in the previous procedure.
   d. In **Authentication type**, set the authentication type that you configured for the Certificate Enrollment Web Policy Service.
   e. Click **Validate Server**, and when the server is validated, click **Add**. Click **OK**.

      Note

      You can only validate the server if you have the appropriate credentials. This could be an issue if you have selected client certificate validation and you do not already have a certificate for the computer. If this is the case, you will first have to obtain a certificate for the computer. You will need a computer certificate with the following characteristics: Enhanced Key Usage Client Authentication 1.3.6.1.5.5.7.3.2.
8. To distribute certificates for users, in the console pane, under **User Configuration**, click **Policies**, click **Windows Settings**, click **Security Settings**, and then click **Public Key Policies**.
   a. In the details pane, double-click **Certificate Services Client - Certificate Enrollment Policy**.

b. Set **Configuration Model** to **Enabled**, and then click **Add**.

c. In the **Certificate Enrollment Policy Server** dialog box, under **Enter enrollment policy server URI**, enter the URI that you copied in the previous procedure.

d. In **Authentication type**, set the authentication type that you configured for the Certificate Enrollment Web Policy Service.

e. Click **Validate Server**, and when the server is validated, click **Add**. Click **OK**.

Note
You can only validate the server if you have the appropriate credentials. This could be an issue if you have selected client certificate validation and you do not already have a certificate for the user. If this is the case, you will first have to obtain a certificate for the user. You will need a user certificate that includes an enhanced key usage (EKU) of Client Authentication with object ID (OID) 1.3.6.1.5.5.7.3.2.

9. Close the Group Policy Management Editor and the Group Policy Management Console.

Related content

1. Certificate Enrollment Web Service Guidance
2. Windows Server Security Forum
3. Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
4. Windows PKI Documentation Reference and Library
5. Windows PKI Blog

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

Network Device Enrollment Service Guidance

The Network Device Enrollment Service (NDES) allows software on routers and other network devices running without domain credentials to obtain certificates based on the Simple Certificate Enrollment Protocol (SCEP).

Note
SCEP was developed to support the secure, scalable issuance of certificates to network devices by using existing certification authorities (CAs). The protocol supports CA and registration authority public key distribution, certificate enrollment, certificate revocation, certificate queries, and certificate revocation queries.

The Network Device Enrollment Service performs the following functions:

1. Generates and provides one-time enrollment passwords to administrators
2. Submits enrollment requests to the CA
3. Retrieves enrolled certificates from the CA and forwards them to the network device

NDES configuration settings

The following sections describe the configuration options that you can select after installing the NDES binary installation files.

Configure a service account for NDES

NDES can be configured to run as either of the following:

- A user account that is specified as a service account
- The built-in application pool identity of the Internet Information Services (IIS) computer

If you select the built-in application pool identity, there is no additional configuration required. However, the recommended configuration is to specify a user account, which requires additional configuration. The user account that is specified as the NDES service account must meet the following requirements:

- Be a domain user account
- Be a member of the local IIS_IUSRS group
- Have Request permissions on the configured CA
- Have Read and Enroll permissions on the NDES certificate template, which is configured automatically
- Have a service principal name (SPN) set in Active Directory

To create a domain user account to act as the NDES service account

1. Sign in to the domain controller or administrative computer with Active Directory Domain Services Remote Server Administration Tools installed. Open **Active Directory Users and Computers** by using an account that has permissions to add users to the domain.
2. In the console tree, expand the structure until you see the container where you want to create the user account. For example, some organizations have a Services OU or similar account. Right-click the container, click **New**, and then click **User**.
3. In the **New Object - User** text boxes, enter appropriate names for all the fields so that it is clear that you are creating a user account. Be sure to follow your organization's policy for creating a service account, if such a policy exists. As an example, you could enter the following, and then click **Next**.
    a. **First name**: **Ndes**
    b. **Last name**: **Service**
    c. **User logon name**: **NdesService**
4. Ensure that you set a complex password for the account and confirm the password.

Configure the password options to correspond to your organization's security policies regarding service accounts. If the password is configured to expire, you should have a process in place to ensure that you reset the password at the required intervals.

5. Click **Next**, and then click **Finished**.

Tip

- You can also use the New-ADUser Windows PowerShell® cmdlet to add a domain user account.
- Depending upon your Active Directory Domain Service (AD DS) configuration, you may be able to implement a Managed Service Account or Group Managed Service Account for NDES. For more information about Managed Service Accounts, see Managed Service Accounts. For more information about Group Managed Service Accounts, see Group Managed Service Accounts Overview.

To add the NDES service account to the local IIS_IUSERS group

1. On the server that is hosting the NDES service, open **Computer Management** (compmgmt.msc).
2. In the Computer Management console tree, under **System Tools**, expand **Local User and Groups**. Click **Groups**.
3. In the details pane, double-click **IIS_IUSRS**.
4. In the **General** tab, click **Add**.
5. In the **Select Users, Computers, Service Accounts, or Groups** text box, type the user sign-in name for the account that you configured to be the service account.
6. Click **Check Names**, click **OK** twice, and then close **Computer Management**.

Tip

You can also use `net localgroup IIS_IUSRS <domain>\<username> /Add` to add the NDES service account to the local IIS_IUSRS group. The command prompt or Windows PowerShell must be run as Administrator. For more information, see Add a member to a local group.

To configure the NDES service account with request permission on the CA

1. On the CA that is to be used by NDES, open the Certification Authority console with an account that has Manage CA permissions.
2. Open the Certification Authority console. Right-click the certification authority, and then click **Properties**.
3. On the **Security** tab, you can see the accounts that have Request Certificates permissions. By default the group **Authenticated Users** has this permission. The service account that you created will be a member of **Authenticated Users** when it is in use. You do not need to grant additional permissions, if **Authenticated Users** has the Request Certificates permission. However, if that is not the case, you should grant the NDES service account Request Certificates permission on the CA. To do so:

- Click **Add**.
- In the **Select Users, Computers, Service Accounts, or Groups** text box, type the name of the NDES service account, and click **Check Names**, and then click **OK**.
- Ensure that NDES service account is selected. Ensure that the **Allow** check box that corresponds to **Request Certificates** is selected. Click **OK**.

To set a service principal name for the NDES service account

1. Ensure that you are using an account that is a member of the Domain Admins group. Open Windows PowerShell or a command prompt as an administrator.
2. Use the following command syntax to register the server principal name (SPN) for the NDES service account: `setspn -s http/<computername> <domainname>\<accountname>`. For example, to register a service account with the sign-in name NdesService in the cpandl.com domain that is running on a computer named CA1, you would run the following command: `setspn -s http/CA1.cpandl.com cpandl\NdesService`

Select a CA for NDES

You must select a CA for the NDES service to use when issuing certificates to clients. If NDES is installed on a CA, you do not have the opportunity to select a CA because the local CA is used. When you install NDES on a computer that is not a CA, you must select the target CA. You can select the CA by the CA name or by the computer name. Click **CA name** or **Computer name**, and then click **Select**. The option you choose will determine the type of dialog box that is presented next.

- If you clicked **CA name**, you will be presented with the **Select Certification Authority** dialog box, which has a list of CAs from which you can choose.
- If you clicked **Computer name**, you see the **Select Computer** dialog box where you can set the **Locations** and enter the computer name that you want to specify as the CA.

Set RA information

On the **RA Information** page, all the required and optional fields for setting up the service as the RA are collected. The information that you provide here will be used to construct the signing certificate that is issued to the service.

Configure cryptography for NDES

The Network Device Enrollment Service uses two certificates and their keys to enable device enrollment. Organizations might want to use different Cryptographic Service Providers (CSPs) to store these keys, or they may want to change the length of the keys that is used by the service. Only Cryptographic Application Programming Interface (CryptoAPI) Service Providers are supported for the RA keys—Cryptography API: Next Generation (CNG) providers are not supported.

Complete NDES configuration

You can learn more about NDES configuration and operation in the following article Network Device Enrollment Service (NDES) in Active Directory Certificate Services (AD CS). on Microsoft TechNet.

If you require over-the-air enrollment for mobile devices, see Using a Policy Module with the Network Device Enrollment Service.

Related content

- You cannot download CA certificate from web enrollment pages
- AD CS: Web Enrollment
- Windows Server Security Forum
- Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)
- Windows PKI Documentation Reference and Library
- Windows PKI Blog

Using a Policy Module with the Network Device Enrollment Service

In Windows Server 2012 R2 the Active Directory Certificate Services (AD CS) Network Device Enrollment Service (NDES) supports a policy module that provides additional security for the Simple Certificate Enrollment Protocol (SCEP). This enhancement lets an organization or mobile device management solution address the issue described in CERT Vulnerability Note VU#971035 titled, “Simple Certificate Enrollment Protocol (SCEP) does not strongly authenticate certificate requests.”

Without this policy module, when a user or a device requests a certificate, the SCEP implementation might require a unique or shared password. Then, to obtain a certificate without the policy module, only the password is required. In addition, the certificate subject name value is provided by the user at request time. Using a legitimately obtained password, a rogue user could request a certificate that has the following security problems:

- The subject name value is for another user and therefore the security risk is impersonation.
- The certificate purpose is changed and therefore the security risk is an elevation of privileges.

When you use a policy module with the Network Device Enrollment Service, this module addresses these security risks by implementing additional authentication. For example, the module can verify that the requested certificate is for a specific user and for a specific purpose, and it can enforce whether to deploy a user certificate or computer certificate.

Windows Server 2012 R2 AD CS Network Device Enrollment Service does not ship with a policy module. You must create it yourself or obtain it as part of a software solution from a vendor. For

example, it might be included in a mobile device management solution. System Center 2012 R2 Configuration Manager provides a policy module that is required to deploy certificate profiles.

Use the following sections to help you understand how the policy module works, your deployment options, and how to install and uninstall the policy module.

- How the policy module works
- Deployment options for the Network Device Enrollment Service and a policy module
- How to install and uninstall the policy module

For more information about using the policy module with Configuration Manager, see Certificate Profiles in Configuration Manager in the System Center 2012 Configuration Manager Documentation Library.

For developer information about how to create your own policy module, see INDESPolicy interface on MSDN.

How the policy module works

The most typical scenario for using a policy module is to support the enrollment of user and computer certificates for mobile devices that use a cloud service. This is sometimes referred to as over-the-air enrollment. The following diagram shows the process flow after a mobile device administrator configures policy for certificate enrollment for a mobile device that is being used by an information worker.

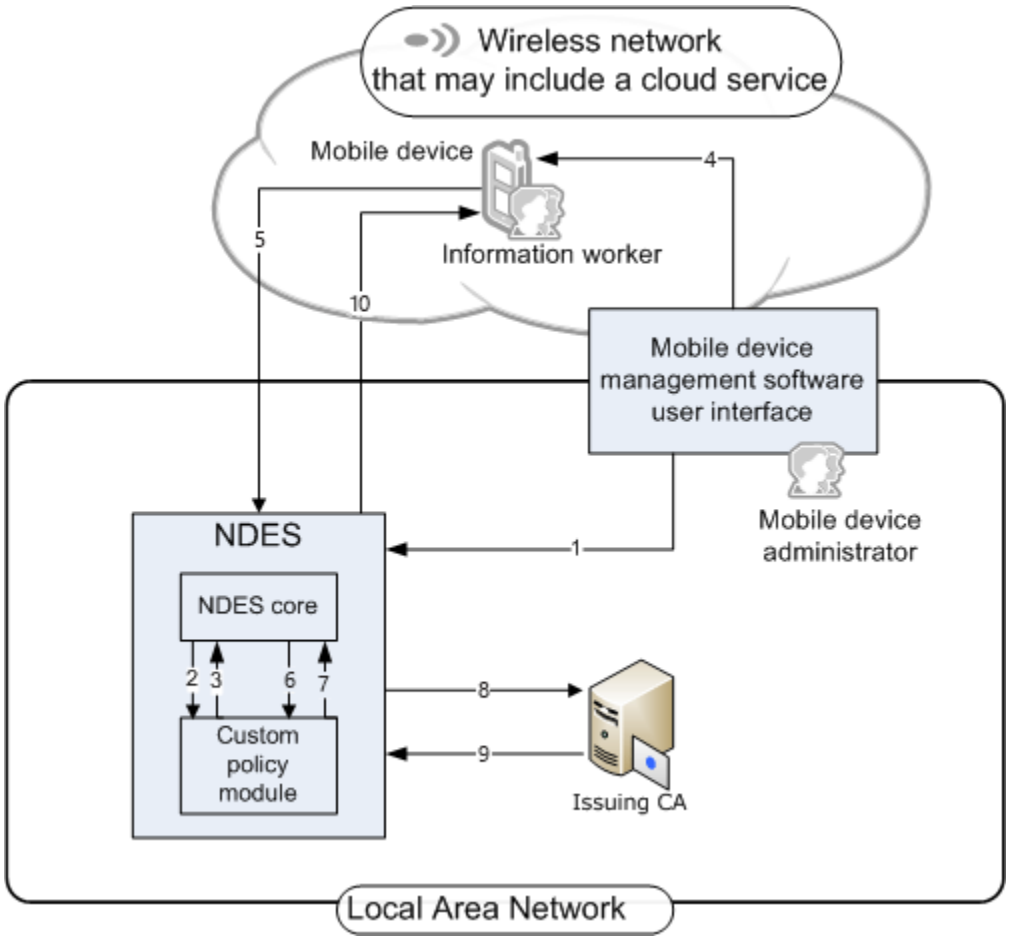

Process flow for a policy module and the Network Device Enrollment Service

1. The mobile device management (MDM) software requests a challenge password from the Network Device Enrollment Service.

   Note
   The mobile device management software contacts the Network Device Enrollment Service by using the NDES mscep_admin interface.

2. The Network Device Enrollment Service delegates the challenge password request to the policy module.
3. The policy module creates a challenge password that requires that the certificate request includes the following items and then sends this instruction to the Network Device Enrollment Service:
   - Specific user
   - Specific purpose
   - Type of certificate (user or computer)

4. When the mobile device management software receives the challenge password, this software sends the uniform resource indicator (URI) for contacting the Network Device Enrollment Service and the challenge password to the mobile device.
5. The mobile device contacts the Network Device Enrollment Service to enroll a certificate.

   Note
   The mobile device contacts the Network Device Enrollment Service by using the NDES mscep interface.
6. The Network Device Enrollment Service delegates the request to the policy module.
7. The policy module verifies the challenge password and certificate request. Then, the policy module returns the result of the verification to the Network Device Enrollment Service.

   If the challenge password and certificate request is not successfully verified by the policy module, the Network Device Enrollment Service returns an error to the mobile device. If the verification is successful, the Network Device Enrollment Service forwards the request to the certification authority.
8. When the request is approved by the certification authority, the certificate is issued to the Network Device Enrollment Service.
9. The Network Device Enrollment Service sends the certificate to the mobile device.

Deployment options for the Network Device Enrollment Service and a policy module

You can choose between the following deployment options when you use the Network Device Enrollment Service and the policy module:

- Deployment in a separate forest
- Deployment in an isolated network
- Deployment on an internal domain

Deployment in a separate forest

This deployment design is the most secure but requires more infrastructure to support additional computers in another forest.

Create a new forest for the server running the Network Device Enrollment Service and the policy module, and the issuing CA. This design creates a security boundary from your internal domain controllers and domain accounts on the intranet, thereby reducing the risk of exposure. See the following diagram for an example of this deployment design that has the following characteristics:

- The root CA is offline and not a domain member.
- The issuing CAs in both forests are subordinate to the offline root CA.
- The server running the Network Device Enrollment Service has the policy module installed.
- The mobile device management software could be deployed in the perimeter network or in the internal network, depending on the technical requirements for the mobile device

management solution. For example, for Configuration Manager, the site server and most site systems are deployed in the internal network, and the certificate registration point in the internal network communicates with the policy module on the server running the Network Device Enrollment Service in the perimeter network.

Notes

When you install the Network Device Enrollment Service on a domain member, you must install an issuing CA in the same domain.

If you install the Network Device Enrollment Service on a workgroup computer, you must install the issuing CA on the same computer.

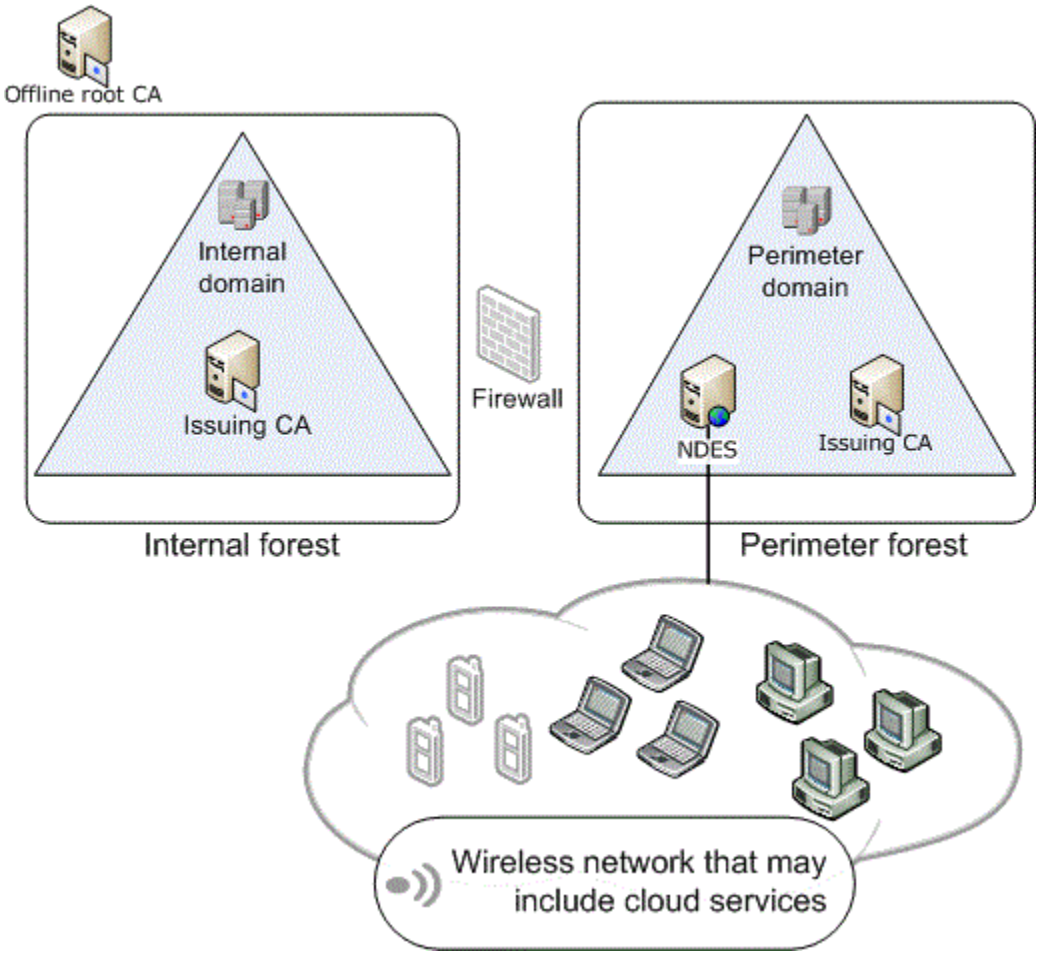

This design provides protection for the internal network by isolating the domain and forest accounts that are used on the perimeter network and the internal network. Because the issuing CA on the perimeter network is subordinate to the root CA, the certificates that are issued in the perimeter network can be trusted on the internal network. If the perimeter network is compromised, you can use the root CA to revoke the certificate of the issuing CA in the perimeter

network, which invalidates all certificates from the issuing CA in the perimeter network. The perimeter network can then be rebuilt and new certificates issued.

Deployment in an isolated network

This deployment design offers a compromise between the most and least secure designs. It requires some infrastructure changes and configuration but without a separate forest.

This deployment design requires a single computer to run the Network Device Enrollment Service and the mobile device management solution. This computer is a member of the internal domain, so you do not have to install a separate issuing CA. The computer running the Network Device Enrollment Service and the mobile device management solution is attached by using a VPN or DirectAccess connection, which provides the isolated network environment. Mobile devices authenticate to the mobile device management software by using a user name and password before they obtain a certificate that then lets them access resources on the internally secured wireless network. See the following diagram for an example of this deployment design.

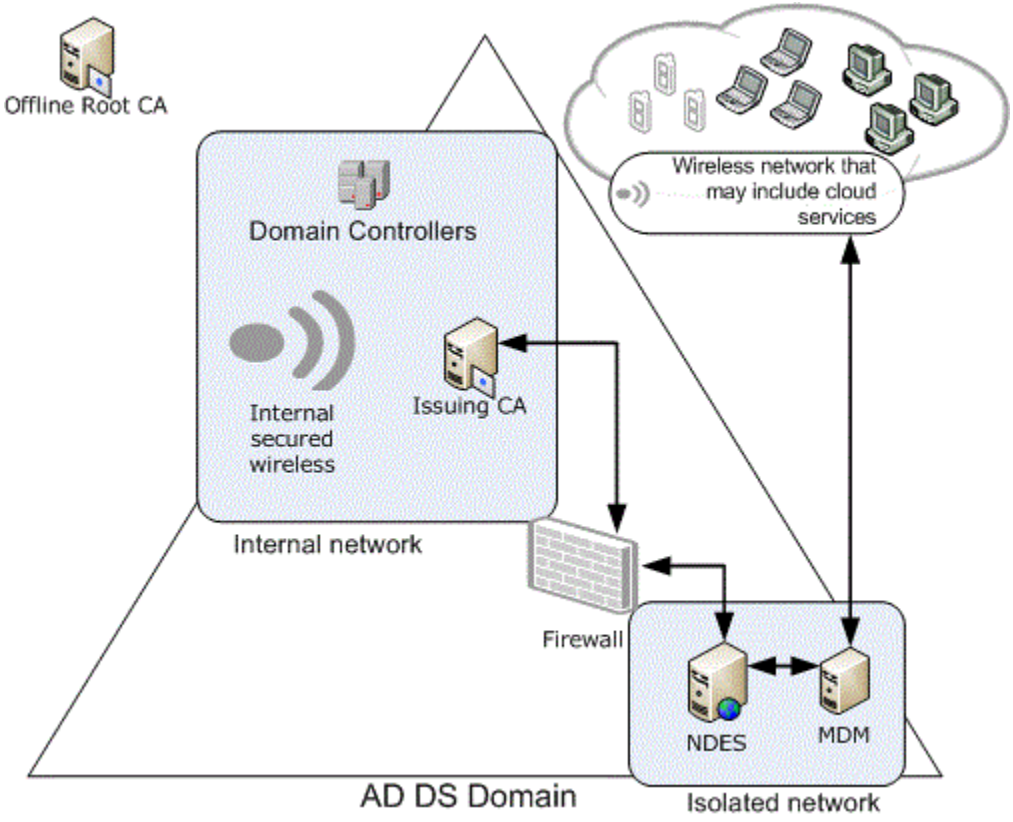

In this design, the server running the Network Device Enrollment Service and the mobile device management solution can be at risk from attackers to compromise the internal network. To help mitigate this risk, install and configure the server running the Network Device Enrollment Service in the internal domain by using a temporary Domain Admin or Enterprise Admin account, then

delete this account, and then move the server to the isolated network. Use additional security controls to protect this server and monitor it carefully for suspicious activity or signs that it might be compromised.

Deployment on an internal domain

This deployment design is the least secure but lets you deploy this enrollment solution by using your existing internal domain structure.

In this design, there is no isolation for the server running the Network Device Enrollment Service and it is a member of the internal domain. The mobile device management software must authenticate all the requests for certificates from the wireless network. For example, this authentication might be to prompt for user name and password. See the following diagram for an example of this deployment design.

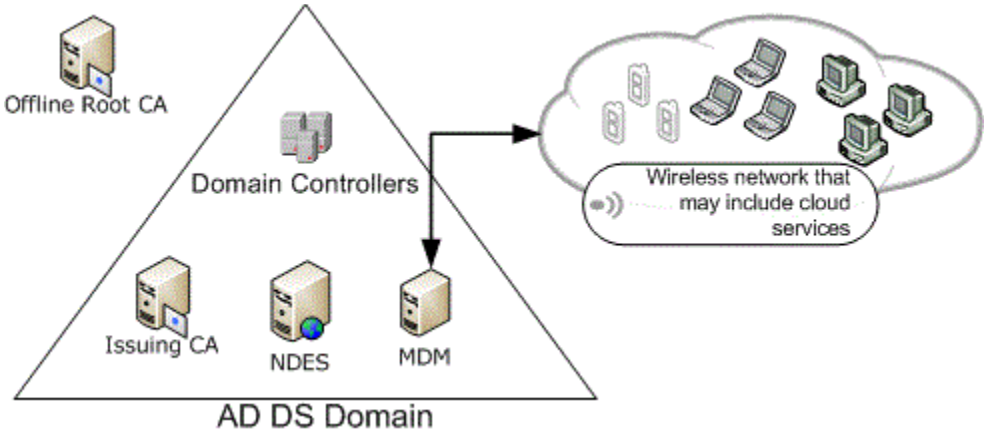

This deployment design does not use a security boundary when the wireless network that is used to issue certificates is connected to the internal domain. In this scenario, the external wireless network could be used directly by an attacker in an attempt to compromise the internal domain. If you use this design, all computers on the internal network become potential targets for attack if the wireless network is compromised, so make sure that you use additional security controls to protect the entire internal network and all computers, and monitor them carefully for suspicious activity or signs that they might be compromised.

How to install and uninstall the policy module

Before you can install the policy module, you must first install and configure the Network Device Enrollment Service, and verify that this can communicate successfully with an issuing certification authority. For more information, see Network Device Enrollment Service Guidance.

If you have obtained a software solution that includes the policy module, follow the installation and configuration instructions that accompany this software rather than the instructions in this section. For example, System Center 2012 R2 Configuration Manager provides its own wizard,

the Configuration Manager Policy Module Setup wizard, to install and configure the policy module.

If there are no accompanying instructions, or you have developed your own policy module, use the following procedures.

To install and configure the policy module

1. Copy the .dll file for the policy module onto the server that runs the Network Device Enrollment Service.
2. Register the policy module by using the Regsvr32.exe command from an elevated command prompt.

   For example, if your policy module is named ndespol.dll, type **Regsvr32 ndespol.dll** and press ENTER.
3. Run the following Windows PowerShell cmdlets to update the registry path HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography\MSCEP:

   ```
   Item -Path
   HKLM:\SOFTWARE\Microsoft\Cryptography\MSCEP\Modules
   ```

   This command creates the new \Modules key.

   ```
   New-ItemProperty
   HKLM:\SOFTWARE\Microsoft\Cryptography\MSCEP\Modules\ -Name
   Policy -PropertyType String -Value "NDESPolicy.OTA.1"
   ```

   This command specifies the policy module name, using *NDESPolicy.OTA.1* as an example. Replace *NDESPolicy.OTA.1* with the name of your own policy module.
4. Restart Internet Information Services (IIS) by typing **IISReset** from an elevated command prompt.

The Network Device Enrollment Service supports only one registered policy module and after it is registered and configured, you cannot choose whether to use it with the Network Device Enrollment Service; it is always used. To remove the policy module, you must manually uninstall it by updating the registry and unregistering the policy module.

Important
Removing the Network Device Enrollment Service role service does not uninstall the policy module.

To uninstall the policy module

1. Run the following Windows PowerShell cmdlet to remove the registry settings:

   ```
   Remove-Item -Path
   HKLM:\SOFTWARE\Microsoft\Cryptography\MSCEP\Modules -Recurse
   ```
2. From an elevated command prompt, run the regsvr32.exe command with the /u option to unregister the policy module.

   For example, if the policy module name is ndespol.dll, type **regsvr32 ndespol.dll /u** and

press ENTER.

Related content

Windows Server Security Forum

Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)

Windows PKI Documentation Reference and Library

Windows PKI Blog

Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy

The purpose of this Test Lab Guide (TLG) is to enable you to create a two-tier public key infrastructure (PKI) hierarchy using Windows Server® 2012 and Active Directory Certificate Services (AD CS).

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

In this guide

This document contains instructions for extending the Windows Server 2012 Base Configuration Test Lab Guide (TLG) to include an offline root certification authority and install an online enterprise subordinate certification authority on the computer APP1 from the Base Configuration TLG. In this guide you will deploy a two-tier PKI hierarchy, configure a certificate revocation list (CRL) distribution point (CDP), automatically deploy certificates to the domain, and utilize a certificate to enable Secure Sockets Layer (SSL) communication with the APP1 web site.

Important
The configuration of the computers and network in this guide was designed to give you hands-on practice in creating a two-tier certification authority PKI hierarchy. The design decisions made in this guide were geared toward increasing your hands-on experience and do not reflect a best practices configuration. For best practice information, see Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure (http://technet.microsoft.com/library/cc772670.aspx) and PKI Design Brief Overview (http://social.technet.microsoft.com/wiki/contents/articles/pki-design-brief-overview.aspx).

Test lab overview

The test lab configuration demonstrated in this guide extends the Windows Server 2012 or Windows Server 2012 R2 Base Configuration TLG by one server computer. The additional computer will serve as an offline root CA and be named ORCA1. There are six major steps in this test lab guide to complete that include multiple subordinate procedures.

1. Complete the Base TLG Configuration
2. Configure ORCA1
3. Configure APP1 to distribute certificates and CRLs
4. Configure APP1 as an enterprise subordinate CA
5. Enable certificate auto-enrollment
6. Configure SSL for APP1

AD CS Two Tier PKI Hierarchy Network Configuration

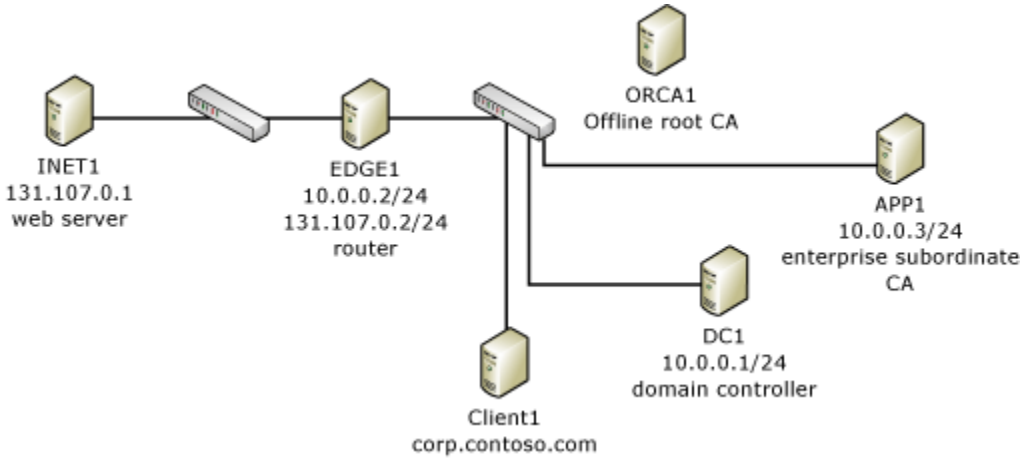

Important

Although EDGE1 and INET1 are pictured in the figure, they are not utilized in the lab.

Hardware and software requirements

The following are the minimum required components of the test lab:

1. The product disc or files for Windows Server 2012 or Windows Server 2012 R2.
2. Three computers that meet the minimum hardware requirements for Windows Server 2012 or Windows Server 2012 R2.

   Note

   You will need only the DC1, APP1, and CLIENT1 computers from the Base Test Lab configuration to complete this lab. You will also build the ORCA1 computer during this lab. As previously mentioned, INET1 and EDGE1 are not utilized in this lab.
3. The product disc or files for Windows® 8 or Windows® 8.1.

4. One computer that meets the minimum hardware requirements for Windows 8 or Windows 8.1.
5. One removable media with enough free space to hold a few certificates and certificate revocation lists (about 10 kilobytes). This can be either physical or virtual removable media depending on whether your lab is using physical or virtual computers.

   Note
   For instructions on transferring files using a virtual floppy disk using Microsoft Windows Server™ Hyper-V, see Creating, Using, and Transferring Files using Virtual Floppy Disks (http://social.technet.microsoft.com/wiki/contents/articles/4272.aspx).
6. If you wish to deploy the Base Configuration test lab in a virtualized environment, your virtualization solution must support Windows Server 2012 or Windows Server 2012 R2 64-bit virtual machines. The server hardware must support the amount of RAM required to run the virtual operating systems included in the Base Configuration test lab and any other virtual machines that may be required by additional TLGs.

Important
Run Windows Update on all computers or virtual machines either during the installation or immediately after installing the operating systems. After running Windows Update, you can isolate your physical or virtual test lab from your production network.

Step 1: Complete the Base TLG Configuration

The Windows Server 2012 Base Configuration Test Lab Guide (TLG) is located at http://go.microsoft.com/fwlink/p/?LinkId=236358.

Tip
See Test Lab Guides for information on the location of other test lab guide files.

Step 2: Configure ORCA1

The procedures to complete the configuration of the offline root CA, named ORCA1, include:

- Install the Operating system
- Rename the computer
- Prepare the CAPolicy.inf for the standalone root CA
- Install the standalone root CA
- Configure the root CA settings
- Copy the root CA certificate and CRL to removable media
- Distribute the root CA via GPO
- Create an internal contoso.com DNS zone and www host record

To install the operating system on ORCA1

1. Do not connect this computer to a network.

2. Start the installation of Windows Server 2012 or Windows Server 2012 R2.
3. Follow the instructions to complete the installation, specifying Windows Server 2012 or Windows Server 2012 R2 (full installation) and a strong password for the local Administrator account. Sign in using the local Administrator account.

To rename the computer

1. Open Windows PowerShell®.
2. Type **rename-computer orca1** and then press ENTER.
3. Type **restart-computer** and then press ENTER.

   After the computer restarts, sign in using the local Administrator account.

To prepare the CAPolicy.inf for the standalone root CA

1. Open Windows PowerShell, type **notepad c:\Windows\CAPolicy.inf** and press ENTER.
2. When prompted to create a new file, click **Yes**.
3. Enter the following as the contents of the file:

```
[Version]
Signature="$Windows NT$"
[PolicyStatementExtension]
Policies=InternalPolicy
[InternalPolicy]
OID= 1.2.3.4.1455.67.89.5
Notice="Legal Policy Statement"
URL=http://www.contoso.com/pki/cps.txt
[Certsrv_Server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=20
CRLPeriod=weeks
CRLPeriodUnits=26
CRLDeltaPeriod=Days
CRLDeltaPeriodUnits=0
LoadDefaultTemplates=0
AlternateSignatureAlgorithm=1
```

Caution

Windows XP and Windows Server 2003 certificate clients do not support the

Alternate Signature Algorithm. If you want these clients to be able to enroll for certificates, do not add the line AlternateSignatureAlgorithm=1 to the CAPolicy.inf. For more information, see Guidelines for Using Alternate Signature Formats.

Note

The OID shown in the example is the Microsoft OID. Individual organizations should obtain their own OIDs. For more information about OIDs, see Obtaining a Root OID from an ISO Name Registration Authority.

Tip

Setting the CRLDeltaPeriodUnits=0 in the CAPolicy.inf disables Delta CRL publishing, which is the appropriate setting for an offline Root CA.

4. Click **Save As**. Ensure the following:
   - **File name** is set to **CAPolicy.inf**
   - **Save as type** is set to **All Files**
   - **Encoding** is **ANSI**
5. When you are prompted to overwrite the file, click **Yes**.

Ensure CAPolicy.inf file has appropriate settings

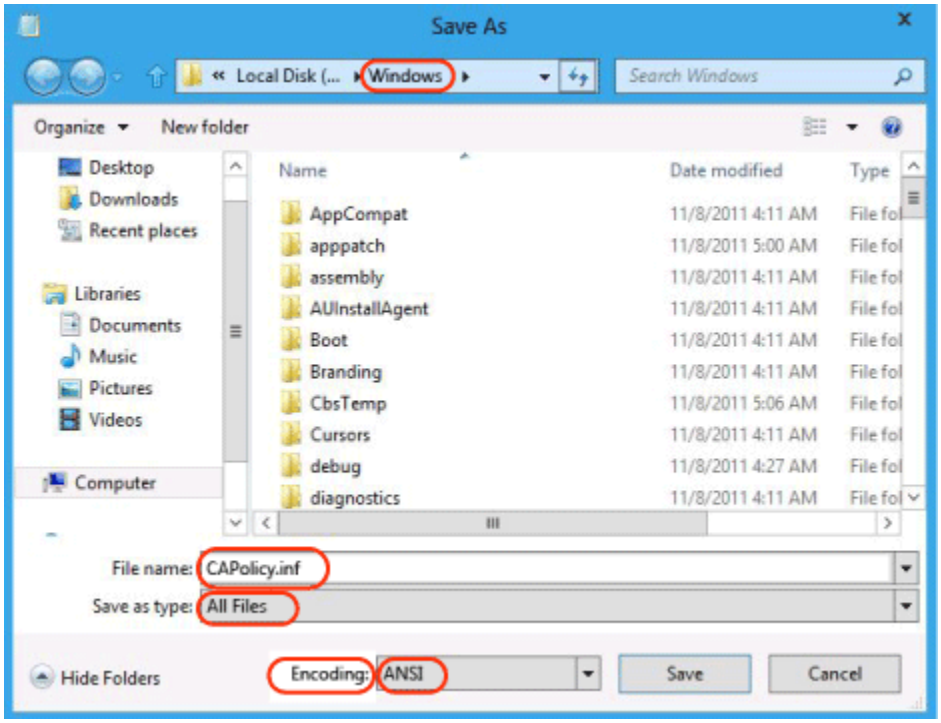

Caution

Be sure to save the CAPolicy.inf with the inf extension. If you do not specifically type **.inf** at the end of the file name and select the options as described, the file

will be saved as a text file and will not be used during CA installation.

6. Close Notepad.

Important

In the CAPolicy.inf, you can see there is a line specifying the URL http://www.contoso.com/pki/cps.txt. The Internal Policy section of the CAPolicy.inf is just shown as an example of how you would specify the location of a certificate practice statement (CPS). To learn more about policy statements including CPS, see Creating Certificate Policies and Certificate Practice Statements (http://technet.microsoft.com/library/cc780454.aspx) and RFC 2527 (http://www.ietf.org/rfc/rfc2527.txt). For more information about CAPolicy.inf file syntax and purposes, see CA Policy.inf Syntax (http://technet.microsoft.com/library/cc728279.aspx).

To install the standalone root CA

1. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. On the **Before you begin** screen, click **Next**.
3. On the **Select installation type** screen, ensure the default selection of **Role-based or feature-based installation** is selected. Click **Next**.
4. On the **Select destination server** screen, ensure that **orca1** is selected and then click **Next**.
5. On the **Select server roles** screen, select the **Active Directory Certificate Services** role.
6. When prompted to install **Remote Server Administration Tools** click **Add Features**. Click **Next**.
7. On the **Select features** screen, click **Next**.
8. On the **Active Directory Certificate Services** screen, click **Next**.
9. On the **Select role services** screen, the **Certification Authority** role is selected by default. Click **Next**.
10. On the **Confirm installation selections** screen, verify the information and then click **Install**.
11. Wait for the installation to complete. The installation progress screen is displayed while the binary files for the CA are installed. When the binary file installation is complete, click the **Configure Active Directory Certificate Services on the destination server** link.

    Click Configure Active Directory Certificate Services on destination server

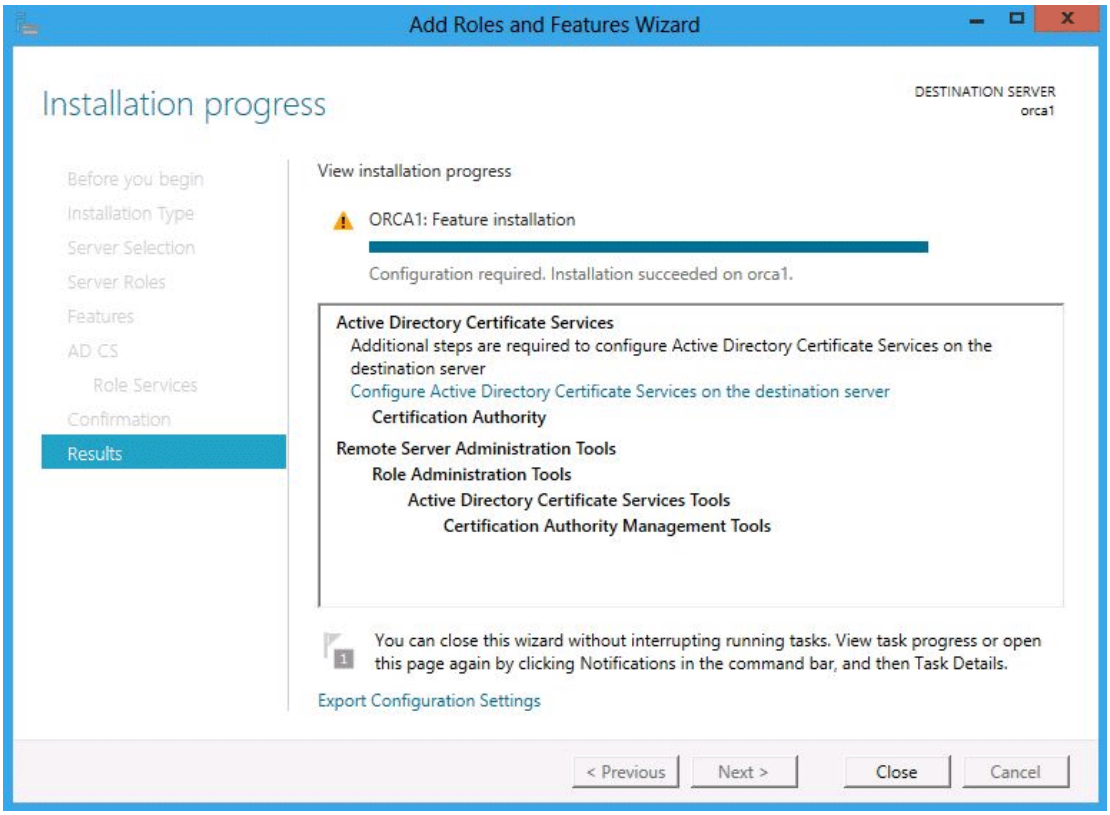

Tip

If you were to click **Close** before the installation completed, you could complete the configuration of the role service by through a link to complete the configuration in the notifications icon of Server Manager.

12. On the **Credentials** screen, you should see that the **ORCA1\Administrator** is displayed in the **Credentials** box. Click **Next**.

Note

When installing a Standalone CA, you must use an account that is a member of the local **Administrators** group.

13. On the **Role Services** screen, select **Certification Authority**. This is the only available selection when only the binary files for the certification authority role are installed on the server. Click **Next**.
14. The only selection available on the **Setup Type** screen is **Standalone CA**. This is because the account used to install is a member of the local Administrators group and the server is not a member of an Active Directory Domain Services (AD DS) domain. Click **Next**.
15. On the **CA Type** screen, **Root CA** is selected by default. Click **Next**.
16. On the **Private Key** screen, leave the default selection to **Create a new private key** selected. Click **Next**.

17. On the **Cryptography for CA** screen, ensure that the cryptographic provider is **RSA#Microsoft Software Key Storage Provider**, the key length is set to **2048** and the hash algorithm is set to **SHA1** then click **Next**.

    Note

    Do not select the **Allow administrator interaction when the private key is accessed by the CA** checkbox. This setting is typically used with Hardware Security Modules (HSMs) and similar key protection devices prompt for additional information when the private key is accessed.

18. On the **CA Name** screen, in the **Common name for this CA** text box, type **ContosoRootCA** and then click **Next**.
19. On the **Validity Period** screen, enter **20** for the number of years for the certificate to be valid.
20. On the **CA Database** screen, leave the default locations for the database and database log files. Click **Next**.
21. On the **Confirmation** screen, click **Configure**.
22. The **Progress** screen is displayed during the configuration processing, then the **Results** screen appears. Click **Close**. If the **Installation progress** screen is still open, click **Close** on that screen as well.

Tip

The following Windows PowerShell commands would perform the same action as shown above

```
Add-WindowsFeature Adcs-Cert-Authority -IncludeManagementTools

Install-AdcsCertificationAuthority –CAType StandaloneRootCA –CACommonName
"ContosoRootCA" –KeyLength 2048 –HashAlgorithm SHA1 –CryptoProviderName
"RSA#Microsoft Software Key Storage Provider"
```

To configure the root CA settings

1. In Server Manager, click **Tools** and then click **Certification Authority**.
2. In the Certification Authority console tree, expand **ORCA1-ContosoRootCA**. Right-click **Revoked Certificates** and then click **Properties**.
3. On the **CRL Publishing Parameters** tab, ensure that **Publish Delta CRLs** is cleared (not selected). Click **OK**.
4. In the Certification Authority console tree, right-click **ORCA1-ContosoRootCA** and then click **Properties**.
5. Click the **Extensions** tab. Ensure that **Select extensions** is set to **CRL Distribution Point (CDP)** and in the **Specify locations from which users can obtain a certificate revocation list (CRL)**, review the default settings.
6. Change **Select extension** to **Authority Information Access (AIA)** and review the default settings. Click **OK**. If you are prompted to restart Active Directory Certificate Services, click **No**. You will restart the service after modifying the default paths in the next

step.

7. From Windows PowerShell run the following commands:

```
certutil -setreg CA\CRLPublicationURLs
"1:C:\Windows\system32\CertSrv\CertEnroll\%3%8.crl\n2:http://www.contoso.com/pki/%
3%8.crl"

certutil –setreg CA\CACertPublicationURLs
"2:http://www.contoso.com/pki/%1_%3%4.crt"

Certutil -setreg CA\CRLOverlapPeriodUnits 12

Certutil -setreg CA\CRLOverlapPeriod "Hours"

Certutil -setreg CA\ValidityPeriodUnits 10

Certutil -setreg CA\ValidityPeriod "Years"

certutil -setreg CA\DSConfigDN CN=Configuration,DC=corp,DC=contoso,DC=com

restart-service certsvc

certutil -crl
```

Notes

The certutil commands above set the CDP and AIA paths respectively for the Root CA. The overlap period for CRLs is the amount of time at the end of a published CRLs lifetime that a client can use to obtain a new CRL before the old CRL is considered unusable, which is set for 12 hours. The default setting for this value is 10% of the CRL lifetime. The validity period settings are to define the number of days, weeks, months, or years that a certificate issued by the CA will be valid, which is set for 10 years in the commands above. The validity period for a certificate cannot be greater than the validity period of the CA that issued the certificate. The default value depends on the type of certificate. The default location of the CDP in also established for eventual use with Active Directory. The same configuration can be accomplished by using the following Windows PowerShell® and certutil commands:

```
$crllist = Get-CACrlDistributionPoint; foreach ($crl in $crllist) {Remove-
CACrlDistributionPoint $crl.uri -Force};

Add-CACRLDistributionPoint -Uri C:\Windows\System32\CertSrv\CertEnroll\%3%8.crl -
PublishToServer -Force

Add-CACRLDistributionPoint -Uri http://www.contoso.com/pki/%3%8.crl -
AddToCertificateCDP -Force

$aialist = Get-CAAuthorityInformationAccess; foreach ($aia in $aialist) {Remove-
CAAuthorityInformationAccess $aia.uri -Force};

Certutil -setreg CA\CRLOverlapPeriodUnits 12

Certutil -setreg CA\CRLOverlapPeriod "Hours"

Certutil -setreg CA\ValidityPeriodUnits 10

Certutil -setreg CA\ValidityPeriod "Years"
```

```
restart-service certsvc

certutil -crl
```

To view the AIA and CDP, you can run the following commands: `Get-CAAuthorityInformationAccess | format-list` and `Get-CACRLDistributionPoint | format-list`. You can also return to the **Extensions** tab in certification authority properties dialog box and see the changes made to the AIA and CDP.

To copy the root CA certificate and CRL to removable media

1. From Windows PowerShell, run the command **dir C:\Windows\system32\certsrv\certenroll*.cr***, which displays the certificates and CRLs in the default certificate store.
2. Copy the CA certificate file and CRL to removable media. For example, if you were running commands to copy the certificate and CRL to a floppy disk drive (A:), you would run the following commands:
   a. **copy C:\Windows\system32\certsrv\certenroll*.cr* A:**
   b. **dir A:**

   Tip
   Substitute the drive letter of your removable media for A: in the commands shown above. The removable media can be either physical or virtual, as discussed in Hardware and software requirements. Also, if you see an error that reads “The volume does not contain a recognized file system.” You may need to format the media. For example, if it is a floppy disk, you might need to type **format a:** and then press ENTER.

To distribute the root CA certificate

1. On APP1, sign in using the User1 account, which is a member of both **Domain Admins** and **Enterprise Admins**. Open Windows PowerShell as administrator. To do so, right-click the Windows PowerShell icon and then click **Run as administrator**. When prompted by User Account Control, click **Yes**.
2. Insert the removable media containing the offline root CA certificate into APP1.
3. From Windows PowerShell change to the removable media drive using the cd command (as in run **cd a:** to change to the root of drive A).
4. From the Windows PowerShell on the removable media drive, run the following commands:

```
certutil –dspublish –f orca1_ContosoRootCA.crt RootCA

certutil –addstore –f root orca1_ContosoRootCA.crt

certutil –addstore –f root ContosoRootCA.crl
```

Note

The first command places the root CA public certificate into the Configuration container of Active Directory. Doing so allows domain client computers to automatically trust the root CA certificate and there is no additional need to distribute that certificate in Group Policy. The second and third commands place the root CA certificate and CRL into the local store of APP1. This provides APP1 immediate trust of root CA public certificate and knowledge of the root CA CRL. APP1 could obtain the certificate from Group Policy and the CRL from the CDP location, but publishing these two items to the local store on APP1 is helpful to speed the configuration of APP1 as a subordinate CA.

The public certificates, certificate revocation lists, and certificate practices statement are all to be placed in the location http://www.contoso.com/pki. Internal client computers will not be able to resolve this computer name to the internal web site (APP1) unless an appropriate DNS entry is placed on the DNS server.

To create a contoso.com DNS zone and www host record

1. On DC1, open the DNS console. In Server Manager, click **Tools**, then click **DNS**.
2. In the DNS console, expand the following in the console tree: **DC1**, **Forward Lookup Zones**.
3. Right-click the **Forward Lookup Zones** and then click **New Zone**.
4. On the **Welcome to the New Zone Wizard** screen, click **Next**.
5. By default you will see that **Primary zone** is selected and that the zone will be stored in Active Directory. To accept these defaults, click **Next**.
6. Leave the default setting and then click **Next**.
7. On **Zone name** screen, type **contoso.com** and then click **Next**.
8. On the **Dynamic Update** screen, leave the default setting and then click **Next**.
9. On the **Completing the New Zone Wizard**, click **Finish**.
10. In the console tree of the DNS console, right-click the **contoso.com** zone and then click **New Host (A or AAAA)**.

    Tip
    You may have to click the corp.contoso.com zone one time before you are able to access the right-click options.
11. In **Name (uses parent domain if left blank)**, type **www**.
12. In **IP Address**, type **10.0.0.3**. This zone and record will direct communications from internal clients for www.contoso.com to the address of APP1. Click **Add Host**.
13. Click **OK** to confirm that the record was created. Click **Done**.
14. Close the DNS console

Step 3: Configure APP1 to distribute certificates and CRLs

In the extensions of the root CA, it was stated that the CRL from the root CA would be available via http://www.contoso.com/pki. Currently, there is not a PKI virtual directory on APP1, so one must be created. In a production environment, you would typically separate the issuing CA role from the role of hosting the AIA and CDP. However, this lab combines both in order to reduce the number of resources needed to complete the lab.

Tip

If a CA cannot find the CRLs of its parent CA, the AD DS service (certsvc) will fail to start on the subordinate CA. This can only be remedied by resolving the CRL distribution issue (recommended) or by changing the CA log level from the default of 3 to level 2. For more information on CA log levels, see Microsoft Knowledge Base article 305018 http://support.microsoft.com/kb/305018.

To configure APP1 to distribute certificates and CRLs

1. Ensure that you sign in using the User1 account. Run Windows PowerShell as Administrator and then run the following commands:

```
New-item -path c:\pki -type directory

write-output "Example CPS statement" | out-file c:\pki\cps.txt

new-smbshare -name pki c:\pki -FullAccess SYSTEM,"CORP\Domain Admins" -
ChangeAccess "CORP\Cert Publishers"
```

2. Open the IIS console. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
3. In the Internet Information Services (IIS) Manager console tree, expand **APP1**. If you are invited to get started with Microsoft Web Platform, click **Cancel**.
4. Expand **Sites** and then right-click the **Default Web Site** and then click **Add Virtual Directory**.
5. In **Alias**, type **pki** and then in physical path type **C:\pki**, then click **OK**.
6. Enable Anonymous access to the pki virtual directory. To do so:
    a. In the **Connections** pane, expand **Default Web Site**, ensure that **pki** is selected.
    b. On **pki Home** click **Authentication**.
    c. In the **Actions** pane, click **Edit Permissions**.
    d. On the **Security** tab, click **Edit**
    e. On the **Permissions for pki** dialog box, click **Add**.
    f. On **Select Users, Computers, Service Accounts, or Groups**, type **Cert Publishers** and then click **Check Names**.
    g. On **Select Users, Computers, Service Accounts, or Groups**, click **Object Types**.
    h. On **Object Types**, select **Service Accounts** and then click **OK**.

i. On **Select Users, Computers, Service Accounts, or Groups**, click **Locations**.

j. On **Locations**, click **APP1** and then click **OK**.

k. On **Select Users, Computers, Service Accounts, or Groups** after **Cert Publishers**, type **;IIS AppPool\DefaultAppPool** and then click **Check Names**. Click **OK**.

> **Note**
> These steps have granted the IIS default application pool **Read & execute**, **List folder contents**, and **Read** permissions. IIS uses the default application pool to allow anonymous access. This will allow users to check the AIA and CDP hosted on IIS.

l. On **Permissions for pki** select **Cert Publishers (CORP\Cert Publishers)**. Under **Permissions for Cert Publishers**, select the **Modify** checkbox in the **Allow** column and then click **OK** twice.

> **Note**
> Granting modify permissions to the pki folder to **Cert Publishers** allows for the publishing of certificates and CRLs by CAs in the enterprise to the folder.

7. In the **pki Home** pane, double-click **Request Filtering**.
8. The **File Name Extensions** tab is selected by default in the **Request Filtering** pane. In the **Actions** pane, click **Edit Feature Settings**.
9. In **Edit Request Filtering Settings**, select **Allow double escaping** and then click **OK**. Close Internet Information Services (IIS) Manager.

> **Note**
> Allowing double escaping is needed if you are publishing Delta CRLs to IIS because the Delta CRL file contains a + symbol. For more information, see Microsoft Knowledge Base article 942076 (http://support.microsoft.com/kb/942076).

10. Run Windows PowerShell as an administrator. From Windows PowerShell, run the command `iisreset`

Step 4: Configure APP1 as an Enterprise Subordinate CA

The steps to configure APP1 as an Enterprise Subordinate CA include the following procedures:

1. Configure the CAPolicy.inf
2. Install the enterprise subordinate CA role
3. To configure the AIA and CDP

To configure the CAPolicy.inf

1. On APP1, as User1, open Windows PowerShell as Administrator and then type **notepad**

c:\Windows\CAPolicy.inf and press ENTER.

2. When asked if you want to create the file. Click **Yes**.
3. Use the following information for the enterprise subordinate CA CAPolicy.inf file.

```
[Version]
Signature="$Windows NT$"
[PolicyStatementExtension]
Policies=InternalPolicy
[InternalPolicy]
OID= 1.2.3.4.1455.67.89.5
Notice="Legal Policy Statement"
URL=http://www.contoso.com/pki/cps.txt
[Certsrv_Server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=5
LoadDefaultTemplates=0
AlternateSignatureAlgorithm=1
```

Caution
Windows XP and Windows Server 2003 certificate clients do not support the Alternate Signature Algorithm. If you want these clients to be able to enroll for certificates, do not add the line AlternateSignatureAlgorithm=1 to the CAPolicy.inf. For more information, see Guidelines for Using Alternate Signature Formats.

4. Click **File**, **Save As** and ensure that you are saving an **ANSI** file named **CAPolicy.inf** in the **C:\Windows** folder. You will have to switch the **Save as type** to **All Files** in order to get the inf extension instead of txt extension. When prompted to replace CAPolicy.inf, click **Yes**.
5. Close Notepad.

To install the enterprise subordinate CA role

1. On APP1, as User1, run Windows PowerShell as Administrator, and then run the following command `gpupdate /force`. This action ensures that the GPO for the trusted root certification authority is applied to APP1.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin**, click **Next**.
4. On the **Select installation type** screen, ensure the default selection of **Role or Feature**

Based Install is selected. Click **Next**.

5. On the **Select destination server** screen, ensure that **APP1** is selected and then click **Next**.
6. On the **Select server roles** screen, select the **Active Directory Certificate Services** role.
7. When prompted to install **Remote Server Administration Tools** click **Add Features**. Click **Next**.
8. On the **Select features** screen, click **Next**.
9. On the **Active Directory Certificate Services** screen, click **Next**.
10. On the **Select role services** screen, ensure **Certification Authority** is selected and then click **Next**.
11. On the **Confirm installation selections** screen, verify the information and then click **Install**.
12. Wait for the installation to complete. The installation progress screen is displayed while the binary files for the CA are installed. When the binary file installation is complete, click the **Configure Active Directory Certificate Services on the destination server** link.

    Tip

    If you clicked **Close** before the installation completed, you could complete the configuration of the role service by through a link to complete the configuration in the notifications icon of Server Manager.
13. On the **Credentials** screen, the credentials for User1 appear. Click **Next**.
14. On the **Role Services** screen, select **Certification Authority**.
15. On the **Setup Type** screen, ensure that **Enterprise CA** is selected and then click **Next**.

    Note

    If the computer is a domain member and the credentials supplied previously were for an account that is a member of the **Enterprise Admins** group, you can select **Enterprise CA** or **Standalone CA**. If the computer is not a domain member or credentials were entered for an account that is not a member of **Enterprise Admins**, then only the **Standalone CA** selection is available.
16. On the **CA Type** screen, select **Subordinate CA** to install an Enterprise Subordinate CA. Click **Next**.
17. On the **Private Key** screen, ensure the **Create a new private key** option is selected and then click **Next**.
18. The **Cryptography for CA** screen, ensure that the cryptographic provider is **RSA#Microsoft Software Key Storage Provider**, key length is **2048**, and the hash algorithm is set to **SHA1**. Click **Next**.
19. On the **CA Name** screen, in **Common name for this CA**, type **IssuingCA-APP1**. You will see that the distinguished name changes to **CN=IssuingCA-APP1,DC=corp,DC=contoso,DC=com**. Click **Next**.
20. On the **Certificate Request** screen, notice that **Save a certificate request to file on the target machine** is selected. This is the correct option because we are using an offline

parent CA (the root CA) in this configuration. Leave the default and click **Next**.

21. On the **CA Database** screen, leave the default database and log locations and then click **Next**.
22. On the **Confirmation** screen, click **Configure**.
23. On the **Results** screen, you see that you must take the certificate request to the ContosoRootCA in order to complete the configuration. Click **Close**

    Notes

    The Windows PowerShell commands to perform the installation of the Enterprise Subordinate CA as shown in this section are:

    ```
    Add-WindowsFeature Adcs-Cert-Authority -IncludeManagementTools

    Install-AdcsCertificationAuthority -CAType EnterpriseSubordinateCA -
    CACommonName "IssuingCA-APP1" -KeyLength 2048 -HashAlgorithm SHA1 -
    CryptoProviderName "RSA#Microsoft Software Key Storage Provider"
    ```

24. Copy the certificate request to removable media to take to the ORCA1. For example, if you wanted to copy the file from the C:\ drive to a floppy drive with drive letter A:\, then you could run the following command from Windows PowerShell: **copy C:*.req A:**
25. Take the removable media with the certificate request file to the ORCA1. Sign on to the root CA using an account that is a member of local **Administrators**.
26. On ORCA1, from Windows PowerShell, submit the request using the following command (assuming that A:\ is your removable media drive letter):
    certreq -submit A:\APP1.corp.contoso.com_IssuingCA-APP1.req

    Note

    If the removable media has a different drive letter, then substitute that letter for A:\.

27. On **Certification Authority List**, ensure that **ContosoRootCA (Kerberos)** CA is selected and then click **OK**. You see that the certificate request is pending and the request identification number. Ensure that you note the request ID number.
28. On ORCA1, you must approve the request. You can do this using Server Manager or by using certutil from the command line.
    - To use Server Manager, click **Tools**, and then click **Certification Authority**. Expand the **ContosoRootCA** object and then click **Pending Requests**.

      Right-click the Request ID that corresponds with the one you saw when you submitted the request in the previous step. Click **All Tasks** and then click **Issue**.

      Click **Issued Certificates** and see the issued certificate in the **Details** pane.
    - To use certutil, enter `Certutil resubmit` <*RequestId*>, replace the actual request number for <RequestId>. For example, if the Request ID is 2, you would enter `Certutil resubmit 2`
29. From the command prompt on ORCA1, retrieve the issued certificate by running the command
    certreq –retrieve <*RequestId*> <*drive*>**:\APP1.corp.contoso.com_corp-APP1-CA.crt**.
    Substitute the actual number of the request when it was submitted for <RequestId> and

the actual drive letter of the removable media for <drive>. For example, if the request ID where 2 and the removable media was drive A, then the request would be: `certreq -retrieve 2 a:\APP1.corp.contoso.com_IssuingCA-APP1.crt`. When prompted to select the CA, ensure that ORCA1-ContosoRootCA is selected and then click **OK**.

30. On ORCA1, run the command **dir A:** (assuming that A is the removable media drive letter, if not substitute the correct drive letter for A). You see that ContosoRootCA.crl, orca1_ORCA1-ContosoRootCA.crt, and APP1.corp.contoso.com_corp-APP1-CA.crt are now saved to the removable media. Move the removable media to APP1.
31. On APP1, in Windows PowerShell, run the following commands to copy the root CA certificate and CRL to the PKI folder (assuming that A: is the removable media drive, if not substitute the correct drive letter), install the subordinate CA certificate, start the certificate service, and copy the subordinate CA certificate and CRLs to the PKI folder:
    - `copy a:\*.cr* c:\pki\`
    - `certutil –installcert a:\APP1.corp.contoso.com_corp-APP1-CA.crt`
    - `start-service certsvc`
    - `copy c:\Windows\system32\certsrv\certenroll\*.cr* c:\pki\`

Tip

ORCA1 is no longer needed for this lab, so you can turn it off. To turn off a computer from Windows PowerShell, you can run the command `stop-computer`.

To configure the AIA and CDP Settings

1. On APP1, as User1, right-click Windows PowerShell, click **Run as Administrator**. Click **Yes** to confirm that you want to run Windows PowerShell as an Administrator.
2. From Windows PowerShell run the following commands:

```
certutil -setreg CA\CRLPublicationURLs
"1:C:\Windows\system32\CertSrv\CertEnroll\%3%8.crl\n2:http://www.contoso.com/pki/%
3%8.crl"

certutil -setreg CA\CACertPublicationURLs
"2:http://www.contoso.com/pki/%1_%3%4.crt\n1:file://\\App1.corp.contoso.com\pki\%1
_%3%4.crt"

Certutil -setreg CA\CRLPeriodUnits 2

Certutil -setreg CA\CRLPeriod "Weeks"

Certutil -setreg CA\CRLDeltaPeriodUnits 1

Certutil -setreg CA\CRLDeltaPeriod "Days"

Certutil -setreg CA\CRLOverlapPeriodUnits 12

Certutil -setreg CA\CRLOverlapPeriod "Hours"

Certutil -setreg CA\ValidityPeriodUnits 5

Certutil -setreg CA\ValidityPeriod "Years"

restart-service certsvc
```

```
certutil -crl
```

Notes

The same configuration can be accomplished using the following Windows PowerShell and certutil commands:

```
$crllist = Get-CACrlDistributionPoint; foreach ($crl in $crllist) {Remove-CACrlDistributionPoint $crl.uri -Force};

Add-CACRLDistributionPoint -Uri C:\Windows\System32\CertSrv\CertEnroll\%3%8%9.crl -PublishToServer -PublishDeltaToServer -Force

Add-CACRLDistributionPoint -Uri http://www.contoso.com/pki/%3%8%9.crl -AddToCertificateCDP -Force

Add-CACRLDistributionPoint -Uri file://\\App1.corp.contoso.com\pki\%3%8%9.crl -PublishToServer -PublishDeltaToServer -Force

$aialist = Get-CAAuthorityInformationAccess; foreach ($aia in $aialist) {Remove-CAAuthorityInformationAccess $aia.uri -Force};

Add-CAAuthorityInformationAccess -AddToCertificateAia http://www.contoso.com/pki/%1_%3%4.crt -Force

Certutil -setreg CA\CRLPeriodUnits 2

Certutil -setreg CA\CRLPeriod "Weeks"

Certutil -setreg CA\CRLDeltaPeriodUnits 1

Certutil -setreg CA\CRLDeltaPeriod "Days"

Certutil -setreg CA\CRLOverlapPeriodUnits 12

Certutil -setreg CA\CRLOverlapPeriod "Hours"

Certutil -setreg CA\ValidityPeriodUnits 5

Certutil -setreg CA\ValidityPeriod "Years"

restart-service certsvc

certutil -crl
```

By sharing the pki folder and including the file path file://\\App1.corp.contoso.com\pki\%3%8%9.crl as a CDP extension, the CRLs and Delta CRLs will be copied to the share when you run the command `certutil -crl`. If you want to further restrict access to the share, you could create a separate group and include only the CAs that you want to authorize to publish to the share in that group. Then, share the pki folder only to that specific group and the SYSTEM account.

Important

A configuration item that is typically performed on production CAs that is not part of this lab is to enable Audit Object Access (http://technet.microsoft.com/library/cc776774.aspx) and then to enable all auditing events by running the following command: **certutil -setreg CA\AuditFilter 127**. After doing so, ensure that you regularly archive the Security Event

Log and follow the Auditing Security Events Best Practices (http://technet.microsoft.com/library/cc778162.aspx).

Step 5: Configure computer certificate autoenrollment

There are two procedures in order to configure computer certificate autoenrollment:

1. Enable certificate autoenrollment through Group Policy
2. Configure a client and server authentication certificate template for autoenrollment

To enable certificate autoenrollment through Group Policy

1. On DC1, sign in as User1. In Server Manager, click **Tools**, and then click **Group Policy Management**.
2. On the console tree, expand the following objects: **Forest: corp.contoso.com**, **Domains**, **corp.contoso.com**.

   Note
   You might see a warning that any policies linked to the domain will affect all computers to which the policy is linked. If so, read it and then click **OK**.
3. In the console tree, right-click **Default Domain Policy**, and then click **Edit**.
4. In the console tree of the Group Policy Management Editor, under **Computer Configuration**, expand the following objects: **Policies**, **Windows Settings**, **Security Settings**, and then click **Public Key Policies**.
5. In the details pane, double-click **Certificate Services Client - Auto-Enrollment**. In **Configuration Model**, select **Enabled**.
6. Select **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**. Click **OK**.
7. Close Group Policy Management Editor and Group Policy Management Console.

To configure a client server authentication certificate template for autoenrollment

1. On APP1, in the Certification Authority console pane, ensure that **IssuingCA-APP1** is expanded.
2. Right-click **Certificate Templates** and then click **Manage**.
3. In the details pane, right-click **Workstation Authentication** and then click **Duplicate Template**.
4. Click the **General** tab, in **Template display name**, type **Client-Server Authentication**.
5. Click the **Extensions** tab, ensure **Application Policies** is selected, and then click **Edit**.
6. Click **Add** then click **Server Authentication**. Click **OK** twice.
7. On the **Properties of New Template** dialog, click the **Security** tab.
8. In **Group or user names**, click **Domain Computers (CORP\Domain Computers)**.
9. In the **Autoenroll** row, select the **Allow** checkbox. This will cause all domain computers

to automatically enroll for certificates using this template.

Notes

- You would typically not assign a template both the Client Authentication and the Server Authentication enhanced key usage (EKU). Also, Server Authentication EKU are typically not configured for autoenrollment. This is done in this lab only for convenience and compatibility with other labs.
- The computers also need **Read** permission for the template in order to enroll. However, this permission is already granted to the **Authenticated Users** group. All computer accounts in the domain are members of **Authenticated Users**, so they already have the permission to **Read** the template.

10. Click **OK**. Close the **Certificate Templates Console**.
11. Right-click **Certificate Templates**, click **New**, click **Certificate Template to Issue**.
12. In the **Enable Certificate Templates** dialog box, click **Client-Server Authentication** and then click **OK**. Close the Certification Authority console.

Step 6: Configuring SSL for APP1

To demonstrate how the certificates deployed through AD DS and AD CS can be used, you will secure the APP1 Web site using SSL and then connect to that secure site with CLIENT1.

Note

This part of the lab is done to demonstrate using a certificate to secure a Web site.

There are two procedures in this step:

1. Secure the APP1 Default Web Site
2. Connect to the secure web site

To secure the APP1 Default Web Site

1. On APP1, as User1, run Windows PowerShell as Administrator. Then, run the following commands:

   `Gpupdate /force`. Wait for the update of Group Policy to complete and then close the Command Prompt. This ensures that the autoenrollment certificate distributed through Group Policy is issued to APP1.

   ```
   cd cert:\LocalMachine\My
   ```

   ```
   dir | format-list
   ```

   You should see that you have two certificates. One was issued by ContosoRootCA, which is the APP1 CA certificate. The other certificate was issued by IssuingCA-APP1 and it can be used to secure the APP1 default web site.

2. Open the Internet Information Services (IIS) Manager console. To do so, in Server Manager, click **Tools** and then click **Internet Information Services (IIS) Manager**. In the contents pane, expand the following path **APP1**, **Sites**, and **Default Web Site**.

Note

If you see an Internet Information Services (IIS) Manager prompt asking if you want to get started with Microsoft Web Platform, click **Cancel**.

3. Click **Default Web Site**. In the **Actions** pane click **Bindings**.
4. In the **Site Bindings** dialog box, click **Add**.
5. In the **Add Site Binding** dialog box, in Type, select **https**.
6. Under **SSL certificate**, click **Select**.
7. In **Select Certificate** use the selection box to select the certificate that was issued by the IssuingCA-APP1 through the Group Policy. This will be a certificate with a long alphanumeric, as opposed one that reads IssuingCA-APP1. To verify you have the correct certificate, click **View**. Ensure the certificate you select shows that it was issued to **APP1.corp.contoso.com** and issued by **IssuingCA-APP1**. Once you have the correct certificate, click **OK** on the **Certificate** dialog box.
8. On **Add Site Binding** dialog box, click **OK**.
9. In the **Site Bindings** dialog box, click **Close**.

To connect to the secure web site

1. Connect CLIENT1 to the Corporate network.
2. Log on to CLIENT1 as User1.
3. Open Internet Explorer on CLIENT1.
4. In Internet Explorer, enter the address **https://app1.corp.contoso.com** and press ENTER. When you see the default IIS 8 web page, you are confirming that https and the SSL binding are working for the Default Web Site on APP1.

Tip

If instead you see that there is a problem with the certificate, then you probably selected an incorrect certificate in the previous procedure. You must select the certificate that was issued for the name APP1.corp.contoso.com. Also, it could be that Group Policy has not yet updated the Trusted Root Certification authorities. To ensure that the Group Policy updates are in place, open Explorer, then type **cmd** in the Explorer address bar. Then type **gpupdate /force** and press ENTER.

Important

The ORCA1 certificate revocation list (CRL) is valid for 26 weeks, which was configured using the CAPolicy.inf. The APP1 CRL must be updated weekly by default. To update the CRL, use the command:

`Certutil -crl`, which publishes the CRL to the locations that you specified in the CA Properties Extensions tab.

Note

To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

See Also

Windows Server Security Forum

Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)

Windows PKI Documentation Reference and Library

Windows PKI Blog

Test Lab Guide: Demonstrating Certificate Key-Based Renewal

The purpose of this Test Lab Guide (TLG) is to give you hands-on experience configuring the Certificate Enrollment Web Service and the Certificate Enrollment Policy Web Service role services. These roles services are part of the Active Directory Certificate Services (AD CS) server role in Windows Server® 2012 and Windows Server® 2012 R2.

Note
To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

In this guide

This document provides instructions that explain how to extend the Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy to provide Certificate Enrollment Web Services. In Windows Server® 2012, Windows Server® 2012 R2, and Windows® 8 and Windows® 8.1, you can configure certificate autorenewal for computers outside the domain. This includes computers from other forests, domains, and workgroups. This lab demonstrates the steps to issue a certificate to a computer that is not joined to your domain, and then configure that certificate for autorenewal.

Important
The configuration of the computers and network in this guide is designed to give you hands-on practice using Certificate Enrollment Web Services. The design decisions made in this guide were aimed at increasing your hands-on experience, and they do not reflect a best practices configuration. For best practice information, see Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure and PKI Design Brief Overview.

Test lab overview

The following test lab configuration adds three computers to the configuration that is outlined in Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy. One of the additional computers will be a Certificate Enrollment Web Services server. Another computer will be a Certificate

Enrollment Policy Web Services server. The third computer will be a web server that is not joined to the domain. There are seven major steps to complete, which include multiple subordinate procedures.

1. Step 1: Complete the Base configuration Test Lab
2. Step 2: Complete the Test Lab Guide: Deploying an AD CS Two Tier PKI Hierarchy
3. Step 3: Configure the CEP1 server
4. Step 4: Configure the CES1 server
5. Step 5: Prepare an appropriate certificate template
6. Step 6: Configure WEB1
7. Step 7: Obtain a certificate and test automatic renewal

Hardware and software requirements

The following are the minimum required components for this test lab:

1. The product disc or files for Windows Server 2012 or Windows Server 2012 R2.
2. The product disc or files for Windows 8 or Windows 8.1.
3. Six computers that meet the minimum hardware requirements for Windows Server 2012 or Windows Server 2012 R2.

   Tip
   The CEP1 and CES1 servers should be allocated at least 1.5 GB of RAM, if possible.
4. One computer that meets the minimum hardware requirements for Windows 8 or Windows 8.1.
5. One removable storage device with enough free space to hold a few certificates and certificate revocation lists (about 10 kilobytes). This can be a physical or virtual removable storage device depending on whether your lab is using physical or virtual computers.

   Note
   For instructions about how to transfer files by using a virtual floppy disk in a server running Hyper-V , see Creating, Using, and Transferring Files using Virtual Floppy Disks.
6. If you want to deploy the Base Configuration Test Lab in a virtualized environment, your virtualization solution must support Windows Server 2012 or Windows Server 2012 R2 and Windows 8 or Windows 8.1 64-bit virtual machines. The server hardware must support the amount of RAM required to run the virtual operating systems included in the Base Configuration Test Lab and any other virtual machines that may be required by additional TLGs.

Important
Run Windows Update on all computers or virtual machines either during the installation or immediately after installing the operating systems. After running Windows Update, you can isolate your physical or virtual test lab from your production network.

Note
Run Windows PowerShell® commands as an administrator. When you are not signed in as the default administrator, you can right-click the Windows PowerShell program icon and then select **Run as administrator**.

Step 1: Complete the Base Configuration Test Lab

Before you begin the instructions in this guide, you must complete the Base Configuration Test Lab. For more information, see Test Lab Guide: Base Test Lab Guide for Windows Server 2012.

Step 2: Complete the Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy

The two-tier PKI hierarchy provides the basis for the lab explained in this topic. For more information, see Test Lab Guide: Deploying an AD CS Two-Tier PKI Hierarchy .

Important
The ORCA1 certificate revocation list (CRL) for this lab was configured by using the CAPolicy.inf, and it is valid for 26 weeks. The APP1 CRL must be updated weekly. To update the CRL, run the following command on APP1 from Windows PowerShell:

```
certutil -crl
```

Step 3: Configure the CEP1 server

The configuration of the CEP1 server allows two methods for client computers to obtain certificate enrollment policies:

1. User name and password authentication
2. Certificate authentication

Tip
For the remainder of this lab, only the domain controller (DC1) and the, APP1 server are needed from the Base Configuration Test Lab Guide. You will be installing three additional servers: CEP1, CES1, and WEB1. Before installing any new servers, ensure that DC1 and APP1 are running.

The procedures to configure the CEP1 server to support the configuration that is demonstrated in this guide are as follows:

1. Install the operating system
2. Configure TCP/IP
3. Join the computer to the domain
4. Install the Certificate Enrollment Policy Web Service to use user name and password authentication
5. Install the Certificate Enrollment Policy Web Service to use certificate authentication

To install the operating system

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation. Specify Windows Server 2012 (full installation), and create a strong password for the local Administrator account. Then sign in by using the local Administrator account.
3. Connect the computer to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.
4. Connect the computer to the Corpnet subnet.

To configure TCP/IP

1. From Windows PowerShell, run **ncpa.cpl**.
2. In **Network Connections**, right-click **Ethernet**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.4**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**.
7. Close the **Network Connections** window.
8. In Windows PowerShell, run **sysdm.cpl**.
9. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
10. In **Computer name**, type **CEP1** as the new name for the computer, and then click **OK**.
11. When you are prompted that you must restart the computer, click **OK**.
12. In the **System Properties** dialog box, click **Close**.
13. When you are prompted to restart the computer, click **Restart Now**.
14. After restarting, sign in by using the local Administrator account.

To join the computer to the domain

1. From Windows PowerShell, run **sysdm.cpl**.
2. In the **System Properties** dialog box, click the **Computer Name** tab, and then click **Change**.
3. In **Member of**, select **Domain**, type **corp.contoso.com**, and then click **OK**.
4. When you are prompted for a user name and password, enter the credentials for User1, and then click **OK**.
5. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. In the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

9. After the computer restarts, click **Switch User**, and then click **Other User**. Sign in to the CORP domain by using the User1 account, which is a member of Domain Admins and Enterprise Admins groups.

Tip

The Windows PowerShell commands to configure the IP address, rename the computer, and join the computer to the domain are:

```
$NetIP = Get-NetIPAddress | where {$_.Addressfamily -eq "IPv4" -and
$_.InterfaceAlias -like "*Ethernet*"}

$NetAlias = $NetIP.InterfaceAlias

New-NetIPAddress -InterfaceAlias $NetAlias -IPAddress 10.0.0.4 -PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias $NetAlias -ServerAddresses 10.0.0.1

Set-DnsClient -InterfaceAlias $NetAlias -ConnectionSpecificSuffix corp.contoso.com

Add-Computer -NewName CEP1 -DomainName corp.contoso.com -Credential CORP\User1

Restart-computer
```

To install the Certificate Enrollment Policy Web Service to use user name and password authentication

1. On the CEP1 server, ensure you are signed in as User1. Right-click Windows PowerShell and then click **Run as Administrator**, and then run the following commands:

   ```
   gpupdate /force

   cd cert:\LocalMachine\My

   dir | format-list
   ```

   Important

   You need a Server Authentication certificate for the CEP1 server to perform the following procedure. The certificate should be automatically distributed to your computer through Group Policy and the certification authority (CA) that is running on APP1, which was configured in the Test Lab Guide: Deploying a Two-Tier PKI Hierarchy. The `gpupdate` command forces the Group Policy to update and download the certificate. You should see that you have a certificate issued by IssuingCA-APP1, which you will be using to install the Certificate Enrollment Policy Web Service. If you do not see the certificate immediately after running these commands, wait a couple of minutes, and then run `dir | format-list` command again.

2. In Server Manager, click **Manage**, and then click **Add Roles and Features**. On the **Before you begin** screen, click **Next**.
3. On the **Select installation type** screen, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
4. On the **Select destination server** screen, ensure that **CEP1.corp.contoso.com** is selected, and then click **Next**.

5. On the **Select server roles** screen, select **Active Directory Certificate Services**. When you are prompted to add the Remote Server Administration Tools, click **Add Features**, and then click **Next**.
6. On the **Select features** screen, click **Next**.
7. On the **Active Directory Certificate Services** screen, click **Next**.
8. On the **Select role service** screen, clear the **Certification Authority** role, and select the **Certificate Enrollment Policy Web Service**. When you are prompted to add roles and features, click **Add Features**, and then click **Next**.
9. On the **Web Server Role (IIS)** screen, click **Next**.
10. On the **Select role services** screen, click **Next**.
11. On the **Confirm installation selections** screen, click **Install**.
12. When the installation is complete, click **Configure Active Directory Certificate Services on the destination computer**.

    Tip
    If you clicked **Close** before the installation completed, you can complete the role service configuration through a link in the notifications icon in Server Manager.
13. On the **Credentials** screen, click **Next**.
14. On the **Role Service** screen, select **Certificate Enrollment Policy Web Service**, and then click **Next**.
15. On the **Authentication Type for CEP** screen, select **User name and password**, and then click **Next**.
16. On **Enable Key-Based Renewal for CEP** screen, select the **Enable key-based renewal** check box, and then click **Next**.
17. On the **Server Certificate** screen, select the **CEP1.corp.contoso.com** certificate that was issued by IssuingCA-APP1, and then click **Next**.
18. On the **Confirmation** screen, click **Configure**.
19. After the configuration is complete, on the **Results** screen, click **Close**, and then in the **Add Roles and Features Wizard**, click **Close**.

    Tip
    The following Windows PowerShell commands run from the **Cert:\LocalMachine\My** path as an Administrator will also perform the installation that was described in the previous steps:

    ```
    Install-WindowsFeature Web-WebServer -IncludeManagementTools

    Add-WindowsFeature Adcs-Enroll-Web-Pol

    Install-AdcsEnrollmentPolicyWebService -AuthenticationType Username -
    KeyBasedRenewal -SSLCertThumbprint (dir -dnsname
    cep1.corp.contoso.com).Thumbprint
    ```
20. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
21. In the **Connections** pane of the Internet Information Services (IIS) Manager console,

expand the **CEP1** server.

Note
If you are prompted to get started with the Microsoft Web Platform, click **Cancel**.

22. Expand **Sites**, and then expand the **Default Web Site**.
23. Click the **KeyBasedRenewal_ADPolicyProvider_CEP_UsernamePassword** application.
24. In the center pane, double-click **Application Settings**.
25. In **Application Settings**, double-click **FriendlyName**. In the **Value** text box, type **SSL Server Certificates**, and then click **OK**.
26. In **Application Settings**, double-click **URI**, and ensure that the URI value is **https://cep1.corp.contoso.com/KeybasedRenewal_ADPolicyProvider_CEP_UsernamePassword/service.svc/CEP**.

Note
This URI will be used in a Windows PowerShell command later from WEB1to contact the CEP1 server for certificate enrollment.

27. Click **OK**, and then close Internet Information Services (IIS) Manager.

To install a Certificate Enrollment Policy Web Service that uses certificate authentication

1. To install a second instance of the Certificate Enrollment Policy Web Service on the CEP1 server, you must use Windows PowerShell. Open Windows PowerShell as an Administrator, and run the following command:

```
cd cert:\LocalMachine\My
Install-AdcsEnrollmentPolicyWebService -AuthenticationType
Certificate -KeyBasedRenewal -SSLCertThumbprint (dir -dnsname
cep1.corp.contoso.com).Thumbprint
```

2. When you are prompted for confirmation, type **Y**, and then press ENTER.

Note
You will see a confirmation that reads **ErrorString**. If the confirmation is blank under **ErrorString**, the installation succeeded. Otherwise, review your command for errors, correct them, and try again.

3. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
4. In the **Connections** pane of the Internet Information Services (IIS) Manager console, expand the **CEP1** server.

Note
If you are prompted to get started with the Microsoft Web Platform, click **Cancel**.

5. Expand **Sites**, and then expand **Default Web Site**.
6. Click the **KeyBasedRenewal_ADPolicyProvider_CEP_Certificate** application.
7. In the center pane, double-click **Application Settings**.
8. In **Application Settings**, double-click **FriendlyName**, and then in the **Value** text box, type **SSL Server Certificates**. Click **OK**.
9. In **Application Settings**, double-click **URI**, and ensure that the URI value is **https://cep1.corp.contoso.com/KeybasedRenewal_ADPolicyProvider_CEP_Certificate/service.svc/CEP**. This URI will be used to configure WEB1 to contact the CEP1 server for certificate renewal.

Step 4: Configure the CES1 server

The procedures to configure the CES1 server to support the configuration that is demonstrated in this guide are as follows:

1. Install the operating system
2. Configure TCP/IP
3. Join the computer to the domain
4. Configure the service account
5. Install the Certificate Enrollment Web Service to use user name and password authentication
6. Install the Certificate Enrollment Web Service to use certificate authentication
7. Grant the service account Read permission on the CA
8. Trust the service account for delegation

The Certificate Enrollment Web Services server is used to submit certificate requests to the CA. The certificate requests are submitted by the Certificate Enrollment Policy Web Services service account to APP1 on behalf of the users, computers, and devices that request them. In addition to the configuration for accepting user name and password authentication and certificate authentication, the service account requires Read permission to the CA.

To install the operating system

1. Start the installation of Windows Server 2012
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Sign in by using the local Administrator account.
3. Connect the computer to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.
4. Connect the computer to the Corpnet subnet.

To configure TCP/IP

1. From Windows PowerShell, run **ncpa.cpl**.
2. In **Network Connections**, right-click **Ethernet**, and then click **Properties**.

3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.5**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**.
7. Close the **Network Connections** window.
8. From Windows PowerShell, run **sysdm.cpl**.
9. On the **System Properties** dialog box on the **Computer Name** tab, click **Change**.
10. In **Computer name**, type **CES1** as the new name for the computer, and then click **OK**.
11. When you are prompted that you must restart the computer, click **OK**.
12. On the **System Properties** dialog box, click **Close**.
13. When you are prompted to restart the computer, click **Restart Now**.
14. After restarting, sign in using the local Administrator account.

To join the computer to the domain

1. From Windows PowerShell, run **sysdm.cpl**.
2. In the **System Properties** dialog box, click the **Computer Name** tab, click **Change**.
3. In **Member of**, select **Domain**, and then type **corp.contoso.com**. Click **OK**.
4. When you are prompted for a user name and password, type a domain user name and password (you can use any valid user account, including the default administrator), and then click **OK**.
5. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.
9. After the computer restarts, click **Switch User**, and then click **Other User** and sign in to the CORP domain using an account that is a member of **Enterprise Admins**.

Tip

The Windows PowerShell commands to configure the IP address, rename the computer, and join the computer to the domain are:

```
$NetIP = Get-NetIPAddress | where {$_.Addressfamily -eq "IPv4" -and
$_.InterfaceAlias -like "*Ethernet*"}

$NetAlias = $NetIP.InterfaceAlias

New-NetIPAddress -InterfaceAlias $NetAlias -IPAddress 10.0.0.5 -PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias $NetAlias -ServerAddresses 10.0.0.1

Set-DnsClient -InterfaceAlias $NetAlias -ConnectionSpecificSuffix corp.contoso.com

Add-Computer -NewName CES1 -DomainName corp.contoso.com -Credential CORP\User1
```

```
Restart-computer
```

To configure the service account

1. On the domain controller, DC1, as User1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. In the console tree, click the **corp (local)**.
3. In the **Tasks** pane, click **New**, and then click **User**.
4. In the **Create User** dialog box, in **Full name**, type **CES**, and in **User SamAccountName logon:**, ensure that **corp** is the displayed as the domain, and then type **CES** as the account name.
5. In **Password**, type the password that you want to use for this account, and in **Confirm password**, type that password again.
6. Under Password options, select **Other password options**, and select **Password never expires**.
7. Click **OK** to create the user account, and then close the **Create User** dialog box.

   Tip

   Alternatively, you could create the service account using Windows PowerShell. You can run the following command to add the CES user account to Active Directory Domain Services (AD DS):

   ```
   New-ADUser -SamAccountName ces -AccountPassword (read-host "Set user password" -assecurestring) -name "ces" -enabled $true -PasswordNeverExpires $true -ChangePasswordAtLogon $false
   ```

8. The CES user account requires a service principal name (SPN) to be delegated web enrollment permissions. To create the SPN, open Windows PowerShell and run the following command: `setspn -s http/ces1.corp.contoso.com corp\ces`
9. In CES1, sign in as User1. In **Server Manager**, click **Tools**, and then click **Computer Management**.
10. In the console tree, expand **Local Users and Groups**, and then click **Groups**.
11. In the details pane, double-click **IIS_IUSRS**.
12. In **IIS_IUSRS Properties**, click **Add**.
13. In **Select Users, Computers, Service Accounts, or Groups**, type **CES**, and then click **Check Names**. Click **OK** twice.

    Tip

    Alternatively, you can use Windows PowerShell to add the CES user account to the IIS_IUSRS local group. To do so, run the following command:

    ```
    Net localgroup IIS_IUSRS corp\ces /Add
    ```

Important

You need a Server Authentication certificate for the CES1 server to perform the following procedure. The certificate should be automatically distributed to your computer through

Group Policy and the CA that is running on APP1, which was configured in the Test Lab Guide: Deploying a Two-Tier PKI Hierarchy. The first command that you are asked to run in the following procedure is a command to update Group Policy to ensure that the certificate is distributed to CES1.

To install the Certificate Enrollment Web Service to use user name and password authentication

1. On CES1, as User1, open Windows PowerShell as an Administrator, and run the following command:

   ```
   gpupdate /force
   cd cert:\LocalMachine\My
   dir | format-list
   ```

   You should see that you have a certificate issued by IssuingCA-APP1, which you will be using to install the Certificate Enrollment Web Service. If this is not the case, try restarting your computer.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**. If the **Before you begin** screen appears, click **Next**.
3. On the **Select installation type** screen, select **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** screen, select **CES1.corp.contoso.com**, and then click **Next**.
5. On the **Select server roles** screen, select **Active Directory Certificate Services**. When you are prompted to add the Remote Server Administration Tools, click **Add Features**, and then click **Next**.
6. On the **Select features** screen, click **Next**.
7. On the **Active Directory Certificate Services** screen, click **Next**.
8. On the **Select role service** screen, clear the **Certification Authority** role, and then select the **Certificate Enrollment Web Service**. When you are prompted to add roles and features, click **Add Features**, and then click **Next**.
9. On the **Web Server Role (IIS)** screen, click **Next**.
10. On the **Select role services** screen, click **Next**.
11. On the **Confirm installation selections** screen, click **Install**.
12. When the installation is complete, click **Configure Active Directory Certificate Services on the destination computer**.

    Tip

    If you clicked **Close** before the installation completed, you can complete the role service configuration through a link in the notifications icon in Server Manager.
13. On the **Credentials** screen, ensure that you see CORP\User1 as the account to use for installation, and then click **Next**.
14. On the **Role Service** screen, select **Certificate Enrollment Web Service**, and then click

Next.

15. On the **CA for CES** screen, click **Select**. In **Select Certification Authority**, select **IssuingCA-APP1**, click **OK**, and then click **Next**.
16. On the **Authentication Type for CES** screen, select **User name and password**, and then click **Next**.
17. On the **Service Account for CES** screen, ensure **Specify service account (recommended)** is selected, and then click **Select**.
18. In AD CS Configuration enter the **CORP\CES** as the user name. Enter the password for the account and then click **OK**.
19. On the **Server Certificate** screen, select the **CES1.corp.contoso.com** certificate that was issued by IssuingCA-APP1, and then click **Next**.
20. On the **Confirmation** screen, click **Configure**.
21. After the configuration is complete, on the **Results** screen, click **Close**, and then in the **Add Roles and Features Wizard**, click **Close**.

Tip

Alternatively, the following Windows PowerShell commands can be run from the **Cert:\LocalMachine\My** path as an Administrator to perform the installation and configuration described in the previous steps:

```
Install-WindowsFeature Web-WebServer -IncludeManagementTools

Add-WindowsFeature Adcs-Enroll-Web-Svc

Install-AdcsEnrollmentWebService -ServiceAccountName "CORP\CES" -CAConfig
"APP1.corp.contoso.com\IssuingCA-APP1" -SSLCertThumbprint (dir -dnsname
ces1.corp.contoso.com).Thumbprint -AuthenticationType Username
```

To install the Certificate Enrollment Web Service to use certificate authentication

1. To install a second instance of the Certificate Enrollment Web Service on CES1, you must use Windows PowerShell. Open Windows PowerShell.
2. Type **cd cert:\LocalMachine\My**, and then press ENTER.
3. Type **certutil**, and then press ENTER. Take note of the line that reads **Config**. This is the configuration that you will use when you install the Certificate Enrollment Web Service. For this lab, the configuration is APP1.corp.contoso.com\IssuingCA-APP1.

   Note

   The last line of the configuration output displays Web Enrollment Servers and shows https://ces1.corp.contoso.com/IssuingCA-APP1_CES_UsernamePassword/service.svc/CES, which is the URI that the Certificate Enrollment Policy Web Service will pass to the client during certificate enrollment.
4. Type the following command to install the Certificate Enrollment Web Service:

   ```
   Install-AdcsEnrollmentWebService -CAConfig
   ```

```
"APP1.corp.contoso.com\IssuingCA-APP1" -SSLCertThumbprint
(dir -dnsname ces1.corp.contoso.com).Thumbprint -
AuthenticationType Certificate -RenewalOnly -
AllowKeyBasedRenewal
```

5. Enter the password for the CES user account when you are prompted, and then press ENTER.
6. When you are prompted for confirmation, type **Y**, and then press ENTER.

Notes

a. You will see a confirmation that reads **ErrorString**. If the confirmation is blank under **ErrorString**, the installation succeeded. Otherwise, review your command for errors, correct them, and try again.

b. After the service is configured, type **certutil** again. You willsee that there are now two Web Enrollment Servers. The URI added in this procedure is https://ces1.corp.contoso.com/IssuingCA-APP1_CES_Certificate/services.svc/CES. This is the URI that the Certificate Enrollment Web Service will pass to the client during renewal.

Grant service account Read permission on the CA

1. On APP1, open the Certification Authority console as CORP\User1.
2. In the navigation pane, right-click **IssuingCA-APP1**, and then click **Properties**.
3. On the **Security** tab, click **Add**.
4. In **Enter the object names to select**, type **CES**, click **Check Names**, and then click **OK**.
5. Select **CES**, and then in **Permissions for CES**, select the check boxes that correspond to the **Allow** and **Read** permissions. Clear the check box that corresponds to **Allow** and **Request Certificates**, and then click **OK**.

Note

The **Authenticated Users** group has the **Request Certificates** permission set by default, and the **Authenticated Users** group includes all the computer accounts in the domain. This means that CES has **Request Certificates** permission through its membership in **Authenticated Users**.

Trust the service account for delegation

1. On DC1, open **Active Directory Users and Computers** as User1.
2. In the navigation pane, expand **corp.contoso.com**, and then click **Users**.
3. In the details pane, right-click the **CES** user account, and then click **Properties**.
4. On the **Delegation** tab, select **Trust this user for delegation to specified services only**. Select **Use any authentication protocol**, and then click **Add**.
5. In **Add Services**, click **Users or Computers**.
6. In **Select Users or Computers**, under **Enter the object names to select**, type **APP1**, click **Check Names**, and then click **OK**.

7. From the list of available services, select the **HOST** and **rpcss** services. Click **OK** twice.

   Tip
   You can hold the CTRL key to select multiple services in this interface.

8. Close **Active Directory Users and Computers**.

Step 5: Prepare an appropriate certificate template

For Certificate Enrollment Web Services to provide certificates to clients, an appropriate certificate template must be configured and published.

To prepare a certificate template

1. In the navigation pane of the Certification Authority console on APP1, expand IssuingCA-APP1.
2. Right-click **Certificate Templates**, and then click **Manage**. The Certificates Templates Console opens.
3. In **Template Display Name**, right-click the **Web Server** template, and then click **Duplicate Template**.
4. On the **Compatibility** tab, in **Compatibility Settings**, set **Certification Authority** to Windows Server 2012. When the **Resulting changes** dialog box appears, click **OK**.
5. Set **Certificate recipient** to Windows 8 / Windows Server 2012. When the **Resulting changes** dialog box appears, click **OK**.

   Note
   Setting the certification authority and the certificate client to Windows Server 2012 / Windows 8 allows key-based renewal, which enables the client to renew its certificate by using the existing certificate.

6. On the **General** tab, in **Template display name**, type **Internet Server** to rename the template. Set the **Validity period** to **1 years** and ensure the **Renewal period** is set to **6 weeks**.
7. On the **Security** tab, under **Group or user names**, select **Authenticated Users**, and then select the check box that corresponds to **Allow** and **Enroll** permission. This ensures that the template is visible to all members of the **Authenticated Users** group, which includes any account (user, computer, or device) that successfully authenticates to the domain.

   Note
   In a production environment, you may elect to further secure this template so that only members of a specific group can access the template.

8. On the **Extensions** tab, under **Extensions included in this template**, select **Application Policies**, and then click **Edit**.
9. In **Edit Application Policies Extension**, click **Add**.
10. In **Add Application Policy**, under **Application Policies**, double-click **Client**

Authentication. In **Edit Application Policies Extension**, click **OK**.

> **Note**
> Under **Description of Application Policies**, **Client Authentication** and **Server Authentication** should appear. **Client Authentication** allows a certificate to prove the identity of the certificate services client. **Server Authentication** allows a certificate to prove the identity of a web server.

11. On the **Subject Name** tab, select **Supply in the request**, and then select **Use subject information from existing certificates for autoenrollment and renewal request**.
12. On the **Issuance Requirements** tab, under **Require the following for enrollment**, select **CA certificate manager approval**. Under **Require the following for reenrollment**, select **Valid existing certificate**, and then select **Allow key based renewal**. Click **OK**.
13. Open Windows PowerShell as an Administrator. Type **certutil**, and then press ENTER. This shows you the CA configuration that is used in the following command.
14. Run the following command:

```
Certutil -config "APP1.corp.contoso.com\IssuingCA-APP1" -setreg policy\EditFlags
+EDITF_ENABLERENEWONBEHALFOF
```

15. Run the following command to restart the CA service to ensure that the configuration change is complete:

```
Restart-service certsvc
```

16. Close Windows PowerShell.
17. Close the Certificate Templates Console.
18. In the Certification Authority console, in the navigation console tree, click **Certificate Templates**. The details pane displays the issued certificate templates.
19. In the console tree, right-click **Certificate Templates**, click **New**, and then click **Certificate Template to Issue**.
20. In the **Enable Certificate Templates** dialog box, under **Name**, click **Internet Server**, and then click **OK**.

> **Notes**
> You can also enable the Internet Server certificate template in Windows PowerShell by running the following command:
>
> ```
> Add-CATemplate InternetServer
> ```

Step 6: Configure WEB1

In this step, you will configure WEB1 as a member of a workgroup that is connected to the CorpNet subnet. You will configure WEB1 to trust the root CA of corp.contoso.com. The procedures to complete this step are as follows:

1. Install the operating system
2. Configure TCP/IP

3. Configure WEB1 to trust the root CA

To install the operating system

1. Start the installation of Windows Server 2012
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Sign in by using the local Administrator account.
3. Connect the computer to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.
4. Connect the computer to the Corpnet subnet.

To configure TCP/IP

1. From Windows PowerShell, run **ncpa.cpl**.
2. In **Network Connections**, right-click **Ethernet**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.6**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**.
7. Close the **Network Connections** window.
8. In Windows PowerShell, run **sysdm.cpl**.
9. On the **System Properties** dialog box on the **Computer Name** tab, click **Change**.
10. In **Computer name**, type **WEB1** as the new name for the computer, and then click **OK**.
11. When you are prompted that you must restart the computer, click **OK**.
12. On the **System Properties** dialog box, click **Close**.
13. When you are prompted to restart the computer, click **Restart Now**.
14. After restarting, sign in by using the local Administrator account.

Tip

The Windows PowerShell commands change the IP address and rename the computer are as follows:

```
$NetIP = Get-NetIPAddress | where {$_.Addressfamily -eq "IPv4" -and
$_.InterfaceAlias -like "*Ethernet*"}

$NetAlias = $NetIP.InterfaceAlias

New-NetIPAddress -InterfaceAlias $NetAlias -IPAddress 10.0.0.6 -PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias $NetAlias -ServerAddresses 10.0.0.1

Set-DnsClient -InterfaceAlias $NetAlias -ConnectionSpecificSuffix corp.contoso.com

Rename-computer WEB1
```

```
Restart-computer
```

To configure WEB1 to trust the root CA

1. On WEB1, insert the removable storage device that contains the certificates for APP1 and ORCA1.

   Tip
   If you no longer have the removable storage device, you can copy the orca1_ORCA-CA.crt files to a removable storage device from the c:\pki folder on APP1. The storage device can be physical or virtual, as discussed in **Hardware and software requirements** earlier in this document.
2. On WEB1, open Windows PowerShell as an Administrator, and then type `mmc`.
3. Click **File**, and then click **Add/Remove Snap-in**.
4. In **Add or Remove Snap-ins**, click **Certificates**, and then click **Add**.
5. In **Certificates snap-in** select **Computer account**. Click **Next**.
6. In **Select Computer**, leave **Local computer** selected, click **Finish**, and then click **OK**.
7. In the navigation pane, expand **Certificates (Local Computer)**.
8. Right-click **Trusted Root Certification Authorities**, click **All Tasks**, and then click **Import**.
9. In the **Certificate Import Wizard**, click **Next**.
10. On the **File to Import** screen, in **File name**, type the path to the ORCA1 certificate that is on your removable storage device. For example, if the ORCA1 certificate is named orca1_ORCA1-ContosoRootCA.crt and on a floppy disk, you would type **A:\orca1_ORCA1-ContosoRootCA.crt**. You can alternatively use the **Browse** button to search for the certificate. Select orca1_ORCA1-ContosoRootCA.crt, and then click **Next**.
11. On the **Certificate Store** screen, select **Place all certificates in the following store** and set **Certificate store** to **Trusted Root Certification Authorities**. Click **Next**, and then click **Finish**. When the **Certificate Import Wizard** shows that the import was successful, click **OK**.
12. On **Console1**, click **File**, and then click **Save**. Ensure that **Save in** is set to **Desktop** (to save the console on the current user account's Desktop). In **File name**, type **Certificates** to change the console name from Console1 to Certificates. Click **OK**.

Step 7: Obtain a certificate and test automatic renewal

You will use user name and password authentication through Certificate Enrollment Web Services to request an initial certificate. Then, you will simulate the automatic renewal of that certificate by using the existing certificate. The procedures to complete this step are as follows:

1. Request a certificate
2. Approve the certificate request

3. Install the certificate
4. Configure WEB1 for automated certificate renewal
5. Test the certificate renewal

To request a certificate

1. To use certificate-based authentication, you must first obtain a certificate. On WEB1, sign in as the local Administrator. Run Windows PowerShell as Administrator. Run the following command to obtain the initial certificate by using user name and password authentication:

   ```
   Get-Certificate -template InternetServer -Url
   "https://cep1.corp.contoso.com/KeybasedRenewal_ADPolicyProvider_CEP_UsernamePasswo
   rd/service.svc/CEP" -SubjectName "CN=WEB1" -DnsName "web1.treyresearch.com" -
   Credential (Get-Credential) -CertStoreLocation "cert:\LocalMachine\My"
   ```

   Enter Corp\User1 credentials when you are prompted.

   Note
   The request will be pending, and you must approve the request on APP1. If the request does not complete successfully, check your command syntax and try again. If the request times out, try again.

To approve the certificate request

1. On APP1, as User1, open the Certification Authority console. In the navigation pane, click **Pending Requests**.
2. In the details pane, make a note of the Request ID number, and then right-click the pending request. Click **All Tasks**, and then click **Issue**.

To install the certificate

1. On WEB1, run the following Windows PowerShell commands to retrieve the certificate

   ```
   Cd Cert:\LocalMachine\Request
   ```

   ```
   Dir | Get-Certificate -Credential (Get-Credential)
   ```

   Enter your Corp\User1 credentials when you are prompted. If the request does not complete successfully the first time, check your command syntax and try again. If the request times out, try again.

To configure WEB1 for automated certificate renewal

1. On WEB1, open the Local Group Policy Editor console. To do so, open Windows PowerShell as an Administrator and type the following command `gpedit.msc`.
2. In the Local Group Policy Editor navigation pane, expand **Local Computer Policy**, expand **Computer Configuration**, expand **Windows Settings**, expand **Security Settings**, and then click **Public Key Policies**.

3. In the details pane, double-click **Certificate Services Client – Auto-Enrollment**.
4. On the **Enrollment Policy Configuration** tab, set **Configuration Model** to **Enabled**. Select **Renew expired certificates, update pending certificates, and remove revoked certificates**, select **Update certificates that use certificate templates**, and then click **OK**.
5. In the Local Group Policy Editor console details pane, double-click **Certificate Services Client – Certificate Enrollment Policy**.
6. Set **Configuration Model** to **Enabled**.
7. On the **Enrollment Policy** tab, click **Add**.
8. In **Certificate Enrollment Policy Server**, in the **Enter enrollment policy server URI** text box, enter the following URI:

   https://cep1.corp.contoso.com/KeyBasedRenewal_ADPolicyProvider_CEP_Certificate/service.svc/CEP
9. Set **Authentication type** to **X.509 Certificate**.
10. Click **Validate Server**. **Windows Security** displays the web1.treyresearch.com certificate. Click **OK**.

    Note

    If you receive an operation timeout, ensure that the CEP1 and CES1 servers are online and then retry.
11. When the path is successfully validated, click **Add**.
12. On the **Enrollment Policy** tab, in the **Certificate enrollment policy** list, select the **Default** check box for **SSL Server Certificates**, and then click **OK**.

To test the certificate renewal

1. On WEB1, run the following command from Windows PowerShell as an Administrator:

   ```
   Cd Cert:\LocalMachine\My
   Dir | format-list
   ```

   Copy the certificate thumbprint from the output. (You can copy by selecting the text and right-clicking.)
2. Run the following command in to delete the policy cache:

   ```
   certutil -f -policyserver * -policycache delete
   ```
3. Run the following command to renew the certificate. Replace <thumbprint> with the actual characters of the certificate thumbprint that you copied. (You can paste by right-clicking.)

   ```
   certreq -machine -q -enroll -cert <thumbprint> renew
   ```

   Note

   If the operation times out, try again. If you run into other errors, ensure that the CEP1 and CES1 servers are online by running the command **iisreset** from Windows PowerShell on the CES1 and CEP1 servers and then retry.

4. On WEB1, run the following command from Windows PowerShell as an Administrator:

```
Cd Cert:\LocalMachine\My
Dir | format-list
```

Note that the certificate thumbprint has changed. This demonstrates that the certificate was successfully renewed.

Note

To comment on this content or ask questions about the information presented here, please use our Feedback guidance.

See Also

Windows Server Security Forum

Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Frequently Asked Questions (FAQ)

Windows PKI Documentation Reference and Library

Windows PKI Blog

Active Directory Domain Services Overview

By using the Active Directory® Domain Services (AD DS) server role, you can create a scalable, secure, and manageable infrastructure for user and resource management, and provide support for directory-enabled applications such as Microsoft® Exchange Server.

The rest of this topic explains a high-level overview of the AD DS server role. For more information about new features in AD DS in Windows Server 2012, see What's New in Active Directory Domain Services (AD DS).

AD DS provides a distributed database that stores and manages information about network resources and application-specific data from directory-enabled applications. A server that is running AD DS is called a domain controller. Administrators can use AD DS to organize elements of a network, such as users, computers, and other devices, into a hierarchical containment structure. The hierarchical containment structure includes the Active Directory forest, domains in the forest, and organizational units (OUs) in each domain.

Organizing network elements into a hierarchical containment structure provides the following benefits:

- The forest acts as a security boundary for an organization and defines the scope of authority for administrators. By default, a forest contains a single domain, which is known as the forest root domain.
- Additional domains can be created in the forest to provide partitioning of AD DS data, which enables organizations to replicate data only where it is needed. This makes it possible for AD DS to scale globally over a network that has limited available bandwidth. An

Active Directory domain also supports a number of other core functions that are related to administration, including network-wide user identity, authentication, and trust relationships.

- OUs simplify the delegation of authority to facilitate the management of large numbers of objects. Through delegation, owners can transfer full or limited authority over objects to other users or groups. Delegation is important because it helps to distribute the management of large numbers of objects to a number of people who are trusted to perform management tasks.

Security is integrated with AD DS through logon authentication and access control to resources in the directory. With a single network logon, administrators can manage directory data and organization throughout their network. Authorized network users can also use a single network logon to access resources anywhere in the network. Policy-based administration eases the management of even the most complex network.

Additional AD DS features include the following:

- A set of rules, the schema, that defines the classes of objects and attributes that are contained in the directory, the constraints and limits on instances of these objects, and the format of their names.
- A global catalog that contains information about every object in the directory. Users and administrators can use the global catalog to find directory information, regardless of which domain in the directory actually contains the data.
- A query and index mechanism, so that objects and their properties can be published and found by network users or applications.
- A replication service that distributes directory data across a network. All writable domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain.
- Operations master roles (also known as flexible single master operations or FSMO). Domain controllers that hold operations master roles are designated to perform specific tasks to ensure consistency and eliminate conflicting entries in the directory.

Requirements for running Active Directory Domain Services

What hardware, software, or settings configurations are required for running this feature? What prerequisites are there for running the role? Does this role/feature require special hardware?

Requirement	Description
TCP/IP	Configure appropriate TCP/IP and DNS server addresses.
NTFS	The drives that store the database, log files, and SYSVOL folder for Active Directory Domain Services (AD DS) must be placed on a

Requirement	Description
	local fixed volume. SYSVOL must be placed on a volume that is formatted with the NTFS file system. For security purposes, the Active Directory database and log files should be placed on a volume that is formatted with NTFS.
Credentials	To install a new AD DS forest, you need to be local Administrator on the server. To install an additional domain controller in an existing domain, you need to be a member of the Domain Admins group.
Domain Name System (DNS) infrastructure	Verify that a DNS infrastructure is in place. When you install AD DS, you can include DNS server installation, if it is needed. When you create a new domain, a DNS delegation is created automatically during the installation process. Creating a DNS delegation requires credentials that have permissions to update the parent DNS zones. For more information, see DNS Options wizard page.
Adprep	To add the first domain controller that runs Windows Server 2012 to an existing Active Directory, adprep.exe commands run automatically as needed. These commands have additional credential and connectivity requirements. For more information, see Running Adprep.exe.
Read-only domain controllers (RODCs)	Additional requirements to install RODCs: • Forest functional level must be at least Windows Server 2003 • At least one writable domain controller that runs Windows Server 2008 or later must be installed in the same domain. For more information, see Prerequisites for Deploying an RODC.

Note

With the exception of DNS server, domain controllers generally should not host other server roles.

Running Active Directory Domain Services

How do I deploy and configure this role by using Windows PowerShell?

For step-by-step instructions for how to install and configure AD DS by using the ADDSDeployment module for Windows PowerShell® command-line interface, see Active Directory Domain Services Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=222597).

How do I deploy and configure this role in a multi-server environment?

AD DS is a distributed service that is designed to run on multiple domain controllers. For step-by-step instructions for how to install and configure AD DS on multiple domain controllers, see Active Directory Domain Services Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=222597).

How can I run this role on virtual machines?

AD DS in Windows Server 2012 includes safeguards for running on virtual machines to ensure safety and consistency of virtualized AD DS environments. For more information about how to run AD DS on virtual machines, see Running Domain Controllers in Hyper-V (http://go.microsoft.com/fwlink/?LinkID=213293).

Security considerations for running this role

After installation, AD DS is designed to be secure by default. For more information about default security settings for domain controllers, risks, and how to operate domain controllers securely, see Best Practice Guide for Securing Active Directory Installations.

Special considerations for managing this role remotely

To manage AD DS remotely, install the Remote Server Administration Tools (RSAT). There is a 32-bit version and a 64-bit version of RSAT. For more information, see Remote Server Administration Tools (http://go.microsoft.com/fwlink/?LinkId=222628).

Special considerations for managing the role on the Server Core installation option

AD DS can be installed on a Server Core installation or a server with a Minimal Server Interface, and is recommended in cases where reducing the footprint of the operating system installation is advantageous, such as for a dedicated server role in a datacenter, for virtualization guests, or RODCs in remote offices. Beginning with Windows Server 2012, a domain controller that runs on

a Server Core installation can be converted to server installation with a GUI (also known as a full installation) and vice versa.

Upgrade from a Server Core installation running on a previous version of Windows Server is supported, but there is no way to upgrade directly from a Server Core installation of a previous version of Windows Server to a server installation with a GUI or directly from a server installation with a GUI to a Server Core installation. In this case, you need to upgrade directly to the same installation type on Windows Server 2012 and then convert to a different installation after the upgrade as needed.

For more information, see Windows Server Installation Options.

Role services for Active Directory Domain Services

Identity Management for UNIX is a role service of AD DS that can be installed only on domain controllers. Two Identity Management for UNIX technologies, Server for NIS and Password Synchronization, make it easier to integrate computers running Windows® into your existing UNIX enterprise. AD DS administrators can use Server for NIS to manage Network Information Service (NIS) domains. Password Synchronization automatically synchronizes passwords between Windows and UNIX operating systems.

Role service technologies	Role service description
Server for NIS	Enables a Microsoft Windows–based Active Directory domain controller to administer UNIX Network Information Service (NIS) networks. For more information, see Overview of Server for NIS (http://go.microsoft.com/fwlink/?LinkId=222677).
Password Synchronization	Helps integrate Windows and UNIX networks by simplifying the process of maintaining secure passwords in both environments. For more information, see Overview of Password Synchronization (http://go.microsoft.com/fwlink/?LinkId=222676).

Additional references

- Installing AD DS by Using Windows PowerShell
- Active Directory Domain Services Command Reference
- Best Practices Analyzer for Active Directory Domain Services
- Troubleshooting Active Directory Domain Services

What's New in Active Directory Domain Services (AD DS)

You can use Active Directory Domain Services (AD DS) in Windows Server 2012 to more rapidly and easily deploy domain controllers (on-premises and in the cloud), increase flexibility when auditing and authorizing access to files, and more easily perform administrative tasks at scale (locally or remotely) through consistent graphical and scripted management experiences. AD DS improvements in Windows Server 2012 include:

- Virtualization that just works

  Windows Server 2012 provides greater support for the capabilities of public and private clouds through virtualization-safe technologies and the rapid deployment of virtual domain controllers through cloning.

- Simplified deployment and upgrade preparation

  The upgrade and preparation processes (dcpromo and adprep) have been replaced with a new streamlined domain controller promotion wizard that is integrated with Server Manager and built on Windows PowerShell. It validates prerequisites, automates forest and domain preparation, requires only a single set of logon credentials, and it can remotely install AD DS on a target server.

- Simplified management

  Examples of simplified management include the integration of claims-based authorization into AD DS and the Windows platform, two critical components of a broader feature known as Dynamic Access Control (DAC). DAC comprises central access policies, directory attributes, the Windows file-classification engine, and compound-identities that combine user and machine identity into one. In addition, the Active Directory Administrative Center (ADAC) now allows you to perform graphical tasks that automatically generate the equivalent Windows PowerShell commands. The commands can be easily copied and pasted into a script simplifying the automation of repetitive administrative actions.

- AD DS Platform Changes

  The AD DS platform comprises core functionality, including the “under-the-covers” behaviors that govern the components upon which the rest of the directory service is built. Updates to the AD DS platform include improved allocation and scale of RIDs (relative identifiers), deferred index creation, various Kerberos enhancements and support for Kerberos claims (see Dynamic Access Control) in AD FS.

Active Directory and AD DS has been at the center of IT infrastructure for over 10 years, and its features, adoption, and business-value have grown release over release. Today, the majority of that Active Directory infrastructure remains on the premises, but there is an emerging trend toward cloud computing. The adoption of cloud computing, however, will not occur overnight, and migrating suitable on-premises workloads or applications is an incremental and long-term exercise. New hybrid infrastructures will emerge, and it is essential that AD DS support the needs of these new and unique deployment models that include services hosted entirely in the cloud,

services that comprise cloud and on-premises components, and services that remain exclusively on the premises. These hybrid models will increase the importance, visibility, and emphasis around security and compliance, and they will compound the already complex and time-consuming exercise of ensuring that access to corporate data and services is appropriately audited and accurately expresses the business intent.

The following sections describe how AD DS in Windows Server 2012 addresses these emerging needs.

For more information about installing AD DS, see Deploy Active Directory Domain Services (AD DS) in Your Enterprise and Upgrade Domain Controllers to Windows Server 2012.

Virtualization that just works

Rapid deployment with cloning

AD DS in Windows Server 2012 allows you to deploy replica virtual domain controllers by "cloning" existing virtual domain controllers. You can promote a single virtual domain controller by using the domain controller promotion interface in Server Manager, and then rapidly deploy additional virtual domain controllers within the same domain, through cloning.

The process of cloning involves creating a copy of an existing virtual domain controller, authorizing the source domain controller to be cloned in AD DS, and running Windows PowerShell cmdlets to create a configuration file that contains detailed promotion instructions (name, IP address, Domain Name System [DNS] servers, and so on). Or you can leave the configuration file empty, which allows the system to automatically fill in the information. Cloning reduces the number of steps and time involved by eliminating repetitive deployment tasks, and it enables you to fully deploy additional domain controllers that are authorized and configured for cloning by the Active Directory domain administrator.

For detailed information about virtualized domain controller cloning, see Active Directory Domain Services (AD DS) Virtualization.

Safer virtualization of domain controllers

AD DS has been virtualized for several years, but features present in most hypervisors can invalidate strong assumptions made by the Active Directory replication algorithms. Primarily, the logical clocks that are used by domain controllers to determine relative levels of convergence only go forward in time. In Windows Server 2012, a virtual domain controller uses a unique identifier that is exposed by the hypervisor. This is called the virtual machine GenerationID. The virtual machine GenerationID changes whenever the virtual machine experiences an event that affects its position in time. The virtual machine GenerationID is exposed to the virtual machine's address space within its BIOS, and it is made available to the operating system and applications through a driver in Windows Server 2012.

During boot and before completing any transaction, a virtual domain controller running Windows Server 2012 compares the current value of the virtual machine GenerationID against the value that it stored in the directory. A mismatch is interpreted as a "rollback" event, and the domain

controller employs AD DS safeguards that are new in Windows Server 2012. These safeguards allow the virtual domain controller to converge with other domain controllers, and they prevent the virtual domain controller from creating duplicate security principals. For Windows Server 2012 virtual domain controllers to gain this extra level of protection, the virtual domain controller must be hosted on a virtual machine GenerationID–aware hypervisor such as Windows Server 2012 with the Hyper-V role.

For detailed information about the virtualization-safe technology feature, see Active Directory Domain Services (AD DS) Virtualization.

Simplified deployment and upgrade preparation

AD DS deployment in Windows Server 2012 integrates all the required steps to deploy new domain controllers into a single graphical interface. It requires only one enterprise-level credential, and it can prepare the forest or domain by remotely targeting the appropriate operations master roles. The new deployment process conducts extensive prerequisite validation tests that minimize the opportunity for errors that might have otherwise blocked or slowed the installation. The AD DS installation process is built on Windows PowerShell, integrated with Server Manager, able to target multiple servers, and remotely deploy domain controllers, which results in a deployment experience that is simpler, more consistent, and less time consuming. The following figure shows the AD DS Configuration Wizard in Windows Server 2012.

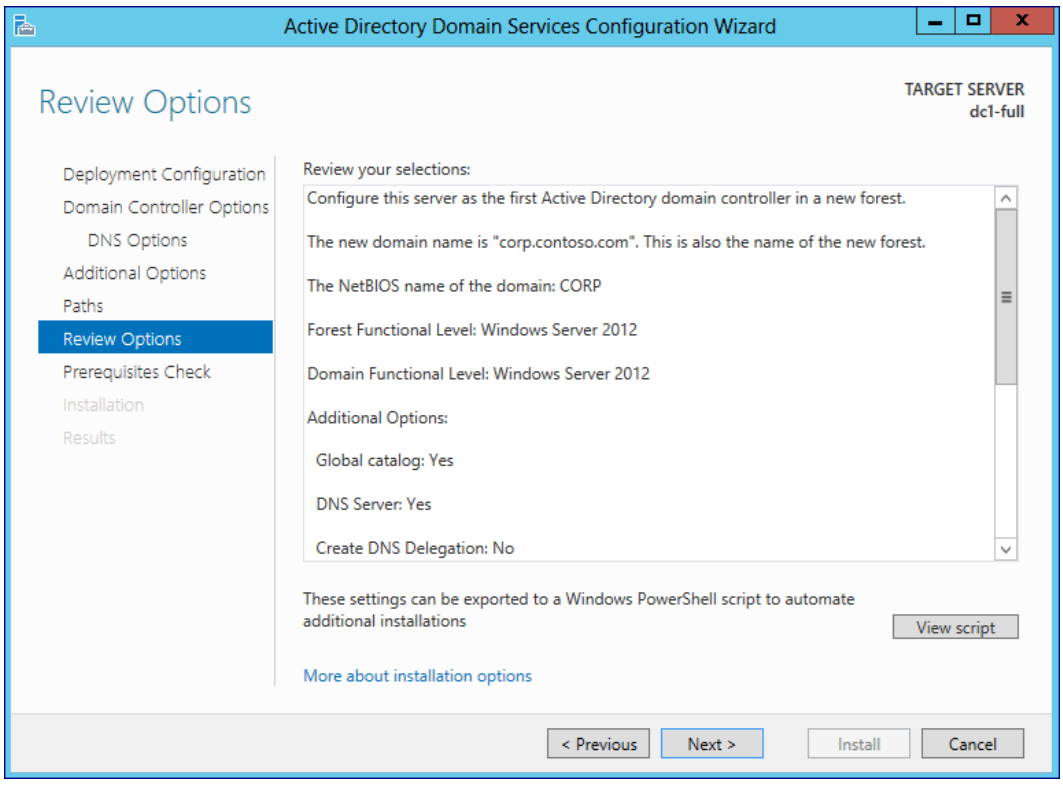

Figure 1 AD DS Configuration Wizard

An AD DS installation includes the following features:

- **Adprep.exe integration into the AD DS installation process.** Reduces the time required to install AD DS and reduces the chances for errors that might block domain controller promotion.
- **The AD DS server role installation, which is built on Windows PowerShell and can be run remotely on multiple servers.** Reduces the likelihood of administrative errors and the overall time that is required for installation, especially when you are deploying multiple domain controllers across global regions and domains.
- **Prerequisite validation in the AD DS Configuration Wizard.** Identifies potential errors before the installation begins. You can correct error conditions before they occur without the concerns that result from a partially complete upgrade.
- **Configuration pages grouped in a sequence that mirror the requirements of the most common promotion options, with related options grouped in fewer wizard pages.** Provides better context for making installation choices and reduces the number of steps and time that are required to complete the domain controller installation.
- **A wizard that exports a Windows PowerShell script that contains all the options that were specified during the graphical installation.** Simplifies the process by automating subsequent AD DS installations through automatically generated Windows PowerShell scripts.

For detailed information about AD DS integration with Server Manager see Deploy Active Directory Domain Services (AD DS) in Your Enterprise.

Simplified management

Numerous areas were addressed with a view towards simplifying AD DS management experience. These areas include:

- Dynamic Access Control
- DirectAccess Offline Domain Join
- Active Directory Federation Services (AD FS)
- Windows PowerShell History Viewer
- Active Directory Recycle Bin User Interface
- Fine-Grained Password Policy User Interface
- Active Directory Replication and Topology Windows PowerShell cmdlets
- Active Directory Based Activation (AD BA)
- Group Managed Service Accounts (gMSA)

Dynamic Access Control

Today, it is difficult to translate business-intent using the existing authorization model. The existing capabilities of access control entries (ACEs) make it hard or impossible to fully express requirements. In addition, there are no central administration capabilities. Finally, modern-day

increases in regulatory and business requirements around compliance further compound the problem.

Windows Server 2012 AD DS addresses these challenges by introducing:

- A new claims-based authorization platform that enhances, not replaces, the existing model, which includes:
  - User-claims and device-claims
  - User + device claims (also known as compound identity)
- New central access policies (CAP) model
- Use of file-classification information in authorization decisions
- Easier access-denied remediation experience
- Access policies and audit policies can be defined flexibly and simply:
  - IF `resource.Confidentiality` = `high` THEN `audit.Success` WHEN `user.EmployeeType` = `vendor`

Requirements

- One or more Windows Server 2012 domain controllers
- Windows Server 2012 file server
- Enable the claims-policy in the Default Domain Controllers Policy
- Windows Server 2012 Active Directory Administrative Center
- For device-claims, compound ID must be switched on at the target service account by using Group Policy or editing the object directly

For more information about Dynamic Access Control see the Dynamic Access Control section of the technical library.

DirectAccess Offline Domain Join

The offline domain-join feature that was added to AD DS in Windows Server 2008 R2 effectively allows client computers to be joined to a domain without requiring network connectivity to a domain controller, but the client computer could not also be preconfigured for DirectAccess as part of the domain join.

Windows Server 2012 AD DS provides the following improvements:

- Extends offline domain-join by allowing the blob to accommodate DirectAccess prerequisites
  - Certs
  - Group Policies
- What does this mean?
  - A computer can now be domain-joined over the Internet if the domain is DirectAccess enabled
  - Getting the blob to the non-domain-joined machine is an offline process and the responsibility of the administrator

Requirements

- Windows Server 2012 domain controllers

For more information, see DirectAccess Offline Domain Join.

Active Directory Federation Services (AD FS)

AD FS v2.0 shipped out-of-band of the Windows Server release. In Windows Server 2012, AD FS (v2.1) ships in-the-box as a server role. This provides:

- Simplified trust-setup and automatic trust management
- SAML-protocol support
- Extensible attribute store
- Allows claims to be sourced from anywhere in the enterprise
- Active Directory Lightweight Directory Service (AD LDS) and SQL attribute-store providers supplied out-of-the-box

Requirements

- Windows Server 2012

For detailed information about AD FS in Windows Server 2012, see AD FS.

Windows PowerShell History Viewer

Windows PowerShell is a key technology in creating a consistent experience between the command-line and the graphical user interface. Windows PowerShell increases productivity, but also requires investment in learning how to use it.

To minimize the learning investment, Windows Server 2012 includes the new Windows PowerShell History Viewer. The benefits include:

- Allow administrators to view the Windows PowerShell commands executed when using the Active Directory Administrative Center. For example:
  - The administrator adds a user to a group
  - The UI displays the equivalent Windows PowerShell for Active Directory command
  - The administrator copies the resulting syntax and integrates it into a script
  - Reduces Windows PowerShell learning-curve
  - Increases confidence in scripting
  - Further enhances Windows PowerShell discoverability

Requirements

- Windows Server 2012 Active Directory Administrative Center

For more information about the Windows PowerShell History Viewer, see Active Directory Administrative Center Enhancements.

Active Directory Recycle Bin User Interface

The Active Directory Recycle Bin feature introduced with Windows Server® 2008 R2 provided an architecture permitting complete object recovery. Scenarios that require object recovery by using the Active Directory Recycle Bin are typically high-priority, such as recovery from accidental

deletions, for example, resulting in failed logons or work stoppages. But the absence of a rich, graphical user interface complicated its usage and slowed recovery.

To address this challenge, Windows Server 2012 AD DS has a user interface for the Active Directory Recycle Bin that provides the following advantages:

- Simplifies object recovery through the inclusion of a **Deleted Objects** node in the Active Directory Administrative Center (ADAC)
  - Deleted objects can now be recovered within the graphical user interface
- Reduces recovery-time by providing a discoverable, consistent view of deleted object

Requirements

- Recycle Bin requirements must be met:
  - Windows Server 2008 R2 forest functional level
  - Recycle Bin optional-feature must be enabled
- Windows Server 2012 Active Directory Administrative Center
- Objects requiring recovery must have been deleted within Deleted Object Lifetime (DOL)
  - By default, DOL is set to 180 days

For more information about the user interface for AD DS Recycle Bin, see Active Directory Administrative Center Enhancements.

Fine-Grained Password Policy User Interface

The Fine-Grained Password Policy (FGPP) introduced with Windows Server 2008 provided more precise management of password-policies. In order to leverage the feature, administrators had to manually create password-settings objects (PSOs). It proved difficult to ensure that the manually defined policy-values behaved as desired, which resulted in time-consuming, trial and error administration.

In Windows Server 2012:

- Creating, editing and assigning PSOs now managed through the Active Directory Administrative Center
- Greatly simplifies management of password-settings objects

Requirements

- FGPP requirements must be met:
  - Windows Server® 2008 domain functional level
- Windows Server 2012 Active Directory Administrative Center

For more information about the user interface for fine-grained password policies, see Active Directory Administrative Center Enhancements.

Active Directory Replication and Topology Windows PowerShell cmdlets

Administrators require a variety of tools to manage Active Directory's site topology

- repadmin
- ntdsutil

- Active Directory Sites and Services

The usage of multiple tools results in an inconsistent experience that is difficult to automate.

Using Windows Server 2012 AD DS, administrators can:

- Manage replication and site-topology with Windows PowerShell
  - Create and manage sites, site-links, site-link bridges, subnets and connections
  - Replicate objects between domain controllers
  - View replication metadata on object attributes
  - View replication failures
- Take advantage of a consistent and easily scriptable experience
- Compatible and interoperable with other Windows PowerShell cmdlets

Requirements

- Active Directory Web Service (also known as Active Directory Management Gateway for Windows Server 2003 or Windows Server 2008)
- Windows Server 2012 domain controller or Windows Server 2012 with the Role Administration Tools (RSAT) for AD DS and AD LDS installed

For more information about the Windows PowerShell cmdlets to manage Active Directory topology and replication, see Active Directory Replication and Topology Management Using Windows PowerShell.

Active Directory Based Activation (AD BA)

Today, Volume Licensing for Windows and Office requires Key Management Service (KMS) servers. That solution requires minimal training, and is a turnkey solution that covers about 90% of deployments.

But there is complexity caused by the lack of a graphical administration console. The solution requires RPC traffic on the network, which complicates matters, and it does not support any kind of authentication. The end-user licensing agreement (EULA) prohibits the customer from connecting the KMS server to any external network. For example, connectivity-alone to the service equates to *activated*.

In Windows Server 2012, the Active Directory-based activation provides the following improvements:

- Uses your existing Active Directory infrastructure to activate your clients
  - No additional machines required
  - No RPC requirement; uses LDAP exclusively
  - Includes RODCs
- Beyond installation and service-specific requirements, no data is written back to the directory
  - Activating initial CSVLK (customer-specific volume license key) requires:
    - One-time contact with Microsoft Activation Services over the Internet (identical to retail activation)
    - Key entered using volume activation server role or using command line.

    - Repeat the activation process for additional forests up to 6 times by default
- Activation-object maintained in configuration partition
    - Represents proof of purchase
    - Computers can be member of any domain in the forest
- All Windows 8 computers will automatically activate

Requirements

- Only Windows 8 computers can leverage AD BA
- KMS and AD BA can coexist
    - You still need KMS if you require down-level volume-licensing
- Requires Windows Server 2012 Active Directory schema, not Windows Server 2012 domain controllers

For more information about AD BA see the following:

- Volume Activation Overview
- Test Lab Guide: Demonstrate Volume Activation Services

Group Managed Service Accounts (gMSA)

Managed Service Accounts (MSAs) were introduced with Windows Server 2008 R2. Clustered or load-balanced services that needed to share a single security-principal were unsupported. As a result, MSAs were not able to be used in many desirable scenarios.

Windows Server 2012 includes the following changes:

- Introduces a new security principal type known as a gMSA
- Services running on multiple hosts can run under the same gMSA account
- One or more Windows Server 2012 domain controllers required
    - gMSAs can authenticate against any domain controllers that run any version of Windows Server
    - Passwords computed by Group Key Distribution Service (GKDS) running on all Windows Server 2012 domain controllers
- Windows Server 2012 hosts using gMSAs obtain password and password-updates from GKDS
    - Password retrieval limited to authorized computers
- Password-change interval defined at gMSA account creation (30 days by default)
- Like MSAs, gMSAs are supported only by the Windows Service Control Manager (SCM) and IIS application pools

Requirements

- Windows Server 2012 Active Directory schema updated in forests containing gMSAs
- One or more Windows Server 2012 domain controllers to provide password computation and retrieval
- Only services running on Windows Server 2012 can use gMSAs

For more information about group managed service accounts. see Managed Service Accounts.

AD DS Platform Changes

Numerous platform changes were made around scalability, throttling, and security. These areas include:

- AD DS Claims in AD FS
- Relative ID (RID) Improvements
- Deferred Index Creation
- Kerberos Enhancements

AD DS Claims in AD FS

AD FS v2.0 is able to generate user-claims directly from Windows NT tokens. AD FS v2.0 was also capable of further expanding claims based on attributes in AD DS and other attribute stores.

In Windows Server 2012, Kerberos tickets can be populated with user and device attributes serving as claims. AD FS 2.0 cannot read claims from Kerberos tickets. Therefore, a separate LDAP call to Active Directory must be made to source user-attribute claims, and AD FS 2.0 cannot leverage device-attribute claims at all.

AD FS v2.1 in Windows Server 2012 is able to populate SAML tokens with user- and device-claims taken directly from the Kerberos ticket.

Requirements

- Dynamic Access Control enabled and configured
- Compound ID must be switched on for the AD FS service account
- Windows Server 2012 AD FS v2.1

For detailed information about AD FS in Windows Server 2012, see AD FS.

Relative ID (RID) Improvements

The following RID improvements in Windows Server 2012 provide greater ability to react to any potential exhaustion of the global RID pool space:

- **Periodic RID consumption warning**
  - At 10% of remaining global space, system logs informational event
    - First event at 100,000,000 RIDs used, second event logged at 10% of remainder
      - Remainder = 900,000,000
      - 10% of remainder = 90,000,000
    - Second event logged at 190,000,000
      - Existing RID consumption plus 10% of remainder
  - Events become more frequent as the global space is further depleted
- **RID Manager artificial ceiling protection mechanism**
  - A soft ceiling that is 90% of the global RID space and is not configurable
  - The soft ceiling is deemed as ”reached” when a RID pool containing the 90% RID is issued

- Blocks further allocations of RID pools
  - When the ceiling is reached, system sets **msDS-RIDPoolAllocationEnabled** attribute of the **RID Manager$** object to **FALSE**. An administrator must set it back to **TRUE** to override.
- Log an event indicating that the ceiling is reached
  - An initial warning is logged when the global RID spaces reaches 80%
- The attribute can only be set to **FALSE** by the **SYSTEM** and is mastered by the RID master (for example, write it against the RID master)
  - Domain Admin can set it back to TRUE

    Note
    It is set to TRUE by default

- **Increased the global RID space per domain, doubling the number of security principals that can be created throughout the lifetime of a domain from 1 billion to 2 billion.**

Requirements

- Windows Server 2012 RID master
- Windows Server 2012 Domain Controllers

For more information on RID improvements, see Managing RID Issuance.

Deferred Index Creation

In the past, index creation could adversely impact domain controller performance. Windows Server 2012 introduces a new capability that allows forest administrators to defer index creation to a point in time they choose. By default, domain controllers create indices when they receive the appropriate schema change through replication. In Windows Server 2012, a new DSheuristic was introduced to control whether or not domain controllers defer index creation. The details are as follows:

- Setting the 19th byte to 1 causes any Windows Server 2012 DC (DCs that run earlier operating systems will ignore the setting) to defer building indices until:
  - It receives the **UpdateSchemaNow** rootDSE mod (triggers rebuild of the schema cache)
  - It is rebooted (which requires that the schema cache be rebuilt and, in turn, the deferred indices)
- Any attribute that is in a deferred index state will be logged in the Event Log every 24 hours
  - **2944: Index deferred** – logged once
  - **2945: Index still pending** – logged every 24 hours
  - **1137: Index created** – logged once (not a new event)

Requirements

- Windows Server 2012 domain controllers

Kerberos Enhancements

- Kerberos Constrained Delegation across domains

- Flexible Authentication Secure Tunneling (FAST)

Kerberos Constrained Delegation across domains

Kerberos Constrained Delegation (KCD) was introduced with Windows Server 2003. KCD permits a service's account (*front-end*) to act on the behalf of users in multi-tier applications for a limited set of *back-end* services. For example:

1. User accesses web site as *user1*
2. User requests information from web site (*front-end*) that requires the web server to query a SQL database *(back-end*)
3. Access to this data is authorized according to who accessed the *front-end*
4. In this case, the web service must impersonate *user1* when making the request to SQL

The front-end needed to be configured with the services (by SPN) to which it can impersonate users. Setup and administration requires Domain Admin credentials. KCD delegation only works for back-end services in the same domain as the front-end service-accounts.

The KCD in Windows Server 2012 moves the authorization decision to the resource-owners, which provides these advantages:

- Permits back-end to authorize which front-end service-accounts can impersonate users against their resources
- Supports across-domain, across-forest scenarios
- No longer requires Domain Admin privileges
    - Requires only administrative permission to the back-end service-account

Requirements

- Clients run Windows XP or later
- Client domain's domain controllers running Windows Server 2003 or later
- Front-end server running Windows Server 2012
- One or more domain controllers in front-end domain running Windows Server 2012
- One or more domain controllers in back-end domain running Windows Server 2012
- Back-end server account configured with the accounts that are permitted for impersonation
    - Not exposed through Active Directory Administrative Center
    - Configured through Windows PowerShell:
        - `New/Set-ADComputer [-name] <string> [-PrincipalsAllowedToDelegateToAccount <ADPrincipal[]>]`
        - `New/Set-ADServiceAccount [-name] <string> [-PrincipalsAllowedToDelegateToAccount <ADPrincipal[]>]`
- Windows Server 2012 schema update in back-end server's forest
- Back-end application server running Windows Server 2003 or later

For more information about Kerberos constrained delegation see the Kerberos section of the technical library.

Flexible Authentication Secure Tunneling (FAST)

Today, offline dictionary attack against password-based logons is possible. There is a relatively well-known concern around Kerberos errors being spoofed. Clients may:

- Fallback to less-secure legacy protocols
- Weaken their cryptographic key strength and/or ciphers

Kerberos in Windows Server 2012 supports Flexible Authentication Secure Tunneling (FAST)

- Defined by RFC 6113
- Sometimes referred to as Kerberos armoring
- Provides a protected channel between a domain-joined client and DC
    - Protects pre-authentication data for user's AS_REQs
        - Uses LSK (logon session key) from computer's TGT as shared secret
        - Note that computer authentication is NOT armored
    - Allows DCs to return authenticated Kerberos errors thereby protecting them from spoofing
- Once all Kerberos clients and DCs support FAST (the admin's decision to make)
    - The domain can be configured to either require Kerberos armoring or use it upon request
        - Must first ensure all or enough DCs are running Windows Server 2012
        - Enable the appropriate policy
            - "Support CBAC and Kerberos armoring"
            - "All DCs can support CBAC and Require Kerberos armoring"

Requirements

- Windows Server 2012 servers
- Ensure that all domains the client uses including transited referral domains:
    - Enable the "Support CBAC and Kerberos armoring" policy for all Windows Server 2012 DCs
    - Have a sufficient number of Windows Server 2012 DCs to support FAST
- Enable "Require FAST" policy on supported clients
- RFC-compliant FAST interoperability requires Windows Server 2012 domain functional level

Active Directory Domain Services (AD DS) Virtualization

This topic lists resources that are available for using virtualized domain controllers.

- Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100)
- Virtualized Domain Controller Technical Reference (Level 300)
- Virtualized Domain Controller Cloning Test Guidance for Application Vendors
- Support for using Hyper-V Replica for virtualized domain controllers

Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100)

Virtualization of Active Directory Domain Services (AD DS) environments has been ongoing for a number of years. Beginning with Windows Server® 2012, AD DS provides greater support for virtualizing domain controllers by introducing virtualization-safe capabilities and enabling rapid deployment of virtual domain controllers through cloning. These new virtualization features provide greater support for public and private clouds, hybrid environments where portions of AD DS exist on-premises and in the cloud, and AD DS infrastructures that reside completely on-premises.

In this document

- Safe virtualization of domain controllers
- Virtualized domain controller cloning
- Steps for deploying a clone virtualized domain controller
- Troubleshooting

Safe virtualization of domain controllers

Virtual environments present unique challenges to distributed workloads that depend upon a logical clock-based replication scheme. AD DS replication, for example, uses a monotonically increasing value (known as a USN or Update Sequence Number) assigned to transactions on each domain controller. Each domain controller's database instance is also given an identity, known as an InvocationID. The InvocationID of a domain controller and its USN together serve as a unique identifier associated with every write-transaction performed on each domain controller and must be unique within the forest.

AD DS replication uses InvocationID and USNs on each domain controller to determine what changes need to be replicated to other domain controllers. If a domain controller is rolled back in time outside of the domain controller's awareness and a USN is reused for an entirely different transaction, replication will not converge because other domain controllers will believe they have already received the updates associated with the re-used USN under the context of that InvocationID.

For example, the following illustration shows the sequence of events that occurs in Windows Server 2008 R2 and earlier operating systems when USN rollback is detected on VDC2, the destination domain controller that is running on a virtual machine. In this illustration, the detection of USN rollback occurs on VDC2 when a replication partner detects that VDC2 has sent an up-to-dateness USN value that was seen previously by the replication partner, which indicates that VDC2's database has rolled back in time improperly.

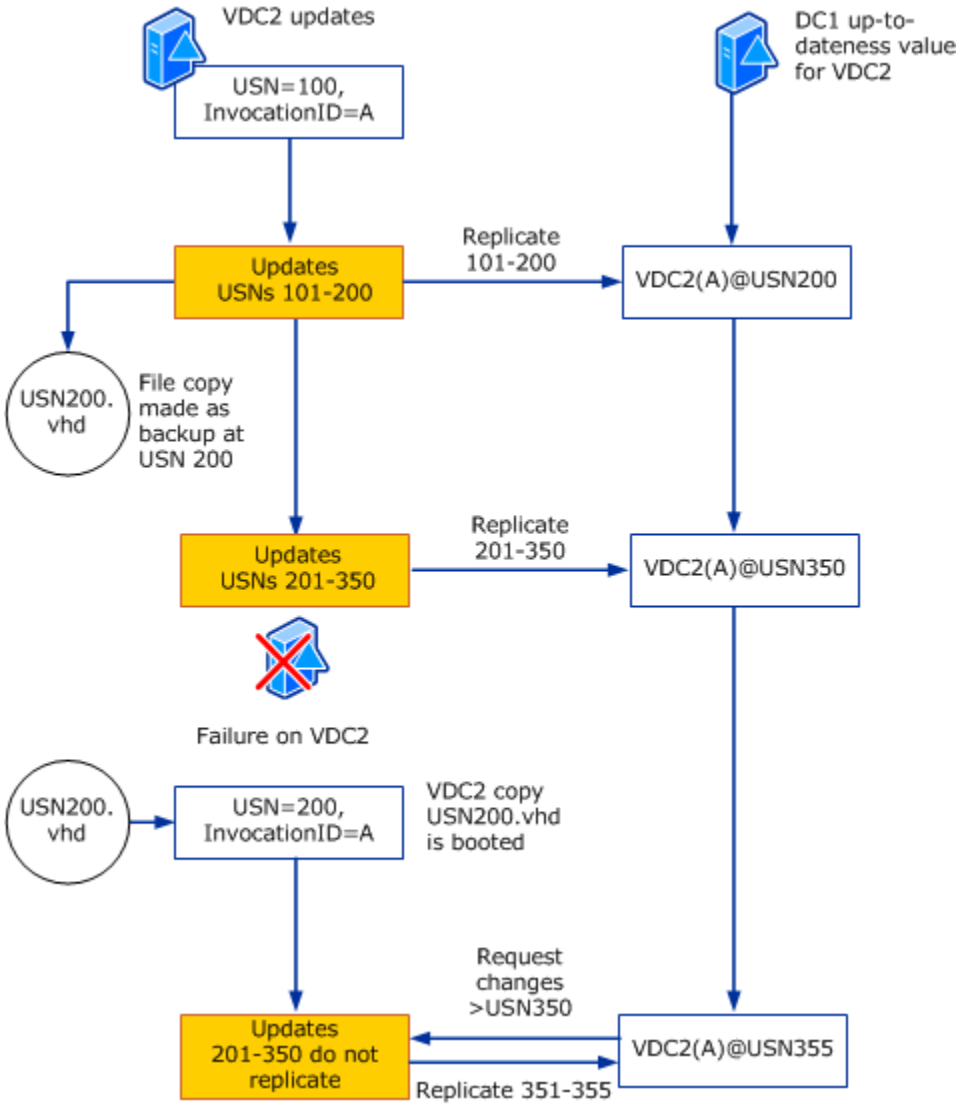

A virtual machine (VM) makes it easy for hypervisor administrators to roll back a domain controller's USNs (its logical clock) by, for example, applying a snapshot outside of the domain controller's awareness. For more information about USN and USN rollback, including another illustration to demonstrate undetected instances of USN rollback, see USN and USN Rollback.

Beginning with Windows Server 2012, AD DS virtual domain controllers hosted on hypervisor platforms that expose an identifier called VM-Generation ID can detect and employ necessary safety measures to protect the AD DS environment if the virtual machine is rolled back in time by the application of a VM snapshot. The VM-GenerationID design uses a hypervisor-vendor independent mechanism to expose this identifier in the address space of the guest virtual machine, so the safe virtualization experience is consistently available of any hypervisor that supports VM-GenerationID. This identifier can be sampled by services and applications running inside the virtual machine to detect if a virtual machine has been rolled back in time.

How do these virtualization safeguards work?

During domain controller installation, AD DS initially stores the VM GenerationID identifier as part of the msDS-GenerationID attribute on the domain controller's computer object in its database (often referred to as the directory information tree, or DIT). The VM GenerationID is independently tracked by a Windows driver inside the virtual machine.

When an administrator restores the virtual machine from a previous snapshot, the current value of the VM GenerationID from the virtual machine driver is compared against a value in the DIT.

If the two values are different, the invocationID is reset and the RID pool discarded thereby preventing USN re-use. If the values are the same, the transaction is committed as normal.

AD DS also compares the current value of the VM GenerationID from the virtual machine against the value in the DIT each time the domain controller is rebooted and, if different, it resets the invocationID, discards the RID pool and updates the DIT with the new value. It also non-authoritatively synchronizes the SYSVOL folder in order to complete safe restoration. This enables the safeguards to extend to the application of snapshots on VMs that were shutdown. These safeguards introduced in Windows Server 2012 enable AD DS administrators to benefit from the unique advantages of deploying and managing domain controllers in a virtualized environment.

The following illustration shows how virtualization safeguards are applied when the same USN rollback is detected on a virtualized domain controller that runs Windows Server 2012 on a hypervisor that supports VM-GenerationID.

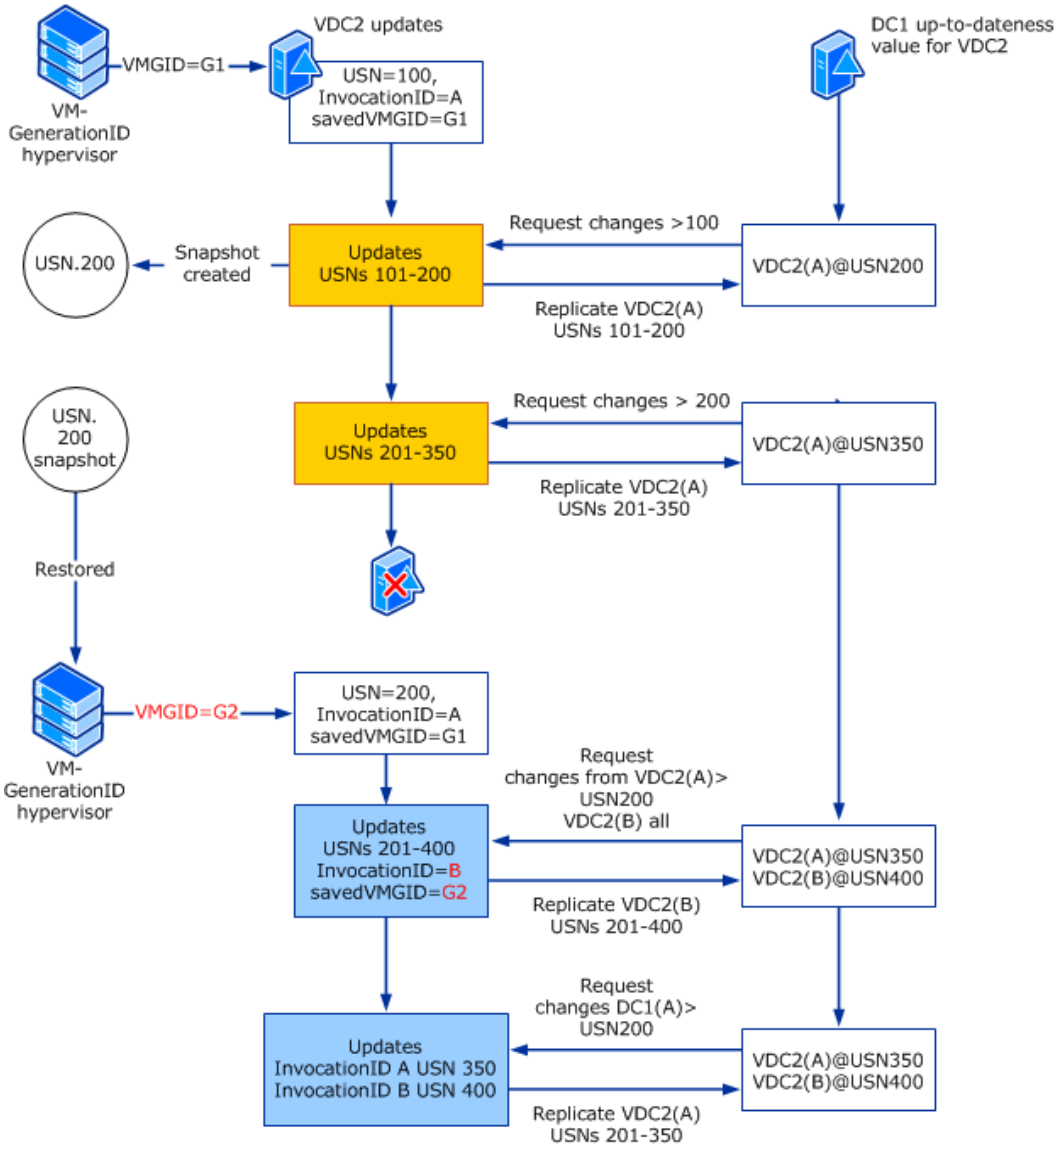

In this case, when the hypervisor detects a change to VM-GenerationID value, virtualization safeguards are triggered, including the reset of the InvocationID for the virtualized DC (from A to B in the preceding example) and updating the VM-GenerationID value saved on the VM to match the new value (G2) stored by the hypervisor. The safeguards ensure that replication converges for both domain controllers.

With Windows Server 2012, AD DS employs safeguards on virtual domain controllers hosted on VM-GenerationID aware hypervisors and ensures that the accidental application of snapshots or other such hypervisor-enabled mechanisms that could ‘rollback’ a virtual machine’s state does not disrupt the AD DS environment (by preventing replication problems such as a USN bubble or lingering objects). However, restoring a domain controller by applying a virtual machine snapshot is not recommended as an alternative mechanism to backing up a domain controller. It is

recommended that you continue to use Windows Server Backup or other VSS-writer based backup solutions.

Caution

If a domain controller in a production environment is accidentally reverted to a snapshot, it's advised that you consult the vendors for the applications, and services hosted on that virtual machine, for guidance on verifying the state of these programs after snapshot restore.

For more information, see Virtualized domain controller safe restore architecture.

Virtualized domain controller cloning

Beginning with Windows Server 2012, administrators can easily and safely deploy replica domain controllers by copying an existing virtual domain controller. In a virtual environment, administrators no longer have to repeatedly deploy a server image prepared by using sysprep.exe, promote the server to a domain controller and then complete additional configuration requirements for deploying each replica domain controller.

Note

Administrators need to follow existing processes to deploy the first domain controller in a domain, such as using a sysprep.exe to prepare a server virtual hard disk (VHD), promote the server to a domain controller and then complete any additional configuration requirements. In a disaster recovery scenario, use the latest server backup to restore the first domain controller in a domain.

Scenarios that benefit from virtual domain controller cloning

- Rapid deployment of additional domain controllers in a new domain
- Quickly restore business continuity during disaster recovery by restoring AD DS capacity via rapid deployment of domain controllers using cloning
- Optimize private cloud deployments by leveraging elastic provisioning of domain controllers to accommodate increased scale requirements
- Rapid provisioning of test environments enabling deployment and testing of new features and capabilities before production rollout
- Quickly meet increased capacity needs in branch offices by cloning existing domain controllers in branch offices

When rapidly deploying a large number of domain controllers, continue to follow your existing procedures for validating the health of each domain controller after installation finishes. Deploy domain controllers in reasonably sized batches so you can validate their health after each batch of installations is complete. The recommended batch size is 10. For more information, see Steps for deploying a clone virtualized domain controller.

Clear separation of responsibilities

The authorization to clone virtualized domain controllers is under the control of the AD DS administrator. In order for hypervisor administrators to deploy additional domain controllers by copying virtual domain controllers, the AD DS administrator has to select and authorize a domain controller and then run preparatory steps to enable it as a source for cloning.

With the virtual machine provisioning typically under the purview of the hypervisor administrator, hypervisor administrators can provision replica domain controller virtual machines by copying virtualized domain controllers that are authorized and prepared for cloning by the AD DS administrator.

Warning
Anyone allowed to administer the hypervisor that hosts a virtual domain controller must be highly trusted and audited in the environment.

How does virtual domain controller cloning work?

The process of cloning involves making a copy of an existing virtual domain controller's VHD (or, for more complex configurations, the domain controller VM), authorizing it for cloning in AD DS and creating a clone configuration file. This reduces the number of steps and time involved in deploying a replica virtual domain controller by eliminating otherwise repetitive deployment tasks.

The clone domain controller uses the following criteria to detect that it is a copy of another domain controller:

1. The value of the VM-Generation ID supplied by the virtual machine is different than the value of the VM-Generation ID stored in the DIT.

   Note
   The hypervisor platform must support VM-Generation ID (Windows Server 2012 Hyper-V supports VM-Generation ID).
2. Presence of a file called DCCloneConfig.xml in one of the following locations:
   - The directory where the DIT resides
   - %windir%\NTDS
   - The root of a removable media drive

Once the criteria are met, it goes through the process of cloning to provision itself as a replica domain controller.

The clone domain controller uses the security context of the source domain controller (the domain controller whose copy it represents) to contact the Windows Server 2012 Primary Domain Controller (PDC) emulator operations master role holder (also known as flexible single master operations, or FSMO). The PDC emulator must be running Windows Server 2012, but it does not have to be running on a hypervisor.

Note
If you have a schema extension with attributes that reference the source domain controller and the attribute is on one of the objects copied (computer object, NTDS

settings object) to create the clone, that attribute will not be copied or updated to reference the clone domain controller.

After verifying that the requesting domain controller is authorized for cloning, the PDC emulator will create a new machine identity including new account, SID, name, and password that identifies this machine as a replica domain controller and send this information back to the clone. The clone domain controller will then prepare the AD DS database files to serve as a replica and it will also clean up the machine state.

For more information, see Virtualized domain controller cloning architecture.

Cloning components

The cloning components include new cmdlets in the Active Directory module for Windows PowerShell and associated XML files:

- **New-ADDCCloneConfigFile** – This cmdlet creates and places DCCloneConfig.xml at the right location to ensure it is available to trigger cloning. It also performs prerequisite checks to ensure successful cloning. It is included in the Active Directory module for Windows PowerShell. You can run it locally on a virtualized domain controller that is being prepared for cloning, or you can run it remotely using the -offline option. You can specify settings for the clone domain controller, such as its name, site, and IP address.

  The prerequisite checks that it performs are:

  Note
  The prerequisite checks are not performed when the –offline option is used. For more information, see Running New-ADDCCloneConfigFile in offline mode.

  - The DC being prepared is authorized for cloning (is a member of the **Cloneable Domain Controllers** group)
  - The PDC emulator runs Windows Server 2012.
  - Any programs or services listed from running **Get-ADDCCloningExcludedApplicationList** are included in CustomDCCloneAllowList.xml (explained in more detail at the end of this list of cloning components).
- **DCCloneConfig.xml** – To successfully clone a virtualized domain controller, this file must be present in the directory where the DIT resides, *%windir%\NTDS*, or the root of a removable media drive. Besides being used as one of the triggers to detect and initiate cloning, it also provides a means to specify configuration settings for the clone domain controller.

  The schema and a sample file for the DCCloneConfig.xml file are stored on all Windows Server 2012 computers at:

  - %windir%\system32\DCCloneConfigSchema.xsd
  - %windir%\system32\SampleDCCloneConfig.xml

  It is recommended that you use the New-ADDCCloneConfigFile cmdlet to create the DCCloneConfig.xml file. Although you could also use the schema file with an XML-aware editor to create this file, manually editing the file increases the likelihood of errors. If you edit the file, it must be done by using XML-aware editors, such as Visual Studio, XML Notepad, or third-party applications (do not use Notepad).

- **Get-ADDCCloningExcludedApplicationList** – This cmdlet is run on the source domain controller before beginning the cloning process to determine which services or installed programs are not on the default supported list, DefaultDCCloneAllowList.xml, or a user-defined inclusion list named CustomDCCloneAllowList.xml file, and thereby have not been evaluated for cloning impact.

  This cmdlet searches the source domain controller for services in the Services Control Manager, and installed programs listed under **HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall** that are not specified in the default list (DefaultDCCloneAllowList.xml) or, if one is provided, the user-defined inclusion list (CustomDCCloneAllowList.xml file). The list of applications and services that is returned by running the cmdlet is the difference between what has already been provided in the DefaultDCCloneAllowList.xml or the CustomDCCloneAllowList.xml file and the list that is constructed at run time, based on what is installed on the source DC. The services and programs output from Get-ADDCCloningExcludedApplicationList can be added to the CustomDCCloneAllowList.xml file if you determine that the services and programs can be safely cloned. To determine if a service or installed program can be safely cloned, evaluate the following conditions:

  - Is the service or installed program affected by the machine identity, such as name, SID, password, and so on?
  - Does the service or installed program store any state locally on the computer that might affect its functionality on the clone?

  You must work with the software vendor of the application to determine if the service or program can be safely cloned.

  Note

  Before provisioning additional services or programs in the CustomDCCloneAllowList.xml file, verify whether you have the necessary license to copy that software contained on that virtual machine.

  If the applications are not cloneable, remove them from the source domain controller before you create the clone media. If an application appears in the cmdlet output, but is not included in the CustomDCCloneAllowList.xml file, cloning will fail. For cloning to succeed, the cmdlet output should not list any services or programs. In other words, an application should either be included in the CustomDCCloneAllowList.xml file or removed from the source domain controller.

  The following table explains the options for running Get-ADDCCloningExcludedApplicationList.

Argument	Explanation
<no argument specified>	Displays a list of services or programs on the console that have not been accounted for cloning. If there is already a CustomDCCloneAllowList.XML in any of the

	permissible locations, it uses that file to displays the remaining services and programs (which may be nothing if the lists match).
-GenerateXml	Creates the CustomDCCloneAllowList.XML file populated with the services and programs listed on the console.
-Force	Overwrites an existing CustomDCCloneAllowList.XML file.
-Path	Folder path to create the CustomDCCloneAllowList.XML.

- **DefaultDCCloneAllowList.xml** – This file is present by default on every Windows Server 2012 domain controller in the *%windir%\system32*. It lists the services and installed programs that can be safely cloned by default. You must not change the location or contents of this file or cloning will fail.
- **CustomDCCloneAllowList.xml** – If you have services or installed programs that reside on your source domain controller that are outside of those listed in the DefaultDCCloneAllowList.xml file, those services and programs must be included in this file. To find the services or installed programs that are not listed in the in the DefaultDCCloneAllowList.xml file, run the **Get-ADDCCloningExcludedApplicationList** cmdlet. You should use the **–GenerateXml** argument to generate the XML file.

  The cloning process checks the following locations in order for this file and uses the first XML file found, regardless of the other folder's contents:

  a. The following registry key:

  ```
  HKey_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters

  AllowListFolder (REG_SZ)
  ```

  b. DSA Working Directory

  c. %systemroot%\NTDS

  d. Removable read/write media, in order of drive letter, at the root of the drive

Deployment scenarios

The following deployment scenarios are supported for virtual domain controller cloning:

- Deploy a clone domain controller by making a copy of a source domain controller's virtual hard disk (vhd) file.
- Deploy a clone domain controller by copying the virtual machine of a source domain controller using the export/import semantics exposed by the hypervisor.

Note

The steps in the section Steps for deploying a clone virtualized domain controller demonstrate copying a virtual machine using the export/import feature of Windows Server 2012 Hyper-V.

Steps for deploying a clone virtualized domain controller

- Prerequisites
- Step 1: Grant the source virtualized domain controller the permission to be cloned
- Step 2: Run Get-ADDCCloningExcludedApplicationList cmdlet
- Step 3: Run New-ADDCCloneConfigFile
- Step 4: Export and then import the virtual machine of the source domain controller

Prerequisites

- To complete the steps in the following procedures, you must be a member of the Domain Admins group or have the equivalent permissions assigned to it.
- The Windows PowerShell commands used in this guide must be run from an elevated command prompt. To do this, right click the **Windows PowerShell** icon, and then click **Run as administrator**.
- A Windows Server 2012 server with the Hyper-V server role installed (**HyperV1**).
- A second Windows Server 2012 server with the Hyper-V server role installed (**HyperV2**).

Notes

- If you are using another hypervisor, you should contact the vendor of that hypervisor to verify if the hypervisor supports VM-Generation ID. If the hypervisor does not support VM-Generation ID and you have provided a DCCloneConfig.xml, the new VM will boot into Directory Services Restore Mode (DSRM).
- To increase the availability of the AD DS service, this guide recommends and provides instructions using two different Hyper-V hosts, which helps prevent a potentially single point of failure. However, you do not need two Hyper-V hosts to perform virtual domain controller cloning.
- You need to be a member of the local Administrators group on each Hyper-V server (**HyperV1** and **HyperV2**).
- In order to successfully import and export a VHD file using Hyper-V, the virtual network switches on both Hyper-V hosts should have the same name. For example, if you have a virtual network switch on **HyperV1** named VNet then there needs to be a virtual network switch on **HyperV2** named VNet.
- If the two Hyper-V hosts (**HyperV1** and **HyperV2**) have different processors, shut down the virtual machine (**VirtualDC1**) that you plan to export, right-click the VM, click **Settings**, click **Processor**, and under **Processor compatibility** select **Migrate to a physical computer with a different processor version** and click **OK**.

- A deployed Windows Server 2012 domain controller (virtualized or physical) that hosts the PDC emulator role (**DC1**). To verify whether the PDC emulator role is hosted on a Windows Server 2012 domain controller, run the following Windows PowerShell command:

  ```
  Get-ADComputer (Get-ADDomainController –Discover –Service
  "PrimaryDC").name –Property operatingsystemversion | fl
  ```

  The OperatingSystemVersion value should return as a version 6.2. If necessary, you can transfer the PDC emulator role to a domain controller that runs Windows Server 2012. For more information, see Using Ntdsutil.exe to transfer or seize FSMO roles to a domain controller.

- A deployed Windows Server 2012 guest virtualized domain controller (**VirtualDC1**) that is in the same domain as the Windows Server 2012 domain controller hosting the PDC emulator role (**DC1**). This will be the source domain controller used for cloning. The guest virtual domain controller will be hosted on a Windows Server 2012 Hyper-V server (**HyperV1**).

  Notes

  - For cloning to succeed, the source domain controller that is used to create the clone cannot be from a DC that has been demoted since the source VHD media was created.
  - Shut down the source domain controller prior to copying the VM or its VHD.
  - You should not clone a VHD or restore a snapshot that is older than the tombstone lifetime value (or the deleted object lifetime value if Active Directory Recycle Bin is enabled). If you are copying a VHD of an existing domain controller, be sure the VHD file is not older that the tombstone lifetime value (by default, 60 days). You should not copy a VHD of a running domain controller to create clone media.

  Eject any virtual floppy drive (VFD) the source DC may have. This can cause a sharing problem when trying to import the new VM.

  Only Windows Server 2012 domain controllers hosted on a VM-GenerationID hypervisor can be used as a source for cloning. The source Windows Server 2012 domain controller used for cloning should be in a healthy state. To determine the state of the source domain controller run dcdiag. To gain a better understanding of the output returned by dcdiag, see What does DCDIAG actually…do?.

  If the source domain controller is a DNS server, the cloned domain controller will also be a DNS server. You should choose a DNS server that hosts only Active Directory-integrated zones.

  DNS client settings are not cloned but are instead specified in the DCCloneConfig.xml file. If they are not specified, the cloned domain controller will point to itself as Preferred DNS server by default. The cloned domain controller will not have a DNS delegation. The administrator of the parent DNS zone should update the DNS delegation for the cloned domain controller as needed.

  Warning

  The virtualization safeguards do not extend to Active Directory Lightweight Directory Services (AD LDS). Therefore you should not attempt to clone an AD DS domain controller that hosts an AD LDS instance by adding this AD LDS instance to the

CustomDCCloneAllowList.xml. Because AD LDS is not VM-Generation ID aware, cloning a domain controller with AD LDS can cause USN rollback-induced divergence on that AD LDS configuration set.

The following server roles are not supported for cloning:

- Dynamic Host Configuration Protocol (DHCP)
- Active Directory Certificate Services (AD CS)
- Active Directory Lightweight Directory Services (AD LDS)

Step 1: Grant the source virtualized domain controller the permission to be cloned

In this procedure, you grant the source domain controller the permission to be cloned by using **Active Directory Administrative Center** to add the source domain controller to the **Cloneable Domain Controllers** group.

To grant the source virtualized domain controller the permission to be cloned

1. On any domain controller in the same domain as the domain controller being prepared for cloning (**VirtualDC1**), open **Active Directory Administrative Center** (ADAC), locate the virtualized domain controller object (domain controllers are usually located under the **Domain Controllers** container in ADAC), right click it, choose **Add to group** and under **Enter the object name to select** type **Cloneable Domain Controllers** and then click **OK**.

   The group membership update performed in this step must replicate to PDC emulator before cloning can be performed. If the **Cloneable Domain Controllers** group is not found, the PDC emulator role might not be hosted on a domain controller that runs Windows Server 2012.

   Note
   To open ADAC on a Windows Server 2012 domain controller, open Windows PowerShell and type **dsac.exe**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet performs the same function as the preceding procedure:

```
Add-ADGroupMember -Identity "CN=Cloneable Domain Controllers,CN=Users,
DC=Fabrikam,DC=Com" -Member "CN=VirtualDC1,OU=Domain Controllers,DC=Fabrikam,DC=com"
```

Step 2: Run Get-ADDCCloningExcludedApplicationList cmdlet

In this procedure, run the `Get-ADDCCloningExcludedApplicationList` cmdlet on the source virtualized domain controller to identify any programs or services that are not evaluated for cloning. You need to run the Get-ADDCCloningExcludedApplicationList cmdlet before the New-

ADDCCloneConfigFile cmdlet because if the New-ADDCCloneConfigFile cmdlet detects an excluded application, it will not create a DCCloneConfig.xml file.

To identify applications or services that run on a source domain controller which have not been evaluated for cloning

1. On the source domain controller (**VirtualDC1**), click **Server Manager**, click **Tools**, click **Active Directory Module for Windows PowerShell** and then type the following command:

   ```
   Get-ADDCCloningExcludedApplicationList
   ```

2. Vet the list of the returned services and installed programs with the software vendor to determine whether they can be safely cloned. If applications or services in the list cannot be safely cloned, you must remove them from the source domain controller or cloning will fail.
3. For the set of services and installed programs that were determined to be safely cloned, run the command again with the **–GenerateXML** switch to provision these services and programs in the **CustomDCCloneAllowList.xml** file.

   ```
   Get-ADDCCloningExcludedApplicationList -GenerateXml
   ```

Step 3: Run New-ADDCCloneConfigFile

Run New-ADDCCloneConfigFile on the source domain controller, and optionally specify configuration settings for the clone domain controller, such as the name, the IP address, and DNS resolver.

For example, to create a clone domain controller named VirtualDC2 with a static IPv4 address, type:

```
New-ADDCCloneConfigFile -Static -IPv4Address "10.0.0.2" -IPv4DNSResolver "10.0.0.1" -
IPv4SubnetMask "255.255.255.0" -CloneComputerName "VirtualDC2" -IPv4DefaultGateway
"10.0.0.3" -SiteName "REDMOND"
```

Note

The clone domain controller will be located in the same site as the source domain controller unless a different site is specified in the DCCloneConfig.xml file. It is recommended that you specify a suitable site in the DCCloneConfig.xml file for the clone domain controller based on its IP address.

The computer name is optional. If you do not specify one, a unique name will be generated based on the following algorithm:

- The prefix is the first 8 characters of the source domain controller computer name. For example, a source computer name of SourceComputer is truncated to a prefix string of SourceCo.
- A unique naming suffix of the format "–CL*nnnn*" is appended to the prefix string where *nnnn* is the next available value from 0001-9999 that the PDC determines is not currently in use.

For example, if 0047 is the next available number in the allowed range, using the preceding example of the computer name prefix SourceCo, the derived name to use for the clone computer will be set as SourceCo-CL0047.

Note

A global catalog server (GC) is required for the New-ADDCCloneConfigFile cmdlet to work successfully. The source domain controller's membership in the **Cloneable Domain Controllers** group must be reflected on the GC. The GC does not need to be the same domain controller as the PDC emulator, but preferably it should be in the same site. If a GC is not available, the command fails with the error "The server is not operational." For more information, see Virtualized Domain Controller Troubleshooting.

To create a clone domain controller named Clone1 with static IPv4 settings and specify preferred and alternate WINS servers, type:

```
New-ADDCCloneConfigFile -CloneComputerName "Clone1" -Static -IPv4Address "10.0.0.5" -
IPv4DNSResolver "10.0.0.1" -IPv4SubnetMask "255.255.0.0" -PreferredWinsServer "10.0.0.1"
-AlternateWinsServer "10.0.0.2"
```

Note

If you specify WINS servers, you must specify both **–PreferredWINSServer** and **–AlternateWINSServer**. If you specify only of those arguments, cloning fails with error code 0x80041005 appearing in the dcpromo.log.

To create a clone domain controller named Clone2 with dynamic IPv4 settings, type:

```
New-ADDCCloneConfigFile -CloneComputerName "Clone2" -IPv4DNSResolver "10.0.0.1"
```

Note

In this case, there should be a DHCP server in the environment that the clone can reach and obtain IP address and other relevant network settings.

To create a clone domain controller named Clone2 with dynamic IPv4 settings and specify preferred and alternate WINS servers, type:

```
New-ADDCCloneConfigFile -CloneComputerName "Clone2" -IPv4DNSResolver "10.0.0.1" -SiteName
"REDMOND" -PreferredWinsServer "10.0.0.1" -AlternateWinsServer "10.0.0.2"
```

To create a clone domain controller with dynamic IPv6 settings, type:

```
New-ADDCCloneConfigFile -IPv6DNSResolver "2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc"
```

To create a clone domain controller with static IPv6 settings, type:

```
New-ADDCCloneConfigFile -Static -IPv6DNSResolver "2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc"
```

Note

When specifying IPv6 settings, the only difference between the static and dynamic settings is the inclusion of **–Static** switch. The inclusion of the **–Static** switch makes it

mandatory to specify at least one **IPv6DNSResolver**.The static IPv6 address is expected to be configured via stateless address auto configuration (SLAAC) with router assigned prefixes. With dynamic IPv6, the DNS resolvers are optional, but it's expected that the clone can reach an IPv6-enabled DHCP server on the subnet to obtain IPv6 address and DNS configuration information.

Running New-ADDCCloneConfigFile in offline mode

If you have multiple copies of source domain controller media that have been prepared for cloning (meaning the source domain controller is authorized for cloning, the Get-ADDCCloningExcludedApplicationList cmdlet has been run, and so on) and you want to specify different settings for each copy of the media, you can run New-ADDCCloneConfigFile in offline mode. This can be more efficient than individually preparing each VM, for example, by importing each copy.

In this case, domain administrators can mount the offline disk and use Remote Server Administration Tools (RSAT) to run the New-ADDCCloneConfigFile cmdlet with the –offline argument in order to add the XML files, which allows for factory-like automation using new Windows PowerShell options included in Windows Server 2012. For more information about how to mount the offline disk in order to run the New-ADDCCloneConfigFile cmdlet in offline mode, see Adding XML to the Offline System Disk.

You should first run the cmdlet locally on the source media to ensure that prerequisite checks pass. The prerequisite checks are not performed in offline mode because the cmdlet could be run from a machine that may not be from the same domain or from a domain-joined computer. After you run the cmdlet locally, it will create a DCCloneConfig.xml file. You may delete the DCCloneConfig.xml that is created locally if you plan to use the offline mode subsequently.

To create a clone domain controller named CloneDC1 in offline mode, in a site called "REDMOND" with static IPv4 address, type:

```
New-ADDCCloneConfigFile –Offline –CloneComputerName CloneDC1 –SiteName REDMOND -
IPv4Address "10.0.0.2" -IPv4DNSResolver "10.0.0.1" -IPv4SubnetMask "255.255.0.0" -
IPv4DefaultGateway "10.0.0.1" –Static –Path F:\Windows\NTDS
```

To create a clone domain controller named Clone2 in offline mode with static IPv4 and static IPv6 settings, type:

```
New-ADDCCloneConfigFile –Offline -IPv4Address "10.0.0.2" -IPv4DNSResolver "10.0.0.1" -
IPv4SubnetMask "255.255.0.0" –Static -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -CloneComputerName "Clone2" -PreferredWINSServer
"10.0.0.1" -AlternateWINSServer "10.0.0.3" –Path F:\Windows\NTDS
```

To create a clone domain controller in offline mode with static IPv4 and dynamic IPv6 settings and specify multiple DNS servers for the DNS resolver settings, type:

```
New-ADDCCloneConfigFile -Offline -IPv4Address "10.0.0.10" -IPv4SubnetMask "255.255.0.0" -
IPv4DefaultGateway "10.0.0.1" -IPv4DNSResolver @( "10.0.0.1","10.0.0.2" ) -Static -
IPv6DNSResolver "2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -Path F:\Windows\NTDS
```

To create a clone domain controller named Clone1 in offline mode with dynamic IPv4 and static IPv6 settings, type:

```
New-ADDCCloneConfigFile -Offline -Static -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -CloneComputerName "Clone1" -PreferredWINSServer
"10.0.0.1" -AlternateWINSServer "10.0.0.3" -SiteName "REDMOND" -Path F:\Windows\NTDS
```

To create a clone domain controller in offline mode with dynamic IPv4 and dynamic IPv6 settings, type:

```
New-ADDCCloneConfigFile -Offline -IPv4DNSResolver "10.0.0.1" -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -Path F:\Windows\NTDS
```

Step 4: Export and then import the virtual machine of the source domain controller

In this procedure, export the virtual machine of the source virtualized domain controller and then import the virtual machine. This action creates a clone virtualized domain controller in your domain.

You need to be a member of the local Administrators group on each Hyper-V host. If you use different credentials for each server, run the Windows PowerShell cmdlets to export and import the VM in different Windows PowerShell sessions.

If there are snapshots on the source domain controller, they should be deleted before the source domain controller is exported because the VM will not import if a snapshot has processor settings that are incompatible with the target hyper-v host. If the processor settings are compatible between the source and target hyper-v hosts, you may export and copy the source without deleting snapshots beforehand. After import, however, the snapshots must be deleted from the clone VM before it starts.

To copy a virtual domain controller by exporting and then importing the virtualized source domain controller

1. On **HyperV1**, shutdown the source domain controller (**VirtualDC1**).

   Windows PowerShell equivalent commands

   ```
   Stop-VM –Name VirtualDC1 –ComputerName HyperV1
   ```

2. On **HyperV1**, delete snapshots and then export the source domain controller (VirtualDC1) to the c:\CloneDCs directory.

   Note

   You should delete all the associated snapshots because each time a snapshot is

taken, a new AVHD file is created that acts as differencing disk. This creates a chain affect. If you have taken snapshots and insert the DCCLoneConfig.xml file into the VHD, you may end up creating a clone from an older DIT version or inserting the configuration file into the wrong VHD file. Deleting the snapshot merges all these AVHDs into the base VHD.

Windows PowerShell equivalent commands

```
Get-VMSnapshot VirtualDC1 | Remove-VMSnapshot –
IncludeAllChildSnapshots

Export-VM –Name VirtualDC1 –ComputerName HyperV1 -Path
c:\CloneDCs\VirtualDC1
```

3. Copy the folder **virtualdc1** to the c:\Import directory of **HyperV2**.
4. On **HyperV2**, using **Hyper-V Manager**, import the virtual machine (using the **Import Virtual Machine** wizard in **Hyper-V Manager**) from the folder **c:\Import\virtualdc1** and delete all associated **Snapshots**.

   Use the **Copy the virtual machine (create new unique ID)** option when importing the virtual machine.

Windows PowerShell equivalent commands

```
$path = Get-ChildItem
"C:\CloneDCs\VirtualDC1\VirtualDC1\Virtual Machines"

$vm = Import-VM –Path $path.fullname –Copy -GenerateNewId

Rename-VM $vm VirtualDC2
```

To create multiple clone domain controllers from the same source domain controller:

- UI: in the **Import Virtual Machine** wizard, specify new locations for **Virtual machine configuration folder**, **Snapshot store**, **Smart Paging folder**, and a different **Location** for the virtual hard disks for the virtual machine.
- Windows PowerShell: specify new locations for the virtual machine by using the following parameters for the `Import-VM` cmdlet:

```
$path = Get-ChildItem
"C:\CloneDCs\VirtualDC1\VirtualDC1\Virtual Machines"

Import-VM –Path $path.fullname –Copy –GenerateNewId –
ComputerName HyperV2 –VhdDestinationPath "path" –
SnapshotFilePath "path" –SmartPagingFilePath "path" –
VirtualMachinePath "path"
```

Note

The recommended batch size for creating multiple clone domain controllers simultaneously is 10. The maximum number is restricted by the maximum number of outbound replication connections, which by default is 16 for Distributed File System Replication (DFSR) and 10 for File Replication Service

(FRS). You should not deploy more than the recommended number of clone domain controllers simultaneously unless you have thoroughly tested that number for your environment.

5. On **HyperV1**, restart the source domain controller (**(VirtualDC1**) to bring it back online.

Windows PowerShell equivalent commands

```
Start-VM –Name VirtualDC1 –ComputerName HyperV1
```

6. On **HyperV2**, start the virtual machine (**VirtualDC2**) to bring it online as a clone domain controller in the domain.

Windows PowerShell equivalent commands

```
Start-VM –Name VirtualDC2 –ComputerName HyperV2
```

Note

The PDC emulator must be running for cloning to succeed. If it was shutdown, make sure it has started and performed initial synchronization so it is aware that is holds the PDC emulator role. For more information, see Microsoft KB article 305476.

After cloning completes, verify the name of the clone computer to ensure the cloning operation succeeded. Verify that the VM did not start in Directory Services Restore Mode (DSRM). If you try to log on and receive an error indicating no logon servers are available, try logging on in DSRM. If the DC did not clone successfully and it is booted in DSRM, check the logs in Event Viewer and dcpromo logs in the %systemroot%/debug folder.

The cloned domain controller will be a member of the **Cloneable Domain Controllers** group because it copies the membership from the source domain controller. As a best practice, you should leave the **Cloneable Domain Controllers** group empty until you are ready to perform cloning operations, and you should remove members after cloning operations are complete.

If the source domain controller stores a backup media, the cloned domain controller will also store the backup media. You can run `wbadmin get versions` to show the backup media on the cloned domain controller. A member of the Domain Admins group should delete the backup media on the cloned domain controller to prevent it from being accidentally restored. For more information about how to delete a system state backup using wbadmin.exe, see Wbadmin delete systemstatebackup.

Troubleshooting

If the clone domain controller (**VirtualDC2**) starts in Directory Services Restore Mode (DSRM), it does not return to a normal mode on its own on the next reboot. To log on to a domain controller that is started in DSRM, use **.\Administrator** and specify the DSRM password.

Correct the cause for cloning failure and verify that the dcpromo.log does not indicate that cloning cannot be re-tried. If cloning cannot be re-tried, safely discard the media. If cloning can be re-tried, you must remove the DS Restore Mode boot flag in order to try cloning again.

1. Open Windows Server 2012 with an elevated command (right click Windows Server 2012 and choose Run as Administrator), and then type **msconfig**.
2. On the **Boot** tab, under **Boot Options**, clear **Safe boot** (it is already selected with the option **Active Directory repair enabled**).
3. Click **OK** and restart when prompted.

For more troubleshooting information about virtualized domain controllers, see Virtualized Domain Controller Troubleshooting.

Virtualized Domain Controller Technical Reference (Level 300)

The virtualized domain controller (VDC) technical reference consists of the following topics:

- Virtualized Domain Controller Architecture
- Virtualized Domain Controller Deployment and Configuration
- Virtualized Domain Controller Troubleshooting
- Virtualized Domain Controller Technical Reference Appendix
- Virtualized Domain Controller Additional Resources

Virtualized Domain Controller Architecture

This topic covers the architecture of virtualized domain controller cloning and safe restore. It shows the processes cloning and safe restore with flowcharts and then provides a detailed explanation of each step in the process.

- Virtualized domain controller cloning architecture
- Virtualized domain controller safe restore architecture

Virtualized domain controller cloning architecture

Overview

Virtualized domain controller cloning relies on the hypervisor platform to expose an identifier called **VM-Generation ID** to detect creation of a virtual machine. AD DS initially stores the value of this identifier in its database (NTDS.DIT) during domain controller promotion. When the virtual machine boots up, the current value of the VM-Generation ID from the virtual machine is compared against the value in the database. If the two values are different, the domain controller resets the Invocation ID and discards the RID pool, thereby preventing USN re-use or the potential creation of duplicate security-principals. The domain controller then looks for a DCCloneConfig.xml file in the locations called out in Step 3 in Cloning Detailed Processing. If it

finds a DCCloneConfig.xml file, it concludes that it is being deployed as a clone, so it initiates cloning to provision itself as an additional domain controller by re-promoting using the existing NTDS.DIT and SYSVOL contents copied from source media.

In a mixed environment where some hypervisors support VM-GenerationID and others do not, it is possible for a clone media to be accidentally deployed on a hypervisor that does not support VM-GenerationID. The presence of DCCloneConfig.xml file indicates administrative intent to clone a DC. Therefore, if a DCCloneConfig.xml file is found during boot but a VM-GenerationID is not provided from the host, the clone DC is booted into Directory Services Restore Mode (DSRM) to prevent any impact to the rest of the environment. The clone media can be subsequently moved to a hypervisor that supports VM-GenerationID and then cloning can be retried.

If the clone media is deployed on a hypervisor that supports VM-GenerationID but a DCCloneConfig.xml file is not provided, as the DC detects a VM-GenerationID change between its DIT and the one from the new VM, it will trigger safeguards to prevent USN re-use and avoid duplicate SIDs. However, cloning will not be initiated, so the secondary DC will continue to run under the same identity as the source DC. This secondary DC should be removed from the network at the earliest possible time to avoid any inconsistencies in the environment. For more information about how to reclaim this secondary DC while ensuring that updates get replicated outbound, see Microsoft KB article 2742970.

Cloning Detailed Processing

The following diagram shows the architecture for an initial cloning operation and for a cloning retry operation. These processes are explained in more detail later in this topic.

Initial Cloning Operation

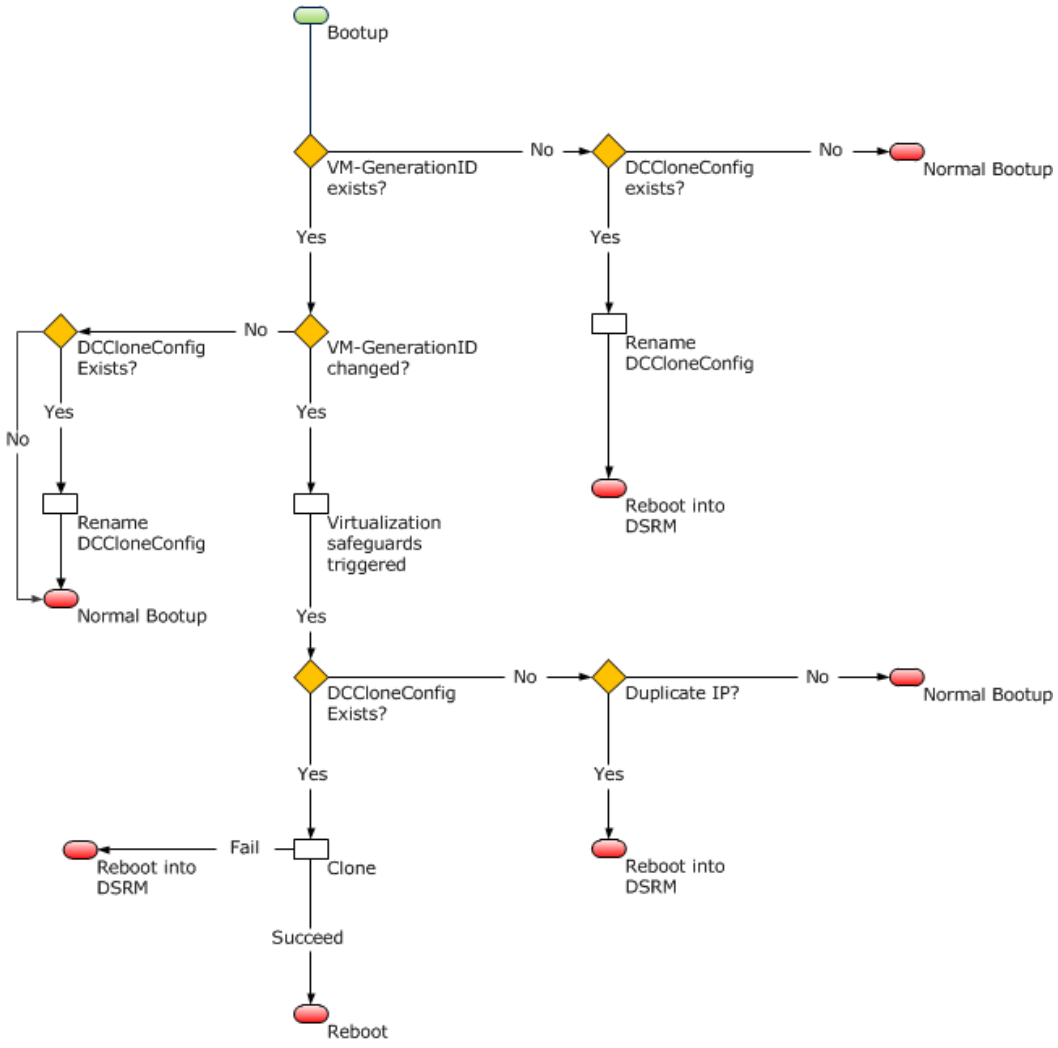

Cloning retry operation

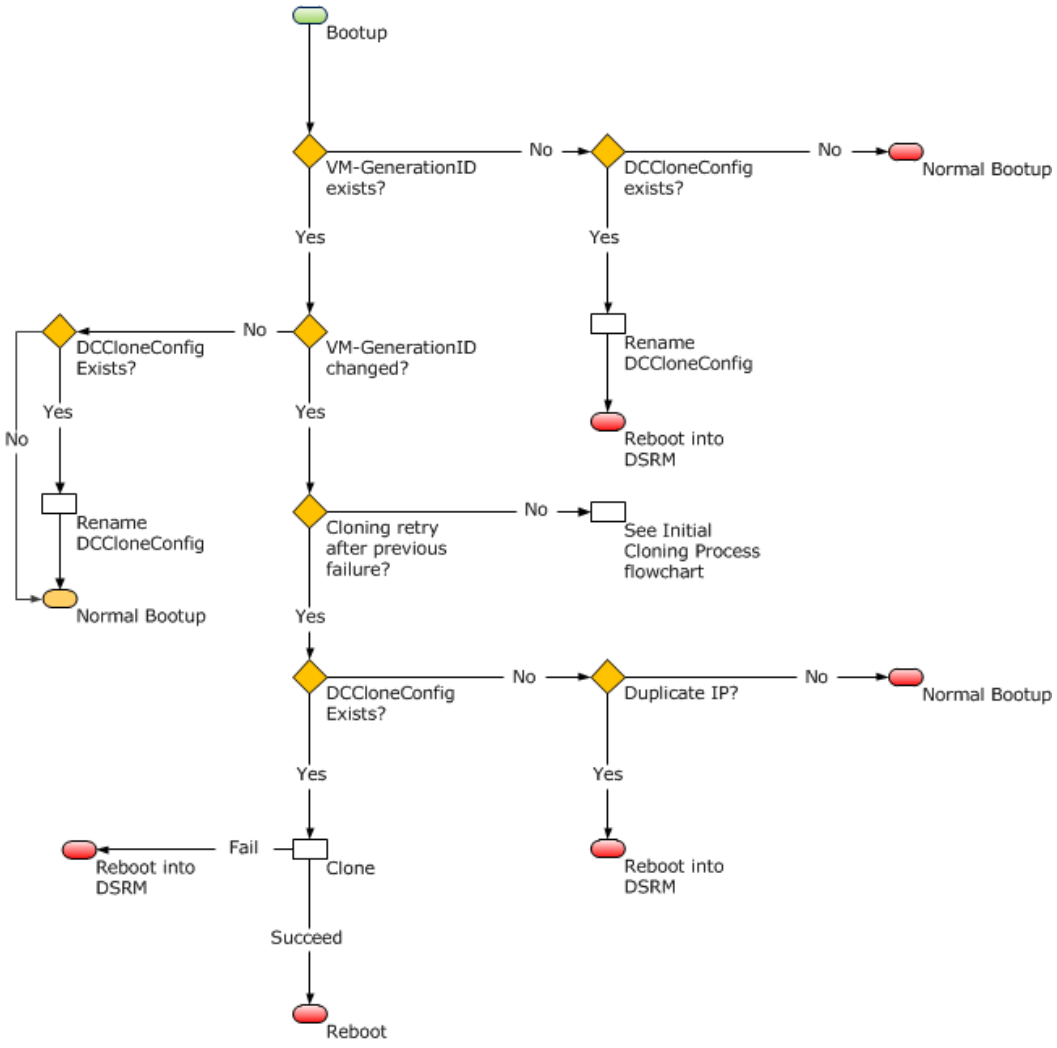

The following steps explain the process in more detail:

1. An existing virtual machine domain controller boots up in a hypervisor that supports VM-Generation ID.
    a. This VM has no existing VM Generation-ID value set on its AD DS computer object after promotion.
    b. Even if it is null, the next computer creation will mean it still clones, as a new VM Generation-ID will not match.
    c. The VM Generation-ID is set after the next reboot of the DC, and does not replicate.
2. The virtual machine then reads the VM-Generation ID provided by the VMGenerationCounter driver. It compares the two VM-Generation IDs.
    a. If the IDs match, this is not a new virtual machine and cloning will not proceed. If a DCCloneConfig.xml file exists, the domain controller renames the file with a time-date stamp to prevent cloning. The server continues booting normally. This is how every reboot of any virtual domain controller operates in Windows Server 2012.

b. If the two IDs do not match, this is a new virtual machine that contains an NTDS.DIT from a previous domain controller (or it is a restored snapshot). If a DCCloneConfig.xml file exists, the domain controller proceeds with cloning operations. If not, it continues with snapshot restoration operations. See Virtualized domain controller safe restore architecture.

c. If the hypervisor does not provide a VM-Generation ID for comparison but there is a DCCloneConfig.xml file, the guest renames the file and then boots into DSRM to protect the network from a duplicate domain controller. If there is no dccloneconfig.xml file, the guest boots normally (with the potential for a duplicate domain controller on the network). For more information about how to reclaim this duplicate domain controller, see Microsoft KB article 2742970.

3. The NTDS service checks the value of the VDCisCloning DWORD registry value name (under HKEY_Local_Machine\System\CurrentControlSet\Services\Ntds\Parameters).
   a. If it does not exist, this is a first attempt at cloning for this virtual machine. The guest implements the VDC object duplication safeguards of invalidating the local RID pool and setting a new replication invocation ID for the domain controller
   b. If it is already set to 0x1, this is a "retry" cloning attempt, where a previous cloning operation failed. The VDC object duplication safety measures are not taken as they had to have already run once before and would unnecessarily alter the guest multiple times.
4. The IsClone DWORD registry value name is written (under Hkey_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters)
5. The NTDS service changes the guest boot flag to start in DS Repair Mode for any further reboots.
6. The NTDS service attempts to read the DcCloneConfig.xml in one of the three accepted locations (DSA Working Directory, %windir%\NTDS, or removable read/write media, in order of drive letter, at the root of the drive).
   a. If the file does not exist in any valid location, the guest checks the IP address for duplication. If the IP address is not duplicated, the server boots up normally. If there is a duplicate IP address, the computer boots into DSRM to protect the network from a duplicate domain controller.
   b. If the file does exist in a valid location, the NTDS service validates its settings. If the file is blank (or any particular settings are blank) then NTDS configures automatic values for those settings.
   c. If the DcCloneConfig.xml exists but contains any invalid entries or is unreadable, cloning fails and the guest boots into Directory Services Restore Mode (DSRM).
7. The guest disables all DNS auto-registration to prevent accidental hijacking of the source computer name and IP addresses.
8. The guest stops the Netlogon service to prevent any advertising or answering of network AD DS requests from clients.
9. NTDS validates that there are no services or programs installed that are not part of the DefaultDCCloneAllowList.xml or CustomDCCloneAllowList.xml

a. If there are services or programs installed that are not in the default exclusion allow list or the custom exclusion allow list, cloning fails and the guest boots into DSRM to protect the network from a duplicate domain controller.
b. If there are no incompatibilities, cloning continues.

10. If automatic IP addressing will be used due to blank DCCloneConfig.xml network settings, the guest enables DHCP on the network adapters to gain an IP address lease, network routing, and name resolution information.
11. The guest locates and contacts the domain controller running the PDC emulator FSMO role. This uses DNS and the DCLocator protocol. It makes an RPC connection and calls the method IDL_DRSAddCloneDC to clone the domain controller computer object.
    a. If the guest's source computer object holds domain head extended permission of "'Allow a DC to create a clone of itself" then cloning proceeds.
    b. If the guest's source computer object does not hold that extended permission, cloning fails and the guest boots into DSRM to protect the network from a duplicate domain controller.
12. The AD DS computer object name is set to match the name specified in the DCCloneConfig.xml, if any, or else automatically generated on the PDCE. NTDS creates the correct NTDS setting object for the appropriate Active Directory logical site.
    a. If this is a PDC cloning, then the guest renames the local computer and reboots. After reboot, it goes through step 1 – 10 again, then goes to step 13.
    b. If this is a replica DC cloning, there is no reboot at this stage.
13. The guest provides the promotion settings to the DS Role Server service, which commences promotion.
14. The DS Role Server service stops all of the AD DS-related services (NTDS, NTFRS/DFSR, KDC, DNS).
15. The guest forces NT5DS (Windows NTP) time synchronization with another domain controller (in a default Windows Time Service hierarchy, this means using the PDCE). The guest contacts the PDCE. All existing Kerberos tickets flush.
16. The guest configures the DFSR or NTFRS services to run automatically. The guest deletes all existing DFSR and NTFRS database files (default: c:\windows\ntfrs and c:\system volume information\dfsr\<*database_GUID*>), in order to force non-authoritative synchronization of SYSVOL when the service is next started. The guest does not delete the file contents of SYSVOL, to pre-seed the SYSVOL when the synchronization starts later.
17. The guest is renamed. The DS Role Server service on the guest begins AD DS configuration (promotion), using the existing NTDS.DIT database file as a source, rather than the template database included in c:\windows\system32 like a promotion normally does.
18. The guest contacts the RID Master FSMO role holder to get a new RID pool allocation.
19. The promotion process creates a new invocation ID and recreates the NTDS Settings object for the cloned domain controller (irrespective of cloning, this is part of domain promotion when using an existing NTDS.DIT database).
20. NTDS replicates in objects that are missing, newer, or have a higher version from a partner domain controller. The NTDS.DIT already contains objects from the time the source domain

controller went offline, and those are used as possible in order to minimize replication traffic inbound. The global catalog partitions are populated.

21. The DFSR or FRS service starts and because there is no database, SYSVOL non-authoritatively synchronizes inbound from a replication partner. This process re-uses pre-existing data in the SYSVOL folder, in order to minimize network replication traffic.
22. The guest re-enables DNS client registration now that the computer is uniquely named and networked.
23. The guest runs the SYSPREP modules specified by the DefaultDCCloneAllowList.xml <SysprepInformation> element in order to scrub out references to the previous computer name and SID.
24. Cloning promotion is complete.
    a. The guest removes the DSRM boot flag so the next reboot will be normal.
    b. The guest renames the DCCloneConfig.xml with an appended date-time stamp, so that it is not read again at next boot up.
    c. The guest removes the VdcIsCloning DWORD registry value name under HKEY_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters.
    d. The guest sets the "VdcCloningDone" DWORD registry value name under HKEY_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters to 0x1. Windows does not use this value, but instead provides it as a marker for third parties.
25. The guest updates the msDS-GenerationID attribute on its own cloned domain controller object to match the current guest VM-Generation ID.
26. The guest restarts. It is now a normal, advertising domain controller.

Virtualized domain controller safe restore architecture

Overview

AD DS relies on the hypervisor platform to expose an identifier called **VM-Generation ID** to detect the snapshot restore of a virtual machine. AD DS initially stores the value of this identifier in its database (NTDS.DIT) during domain controller promotion. When an administrator restores the virtual machine from a previous snapshot, the current value of the VM-Generation ID from the virtual machine is compared against the value in the database. If the two values are different, the domain controller resets the Invocation ID and discards the RID pool, thereby preventing USN re-use or the potential creation of duplicate security-principals. There are two scenarios where safe restore can occur:

- When a virtual domain controller is started after a snapshot has been restored while it was shut down
- When a snapshot is restored on a running virtual domain controller

  If the virtualized domain controller in the snapshot is in a suspended state rather than shutdown, then you need to restart the AD DS service to trigger a new RID pool request. You

can restart the AD DS service by using the Services snap-in or using Windows PowerShell (Restart-Service NTDS -force).

The following sections explain safe restore in detail for each scenario.

Safe Restore Detailed Processing

The following flowchart shows how safe restore occurs when a virtual domain controller is started after a snapshot has been restored while it was shut down.

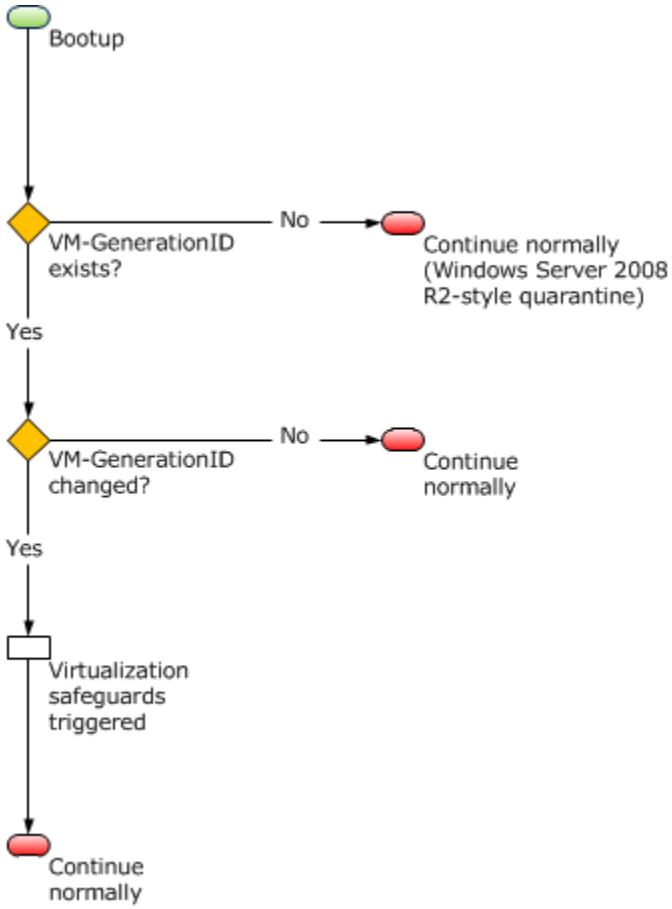

1. When the virtual machine boots up after a snapshot restore, it will have new VM-Generation ID provided by the hypervisor host because of the snapshot restore.
2. The new VM-Generation ID from the virtual machine is compared to the VM-Generation ID in the database. Because the two IDs do not match, it employs virtualization safeguards (see step 3 in the previous section). After the restore finishes applying, the VM-GenerationID set on its AD DS computer object is updated to match the new ID provide by the hypervisor host.
3. The guest employs virtualization safeguards by:
   a. Invalidating the local RID pool.
   b. Setting a new invocation ID for the domain controller database.

Note

This part of the safe restore overlaps with the cloning process. Although this process is about safe restore of a virtual domain controller after it boots up following a snapshot restore, the same steps happen during the cloning process.

The following diagram shows how virtualization safeguards prevent divergence induced by USN rollback when a snapshot is restored on a running virtual domain controller.

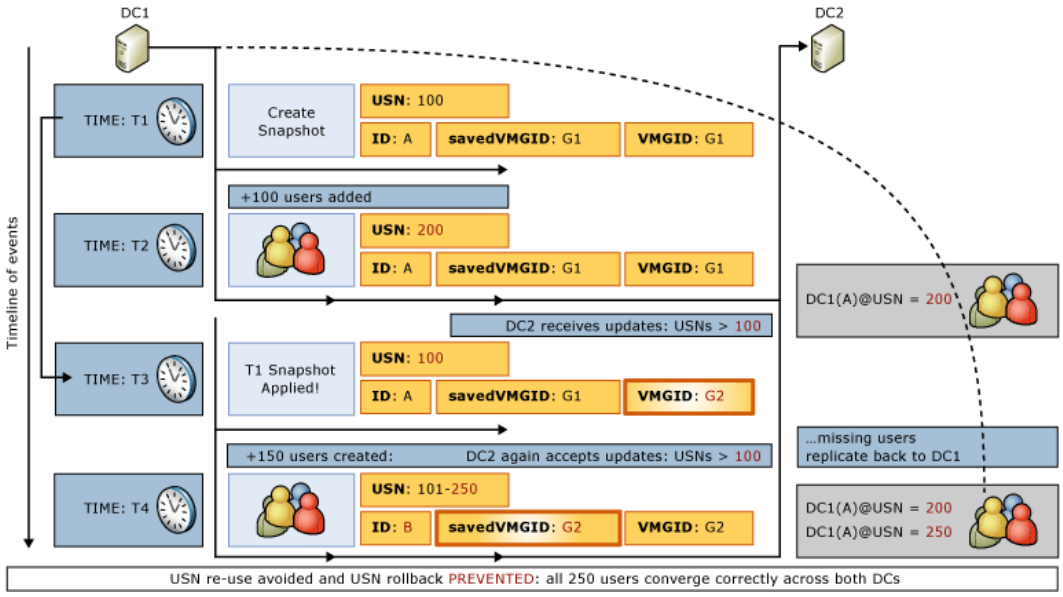

Note

The preceding illustration is simplified to explain the concepts.

1. At time T1, the hypervisor administrator takes a snapshot of virtual DC1. DC1 at this time has a USN value (**highestCommittedUsn** in practice) of 100, InvocationId (represented as ID in the preceding diagram) value of A (in practice this would be GUID). The savedVMGID value is the VM-GenerationID in the DIT file of the DC (stored against the computer object of the DC in an attribute named **msDS-GenerationId**). The VMGID is the current value of the VM-GenerationId available from the virtual machine driver. This value is supplied by the hypervisor.
2. At a later time T2, 100 users are added to this DC (consider users as an example of updates that could have been performed on this DC between time T1 and T2; these updates could actually be a mix of user creations, group creations, password updates, attribute updates, and so on). In this example, each update consumes one unique USN (though in practice a user creation may consume more than one USN). Before committing these updates, DC1 checks if the value of VM-GenerationID in its database (savedVMGID) is the same as the current value available from the driver (VMGID). They are same, as no rollback has happened yet, so the updates are committed and USN moves up to 200, indicating that the next update can use USN 201. There is no change in InvocationId, savedVMGID, or VMGID. These updates replicate out to DC2 at the next replication cycle. DC2 updates it high watermark (and **UptoDatenessVector**) represented here simply as DC1(A) @USN = 200.

That is, DC2 is aware of all updates from DC1 in the context of InvocationId A through USN 200.

3. At time T3, the snapshot taken at time T1 is applied to DC1. DC1 has been rolled back, so its USN rolls back to 100, indicating it could use USNs from 101 to associate with subsequent updates. However, at this point, the value of VMGID would be different on hypervisors that support VM-GenerationID.
4. Subsequently, when DC1 performs any update, it checks whether the value of VM-GenerationId that it has in its database (savedVMGID) is the same as the value from the virtual machine driver (VMGID). In this case, it is not the same, so DC1 infers this as indicative of a rollback, and it triggers virtualization safeguards; in other words, it resets its InvocationId (ID = B) and discards the RID pool (not shown in the preceding diagram). It then saves the new value of VMGID in its database and commits those updates (USN 101 – 250) in the context of the new InvocationId B. At the next replication cycle, DC2 knows nothing from DC1 in the context of InvocationId B, so it requests everything from DC1 associated with InvocationID B. As a result, the updates performed on DC1 subsequent to the application of snapshot will safely converge. In addition, the set of updates that were performed on DC1 at T2 (which were lost on DC1 after the restore of the snapshot) would replicate back into DC1 at the next scheduled replication because they had replicated out to DC2 (as indicated by the dotted line back to DC1).

After the guest employs virtualization safeguards, NTDS replicates Active Directory object differences inbound non-authoritatively from a partner domain controller. The up-to-dateness vector of the destination directory service is updated accordingly. Then the guest synchronizes SYSVOL:

- If using FRS, the guest stops the NTFRS service and sets D2 BURFLAGS registry value. It then starts the NTFRS service, which non-authoritatively replicates inbound, re-using existing unchanged SYSVOL data when possible.
- If using DFSR, the guest stops the DFSR service and deletes the DFSR database files (default location: %systemroot%\system volume information\dfsr\<*database GUID*>). It then starts the DFSR service, which non-authoritatively replicates inbound, re-using existing unchanged SYSVOL data when possible.

Note

- If the hypervisor does not provide a VM-Generation ID for comparison, the hypervisor does not support virtualization safeguards and the guest will operate like a virtualized domain controller that runs Windows Server 2008 R2 or earlier. The guest implements USN rollback quarantine protection if there is an attempt to start replicating with USNs that have not advanced past the last highest USN seen by the partner DC. For more information about USN rollback quarantine protection, see USN and USN Rollback

Virtualized Domain Controller Deployment and Configuration

This topic covers:

- Installation Considerations

  This includes platform requirements and other important constraints.
- Virtualized Domain Controller Cloning

  This explains in detail the entire virtualized domain controller cloning process.
- Virtualization safeguards

  This explains in detail the validations that are made during virtualized domain controller safe restore.

Installation Considerations

There is no special role or feature installation for virtualized domain controllers; all domain controllers automatically contain cloning and safe restore capabilities. You cannot remove or disable these capabilities.

Use of Windows Server 2012 domain controllers requires a Windows Server 2012 AD DS Schema version 56 or higher and forest functional level equal to Windows Server 2003 Native or higher.

Both writable and read-only domain controllers support all aspects of virtualized DC, as do Global Catalogs and FSMO roles.

Important
The PDC Emulator FSMO role holder must be online when cloning begins.

Platform Requirements

Virtualized Domain Controller cloning requires:

- PDC emulator FSMO role hosted on a Windows Server 2012 DC
- PDC emulator available during cloning operations

Both cloning and safe restore require:

- Windows Server 2012 virtualized guests
- Virtualization host platform supports VM-Generation ID (VMGID)

Review the table below for virtualization products and whether they support virtualized domain controllers and VM-Generation ID.

Virtualization Product	Supports virtualized domain controllers and VMGID
Microsoft Windows Server 2012 server with Hyper-V Feature	Yes
Microsoft Windows Server 2012 Hyper-V Server	Yes
Microsoft Windows 8 with Hyper-V Client	Yes

Feature	
Windows Server 2008 R2 and Windows Server 2008	No
Non-Microsoft virtualization solutions	Contact vendor

Even though Microsoft supports Windows 7 Virtual PC, Virtual PC 2007, Virtual PC 2004, and Virtual Server 2005, they cannot run 64-bit guests, nor do they support VM-GenerationID.

For help with third party virtualization products and their support stance with virtualized domain controllers, contact that vendor directly.

For more information, review Support policy for Microsoft software running in non-Microsoft hardware virtualization software.

Critical Caveats

Virtualized domain controllers do *not* support safe restore of the following:

- VHD and VHDX files manually copied over existing VHD files
- VHD and VHDX files restored using file backup or full disk backup software

Note
VHDX files are new to Windows Server 2012 Hyper-V.

Neither of these operations is covered under VM-GenerationID semantics and therefore do not change the VM-Generation ID. Restoring domain controllers using these methods could either result in a USN rollback and either quarantine the domain controller or introduce lingering objects and the need for forest wide cleanup operations.

Warning
Virtualized domain controller safe restore is not a replacement for system state backups and the AD DS Recycle Bin.

After restoring a snapshot, the deltas of previously un-replicated changes originating from that domain controller after the snapshot are permanently lost. Safe restore implements automated non-authoritative restoration to prevent accidental domain controller quarantine *only*.

For more information about USN bubbles and lingering objects, see Troubleshooting Active Directory operations that fail with error 8606: "Insufficient attributes were given to create an object".

Virtualized Domain Controller Cloning

There are a number of stages and steps to cloning a virtualized domain controller, regardless of using graphical tools or Windows PowerShell. At a high level, the three stages are:

Prepare the environment

- Step 1: Validate that the hypervisor supports VM-Generation ID and therefore, cloning
- Step 2: Verify the PDC emulator role is hosted by a domain controller that runs Windows Server 2012 and that it is online and reachable by the cloned domain controller during cloning.

Prepare the source domain controller

- Step 3: Authorize the source domain controller for cloning
- Step 4: Remove incompatible services or programs or add them to the CustomDCCloneAllowList.xml file.
- Step 5: Create DCCloneConfig.xml
- Step 6: Take the source domain controller offline

Create the cloned domain controller

- Step 7: Copy or export the source VM and add the XML if not already copied
- Step 8: Create a new virtual machine from the copy
- Step 9: Start the new virtual machine to commence cloning

There are no procedural differences in the operation when using graphical tools such as the Hyper-V Management Console or command-line tools such as Windows PowerShell, so the steps are presented only once with both interfaces. This topic provides Windows PowerShell samples for you to explore end-to-end automation of the cloning process; they are not required for any steps. There is no graphical management tool for virtualized domain controllers included in Windows Server 2012.

There are several points in the procedure where you have choices for how to create the cloned computer and how you add the xml files; these steps are noted in the details below. The process is otherwise unalterable.

The following diagram illustrates the virtualized domain controller cloning process, where the domain already exists.

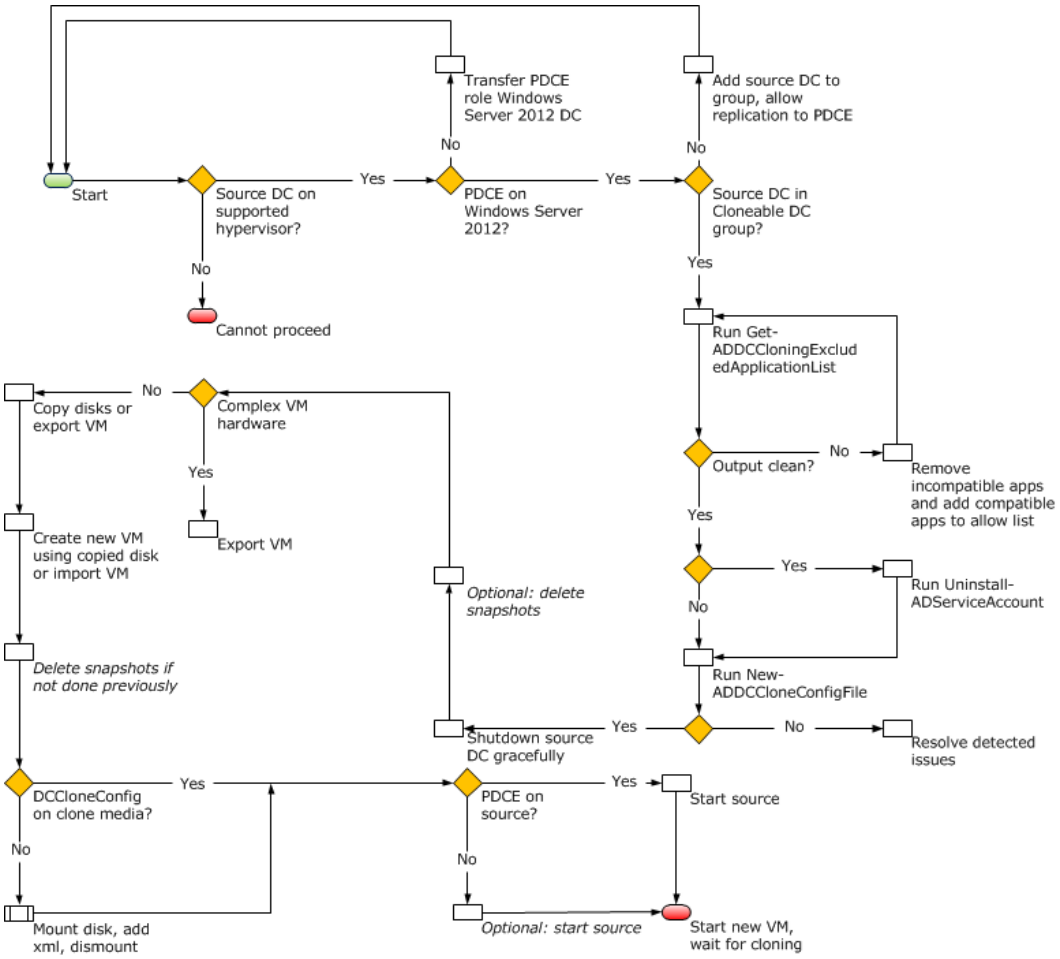

Step 1 - Validate the Hypervisor

Ensure the source domain controller is running on a supported hypervisor by reviewing vendor documentation. Virtualized domain controllers are hypervisor-independent and do not require Hyper-V.

If the hypervisor is Microsoft Hyper-V, ensure it is running on Windows Server 2012. You can validate this using Device Management

Open **Devmgmt.msc** and examine **System Devices** for installed Microsoft Hyper-V devices and drivers. The specific system device required for a virtualized domain controller is the **Microsoft Hyper-V Generation Counter** (driver: vmgencounter.sys).

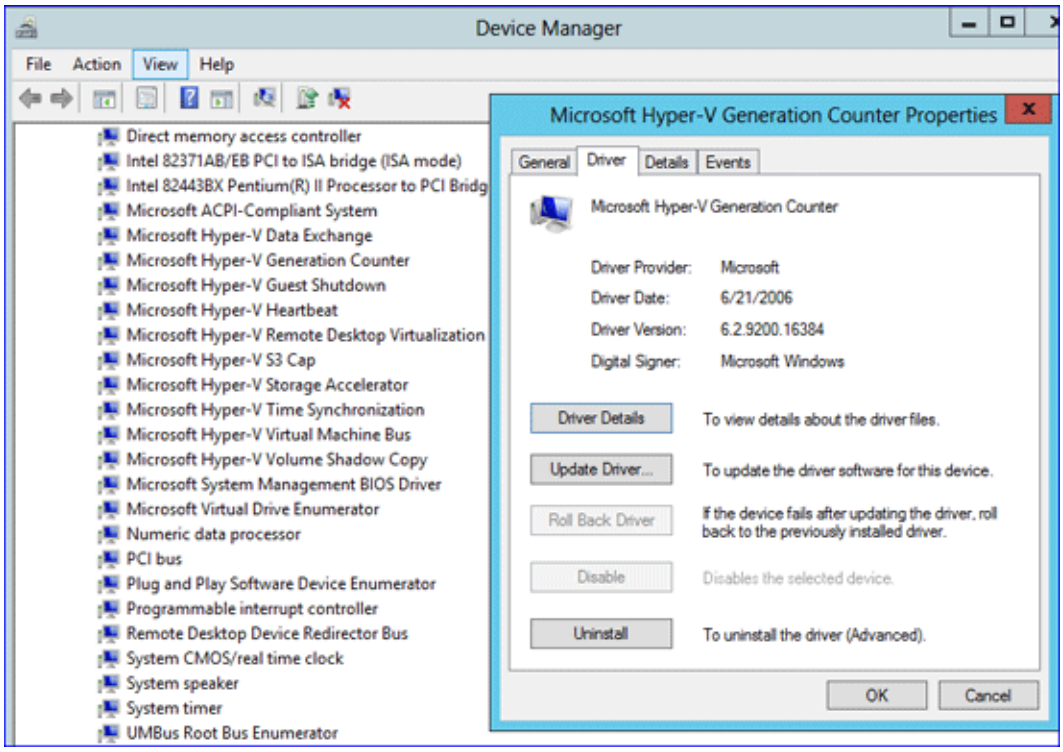

Step 2 - Verify the PDCE FSMO role

Before you attempt to clone a DC, you must validate that the domain controller hosting the Primary Domain Controller Emulator FSMO runs Windows Server 2012. The PDC emulator (PDCE) is required for several reasons:

1. The PDCE creates the special **Cloneable Domain Controllers** group and sets its permission on the root of the domain to allow a domain controller to clone itself.
2. The cloning domain controller contacts the PDCE directly using the DRSUAPI RPC protocol, in order to create computer objects for the clone DC.

Notes

Windows Server 2012 extends the existing Directory Replication Service (DRS) Remote Protocol (UUID **E3514235-4B06-11D1-AB04-00C04FC2DCD2**) to include a new RPC method **IDL_DRSAddCloneDC** (Opnum **28**). The **IDL_DRSAddCloneDC** method creates a new domain controller object by copying attributes from an existing domain controller object.

The states of a domain controller are composed of computer, server, NTDS settings, FRS, DFSR, and connection objects maintained for each domain controller. When duplicating an object, this RPC method replaces all references to the original domain controller with corresponding objects of the new domain controller. The caller must have the control access right DS-Clone-Domain-Controller on the domain naming context.

Use of this new method always requires direct access to the PDC emulator domain controller from the caller.

Because this RPC method is new, your network analysis software requires updated parsers to include fields for the new Opnum 28 in the existing UUID E3514235-4B06-11D1-AB04-00C04FC2DCD2. Otherwise, you cannot parse this traffic.

For more information, see 4.1.29 IDL_DRSAddCloneDC (Opnum 28).

This also means when using non-fully routed networks, virtualized domain controller cloning requires network segments with access to the PDCE. It is acceptable to move a cloned domain controller to a different network after cloning - just like a physical domain controller - as long as you are careful to update the AD DS logical site information.

Important
When cloning a domain that contains only a single domain controller, you must ensure the source DC is back online before starting the clone copies. A production domain should always contain at least two domain controllers.

Active Directory Users and Computers Method

1. Using the Dsa.msc snap-in, right click the domain and click **Operations Masters**. Note the domain controller named on the PDC tab and close the dialog.
2. Right-click that DC's computer object and click **Properties**, and then validate the Operating System info.

Windows PowerShell Method

You can combine the following Active Directory Windows PowerShell Module cmdlets to return the version of the PDC emulator:

```
Get-adddomaincontroller

Get-adcomputer
```

If not provided the domain, these cmdlets assume the domain of the computer where run.

The following command returns PDCE and Operating System info:

```
get-adcomputer(Get-ADDomainController -Discover -Service "PrimaryDC").name -property * |
format-list dnshostname,operatingsystem,operatingsystemversion
```

This example below demonstrates specifying the domain name and filtering the returned properties before the Windows PowerShell pipeline:

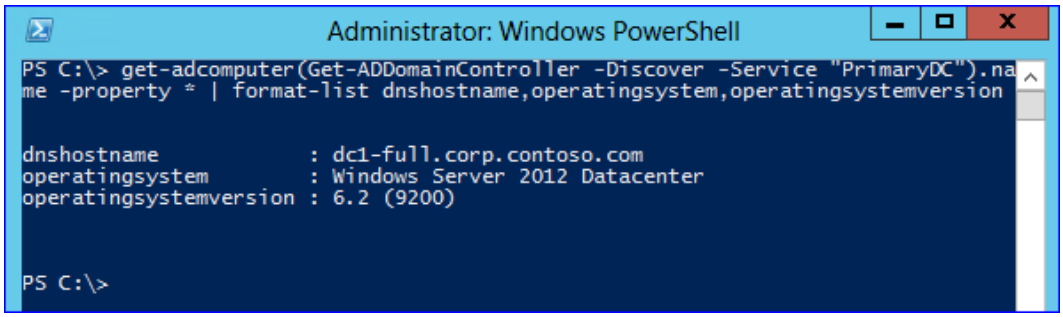

Step 3 - Authorize a Source DC

The source domain controller must have the control access right (CAR) **Allow a DC to create a clone of itself** on the domain NC head. By default, the well-known group **Cloneable Domain Controllers** has this permission and contains no members. The PDCE creates this group when that FSMO role transfers to a Windows Server 2012 domain controller.

Active Directory Administrative Center Method

1. Start Dsac.exe and navigate to the source DC, then open its detail page.
2. In the **Member Of** section, add the **Cloneable Domain Controllers** group for that domain.

Windows PowerShell Method

You can combine the following Active Directory Windows PowerShell Module cmdlets **get-adcomputer** and **add-adgroupmember** to add a domain controller to the **Cloneable Domain Controllers** group:

```
Get-adcomputer <dc name> | %{add-adgroupmember "cloneable domain controllers"
$_.samaccountname}
```

For instance, this adds server DC1 to the group, without the need to specify the distinguished name of the group member:

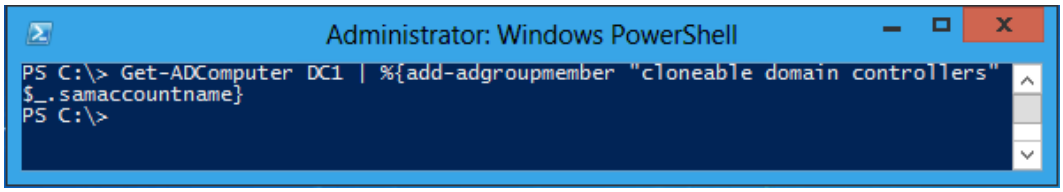

Rebuilding Default Permissions

If you remove this permission from the domain head, cloning fails. You can recreate the permission using the Active Directory Administrative Center or Windows PowerShell.

Active Directory Administrative Center Method

1. Open **Active Directory Administrative Center**, right-click the domain head, click **Properties**, click the **Extensions** tab, click **Security**, and then click **Advanced**. Click **This Object Only**.
2. Click **Add**, under **Enter the object name to select**, type the group name **Cloneable Domain Controllers.**
3. Under Permissions, click **Allow a DC to create a clone of itself**, and then click **OK**.

Note

You can also remove the default permission and add individual domain controllers. Doing so is likely to cause ongoing maintenance problems however, where new administrators are unaware of this customization. Changing the default setting does not increase security and is discouraged.

Windows PowerShell Method

Use the following commands in an administrator-elevated Windows PowerShell console prompt. These commands detect the domain name and add back in the default permissions:

```
import-module activedirectory

cd ad:

$domainNC = get-addomain

$dcgroup = get-adgroup "Cloneable Domain Controllers"

$sid1 = (get-adgroup $dcgroup).sid

$acl = get-acl $domainNC

$objectguid = new-object Guid 3e0f7e18-2c7a-4c10-ba82-4d926db99a3e

$ace1 = new-object System.DirectoryServices.ActiveDirectoryAccessRule
$sid1,"ExtendedRight","Allow",$objectguid

$acl.AddAccessRule($ace1)

set-acl -aclobject $acl $domainNC

cd c:
```

Alternatively, run the sample FixVDCPermissions.ps1 in a Windows PowerShell console, where the console starts as an elevated administrator on a domain controller in the affected domain. It automatically set the permissions. The sample is located in the appendix of this module.

Step 4 - Remove Incompatible applications or services (if not using CustomDCCloneAllowList.xml)

Any programs or services previously returned by Get-ADDCCloningExcludedApplicationList - *and not added to the CustomDCCloneAllowList.xml* - must be removed prior to cloning. Uninstalling the application or service is the recommended method.

Warning

Any incompatible program or service not uninstalled or added to the CustomDCCloneAllowList.xml prevents cloning.

Use the Get-AdComputerServiceAccount cmdlet to locate any standalone Managed Service Accounts (MSAs) in the domain and if this computer is using any of them. If any MSA is installed, use the Uninstall-ADServiceAccount cmdlet to remove the locally installed service account. Once you are done with taking the source domain controller offline in step 6, you can re-add the MSA using Install-ADServiceAccount when the server is back online. For more information, see Uninstall-ADServiceAccount.

Important
Standalone MSAs - first released in Windows Server 2008 R2 - were replaced in Windows Server 2012 with group MSAs. Group MSAs support cloning.

Step 5 - Create DCCloneConfig.xml

The DcCloneConfig.xml file is required for cloning Domain controllers. Its contents allow you to specify unique details like the new computer name and IP address.

The CustomDCCloneAllowList.xml file is optional unless you install applications or potentially incompatible Windows services on the source domain controller. The files require precise naming, formatting, and placement; otherwise, cloning fails.

For that reason, you should always use the Windows PowerShell cmdlets to create the XML files and place them in the correct location.

Generating with New-ADDCCloneConfigFile

The Active Directory Windows PowerShell module contains a new cmdlet in Windows Server 2012:

```
New-ADDCCloneConfigFile
```

You run the cmdlet on the proposed source domain controller that you intend to clone. The cmdlet supports multiple arguments and when used, always tests the computer and environment where it is run unless you specify the -offline argument.

ActiveDirectory Cmdlet	Arguments	Explanation
New-ADDCCloneConfigFile	*<no argument specified>*	Creates a blank DcCloneConfig.xml file in the DSA Working Directory (default: %systemroot%\ntds)
	-CloneComputerName	Specifies the clone DC computer name. String data type.

	-Path	Specifies the folder to create the DcCloneConfig.xml. If not specified, writes to the DSA Working Directory (default: %systemroot%\ntds). String data type.
	-SiteName	Specifies the AD logical site name to join during cloned computer account creation. String data type.
	-IPv4Address	Specifies the static IPv4 address of the cloned computer. String data type.
	-IPv4SubnetMask	Specifies the static IPv4 subnet mask of the cloned computer. String data type.
	-IPv4DefaultGateway	Specifies the static IPv4 default gateway address of the cloned computer. String data type.
	-IPv4DNSResolver	Specifies the static IPv4 DNS entries of the cloned computer in a comma-separated list. Array data type. Up to four entries can be provided.
	-PreferredWINSServer	Specifies the static IPv4 address of the primary WINS server. String data type.
	-AlternateWINSServer	Specifies the static IPv4 address of the secondary WINS server. String data type.
	-IPv6DNSResolver	Specifies the static IPv6 DNS entries of the cloned computer in a comma-separated list. There is no way to set Ipv6 static information in virtualized domain controller cloning. Array data type.

	-Offline	Does not perform the validation tests and overwrites any existing dccloneconfig.xml. Has no parameters. For more information, see Running New-ADDCCloneConfigFile in offline mode.
	-Static	Required if specifying static IP arguments IPv4SubnetMask, IPv4SubnetMask, or IPv4DefaultGateway. Has no parameters.

Tests performed when run in online mode:

- PDC Emulator is Windows Server 2012 or later
- Source domain controller is a member of Cloneable Domain Controllers group
- Source domain controller does not include any excluded applications or services
- Source domain controller does not already contain a DcCloneConfig.xml at the specified path

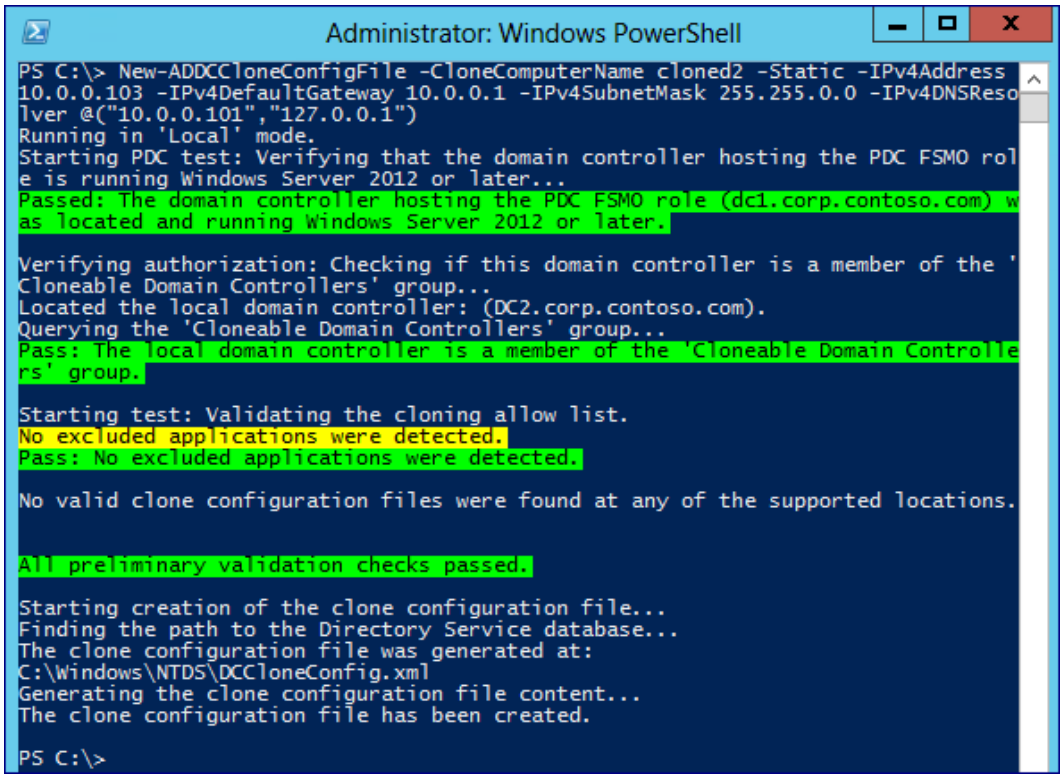

Step 6 - Take the Source Domain Controller Offline

You cannot copy a running source DC; it must be shutdown gracefully. Do not clone a domain controller stopped by graceless power loss.

Graphical Method

Use the shutdown button within the running DC, or the Hyper-V Manager shutdown button.

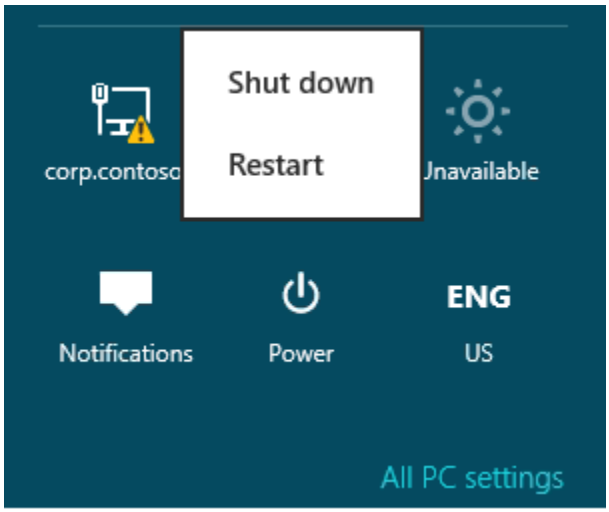

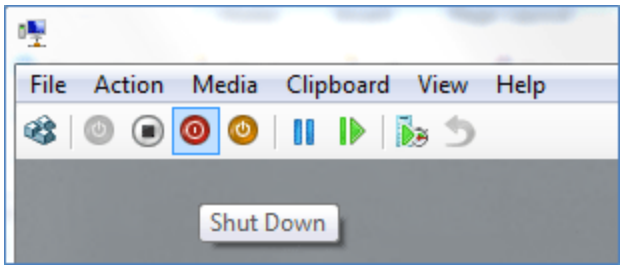

Windows PowerShell Method

You can shut down a virtual machine using either of the following cmdlets:

```
Stop-computer
```

```
Stop-vm
```

Stop-computer is a cmdlet that supports shutting down computers regardless of virtualization, and is analogous to the legacy Shutdown.exe utility. Stop-vm is a new cmdlet in the Windows Server 2012 Hyper-V Windows PowerShell module, and is equivalent to the power options in Hyper-V Manager. The latter is useful in lab environments where the domain controller often operates on a private virtualized network.

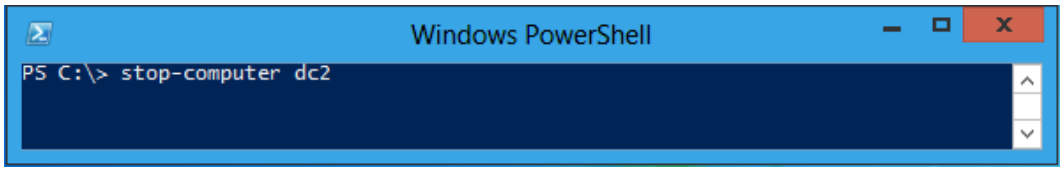

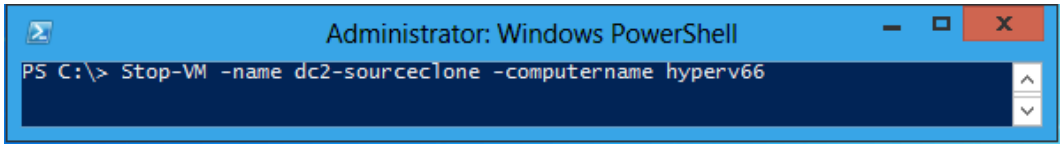

Step 7 - Copy Disks

An administrative choice is required in the copying phase:

- Copy the disks manually, without Hyper-V
- Export the VM, using Hyper-V
- Export the merged disks, using Hyper-V

All of a virtual machine's disks must be copied, not just the system drive. If the source domain controller uses differencing disks and you plan to move your cloned domain controller to another Hyper-V host, you must export.

Copying disks manually is recommended if the source domain controller has only *one* drive. Export/Import is recommended for VMs with *more than one* drive or other complex virtualized hardware customizations like multiple NICs.

If copying files manually, delete any snapshots prior to copying. If exporting the VM, delete snapshots prior to exporting or delete them from the new VM after importing.

Warning

Snapshots are differencing disks that can return a domain controller to previous state. If you were to clone a domain controller and then restore its pre-cloning snapshot, you would end up with duplicate domain controllers in the forest. There is no value in prior snapshots on a newly cloned domain controller.

Manually Copying Disks

Hyper-V Manager Method

Use the Hyper-V Manager snap-in to determine which disks are associated with the source domain controller. Use the Inspect option to validate if the domain controller uses differencing disks (which requires that you copy the parent disk also)

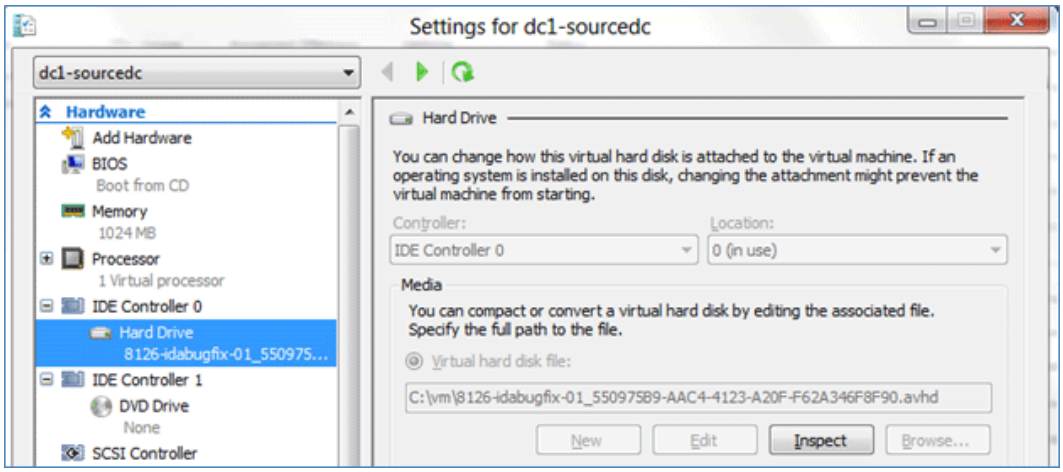

To delete snapshots, select a VM and delete the snapshot subtree.

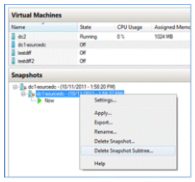

You can then manually copy the VHD or VHDX files using Windows Explorer, Xcopy.exe, or Robocopy.exe. No special steps are required. It is a best practice to change the file names even if moving to another folder.

Note

If copying between host computers on a LAN (1-Gbit or greater), the **Xcopy.exe /J** option copies VHD/VHDX files considerably faster than any other tool, at the cost of much greater bandwidth usage.

Windows PowerShell Method

To determine the disks using Windows PowerShell, use the Hyper-V Modules:

```
Get-vmidecontroller

Get-vmscsicontroller

Get-vmfibrechannelhba

Get-vmharddiskdrive
```

For example, you can return all IDE hard drives from a VM named **DC2** with the following sample:

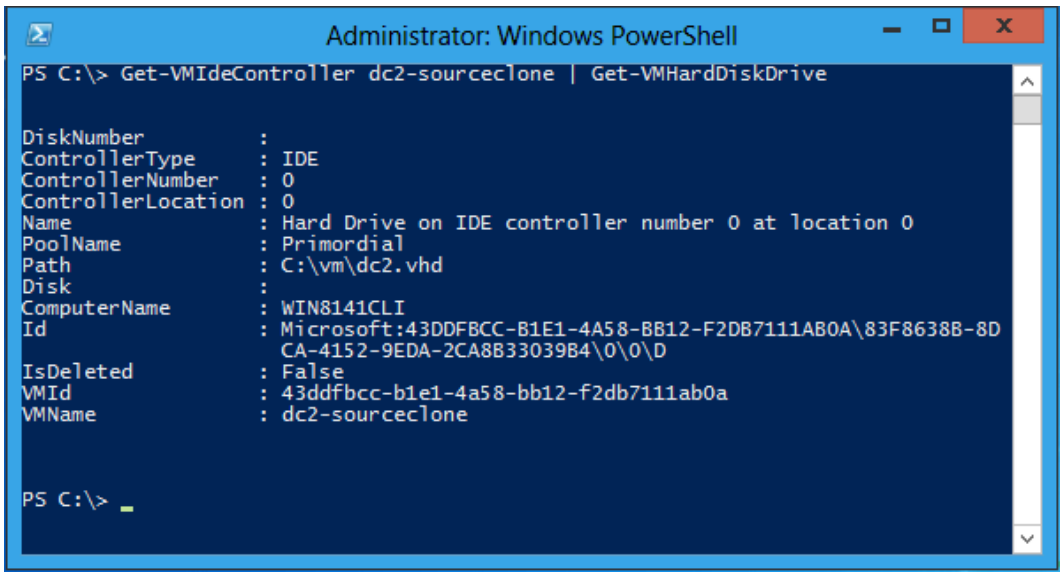

If the disk path points to an AVHD or AVHDX file, it is a snapshot. To delete the snapshots associated with a disk and merge in the real VHD or VHDX, use cmdlets:

```
Get-VMSnapshot
```

```
Remove-VMSnapshot
```

For example, to delete all snapshots from a VM named DC2-SOURCECLONE:

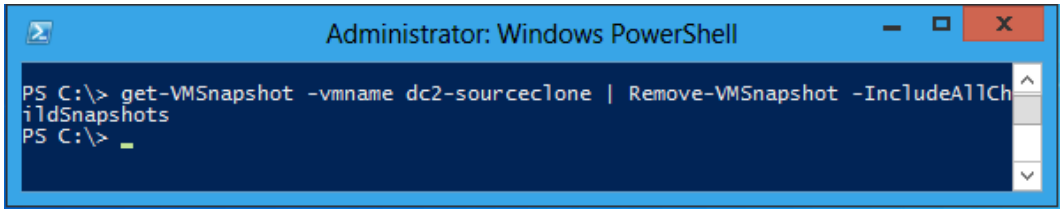

To copy the files using Windows PowerShell, use the following cmdlet:

```
Copy-Item
```

Combine with VM cmdlets in pipelines to aid automation. The pipeline is a channel used between multiple cmdlets to pass data. For example, to copy the drive of an offline source domain controller named DC2-SOURCECLONE to a new disk called c:\temp\copy.vhd without the need to know the exact path to its system drive:

```
Get-VMIdeController dc2-sourceclone | Get-VMHardDiskDrive | select-Object {copy-item -
path $_.path -destination c:\temp\copy.vhd}
```

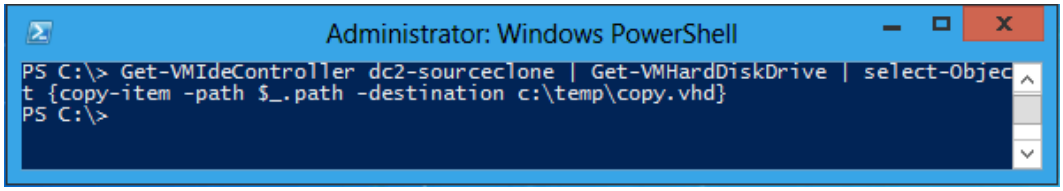

Important

You cannot use passthru disks with cloning, as they do not use a virtual disk file but instead an actual hard disk.

Note

For more information about more Windows PowerShell operations with pipelines, see Piping and the Pipeline in Windows PowerShell.

Exporting the VM

As an alternative to copying the disks, you can export the entire Hyper-V VM as a copy. Exporting automatically creates a folder named for the VM and containing all disks and configuration information.

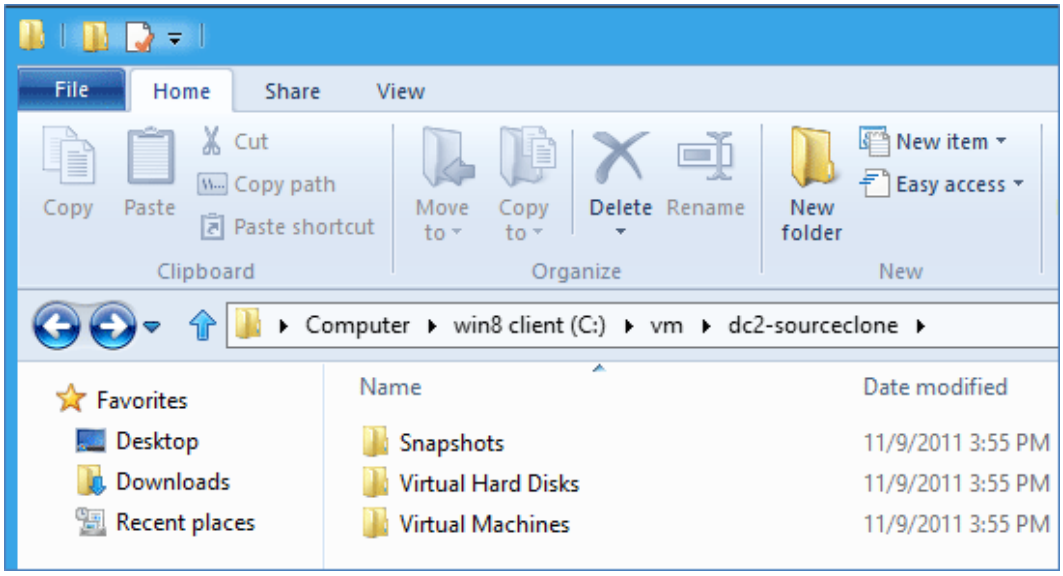

Hyper-V Manager Method

To export a VM with Hyper-V Manager:

1. Right-click the source domain controller and click **Export**.
2. Select an existing folder as the export container.
3. Wait for the Status column to stop showing **Exporting**.

Windows PowerShell Method

To export a VM using the Hyper-V Windows PowerShell module, use cmdlet:

```
Export-vm
```

For example, to export a VM named DC2-SOURCECLONE to a folder named C:\VM:

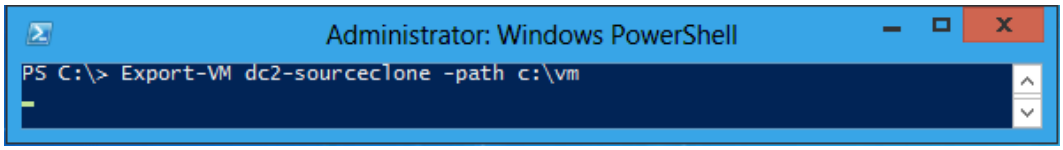

Note

Windows Server 2012 Hyper-V supports new export and import capabilities that are outside the scope of this training. Review TechNet for more information.

Exporting merged disks, using Hyper-V

The final option is to use the disk merge and conversion options within Hyper-V. These allow you to make a copy of an existing disk structure - even when including snapshot AVHD/AVHDX files - into a single new disk. Like the manual disk copy scenario, this is primarily intended for simpler virtual machines that only use a single drive, such as C:\. Its lone advantage is that, unlike manually copying, it does not require you to first delete snapshots. This operation is necessarily slower than simply deleting the snapshots and copying disks.

Hyper-V Manager Method

To create a merged disk using Hyper-V Manager:

1. Click **Edit Disk…**.
2. Browse for the lowest child disk. For example, if you are using a differencing disk, the child disk is the lowest child. If the virtual machine has a snapshot (or multiple ones), the currently selected snapshot is the lowest child disk.
3. Select the **Merge** option to create a single disk out of the entire parent-child structure.
4. Select a new virtual hard disk and provide a path. This reconciles the existing VHD/VHDX files into a single new portable unit that is not at risk of restoring previous snapshots.

Windows PowerShell Method

To create a merged disk from a complex set of parents using the Hyper-V Windows PowerShell module, use cmdlet:

```
Convert-vm
```

For example, to export the entire chain of a VM's disk snapshots (this time not including any differencing disks) and parent disk into a new single disk named DC4-CLONED.VHDX:

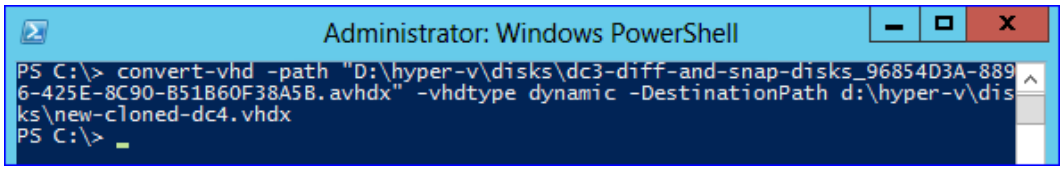

Adding XML to the Offline System Disk

If you did copy the Dccloneconfig.xml to the running source DC, you must copy the updated dccloneconfig.xml file to the offline copied/exported system disk now. Depending on installed applications detected with Get-ADDCCloningExcludedApplicationList earlier, you may also need to copy the CustomDCCloneAllowList.xml file to the disk.

The following locations can contain the DcCloneConfig.xml file:

1. DSA Working Directory
2. %windir%\NTDS
3. Removable read/write media, in order of drive letter, at the root of the drive

These paths are not configurable. After cloning begins, the cloning checks these locations in that specific order and uses the first DcCloneConfig.xml file found, regardless of the other folder's contents.

The following locations can contain the CustomDCCloneAllowList.xml file:

1. HKey_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters
   AllowListFolder (*REG_SZ*)
2. DSA Working Directory
3. %windir%\NTDS
4. Removable read/write media, in order of drive letter, at the root of the drive

You can run New-ADDCCloneConfigFile with the **-offline** argument (also known as offline mode) to create the DcCloneConfig.xml file and place it in a correct location. The following examples show how to run New-ADDCCloneConfigFile in offline mode.

To create a clone domain controller named CloneDC1 in offline mode, in a site called “REDMOND” with static IPv4 address, type:

```
New-ADDCCloneConfigFile –Offline –CloneComputerName CloneDC1 –SiteName REDMOND -
IPv4Address "10.0.0.2" -IPv4DNSResolver "10.0.0.1" -IPv4SubnetMask "255.255.0.0" -
IPv4DefaultGateway "10.0.0.1" –Static –Path F:\Windows\NTDS
```

To create a clone domain controller named Clone2 in offline mode with static IPv4 and static IPv6 settings, type:

```
New-ADDCCloneConfigFile –Offline -IPv4Address "10.0.0.2" -IPv4DNSResolver "10.0.0.1" -
IPv4SubnetMask "255.255.0.0" –Static -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -CloneComputerName "Clone2" -PreferredWINSServer
"10.0.0.1" -AlternateWINSServer "10.0.0.3" –Path F:\Windows\NTDS
```

To create a clone domain controller in offline mode with static IPv4 and dynamic IPv6 settings and specify multiple DNS servers for the DNS resolver settings, type:

```
New-ADDCCloneConfigFile –Offline -IPv4Address "10.0.0.10" -IPv4SubnetMask "255.255.0.0" -
IPv4DefaultGateway "10.0.0.1" -IPv4DNSResolver @( "10.0.0.1","10.0.0.2" ) –Static -
IPv6DNSResolver "2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" –Path F:\Windows\NTDS
```

To create a clone domain controller named Clone1 in offline mode with dynamic IPv4 and static IPv6 settings, type:

```
New-ADDCCloneConfigFile –Offline -Static -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" -CloneComputerName "Clone1" –PreferredWINSServer
"10.0.0.1" -AlternateWINSServer "10.0.0.3" -SiteName "REDMOND" –Path F:\Windows\NTDS
```

To create a clone domain controller in offline mode with dynamic IPv4 and dynamic IPv6 settings, type:

```
New-ADDCCloneConfigFile –Offline -IPv4DNSResolver "10.0.0.1" -IPv6DNSResolver
"2002:4898:e0:31fc:d61:2b0a:c9c9:2ccc" –Path F:\Windows\NTDS
```

Windows Explorer Method

Windows Server 2012 now offers a graphical option for mounting VHD and VHDX files. This requires installation of the Desktop Experience feature on Windows Server 2012.

1. Click the newly copied VHD/VHDX file that contains the source DC's system drive or DSA Working Directory location folder, and then click **Mount** from the **Disc Image Tools** menu.
2. In the now-mounted drive, copy the XML files to a valid location. You may be prompted for permissions to the folder.
3. Click the mounted drive and click **Eject** from the **Disk Tools** menu.

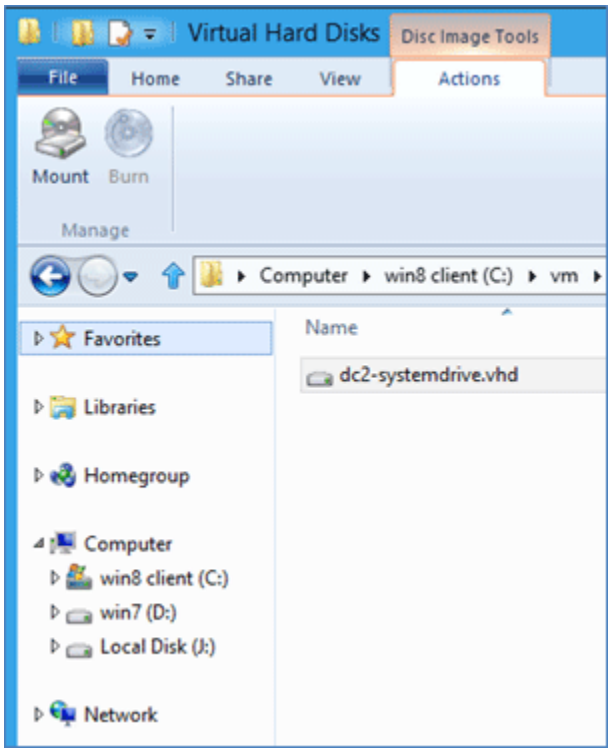
Virtual Hard Disks
Disc Image Tools
File
Home
Share
View
Actions
Mount
Burn
Manage
Computer ▸ win8 client (C:) ▸ vm ▸
Favorites
Libraries
Homegroup
Computer
win8 client (C:)
win7 (D:)
Local Disk (J:)
Network
Name
dc2-systemdrive.vhd

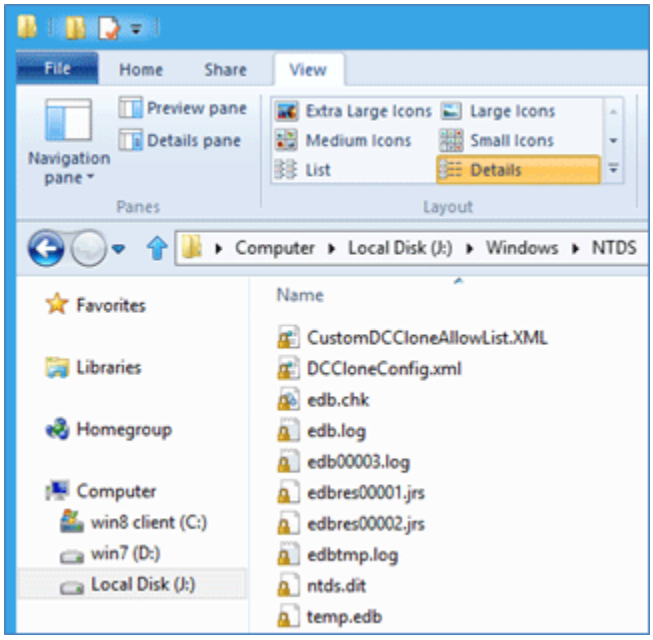
File
Home
Share
View
Navigation pane
Preview pane
Details pane
Panes
Extra Large Icons
Large Icons
Medium Icons
Small Icons
List
Details
Layout
Computer ▸ Local Disk (J:) ▸ Windows ▸ NTDS
Favorites
Libraries
Homegroup
Computer
win8 client (C:)
win7 (D:)
Local Disk (J:)
Name
CustomDCCloneAllowList.XML
DCCloneConfig.xml
edb.chk
edb.log
edb00003.log
edbres00001.jrs
edbres00002.jrs
edbtmp.log
ntds.dit
temp.edb

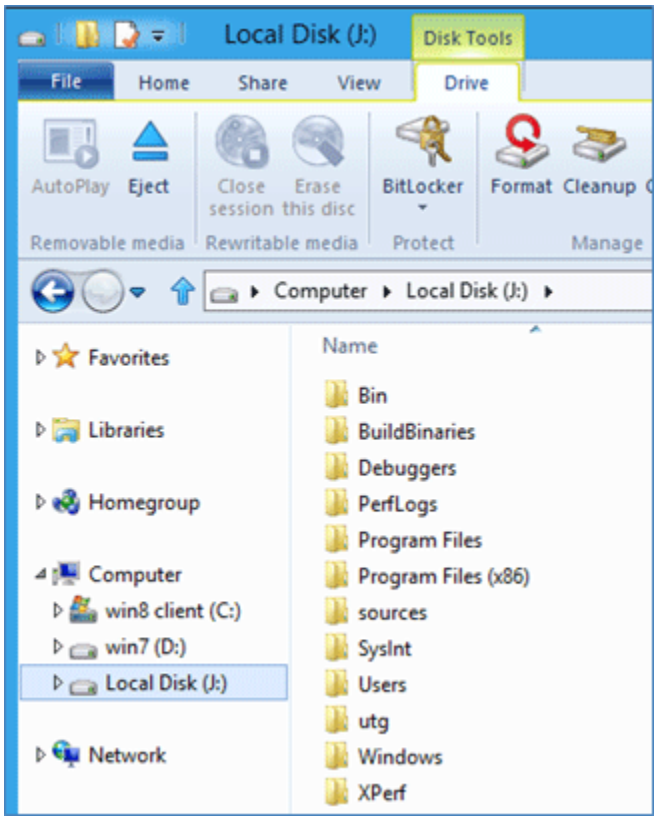

Windows PowerShell Method

Alternatively, you can mount the offline disk and copy the XML file using the Windows PowerShell cmdlets:

```
mount-vhd

get-disk

get-partition

get-volume

Add-PartitionAccessPath

Copy-Item
```

This allows you complete control over the process. For instance, the drive can be mounted with a specific drive letter, the file copied, and the drive dismounted.

```
mount-vhd <disk path> -passthru -nodriveletter | get-disk | get -partition | get-volume |
get-partition | where {$_.partition number -eq 2} | Add-PartitionAccessPath -accesspath
<drive letter>


copy-item <xml file path><destination path>\dccloneconfig.xml
```

`dismount-vhd <disk path>`

For example:

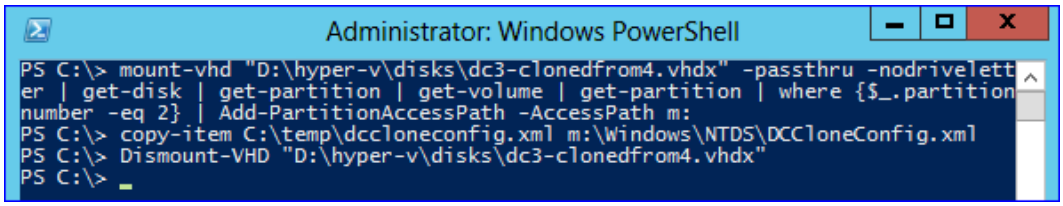

Alternatively, you can use the new **Mount-DiskImage** cmdlet to mount a VHD (or ISO) file.

Step 8 - Create the New Virtual Machine

The final configuration step before starting the cloning process is creating a new VM that uses the disks from the copied source domain controller. Depending on the selection made in the copying disks phase, you have two options:

1. Associate a new VM with the copied disk
2. Import the exported VM

Associating a New VM with Copied Disks

If you copied the system disk manually, you must create a new virtual machine using the copied disk. The hypervisor automatically sets the VM-Generation ID when a new VM is created; no configuration changes are required in the VM or Hyper-V host.

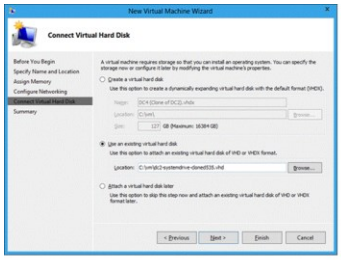

Hyper-V Manager Method

1. Create a new virtual machine.
2. Specify the VM name, memory, and network.
3. On the Connect Virtual Hard Disk page, specify the copied system disk.
4. Complete the wizard to create the VM.

If there were multiple disks, network adapters, or other customizations, configure them before starting the domain controller. The "Export-Import" method of copying disks is recommended for complex VMs.

Windows PowerShell Method

You can use the Hyper-V Windows PowerShell module to automate VM creation in Windows Server 2012, using the following cmdlet:

```
New-VM
```

For example, here the DC4-CLONEDFROMDC2 VM is created, using 1GB of RAM, booting from the c:\vm\dc4-systemdrive-clonedfromdc2.vhd file, and using the 10.0 virtual network:

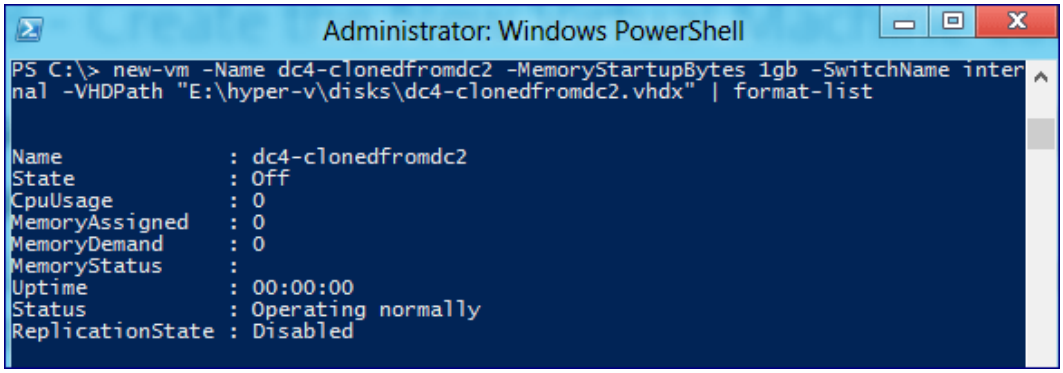

Import VM

If you previously exported your VM, you now need to import it back in as a copy. This uses the exported XML to recreate the computer using all the previous settings, drives, networks, and memory settings.

If you intend to create additional copies from the same exported VM, make as many copies of the exported VM as necessary. Then use Import for each copy.

Important
It is important to use the **Copy** option, as export preserves all information from the source; importing the server with **Move** or **In Place** causes information collision if done on the same Hyper-V host server.

Hyper-V Manager Method

To import using the Hyper-V Manager snap-in:

1. Click **Import Virtual Machine**
2. On the **Locate Folder** page, select the exported VM definition file using the Browse button
3. On the **Select Virtual Machine** page, click the source computer.
4. On the **Choose Import Type** page, click **Copy the virtual machine (create a new unique ID)**, then click **Finish**.

5. Rename the imported VM if importing on the same Hyper-V host; it will have the same name as the exported source domain controller.

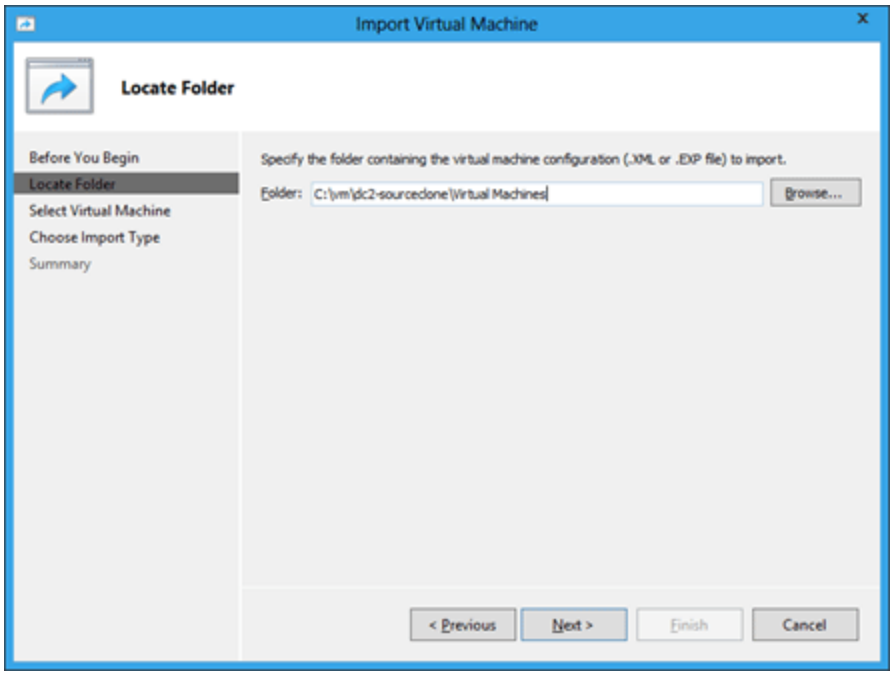

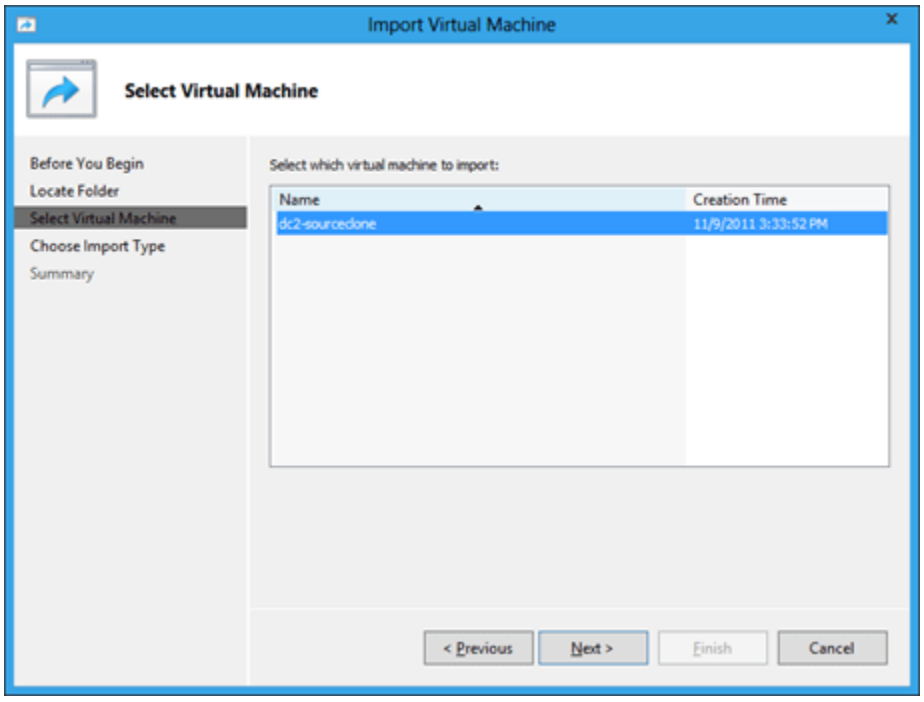

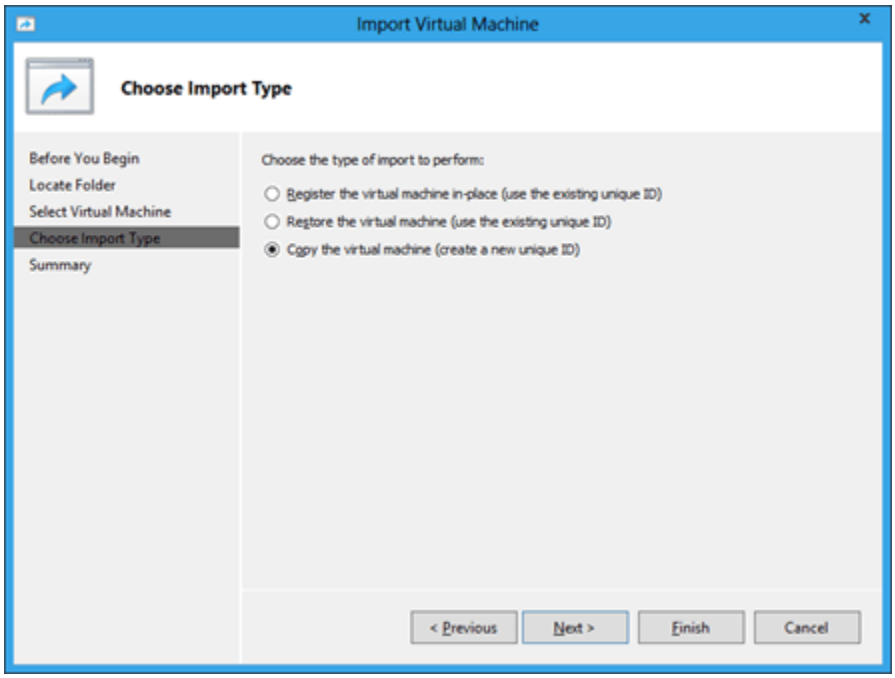

Remember to remove any imported snapshots, using the Hyper-V Management snap-in:

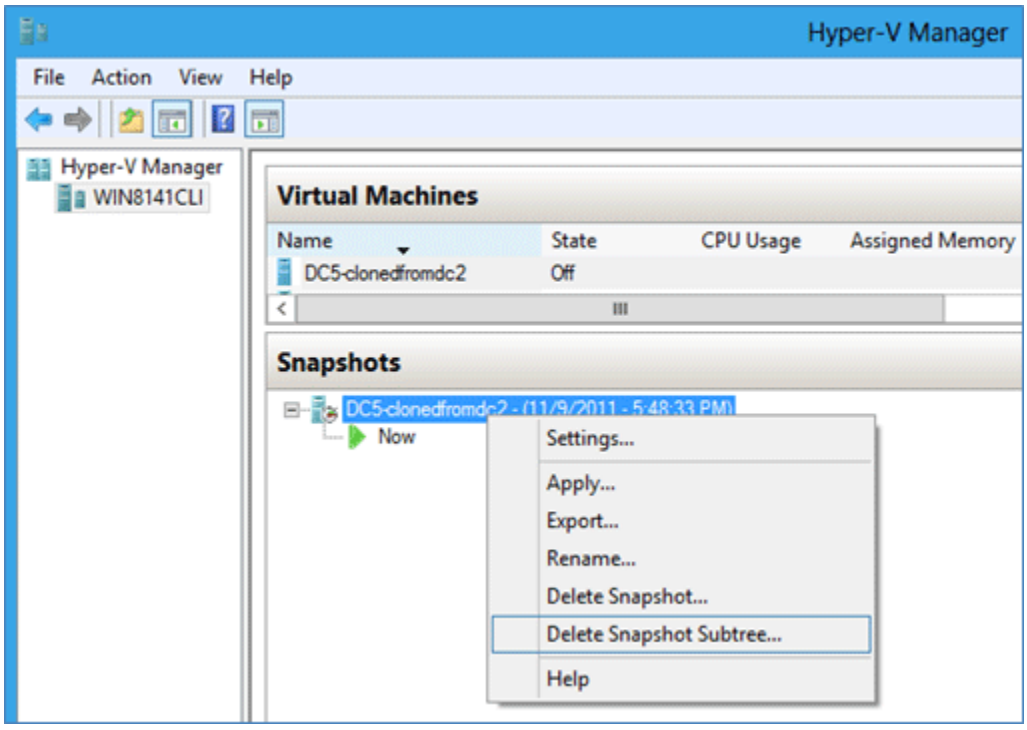

Warning

Deleting any imported snapshots is critically important; if applied, they would return the cloned domain controller to the state of a previous - and possibly live - DC, leading to replication failure, duplicate IP information, and other disruptions.

Windows PowerShell Method

You can use the Hyper-V Windows PowerShell module to automate VM import in Windows Server 2012, using the following cmdlets:

```
Import-VM
```

```
Rename-VM
```

For example, here the exported VM DC2-CLONED is imported using its automatically determined XML file, then renamed immediately to its new VM name DC5-CLONEDFROMDC2:

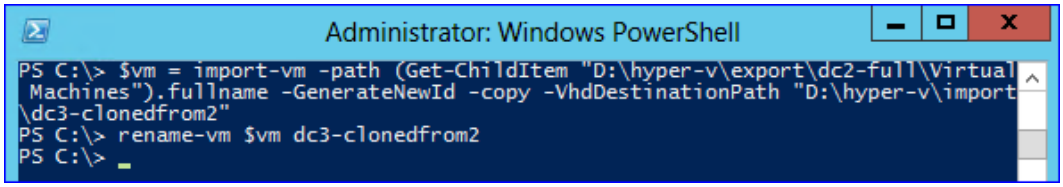

Remember to remove any imported snapshots, using the following cmdlets:

```
Get-VMSnapshot
```

```
Remove-VMSnapshot
```

For example:

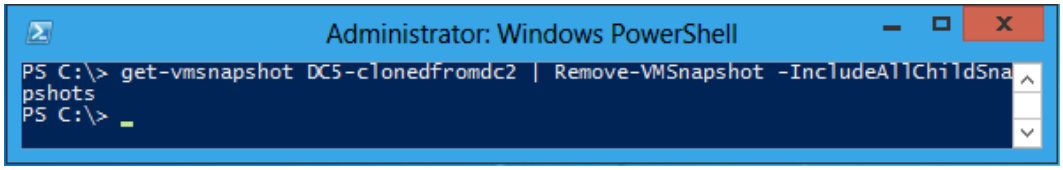

Warning

Ensure that, when importing the computer, static MAC addresses were not assigned to the source domain controller. If a source computer with a static MAC is cloned, those copied computers will not correctly send or receive any network traffic. Set a new unique static or dynamic MAC address if this is the case. You can see if a VM uses static MAC addresses with the command:

Get-VM –VMName *test-vm* | Get-VMNetworkAdapter | fl *

Step 9 - Clone the New Virtual Machine

Optionally, before you begin cloning, restart the offline clone source domain controller. Ensure that the PDC emulator is online, regardless.

To begin cloning, simply start the new virtual machine. The process initiates automatically and the domain controller reboots automatically after cloning is complete.

Important

Keeping domain controllers turned off for an extended period of time is not recommended and if the clone is joining the same site as its source DC, the initial intra and inter-site replication topology may take longer to build if the source domain controller is offline.

If using Windows PowerShell to start a VM, the new Hyper-V Module cmdlet is:

```
Start-VM
```

For example:

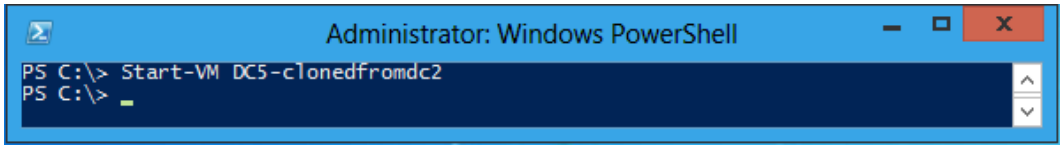

Once the computer restarts after cloning completes, it is a domain controller and you can logon on normally to confirm normal operation. If there are any errors, the server is set to start in Directory Services Restore Mode for investigation.

Virtualization safeguards

Unlike virtualized domain controller cloning, Windows Server 2012 virtualization safeguards have no configuration steps. The feature works without intervention as long as you meet some simple conditions:

- The hypervisor supports VM-Generation ID
- There is a valid partner domain controller that a restored domain controller can replicate changes from non-authoritatively.

Validate the Hypervisor

Ensure the source domain controller is running on a supported hypervisor by reviewing vendor documentation. Virtualized domain controllers are hypervisor-independent and do not require Hyper-V.

Review the previous Platform Requirements section for known VM-Generation ID support.

If you are migrating VMs from a source hypervisor to a different target hypervisor, virtualization safeguards may or may not be triggered depending on whether the hypervisors support VM-Generation ID, as explained in the following table.

Source hypervisor	Target hypervisor	Result
Supports VM-Generation ID	Does not support VM-Generation ID	Safeguards not triggered (if a DCCloneConfigFile.xml is present, DC will boot into DSRM)
Does not support VM-	Supports VM-Generation ID	Safeguards triggered

Source hypervisor	Target hypervisor	Result
Generation ID		
Supports VM-Generation ID	Supports VM-Generation ID	Safeguards not triggered because VM definition has not changed, which means so VM-Generation ID remains the same

Validate the Replication Topology

Virtualization safeguards initiate non-authoritative inbound replication for the delta of Active Directory replication as well as non-authoritative resynchronization of all SYSVOL contents. This ensures the domain controller returns from a snapshot with full functionality and is eventually consistent with the rest of the environment.

With this new capability come several requirements and limitations:

- A restored domain controller must be able to contact a writable DC
- All domain controllers in a domain must not be restored simultaneously
- Any changes originating from a restored domain controller that have not yet replicated outbound since the snapshot was taken are lost forever

While the troubleshooting section covers these scenarios, details below ensure you do not create a topology that could cause problems.

Writable Domain Controller Availability

If restored, a domain controller must have connectivity to a writable domain controller; a read-only domain controller cannot send the delta of updates. The topology is likely correct for this already, as a writable domain controller always needed a writable partner. However, if all writable domain controllers are restoring simultaneously, none of them can find a valid source. The same goes if the writable domain controllers are offline for maintenance or otherwise unreachable through the network.

Simultaneous Restore

Do not restore all domain controllers in a single domain simultaneously. If all snapshots restore at once, Active Directory replication works normally but SYSVOL replication halts. The restore architecture of FRS and DFSR require setting their replica instance to non-authoritative sync mode. If all domain controllers restore at once, and each domain controller marks itself non-authoritative for SYSVOL, they all will then try to synchronize group policies and scripts from an authoritative partner; at that point, though, all partners are also non-authoritative.

Important

If all domain controllers are restored at once, use the following articles to set one domain controller - typically the PDC emulator - as authoritative, so that the other domain controllers can return to normal operation:

Using the BurFlags registry key to reinitialize File Replication Service replica sets

How to force an authoritative and non-authoritative synchronization for DFSR-replicated SYSVOL (like "D4/D2" for FRS)

Warning

Do not run all domain controllers in a forest or domain on the same hypervisor host. That introduces a single point of failure that cripples AD DS, Exchange, SQL, and other enterprise operations each time the hypervisor goes offline. This is no different from using only one domain controller for an entire domain or forest. Multiple domain controllers on multiple platforms help provide redundancy and fault tolerance.

Post-Snapshot Replication

Do not restore snapshots until all locally originating changes made since snapshot creation have replicated outbound. Any originating changes are lost forever if other domain controllers did not already receive them through replication.

Use Repadmin.exe to show any un-replicated outbound changes between a domain controller and its partners:

1. Return the DC's partner names and DSA Object GUIDs with:

   ```
   Repadmin.exe /showrepl <DC Name of the partner> /repsto
   ```

2. Return the pending inbound replication of the partner domain controller to the domain controller to be restored:

   ```
   Repadmin.exe /showchanges < Name of partner DC> <DSA Object GUID of the domain controller being restored> <naming context to compare>
   ```

Alternatively, just to see the count of un-replicated changes:

```
Repadmin.exe /showchanges <Name of partner DC> <DSA Object GUID of the domain controller being restored> <naming context to compare> /statistics
```

For example (with output modified for readability and important entries ***italicized***), here you look at the replication partnerships of DC4:

```
C:\>repadmin.exe /showrepl dc4.corp.contoso.com /repsto


Default-First-Site-Name\DC4

DSA Options: IS_GC

Site Options: (none)

DSA object GUID: 5d083398-4bd3-48a4-a80d-fb2ebafb984f

DSA invocationID: 730fafec-b6d4-4911-88f2-5b64e48fc2f1
```

```
==== OUTBOUND NEIGHBORS FOR CHANGE NOTIFICATIONS ============


DC=corp,DC=contoso,DC=com

    Default-First-Site-Name\DC3 via RPC

        DSA object GUID: f62978a8-fcf7-40b5-ac00-40aa9c4f5ad3

        Last attempt @ 2011-11-11 15:04:12 was successful.

    Default-First-Site-Name\DC2 via RPC

        DSA object GUID: 3019137e-d223-4b62-baaa-e241a0c46a11

        Last attempt @ 2011-11-11 15:04:15 was successful.
```

Now you know that it is replicating with DC2 and DC3. You then show the list of changes that DC2 states it still does not have from DC4, and see that there is one new group:

```
C:\>repadmin /showchanges dc2.corp.contoso.com 5d083398-4bd3-48a4-a80d-fb2ebafb984f
dc=corp,dc=contoso,dc=com


==== SOURCE DSA: (null) ====

Objects returned: 1

(0) add CN=newgroup4,CN=Users,DC=corp,DC=contoso,DC=com

    1> parentGUID: 55fc995a-04f4-4774-b076-d6a48ac1af99

    1> objectGUID: 96b848a2-df1d-433c-a645-956cfbf44086

    2> objectClass: top; group

    1> instanceType: 0x4 = ( WRITE )

    1> whenCreated: 11/11/2011 3:03:57 PM Eastern Standard Time
```

You would also test the other partner to ensure that it had not already replicated.

Alternatively, if you did not care which objects had not replicated and only cared that any objects were outstanding, you can use the **/statistics** option:

```
C:\>repadmin /showchanges dc2.corp.contoso.com 5d083398-4bd3-48a4-a80d-fb2ebafb984f
dc=corp,dc=contoso,dc=com /statistics


**********************************************

********* Grand total ************************

Packets:              1

Objects:              1

Object Additions:     1

Object Modifications: 0
```

```
Object Deletions:      0
Object Moves:          0
Attributes:            12
Values:                13
```

Important

Test all writable partners if you see any failures or outstanding replication. As long as at least one is converged, it is generally safe to restore the snapshot, as transitive replication eventually reconciles the other servers.

Be sure to note any errors in replication shown by /showchanges and do not proceed until they are fixed.

Windows PowerShell Snapshot Cmdlets

The following Windows PowerShell Hyper-V module cmdlets provide snapshot capabilities in Windows Server 2012:

```
Checkpoint-VM

Export-VMSnapshot

Get-VMSnapshot

Remove-VMSnapshot

Rename-VMSnapshot

Restore-VMSnapshot
```

Virtualized Domain Controller Troubleshooting

This topic provides detailed methodology on troubleshooting the virtualized domain controller feature.

- Troubleshooting virtualized domain controller cloning
- Troubleshooting virtualized domain controller safe restore

Introduction

The most important way to improve your troubleshooting skills is build a test lab and rigorously examine normal, working scenarios. If you encounter errors, they are more obvious and easy to understand, since you then have a solid foundation of how domain controller promotion works. This also allows you to build your analysis and network analysis skills. This goes for all distributed systems technologies, not just virtualized domain controller deployment.

The critical elements to advanced troubleshooting of domain controller configuration are:

1. Linear analysis combined with focus and attention to detail.
2. Understanding network capture analysis
3. Understanding the built-in logs

The first and second are beyond the scope of this topic, but the third can be explained in some detail. Virtualized domain controller troubleshooting requires a logical and linear method. The key is to approach the issue using the data provided and only resort to complex tools and analysis when you have exhausted the provided output and logging.

Troubleshooting virtualized domain controller cloning

This sections covers:

- Tools for Troubleshooting
- Logging Options
- General Methodology for Troubleshooting Domain Controller Cloning
- Server Core and the Event Log
- Troubleshooting Specific Problems

The troubleshooting strategy for virtualized domain controller cloning follows this general format:

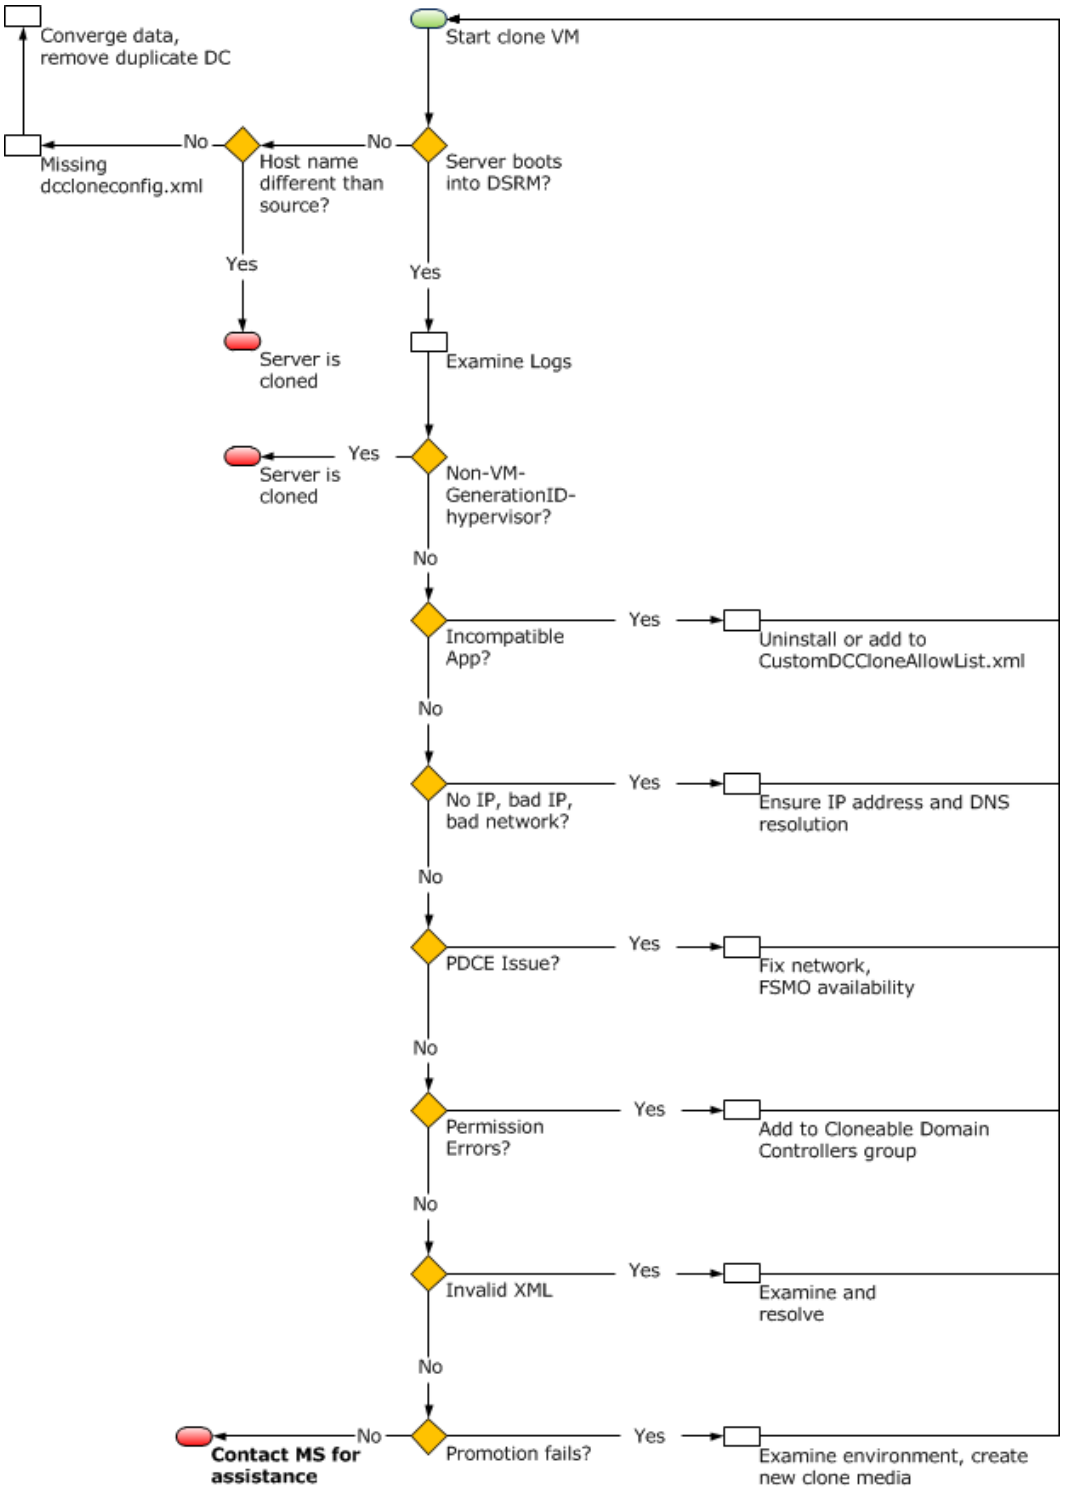
Converge data, remove duplicate DC
Start clone VM
No
Missing dccloneconfig.xml
No
Host name different than source?
Server boots into DSRM?
Yes
Server is cloned
Yes
Examine Logs
Yes
Server is cloned
Non-VM-GenerationID-hypervisor?
No
Incompatible App?
Yes
Uninstall or add to CustomDCCloneAllowList.xml
No
No IP, bad IP, bad network?
Yes
Ensure IP address and DNS resolution
No
PDCE Issue?
Yes
Fix network, FSMO availability
No
Permission Errors?
Yes
Add to Cloneable Domain Controllers group
No
Invalid XML
Yes
Examine and resolve
No
No
Contact MS for assistance
Promotion fails?
Yes
Examine environment, create new clone media

Tools for Troubleshooting

Logging Options

The built-in logs are the most important tool for troubleshooting issues with domain controller cloning. All of these logs are enabled and configured for maximum verbosity, by default.

Operation	Log
Cloning	• Event viewer\Windows logs\System • Event viewer\Applications and services logs\Directory Service • %systemroot%\debug\dcpromo.log
Promotion	• %systemroot%\debug\dcpromo.log • Event viewer\Applications and services logs\Directory Service • Event viewer\Windows logs\System • Event viewer\Applications and services logs\File Replication Service • Event viewer\Applications and services logs\DFS Replication

Tools and Commands for Troubleshooting Domain Controller Configuration

To troubleshoot issues not explained by the logs, use the following tools as a starting point:

- Dcdiag.exe
- Repadmin.exe
- Network Monitor 3.4

General Methodology for Troubleshooting Domain Controller Cloning

1. Is the VM booting into DS Repair Mode (DSRM)? This indicates troubleshooting is necessary. To log on in DSRM, use **.\Administrator** account and specify the DSRM password.
    a. Examine the Dcpromo.log.
        i. Did initial cloning steps succeed but domain controller promotion fail?
        ii. Do errors indicate issues with the local domain controller or with the AD DS environment, such as errors returned from the PDC emulator?
    b. Examine the System and Directory Services event logs and the dccloneconfig.xml and CustomDCCloneAllowList.xml

i. Does an incompatible application need to be in the CustomDCCloneAllowList.xml allow list?
ii. Is the IP address or computer name either duplicated or invalid in the dccloneconfig.xml?
iii. Is the Active Directory site invalid in the dccloneconfig.xml?
iv. Is the IP address not set in the dccloningconfig.xml and there is no DHCP server available?
v. Is the PDC emulator online and available through the RPC protocol?
vi. Is the domain controller a member of the Cloneable Domain Controllers group? Is the permission **Allow a DC to create a clone of itself** set on the domain root for that group?
vii. Does the Dccloneconfig.xml file contain syntax errors that prevent correct parsing?
viii. Is the hypervisor supported?
ix. Did domain controller promotion fail after cloning began successfully?
x. Was the maximum number of auto-generated domain controller names (9999) exceeded?
xi. Is the MAC address duplicated?

2. Is host name of the clone the same as the source DC?
   a. Is there a Dccloneconfig.xml file in one of the allowed locations?
3. Is the VM booting into normal mode and cloning completed, but the domain controller is not functioning correctly?
   a. First check if the host name is changed on the clone. If the host name is different, cloning has at least partially completed.
   b. Does the domain controller have a duplicate IP address of the source domain controller from the dccloneconfig.xml, but the source domain controller was offline during cloning?
   c. If the domain controller is advertising, treat the issue as any normal post-promotion issue you would have without cloning.
   d. If the domain controller is not advertising, examine the Directory Service, System, Application, File Replication and DFS Replication event logs for post-promotion errors.

Disabling DSRM Boot

Once booted into DSRM due to any error, diagnose the cause for failure and if the dcpromo.log does not indicate that cloning cannot be retried, fix the cause for failure and reset the DSRM flag. A failed clone does not return to normal mode on its own on the next reboot; you must remove the DS Restore Mode boot flag in order to try cloning again. All of these steps require running as an elevated administrator.

Removing DSRM with Msconfig.exe

To turn DSRM boot off using a GUI, use the System Configuration tool:

1. Run msconfig.exe

2. On the **Boot** tab, under **Boot Options**, de-select **Safe boot** (it is already selected with the option **Active Directory repair** enabled)
3. Click OK and restart when prompted

Removing DSRM with Bcdedit.exe

To turn DSRM boot off from the command-line, use the Boot Configuration Data Store Editor:

1. Open a CMD prompt and run:

```
Bcdedit.exe /deletevalue safeboot
```

2. Restart the computer with:

```
Shutdown.exe /t /0 /r
```

Notes

Bcdedit.exe also works in a Windows PowerShell console. The commands there are:

Bcdedit.exe /deletevalue safeboot

Restart-computer

Server Core and the Event Log

The event logs contain much of the useful information about virtualized domain controller cloning operations. By default, a Windows Server 2012 computer installation is a Server Core installation, which means there is no graphical interface and therefore, no way to run the local Event Viewer snap-in.

To review the event logs on a server running a Server Core installation:

- Run the Wevtutil.exe tool locally
- Run PowerShell cmdlet Get-WinEvent locally
- If you have enabled the Windows Advanced Firewall rules for the “Remote Event Log Management” groups (or equivalent ports) to allow inbound communication, you can manage the event log remotely using Eventvwr.exe, wevtutil.exe, or Get-Winevent. This can be done on Server Core installation using NETSH.exe, Group Policy, or the new Set-NetFirewallRule cmdlet in Windows PowerShell 3.0.

Warning

Do not attempt to add the graphical shell back to the computer while it is in DSRM. Windows servicing stack (CBS) cannot operate correctly while in Safe Mode or DSRM. Attempts to add features or roles while in DSRM will not complete and leave the computer in an unstable state until it is booted normally. Since a virtualized domain controller clone in DSRM cannot boot normally, and should not be booted normally under most circumstances, it is impossible to safely add the graphical shell. Doing so is unsupported and may leave you with an unusable server.

Troubleshooting Specific Problems

Events

All virtualized domain controller cloning events write to the Directory Services event log of the clone domain controller VM. The Application, File Replication Service, and DFS Replication event logs may also contain useful troubleshooting information for failed cloning. Failures during the RPC call to the PDC emulator may be available in the event log on the PDC emulator.

Below are the Windows Server 2012 cloning-specific events in the Directory Services event log, with notes and suggested resolutions for errors.

Directory Services Event Log

Event ID	**2160**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The local <*COMPUTERNAME*> has found a virtual domain controller cloning configuration file. The virtual domain controller cloning configuration file is found at: %1 The existence of the virtual domain controller cloning configuration file indicates that the local virtual domain controller is a clone of another virtual domain controller. The <*COMPUTERNAME*> will start to clone itself.
Notes and resolution	This is a success event and only an issue if unexpected. Examine the DSA Working Directory, %systemroot%\ntds, and root of any local or removable disks for the dcclconeconfig.xml file.

Event ID	**2161**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The local <*COMPUTERNAME*> did not find the

	virtual domain controller cloning configuration file. The local machine is not a cloned DC.
Notes and resolution	This is a success event and only an issue if unexpected. Examine the DSA Working Directory, %systemroot%\ntds, and root of any local or removable disks for the dcclconeconfig.xml file.

Event ID	**2162**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Virtual domain controller cloning failed. Please check events logged in System event logs and %systemroot%\debug\dcpromo.log for more information on errors that correspond to the virtual domain controller cloning attempt. Error code: %1
Notes and resolution	Follow message instructions, this error is a catchall.

Event ID	**2163**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	DsRoleSvc service was started to clone the local virtual domain controller.
Notes and resolution	This is a success event and only an issue if unexpected. Examine the DSA Working Directory, %systemroot%\ntds, and root of any local or removable disks for the dcclconeconfig.xml file.

Event ID	**2164**

Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to start the DsRoleSvc service to clone the local virtual domain controller.
Notes and resolution	Examine the service settings for the DS Role Server service (DsRoleSvc) and ensure its start type is set to manual. Validate that no third party program is preventing the start of this service.

Event ID	**2165**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to start a thread during the cloning of the local virtual domain controller. Error code:%1 Error message:%2 Thread name:%3
Notes and resolution	Contact Microsoft Product Support

Event ID	**2166**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> needs RPCSS service to initiate rebooting into DSRM. Waiting for RPCSS to initialize into a running state failed. Error code:%1
Notes and resolution	Examine the System event log and service settings for the RPC Server service (Rpcss)

Event ID	2167
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> could not initialize virtual domain controller knowledge. See previous event log entry for details. Additional Data Failure code:%1
Notes and resolution	Follow message instructions, this error is a catchall.

Event ID	2168
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Microsoft-Windows-ActiveDirectory_DomainService The DC is running on a supported hypervisor. VM Generation ID is detected. Current value of VM Generation ID: %1
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2169
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	There is no VM Generation ID detected. The DC is hosted on a physical machine, a down-level version of Hyper-V, or a hypervisor that does not support the VM Generation ID. Additional Data

	Failure code returned when checking VM Generation ID:%1
Notes and resolution	This is a success event if not intending to clone. Otherwise, examine the System event log and review hypervisor product support documentation.

Event ID	**2170**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Warning
Message	A Generation ID change has been detected. Generation ID cached in DS (old value):%1 Generation ID currently in VM (new value):%2 The Generation ID change occurs after the application of a virtual machine snapshot, after a virtual machine import operation or after a live migration operation. <*COMPUTERNAME*> will create a new invocation ID to recover the domain controller. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services aware backup application.
Notes and resolution	This is a success event if intending to clone. Otherwise, examine the System event log.

Event ID	**2171**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	No Generation ID change has been detected. Generation ID cached in DS (old value):%1

	Generation ID currently in VM (new value):%2
Notes and resolution	This is a success event if not intending to clone, and should be seen at every reboot of a virtualized DC. Otherwise, examine the System event log.

Event ID	**2172**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Read the msDS-GenerationId attribute of the Domain Controller's computer object. msDS-GenerationId attribute value:%1
Notes and resolution	This is a success event if intending to clone. Otherwise, examine the System event log.

Event ID	**2173**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Failed to read the msDS-GenerationId attribute of the Domain Controller's computer object. This may be caused by database transaction failure, or the generation id does not exist in the local database. The msDS-GenerationId does not exist during the first reboot after dcpromo or the DC is not a virtual domain controller. Additional Data Failure code:%1
Notes and resolution	This is a success event if intending to clone and it is the first VM reboot after cloning has completed. It can also be ignored on non-virtual Domain controllers. Otherwise, examine the System event log.

Event ID	2174
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The DC is neither a virtual domain controller clone nor a restored virtual domain controller snapshot.
Notes and resolution	This is a success event if not intending to clone. Otherwise, examine the System event log.

Event ID	2175
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Virtual domain controller clone configuration file exists on an unsupported platform.
Notes and resolution	This occurs when a dccloneconfig.xml is found but a VM Generation-ID could not be found, such as when a dccloneconfig.xml file is found on a physical computer or on a hypervisor that does not support VM Generation-ID.

Event ID	2176
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Renamed virtual domain controller clone configuration file. Additional Data Old file name:%1 New file name:%2
Notes and resolution	Rename expected when booting a source VM back up, because the VM Generation ID has not changed. This prevents the source domain

	controller from trying to clone.

Event ID	**2177**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Renaming virtual domain controller clone configuration file failed. Additional Data File name:%1 Failure code:%2 %3
Notes and resolution	Rename attempt expected when booting a source VM back up, because the VM Generation ID has not changed. This prevents the source domain controller from trying to clone. Manually rename the file and investigate installed third party products that may be preventing the file rename.

Event ID	**2178**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Detected virtual domain controller clone configuration file, but VM Generation ID has not been changed. The local DC is the clone source DC. Rename the clone configuration file.
Notes and resolution	Expected when booting a source VM back up, because the VM Generation ID has not changed. This prevents the source domain controller from trying to clone.

Event ID	**2179**

Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The msDS-GenerationId attribute of the Domain Controller's computer object has been set to the following parameter: GenerationID attribute:%1
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	**2180**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Warning
Message	Failed to set the msDS-GenerationId attribute of the Domain Controller's computer object. Additional Data Failure code:%1
Notes and resolution	Examine the System event log and Dcpromo.log. Lookup the specific error in MS TechNet, MS Knowledgebase, and MS blogs to determine its usual meaning, and then troubleshoot based on those results.

Event ID	**2182**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Internal event: The Directory Service has been asked to clone a remote DSA:
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2183
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Internal event: <*COMPUTERNAME*> completed the request to clone the remote Directory System Agent. Original DC name:%3 Request clone DC name:%4 Request clone DC site:%5 Additional Data Error value:%1 %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2184
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to create a domain controller account for the cloned DC. Original DC name:%1 Allowed number of cloned DC:%2 The limit on the number of domain controller accounts that can be generated by cloning <*COMPUTERNAME*>was exceeded.
Notes and resolution	A single source domain controller name can only automatically generate 9999 times if domain controllers are not demoted, based on the naming convention. Use the <computername> element in the XML to generate a new unique name or clone from a differently named DC.

Event ID	2191

Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> set the following registry value to disable DNS updates. Registry Key:%1 Registry Value: %2 Registry Value data: %3 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning. The cloning process will enable DNS updates again after cloning is completed.
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	**2192**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to set the following registry value to disable DNS updates. Registry Key:%1 Registry Value: %2 Registry Value data: %3 Error code:%4 Error message:%5 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing

	cloning.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking registry updates.

Event ID	**2193**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> set the following registry value to enable DNS updates. Registry Key:%1 Registry Value: %2 Registry Value data: %3 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning.
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	**2194**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to set the following registry value to enable DNS updates. Registry Key:%1 Registry Value: %2 Registry Value data: %3 Error code:%4 Error message:%5

	During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking registry updates.

Event ID	**2195**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Failed to set DSRM boot. Error code:%1 Error message:%2 When virtual domain controller cloning failed or virtual domain controller clone configuration file appears on a non-supported hypervisor, the local machine will reboot into DSRM for troubleshooting. Setting DSRM boot failed.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking registry updates.

Event ID	**2196**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Failed to enable shutdown privilege. Error code:%1 Error message:%2 When virtual domain controller cloning failed or virtual domain controller clone configuration file

	appears on a non-supported hypervisor, the local machine will reboot into DSRM for troubleshooting. Enabling shutdown privilege failed.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking privilege usage.

Event ID	**2197**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Failed to initiate system shutdown. Error code:%1 Error message:%2 When virtual domain controller cloning failed or virtual domain controller clone configuration file appears on a non-supported hypervisor, the local machine will reboot into DSRM for troubleshooting. Initiating system shutdown failed.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking privilege usage.

Event ID	**2198**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to create or modify the following cloned DC object. Additional data: Object: %1 Error value: %2

	%3
Notes and resolution	Lookup the specific error in MS TechNet, MS Knowledgebase, and MS blogs to determine its usual meaning, and then troubleshoot based on those results.

Event ID	**2199**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to create the following cloned DC object because the object already exists. Additional data: Source DC: %1 Object: %2
Notes and resolution	Validate the dccloneconfig.xml did not specify an existing domain controller or that copies of the dccloneconfig.xml have been used on multiple clones without editing the name. If the collision is still unexpected, determine which administrator promoted it; contact them to discuss if the existing domain controller should be demoted, the existing domain controller metadata cleaned, or if the clone should use a different name.

Event ID	**2203**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Last virtual domain controller cloning failed. This is the first reboot since then so this should be a re-try of the cloning. However, neither

	virtual domain controller clone configuration file exists nor virtual machine generation ID change is detected. Boot into DSRM. Last virtual domain controller cloning failed:%1 Virtual domain controller clone configuration file exists:%2 Virtual machine generation ID change is detected:%3
Notes and resolution	Expected if cloning failed previously, due to missing or invalid dccloneconfig.xml

Event ID	2210
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<COMPUTERNAME> failed to create objects for clone domain controller. Additional data: Clone Id: %6 Clone domain controller name: %1 Retry loop: %2 Exception value: %3 Error value: %4 DSID: %5
Notes and resolution	Review the System and Directory Services event logs and the dcpromo.log for further details on why cloning failed.

Event ID	2211
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> has created objects for clone domain controller.

	Additional data: Clone Id: %3 Clone domain controller name: %1 Retry loop: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2212
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> started to create objects for the clone domain controller. Additional data: Clone Id: %1 Clone name: %2 Clone site: %3 Clone RODC: %4
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2213
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> created a new KrbTgt object for Read-Only domain controller cloning. Additional data: Clone Id: %1 New KrbTgt Object Guid: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2214
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will create a computer object for the clone domain controller. Additional data: Clone Id: %1 Original domain controller: %2 Clone domain controller: %3
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2215
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will add the clone domain controller in the following site. Additional data: Clone Id: %1 Site: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2216
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will create a servers container for the clone domain controller. Additional data: Clone Id: %1

	Servers Container: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2217
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will create a server object for the clone domain controller. Additional data: Clone Id: %1 Server Object: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2218
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will create a NTDS Settings object for the clone domain controller. Additional data: Clone Id: %1 Object: %2
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2219
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational

Message	<COMPUTERNAME> will create connection objects for the clone Read-Only domain controller. Additional data: Clone Id: %1
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2220
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> will create SYSVOL objects for the clone Read-Only domain controller. Additional data: Clone Id: %1
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2221
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<COMPUTERNAME> failed to generate a random password for the cloned domain controller. Additional data: Clone Id: %1 Clone domain controller name: %2 Error: %3 %4
Notes and resolution	Examine the system event log for further details on why the machine account password could not be created.

Event ID	2222
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<COMPUTERNAME> failed to set password for the cloned domain controller. Additional data: Clone Id: %1 Clone domain controller name: %2 Error: %3 %4
Notes and resolution	Examine the system event log for further details on why the machine account password could not be set.

Event ID	2223
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<COMPUTERNAME> successfully set machine account password for the cloned domain controller. Additional data: Clone Id: %1 Clone domain controller name: %2 Total retry times: %3
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2224
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Virtual domain controller cloning failed. The

	following %1 Managed Service Account(s) exist on the cloned machine: %2 For cloning to succeed, all Managed Service Accounts must be removed. This can be done using the Remove-ADComputerServiceAccount PowerShell cmdlet.
Notes and resolution	Expected when using standalone MSAs (not group MSA). Do *not* follow the event advice to remove the account - it is incorrectly written. Use Uninstall-AdServiceAccount - http://technet.microsoft.com/en-us/library/hh852310. Standalone MSAs - first released in Windows Server 2008 R2 - were replaced in Windows Server 2012 with group MSAs (gMSA). GMSAs support cloning.

Event ID	2225
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The cached secrets of the following security principal have been successfully removed from local domain controller: %1 After cloning a read-only domain controller, secrets which were previously cached on the cloning source read-only domain controller will be removed on the cloned domain controller.
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	2226
Source	Microsoft-Windows-ActiveDirectory_DomainService

Severity	Error
Message	Failed to remove cached secrets of the following security principal from local domain controller: %1 Error: %2 (%3) After cloning a read-only domain controller, secrets which were previously cached on the cloning source read-only domain controller need to be removed on the clone in order to decrease the risk that an attacker can obtain those credentials from stolen or compromised clone. If the security principal is a highly privileged account and should be protected against this, please use rootDSE operation rODCPurgeAccount to manually clear its secrets on local domain controller.
Notes and resolution	Examine the System and Directory Services event logs for further information.

Event ID	2227
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Exception is raised while trying to remove cached secrets from local domain controller. Additional data: Exception value: %1 Error value: %2 DSID: %3 After cloning a read-only domain controller, secrets which were previously cached on the cloning source read-only domain controller need to be removed on the clone in order to decrease the risk that an attacker can obtain those credentials from stolen or compromised clone. If any of these security principals is a highly privileged account and should be

	protected against this, please use rootDSE operation rODCPurgeAccount to manually clear its secrets on local domain controller.
Notes and resolution	Examine the System and Directory Services event logs for further information.

Event ID	2228
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	The Virtual machine generation ID in the Active Directory database of this domain controller is different from the current value of this virtual machine. However, a virtual domain controller clone configuration file (DCCloneConfig.xml) could not be located so domain controller cloning was not attempted. If a domain controller cloning operation was intended, please ensure that a DCCloneConfig.xml is provided in any one of the supported locations. In addition, the IP address of this domain controller conflicts with another domain controller's IP address. To ensure no disruptions in service occur, the domain controller has been configured to boot into DSRM. Additional data: The duplicate IP address: %1
Notes and resolution	This protection mechanism stops duplicate domain controllers when possible (it will not when using DHCP, for example). Add a valid DcCloneConfig.xml file, remove the DSRM flag, and re-attempt cloning

Event ID	29218
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server

Severity	Error
Message	Virtual domain controller cloning failed. The cloning operation could not be completed and the cloned domain controller was rebooted into Directory Services Restore Mode (DSRM). Please check previously logged events and %systemroot%\debug\dcpromo.log for more information on errors that correspond to the virtual domain controller cloning attempt and whether or not this clone image can be reused. If one or more log entries indicate that the cloning process cannot be retried, the image must be securely destroyed. Otherwise you may fix the errors, clear the DSRM boot flag, and reboot normally; upon reboot, the cloning operation will be retried.
Notes and resolution	Review the System and Directory Services event logs and the dcpromo.log for further details on why cloning failed.

Event ID	29219
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Informational
Message	Virtual domain controller cloning succeeded.
Notes and resolution	This is a success event and only an issue if unexpected.

Event ID	29248
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed to obtain Winlogon Notification. The returned error code is %1 (%2). For more information on this error, please

	review %systemroot%\debug\dcpromo.log for errors that correspond to the virtual domain controller cloning attempt.
Notes and resolution	Contact Microsoft Product Support

Event ID	29249
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed to parse virtual domain controller configuration file. The returned HRESULT code is %1. The configuration file is:%2 Please fix the errors in the configuration file and retry the cloning operation. For more information about this error, please see %systemroot%\debug\dcpromo.log.
Notes and resolution	Examine the dcloneconfig.xml file for syntax errors using an XML editor and the DCCloneConfigSchema.xsd schema file.

Event ID	29250
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed. There are software or services currently enabled on the cloned virtual domain controller that are not present in the allowed application list for virtual domain controller cloning. Following are the missing entries: %2 %1 (if any) was used as the defined inclusion list. The cloning operation cannot be completed if there are non-cloneable applications installed. Please run Active Directory PowerShell Cmdlet Get-ADDCCloningExcludedApplicationList to check which applications are installed on the cloned machine, but not included in the allow list, and add them to the allow list if they

	are compatible with virtual domain controller cloning. If any of these applications are not compatible with virtual domain controller cloning, please uninstall them before re-trying the cloning operation. The virtual domain controller cloning process searches for the allowed application list file, CustomDCCloneAllowList.xml, based on the following search order; the first file found is used and all others are ignored: 1. The registry value name: HKey_Local_Machine\System\CurrentControlSet\Services\NTDS\Parameters\AllowListFolder 2. The same directory where the DSA Working Directory folder resides 3. %windir%\NTDS 4. Removable read/write media in order of drive letter at the root of the drive
Notes and resolution	Follow the message instructions

Event ID	29251
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed to reset the IP addresses of the clone machine. The returned error code is %1 (%2). This error might be caused by misconfiguration in network configuration sections in the virtual domain controller configuration file. Please see %systemroot%\debug\dcpromo.log for more information about errors that correspond to IP addresses resetting during virtual domain controller cloning attempts. Details on resetting machine IP addresses on the cloned machine can be found at http://go.microsoft.com/fwlink/?LinkId=208030
Notes and resolution	Verify the IP information set in the dccloneconfig.xml is valid and does not duplicate the original source machine.

Event ID	29253
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed. The clone domain controller was unable to locate the primary domain controller (PDC) operations master in the cloned computer's home domain of the cloned machine. The returned error code is %1 (%2). Please verify that the primary domain controller in the home domain of the cloned machine is assigned to a live domain controller, is online, and is operational. Verify that the cloned machine has LDAP/RPC connectivity to the primary domain controller over the required ports and protocols.
Notes and resolution	Validate the cloned domain controller IP and DNS information is set. Use Dcdiag.exe /test:locatorcheck to validate if the PDCE is online, use Nltest.exe /server:<*PDCE*> /dclist:<*domain*> to valid RPC, obtain a network capture from the PDCE while cloning fails and analyze the traffic.

Event ID	29254
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed to bind to the primary domain controller %1. The returned error code is %2 (%3). Please verify that the primary domain controller %1 is online and is operational. Verify that the cloned machine has LDAP/RPC connectivity to the primary domain controller over the required ports and protocols.

Notes and resolution	Validate the cloned domain controller IP and DNS information is set. Use Dcdiag.exe /test:locatorcheck to validate if the PDCE is online, use Nltest.exe /server:<*PDCE*> /dclist:<*domain*> to valid RPC, obtain a network capture from the PDCE while cloning fails and analyze the traffic.

Event ID	29255
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed. An attempt to create objects on the primary domain controller %1 required for the image being cloned returned error %2 (%3). Please verify that the cloned domain controller has privilege to clone itself. Check for related events in the Directory Service event log on primary domain controller %1.
Notes and resolution	Lookup the specific error in MS TechNet, MS Knowledgebase, and MS blogs to determine its typical meaning, and then troubleshoot based on those results.

Event ID	29256
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	An attempt to set the Boot into Directory Services Restore Mode flag failed with error code %1. Please see %systemroot%\debug\dcpromo.log for more information about errors.
Notes and resolution	Examine the Directory Services log and dcpromo.log for details. Examine Application

	and System event logs. Investigate third party application that may be blocking privilege usage.

Event ID	29257
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning has done. An attempt to reboot the machine failed with error code %1. Please reboot the machine to finish the cloning operation.
Notes and resolution	Examine Application and System event logs. Investigate third party application that may be blocking privilege usage.

Event ID	29264
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	An attempt to clear the Boot into Directory Services Restore Mode flag failed with error code %1. Please see %systemroot%\debug\dcpromo.log for more information about errors.
Notes and resolution	Examine the Directory Services log and dcpromo.log for details. Examine Application and System event logs. Investigate third party application that may be blocking privilege usage.

Event ID	29265
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server

Severity	Informational
Message	Virtual domain controller cloning succeeded. The virtual domain controller cloning configuration file %1 has been renamed to %2.
Notes and resolution	N/A, this is a success event.

Event ID	29266
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning succeeded. The attempt to rename virtual domain controller cloning configuration file %1 failed with error code %2 (%3).
Notes and resolution	Manually rename the dccloneconfig.xml file.

Event ID	29267
Source	Microsoft-Windows-DirectoryServices-DSROLE-Server
Severity	Error
Message	Virtual domain controller cloning failed to check the virtual domain controller cloning allowed application list. The returned error code is %1 (%2). This error might be caused by a syntax error in the clone allow list file (The file currently being checked is: %3). For more information about this error, please see %systemroot%\debug\dcpromo.log.
Notes and resolution	Follow the event instructions

Error Messages

There are no direct interactive errors for failed virtualized domain controller cloning; all cloning information logs in the System and Directory Services logs and the domain controller promotion

logs in dcpromo.log. However, if the server boots into DS Restore Mode, investigate immediately, as promotion or cloning failed.

The dcpromo.log is the first place to check for cloning failure. Depending on the failure listed, it may be necessary to subsequently review Directory Services and System logs for further diagnosis.

Known Issues and Support Scenarios

The following are common issues seen during the Windows Server 2012 development process. All of these issues are "by design" and have either a valid workaround or more appropriate technique to avoid them in the first place. Some may be resolved in later releases of Windows Server 2012.

Issue	**Cloning fails, DSRM**
Symptoms	Clone boots into Directory Services Restore Mode
Resolution and Notes	Validate all steps followed from sections Deploying Virtualized Domain Controller section and General Methodology for Troubleshooting Domain Controller Cloning Described in KB 2742844.

Issue	**Extra IP leases when using DHCP to clone**
Symptoms	After successfully cloning a DC and using DHCP, the first boot of the clone takes a DHCP lease. Then when the server is renamed and restarted as a DC, it takes a second DHCP lease. The first IP address is not released and you end up with a "phantom" lease
Resolution and Notes	Manually delete the unused address lease in DHCP or allow it to expire normally. Described in KB 2742836.

Issue	**Cloning fails into DSRM after very long delay**
Symptoms	Cloning appears to pause at "Domain controller cloning is at X% completion" for between 8 and 15 minutes. After this, the cloning fails and

	boots into DSRM.
Resolution and Notes	The cloned computer cannot get a dynamic IP address from DHCP or SLAAC, or is using a duplicate IP address, or cannot find the PDC. Multiple retry attempts performed by cloning lead to the delay. Resolve the networking issue to allow cloning. Described in KB 2742844.

Issue	**Cloning does not recreate all service principal names**
Symptoms	If a set of *three-part* service principal names (SPN) includes both a NetBIOS name with a port and an otherwise identical NetBIOS name without a port, the non-port entry is not recreated with the new computer name. For example: customspn/DC1:200/app1 INVALID USE OF SYMBOLS *this is recreated with the new computer name* customspn/DC1/app1 INVALID USE OF SYMBOLS *this is not recreated with the new computer name* Fully qualified names are recreated and SPNs without three parts are recreated, regardless of ports. For example, these are recreated successfully on the clone: customspn/DC1:202 INVALID USE OF SYMBOLS *this is recreated* customspn/DC1 INVALID USE OF SYMBOLS *this is recreated* customspn/DC1.corp.contoso.com:202 INVALID USE OF SYMBOLS *this is recreated name* customspn/DC1.corp.contoso.com INVALID USE OF SYMBOLS *this is recreated*
Resolution and Notes	This is a limitation of the domain controller rename process in Windows, not just in cloning. Three-part SPNS are not handled by the

	renaming logic in any scenario. Most included Windows services are unaffected by this, as they recreate any missing SPNs as needed. Other applications may require manually entering the SPN to resolve the issue. Described in KB 2742874.

Issue	**Cloning fails, boots into DSRM, general networking errors**
Symptoms	Clone boots into Directory Services Repair Mode. There are general networking errors.
Resolution and Notes	Ensure that the new clone does not have a duplicate static MAC address assigned from the source domain controller; you can see if a VM uses static MAC addresses by running this command on the hypervisor host for both the source and clone virtual machines: Get-VM –VMName *test-vm* | Get-VMNetworkAdapter | fl * Change the MAC address to a unique static address or switch to using dynamic MAC addresses. Described in KB 2742844

Issue	**Cloning fails, boots into DSRM as a duplicate of the source DC**
Symptoms	A new clone boots up without cloning. The dccloneconfig.xml is not renamed and the server starts in DS Restore Mode. The Directory Services event log shows Error 2164 `<COMPUTERNAME> failed to start the DsRoleSvc service to clone the local virtual domain controller.`
Resolution and Notes	Examine the service settings for the DS Role Server service (DsRoleSvc) and ensure its start type is set to Manual. Validate that no third

	party program is preventing the start of this service. For more information about how to reclaim this secondary DC while ensuring that updates get replicated outbound, see Microsoft KB article 2742970.

Issue	**Cloning fails, boots into DSRM, error 8610**
Symptoms	Clone boots into Directory Services Restore Mode. Dcpromo .log shows 8610 error (which is ERROR_DS_ROLE_NOT_VERIFIED 8610 or 0x21A2)
Resolution and Notes	Will happen if the PDC can be discoverable but it has not performed sufficient replication to allow itself to assume the role. For example, if cloning is started and another administrator moves the PDCE FSMO role to a new DC. Described in KB 2742916.

Issue	**Cloning fails, boots into DSRM, general networking errors**
Symptoms	Clone boots into Directory Services Restore Mode. There are general networking errors.
Resolution and Notes	Ensure that the new clone does not have a duplicate static MAC address assigned from the source domain controller; you can see if a VM uses static MAC addresses by running this command on the Hyper-V host for both the source and clone virtual machines: `Get-VM –VMName test-vm | Get-VMNetworkAdapter | fl *` Change the MAC address to a unique static address or switch to using dynamic MAC addresses. Described in KB 2742844.

Issue	**Cloning fails, boots into DSRM**
Symptoms	Clone boots into Directory Services Repair Mode
Resolution and Notes	Ensure that the dccloneconfig.xml contains the schema definition (see sampledccloneconfig.xml, line 2): **<d3c:DCCloneConfig xmlns:d3c="uri:microsoft.com:schemas:DCCloneConfig">** Described in KB 2742844

Issue	**No logon servers are available error logging into DSRM**
Symptoms	Clone boots into Directory Services Repair Mode. You attempt to logon and receive error: **There are currently no logon servers are available to service the logon request**
Resolution and Notes	Ensure you logon with the DSRM administrator account, and not the domain account. Use the left arrow and type a user name of: **.\administrator** Described in KB 2742908

Issue	**Clone Source fails into DSRM, error**
Symptoms	During cloning, fails 8437 "Create clone DC objects on PDC failed" (0x20f5)
Resolution and Notes	Duplicate computer name was set in DCCloneConfig.xml as the source DC or an existing DC. The computer name also needs to be in the NetBIOS computer name format (15 characters or fewer, not an FQDN). Fix the dccloneconfig.xml file by setting a unique, valid name. Described in KB 2742959

Issue	**New-addccloneconfigfile error "index was out of range"**

Symptoms	When running the new-addccloneconfigfile cmdlet, you receive error: Index was out of range. Must be non-negative and less than the size of the collection.
Resolution and Notes	You must run the cmdlet in an administrator-elevated Windows PowerShell console. This error is caused by lack of local administrator group membership on the computer. Described in KB 2742927

Issue	**Cloning fails, duplicate DC**
Symptoms	Clone boots without cloning, duplicates existing source DC
Resolution and Notes	The computer was copied and started but does not contain a DcCloneConfig.xml file in any of the supported locations, and did not have a duplicate IP address with the source domain controller. The DC must be correctly removed in order to avoid data loss. Described in KB 2742970

Issue	**New-ADDCCloneConfigFile fails with The server is not operational error when it checks if the source domain controller is a member of the Cloneable Domain controllers group if a GC is not available.**
Symptoms	When running New-ADDCCloneConfigFile to create a dccloneconfig.xml file, you receive error: `The server is not operational`
Resolution and Notes	Verify connectivity to a GC from the server where you run New-ADDCCloneConfigFile and verify that the membership of the source domain controller in the Cloneable Domain Controllers group has replicated to that GC. Run the following command as a means of flushing the DC locator cache for cases where

	a GC or DC may have been taken offline recently: `nltest /dsgetdc: /GC /FORCE`

Advanced Troubleshooting

This module seeks to teach advanced troubleshooting by using *working* logs as samples, with some explanation of what occurred. If you understand what a successful virtualized domain controller operation looks like, failures become obvious in your environment. These logs are presented by their source, with the ascending order of *expected* events (even when they are warnings and errors) related to a cloned domain controller within each log.

Cloning a Domain Controller

In this example, the clone domain controller uses DHCP to get an IP address, replicates SYSVOL using FRS or DFSR (see the appropriate log as necessary), is a global catalog, and uses a blank dccloneconfig.xml file.

Directory Services Event Log

The Directory Services log contains the majority of event-based cloning operational information. The hypervisor changes the VM-Generation ID and the NTDS service notes it, then invalidates the RID pool and changes the invocation ID. The new VM-Generation ID is set and the server replicates Active Directory data inbound. The DFSR service is stopped and its database that hosts SYSVOL is deleted, forcing a non-authoritative sync inbound. The USN high watermark is adjusted.

Event ID	Source	Message
2160	ActiveDirectory_DomainService	The local Active Directory Domain Services has found a virtual domain controller cloning configuration file. The virtual domain controller cloning configuration file is found at: <*path*>\DCCloneConfig.xml The existence of the virtual domain controller cloning configuration file indicates that the local virtual domain controller is a clone of another virtual domain controller. The Active Directory Domain Services will start to clone itself.
2191	ActiveDirectory_DomainService	Active Directory Domain Services set the following registry value to disable DNS updates. Registry Key:

		SYSTEM\CurrentControlSet\Services\Netlogon\Parameters Registry Value: UseDynamicDns Registry Value data: 0 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning. The cloning process will enable DNS updates again after cloning is completed.
2191	ActiveDirectory_DomainService	Active Directory Domain Services set the following registry value to disable DNS updates. Registry Key: SYSTEM\CurrentControlSet\Services\Dnscache\Parameters Registry Value: RegistrationEnabled Registry Value data: 0 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning. The cloning process will enable DNS updates again after cloning is completed. “Information 2/7/2012 3:12:49 PM Microsoft-Windows-ActiveDirectory_DomainService 2191 Internal Configuration” Active Directory Domain Services set the following registry value to disable DNS updates. Registry Key: SYSTEM\CurrentControlSet\Services\Tcpip\Parameters Registry Value: DisableDynamicUpdate Registry Value data:

		1 During the cloning process, the local machine may have the same computer name as the clone source machine for a short time. DNS A and AAAA record registration are disabled during this period so clients cannot send requests to the local machine undergoing cloning. The cloning process will enable DNS updates again after cloning is completed.
2172	ActiveDirectory_DomainService	Read the msDS-GenerationId attribute of the Domain Controller's computer object. msDS-GenerationId attribute value: *<Number>*
2170	ActiveDirectory_DomainService	A Generation ID change has been detected. Generation ID cached in DS (old value): *<Number>* Generation ID currently in VM (new value): *<Number>* The Generation ID change occurs after the application of a virtual machine snapshot, after a virtual machine import operation or after a live migration operation. Active Directory Domain Services will create a new invocation ID to recover the domain controller. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services aware backup application.
1109	ActiveDirectory_DomainService	The invocationID attribute for this directory server has been changed. The highest update sequence number at the time the backup was created is as follows: InvocationID attribute (old value): *<GUID>* InvocationID attribute (new value): *<GUID>* Update sequence number: *<Number>* The invocationID is changed when a directory server is

		restored from backup media, is configured to host a writeable application directory partition, has been resumed after a virtual machine snapshot has been applied, after a virtual machine import operation, or after a live migration operation. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services-aware backup application.
1000	ActiveDirectory_DomainService	Microsoft Active Directory Domain Services startup complete.
1394	ActiveDirectory_DomainService	All problems preventing updates to the Active Directory Domain Services database have been cleared. New updates to the Active Directory Domain Services database are succeeding. The Net Logon service has restarted
2163	ActiveDirectory_DomainService	DsRoleSvc service was started to clone the local virtual domain controller.
326	NTDS ISAM	NTDS (536) NTDSA: The database engine attached a database (1, C:\Windows\NTDS\ntds.dit). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.016, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.000, [12] 0.000. Saved Cache: 1
103	NTDS ISAM	NTDS (536) NTDSA: The database engine stopped the instance (0). Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.032, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.031, [10] 0.000, [11] 0.000, [12] 0.000, [13] 0.000, [14] 0.000, [15] 0.000.
102	NTDS ISAM	NTDS (536) NTDSA: The database engine (6.02.8225.0000) is starting a new instance (0).
105	NTDS ISAM	NTDS (536) NTDSA: The database engine started a new instance (0). (Time=0 seconds) Internal Timing Sequence: [1] 0.016, [2] 0.000, [3] 0.015, [4] 0.078, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9]

		0.046, [10] 0.000, [11] 0.000.
1004	ActiveDirectory_DomainService	Active Directory Domain Services was shut down successfully.
102	NTDS ISAM	NTDS (536) NTDSA: The database engine (6.02.8225.0000) is starting a new instance (0).
326	NTDS ISAM	NTDS (536) NTDSA: The database engine attached a database (1, C:\Windows\NTDS\ntds.dit). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.015, [3] 0.016, [4] 0.000, [5] 0.031, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.000, [12] 0.000. Saved Cache: 1
105	NTDS ISAM	NTDS (536) NTDSA: The database engine started a new instance (0). (Time=1 seconds) Internal Timing Sequence: [1] 0.031, [2] 0.000, [3] 0.000, [4] 0.391, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.031, [10] 0.000, [11] 0.000.
1109	ActiveDirectory_DomainService	The invocationID attribute for this directory server has been changed. The highest update sequence number at the time the backup was created is as follows: InvocationID attribute (old value): *<GUID>* InvocationID attribute (new value): *<GUID>* Update sequence number: *<Number>* The invocationID is changed when a directory server is restored from backup media, is configured to host a writeable application directory partition, has been resumed after a virtual machine snapshot has been applied, after a virtual machine import operation, or after a live migration operation. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services-aware backup application.
1168	ActiveDirectory_DomainServi	Internal error: An Active Directory Domain Services error

	ce	has occurred. Additional Data Error value (decimal): 2 Error value (hexadecimal): 2 Internal ID: 7011658
1110	ActiveDirectory_DomainService	Promotion of this domain controller to a global catalog will be delayed for the following interval. Interval (minutes): 5 This delay is necessary so that the required directory partitions can be prepared before the global catalog is advertised. In the registry, you can specify the number of seconds that the directory system agent will wait before promoting the local domain controller to a global catalog. For more information about the Global Catalog Delay Advertisement registry value, see the Resource Kit Distributed Systems Guide
103	NTDS ISAM	NTDS (536) NTDSA: The database engine stopped the instance (0). Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.047, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.016, [10] 0.000, [11] 0.000, [12] 0.000, [13] 0.000, [14] 0.000, [15] 0.000.
1004	ActiveDirectory_DomainService	Active Directory Domain Services was shut down successfully.
1539	ActiveDirectory_DomainService	Active Directory Domain Services could not disable the software-based disk write cache on the following hard disk. Hard disk: c: Data might be lost during system failures
2179	ActiveDirectory_DomainService	The msDS-GenerationId attribute of the Domain Controller's computer object has been set to the following

		parameter: GenerationID attribute: <*Number*>
2173	ActiveDirectory_DomainService	Failed to read the msDS-GenerationId attribute of the Domain Controller's computer object. This may be caused by database transaction failure, or the generation id does not exist in the local database. The msDS-GenerationId does not exist during the first reboot after dcpromo or the DC is not a virtual domain controller. Additional Data Failure code: 6
1000	ActiveDirectory_DomainService	Microsoft Active Directory Domain Services startup complete, version 6.2.8225.0
1394	ActiveDirectory_DomainService	All problems preventing updates to the Active Directory Domain Services database have been cleared. New updates to the Active Directory Domain Services database are succeeding. The Net Logon service has restarted.
1128	ActiveDirectory_DomainService	1128 Knowledge Consistency Checker "A replication connection was created from the following source directory service to the local directory service. Source directory service: CN=NTDS Settings,<*Domain Controller DN*> Local directory service: CN=NTDS Settings, <*Domain Controller DN*> Additional Data Reason Code: 0x2 Creation Point Internal ID: f0a025d
1999	ActiveDirectory_DomainService	The source directory service has optimized the update sequence number (USN) presented by the destination directory service. The source and destination directory services have a common replication partner. The destination directory service is up to date with the common replication partner, and the source directory

		service was installed using a backup of this partner. Destination directory service ID: *<GUID> (<FQDN>)* Common directory service ID: *<GUID>* Common property USN: *<Number>* As a result, the up-to-dateness vector of the destination directory service has been configured with the following settings. Previous object USN: 0 Previous property USN: 0 Database GUID: *<GUID>* Object USN: *<Number>* Property USN: *<Number>*

System Event Log

The next indications of cloning operations are in the System Event log. As the hypervisor tells the guest computer that it was cloned or restored from a snapshot, the domain controller immediately invalidates its RID pool to avoid duplicating security principals later. As cloning proceeds, various expected operations and messages appear, mostly around services starting and stopping and some expected errors caused by this. When completed the System event log notes overall cloning success.

Event ID	Source	Message
16654	Directory-Services-SAM	A pool of account-identifiers (RIDs) has been invalidated. This may occur in the following expected cases: 1. A domain controller is restored from backup. 2. A domain controller running on a virtual machine is restored from snapshot. 3. An administrator has manually invalidated the

		pool
7036	Service Control Manager	The Active Directory Domain Services service entered the running state.
7036	Service Control Manager	The Kerberos Key Distribution Center service entered the running state.
3096	Netlogon	The primary Domain Controller for this domain could not be located.
7036	Service Control Manager	The Security Accounts Manager service entered the running state.
7036	Service Control Manager	The Server service entered the running state.
7036	Service Control Manager	The Netlogon service entered the running state.
7036	Service Control Manager	The Active Directory Web Services service entered the running state.
7036	Service Control Manager	The DFS Replication service entered the running state.
7036	Service Control Manager	The File Replication Service service entered the running state.
14533	Microsoft-Windows-DfsSvc	DFS has finished building all namespaces.
14531	Microsoft-Windows-DfsSvc	DFS server has finished initializing.
7036	Service Control Manager	The DFS Namespace service entered the running state.
7023	Service Control Manager	The Intersite Messaging service terminated with the following error: The specified server cannot perform the requested operation.
7036	Service Control Manager	The Intersite Messaging service entered the stopped state.
5806	Netlogon	Dynamic updates have been manually disabled on this domain controller. USER ACTION Reconfigure this domain controller to use dynamic updates or manually add the DNS records from the file

		'%SystemRoot%\System32\Config\Netlogon.dns' to the DNS database."
16651	Directory-Services-SAM	The request for a new account-identifier pool failed. The operation will be retried until the request succeeds. The error is The requested FSMO operation failed. The current FSMO holder could not be contacted.
7036	Service Control Manager	The DNS Server service entered the running state.
7036	Service Control Manager	The DS Role Server service entered the running state.
7036	Service Control Manager	The Netlogon service entered the stopped state.
7036	Service Control Manager	The File Replication Service service entered the stopped state.
7036	Service Control Manager	The Kerberos Key Distribution Center service entered the stopped state.
7036	Service Control Manager	The DNS Server service entered the stopped state.
7036	Service Control Manager	The Active Directory Domain Services service entered the stopped state.
7036	Service Control Manager	The Netlogon service entered the running state.
7040	Service Control Manager	The start type of the Active Directory Domain Services service was changed from auto start to disabled.
7036	Service Control Manager	The Netlogon service entered the stopped state.
7036	Service Control Manager	The File Replication Service service entered the running state.
29219	DirectoryServices-DSROLE-Server	Virtual domain controller cloning succeeded.
29223	DirectoryServices-DSROLE-Server	This server is now a Domain Controller.
29265	DirectoryServices-DSROLE-Server	Virtual domain controller cloning succeeded. The virtual domain controller cloning configuration file C:\Windows\NTDS\DCCloneConfig.xml has been renamed to C:\Windows\NTDS\DCCloneConfig.20120207-

		151533.xml.
1074	User32	The process C:\Windows\system32\lsass.exe (DC2) has initiated the restart of computer DC2 on behalf of user NT AUTHORITY\SYSTEM for the following reason: Operating System: Reconfiguration (Planned) Reason Code: 0x80020004 Shutdown Type: restart Comment: "

DCPROMO.LOG

The Dcpromo.log contains the actual promotion portion of cloning that the Directory Services event log does not describe. Since the log does not provide the level of explanation that the event log entries impart, this section of the module contains additional annotation.

The promotion process means that the cloning starts, the DC is scrubbed of its current configuration and re-promoted using the existing AD database (much like an IFM promotion), then the DC replicates inbound change deltas of AD and SYSVOL, and cloning is complete.

Note
The log has been modified in this module for readability, by removing the date column.

Notes
For further explanation of the dcpromo.log see the Understand and Troubleshoot AD DS Simplified Administration in Windows Server 2012.

http://go.microsoft.com/fwlink/p/?LinkId=237244

- Start clone-based promotion
- Set the Directory Services Restore Mode flag so that the server does not boot back up normally as the original clone and cause naming or Directory Service collisions
- Update the Directory Services event log

```
15:14:01 [INFO] vDC Cloneing: Setting Boot into DSRM flag succeeded.

15:14:01 [WARNING] Cannot get user Token for Format Message: 17251

15:14:01 [INFO] vDC Cloning: Created vDCCloningUpdate event.

15:14:01 [INFO] vDC Cloning: Created vDCCloningComplete event.
```

- Stop the NetLogon service so that the domain controller does not advertise

```
15:14:01 [INFO] Stopping service NETLOGON

15:14:01 [INFO] ControlService(STOP) on NETLOGON returned 1(gle=0)

15:14:01 [INFO] DsRolepWaitForService: waiting for NETLOGON to enter one of 7 states

15:14:01 [INFO] DsRolepWaitForService: QueryServiceStatus on NETLOGON returned 1 (gle=0),
SvcStatus.dwCS=3
```

```
15:14:02 [INFO] DsRolepWaitForService: QueryServiceStatus on NETLOGON returned 1 (gle=0),
SvcStatus.dwCS=1

15:14:02 [INFO] DsRolepWaitForService: exiting because NETLOGON entered STOPPED state

15:14:02 [INFO] DsRolepWaitForService(for any end state) on NETLOGON service returned 0

15:14:02 [INFO] ControlService(STOP) on NETLOGON returned 0(gle=1062)

15:14:02 [INFO] Exiting service-stop loop after service NETLOGON entered STOPPED state

15:14:02 [INFO] StopService on NETLOGON returned 0

15:14:02 [INFO] Configuring service NETLOGON to 1 returned 0

15:14:02 [INFO] Updating service status to 4

15:14:02 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Examine the dccloneconfig.xml file for administrator-specified customizations.
- In this sample case it is a blank file, so all settings are automatically generated and automatic IP addressing is required from the network

```
15:14:02 [INFO] vDC Cloning: Clone config file C:\Windows\NTDS\DCCloneConfig.xml is
considered to be a blank file (containing 0 bytes)

15:14:02 [INFO] vDC Cloning: Parsing clone config file C:\Windows\NTDS\DCCloneConfig.xml
returned HRESULT 0x0
```

- Validate that there are no services or programs installed that are not part of the DefaultDCCloneAllowList.xml or CustomDCCloneAllowList.xml

```
15:14:02 [INFO] vDC Cloning: Checking allowed list:

15:14:03 [INFO] vDC Cloning: Completed checking allowed list:

15:14:03 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Enable DHCP on the network adapters, since IP information was not specified by the administrator

```
15:14:03 [INFO] vDC Cloning: Enable DHCP:

15:14:03 [INFO] WMI Instance: Win32_NetworkAdapterConfiguration.Index=12

15:14:03 [INFO] Method: EnableDHCP

15:14:03 [INFO] HRESULT code: 0x0 (0)

15:14:03 [INFO] Return Value: 0x0 (0)

15:14:03 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:03 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Locate the PDC emulator
- Set the clone's site (automatically generated in this case)
- Set the clone's name (automatically generated in this case)

```
15:14:03 [INFO] vDC Cloning: Found PDC. Name: DC1.root.fabrikam.com

15:14:04 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

```
15:14:04 [INFO] vDC Cloning: Winlogon UI Notification #1: Domain Controller cloning is at
5% completion...

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #2: Domain Controller cloning is at
10% completion...

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:05 [INFO] Site of the cloned DC: Default-First-Site-Name
```

- Create the new clone computer object
- Rename the clone to match the new name

```
15:14:05 [INFO] vDC Cloning: Clone DC objects are created on PDC.

15:14:05 [INFO] Name of the cloned DC: DC2-CL0001

15:14:05 [INFO] DsRolepSetRegStringValue on
System\CurrentControlSet\Services\NTDS\Parameters\CloneMachineName to DC2-CL0001 returned
0

15:14:05 [INFO] vDC Cloning: Save CloneMachineName in registry: 0x0 (0)
```

- Provide the promotion settings, based on previous dccloneconfig.xml or automatic generation rules

```
15:14:05 [INFO] vDC Cloning: Promotion parameters setting:

15:14:05 [INFO] DNS Domain Name: root.fabrikam.com

15:14:05 [INFO] Replica Partner: \\DC1.root.fabrikam.com

15:14:05 [INFO] Site Name: Default-First-Site-Name

15:14:05 [INFO] DS Database Path: C:\Windows\NTDS

15:14:05 [INFO] DS Log Path: C:\Windows\NTDS

15:14:05 [INFO] SysVol Root Path: C:\Windows\SYSVOL

15:14:05 [INFO] Account: root.fabrikam.com\DC2-CL0001$

15:14:05 [INFO] Options: DSROLE_DC_CLONING (0x800400)
```

- Start promotion

```
15:14:05 [INFO] Promote DC as a clone

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #3: Domain Controller cloning is at
15% completion...

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #4: Domain Controller cloning is at
16% completion...

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:05 [INFO] Validate supplied paths

15:14:05 [INFO] Validating path C:\Windows\NTDS.
```

```
15:14:05 [INFO] Path is a directory

15:14:05 [INFO] Path is on a fixed disk drive.

15:14:05 [INFO] Validating path C:\Windows\NTDS.

15:14:05 [INFO] Path is a directory

15:14:05 [INFO] Path is on a fixed disk drive.

15:14:05 [INFO] Validating path C:\Windows\SYSVOL.

15:14:05 [INFO] Path is on a fixed disk drive.

15:14:05 [INFO] Path is on an NTFS volume

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #5: Domain Controller cloning is at
17% completion...

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:05 [INFO] Start the worker task

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #6: Domain Controller cloning is at
20% completion...

15:14:05 [INFO] Request for promotion returning 0

15:14:05 [INFO] vDC Cloning: Winlogon UI Notification #7: Domain Controller cloning is at
21% completion...

15:14:05 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Stop and configure all of the AD DS-related services (NTDS, NTFRS/DFSR, KDC, DNS)

Note

The DNS service taking a long time to shutdown is expected in this scenario, as it is using AD-integrated zones that were no longer available even before the NTDS service stopped - see the DNS events described later in this section of the module.

```
15:14:15 [INFO] Stopping service NTDS

15:14:15 [INFO] Stopping service NtFrs

15:14:15 [INFO] ControlService(STOP) on NtFrs returned 1(gle=0)

15:14:15 [INFO] DsRolepWaitForService: waiting for NtFrs to enter one of 7 states

15:14:15 [INFO] DsRolepWaitForService: QueryServiceStatus on NtFrs returned 1 (gle=0),
SvcStatus.dwCS=3

15:14:16 [INFO] DsRolepWaitForService: QueryServiceStatus on NtFrs returned 1 (gle=0),
SvcStatus.dwCS=1

15:14:16 [INFO] DsRolepWaitForService: exiting because NtFrs entered STOPPED state

15:14:16 [INFO] DsRolepWaitForService(for any end state) on NtFrs service returned 0

15:14:16 [INFO] ControlService(STOP) on NtFrs returned 0(gle=1062)
```

15:14:16 [INFO] Exiting service-stop loop after service NtFrs entered STOPPED state

15:14:16 [INFO] StopService on NtFrs returned 0

15:14:16 [INFO] Configuring service NtFrs to 1 returned 0

15:14:16 [INFO] Stopping service Kdc

15:14:16 [INFO] ControlService(STOP) on Kdc returned 1(gle=0)

15:14:16 [INFO] DsRolepWaitForService: waiting for Kdc to enter one of 7 states

15:14:16 [INFO] DsRolepWaitForService: QueryServiceStatus on Kdc returned 1 (gle=0), SvcStatus.dwCS=3

15:14:17 [INFO] DsRolepWaitForService: QueryServiceStatus on Kdc returned 1 (gle=0), SvcStatus.dwCS=1

15:14:17 [INFO] DsRolepWaitForService: exiting because Kdc entered STOPPED state

15:14:17 [INFO] DsRolepWaitForService(for any end state) on Kdc service returned 0

15:14:17 [INFO] ControlService(STOP) on Kdc returned 0(gle=1062)

15:14:17 [INFO] Exiting service-stop loop after service Kdc entered STOPPED state

15:14:17 [INFO] StopService on Kdc returned 0

15:14:17 [INFO] Configuring service Kdc to 1 returned 0

15:14:17 [INFO] Stopping service DNS

15:14:17 [INFO] ControlService(STOP) on DNS returned 1(gle=0)

15:14:17 [INFO] DsRolepWaitForService: waiting for DNS to enter one of 7 states

15:14:17 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:18 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:19 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:20 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:21 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:22 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:23 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:24 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:25 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:26 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:27 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:28 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:29 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:30 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:31 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:32 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:33 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:34 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:35 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:36 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:37 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:38 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:39 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:40 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:41 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

15:14:42 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0), SvcStatus.dwCS=3

```
15:14:43 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:44 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:45 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:46 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:47 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:48 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:49 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:50 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:51 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:52 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:53 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:54 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:55 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:56 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:57 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:58 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:14:59 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=3
15:15:00 [INFO] DsRolepWaitForService: QueryServiceStatus on DNS returned 1 (gle=0),
SvcStatus.dwCS=1
15:15:00 [INFO] DsRolepWaitForService: exiting because DNS entered STOPPED state
```

```
15:15:00 [INFO] DsRolepWaitForService(for any end state) on DNS service returned 0
15:15:00 [INFO] ControlService(STOP) on DNS returned 0(gle=1062)
15:15:00 [INFO] Exiting service-stop loop after service DNS entered STOPPED state
15:15:00 [INFO] StopService on DNS returned 0
15:15:00 [INFO] Configuring service DNS to 1 returned 0
15:15:00 [INFO] ControlService(STOP) on NTDS returned 1(gle=1062)
15:15:00 [INFO] DsRolepWaitForService: waiting for NTDS to enter one of 7 states
15:15:00 [INFO] DsRolepWaitForService: QueryServiceStatus on NTDS returned 1 (gle=0),
SvcStatus.dwCS=3
15:15:01 [INFO] DsRolepWaitForService: QueryServiceStatus on NTDS returned 1 (gle=0),
SvcStatus.dwCS=1
15:15:01 [INFO] DsRolepWaitForService: exiting because NTDS entered STOPPED state
15:15:01 [INFO] DsRolepWaitForService(for any end state) on NTDS service returned 0
15:15:01 [INFO] ControlService(STOP) on NTDS returned 0(gle=1062)
15:15:01 [INFO] Exiting service-stop loop after service NTDS entered STOPPED state
15:15:01 [INFO] StopService on NTDS returned 0
15:15:01 [INFO] Configuring service NTDS to 1 returned 0
15:15:01 [INFO] Configuring service NTDS
15:15:01 [INFO] Configuring service NTDS to 64 returned 0
15:15:01 [INFO] vDC Cloning: Winlogon UI Notification #8: Domain Controller cloning is at
22% completion...
15:15:01 [INFO] vDC Cloning: Set vDCCloningUpdate event.
15:15:01 [INFO] vDC Cloning: Winlogon UI Notification #9: Domain Controller cloning is at
25% completion...
15:15:01 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Force NT5DS (NTP) time synchronization with another domain controller (typically the PDCE)

```
15:15:02 [INFO] Forcing time sync
```

- Contact a domain controller that holds the source domain controller account of the clone
- Flush any existing Kerberos tickets

```
15:15:02 [INFO] Searching for a domain controller for the domain root.fabrikam.com that
contains the account DC2$
15:15:02 [INFO] Located domain controller DC1.root.fabrikam.com for domain
root.fabrikam.com
15:15:02 [INFO] vDC Cloning: Winlogon UI Notification #10: Domain Controller cloning is
at 26% completion...
```

```
15:15:02 [INFO] vDC Cloning: Set vDCCloningUpdate event.
15:15:02 [INFO] Directing kerberos authentication to DC1.root.fabrikam.com returns 0
15:15:02 [INFO] DsRolepFlushKerberosTicketCache() successfully flushed the Kerberos
ticket cache
15:15:02 [INFO] vDC Cloning: Winlogon UI Notification #11: Domain Controller cloning is
at 27% completion...
15:15:02 [INFO] vDC Cloning: Set vDCCloningUpdate event.
15:15:02 [INFO] Using site Default-First-Site-Name for server \\DC1.root.fabrikam.com
15:15:02 [INFO] vDC Cloning: Set vDCCloningUpdate event.
15:15:02 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Stop the NetLogon service and set its start type

```
15:15:02 [INFO] Stopping service NETLOGON
15:15:02 [INFO] Stopping service NETLOGON
15:15:02 [INFO] vDC Cloning: Winlogon UI Notification #12: Domain Controller cloning is
at 29% completion...
15:15:02 [INFO] ControlService(STOP) on NETLOGON returned 1(gle=0)
15:15:02 [INFO] DsRolepWaitForService: waiting for NETLOGON to enter one of 7 states
15:15:02 [INFO] DsRolepWaitForService: QueryServiceStatus on NETLOGON returned 1 (gle=0),
SvcStatus.dwCS=3
15:15:03 [INFO] DsRolepWaitForService: QueryServiceStatus on NETLOGON returned 1 (gle=0),
SvcStatus.dwCS=1
15:15:03 [INFO] DsRolepWaitForService: exiting because NETLOGON entered STOPPED state
15:15:03 [INFO] DsRolepWaitForService(for any end state) on NETLOGON service returned 0
15:15:03 [INFO] ControlService(STOP) on NETLOGON returned 0(gle=1062)
15:15:03 [INFO] Exiting service-stop loop after service NETLOGON entered STOPPED state
15:15:03 [INFO] StopService on NETLOGON returned 0
15:15:03 [INFO] Configuring service NETLOGON to 1 returned 0
15:15:03 [INFO] Stopped NETLOGON
15:15:03 [INFO] vDC Cloning: Set vDCCloningUpdate event.
15:15:03 [INFO] vDC Cloning: Winlogon UI Notification #13: Domain Controller cloning is
at 30% completion...
```

- Configure the DFSR/NTFRS services to run automatically
- Delete their existing database files to force non-authoritative sync of SYSVOL when the service next starts

```
15:15:03 [INFO] Configuring service DFSR
```

```
15:15:03 [INFO] Configuring service DFSR to 256 returned 0

15:15:03 [INFO] Configuring service NTFRS

15:15:03 [INFO] Configuring service NTFRS to 256 returned 0

15:15:03 [INFO] Removing DFSR Database files for SysVol

15:15:03 [INFO] Removing FRS Database files in C:\Windows\ntfrs\jet

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\log\edb.log

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\log\edbres00001.jrs

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\log\edbres00002.jrs

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\log\edbtmp.log

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\ntfrs.jdb

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\sys\edb.chk

15:15:03 [INFO] Removed C:\Windows\ntfrs\jet\temp\tmp.edb

15:15:04 [INFO] Created system volume path

15:15:04 [INFO] Configuring service DFSR

15:15:04 [INFO] Configuring service DFSR to 128 returned 0

15:15:04 [INFO] Configuring service NTFRS

15:15:04 [INFO] Configuring service NTFRS to 128 returned 0

15:15:04 [INFO] vDC Cloning: Winlogon UI Notification #14: Domain Controller cloning is
at 40% completion...

15:15:04 [INFO] vDC Cloning: Set vDCCloningUpdate event.
```

- Start the promotion process using the existing NTDS database file
- Contact the RID Master

Note

The AD DS service is not actually installed here, this is legacy instrumentation in the log

```
15:15:04 [INFO] Installing the Directory Service

15:15:04 [INFO] Calling NtdsInstall for root.fabrikam.com

15:15:04 [INFO] Starting Active Directory Domain Services installation

15:15:04 [INFO] Validating user supplied options

15:15:04 [INFO] Determining a site in which to install

15:15:04 [INFO] Examining an existing forest...

15:15:04 [INFO] Starting a replication cycle between DC1.root.fabrikam.com and the RID
operations master (2008r2-01.root.fabrikam.com), so that the new replica will be able to
create users, groups, and computer objects...

15:15:04 [INFO] Configuring the local computer to host Active Directory Domain Services
```

```
15:15:04 [INFO] EVENTLOG (Warning): NTDS General / Service Control : 1539
Active Directory Domain Services could not disable the software-based disk write cache on
the following hard disk.
Hard disk:
c:
Data might be lost during system failures.
15:15:10 [INFO] EVENTLOG (Informational): NTDS General / Internal Processing : 2041
Duplicate event log entries were suppressed.
See the previous event log entry for details. An entry is considered a duplicate if
the event code and all of its insertion parameters are identical. The time period for
this run of duplicates is from the time of the previous event to the time of this event.
Event Code:
80000603
Number of duplicate entries:
2
15:15:10 [INFO] EVENTLOG (Informational): NTDS General / Internal Configuration : 2121
This Active Directory Domain Services server is disabling the Recycle Bin. Deleted
objects may not be undeleted at this time.
```

- Change the existing invocation ID that existed in the source computers database
- Create a new NTDS Settings object for this clone
- Replicate in AD object delta from the partner domain controller

Note

Even though all objects are listed as replicated, this is just metadata needed to subsume the updates. All the unchanged objects in the cloned NTDS database already exist and do not require replication again, just like using IFM-based promotion.

```
15:15:10 [INFO] EVENTLOG (Informational): NTDS Replication / Replication : 1109
The invocationID attribute for this directory server has been changed. The highest update
sequence number at the time the backup was created is as follows:
InvocationID attribute (old value):
24e7b22f-4706-402d-9b4f-f2690f730b40
InvocationID attribute (new value):
f74cefb2-89c2-442c-b1ba-3234b0ed62f8
Update sequence number:
20520
```

The invocationID is changed when a directory server is restored from backup media, is configured to host a writeable application directory partition, has been resumed after a virtual machine snapshot has been applied, after a virtual machine import operation, or after a live migration operation. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services-aware backup application.

15:15:10 [INFO] EVENTLOG (Error): NTDS General / Internal Processing : 1168

Internal error: An Active Directory Domain Services error has occurred.

Additional Data

Error value (decimal):

2

Error value (hexadecimal):

2

Internal ID:

7011658

15:15:11 [INFO] Creating the NTDS Settings object for this Active Directory Domain Controller on the remote AD DC DC1.root.fabrikam.com...

15:15:11 [INFO] Replicating the schema directory partition

15:15:11 [INFO] Replicated the schema container.

15:15:12 [INFO] Active Directory Domain Services updated the schema cache.

15:15:12 [INFO] Replicating the configuration directory partition

15:15:12 [INFO] Replicating data CN=Configuration,DC=root,DC=fabrikam,DC=com: Received 2612 out of approximately 2612 objects and 94 out of approximately 94 distinguished name (DN) values...

15:15:12 [INFO] Replicated the configuration container.

15:15:13 [INFO] Replicating critical domain information...

15:15:13 [INFO] Replicating data DC=root,DC=fabrikam,DC=com: Received 109 out of approximately 109 objects and 35 out of approximately 35 distinguished name (DN) values...

15:15:13 [INFO] Replicated the critical objects in the domain container.

- Populate the GC partitions as needed with any missing updates
- Complete the critical AD DS portion of the promotion

15:15:13 [INFO] EVENTLOG (Informational): NTDS General / Global Catalog : 1110

Promotion of this domain controller to a global catalog will be delayed for the following interval.

Interval (minutes):

5

This delay is necessary so that the required directory partitions can be prepared before the global catalog is advertised. In the registry, you can specify the number of seconds that the directory system agent will wait before promoting the local domain controller to a global catalog. For more information about the Global Catalog Delay Advertisement registry value, see the Resource Kit Distributed Systems Guide.

15:15:14 [INFO] EVENTLOG (Informational): NTDS General / Service Control : 1000

Microsoft Active Directory Domain Services startup complete, version 6.2.8225.0

15:15:15 [INFO] Creating new domain users, groups, and computer objects

15:15:16 [INFO] Completing Active Directory Domain Services installation

15:15:16 [INFO] NtdsInstall for root.fabrikam.com returned 0

15:15:16 [INFO] DsRolepInstallDs returned 0

15:15:16 [INFO] Installed Directory Service

- Complete the inbound replication of SYSVOL

15:15:16 [INFO] vDC Cloning: Winlogon UI Notification #15: Domain Controller cloning is at 60% completion...

15:15:16 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] Completed system volume replication

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #16: Domain Controller cloning is at 70% completion...

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] SetProductType to 2 [LanmanNT] returned 0

15:15:18 [INFO] Set the product type

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #17: Domain Controller cloning is at 71% completion...

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #18: Domain Controller cloning is at 72% completion...

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] Set the system volume path for NETLOGON

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #19: Domain Controller cloning is at 73% completion...

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] Replicating non critical information

```
15:15:18 [INFO] User specified to not replicate non-critical data

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #20: Domain Controller cloning is
at 80% completion...

15:15:18 [INFO] Stopped the DS

15:15:18 [INFO] vDC Cloning: Set vDCCloningUpdate event.

15:15:18 [INFO] vDC Cloning: Winlogon UI Notification #21: Domain Controller cloning is
at 90% completion...

15:15:18 [INFO] Configuring service NTDS

15:15:18 [INFO] Configuring service NTDS to 16 returned 0
```

- Enable client DNS registration

```
15:15:18 [INFO] vDC Cloning: Set DisableDynamicUpdate reg value to 0 to enable dynamic
update records registration.

15:15:18 [INFO] vDC Cloning: Set UseDynamicDns reg value to 1 to enable dynamic update
records registration.

15:15:18 [INFO] vDC Cloning: Set RegistrationEnabled reg value to 1 to enable dynamic
update records registration.
```

- Run the SYSPREP modules specified by the DefaultDCCloneAllowList.xml <SysprepInformation> element.

```
15:15:18 [INFO] vDC Cloning: Running sysprep providers.

15:15:32 [INFO] vDC Cloning: Completed running sysprep providers.
```

- Cloning promotion is complete
- Remove the DSRM boot flag so the server boots normally next time
- Rename the dccloneconfig.xml so that it is not read again at next bootup
- Restart the computer

```
15:15:32 [INFO] The attempted domain controller operation has completed

15:15:32 [INFO] Updating service status to 4

15:15:32 [INFO] DsRolepSetOperationDone returned 0

15:15:32 [INFO] vDC Cloning: Set vDCCloningComplete event.

15:15:32 [INFO] vDC Cloneing: Clearing Boot into DSRM flag succeeded.

15:15:32 [INFO] vDC Cloning: Winlogon UI Notification #22: Cloning Domain Controller
succeeded. Now rebooting...

15:15:33 [INFO] vDC Cloning: Renamed vDC clone configuration file.

15:15:33 [INFO] vDC Cloning: The old name is: C:\Windows\NTDS\DCCloneConfig.xml

15:15:33 [INFO] vDC Cloning: The new name is: C:\Windows\NTDS\DCCloneConfig.20120207-
151533.xml
```

```
15:15:34 [INFO] vDC Cloning: Release Ipv4 on interface 'Wired Ethernet Connection 2',
result=0.
15:15:34 [INFO] vDC Cloning: Release Ipv6 on interface 'Wired Ethernet Connection 2',
result=0.
15:15:34 [INFO] Rebooting machine
```

Active Directory Web Services Event Log

While cloning is occurring, the NTDS.DIT database is often offline for extended periods. The ADWS service logs at least one event for this. After cloning is complete, the ADWS service starts, notes that there is not yet a valid computer certificate yet (there may or may not be, depending on your environment deploying a Microsoft PKI with auto-enrollment or not) and then starts the instance for the new domain controller.

Event ID	Source	Message
1202	ADWS Instance Events	This computer is now hosting the specified directory instance, but Active Directory Web Services could not service it. Active Directory Web Services will retry this operation periodically. Directory instance: NTDS Directory instance LDAP port: 389 Directory instance SSL port: 636
1000	ADWS Instance Events	Active Directory Web Services is starting
1008	ADWS Instance Events	Active Directory Web Services has successfully reduced its security privileges
1100	ADWS Instance Events	The values specified in the <appsettings> section of the configuration file for Active Directory Web Services have been loaded without errors.
1400	ADWS Instance Events	ADWS Certificate Events"Active Directory Web Services could not find a server certificate with the specified certificate name. A certificate is required to use

		SSL/TLS connections. To use SSL/TLS connections, verify that a valid server authentication certificate from a trusted Certification Authority (CA) is installed on the machine. Certificate name: <*Server FQDN*>
1100	ADWS Instance Events	The values specified in the <appsettings> section of the configuration file for Active Directory Web Services have been loaded without errors.
1200	ADWS Instance Events	Active Directory Web Services is now servicing the specified directory instance. Directory instance: NTDS Directory instance LDAP port: 389 Directory instance SSL port: 636

DNS Server Event Log

The DNS service will experience brief expected outages while cloning occurs, as the DNS service is still running while the AD DS database is offline. This occurs if using Active Directory Integrated DNS, but not if using Standard Primary or Secondary DNS. These errors log multiple times. After cloning completes, DNS comes back online normally.

Event ID	Source	Message
4013	DNS-Server-Service	The DNS server is waiting for Active Directory Domain Services (AD DS) to signal that the initial synchronization of the directory has been completed. The DNS server service cannot start until the initial synchronization is complete because critical DNS data might not yet be replicated onto this domain controller. If events in the AD DS event log indicate that there is a problem

		with DNS name resolution, consider adding the IP address of another DNS server for this domain to the DNS server list in the Internet Protocol properties of this computer. This event will be logged every two minutes until AD DS has signaled that the initial synchronization has successfully completed.
4015	DNS-Server-Service	The DNS server has encountered a critical error from the Active Directory. Check that the Active Directory is functioning properly. The extended error debug information (which may be empty) is """". The event data contains the error.
4000	DNS-Server-Service	The DNS server was unable to open Active Directory. This DNS server is configured to obtain and use information from the directory for this zone and is unable to load the zone without it. Check that the Active Directory is functioning properly and reload the zone. The event data is the error code.
4013	DNS-Server-Service	The DNS server is waiting for Active Directory Domain Services (AD DS) to signal that the initial synchronization of the directory has been completed. The DNS server service cannot start until the initial synchronization is complete because critical DNS data might not yet be replicated onto this domain controller. If events in the AD DS event log indicate that there is a problem with DNS name resolution, consider adding the IP address of another DNS server for this

		domain to the DNS server list in the Internet Protocol properties of this computer. This event will be logged every two minutes until AD DS has signaled that the initial synchronization has successfully completed.
2	DNS-Server-Service	The DNS server has started.
4	DNS-Server-Service	The DNS server has finished the background loading of zones. All zones are now available for DNS updates and zone transfers, as allowed by their individual zone configuration.

File Replication Service Event Log

The File Replication Service synchronizes non-authoritatively from a partner during cloning. Cloning accomplishes this by deleting the NTFRS database files and leaving the contents of SYSVOL untouched, for use as pre-seeded data. The two attempts to synchronize are expected.

Event ID	Source	Message
13562	NtFrs	Following is the summary of warnings and errors encountered by File Replication Service while polling the Domain Controller DC2.root.fabrikam.com for FRS replica set configuration information. Could not bind to a Domain Controller. Will try again at next polling cycle
13502	NtFrs	The File Replication Service is stopping.
13565	NtFrs	File Replication Service is initializing the system volume with data from another domain controller. Computer DC2 cannot become a domain controller until this process is complete. The system volume will then be shared as SYSVOL. To check for the SYSVOL share, at the command prompt, type: net share When File Replication Service completes the

		initialization process, the SYSVOL share will appear. The initialization of the system volume can take some time. The time is dependent on the amount of data in the system volume, the availability of other domain controllers, and the replication interval between domain controllers.
13501	NtFrs	The File Replication Service is starting
13502	NtFrs	The File Replication Service is stopping.
13503	NtFrs	The File Replication Service has stopped.
13565	NtFrs	File Replication Service is initializing the system volume with data from another domain controller. Computer DC2 cannot become a domain controller until this process is complete. The system volume will then be shared as SYSVOL. To check for the SYSVOL share, at the command prompt, type: net share When File Replication Service completes the initialization process, the SYSVOL share will appear. The initialization of the system volume can take some time. The time is dependent on the amount of data in the system volume, the availability of other domain controllers, and the replication interval between domain controllers.
13501	NtFrs	The File Replication Service is starting.
13553	NtFrs	The File Replication Service successfully added this computer to the following replica set: "DOMAIN SYSTEM VOLUME (SYSVOL SHARE)" Information related to this event is shown below: Computer DNS name is *<Domain Controller*

		FQDN> Replica set member name is <*Domain Controller*> Replica set root path is <*path*> Replica staging directory path is <*path*> Replica working directory path is <*path*>
13520	NtFrs	The File Replication Service moved the preexisting files in <path>to <*path*>\NtFrs_PreExisting___See_EventLog. The File Replication Service may delete the files in <*path*>\NtFrs_PreExisting___See_EventLog at any time. Files can be saved from deletion by copying them out of <*path*>\NtFrs_PreExisting___See_EventLog. Copying the files into c:\windows\sysvol\domain may lead to name conflicts if the files already exist on some other replicating partner. In some cases, the File Replication Service may copy a file from <*path*>\NtFrs_PreExisting___See_EventLog into <*path*> instead of replicating the file from some other replicating partner. Space can be recovered at any time by deleting the files in <*path*>\NtFrs_PreExisting___See_EventLog."
13508	NtFrs	he File Replication Service is having trouble enabling replication from \\<*Domain Controller FQDN*> to <*Domain Controller*> for <*path*> using the DNS name \\<*Domain Controller FQDN*>. FRS will keep retrying. Following are some of the reasons you would see this warning. [1] FRS cannot correctly resolve the DNS name \\<*Domain Controller FQDN*> from this computer. [2] FRS is not running on \\<*Domain Controller FQDN*>.

		[3] The topology information in the Active Directory Domain Services for this replica has not yet replicated to all the Domain Controllers. This event log message will appear once per connection, After the problem is fixed you will see another event log message indicating that the connection has been established.
13509	NtFrs	The File Replication Service has enabled replication from \\<*Domain Controller FQDN*> to <*Domain Controller*> for <*Path*> after repeated retries.
13516	NtFrs	The File Replication Service is no longer preventing the computer <*Domain Controller*> from becoming a domain controller. The system volume has been successfully initialized and the Netlogon service has been notified that the system volume is now ready to be shared as SYSVOL. Type "net share" to check for the SYSVOL share."

DFS Replication Event Log

The DFSR services synchronizes non-authoritatively from a partner during cloning. Cloning accomplishes this by deleting the DFSR database files and leaving the contents of SYSVOL untouched, for use as pre-seeded data. The two attempts to synchronize are expected.

Event ID	Source	Message
1004	DFSR	The DFS Replication service has started.
1314	DFSR	The DFS Replication service successfully configured the debug log files. Additional Information: Debug Log File Path: C:\Windows\debug
6102	DFSR	The DFS Replication service has successfully registered the WMI provider

1206	DFSR	The DFS Replication service successfully contacted domain controller DC2.corp.contoso.com to access configuration information.
1210	DFSR	The DFS Replication service successfully set up an RPC listener for incoming replication requests. Additional Information: Port: 0"
4614	DFSR	The DFS Replication service initialized SYSVOL at local path C:\Windows\SYSVOL\domain and is waiting to perform initial replication. The replicated folder will remain in the initial synchronization state until it has replicated with its partner. If the server was in the process of being promoted to a domain controller, the domain controller will not advertise and function as a domain controller until this issue is resolved. This can occur if the specified partner is also in the initial synchronization state, or if sharing violations are encountered on this server or the synchronization partner. If this event occurred during the migration of SYSVOL from File Replication Service (FRS) to DFS Replication, changes will not replicate out until this issue is resolved. This can cause the SYSVOL folder on this server to become out of sync with other domain controllers. Additional Information: Replicated Folder Name: SYSVOL Share Replicated Folder ID: <*GUID*> Replication Group Name: Domain System Volume Replication Group ID: <*GUID*>

		Member ID: <*GUID*> Read-Only: 0
4604	DFSR	The DFS Replication service successfully initialized the SYSVOL replicated folder at local path C:\Windows\SYSVOL\domain. This member has completed initial synchronization of SYSVOL with partner dc1.corp.contoso.com. To check for the presence of the SYSVOL share, open a command prompt window and then type ""net share"". Additional Information: Replicated Folder Name: SYSVOL Share Replicated Folder ID: <*GUID*> Replication Group Name: Domain System Volume Replication Group ID: <*GUID*> Member ID: <*GUID*> Sync partner: <*domain controller FQDN*>

Troubleshooting virtualized domain controller safe restore

Tools for Troubleshooting

Logging Options

The built-in logs are the most important tool for troubleshooting issues with domain controller safe snapshot restore. All of these logs are enabled and configured for maximum verbosity, by default.

Operation	Log
Snapshot creation	• Event viewer\Applications and services logs\Microsoft\Windows\Hyper-V-Worker
Snapshot restore	• Event viewer\Applications and services logs\Directory Service

	• Event viewer\Windows logs\System • Event viewer\Windows logs\Application • Event viewer\Applications and services logs\File Replication Service • Event viewer\Applications and services logs\DFS Replication • Event viewer\Applications and services logs\DNS • Event viewer\Applications and services logs\Microsoft\Windows\Hyper-V-Worker

Tools and Commands for Troubleshooting Domain Controller Configuration

To troubleshoot issues not explained by the logs, use the following tools as a starting point:

- Dcdiag.exe
- Repadmin.exe
- Network Monitor 3.4

General Methodology for Troubleshooting Domain Controller Safe Restore

1. Is the safe snapshot restore expected, but having issues?
    a. Examine the Directory Services event log
        i. Are there snapshot restore errors?
        ii. Are there AD replication errors?
    b. Examine the System event log
        i. Are there communication errors?
        ii. Are there AD errors?
2. Is the safe snapshot restore unexpected?
    a. Examine the hypervisor audit logs to determine who or what caused a rollback
    b. Contact all administrators of the hypervisor and interrogate them as to who rolled back the VM without notification
3. Is the server implementing USN rollback protection and not safely restoring?
    a. Examine the Directory Services event log for an unsupported hypervisor or integration services
    b. Examine the operating system and validate running Windows Server 2012?

Troubleshooting Specific Problems

Events

All virtualized domain controller safe snapshot restore events write to the Directory Services event log of the restored domain controller VM. The Application, System, File Replication Service, and DFS Replication event logs may also contain useful troubleshooting information for failed restores.

Below are the Windows Server 2012 safe restore-specific events in the Directory Services event log.

Event ID	**2170**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Warning
Message	A Generation ID change has been detected. Generation ID cached in DS (old value):%1 Generation ID currently in VM (new value):%2 The Generation ID change occurs after the application of a virtual machine snapshot, after a virtual machine import operation or after a live migration operation. <*COMPUTERNAME*> will create a new invocation ID to recover the domain controller. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services aware backup application.
Notes and resolution	This is a success event if the snapshot was expected. If not, examine the Hyper-V-Worker event log or contact the hypervisor administrator.

Event ID	**2174**
Source	Microsoft-Windows-ActiveDirectory_DomainService

Severity	Informational
Message	The DC is neither a virtual domain controller clone nor a restored virtual domain controller snapshot.
Notes and resolution	Expected event when starting physical domain controllers or virtualized domain controllers not restored from snapshot

Event ID	**2181**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	The transaction was aborted due to the virtual machine being reverted to a previous state. This occurs after the application of a virtual machine snapshot, after a virtual machine import operation, or after a live migration operation.
Notes and resolution	Expected when restoring a snapshot. Transactions track the VM Generation ID changing

Event ID	**2185**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> stopped the FRS or DFSR service used to replicate the SYSVOL folder. Service name:%1 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> must initialize a non-authoritative restore on the local SYSVOL replica. This is performed by stopping the FRS

	or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore. Event 2187 will be logged when FRS or DFSR service is restarted.
Notes and resolution	Expected when restoring a snapshot. All SYSVOL data on this domain controller is replaced with a partner DC's copy.
Event ID	2186
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to stop the FRS or DFSR service used to replicate the SYSVOL folder. Service name:%1 Error code:%2 Error message:%3 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> must initialize a non-authoritative restore on the local SYSVOL replica. This is done by stopping the FRS or DFSR replication service used to replicate the SYSVOL folder and then starting it with the appropriate registry keys and values to trigger the restore. <*COMPUTERNAME*> failed to stop the current running service and cannot complete the non-authoritative restore. Please perform a non-authoritative restore manually.
Notes and resolution	Examine the System, FRS and DFSR event logs for further information.

Event ID	**2187**
Source	Microsoft-Windows-ActiveDirectory_DomainService

Severity	Informational
Message	<*COMPUTERNAME*> started the FRS or DFSR service used to replicate the SYSVOL folder. Service name:%1 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> needed to initialize a non-authoritative restore on the local SYSVOL replica. This was done by stopping the FRS or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore.
Notes and resolution	Expected when restoring a snapshot. All SYSVOL data on this domain controller is replaced with a partner DC's copy.

Event ID	**2188**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to start the FRS or DFSR service used to replicate the SYSVOL folder. Service name:%1 Error code:%2 Error message:%3 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> needs to initialize a non-authoritative restore on the local SYSVOL replica. This is done by stopping the FRS or DFSR service used to replicate the SYSVOL and starting it with appropriate registry keys and values to trigger the restore. <*COMPUTERNAME*> failed to start the FRS or

	DFSR service used to replicate the SYSVOL folder and cannot complete the non-authoritative restore. Please perform a non-authoritative restore manually and restart the service.
Notes and resolution	Examine the System, FRS and DFSR event logs for further information.

Event ID	**2189**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> set the following registry values to initialize SYSVOL replica during a non-authoritative restore: Registry Key:%1 Registry Value: %2 Registry Value data: %3 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> needs to initialize a non-authoritative restore on the local SYSVOL replica. This is done by stopping the FRS or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore.
Notes and resolution	Expected when restoring a snapshot. All SYSVOL data on this domain controller is replaced with a partner DC's copy.

Event ID	**2190**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to set the following

	registry values to initialize the SYSVOL replica during a non-authoritative restore: Registry Key:%1 Registry Value: %2 Registry Value data: %3 Error code:%4 Error message:%5 Active Directory detected that the virtual machine that hosts the domain controller role was reverted to a previous state. <*COMPUTERNAME*> needs to initialize a non-authoritative restore on the local SYSVOL replica. This is done by stopping the FRS or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore. <*COMPUTERNAME*> failed to set the above registry values and cannot complete the non-authoritative restore. Please perform a non-authoritative restore manually.
Notes and resolution	Examine Application and System event logs. Investigate third party applications that may be blocking registry updates.

Event ID	**2200**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> initializes replication to bring the domain controller current. Event 2201 will be logged when the replication is finished.
Notes and resolution	Expected when restoring a snapshot. Marks the beginning of inbound AD replication.

Event ID	2201
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> has finished replication to bring the domain controller current.
Notes and resolution	Expected when restoring a snapshot. Marks the end of inbound AD replication.

Event ID	2202
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> failed replication to bring the domain controller up-to-date. The domain controller will be updated after next periodic replication.
Notes and resolution	Examine the Directory Services and System event logs. Use repadmin.exe to attempt forcing replication and note any failures.

Event ID	2204
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> has detected a change of virtual machine generation ID. The change means that the virtual domain controller has been reverted to a previous state.

	<*COMPUTERNAME*> will perform the following operations to protect the reverted domain controller against possible data divergence and to protect creation of security principals with duplicate SIDs: Create a new invocation ID Invalidate current RID pool Ownership of the FSMO roles will be validated at next inbound replication. During this window if the domain controller held a FSMO role, that role will be unavailable. Start SYSVOL replication service restore operation. Start replication to bring the reverted domain controller to the most current state. Request a new RID pool.
Notes and resolution	Expected when restoring a snapshot. This explains all the various reset operations that will occur as part of the safe restore process.

Event ID	**2205**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> invalidated current RID pool after virtual domain controller was reverted to previous state.
Notes and resolution	Expected when restoring a snapshot. The local RID pool must be destroyed as the domain controller has time travelled and they may have already been issued.

Event ID	**2206**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	ERROR

Message	<*COMPUTERNAME*> failed to invalidate current RID pool after virtual domain controller was reverted to previous state. Additional data: Error code: %1 Error value: %2
Notes and resolution	Examine the Directory Services and System event logs. Validate that the RID Master is online can be reached from this server using Dcdiag.exe /test:ridmanager

Event ID	**2207**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	ERROR
Message	<*COMPUTERNAME*> failed to restore after virtual domain controller was reverted to previous state. A reboot into DSRM was requested. Please check previous events for more information.
Notes and resolution	Examine the Directory Services and System event logs.

Event ID	**2208**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Informational
Message	<*COMPUTERNAME*> deleted DFSR databases to initialize SYSVOL replica during a non-authoritative restore.
Notes and resolution	Expected when restoring a snapshot. This guarantees DFSR non-authoritatively synchronizes SYSVOL from a partner DC. Note that any other DFSR Replicated Folders on the same volume as SYSVOL will also non-authoritatively sync (domain controllers are not

	recommended to host custom DFSR sets on the same volume as SYSVOL).

Event ID	**2209**
Source	Microsoft-Windows-ActiveDirectory_DomainService
Severity	Error
Message	<*COMPUTERNAME*> failed to delete DFSR databases. Additional data: Error code: %1 Error value: %2 Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. <*COMPUTERNAME*> needs to initialize a non-authoritative restore on the local SYSVOL replica. For DFSR, this is done by stopping the DFSR service, deleting DFSR databases, and re-starting the service. Upon restarting DFSR will rebuild the databases and start the initial sync.
Notes and resolution	Examine the DFSR event log.

Error Messages

There are no direct interactive errors for failed virtualized domain controller safe snapshot restore; all cloning information logs in the Directory Services event logs. Naturally, any critical replication or server advertising errors manifest themselves as symptoms elsewhere.

Known Issues and Support Scenarios

The General Methodology for Troubleshooting Domain Controller Safe Restore section and events listed in Troubleshooting Specific Problems are usually adequate to troubleshoot most issues.

Issue	**Cannot create new security principals on recently safe restored domain controller**

Symptoms	After restoring a snapshot, attempts to create a new security principal (user, computer, group) on that domain controller fail with: Error 0x2010 The directory service was unable to allocate a relative identifier.
Resolution and Notes	This issue is caused by the restored computer's stale knowledge of the RID Master FSMO role. If the role moved to this or another domain controller after a snapshot was taken and then later restored, the restored domain controller will not have knowledge of the RID master until initial replication has completed. To resolve the issue, allow AD replication to complete inbound to the restored domain controller. If still not working, validate that all domain controllers have the same correct knowledge of which DC hosts the RID Master.

Issue	**Restored domain controllers do not share SYSVOL, advertise**
Symptoms	After restoring a snapshot, one or more DCs do not advertise, do not share sysvol, and do not have up to date SYSVOL contents
Resolution and Notes	The DC's upstream partners do not have a working SYSVOL replica that is correctly replicating with DFSR or FRS. This issue is unrelated to safe restore but is likely to manifest as a safe restore issue, because the customer was unaware of the other replication issue affecting un-restored DCs

Advanced Troubleshooting

This module seeks to teach advanced troubleshooting by using *working* logs as samples, with some explanation of what occurred. If you understand what a successful virtualized domain controller operation looks like, failures become obvious in your environment. These logs are presented by their source, with the ascending order of *expected* events related to a cloned domain controller within each log.

Restoring a Domain Controller that Replicates SYSVOL Using DFSR

Directory Services Event Log

The Directory Services log contains the majority of safe restore operational information. The hypervisor changes the VM-Generation ID and the NTDS service notes it, then invalidates the RID pool and changes the invocation ID. The new VM-Generation ID is set and the servers replicates AD data inbound. The DFSR service is stopped and its database that hosts SYSVOL is deleted, forcing a non-authoritative sync inbound. The USN high watermark is adjusted.

Event ID	Source	Message
2170	ActiveDirectory_DomainService	A Generation ID change has been detected. Generation ID cached in DS (old value): *<number>* Generation ID currently in VM (new value): *<number>* The Generation ID change occurs after the application of a virtual machine snapshot, after a virtual machine import operation or after a live migration operation. Active Directory Domain Services will create a new invocation ID to recover the domain controller. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services aware backup application."
2181	ActiveDirectory_DomainService	The transaction was aborted due to the virtual machine

		being reverted to a previous state. This occurs after the application of a virtual machine snapshot, after a virtual machine import operation, or after a live migration operation.
2204	ActiveDirectory_DomainService	Active Directory Domain Services has detected a change of virtual machine generation ID. The change means that the virtual domain controller has been reverted to a previous state. Active Directory Domain Services will perform the following operations to protect the reverted domain controller against possible data divergence and to protect creation of security principals with duplicate SIDs: Create a new invocation ID Invalidate current RID pool Ownership of the FSMO roles will be validated at next inbound replication. During this window if the domain controller held a FSMO role, that role will be unavailable. Start SYSVOL replication service restore operation. Start replication to bring the reverted domain controller to the most current state. Request a new RID pool."
2181	ActiveDirectory_DomainService	The transaction was aborted due to the virtual machine being reverted to a previous state. This occurs after the

		application of a virtual machine snapshot, after a virtual machine import operation, or after a live migration operation.
1109	ActiveDirectory_DomainService	The invocationID attribute for this directory server has been changed. The highest update sequence number at the time the backup was created is as follows: InvocationID attribute (old value): <*GUID*> InvocationID attribute (new value): <*GUID*> Update sequence number: <*number*> The invocationID is changed when a directory server is restored from backup media, is configured to host a writeable application directory partition, has been resumed after a virtual machine snapshot has been applied, after a virtual machine import operation, or after a live migration operation. Virtualized domain controllers should not be restored using virtual machine snapshots. The supported method to restore or rollback the content of an Active Directory Domain Services database is to restore a system state backup made with an Active Directory Domain Services-

		aware backup application."
2179	ActiveDirectory_DomainService	The msDS-GenerationId attribute of the Domain Controller's computer object has been set to the following parameter: GenerationID attribute: <*number*>
2200	ActiveDirectory_DomainService	Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. Active Directory Domain Services initializes replication to bring the domain controller current. Event 2201 will be logged when the replication is finished.
2201	ActiveDirectory_DomainService	Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. Active Directory Domain Services has finished replication to bring the domain controller current.
2185	ActiveDirectory_DomainService	Active Directory Domain Services stopped the FRS or DFSR service used to replicate the SYSVOL folder. Service name: DFSR Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. Active Directory Domain Services must initialize a non-authoritative restore on the local SYSVOL replica.

		This is performed by stopping the FRS or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore. Event 2187 will be logged when FRS or DFSR service is restarted."
2208	ActiveDirectory_DomainService	Active Directory Domain Services deleted DFSR databases to initialize SYSVOL replica during a non-authoritative restore. Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. Active Directory Domain Services needs to initialize a non-authoritative restore on the local SYSVOL replica. For DFSR, this is done by stopping the DFSR service, deleting DFSR databases, and re-starting the service. Upon restarting DFSR will rebuild the databases and start the initial sync. "
2187	ActiveDirectory_DomainService	Active Directory Domain Services started the FRS or DFSR service used to replicate the SYSVOL folder. Service name: DFSR Active Directory detected that the virtual machine that hosts the domain controller was reverted to a previous state. Active Directory Domain Services needed to initialize

		a non-authoritative restore on the local SYSVOL replica. This was done by stopping the FRS or DFSR service used to replicate the SYSVOL folder and starting it with the appropriate registry keys and values to trigger the restore. "
1587	ActiveDirectory_DomainService	This directory service has been restored or has been configured to host an application directory partition. As a result, its replication identity has changed. A partner has requested replication changes using our old identity. The starting sequence number has been adjusted. The destination directory service corresponding to the following object GUID has requested changes starting at a USN that precedes the USN at which the local directory service was restored from backup media. Object GUID: *<GUID> (<FQDN of partner domain controller>)* USN at the time of restore: *<number>* As a result, the up-to-dateness vector of the destination directory service has been configured with the following settings. Previous database GUID: *<GUID>*

		Previous object USN: *<number>* Previous property USN: *<number>* New database GUID: *<GUID>* New object USN: *<number>* New property USN: *<number>*

System Event Log

The System event log notes that the machine time that occurs when bringing an offline virtual machine back online and synchronizing with host time. The RID pool invalidates and the DFSR or FRS services are restarted.

Event ID	Source	Message
1	Kernel-General	The system time has changed to *<now>* from *<snapshot time/date>*. Change Reason: An application or system component changed the time.
16654	Directory-Services-SAM	A pool of account-identifiers (RIDs) has been invalidated. This may occur in the following expected cases: 1. A domain controller is restored from backup. 2. A domain controller running on a virtual machine is restored from snapshot. 3. An administrator has manually invalidated the pool. See http://go.microsoft.com/fwlink/?LinkId=226247 for more information.
7036	Service Control Manager	The DFS Replication service entered the stopped state.
7036	Service Control Manager	The DFS Replication service entered the running state.

Application Event Log

The Application event log notes the DFSR database stopping and starting.

Event ID	Source	Message
103	ESENT	DFSRs (1360) \\.\C:\System Volume Information\DFSR\database_<*GUID*>\dfsr.db: The database engine stopped the instance (0). Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.141, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.016, [12] 0.000, [13] 0.000, [14] 0.000, [15] 0.000.
102	ESENT	DFSRs (532) \\.\C:\System Volume Information\DFSR\database_<*GUID*>\dfsr.db: The database engine (6.02.8189.0000) is starting a new instance (0).
105	ESENT	DFSRs (532) \\.\C:\System Volume Information\DFSR\database_<*GUID*>\dfsr.db: The database engine started a new instance (0). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.031, [10] 0.000, [11] 0.000.
		DFSRs (532) \\.\C:\System Volume Information\DFSR\database_<*GUID*>\dfsr.db: The database engine created a new database (1, \\.\C:\System Volume Information\DFSR\database_<*GUID*>\dfsr.db). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.016, [4] 0.062, [5] 0.000, [6] 0.016, [7] 0.000, [8] 0.000, [9] 0.015, [10] 0.000, [11] 0.000.

DFS Replication Event Log

The DFSR service is stopped and the database that contains SYSVOL is deleted, forcing a non-authoritative synchronization inbound.

Event ID	Source	Message
1006	DFSR	The DFS Replication service is stopping.
1008	DFSR	The DFS Replication service has stopped.
1002	DFSR	The DFS Replication service is starting.
1004	DFSR	The DFS Replication service has started.
1314	DFSR	The DFS Replication service successfully configured the debug log files. Additional Information: Debug Log File Path: C:\Windows\debug
6102	DFSR	The DFS Replication service has successfully registered the WMI provider.
1206	DFSR	The DFS Replication service successfully contacted domain controller *<domain controller FQDN>* to access configuration information.
1210	DFSR	The DFS Replication service successfully set up an RPC listener for incoming replication requests. Additional Information: Port: 0
4614	DFSR	The DFS Replication service initialized SYSVOL at local path C:\Windows\SYSVOL\domain and is waiting to perform initial replication. The replicated folder will remain in the initial synchronization state until it has replicated with its partner. If the server

		was in the process of being promoted to a domain controller, the domain controller will not advertise and function as a domain controller until this issue is resolved. This can occur if the specified partner is also in the initial synchronization state, or if sharing violations are encountered on this server or the synchronization partner. If this event occurred during the migration of SYSVOL from File Replication Service (FRS) to DFS Replication, changes will not replicate out until this issue is resolved. This can cause the SYSVOL folder on this server to become out of sync with other domain controllers. Additional Information: Replicated Folder Name: SYSVOL Share Replicated Folder ID: <*GUID*> Replication Group Name: Domain System Volume Replication Group ID: <*GUID*> Member ID: <*GUID*> Read-Only: 0
4604	DFSR	The DFS Replication service successfully initialized the SYSVOL replicated folder at local path C:\Windows\SYSVOL\domain. This member has completed initial synchronization of SYSVOL with partner dc1.corp.contoso.com. To check for the presence of the SYSVOL share, open a command prompt window and then type "net share". Additional Information: Replicated Folder Name: SYSVOL Share Replicated Folder ID: <*GUID*>

		Replication Group Name: Domain System Volume Replication Group ID: <*GUID*> Member ID: <*GUID*> Sync partner: <*partner domain controller FQDN*>

Restoring a Domain Controller that Replicates SYSVOL Using FRS

The File Replication Event log is used instead of the DFSR event log in this case. The Application event log also writes different FRS-related events. Otherwise, the Directory Services and System Event log messages are generally the same and in the same order as previously described.

File Replication Service Event Log

The FRS service is stopped and restarted with a D2 BURFLAGS value to non-authoritatively synchronize SYSVOL.

Event ID	Source	Message
13502	NTFRS	The File Replication Service is stopping.
13503	NTFRS	The File Replication Service has stopped.
13501	NTFRS	The File Replication Service is starting
13512	NTFRS	The File Replication Service has detected an enabled disk write cache on the drive containing the directory c:\windows\ntfrs\jet on the computer DC4. The File Replication Service might not recover when power to the drive is interrupted and critical updates are lost.
13565	NTFRS	File Replication Service is initializing the system volume with data from another domain controller. Computer DC4 cannot become a domain controller until this process is complete. The system volume will then be shared as SYSVOL. To check for the SYSVOL share, at the command prompt, type: net share When File Replication Service completes the initialization process, the SYSVOL share will

		appear. The initialization of the system volume can take some time. The time is dependent on the amount of data in the system volume, the availability of other domain controllers, and the replication interval between domain controllers."
13520	NTFRS	The File Replication Service moved the preexisting files in <*path*> to <*path*>\NtFrs_PreExisting___See_EventLog. The File Replication Service may delete the files in <*path*>\NtFrs_PreExisting___See_EventLog at any time. Files can be saved from deletion by copying them out of <*path*>\NtFrs_PreExisting___See_EventLog. Copying the files into <*path*> may lead to name conflicts if the files already exist on some other replicating partner. In some cases, the File Replication Service may copy a file from <*path*>\NtFrs_PreExisting___See_EventLog into <*path*> instead of replicating the file from some other replicating partner. Space can be recovered at any time by deleting the files in <*path*>\NtFrs_PreExisting___See_EventLog.
13553	NTFRS	The File Replication Service successfully added this computer to the following replica set: "DOMAIN SYSTEM VOLUME (SYSVOL SHARE)" Information related to this event is shown below: Computer DNS name is "<*domain controller FQDN*>" Replica set member name is "<*domain controller name*>" Replica set root path is "<*path*>" Replica staging directory path is "<*path*> "

		Replica working directory path is "<*path*>"
13554	NTFRS	The File Replication Service successfully added the connections shown below to the replica set: "DOMAIN SYSTEM VOLUME (SYSVOL SHARE)" Inbound from "<*partner domain controller FQDN*>" Outbound to "<*partner domain controller FQDN*>" More information may appear in subsequent event log messages.
13516	NTFRS	The File Replication Service is no longer preventing the computer DC4 from becoming a domain controller. The system volume has been successfully initialized and the Netlogon service has been notified that the system volume is now ready to be shared as SYSVOL. Type "net share" to check for the SYSVOL share.

Application Event Log

The FRS database stops and starts, and is purged due to the D2 BURFLAGS operation.

Event ID	Source	Message
327	ESENT	ntfrs (1424) The database engine detached a database (1, c:\windows\ntfrs\jet\ntfrs.jdb). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.015, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.516, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.063, [12] 0.000. Revived Cache: 0
103	ESENT	ntfrs (1424) The database engine stopped the instance (0).

		Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.031, [10] 0.000, [11] 0.016, [12] 0.000, [13] 0.000, [14] 0.047, [15] 0.000.
102	ESENT	ntfrs (3000) The database engine (6.02.8189.0000) is starting a new instance (0).
105	ESENT	ntfrs (3000) The database engine started a new instance (0). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.062, [10] 0.000, [11] 0.141.
103	ESENT	ntfrs (3000) The database engine stopped the instance (0). Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.000, [12] 0.000, [13] 0.015, [14] 0.000, [15] 0.000.
102	ESENT	ntfrs (3000) The database engine (6.02.8189.0000) is starting a new instance (0).
105	ESENT	ntfrs (3000) The database engine started a new instance (0). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.078, [10] 0.000, [11] 0.109.
325	ESENT	ntfrs (3000) The database engine created a new database (1, c:\windows\ntfrs\jet\ntfrs.jdb). (Time=0

		seconds) Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.016, [4] 0.016, [5] 0.000, [6] 0.015, [7] 0.000, [8] 0.000, [9] 0.078, [10] 0.016, [11] 0.000.
103	ESENT	ntfrs (3000) The database engine stopped the instance (0). Dirty Shutdown: 0 Internal Timing Sequence: [1] 0.000, [2] 0.000, [3] 0.000, [4] 0.000, [5] 0.078, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.125, [10] 0.016, [11] 0.000, [12] 0.000, [13] 0.000, [14] 0.000, [15] 0.000.
102	ESENT	ntfrs (3000) The database engine (6.02.8189.0000) is starting a new instance (0).
105	ESENT	ntfrs (3000) The database engine started a new instance (0). (Time=0 seconds) Internal Timing Sequence: [1] 0.016, [2] 0.000, [3] 0.000, [4] 0.094, [5] 0.000, [6] 0.000, [7] 0.000, [8] 0.000, [9] 0.032, [10] 0.000, [11] 0.000.
326	ESENT	ntfrs (3000) The database engine attached a database (1, c:\windows\ntfrs\jet\ntfrs.jdb). (Time=0 seconds) Internal Timing Sequence: [1] 0.000, [2] 0.015, [3] 0.000, [4] 0.000, [5] 0.016, [6] 0.015, [7] 0.000, [8] 0.000, [9] 0.000, [10] 0.000, [11] 0.000, [12] 0.000. Saved Cache: 1

Virtualized Domain Controller Technical Reference Appendix

This topic covers:

- Terminology
- FixVDCPermissions.ps1

Terminology

- **Snapshot** – The state of a virtual machine at a particular point in time. It is dependent on the chain of previous snapshots taken, on the hardware, and on the virtualization platform.
- **Clone** – A complete and separate copy of a virtual machine. It is dependent on the virtual hardware (hypervisor).
- **Full Clone** – A full clone is an independent copy of a virtual machine that shares no resources with the parent virtual machine after the cloning operation. Ongoing operation of a full clone is entirely separate from the parent virtual machine.
- **Differencing disk** - A copy of a virtual machine that shares virtual disks with the parent virtual machine in an ongoing manner. This usually conserves disk space and allows multiple virtual machines to use the same software installation.
- **VM Copy**- A file system copy of all the related files and folders of a virtual machine.
- **VHD File Copy** – A copy of a virtual machine's VHD
- **VM Generation ID** – a 128-bit integer given to the virtual machine by the hypervisor. This ID is stored in memory and reset every time a snapshot is applied. The design uses a hypervisor-agnostic mechanism for surfacing the VM-Generation ID in the virtual machine. The Hyper-V implementation exposes the ID in the ACPI table of the virtual machine.
- **Import/Export** – A Hyper-V feature that allows the user to save the entire virtual machine (VM files, VHD and the machine configuration). It then allows users to using that set of files to bring the machine back on the same machine as the same VM (Restore), on a different machine as the same VM (Move), or a new VM (copy)

FixVDCPermissions.ps1

```
# Unsigned script, requires use of set-executionpolicy remotesigned -force

# You must run the Windows PowerShell console as an elevated administrator


# Load Active Directory Windows PowerShell Module and switch to AD DS drive

import-module activedirectory

cd ad:


## Get Domain NC
```

```
$domainNC = get-addomain

## Get groups and obtain their SIDs
$dcgroup = get-adgroup "Cloneable Domain Controllers"

$sid1 = (get-adgroup $dcgroup).sid

## Get the DACL of the domain
$acl = get-acl $domainNC

## The following object specific ACE grants extended right 'Allow a DC to create a clone
of itself' for the CDC group to the Domain NC
## 3e0f7e18-2c7a-4c10-ba82-4d926db99a3e is the schemaIDGuid for 'DS-Clone-Domain-
Controller"

$objectguid = new-object Guid 3e0f7e18-2c7a-4c10-ba82-4d926db99a3e
$ace1 = new-object System.DirectoryServices.ActiveDirectoryAccessRule
$sid1,"ExtendedRight","Allow",$objectguid

## Add the ACE in the ACL and set the ACL on the object

$acl.AddAccessRule($ace1)
set-acl -aclobject $acl $domainNC
write-host "Done writing new VDC permissions."
cd c:
```

Virtualized Domain Controller Additional Resources

- AD DS Virtualization (Cloning and Virtualization safe improvements)
- Active Directory Replication Model Technical Reference
- Running Domain Controllers in Hyper-V (Windows Server 2008 R2 Behaviors)
- USN and USN Rollback Protection (Windows Server 2008 R2)
- Active Directory Administration with Windows PowerShell (Windows Server 2008 R2)

- Hyper-V in Windows Server 2012
- Ask the Directory Services Team (Official Microsoft Commercial Technical Support Blog)

Virtualized Domain Controller Cloning Test Guidance for Application Vendors

This topic explains what application vendors should consider to help ensure their application continues to work as expected after the virtualized domain controller (DC) cloning process completes. It covers those aspects of the cloning process that interest application vendors and scenarios that may warrant additional testing. Application vendors who have validated that their application works on virtualized domain controllers that have been cloned are encouraged to list the name of the application in the Community Content at the bottom of this topic, along with a link to your organization's web site where users can learn more about the validation.

Overview of virtualized DC cloning

The virtualized domain controller cloning process is described in detail in Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100) and Virtualized Domain Controller Technical Reference (Level 300). From an application vendor's perspective, these are some considerations to take into account when assessing the impact of cloning to your application:

- The original computer is not destroyed. It remains on the network, interacting with clients. Unlike a rename where the DNS records of the original computer are removed, the original records for the source domain controller remain.
- During the cloning process, the new computer is initially running for a brief period of time under the identity of the old computer until the cloning process is initiated and makes the necessary changes. Applications that create records about the host should ensure that the cloned computer does not overwrite records about the original host during the cloning process.
- Cloning is a specific deployment capability for only virtualized domain controllers, not a general purpose extension to clone other server roles. Some server roles are specifically not supported for cloning:
  - Dynamic Host Configuration Protocol (DHCP)
  - Active Directory Certificate Services (AD CS)
  - Active Directory Lightweight Directory Services (AD LDS)
- As part of the cloning process, the entire VM that represents the original DC is copied, so any application state on that VM is also copied. Validate that the application adapts to this change in state of the local host on the cloned DC, or if any intervention is required, such as a service restart.
- As part of cloning, the new DC gets a new machine identity and provisions itself as a replica DC in the topology. Validate whether the application depends on the machine identity, such as its name, account, SID, and so on. Does it automatically adapt to the change of machine

identity on the clone? If that application caches data, ensure that it does not rely on machine identity data that may be cached.

What is interesting for Application Vendors?

CustomDCCloneAllowList.xml

A domain controller that runs your application or service cannot be cloned until the application or service is either:

- Added to the CustomDCCloneAllowList.xml file by using the Get-ADDCCloningExcludedApplicationList Windows PowerShell cmdlet

-Or-

- Removed from the domain controller

The first time the user runs the Get-ADDCCloningExcludedApplicationList cmdlet, it returns a list of services and applications that are running on the domain controller but are not in the default list of services and applications that are supported for cloning. By default, your service or application will not be listed. To add your service or application to the list of applications and services that can be safely cloned, the user runs Get-ADDCCloningExcludedApplicationList cmdlet again with the –GenerateXML option in order to add it to the CustomDCCloneAllowList.xml file. For more information, see Step 2: Run Get-ADDCCloningExcludedApplicationList cmdlet.

Distributed System Interactions

Usually services isolated to the local computer either pass or fail when participating in cloning. Distributed services have to be concerned about having two instances of the host computer on the network simultaneously for a brief period of time. This may manifest as a service instance trying to pull information from a partner system where the clone has registered as the new vendor of the identity. Or both instances of the service may push information into the AD DS database at the same time with different results. For example, it is not deterministic which computer will be communicated with when two computers that have Windows Testing Technologies (WTT) service are on the network with the domain controller.

For the distributed DNS Server service, the cloning process carefully avoids overwriting the DNS records of the source domain controller when the clone domain controller starts with a new IP address.

You should not rely on the computer to remove all of the old identity until the end of cloning. After the new domain controller is promoted inside the new context, select Sysprep providers are run to clean up the additional state of the computer. For example, it is at this point the old certificates of the computer are removed and the cryptography secrets that the computer can access are changed.

The greatest factor that varies the timing of the cloning is how many objects there are to replicate from the PDC. Older media increases the time required to complete cloning.

Because your service or application is unknown, it is left running. The cloning process does not change the state of non-Windows services.

Additionally, the new computer has a different IP address than the original computer. These behaviors may cause side effects to your service or application depending on how the service or application behaves in this environment.

Additional scenarios suggested for testing

Cloning Failure

Service vendors should test this scenario because when cloning fails the computer boots into Directory Services Repair Mode (DSRM), a form of Safe Mode. At this point the computer has not completed cloning. Some state may have changed and some state may remain from the original domain controller. Test this scenario to understand what impact it can have on your application.

To induce a cloning failure, try to clone a domain controller without granting it permission to be cloned. In this case, the computer will have only changed the IP addresses and still have the majority of its state from the original domain controller. For more information about granting a domain controller permission to be cloned, see Step 1: Grant the source virtualized domain controller the permission to be cloned.

PDC emulator cloning

Service and application vendors should test this scenario because there is an additional reboot when the PDC emulator is cloned. In addition, the majority of cloning is performed under a temporary identity to allow the new clone to interact with the PDC emulator during the cloning process.

Writable versus read-only domain controllers

Service and application vendors should test cloning by using the same type of domain controller (that is, on a writable or read-only domain controller) that service is planned to run on.

Support for using Hyper-V Replica for virtualized domain controllers

This topic explains the supportability of using Hyper-V Replica to replicate a virtual machine (VM) that runs as a domain controller (DC). Hyper-V Replica is a new capability of Hyper-V beginning with Windows Server 2012 that provides a built-in replication mechanism at a VM level.

Hyper-V Replica asynchronously replicates selected VMs from a primary Hyper-V host to a replica Hyper-V host across either LAN or WAN links. After initial replication is complete, subsequent changes are replicated at an interval defined by the administrator.

Failover can be either planned or unplanned. A planned failover is initiated by an administrator on the primary VM, and any un-replicated changes are copied over to the replica VM to prevent any data loss. An unplanned failover is initiated on the replica VM in response to an unexpected failure of the primary VM. Data loss is possible because there is no opportunity to transmit changes on the primary VM that might not have been replicated yet.

For more information about Hyper-V Replica, see Hyper-V Replica Overview and Deploy Hyper-V Replica.

Note

Hyper-V Replica can be run only on Windows Server Hyper-V, not the version of Hyper-V that runs on Windows 8.

Windows Server 2012 domain controllers required

Windows Server 2012 Hyper-V also introduces VM-GenerationID (VMGenID). VMGenID provides a way for the hypervisor to communicate to the guest OS when significant changes have occurred. For example, the hypervisor can communicate to a virtualized DC that a restore from snapshot has occurred (Hyper-V snapshot restore technology, not backup restore). AD DS in Windows Server 2012 is aware of VMGenID VM technology and uses it to detect when hypervisor operations are performed, such as snapshot restore, which allows it to better protect itself.

Note

To reinforce the point, only AD DS on Windows Server 2012 DCs provides these safety measures resulting from VMGenID; DCs that run all previous releases of Windows Server are subject to problems such as USN rollback that can occur when a virtualized DC is restored using an unsupported mechanism, such as snapshot restore. For more information about these safeguards and when they are triggered, see Virtualized Domain Controller Architecture.

When a Hyper-V replica failover occurs (planned or unplanned), the Windows Server 2012 virtualized DC detects a VMGenID reset, triggering the aforementioned safety features. Active Directory operations then proceed as normal. The replica VM runs in place of the primary VM.

Note

Given that now there are now two instances of the same DC identity, there is a potential for both the primary instance and the replicated instance to run. While Hyper-V Replica has control mechanisms in place to ensure the primary and replica VMs do not run simultaneously, it is possible for them to run at the same time in the event the link between them fails after replication of the VM. In the event of this unlikely occurrence, virtualized DCs that run Windows Server 2012 have safeguards to help protect AD DS, whereas virtualized DCs that run earlier versions of Windows Server do not.

When using Hyper-V Replica, ensure that you follow best practices for running virtual domain controllers on Hyper-V. This discusses, for example, recommendations for storing Active Directory files on virtual SCSI disks, which provides stronger guarantees of data durability.

Supported and unsupported scenarios

Only VMs that run Windows Server 2012 are supported for unplanned failover and for testing failover. Even for planned failover, Windows Server 2012 is recommended for the virtualized DC in order to mitigate risks in the event that an administrator inadvertently starts both the primary VM and the replicated VM at the same time.

VMs that run earlier versions of Windows Server are supported for planned failover but unsupported for unplanned failover because of the potential for USN rollback. For more information about USN rollback, see USN and USN Rollback.

Note
There are no functional level requirements for the domain or forest; there are only operating system requirements for the DCs that run as VMs that are replicated using Hyper-V Replica. The VMs can be deployed in a forest that contains other physical or virtual DCs that run earlier versions of Windows Server and may or may not also be replicated using Hyper-V Replica.

This support statement is based on tests that were performed in a single domain-forest, though multi-domain forest configurations are also supported. For these tests, virtualized domain controllers DC1 and DC2 are Active Directory replication partners in the same site, hosted on a server that runs Hyper-V on Windows Server 2012. The VM guest that runs DC2 has Hyper-V Replica enabled. The Replica server is hosted in another geographically distant datacenter. To help explain the test case processes outlined below, the VM running on the replica server is referred to as DC2-Rec (although in practice it retains the same name as the original VM).

Windows Server 2012

The following table explains support for virtualized DCs that run Windows Server 2012 and test cases.

Planned Failover	Unplanned Failover
Supported	Supported
Test case: • DC1 and DC2 are running Windows Server 2012. • DC2 is shut down and a failover is performed on DC2-Rec. The failover can be either planned or unplanned. • After DC2-Rec starts, it checks whether the value of VMGenID that it has in its database is the same as the value from the virtual machine driver saved by the Hyper-V Replica server.	The test case is the same as for a planned failover, with these exceptions: • Any AD updates received on DC2 but not yet replicated by AD to a replication partner before the failover event will be lost. • AD updates received on DC2 after the time of the recovery point that were replicated by AD to DC1 will be replicated from DC1 back to DC2-Rec.

• As a result, DC2-Rec triggers virtualization safeguards; in other words, it resets its InvocationID, discards its RID pool, and sets an initial synchronization requirement before it will assume an operations master role. For more information about initial synchronization requirement, see . • DC2-Rec then saves the new value of VMGenID in its database and commits any subsequent updates in the context of the new InvocationID. • As a result of the InvocationID reset, DC1 will converge on all AD changes introduced by DC2-Rec even if it was rolled back in time, meaning any AD updates performed on DC2-Rec after the failover will safely converge	

Windows Server 2008 R2 and earlier versions

The following table explains support for virtualized DCs that run Windows Server 2008 R2 and earlier versions.

Planned Failover	Unplanned Failover
Supported but not recommended because DCs that run these versions of Windows Server do not support VMGenID or use associated virtualization safeguards. This places them at risk for USN rollback. For more information, see USN and USN Rollback.	Not supported **Note** Unplanned failover would be supported where USN rollback is not a risk, such as a single DC in the forest (a configuration that is not recommended).
Test case: • DC1 and DC2 are running Windows Server 2008 R2. • DC2 is shut down and a planned failover is performed on DC2-Rec. All data on DC2 is replicated to DC2-Rec before the shutdown is complete. • After DC2-Rec starts, it resumes replication with DC1 using the same invocationID as DC2.	N/A

Deploy Active Directory Domain Services (AD DS) in Your Enterprise

This guide covers how to install and remove Active Directory Domain Services (AD DS) in Windows Server 2012, and important issues to be aware of when you add new domain controllers to an existing Active Directory environment.

- What's New in Active Directory Domain Services Installation and Removal
- Upgrade Domain Controllers to Windows Server 2012
- Install Active Directory Domain Services (Level 100)
- Remove Active Directory Domain Services (Level 100)
- AD DS Installation and Removal Wizard Page Descriptions
- Windows Server 2012: Changes Made by Adprep.exe
- Troubleshooting Domain Controller Deployment

What's New in Active Directory Domain Services Installation and Removal

This topic covers:

- The Active Directory Domain Services Configuration Wizard
- Adprep.exe integration
- AD DS installation prerequisite validation
- System requirements
- Known issues

Active Directory Domain Services (AD DS) deployment in Windows Server 2012 is simpler and faster than previous versions of Windows Server. The AD DS installation process is now built on Windows PowerShell and is integrated with Server Manager. The number of steps required to introduce domain controllers into an existing Active Directory environment is reduced. This makes the process for creating a new Active Directory environment simpler and more efficient. The new AD DS deployment process minimizes the chances of errors that would have otherwise blocked installation.

In addition, you can install the AD DS server role binaries (that is the AD DS server role) on multiple servers at the same time. You can also run the AD DS installation wizard remotely on an individual server. These improvements provide more flexibility for deploying domain controllers that run Windows Server 2012, especially for large-scale, global deployments where many domain controllers need to be deployed to offices in different regions.

AD DS installation includes the following features:

- **Adprep.exe integration into the AD DS installation process.** The cumbersome steps required to prepare an existing Active Directory, such as the need to use a variety of different credentials, copy the Adprep.exe files, or log on to specific domain controllers, are all simplified or occur automatically. This reduces the time required to install AD DS and reduces the chances for errors that might otherwise block domain controller promotion.

  For environments where it is preferable to run adprep.exe commands in advance of a new domain controller installation, you can still execute adprep.exe commands separately from the AD DS installation. The Windows Server 2012 version of adprep.exe runs remotely, so you can execute all necessary commands from a server that runs a 64-bit version of Windows Server 2008 or later.
- **The new AD DS installation is built on Windows PowerShell and can be invoked remotely.** The new AD DS installation is integrated with Server Manager, so you can use the same interface to install AD DS that you use when installing other server roles. For Windows PowerShell users, the AD DS deployment cmdlets provide greater functionality and flexibility. There is functional parity between command-line and GUI installation options.
- **The new AD DS installation includes prerequisite validation.** Any potential errors are identified before the installation begins. You can correct error conditions before they occur without the concerns resulting from a partially complete upgrade. For example, if adprep /domainprep needs to be run, the installation wizard verifies that the user has sufficient rights to execute the operation.
- **Configuration pages are grouped in a sequence that mirrors the requirements of the most common promotion options with related options grouped in fewer wizard pages.** This provides better context for making installation choices.
- **You can export a Windows PowerShell script that contains all the options that were specified during the graphical installation.** At the end of an installation or removal, you can export the settings to a Windows PowerShell script for use with automating the same operation.
- **Only critical replication occurs before reboot.** New switch to allow replication of non-critical data before reboot. For more information, see ADDSDeployment cmdlet arguments.

The Active Directory Domain Services Configuration Wizard

Beginning with Windows Server 2012, the Active Directory Domain Services Configuration Wizard replaces the legacy Active Directory Domain Services Installation Wizard as the user interface (UI) option to specify settings when you install a domain controller. The Active Directory Domain Services Configuration Wizard begins after Add Roles Wizard is finished.

Warning

The legacy Active Directory Domain Services Installation Wizard (dcpromo.exe) is deprecated beginning with Windows Server 2012.

In Install Active Directory Domain Services (Level 100), the UI procedures show how to start the Add Roles Wizard to install the AD DS server role binaries and then run the Active Directory Domain Services Configuration Wizard to complete the domain controller installation. The

Windows PowerShell examples show how to complete both steps using an AD DS deployment cmdlet.

Adprep.exe integration

Beginning with Windows Server 2012, there is only one version of Adprep.exe (there is no 32-bit version, adprep32.exe). Adprep commands are run automatically as needed when you install a domain controller that runs Windows Server 2012 to an existing Active Directory domain or forest.

Although adprep operations are run automatically, you can run Adprep.exe separately. For example, if the user who installs AD DS is not a member of the Enterprise Admins group, which is required in order to run Adprep /forestprep, then you might need to run the command separately. But, you only have to run adprep.exe if you are planning to in-place upgrade your first Windows Server 2012 domain controller (in other words, you plan to in-place upgrade the operating system of a domain controller that runs Windows Server 2012).

Adprep.exe is located in the \support\adprep folder of the Windows Server 2012 installation disc. The Windows Server 2012 version of adprep is capable of executing remotely.

The Windows Server 2012 version of adprep.exe can run on any server that runs a 64-bit version of Windows Server 2008 or later. The server needs network connectivity to the schema master for the forest and the infrastructure master of the domain where you want to add a domain controller. If either of those roles is hosted on a server that runs Windows Server 2003, then adprep must be run remotely. The server where you run adprep does not need to be a domain controller. It can be domain joined or in a workgroup.

Notes

If you try to run the Windows Server 2012 version of adprep.exe on a server that runs Windows Server 2003, the following error appears:

Adprep.exe is not a valid Win32 application.

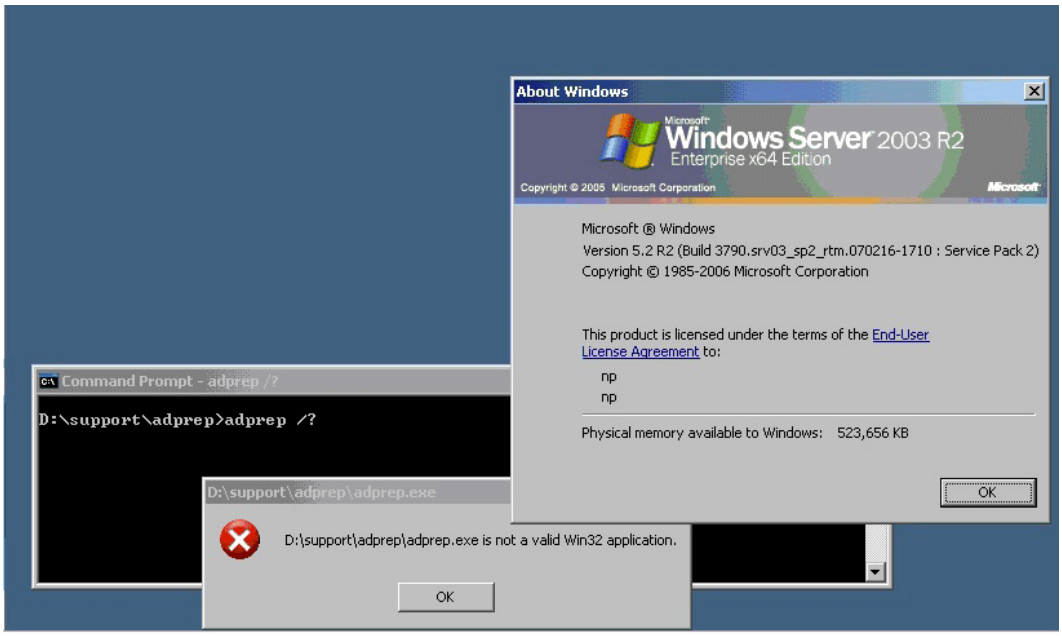

For information about resolving other errors returned by Adprep.exe, see Known issues.

Group membership check against Windows Server 2003 operations master roles

For each command (/forestprep, /domainprep, or /rodcprep), Adprep performs a group membership check to determine whether the specified credential represents an account in certain groups. To perform this check, Adprep contacts the operations master role owner. If the operations master is running Windows Server 2003, you need to specify the /user and /userdomain command line parameters if you run Adprep.exe to ensure the group membership check is performed in all cases.

The /user and /userdomain are new parameters for Adprep.exe in Windows Server 2012. These parameters specify the user account name and user domain, respectively, of the user who runs the adprep command. The Adprep.exe command-line utility blocks specifying one of /userdomain and /user but omitting the other.

However, Adprep operations can also be run as part of an AD DS installation using Windows PowerShell or Server Manager. Those experiences share the same underlying implementation (adprep.dll) as adprep.exe. The Windows PowerShell and Server Manager experiences have their separate credentials input, which does not impose the same requirements as by adprep.exe. Using Windows PowerShell or Server Manager, it is possible to pass a value for /user but not /userdomain to adprep.dll. If /user is specified but /userdomain is not specified, the local machine's domain is used to perform the check. If the machine is not domain joined, group membership cannot be checked.

When group membership cannot be checked, Adprep shows a warning message in the adprep log files and continues:

```
Adprep was unable to check the specified user's group membership. This could happen if
the FSMO role owner <DNS host name of operations master> is running Windows Server 2003
or lower version of Windows.
```

If you run Adprep.exe without specifying the /user and /userdomain parameters and the operations master is running Windows Server 2003, Adprep.exe contacts a domain controller in the domain of the current logon user. If the current logon user is not a domain account, Adprep.exe cannot perform the group membership check. Adprep.exe also cannot perform the group membership check if smartcard credentials are used, even if you do specify both /user and /userdomain.

If Adprep finishes successfully, there is no action required. If Adprep fails during execution with access errors, provide an account with the correct membership. For more information, see Credential requirements to run Adprep.exe and install Active Directory Domain Services.

Syntax for Adprep in Windows Server 2012

Use the following syntax to run adprep separately from an AD DS installation:

```
Adprep.exe /forestprep /forest <forest name> /userdomain <user domain name> /user <user
name> /password *
```

Use /logdsid in the command in order to generate more detailed logging. The adprep.log is located in %windir%\System32\Debug\Adprep\Logs.

Running adprep using smartcard

The Windows Server 2012 version of adprep.exe works using smartcard as credentials, but there is no easy way to specify the smart card credential through the command line. One way to do it is to obtain the smart card credential through PowerShell cmdlet Get-Credential. Then use the user name of the returned PSCredential object, which appears as `@@...`. The password is the PIN of the smart card.

Adprep.exe requires /userdomain if /user is specified. For smartcard credentials, the /userdomain should be the domain of the underlying user account represented by the smartcard.

Adprep /domainprep /gpprep command is not run automatically

The adprep /domainprep /gpprep command is not run as part of AD DS installation. This command sets permissions that are required for Resultant Set of Policy (RSOP) planning mode functionality. For more information about this command, see Microsoft Knowledge Base article 324392. If the command needs to be run in your Active Directory domain, you can run it separately from the AD DS installation. If the command has already been run in preparation of deploying domain controllers that run Windows Server 2003 SP1 or later, the command does not need to be run again.

You can safely add domain controllers that run Windows Server 2012 to an existing domain without running adprep /domainprep /gpprep, but RSOP planning mode will not function properly.

AD DS installation prerequisite validation

The AD DS installation wizard checks that the following prerequisites are met before the installation begins. This provides you with a chance to correct issues that can potentially block installation.

For example, Adprep-related prerequisites include:

- Adprep credential verification: If adprep needs to be run, the installation wizard verifies that the user has sufficient rights to execute the required Adprep operations.
- Schema master availability check: If the installation wizard determines that adprep /forestprep needs to be run, it verifies that the schema master is online and fails otherwise.
- Infrastructure master availability check: If the installation wizard determines that adprep /domainprep needs to be run, it verifies that the infrastructure master is online and fails otherwise.

Other prerequisite checks that are carried forward from the legacy Active Directory Installation Wizard (dcpromo.exe) include:

- Forest name verification: Ensures that the forest name is valid and does not currently exist.
- NetBIOS name verification: Checks that provided NetBIOS name is valid and does not conflict with existing names.
- Component path verification: Verifies that paths for the Active Directory database, logs, and SYSVOL are valid and that there is enough disk space available for them.
- Child domain name verification: Ensures that the parent and new child domain names are valid and that they do not conflict with existing domains.
- Tree domain name verification: Ensures that the specified tree name is valid and that it does not currently exist.

System requirements

System requirements for Windows Server 2012 are unchanged from Windows Server 2008 R2. For more information, see Windows Server 2008 R2 with SP1 System Requirements (http://www.microsoft.com/windowsserver2008/en/us/system-requirements.aspx).

Some features can have additional requirements. For example, the virtual domain controller cloning feature requires that the PDC emulator run Windows Server 2012 and a computer running Windows Server 2012 with the Hyper-V role installed.

Known issues

This section lists some of the known issues that affect AD DS installation in Windows Server 2012. For additional known issues, see Troubleshooting Domain Controller Deployment.

- If WMI access to the schema master is blocked by Windows Firewall when you remotely run adprep /forestprep, the following error is logged in the adprep log at %systemroot%\system32\debug\adprep:

  ```
  Adprep encountered a Win32 error.
  Error code: 0x6ba Error message: The RPC server is unavailable.
  ```

  In this case, you can work around the error by either running adprep /forestprep directly on the schema master, or you can run one of the following commands to allow WMI traffic through Windows Firewall.

  For Windows Server 2008 or later:

  ```
  netsh advfirewall firewall set rule group="windows management
  instrumentation (wmi)" new enable=yes
  ```

  For Windows Server 2003:

  ```
  netsh firewall set service RemoteAdmin enable
  ```

  After adprep finishes you can run either of the following commands to block WMI traffic again:

  ```
  netsh advfirewall firewall set rule group="windows management
  instrumentation (wmi)" new enable=no
  ```

  ```
  netsh firewall set service remoteadmin disable
  ```

- You can type Ctrl + C to cancel the Install-ADDSForest cmdlet. The cancellation stops the installation and any changes that were made to the state of the server are reverted. But after the cancellation command is issued, control is not returned to Windows PowerShell, and the cmdlet can hang indefinitely.
- **Installation of an additional domain controller using smart card credentials fails if the target server is not joined to the domain before installation.**

  The error message returned in this case is:

  Unable to connect to the replication source domain controller *source domain controller name*. (Exception: Logonfailure: unknown user name or bad password)

  If you join the target server to the domain and then perform the installation using a smart card, the installation succeeds.
- **The ADDSDeployment module does not run under 32-bit processes.** If you are automating deployment and configuration of Windows Server 2012 using a script that includes an ADDSDeployment cmdlet and any other cmdlet that does not support native 64-bit processes, the script can fail with an error that indicates the ADDSDeployment cmdlet cannot be found.

In this case, you need to run the ADDSDeployment cmdlet separately from the cmdlet that does not support native 64-bit processes.

- There is a new file system in Windows Server 2012 named Resilient File System. Do not store the Active Directory database, log files, or SYSVOL on a data volume formatted with Resilient File System (ReFS). For more information about ReFS, see Building the next generation file system for Windows: ReFS.
- In Server Manager, servers that run AD DS or other server roles on a Server Core installation and have been upgraded to Windows Server 2012, the server role can appear with red status, even though events and status are collected as expected. Servers that run a Server Core installation of a preliminary release Windows Server 2012 can also be impacted.

Active Directory Domain Services installation hangs if an error prevents critical replication

If the AD DS installation encounters an error during the critical replication phase, the installation can hang indefinitely. For example, if networking errors prevent critical replication from completing, the installation will not proceed.

If you are installing using Server Manager, you may see the installation progress page remain open, but there is no error reported on screen, and the progress may not change for about 15 minutes. If you are using Windows PowerShell, the progress shown in the Windows PowerShell window will not change for more than 15 minutes.

If you experience this problem, check the dcpromo.log file in the %systemroot%\debug folder on the target server. The log file will typically indicate repeated failures to replicate. Some known causes for this problem are:

- Networking problems prevent critical replication between the target server being promoted and the replication source domain controller.

  For example, the dcpromo.log shows:

```
05/02/2012 14:16:46 [INFO] EVENTLOG (Error): NTDS Replication / DS RPC Client :
1963

Internal event: The following local directory service received an exception from
a remote procedure call (RPC) connection. Extensive RPC information was
requested. This is intermediate information and might not contain a possible
cause.

Process ID:

500

Reported error information:

Error value:

Could not find the domain controller for this domain. (1908)

directory service:
```

```
<domain>.com
Extensive error information:
Error value:
A security package specific error occurred. 1825
directory service:
<DC Name>
```

Because the installation process retries critical replication indefinitely, the domain controller installation proceeds if the underlying network problems are resolved. Investigate the networking problem using tools such as ipconfig, nslookup, and netmon as needed. Ensure connectivity exists between the domain controller you are promoting and the replication partner selected during the AD DS installation. Also make sure name resolution is working.

AD DS installation requirements for network connectivity and name resolution are validated during the prerequisite check before the installation begins. But some error conditions can arise in the time after prerequisite validation occurs and before the installation completes, such as if the replication partner becomes unavailable during installation.

- During replica domain controller installation, the local Administrator account of the target server is specified for the installation credentials and the password of the local Administrator account matches the password of a Domain Admin account. In this case, you can complete the installation wizard and begin the installation before you encounter the “Access is denied” failure.

  For example, the dcpromo.log shows:

```
03/30/2012 11:36:51 [INFO] Creating the NTDS Settings object for
this Active Directory Domain Controller on the remote AD DC
DC2.contoso.com...
03/30/2012 11:36:51 [INFO] EVENTLOG (Error): NTDS Replication / DS RPC Client :
1963
Internal event: The following local directory service received an exception from
a remote procedure call (RPC) connection. Extensive RPC information was
requested. This is intermediate information and might not contain a possible
cause.
Process ID:
508
Reported error information:
Error value:
Access is denied. (5)
directory service:
```

```
DC2.contoso.com
```

If the error is caused by specifying a local Administrator account and password, in order to recover you need to reinstall the operating system, perform metadata cleanup of the account for the domain controller that failed to complete installation, and then retry the AD DS installation using Domain Admin credentials. Restarting the server will not correct this error condition because the server will indicate that AD DS is installed even though the installation did not finish successfully.

Active Directory Domain Services Configuration Wizard warns when a non-normalized DNS name is specified

If you create a new domain or forest and you specify a DNS domain name that includes internationalized characters that are not normalized, then the Active Directory Domain Services Configuration Wizard displays a warning that DNS queries for the name can fail. Although the DNS domain name is specified in the Deployment Configuration page, the warning appears on the Prerequisites Check page later in the wizard.

If a DNS domain name is specified using an un-normalized name like füßball.com or BΣTA.com (the normalized versions are: füssball.com and βστα.com), client applications that try to access it with WinHTTP will normalize the name before calling name resolution APIs. If the user types “BΣTA.com” on some dialog, the DNS query will be sent as “βστα.com” and no DNS server will match it with a resource record for “BΣTA.com”. The user will be unable to resolve name.

The following example explains one of the issues that can happen when using an IDN name that is not normalized:

1. The domain using a non-normalized name is created and registered on dns server: füßball.com
2. Machine “nps” is joined to the domain and gets its name registered: nps.füßball.com
3. A client application tries to connect to the server nps.füßball.com
4. The client application tries to resolve the name nps.füßball.com calling name resolution APIs.
5. Due to normalization, the name gets converted to nps.füssball.com and is queried over the wire as nps.füßball.com
6. The client application is unable to resolve the name since the registered name is nps.füßball.com

If the warning appears in the Prerequisites Check page in the Active Directory Domain Services Configuration Wizard, return to the Deployment Configuration page and specify a normalized DNS domain name. If you are installing a new domain using Windows PowerShell, specify a normalized DNS name for the -DomainName option.

For more information about IDNs, see Handling Internationalized Domain Names (IDNs).

Upgrade Domain Controllers to Windows Server 2012

This topic provides background information about Active Directory Domain Services in Windows Server 2012 and explains the process for upgrading domain controllers from Windows Server 2008 or Windows Server 2008 R2.

- What's new in Windows Server 2012?
- What's new in AD DS in Windows Server 2012?
- AD DS server role installation changes
- Deprecated features and behavior changes related to AD DS in Windows Server 2012
- Operating system requirements
- Supported in-place upgrade paths
- Functional level features and requirements
- AD DS interoperability with other server roles and Windows operating systems
- Operations master roles
- Virtualizing domain controllers that run Windows Server 2012
- Administration of Windows Server 2012 servers
- Application compatibility
- Known issues
- Domain controller upgrade workflow

What's new in Windows Server 2012?

New features listed by server role and technology area are listed in the following table. For more whitepapers, video demonstrations, and presentations about other features in Windows Server 2012, see Server and Cloud Platform.

Active Directory Certificate Services (AD CS)	Active Directory Rights Management Services (AD RMS)	BitLocker Drive Encryption
BranchCache	Dynamic Host Configuration protocol (DHCP)	Domain Name System (DNS)
Failover Clustering	File Server Resource Manager	Group Policy
Hyper-V	IP Address Management (IPAM)	Kerberos Authentication
Managed Service Accounts	Networking	Remote Desktop Services
Security Auditing	Server Manager	Smart Cards

TLS/SSL (Schannel SSP)	Windows Deployment Services	Windows PowerShell 3.0

Automatic Maintenance and changes to restart behavior after updates are applied by Windows Update

Prior to the release of Windows 8, Windows Update managed its own internal schedule to check for updates, and to download and install them. It required that the Windows Update Agent was always running in the background, consuming memory and other system resources.

Windows 8 and Windows Server 2012 introduce a new feature called Automatic Maintenance. Automatic Maintenance consolidates many different features that each used to manage its own scheduling and execution logic. This consolidation allows for all these components to use far less system resources, work consistently, respect the new Connected Standby state for new device types, and consume less battery on portable devices.

Because Windows Update is a part of Automatic Maintenance in Windows 8 and Windows Server 2012, its own internal schedule for setting a day and time to install updates is no longer effective. To help ensure consistent and predictable restart behavior for all devices and computers in your enterprise, including those that run Windows 8 and Windows Server 2012, see Microsoft KB article 2885694 (or see October 2013 cumulative rollup 2883201), then configure policy settings described in the WSUS blog post Enabling a more predictable Windows Update experience for Windows 8 and Windows Server 2012 (KB 2885694).

What's new in AD DS in Windows Server 2012?

The following table summarizes the new features for AD DS, with a link to more detailed information where it is available. For a more detailed explanation of some features, including their requirements, see What's New in Active Directory Domain Services (AD DS).

Feature	Description
Active Directory-Based Activation (AD BA); see Volume Activation Overview	Simplifies the task of configuring the distribution and management of volume software licenses.
Active Directory Federation Services (AD FS)	Adds role install via Server Manager, simplified trust-setup, automatic trust management, SAML-protocol support, and more.
Active Directory lost page flush events	NTDS ISAM event 530 with jet error -1119 is logged to detect lost page flush events to Active Directory databases.
Active Directory Recycle Bin User Interface	Active Directory Administrative Center (ADAC) adds GUI management of recycle bin feature originally introduced in Windows Server 2008

Feature	Description
	R2.
Active Directory Replication and Topology Windows PowerShell cmdlets	Supports the creation and management of Active Directory sites, site-links, connection objects, and more using Windows PowerShell.
Dynamic Access Control	New claims-based authorization platform that enhances the legacy access control model.
Fine-Grained Password Policy User Interface	ADAC adds GUI support for the creating, editing and assignment of PSOs originally added in Windows Server 2008.
Group Managed Service Accounts (gMSA)	A new security principal type known as a gMSA. Services running on multiple hosts can run under the same gMSA account.
DirectAccess Offline Domain Join	Extends offline domain-join by including DirectAccess prerequisites.
Rapid deployment via virtual domain controller (DC) cloning	Virtualized DCs can be rapidly deployed by cloning existing virtual domain controllers using Windows PowerShell cmdlets.
RID pool changes	Adds new monitoring events and quotas to safeguard against excessive consumption of the global RID pool. Optionally doubles the size of the global RID pool if the original pool becomes exhausted.
Secure Time service	Enhances security for W32tm by removing secrets from the wire, removing the MD5 hash functions and requiring the server to authenticate with Windows 8 time clients
USN rollback protection for virtualized DCs	Accidentally restoring snapshot backups of virtualized DCs no longer causes USN rollback.
Windows PowerShell History Viewer	Allow administrators to view the Windows PowerShell commands executed when using ADAC.

Automatic Maintenance and changes to restart behavior after updates are applied by Windows Update

Prior to the release of Windows 8, Windows Update managed its own internal schedule to check for updates, and to download and install them. It required that the Windows Update Agent was always running in the background, consuming memory and other system resources.

Windows 8 and Windows Server 2012 introduce a new feature called Automatic Maintenance. Automatic Maintenance consolidates many different features that each used to manage its own scheduling and execution logic. This consolidation allows for all these components to use far less system resources, work consistently, respect the new Connected Standby state for new device types, and consume less battery on portable devices.

Because Windows Update is a part of Automatic Maintenance in Windows 8 and Windows Server 2012, its own internal schedule for setting a day and time to install updates is no longer effective. To help ensure consistent and predictable restart behavior for all devices and computers in your enterprise, including those that run Windows 8 and Windows Server 2012, you can configure the following Group Policy settings:

- **Computer Configuration|Policies|Administrative Templates|Windows Components|Windows Update|Configure Automatic Updates**
- **Computer Configuration|Policies|Administrative Templates|Windows Components|Windows Update|No auto-restart with logged on users**
- **Computer Configuration|Policies|Administrative Templates|Windows Components|Maintenance Scheduler|Maintenance Random Delay**

The following table lists some examples of how to configure these settings to provide desired restart behavior.

Scenario	Recommended configuration(s)
WSUS managed • Install updates once per week • Reboot Fridays at 11PM	Set machines to auto-install, prevent auto-reboot until desired time **Policy**: Configure Automatic Updates (Enabled) Configure automatic updating: 4 – Auto download and schedule the install **Policy**: No auto-restart with logged-on users… (Disabled) **WSUS deadlines**: set to Fridays at 11PM
WSUS managed • Stagger installs across different hours/days	Set target groups for different groups of machines that should be updated together Use above steps for previous scenario Set different deadlines for different target groups
Not WSUS-managed – no support for	**Policy**: Configure Automatic Updates (Enabled)

deadlines • Stagger installs at different times	Configure automatic updating: 4 – Auto download and schedule the install **Registry key:** Enable the registry key discussed in Microsoft KB article 2835627 **Policy:** Automatic Maintenance Random Delay (Enabled) Set **Regular maintenance random delay** to PT6H for 6-hour random delay to provide the following behavior: • Updates will install at the configured maintenance time plus a random delay • Restart for each machine will take place exactly 3 days later Alternatively, set a different maintenance time for each group of machines

For more information about why the Windows engineering team implemented these changes, see Minimizing restarts after automatic updating in Windows Update.

AD DS server role installation changes

In Windows Server 2003 through Windows Server 2008 R2, you ran the x86 or X64 version of the Adprep.exe command-line tool before running the Active Directory Installation Wizard, Dcpromo.exe, and Dcpromo.exe had optional variants to install from media or for unattended installation.

In Windows Server 2012, command-line installations are performed by using the ADDSDeployment Module in Windows PowerShell. GUI-based promotions are performed in Server Manager using a completely new AD DS Configuration Wizard. To simplify the installation process, ADPREP has been integrated into the AD DS installation and runs automatically as needed. The Windows PowerShell–based AD DS Configuration Wizard automatically targets the schema and infrastructure master roles in the domains where Windows Server 2012 DCs are being added, then remotely runs the required ADPREP commands on the relevant domain controllers.

Prerequisite checks in the AD DS Installation Wizard identify potential errors before the installation begins. Error conditions can be corrected to eliminate concerns from a partially complete upgrade. The wizard also exports a Windows PowerShell script that contains all the options that were specified during the graphical installation.

Taken together, the AD DS installation changes in Windows Server 2012 simplify the DC role installation process and reduce the likelihood of administrative errors, especially when you are deploying multiple domain controllers across global regions and domains. More detailed information on GUI and Windows PowerShell-based installations, including command line syntax

and step-by-step wizard instructions, see Install Active Directory Domain Services. For administrators that want to control the introduction of schema changes in an Active Directory forest independent of the installation of Windows Server 2012 DCs in an existing forest, Adprep.exe commands can still be run at an elevated command prompt.

Deprecated features and behavior changes related to AD DS in Windows Server 2012

There are some changes related to AD DS:

- **Deprecation of Adprep32.exe**

  There is only one version of Adprep.exe and it can be run as needed on 64-bit servers that run Windows Server 2008 or later. It can be run remotely, and must be run remotely if that targeted operations master role is hosted on a 32-bit operating system or Windows Server 2003.

- **Deprecation of Dcpromo.exe**

  Dcpromo is deprecated although in Windows Server 2012 only it can still be run with an answer file or command line parameters to give organizations time to transition existing automation to the new Windows PowerShell installation options.

- **LMHash is disabled on user accounts**

  Secure defaults in Security templates on Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012 enable the NoLMHash policy which is disabled in the security templates of Windows 2000 and Windows Server 2003 domain controllers. Disable the NoLMHash policy for LMHash-dependent clients as required, using the steps in KB article 946405.

Beginning with Windows Server 2008, domain controllers also have the following secure default settings, compared to domain controllers that run Windows Server 2003 or Windows 2000.

Encryption type or policy	Windows Server 2008 default	Windows Server 2012 and Windows Server 2008 R2 default	Comment
AllowNT4Crypto	Disabled	Disabled	Third-party Server Message Block (SMB) clients may be incompatible with the secure default settings on domain controllers. In all cases, these settings can be relaxed to allow interoperability, but only at the expense of security. For more information, see article 942564 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=164558) .

DES	Enabled	Disabled	Article 977321 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=177717)
CBT/Extended Protection for Integrated Authentication	N/A	Enabled	See Microsoft Security Advisory (937811) (http://go.microsoft.com/fwlink/?LinkId=164559) and article 976918 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=178251) . Review and install the hotfix in article 977073 (http://go.microsoft.com/fwlink/?LinkId=186394) in the Microsoft Knowledge Base as required.
LMv2	Enabled	Disabled	Article 976918 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=178251)

Operating system requirements

The minimum system requirements for Windows Server 2012 are listed in the following table. For more information about system requirements and pre-installation information, see Installing Windows Server 2012. There are no additional system requirements to install a new Active Directory forest, but you should add sufficient memory to cache the contents of Active Directory database in order to improve performance for domain controllers, LDAP client requests, and Active Directory-enabled applications. If you are upgrading an existing domain controller or adding a new domain controller that runs Windows Server 2012 to an existing forest, review the next section to ensure the server meets disk space requirements.

Processor	1.4 Ghz 64-bit processor
RAM	512 MB
Free disk space requirements	32 GB
Screen resolution	800 x 600 or higher
Miscellaneous	DVD drive, keyboard, Internet access

Disk space requirements for upgrading domain controllers to Windows Server 2012

This section covers disk space requirements only for upgrading domain controllers from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012. For more information about

disk space requirements for upgrading domain controllers to earlier versions of Windows Server, see Disk space requirements for upgrading to Windows Server 2008 or Disk space requirements for upgrading to Windows Server 2008 R2.

Size the disk that hosts the Active Directory database and log files in order to accommodate the custom and application-driven schema extensions, application and administrator-initiated indexes, plus space for the objects and attributes that you will be added to the directory over deployment life of the domain controller (typically 5 to 8 years). Right sizing at deployment time is typically a good investment compared to greater touch costs required to expand disk storage after deployment. For more information, see Capacity Planning for Active Directory Domain Services.

On domain controllers that you plan to upgrade to Windows Server 2012, make sure that the drive that hosts the Active Directory database (NTDS.DIT) has free disk space that represents at least 20% of the NTDS.DIT file before you begin the operating system upgrade. If there is insufficient free disk space on the volume, the upgrade can fail and the upgrade compatibility report returns an error indicating insufficient free disk space:

In this case, you can try an offline defragmentation of the Active Directory database to recapture additional space, and then retry the upgrade. For more information, see Compact the Directory Database File (Offline Defragmentation).

Available SKUs

There are 4 editions of Windows Server 2012: Foundation, Essentials, Standard and Datacenter. The two editions that support the AD DS role are Standard and Datacenter.

In previous releases, Windows Server editions differed in their support of server roles, processor counts and large memory support. The Standard and Datacenter editions of Windows Server 2012 support all features and underlying hardware but vary in their virtualization rights – two virtual instances are allowed for Standard edition and unlimited virtual instances are allowed for Datacenter edition.

Windows client and Windows Server operating systems that are supported to join Windows Server 2012 domains

The following Windows client and Windows Server operating systems are supported for domain member computers with domain controllers that run Windows Server 2012:

- Client operating systems: Windows 8, Windows 7, Windows Vista, Windows XP

  Computers that run Windows 8 are also able to join domains that have domain controllers that run earlier version of Windows Server, including Windows Server 2003 or later. In this case however, some Windows 8 features may require additional configuration or may not be available. For more information about those features and other recommendations for managing Windows 8 clients in downlevel domains, see Running Windows 8 member computers in Windows Server 2003 domains.
- Server operating systems: Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 R2, Windows Server 2003

Supported in-place upgrade paths

Domain controllers that run 64-bit versions of Windows Server 2008 or Windows Server 2008 R2 can be upgraded to Windows Server 2012. You cannot upgrade domain controllers that run Windows Server 2003 or 32-bit versions of Windows Server 2008. To replace them, install domain controllers that run Windows Server 2012 in the domain, and then remove the domain controllers that Windows Server 2003.

If you are running these editions	You can upgrade to these editions
Windows Server 2008 Standard with SP2 OR Windows Server 2008 Enterprise with SP2	Windows Server 2012 Standard OR Windows Server 2012 Datacenter
Windows Server 2008 Datacenter with SP2	Windows Server 2012 Datacenter
Windows Web Server 2008	Windows Server 2012 Standard
Windows Server 2008 R2 Standard with SP1 OR Windows Server 2008 R2 Enterprise with SP1	Windows Server 2012 Standard OR Windows Server 2012 Datacenter
Windows Server 2008 R2 Datacenter with SP1	Windows Server 2012 Datacenter
Windows Web Server 2008 R2	Windows Server 2012 Standard

For more information about supported upgrade paths, see Evaluation Versions and Upgrade Options for Windows Server 2012. Note that you cannot convert a domain controller that runs an evaluation version of Windows Server 2012 directly to a retail version. Instead, install an additional domain controller on a server that runs a retail version and remove AD DS from the domain controller that runs on the evaluation version.

Due to a known issue, you cannot upgrade a domain controller that runs a Server Core installation of Windows Server 2008 R2 to a Server Core installation of Windows Server 2012. The upgrade will hang on a solid black screen late in the upgrade process. Rebooting such DCs exposes an option in boot.ini file to roll back to the previous operating system version. An additional reboot triggers the automatic rollback to the previous operating system version. Until a solution is available, it is recommended that you install a new domain controller running a Server Core installation of Windows Server 2012 instead of in-place upgrading an existing domain controller that runs a Server Core installation of Windows Server 2008 R2. For more information, see KB article 2734222.

Functional level features and requirements

Windows Server 2012 requires a Windows Server 2003 forest functional level. That is, before you can add a domain controller that runs Windows Server 2012 to an existing Active Directory forest,

the forest functional level must be Windows Server 2003 or higher. This means that domain controllers that run Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 can operate in the same forest, but domain controllers that run Windows 2000 Server are not supported and will block installation of a domain controller that runs Windows Server 2012. If the forest contains domain controllers running Windows Server 2003 or later but the forest functional level is still Windows 2000, the installation is also blocked.

Windows 2000 domain controllers must be removed prior to adding Windows Server 2012 domain controllers to your forest. In this case, consider the following workflow:

1. Install domain controllers that run Windows Server 2003 or later. These domain controllers can be deployed on an evaluation version of Windows Server. This step also requires running adprep.exe for that operating system release as a prerequisite.
2. Remove the Windows 2000 domain controllers. Specifically, gracefully demote or forcibly remove Windows Server 2000 domain controllers from the domain and used Active Directory Users and Computers to remove the domain controller accounts for all removed domain controllers.
3. Raise the forest functional level to Windows Server 2003 or higher.
4. Install domain controllers that run Windows Server 2012.
5. Remove domain controllers that run earlier versions of Windows Server.

The new Windows Server 2012 domain functional level enables one new feature: the **KDC support for claims, compound authentication, and Kerberos armoring** KDC administrative template policy has two settings (**Always provide claims** and **Fail unarmored authentication requests**) that require Windows Server 2012 domain functional level. For more information, see Support for claims, compound authentication, and Kerberos armoring.

The Windows Server 2012 forest functional level does not provide any new features, but it ensures that any new domain created in the forest will automatically operate at the Windows Server 2012 domain functional level. The Windows Server 2012 domain functional level does not provide other new features beyond KDC support for claims, compound authentication, and Kerberos armoring. But it ensures that any domain controller in the domain runs Windows Server 2012. For more information about other features that are available at different functional levels, see Understanding Active Directory Domain Services (AD DS) Functional Levels.

After you set the forest functional level to a certain value, you cannot roll back or lower the forest functional level, with the following exceptions: after you raise the forest functional level to Windows Server 2012, you can lower it to Windows Server 2008 R2. If Active Directory Recycle Bin has not been enabled, you can also lower the forest functional level from Windows Server 2012 to Windows Server 2008 R2 or Windows Server 2008 or from Windows Server 2008 R2 back to Windows Server 2008. If the forest functional level is set to Windows Server 2008 R2, it cannot be rolled back, for example, to Windows Server 2003.

After you set the domain functional level to a certain value, you cannot roll back or lower the domain functional level, with the following exceptions: when you raise the domain functional level to Windows Server 2008 R2 or Windows Server 2012, and if the forest functional level is Windows Server 2008 or lower, you have the option of rolling the domain functional level back to Windows Server 2008 or Windows Server 2008 R2. You can lower the domain functional level

only from Windows Server 2012 to Windows Server 2008 R2 or Windows Server 2008 or from Windows Server 2008 R2 to Windows Server 2008. If the domain functional level is set to Windows Server 2008 R2, it cannot be rolled back, for example, to Windows Server 2003.

For more information about features that are available at lower functional levels, see Understanding Active Directory Domain Services (AD DS) Functional Levels.

Beyond functional levels, a domain controller that runs Windows Server 2012 provides additional features that are not available on a domain controller that runs an earlier version of Windows Server. For example, a domain controller that runs Windows Server 2012 can be used for virtual domain controller cloning, whereas a domain controller that runs an earlier version of Windows Server cannot. But virtual domain controller cloning and virtual domain controller safeguards in Windows Server 2012 do not have any functional level requirements.

Note
Microsoft Exchange Server 2013 requires a forest functional level of Windows server 2003 or higher.

AD DS interoperability with other server roles and Windows operating systems

AD DS is not supported on the following Windows operating systems:

- Windows MultiPoint Server
- Windows Server 2012 Essentials

AD DS cannot be installed on a server that also runs the following server roles or role services:

- Hyper-V Server
- Remote Desktop Connection Broker

Operations master roles

Some new features in Windows Server 2012 affect operations master roles:

- The PDC emulator must be running Windows Server 2012 to support cloning virtual domain controllers. There are additional prerequisites for cloning DCs. For more information, see Active Directory Domain Services (AD DS) Virtualization.
- New security principals are created when the PDC emulator runs Windows Server 2012.
- The RID Master has new RID issuance and monitoring functionality. The improvements include better event logging, more appropriate limits, and the ability to - in an emergency - increase the overall RID pool allocation by one bit. For more information, see Managing RID Issuance.

Note
Though they are not operations master roles, another change in AD DS installation is that DNS server role and the global catalog are installed by default on all domain controllers that run Windows Server 2012.

Virtualizing domain controllers that run Windows Server 2012

Improvements in AD DS in Windows Server 2012 enable safer virtualization of domain controllers and the ability to clone domain controllers. Cloning domain controllers in turn enables rapid deployment of additional domain controllers in a new domain and other benefits. For more information, see Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100).

Administration of Windows Server 2012 servers

Use the Remote Server Administration Tools for Windows 8 to manage domain controllers and other servers that run Windows Server 2012. You can run the Windows Server 2012 Remote Server Administration Tools on a computer that runs Windows 8.

Application compatibility

The following table covers common Active Directory-integrated Microsoft applications. The table covers what versions of Windows Server that the applications can be installed on and whether the introduction of Windows Server 2012 DCs affects application compatibility.

Product	Notes
Microsoft Configuration Manager 2007	Configuration Manager 2007 with SP2 (includes Configuration Manager 2007 R2 and Configuration Manager 2007 R3): • Windows 8 Pro • Windows 8 Enterprise • Windows Server 2012 Standard • Windows Server 2012 Datacenter **Note** Though these will be fully supported as clients, there is no plan to add support for deploying these as operating systems by using the Configuration Manager 2007 operating system deployment feature. Also, no site servers or site systems will be supported on any SKU of Windows Server 2012.
Microsoft SharePoint 2007	Microsoft Office SharePoint Server 2007 is not supported for installation on Windows Server

Product	Notes
	2012.
Microsoft SharePoint 2010	SharePoint 2010 Service Pack 2 is required to install and operate SharePoint 2010 on Windows Server 2012 Servers SharePoint 2010 Foundation Service Pack 2 is required to install and operate SharePoint 2010 Foundation on Windows Server 2012 Servers The SharePoint Server 2010 (without service packs) installation process fails on Windows Server 2012 The SharePoint Server 2010 prerequisite installer (PrerequisiteInstaller.exe) fails with error “This program has compatibility issues.” Clicking “Run the program without getting help” displays the error “| Verifying if SharePoint can be installed | SharePoint Server 2010 (without service packs) cannot be installed on Windows Server 2012.
Microsoft SharePoint 2013	Minimum requirements for a database server in a farm The 64-bit edition of Windows Server 2008 R2 Service Pack 1 (SP1) Standard, Enterprise, or Datacenter or the 64-bit edition of Windows Server 2012 Standard or Datacenter Minimum requirements for a single server with built-in database: The 64-bit edition of Windows Server 2008 R2 Service Pack 1 (SP1) Standard, Enterprise, or Datacenter or the 64-bit edition of Windows Server 2012 Standard or Datacenter Minimum requirements for front-end web servers and application servers in a farm: The 64-bit edition of Windows Server 2008 R2 Service Pack 1 (SP1) Standard, Enterprise, or Datacenter or the 64-bit edition of Windows Server 2012 Standard or Datacenter.
Microsoft System Center Configuration	System Center 2012 Configuration Manager

Product	Notes
Manager 2012	Service Pack 1: Microsoft will add the following operating systems to our client support matrix with the release of Service Pack 1: • Windows 8 Pro • Windows 8 Enterprise • Windows Server 2012 Standard • Windows Server 2012 Datacenter All site server roles - including site servers, SMS providers, and management points - can be deployed to servers with the following operating system editions: • Windows Server 2012 Standard • Windows Server 2012 Datacenter
Microsoft Lync Server 2013	Lync Server 2013 requires with Windows Server 2008 R2 or Windows Server 2012. It cannot be run on a Server Core installation. It can be run on virtual servers.
Lync Server 2010	Lync Server 2010 can be installed on a new (not upgraded) installation Windows Server 2012 if October 2012 cumulative updates for Lync Server are installed. Upgrading the operating system to Windows Server 2012 for an existing installation of Lync Server 2010 is not supported. Microsoft Lync Server 2010 Group Chat Server is also not supported on Windows Server 2012.
System Center 2012 Endpoint Protection	System Center 2012 Endpoint Protection Service Pack 1 will update the client support matrix to include the following operating systems • Windows 8 Pro • Windows 8 Enterprise • Windows Server 2012 Standard • Windows Server 2012 Datacenter
System Center 2012 Forefront Endpoint Protection	FEP 2010 with Update Rollup 1 will update the client support matrix to include the following operating systems:

Product	Notes
	• Windows 8 Pro • Windows 8 Enterprise • Windows Server 2012 Standard • Windows Server 2012 Datacenter
Forefront Threat Management Gateway (TMG)	TMG is supported to run only on Windows Server 2008 and Windows Server 2008 R2. For more information, see System requirements for Forefront TMG.
Windows Server Update Services	This release of WSUS already supports Windows 8-based computers or Windows Server 2012-based computers as clients.
Windows Server Update Services 3.0	Update KB article 2734608 lets servers that are running Windows Server Update Services (WSUS) 3.0 SP2 provide updates to computers that are running Windows 8 or Windows Server 2012: **Note** Customers with standalone WSUS 3.0 SP2 environments or System Center Configuration Manager 2007 Service Pack 2 environments with WSUS 3.0 SP2 require 2734608 to properly manage Windows 8-based computers or Windows Server 2012-based computers as clients.
Exchange 2013	Windows Server 2012 Standard and Datacenter are supported for the following roles: schema master, global catalog server, domain controller, mailbox and client access server role Forest Functional Level: Windows Server 2003 or higher Source: Exchange 2013 System Requirements
Exchange 2010	Source: Exchange 2010 Service Pack 3 Exchange 2010 with Service Pack 3 can be installed on Windows Server 2012 member

Product	Notes
	servers. Exchange 2010 System Requirements lists the latest supported schema master, global catalog and domain controller as Windows Server 2008 R2. Forest Functional Level: Windows Server 2003 or higher
SQL Server 2012	Source: KB 2681562 SQL Server 2012 RTM is supported on Windows Server 2012.
SQL Server 2008 R2	Source: KB 2681562 Requires SQL Server 2008 R2 with Service Pack 1 or later to install on Windows Server 2012.
SQL Server 2008	Source: KB 2681562 Requires SQL Server 2008 with Service Pack 3 or later to install on Windows Server 2012.
SQL Server 2005	Source: KB 2681562 Not supported to install on Windows Server 2012.

Known issues

The following table lists known issues related to AD DS installation.

KB article number and title	Technology area impacted	Issue/description
2830145: SID S-1-18-1 and SID S-1-18-2 can't be mapped on Windows 7 or Windows Server 2008 R2-based computers in a domain environment	AD DS Management/App compat	Applications that map SID S-1-18-1 and SID S-1-18-2, which are new in Windows Server 2012, may fail because the SIDs cannot be resolved on Windows 7-based or Windows Server 2008 R2-based computers. To resolve this issue, install the hotfix on the Windows 7-based and Windows Server 2008 R2-based computers in the domain.

2737129: Group Policy preparation is not performed when you automatically prepare an existing domain for Windows Server 2012	AD DS Installation	Adprep /domainprep /gpprep is not automatically run as part of installing the first DC that runs Windows Server 2012 in a domain. If it has never been run previously in the domain, it must be run manually.
2737416: Windows PowerShell-based domain controller deployment repeats warnings	AD DS Installation	Warnings can appear during prerequisite validation and then reappear during the installation.
2737424: "Format of the specified domain name is invalid" error when you try to remove Active Directory Domain Services from a domain controller	AD DS Installation	This error appears if you are removing the last DC in a domain where pre-created RODC accounts still exist. This affects Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 R2.
2737463: Domain controller does not start, c00002e2 error occurs, or "Choose an option" is displayed	AD DS Installation	A DC does not start because an administrator used Dism.exe, Pkgmgr.exe, or Ocsetup.exe to remove the DirectoryServices-DomainController role.
2737516: IFM verification limitations in Windows Server 2012 Server Manager	AD DS Installation	IFM verification can have limitations as explained in the KB article.
2737535: Install-AddsDomainController cmdlet returns parameter set error for RODC	AD DS Installation	You can receive an error when you try to attach a server to an RODC account if you specify arguments that are already populated on the pre-created RODC account.
2737560: "Unable to perform Exchange schema conflict check" error, and prerequisites check fails	AD DS Installation	Prerequisite check fails when you configure the first Windows Server 2012 DC in an existing domain because DCs are missing the SeServiceLogonRight for Network Service or because WMI or DCOM protocols are blocked.

2737797: AddsDeployment module with the -Whatif argument shows incorrect DNS results	AD DS Installation	The –WhatIf parameter shows DNS server will not be installed but it will be.
2737807: The Next button is not available on the Domain Controller Options page	AD DS Installation	The Next button is disabled on the Domain Controller Options page because the IP address of the target DC does not map to an existing subnet or site, or because the DSRM password is not typed and confirmed correctly.
2737935: Active Directory installation stalls at the "Creating the NTDS settings object" stage	AD DS Installation	The installation hangs because the local Administrator password matches the domain Administrator password, or because networking problems prevent critical replication from completing.
2738060: "Access is denied" error message when you create a child domain remotely by using Install-AddsDomain	AD DS Installation	You receive the error when you run Install-ADDSDomain with the Invoke-Command cmdlet if the DNSDelegationCredential has a bad password.
2738697: "The server is not operational" domain controller configuration error when you configure a server by using Server Manager	AD DS Installation	You receive this error when you try to install AD DS on a workgroup computer because NTLM authentication is disabled.
2738746: You receive access denied errors after you log on to a local administrator domain account	AD DS Installation	When you log on using a local Administrator account rather than the built-in Administrator account and then create a new domain, the account is not added to the Domain Admins group.
2743345: "The system cannot find the file specified" Adprep /gpprep error, or tool crashes	AD DS Installation	You receive this error when you run adprep /gpprep because the infrastructure master is implements a disjoint namespace

2743367: Adprep "not a valid Win32 application" error on Windows Server 2003, 64-bit version	AD DS Installation	You receive this error because Windows Server 2012 Adprep cannot be run on Windows Server 2003.
2753560: ADMT 3.2 and PES 3.1 installation errors on Windows Server 2012	ADMT	ADMT 3.2 cannot be installed on Windows Server 2012 by design.
2750857: DFS Replication diagnostic reports do not display correctly in Internet Explorer 10	DFS Replication	DFS Replication diagnostic report does not display correctly because of changes in Internet Explorer 10.
2741537: Remote Group Policy updates are visible to users	Group Policy	This is due to scheduled tasks run in the context of each user who is logged on. The Windows Task Scheduler design requires an interactive prompt in this scenario.
2741591: ADM files are not present in SYSVOL in the GPMC Infrastructure Status option	Group Policy	GP replication can report “replication in progress” because GPMC Infrastructure Status does not follow customized filtering rules.
2737880: "The service cannot be started" error during AD DS configuration	Virtual DC cloning	You receive this error while installing or removing AD DS, or cloning, because the DS Role Server service is disabled.
2742836: Two DHCP leases are created for each domain controller when you use the VDC cloning feature	Virtual DC cloning	This happens because the cloned domain controller received a lease before cloning and again when cloning was complete.
2742844: Domain controller cloning fails and the server restarts in DSRM in Windows Server 2012	Virtual DC cloning	The cloned DC starts in DSRM because cloning failed for any of a variety of reasons listed in the KB article.
2742874: Domain controller cloning does not re-create all service principal names	Virtual DC cloning	Some three-part SPNs are not recreated on the cloned DC because of a limitation of the domain rename process.

2742908: "No logon servers are available" error after cloning domain controller	Virtual DC cloning	You receive this error when you try to log on after cloning a virtualized DC because cloning failed and the DC is started in DSRM. Log on as .\administrator to troubleshoot the cloning failure.
2742916: Domain controller cloning fails with error 8610 in dcpromo.log	Virtual DC cloning	Cloning fails because the PDC emulator has not performed inbound replication of the domain partition, likely because the role was transferred.
2742927: "Index was out of range" New-AdDcCloneConfig error	Virtual DC cloning	You receive the error after you run New-ADDCCloneConfigFile cmdlet while cloning virtual DCs, either because the cmdlet was not run from an elevated command prompt or because your access token does not contain the Administrators group.
2742959: Domain controller cloning fails with error 8437: "invalid parameter was specified for this replication operation"	Virtual DC cloning	Cloning failed because an invalid clone name or a duplicate NetBIOS name was specified.
2742970: DC Cloning fails with no DSRM, duplicate source and clone computer	Virtual DC cloning	The cloned virtual DC boots in Directory Services Repair Mode (DSRM), using a duplicate name as the source DC because the DCCloneConfig.xml file was not created in the correct location or because the source DC was rebooted before cloning.
2743278: Domain controller cloning error 0x80041005	Virtual DC cloning	The cloned DC boots into DSRM because only one WINS server was specified. If any WINS server is specified, both Preferred and Alternate WINS servers must be specified.
2745013: "Server is not	Virtual DC cloning	You receive this error after you

operational" error message if you run New-AdDcCloneConfigFile in Windows Server 2012		run the New-ADDCCloneConfigFile cmdlet because the server cannot contact a global catalog server.
2747974: Domain controller cloning event 2224 provides incorrect guidance	Virtual DC cloning	Event ID 2224 incorrectly states that managed service accounts must be removed before cloning. Standalone MSAs must be removed but Group MSAs do not block cloning.
2748266: You cannot unlock a BitLocker-encrypted drive after you upgrade to Windows 8	BitLocker	You receive an “Application not found” error when you try to unlock a drive on a computer that was upgraded from Windows 7.

Domain controller upgrade workflow

1. Verify the target server meets system requirements.
2. Verify Application compatibility.
3. Verify security settings. For more information, see Deprecated features and behavior changes related to AD DS in Windows Server 2012 and Secure default settings in Windows Server 2008 and Windows Server 2008 R2.
4. Check connectivity to the target server from the computer where you plan to run the installation.
5. Check for availability of necessary operation master roles. To install the first DC that runs Windows Server 2012 in an existing domain and forest, the machine where you run the installation needs connectivity to the schema master in order to run adprep /forestprep and the infrastructure master in order to run adprep /domainprep. To install the first DC in a domain where the forest schema is already extended, you only need connectivity to infrastructure master. To install or remove a domain in an existing forest, you need connectivity to the domain naming master. Any domain controller installation also requires connectivity to the RID master. If you are installing the first read-only domain controller in an existing forest, you need connectivity to the infrastructure master for each application directory partition, also known as a non-domain naming context or NDNC.
6. Be sure to supply the necessary credentials to run the AD DS installation.

Installation action	Credential requirements
Install a new forest	Local Administrator on the target server
Install a new domain in an existing forest	Enterprise Admins
Install an additional DC in an existing domain	Domain Admins
Run adprep /forestprep	Schema Admins, Enterprise Admins, and Domain Admins
Run adprep /domainprep	Domain Admins
Run adprep /domainprep /gpprep	Domain Admins
Run adprep /rodcprep	Enterprise Admins

You can delegate permissions to install AD DS. For more information, see Installation Management Tasks.

Steps-by-step instructions to promote new and replica Windows Server 2012 domain controllers using Windows PowerShell cmdlets and Server Manager can be found in the following links:

- Install Active Directory Domain Services (Level 100)
- Install a New Windows Server 2012 Active Directory Forest (Level 200)
- Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200)
- Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200)
- Install a Windows Server 2012 Active Directory Read-Only Domain Controller (RODC) (Level 200)
- Step-by-Step Guide for Setting Up Windows Server 2012 Domain Controller (en-US)

See Also

Windows Server 2012 Evaluation Resources

Windows Server 2012 Evaluation Guide

Install and Deploy Windows Server 2012

AD DS Simplified Administration

This topic explains the new capabilities and benefits of Windows Server 2012 domain controller deployment and administration, and the differences between previous operating system DC deployment and the new Windows Server 2012 implementation.

Windows Server 2012 introduces the next generation of Active Directory Domain Services Simplified Administration, and is the most radical domain re-envisioning since Windows 2000 Server. AD DS Simplified Administration takes lessons learned from twelve years of Active

Directory and makes a more supportable, more flexible, more intuitive administrative experience for architects and administrators. This meant creating new versions of existing technologies as well as extending the capabilities of components released in Windows Server 2008 R2.

AD DS Simplified Administration is a reimagining of domain deployment.

- AD DS role deployment is now part of the new Server Manager architecture and allows remote installation
- The AD DS deployment and configuration engine is now Windows PowerShell, even when using the new AD DS Configuration Wizard
- Schema extension, forest preparation, and domain preparation are automatically part of domain controller promotion and no longer require separate tasks on special servers such as the Schema Master
- Promotion now includes prerequisite checking that validates forest and domain readiness for the new domain controller, lowering the chance of failed promotions
- Active Directory module for Windows PowerShell now includes cmdlets for replication topology management, Dynamic Access Control, and other operations
- The Windows Server 2012 forest functional level does not implement new features and domain functional level is required only for a subset of new Kerberos features, relieving administrators of the frequent need for a homogenous domain controller environment
- Full support added for Virtualized Domain Controllers, to include automated deployment and rollback protection

For more information about virtualized domain controllers, see Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100).

In addition, there are many administrative and maintenance improvements:

- The Active Directory Administrative Center includes a graphical Active Directory Recycle Bin, Fine-Grained Password Policy management, and Windows PowerShell history viewer
- The new Server Manager has AD DS-specific interfaces into performance monitoring, best practice analysis, critical services, and the event logs
- Group Managed Service Accounts support multiple computers using the same security principals
- Improvements in Relative Identifier (RID) issuance and monitoring for better manageability in mature Active Directory domains

Finally, AD DS profits from other new features included in Windows Server 2012, such as:

- NIC teaming and Datacenter Bridging
- DNS Security and faster AD-integrated zone availability after boot
- Hyper-V reliability and scalability improvements
- BitLocker Network Unlock
- Additional Windows PowerShell component administration modules

Technical Overview

ADPREP Integration

Active Directory forest schema extension and domain preparation now integrate into the domain controller configuration process. If you promote a new domain controller into an existing forest, the process detects upgrade status and the schema extension and domain preparation phases occur automatically. The user installing the first Windows Server 2012 domain controller must still be an Enterprise Admin and Schema Admin or provide valid alternate credentials.

Adprep.exe remains on the DVD for separate forest and domain preparation. The version of the tool included with Windows Server 2012 is backwards compatible to Windows Server 2008 x64 and Windows Server 2008 R2. Adprep.exe also supports remote forestprep and domainprep, just like the ADDSDeployment-based domain controller configuration tools.

For information about Adprep and previous operating system forest preparation, see Running Adprep (Windows Server 2008 R2).

Server Manager AD DS Integration

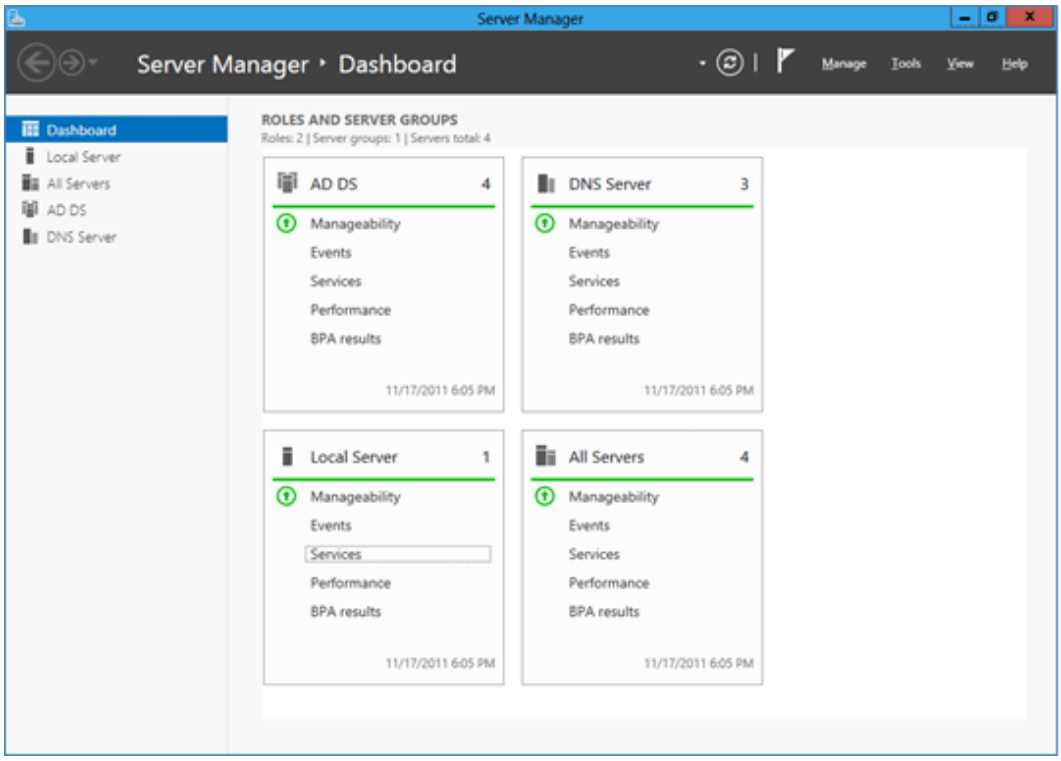

Server Manager acts as a hub for server management tasks. Its dashboard-style appearance periodically refreshes views of installed roles and remote server groups. Server Manager provides centralized management of local and remote servers, without the need for console access.

Active Directory Domain Services is one of those hub roles; by running Server Manager on a domain controller or the Remote Server Administration Tools on a Windows 8, you see important recent issues on domain controllers in your forest.

These views include:

- Server availability
- Performance monitor alerts for high CPU and memory usage
- The status of Windows services specific to AD DS
- Recent Directory Services-related warning and error entries in the event log
- Best Practice analysis of a domain controller against a set of Microsoft-recommended rules

Active Directory Administrative Center Recycle Bin

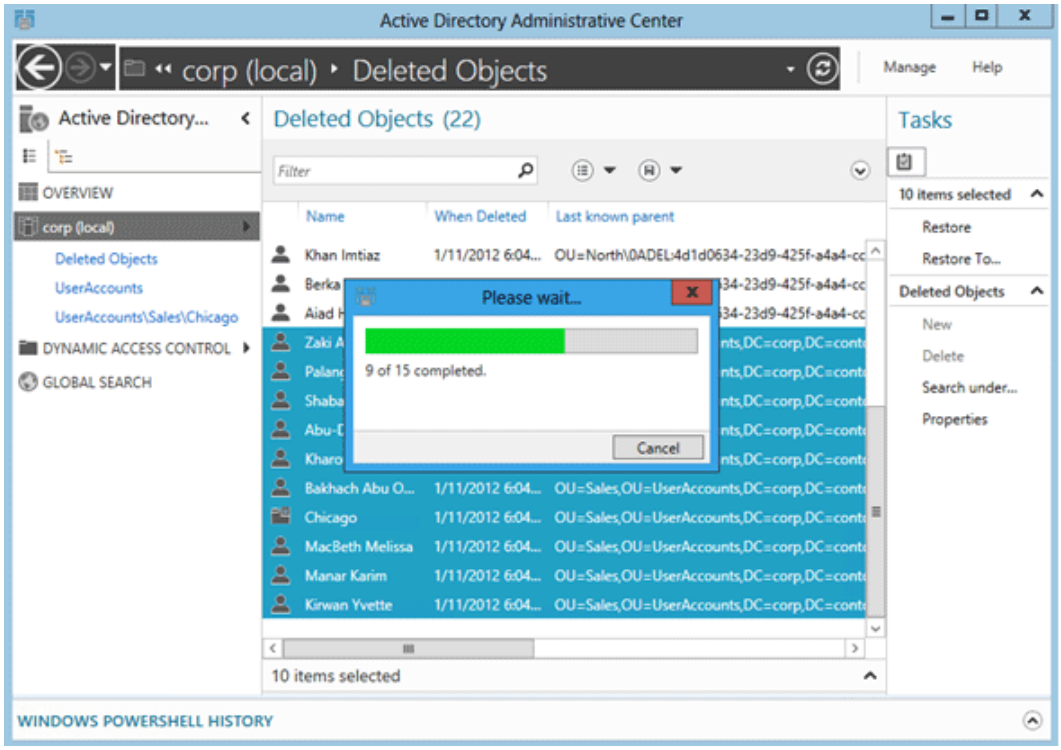

Windows Server 2008 R2 introduced the Active Directory Recycle Bin, which recovers deleted Active Directory objects without restoring from backup, restarting the AD DS service, or rebooting domain controllers.

Windows Server 2012 enhances the existing Windows PowerShell-based restore capabilities with a new graphical interface in the Active Directory Administrative Center. This allows administrators to enable the Recycle Bin and locate or restore deleted objects in the domain contexts of the forest, all without directly running Windows PowerShell cmdlets. The Active Directory Administrative Center and Active Directory Recycle Bin still use Windows PowerShell under the covers, so previous scripts and procedures are still valuable.

For information about the Active Directory Recycle Bin, see Active Directory Recycle Bin Step-by-Step Guide (Windows Server 2008 R2).

Active Directory Administrative Center Fine-Grained Password Policy

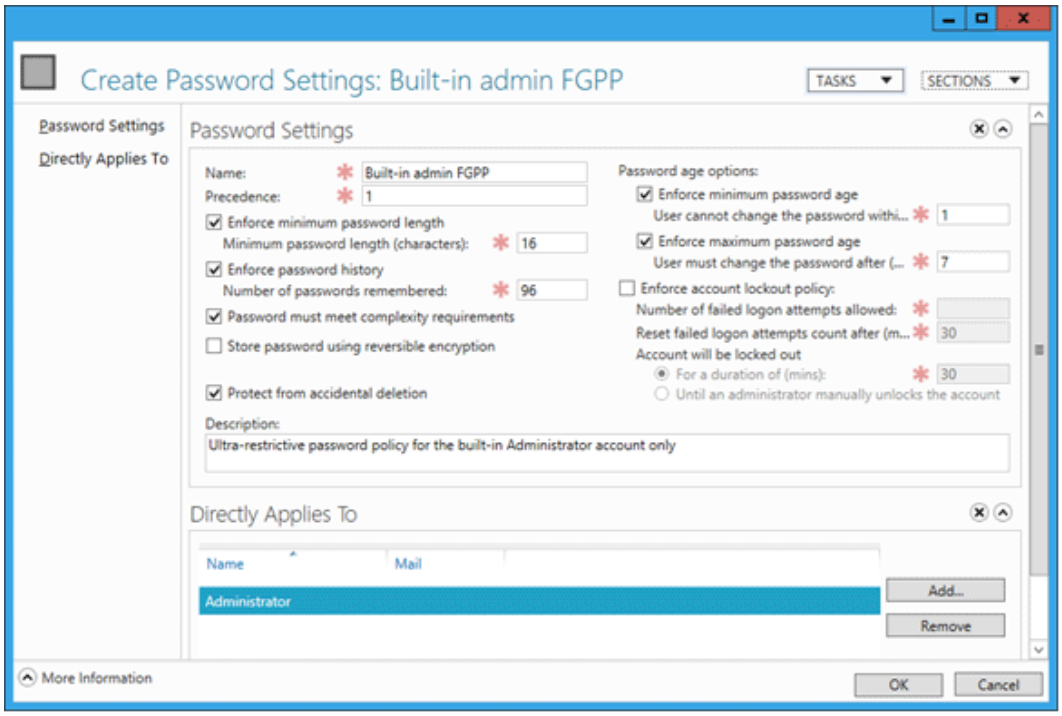

Windows Server 2008 introduced the Fine-Grained Password policy, which allows administrators to configure multiple password and account lockout policies per domain. This allows domains a flexible solution to enforce more or less restrictive password rules, based on users and groups. It had no managerial interface and required administrators to configure it using Ldp.exe or Adsiedit.msc. Windows Server 2008 R2 introduced the Active Directory module for Windows PowerShell, which granted administrators a command-line interface to FGPP.

Windows Server 2012 brings a graphical interface to Fine-Grained Password Policy. The Active Directory Administrative Center is the home of this new dialog, which brings simplified FGPP management to all administrators.

For information about the Fine-Grained Password Policy, see AD DS Fine-Grained Password and Account Lockout Policy Step-by-Step Guide (Windows Server 2008 R2).

Active Directory Administrative Center Windows PowerShell History Viewer

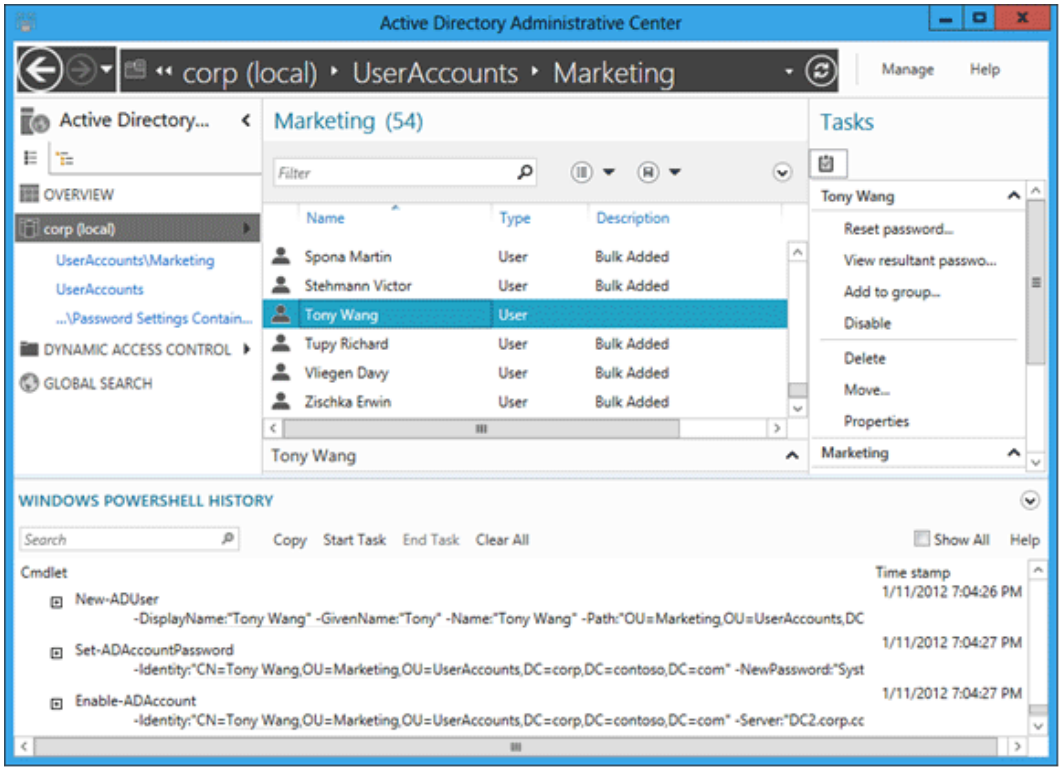

Windows Server 2008 R2 introduced the Active Directory Administrative Center, which superseded the older Active Directory Users and Computers snap-in created in Windows 2000. The Active Directory Administrative Center creates a graphical administrative interface to the then-new Active Directory module for Windows PowerShell.

While the Active Directory module contains over a hundred cmdlets, the learning curve for an administrator can be steep. Since Windows PowerShell integrates heavily into the strategy of Windows administration, the Active Directory Administrative Center now includes a viewer that enables you to see the cmdlet execution in the graphical interface. You can search, copy, clear history, and add notes with a simple interface. The intent is for an administrator to use the graphical interface to create and modify objects, and then review them in the history viewer to learn more about Windows PowerShell scripting and modify the examples.

AD Replication Windows PowerShell

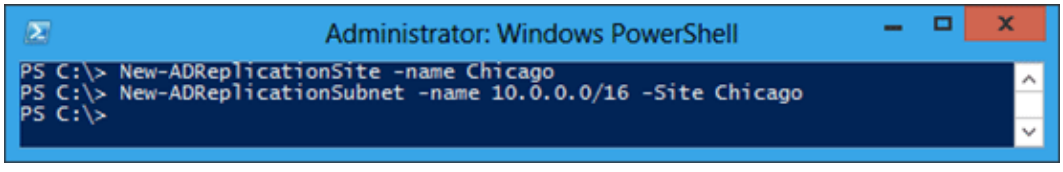

Windows Server 2012 adds additional Active Directory replication cmdlets to the Active Directory Windows PowerShell module. These allow configuration of new or existing sites, subnets, connections, site links, and bridges. They also return Active Directory replication metadata, replication status, queuing, and up-to-dateness version vector information. The introduction of the replication cmdlets - combined with the deployment and other existing AD DS cmdlets - makes it possible to administer a forest using Windows PowerShell alone. This creates new opportunities for administrators wishing to provision and manage Windows Server 2012 without a graphical interface, which then reduces the operating system's attack surface and servicing requirements. This is especially important when deploying servers into high security networks such as Secret Internet Protocol Router (SIPR) and corporate DMZs.

For more information about AD DS site topology and replication, see the Windows Server Technical Reference.

RID Management and Issuance Improvements

Windows 2000 Active Directory introduced the RID Master, which issues pools of relative identifiers to domain controllers, in order to create security identifiers (SIDs) of security trustees like users, groups, and computers. By default, this global RID space is limited to 2^{30} (or 1,073,741,823) total SIDs created in a domain. SIDs cannot return to the pool or reissue. Over time, a large domain may begin to run low on RIDs, or accidents may lead to unnecessary RID depletion and eventual exhaustion.

Windows Server 2012 addresses a number of RID issuance and management issues uncovered by customers and Microsoft Customer Support as AD DS matured since the creation of the first Active Directory domains in 1999. These include:

- Periodic RID consumption warnings are written to the event log
- Events log when an administrator invalidates a RID pool
- A maximum cap on the RID policy RID Block Size is now enforced
- Artificial RID ceilings are now enforced and logged when the global RID space is low, allowing an administrator to take action before the global space is exhausted
- The global RID space can now be increased by one bit, doubling the size to 2^{31} (2,147,483,648 SIDs)

For more information about RIDs and the RID Master, review How Security Identifiers Work.

New AD DS Deployment Architecture

AD DS Role Deployment and Management Architecture

Server Manager and ADDSDeployment Windows PowerShell rely on the following core assemblies for functionality when deploying or managing the AD DS role:

- Microsoft.ADroles.Aspects.dll
- Microsoft.ADroles.Instrumentation.dll
- Microsoft.ADRoles.ServerManager.Common.dll

- Microsoft.ADRoles.UI.Common.dll
- Microsoft.DirectoryServices.Deployment.Types.dll
- Microsoft.DirectoryServices.ServerManager.dll
- Addsdeployment.psm1
- Addsdeployment.psd1

Both rely on Windows PowerShell and its remote invoke-command for remote role installation and configuration.

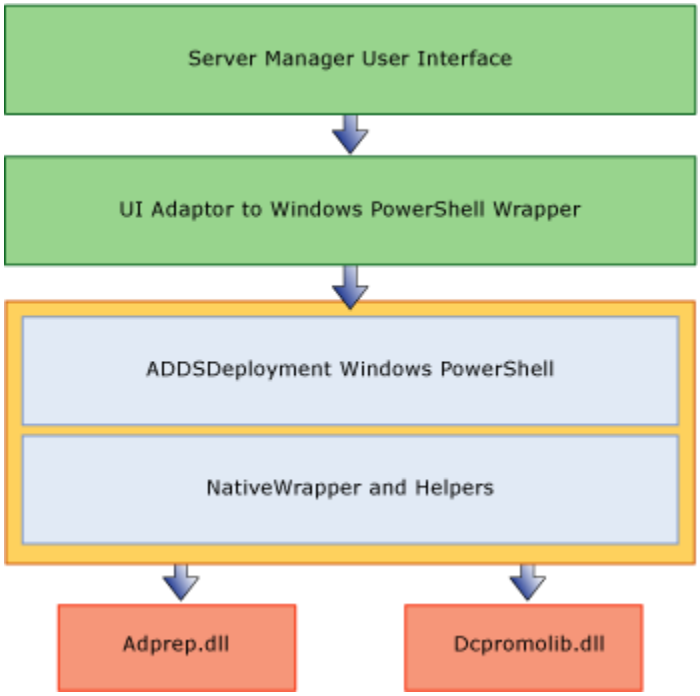

Windows Server 2012 also refactors a number of previous promotion operations out of LSASS.EXE, as part of:

- DS Role Server Service (DsRoleSvc)
- DSRoleSvc.dll (loaded by DsRoleSvc service)

This service must be present and running in order to promote, demote, or clone virtual domain controllers. AD DS role installation adds this service and sets a start type of Manual, by default. Do not disable this service.

ADPrep and Prerequisite Checking Architecture

Adprep no longer requires running on the schema master. It can be run remotely from a computer that runs Windows Server 2008 x64 or later.

Note

Adprep uses LDAP to import Schxx.ldf files and does not automatically reconnect if the connection to the schema master is lost during import. As part of the import process, the

schema master is set in a specific mode and automatic reconnection is disabled because if LDAP reconnects after the connection is lost, the re-established connection would not be in the specific mode. In that case, the schema would not be updated correctly.

Prerequisite checking ensures that certain conditions are true. These conditions are required for successful AD DS installation. If some required conditions are not true, they can be resolved before continuing the installation. It also detects that a forest or domain are not yet prepared, so that the Adprep deployment code runs automatically.

ADPrep Executables, DLLs, LDFs, files

- ADprep.dll
- Ldifde.dll
- Csvde.dll
- Sch14.ldf - Sch56.ldf
- Schupgrade.cat
- *dcpromo.csv

The AD Preparation code formerly housed in ADprep.exe is refactored into adprep.dll. This allows both ADPrep.exe and the ADDSDeployment Windows PowerShell module to use the library for the same tasks and have the same capabilities. Adprep.exe is included with the installation media but automated processes do not call it directly - only an Administrator runs it manually. It can only run on Windows Server 2008 x64 and later operating systems. Ldifde.exe and csvde.exe also have refactored versions as DLLs that are loaded by the preparation process. Schema extension still uses the signature-verified LDF files, like in previous operating system versions.

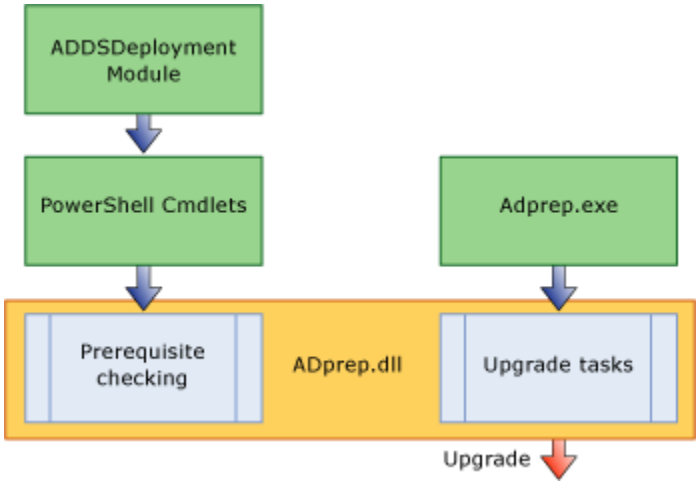

Important

There is no 32-bit Adprep32.exe tool for Windows Server 2012. You must have at least one Windows Server 2008 x64, Windows Server 2008 R2, or Windows Server 2012 computer, running as a domain controller, member server, or in a workgroup, to prepare the forest and domain. Adprep.exe does not run on Windows Server 2003 x64.

Prerequisite Checking

The prerequisite checking system built into ADDSDeployment Windows PowerShell managed code works in different modes, based on the operation. The tables below describe each test, when it is used, and an explanation of how and what it validates. These tables may be useful if there are issues where the validation fails and the error is not sufficient to troubleshoot the problem.

These tests log in the **DirectoryServices-Deployment** operational event log channel under the Task Category **Core**, always as Event ID **103**.

Prerequisite Windows PowerShell

There are ADDSDeployment Windows PowerShell cmdlets for all of the domain controller deployment cmdlets. They have approximately the same arguments as their associated cmdlets.

- Test-ADDSDomainControllerInstallation
- Test-ADDSDomainControllerUninstallation
- Test-ADDSDomainInstallation
- Test-ADDSForestInstallation
- Test-ADDSReadOnlyDomainControllerAccountCreation

There is no need to run these cmdlets, ordinarily; they already automatically execute with the deployment cmdlets by default.

Prerequisite Tests

Test Name	Protocols used	Explanation and notes
VerifyAdminTrusted ForDelegationProvider	LDAP	Validates that you have the "Enable computer and user accounts to be trusted for delegation" (SeEnableDelegationPrivilege) privilege on the existing partner domain controller. This requires access to your constructed tokenGroups attribute. Not used when contacting Windows Server 2003 domain controllers. You must manually confirm this privilege prior to promotion
VerifyADPrep Prerequisites (forest)	LDAP	Discovers and contacts the Schema Master using the rootDSE namingContexts attribute and Schema naming context fsmoRoleOwner attribute. Determines which preparatory operations (forestprep, domainprep, or rodcprep) are required for AD DS installation. Validates the schema objectVersion is expected

		and if it requires further extension.
VerifyADPrep Prerequisites (domain and RODC)	LDAP	Discovers and contacts the Infrastructure Master using the rootDSE namingContexts attribute and the Infrastructure container fsmoRoleOwner attribute. In the case of an RODC installation, this test discovers the domain naming master and make sure it is online.
CheckGroup Membership	LDAP, RPC over SMB (LSARPC)	Validate the user is a member of Domain Admins or Enterprise Admins group, depending on the operation (DA for adding or demoting a domain controller, EA for adding or removing a domain)
CheckForestPrep GroupMembership	LDAP, RPC over SMB (LSARPC)	Validate the user is a member of Schema Admins and Enterprise Admins groups and has the Manage Audit and Security Event Logs (SesScurityPrivilege) privilege on the existing domain controllers
CheckDomainPrep GroupMembership	LDAP, RPC over SMB (LSARPC)	Validate the user is a member of Domain Admins group and has the Manage Audit and Security Event Logs (SesScurityPrivilege) privilege on the existing domain controllers
CheckRODCPrep GroupMembership	LDAP, RPC over SMB (LSARPC)	Validate the user is a member of Enterprise Admins group and has the Manage Audit and Security Event Logs (SesScurityPrivilege) privilege on the existing domain controllers
VerifyInitSync AfterReboot	LDAP	Validate that the Schema Master has replicated at least once since it restarted by setting a dummy value on rootDSE attribute becomeSchemaMaster
VerifySFUHotFix Applied	LDAP	Validate the existing forest schema does not contain known problem SFU2 extension for the UID attribute with OID 1.2.840.113556.1.4.7000.187.102 (http://support.microsoft.com/kb/821732)
VerifyExchange SchemaFixed	LDAP, WMI, DCOM, RPC	Validate the existing forest schema does not still contain problem Exchange 2000 extensions ms-Exch-Assistant-Name, ms-Exch-LabeledURI, and ms-Exch-House-Identifier (http://support.microsoft.com/kb/314649)

VerifyWin2KSchema Consistency	LDAP	Validate the existing forest schema has consistent (not incorrectly modified by a third party) core attributes and classes.
DCPromo	DRSR over RPC, LDAP, DNS RPC over SMB (SAMR)	Validate the command-line syntax passed to the promotion code and test promotion. Validate the forest or domain does not already exist if creating new.
VerifyOutbound ReplicationEnabled	LDAP, DRSR over SMB, RPC over SMB (LSARPC)	Validate the existing domain controller specified as the replication partner has outbound replication enabled by checking the NTDS Settings object's options attribute for NTDSDSA_OPT_DISABLE_OUTBOUND_REPL (0x00000004)
VerifyMachineAdmin Password	DRSR over RPC, LDAP, DNS RPC over SMB (SAMR)	Validate the safe mode password set for DSRM meets domain complexity requirements.
VerifySafeModePassword	*N/A*	Validate the local Administrator password set meets computer security policy complexity requirements.

Simplified Administration Appendix

- Server Manager Add Servers Dialog (Active Directory)
- Server Manager Remote Server Status
- Windows PowerShell Module Loading
- RID Issuance Hotfixes for Previous Operating Systems
- Ntdsutil.exe Install from Media Changes

Server Manager Add Servers Dialog (Active Directory)

The **Add Servers** dialog allows searching Active Directory for servers, by operating system, using wildcards, and by location. The dialog also allows using DNS queries by fully qualified domain name or prefix name. These searches use native DNS and LDAP protocols implemented through .NET, not AD Windows PowerShell against the AD Management Gateway through SOAP – meaning that the domain controllers contacted by Server Manager can even run Windows Server 2003. You can also import a file with server names for provisioning purposes.

The Active Directory search uses the following LDAP filters:

```
(&(ObjectCategory=computer)


(&(ObjectCategory=computer)(cn=dc*)(OperatingSystemVersion=6.2*))


(&(ObjectCategory=computer)(OperatingSystemVersion=6.1*))


(&(ObjectCategory=computer)(OperatingSystemVersion=6.0*))


(&(ObjectCategory=computer)(|(OperatingSystemVersion=5.2*)(OperatingSystemVersion=5.1*)))
```

The Active Directory search returns the following attributes:

```
( dnsHostName )( operatingSystem )( cn )
```

Server Manager Remote Server Status

Server Manager tests remote server accessibility using Address Routing Protocol. Any servers not responding to ARP requests are not listed, even if they are in the pool.

If ARP responds, then DCOM and WMI connections are made to the server to return status information. If RPC, DCOM, and WMI are unreachable, server manager cannot fully manage the server.

Windows PowerShell Module Loading

Windows PowerShell 3.0 implements dynamic module loading. Using the **Import-Module** cmdlet is typically no longer required; instead, simply invoking the cmdlet, alias, or function automatically loads the module.

To see loaded modules, use the **Get-Module** cmdlet.

```
Get-Module
```

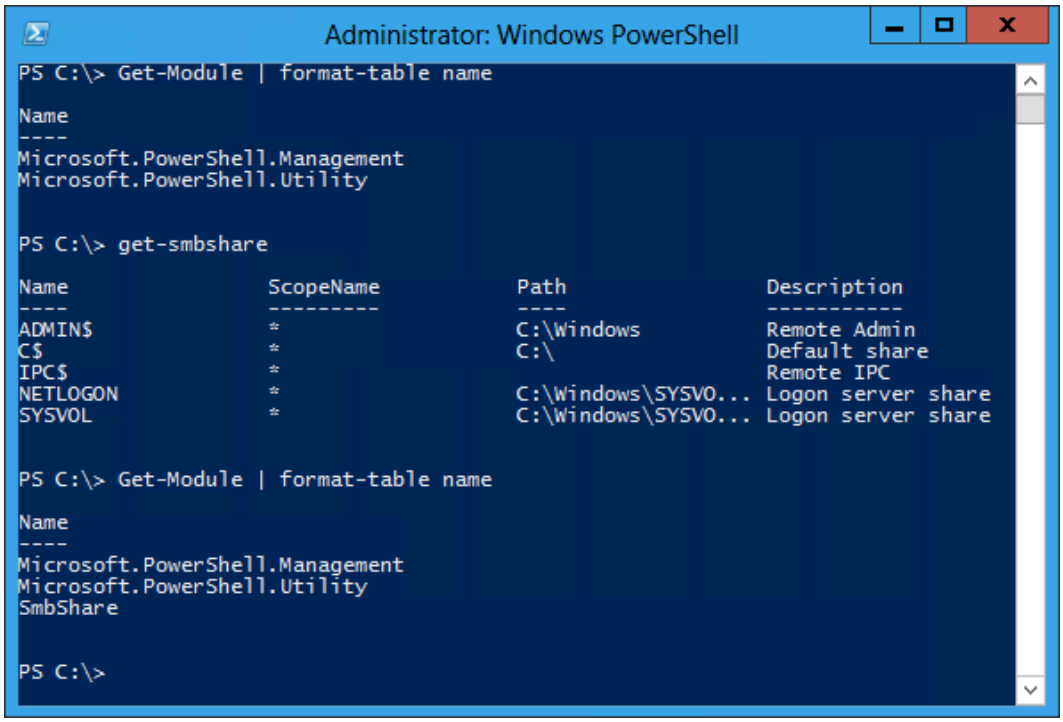

To see all installed modules with their exported functions and cmdlets, use:

```
Get-Module -ListAvailable
```

The main case for using the **import-module** command is when you need access to the "AD:" Windows PowerShell virtual drive and nothing else has already loaded the module. For example, using the following commands:

```
import-module activedirectory

cd ad:

dir
```

RID Issuance Hotfixes for Previous Operating Systems

See [An update is available to detect and prevent too much consumption of the global RID pool on a domain controller that is running Windows Server 2008 R2](#).

Ntdsutil.exe Install from Media Changes

Windows Server 2012 adds two additional options to the Ntdsutil.exe command-line tool for the **IFM (IFM Media Creation)** menu. These allow you to create IFM stores without first performing an offline defrag of the exported NTDS.DIT database file. When disk space is not a premium, this saves time creating the IFM.

The following table describes the two new menu items:

Menu Item	Explanation
Create Full NoDefrag %s	Create IFM media without defragmenting for a full AD DC or an AD/LDS instance into folder %s
Create Sysvol Full NoDefrag %s	Create IFM media with SYSVOL and without defragmenting for a full AD DC into folder %s

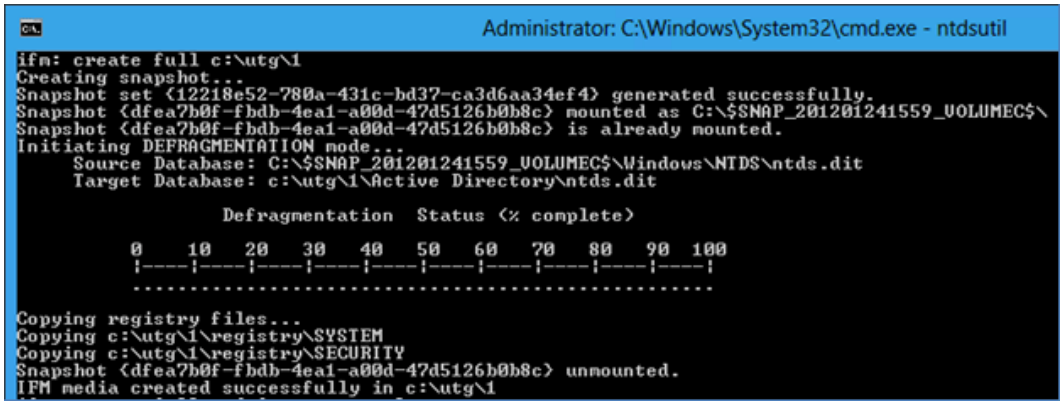

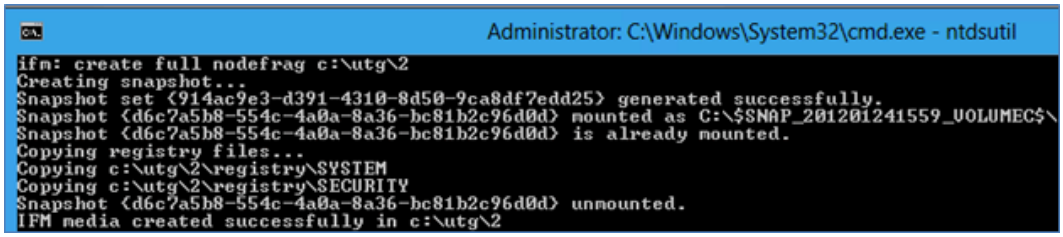

Install Active Directory Domain Services (Level 100)

This topic explains how to install AD DS in Windows Server 2012 by using any of the following methods:

- Credential requirements to run Adprep.exe and install Active Directory Domain Services

- Installing AD DS by Using Windows PowerShell
- Installing AD DS by using Server Manager
- Performing a Staged RODC Installation using the Graphical User Interface

Credential requirements to run Adprep.exe and install Active Directory Domain Services

The following credentials are required to run Adprep.exe and install AD DS.

- To install a new forest, you must be logged on as the local Administrator for the computer.
- To install a new child domain or new domain tree, you must be logged on as a member of the Enterprise Admins group.
- To install an additional domain controller in an existing domain, you must be a member of the Domain Admins group.

Notes

If you do not run adprep.exe command separately and you are installing the first domain controller that runs Windows Server 2012 in an existing domain or forest, you will be prompted to supply credentials to run Adprep commands. The credential requirements are as follows:

- To introduce the first Windows Server 2012 domain controller in the forest, you need to supply credentials for a member of Enterprise Admins group, the Schema Admins group, and the Domain Admins group in the domain that hosts the schema master.
- To introduce the first Windows Server 2012 domain controller in a domain, you need to supply credentials for a member of the Domain Admins group.
- To introduce the first read-only domain controller (RODC) in the forest, you need to supply credentials for a member of the Enterprise Admins group.

Note

If you have already run adprep /rodcprep in Windows Server 2008 or Windows Server 2008 R2, you do not need to run it again for Windows Server 2012.

Installing AD DS by Using Windows PowerShell

Beginning with Windows Server 2012, you can install AD DS using Windows PowerShell. Dcpromo.exe is deprecated beginning with Windows Server 2012, but you can still run dcpromo.exe by using an answer file (dcpromo /unattend:<answerfile> or dcpromo /answer:<answerfile>). The ability to continue running dcpromo.exe with an answer file provides organizations that have resources invested in existing automation time to convert the automation from dcpromo.exe to Windows PowerShell. For more information about running dcpromo.exe with an answer file, see http://support.microsoft.com/kb/947034.

For more information about removing AD DS using Windows PowerShell, see Remove AD DS using Windows PowerShell.

Start with adding the role using Windows PowerShell. This command installs the AD DS server role and installs the AD DS and AD LDS server administration tools, including GUI-based tools such as Active Directory Users and Computers and command-line tools such as dcdia.exe. Server administration tools are not installed by default when you use Windows PowerShell. You need to specify **–IncludeManagementTools** to manage the local server or install Remote Server Administration Tools to manage a remote server.

```
Install-windowsfeature -name AD-Domain-Services –IncludeManagementTools
```

```
<<Windows PowerShell cmdlet and arguments>>
```

There is no reboot required until after the AD DS installation is complete.

You can then run this command to see the available cmdlets in the ADDSDeployment module.

```
Get-command –module ADDSDeployment
```

To see the list of arguments that can be specified for a cmdlets and syntax:

```
Get-help <cmdlet name>
```

For example, to see the arguments for creating an unoccupied read-only domain controller (RODC) account, type

```
Get-help Add-ADDSReadOnlyDomainControllerAccount
```

Optional arguments appear in square brackets.

You can also download the latest Help examples and concepts for Windows PowerShell cmdlets. For more information, see about_Updatable_Help.

You can run Windows PowerShell cmdlets against remote servers:

- In Windows PowerShell, use invoke-command with the ADDSDeployment cmdlet. For example, to install AD DS on a remote server named ConDC3 in the contoso.com domain, type:

```
invoke-command {install-addsdomaincontroller –domainname
contoso.com –credential (get-credential) –computername condc3
```

-or-

- In Server Manager, create a server group that includes the remote server. Right-click the name of the remote server and click **Windows PowerShell**.

The next sections explain how to run ADDSDeployment module cmdlets to install AD DS.

- ADDSDeployment cmdlet arguments
- Specifying Windows PowerShell Credentials
- Using test cmdlets
- Installing a new forest root domain using Windows PowerShell
- Installing a new child or tree domain using Windows PowerShell
- Installing an additional (replica) domain controller using Windows PowerShell

ADDSDeployment cmdlet arguments

The following table lists arguments for the ADDSDeployment cmdlets in Windows PowerShell. Arguments in bold are required. Equivalent arguments for dcpromo.exe are listed in parentheses if they are named different in Windows PowerShell.

Windows PowerShell switches accept $TRUE or $FALSE arguments. Arguments that are $True by default do not need to be specified.

To override default values, you can specify the argument with a $False value. For example, because **-installdns** is automatically run for a new forest installation if it is not specified, the only way to *prevent* DNS installation when you install a new forest is to use:

```
-InstallDNS:$false
```

Similarly, because **–installdns** has a default value of $False if you install a domain controller in an environment that does not host Windows Server DNS server, you need to specify the following argument in order to install DNS server:

```
-InstallDNS:$true
```

Argument	Description
ADPrepCredential <PS Credential> **Note** Required if you are installing the first Windows Server 2012 domain controller in a domain or forest and the credentials of the current user are insufficient to perform the operation.	Specifies the account with Enterprise Admins and Schema Admins group membership that can prepare the forest, according to the rules of Get-Credential and a PSCredential object. If no value is specified, the value of the **–credential** argument is used.
AllowDomainControllerReinstall	Specifies whether to continue installing this writable domain controller, despite the fact that another writable domain controller account with the same name is detected. Use **$True** only if you are sure that the account is not currently used by another writable domain controller. The default is **$False**. This argument is not valid for an RODC.
AllowDomainReinstall	Specifies whether an existing domain is recreated. The default is **$False**.
AllowPasswordReplicationAccountName <string []>	Specifies the names of user accounts, group accounts, and computer accounts whose passwords can be replicated to this

Argument	Description
	RODC. Use an empty string "" if you want to keep the value empty. By default, only the Allowed RODC Password Replication Group is allowed, and it is originally created empty. Supply values as a string array. For example: `-AllowPasswordReplicationAccountName "JSmith","JSmithPC","Branch Users"`
ApplicationPartitionsToReplicate <string []> **Note** There is no equivalent option in the UI. If you install using the UI, or using IFM, then all application partitions will be replicated.	Specifies the application directory partitions to replicate. This argument is applied only when you specify the **-InstallationMediaPath** argument to install from media (IFM). By default, all application partitions will replicate based on their own scopes. Supply values as a string array. For example: `-ApplicationPartitionsToReplicate "partition1","partition2","partition3"`
Confirm	Prompts you for confirmation before running the cmdlet.
CreateDnsDelegation **Note** You cannot specify this argument when you run the Add-ADDSReadOnlyDomainController cmdlet.	Indicates whether to create a DNS delegation that references the new DNS server that you are installing along with the domain controller. Valid for Active Directory–integrated DNS only. Delegation records can be created only on Microsoft DNS servers that are online and accessible. Delegation records cannot be created for domains that are immediately subordinate to top-level domains such as .com, .gov, .biz, .edu or two-letter country code domains such as .nz and .au. The default is computed automatically based on the environment.
Credential <PS Credential> **Note**	Specifies the domain account that can logon to the domain, according to the rules

Argument	Description
Required only if the credentials of the current user are insufficient to perform the operation.	of Get-Credential and a PSCredential object. If no value is specified, the credentials of the current user are used.
CriticalReplicationOnly	Specifies whether the AD DS installation operation performs only critical replication before reboot and then continues. The noncritical replication happens after the installation finishes and the computer reboots. Using this argument is not recommended. There is no equivalent for this option in the user interface (UI).
DatabasePath <string>	Specifies the fully qualified, non–Universal Naming Convention (UNC) path to a directory on a fixed disk of the local computer that contains the domain database, for example, **C:\Windows\NTDS.** The default is **%SYSTEMROOT%\NTDS**. **Important** While you can store the AD DS database and log files on volume formatted with Resilient File System (ReFS), there are no specific benefits for hosting AD DS on ReFS, other than the normal benefits of resiliency you get for hosting any data on ReFS.
DelegatedAdministratorAccountName <string>	Specifies the name of the user or group that can install and administer the RODC. By default, only members of the Domain Admins group can administer an RODC.
DenyPasswordReplicationAccountName <string []>	Specifies the names of user accounts, group accounts, and computer accounts whose passwords are not to be replicated to this RODC. Use an empty string "" if you do not want to deny the replication of

Argument	Description
	credentials of any users or computers. By default, Administrators, Server Operators, Backup Operators, Account Operators, and the Denied RODC Password Replication Group are denied. By default, the Denied RODC Password Replication Group includes Cert Publishers, Domain Admins, Enterprise Admins, Enterprise Domain Controllers, Enterprise Read-Only Domain Controllers, Group Policy Creator Owners, the krbtgt account, and Schema Admins. Supply values as a string array. For example: `-DenyPasswordReplicationAccountName "RegionalAdmins","AdminPCs"`
DnsDelegationCredential <PS Credential> **Note** You cannot specify this argument when you run the Add-ADDSReadOnlyDomainController cmdlet.	Specifies the user name and password for creating DNS delegation, according to the rules of Get-Credential and a PSCredential object.
DomainMode <DomainMode> {Win2003 | Win2008 | Win2008R2 | Win2012} Or DomainMode <DomainMode> {2 | 3 | 4 | 5}	Specifies the domain functional level during the creation of a new domain. The domain functional level cannot be lower than the forest functional level, but it can be higher. The default value is automatically computed and set to the existing forest functional level or the value that is set for **-ForestMode**.
DomainName Required for Install-ADDSForest and Install-ADDSDomainController cmdlets.	Specifies the FQDN of the domain in which you want to install an additional domain controller.
DomainNetbiosName <string> Required for Install-ADDSForest if FQDN prefix name is longer than 15 characters.	Use with Install-ADDSForest. Assigns a NetBIOS name to the new forest root domain.
DomainType <DomainType> {ChildDomain | TreeDomain} or {child | tree}	Indicates the type of domain that you want to create: a new domain tree in an existing

Argument	Description
	forest, a child of an existing domain, or a new forest. The default for DomainType is ChildDomain.
Force	When this parameter is specified any warnings that might normally appear during the installation and addition of the domain controller will be suppressed to allow the cmdlet to complete its execution. This parameter can be useful to include when scripting installation.
ForestMode <ForestMode> {Win2003 | Win2008 | Win2008R2 | Win2012} Or ForestMode <ForestMode> {2 | 3 | 4 | 5}	Specifies the forest functional level when you create a new forest. The default value is Win2012.
InstallationMediaPath	Indicates the location of the installation media that will be used to install a new domain controller.
InstallDns	Specifies whether the DNS Server service should be installed and configured on the domain controller. For a new forest, the default is **$True** and DNS Server is installed. For a new child domain or domain tree, if the parent domain (or forest root domain for a domain tree) already hosts and stores the DNS names for the domain, then the default for this parameter is $True. For a domain controller installation in an existing domain, if this parameter is left unspecified and the current domain already hosts and stores the DNS names for the domain, then the default for this parameter is **$True**. Otherwise, if DNS domain names are hosted outside of Active Directory, the default is **$False** and no DNS Server is installed.
LogPath <string>	Specifies the fully qualified, non-UNC path

Argument	Description
	to a directory on a fixed disk of the local computer that contains the domain log files, for example, **C:\Windows\Logs**. The default is **%SYSTEMROOT%\NTDS**. **Important** Do not store the Active Directory log files on a data volume formatted with Resilient File System (ReFS).
MoveInfrastructureOperationMasterRoleIfNecessary	Specifies whether to transfer the infrastructure master operations master role (also known as flexible single master operations or FSMO) to the domain controller that you are creating—in case it is currently hosted on a global catalog server—and you do not plan to make the domain controller that you are creating a global catalog server. Specify this parameter to transfer the infrastructure master role to the domain controller that you are creating in case the transfer is needed; in this case, specify the **NoGlobalCatalog** option if you want the infrastructure master role to remain where it currently is.
NewDomainName <string> **Note** Required only for Install-ADDSDomain.	Specifies the single domain name for the new domain. For example, if you want to create a new child domain named **emea.corp.fabrikam.com**, you should specify **emea** as the value of this argument.
NewDomainNetbiosName <string> Required for Install-ADDSDomain if FQDN prefix name is longer than 15 characters.	Use with Install-ADDSDomain. Assigns a NetBIOS name to the new domain. The default value is derived from the value of **–NewDomainName**.
NoDnsOnNetwork	Specifies that DNS service is not available on the network. This parameter is used only when the IP setting of the network

Argument	Description
	adapter for this computer is not configured with the name of a DNS server for name resolution. It indicates that a DNS server will be installed on this computer for name resolution. Otherwise, the IP settings of the network adapter must first be configured with the address of a DNS server. Omitting this parameter (the default) indicates that the TCP/IP client settings of the network adapter on this server computer will be used to contact a DNS server. Therefore, if you are not specifying this parameter, ensure that TCP/IP client settings are first configured with a preferred DNS server address.
NoGlobalCatalog	Specifies that you do not want the domain controller to be a global catalog server. Domain controllers that run Windows Server 2012 are installed with the global catalog by default. In other words, this runs automatically without computation, unless you specify: `-NoGlobalCatalog`
NoRebootOnCompletion	Specifies whether to restart the computer upon completion of the command, regardless of success. By default, the computer will restart. To prevent the server from restarting, specify: `-NoRebootOnCompletion:$True` There is no equivalent for this option in the user interface (UI).
ParentDomainName <string> **Note** Required for Install-ADDSDomain cmdlet	Specifies the FQDN of an existing parent domain. You use this argument when you install a child domain or new domain tree. For example, if you want to create a new child domain named **emea.corp.fabrikam.com**, you should specify **corp.fabrikam.com** as the value of this argument.

Argument	Description
ReadOnlyReplica	Specifies whether to install a read-only domain controller (RODC).
ReplicationSourceDC <string>	Indicates the FQDN of the partner domain controller from which you replicate the domain information. The default is automatically computed.
SafeModeAdministratorPassword <securestring>	Supplies the password for the administrator account when the computer is started in Safe Mode or a variant of Safe Mode, such as Directory Services Restore Mode. The default is an empty password. You must supply a password. The password must be supplied in a System.Security.SecureString format, such as that provided by read-host -assecurestring or ConvertTo-SecureString. The SafeModeAdministratorPassword argument's operation is special:If not specified as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.If specified without a value, and there are no other arguments specified to the cmdlet, the cmdlet prompts you to enter a masked password without confirmation. This is not the preferred usage when running the cmdlet interactively.If specified with a value, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.For example, you can manually prompt for a password by using the Read-Host cmdlet to prompt the user for a secure string:-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)You can also provide a secure string as a converted clear-text

Argument	Description
	variable, although this is highly discouraged. -safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
SiteName <string> Required for the Add-addsreadonlydomaincontrolleraccount cmdlet	Specifies the site where the domain controller will be installed. There is no **–sitename** argument when you run **Install-ADDSForest** because the first site created is Default-First-Site-Name. The site name must already exist when provided as an argument to **-sitename**. The cmdlet will not create the site.
SkipAutoConfigureDNS	Skips automatic configuration of DNS client settings, forwarders, and root hints. This argument is in effect only if the DNS Server service is already installed or automatically installed with **-InstallDNS**.
SystemKey <string>	Specifies the system key for the media from which you replicate the data. The default is **none**. Data must be in format provided by read-host -assecurestring or ConvertTo-SecureString.
SysvolPath <string>	Specifies the fully qualified, non-UNC path to a directory on a fixed disk of the local computer, for example, **C:\Windows\SYSVOL**. The default is **%SYSTEMROOT%\SYSVOL**. **Important** SYSVOL cannot be stored on a data volume formatted with Resilient File System (ReFS).
SkipPreChecks	Does not run the prerequisite checks before starting installation. It is not advisable to use this setting.

Argument	Description
WhatIf	Shows what would happen if the cmdlet runs. The cmdlet is not run.

Specifying Windows PowerShell Credentials

You can specify credentials without revealing them in plain text on screen by using Get-credential.

The operation for the -SafeModeAdministratorPassword and LocalAdministratorPassword arguments is special:

- If not specified as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.
- If specified with a value, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string

```
-safemodeadministratorpassword (read-host -prompt "DSRM Password:" -assecurestring)
```

Warning
As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged:

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Warning
Providing or storing a clear text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM password of that domain controller. With that knowledge, they can impersonate the domain controller itself and elevate their privilege to the highest level in an Active Directory forest.

Using test cmdlets

Each ADDSDeployment cmdlet has a corresponding test cmdlet. The test cmdlets runs only the prerequisite checks for the installation operation; no installation settings are configured. The arguments for each test cmdlet are the same as for the corresponding installation cmdlet, but **–SkipPreChecks** is not available for test cmdlets.

Test cmdlet	Description
Test-ADDSForestInstallation	Runs the prerequisites for installing a new

Test cmdlet	Description
	Active Directory forest.
Test-ADDSDomainInstallation	Runs the prerequisites for installing a new domain in Active Directory.
Test-ADDSDomainControllerInstallation	Runs the prerequisites for installing a domain controller in Active Directory.
Test-ADDSReadOnlyDomainControllerAccountCreation	Runs the prerequisites for adding a read-only domain controller (RODC) account.

Installing a new forest root domain using Windows PowerShell

The command syntax for installing a new forest is as follows. Optional arguments appear within square brackets.

```
Install-ADDSForest [-SkipPreChecks] -DomainName <string> -SafeModeAdministratorPassword
<SecureString> [-CreateDNSDelegation] [-DatabasePath <string>] [-DNSDelegationCredential
<PS Credential>] [-NoDNSOnNetwork] [-DomainMode <DomainMode> {Win2003 | Win2008 |
Win2008R2 | Win2012}] [-DomainNetBIOSName <string>] [-ForestMode <ForestMode> {Win2003 |
Win2008 | Win2008R2 | Win2012}] [-InstallDNS] [-LogPath <string>] [-NoRebootOnCompletion]
[-SkipAutoConfigureDNS] [-SYSVOLPath] [-Force] [-WhatIf] [-Confirm] [<CommonParameters>]
```

Note

The -DomainNetBIOSName argument is required if you want to change the 15-character name that is automatically generated based on the DNS domain name prefix or if the name exceeds 15 characters.

For example, to install a new forest named corp.contoso.com and be securely prompted to provide the DSRM password, type:

```
Install-ADDSForest -domainname "corp.contoso.com"
```

Note

DNS server is installed by default when you run Install-ADDSForest.

To install a new forest named corp.contoso.com, create a DNS delegation in the contoso.com domain, set domain functional level to Windows Server 2008 R2 and set forest functional level to Windows Server 2008, install the Active Directory database and SYSVOL on the D:\ drive, install the log files on the E:\ drive, and be prompted to provide the Directory Services Restore Mode password and type:

```
Install-ADDSForest -DomainName corp.contoso.com -CreateDNSDelegation -DomainMode Win2008
-ForestMode Win2008R2 -DatabasePath "d:\NTDS" -SYSVOLPath "d:\SYSVOL" -LogPath "e:\Logs"
```

Installing a new child or tree domain using Windows PowerShell

The command syntax for installing a new domain is as follows. Optional arguments appear within square brackets.

```
Install-ADDSDomain [-SkipPreChecks] -NewDomainName <string> -ParentDomainName <string> -
SafeModeAdministratorPassword <SecureString> [-ADPrepCredential <PS Credential>] [-
AllowDomainReinstall] [-CreateDNSDelegation] [-Credential <PS Credential>] [-DatabasePath
<string>] [-DNSDelegationCredential <PS Credential>] [-NoDNSOnNetwork] [-DomainMode
<DomainMode> {Win2003 | Win2008 | Win2008R2 | Win2012}] [DomainType <DomainType> {Child
Domain | TreeDomain} [-InstallDNS] [-LogPath <string>] [-NoGlobalCatalog] [-
NewDomainNetBIOSName <string>] [-NoRebootOnCompletion] [-ReplicationSourceDC <string>] [-
SiteName <string>] [-SkipAutoConfigureDNS] [-Systemkey <SecureString>] [-SYSVOLPath] [-
Force] [-WhatIf] [-Confirm] [<CommonParameters>]
```

Notes

The **-credential** argument is only required when you are not currently logged on as a member of the Enterprise Admins group.

The **-NewDomainNetBIOSName** argument is required if you want to change the automatically generated 15-character name based on the DNS domain name prefix or if the name exceeds 15 characters.

For example, to use credentials of corp\EnterpriseAdmin1 to create a new child domain named child.corp.contoso.com, install DNS server, create a DNS delegation in the corp.contoso.com domain, set domain functional level to Windows Server 2003, make the domain controller a global catalog server in a site named Houston, use DC1.corp.contoso.com as the replication source domain controller, install the Active Directory database and SYSVOL on the D:\ drive, install the log files on the E:\ drive, and be prompted to provide the Directory Services Restore Mode password but not prompted to confirm the command, type:

```
Install-ADDSDomain –SafeModeAdministratorPassword –credential (get-credential
corp\EnterpriseAdmin1) –NewDomainName child –ParentDomainName corp.contoso.com –
InstallDNS –CreateDNSDelegation –DomainMode Win2003 –ReplicationSourceDC
DC1.corp.contoso.com –SiteName Houston –DatabasePath "d:\NTDS" –SYSVOLPath "d:\SYSVOL" –
LogPath "e:\Logs" –Confirm:$False
```

Installing an additional (replica) domain controller using Windows PowerShell

The command syntax for installing an additional domain controller is as follows. Optional arguments appear within square brackets.

```
Install-ADDSDomainController -DomainName <string> [-SkipPreChecks] –
SafeModeAdministratorPassword <SecureString> [-ADPrepCredential <PS Credential>] [-
AllowDomainControllerReinstall] [-ApplicationPartitionsToReplicate <string[]>] [-
CreateDNSDelegation] [-Credential <PS Credential>] [-CriticalReplicationOnly] [-
```

```
DatabasePath <string>] [-DNSDelegationCredential <PS Credential>] [-NoDNSOnNetwork] [-
NoGlobalCatalog] [-InstallationMediaPath <string>] [-InstallDNS] [-LogPath <string>] [-
MoveInfrastructureOperationMasterRoleIfNecessary] [-NoRebootOnCompletion] [-
ReplicationSourceDC <string>] [-SiteName <string>] [-SkipAutoConfigureDNS] [-SystemKey
<SecureString>] [-SYSVOLPath <string>] [-Force] [-WhatIf] [-Confirm] [<CommonParameters>]
```

To install a domain controller and DNS server in the corp.contoso.com domain and be prompted to supply the domain Administrator credentials and the DSRM password, type:

```
Install-ADDSDomainController -credential (get-credential corp\administrator) -domainname
"corp.contoso.com"
```

If the computer is already domain joined and you are a member of the Domain Admins group, you can use:

```
Install-ADDSDomainController -domainname "corp.contoso.com"
```

To be prompted for the domain name, type:

```
Install-ADDSDomainController -credential (get-credential) -domainname (read-host "Domain
to promote into")
```

The following command will use credentials of Contoso\EnterpriseAdmin1 to install a writable domain controller and a global catalog server in a site named Boston, install DNS server, create a DNS delegation in the contoso.com domain, install from media that is stored in the c:\ADDS IFM folder, install the Active Directory database and SYSVOL on the D:\ drive, install the log files on the E:\ drive, have the server automatically restart after AD DS installation is complete, and be prompted to provide the Directory Services Restore Mode password:

```
Install-ADDSDomainController –Credential (get-credential contoso\EnterpriseAdmin1) –
CreateDNSDelegation –DomainName corp.contoso.com –SiteName Boston –InstallationMediaPath
"c:\ADDS IFM" –DatabasePath "d:\NTDS" –SYSVOLPath "d:\SYSVOL" –LogPath "e:\Logs"
```

Performing a staged RODC installation using Windows PowerShell

The command syntax to create an RODC account is as follows. Optional arguments appear within square brackets.

```
Add-ADDSReadOnlyDomainControllerAccount [-SkipPreChecks] –DomainControllerAccuntName
<string> -DomainName <string> -SiteName <string> [-AllowPasswordReplicationAccountName
<string []>] [-NoGlobalCatalog] [-Credential <PS Credential>] [-
DelegatedAdministratorAccountName <string>] [-DenyPasswordReplicationAccountName <string
[]>] [-InstallDNS] [-ReplicationSourceDC <string>] [-Force] [-WhatIf] [-Confirm] [<Common
Parameters>]
```

The command syntax to attach a server to an RODC account is as follows. Optional arguments appear within square brackets.

```
Install-ADDSDomainController -DomainName <string> [-SkipPreChecks] –
SafeModeAdministratorPassword <SecureString> [-ADPrepCredential <PS Credential>] [-
```

```
ApplicationPartitionsToReplicate <string[]>] [-Credential <PS Credential>] [-
CriticalReplicationOnly] [-DatabasePath <string>] [-NoDNSOnNetwork] [-
InstallationMediaPath <string>] [-InstallDNS] [-LogPath <string>] [-
MoveInfrastructureOperationMasterRoleIfNecessary] [-NoRebootOnCompletion] [-
ReplicationSourceDC <string>] [-SkipAutoConfigureDNS] [-SystemKey <SecureString>] [-
SYSVOLPath <string>] [-UseExistingAccount] [-Force] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

For example, to create an RODC account named RODC1:

```
Add-ADDSReadOnlyDomainControllerAccount -DomainControllerAccountName RODC1 -DomainName
corp.contoso.com -SiteName Boston DelegatedAdministratoraccountName PilarA
```

Then run the following commands on the server that you want to attach to the RODC1 account. The server cannot be joined to the domain. First, install the AD DS server role and management tools:

```
install-windowsfeature -name AD-Domain-Services -includemanagementtools
```

The run the following command to create the RODC:

```
Install-ADDSDomainController -DomainName corp.contoso.com -SafeModeAdministratorPassword
(read-host -prompt "DSRM Password:" -assecurestring) -credential (get-credential
Corp\PilarA) -useexistingaccount
```

Press **Y** to confirm or include the **–confirm** argument to prevent the confirmation prompt.

Installing AD DS by using Server Manager

AD DS can be installed in Windows Server 2012 by using the Add Roles Wizard in Server Manager, followed by the Active Directory Domain Services Configuration Wizard, which is new beginning in Windows Server 2012. The Active Directory Domain Services Installation Wizard (dcpromo.exe) is deprecated beginning in Windows Server 2012.

The following sections explain how to create server pools in order to install and manage AD DS on multiple servers, and how to use the wizards to install AD DS.

Creating server pools

Server Manager can pool other servers on the network as long as they are accessible from the computer running Server Manager. Once pooled, you choose those servers for remote installation of AD DS or any other configuration options possible within Server Manager. The computer running Server Manager automatically pools itself. For more information about server pools, see Add Servers to Server Manager.

Note

In order to manage a domain-joined computer using Server Manager on a workgroup server, or vice-versa, additional configuration steps are needed. For more information, see “Add and manage servers in workgroups” in Add Servers to Server Manager.

Installing AD DS

Administrative credentials

The credential requirements to install AD DS vary depending on which deployment configuration you choose. For more information, see Credential requirements to run Adprep.exe and install Active Directory Domain Services.

Use the following procedures to install AD DS using the GUI method. The steps can be performed locally or remotely. For more detailed explanation of these steps, see the following topics:

- Deploying a Forest with Server Manager
- Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200)
- Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200)
- Install a Windows Server 2012 Active Directory Read-Only Domain Controller (RODC) (Level 200)

To install AD DS by using Server Manager

1. In Server Manager, click **Manage** and click **Add Roles and Features** to start the Add Roles Wizard.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation** and then click **Next**.
4. On the **Select destination server** page, click **Select a server from the server pool**, click the name of the server where you want to install AD DS and then click **Next**.

   To select remote servers, first create a server pool and add the remote servers to it. For more information about creating server pools, see Add Servers to Server Manager.
5. On the **Select server roles** page, click **Active Directory Domain Services**, then on the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
6. On the **Select features** page, select any additional features you want to install and click **Next**.
7. On the **Active Directory Domain Services** page, review the information and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. On the **Results** page, verify that the installation succeeded, and click **Promote this server to a domain controller** to start the Active Directory Domain Services Configuration Wizard.

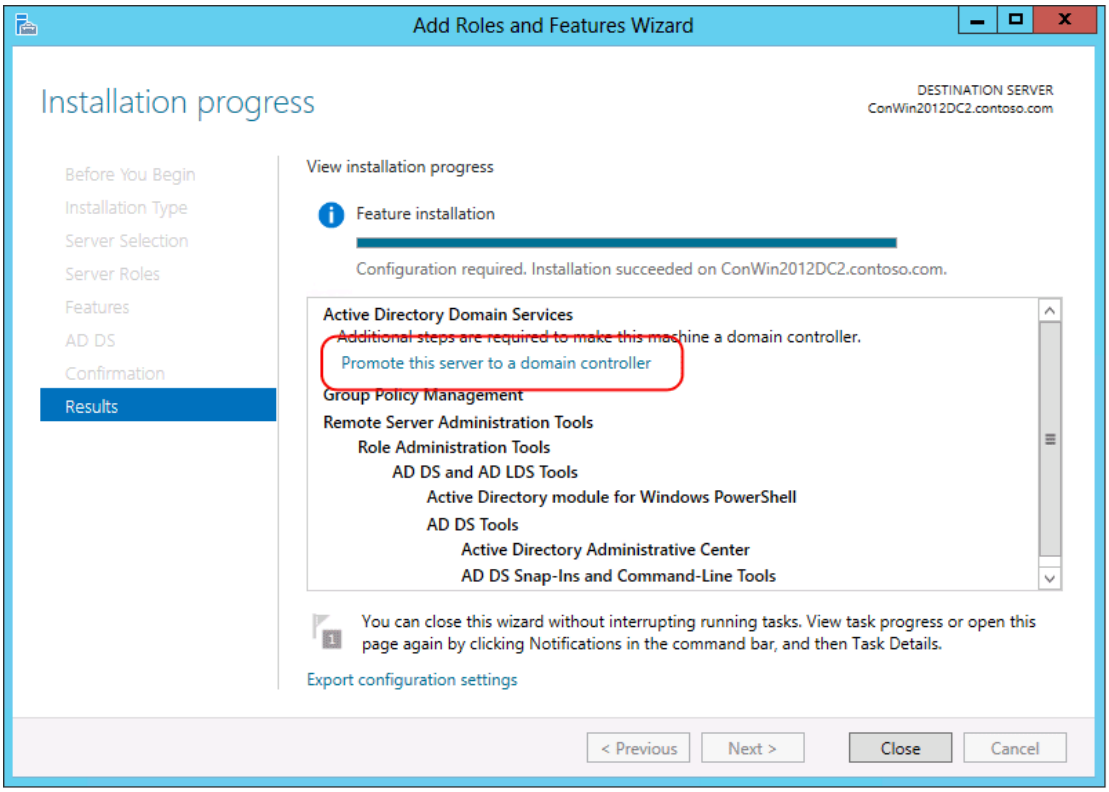

Important

If you close Add Roles Wizard at this point without starting the Active Directory Domain Services Configuration Wizard, you can restart it by clicking Tasks in Server Manager.

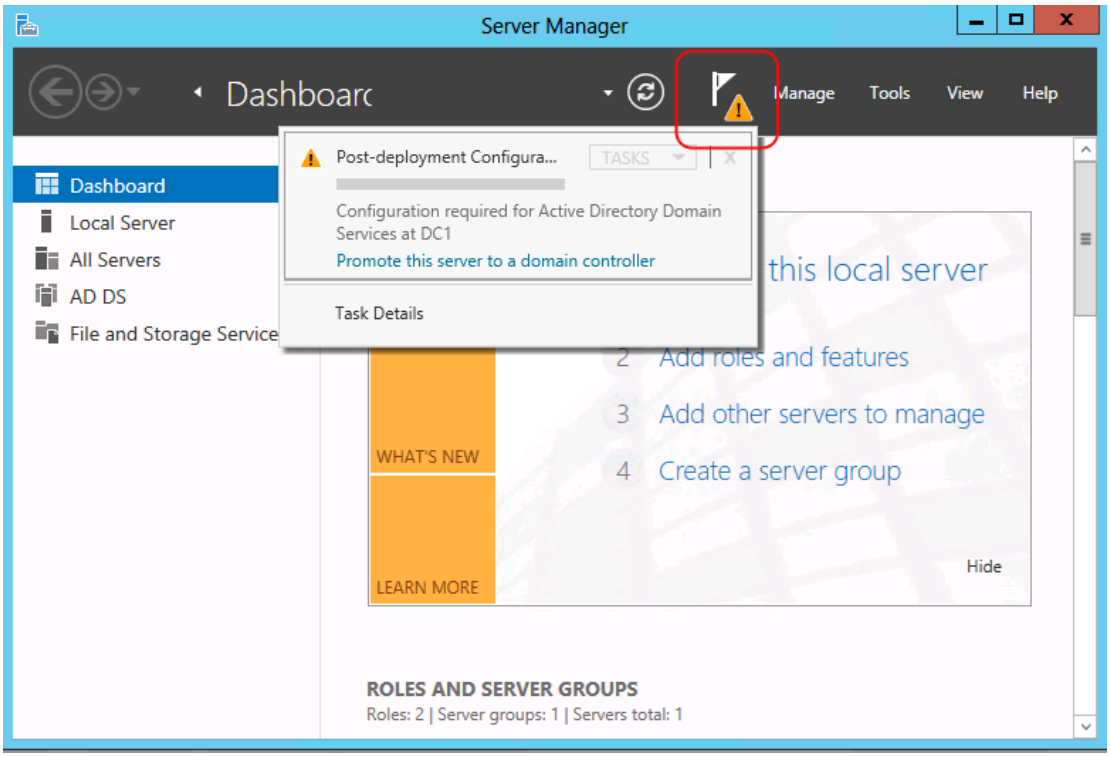

10. On the **Deployment Configuration** page, choose one of the following options:
    - If you are installing an additional domain controller in an existing domain, click **Add a domain controller to an existing domain**, and type the name of the domain (for example, emea.corp.contoso.com) or click **Select…** to choose a domain, and credentials (for example, specify an account that is a member of the Domain Admins group) and then click **Next**.

      Note
      The name of the domain and current user credentials are supplied by default only if the machine is domain-joined and you are performing a local installation. If you are installing AD DS on a remote server, you need to specify the credentials, by design. If current user credentials are not sufficient to perform the installation, click **Change…** in order to specify different credentials.

      For more information, see Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200).
    - If you are installing a new child domain, click **Add a new domain to an existing forest**, for **Select domain type**, select **Child Domain**, type or browse to the name of the parent domain DNS name (for example, corp.contoso.com), type the relative name of the new child domain (for example emea), type credentials to use to create the new domain, and then click **Next**.

      For more information, see Install a New Windows Server 2012 Active Directory Child

or Tree Domain (Level 200).

- If you are installing a new domain tree, click **Add new domain to an existing forest**, for **Select domain type**, choose **Tree Domain**, type the name of the root domain (for example, corp.contoso.com), type the DNS name of the new domain (for example, fabrikam.com), type credentials to use to create the new domain, and then click **Next**.

  For more information, see Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200).

- If you are installing a new forest, click **Add a new forest** and then type the name of the root domain (for example, corp.contoso.com).

  For more information, see Install a New Windows Server 2012 Active Directory Forest (Level 200).

11. On the **Domain Controller Options** page, choose one of the following options:
    - If you are creating a new forest or domain, select the domain and forest functional levels, click **Domain Name System (DNS) server**, specify the DSRM password, and then click **Next**.
    - If you are adding a domain controller to an existing domain, click **Domain Name System (DNS) server**, **Global Catalog (GC)**, or **Read Only Domain Controller (RODC)** as needed, choose the site name, and type the DSRM password and then click **Next**.

    For more information about which options on this page are available or not available under different conditions, see Domain Controller Options.
12. On the **DNS Options** page (which appears only if you install a DNS server), click **Update DNS delegation** as needed. If you do, provide credentials that have permission to create DNS delegation records in the parent DNS zone.

    If a DNS server that hosts the parent zone cannot be contacted, the **Update DNS Delegation** option is not available.

    For more information about whether you need to update the DNS delegation, see Understanding Zone Delegation. If you attempt to update the DNS delegation and encounter an error, see DNS Options.
13. On the **RODC Options** page (which appears only if you install an RODC), specify the name of a group or user who will manage the RODC, add accounts to or remove accounts from the Allowed or Denied password replication groups, and then click **Next**.

    For more information, see Password Replication Policy.
14. On the **Additional Options** page, choose one of the following options:
    - If you are creating a new domain, type a new NetBIOS name or verify the default NetBIOS name of the domain, and then click **Next**.
    - If you are adding a domain controller to an existing domain, select the domain controller that you want to replicate the AD DS installation data from (or allow the wizard to select any domain controller). If you are installing from media, click **Install from media path** type and verify the path to the installation source files, and then click **Next**.

      You cannot use install from media (IFM) to install the first domain controller in a

domain. IFM does not work across different operating system versions. In other words, in order to install an additional domain controller that runs Windows Server 2012 by using IFM, you must create the backup media on a Windows Server 2012 domain controller. For more information about IFM, see Installing an Additional Domain Controller by Using IFM.

15. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder (or accept default locations), and click **Next**.

    Important
    Do not store the Active Directory database, log files, or SYSVOL folder on a data volume formatted with Resilient File System (ReFS).

16. On the **Preparation Options** page, type credentials that are sufficient to run adprep. For more information, see Credential requirements to run Adprep.exe and install Active Directory Domain Services.
17. On the **Review Options** page, confirm your selections, click **View script** if you want to export the settings to a Windows PowerShell script, and then click **Next**.
18. On the **Prerequisites Check** page, confirm that prerequisite validation completed and then click **Install**.
19. On the **Results** page, verify that the server was successfully configured as a domain controller. The server will be restarted automatically to complete the AD DS installation.

Performing a Staged RODC Installation using the Graphical User Interface

A staged RODC installation allows you to create an RODC in two stages. In the first stage, a member of the Domain Admins group creates an RODC account. In the second stage, a server is attached to the RODC account. The second stage can be completed by a member of the Domain Admins group or a delegated domain user or group.

To create an RODC account by using the Active Directory management tools

1. You can create the RODC account using Active Directory Administrative Center or Active Directory Users and Computers.
    a. Click **Start**, click **Administrative Tools**, and then click **Active Directory Administrative Center**.
    b. In the navigation pane (left pane), click the name of the domain.
    c. In the Management list (center pane), click the **Domain Controllers** OU.
    d. In the Tasks Pane (right pane), click **Pre-create a read-only domain controller account**.

    -Or-

    a. Click **Start**, click **Administrative Tools**, and then click **Active Directory Users and**

Computers.

b. Either right-click the **Domain Controllers** organizational unit (OU) or click the **Domain Controllers** OU, and then click **Action**.

c. Click **Pre-create Read-only Domain Controller account**.

2. On the **Welcome to the Active Directory Domain Services Installation Wizard** page, if you want to modify the default the Password Replication Policy (PRP), select **Use advanced mode installation**, and then click **Next**.
3. On the **Network Credentials** page, under **Specify the account credentials to use to perform the installation**, click **My current logged on credentials** or click **Alternate credentials**, and then click **Set**. In the **Windows Security** dialog box, provide the user name and password for an account that can install the additional domain controller. To install an additional domain controller, you must be a member of the Enterprise Admins group or the Domain Admins group. When you are finished providing credentials, click **Next**.
4. On the **Specify the Computer Name** page, type the computer name of the server that will be the RODC.
5. On the **Select a Site** page, select a site from the list or select the option to install the domain controller in the site that corresponds to the IP address of the computer on which you are running the wizard, and then click **Next**.
6. On the **Additional Domain Controller Options** page, make the following selections, and then click **Next**:
   - **DNS server**: This option is selected by default so that your domain controller can function as a Domain Name System (DNS) server. If you do not want the domain controller to be a DNS server, clear this option. However, if you do not install the DNS server role on the RODC and the RODC is the only domain controller in the branch office, users in the branch office will not be able to perform name resolution when the wide area network (WAN) to the hub site is offline.
   - **Global catalog**: This option is selected by default. It adds the global catalog, read-only directory partitions to the domain controller, and it enables global catalog search functionality. If you do not want the domain controller to be a global catalog server, clear this option. However, if you do not install a global catalog server in the branch office or enable universal group membership caching for the site that includes the RODC, users in the branch office will not be able to log on to the domain when the WAN to the hub site is offline.
   - **Read-only domain controller**. When you create an RODC account, this option is selected by default and you cannot clear it.
7. If you selected the **Use advanced mode installation** check box on the **Welcome** page, the **Specify the Password Replication Policy** page appears. By default, no account passwords are replicated to the RODC, and security-sensitive accounts (such as members of the Domain Admins group) are explicitly denied from ever having their passwords replicated to the RODC.

   To add other accounts to policy, click **Add**, then click **Allow passwords for the account to replicate to this RODC** or click **Deny passwords for the account from replicating**

to this RODC and then select the accounts.

When complete (or to accept the default setting), click **Next**.

8. On the **Delegation of RODC Installation and Administration** page, type the name of the user or the group who will attach the server to the RODC account that you are creating. You can type the name of only one security principal.

   To search the directory for a specific user or group, click **Set**. In **Select User or Group**, type the name of the user or group. We recommend that you delegate RODC installation and administration to a group.

   This user or group will also have local administrative rights on the RODC after the installation. If you do not specify a user or group, only members of the Domain Admins group or the Enterprise Admins group will be able to attach the server to the account.

   When you are finished, click **Next**.

9. On the **Summary** page, review your selections. Click **Back** to change any selections, if necessary.

   To save the settings that you selected to an answer file that you can use to automate subsequent AD DS operations, click **Export settings**. Type a name for your answer file, and then click **Save**.

   When you are sure that your selections are accurate, click **Next** to create the RODC account.

10. On the **Completing the Active Directory Domain Services Installation Wizard** page, click **Finish**.

After an RODC account is created, you can attach a server to account to complete the RODC installation. This second stage can be completed in the branch office where the RODC will be located. The server where you perform this procedure must not be joined to the domain. Beginning in Windows Server 2012, you use the Add Roles Wizard in Server Manager to attach a server to an RODC account.

To attach a server to an RODC account using Server Manager

1. Log on as local Administrator.
2. In Server Manager, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or feature-based installation** and then click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, click the name of the server where you want to install AD DS and then click **Next**.
6. On the **Select server roles** page, click **Active Directory Domain Services**, click **Add Features** and then click **Next**.
7. On the **Select features** page, select any additional features that you want to install and click **Next**.
8. On the **Active Directory Domain Services** page, review the information and then click

Next.

9. On the **Confirm installation selections** page, click **Install**.
10. On the **Results** page, verify **Installation succeeded**, and click **Promote this server to a domain controller** to start the Active Directory Domain Services Configuration Wizard.

> **Important**
> If you close Add Roles Wizard at this point without starting the Active Directory Domain Services Configuration Wizard, you can restart it by clicking Tasks in Server Manager.

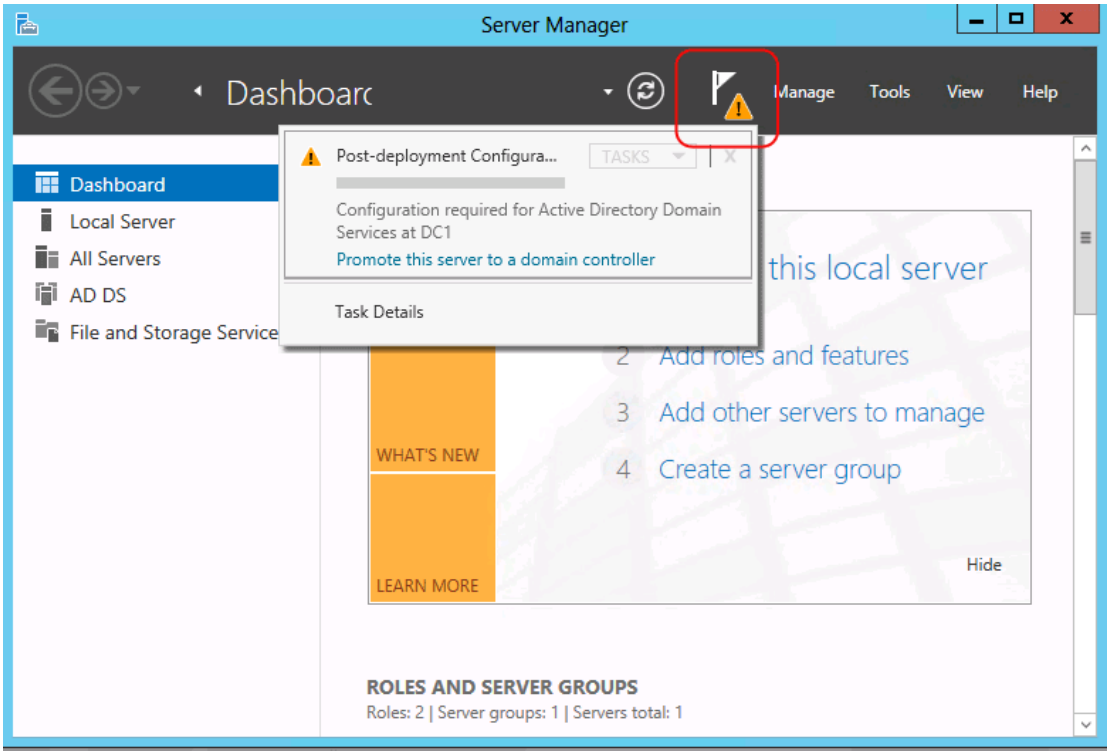

11. On the **Deployment Configuration** page, click **Add a domain controller to an existing domain**, type the name of the domain (for example, emea.contoso.com) and credentials (for example, specify an account that is delegated to manage and install the RODC), and then click **Next**.
12. On the **Domain Controller Options** page, click **Use existing RODC account**, type and confirm the Directory Services Restore Mode password, and then click **Next**.
13. On the **Additional Options** page, if you are installing from media, click **Install from media path** type and verify the path to the installation source files, select the domain controller that you want to replicate the AD DS installation data from (or allow the wizard to select any domain controller) and then click **Next**.
14. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder, or accept default locations, and then click **Next**.
15. On the **Review Options** page, confirm your selections, click **View Script** to export the

settings to a Windows PowerShell script, and then click **Next**.

16. On the **Prerequisites Check** page, confirm that prerequisite validation completed and then click **Install**.

    To complete the AD DS installation, the server will restart automatically.

See Also

Troubleshooting Domain Controller Deployment

Install a New Windows Server 2012 Active Directory Forest (Level 200)

Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200)

Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200)

Install a New Windows Server 2012 Active Directory Forest (Level 200)

This topic explains the new Windows Server 2012 Active Directory Domain Services domain controller promotion feature at an introductory level. In Windows Server 2012, AD DS replaces the Dcpromo tool with a Server Manager and Windows PowerShell-based deployment system.

- Active Directory Domain Services Simplified Administration
- Technical Overview
- Deploying a Forest with Server Manager
- Deploying a Forest with Windows PowerShell

Active Directory Domain Services Simplified Administration

Windows Server 2012 introduces the next generation of Active Directory Domain Services Simplified Administration, and is the most radical domain re-envisioning since Windows 2000 Server. AD DS Simplified Administration takes lessons learned from twelve years of Active Directory and makes a more supportable, more flexible, more intuitive administrative experience for architects and administrators. This meant creating new versions of existing technologies as well as extending the capabilities of components released in Windows Server 2008 R2.

What Is AD DS Simplified Administration?

AD DS Simplified Administration is a reimagining of domain deployment. Some of those features include:

- AD DS role deployment is now part of the new Server Manager architecture and allows remote installation.

- The AD DS deployment and configuration engine is now Windows PowerShell, even when using a graphical setup.
- Promotion now includes prerequisite checking that validates forest and domain readiness for the new domain controller, lowering the chance of failed promotions.
- The Windows Server 2012 forest functional level does not implement new features and domain functional level is required only for a subset of new Kerberos features, relieving administrators of the frequent need for a homogenous domain controller environment.

Purpose and Benefits

These changes may appear more complex, not simpler. In redesigning the AD DS deployment process though, there was opportunity to coalesce many steps and best practices into fewer, easier actions. This means, for example, that the graphical configuration of a new replica domain controller is now eight dialogs rather than the previous twelve. Creating a new Active Directory forest requires a *single* Windows PowerShell command with only *one* argument: the name of the domain.

Why is there such an emphasis on Windows PowerShell in Windows Server 2012? As distributed computing evolves, Windows PowerShell allows a single engine for configuration and maintenance from both graphical and command-line interfaces. It permits fully featured scripting of any component with the same first class citizenship for an IT Professional that an API grants to developers. As cloud-based computing becomes ubiquitous, Windows PowerShell also finally brings the ability to remotely administer a server, where a computer with no graphical interface has the same management capabilities as one with a monitor and mouse.

A veteran AD DS administrator should find their previous knowledge highly relevant. A beginning administrator will find a far shallower learning curve.

Technical Overview

What You Should Know Before You Begin

This topic assumes familiarity with previous releases of Active Directory Domain Services, and does not provide foundational detail around their purpose and functionality. For more information about AD DS, see the TechNet Portal pages linked below:

- Active Directory Domain Services for Windows Server 2008 R2
- Active Directory Domain Services for Windows Server 2008
- Windows Server Technical Reference

Functional Descriptions

AD DS Role Installation

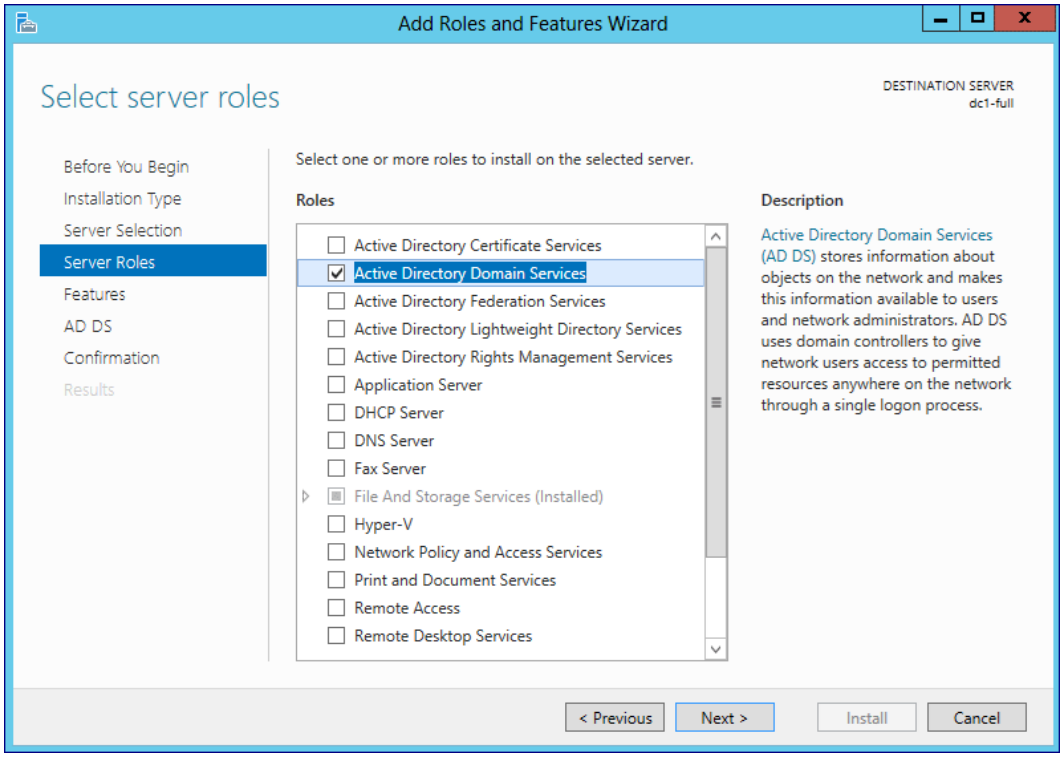

Active Directory Domain Services installation uses Server Manager and Windows PowerShell, like all other server roles and features in Windows Server 2012. The Dcpromo.exe program no longer provides GUI configuration options.

You use a graphical wizard in Server Manager or the ServerManager module for Windows PowerShell in both local and remote installations. By running multiple instances of those wizards or cmdlets and targeting different servers, you can deploy AD DS to multiple domain controllers simultaneously, all from one single console. Although these new features are not backwards compatible with Windows Server 2008 R2 or earlier operating systems, you can also still use the Dism.exe application introduced in Windows Server 2008 R2 for local role installation from the classic command-line.

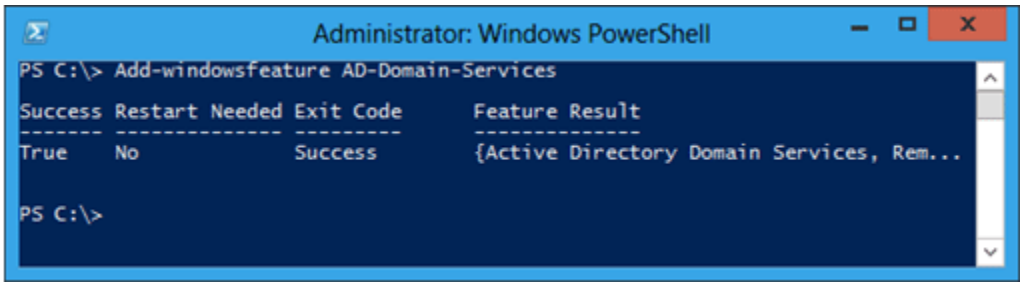

AD DS Role Configuration

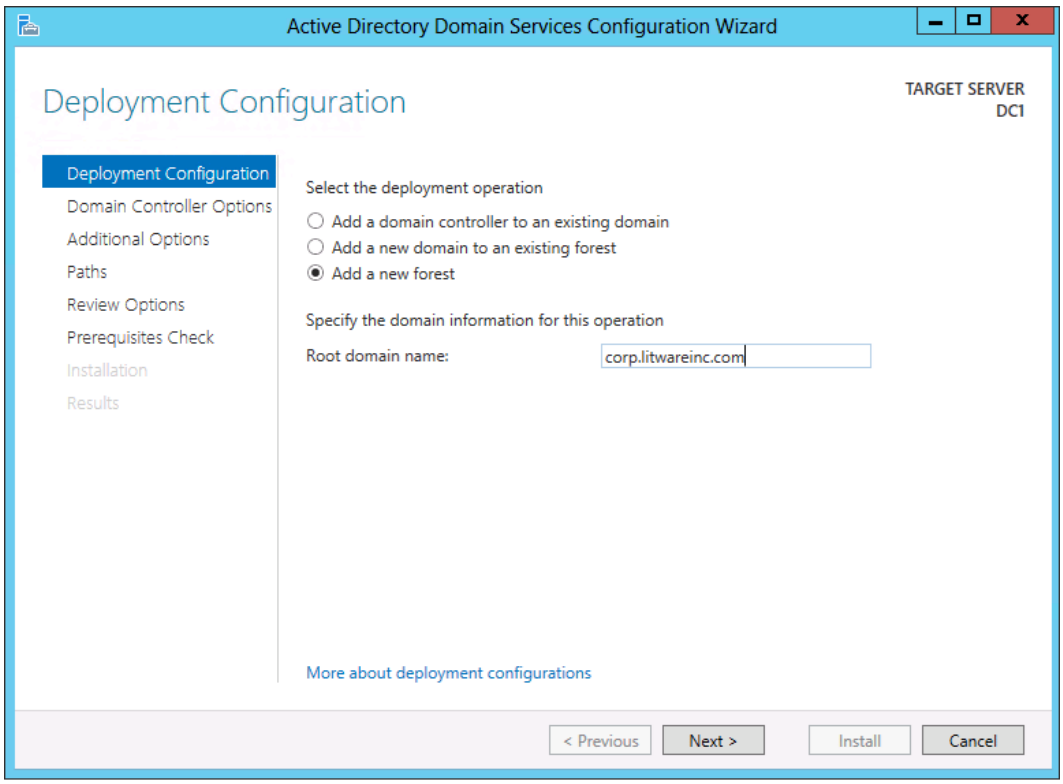

Active Directory Domain Services configuration – previously known as DCPROMO – is a now a discrete operation from role installation. After installing the AD DS role, an administrator configures the server as a domain controller using a separate wizard within Server Manager or using the ADDSDeployment Windows PowerShell module.

AD DS role configuration builds on twelve years of field experience and now configures domain controllers based on the most recent Microsoft best practices. For example, Domain Name System and Global Catalogs install by default on every domain controller.

The Server Manager AD DS configuration wizard merges many individual dialogs into fewer prompts and no longer hides settings in an “advanced” mode. The entire promotion process is in one expanding dialog box during installation. The wizard and the ADDSDeployment Windows PowerShell module show you notable changes and security concerns, with links to further information.

The Dcpromo.exe remains in Windows Server 2012 for command-line unattended installations only, and no longer runs the graphical installation wizard. It is highly recommended that you discontinue use of Dcpromo.exe for unattended installs and replace it with the ADDSDeployment module, as the now-deprecated executable will not be included in the next version of Windows.

These new features are not backwards compatible to Windows Server 2008 R2 or older operating systems.

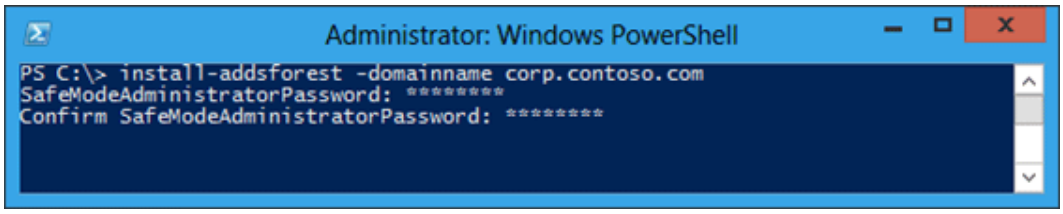

Important

Dcpromo.exe no longer contains a graphical wizard and no longer installs role or feature binaries. Attempting to run Dcpromo.exe from the Explorer shell returns:

"The Active Directory Domain Services Installation Wizard is relocated in Server Manager. For more information, see http://go.microsoft.com/fwlink/?LinkId=220921."

Attempting to run Dcpromo.exe /unattend still installs the binaries, as in previous operating systems, but warns:

"The dcpromo unattended operation is replaced by the ADDSDeployment module for Windows PowerShell. For more information, see http://go.microsoft.com/fwlink/?LinkId=220924."

Windows Server 2012 deprecates dcpromo.exe and it will not be included with future versions of Windows, nor will it receive further enhancements in this operating system. Administrators should discontinue its use and switch to the supported Windows PowerShell modules if they wish to create domain controllers from the command-line.

Prerequisite Checking

Domain controller configuration also implements a prerequisite checking phase that evaluates the forest and domain prior to continuing with domain controller promotion. This includes FSMO role availability, user privileges, extended schema compatibility and other requirements. This new design alleviates issues where domain controller promotion starts and then halts midway with a fatal configuration error. This lessens the chance of orphaned domain controller metadata in the forest or a server that incorrectly believes it is a domain controller.

Deploying a Forest with Server Manager

This section explains how to install the first domain controller in a forest root domain using Server Manager on a graphical Windows Server 2012 computer.

Server Manager AD DS Role Installation Process

The diagram below illustrates the Active Directory Domain Services role installation process, beginning with you running ServerManager.exe and ending right before the promotion of the domain controller.

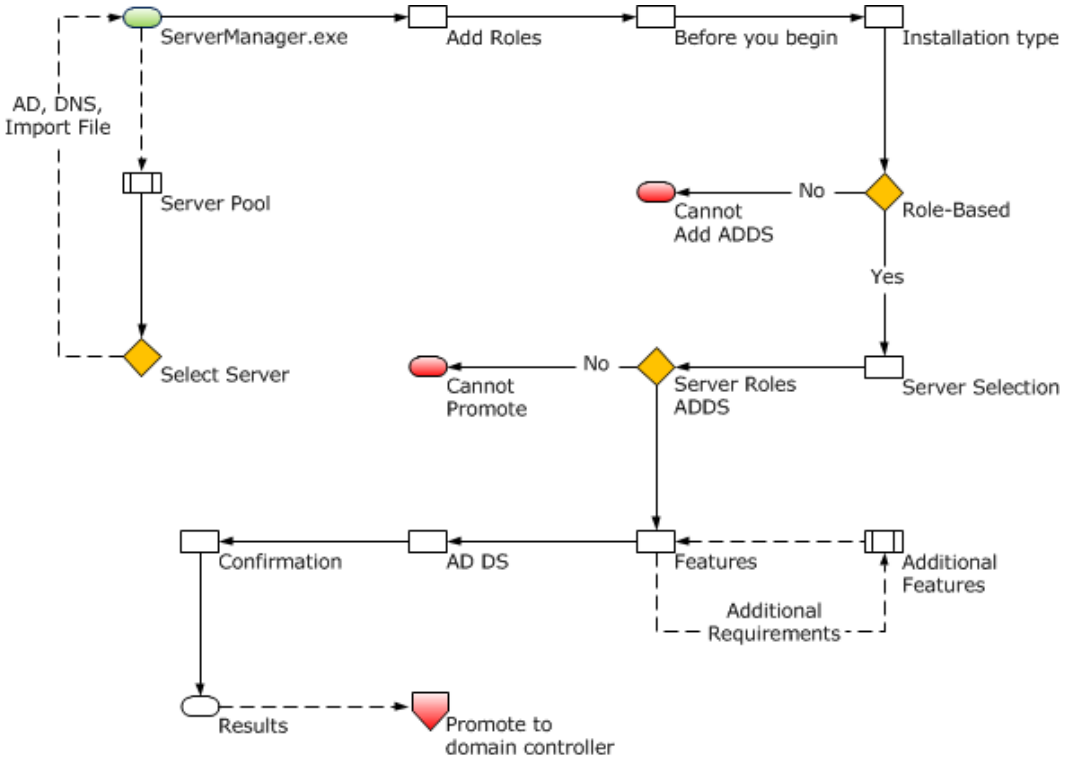

Server Pool and Add Roles

Any Windows Server 2012 computers accessible from the computer running Server Manager are eligible for pooling. Once pooled, you select those servers for remote installation of AD DS or any other configuration options possible within Server Manager.

To add servers, choose one of the following:

- Click **Add Other Servers to Manage** on the dashboard welcome tile
- Click the **Manage** menu and select **Add Servers**
- Right-click **All Servers** and choose **Add Servers**

This brings up the Add Servers dialog:

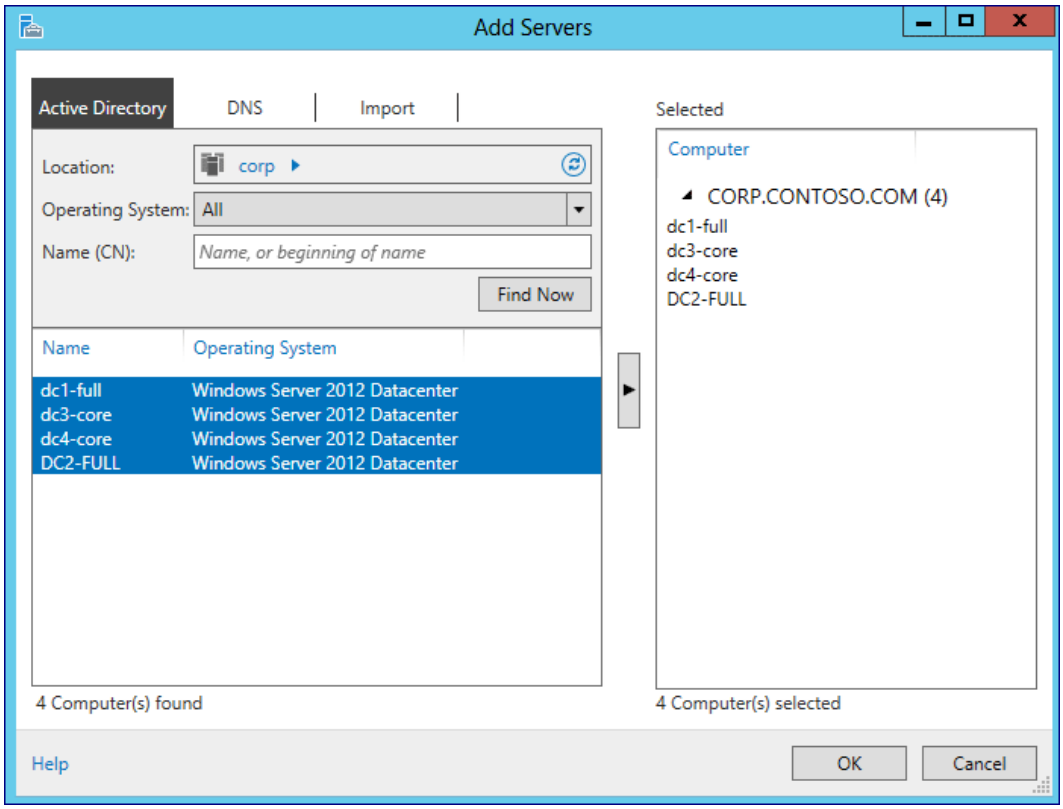

This gives you three ways to add servers to the pool for use or grouping:

- Active Directory search (uses LDAP, requires that the computers belong to a domain, allows operating system filtering and supports wildcards)
- DNS search (uses DNS alias or IP address via ARP or NetBIOS broadcast or WINS lookup, does not allow operating system filtering or support wildcards)
- Import (uses a text file list of servers separated by CR/LF)

Click **Find Now** to return a list of servers from that same Active Directory domain that the computer is joined to, Click one or more server names from the list of servers. Click the right arrow to add the servers to the **Selected** list. Use the **Add Servers** dialog to add selected servers to dashboard role groups. Or Click **Manage**, and then click **Create Server Group**, or click **Create Server Group** on the dashboard **Welcome to Server Manager** tile to create custom server groups.

Note

The Add Servers procedure does not validate that a server is online or accessible. However, any unreachable servers flag in the Manageability view in Server Manager at the next refresh

You can install roles remotely on any Windows Server 2012 computers added the pool, as shown:

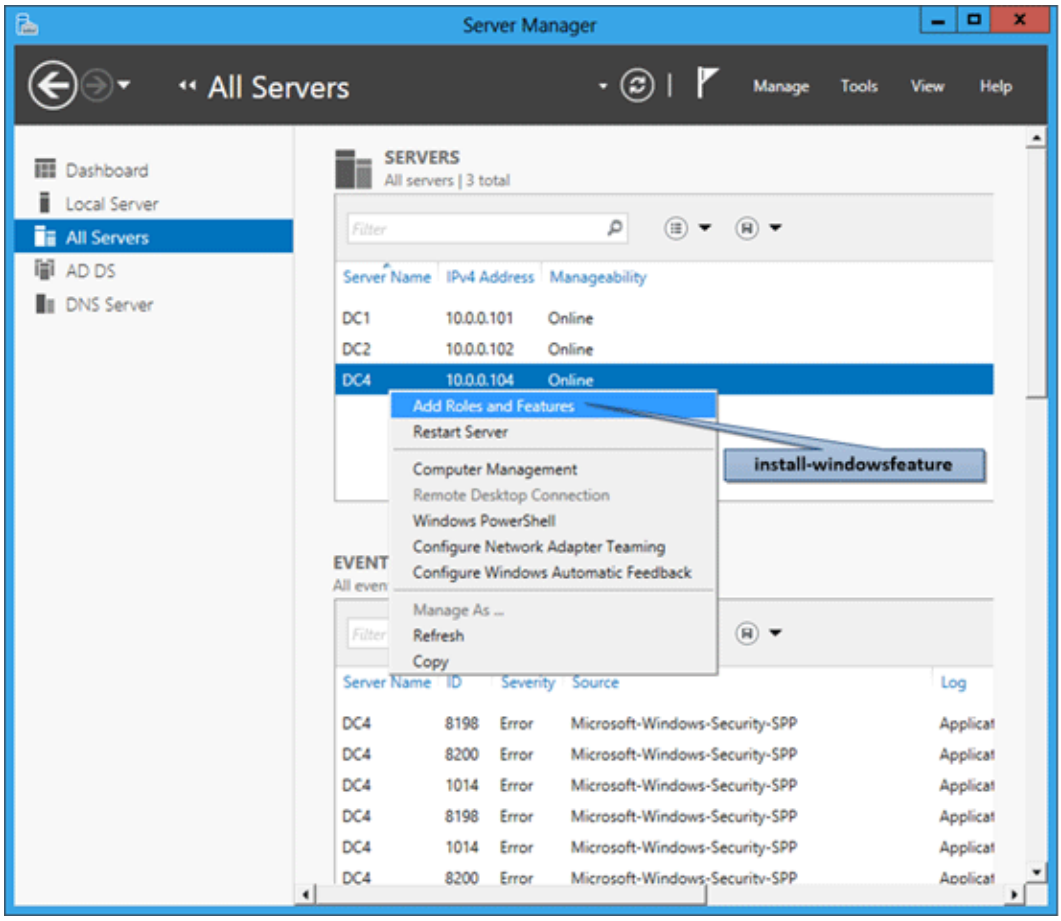

You cannot fully manage servers running operating systems older than Windows Server 2012. The **Add Roles and Features** selection is running ServerManager Windows PowerShell Module **Install-WindowsFeature**.

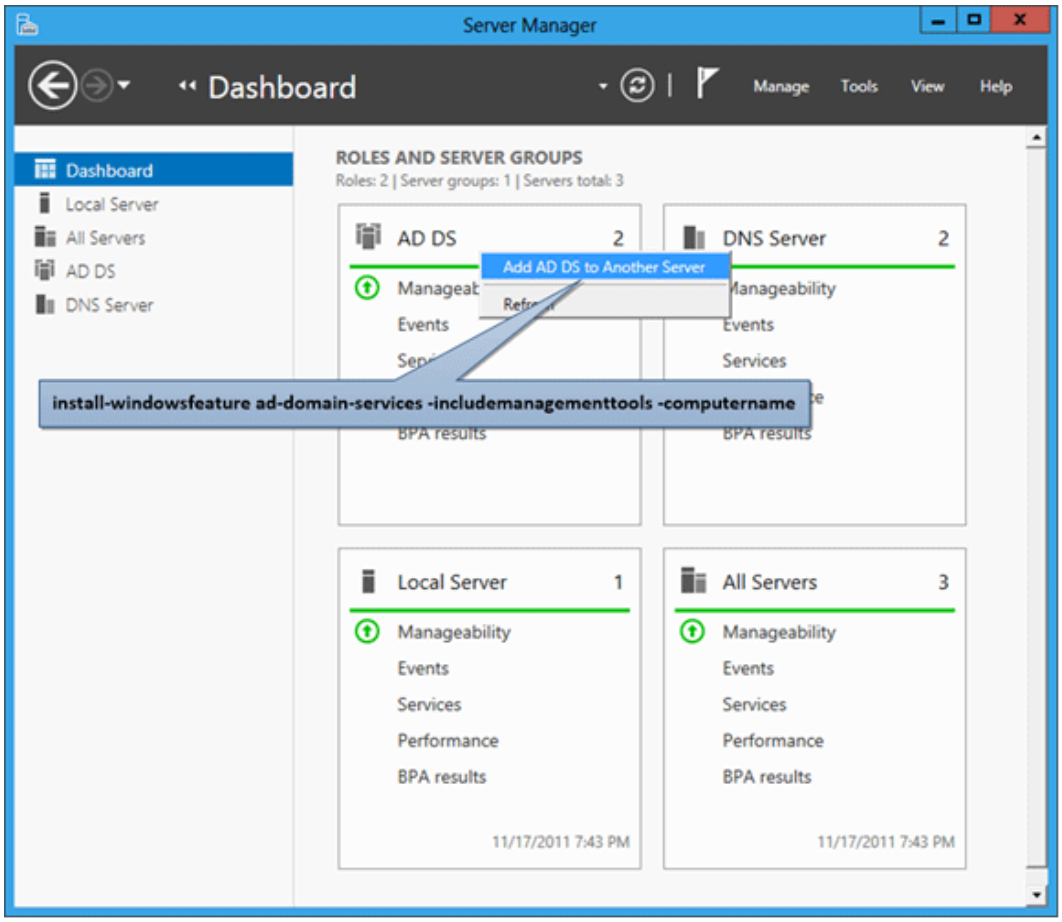

You can also use the Server Manager Dashboard on an existing domain controller to select remote server AD DS installation with the role already preselected by right clicking the AD DS dashboard tile and selecting **Add AD DS to Another Server**. This is invoking **Install-WindowsFeature AD-Domain-Services**.

The computer you are running Server Manager on pools itself automatically. To install the AD DS role here, simply click the **Manage** menu and click **Add Roles and Features**.

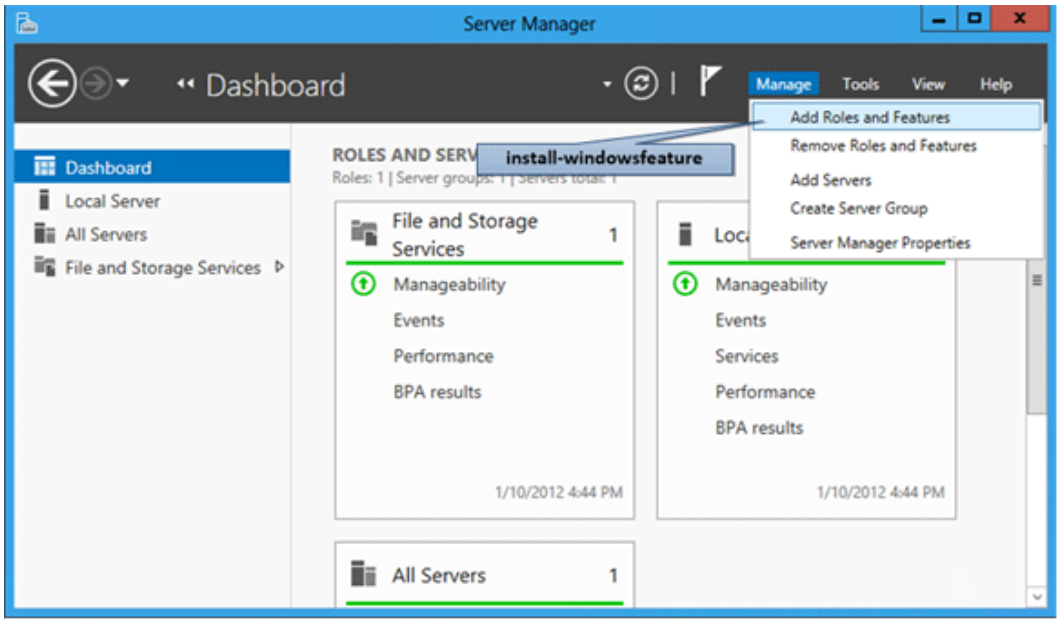

Installation Type

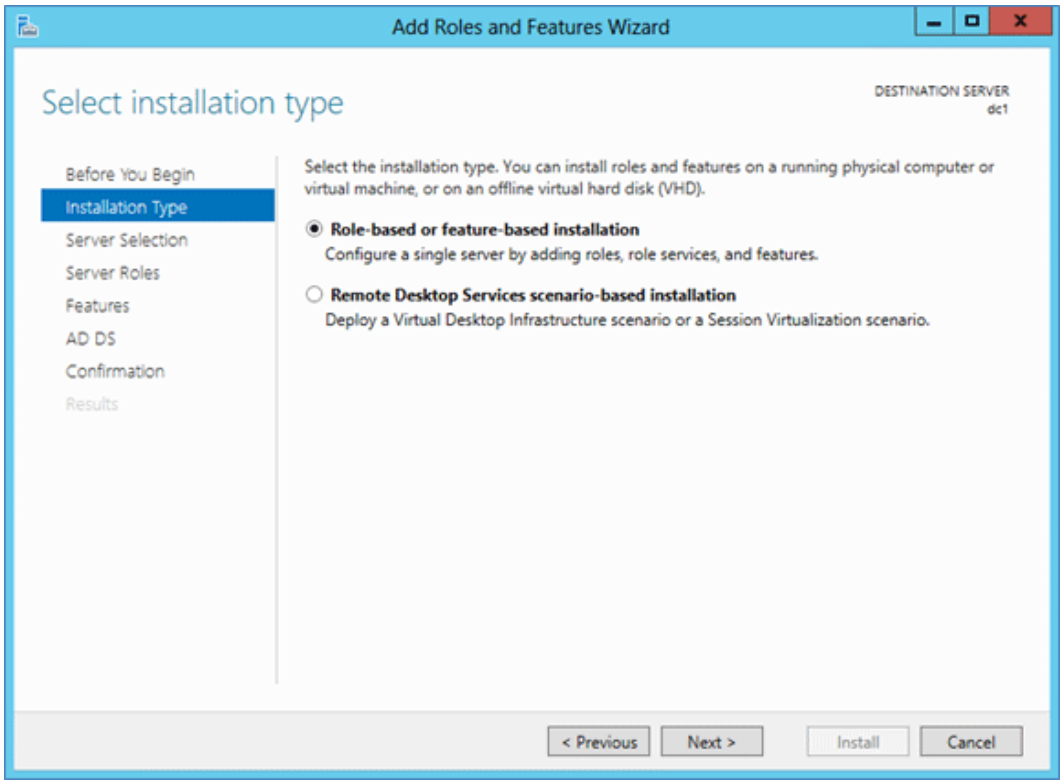

The **Installation Type** dialog provides an option that does not support Active Directory Domain Services: the **Remote Desktop Services scenario based-installation**. That option only allows

Remote Desktop Service in a multi-server distributed workload. If you select it, AD DS cannot install.

Always leave the default selection in place when installing AD DS: **Role-based or Feature-based Installation**.

Server Selection

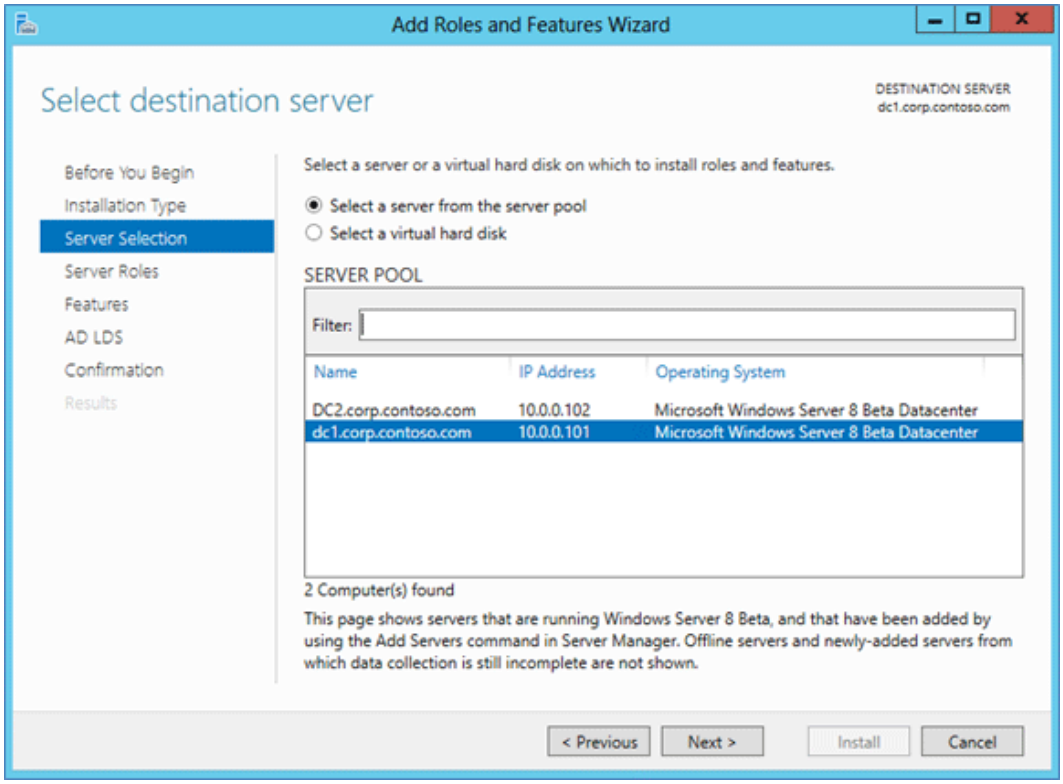

The **Server Selection** dialog enables you to choose from one of the servers previously added to the pool, as long as it is accessible. The local server running Server Manager is automatically available.

In addition, you can select offline Hyper-V VHD files with the Windows Server 2012 operating system and Server Manager adds the role to them directly through component servicing. This allows you to provision virtual servers with the necessary components before further configuring them.

Server Roles and Features

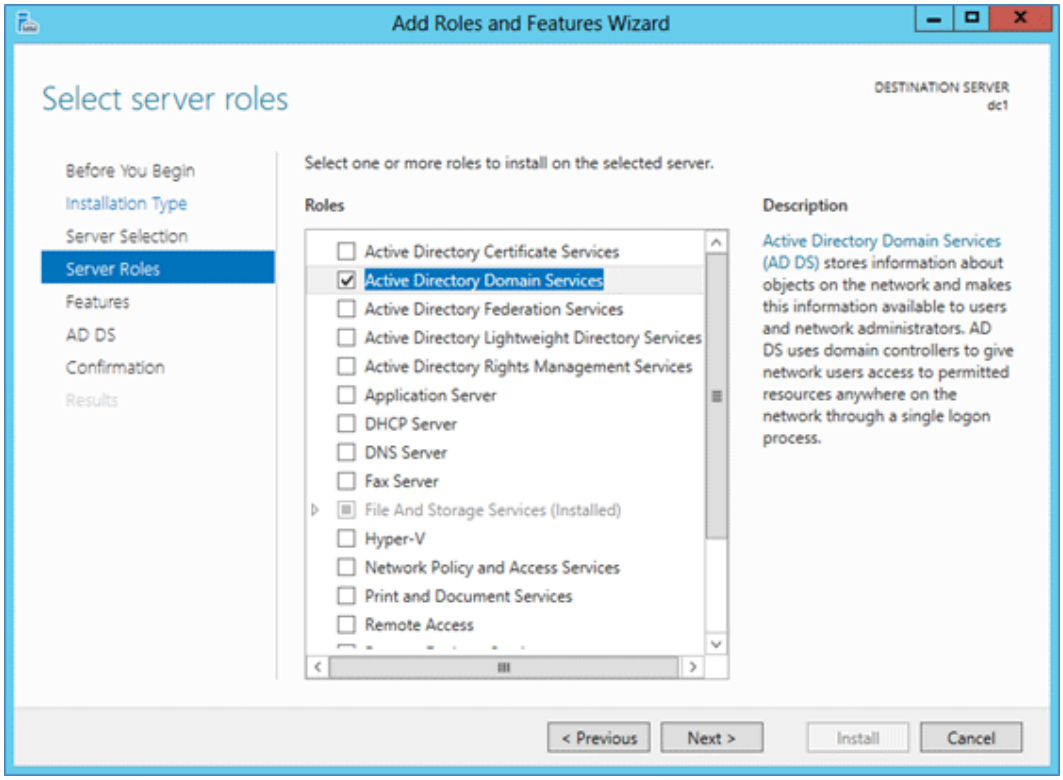

Select the **Active Directory Domain Services** role if you intend to promote a domain controller. All Active Directory administration features and required services install automatically, even if they are ostensibly part of another role or do not appear selected in the Server Manager interface.

Server Manager also presents an informational dialog that shows which management features this role implicitly installs; this is equivalent to the **-IncludeManagementTools** argument.

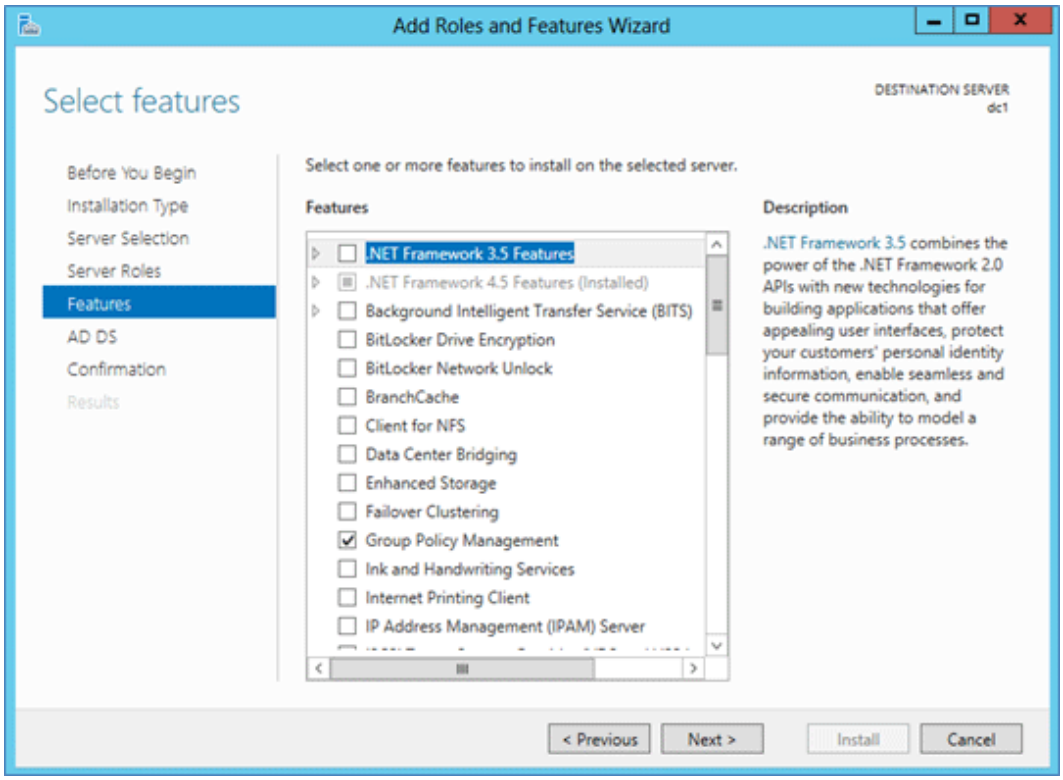

Additional **Features** can be added here as desired.

Active Directory Domain Services

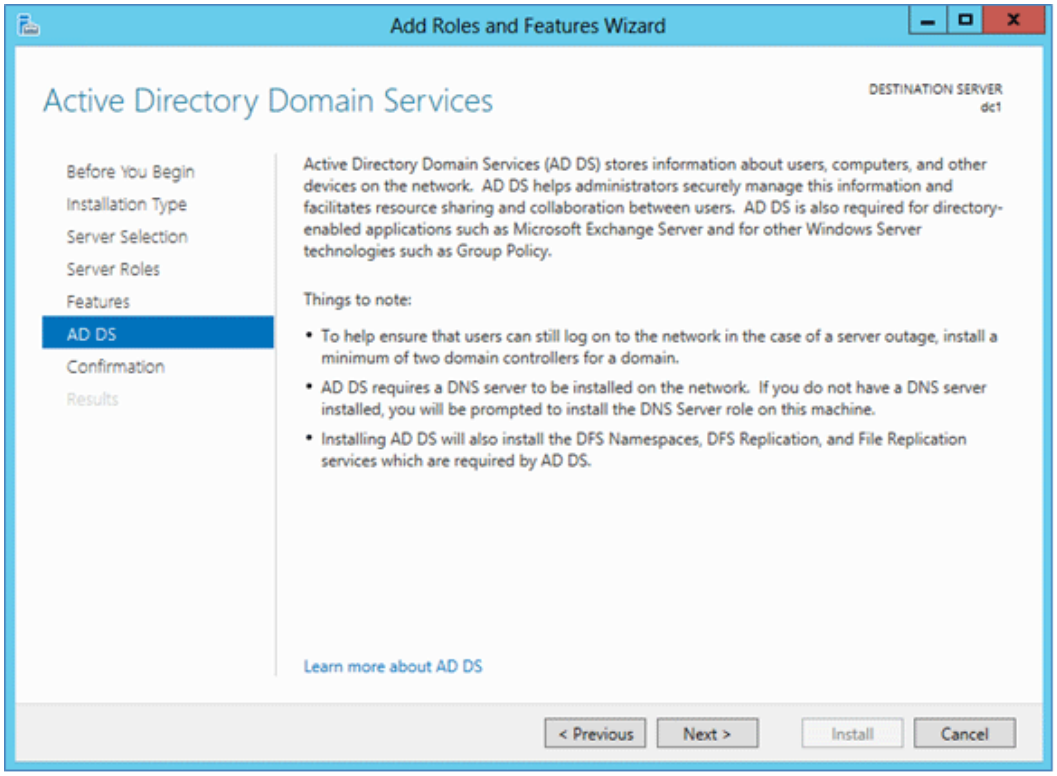

The **Active Directory Domain Services** dialog provides limited information on requirements and best practices. It mainly acts as a confirmation that you chose the AD DS role – if this screen does not appear, you did not select AD DS.

Confirmation

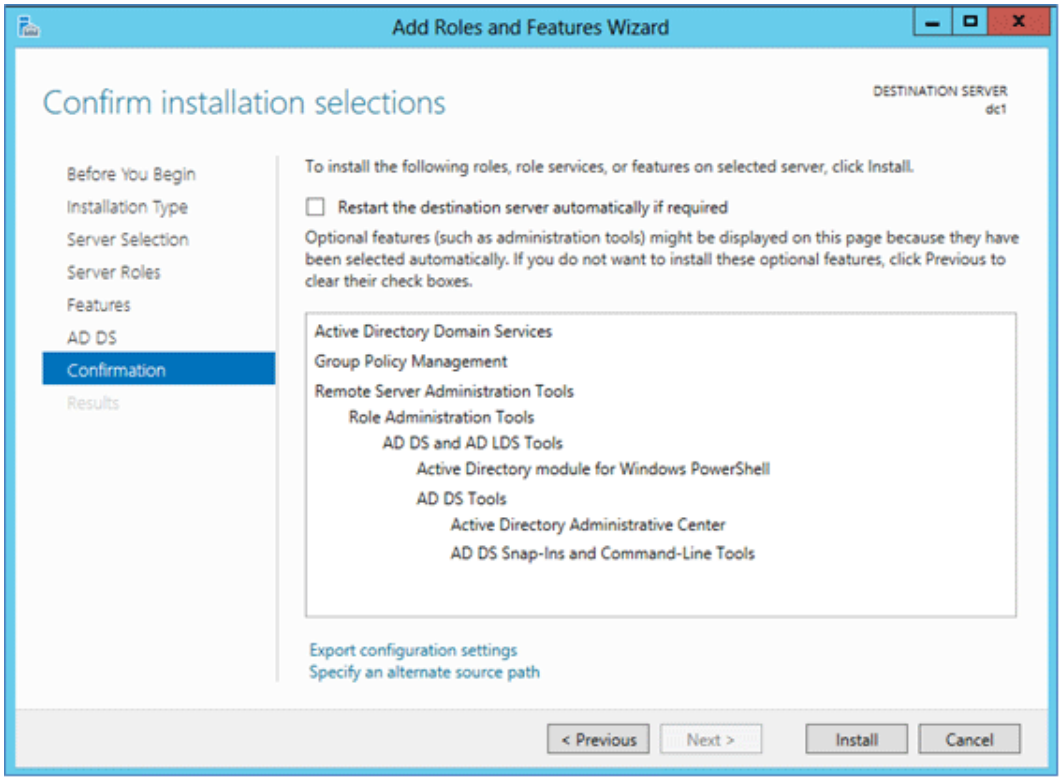

The **Confirmation** dialog is the final checkpoint before role installation starts. It offers an option to restart the computer as needed after role installation, but AD DS installation does not require a reboot.

By clicking **Install**, you confirm you are ready to begin role installation. You cannot cancel a role installation once it begins.

Results

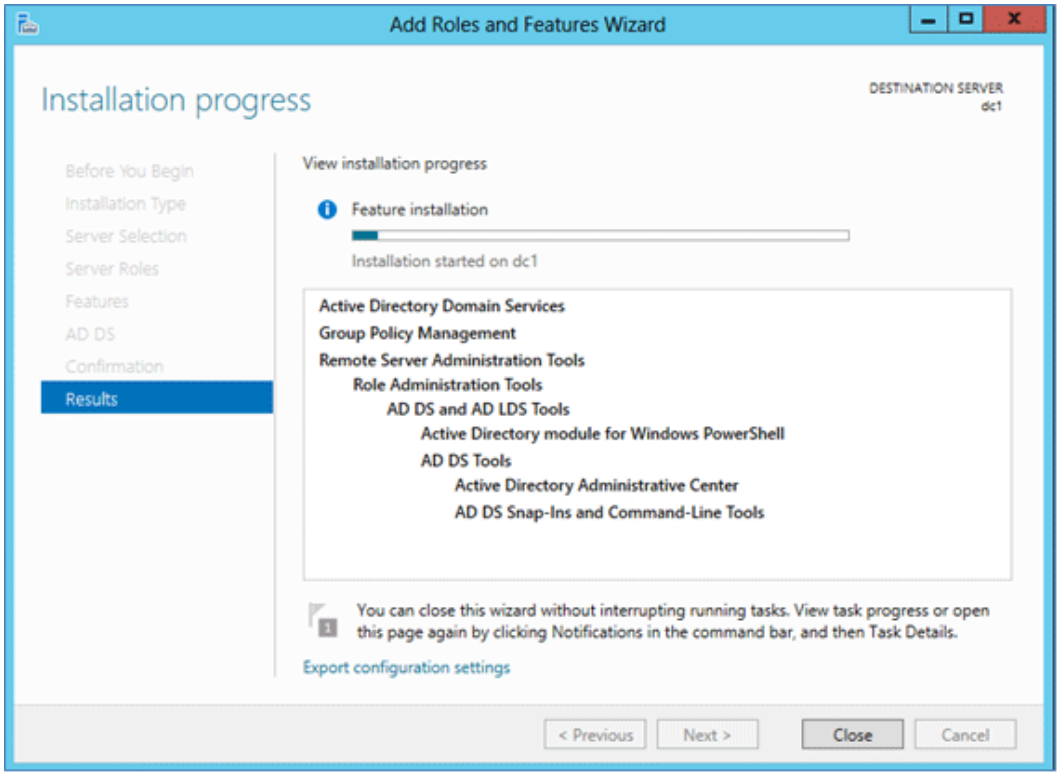

The **Results** dialog shows the current installation progress and current installation status. Role installation continues regardless of whether Server Manager is closed.

Verifying the installation results is still a best practice. If you close the **Results** dialog before installation completes, you can check the results using the Server Manager notification flag. Server Manager also shows a warning message for any servers that have installed the AD DS role but not been further configured as domain controllers.

Task Notifications

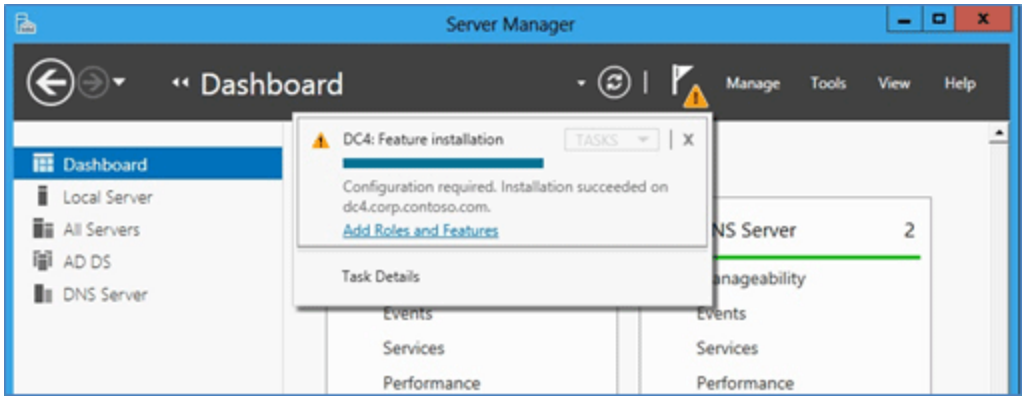

AD DS Details

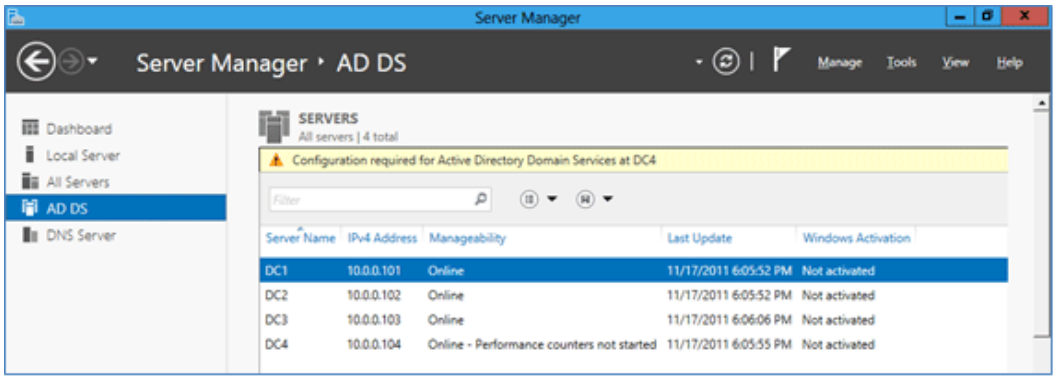

Task Details

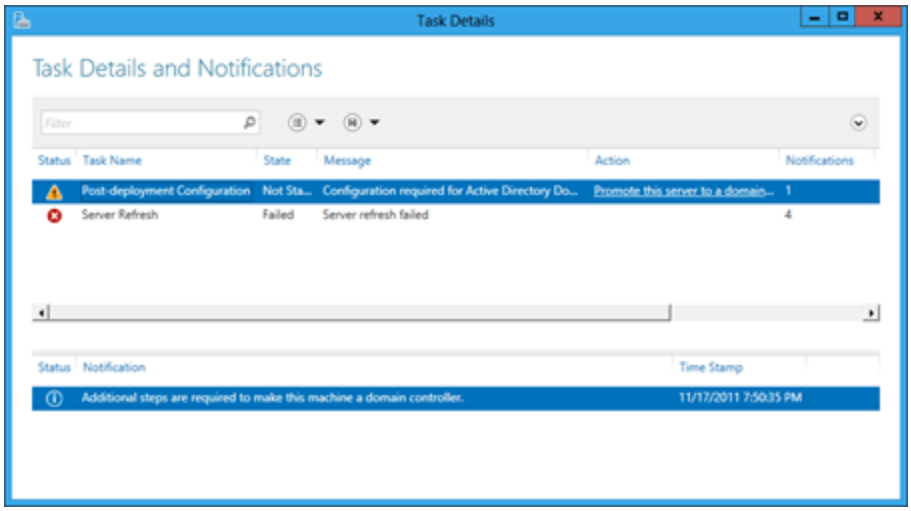

Promote to Domain Controller

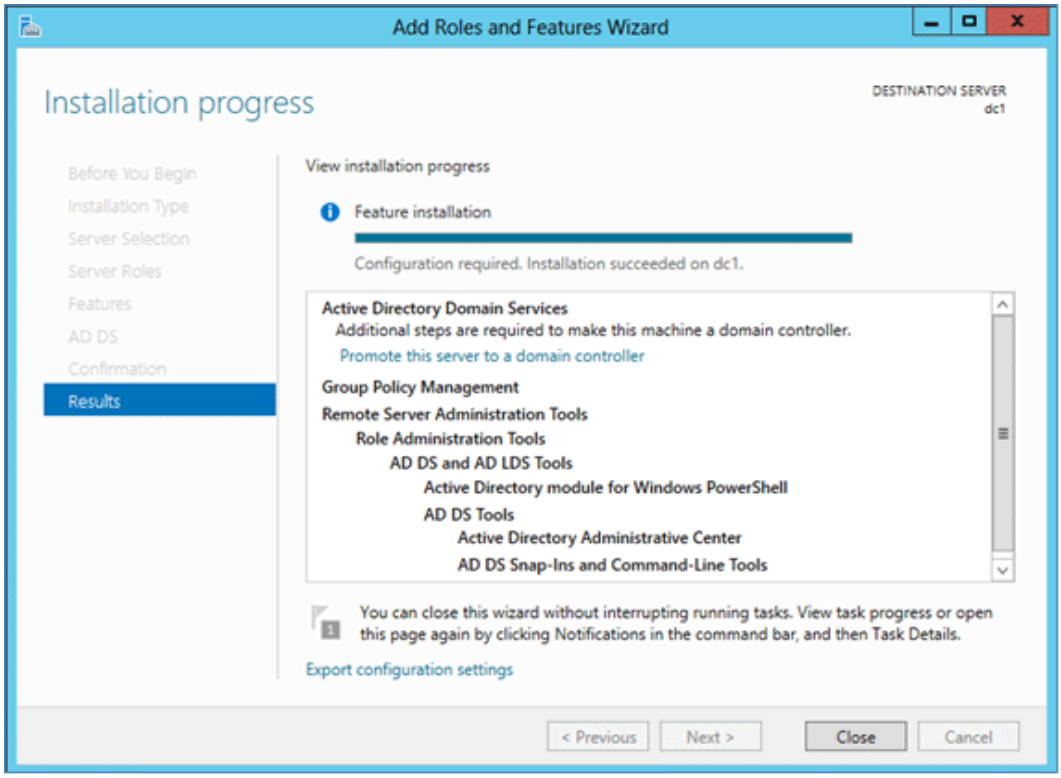

At the end of the AD DS role installation, you can continue with configuration by using the **Promote this server to a domain controller** link. This is required to make the server a domain controller, but is not necessary to run the configuration wizard immediately. For example, you may only want to provision servers with the AD DS binaries before sending them to another branch office for later configuration. By adding the AD DS role before shipping, you save time when it reaches its destination. You also follow the best practice of not keeping a domain controller offline for days or weeks. Finally, this enables you to update components before domain controller promotion, saving you at least one subsequent reboot.

Selecting this link later invokes the ADDSDeployment cmdlets: **install-addsforest**, **install-addsdomain**, or **install-addsdomaincontroller**.

Uninstalling/Disabling

You remove the AD DS role like any other role, regardless of whether you promoted the server to a domain controller. However, removing the AD DS role requires a restart on completion.

Active Directory Domain Services role removal is different from installation, in that it requires domain controller demotion before it can complete. This is necessary to prevent a domain controller from having its role binaries uninstalled without proper metadata cleanup in the forest. For more information, see Demoting Domain Controllers and Domains (Level 200).

Warning

Removing the AD DS roles with Dism.exe or the Windows PowerShell DISM module after promotion to a Domain Controller is not supported and will prevent the server from booting normally.

Unlike Server Manager or the AD DS Deployment module for Windows PowerShell, DISM is a native servicing system that has no inherent knowledge of AD DS or its configuration. Do not use Dism.exe or the Windows PowerShell DISM module to uninstall the AD DS role unless the server is no longer a domain controller.

Create an AD DS Forest Root Domain with Server Manager

The following diagram illustrates the Active Directory Domain Services configuration process, in the case where you have previously installed the AD DS role and started the **Active Directory Domain Services Configuration Wizard** using Server Manager.

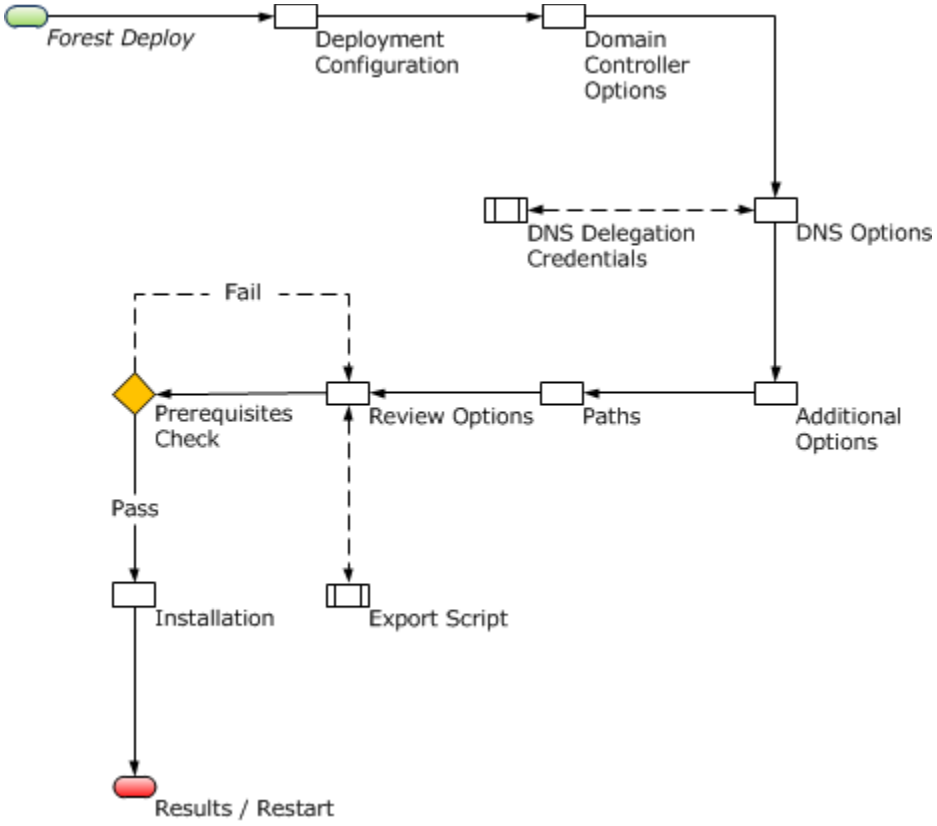

Deployment Configuration

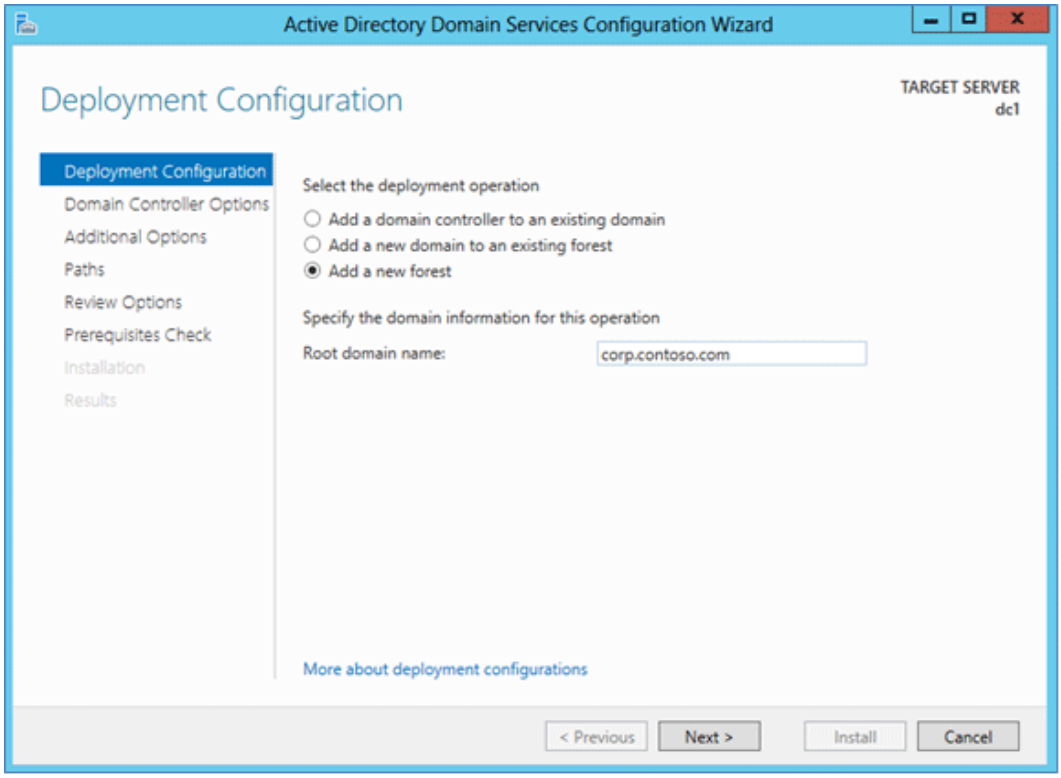

Server Manager begins every domain controller promotion with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select.

To create a new Active Directory forest, click **Add a new forest**. You must provide a valid root domain name; the name cannot be single-labeled (for example, the name must be *contoso.com* or similar and not just *contoso*) and must use allowed DNS domain naming requirements.

For more information on valid domain names, see KB article Naming conventions in Active Directory for computers, domains, sites, and OUs.

Warning

Do not create new Active Directory forests with the same name as an external DNS name. For example, if your Internet DNS URL is http://contoso.com, you must choose a different name for your internal forest to avoid future compatibility issues. That name should be unique and unlikely for web traffic. For example: corp.contoso.com.

A new forest does not need new credentials for the domain's Administrator account. The domain controller promotion process uses the credentials of the built-in Administrator account from the first domain controller used to create the forest root. There is no way (by default) to disable or lock out the built-in Administrator account and it may be the only entry point into a forest if the other administrative domain accounts are unusable. It is critical to know the password before deploying a new forest.

DomainName requires a valid fully qualified domain DNS name and is required.

Domain Controller Options

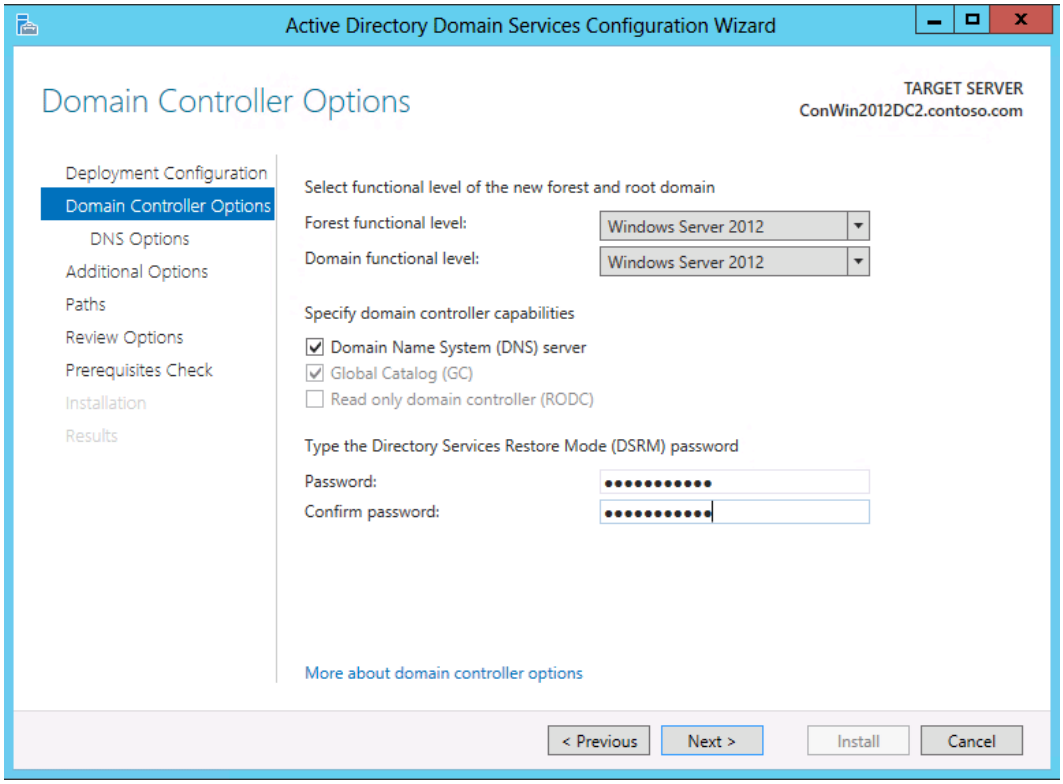

The **Domain Controller Options** enables you to configure the **forest functional level** and **domain functional level** for the new forest root domain. By default, these settings are Windows Server 2012 in a new forest root domain. The Windows Server 2012 forest functional level does not provide any new functionality over the Windows Server 2008 R2 forest functional level. The Windows Server 2012 domain functional level is required only in order to implement the new Kerberos settings "always provide claims" and "Fail unarmored authentication requests." A primary use for functional levels in Windows Server 2012 is to restrict participation in the domain to domain controllers that meet minimum-allowed operating system requirements. In other words, you can specify Windows Server 2012 domain functional level only domain controllers that run Windows Server 2012 can host the domain. Windows Server 2012 implements a new domain controller flag called **DS_WIN8_REQUIRED** in the **DSGetDcName** function of NetLogon that exclusively locates Windows Server 2012 domain controllers. This allows you the flexibility of a more homogeneous or heterogeneous forest in terms of which operating systems are permitted to be run on domain controllers.

For more information about domain controller Location, review Directory Service Functions.

The only configurable domain controller capability is the DNS server option. Microsoft recommends that all domain controllers provide DNS services for high availability in distributed environments, which is why this option is selected by default when installing a domain controller

in any mode or domain. The Global Catalog and read only domain controller options are unavailable when creating a new forest root domain; the first domain controller must be a GC, and cannot be a read only domain controller (RODC).

The specified **Directory Services Restore Mode Password** must adhere to the password policy applied to the server, which by default does not require a strong password; only a non-blank one. Always choose a strong, complex password or preferably, a passphrase.

DNS Options and DNS Delegation Credentials

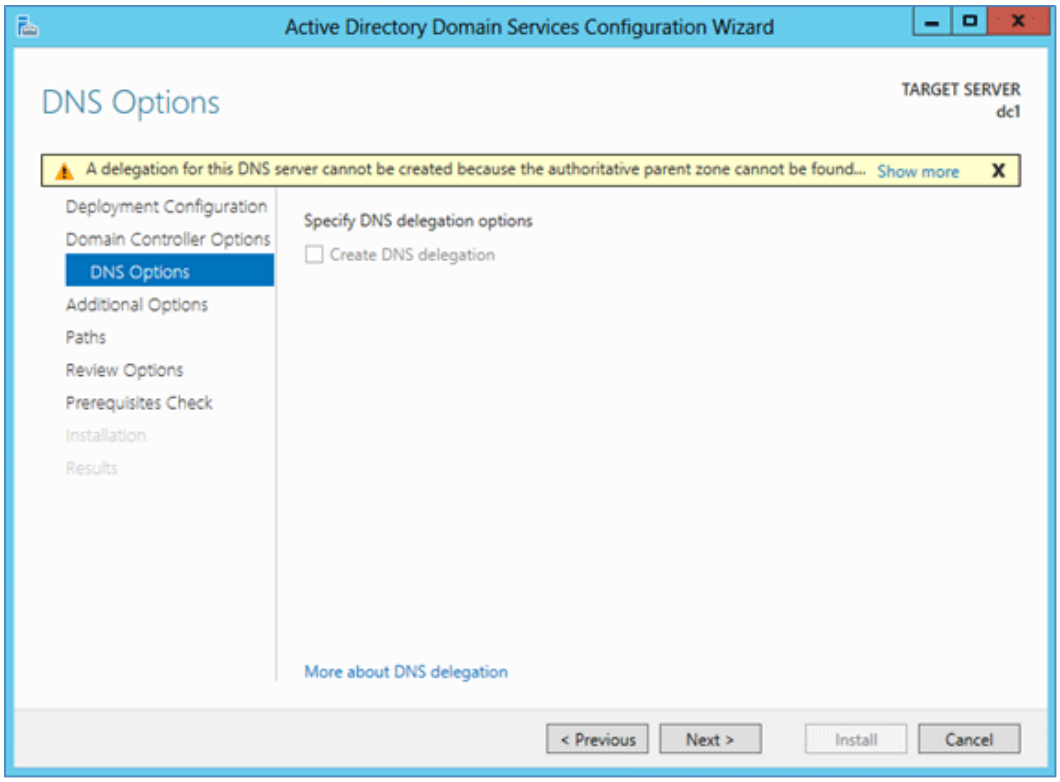

The **DNS Options** page enables you to configure DNS delegation and provide alternate DNS administrative credentials.

You cannot configure DNS options or delegation in the Active Directory Domain Services Configuration Wizard when installing a new Active Directory Forest Root Domain where you selected the **DNS server** on the **Domain Controller Options** page. The **Create DNS delegation** option is available when creating a new forest root DNS zone in an existing DNS server infrastructure. This option enables you to provide alternate DNS administrative credentials that have the rights to update DNS zone.

For more information about whether you need to create a DNS delegation, see Understanding Zone Delegation.

Additional Options

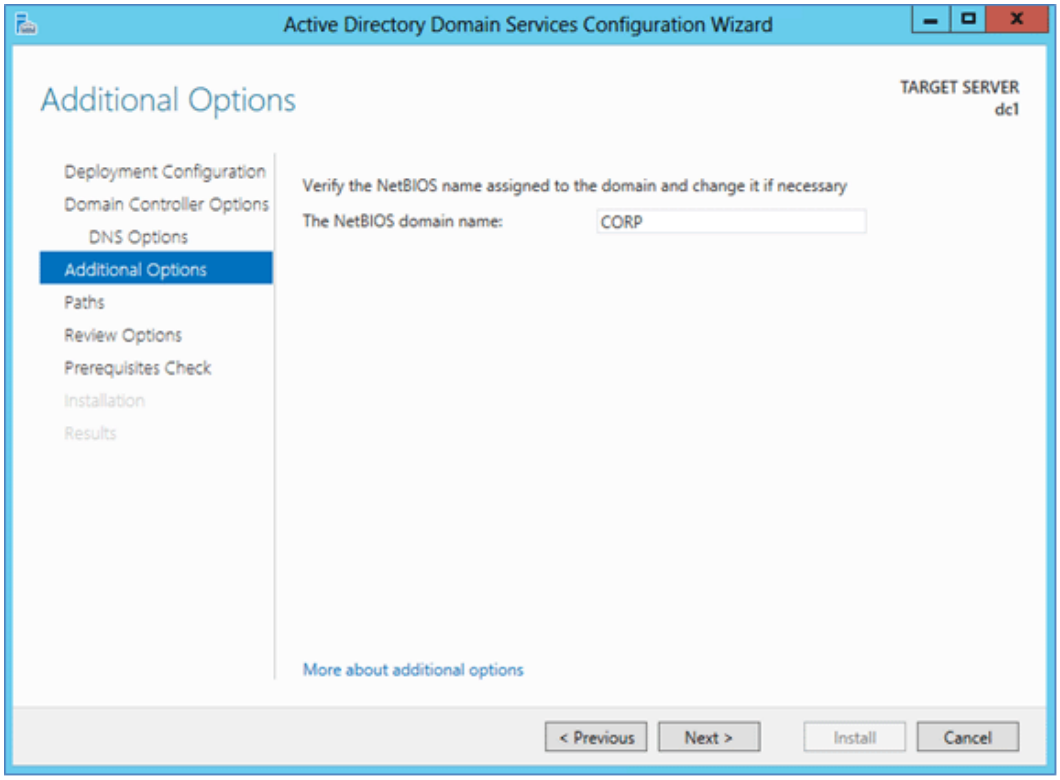

The **Additional Options** page shows the NetBIOS name of the domain and enables you to override it. By default, the NetBIOS domain name matches the left-most label of the fully qualified domain name provided on the **Deployment Configuration** page. For example, if you provided the fully qualified domain name of corp.contoso.com, the default NetBIOS domain name is CORP.

If the name is 15 characters or less and does not conflict with another NetBIOS name, it is unaltered. If it does conflict with another NetBIOS name, a number is appended to the name. If the name is more than 15 characters, the wizard provides a unique, truncated suggestion. In either case, the wizard first validates the name is not already in use via a WINS lookup and NetBIOS broadcast.

For more information on valid domain names, see KB article Naming conventions in Active Directory for computers, domains, sites, and OUs.

Paths

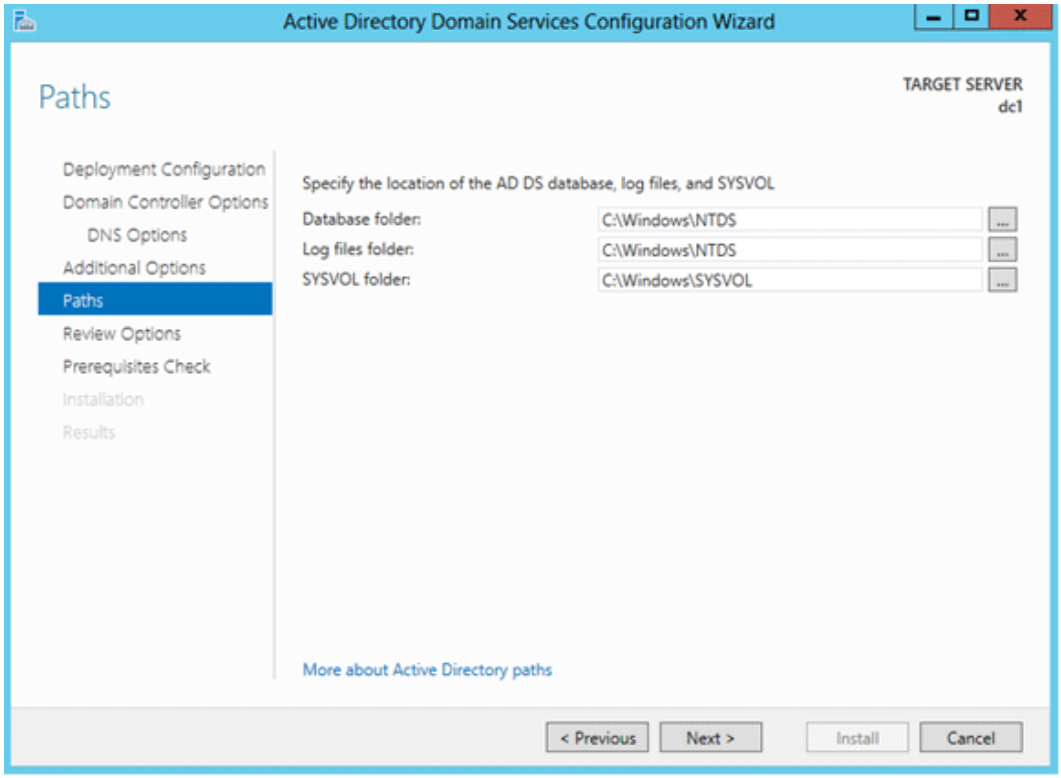

The **Paths** page enables you to override the default folder locations of the AD DS database, the database transaction logs, and the SYSVOL share. The default locations are always in subdirectories of %systemroot% (i.e. C:\Windows).

Review Options and View Script

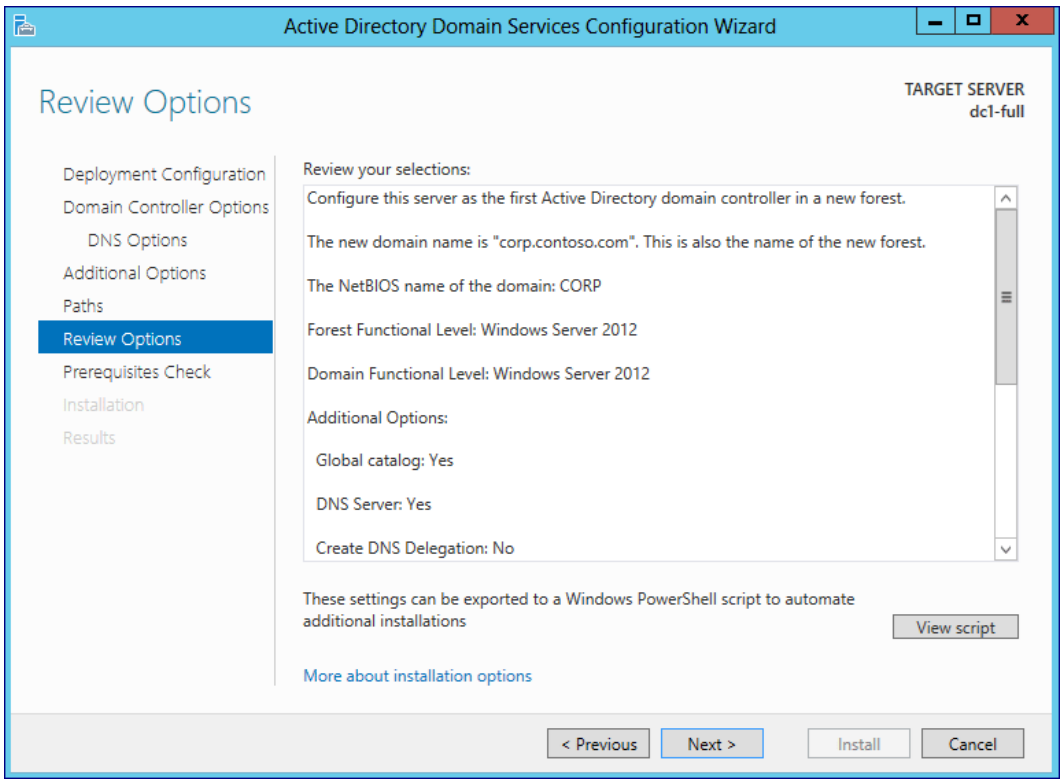

The **Review Options** page enables you to validate your settings and ensure they meet your requirements before you start the installation. This is not the last opportunity to stop the installation when using Server Manager. This is simply an option to confirm your settings before continuing the configuration

The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use. For example:

```
#

# Windows PowerShell Script for AD DS Deployment

#


Import-Module ADDSDeployment

Install-ADDSForest `

-CreateDNSDelegation `
```

```
-DatabasePath "C:\Windows\NTDS" `
-DomainMode "Win2012" `
-DomainName "corp.contoso.com" `
-DomainNetBIOSName "CORP" `
-ForestMode "Win2012" `
-InstallDNS:$true `
-LogPath "C:\Windows\NTDS" `
-NoRebootOnCompletion:$false `
-SYSVOLPath "C:\Windows\SYSVOL"
-Force:$true
```

Note

Server Manager generally fills in all arguments with values when promoting and does not rely on defaults (as they may change between future versions of Windows or service packs). The one exception to this is the **-safemodeadministratorpassword** argument (which is deliberately omitted from the script). To force a confirmation prompt, omit the value when running cmdlet interactively.

Prerequisites Check

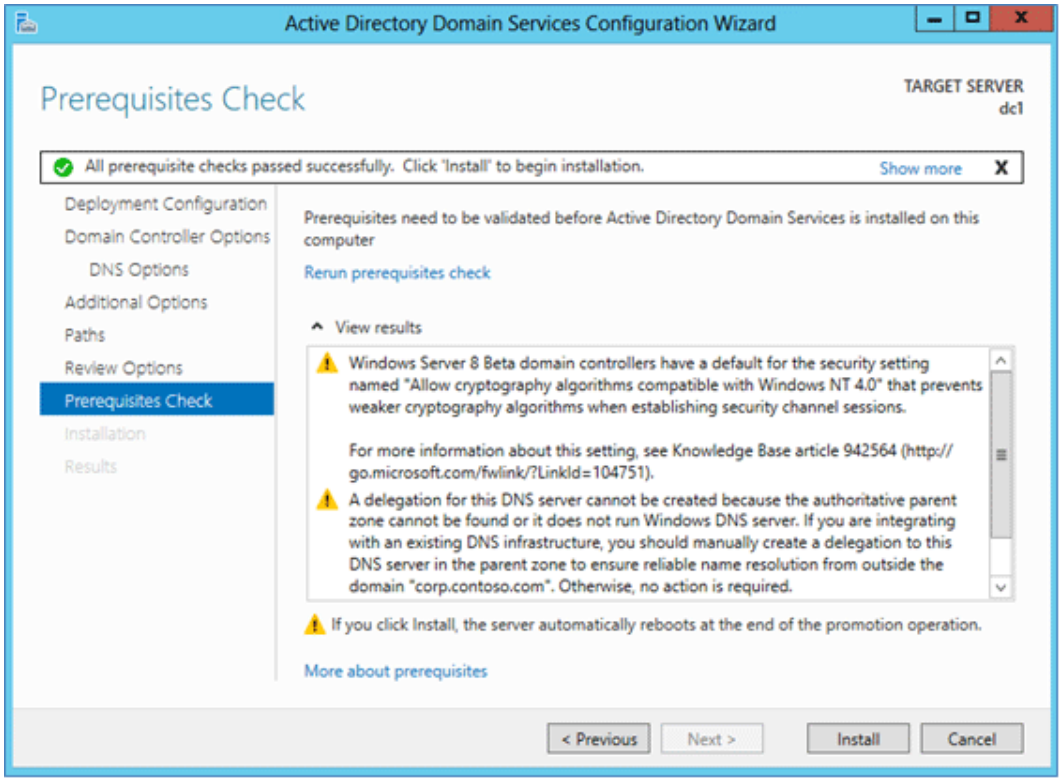

The **Prerequisites Check** is a new feature in AD DS domain configuration. This new phase validates that the server configuration is capable of supporting a new AD DS forest.

When installing a new forest root domain, the Server Manager Active Directory Domain Services Configuration Wizard invokes a series of modular tests. These tests alert you with suggested repair options. You can run the tests as many times as required. The domain controller process cannot continue until all prerequisite tests pass.

The **Prerequisites Check** also surfaces relevant information such as security changes that affect older operating systems.

For more information on the specific prerequisite checks, see Prerequisite Checking.

Installation

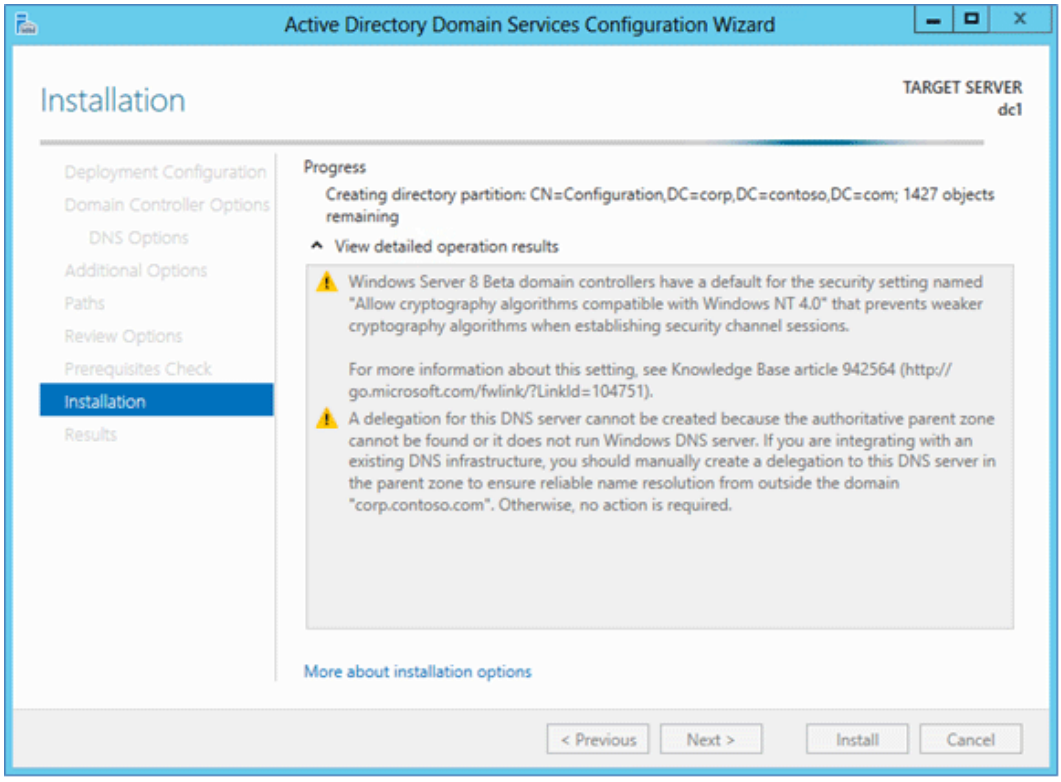

When the **Installation** page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and are written to logs:

- %systemroot%\debug\dcpromo.log
- %systemroot%\debug\dcpromoui.log

Note

You can run multiple role installation and AD DS configuration wizards from the same Server Manager console simultaneously.

Results

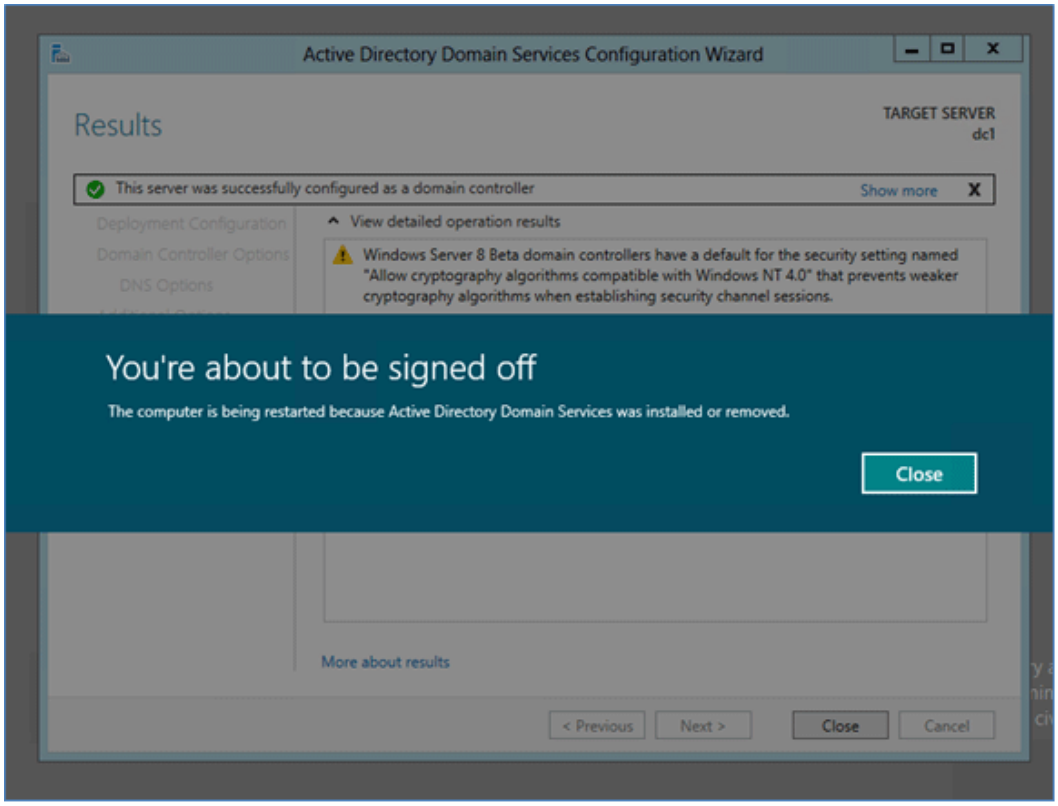

The **Results** page shows the success or failure of the promotion and any important administrative information. The domain controller will automatically reboot after 10 seconds.

Deploying a Forest with Windows PowerShell

This section explains how to install the first domain controller in a forest root domain using Windows PowerShell on a Core Windows Server 2012 computer.

Windows PowerShell AD DS Role Installation Process

By implementing a few straightforward ServerManager deployment cmdlets into your deployment processes, you further realize the vision of AD DS simplified administration.

The next figure illustrates the Active Directory Domain Services role installation process, beginning with you running **PowerShell.exe** and ending right before the promotion of the domain controller.

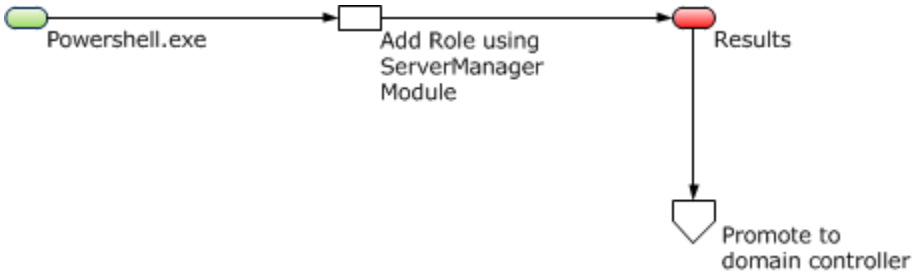

ServerManager Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-WindowsFeature/Add-WindowsFeature	***-Name*** *-Restart* *-IncludeAllSubFeature* *-IncludeManagementTools* -Source *-ComputerName* -Credential -LogPath *-Vhd* *-ConfigurationFilePath*

Note

While not required, the argument **-IncludeManagementTools** is highly recommended when installing the AD DS role binaries

The ServerManager module exposes role installation, status, and removal portions of the new DISM module for Windows PowerShell. This layering simplifies the most tasks and reduces need for direct usage of the powerful (but dangerous when misused) DISM module.

Use **Get-Command** to export the aliases and cmdlets in ServerManager.

```
Get-Command -module ServerManager
```

For example:

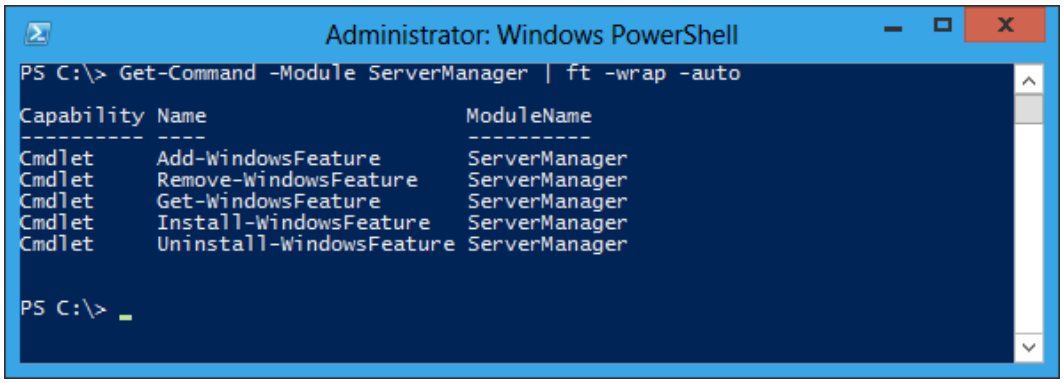

To add the Active Directory Domain Services role, simply run the **Install-WindowsFeature** with the AD DS role name as an argument. Like Server Manager, all required services implicit to the AD DS role install automatically.

```
Install-WindowsFeature -name AD-Domain-Services
```

If you also want the AD DS management tools installed - and this is highly recommended - then provide the **-IncludeManagementTools** argument:

```
Install-WindowsFeature -name AD-Domain-Services -IncludeManagementTools
```

For example:

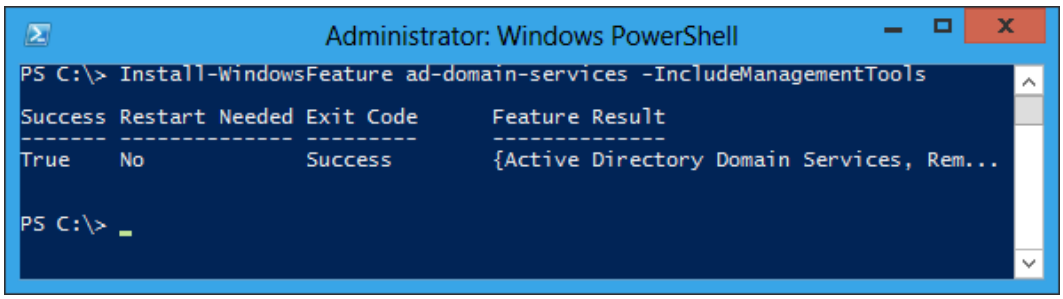

To list all features and roles with their installation status, use **Get-WindowsFeature** without arguments. Specify **-ComputerName** argument for the installation status from a remote server.

```
Get-WindowsFeature
```

Because **Get-WindowsFeature** does not have a filtering mechanism, you must use **Where-Object** with a pipeline to find specific features. The pipeline is a channel used between multiple cmdlets to pass data and the Where-Object cmdlet acts as a filter. The built-in **$_** variable acts as the current object passing through the pipeline with any properties it may contain.

```
Get-WindowsFeature | where-object <options>
```

For example, to find all features containing "Active Dir" in their **Display Name** property, use:

```
Get-WindowsFeature | where displayname -like "*active dir*"
```

Further examples illustrated below:

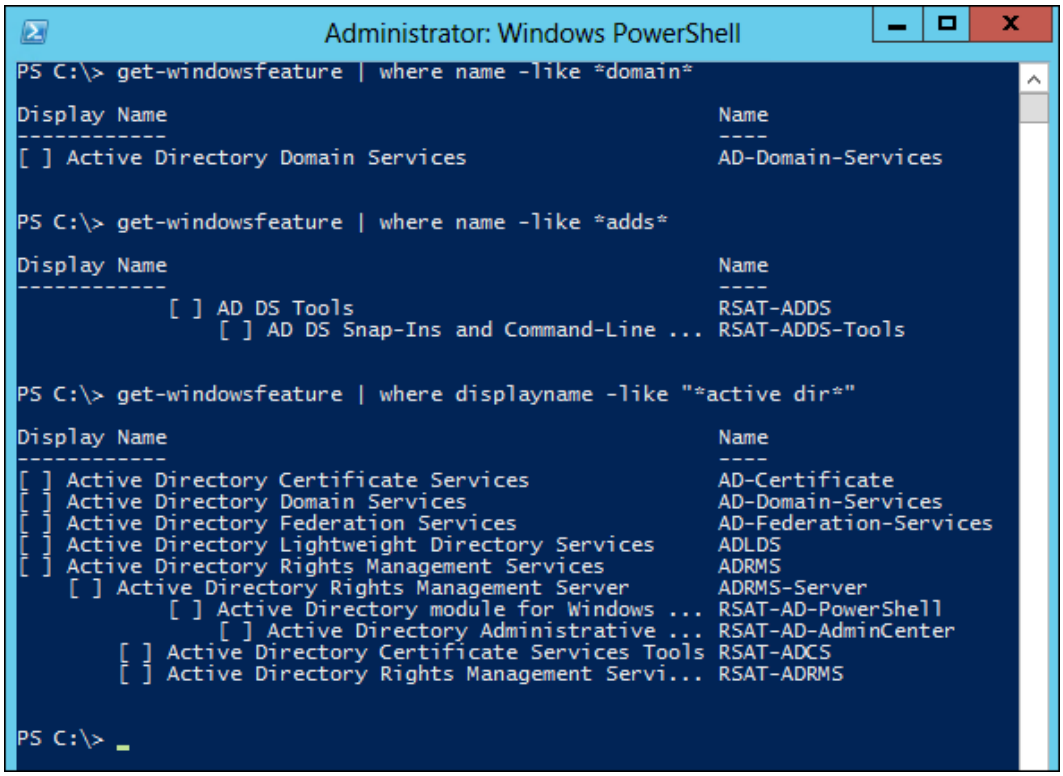

For more information about more Windows PowerShell operations with pipelines and Where-Object, see [Piping and the Pipeline in Windows PowerShell].

Note also that Windows PowerShell 3.0 significantly simplified the command-line arguments needed in this pipeline operation. Windows PowerShell 2.0 would have required:

```
Get-WindowsFeature | where {$_.displayname - like "*active dir*"}
```

By using the Windows PowerShell pipeline, you can create readable results. For example:

```
Install-WindowsFeature | Format-List

Install-WindowsFeature | select-object | Format-List
```

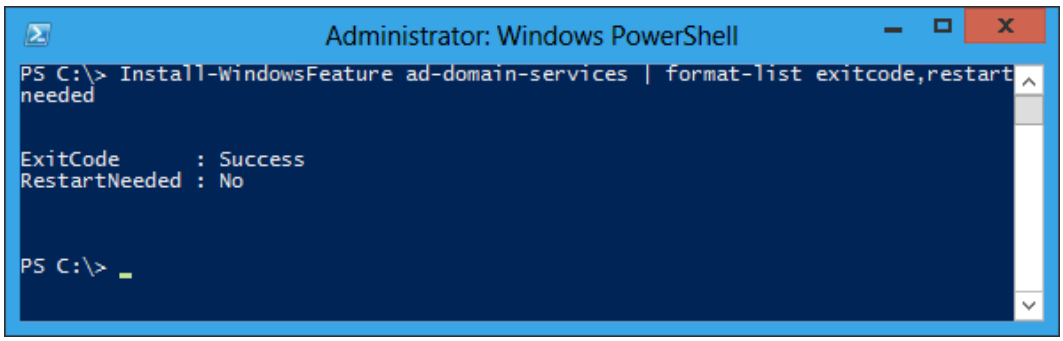

Note how using the **Select-Object** cmdlet with the **-expandproperty** argument returns interesting data:

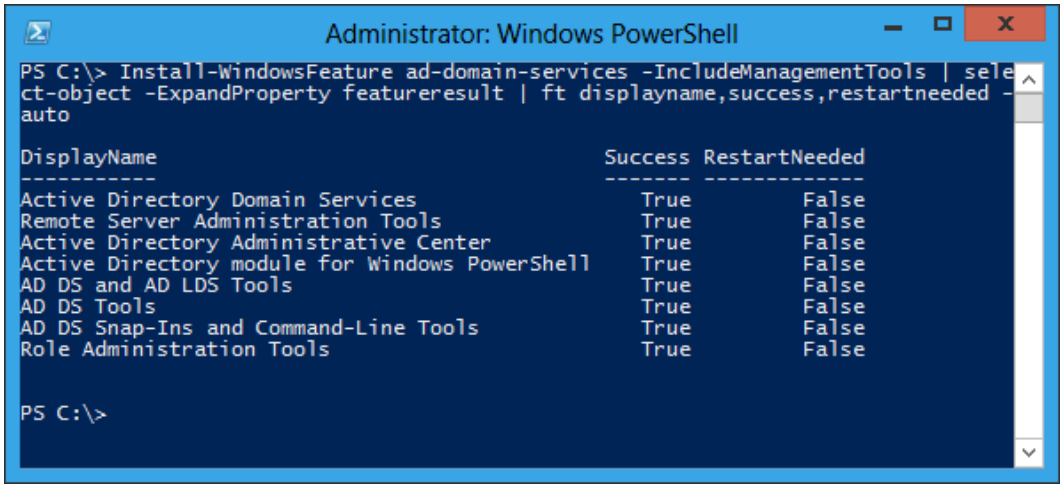

Note

The **Select-Object -expandproperty** argument slows down overall installation performance slightly.

Create an AD DS Forest Root Domain with Windows PowerShell

To install a new Active Directory forest using the ADDSDeployment module, use the following cmdlet:

```
Install-addsforest
```

The **Install-AddsForest** cmdlet only has two phases (prerequisite checking and installation). The two figures below show the installation phase with the minimum required argument of **-domainname**.

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-Addsforest	-Confirm *-CreateDNSDelegation* *-DatabasePath* *-DomainMode* ***-DomainName*** ***-DomainNetBIOSName*** *-DNSDelegationCredential*

	-ForestMode -Force *-InstallDNS* *-LogPath* -NoDnsOnNetwork -NoRebootOnCompletion *-SafeModeAdministratorPassword* -SkipAutoConfigureDNS -SkipPreChecks *-SYSVOLPath* *-Whatif*

Note

The **-DomainNetBIOSName** argument is required if you want to change the automatically generated 15-character name based on the DNS domain name prefix or if the name exceeds 15 characters.

The equivalent Server Manager **Deployment Configuration** ADDSDeployment cmdlet and arguments are:

```
Install-ADDSForest

-DomainName <string>
```

The equivalent Server Manager Domain Controller Options ADDSDeployment cmdlet arguments are:

```
-ForestMode <{Win2003 | Win2008 | Win2008R2 | Win2012 | Default}>

-DomainMode <{Win2003 | Win2008 | Win2008R2 | Win2012 | Default}>

-InstallDNS <{$false | $true}>

-SafeModeAdministratorPassword <secure string>
```

The **Install-ADDSForest** arguments follow the same defaults as Server Manager if not specified.

The **SafeModeAdministratorPassword** argument's operation is special:

- If *not specified* as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.

  For example, to create a new forest named corp.contoso.com and be prompted to enter and confirm a masked password:

  ```
  Install-ADDSForest -DomainName corp.contoso.com
  ```

- If specified *with a value*, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string:

```
-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged.

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Finally, you could store the obfuscated password in a file, and then reuse it later, without the clear text password ever appearing. For example:

```
$file = "c:\pw.txt"

$pw = read-host -prompt "Password:" -assecurestring

$pw | ConvertFrom-SecureString | Set-Content $file


-safemodeadministratorpassword (Get-Content $File | ConvertTo-SecureString)
```

Warning

Providing or storing a clear or obfuscated text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM password of that domain controller. Anyone with access to the file could reverse that obfuscated password. With that knowledge, they can logon to a DC started in DSRM and eventually impersonate the domain controller itself, elevating their privileges to the highest level in an Active Directory forest. An additional set of steps using **System.Security.Cryptography** to encrypt the text file data is advisable but out of scope. The best practice is to totally avoid password storage.

The ADDSDeployment cmdlet offers an additional option to skip automatic configuration of DNS client settings, forwarders, and root hints. You cannot skip this configuration option when using Server Manager. This argument matters only if you installed the DNS Server role prior to configuring the domain controller:

```
-SkipAutoConfigureDNS
```

The **DomainNetBIOSName** operation is also special:

- If the **DomainNetBIOSName** argument is not specified with a NetBIOS domain name and the single-label prefix domain name in the **DomainName** argument is 15 characters or fewer, then promotion continues with an automatically generated name.
- If the **DomainNetBIOSName** argument is not specified with a NetBIOS domain name and the single-label prefix domain name in the **DomainName** argument is 16 characters or more, then promotion fails.

- If the **DomainNetBIOSName** argument is specified with a NetBIOS domain name of 15 characters or fewer, then promotion continues with that specified name.
- If the **DomainNetBIOSName** argument is specified with a NetBIOS domain name of 16 characters or more, then promotion fails.

The equivalent Server Manager Additional Options ADDSDeployment cmdlet argument is:

```
-domainnetbiosname <string>
```

The equivalent Server Manager **Paths** ADDSDeployment cmdlet arguments are:

```
-databasepath <string>
```

```
-logpath <string>
```

```
-sysvolpath <string>
```

Use the optional **Whatif** argument with the **Install-ADDSForest** cmdlet to review configuration information. This enables you to see the explicit and implicit values of a cmdlet's arguments.

For example:

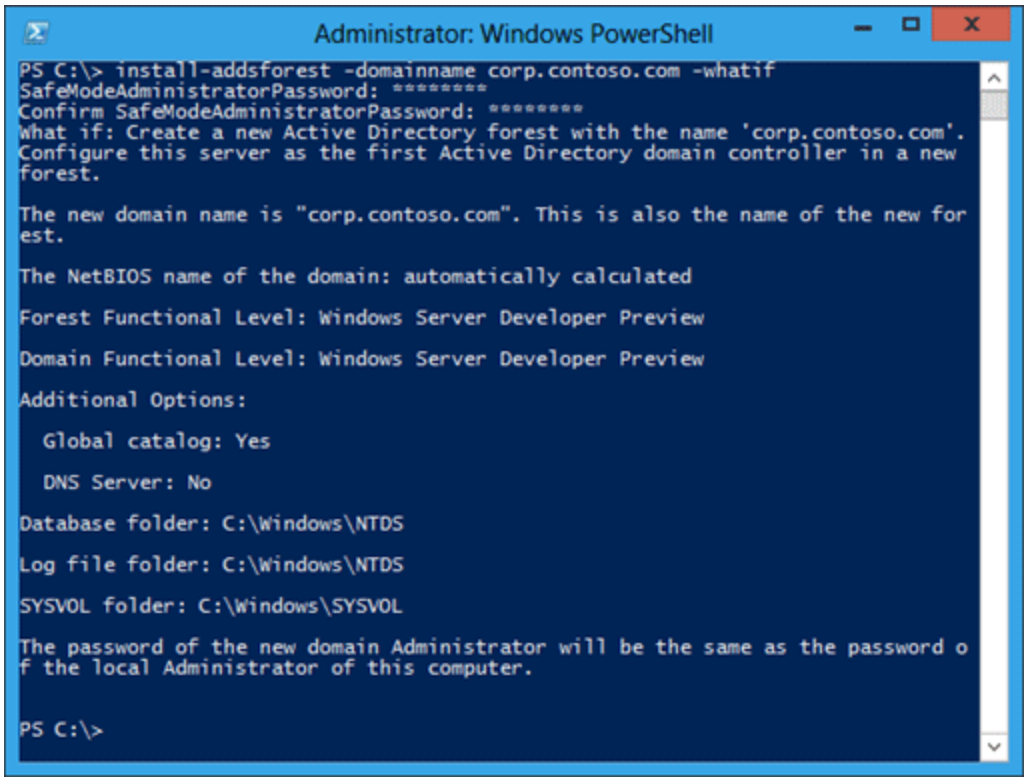

You cannot bypass the **Prerequisite Check** when using Server Manager, but you can skip the process when using the AD DS Deployment cmdlet using the following argument:

```
-skipprechecks
```

Warning

Microsoft discourages skipping the prerequisite check as it can lead to a partial domain controller promotion or damaged AD DS forest.

Note how, just like Server Manager, **Install-ADDSForest** reminds you that promotion will reboot the server automatically.

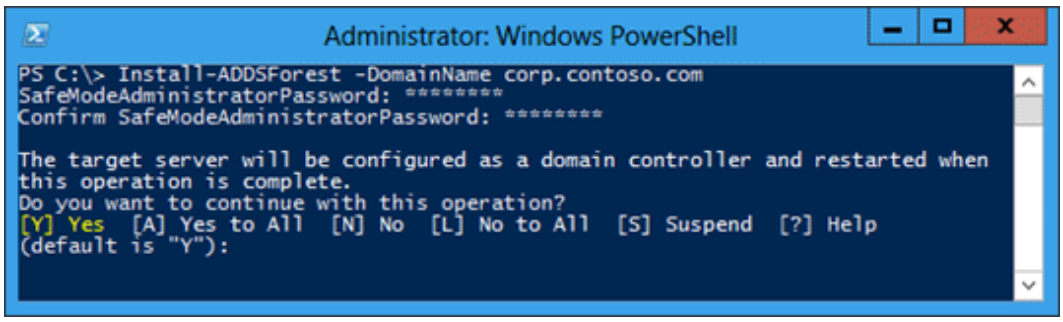

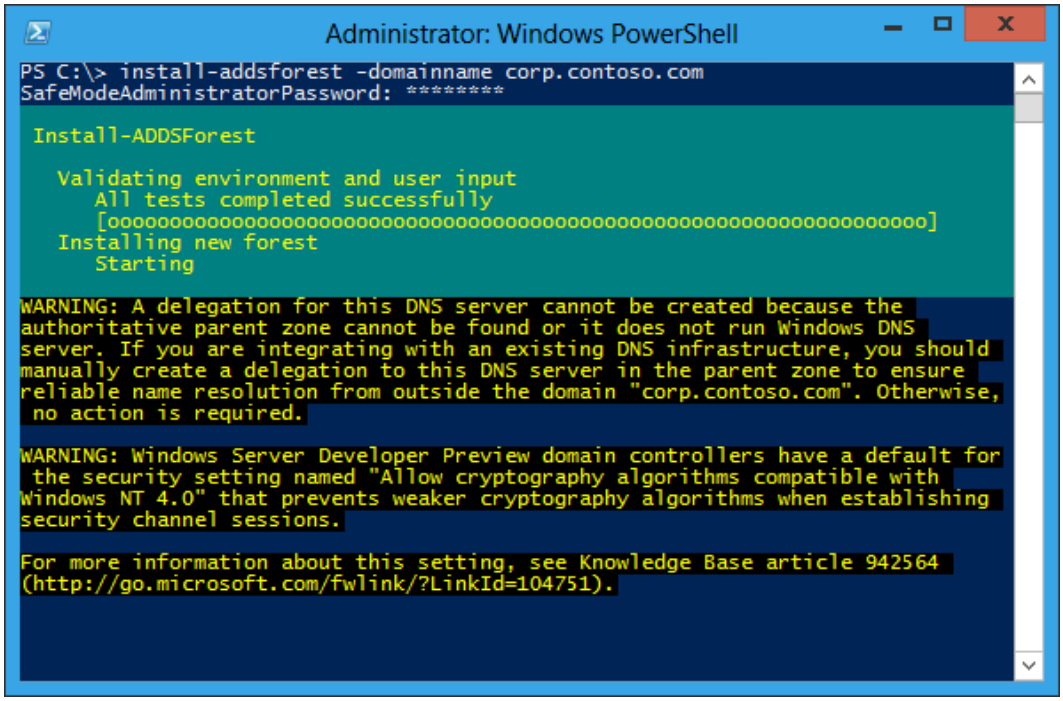

To accept the reboot prompt automatically, use the **-force** or **-confirm:$false** arguments with any ADDSDeployment Windows PowerShell cmdlet. To prevent the server from automatically rebooting at the end of promotion, use the **-norebootoncompletion** argument.

Warning

Overriding the reboot is discouraged. The domain controller must reboot to function correctly.

See Also

Active Directory Domain Services (TechNet Portal)

Active Directory Domain Services for Windows Server 2008 R2

Active Directory Domain Services for Windows Server 2008

Windows Server Technical Reference (Windows Server 2003)

Active Directory Administrative Center: Getting Started (Windows Server 2008 R2)

Active Directory Administration with Windows PowerShell (Windows Server 2008 R2)

Ask the Directory Services Team (Official Microsoft Commercial Technical Support Blog)

Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200)

This topic covers the steps necessary to upgrade an existing forest or domain to Windows Server 2012, using either Server Manager or Windows PowerShell. It covers how to add domain controllers that run Windows Server 2012 to an existing domain.

- Upgrade and Replica Workflow
- Upgrade and Replica Windows PowerShell
- Deployment

Upgrade and Replica Workflow

The following diagram illustrates the Active Directory Domain Services configuration process when you previously installed the AD DS role and you have started the Active Directory Domain Services Configuration Wizard using Server Manager to create a new domain controller in an existing domain.

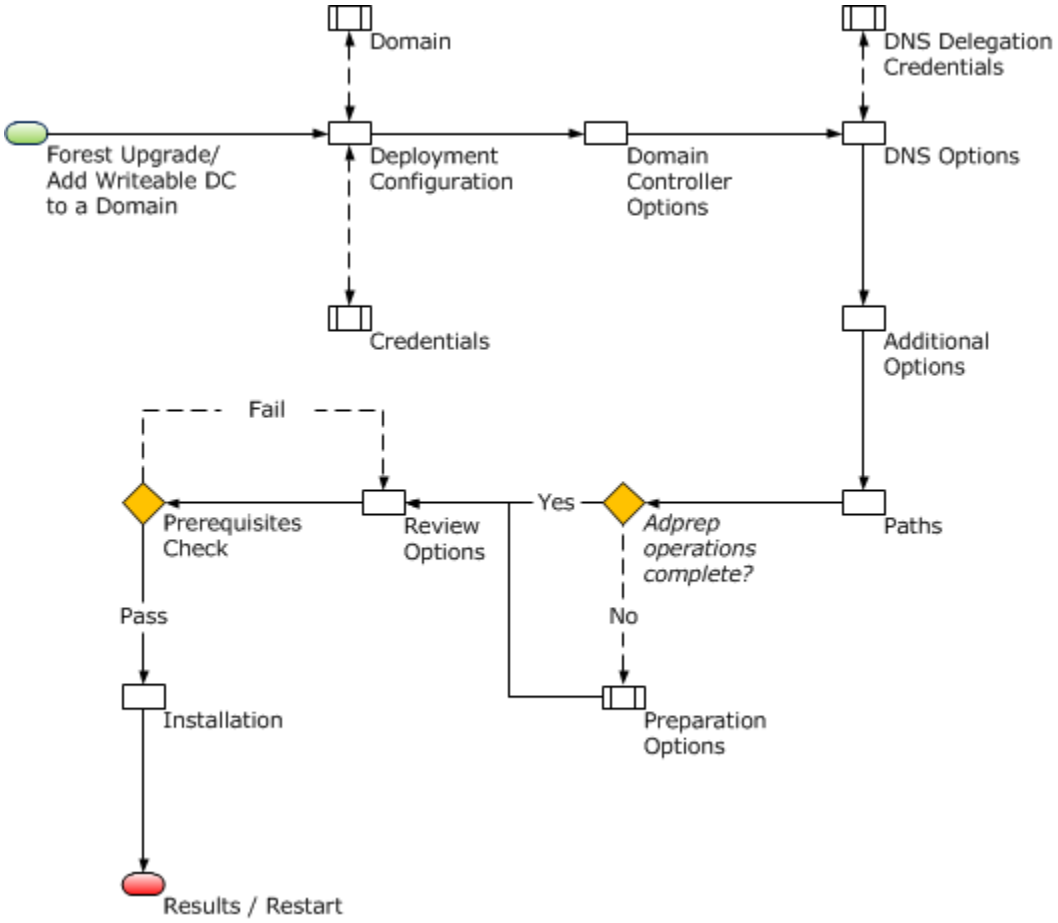

Upgrade and Replica Windows PowerShell

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-AddsDomainController	-SkipPreChecks ***-DomainName*** *-SafeModeAdministratorPassword* *-SiteName* *-ADPrepCredential* -ApplicationPartitionsToReplicate *-AllowDomainControllerReinstall* -Confirm

	-CreateDNSDelegation ***-Credential*** -CriticalReplicationOnly *-DatabasePath* *-DNSDelegationCredential* -Force *-InstallationMediaPath* *-InstallDNS* *-LogPath* -MoveInfrastructureOperationMasterRoleIfNecessary -NoDnsOnNetwork *-NoGlobalCatalog* -Norebootoncompletion *-ReplicationSourceDC* -SkipAutoConfigureDNS -SiteName *-SystemKey* *-SYSVOLPath* *-UseExistingAccount* *-Whatif*

Note

The **-credential** argument is only required if you are not already logged on as a member of the Enterprise Admins and Schema Admins groups (if you are upgrading the forest) or the Domain Admins group (if you are adding a new DC to an existing domain).

Deployment

Deployment Configuration

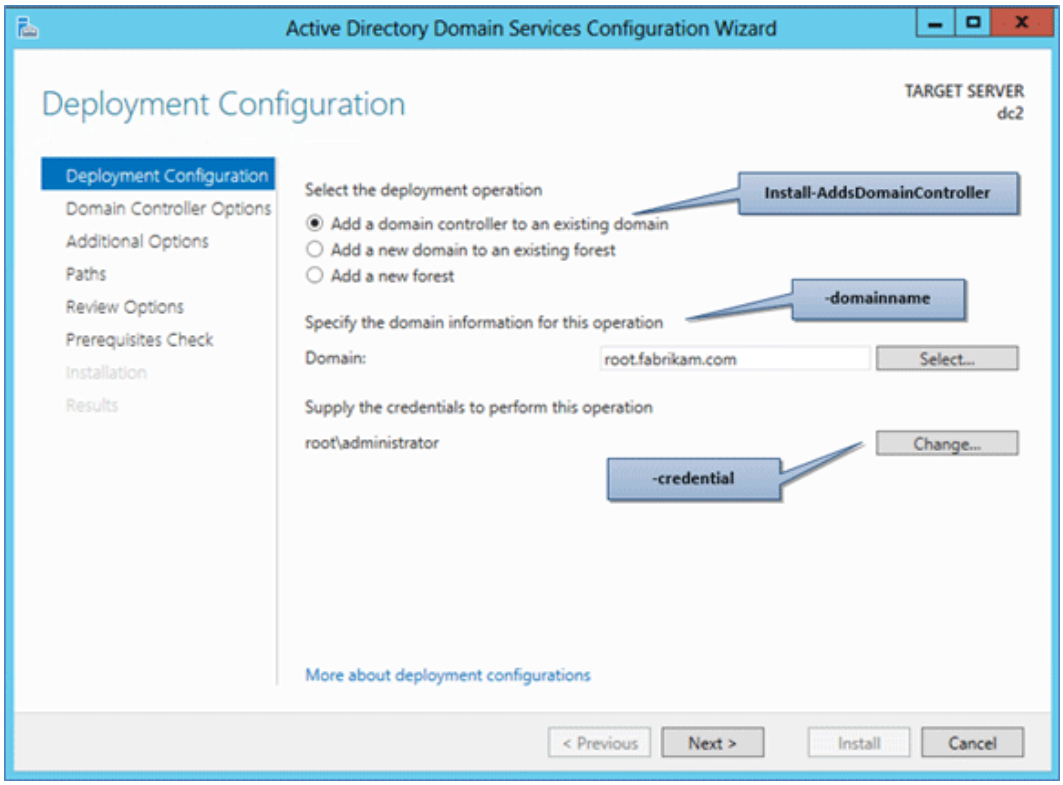

Server Manager begins every domain controller promotion with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select.

To upgrade an existing forest or add a writable domain controller to an existing domain, click **Add a domain controller to an existing domain** and click **Select** to **Specify the domain information for this domain**. Server Manager prompts you for valid credentials if needed.

Upgrading the forest requires credentials that include group memberships in both the Enterprise Admins and Schema Admins groups in Windows Server 2012. The Active Directory Domain Services Configuration Wizard prompts you later if your current credentials do not have adequate permissions or group memberships.

The automatic Adprep process is the only operational difference between adding a domain controller to an existing Windows Server 2012 domain and a domain where domain controllers run an earlier version of Windows Server.

The Deployment Configuration ADDSDeployment cmdlet and arguments are:

```
Install-AddsDomainController

-domainname <string>
```

`-credential <pscredential>`

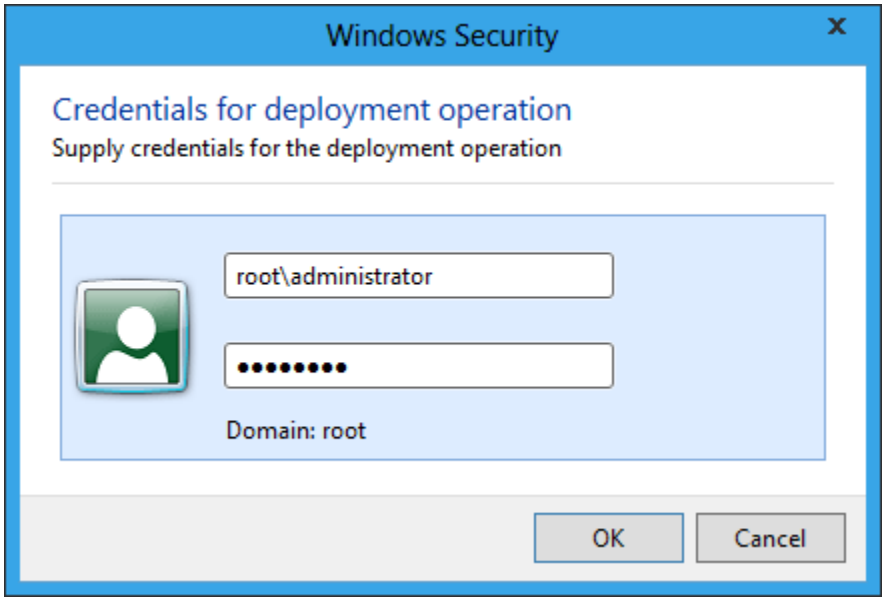

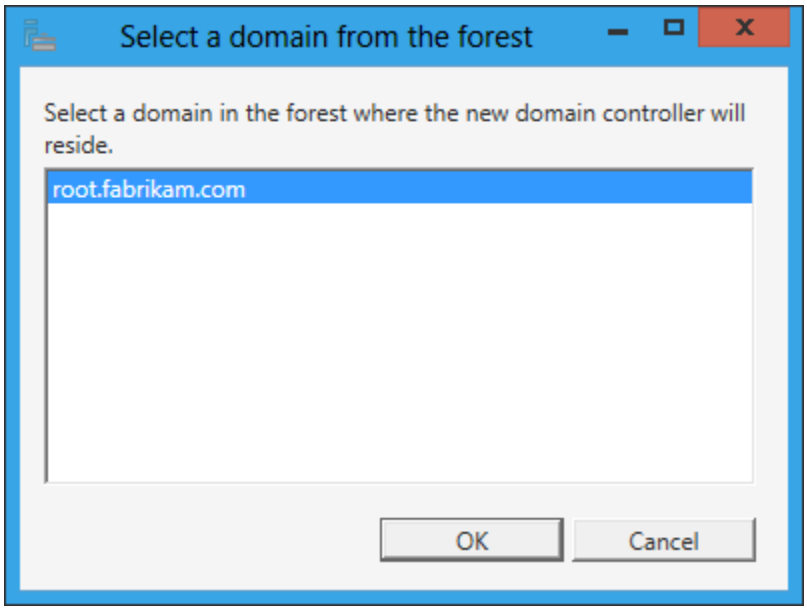

Certain tests perform at each page, some of which repeat later as discrete prerequisite checks. For instance, if the selected domain does not meet the minimal functional levels, you do not have to go all the way through promotion to the prerequisite check to find out:

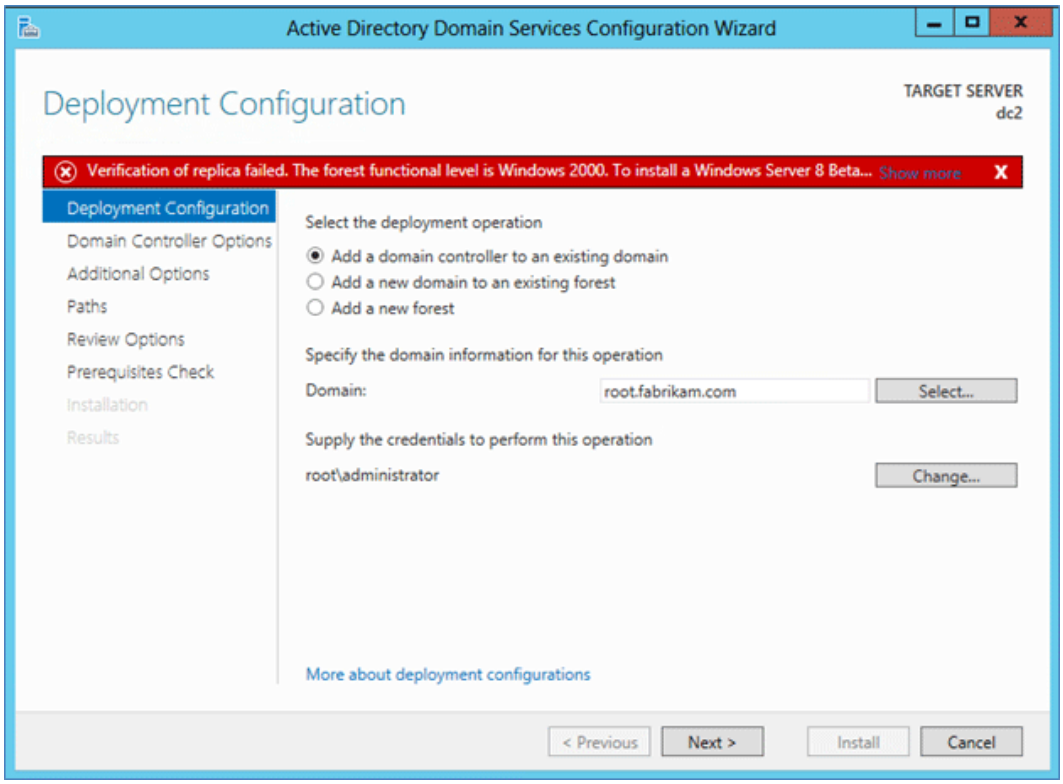
Active Directory Domain Services Configuration Wizard
Deployment Configuration
TARGET SERVER
dc2
Verification of replica failed. The forest functional level is Windows 2000. To install a Windows Server 8 Beta... Show more
Deployment Configuration
Domain Controller Options
Additional Options
Paths
Review Options
Prerequisites Check
Installation
Results
Select the deployment operation
Add a domain controller to an existing domain
Add a new domain to an existing forest
Add a new forest
Specify the domain information for this operation
Domain:
root.fabrikam.com
Select...
Supply the credentials to perform this operation
root\administrator
Change...
More about deployment configurations
< Previous
Next >
Install
Cancel

Domain Controller Options

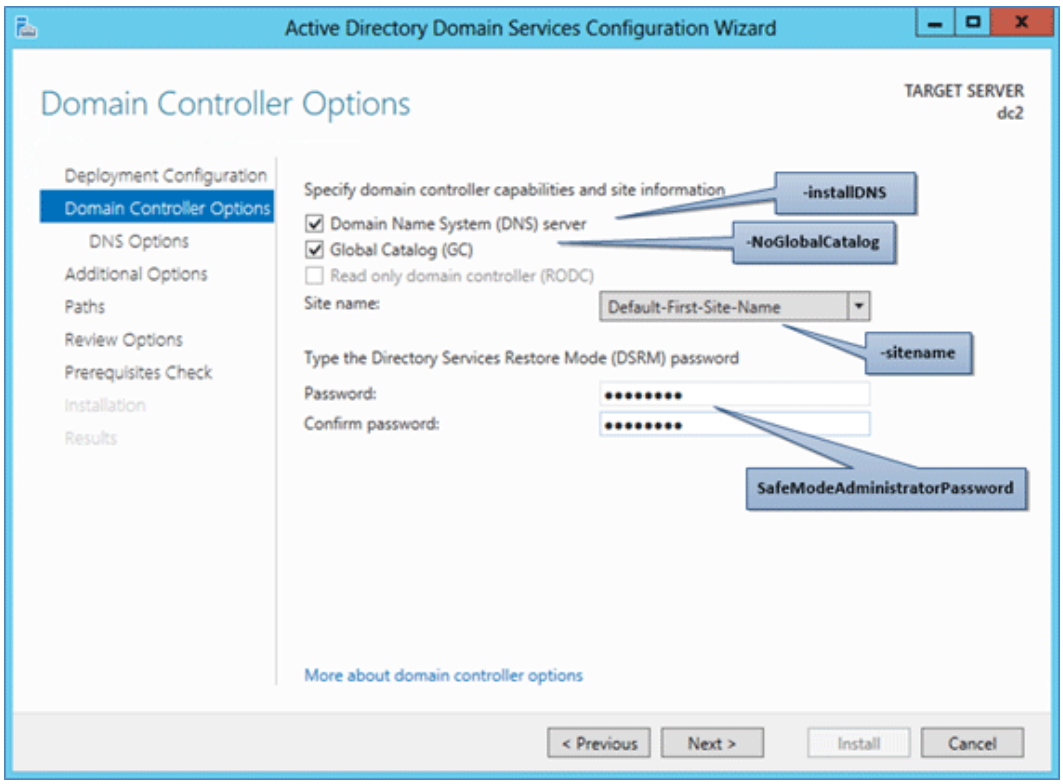

The **Domain Controller Options** page specifies the domain controller capabilities for the new domain controller. The configurable domain controller capabilities are **DNS server**, **Global Catalog**, and **Read-only domain controller**. Microsoft recommends that all domain controllers provide DNS and GC services for high availability in distributed environments. GC is always selected by default and DNS server is selected by default if the current domain hosts DNS already on its DCs based on Start of Authority query. The **Domain Controller Options** page also enables you to choose the appropriate Active Directory logical **site name** from the forest configuration. By default, it selects the site with the most correct subnet. If there is only one site, it selects automatically.

Note

If the server does not belong to an Active Directory subnet and there is more than one Active Directory site, nothing is selected and the **Next** button is unavailable until you choose a site from the list.

The specified **Directory Services Restore Mode Password** must adhere to the password policy applied to the server. Always choose a strong, complex password or preferably, a passphrase.

The **Domain Controller Options** ADDSDeployment arguments are:

```
-InstallDNS <{$false | $true}>
```

```
-NoGlobalCatalog <{$false | $true}>
```

```
-sitename <string>

-SafeModeAdministratorPassword <secure string>
```

Important

The site name must already exist when provided as an argument to **-sitename**. The **install-AddsDomainController** cmdlet does not create sites. You can use cmdlet **new-adreplicationsite** to create new sites.

The **SafeModeAdministratorPassword** argument's operation is special:

If *not specified* as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.

For example, to create an additional domain controller in treyresearch.net domain and be prompted to enter and confirm a masked password:

```
Install-ADDSDomainController -DomainName treyresearch.net -
credential (get-credential)
```

If specified *with a value*, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string:

```
-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged.

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Finally, you could store the obfuscated password in a file, and then reuse it later, without the clear text password ever appearing. For example:

```
$file = "c:\pw.txt"

$pw = read-host -prompt "Password:" -assecurestring

$pw | ConvertFrom-SecureString | Set-Content $file


-safemodeadministratorpassword (Get-Content $File | ConvertTo-SecureString)
```

Warning

Providing or storing a clear or obfuscated text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM password of that domain controller. Anyone with access to the file could reverse that obfuscated password. With that knowledge, they can logon to a DC started in DSRM and eventually impersonate the domain controller itself, elevating their privileges to the highest level in an Active Directory forest. An additional set of steps using **System.Security.Cryptography** to encrypt the text file data is advisable but out of scope. The best practice is to totally avoid password storage.

The ADDSDeployment cmdlet offers an additional option to skip automatic configuration of DNS client settings, forwarders, and root hints. You cannot skip this configuration option when using Server Manager. This argument matters only if you installed the DNS Server role prior to configuring the domain controller:

```
-SkipAutoConfigureDNS
```

The **Domain Controller Options** page warns that you cannot create read only domain controllers if your existing domain controllers run Windows Server 2003. This is expected, and you can dismiss the warning.

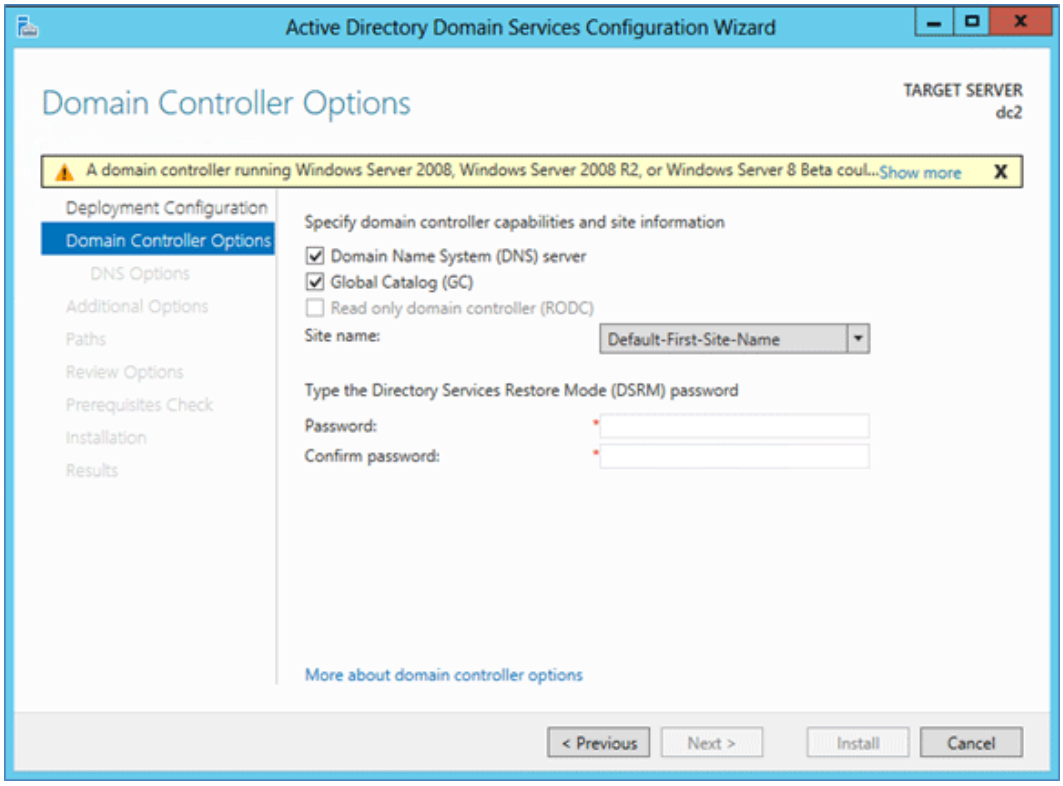

DNS Options and DNS Delegation Credentials

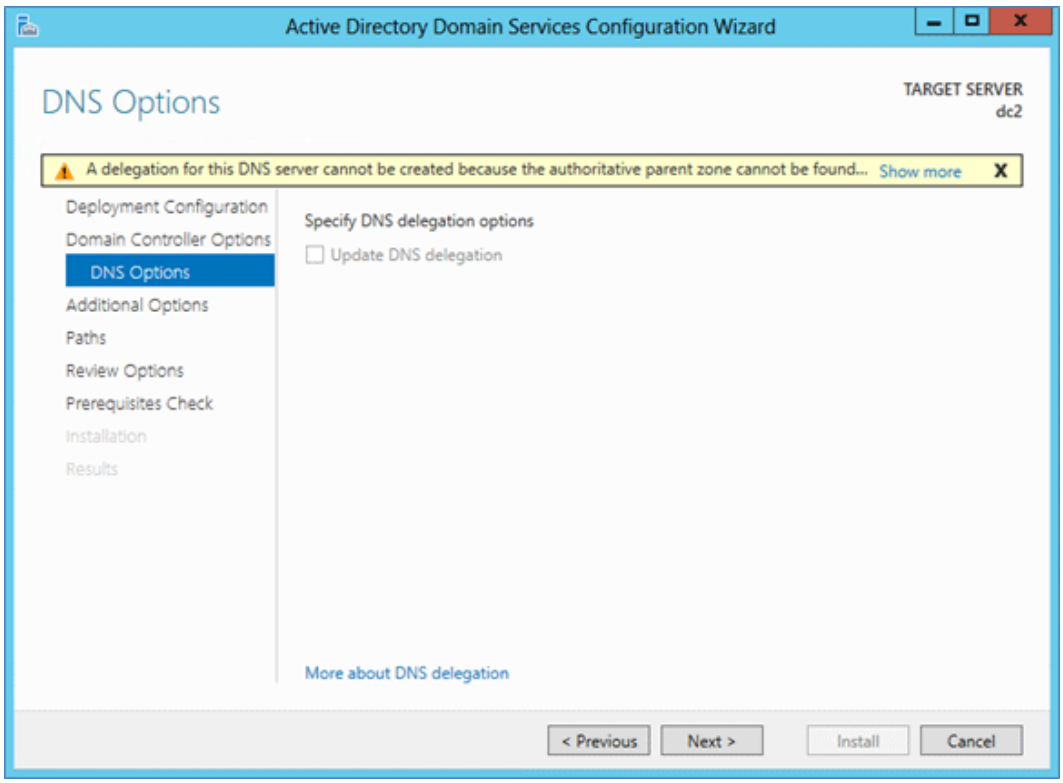

The **DNS Options** page enables you to configure DNS delegation if you selected the **DNS server** option on the *Domain Controller Options* page and if pointing to a zone where DNS delegations are allowed. You may need to provide alternate credentials of a user that is a member of the **DNS Admins** group.

The **DNS Options** ADDSDeployment cmdlet arguments are:

```
-creatednsdelegation
```

```
-dnsdelegationcredential <pscredential>
```

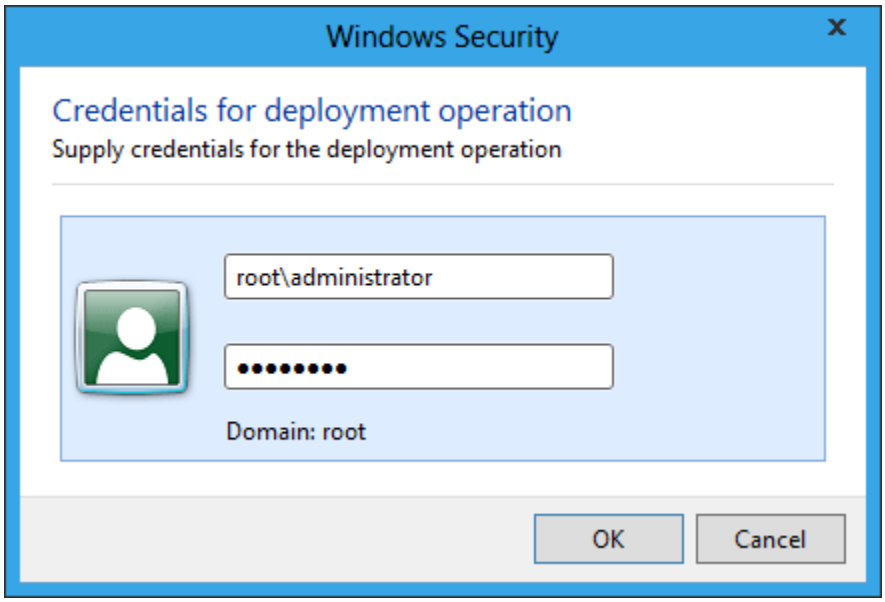

For more information about whether you need to create a DNS delegation, see [Understanding Zone Delegation].

Additional Options

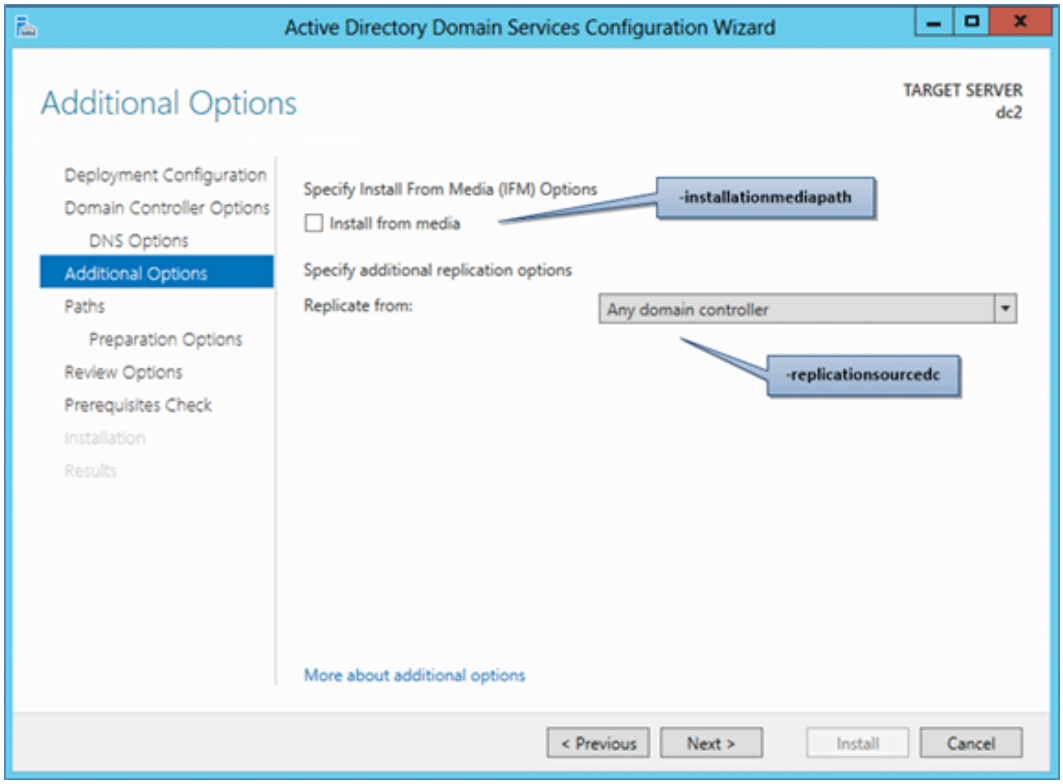

The **Additional Options** page provides the configuration option to name a domain controller as the replication source, or you can use any domain controller as the replication source.

You can also choose to install the domain controller using backed up media using the Install from media (IFM) option. The **Install from media** checkbox provides a browse option once selected and you must click **Verify** to ensure the provided path is valid media. Media used by the IFM option is created with Windows Server Backup or Ntdsutil.exe from another existing Windows Server 2012 computer only; you cannot use a Windows Server 2008 R2 or previous operating system to create media for a Windows Server 2012 domain controller. For more information about changes in IFM, see Simplified Administration Appendix. If using media protected with a SYSKEY, Server Manager prompts for the image's password during verification.

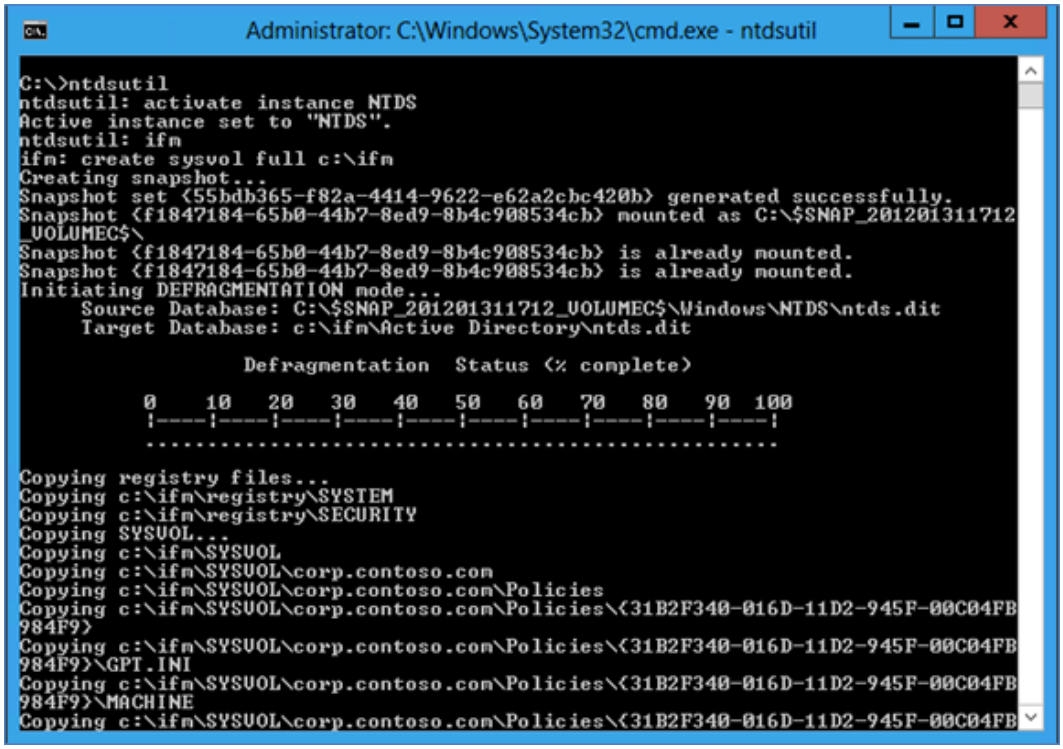

The **Additional Options** ADDSDeployment cmdlet arguments are:

```
-replicationsourcedc <string>

-installationmediapath <string>

-syskey <secure string>
```

Paths

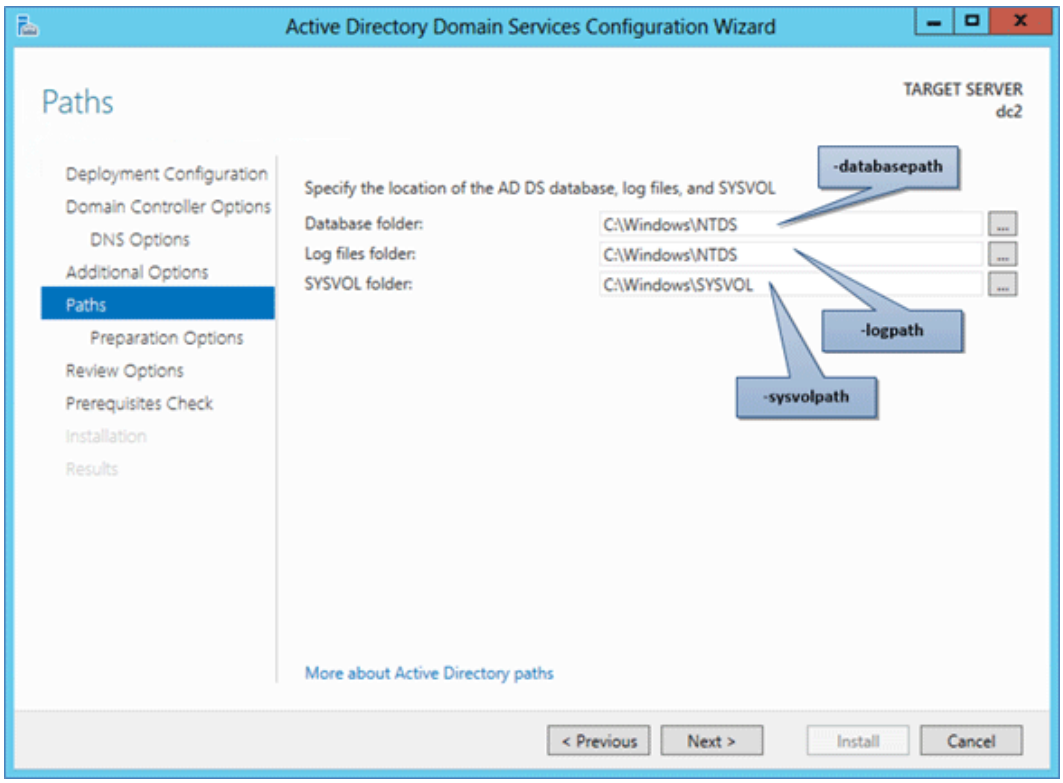

The **Paths** page enables you to override the default folder locations of the AD DS database, the database transaction logs, and the SYSVOL share. The default locations are always in subdirectories of %systemroot%.

The Active Directory Paths ADDSDeployment cmdlet arguments are:

`-databasepath <string>`

`-logpath <string>`

`-sysvolpath <string>`

Preparation Options

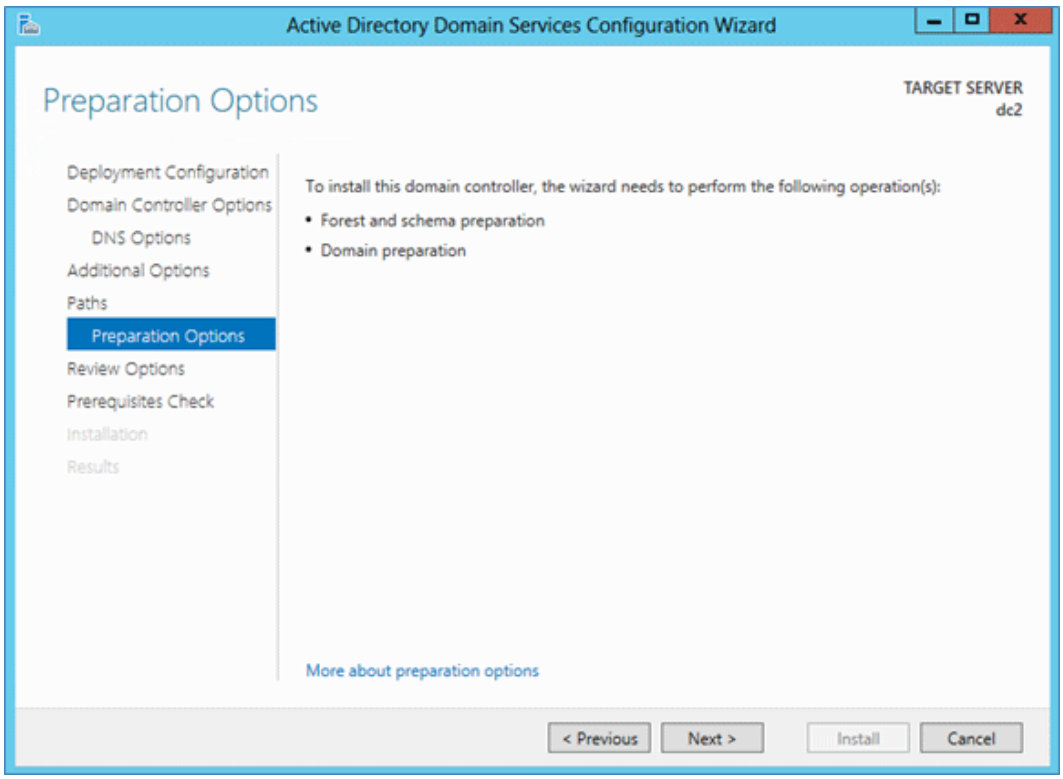

The **Preparation Options** page alerts you that the AD DS configuration includes extending the Schema (forestprep) and updating the domain (domainprep). You only see this page when the forest and domain have not been prepared by previous Windows Server 2012 domain controller installation or from manually running Adprep.exe. For example, the Active Directory Domain Services Configuration Wizard suppresses this page if you add a new domain controller to an existing Windows Server 2012 forest root domain.

Extending the Schema and updating the domain do not occur when you click **Next**. These events occur only during the installation phase. This page simply brings awareness about the events that will occur later in the installation.

This page also validates that the current user credentials are members of the Schema Admin and Enterprise Admins groups, as you need membership in these groups to extend the schema or prepare a domain. Click **Change** to provide the adequate user credentials if the page informs you that the current credentials do not provide sufficient permissions.

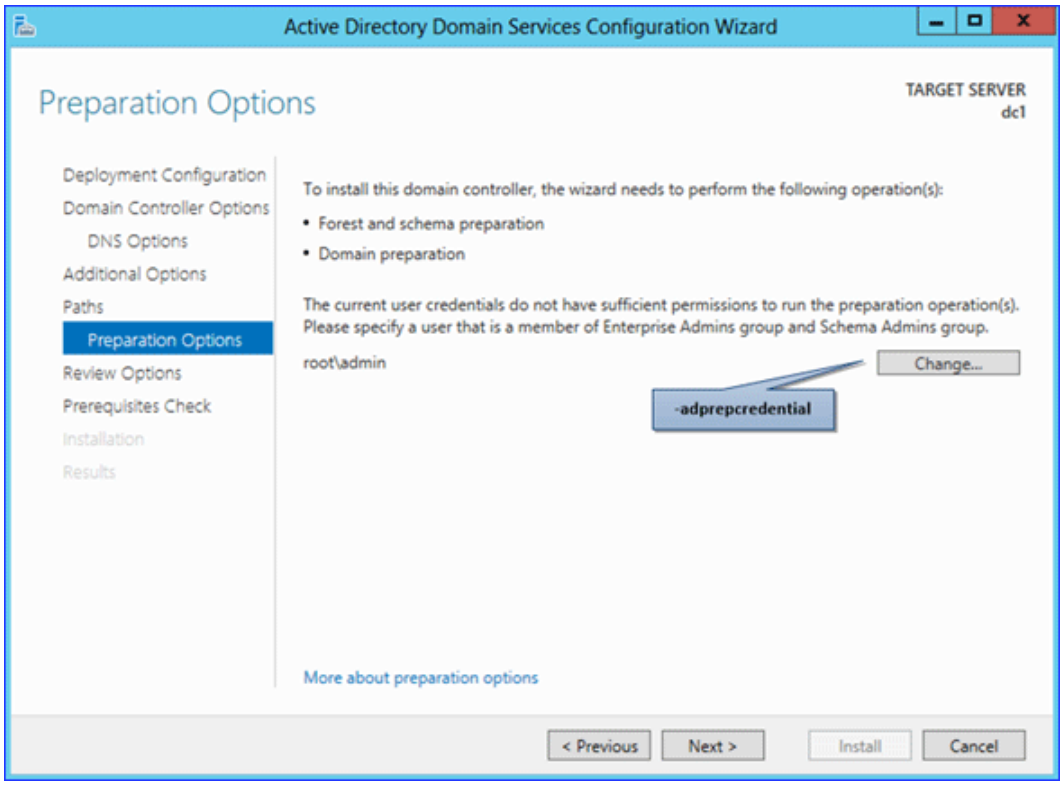

The Additional Options ADDSDeployment cmdlet argument is:

```
-adprepcredential <pscredential>
```

Important

As with previous versions of Windows Server, automated domain preparation for domain controllers that run Windows Server 2012 does not run GPPREP. Run **adprep.exe /gpprep** manually for all domains that were not previously prepared for Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2. You should run GPPrep only once in the history of a domain, not with every upgrade. Adprep.exe does not run /gpprep automatically because its operation can cause all files and folders in the SYSVOL folder to re-replicate on all domain controllers.

Automatic RODCPrep runs when you promote the first un-staged RODC in a domain. It does not occur when you promote the first writeable Windows Server 2012 domain controller. You can also still manually **adprep.exe /rodcprep** if you plan to deploy read-only domain controllers.

Review Options and View Script

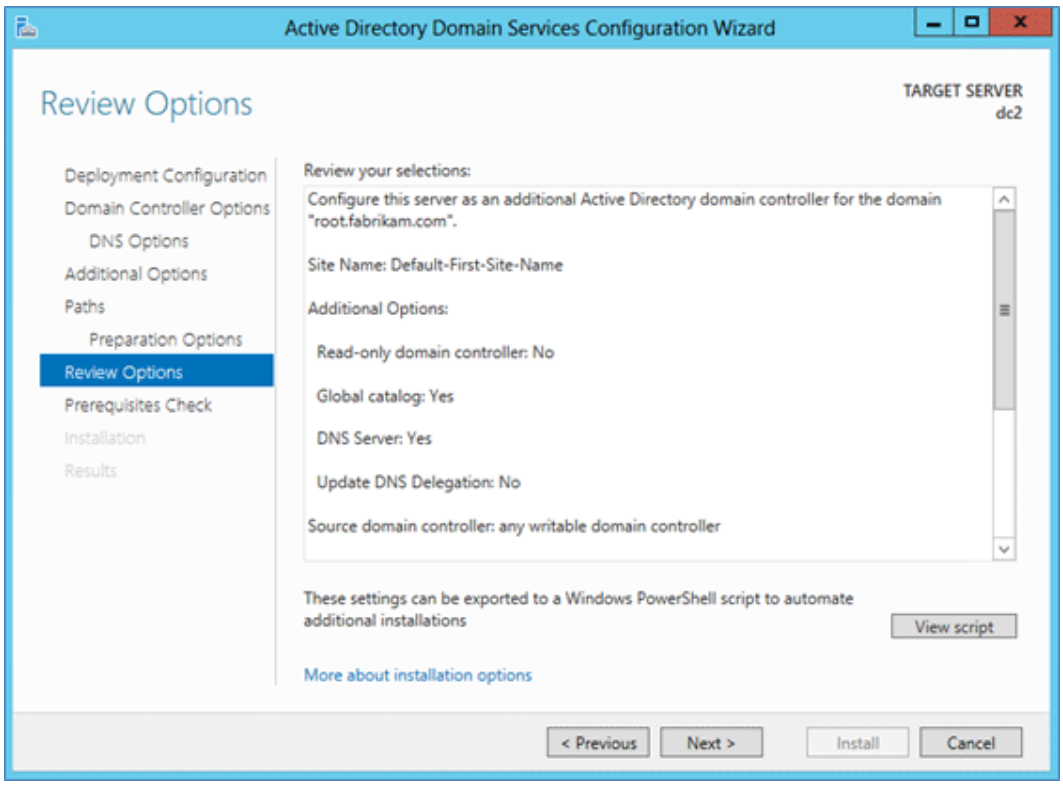

The **Review Options** page enables you to validate your settings and ensure that they meet your requirements before you start the installation. This is not the last opportunity to stop the installation using Server Manager. This page simply enables you to review and confirm your settings before continuing the configuration.

The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use.

For example:

```
#

# Windows PowerShell Script for AD DS Deployment

#

Import-Module ADDSDeployment

Install-ADDSDomainController `

-CreateDNSDelegation `

-Credential (Get-Credential) `
```

```
-CriticalReplicationOnly:$false `
-DatabasePath "C:\Windows\NTDS" `
-DomainName "root.fabrikam.com" `
-InstallDNS:$true `
-LogPath "C:\Windows\NTDS" `
-SiteName "Default-First-Site-Name" `
-SYSVOLPath "C:\Windows\SYSVOL"
-Force:$true
```

Notes

Server Manager generally fills in all arguments with values when promoting and does not rely on defaults (as they may change between future versions of Windows or service packs). The one exception to this is the **-safemodeadministratorpassword** argument. To force a confirmation prompt omit the value when running cmdlet interactively

Use the optional **Whatif** argument with the **Install-ADDSDomainController** cmdlet to review configuration information. This enables you to see the explicit and implicit values of the arguments for a cmdlet.

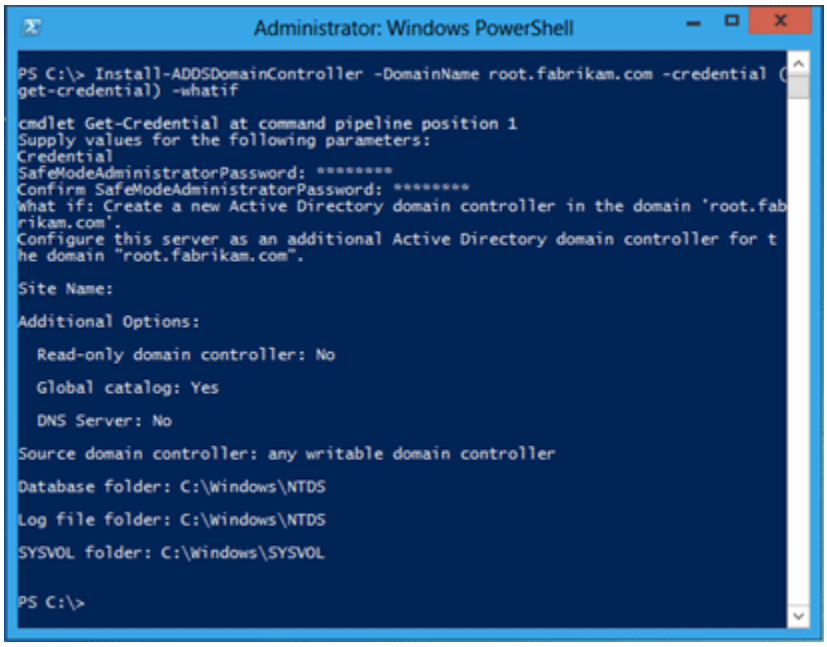

Prerequisites Check

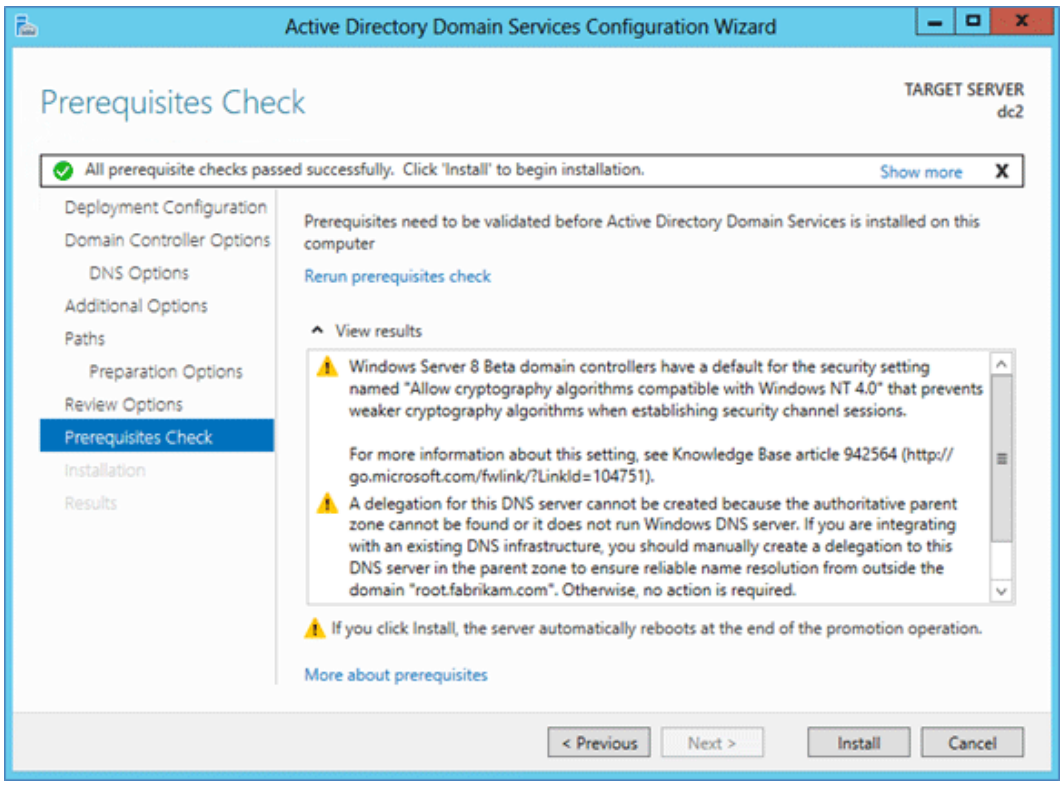

The **Prerequisites Check** is a new feature in AD DS domain configuration. This new phase validates that the domain and forest are capable of supporting a new Windows Server 2012 domain controller.

When installing a new domain controller, the Server Manager Active Directory Domain Services Configuration Wizard invokes a series of serialized modular tests. These tests alert you with suggested repair options. You can run the tests as many times as required. The domain controller process cannot continue until all prerequisite tests pass.

The **Prerequisites Check** also surfaces relevant information such as security changes that affect older operating systems.

For more information about the specific prerequisite checks, see Prerequisite Checking.

You cannot bypass the **Prerequisite Check** when using Server Manager, but you can skip the process when using the AD DS Deployment cmdlet using the following argument:

```
-skipprechecks
```

Warning

Microsoft discourages skipping the prerequisite check as it can lead to a partial domain controller promotion or damaged AD DS forest.

Click **Install** to begin the domain controller promotion process. This is last opportunity to cancel the installation. You cannot cancel the promotion process once it begins. The computer will reboot automatically at the end of promotion, regardless of the promotion results.The **Prerequisites Check** page displays any issues it encountered during the process and guidance for resolving the issue.

Installation

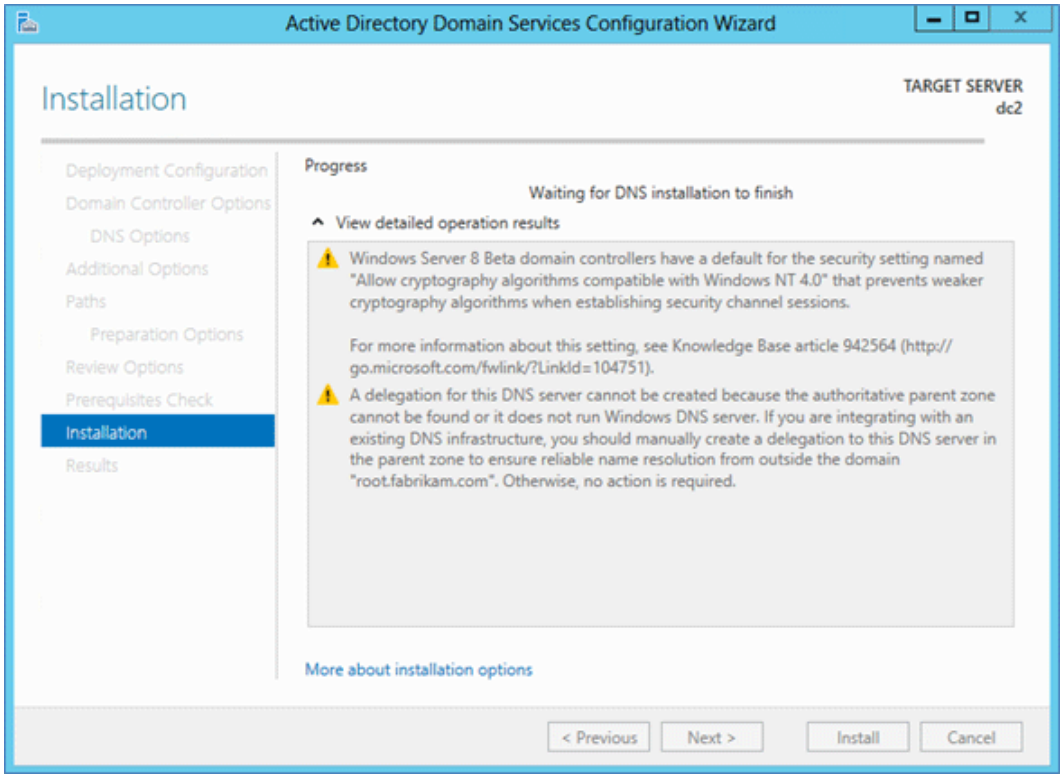

When the **Installation** page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and are written to logs:

- %systemroot%\debug\dcpromo.log
- %systemroot%\debug\dcpromoui.log
- %systemroot%\debug\adprep\logs
- %systemroot%\debug\netsetup.log (if server is in a workgroup)

To install a new Active Directory forest using the ADDSDeployment module, use the following cmdlet:

```
Install-addsdomaincontroller
```

See Upgrade and Replica Windows PowerShell for required and optional arguments.

The **Install-AddsDomainController** cmdlet only has two phases (prerequisite checking and installation). The two figures below show the installation phase with the minimum required arguments of **-domainname** and **-credential**. Note how the Adprep operation happens

automatically as part of adding the first Windows Server 2012 domain controller to an existing Windows Server 2003 forest:

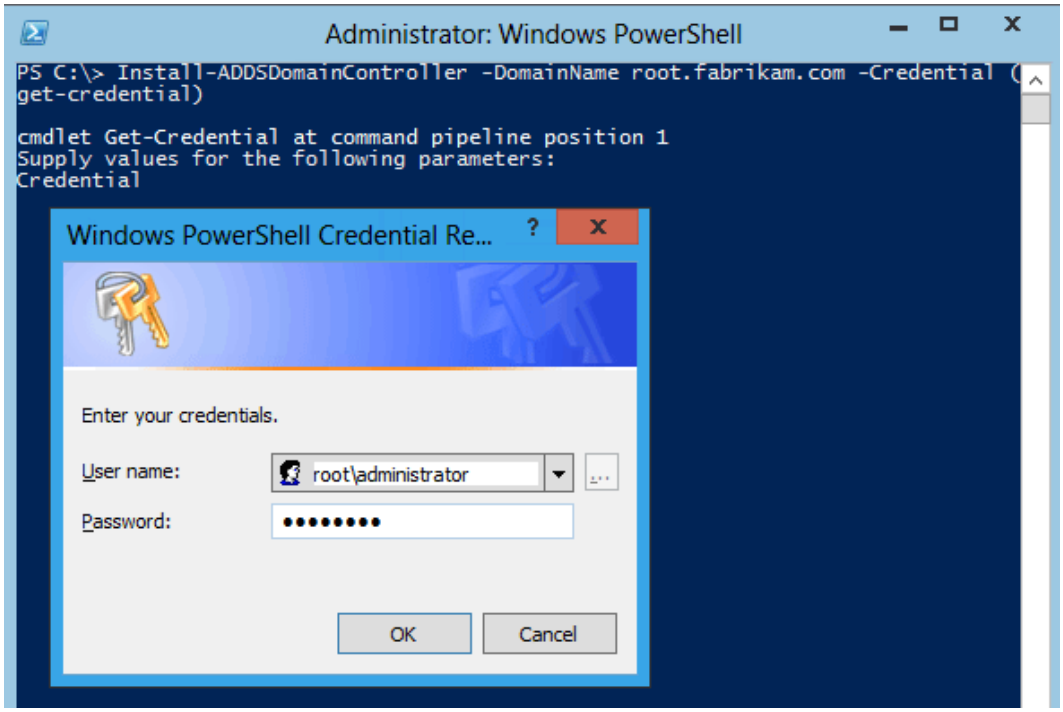

Note how, just like Server Manager, **Install-ADDSDomainController** reminds you that promotion will reboot the server automatically. To accept the reboot prompt automatically, use the **-force** or **-confirm:$false** arguments with any ADDSDeployment Windows PowerShell cmdlet. To prevent the server from automatically rebooting at the end of promotion, use the **-norebootoncompletion** argument.

Warning

Overriding the reboot is discouraged. The domain controller must reboot to function correctly.

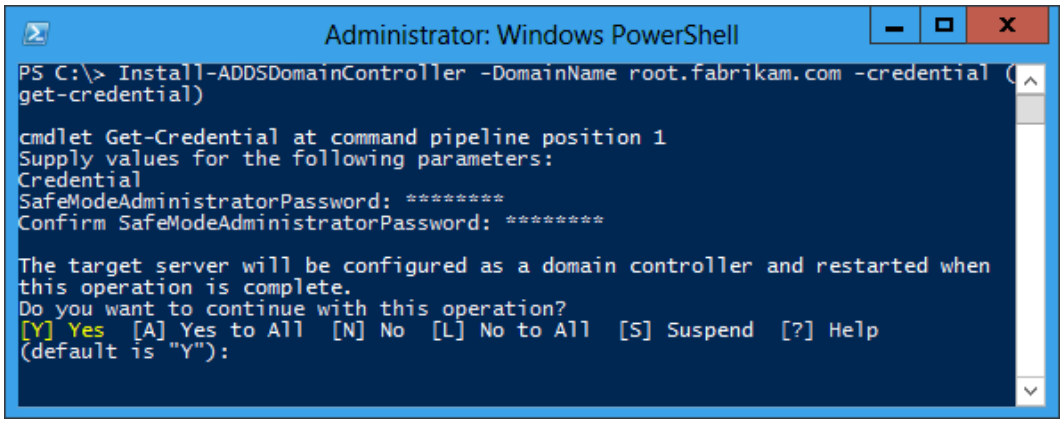

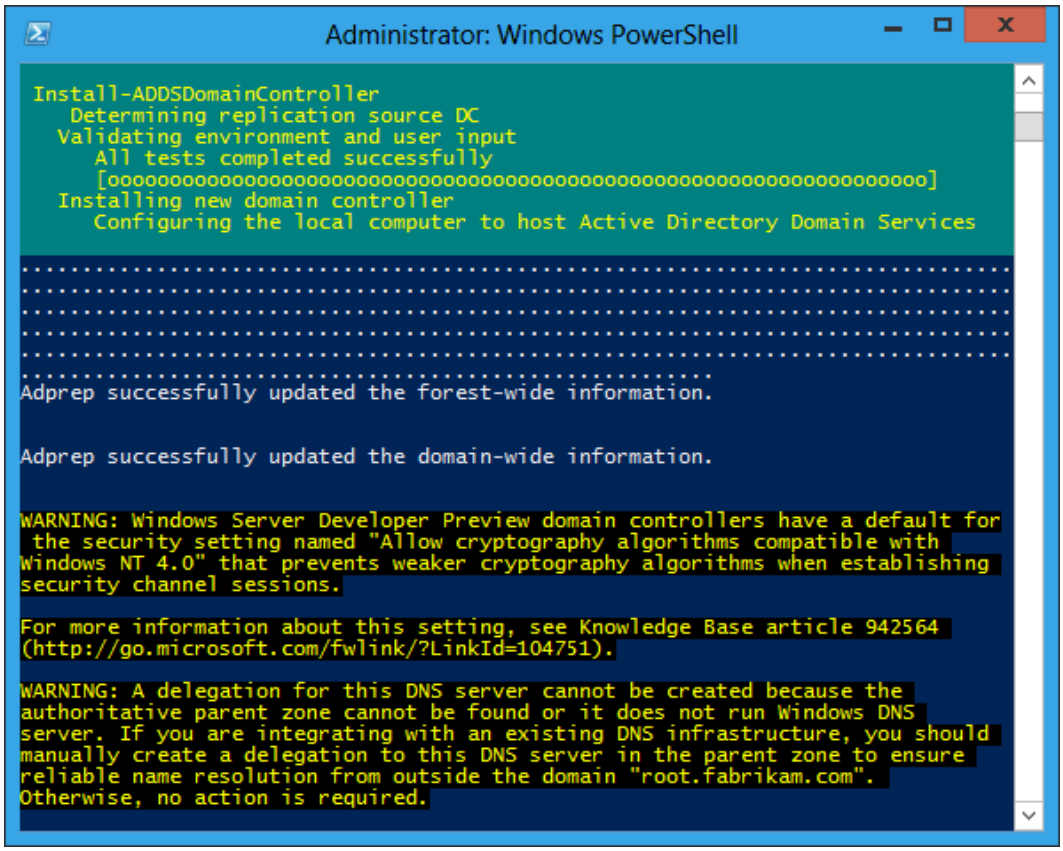
Administrator: Windows PowerShell
Install-ADDSDomainController
Determining replication source DC
Validating environment and user input
All tests completed successfully
[ooooooo

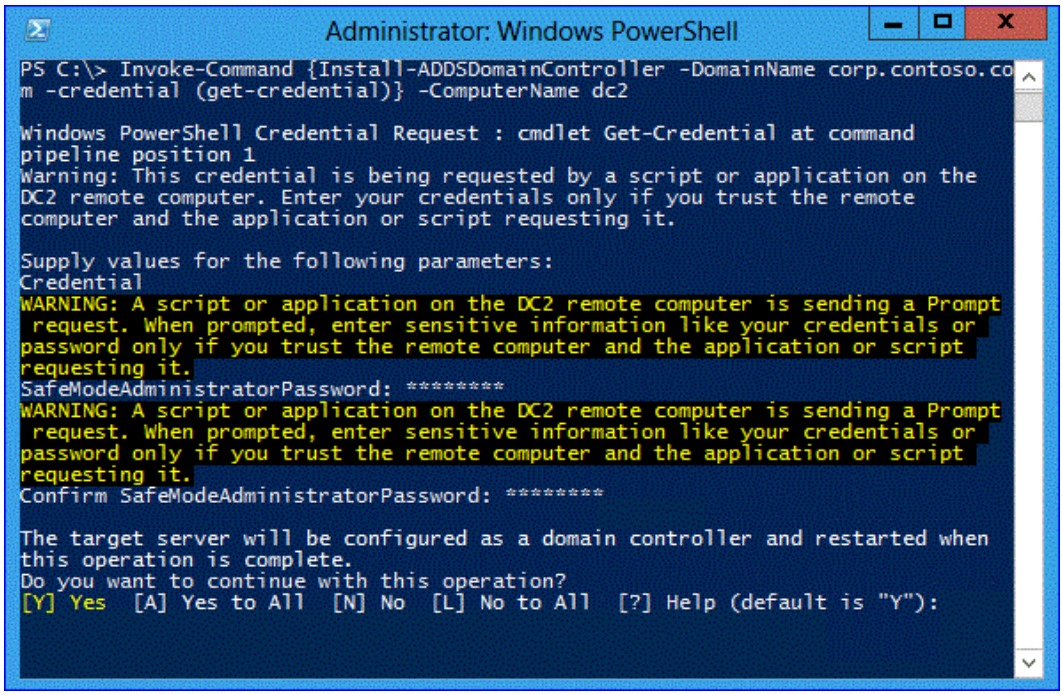

Note

For more information on how the installation and Adprep process works, see the Troubleshooting Domain Controller Deployment.

Results

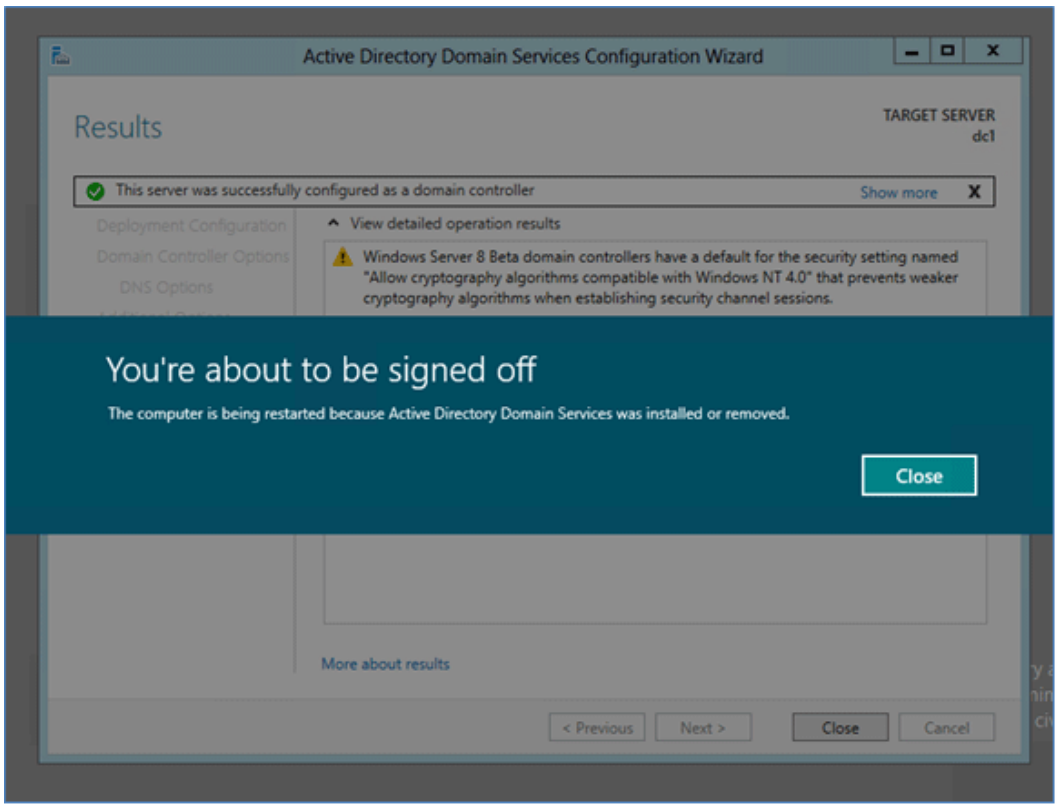

The **Results** page shows the success or failure of the promotion and any important administrative information. If successful, the domain controller will automatically reboot after 10 seconds.

As with previous versions of Windows Server, automated domain preparation for domain controllers that run Windows server 2012 does not run GPPREP. Run **adprep.exe /gpprep** manually for all domains that were not previously prepared for Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2. You should run GPPrep only once in the history of a domain, not with every upgrade. Adprep.exe does not run /gpprep automatically because its operation can cause all files and folders in the SYSVOL folder to re-replicate on all domain controllers.

Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200)

This topic explains how to add child and tree domains to an existing Windows Server 2012 forest, using Server Manager or Windows PowerShell.

- Child and Tree Domain Workflow
- Child and Tree Domain Windows PowerShell

- Deployment

Child and Tree Domain Workflow

The following diagram illustrates the Active Directory Domain Services configuration process when you previously installed the AD DS role and you have started the Active Directory Domain Services Configuration Wizard using Server Manager to create a new domain in an existing forest.

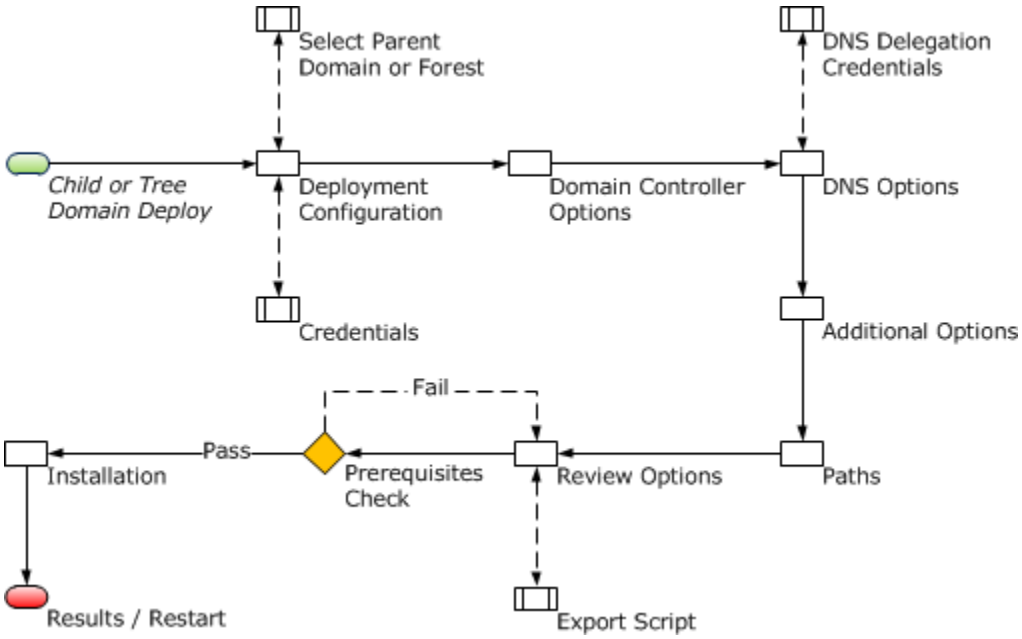

Child and Tree Domain Windows PowerShell

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-AddsDomain	-SkipPreChecks ***-NewDomainName*** ***-ParentDomainName*** ***-SafeModeAdministratorPassword*** *-ADPrepCredential* -AllowDomainReinstall -Confirm

	-CreateDNSDelegation ***-Credential*** *-DatabasePath* *-DNSDelegationCredential* -NoDNSOnNetwork *-DomainMode* ***-DomainType*** -Force *-InstallDNS* *-LogPath* *-NewDomainNetBIOSName* *-NoGlobalCatalog* -NoNorebootoncompletion *-ReplicationSourceDC* *-SiteName* -SkipAutoConfigureDNS *-SYSVOLPath* *-Whatif*

Note
The **-credential** argument is only required when you are not currently logged on as a member of the Enterprise Admins group.The **-NewDomainNetBIOSName** argument is required if you want to change the automatically generated 15-character name based on the DNS domain name prefix or if the name exceeds 15 characters.

Deployment

Deployment Configuration

The following screenshot shows the options for adding a child domain:

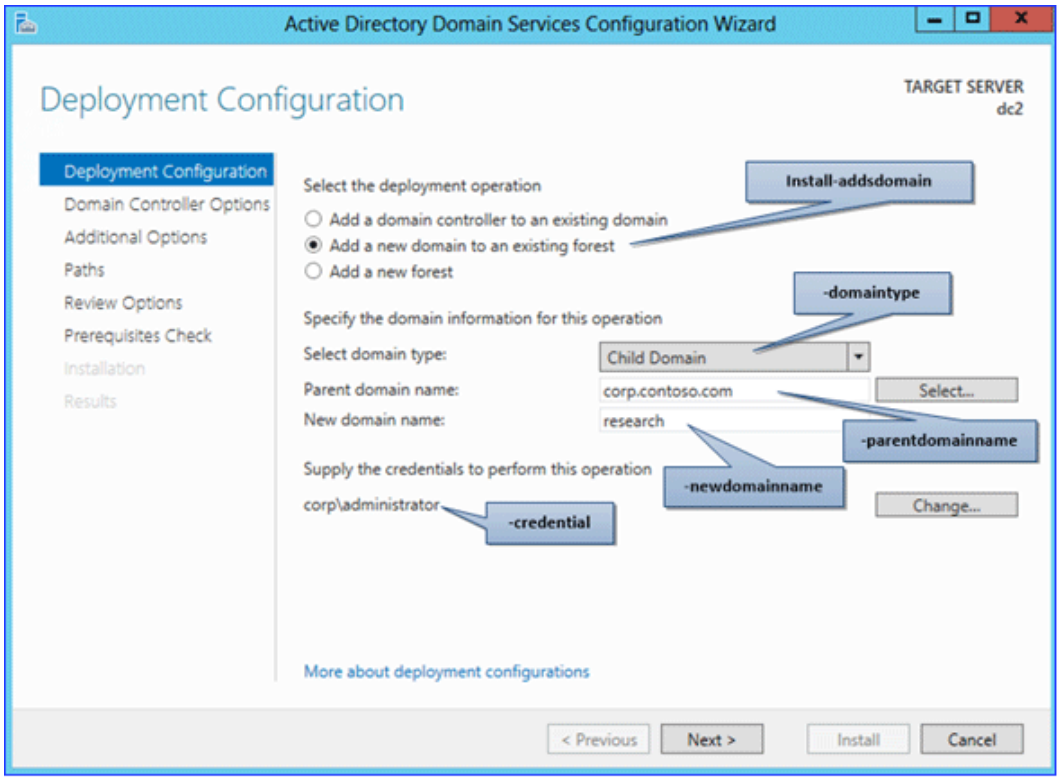

The following screenshot shows the options for adding a tree domain:

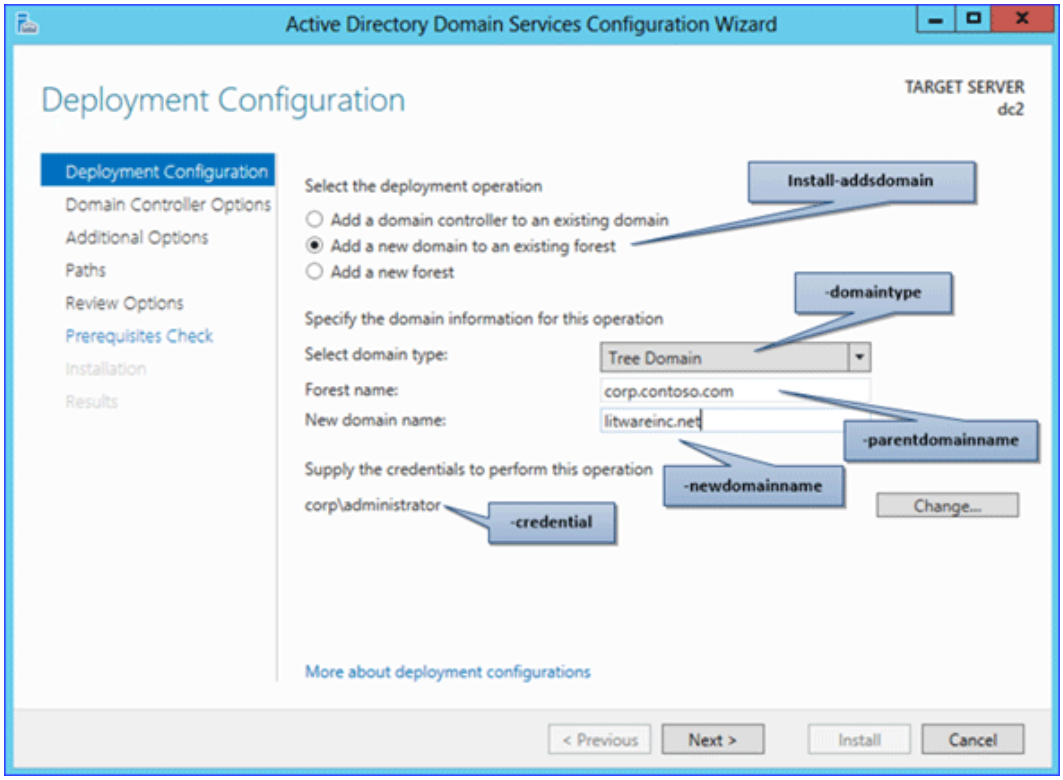

Server Manager begins every domain controller promotion with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select.

This topic combines two discrete operations: child domain promotion and tree domain promotion. The only difference between the two operations is the domain type that you choose to create. All of the other steps are identical between the two operations.

- To create a new child domain, click **Add a domain to an existing Forest** and choose **Child Domain**. For **Parent domain name**, type or select the name of the parent domain. Then type the name of the new domain in the **New domain name** box. Provide a valid, single-label child domain name; the name must use DNS domain name requirements.
- To create a tree domain within an existing forest, click **Add a domain to an existing Forest** and choose **Tree Domain**. Type the name of the forest root domain, and then type the name of the new domain. Provide a valid, fully qualified root domain name; the name cannot be single-labeled and must use DNS domain name requirements.

For more information about DNS names, see Naming conventions in Active Directory for computers, domains, sites, and OUs.

The Server Manager Active Directory Domain Services Configuration Wizard prompts you for domain credentials if your current credentials are not from the domain. Click **Change** to provide domain credentials for the promotion operation.

The Deployment Configuration ADDSDeployment cmdlet and arguments are:

```
Install-AddsDomain

-domaintype <{childdomain | treedomain}>

-parentdomainname <string>

-newdomainname <string>

-credential <pscredential>
```

Domain Controller Options

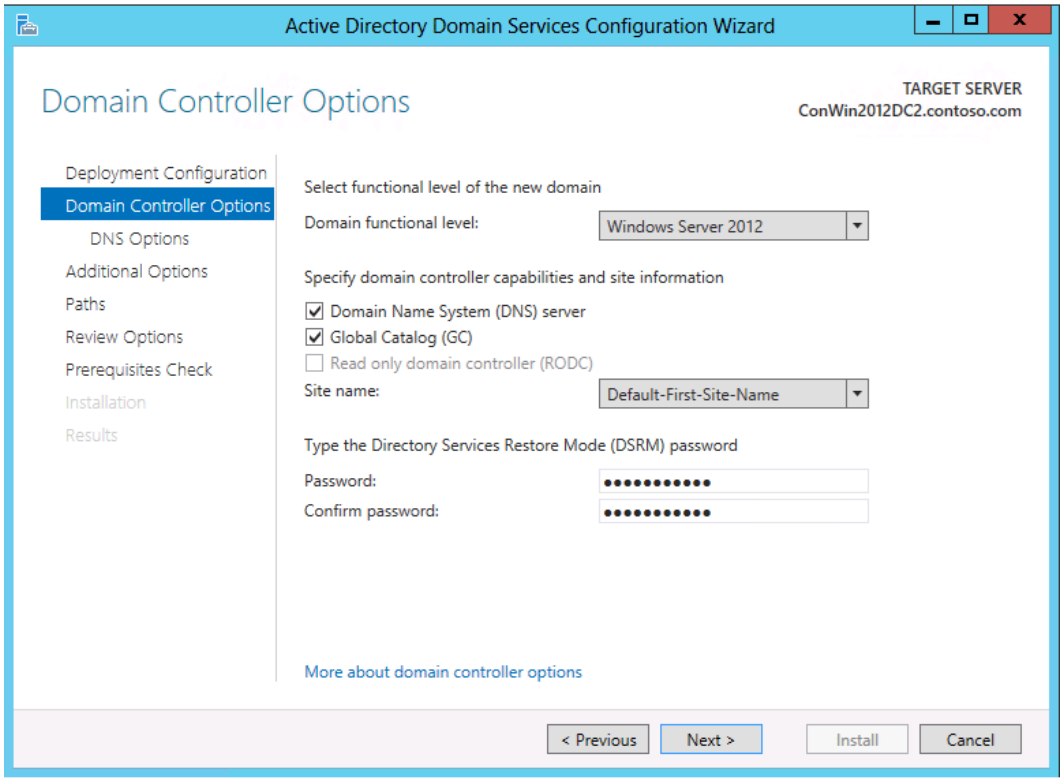

The **Domain Controller Options** page specifies the domain controller options for the new domain controller. The configurable domain controller options include **DNS server** and **Global Catalog**; you cannot configure read-only domain controller as the first domain controller in a new domain.

Microsoft recommends that all domain controllers provide DNS and GC services for high availability in distributed environments. GC is always selected by default and DNS is selected by default if the current domain hosts DNS already on its DCs, based on a Start-of-Authority query. You must also specify a **Domain functional level**. The default functional level is Windows Server 2012, and you can choose any other value that is equal to or greater than the current forest functional level.

The **Domain Controller Options** page also enables you to choose the appropriate Active Directory logical **site name** from the forest configuration. By default, the site with the most correct subnet is selected. If there is only one site, it is selected automatically.

Important

If the server does not belong to an Active Directory subnet and there is more than one Active Directory site, nothing is selected and the **Next** button is unavailable until you choose a site from the list.

The specified **Directory Services Restore Mode Password** must adhere to the password policy applied to the server. Always choose a strong, complex password or preferably, a passphrase.

The **Domain Controller Options** ADDSDeployment cmdlet arguments are:

```
-InstallDNS <{$false | $true}>

-NoGlobalCatalog <{$false | $true}>

-DomainMode <{Win2003 | Win2008 | Win2008R2 | Win2012 | Default}>

-Sitename <string>

-SafeModeAdministratorPassword <secure string>

-Credential <pscredential>
```

Important

The site name must already exist when provided as a value to the **sitename** argument. The **install-AddsDomainController** cmdlet does not create site names. You can use the **new-adreplicationsite** cmdlet to create new sites.

The **Install-ADDSDomainController** cmdlet arguments follow the same defaults as Server Manager if not specified.

The **SafeModeAdministratorPassword** argument's operation is special:

- If *not specified* as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.

  For example, to create a new child domain named NorthAmerica in the Contoso.com forest and be prompted to enter and confirm a masked password:

  ```
  Install-ADDSDomain -NewDomainName NorthAmerica -ParentDomainName
  Contoso.com -DomainType Child
  ```

- If specified *with a value*, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string:

```
-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged.

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Finally, you could store the obfuscated password in a file, and then reuse it later, without the clear text password ever appearing. For example:

```
$file = "c:\pw.txt"
```

```
$pw = read-host -prompt "Password:" -assecurestring
$pw | ConvertFrom-SecureString | Set-Content $file


-safemodeadministratorpassword (Get-Content $File | ConvertTo-SecureString)
```

Warning

Providing or storing a clear or obfuscated text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM password of that domain controller. Anyone with access to the file could reverse that obfuscated password. With that knowledge, they can logon to a DC started in DSRM and eventually impersonate the domain controller itself, elevating their privileges to the highest level in an AD forest. An additional set of steps using **System.Security.Cryptography** to encrypt the text file data is advisable but out of scope. The best practice is to totally avoid password storage.

The ADDSDeployment module offers an additional option to skip automatic configuration of DNS client settings, forwarders, and root hints. This is not configurable when using Server Manager. This argument matters only if you already installed the DNS Server service prior to configuring the domain controller:

```
-SkipAutoConfigureDNS
```

DNS Options and DNS Delegation Credentials

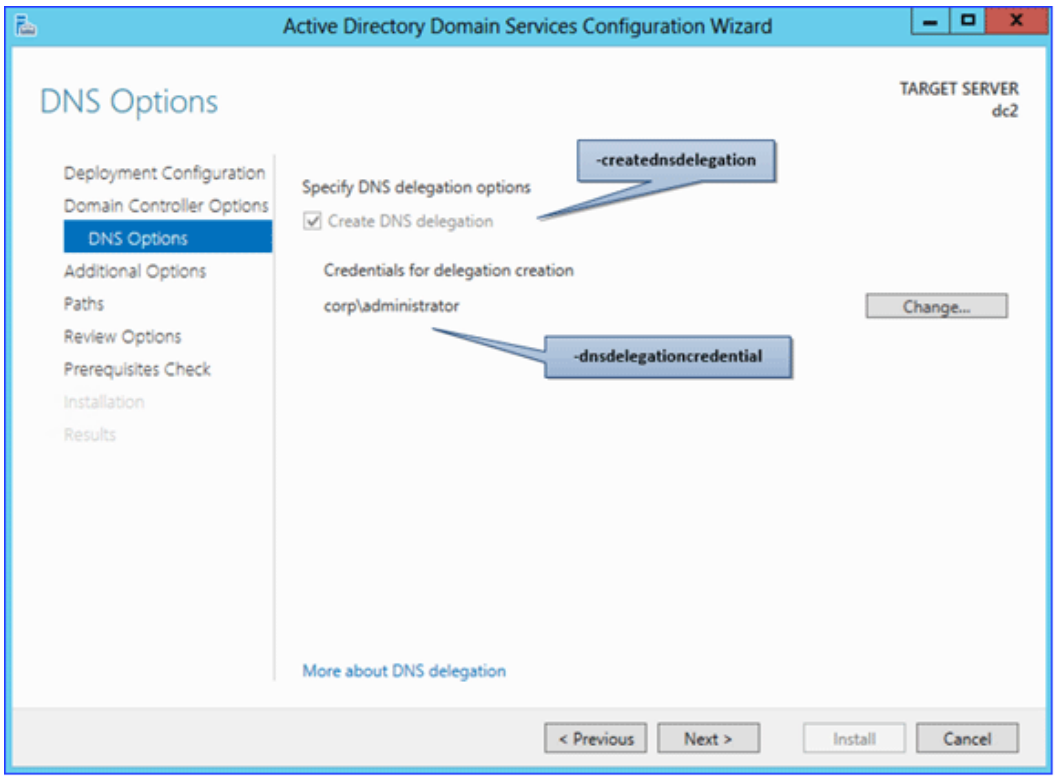

The **DNS Options** page enables you to provide alternate DNS Admin credentials for delegation.

When installing a new domain in an existing forest - where you selected DNS installation on the **Domain Controller Options** page - you cannot configure any options; the delegation happens automatically and irrevocably. You have the option to provide alternate DNS administrative credentials with rights to update that structure.

The **DNS Options** ADDSDeployment Windows PowerShell arguments are:

```
-creatednsdelegation
```

```
-dnsdelegationcredential <pscredential>
```

For more information about DNS delegation, see Understanding Zone Delegation.

Additional Options

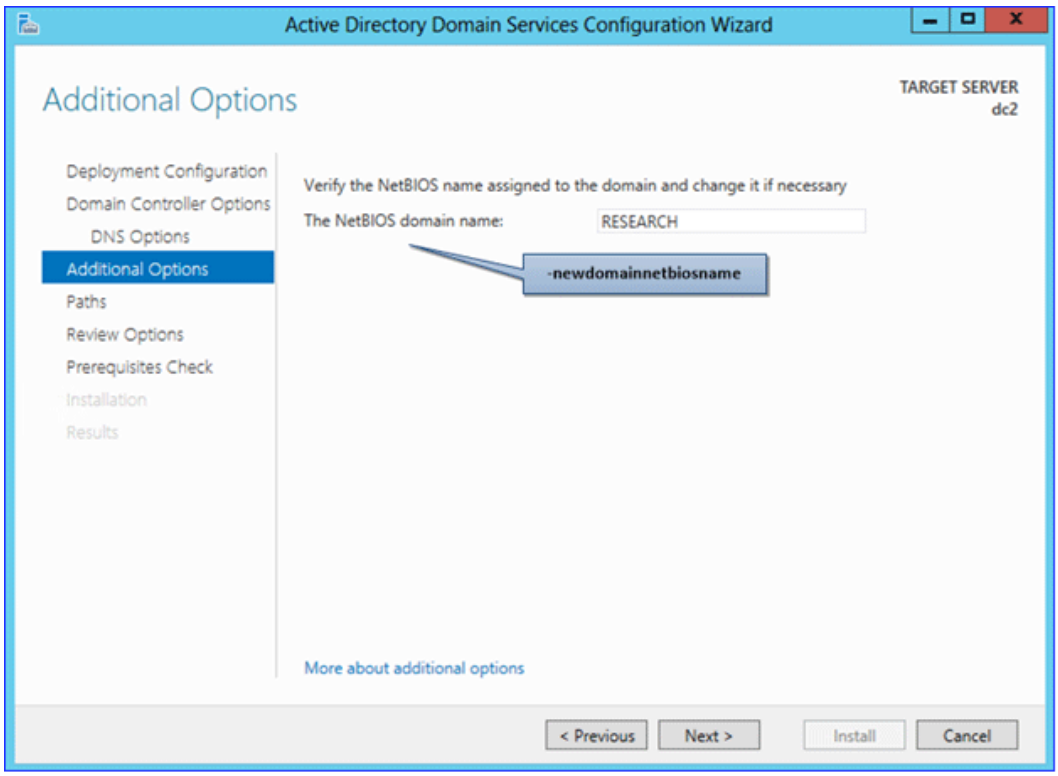

The **Additional Options** page shows the NetBIOS name of the domain and enables you to override it. By default, the NetBIOS domain name matches the left-most label of the fully qualified domain name provided on the **Deployment Configuration** page. For example, if you provided the fully qualified domain name of corp.contoso.com, the default NetBIOS domain name is CORP.

If the name is 15 characters or less and does not conflict with another NetBIOS name, it is unaltered. If it does conflict with another NetBIOS name, a number is appended to the name. If the name is more than 15 characters, the wizard provides a unique, truncated suggestion. In either case, the wizard first validates the name is not already in use via a WINS lookup and NetBIOS broadcast.

For more information about DNS names, see Naming conventions in Active Directory for computers, domains, sites, and OUs.

The **Install-AddsDomain** arguments follow the same defaults as Server Manager if not specified. The **DomainNetBIOSName** operation is special:

1. If the **NewDomainNetBIOSName** argument is not specified with a NetBIOS domain name and the single-label prefix domain name in the **DomainName** argument is 15 characters or fewer, then promotion continues with an automatically generated name.
2. If the **NewDomainNetBIOSName** argument is not specified with a NetBIOS domain name and the single-label prefix domain name in the **DomainName** argument is 16 characters or more, then promotion fails.

3. If the **NewDomainNetBIOSName** argument is specified with a NetBIOS domain name of 15 characters or fewer, then promotion continues with that specified name.
4. If the **NewDomainNetBIOSName** argument is specified with a NetBIOS domain name of 16 characters or more, then promotion fails.

The **Additional Options** ADDSDeployment cmdlet argument is:

```
-newdomainnetbiosname <string>
```

Paths

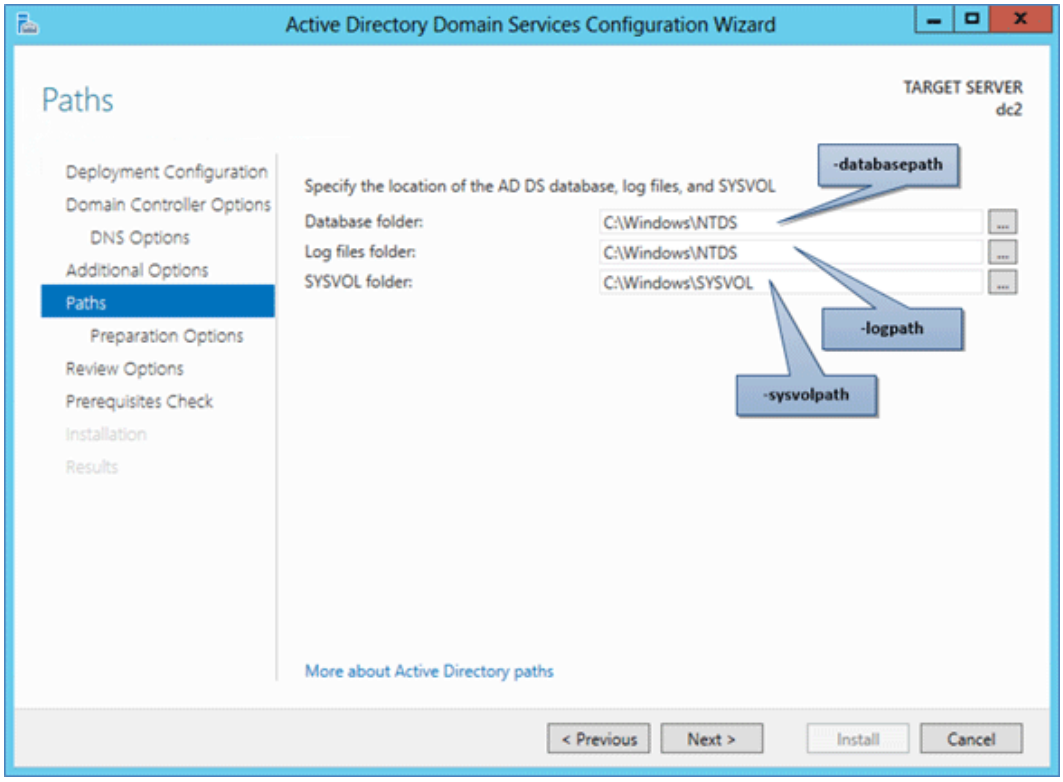

The **Paths** page enables you to override the default folder locations of the AD DS database, the data base transaction logs, and the SYSVOL share. The default locations are always in subdirectories of %systemroot%.

The **Paths** ADDSDeployment cmdlet arguments are:

```
-databasepath <string>
```

```
-logpath <string>
```

```
-sysvolpath <string>
```

Review Options and View Script

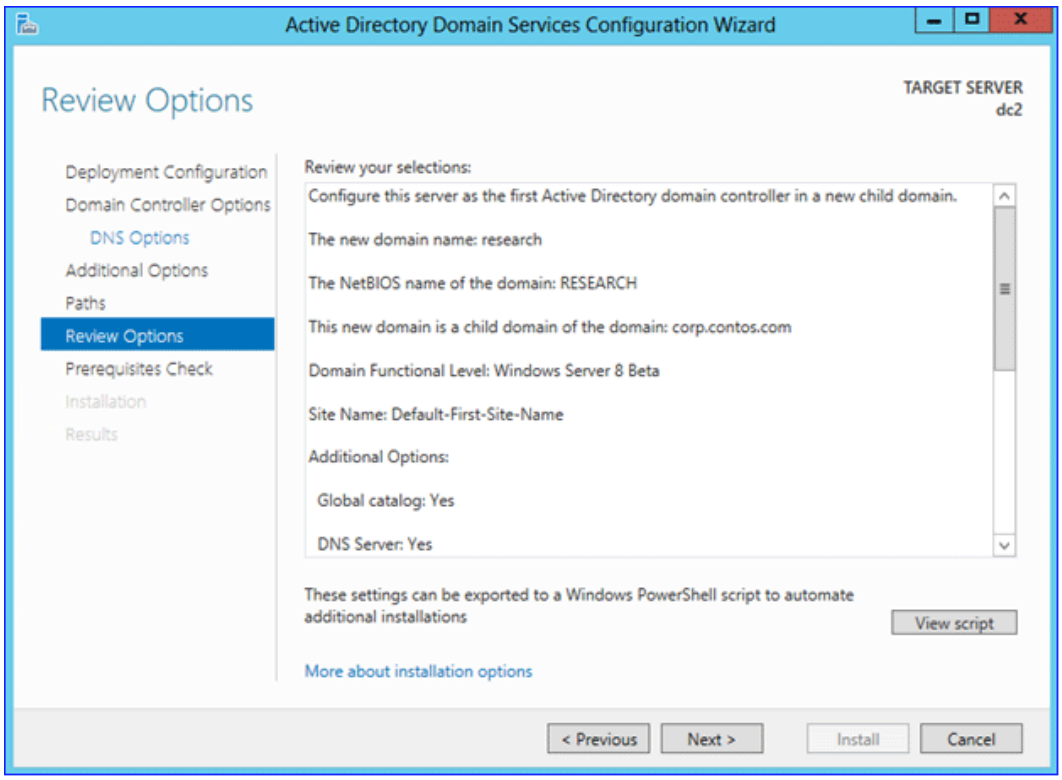

The **Review Options** page enables you to validate your settings and ensure they meet your requirements before you start the installation. This is not the last opportunity to stop the installation when using Server Manager. This is simply an option to confirm your settings before continuing the configuration

The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use. For example:

```
#

# Windows PowerShell Script for AD DS Deployment

#


Import-Module ADDSDeployment

Install-ADDSDomain `

-NoGlobalCatalog:$false `
```

```
-CreateDNSDelegation `
-Credential (Get-Credential) `
-DatabasePath "C:\Windows\NTDS" `
-DomainMode "Win2012" `
-DomainType "ChildDomain" `
-InstallDNS:$true `
-LogPath "C:\Windows\NTDS" `
-NewDomainName "research" `
-NewDomainNetBIOSName "RESEARCH" `
-ParentDomainName "corp.contoso.com" `
-Norebootoncompletion:$false `
-SiteName "Default-First-Site-Name" `
-SYSVOLPath "C:\Windows\SYSVOL"
-Force:$true
```

Note

Server Manager generally fills in all arguments with values when promoting and does not rely on defaults (as they may change between future versions of Windows or service packs). The one exception to this is the **-safemodeadministratorpassword** argument (which is deliberately omitted from the script). To force a confirmation prompt, omit the value when running cmdlet interactively.

Use the optional **Whatif** argument with the **Install-ADDSForest** cmdlet to review configuration information. This enables you to see the explicit and implicit values of the arguments for a cmdlet.

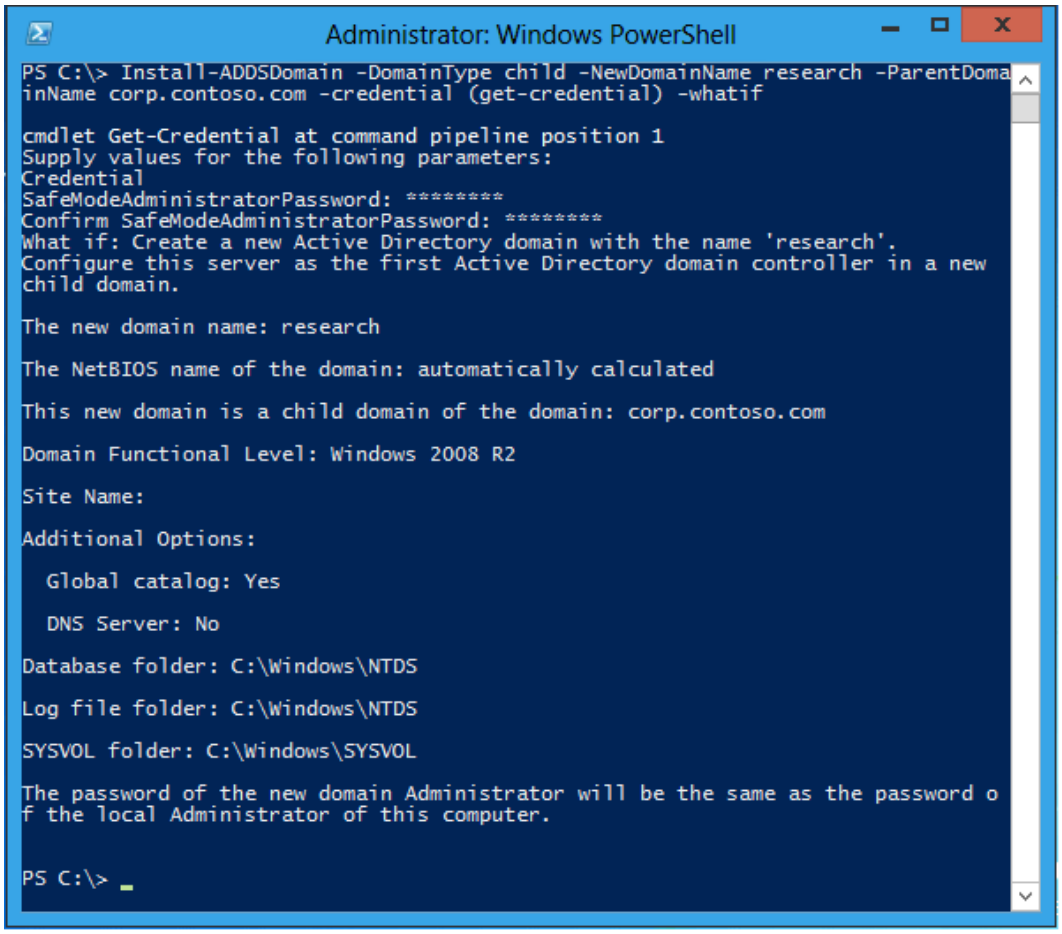
Administrator: Windows PowerShell
PS C:\> Install-ADDSDomain -DomainType child -NewDomainName research -ParentDomainName corp.contoso.com -credential (get-credential) -whatif
cmdlet Get-Credential at command pipeline position 1
Supply values for the following parameters:
Credential
SafeModeAdministratorPassword: ********
Confirm SafeModeAdministratorPassword: ********
What if: Create a new Active Directory domain with the name 'research'.
Configure this server as the first Active Directory domain controller in a new child domain.
The new domain name: research
The NetBIOS name of the domain: automatically calculated
This new domain is a child domain of the domain: corp.contoso.com
Domain Functional Level: Windows 2008 R2
Site Name:
Additional Options:
Global catalog: Yes
DNS Server: No
Database folder: C:\Windows\NTDS
Log file folder: C:\Windows\NTDS
SYSVOL folder: C:\Windows\SYSVOL
The password of the new domain Administrator will be the same as the password of the local Administrator of this computer.
PS C:\>

Prerequisites Check

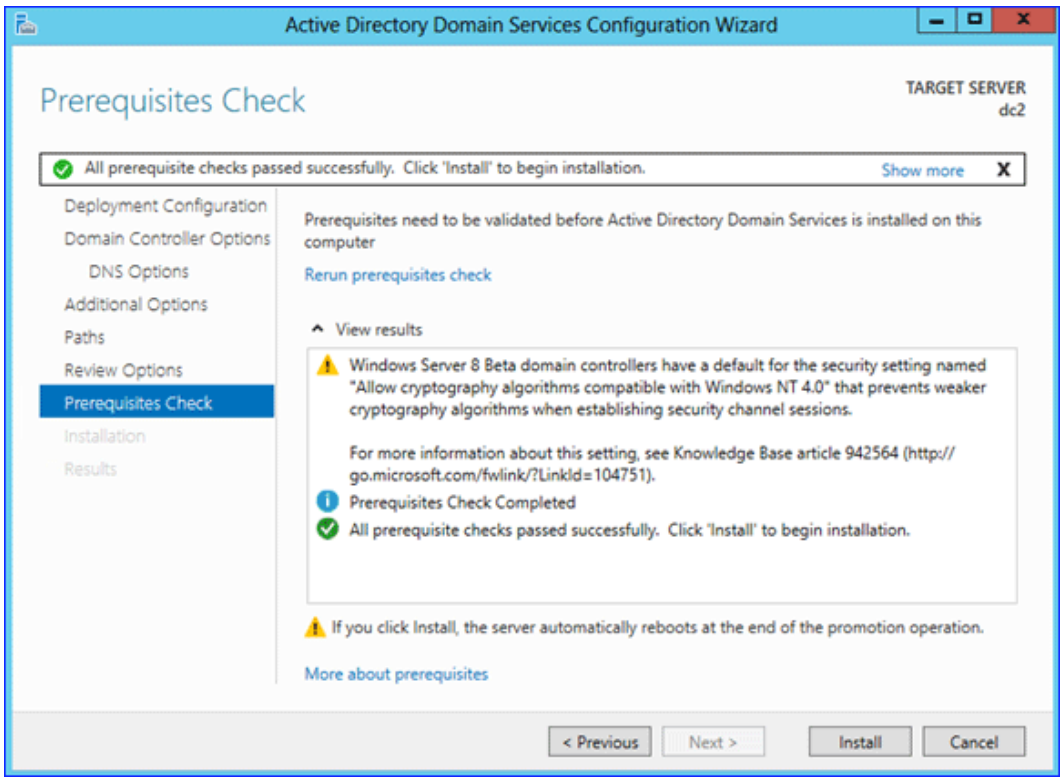

The **Prerequisites Check** is a new feature in AD DS domain configuration. This new phase validates that the server configuration is capable of supporting a new AD DS domain.

When installing a new forest root domain, the Server Manager Active Directory Domain Services Configuration Wizard invokes a series of serialized modular tests. These tests alert you with suggested repair options. You can run the tests as many times as required. The domain controller process cannot continue until all prerequisite tests pass.

The **Prerequisites Check** also surfaces relevant information such as security changes that affect older operating systems.

For more information on the specific prerequisite checks, see Prerequisite Checking.

You cannot bypass the **Prerequisite Check** when using Server Manager, but you can skip the process when using the AD DS Deployment cmdlet using the following argument:

```
-skipprechecks
```

> **Warning**
> Microsoft discourages skipping the prerequisite check as it can lead to a partial domain controller promotion or damaged AD DS forest.

Click **Install** to begin the domain controller promotion process. This is last opportunity to cancel the installation. You cannot cancel the promotion process once it begins. The computer will reboot automatically at the end of promotion, regardless of the promotion results.

Installation

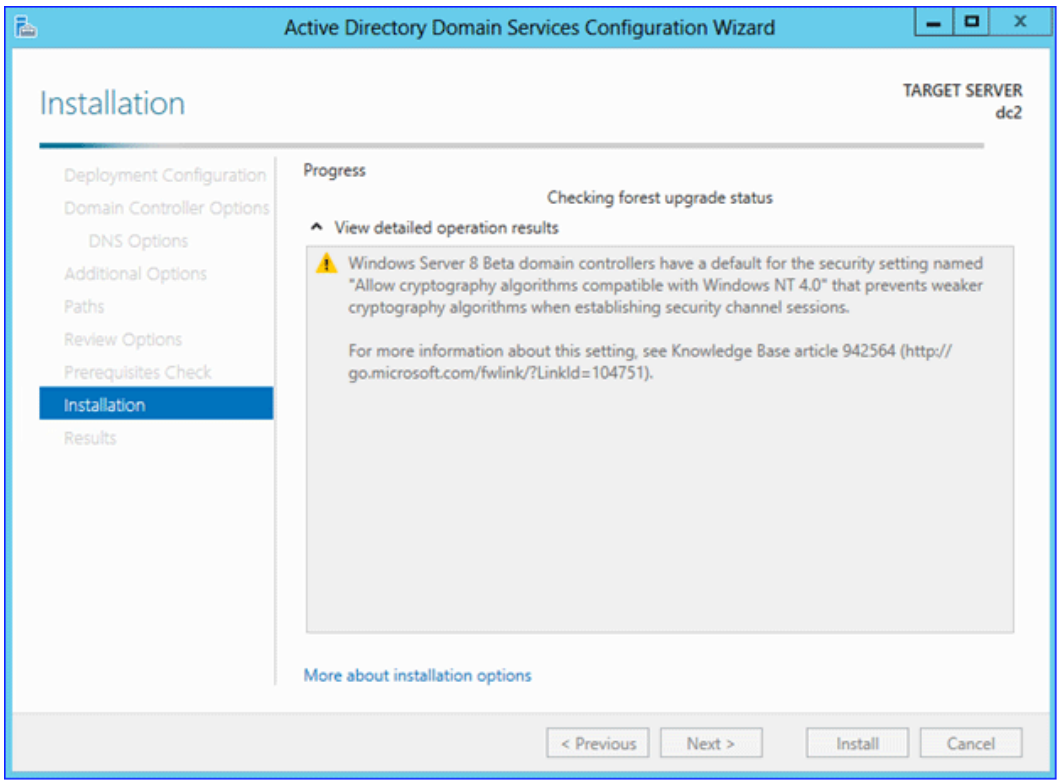

When the **Installation** page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and are written to logs:

- %systemroot%\debug\dcpromo.log
- %systemroot%\debug\dcpromoui.log

To install a new Active Directory domain using the ADDSDeployment module, use the following cmdlet:

```
Install-addsdomain
```

See Child and Tree Domain Windows PowerShell for required and optional arguments.The **Install-addsdomain** cmdlet only has two phases (prerequisite checking and installation). The two figures below show the installation phase with the minimum required arguments of **-domaintype**, **-newdomainname**, **-parentdomainname**, and **-credential**. Note how, just like Server Manager, **Install-ADDSDomain** reminds you that promotion will reboot the server automatically.

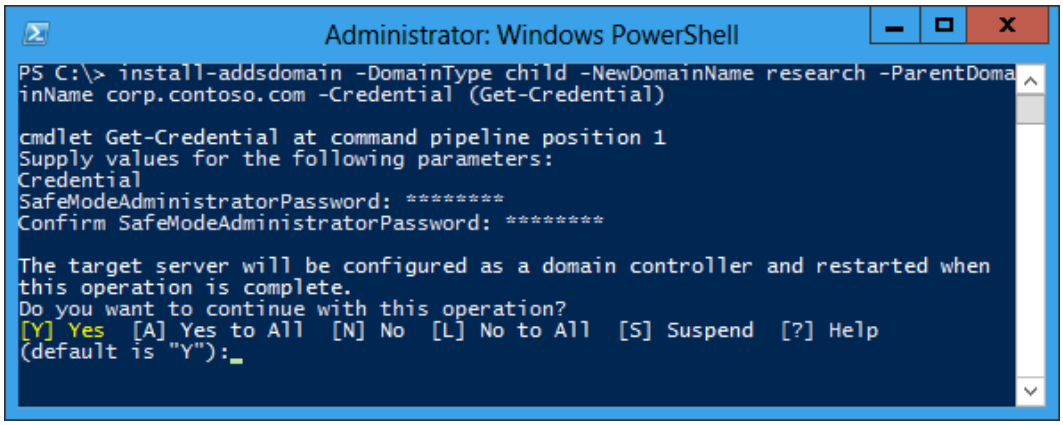

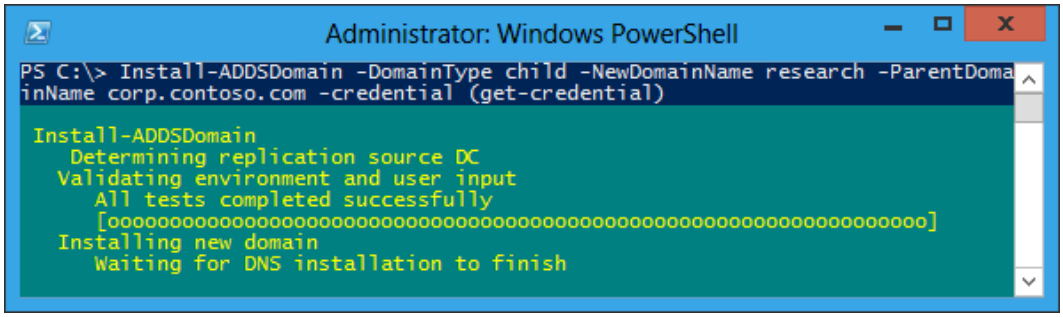
Administrator: Windows PowerShell
PS C:\> Install-ADDSDomain -DomainType child -NewDomainName research -ParentDomainName corp.contoso.com -credential (get-credential)
Install-ADDSDomain
Determining replication source DC
Validating environment and user input
All tests completed successfully
[ooooooo

Results

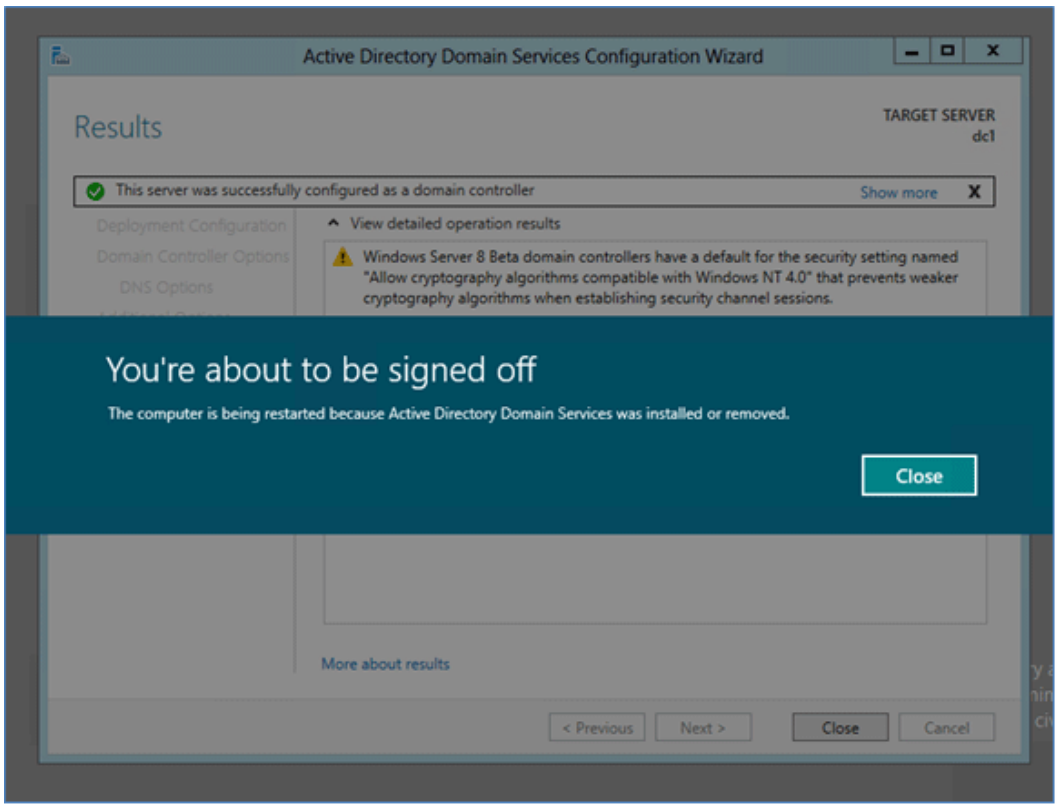

The **Results** page shows the success or failure of the promotion and any important administrative information. The domain controller will automatically reboot after 10 seconds.

Install a Windows Server 2012 Active Directory Read-Only Domain Controller (RODC) (Level 200)

This topic explains how to create a staged RODC account and then attach a server to that account during RODC installation. This topic also explains how to install an RODC without performing a staged installation.

Stage RODC Workflow

A staged read only domain controller (RODC) installation works in two discrete phases:

1. Staging an unoccupied computer account
2. Attaching an RODC to that account during promotion

The following diagram illustrates the Active Directory Domain Services Read-Only Domain Controller staging process, where you create an empty RODC computer account in the domain using the Active Directory Administrative Center (Dsac.exe).

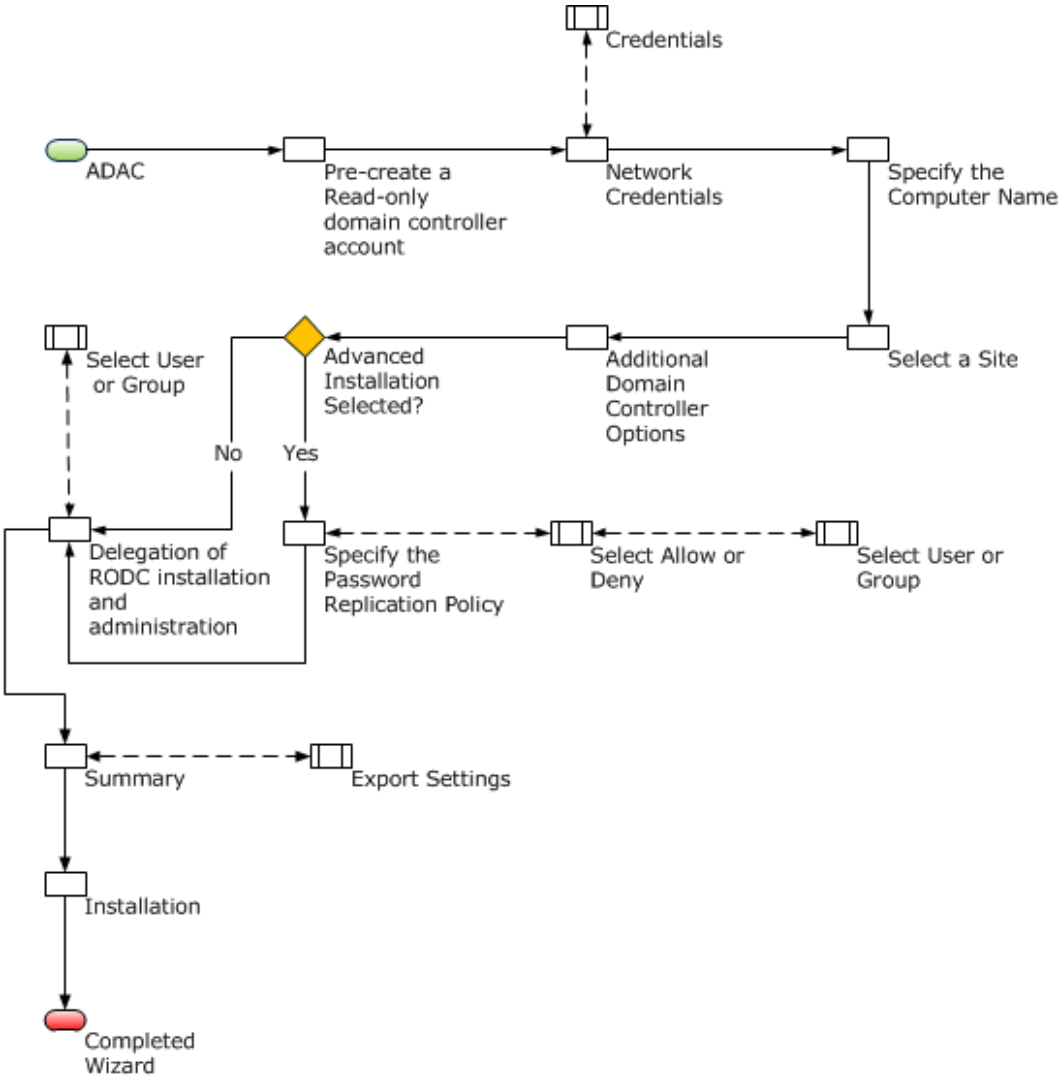

Stage RODC Windows PowerShell

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Add-addsreadonlydomaincontrolleraccount	-SkipPreChecks ***-DomainControllerAccountName***

	-DomainName ***-SiteName*** *-AllowPasswordReplicationAccountName* ***-Credential*** *-DelegatedAdministratorAccountName* *-DenyPasswordReplicationAccountName* *-NoGlobalCatalog* *-InstallDNS* -ReplicationSourceDC

Note

The **-credential** argument is only required if you are not already logged on as a member of the Domain Admins group.

Attach RODC Workflow

The diagram below illustrates the Active Directory Domain Services configuration process, where you already installed the AD DS role, you staged the RODC account, and started **Promote this Server to a Domain Controller** using Server Manager to create a new RODC in an existing domain, attaching it to the staged computer account.

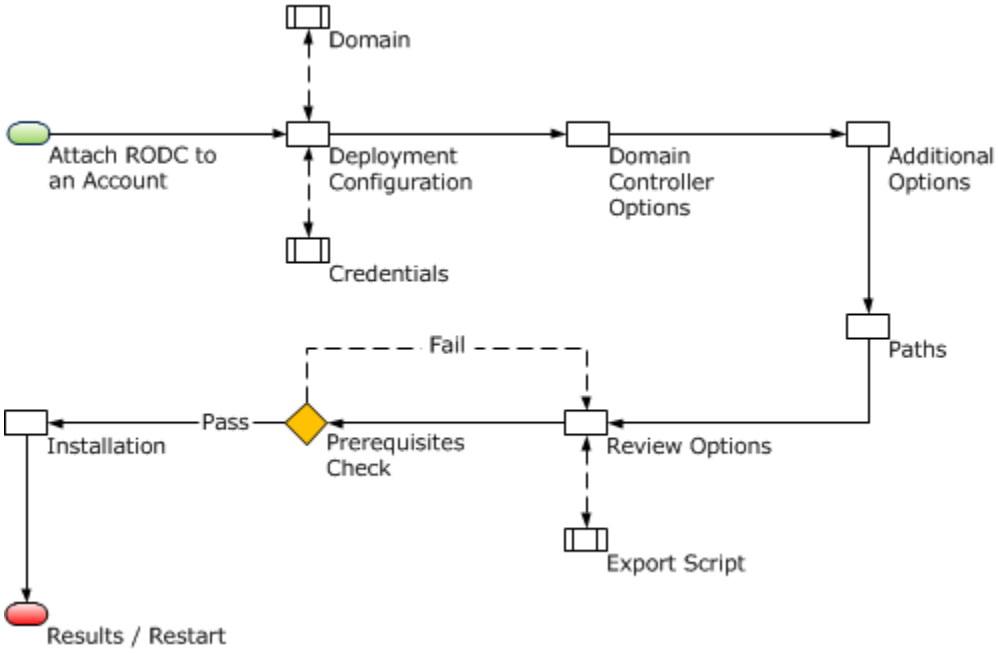

Attach RODC Windows PowerShell

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-AddsDomaincontroller	-SkipPreChecks ***-DomainName*** *-SafeModeAdministratorPassword* *-ApplicationPartitionsToReplicate* *-CreateDNSDelegation* ***-Credential*** -CriticalReplicationOnly *-DatabasePath* *-DNSDelegationCredential* *-InstallationMediaPath* *-LogPath* -Norebootoncompletion *-ReplicationSourceDC* *-SystemKey* *-SYSVOLPath* ***-UseExistingAccount***

Note

The **-credential** argument is only required if you are not already logged on as a member of the Domain Admins group.

Staging

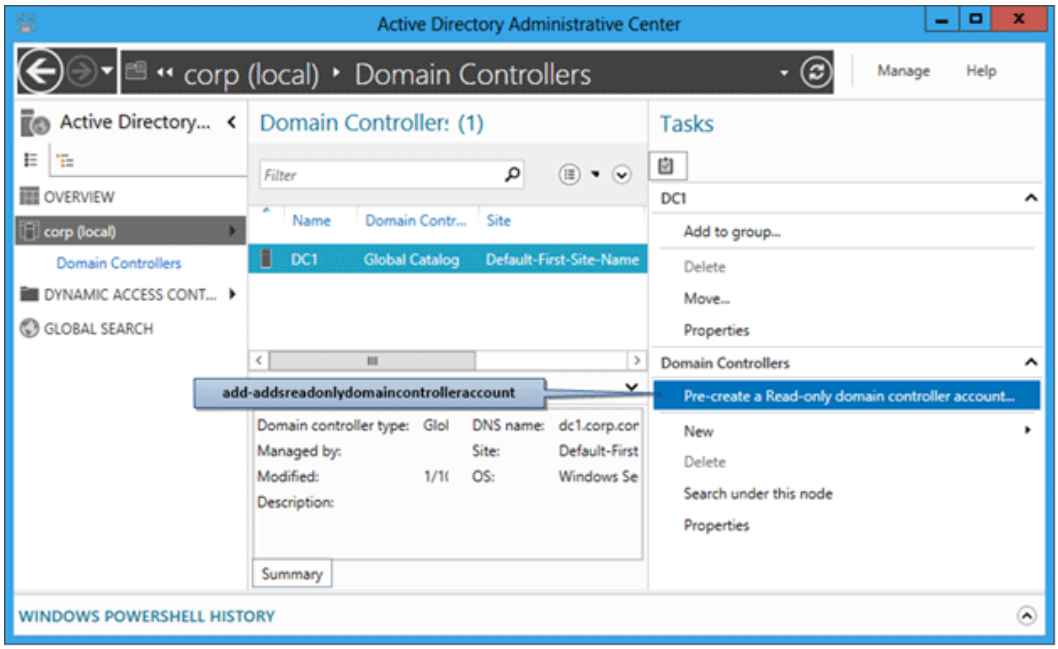

You perform the staging operation of a read-only domain controller computer account by opening the Active Directory Administrative Center (**Dsac.exe**). Click the name of the domain in the navigation pane. Double-click **Domain Controllers** in the management list. Click **Pre-create a Read-only domain controller account** in the tasks pane.

For more information about the Active Directory Administrative Center, see Advanced AD DS Management Using Active Directory Administrative Center (Level 200) and review Active Directory Administrative Center: Getting Started.

If you have experience creating read-only domain controllers, you will discover that the installation wizard has the same graphical interface as seen when using the older Active Directory Users and Computers snap-in from Windows Server 2008 and uses the same code, which includes exporting the configuration in the unattend file format used by the obsolete dcpromo.

Windows Server 2012 introduces a new ADDSDeployment cmdlet to stage RODC computer accounts, but the wizard does not use the cmdlet for its operation. The following sections display the equivalent cmdlet and arguments in order to make the information associated with each easier to understand.

The **Pre-create a Read-only domain controller account** link in the Active Directory Administrative Center's task pane is equivalent to the ADDSDeployment Windows PowerShell cmdlet:

```
Add-addsreadonlydomaincontrolleraccount
```

Welcome

The **Welcome to the Active Directory Domain Services Installation Wizard** dialog has one option named **Use advanced mode installation**. Select this option and click **Next** to show password replication policy options. Clear this option to use the default values for password replication policy options (this is discussed in further detail later in this section).

Network Credentials

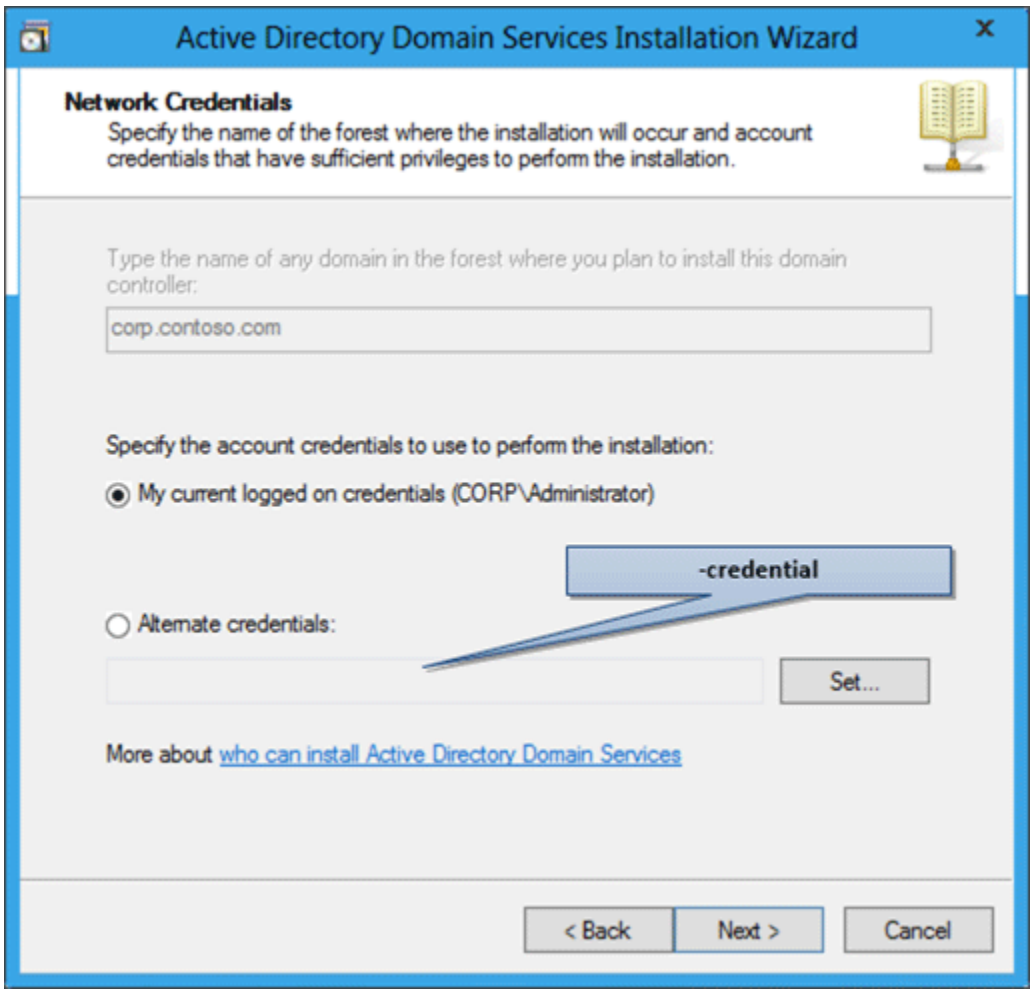

The domain name option in the **Network Credentials** dialog displays the domain targeted by the Active Directory Administrative Center by default. Your current credentials are used by default. If they do not include membership in the Domain Admins group, click **Alternate Credentials**, and click **Set** to provide the wizard with a user name and password that is a member of Domain Admins.

The equivalent ADDSDeployment Windows PowerShell argument is:

```
-credential <pscredential>
```

Keep in mind that the staging system is a direct port from Windows Server 2008 R2 and does not provide the new Adprep functionality. If you plan to deploy staged RODC accounts, you must either first deploy an un-staged RODC in that domain so that the automatic rodcprep operation runs, or manually run adprep.exe /rodcprep first.

Otherwise, you will receive error "You will not be able to install a read-only domain controller in this domain because "adprep /rodcprep" was not yet run".

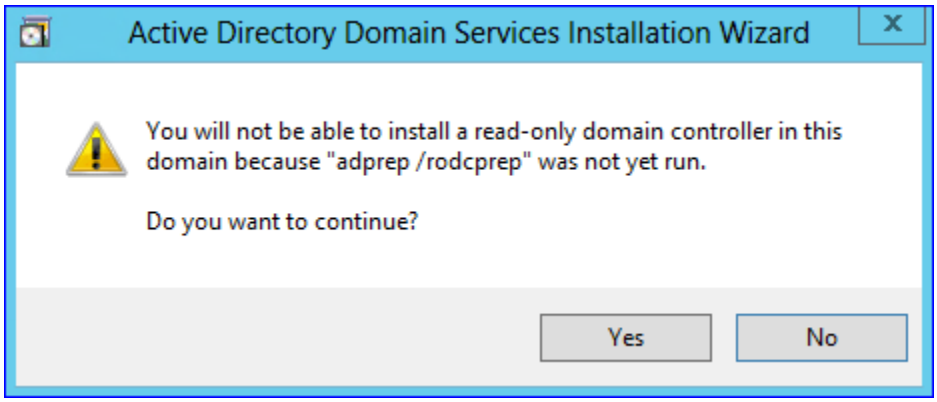

Specify the Computer Name

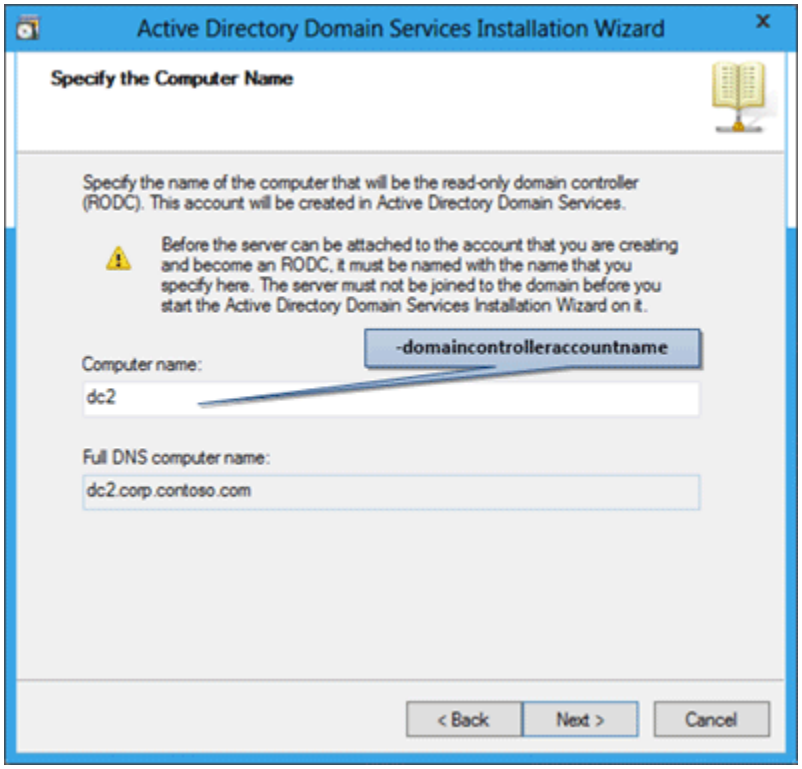

The **Specify the Computer Name** dialog requires you to enter the single-label **Computer name** of a domain controller that does not exist. The domain controller you configure and attach to this account later must have the same name, or the promotion operation will not detect the staged account.

The equivalent ADDSDeployment Windows PowerShell argument is:

```
-domaincontrolleraccountname <string>
```

Select a Site

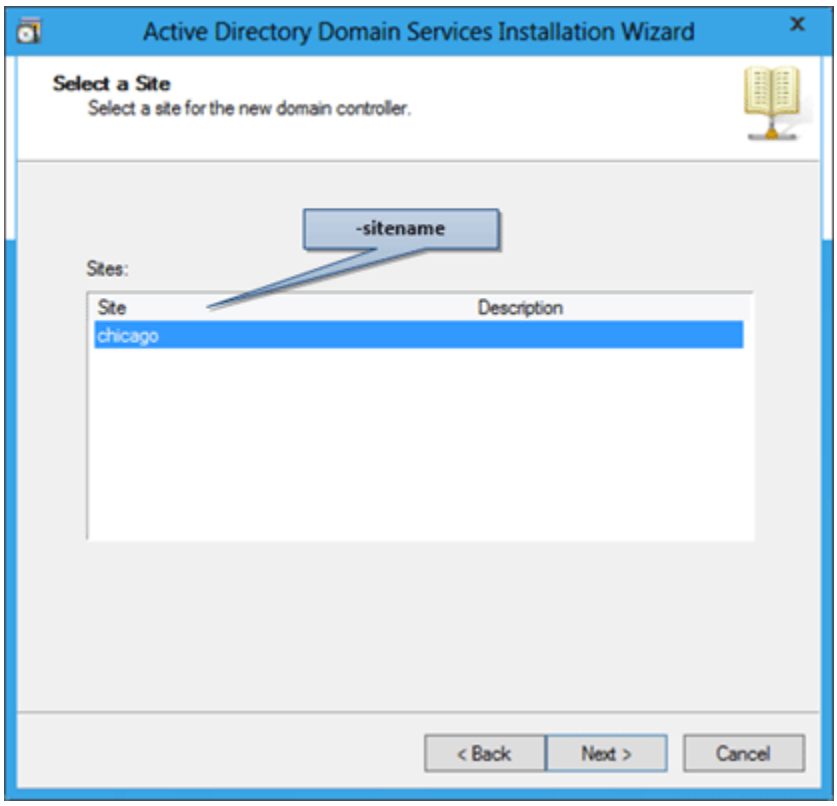

The **Select a Site** dialog shows a list of Active Directory sites for the current forest. The staged read-only domain controller operation requires you to select a single site from the list. The RODC uses this information to create its NTDS Settings object in the Configuration partition and join itself to the correct site when it starts for the first time after being deployed.

The equivalent ADDSDeployment Windows PowerShell argument is:

```
-sitename <string>
```

Additional Domain Controller Options

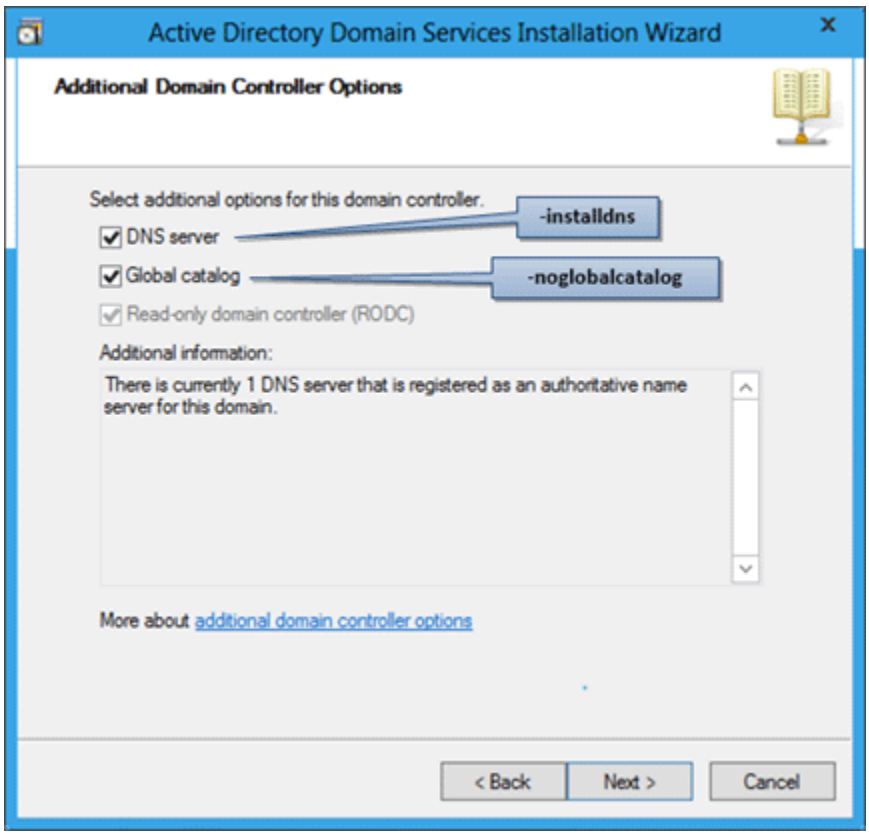

The **Additional Domain Controller Options** dialog enables you to specify that a domain controller include running as a **DNS Server** and a **Global Catalog**. Microsoft recommends that read-only domain controllers provide DNS and GC services, so both are installed by default; one intention of the RODC role is branch office scenarios where the wide area network may not be available and without those DNS and global catalog services, computers in the branch will not be able to use AD DS resources and functionality.

The **Read-only domain controller (RODC)** option is pre-selected and cannot be disabled. The equivalent ADDSDeployment Windows PowerShell arguments are:

```
-installdns <string>
```

```
-NoGlobalCatalog <{$true | $false}>
```

Note

By default, the **–NoGlobalCatalog** value is $false, which means the domain controller will be a global catalog server if the argument is not specified.

Specify the Password Replication Policy

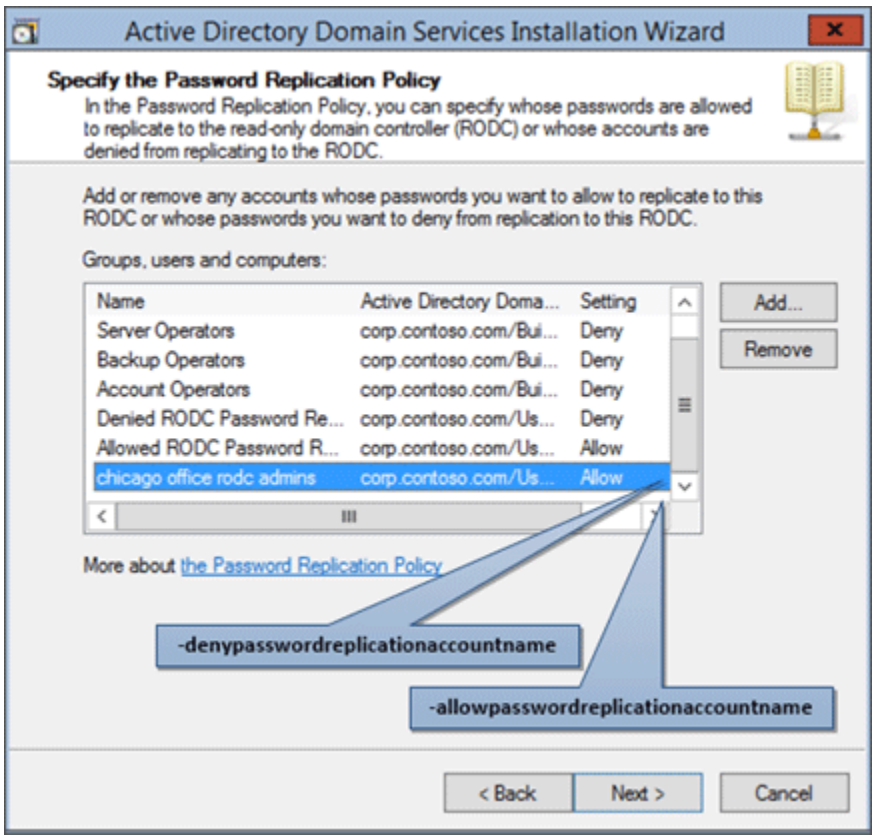

The **Specify the Password Replication Policy** dialog enables you to modify the default list of accounts that are allowed to cache their passwords on this read-only domain controller. Accounts in the list configured with **Deny** or that are not in the list (implicit) do not cache their password. Accounts that are not allowed to cache passwords on the RODC and cannot connect and authenticate to a writable domain controller cannot access resources or functionality provided by Active Directory.

Important
The wizard shows this dialog only if you select the **Use Advanced Mode Installation** check box on the welcome screen. If you clear this check box, then the wizard uses following default groups and values:

- Administrators - Deny
- Server Operators - Deny
- Backup Operators - Deny
- Account Operators - Deny
- Denied RODC Password Replication Group - Deny
- Allowed RODC Password Replication Group - Allow

The equivalent ADDSDeployment Windows PowerShell arguments are:

```
-allowpasswordreplicationaccountname <string []>

-denypasswordreplicationaccountname <string []>
```

Delegation of RODC Installation and Administration

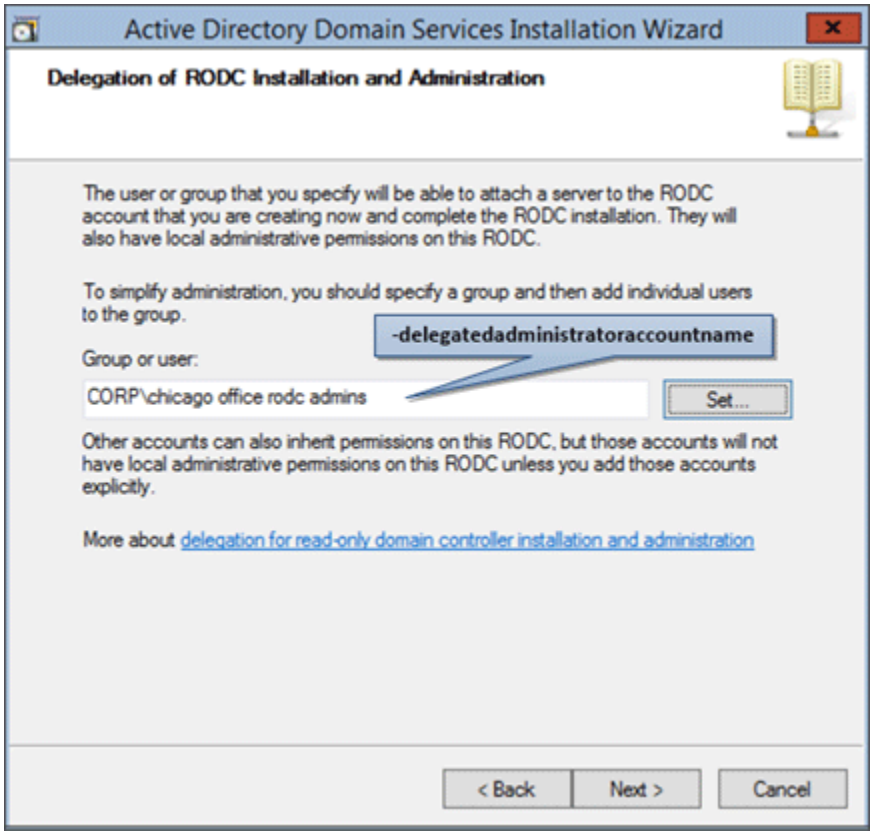

The **Delegation of RODC Installation and Administration** dialog enables you to configure a user or group containing users who are allowed to attach the server to the RODC computer

account. Click **Set** to browse the domain for a user or group. The user or group specified in this dialog gains local administrative permissions to the RODC. The specified user or members of the specified group can perform operations on the RODC with privileges equivalent to the computer's Administrators group. They are *not* members of the Domain Admins or domain built-in Administrators groups.

Use this option to delegate branch office administration without granting the branch administrator membership to the Domain Admins group. Delegating RODC administration is not required.

The equivalent ADDSDeployment Windows PowerShell argument is:

```
-delegatedadministratoraccountname <string>
```

Summary

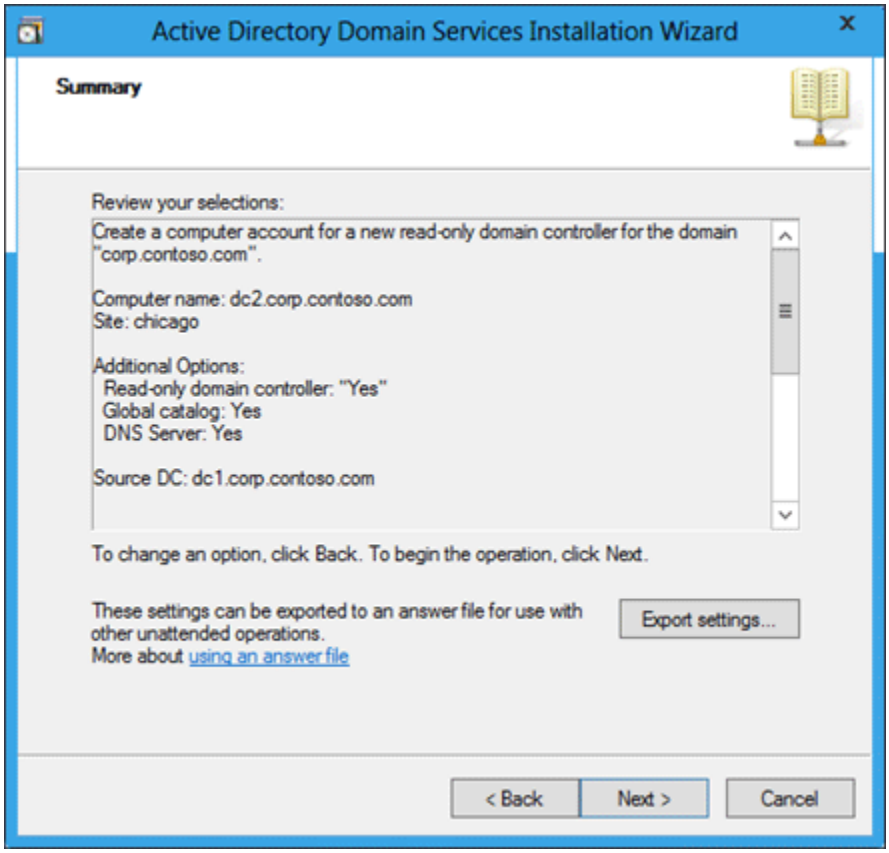

The **Summary** dialog enables you to confirm your settings. This is the last opportunity to stop the installation before the wizard creates the staged account. Click **Next** when you are ready to create the staged RODC computer account. Click **Export Settings** to save an answer file in the obsolete dcpromo unattend file format.

Creation

The **Active Directory Domain Services Installation Wizard** creates the staged read-only domain controller in Active Directory. You cannot cancel this operation after it starts.

Use the following cmdlet to stage a read-only domain controller computer account using the ADDSDeployment Windows PowerShell module:

```
Add-addsreadonlydomaincontrolleraccount
```

See Stage RODC Windows PowerShell for required and optional arguments.

Because **Add-addsreadonlydomaincontrolleraccount** only has one action with two phases (prerequisite checking and installation), the following screen shots show the installation phase with the minimum required arguments.

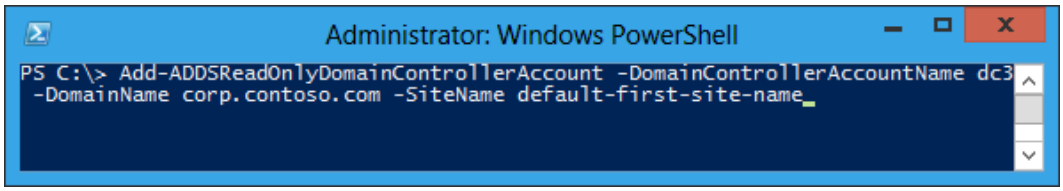

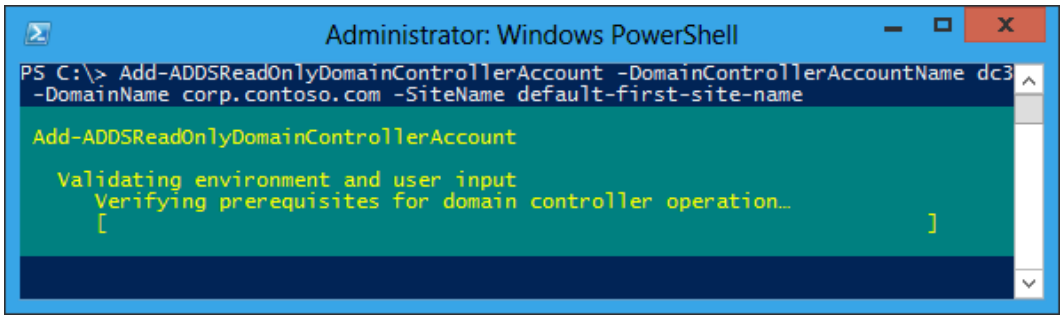

The stage RODC operation creates the RODC computer account in Active Directory. The Active Directory Administrative Center shows the **Domain Controller Type** as an **Unoccupied Domain Controller Account**. This domain controller types indicates that staged RODC account is ready for a server to attach to it as a read only domain controller.

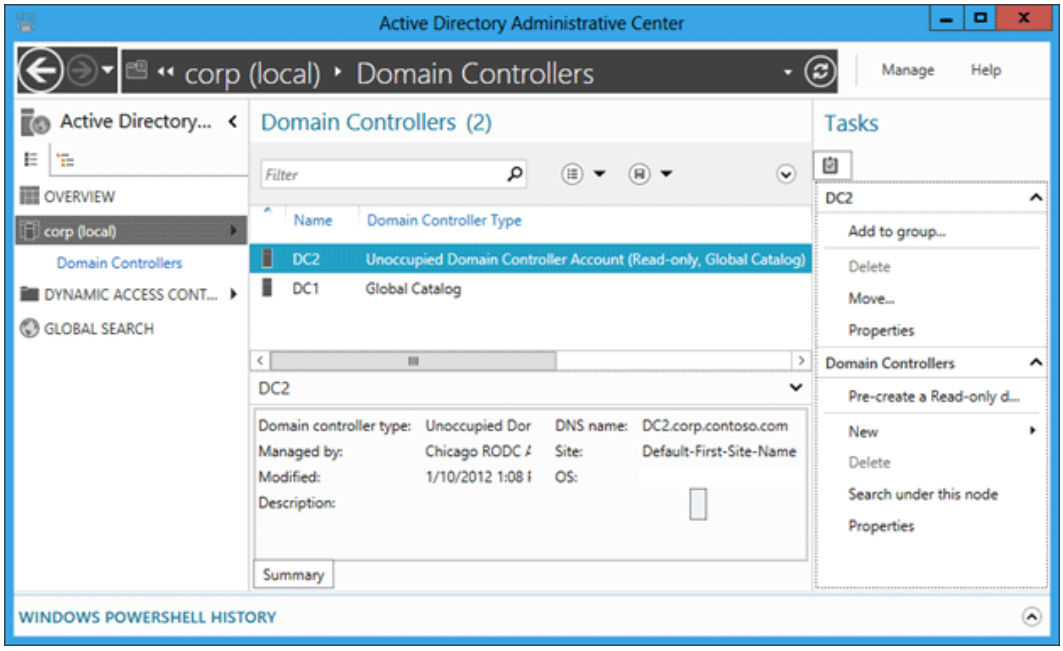

Important

The Active Directory Administrative Center is no longer required to attach a server to a read-only domain controller computer account. Use Server Manager and the Active Directory Domain Services Configuration Wizard or the ADDSDeployment Windows PowerShell module cmdlet **Install-AddsDomainController** to attach a new RODC to its staged account. The steps are similar to adding a new writable domain controller to an existing domain, with the exception that the staged RODC computer account contains configuration options decided at the time you staged the RODC computer account.

Attaching

Deployment Configuration

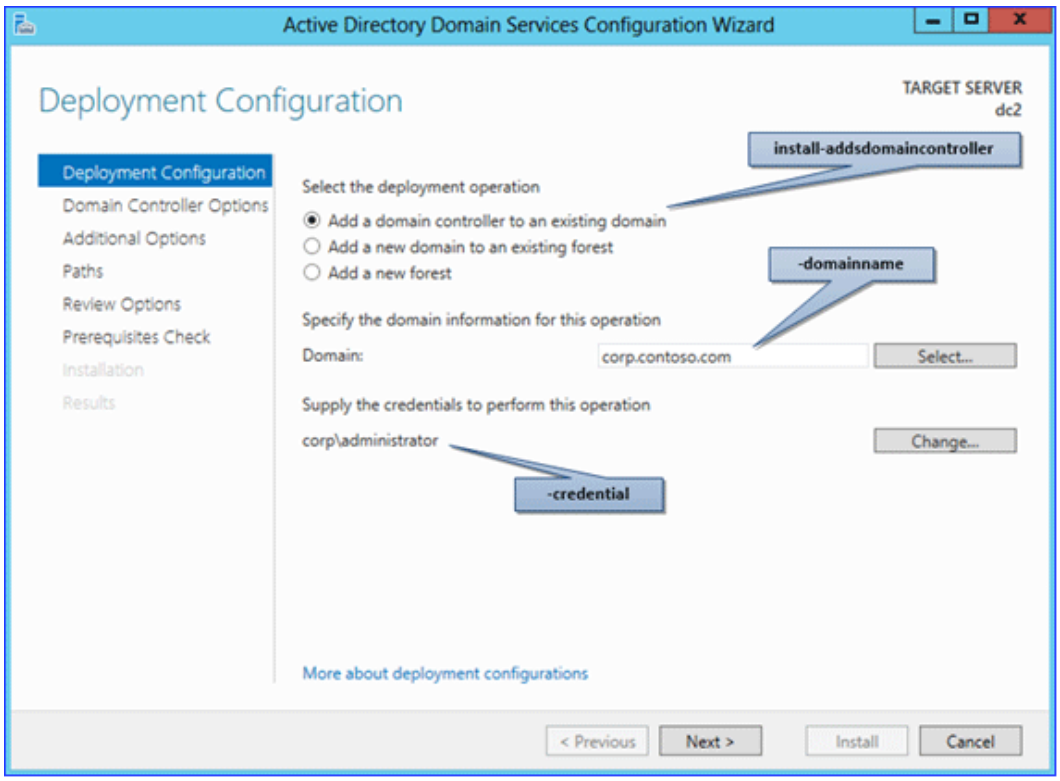

Server Manager begins every domain controller promotion with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select.

To add a read-only domain controller to an existing domain, select **Add a domain controller to an existing domain** and click the **Select** button to **Specify the domain information for this domain**. Server Manager automatically prompts you for valid credentials, or you can click **Change**.

Attaching an RODC requires membership in the Domain Admins groups in Windows Server 2012. The Active Directory Domain Services Configuration Wizard prompts you later if your current credentials do not have adequate permissions or group memberships.

The **Deployment Configuration** ADDSDeployment Windows PowerShell cmdlet and arguments are:

```
Install-AddsDomainController

-domainname <string>

-credential <pscredential>
```

Domain Controller Options

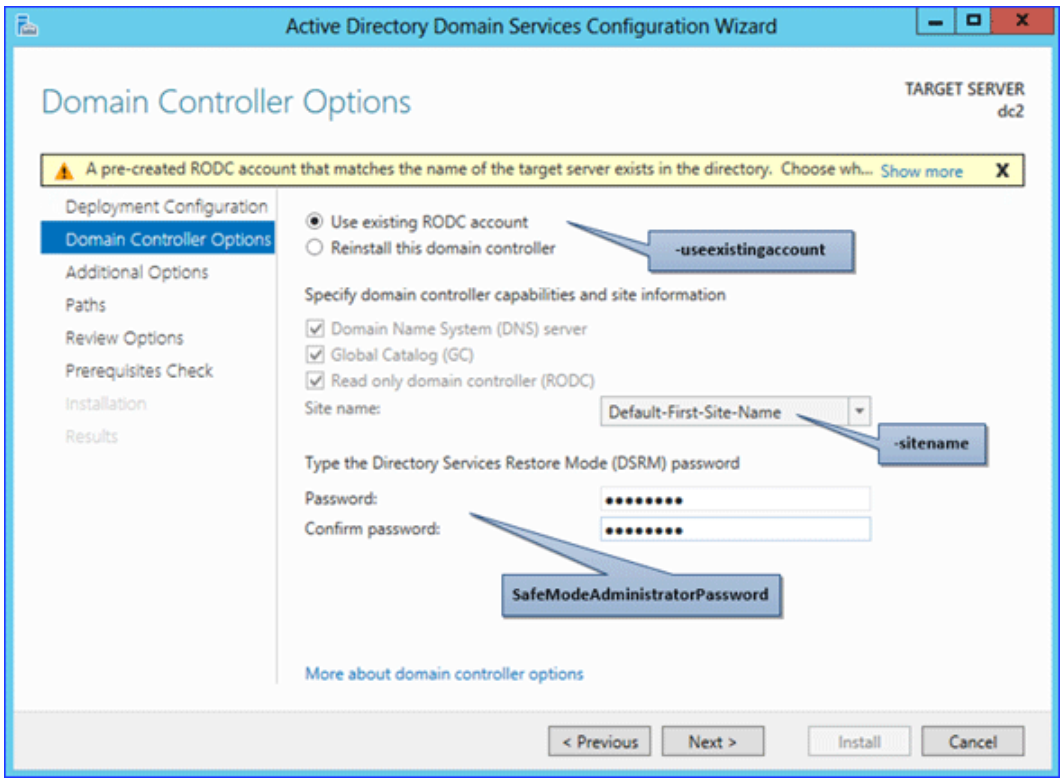

The **Domain Controller Options** page shows the domain controller options for the new domain controller. When this page loads, the Active Directory Domain Services Configuration Wizard sends an LDAP query to an existing domain controller to check for unoccupied accounts. If the query finds an unoccupied domain controller computer account that shares the same name as the current computer, then the wizard displays an informational message at the top of the page that reads "**A Pre-created RODC account that matches the name of the target server exists in the directory. Choose whether to use this existing RODC account or reinstall this domain controller**." The wizard uses the **Use existing RODC account** as the default configuration.

Important

You can use the **Reinstall this domain controller** option when a domain controller has suffered a physical problem and cannot return to functionality. This saves time when configuring the replacement domain controller, by leaving the domain controller computer account and object metadata in Active Directory. Install the new computer with the *same name*, and promote it as a domain controller in the domain. The **Reinstall this domain controller** option is unavailable if you removed the domain controller object's metadata from Active Directory (metadata cleanup).

You cannot configure domain controller options when you are attaching a server to an RODC computer account. You configure domain controller options when you create the staged RODC computer account.

The specified **Directory Services Restore Mode Password** must adhere to the password policy applied to the server. Always choose a strong, complex password or preferably, a passphrase.

The **Domain Controller Options** ADDSDeployment Windows PowerShell arguments are:

```
-UseExistingAccount <{$true | $false}>

-SafeModeAdministratorPassword <secure string>
```

Important

The site name must already exist when provided as an argument to **-sitename**. The **install-AddsDomainController** cmdlet does not create site names. You can use cmdlet **new-adreplicationsite** to create new sites.

The **Install-ADDSDomainController** arguments follow the same defaults as Server Manager if not specified.

The **SafeModeAdministratorPassword** argument's operation is special:

- If *not specified* as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.

  For example, to create a new RODC in the corp.contoso.com and be prompted to enter and confirm a masked password:

  ```
  Install-ADDSDomainController –DomainName corp.contoso.com –
  credential (get-credential)
  ```

- If specified *with a value*, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string:

```
-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged.

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Finally, you could store the obfuscated password in a file, and then reuse it later, without the clear text password ever appearing. For example:

```
$file = "c:\pw.txt"

$pw = read-host -prompt "Password:" -assecurestring

$pw | ConvertFrom-SecureString | Set-Content $file


-safemodeadministratorpassword (Get-Content $File | ConvertTo-SecureString)
```

Warning

Providing or storing a clear or obfuscated text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM password of that domain controller. Anyone with access to the file could reverse that obfuscated password. With that knowledge, they can logon to a DC started in DSRM and eventually impersonate the domain controller itself, elevating their privileges to the highest level in an AD forest. An additional set of steps using **System.Security.Cryptography** to encrypt the text file data is advisable but out of scope. The best practice is to totally avoid password storage.

Additional Options

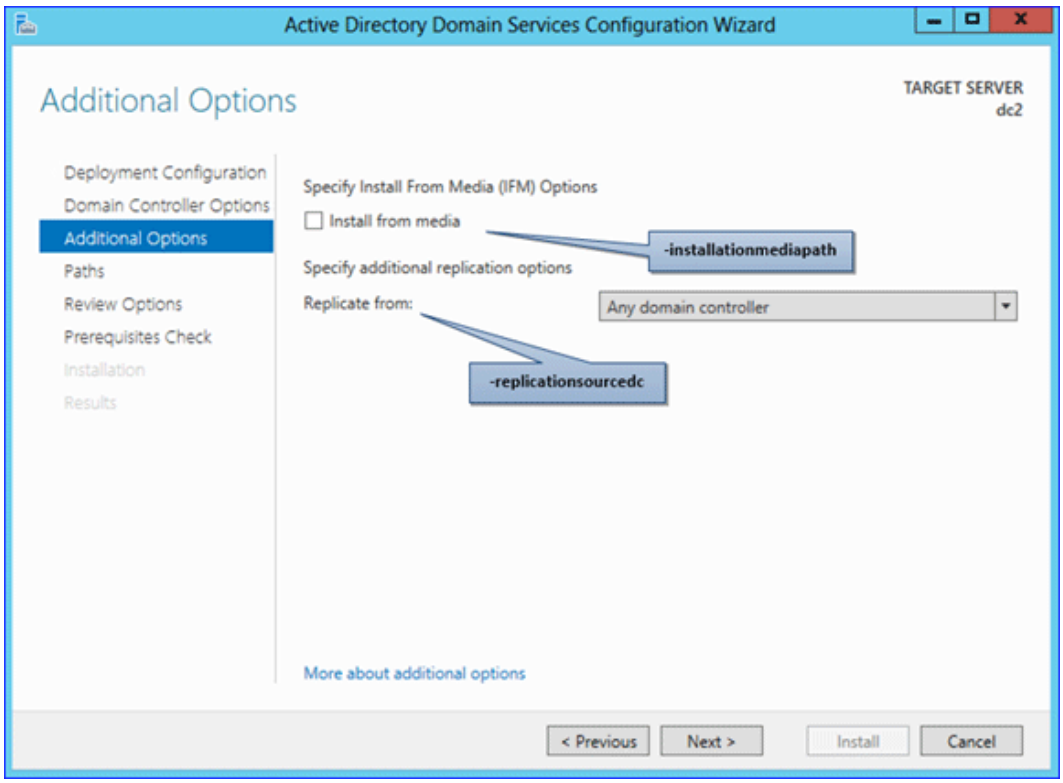

The **Additional Options** page provides configuration options to name a domain controller as the replication source, or you can use any domain controller as the replication source.

You can also choose to install the domain controller using backed up media using the Install from media (IFM) option. The **Install from media** checkbox provides a browse option once selected and you must click **Verify** to ensure the provided path is valid media. Media used by the IFM option is created with Windows Server Backup or Ntdsutil.exe from another existing Windows Server 2012 computer only; you cannot use a Windows Server 2008 R2 or previous operating system to create media for a Windows Server 2012 domain controller. For more information about changes in IFM, see Ntdsutil.exe Install from Media Changes. If using media protected with a SYSKEY, Server Manager prompts for the image's password during verification.

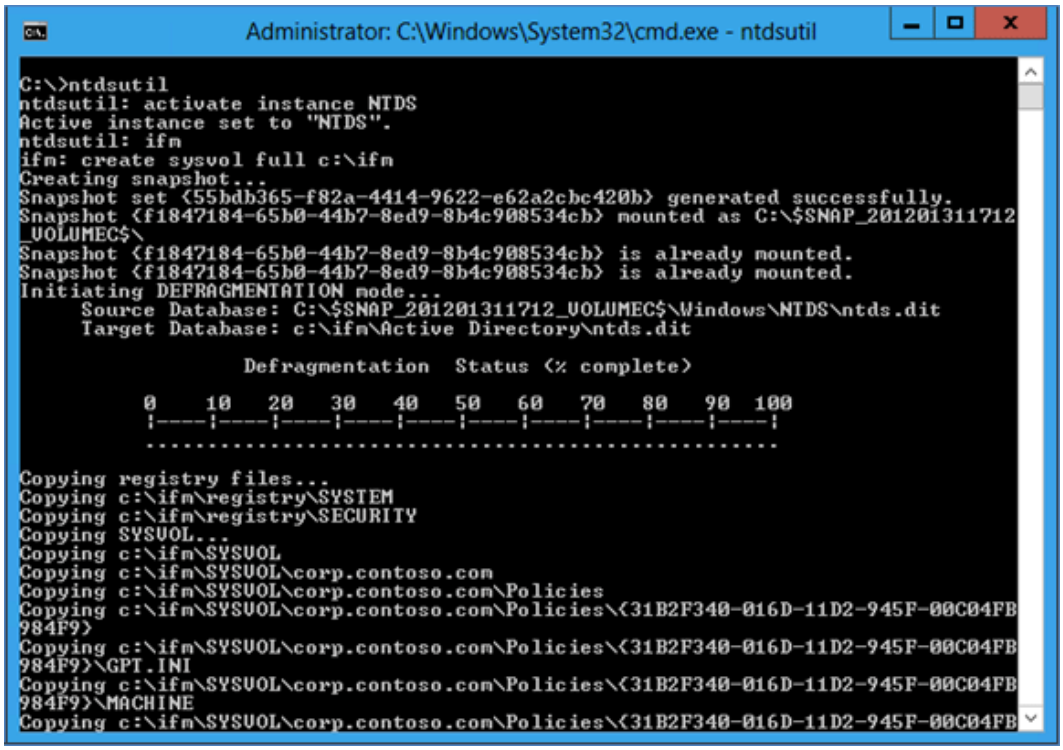

The **Additional Options** ADDSDeployment cmdlet arguments are:

`-replicationsourcedc` *<string>*

`-installationmediapath` *<string>*

`-systemkey` *<secure string>*

Paths

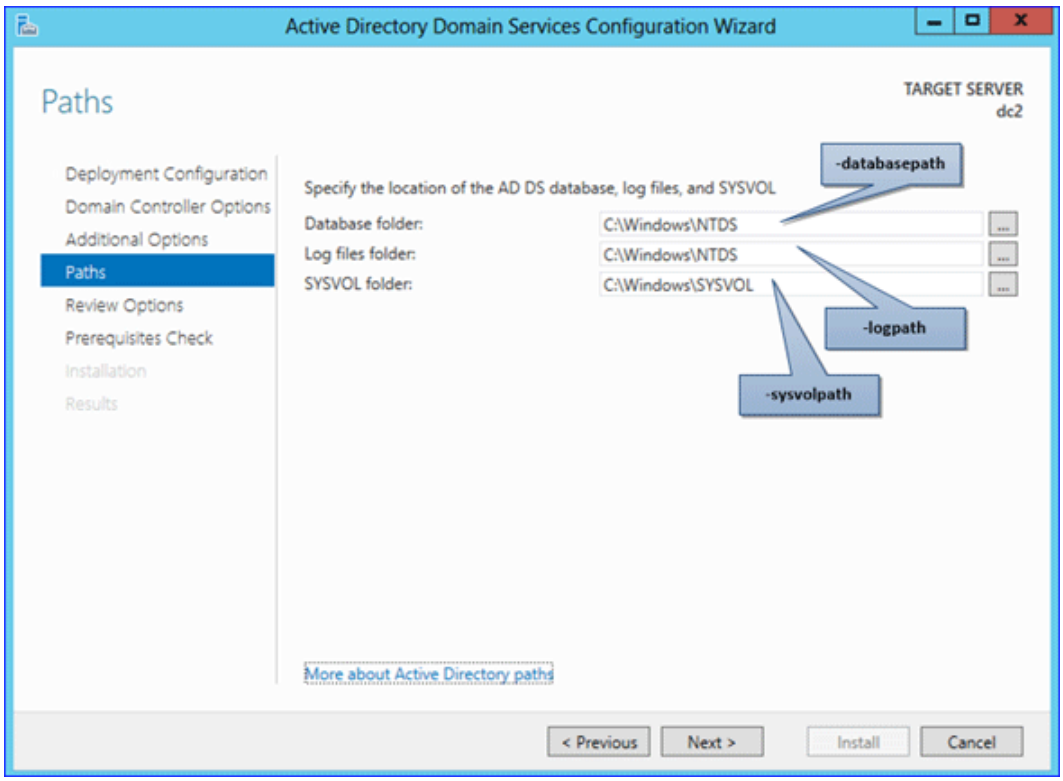

The **Paths** page enables you to override the default folder locations of the AD DS database, the database transaction logs, and the SYSVOL share. The default locations are always in subdirectories of %systemroot%. The **Paths** ADDSDeployment cmdlet arguments are:

`-databasepath <string>`

`-logpath <string>`

`-sysvolpath <string>`

Review Options and View Script

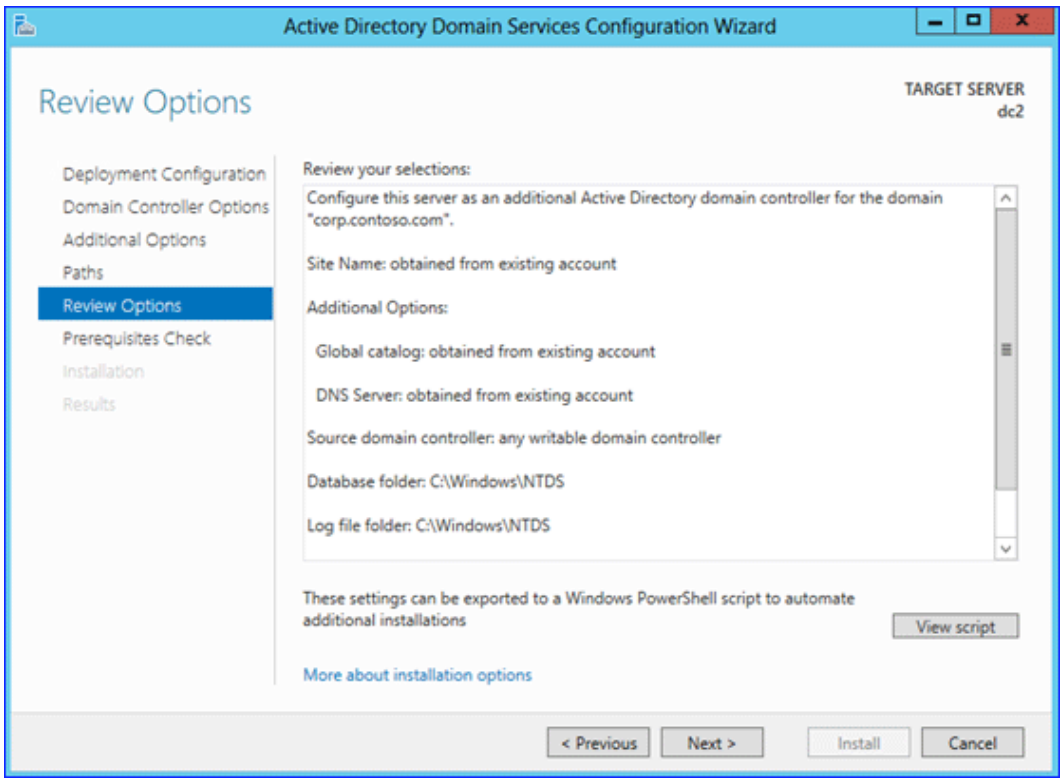

The **Review Options** page enables you to validate your settings and ensure that they meet your requirements before you start the installation. This is not the last opportunity to stop the installation using Server Manager. This page simply enables you to review and confirm your settings before continuing the configuration. The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use. For example:

```
#

# Windows PowerShell Script for AD DS Deployment

#


Import-Module ADDSDeployment

Install-ADDSDomainController `

-Credential (Get-Credential) `

-CriticalReplicationOnly:$false `
```

```
-DatabasePath "C:\Windows\NTDS" `
-DomainName "corp.contoso.com" `
-LogPath "C:\Windows\NTDS" `
-SYSVOLPath "C:\Windows\SYSVOL" `
-UseExistingAccount:$true `
-Norebootoncompletion:$false
-Force:$true
```

Note

Server Manager generally fills in all arguments with values when promoting and does not rely on defaults (as they may change between future versions of Windows or service packs). The one exception to this is the **-safemodeadministratorpassword** argument. To force a confirmation prompt omit the value when running cmdlet interactively

Use the optional **Whatif** argument with the **Install-ADDSDomainController** cmdlet to review configuration information. This enables you to see the explicit and implicit values of the arguments for a cmdlet.

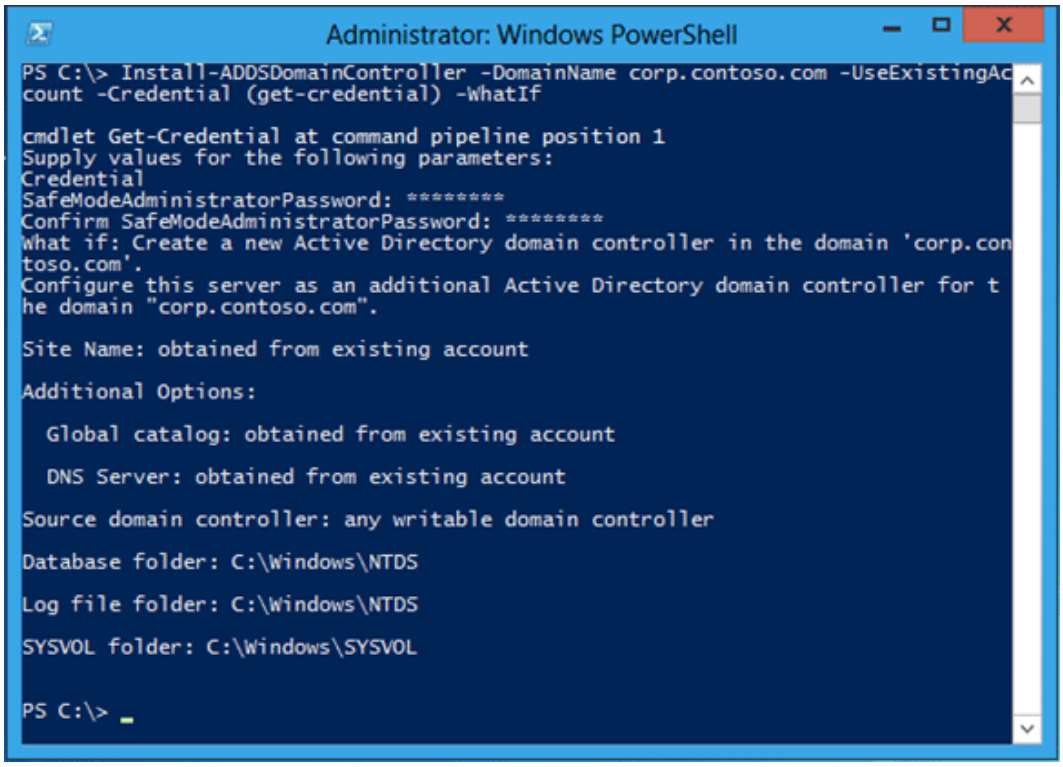

Prerequisites Check

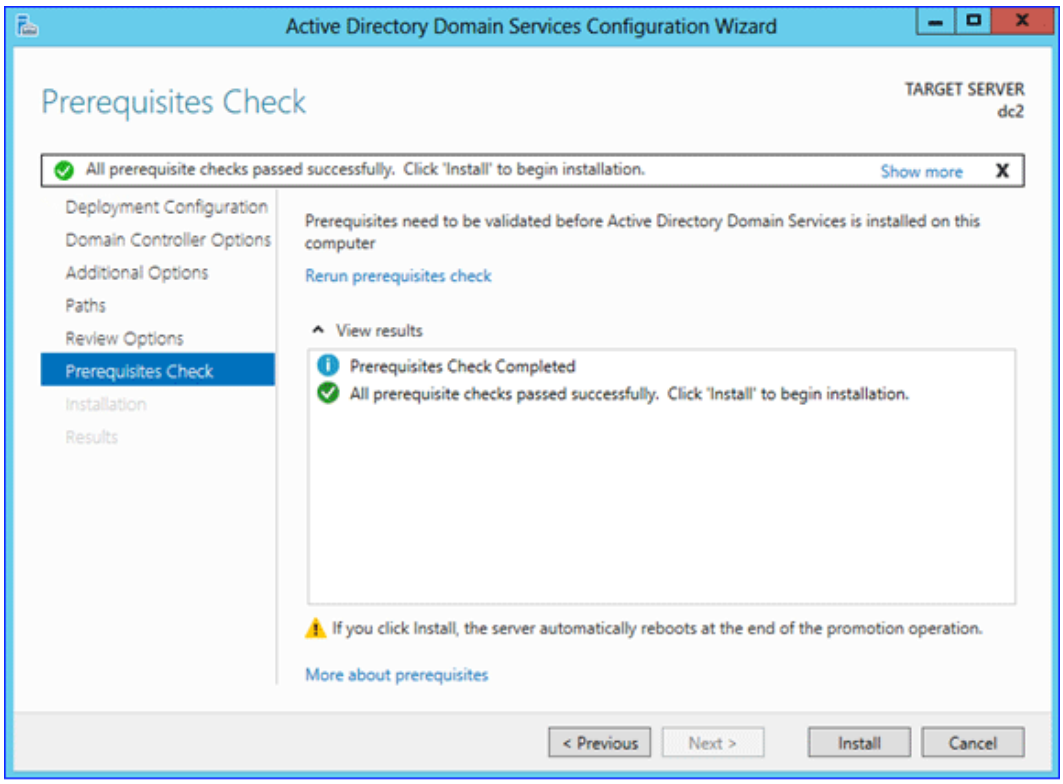

The **Prerequisites Check** is a new feature in AD DS domain configuration. This new phase validates that the server configuration is capable of supporting a new AD DS forest.

When installing a new forest root domain, the Server Manager Active Directory Domain Services Configuration Wizard invokes a series of serialized modular tests. These tests alert you with suggested repair options. You can run the tests as many times as required. The domain controller installation process cannot continue until all prerequisite tests pass.

The **Prerequisites Check** also surfaces relevant information such as security changes that affect older operating systems. For more information about the prerequisite checks, see Prerequisite Checking.

You cannot bypass the **Prerequisite Check** when using Server Manager, but you can skip the process when using the AD DS Deployment cmdlet using the following argument:

```
-skipprechecks
```

Warning

Microsoft discourages skipping the prerequisite check as it can lead to a partial domain controller promotion or damaged AD DS forest.

Click **Install** to begin the domain controller promotion process. This is last opportunity to cancel the installation. You cannot cancel the promotion process once it begins. The computer will reboot automatically at the end of promotion, regardless of the promotion results.

Installation

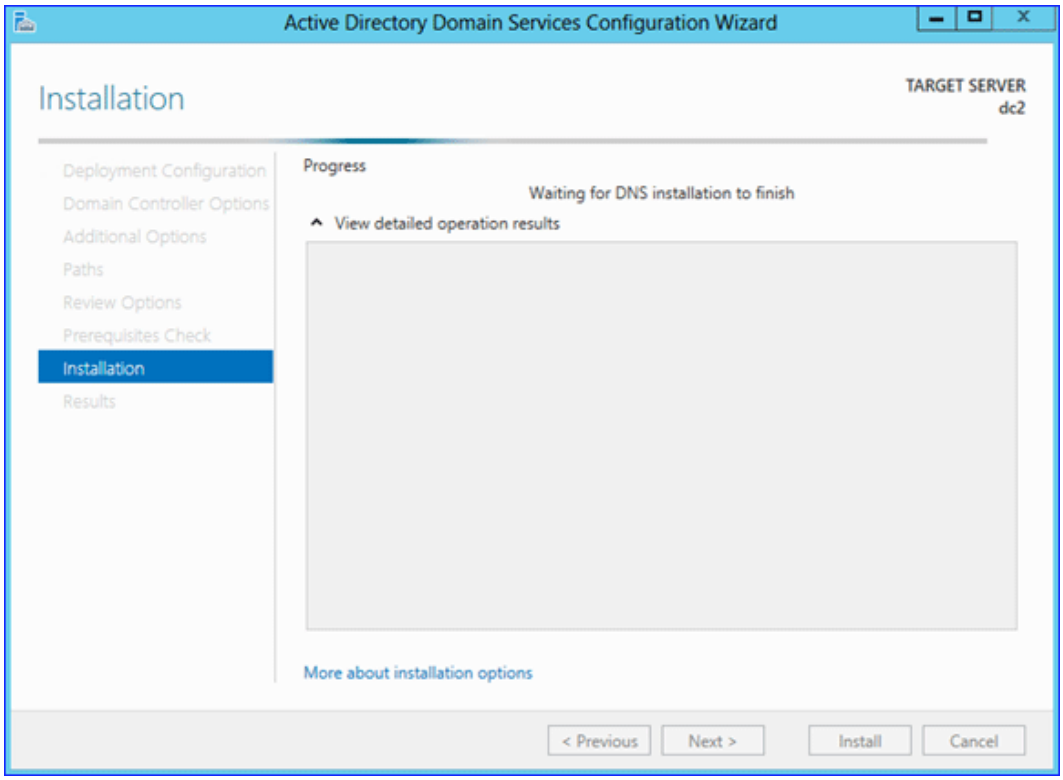

When the Installation page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and are written to logs:

- %systemroot%\debug\dcpromo.log
- %systemroot%\debug\dcpromoui.log

To install a new Active Directory forest using the ADDSDeployment module, use the following cmdlet:

```
Install-addsdomaincontroller
```

See Attach RODC Windows PowerShell for required and optional arguments.

The **Install-addsdomaincontroller** cmdlet only has two phases (prerequisite checking and installation). The two figures below show the installation phase with the minimum required arguments of **-domainname**, **-useexistingaccount**, and **-credential**. Note how, just like Server Manager, **Install-ADDSDomainController** reminds you that promotion will reboot the server automatically:

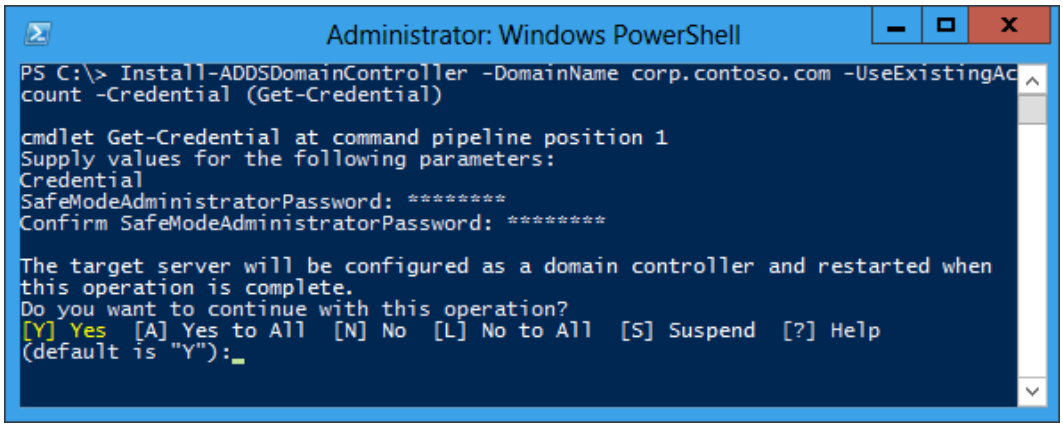

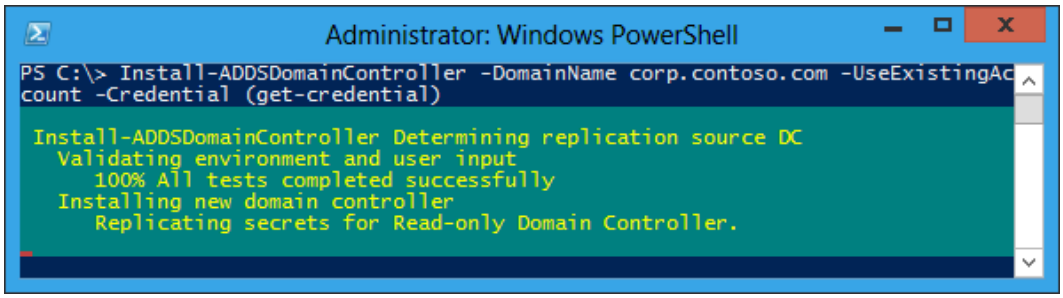

To accept the reboot prompt automatically, use the **-force** or **-confirm:$false** arguments with any ADDSDeployment Windows PowerShell cmdlet. To prevent the server from automatically rebooting at the end of promotion, use the **-norebootoncompletion** argument.

Warning

Overriding the reboot is discouraged. The domain controller must reboot to function correctly.

Results

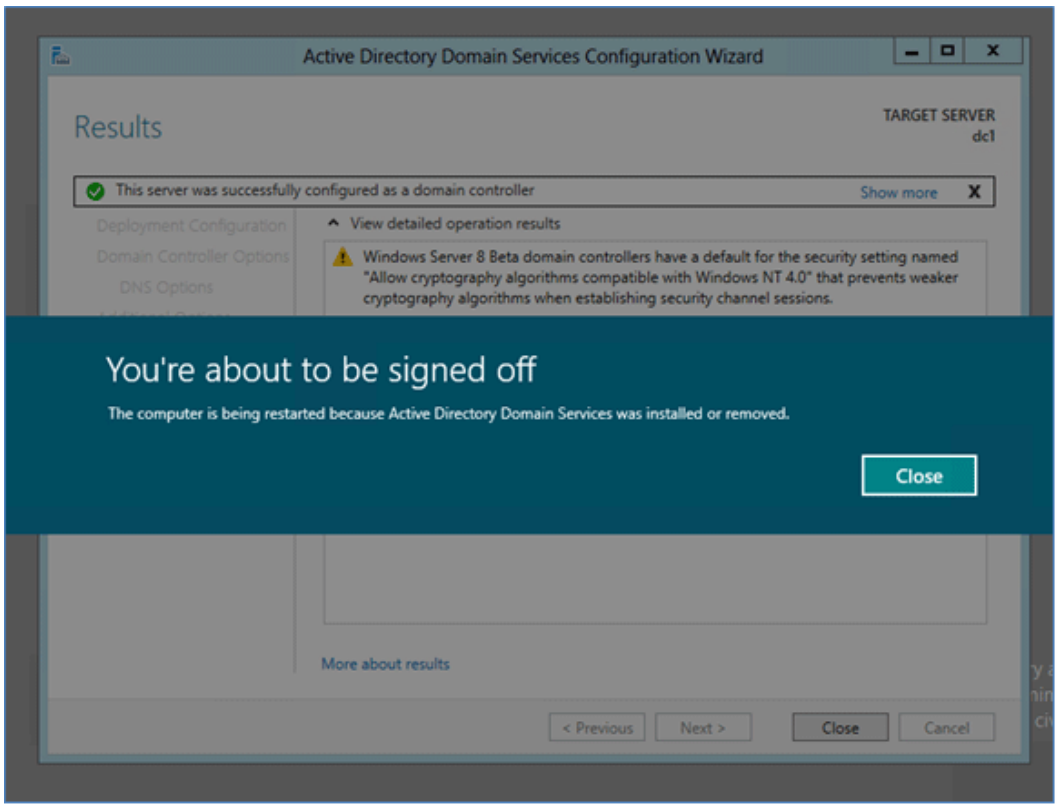

The **Results** page shows the success or failure of the promotion and any important administrative information. The domain controller will automatically reboot after 10 seconds.

RODC without Staging Workflow

The following diagram illustrates the Active Directory Domain Services configuration process, when you previously installed the AD DS role and you have started the Active Directory Domain Services Configuration Wizard using Server Manager to create a new non-staged read-only domain controller in an existing Windows Server 2012 domain.

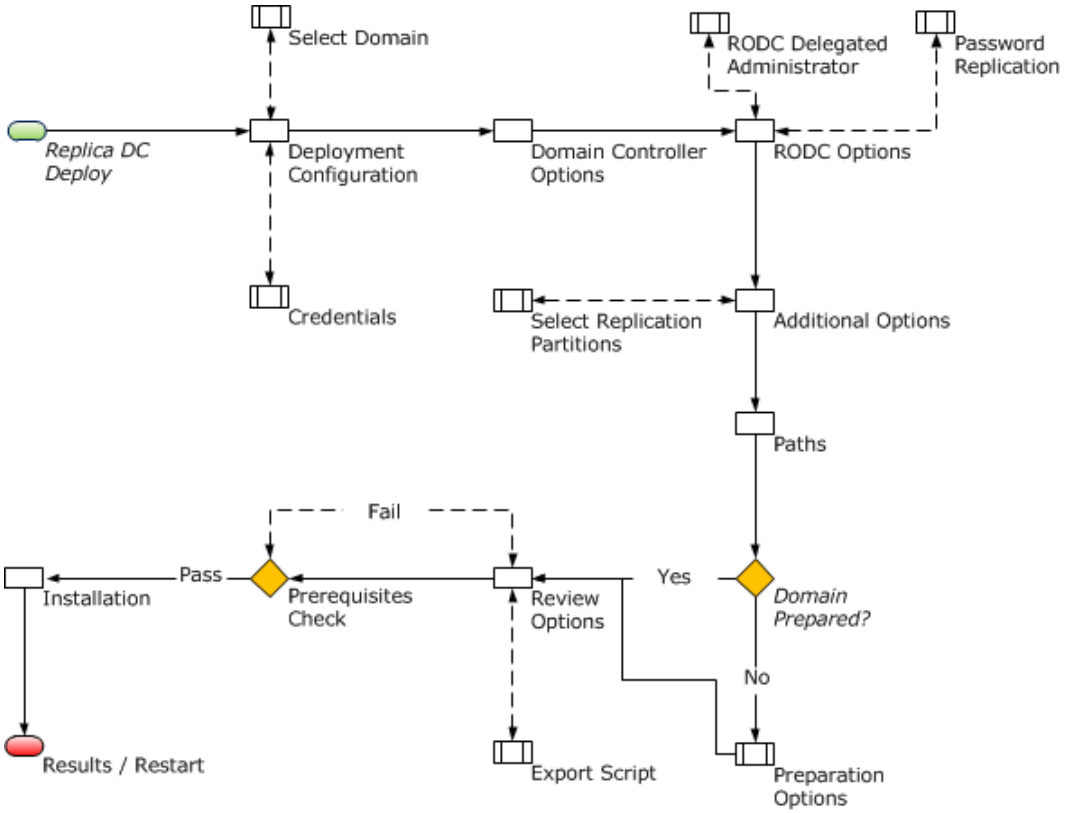

RODC without Staging Windows PowerShell

ADDSDeployment Cmdlet	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Install-AddsDomainController	-SkipPreChecks ***-DomainName*** *-SafeModeAdministratorPassword* ***-SiteName*** *-ApplicationPartitionsToReplicate* *-CreateDNSDelegation* ***-Credential*** *-CriticalReplicationOnly* *-DatabasePath* *-DNSDelegationCredential*

	-DNSOnNetwork *-InstallationMediaPath* *-InstallDNS* *-LogPath* -MoveInfrastructureOperationMasterRoleIfNecessary *-NoGlobalCatalog* -Norebootoncompletion *-ReplicationSourceDC* -SkipAutoConfigureDNS *-SystemKey* *-SYSVOLPath* *-AllowPasswordReplicationAccountName* *-DelegatedAdministratorAccountName* *-DenyPasswordReplicationAccountName* ***-ReadOnlyReplica***

Note
The **-credential** argument is only required if you are not already logged on as a member of the Domain Admins group.

RODC without Staging Deployment

Deployment Configuration

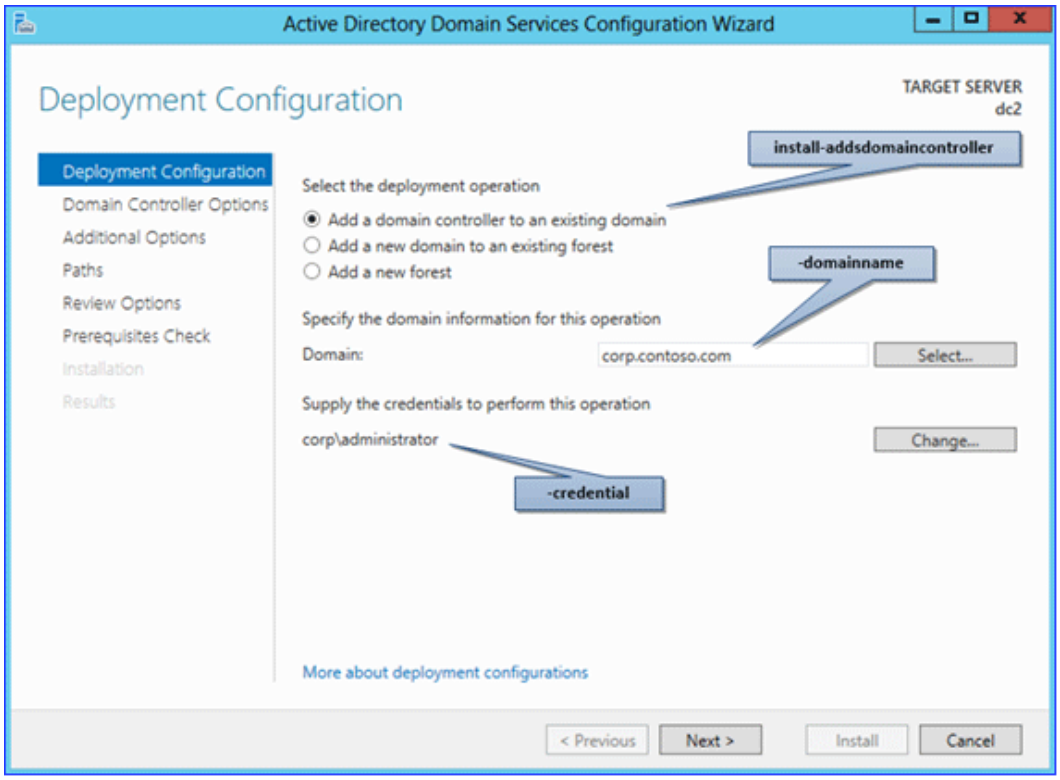

Server Manager begins every domain controller promotion with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select.

To add an un-staged read-only domain controller to an existing Windows Server 2012 domain, select **Add a domain controller to an existing domain** and click the **Select** button to **Specify the domain information for this domain**. Server Manager automatically prompts you for valid credentials, or you can click **Change**.

Attaching an RODC requires membership in the Domain Admins groups in Windows Server 2012. The Active Directory Domain Services Configuration Wizard prompts you later if your current credentials do not have adequate permissions or group memberships.

The **Deployment Configuration** ADDSDeployment Windows PowerShell cmdlet and arguments are:

```
Install-AddsDomainController

-domainname <string>

-credential <pscredential>
```

Domain Controller Options

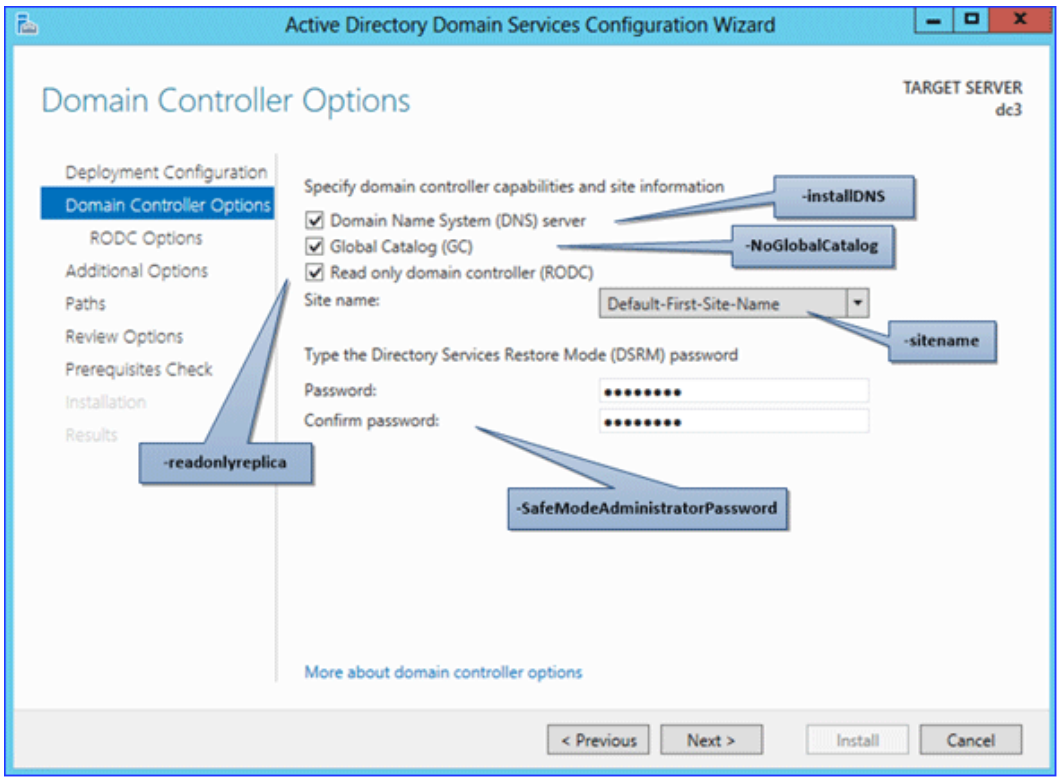

The **Domain Controller Options** page specifies the domain controller capabilities for the new domain controller. The configurable domain controller capabilities are **DNS server**, **Global Catalog**, and **Read-only domain controller**. Microsoft recommends that all domain controllers provide DNS and GC services for high availability in distributed environments. GC is always selected by default and DNS server is selected by default if the current domain hosts DNS already on its DCs based on Start of Authority query.

The **Domain Controller Options** page also enables you to choose the appropriate Active Directory logical **site name** from the forest configuration. By default, it selects the site with the most correct subnet. If there is only one site, it selects that site automatically.

Important

If the server does not belong to an Active Directory subnet and there is more than one Active Directory site, nothing is selected and the **Next** button is unavailable until you choose a site from the list.

The specified **Directory Services Restore Mode Password** must adhere to the password policy applied to the server. Always choose a strong, complex password or preferably, a passphrase.The **Domain Controller Options** ADDSDeployment Windows PowerShell arguments are:

```
-UseExistingAccount <{$true | $false}>

-SafeModeAdministratorPassword <secure string>
```

Important

The site name must already exist when provided as an argument to **-sitename**. The **install-AddsDomainController** cmdlet does not create site names. You can use cmdlet **new-adreplicationsite** to create new sites.

The **Install-ADDSDomainController** arguments follow the same defaults as Server Manager if not specified.

The **SafeModeAdministratorPassword** argument's operation is special:

- If *not specified* as an argument, the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively.

  For example, to create a new RODC in the corp.contoso.com and be prompted to enter and confirm a masked password:

  ```
  Install-ADDSDomainController –DomainName corp.contoso.com –
  credential (get-credential)
  ```

- If specified *with a value*, the value must be a secure string. This is not the preferred usage when running the cmdlet interactively.

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string:

```
-safemodeadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous option does not confirm the password, use extreme caution: the password is not visible.

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged.

```
-safemodeadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Finally, you could store the obfuscated password in a file, and then reuse it later, without the clear text password ever appearing. For example:

```
$file = "c:\pw.txt"

$pw = read-host -prompt "Password:" -assecurestring

$pw | ConvertFrom-SecureString | Set-Content $file


-safemodeadministratorpassword (Get-Content $File | ConvertTo-SecureString)
```

Warning

Providing or storing a clear or obfuscated text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the DSRM

password of that domain controller. Anyone with access to the file could reverse that obfuscated password. With that knowledge, they can logon to a DC started in DSRM and eventually impersonate the domain controller itself, elevating their privileges to the highest level in an AD forest. An additional set of steps using **System.Security.Cryptography** to encrypt the text file data is advisable but out of scope. The best practice is to totally avoid password storage.

RODC Options

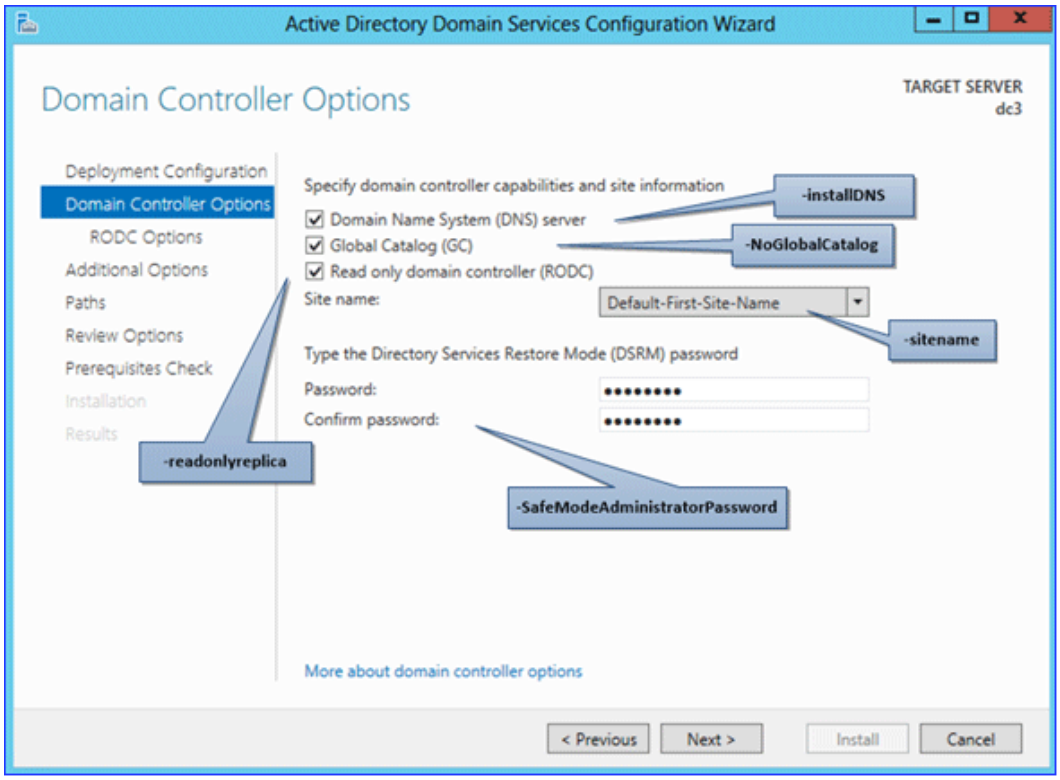

The **RODC Options** page enables you to modify the settings:

- Delegated Administrator Account
- Accounts that are allowed to replicate passwords to the RODC
- Accounts that are denied from replicating passwords to the RODC

Delegated administrator accounts gain local administrative permissions to the RODC. These users can operate with privileges equivalent to the local computer's Administrators group. They are not members of the Domain Admins or the domain built-in Administrators groups. This option is useful for delegating branch office administration without giving out domain administrative permissions. Configuring delegation of administration is not required.

The equivalent ADDSDeployment Windows PowerShell argument is:

```
-delegatedadministratoraccountname <string>
```

Accounts that are not allowed to cache passwords on the RODC and cannot connect and authenticate to a writable domain controller cannot access resources or functionality provided by Active Directory.

Important

If not modified, the default groups and settings are used:

- Administrators - Deny
- Server Operators - Deny
- Backup Operators - Deny
- Account Operators - Deny
- Denied RODC Password Replication Group - Deny
- Allowed RODC Password Replication Group - Allow

The equivalent ADDSDeployment Windows PowerShell arguments are:

```
-allowpasswordreplicationaccountname <string []>

-denypasswordreplicationaccountname <string []>
```

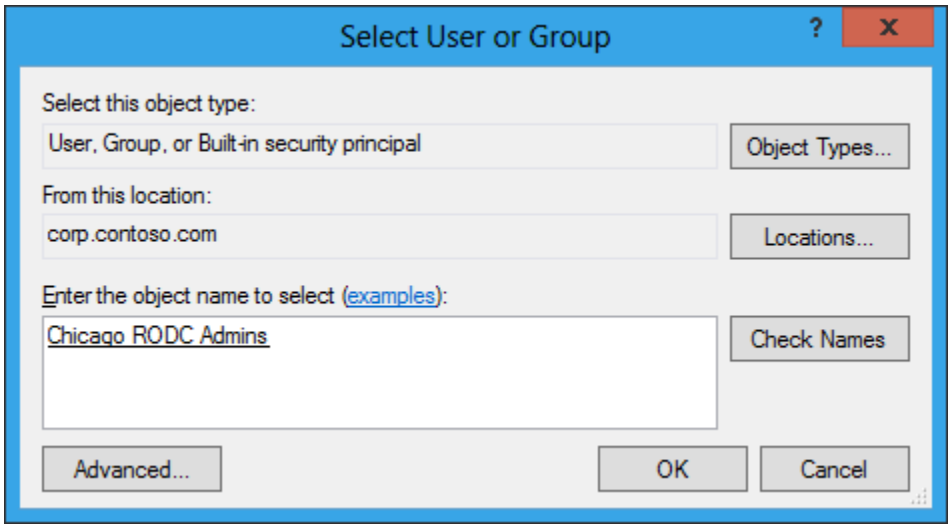

Additional Options

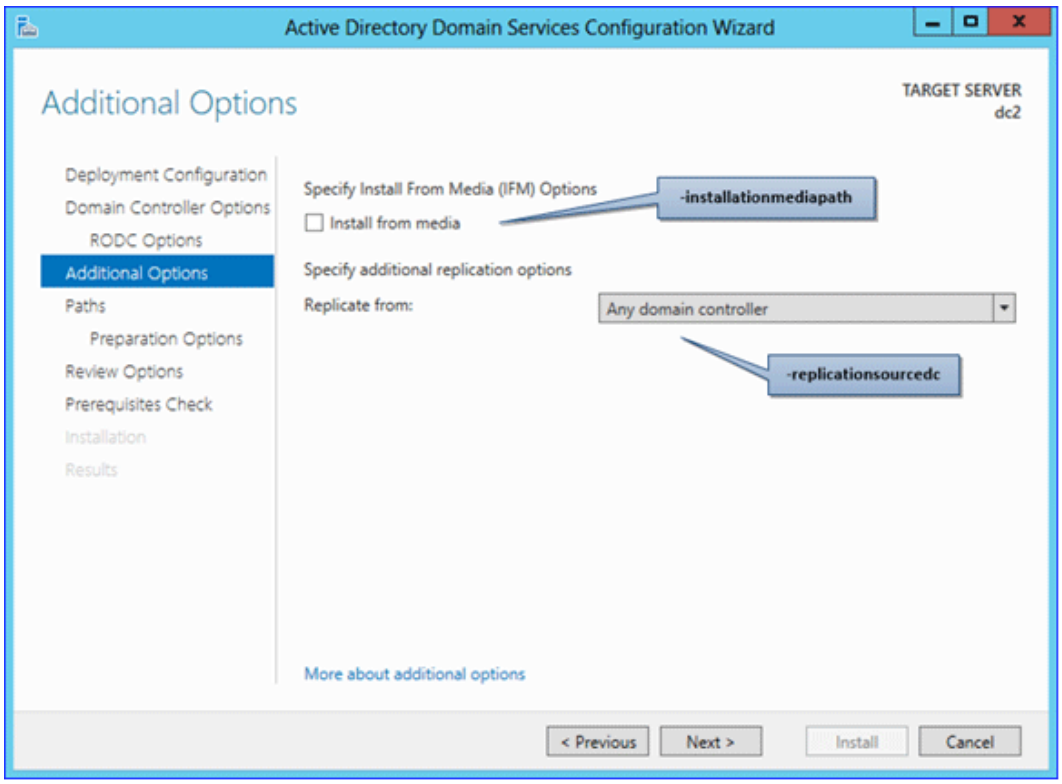

The **Additional Options** page provides configuration options to name a domain controller as the replication source, or you can use any domain controller as the replication source.

You can also choose to install the domain controller using backed up media using the Install from media (IFM) option. The **Install from media** checkbox provides a browse option once selected and you must click **Verify** to ensure the provided path is valid media. Media used by the IFM option is created with Windows Server Backup or Ntdsutil.exe from another existing Windows Server 2012 computer only; you cannot use a Windows Server 2008 R2 or previous operating system to create media for a Windows Server 2012 domain controller. The Appendices provides more information on changes in IFM. If using media protected with a SYSKEY, Server Manager prompts for the image's password during verification.

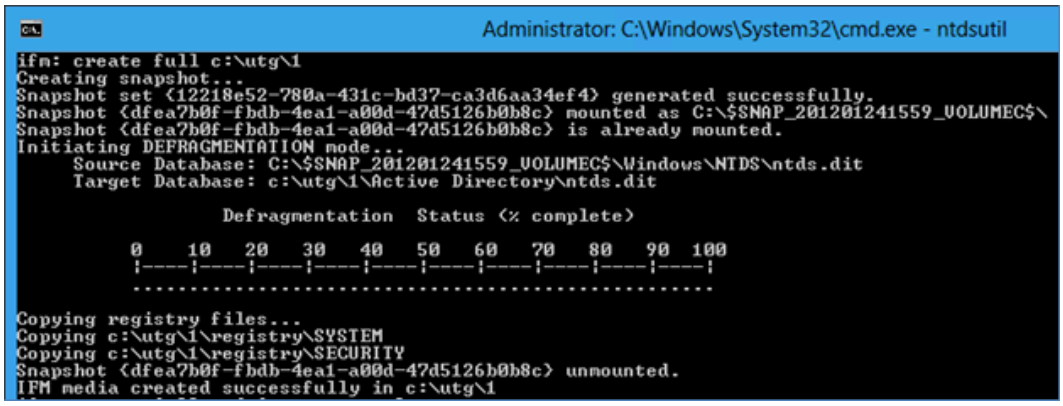

The Additional Options ADDSDeployment cmdlet arguments are:

`-replicationsourcedc <string>`

`-installationmediapath <string>`

`-systemkey <secure string>`

Paths

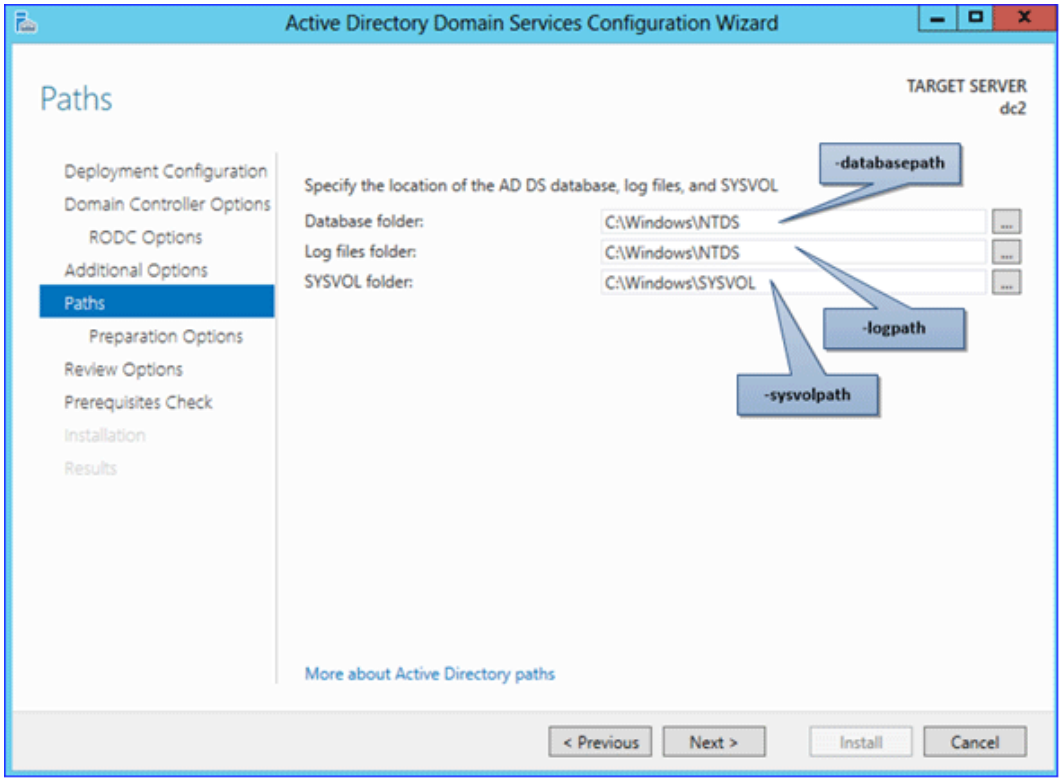

The **Paths** page enables you to override the default folder locations of the AD DS database, the database transaction logs, and the SYSVOL share. The default locations are always in subdirectories of %systemroot%. The **Paths** ADDSDeployment cmdlet arguments are:

```
-databasepath <string>

-logpath <string>

-sysvolpath <string>
```

Preparation Options

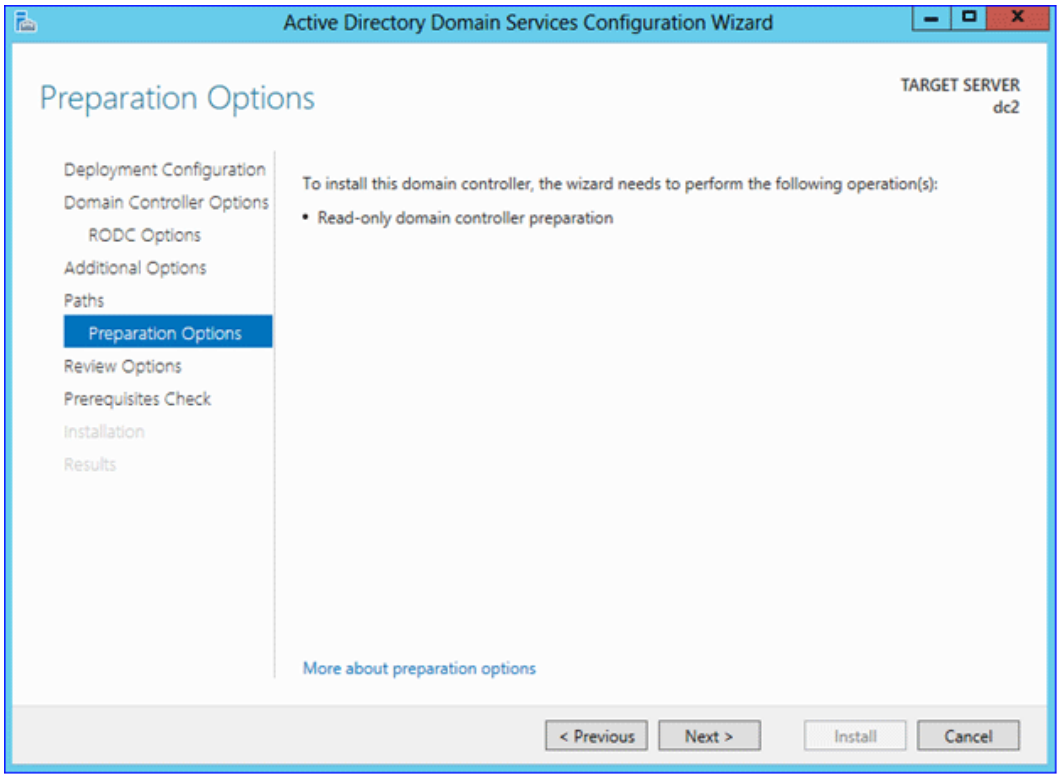

The **Preparation Options** page alerts you that the AD DS configuration includes extending the Schema (forestprep) and updating the domain (domainprep). You only see this page when the forest or domain has not been prepared by previous Windows Server 2012 domain controller installation or from manually running Adprep.exe. For example, the Active Directory Domain Services Configuration Wizard suppresses this page if you add a new replica domain controller to an existing Windows Server 2012 forest root domain.

Extending the Schema and updating the domain do not occur when you click **Next**. These events occur only during the installation phase. This page simply brings awareness about the events that will occur later in the installation.

This page also validates that the current user credentials are members of the Schema Admin and Enterprise Admins groups, as you need membership in these groups to extend the schema or

prepare a domain. Click **Change** to provide the adequate user credentials if the page informs you that the current credentials do not provide sufficient permissions.

The Additional Options ADDSDeployment cmdlet argument is:

```
-adprepcredential <pscredential>
```

Important

As with previous versions of Windows Server, Windows Server 2012's automated domain preparation does not run GPPREP. Run **adprep.exe /gpprep** manually for all domains that were not previously prepared for Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2. You should run GPPrep only once in the history of a domain, not with every upgrade. Adprep.exe does not run /gpprep automatically because its operation can cause all files and folders in the SYSVOL folder to re-replicate on all domain controllers.

Automatic RODCPrep runs when you promote the first un-staged RODC in a domain. It does not occur when you promote the first writeable Windows Server 2012 domain controller. You can also still manually run **adprep.exe /rodcprep** if you plan to deploy read-only domain controllers.

Review Options and View Script

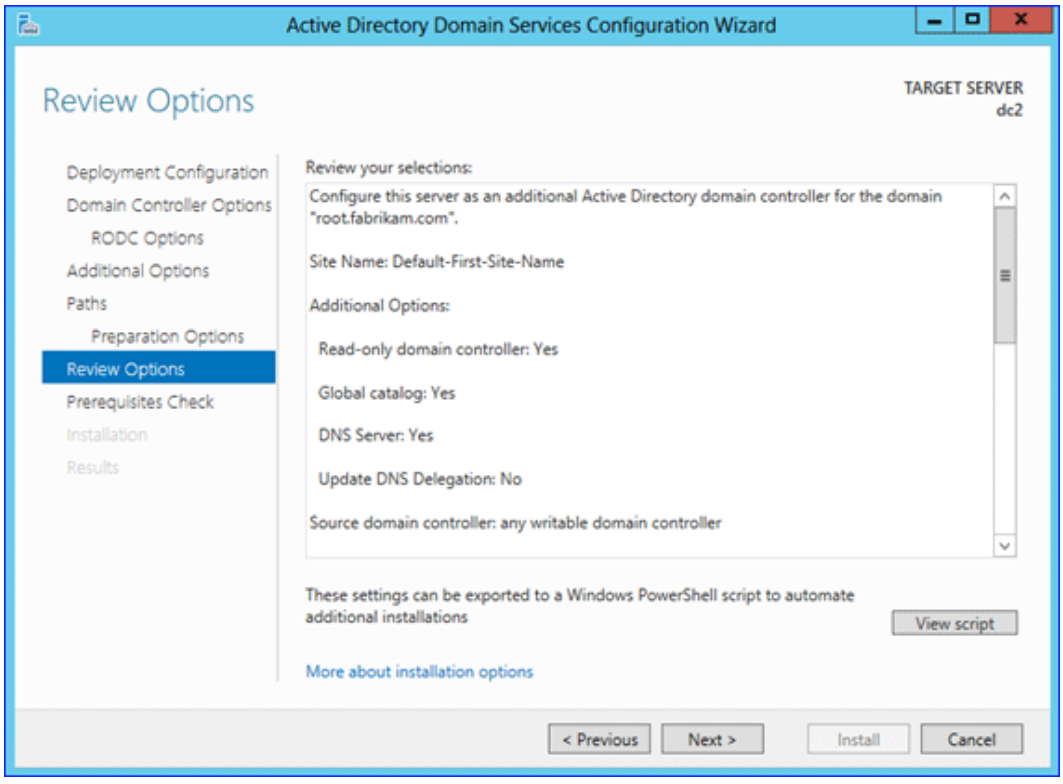

The **Review Options** page enables you to validate your settings and ensure that they meet your requirements before you start the installation. This is not the last opportunity to stop the installation using Server Manager. This page simply enables you to review and confirm your settings before continuing the configuration.

The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use. For example:

```
#
# Windows PowerShell Script for AD DS Deployment
#


Import-Module ADDSDeployment
Install-ADDSDomainController `
-AllowPasswordReplicationAccountName @("CORP\Allowed RODC Password Replication Group",
"CORP\Chicago RODC Admins", "CORP\Chicago RODC Users and Computers") `
-Credential (Get-Credential) `
-CriticalReplicationOnly:$false `
-DatabasePath "C:\Windows\NTDS" `
-DelegatedAdministratorAccountName "CORP\Chicago RODC Admins" `
-DenyPasswordReplicationAccountName @("BUILTIN\Administrators", "BUILTIN\Server
Operators", "BUILTIN\Backup Operators", "BUILTIN\Account Operators", "CORP\Denied RODC
Password Replication Group") `
-DomainName "corp.contoso.com" `
-InstallDNS:$true `
-LogPath "C:\Windows\NTDS" `
-ReadOnlyReplica:$true `
-SiteName "Default-First-Site-Name" `
-SYSVOLPath "C:\Windows\SYSVOL"
-Force:$true
```

Note

Server Manager generally fills in all arguments with values when promoting and does not rely on defaults (as they may change between future versions of Windows or service packs). The one exception to this is the **-safemodeadministratorpassword** argument. To force a confirmation prompt, omit the value when running cmdlet interactively.

Use the optional Whatif argument with the Install-ADDSDomainController cmdlet to review configuration information. This enables you to see the explicit and implicit values of the arguments for a cmdlet.

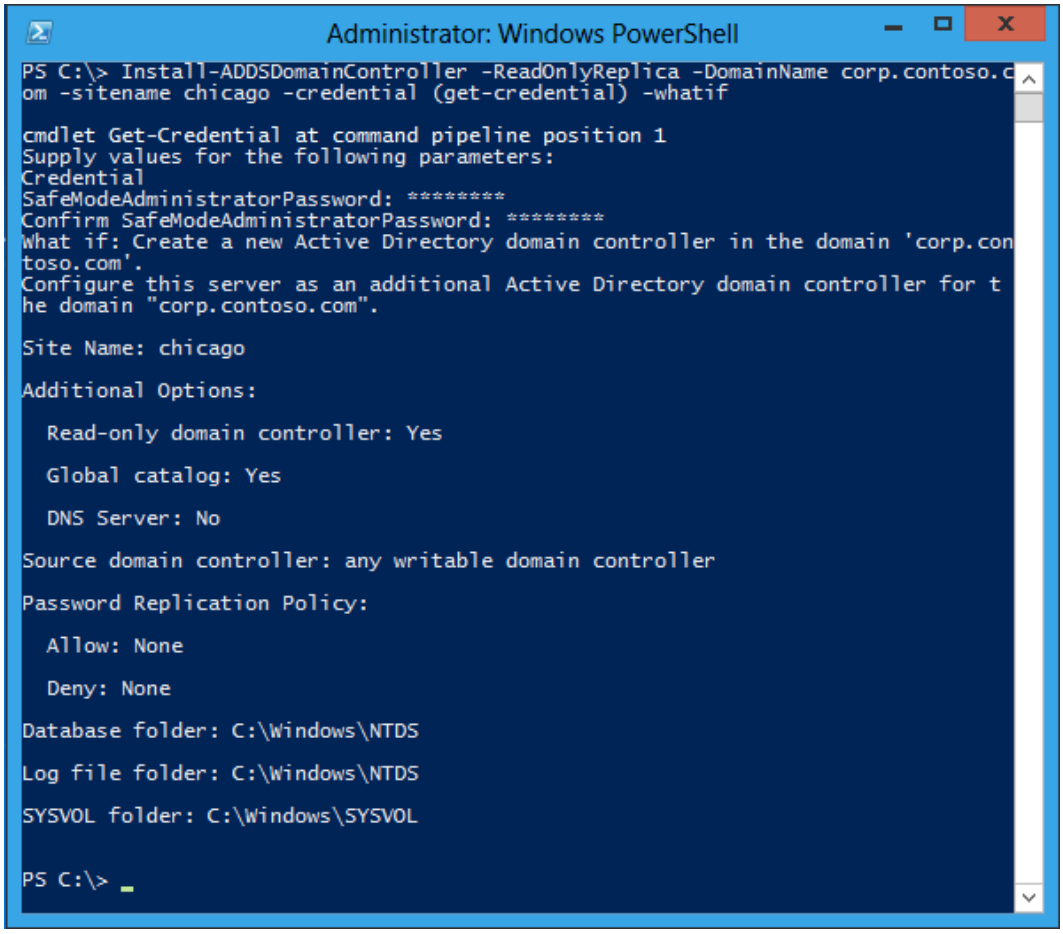

Prerequisites Check

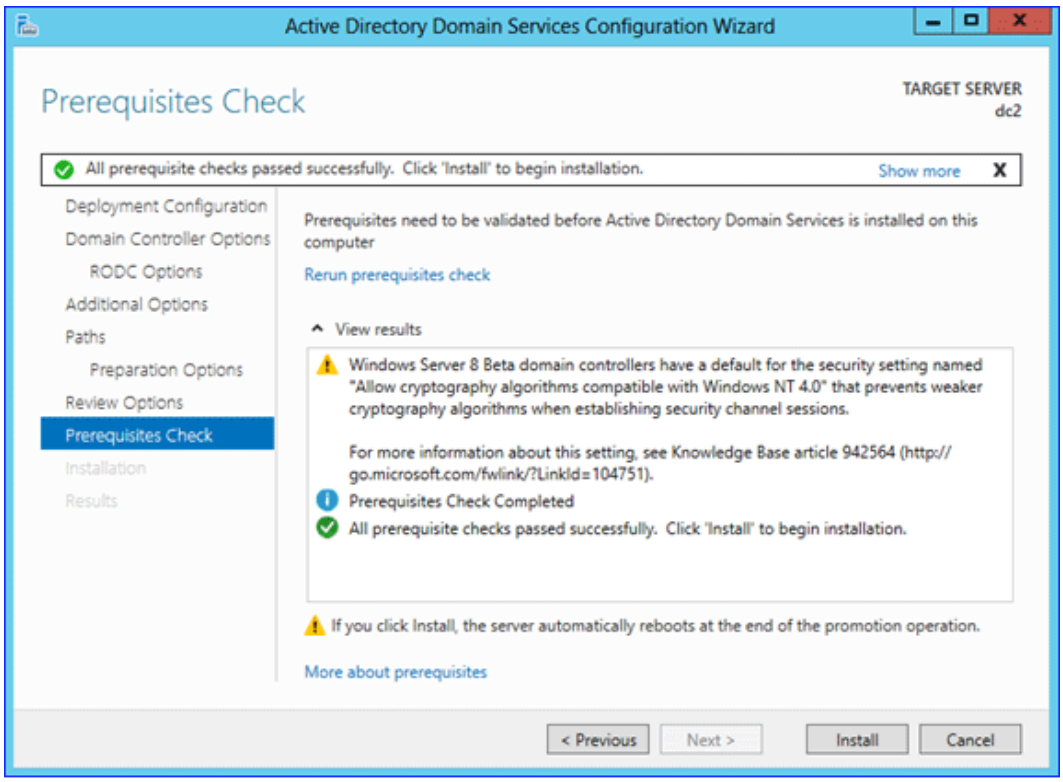

The **Prerequisites Check** is a new feature in AD DS domain configuration. This new phase validates that the server configuration is capable of supporting a new AD DS forest.

When installing a new forest root domain, the Server Manager Active Directory Domain Services Configuration Wizard invokes a series of serialized modular tests. These tests alert you with suggested repair options. You can run the tests as many times as required. The domain controller process cannot continue until all prerequisite tests pass.

The **Prerequisites Check** also surfaces relevant information such as security changes that affect older operating systems.

You cannot bypass the **Prerequisite Check** when using Server Manager, but you can skip the process when using the AD DS Deployment cmdlet using the following argument:

```
-skipprechecks
```

Click **Install** to begin the domain controller promotion process. This is last opportunity to cancel the installation. You cannot cancel the promotion process once it begins. The computer will reboot automatically at the end of promotion, regardless of the promotion results.

Installation

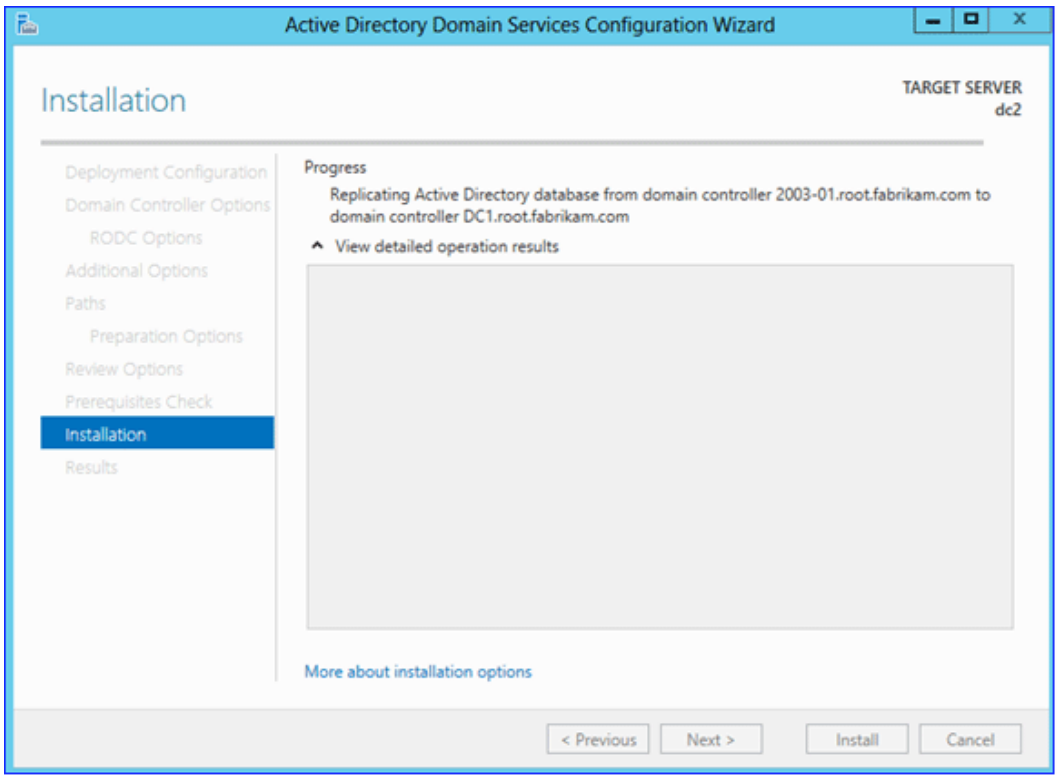

When the **Installation** page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and are written to logs:

- %systemroot%\debug\dcpromo.log
- %systemroot%\debug\dcpromoui.log

To install a new Active Directory forest using the ADDSDeployment module, use the following cmdlet:

```
Install-addsdomaincontroller
```

See the **ADDSDeployment Cmdlet** table at the begininng of this section for required and optional arguments.

The **Install-addsdomaincontroller** cmdlet only has two phases (prerequisite checking and installation). The two figures below show the installation phase with the minimum required arguments of **-domainname**, **-readonlyreplica**, **-sitename**, and **-credential**. Note how, just like Server Manager, **Install-ADDSDomainController** reminds you that promotion will reboot the server automatically:

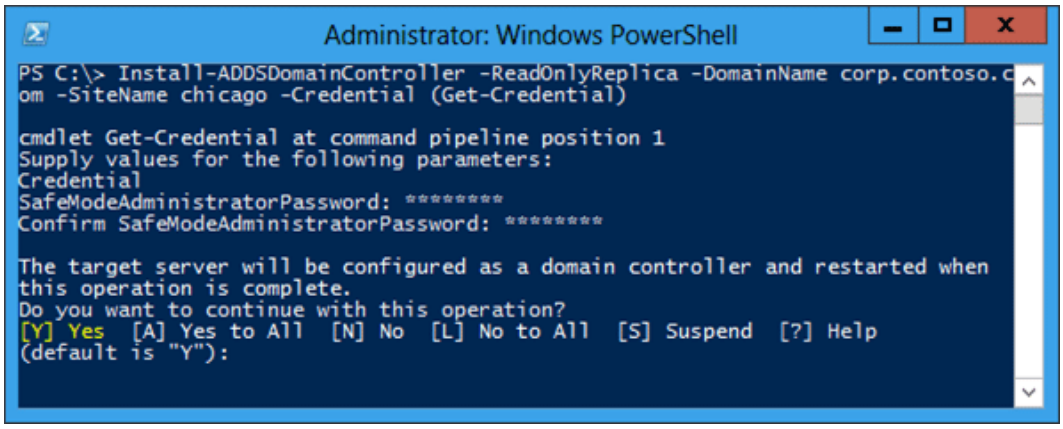

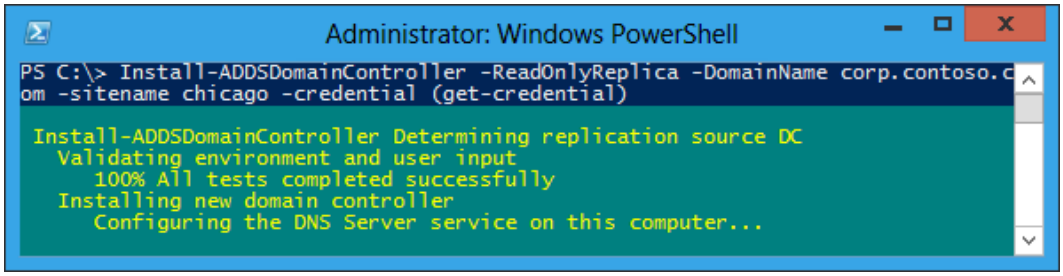

To accept the reboot prompt automatically, use the **-force** or **-confirm:$false** arguments with any ADDSDeployment Windows PowerShell cmdlet. To prevent the server from automatically rebooting at the end of promotion, use the **-norebootoncompletion** argument.

Warning

Overriding the reboot is not recommended. The domain controller must reboot to function correctly. If you log off the domain controller, you cannot log back on interactively until you restart it.

Results

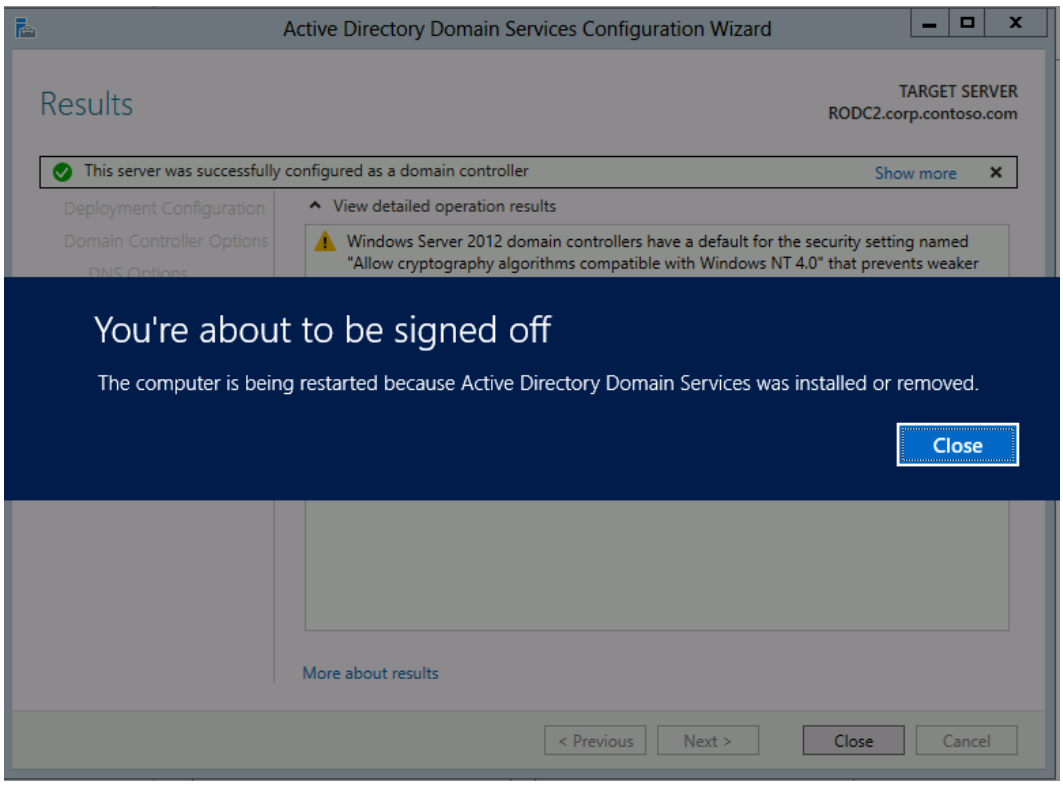

The **Results** page shows the success or failure of the promotion and any important administrative information. The domain controller will automatically reboot after 10 seconds.

Remove Active Directory Domain Services (Level 100)

This topic covers procedures for removing Active Directory Domain Services (AD DS). You can remove AD DS using these methods:

- Remove AD DS using Windows PowerShell
- Remove AD DS using the Remove Roles Wizard in Server Manager

After AD DS is removed, the DNS server role remains installed and running if it was previously installed on the domain controller. But any Active Directory–integrated DNS zones that were installed are removed. By default, the AD DS removal process also attempts to remove the Domain Name System (DNS) delegations for the zones that point to the domain controller.

If the DNS server no longer serves any purpose after you remove AD DS, use Remove Roles Wizard to remove DNS server role. If you remove the DNS server role, you must reconfigure any DHCP scopes and DNS clients that resolved against this DNS server to use a suitable alternative (typically, another DNS server running on a domain controller within the same domain).

If you are removing the last domain controller in a domain, the wizard displays a list of all the application directory partitions that are stored on the domain controller. If the application directory partitions were created by an application other than AD DS, first try to use an appropriate tool that is provided by the application to remove these directory partitions. If the application does not provide such a tool, you can let the wizard remove the directory partitions. Application directory partitions that are created by AD DS, such as the DomainDNSZones application directory partition, cannot be retained if you remove AD DS.

Remove AD DS using Windows PowerShell

Windows PowerShell allows you to script the removal of AD DS. If you plan to demote a domain controller temporarily, you do not need to remove and then reinstall the AD DS server role.

To view the syntax and parameter options for removing AD DS using the ADDSDeployment module in Windows PowerShell, type the following command:

```
get-help Uninstall-ADDSDomainController
```

The command syntax for removing AD DS is as follows. Optional parameters appear within brackets.

```
Uninstall-ADDSDomainController [-skipprechecks] -LocalAdministratorPassword
<SecureString> [-Credential <PS Credential>] [-DemoteOperationsMasterRole] [-
DNSDelegationRemovalCredential <PS Credential>] [-ForceRemoval] [-
IgnoreLastDCInDomainMismatch] [-IgnoreLastDNSServerForZone] [-
LastDomainControllerInDomain] [-NoRebootOnCompletion] [-RemoveApplicationPartitions] [-
RemoveDNSDelegation] [-RetainDCMetadata] [-Force] [-WhatIf] [-Confirm] [<common
parameters>]

Uninstall-ADDSDomainController -ForceRemoval [-skipprechecks] -LocalAdministratorPassword
<SecureString> [-Credential <PS Credential>] [-DemoteOperationsMasterRole] [-
NoRebootOnCompletion] [-Force] [-WhatIf] [-Confirm] [<common parameters>]
```

Tip

Normally on demotion, any warnings with high importance will require explicit acknowledgment of the user. To ensure that those warnings do not wait for user acknowledgment, specify –force option. The UI always passes the –Force flag.

The –Forceremoval option is typically used to remove AD DS when the domain controller has no connectivity with other domain controllers.

For example, to remove AD DS from an additional domain controller in a domain and be prompted to set the local Administrator password, type the following command:

Note

You will be prompted to supply the local Administrator password.

```
Uninstall-ADDSDomainController
```

To remove AD DS from an additional domain controller in a domain and be prompted to set the local Administrator password but not prompted to confirm the command, type the following command:

```
Uninstall-ADDSDomainController -confirm:$false
```

Here is an example of forcibly demoting with its minimal required arguments of **-forceremoval** and **-demoteoperationmasterrole**. The **-credential** argument is not required because the user logged on as a member of the Enterprise Admins group:

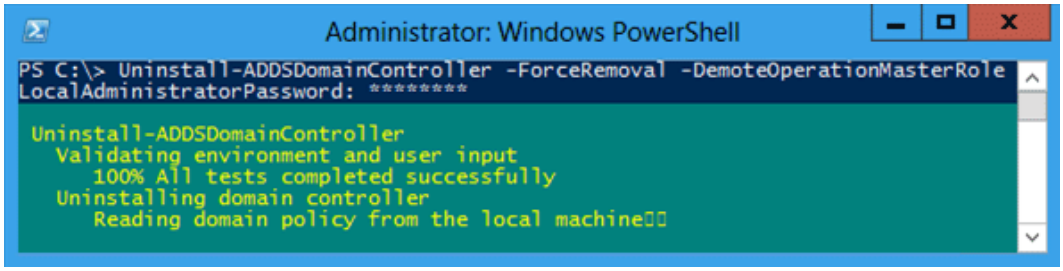

Here is an example of removing the last domain controller in the domain with its minimal required arguments of **-lastdomaincontrollerindomain** and **–removeapplicationpartitions**:

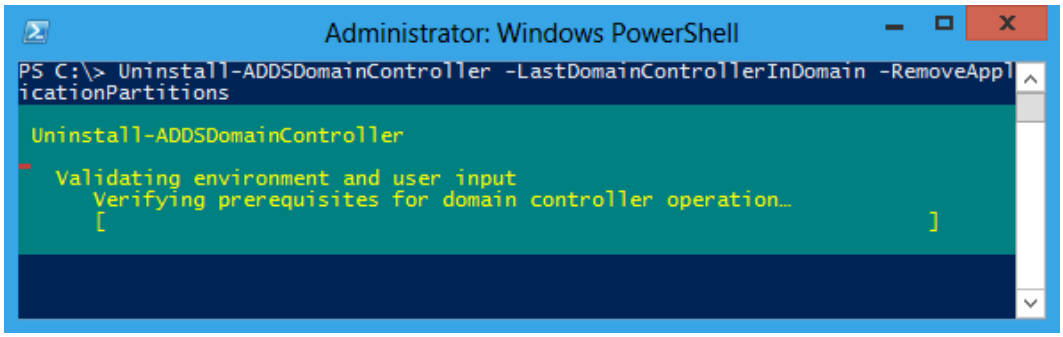

Using Test-ADDSDomainControllerUninstallation cmdlet

The Test-ADDSDomainControllerUninstallation cmdlet runs those prerequisite checks (only) which would be performed if you were to use the Uninstall-ADDSDomainController cmdlet to uninstall a domain controller in Active Directory. It differs from using the **-WhatIf** parameter with the Uninstall-ADDSDomainController cmdlet in that instead of summarizing the changes that would occur during the uninstallation process, this cmdlet actually tests whether those changes are possible given the current environment.

For more information on the scope of these prerequisite checks that the ADDSDeployment module performs when using this cmdlet, see Prerequisite Checking.

The syntax for the Test-ADDSDomainControllerUninstallation cmdlet is:

```
Test-ADDSDomainControllerUninstallation [-Credential <PSCredential>] [-
DemoteOperationMasterRole] [-DnsDelegationRemovalCredential <PSCredential>] [-Force] [-
IgnoreLastDCInDomainMismatch] [-IgnoreLastDnsServerForZone] [-
```

```
LastDomainControllerInDomain] [-LocalAdministratorPassword <SecureString>] [-
NoRebootOnCompletion] [-RemoveApplicationPartitions] [-RemoveDnsDelegation] [-
RetainDCMetadata] [<CommonParameters>]


Test-ADDSDomainControllerUninstallation -ForceRemoval [-Credential <PSCredential>] [-
DemoteOperationMasterRole] [-Force] [-LocalAdministratorPassword <SecureString>] [-
NoRebootOnCompletion] [<CommonParameters>]
```

Remove AD DS using the Remove Roles Wizard in Server Manager

The following procedure shows the steps to remove AD DS using the Remove Roles Wizard in Server Manager.

Administrative credentials

To remove a domain controller, you must be a member of the Domain Admins group in the domain. To remove the last domain controller in a domain or forest, you must be a member of the Enterprise Admins group.

Note

If you force the removal of AD DS while the domain controller is started in Directory Services Restore Mode (DSRM), credentials are optional.

To Remove AD DS using the Remove Roles Wizard

1. In Server Manager, click **Manage**, and then click **Remove Roles and Features**.
2. On the **Before you begin** page, review the information and then click **Next**.
3. On the **Select destination server** page, click the name of the server that you want to remove AD DS from and then click **Next**.
4. On the **Remove server roles** page, clear the check box for **Active Directory Domain Services** and any other server roles that you want to remove, such as **DNS Server**, and then on the **Remove Roles and Features Wizard** dialog box, click **Remove Features**, and then click **Next**.
5. The Remove Roles and Features Wizard returns the following validation error:

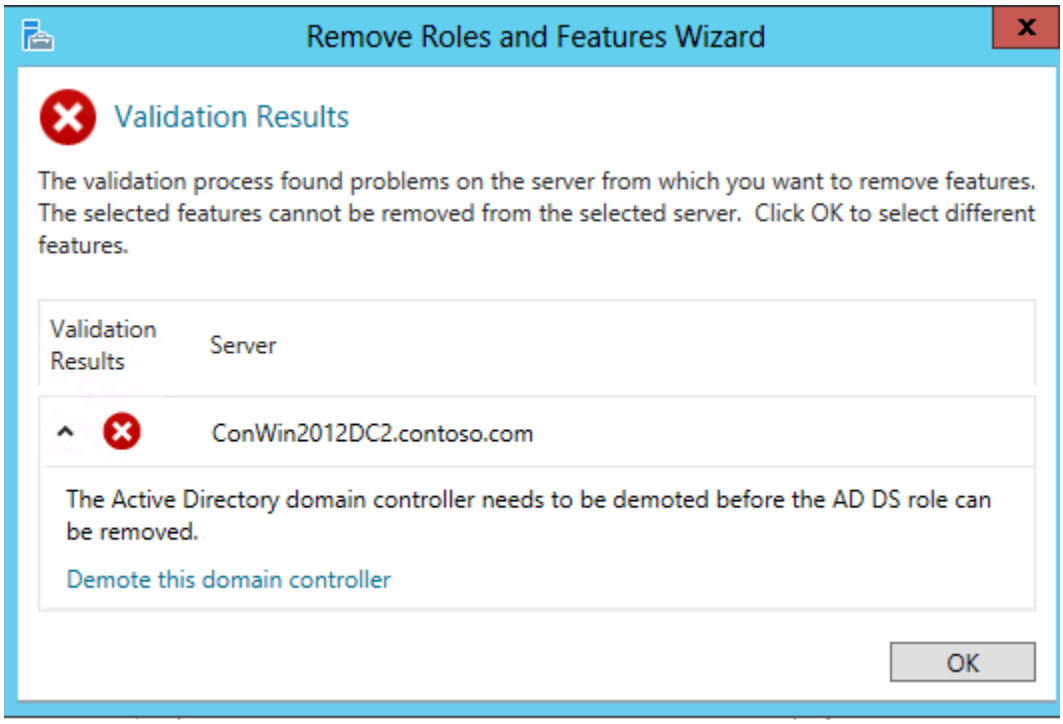

The validation error appears by design because the AD DS server role binaries cannot be removed while the server is running as a domain controller. Click **Demote this domain controller**.

6. On the **Credentials** page, specify credentials to remove AD DS. If previous attempts to remove AD DS on this domain controller have failed, then you can select the **Force the removal of this domain controller** check box. If you are removing the last domain controller in the domain, click **Last domain controller in the domain** check box. Click **Next**.
7. On the **Warnings** page, review the information about the roles hosted by the domain controller, click **Proceed with removal**, and then click **Next**.
8. On the **Removal Options** page:

Note
This page does not appear if you chose **Force the removal of this domain controller**.

- If you plan to reinstall the domain controller using the same domain controller account, click **Retain the domain controller metadata**.

In addition, if either of the following two options appears, it must be selected before you can proceed.

- If you are removing the last DNS server that hosts the zones hosted on this domain controller, click **Remove this DNS zone (this is the last server that hosts the zone)**.

Important
This option will be shown only if the domain controller is the last DNS server for the zone.

- If you want to delete the application partitions, click **Remove application partitions**.

Click **Next**.

9. On the **New Administrator Password** page, type and confirm the password for the local Administrator account for the server, and then click **Next**.
10. On the **Review Options** page, click **Demote**.

    The server will restart automatically to complete the domain controller demotion. Continue with the next steps, which are needed to fully remove the AD DS server role binaries after the machine restarts to complete the demotion.
11. In Server Manager, click **Manage**, and then click **Remove Roles and Features**.
12. On the **Before you begin** page, review the information and then click **Next**.
13. On the **Select destination server** page, click the name of the server that you want to remove AD DS from and then click **Next**.
14. On the **Remove server roles** page, clear the check box for **Active Directory Domain Services** and any other server roles that you want to remove, such as **DNS Server**. On the **Remove Roles and Features Wizard** dialog box, clear the **Remove management tools (if applicable)** check box if you plan to administer another domain controller from this server, and click **Remove Features**, and then click **Next**.
15. On the **Remove features** page, clear the check box for any feature you want to remove and then click **Next**.
16. On the **Confirm removal selections** page, click **Remove**.
17. After the role is removed, click **Close**, and restart the server.

See Also

Demoting Domain Controllers and Domains (Level 200)

Demoting Domain Controllers and Domains (Level 200)

This topic explains how to remove AD DS, using Server Manager or Windows PowerShell.

- AD DS Removal Workflow
- Demotion and Role Removal Windows PowerShell
- Demote

AD DS Removal Workflow

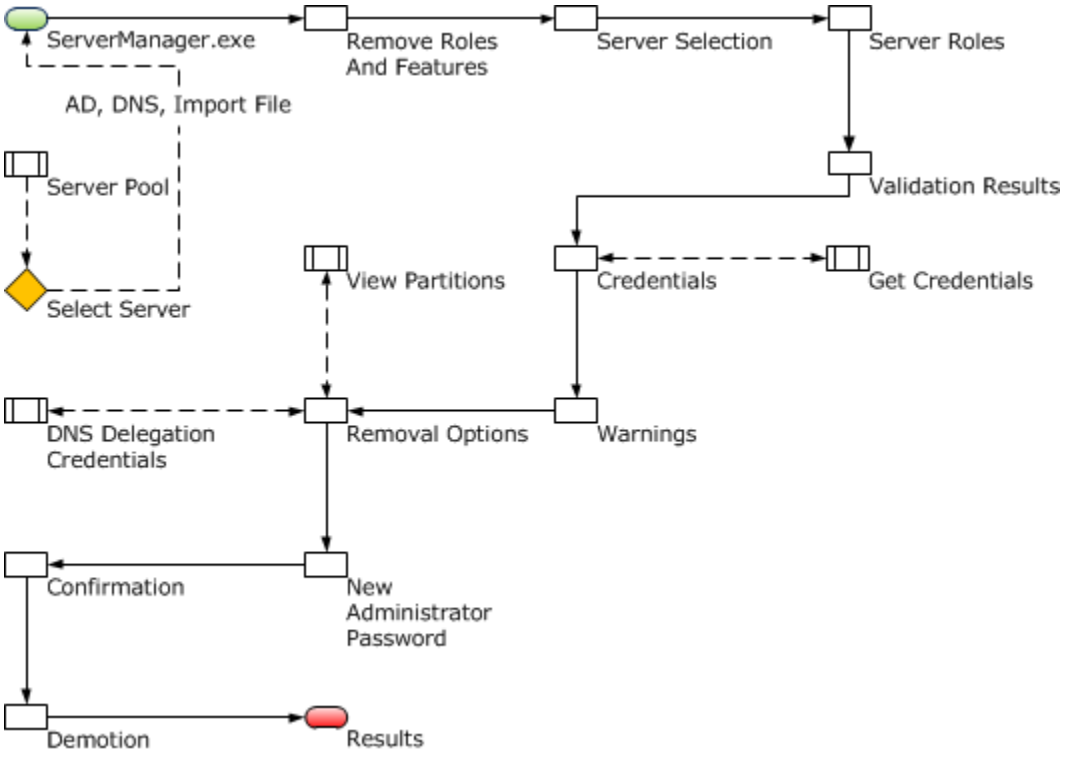

Caution

Removing the AD DS roles with Dism.exe or the Windows PowerShell DISM module after promotion to a Domain Controller is not supported and will prevent the server from booting normally.

Unlike Server Manager or the ADDSDeployment module for Windows PowerShell, DISM is a native servicing system that has no inherent knowledge of AD DS or its configuration. Do not use Dism.exe or the Windows PowerShell DISM module to uninstall the AD DS role unless the server is no longer a domain controller.

Demotion and Role Removal Windows PowerShell

ADDSDeployment and ServerManager Cmdlets	Arguments (**Bold** arguments are required. *Italicized* arguments can be specified by using Windows PowerShell or the AD DS Configuration Wizard.)
Uninstall-AddsDomainController	-SkipPreChecks *-LocalAdministratorPassword* -Confirm

	-Credential -DemoteOperationMasterRole *-DNSDelegationRemovalCredential* -Force *-ForceRemoval* *-IgnoreLastDCInDomainMismatch* *-IgnoreLastDNSServerForZone* *-LastDomainControllerInDomain* -Norebootoncompletion *-RemoveApplicationPartitions* *-RemoveDNSDelegation* -RetainDCMetadata
Uninstall-WindowsFeature/Remove-WindowsFeature	***-Name*** ***-IncludeManagementTools*** *-Restart* -Remove -Force -ComputerName -Credential -LogPath -Vhd

Note

The **-credential** argument is only required if you are not already logged on as a member of the Enterprise Admins group (demoting last DC in a domain) or the Domain Admins group (demoting a replica DC).The **-includemanagementtools** argument is only required if you want to remove all of the AD DS management utilities.

Demote

Remove Roles and Features

Server Manager offers two interfaces to removing the Active Directory Domain Services role:

- The **Manage** menu on the main dashboard, using **Remove Roles and Features**

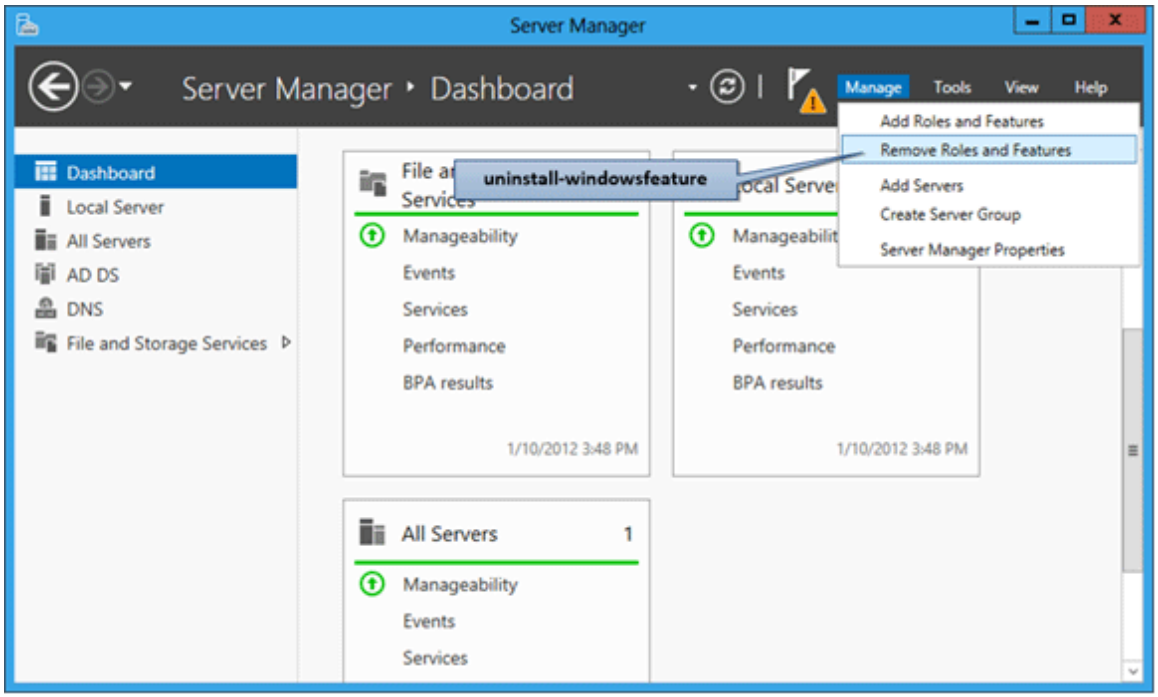

- Click **AD DS** or **All Servers** on the navigation pane. Scroll down to the **Roles and Features** section. Right-click **Active Directory Domain Services** in the **Roles and Features** list and click **Remove Role or Feature**. This interface skips the **Server Selection** page.

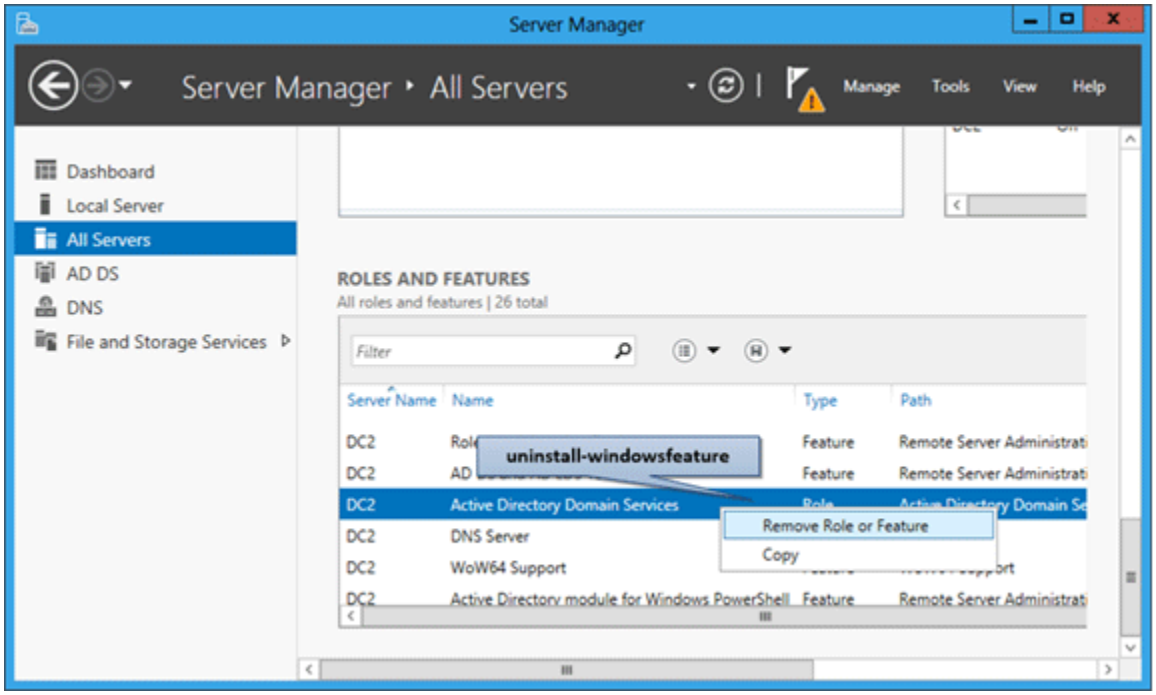

The ServerManager cmdlets **Uninstall-WindowsFeature** and **Remove-WindowsFeature** prevent you from removing the AD DS role until you demote the domain controller.

Server Selection

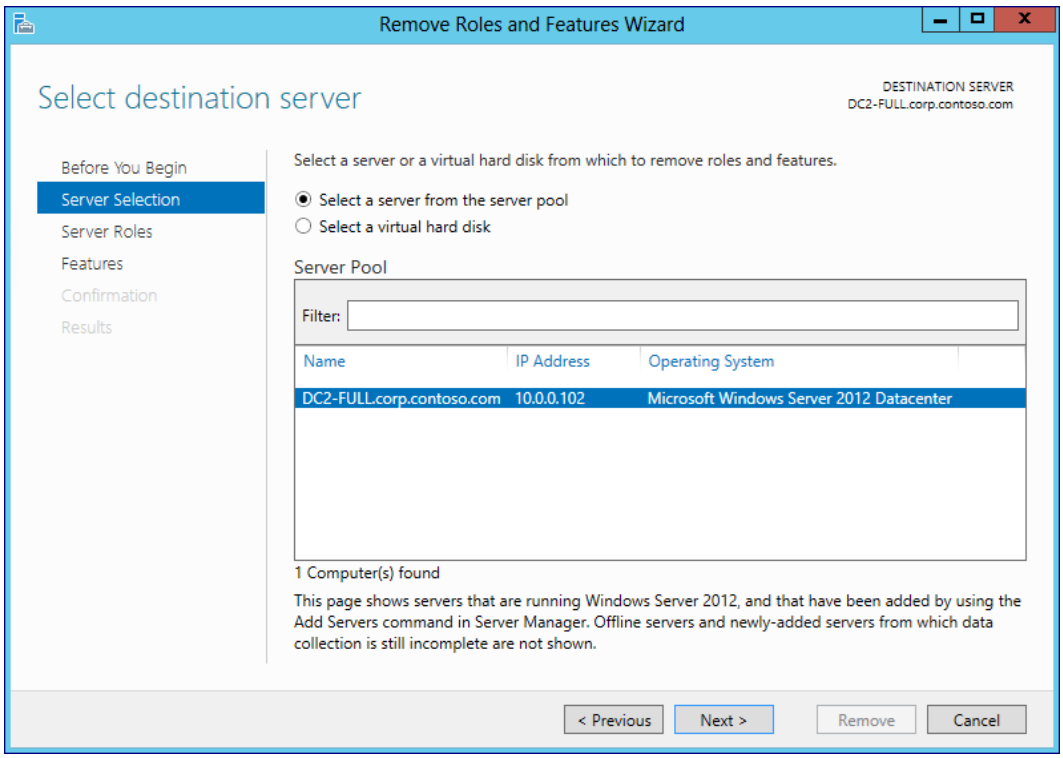

The **Server Selection** dialog enables you to choose from one of the servers previously added to the pool, as long as it is accessible. The local server running Server Manager is always automatically available.

Server Roles and Features

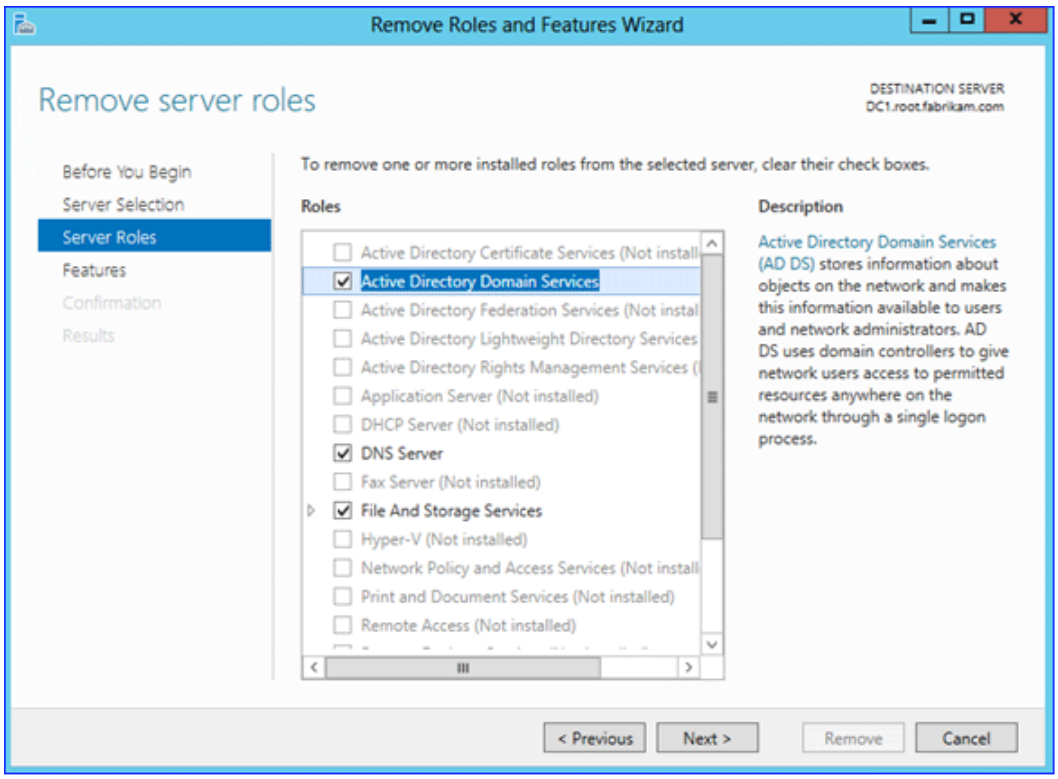

Clear the **Active Directory Domain Services** check box to demote a domain controller; if the server is currently a domain controller, this does not remove the AD DS role and instead switches to a **Validation Results** dialog with the offer to demote. Otherwise, it simply removes the binaries like any other role feature.

- Do not remove any other AD DS-related roles or features - such as DNS, GPMC, or the RSAT tools - if you intend to promote the domain controller again immediately. Removing additional roles and feature increases the time to re-promote, as Server Manager reinstalls these features when you reinstall the role.
- Remove unneeded AD DS roles and features at your own discretion if you intend to demote the domain controller permanently. This requires clearing the check boxes for those roles and features.

  The full list of AD DS-related roles and features include:

  - Active Directory Module for Windows PowerShell feature
  - AD DS and AD LDS Tools feature
  - Active Directory Administrative Center feature
  - AD DS Snap-ins and Command-line Tools feature
  - DNS Server
  - Group Policy Management Console

The equivalent ADDSDeployment and ServerManager Windows PowerShell cmdlets are:

```
Uninstall-addsdomaincontroller

Uninstall-windowsfeature
```

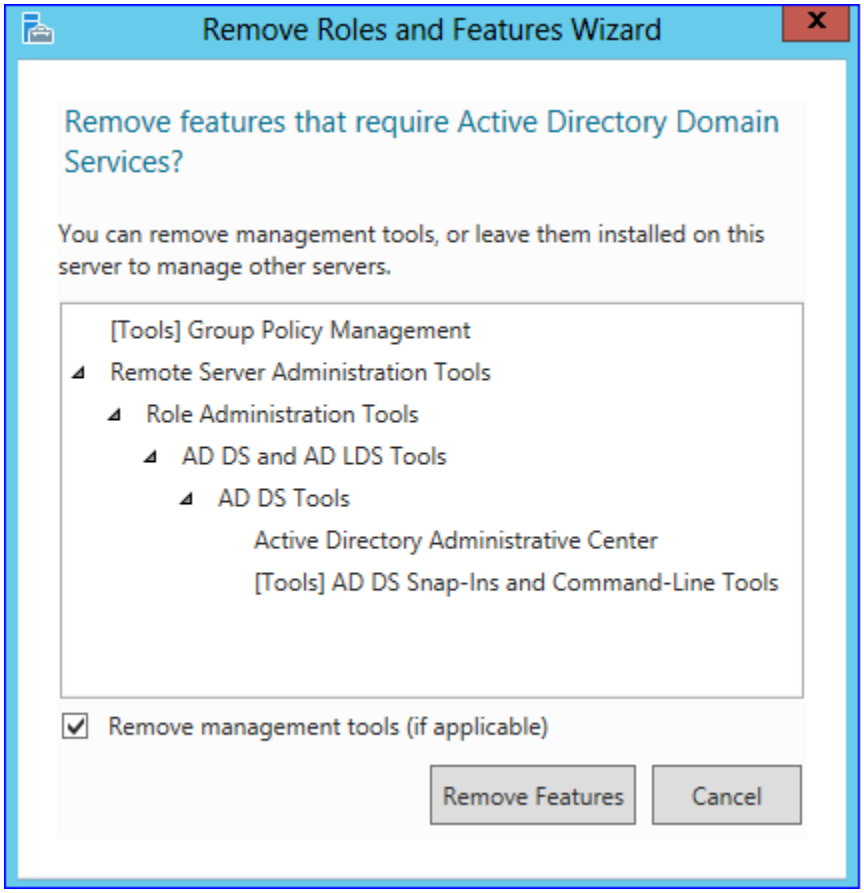

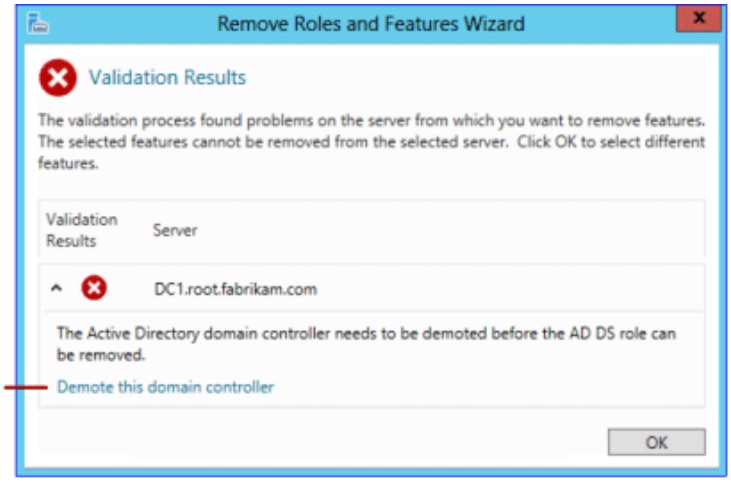

Uninstall-addsdomaincontroller

Credentials

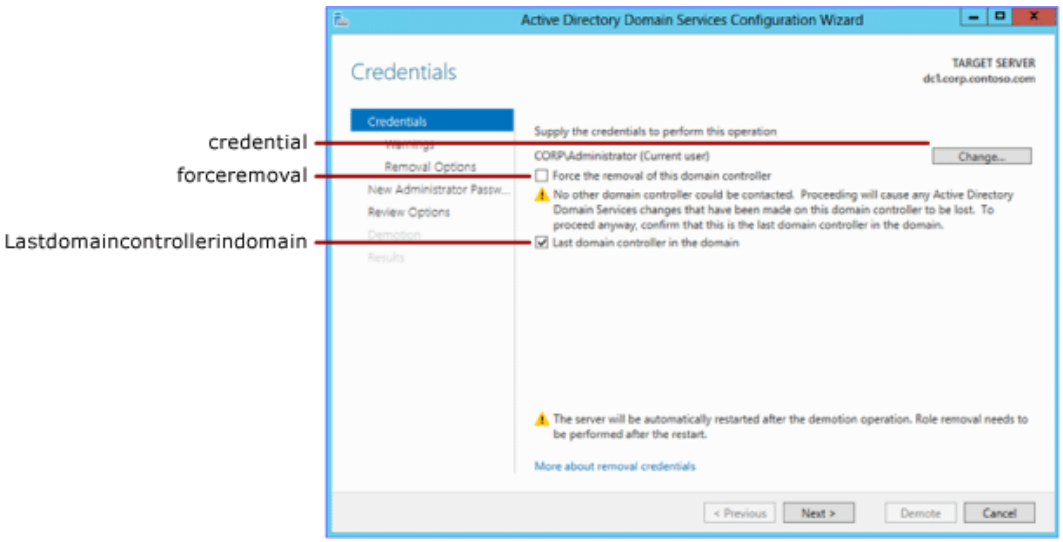

You configure demotion options on the **Credentials** page. Provide the credentials necessary to perform the demotion from the following list:

- Demoting an additional domain controller requires Domain Admin credentials. Selecting **Force removal of the domain controller** demotes the domain controller without removing the domain controller object's metadata from Active Directory.

  Warning

  Do not select this option unless the domain controller cannot contact other domain controllers and there is *no reasonable way* to resolve that network issue. Forced demotion leaves orphaned metadata in Active Directory on the remaining domain controllers in the forest. In addition, all un-replicated changes on that domain controller, such as passwords or new user accounts, are lost forever. Orphaned metadata is the root cause in a significant percentage of Microsoft Customer Support cases for AD DS, Exchange, SQL, and other software.

  If you forcibly demote a domain controller, you *must* manually perform metadata cleanup immediately. For steps, review Clean Up Server Metadata.

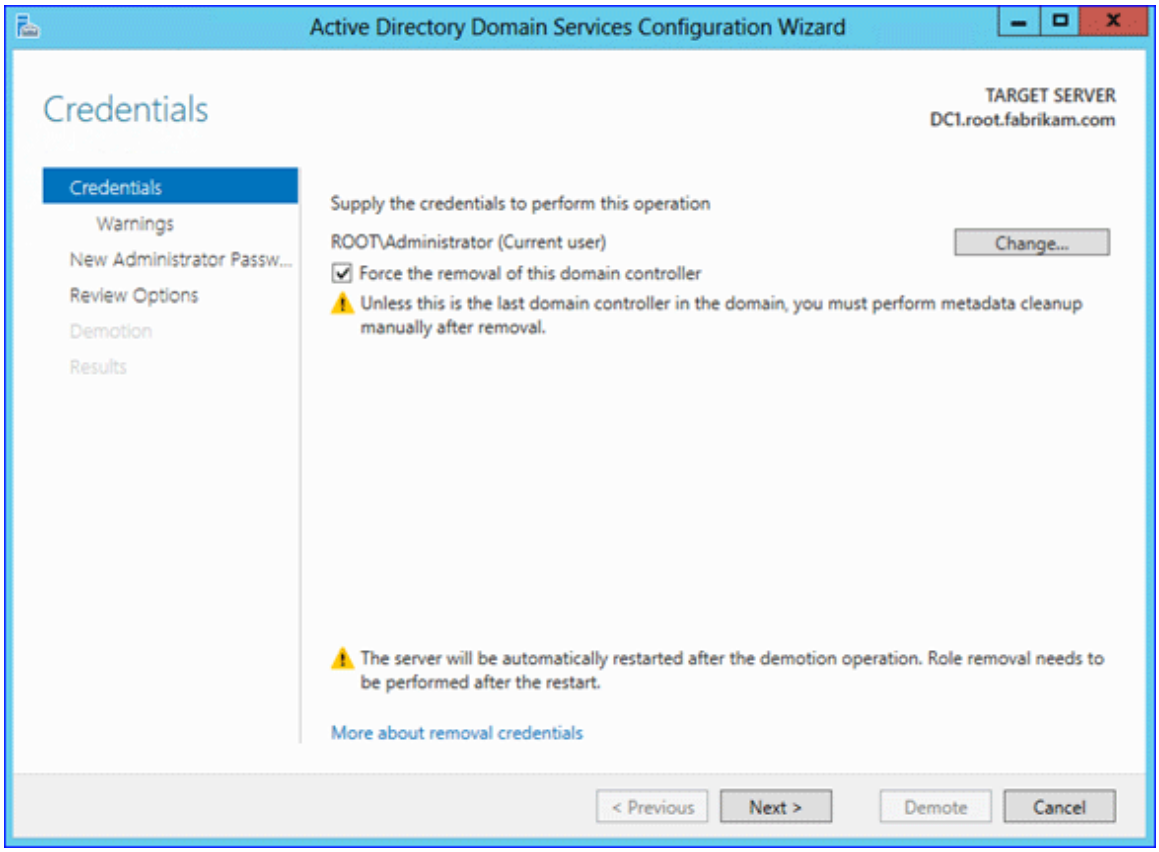

- Demoting the last domain controller in a domain requires Enterprise Admins group membership, as this removes the domain itself (if the last domain in the forest, this removes the forest). Server Manager informs you if the current domain controller is the last domain controller in the domain. Select the **Last domain controller in the domain** check box to confirm the domain controller is the last domain controller in the domain.

The equivalent ADDSDeployment Windows PowerShell arguments are:

```
-credential <pscredential>

-forceremoval <{ $true | false }>

-lastdomaincontrollerindomain <{ $true | false }>
```

Warnings

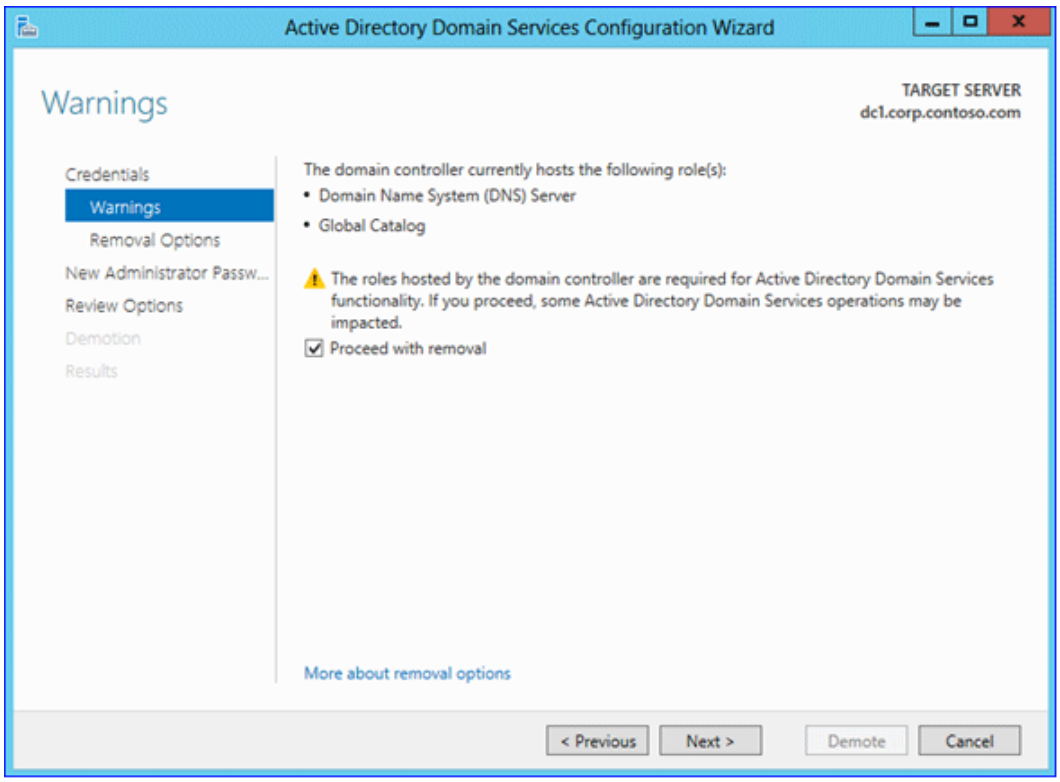

The **Warnings** page alerts you to the possible consequences of removing this domain controller. To continue, you must select **Proceed with removal**.

Warning

If you previously selected **Force the removal of this domain controller** on the **Credentials** page, then the **Warnings** page shows all Flexible Single Master Operations roles hosted by this domain controller. You *must* seize the roles from another domain controller *immediately* after demoting this server. For more information on seizing FSMO roles, see Seize the Operations Master Role.

This page does not have an equivalent ADDSDeployment Windows PowerShell argument.

Removal Options

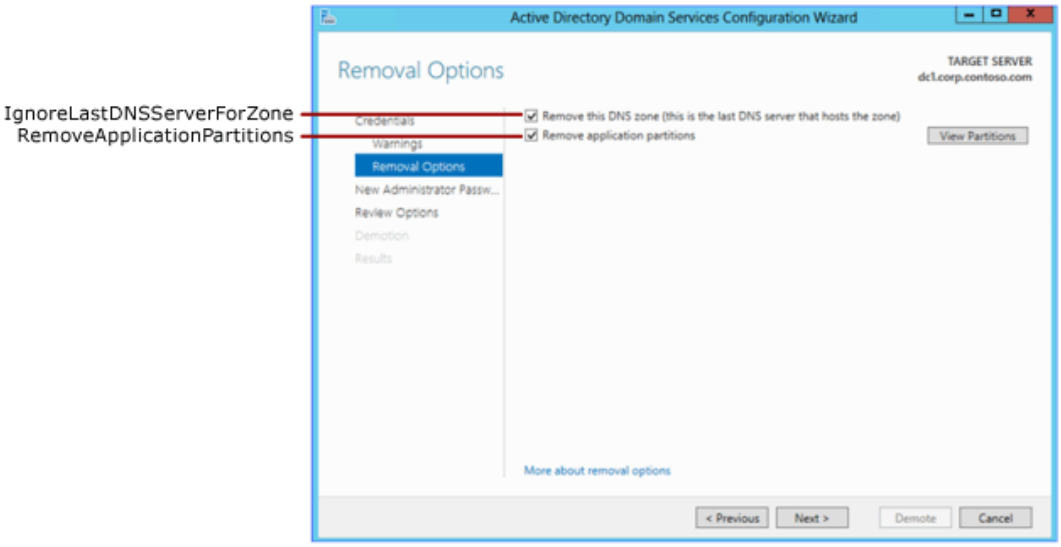

The **Removal Options** page appears depending on previously selecting **Last domain controller in the domain** on the **Credentials** page. This page enables you to configure additional removal options. Select **Ignore last DNS server for zone**, **Remove application partitions**, and **Remove DNS Delegation** to expose the **Next** button.

The options only appear if applicable to this domain controller. For instance, if there is no DNS delegation for this server then that checkbox will not display.

Click **Change** to specify alternate DNS administrative credentials. Click **View Partitions** to view additional partitions the wizard removes during the demotion. By default, the only additional partitions are Domain DNS and Forest DNS Zones. All other partitions are non-Windows partitions.

The equivalent ADDSDeployment cmdlet arguments are:

```
-ignorelastdnsserverforzone <{ $true | false }>

-removeapplicationpartitions <{ $true | false }>

-removednsdelegation <{ $true | false }>

-dnsdelegationremovalcredential <pscredential>
```

New Administrator Password

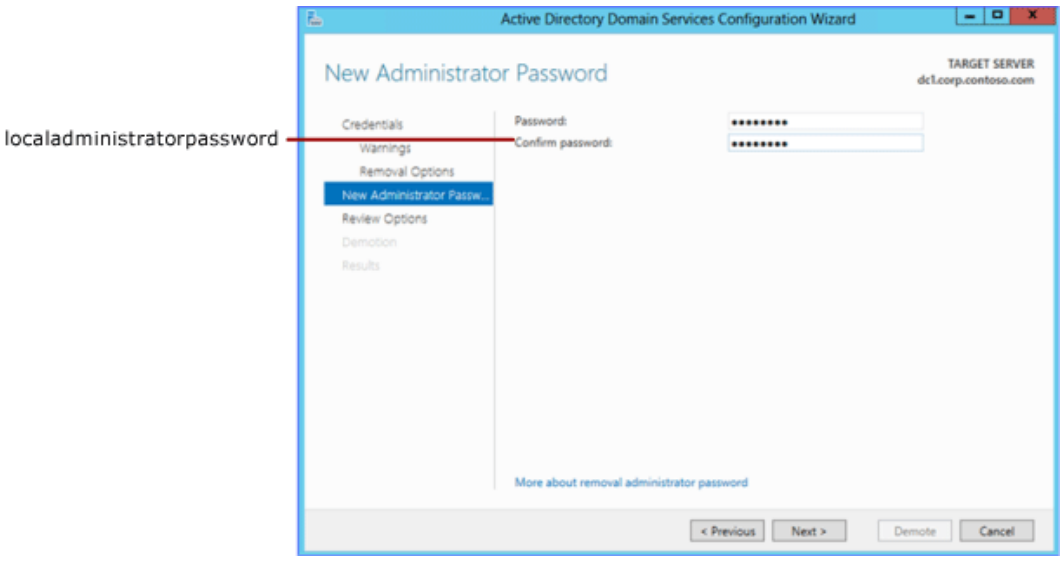

The **New Administrator Password** page requires you to provide a password for the built-in local computer's Administrator account, once the demotion completes and the computer becomes a domain member server or workgroup computer.

The **Uninstall-ADDSDomainController** cmdlet and arguments follow the same defaults as Server Manager if not specified.

The **LocalAdministratorPassword** argument is special:

- If *not specified* as an argument, then the cmdlet prompts you to enter and confirm a masked password. This is the preferred usage when running the cmdlet interactively
- If specified *with a value*, then the value must be a secure string. This is not the preferred usage when running the cmdlet interactively

For example, you can manually prompt for a password by using the **Read-Host** cmdlet to prompt the user for a secure string

```
-localadministratorpassword (read-host -prompt "Password:" -assecurestring)
```

Warning

As the previous two options do not confirm the password, use extreme caution: the password is not visible

You can also provide a secure string as a converted clear-text variable, although this is highly discouraged. For example:

```
-localadministratorpassword (convertto-securestring "Password1" -asplaintext -force)
```

Warning

Providing or storing a clear text password is not recommended. Anyone running this command in a script or looking over your shoulder knows the local administrator password of that computer. With that knowledge, they have access to all of its data and can impersonate the server itself.

Confirmation

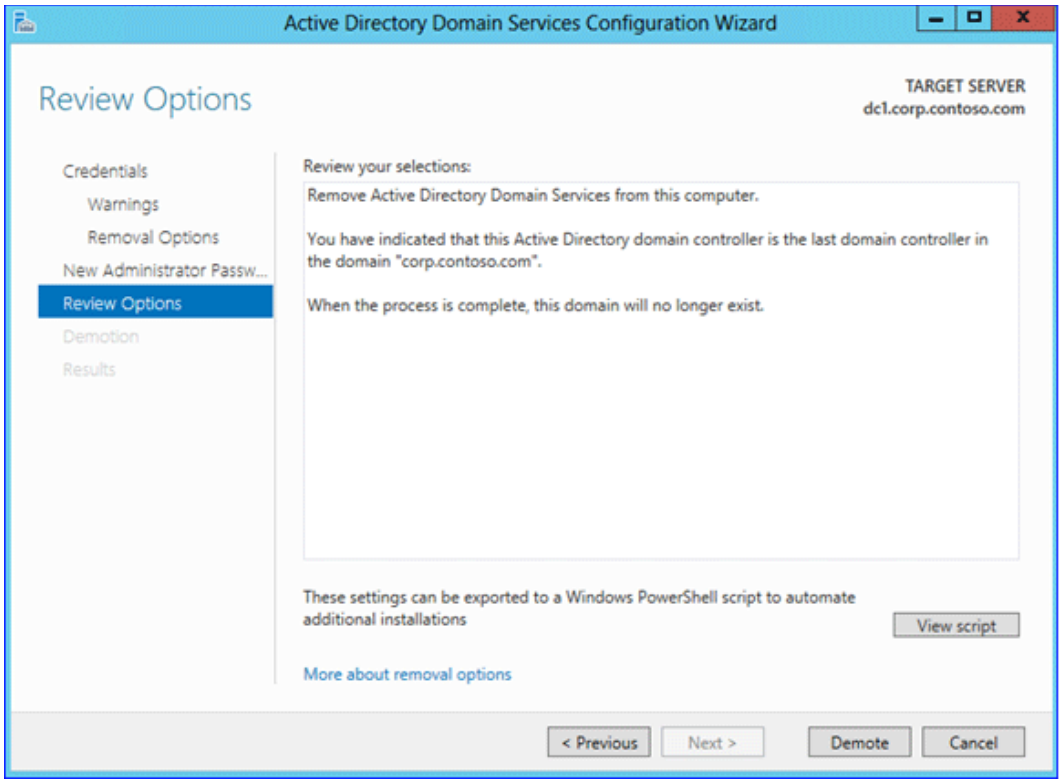

The **Confirmation** page shows the planned demotion; the page does not list demotion configuration options. This is the last page the wizard shows before the demotion begins. The View Script button creates a Windows PowerShell demotion script.

Click **Demote** to run the following AD DS Deployment cmdlet:

```
Uninstall-DomainController
```

Use the optional **Whatif** argument with the **Uninstall-ADDSDomainController** and cmdlet to review configuration information. This enables you to see the explicit and implicit values of a cmdlet's arguments.

For example:

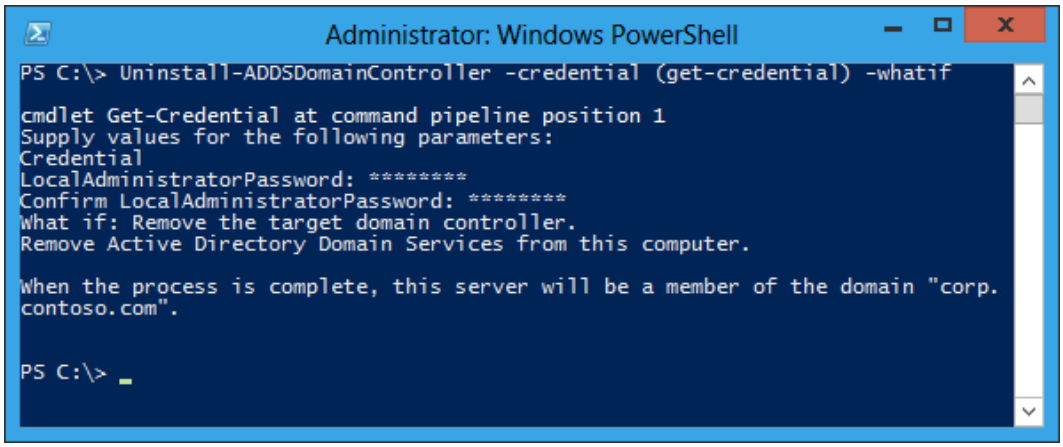

The prompt to restart is your last opportunity to cancel this operation when using ADDSDeployment Windows PowerShell. To override that prompt, use the **-force** or **confirm:$false** arguments.

Demotion

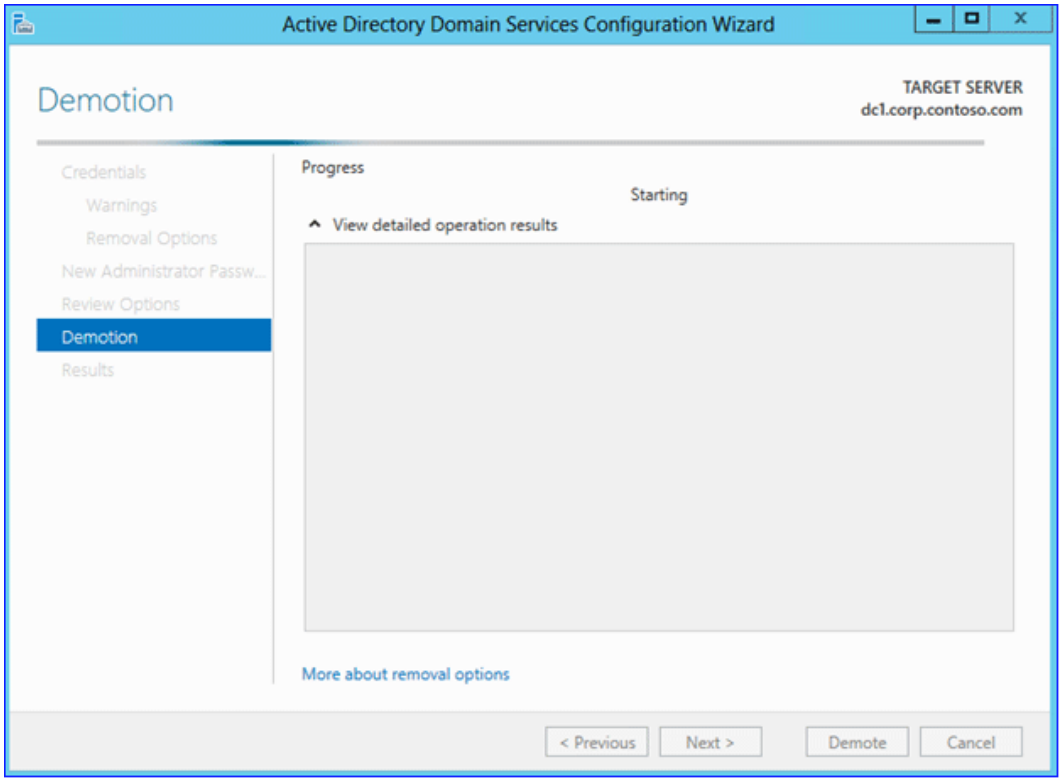

When the **Demotion** page displays, the domain controller configuration begins and cannot be halted or canceled. Detailed operations display on this page and write to logs:

- %systemroot%\debug\dcpromo.log

- %systemroot%\debug\dcpromoui.log

Since **Uninstall-AddsDomainController** and **Uninstall-WindowsFeature** only have one action apiece, they are shown here in the Confirmation phase with the minimum required arguments. Pressing ENTER starts the irrevocable demotion process and restarts the computer.

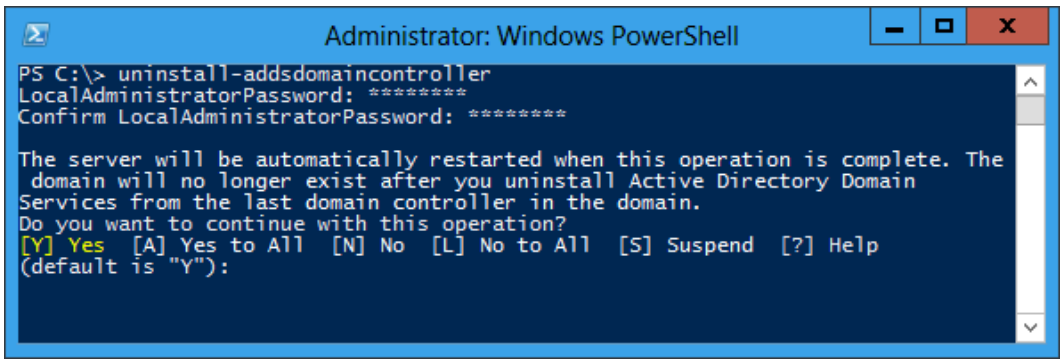

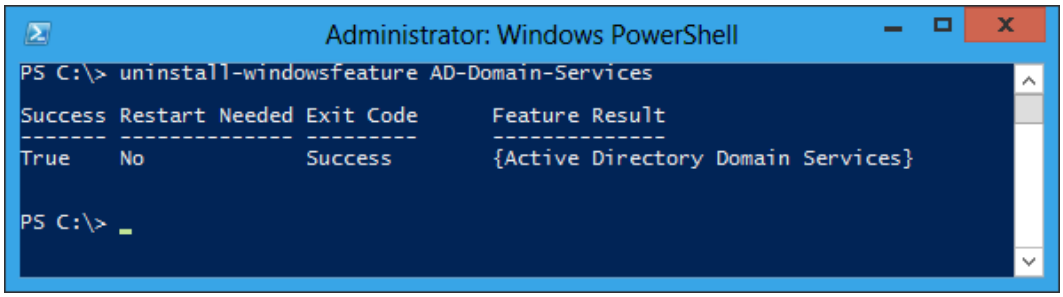

To accept the reboot prompt automatically, use the **-force** or **-confirm:$false** arguments with any ADDSDeployment Windows PowerShell cmdlet. To prevent the server from automatically rebooting at the end of promotion, use the **-norebootoncompletion:$false** argument.

Warning

Overriding the reboot is discouraged. The member server must reboot to function correctly.

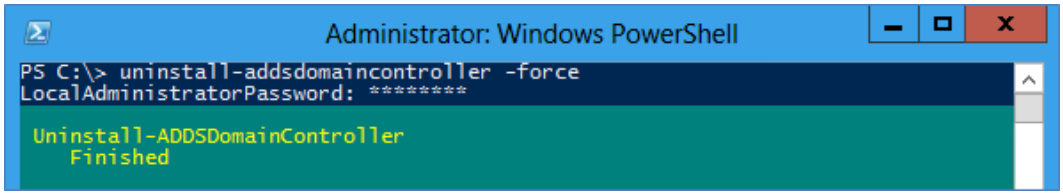

Here is an example of forcibly demoting with its minimal required arguments of **-forceremoval** and **-demoteoperationmasterrole**. The **-credential** argument is not required because the user logged on as a member of the Enterprise Admins group:

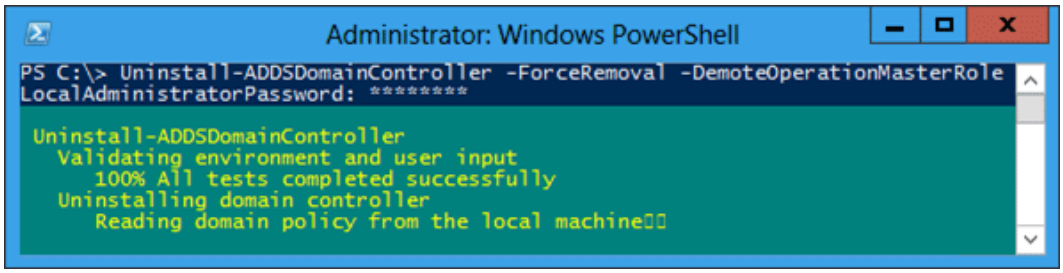

Here is an example of removing the last domain controller in the domain with its minimal required arguments of **-lastdomaincontrollerindomain** and **-removeapplicationpartitions**:

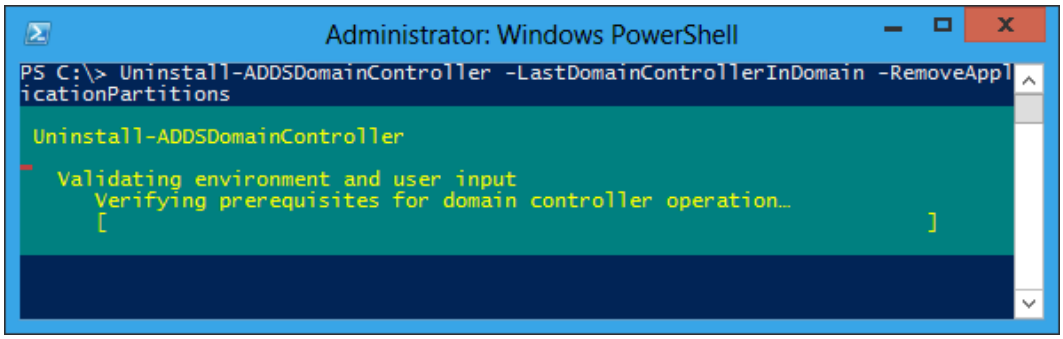

If you attempt to remove the AD DS role before demoting the server, Windows PowerShell blocks you with an intentional error:

```
Uninstall-WindowsFeature : An uninstallation prerequisite step failed duringthe removal
of AD-Domain-Services, and uninstallation cannot continue.1. The domain controller needs
to be demoted before the Active DirectoryDomain Services Role can be uninstalled.
```

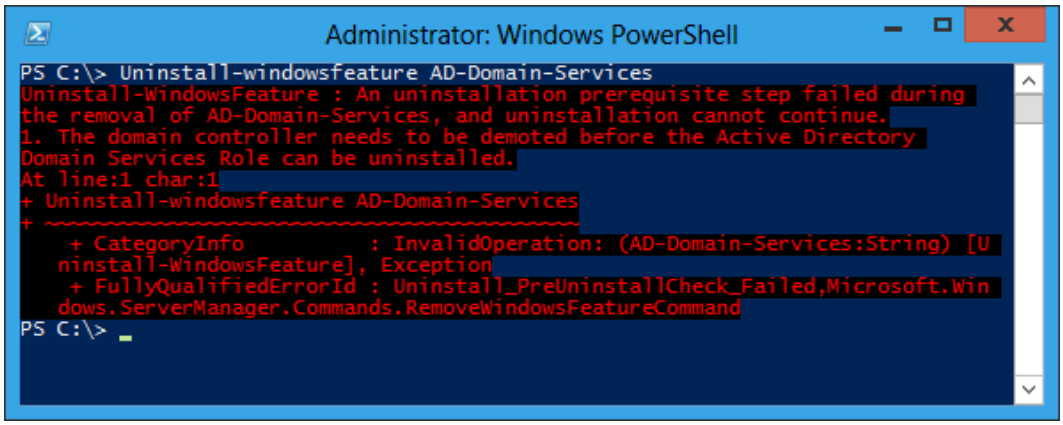

Important

You must restart the computer after demoting the server before you can remove the AD-Domain-Services role binaries.

Results

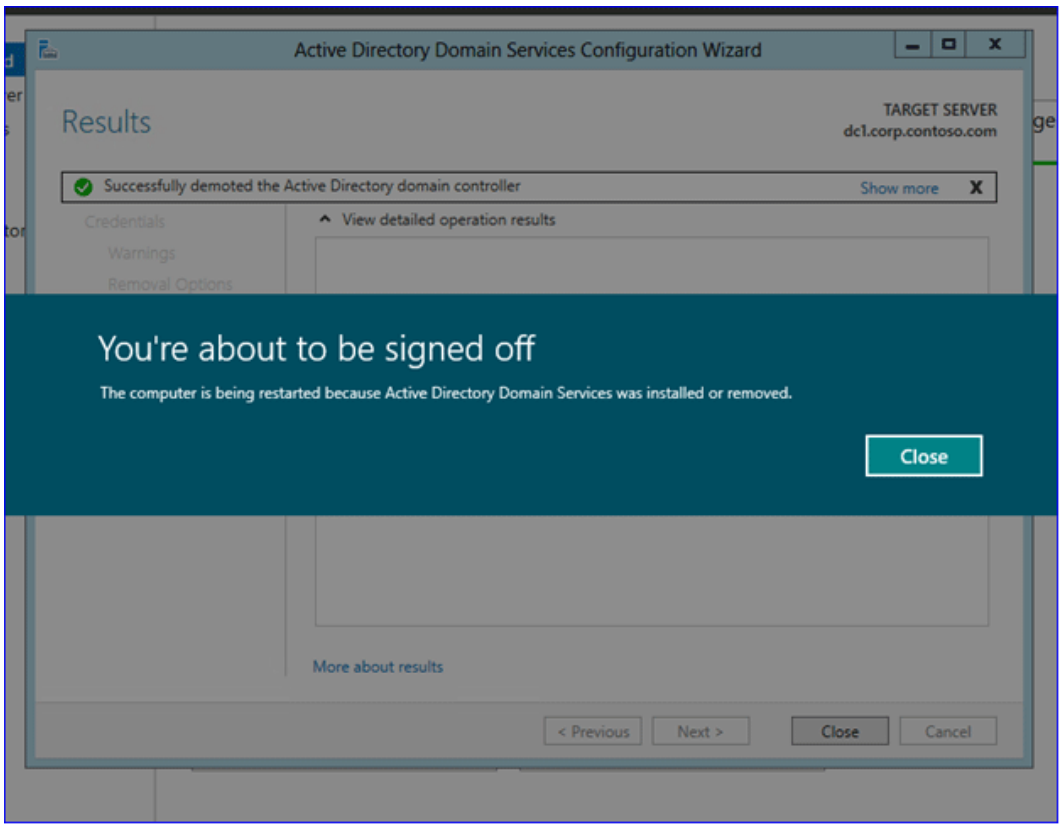

The **Results** page shows the success or failure of the promotion and any important administrative information. The domain controller will automatically reboot after 10 seconds.

AD DS Installation and Removal Wizard Page Descriptions

This is pre-release documentation that is subject to change in future releases. Blank sections are included as placeholders.

This topic provides descriptions for the controls on the following wizard pages that comprise the AD DS server role installation and removal in Server Manager.

- Deployment Configuration
- Domain Controller Options
- DNS Options
- RODC Options
- Additional Options
- Paths

- Preparation Options
- Review Options
- Prerequisites Check
- Results
- Role Removal credentials
- AD DS Removal Options and Warnings
- New Administrator Password
- Review Options

Deployment Configuration

Server Manager begins every domain controller installation with the **Deployment Configuration** page. The remaining options and required fields change on this page and subsequent pages, depending on which deployment operation you select. For example, if you create a new forest, the **Preparation Options** page does not appear, but it does if you install the first domain controller that runs Windows Server 2012 in an existing forest or domain.

Some validations tests are performed on this page, and again later as part of prerequisite checks. For example, if you try to install the first Windows Server 2012 domain controller in a forest that has Windows 2000 functional level, an error appears on this page.

The following options appear when you create a new forest.

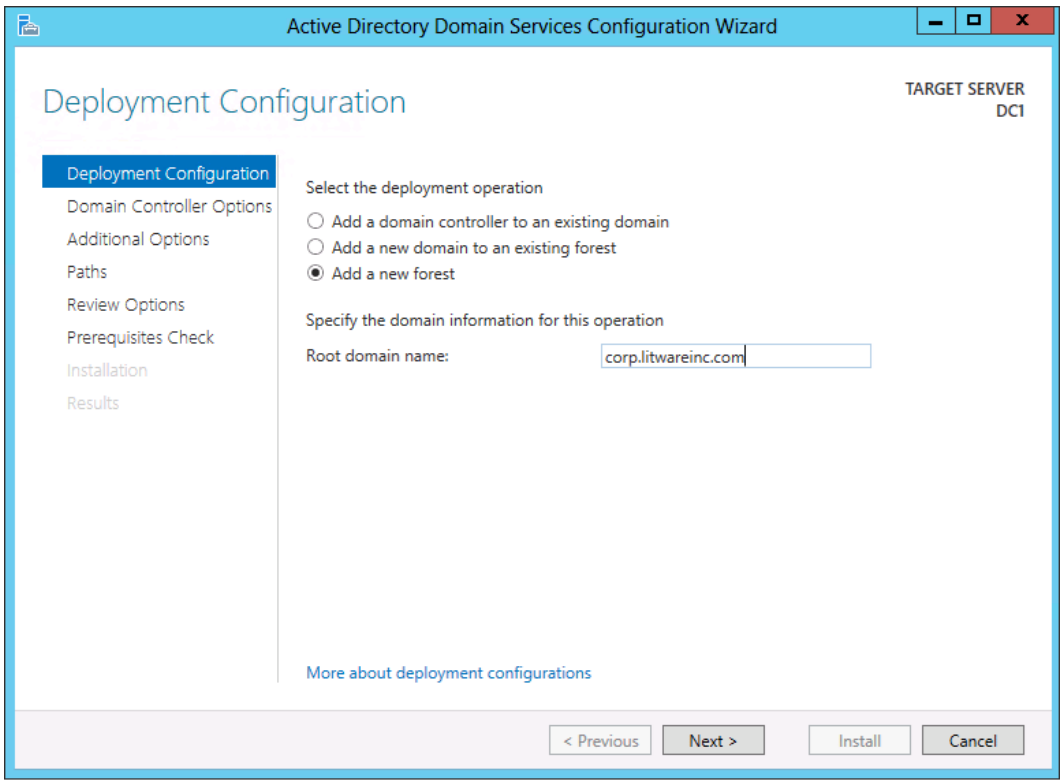

- When you create a new forest, you must specify a name for the forest root domain. The forest root domain name cannot be single-labeled (for example, it must be “contoso.com” instead of “contoso”). It must use allowed DNS domain naming conventions. You can specify an Internationalized Domain Name (IDN). For more information about DNS domain naming conventions, see KB 909264.
- Do not create new Active Directory forests with the same name as your external DNS name. For example, if your Internet DNS URL is http://contoso.com, you must choose a different name for your internal forest to avoid future compatibility issues. That name should be unique and unlikely for web traffic, such as corp.contoso.com.
- You must be a member of Administrators group on the server where you want to create a new forest.

For more information about how to create a forest, see Install a New Windows Server 2012 Active Directory Forest (Level 200).

The following options appear when you create a new domain.

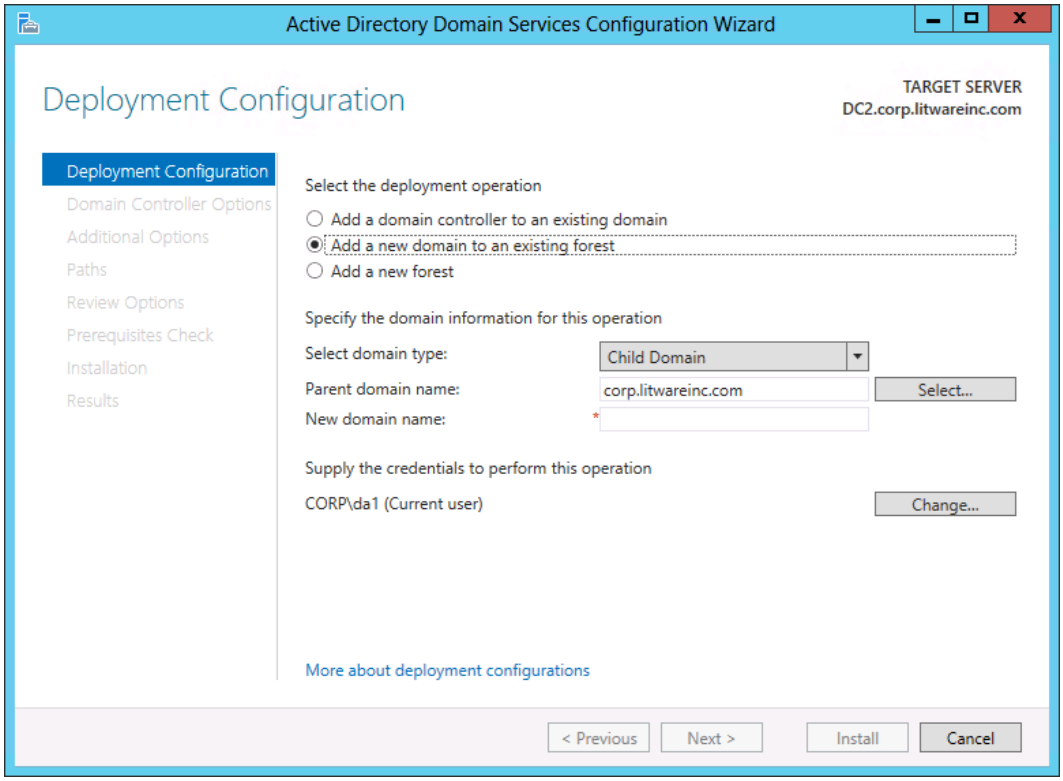

Note

If you create a new tree domain, you need to specify the name of the forest root domain instead of the parent domain, but the remaining wizard pages and options are the same.

- Click **Select** to browse to the parent domain or Active Directory tree, or type a valid parent domain or tree name. Then type the name of the new domain in **New domain name**.
- Tree domain: provide a valid, fully qualified root domain name; the name cannot be single-labeled and must use DNS domain name requirements.

- Child domain: provide a valid, single-label child domain name; the name must use DNS domain name requirements.
- The Active Directory Domain Services Configuration Wizard prompts you for domain credentials if your current credentials are not from the domain. Click **Change** to provide domain credentials.

For more information about how to create a domain, see Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200).

The following options appear when you add a new domain controller to an existing domain.

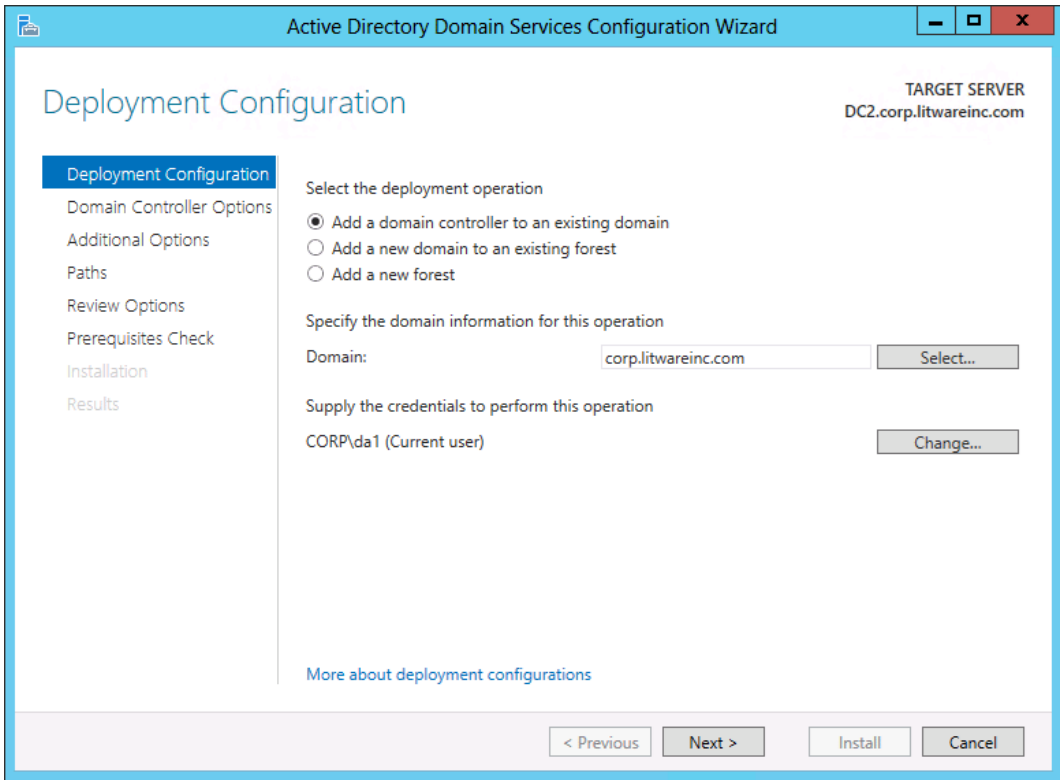

- Click **Select** to browse to the domain, or type a valid domain name.
- Server Manager prompts you for valid credentials if needed. Installing an additional domain controller requires membership in the Domain Admins group.

  In addition, installing the first domain controller that runs Windows Server 2012 in a forest requires credentials that include group memberships in both the Enterprise Admins and Schema Admins groups. The Active Directory Domain Services Configuration Wizard prompts you later if your current credentials do not have adequate permissions or group memberships.

For more information about how to add a domain controller to an existing domain, see Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200).

Domain Controller Options

If you are creating a new forest, the Domain Controller Options page has these options:

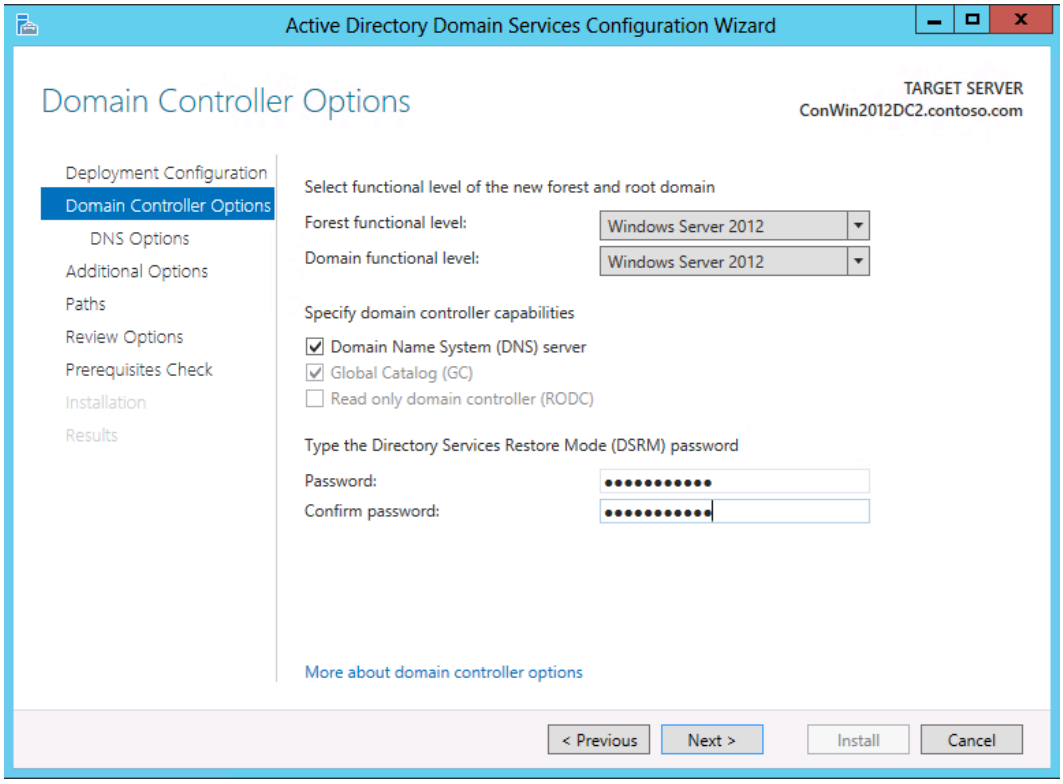

- The forest and domain functional levels are set to Windows Server 2012 by default.

  There is one new feature available at the Windows Server 2012 domain functional level: the Support for Dynamic Access Control and Kerberos armoring KDC administrative template policy has two settings (Always provide claims and Fail unarmored authentication requests) that require Windows Server 2012 domain functional level. For more information, see “Support for claims, compound authentication and Kerberos armoring” in What's new in Kerberos Authentication.

  The Windows Server 2012 forest functional level does not provide any new features, but it ensures that any new domain created in the forest will automatically operate at the Windows Server 2012 domain functional level. The Windows Server 2012 domain functional level does not provide any new other features beside support for Dynamic Access Control and Kerberos armoring, but it ensures that any domain controller in the domain runs Windows Server 2012. For more information about other features that are available at different functional levels, see Understanding Active Directory Domain Services (AD DS) Functional Levels.

  Beyond functional levels, a domain controller that runs Windows Server 2012 provides additional features that are not available on a domain controller that runs an earlier version of Windows Server. For example, a domain controller that runs Windows Server 2012 can be

used for virtual domain controller cloning, whereas a domain controller that runs an earlier version of Windows Server cannot.

- DNS server is selected by default when you create a new forest. The first domain controller in the forest must be a global catalog (GC) server, and it cannot be a read only domain controller (RODC).
- The Directory Services Restore Mode (DSRM) password is needed in order to log on to a domain controller where AD DS is not running. The password you specify must adhere to the password policy applied to the server, which by default does not require a strong password; only a non-blank password. Always choose a strong, complex password or preferably, a passphrase. For information about how to synchronize the DSRM password with the password of a domain user account, see KB 961320.

For more information about how to create a forest, see Install a New Windows Server 2012 Active Directory Forest (Level 200).

If you are creating a child domain, the Domain Controller Options page has these options:

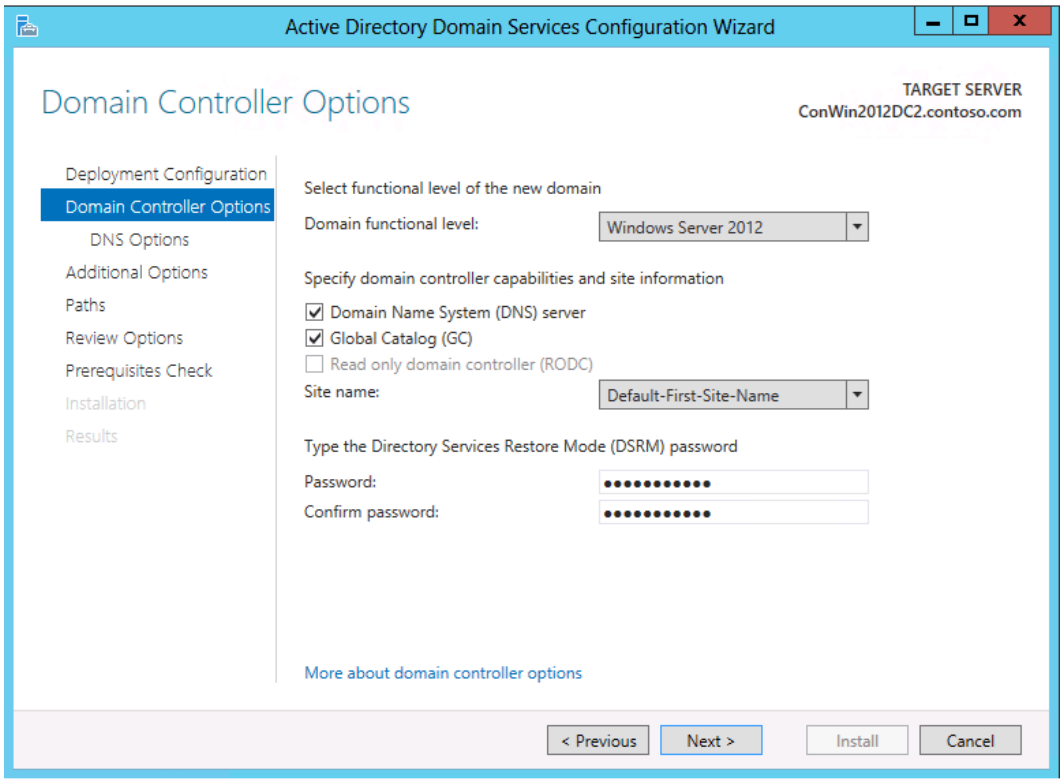

- The domain functional level is set to Windows Server 2012 by default. You can specify any other value that is at least the value of the forest functional level or higher.
- The configurable domain controller options include **DNS server** and **Global Catalog**; you cannot configure read-only domain controller as the first domain controller in a new domain.

  Microsoft recommends that all domain controllers provide DNS and global catalog services for high availability in distributed environments, which is why the wizard enables these options by default when creating a new domain.

- The **Domain Controller Options** page also enables you to choose the appropriate Active Directory logical **site name** from the forest configuration. By default, it selects the site with the most correct subnet. If there is only one site, it selects that site automatically.

  Important
  If the server does not belong to an Active Directory subnet and there is more than one site, nothing is selected and the **Next** button is unavailable until you choose a site from the list.

For more information about how to create a domain, see Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200).

If you are adding a domain controller to a domain, the Domain Controller Options page has these options:

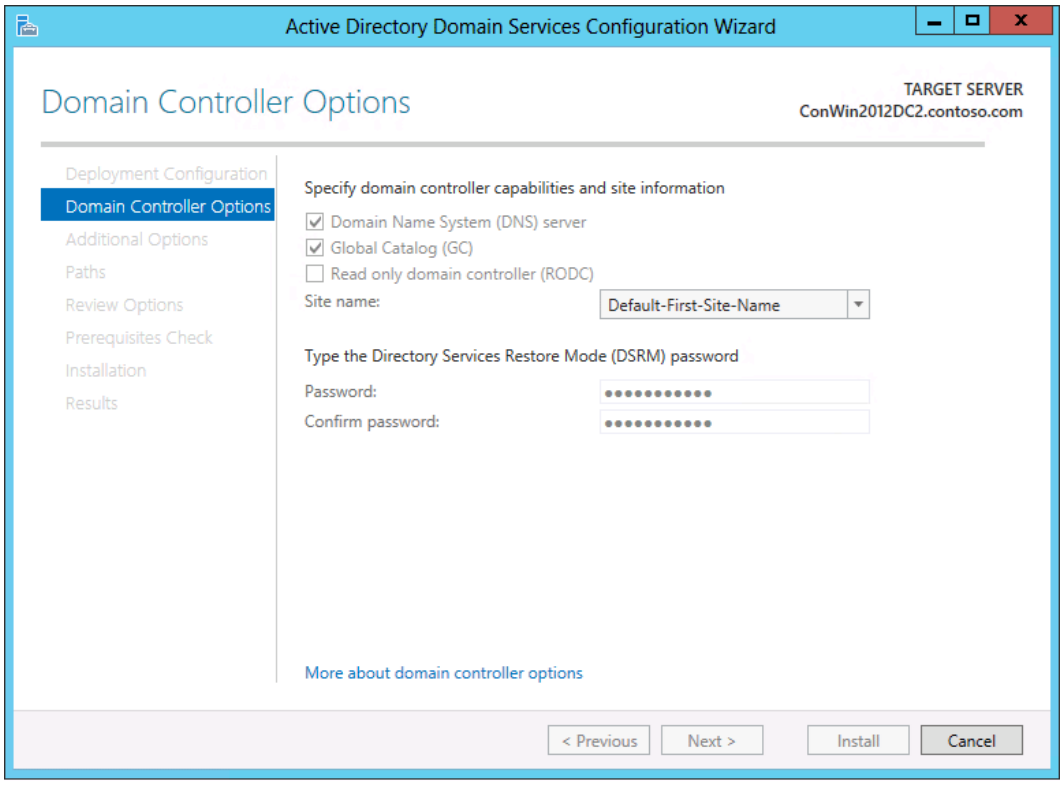

- The configurable domain controller options include **DNS server** and **Global Catalog**, and **Read-only domain controller**.

  Microsoft recommends that all domain controllers provide DNS and global catalog services for high availability in distributed environments, which is why the wizard enables these options by default. For more information about deploying RODCs, see Read-Only Domain Controller Planning and Deployment Guide.

For more information about how to add a domain controller to an existing domain, see Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200).

DNS Options

If you install DNS server, the following **DNS Options** page appears:

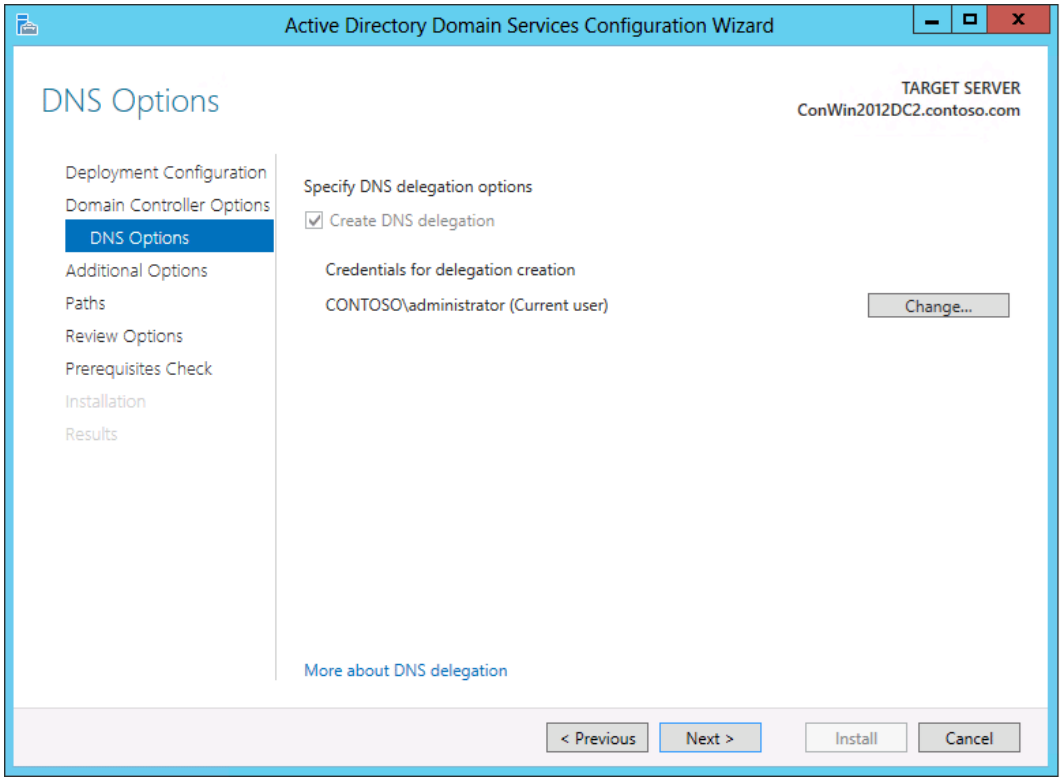

When you install DNS server, delegation records that point to the DNS server as authoritative for the zone should be created in the parent Domain Name System (DNS) zone. Delegation records transfer name resolution authority and provide correct referral to other DNS servers and clients of the new servers that are being made authoritative for the new zone. These resource records include the following:

- A name server (NS) resource record to effect the delegation. This resource record advertises that the server named ns1.na.example.microsoft.com is an authoritative server for the delegated subdomain.
- A host (A or AAAA) resource record—also known as a glue record—must be present to resolve the name of the server that is specified in the name server (NS) resource record to its IP address. The process of resolving the host name in this resource record to the delegated DNS server in the name server (NS) resource record is sometimes referred to as "glue chasing."

You can have the Active Directory Domain Services Configuration Wizard create them automatically. The wizard verifies that the appropriate records exist in the parent DNS zone after you click **Next** on the **Domain Controller Options** page. If the wizard cannot verify that the records exist in the parent domain, the wizard provides you with the option to create a new DNS delegation for a new domain (or update the existing delegation) automatically and continue with the new domain controller installation.

Alternatively, you can create these DNS delegation records before you install DNS server. To create a zone delegation, open **DNS Manager**, right-click the parent domain, and then click **New Delegation**. Follow the steps in the New Delegation Wizard to create the delegation.

The installation process tries to create the delegation to ensure that computers in other domains can resolve DNS queries for hosts, including domain controllers and member computers, in the DNS subdomain. Note that the delegation records can be automatically created only on Microsoft DNS servers. If the parent DNS domain zone resides on third party DNS servers such as BIND, a warning about the failure to create DNS delegation records appears on the Prerequisites check page. For more information about the warning, see Known issues for installing AD DS.

Delegations between the parent domain and the subdomain being promoted can be created and validated before or after the installation. There is no reason to delay the installation of a new domain controller because you cannot create or update the DNS delegation.

For more information about delegation, see Understanding Zone Delegation (http://go.microsoft.com/fwlink/?LinkId=164773). If zone delegation is not possible in your situation, you might consider other methods for providing name resolution from other domains to the hosts in your domain. For example, the DNS administrator of another domain could configure conditional forwarding, stub-zones, or secondary zones in order to resolve names in your domain. For more information, see the following topics:

- Understanding zone types (http://go.microsoft.com/fwlink/?LinkID=157399)
- Understanding stub zones (http://go.microsoft.com/fwlink/?LinkId=164776)
- Understanding forwarders (http://go.microsoft.com/fwlink/?LinkId=164778)

RODC Options

The following options appear when you install a read-only domain controller (RODC).

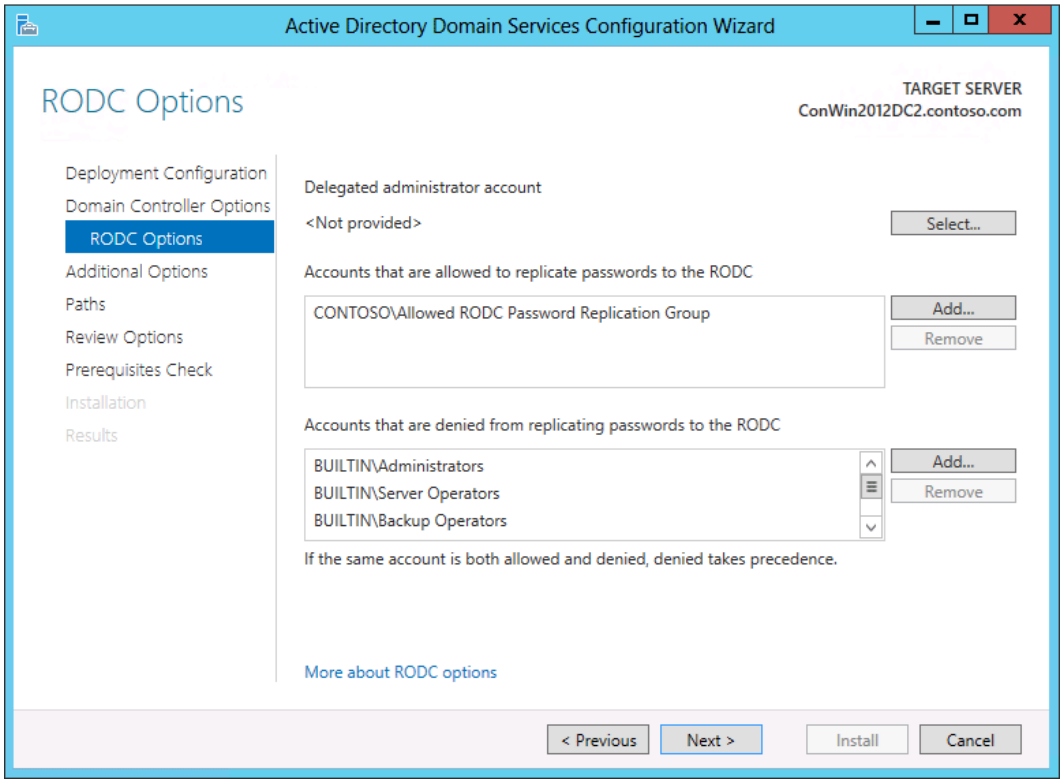

- Delegated administrator accounts gain local administrative permissions to the RODC. These users can operate with privileges equivalent to the local computer's Administrators group. They are not members of the Domain Admins or the domain built-in Administrators groups. This option is useful for delegating branch office administration without giving out domain administrative permissions. Configuring delegation of administration is not required. For more information, see Administrator Role Separation.
- The Password Replication Policy acts as an access control list (ACL). It determines if an RODC should be permitted to cache a password. After the RODC receives an authenticated user or computer logon request, it refers to the Password Replication Policy to determine if the password for the account should be cached. The same account can then perform subsequent logons more efficiently.

  The Password Replication Policy (PRP) lists the accounts whose passwords are allowed to be cached, and accounts whose passwords are explicitly denied from being cached. The list of user and computer accounts that are permitted to be cached does not imply that the RODC has necessarily cached the passwords for those accounts. An administrator can, for example, specify in advance any accounts that an RODC will cache. This way, the RODC can authenticate those accounts, even if the WAN link to the hub site is offline.

  Any users or computers who are not allowed (including implicit) or denied do not cache their password. If those users or computers do not have access to a writable domain controller, they cannot access AD DS-provided resources or functionality. For more information about the PRP, see Password Replication Policy. For more information about managing the PRP, see Administering the Password Replication Policy.

For more information about installing RODCs, see [Install a Windows Server 2012 Active Directory Read-Only Domain Controller (RODC) (Level 200)](#).

Additional Options

The following option appears on the **Additional Options** page if you are creating a new domain:

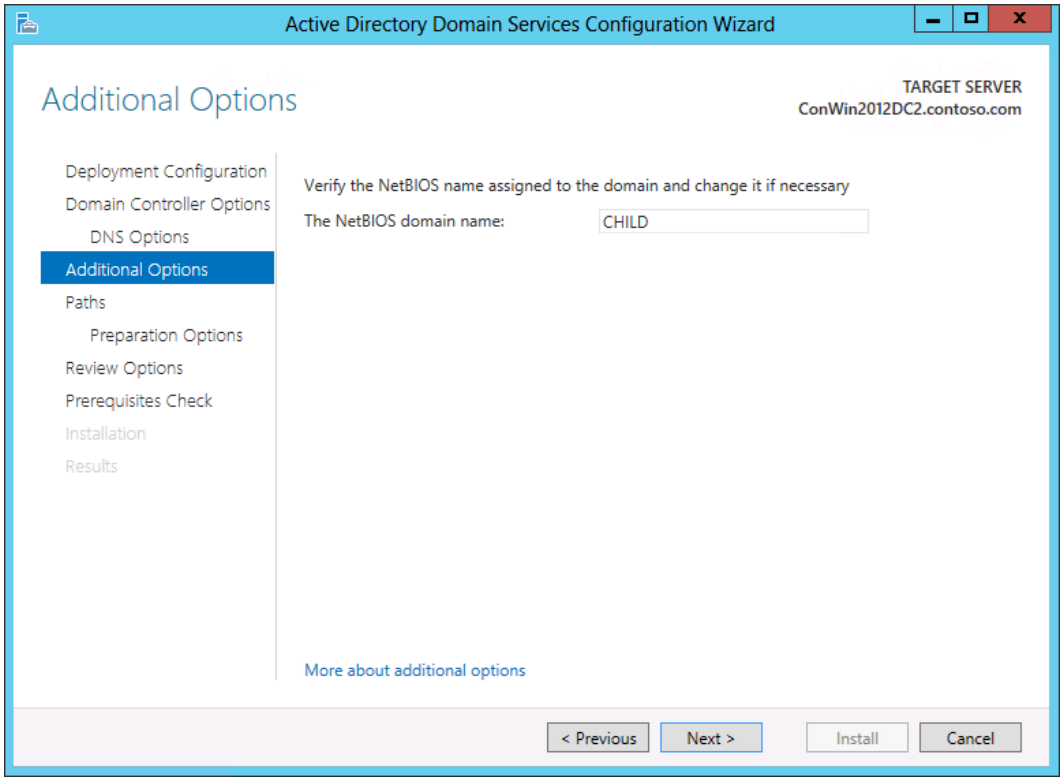

The following options appear on the **Additional Options** page if you install an additional domain controller in an existing domain:

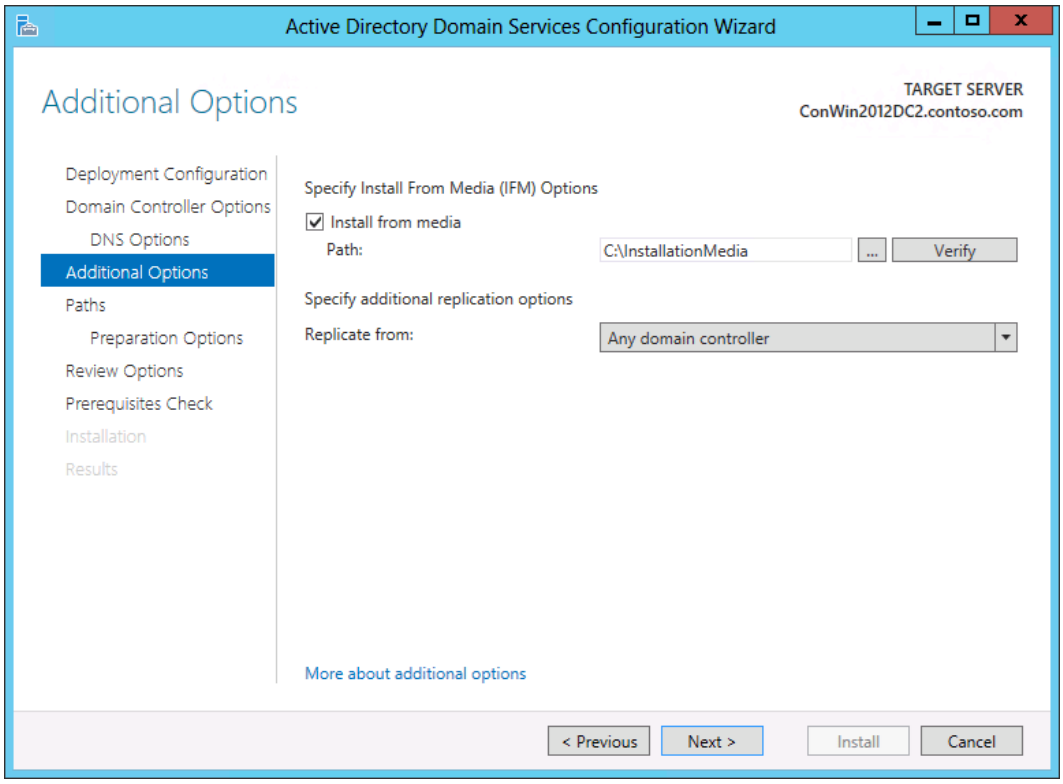

- You can either specify a domain controller as the replication source, or allow the wizard to choose any domain controller as the replication source.
- You can also choose to install the domain controller using backed up media using the Install from media (IFM) option. If the installation media is stored locally, the **Install from media Path** option allows you to browse to the file location. The browse option is not available for a remote installation. You can click **Verify** to ensure the provided path is valid media. Media used by the IFM option must be created with Windows Server Backup or Ntdsutil.exe from another existing Windows Server 2012 computer only; you cannot use a Windows Server 2008 R2 or previous operating system to create media for a Windows Server 2012 domain controller. If the media is protected with a SYSKEY, Server Manager prompts for the image's password during verification.

For more information about how to create a domain, see Install a New Windows Server 2012 Active Directory Child or Tree Domain (Level 200). For more information about how to add a domain controller to an existing domain, see Install a Replica Windows Server 2012 Domain Controller in an Existing Domain (Level 200).

Paths

The following options appear on the **Paths** page.

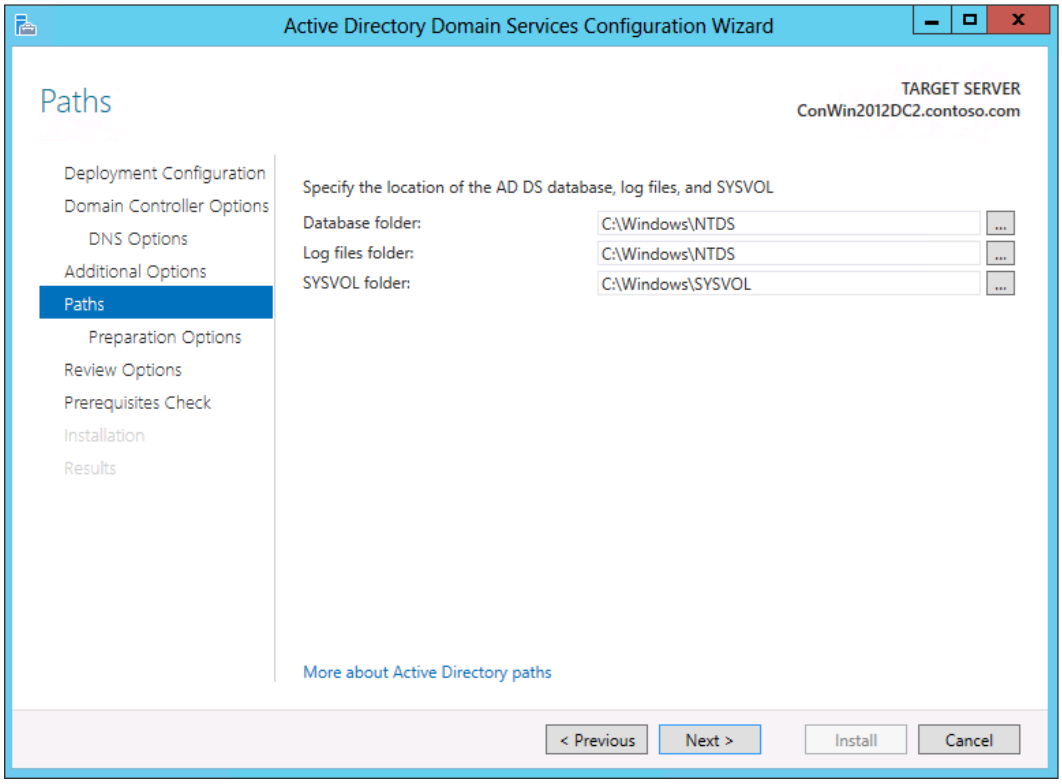

- The **Paths** page enables you to override the default folder locations of the AD DS database, the database transaction logs, and the SYSVOL share. The default locations are always in %systemroot%.

Specify the location for the AD DS database (NTDS.DIT), log files, and SYSVOL. For a local installation, you can browse to the location where you want to store the files.

Preparation Options

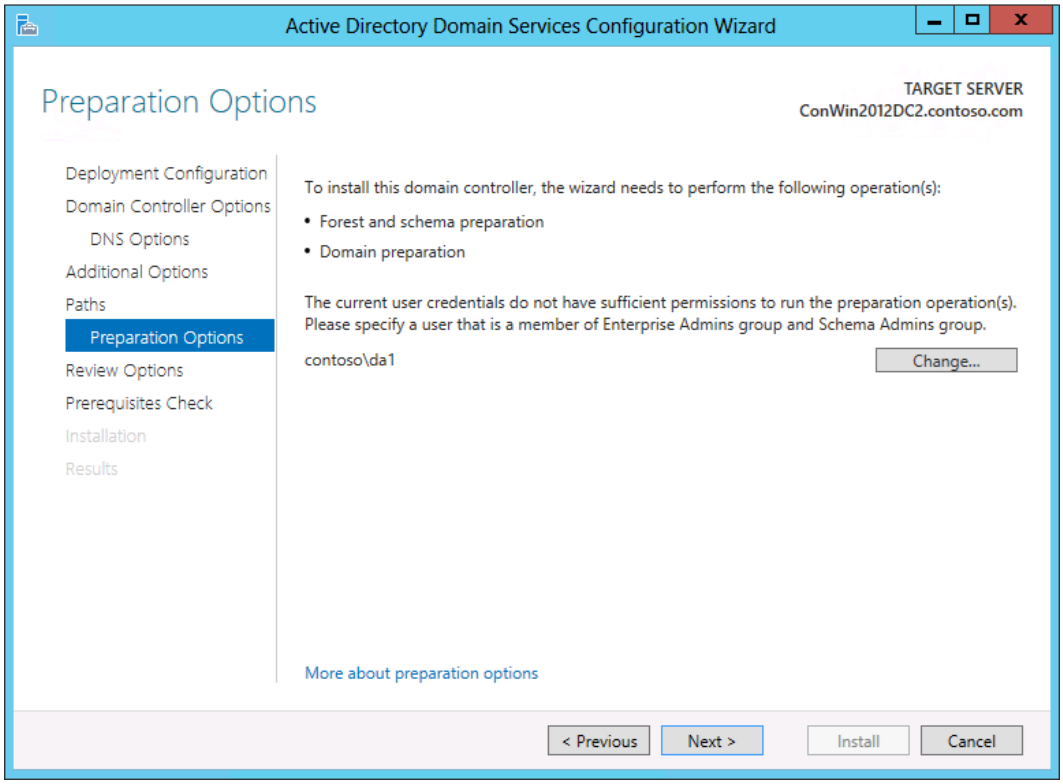

If you are not currently logged on with sufficient credentials to run adprep.exe commands and adprep is required to run in order to complete the AD DS installation, you are prompted to supply credentials to run adprep.exe. Adprep is required to run in order to add the first domain controller that runs Windows Server 2012 to an existing domain or forest. More specifically:

- Adprep /forestprep must be run to add the first domain controller that runs Windows Server 2012 to an existing forest. This command must be run by a member of the Enterprise Admins group, the Schema Admins group, and the Domain Admins group of the domain that hosts the schema master. For this command to complete successfully, there must be connectivity between the computer where you run the command and the schema master for the forest.
- Adprep /domainprep must be run to add the first domain controller that runs Windows Server 2012 to an existing domain. This command must be run by a member of the Domain Admins group of the domain where you are installing the domain controller that runs Windows Server 2012. For this command to complete successfully, there must be connectivity between the computer where you run the command and the infrastructure master for the domain.
- Adprep /rodcprep must be run to add the first RODC to an existing forest. This command must be run by a member of the Enterprise Admins group. For this command to complete successfully, there must be connectivity between the computer where you run the command and the infrastructure master for each application directory partition in the forest.

For more information about Adprep.exe, see Adprep.exe integration and see Running Adprep.exe.

Review Options

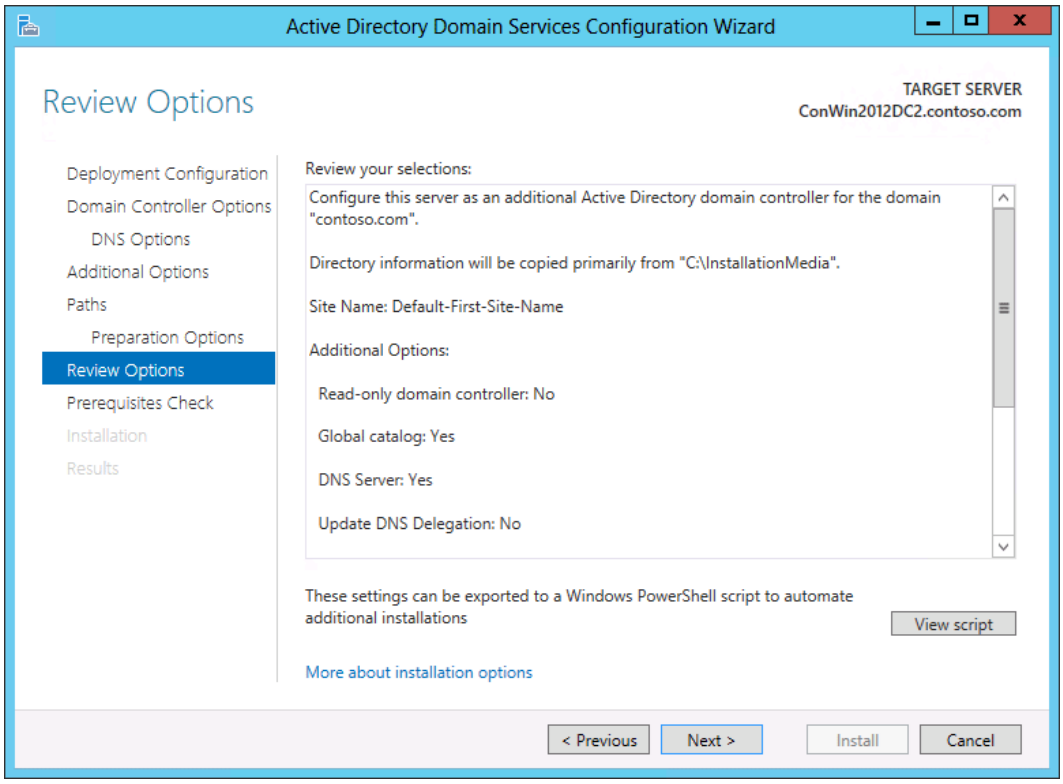

- The **Review Options** page enables you to validate your settings and ensure that they meet your requirements before you start the installation. This is not the last opportunity to stop the installation using Server Manager. This page simply enables you to review and confirm your settings before continuing the configuration.
- The **Review Options** page in Server Manager also offers an optional **View Script** button to create a Unicode text file that contains the current ADDSDeployment configuration as a single Windows PowerShell script. This enables you to use the Server Manager graphical interface as a Windows PowerShell deployment studio. Use the Active Directory Domain Services Configuration Wizard to configure options, export the configuration, and then cancel the wizard. This process creates a valid and syntactically correct sample for further modification or direct use.

Prerequisites Check

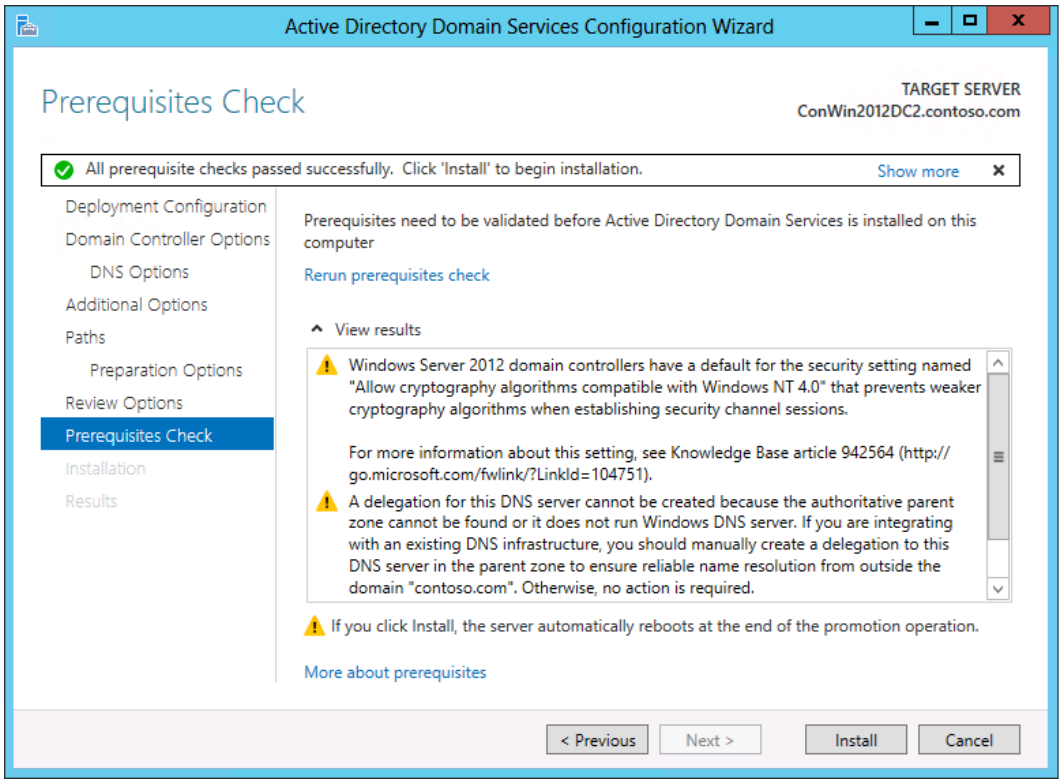

Some of the warnings that appear on this page include:

- Domain controllers that run Windows Server 2008 or later have a default setting for “Allow cryptography algorithms compatible with Windows NT 4” that prevents weaker cryptography algorithms when establishing secure channel sessions. For more information about the potential impact and a workaround, see KB article 942564.
- DNS delegation could not be created or updated. For more information, see DNS Options.
- The prerequisite check requires WMI calls. They can fail if they are blocked firewall rules block, and return an RPC server unavailable error.

For more information about the specific prerequisite checks that are performed for AD DS installation, see Prerequisite Tests.

Results

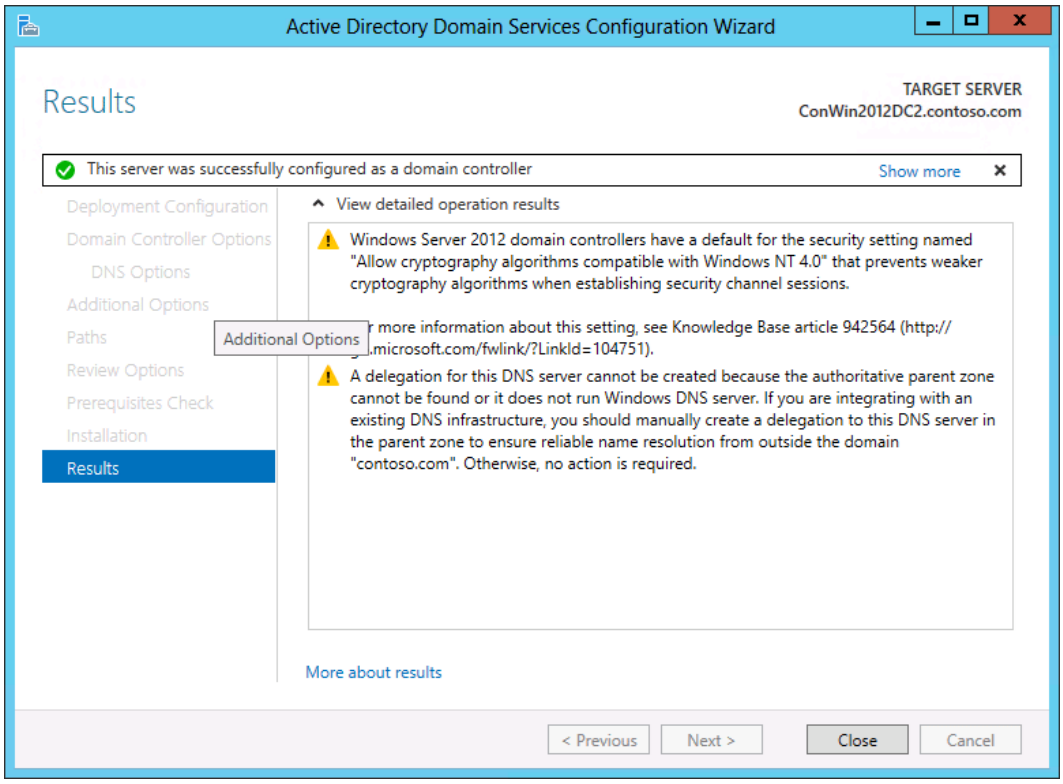

On this page, you can review the results of the installation.

You can also select to restart the target server after the wizard completes, but if the installation succeeds, the server will always restart regardless of whether you select that option. In some cases after the wizard completes on a target server that was not joined to the domain before the installation, the system state of the target server can make the server unreachable on the network, or the system state can prevent you from having permissions to manage the remote server.

If the target server fails to restart in this case, you must manually restart it. Tools such as shutdown.exe or Windows PowerShell cannot restart it. You can use Remote Desktop Services to log on and remotely shut down the target server.

Role Removal credentials

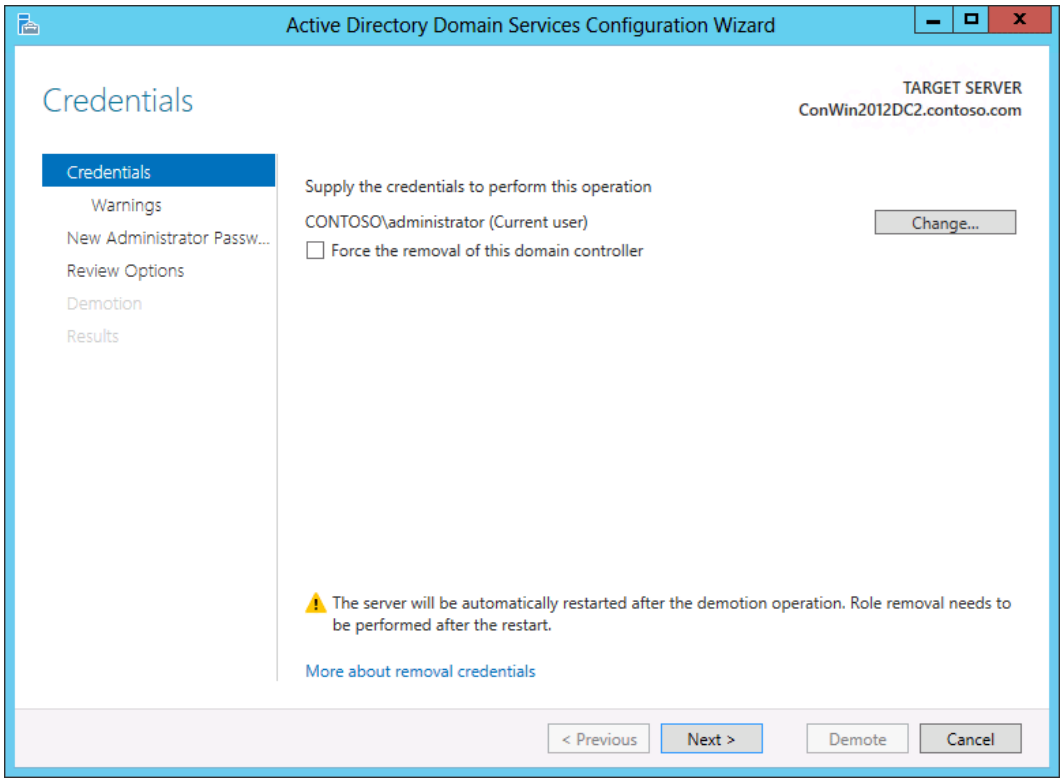

You configure demotion options on the **Credentials** page. Provide the credentials necessary to perform the demotion from the following list:

- Demoting an additional domain controller requires Domain Admin credentials. Selecting **Force removal of the domain controller** demotes the domain controller without removing the domain controller object's metadata from Active Directory.

  Important
  Do not select this option unless the domain controller cannot contact other domain controllers and there is *no reasonable way* to resolve that network issue. Forced demotion leaves orphaned metadata in Active Directory on the remaining domain controllers in the forest. In addition, all un-replicated changes on that domain controller, such as passwords or new user accounts, are lost forever. Orphaned metadata is the root cause in a significant percentage of Microsoft Customer Support cases for AD DS, Exchange, SQL, and other software. If you forcibly demote a domain controller, you *must* manually perform metadata cleanup immediately. For steps, review Clean Up Server Metadata.

- Demoting the last domain controller in a domain requires Enterprise Admins group membership, as this removes the domain itself (if this is the last domain in the forest, this removes the forest). Server Manager informs you if the current domain controller is the last

domain controller in the domain. Select **Last domain controller in the domain** to confirm the domain controller is the last domain controller in the domain.

For more information about removing AD DS, see Remove Active Directory Domain Services (Level 100) and Demoting Domain Controllers and Domains (Level 200).

AD DS Removal Options and Warnings

If you need help with the Review Options page, see Review Options

If the domain controller hosts additional roles, such as DNS server role or global catalog server, the following Warning page appears:

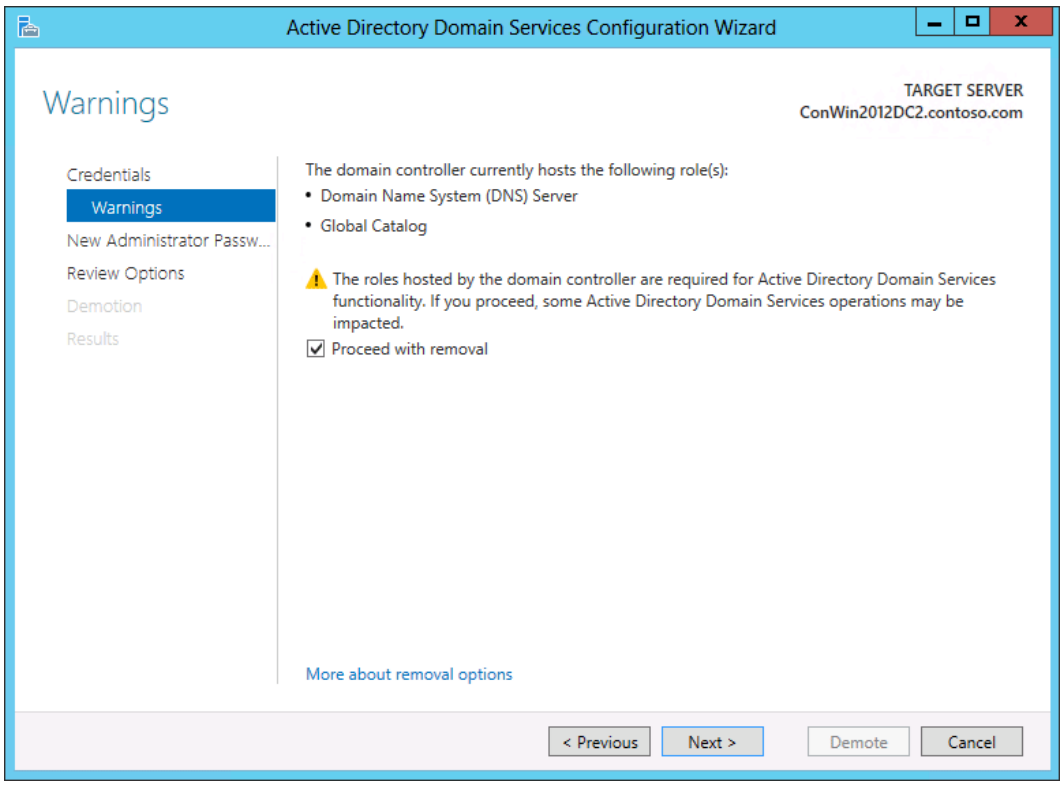

You must click **Proceed with removal** in order to acknowledge that the additional roles will no longer be available before you can click **Next** to continue.

If you force the removal of a domain controller, any Active Directory object changes that have not replicated to other domain controllers in the domain will be lost. Additionally, if the domain controller hosts operation master roles, the global catalog, or DNS server role, critical operations in the domain and forest may be impacted as follows. Before you remove a domain controller that hosts any operations master role, try to transfer the role to another domain controller. If it is not possible to transfer the role, first remove Active Directory Domain Services from this computer, and then use Ntdsutil.exe to seize the role. Use Ntdsutil on the domain controller that you plan to seize the role to; if possible, use a recent replication partner in the same site as this domain controller. For more information about transferring and seizing operations master roles, see

article 255504 in the Microsoft Knowledge Base. If the wizard cannot determine if the domain controller host an operations master role, run netdom.exe command to determine whether this domain controller performs any operations master roles.

- Global catalog: Users might have trouble logging on to domains in the forest. Before you remove a global catalog server, ensure that enough global catalog servers are in this forest and site to service user logons. If necessary, designate another global catalog server and update clients and applications with the new information.
- DNS server: All of the DNS data that is stored in Active Directory-integrated zones will be lost. After you remove AD DS, this DNS server will not be able to perform name resolution for the DNS zones that were Active Directory-integrated. Therefore, we recommend that you update the DNS configuration of all computers that currently refer to the IP address of this DNS server for name resolution with the IP address of a new DNS server.
- Infrastructure master: clients in the domain might have difficulty locating objects in other domains. Before you continue, transfer the infrastructure master role to a domain controller that is not a global catalog server.
- RID master: you might have problems creating new user accounts, computer accounts, and security groups. Before you continue, transfer the RID master role to a domain controller in the same domain as this domain controller.
- Primary domain controller (PDC) emulator: operations that are performed by the PDC emulator, such as Group Policy updates and password resets for non-AD DS accounts, will not function properly. Before you continue, transfer the PDC emulator master role to a domain controller that is in the same domain as this domain controller.
- Schema master: you will no longer be able to modify the schema for this forest. Before you continue, transfer the schema master role to a domain controller in the root domain in the forest.
- Domain naming master: you will no longer be able to add domains to or remove domains from this forest. Before you continue, transfer the domain naming master role to a domain controller in the root domain in the forest.
- All application directory partitions on this Active Directory domain controller will be removed. If a domain controller holds the last replica of one or more application directory partitions, when the removal operation is complete, those partitions will no longer exist.

Be aware that the domain will no longer exist after you uninstall Active Directory Domain Services from the last domain controller in the domain.

If the domain controller is a DNS server that is delegated to host the DNS zone, the following page will provide the option to remove the DNS server from the DNS zone delegation.

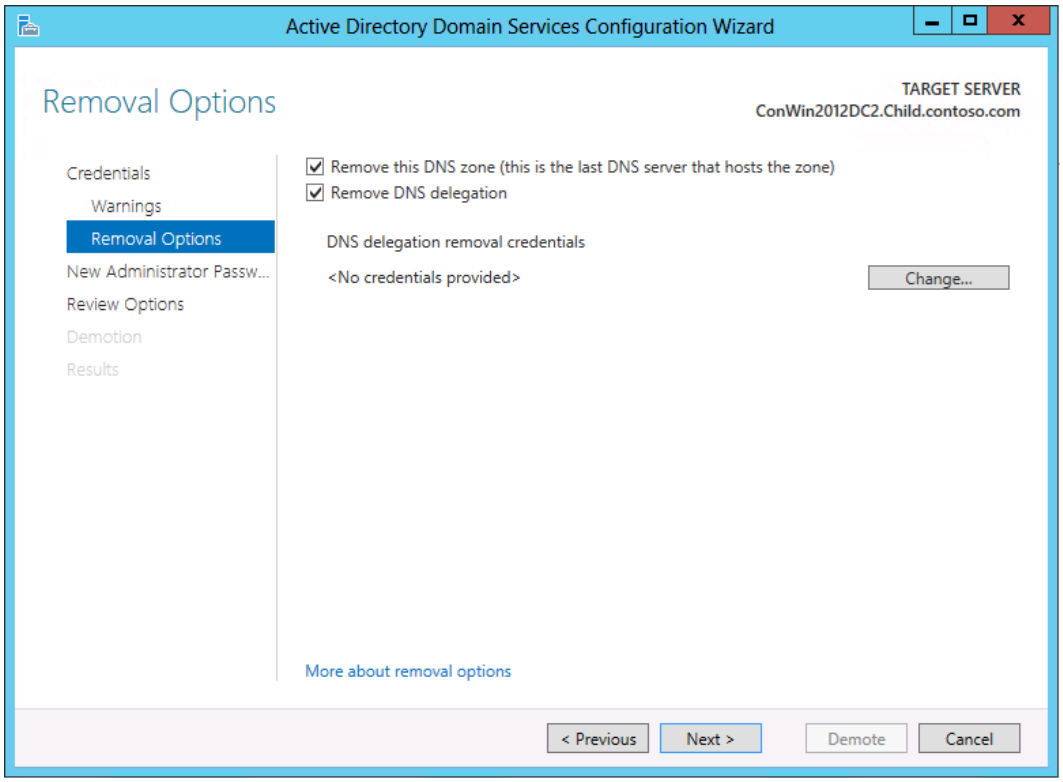

For more information about removing AD DS, see Remove Active Directory Domain Services (Level 100) and Demoting Domain Controllers and Domains (Level 200).

New Administrator Password

The **New Administrator Password** page requires you to provide a password for the built-in local computer's Administrator account, once the demotion completes and the computer becomes a domain member server or workgroup computer.

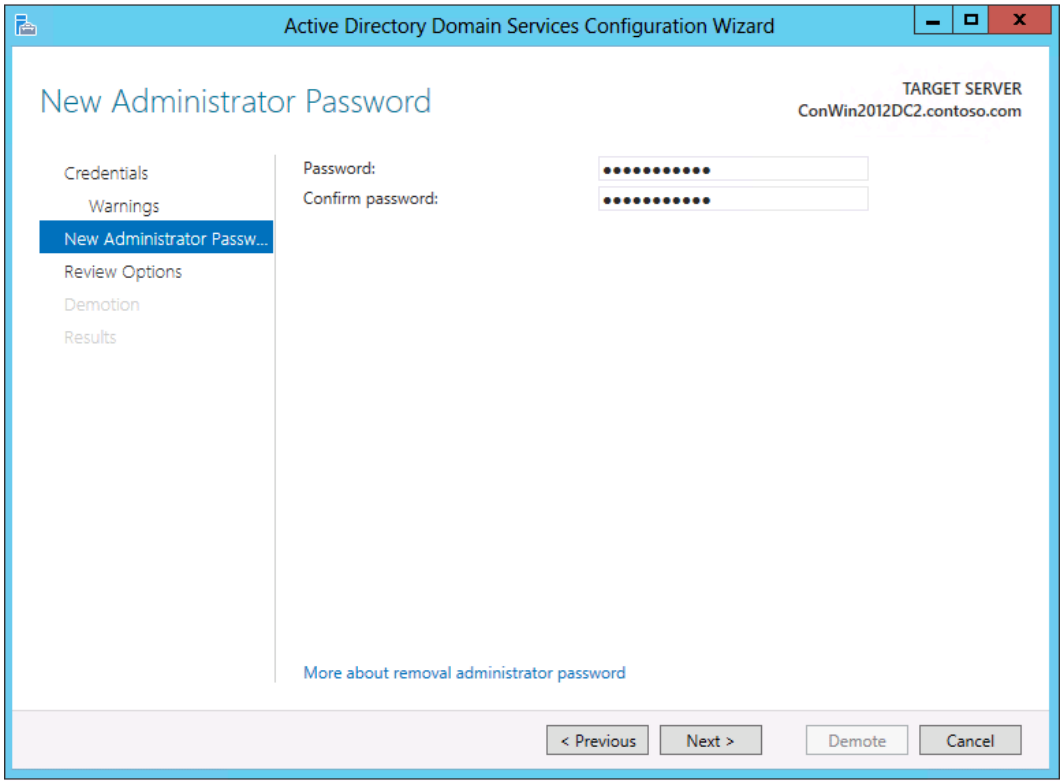

For more information about removing AD DS, see Remove Active Directory Domain Services (Level 100) and Demoting Domain Controllers and Domains (Level 200).

Review Options

The **Review Options** page provides you the chance to export the configuration settings for demotion to a Windows PowerShell script so you can automate additional demotions. Click **Demote** to remove AD DS.

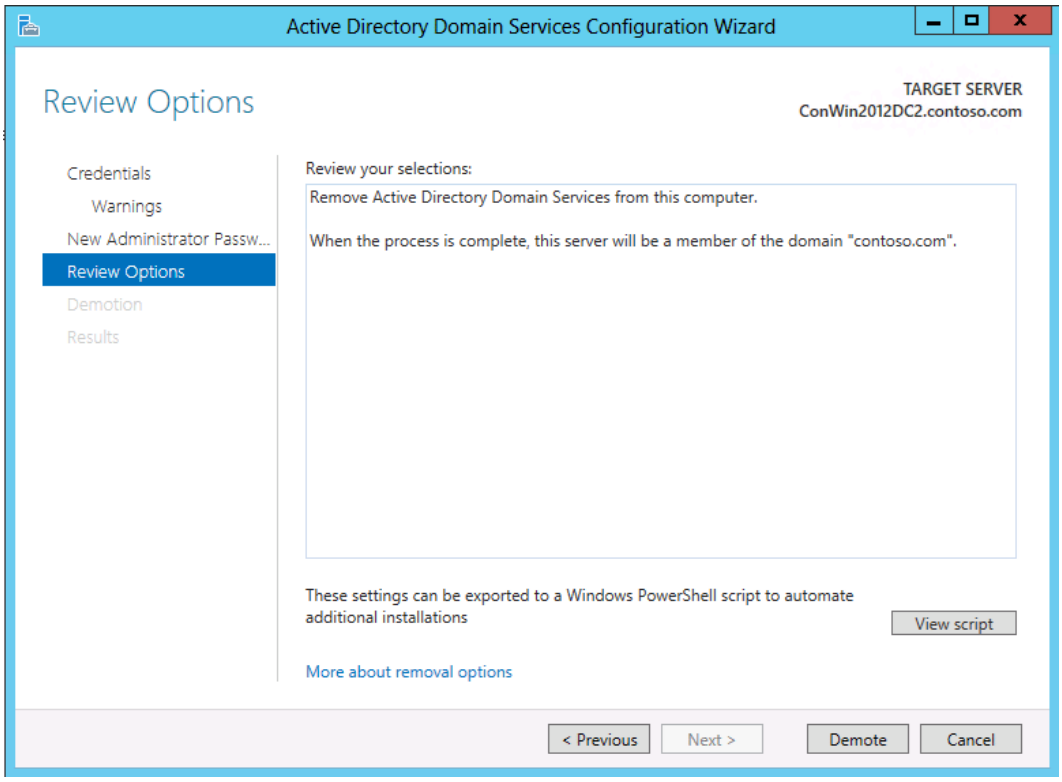

Windows Server 2012: Changes Made by Adprep.exe

This topic describes the changes that Adprep.exe makes in Windows Server 2012.

- Windows Server 2012: Forest-Wide Updates
- Windows Server 2012: Domain-Wide Updates
- Windows Server 2012: Read-Only Domain Controller Updates
- Windows Server 2012: Schema Updates

See Also

Windows Server 2008 R2: Appendix of Changes to Adprep.exe to Support AD DS

Windows Server 2008: Appendix of Changes to Adprep.exe to Support AD DS

Windows Server 2012: Schema Updates

This topic describes the changes that Adprep.exe makes in Windows Server 2012.

- Sch48.ldf
- Sch49.ldf
- Sch50.ldf
- Sch51.ldf
- Sch52.ldf
- Sch53.ldf
- Sch54.ldf
- Sch55.ldf
- Sch56.ldf

Schema Updates

Sch48.ldf

```
dn: CN=ms-DS-Members-Of-Resource-Property-List,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-MembersOfResourcePropertyList

adminDisplayName: ms-DS-Members-Of-Resource-Property-List

adminDescription: For a resource property list object, this multi-valued link attribute
points to one or more resource property objects.

attributeId: 1.2.840.113556.1.4.2103

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: ERw3Ta1MQUyK0rGAqyvRPA==

linkID: 2180

showInAdvancedViewOnly: TRUE

systemFlags: 16


dn: CN=ms-DS-Members-Of-Resource-Property-List-BL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd
```

objectClass: attributeSchema

ldapDisplayName: msDS-MembersOfResourcePropertyListBL

adminDisplayName: ms-DS-Members-Of-Resource-Property-List-BL

adminDescription: Backlink for ms-DS-Members-Of-Resource-Property-List. For a resource property object, this attribute references the resource property list object that it is a member of.

attributeId: 1.2.840.113556.1.4.2104

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: BLdpdLDtaEWlpVn0hix1pw==

linkID: 2181

showInAdvancedViewOnly: TRUE

systemFlags: 17

dn: CN=ms-DS-Claim-Value-Type,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimValueType

adminDisplayName: ms-DS-Claim-Value-Type

adminDescription: For a claim type object, specifies the value type of the claims issued.

attributeId: 1.2.840.113556.1.4.2098

attributeSyntax: 2.5.5.16

omSyntax: 65

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 0

schemaIdGuid:: uRdixo7k90e31WVSuK/WGQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

```
dn: CN=ms-DS-Claim-Possible-Values,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-ClaimPossibleValues
adminDisplayName: ms-DS-Claim-Possible-Values
adminDescription: For a claim type or resource property object, this attribute describes
the values suggested to a user when the he/she use the claim type or resource property in
applications.
attributeId: 1.2.840.113556.1.4.2097
attributeSyntax: 2.5.5.12
omSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeUpper: 1048576
schemaIdGuid:: 7u0oLnztP0Wv5JO9hvIXTw==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-Claim-Attribute-Source,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-ClaimAttributeSource
adminDisplayName: ms-DS-Claim-Attribute-Source
adminDescription: For a claim type object, this attribute points to the attribute that
will be used as the source for the claim type.
attributeId: 1.2.840.113556.1.4.2099
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: PhK87ua6ZkGeWymISot2sA==
```

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Type-Applies-To-Class,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimTypeAppliesToClass

adminDisplayName: ms-DS-Claim-Type-Applies-To-Class

adminDescription: For a claim type object, this linked attribute points to the AD security principal classes that for which claims should be issued. (For example, a link to the user class).

attributeId: 1.2.840.113556.1.4.2100

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: TA77anbYfEOutsPkFFTCcg==

linkID: 2176

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Shares-Possible-Values-With,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimSharesPossibleValuesWith

adminDisplayName: ms-DS-Claim-Shares-Possible-Values-With

adminDescription: For a resource property object, this attribute indicates that the suggested values of the claims issued are defined on the object that this linked attribute points to. Overrides ms-DS-Claim-Possible-Values on itself, if populated.

attributeId: 1.2.840.113556.1.4.2101

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: OtHIUgvOV0+JKxj1pDokAA==

linkID: 2178

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Shares-Possible-Values-With-BL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimSharesPossibleValuesWithBL

adminDisplayName: ms-DS-Claim-Shares-Possible-Values-With-BL

adminDescription: For a claim type object, this attribute indicates that the possible values described in ms-DS-Claim-Possible-Values are being referenced by other claim type objects.

attributeId: 1.2.840.113556.1.4.2102

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: 2yLVVJXs9UibvRiA67shgA==

linkID: 2179

showInAdvancedViewOnly: TRUE

systemFlags: 17

dn: CN=ms-DS-Is-Used-As-Resource-Security-Attribute,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-IsUsedAsResourceSecurityAttribute

adminDisplayName: ms-DS-Is-Used-As-Resource-Security-Attribute

adminDescription: For a resource property, this attribute indicates whether it is being used as a secure attribute.

attributeId: 1.2.840.113556.1.4.2095

attributeSyntax: 2.5.5.8

omSyntax: 1

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: nfjJUTBHjUaitR1JMhLRfg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-SPP-KMS-Ids,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msSPP-KMSIds

adminDisplayName: ms-SPP-KMS-Ids

adminDescription: KMS IDs enabled by the Activation Object

attributeId: 1.2.840.113556.1.4.2082

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 1

rangeLower: 16

rangeUpper: 16

schemaIdGuid:: 2j5mm0I11kad8DFAJa8rrA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-SPP-CSVLK-Pid,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msSPP-CSVLKPid

```
adminDisplayName: ms-SPP-CSVLK-Pid
adminDescription: ID of CSVLK product-key used to create the Activation Object
attributeId: 1.2.840.113556.1.4.2105
attributeSyntax: 2.5.5.12
omSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeUpper: 512
schemaIdGuid:: DVF/tFBr4Ue1VncseeT/xA==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-SPP-CSVLK-Sku-Id,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msSPP-CSVLKSkuId
adminDisplayName: ms-SPP-CSVLK-Sku-Id
adminDescription: SKU ID of CSVLK product-key used to create the Activation Object
attributeId: 1.2.840.113556.1.4.2081
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeLower: 16
rangeUpper: 16
schemaIdGuid:: OfeElnh7bUeNdDGtdpLu9A==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-SPP-Phone-License,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
```

```
ldapDisplayName: msSPP-PhoneLicense
adminDisplayName: ms-SPP-Phone-License
adminDescription: License used during phone activation of the Active Directory forest
attributeId: 1.2.840.113556.1.4.2086
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeUpper: 5242880
schemaIdGuid:: EtnkZ2LzUkCMeUL0W6eyIQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-SPP-Config-License,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msSPP-ConfigLicense
adminDisplayName: ms-SPP-Config-License
adminDescription: Product-key configuration license used during online/phone activation
of the Active Directory forest
attributeId: 1.2.840.113556.1.4.2087
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeUpper: 5242880
schemaIdGuid:: tcRTA5nRsECzxd6zL9nsBg==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-SPP-Online-License,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
```

objectClass: attributeSchema

ldapDisplayName: msSPP-OnlineLicense

adminDisplayName: ms-SPP-Online-License

adminDescription: License used during online activation of the Active Directory forest

attributeId: 1.2.840.113556.1.4.2085

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 5242880

schemaIdGuid:: jjaPCRJIzUivt6E2uWgH7Q==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-SPP-Confirmation-Id,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msSPP-ConfirmationId

adminDisplayName: ms-SPP-Confirmation-Id

adminDescription: Confirmation ID (CID) used for phone activation of the Active Directory forest

attributeId: 1.2.840.113556.1.4.2084

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 512

schemaIdGuid:: xJeHbtqsSUqHQLC9Bam4MQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-SPP-Installation-Id,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msSPP-InstallationId

adminDisplayName: ms-SPP-Installation-Id

adminDescription: Installation ID (IID) used for phone activation of the Active Directory forest

attributeId: 1.2.840.113556.1.4.2083

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 512

schemaIdGuid:: FLG/aXtAOUeiE8ZjgCs+Nw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-SPP-Issuance-License,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msSPP-IssuanceLicense

adminDisplayName: ms-SPP-Issuance-License

adminDescription: Issuance license used during online/phone activation of the Active Directory forest

attributeId: 1.2.840.113556.1.4.2088

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 5242880

schemaIdGuid:: obN1EK+70kmujcTyXIIzAw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

```
dn: CN=ms-SPP-CSVLK-Partial-Product-Key,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msSPP-CSVLKPartialProductKey
adminDisplayName: ms-SPP-CSVLK-Partial-Product-Key
adminDescription: Last 5 characters of CSVLK product-key used to create the Activation
Object
attributeId: 1.2.840.113556.1.4.2106
attributeSyntax: 2.5.5.12
omSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeLower: 5
rangeUpper: 5
schemaIdGuid:: kbABplKGOkWzhoetI5t8CA==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-TPM-Srk-Pub-Thumbprint,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msTPM-SrkPubThumbprint
adminDisplayName: TPM-SrkPubThumbprint
adminDescription: This attribute contains the thumbprint of the SrkPub corresponding to a
particular TPM. This helps to index the TPM devices in the directory.
attributeId: 1.2.840.113556.1.4.2107
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 11
rangeUpper: 20
```

schemaIdGuid:: 6wbXGXZNokSF1hw0K+O+Nw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-TPM-Owner-Information-Temp,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msTPM-OwnerInformationTemp

adminDisplayName: TPM-OwnerInformationTemp

adminDescription: This attribute contains temporary owner information for a particular TPM.

attributeId: 1.2.840.113556.1.4.2108

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 640

rangeUpper: 128

schemaIdGuid:: nYCUyBO1+E+IEfT0P1rHvA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-TPM-Tpm-Information-For-Computer,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msTPM-TpmInformationForComputer

adminDisplayName: TPM-TpmInformationForComputer

adminDescription: This attribute links a Computer object to a TPM object.

attributeId: 1.2.840.113556.1.4.2109

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 16

```
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: k3sb6khe1Ua8bE30/aeKNQ==
linkID: 2182
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-TPM-Tpm-Information-For-Computer-BL,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msTPM-TpmInformationForComputerBL
adminDisplayName: TPM-TpmInformationForComputerBL
adminDescription: This attribute links a TPM object to the Computer objects associated
with it.
attributeId: 1.2.840.113556.1.4.2110
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: FALSE
systemOnly: TRUE
searchFlags: 0
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: yYT6FM2OSEO8kW087Ucqtw==
linkID: 2183
showInAdvancedViewOnly: TRUE
systemFlags: 17

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn: CN=ms-DS-Claim-Types,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
```

ldapDisplayName: msDS-ClaimTypes

adminDisplayName: ms-DS-Claim-Types

adminDescription: A container of this class can contain claim type objects.

governsId: 1.2.840.113556.1.5.270

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: NTIJNhXHIUirarVvsoBaWA==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Claim-Types,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Resource-Property-List,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ResourcePropertyList

adminDisplayName: ms-DS-Resource-Property-List

adminDescription: An object of this class contains a list of resource properties.

governsId: 1.2.840.113556.1.5.274

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMayContain: 1.2.840.113556.1.4.2103

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: etTjckKzRU2PVrr/gDyr+Q==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Resource-Property-List,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Resource-Properties,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ResourceProperties

adminDisplayName: ms-DS-Resource-Properties

adminDescription: A container of this class can contain resource properties.

governsId: 1.2.840.113556.1.5.271

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: hEVKelCzj0es1rS4UtgswA==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Resource-Properties,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Claim-Type-Property-Base,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ClaimTypePropertyBase

adminDisplayName: ms-DS-Claim-Type-Property-Base

adminDescription: An abstract class that defines the base class for claim type or
resource property classes.

governsId: 1.2.840.113556.1.5.269

objectClassCategory: 2

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMayContain: 1.2.840.113556.1.4.2101

systemMayContain: 1.2.840.113556.1.2.557

systemMayContain: 1.2.840.113556.1.4.2097

schemaIdGuid:: WC9EuJDEh0SKndgLiDJxrQ==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Claim-Type-Property-Base,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Resource-Property,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ResourceProperty

adminDisplayName: ms-DS-Resource-Property

adminDescription: An instance of this class holds the definition of a property on
resources.

governsId: 1.2.840.113556.1.5.273

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 1.2.840.113556.1.5.269

systemMayContain: 1.2.840.113556.1.4.2095

systemPossSuperiors: 1.2.840.113556.1.5.271

schemaIdGuid:: Xj0oWwSElUGTOYRQGIxQGg==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Resource-Property,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Claim-Type,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ClaimType

adminDisplayName: ms-DS-Claim-Type

adminDescription: An instance of this class holds the definition of a claim type that can be defined on security principals.

governsId: 1.2.840.113556.1.5.272

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 1.2.840.113556.1.5.269

systemMayContain: 1.2.840.113556.1.4.2100

systemMayContain: 1.2.840.113556.1.4.2099

systemPossSuperiors: 1.2.840.113556.1.5.270

schemaIdGuid:: fIWjgWlUj02q5sJ2mXYmBA==

defaultSecurityDescriptor: D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Claim-Type,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-SPP-Activation-Objects-Container,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msSPP-ActivationObjectsContainer

adminDisplayName: ms-SPP-Activation-Objects-Container

```
adminDescription: Container for Activation Objects used by Active Directory based
activation
governsId: 1.2.840.113556.1.5.266
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemPossSuperiors: 1.2.840.113556.1.3.23
schemaIdGuid:: K4YvtyW7XU2qUWLFm9+Qrg==
defaultSecurityDescriptor: O:BAG:BAD:
(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPLCLORC;;;AU)
showInAdvancedViewOnly: FALSE
defaultHidingValue: TRUE
systemOnly: FALSE
defaultObjectCategory: CN=ms-SPP-Activation-Objects-
Container,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=ms-SPP-Activation-Object,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msSPP-ActivationObject
adminDisplayName: ms-SPP-Activation-Object
adminDescription: Activation Object used in Active Directory based activation
governsId: 1.2.840.113556.1.5.267
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemMustContain: 1.2.840.113556.1.4.2082
systemMustContain: 1.2.840.113556.1.4.2081
systemMustContain: 1.2.840.113556.1.4.2106
systemMustContain: 1.2.840.113556.1.4.2105
systemMayContain: 1.2.840.113556.1.4.2088
systemMayContain: 1.2.840.113556.1.4.2087
systemMayContain: 1.2.840.113556.1.4.2086
```

systemMayContain: 1.2.840.113556.1.4.2085

systemMayContain: 1.2.840.113556.1.4.2084

systemMayContain: 1.2.840.113556.1.4.2083

systemPossSuperiors: 1.2.840.113556.1.5.266

schemaIdGuid:: jOagUcUNykOTXcHJEb8u5Q==

defaultSecurityDescriptor: O:BAG:BAD:
(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPLCLORC;;;AU)

showInAdvancedViewOnly: FALSE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-SPP-Activation-Object,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-TPM-Information-Objects-Container,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msTPM-InformationObjectsContainer

adminDisplayName: TPM-InformationObjectsContainer

adminDescription: Container for TPM objects.

governsId: 1.2.840.113556.1.5.276

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMustContain: 2.5.4.3

systemPossSuperiors: 1.2.840.113556.1.5.67

systemPossSuperiors: 1.2.840.113556.1.5.66

schemaIdGuid:: vagn4FZk3kWQozhZOHfudA==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;LOLCCCRP;;;
DC)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

```
defaultObjectCategory: CN=ms-TPM-Information-Objects-
Container,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=ms-TPM-Information-Object,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msTPM-InformationObject
adminDisplayName: TPM-InformationObject
adminDescription: This class contains recovery information for a Trusted Platform Module
(TPM) device.
governsId: 1.2.840.113556.1.5.275
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemMustContain: 1.2.840.113556.1.4.1966
systemMayContain: 1.2.840.113556.1.4.2108
systemMayContain: 1.2.840.113556.1.4.2107
systemPossSuperiors: 1.2.840.113556.1.5.276
schemaIdGuid:: alsEhaZHQ0KnzGiQcB9mLA==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLO;;;DC)(
A;;WP;;;CO)
showInAdvancedViewOnly: TRUE
defaultHidingValue: TRUE
systemOnly: FALSE
defaultObjectCategory: CN=ms-TPM-Information-Object,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=Top,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2102
systemMayContain: 1.2.840.113556.1.4.2104
```

```
-

dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2109
-

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 48
-
```

Sch49.ldf

```
dn: CN=ms-DNS-Is-Signed,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-IsSigned
adminDisplayName: ms-DNS-Is-Signed
adminDescription: An attribute used to define whether or not the DNS zone is signed.
attributeId: 1.2.840.113556.1.4.2130
attributeSyntax: 2.5.5.8
omSyntax: 1
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 8
schemaIdGuid:: TIUSqvzYXk2RyjaLjYKb7g==
```

```
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DNS-NSEC3-OptOut,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-NSEC3OptOut
adminDisplayName: ms-DNS-NSEC3-OptOut
adminDescription: An attribute used to define whether or not the DNS zone should be
signed using NSEC opt-out.
attributeId: 1.2.840.113556.1.4.2132
attributeSyntax: 2.5.5.8
omSyntax: 1
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 8
schemaIdGuid:: iCDqe+KMPEKxkWbsUGsVlQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DNS-Signing-Keys,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-SigningKeys
adminDisplayName: ms-DNS-Signing-Keys
adminDescription: An attribute that contains the set of encrypted DNSSEC signing keys
used by the DNS server to sign the DNS zone.
attributeId: 1.2.840.113556.1.4.2144
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: FALSE
systemOnly: FALSE
searchFlags: 8
rangeUpper: 10000
```

schemaIdGuid:: bT5nt9nKnk6zGmPoCY/dYw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-Sign-With-NSEC3,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-SignWithNSEC3

adminDisplayName: ms-DNS-Sign-With-NSEC3

adminDescription: An attribute used to define whether or not the DNS zone is signed with NSEC3.

attributeId: 1.2.840.113556.1.4.2131

attributeSyntax: 2.5.5.8

omSyntax: 1

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: mSGfx6Ft/0aSPB8/gAxyHg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-NSEC3-User-Salt,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-NSEC3UserSalt

adminDisplayName: ms-DNS-NSEC3-User-Salt

adminDescription: An attribute that defines a user-specified NSEC3 salt string to use when signing the DNS zone. If empty, random salt will be used.

attributeId: 1.2.840.113556.1.4.2148

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 510

schemaIdGuid:: cGfxryKWvE+hKDCId3YFuQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-DNSKEY-Records,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-DNSKEYRecords

adminDisplayName: ms-DNS-DNSKEY-Records

adminDescription: An attribute that contains the DNSKEY record set for the root of the DNS zone and the root key signing key signature records.

attributeId: 1.2.840.113556.1.4.2145

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 8

rangeUpper: 10000

schemaIdGuid:: 9VjEKC1gyUqnfLPxvlA6fg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-DS-Record-Set-TTL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-DSRecordSetTTL

adminDisplayName: ms-DNS-DS-Record-Set-TTL

adminDescription: An attribute that defines the time-to-live (TTL) value assigned to DS records when signing the DNS zone.

attributeId: 1.2.840.113556.1.4.2140

attributeSyntax: 2.5.5.9

omSyntax: 2

```
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 8
rangeLower: 0
rangeUpper: 2592000
schemaIdGuid:: fJuGKcRk/kKX1fvC+hJBYA==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DNS-Keymaster-Zones,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-KeymasterZones
adminDisplayName: ms-DNS-Keymaster-Zones
adminDescription: A list of Active Directory-integrated zones for which the DNS server is
the keymaster.
attributeId: 1.2.840.113556.1.4.2128
attributeSyntax: 2.5.5.12
omSyntax: 64
isSingleValued: FALSE
systemOnly: FALSE
searchFlags: 0
schemaIdGuid:: O93gCxoEjEGs6S8X0j6dQg==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DNS-NSEC3-Iterations,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-NSEC3Iterations
adminDisplayName: ms-DNS-NSEC3-Iterations
adminDescription: An attribute that defines how many NSEC3 hash iterations to perform
when signing the DNS zone.
attributeId: 1.2.840.113556.1.4.2138
```

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 10000

schemaIdGuid:: qwq3gFmJwE6OkxJudt86yg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-Propagation-Time,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-PropagationTime

adminDisplayName: ms-DNS-Propagation-Time

adminDescription: An attribute used to define in seconds the expected time required to propagate zone changes through Active Directory.

attributeId: 1.2.840.113556.1.4.2147

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: Rw00uoEhoEyi9vrkR52rKg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-NSEC3-Current-Salt,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-NSEC3CurrentSalt

adminDisplayName: ms-DNS-NSEC3-Current-Salt

adminDescription: An attribute that defines the current NSEC3 salt string being used to sign the DNS zone.

attributeId: 1.2.840.113556.1.4.2149

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 510

schemaIdGuid:: MpR9ONGmdESCzQqJquCErg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-RFC5011-Key-Rollovers,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-RFC5011KeyRollovers

adminDisplayName: ms-DNS-RFC5011-Key-Rollovers

adminDescription: An attribute that defines whether or not the DNS zone should be maintained using key rollover procedures defined in RFC 5011.

attributeId: 1.2.840.113556.1.4.2135

attributeSyntax: 2.5.5.8

omSyntax: 1

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: QDzZJ1oGwEO92M3yx9Egqg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-NSEC3-Hash-Algorithm,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-NSEC3HashAlgorithm

adminDisplayName: ms-DNS-NSEC3-Hash-Algorithm

adminDescription: An attribute that defines the NSEC3 hash algorithm to use when signing the DNS zone.

attributeId: 1.2.840.113556.1.4.2136

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: UlWe/7d9OEGIiAXOMgoDIw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-DS-Record-Algorithms,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-DSRecordAlgorithms

adminDisplayName: ms-DNS-DS-Record-Algorithms

adminDescription: An attribute used to define the algorithms used when writing the dsset file during zone signing.

attributeId: 1.2.840.113556.1.4.2134

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: 0npbXPogu0S+szS5wPZVeQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-DNSKEY-Record-Set-TTL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-DNSKEYRecordSetTTL

adminDisplayName: ms-DNS-DNSKEY-Record-Set-TTL

adminDescription: An attribute that defines the time-to-live (TTL) value assigned to DNSKEY records when signing the DNS zone.

attributeId: 1.2.840.113556.1.4.2139

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 2592000

schemaIdGuid:: fzFOj9coLESm3x9JH5ezJg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-Maintain-Trust-Anchor,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-MaintainTrustAnchor

adminDisplayName: ms-DNS-Maintain-Trust-Anchor

adminDescription: An attribute used to define the type of trust anchor to automatically publish in the forest-wide trust anchor store when the DNS zone is signed.

attributeId: 1.2.840.113556.1.4.2133

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

schemaIdGuid:: wWPADdlSVkSeFZwkNKr9lA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-NSEC3-Random-Salt-Length,CN=Schema,CN=Configuration,DC=X

```
changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-NSEC3RandomSaltLength

adminDisplayName: ms-DNS-NSEC3-Random-Salt-Length

adminDescription: An attribute that defines the length in bytes of the random salt used
when signing the DNS zone.

attributeId: 1.2.840.113556.1.4.2137

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 255

schemaIdGuid:: ZRY2E2yR502lnbHrvQ3hKQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16


dn: CN=ms-DNS-Signing-Key-Descriptors,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-SigningKeyDescriptors

adminDisplayName: ms-DNS-Signing-Key-Descriptors

adminDescription: An attribute that contains the set of DNSSEC Signing Key Descriptors
(SKDs) used by the DNS server to generate keys and sign the DNS zone.

attributeId: 1.2.840.113556.1.4.2143

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 8

rangeUpper: 10000

schemaIdGuid:: zdhDNLblO0+wmGWaAhSgeQ==

showInAdvancedViewOnly: TRUE
```

```
systemFlags: 16

dn: CN=ms-DNS-Signature-Inception-Offset,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-SignatureInceptionOffset
adminDisplayName: ms-DNS-Signature-Inception-Offset
adminDescription: An attribute that defines in seconds how far in the past DNSSEC
signature validity periods should begin when signing the DNS zone.
attributeId: 1.2.840.113556.1.4.2141
attributeSyntax: 2.5.5.9
omSyntax: 2
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 8
rangeLower: 0
rangeUpper: 2592000
schemaIdGuid:: LsPUAxfiYUqWmXu8RymgJg==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DNS-Parent-Has-Secure-Delegation,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDNS-ParentHasSecureDelegation
adminDisplayName: ms-DNS-Parent-Has-Secure-Delegation
adminDescription: An attribute used to define whether the parental delegation to the DNS
zone is secure.
attributeId: 1.2.840.113556.1.4.2146
attributeSyntax: 2.5.5.8
omSyntax: 1
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 8
```

schemaIdGuid:: ZGlcKBrBnkmW2L98daIjxg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DNS-Secure-Delegation-Polling-Period,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDNS-SecureDelegationPollingPeriod

adminDisplayName: ms-DNS-Secure-Delegation-Polling-Period

adminDescription: An attribute that defines in seconds the time between polling attempts for child zone key rollovers.

attributeId: 1.2.840.113556.1.4.2142

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 8

rangeLower: 0

rangeUpper: 2592000

schemaIdGuid:: vvCw9uSoaESP2cPEe4ci+Q==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Member-Rules-In-Central-Access-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-MemberRulesInCentralAccessPolicy

adminDisplayName: ms-Authz-Member-Rules-In-Central-Access-Policy

adminDescription: For a central access policy, this attribute identifies the central access rules that comprise the policy.

attributeId: 1.2.840.113556.1.4.2155

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: ei/yV343w0KYcs7G8h0uPg==

linkID: 2184

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Member-Rules-In-Central-Access-Policy-BL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-MemberRulesInCentralAccessPolicyBL

adminDisplayName: ms-Authz-Member-Rules-In-Central-Access-Policy-BL

adminDescription: Backlink for ms-Authz-Member-Rules-In-Central-Access-Policy. For a central access rule object, this attribute references one or more central access policies that point to it.

attributeId: 1.2.840.113556.1.4.2156

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: z2duUd3+lES7OrxQapSIkQ==

linkID: 2185

showInAdvancedViewOnly: TRUE

systemFlags: 17

dn: CN=ms-DS-Claim-Source,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimSource

adminDisplayName: ms-DS-Claim-Source

adminDescription: For a claim type, this attribute indicates the source of the claim type. For example, the source can be certificate.

attributeId: 1.2.840.113556.1.4.2157

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: pvIy+ovy0Ee/kWY+j5EKcg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Proposed-Security-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-ProposedSecurityPolicy

adminDisplayName: ms-Authz-Proposed-Security-Policy

adminDescription: For a Central Access Policy Entry, defines the proposed security policy of the objects the CAPE is applied to.

attributeId: 1.2.840.113556.1.4.2151

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: zr5GubUJakuyWktjozDoDg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Source-Type,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimSourceType

adminDisplayName: ms-DS-Claim-Source-Type

adminDescription: For a security principal claim type, lists the type of store the issued claim is sourced from

attributeId: 1.2.840.113556.1.4.2158

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: BZzxkvqNIkK70SxPAUh3VA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Effective-Security-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-EffectiveSecurityPolicy

adminDisplayName: ms-Authz-Security-Policy

adminDescription: For a central access rule, this attribute defines the permission that is applying to the target resources on the central access rule.

attributeId: 1.2.840.113556.1.4.2150

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: GRmDB5SPtk+KQpFUXcza0w==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Is-Single-Valued,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimIsSingleValued

adminDisplayName: ms-DS-Claim-Is-Single-Valued

adminDescription: For a claim type object, this attribute identifies if the claim type or resource property can only contain single value.

attributeId: 1.2.840.113556.1.4.2160

attributeSyntax: 2.5.5.8

omSyntax: 1

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 0

schemaIdGuid:: uZ94zbSWSEaCGco3gWGvOA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Last-Effective-Security-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-LastEffectiveSecurityPolicy

adminDisplayName: ms-Authz-Last-Effective-Security-Policy

adminDescription: For a Central Access Policy Entry, defines the security policy that was last applied to the objects the CAPE is applied to.

attributeId: 1.2.840.113556.1.4.2152

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: xoUWji8+okiljVrw6nifoA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Resource-Condition,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-ResourceCondition

adminDisplayName: ms-Authz-Resource-Condition

adminDescription: For a central access rule, this attribute is an expression that identifies the scope of the target resource to which the policy applies.

attributeId: 1.2.840.113556.1.4.2153

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: d3iZgHT4aEyGTW5QioO9vQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Claim-Is-Value-Space-Restricted,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ClaimIsValueSpaceRestricted

adminDisplayName: ms-DS-Claim-Is-Value-Space-Restricted

adminDescription: For a claim type, this attribute identifies whether a user can input values other than those described in the msDS-ClaimPossibleValues in applications.

attributeId: 1.2.840.113556.1.4.2159

attributeSyntax: 2.5.5.8

omSyntax: 1

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: x+QsDMPxgkSFeMYNS7dEIg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Authz-Central-Access-Policy-ID,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msAuthz-CentralAccessPolicyID

adminDisplayName: ms-Authz-Central-Access-Policy-ID

adminDescription: For a Central Access Policy, this attribute defines a GUID that can be used to identify the set of policies when applied to a resource.

attributeId: 1.2.840.113556.1.4.2154

attributeSyntax: 2.5.5.17

omSyntax: 4

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: YJvyYnS+MEaUVi9mkZk6hg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Generation-Id,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-GenerationId

adminDisplayName: ms-DS-Generation-Id

adminDescription: For virtual machine snapshot resuming detection. This attribute represents the VM Generation ID.

attributeId: 1.2.840.113556.1.4.2166

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

rangeLower: 16

rangeUpper: 16

systemOnly: TRUE

searchFlags: 0

schemaIdGuid:: PTldHreMT0uECpc7NswJww==

showInAdvancedViewOnly: TRUE

systemFlags: 17

dn: CN=ms-DS-Claim-Shares-Possible-Values-With,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaModify

replace: adminDescription

```
adminDescription: For a claim type object, indicates that the possible values of the
claims issued are defined on the object this linked attribute points to; overrides msDS-
ClaimPossibleValues, msDS-ClaimValueType, and msDS-ClaimIsValueSpaceRestricted, if
populated.
-
replace: isSingleValued
isSingleValued: TRUE
-

dn: CN=ms-DNS-Server-Settings,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msDNS-ServerSettings
adminDisplayName: ms-DNS-Server-Settings
adminDescription: A container for storing DNS server settings.
governsId: 1.2.840.113556.1.4.2129
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemMayContain: 1.2.840.113556.1.4.2128
systemPossSuperiors: 1.2.840.113556.1.5.17
schemaIdGuid:: 7cMv7xhuW0GZ5DEUqMsSSw==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLOLORCWOWDSDDTDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLOR
C;;;AU)
showInAdvancedViewOnly: TRUE
defaultHidingValue: FALSE
systemOnly: FALSE
defaultObjectCategory: CN=ms-DNS-Server-Settings,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=ms-Authz-Central-Access-Policies,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msAuthz-CentralAccessPolicies
```

adminDisplayName: ms-Authz-Central-Access-Policies

adminDescription: A container of this class can contain Central Access Policy objects.

governsId: 1.2.840.113556.1.4.2161

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: wyFcVTahWkWTl3lrvTWOJQ==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;
AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-Authz-Central-Access-
Policies,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-Authz-Central-Access-Rules,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msAuthz-CentralAccessRules

adminDisplayName: ms-Authz-Central-Access-Rules

adminDescription: A container of this class can contain Central Access Policy Entry
objects.

governsId: 1.2.840.113556.1.4.2162

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: ehu7mW1gi0+ADuFb5VTKjQ==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;
AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-Authz-Central-Access-Rules,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-Authz-Central-Access-Rule,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msAuthz-CentralAccessRule

adminDisplayName: ms-Authz-Central-Access-Rule

adminDescription: A class that defines Central Access Rules used to construct a central access policy.

governsId: 1.2.840.113556.1.4.2163

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMayContain: 1.2.840.113556.1.4.2153

systemMayContain: 1.2.840.113556.1.4.2152

systemMayContain: 1.2.840.113556.1.4.2151

systemMayContain: 1.2.840.113556.1.4.2150

systemMayContain: 1.2.840.113556.1.2.557

systemPossSuperiors: 1.2.840.113556.1.4.2162

schemaIdGuid:: 3AZKWxwl206IEwvdcTJyJg==

defaultSecurityDescriptor: D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;AU)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-Authz-Central-Access-Rule,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-Authz-Central-Access-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

```
objectClass: classSchema
ldapDisplayName: msAuthz-CentralAccessPolicy
adminDisplayName: ms-Authz-Central-Access-Policy
adminDescription: A class that defines Central Access Policy objects.
governsId: 1.2.840.113556.1.4.2164
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemMayContain: 1.2.840.113556.1.4.2155
systemMayContain: 1.2.840.113556.1.4.2154
systemPossSuperiors: 1.2.840.113556.1.4.2161
schemaIdGuid:: sJxnpZ1vLEOLdR4+g08Cqg==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;
AU)
showInAdvancedViewOnly: TRUE
defaultHidingValue: FALSE
systemOnly: FALSE
defaultObjectCategory: CN=ms-Authz-Central-Access-Policy,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=ms-DS-Claim-Types,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
replace: defaultHidingValue
defaultHidingValue: TRUE
-

dn: CN=ms-DS-Resource-Properties,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
replace: defaultHidingValue
defaultHidingValue: TRUE
-

dn: CN=ms-DS-List-Of-Claim-Types,CN=Schema,CN=Configuration,DC=X
```

```
changetype: ntdsSchemaModify
replace: defaultHidingValue
defaultHidingValue: TRUE
-

dn: CN=ms-DS-Claim-Type,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2157
systemMayContain: 1.2.840.113556.1.4.2158
systemMayContain: 1.2.840.113556.1.4.2098
systemMayContain: 1.2.840.113556.1.4.2159
systemMayContain: 1.2.840.113556.1.4.2160
-

dn: CN=Dns-Zone,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2130
systemMayContain: 1.2.840.113556.1.4.2131
systemMayContain: 1.2.840.113556.1.4.2132
systemMayContain: 1.2.840.113556.1.4.2133
systemMayContain: 1.2.840.113556.1.4.2134
systemMayContain: 1.2.840.113556.1.4.2135
systemMayContain: 1.2.840.113556.1.4.2136
systemMayContain: 1.2.840.113556.1.4.2137
systemMayContain: 1.2.840.113556.1.4.2138
systemMayContain: 1.2.840.113556.1.4.2139
systemMayContain: 1.2.840.113556.1.4.2140
systemMayContain: 1.2.840.113556.1.4.2141
systemMayContain: 1.2.840.113556.1.4.2142
systemMayContain: 1.2.840.113556.1.4.2143
systemMayContain: 1.2.840.113556.1.4.2144
systemMayContain: 1.2.840.113556.1.4.2145
```

```
systemMayContain: 1.2.840.113556.1.4.2146
systemMayContain: 1.2.840.113556.1.4.2147
systemMayContain: 1.2.840.113556.1.4.2148
systemMayContain: 1.2.840.113556.1.4.2149
-

dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2166
-

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn: CN=DS-Clone-Domain-Controller,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: controlAccessRight
displayName: Allow a DC to create a clone of itself
rightsGuid: 3e0f7e18-2c7a-4c10-ba82-4d926db99a3e
appliesTo: 19195a5b-6da0-11d0-afd3-00c04fd930c9
validAccesses: 256
localizationDisplayId: 80

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 49
-
```

Sch50.ldf

```
dn: CN=ms-DS-Allowed-To-Act-On-Behalf-Of-Other-Identity,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-AllowedToActOnBehalfOfOtherIdentity

adminDisplayName: ms-DS-Allowed-To-Act-On-Behalf-Of-Other-Identity

adminDescription: This attribute is used for access checks to determine if a requester
has permission to act on the behalf of other identities to services running as this
account.

attributeId: 1.2.840.113556.1.4.2182

attributeSyntax: 2.5.5.15

omSyntax: 66

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 0

rangeLower: 0

rangeUpper: 132096

schemaIdGuid:: 5cN4P5r3vUaguJ0YEW3ceQ==

attributeSecurityGuid:: hri1d0qU0RGuvQAA+ANnwQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16


dn: CN=ms-Kds-Version,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msKds-Version

adminDisplayName: ms-Kds-Version

adminDescription: Version number of this root key.

attributeId: 1.2.840.113556.1.4.2176

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 640
```

```
schemaIdGuid:: QHPw1bDmSh6Xvg0zGL2dsQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-DomainID,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-DomainID
adminDisplayName: ms-Kds-DomainID
adminDescription: Distinguished name of the Domain Controller which generated this root
key.
attributeId: 1.2.840.113556.1.4.2177
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: ggRAlgfPTOmQ6PLvxPBJXg==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-KDF-Param,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-KDFParam
adminDisplayName: ms-Kds-KDF-Param
adminDescription: Parameters for the key derivation algorithm.
attributeId: 1.2.840.113556.1.4.2170
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
```

```
rangeUpper: 2000
schemaIdGuid:: cgeAirj0TxW0HC5Cce/3pw==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-CreateTime,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-CreateTime
adminDisplayName: ms-Kds-CreateTime
adminDescription: The time when this root key was created.
attributeId: 1.2.840.113556.1.4.2179
attributeSyntax: 2.5.5.16
omSyntax: 65
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
schemaIdGuid:: nxEYrpBjRQCzLZfbxwGu9w==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-RootKeyData,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-RootKeyData
adminDisplayName: ms-Kds-RootKeyData
adminDescription: Root key.
attributeId: 1.2.840.113556.1.4.2175
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
rangeUpper: 128
```

schemaIdGuid:: J3xiJqIIQAqhsY3OhbQpkw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Primary-Computer,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-PrimaryComputer

adminDisplayName: ms-DS-Primary-Computer

adminDescription: For a user or group object, identifies the primary computers.

attributeId: 1.2.840.113556.1.4.2167

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 1

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: 4vQ9obDb60yCi4suFD6egQ==

linkID: 2186

showInAdvancedViewOnly: TRUE

isMemberOfPartialAttributeSet: TRUE

systemFlags: 16

dn: CN=ms-Kds-UseStartTime,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msKds-UseStartTime

adminDisplayName: ms-Kds-UseStartTime

adminDescription: The time after which this root key may be used.

attributeId: 1.2.840.113556.1.4.2178

attributeSyntax: 2.5.5.16

omSyntax: 65

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 640

schemaIdGuid:: fwTcbCL1SreanNlayM39og==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Imaging-Hash-Algorithm,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msImaging-HashAlgorithm

adminDisplayName: ms-Imaging-Hash-Algorithm

adminDescription: Contains the name of the hash algorithm used to create the Thumbprint Hash for the Scan Repository/Secure Print Device.

attributeId: 1.2.840.113556.1.4.2181

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 64

schemaIdGuid:: tQ3nigZklkGS/vO7VXUgpw==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Kds-KDF-AlgorithmID,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msKds-KDFAlgorithmID

adminDisplayName: ms-Kds-KDF-AlgorithmID

adminDescription: The algorithm name of the key derivation function used to compute keys.

attributeId: 1.2.840.113556.1.4.2169

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 640

rangeUpper: 200

schemaIdGuid:: skgs203RTuyfWK1XnYtEDg==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Imaging-Thumbprint-Hash,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msImaging-ThumbprintHash

adminDisplayName: ms-Imaging-Thumbprint-Hash

adminDescription: Contains a hash of the security certificate for the Scan Repository/Secure Print Device.

attributeId: 1.2.840.113556.1.4.2180

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

rangeUpper: 1024

schemaIdGuid:: xdvfnAQDaUWV9sT2Y/5a5g==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-Kds-PublicKey-Length,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msKds-PublicKeyLength

adminDisplayName: ms-Kds-PublicKey-Length

adminDescription: The length of the secret agreement public key.

attributeId: 1.2.840.113556.1.4.2173

attributeSyntax: 2.5.5.9

omSyntax: 2

isSingleValued: TRUE

```
systemOnly: FALSE
searchFlags: 640
schemaIdGuid:: cPQ44805SUWrW/afnlg/4A==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-PrivateKey-Length,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-PrivateKeyLength
adminDisplayName: ms-Kds-PrivateKey-Length
adminDescription: The length of the secret agreement private key.
attributeId: 1.2.840.113556.1.4.2174
attributeSyntax: 2.5.5.9
omSyntax: 2
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
schemaIdGuid:: oUJfYec3SBGg3TAH4Jz8gQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-Is-Primary-Computer-For,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-IsPrimaryComputerFor
adminDisplayName: ms-DS-Is-Primary-Computer-For
adminDescription: Backlink atribute for msDS-IsPrimaryComputer.
attributeId: 1.2.840.113556.1.4.2168
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: FALSE
systemOnly: FALSE
searchFlags: 0
```

```
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: rAaMmYc/TkSl3xGwPcilDA==
linkID: 2187
showInAdvancedViewOnly: TRUE
systemFlags: 17

dn: CN=ms-Kds-SecretAgreement-Param,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-SecretAgreementParam
adminDisplayName: ms-Kds-SecretAgreement-Param
adminDescription: The parameters for the secret agreement algorithm.
attributeId: 1.2.840.113556.1.4.2172
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 640
rangeUpper: 2000
schemaIdGuid:: MLCZ2e3+dUm4B+ukRNp56Q==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-Kds-SecretAgreement-AlgorithmID,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msKds-SecretAgreementAlgorithmID
adminDisplayName: ms-Kds-SecretAgreement-AlgorithmID
adminDescription: The name of the secret agreement algorithm to be used with public keys.
attributeId: 1.2.840.113556.1.4.2171
attributeSyntax: 2.5.5.12
omSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
```

searchFlags: 640

rangeUpper: 200

schemaIdGuid:: XZcCF14iSsuxXQ2uqLXpkA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Value-Type-Reference,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ValueTypeReference

adminDisplayName: ms-DS-Value-Type-Reference

adminDescription: This attribute is used to link a resource property object to its value type.

attributeId: 1.2.840.113556.1.4.2187

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: hF38eNzBSDGJhFj3ktQdPg==

linkID: 2188

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Value-Type-Reference-BL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ValueTypeReferenceBL

adminDisplayName: ms-DS-Value-Type-Reference-BL

adminDescription: This is the back link for ms-DS-Value-Type-Reference. It links a value type object back to resource properties.

attributeId: 1.2.840.113556.1.4.2188

attributeSyntax: 2.5.5.1

```
omSyntax: 127
isSingleValued: FALSE
systemOnly: TRUE
searchFlags: 0
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: rUNVq6EjRTu5N5sxPVR0qA==
linkID: 2189
showInAdvancedViewOnly: TRUE
systemFlags: 17

dn: CN=ms-DS-Is-Possible-Values-Present,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-IsPossibleValuesPresent
adminDisplayName: ms-DS-Is-Possible-Values-Present
adminDescription: This attribute identifies if ms-DS-Claim-Possible-Values on linked
resource property must have value or must not have value.
attributeId: 1.2.840.113556.1.4.2186
attributeSyntax: 2.5.5.8
omSyntax: 1
isSingleValued: TRUE
systemOnly: TRUE
searchFlags: 0
schemaIdGuid:: 2tyrb1OMTyCxpJ3wxnwetA==
showInAdvancedViewOnly: TRUE
systemFlags: 16


dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-
```

```
dn: CN=ms-Kds-Prov-RootKey,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msKds-ProvRootKey
adminDisplayName: ms-Kds-Prov-RootKey
adminDescription: Root keys for the Group Key Distribution Service.
governsId: 1.2.840.113556.1.5.278
objectClassCategory: 1
rdnAttId: 2.5.4.3
subClassOf: 2.5.6.0
systemMustContain: 1.2.840.113556.1.4.2179
systemMustContain: 1.2.840.113556.1.4.2175
systemMustContain: 1.2.840.113556.1.4.2174
systemMustContain: 1.2.840.113556.1.4.2173
systemMustContain: 1.2.840.113556.1.4.2171
systemMustContain: 1.2.840.113556.1.4.2169
systemMustContain: 1.2.840.113556.1.4.2178
systemMustContain: 1.2.840.113556.1.4.2177
systemMustContain: 1.2.840.113556.1.4.2176
systemMustContain: 2.5.4.3
systemMayContain: 1.2.840.113556.1.4.2172
systemMayContain: 1.2.840.113556.1.4.2170
systemPossSuperiors: 1.2.840.113556.1.3.23
schemaIdGuid:: Qf0CquAXGE+Gh7Ijlklzaw==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPWPCRCCDCL
CLORCWOWDSDDTSW;;;DA)
showInAdvancedViewOnly: TRUE
defaultHidingValue: TRUE
systemOnly: FALSE
defaultObjectCategory: CN=ms-Kds-Prov-RootKey,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=ms-Kds-Prov-ServerConfiguration,CN=Schema,CN=Configuration,DC=X
```

```
changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msKds-ProvServerConfiguration

adminDisplayName: ms-Kds-Prov-ServerConfiguration

adminDescription: Configuration for the Group Key Distribution Service.

governsId: 1.2.840.113556.1.5.277

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMustContain: 1.2.840.113556.1.4.2176

systemMayContain: 1.2.840.113556.1.4.2174

systemMayContain: 1.2.840.113556.1.4.2173

systemMayContain: 1.2.840.113556.1.4.2172

systemMayContain: 1.2.840.113556.1.4.2171

systemMayContain: 1.2.840.113556.1.4.2170

systemMayContain: 1.2.840.113556.1.4.2169

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: qEPyXiUqpkWLcwinGuZ3zg==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPWPCRCCDCL
CLORCWOWDSDDTSW;;;DA)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-Kds-Prov-ServerConfiguration,CN=Schema,CN=Configuration,DC=X

systemFlags: 16


dn: CN=Top,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaModify

add: systemMayContain

systemMayContain: 1.2.840.113556.1.4.2168

systemMayContain: 1.2.840.113556.1.4.2188

-
```

```
dn: CN=Group,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2167
-

dn: CN=User,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2167
-

dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2180
systemMayContain: 1.2.840.113556.1.4.2181
-

dn: CN=Organizational-Person,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2182
-

dn: CN=ms-DS-Resource-Property,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMustContain
systemMustContain: 1.2.840.113556.1.4.2187
-

dn: CN=ms-DS-Value-Type,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
```

ldapDisplayName: msDS-ValueType

adminDisplayName: ms-DS-Value-Type

adminDescription: An value type object holds value type information for a resource property.

governsId: 1.2.840.113556.1.5.279

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMustContain: 1.2.840.113556.1.4.2186

systemMustContain: 1.2.840.113556.1.4.2160

systemMustContain: 1.2.840.113556.1.4.2159

systemMustContain: 1.2.840.113556.1.4.2098

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: 33/C4x2wTk+H5wVu7w65Ig==

defaultSecurityDescriptor: D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Value-Type,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn:

changetype: modify

add: schemaUpdateNow

schemaUpdateNow: 1

-

dn: CN=Validated-MS-DS-Behavior-Version,CN=Extended-Rights,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: controlAccessRight

rightsGuid: d31a8757-2447-4545-8081-3bb610cacbf2

```
appliesTo: f0f8ffab-1191-11d0-a060-00aa006c33ed
displayName: Validated write to MS DS behavior version
localizationDisplayId: 81
validAccesses: 8
showInAdvancedViewOnly: TRUE


dn: CN=Validated-MS-DS-Additional-DNS-Host-Name,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: controlAccessRight
rightsGuid: 80863791-dbe9-4eb8-837e-7f0ab55d9ac7
appliesTo: bf967a86-0de6-11d0-a285-00aa003049e2
displayName: Validated write to MS DS Additional DNS Host Name
localizationDisplayId: 82
validAccesses: 8
showInAdvancedViewOnly: TRUE


dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 50
-
```

Sch51.ldf

```
dn: CN=ms-DS-Transformation-Rules,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-TransformationRules
adminDisplayName: ms-DS-Transformation-Rules
adminDescription: Specifies the Transformation Rules for Across-Forest Claims
Transformation.
attributeId: 1.2.840.113556.1.4.2189
attributeSyntax: 2.5.5.12
```

omSyntax: 64

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: cSuHVbLESDuuUUCV+R7GAA==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Applies-To-Resource-Types,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-AppliesToResourceTypes

adminDisplayName: ms-DS-Applies-To-Resource-Types

adminDescription: For a resource property, this attribute indicates what resource types this resource property applies to.

attributeId: 1.2.840.113556.1.4.2195

attributeSyntax: 2.5.5.12

omSyntax: 64

isSingleValued: FALSE

systemOnly: FALSE

searchFlags: 0

schemaIdGuid:: BiA/aWRXSj2EOVjwSqtLWQ==

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Transformation-Rules-Compiled,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-TransformationRulesCompiled

adminDisplayName: ms-DS-Transformation-Rules-Compiled

adminDescription: Blob containing compiled transformation rules.

attributeId: 1.2.840.113556.1.4.2190

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 128

schemaIdGuid:: EJq0C2tTTbyicwurDdS9EA==

showInAdvancedViewOnly: TRUE

systemFlags: 17

dn: CN=ms-DS-Egress-Claims-Transformation-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-EgressClaimsTransformationPolicy

adminDisplayName: ms-DS-Egress-Claims-Transformation-Policy

adminDescription: This is a link to a Claims Transformation Policy Object for the egress claims (claims leaving this forest) to the Trusted Domain. This is applicable only for an incoming or bidirectional Across-Forest Trust. When this link is not present, all claims are allowed to egress as-is.

attributeId: 1.2.840.113556.1.4.2192

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: fkI3wXOaQLCRkBsJW7QyiA==

linkID: 2192

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Ingress-Claims-Transformation-Policy,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-IngressClaimsTransformationPolicy

adminDisplayName: ms-DS-Ingress-Claims-Transformation-Policy

adminDescription: This is a link to a Claims Transformation Policy Object for the ingress claims (claims entering this forest) from the Trusted Domain. This is applicable only for an outgoing or bidirectional Across-Forest Trust. If this link is absent, all the ingress claims are dropped.

attributeId: 1.2.840.113556.1.4.2191

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: TRUE

systemOnly: FALSE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: CEwohm4MQBWLFXUUfSPSDQ==

linkID: 2190

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-TDO-Egress-BL,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-TDOEgressBL

adminDisplayName: ms-DS-TDO-Egress-BL

adminDescription: Backlink to TDO Egress rules link on object.

attributeId: 1.2.840.113556.1.4.2194

attributeSyntax: 2.5.5.1

omSyntax: 127

isSingleValued: FALSE

systemOnly: TRUE

searchFlags: 0

omObjectClass:: KwwCh3McAIVK

schemaIdGuid:: KWIA1ROZQiKLF4N2HR4OWw==

linkID: 2193

showInAdvancedViewOnly: TRUE

systemFlags: 17

```
dn: CN=ms-DS-TDO-Ingress-BL,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-TDOIngressBL
adminDisplayName: ms-DS-TDO-Ingress-BL
adminDescription: Backlink to TDO Ingress rules link on object.
attributeId: 1.2.840.113556.1.4.2193
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: FALSE
systemOnly: TRUE
searchFlags: 0
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: oWFWWsaXS1SAVuQw/nvFVA==
linkID: 2191
showInAdvancedViewOnly: TRUE
systemFlags: 17

dn: CN=ms-DS-ManagedPassword,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-ManagedPassword
adminDisplayName: msDS-ManagedPassword
adminDescription: This attribute is the managed password data for a group MSA.
attributeId: 1.2.840.113556.1.4.2196
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
schemaIdGuid:: hu1i4yi3QgiyfS3qep3yGA==
showInAdvancedViewOnly: TRUE
systemFlags: 20
```

```
dn: CN=ms-DS-ManagedPasswordId,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-ManagedPasswordId
adminDisplayName: msDS-ManagedPasswordId
adminDescription: This attribute is the identifier for the current managed password data
for a group MSA.
attributeId: 1.2.840.113556.1.4.2197
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
systemOnly: TRUE
searchFlags: 0
rangeUpper: 1024
schemaIdGuid:: Wil4DtPGQAq0kdYiUf+gpg==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-GroupMSAMembership,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-GroupMSAMembership
adminDisplayName: msDS-GroupMSAMembership
adminDescription: This attribute is used for access checks to determine if a requester
has permission to retrieve the password for a group MSA.
attributeId: 1.2.840.113556.1.4.2200
attributeSyntax: 2.5.5.15
omSyntax: 66
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 0
rangeUpper: 132096
schemaIdGuid:: 1u2OiATOQN+0YrilDkG6OA==
showInAdvancedViewOnly: TRUE
```

```
systemFlags: 16

dn: CN=ms-DS-GeoCoordinates-Altitude,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-GeoCoordinatesAltitude
adminDisplayName: ms-DS-GeoCoordinates-Altitude
adminDescription: ms-DS-GeoCoordinates-Altitude
attributeId: 1.2.840.113556.1.4.2183
attributeSyntax: 2.5.5.16
omSyntax: 65
isSingleValued: TRUE
searchFlags: 1
schemaIdGuid:: twMXoUFWnE2GPl+zMl504A==
attributeSecurityGuid:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-GeoCoordinates-Latitude,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-GeoCoordinatesLatitude
adminDisplayName: ms-DS-GeoCoordinates-Latitude
adminDescription: ms-DS-GeoCoordinates-Latitude
attributeId: 1.2.840.113556.1.4.2184
attributeSyntax: 2.5.5.16
omSyntax: 65
isSingleValued: TRUE
searchFlags: 1
schemaIdGuid:: TtRm3EM99UCFxTwS4WmSfg==
attributeSecurityGuid:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16
```

```
dn: CN=ms-DS-GeoCoordinates-Longitude,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-GeoCoordinatesLongitude
adminDisplayName: ms-DS-GeoCoordinates-Longitude
adminDescription: ms-DS-GeoCoordinates-Longitude
attributeId: 1.2.840.113556.1.4.2185
attributeSyntax: 2.5.5.16
omSyntax: 65
isSingleValued: TRUE
searchFlags: 1
schemaIdGuid:: ECHElOS66kyFd6+BOvXaJQ==
attributeSecurityGuid:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-ManagedPasswordInterval,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-ManagedPasswordInterval
adminDisplayName: msDS-ManagedPasswordInterval
adminDescription: This attribute is used to retrieve the number of days before a managed
password is automatically changed for a group MSA.
attributeId: 1.2.840.113556.1.4.2199
attributeSyntax: 2.5.5.9
omSyntax: 2
isSingleValued: TRUE
systemOnly: TRUE
searchFlags: 0
schemaIdGuid:: 9451+HasQ4ii7qJrTcr0CQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-ManagedPasswordPreviousId,CN=Schema,CN=Configuration,DC=X
```

```
changetype: ntdsSchemaAdd

objectClass: attributeSchema

ldapDisplayName: msDS-ManagedPasswordPreviousId

adminDisplayName: msDS-ManagedPasswordPreviousId

adminDescription: This attribute is the identifier for the previous managed password data
for a group MSA.

attributeId: 1.2.840.113556.1.4.2198

attributeSyntax: 2.5.5.10

omSyntax: 4

isSingleValued: TRUE

systemOnly: TRUE

searchFlags: 0

rangeUpper: 1024

schemaIdGuid:: MSHW0EotT9CZ2RxjZGIppA==

showInAdvancedViewOnly: TRUE

systemFlags: 16



dn:

changetype: modify

add: schemaUpdateNow

schemaUpdateNow: 1

-


dn: CN=ms-DS-Claims-Transformation-Policies,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ClaimsTransformationPolicies

adminDisplayName: ms-DS-Claims-Transformation-Policies

adminDescription: An object of this class holds the one set of Claims Transformation
Policy for Across-Forest Claims Transformation.

governsId: 1.2.840.113556.1.5.281

objectClassCategory: 1

rdnAttId: 2.5.4.3
```

subClassOf: 2.5.6.0

systemPossSuperiors: 1.2.840.113556.1.3.23

schemaIdGuid:: san8yIh9T7uCekSJJ3EHYg==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Claims-Transformation-
Policies,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

dn: CN=ms-DS-Claims-Transformation-Policy-Type,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-ClaimsTransformationPolicyType

adminDisplayName: ms-DS-Claims-Transformation-Policy-Type

adminDescription: An object of this class holds the one set of Claims Transformation
Policy for Across-Forest Claims Transformation.

governsId: 1.2.840.113556.1.5.280

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

systemMayContain: 1.2.840.113556.1.4.2190

systemMayContain: 1.2.840.113556.1.4.2189

systemPossSuperiors: 1.2.840.113556.1.5.281

schemaIdGuid:: s2LrLnMTRf6BATh/Fnbtxw==

defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;EA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)

showInAdvancedViewOnly: TRUE

defaultHidingValue: TRUE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Claims-Transformation-Policy-
Type,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

```
dn: CN=Top,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2193
systemMayContain: 1.2.840.113556.1.4.2194
-

dn: CN=Trusted-Domain,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2191
systemMayContain: 1.2.840.113556.1.4.2192
-

dn: CN=ms-DS-Resource-Property,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2195
-

dn: CN=Mail-Recipient,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: mayContain
mayContain: 1.2.840.113556.1.4.2183
mayContain: 1.2.840.113556.1.4.2184
mayContain: 1.2.840.113556.1.4.2185
-

dn: CN=ms-DS-Group-Managed-Service-Account,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: classSchema
ldapDisplayName: msDS-GroupManagedServiceAccount
adminDisplayName: msDS-Group-Managed-Service-Account
```

adminDescription: The group managed service account class is used to create an account which can be shared by different computers to run Windows services.

governsId: 1.2.840.113556.1.5.282

objectClassCategory: 1

rdnAttId: 2.5.4.3

subClassOf: 1.2.840.113556.1.3.30

systemMustContain: 1.2.840.113556.1.4.2199

systemMayContain: 1.2.840.113556.1.4.2200

systemMayContain: 1.2.840.113556.1.4.2198

systemMayContain: 1.2.840.113556.1.4.2197

systemMayContain: 1.2.840.113556.1.4.2196

systemPossSuperiors: 1.2.840.113556.1.3.30

systemPossSuperiors: 1.2.840.113556.1.3.23

systemPossSuperiors: 2.5.6.5

systemPossSuperiors: 1.2.840.113556.1.5.67

schemaIdGuid:: ilWLe6WT90qtysAX5n8QVw==

defaultSecurityDescriptor: D:(OD;;CR;00299570-246d-11d0-a768-00aa006e0529;;WD)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;AO)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPCRLCLORCSDDT;;;CO)(OA;;WP;4c164200-20c0-11d0-a768-00aa006e0529;;CO)(OA;;SW;72e39547-7b18-11d1-adef-00c04fd8d5cd;;CO)(OA;;SW;f3a64788-5306-11d1-a9c5-0000f80367c1;;CO)(OA;;WP;3e0abfd0-126a-11d0-a060-00aa006c33ed;bf967a86-0de6-11d0-a285-00aa003049e2;CO)(OA;;WP;5f202010-79a5-11d0-9020-00c04fc2d4cf;bf967a86-0de6-11d0-a285-00aa003049e2;CO)(OA;;WP;bf967950-0de6-11d0-a285-00aa003049e2;bf967a86-0de6-11d0-a285-00aa003049e2;CO)(OA;;WP;bf967953-0de6-11d0-a285-00aa003049e2;bf967a86-0de6-11d0-a285-00aa003049e2;CO)(OA;;SW;f3a64788-5306-11d1-a9c5-0000f80367c1;;PS)(OA;;RPWP;77B5B886-944A-11d1-AEBD-0000F80367C1;;PS)(OA;;SW;72e39547-7b18-11d1-adef-00c04fd8d5cd;;PS)(A;;RPLCLORC;;;AU)(OA;;RPWP;bf967a7f-0de6-11d0-a285-00aa003049e2;;CA)(OA;;RP;46a9b11d-60ae-405a-b7e8-ff8a58d456d2;;S-1-5-32-560)(OA;;RP;e362ed86-b728-0842-b27d-2dea7a9df218;;WD)

showInAdvancedViewOnly: TRUE

defaultHidingValue: FALSE

systemOnly: FALSE

defaultObjectCategory: CN=ms-DS-Group-Managed-Service-Account,CN=Schema,CN=Configuration,DC=X

systemFlags: 16

```
dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 51
-
```

Sch52.ldf

```
dn: CN=ms-DS-RID-Pool-Allocation-Enabled,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
ldapDisplayName: msDS-RIDPoolAllocationEnabled
adminDisplayName: ms-DS-RID-Pool-Allocation-Enabled
adminDescription: This attribute indicates whether RID pool allocation is enabled or not.
attributeId: 1.2.840.113556.1.4.2213
attributeSyntax: 2.5.5.8
omSyntax: 1
instanceType: 4
isSingleValued: TRUE
systemOnly: TRUE
searchFlags: 0
schemaFlagsEx: 1
schemaIdGuid:: jHyXJLfBQDO09is3XrcR1w==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=RID-Set-References,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
replace: searchFlags
```

```
searchFlags: 8
-

dn: CN=Netboot-DUID,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: Netboot-DUID
ldapDisplayName: netbootDUID
adminDisplayName: Netboot-DUID
adminDescription: This attribute is used to store DHCPv6 DUID device ID.
attributeId: 1.2.840.113556.1.4.2234
attributeSyntax: 2.5.5.10
omSyntax: 4
instanceType: 4
isSingleValued: TRUE
searchFlags: 1
systemFlags: 16
isMemberOfPartialAttributeSet: TRUE
systemOnly: FALSE
rangeLower: 2
rangeUpper: 128
schemaIdGuid:: vXAlU3c9T0KCLw1jbcbarQ==
showInAdvancedViewOnly: TRUE

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

dn: CN=RID-Manager,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2213
```

```
-

dn: CN=domainDNS-Display,CN=409,CN=DisplaySpecifiers,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: adminContextMenu
adminContextMenu: 3,{2fb1b669-59ea-4f64-b728-05309f2c11c8}
-

dn: CN=computer-Display,CN=409,CN=DisplaySpecifiers,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: adminPropertyPages
adminPropertyPages: 13,{2fb1b669-59ea-4f64-b728-05309f2c11c8}
-

dn: CN=Certificate-AutoEnrollment,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: controlAccessRight
showInAdvancedViewOnly: TRUE
appliesTo: e5209ca2-3bba-11d2-90cc-00c04fd91ab1
displayname: AutoEnrollment
localizationDisplayId: 83
rightsGuid: a05b8cc2-17bc-4802-a710-e7c15ab866a2
validAccesses: 256

# Update element: computer
dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: mayContain
mayContain: 1.2.840.113556.1.4.2234
-

dn: CN=ms-DS-cloudExtensionAttribute1,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
```

```
cn: ms-DS-cloudExtensionAttribute1
lDAPDisplayName: msDS-cloudExtensionAttribute1
adminDisplayName: ms-DS-cloudExtensionAttribute1
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2214
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: r+oJl9pJsk2QigRG5eq4RA==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute2,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute2
lDAPDisplayName: msDS-cloudExtensionAttribute2
adminDisplayName: ms-DS-cloudExtensionAttribute2
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2215
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: rOBO88HAqUuCyRqQdS8WpQ==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16
```

```
dn: CN=ms-DS-cloudExtensionAttribute3,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute3
lDAPDisplayName: msDS-cloudExtensionAttribute3
adminDisplayName: ms-DS-cloudExtensionAttribute3
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2216
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: Gsj2gtr6DUqw93BtRoOOtQ==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute4,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute4
lDAPDisplayName: msDS-cloudExtensionAttribute4
adminDisplayName: ms-DS-cloudExtensionAttribute4
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2217
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
```

```
schemaIDGUID:: NzS/nG5OW0iykSKwJVQnPw==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute5,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute5
lDAPDisplayName: msDS-cloudExtensionAttribute5
adminDisplayName: ms-DS-cloudExtensionAttribute5
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2218
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: W+gVKUfjUkiquyLlplHIZA==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute6,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute6
lDAPDisplayName: msDS-cloudExtensionAttribute6
adminDisplayName: ms-DS-cloudExtensionAttribute6
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2219
attributeSyntax: 2.5.5.12
oMSyntax: 64
```

```
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: eSZFYOEo7Eus43EoMzYUVg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute7,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute7
lDAPDisplayName: msDS-cloudExtensionAttribute7
adminDisplayName: ms-DS-cloudExtensionAttribute7
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2220
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: GRN8Sk7jwkCdAGD/eJDyBw==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute8,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute8
lDAPDisplayName: msDS-cloudExtensionAttribute8
adminDisplayName: ms-DS-cloudExtensionAttribute8
```

```
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2221
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: FMXRPEmEykSBwAIXgYANKg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute9,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute9
lDAPDisplayName: msDS-cloudExtensionAttribute9
adminDisplayName: ms-DS-cloudExtensionAttribute9
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2222
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: LOFjCkAwQUSuJs2Vrw0kfg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute10,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
```

```
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute10
lDAPDisplayName: msDS-cloudExtensionAttribute10
adminDisplayName: ms-DS-cloudExtensionAttribute10
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2223
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: s/wKZ70T/EeQswpSftgatw==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute11,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute11
lDAPDisplayName: msDS-cloudExtensionAttribute11
adminDisplayName: ms-DS-cloudExtensionAttribute11
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2224
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: yLuenqV9pkKJJSROEqVuJA==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
```

```
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute12,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute12
lDAPDisplayName: msDS-cloudExtensionAttribute12
adminDisplayName: ms-DS-cloudExtensionAttribute12
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2225
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: PcQBPAvhyk+Sskz2FdWwmg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute13,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute13
lDAPDisplayName: msDS-cloudExtensionAttribute13
adminDisplayName: ms-DS-cloudExtensionAttribute13
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2226
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
```

```
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: S0a+KJCreUumsN9DdDHQNg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute14,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute14
lDAPDisplayName: msDS-cloudExtensionAttribute14
adminDisplayName: ms-DS-cloudExtensionAttribute14
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2227
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: ura8zoBuJ0mFYJj+yghqnw==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute15,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute15
lDAPDisplayName: msDS-cloudExtensionAttribute15
adminDisplayName: ms-DS-cloudExtensionAttribute15
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2228
attributeSyntax: 2.5.5.12
```

```
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: N9XkqvCKqk2cxmLq24T/Aw==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute16,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute16
lDAPDisplayName: msDS-cloudExtensionAttribute16
adminDisplayName: ms-DS-cloudExtensionAttribute16
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2229
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: WyGBlZZRU0ChHm/8r8YsTQ==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute17,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute17
lDAPDisplayName: msDS-cloudExtensionAttribute17
```

```
adminDisplayName: ms-DS-cloudExtensionAttribute17
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2230
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: 2m08PehrKUKWfi/1u5O0zg==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute18,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute18
lDAPDisplayName: msDS-cloudExtensionAttribute18
adminDisplayName: ms-DS-cloudExtensionAttribute18
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2231
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: NDvniKYKaUSYQm6wGzKltQ==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute19,CN=Schema,CN=Configuration,dc=X
```

```
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute19
lDAPDisplayName: msDS-cloudExtensionAttribute19
adminDisplayName: ms-DS-cloudExtensionAttribute19
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2232
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: mf51CQeWikaOGMgA0zhzlQ==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
showInAdvancedViewOnly: TRUE
systemFlags: 16

dn: CN=ms-DS-cloudExtensionAttribute20,CN=Schema,CN=Configuration,dc=X
changetype: ntdsSchemaAdd
objectClass: attributeSchema
cn: ms-DS-cloudExtensionAttribute20
lDAPDisplayName: msDS-cloudExtensionAttribute20
adminDisplayName: ms-DS-cloudExtensionAttribute20
adminDescription: An attribute used to house an arbitrary cloud-relevant string
attributeID: 1.2.840.113556.1.4.2233
attributeSyntax: 2.5.5.12
oMSyntax: 64
isSingleValued: TRUE
systemOnly: FALSE
searchFlags: 1
isMemberOfPartialAttributeSet: TRUE
schemaIDGUID:: KGNE9W6LjUmVqCEXSNWs3A==
attributeSecurityGUID:: hri1d0qU0RGuvQAA+ANnwQ==
```

showInAdvancedViewOnly: TRUE

systemFlags: 16

dn: CN=ms-DS-Cloud-Extensions,CN=Schema,CN=Configuration,DC=X

changetype: ntdsSchemaAdd

objectClass: classSchema

ldapDisplayName: msDS-CloudExtensions

adminDisplayName: ms-DS-Cloud-Extensions

adminDescription: A collection of attributes used to house arbitrary cloud-relevant strings.

governsId: 1.2.840.113556.1.5.283

objectClassCategory: 3

rdnAttId: 2.5.4.3

subClassOf: 2.5.6.0

MayContain: 1.2.840.113556.1.4.2214

MayContain: 1.2.840.113556.1.4.2215

MayContain: 1.2.840.113556.1.4.2216

MayContain: 1.2.840.113556.1.4.2217

MayContain: 1.2.840.113556.1.4.2218

MayContain: 1.2.840.113556.1.4.2219

MayContain: 1.2.840.113556.1.4.2220

MayContain: 1.2.840.113556.1.4.2221

MayContain: 1.2.840.113556.1.4.2222

MayContain: 1.2.840.113556.1.4.2223

MayContain: 1.2.840.113556.1.4.2224

MayContain: 1.2.840.113556.1.4.2225

MayContain: 1.2.840.113556.1.4.2226

MayContain: 1.2.840.113556.1.4.2227

MayContain: 1.2.840.113556.1.4.2228

MayContain: 1.2.840.113556.1.4.2229

MayContain: 1.2.840.113556.1.4.2230

MayContain: 1.2.840.113556.1.4.2231

MayContain: 1.2.840.113556.1.4.2232

MayContain: 1.2.840.113556.1.4.2233

```
schemaIdGuid:: pIceZCaDcUe6LccG3zXjWg==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLCLORC;;;
AU)
showInAdvancedViewOnly: TRUE
defaultHidingValue: TRUE
systemOnly: FALSE
defaultObjectCategory: CN=ms-DS-Cloud-Extensions,CN=Schema,CN=Configuration,DC=X
systemFlags: 16

dn: CN=User,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemAuxiliaryClass
systemAuxiliaryClass: 1.2.840.113556.1.5.283
-

dn: CN=Personal-Information,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: appliesTo
appliesTo: 641E87A4-8326-4771-BA2D-C706DF35E35A
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 52
-
```

Sch53.ldf

```
dn: CN=ms-Authz-Central-Access-Rule,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2156
-
```

```
dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 53
-
```

Sch54.ldf

```
dn: CN=User-Account-Restrictions,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: appliesTo
appliesTo: 7b8b558a-93a5-4af7-adca-c017e67f1057
-

dn: CN=ms-DS-Allowed-To-Act-On-Behalf-Of-Other-Identity,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
replace: attributeSecurityGuid
attributeSecurityGuid:: AEIWTMAg0BGnaACqAG4FKQ==
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 54
-
```

Sch55.ldf

```
dn: CN=DNS-Host-Name-Attributes,CN=Extended-Rights,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: appliesTo
appliesTo: 7b8b558a-93a5-4af7-adca-c017e67f1057
-

dn: CN=Validated-DNS-Host-Name,CN=Extended-Rights,CN=Configuration,DC=X
```

```
changetype: ntdsSchemaModify
add: appliesTo
appliesTo: 7b8b558a-93a5-4af7-adca-c017e67f1057
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 55
-
```

Sch56.ldf

```
# Update element: computer. Remove netboot-DUID from mayContain
dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
delete: mayContain
mayContain: 1.2.840.113556.1.4.2234
-

# Update element: computer. Add netboot-DUID to SystemMayContain
dn: CN=Computer,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaModify
add: systemMayContain
systemMayContain: 1.2.840.113556.1.4.2234
-

dn: CN=Schema,CN=Configuration,DC=X
changeType: ntdsSchemaModify
replace: objectVersion
objectVersion: 56
-
```

Windows Server 2012: Read-Only Domain Controller Updates

There are no changes to adprep /rodcprep in Windows Server 2012.

Windows Server 2012: Domain-Wide Updates

You can review the following set of changes to help understand the schema updates that are performed by adprep /domainprep in Windows Server 2012 and prepare for them. After the operations that are performed by the **domainprep** command in Windows Server 2012 (operations 78, 79, 80, and 81) complete, the **revision** attribute for the CN=ActiveDirectoryUpdate,CN=DomainUpdates,CN=System,DC=ForestRootDomain object is set to **9**:

In Windows Server 2012, Adprep commands run automatically as needed during AD DS installation. They can also be run separately in advance of AD DS installation. For more information, see Running Adprep.exe.

For more information about how to interpret the access control entry (ACE) strings, see ACE strings

For more information about how to interpret the security ID (SID) strings, see SID strings

Operations number and GUID	Description	Attributes	Permissions
Operation 78: {c3c927a6-cc1d-47c0-966b-be8f9b63d991}	Create a new object CN=TPM Devices in the Domain partition.	• Object class: msTPM-InformationObjectsContainer	N/A
Operation 79: {54afcfb9-637a-4251-9f47-4d50e7021211}	Created an access control entry for the TPM service.	N/A	• (OA;CIIO;WP;ea1b7b93-5e48-46d5-bc6c-4df4fda78a35;bf967a86-0de6-11d0-a285-00aa003049e2;PS)
Operation 80: {f4728883-84dd-	Grant "Clone DC"	N/A	• (OA;;CR;3e0f7e18-2c7a-4c10-ba82-

Operations number and GUID	Description	Attributes	Permissions
483c-9897-274f2ebcf11e}	extended right to **Cloneable Domain Controllers** group		4d926db99a3e;;*domain SID*-522)
Operation 81: {ff4f9d27-7157-4cb0-80a9-5d6f2b14c8ff}	Grant ms-DS-Allowed-To-Act-On-Behalf-Of-Other-Identity to Principal Self on all objects.	N/A	• (OA;CIOI;RPWP;3f78c3e5-f79a-46bd-a0b8-9d18116ddc79;;PS)

Windows Server 2012: Forest-Wide Updates

You can review the following set of changes to help understand the schema updates that are performed by adprep /forestprep in Windows Server 2012 and prepare for them. After the operations that are performed by the **forestprep** command in Windows Server 2012 (operations 84-130) are complete, the **revision** attribute for the CN=ActiveDirectoryUpdate,CN=ForestUpdates,CN=Configuration,DC=ForestRootDomain object is set to **11**.

In Windows Server 2012, Adprep commands run automatically as needed during AD DS installation. They can also be run separately in advance of AD DS installation. For more information, see Running Adprep.exe.

For more information about how to interpret the access control entry (ACE) strings, see ACE strings

For more information about how to interpret the security ID (SID) strings, see SID strings

Operation number and GUID	Description	Attributes	Permissions
Operation 84: {4664e973-cb20-4def-b3d5-559d6fe123e0}	Created a new container CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: container	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 85: {2972d92d-a07a-44ac-9cb0-bf243356f345}	Created a new object CN=Claim Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ClaimTypes • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 86: {0x09a49cb3-6c54-4b83-ab20-8370838ba149}	Created a new object CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperties • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 87: {77283e65-ce02-4dc3-8c1e-bf-99b22527c2}	Created a new container CN=Resource Property Lists,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: container • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operati	Created a new object	N/A	Created the following

Operation number and GUID	Description	Attributes	Permissions
on 88: {0afb7f53-96bd-404b-a659-89e65c269420}	CN=Sam-Domain in the Schema partition.		access control entry (ACE) to grant Write Property to Principal Self on the object: • (OA;CIIO;WP;ea1b7b93-5e48-46d5-bc6c-4df4fda78a35;bf967a86-0de6-11d0-a285-00aa003049e2;PS)
Operation 89: {c7f717ef-fdbe-4b4b-8dfc-fa8b839fbcfa}	Created a new object CN=Domain-DNS in the Schema partition.	N/A	Created the following access control entry (ACE) to grant Write Property to Principal Self on the object: • (OA;CIIO;WP;ea1b7b93-5e48-46d5-bc6c-4df4fda78a35;bf967a86-0de6-11d0-a285-00aa003049e2;PS)
Operation 90: {00232167-f3a4-43c6-b503-9acb7a81b01c}	Call back function to upgrade display specifiers.	N/A	N/A
Operation 91: {73a9515b-511c-44d2-822b-444a33d3bd33}	Created a new container CN=Microsoft SPP,CN=Services in the Configuration partition.	• objectClass: container • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operati	Created a new Activation	• objectClass: msSPP-	• (A;;RPLCLORC;;;AU)

Operation number and GUID	Description	Attributes	Permissions
on 92: {e0c60003-2ed7-4fd3-8659-7655a7e79397}	Objects container CN=Activation Objects,CN=Microsoft SPP,CN=Services in the Configuration partition.	ActivationObjectsContainer • showInAdvancedViewOnly: True	• (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 93: {ed0c8ca-80ab-4b6b-ac5a-59b1d317e11f}	Created a new Central Access Policies container CN=Central Access Policies,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msAuthz-CentralAccessPolicies • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 94: {b6a6c19a-afc9-476b-8994-61f5b14b3f05}	Created a new Central Access Policy Entries container CN=Central Access Rules,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msAuthz-CentralAccessRules • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 95: {defc28cd-6cb6-4479-8bcb-aabfb41e9713}	Created a new Group Key Distribution service container CN=Group Key Distribution Service,CN=Services in the Configuration partition.	• objectClass: container • description: The container contains configuration and data for Group Key Distribution Service. • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 96:	Created a new Master Root Keys container	• objectClass: container • description: The	• (A;;RPLCLORC;;;AU) •

Operation number and GUID	Description	Attributes	Permissions
{d6bd96d4-e66b-4a38-9c6b-e976ff58c56d}	CN=Master Root Keys,CN=Group Key Distribution Service,CN=Services in the Configuration partition.	container contains master root keys for Group Key Distribution Service. • showInAdvancedViewOnly: True	(A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 97: {bb8efc40-3090-4fa2-8a3f-7cd1d380e695}	Created a new Server Configuration container CN=Server Configuration,CN=Group Key Distribution Service,CN=Services in the Configuration partition.	• objectClass: container • description: The container contains Group Key Distribution Service configurations. • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 98: {2d6abe1b-4326-489e-920c-76d5337d2dc5}	Created a new Empty server configuration objects container CN=Group Key Distribution Service Server Configuration,CN=Server Configuration,CN=Group Key Distribution Service,CN=Services in the Configuration partition.	• objectClass: msKds-ProvServerConfiguration • description: The configuration of cryptography algorithms used by Group Key Distribution Service. • msKds-Version: 1 • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 99: {6b13dfb5-cecc-4fb8-b28d-0505cea24175}	Created a new Claims Transformation Policies configuration container CN=Claims Transformation Policies,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ClaimsTransformationPolicies • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCDCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
Operation 100: {92e73422-c68b-46c9-b0d5-b55f9c741410}	Created a new Value Types configuration container CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: container • showInAdvancedViewOnly: True	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 101: {c0ad80b4-8e84-4cc4-9163-2f84649bcc42}	Created a new SinglevaluedChoice value type configuration object CN=MS-DS-SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author a resource property. When assigning value to a resource property of this value type, a user can choose only one entry from a list of suggested values. • displayname: Single-valued Choice • msDS-ClaimValueType: 3 • msDS-ClaimIsValueSpaceRestricted: True • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent: True • showInAdvancedViewOnly: True	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 102:	Created a new YesNo value type configuration	• objectClass: msDS-ValueType	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU)

Operation number and GUID	Description	Attributes	Permissions
{992fe1d0-6591-4f24-a163-c820fcb7f308}	object CN=MS-DS-YesNo,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• description: The valid values for this type are Yes or No. • displayname: Yes/No • msDS-ClaimValueType: 6 • msDS-ClaimIsValueSpaceRestricted: False • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent: False • showInAdvancedViewOnly: True	• (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 103: {ede85f960-7061-47bf-b11b-0c0d999595b5}	Created a new Number value type configuration object CN=MS-DS-Number,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that contain a single number. • displayname: Number • msDS-ClaimValueType: 1 • msDS-ClaimIsValueSpaceRestricted: False • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent: False •	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		showInAdvancedViewOnly: True	
Operation 104: {ee0f3271-eb51-414a-bdac-8f9ba6397a39}	Created a new DateTime value type configuration object CN=MS-DS-DateTime,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that are of the date and time format. • displayname: Date Time • msDS-ClaimValueType: 1 • msDS-ClaimIsValueSpaceRestricted: False • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent: False • showInAdvancedViewOnly: True	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 105: {587d52e0-507e-440e-9d67-e6129f33bb68}	Created a new OrderedList value type configuration object CN=MS-DS-OrderedList,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that contain a single choice entry that can be compared to other resource properties of the same type. A user typically chooses the entry from a list of ordered suggested values that are provided	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		by ms-DS-Claim-Possible-Values on the resource properties. • displayname: Ordered List • msDS-ClaimValueType: 1 • msDS-ClaimIsValueSpaceRestricted: True • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent: True • showInAdvancedViewOnly: True	
Operation 106: {}	Created a new Text value type configuration object CN=MS-DS-Text,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that contain a single text entry. • displayname: Text • msDS-ClaimValueType: 3 • msDS-ClaimIsValueSpaceRestricted: False • msDS-ClaimIsSingleValued: True • msDS-IsPossibleValuesPresent	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		: False • showInAdvancedViewOnly: True	
Operation 107: {7f77d431-dd6a-434f-ae-4dce-82928e498f}	Created a new MultivaluedText value type configuration object CN=MS-DS-MultivaluedText,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that can have multiple text entries. • displayname: Multi-valued Text • msDS-ClaimValueType: 3 • msDS-ClaimIsValueSpaceRestricted: False • msDS-ClaimIsSingleValued: False • msDS-IsPossibleValuesPresent: False • showInAdvancedViewOnly: True	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 108: {ba14e1f6-7cd1-4739-804f-57d0ea74edf4}	Created a new MultivaluedChoice value type configuration object CN=MS-DS-MultivaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ValueType • description: You can use this type to author resource properties that can have multiple entries that cannot be compared. A user typically chooses each entry from a list of	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		suggested values that are provided by ms-DS-Claim-Possible-Values on the resource properties. • displayname: Multi-valued Choice • msDS-ClaimValueType: 3 • msDS-ClaimIsValueSpaceRestricted: True • msDS-ClaimIsSingleValued: False • msDS-IsPossibleValuesPresent: True • showInAdvancedViewOnly: True	
Operation 109: {156ffa2a-e07c-46fb-a5c4-fbd84a4e5cce}	Created a new Personally Identifiable Information resource property object CN=PII_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Personally Identifiable Information (PII) property specifies whether the resource contains PII and if it does, what the sensitivity level of that information is. • displayname: Personally Identifiable Information • Enabled: False • msDS-IsUsedAsResourceSecu	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		rityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-OrderedList,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 110: {7771d7dd-2231-4470-aa74-84a6f56fc3b6}	Created a new Protected Health Information resource property object CN=ProtectedHealthInformation_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Protected Health Information (PHI) property specifies whether the resource contains any data related to an individual's medical record or medical payment history. • displayname: Protected Health Information • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-YesNo,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 111: {49b2ae86-839a-	Created a new Required Clearance resource property object CN=RequiredClearance_	• objectClass: msDS-ResourceProperty • description: The Required Clearance	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLO

Operation number and GUID	Description	Attributes	Permissions
4ea0-81fe-9171c1b98e83}	MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	property specifies the level of clearance a user should possess before attempting to access the resource. • displayname: Required Clearance • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-OrderedList,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	CCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 112: {1b1de989-57ec-4e96-b933-8279a8119da4}	Created a new Confidentiality resource property object CN=Confidentiality_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Confidentiality property specifies the level of confidentiality of the resource, and the potential impact of inadvertent access or disclosure. • displayname: Confidentiality • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		OrderedList,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 113: {281c63f0-2c9a-4cce-9256-a238c23c0db9}	Created a new Compliancy resource property object CN=Compliancy_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Compliancy property specifies the compliance frameworks that apply to the resource. • displayname: Compliancy • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-MultivaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 114: {4c47881a-f15a-4f6c-9f49-2742f7a11f4b}	Created a new Discoverability resource property object CN=Discoverability_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Discoverability property specifies whether the resource contains potential evidence that might require disclosure to opposing legal counsel during the course of current or future litigation.	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		• displayname: Discoverability • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 115: {2aea2dc6-d1d3-4f0c-9994-66c1da21de0f}	Created a new Immutable resource property object CN=Immutable_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Immutable property specifies whether a user should be allowed to delete a resource or change its contents. • displayname: Immutable • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-YesNo,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operati	Created a new	• objectClass: msDS-	• (D;;SDDT;;;WD)

Operation number and GUID	Description	Attributes	Permissions
on 116: {ae7824 0c-43b9-499e-ae65-2b6e0f0 e202a}	Intellectual Property resource property object CN=IntellectualProperty_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	ResourceProperty • description: The Intellectual Property (IP) property specifies whether the resource contains IP, and if so, what kind. • displayname: Intellectual Property • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 117: {261b5bba-3438-4d5c-a3e9-7b871e5f57f0}	Created a new Department resource property object CN=Department_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Department property specifies the name of the department to which the resource belongs. • displayname: Department • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 118: {3fb79c05-8ea1-438c-8c7a-81f213aa61c2}	Created a new Impact resource property object CN=Impact_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Department property specifies the name of the department to which the resource belongs. • displayname: Impact • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-OrderedList,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 119: {0b2be39a-d463-4c23-8290-32186759d3b1}	Created a new Personal Use resource property object CN=PersonalUse_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Personal Use property specifies whether the file is for personal use (not business related). • displayname: Personal Use • Enabled: False • msDS-	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-YesNo,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 120: {f0842b44-bc03-46a1-a860-006e8527fccd}	Created a new Project resource property object CN=Project_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Project property specifies the names of one or more projects that are relevant to the resource. • displayname: Project • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-MultivaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 121: {93efec15-4dd9-4850-bc86-a1f2c8e	Created a new Retention Period resource property object CN=RetentionPeriod_MS,CN=Resource Properties,CN=Claims Configuration,CN=Servic	• objectClass: msDS-ResourceProperty • description: The Retention Period property specifies the maximum period for which the file should be retained.	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLO

Operation number and GUID	Description	Attributes	Permissions
2ebb9}	es in the Configuration partition.	• displayname: Retention Period • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	CCDCRCWDWOSDDTSW;;;SY)
Operation 122: {9e108d96-672f-40f0-b6bd-69ee1f0b7ac4}	Created a new Retention Start Date resource property object CN=RetentionStartDate_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Retention Start Date property defines the starting date for a Retention Period. The retention period would begin on the Retention Start Date. • displayname: Retention Start Date • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: False • msDS-ValueTypeReference: CN=MS-DS-DateTime,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		=<forest root domain>	
Operation 123: {1e269508-f862-4c4a-b01f-420d26c4ff8c}	Created a new Company resource property object CN=Company_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Company property specifies which company the resource belongs to. • displayname: Company • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: True • msDS-ValueTypeReference: CN=MS-DS-SinglevaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)
Operation 125: {e1ab17ed-5efb-4691-ad2d-0424592c5755} **Note** Opera	Created a new Folder Usage resource property object CN=FolderUsage_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourceProperty • description: The Folder Usage property specifies the purpose of the folder and the kind of files stored in it. • displayname: Folder Usage • Enabled: False • msDS-IsUsedAsResourceSecurityAttribute: False • msDS-AppliestoResourceTypes: MS-DS-Container	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
tion 124 was deleted.		• msDS-ValueTypeReference: CN=MS-DS-MultivaluedChoice,CN=Value Types,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 126: {0e848bd4-7c70-48f2-b8fc-00fbaa82e360}	Created a new Global Resource Property List configuration object CN=Global Resource Property List,CN=Resource Property Lists,CN=Claims Configuration,CN=Services in the Configuration partition.	• objectClass: msDS-ResourcePropertyList • description: This is a global out of box resource property list that contains all resource properties that can be consumed by applications. • displayname: Folder Usage • showInAdvancedViewOnly: True • msDS-MembersOfResourcePro	• (D;;SDDT;;;WD) • (A;;RPLCLORC;;;AU) • (A;;RPWPCRLCLOCCRCWDWOSW;;;EA) • (A;;RPWPCRLCLOCCDCRCWDWOSDDTSW;;;SY)

Operation number and GUID	Description	Attributes	Permissions
		pertyList: CN=PII_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=ProtectedHealthInformation_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=RequiredClearance_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Confidentiality_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Compliancy_MS,CN=Resource	

Operation number and GUID	Description	Attributes	Permissions
		Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Discoverability_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Immutable_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=IntellectualProperty_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Department_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN	

Operation number and GUID	Description	Attributes	Permissions
		=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Impact_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=PersonalUse_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Project_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=CN=RetentionPeriod_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePro	

Operation number and GUID	Description	Attributes	Permissions
		pertyList: CN=RetentionStartDate_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=Company_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain> • msDS-MembersOfResourcePropertyList: CN=FolderUsage_MS,CN=Resource Properties,CN=Claims Configuration,CN=Services,CN=Configuration,CN=<forest root domain>	
Operation 127: {016f23f7-077d-41fa-a356-de7cfdb01797}	Call back function to upgrade display specifiers.	N/A	N/A
Operation 128: {49c140db-2de3-44c2-	Updated strings for Folder Usage resource property object CN=FolderUsage_MS,CN=Resource	• description: The Folder Usage property specifies the purpose of the folder and the kind of files stored in it.	N/A

Operation number and GUID	Description	Attributes	Permissions
a99a-bab2e6d2ba81}	Properties,CN=Claims Configuration,CN=Services in the Configuration partition.		
Operation 129: {e0b11c80-62c5-47f7-ad0d-3734a71b8312}	Added ACE to grant Principal Self Write Property and Read Property on CN=Sam-Domain object.	N/A	(OA;CIOI;RPWP;3f78c3e5-f79a-46bd-a0b8-9d18116ddc79;;PS)
Operation 130: {2ada1a2d-b02f-4731-b4fe-59f955e24f71}	Added ACE to grant Principal Self Write Property and Read Property on CN=Domain-DNS object.	N/A	(OA;CIOI;RPWP;3f78c3e5-f79a-46bd-a0b8-9d18116ddc79;;PS)

Troubleshooting Domain Controller Deployment

This topic covers detailed methodology on troubleshooting domain controller configuration and deployment.

- Introduction to Troubleshooting
- Troubleshooting Options

Introduction to Troubleshooting

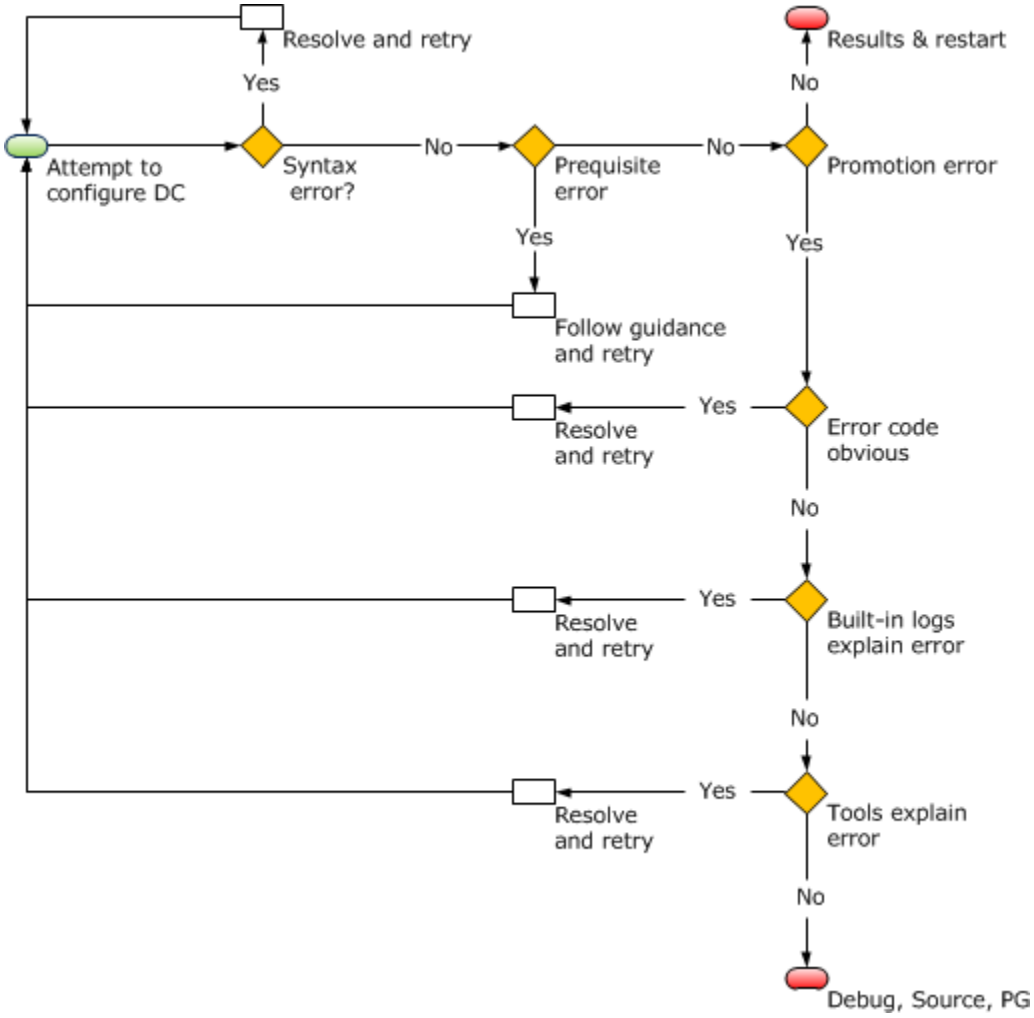

Troubleshooting Options

Logging Options

The built-in logs are the most important instrument for troubleshooting issues with domain controller promotion and demotion. All of these logs are enabled and configured for maximum verbosity by default.

Phase	Log
Server Manager or ADDSDeployment	• %systemroot%\debug\dcpromoui.log

Windows PowerShell operations	• %systemroot%\debug\dcpromoui*.log
Installation/Promotion of the domain controller	• %systemroot%\debug\dcpromo.log • %systemroot%\debug\dcpromo*.log • Event viewer\Windows logs\System • Event viewer\Windows logs\Application • Event viewer\Applications and services logs\Directory Service • Event viewer\Applications and services logs\File Replication Service • Event viewer\Applications and services logs\DFS Replication
Forest or domain upgrade	• %systemroot%\debug\adprep\<datetime>\adprep.log • %systemroot%\debug\adprep\<datetime>\csv.log • %systemroot%\debug\adprep\<datetime>\dspecup.log • %systemroot%\debug\adprep\<datetime>\ldif.log*
Server Manager ADDSDeployment Windows PowerShell deployment engine	• Event viewer\Applications and services logs\Microsoft\Windows\DirectoryServices-Deployment\Operational
Windows Servicing	• %systemroot%\Logs\CBS* • %systemroot%\servicing\sessions\sessions.xml • %systemroot%\winsxs\poqexec.log • %systemroot%\winsxs\pending.xml

Tools and Commands for Troubleshooting Domain Controller Configuration

To troubleshoot issues not explained by the logs, use the following tools as a starting point:

- Dcdiag.exe
- Repadmin.exe
- AutoRuns.exe, Task Manager, and MSInfo32.exe
- Network Monitor 3.4 (or a third party network capture and analysis tool)

General Methodology for Troubleshooting Domain Controller Configuration

1. Did a simple syntax issue cause the error?

a. Did you mistype or forget to provide an argument to ADDSDeployment Windows PowerShell? For example, if using ADDSDeployment Windows PowerShell, did you forget to add required argument **-domainname** with a valid name?
b. Examine the Windows PowerShell console output carefully to see exactly why it is failing to parse the command-line provided.

2. Is the error a prerequisite failure?
   a. Many errors that used to appear as fatal promotion results are now prevented by the prerequisite checker.
   b. Examine the text of the prerequisite errors carefully, they provide the necessary guidance to resolve most issues, as they are controlled scenarios.
3. Is the error in promotion and therefore fatal?
   a. Examine the results carefully: many errors have simple explanations such as bad passwords, network name resolution, or critical offline domain controllers.
   b. Examine the Dcpromoui.log and dcpromo.log for the errors shown in the output, then work backwards from them to see indications of why the failure occurred.
      i. Always compare to a working sample log
      ii. Examine the ADPrep logs for errors only if the results indicate a problem extending the schema or preparing the forest or domain.
      iii. Examine the DirectoryServices-Deployment event log for errors only if the Dcpromoui.log lacks detail or ends arbitrarily due to an unhandled exception in the configuration process.
   c. Examine the Directory Services, System, and Application event logs for other indicators of a configuration issue. Often times, the domain controller promotion is just a symptom of other network misconfiguration that would affect all distributed systems.
   d. Use dcdiag.exe and repadmin.exe to validate the overall forest health and indicate subtle misconfigurations that may prevent further domain controller promotion.
   e. Use AutoRuns.exe, Task Manager, or MSinfo32.exe to examine the computer for third party software that may be interfering.
      i. Remove third party software (do not simply disable the software; that does not prevent drivers loading).
   f. Install NetMon 3.4 on the computer that fails to promote as well the replication partner domain controller and analyze the promotion process with double-sided network captures.
      i. Compare this to your working lab environment to understand what a healthy promotion looks like and where it is failing.
      ii. At this point, the errors are likely with the forest objects, non-default security changes, or the network, and this new domain controller is a victim of misconfigurations in DNS, firewalls, host intrusion protection software, or other outside factors.

Troubleshooting Specific Problems

Events and Error Messages

Domain controller promotion and demotion always returns a code at the end of operation and unlike most programs, do not return zero for success. To see the code at the end of a domain controller configuration, you have several options:

1. When using Server Manager, examine the promotion results in the ten seconds prior to automatic reboot.
2. When using ADDSDeployment Windows PowerShell, examine the promotion results in the ten seconds prior to automatic reboot. Alternatively, choose not to restart automatically on completion. You should add the **Format-List** pipeline to make the output easier to read. For example:

   ```
   Install-addsdomaincontroller <options> -
   norebootoncompletion:$true | format-list
   ```

   Errors in prerequisite validation and verification do not continue on to a reboot, so they are visible in all cases. For example:

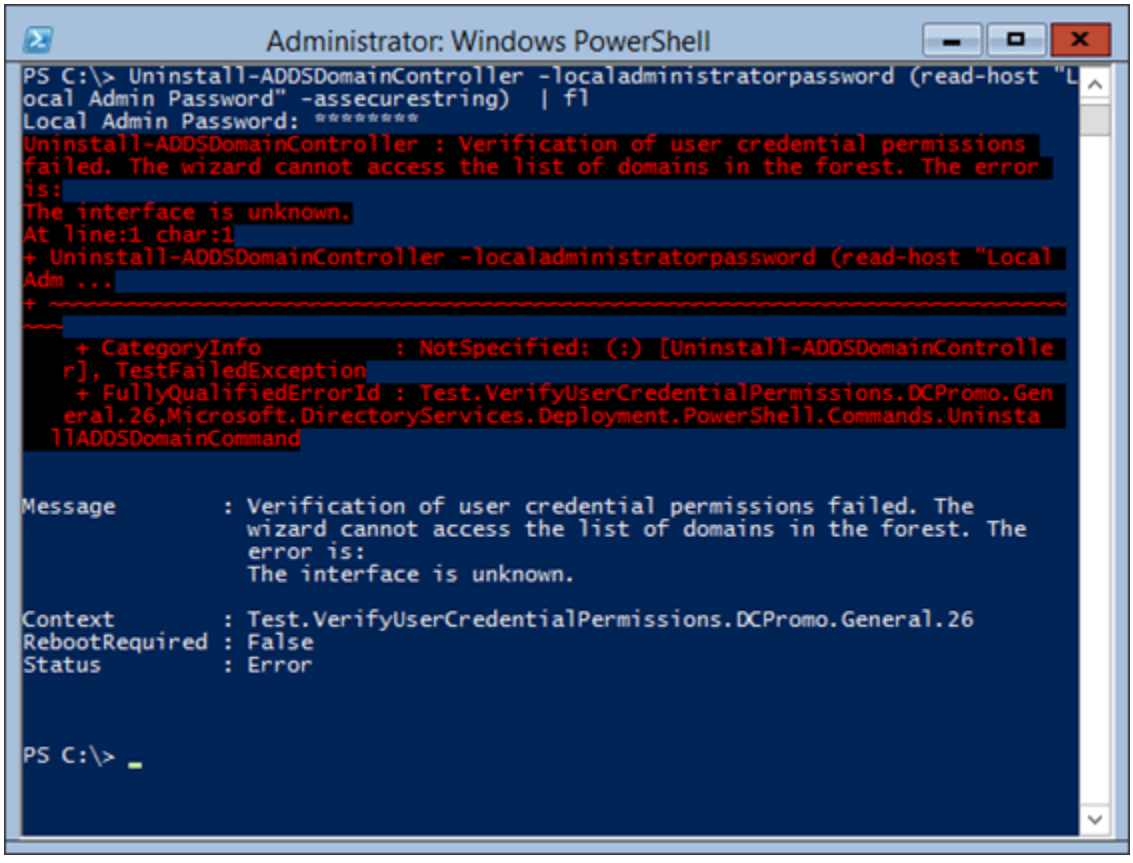
Administrator: Windows PowerShell
PS C:\> Uninstall-ADDSDomainController -localadministratorpassword (read-host "Local Admin Password" -assecurestring) | fl
Local Admin Password: ********
Uninstall-ADDSDomainController : Verification of user credential permissions failed. The wizard cannot access the list of domains in the forest. The error is:
The interface is unknown.
At line:1 char:1
+ Uninstall-ADDSDomainController -localadministratorpassword (read-host "Local Adm ...
+ ~~~

Note

Some of the errors listed below are no longer possible due to operating system and domain controller configuration changes in later operating systems. The new ADDSDeployment Windows PowerShell codes also prevents certain errors, but the dcpromo.exe /unattend does not; this is another compelling reason to switch all of your current automation from the deprecated DCPromo to ADDSDeployment Windows PowerShell.

Promotion and demotion return the following success message codes.

Error Code	Explanation	Note
1	Exit, success	You still must reboot, this just notes that the automatic restart flag was removed
2	Exit, success, need to reboot	
3	Exit, success, with a non-critical failure	Typically seen when returning the DNS Delegation warning. If not configuring DNS delegation, use: -createdns delegation:$false
4	Exit, success, with a non-critical failure, need to reboot	Typically seen when returning the DNS Delegation warning. If not configuring DNS delegation, use: -createdns delegation:$false

Promotion and demotion return the following failure message codes. There is also likely to be an extended error message; always read the entire error carefully, not just the numeric portion.

Error Code	Explanation	Suggested resolution
11	Domain controller promotion is already running	Do not run than one instance of domain controller promotion at the same time for the same target computer
12	User must be administrator	Logon as a member of the built-in Administrators group and ensure you are elevating with UAC
13	Certification Authority is installed	You cannot demote this domain controller, as it is also a Certification Authority. Do not remove the CA before

		you carefully inventory its usage - if it is issuing certificates, removing the role will cause an outage. Running CAs on domain controllers is discouraged
14	Running in safe-boot mode	Boot the server into normal mode
15	Role change is in progress or needs reboot	You must restart the server (due to prior configuration changes) before promotion
16	Running on wrong platform	*Not likely to get this error*
17	No NTFS 5 drives exist	This error is not possible in Windows Server 2012, which requires at least the %systemdrive% be formatted with NTFS
18	Not enough space in windir	Free up space on the %systemdrive% volume using cleanmgr.exe
19	Name change pending, needs reboot	Reboot the server
20	Computer name is invalid syntax	Rename the computer with a valid name
21	This domain controller holds FSMO roles, is a GC, and/or is a DNS server	Add **-demoteoperationmasterrole** when using **-forceremoval**.
22	TCP/IP needs to be installed or isn't functioning	Verify computer has TCP/IP configured, bound, and working correctly
23	DNS client needs to be configured first	Set a primary DNS server when adding a new domain controller to a domain
24	Supplied credentials are invalid or missing required elements	Verify your user name and password is correct
25	Domain controller for the specified domain could not be located	Validate DNS client settings, firewall rules
26	List of domains could not be read from the forest	Validate DNS client settings, LDAP functionality, firewall rules
27	Missing domain name	Specify a domain when promoting or demoting

28	Bad domain name	Choose a different, valid DNS domain name when promoting
29	Parent domain does not exist	Verify the parent domain specified when creating a new child domain or tree domain
30	Domain not in forest	Verify the domain name provided
31	Child Domain already exists	Specify a different domain name
32	Bad NetBIOS domain name	Specify a valid NetBIOS domain name
33	Path to the IFM files is invalid	Validate your path to the Install From Media folder
34	The IFM database is bad	Use the correct Install From Media for this operating system and role (same operating system version, same type of domain controller - RODC versus RWDC)
35	Missing SYSKEY	The Install from Media is encrypted and you must provide a valid SYSKEY to use it
37	Path for NTDS Database or its logs is invalid	Change path of Database and Logs to a fixed NTFS volume, not a mapped drive or UNC path
38	Volume does not have enough space for NTDS database or logs	Free up space using cleanmgr.exe, add more disk space, manually clear space by moving unnecessary data elsewhere
39	Path for SYSVOL is invalid	Change path of SYSVOL folder to a fixed NTFS volume, not a mapped drive or UNC path
40	Invalid site name	Provide a site name that exists
41	Need to specify a password for safe-mode	Provide a password for the DSRM account, it cannot be blank no matter how the password policy is configured
42	Safe-mode password does not meet criteria (promotion only)	Provide a password for the DSRM account that meets the password policy's configured rules
43	Admin password does not meet criteria (demotion only)	Provide a password for the local administrator account that meets the

		password policy's configured rules
44	The specified name for the forest is invalid	Specify a valid forest root DNS domain name
45	A forest with the specified name already exists	Choose a different forest root DNS domain name
46	The specified name for the tree is invalid	Specify a valid tree DNS domain name
47	A tree with the specified name already exists	Choose a different tree DNS domain name
48	The tree name does not fit into the forest structure	Choose a different tree DNS domain name
49	The specified domain does not exist	Verify your typed domain name
50	During demote, last domain controller was detected even though it is not, or last domain controller was specified, but it is not	Do not specify **Last Domain Controller in the Domain** (**-lastdomaincontrollerindomain**) unless it is true. Use **-ignorelastdcindomainmismatch** to override if this is truly the last domain controller and there is phantom domain controller metadata
51	App partitions exist on this domain controller	Specify to **Remove Application Partitions** (**-removeapplicationpartitions**)
52	Required command-line argument is missing (that is, an answer file must be specified on the command-line)	*Only seen with dcpromo /unattend, which is deprecated. See older documentation*
53	The promotion/demotion failed, machine must be rebooted to clean up	Examine the extended error and logs
54	The promotion/demotion failed	Examine the extended error and logs
55	The promotion/demotion was canceled by the user	Examine the extended error and logs
56	The promotion/demotion was canceled by the user, machine must be rebooted to clean up	Examine the extended error and logs

58	A site name must be specified during RODC promotion	You must specify a site for an RODC, it will not automatically detect one like an RWDC
59	During demote, this domain controller is the last DNS server for one of its zones	Specify that this is the **Last DNS Server in the Domain** or use **-ignorelastdnsserverfordomain**
60	A domain controller running Windows Server 2008 or later must be present in the domain in order to promote RODC	Promote at least one Windows Server 2008 or later model writable domain controller
61	You cannot install Active Directory Domain Services with DNS in an existing domain that does not already host DNS	*Not possible to get this error*
62	Answer file does not have a [DCInstall] section	*Only seen with dcpromo /unattend, which is deprecated. See older documentation.*
63	Forest functional level is below windows server 2003	Raise the forest functional level to at least Windows Server 2003 Native. Windows 2000 and Windows NT 4.0 are no longer supported operating systems
64	Promo failed because component binary detection failed	Install the AD DS role
65	Promo failed because component binary installation failed	Install the AD DS role
66	Promo failed because operating system detection failed	Examine the extended error and logs; the server is failing to return its operating system version. It is likely that the computer will need to be re-installed, as its overall health is highly suspect
68	Replication partner is invalid	Use repadmin.exe or the **Get-ADReplication*** Windows PowerShell to validate partner domain controller

		health
69	Required Port is already in use by some other application	Use **netstat.exe -anob** to locate processes that are incorrectly assigned to reserved AD DS ports
70	The forest root domain controller must be a GC	*Only seen with dcpromo /unattend, which is deprecated. See older documentation*
71	DNS server is already installed	Do not specify to install DNS (**-installDNS**) if the DNS service is already installed
72	Computer is running Remote Desktop Services in non-admin mode	You cannot promote this domain controller, as it is also a RDS server configured for more than two admin users. Do not remove RDS before you carefully inventory its usage - if it is being used by applications or end-users, removal will cause an outage
73	The specified forest functional level is invalid.	Specify a valid forest functional level
74	The specified domain functional level is invalid.	Specify a valid domain functional level
75	Unable to determine the default password replication policy.	Validate that the RODC password replication policy exists and is accessible
76	Specified replicated/non-replicated security groups are invalid	Validate that you have typed in valid domain and user accounts when specifying a password replication policy
77	The specified argument is invalid	Examine the extended error and logs
78	Failed to examine Active Directory Forest	Examine the extended error and logs
79	RODC cannot be promoted because rodcprep has not been performed	Use Windows Server 2012 to prepare the forest or use **adprep.exe /rodcprep**
80	Domainprep has not been performed	Use Windows Server 2012 to prepare the domain or use **adprep.exe /domainprep**

81	Forestprep has not been performed	Use Windows Server 2012 to prepare the forest or use **adprep.exe /forestprep**
82	Forest schema mismatch	Use Windows Server 2012 to prepare the forest or use **adprep.exe /forestprep**
83	Unsupported SKU	*Not likely to get this error*
84	Unable to detect a domain controller account	Validate that existing domain controllers have correct user account control attribute set.
85	Unable to select a domain controller account for stage 2	Returned if you specify "Use Existing Account" but either no account found or there is an error during account lookup. Ensure you provided the correct RODC staged account
86	Need to run stage 2 promotion	Returned if you promote an additional domain controller but an existing account exists and "Allow Reinstall" was not specified
87	A domain controller account of conflicting type exists	Rename the computer before promoting, if not trying to attach to an unoccupied domain controller. You must attach to the unoccupied domain controller account using **-useexistingaccount** and the correct read-only or writable argument, depending on account type
88	The specified server admin is not valid	You specified an invalid account for RODC admin delegation. Verify that the account specified is a valid user or group
89	RID master for the specified domain is offline.	Use **netdom.exe query fsmo** to detect the RID master. Bring it online and make it accessible to the domain controller you are promoting
90	Domain naming master is offline.	Use **netdom.exe query fsmo** to detect the domain naming master. Bring it online and make it accessible to the

		domain controller you are promoting
91	Failed to detect if the process is wow64	*Not possible to get this error anymore, the operating system is 64-bit*
92	Wow64 process is not supported	*Not possible to get this error anymore, the operating system is 64-bit*
93	Domain controller service is not running for non-forceful demotion	Start the AD DS service
94	Local admin password does not meet requirement: either blank or not required	Provide a non-blank password and ensure that the local password policy requires a password
95	Cannot demote last Windows Server 2008 or later domain controller in the domain where live RODCs exist	You must first demote all RODCs before you can demote all Windows Server 2008 or later writable domain controllers
96	Unable to uninstall DS binaries	*Only seen with dcpromo /unattend, which is deprecated. See older documentation*
97	Forest functional level version higher than that of the child domain operating system	Provide a child domain functional the same or higher than the forest functional level
98	Component binary install/uninstall is in progress.	*Only seen with dcpromo /unattend, which is deprecated. See older documentation*
99	Forest functional level is too low (error is Windows Server 2012 only)	Raise the forest functional level to at least Windows Server 2003 native. Windows 2000 and Windows NT 4.0 are no longer supported operating systems
100	Domain functional level is too low (error is Windows Server 2012 only)	Raise the domain functional level to at least Windows Server 2003 native. Windows 2000 and Windows NT 4.0 are no longer supported operating systems

Known/Likely Issues and Support Scenarios

The following are common issues seen during the Windows Server 2012 development process. All of these issues are "by design" and have either a valid workaround or more appropriate technique to avoid them in the first place. Many of these behaviors are identical in Windows Server 2008 R2 and older operating systems, but the rewrite of AD DS deployment brings heightened sensitivity to issues.

Issue	Demoting a domain controller leaves DNS running with no zones
Symptoms	Server still responds to DNS requests but has no zone information

Resolution and Notes	When removing the AD DS role, also remove the DNS Server role or set the DNS Server service to disabled. Remember to point the DNS client to another server than itself. If using Windows PowerShell, run the following after you demote the server:
	`uninstall-windowsfeature dns`
	or
	`set-service dns -starttype disabled`
	`stop-service dns`

Issue	Promoting a Windows Server 2012 into an existing single-label domain does not configure updatetopleveldomain=1 or allowsinglelabeldnsdomain=1
Symptoms	DNS dynamic record registration does not occur
Resolution and Notes	Set these values using the Netlogon and DNS group policies. Microsoft began blocking single-label domain creation in Windows Server 2008; you can use ADMT or the Domain Rename Tool to change to an approved DNS domain structure.

Issue	Demotion of last domain controller in a domain fails if there are pre-created, unoccupied RODC accounts

Symptoms	Demotion fails with message: **Dcpromo.General.54** Active Directory Domain Services could not find another Active Directory Domain Controller to transfer the remaining data in directory partition CN=Schema,CN=Configuration,DC=corp,DC=contoso,DC=com. "The format of the specified domain name is invalid."
Resolution and Notes	Remove any remaining pre-created RODC accounts before demoting a domain, using **Dsa.msc** or **Ntdsutil.exe metadata cleanup**.

Issue	Automated forest and domain preparation does not run GPPREP
Symptoms	Cross-domain planning functionality for Group Policy, Resultant Set of Policy (RSOP) Planning Mode, requires updated file system and Active Directory permissions for existing GP. Without Gpprep, you cannot use RSOP Planning across domains.
Resolution and Notes	Run **adprep.exe /gpprep** manually for all domains that were not previously prepared for Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2. Administrators should run GPPrep only once in the history of a domain, not with every upgrade. It is not run by automatic adprep because if you have already set adequate custom permissions, it would cause all SYSVOL contents to re-replicate on all domain controllers.

Issue	Install from media fails to verify when pointing to a UNC path
Symptoms	Error returned: `Could not validate media path. Exception calling "GetDatabaseInfo" with "2" arguments. The folder is not valid.`

Resolution and Notes	You must store IFM files on a local disk, not a remote UNC path. This intentional block prevents partial server promotion due to a network interruption.

Issue	DNS delegation warning shown twice during domain controller promotion
Symptoms	Warning returned *twice* when promoting using ADDSDeployment Windows PowerShell: `"A delegation for this DNS server cannot be created because the authoritative parent zone cannot be found or it does not run Windows DNS server. If you are integrating with an existing DNS infrastructure, you should manually create a delegation to this DNS server in the parent zone to ensure reliable name resolution from outside the domain. Otherwise, no action is required."`
Resolution and Notes	Ignore. ADDSDeployment Windows PowerShell shows the warning first during prerequisite checking, then again during configuration of the domain controller. If you do not wish to configure DNS delegation, use argument: `-creatednsdelegation:$false` Do *not* skip the prerequisite checks in order to suppress this message

Issue	Specifying UPN or non-domain credentials during configuration returns misleading errors

Symptoms	Server Manager returns error:
	`Exception calling "DNSOption" with "6" Arguments`
	ADDSDeployment Windows PowerShell returns error:
	`Verification of user permissions failed. You must supply the name of the domain to which this user account belongs.`

Resolution and Notes	Ensure you are providing valid domain credentials in the form of **domain\user**.

Issue	Removing the DirectoryServices-DomainController role using Dism.exe leads to unbootable server

Symptoms	If using Dism.exe to remove the AD DS role before demoting a domain controller gracefully, the server no longer boots normally and shows error:
	`Status: 0x000000000`
	`Info: An unexpected error has occurred.`

Resolution and Notes	Boot into Directory Services Repair Mode using *Shift+F8*. Add the AD DS role back, and then forcibly demote the domain controller. Alternatively, restore the System State from backup. Do not use Dism.exe for AD DS role removal; the utility has no knowledge of domain controllers.

Issue	Installing a new forest fails when setting forestmode to Win2012
Symptoms	Promotion using ADDSDeployment Windows PowerShell returns error: **`Test.VerifyDcPromoCore.DCPromo.General.74`** `Verification of prerequisites for Domain Controller promotion failed. The specified domain functional level is invalid`
Resolution and Notes	Do not specify a forest functional mode of Win2012 without *also* specifying a domain functional mode of Win2012. Here is an example that will work without errors: `-forestmode Win2012 -domainmode Win2012`

Issue	Clicking Verify in the Install from Media selection area appears to do nothing
Symptoms	When you specify a path to an IFM folder, clicking the **Verify** button never returns a message or appears to do anything.
Resolution and Notes	The **Verify** button only returns errors if there are issues. Otherwise, it makes the **Next** button selectable if you have provided an IFM path. You must click **Verify** to proceed if you have selected IFM.

Issue	Demoting with Server Manager does not provide feedback until completed.
Symptoms	When using Server Manager to remove the AD DS role and demote a domain controller, there is no ongoing feedback given until the demotion completes or fails.
Resolution and Notes	This is a limitation of Server Manager. For feedback, use ADDSDeployment Windows PowerShell cmdlet: `Uninstall-addsdomaincontroller`

Issue	Install from Media Verify does not detect that RODC media provided for writable domain controller, or vice versa.

Symptoms	When promoting a new domain controller using IFM and providing incorrect media to IFM - such as RODC media for a writable domain controller, or RWDC media for an RODC - the Verify button does not return an error. Later, promotion fails with error:
	`An error occurred while trying to configure this machine as a domain controller.`
	`The Install-From-Media promotion of a Read-Only DC cannot start because the specified source database is not allowed. Only databases from other RODCs can be used for`

	`IFM promotion of a RODC.`
Resolution and Notes	Verify only validates the overall integrity of IFM. Do not provide the wrong IFM type to a server. Restart the server before you attempt promotion again with the correct media.

Issue	Promoting an RODC into a pre-created computer account fails

Symptoms	When using ADDSDeployment Windows PowerShell to promote a new RODC with a staged computer account, receive error:
	`Parameter set cannot be resolved using the specified named parameters.`
	`InvalidArgument: ParameterBindingException`
	`+ FullyQualifiedErrorId : AmbiguousParameterSet,Microsoft.DirectoryServices.Deployment.PowerShell.Commands.Install`

Resolution and Notes	Do not provide parameters already defined already on a pre-created RODC account. These include:
	`-readonlyreplica`
	`-installdns`
	`-donotconfigureglobalcatalog`
	`-sitename`
	`-installdns`

Issue	Deselecting/selecting "Restart each destination server automatically if required" does nothing
Symptoms	If selecting (or not selecting) the Server Manager option **Restart each destination server automatically if required** when demoting a domain controller through role removal, the server always restarts, regardless of choice.
Resolution and Notes	This is intentional. The demotion process restarts the server regardless of this setting.

Issue	Dcpromo.log shows "[error] setting security on

	server files failed with 2"
Symptoms	Demotion of a domain controller completes without issues, but examination of the dcpromo log shows error: `[error] setting security on server files failed with 2`
Resolution and Notes	Ignore, error is expected and cosmetic.

Issue	Prerequisite adprep check fails with error "Unable to perform Exchange schema conflict check"

Symptoms	When attempting to promote a Windows Server 2012 domain controller into an existing Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2 forest, prerequisite check fails with error:
	`Verification of prerequisites for AD prep failed. Unable to perform Exchange schema conflict check for domain`
	`<domain name>`
	`(Exception: the RPC server is unavailable)`
	The adprep.log shows error:
	`Adprep could not retrieve data from the server`
	`<domain controller>`
	`through Windows Management Instrumentation (WMI).`

Resolution and Notes	The new domain controller cannot access WMI through DCOM/RPC protocols against the existing domain controllers. To date, there have been three causes for this: • A firewall rule blocks access to the existing domain controllers • The NETWORK SERVICE account is missing from the "Logon as a service" (SeServiceLogonRight) privilege on the existing domain controllers

	• NTLM is disabled on domain controllers, using security policies described in Introducing the Restriction of NTLM Authentication

Issue	Creating a new AD DS forest always shows DNS warning

Symptoms	When creating a new AD DS forest and creating the DNS zone on the new domain controller for itself, you always receive warning message:
	`An error was detected in the DNS configuration.`
	`None of the DNS servers used by this computer responded within the timeout interval.`
	`(error code 0x000005B4 "ERROR_TIMEOUT")`

Resolution and Notes	Ignore. This warning is intentional on the first domain controller in the root domain of a new forest, in case you intended to point to an existing DNS server and zone.

Issue	Windows PowerShell -whatif argument returns incorrect DNS server information
Symptoms	If you use the **-whatif** argument when configuring a domain controller with implicit or explicit **-installdns:$true**, the resulting output shows: `"DNS Server: No"`
Resolution and Notes	Ignore. DNS is installed and configured correctly.

Issue	After promotion, logon fails with " Not enough storage is available to process this command"
Symptoms	After you promote a new domain controller and then log off and attempt to log on interactively,

	you receive error: `Not enough storage is available to process this command`
Resolution and Notes	The domain controller was not rebooted after promotion, either due to an error or because you specified the ADDSDeployment Windows PowerShell argument **-norebootoncompletion**. Restart the domain controller.

Issue	The Next button is not available on the Domain Controller Options page
Symptoms	Even though you have set a password, the **Next** button on the **Domain Controller Options** page in Server Manager is not available. There is no site listed in the **Site name** menu.
Resolution and Notes	You have multiple AD DS sites and at least one is missing subnets; this future domain controller belongs to one of those subnets. You must manually select the subnet from the Site name dropdown menu. You should also review all AD sites using DSSITE.MSC or use the following Windows PowerShell command to find all sites missing subnets: `get-adreplicationsite -filter * -property subnets | where-object {!$_.subnets -eq "*"} | format-table name`

Issue	Promotion or demotion fails with message "the service cannot be started"
Symptoms	If you attempt promotion, demotion, or cloning of a domain controller you receive error: `The service cannot be started, either because it is disabled or it has no enabled devices associated with it" (0x80070422)` The error may be interactive, an event, or

	written to a log like dcpromoui.log or dcpromo.log
Resolution and Notes	The DS Role Server service (DsRoleSvc) is disabled. By default, this service is installed during AD DS role installation and set to a Manual start type. Do not disable this service. Set it back to Manual and allow the DS role operations to start and stop it on demand. This behavior is by design.

Issue	Server Manager still warns that you need to promote DC
Symptoms	If you promote a domain controller using the deprecated dcpromo.exe /unattend or upgrade an existing Windows Server 2008 R2 domain controller in place to Windows Server 2012, Server Manager still shows the post-deployment configuration task **Promote this server to a domain controller**.
Resolution and Notes	Click the post-deployment warning link and the message will disappear for good. This behavior is cosmetic and expected.

Issue	Server Manager deployment script missing role installation
Symptoms	If you promote a domain controller using Server Manager and save the Windows PowerShell deployment script, it does not include the role installation cmdlet and arguments (install-windowsfeature -name ad-domain-services -includemanagementtools). Without the role, the DC cannot be configured.
Resolution and Notes	Manually add that cmdlet and arguments to any scripts. This behavior is expected and by design.

Issue	Server Manager deployment script is not

	named PS1
Symptoms	If you promote a domain controller using Server Manager and save the Windows PowerShell deployment script, the file is named with a random temporary name and not as a PS1 file.
Resolution and Notes	Manually rename the file. This behavior is expected and by design.

Issue	Dcpromo /unattend allows unsupported functional levels

Symptoms	If you promote a domain controller using dcpromo /unattend with the following sample answer file:
	`[DCInstall]`
	`NewDomain=Forest`
	`ReplicaOrNewDomain=Domain`
	`NewDomainDNSName=corp.contoso.com`
	`SafeModeAdminPassword=Safepassword@6`
	`DomainNetbiosName=corp`
	`DNSOnNetwork=Yes`
	`AutoConfigDNS=Yes`
	`RebootOnSuccess=NoAndNoPromptEither`
	`RebootOnCompletion=No`
	`DomainLevel=0`
	`ForestLevel=0`
	
	Promotion fails with the following errors in the dcpromoui.log:
	`dcpromoui EA4.5B8 0089 13:31:50.783 Enter CArgumentsSpec::ValidateArgument DomainLevel`
	`dcpromoui EA4.5B8 008A 13:31:50.783 Value for DomainLevel is 0`
	`dcpromoui EA4.5B8 008B 13:31:50.783 Exit code is 77`

	dcpromoui EA4.5B8 008C 13:31:50.783 The specified argument is invalid. dcpromoui EA4.5B8 008D 13:31:50.783 closing log dcpromoui EA4.5B8 0032 13:31:50.830 Exit code is 77 Level 0 is Windows 2000, which is not supported in Windows Server 2012.
Resolution and Notes	Do not use the deprecated dcpromo /unattend and understand that it allows you to specify invalid settings that later fail. This behavior is expected and by design.

Issue	Promotion “hangs” at creating NTDS settings object, never completes
Symptoms	If you promote a replica DC or RODC, the promotion reaches “creating NTDS settings object” and never proceeds or completes. The logs stop updating as well.
Resolution and Notes	This is a known issue caused by providing credentials of the built-in local Administrator account with a matching password to the built-in domain Administrator account. This causes a failure down in the core setup engine that does not error, but instead waits indefinitely (quasi-loop). This is expected – albeit undesirable – behavior. To fix the server: 1. Reboot it. 1. In AD, delete that server’s member computer account (it will not yet be a DC account) 1. On that server, forcibly disjoin it from the domain 1. On that server, remove the AD DS role. 1. Reboot 1. Re-add the AD DS role and reattempt

	promotion, ensuring that you always provide the ***domain\admin*** formatted credentials to DC promotion and not just the built-in local administrator account

Active Directory Administrative Center Enhancements

The Active Directory Administrative Center (ADAC) in Windows Server® 2012 includes enhanced management experience features. These new features in Windows Server 2012 ease the administrative burden for managing Active Directory Domain Services (AD DS). The following topics provide an introduction and additional details:

- Introduction to Active Directory Administrative Center Enhancements (Level 100)
- Advanced AD DS Management Using Active Directory Administrative Center (Level 200)

Introduction to Active Directory Administrative Center Enhancements (Level 100)

ADAC in Windows Server 2012 includes management features for the following:

- Active Directory Recycle Bin
- Fine-Grained Password Policy
- Windows PowerShell History Viewer

Active Directory Recycle Bin

Accidental deletion of Active Directory objects is a common occurrence for users of Active Directory Domain Services (AD DS) and Active Directory Lightweight Directory Services (AD LDS). In past versions of Windows Server, prior to Windows Server 2008 R2, one could recover accidentally deleted objects in Active Directory, but the solutions had their drawbacks.

In Windows Server 2008, you could use the Windows Server Backup feature and **ntdsutil** authoritative restore command to mark objects as authoritative to ensure that the restored data was replicated throughout the domain. The drawback to the authoritative restore solution was that it had to be performed in Directory Services Restore Mode (DSRM). During DSRM, the domain controller being restored had to remain offline. Therefore, it was not able to service client requests.

In Windows Server 2003 Active Directory and Windows Server 2008 AD DS, you could recover deleted Active Directory objects through tombstone reanimation. However, reanimated objects' link-valued attributes (for example, group memberships of user accounts) that were physically removed and non-link-valued attributes that were cleared were not recovered. Therefore, administrators could not rely on tombstone reanimation as the ultimate solution to accidental deletion of objects. For more information about tombstone reanimation, see Reanimating Active Directory Tombstone Objects.

Active Directory Recycle Bin, starting in Windows Server 2008 R2, builds on the existing tombstone reanimation infrastructure and enhances your ability to preserve and recover accidentally deleted Active Directory objects.

When you enable Active Directory Recycle Bin, all link-valued and non-link-valued attributes of the deleted Active Directory objects are preserved and the objects are restored in their entirety to the same consistent logical state that they were in immediately before deletion. For example, restored user accounts automatically regain all group memberships and corresponding access rights that they had immediately before deletion, within and across domains. Active Directory Recycle Bin works for both AD DS and AD LDS environments. For a detailed description of Active Directory Recycle Bin, see What's New in AD DS: Active Directory Recycle Bin.

What's new? In Windows Server 2012, the Active Directory Recycle Bin feature has been enhanced with a new graphical user interface for users to manage and restore deleted objects. Users can now visually locate a list of deleted objects and restore them to their original or desired locations.

If you plan to enable Active Directory Recycle Bin in Windows Server 2012, consider the following:

- By default, Active Directory Recycle Bin is disabled. To enable it, you must first raise the forest functional level of your AD DS or AD LDS environment to Windows Server 2008 R2 or higher. This in turn requires that all domain controllers in the forest or all servers that host instances of AD LDS configuration sets be running Windows Server 2008 R2 or higher.
- The process of enabling Active Directory Recycle Bin is irreversible. After you enable Active Directory Recycle Bin in your environment, you cannot disable it.
- To manage the Recycle Bin feature through a user interface, you must install the version of Active Directory Administrative Center in Windows Server 2012.

Notes

You can use **Server Manager** to install Remote Server Administration Tools (RSAT) on Windows Server 2012 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

You can use RSAT on Windows® 8 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

Active Directory Recycle Bin step-by-step

In the following steps, you will use ADAC to perform the following Active Directory Recycle Bin tasks in Windows Server 2012:

- Step 1: Raise the forest functional level
- Step 2: Enable Recycle Bin
- Step 3: Create test users, group and organizational unit
- Step 4: Restore deleted objects

Note
Membership in the Enterprise Admins group or equivalent permissions is required to perform the following steps.

Step 1: Raise the forest functional level

In this step, you will raise the forest functional level. You must first raise the functional level on the target forest to be Windows Server 2008 R2 at a minimum before you enable Active Directory Recycle Bin.

To raise the functional level on the target forest

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Click the target domain in the left navigation pane and in the **Tasks** pane, click **Raise the forest functional level**. Select a forest functional level that is at least Windows Server 2008 R2 or higher and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADForestMode -Identity contoso.com -ForestMode Windows2008R2Forest -Confirm:$false
```

For the **–Identity** argument, specify the fully qualified DNS name.

Step 2: Enable Recycle Bin

In this step, you will enable the Recycle Bin to restore deleted objects in AD DS.

To enable Active Directory Recycle Bin in ADAC on the target domain

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the **Tasks** pane, click **Enable Recycle Bin ...** in the **Tasks** pane, click **OK** on the warning message box, and then click **OK** to the refresh ADAC message.

4. Press F5 to refresh ADAC.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Enable-ADOptionalFeature –Identity 'CN=Recycle Bin Feature,CN=Optional
Features,CN=Directory Service,CN=Windows
NT,CN=Services,CN=Configuration,DC=contoso,DC=com' –Scope ForestOrConfigurationSet –
Target 'contoso.com'
```

Step 3: Create test users, group and organizational unit

In the following procedures, you will create two test users. You will then create a test group and add the test users to the group. In addition, you will create an OU.

To create test users

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the **Tasks** pane, click **New** and then click **User**.

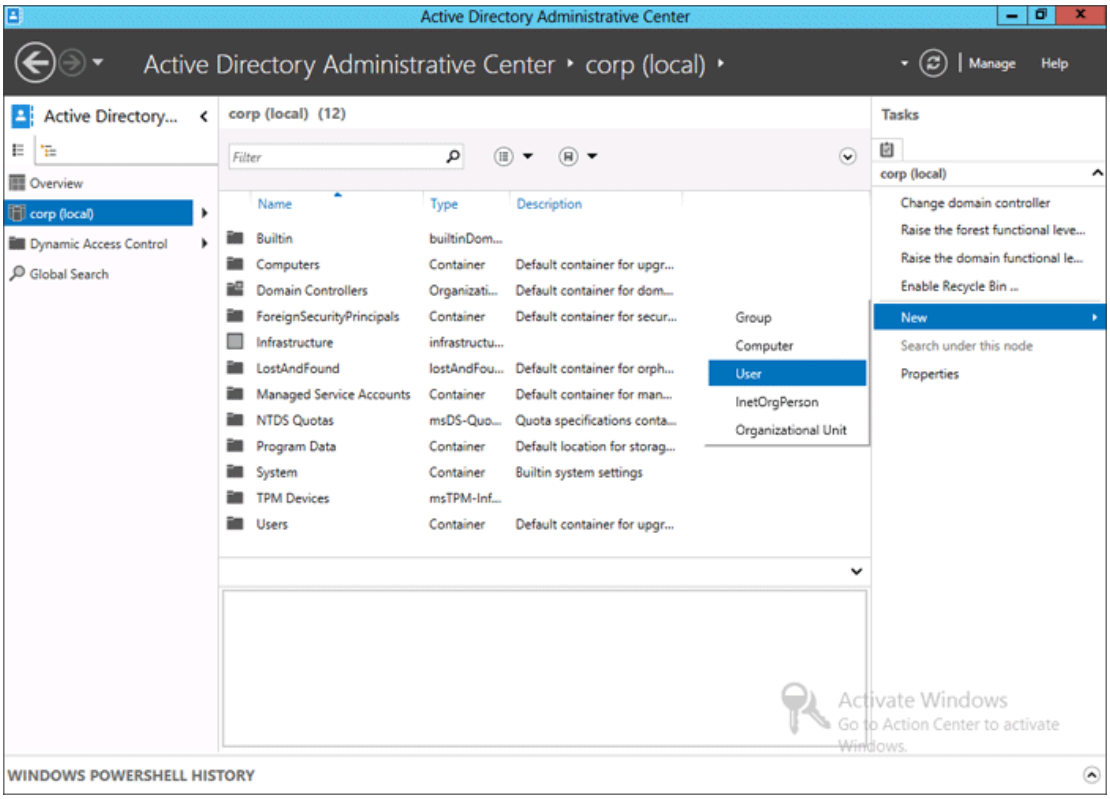

4. Enter the following information under **Account** and then click OK:
   - Full name: test1
   - User SamAccountName logon: test1
   - Password: p@ssword1
   - Confirm password: p@ssword1
5. Repeat the previous steps to create a second user, test2.

To create a test group and add users to the group

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the **Tasks** pane, click **New** and then click **Group**.
4. Enter the following information under **Group** and then click **OK**:
   - **Group name: group1**
5. Click **group1**, and then under the **Tasks** pane, click **Properties**.
6. Click **Members**, click **Add**, type **test1;test2**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-ADGroupMember -Identity group1 -Member test1
```

To create an organizational unit

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the **Tasks** pane, click **New** and then click **Organizational Unit**.
4. Enter the following information under **Organizational Unit** and then click **OK**:
   - **Name OU1**

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
1..2 | ForEach-Object {New-ADUser -SamAccountName test$_ -Name "test$_" -Path
"DC=fabrikam,DC=com" -AccountPassword (ConvertTo-SecureString -AsPlainText "p@ssword1" -
Force) -Enabled $true}

New-ADGroup -Name "group1" -SamAccountName group1 -GroupCategory Security -GroupScope
Global -DisplayName "group1"

New-ADOrganizationalUnit -Name OU1 -Path "DC=fabrikam,DC=com"
```

Step 4: Restore deleted objects

In the following procedures, you will restore deleted objects from the **Deleted Objects** container to their original location and to a different location.

To restore deleted objects to their original location

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Select users **test1** and **test2**, click **Delete** in the **Tasks** pane and then click **Yes** to confirm the deletion.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Get-ADUser –Filter 'Name –Like "*test*"'|Remove-ADUser -
Confirm:$false
```

4. Navigate to the **Deleted Objects** container, select **test2** and **test1** and then click **Restore** in the **Tasks** pane.
5. To confirm the objects were restored to their original location, navigate to the target domain and verify the user accounts are listed.

Note

If you navigate to the **Properties** of the user accounts **test1** and **test2** and then click **Member Of**, you will see that their group membership was also restored.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Windows PowerShell equivalent commands

```
Get-ADObject –Filter 'Name –Like "*test*"' –IncludeDeletedObjects | Restore-ADObject
```

To restore deleted objects to a different location

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Select users **test1** and **test2**, click **Delete** in the **Tasks** pane and then click **Yes** to confirm the deletion.
4. Navigate to the **Deleted Objects** container, select **test2** and **test1** and then click **Restore To** in the **Tasks** pane.
5. Select **OU1** and then click **OK**.
6. To confirm the objects were restored to **OU1**, navigate to the target domain, double click **OU1** and verify the user accounts are listed.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Get-ADObject –Filter 'Name –Like "*test*"' –IncludeDeletedObjects | Restore-ADObject –
TargetPath "OU=OU1,DC=contoso,DC=com"
```

Fine-Grained Password Policy

The Windows Server 2008 operating system provides organizations with a way to define different password and account lockout policies for different sets of users in a domain. In Active Directory domains prior to Windows Server 2008, only one password policy and account lockout policy could be applied to all users in the domain. These policies were specified in the Default Domain Policy for the domain. As a result, organizations that wanted different password and account lockout settings for different sets of users had to either create a password filter or deploy multiple domains. Both are costly options.

You can use fine-grained password policies to specify multiple password policies within a single domain and apply different restrictions for password and account lockout policies to different sets of users in a domain. For example, you can apply stricter settings to privileged accounts and less strict settings to the accounts of other users. In other cases, you might want to apply a special password policy for accounts whose passwords are synchronized with other data sources. For a detailed description of Fine-Grained Password Policy, see AD DS: Fine-Grained Password Policies

What's new? In Windows Server 2012, fine-grained password policy management is made easier and more visual by providing a user interface for AD DS administrators to manage them in ADAC. Administrators can now view a given user's resultant policy, view and sort all password policies within a given domain, and manage individual password policies visually.

If you plan to use fine-grained password policies in Windows Server 2012, consider the following:

- Fine-grained password policies apply only global security groups and user objects (or inetOrgPerson objects if they are used instead of user objects). By default, only members of the Domain Admins group can set fine-grained password policies. However, you can also delegate the ability to set these policies to other users. The domain functional level must be Windows Server 2008 or higher.
- You must use the Windows Server 2012 version of Active Directory Administrative Center to administer fine-grained password policies through a graphical user interface.

Notes

You can use **Server Manager** to install Remote Server Administration Tools (RSAT) on Windows Server 2012 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

You can use RSAT on Windows® 8 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

Fine-Grained Password Policy step-by-step

In the following steps, you will use ADAC to perform the following fine-grained password policy tasks:

- Step 1: Raise the domain functional level
- Step 2: Create test users, group, and organizational unit
- Step 3: Create a new fine-grained password policy

- Step 4: View a resultant set of policies for a user
- Step 5: Edit a fine-grained password policy
- Step 6: Delete a fine-grained password policy

Note
Membership in the Domain Admins group or equivalent permissions is required to perform the following steps.

Step 1: Raise the domain functional level

In the following procedure, you will raise the domain functional level of the target domain to Windows Server 2008 or higher. A domain functional level of Windows Server 2008 or higher is required to enable fine-grained password policies.

To raise the domain functional level

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Click the target domain in the left navigation pane and in the **Tasks** pane, click **Raise the domain functional level**. Select a forest functional level that is at least Windows Server 2008 or higher and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADDomainMode -Identity contoso.com -DomainMode 3
```

Step 2: Create test users, group, and organizational unit

To create the test users and group need for this step, follow the procedures located here: Step 3: Create test users, group and organizational unit (you do not need to create the OU to demonstrate fine-grained password policy).

Step 3: Create a new fine-grained password policy

In the following procedure you will create a new fine-grained password policy using the UI in ADAC.

To create a new fine grained password policy

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in

the **Add Navigation Nodes** dialog box and then click **OK**.

3. In the ADAC navigation pane, open the **System** container and then click **Password Settings Container**.
4. In the **Tasks** pane, click **New**, and then click **Password Settings**.

   Fill in or edit fields inside the property page to create a new **Password Settings** object. The **Name** and **Precedence** fields are required.

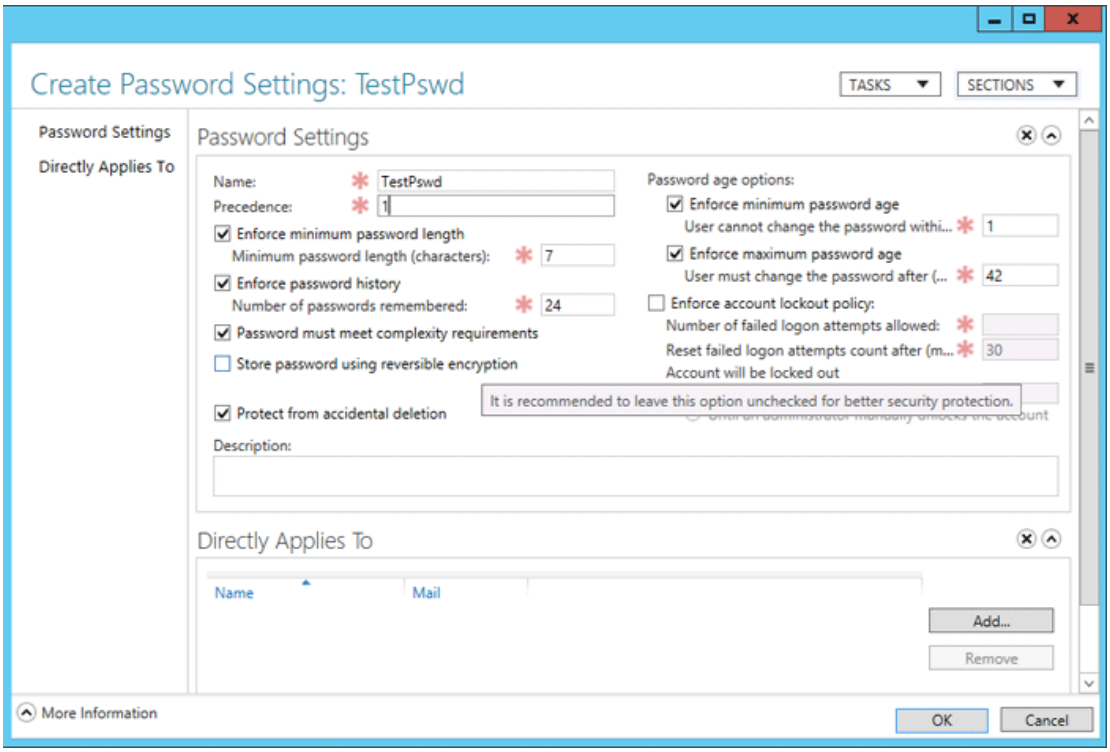

5. Under **Directly Applies To**, click **Add**, type **group1**, and then click **OK**.

   This associates the Password Policy object with the members of the global group you created for the test environment.
6. Click **OK** to submit the creation.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADFineGrainedPasswordPolicy TestPswd -ComplexityEnabled:$true -
LockoutDuration:"00:30:00" -LockoutObservationWindow:"00:30:00" -LockoutThreshold:"0" -
MaxPasswordAge:"42.00:00:00" -MinPasswordAge:"1.00:00:00" -MinPasswordLength:"7" -
PasswordHistoryCount:"24" -Precedence:"1" -ReversibleEncryptionEnabled:$false -
ProtectedFromAccidentalDeletion:$true

Add-ADFineGrainedPasswordPolicySubject TestPswd -Subjects group1
```

Step 4: View a resultant set of policies for a user

In the following procedure, you will view the resultant password settings for a user that is a member of the group to which you assigned a fine grained password policy in Step 3: Create a new fine-grained password policy.

To view a resultant set of policies for a user

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Select a user, **test1** that belongs to the group, **group1** that you associated a fine-grained password policy with in Step 3: Create a new fine-grained password policy.
4. Click **View Resultant Password Settings** in the **Tasks** pane.
5. Examine the password setting policy and then click **Cancel**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Get-ADUserResultantPasswordPolicy test1
```

Step 5: Edit a fine-grained password policy

In the following procedure, you will edit the fine grained password policy you created in Step 3: Create a new fine-grained password policy

To edit a fine-grained password policy

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the ADAC **Navigation Pane**, expand **System** and then click **Password Settings Container**.
4. Select the fine grained password policy you created in Step 3: Create a new fine-grained password policy and click **Properties** in the **Tasks** pane.
5. Under **Enforce password history**, change the value of **Number of passwords remembered** to **30**.
6. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADFineGrainedPasswordPolicy TestPswd -PasswordHistoryCount:"30"
```

Step 6: Delete a fine-grained password policy

To delete a fine-grained password policy

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. In the ADAC Navigation Pane, expand **System** and then click **Password Settings Container**.
4. Select the fine grained password policy you created in Step 3: Create a new fine-grained password policy and in the **Tasks** pane click **Properties**.
5. Clear the **Protect from accidental deletion** checkbox and click **OK**.
6. Select the fine grained password policy, and in the **Tasks** pane click **Delete**.
7. Click **OK** in the confirmation dialog.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADFineGrainedPasswordPolicy -Identity TestPswd -ProtectedFromAccidentalDeletion
$False

Remove-ADFineGrainedPasswordPolicy TestPswd -Confirm
```

Windows PowerShell History Viewer

ADAC is a user interface tool built on top of Windows PowerShell. In Windows Server 2012, IT administrators can leverage ADAC to learn Windows PowerShell for Active Directory cmdlets by using the Windows PowerShell History Viewer. As actions are executed in the user interface, the equivalent Windows PowerShell command is shown to the user in Windows PowerShell History Viewer. This allows administrators to create automated scripts and reduce repetitive tasks, thus increasing IT productivity. Also, this feature reduces the time to learn Windows PowerShell for Active Directory and increases the users' confidence in the correctness of their automation scripts.

When using the Windows PowerShell History Viewer in Windows Server 2012 consider the following:

- To use Windows PowerShell Script Viewer, you must use the Windows Server 2012 version of ADAC

  Notes

  You can use **Server Manager** to install Remote Server Administration Tools (RSAT) on Windows Server 2012 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

  You can use RSAT on Windows® 8 computers to use the correct version of Active Directory Administrative Center to manage Recycle Bin through a user interface.

- Have some basic Windows PowerShell knowledge. For example, you need to know how piping in Windows PowerShell works. For more information about piping in Windows PowerShell, see Piping and the Pipeline in Windows PowerShell.

Windows PowerShell History Viewer step-by-step

In the following procedure, you will use the Windows PowerShell History Viewer in ADAC to construct a Windows PowerShell script. Before you begin this procedure, remove user, **test1** from the group, **group1**.

To construct a script using PowerShell History Viewer

1. Right click the Windows PowerShell icon, click **Run as Administrator** and type **dsac.exe** to open ADAC.
2. Click **Manage**, click **Add Navigation Nodes** and select the appropriate target domain in the **Add Navigation Nodes** dialog box and then click **OK**.
3. Expand the **Windows PowerShell History** pane at the bottom of the ADAC screen.
4. Select user, **test1**.
5. Click **Add to group…** in the **Tasks** pane.
6. Navigate to **group1** and click **OK** in the dialog box.
7. Navigate to the **Windows PowerShell History** pane and locate the command just generated.
8. Copy the command and paste it into your desired editor to construct your script.

   For example, you can modify the command to add a different user to **group1**, or add **test1** to a different group.

See Also

Advanced AD DS Management Using Active Directory Administrative Center (Level 200)

Advanced AD DS Management Using Active Directory Administrative Center (Level 200)

This topic covers the updated Active Directory Administrative Center with its new Active Directory Recycle Bin, Fine-grained Password policy, and Windows PowerShell History Viewer in more detail, including architecture, examples for common tasks, and troubleshooting information. For an introduction, see Introduction to Active Directory Administrative Center Enhancements (Level 100).

- Active Directory Administrative Center Architecture
- Enabling and Managing the Active Directory Recycle Bin Using Active Directory Administrative Center
- Configuring and Managing Fine-Grained Password Policies Using Active Directory Administrative Center
- Using the Active Directory Administrative Center Windows PowerShell History Viewer
- Troubleshooting AD DS Management

Active Directory Administrative Center Architecture

ADPrep Executables, DLLs

The module and underlying architecture of Active Directory Administrative Center has not changed with the new recycle bin, FGPP, and history viewer capabilities.

- Microsoft.ActiveDirectory.Management.UI.dll
- Microsoft.ActiveDirectory.Management.UI.resources.dll
- Microsoft.ActiveDirectory.Management.dll
- Microsoft.ActiveDirectory.Management.resources.dll
- ActiveDirectoryPowerShellResources.dll

The underlying Windows PowerShell and layer of operations for the new Recycle Bin functionality are illustrated below:

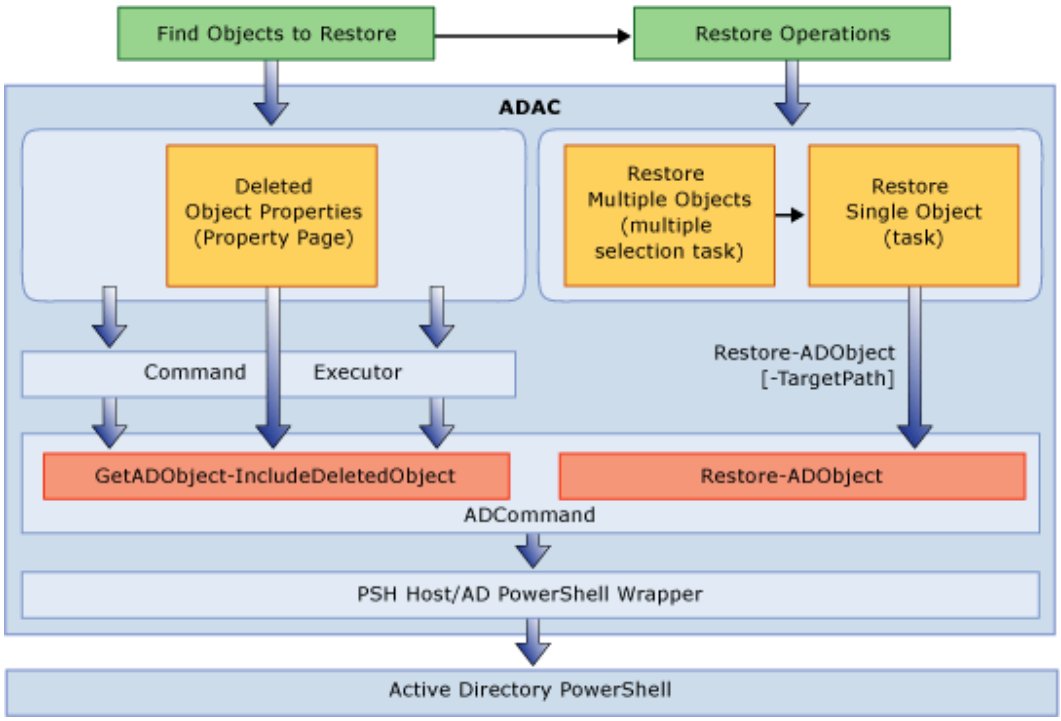

Enabling and Managing the Active Directory Recycle Bin Using Active Directory Administrative Center

Capabilities

- The Windows Server 2012 Active Directory Administrative Center enables you to configure and manage the Active Directory Recycle Bin for any domain partition in a forest. There is no longer a requirement to use Windows PowerShell or Ldp.exe to enable the Active Directory Recycle Bin or restore objects in domain partitions.
- The Active Directory Administrative Center has advanced filtering criteria, making targeted restoration easier in large environments with many intentionally deleted objects.

Limitations

- Because the Active Directory Administrative Center can only manage domain partitions, it cannot restore deleted objects from the Configuration, Domain DNS, or Forest DNS partitions (you cannot delete objects from the Schema partition). To restore objects from non-domain partitions, use Restore-ADObject.
- The Active Directory Administrative Center cannot restore sub-trees of objects in a single action. For example, if you delete an OU with nested OUs, users, groups, and computers, restoring the base OU does not restore the child objects.

Note

The Active Directory Administrative Center batch restore operation does a “best effort” sort of the deleted objects *within the selection only* so parents are ordered before the children for the restore list. In simple test cases, sub-trees of objects may be restored in a single action. But corner cases, such as a selection that contains partial trees – trees with some of the deleted parent nodes missing - or error cases, such as skipping the child objects when parent restore fails, may not work as expected. For this reason, you should always restore sub-trees of objects as a separate action after you restore the parent objects.

Active Directory Recycle Bin requires a Windows Server 2008 R2 Forest Functional Level and you must be a member of the Enterprise Admins group. Once enabled, you cannot disable Active Directory Recycle Bin. Active Directory Recycle Bin increases the size of the Active Directory database (NTDS.DIT) on every domain controller in the forest. Disk space used by the recycle bin continues to increase over time as it preserves objects and all their attribute data.

For more information, see Active Directory Recycle Bin Step-by-Step Guide and Assess Hardware Requirements (for Active Directory Recycle Bin).

You can *lower* the forest and domain functional levels from Windows Server 2012 to Windows Server 2008 R2, even after enabling the Active Directory Recycle Bin. This requires using the **Set-ADForestMode** and **Set-ADDomainMode** Active Directory cmdlets.

For example:

```
Set-AdForestMode -identity corp.contoso.com -server dc1.corp.contoso.com -forestmode
Windows2008R2Forest

Set-AdDomainMode -identity research.corp.contoso.com -server
dc3.research.corp.contoso.com -domainmode Windows2008R2Domain
```

You cannot use the Active Directory Administrative Center to make this change - it only raises functional levels.

Enabling Active Directory Recycle Bin using Active Directory Administrative Center

To enable the Active Directory Recycle Bin, open the **Active Directory Administrative Center** and click the name of your forest in the navigation pane. From the **Tasks** pane, click **Enable Recycle Bin**.

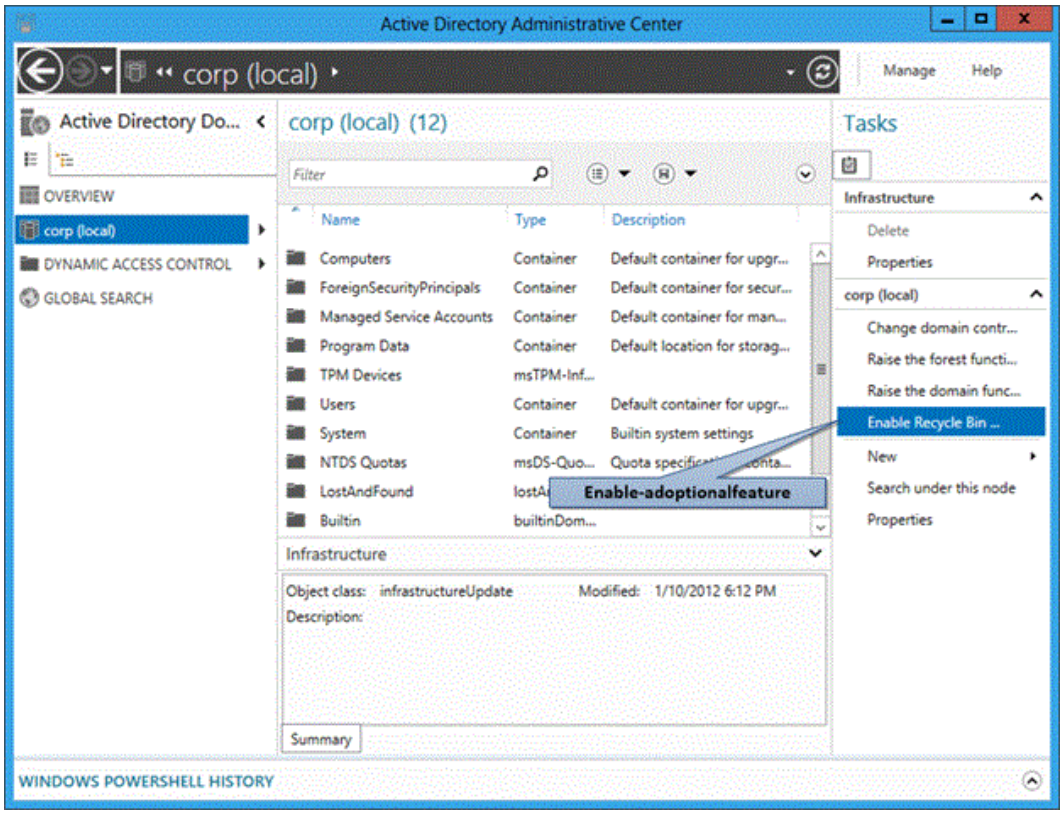

The Active Directory Administrative Center shows the **Enable Recycle Bin Confirmation** dialog. This dialog warns you that enabling the recycle bin is irreversible. Click **OK** to enable the Active Directory Recycle Bin. The Active Directory Administrative Center shows another dialog to remind you that the Active Directory Recycle Bin is not fully functional until all domain controllers replicate the configuration change.

Important

The option to enable the Active Directory Recycle Bin is unavailable if:

- The forest functional level is less than Windows Server 2008 R2
- It is already enabled

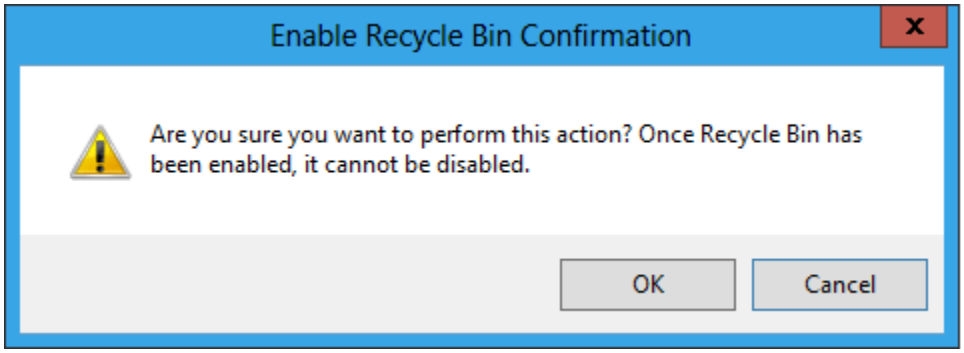

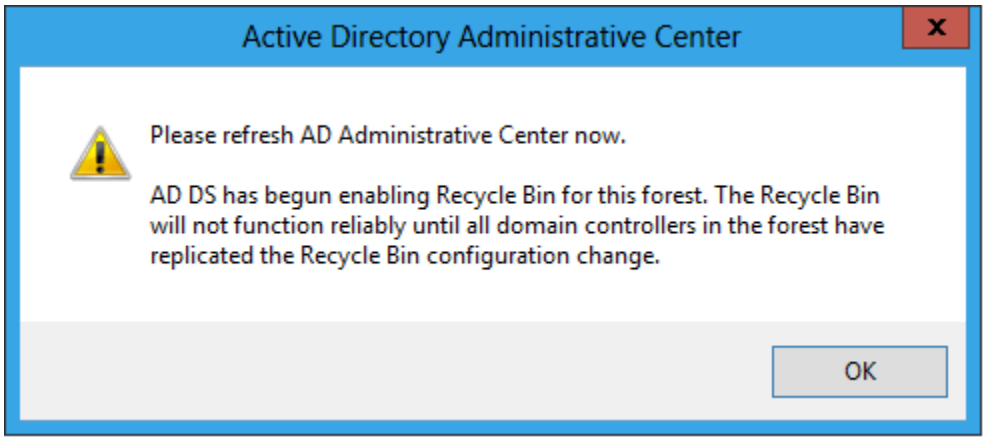

The equivalent Active Directory Windows PowerShell cmdlet is still:

```
Enable-ADOptionalFeature
```

For more information about using Windows PowerShell to enable the Active Directory Recycle Bin, see the Active Directory Recycle Bin Step-by-Step Guide.

Managing Active Directory Recycle Bin using Active Directory Administrative Center

This section uses the example of an existing domain named **corp.contoso.com**. This domain organizes users into a parent OU named **UserAccounts**. The **UserAccounts** OU contains three child OUs named by department, which each contain further OUs, users, and groups.

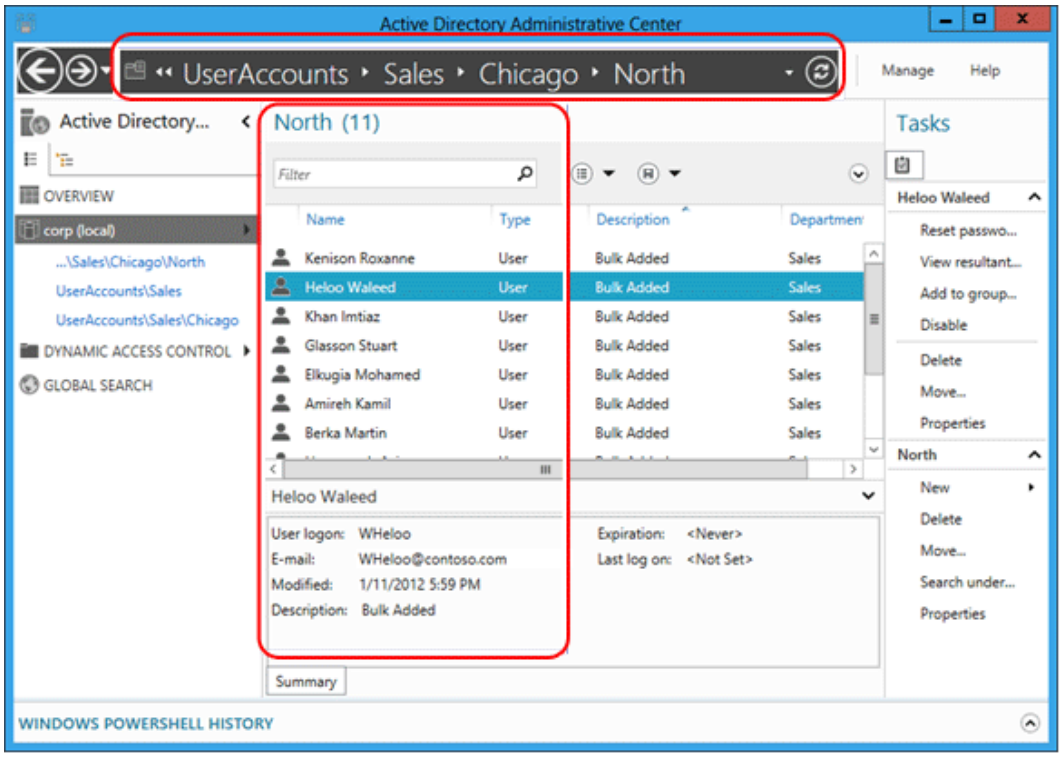

Storage and Filtering

The Active Directory Recycle Bin preserves all objects deleted in the forest. It saves these objects according to the **msDS-deletedObjectLifetime** attribute, which by default is set to match the **tombstoneLifetime** attribute of the forest. In any forest created using Windows Server 2003 SP1 or later, the value of **tombstoneLifetime** is set to 180 days by default. In any forest upgraded from Windows 2000 or installed with Windows Server 2003 (no service pack), the default tombstoneLifetime attribute is NOT SET and Windows therefore uses the internal default of 60 days. All of this is configurable.You can use the Active Directory Administrative Center to restore any objects deleted from the domain partitions of the forest. You must continue to use the cmdlet **Restore-ADObject** to restore deleted objects from other partitions, such as Configuration.Enabling the Active Directory Recycle Bin makes the **Deleted Objects** container visible under every domain partition in the Active Directory Administrative Center.

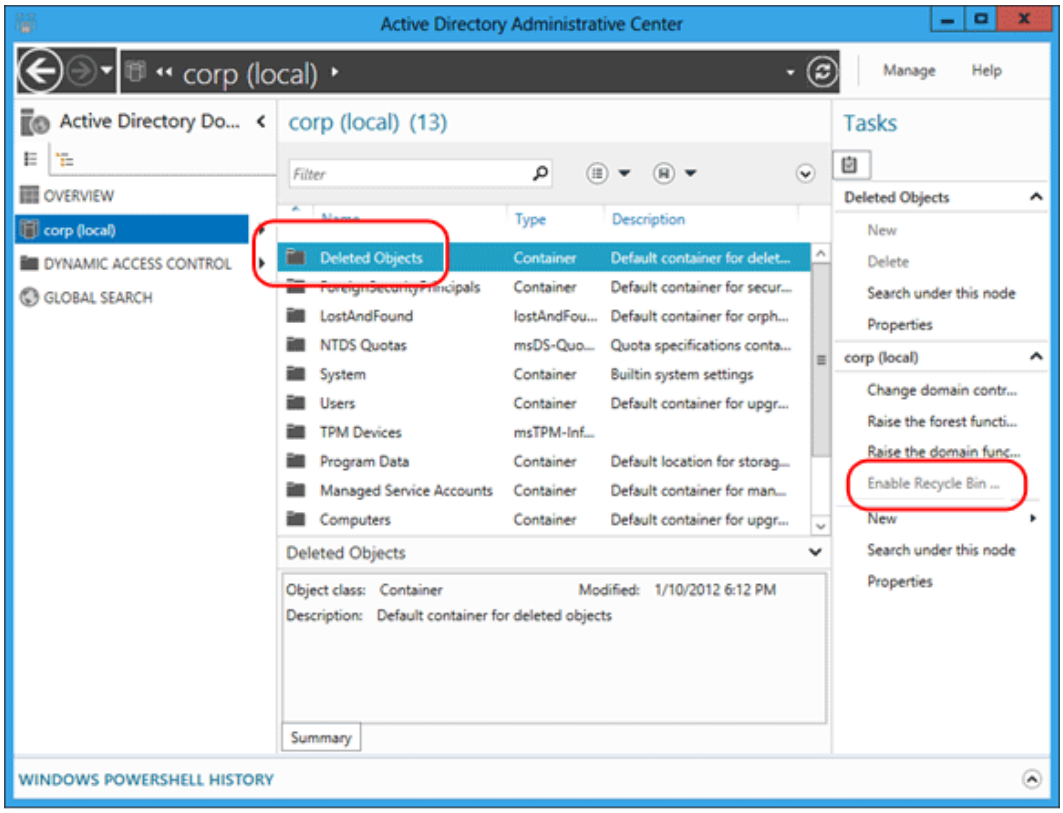

The **Deleted Objects** container shows you all the restorable objects in that domain partition. Deleted objects older than **msDS-deletedObjectLifetime** are known as recycled objects. The Active Directory Administrative Center does not show recycled objects and you cannot restore these objects using Active Directory Administrative Center.

For a deeper explanation of the recycle bin's architecture and processing rules, see [The AD Recycle Bin: Understanding, Implementing, Best Practices, and Troubleshooting]().

The Active Directory Administrative Center artificially limits the default number of objects returned from a container to 20,000 objects. You can raise this limit as high as 100,000 objects by clicking the **Manage** menu, then **Management List Options**.

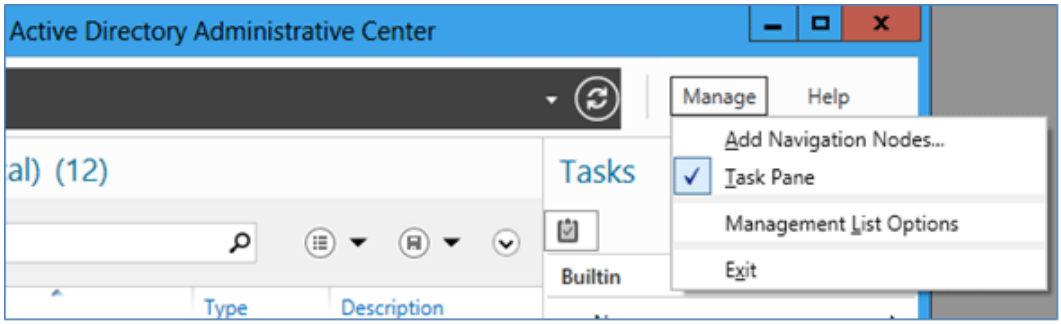

Restoration

Filtering

Active Directory Administrative Center offers powerful criteria and filtering options that you should become familiar with before you need to use them in a real-life restoration. Domains intentionally delete many objects over their lifetime .With a likely deleted object lifetime of 180 days, you cannot simply restore all objects when an accident occurs.

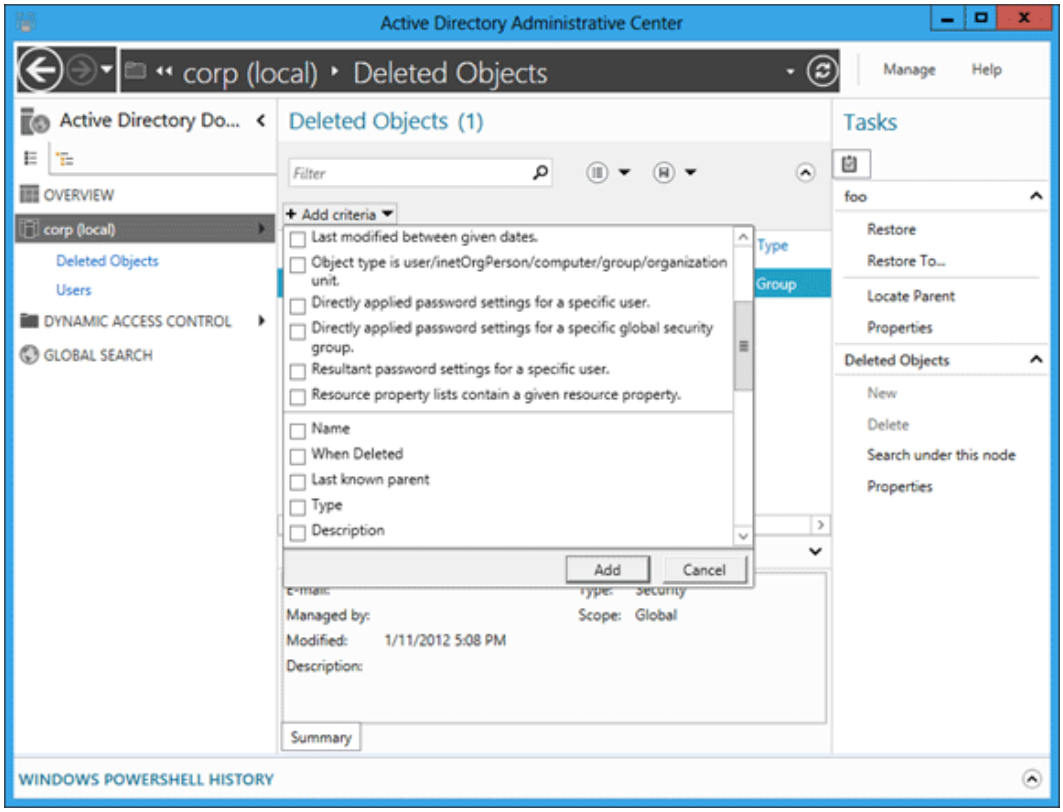

Rather than writing complex LDAP filters and converting UTC values into dates and times, use the basic and advanced **Filter** menu to list only the relevant objects. If you know the day of deletion, the names of objects, or any other key data, use that to your advantage when filtering. Toggle the advanced filter options by clicking the chevron to the right of the search box.

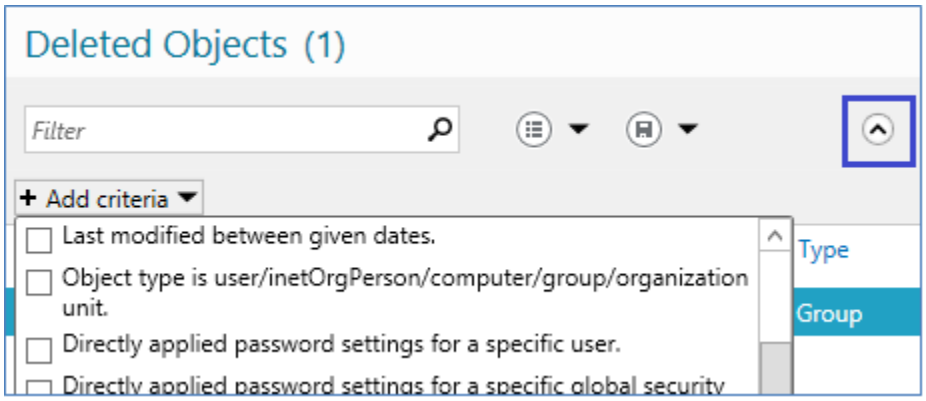

The restore operation supports all the standard filter criteria options, the same as any other search. Of the built-in filters, the important ones for restoring objects are typically:

- *ANR (ambiguous name resolution - not listed in the menu, but what is used when you type in the* **Filter** *box)*
- Last modified between given dates
- Object is user/inetorgperson/computer/group/organization unit
- Name
- When deleted
- Last known parent
- Type
- Description
- City
- Country /region
- Department
- Employee ID
- First name
- Job title
- Last name
- SAMaccountname
- State/Province
- Telephone number
- UPN
- ZIP/Postal code

You can add multiple criteria. For example, you can find all user objects deleted on September 24th 2012 from Chicago, Illinois with a job title of Manager.

You can also add, modify, or reorder the column headers to provide more detail when evaluating which objects to recover.

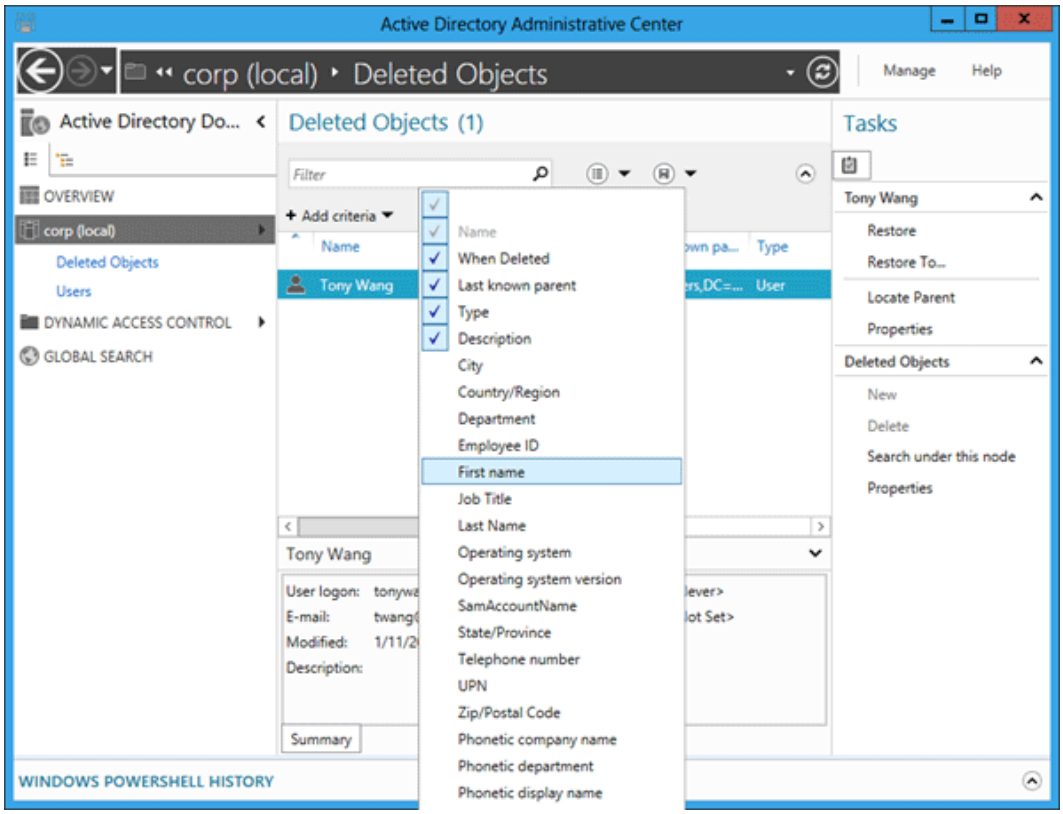

For more information about Ambiguous Name Resolution, see ANR Attributes.

Single Object

Restoring deleted objects has always been a single operation. The Active Directory Administrative Center makes that operation easier. To restore a deleted object, such as a single user:

1. Click the domain name in the navigation pane of the Active Directory Administrative Center.
2. Double-click **Deleted Objects** in the management list.
3. Right-click the object and then click **Restore**, or click **Restore** from the **Tasks** pane.

The object restores to its original location.

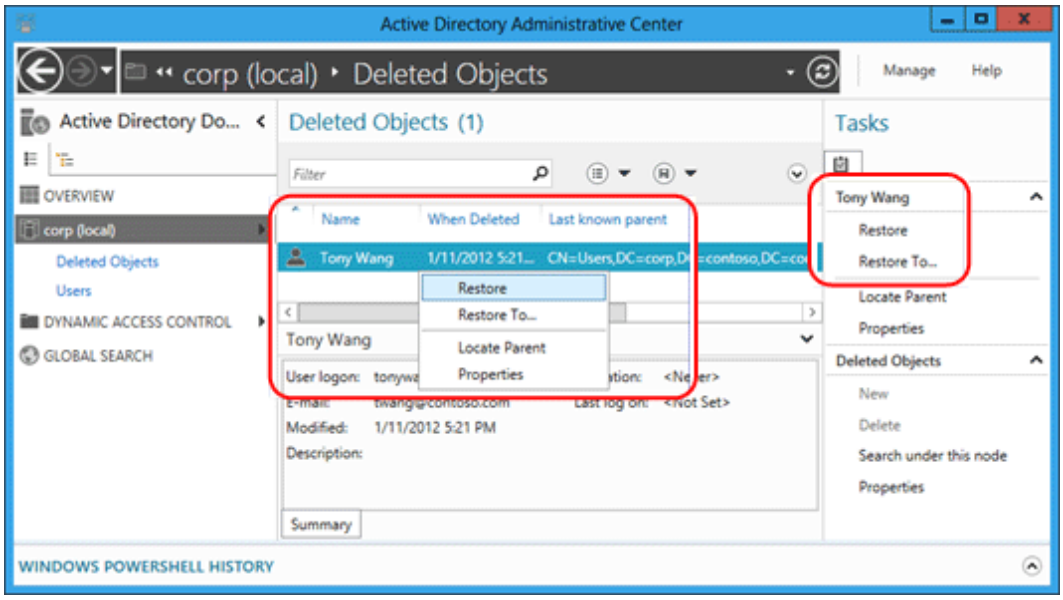

Click **Restore To…** to change the restore location. This is useful if the deleted object's parent container was also deleted but you do not want to restore the parent.

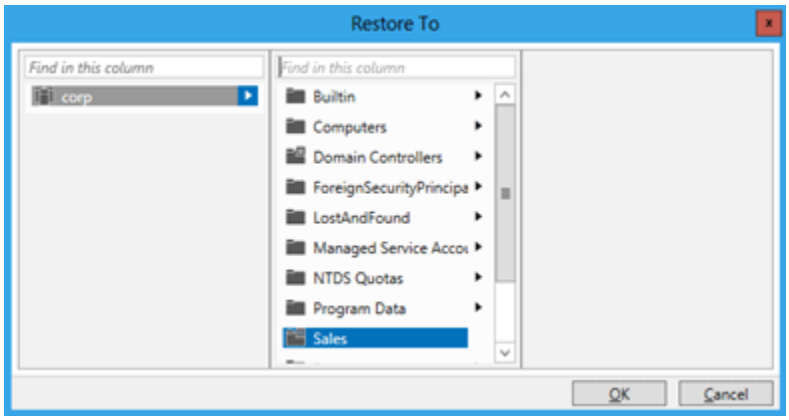

Multiple Peer Objects

You can restore multiple peer-level objects, such as all the users in an OU. Hold down the CTRL key and click one or more deleted objects you want to restore. Click **Restore** from the Tasks pane. You can also select all displayed objects by holding down the CTRL and A keys, or a range of objects using SHIFT and clicking.

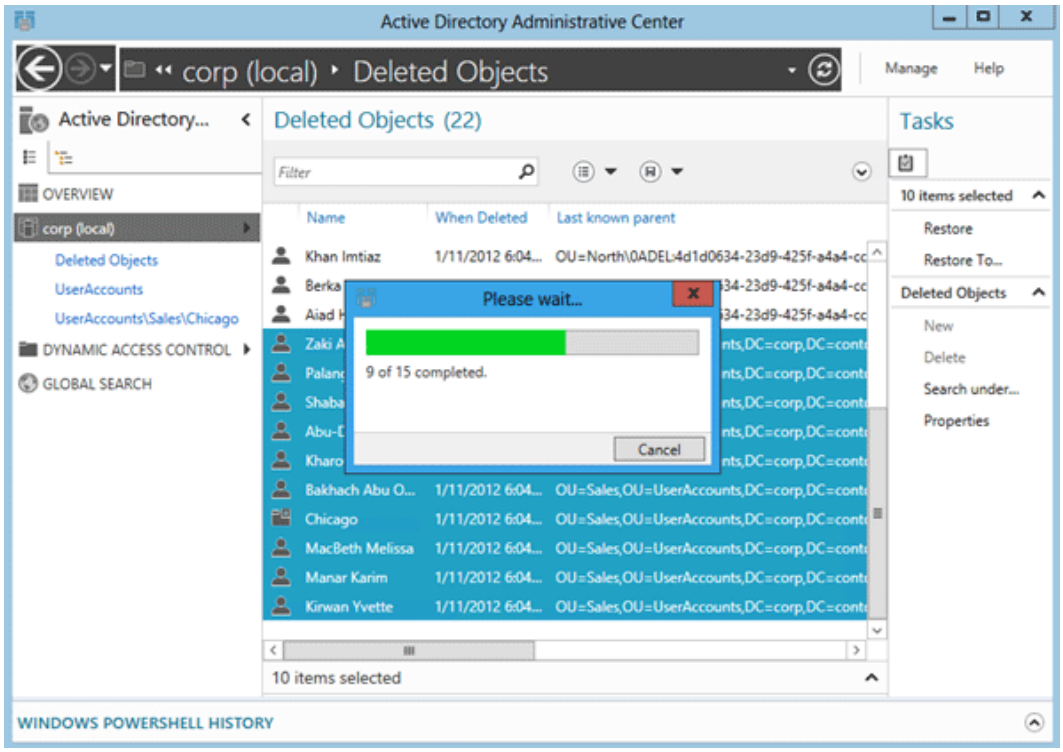

Multiple Parent and Child Objects

It is critical to understand the restoration process for a multi-parent-child restoration because the Active Directory Administrative Center cannot restore a nested tree of deleted objects with a single action.

1. Restore the top-most deleted object in a tree.
2. Restore the immediate children of that parent object.
3. Restore the immediate children of those parent objects.
4. Repeat as necessary until all objects restore.

You cannot restore a child object before restoring its parent. Attempting this restoration returns the following error:

The operation could not be performed because the object's parent is either uninstantiated or deleted.

The **Last Known Parent** attribute shows the parent relationship of each object. The **Last Known Parent** attribute changes from the deleted location to the restored location when you refresh the Active Directory Administrative Center after restoring a parent. Therefore, you can restore that child object when a parent object's location no longer shows the distinguished name of the deleted objects container.

Consider the scenario where an administrator accidentally deletes the Sales OU, which contains child OUs and users.

First, observe the value of the **Last Known Parent** attribute for all the deleted users and how it reads **OU=Sales\0ADEL:<guid+deleted objects container distinguished name>**:

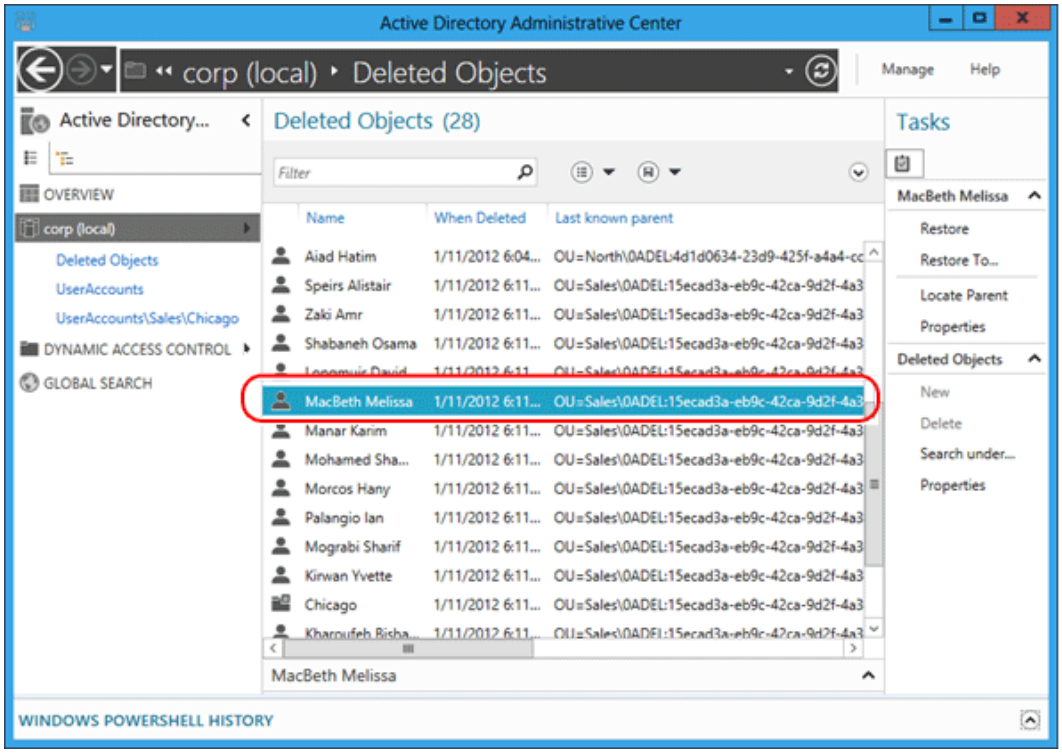

Filter on the ambiguous name Sales to return the deleted OU, which you then restore:

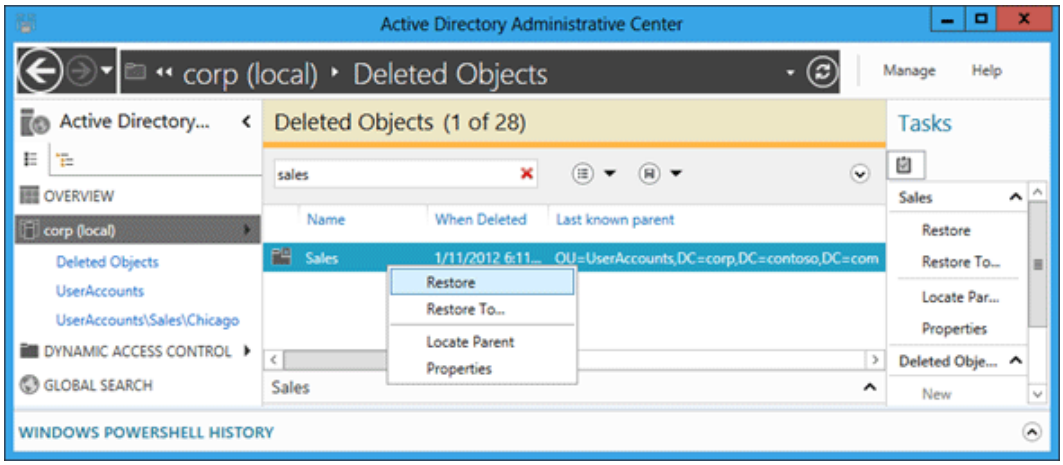

Refresh the Active Directory Administrative Center to see the deleted user object's Last Known Parent attribute change to the restored Sales OU distinguished name:

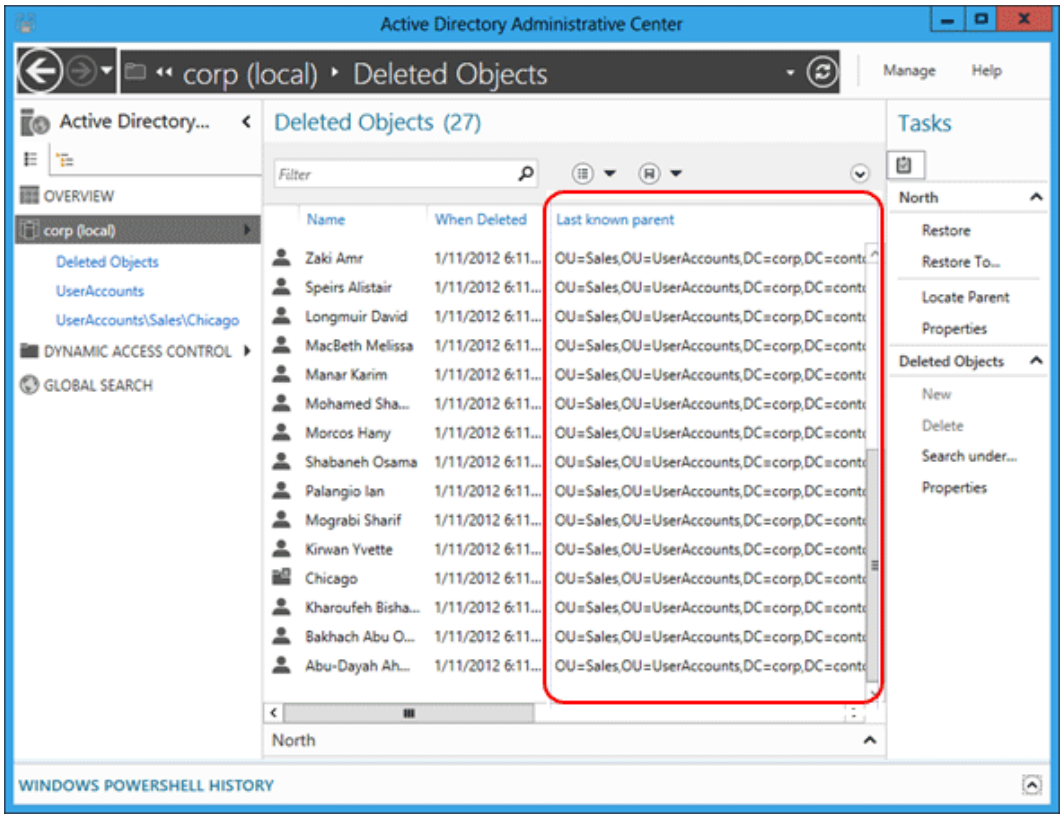

Filter on all the Sales users. Hold down the CTRL and A keys to select all the deleted Sales users. Click **Restore** to move the objects from the **Deleted Objects** container to the Sales OU with their group memberships and attributes intact.

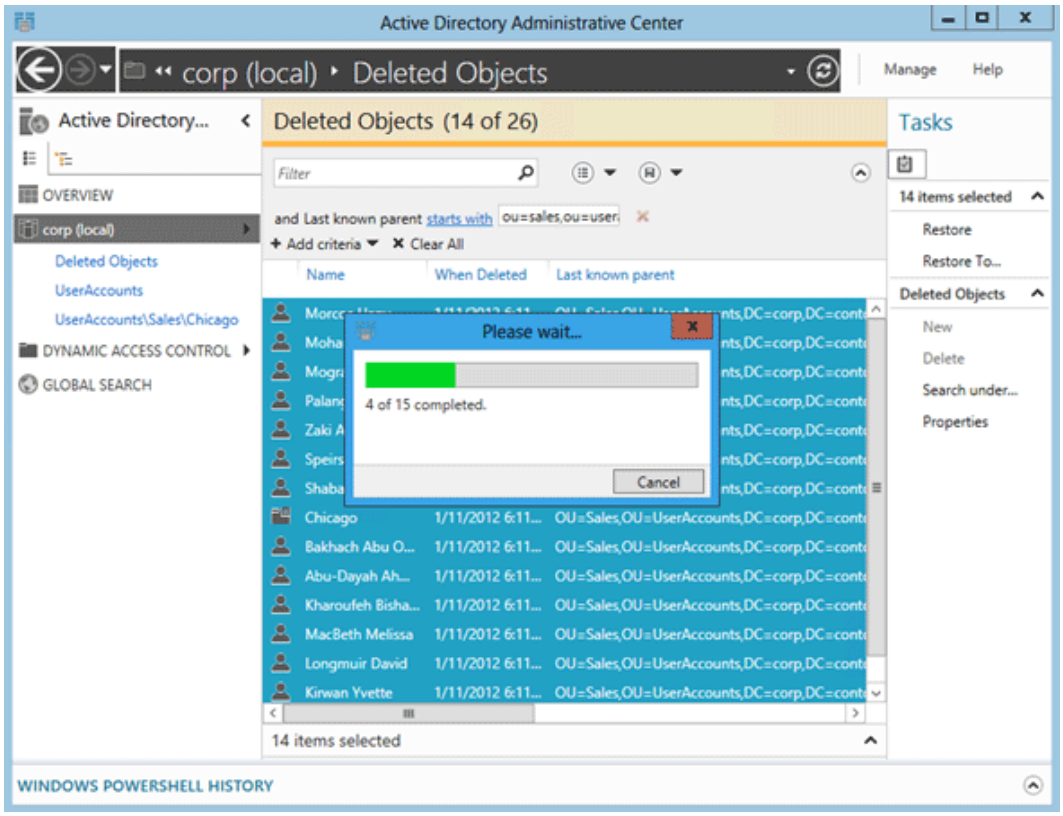

If the **Sales** OU contained child OUs of its own, then you would restore the child OUs first before restoring their children, and so on.

To restore all nested deleted objects by specifying a deleted parent container, see Appendix B: Restore Multiple, Deleted Active Directory Objects (Sample Script).

The Active Directory Windows PowerShell cmdlet for restoring deleted objects is:

```
Restore-adobject
```

The **Restore-ADObject** cmdlet functionality did not change between Windows Server 2008 R2 and Windows Server 2012.

Server-side Filtering

It is possible that over time, the Deleted Objects container will accumulate over 20,000 (or even 100,000) objects in medium and large enterprises and have difficulty showing all objects. Since the filter mechanism in Active Directory Administrative Center relies on client-side filtering, it cannot show these additional objects. To work around this limitation, use the following steps to perform a server-side search:

1. Right click the **Deleted Objects** container and click **Search under this node**.
2. Click the chevron to expose the **+Add criteria** menu, select and add **Last modified between given dates**. The Last Modified time (the **whenChanged** attribute) is a close approximation

of the deletion time; in most environments, they are identical. This query performs a server-side search.

3. Locate the deleted objects to restore by using further display filtering, sorting, and so on in the results, and then restore them normally.

Configuring and Managing Fine-Grained Password Policies Using Active Directory Administrative Center

Configuring Fine-Grained Password Policies

The Active Directory Administrative Center enables you to create and manage Fine-Grained Password Policy (FGPP) objects. Windows Server 2008 introduced the FGPP feature but Windows Server 2012 has the first graphical management interface for it. You apply Fine-Grained Password Policies at a domain level and it enables overriding the single domain password required by Windows Server 2003. By creating different FGPP with different settings, individual users or groups get differing password policies in a domain.

For information about the Fine-Grained Password Policy, see [AD DS Fine-Grained Password and Account Lockout Policy Step-by-Step Guide (Windows Server 2008 R2)].

In the Navigation pane, click Tree View, click your domain, click **System**, click **Password Settings Container**, and then in the Tasks pane, click **New** and **Password Settings**.

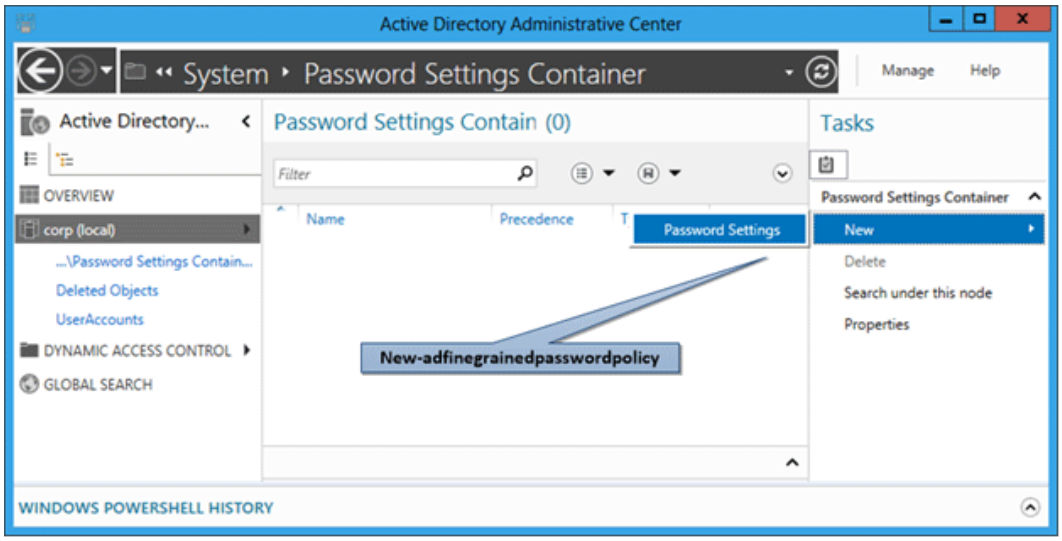

Managing Fine-Grained Password Policies

Creating a new FGPP or editing an existing one brings up the **Password Settings** editor. From here, you configure all desired password policies, as you would have in Windows Server 2008 or Windows Server 2008 R2, only now with a purpose-built editor.

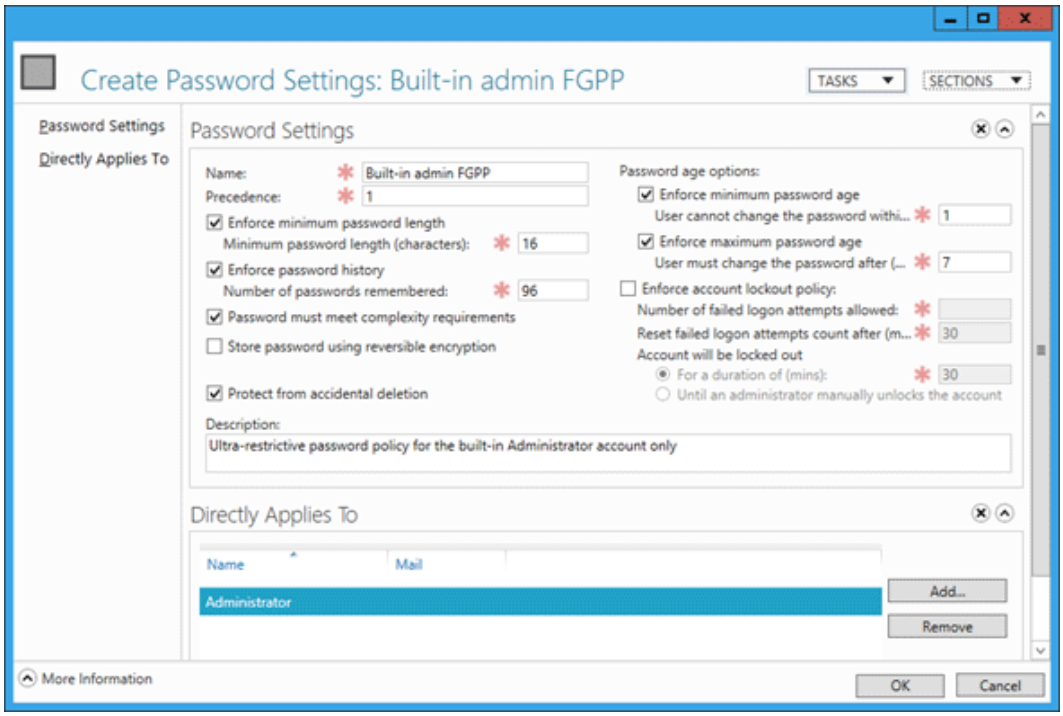

Fill out all required (red asterisk) fields and any optional fields, and then click **Add** to set the users or groups that receives this policy. FGPP overrides default domain policy settings for those specified security principals. In the figure above, an extremely restrictive policy applies only to the built-in Administrator account, to prevent compromise. The policy is far too complex for standard users to comply with, but is perfect for a high-risk account used only by IT professionals.

You also set precedence and to which users and groups the policy applies within a given domain.

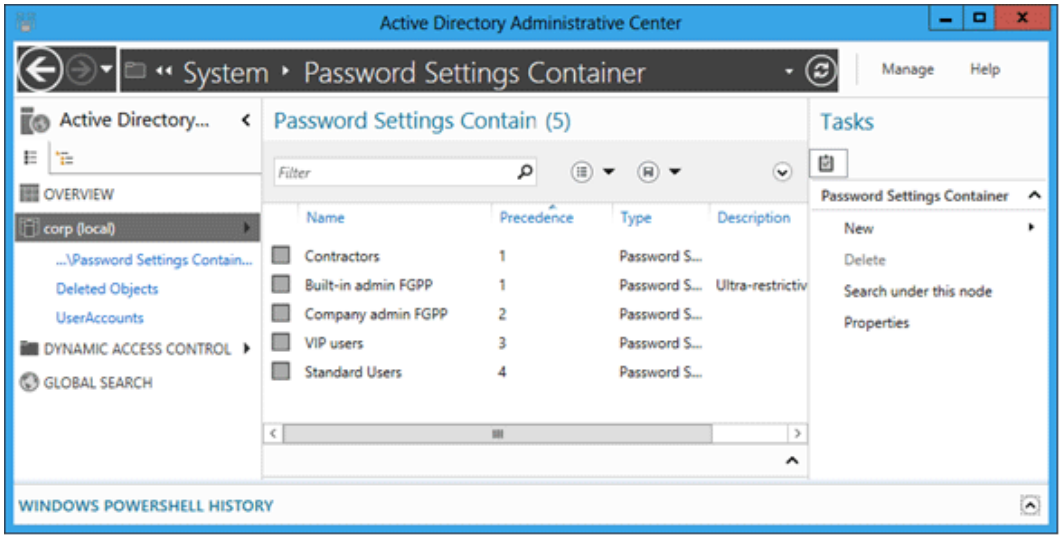

The Active Directory Windows PowerShell cmdlets for Fine-Grained Password Policy are:

```
Add-ADFineGrainedPasswordPolicySubject
```

```
Get-ADFineGrainedPasswordPolicy

Get-ADFineGrainedPasswordPolicySubject

New-ADFineGrainedPasswordPolicy

Remove-ADFineGrainedPasswordPolicy

Remove-ADFineGrainedPasswordPolicySubject

Set-ADFineGrainedPasswordPolicy
```

Fine-Grained Password Policy cmdlet functionality did not change between the Windows Server 2008 R2 and Windows Server 2012. As a convenience, the following diagram illustrates the associated arguments for cmdlets:

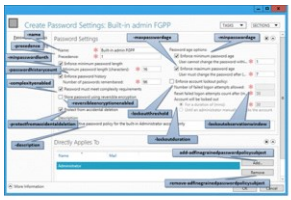

The Active Directory Administrative Center also enables you to locate the resultant set of applied FGPP for a specific user. Right click any user and click **View resultant password settings…** to open the *Password Settings* page that applies to that user through implicit or explicit assignment:

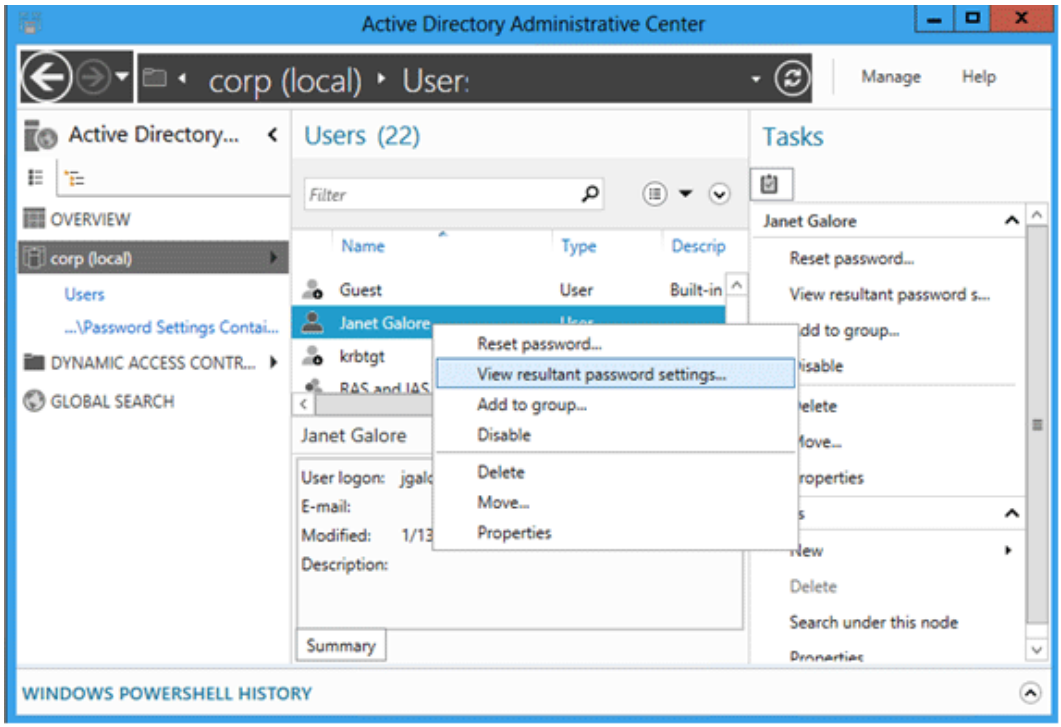

Examining the **Properties** of any user or group shows the **Directly Associated Password Settings**, which are the explicitly assigned FGPPs:

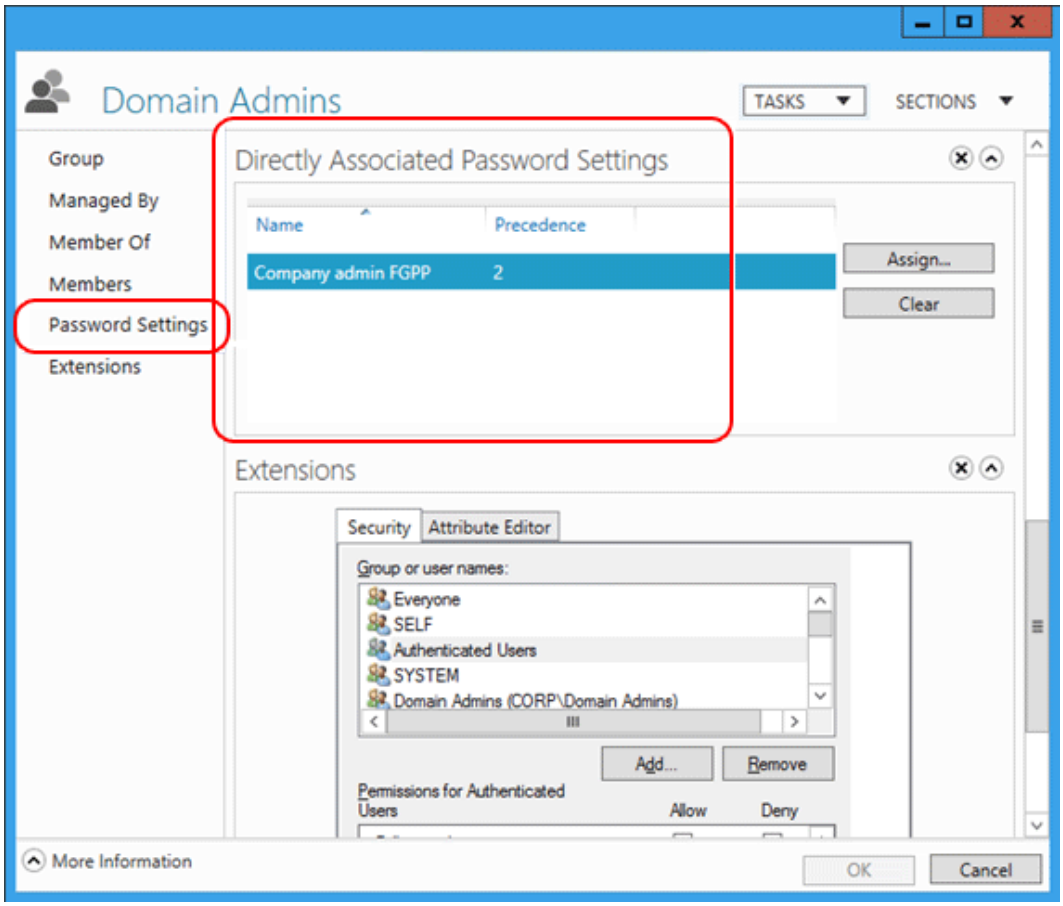

Implicit FGPP assignment does not display here; for that, you must use the **View resultant password settings…** option.

Using the Active Directory Administrative Center Windows PowerShell History Viewer

The future of Windows management is Windows PowerShell. By layering graphical tools on top of a task automation framework, management of the most complex distributed systems becomes consistent and efficient. You need to understand how Windows PowerShell works in order to reach your full potential and maximize your computing investments.

The Active Directory Administrative Center now provides a complete history of all the Windows PowerShell cmdlets it runs and their arguments and values. You can copy the cmdlet history elsewhere for study or modification and re-use. You can create Task notes to assist in isolating what your Active Directory Administrative Center commands resulted in Windows PowerShell. You can also filter the history to find points of interest.

The Active Directory Administrative Center Windows PowerShell History Viewer's purpose is for you to learn through practical experience.

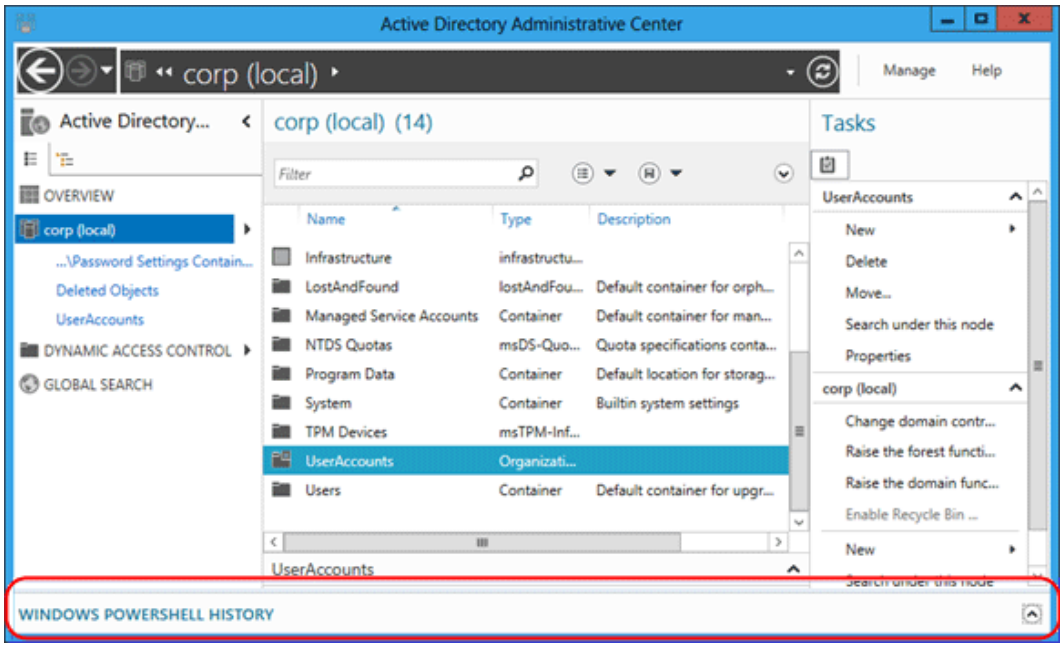

Click the chevron (arrow) to show Windows PowerShell History Viewer.

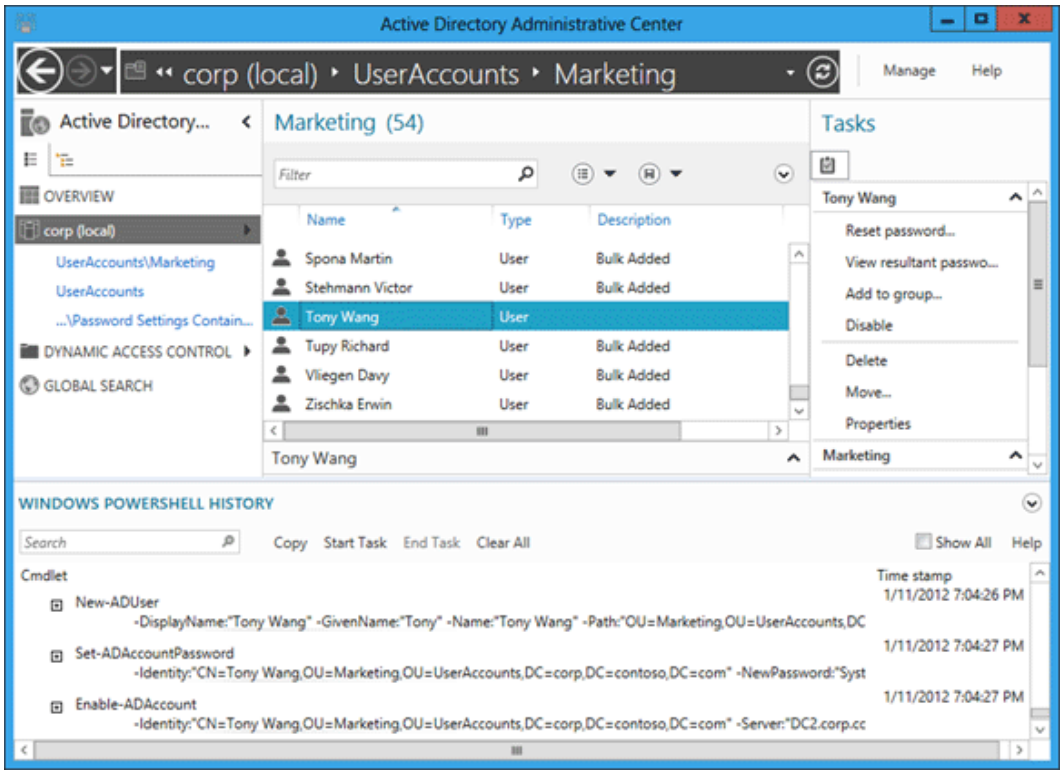

Then, create a user or modify a group's membership. The history viewer continually updates with a collapsed view of each cmdlet that the Active Directory Administrative Center ran with the arguments specified.

Expand any line item of interest to see all values provided to the cmdlet's arguments:

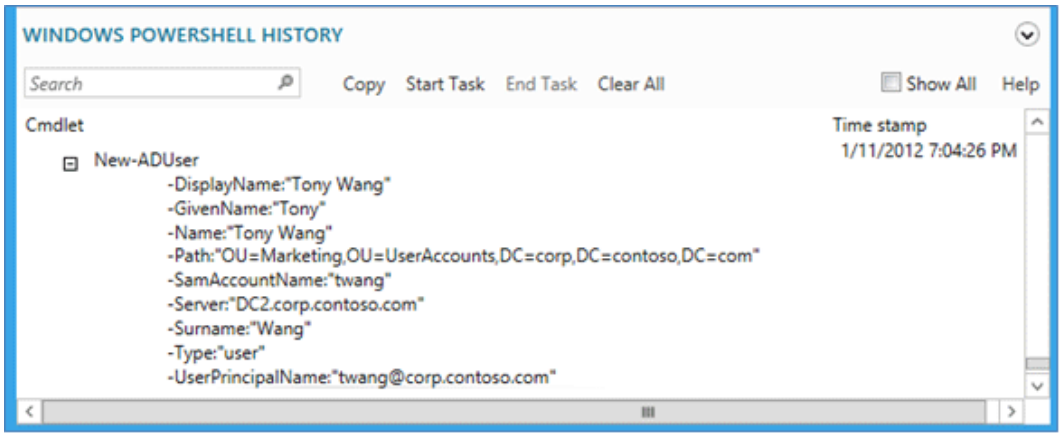

Click the **Start Task** menu to create a manual notation before you use Active Directory Administrative Center to create, modify, or delete an object. Type in what you were doing. When done with your change, select **End Task**. The task note groups all of those actions performed into a collapsible note you can use for better understanding.

For example, to see the Windows PowerShell commands used to change a user's password and remove him from a group:

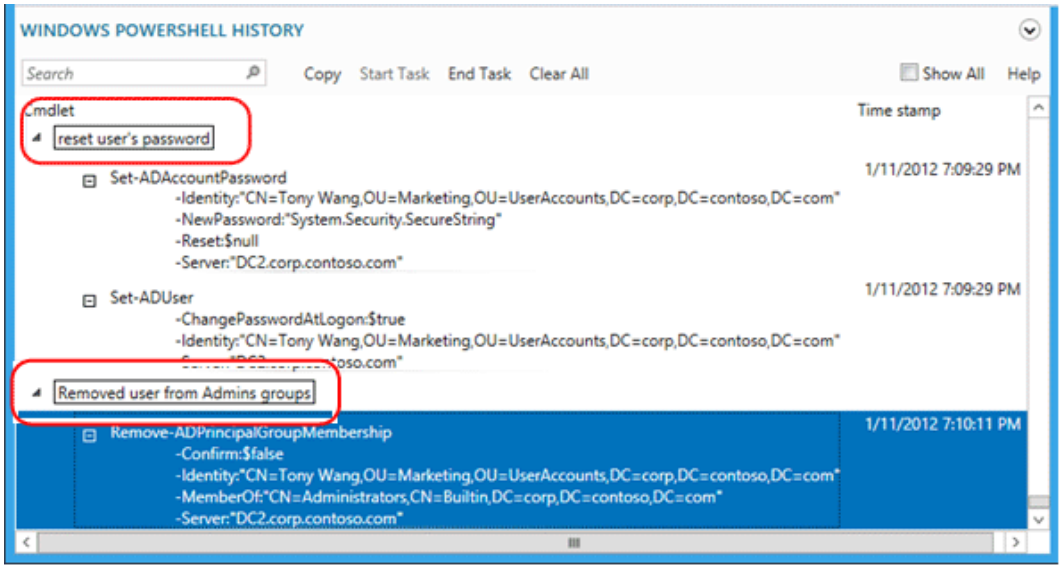

Selecting the Show All check box also shows the Get-* verb Windows PowerShell cmdlets that only retrieve data.

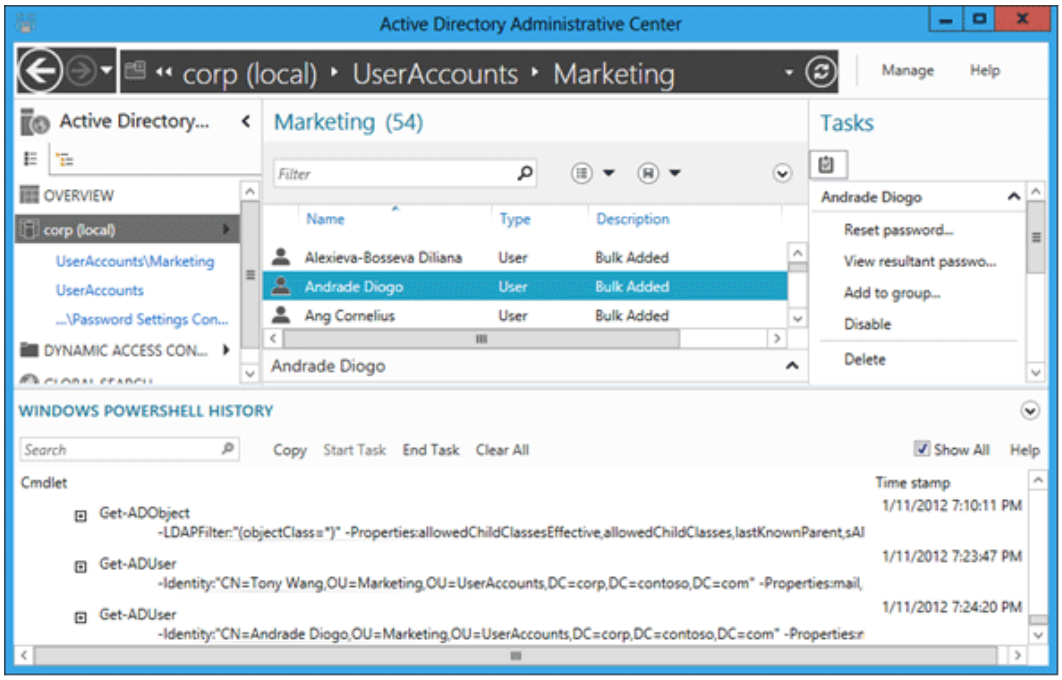

The history viewer shows the literal commands run by the Active Directory Administrative Center and you might note that some cmdlets appear to run unnecessarily. For example, you can create a new user with:

```
new-aduser
```

and do not need to use:

```
set-adaccountpassword
```

```
enable-adaccount
```

```
set-aduser
```

The Active Directory Administrative Center's design required minimal code usage and modularity. Therefore, instead of a set of functions that create new users and another set that modify existing users, it minimally does each function and then chains them together with the cmdlets. Keep this in mind when you are learning Active Directory Windows PowerShell. You can also use that as a learning technique, where you see how simply you can use Windows PowerShell to complete a single task.

Troubleshooting AD DS Management

Introduction to Troubleshooting

Because of its relative newness and lack of usage in existing customer environments, the Active Directory Administrative Center has limited troubleshooting options.

Troubleshooting Options

Logging Options

The Active Directory Administrative Center now contains built-in logging in Windows Server 2012, as part of a tracing config file. Create/modify the following file in the same folder as dsac.exe:

dsac.exe.config

Create the following contents:

```
<appSettings>
  <add key="DsacLogLevel" value="Verbose" />
</appSettings>
<system.diagnostics>
 <trace autoflush="false" indentsize="4">
  <listeners>
   <add name="myListener"
    type="System.Diagnostics.TextWriterTraceListener"
    initializeData="dsac.trace.log" />
   <remove name="Default" />
  </listeners>
 </trace>
</system.diagnostics>
```

The verbosity levels for **DsacLogLevel** are **None**, **Error**, **Warning**, **Info**, and **Verbose**. The output file name is configurable and writes to the same folder as dsac.exe. The output can tell you more about how ADAC is operating, which domain controllers it contacted, what Windows PowerShell commands executed, what the responses were, and further details.

For example, while using the INFO level, which returns all results except the trace-level verbosity:

- DSAC.exe starts
- Logging starts
- Domain Controller requested to return initial domain information

```
[12:42:49][TID 3][Info] Command Id, Action, Command, Time,
Elapsed Time ms (output), Number objects (output)
[12:42:49][TID 3][Info] 1, Invoke, Get-ADDomainController, 2012-
04-16T12:42:49
[12:42:49][TID 3][Info] Get-ADDomainController
-Discover:$null
-DomainName:"CORP"
```

-ForceDiscover:$null

-Service:ADWS

-Writable:$null

- Domain controller DC1 returned from domain Corp
- PS AD virtual drive loaded

```
[12:42:49][TID 3][Info] 1, Output, Get-ADDomainController, 2012-04-16T12:42:49, 1
[12:42:49][TID 3][Info] Found the domain controller 'DC1' in the domain 'CORP'.
[12:42:49][TID 3][Info] 2, Invoke, New-PSDrive, 2012-04-16T12:42:49
[12:42:49][TID 3][Info] New-PSDrive
-Name:"ADDrive0"
-PSProvider:"ActiveDirectory"
-Root:""
-Server:"dc1.corp.contoso.com"
[12:42:49][TID 3][Info] 2, Output, New-PSDrive, 2012-04-16T12:42:49, 1
[12:42:49][TID 3][Info] 3, Invoke, Get-ADRootDSE, 2012-04-16T12:42:49
```

- Get domain Root DSE Information

```
[12:42:49][TID 3][Info] Get-ADRootDSE
-Server:"dc1.corp.contoso.com"
[12:42:49][TID 3][Info] 3, Output, Get-ADRootDSE, 2012-04-16T12:42:49, 1
[12:42:49][TID 3][Info] 4, Invoke, Get-ADOptionalFeature, 2012-04-16T12:42:49
```

- Get domain AD recycle bin information

```
[12:42:49][TID 3][Info] Get-ADOptionalFeature
-LDAPFilter:"(msDS-OptionalFeatureFlags=1)"
-Server:"dc1.corp.contoso.com"
[12:42:49][TID 3][Info] 4, Output, Get-ADOptionalFeature, 2012-04-16T12:42:49, 1
```

```
[12:42:49][TID 3][Info] 5, Invoke, Get-ADRootDSE, 2012-04-
16T12:42:49
[12:42:49][TID 3][Info] Get-ADRootDSE
-Server:"dc1.corp.contoso.com"
[12:42:49][TID 3][Info] 5, Output, Get-ADRootDSE, 2012-04-
16T12:42:49, 1
[12:42:49][TID 3][Info] 6, Invoke, Get-ADRootDSE, 2012-04-
16T12:42:49
[12:42:49][TID 3][Info] Get-ADRootDSE
-Server:"dc1.corp.contoso.com"
[12:42:49][TID 3][Info] 6, Output, Get-ADRootDSE, 2012-04-
16T12:42:49, 1
[12:42:49][TID 3][Info] 7, Invoke, Get-ADOptionalFeature, 2012-
04-16T12:42:49
[12:42:49][TID 3][Info] Get-ADOptionalFeature
-LDAPFilter:"(msDS-OptionalFeatureFlags=1)"
-Server:"dc1.corp.contoso.com"
[12:42:50][TID 3][Info] 7, Output, Get-ADOptionalFeature, 2012-
04-16T12:42:50, 1
[12:42:50][TID 3][Info] 8, Invoke, Get-ADForest, 2012-04-
16T12:42:50
```

- Get AD forest

```
[12:42:50][TID 3][Info] Get-ADForest
-Identity:"corp.contoso.com"
-Server:"dc1.corp.contoso.com"
[12:42:50][TID 3][Info] 8, Output, Get-ADForest, 2012-04-
16T12:42:50, 1
[12:42:50][TID 3][Info] 9, Invoke, Get-ADObject, 2012-04-
16T12:42:50
```

- Get Schema information for supported encryption types, FGPP, certain user information

```
[12:42:50][TID 3][Info] Get-ADObject
-LDAPFilter:"(|(ldapdisplayname=msDS-
PhoneticDisplayName)(ldapdisplayname=msDS-
PhoneticCompanyName)(ldapdisplayname=msDS-
```

```
PhoneticDepartment)(ldapdisplayname=msDS-
PhoneticFirstName)(ldapdisplayname=msDS-
PhoneticLastName)(ldapdisplayname=msDS-
SupportedEncryptionTypes)(ldapdisplayname=msDS-
PasswordSettingsPrecedence))"
-Properties:lDAPDisplayName
-ResultPageSize:"100"
-ResultSetSize:$null
-
SearchBase:"CN=Schema,CN=Configuration,DC=corp,DC=contoso,DC=com
"
-SearchScope:"OneLevel"
-Server:"dc1.corp.contoso.com"
[12:42:50][TID 3][Info] 9, Output, Get-ADObject, 2012-04-
16T12:42:50, 7
[12:42:50][TID 3][Info] 10, Invoke, Get-ADObject, 2012-04-
16T12:42:50
```

- Get all information about the domain object to display to administrator who clicked on the domain head.

```
[12:42:50][TID 3][Info] Get-ADObject
-IncludeDeletedObjects:$false
-LDAPFilter:"(objectClass=*)"
-
Properties:allowedChildClassesEffective,allowedChildClasses,last
KnownParent,sAMAccountType,systemFlags,userAccountControl,displa
yName,description,whenChanged,location,managedBy,memberOf,primar
yGroupID,objectSid,msDS-User-Account-Control-
Computed,sAMAccountName,lastLogonTimestamp,lastLogoff,mail,accou
ntExpires,msDS-PhoneticCompanyName,msDS-PhoneticDepartment,msDS-
PhoneticDisplayName,msDS-PhoneticFirstName,msDS-
PhoneticLastName,pwdLastSet,operatingSystem,operatingSystemServi
cePack,operatingSystemVersion,telephoneNumber,physicalDeliveryOf
ficeName,department,company,manager,dNSHostName,groupType,c,l,em
ployeeID,givenName,sn,title,st,postalCode,managedBy,userPrincipa
lName,isDeleted,msDS-PasswordSettingsPrecedence
-ResultPageSize:"100"
```

```
-ResultSetSize:"20201"
-SearchBase:"DC=corp,DC=contoso,DC=com"
-SearchScope:"Base"
-Server:"dc1.corp.contoso.com"
```

Setting the Verbose level also shows the .NET stacks for each function, but these do not include enough data to be particularly useful except when troubleshooting the Dsac.exe suffering an access violation or crash.

Important
There is also an out-of-band version of the service called the Active Directory Management Gateway, which runs on Windows Server 2008 SP2 and Windows Server 2003 SP2.

The two likely causes of this issue are:

- The ADWS service is not running on any accessible domain controllers.
- Network communications are blocked to the ADWS service from the computer running the Active Directory Administrative Center

The errors shown when no Active Directory Web Services instances are available are:

Error	Operation
"Cannot connect to any domain. Refresh or try again when connection is available"	Shown at start of the Active Directory Administrative Center application
"Cannot find an available server in the *<NetBIOS domain name>* domain that is running the Active Directory Web Service (ADWS)"	Shown when trying to select a domain node in the Active Directory Administrative Center application

To troubleshoot this issue, use these steps:

1. Validate the Active Directory Web Services service is started on at least one domain controller in the domain (and preferably all domain controllers in the forest). Ensure that it is set to start automatically on all domain controllers as well.
2. From the computer running the Active Directory Administrative Center, validate that you can locate a server running ADWS by running these NLTest.exe commands:

   ```
   nltest /dsgetdc:<domain NetBIOS name> /ws /force
   nltest /dsgetdc:<domain fully qualified DNS name> /ws /force
   ```

   If those tests fail even though the ADWS service is running, the issue is with name resolution or LDAP and not ADWS or Active Directory Administrative Center. This test fails with error

"1355 0x54B ERROR_NO_SUCH_DOMAIN" if ADWS is not running on any domain controllers though, so double-check before reaching any conclusions.

3. On the domain controller returned by NLTest, dump the listening port list with command:

```
Netstat -anob > ports.txt
```

Examine the ports.txt file and validate that the ADWS service is listening on port 9389. Example:

```
TCP    0.0.0.0:9389    0.0.0.0:0    LISTENING    1828
[Microsoft.ActiveDirectory.WebServices.exe]

TCP    [::]:9389       [::]:0       LISTENING    1828
[Microsoft.ActiveDirectory.WebServices.exe]
```

If listening, validate the Windows Firewall rules and ensure that they allow 9389 TCP inbound. By default, domain controllers enable firewall rule "Active Directory Web Services (TCP-in)". If not listening, validate again that the service is running on this server and restart it. Validate that no other process is already listening on port 9389.

4. Install NetMon or another network capture utility on the computer running Active Directory Administrative Center and on the domain controller returned by NLTEST. Gather simultaneous network captures from both computers, where you start Active Directory Administrative Center and see the error before stopping the captures. Validate that the client is able to send to and receive from the domain controller on port TCP 9389. If packets are sent but never arrive, or arrive and the domain controller replies but they never reach the client, it is likely there is a firewall in between the computers on the network dropping packets on that port. This firewall may be software or hardware, and may be part of third party endpoint protection (antivirus) software.

Known/Likely Issues and Support Scenarios

The only likely issue with the Active Directory Administrative Center is an inability to connect to the Active Directory Web Service (ADWS) running on a Windows Server 2012 or Windows Server 2008 R2 domain controller.

See Also

Introduction to Active Directory Administrative Center Enhancements (Level 100)

Active Directory Replication and Topology Management Using Windows PowerShell

Windows PowerShell for Active Directory now includes support for replication and topology management. The following topics provide an introduction and additional details:

- Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100)
- Advanced Active Directory Replication and Topology Management Using Windows PowerShell (Level 200)

Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100)

Windows PowerShell for Active Directory includes the ability to manage replication, sites, domains and forests, domain controllers, and partitions. Users of prior management tools such as the Active Directory Sites and Services snap-in and repadmin.exe will notice that similar functionality is now available from within the Windows PowerShell for Active Directory context. In addition, the cmdlets are compatible with the existing Windows PowerShell for Active Directory cmdlets, thus creating a streamlined experience and allowing customers to easily create automation scripts.

Notes

The Windows PowerShell for Active Directory replication and topology cmdlets are available in the following environments:

- Windows Server 2012 domain controller
- Windows Server 2012 with the Remote Server Administration Tools for AD DS and AD LDS installed.
- Windows® 8 with the Remote Server Administration Tools for AD DS and AD LDS installed.

Installing the Active Directory Module for Windows PowerShell

The Active Directory Module for Windows PowerShell is installed by default when the AD DS server role is installed on a server that runs Windows Server 2012. No additional steps are required other than adding the server role. You can also install the Active Directory Module on a server that runs Windows Server 2012 by installing the Remote Server Administration Tools, and you can install the Active Directory Module on a computer running Windows 8 by downloading

and installing the Remote Server Administrative Tools (RSAT). See Instructionsfor installation steps.

Scenarios for testing Windows PowerShell for Active Directory replication and topology management cmdlets

The following scenarios are designed for administrators to familiarize themselves with the new management cmdlets:

- Get a list of all domain controllers and their corresponding sites
- Manage replication topology
- View replication status and information

Lab Requirements

- Two Windows Server 2012 domain controllers: **DC1** and **DC2** that are part of the contoso.com domain and reside in the CORPORATE site within that domain.

View domain controllers and their sites

In this step, you will use the Active Directory Module for Windows PowerShell to view the existing domain controllers and the replication topology for the domain.

To complete the steps in the following procedures, you must be a member of the Domain Admins group or have equivalent permissions.

To view all Active Directory sites

1. On **DC1**, click **Windows PowerShell** on the taskbar.
2. Type the following command:

   ```
   Get-ADReplicationSite -Filter *
   ```

   This returns detailed information about each site. The `Filter` parameter is used throughout Active Directory PowerShell cmdlets to limit the list of objects returned. In this case, the asterisk (*) indicates all site objects.

   Tip

   You can use the Tab key to auto-complete commands in Windows PowerShell.

   Example: Type `Get-ADRep` and press Tab multiple times to skip through the matching commands until you reach `Get-ADReplicationSite`. Auto-complete also works for parameter names such as `Filter`.

   To format the output from the `Get-ADReplicationSite` command as a table and limit the display to specific fields, you can pipe the output to the `Format-Table` command (or "`ft`" for short):

```
Get-ADReplicationSite -Filter * | ft Name
```

This returns a shorter version of the site list, including only the Name field.

To produce a table of all domain controllers

- Type the following command at the **Active Directory module for Windows PowerShell** prompt:

  ```
  Get-ADDomainController -Filter * | ft Hostname,Site
  ```

  This command returns the domain controllers host name as well as their site associations.

Manage replication topology

In the previous step, after running the command, `Get-ADDomainController -Filter * | ft Hostname,Site`, **DC2** was listed as part of the **CORPORATE** site. In the procedures below, you will create a new branch office site, **BRANCH1**, create a new site link, set the site link cost and replication frequency and then move **DC2** to **BRANCH1**.

To complete the steps in the following procedures, you must be a member of the Domain Admins group or have equivalent permissions.

To create a new site

- Type the following command at the **Active Directory module for Windows PowerShell** prompt:

  ```
  New-ADReplicationSite BRANCH1
  ```

  This command creates the new branch office site, branch1.

To create a new site link

- Type the following command at the **Active Directory module for Windows PowerShell** prompt:

  ```
  New-ADReplicationSiteLink 'CORPORATE-BRANCH1'  -SitesIncluded CORPORATE,BRANCH1 -
  OtherAttributes @{'options'=1}
  ```

  This command created the site link to **BRANCH1** and turned on the change notification process.

  Tip
  Use Tab to auto-complete parameter names such as `-SitesIncluded` and `-OtherAttributes` rather than typing them out manually.

To set the site link cost and replication frequency

- Type the following command at the **Active Directory module for Windows PowerShell** prompt:

```
Set-ADReplicationSiteLink CORPORATE-BRANCH1 -Cost 100 -
ReplicationFrequencyInMinutes 15
```

This command sets the site link cost to **BRANCH1** at **100** and set the replication frequency with the site to **15 minutes**.

To move a domain controller to a different site

- Type the following command at the **Active Directory module for Windows PowerShell** prompt:

  ```
  Get-ADDomainController DC2 | Move-ADDirectoryServer -Site BRANCH1
  ```

  This command moves the domain controller, **DC2** to the **BRANCH1** site.

Verification

To verify site creation, new site link, and cost and replication frequency

- Click **Server Manager**, click **Tools** and then click **Active Directory Sites and Services** and verify the following:

  Verify that the **BRANCH1** site contains all of the correct values from the Windows PowerShell commands.

  Verify the **CORPORATE-BRANCH1** site link is created and connects the **BRANCH1** and **CORPORATE** sites.

  Verify **DC2** is now in the **BRANCH1** site. Alternatively, you can open the **Active Directory Module for Windows PowerShell** and type the following command to verify **DC2** is now in the **BRANCH1** site: `Get-ADDomainController -Filter * | ft Hostname,Site`.

View replication status information

In the following procedures, you will use one of the Windows PowerShell for Active Directory replication and management cmdlets, `Get-ADReplicationUpToDatenessVectorTable DC1`, to produce a simple replication report using the up-to-dateness vector table maintained by each domain controller. This up-to-dateness vector table keeps track of the highest originating write USN seen from each domain controller in the forest.

To complete the steps in the following procedures, you must be a member of the Domain Admins group or have equivalent permissions.

To view the up-to-dateness vector table for a single domain controller

1. Type the following command at the **Active Directory module for Windows PowerShell** prompt:

   ```
   Get-ADReplicationUpToDatenessVectorTable DC1
   ```

This shows a list of the highest USNs seen by **DC1** for every domain controller in the forest. The **Server** value refers to the server maintaining the table, in this case **DC1**. The **Partner** value refers to the replication partner (direct or indirect) on which changes were made. The UsnFilter value is the highest USN seen by **DC1** from Partner. If a new domain controller is added to the forest, it will not appear in **DC1**'s table until **DC1** receives a change that originated from the new domain.

To view the up-to-dateness vector table for all domain controllers in a domain

1. Type the following command at the Active Directory module for Windows PowerShell prompt:

   ```
   Get-ADReplicationUpToDatenessVectorTable * | sort Partner,Server | ft
   Partner,Server,UsnFilter
   ```

   This command replaces **DC1** with `*`, thus collecting the up-to-dateness vector table data from all domain controllers. The data is sorted by **Partner** and **Server** and then displayed in a table.

   The sorting allows you to easily compare the last USN seen by each domain controller for a given replication partner. This is a quick way to check that replication is occurring across your environment. If replication is working correctly, the UsnFilter values reported for a given replication partner should be fairly similar across all domain controllers.

See Also

Advanced Active Directory Replication and Topology Management Using Windows PowerShell (Level 200)

Advanced Active Directory Replication and Topology Management Using Windows PowerShell (Level 200)

This topic explains the new AD DS replication and topology management cmdlets in more detail, and provides additional examples. For an introduction, see Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100).

1. Introduction
2. Replication and Metadata
3. Get-ADReplicationAttributeMetadata
4. Get-ADReplicationPartnerMetadata
5. Get-ADReplicationFailure
6. Get-ADReplicationQueueOperation and Get-ADReplicationUpToDatenessVectorTable
7. Sync-ADObject

8. Topology

Introduction

Windows Server 2012 extends the Active Directory module for Windows PowerShell with twenty-five new cmdlets to manage replication and forest topology. Prior to this, you were forced to use the generic ***-AdObject** nouns or call .NET functions.

Like all Active Directory Windows PowerShell cmdlets, this new functionality requires installing the Active Directory Management Gateway Service on at least one domain controller (and preferably, all domain controllers).

The following table lists new replication and topology cmdlets added to the Active Directory Windows PowerShell module.

Cmdlet	Explanation
Get-ADReplicationAttributeMetadata	Returns attribute replication metadata for an object
Get-ADReplicationConnection	Returns domain controller connection object details
Get-ADReplicationFailure	Returns the most replication recent failure for a domain controller
Get-ADReplicationPartnerMetadata	Returns replication configuration of a domain controller
Get-ADReplicationQueueOperation	Returns the current replication queue backlog
Get-ADReplicationSite	Returns site information
Get-ADReplicationSiteLink	Returns site link information
Get-ADReplicationSiteLinkBridge	Returns site link bridge information
Get-ADReplicationSubnet	Returns AD subnet information
Get-ADReplicationUpToDatenessVectorTable	Returns the UTD vector for a domain controller
Get-ADTrust	Returns information about an inter-domain or inter-forest trust
New-ADReplicationSite	Creates a new site
New-ADReplicationSiteLink	Creates a new site link
New-ADReplicationSiteLinkBridge	Creates a new site link bridge
New-ADReplicationSubnet	Creates a new AD subnet
Remove-ADReplicationSite	Deletes a site

Remove-ADReplicationSiteLink	Deletes a site link
Remove-ADReplicationSiteLinkBridge	Deletes a site link bridge
Remove-ADReplicationSubnet	Deletes an AD subnet
Set-ADReplicationConnection	Modifies a connection
Set-ADReplicationSite	Modifies a site
Set-ADReplicationSiteLink	Modifies a site link
Set-ADReplicationSiteLinkBridge	Modifies a site link bridge
Set-ADReplicationSubnet	Modifies an AD subnet
Sync-ADObject	Forces replication of a single object

Most of these cmdlets have their basis in Repadmin.exe. Other cmdlets (not listed) handle features like Dynamic Access Control and Group Managed Service Accounts.

For a complete list of all Active Directory Windows PowerShell cmdlets, run:

```
Get-command -module ActiveDirectory
```

For a complete list of all Active Directory Windows PowerShell cmdlet arguments, reference the help. For example:

```
Get-help New-ADReplicationSite
```

Use the `Update-Help` cmdlet to download and install help files

Replication and Metadata

Repadmin.exe validates the health and consistency of Active Directory replication. Repadmin.exe offers simple data manipulation options - some arguments support CSV outputs, for example - but automation generally required parsing through text file outputs. The Active Directory module for Windows PowerShell is the first attempt at offering an option that allows real control over the returned data; prior to this, you had to create scripts or use third party tools.

Additionally, the following cmdlets implement a new parameter set of **Target**, **Scope**, and **EnumerationServer**:

- **Get-ADReplicationFailure**
- **Get-ADReplicationPartnerMetadata**
- **Get-ADReplicationUpToDatenessVectorTable**

The **Target** argument accepts a comma-separated list of strings that identify the target servers, sites, domains, or forests specified by the **Scope** argument. An asterisk (*) is also permissible and means all servers within the specified scope. If no scope is specified, it implies all servers in the current user's forest. The **Scope** argument specifies the latitude of the search. Acceptable values are **Server**, **Site**, **Domain**, and **Forest**. The **EnumerationServer** specifies the server that

enumerates the list of domain controllers specified in **Target** and **Scope**. It operates the same as the **Server** argument and requires the specified server run the Active Directory Web Service.

To introduce the new cmdlets, here are some sample scenarios showing capabilities impossible to repadmin.exe; armed with these illustrations, the administrative possibilities become obvious. Review the cmdlet help for specific usage requirements.

Get-ADReplicationAttributeMetadata

This cmdlet is similar to **repadmin.exe /showobjmeta**. It enables you to return replication metadata, such as when an attribute changed, the originating domain controller, the version and USN information, and attribute data. This cmdlet is useful for auditing where and when a change occurred.

Unlike Repadmin, Windows PowerShell gives flexible search and output control. For example, you can output the metadata of the Domain Admins object, ordered as a readable list:

```
Get-ADReplicationAttributeMetadata -object "cn=domain
admins,cn=users,dc=corp,dc=contoso,dc=com" -server dc1.corp.contoso.com -
showalllinkedvalues | format-list
```

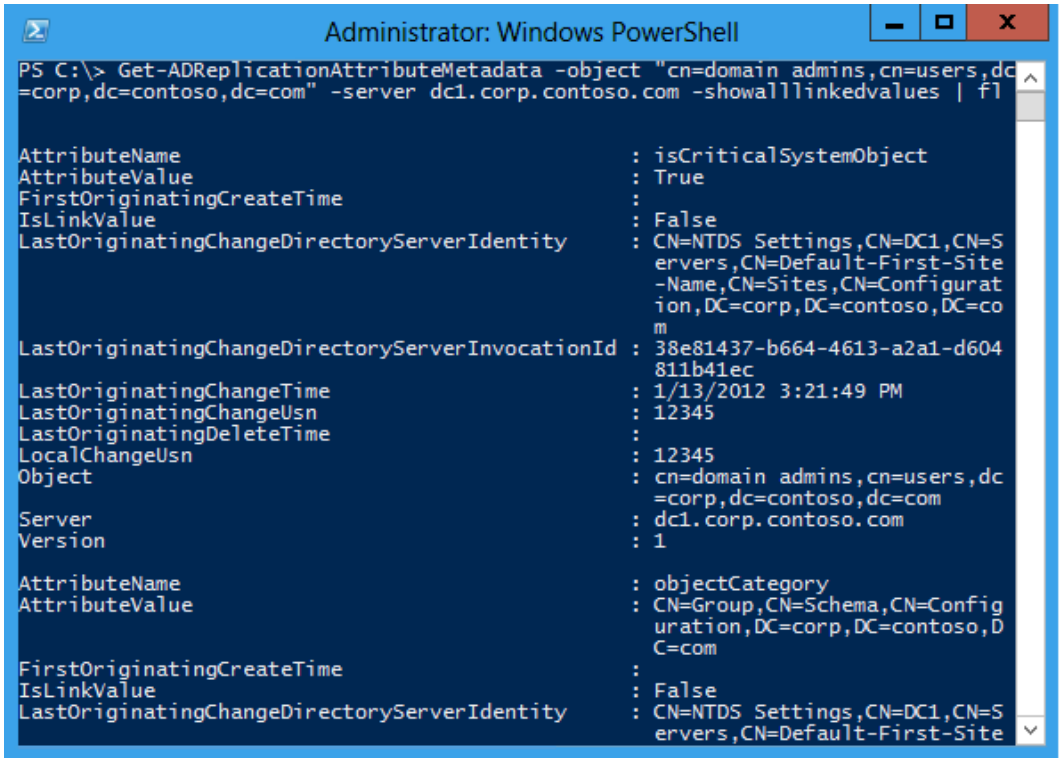

Alternatively, you can arrange the data to look like repadmin, in a table:

```
Get-ADReplicationAttributeMetadata -object "cn=domain
admins,cn=users,dc=corp,dc=contoso,dc=com" -server dc1.corp.contoso.com -
showalllinkedvalues | format-table -wrap
```

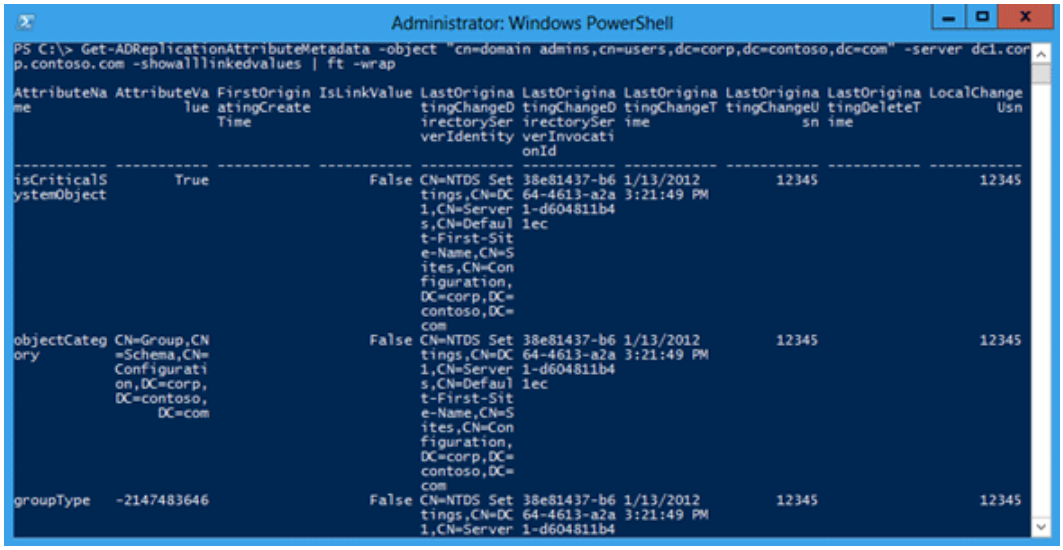

Alternatively, you can get metadata for an entire class of objects, by pipelining the **Get-Adobject** cmdlet with a filter, such as all groups - then combine that with a specific date. The pipeline is a channel used between multiple cmdlets to pass data. To see all groups modified in some fashion on January 13th, 2012:

```
get-adobject -filter 'objectclass -eq "group"' | Get-ADReplicationAttributeMetadata -
server dc1.corp.contoso.com | where-object {$_.lastoriginatingchangetime -like
"*1/13/2012*" -and $_.attributename -eq "name"} | format-table object
```

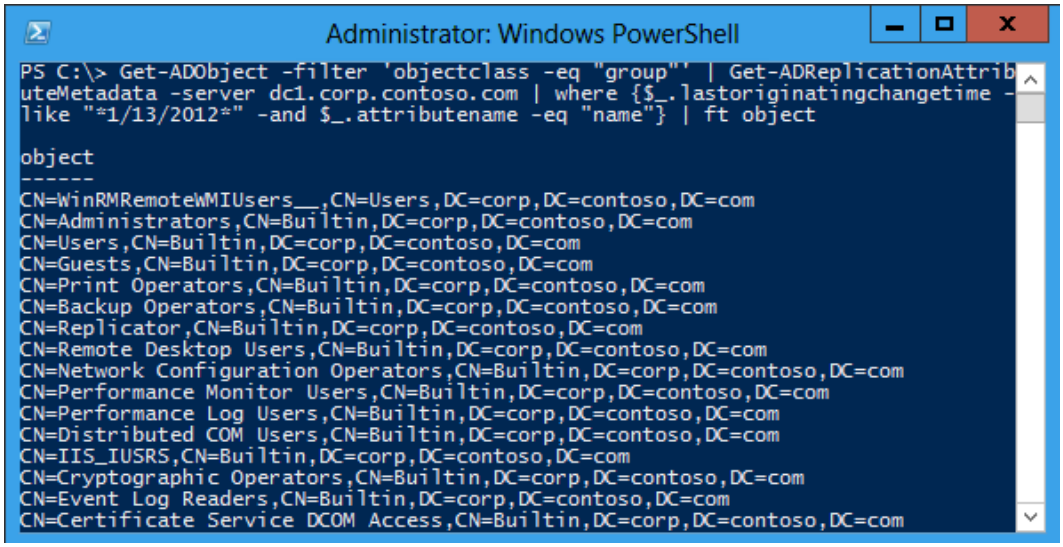

For more information about more Windows PowerShell operations with pipelines, see Piping and the Pipeline in Windows PowerShell.

Alternatively, to find out every group that has Tony Wang as a member and when the group was last modified:

```
get-adobject -filter 'objectclass -eq "group"' | Get-ADReplicationAttributeMetadata -server dc1.corp.contoso.com -showalllinkedvalues | where-object {$_.attributevalue -like "*tony wang*"} | format-table object,LastOriginatingChangeTime,version -auto
```

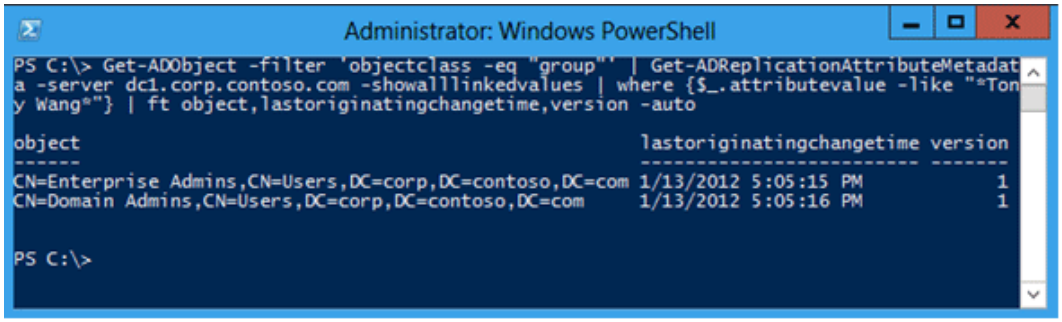

Alternatively, to find all objects authoritatively restored using a system state backup in the domain, based on their artificially high version:

```
get-adobject -filter 'objectclass -like "*"' | Get-ADReplicationAttributeMetadata -server dc1.corp.contoso.com | where-object {$_.version -gt "100000" -and $_.attributename -eq "name"} | format-table object,LastOriginatingChangeTime
```

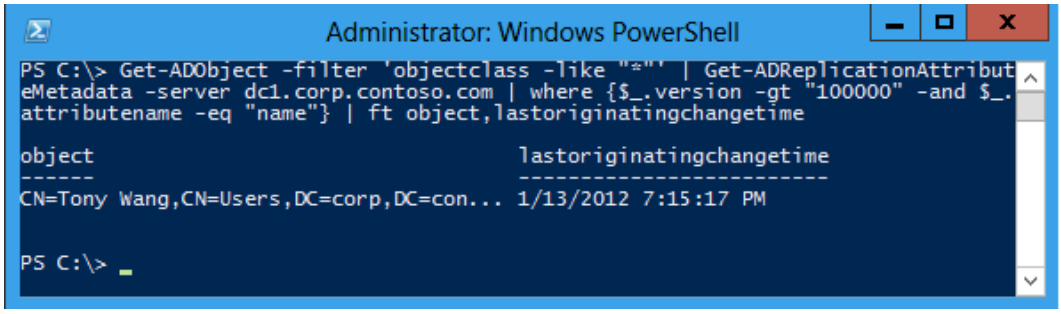

Alternatively, send all user metadata to a CSV file for later examination in Microsoft Excel:

```
get-adobject -filter 'objectclass -eq "user"' | Get-ADReplicationAttributeMetadata -server dc1.corp.contoso.com -showalllinkedvalues | export-csv allgroupmetadata.csv
```

Get-ADReplicationPartnerMetadata

This cmdlet returns information about the configuration and state of replication for a domain controller, allowing you to monitor, inventory, or troubleshoot. Unlike Repadmin.exe, using Windows PowerShell means you see only the data that is important to you, in the format you want.

For example, the readable replication state of a single domain controller:

```
Get-ADReplicationPartnerMetadata -target dc1.corp.contoso.com
```

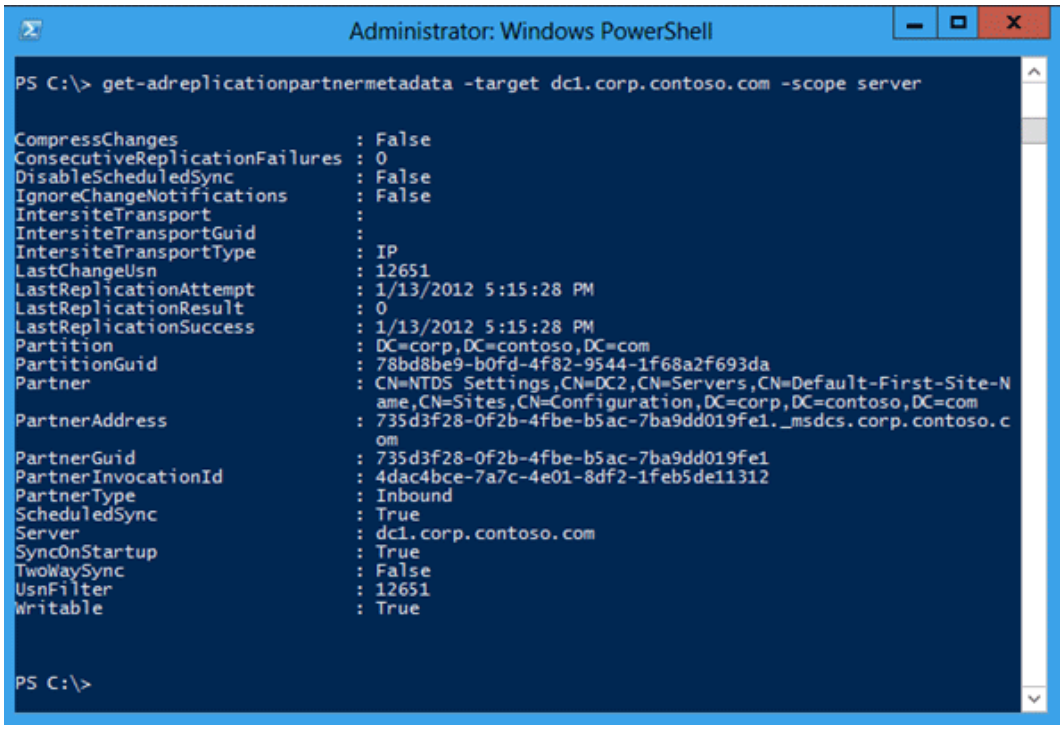

Alternatively, the last time a domain controller replicated inbound and its partners, in a table format:

```
Get-ADReplicationPartnerMetadata -target dc1.corp.contoso.com | format-table
lastreplicationattempt,lastreplicationresult,partner -auto
```

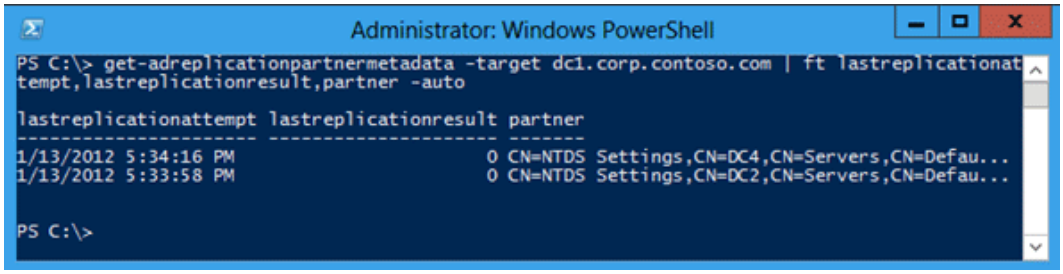

Alternatively, contact all domain controllers in the forest and display any whose last attempted replication failed for any reason:

```
Get-ADReplicationPartnerMetadata -target * -scope server | where
{$_.lastreplicationresult -ne "0"} | ft
server,lastreplicationattempt,lastreplicationresult,partner -auto
```

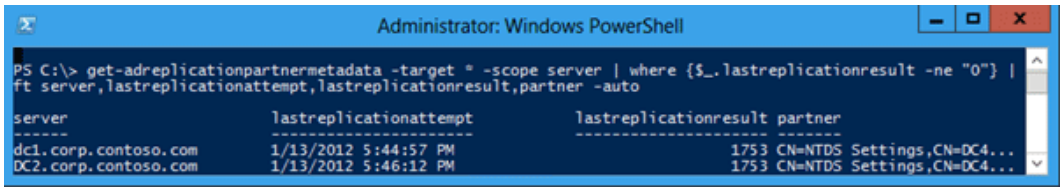

Get-ADReplicationFailure

This cmdlet can be used to returns information about recent errors in replication. It is analogous to **Repadmin.exe /showreplsum**, but again, with much more control thanks to Windows PowerShell.

For example, you can return a domain controller's most recent failures and the partners he failed contacting:

```
Get-ADReplicationFailure dc1.corp.contoso.com
```

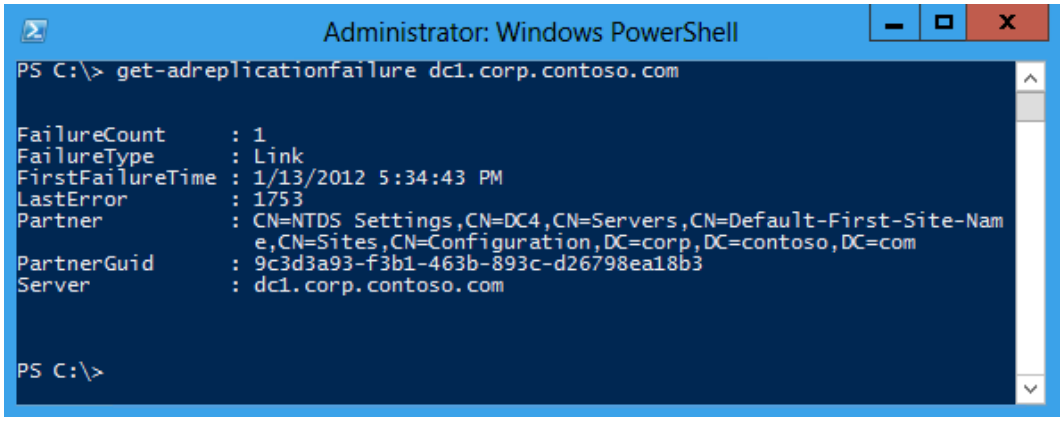

Alternatively, return a table view for all servers in a specific AD logical site, ordered for easier viewing and containing only the most critical data:

```
Get-ADReplicationFailure -scope site -target default-first-site-name | format-table server,firstfailuretime,failurecount,lasterror,partner -auto
```

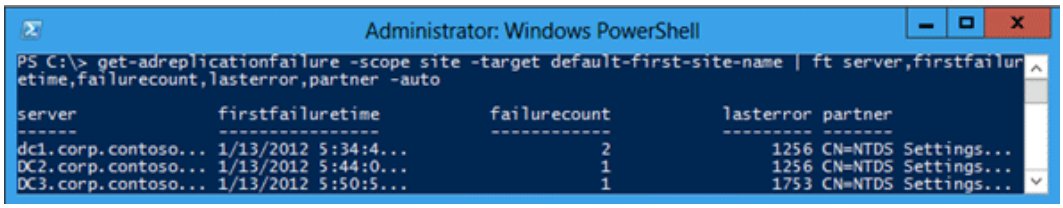

Get-ADReplicationQueueOperation and Get-ADReplicationUpToDatenessVectorTable

Both of these cmdlets returns further aspects of domain controller "up to dateness", which includes pending replication and version vector information.

Sync-ADObject

This cmdlet is analogous to running **Repadmin.exe /replsingleobject**. It is very useful when you make changes that require out of band replication, especially to fix an issue.

For example, if someone deleted the CEO's user account and then restored it with the Active Directory Recycle Bin, you probably want it replicated to all domain controllers immediately. You also probably want to do this without forcing replication of all the other object changes made ; after all, that is why you have a replication schedule - to avoid overloading WAN links.

```
Get-ADDomainController -filter * | foreach {Sync-ADObject -object "cn=tony
wang,cn=users,dc=corp,dc=contoso,dc=com" -source dc1 -destination $_.hostname}
```

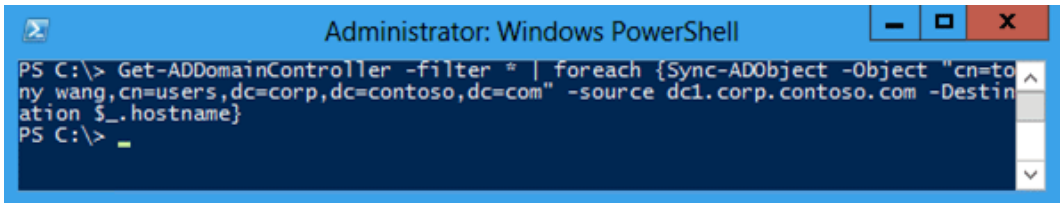

Topology

While Repadmin.exe is good at returning information about replication topology like sites, site links, site link bridges, and connections, it does not have a comprehensive set of arguments to make changes. In fact, there has never been scriptable, in-box Windows utility designed specifically for administrators to create and modify AD DS topology. As Active Directory has matured in millions of customer environments, the need to bulk modify Active Directory logical information becomes apparent.

For example, after a rapid expansion of new branch offices, combined with the consolidation of others, you might have a hundred site changes to make based on physical locations, network changes, and new capacity requirements. Rather than using Dssites.msc and Adsiedit.msc to make changes, you can automate. This is especially compelling when you start with a spreadsheet of data provided by your network and facilities teams.

The **Get-Adreplication*** cmdlets return information about replication topology and are useful for pipelining into the **Set-Adreplication*** cmdlets in bulk. **Get** cmdlets do not change data, they only show data or to create Windows PowerShell session objects that can be pipelined to **Set-Adreplication*** cmdlets. The **New** and **Remove** cmdlets are useful for creating or removing Active Directory topology objects.

For example, you can create new sites using a CSV file:

```
import-csv -path C:\newsites.csv | new-adreplicationsite
```

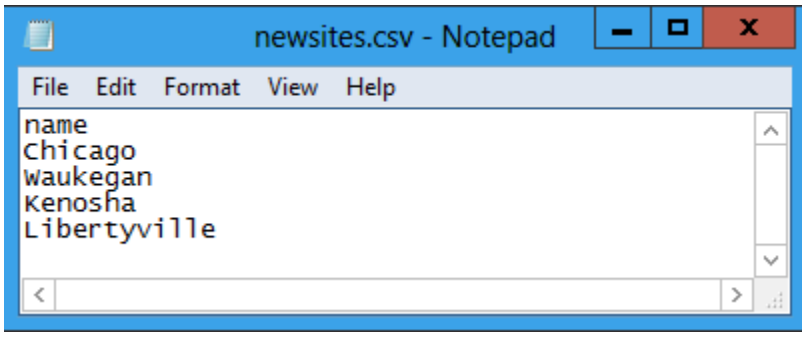

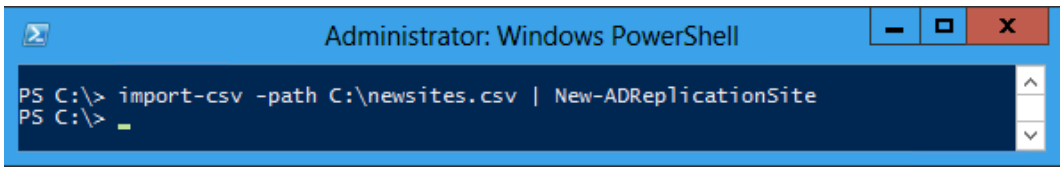

Alternatively, create a new site link between two existing sites with a custom replication interval and site cost:

```
new-adreplicationsitelink -name "chicago<-->waukegan" -sitesincluded chicago,waukegan -
cost 50 -replicationfrequencyinminutes 15
```

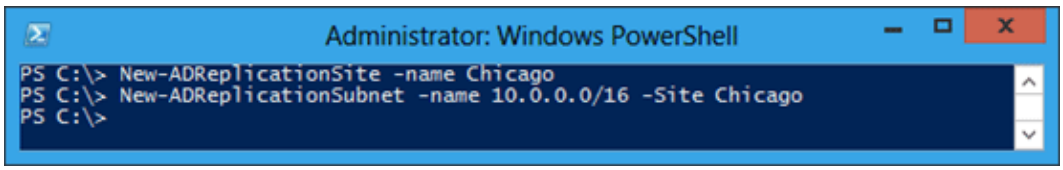

Alternatively, find every site in the forest and replace their **Options** attributes with the flag to enable inter-site change notification, in order to replicate at maximum speed with compression:

```
get-adreplicationsitelink -filter * | set-adobject -replace @{options=$($_.options -bor
1)}
```

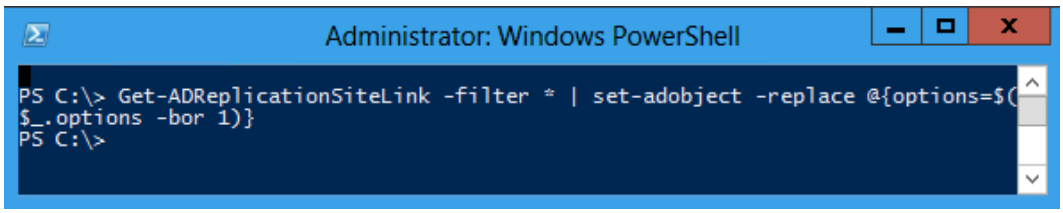

Important

Set **-bor 5** to disable compression on those site links as well.

Alternatively, find all sites missing subnet assignments, in order to reconcile the list with the actual subnets of those locations:

```
get-adreplicationsite -filter * -property subnets | where-object {!$_.subnets -eq "*"} |
format-table name
```

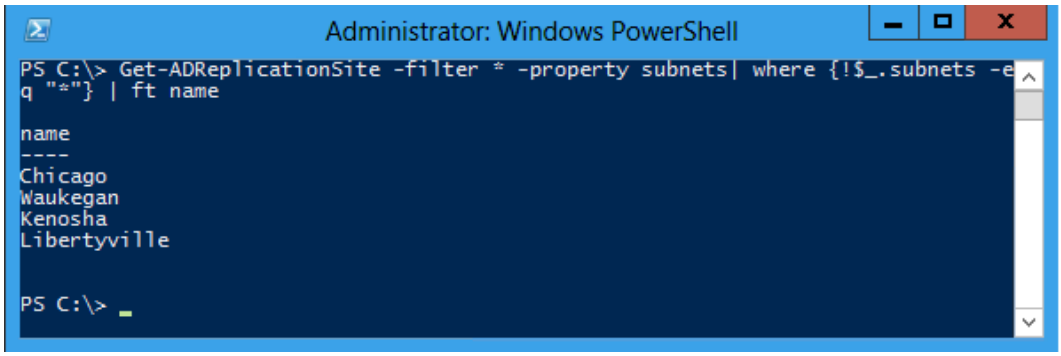

See Also

Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100)

Managing RID Issuance

This topic explains the change to the RID master FSMO role, including the new issuance and monitoring functionality in the RID master and how to analyze and troubleshoot RID issuance.

- Managing RID Issuance
- Troubleshooting RID Issuance

More information is available at the AskDS Blog.

Managing RID Issuance

By default, a domain has capacity for roughly one billion security principals, such as users, groups, and computers. Naturally, there are no domains with that many actively used objects. However, Microsoft Customer Support has found cases where:

- Provisioning software or administrative scripts accidentally bulk created users, groups, and computers.
- Many unused security and distribution groups were created by delegated users
- Many domain controllers were demoted, restored, or metadata cleaned
- Forest recoveries were performed
- The InvalidateRidPool operation was performed frequently
- The RID Block Size registry value was increased incorrectly

All of these situations use up RIDs unnecessarily, often by mistake. Over many years, a few environments ran out of RIDs and this forced them to migrate to a new domain or perform forest recoveries.

Windows Server 2012 addresses issues with RID allocation that have only become problematic with the age and ubiquity of Active Directory. These include better event logging, more appropriate limits, and the ability to - in an emergency – to double the overall size of the global RID space for a domain.

Periodic Consumption Warnings

Windows Server 2012 adds global RID space event tracking that provides early warning when major milestones are crossed. The model computes the ten (10) percent used mark in the global pool and logs an event when reached. Then it computes the next ten percent used of the remaining and the event cycle continues. As the global RID space is exhausted, events will accelerate as ten percent hits faster in a decreasing pool (but event log dampening will prevent more than one entry per hour). The System event log on every domain controller writes Directory-Services-SAM warning event 16658.

Assuming a default 30-bit global RID space, the first event logs when allocating the pool containing the 107,374,182nd RID. The event rate accelerates naturally until the last checkpoint of 100,000, with 110 events generated in total. The behavior is similar for an unlocked 31-bit global RID space: starting at 214,748,365 and completing in 117 events.

Important

This event is not expected; investigate the user, computer, and group creation processes immediately in the domain. Creating more than 100 million AD DS objects is quite out of the ordinary.

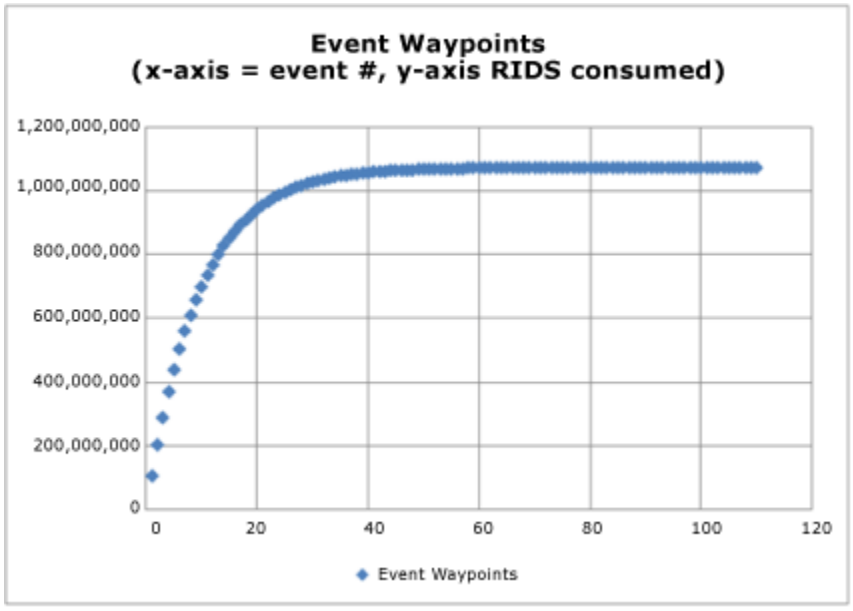

RID Pool Invalidation Events

There are new event alerts that a local DC RID pool was discarded. These are Informational and could be expected, especially due to the new VDC functionality. See the event list below for details on the event.

RID Block Size Limit

Ordinarily, a domain controller requests RID allocations in blocks of 500 RIDs at one time. You can override this default using the following registry REG_DWORD value on a domain controller:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\RID Values

RID Block Size
```

Prior to Windows Server 2012, there was no maximum value enforced in that registry key, except the implicit DWORD maximum (which has a value of 0xffffffff or 4294967295). This value is considerably larger than the total global RID space. Administrators sometimes inappropriately or accidentally configured RID Block Size with values that exhausted the global RID at a massive rate.

In Windows Server 2012, you cannot set this registry value higher than 15,000 decimal (0x3A98 hexadecimal). This prevents massive unintended RID allocation.

If you set the value *higher* than 15,000, the value is treated as 15,000 and the domain controller logs event 16653 in the Directory Services event log at every reboot until the value is corrected.

Global RID Space Size Unlock

Prior to Windows Server 2012, the global RID space was limited to 2^{30} (or 1,073,741,823) total RIDs. Once reached, only a domain migration or forest recovery to an older timeframe allowed new SIDs creation - disaster recovery, by any measure. Starting in Windows Server 2012, the 2^{31} bit can be unlocked in order to increase the global pool to 2,147,483,648 RIDs.

AD DS stores this setting in a special hidden attribute named **SidCompatibilityVersion** on the RootDSE context of all domain controllers. This attribute is not readable using ADSIEdit, LDP, or other tools. To see an increase in the global RID space, examine the System event log for warning event 16655 from Directory-Services-SAM or use the following Dcdiag command:

```
Dcdiag.exe /TEST:RidManager /v | find /i "Available RID Pool for the Domain"
```

If you increase the global RID pool, the available pool will change to 2,147,483,647 instead of the default 1,073,741,823. For example:

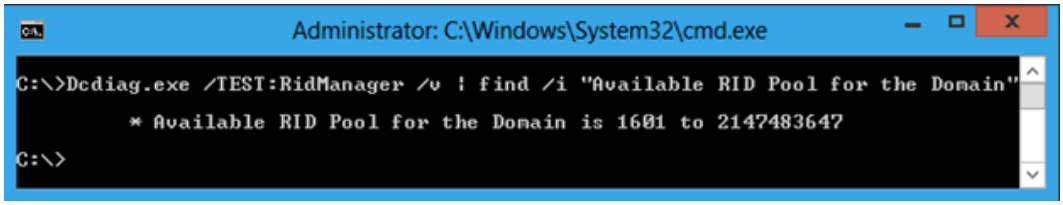

Warning

This unlock is intended *only* to prevent running out of RIDS and is to be used *only* in conjunction with RID Ceiling Enforcement (see next section). Do not "preemptively" set this in environments that have millions of remaining RIDs and low growth, as application compatibility issues potentially exist with SIDs generated from the unlocked RID pool.

This unlock operation cannot be reverted or removed, except by a complete forest recovery to earlier backups.

Important Caveats

Windows Server 2003 and Windows Server 2008 Domain Controllers cannot issue RIDs when the global RID pool 31st bit is unlocked. Windows Server 2008 R2 domain controllers *can* use 31st bit RIDs *but only if* they have hotfix KB 2642658 installed. Unsupported and unpatched domain controllers treat the global RID pool as exhausted when unlocked.

This feature is not enforced by any domain functional level; take great care that only Windows Server 2012 or updated Windows Server 2008 R2 domain controllers exist in the domain.

Implementing Unlocked Global RID space

To unlock the RID pool to the 31st bit after receiving the RID ceiling alert (see below) perform the following steps:

1. Ensure that the RID Master role is running on a Windows Server 2012 domain controller. If not, transfer it to a Windows Server 2012 domain controller.
2. Run LDP.exe
3. Click the **Connection** menu and click **Connect** for the Windows Server 2012 RID Master on port 389, and then click **Bind** as a domain administrator.
4. Click the **Browse** menu and click **Modify**.
5. Ensure that **DN** is blank.
6. In **Edit Entry Attribute**, type:

   ```
   SidCompatibilityVersion
   ```

7. In **Values**, type:

   ```
   1
   ```

8. Ensure that **Add** is selected in **Operation** and click **Enter**. This updates the **Entry List**.
9. Select the **Synchronous** and **Extended** options, then click **Run**.

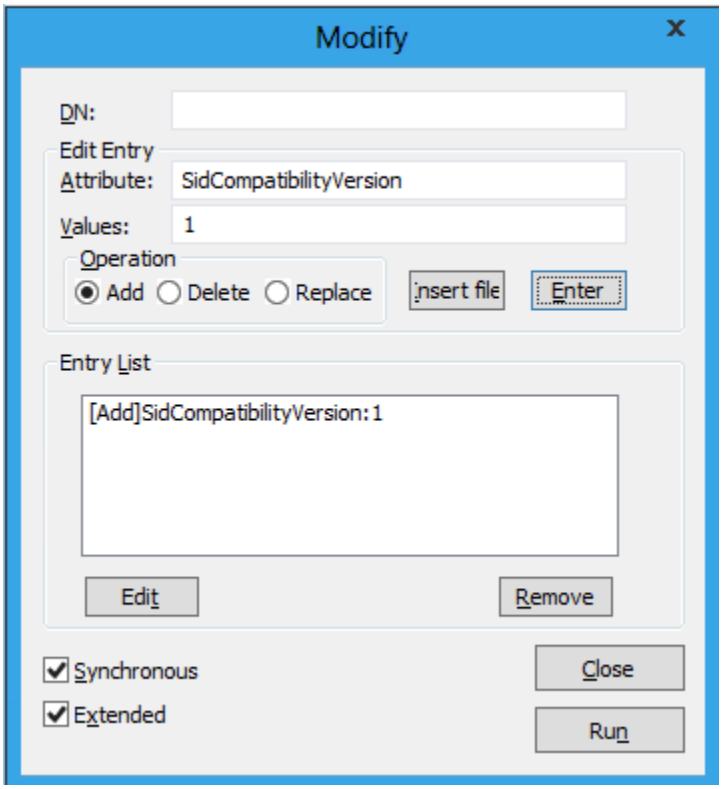

10. If successful, the LDP output window shows:

```
***Call Modify...
 ldap_modify_ext_s(Id, '(null)',[1] attrs, SvrCtrls, ClntCtrls);
modified "".
```

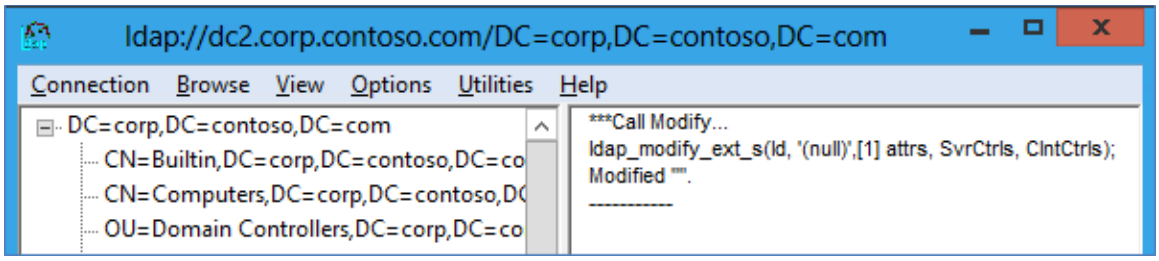

11. Confirm the global RID pool increased by examining the System Event Log on that domain controller for Directory-Services-SAM Informational event 16655.

RID Ceiling Enforcement

To afford a measure of protection and elevate administrative awareness, Windows Server 2012 introduces an artificial ceiling on the global RID range at ten (10) percent remaining RIDs in the global space. When within one (1) percent of the artificial ceiling, domain controllers requesting

RID pools write Directory-Services-SAM warning event 16656 to their System event log. When reaching the ten percent ceiling on the RID Master FSMO, it writes Directory-Services-SAM event 16657 to its System event log and will not allocate any further RID pools until overriding the ceiling. This forces you to assess the state of the RID master in the domain and address potential runaway RID allocation; this also protects domains from exhausting the entire RID space.

This ceiling is hard-coded at ten percent remaining of the available RID space. That is, the ceiling activates when the RID master allocates a pool that includes the RID corresponding to ninety (90) percent of the global RID space.

- For default domains, the first trigger point is 2^{30}-1 * 0.90 = 966,367,640 (or 107,374,183 RIDs remaining).
- For domains with an unlocked 31-bit RID space, the trigger point is 2^{31}-1 * 0.90 = 1,932,735,282 RIDs (or 214,748,365 RIDs remaining).

When triggered, the RID master sets Active Directory attribute **msDS-RIDPoolAllocationEnabled** (common name **ms-DS-RID-Pool-Allocation-Enabled**) to FALSE on the object:

CN=RID Manager$,CN=System,DC=<*domain*>

This writes the 16657 event and prevents further RID block issuance to all domain controllers. Domain controllers continue to consume any outstanding RID pools already issued to them.

To remove the block and allow RID pool allocation to continue, set that value to TRUE. On the next RID allocation performed by the RID master, the attribute will return to its default NOT SET value. After that, there are no further ceilings and eventually, the global RID space runs out, requiring forest recovery or domain migration.

Removing the Ceiling Block

To remove the block once reaching the artificial ceiling, perform the following steps:

1. Ensure that the RID Master role is running on a Windows Server 2012 domain controller. If not, transfer it to a Windows Server 2012 domain controller.
2. Run LDP.exe.
3. Click the **Connection** menu and click *Connect* for the Windows Server 2012 RID Master on port 389, and then click **Bind** as a domain administrator.
4. Click the **View** menu and click **Tree**, then for the **Base DN** select the RID Master's own domain naming context. Click **Ok**.
5. In the navigation pane, drill down into the **CN=System** container and click the **CN=RID Manager$** object. Right click it and click **Modify**.
6. In Edit Entry Attribute, type:

   ```
   MsDS-RidPoolAllocationEnabled
   ```

7. In **Values**, type (in upper case):

   ```
   TRUE
   ```

8. Select **Replace** in **Operation** and click **Enter**. This updates the **Entry List**.
9. Enable the **Synchronous** and **Extended** options, then click **Run**:

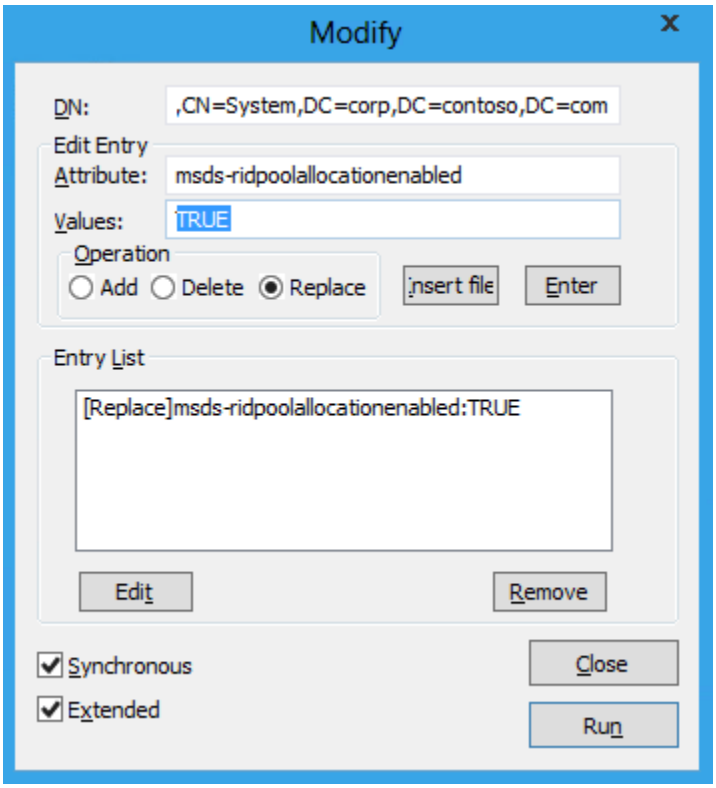

10. If successful, the LDP output window shows:

```
***Call Modify...
ldap_modify_ext_s(ld, 'CN=RID
Manager$,CN=System,DC=<domain>',[1] attrs, SvrCtrls, ClntCtrls);
Modified "CN=RID Manager$,CN=System,DC=<domain>".
```

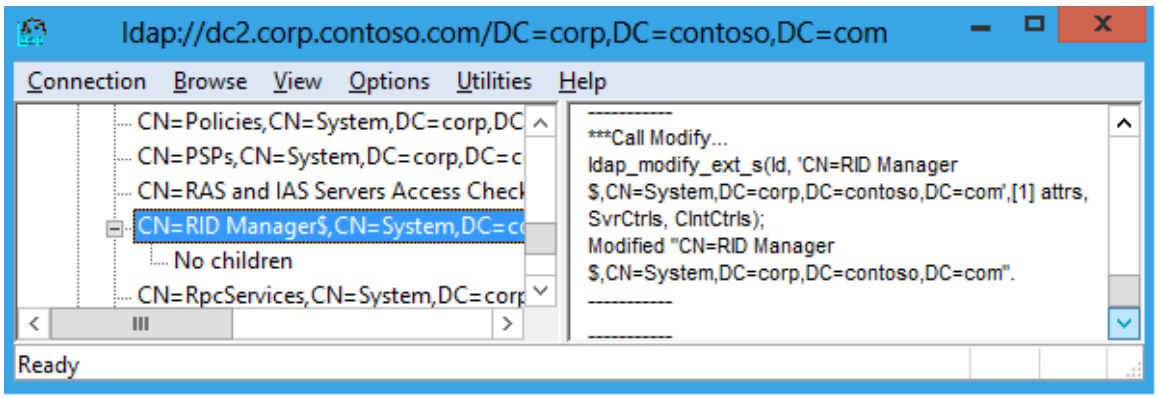

Other RID Fixes

Previous Windows Server operating systems had a RID pool leak when missing rIDSetReferences attribute. To resolve this problem on domain controllers that run Windows Server 2008 R2, install the hotfix from KB 2618669.

Unfixed RID Issues

There has historically been a RID leak on account creation failure; when creating an account, failure still uses up a RID. The common example is to create a user with a password that does not meet complexity.

RID Fixes for earlier versions of Windows Server

All of the fixes and changes above have Windows Server 2008 R2 hotfixes released. There are currently no Windows Server 2008 hotfixes planned or in progress.

Troubleshooting RID Issuance

Introduction to Troubleshooting

RID issuance troubleshooting requires a logical and linear method. Unless you are monitoring your event logs carefully for RID-triggered warnings and errors, your first indications of a problem are likely to be failed account creations. The key to troubleshooting RID issuance is to understand when the symptom is expected or not; many RID issuance issues may affect only one domain controller and have nothing to do with component improvements. This simple diagram below helps make those decisions more clear:

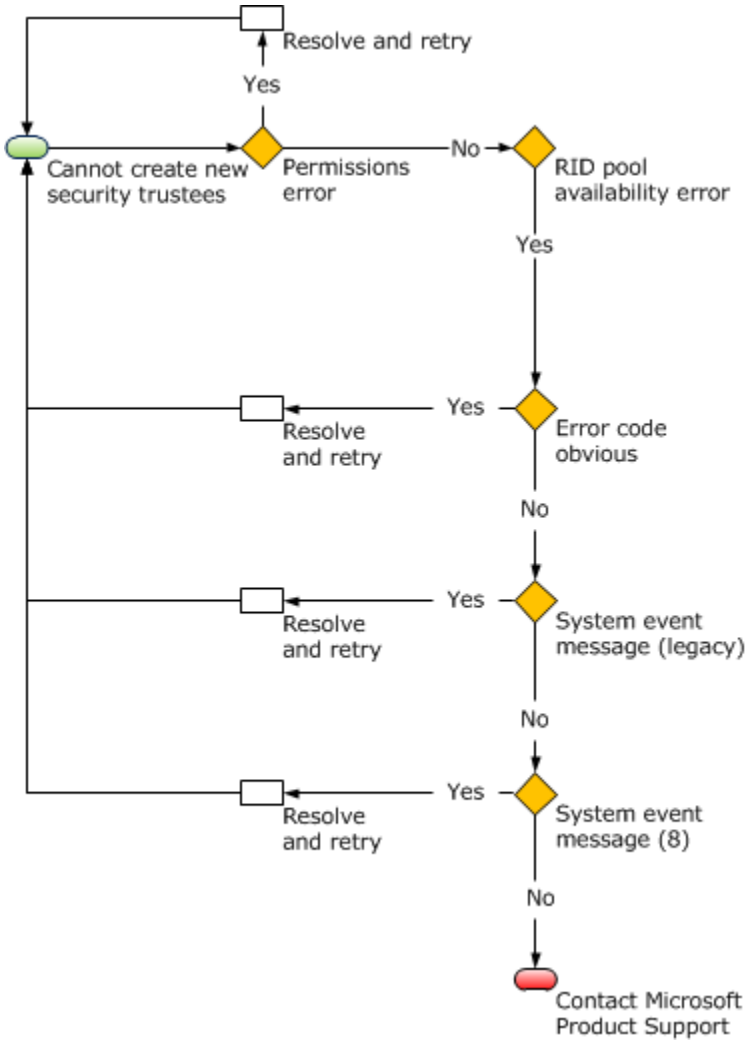

Troubleshooting Options

Logging Options

All logging in RID issuance occurs in the System Event log, under source Directory-Services-SAM. Logging is enabled and configured for maximum verbosity, by default. If no entries are logged for the new component changes in Windows Server 2012, treat the issue as a classic (aka legacy, pre-Windows Server 2012) RID issuance problem seen in Windows 2008 R2 or older operating systems.

Utilities and Commands for Troubleshooting

To troubleshoot issues not explained by the aforementioned logs - especially older RID issuance issues - use the following list of tools as a starting point:

- Dcdiag.exe
- Repadmin.exe
- Network Monitor 3.4

General Methodology for Troubleshooting Domain Controller Configuration

1. Is the error caused by a simple permissions or domain controller availability issue?
   a. Are you trying to create a security principal without the necessary permissions? Examine the output for access denied errors.
   b. Is a domain controller available? Examine the returned error or LDAP or domain controller availability messages.
2. Does the error returned specifically mention RIDs, and is specific enough to use as guidance? If so, follow the guidance.
3. Does the error returned specifically mention RIDs but is otherwise non-specific? For example, "Windows cannot create the object because the Directory Service was unable to allocate a relative identifier."
   a. Examine the System Event log on the domain controller for "legacy" (pre-Windows Server 2012) RID events detailed in RID Pool Request (16642, 16643, 16644, 16645, 16656).
   b. Examine the System Event on the domain controller and the RID Master for new block-indicating events detailed below in this topic (16655, 16656, 16657).
   c. Validate Active Directory replication health with Repadmin.exe and RID Master availability with **Dcdiag.exe /test:ridmanager /v**. Enable double-sided network captures between the domain controller and the RID Master if these tests are inconclusive.

Troubleshooting Specific Problems

The following new messages log in the System event log on Windows Server 2012 domain controllers. Automated AD health tracking systems, such as System Center Operations Manager, should monitor for these events; all are notable, and some are indicators of critical domain issues.

Event ID	16653
Source	Directory-Services-SAM
Severity	Warning
Message	A pool size for account-identifiers (RIDs) that was configured by an Administrator is greater than the supported maximum. The maximum value of %1 will be used when the domain controller is the RID master. For more information, see RID Block Size Limit.

Notes and resolution	The maximum value for the RID Block Size is now 15000 decimal (3A98 hexadecimal). A domain controller cannot request more than 15,000 RIDs. This event logs at every boot until the value is set to a value at or below this maximum.

Event ID	16654
Source	Directory-Services-SAM
Severity	Informational
Message	A pool of account-identifiers (RIDs) has been invalidated. This may occur in the following expected cases: 1. A domain controller is restored from backup. 2. A domain controller running on a virtual machine is restored from snapshot. 3. An administrator has manually invalidated the pool. See http://go.microsoft.com/fwlink/?LinkId=226247 for more information.
Notes and resolution	If this event is unexpected, contact all domain administrators and determine which of them performed the action. The Directory Services event log also contains further information on when one of these steps was performed.

Event ID	16655
Source	Directory-Services-SAM
Severity	Informational
Message	The global maximum for account-identifiers (RIDs) has been increased to %1.
Notes and resolution	If this event is unexpected, contact all domain administrators and determine which of them performed the action. This event notes the increase of the overall RID pool size beyond

	the default of 2^{30} and will not happen automatically; only by administrative action.

Event ID	16656
Source	Directory-Services-SAM
Severity	Warning
Message	The global maximum for account-identifiers (RIDs) has been increased to %1.
Notes and resolution	Action required! An account-identifier (RID) pool was allocated to this domain controller. The pool value indicates this domain has consumed a considerable portion of the total available account-identifiers. A protection mechanism will be activated when the domain reaches the following threshold of total available account-identifiers remaining: %1. The protection mechanism will prevent account creation until you manually re-enable account-identifier allocation on the RID master domain controller. See http://go.microsoft.com/fwlink/?LinkId=228610 for more information.

Event ID	16657
Source	Directory-Services-SAM
Severity	Error
Message	Action required! This domain has consumed a considerable portion of the total available account-identifiers (RIDs). A protection mechanism has been activated because the total available account-identifiers remaining is less than: X% [artificial ceiling argument]. The protection mechanism prevents account creation until you manually re-enable account-identifier allocation on the RID master domain controller.

	It is extremely important that certain diagnostics are performed prior to re-enabling account creation to ensure this domain is not consuming account-identifiers at an abnormally high rate. Any issues identified should be resolved prior to re-enabling account creation. Failure to diagnose and fix any underlying issue causing an abnormally high rate of account-identifier consumption can lead to account-identifier exhaustion in the domain after which account creation will be permanently disabled in this domain. See http://go.microsoft.com/fwlink/?LinkId=228610 for more information.
Notes and resolution	Contact all domain administrators and inform them that no further security principals can be created in this domain until this protection is overridden. For more information about how to override the protection and possibly increase the overall RID pool, see Global RID Space Size Unlock.

Event ID	16658
Source	Directory-Services-SAM
Severity	Warning
Message	This event is a periodic update on the remaining total quantity of available account-identifiers (RIDs). The number of remaining account-identifiers is approximately: %1. Account-identifiers are used as accounts are created, when they are exhausted no new accounts may be created in the domain. See http://go.microsoft.com/fwlink/?LinkId=228745 for more information.
Notes and resolution	Contact all domain administrators and inform them that RID consumption has crossed a major milestone; determine if this is expected

	behavior or not by reviewing security trustee creation patterns. To ever see this event would be highly unusual, as it means that at least ~100 million RIDS have been allocated.

See Also

Managing RID Issuance in Windows Server 2012

Dynamic Access Control: Scenario Overview

In Windows Server 2012, you can apply data governance across your file servers to control who can access information and to audit who has accessed information. Dynamic Access Control lets you:

- Identify data by using automatic and manual classification of files. For example, you could tag data in file servers across the organization.
- Control access to files by applying safety-net policies that use central access policies. For example, you could define who can access health information within the organization.
- Audit access to files by using central audit policies for compliance reporting and forensic analysis. For example, you could identify who accessed highly sensitive information.
- Apply Rights Management Services (RMS) protection by using automatic RMS encryption for sensitive Microsoft Office documents. For example, you could configure RMS to encrypt all documents that contain Health Insurance Portability and Accountability Act (HIPAA) information.

The Dynamic Access Control feature set is based on infrastructure investments that can be used further by partners and line-of-business applications, and the features can provide great value for organizations that use Active Directory. This infrastructure includes:

- A new authorization and audit engine for Windows that can process conditional expressions and central policies.
- Kerberos authentication support for user claims and device claims.
- Improvements to the File Classification Infrastructure (FCI).
- RMS extensibility support so partners can provide solutions that encrypt non-Microsoft files.

In this scenario

The following scenarios and guidance are included as part of this content set:

Dynamic Access Control Content Roadmap

Scenario	Evaluate	Plan	Deploy	Operate
Scenario: Central Access Policy Creating Central access policies for files allow organizations to centrally deploy and manage authorization policies that include conditional expressions using user claims, device claims, and resource properties. These polices are based on compliance and business regulatory requirements. These policies are created and hosted in Active Directory, therefore making it easier to manage and deploy. **Deploying Claims Across Forests** In Windows Server 2012, the AD DS maintains a 'claims dictionary' in each forest and all claim types in use within the forest are defined at the Active Directory forest level. There are many scenarios where a principal may need to	Dynamic Access Control: Scenario Overview Deploy Claims Across Forests	Plan for a Central Access Policy Deployment • Process to map a business request to a central access policy • Delegating of administration for Dynamic Access Control • Exception Mechanisms for Planning Central Access Policies Best Practices for Using User Claims • Choosing the right configuration to enable claims in your user domain • Operations to enable user claims • Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies Using Device Claims and Device Security Groups • Considerations for using static device claims	Deploy a Central Access Policy (Demonstration Steps) Deploy Claims Across Forests (Demonstration Steps)	• Modeling a central access policy

Scenario	Evaluate	Plan	Deploy	Operate
traverse a trust boundary. This scenario describes how a claim traverses a trust boundary.		• Operations to enable device claims Tools for Deployment • Data Classification Toolkit		
Scenario: File Access Auditing Security auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is regulatory compliance. For example, industry standards such as Sarbanes Oxley, HIPAA, and Payment Card Industry (PCI) require enterprises to follow a strict set of rules related to data security and privacy. Security audits help establish the presence or absence of such policies; thereby, they prove compliance or noncompliance with these standards. Additionally,	Scenario: File Access Auditing	Plan for File Access Auditing	Deploy Security Auditing with Central Audit Policies (Demonstration Steps)	• Monitor the Central Access Policies that Apply on a File Server • Monitor the Central Access Policies Associated with Files and Folders • Monitor the Resource Attributes on Files and Folders • Monitor Claim Types • Monitor User and Device Claims During Sign-in • Monitor Central

Scenario	Evaluate	Plan	Deploy	Operate
security audits help detect anomalous behavior, identify and mitigate gaps in security policy, and deter irresponsible behavior by creating a record of user activity that can be used for forensic analysis.				• Access Policy and Rule Definitions • Monitor Resource Attribute Definitions • Monitor the Use of Removable Storage Devices.
Scenario: Access-Denied Assistance Today, when users try to access a remote file on the file server, the only indication that they would get is that access is denied. This generates requests to helpdesk or IT administrators that need to figure out what the issue is and often the administrators have a hard time getting the appropriate context from users which makes it harder to resolve the issue. In Windows Server 2012, the goal is to try and help the information worker and business owner	Scenario: Access-Denied Assistance	Plan for Access-Denied Assistance • Determine the access-denied assistance model • Determine who should handle access requests • Customize the access-denied assistance message • Plan for exceptions • Determine how access-denied assistance is deployed	Deploy Access-Denied Assistance (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
of the data to deal with the access denied issue before IT gets involved and when IT gets involved, provide all the right information for a quick resolution. One of the challenges in achieving this goal is that there is no central way to deal with access denied and every application deals with it differently and thus in Windows Server 2012, one of the goals is to improve the access-denied experience for Windows Explorer.				
Scenario: Classification-Based Encryption for Office Documents Protection of sensitive information is mainly about mitigating risk for the organization. Various compliance regulations, such as HIPAA or Payment Card Industry Data Security Standard (PCI-DSS), dictate	Scenario: Classification-Based Encryption for Office Documents	Planning Considerations for Encryption of Office Documents	Deploy Encryption of Office Files (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
encryption of information, and there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it might impair business productivity. Thus, organizations tend to have different approaches and priorities for encrypting their information. To support this scenario, Windows Server 2012 provides the ability to automatically encrypt sensitive Windows Office files based on their classification. This is done through file management tasks that invoke Active Directory Rights Management Server (AD RMS) protection for sensitive documents a few seconds after the file is identified as being a sensitive file on the file				

Scenario	Evaluate	Plan	Deploy	Operate
server.				
Scenario: Get Insight into Your Data by Using Classification Reliance on data and storage resources has continued to grow in importance for most organizations. IT administrators face the growing challenge of overseeing larger and more complex storage infrastructures while simultaneously being tasked with the responsibility to ensure total cost of ownership is maintained at reasonable levels. Managing storage resources is not just about the volume or availability of data anymore, but also about the enforcement of company policies and knowing how storage is consumed to enable efficient utilization and compliance to mitigate risk. File Classification	Scenario: Get Insight into Your Data by Using Classification	Plan for Automatic File Classification	Deploy Automatic File Classification (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
Infrastructure provides insight into your data by automating classification processes so that you can manage your data more effectively. The following classification methods are available with File Classification Infrastructure: manual, programmatically, and automatic. This scenario focuses on the automatic file classification method.				
Scenario: Implement Retention of Information on File Servers A retention period is the amount of time that a document should be kept before it is expired. Depending on the organization, the retention period can be different. You can classify files in a folder as having a short, medium, or long-term retention period and then	Scenario: Implement Retention of Information on File Servers	Plan for Retention of Information on File Servers	Deploy Implementing Retention of Information on File Servers (Demonstration Steps)	

Scenario	Evaluate	Plan	Deploy	Operate
assign the timeframe for each period. You may want to keep a file indefinitely by putting it on legal hold. File Classification Infrastructure and File Server Resource Manager uses file management tasks and file classification to apply retention periods for a set of files. You can assign a retention period on a folder and then use a file management task to configure how long an assigned retention period is to last. When the files in the folder are about to expire, the owner of the file gets a notification email. You can also classify a file as being on legal hold so that the file management task will not expire the file.				

Note
Dynamic Access Control is not supported on ReFS (Resilient File System).

See also

Content type	References
Product evaluation	• Dynamic Access Control Reviewers Guide • Dynamic Access Control Developer Guidance
Planning	• Plan for a Central Access Policy Deployment • Plan for File Access Auditing
Deployment	• Active Directory Deployment • File and Storage Services Deployment
Operations	Dynamic Access Control PowerShell Reference
Tools and settings	Data Classification Toolkit
Community resources	Directory Services Forum

Scenario: Central Access Policy

Central access policies for files enable organizations to centrally deploy and manage authorization policies that include conditional expressions that use user groups, user claims, device claims, and resource properties. (Claims are assertions about the attributes of the object with which they are associated). For example, to access high-business-impact (HBI) data, a user must be a full-time employee, obtain access from a managed device, and log on with a smart card. These policies are defined and hosted in Active Directory Domain Services (AD DS).

Organizational access policies are driven by compliance and business regulatory requirements. For example, if an organization has a business requirement to restrict access to personally identifiable information (PII) in files to only the file owner and members of the human resources (HR) department who are allowed to view PII information, this policy applies to PII files wherever they are located on file servers across the organization. In this example, you need to be able to:

- Identify and mark the files that contain PII.
- Identify the group of HR members who are allowed to view PII information.
- Create a central access policy that applies to all files that contain PII wherever they are located on file servers across the organization.

The initiative to deploy and enforce an authorization policy can come for many reasons and apply to multiple levels of the organization. The following are some example policy types:

- **Organization-wide authorization policy.** Most commonly initiated from the information security office, this authorization policy is driven by compliance or a high-level organization requirements, and it is relevant across the organization. For example, HBI files are accessible to only full-time employees.
- **Departmental authorization policy.** Each department in an organization has some special data-handling requirements that they want to enforce. For example, the finance department might want to limit access to finance servers to the finance employees.
- **Specific data-management policy.** This policy usually relates to compliance and business requirements, and it is targeted at protecting the correct access to the information that is being managed. For example, financial institutions might implement information walls so that analysts do not access brokerage information and brokers do not access analysis information.
- **Need-to-know policy.** This authorization policy type is typically used in conjunction with the previous policy types. For example, vendors should be able to access and edit only files that pertain to a project they are working on.

Real-life environments also teach us that every authorization policy needs to have exceptions so that organizations can quickly react when important business needs arise. For example, executives who cannot find their smart cards and need quick access to HBI information can call the Help Desk to get a temporary exception to access that information.

Central access policies act as security umbrellas that an organization applies across its servers. These policies enhance (but do not replace) the local access policies or discretionary access control lists (DACL) that are applied to files and folders. For example, if a DACL on a file allows access to a specific user, but a central policy that is applied to the file restricts access to the same user, the user cannot obtain access to the file. If the central access policy allows access, but the DACL does not allow access, the user cannot obtain access to the file.

A central access policy rule has the following logical parts:

- **Applicability.** A condition that defines which data the policy applies to, such as Resource.BusinessImpact=High.
- **Access conditions.** A list of one or more access control entries (ACEs) that define who can access the data, such as Allow | Full Control | User.EmployeeType=FTE.
- **Exceptions.** An additional list of one or more ACEs that define an exception for the policy, such as MemberOf(HBIExceptionGroup).

The following two figures show the workflow in central access and audit policies.

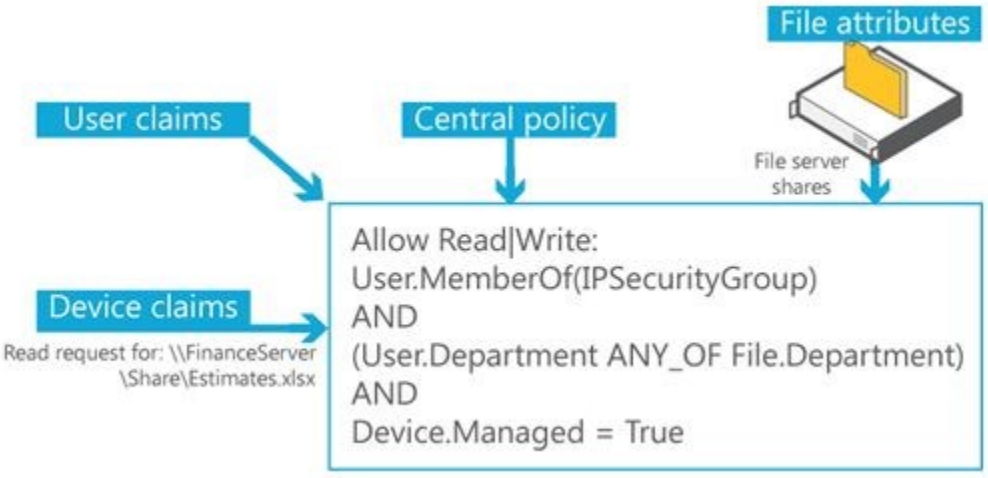

Figure 1 Central access and audit policy concepts

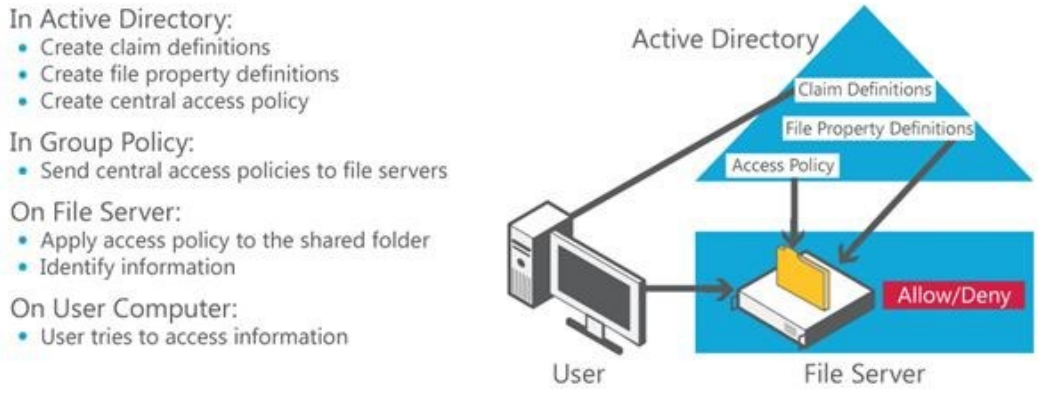

Figure 2 Central access policy workflow

The central authorization policy combines the following components:

- A list of centrally defined access rules that target specific types of information, such as HBI or PII.
- A centrally defined policy that contains a list of rules.
- A policy identifier that is assigned to each file on the file servers to point to a specific central access policy that should be applied during the access authorization.

The following figure demonstrates how you can combine policies into policy lists to centrally control access to files.

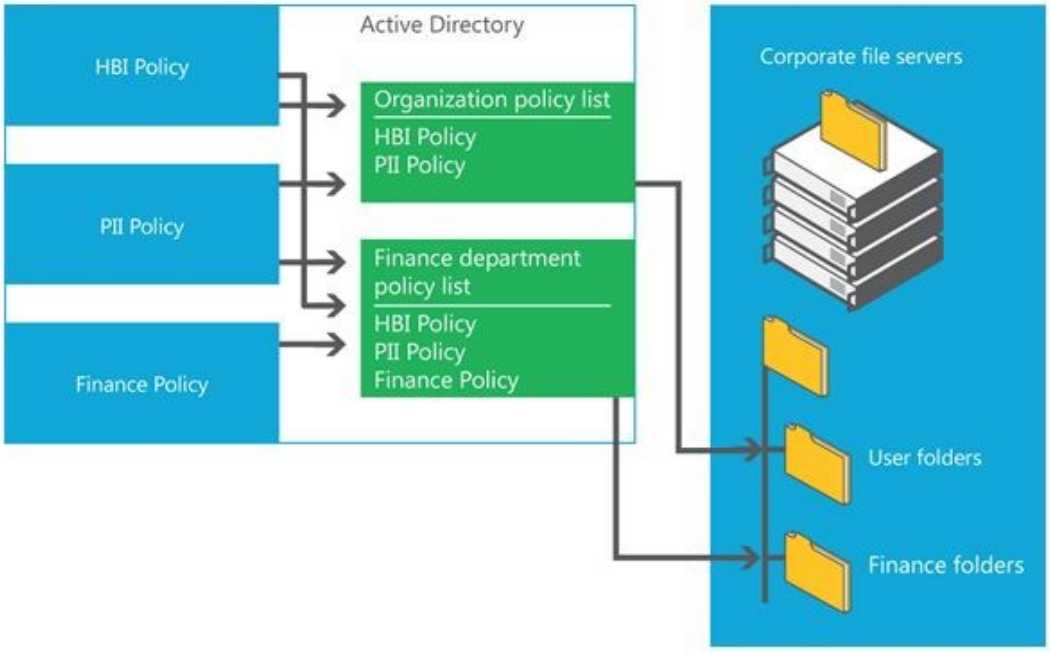

Figure 3 Combining policies

In this scenario

The following guidance is available to you for central access policies:

- Plan for a Central Access Policy Deployment
- Deploy a Central Access Policy (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services role	AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables the creation of user claims and device claims, compound identity, (user plus device claims), new central access policy (CAP) models, and the use of file-classification information in authorization decisions.

Role/feature	How it supports this scenario
File and Storage Services Server role	File and Storage Services provides technologies that help you set up and manage one or more file servers that provide central locations on your network where you can store files and share them with users. If your network users need access to the same files and applications, or if centralized backup and file management are important to your organization, you should set up one or more computers as a file server by adding the File and Storage Services role and the appropriate role services to the computers.
Windows client computer	Users can access files and folders on the network through the client computer.

Plan for a Central Access Policy Deployment

The need to control the information in enterprise-level organizations for compliance and business regulations is one of the drivers in the consolidation trend where large amounts of information from users' desktops and departmental file shares are moved into centrally managed file servers. The initiative to deploy and enforce an authorization policy may come for different reasons and from multiple levels of the organization:

- **Organization-wide authorization policy**: most commonly initiated from the Information Security office, this authorization policy is driven from compliance or very high level organization requirement and would be relevant across the organization. For example: High Business Impact files should be accessible by full time employees only
- **Departmental authorization policy**: each department in an organization has some special data handling requirements that they would like to enforce. This is very common in distributed organization. For example: the finance department might want to limit all access to finance servers to the finance employees
- **Specific data management policy**: this policy usually relates to compliance and business requirements and is targeted at protecting the right access to information that is being managed. For example: Preventing modification or deletion of files that are under retention or files that are under eDiscovery
- **Need to know policy**: This is a catch all authorization policy type and most probably used in conjunction with the policy types mentioned above. Examples include: Vendors should be able to access and edit only files that pertain to a project that they are working on.
  In financial institutions, information walls are important so that analysts do not access brokerage information and brokers do not access analysis information

Process to map a business request to a central access policy

1. Understand and translate business intent
2. Express access policy in Windows Server 2012 constructs
3. Determine the user groups, resource properties and claim types
4. Determine the servers where this policy should be applied to

Understand and translate business intent

Business determines that a central access policy is needed. The next step is to decide on the content (resources) that you want to apply policies. Then create a list of all the policies you want to apply to your content. For example some of the common policies that apply to the finance department of an organization would be:

- Archived finance documents should only be read by members of the Finance department.
- Members of the Finance department should only be able to access documents in their own country.
- Only Finance Administrators should have write access. An exception will be allowed for members of the FinanceException group. This group will have read access.

Express access policy in Windows Server 2012 constructs

The next step in the planning process is to translate the policies you require into expressions. A central access policy is targeted at providing an easy interpretation from a business requirement language to an authorization language.
A central policy in Windows Server 2012has the following distinct parts:

- Applicability: A condition that defines which data the policy applies to
  (Example: Resource.BusinessImpact=High)
- Access conditions: A list of one or more Access Control Entries (ACE) that defines who can access the data
  (Example: Allow | Full Control | User.EmployeeType=FTE)
- Exception: An additional list of one or more access control entries that define an exception for the policy
  (Example: MemberOf(HBIExceptionGroup)

Determine the user groups, resource properties and claim types

In this step, the expressions that were created for the policies are broken down and analyzed to figure out what resource properties, security groups as well as potential user claims need to be created to deploy the specific policies. Note that before using user claims, you should review the relevant section on user claims in this topic. Remember that you can always use security groups in case that you do not have the appropriate user claim attribute available in your environment.

Determine the servers where this policy should be applied to

The next step is to determine what file servers you want to deploy the access policies you have decided on. For example you can have a finance access policy that you want to roll out only to the finance file servers.

Planning Guidelines for Deploying Central Access Policies

There are multiple ways to deploy central access policies, based on different configurations in your environment. You can choose to have a simple deployment with security groups or have more advanced central access policies using user claims and device claims. The following sections in this topic discuss the different deployment options available to you base on the configuration you choose. The design and deployment guidance for each of these is discussed in further detail:

- Using Security Groups for Dynamic Access Control
- Using User Claims
- Using Device Claims and Device Security Groups

For a list of all configuration options and guidance on which configuration to choose, see the section, Configuring central access policies with different options

For a high-level view of the different deployment options, requirements and configurations, see the Appendix: Deployment Configurations for Central Access Policies

Using Security Groups for Dynamic Access Control

You don't have to use user claims in order to implement Dynamic Access Control in your organization. In fact, you can use security groups with very minimal upgrade requirements to your current environment. You can use security groups in conjunction with central access policies and conditional expressions in Windows Server 2012. This way you can use Dynamic Access Control to limit access to specific data using the existing security group mechanism.

To use Dynamic Access Control with security groups you need the following in your environment:

- A Windows Server 2012 File Server
- A domain with a Windows Server 2012 schema (so that you can define central access policies)

Using security groups to limit access to data

Dynamic Access Control can help you limit access to data to the specific groups of people. The steps to achieve this are:

1. Tagging the data by marking the folders that contain confidential data

2. Configuring a Central Access Rule that specifies that only specific security groups can access data that are tagged in a specific way
3. Applying a Central Access Policy to the appropriate Windows Server 2012 File Servers in your organization

For a detailed walkthrough of using security groups, see the Dynamic Access Control Blog.

Using conditional expressions to reduce complexity of security groups

You can use Dynamic Access Control to considerably reduce the complexity of a combinatorial number of security groups (for example, from 2,000 to less than a 100) so that you can have a clear understanding of who can access what data and you are able to easily adjust access when people move between different roles in the company.

In a typical enterprise environment there are a large number of users, departments, security groups, and thus, a large number of access control lists (ACLs). Also, if a person moves from one department to the other, that involves updating a large number of security groups. This results in a lot of IT overhead and Dynamic Access Control can help with reducing the workload of IT admins for managing security group. The steps required for this are:

1. Tag all the folders with the appropriate values, for example, department, country, and sensitive.
2. Decide on the combinations of expressions to be used in in Windows ACL. For example, you would use **MemberOf (Spain_Security_Group) AND MemberOf (Finance_Security_Group) AND MemberOf(Sensitive_Security_Group)** to limit access to Spain's finance department sensitive information.
3. Create specific central access rules with these expressions that target certain security groups and specific folders on the files servers.

For a detailed walkthrough of using Dynamic Access Control to reduce the complexity of security groups, see the Dynamic Access Control Blog.

Using User Claims

As a rule of thumb, you should use user claims (vs. security group) when:

- You want to be able to use conditions such as User.Project = File.Project in your policy so that you can compress thousands of conditions to a simple expression (avoiding conditions such as File.Project=Cosmos AND User.Memberof(Cosmos_security_group)).
- The user attribute in Active Directory that you are sourcing the user claim from has the appropriate security setting on who or what can set that attribute.
- High integrity of the attribute value in Active Directory and the system that sets this value has operational procedures that take into consideration the use of that value for authorization decisions.
- No foreseeable changes to the values in the attribute. For example if the attribute is a department name and these often change due to re-organization, then it is not fit to be used as a user claim.

When using user claims, you should make sure that you have the appropriate Windows Server 2012 domain controller environment to support these user claims. The two main domains of interest are the User domain and the File Server domain.

- Both domains cannot include Windows Server 2003 (or earlier version) domain controllers.
- Have adequate number of Windows Server 2012 domain controllers in the user domain to support the number of Windows 8 clients deployed in that domain or alternatively set the Windows 8 clients to not require claims in which case having at least one Windows Server 2012 domain controller in the user domain will be sufficient (see more details below
- If the User domain and File Server domain are in a different forest, you need to:
  - Have two-way trust relations between the two forests.
  - Have all forest root domain controllers for both domains be Windows Server 2012

Operations to enable user claims

- Enable the domain controllers to provide claims and compound authentication on request

Enable the domain controllers to provide claims and compound authentication on request

To enable the domain controllers to provide claims and compound authentication on request

1. Open Group Policy Management and navigate to **Domain Controllers** OU in the domain.
2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **KDC**.
4. Select **KDC support for claims, compound authentication, and Kerberos armoring**.
5. Under **Claims, compound authentication for Dynamic Access Control and Kerberos armoring** options select **Supported**.
6. Click **OK**. Close Group Policy Management.

Important

Until the Kerberos Group Policy **Kerberos client support for claims, compound authentication and Kerberos armoring** is enabled on all Windows 8 devices, claims and compound authentication will not be requested and claims-based and device-based access will fail.

Warning

Windows 8 devices enabled to support requesting claims and compound authentication will fail authentication when a Windows Server 2012 DC cannot be found in domains configured to support claims and compound authentication.

Considerations for using user claims in the file server discretionary ACLs without using Central Access Policies

Windows Server 2012 File servers have a group policy setting (On/Off/Automatic) that specifies whether it needs to get user claims for user tokens that do not carry claims. This setting is by default set to “automatic” which results in this group policy setting to be turned **On** if there is a central policy that contains user and/or device claims for that file server. If the file server contains local access policy (Discretionary ACLs) that uses user claims, you need to set this group policy to “On” so that the server knows to request claims on behalf of users that do not provide claims when accessing the server (For example, non-Windows 8 clients.

Using Device Claims and Device Security Groups

Windows 8 introduces the concept of compound identity that allows for the user token to include both the user details (ID, security groups and claims) as well as the device details (ID, security groups and claims) so that you can use these details in the authorization decision such as: Only specific users using a specific device can access highly sensitive data.

Much like user claims, device claims are sourced from the device object attribute in Active Directory (claim source) that contains the value of the claim.

Important
Device claims are supported in Windows 8clients only.

Considerations for using static device claims

To configure static device claims the administrator needs to:

- Apply the appropriate security on the device attribute
- Enable the claim definition – This operation will result in the claim to be available in the Kerberos ticket with the value sourced from the device attribute.
- Enable domain controller support for Dynamic Access Control – This operation will result in the directory service returning claims to the KDC and the KDC creating TGTs containing claims
- Apply the appropriate mechanisms to populate the device attribute in Active Directory.
- Set Windows 8 clients and Windows Server 2012 to request compound tokens.

Operations to enable device claims

- Enable the Windows 8 devices in domain to request claims and compound authentication
- Enable the Windows 8 devices to request claims and compound authentication using custom policy
- Enable the Windows 8 device to receive compound authentication

Enable the Windows 8 devices in domain to request claims and compound authentication

To enable the Windows 8 devices to request claims and compound authentication

1.

Note

Apply the group policy on the domain.

Open Group Policy Management and navigate to **Domain**.

2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **Kerberos**.
4. Select **Enable KDC support for claims, compound authentication, and Kerberos armoring**.
5. Click **OK**. Close Group Policy Management.

Enable the Windows 8 devices to request claims and compound authentication using custom policy

To enable Windows 8 devices to request claims and compound authentication using custom policy

1.

Note

Apply the group policy on the domain, site, OU or object scope.

Open Group Policy Management and navigate the appropriate policy.

2. Right-click the **Default Domain Controllers Policy** and select **Edit**.
3. In the Group Policy Management Editor window, navigate to **Computer Configuration**, > **Administrative Templates**, **System**, and **Kerberos**.
4. Select **Enable KDC support for claims, compound authentication, and Kerberos armoring**.
5. Click **OK**. Close Group Policy Management.

Enable the Windows 8 device to receive compound authentication

This group policy setting must be applied on the local group policy.

To enable the Windows 8 device to receive compound authentication

1. Open Local Group Policy Editor.
2. In the Local Group Policy window, expand to **Computer Configuration**, expand **Administrative Templates**, expand **System** and click **Kerberos**. Select **Enable** and **Support compound authentication**.
3. Click **OK** and close the Group Policy Management Editor.

Configuring central access policies with different options

As mentioned before, if you are using security groups, you do not have any specific domain consideration other than the need to upgrade your schema to the Windows Server 2012 schema.

If you decide to use user claims or you would like to use user/device compound identity, your deployment design for Windows Server 2012 can be any of the following three configurations:

- **Configuration 1**: Domains providing claims and compound authentication have all Windows Server 2012domain controllers.
- **Configuration 2**: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 DCs in all the file server sites.
- **Configuration 3**: Device-based access control needed, but cannot wait until all domain controllers can be upgraded.

Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs

This configuration provides complete infrastructure for claims and compound authentication for Dynamic Access Control and is the easiest to support. It requires that:

- If a cross-forest trusts exist, then root domain have all Windows Server 2012 domain controllers.
- Domains that provide claims and compound authentication have all Windows Server 2012 domain controllers
- All Windows 8 devices must be enabled to support requesting claims and compound authentication

Configuring forest root DCs

Simply upgrade all the domain controllers in the forest root to Windows Server 2012 and configure egress and ingress claims transformation.

Note

This will help ensure that claims are not lost from trusted forests. If pre-Windows Server 2012 domain controllers exist, those domain controllers will discard claims which will result in intermittent access control failures. Additionally pre-Windows Server 2012 domain controllers do not transform claims so the claims data is disclosed to all trusting forests

Configuring domains which provide claims and compound authentication

First identify which domains will be providing claims and compound authentication. These will be the account and resource domains. If all the domains in the environment have Windows Server

2012DCs then configure all.If claims-based access is required then claims will need to be provisioned.

Configure each domain which provides compound authentication and claims to DFL and Enable the DCs to support compound authentication and always provide claims.

Configuring devices to request claims and compound authentication

Update the default domain policy to Enable the Windows 8 devices in domain to request claims and compound authentication or Enable the Windows 8 devices to request claims and compound authentication using custom policy which applies to all the Windows 8 computer objects.

Note

This configuration that you set up to enable Windows 8 devices to request claims and compound authentication is ignored by versions of Windows which do not support it.

Configuring resources to receive compound authentication

If the resources are using CAP, then simply apply the CAP. Otherwise, Enable the Windows 8 device to receive compound authentication.

Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites

This configuration provides limited infrastructure for only user claims for Dynamic Access Control. It requires:

- If a cross-forest trusts exist, then root domain has all Windows Server 2012 DCs
- For each domain which provides user claims on request has Windows Server 2012 DCs in the sites with file servers and no Windows Server 2003 domain controllers.
- Enable all Windows 8 file servers to support requesting claims on the behalf of users.

Important

Authentication mechanism assurance-based universal groups and device-based access control are incompatible with file server retrieval of user claims.

Configuring forest root DCs

Simply upgrade all the domain controllers in the forest root to Windows Server 2012 and configure egress and ingress claims transformation.

Note

This will help ensure that claims are not lost from trusted forests. If pre-Windows Server 2012 domain controllers exist, those domain controllers will discard claims which will result in intermittent access control failures. Additionally pre-Windows Server 2012

domain controllers do not transform claims so the claims data is disclosed to all trusting forests

Configuring domains which provides claims and compound authentication

1. First identify which domains will be providing claims and compound authentication. These will be the account and resource domains.
2. Configure and provision claims as explained in Deploy a Central Access Policy (Demonstration Steps)
3. Configure each domain to Enable the domain controllers to provide claims and compound authentication on request.

Configuring file servers to request claims on the behalf of users

If the resources are using a central access policy, then simply apply the policy. Otherwise, Enable the Windows 8 devices to request claims and compound authentication using custom policy which applies to Windows 8 file server computer objects.

Note
This setting is ignored by versions of Windows which do not support it.

Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded

This configuration will be unique to your environment and can be difficult to support when Windows 8 devices have different configurations.

General requirements for all environments:

- If across-forest trusts exist, then root domain must have all Windows Server 2012domain controllers
- For each domain which provides claims and compound authentication on request, there cannot be Windows Server 2003 domain controllers
- For resources using device-based access control, receiving compound authentication must be enabled unless a central access policy is being used.

Considerations for using smartcards for Central Access Polices

In Windows 7 and Windows 8 clients, smart card logon is mapped through security groups. In order to create this mapping, the domain administrator needs to:

- Create a security group that will represent the smart card logon.
- Create a mapping from the smart card certificate OID to and smart card security group.
- Domain controllers need to have Windows Server 2008 R2 Domain Functional Level

Note

This group membership will be lost if the server performs S4U2Self on behalf of the user

Best Practices for Deploying Central Access Policies

The following sections provide additional guidance for best practices for delegation of administration, setting up exception mechanisms, and so on.

Delegating of administration for Dynamic Access Control

Ideally, you should delegate permissions for all Dynamic Access Control containers in the forest root of your domain controller. You would need to grant read-write permissions for to specific security groups that have access to certain objects. For example, you can create universal groups like:

- DAC Claim Admins
- DAC Resource Property Admins
- DAC Central Access Rule Admins
- DAC Central Access Policy Admins

You can then delegate the corresponding rights to these groups.

Following are the built-in Dynamic Access Control containers in Windows Server 2012. To access these containers, browse to configuration partition-> system->services->claims configuration.

- Claim Types
- Resource Properties
- Central Access Rules
- Central Access Policies

You can set up delegation of policies from Active Directory Administrative Center (ADAC).

To set up delegation of permissions from ADAC

1. From Server Manager, on the **Tools** menu, select **Active Directory Administrative Center**.
2. Select Dynamic Access Control on the left pane, and select the container that has the object that you want to delegate permissions to. For example, select **Central Access Policies** and select a policy from the list.
3. Right-click and select **Properties**. In the Policy Properties window, select the **Extensions** tab.
4. Select the **Security** tab. Click **Advanced**. In the Advanced Security Settings dialog box, click **Add**. In the Permission entry dialog box, click **Select a Principal** and type the name of the security group to which you want to grant access to. Click OK.
5. In the Permission entry dialog box, select the permissions you want to grant for the group. Click **OK** three times.

Exception Mechanisms for Planning Central Access Policies

Exceptions to the common access rules are a key component of every access policy and in particular a central access policy. For example, if a central policy for the organization protects access to high business impact data (HBI) so that only full-time employees can access this data, what happens if a vendor in the finance department requires access to financial information that is considered to be HBI so that he can do his job?

There are several valid answers ranging from using one security group in the central access policy to grant access to HBI data (which would then allow the vendor in this example to access both financial HBI data and other HBI data such as HR or engineering) to having multiple security groups for exceptions to having per file- or user-specific exception mechanisms.

Exception Mechanism	Central Access Policy scenario	Description	Pros	Cons
Security group	Any access rule where we are OK with providing access to all the information covered by the central access policy when granting the exception Example is a central access rule for department data that implements a safety net over ALL of the departments data (e.g.: Resource.Department=Finance)	One security group managed by the department IT or Information Security personal with a delegation to allow department senior personal to add / remove (manage) users from the group. Content owners from the department will be able to manage the exception group so that they control	Simple to implement Provides a good delegation model with existing tools to manage groups	Exception is wide and grants access to information that might not be relevant for this specific user

Exception Mechanism	Central Access Policy scenario	Description	Pros	Cons
		who can access information to the department data		
Multiple security groups	Central access rule for access to a range of countries with a finite (e.g.: <20) number of permutations where we would like to have each one controlled by a different content owner Example is a central access rule for access to country data Resource.Country = User.Country	A security group for each of the countries. Each security group is managed by the content owners of that country	Simple to implement Provides a good delegation model with existing tools to manage groups	The more groups there are, the more complex the access conditions becomes. Each time a new instance is added (for example, a new country), the administrator needs to create a new group to manage the exception for it

Tools for Deployment

Data Classification Toolkithttp://go.microsoft.com/fwlink/?LinkId=217654: The Data Classification Toolkit for Windows Server 2012 is designed to help organizations identify, classify, and protect data on their file servers. The out-of-the-box classification and rule examples help organizations build and deploy their policies to protect critical information in a cost-effective manner.The toolkit supports both Windows Server 2012 file servers and Windows Server 2008 R2 file servers. In addition to configuring file classification infrastructure, the latest version of the toolkit allows you to manage a Central Access Policy across the file servers in your organization. It provides tools for you to provision user and device claim values, and manage Central Access Policy across the forest to help simplify the configuration process of Dynamic Access Controls in Windows Server 2012. The toolkit also provides a new report template that you can use to review existing Central Access Policy on file shares.

Appendix: Deployment Configurations for Central Access Policies

Deployment Option	Requirements and Configuration Options	Guidance
Using Security Groups	• Windows Server 2012 File Server • A domain with a Windows Server 2012 schema	Using Security Groups for Dynamic Access Control
Using User Claims	• Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites For cross-forest claims deployment, you need a two-way trust been the user domain and a resource domain. For more information, see Deploy Claims Across Forests • Configuration 1: Domains providing claims and compound authentication have all Windows Server 2012 DCs	Using User Claims
Using Device Claims	• Configuration 2: Only user claim-based access control, so file servers retrieve user claims and domains providing claims have Windows Server 2012 domain controllers in all the file server sites • Configuration 1: Domains providing claims and compound authentication have all Windows Server	Using Device Claims and Device Security Groups

Deployment Option	Requirements and Configuration Options	Guidance
	2012 DCs • Configuration 3: Device-based access control needed, but cannot wait until all domain controllers can be upgraded	

Deploy a Central Access Policy (Demonstration Steps)

In this scenario, the finance department security operations is working with central information security to specify the need for a central access policy so that they can protect archived finance information stored on file servers. The archived finance information from each country can be accessed as read-only by finance employees from the same country. A central finance admin group can access the finance information from all countries.

Deploying a central access policy includes the following phases:

Phase	Description
Plan: Identify the need for policy and the configuration required for deployment	Identify the need for a policy and the configuration required for deployment.
Implement: Configure the components and policy	Configure the components and policy.
Deploy the central access policy	Deploy the policy.
Maintain: Change and stage the policy	Policy changes and staging

Set up a test environment

Before you begin, you need to set up lab to test this scenario. The steps for setting up the lab are explained in detail in Appendix B: Setting Up the Test Environment.

Plan: Identify the need for policy and the configuration required for deployment

This section provides the high-level series of steps that aid in the planning phase of your deployment.

<table>
<tr><th></th><th>Step</th><th>Example</th></tr>
<tr><td>1.1</td><td>Business determines that a central access policy is needed</td><td>To protect finance information that is stored on file servers, the finance department security operations is working with central information security to specify the need for a central access policy.</td></tr>
<tr><td>1.2</td><td>Express the access policy</td><td>Finance documents should only be read by members of the Finance department. Members of the Finance department should only access documents in their own country. Only Finance Administrators should have write-access. An exception will be allowed for members of the FinanceException group. This group will have Read access.</td></tr>
<tr><td>1.3</td><td>Express the access policy in Windows Server 2012 constructs</td><td><table>
<tr><td>Targeting:</td><td>• Resource.Department Contains Finance</td></tr>
<tr><td>Access rules:</td><td>• Allow read User.Country=Resource.Country AND User.department = Resource.Department
• Allow Full control User.MemberOf(FinanceAdmin)</td></tr>
<tr><td>Exception:</td><td>Allow read memberOf(FinanceException)</td></tr>
</table></td></tr>
<tr><td>1.4</td><td>Determine the file properties required for the policy</td><td>Tag files with:
• Department
• Country</td></tr>
<tr><td>1.5</td><td>Determine the claim types and groups required for the policy</td><td>Claim types:
• Country
• Department</td></tr>
</table>

	Step	Example
		User groups: • FinanceAdmin • FinanceException
1.6	Determine the servers on which to apply this policy	Apply the policy on all finance file servers.

Implement: Configure the components and policy

This section presents an example that deploys a central access policy for finance documents.

No	Step	Example
2.1	Create claim types	Create the following claim types: • Department • Country
2.2	Create resource properties	Create and enable the following resource properties: • Department • Country
2.3	Configure a central access rule	Create a Finance Documents rule that includes the policy determined in the previous section.
2.4	Configure a central access policy (CAP)	Create a CAP called Finance Policy and add the Finance Documents rule to that CAP.
2.5	Target central access policy to the file servers	Publish the Finance Policy CAP to the file servers.
2.6	Enable KDC Support for claims, compound authentication and Kerberos armoring.	Enable KDC Support for claims, compound authentication and Kerberos armoring for contoso.com.

In the following procedure, you create two claim types: Country and Department.

To create claim types

1. Open Server DC1 in Hyper-V Manager and log on as contoso\administrator, with the

password **pass@word1**.

2. Open Active Directory Administrative Center.
3. Click the **Tree View icon**, expand **Dynamic Access Control**, and then select **Claim Types**.

   Right-click **Claim Types**, click **New**, and then click **Claim Type**.

   Tip

   You can also open a **Create Claim Type:** window from the **Tasks** pane. On the **Tasks** pane, click **New**, and then click **Claim Type**.
4. In the **Source Attribute** list, scroll down the list of attributes, and click **department**. This should populate the **Display name** field with **department**. Click **OK**.
5. In **Tasks** pane, click **New**, and then click **Claim Type**.
6. In the **Source Attribute** list, scroll down the list of attributes, and then click the **c** attribute (Country-Name). In the **Display name** field, type **country**.
7. In the **Suggested Values** section, select **The following values are suggested:**, and then click **Add**.
8. In the **Value** and **Display name** fields, type **US**, and then click **OK**.
9. Repeat the above step. In the **Add a suggest value** dialog box, type **JP** in the **Value** and **Display name** fields, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADClaimType country -SourceAttribute c -SuggestedValues:@((New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("US","US","")), (New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("JP","JP","")))
New-ADClaimType department -SourceAttribute department
```

Tip

You can use the Windows PowerShell History Viewer in Active Directory Administrative Center to look up the Windows PowerShell cmdlets for each procedure you perform in Active Directory Administrative Center. For more information, see Windows PowerShell History Viewer

The next step is to create resource properties. In the following procedure you create a resource property that is automatically added to the Global Resource Properties list on the domain controller, so that it is available to the file server.

To create and enable pre-created resource properties

1. In the left pane of Active Directory Administrative Center, click **Tree View**. Expand **Dynamic Access Control**, and then select **Resource Properties**.
2. Right-click **Resource Properties**, click **New**, and then click **Reference Resource Property**.

   Tip
   You can also choose a resource property from the **Tasks** pane. Click **New** and then click **Reference Resource Property**.
3. In **Select a claim type to share its suggested values list**, click **country**.
4. In the **Display name** field, type **country**, and then click **OK**.
5. Double-click the **Resource Properties** list, scroll down to the **Department** resource property. Right-click, and then click **Enable**. This will enable the built-in **Department** resource property.
6. In the **Resource Properties** list on the navigation pane of the Active Directory Administrative Center, you will now have two enabled resource properties:
   - Country
   - Department

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADResourceProperty Country -IsSecured $true -ResourcePropertyValueType MS-DS-MultivaluedChoice -SharesValuesWith country

Set-ADResourceProperty Department_MS -Enabled $true

Add-ADResourcePropertyListMember "Global Resource Property List" -Members Country

Add-ADResourcePropertyListMember "Global Resource Property List" -Members Department_MS
```

The next step is to create central access rules that define who can access resources. In this scenario the business rules are:

- Finance documents can be read only by members of the Finance department.
- Members of the Finance department can access only documents in their own country.
- Only Finance Administrators can have Write access.
- We will allow an exception for members of the FinanceException group. This group will have Read access.
- The administrator and document owner will still have full access.

Or to express the rules with Windows Server 2012 constructs:

Targeting: Resource.Department Contains Finance

Access Rules:

Allow Read User.Country=Resource.Country AND User.department = Resource.Department

Allow Full control User.MemberOf(FinanceAdmin)

Allow Read User.MemberOf(FinanceException)

To create a central access rule

1. In the left pane of the Active Directory Administrative Center, click **Tree View**, select **Dynamic Access Control**, and then click **Central Access Rules**.
2. Right-click **Central Access Rules**, click **New**, and then click **Central Access Rule**.
3. In the **Name** field, type **Finance Documents Rule**.
4. In the **Target Resources** section, click **Edit**, and in the **Central Access Rule** dialog box, click **Add a condition**. Add the following condition:
   [**Resource**] [**Department**] [**Equals**] [**Value**] [**Finance**], and then click **OK**.
5. In the **Permissions** section, select **Use following permissions as current permissions**, click **Edit**, and in the **Advanced Security Settings for Permissions** dialog box click **Add**.

   Note
   Use the following permissions as proposed permissions option lets you create the policy in staging. For more information on how to do this refer to the Maintain: Change and stage the policy section in this topic.
6. In the **Permission entry for Permissions** dialog box, click **Select a principal**, type **Authenticated Users**, and then click **OK**.
7. In the **Permission Entry for Permissions** dialog box, click **Add a condition**, and add the following conditions:
   [**User**] [**country**] [**Any of**] [**Resource**] [**country**]
   Click **Add a condition**.
   [**And**]
   Click [**User**] [**Department**] [**Any of**] [**Resource**] [**Department**]. Set the **Permissions** to **Read**.
8. Click **OK**, and then click **Add**. Click **Select a principal**, type **FinanceAdmin**, and then click **OK**.
9. Select the **Modify, Read and Execute, Read, Write** permissions, and then click **OK**.
10. Click **Add**, click **Select a principal**, type **FinanceException**, and then click **OK**. Select the permissions to be **Read** and **Read and Execute**.
11. Click **OK** three times to finish and return to Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$countryClaimType = Get-ADClaimType country
$departmentClaimType = Get-ADClaimType department
$countryResourceProperty = Get-ADResourceProperty Country
```

```
$departmentResourceProperty = Get-ADResourceProperty
Department
$currentAcl =
"O:SYG:SYD:AR(A;;FA;;;OW)(A;;FA;;;BA)(A;;0x1200a9;;;S-1-5-21-
1787166779-1215870801-2157059049-1113)(A;;0x1301bf;;;S-1-5-
21-1787166779-1215870801-2157059049-
1112)(A;;FA;;;SY)(XA;;0x1200a9;;;AU;((@USER." +
$countryClaimType.Name + " Any_of @RESOURCE." +
$countryResourceProperty.Name + ") && (@USER." +
$departmentClaimType.Name + " Any_of @RESOURCE." +
$departmentResourceProperty.Name + ")))"
$resourceCondition = "(@RESOURCE." +
$departmentResourceProperty.Name + " Contains {`"Finance`"})"
New-ADCentralAccessRule "Finance Documents Rule" -CurrentAcl
$currentAcl -ResourceCondition $resourceCondition
```

Important

In the above cmdlet example, the security identifiers (SIDs) for the group FinanceAdmin and users is determined at creation time and will be different in your example. For example, the provided SID value (S-1-5-21-1787166779-1215870801-2157059049-1113) for the FinanceAdmins needs to be replaced with the actual SID for the FinanceAdmin group that you would need to create in your deployment. You can use Windows PowerShell to look up the SID value of this group, assign that value to a variable, and then use the variable here. For more information, see Windows PowerShell Tip: Working with SIDs.

You should now have a central access rule that allows people to access documents from the same country and the same department. The rule allows the FinanceAdmin group to edit the documents, and it allows the FinanceException group to read the documents. This rule targets only documents classified as Finance.

To add a central access rule to a central access policy

1. In the left pane of the Active Directory Administrative Center, click **Dynamic Access Control**, and then click **Central Access Policies**.
2. In the **Tasks** pane, click **New**, and then click **Central Access Policy**.
3. In **Create Central Access Policy:**, type **Finance Policy** in the **Name** box.
4. In **Member central access rules**, click **Add**.
5. Double-click the **Finance Documents Rule** to the add it to the **Add the following central access rules** list , and then click **OK**.

6. Click **OK** to finish. You should now have a central access policy named Finance Policy.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADCentralAccessPolicy "Finance Policy" Add-
ADCentralAccessPolicyMember
-Identity "Finance Policy"
-Member "Finance Documents Rule"
```

To apply the central access policy across file servers by using Group Policy

1. On the **Start** screen, in the **Search** box, type **Group Policy Management**. Double-click **Group Policy Management**.

   Tip
   If the **Show Administrative tools** setting is disabled, the **Administrative Tools** folder and its contents will not appear in the **Settings** results.

   Tip
   In your production environment, you should create a File Server Organization Unit (OU) and add all your file servers to this OU, to which you want to apply this policy. You can then create a group policy and add this OU to that policy..
2. In this step, you edit the group policy object you created in Build the domain controller section in the Test Environment to include the central access policy that you created. In the Group Policy Management Editor, navigate to and select the organizational unit in the domain (contoso.com in this example): **Group Policy Management**, **Forest: contoso.com**, **Domains**, **contoso.com**, **Contoso**, **FileServerOU**.
3. Right-click **FlexibleAccessGPO**, and then click **Edit**.
4. In the Group Policy Management Editor window, navigate to **Computer Configuration**, expand **Policies**, expand **Windows Settings**, and click **Security Settings**.
5. Expand **File System**, right-click **Central Access Policy**, and then click **Manage Central access policies**.
6. In the **Central Access Policies Configuration** dialog box, add **Finance Policy**, and then click **OK**.
7. Scroll down to **Advanced Audit Policy Configuration**, and expand it.
8. Expand **Audit Policies**, and select **Object Access**.
9. Double-click **Audit Central Access Policy Staging**. Select all three check boxes and then click **OK**. This step allows the system to receive audit events related to Central Access Staging Policies.
10. Double-click **Audit File System Properties**. Select all three check boxes then click **OK**.

11. Close the Group Policy Management Editor. You have now included the central access policy to the Group Policy.

For a domain's domain controllers to provide claims or device authorization data, the domain controllers need to be configured to support dynamic access control.

To enable support for claims and compound authentication for contoso.com

1. Open Group Policy Management, click **contoso.com**, and then click **Domain Controllers**.
2. Right-click **Default Domain Controllers Policy**, and then click **Edit**.
3. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.
4. Double-click **KDC Support for claims, compound authentication and Kerberos armoring**. In the **KDC Support for claims, compound authentication and Kerberos armoring** dialog box, click **Enabled** and select **Supported** from the **Options** drop-down list. (You need to enable this setting to use user claims in central access policies.)
5. Close **Group Policy Management**.
6. Open a command prompt and type `gpupdate /force`.

Deploy the central access policy

	Step	Example
3.1	Assign the CAP to the appropriate shared folders on the file server.	Assign the central access policy to the appropriate shared folder on the file server.
3.2	Verify that access is appropriately configured.	Check the access for users from different countries and departments.

In this step you will assign the central access policy to a file server. You will log onto a file server that is receiving the central access policy that you created the previous steps and assign the policy to a shared folder.

To assign a central access policy to a file server

1. In Hyper-V Manager, connect to server FILE1. Log on to the server by using contoso\administrator with the password: **pass@word1**.
2. Open an elevated command prompt and type: **gpupdate /force**. This ensures that your Group Policy changes take effect on your server.
3. You also need to refresh the Global Resource Properties from Active Directory. Open an

elevated Windows PowerShell window and type `Update-FSRMClassificationpropertyDefinition`. Click ENTER, and then close Windows PowerShell.

Tip

You can also refresh the Global Resource Properties by logging on to the file server. To refresh the Global Resource Properties from the file server, do the following

a. Logon to File Server FILE1 as contoso\administrator, using the password **pass@word1**.

b. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.

c. In the File Server Resource Manager, click **File Classification Management** , right-click **Classification Properties** and then click **Refresh**.

4. Open Windows Explorer, and in the left pane, click drive D. Right-click the **Finance Documents** folder, and click **Properties**.
5. Click the **Classification** tab, click **Country**, and then select **US** in the **Value** field.
6. Click **Department**, then select **Finance** in the **Value** field and then click **Apply**.

Note

Remember that the central access policy was configured to target files for the Department of Finance. The previous steps mark all documents in the folder with the Country and Department attributes.

7. Click the **Security** tab, and then click **Advanced**. Click the **Central Policy** tab.

8. Click **Change**, select **Finance Policy** from the drop-down menu, and then click **Apply**. You can see the **Finance Documents Rule** listed in the policy. Expand the item to view all of the permissions that you set when you created the rule in Active Directory.
9. Click **OK** to return to Windows Explorer.

In the next step, you ensure that access is appropriately configured. User accounts need to have the appropriate Department attribute set (set this using Active Directory Administrative Center). The simplest way to view the effective results of the new policy is to use the **Effective Access** tab in Windows Explorer. The **Effective Access** tab shows the access rights for a given user account.

To examine the access for various users

1. In Hyper-V Manager, connect to server FILE1. Log on to the server by using contoso\administrator. Navigate to D:\ in Windows Explorer. Right-click the **Finance Documents** folder, and then click **Properties**.
2. Click the **Security** tab, click **Advanced**, and then click the **Effective Access** tab.
3. To examine the permissions for a user, click **Select a user**, type the user's name, and

then click **View effective access** to see the effective access rights. For example:

- Myriam Delesalle (MDelesalle) is in the Finance department and should have Read access to the folder.
- Miles Reid (MReid) is a member of the FinanceAdmin group and should have Modify access to the folder.
- Esther Valle (EValle) is not in the Finance department; however, she is a member of the FinanceException group and should have Read access.
- Maira Wenzel (MWenzel) is not in the Finance department and is not a member of either the FinanceAdmin or FinanceException group. She should not have any access to the folder.

Notice that the last column named **Access limited by** in the effective access window. This column tells you which gates are effecting the person's permissions. In this case, the Share and NTFS permissions allow all users full control. However, the central access policy restricts access based on the rules you configured earlier.

Maintain: Change and stage the policy

Number	Step	Example
4.1	Configure Device Claims for Clients	Set the group policy setting to enable device claims
4.2	Enable a claim for devices.	Enable the country claim type for devices.
4.3	Add a staging policy to the existing central access rule that you would like to modify.	Modify the Finance Documents Rule to add a staging policy.
4.4	View the results of the staging policy.	Check for Ester Velle's permissions.

To set up group policy setting to enable claims for devices

1. Log on to DC1, open Group Policy Management, click **contoso.com**, click **Default Domain Policy**, right-click and select **Edit**.
2. In the Group Policy Management Editor window, navigate to **Computer Configuration**, **Policies**, **Administrative Templates**, **System**, **Kerberos**.
3. Select **Kerberos client support for claims, compound authentication and Kerberos armoring** and click **Enable**.

To enable a claim for devices

1. Open Server DC1 in Hyper-V Manager and log on as contoso\Administrator, with the password **pass@word1**.
2. From the **Tools** menu, open Active Directory Administrative Center.
3. Click **Tree View**, expand **Dynamic Access Control**, double-click **Claim Types**, and double-click the **country** claim.
4. In **Claims of this type can be issued for the following classes**, select the **Computer** check box. Click **OK**.
   Both the **User** and **Computer** check boxes should now be selected. The country claim can now be used with devices in addition to users.

The next step is to create a staging policy rule. Staging policies can be used to monitor the effects of a new policy entry before you enable it. In the following step, you will create a staging policy entry and monitor the effect on your shared folder.

To create a staging policy rule and add it to the central access policy

1. Open Server DC1 in Hyper-V Manager and log on as contoso\Administrator, with the password **pass@word1**.
2. Open Active Directory Administrative Center.
3. Click **Tree View**, expand **Dynamic Access Control**, and select **Central Access Rules**.
4. Right-click **Finance Documents Rule**, and then click **Properties**.
5. In the **Proposed Permissions** section, select the **Enable permission staging configuration** check box, click **Edit**, and then click **Add**. In the **Permission Entry for Proposed Permissions** window, click the **Select a Principal** link, type **Authenticated Users**, and then click **OK**.
6. Click the **Add a condition** link and add the following condition:
   [**User**] [**country**] [**Any of**] [**Resource**] [**Country**].
7. Click **Add a condition** again, and add the following condition: [**And**]
   [**Device**] [**country**] [**Any of**] [**Resource**] [**Country**]
8. Click **Add a condition** again, and add the following condition. [And]
   [**User**] [**Group**] [**Member of any**] [**Value**](**FinanceException**)
9. To set the FinanceException, group, click **Add items** and in the **Select User, Computer, Service Account, or Group** window, type **FinanceException**.
10. Click **Permissions**, select **Full Control**, and click **OK**.
11. In the Advance Security Settings for Proposed Permissions window, select **FinanceException** and click **Remove**.
12. Click **OK** two times to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADCentralAccessRule
-Identity:
"CN=FinanceDocumentsRule,CN=CentralAccessRules,CN=ClaimsConfiguration,CN=Configuration,DC
=Contoso.com"
-ProposedAcl: "O:SYG:SYD:AR(A;;FA;;;BA)(A;;FA;;;SY)(A;;0x1301bf;;;S-1-21=1426421603-
1057776020-1604)"
-Server: "WIN-2R92NN8VKFP.Contoso.com"
```

Note

In the above cmdlet example, the Server value reflects the Server in the test lab environment. You can use the Windows PowerShell History Viewer to look up the Windows PowerShell cmdlets for each procedure you perform in Active Directory Administrative Center. For more information, see Windows PowerShell History Viewer

In this proposed permissions set, members of the FinanceException group will have Full Access to files from their own country when they access them through a device from the same country as the document. Audit entries are available in the File Servers security log when someone from the Finance department attempts to access files. However, security settings are not enforced until the policy is promoted from staging.

In the next procedure, you verify the results of the staging policy. You access the shared folder with a user name that has permissions based on the current rule. Esther Valle (EValle) is a member of FinanceException, and she currently has Read rights. According to our staging policy, EValle should not have any rights.

To verify the results of the staging policy

1. Connect to the File Server FILE1 in Hyper-V Manager and log on as contoso\administrator, with the password **pass@word1**.
2. Open a Command Prompt window and type **gpupdate /force**. This ensures that your Group Policy changes will take effect on your server.
3. In Hyper-V Manager, connect to server CLIENT1. Log off the user who is currently logged on. Restart the virtual machine, CLIENT1. Then log on to the computer by using contoso\EValle pass@word1.
4. Double-click the desktop shortcut to \\FILE1\Finance Documents. EValle should still have access to the files. Switch back to FILE1.
5. Open **Event Viewer** from the shortcut on the desktop. Expand **Windows Logs**, and then select **Security**. Open the entries with **Event ID 4818** under the **Central Access Policy Staging** task category. You will see that EValle was allowed access; however, according to the staging policy, the user would have been denied access.

Next Steps

If you have a central server management system such as System Center Operations Manager, you can also configuring monitoring for events. This allows Administrators to monitor the effects of central access policies before enforcing them.

Scenario: File Access Auditing

Security Auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is regulatory compliance. Industry standards such as Sarbanes Oxley, Health Insurance Portability and Accountability Act (HIPAA), and Payment Card Industry (PCI) require enterprises to follow a strict set of rules related to data security and privacy. Security audits help establish the presence of such policies and prove compliance with these standards. Additionally, security audits help detect anomalous behavior, identify and mitigate gaps in security policies, and deter irresponsible behavior by creating a trail of user activity that can be used for forensic analysis.

Audit policy requirements are typically driven at the following levels:

- **Information security.** File access audit trails are often used for forensic analysis and intrusion detection. Being able to get targeted events about access to high-value information lets organizations considerably improve their response time and investigation accuracy.
- **Organizational policy.** For example, organizations regulated by PCI standards could have a central policy to monitor access to all files that are marked as containing credit card information and personally identifiable information (PII).
- **Departmental policy.** For example, the finance department may require that the ability to modify certain finance documents (such as a quarterly earnings report) be restricted to the finance department, and thus the department would want to monitor all other attempts to change these documents.
- **Business policy.** For example, business owners may want to monitor all unauthorized attempts to view data that belongs to their projects.

Additionally, the compliance department may want to monitor all changes to central authorization policies and policy constructs such as user, computer, and resource attributes.

One of the biggest considerations of security audits is the cost of collecting, storing, and analyzing audit events. If the audit policies are too broad, the volume of audit events collected rises, and this increases costs. If the audit policies are too narrow, you risk missing important events.

With Windows Server 2012, you can author audit policies by using claims and resource properties. This leads to richer, more targeted, and easier-to-manage audit policies. It enables scenarios that, until now, were impossible or too difficult to perform. The following are examples of audit policies that administrators can author:

- Audit everyone who does not have a high-security clearance and tries to access an HBI document. For example, Audit | Everyone | All-Access | Resource.BusinessImpact=HBI AND User.SecurityClearance!=High.
- Audit all vendors when they try to access documents that are related to projects that they are not working on. For example, Audit | Everyone | All-Access | User.EmploymentStatus=Vendor AND User.Project Not_AnyOf Resource.Project.

These policies help regulate the volume of audit events and limit them to only the most relevant data or users.

After administrators have created and applied the audit policies, the next consideration for them is gleaning meaningful information from the audit events that they collected. Expression-based audit events help reduce the volume of audits. However, users need a way to query these events for meaningful information and ask questions such as, “Who is accessing my HBI data?” or “Was there an unauthorized attempt to access sensitive data?”

Windows Server 2012 enhances existing data access events with user, computer, and resource claims. These events are generated on a per-server basis. To provide a full view of events across the organization, Microsoft is working with partners to provide event collection and analysis tools, such as the Audit Collection Services in System Center Operation Manager .

Figure 4 shows an overview of a central audit policy.

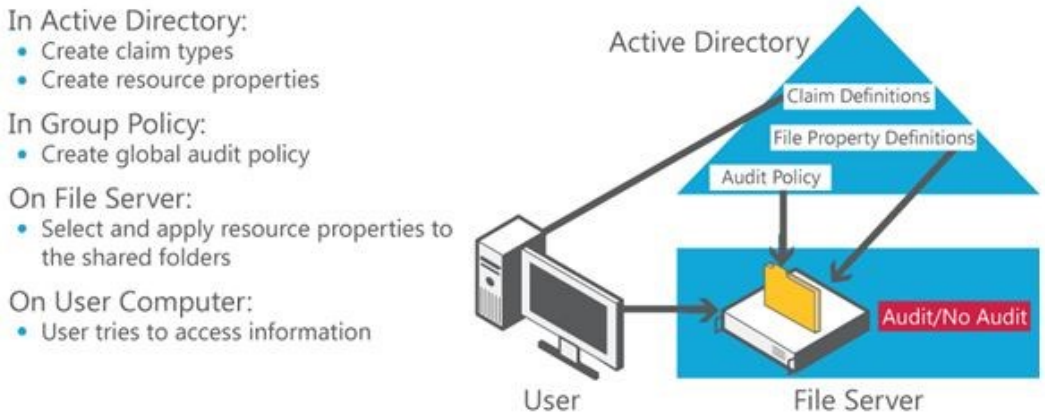

Figure 4 Central auditing experiences

Setting up and consuming security audits typically involves the following general steps:

1. Identify the correct set of data and users to monitor
2. Create and apply appropriate audit policies
3. Collect and analyze audit events
4. Manage and monitor the policies that were created

In this scenario

The following topics provide additional guidance for this scenario:

- Plan for File Access Auditing

- Deploy Security Auditing with Central Audit Policies (Demonstration Steps)

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Doman Services role	AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables creating user claims and device claims, compound identity, (user plus device claims), new central access policies (CAP) model, and the use of file classification information in authorization decisions.
File and Storage Services role	File servers in Windows Server 2012 provide a user interface where administrators can view the effective permissions for users for a file or folder and troubleshoot access issues and grant access as required.

Plan for File Access Auditing

The information in this topic explains the security auditing enhancements that are introduced in Windows Server 2012 and new audit settings that you should consider as you deploy Dynamic Access Control in your enterprise. The actual audit policy settings that you deploy will depend on your goals, which can include regulatory compliance, monitoring, forensic analysis, and troubleshooting.

Note
Detailed information about how to plan and deploy an overall security auditing strategy for your enterprise is explained in Planning and Deploying Advanced Security Audit Policies. For more information about configuring and deploying a security audit policy, see the Advanced Security Audit Policy Step-by-Step Guide.

The following security auditing capabilities in Windows Server 2012 can be used with Dynamic Access Control to extend your overall security auditing strategy.

- **Expression-based audit policies**. Dynamic Access Control enables you to create targeted audit policies by using expressions based on user, computer, and resource claims. For example, you could create an audit policy to track all Read and Write operations on files classified as high-business impact by employees who do not have a high-security clearance.

Expression-based audit policies can be authored directly for a file or folder or centrally through Group Policy. For more information, see Group Policy using Global Object Access Auditing.

- **Additional information from object access auditing**. File access auditing is not new to Windows Server 2012. With the right audit policy in place, the Windows and Windows Server operating systems generate an audit event each time a user accesses a file. Existing File Access events (4656, 4663) contain information about the attributes of the file that was accessed. This information can be used by event log filtering tools to help you identify the most relevant audit events. For more information, see Audit Handle Manipulation and Audit Security Accounts Manager.
- **More information from user logon events**. With the right audit policy in place, Windows operating systems generate an audit event every time a user signs in to a computer locally or remotely. In Windows Server 2012 or Windows 8, you can also monitor user and device claims associated with a user's security token. Examples can include Department, Company, Project, and Security clearances.Event 4626 contains information about these user claims and device claims, which can be leveraged by audit log management tools to correlate user logon events with object access events to enable event filtering based on file attributes and user attributes. For more information about user logon auditing, see Audit Logon.
- **Change tracking for new types of securable objects**. Tracking changes to securable objects can be important in the following scenarios:
  - **Change tracking for central access policies and central access rules**. Central access policies and central access rules define the central policy that can be used to control access to critical resources. Any change to these can directly impact the file access permissions that are granted to users on multiple computers. Therefore, tracking changes to central access policies and central access rules can be important for your organization. Because central access policies and central access rules are stored in Active Directory Domain Services (AD DS), you can audit attempts to modify them, like auditing changes to any other securable object in AD DS. For more information, see Audit Directory Service Access.
  - **Change tracking for definitions in the claim dictionary**. Claim definitions include the claim name, description, and possible values. Any change to the claim definition can impact the access permissions on critical resources. Therefore, tracking changes to claim definitions can be important to your organization. Like central access policies and central access rules, claim definitions are stored in AD DS; therefore, they can be audited like any another securable object in AD DS. For more information, see Audit Directory Service Access.
  - **Change tracking for file attributes**. File attributes determine which central access rule applies to the file. A change to the file attributes can potentially impact the access restrictions on the file. Therefore, it can be important to track changes to file attributes. You can track changes to file attributes on any computer by configuring the authorization policy change auditing policy. For more information, see Authorization Policy Change auditing and Object Access auditing for File Systems. In Windows Server 2012, Event 4911 differentiates file attribute policy changes from other authorization policy change events.

- **Chang tracking for the central access policy associated with a file.** Event 4913 displays the security identifiers (SIDs) of the old and new central access policies. Each central access policy also has a user friendly name that can be looked up using this security identifier. For more information, see Authorization Policy Change auditing.
- **Change tracking for user and computer attributes**. Like files, user and computer objects can have attributes, and changes to these attributes can impact the user's ability to access files. Therefore, it can be valuable to track changes to user or computer attributes. User and computer objects are stored in AD DS; therefore, changes to their attributes can be audited. For more information, see DS Access .

- **Policy change staging**. Changes to central access policies can impact the access control decisions on all computers where the policies are enforced. A loose policy could grant more access than desired, and an overly restrictive policy could generate an excessive number of Help Desk calls. As a result, it can be extremely valuable to verify changes to a central access policy before enforcing the change. For that purpose, Windows Server 2012 introduces the concept of "staging." Staging enables users to verify their proposed policy changes before enforcing them. To use policy staging, proposed policies are deployed with the enforced policies, but staged policies do not actually grant or deny permissions. Instead, Windows Server 2012 logs an audit event (4818) any time the result of the access check that uses the staged policy is different from the result of an access check that uses the enforced policy.

Deploy Security Auditing with Central Audit Policies (Demonstration Steps)

In this scenario, you will audit access to files in the Finance Documents folder by using the Finance Policy that you created in Deploy a Central Access Policy (Demonstration Steps). If a user who is not authorized to access the folder attempts to access it, the activity is captured in the event viewer.

The following steps are required to test this scenario.

Task	Description
Configure global object access policy	In this step, you configure the global object access policy on the domain controller.
Update Group Policy settings	Sign in to the file server and apply the Group Policy update.
Verify that the global object access policy has been applied	View the relevant events in the event viewer. The events should include metadata for the country and document type.

Configure global object access policy

In this step, you configure the global object access policy in the domain controller.

To configure a global object access policy

1. Sign in to the domain controller DC1 as contoso\administrator with the password **pass@word1**.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Domains**, double-click **contoso.com**, click **Contoso**, and then double-click **File Servers**.
4. Right-click **FlexibleAccessGPO**, and click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **Audit Policies**.
7. Double-click **Object Access**, and then double-click **Audit File System**.
8. Select the **Configure the following events** check box, select the **Success** and **Failure** check boxes, and then click **OK**.
9. In the navigation pane, double-click **Global Object Access Auditing**, and then double-click **File system**.
10. Select the **Define this policy setting** check box, and click **Configure**.
11. In the **Advanced Security Settings for Global File SACL** box, click **Add**, then click **Select a principal**, type **Everyone**, and then click **OK**.
12. In the **Auditing Entry for Global File SACL** box, select **Full control** in the **Permissions** box.
13. In the **Add a condition:** section, click **Add a condition** and in the drop-down lists select [**Resource**] [**Department**] [**Any of**] [**Value**] [**Finance**].
14. Click **OK** three times to complete the configuration of the global object access audit policy setting.
15. In the navigation pane, click **Object Access**, and in the results pane, double-click **Audit Handle Manipulation**. Click **Configure the following audit events**, **Success**, and **Failure**, click **OK**, and then close the flexible access GPO.

Update Group Policy settings

In this step, you update the Group Policy settings after you have created the audit policy.

To update Group Policy settings

1. Sign in to the file server, FILE1 as contoso\Administrator, with the password **pass@word1**.
2. Press the Windows key+R, then type **cmd** to open a Command Prompt window.

Note

If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

3. Type **gpupdate /force** and then press ENTER.

Verify that the global object access policy has been applied

After the Group Policy settings have been applied, you can verify that the audit policy settings were applied correctly.

To verify that the global object access policy has been applied

1. Sign in to client computer, CLIENT1 as Contoso\MReid. Browse to the folder \\FILE1\Finance Documents, and modify Word Document 2.
2. Sign in to the file server, FILE1 as contoso\administrator. Open Event Viewer, browse to **Windows Logs**, select **Security**, and confirm that your activities resulted in audit events **4656** and **4663** (even though you did not set explicit auditing SACLs on the files or folders that you created, modified, and deleted).

Important

A new logon event is generated on the computer where the resource is located, on behalf of the user for whom effective access is being checked. When analyzing security audit logs for user sign-in activity, to differentiate between logon events that are generated because of effective access and those generated because of an interactive network user sign in, the Impersonation Level information is included. When the logon event is generated because of effective access, the Impersonation Level will be Identity. A network interactive user sign in typically generates a logon event with the Impersonation Level = Impersonation or Delegation.

See also

- Scenario: File Access Auditing
- Plan for File Access Auditing
- Dynamic Access Control: Scenario Overview

Scenario: Access-Denied Assistance

Users will get an access-denied message when they try to access shared files and folders on a file server for which they do not have permissions. Administrators often do not have the appropriate context to troubleshoot the access issue, which makes it hard to resolve the issue.

Scenario description

Access-denied assistance is a new feature in Windows Server 2012, which provides the following ways to troubleshoot issues that are related to access to files and folders:

- **Self-assistance.** If a user can determine the issue and remediate the problem so that they can get the requested access, the impact to the business is low, and no special exceptions are needed in the central access policy. Access-denied assistance provides an access-denied message that file server administrators can customize with information specific to their organizations. For example, an administrator could set the message so that users can request access from a data owner without involving the file server administrator.
- **Assistance by the data owner.** You can define a distribution list for shared folders, and configure it so that the folder owner receives an email notification when a user needs access. If the data owner does not know how to help the user get access, the owner can forward this information to the file server administrator.
- **Assistance by the file server administrator.** This type of assistance is available when the user cannot fix an issue and the data owner cannot help. Windows Server 2012 provides a user interface where file server administrators can view the effective permissions for a user on a file or folder so that it is easier to troubleshoot access issues.

Access-denied assistance in Windows Server 2012 provides file server administrators the relevant access details so that they can determine the issue and appropriate tools so that they can make configuration changes to satisfy the access request. For example, a user might follow this process to access a file that they currently do not have access to:

- The user attempts to read a file in the \\financeshares shared folder, but the server displays an access-denied message.
- Windows Server 2012 displays the access-denied assistance information to the user with an option to request assistance.
- If the user requests access to the resource, the server sends an email with the access request information to the folder owner.

You can find planning information for configuring access-denied assistance in Plan for Access-Denied Assistance.

You can find steps about configuring access-denied assistance in Deploy Access-Denied Assistance (Demonstration Steps).

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Practical applications

Access-denied assistance in Windows Server 2012 contributes to Dynamic Access Control by giving users the ability to request access to shared files and folders directly from an access-denied message.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	Access-denied assistance can be configured by using the File Server Resource Manager console on the file server.
File and Storage Services Overview	File Server Resource Manager is a File and Storage Services role service, and it is comprised of a set of features that can be used to administer the file servers on your network.

Deploy Access-Denied Assistance (Demonstration Steps)

This topic explains how to configure access-denied assistance, and verify that it is working properly.

In this document

- Step 1: Configure access-denied assistance
- Step 2: Configure the email notification settings
- Step 3: Verify that access-denied assistance is configured correctly

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Configure access-denied assistance

You can configure access-denied assistance within a domain by using Group Policy, or you can configure the assistance individually on each file server by using the File Server Resource

Manager console. You can also change the access-denied message for a specific shared folder on a file server.

You can configure access-denied assistance for the domain by using Group Policy as follows:

Do this step using Windows PowerShell

To configure access-denied assistance by using Group Policy

1. Open Group Policy Management. In Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Right-click the appropriate Group Policy, and then click **Edit**.
3. Click **Computer Configuration**, click **Policies**, click **Administrative Templates**, click **System**, and then click **Access-Denied Assistance**.
4. Right-click **Customize message for Access Denied errors**, and then click **Edit**.
5. Select the **Enabled** option.
6. Configure the following options:
    a. In the **Display the following message to users who are denied access** box, type a message that users will see when they are denied access to a file or folder.

       You can add macros to the message that will insert customized text. The macros include:

       - **[Original File Path]** The original file path that was accessed by the user.
       - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
       - **[Admin Email]** The administrator email recipient list.
       - **[Data Owner Email]** The data owner email recipient list.

    b. Select the **Enable users to request assistance** check box.
    c. Leave the remaining default settings.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-GPRegistryValue –Name "Name of GPO" –key
"HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied" –ValueName AllowEmailRequests
–Type DWORD –value 1

Set-GPRegistryValue –Name "Name of GPO" –key
"HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied" –ValueName GenerateLog –Type
DWORD –value 1

Set-GPRegistryValue –Name "Name of GPO" –key
"HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied" –ValueName
IncludeDeviceClaims –Type DWORD –value 1
```

```
Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName IncludeUserClaims
–Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName PutAdminOnTo –Type
DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName PutDataOwnerOnTo –
Type DWORD –value 1

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName ErrorMessage –Type
MultiString –value “Type the text that the user will see in the error message dialog
box.”

Set-GPRegistryValue –Name “Name of GPO” –key
“HKLM\Software\Policies\Microsoft\Windows\ADR\AccessDenied” –ValueName Enabled –Type
DWORD –value 1
```

Alternatively, you can configure access-denied assistance individually on each file server by using the File Server Resource Manager console.

Do this step using Windows PowerShell

To configure access-denied assistance by using File Server Resource Manager

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
3. Click the **Access-Denied Assistance** tab.
4. Select the **Enable access-denied assistance** check box.
5. In the **Display the following message to users who are denied access to a folder or file** box, type a message that users will see when they are denied access to a file or folder.

   You can add macros to the message that will insert customized text. The macros include:

   - **[Original File Path]** The original file path that was accessed by the user.
   - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
   - **[Admin Email]** The administrator email recipient list.
   - **[Data Owner Email]** The data owner email recipient list.
6. Click **Configure email requests**, select the **Enable users to request assistance** check box, and then click **OK**.
7. Click **Preview** if you want to see how the error message will look to the user.

8. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FSRMAdrSetting -Event "AccessDenied" –DisplayMessage "Type the text that the user
will see in the error message dialog box." -Enabled:$true -AllowRequests:$true
```

After you configure the access-denied assistance, you must enable it for all file types by using Group Policy.

Do this step using Windows PowerShell

To configure access-denied assistance for all file types by using Group Policy

1. Open Group Policy Management. In Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Right-click the appropriate Group Policy, and then click **Edit**.
3. Click **Computer Configuration**, click **Policies**, click **Administrative Templates**, click **System**, and then click **Access-Denied Assistance**.
4. Right-click **Enable access-denied assistance on client for all file types**, and then click **Edit**.
5. Click **Enabled**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-GPRegistryValue –Name "Name of GPO" –key
"HKLM\SOFTWARE\Policies\Microsoft\Windows\Explorer" –ValueName
EnableShellExecuteFileStreamCheck –Type DWORD –value 1
```

You can also specify a separate access-denied message for each shared folder on a file server by using the File Server Resource Manager console.

Do this step using Windows PowerShell

To specify a separate access-denied message for a shared folder by using File Server Resource Manager

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Expand **File Server Resource Manager (Local)**, and then click **Classification Management**.
3. Right-click **Classification Properties**, and then click **Set Folder Management**

Properties.

4. In the **Property** box, click **Access-Denied Assistance Message**, and then click **Add**.
5. Click **Browse**, and then choose the folder that should have the custom access-denied message.
6. In the **Value** box, type the message that should be presented to the users when they cannot access a resource within that folder.

   You can add macros to the message that will insert customized text. The macros include:

   - **[Original File Path]** The original file path that was accessed by the user.
   - **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
   - **[Admin Email]** The administrator email recipient list.
   - **[Data Owner Email]** The data owner email recipient list.
7. Click **OK**, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FSRMMgmtProperty -Namespace "folder path" -Name "AccessDeniedMessage_MS" -Value "Type the text that the user will see in the error message dialog box."
```

Step 2: Configure the email notification settings

You must configure the email notification settings on each file server that will send the access-denied assistance messages.

Do this step using Windows PowerShell

1. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. Right-click **File Server Resource Manager (Local)**, and then click **Configure Options**.
3. Click the **Email Notifications** tab.
4. Configure the following settings:
   - In the **SMTP server name or IP address** box, type the name of IP address of the SMTP server in your organization.
   - In the **Default administrator recipients** and **Default "From" e-mail address** boxes, type the email address of the file server administrator.
5. Click **Send Test E-mail** to ensure that the email notifications are configured correctly.
6. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
set-FSRMSetting -SMTPServer "server1" -AdminEmailAddress "fileadmin@contoso.com" -
FromEmailAddress "fileadmin@contoso.com"
```

Step 3: Verify that access-denied assistance is configured correctly

You can verify that the access-denied assistance is configured correctly by having a user who is running Windows 8 try to access a share or a file in that share that they do not have access to. When the access-denied message appears, the user should see a **Request Assistance** button. After clicking the Request Assistance button, the user can specify a reason for access and then send an email to the folder owner or file server administrator. The folder owner or file server administrator can verify for you that the email arrived and contains the appropriate details.

Important
If you want to verify access-denied assistance by having a user who is running Windows Server 2012, you must install the Desktop Experience before connecting to the file share.

See also

- Scenario: Access-Denied Assistance
- Plan for Access-Denied Assistance
- Dynamic Access Control: Scenario Overview

Scenario: Classification-Based Encryption for Office Documents

Protection of sensitive information is mainly about mitigating risk for the organization. Various compliance regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and Payment Card Industry Data Security Standard (PCI-DSS), dictate encryption of information, and there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it might impair business productivity. Thus, organizations tend to have different approaches and priorities for encrypting their information.

Scenario description

Windows Server 2012 provides the ability to automatically encrypt sensitive Microsoft Office files, based on their classification. This is done through file management tasks that invoke Active Directory Rights Management Services (AD RMS) protection for sensitive documents a few

seconds after the file is identified as being a sensitive file on the file server. This is facilitated by continuous file management tasks on the file server.

AD RMS encryption provides another layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the file is protected by the AD RMS encryption. Users who want to access the file must first authenticate themselves to an AD RMS server to receive the decryption key. The following figure shows this process.

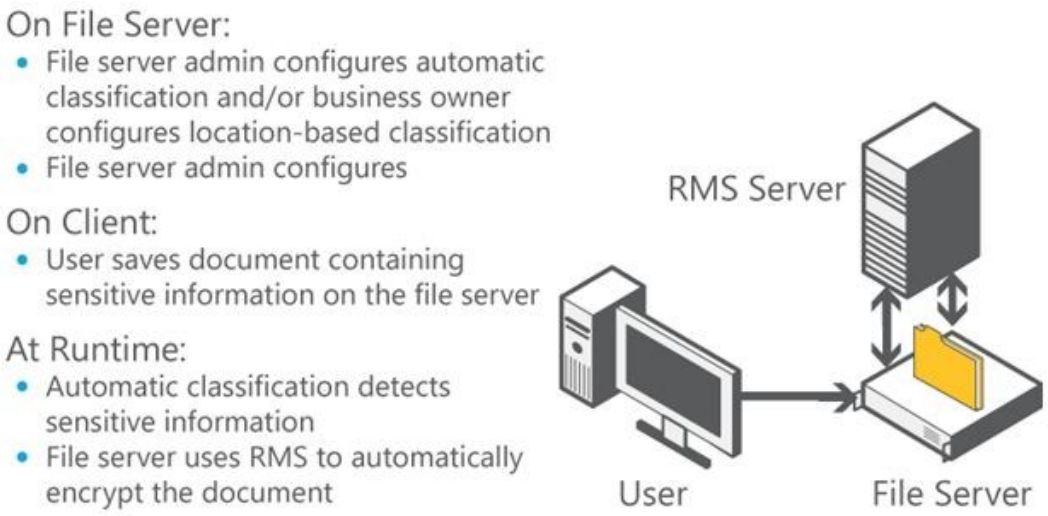

Figure 6 Classification-based RMS protection

Support for non-Microsoft file formats is available through non-Microsoft vendors. After a file has been protected by AD RMS encryption, data management features such as search- or content-based classification are no longer available for that file.

In this scenario

Following is the guidance that is available for this scenario:

- Planning Considerations for Encryption of Office Documents
- Deploy Encryption of Office Files (Demonstration Steps)
- Dynamic Access Control: Scenario Overview

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services role (AD DS)	AD DS provides a distributed database that stores and manages information about network resources and application-specific data from directory-enabled applications. In this scenario, AD DS in Windows Server 2012 introduces a claims-based authorization platform that enables the creation of user claims and device claims, compound identity (user plus device claims), a new central access policies model, and the use of file-classification information in authorization decisions.
File and Storage Services role File Server Resource Manager	File and Storage Services provides technologies to help you set up and manage one or more file servers that provide central locations on your network where you can store files and share them with users. If your network users need access to the same files and applications, or if centralized backup and file management are important to your organization, you should set up one or more computers as a file server by adding the File and Storage Services role and the appropriate role services to the computers. In this scenario, file server administrators can configure file management tasks that invoke AD RMS protection for sensitive documents a few seconds after the file is identified as being a sensitive file on the file server (continuous file management tasks on the file server).
Active Directory Rights Management Services (AD RMS) role	AD RMS enables individuals and administrators (through Information Rights Management (IRM) policies) to specify access permissions to documents, workbooks, and presentations. This helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. After permission for a file has been restricted by using IRM, the access and usage restrictions are enforced no matter where the information is, because the permission to a file is stored in the document

Role/feature	How it supports this scenario
	file itself. In this scenario, AD RMS encryption provides another layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the file is protected by the AD RMS encryption. Users who want to access the file must first authenticate themselves to an AD RMS server to receive the decryption key.

Deploy Encryption of Office Files (Demonstration Steps)

Contoso's Finance Department has a number of file servers that store their documents. These documents can be general documentation or they can have a high-business impact (HBI). For example, any document that contains confidential information is deemed, by Contoso, to have a high-business impact. Contoso wants to ensure that all their documentation has a minimum amount of protection and that their HBI documentation is restricted to the appropriate people. To accomplish this, Contoso is exploring using the File Classification Infrastructure (FCI) and AD RMS that is available in Windows Server 2012. By using FCI, Contoso will classify all of the documents on their file server, based on the content, and then use AD RMS to apply the appropriate rights policy.

In this scenario, you'll perform the following steps:

Task	Description
Step 1: Enable resource properties	Enable the **Impact** and **Personally Identifiable Information** resource properties.
Step 2: Create classification rules	Create the following classification rules: **HBI Classification Rule** and **PII Classification Rule**.
Step 3: Use file management tasks to automatically protect documents with AD RMS	Create a file management task that automatically used AD RMS to protect documents with high personally identifiable information (PII). Only members of the FinanceAdmin group will have access to documents that contain high PII.
Step 4: View the results	Examine the classification of documents and

Task	Description
	observe how they change as you change the content in the document. Also verify how the document gets protected by AD RMS.
Step 5: Verify protection with AD RMS	Verify that the document is protected with AD RMS.

Step 1: Enable resource properties

To enable resource properties

1. In Hyper-V Manager, connect to server ID_AD_DC1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. Open Active Directory Administrative Center, and click **Tree View**.
3. Expand **DYNAMIC ACCESS CONTROL**, and select **Resource Properties**.
4. Scroll down to the **Impact** property in the **Display name** column. Right-click **Impact**, and then click **Enable**.
5. Scroll down to the **Personally Identifiable Information** property in the **Display name** column. Right-click **Personally Identifiable Information**, and then click **Enable**.
6. To publish the resource properties in the **Global Resource List**, in the left pane, click **Resource Property Lists**, and then double-click **Global Resource Property List**.
7. Click **Add**, and then scroll down to and click **Impact** to add it to the list. Do the same for **Personally Identifiable Information**. Click **OK** twice to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Impact_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"

Set-ADResourceProperty -Enabled:$true -Identity:"CN=PII_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Create classification rules

This step explains how to create the **High Impact** classification rule. This rule will search the content of documents and if the string “Contoso Confidential” is found, it will classify this document as having high-business impact. This classification will override any previously assigned classification of low-business impact.

You will also create a **High PII** rule. This rule searches the content of documents, and if a Social Security number is found, it classifies the document as having high PII.

To create the high-impact classification rule

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. You need to refresh the Global Resource Properties from Active Directory. Open Windows PowerShell and type: `Update-FSRMClassificationPropertyDefinition`, and then press ENTER. Close Windows PowerShell.
3. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
4. In the left pane of File Server Resource Manager, expand **Classification Management**, and then select **Classification Rules**.
5. In the **Actions** pane, click **Configure Classification Schedule**. On the **Automatic Classification** tab, select **Enable fixed schedule**, select a **Day of the week**, and then select the **Allow continuous classification for new files** check box. Click **OK**.
6. In the **Actions** pane, click **Create Classification Rule**. This opens the **Create Classification Rule** dialog box.
7. In the **Rule name** box, type **High Business Impact**.
8. In the **Description** box, type **Determines if the document has a high business impact based on the presence of the string "Contoso Confidential"**
9. On the **Scope** tab, click **Set Folder Management Properties**, select **Folder Usage**, click **Add**, then click **Browse**, browse to D:\Finance Documents as the path, click **OK**, and then choose a property value named **Group Files** and click **Close**. Once management properties are set, on the **Rule Scope** tab select **Group Files**.
10. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list.
11. Under **Choose a property to assign to files**, select **Impact** from the drop-down list.
12. Under **Specify a value**, select **High** from the drop-down list.
13. Click **Configure** under **Parameters**. In the **Classification Parameters** dialog box, in the **Expression Type** list, select **String**. In the **Expression** box, type: **Contoso Confidential**, and then click **OK**.
14. Click the **Evaluation Type** tab. Click **Re-evaluate existing property values**, click **Overwrite** the existing value, and then click **OK** to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Update-FSRMClassificationPropertyDefinition

$date = Get-Date
```

```
$AutomaticClassificationScheduledTask = New-FsrmScheduledTask -Time $date -Weekly @(3, 2, 4, 5,1,6,0) -RunDuration 0;

Set-FsrmClassification -Continuous -schedule $AutomaticClassificationScheduledTask

New-FSRMClassificationRule -Name "High Business Impact" -Property "Impact_MS" -Description "Determines if the document has a high business impact based on the presence of the string 'Contoso Confidential'" -PropertyValue "3000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("StringEx=Min=1;Expr=Contoso Confidential") -ReevaluateProperty Overwrite
```

To create the high-PII classification rule

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. On the desktop, open the folder named **Regular Expressions**, and then open the text document named **RegEx-SSN**. Highlight and copy the following regular expression string: **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**. This string will be used later in this step so keep it on your clipboard.
3. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
4. In the left pane of File Server Resource Manager, expand **Classification Management**, and then select **Classification Rules**.
5. In the **Actions** pane, click **Configure Classification Schedule**. On the **Automatic Classification** tab, select **Enable fixed schedule**, select a **Day of the week**, and then select the **Allow continuous classification for new files** check box. Click OK.
6. In the **Rule name** box, type **High PII**. In the **Description** box, type **Determines if the document has a high PII based on the presence of a Social Security Number.**
7. Click the **Scope** tab, select the **Group Files** check box.
8. Click the **Classification** tab. Under **Choose a method to assign the property to files**, select **Content Classifier** from the drop-down list.
9. Under **Choose a property to assign to files**, select **Personally Identifiable Information** from the drop-down list.
10. Under **Specify a value**, select **High** from the drop-down list.
11. Click **Configure** under **Parameters**.
    In the **Classification Parameters** window, in the **Expression Type** list, select **Regular Expression**. In the **Expression** box, paste the text from your clipboard: **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**, and then click **OK**.

    Note
    This expression will allow invalid Social Security numbers. This allows us to use fictitious Social Security numbers in the demonstration.
12. Click the **Evaluation Type** tab. Select **Re-evaluate existing property values**,

Overwrite the existing value, and then click **OK** to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-FSRMClassificationRule -Name "High PII" -Description "Determines if the document has a high PII based on the presence of a Social Security Number." -Property "PII_MS" -PropertyValue "5000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("RegularExpressionEx=Min=1;Expr=^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -]?)(?!00)\d\d\3(?!0000)\d{4}$") -ReevaluateProperty Overwrite
```

You should now have two classification rules:

- High Business Impact
- High PII

Step 3: Use file management tasks to automatically protect documents with AD RMS

Now that you've created rules to automatically classify documents based on content, the next step is to create a file management task that uses AD RMS to automatically protect certain documents based on their classification. In this step, you will create a file management task that automatically protects any documents with a high PII. Only members of the FinanceAdmin group will have access to documents that contain high PII.

To protect documents with AD RMS

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
3. In the left pane, select **File Management Tasks**. In the **Actions** pane, select **Create File Management Task**.
4. In the **Task name:** field, type **High PII**. In the **Description** field, type **Automatic RMS protection for high PII documents**.
5. Click the **Scope** tab, select the **Group Files** check box.
6. Click the **Action** tab. Under **Type**, select **RMS Encryption**. Click **Browse** to select a template, and then select the **Contoso Finance Admin Only** template.
7. Click the **Condition** tab, and then click **Add**. Under **Property**, select **Personally Identifiable Information**. Under **Operator**, select **Equal**. Under **Value**, select **High**. Click **OK**.

8. Click the **Schedule** tab. In the **Schedule** section, click **Weekly**, and then select **Sunday**. Running the task once-a-week will ensure that you catch any documents that may have been missed due to a service outage or other disruptive event.
9. In the **Continuous operation** section, select **Run task continuously on new files**, and then click **OK**. You should now have a file management task named High PII.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$fmjRmsEncryption = New-FSRMFmjAction -Type 'Rms' -RmsTemplate 'Contoso Finance Admin Only'

$fmjCondition1 = New-FSRMFmjCondition -Property 'PII_MS' -Condition 'Equal' -Value '5000'

$date = get-date

$schedule = New-FsrmScheduledTask -Time $date -Weekly @('Sunday')

$fmj1=New-FSRMFileManagementJob -Name "High PII" -Description "Automatic RMS protection for high PII documents" -Namespace @('D:\Finance Documents') -Action $fmjRmsEncryption -Schedule $schedule -Continuous -Condition @($fmjCondition1)
```

Step 4: View the results

It's time to take a look at your new automatic classification and AD RMS protection rules in action. In this step you will examine the classification of documents and observe how they change as you change the content in the document.

To view the results

1. In Hyper-V Manager, connect to server ID_AD_FILE1. Sign in to the server by using Contoso\Administrator with the password **pass@word1**.
2. In Windows Explorer, navigate to D:\Finance Documents.
3. Right-click the Finance Memo document and click **Properties**.Click the **Classification** tab, and notice that the Impact property currently has no value. Click **Cancel**.
4. Right-click the **Request for Approval to Hire document**, and then select **Properties**.
5. Click the **Classification** tab, and notice that **the Personally Identifiable Information property** currently has no value. Click **Cancel**.
6. Switch to CLIENT1. Sign off any user who is signed in, and then sign in as Contoso\MReid with the password **pass@word1**.
7. From the Desktop, open the **Finance Documents** shared folder.
8. Open the **Finance Memo** document. Near the bottom of the document, you will see the word **Confidential**. Modify it to read: **Contoso Confidential**. Save the document and close it.
9. Open the **Request for Approval to Hire** document. In the **Social Security#:** section,

type: 777-77-7777. Save the document and close it.

Note
You may need to wait 30 seconds for the classification to occur.

10. Switch back to ID_AD_FILE1. In Windows Explorer, navigate to D:\Finance Documents.
11. Right-click the Finance Memo document, and click **Properties**. Click the **Classification** tab. Notice that the **Impact** property is now set to **High**. Click **Cancel**.
12. Right-click the Request for Approval to Hire document and click **Properties**.
13. . Click the **Classification** tab. Notice that the **Personally Identifiable Information** property is now set to **High**. Click **Cancel**.

Step 5: Verify protection with AD RMS

To verify that the document is protected

1. Switch back to ID_AD_CLIENT1.
2. Open the **Request for approval to Hire** document.
3. Click **OK** to allow the document to connect to your AD RMS server.
4. You can now see that the document has been protected by AD RMS because it contains a Social Security number.

Scenario: Get Insight into Your Data by Using Classification

Reliance on data and storage resources has continued to grow in importance for most organizations. IT administrators face the growing challenge of overseeing larger and more complex storage infrastructures, while simultaneously being tasked with the responsibility to ensure that total cost-of-ownership is maintained at reasonable levels. Managing storage resources is not only about the volume or availability of data; it is also about enforcing company policies and knowing how storage is consumed to enable efficient utilization and compliance to mitigate risk. File Classification Infrastructure provides insight into your data by automating classification processes so that you can manage your data more effectively. The following classification methods are available with File Classification Infrastructure: manual, programmatic, and automatic. This topic focuses on the automatic file classification method.

Scenario description

File Classification Infrastructure uses classification rules to automatically scan files and classify them according to the contents of the file. Classification properties are defined centrally in Active

Directory so that these definitions can be shared across file servers in the organization. You can create classification rules that scan files for a standard string or for a string that matches a pattern (regular expression). When a configured classification parameter is found in a file, that file is classified as configured in the classification rule. Some examples of classification rules include:

- Classify any file that contains the string “Contoso Confidential” as having high business impact
- Classify any file that contains at least 10 social security numbers as having personally identifiable information

When a file is classified, you can use a file management task to take action on any files that are classified a specific way. The actions in a file management task include protecting the rights associated with the file, expiring the file, and running a custom action (such as posting information to a web service).

You can find planning information for configuring automatic file classification in Plan for Automatic File Classification.

You can find steps for how to automatically classify files in Deploy Automatic File Classification (Demonstration Steps).

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Practical applications

File Classification Infrastructure in Windows Server 2012 contributes to Dynamic Access Control by enabling business data owners to easily classify and label data. The classification information that is stored in the central access policy allows you to define access policies for data classes that are critical to business.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	File Classification Infrastructure is a feature that is included in File Server Resource Manager.
File and Storage Services Overview	File Server Resource Manager is a feature that is included with the File Services server role.

Deploy Automatic File Classification (Demonstration Steps)

This topic explains how to enable resource properties in Active Directory, create classification rules on the file server, and then assign values to the resource properties for files on the file server. For this example, the following classification rules are created:

- A content classification rule that searches a set of files for the string “Contoso Confidential.” If the string is found in a file, the Impact resource property is set to High on the file.
- A content classification rule that searches a set of files for a regular expression that matches a social security number at least 10 times in one file. If the pattern is found, the file is classified as having personally identifiable information and the Personally Identifiable Information resource property is set to High.

In this document

- **Step 1: Create resource property definitions**
- **Step 2: Create a string content classification rule**
- **Step 3: Create a regular expression content classification rule**
- Step 4: Verify that the files are classified correctly

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Create resource property definitions

The Impact and Personally Identifiable Information resource properties are enabled so that File Classification Infrastructure can use these resource properties to tag the files that are scanned on a network shared folder.

Do this step using Windows PowerShell

To create resource property definitions

1. On the domain controller, sign in to the server as a member of the Domain Admins security group.
2. Open Active Directory Administrative Center. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Expand **Dynamic Access Control**, and then click **Resource Properties**.
4. Right-click **Impact**, and then click **Enable**.
5. Right-click **Personally Identifiable Information**, and then click **Enable**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Impact_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=PII_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Create a string content classification rule

A string content classification rule scans a file for a specific string. If the string is found, the value of a resource property can be configured. In this example, we will scan each file on a network shared folder and look for the string "Contoso Confidential." If the string is found, the associated file is classified as having high business impact.

Do this step using Windows PowerShell

To create a string content classification rule

1. Log on to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition** and then press ENTER. This will synchronize the property definitions created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Expand **Classification Management**, right-click **Classification Rules**, and then click **Configure Classification Schedule**.
5. Select the **Enable fixed schedule** check box, select the **Allow continuous classification for new files** check box, choose a day of the week to run the classification, and then click **OK**.
6. Right-click **Classification Rules**, and then click **Create Classification Rule**.
7. On the **General** tab, in the **Rule name** box, type a rule name such as **Contoso Confidential**.
8. On the **Scope** tab, click **Add**, and choose the folders that should be included in this rule, such as D:\Finance Documents.

   Note
   You can also choose a dynamic name space for the scope. For more information about dynamic name spaces for classification rules, see What's New in File Server Resource Manager in Windows Server 2012.
9. On the **Classification** tab, configure the following:
   - In the **Choose a method to assign a property to files** box, ensure that **Content Classifier** is selected.

- In the **Choose a property to assign to files** box, click **Impact**.
- In the **Specify a value** box, click **High**.

10. Under the **Parameters** heading, click **Configure**.
11. In the **Expression Type** column, select **String**.
12. In the **Expression** column, type **Contoso Confidential**, and then click **OK**.
13. On the **Evaluation Type** tab, select the **Re-evaluate existing property values** check box, click **Overwrite the existing value**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$date = Get-Date

$AutomaticClassificationScheduledTask = New-FsrmScheduledTask -Time $date -Weekly @(3, 2, 4, 5,1,6,0) -RunDuration 0;$AutomaticClassificationScheduledTask

Set-FsrmClassification -Continuous -schedule $AutomaticClassificationScheduledTask

New-FSRMClassificationRule -Name "Contoso Confidential" -Property "Impact_MS" -PropertyValue "3000" -Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -Parameters @("StringEx=Min=1;Expr=Contoso Confidential") -ReevaluateProperty Overwrite
```

Step 3: Create a regular expression content classification rule

A regular expression classification rule scans a file for a pattern that matches the regular expression. If a string that matches the regular expression is found, the value of a resource property can be configured. In this example, we will scan each file on a network shared folder and look for a string that matches the pattern of a social security number (XXX-XX-XXXX). If the pattern is found, the associated file is classified as having personally identifiable information.

Do this step using Windows PowerShell

To create a regular expression content classification rule

1. Sign in to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition**, and then press ENTER. This will synchronize the property definitions that are created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Right-click **Classification Rules**, and then click **Create Classification Rule**.
5. On the **General** tab, in the **Rule name** box, type a name for the classification rule, such

as PII Rule.

6. On the **Scope** tab, click **Add**, and then choose the folders that should be included in this rule, such as D:\Finance Documents.
7. On the **Classification** tab, configure the following:
   - In the **Choose a method to assign a property to files** box, ensure that **Content Classifier** is selected.
   - In the **Choose a property to assign to files** box, click **Personally Identifiable Information**.
   - In the **Specify a value** box, click **High**.
8. Under the **Parameters** heading, click **Configure**.
9. In the **Expression Type** column, select **Regular expression**.
10. In the **Expression** column, type **^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$**
11. In the **Minimum Occurrences** column, type **10**, and then click **OK**.
12. On the **Evaluation Type** tab, select the **Re-evaluate existing property values** check box, click **Overwrite the existing value**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-FSRMClassificationRule -Name "PII Rule" -Property "PII_MS" -PropertyValue "5000" -
Namespace @("D:\Finance Documents") -ClassificationMechanism "Content Classifier" -
Parameters @("RegularExpressionEx=Min=10;Expr=^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([ -
]?)(?!00)\d\d\3(?!0000)\d{4}$") -ReevaluateProperty Overwrite
```

Step 4: Verify that the files are classified correctly

You can verify that the files are properly classified by viewing the properties of a file that was created in the folder specified in the classification rules.

To verify that the files are classified correctly

1. On the file server, run the classification rules by using File Server Resource Manager.
   a. Click **Classification Management**, right-click **Classification Rules**, and then click **Run Classification With All Rules Now**.
   b. Click the **Wait for classification to complete** option, and then click **OK**.
   c. Close the Automatic Classification Report.
   d. You can do this by using Windows PowerShell with the following command: **Start-FSRMClassification –RunDuration 0 -Confirm:$false**
2. Navigate to the folder that was specified in the classification rules, such as D:\Finance Documents.

3. Right-click a file in that folder, and then click **Properties**.
4. Click the **Classification** tab, and verify that the file is classified correctly.

See also

- Scenario: Get Insight into Your Data by Using Classification
- Plan for Automatic File Classification
- Dynamic Access Control: Scenario Overview

Scenario: Implement Retention of Information on File Servers

A retention period is the amount of time that a document should be kept before it is expired. Depending on the organization, the retention period can be different. You can classify files in a folder as having a short, medium, or long-term retention period, and then assign a timeframe for each period. You may want to keep a file indefinitely by putting it on legal hold.

Scenario description

File Classification Infrastructure and File Server Resource Manager uses file management tasks and file classification to apply retention periods for a set of files. You can assign a retention period on a folder and then use a file management task to configure how long an assigned retention period is to last. When the files in the folder are about to expire, the owner of the file receives a notification email. You can also classify a file as being on legal hold so that the file management task will not expire the file.

You can find planning information for configuring retention in Plan for Retention of Information on File Servers.

You can find steps for classifying files for legal hold and configuring a retention period in Deploy Implementing Retention of Information on File Servers (Demonstration Steps).

Note
That scenario only discusses how to manually classify a document for legal hold. However, it is possible in Windows Server 2012 to automatically classify documents for legal hold. One way to do this is to create a Windows PowerShell classifier that compares the file owner to a list of user accounts that are under legal hold. If the file owner is a part of the user account list, the file is classified for legal hold.

In this scenario

This scenario is part of the Dynamic Access Control scenario. For additional information about Dynamic Access Control, see:

- Dynamic Access Control: Scenario Overview

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
File Server Resource Manager Overview	File Management Tasks is a feature that is included in File Server Resource Manager.
File and Storage Services Overview	File Server Resource Manager is a feature that is included with the File Services server role.

Deploy Implementing Retention of Information on File Servers (Demonstration Steps)

You can set retention periods for folders and put files on legal hold by using File Classification Infrastructure and File Server Resource Manager.

In this document

- **Prerequisites**
- **Step 1: Create resource property definitions**
- Step 2: Configure notifications
- **Step 3: Create a file management task**
- Step 4: Classify a file manually

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

The steps in this topic assume you have a SMTP server configured for file expiration notifications.

Step 1: Create resource property definitions

In this step, we enable the Retention Period and Discoverability resource properties so that File Classification Infrastructure can use these resource properties to tag the files that are scanned in a network shared folder.

Do this step using Windows PowerShell

To create resource property definitions

1. On the domain controller, sign in to the server as a member of the Domain Admins security group.
2. Open Active Directory Administrative Center. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Expand **Dynamic Access Control**, and then click **Resource Properties**.
4. Right-click **Retention Period**, and then click **Enable**.
5. Right-click **Discoverability**, and then click **Enable**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADResourceProperty -Enabled:$true -Identity:"CN=RetentionPeriod_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
Set-ADResourceProperty -Enabled:$true -Identity:"CN=Discoverability_MS,CN=Resource
Properties,CN=Claims Configuration,CN=Services,CN=Configuration,DC=contoso,DC=com"
```

Step 2: Configure notifications

In this step, we use the File Server Resource Manager console to configure the SMTP server, the default administrator email address, and the default email address that the reports are sent from.

Do this step using Windows PowerShell

To configure notifications

1. Sign in to the file server as a member of the Administrators security group.
2. From the Windows PowerShell command prompt, type **Update-FsrmClassificationPropertyDefinition**, and then press ENTER. This will synchronize the property definitions that are created on the domain controller to the file server.
3. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
4. Right-click **File Server Resource Manager (local)**, and then click **Configure Options**.
5. On the **Email Notifications** tab, configure the following:
   - In the **SMTP server name or IP address** box, type the name of the SMTP server on

your network.

- In the **Default administrator recipients** box, type the email address of the administrator who should get the notification.
- In the **Default "From" e-mail address** box, type the email address that should be used to send the notifications.

6. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-FsrmSetting –SmtpServer IP address of SMTP server -FromEmailAddress
"FromEmailAddress" -AdminEmailAddress "AdministratorEmailAddress"
```

Step 3: Create a file management task

In this step, we use the File Server Resource Manager console to create a file management task that will run on the last day of the month and expire any files with the following criteria:

- The file is not classified as being on legal hold.
- The file is classified as having a long-term retention period.
- The file has not been modified in the last 10 years.

Do this step using Windows PowerShell

To create a file management task

1. Sign in to the file server as a member of the Administrators security group.
2. Open File Server Resource Manager. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
3. Right-click **File Management Tasks**, and then click **Create File Management Task**.
4. On the **General** tab, in the **Task name** box, type a name for the file management task, such as Retention Task.
5. On the **Scope** tab, click **Add**, and choose the folders that should be included in this rule, such as D:\Finance Documents.
6. On the **Action** tab, in the **Type** box, click **File expiration**. In the **Expiration directory** box, type a path to a folder on the local file server where the expired files will be moved. This folder should have an access control list that grants only file server administrators access.
7. On the **Notification** tab, click **Add**.
   - Select the **Send e-mail to the following administrators** check box.
   - Select the **Send an email to users with affected files** check box, and then click **OK**.
8. On the **Condition** tab, click **Add**, and add the following properties:
   - In the **Property** list, click **Discoverability**. In the **Operator** list, click **Not equal**. In the

Value list, click **Hold**.

- In the **Property** list, click **Retention Period**. In the **Operator** list, click **Equal**. In the **Value** list, click **Long-Term**.

9. On the **Condition** tab, select the **Days since file was last modified** check box, and then set the value to **3650**.
10. On the **Schedule** tab, click the **Monthly** option, and then select the **Last** check box.
11. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$fmjexpiration = New-FSRMFmjAction -Type 'Expiration' -ExpirationFolder folder

$fmjNotificationAction = New-FsrmFmjNotificationAction -Type Email -MailTo
"[FileOwner],[AdminEmail]"

$fmjNotification = New-FsrmFMJNotification -Days 10 -Action @($fmjNotificationAction)

$fmjCondition1 = New-FSRMFmjCondition -Property 'Discoverability_MS' -Condition
'NotEqual' -Value "Hold"

$fmjCondition2 = New-FSRMFmjCondition -Property 'RetentionPeriod_MS' -Condition 'Equal' -
Value "Long-term"

$fmjCondition3 = New-FSRMFmjCondition -Property 'File.DateLastAccessed' -Condition
'Equal' -Value 3650

$date = get-date
$schedule = New-FsrmScheduledTask -Time $date -Monthly @(-1)

$fmj1=New-FSRMFileManagementJob -Name "Retention Task" -Namespace @('D:\Finance
Documents') -Action $fmjexpiration -Schedule $schedule -Notification @($fmjNotification)
-Condition @( $fmjCondition1, $fmjCondition2, $fmjCondition3)
```

Step 4: Classify a file manually

In this step, we manually classify a file to be on legal hold. The parent folder of this file will be classified with a long-term retention period.

To manually classify a file

1. Sign in to the file server as a member of the Administrators security group.
2. Navigate to the folder that was configured in the scope of the file management task created in Step 3.
3. Right-click the folder, and then click **Properties**.
4. On the **Classification** tab, click **Retention Period**, click **Long-Term**, and then click **OK**.

5. Right-click a file within that folder, and then click **Properties**.
6. On the **Classification** tab, click **Discoverability**, click **Hold**, click **Apply**, and then click **OK**.
7. On the file server, run the file management task by using the File Server Resource Manager console. After the file management task completes, check the folder and ensure the file was not moved to the expiration directory.
8. Right-click the same file within that folder, and then click **Properties**.
9. On the **Classification** tab, click **Discoverability**, click **Not Applicable**, click **Apply**, and then click **OK**.
10. On the file server, run the file management task again by using the File Server Resource Manager console. After the file management task completes, check the folder and ensure that file was moved to the expiration directory.

See also

- Scenario: Implement Retention of Information on File Servers
- Plan for Retention of Information on File Servers
- Dynamic Access Control: Scenario Overview

Deploy Claims Across Forests

In Windows Server 2012, a claim type is an assertion about the object with which it's associated. Claim types are defined per forest in Active Directory. There are many scenarios where a security principal may need to traverse a trust boundary to access resources in a trusted forest. Cross-forest claims transformation in Windows Server 2012 enables you to transform egress and ingress claims that traverse forests so that the claims are recognized and accepted in the trusting and trusted forests. Some of the real-world scenarios for transformation of claims are:

- Trusting forests can use claim transformation as a guard against elevation of privilege by filtering the incoming claims with specific values.

  Trusting forests can also issue claims for principals coming over a trust boundary if the trusted forest does not support or issue any claims.
- Trusted forests can use claim transformation to prevent certain claim types and claims with certain values from going out to the trusting forest.
- You can also use claim transformation to map different claim types between trusting and trusted forests. This can be used to generalize the claim-type, the claim value, or both. Without this, you need to standardize the data between the forests before you can use the claims. Generalizing claims between the trusting and trusted forests reduces the IT costs.

Claim transformation rules

The transformation rule language syntax divides a single rule into two main parts: a series of condition statements and the issue statement. Each condition statement has two subcomponents: the claim identifier and the condition. The issue statement contains keywords, delimiters, and an issue expression. The condition statement optionally begins with a claim identifier variable, which represents the matched input claim. The condition checks for the expression. If the input claim does not match the condition, then the transformation engine ignores the issue statement and evaluates the next input claim against the transformation rule. If all conditions match the input claim, it processes the issue statement.

For detailed information on claim rules language, see Claims Transformation Rules Language.

Linking claim transformation policies to forests

There are two components involved in setting up claim transformation policies: claim transformation policy objects and the transformation link. The policy objects live in the configuration naming context in a forest, and they contain mapping information for the claims. The link specifies which trusting and trusted forests the mapping applies to.

It is important to understand if the forest is the trusting or trusted forest because this is basis for linking transformation policy objects. For example, the trusted forest is the forest that contains user accounts that require access. The trusting forest is the forest that contains resources that you want to give users access to. Claims travel in the same direction as the security principal that requires access. For example, if there is a one-way trust from the contoso.com forest to the adatum.com forest, the claims will flow from adatum.com to contoso.com, which allows users from adatum.com to access resources in contoso.com.

By default, a trusted forest allows all outgoing claims to pass, and a trusting forest drops all incoming claims that it receives.

In this scenario

The following guidance is available for this scenario:

- Deploy Claims Across Forests (Demonstration Steps)
- Claims Transformation Rules Language

Roles and features included in this scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services	In this scenario, you are required to set up two Active Directory forests with a two-way trust.

Role/feature	How it supports this scenario
	You have claims in both forests. You also set central access policies on the trusting forest where the resources reside.
File and Storage Services role	In this scenario, the data classification is applied to the resources on the file servers. The central access policy is applied to the folder where you want to grant user access. After transformation, the claim grants user access to resources based on the central access policy that is applied to the folder on the file server.

Deploy Claims Across Forests (Demonstration Steps)

In this topic, we'll cover a basic scenario that explains how to configure claims transformations between trusting and trusted forests. You will learn how claims transformation policy objects can be created and linked to the trust on the trusting forest and the trusted forest. You will then validate the scenario.

Scenario overview

Adatum Corporation provides financial services to Contoso, Ltd. Each quarter, Adatum accountants copy their account spreadsheets to a folder on a file server located at Contoso, Ltd. There is a two-way trust set up from Contoso to Adatum. Contoso, Ltd. wants to protect the share so that only Adatum employees can access the remote share.

In this scenario:

1. Set up the prerequisites and the test environment
2. Set up claims transformation on trusted forest (Adatum)
3. Set up claims transformation in the trusting forest (Contoso)
4. Validate the scenario

Set up the prerequisites and the test environment

The test configuration involves setting up two forests: Adatum Corporation and Contoso, Ltd, and having a two-way trust between Contoso and Adatum. "adatum.com" is the trusted forest and "contoso.com" is the trusting forest.

The claims transformation scenario demonstrates transformation of a claim in the trusted forest to a claim in the trusting forest. To do this, you need to set up a new forest called adatum.com and populate the forest with a test user with a company value of “Adatum”. You then have to set up a two-way trust between contoso.com and adatum.com.

Important
When setting up the Contoso and Adatum forests, you must ensure that both the root domains are at the Windows Server 2012 Domain Functional Level for claims transformation to work.

You need to set up the following for the lab. These procedures are explained in detail in Appendix B: Setting Up the Test Environment

You need to implement the following procedures to set up the lab for this scenario:

1. **Set Adatum as trusted forest to Contoso**
2. Create the “Company” claim type on Contoso
3. Enable the “Company” resource property on Contoso
4. Create the central access rule
5. Create the central access policy
6. Publish the new policy through Group Policy
7. Create the Earnings folder on the file server
8. Set classification and apply the central access policy on the new folder

Use the following information to complete this scenario:

Objects	Details
Users	Jeff Low, Contoso
User claims on Adatum and Contoso	ID: ad://ext/Company:ContosoAdatum, Source attribute: company Suggested values: Contoso, Adatum **Important** You must set the ID on the “Company” claim type on both Contoso and Adatum to be the same for the claims transformation to work.
Central access rule on Contoso	AdatumEmployeeAccessRule
Central access policy on Contoso	Adatum Only Access Policy
Claims Transformation policies on Adatum and Contoso	DenyAllExcept Company
File folder on Contoso	D:\EARNINGS

Set up claims transformation on trusted forest (Adatum)

In this step you create a transformation policy in Adatum to deny all claims except “Company” to pass to Contoso.

The Active Directory module for Windows PowerShell provides the **DenyAllExcept** argument, which drops everything except the specified claims in the transformation policy.

To set up a claims transformation, you need to create a claims transformation policy and link it between the trusted and trusting forests.

Create a claims transformation policy in Adatum

To create a transformation policy Adatum to deny all claims except “Company”

1. Sign in to the domain controller, adatum.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell, and type the following:

```
New-ADClaimTransformPolicy `
-Description:"Claims transformation policy to deny all claims except Company"`
-Name:"DenyAllClaimsExceptCompanyPolicy" `
-DenyAllExcept:company `
-Server:"adatum.com" `
```

Set a claims transformation link on Adatum’s trust domain object

In this step, you apply the newly created claims transformation policy on Adatum’s trust domain object for Contoso.

To apply the claims transformation policy

1. Sign in to the domain controller, adatum.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell, and type the following:

```
Set-ADClaimTransformLink `
-Identity:"contoso.com" `
```

```
-Policy:"DenyAllClaimsExceptCompanyPolicy" `
-TrustRole:Trusted `
```

Set up claims transformation in the trusting forest (Contoso)

In this step you create a claims transformation policy in Contoso (the trusting forest) to deny all claims except “Company.” You need to create a claims transformation policy and link it to the forest trust.

Create a claims transformation policy in Contoso

To create a transformation policy Adatum to deny all except “Company”

1. Sign in to the domain controller, contoso.com as Administrator with the password **pass@word1**.
2. Open an elevated command prompt in Windows PowerShell and type the following:

```
New-ADClaimTransformPolicy `
-Description:"Claims transformation policy to deny all claims
except company" `
-Name:"DenyAllClaimsExceptCompanyPolicy" `
-DenyAllExcept:company `
-Server:"contoso.com" `
```

Set a claims transformation link on Contoso’s trust domain object

In this step, you apply the newly created claims transformation policy on the contoso.com trust domain object for Adatum to allow “Company” be passed through to contoso.com. The trust domain object is named adatum.com.

To set the claims transformation policy

1. Sign in to the domain controller, contoso.com as Administrator with the password **pass@word1**.

2. Open an elevated command prompt in Windows PowerShell and type the following:

```
Set-ADClaimTransformLink
-Identity:"adatum.com" `
-Policy:"DenyAllClaimsExceptCompanyPolicy" `
-TrustRole:Trusting `
```

Validate the scenario

In this step you try to access the D:\EARNINGS folder that was set up on the file server FILE1 to validate that the user has access to the shared folder.

To ensure that the Adatum user can access the shared folder

1. Sign in to the Client machine, CLIENT1 as Jeff Low with the password **pass@word1**.
2. Browse to the folder \\FILE1.contoso.com\Earnings.
3. Jeff Low should be able to access the folder.

Additional scenarios for claims transformation policies

Following is a list of additional common cases in claims transformation.

Scenario	Policy

Allow all claims that come from Adatum to go through to Contoso Adatum	`New-ADClaimTransformPolicy ` `-Description:"Claims transformation policy to allow all claims" ` `-Name:"AllowAllClaimsPolicy" ` `-AllowAll ` `-Server:"contoso.com" `
	
	`Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"AllowAllClaimsPolicy" ` `-TrustRole:Trusting `

Scenario	Policy
	`-Server:"contoso.com" `
Deny all claims that come from Adatum to go through to Contoso Adatum	`New-ADClaimTransformPolicy ` `-Description:"Claims transformation policy to deny all claims" ` `-Name:"DenyAllClaimsPolicy" ` `-DenyAll ` `-Server:"contoso.com" ` `Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"DenyAllClaimsPolicy" ` `-TrustRole:Trusting ` `-Server:"contoso.com"`
Allow all claims that come from Adatum except “Company” and “Department” to go through to Contoso Adatum	`New-ADClaimTransformationPolicy ` `-Description:"Claims transformation policy to allow all claims except company and department" ` `-Name:"AllowAllClaimsExceptCompanyAndDepartmentPolicy" ` `-AllowAllExcept:company,department ` `-Server:"contoso.com" ` `Set-ADClaimTransformLink ` `-Identity:"adatum.com" ` `-Policy:"AllowAllClaimsExceptCompanyAndDepartmentPolicy" ` `-TrustRole:Trusting ` `-Server:"contoso.com" `

See also

- For a list of all Windows PowerShell cmdlets that are available for claims transformation, see Active Directory PowerShell Cmdlet Reference.
- For advanced tasks that involve export and import of DAC configuration information between two forests, use the Dynamic Access Control PowerShell Reference
- Deploy Claims Across Forests
- Claims Transformation Rules Language
- Dynamic Access Control: Scenario Overview

Claims Transformation Rules Language

The across-forest claims transformation feature enables you to bridge claims for Dynamic Access Control across forest boundaries by setting claims transformation policies on across-forest trusts. The primary component of all policies is rules that are written in claims transformation rules language. This topic provides details about this language and provides guidance about authoring claims transformation rules.

The Windows PowerShell cmdlets for transformation policies on across-forest trusts have options to set simple policies that are required in common scenarios. These cmdlets translate the user input into policies and rules in the claims transformation rules language, and then store them in Active Directory in the prescribed format. For more information about cmdlets for claims transformation, see the AD DS Cmdlets for Dynamic Access Control.

Depending on the claims configuration and the requirements placed on the across-forest trust in your Active Directory forests, your claims transformation policies may have to be more complex than the policies supported by the Windows PowerShell cmdlets for Active Directory. To effectively author such policies, it is essential to understand the claims transformation rules language syntax and semantics. This claims transformation rules language (“the language”) in Active Directory is a subset of the language that is used by Active Directory Federation Services for similar purposes, and it has a very similar syntax and semantics. However, there are fewer operations allowed, and additional syntax restrictions are placed in the Active Directory version of the language.

This topic briefly explains the syntax and semantics of the claims transformation rules language in Active Directory and considerations to be made when authoring policies. It provides several sets of example rules to get you started, and examples of incorrect syntax and the messages they generate, to help you decipher error messages when you author the rules.

Tools for authoring claims transformation policies

Windows PowerShell cmdlets for Active Directory: This is the preferred and recommended way to author and set claims transformation policies. These cmdlets provide switches for simple policies and verify rules that are set for more complex policies.

LDAP: Claims transformation policies can be edited in Active Directory through Lightweight Directory Access Protocol (LDAP). However, this is not recommended because the policies have several complex components, and the tools you use may not validate the policy before writing it to Active Directory. This may subsequently require a considerable amount of time to diagnose problems.

Active Directory claims transformation rules language

Syntax overview

Here is a brief overview of the syntax and semantics of the language:

- The claims transformation rule set consists of zero or more rules. Each rule has two active parts: **Select Condition List** and **Rule Action**. If the **Select Condition List** evaluates to TRUE, the corresponding rule action is executed.
- **Select Condition List** has zero or more **Select Conditions**. All of the **Select Conditions** must evaluate to TRUE for the **Select Condition List** to evaluate to TRUE.
- Each **Select Condition** has a set of zero or more **Matching Conditions**. All the **Matching Conditions** must evaluate to TRUE for the Select Condition to evaluate to TRUE. All of these conditions are evaluated against a single claim. A claim that matches a **Select Condition** can be tagged by an **Identifier** and referred to in the **Rule Action**.
- Each **Matching Condition** specifies the condition to match the **Type** or **Value** or **ValueType** of a claim by using different **Condition Operators** and **String Literals**.
    - When you specify a **Matching Condition** for a **Value**, you must also specify a **Matching Condition** for a specific **ValueType** and vice versa. These conditions must be next to each other in the syntax.
    - **ValueType** matching conditions must use specific **ValueType** literals only.
- A **Rule Action** can copy one claim that is tagged with an **Identifier** or issue one claim based on a claim that is tagged with an Identifier and/or given String Literals.

Example rule

This example shows a rule that can be used to translate the claims Type between two forests, provided that they use the same claims ValueTypes and have the same interpretations for claims Values for this type. The rule has one matching condition and an Issue statement that uses String Literals and a matching claims reference.

```
C1: [TYPE=="EmployeeType"]
                => ISSUE (TYPE= "EmpType", VALUE = C1.VALUE, VALUETYPE = C1.VALUETYPE);
```

`[TYPE=="EmployeeType"] == Select Condition List with one Matching Condition for claims Type.`

`ISSUE (TYPE= "EmpType", VALUE = C1.VALUE, VALUETYPE = C1.VALUETYPE) == Rule Action that issues a claims using string literal and matching claim referred with the Identifier.`

Runtime operation

It is important to understand the runtime operation of claims transformations to author the rules effectively. The runtime operation uses three sets of claims:

1. **Input claims set**: The input set of claims that are given to the claims transformation operation.
2. **Working claims set**: Intermediate claims that are read from and written to during the claims transformation.
3. **Output claims set**: Output of the claims transformation operation.

Here is a brief overview of the runtime claims transformation operation:

1. Input claims for claims transformation are used to initialize the working claims set.
    a. When processing each rule, the working claims set is used for the input claims.
    b. The Selection Condition List in a rule is matched against all possible sets of claims from the working claims set.
    c. Each set of matching claims is used to run the action in that rule.
    d. Running a rule action results in one claim, which is appended to the output claims set and the working claims set. Thus, the output from a rule is used as input for subsequent rules in the rule set.
2. The rules in the rule set are processed in sequential order starting with the first rule.
3. When the entire rule set is processed, the output claims set is processed to remove duplicate claims and for other security issues. The resulting claims are the output of the claims transformation process.

It is possible to write complex claims transformations based on the previous runtime behavior.

Example: Runtime operation

This example shows the runtime operation of a claims transformation that uses two rules.

```
     C1:[Type==”EmpType”, Value==”FullTime”,ValueType==”string”] =>
                Issue(Type==”EmployeeType”, Value==”FullTime”,ValueType==”string”);
     [Type==”EmployeeType”] =>
               Issue(Type==”AccessType”, Value==”Privileged”, ValueType==”string”);
Input claims and Initial Evaluation Context:
  {(Type= “EmpType”),(Value=”FullTime”),(ValueType=”String”)}
{(Type= “Organization”),(Value=”Marketing”),(ValueType=”String”)}
After Processing Rule 1:
 Evaluation Context:
```

```
  {(Type= "EmpType"),(Value="FullTime"),(ValueType="String")}
{(Type= "Organization"), (Value="Marketing"),(ValueType="String")}
  {(Type= "EmployeeType"),(Value="FullTime"),(ValueType="String")}
Output Context:
  {(Type= "EmployeeType"),(Value="FullTime"),(ValueType="String")}

After Processing Rule 2:
Evaluation Context:
  {(Type= "EmpType"),(Value="FullTime"),(ValueType="String")}
{(Type= "Organization"),(Value="Marketing"),(ValueType="String")}
  {(Type= "EmployeeType"),(Value="FullTime"),(ValueType="String")}
  {(Type= "AccessType"),(Value="Privileged"),(ValueType="String")}
Output Context:
  {(Type= "EmployeeType"),(Value="FullTime"),(ValueType="String")}
  {(Type= "AccessType"),(Value="Privileged"),(ValueType="String")}

Final Output:
  {(Type= "EmployeeType"),(Value="FullTime"),(ValueType="String")}
  {(Type= "AccessType"),(Value="Privileged"),(ValueType="String")}
```

Special rules semantics

The following are special syntax for rules:

1. Empty Rule Set == No Output Claims
2. Empty Select Condition List == Every Claim matches the Select Condition List

   Example: Empty Select Condition List

   The following rule matches every claim in the working set.

   ```
   => Issue (Type = "UserType", Value = "External", ValueType =
   "string")
   ```
3. Empty Select Matching List == Every claim matches the Select Condition List

   Example: Empty Matching Conditions

   The following rule matches every claim in the working set. This is the basic "Allow-all" rule if it is used alone.

```
C1:[] => Issule (claim = C1);
```

Security considerations

Claims that enter a forest

The claims presented by principals that are incoming to a forest need to be inspected thoroughly to ensure that we allow or issue only the correct claims. Improper claims can compromise the forest security, and this should be a top consideration when authoring transformation policies for claims that enter a forest.

Active Directory has the following features to prevent misconfiguration of claims that enter a forest:

- If a forest trust has no claims transformation policy set for the claims that enter a forest, for security purposes, Active Directory drops all the principal claims that enter the forest.
- If running the rule set on claims that enters a forest results in claims that are not defined in the forest, the undefined claims are dropped from the output claims.

Claims that leave a forest

Claims that leave a forest present a lesser security concern for the forest than the claims that enter the forest. Claims are allowed to leave the forest as-is even when there is no corresponding claims transformation policy in place. It is also possible to issue claims that are not defined in the forest as part of transforming claims that leave the forest. This is to easily set up across-forest trusts with claims. An administrator can determine if claims that enter the forest need to be transformed, and set up the appropriate policy. For example, an administrator could set a policy if there is a need to hide a claim to prevent information disclosure.

Syntax errors in claims transformation rules

If a given claims transformation policy has a rules set that is syntactically incorrect or if there are other syntax or storage issues, the policy is considered invalid. This is treated differently than the default conditions mentioned earlier.

Active Directory is unable to determine the intent in this case and goes into a fail-safe mode, where no output claims are generated on that trust+direction of traversal. Administrator intervention is required to correct the issue. This could happen if LDAP is used to edit the claims transformation policy. Windows PowerShell cmdlets for Active Directory have validation in place to prevent writing a policy with syntax issues.

Other language considerations

1. There are several key words or characters that are special in this language (referred to as terminals). These are presented in the Language terminals table later in this topic. The error messages use the tags for these terminals for disambiguation.
2. Terminals can sometimes be used as string literals. However, such usage may conflict with the language definition or have unintended consequences. This kind of usage is not recommended.

3. The rule action cannot perform any type conversions on claim Values, and a rule set that contains such a rule action is considered invalid. This would cause a runtime error, and no output claims are produced.
4. If a rule action refers to an Identifier that was not used in the Select Condition List portion of the rule, it is an invalid usage. This would cause a syntax error.

   Example: Incorrect Identifier reference
   The following rule illustrates an incorrect Identifier used in rule action.

```
C1:[] => Issue (claim = C2);
```

Sample transformation rules

- **Allow all claims of a certain type**

  Exact type

```
C1:[type==”XYZ”] => Issue (claim = C1);
```

  Using Regex

```
C1: [type =~ “XYZ*”] => Issue (claim = C1);
```

- **Disallow a certain claim type**

  Exact type

```
C1:[type != “XYZ”] => Issue (claim=C1);
```

  Using Regex

```
C1:[Type !~ “XYZ?”] => Issue (claim=C1);
```

Examples of rules parser errors

Claims transformation rules are parsed by a custom parser to check for syntax errors. This parser is run by related Windows PowerShell cmdlets before storing rules in Active Directory. Any errors in parsing the rules, including syntax errors, are printed on the console. Domain controllers also run the parser before using the rules for transforming claims, and they log errors in the event log (add event log numbers).

This section illustrates some examples of rules that are written with incorrect syntax and the corresponding syntax errors that are generated by the parser.

1. Example:

```
c1;[]=>Issue(claim=c1);
```

   This example has an incorrectly used semicolon in place of a colon.
   Error message:
   POLICY0002: Could not parse policy data.
   Line number: 1, Column number: 2, Error token: ;. Line: 'c1;[]=>Issue(claim=c1);'.

Parser error: 'POLICY0030: Syntax error, unexpected ';', expecting one of the following: ':' .'

2. Example:

```
c1:[]=>Issue(claim=c2);
```

In this example, the Identifier tag in the copy issuance statement is undefined.
Error message:
POLICY0011: No conditions in the claim rule match the condition tag specified in the CopyIssuanceStatement: 'c2'.

3. Example:

```
c1:[type=="x1", value=="1", valuetype=="bool"]=>Issue(claim=c1)
```

“bool” is not a Terminal in the language, and it is not a valid ValueType. Valid terminals are listed in the following error message.
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 39, Error token: "bool". Line: 'c1:[type=="x1", value=="1",valuetype=="bool"]=>Issue(claim=c1);'.
Parser error: 'POLICY0030: Syntax error, unexpected 'STRING', expecting one of the following: 'INT64_TYPE' 'UINT64_TYPE' 'STRING_TYPE' 'BOOLEAN_TYPE' 'IDENTIFIER'

4. Example:

```
c1:[type=="x1", value==1,
valuetype=="boolean"]=>Issue(claim=c1);
```

The numeral **1** in this example is not a valid token in the language, and such usage is not allowed in a matching condition. It has to be enclosed in double quotes to make it a string.
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 23, Error token: 1. Line: 'c1:[type=="x1", value==1, valuetype=="bool"]=>Issue(claim=c1);'.Parser error: 'POLICY0029: Unexpected input.

5. Example:

```
c1:[type == "x1", value == "1", valuetype == "boolean"] =>

      Issue(type = c1.type, value="0", valuetype == "boolean");
```

This example used a double equal sign (==) instead of a single equal sign (=).
Error message:
POLICY0002: Could not parse policy data.
Line number: 1, Column number: 91, Error token: ==. Line: 'c1:[type=="x1", value=="1", valuetype=="boolean"]=>Issue(type=c1.type, value="0", valuetype=="boolean");'.

Parser error: 'POLICY0030: Syntax error, unexpected '==', expecting one of the following: '='

6. Example:

```
c1:[type=="x1", value=="boolean", valuetype=="string"] =>

      Issue(type=c1.type, value=c1.value, valuetype = "string");
```

This example is syntactically and semantically correct. However, using “boolean” as a string value is bound to cause confusion, and it should be avoided. As previously mentioned, using language terminals as claims values should be avoided where possible.

Language terminals

The following table lists the complete set of terminal strings and the associated language terminals that are used in the claims transformation rules language. These definitions use case-insensitive UTF-16 strings.

String	Terminal
"=>"	IMPLY
";"	SEMICOLON
":"	COLON
","	COMMA
"."	DOT
"["	O_SQ_BRACKET
"]"	C_SQ_BRACKET
"("	O_BRACKET
")"	C_BRACKET
"=="	EQ
"!="	NEQ
"=~"	REGEXP_MATCH
"!~"	REGEXP_NOT_MATCH
"="	ASSIGN
"&&"	AND
"issue"	ISSUE
"type"	TYPE

String	Terminal
"value"	VALUE
"valuetype"	VALUE_TYPE
"claim"	CLAIM
"[_A-Za-z][_A-Za-z0-9]*"	IDENTIFIER
"\"[^\"\n]*\""	STRING
"uint64"	UINT64_TYPE
"int64"	INT64_TYPE
"string"	STRING_TYPE
"boolean"	BOOLEAN_TYPE

Language syntax

The following claims transformation rules language is specified in ABNF form. This definition uses the terminals that are specified in the previous table in addition to the ABNF productions defined here. The rules must be encoded in UTF-16, and the string comparisons must be treated as case insensitive.

```
Rule_set        = ;/*Empty*/
             / Rules
Rules         = Rule
             / Rule Rules
Rule          = Rule_body
Rule_body       = (Conditions IMPLY Rule_action SEMICOLON)
Conditions       = ;/*Empty*/
             / Sel_condition_list
Sel_condition_list   = Sel_condition
             / (Sel_condition_list AND Sel_condition)
Sel_condition     = Sel_condition_body
             / (IDENTIFIER COLON Sel_condition_body)
Sel_condition_body   = O_SQ_BRACKET Opt_cond_list C_SQ_BRACKET
Opt_cond_list     = /*Empty*/
             / Cond_list
Cond_list       = Cond
```

```
              / (Cond_list COMMA Cond)
Cond           = Value_cond
              / Type_cond
Type_cond       = TYPE Cond_oper Literal_expr
Value_cond       = (Val_cond COMMA Val_type_cond)
              /(Val_type_cond COMMA Val_cond)
Val_cond        = VALUE Cond_oper Literal_expr
Val_type_cond     = VALUE_TYPE Cond_oper Value_type_literal
claim_prop       = TYPE
              / VALUE
Cond_oper       = EQ
              / NEQ
              / REGEXP_MATCH
              / REGEXP_NOT_MATCH
Literal_expr      = Literal
              / Value_type_literal

Expr           = Literal
              / Value_type_expr
              / (IDENTIFIER DOT claim_prop)
Value_type_expr    = Value_type_literal
              /(IDENTIFIER DOT VALUE_TYPE)
Value_type_literal   = INT64_TYPE
              / UINT64_TYPE
              / STRING_TYPE
              / BOOLEAN_TYPE
Literal        = STRING
Rule_action      = ISSUE O_BRACKET Issue_params C_BRACKET
Issue_params      = claim_copy
              / claim_new
claim_copy       = CLAIM ASSIGN IDENTIFIER
claim_new       = claim_prop_assign_list
claim_prop_assign_list = (claim_value_assign COMMA claim_type_assign)
              /(claim_type_assign COMMA claim_value_assign)
```

```
claim_value_assign   = (claim_val_assign COMMA claim_val_type_assign)
             /(claim_val_type_assign COMMA claim_val_assign)
claim_val_assign    = VALUE ASSIGN Expr
claim_val_type_assign = VALUE_TYPE ASSIGN Value_type_expr
Claim_type_assign   = TYPE ASSIGN Expr
```

Appendix A: Dynamic Access Control Glossary

Following are the list of terms and definitions that are included in the Dynamic Access Control scenario.

Term	Definition
Automatic classification	Classification that occurs based on classification properties that are determined by classification rules configured by an administrator.
CAPID	Central access policy ID. This ID references a specific central access policy, and it is used to reference the policy from the security descriptor of files and folders.
Central access rule	A rule that includes a condition and an access expression.
Central access policy	Policies that are authored and hosted in Active Directory.
Claims-based access control	A paradigm that utilizes claims to make access control decisions to resources.
Classification	The process of determining the classification properties of resources and assigning these properties to the metadata that is associated with the resources. See also Automatic classification, Inherited classification, and Manual classification.
Device claim	A claim that is associated with the system. With user claims, it is included in the token of a

Term	Definition
	user attempting to access a resource.
Discretionary access control list (DACL)	An access control list that identifies trustees who are allowed or denied access to a securable resource. It can be modified at the discretion of the resource owner.
Resource property	Properties (such as labels) that describe a file and are assigned to files by using automatic classification or manual classification. Examples include: Sensitivity, Project, and Retention period.
File Server Resource Manager	A feature in the Windows Server operating system that offers management of folder quotas, file screening, storage reports, file classification, and file management jobs on a file server.
Folder properties and labels	Properties and labels that describe a folder and are assigned manually by administrators and folder owners. These properties assign default property values to the files within these folders, for example, Secrecy or Department.
Group Policy	A set of rules and policies that controls the working environment of users and computers in an Active Directory environment.
Near real time classification	Automatic classification that is performed shortly after a file is created or modified.
Near real-time file management tasks	File management tasks that are performed shortly after (a file is created or modified. These tasks are triggered by the Near real-time classification.
Organizational Unit (OU)	An Active Directory container that represents hierarchical, logical structures within an organization. It is the smallest scope to which Group Policy settings are applied.
Secure property	A classification property that the authorization runtime can trust to be a valid assertion about the resource at a certain point-in-time. In claims-based access control, a secure property

Term	Definition
	that is assigned to a resource is treated as a resource claim.
Security descriptor	A data structure that contains security information associated with a securable resource, such as access control lists.
Security descriptor definition language	A specification that describes the information in a security descriptor as a text string.
Staging policy	A central access policy that is not yet in effect.
System access control list (SACL)	An access control list that specifies the types of access attempts by particular trustees for which audit records need to be generated.
User claim	Attributes of a user that are provided within the user security token. Examples include: Department, Company, Project, and Security clearance. Information in the user token from systems prior to Windows Server 2012, such as the security groups that the user is part of, can also be considered user claims. Some user claims are provided through Active Directory and others are calculated dynamically, such as whether the user logged in with a smart card.
User token	A data object that identifies a user and the user claims and device claims that are associated with that user. It is used to authorize the user's access to resources.

See Also

Dynamic Access Control: Scenario Overview

Appendix B: Setting Up the Test Environment

This topic outlines the steps to build a hands-on lab to test Dynamic Access Control. The instructions are meant to be followed sequentially because there are many components that have dependencies.

Prerequisites

Hardware and software requirements

Requirements for setting up the test lab:

- A host server running Windows Server 2008 R2 with SP1 and Hyper-V
- A copy of the Windows Server 2012 ISO
- A copy of the Windows 8 ISO
- Microsoft Office 2010
- A server running Microsoft Exchange Server 2003 or later

You need to build the following virtual machines to test the Dynamic Access Control scenarios:

- DC1 (domain controller)
- DC2 (domain controller)
- FILE1 (file server and Active Directory Rights Management Services)
- SRV1 (POP3 and SMTP server)
- CLIENT1 (client computer with Microsoft Outlook)

The passwords for the virtual machines should be as follows:

- BUILTIN\Administrator: pass@word1
- Contoso\Administrator: pass@word1
- All other accounts: pass@word1

Build the test lab virtual machines

Install the Hyper-V role

You need to install the Hyper-V role on a computer running Windows Server 2008 R2 with SP1.

To install the Hyper-V Role

1. Click **Start**, and then click Server Manager.
2. In the Roles Summary area of the Server Manager main window, click **Add Roles**.
3. On the **Select Server Roles** page, click **Hyper-V**.
4. On the **Create Virtual Networks** page, click one or more network adapters if you want to make their network connection available to virtual machines.
5. On the **Confirm Installation Selections** page, click **Install**.
6. The computer must be restarted to complete the installation. Click **Close** to finish the wizard, and then click **Yes** to restart the computer.
7. After you restart the computer, sign in with the same account you used to install the role. After the Resume Configuration Wizard completes the installation, click **Close** to finish the wizard.

Create an internal virtual network

Now you will create an internal virtual network called ID_AD_Network.

To create a virtual network

1. Open Hyper-V Manager.
2. From the **Actions** menu, click **Virtual Network Manager**.
3. Under **Create virtual network**, select the **Internal**.
4. Click **Add**. The **New Virtual Network** page appears.
5. Type **ID_AD_Network** as the name for the new network. Review the other properties and modify them if necessary.
6. Click **OK** to create the virtual network and close Virtual Network Manager, or click **Apply** to create the virtual network and continue using Virtual Network Manager.

Build the domain controller

Build a virtual machine to be used as the domain controller (DC1). Install the virtual machine using Windows Server 2012 ISO, and name it DC1.

To install Active Directory Domain Services

1. Connect the virtual machine to the ID_AD_Network. Sign in to the DC1 as Administrator with the password **pass@word1**.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or Feature-based Install**, and then click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, click **Active Directory Domain Services**. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Domain Services** page, review the information, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**. The Feature installation progress bar on the Results page indicates that the role is being installed.
10. On the **Results** page, verify that the installation succeeded, and click **Close**. In Server Manager, click the warning icon with an exclamation mark on top right corner of the screen, next to **Manage**. In the Tasks list, click the **Promote this server to a domain controller** link.
11. On the **Deployment Configuration** page, click **Add a new forest**, type the name of the root domain, **contoso.com**, and then click **Next**.
12. On the **Domain Controller Options** page, select the domain and forest functional levels as Windows Server 2012, specify the DSRM password **pass@word1**, and then click

Next.

13. On the **DNS Options** page, click **Next**.
14. On the **Additional Options** page, click **Next**.
15. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder (or accept default locations), and then click **Next**.
16. On the **Review Options** page, confirm your selections, and then click **Next**.
17. On the **Prerequisites Check** page, confirm that the prerequisites validation is completed, and then click **Install**.
18. On the **Results** page, verify that the server was successfully configured as a domain controller, and then click **Close**.
19. Restart the server to complete the AD DS installation. (By default, this happens automatically.)

Create the following users by using Active Directory Administrative Center.

Create users and groups on DC1

1. Sign in to contoso.com as Administrator. Launch Active Directory Administrative Center.
2. Create the following security groups:

Group Name	Email Address
FinanceAdmin	financeadmin@contoso.com
FinanceException	financeexception@contoso.com

3. Create the following organizational unit (OU):

OU Name	Computers
FileServerOU	FILE1

4. Create the following users with the attributes indicated:

User	Username	Email address	Department	Group	Country/Region
Myriam Delesalle	MDelesalle	MDelesalle@contoso.com	Finance		US
Miles Reid	MReid	MReid@contoso.com	Finance	FinanceAdmin	US

Esther Valle	EValle	EValle@contoso.com	Operations	FinanceException	US
Maira Wenzel	MWenzel	MWenzel@contoso.com	HR		US
Jeff Low	JLow	JLow@contoso.com	HR		US
RMS Server	rms	rms@contoso.com			

For more information about creating security groups, see Create a New Group on the Windows Server website.

To create a Group Policy Object

1. Hover the cursor on the upper right corner of screen and click the search icon. In the Search box, type **group policy management**, and click **Group Policy Management**.
2. Expand **Forest: contoso.com**, and then expand **Domains**, navigate to **contoso.com**, expand **(contoso.com)**, and then select **FileServerOU**. Right-click **Create a GPO in this domain and Link it here…**
3. Type a descriptive name for the GPO, such as **FlexibleAccessGPO**, and then click **OK**.

To enable Dynamic Access Control for contoso.com

1. Open the Group Policy Management Console, click **contoso.com**, and then double-click **Domain Controllers**.
2. Right-click **Default Domain Controllers Policy**, and select **Edit**.
3. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.
4. Double-click **KDC support for claims, compound authentication, and Kerberos armoring** and select the option next to **Enabled**. You need to enable this setting to use Central Access Policies.
5. Open an elevated command prompt, and run the following command:

```
gpupdate /force
```

Build the file server and AD RMS server (FILE1)

1. Build a virtual machine with the name FILE1 from the Windows Server 2012 ISO.
2. Connect the virtual machine to the ID_AD_Network.

3. Join the virtual machine to the contoso.com domain, and then sign in to FILE1 as contoso\administrator using the password **pass@word1**.

Install File Services Resource Manager

To install the File Services role and the File Server Resource Manager

1. In Server Manager, click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select Server Roles** page, expand **File and Storage Services**, select the check-box next to **File and iSCSI Services**, expand, and select **File Server Resource Manager**.

   In the Add Roles and Features Wizard, click **Add Features**, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. On the **Installation progress** page, click **Close**.

Install the Microsoft Office Filter Packs on the file server

You should install the Microsoft Office Filter Packs on Windows Server 2012 to enable IFilters for a wider array of Office files than are provided by default. Windows Server 2012 does not have any IFilters for Microsoft Office Files installed by default, and the file classification infrastructure uses IFilters to perform content analysis.

To download and install the IFilters, see Microsoft Office 2010 Filter Packs.

Configure email notifications on FILE1

When you create quotas and file screens, you have the option of sending email notifications to users when their quota limit is approaching or after they have attempted to save files that have been blocked. If you want to routinely notify certain administrators of quota and file screening events, you can configure one or more default recipients. To send these notifications, you must specify the SMTP server to be used for forwarding the email messages.

To configure email options in File Server Resource Manager

1. Open File Server Resource Manager. To open File Server Resource Manager, click **Start**, type **file server resource manager**, and then click **File Server Resource Manager**.
2. In the File Server Resource Manager interface, right-click **File Server Resource Manager**, and then click **Configure options**. The **File Server Resource Manager Options** dialog box opens.
3. On the **E-mail Notifications** tab, under SMTP server name or IP address, type the host

name or the IP address of the SMTP server that will forward email notifications.

4. If you want to routinely notify certain administrators of quota or file screening events, under **Default administrator recipients**, type each email address such as fileadmin@contoso.com. Use the format account@domain, and use semicolons to separate multiple accounts.

Create groups on FILE1

To create security groups on FILE1

1. Sign in to FILE1 as contoso\administrator, with the password: **pass@word1**.
2. Add NT AUTHORITY\Authenticated Users to the **WinRMRemoteWMIUsers__** group.

Create files and folders on FILE1

1. Create a new NTFS volume on FILE1 and then create the following folder: D:\Finance Documents.
2. Create the following files with the details specified:
   - **Finance Memo.docx**: Add some finance related text in the document. For example, “The business rules about who can access finance documents have changed. Finance documents are now only accessed by members of the FinanceExpert group. No other departments or groups have access.” You need to evaluate the impact of this change before implementing it in the environment. Ensure that this document has CONTOSO CONFIDENTIAL as the footer on every page.
   - **Request for Approval to Hire.docx**: Create a form in this document that collects applicant information. You must have the following fields in the document: **Applicant Name, Social Security number, Job Title, Proposed Salary, Starting Date, Supervisor name, Department**. Add an additional section in the document that has a form for **Supervisor Signature, Approved Salary, Conformation of Offer**, and **Status of Offer**.
     Make the document rights-management enabled.
   - **Word Document1.docx**: Add some test content to this document.
   - **Word Document2.docx**: Add test content to this document.
   - **Workbook1.xlsx**
   - **Workbook2.xlsx**
   - Create a folder on the desktop called Regular Expressions. Create a text document under the folder called **RegEx-SSN**. Type the following content in the file, and then save and close the file:
     ^(?!000)([0-7]\d{2}|7([0-7]\d|7[012]))([-]?)(?!00)\d\d\3(?!0000)\d{4}$
3. Share the folder D:\Finance Documents as Finance Documents and allow everyone to have Read and Write access to the share.

Note
Central access policies are not enabled by default on the system or boot volume C:.

Install Active Directory Rights Management Services

Add the Active Directory Rights Management Services (AD RMS) and all required features through Server Manager. Choose all the defaults.

To install Active Directory Rights Management Services

1. Sign in to the FILE1 as CONTOSO\Administrator or as a member of the Domain Admins group.

   Important
   In order to install the AD RMS server role the installer account (in this case, CONTOSO\Administrator) will have to be given membership in both the local Administrators group on the server computer where AD RMS is to be installed as well as membership in the Enterprise Admins group in Active Directory.
2. In Server Manager, click **Add Roles and Features**. The Add Roles and Features Wizard appears.
3. On the **Before you Begin** screen, click **Next**.
4. On the **Select Installation Type** screen, click **Role/Feature Based Install**, and then click **Next**.
5. On the **Select Server Targets** screen, click **Next**.
6. On the **Select Server Roles** screen, select the box next to **Active Directory Rights Management Services**, and then click **Next**.
7. In the **Add features that are required for Active Directory Rights Management Services?** dialog box, click **Add Features**.
8. On the **Select Server Roles** screen, click **Next**.
9. On the **Select Features to Install** screen, click **Next**.
10. On the **Active Directory Rights Management Services** screen, click Next.
11. On the **Select Role Services** screen, click **Next**.
12. On the **Web Server Role (IIS)** screen, click **Next**.
13. On the **Select Role Services** screen, click **Next**.
14. On the **Confirm Installation Selections** screen, click **Install**.
15. After the installation has completed, on the **Installation Progress** screen, click **Perform additional configuration**. The AD RMS Configuration Wizard appears.
16. On the **AD RMS** screen, click **Next**.
17. On the **AD RMS Cluster** screen, select **Create a new AD RMS root cluster** and then click **Next**.
18. On the **Configuration Database** screen, click **Use Windows Internal Database on this server**, and then click **Next**.

    Note
    Using the Windows Internal Database is recommended for test environments only because it does not support more than one server in the AD RMS cluster.

Production deployments should use a separate database server.

19. On the **Service Account** screen, in **Domain User Account**, click **Specify** and then specify the user name (**contoso\rms**), and Password (**pass@word1**) and click **OK**, and then click **Next**.
20. On the **Cryptographic Mode** screen, click **Cryptographic Mode 2**.
21. On the **Cluster Key Storage** screen, click **Next**.
22. On the **Cluster Key Password** screen, in the **Password** and **Confirm password** boxes, type **pass@word1**, and then click **Next**.
23. On the **Cluster Web Site** screen, make sure that **Default Web Site** is selected, and then click **Next**.
24. On the **Cluster Address** screen, select the **Use an unencrypted connection** option, in the **Fully Qualified Domain Name** box, type **FILE1.contoso.com**, and then click **Next**.
25. On the **Licensor Certificate Name** screen, accept the default name (**FILE1**) in the text box and click **Next**.
26. On the **SCP Registration** screen, select **Register SCP now**, and then click **Next**.
27. On the **Confirmation** screen, click **Install**.
28. On the **Results** screen, click **Close**, and then click **Close** on **Installation Progress** screen. When complete, log off and log on as contoso\rms using the password provided (**pass@word1**).
29. Launch the AD RMS console and navigate to **Rights Policy Templates**.

    To open the AD RMS console, in Server Manager, click **Local Server** in the console tree, then click **Tools**, and then click **Active Directory Rights Management Services**.
30. Click the **Create Distributed Rights Policy** template located on the right panel, click **Add**, and select the following information:
    - Language: US English
    - Name: Contoso Finance Admin Only
    - Description: Contoso Finance Admin Only

    Click **Add**, and then click **Next**.
31. Under the Users and Rights section, click **Users and rights**, click **Add**, type **financeadmin@contoso.com**, and click **OK**.
32. Select **Full Control**, and leave **Grant owner (author) full control right with no expiration** selected.
33. Click though the remaining tabs with no changes, and then click **Finish**. Sign in as CONTOSO\Administrator.
34. Browse to the folder, C:\inetpub\wwwroot_wmcs\certification, select the ServerCertification.asmx file, and add Authenticated Users to have Read and Write permissions to the file.
35. Open Windows PowerShell and run `Get-FsrmRmsTemplate`. Verify that you are able to see the RMS template you created in the previous steps in this procedure with this command.

Important

If you want your file servers to immediately change so you can test them, you need to do the following:

1. On the file server, FILE1, open an elevated command prompt, and run the following commands:
   - gpupdate /force.
   - NLTEST /SC_RESET:contoso.com
2. On the domain controller (DC1), replicate Active Directory.

   For more information about steps to force the replication of Active Directory, see Active Directory Replication

Optionally, instead of using the Add Roles and Features Wizard in Server Manager, you can use Windows PowerShell to install and configure the AD RMS server role as show in the following procedure.

To install and configure an AD RMS cluster in Windows Server 2012 using Windows PowerShell

1. Logon on as CONTOSO\Administrator with the password: **pass@word1**.

   Important
   In order to install the AD RMS server role the installer account (in this case, CONTOSO\Administrator) will have to be given membership in both the local Administrators group on the server computer where AD RMS is to be installed as well as membership in the Enterprise Admins group in Active Directory.
2. On the Server desktop, right-click the Windows PowerShell icon on the taskbar and select **Run as Administrator** to open a Windows PowerShell prompt with administrative privileges.
3. To use Server Manager cmdlets to install the AD RMS server role, type:

   ```
   Add-WindowsFeature ADRMS –IncludeAllSubFeature –
   IncludeManagementTools
   ```
4. Create the Windows PowerShell drive to represent the AD RMS server you are installing.

   For example, to create a Windows PowerShell drive named RC to install and configure the first server in an AD RMS root cluster, type:

   ```
   Import-Module ADRMS
   New-PSDrive -PSProvider ADRMSInstall -Name RC -Root
   RootCluster
   ```
5. Set properties on objects in the drive namespace that represent required configuration settings.

   For example, to set the AD RMS service account, at the Windows PowerShell command prompt, type:

   ```
   $svcacct = Get-Credential
   ```

   When the Windows security dialog box appears, type the AD RMS service account

domain user name CONTOSO\RMS and the assigned password.

Next, to assign the AD RMS service account to the AD RMS cluster settings, type the following:

```
Set-ItemProperty –Path RC:\ -Name ServiceAccount -Value
$svcacct
```

Next, to set the AD RMS server to use the Windows Internal Database, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty –Path RC:\ClusterDatabase -Name
UseWindowsInternalDatabase -Value $true
```

Next, to securely store the cluster key password in a variable, at the Windows PowerShell command prompt, type:

```
$password = Read-Host -AsSecureString -Prompt "Password:"
```

Type the cluster key password, and then press the ENTER key.

Next, to assign the password to your AD RMS installation, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ClusterKey -Name
CentrallyManagedPassword -Value $password
```

Next, to set the AD RMS cluster address, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name ClusterURL -Value
"http://file1.contoso.com:80"
```

Next, to assign the SLC name for your AD RMS installation, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name SLCName -Value "FILE1"
```

Next, to set the service connection point (SCP) for the AD RMS cluster, at the Windows PowerShell command prompt, type:

```
Set-ItemProperty -Path RC:\ -Name RegisterSCP -Value $true
```

6. Run the **Install-ADRMS** cmdlet. In addition to installing the AD RMS server role and configuring the server, this cmdlet also installs other features required by AD RMS if necessary.

   For example, to change to the Windows PowerShell drive named RC and install and configure AD RMS, type:

   ```
   Set-Location RC:\
   Install-ADRMS –Path.
   ```

   Type "Y" when the cmdlet prompts you to confirm you want to start the installation.

7. Log out as CONTOSO\Administrator and log on as CONTOSO\RMS using the provided password ("pass@word1").

Important

In order to manage the AD RMS server the account you are logged on to and using to manage the server (in this case, CONTOSO\RMS) will have to be given membership in both the local Administrators group on the AD RMS server computer as well as membership in the Enterprise Admins group in Active Directory.

8. On the Server desktop, right-click the Windows PowerShell icon on the taskbar and select **Run as Administrator** to open a Windows PowerShell prompt with administrative privileges.
9. Create the Windows PowerShell drive to represent the AD RMS server you are configuring.

   For example, to create a Windows PowerShell drive named RC to configure the AD RMS root cluster, type:

   ```
   Import-Module ADRMSAdmin `
   New-PSDrive -PSProvider ADRMSAdmin -Name RC -Root
   http://localhost -Force -Scope Global
   ```

10. To create new rights template for the Contoso finance administrator and assign it user rights with full control in your AD RMS installation, at the Windows PowerShell command prompt, type:

    ```
    New-Item -Path RC:\RightsPolicyTemplate –LocaleName en-us -
    DisplayName "Contoso Finance Admin Only" -Description
    "Contoso Finance Admin Only" -UserGroup
    financeadmin@contoso.com  -Right ('FullControl')
    ```

11. To verify that you can see the new rights template for the Contoso finance administrator, at the Windows PowerShell command prompt:

    ```
    Get-FsrmRmsTemplate
    ```

    Review the output of this cmdlet to confirm the RMS template you created in the previous step is present.

Build the mail server (SRV1)

SRV1 is the SMTP/POP3 mail server. You need to set it up so that you can send email notifications as part of the Access-Denied assistance scenario.

Configure Microsoft Exchange Server on this computer. For more information, see How to Install Exchange Server.

Build the client virtual machine (CLIENT1)

To build the client virtual machine

1. Connect the CLIENT1 to the ID_AD_Network.

2. Install Microsoft Office 2010.
3. Sign in as Contoso\Administrator, and use the following information to configure Microsoft Outlook.
    - Your name: File Administrator
    - Email address: fileadmin@contoso.com
    - Account type: POP3
    - Incoming mail server: Static IP address of SRV1
    - Outgoing mail server: Static IP address of SRV1
    - User name: fileadmin@contoso.com
    - Remember password: Select
4. Create a shortcut to Outlook on the contoso\administrator desktop.
5. Open Outlook and address all the “first time launched” messages.
6. Delete any test messages that were generated.
7. Create a new short cut on desktop for all users on the client virtual machine that points to \\FILE1\Finance Documents.
8. Reboot as needed.

Enable Access-Denied assistance on the client virtual machine

1. Open Registry Editor, and navigate to **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\Explorer**.
    - Set **EnableShellExecuteFileStreamCheck** to **1**.
    - Value: DWORD

Lab setup for deploying claims across forests scenario

Build a virtual machine for DC2

- Build a virtual machine from the Windows Server 2012 ISO.

- Create the virtual machine name as DC2.
- Connect the virtual machine to the ID_AD_Network.

Important
Joining virtual machines to a domain and deploying claim types across forests require that the virtual machines be able to resolve the FQDNs of the relevant domains. You may have to manually configure the DNS settings on the virtual machines to accomplish this. For more information, see Configuring a virtual network.

All the virtual machine images (servers and clients) must be reconfigured to use a static IP version 4 (IPv4) address and Domain Name System (DNS) client settings. For more information, see Configure a DNS Client for Static IP Address.

Set up a new forest called adatum.com

To install Active Directory Domain Services

1. Connect the virtual machine to the ID_AD_Network. Sign in to the DC2 as Administrator with the password **Pass@word1**.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select Installation Type** page, click **Role-based or Feature-based Install**, and then click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, click the names of the server where you want to install Active Directory Domain Services (AD DS), and then click **Next**.
6. On the **Select Server Roles** page, click **Active Directory Domain Services**. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
7. On the **Select Features** page, click **Next**.
8. On the **AD DS** page, review the information, and then click **Next**.
9. On the **Confirmation** page, click **Install**. The Feature installation progress bar on the Results page indicates that the role is being installed.
10. On the **Results** page, verify that the installation succeeded, and then click the warning icon with an exclamation mark on top right corner of the screen, next to **Manage**. In the Tasks list, click the **Promote this server to a domain controller** link.

    Important
    If you close the installation wizard at this point rather than click **Promote this server to a domain controller**, you can continue the AD DS installation by clicking **Tasks** in Server Manager.
11. On the **Deployment Configuration** page, click **Add a new forest**, type the name of the root domain, **adatum.com**, and then click **Next**.
12. On the **Domain Controller Options** page, select the domain and forest functional levels as Windows Server 2012, specify the DSRM password **pass@word1**, and then click **Next**.
13. On the **DNS Options** page, click **Next**.
14. On the **Additional Options** page, click **Next**.
15. On the **Paths** page, type the locations for the Active Directory database, log files, and SYSVOL folder (or accept default locations), and then click **Next**.
16. On the **Review Options** page, confirm your selections, and then click **Next**.
17. On the **Prerequisites Check** page, confirm that the prerequisites validation is completed,

and then click **Install**.

18. On the **Results** page, verify that the server was successfully configured as a domain controller, and then click **Close**.
19. Restart the server to complete the AD DS installation. (By default, this happens automatically.)

Important

To ensure that the network is configured properly, after you have set up both the forests, you must do the following:

- Sign in to adatum.com as adatum\administrator. Open a Command Prompt window, type **nslookup contoso.com**, and then press ENTER.
- Sign in to contoso.com as contoso\administrator. Open a Command Prompt window, type **nslookup adatum.com**, and then press ENTER.

  If these commands execute without errors, the forests can communicate with each other. For more information on nslookup errors, see the troubleshooting section in the topic Using NSlookup.exe

Set contoso.com as a trusting forest to adatum.com

In this step, you create a trust relationship between the Adatum Corporation site and the Contoso, Ltd. site.

To set Contoso as a trusting forest to Adatum

1. Sign in to DC2 as administrator. On the **Start** screen, type domain.msc.
2. In the console tree, right-click adatum.com, and then click Properties.
3. On the **Trusts** tab, click **New Trust**, and then click **Next**.
4. On the **Trust Name** page, type **contoso.com**, in the Domain Name System (DNS) name field, and then click **Next**.
5. On the **Trust Type** page, click **Forest Trust**, and then click **Next**.
6. On the **Direction of Trust** page, click **Two-way**.
7. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.
8. Continue to follow the instructions in the wizard.

Create additional users in the Adatum forest

Create the user Jeff Low with the password **pass@word1**, and assign the company attribute with the value **Adatum**.

To create a user with the Company attribute

1. Open an elevated command prompt in Windows PowerShell, and paste the following code:

```
New-ADUser `
-SamAccountName jlow `
-Name "Jeff Low" `
-UserPrincipalName jlow@adatum.com `
-AccountPassword (ConvertTo-SecureString `
-AsPlainText "pass@word1" -Force) `
-Enabled $true `
-PasswordNeverExpires $true `
-Path 'CN=Users,DC=adatum,DC=com' `
-Company Adatum`
```

Create the Company claim type on adataum.com

To create a claim type by using Windows PowerShell

1. Sign in to adatum.com as an administrator.
2. Open an elevated command prompt in Windows PowerShell, and type the following code:

```
New-ADClaimType `
-AppliesToClasses:@('user') `
-Description:"Company" `
-DisplayName:"Company" `
-ID:"ad://ext/Company:ContosoAdatum" `
-IsSingleValued:$true `
-Server:"adatum.com" `
-SourceAttribute:Company `
-SuggestedValues:@((New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("C
ontoso", "Contoso", "")), (New-Object
Microsoft.ActiveDirectory.Management.ADSuggestedValueEntry("A
datum", "Adatum", ""))) `
```

Enable the Company resource property on contoso.com

To enable the Company resource property on contoso.com

1. Sign in to contoso.com as an administrator.
2. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. In the left pane of Active Directory Administrative Center, click **Tree View**. In the left pane, click **Dynamic Access Control**, and then double-click **Resource Properties**.
4. Select **Company** from the **Resource Properties** list, right-click and select **Properties**. In the **Suggested Values** section, click **Add** to add the suggested values: Contoso and Adatum, and then click **OK** twice.
5. Select **Company** from the **Resource Properties** list, right-click and select **Enable**.

Enable Dynamic Access Control on adatum.com

To enable Dynamic Access Control for adatum.com

1. Sign in to adatum.com as an administrator.
2. Open the Group Policy Management Console, click **adatum.com**, and then double-click **Domain Controllers**.
3. Right-click **Default Domain Controllers Policy**, and select **Edit**.
4. In the Group Policy Management Editor window, double-click **Computer Configuration**, double-click **Policies**, double-click **Administrative Templates**, double-click **System**, and then double-click **KDC**.
5. Double-click **KDC support for claims, compound authentication, and Kerberos armoring** and select the option next to **Enabled**. You need to enable this setting to use Central Access Policies.
6. Open an elevated command prompt, and run the following command:

```
gpupdate /force
```

Create the Company claim type on contoso.com

To create a claim type by using Windows PowerShell

1. Sign in to contoso.com as an administrator.
2. Open an elevated command prompt in Windows PowerShell, then type the following code:

```
New-ADClaimType -SourceTransformPolicy `
-DisplayName "Company" `
-ID "ad://ext/Company:ContosoAdatum" `
-IsSingleValued $true `
```

```
-ValueType "string" `
```

Create the central access rule

To create a central access rule

1. In the left pane of Active Directory Administrative Center, click **Tree View**. In the left pane, click **Dynamic Access Control**, and then click **Central Access Rules**.
2. Right-click **Central Access Rules**, click **New**, and then **Central Access Rule**.
3. In the **Name** field, type **AdatumEmployeeAccessRule**.
4. In the **Permissions** section, select the **Use following permissions as current permissions** option, click **Edit**, and then click **Add**. Click the **Select a principal** link, type **Authenticated Users**, and then click **OK**.
5. In the **Permission Entry for Permissions** dialog box, click **Add a condition**, and enter the following conditions: [**User**] [**Company**] [**Equals**] [**Value**] [**Adatum**]. Permissions should be **Modify, Read and Execute, Read, Write**.
6. Click **OK**.
7. Click **OK** three times to finish and return to Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADCentralAccessRule `
-
CurrentAcl:"O:SYG:SYD:AR(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)(
XA;;0x1301bf;;;AU;(@USER.ad://ext/Company:ContosoAdatum ==
`"Adatum`"))" `
-Name:"AdatumEmployeeAccessRule" `
-ProposedAcl:$null `
-ProtectedFromAccidentalDeletion:$true `
-Server:"contoso.com" `
```

Create the central access policy

To create a central access policy

1. Sign in to contoso.com as an administrator.

2. Open an elevated command prompt in Windows PowerShell, and then paste the following code:

```
New-ADCentralAccessPolicy "Adatum Only Access Policy"
Add-ADCentralAccessPolicyMember "Adatum Only Access Policy" `
-Member "AdatumEmployeeAccessRule" `
```

Publish the new policy through Group Policy

To apply the central access policy across file servers through Group Policy

1. On the **Start** screen, type **Administrative Tools**, and in the **Search** bar, click **Settings**. In the **Settings** results, click **Administrative Tools**. Open the Group Policy Management Console from the **Administrative Tools** folder.

   Tip
   If the **Show Administrative tools** setting is disabled, the Administrative Tools folder and its contents will not appear in the **Settings** results.
2. Right-click the contoso.com domain, click **Create a GPO in this domain and Link it here…**
3. Type a descriptive name for the GPO, such as **AdatumAccessGPO**, and then click **OK**.

To apply the central access policy to the file server through Group Policy

1. On the **Start** screen, type **Group Policy Management**, in the **Search** box. Open **Group Policy Management** from the Administrative Tools folder.

   Tip
   If the **Show Administrative tools** setting is disabled, the Administrative Tools folder and its contents will not appear in the Settings results.
2. Navigate to and select Contoso as follows: Group Policy Management\Forest: contoso.com\Domains\contoso.com.
3. Right-click the **AdatumAccessGPO** policy, and select **Edit**.
4. In Group Policy Management Editor, click **Computer Configuration**, expand **Policies**, expand **Windows Settings**, and then click **Security Settings**.
5. Expand **File System**, right-click **Central Access Policy**, and then click **Manage Central access policies**.
6. In the **Central Access Policies Configuration** dialog box, click **Add**, select **Adatum Only Access Policy**, and then click **OK**.
7. Close the Group Policy Management Editor. You have now added the central access policy to Group Policy.

Create the Earnings folder on the file server

Create a new NTFS volume on FILE1, and create the following folder: D:\Earnings.

Note
Central access policies are not enabled by default on the system or boot volume C:.

Set classification and apply the central access policy on the Earnings folder

To assign the central access policy on the file server

1. In Hyper-V Manager, connect to server FILE1. Sign in to the server by using Contoso\Administrator, with the password **pass@word1**.
2. Open an elevated command prompt and type: **gpupdate /force**. This will ensure that your Group Policy changes will take effect on your server.
3. You also need to refresh the Global Resource Properties from Active Directory. Open Windows PowerShell, type `Update-FSRMClassificationpropertyDefinition`, and then press ENTER. Close Windows PowerShell.
4. Open Windows Explorer, and navigate to D:\EARNINGS. Right-click the **Earnings** folder, and click **Properties**.
5. Click the **Classification** tab. Select **Company**, and then select **Adatum** in the **Value** field.
6. Click **Change**, select **Adatum Only Access Policy** from the drop-down menu, and then click **Apply**.
7. Click the **Security** tab, click **Advanced**, and then click the **Central Policy** tab. You should see the **AdatumEmployeeAccessRule** listed. You can expand the item to view all of the permissions that you set when you created the rule in Active Directory.
8. Click **OK** to return to Windows Explorer.

Best Practices for Securing Active Directory

This document provides a practitioner's perspective and contains a set of practical techniques to help IT executives protect an enterprise Active Directory environment. Active Directory plays a critical role in the IT infrastructure, and ensures the harmony and security of different network resources in a global, interconnected environment. The methods discussed are based largely on the Microsoft Information Security and Risk Management (ISRM) organization's experience, which is accountable for protecting the assets of Microsoft IT and other Microsoft Business Divisions, in addition to advising a selected number of Microsoft Global 500 customers.

- Foreword

- Acknowledgements
- Executive Summary
- Introduction
- Avenues to Compromise
- Attractive Accounts for Credential Theft
- Reducing the Active Directory Attack Surface
- Implementing Least-Privilege Administrative Models
- Implementing Secure Administrative Hosts
- Securing Domain Controllers Against Attack
- Monitoring Active Directory for Signs of Compromise
- Audit Policy Recommendations
- Planning for Compromise
- Maintaining a More Secure Environment
- Summary of Best Practices
- Appendices
- Appendix A: Patch and Vulnerability Management Software
- Appendix B: Privileged Accounts and Groups in Active Directory
- Appendix C: Protected Accounts and Groups in Active Directory
- Appendix D: Securing Built-In Administrator Accounts in Active Directory
- Appendix E: Securing Enterprise Admins Groups in Active Directory
- Appendix F: Securing Domain Admins Groups in Active Directory
- Appendix G: Securing Administrators Groups in Active Directory
- Appendix H: Securing Local Administrator Accounts and Groups
- Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory
- Appendix J: Third-Party RBAC Vendors
- Appendix K: Third-Party PIM Vendors
- Appendix L: Events to Monitor
- Appendix M: Document Links and Recommended Reading

Foreword

In today's information rich environment, senior executives are faced with the challenge of harnessing information technology to help their business:

- Execute its strategy
- Improve its operations
- Enhance the perceived value of its own products and services

In support of these challenges, consumer-centric computing models that leverage highly scalable platforms (the cloud) and a plethora of devices are being utilized. The new business paradigm is obsessed with speed, agility, and execution. The evolution of these models requires even more comprehensive and agile security and risk management programs to ensure success.

This document provides a practitioner's perspective and contains a set of practical techniques to help IT executives protect an enterprise Active Directory® environment. Active Directory plays a critical role in the IT infrastructure, and ensures the harmony and security of different network resources in a global, interconnected environment. The methods discussed are based largely on the Microsoft® Information Security and Risk Management (ISRM) organization's experience, which is accountable for protecting the assets of Microsoft IT and other Microsoft Business Divisions, in addition to advising a selected number of Microsoft Global 500 customers.

Key tenets of this paper are understanding the avenues for establishing a healthy Active Directory, implementing monitoring systems, actions to reduce the attack surface, and managing a resilient environment. This risk-based approach assumes that the corporate infrastructure, and more specifically the Active Directory, is a critical target. With this mindset, resiliency and recovery become critical components of an Active Directory protection program

This document encompasses experience from several hundred Active Directory Security Assessments, critical incident responses, and recovery engagements, and proven techniques for mitigating IT risks. Based on numerous requests from Microsoft customers and partners, this document reflects a comprehensive guide, and it contains best practices for protecting Active Directory. Information security and risk management executives will find the techniques explained

in this document to be a significant contribution to their understanding of best practices, in addition to practical implementation programs for their Active Directory environments.

Bret Arsenault

Microsoft Chief Information Security Officer

Acknowledgements

Authors	Dean Wells
Laura A. Robinson, Lead Author	Fernando Cima
Khurram Chaudhary	Georgeo Pulikkathara
Roger Grimes	Joost van Haaren
Eric Leonard	Joseph Lindstrom
Contributors	Laura Hunter
Ahmad Mahdi	Lesley Kipling
Ashish Popli	Mark Simos
Dan Kaufman	Michiko Short
Jenn LeMond	Neil Carpenter
Joost van Haaren	Patrick Jungles
Rob Labbé	Paul Rich
Reviewers	Rob Labbé
Alan von Weltin	Sean Finnegan
Andrew Idell	Tim Rains
Bret Arsenault	Todd Thompson
David Shaw	

Executive Summary

No organization with an information technology (IT) infrastructure is immune from attack, but if appropriate policies, processes, and controls are implemented to protect key segments of an organization's computing infrastructure, it might be possible to prevent a breach event from growing to a wholesale compromise of the computing environment.

This executive summary is intended to be useful as a standalone document summarizing the content of the document, which contains recommendations that will assist organizations in enhancing the security of their Active Directory installations. By implementing these recommendations, organizations will be able to identify and prioritize security activities, protect

key segments of their organization's computing infrastructure, and create controls that significantly decrease the likelihood of successful attacks against critical components of the IT environment.

Although this document discusses the most common attacks against Active Directory and countermeasures to reduce the attack surface, it also contains recommendations for recovery in the event of complete compromise. The only sure way to recover in the event of a complete compromise of Active Directory is to be prepared for the compromise before it happens.

The major sections of this document are:

- Avenues to Compromise
- Reducing the Active Directory Attack Surface
- Monitoring Active Directory for Signs of Compromise
- Planning for Compromise

Avenues to Compromise

This section provides information about some of the most commonly leveraged vulnerabilities used by attackers to compromise customers' infrastructures. It contains general categories of vulnerabilities and how they're used to initially penetrate customers' infrastructures, propagate compromise across additional systems, and eventually target Active Directory and domain controllers to obtain complete control of the organizations' forests. It does not provide detailed recommendations about addressing each type of vulnerability, particularly in the areas in which the vulnerabilities are not used to directly target Active Directory. However, for each type of vulnerability, we have provided links to additional information to use to develop countermeasures and reduce the organization's attack surface.

Included are the following subjects:

- **Initial breach targets** - Most information security breaches start with the compromise of small pieces of an organization's infrastructure—often one or two systems at a time. These initial events, or entry points into the network, often exploit vulnerabilities that could have been fixed, but weren't. Commonly seen vulnerabilities are:
  - Gaps in antivirus and antimalware deployments
  - Incomplete patching
  - Outdated applications and operating systems
  - Misconfiguration
  - Lack of secure application development practices
- **Attractive Accounts for Credential Theft** - Credential theft attacks are those in which an attacker initially gains privileged access to a computer on a network and then uses freely available tooling to extract credentials from the sessions of other logged-on accounts.

  Included in this section are the following:

  - **Activities that Increase the Likelihood of Compromise** - Because the target of credential theft is usually highly privileged domain accounts and "very important person"

(VIP) accounts, it is important for administrators to be conscious of activities that increase the likelihood of a success of a credential-theft attack. These activities are:

- Logging on to unsecured computers with privileged accounts
- Browsing the Internet with a highly privileged account
- Configuring local privileged accounts with the same credentials across systems
- Overpopulation and overuse of privileged domain groups
- Insufficient management of the security of domain controllers.

- **Privilege Elevation and Propagation** - Specific accounts, servers, and infrastructure components are usually the primary targets of attacks against Active Directory. These accounts are:
  - Permanently privileged accounts
  - VIP accounts
  - “Privilege-Attached” Active Directory accounts
  - Domain controllers
  - Other infrastructure services that affect identity, access, and configuration management, such as public key infrastructure (PKI) servers and systems management servers

Reducing the Active Directory Attack Surface

This section focuses on technical controls to reduce the attack surface of an Active Directory installation. Included in this section are the following subjects:

- The **Privileged Accounts and Groups in Active Directory** section discusses the highest privileged accounts and groups in Active Directory and the mechanisms by which privileged accounts are protected. Within Active Directory, three built-in groups are the highest privilege groups in the directory (Enterprise Admins, Domain Admins, and Administrators), although a number of additional groups and accounts should also be protected.
- The **Implementing Least-Privilege Administrative Models** section focuses on identifying the risk that the use of highly privileged accounts for day-to-day administration presents, in addition to providing recommendations to reduce that risk.

Excessive privilege isn’t only found in Active Directory in compromised environments. When an organization has developed the habit of granting more privilege than is required, it is typically found throughout the infrastructure:

- In Active Directory
- On member servers
- On workstations
- In applications
- In data repositories
- The **Implementing Secure Administrative Hosts** section describes secure administrative hosts, which are computers that are configured to support administration of Active Directory and connected systems. These hosts are dedicated to administrative functionality and do not

run software such as email applications, web browsers, or productivity software (such as Microsoft Office).

Included in this section are the following:

- **Principles for Creating Secure Administrative Hosts** - The general principles to keep in mind are:
  - Never administer a trusted system from a less-trusted host.
  - Do not rely on a single authentication factor when performing privileged activities.
  - Do not forget physical security when designing and implementing secure administrative hosts.
- **Securing Domain Controllers Against Attack** - If a malicious user obtains privileged access to a domain controller, that user can modify, corrupt, and destroy the Active Directory database, and by extension, all of the systems and accounts that are managed by Active Directory.

Included in this section are the following subjects:

- **Physical Security for Domain Controllers** – Contains recommendations for providing physical security for domain controllers in datacenters, branch offices, and remote locations.
- **Domain Controller Operating Systems** – Contains recommendations for securing the domain controller operating systems.
- **Secure Configuration of Domain Controllers** – Native and freely available configuration tools and settings can be used to create security configuration baselines for domain controllers that can subsequently be enforced by Group Policy Objects (GPOs).

Monitoring Active Directory for Signs of Compromise

This section provides information about legacy audit categories and audit policy subcategories (which were introduced in Windows Vista and Windows Server® 2008), and Advanced Audit Policy (which was introduced in Windows Server 2008 R2). Also provided is information about events and objects to monitor that can indicate attempts to compromise the environment and some additional references that can be used to construct a comprehensive audit policy for Active Directory.

Included in this section are the following subjects:

- **Windows Audit Policy** – Windows security event logs have categories and subcategories that determine which security events are tracked and recorded.
- **Audit Policy Recommendations** - This section describes the Windows default audit policy settings, audit policy settings that are recommended by Microsoft, and more aggressive recommendations for organizations to use to audit critical servers and workstations.

Planning for Compromise

This section contains recommendations that will help organizations prepare for a compromise before it happens, implement controls that can detect a compromise event before a full breach

has occurred, and provide response and recovery guidelines for cases in which a complete compromise of the directory is achieved by attackers. Included in this section are the following subjects:

- **Rethinking the Approach** - Contains principles and guidelines to create secure environments into which an organization can place their most critical assets. These guidelines are as follows:
  - Identifying principles for segregating and securing critical assets
  - Defining a limited, risk-based migration plan
  - Leveraging “nonmigratory” migrations where necessary
  - Implementing “creative destruction”
  - Isolating legacy systems and applications
  - Simplifying security for end users
- **Maintaining a More Secure Environment** – Contains high-level recommendations meant to be used as guidelines to use in developing not only effective security, but effective lifecycle management. Included in this section are the following subjects:
  - **Creating Business-Centric Security Practices for Active Directory** - To effectively manage the lifecycle of the users, data, applications and systems managed by Active Directory, follow these principles.
    - **Assign a Business Ownership to Active Directory Data** - Assign ownership of infrastructure components to IT; for data that is added to Active Directory Domain Services (AD DS) to support the business, for example, new employees, new applications, and new information repositories, a designated business unit or user should be associated with the data.
    - **Implement Business-Driven Lifecycle Management** - Lifecycle management should be implemented for data in Active Directory.
    - **Classify all Active Directory Data** - Business owners should provide classification for data in Active Directory. Within the data classification model, classification for the following Active Directory data should be included:
      - **Systems** - Classify server populations, their operating system their role, the applications running on them, and the IT and business owners of record.
      - **Applications** - Classify applications by functionality, user base, and their operating system.
      - **Users** - The accounts in the Active Directory installations that are most likely to be targeted by attackers should be tagged and monitored.

Summary of Best Practices for Securing Active Directory Domain Services

The following table provides a summary of the recommendations provided in this document for securing an AD DS installation. Some best practices are strategic in nature and require

comprehensive planning and implementation projects; others are tactical and focused on specific components of Active Directory and related infrastructure.

Practices are listed in approximate order of priority, that is., lower numbers indicate higher priority. Where applicable, best practices are identified as preventative or detective in nature. All of these recommendations should be thoroughly tested and modified as needed for your organization's characteristics and requirements.

	Best Practice	**Tactical or Strategic**	**Preventative or Detective**
1	Patch applications.	Tactical	Preventative
2	Patch operating systems.	Tactical	Preventative
3	Deploy and promptly update antivirus and antimalware software across all systems and monitor for attempts to remove or disable it.	Tactical	Both
4	Monitor sensitive Active Directory objects for modification attempts and Windows for events that may indicate attempted compromise.	Tactical	Detective
5	Protect and monitor accounts for users who have access to sensitive data	Tactical	Both
6	Prevent powerful accounts from being used on unauthorized systems.	Tactical	Preventative
7	Eliminate permanent membership in highly privileged groups.	Tactical	Preventative
8	Implement controls to grant temporary membership in privileged groups when needed.	Tactical	Preventative
9	Implement secure	Tactical	Preventative

	administrative hosts.		
10	Use application whitelisting on domain controllers, administrative hosts, and other sensitive systems.	Tactical	Preventative
11	Identify critical assets, and prioritize their security and monitoring.	Tactical	Both
12	Implement least-privilege, role-based access controls for administration of the directory, its supporting infrastructure, and domain-joined systems.	Strategic	Preventative
13	Isolate legacy systems and applications.	Tactical	Preventative
14	Decommission legacy systems and applications.	Strategic	Preventative
15	Implement secure development lifecycle programs for custom applications.	Strategic	Preventative
16	Implement configuration management, review compliance regularly, and evaluate settings with each new hardware or software version.	Strategic	Preventative
17	Migrate critical assets to pristine forests with stringent security and monitoring requirements.	Strategic	Both
18	Simplify security for end users.	Strategic	Preventative
19	Use host-based firewalls to control and secure communications.	Tactical	Preventative
20	Patch devices.	Tactical	Preventative

21	Implement business-centric lifecycle management for IT assets.	Strategic	N/A
22	Create or update incident recovery plans.	Strategic	N/A

Introduction

Attacks against computing infrastructures, whether simple or complex, have existed as long as computers have. However, within the past decade, increasing numbers of organizations of all sizes, in all parts of the world have been attacked and compromised in ways that have significantly changed the threat landscape. Cyber-warfare and cybercrime have increased at record rates. “Hacktivism,” in which attacks are motivated by activist positions, has been claimed as the motivation for a number of breaches intended to expose organizations’ secret information, to create denials-of-service, or even to destroy infrastructure. Attacks against public and private institutions with the goal of exfiltrating the organizations’ intellectual property (IP) have become ubiquitous.

No organization with an information technology (IT) infrastructure is immune from attack, but if appropriate policies, processes, and controls are implemented to protect key segments of an organization’s computing infrastructure, escalation of attacks from penetration to complete compromise might be preventable. Because the number and scale of attacks originating from outside an organization has eclipsed insider threat in recent years, this document often discusses external attackers rather than misuse of the environment by authorized users. Nonetheless, the principles and recommendations provided in this document are intended to help secure your environment against external attackers and misguided or malicious insiders.

The information and recommendations provided in this document are drawn from a number of sources and derived from practices designed to protect Active Directory installations against compromise. Although it is not possible to prevent attacks, it is possible to reduce the Active Directory attack surface and to implement controls that make compromise of the directory much more difficult for attackers. This document presents the most common types of vulnerabilities we have observed in compromised environments and the most common recommendations we have made to customers to improve the security of their Active Directory installations.

Account and Group Naming Conventions

The following table provides a guide to the naming conventions used in this document for the groups and accounts referenced throughout the document. Included in the table is the location of each account/group, its name, and how these accounts/groups are referenced in this document.

Account/Group Location	Name of Account/Group	How It is Referenced in this Document
Active Directory - each domain	Administrator	Built-in Administrator account
Active Directory - each domain	Administrators	Built-in Administrators (BA) group
Active Directory - each domain	Domain Admins	Domain Admins (DA) group
Active Directory - forest root domain	Enterprise Admins	Enterprise Admins (EA) group
Local computer security accounts manager (SAM) database on computers running Windows Server and workstations that are not domain controllers	Administrator	Local Administrator account
Local computer security accounts manager (SAM) database on computers running Windows Server and workstations that are not domain controllers	Administrators	Local Administrators group

About This Document

The Microsoft Information Security and Risk Management (ISRM) organization, which is part of Microsoft Information Technology (MSIT), works with internal business units, external customers, and industry peers to gather, disseminate, and define policies, practices, and controls. This information can be used by Microsoft and our customers to increase the security and reduce the attack surface of their IT infrastructures. The recommendations provided in this document are based on a number of information sources and practices used within MSIT and ISRM. The following sections present more information about the origins of this document.

Microsoft IT and ISRM

A number of practices and controls have been developed within MSIT and ISRM to secure the Microsoft AD DS forests and domains. Where these controls are broadly applicable, they have been integrated into this document. SAFE-T (Solution Accelerators for Emerging Technologies) is a team within ISRM whose charter is to identify emerging technologies, and to define security requirements and controls to accelerate their adoption.

Active Directory Security Assessments

Within Microsoft ISRM, the Assessment, Consulting, and Engineering (ACE) Team works with internal Microsoft business units and external customers to assess application and infrastructure security and to provide tactical and strategic guidance to increase the organization's security posture. One ACE service offering is the Active Directory Security Assessment (ADSA), which is a holistic assessment of an organization's AD DS environment that assesses people, process, and technology and produces customer-specific recommendations. Customers are provided with recommendations that are based on the organization's unique characteristics, practices, and risk appetite. ADSAs have been performed for Active Directory installations at Microsoft in addition to those of our customers. Over time, a number of recommendations have been found to be applicable across customers of varying sizes and industries.

Content Origin and Organization

Much of the content of this document is derived from the ADSA and other ACE Team assessments performed for compromised customers and customers who have not experienced significant compromise. Although individual customer data was not used to create this document, we have collected the most commonly exploited vulnerabilities we have identified in our assessments and the recommendations we have made to customers to improve the security of their AD DS installations. Not all vulnerabilities are applicable to all environments, nor are all recommendations feasible to implement in every organization.

This document is organized as follows:

Executive Summary

The Executive Summary, which can be read as a standalone document or in combination with the full document, provides a high-level summary of this document. Included in the Executive Summary are the most common attack vectors we have observed used to compromise customer environments, summary recommendations for securing Active Directory installations, and basic objectives for customers who plan to deploy new AD DS forests now or in the future.

Introduction

This is the section you are reading now.

Avenues to Compromise

This section provides information about some of the most commonly leveraged vulnerabilities we have found to be used by attackers to compromise customers' infrastructures. This section begins with general categories of vulnerabilities and how they are leveraged to initially penetrate customers' infrastructures, propagate compromise across additional systems, and eventually target AD DS and domain controllers to obtain complete control of organizations' forests.

This section does not provide detailed recommendations about addressing each type of vulnerability, particularly in the areas in which the vulnerabilities are not used to directly target

Active Directory. However, for each type of vulnerability, we have provided links to additional information that you can use to develop countermeasures and reduce your organization's attack surface.

Reducing the Active Directory Attack Surface

This section begins by providing background information about privileged accounts and groups in Active Directory to provide the information that helps clarify the reasons for the subsequent recommendations for securing and managing privileged groups and accounts. We then discuss approaches to reduce the need to use highly privileged accounts for day-to-day administration, which does not require the level of privilege that is granted to groups such as the Enterprise Admins (EA), Domain Admins (DA), and Built-in Administrators (BA) groups in Active Directory. Next, we provide guidance for securing the privileged groups and accounts and for implementing secure administrative practices and systems.

Although this section provides detailed information about these configuration settings, we have also included appendices for each recommendation that provide step-by-step configuration instructions that can be used "as is" or can be modified for the organization's needs. This section finishes by providing information to securely deploy and manage domain controllers, which should be among the most stringently secured systems in the infrastructure.

Monitoring Active Directory for Signs of Compromise

Whether you have implemented robust security information and event monitoring (SIEM) in your environment or are using other mechanisms to monitor the security of the infrastructure, this section provides information that can be used to identify events on Windows systems that may indicate that an organization is being attacked. We discuss traditional and advanced audit policies, including effective configuration of audit subcategories in the Windows 7 and Windows Vista operating systems. This section includes comprehensive lists of objects and systems to audit, and an associated appendix lists events for which you should monitor if the goal is to detect compromise attempts.

Planning for Compromise

This section begins by "stepping back" from technical detail to focus on principles and processes that can be implemented to identify the users, applications, and systems that are most critical not only to the IT infrastructure, but to the business. After identifying what is most critical to the stability and operations of your organization, you can focus on segregating and securing these assets, whether they are intellectual property, people, or systems. In some cases, segregating and securing assets may be performed in your existing AD DS environment, while in other cases, you should consider implementing small, separate "cells" that allow you to establish a secure boundary around critical assets and monitor those assets more stringently than less-critical components. A concept called "creative destruction," which is a mechanism by which legacy applications and systems can be eliminated by creating new solutions is discussed, and the section ends with recommendations that can help to maintain a more secure environment by

combining business and IT information to construct a detailed picture of what is a normal operational state. By knowing what is normal for an organization, abnormalities that may indicate attacks and compromises can be more easily identified.

Summary of Best Practice Recommendations

This section provides a table that summarizes the recommendations made in this document and orders them by relative priority, in addition to providing links to where more information about each recommendation can be found in the document and its appendices.

Appendices

Appendices are included in this document to augment the information contained in the body of the document. The list of appendices and a brief description of each is included the following table.

Appendix	Description
Appendix A: Patch and Vulnerability Management Software	Contains a list of companies that produce patch and vulnerability management software.
Appendix B: Privileged Accounts and Groups in Active Directory	Provides background information that helps you to identify the users and groups you should focus on securing because they can be leveraged by attackers to compromise and even destroy your Active Directory installation.
Appendix C: Protected Accounts and Groups in Active Directory	Contains information about protected groups in Active Directory. It also contains information for limited customization (removal) of groups that are considered protected groups and are affected by AdminSDHolder and SDProp.
Appendix D: Securing Built-In Administrator Accounts in Active Directory	Contains guidelines to help secure the Administrator account in each domain in the forest.
Appendix E: Securing Enterprise Admins Groups in Active Directory	Contains guidelines to help secure the Enterprise Admins group in the forest.
Appendix F: Securing Domain Admins Groups in Active Directory	Contains guidelines to help secure the Domain Admins group in each domain in the forest.
Appendix G: Securing Administrators Groups in Active Directory	Contains guidelines to help secure the Built-in Administrators group in each domain in the forest.

Appendix H: Securing Local Administrator Accounts and Groups	Contains guidelines to help secure local Administrator accounts and Administrators groups on domain-joined servers and workstations.
Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory	Provides information to create accounts that have limited privileges and can be stringently controlled, but can be used to populate privileged groups in Active Directory when temporary elevation is required.
Appendix J: Third-Party RBAC Vendors	Contains a list of third-party RBAC vendors and the RBAC solutions they offer.
Appendix K: Third-Party PIM Vendors	Contains a list of third-party PIM vendors and the PIM solutions they offer.
Appendix L: Events to Monitor	Lists events for which you should monitor in your environment.
Appendix M: Document Links and Recommended Reading	Contains a list of recommended reading. Also contains a list of links to external documents and their URLs so that readers of hard copies of this document can access this information.

Avenues to Compromise

Law Number Seven: The most secure network is a well-administered one. – 10 Immutable Laws of Security Administration

In organizations that have experienced catastrophic compromise events, assessments usually reveal that the organizations have limited visibility into the actual state of their IT infrastructures, which may differ significantly from their "as documented" states. These variances introduce vulnerabilities that expose the environment to compromise, often with little risk of discovery until the compromise has progressed to the point at which the attackers effectively "own" the environment.

Detailed assessments of these organizations' AD DS configuration, public key infrastructures (PKIs), servers, workstations, applications, access control lists (ACLs), and other technologies reveal misconfigurations and vulnerabilities that, if remediated, could have prevented the initial compromise.

Analysis of IT documentation, processes, and procedures identifies vulnerabilities introduced by gaps in administrative practices that were leveraged by attackers to eventually obtain privileges that were used to fully compromise the Active Directory forest. A fully compromised forest is one in which attackers compromise not only individual systems, applications, or user accounts, but

escalate their access to obtain a level of privilege in which they can modify or destroy all aspects of the forest. When an Active Directory installation has been compromised to that degree, attackers can make changes that allow them to maintain a presence throughout the environment, or worse, to destroy the directory and the systems and accounts it manages.

Although a number of the commonly exploited vulnerabilities in the descriptions that follow are not attacks against Active Directory, they allow attackers to establish a foothold in an environment that can be used to run privilege escalation (also called privilege elevation) attacks and to eventually target and compromise AD DS.

This section of this document focuses on describing the mechanisms that attackers typically use to gain access to the infrastructure and eventually to launch privilege elevation attacks. Also see the following sections:

- Reducing the Active Directory Attack Surface Detailed recommendations for the secure configuration of Active Directory.
- Monitoring Active Directory for Signs of Compromise Recommendations to help detect compromise
- Planning for Compromise High-level approaches to help prepare for attacks against the infrastructure from IT and business perspectives

Notes

Although this document focuses on Active Directory and Windows systems that are part of an AD DS domain, attackers rarely focus solely on Active Directory and Windows. In environments with a mixture of operating systems, directories, applications, and data repositories, it is common to find that non-Windows systems have also been compromised. This is particularly true if the systems provide a “bridge” between Windows and non-Windows environments, such as file servers accessed by Windows and UNIX or Linux clients, directories that provide authentication services to multiple operating systems, or metadirectories that synchronize data across disparate directories.

AD DS is targeted because of the centralized access and configuration management capabilities it provides not only to Windows systems, but to other clients. Any other directory or application that provides authentication and configuration management services can, and will be targeted by determined attackers. Although this document is focused on protections that can reduce the likelihood of a compromise of Active Directory installations, every organization that includes non-Windows computers, directories, applications, or data repositories should also prepare for attacks against those systems.

Initial Breach Targets

Nobody intentionally builds an IT infrastructure that exposes the organization to compromise. When an Active Directory forest is first constructed, it is usually pristine and current. As years pass and new operating systems and applications are acquired, they’re added to the forest. As the manageability benefits that Active Directory provides are recognized, more and more content

is added to the directory, more people integrate their computers or applications with AD DS, and domains are upgraded to support new functionality offered by the most current versions of the Windows operating system. What also happens over time, however, is that even as a new infrastructure is being added, other parts of the infrastructure might not be maintained as well as they initially were, systems and applications are functioning properly and therefore are not receiving attention, and organizations begin to forget that they have not eliminated their legacy infrastructure. Based on what we see in assessing compromised infrastructures, the older, larger, and more complex the environment, the more likely it is that there are numerous instances of commonly exploited vulnerabilities.

Regardless of the motivation of the attacker, most information security breaches start with the compromise of one or two systems at a time. These initial events, or entry points into the network, often leverage vulnerabilities that could have been fixed, but were not. The 2012 Data Breach Investigations Report (DBIR), which is an annual study produced by the Verizon RISK Team in cooperation with a number of national security agencies and other companies, states that 96 percent of attacks were “not highly difficult,” and that “97 percent of breaches were avoidable through simple or intermediate controls.” These findings may be a direct consequence of the commonly exploited vulnerabilities that follow.

Gaps in Antivirus and Antimalware Deployments

Law Number Eight: An out-of-date malware scanner is only marginally better than no scanner at all. - Ten Immutable Laws of Security (Version 2.0)

Analysis of organizations’ antivirus and antimalware deployments often reveals an environment in which most workstations are configured with antivirus and antimalware software that is enabled and current. Exceptions are usually workstations that connect infrequently to the corporate environment or employee devices for which antivirus and antimalware software can be difficult to deploy, configure, and update.

Server populations, however, tend to be less consistently protected in many compromised environments. As reported in the 2012 Data Breach Investigations , 94 percent of all data compromises involved servers, which represents an 18 percent increase over the previous year, and 69 percent of attacks incorporated malware. In server populations, it is not uncommon to find that antivirus and antimalware installations are inconsistently configured, outdated, misconfigured, or even disabled. In some cases, the antivirus and antimalware software is disabled by administrative staff, but in other cases, attackers disable the software after compromising a server via other vulnerabilities. When the antivirus and antimalware software is disabled, the attackers then plant malware on the server and focus on propagating compromise across the server population.

It is important not only to ensure that your systems are protected with current, comprehensive malware protection, but also to monitor systems for disabling or removal of antivirus and antimalware software and to automatically restart protection when it is manually disabled. Although no antivirus and antimalware software can guarantee prevention and detection of all infections, a properly configured and deployed antivirus and antimalware implementation can reduce the likelihood of infection.

Incomplete Patching

Law Number Three: If you don't keep up with security fixes, your network won't be yours for long. – 10 Immutable Laws of Security Administration

Microsoft releases security bulletins on the second Tuesday of each month, although on rare occasions security updates are released between the monthly security updates (these are also known as "out-of-band" updates) when the vulnerability is determined to pose an urgent risk to customer systems. Whether a small business configures its Windows computers to use Windows Update to manage system and application patching or a large organization uses management software such as System Center Configuration Manager (SCCM) to deploy patches according to detailed, hierarchical plans, many customers patch their Windows infrastructures in a relatively timely manner.

However, few infrastructures include only Windows computers and Microsoft applications, and in compromised environments, it is common to find that the organization's patch management strategy contains gaps. Windows systems in these environments are inconsistently patched. Non-Windows operating systems are patched sporadically, if at all. Commercial off-the-shelf (COTS) applications contain vulnerabilities for which patches exist, but have not been applied. Networking devices are often configured with factory-default credentials and no firmware updates years after their installation. Applications and operating systems that are no longer supported by their vendors are often kept running, despite the fact that they can no longer be patched against vulnerabilities. Each of these unpatched systems represents another potential entry point for attackers.

The consumerization of IT has introduced additional challenges in that employee owned devices are being used to access corporate owned data, and the organization may have little to no control over the patching and configuration of employees' personal devices. Enterprise-class hardware typically ships with enterprise-ready configuration options and management capabilities, at the cost of less choice in individual customization and device selection. Employee-focused hardware offers a broader range of manufacturers, vendors, hardware security features, software security features, management capabilities and configuration options, and many enterprise features may be absent altogether.

Patch and Vulnerability Management Software

If an effective patch management system is in place for the Windows systems and Microsoft applications, part of the attack surface that unpatched vulnerabilities create has been addressed. However, unless the non-Windows systems, non-Microsoft applications, network infrastructure, and employee devices are also kept up-to-date on patches and other fixes, the infrastructure remains vulnerable. In some cases, an application's vendor may offer automatic update capabilities; in others, there may be a need to devise an approach to regularly retrieve and apply patches and other fixes. Although specific product recommendations cannot be made here, Appendix A: Patch and Vulnerability Management Software includes information about commercially available patch and vulnerability management software provided by Microsoft Partners.

Outdated Applications and Operating Systems

"You can't expect a six-year-old operating system to protect you against a six-month-old attack." – Information Security Professional with 10 years of experience securing enterprise installations

Although "get current, stay current" may sound like a marketing phrase, outdated operating systems and applications create risk in many organizations' IT infrastructures. An operating system that was released in 2003 might still be supported by the vendor and provided with updates to address vulnerabilities, but that operating system might not contain security features added in newer versions of the operating system. Outdated systems can even require weakening of certain AD DS security configuration to support the lesser capabilities of those computers.

Applications that were written to use legacy authentication protocols by vendors who are no longer supporting the application usually cannot be retooled to support stronger authentication mechanisms. However, an organization's Active Directory domain may still be configured to store LAN Manager hashes or reversibly encrypted passwords to support such applications. Applications written prior to the introduction of newer operating systems may not function well or at all on current operating systems, requiring organizations to maintain older and older systems, and in some cases, completely unsupported hardware and software.

Even in cases in which organizations have updated their domain controllers to Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, it is typical to find significant portions of the member server population to be running Windows Server 2003 (which is no longer in mainstream support), or even Windows 2000 Server or Windows NT Server 4.0 (which are completely unsupported). The longer an organization maintains aging systems, the more the disparity between feature sets grows, and the more likely it becomes that production systems will be unsupported. Additionally, the longer an Active Directory forest is maintained, the more we observe that legacy systems and applications are missed in upgrade plans. This can mean that a single computer running a single application can introduce domain- or forest-wide vulnerabilities because Active Directory is configured to support its legacy protocols and authentication mechanisms.

To eliminate legacy systems and applications, you should first focus on identifying and cataloging them, then on determining whether to upgrade or replace the application or host. Although it can be difficult to find replacements for highly specialized applications for which there is neither support nor an upgrade path, you may be able to leverage a concept called "creative destruction" to replace the legacy application with a new application that provides the necessary functionality. Planning for Compromise is described in more depth in "Planning for Compromise" later in this document.

Misconfiguration

Law Number Four: It doesn't do much good to install security fixes on a computer that was never secured to begin with. – 10 Immutable Laws of Security Administration

Even in environments where systems are generally kept current and patched, we commonly identify gaps or misconfigurations in the operating system, applications running on computers, and Active Directory. Some misconfigurations expose only the local computer to compromise, but

after a computer is “owned,” attackers usually focus on further propagating the compromise across other systems and eventually to Active Directory. Following are some of the common areas in which we identify configurations that introduce risk.

In Active Directory

The accounts in Active Directory that are most commonly targeted by attackers are those that are members of the most-highly privileged groups, such as members of the Domain Admins (DA), Enterprise Admins (EA), or built-in Administrators (BA) groups in Active Directory. The membership of these groups should be reduced to the smallest number of accounts possible so that the attack surface of these groups is limited. It is even possible to eliminate “permanent” membership in these privileged groups; that is, you can implement settings that allow you to temporarily populate these groups only when their domain- and forest-wide privileges are needed. When highly privileged accounts are used, they should be used only on designated, secure systems such as domain controllers or secure administrative hosts. Detailed information to help implement all of these configurations is provided in Reducing the Active Directory Attack Surface.

When we evaluate the membership of the highest privileged groups in Active Directory, we commonly find excessive membership in all three of the most- privileged groups. In some cases, organizations have dozens, even hundreds of accounts in DA groups. In other cases, organizations place accounts directly into built-in Administrators groups, thinking that that group is “less privileged” than the DAs group. It is not. We often find a handful of permanent members of the EA group in the forest root domain, despite the fact that EA privileges are rarely and temporarily required. Finding an IT user’s day-to-day administrative account in all three groups is also common, even though this is an effectively redundant configuration. As described in Reducing the Active Directory Attack Surface, whether an account is a permanent member of one of these groups or all of them, the account can be used to compromise, and even destroy the AD DS environment and the systems and accounts managed by it. Recommendations for the secure configuration and use of privileged accounts in Active Directory are provided in Reducing the Active Directory Attack Surface.

On Domain Controllers

When we assess domain controllers, we find often find them configured and managed no differently than member servers. Domain controllers sometimes run the same applications and utilities installed on member servers, not because they’re needed on the domain controllers, but because the applications are part of a standard build. These applications may provide minimal functionality on the domain controllers but add significantly to its attack surface by requiring configuration setting that open ports, create highly privileged service accounts, or grant access to the system by users who should not connect to a domain controller for any purpose other than authentication and Group Policy application. In some breaches, attackers have used tools that were already installed on domain controllers not only to gain access to the domain controllers, but to modify or damage the AD DS database.

When we extract the Internet Explorer® configuration settings on domain controllers, we find that users have logged on with accounts that have high levels of privilege in Active Directory and have used the accounts to access the Internet and intranet from the domain controllers. In some cases, the accounts have configured Internet Explorer settings on the domain controllers to allow download of Internet content, and freeware utilities have been downloaded from Internet sites and installed on the domain controllers. Internet Explorer Enhanced Security Configuration is enabled for Users and Administrators by default, yet we often observe that is has been disabled for Administrators. When a highly privileged account accesses the Internet and downloads content to any computer, that computer is put at severe risk. When the computer is a domain controller, the entire AD DS installation is put at risk.

Protecting Domain Controllers

Domain controllers should be treated as critical infrastructure components, secured more stringently and configured more rigidly than file, print, and application servers. Domain controllers should not run any software that is not required for the domain controller to function or doesn't protect the domain controller against attacks. Domain controllers should not be permitted to access the Internet, and security settings should be configured and enforced by Group Policy Objects (GPOs). Detailed recommendations for the secure installation, configuration, and management of domain controllers are provided in Securing Domain Controllers Against Attack.

Within the Operating System

Law Number Two: If a bad guy can alter the operating system on your computer, it's not your computer anymore. – Ten Immutable Laws of Security (Version 2.0)

Although some organizations create baseline configurations for servers of different types and allow limited customization of the operating system after it's installed, analysis of compromised environments often uncovers large numbers of servers deployed in an ad hoc fashion, and configured manually and independently. Configurations between two servers performing the same function may be completely different, where neither server is configured securely. Conversely, server configuration baselines may be consistently enforced, but also consistently misconfigured; that is, servers are configured in a manner that creates the same vulnerability on all servers of a given type. Misconfiguration includes practices such as disabling of security features, granting excessive rights and permissions to accounts (particularly service accounts), use of identical local credentials across systems, and permitting installation of unauthorized applications and utilities that create vulnerabilities of their own.

Disabling Security Features

Organizations sometimes disable Windows Firewall with Advanced Security (WFAS) out of a belief that WFAS is difficult to configure or requires work-intensive configuration. However, beginning with Windows Server 2008, when any role or feature is installed on a server, it is configured by default with the least privileges required for the role or feature to function, and the Windows Firewall is automatically configured to support the role or feature. By disabling WFAS (and not using another host-based firewall in its place), organizations increase the attack surface of the entire Windows environment. Perimeter firewalls provide some protection against attacks

that directly target an environment from the Internet, but they provide no protection against attacks that exploit other attack vectors such as drive-by download attacks, or attacks that originate from other compromised systems on the intranet.

User Account Control (UAC) settings are sometimes disabled on servers because administrative staff find the prompts intrusive. Although Microsoft Support article 2526083 describes scenarios in which UAC may be disabled on Windows Server, unless you are running a server core installation (where UAC is disabled by design), you should not disable UAC on servers without careful consideration and research.

In other cases, server settings are configured to less-secure values because organizations apply outdated server configuration settings to new operating systems, such as applying Windows Server 2003 baselines to computers running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, without changing the baselines to reflect the changes in the operating system. Rather than carrying old server baselines to new operating systems, when deploying a new operating system, review security changes and configuration settings to ensure that the settings implemented are applicable and appropriate for the new operating system.

Granting Excessive Privilege

In nearly every environment we have assessed, excessive privilege is granted to local and domain-based accounts on Windows systems. Users are granted local Administrator rights on their workstations, member servers run services that are configured with rights beyond what they need to function, and local Administrators groups across the server population contain dozens or even hundreds of local and domain accounts. Compromise of only one privileged account on a computer allows attackers to compromise the accounts of every user and service that logs on to the computer, and to harvest and leverage credentials to propagate the compromise to other systems.

Although pass-the-hash (PTH) and other credential theft attacks are ubiquitous today, it is because there is freely available tooling that makes it simple and easy to extract the credentials of other privileged accounts when an attacker has gained Administrator- or SYSTEM-level access to a computer. Even without tooling that allows harvesting of credentials from logon sessions, an attacker with privileged access to a computer can just as easily install keystroke loggers that capture keystrokes, screenshots, and clipboard contents. An attacker with privileged access to a computer can disable antimalware software, install rootkits, modify protected files, or install malware on the computer that automates attacks or turns a server into a drive-by download host.

The tactics used to extend a breach beyond a single computer vary, but the key to propagating compromise is the acquisition of highly privileged access to additional systems. By reducing the number of accounts with privileged access to any system, you reduce the attack surface not only of that computer, but the likelihood of an attacker harvesting valuable credentials from the computer.

Standardizing Local Administrator Credentials

There has long been debate among security specialists as to whether there is value in renaming local Administrator accounts on Windows computers. What is actually important about local

Administrator accounts is whether they are configured with the same user name and password across multiple computers.

If the local Administrator account is named to the same value across servers and the password assigned to the account is also configured to the same value, attackers can extract the account's credentials on one computer on which Administrator or SYSTEM-level access has been obtained. The attacker does not have to initially compromise the Administrator account; they need only compromise the account of a user who is a member of the local Administrators group, or of a service account that is configured to run as LocalSystem or with Administrator privileges. The attacker can then extract the credentials for the Administrator account and replay those credentials in network logons to other computers on the network.

As long as another computer has a local account with the same user name and password (or password hash) as the account credentials that are being presented, the logon attempt succeeds and the attacker obtains privileged access to the targeted computer. In current versions of Windows, the built-in Administrator account is disabled by default, but in legacy operating systems, the account is enabled by default.

Note

Some organizations have intentionally configured local Administrator accounts to be enabled in the belief that this provides a "failsafe" in case all other privileged accounts are locked out of a system. However, even if the local Administrator account is disabled and there are no other accounts available that can enable the account or log on to the system with Administrator privileges, the system can be booted into safe mode and the built-in local Administrator account can be re-enabled, as described in Microsoft Support article 814777. Additionally, if the system still successfully applies GPOs, a GPO can be modified to (temporarily) re-enable the Administrator account, or Restricted Groups can be configured to add a domain-based account to the local Administrators group. Repairs can be performed and the Administrator account can again be disabled. To effectively prevent a lateral compromise that uses built-in local Administrator account credentials, unique user names and passwords must be configured for local Administrator accounts. To deploy unique passwords for local Administrator accounts via a GPO, see Solution for management of built-in Administrator account's password via GPO on the MSDN website.

Permitting Installation of Unauthorized Applications

Law Number One: If a bad guy can persuade you to run his program on your computer, it's not solely your computer anymore. – Ten Immutable Laws of Security (Version 2.0)

Whether an organization deploys consistent baseline settings across servers, the installation of applications that are not part of a server's defined role should not be permitted. By allowing software to be installed that is not part of a server's designated functionality, servers are exposed to inadvertent or malicious installation of software that increases the server's attack surface, introduces application vulnerabilities, or causes system instability.

Applications

As described earlier, applications are often installed and configured to use accounts that are granted more privilege than the application actually requires. In some cases, the application's documentation specifies that service accounts must be members of a server's local Administrators group or must be configured to run in the context of the LocalSystem. This is often not because the application requires those rights, but because determining what rights and permissions an application's service accounts need requires investment in additional time and effort. If an application does not install with the minimum privileges required for the application and its configured features to function, the system is exposed to attacks that leverage application privileges without any attack against the operating system itself.

Lack of Secure Application Development Practices

Infrastructure exists to support business workloads. Where these workloads are implemented in custom applications, it is critical to ensure that the applications are developed using secure best practices. Root-cause analysis of enterprise-wide incidents often reveals that an initial compromise is effected through custom applications–particularly those that are Internet facing. Most of these compromises are accomplished via compromise of well-known attacks such as SQL injection (SQLi) and cross-site scripting (XSS) attacks.

SQL Injection is an application vulnerability that allows user-defined input to modify a SQL statement that is passed to the database for execution. This input can be provided via a field in the application, a parameter (such as the query string or a cookie), or other methods. The result of this injection is that the SQL statement provided to the database is fundamentally different than what the developer intended. Take, for example, a common query used in the evaluation of a user name/password combination:

```
SELECT userID FROM users WHERE username = 'sUserName' AND password = 'sPassword'
```

When this is received by the database server, it instructs the server to look through the users table and return any userID record where the user name and password match those provided by the user (presumably via a login form of some kind). Naturally the intent of the developer in this case is to only return a valid record if a correct user name and password can be provided by the user. If either is incorrect, the database server will be unable to find a matching record and return an empty result.

The issue occurs when an attacker does something unexpected such as providing their own SQL in place of valid data. Because SQL is interpreted on-the-fly by the database server, the injected code would be processed as if the developer had put it in himself. For example, if the attacker entered **administrator** for the user ID and **xyz** OR **1=1** as the password, the resulting statement processed by the database would be:

```
SELECT userID FROM users WHERE username = 'administrator' AND password = 'xyz' OR 1=1
```

When this query is processed by the database server, all rows in the table will be returned in the query because 1=1 will always evaluate to True, thus it doesn't matter if the correct username and password is known or provided. The net result in most cases is that the user will be logged on as the first user in the user's database; in most cases, this will be the administrative user.

In addition to simply logging on, malformed SQL statements such as this can be used to add, delete, or change data, or even drop (delete) entire tables from a database. In the most extreme cases where SQLi is combined with excessive privilege, operating system commands can be run to enable the creation of new users, to download attack tools, or to take any other actions of the attackers choosing.

In cross-site scripting, the vulnerability is introduced in the application's output. An attack begins with an attacker providing malformed data to the application, but in this case the malformed data is in the form of scripting code (such as JavaScript) that will be run by the victim's browser. Exploit of an XSS vulnerability can allow an attacker to run any functions of the target application in the context of the user who launched the browser. XSS attacks are typically initiated by a phishing email encouraging the user to click a link that connects to the application and runs the attack code.

XSS is often exploited in online banking and e-commerce scenarios where an attacker can make purchases or transfer money in the context of the exploited user. In the case of a targeted attack on a custom web-based identity management application, it can allow an attacker to create their own identities, modify permissions and rights, and lead to a systemic compromise.

Although a full discussion of cross-site scripting and SQL injection is outside the scope of this document, the Open Web Application Security Project (OWASP) publishes a top 10 list with in-depth discussion of the vulnerabilities and countermeasures.

Regardless of the investment in infrastructure security, if poorly designed and written applications are deployed within that infrastructure, the environment is made vulnerable to attacks. Even well-secured infrastructures often cannot provide effective countermeasures to these application attacks. Compounding the problem, poorly designed applications may require that service accounts be granted excessive permissions for the application to function.

The Microsoft Security Development Lifecycle (SDL) is a set of structural process controls that work to improve security beginning early in requirements gathering and extending through the lifecycle of the application until it is decommissioned. This integration of effective security controls is not only critical from a security perspective, it is critical to ensure that application security is cost and schedule effective. Assessing an application for security issues when it is effectively code complete requires organizations to make decisions about application security only before or even after the application has been deployed. An organization can choose to address the application flaws before deploying the application in production, incurring costs and delays, or the application can be deployed in production with known security flaws, exposing the organization to compromise.

Some organizations place the full cost of fixing a security issue in production code above $10,000 per issue, and applications developed without an effective SDL can average more than ten high-severity issues per 100,000 lines of code. In large applications, the costs escalate quickly. By contrast, many companies set a benchmark of less than one issue per 100,000 lines of code at the final code review stage of the SDL, and aim for zero issues in high-risk applications in production.

Implementing the SDL improves security by including security requirements early in requirements gathering and design of an application provides threat modeling for high-risk applications;

requires effective training and monitoring of developers; and requires clear, consistent code standards and practices. The net effect of an SDL is significant improvements in application security while reducing the cost to develop, deploy, maintain, and decommission an application. Although a detailed discussion of the design and implementation of SDL is beyond the scope of this document, refer to the Microsoft Security Development Lifecycle for detailed guidance and information.

Attractive Accounts for Credential Theft

Credential theft attacks are those in which an attacker initially gains highest-privilege (root, Administrator, or SYSTEM, depending on the operating system in use) access to a computer on a network and then uses freely available tooling to extract credentials from the sessions of other logged-on accounts. Depending on the system configuration, these credentials can be extracted in the form of hashes, tickets, or even plaintext passwords. If any of the harvested credentials are for local accounts that are likely to exist on other computers on the network (for example, Administrator accounts in Windows, or root accounts in OSX, UNIX, or Linux), the attacker presents the credentials to other computers on the network to propagate compromise to additional computers and to try to obtain the credentials of two specific types of accounts:

1. Privileged domain accounts with both broad and deep privileges (that is, accounts that have administrator-level privileges on many computers and in Active Directory). These accounts may not be members of any of the highest-privilege groups in Active Directory, but they may have been granted Administrator-level privilege across many servers and workstations in the domain or forest, which makes them effectively as powerful as members of privileged groups in Active Directory. In most cases, accounts that have been granted high levels of privilege across broad swaths of the Windows infrastructure are service accounts, so service accounts should always be assessed for breadth and depth of privilege.
2. “Very Important Person” (VIP) domain accounts. In the context of this document, a VIP account is any account that has access to information an attacker wants (intellectual property and other sensitive information), or any account that can be used to grant the attacker access to that information. Examples of these user accounts include:
    a. Executives whose accounts have access to sensitive corporate information
    b. Accounts for Help Desk staff who are responsible for maintaining the computers and applications used by executives
    c. Accounts for legal staff who have access to an organization’s bid and contract documents, whether the documents are for their own organization or client organizations
    d. Product planners who have access to plans and specifications for products in an company’s development pipeline, regardless of the types of products the company makes
    e. Researchers whose accounts are used to access study data, product formulations, or any other research of interest to an attacker

Because highly privileged accounts in Active Directory can be used to propagate compromise and to manipulate VIP accounts or the data that they can access, the most useful accounts for

credential theft attacks are accounts that are members of Enterprise Admins, Domain Admins, and Administrators groups in Active Directory.

Because domain controllers are the repositories for the AD DS database and domain controllers have full access to all of the data in Active Directory, domain controllers are also targeted for compromise, whether in parallel with credential theft attacks, or after one or more highly privileged Active Directory accounts have been compromised. Although numerous publications (and many attackers) focus on the Domain Admins group memberships when describing pass-the-hash and other credential theft attacks (as is described in Reducing the Active Directory Attack Surface), an account that is a member of any of the groups listed here can be used to compromise the entire AD DS installation.

Note

For comprehensive information about pass-the-hash and other credential theft attacks, please see the Mitigating Pass-the-Hash (PTH) Attacks and Other Credential Theft Techniques whitepaper listed in Appendix M: Document Links and Recommended Reading. For more information about attacks by determined adversaries, which are sometimes referred to as “advanced persistent threats” (APTs), please see Determined Adversaries and Targeted Attacks.

Activities that Increase the Likelihood of Compromise

Because the target of credential theft is usually highly privileged domain accounts and VIP accounts, it is important for administrators to be conscious of activities that increase the likelihood of success of a credential-theft attack. Although attackers also target VIP accounts, if VIPs are not given high levels of privilege on systems or in the domain, theft of their credentials requires other types of attacks, such as socially engineering the VIP to provide secret information. Or the attacker must first obtain privileged access to a system on which VIP credentials are cached. Because of this, activities that increase the likelihood of credential theft described here are focused primarily on preventing the acquisition of highly privileged administrative credentials. These activities are common mechanisms by which attackers are able to compromise systems to obtain privileged credentials.

Logging on to Unsecured Computers with Privileged Accounts

The core vulnerability that allows credential theft attacks to succeed is the act of logging on to computers that are not secure with accounts that are broadly and deeply privileged throughout the environment. These logons can be the result of various misconfigurations described here.

Not Maintaining Separate Administrative Credentials

Although this is relatively uncommon, in assessing various AD DS installations, we have found IT employees using a single account for all of their work. The account is a member of at least one of the most highly privileged groups in Active Directory and is the same account that the employees

use to log on to their workstations in the morning, check their email, browse Internet sites, and download content to their computers. When users run with accounts that are granted local Administrator rights and permissions, they expose the local computer to complete compromise. When those accounts are also members of the most privileged groups in Active Directory, they expose the entire forest to compromise, making it trivially easy for an attacker to gain complete control of the Active Directory and Windows environment.

Similarly, in some environments, we've found that the same user names and passwords are used for root accounts on non-Windows computers as are used in the Windows environment, which allows attackers to extend compromise from UNIX or Linux systems to Windows systems and vice versa.

Logons to Compromised Workstations or Member Servers with Privileged Accounts

When a highly privileged domain account is used to log on interactively to a compromised workstation or member server, that compromised computer may harvest credentials from any account that logs on to the system.

Unsecured Administrative Workstations

In many organizations, IT staff use multiple accounts. One account is used for logon to the employee's workstation, and because these are IT staff, they often have local Administrator rights on their workstations. In some cases, UAC is left enabled so that the user at least receives a split access token at logon and must elevate when privileges are required. When these users are performing maintenance activities, they typically use locally installed management tools and provide the credentials for their domain-privileged accounts, by selecting the **Run as Administrator** option or by providing the credentials when prompted. Although this configuration may seem appropriate, it exposes the environment to compromise because:

- The "regular" user account that the employee uses to log on to their workstation has local Administrator rights, the computer is vulnerable to drive-by download attacks in which the user is convinced to install malware.
- The malware is installed in the context of an administrative account, the computer can now be used to capture keystrokes, clipboard contents, screenshots, and memory-resident credentials, any of which can result in exposure of the credentials of a powerful domain account.

The problems in this scenario are twofold. First, although separate accounts are used for local and domain administration, the computer is unsecured and does not protect the accounts against theft. Second, the regular user account and the administrative account have been granted excessive rights and permissions.

Browsing the Internet with a Highly Privileged Account

Users who log on to computers with accounts that are members of the local Administrators group on the computer, or members of privileged groups in Active Directory, and who then browse the Internet (or a compromised intranet) expose the local computer and the directory to compromise.

Accessing a maliciously crafted website with a browser running with administrative privileges can allow an attacker to deposit malicious code on the local computer in the context of the privileged user. If the user has local Administrator rights on the computer, attackers may deceive the user into downloading malicious code or opening email attachments that leverage application vulnerabilities and leverage the user's privileges to extract locally cached credentials for all active users on the computer. If the user has administrative rights in the directory by membership in the Enterprise Admins, Domain Admins, or Administrators groups in Active Directory, the attacker can extract the domain credentials and use them to compromise the entire AD DS domain or forest, without needing to compromise any other computer in the forest.

Configuring Local Privileged Accounts with the Same Credentials across Systems

Configuring the same local Administrator account name and password on many or all computers enables credentials stolen from the SAM database on one computer to be used to compromise all other computers that use the same credentials. At a minimum, you should use different passwords for local Administrator accounts across each domain-joined system. Local Administrator accounts may also be uniquely named, but using different passwords for each system's privileged local accounts is sufficient to ensure that credentials cannot be used on other systems.

Overpopulation and Overuse of Privileged Domain Groups

Granting membership in the EA, DA, or BA groups in a domain creates a target for attackers. The greater the number of members of these groups, the greater the likelihood that a privileged user may inadvertently misuse the credentials and expose them to credential theft attacks. Every workstation or server to which a privileged domain user logs on presents a possible mechanism by which the privileged user's credentials may be harvested and used to compromise the AD DS domain and forest.

Poorly Secured Domain Controllers

Domain controllers house a replica of a domain's AD DS database. In the case of read-only domain controllers, the local replica of the database contains the credentials for only a subset of the accounts in the directory, none of which are privileged domain accounts by default. On read-write domain controllers, each domain controller maintains a full replica of the AD DS database, including credentials not only for privileged users like Domain Admins, but privileged accounts such as domain controller accounts or the domain's Krbtgt account, which is the account that is associated with the KDC service on domain controllers. If additional applications that are not necessary for domain controller functionality are installed on domain controllers, or if domain

controllers are not stringently patched and secured, attackers may compromise them via unpatched vulnerabilities, or they may leverage other attack vectors to install malicious software directly on them.

Privilege Elevation and Propagation

Regardless of the attack methods used, Active Directory is always targeted when a Windows environment is attacked, because it ultimately controls access to whatever the attackers want. This does not mean that the entire directory is targeted, however. Specific accounts, servers, and infrastructure components are usually the primary targets of attacks against Active Directory. These accounts are described as follows.

Permanent Privileged Accounts

Because the introduction of Active Directory, it has been possible to use highly privileged accounts to build the Active Directory forest and then to delegate rights and permissions required to perform day-to-day administration to less-privileged accounts. Membership in the Enterprise Admins, Domain Admins, or Administrators groups in Active Directory is required only temporarily and infrequently in an environment that implements least-privilege approaches to daily administration.

Permanent privileged accounts are accounts that have been placed in privileged groups and left there from day to day. If your organization places five accounts into the Domain Admins group for a domain, those five accounts can be targeted 24-hours a day, seven days a week. However, the actual need to use accounts with Domain Admins privileges is typically only for specific domain-wide configuration, and for short periods of time.

VIP Accounts

An often overlooked target in Active Directory breaches is the accounts of “very important persons” (or VIPs) in an organization. Privileged accounts are targeted because those accounts can grant access to attackers, which allows them to compromise or even destroy targeted systems, as described earlier in this section.

“Privilege-Attached” Active Directory Accounts

“Privilege-attached” Active Directory accounts are domain accounts that have not been made members of any of the groups that have the highest levels of privilege in Active Directory, but have instead been granted high levels of privilege on many servers and workstations in the environment. These accounts are most often domain-based accounts that are configured to run services on domain-joined systems, typically for applications running on large sections of the infrastructure. Although these accounts have no privileges in Active Directory, if they are granted high privilege on large numbers of systems, they can be used to compromise or even destroy large segments of the infrastructure, achieving the same effect as compromise of a privileged Active Directory account.

Reducing the Active Directory Attack Surface

This section focuses on technical controls to implement to reduce the attack surface of the Active Directory installation. The section contains the following information:

- Implementing Least-Privilege Administrative Models focuses on identifying the risk that the use of highly privileged accounts for day-to-day administration presents, in addition to providing recommendations to implement to reduce the risk that privileged accounts present.
- Implementing Secure Administrative Hosts describes principles for deployment of dedicated, secure administrative systems, in addition to some sample approaches to a secure administrative host deployment.
- Securing Domain Controllers Against Attack discusses policies and settings that, although similar to the recommendations for the implementation of secure administrative hosts, contain some domain controller-specific recommendations to help ensure that the domain controllers and the systems used to manage them are well-secured.

Privileged Accounts and Groups in Active Directory

This section provides background information about privileged accounts and groups in Active Directory intended to explain the commonalities and differences between privileged accounts and groups in Active Directory. By understanding these distinctions, whether you implement the recommendations in Implementing Least-Privilege Administrative Models verbatim or choose to customize them for your organization, you have the tools you need to secure each group and account appropriately.

Built-in Privileged Accounts and Groups

Active Directory facilitates delegation of administration and supports the principle of least privilege in assigning rights and permissions. “Regular” users who have accounts in a domain are, by default, able to read much of what is stored in the directory, but are able to change only a very limited set of data in the directory. Users who require additional privilege can be granted membership in various “privileged” groups that are built into the directory so that they may perform specific tasks related to their roles, but cannot perform tasks that are not relevant to their duties. Organizations can also create groups that are tailored to specific job responsibilities and are granted granular rights and permissions that allow IT staff to perform day-to-day administrative functions without granting rights and permissions that exceed what is required for those functions.

Within Active Directory, three built-in groups are the highest privilege groups in the directory: Enterprise Admins, Domain Admins, and Administrators. The default configuration and capabilities of each of these groups are described in the following sections:

Highest Privilege Groups in Active Directory

Enterprise Admins

Enterprise Admins (EA) is a group that exists only in the forest root domain, and by default, it is a member of the Administrators group in all domains in the forest. The built-in Administrator account in the forest root domain is the only default member of the EA group. EAs are granted rights and permissions that allow them to implement forest-wide changes (that is, changes that affect all domains in the forest), such as adding or removing domains, establishing forest trusts, or raising forest functional levels. In a properly designed and implemented delegation model, EA membership is required only when first constructing the forest or when making certain forest-wide changes such as establishing an outbound forest trust. Most of the rights and permissions granted to the EA group can be delegated to lesser-privileged users and groups.

Domain Admins

Each domain in a forest has its own Domain Admins (DA) group, which is a member of that domain's Administrators group and a member of the local Administrators group on every computer that is joined to the domain. The only default member of the DA group for a domain is the built-in Administrator account for that domain. DAs are "all-powerful" within their domains, while EAs have forest-wide privilege. In a properly designed and implemented delegation model, Domain Admins membership should be required only in "break glass" scenarios (such as situations in which an account with high levels of privilege on every computer in the domain is needed). Although native Active Directory delegation mechanisms allow delegation to the extent that it is possible to use DA accounts only in emergency scenarios, constructing an effective delegation model can be time consuming, and many organizations leverage third-party tools to expedite the process.

Administrators

The third group is the built-in domain local Administrators (BA) group into which DAs and EAs are nested. This group is granted many of the direct rights and permissions in the directory and on domain controllers. However, the Administrators group for a domain has no privileges on member servers or on workstations. It is via membership in the computers' local Administrators group that local privilege is granted.

Note

Although these are the default configurations of these privileged groups, a member of any of the three groups can manipulate the directory to gain membership in any of the other groups. In some cases, it is trivial to obtain membership in the other groups, while in others it is more difficult, but from the perspective of potential privilege, all three groups should be considered effectively equivalent.

Schema Admins

A fourth privileged group, Schema Admins (SA), exists only in the forest root domain and has only that domain's built-in Administrator account as a default member, similar to the Enterprise

Admins group. The Schema Admins group is intended to be populated only temporarily and occasionally (when modification of the AD DS schema is required).

Although the SA group is the only group that can modify the Active Directory schema (that is., the directory's underlying data structures such as objects and attributes), the scope of the SA group's rights and permissions is more limited than the previously described groups. It is also common to find that organizations have developed appropriate practices for the management of the membership of the SA group because membership in the group is typically infrequently needed, and only for short periods of time. This is technically true of the EA, DA, and BA groups in Active Directory, as well, but it is far less common to find that organizations have implemented similar practices for these groups as for the SA group.

Protected Accounts and Groups in Active Directory

Within Active Directory, a default set of privileged accounts and groups called "protected" accounts and groups are secured differently than other objects in the directory. Any account that has direct or transitive membership in any protected group (regardless of whether the membership is derived from security or distribution groups) inherits this restricted security.

For example, if a user is a member of a distribution group that is, in turn, a member of a protected group in Active Directory, that user object is flagged as a protected account. When an account is flagged as a protected account, the value of the adminCount attribute on the object is set to 1.

Note

Although transitive membership in a protected group includes nested distribution and nested security groups, accounts that are members of nested distribution groups will not receive the protected group's SID in their access tokens. However, distribution groups can be converted to security groups in Active Directory, which is why distribution groups are included in protected group member enumeration. Should a protected nested distribution group ever be converted to a security group, the accounts that are members of the former distribution group will subsequently receive the parent protected group's SID in their access tokens at the next logon.

The following table lists the default protected accounts and groups in Active Directory by operating system version and service pack level.

Default Protected Accounts and Groups in Active Directory by Operating System and Service Pack (SP) Version

Windows 2000 <SP4	Windows 2000 SP4 – Windows Server 2003	Windows Server 2003 SP1+	Windows Server 2008 – Windows Server 2012
Administrators	Account Operators	Account Operators	Account Operators
	Administrator	Administrator	Administrator

	Administrators	Administrators	Administrators
Domain Admins	Backup Operators	Backup Operators	Backup Operators
	Cert Publishers		
	Domain Admins	Domain Admins	Domain Admins
Enterprise Admins	Domain Controllers	Domain Controllers	Domain Controllers
	Enterprise Admins	Enterprise Admins	Enterprise Admins
	Krbtgt	Krbtgt	Krbtgt
	Print Operators	Print Operators	Print Operators
			Read-only Domain Controllers
	Replicator	Replicator	Replicator
Schema Admins	Schema Admins		Schema Admins

AdminSDHolder and SDProp

In the System container of every Active Directory domain, an object called AdminSDHolder is automatically created. The purpose of the AdminSDHolder object is to ensure that the permissions on protected accounts and groups are consistently enforced, regardless of where the protected groups and accounts are located in the domain.

Every 60 minutes (by default), a process known as Security Descriptor Propagator (SDProp) runs on the domain controller that holds the domain's PDC Emulator role. SDProp compares the permissions on the domain's AdminSDHolder object with the permissions on the protected accounts and groups in the domain. If the permissions on any of the protected accounts and groups do not match the permissions on the AdminSDHolder object, the permissions on the protected accounts and groups are reset to match those of the domain's AdminSDHolder object.

Permissions inheritance is disabled on protected groups and accounts, which means that even if the accounts or groups are moved to different locations in the directory, they do not inherit permissions from their new parent objects. Inheritance is also disabled on the AdminSDHolder object so that permissions changes to the parent objects do not change the permissions of AdminSDHolder.

Note

When an account is removed from a protected group, it is no longer considered a protected account, but its adminCount attribute remains set to 1 if it is not manually changed. The result of this configuration is that the object's ACLs are no longer updated by SDProp, but the object still does not inherit permissions from its parent object. Therefore, the object may reside in an organizational unit (OU) to which permissions have been delegated, but the formerly protected object will not inherit these delegated

permissions. A script to locate and reset formerly protected objects in the domain can be found in the Microsoft Support article 817433.

AdminSDHolder Ownership

Most objects in Active Directory are owned by the domain's BA group. However, the AdminSDHolder object is, by default, owned by the domain's DA group. (This is a circumstance in which DAs do not derive their rights and permissions via membership in the Administrators group for the domain.)

In versions of Windows earlier than Windows Server 2008, owners of an object can change permissions of the object, including granting themselves permissions that they did not originally have. Therefore, the default permissions on a domain's AdminSDHolder object prevent users who are members of BA or EA groups from changing the permissions for a domain's AdminSDHolder object. However, members of the Administrators group for the domain can take ownership of the object and grant themselves additional permissions, which means that this protection is rudimentary and only protects the object against accidental modification by users who are not members of the DA group in the domain. Additionally, the BA and EA (where applicable) groups have permission to change the attributes of the AdminSDHolder object in the local domain (root domain for EA).

Note

An attribute on the AdminSDHolder object, dSHeuristics, allows limited customization (removal) of groups that are considered protected groups and are affected by AdminSDHolder and SDProp. This customization should be carefully considered if it is implemented, although there are valid circumstances in which modification of dSHeuristics on AdminSDHolder is useful. More information about modification of the dSHeuristics attribute on an AdminSDHolder object can be found in the Microsoft Support articles 817433 and 973840, and in Appendix C: Protected Accounts and Groups in Active Directory.

Although the most privileged groups in Active Directory are described here, there are a number of other groups that have been granted elevated levels of privilege. For more information about all of the default and built-in groups in Active Directory and the user rights assigned to each, see Appendix B: Privileged Accounts and Groups in Active Directory.

Implementing Least-Privilege Administrative Models

The following excerpt is from The Administrator Accounts Security Planning Guide, first published on April 1, 1999:

The following excerpt is from the Microsoft Windows Security Resource Kit, first published in 2005:

The Privilege Problem

The principles described in the preceding excerpts have not changed, but in assessing Active Directory installations, we invariably find excessive numbers of accounts that have been granted rights and permissions far beyond those required to perform day-to-day work. The size of the environment affects the raw numbers of overly privileged accounts, but not the proportion—midsized directories may have dozens of accounts in the most highly privileged groups, while large installations may have hundreds or even thousands. With few exceptions, regardless of the sophistication of an attacker's skills and arsenal, attackers typically follow the path of least resistance. They increase the complexity of their tooling and approach only if and when simpler mechanisms fail or are thwarted by defenders.

Unfortunately, the path of least resistance in many environments has proven to be the overuse of accounts with broad and deep privilege. Broad privileges are rights and permissions that allow an account to perform specific activities across a large cross-section of the environment- for example, Help Desk staff may be granted permissions that allow them to reset the passwords on many user accounts.

Deep privileges are powerful privileges that are applied to a narrow segment of the population, such giving an engineer Administrator rights on a server so that they can perform repairs. Neither broad privilege nor deep privilege is necessarily dangerous, but when many accounts in the domain are permanently granted broad and deep privilege, if only one of the accounts is compromised, it can quickly be used to reconfigure the environment to the attacker's purposes or even to destroy large segments of the infrastructure.

Pass-the-hash attacks, which are a type of credential theft attack, are ubiquitous because the tooling to perform them is freely available and easy-to-use, and because many environments are vulnerable to the attacks. Pass-the-hash attacks, however, are not the real problem. The crux of the problem is twofold:

1. It is usually easy for an attacker to obtain deep privilege on a single computer and then propagate that privilege broadly to other computers.
2. There are usually too many permanent accounts with high levels of privilege across the computing landscape.

Even if pass-the-hash attacks are eliminated, attackers would simply use different tactics, not a different strategy. Rather than planting malware that contains credential theft tooling, they might plant malware that logs keystrokes, or leverage any number of other approaches to capture credentials that are powerful across the environment. Regardless of the tactics, the targets remain the same: accounts with broad and deep privilege.

Granting of excessive privilege isn't only found in Active Directory in compromised environments. When an organization has developed the habit of granting more privilege than is required, it is typically found throughout the infrastructure as discussed in the following sections.

In Active Directory

In Active Directory, it is common to find that the EA, DA and BA groups contain excessive numbers of accounts. Most commonly, an organization's EA group contains the fewest members, DA groups usually contain a multiplier of the number of users in the EA group, and Administrators groups usually contain more members than the populations of the other groups combined. This is often due to a belief that Administrators are somehow "less privileged" than DAs or EAs. While the rights and permissions granted to each of these groups differ, they should be effectively considered equally powerful groups because a member of one can make himself or herself a member of the other two.

On Member Servers

When we retrieve the membership of local Administrators groups on member servers in many environments, we find membership ranging from a handful of local and domain accounts, to dozens of nested groups that, when expanded, reveal hundreds, even thousands, of accounts with local Administrator privilege on the servers. In many cases, domain groups with large memberships are nested in member servers' local Administrators groups, without consideration to the fact that any user who can modify the memberships of those groups in the domain can gain administrative control of all systems on which the group has been nested in a local Administrators group.

On Workstations

Although workstations typically have significantly fewer members in their local Administrators groups than member servers do, in many environments, users are granted membership in the local Administrators group on their personal computers. When this occurs, even if UAC is enabled, those users present an elevated risk to the integrity of their workstations.

Important

You should consider carefully whether users require administrative rights on their workstations, and if they do, a better approach may be to create a separate local account on the computer that is a member of the Administrators group. When users require elevation, they can present the credentials of that local account for elevation, but because the account is local, it cannot be used to compromise other computers or access domain resources. As with any local accounts, however, the credentials for the local privileged account should be unique; if you create a local account with the same credentials on multiple workstations, you expose the computers to pass-the-hash attacks.

In Applications

In attacks in which the target is an organization's intellectual property, accounts that have been granted powerful privileges within applications can be targeted to allow exfiltration of data. Although the accounts that have access to sensitive data may have been granted no elevated privileges in the domain or the operating system, accounts that can manipulate the configuration of an application or access to the information the application provides present risk.

In Data Repositories

As is the case with other targets, attackers seeking access to intellectual property in the form of documents and other files can target the accounts that control access to the file stores, accounts that have direct access to the files, or even groups or roles that have access to the files. For example, if a file server is used to store contract documents and access is granted to the documents by the use of an Active Directory group, an attacker who can modify the membership of the group can add compromised accounts to the group and access the contract documents. In cases in which access to documents is provided by applications such as SharePoint, attackers can target the applications as described earlier.

Reducing Privilege

The larger and more complex an environment, the more difficult it is to manage and secure. In small organizations, reviewing and reducing privilege may be a relatively simple proposition, but each additional server, workstation, user account, and application in use in an organization adds another object that must be secured. Because it can be difficult or even impossible to properly secure every aspect of an organization's IT infrastructure, you should focus efforts first on the accounts whose privilege create the greatest risk, which are typically the built-in privileged accounts and groups in Active Directory, and privileged local accounts on workstations and member servers.

Securing Local Administrator Accounts on Workstations and Member Servers

Although this document focuses on securing Active Directory, as has been previously discussed, most attacks against the directory begin as attacks against individual hosts. Full guidelines for securing local groups on member systems cannot be provided, but the following recommendations can be used to help you secure the local Administrator accounts on workstations and member servers.

Securing Local Administrator Accounts

On all versions of Windows currently in mainstream support, the local Administrator account is disabled by default, which makes the account unusable for pass-the-hash and other credential theft attacks. However, in domains containing legacy operating systems or in which local

Administrator accounts have been enabled, these accounts can be used as previously described to propagate compromise across member servers and workstations. For this reason, the following controls are recommended for all local Administrator accounts on domain-joined systems.

Detailed instructions for implementing these controls are provided in Appendix H: Securing Local Administrator Accounts and Groups. Before implementing these settings, however, ensure that local Administrator accounts are not currently used in the environment to run services on computers or perform other activities for which these accounts should not be used. Test these settings thoroughly before implementing them in a production environment.

Controls for Local Administrator Accounts

Built-in Administrator accounts should never be used as service accounts on member servers, nor should they be used to log on to local computers (except in Safe Mode, which is permitted even if the account is disabled). The goal of implementing the settings described here is to prevent each computer's local Administrator account from being usable unless protective controls are first reversed. By implementing these controls and monitoring Administrator accounts for changes, you can significantly reduce the likelihood of success of an attack that targets local Administrator accounts.

Configuring GPOs to Restrict Administrator Accounts on Domain-Joined Systems

In one or more GPOs that you create and link to workstation and member server OUs in each domain, add the Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:

- Deny access to this computer from the network
- Deny log on as a batch job
- Deny log on as a service
- Deny log on through Remote Desktop Services

When you add Administrator accounts to these user rights, specify whether you are adding the local Administrator account or the domain's Administrator account by the way that you label the account. For example, to add the NWTRADERS domain's Administrator account to these deny rights, you would type the account as **NWTRADERS\Administrator**, or browse to the Administrator account for the NWTRADERS domain. To ensure that you restrict the local Administrator account, type **Administrator** in these user rights settings in the Group Policy Object Editor.

Note

Even if local Administrator accounts are renamed, the policies will still apply.

These settings will ensure that a computer's Administrator account cannot be used to connect to the other computers, even if it is inadvertently or maliciously enabled. Local logons using the local Administrator account cannot be completely disabled, nor should you attempt to do so, because a computer's local Administrator account is designed to be used in disaster recovery scenarios.

Should a member server or workstation become disjoined from the domain with no other local accounts granted administrative privileges, the computer can be booted into safe mode, the Administrator account can be enabled, and the account can then be used to effect repairs on the computer. When repairs are completed, the Administrator account should again be disabled.

Securing Local Privileged Accounts and Groups in Active Directory

Law Number Six: A computer is only as secure as the administrator is trustworthy. – Ten Immutable Laws of Security (Version 2.0)

The information provided here is intended to give general guidelines for securing the highest privilege built-in accounts and groups in Active Directory. Detailed step-by-step instructions are also provided in Appendix D: Securing Built-In Administrator Accounts in Active Directory, Appendix E: Securing Enterprise Admins Groups in Active Directory, Appendix F: Securing Domain Admins Groups in Active Directory, and in Appendix G: Securing Administrators Groups in Active Directory.

Before you implement any of these settings, you should also test all settings thoroughly to determine if they are appropriate for your environment. Not all organizations will be able to implement these settings.

Securing Built-in Administrator Accounts in Active Directory

In each domain in Active Directory, an Administrator account is created as part of the creation of the domain. This account is by default a member of the Domain Admins and Administrator groups in the domain, and if the domain is the forest root domain, the account is also a member of the Enterprise Admins group. Use of a domain's local Administrator account should be reserved only for initial build activities and, possibly, disaster-recovery scenarios. To ensure that a built-in Administrator account can be used to effect repairs in the event that no other accounts can be used, you should not change the default membership of the Administrator account in any domain in the forest. Instead, you should following guidelines to help secure the Administrator account in each domain in the forest. Detailed instructions for implementing these controls are provided in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

Controls for Built-in Administrator Accounts

The goal of implementing the settings described here is to prevent each domain's Administrator account (not a group) from being usable unless a number of controls are reversed. By implementing these controls and monitoring the Administrator accounts for changes, you can significantly reduce the likelihood of a successful attack by leveraging a domain's Administrator account. For the Administrator account in each domain in your forest, you should configure the following settings.

Enable the "Account is sensitive and cannot be delegated" flag on the account

By default, all accounts in Active Directory can be delegated. Delegation allows a computer or service to present the credentials for an account that has authenticated to the computer or service to other computers to obtain services on behalf of the account. When you enable the **Account is sensitive and cannot be delegated** attribute on a domain-based account, the account's credentials cannot be presented to other computers or services on the network, which limits attacks that leverage delegation to use the account's credentials on other systems.

Enable the "Smart card is required for interactive logon" flag on the account

When you enable the **Smart card is required for interactive logon** attribute on an account, Windows resets the account's password to a 120-character random value. By setting this flag on built-in Administrator accounts, you ensure that the password for the account is not only long and complex, but is not known to any user. It is not technically necessary to create smart cards for the accounts before enabling this attribute, but if possible, smart cards should be created for each Administrator account prior to configuring the account restrictions and the smart cards should be stored in secure locations.

Although setting the **Smart card is required for interactive logon** flag resets the account's password, it does not prevent a user with rights to reset the account's password from setting the account to a known value and using the account's name and new password to access resources on the network. Because of this, you should implement the following additional controls on the account.

Disable the Account

If the Administrator account is not already disabled, disable it when you have completed configuration of the account's properties. This prevents the account from being used for any purpose unless it is first enabled. In a disaster recovery scenario in which no accounts are available to perform repairs of the AD DS environment, you can boot a domain controller into safe mode, log on locally with the built-in Administrator account (which is never blocked from local logon), and enable the domain's Administrator account, if necessary.

Configuring GPOs to Restrict Domains' Administrator Accounts on Domain-Joined Systems

Although disabling the Administrator account in a domain makes the account effectively unusable, you should implement additional restrictions on the account in case the account is inadvertently or maliciously enabled. Although these controls can ultimately be reversed by the Administrator account, the goal is to create controls that slow an attacker's progress and limit the damage the account can inflict.

In one or more GPOs that you create and link to workstation and member server OUs in each domain, add each domain's Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:

- Deny access to this computer from the network
- Deny log on as a batch job

- Deny log on as a service
- Deny log on through Remote Desktop Services

Notes

When you add local Administrator accounts to this setting, you must specify whether you are configuring local Administrator accounts or domain Administrator accounts. For example, to add the NWTRADERS domain's local Administrator account to these deny rights, you must either type the account as **NWTRADERS\Administrator**, or browse to the local Administrator account for the NWTRADERS domain. If you type **Administrator** in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied.

We recommend restricting local Administrator accounts on member servers and workstations in the same manner as domain-based Administrator accounts. Therefore, you should generally add the Administrator account for each domain in the forest and the Administrator account for the local computers to these user rights settings. The following screenshot shows an example of configuring these user rights to block local Administrator accounts and a domain's Administrator account from performing logons that should not be needed for these accounts.

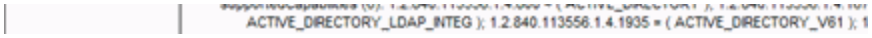

Configuring GPOs to Restrict Administrator Accounts on Domain Controllers

In each domain in the forest, the Default Domain Controllers policy or a policy linked to the Domain Controllers OU should be modified to add each domain's Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:

- Deny access to this computer from the network
- Deny log on as a batch job
- Deny log on as a service
- Deny log on through Remote Desktop Services

Note

These settings will ensure that the local Administrator account cannot be used to connect to a domain controller, although the account, if enabled, can log on locally to domain controllers. Because this account should only be enabled and used in disaster-recovery scenarios, it is anticipated that physical access to at least one domain controller will be available, or that other accounts with permissions to access domain controllers remotely can be used.

Configure Auditing of Built-in Administrator Accounts

When you have secured each domain's Administrator account and disabled it, you should configure auditing to monitor for changes to the account. If the account is enabled, its password is reset, or any other modifications are made to the account, alerts should be sent to the users or teams responsible for administration of AD DS, in addition to incident response teams in your organization.

Securing Administrators, Domain Admins and Enterprise Admins Groups

Securing Enterprise Admin Groups

The Enterprise Admins group, which is housed in the forest root domain, should contain no users on a day-to-day basis, with the possible exception of the domain's local Administrator account, provided it is secured as described earlier and in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

When EA access is required, the users whose accounts require EA rights and permissions should be temporarily placed into the Enterprise Admins group. Although users are using the highly privileged accounts, their activities should be audited and preferably performed with one user performing the changes and another user observing the changes to minimize the likelihood of inadvertent misuse or misconfiguration. When the activities have been completed, the accounts should be removed from the EA group. This can be achieved via manual procedures and documented processes, third-party privileged identity/access management (PIM/PAM) software, or a combination of both. Guidelines for creating accounts that can be used to control the membership of privileged groups in Active Directory are provided in Attractive Accounts for Credential Theft and detailed instructions are provided in Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory.

Enterprise Admins are, by default, members of the built-in Administrators group in each domain in the forest. Removing the Enterprise Admins group from the Administrators groups in each domain is an inappropriate modification because in the event of a forest disaster-recovery scenario, EA rights will likely be required. If the Enterprise Admins group has been removed from Administrators groups in a forest, it should be added to the Administrators group in each domain and the following additional controls should be implemented:

- As described earlier, the Enterprise Admins group should contain no users on a day-to-day basis, with the possible exception of the forest root domain's Administrator account, which should be secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.
- In GPOs linked to OUs containing member servers and workstations in each domain, the EA group should be added to the following user rights:
  - Deny access to this computer from the network
  - Deny log on as a batch job
  - Deny log on as a service

- Deny log on locally
- Deny log on through Remote Desktop Services.

This will prevent members of the EA group from logging on to member servers and workstations. If jump servers are used to administer domain controllers and Active Directory, ensure that jump servers are located in an OU to which the restrictive GPOs are not linked.

- Auditing should be configured to send alerts if any modifications are made to the properties or membership of the EA group. These alerts should be sent, at a minimum, to users or teams responsible for Active Directory administration and incident response. You should also define processes and procedures for temporarily populating the EA group, including notification procedures when legitimate population of the group is performed.

Securing Domain Admins Groups

As is the case with the Enterprise Admins group, membership in Domain Admins groups should be required only in build or disaster-recovery scenarios. There should be no day-to-day user accounts in the DA group with the exception of the local Administrator account for the domain, if it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

When DA access is required, the accounts needing this level of access should be temporarily placed in the DA group for the domain in question. Although the users are using the highly privileged accounts, activities should be audited and preferably performed with one user performing the changes and another user observing the changes to minimize the likelihood of inadvertent misuse or misconfiguration. When the activities have been completed, the accounts should be removed from the Domain Admins group. This can be achieved via manual procedures and documented processes, via third-party privileged identity/access management (PIM/PAM) software, or a combination of both. Guidelines for creating accounts that can be used to control the membership of privileged groups in Active Directory are provided in Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory.

Domain Admins are, by default, members of the local Administrators groups on all member servers and workstations in their respective domains. This default nesting should not be modified because it affects supportability and disaster recovery options. If Domain Admins groups have been removed from the local Administrators groups on the member servers, they should be added to the Administrators group on each member server and workstation in the domain via restricted group settings in linked GPOs. The following general controls, which are described in depth in Appendix F: Securing Domain Admins Groups in Active Directory should also be implemented.

For the Domain Admins group in each domain in the forest:

1. Remove all members from the DA group, with the possible exception of the built-in Administrator account for the domain, provided it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.
2. In GPOs linked to OUs containing member servers and workstations in each domain, the

DA group should be added to the following user rights:

- Deny access to this computer from the network
- Deny log on as a batch job
- Deny log on as a service
- Deny log on locally
- Deny log on through Remote Desktop Services

This will prevent members of the DA group from logging on to member servers and workstations. If jump servers are used to administer domain controllers and Active Directory, ensure that jump servers are located in an OU to which the restrictive GPOs are not linked.

3. Auditing should be configured to send alerts if any modifications are made to the properties or membership of the DA group. These alerts should be sent, at a minimum, to users or teams responsible for AD DS administration and incident response. You should also define processes and procedures for temporarily populating the DA group, including notification procedures when legitimate population of the group is performed.

Securing Administrators Groups in Active Directory

As is the case with the EA and DA groups, membership in the Administrators (BA) group should be required only in build or disaster-recovery scenarios. There should be no day-to-day user accounts in the Administrators group with the exception of the local Administrator account for the domain, if it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

When Administrators access is required, the accounts needing this level of access should be temporarily placed in the Administrators group for the domain in question. Although the users are using the highly privileged accounts, activities should be audited and, preferably, performed with a user performing the changes and another user observing the changes to minimize the likelihood of inadvertent misuse or misconfiguration. When the activities have been completed, the accounts should immediately be removed from the Administrators group. This can be achieved via manual procedures and documented processes, via third-party privileged identity/access management (PIM/PAM) software, or a combination of both.

Administrators are, by default, the owners of most of the AD DS objects in their respective domains. Membership in this group may be required in build and disaster recovery scenarios in which ownership or the ability to take ownership of objects is required. Additionally, DAs and EAs inherit a number of their rights and permissions by virtue of their default membership in the Administrators group. Default group nesting for privileged groups in Active Directory should not be modified, and each domain's Administrators group should be secured as described in Appendix G: Securing Administrators Groups in Active Directory, and in the general instructions below.

1. Remove all members from the Administrators group, with the possible exception of the local Administrator account for the domain, provided it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

2. Members of the domain's Administrators group should never need to log on to member servers or workstations. In one or more GPOs linked to workstation and member server OUs in each domain, the Administrators group should be added to the following user rights:
    - Deny access to this computer from the network
    - Deny log on as a batch job,
    - Deny log on as a service
    - This will prevent members of the Administrators group from being used to log on or connect to member servers or workstations (unless multiple controls are first breached), where their credentials could be cached and thereby compromised. A privileged account should never be used to log on to a less-privileged system, and enforcing these controls affords protection against a number of attacks.
3. At the domain controllers OU in each domain in the forest, the Administrators group should be granted the following user rights (if they do not already have these rights), which will allow the members of the Administrators group to perform functions necessary for a forest-wide disaster recovery scenario:
    - Access this computer from the network
    - Allow log on locally
    - Allow log on through Remote Desktop Services
4. Auditing should be configured to send alerts if any modifications are made to the properties or membership of the Administrators group. These alerts should be sent, at a minimum, to members of the team responsible for AD DS administration. Alerts should also be sent to members of the security team, and procedures should be defined for modifying the membership of the Administrators group. Specifically, these processes should include a procedure by which the security team is notified when the Administrators group is going to be modified so that when alerts are sent, they are expected and an alarm is not raised. Additionally, processes to notify the security team when the use of the Administrators group has been completed and the accounts used have been removed from the group should be implemented.

Notes

When you implement restrictions on the Administrators group in GPOs, Windows applies the settings to members of a computer's local Administrators group in addition to the domain's Administrators group. Therefore, you should use caution when implementing restrictions on the Administrators group. Although prohibiting network, batch and service logons for members of the Administrators group is advised wherever it is feasible to implement, do not restrict local logons or logons through Remote Desktop Services. Blocking these logon types can block legitimate administration of a computer by members of the local Administrators group. The following screenshot shows configuration settings that block misuse of built-in local and domain Administrator accounts, in addition to misuse of built-in local or domain Administrators groups. Note that the **Deny log on through Remote Desktop Services** user right does not include the Administrators group, because including it in this setting would also block these logons for accounts that are members of the local computer's Administrators group. If services on computers are configured to run in the context of any of the privileged groups described in this section,

implementing these settings can cause services and applications to fail. Therefore, as with all of the recommendations in this section, you should thoroughly test settings for applicability in your environment.

Role-Based Access Controls (RBAC) for Active Directory

Generally speaking, role-based access controls (RBAC) are a mechanism for grouping users and providing access to resources based on business rules. In the case of Active Directory, implementing RBAC for AD DS is the process of creating roles to which rights and permissions are delegated to allow members of the role to perform day-to-day administrative tasks without granting them excessive privilege. RBAC for Active Directory can be designed and implemented via native tooling and interfaces, by leveraging software you may already own, by purchasing third-party products, or any combination of these approaches. This section does not provide step-by-step instructions to implement RBAC for Active Directory, but instead discusses factors you should consider in choosing an approach to implementing RBAC in your AD DS installations.

Native Approaches to RBAC for Active Directory

In the simplest RBAC implementation, you can implement roles as AD DS groups and delegate rights and permissions to the groups that allow them to perform daily administration within the designated scope of the role.

In some cases, existing security groups in Active Directory can be used to grant rights and permissions appropriate to a job function. For example, if specific employees in your IT organization are responsible for the management and maintenance of DNS zones and records, delegating those responsibilities can be as simple as creating an account for each DNS administrator and adding it to the DNS Admins group in Active Directory. The DNS Admins group, unlike more highly privileged groups, has few powerful rights across Active Directory, although members of this group have been delegated permissions that allow them to administer DNS.

In other cases, you may need to create security groups and delegate rights and permissions to Active Directory objects, file system objects, and registry objects to allow members of the groups to perform designated administrative tasks. For example, if your Help Desk operators are responsible for resetting forgotten passwords, assisting users with connectivity problems, and troubleshooting application settings, you may need to combine delegation settings on user objects in Active Directory with privileges that allow Help Desk users to connect remotely to users' computers to view or modify the users' configuration settings. For each role you define, you should identify:

1. Which tasks members of the role perform on a day-to-day basis and which tasks are less frequently performed.
2. On which systems and in which applications members of a role should be granted rights and permissions.

3. Which users should be granted membership in a role.
4. How management of role memberships will be performed.

In many environments, manually creating role-based access controls for administration of an Active Directory environment can be challenging to implement and maintain. If you have clearly defined roles and responsibilities for administration of your IT infrastructure, you may want to leverage additional tooling to assist you in creating a manageable native RBAC deployment. For example, if Forefront Identity Manager (FIM) is in use in your environment, you can use FIM to automate the creation and population of administrative roles, which can ease ongoing administration. If you use System Center Configuration Manager (SCCM) and System Center Operations Manager (SCOM), you can use application-specific roles to delegate management and monitoring functions, and also enforce consistent configuration and auditing across systems in the domain. If you have implemented a public key infrastructure (PKI), you can issue and require smart cards for IT staff responsible for administering the environment. With FIM Credential Management (FIM CM), you can even combine management of roles and credentials for your administrative staff.

In other cases, it may be preferable for an organization to consider deploying third-party RBAC software that provides "out-of-box" functionality. Commercial, off-the-shelf (COTS) solutions for RBAC for Active Directory, Windows, and non-Windows directories and operating systems are offered by a number of vendors. A list of third-party RBAC solutions is provided in Appendix J: Third-Party RBAC Vendors. When choosing between native solutions and third-party products, you should consider the following factors:

1. Budget: By investing in development of RBAC using software and tools you may already own, you can reduce the software costs involved in deploying a solution. However, unless you have staff who are experienced in creating and deploying native RBAC solutions, you may need to engage consulting resources to develop your solution. You should carefully weigh the anticipated costs for a custom-developed solution with the costs to deploy an "out-of-box" solution, particularly if your budget is limited.
2. Composition of the IT environment: If your environment is comprised primarily of Windows systems, or if you are already leveraging Active Directory for management of non-Windows systems and accounts, custom native solutions may provide the optimal solution for your needs. If your infrastructure contains many systems that are not running Windows and are not managed by Active Directory, you may need to consider options for management of non-Windows systems separately from the Active Directory environment.
3. Privilege model in the solution: If a product relies on placement of its service accounts into highly privileged groups in Active Directory and does not offer options that do not require excessive privilege be granted to the RBAC software, you have not really reduced your Active Directory attack surface—you've only changed the composition of the most privileged groups in the directory. Unless an application vendor can provide controls for service accounts that minimize the probability of the accounts being compromised and maliciously used, you may want to consider other options.

Privileged Identity Management

Privileged identity management (PIM), sometimes referred to as privileged account management (PAM) or privileged credential management (PCM) is the design, construction, and implementation of approaches to managing privileged accounts in your infrastructure. Generally speaking, PIM provides mechanisms by which accounts are granted temporary rights and permissions required to perform build-or-break fix functions, rather than leaving privileges permanently attached to accounts. Whether PIM functionality is manually created or is implemented via the deployment of third-party software (see Appendix K: Third-Party PIM Vendors), one or more of the following features may be available:

- Credential "vaults," where passwords for privileged accounts are "checked out" and assigned an initial password, then "checked in" when activities have been completed, at which time passwords are again reset on the accounts.
- Time-bound restrictions on the use of privileged credentials
- One-time-use credentials
- Workflow-generated granting of privilege with monitoring and reporting of activities performed and automatic removal of privilege when activities are completed or allotted time has expired
- Replacement of hard-coded credentials such as user names and passwords in scripts with application programming interfaces (APIs) that allow credentials to be retrieved from vaults as needed
- Automatic management of service account credentials

Creating Unprivileged Accounts to Manage Privileged Accounts

One of the challenges in managing privileged accounts is that, by default, the accounts that can manage privileged and protected accounts and groups are privileged and protected accounts. If you implement appropriate RBAC and PIM solutions for your Active Directory installation, the solutions may include approaches that allow you to effectively depopulate the membership of the most privileged groups in the directory, populating the groups only temporarily and when needed.

If you implement native RBAC and PIM, however, you should consider creating accounts that have no privilege and with the only function of populating and depopulating privileged groups in Active Directory when needed. Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory provides step-by-step instructions that you can use to create accounts for this purpose.

Implementing Robust Authentication Controls

Law Number Six: There really is someone out there trying to guess your passwords. – 10 Immutable Laws of Security Administration

Pass-the-hash and other credential theft attacks are not specific to Windows operating systems, nor are they new. The first pass-the-hash attack was created in 1997. Historically, however, these attacks required customized tools, were hit-or-miss in their success, and required attackers to have a relatively high degree of skill. The introduction of freely available, easy-to-use tooling that natively extracts credentials has resulted in an exponential increase in the number and success

of credential theft attacks in recent years. However, credential theft attacks are by no means the only mechanisms by which credentials are targeted and compromised.

Although you should implement controls to help protect you against credential theft attacks, you should also identify the accounts in your environment that are most likely to be targeted by attackers, and implement robust authentication controls for those accounts. If your most privileged accounts are using single factor authentication such as user names and passwords (both are “something you know,” which is one authentication factor), those accounts are weakly protected. All that an attacker needs is knowledge of the user name and knowledge of the password associated with the account, and pass-the-hash attacks are not required—the attacker can authenticate as the user to any systems that accept single factor credentials.

Although implementing multifactor authentication does not protect you against pass-the-hash attacks, implementing multifactor authentication in combination with protected systems can. More information about implementing protected systems is provided in Implementing Secure Administrative Hosts, and authentication options are discussed in the following sections.

General Authentication Controls

If you have not already implemented multifactor authentication such as smart cards, consider doing so. Smart cards implement hardware-enforced protection of private keys in a public-private key pair, preventing a user’s private key from being accessed or used unless the user presents the proper PIN, passcode, or biometric identifier to the smart card. Even if a user’s PIN or passcode is intercepted by a keystroke logger on a compromised computer, for an attacker to reuse the PIN or passcode, the card must also be physically present.

In cases in which long, complex passwords have proven difficult to implement because of user resistance, smart cards provide a mechanism by which users may implement relatively simple PINs or passcodes without the credentials being susceptible to brute force or rainbow table attacks. Smart card PINs are not stored in Active Directory or in local SAM databases, although credential hashes may still be stored in LSASS protected memory on computers on which smart cards have been used for authentication.

Additional Controls for VIP Accounts

Another benefit of implementing smart cards or other certificate-based authentication mechanisms is the ability to leverage Authentication Mechanism Assurance to protect sensitive data that is accessible to VIP users. Authentication Mechanism Assurance is available in domains in which the functional level is set to Windows Server 2012 or Windows Server 2008 R2. When it is enabled, Authentication Mechanism Assurance adds an administrator-designated global group membership to a user’s Kerberos token when the user’s credentials are authenticated during logon using a certificate-based logon method.

This makes it possible for resource administrators to control access to resources, such as files, folders, and printers, based on whether the user logs on using a certificate-based logon method, in addition to the type of certificate used. For example, when a user logs on by using a smart card, the user’s access to resources on the network can be specified as different from what the

access is when the user does not use a smart card (that is, when the user logs on by entering a user name and password). For more information about Authentication Mechanism Assurance, see the Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide.

Configuring Privileged Account Authentication

In Active Directory for all administrative accounts, enable the **Require smart card for interactive logon** attribute, and audit for changes to (at a minimum), any of the attributes on the **Account** tab for the account (for example, cn, name, sAMAccountName, userPrincipalName, and userAccountControl) administrative user objects.

Although setting the **Require smart card for interactive logon** on accounts resets the account's password to a 120-character random value and requires smart cards for interactive logons, the attribute can still be overwritten by users with permissions that allow them to change passwords on the accounts, and the accounts can then be used to establish noninteractive logons with only user name and password.

In other cases, depending on the configuration of accounts in Active Directory and certificate settings in Active Directory Certificate Services (AD CS) or a third-party PKI, User Principal Name (UPN) attributes for administrative or VIP accounts can be targeted for a specific kind of attack, as described here.

UPN Hijacking for Certificate Spoofing

Although a thorough discussion of attacks against public key infrastructures (PKIs) is outside the scope of this document, attacks against public and private PKIs have increased exponentially since 2008. Breaches of public PKIs have been broadly publicized, but attacks against an organization's internal PKI are perhaps even more prolific. One such attack leverages Active Directory and certificates to allow an attacker to spoof the credentials of other accounts in a manner that can be difficult to detect.

When a certificate is presented for authentication to a domain-joined system, the contents of the Subject or the Subject Alternative Name (SAN) attribute in the certificate are used to map the certificate to a user object in Active Directory. Depending on the type of certificate and how it is constructed, the Subject attribute in a certificate typically contains a user's common name (CN), as shown in the following screenshot.

By default, Active Directory constructs a user's CN by concatenating the account's first name + " "+ last name. However, CN components of user objects in Active Directory are not required or guaranteed to be unique, and moving a user account to a different location in the directory changes the account's distinguished name (DN), which is the full path to the object in the directory, as shown in the bottom pane of the previous screenshot.

Because certificate subject names are not guaranteed to be static or unique, the contents of the Subject Alternative Name are often used to locate the user object in Active Directory. The SAN attribute for certificates issued to users from enterprise certification authorities (Active Directory integrated CAs) typically contains the user's UPN or email address. Because UPNs are

guaranteed to be unique in an AD DS forest, locating a user object by UPN is commonly performed as part of authentication, with or without certificates involved in the authentication process.

The use of UPNs in SAN attributes in authentication certificates can be leveraged by attackers to obtain fraudulent certificates. If an attacker has compromised an account that has the ability to read and write UPNs on user objects, the attack is implemented as follows:

The UPN attribute on a user object (such as a VIP user) is temporarily changed to a different value. The SAM account name attribute and CN can also be changed at this time, although this is usually not necessary for the reasons described earlier.

When the UPN attribute on the target account has been changed, a stale, enabled user account or a freshly created user account's UPN attribute is changed to the value that was originally assigned to the target account. Stale, enabled user accounts are accounts that have not logged on for long periods of time, but have not been disabled. They are targeted by attackers who intend to "hide in plain sight" for the following reasons:

1. Because the account is enabled, but hasn't been used recently, using the account is unlikely to trigger alerts the way that enabling a disabled user account might.
2. Use of an existing account doesn't require the creation of a new user account that might be noticed by administrative staff.
3. Stale user accounts that are still enabled are usually members of various security groups and are granted access to resources on the network, simplifying access and "blending in" to an existing user population.

The user account on which the target UPN has now been configured is used to request one or more certificates from Active Directory Certificate Services.

When certificates have been obtained for the attacker's account, the UPNs on the "new" account and the target account are returned to their original values.

The attacker now has one or more certificates that can be presented for authentication to resources and applications as if the user is the VIP user whose account was temporarily modified. Although a full discussion of all of the ways in which certificates and PKI can be targeted by attackers is outside the scope of this document, this attack mechanism is provided to illustrate why you should monitor privileged and VIP accounts in AD DS for changes, particularly for changes to any of the attributes on the **Account** tab for the account (for example, cn, name, sAMAccountName, userPrincipalName, and userAccountControl). In addition to monitoring the accounts, you should restrict who can modify the accounts to as small a set of administrative users as possible. Likewise, the accounts of administrative users should be protected and monitored for unauthorized changes.

Implementing Secure Administrative Hosts

Secure administrative hosts are workstations or servers that have been configured specifically for the purposes of creating secure platforms from which privileged accounts can perform administrative tasks in Active Directory or on domain controllers, domain-joined systems, and

applications running on domain-joined systems. In this case, “privileged accounts” refers not only to accounts that are members of the most privileged groups in Active Directory, but to any accounts that have been delegated rights and permissions that allow administrative tasks to be performed.

These accounts may be Help Desk accounts that have the ability to reset passwords for most of the users in a domain, accounts that are used to administer DNS records and zones, or accounts that are used for configuration management. Secure administrative hosts are dedicated to administrative functionality, and they do not run software such as email applications, web browsers, or productivity software such as Microsoft Office.

Although the “most privileged” accounts and groups should accordingly be the most stringently protected, this does not eliminate the need to protect any accounts and groups to which privileges above those of standard user accounts have been granted.

A secure administrative host can be a dedicated workstation that is used only for administrative tasks, a member server that runs the Remote Desktop Gateway server role and to which IT users connect to perform administration of destination hosts, or a server that runs the Hyper-V® role and provides a unique virtual machine for each IT user to use for their administrative tasks. In many environments, combinations of all three approaches may be implemented.

Implementing secure administrative hosts requires planning and configuration that is consistent with your organization’s size, administrative practices, risk appetite, and budget. Considerations and options for implementing secure administrative hosts are provided here for you to use in developing an administrative strategy suitable for your organization.

Principles for Creating Secure Administrative Hosts

To effectively secure systems against attacks, a few general principles should be kept in mind:

1. You should never administer a trusted system (that is, a secure server such as a domain controller) from a less-trusted host (that is, a workstation that is not secured to the same degree as the systems it manages).
2. You should not rely on a single authentication factor when performing privileged activities; that is, user name and password combinations should not be considered acceptable authentication because only a single factor (something you know) is represented. You should consider where credentials are generated and cached or stored in administrative scenarios.
3. Although most attacks in the current threat landscape leverage malware and malicious hacking, do not omit physical security when designing and implementing secure administrative hosts.

Account Configuration

Even if your organization does not currently use smart cards, you should consider implementing them for privileged accounts and secure administrative hosts. Administrative hosts should be

configured to require smart card logon for all accounts by modifying the following setting in a GPO that is linked to the OUs containing administrative hosts:

Computer Configuration\Policies\Windows Settings\Local Policies\Security Options\Interactive logon: Require smart card

This setting will require all interactive logons to use a smart card, regardless of the configuration on an individual account in Active Directory.

You should also configure secure administrative hosts to permit logons only by authorized accounts, which can be configured in:

Computer Configuration\Policies\Windows Settings\Local Policies\Security Settings\Local Policies\User Rights Assignment

This grants interactive (and, where appropriate, Remote Desktop Services) logon rights only to authorized users of the secure administrative host.

Physical Security

For administrative hosts to be considered trustworthy, they must be configured and protected to the same degree as the systems they manage. Most of the recommendations provided in Securing Domain Controllers Against Attack are also applicable to the hosts that are used to administer domain controllers and the AD DS database. One of the challenges of implementing secure administrative systems in most environments is that physical security can be more difficult to implement because these computers often reside in areas that are not as secure as servers hosted in datacenters, such as administrative users' desktops.

Physical security includes controlling physical access to administrative hosts. In a small organization, this may mean that you maintain a dedicated administrative workstation that is kept locked in an office or a desk drawer when not in use. Or it may mean that when you need to perform administration of Active Directory or your domain controllers, you log on to the domain controller directly.

In medium-sized organizations, you may consider implementing secure administrative "jump servers" that are located in a secured location in an office and are used when management of Active Directory or domain controllers is required. You may also implement administrative workstations that are locked in secure locations when not in use, with or without jump servers.

In large organizations, you can deploy datacenter-housed jump servers that provide strictly controlled access to Active Directory; domain controllers; and file, print, or application servers. Implementation of a jump server architecture is most likely to include a combination of secure workstations and servers in large environments.

Regardless of the size of your organization and the design of your administrative hosts, you should secure physical computers against unauthorized access or theft, and should use BitLocker Drive Encryption to encrypt and protect the drives on administrative hosts. By implementing BitLocker on administrative hosts, even if a host is stolen or its disks are removed, you can ensure that the data on the drive is inaccessible to unauthorized users.

Operating System Versions and Configuration

All administrative hosts, whether servers or workstations, should run the newest operating system in use in your organization for the reasons described earlier in this document. By running current operating systems, your administrative staff benefits from new security features, full vendor support, and additional functionality introduced in the operating system. Moreover, when you are evaluating a new operating system, by deploying it first to administrative hosts, you will need to familiarize yourself with the new features, settings and management mechanisms it offers, which can subsequently be leveraged in planning broader deployment of the operating system. By then, the most sophisticated users in your organization will also be the users who are familiar with the new operating system and best positioned to support it.

Microsoft Security Configuration Wizard

If you implement jump servers as part of your administrative host strategy, you should use the built-in Security Configuration Wizard to configure service, registry, audit, and firewall settings to reduce the server's attack surface. When the Security Configuration Wizard configuration settings have been collected and configured, the settings can be converted to a GPO that is used to enforce a consistent baseline configuration on all jump servers. You can further edit the GPO to implement security settings specific to jump servers, and can combine all of the settings with additional baseline settings extracted from the Microsoft Security Compliance Manager.

Microsoft Security Compliance Manager

The Microsoft Security Compliance Manager is a freely available tool that integrates security configurations that are recommended by Microsoft, based on operating system version and role configuration, and collects them in a single tool and UI that can be used to create and configure baseline security settings for domain controllers. Microsoft Security Compliance Manager templates can be combined with Security Configuration Wizard settings to produce comprehensive configuration baselines for jump servers that are deployed and enforced by GPOs deployed at the OUs in which jump servers are located in Active Directory.

Note

As of this writing, the Microsoft Security Compliance Manager does not include settings specific to jump servers or other secure administrative hosts, but Security Compliance Manager (SCM) can still be used to create initial baselines for your administrative hosts. To properly secure the hosts, however, you should apply additional security settings appropriate to highly secured workstations and servers.

AppLocker

Administrative hosts and virtual machinesshould be configured with script, tool, and application whitelists via AppLocker or a third-party application restriction software. Any administrative applications or utilities that do not adhere to secure settings should be upgraded or replaced with tooling that adheres to secure development and administrative practices. When new or additional

tooling is needed on an administrative host, applications and utilities should be thoroughly tested, and if the tooling is suitable for deployment on administrative hosts, it can be added to the systems' whitelists.

RDP Restrictions

Although the specific configuration will vary depending on the architecture of your administrative systems, you should include restrictions on which accounts and computers can be used to establish Remote Desktop Protocol (RDP) connections to managed systems, such as using Remote Desktop Gateway (RD Gateway) jump servers to control access to domain controllers and other managed systems from authorized users and systems.

You should allow interactive logons by authorized users and should remove or even block other logon types that are not needed for server access.

Patch and Configuration Management

Smaller organizations may rely on offerings such as Windows Update or Windows Server Update Services (WSUS) to manage deployment of updates to Windows systems, while larger organizations may implement enterprise patch and configuration management software such as System Center Configuration Manager. Regardless of the mechanisms you use to deploy updates to your general server and workstation population, you should consider separate deployments for highly secure systems such as domain controllers, certification authorities, and administrative hosts. By segregating these systems from the general management infrastructure, if your management software or service accounts are compromised, the compromise cannot be easily extended to the most secure systems in your infrastructure.

Although you should not implement manual update processes for secure systems, you should configure a separate infrastructure for updating secure systems. Even in very large organizations, this infrastructure can usually be implemented via dedicated WSUS servers and GPOs for secured systems.

Blocking Internet Access

Administrative hosts should not be permitted to access the Internet, nor should they be able to browse an organization's intranet. Web browsers and similar applications should not be permitted on administrative hosts. You can block Internet access for secure hosts via a combination of perimeter firewall settings, WFAS configuration, and "black hole" proxy configuration on secure hosts. You can also use application whitelisting to prevent web browsers from being used on administrative hosts.

Virtualization

Where possible, consider implementing virtual machines as administrative hosts. Using virtualization, you can create per-user administrative systems that are centrally stored and managed, and which can be easily shut down when not in use, ensuring that credentials are not

left active on the administrative systems. You can also require that virtual administrative hosts are reset to an initial snapshot after each use, ensuring that the virtual machines remain pristine. More information about options for virtualization of administrative hosts is provided in the following section.

Sample Approaches to Implementing Secure Administrative Hosts

Regardless of how you design and deploy your administrative host infrastructure, you should keep in mind the guidelines provided in “Principles for Creating Secure Administrative Hosts” earlier in this topic. Each of the approaches described here provides general information about how you can separate “administrative” and “productivity” systems used by your IT staff. Productivity systems are computers that IT administrators employ to check email, browse the Internet, and to use general productivity software such as Microsoft Office. Administrative systems are computers that are hardened and dedicated to use for day-to-day administration of an IT environment.

The simplest way to implement secure administrative hosts is to provide your IT staff with secured workstations from which they can perform administrative tasks. In a workstation-only implementation, each administrative workstation is used to launch management tools and RDP connections to manage servers and other infrastructure. Workstation-only implementations can be effective in smaller organizations, although larger, more complex infrastructures may benefit from a distributed design for administrative hosts in which dedicated administrative servers and workstations are used, as described in “Implementing Secure Administrative Workstations and Jump Servers” later in this topic.

Implementing Separate Physical Workstations

One way that you can implement administrative hosts is to issue each IT user two workstations. One workstation is used with a “regular” user account to perform activities such as checking email and using productivity applications, while the second workstation is dedicated strictly to administrative functions.

For the productivity workstation, the IT staff can be given regular user accounts rather than using privileged accounts to log on to unsecured computers. The administrative workstation should be configured with a stringently controlled configuration and the IT staff should use a different account to log on to the administrative workstation.

If you have implemented smart cards, administrative workstations should be configured to require smart card logons, and IT staff should be given separate accounts for administrative use, also configured to require smart cards for interactive logon. The administrative host should be hardened as previously described, and only designated IT users should be allowed to log on locally to the administrative workstation.

Pros

By implementing separate physical systems, you can ensure that each computer is configured appropriately for its role and that IT users cannot inadvertently expose administrative systems to risk.

Cons

- Implementing separate physical computers increases hardware costs.
- Logging on to a physical computer with credentials that are used to administer remote systems caches the credentials in memory.
- If administrative workstations are not stored securely, they may be vulnerable to compromise via mechanisms such as physical hardware key loggers or other physical attacks.

Implementing a Secure Physical Workstation with a Virtualized Productivity Workstation

In this approach, IT users are given a secured administrative workstation from which they can perform day-to-day administrative functions, using Remote Server Administration Tools (RSAT) or RDP connections to servers within their scope of responsibility. When IT users need to perform productivity tasks, they can connect via RDP to a remote productivity workstation running as a virtual machine. Separate credentials should be used for each workstation, and controls such as smart cards should be implemented.

Pros

- Administrative workstations and productivity workstations are separated.
- IT staff using secure workstations to connect to productivity workstations can use separate credentials and smart cards, and privileged credentials are not deposited on the less-secure computer.

Cons

- Implementing the solution requires design and implementation work and robust virtualization options.
- If the physical workstations are not stored securely, they may be vulnerable to physical attacks that compromise the hardware or the operating system and make them susceptible to communications interception.

Implementing a Single Secure Workstation with Connections to Separate “Productivity” and “Administrative” Virtual Machines

In this approach, you can issue IT users a single physical workstation that is locked down as previously described, and on which IT users do not have privileged access. You can provide Remote Desktop Services connections to virtual machines hosted on dedicated servers,

providing IT staff with one virtual machine that runs email and other productivity applications, and a second virtual machine that is configured as the user's dedicated administrative host.

You should require smart card or other multifactor logon for the virtual machines, using separate accounts other than the account that is used to log on to the physical computer. After an IT user logs on to a physical computer, they can use their productivity smart card to connect to their remote productivity computer and a separate account and smart card to connect to their remote administrative computer.

Pros

- IT users can use a single physical workstation.
- By requiring separate accounts for the virtual hosts and using Remote Desktop Services connections to the virtual machines, IT users' credentials are not cached in memory on the local computer.
- The physical host can be secured to the same degree as administrative hosts, reducing the likelihood of compromise of the local computer.
- In cases in which an IT user's productivity virtual machine or their administrative virtual machine may have been compromised, the virtual machine can easily be reset to a "known good" state.
- If the physical computer is compromised, no privileged credentials will be cached in memory, and the use of smart cards can prevent compromise of credentials by keystroke loggers.

Cons

- Implementing the solution requires design and implementation work and robust virtualization options.
- If the physical workstations are not stored securely, they may be vulnerable to physical attacks that compromise the hardware or the operating system and make them susceptible to communications interception.

Implementing Secure Administrative Workstations and Jump Servers

As an alternative to secure administrative workstations, or in combination with them, you can implement secure jump servers, and administrative users can connect to the jump servers using RDP and smart cards to perform administrative tasks.

Jump servers should be configured to run the Remote Desktop Gateway role to allow you to implement restrictions on connections to the jump server and to destination servers that will be managed from it. If possible, you should also install the Hyper-V role and create Personal Virtual Desktops or other per-user virtual machines for administrative users to use for their tasks on the jump servers.

By giving the administrative users per-user virtual machines on the jump server, you provide physical security for the administrative workstations, and administrative users can reset or shut down their virtual machines when not in use. If you prefer not to install the Hyper-V role and the

Remote Desktop Gateway role on the same administrative host, you can install them on separate computers.

Wherever possible, remote administration tools should be used to manage servers. The Remote Server Administration Tools (RSAT) feature should be installed on the users' virtual machines (or the jump server if you are not implementing per-user virtual machines for administration), and administrative staff should connect via RDP to their virtual machines to perform administrative tasks.

In cases when an administrative user must connect via RDP to a destination server to manage it directly, RD Gateway should be configured to allow the connection to be made only if the appropriate user and computer are used to establish the connection to the destination server. Execution of RSAT (or similar) tools should be prohibited on systems that are not designated management systems, such as general-use workstations and member servers that are not jump servers.

Pros

- Creating jump servers allows you to map specific servers to "zones" (collections of systems with similar configuration, connection, and security requirements) in your network and to require that the administration of each zone is achieved by administrative staff connecting from secure administrative hosts to a designated "zone" server.
- By mapping jump servers to zones, you can implement granular controls for connection properties and configuration requirements, and can easily identify attempts to connect from unauthorized systems.
- By implementing per-administrator virtual machines on jump servers, you enforce shutdown and resetting of the virtual machines to a known clean state when administrative tasks are completed. By enforcing shutdown (or restart) of the virtual machines when administrative tasks are completed, the virtual machines cannot be targeted by attackers, nor are credential theft attacks feasible because memory-cached credentials do not persist beyond a reboot.

Cons

- Dedicated servers are required for jump servers, whether physical or virtual.
- Implementing designated jump servers and administrative workstations requires careful planning and configuration that maps to any security zones configured in the environment.

Securing Domain Controllers Against Attack

Law Number Three: If a bad guy has unrestricted physical access to your computer, it's not your computer anymore. – Ten Immutable Laws of Security (Version 2.0)

Domain controllers provide the physical storage for the AD DS database, in addition to providing the services and data that allow enterprises to effectively manage their servers, workstations, users, and applications. If privileged access to a domain controller is obtained by a malicious

user, that user can modify, corrupt, or destroy the AD DS database and, by extension, all of the systems and accounts that are managed by Active Directory.

Because domain controllers can read from and write to anything in the AD DS database, compromise of a domain controller means that your Active Directory forest can never be considered trustworthy again unless you are able to recover using a known good backup and to close the gaps that allowed the compromise in the process.

Depending on an attacker's preparation, tooling, and skill, modification or even irreparable damage to the AD DS database can be completed in minutes to hours, not days or weeks. What matters isn't how long an attacker has privileged access to Active Directory, but how much the attacker has planned for the moment when privileged access is obtained. Compromising a domain controller can provide the most expedient path to wide scale propagation of access, or the most direct path to destruction of member servers, workstations, and Active Directory. Because of this, domain controllers should be secured separately and more stringently than the general Windows infrastructure.

Physical Security for Domain Controllers

This section provides information about physically securing domain controllers, whether the domain controllers are physical or virtual machines, in datacenter locations, branch offices, and even remote locations with only basic infrastructure controls.

Datacenter Domain Controllers

Physical Domain Controllers

In datacenters, physical domain controllers should be installed in dedicated secure racks or cages that are separate from the general server population. When possible, domain controllers should be configured with Trusted Platform Module (TPM) chips and all volumes in the domain controller servers should be protected via BitLocker Drive Encryption. BitLocker generally adds performance overhead in single-digit percentages, but protects the directory against compromise even if disks are removed from the server. BitLocker can also help protect systems against attacks such as rootkits because the modification of boot files will cause the server to boot into recovery mode so that the original binaries can be loaded. If a domain controller is configured to use software RAID, serial-attached SCSI, SAN/NAS storage, or dynamic volumes, BitLocker cannot be implemented, so locally attached storage (with or without hardware RAID) should be used in domain controllers whenever possible.

Virtual Domain Controllers

If you implement virtual domain controllers, you should ensure that domain controllers run on separate physical hosts than other virtual machines in the environment. Even if you use a third-party virtualization platform, consider deploying virtual domain controllers on Hyper-V Server in Windows Server 2012 or Windows Server 2008 R2, which provides a minimal attack surface and

can be managed with the domain controllers it hosts rather than being managed with the rest of the virtualization hosts. If you implement System Center Virtual Machine Manager (SCVMM) for management of your virtualization infrastructure, you can delegate administration for the physical hosts on which domain controller virtual machines reside and the domain controllers themselves to authorized administrators. You should also consider separating the storage of virtual domain controllers to prevent storage administrators from accessing the virtual machine files.

Branch Locations

Physical Domain Controllers

In locations in which multiple servers reside but are not physically secured to the degree that datacenter servers are secured, physical domain controllers should be configured with TPM chips and BitLocker Drive Encryption for all server volumes. If a domain controller cannot be stored in a locked room in branch locations, you should consider deploying RODCs in those locations.

Virtual Domain Controllers

Whenever possible, you should run virtual domain controllers in branch offices on separate physical hosts than the other virtual machines in the site. In branch offices in which virtual domain controllers cannot run on separate physical hosts from the rest of the virtual server population, you should implement TPM chips and BitLocker Drive Encryption on hosts on which virtual domain controllers run at minimum, and all hosts if possible. Depending on the size of the branch office and the security of the physical hosts, you should consider deploying RODCs in branch locations.

Remote Locations with Limited Space and Security

If your infrastructure includes locations in which only a single physical server can be installed, a server capable of running virtualization workloads should be installed in the remote location, and BitLocker Drive Encryption should be configured to protect all volumes in the server. One virtual machine on the server should run an RODC, with other servers running as separate virtual machines on the host. Information about planning for deployment of RODC is provided in the Read-Only Domain Controller Planning and Deployment Guide. For more information about deploying and securing virtualized domain controllers, see Running Domain Controllers in Hyper-V on the TechNet website. For more detailed guidance for hardening Hyper-V, delegating virtual machine management, and protecting virtual machines, see the Hyper-V Security Guide Solution Accelerator on the Microsoft website.

Domain Controller Operating Systems

You should run all domain controllers on the newest version of Windows Server that is supported within your organization and prioritize decommissioning of legacy operating systems in the domain controller population. By keeping your domain controllers current and eliminating legacy

domain controllers, you can often take advantage of new functionality and security that may not be available in domains or forests with domain controllers running legacy operating system. Domain controllers should be freshly installed and promoted rather than upgraded from previous operating systems or server roles; that is, do not perform in-place upgrades of domain controllers or run the AD DS Installation Wizard on servers on which the operating system is not freshly installed. By implementing freshly installed domain controllers, you ensure that legacy files and settings are not inadvertently left on domain controllers, and you simplify the enforcement of consistent, secure domain controller configuration.

Secure Configuration of Domain Controllers

A number of freely available tools, some of which are installed by default in Windows, can be used to create an initial security configuration baseline for domain controllers that can subsequently be enforced by GPOs. These tools are described here.

Security Configuration Wizard

All domain controllers should be locked down upon initial build. This can be achieved using the Security Configuration Wizard that ships natively in Windows Server to configure service, registry, system, and WFAS settings on a “base build” domain controller. Settings can be saved and exported to a GPO that can be linked to the Domain Controllers OU in each domain in the forest to enforce consistent configuration of domain controllers. If your domain contains multiple versions of Windows operating systems, you can configure Windows Management Instrumentation (WMI) filters to apply GPOs only to the domain controllers running the corresponding version of the operating system.

Microsoft Security Compliance Manager

Microsoft Security Compliance Manager domain controller settings can be combined with Security Configuration Wizard settings to produce comprehensive configuration baselines for domain controllers that are deployed and enforced by GPOs deployed at the Domain Controllers OU in Active Directory.

AppLocker

AppLocker or a third-party application whitelisting tool should be used to configure services and applications that are permitted to run on domain controllers, and these permitted applications and services should be comprised only of what is required for the computer to host AD DS and possibly DNS, plus any system security software such as antivirus software. By whitelisting permitted applications on domain controllers, an additional layer of security is added so that even if an unauthorized application is installed on a domain controller, the application cannot run.

RDP Restrictions

Group Policy Objects that link to all domain controllers OUs in a forest should be configured to allow RDP connections only from authorized users and systems—that is, jump servers. This can be achieved through a combination of user rights settings and WFAS configuration and should be implemented in GPOs so that the policy is consistently applied. If it is bypassed, the next Group Policy refresh returns the system to its proper configuration.

Patch and Configuration Management for Domain Controllers

Although it may seem counterintuitive, you should consider patching domain controllers and other critical infrastructure components separately from your general Windows infrastructure. If you leverage enterprise configuration management software for all computers in your infrastructure, compromise of the systems management software can be used to compromise or destroy all infrastructure components managed by that software. By separating patch and systems management for domain controllers from the general population, you can reduce the amount of software installed on domain controllers, in addition to tightly controlling their management.

Blocking Internet Access for Domain Controllers

One of the checks that is performed as part of an Active Directory Security Assessment is the use and configuration of Internet Explorer on domain controllers. Internet Explorer (or any other web browser) should not be used on domain controllers, but analysis of thousands of domain controllers has revealed numerous cases in which privileged users used Internet Explorer to browse the organization's intranet or the Internet.

As previously described in the "Misconfiguration" section of Avenues to Compromise, browsing the Internet (or an infected intranet) from one of the most powerful computers in a Windows infrastructure using a highly privileged account (which are the only accounts permitted to log on locally to domain controllers by default) presents an extraordinary risk to an organization's security. Whether via a drive by download or by download of malware-infected "utilities," attackers can gain access to everything they need to completely compromise or destroy the Active Directory environment.

Although Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and current versions of Internet Explorer offer a number of protections against malicious downloads, in most cases in which domain controllers and privileged accounts had been used to browse the Internet, the domain controllers were running Windows Server 2003, or protections offered by newer operating systems and browsers had been intentionally disabled.

Launching web browsers on domain controllers should be prohibited not only by policy, but by technical controls, and domain controllers should not be permitted to access the Internet. If your domain controllers need to replicate across sites, you should implement secure connections between the sites. Although detailed configuration instructions are outside the scope of this document, you can implement a number of controls to restrict the ability of domain controllers to be misused or misconfigured and subsequently compromised.

Perimeter Firewall Restrictions

Perimeter firewalls should be configured to block outbound connections from domain controllers to the Internet. Although domain controllers may need to communicate across site boundaries, perimeter firewalls can be configured to allow intersite communication by following the guidelines provided in How to configure a firewall for domains and trusts on the Microsoft Support website.

DC Firewall Configurations

As described earlier, you should use the Security Configuration Wizard to capture configuration settings for the Windows Firewall with Advanced Security on domain controllers. You should review the output of Security Configuration Wizard to ensure that the firewall configuration settings meet your organization's requirements, and then use GPOs to enforce configuration settings.

Preventing Web Browsing from Domain Controllers

You can use a combination of AppLocker configuration, "black hole" proxy configuration, and WFAS configuration to prevent domain controllers from accessing the Internet and to prevent the use of web browsers on domain controllers.

Monitoring Active Directory for Signs of Compromise

Law Number Five: Eternal vigilance is the price of security. – 10 Immutable Laws of Security Administration

A solid event log monitoring system is a crucial part of any secure Active Directory design. Many computer security compromises could be discovered early in the event if the victims enacted appropriate event log monitoring and alerting. Independent reports have long supported this conclusion. For example, the 2009 Verizon Data Breach Report states:

"The apparent ineffectiveness of event monitoring and log analysis continues to be somewhat of an enigma. The opportunity for detection is there; investigators noted that 66 percent of victims had sufficient evidence available within their logs to discover the breach had they been more diligent in analyzing such resources."

This lack of monitoring active event logs remains a consistent weakness in many companies' security defense plans. The 2012 Verizon Data Breach report found that even though 85 percent of breaches took several weeks to be noticed, 84 percent of victims had evidence of the breach in their event logs.

Windows Audit Policy

The following are links to the Microsoft official enterprise support blog. The content of these blogs provides advice, guidance, and recommendations about auditing that will assist you in enhancing the security of your Active Directory infrastructure and are a valuable resource when designing an audit policy.

- Global Object Access Auditing is Magic – describes a control mechanism called Advanced Audit Policy Configuration that was added to Windows® 7 and Windows Server 2008 R2 that lets you set what types of data you wanted to audit easily and not juggle scripts and auditpol.exe.
- Introducing Auditing Changes in Windows 2008 – introduces the auditing changes made in Windows Server 2008.
- Cool Auditing Tricks in Vista and 2008 – explains interesting auditing features of Windows Vista and Windows Server 2008 that can be used for troubleshooting problems or seeing what is happening in your environment.
- One-Stop Shop for Auditing in Windows Server 2008 and Windows Vista – contains a compilation of auditing features and information contained in Windows Server 2008 and Windows Vista.

The following links provide information about improvements to Windows auditing in Windows 8 and Windows Server 2012, and information about AD DS auditing in Windows Server 2008.

- What's New in Security Auditing – provides an overview of new security auditing features in Windows 8 and Windows Server 2012.
- AD DS Auditing Step-by-Step Guide – describes the new Active Directory Domain Services (AD DS) auditing feature in Windows Server 2008. It also provides procedures to implement this new feature.

Windows Audit Categories

Prior to Windows Vista and Windows Server 2008, Windows had only nine event log audit policy categories:

- Account Logon Events
- Account Management
- Directory Service Access
- Logon Events
- Object Access
- Policy Change
- Privilege Use
- Process Tracking
- System Events

These nine traditional audit categories comprise an audit policy. Each audit policy category can be enabled for Success, Failure, or Success and Failure events. Their descriptions are included in the next section.

Audit Policy Category Descriptions

The audit policy categories enable the following event log message types.

Audit Account Logon Events

Reports each instance of a security principal (for example, user, computer, or service account) that is logging on to or logging off from one computer in which another computer is used to validate the account. Account logon events are generated when a domain security principal account is authenticated on a domain controller. Authentication of a local user on a local computer generates a logon event that is logged in the local security log. No account logoff events are logged.

This category generates a lot of “noise” because Windows is constantly having accounts logging on to and off of the local and remote computers during the normal course of business. Still, any security plan should include the success and failure of this audit category.

Audit Account Management

This audit setting determines whether to track management of users and groups. For example, users and groups should be tracked when a user or computer account, a security group, or a distribution group is created, changed, or deleted; when a user or computer account is renamed, disabled, or enabled; or when a user or computer password is changed. An event can be generated for users or groups that are added to or removed from other groups.

Audit Directory Service Access

This policy setting determines whether to audit security principal access to an Active Directory object that has its own specified system access control list (SACL). In general, this category should only be enabled on domain controllers. When enabled, this setting generates a lot of “noise.”

Audit Logon Events

Logon events are generated when a local security principal is authenticated on a local computer. Logon Events records domain logons that occur on the local computer. Account logoff events are not generated. When enabled, Logon Events generates a lot of “noise,” but they should be enabled by default in any security auditing plan.

Audit Object Access

Object Access can generate events when subsequently defined objects with auditing enabled are accessed (for example, Opened, Read, Renamed, Deleted, or Closed). After the main auditing category is enabled, the administrator must individually define which objects will have auditing enabled. Many Windows system objects come with auditing enabled, so enabling this category will usually begin to generate events before the administrator has defined any.

This category is very “noisy” and will generate five to ten events for each object access. It can be difficult for administrators new to object auditing to gain useful information. It should only be enabled when needed.

Auditing Policy Change

This policy setting determines whether to audit every incidence of a change to user rights assignment policies, Windows Firewall policies, Trust policies, or changes to the audit policy. This category should be enabled on all computers. It generates very little noise.

Audit Privilege Use

There are dozens of user rights and permissions in Windows (for example, Logon as a Batch Job and Act as Part of the Operating System). This policy setting determines whether to audit each instance of a security principal by exercising a user right or privilege. Enabling this category results in a lot of "noise," but it can be helpful in tracking security principal accounts using elevated privileges.

Audit Process Tracking

This policy setting determines whether to audit detailed process tracking information for events such as program activation, process exit, handle duplication, and indirect object access. It is useful for tracking malicious users and the programs they use.

Enabling Audit Process Tracking generates a large number of events, so typically it is set to **No Auditing**. However, this setting can provide a great benefit during an incident response from the detailed log of the processes started and the time they were launched. For domain controllers and other single-role infrastructure servers, this category can be safely turned on all the time. Single role servers do not generate much process tracking traffic during the normal course of their duties. As such, they can be enabled to capture unauthorized events if they occur.

System Events Audit

System Events is almost a generic catch-all category, registering various events that impact the computer, its system security, or the security log. It includes events for computer shutdowns and restarts, power failures, system time changes, authentication package initializations, audit log clearings, impersonation issues, and a host of other general events. In general, enabling this audit category generates a lot of "noise," but it generates enough very useful events that it is difficult to ever recommend not enabling it.

Advanced Audit Policies

Starting with Windows Vista and Windows Server 2008, Microsoft improved the way event log category selections can be made by creating subcategories under each main audit category. Subcategories allow auditing to be far more granular than it could otherwise by using the main categories. By using subcategories, you can enable only portions of a particular main category, and skip generating events for which you have no use. Each audit policy subcategory can be enabled for Success, Failure, or Success and Failure events.

To list all the available auditing subcategories, review the Advanced Audit Policy container in a Group Policy Object, or type the following at a command prompt on any computer running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, Windows 8, Windows 7, or Windows Vista:

auditpol /list /subcategory:*

To get a list of currently configured auditing subcategories on a computer running Windows Server 2012, Windows Server 2008 R2, or Windows 2008, type the following:

auditpol /get /category:*

The following screenshot shows an example of auditpol.exe listing the current audit policy.

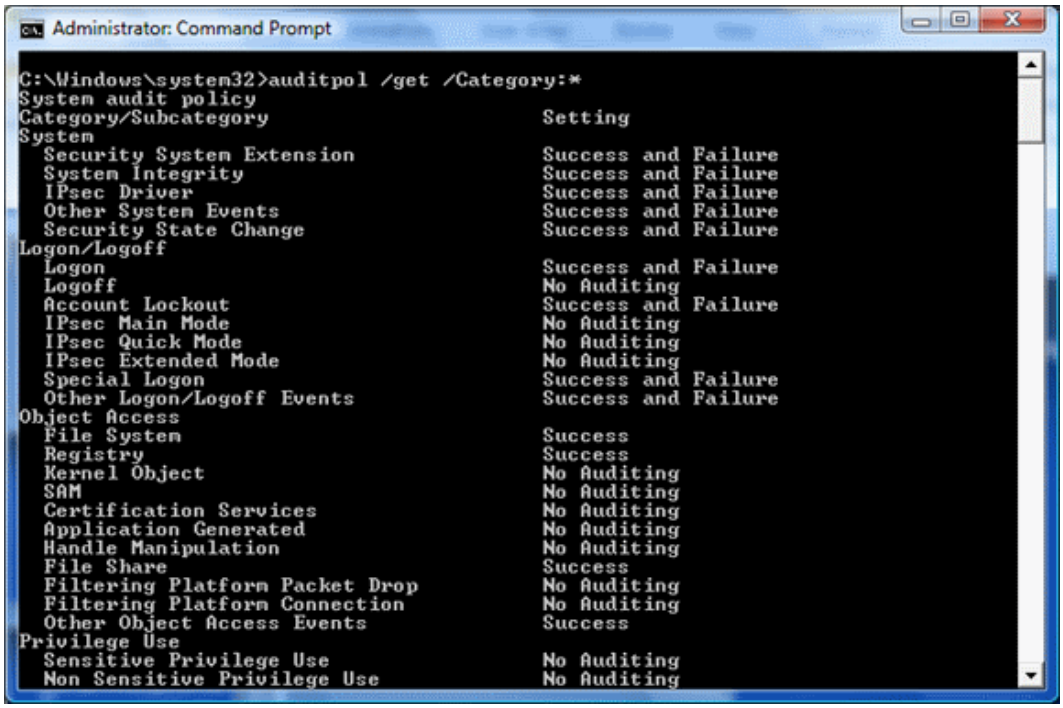

Note

Group Policy does not always accurately report the status of all enabled auditing policies, whereas auditpol.exe does. See Getting the Effective Audit Policy in Windows 7 and 2008 R2 for more details.

Each main category has multiple subcategories. Below is a list of categories, their subcategories, and a description of their functions.

Auditing Subcategories Descriptions

Audit policy subcategories enable the following event log message types:

Account Logon

Credential Validation

This subcategory reports the results of validation tests on credentials submitted for a user account logon request. These events occur on the computer that is authoritative for the

credentials. For domain accounts, the domain controller is authoritative, whereas for local accounts, the local computer is authoritative.

In domain environments, most of the account logon events are logged in the security log of the domain controllers that are authoritative for the domain accounts. However, these events can occur on other computers in the organization when local accounts are used to log on.

Kerberos Service Ticket Operations

This subcategory reports events generated by Kerberos ticket request processes on the domain controller that is authoritative for the domain account.

Kerberos Authentication Service

This subcategory reports events generated by the Kerberos authentication service. These events occur on the computer that is authoritative for the credentials.

Other Account Logon Events

This subcategory reports the events that occur in response to credentials submitted for a user account logon request that do not relate to credential validation or Kerberos tickets. These events occur on the computer that is authoritative for the credentials. For domain accounts, the domain controller is authoritative, whereas for local accounts, the local computer is authoritative.

In domain environments, most account logon events are logged in the security log of the domain controllers that are authoritative for the domain accounts. However, these events can occur on other computers in the organization when local accounts are used to log on. Examples can include the following:

- Remote Desktop Services session disconnections
- New Remote Desktop Services sessions
- Locking and unlocking a workstation
- Invoking a screen saver
- Dismissing a screen saver
- Detection of a Kerberos replay attack, in which a Kerberos request with identical information is received twice
- Access to a wireless network granted to a user or computer account
- Access to a wired 802.1x network granted to a user or computer account

Account Management

User Account Management

This subcategory reports each event of user account management, such as when a user account is created, changed, or deleted; a user account is renamed, disabled, or enabled; or a password is set or changed. If this audit policy setting is enabled, administrators can track events to detect malicious, accidental, and authorized creation of user accounts.

Computer Account Management

This subcategory reports each event of computer account management, such as when a computer account is created, changed, deleted, renamed, disabled, or enabled.

Security Group Management

This subcategory reports each event of security group management, such as when a security group is created, changed, or deleted or when a member is added to or removed from a security group. If this audit policy setting is enabled, administrators can track events to detect malicious, accidental, and authorized creation of security group accounts.

Distribution Group Management

This subcategory reports each event of distribution group management, such as when a distribution group is created, changed, or deleted or when a member is added to or removed from a distribution group. If this audit policy setting is enabled, administrators can track events to detect malicious, accidental, and authorized creation of group accounts.

Application Group Management

This subcategory reports each event of application group management on a computer, such as when an application group is created, changed, or deleted or when a member is added to or removed from an application group. If this audit policy setting is enabled, administrators can track events to detect malicious, accidental, and authorized creation of application group accounts.

Other Account Management Events

This subcategory reports other account management events.

Detailed Process Tracking

Process Creation

This subcategory reports the creation of a process and the name of the user or program that created it.

Process Termination

This subcategory reports when a process terminates.

DPAPI Activity

This subcategory reports encrypt or decrypt calls into the data protection application programming interface (DPAPI). DPAPI is used to protect secret information such as stored password and key information.

RPC Events

This subcategory reports remote procedure call (RPC) connection events.

Directory Service Access

Directory Service Access

This subcategory reports when an AD DS object is accessed. Only objects with configured SACLs cause audit events to be generated, and only when they are accessed in a manner that matches the SACL entries. These events are similar to the directory service access events in earlier versions of Windows Server. This subcategory applies only to domain controllers.

Directory Service Changes

This subcategory reports changes to objects in AD DS. The types of changes that are reported are create, modify, move, and undelete operations that are performed on an object. Directory service change auditing, where appropriate, indicates the old and new values of the changed properties of the objects that were changed. Only objects with SACLs cause audit events to be generated, and only when they are accessed in a manner that matches their SACL entries. Some objects and properties do not cause audit events to be generated due to settings on the object class in the schema. This subcategory applies only to domain controllers.

Directory Service Replication

This subcategory reports when replication between two domain controllers begins and ends.

Detailed Directory Service Replication

This subcategory reports detailed information about the information replicated between domain controllers. These events can be very high in volume.

Logon/Logoff

Logon

This subcategory reports when a user attempts to log on to the system. These events occur on the accessed computer. For interactive logons, the generation of these events occurs on the computer that is logged on to. If a network logon takes place to access a share, these events generate on the computer that hosts the accessed resource. If this setting is configured to **No auditing**, it is difficult or impossible to determine which user has accessed or attempted to access organization computers.

Network Policy Server

This subcategory reports events generated by RADIUS (IAS) and Network Access Protection (NAP) user access requests. These requests can be **Grant**, **Deny**, **Discard**, **Quarantine**, **Lock**, and **Unlock**. Auditing this setting will result in a medium or high volume of records on NPS and IAS servers.

IPsec Main Mode

This subcategory reports the results of Internet Key Exchange (IKE) protocol and Authenticated Internet Protocol (AuthIP) during Main Mode negotiations.

IPsec Extended Mode

This subcategory reports the results of AuthIP during Extended Mode negotiations.

Other Logon/Logoff Events

This subcategory reports other logon and logoff-related events, such as Remote Desktop Services session disconnects and reconnects, using RunAs to run processes under a different account, and locking and unlocking a workstation.

Logoff

This subcategory reports when a user logs off the system. These events occur on the accessed computer. For interactive logons, the generation of these events occurs on the computer that is logged on to. If a network logon takes place to access a share, these events generate on the computer that hosts the accessed resource. If this setting is configured to **No auditing**, it is difficult or impossible to determine which user has accessed or attempted to access organization computers.

Account Lockout

This subcategory reports when a user's account is locked out as a result of too many failed logon attempts.

IPsec Quick Mode

This subcategory reports the results of IKE protocol and AuthIP during Quick Mode negotiations.

Special Logon

This subcategory reports when a special logon is used. A special logon is a logon that has administrator equivalent privileges and can be used to elevate a process to a higher level.

Policy Change

Audit Policy Change

This subcategory reports changes in audit policy including SACL changes.

Authentication Policy Change

This subcategory reports changes in authentication policy.

Authorization Policy Change

This subcategory reports changes in authorization policy including permissions (DACL) changes.

MPSSVC Rule-Level Policy Change

This subcategory reports changes in policy rules used by the Microsoft Protection Service (MPSSVC.exe). This service is used by Windows Firewall.

Filtering Platform Policy Change

This subcategory reports the addition and removal of objects from WFP, including startup filters. These events can be very high in volume.

Other Policy Change Events

This subcategory reports other types of security policy changes such as configuration of the Trusted Platform Module (TPM) or cryptographic providers.

Privilege Use

Sensitive Privilege Use

This subcategory reports when a user account or service uses a sensitive privilege. A sensitive privilege includes the following user rights: act as part of the operating system, back up files and directories, create a token object, debug programs, enable computer and user accounts to be trusted for delegation, generate security audits, impersonate a client after authentication, load and unload device drivers, manage auditing and security log, modify firmware environment values, replace a process-level token, restore files and directories, and take ownership of files or other objects. Auditing this subcategory will create a high volume of events.

Nonsensitive Privilege Use

This subcategory reports when a user account or service uses a nonsensitive privilege. A nonsensitive privilege includes the following user rights: access Credential Manager as a trusted caller, access this computer from the network, add workstations to domain, adjust memory quotas for a process, allow log on locally, allow log on through Remote Desktop Services, bypass traverse checking, change the system time, create a pagefile, create global objects, create permanent shared objects, create symbolic links, deny access this computer from the network, deny log on as a batch job, deny log on as a service, deny log on locally, deny log on through Remote Desktop Services, force shutdown from a remote system, increase a process working set, increase scheduling priority, lock pages in memory, log on as a batch job, log on as a service, modify an object label, perform volume maintenance tasks, profile single process, profile system performance, remove computer from docking station, shut down the system, and synchronize directory service data. Auditing this subcategory will create a very high volume of events.

Other Privilege Use Events

This security policy setting is not currently used.

Object Access

File System

This subcategory reports when file system objects are accessed. Only file system objects with SACLs cause audit events to be generated, and only when they are accessed in a manner matching their SACL entries. By itself, this policy setting will not cause auditing of any events. It

determines whether to audit the event of a user who accesses a file system object that has a specified system access control list (SACL), effectively enabling auditing to take place.

If the audit object access setting is configured to **Success**, an audit entry is generated each time that a user successfully accesses an object with a specified SACL. If this policy setting is configured to **Failure**, an audit entry is generated each time that a user fails in an attempt to access an object with a specified SACL.

Registry

This subcategory reports when registry objects are accessed. Only registry objects with SACLs cause audit events to be generated, and only when they are accessed in a manner matching their SACL entries. By itself, this policy setting will not cause auditing of any events.

Kernel Object

This subcategory reports when kernel objects such as processes and mutexes are accessed. Only kernel objects with SACLs cause audit events to be generated, and only when they are accessed in a manner matching their SACL entries. Typically kernel objects are only given SACLs if the AuditBaseObjects or AuditBaseDirectories auditing options are enabled.

SAM

This subcategory reports when local Security Accounts Manager (SAM) authentication database objects are accessed.

Certification Services

This subcategory reports when Certification Services operations are performed.

Application Generated

This subcategory reports when applications attempt to generate audit events by using the Windows auditing application programming interfaces (APIs).

Handle Manipulation

This subcategory reports when a handle to an object is opened or closed. Only objects with SACLs cause these events to be generated, and only if the attempted handle operation matches the SACL entries. Handle Manipulation events are only generated for object types where the corresponding object access subcategory is enabled (for example, file system or registry).

File Share

This subcategory reports when a file share is accessed. By itself, this policy setting will not cause auditing of any events. It determines whether to audit the event of a user who accesses a file share object that has a specified system access control list (SACL), effectively enabling auditing to take place.

Filtering Platform Packet Drop

This subcategory reports when packets are dropped by Windows Filtering Platform (WFP). These events can be very high in volume.

Filtering Platform Connection

This subcategory reports when connections are allowed or blocked by WFP. These events can be high in volume.

Other Object Access Events

This subcategory reports other object access-related events such as Task Scheduler jobs and COM+ objects.

System

Security State Change

This subcategory reports changes in security state of the system, such as when the security subsystem starts and stops.

Security System Extension

This subcategory reports the loading of extension code such as authentication packages by the security subsystem.

System Integrity

This subcategory reports on violations of integrity of the security subsystem.

IPsec Driver

This subcategory reports on the activities of the Internet Protocol security (IPsec) driver.

Other System Events

This subcategory reports on other system events.

For more information about the subcategory descriptions, refer to the Microsoft Security Compliance Manager tool.

Each organization should review the previous covered categories and subcategories and enable the ones which best fit their environment. Changes to audit policy should always be tested prior to deployment in a production environment.

Configuring Windows Audit Policy

Windows audit policy can be set using group policies, auditpol.exe, APIs, or registry edits. The recommended methods for configuring audit policy for most companies are Group Policy or auditpol.exe. Setting a system's audit policy requires administrator-level account permissions or the appropriate delegated permissions.

Note

The **Manage auditing and security log** privilege must be given to security principals (Administrators have it by default) to allow the modification of object access auditing options of individual resources, such as files, Active Directory objects, and registry keys.

Setting Windows Audit Policy by Using Group Policy

To set audit policy using group policies, configure the appropriate audit categories located under **Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy** (see the following screenshot for an example from the Local Group Policy Editor (gpedit.msc)). Each audit policy category can be enabled for **Success**, **Failure**, or **Success** and Failure events.

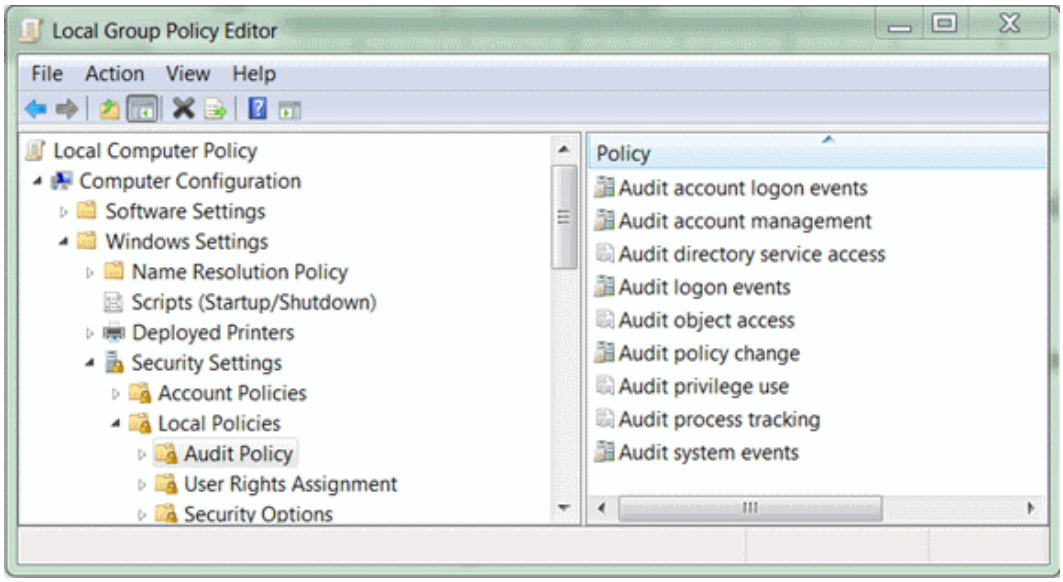

Advanced Audit Policy can be set by using Active Directory or local group policies. To set Advanced Audit Policy, configure the appropriate subcategories located under **Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy** (see the following screenshot for an example from the Local Group Policy Editor (gpedit.msc)). Each audit policy subcategory can be enabled for **Success**, **Failure**, or **Success** and **Failure** events.

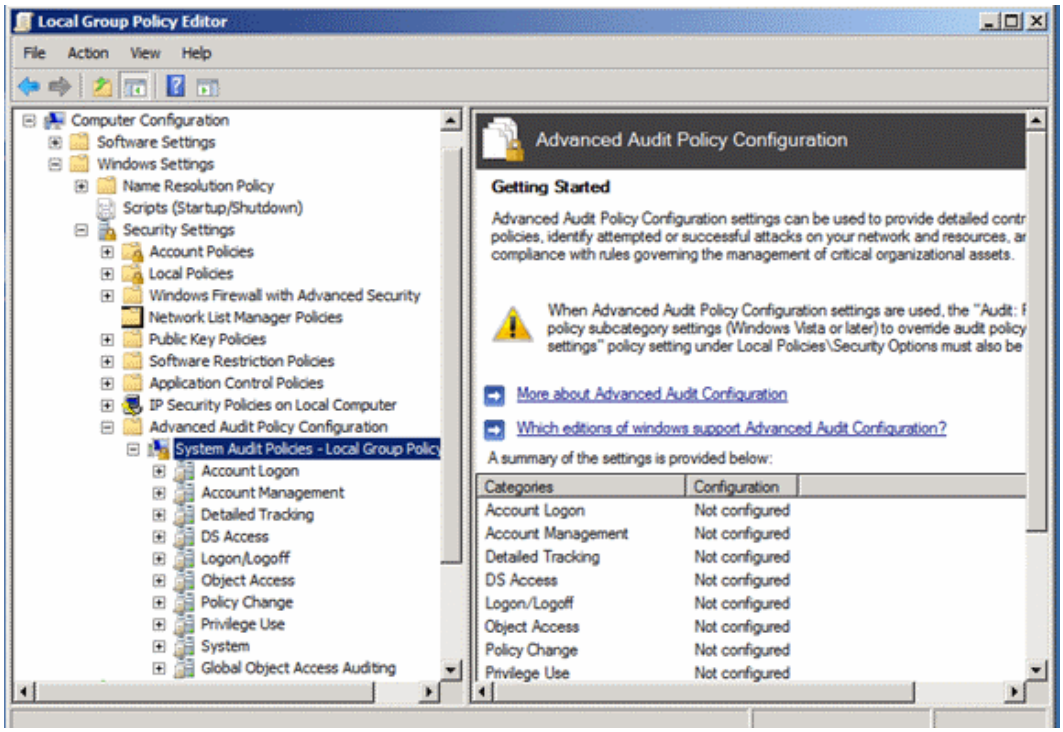

Setting Windows Audit Policy Using Auditpol.exe

Auditpol.exe (for setting Windows audit policy) was introduced in Windows Server 2008 and Windows Vista. Initially, only auditpol.exe could be used to set Advanced Audit Policy, but Group Policy can be used in Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, Windows 8, and Windows 7.

Auditpol.exe is a command-line utility. The syntax is as follows:

auditpol /set /<Category|Subcategory>:<audit category> /<success|failure:> /<enable|disable>

Auditpol.exe syntax examples:

auditpol /set /subcategory:”user account management” /success:enable /failure:enable

auditpol /set /subcategory:”logon” /success:enable /failure:enable

auditpol /set /subcategory:”IPSEC Main Mode” /failure:enable

Note

Auditpol.exe sets Advanced Audit Policy locally. If local policy conflicts with Active Directory or local Group Policy, Group Policy settings usually prevail over auditpol.exe settings. When multiple group or local policy conflicts exist, only one policy will prevail (that is, replace). Audit policies will not merge.

Scripting Auditpol

Microsoft provides a sample script for administrators who want to set Advanced Audit Policy by using a script instead of manually typing in each auditpol.exe command.

Note Group Policy does not always accurately report the status of all enabled auditing policies, whereas auditpol.exe does. See Getting the Effective Audit Policy in Windows 7 and Windows 2008 R2 for more details.

Other Auditpol Commands

Auditpol.exe can be used to save and restore a local audit policy, and to view other auditing related commands. Here are the other **auditpol** commands.

auditpol /clear - Used to clear and reset local audit policies

auditpol /backup /file:<filename> - Used to back up a current local audit policy to a binary file

auditpol /restore /file:<filename> - Used to import a previously saved audit policy file to a local audit policy

auditpol /<get/set> /option:<CrashOnAuditFail> /<enable/disable> - If this audit policy setting is enabled, it causes the system to immediately stop (with STOP: C0000244 {Audit Failed} message) if a security audit cannot be logged for any reason. Typically, an event fails to be logged when the security audit log is full and the retention method specified for the security log is **Do Not Overwrite Events** or **Overwrite Events by Days**. Typically it is only enabled by environments that need higher assurance that the security log is logging. If enabled, administrators must closely watch security log size and rotate logs as required. It can also be set with Group Policy by modifying the security option **Audit: Shut down system immediately if unable to log security audits** (default=disabled).

Auditpol /<get/set> /option:<AuditBaseObjects> /<enable/disable> - This audit policy setting determines whether to audit the access of global system objects. If this policy is enabled, it causes system objects, such as mutexes, events, semaphores, and DOS devices to be created with a default system access control list (SACL). Most administrators consider auditing global system objects to be too “noisy,” and they will only enable it if malicious hacking is suspected. Only named objects are given a SACL. If the audit object access audit policy (or Kernel Object audit subcategory) is also enabled, access to these system objects is audited. When configuring this security setting, changes will not take effect until you restart Windows. This policy can also be set with Group Policy by modifying the security option Audit the access of global system objects (default=disabled).

auditpol /<get/set> /option:<AuditBaseDirectories> /<enable/disable> - This audit policy setting specifies that named kernel objects (such as mutexes and semaphores) are to be given SACLs when they are created. AuditBaseDirectories affects container objects while AuditBaseObjects affects objects that cannot contain other objects.

Auditpol /<get/set> /option:<FullPrivilegeAuditing> /<enable/disable> -

This audit policy setting specifies whether the client generates an event when one or more of these privileges are assigned to a user security token: AssignPrimaryTokenPrivilege, AuditPrivilege, BackupPrivilege, CreateTokenPrivilege, DebugPrivilege,

EnableDelegationPrivilege, ImpersonatePrivilege, LoadDriverPrivilege, RestorePrivilege, SecurityPrivilege, SystemEnvironmentPrivilege, TakeOwnershipPrivilege, and TcbPrivilege. If this option is not enabled (default=Disabled), the BackupPrivilege and RestorePrivilege privileges are not recorded. Enabling this option can make the security log extremely noisy (sometimes hundreds of events a second) during a backup operation. This policy can also be set with Group Policy by modifying the security option **Audit: Audit the use of Backup and Restore privilege**.

Note

Some information provided here was taken from the Microsoft Audit Option Type and the Microsoft SCM tool.

Enforcing Traditional Auditing or Advanced Auditing

In Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista, administrators can choose to enable the nine traditional categories or to use the subcategories. It's a binary choice that must be made in each Windows system. Either the main categories can be enabled or the subcategories—it cannot be both.

To prevent the legacy traditional category policy from overwriting audit policy subcategories, you must enable the **Force audit policy subcategory settings(Windows Vista or later) to override audit policy category settings** policy setting located under **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**.

We recommend that the subcategories be enabled and configured instead of the nine main categories. This requires that a Group Policy setting be enabled (to allow subcategories to override the auditing categories) along with configuring the different subcategories that support auditing policies.

Auditing subcategories can be configured by using several methods, including Group Policy and the command-line program, auditpol.exe.

For more information about Windows auditing, see the following articles:

- Advanced Security Auditing in Windows 7 and Windows Server 2008 R2
- Auditing and Compliance in Windows Server 2008
- How to use Group Policy to configure detailed security auditing settings for Windows Vista-based and Windows Server 2008-based computers in a Windows Server 2008 domain, in a Windows Server 2003 domain, or in a Windows 2000 domain
- Advanced Security Audit Policy Step-by-Step Guide

Audit Policy Recommendations

This section addresses the Windows default audit policy settings, baseline recommended audit policy settings, and the more aggressive recommendations from Microsoft, for workstation and server products.

The SCM baseline recommendations shown here, along with the settings we recommend to help detect compromise, are intended only to be a starting baseline guide to administrators. Each organization must make its own decisions regarding the threats they face, their acceptable risk tolerances, and what audit policy categories or subcategories they should enable. For further information about threats, refer to the Threats and Countermeasures Guide. Administrators without a thoughtful audit policy in place are encouraged to start with the settings recommended here, and then to modify and test, prior to implementing in their production environment.

The recommendations are for enterprise-class computers, which Microsoft defines as computers that have average security requirements and require a high level of operational functionality. Entities needing higher security requirements should consider more aggressive audit policies.

Note

Microsoft Windows defaults and baseline recommendations were taken from the Microsoft Security Compliance Manager tool.

The following baseline audit policy settings are recommended for normal security computers that are not known to be under active, successful attack by determined adversaries or malware.

Recommended Audit Policies by Operating System

This section contains tables that list the audit setting recommendations that apply to the following operating systems:

- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2008
- Windows 8
- Windows 7

These tables contain the Windows default setting, the baseline recommendations, and the stronger recommendations for these operating systems.

Audit Policy Tables Legend	
Notation	**Recommendation**
YES	Enable in general scenarios
NO	Do **not** enable in general scenarios
IF	Enable if needed for a specific scenario, or if a role or feature for which auditing is desired is installed on the machine
DC	Enable on domain controllers

[Blank]	No recommendation

Windows 8 and Windows 7 Audit Settings Recommendations

Audit Policy						
Audit Policy Category or Subcategory	**Windows Default**	**Baseline Recommendation**	**Stronger Recommendation**			
	Success	**Failure**	**Success**	**Failure**	**Success**	**Failure**
Account Logon						
Audit Credential Validation	NO	NO	YES	NO	YES	YES
Audit Kerberos Authentication Service					YES	YES
Audit Kerberos Service Ticket Operations					YES	YES
Audit Other Account Logon Events					YES	YES
Account Management						
Audit Application Group Management						
Audit Computer Account Management			YES	NO	YES	YES
Audit						

Distribution Group Management						
Audit Other Account Management Events			YES	NO	YES	YES
Audit Security Group Management			YES	NO	YES	YES
Audit User Account Management	YES	NO	YES	NO	YES	YES
Detailed Tracking						
Audit DPAPI Activity					YES	YES
Audit Process Creation			YES	NO	YES	YES
Audit Process Termination						
Audit RPC Events						
DS Access						
Audit Detailed Directory Service Replication						
Audit Directory Service Access						
Audit Directory Service Changes						
Audit						

Directory Service Replication						
Logon and Logoff						
Audit Account Lockout	YES	NO			YES	NO
Audit User/Device Claims						
Audit IPsec Extended Mode						
Audit IPsec Main Mode					IF	IF
Audit IPsec Quick Mode						
Audit Logoff	YES	NO	YES	NO	YES	NO
Audit Logon	YES	NO	YES	NO	YES	YES
Audit Network Policy Server	YES	YES				
Audit Other Logon/Logoff Events					YES	YES
Audit Special Logon	YES	NO	YES	NO	YES	YES
Object Access						
Audit Application Generated						
Audit Certification Services						
Audit Detailed File Share						

Audit File Share						
Audit File System						
Audit Filtering Platform Connection						
Audit Filtering Platform Packet Drop						
Audit Handle Manipulation						
Audit Kernel Object						
Audit Other Object Access Events						
Audit Registry						
Audit Removable Storage						
Audit SAM						
Audit Central Access Policy Staging						
Policy Change						
Audit Audit Policy Change	YES	NO	YES	YES	YES	YES
Audit Authentication Policy Change	YES	NO	YES	NO	YES	YES
Audit						

Authorization Policy Change						
Audit Filtering Platform Policy Change						
Audit MPSSVC Rule-Level Policy Change					YES	
Audit Other Policy Change Events						
Privilege Use						
Audit Non Sensitive Privilege Use						
Audit Other Privilege Use Events						
Audit Sensitive Privilege Use						
System						
Audit IPsec Driver			YES	YES	YES	YES
Audit Other System Events	YES	YES				
Audit Security State Change	YES	NO	YES	YES	YES	YES
Audit Security System Extension			YES	YES	YES	YES

Audit System Integrity	YES	YES	YES	YES	YES	YES
Global Object Access Auditing						
Audit IPsec Driver						
Audit Other System Events						
Audit Security State Change						
Audit Security System Extension						
Audit System Integrity						

Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 Audit Settings Recommendations

Audit Policy						
Audit Policy Category or Subcategory	**Windows Default**	**Baseline Recommendation**	**Stronger Recommendation**			
	Success	**Failure**	**Success**	**Failure**	**Success**	**Failure**
Account Logon						
Audit Credential Validation	NO	NO	YES	YES	YES	YES
Audit Kerberos Authenticatio					YES	YES

n Service						
Audit Kerberos Service Ticket Operations					YES	YES
Audit Other Account Logon Events					YES	YES
Account Management						
Audit Application Group Management						
Audit Computer Account Management			YES	DC	YES	YES
Audit Distribution Group Management						
Audit Other Account Management Events			YES	YES	YES	YES
Audit Security Group Management			YES	YES	YES	YES
Audit User Account Management	YES	NO	YES	YES	YES	YES
Detailed Tracking						
Audit DPAPI Activity					YES	YES
Audit Process			YES	NO	YES	YES

Creation						
Audit Process Termination						
Audit RPC Events						
DS Access						
Audit Detailed Directory Service Replication						
Audit Directory Service Access			DC	DC	DC	DC
Audit Directory Service Changes			DC	DC	DC	DC
Audit Directory Service Replication						
Logon and Logoff						
Audit Account Lockout	YES	NO			YES	NO
Audit User/Device Claims						
Audit IPsec Extended Mode						
Audit IPsec Main Mode					IF	IF
Audit IPsec Quick Mode						

Audit Logoff	YES	NO	YES	NO	YES	NO
Audit Logon	YES	NO	YES	YES	YES	YES
Audit Network Policy Server	YES	YES				
Audit Other Logon/Logoff Events					YES	YES
Audit Special Logon	YES	NO	YES	NO	YES	YES
Object Access						
Audit Application Generated						
Audit Certification Services						
Audit Detailed File Share						
Audit File Share						
Audit File System						
Audit Filtering Platform Connection						
Audit Filtering Platform Packet Drop						
Audit Handle Manipulation						
Audit Kernel Object						
Audit Other Object Access						

Events						
Audit Registry						
Audit Removable Storage						
Audit SAM						
Audit Central Access Policy Staging						
Policy Change						
Audit Audit Policy Change	YES	NO	YES	YES	YES	YES
Audit Authentication Policy Change	YES	NO	YES	NO	YES	YES
Audit Authorization Policy Change						
Audit Filtering Platform Policy Change						
Audit MPSSVC Rule-Level Policy Change					YES	
Audit Other Policy Change Events						
Privilege Use						
Audit Non-						

Sensitive Privilege Use						
Audit Other Privilege Use Events						
Audit Sensitive Privilege Use			YES	YES	YES	YES
System						
Audit IPsec Driver			YES	YES	YES	YES
Audit Other System Events	YES	YES				
Audit Security State Change	YES	NO	YES	YES	YES	YES
Audit Security System Extension			YES	YES	YES	YES
Audit System Integrity	YES	YES	YES	YES	YES	YES
Global Object Access Auditing						
Audit IPsec Driver						
Audit Other System Events						
Audit Security State Change						
Audit Security System Extension						
Audit System						

Integrity						

Set Audit Policy on Workstations and Servers

All event log management plans should monitor workstations and servers. A common mistake is to only monitor servers or domain controllers. Because malicious hacking often initially occurs on workstations, not monitoring workstations is ignoring the best and earliest source of information.

Administrators should thoughtfully review and test any audit policy prior to implementation in their production environment.

Events to Monitor

A perfect event ID to generate a security alert should contain the following attributes:

- High likelihood that occurrence indicates unauthorized activity
- Low number of false positives
- Occurrence should result in an investigative/forensics response

Two types of events should be monitored and alerted:

1. Those events in which even a single occurrence indicates unauthorized activity
2. An accumulation of events above an expected and accepted baseline

An example of the first event is:

If Domain Admins (DAs) are forbidden from logging on to computers that are not domain controllers, a single occurrence of a DA member logging on to an end-user workstation should generate an alert and be investigated. This type of alert is easy to generate by using the Audit Special Logon event 4964 (Special groups have been assigned to a new logon). Other examples of single instance alerts include:

- If Server A should never connect to Server B, alert when they connect to each other.
- Alert if a normal end-user account is unexpectedly added to a sensitive security group.
- If employees in factory location A never work at night, alert when a user logs on at midnight.
- Alert if an unauthorized service is installed on a domain controller.
- Investigate if a regular end-user attempts to directly log on to a SQL Server for which they have no clear reason for doing so.
- If you have no members in your DA group, and someone adds themselves there, check it immediately.

An example of the second event is:

An aberrant number of failed logons could indicate a password guessing attack. For an enterprise to provide an alert for an unusually high number of failed logons, they must first understand the normal levels of failed logons within their environment prior to a malicious security event.

For a comprehensive list of events that you should include when you monitor for signs of compromise, please see Appendix L: Events to Monitor.

Active Directory Objects and Attributes to Monitor

The following are the accounts, groups, and attributes that you should monitor to help you detect attempts to compromise your Active Directory Domain Services installation.

- Systems for disabling or removal of antivirus and antimalware software (automatically restart protection when it is manually disabled)
- Administrator accounts for unauthorized changes
- Activities that are performed by using privileged accounts (automatically remove account when suspicious activities are completed or allotted time has expired)
- Privileged and VIP accounts in AD DS. Monitor for changes, particularly changes to attributes on the Account tab (for example, cn, name, sAMAccountName, userPrincipalName, or userAccountControl). In addition to monitoring the accounts, restrict who can modify the accounts to as small a set of administrative users as possible.

Refer to Appendix L: Events to Monitor for a list of recommended events to monitor, their criticality ratings, and an event message summary.

- Group servers by the classification of their workloads, which allows you to quickly identify the servers that should be the most closely monitored and most stringently configured
- Changes to the properties and membership of following AD DS groups: Enterprise Admins (EA), Domain Admins (DA), Administrators (BA), and Schema Admins (SA)
- Disabled privileged accounts (such as built-in Administrator accounts in Active Directory and on member systems) for enabling the accounts
- Management accounts to log all writes to the account
- Built-in Security Configuration Wizard to configure service, registry, audit, and firewall settings to reduce the server's attack surface. Use this wizard if you implement jump servers as part of your administrative host strategy.

Additional Information for Monitoring Active Directory Domain Services

Review the following links for additional information about monitoring AD DS:

- Global Object Access Auditing is Magic – Provides information about configuring and using Advanced Audit Policy Configuration that was added to Windows 7 and Windows Server 2008 R2.
- Introducing Auditing Changes in Windows 2008 – Introduces the auditing changes made in Windows 2008.
- Cool Auditing Tricks in Vista and 2008 – Explains interesting new features of auditing in Windows Vista and Windows Server 2008 that can be used for troubleshooting problems or seeing what's happening in your environment.
- One-Stop Shop for Auditing in Windows Server 2008 and Windows Vista – Contains a compilation of auditing features and information contained in Windows Server 2008 and Windows Vista.

- AD DS Auditing Step-by-Step Guide – Describes the new Active Directory Domain Services (AD DS) auditing feature in Windows Server 2008. It also provides procedures to implement this new feature.

General List of Security Event ID Recommendation Criticalities

All Event ID recommendations are accompanied by a criticality rating as follows:

High: Event IDs with a high criticality rating should always and immediately be alerted and investigated.

Medium: An Event ID with a medium criticality rating could indicate malicious activity, but it must be accompanied by some other abnormality (for example, an unusual number occurring in a particular time period, unexpected occurrences, or occurrences on a computer that normally would not be expected to log the event.). A medium-criticality event may also r be collected as a metric and compared over time.

Low: And Event ID with a low criticality events should not garner attention or cause alerts, unless correlated with medium or high criticality events.

These recommendations are meant to provide a baseline guide for the administrator. All recommendations should be thoroughly reviewed prior to implementation in a production environment.

Refer to Appendix L: Events to Monitor for a list of the recommended events to monitor, their criticality ratings, and an event message summary.

Planning for Compromise

Law Number One: Nobody believes anything bad can happen to them, until it does. – 10 Immutable Laws of Security Administration

Disaster recovery plans in many organizations focus on recovering from regional disasters or failures that result in loss of computing services. However, when working with compromised customers, we often find that recovering from intentional compromise is absent in their disaster recovery plans. This is particularly true when the compromise results in theft of intellectual property or intentional destruction that leverages logical boundaries (such as destruction of all Active Directory domains or all servers) rather than physical boundaries (such as destruction of a datacenter). Although an organization may have incident response plans that define initial activities to take when a compromise is discovered, these plans often omit steps to recover from a compromise that affects the entire computing infrastructure.

Because Active Directory provides rich identity and access management capabilities for users, servers, workstations, and applications, it is invariably targeted by attackers. If an attacker gains highly privileged access to an Active Directory domain or domain controller, that access can be leveraged to access, control, or even destroy the entire Active Directory forest.

This document has discussed some of the most common attacks against Windows and Active Directory and countermeasures you can implement to reduce your attack surface, but the only sure way to recover in the event of a complete compromise of Active Directory is to be prepared for the compromise before it happens. This section focuses less on technical implementation details than previous sections of this document, and more on high-level recommendations that you can use to create a holistic, comprehensive approach to secure and manage your organization's critical business and IT assets.

Whether your infrastructure has never been attacked, has resisted attempted breaches, or has succumbed to attacks and been fully compromised, you should plan for the inevitable reality that you will be attacked again and again. It is not possible to prevent attacks, but it may indeed be possible to prevent significant breaches or wholesale compromise. Every organization should closely evaluate their existing risk management programs, and make necessary adjustments to help reduce their overall level of vulnerability by making balanced investments in prevention, detection, containment, and recovery.

To create effective defenses while still providing services to the users and businesses that depend on your infrastructure and applications, you may need to consider novel ways to prevent, detect, and contain compromise in your environment, and then recover from the compromise. The approaches and recommendations in this document may not help you repair a compromised Active Directory installation, but can help you secure your next one.

Recommendations for recovering an Active Directory forest are presented in Windows Server 2012: Planning for Active Directory Forest Recovery. You may be able to prevent your new environment from being completely compromised, but even if you can't, you will have tools to recover and regain control of your environment.

Rethinking the Approach

Law Number Eight: The difficulty of defending a network is directly proportional to its complexity. – 10 Immutable Laws of Security Administration

It is generally well-accepted that if an attacker has obtained SYSTEM, Administrator, root, or equivalent access to a computer, regardless of operating system, that computer can no longer be considered trustworthy, no matter how many efforts are made to "clean" the system. Active Directory is no different. If an attacker has obtained privileged access to a domain controller or a highly privileged account in Active Directory, unless you have a record of every modification the attacker makes or a known good backup, you can never restore the directory to a completely trustworthy state.

When a member server or a workstation is compromised and altered by an attacker, the computer is no longer trustworthy, but neighboring uncompromised servers and workstations are—compromise of one computer does not imply that all computers are compromised.

However, in an Active Directory domain, all domain controllers host replicas of the same AD DS database. If a single domain controller is compromised and an attacker modifies the AD DS database, those modifications replicate to every other domain controller in the domain, and depending on the partition in which the modifications are made, the forest. Even if you reinstall

every domain controller in the forest, you are simply reinstalling the hosts on which the AD DS database resides. Malicious modifications to Active Directory will replicate to freshly installed domain controllers as easily as they will replicate to domain controllers that have been running for years.

In assessing compromised environments, we commonly find that what was believed to be the first breach "event" was actually triggered after weeks, months, or even years after attackers had initially compromised the environment. Attackers usually obtained the credentials for highly privileged accounts long before a breach was detected, and they leveraged those accounts to compromise the directory, domain controllers, member servers, workstations, and even connected non-Windows systems.

These findings are consistent with several findings in Verizon's 2012 Data Breach Investigations Report, which states that:

- 98 percent of data breaches stemmed from external agents
- 85 percent of data breaches took weeks or more to discover
- 92 percent of incidents were discovered by a third party, and
- 97 percent of breaches were avoidable though simple or intermediate controls.

A compromise to the degree described earlier is effectively irreparable, and the standard advice to "flatten and rebuild" every compromised system is simply not feasible or even possible if Active Directory has been compromised or destroyed. Even restoring to a known good state does not eliminate the flaws that allowed the environment to be compromised in the first place.

Although you must defend every facet of your infrastructure, an attacker only needs to find enough flaws in your defenses to get to their desired goal. If your environment is relatively simple and pristine, and historically well-managed, then implementing the recommendations provided earlier in this document may be a straightforward proposition.

However, we have found that the older, larger, and more complex the environment, the more likely it is that the recommendations in this document will be infeasible or even impossible to implement. It is much harder to secure an infrastructure after the fact than it is to start fresh and to construct an environment that is resistant to attack and compromise. But as previously noted, it is no small undertaking to rebuild an entire Active Directory forest. For these reasons, we recommend a more focused, targeted approach to secure your Active Directory forests.

Rather than focusing on and trying to fix all of the things that are "broken," consider an approach in which you prioritize based on what is most important to your business and in your infrastructure. Instead of trying to remediate an environment filled with outdated, misconfigured systems and applications, consider creating a new small, secure environment into which you can safely port the users, systems, and information that are most critical to your business.

In this section, we describe an approach by which you can create a pristine AD DS forest that serves as a "life boat" or "secure cell" for your core business infrastructure. A pristine forest is simply a newly installed Active Directory forest that is typically limited in size and scope, and which is built by using current operating systems, applications, and with the principles described in Reducing the Active Directory Attack Surface.

By implementing the recommended configuration settings in a newly built forest, you can create an AD DS installation that is built from the ground up with secure settings and practices, and you can reduce the challenges that accompany supporting legacy systems and applications. While detailed instructions for the design and implementation of a pristine AD DS installation are outside the scope of this document, you should follow some general principles and guidelines to create a "secure cell" into which you can house your most critical assets. These guidelines are as follows:

1. Identify principles for segregating and securing critical assets.
2. Define a limited, risk-based migration plan.
3. Leverage "nonmigratory" migrations where necessary.
4. Implement "creative destruction."
5. Isolate legacy systems and applications.
6. Simplify security for end users.

Identifying Principles for Segregating and Securing Critical Assets

The characteristics of the pristine environment that you create to house critical assets can vary widely. For example, you may choose to create a pristine forest into which you migrate only VIP users and sensitive data that only those users can access. You may create a pristine forest in which you migrate not only VIP users, but which you implement as an administrative forest, implementing the principles described in Reducing the Active Directory Attack Surface to create secure administrative accounts and hosts that can be used to manage your legacy forests from the pristine forest. You might implement a "purpose-built" forest that houses VIP accounts, privileged accounts, and systems requiring additional security such as servers running Active Directory Certificate Services (AD CS) with the sole goal of segregating them from less-secure forests. Finally, you might implement a pristine forest that becomes the de facto location for all new users, systems, applications and data, allowing you to eventually decommission your legacy forest via attrition.

Regardless of whether your pristine forest contains a handful of users and systems or it forms the basis for a more aggressive migration, you should follow these principles in your planning:

1. Assume that your legacy forests have been compromised.
2. Do not configure a pristine environment to trust a legacy forest, although you can configure a legacy environment to trust a pristine forest.
3. Do not migrate user accounts or groups from a legacy forest to a pristine environment if there is a possibility that the accounts' group memberships, SID history, or other attributes may have been maliciously modified. Instead, use "nonmigratory" approaches to populate a pristine forest. (Nonmigratory approaches are described later in this section.)
4. Do not migrate computers from legacy forests to pristine forests. Implement freshly installed servers in the pristine forest, install applications on the freshly installed servers, and migrate application data to the newly installed systems. For file servers, copy data to freshly installed servers, set ACLs by using users and groups in the new forest, and then create print servers in a similar fashion.

5. Do not permit the installation of legacy operating systems or applications in the pristine forest. If an application cannot be updated and freshly installed, leave it in the legacy forest and consider creative destruction to replace the application's functionality.

Defining a Limited, Risk-Based Migration Plan

Creating a limited, risk-based migration plan simply means that when deciding which users, applications, and data to migrate into your pristine forest, you should identify migration targets based on the degree of risk to which your organization is exposed if one of the users or systems is compromised. VIP users whose accounts are most likely to be targeted by attackers should be housed in the pristine forest. Applications that provide vital business functions should be installed on freshly built servers in the pristine forest, and highly sensitive data should be moved to secured servers in the pristine forest.

If you do not already have a clear picture of the most business-critical users, systems, applications, and data in your Active Directory environment, work with business units to identify them. Any application required for the business to operate should be identified, as should any servers on which critical applications run or critical data is stored. By identifying the users and resources that are required for your organization to continue to function, you create a naturally prioritized collection of assets on which to focus your efforts.

Leveraging "Nonmigratory" Migrations

Whether you know that your environment has been compromised, suspect that it has been compromised, or simply prefer not to migrate legacy data and objects from a legacy Active Directory installation to a new one, consider migration approaches that do not technically "migrate" objects.

User Accounts

In a traditional Active Directory migration from one forest to another, the SIDHistory (SID history) attribute on user objects is used to store users' SID and the SIDs of groups that users were members of in the legacy forest. If users accounts are migrated to a new forest, and they access resources in the legacy forest, the SIDs in the SID history are used to create an access token that allows the users to access resources to which they had access before the accounts were migrated.

Maintaining SID history, however, has proven problematic in some environments because populating users' access tokens with current and historical SIDs can result in token bloat. Token bloat is an issue in which the number of SIDs that must be stored in a user's access token uses or exceeds the amount of space available in the token.

Although token sizes can be increased to a limited extent, the ultimate solution to token bloat is to reduce the number of SIDs associated with user accounts, whether by rationalizing group memberships, eliminating SID history, or a combination of both. For more information about token bloat, see MaxTokenSize and Kerberos Token Bloat.

Rather than migrating users from a legacy environment (particularly one in which group memberships and SID histories may be compromised) by using SID history, consider leveraging metadirectory applications to “migrate” users, without carrying SID histories into the new forest. When user accounts are created in the new forest, you can use a metadirectory application to map the accounts to their corresponding accounts in the legacy forest.

To provide the new user accounts access to resources in the legacy forest, you can use the metadirectory tooling to identify resource groups into which the users’ legacy accounts were granted access, and then add the users’ new accounts to those groups. Depending on your group strategy in the legacy forest, you may need to create domain local groups for resource access or convert existing groups to domain local groups to allow the new accounts to be added to resource groups. By focusing first on the most critical applications and data and migrating them to the new environment (with or without SID history), you can limit the amount of effort expended in the legacy environment.

Note

Some of the applications listed in Appendix J: Third-Party RBAC Vendors and Appendix K: Third-Party PIM Vendors provide mechanisms to perform “nonmigratory” migrations, as does Microsoft Forefront Identity Manager.

Servers and Workstations

In a traditional migration from one Active Directory forest to another, migrating computers is often relatively simple compared to migrating users, groups, and applications. Depending on the computer role, migrating to a new forest can be as simple as disjoining an old domain and joining a new one. However, migrating computer accounts intact into a pristine forest defeats the purpose of creating a fresh environment. Rather than migrating (potentially compromised, misconfigured, or outdated) computer accounts to a new forest, you should freshly install servers and workstations in the new environment. You can migrate data from systems in the legacy forest to systems in the pristine forest, but not the systems that house the data.

Applications

Applications can present the most significant challenge in any migration from one forest to another, but in the case of a “nonmigratory” migration, one of the most basic principles you should apply is that applications in the pristine forest should be current, supported, and freshly installed. Data can be migrated from application instances in the old forest where possible. In situations in which an application cannot be “recreated” in the pristine forest, you should consider approaches such as creative destruction or isolation of legacy applications as described in the following section.

Implementing Creative Destruction

Creative destruction is an economics term that describes economic development created by the destruction of a prior order. In recent years, the term has been applied to information technology.

It typically refers to mechanisms by which old infrastructure is eliminated, not by upgrading it, but by replacing it with something altogether new. The 2011 Gartner Symposium ITXPO for CIOs and senior IT executives presented creative destruction as one of its key themes for cost reduction and increases in efficiency. Improvements in security are possible as a natural outgrowth of the process.

For example, an organization may be composed of multiple business units that use a different application that performs similar functionality, with varying degrees of modernity and vendor support. Historically, IT might be responsible for maintaining each business unit's application separately, and consolidation efforts would consist of attempting to figure out which application offered the best functionality and then migrating data into that application from the others.

In creative destruction, rather than maintaining outdated or redundant applications, you implement entirely new applications to replace the old, migrate data into the new applications, and decommission the old applications and the systems on which they run. In some cases, you can implement creative destruction of legacy applications by deploying a new application in your own infrastructure, but wherever possible, you should consider porting the application to a cloud-based solution instead.

By deploying cloud-based applications to replace legacy in-house applications, you not only reduce maintenance efforts and costs, but you reduce your organization's attack surface by eliminating legacy systems and applications that present vulnerabilities for attackers to leverage. This approach provides a faster way for an organization to obtain desired functionality while simultaneously eliminating legacy targets in the infrastructure. Although the principle of creative destruction does not apply to all IT assets, it provides an often viable option to eliminating legacy systems and applications while simultaneously deploying robust, secure, cloud-based applications.

Isolating Legacy Systems and Applications

A natural outgrowth of migrating your business-critical users and systems to a pristine, secure environment is that your legacy forest will be contain less valuable information and systems. Although the legacy systems and applications that remain in the less secure environment may present elevated risk of compromise, they also represent a reduced severity of compromise. By rehoming and modernizing your critical business assets, you can focus on deploying effective management and monitoring while not needing to accommodate legacy settings and protocols.

When you have rehomed your critical data to a pristine forest, you can evaluate options to further isolating legacy systems and applications in your "main" AD DS forest. Although you might implement creative destruction to replace one application and the servers on which it runs, in other cases you might consider additional isolation of the least secure systems and applications. For example, an application that is used by a handful of users, but which requires legacy credentials like LAN Manager hashes can be migrated to a small domain you create to support systems for which you have no replacement options.

By removing these systems from domains where they forced implementation of legacy settings, you can subsequently increase the security of the domains by configuring them to support only current operating systems and applications. Although, it is preferable to decommission legacy

systems and applications whenever possible. If decommissioning is simply not feasible for a small segment of your legacy population, segregating it into a separate domain (or forest) allows you to perform incremental improvements in the rest of the legacy installation.

Simplifying Security for End Users

In most organizations, users who have access to the most sensitive information due to the nature of their roles in the organization often have the least amount of time to devote to learning complex access restrictions and controls. Although you should have a comprehensive security education program for all users in your organization, you should also focus on making security as simple to use as possible by implementing controls that are transparent and simplifying principles to which users adhere.

For example, you may define a policy in which executives and other VIPs are required to use secure workstations to access sensitive data and systems, allowing them to use their other devices to access less sensitive data. This is a simple principle for users to remember, but you can implement a number of backend controls to help to enforce the approach.

You can use Authentication Mechanism Assurance to permit the users to access sensitive data only if they've logged on to their secure systems using their smart cards, and can use IPsec and user rights restrictions to control the systems from which they can connect to sensitive data repositories. You can use the Microsoft Data Classification Toolkit to build a robust file classification infrastructure, and you can implement Dynamic Access Control to restrict access to data based on characteristics of an access attempt, translating business rules into technical controls.

From the perspective of the user, accessing sensitive data from a secured system "just works," and attempting to do so from an unsecured system "just doesn't." However, from the perspective of monitoring and managing your environment, you're helping to create identifiable patterns in how users access sensitive data and systems, making it easier for you to detect anomalous access attempts.

In environments in which user resistance to long, complex passwords has resulted in insufficient password policies, particularly for VIP users, consider alternate approaches to authentication, whether via smart cards (which come in a number of form factors and with additional features to strengthen authentication), biometric controls such as finger-swipe readers, or even authentication data that is secured by trusted platform module (TPM) chips in users' computers. Although multifactor authentication does not prevent credential theft attacks if a computer is already compromised, by giving your users easy-to-use authentication controls, you can assign more robust passwords to the accounts of users for whom traditional user name and password controls are unwieldy.

Maintaining a More Secure Environment

Law Number Ten: Technology is not a panacea. – 10 Immutable Laws of Security Administration

When you have created a manageable, secure environment for your critical business assets, your focus should shift to ensuring that it is maintained securely. Although you've been given specific technical controls to increase the security of your AD DS installations, technology alone will not protect an environment in which IT does not work in partnership with the business to maintain a secure, usable infrastructure. The high level recommendations in this section are meant to be used as guidelines that you can use to develop not only effective security, but effective lifecycle management.

In some cases, your IT organization might already have a close working relationship with business units, which will ease implementing these recommendations. In organizations in which IT and business units are not closely tied, you might need to first obtain executive sponsorship for efforts to forge a closer relationship between IT and business units. The Executive Summary is intended to be useful as a standalone document for executive review, and it can be disseminated to decision makers in your organization.

Creating Business-Centric Security Practices for Active Directory

In the past, information technology within many organizations was viewed as a support structure and a cost center. IT departments were often largely segregated from business users, and interactions limited to a request-response model in which the business requested resources and IT responded.

As technology has evolved and proliferated, the vision of "a computer on every desktop" has effectively come to pass for much of the world, and even been eclipsed by the broad range of easily accessible technologies available today. Information technology is no longer a support function, it is a core business function. If your organization could not continue to function if all IT services were unavailable, your organization's business is, at least in part, information technology.

To create effective compromise recovery plans, IT services must work closely with business units in your organization to identify not only the most critical components of the IT landscape, but the critical functions required by the business. By identifying what is important to your organization as a whole, you can focus on securing the components that have the most value. This is not a recommendation to shirk the security of low value systems and data. Rather, like you define levels of service for system uptime, you should consider defining levels of security control and monitoring based on criticality of asset.

When you have invested in creating a current, secure, manageable environment, you can shift focus to managing it effectively and ensuring that you have effective lifecycle management processes that aren't determined only by IT, but by the business. To achieve this, you need not only to partner with the business, but to invest the business in "ownership" of data and systems in Active Directory.

When data and systems are introduced into Active Directory without designated owners, business owners and IT owners, there is no clear chain of responsibility for the provisioning, management, monitoring, updating, and eventually decommissioning the system. This results in infrastructures

in which systems expose the organization to risk but cannot be decommissioned because ownership is unclear. To effectively manage the lifecycle of the users, data, applications, and systems managed by your Active Directory installation, you should follow the principles described in this section.

Assign a Business Owner to Active Directory Data

Data in Active Directory should have an identified business owner, that is, a specified department or user who is the point of contact for decisions about the lifecycle of the asset. In some cases, the business owner of a component of Active Directory will be an IT department or user. Infrastructure components such as domain controllers, DHCP and DNS servers, and Active Directory will most likely be "owned" by IT. For data that is added to AD DS to support the business (for example, new employees, new applications, and new information repositories), a designated business unit or user should be associated with the data.

Whether you use Active Directory to record ownership of data in the directory, or whether you implement a separate database for tracking IT assets, no user account should be created, no server or workstation should be installed, and no application should be deployed without a designated owner of record. Trying to establish ownership of systems after they've been deployed in production can be challenging at best, and impossible in some cases. Therefore, ownership should be established at the time the data is introduced into Active Directory.

Implement Business-Driven Lifecycle Management

Lifecycle management should be implemented for all data in Active Directory. For example, when a new application is introduced into an Active Directory domain, the application's business owner should, at regular intervals, be expected to attest to the continued use of the application. When a new version of an application is released, the application's business owner should be informed and should decide if and when the new version will be implemented.

If a business owner chooses not to approve deployment of a new version of an application, that business owner should also be notified of the date when the current version will no longer be supported and should be responsible for determining whether the application will be decommissioned or replaced. Keeping legacy applications running and unsupported should not be an option.

When user accounts are created in Active Directory, their managers of record should be notified at object creation and required to attest to the validity of the account at regular intervals. By implementing a business driven lifecycle and regular attestation of the validity of the data, the people who are best equipped to identify anomalies in the data are the people who review the data.

For example, attackers might create user accounts that appear to be valid accounts, following your organization's naming conventions and object placement. To detect these account creations, you might implement a daily task that returns all user objects without a designated business owner so that you can investigate the accounts. If attackers create accounts and assign a

business owner, by implementing a task that reports new object creation to the designated business owner, the business owner can quickly identify whether the account is legitimate.

You should implement similar approaches to security and distribution groups. Although some groups may be functional groups created by IT, by creating every group with a designated owner, you can retrieve all groups owned by a designated user and require the user to attest to the validity of their memberships. Similar to the approach taken with user account creation, you can trigger reporting group modifications to the designated business owner. The more routine it becomes for a business owner to attest to the validity or invalidity of data in Active Directory, the more equipped you are to identify anomalies that can indicate process failures or actual compromise.

Classify all Active Directory Data

In addition to recording a business owner for all Active Directory data at the time it is added to the directory, you should also require business owners to provide classification for the data. For example, if an application stores business-critical data, the business owner should label the application as such, in accordance with your organization's classification infrastructure.

Some organizations implement data classification policies that label data according to the damage that exposure of the data would incur if it were stolen or exposed. Other organizations implement data classification that labels data by criticality, by access requirements, and by retention. Regardless of the data classification model in use in your organization, you should ensure that you are able to apply classification to Active Directory data, not only to "file" data. If a user's account is a VIP account, it should be identified in your asset classification database (whether you implement this via the use of attributes on the objects in AD DS, or whether you deploy separate asset classification databases).

Within your data classification model, you should include classification for AD DS data such as the following.

Systems

You should not only classify data, but also their server populations. For each server, you should know what operating system is installed, what general roles the server provides, what applications are running on the server, the IT owner of record, and the business owner of record, where applicable. For all data or applications running on the server, you should require classification, and the server should be secured according to the requirements for the workloads it supports and the classifications applied to the system and data. You can also group servers by the classification of their workloads, which allows you to quickly identify the servers that should be the most closely monitored and most stringently configured.

Applications

You should classify applications by functionality (what they do), user base (who uses the applications), and the operating system on which they run. You should maintain records that

contain version information, patch status, and any other pertinent information. You should also classify applications by the types of data they handle, as previously described.

Users

Whether you call them “VIP” users, critical accounts, or use a different label, the accounts in your Active Directory installations that are most likely to be targeted by attackers should be tagged and monitored. In most organizations, it is simply not feasible to monitor all of the activities of all users. However, if you are able to identify the critical accounts in your Active Directory installation, you can monitor those accounts for changes as described earlier in this document.

You can also begin to build a database of “expected behaviors” for these accounts as you audit the accounts. For example, if you find that a given executive uses his secured workstation to access business-critical data from his office and from his home, but rarely from other locations, if you see attempts to access data by using his account from an unauthorized computer or a location halfway around the planet where you know the executive is not currently located, you can more quickly identify and investigate this anomalous behavior.

By integrating business information with your infrastructure, you can use that business information to help you identify false positives. For example, if executive travel is recorded in a calendar that is accessible to IT staff responsible for monitoring the environment, you can correlate connection attempts with the executives’ known locations.

Let’s say Executive A is normally located in Chicago and uses a secured workstation to access business-critical data from his desk, and an event is triggered by a failed attempt to access the data from an unsecured workstation located in Atlanta. If you are able to verify that the executive is currently in Atlanta, you can resolve the event by contacting the executive or the executive’s assistant to determine if the access failure was the result of the executive forgetting to use the secured workstation to access the data. By constructing a program that uses the approaches described in Planning for Compromise, you can begin to build a database of expected behaviors for the most “important” accounts in your Active Directory installation that can potentially help you more quickly discover and respond to attacks.

Summary of Best Practices

The following table contains a summary of the best practices for securing Active Directory as described in this document, with hyperlinks to the sections in which the recommendations are detailed. Practices and controls that are described as tactical in nature may be implemented more quickly and with less effort than those that are described as strategic or are applicable to discrete components in the larger infrastructure.

Some of the best practices described here are not specific to Active Directory, but are designed to help you implement solutions that can reduce the most commonly exploited vulnerabilities that are used to gain an initial foothold in an organization’s infrastructure, which may then be used to launch attacks directly against Active Directory.

Other recommendations are specific to Active Directory and may be implemented in existing AD DS installations, or implemented as fundamental principles in a new Active Directory installation, whether that installation is an enterprise deployment (housing corporate users, servers, workstations, and applications), or whether the installation is "purpose-built" (designed to house critical accounts and assets that should be separated from other AD DS forests and secured more stringently).

Another version of this table, which provides information about whether each best practice is tactical or strategic in nature, and whether its implementation provides preventative or detective controls can be found in the Executive Summary section of this document. The following table provides each recommended best practice in general order of priority, and links to more information about each.

	Best Practice	**More Information**
1	Patch applications.	"Initial Breach Targets" in Avenues to Compromise Appendix A: Patch and Vulnerability Management Software
2	Patch operating systems.	"Initial Breach Targets" in Avenues to Compromise Appendix A: Patch and Vulnerability Management Software "Principles for Creating Secure Administrative Hosts" in Implementing Secure Administrative Hosts Securing Domain Controllers Against Attack
3	Deploy and promptly update antivirus and antimalware software across all systems and monitor for attempts to remove or disable it.	Avenues to Compromise Appendix A: Patch and Vulnerability Management Software
4	Monitor sensitive Active Directory objects for modification attempts and Windows for events that may indicate attempted compromise.	Monitoring Active Directory for Signs of Compromise "Active Directory Objects and Attributes to Monitor" in Audit Policy Recommendations

		Appendix L: Events to Monitor
5	Protect and monitor accounts for users who have access to sensitive data.	"VIP Accounts" in Attractive Accounts for Credential Theft "Implementing Robust Authentication Controls" in Implementing Least-Privilege Administrative Models "Identifying Principles for Segregating and Securing Critical Assets" in Planning for Compromise "Simplify Security for End Users" in Planning for Compromise "Active Directory Objects and Attributes to Monitor" in Monitoring Active Directory for Signs of Compromise
6	Prevent powerful accounts from being used on unauthorized systems.	Implementing Least-Privilege Administrative Models Implementing Secure Administrative Hosts Securing Domain Controllers Against Attack
7	Eliminate permanent membership in highly privileged groups.	Appendix B: Privileged Accounts and Groups in Active Directory Appendix C: Protected Accounts and Groups in Active Directory Appendix D: Securing Built-In Administrator Accounts in Active Directory Appendix E: Securing Enterprise Admins Groups in Active Directory Appendix F: Securing Domain Admins Groups in Active Directory Appendix G: Securing Administrators Groups in Active Directory

		Appendix H: Securing Local Administrator Accounts and Groups Appendix J: Third-Party RBAC Vendors
8	Implement controls to grant temporary membership in privileged groups when needed.	Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory Appendix K: Third-Party PIM Vendors
9	Implement secure administrative hosts.	Implementing Secure Administrative Hosts
10	Use application whitelisting on domain controllers, administrative hosts, and other sensitive systems.	Implementing Secure Administrative Hosts Securing Domain Controllers Against Attack
11	Identify critical assets, and prioritize their security and monitoring.	Planning for Compromise
12	Implement least-privilege, role-based access controls to administer the directory, its supporting infrastructure, and domain-joined systems.	“Role-Based Access Controls (RBAC) for Active Directory” in Implementing Least-Privilege Administrative Models Appendix J: Third-Party RBAC Vendors
13	Isolate legacy systems and applications.	“Isolating Legacy Systems and Applications” in Planning for Compromise
14	Decommission legacy systems and applications.	“Implementing Creative Destruction” in Planning for Compromise
15	Implement secure development lifecycle programs for custom applications.	“Lack of Secure Application Development Practices” in Avenues to Compromise
16	Implement configuration management, review compliance regularly, and evaluate settings	“Maintaining a More Secure Environment” in Planning for Compromise

	with each new hardware or software version.	
17	Migrate critical assets to pristine forests with stringent security and monitoring requirements.	Planning for Compromise
18	Simplify security for end users.	“Simplify Security for End Users” in Planning for Compromise
19	Use host-based firewalls to control and secure communications.	“Principles for Creating Secure Administrative Hosts” in Implementing Secure Administrative Hosts “Secure Configuration of Domain Controllers” in Securing Domain Controllers Against Attack
20	Patch devices.	Contact your device vendors Appendix A: Patch and Vulnerability Management Software
21	Implement business-centric lifecycle management for IT assets.	“Creating Business-Centric Security Practices for Active Directory” in Planning for Compromise
22	Create or update incident recovery plans.	Planning for Compromise

Appendices

Appendices are included in this document to augment the information contained in the body of the document. The list of appendices and a brief description of each in included the following table.

Appendix	Description
Appendix A: Patch and Vulnerability Management Software	Contains a list of companies that produce patch and vulnerability management software.
Appendix B: Privileged Accounts and Groups in Active Directory	Provides background information that helps you to identify the users and groups you should

	focus on securing because they can be leveraged by attackers to compromise and even destroy your Active Directory installation.
Appendix C: Protected Accounts and Groups in Active Directory	Contains information about protected groups in Active Directory. It also contains information for limited customization (removal) of groups that are considered protected groups and are affected by AdminSDHolder and SDProp.
Appendix D: Securing Built-In Administrator Accounts in Active Directory	Contains guidelines to help secure the Administrator account in each domain in the forest.
Appendix E: Securing Enterprise Admins Groups in Active Directory	Contains guidelines to help secure the Enterprise Admins group in the forest.
Appendix F: Securing Domain Admins Groups in Active Directory	Contains guidelines to help secure the Domain Admins group in each domain in the forest.
Appendix G: Securing Administrators Groups in Active Directory	Contains guidelines to help secure the Built-in Administrators group in each domain in the forest.
Appendix H: Securing Local Administrator Accounts and Groups	Contains guidelines to help secure local Administrator accounts and Administrators groups on domain-joined servers and workstations.
Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory	Provides information to create accounts that have limited privileges and can be stringently controlled, but can be used to populate privileged groups in Active Directory when temporary elevation is required.
Appendix J: Third-Party RBAC Vendors	Contains a list of third-party RBAC vendors and the RBAC solutions they offer.
Appendix K: Third-Party PIM Vendors	Contains a list of third-party PIM vendors and the PIM solutions they offer.
Appendix L: Events to Monitor	Lists events for which you should monitor in your environment.
Appendix M: Document Links and Recommended Reading	Contains a list of recommended reading. Also contains a list of links to external documents and their URLs so that readers of hard copies of this document can access this information.

Appendix A: Patch and Vulnerability Management Software

Appendix A: Patch and Vulnerability Management Software

The following companies produce patch and vulnerability management software that you can use to help keep third-party applications and non-Windows operating systems up to date if you do not already have a comprehensive patch management plan. This list is not intended to serve as a recommendation, only as reference material that you can use to devise a patch management methodology for your environment.

Vendor	Offerings
Absolute Software	Absolute Manage, Absolute Manage MDM
SolarWinds	EminentWare WSUS Extension Pack, EminentWare System Center Configuration Manager Extension Pack
GFI Software	GFI LanGuard
Secunia	Secunia Corporate Software Inspector (CSI), Vulnerability Intelligence Manager
eEye Digital Security	Retina CS Management
Lumension	Lumension Vulnerability Management

Appendix B: Privileged Accounts and Groups in Active Directory

Appendix B: Privileged Accounts and Groups in Active Directory

“Privileged” accounts and groups in Active Directory are those to which powerful rights, privileges, and permissions are granted that allow them to perform nearly any action in Active Directory and on domain-joined systems. This appendix begins by discussing rights, privileges, and permissions, followed by information about the “highest privilege” accounts and groups in Active Directory,that is, the most powerful accounts and groups.

Information is also provided about built-in and default accounts and groups in Active Directory, in addition to their rights. Although specific configuration recommendations for securing the highest privilege accounts and groups are provided as separate appendices, this appendix provides background information that helps you identify the users and groups you should focus on securing. You should do so because they can be leveraged by attackers to compromise and even destroy your Active Directory installation.

Rights, Privileges, and Permissions in Active Directory

The differences between rights, permissions, and privileges can be confusing and contradictory, even within documentation from Microsoft. This section describes some of the characteristics of each as they are used in this document. These descriptions should not be considered authoritative for other Microsoft documentation, because it may use these terms differently.

Rights and Privileges

Rights and privileges are effectively the same thing—system-wide capabilities that are granted to security principals such as users, services, computers, or groups. In interfaces typically used by IT professionals, these are usually referred to as “rights” or “user rights,” and they are often assigned by Group Policy Objects. The following screenshot shows some of the most common user rights that can be assigned to security principals (it represents the Default Domain Controllers GPO in a Windows Server 2012 domain). Some of these rights apply to Active Directory, such as the **Enable computer and user accounts to be trusted for delegation** user right, while other rights apply to the Windows operating system, such as **Change the system time**.

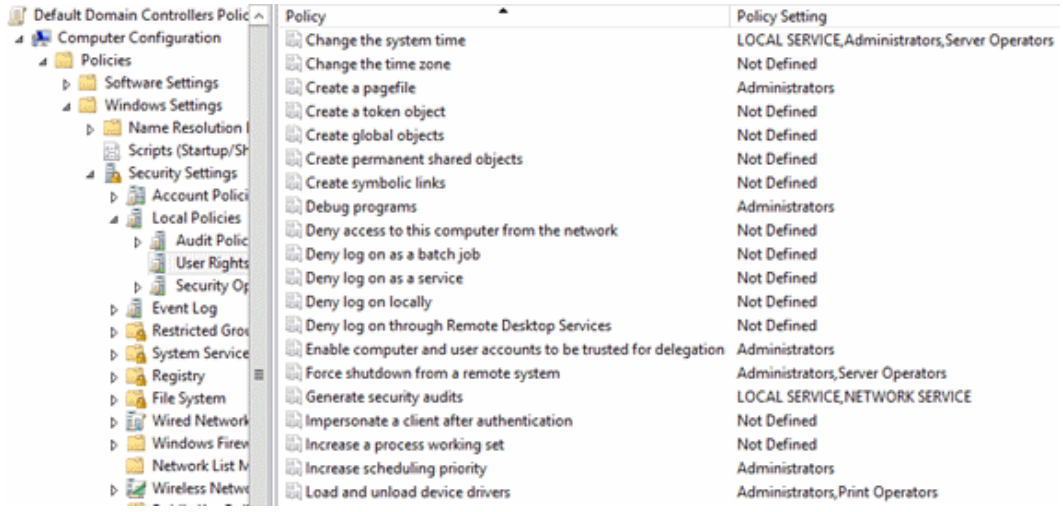

In interfaces such as the Group Policy Object Editor, all of these assignable capabilities are referred to broadly as user rights. In reality however, some user rights are programmatically referred to as rights, while others are programmatically referred to as privileges. Table B-1: User Rights and Privileges provides some of the most common assignable user rights and their

programmatic constants. Although Group Policy and other interfaces refer to all of these as user rights, some are programmatically identified as rights, while others are defined as privileges.

For more information about each of the user rights listed in the following table, use the links in the table or see Threats and Countermeasures Guide: User Rights in the Threats and Vulnerabilities Mitigation guide for Windows Server 2008 R2 on the Microsoft TechNet site. For information applicable to Windows Server 2008, please see User Rights in the Threats and Vulnerabilities Mitigation documentation on the Microsoft TechNet site. As of the writing of this document, corresponding documentation for Windows Server 2012 is not yet published.

Note
For the purposes of this document, the terms “rights” and “user rights” are used to identify rights and privileges unless otherwise specified.

Table B-1: User Rights and Privileges

User Right in Group Policy	Name of Constant
Access Credential Manager as a trusted caller	SeTrustedCredManAccessPrivilege
Access this computer from the network	SeNetworkLogonRight
Act as part of the operating system	SeTcbPrivilege
Add workstations to domain	SeMachineAccountPrivilege
Adjust memory quotas for a process	SeIncreaseQuotaPrivilege
Allow log on locally	SeInteractiveLogonRight
Allow log on through Terminal Services	SeRemoteInteractiveLogonRight
Back up files and directories	SeBackupPrivilege
Bypass traverse checking	SeChangeNotifyPrivilege
Change the system time	SeSystemtimePrivilege
Change the time zone	SeTimeZonePrivilege
Create a pagefile	SeCreatePagefilePrivilege
Create a token object	SeCreateTokenPrivilege
Create global objects	SeCreateGlobalPrivilege
Create permanent shared objects	SeCreatePermanentPrivilege
Create symbolic links	SeCreateSymbolicLinkPrivilege
Debug programs	SeDebugPrivilege
Deny access to this computer from the network	SeDenyNetworkLogonRight

Deny log on as a batch job	SeDenyBatchLogonRight
Deny log on as a service	SeDenyServiceLogonRight
Deny log on locally	SeDenyInteractiveLogonRight
Deny log on through Terminal Services	SeDenyRemoteInteractiveLogonRight
Enable computer and user accounts to be trusted for delegation	SeEnableDelegationPrivilege
Force shutdown from a remote system	SeRemoteShutdownPrivilege
Generate security audits	SeAuditPrivilege
Impersonate a client after authentication	SeImpersonatePrivilege
Increase a process working set	SeIncreaseWorkingSetPrivilege
Increase scheduling priority	SeIncreaseBasePriorityPrivilege
Load and unload device drivers	SeLoadDriverPrivilege
Lock pages in memory	SeLockMemoryPrivilege
Log on as a batch job	SeBatchLogonRight
Log on as a service	SeServiceLogonRight
Manage auditing and security log	SeSecurityPrivilege
Modify an object label	SeRelabelPrivilege
Modify firmware environment values	SeSystemEnvironmentPrivilege
Perform volume maintenance tasks	SeManageVolumePrivilege
Profile single process	SeProfileSingleProcessPrivilege
Profile system performance	SeSystemProfilePrivilege
Remove computer from docking station	SeUndockPrivilege
Replace a process level token	SeAssignPrimaryTokenPrivilege
Restore files and directories	SeRestorePrivilege
Shut down the system	SeShutdownPrivilege
Synchronize directory service data	SeSyncAgentPrivilege
Take ownership of files or other objects	SeTakeOwnershipPrivilege

Permissions

Permissions are access controls that are applied to securable objects such as the file system, registry, service, and Active Directory objects. Each securable object has an associated access control list (ACL), which contains access control entries (ACEs) that grant or deny security principals (users, services, computers, or groups) the ability to perform various operations on the object. For example, the ACLs for many objects in Active Directory contain ACEs that allow Authenticated Users to read general information about the objects, but do not grant them the ability to read sensitive information or to change the objects.

With the exception of each domain's built-in Guest account, every security principal that logs on and is authenticated by a domain controller in an Active Directory forest or a trusted forest has the Authenticated Users Security Identifier (SID) added to its access token by default. Therefore, whether a user, service, or computer account attempts to read general properties on user objects in a domain, the read operation is successful.

If a security principal attempts to access an object for which no ACEs are defined and that contain a SID that is present in the principal's access token, the principal cannot access the object. Moreover, if an ACE in an object's ACL contains a deny entry for a SID that matches the user's access token, the "deny" ACE will generally override a conflicting "allow" ACE. For more information about access control in Windows, see Access Control on the MSDN website.

Within this document, permissions refers to capabilities that are granted or denied to security principals on securable objects. Whenever there is a conflict between a user right and a permission, the user right generally takes precedence. For example, if an object in Active Directory has been configured with an ACL that denies Administrators all read and write access to an object, a user who is a member of the domain's Administrators group will be unable to view much information about the object. However, because the Administrators group is granted the user right "Take ownership of files or other objects," the user can simply take ownership of the object in question, then rewrite the object's ACL to grant Administrators full control of the object.

It is for this reason that this document encourages you to avoid using powerful accounts and groups for day-to-day administration, rather than trying to restrict the capabilities of the accounts and groups. It is not effectively possible to stop a determined user who has access to powerful credentials from using those credentials to gain access to any securable resource.

Built-in Privileged Accounts and Groups

Active Directory is intended to facilitate delegation of administration and the principle of least privilege in assigning rights and permissions. "Regular" users who have accounts in an Active Directory domain are, by default, able to read much of what is stored in the directory, but are able to change only a very limited set of data in the directory. Users who require additional privilege can be granted membership in various privileged groups that are built into the directory so that they may perform specific tasks related to their roles, but cannot perform tasks that are not relevant to their duties.

Within Active Directory, there are three built-in groups that comprise the highest privilege groups in the directory: the Enterprise Admins (EA) group, the Domain Admins (DA) group, and the built-in Administrators (BA) group.

A fourth group, the Schema Admins (SA) group, has privileges that, if abused, can damage or destroy an entire Active Directory forest, but this group is more restricted in its capabilities than the EA, DA, and BA groups.

In addition to these four groups, there are a number of additional built-in and default accounts and groups in Active Directory, each of which is granted rights and permissions that allow specific administrative tasks to be performed. Although this appendix does not provide a thorough discussion of every built-in or default group in Active Directory, it does provide a table of the groups and accounts that you're most likely to see in your installations.

For example, if you install Microsoft Exchange Server into an Active Directory forest, additional accounts and groups may be created in the Built-in and Users containers in your domains. This appendix describes only the groups and accounts that are created in the Built-in and Users containers in Active Directory, based on native roles and features. Accounts and groups that are created by the installation of enterprise software are not included.

Enterprise Admins

The Enterprise Admins (EA) group is located in the forest root domain, and by default, it is a member of the built-in Administrators group in every domain in the forest. The Built-in Administrator account in the forest root domain is the only default member of the EA group. EAs are granted rights and permissions that allow them to affect forest-wide changes. These are changes that affect all domains in the forest, such as adding or removing domains, establishing forest trusts, or raising forest functional levels. In a properly designed and implemented delegation model, EA membership is required only when first constructing the forest or when making certain forest-wide changes such as establishing an outbound forest trust.

The EA group is located by default in the Users container in the forest root domain, and it is a universal security group, unless the forest root domain is running in Windows 2000 Server mixed mode, in which case the group is a global security group. Although some rights are granted directly to the EA group, many of this group's rights are actually inherited by the EA group because it is a member of the Administrators group in each domain in the forest. Enterprise Admins have no default rights on workstations or member servers.

Domain Admins

Each domain in a forest has its own Domain Admins (DA) group, which is a member of that domain's built-in Administrators (BA) group in addition to a member of the local Administrators group on every computer that is joined to the domain. The only default member of the DA group for a domain is the Built-in Administrator account for that domain.

DAs are all-powerful within their domains, while EAs have forest-wide privilege. In a properly designed and implemented delegation model, DA membership should be required only in "break glass" scenarios, which are situations in which an account with high levels of privilege on every

computer in the domain is needed, or when certain domain wide changes must be made. Although native Active Directory delegation mechanisms do allow delegation to the extent that it is possible to use DA accounts only in emergency scenarios, constructing an effective delegation model can be time consuming, and many organizations use third-party applications to expedite the process.

The DA group is a global security group located in the Users container for the domain. There is one DA group for each domain in the forest, and the only default member of a DA group is the domain's Built-in Administrator account. Because a domain's DA group is nested in the domain's BA group and every domain-joined system's local Administrators group, DAs not only have permissions that are specifically granted to Domain Admins, but they also inherit all rights and permissions granted to the domain's Administrators group and the local Administrators group on all systems joined to the domain.

Administrators

The built-in Administrators (BA) group is a domain local group in a domain's Built-in container into which DAs and EAs are nested, and it is this group that is granted many of the direct rights and permissions in the directory and on domain controllers. However, the Administrators group for a domain does not have any privileges on member servers or on workstations. Membership in domain-joined computers' local Administrators group is where local privilege is granted; and of the groups discussed, only DAs are members of all domain-joined computers' local Administrators groups by default.

The Administrators group is a domain-local group in the domain's Built-in container. By default, every domain's BA group contains the local domain's Built-in Administrator account, the local domain's DA group, and the forest root domain's EA group. Many user rights in Active Directory and on domain controllers are granted specifically to the Administrators group, not to EAs or DAs. A domain's BA group is granted full control permissions on most directory objects, and can take ownership of directory objects. Although EA and DA groups are granted certain object-specific permissions in the forest and domains, much of the power of groups is actually "inherited" from their membership in BA groups.

Note

Although these are the default configurations of these privileged groups, a member of any one of the three groups can manipulate the directory to gain membership in any of the other groups. In some cases, it is trivial to achieve, while in others it is more difficult, but from the perspective of potential privilege, all three groups should be considered effectively equivalent.

Schema Admins

The Schema Admins (SA) group is a universal group in the forest root domain and has only that domain's Built-in Administrator account as a default member, similar to the EA group. Although membership in the SA group can allow an attacker to compromise the Active Directory schema,

which is the framework for the entire Active Directory forest, SAs have few default rights and permissions beyond the schema.

You should carefully manage and monitor membership in the SA group, but in some respects, this group is “less privileged” than the three highest privileged groups described earlier because the scope of its privilege is very narrow; that is, SAs have no administrative rights anywhere other than the schema.

Additional Built-in and Default Groups in Active Directory

To facilitate delegating administration in the directory, Active Directory ships with various built-in and default groups that have been granted specific rights and permissions. These groups are described briefly in the following table.

The following table lists the built-in and default groups in Active Directory. Both sets of groups exist by default; however, built-in groups are located (by default) in the Built-in container in Active Directory, while default groups are located (by default) in the Users container in Active Directory. Groups in the Built-in container are all Domain Local groups, while groups in the Users container are a mixture of Domain Local, Global, and Universal groups, in addition to three individual user accounts (Administrator, Guest, and Krbtgt).

In addition to the highest privileged groups described earlier in this appendix, some built-in and default accounts and groups are granted elevated privileges and should also be protected and used only on secure administrative hosts. These groups and accounts can be found in the shaded rows in Table B-1: Built-in and Default Groups and Accounts in Active Directory. Because some of these groups and accounts are granted rights and permissions that can be misused to compromise Active Directory or domain controllers, they are afforded additional protections as described in Appendix C: Protected Accounts and Groups in Active Directory.

Table B-1: Built-in and Default Accounts and Groups in Active Directory

Account or Group	Default Container, Group Scope and Type	Description and Default User Rights
Access Control Assistance Operators (Active Directory in Windows Server 2012)	Built-in container Domain-local security group	Members of this group can remotely query authorization attributes and permissions for resources on this computer. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking

		Increase a process working set
Account Operators	Built-in container Domain-local security group	Members can administer domain user and group accounts. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Administrator account	Users container Not a group	Built-in account for administering the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Adjust memory quotas for a process Allow log on locally Allow log on through Remote Desktop Services Back up files and directories Bypass traverse checking Change the system time Change the time zone Create a pagefile Create global objects Create symbolic links Debug programs Enable computer and user accounts to be trusted for delegation Force shutdown from a remote system Impersonate a client after authentication Increase a process working set

		Increase scheduling priority Load and unload device drivers Log on as a batch job Manage auditing and security log Modify firmware environment values Perform volume maintenance tasks Profile single process Profile system performance Remove computer from docking station Restore files and directories Shut down the system Take ownership of files or other objects
Administrators group	Built-in container Domain-local security group	Administrators have complete and unrestricted access to the domain. **Direct user rights:** Access this computer from the network Adjust memory quotas for a process Allow log on locally Allow log on through Remote Desktop Services Back up files and directories Bypass traverse checking Change the system time Change the time zone Create a pagefile Create global objects Create symbolic links Debug programs Enable computer and user accounts to be trusted for delegation Force shutdown from a remote system Impersonate a client after authentication

		Increase scheduling priority Load and unload device drivers Log on as a batch job Manage auditing and security log Modify firmware environment values Perform volume maintenance tasks Profile single process Profile system performance Remove computer from docking station Restore files and directories Shut down the system Take ownership of files or other objects Inherited user rights: Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Allowed RODC Password Replication Group	Users container Domain-local security group	Members in this group can have their passwords replicated to all read-only domain controllers in the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Backup Operators	Built-in container Domain-local security group	Backup Operators can override security restrictions for the sole purpose of backing up or restoring files. **Direct user rights:** Allow log on locally Back up files and directories

		Log on as a batch job Restore files and directories Shut down the system **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Cert Publishers	Users container Domain-local security group	Members of this group are permitted to publish certificates to the directory. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Certificate Service DCOM Access	Built-in container Domain-local security group	If Certificate Services is installed on a domain controller (not recommended), this group grants DCOM enrollment access to Domain Users and Domain Computers. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Cloneable Domain Controllers (AD DS in Windows Server 2012AD DS)	Users container Global security group	Members of this group that are domain controllers may be cloned. **Direct user rights:** None **Inherited user rights:** Access this computer from the network

		Add workstations to domain Bypass traverse checking Increase a process working set
Cryptographic Operators	Built-in container Domain-local security group	Members are authorized to perform cryptographic operations. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Debugger Users	This is neither a default nor a built-in group, but when present in AD DS, is cause for further investigation.	The presence of a Debugger Users group indicates that debugging tools have been installed on the system at some point, whether via Visual Studio, SQL, Office, or other applications that require and support a debugging environment. This group allows remote debugging access to computers. When this group exists at the domain level, it indicates that a debugger or an application that contains a debugger has been installed on a domain controller.
Denied RODC Password Replication Group	Users container Domain-local security group	Members in this group cannot have their passwords replicated to any read-only domain controllers in the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
DHCP Administrators	Users container	Members of this group have

	Domain-local security group	administrative access to the DHCP Server service. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
DHCP Users	Users container Domain-local security group	Members of this group have view-only access to the DHCP Server service. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Distributed COM Users	Built-in container Domain-local security group	Members of this group are allowed to launch, activate, and use distributed COM objects on this computer. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
DnsAdmins	Users container Domain-local security group	Members of this group have administrative access to the DNS Server service. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain

		Bypass traverse checking Increase a process working set
DnsUpdateProxy	Users container Global security group	Members of this group are DNS clients who are permitted to perform dynamic updates on behalf of clients that cannot themselves perform dynamic updates. Members of this group are typically DHCP servers. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Domain Admins	Users container Global security group	Designated administrators of the domain; Domain Admins is a member of every domain-joined computer's local Administrators group and receives rights and permissions granted to the local Administrators group, in addition to the domain's Administrators group. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Adjust memory quotas for a process Allow log on locally Allow log on through Remote Desktop Services Back up files and directories Bypass traverse checking Change the system time Change the time zone Create a pagefile Create global objects

		Create symbolic links Debug programs Enable computer and user accounts to be trusted for delegation Force shutdown from a remote system Impersonate a client after authentication Increase a process working set Increase scheduling priority Load and unload device drivers Log on as a batch job Manage auditing and security log Modify firmware environment values Perform volume maintenance tasks Profile single process Profile system performance Remove computer from docking station Restore files and directories Shut down the system Take ownership of files or other objects
Domain Computers	Users container Global security group	All workstations and servers that are joined to the domain are by default members of this group. **Default direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Domain Controllers	Users container Global security group	All domain controllers in the domain. Note: Domain controllers are not a member of the Domain Computers group.

		Direct user rights: None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Domain Guests	Users container Global security group	All guests in the domain **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Domain Users	Users container Global security group	All users in the domain **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Enterprise Admins (exists only in forest root domain)	Users container Universal security group	Enterprise Admins have permissions to change forest-wide configuration settings; Enterprise Admins is a member of every domain's Administrators group and receives rights and permissions granted to that group. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Adjust memory quotas for a process

		Allow log on locally Allow log on through Remote Desktop Services Back up files and directories Bypass traverse checking Change the system time Change the time zone Create a pagefile Create global objects Create symbolic links Debug programs Enable computer and user accounts to be trusted for delegation Force shutdown from a remote system Impersonate a client after authentication Increase a process working set Increase scheduling priority Load and unload device drivers Log on as a batch job Manage auditing and security log Modify firmware environment values Perform volume maintenance tasks Profile single process Profile system performance Remove computer from docking station Restore files and directories Shut down the system Take ownership of files or other objects
Enterprise Read-only Domain Controllers	Users container Universal security group	This group contains the accounts for all read-only domain controllers in the forest. **Direct user rights:** None **Inherited user rights:**

		Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Event Log Readers	Built-in container Domain-local security group	Members of this group in can read the event logs on domain controllers. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Group Policy Creator Owners	Users container Global security group	Members of this group can create and modify Group Policy Objects in the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Guest	Users container Not a group	This is the only account in an AD DS domain that does not have the Authenticated Users SID added to its access token. Therefore, any resources that are configured to grant access to the Authenticated Users group will not be accessible to this account. This behavior is not true of members of the Domain Guests and Guests groups, however- members of those groups do have the Authenticated Users SID added to their access tokens. **Direct user rights:** None

		Inherited user rights: Access this computer from the network Bypass traverse checking Increase a process working set
Guests	Built-in container Domain-local security group	Guests have the same access as members of the Users group by default, except for the Guest account, which is further restricted as described earlier. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Hyper-V Administrators (Windows Server 2012)	Built-in container Domain-local security group	Members of this group have complete and unrestricted access to all features of Hyper-V. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
IIS_IUSRS	Built-in container Domain-local security group	Built-in group used by Internet Information Services. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set

Incoming Forest Trust Builders (exists only in forest root domain)	Built-in container Domain-local security group	Members of this group can create incoming, one-way trusts to this forest. (Creation of outbound forest trusts is reserved for Enterprise Admins.) **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Krbtgt	Users container Not a group	The Krbtgt account is the service account for the Kerberos Key Distribution Center in the domain. This account has access to all accounts' credentials stored in Active Directory. This account is disabled by default and should never be enabled **User rights:** N/A
Network Configuration Operators	Built-in container Domain-local security group	Members of this group are granted privileges that allow them to manage configuration of networking features. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Performance Log Users	Built-in container Domain-local security group	Members of this group can schedule logging of performance counters, enable trace providers, and collect event traces locally and via remote access to the computer. **Direct user rights:** Log on as a batch job

		Inherited user rights: Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Performance Monitor Users	Built-in container Domain-local security group	Members of this group can access performance counter data locally and remotely. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Pre-Windows 2000 Compatible Access	Built-in container Domain-local security group	This group exists for backward compatibility with operating systems prior to Windows 2000 Server, and it provides the ability for members to read user and group information in the domain. **Direct user rights:** Access this computer from the network Bypass traverse checking **Inherited user rights:** Add workstations to domain Increase a process working set
Print Operators	Built-in container Domain-local security group	Members of this group can administer domain printers. **Direct user rights:** Allow log on locally Load and unload device drivers Shut down the system **Inherited user rights:** Access this computer from the

		network Add workstations to domain Bypass traverse checking Increase a process working set
RAS and IAS Servers	Users container Domain-local security group	Servers in this group can read remote access properties on user accounts in the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
RDS Endpoint Servers (Windows Server 2012)	Built-in container Domain-local security group	Servers in this group run virtual machines and host sessions where users RemoteApp programs and personal virtual desktops run. This group needs to be populated on servers running RD Connection Broker. RD Session Host servers and RD Virtualization Host servers used in the deployment need to be in this group. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
RDS Management Servers (Windows Server 2012)	Built-in container Domain-local security group	Servers in this group can perform routine administrative actions on servers running Remote Desktop Services. This group needs to be populated on all servers in a Remote Desktop Services deployment. The servers running the RDS Central

		Management service must be included in this group. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
RDS Remote Access Servers (Windows Server 2012)	Built-in container Domain-local security group	Servers in this group enable users of RemoteApp programs and personal virtual desktops access to these resources. In Internet-facing deployments, these servers are typically deployed in an edge network. This group needs to be populated on servers running RD Connection Broker. RD Gateway servers and RD Web Access servers used in the deployment need to be in this group. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Read-only Domain Controllers	Users container Global security group	This group contains all read-only domain controllers in the domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Remote Desktop Services Users	Built-in container	Members of this group are granted

	Domain-local security group	the right to log on remotely using RDP. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Remote Management Servers (Windows Server 2012)	Built-in container Domain-local security group	Members of this group can access WMI resources over management protocols (such as WS-Management via the Windows Remote Management service). This applies only to WMI namespaces that grant access to the user. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Replicator	Built-in container Domain-local security group	Supports legacy file replication in a domain. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Schema Admins (exists only in forest root domain)	Users container Universal security group	Schema admins are the only users who can make modifications to the Active Directory schema, and only if the schema is write-enabled. **Direct user rights:** None

		Inherited user rights: Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Server Operators	Built-in container Domain-local security group	Members of this group can administer domain servers. **Direct user rights:** Allow log on locally Back up files and directories Change the system time Change the time zone Force shutdown from a remote system Restore files and directories Shut down the system **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
Terminal Server License Servers	Built-in container Domain-local security group	Members of this group can update user accounts in Active Directory with information about license issuance, for the purpose of tracking and reporting TS Per User CAL usage **Default direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set

Users	Built-in container Domain-local security group	Users have permissions that allow them to read many objects and attributes in Active Directory, although they cannot change most. Users are prevented from making accidental or intentional system-wide changes and can run most applications. **Direct user rights:** Increase a process working set **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking
Windows Authorization Access Group	Built-in container Domain-local security group	Members of this group have access to the computed tokenGroupsGlobalAndUniversal attribute on User objects **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain Bypass traverse checking Increase a process working set
WinRMRemoteWMIUsers_ (Windows Server 2012)	Users container Domain-local security group	Members of this group can access WMI resources over management protocols (such as WS-Management via the Windows Remote Management service). This applies only to WMI namespaces that grant access to the user. **Direct user rights:** None **Inherited user rights:** Access this computer from the network Add workstations to domain

		Bypass traverse checking Increase a process working set

Appendix C: Protected Accounts and Groups in Active Directory

Appendix C: Protected Accounts and Groups in Active Directory

Within Active Directory, a default set of highly privileged accounts and groups are considered protected accounts and groups. With most objects in Active Directory, delegated administrators (users who have been delegated permissions to manage Active Directory objects) can change permissions on the objects, including changing permissions to allow themselves to change memberships of the groups, for example.

However, with protected accounts and groups, the objects' permissions are set and enforced via an automatic process that ensures the permissions on the objects remains consistent even if the objects are moved the directory. Even if somebody manually changes a protected object's permissions, this process ensures that permissions are returned to their defaults quickly.

Protected Groups

The following table contains the protected groups in Active Directory listed by domain controller operating system.

Protected Accounts and Groups in Active Directory by Operating System

Windows 2000 <SP4	**Windows 2000 SP4 – Windows Server 2003 RTM**	**Windows Server 2003 SP1+**	**Windows Server 2012, Windows Server 2008 R2, Windows Server 2008**
Administrators	Account Operators	Account Operators	Account Operators
	Administrator	Administrator	Administrator
	Administrators	Administrators	Administrators
	Backup Operators	Backup Operators	Backup Operators
	Cert Publishers		
Domain Admins	Domain Admins	Domain Admins	Domain Admins

	Domain Controllers	Domain Controllers	Domain Controllers
Enterprise Admins	Enterprise Admins	Enterprise Admins	Enterprise Admins
	Krbtgt	Krbtgt	Krbtgt
	Print Operators	Print Operators	Print Operators
			Read-only Domain Controllers
	Replicator	Replicator	Replicator
Schema Admins	Schema Admins	Schema Admins	Schema Admins
	Server Operators	Server Operators	Server Operators

AdminSDHolder

The purpose of the AdminSDHolder object is to provide “template” permissions for the protected accounts and groups in the domain. AdminSDHolder is automatically created as an object in the System container of every Active Directory domain. Its path is:
CN=AdminSDHolder,CN=System,DC=<domain_component>,DC=<domain_component>….

Unlike most objects in the Active Directory domain, which are owned by the Administrators group, AdminSDHolder is owned by the Domain Admins group. By default, EAs can make changes to any domain’s AdminSDHolder object, as can the domain’s Domain Admins and Administrators groups. Additionally, although the default owner of AdminSDHolder is the domain’s Domain Admins group, members of Administrators or Enterprise Admins can take ownership of the object.

SDProp

SDProp is a process that runs every 60 minutes (by default) on the domain controller that holds the domain’s PDC Emulator (PDCE). SDProp compares the permissions on the domain’s AdminSDHolder object with the permissions on the protected accounts and groups in the domain. If the permissions on any of the protected accounts and groups do not match the permissions on the AdminSDHolder object, the permissions on the protected accounts and groups are reset to match those of the domain’s AdminSDHolder object.

Additionally, permissions inheritance is disabled on protected groups and accounts, which means that even if the accounts and groups are moved to different locations in the directory, they do not inherit permissions from their new parent objects. Inheritance is disabled on the AdminSDHolder object so that permission changes to the parent objects do not change the permissions of AdminSDHolder.

Changing SDProp Interval

Normally, you should not need to change the interval at which SDProp runs, except for testing purposes. If you need to change the SDProp interval, on the PDCE for the domain, use regedit to

add or modify the AdminSDProtectFrequency DWORD value in HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters.

The range of values is in seconds from 60 to 7200 (one minute to two hours). To reverse the changes, delete AdminSDProtectFrequency key, which will cause SDProp to revert back to the 60 minute interval. You generally should not reduce this interval in production domains as it can increase LSASS processing overhead on the domain controller. The impact of this increase is dependent on the number of protected objects in the domain.

Running SDProp Manually

A better approach to testing AdminSDHolder changes is to run SDProp manually, which causes the task to run immediately but does not affect scheduled execution. Running SDProp manually is performed slightly differently on domain controllers running Windows Server 2008 and earlier than it is on domain controllers running Windows Server 2012 or Windows Server 2008 R2.

Procedures for running SDProp manually on older operating systems are provided in Microsoft Support article 251343, and following are step-by-step instructions for older and newer operating systems. In either case, you must connect to the rootDSE object in Active Directory and perform a modify operation with a null DN for the rootDSE object, specifying the name of the operation as the attribute to modify. For more information about modifiable operations on the rootDSE object, see rootDSE Modify Operations on the MSDN website.

Running SDProp Manually in Windows Server 2008 or Earlier

You can force SDProp to run by using Ldp.exe or by running an LDAP modification script. To run SDProp using Ldp.exe, perform the following steps after you have made changes to the AdminSDHolder object in a domain:

1. Launch **Ldp.exe**.
2. Click **Connection** on the Ldp dialog box, and click **Connect**.

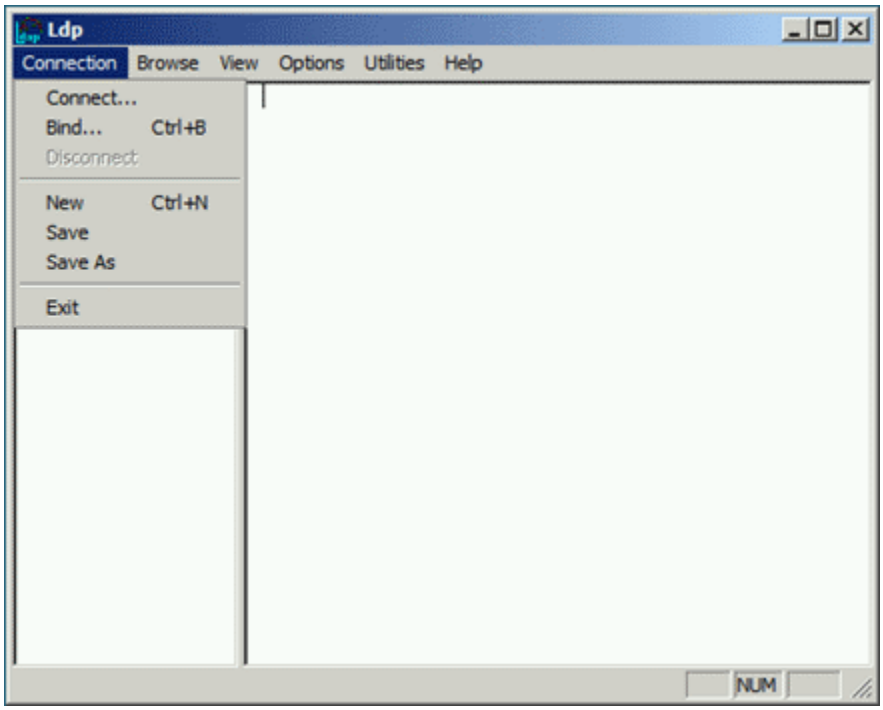

3. In the **Connect** dialog box, type the name of the domain controller for the domain that holds the PDC Emulator (PDCE) role and click **OK**.

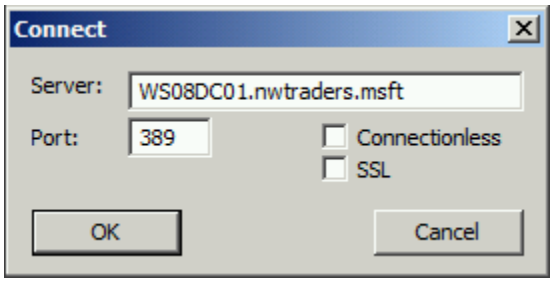

4. Verify that you have connected successfully, as indicated by **Dn: (RootDSE)** in the following screenshot, click **Connection** and click **Bind**.

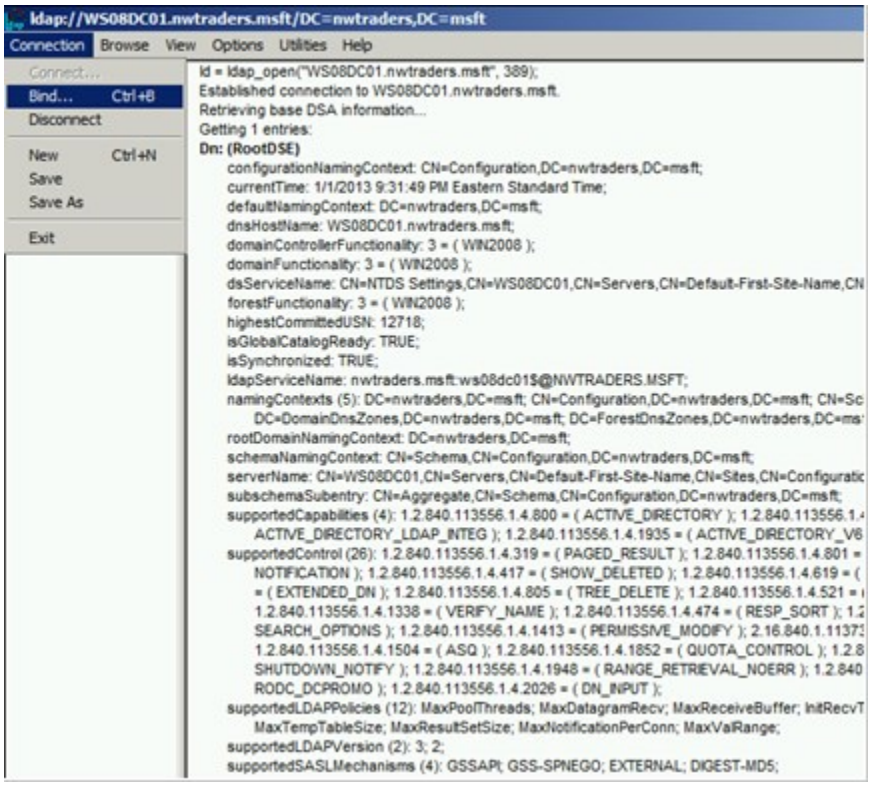

5. In the **Bind** dialog box, type the credentials of a user account that has permission to modify the rootDSE object. (If you are logged on as that user, you can select **Bind as** currently logged on user.) Click **OK**.

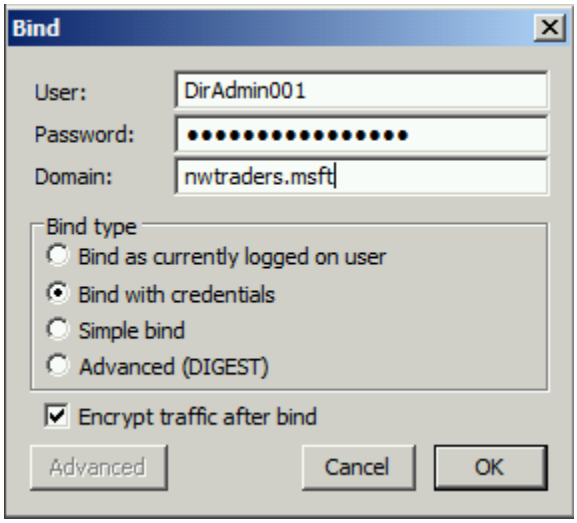

6. After you have completed the bind operation, click **Browse**, and click **Modify**.

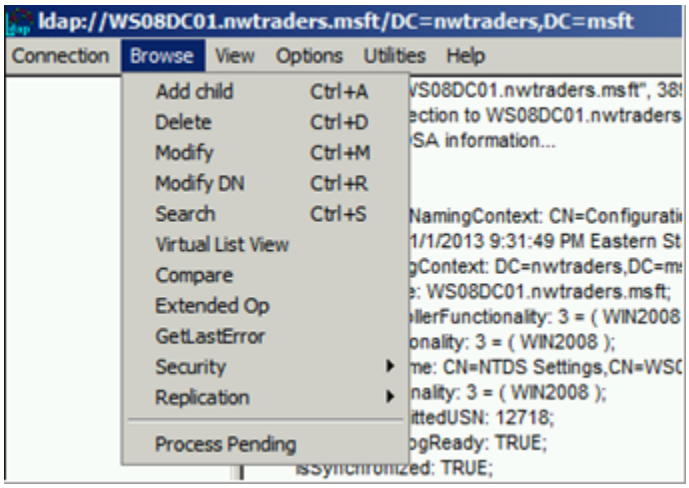

7. In the **Modify** dialog box, leave the **DN** field blank. In the **Edit Entry Attribute** field, type **FixUpInheritance**, and in the **Values** field, type **Yes**. Click **Enter** to populate the **Entry List** as shown in the following screen shot.

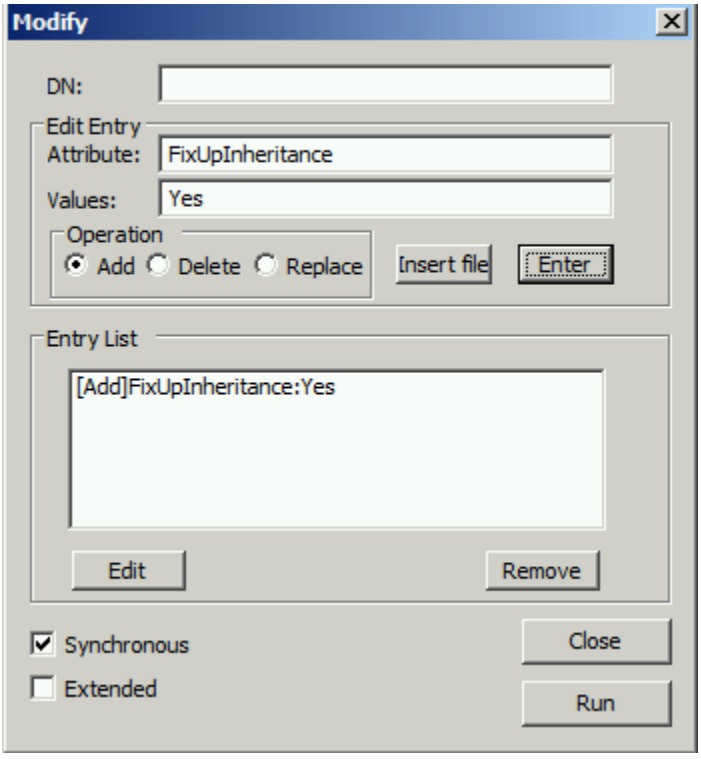

8. In the populated Modify dialog box, click Run, and verify that the changes you made to the AdminSDHolder object have appeared on that object.

Note

For information about modifying AdminSDHolder to allow designated unprivileged accounts to modify the membership of protected groups, see Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory.

If you prefer to run SDProp manually via LDIFDE or a script, you can create a modify entry as shown here:

```
dn:
changetype: modify
add: FixUpInheritance
FixUpInheritance: yes
```

Running SDProp Manually in Windows Server 2012 or Windows Server 2008 R2

You can also force SDProp to run by using Ldp.exe or by running an LDAP modification script. To run SDProp using Ldp.exe, perform the following steps after you have made changes to the AdminSDHolder object in a domain:

1. Launch **Ldp.exe**.
2. In the **Ldp** dialog box, click **Connection**, and click **Connect**.

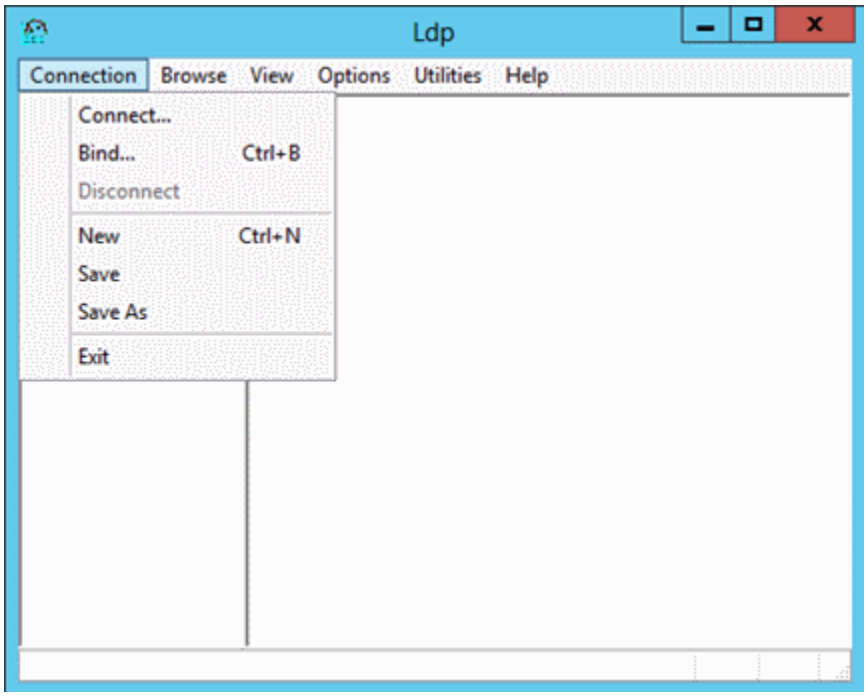

3. In the **Connect** dialog box, type the name of the domain controller for the domain that holds the PDC Emulator (PDCE) role and click **OK**.

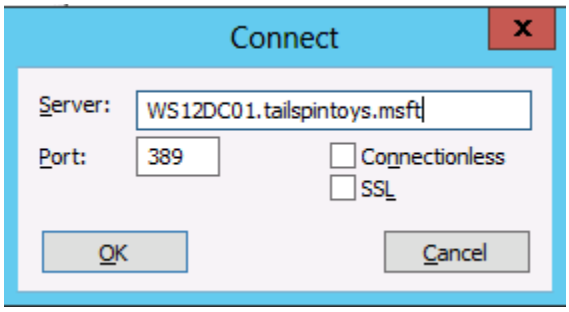

4. Verify that you have connected successfully, as indicated by **Dn: (RootDSE)** in the following screenshot, click **Connection** and click **Bind**.

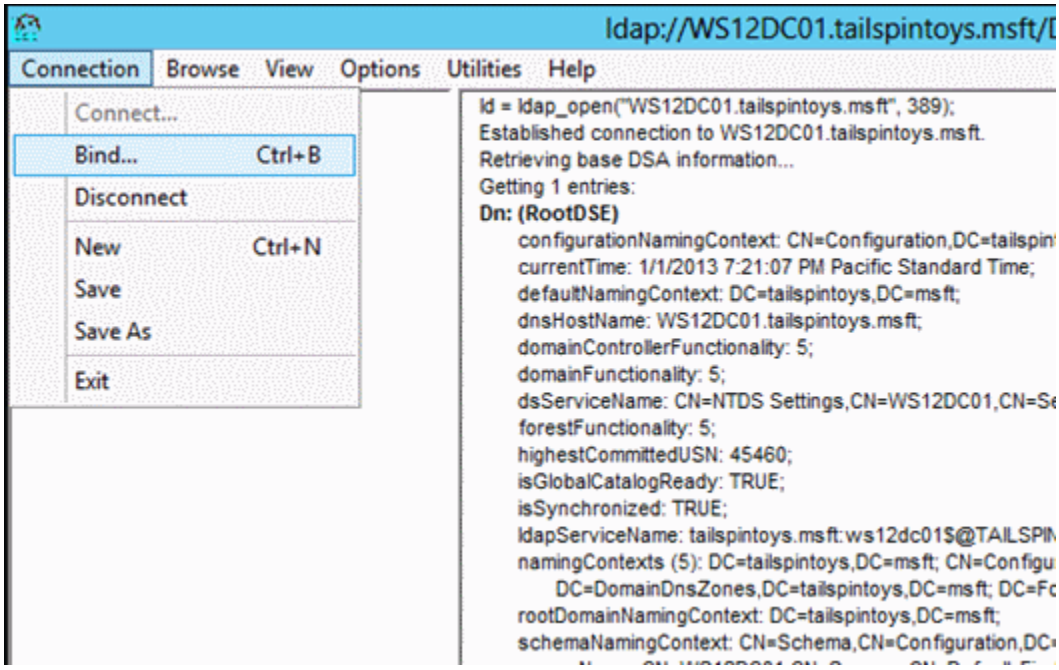

5. In the **Bind** dialog box, type the credentials of a user account that has permission to modify the rootDSE object. (If you are logged on as that user, you can select **Bind as currently logged on user**.) Click **OK**.

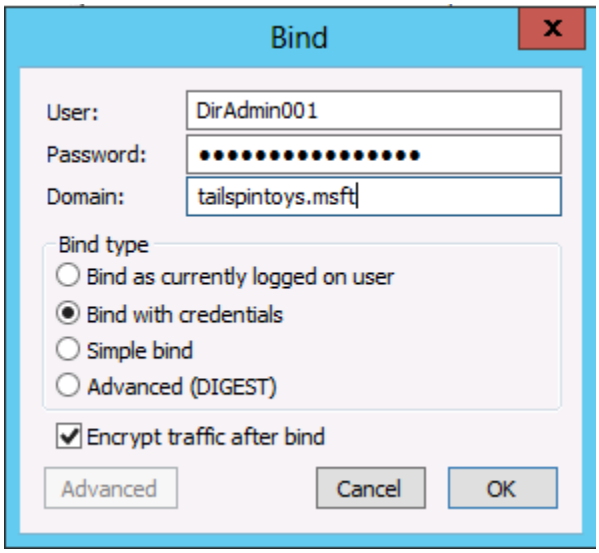

6. After you have completed the bind operation, click **Browse**, and click **Modify**.

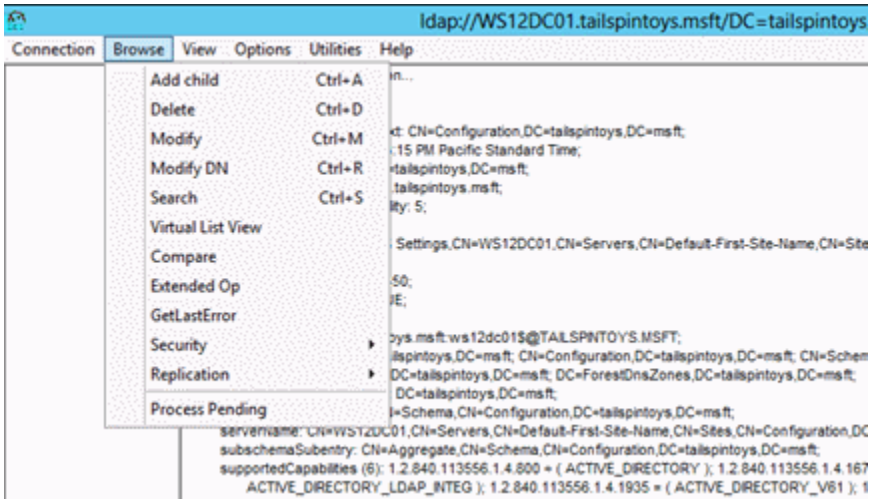

7. In the **Modify** dialog box, leave the **DN** field blank. In the **Edit Entry Attribute** field, type **RunProtectAdminGroupsTask**, and in the **Values** field, type **1**. Click **Enter** to populate the entry list as shown here.

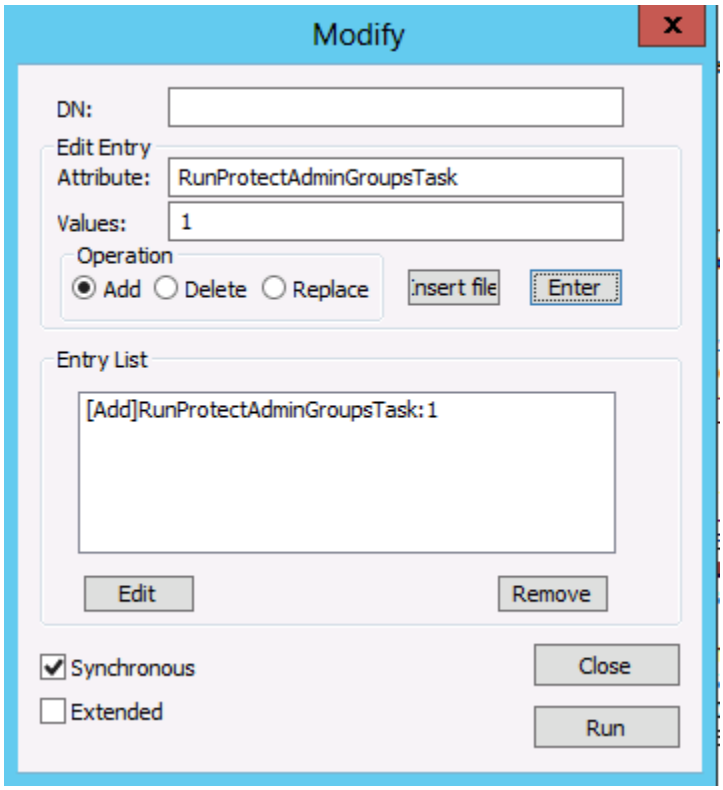

8. In the populated **Modify** dialog box, click **Run**, and verify that the changes you made to the AdminSDHolder object have appeared on that object.

If you prefer to run SDProp manually via LDIFDE or a script, you can create a modify entry as shown here:

```
dn:
changetype: modify
add: runProtectAdminGroupsTask
runProtectAdminGroupsTask: 1
```

Appendix D: Securing Built-In Administrator Accounts in Active Directory

Appendix D: Securing Built-In Administrator Accounts in Active Directory

In each domain in Active Directory, an Administrator account is created as part of the creation of the domain. This account is by default a member of the Domain Admins and Administrators

groups in the domain, and if the domain is the forest root domain, the account is also a member of the Enterprise Admins group.

Use of a domain's Administrator account should be reserved only for initial build activities, and possibly, disaster-recovery scenarios. To ensure that an Administrator account can be used to effect repairs in the event that no other accounts can be used, you should not change the default membership of the Administrator account in any domain in the forest. Instead, you should secure the Administrator account in each domain in the forest as described in the following section and detailed in the step-by-step instructions that follow.

Controls for Built-in Administrator Accounts

For the built-in Administrator account in each domain in your forest, you should configure the following settings:

- Enable the **Account is sensitive and cannot be delegated** flag on the account.
- Enable the **Smart card is required for interactive logon** flag on the account.
- Disable the account.
- Configure GPOs to restrict the Administrator account's use on domain-joined systems:
  - In one or more GPOs that you create and link to workstation and member server OUs in each domain, add each domain's Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:
    - Deny access to this computer from the network
    - Deny log on as a batch job
    - Deny log on as a service
    - Deny log on through Remote Desktop Services

Notes

When you add accounts to this setting, you must specify whether you are configuring local Administrator accounts or domain Administrator accounts. For example, to add the NWTRADERS domain's Administrator account to these deny rights, you must type the account as NWTRADERS\Administrator, or browse to the Administrator account for the NWTRADERS domain. If you type "Administrator" in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied.

We recommend restricting local Administrator accounts on member servers and workstations in the same manner as domain-based Administrator accounts. Therefore you should generally add the Administrator account for each domain in the forest and the Administrator account for the local computers to these user rights settings. The following screenshot shows an example of configuring these user rights to block local Administrator accounts and a domain's Administrator account from performing logons that should not be needed for these accounts.

- Configure GPOs to restrict Administrator accounts on domain controllers
  - In each domain in the forest, the Default Domain Controllers GPO or a policy linked to the domain controllers OU should be modified to add each domain's Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:
    - Deny access to this computer from the network
    - Deny log on as a batch job
    - Deny log on as a service
    - Deny log on through Remote Desktop Services

Note

These settings will ensure that the domain's built-in Administrator account cannot be used to connect to a domain controller, although the account, if enabled, can log on locally to domain controllers. Because this account should only be enabled and used in disaster-recovery scenarios, it is anticipated that physical access to at least one domain controller will be available, or that other accounts with permissions to access domain controllers remotely can be used.

- Configure Auditing of Administrator Accounts

  When you have secured each domain's Administrator account and disabled it, you should configure auditing to monitor for changes to the account. If the account is enabled, its password is reset, or any other modifications are made to the account, alerts should be sent to the users or teams responsible for administration of Active Directory, in addition to incident response teams in your organization.

Step-by-Step Instructions to Secure Built-in Administrator Accounts in Active Directory

1. In **Server Manager**, click **Tools**, and click **Active Directory Users and Computers**.
2. To prevent attacks that leverage delegation to use the account's credentials on other systems, perform the following steps:

   a. Right-click the **Administrator** account and click **Properties**.

b. Click the **Account** tab.

c. Under **Account options**, select **Account is sensitive and cannot be delegated** flag as indicated in the following screenshot, and click **OK**.

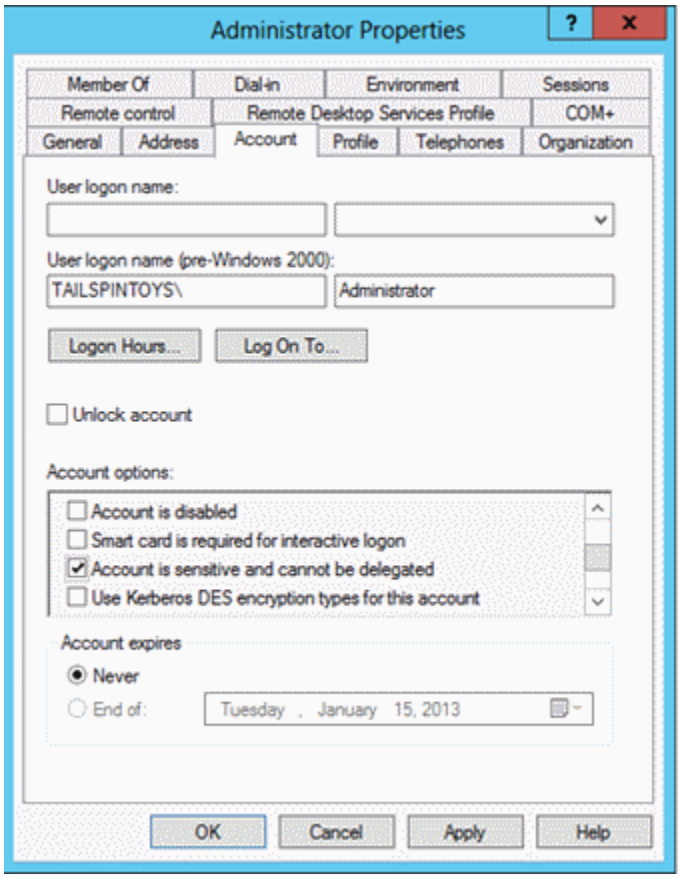

3. To enable the **Smart card is required for interactive logon** flag on the account, perform the following steps:

a. Right-click the **Administrator** account and select **Properties**.

b. Click the **Account** tab.

c. Under **Account** options, select the **Smart card is required for interactive logon** flag as indicated in the following screenshot, and click **OK**.

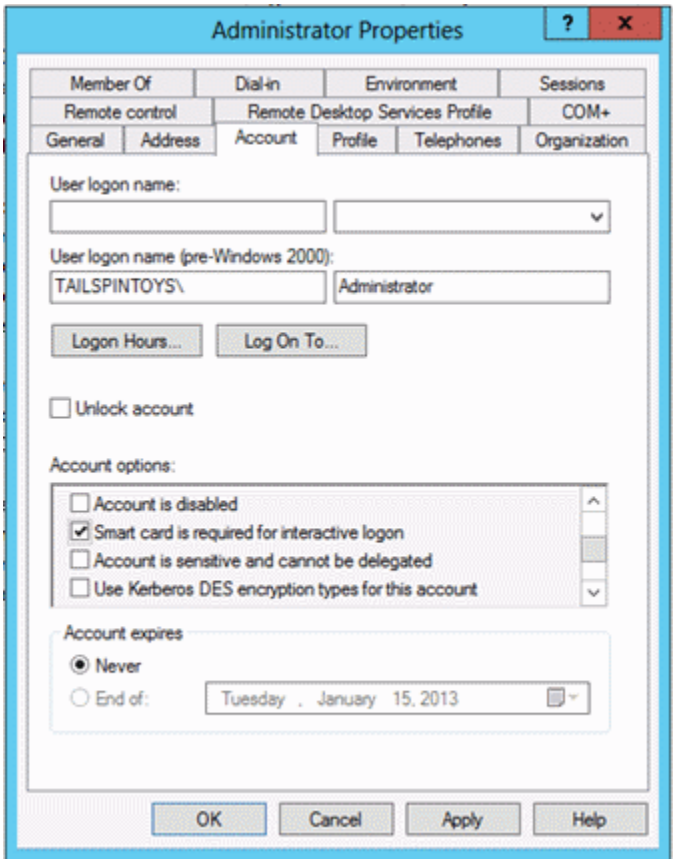

4. To disable the account, perform the following steps:

   a. Right-click the **Administrator** account and click **Properties**.
   b. Click the **Account** tab.
   c. In the **Account** options field, select the **Account is disabled** flag as indicated in the following screenshot, and click **OK**.

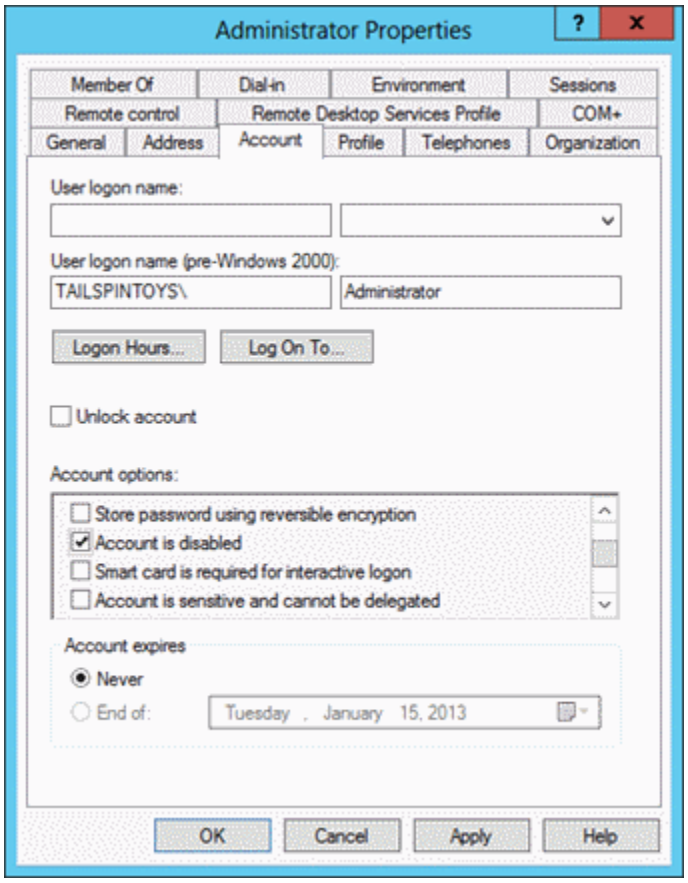

Configuring GPOs to Restrict Administrator Accounts at the Domain-Level

Warning

This GPO should never be linked at the domain-level because it can make the built-in Administrator account unusable, even in disaster recovery scenarios.

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.
2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to create the Group Policy).
3. In the console tree, right-click **Group Policy Objects**, and click **New**.

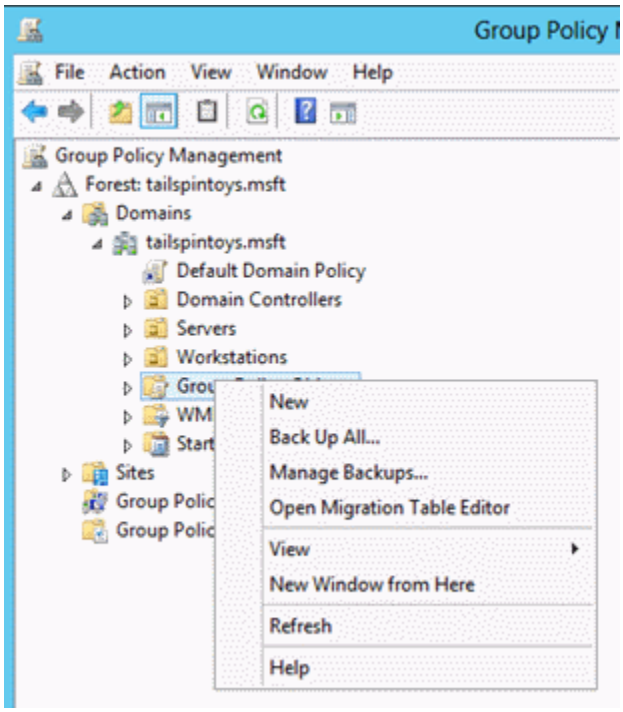

4. In the **New GPO** dialog box, type <GPO Name>, and click **OK** (where <GPO Name> is the name of this GPO) as indicated in the following screenshot.

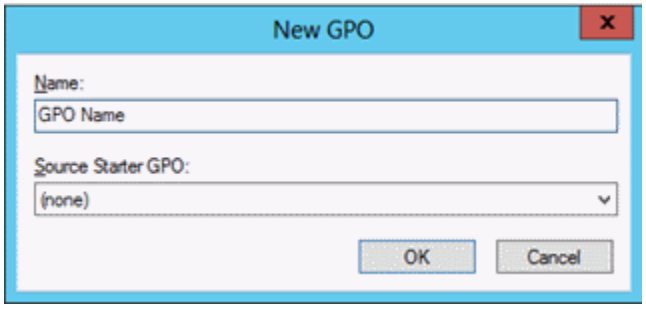

5. In the details pane, right-click <GPO Name>, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

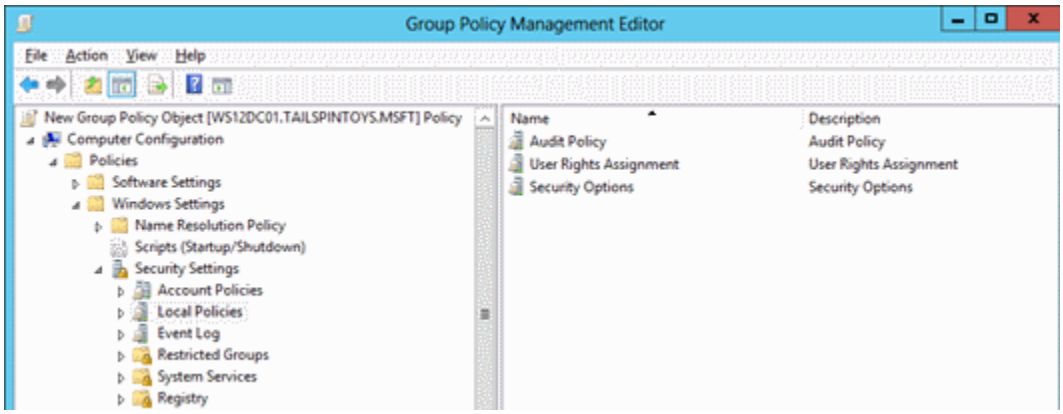

7. Configure the user rights to prevent the Administrator account from accessing members servers and workstations over the network by doing the following:

   a. Double-click **Deny access to this computer from the network** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Administrator**, click **Check Names**, and click **OK**. Verify that the account is displayed in <DomainName>\Username format as indicated in the following screenshot.

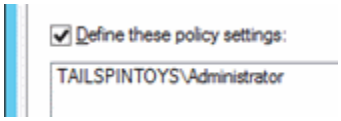

   d. Click **OK**, and **OK** again.

8. Configure the user rights to prevent the Administrator account from logging on as a batch job by doing the following:

   a. Double-click **Deny log on as a batch job** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Administrator**, click **Check Names**, and click **OK**. Verify that the account is displayed in <DomainName>\Username format as indicated in the following screenshot.

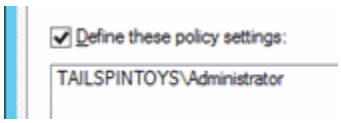

d. Click **OK**, and **OK** again.

9. Configure the user rights to prevent the Administrator account from logging on as a service by doing the following:

a. Double-click **Deny log on as a service** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Administrator**, click **Check Names**, and click **OK**. Verify that the account is displayed in <DomainName>\Username format as indicated in the following screenshot.

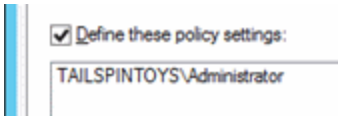

d. Click **OK**, and **OK** again.

10. Configure the user rights to prevent the BA account from accessing member servers and workstations via Remote Desktop Services by doing the following:

a. Double-click **Deny log on through Remote Desktop Services** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Administrator**, click **Check Names**, and click **OK**. Verify that the account is displayed in <DomainName>\Username format as indicated in the following screenshot.

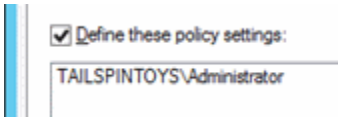

d. Click **OK**, and **OK** again.

11. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
12. In **Group Policy Management**, link the GPO to the member server and workstation OUs by doing the following:

a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set

the Group Policy).

b. Right-click the OU that the GPO will be applied to and click **Link an existing GPO**.

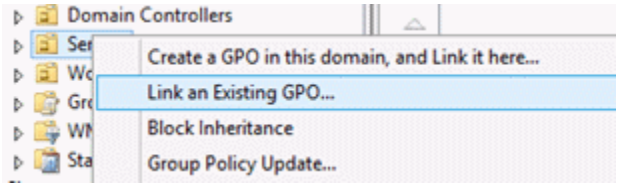

c. Select the GPO that you created and click **OK**.

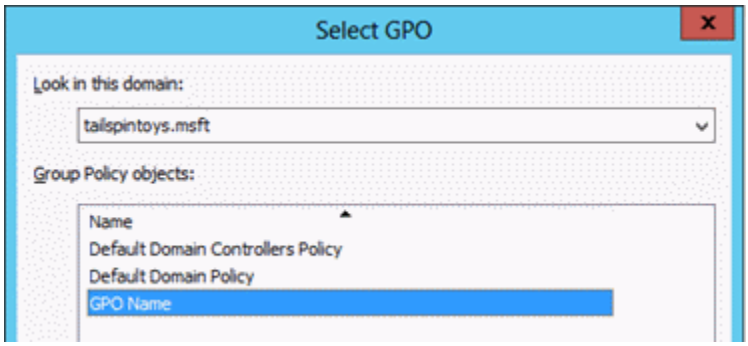

d. Create links to all other OUs that contain workstations.

e. Create links to all other OUs that contain member servers.

Important

When you add the Administrator account to these settings, you specify whether you are configuring a local Administrator account or a domain Administrator account by how you label the accounts. For example, to add the TAILSPINTOYS domain's Administrator account to these deny rights, you would browse to the Administrator account for the TAILSPINTOYS domain, which would appear as TAILSPINTOYS\Administrator. If you type "Administrator" in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied, as described earlier.

Verification Steps

The verification steps outlined here are specific to Windows 8 and Windows Server 2012.

Verify "Smart card is required for interactive logon" Account Option

1. From any member server or workstation affected by the GPO changes, attempt to log on interactively to the domain by using the domain's built-in Administrator account. After attempting to log on, a dialog box similar to the following should appear.

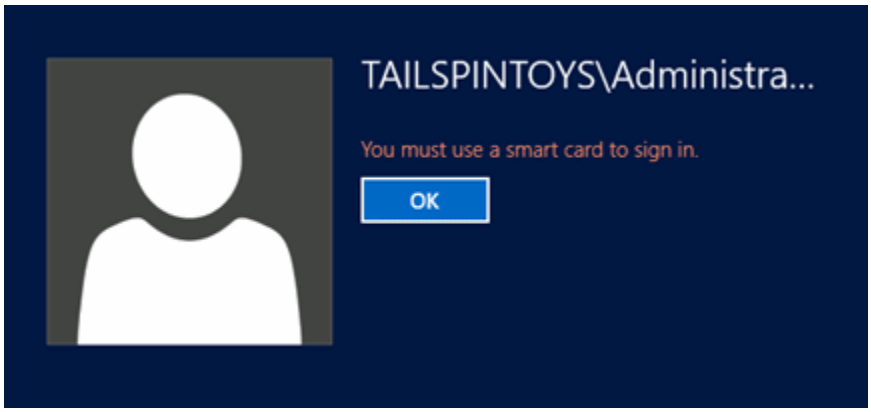

Verify "Account is disabled" Account Option

1. From any member server or workstation affected by the GPO changes, attempt to log on interactively to the domain by using the domain's built-in Administrator account. After attempting to log on, a dialog box similar to the following should appear.

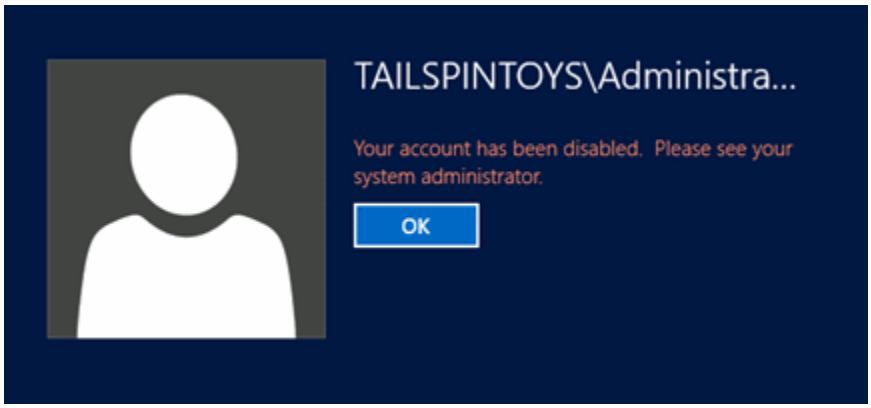

Verify "Deny access to this computer from the network" GPO Settings

From any member server or workstation that is not affected by the GPO changes (such as a jump server), attempt to access a member server or workstation over the network that is affected by the GPO changes. To verify the GPO settings, attempt to map the system drive by using the **NET USE** command by performing the following steps:

1. Log on to the domain using the domain's built-in Administrator account.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **command prompt**, right-click **Command Prompt**, and then click **Run as administrator** to open an elevated command prompt.
4. When prompted to approve the elevation, click **Yes**.

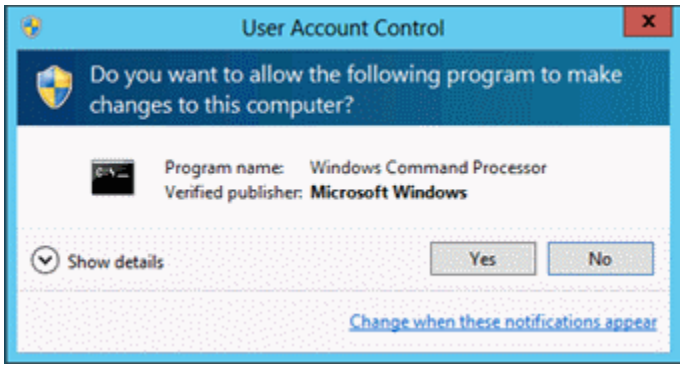

5. In the **Command Prompt** window, type **net use \\<Server Name>\c$**, where <Server Name> is the name of the member server or workstation you are attempting to access over the network.
6. The following screenshot shows the error message that should appear.

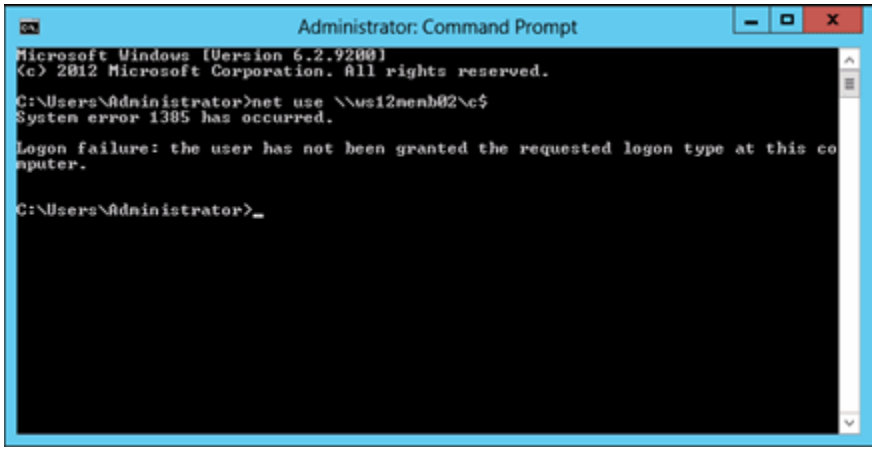

Verify "Deny log on as a batch job" GPO Settings

From any member server or workstation affected by the GPO changes, log on locally.

Create a Batch File

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **notepad**, and click **Notepad**.
3. In **Notepad**, type **dir c:**.
4. Click **File** and click **Save As**.
5. In the **Filename** field, type **<Filename>.bat** (where <Filename> is the name of the new batch file).

Schedule a Task

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **task scheduler**, and click **Task Scheduler**.

   Note
   On computers running Windows 8, in the Search box, type **schedule tasks**, and click **Schedule tasks**.
3. On **Task Scheduler**, click **Action**, and click **Create Task**.
4. In the **Create Task** dialog box, type **<Task Name>** (where **<Task Name>** is the name of the new task).
5. Click the **Actions** tab, and click **New**.
6. Under **Action:**, select **Start a program**.
7. Under **Program/script:**, click **Browse**, locate and select the batch file created in the "Create a Batch File" section, and click **Open**.
8. Click **OK**.
9. Click the **General** tab.
10. Under **Security** options, click **Change User or Group**.
11. Type the name of the BA account at the domain-level, click **Check Names**, and click **OK**.
12. Select **Run whether the user is logged on or not** and **Do not store password**. The task will only have access to local computer resources.
13. Click **OK**.
14. A dialog box should appear, requesting user account credentials to run the task.
15. After entering the credentials, click **OK**.
16. A dialog box similar to the following should appear.

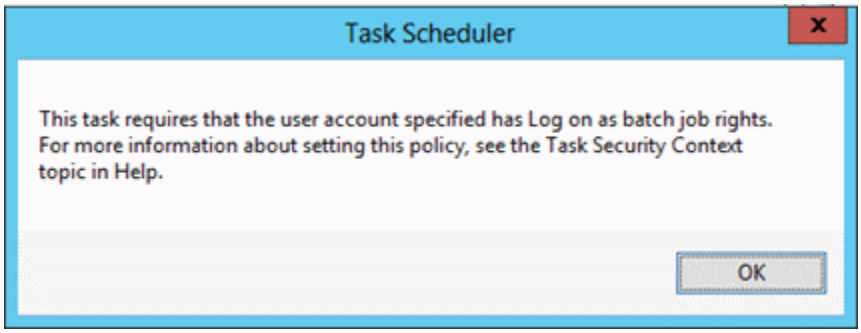

Verify "Deny log on as a service" GPO Settings

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen.

When the **Charms** bar appears, click **Search**.

3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as:**, select **This account**.
7. Click **Browse**, type the name of the BA account at the domain-level, click **Check Names**, and click **OK**.
8. Under **Password:** and **Confirm password:**, type the Administrator account's password, and click **OK**.
9. Click **OK** three more times.
10. Right-click the **Print Spooler service** and select **Restart**.
11. When the service is restarted, a dialog box similar to the following should appear.

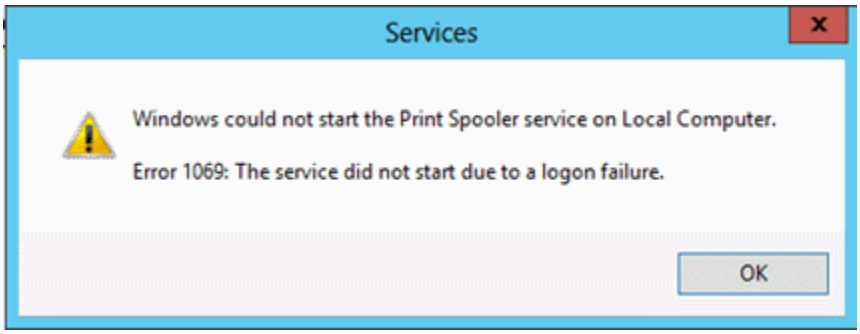

Revert Changes to the Printer Spooler Service

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as:**, select the **Local System** account, and click **OK**.

Verify "Deny log on through Remote Desktop Services" GPO Settings

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **remote desktop connection**, and click **Remote Desktop Connection**.

3. In the **Computer** field, type the name of the computer that you want to connect to, and click **Connect**. (You can also type the IP address instead of the computer name.)
4. When prompted, provide credentials for the name of the BA account at the domain-level.
5. A dialog box similar to the following should appear.

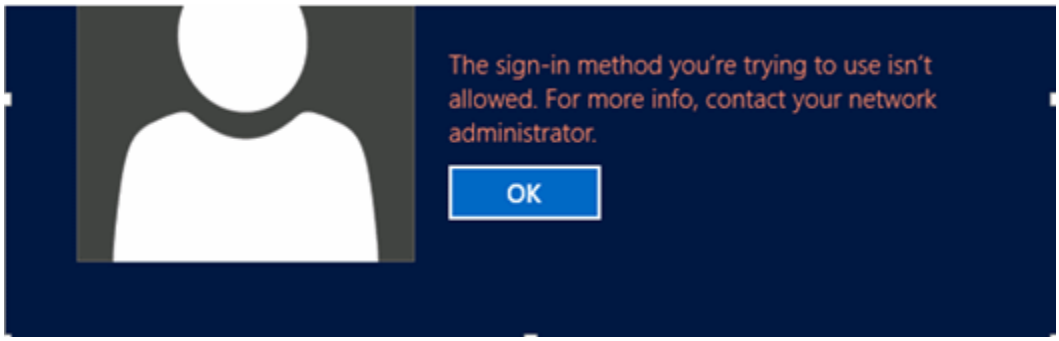

Appendix E: Securing Enterprise Admins Groups in Active Directory

Appendix E: Securing Enterprise Admins Groups in Active Directory

The Enterprise Admins (EA) group, which is housed in the forest root domain, should contain no users on a day-to-day basis, with the possible exception of the root domain's Administrator account, provided it is secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

Enterprise Admins are, by default, members of the Administrators group in each domain in the forest. You should not remove the EA group from the Administrators groups in each domain because in the event of a forest disaster recovery scenario, EA rights will likely be required. The forest's Enterprise Admins group should be secured as detailed in the step-by-step instructions that follow.

For the Enterprise Admins group in the forest:

1. In GPOs linked to OUs containing member servers and workstations in each domain, the Enterprise Admins group should be added to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:
   - Deny access to this computer from the network
   - Deny log on as a batch job
   - Deny log on as a service

- Deny log on locally
- Deny log on through Remote Desktop Services

2. Configure auditing to send alerts if any modifications are made to the properties or membership of the Enterprise Admins group.

Step-by-Step Instructions for Removing All Members from the Enterprise Admins Group

1. In **Server Manager**, click **Tools**, and click **Active Directory Users and Computers**.
2. If you are not managing the root domain for the forest, in the console tree, right-click <Domain>, and then click **Change Domain** (where <Domain> is the name of the domain you're currently administering).

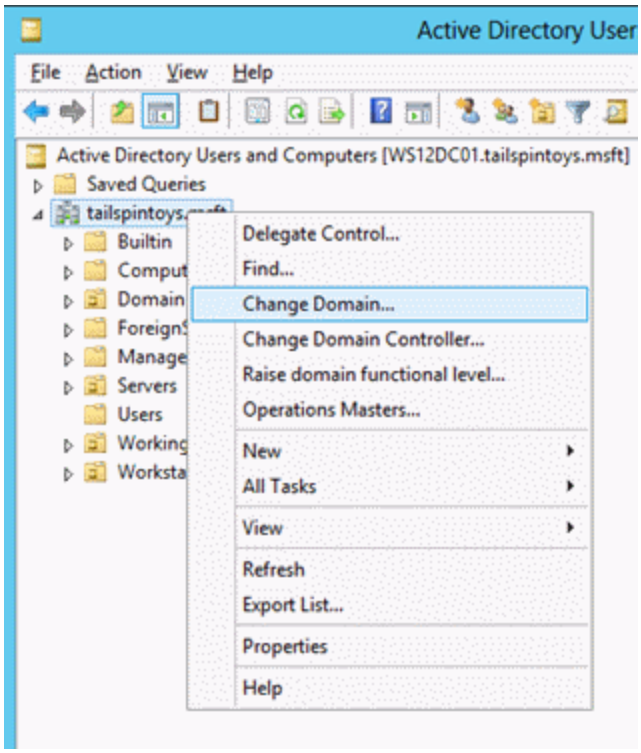

3. In the **Change domain** dialog box, click **Browse**, select the root domain for the forest, and click **OK**.

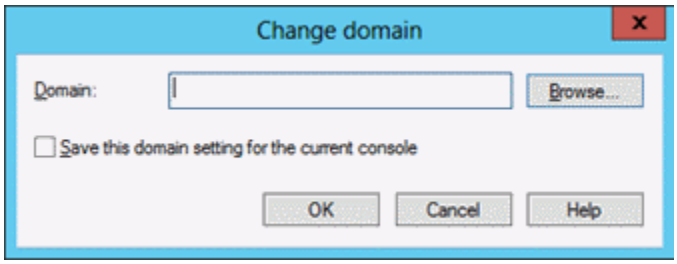

4. To remove all members from the EA group:

   a. Double-click the **Enterprise Admins** group and then click the **Members** tab.

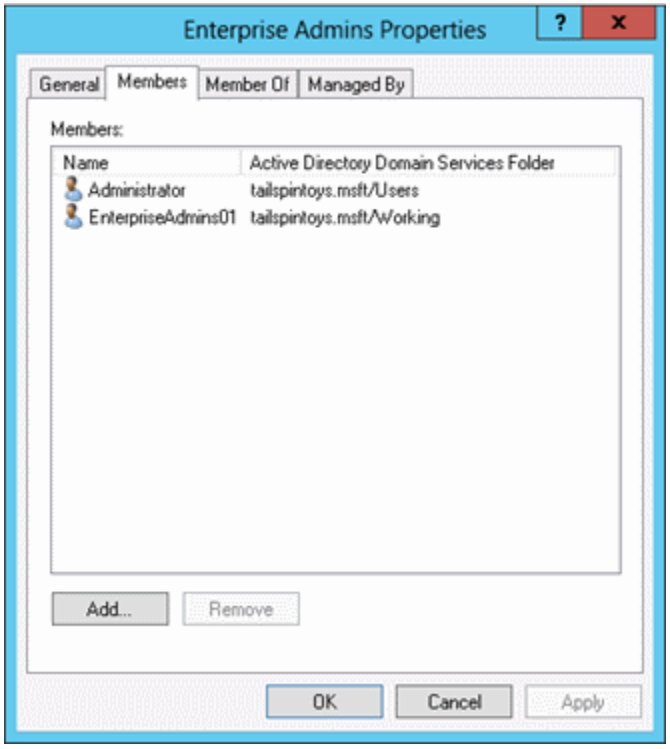

   b. Select a member of the group, click **Remove**, click **Yes**, and click **OK**.

5. Repeat step 2 until all members of the EA group have been removed.

Step-by-Step Instructions to Secure Enterprise Admins in Active Directory

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.

2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).

   Note

   In a forest that contains multiple domains, a similar GPO should be created in each domain that requires that the Enterprise Admins group be secured.

3. In the console tree, right-click **Group Policy Objects**, and click **New**.

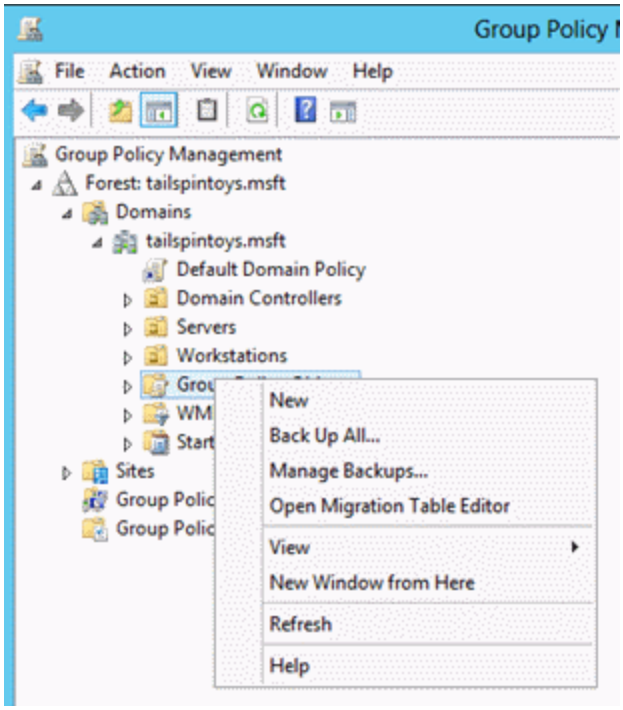

4. In the **New GPO** dialog box, type <GPO Name>, and click **OK** (where <GPO Name> is the name of this GPO).

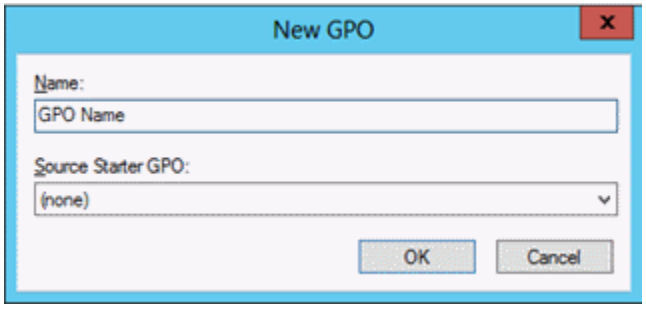

5. In the details pane, right-click <GPO Name>, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

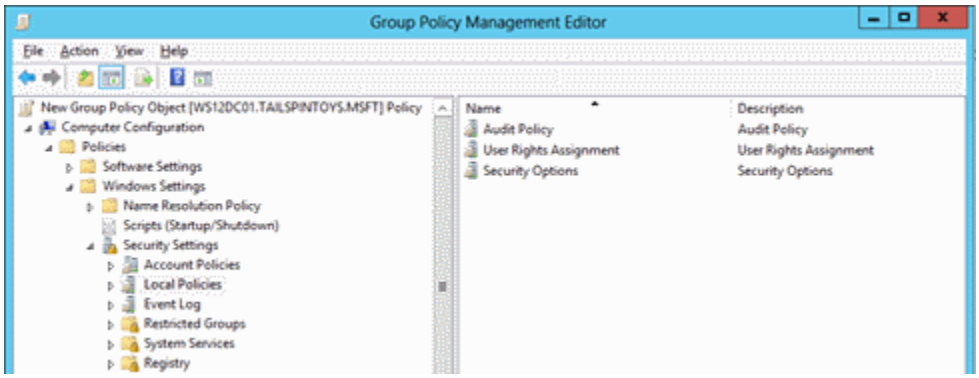

7. Configure the user rights to prevent members of the Enterprise Admins group from accessing member servers and workstations over the network by doing the following:

a. Double-click **Deny access to this computer from the network** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Enterprise Admins**, click **Check Names**, and click **OK**.

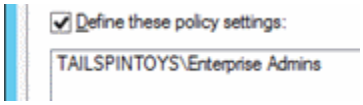

d. Click **OK**, and **OK** again.

8. Configure the user rights to prevent members of the Enterprise Admins group from logging on as a batch job by doing the following:

a. Double-click **Deny log on as a batch job** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.

> **Note**
> In a forest that contains multiple domains, click **Locations** and select the root domain of the forest.

c. Type **Enterprise Admins**, click **Check Names**, and click **OK**.

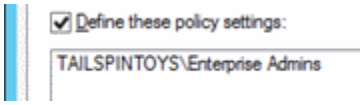

d. Click **OK**, and **OK** again.

9. Configure the user rights to prevent members of the EA group from logging on as a service by doing the following:

   a. Double-click **Deny log as a service** and select **Define these policy settings**.
   b. Click **Add User or Group** and then click **Browse**.

      Note
      In a forest that contains multiple domains, click **Locations** and select the root domain of the forest.

   c. Type **Enterprise Admins**, click **Check Names**, and click **OK**.

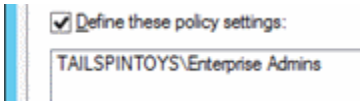

   d. Click **OK**, and **OK** again.

10. Configure user rights to prevent members of the Enterprise Admins group from logging on locally to member servers and workstations by doing the following:

   a. Double-click **Deny log on locally** and select **Define these policy settings**.
   b. Click **Add User or Group** and then click **Browse**.

      Note
      In a forest that contains multiple domains, click **Locations** and select the root domain of the forest.

   c. Type **Enterprise Admins**, click **Check Names**, and click **OK**.

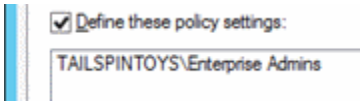

   d. Click **OK**, and **OK** again.

11. Configure the user rights to prevent members of the Enterprise Admins group from accessing member servers and workstations via Remote Desktop Services by doing the following:

   a. Double-click **Deny log on through Remote Desktop Services** and select **Define these policy settings**.

b. Click **Add User or Group** and then click **Browse**.

Note
In a forest that contains multiple domains, click **Locations** and select the root domain of the forest.

c. Type **Enterprise Admins**, click **Check Names**, and click **OK**.

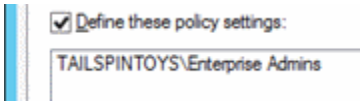

d. Click **OK**, and **OK** again.

12. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
13. In **Group Policy Management**, link the GPO to the member server and workstation OUs by doing the following:

a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).

b. Right-click the OU that the GPO will be applied to and click **Link an existing GPO**.

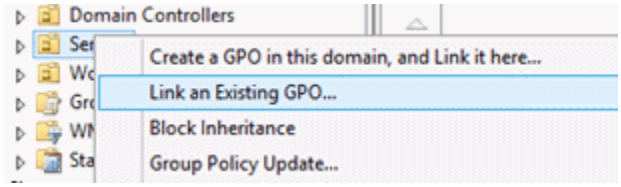

c. Select the GPO that you just created and click **OK**.

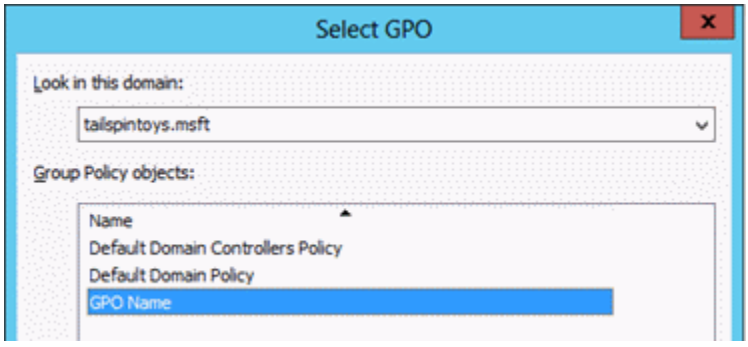

d. Create links to all other OUs that contain workstations.

e. Create links to all other OUs that contain member servers.

f. In a forest that contains multiple domains, a similar GPO should be created in each domain that requires that the Enterprise Admins group be secured.

Important

If jump servers are used to administer domain controllers and Active Directory, ensure that jump servers are located in an OU to which this GPOs is not linked.

Verification Steps

Verify “Deny access to this computer from the network” GPO Settings

From any member server or workstation that is not affected by the GPO changes (such as a “jump server”), attempt to access a member server or workstation over the network that is affected by the GPO changes. To verify the GPO settings, attempt to map the system drive by using the **NET USE** command by performing the following steps:

1. Log on locally using an account that is a member of the EA group.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **command prompt**, right-click **Command Prompt**, and then click **Run as administrator** to open an elevated command prompt.
4. When prompted to approve the elevation, click **Yes**.

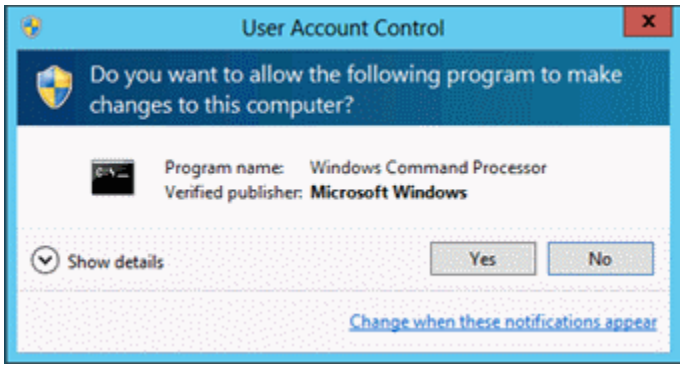

5. In the **Command Prompt** window, type **net use \\<Server Name>\c$**, where <Server Name> is the name of the member server or workstation you’re attempting to access over the network.
6. The following screen shot shows the error message that should appear.

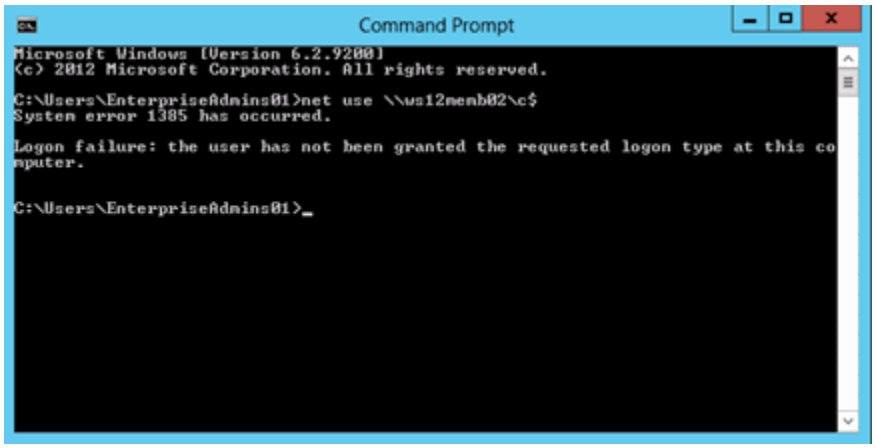

Verify “Deny log on as a batch job” GPO Settings

From any member server or workstation affected by the GPO changes, log on locally.

Create a Batch File

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **notepad**, and click **Notepad**.
3. In **Notepad**, type **dir c:**.
4. Click **File**, and click **Save As**.
5. In the **File** name box, type **<Filename>.bat** (where <Filename> is the name of the new batch file).

Schedule a Task

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **task scheduler**, and click **Task Scheduler**.

   Note
   On computers running Windows 8, in the **Search** box, type **schedule tasks**, and click **Schedule tasks**.
3. Click **Action**, and click **Create Task**.
4. In the **Create Task** dialog box, type **<Task Name>** (where <Task Name> is the name of the new task).
5. Click the **Actions** tab, and click **New**.

6. In the **Action** field, select **Start a program**.
7. Under **Program/script**, click **Browse**, locate and select the batch file created in the **Create a Batch File** section, and click **Open**.
8. Click **OK**.
9. Click the **General** tab.
10. In the **Security options** field, click **Change User or Group**.
11. Type the name of an account that is a member of the EAs group, click **Check Names**, and click **OK**.
12. Select **Run whether the user is logged on or not** and select **Do not store password**. The task will only have access to local computer resources.
13. Click **OK**.
14. A dialog box should appear, requesting user account credentials to run the task.
15. After entering the credentials, click **OK**.
16. A dialog box similar to the following should appear.

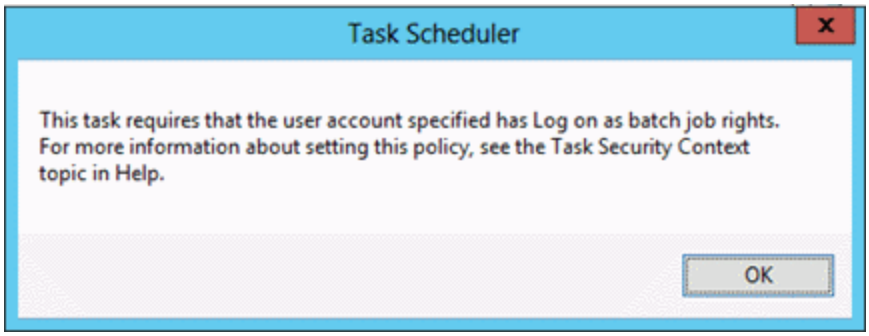

Verify “Deny log on as a service” GPO Settings

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as**, select **This account**.
7. Click **Browse**, type the name of an account that is a member of the EAs group, click **Check Names**, and click **OK**.
8. Under **Password:** and **Confirm password**, type the selected account’s password, and click **OK**.
9. Click **OK** three more times.

10. Right-click the **Print Spooler** service and select **Restart**.
11. When the service is restarted, a dialog box similar to the following should appear.

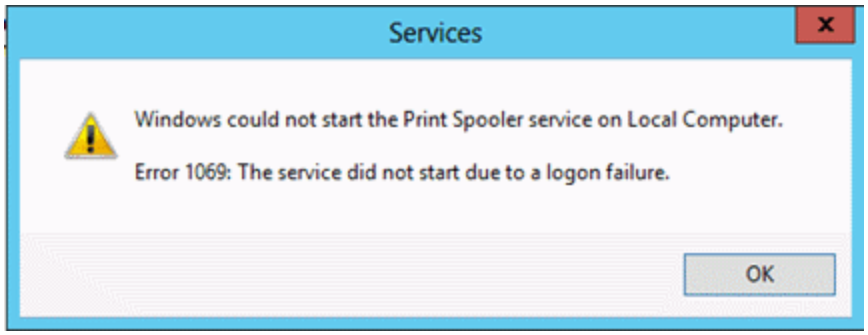

Revert Changes to the Printer Spooler Service

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as**, select the **Local System** account, and click **OK**.

Verify “Deny log on locally” GPO Settings

1. From any member server or workstation affected by the GPO changes, attempt to log on locally using an account that is a member of the EA group. A dialog box similar to the following should appear.

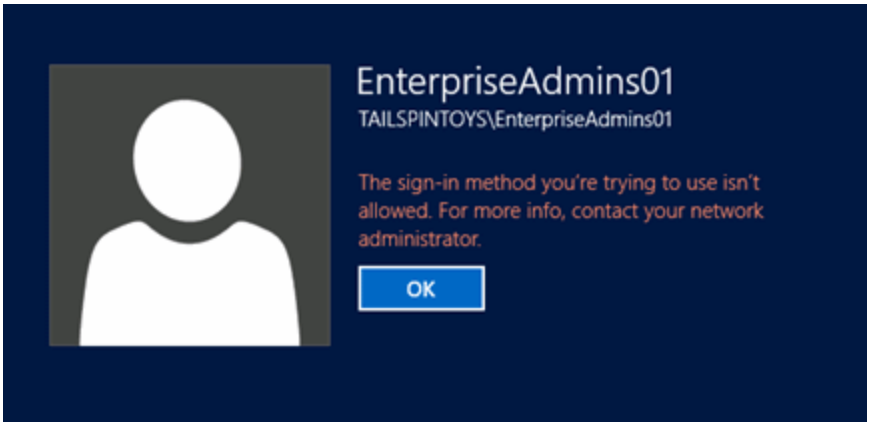

Verify “Deny log on through Remote Desktop Services” GPO Settings

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **remote desktop connection**, and then click **Remote Desktop Connection**.
3. In the **Computer** field, type the name of the computer that you want to connect to, and then click **Connect**. (You can also type the IP address instead of the computer name.)
4. When prompted, provide credentials for an account that is a member of the EA group.
5. A dialog box similar to the following should appear.

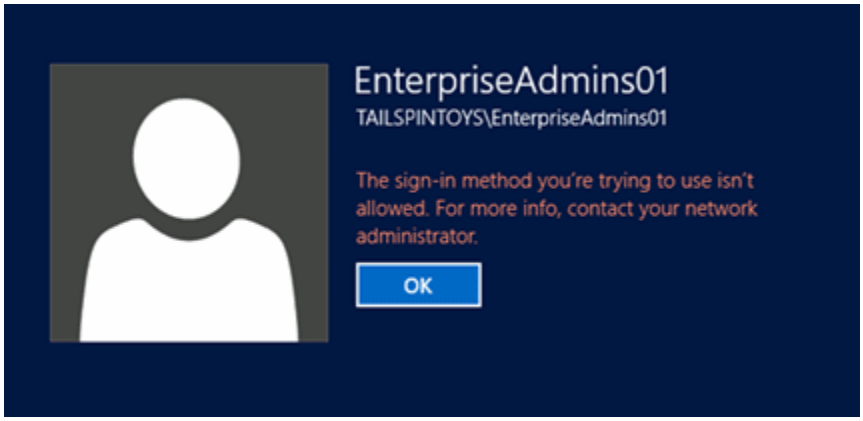

Appendix F: Securing Domain Admins Groups in Active Directory

Appendix F: Securing Domain Admins Groups in Active Directory

As is the case with the Enterprise Admins (EA) group, membership in the Domain Admins (DA) group should be required only in build or disaster recovery scenarios. There should be no day-to-day user accounts in the DA group with the exception of the built-in Administrator account for the domain, if it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

Domain Admins are, by default, members of the local Administrators groups on all member servers and workstations in their respective domains. This default nesting should not be modified for supportability and disaster recovery purposes. If Domain Admins have been removed from the

local Administrators groups on the member servers, the group should be added to the Administrators group on each member server and workstation in the domain. Each domain's Domain Admins group should be secured as described in the step-by-step instructions that follow.

For the Domain Admins group in each domain in the forest:

1. Remove all members from the group, with the possible exception of the built-in Administrator account for the domain, provided it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.
2. In GPOs linked to OUs containing member servers and workstations in each domain, the DA group should be added to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:
   - Deny access to this computer from the network
   - Deny log on as a batch job
   - Deny log on as a service
   - Deny log on locally
   - Deny log on through Remote Desktop Services user rights
3. Auditing should be configured to send alerts if any modifications are made to the properties or membership of the Domain Admins group.

Step-by-Step Instructions for Removing all Members from the Domain Admins Group

1. In **Server Manager**, click **Tools**, and click **Active Directory Users and Computers**.
2. To remove all members from the DA group, perform the following steps:

   a. Double-click the **Domain Admins** group and click the **Members** tab.

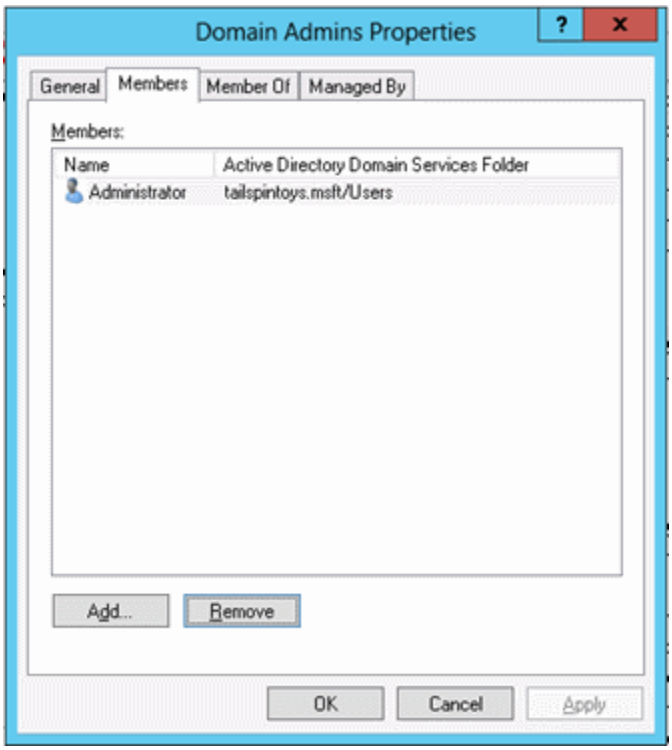

b. Select a member of the group, click **Remove**, click **Yes**, and click **OK**.

3. Repeat step 2 until all members of the DA group have been removed.

Step-by-Step Instructions to Secure Domain Admins in Active Directory

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.
2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
3. In the console tree, right-click **Group Policy Objects**, and click **New**.

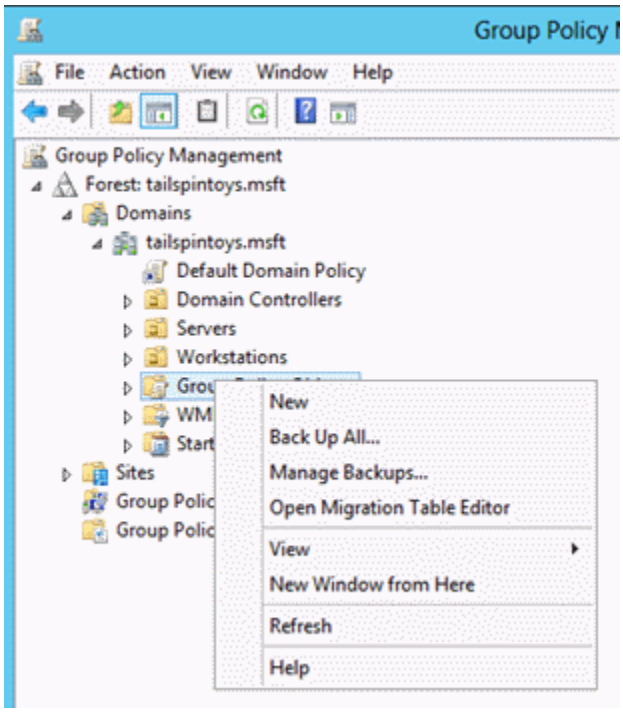

4. In the **New GPO** dialog box, type <GPO Name>, and click **OK** (where GPO Name is the name of this GPO).

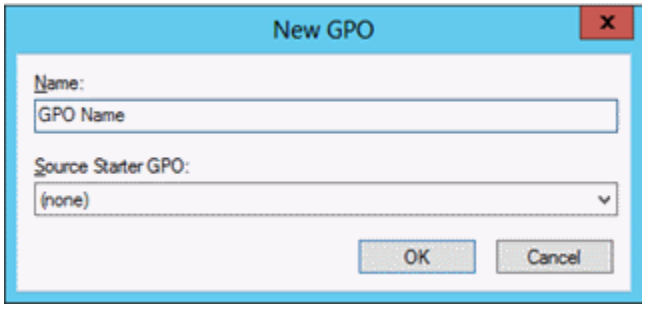

5. In the details pane, right-click <GPO Name>, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

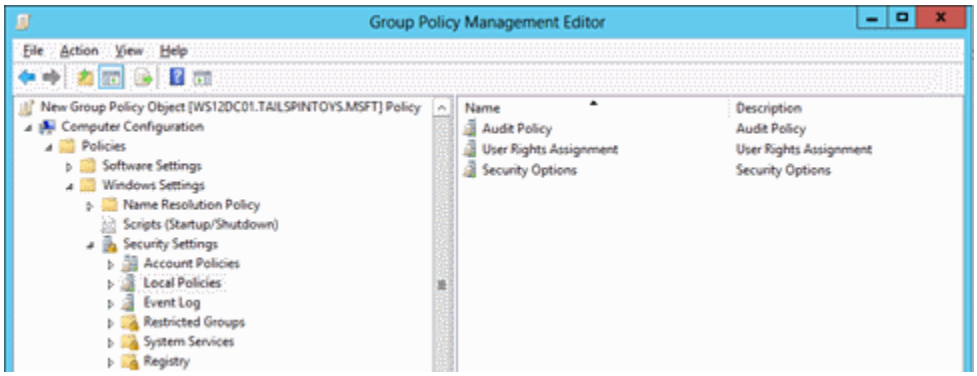

7. Configure the user rights to prevent members of the Domain Admins group from accessing members servers and workstations over the network by doing the following:

   a. Double-click **Deny access to this computer from the network** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Domain Admins**, click **Check Names**, and click **OK**.

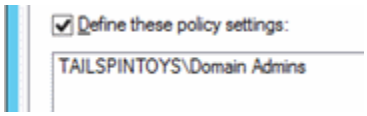

   d. Click **OK**, and **OK** again.

8. Configure the user rights to prevent members of the DA group from logging on as a batch job by doing the following:

   a. Double-click **Deny log on as a batch job** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Domain Admins**, click **Check Names**, and click **OK**.

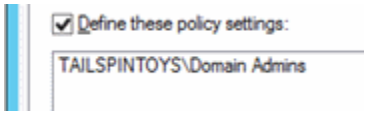

   d. Click **OK**, and **OK** again.

9. Configure the user rights to prevent members of the DA group from logging on as a service by doing the following:

   a. Double-click **Deny log on as a service** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Domain Admins**, click **Check Names**, and click **OK**.

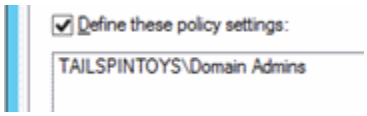

d. Click **OK**, and **OK** again.

10. Configure the user rights to prevent members of the Domain Admins group from logging on locally to member servers and workstations by doing the following:

a. Double-click **Deny log on locally** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Domain Admins**, click **Check Names**, and click **OK**.

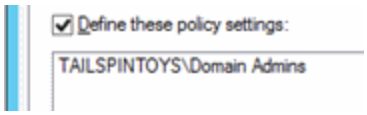

d. Click **OK**, and **OK** again.

11. Configure the user rights to prevent members of the Domain Admins group from accessing member servers and workstations via Remote Desktop Services by doing the following:

a. Double-click **Deny log on through Remote Desktop Services** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Domain Admins**, click **Check Names**, and click **OK**.

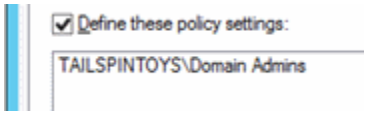

d. Click **OK**, and **OK** again.

12. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
13. In Group Policy Management, link the GPO to the member server and workstation OUs by doing the following:

a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
b. Right-click the OU that the GPO will be applied to and click **Link an existing GPO**.

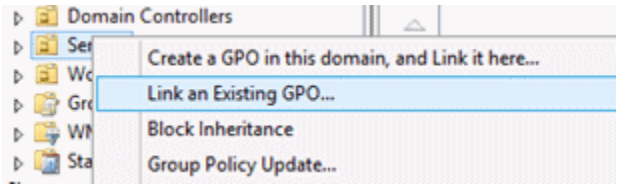

c. Select the GPO that you just created and click **OK**.

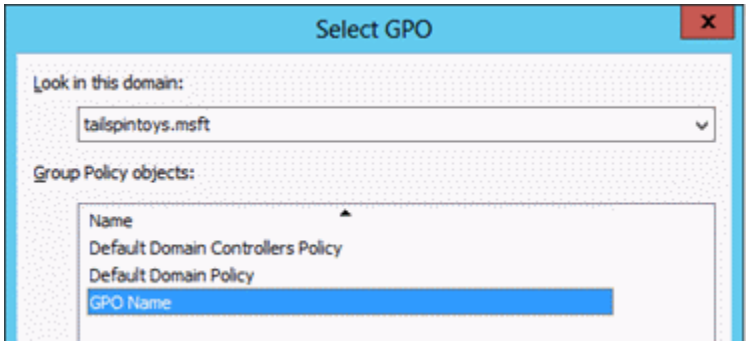

d. Create links to all other OUs that contain workstations.

e. Create links to all other OUs that contain member servers.

Important
If jump servers are used to administer domain controllers and Active Directory, ensure that jump servers are located in an OU to which this GPOs is not linked.

Verification Steps

Verify "Deny access to this computer from the network" GPO Settings

From any member server or workstation that is not affected by the GPO changes (such as a "jump server"), attempt to access a member server or workstation over the network that is affected by the GPO changes. To verify the GPO settings, attempt to map the system drive by using the **NET USE** command.

1. Log on locally using an account that is a member of the Domain Admins group.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **command prompt**, right-click **Command Prompt**, and then click **Run as administrator** to open an elevated command prompt.
4. When prompted to approve the elevation, click **Yes**.

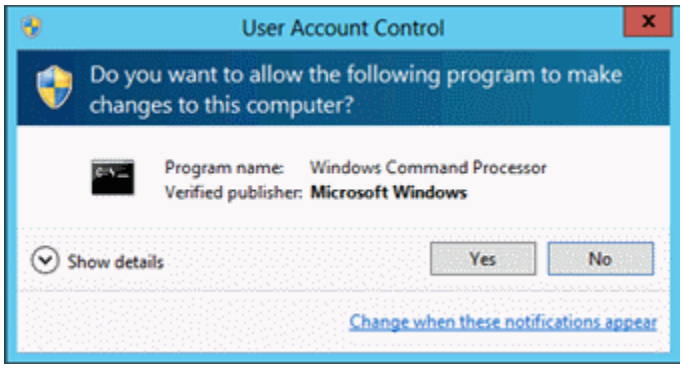

5. In the **Command Prompt** window, type **net use \\<Server Name>\c$**, where <Server Name> is the name of the member server or workstation you're attempting to access over the network.
6. The following screen shot shows the error message that should appear.

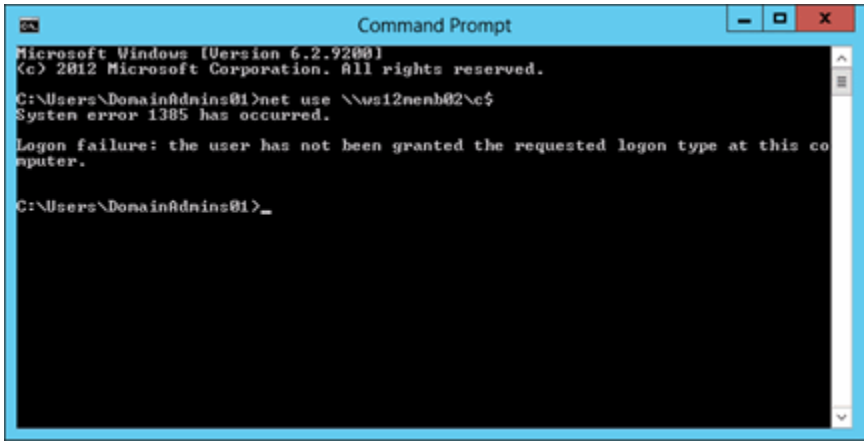

Verify "Deny log on as a batch job" GPO Settings

From any member server or workstation affected by the GPO changes, log on locally.

Create a Batch File

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **notepad**, and click **Notepad**.
3. In **Notepad**, type **dir c:**.
4. Click **File**, and click **Save As**.
5. In the **File** name field, type **<Filename>.bat** (where <Filename> is the name of the new batch file).

Schedule a Task

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **task scheduler**, and click **Task Scheduler**.

   Note
   On computers running Windows 8, in the **Search** box, type **schedule tasks**, and click **Schedule tasks**.
3. In the **Task Scheduler** menu bar, click **Action**, and click **Create Task**.
4. In the **Create Task** dialog box, type **<Task Name>** (where <Task Name> is the name of the new task).
5. Click the **Actions** tab, and click **New**.
6. In the **Action** field, select **Start a program**.
7. Under **Program/script**, click **Browse**, locate and select the batch file created in the **Create a Batch File** section, and click **Open**.
8. Click **OK**.
9. Click the **General** tab.
10. Under **Security** options, click **Change User or Group**.
11. Type the name of an account that is a member of the Domain Admins group, click **Check Names**, and click **OK**.
12. Select **Run whether the user is logged on or not** and select **Do not store password**. The task will only have access to local computer resources.
13. Click **OK**.
14. A dialog box should appear, requesting user account credentials to run the task.
15. After entering the credentials, click **OK**.
16. A dialog box similar to the following should appear.

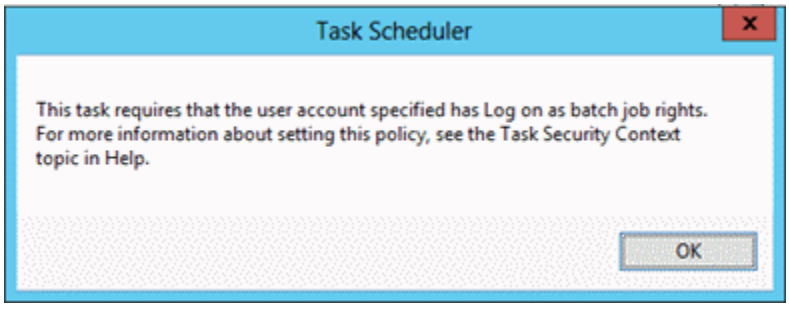

Verify “Deny log on as a service” GPO Settings

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen.

When the **Charms** bar appears, click **Search**.

3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as**, select the **This account** option.
7. Click **Browse**, type the name of an account that is a member of the Domain Admins group, click **Check Names**, and click **OK**.
8. Under **Password** and **Confirm password**, type the selected account's password, and click **OK**.
9. Click **OK** three more times.
10. Right-click **Print Spooler** and click **Restart**.
11. When the service is restarted, a dialog box similar to the following should appear.

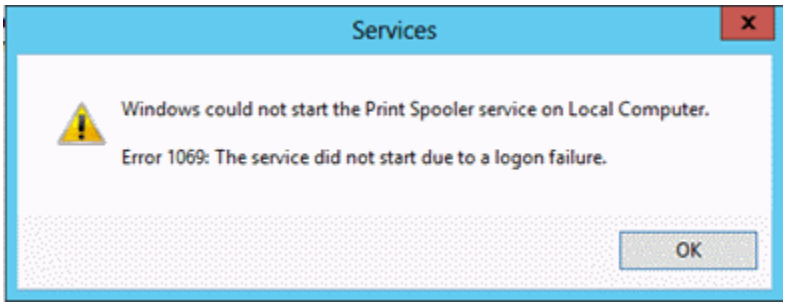

Revert Changes to the Printer Spooler Service

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. Under **Log on as**, select the **Local System** account, and click **OK**.

Verify "Deny log on locally" GPO Settings

1. From any member server or workstation affected by the GPO changes, attempt to log on locally using an account that is a member of the Domain Admins group. A dialog box similar to the following should appear.

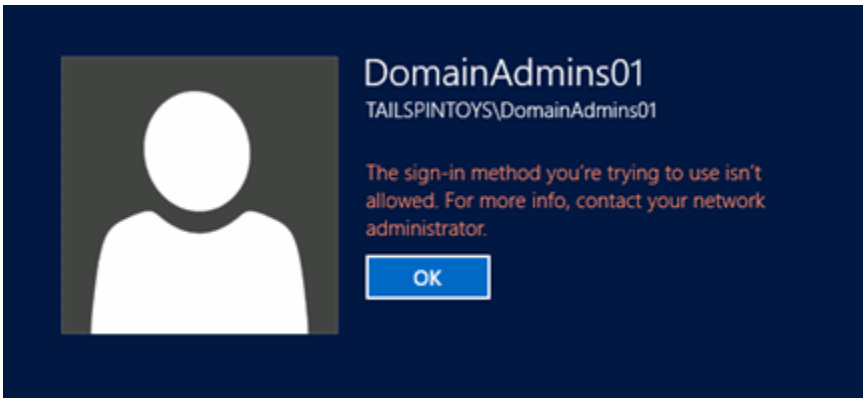

Verify "Deny log on through Remote Desktop Services" GPO Settings

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **remote desktop connection**, and click **Remote Desktop Connection**.
3. In the **Computer** field, type the name of the computer that you want to connect to, and click **Connect**. (You can also type the IP address instead of the computer name.)
4. When prompted, provide credentials for an account that is a member of the Domain Admins group.
5. A dialog box similar to the following should appear.

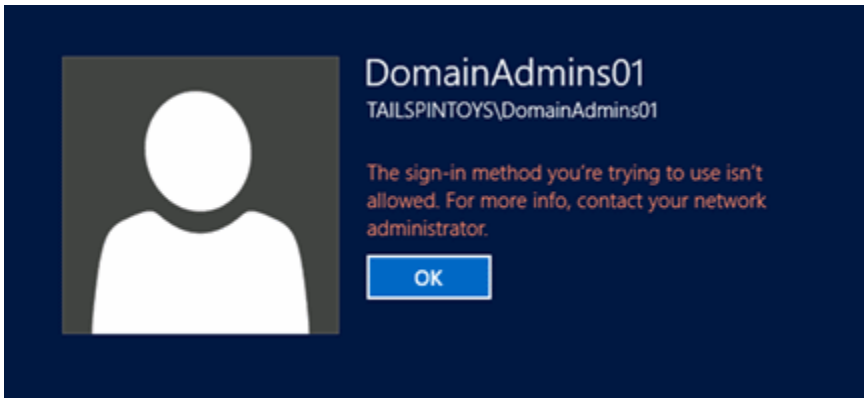

Appendix G: Securing Administrators Groups in Active Directory

Appendix G: Securing Administrators Groups in Active Directory

As is the case with the Enterprise Admins (EA) and Domain Admins (DA) groups, membership in the built-in Administrators (BA) group should be required only in build or disaster recovery scenarios. There should be no day-to-day user accounts in the Administrators group with the exception of the Built-in Administrator account for the domain, if it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

Administrators are, by default, the owners of most of the AD DS objects in their respective domains. Membership in this group may be required in build or disaster recovery scenarios in which ownership or the ability to take ownership of objects is required. Additionally, DAs and EAs inherit a number of their rights and permissions by virtue of their default membership in the Administrators group. Default group nesting for privileged groups in Active Directory should not be modified, and each domain's Administrators group should be secured as described in the step-by-step instructions that follow.

For the Administrators group in each domain in the forest:

1. Remove all members from the Administrators group, with the possible exception of the built-in Administrator account for the domain, provided it has been secured as described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.
2. In GPOs linked to OUs containing member servers and workstations in each domain, the DA group should be added to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\ User Rights Assignment**:
   - Deny access to this computer from the network
   - Deny log on as a batch job
   - Deny log on as a service
3. At the domain controllers OU in each domain in the forest, the Administrators group should be granted the following user rights:
   - Access this computer from the network
   - Allow log on locally
   - Allow log on through Remote Desktop Services
4. Auditing should be configured to send alerts if any modifications are made to the properties or membership of the Administrators group.

Step-by-Step Instructions for Removing All Members from the Administrators Group

1. In **Server Manager**, click **Tools**, and click **Active Directory Users and Computers**.
2. To remove all members from the Administrators group, perform the following steps:

    a. Double-click the **Administrators** group and click the **Members** tab.

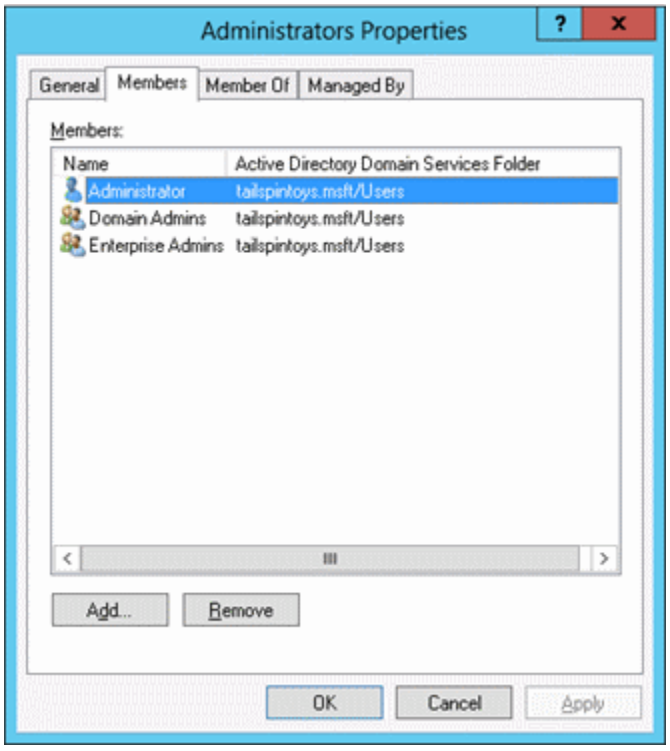

    b. Select a member of the group, click **Remove**, click **Yes**, and click **OK**.

3. Repeat step 2 until all members of the Administrators group have been removed.

Step-by-Step Instructions to Secure Administrators Groups in Active Directory

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.
2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the

domain where you want to set the Group Policy).

3. In the console tree, right-click **Group Policy Objects**, and click **New**.

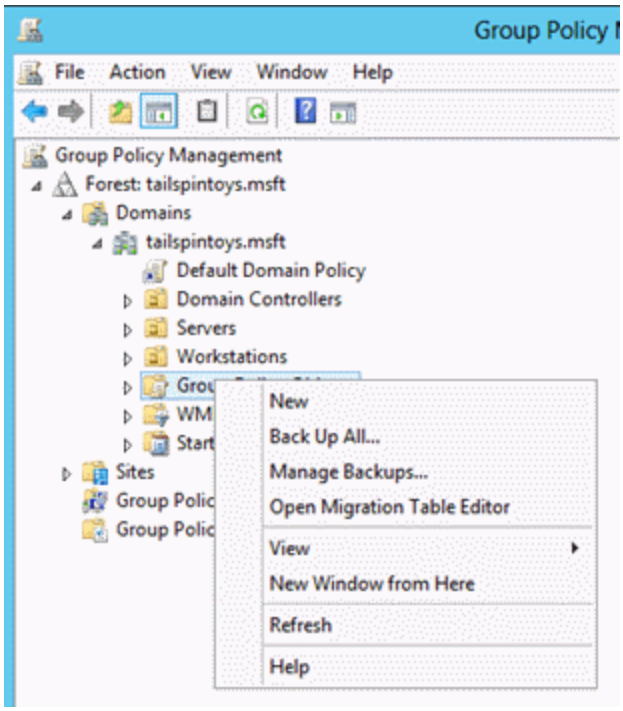

4. In the **New GPO** dialog box, type <GPO Name>, and click **OK** (where *GPO Name* is the name of this GPO).

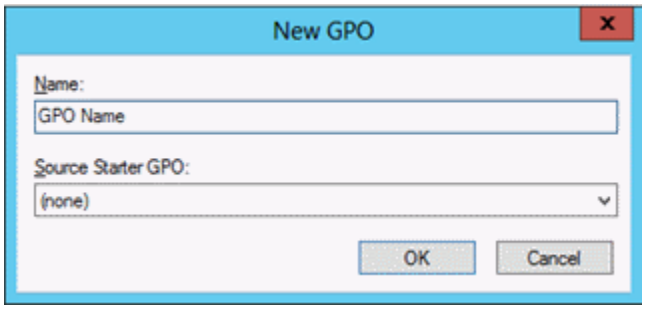

5. In the details pane, right-click **<GPO Name>**, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

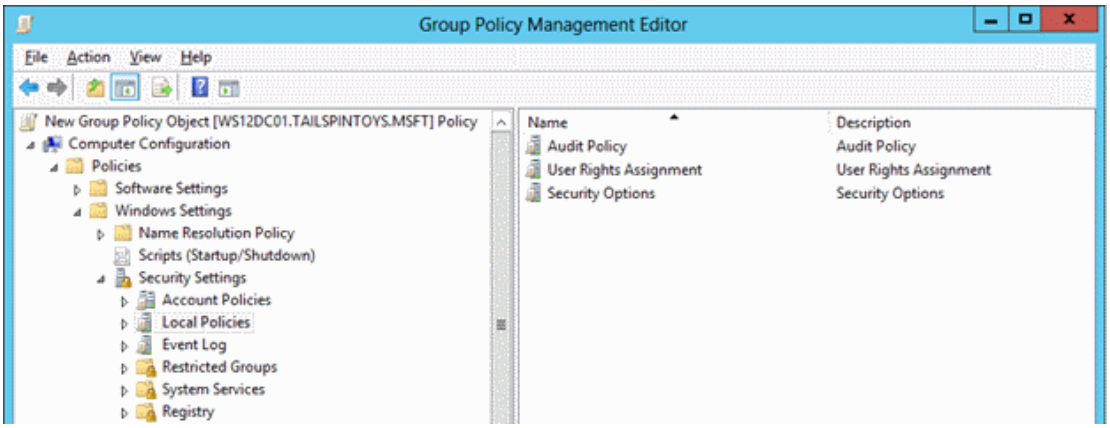

7. Configure the user rights to prevent members of the Administrators group from accessing member servers and workstations over the network by doing the following:

   a. Double-click **Deny access to this computer from the network** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Administrators**, click **Check Names**, and click **OK**.

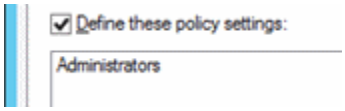

   d. Click **OK**, and **OK** again.

8. Configure the user rights to prevent members of the Administrators group from logging on as a batch job by doing the following:

   a. Double-click **Deny log on as a batch job** and select **Define these policy settings**.
   b. Click **Add User or Group** and click **Browse**.
   c. Type **Administrators**, click **Check Names**, and click **OK**.

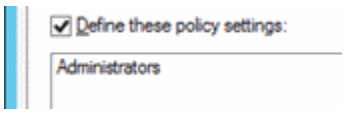

   d. Click **OK**, and **OK** again.

9. Configure the user rights to prevent members of the Administrators group from logging on as a service by doing the following:

a. Double-click **Deny log on as a service** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Administrators**, click **Check Names**, and click **OK**.

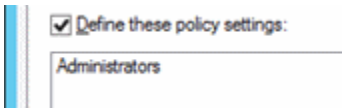

d. Click **OK**, and **OK** again.

10. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
11. In **Group Policy Management**, link the GPO to the member server and workstation OUs by doing the following:

a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
b. Right-click the OU that the GPO will be applied to and click **Link an existing GPO**.

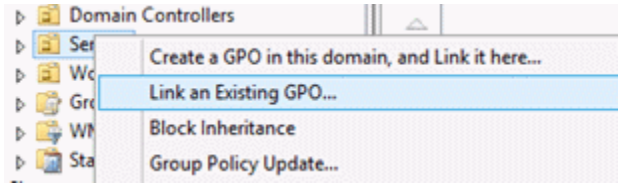

c. Select the GPO that you just created and click **OK**.

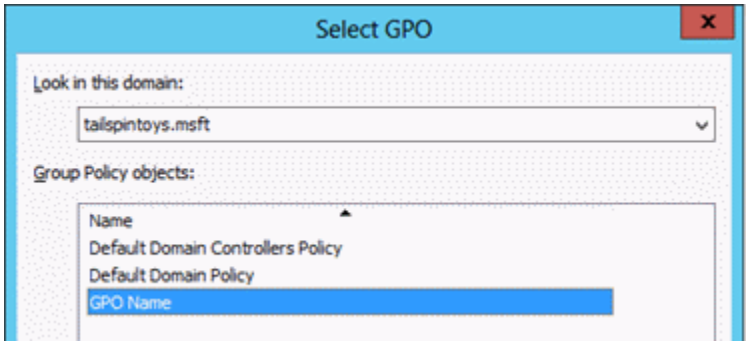

d. Create links to all other OUs that contain workstations.
e. Create links to all other OUs that contain member servers.

Important

If jump servers are used to administer domain controllers and Active Directory, ensure that jump servers are located in an OU to which this GPOs is not linked.

Notes

When you implement restrictions on the Administrators group in GPOs, Windows applies the settings to members of a computer's local Administrators group in addition to the domain's Administrators group. Therefore, you should use caution when implementing restrictions in the Administrators group. Although prohibiting network, batch, and service logons for members of the Administrators group is advised wherever it is feasible to implement, do not restrict local logons or logons through Remote Desktop Services. Blocking these logon types can block legitimate administration of a computer by members of the local Administrators group.

The following screen shot shows configuration settings that block misuse of built-in local and domain Administrator accounts, in addition to misuse of built-in local or domain Administrators groups. Note that the **Deny log on through Remote Desktop Services** user right does not include the Administrators group, because including it in this setting would also block these logons for accounts that are members of the local computer's Administrators group. If services on computers are configured to run in the context of any of the privileged groups described in this section, implementing these settings can cause services and applications to fail. Therefore, as with all of the recommendations in this section, you should thoroughly test settings for applicability in your environment.

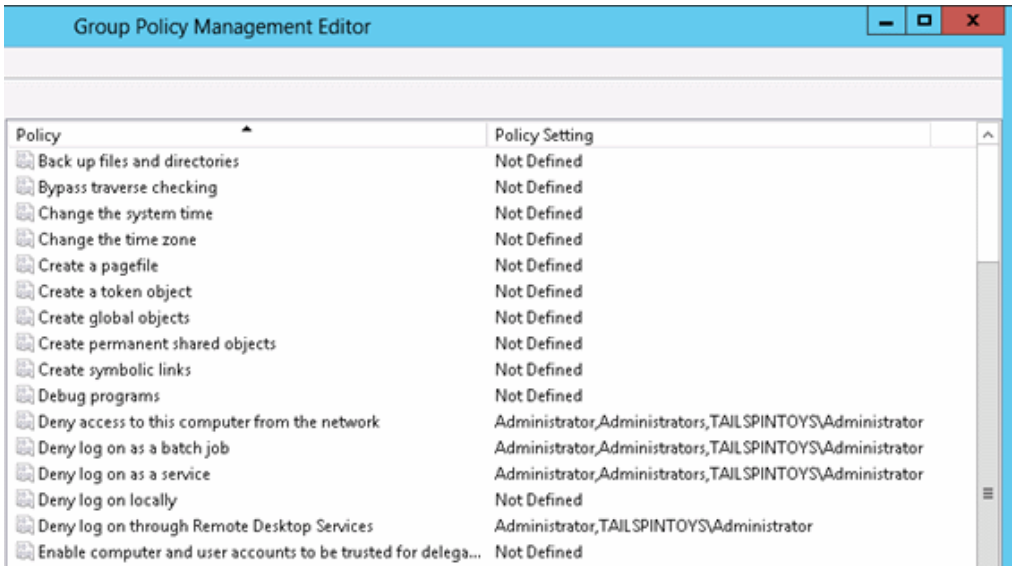

Step-by-Step Instructions to Grant User Rights to the Administrators Group

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.
2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
3. In the console tree, right-click **Group Policy Objects**, and click **New**.

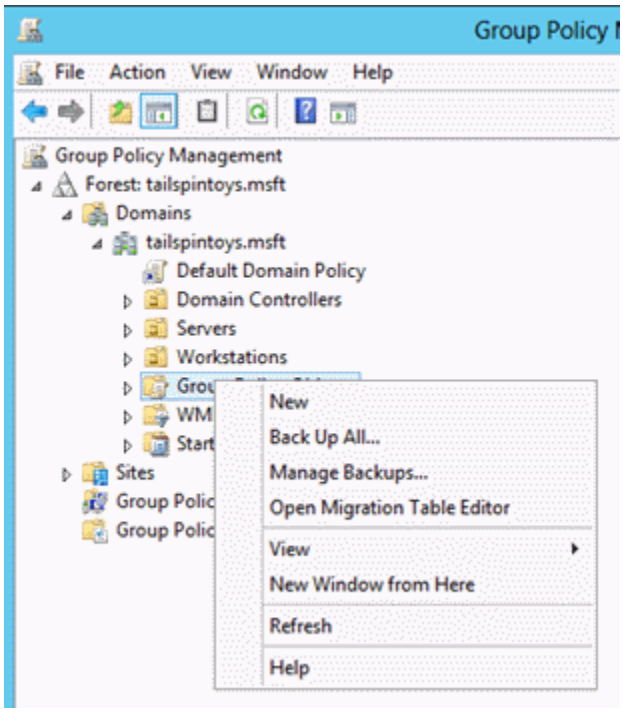

4. In the **New GPO** dialog box, type <GPO Name>, and click **OK** (where <GPO Name> is the name of this GPO).

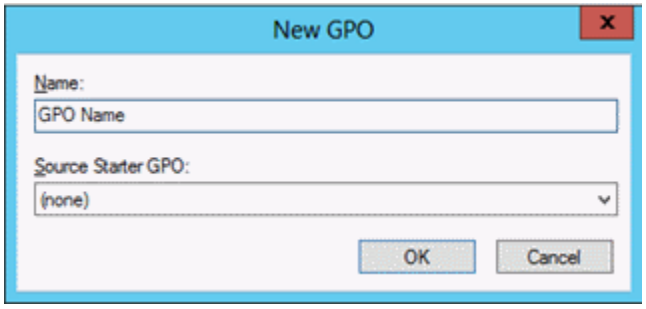

5. In the details pane, right-click **<GPO Name>**, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

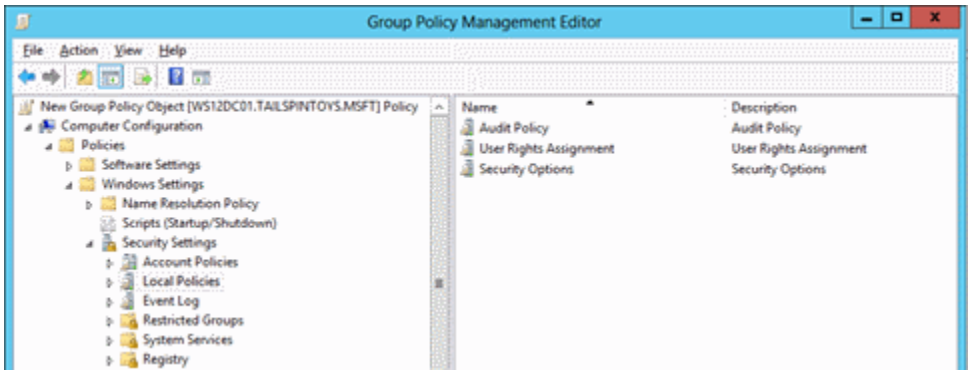

7. Configure the user rights to allow members of the Administrators group to access domain controllers over the network by doing the following:

a. Double-click **Access to this computer from the network** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Click **Add User or Group** and click **Browse**.

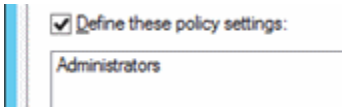

d. Click **OK**, and **OK** again.

8. Configure the user rights to allow members of the Administrators group to log on locally by doing the following:

a. Double-click **Allow log on locally** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Administrators**, click Check **Names**, and click **OK**.

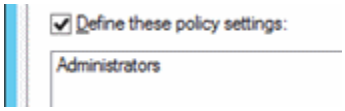

d. Click **OK**, and **OK** again.

9. Configure the user rights to allow members of the Administrators group to log on through Remote Desktop Services by doing the following:

a. Double-click **Allow log on through Remote Desktop Services** and select **Define these policy settings**.
b. Click **Add User or Group** and click **Browse**.
c. Type **Administrators**, click **Check Names**, and click **OK**.

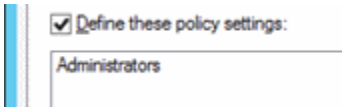

d. Click **OK**, and **OK** again.

10. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
11. In **Group Policy Management**, link the GPO to the domain controllers OU by doing the following:

    a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
    b. Right-click the domain controllers OU and click **Link an existing GPO**.

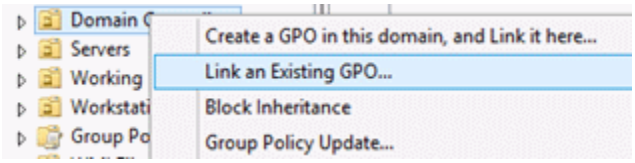

    c. Select the GPO that you just created and click **OK**.

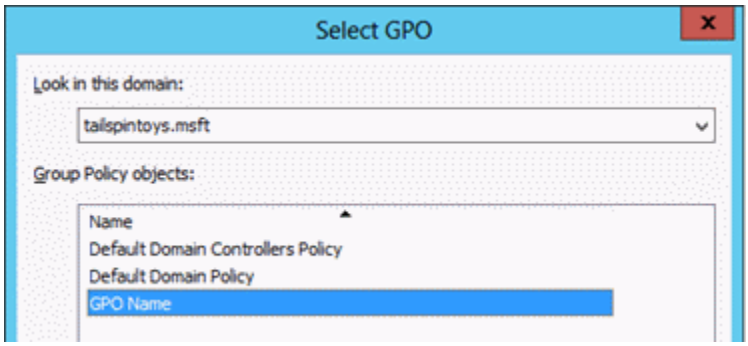

Verification Steps

Verify "Deny access to this computer from the network" GPO Settings

From any member server or workstation that is not affected by the GPO changes (such as a "jump server"), attempt to access a member server or workstation over the network that is affected by the GPO changes. To verify the GPO settings, attempt to map the system drive by using the **NET USE** command.

1. Log on locally using an account that is a member of the Administrators group.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **command prompt**, right-click **Command Prompt**, and then click **Run as administrator** to open an elevated command prompt.

4. When prompted to approve the elevation, click **Yes**.

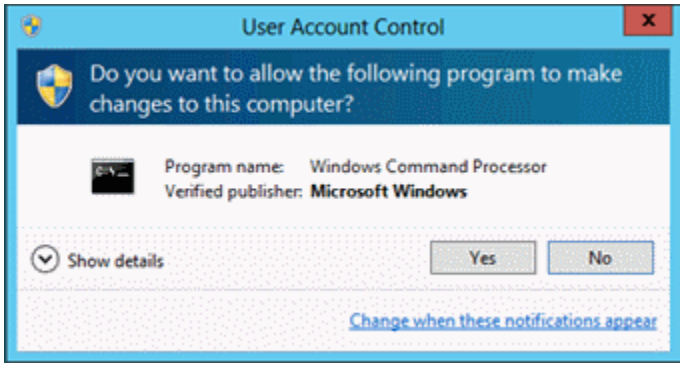

5. In the **Command Prompt** window, type **net use \\<Server Name>\c$**, where <Server Name> is the name of the member server or workstation you're attempting to access over the network.
6. The following screen shot shows the error message that should appear.

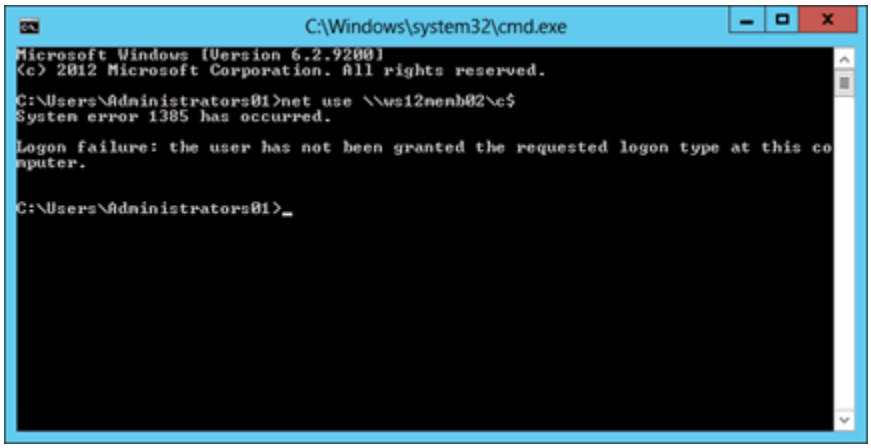

Verify "Deny log on as a batch job" GPO Settings

From any member server or workstation affected by the GPO changes, log on locally.

Create a Batch File

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **notepad**, and click **Notepad**.
3. In **Notepad**, type **dir c:**.
4. Click **File**, and click **Save As**.
5. In the **File name** field, type **<Filename>.bat** (where <Filename> is the name of the new batch file).

Schedule a Task

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **task scheduler**, and click **Task Scheduler**.

   Note
   On computers running Windows 8, in the Search box, type schedule tasks, and click Schedule tasks.
3. Click **Action**, and click **Create Task**.
4. In the **Create Task** dialog box, type **<Task Name>** (where <Task Name> is the name of the new task).
5. Click the **Actions** tab, and click **New**.
6. In the **Action** field, select **Start a program**.
7. In the **Program/script** field, click **Browse**, locate and select the batch file created in the **Create a Batch File** section, and click **Open**.
8. Click **OK**.
9. Click the **General** tab.
10. In the **Security options** field, click **Change User or Group**.
11. Type the name of an account that is a member of the Administrators group, click **Check Names**, and click **OK**.
12. Select **Run whether the user is logged onor not** and **Do not store password**. The task will only have access to local computer resources.
13. Click **OK**.
14. A dialog box should appear, requesting user account credentials to run the task.
15. After entering the password, click **OK**.
16. A dialog box similar to the following should appear.

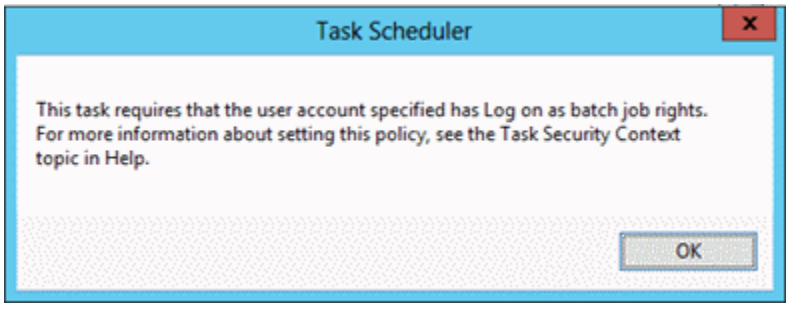

Verify “Deny log on as a service” GPO Settings

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen.

When the **Charms** bar appears, click **Search**.

3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. In the **Log on as** field, select **This account**.
7. Click **Browse**, type the name of an account that is a member of the Administrators group, click **Check Names**, and click **OK**.
8. In the **Password** and **Confirm password** fields, type the selected account's password, and click **OK**.
9. Click **OK** three more times.
10. Right-click **Print Spooler** and click **Restart**.
11. When the service is restarted, a dialog box similar to the following should appear.

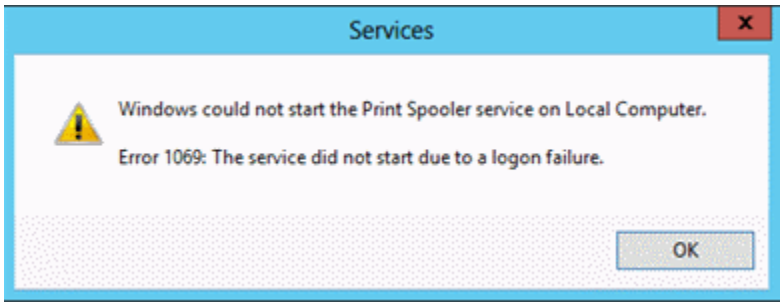

Revert Changes to the Printer Spooler Service

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. In the **Log on as** field, click **Local System** account, and click **OK**.

Appendix H: Securing Local Administrator Accounts and Groups

Appendix H: Securing Local Administrator Accounts and Groups

On all versions of Windows currently in mainstream support, the local Administrator account is disabled by default, which makes the account unusable for pass-the-hash and other credential theft attacks. However, in environments that contain legacy operating systems or in which local Administrator accounts have been enabled, these accounts can be used as previously described to propagate compromise across member servers and workstations. Each local Administrator account and group should be secured as described in the step-by-step instructions that follow.

For detailed information about considerations in securing Built-in Administrator (BA) groups, see Implementing Least-Privilege Administrative Models.

Controls for Local Administrator Accounts

For the local Administrator account in each domain in your forest, you should configure the following settings:

- Configure GPOs to restrict the domain's Administrator account's use on domain-joined systems
  - In one or more GPOs that you create and link to workstation and member server OUs in each domain, add the Administrator account to the following user rights in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Settings\User Rights Assignments**:
    - Deny access to this computer from the network
    - Deny log on as a batch job
    - Deny log on as a service
    - Deny log on through Remote Desktop Services

Step-by-Step Instructions to Secure Local Administrators Groups

Configuring GPOs to Restrict Administrator Account on Domain-Joined Systems

1. In **Server Manager**, click **Tools**, and click **Group Policy Management**.
2. In the console tree, expand <Forest>\Domains\<Domain>, and then **Group Policy Objects** (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
3. In the console tree, right-click **Group Policy Objects**, and click **New**.

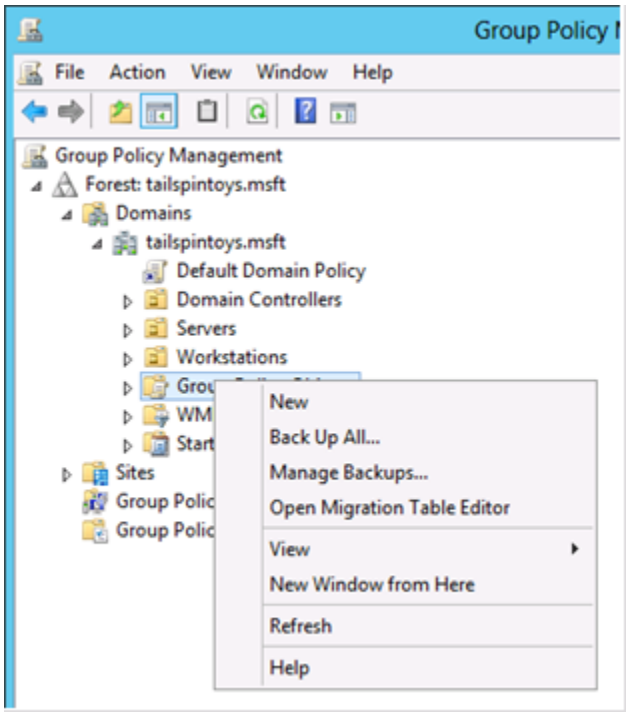

4. In the **New GPO** dialog box, type **<GPO Name>**, and click **OK** (where <GPO Name> is the name of this GPO).

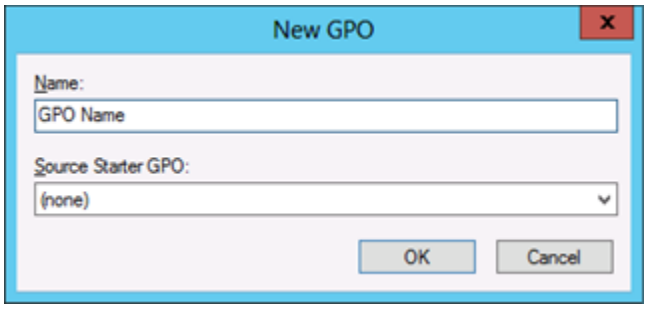

5. In the details pane, right-click **<GPO Name>**, and click **Edit**.
6. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**, and click **User Rights Assignment**.

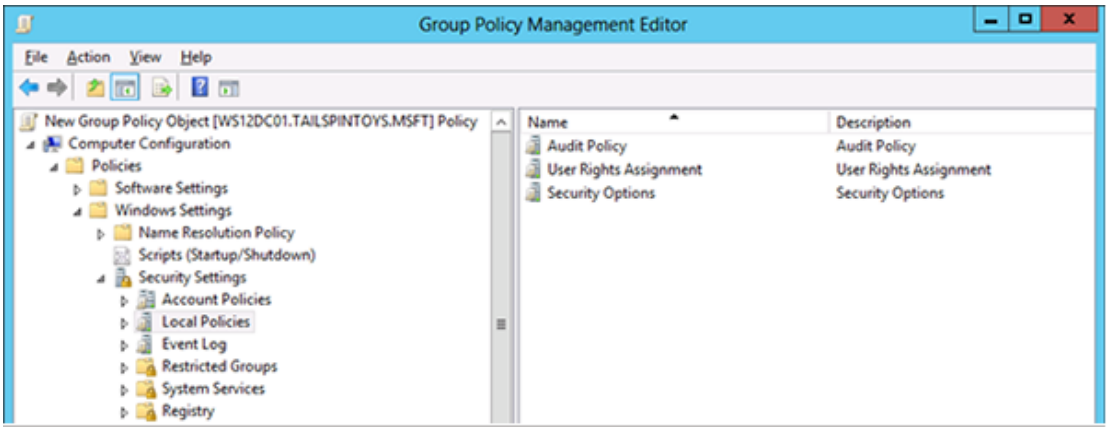

7. Configure the user rights to prevent the local Administrator account from accessing members servers and workstations over the network by doing the following:

   a. Double-click **Deny access to this computer from the network** and select **Define these policy settings**.

   b. Click **Add User or Group**, type the user name of the local Administrator account, and click **OK**. This user name will be **Administrator**, the default when Windows is installed.

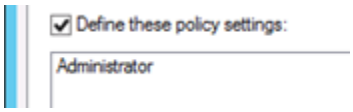

   c. Click OK.

Important

When you add the Administrator account to these settings, you specify whether you are configuring a local Administrator account or a domain Administrator account by how you label the accounts. For example, to add the TAILSPINTOYS domain's Administrator account to these deny rights, you would browse to the Administrator account for the TAILSPINTOYS domain, which would appear as TAILSPINTOYS\Administrator. If you type **Administrator** in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied, as described earlier.

8. Configure the user rights to prevent the local Administrator account from logging on as a batch job by doing the following:

a. Double-click **Deny log on as a batch job** and select **Define these policy settings**.

b. Click **Add User or Group**, type the user name of the local Administrator account, and click **OK**. This user name will be **Administrator**, the default when Windows is installed.

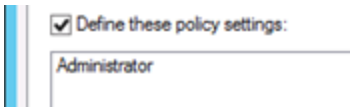

c. Click **OK**.

Important
When you add the Administrator account to these settings, you specify whether you are configuring local Administrator account or domain Administrator account by how you label the accounts. For example, to add the TAILSPINTOYS domain's Administrator account to these deny rights, you would browse to the Administrator account for the TAILSPINTOYS domain, which would appear as TAILSPINTOYS\Administrator. If you type **Administrator** in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied, as described earlier.

9. Configure the user rights to prevent the local Administrator account from logging on as a service by doing the following:

a. Double-click **Deny log on as a service** and select **Define these policy settings**.

b. Click **Add User or Group**, type the user name of the local Administrator account, and click **OK**. This user name will be **Administrator**, the default when Windows is installed.

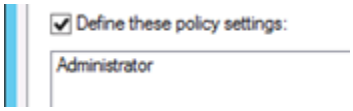

c. Click **OK**.

Important
When you add the Administrator account to these settings, you specify whether you are configuring local Administrator account or domain Administrator account by how you label the accounts. For example, to add the TAILSPINTOYS domain's Administrator account to these deny rights, you would browse to the Administrator account

for the TAILSPINTOYS domain, which would appear as TAILSPINTOYS\Administrator. If you type **Administrator** in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied, as described earlier.

10. Configure the user rights to prevent the local Administrator account from accessing member servers and workstations via Remote Desktop Services by doing the following:

    a. Double-click **Deny log on through Remote Desktop Services** and select **Define these policy settings**.
    b. Click **Add User or Group**, type the user name of the local Administrator account, and click **OK**. This user name will be **Administrator**, the default when Windows is installed.

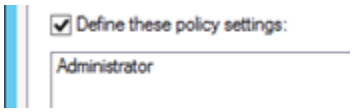

    c. Click **OK**.

Important
When you add the Administrator account to these settings, you specify whether you are configuring local Administrator account or domain Administrator account by how you label the accounts. For example, to add the TAILSPINTOYS domain's Administrator account to these deny rights, you would browse to the Administrator account for the TAILSPINTOYS domain, which would appear as TAILSPINTOYS\Administrator. If you type **Administrator** in these user rights settings in the Group Policy Object Editor, you will restrict the local Administrator account on each computer to which the GPO is applied, as described earlier.

11. To exit **Group Policy Management Editor**, click **File**, and click **Exit**.
12. In **Group Policy Management**, link the GPO to the member server and workstation OUs by doing the following:

    a. Navigate to the <Forest>\Domains\<Domain> (where <Forest> is the name of the forest and <Domain> is the name of the domain where you want to set the Group Policy).
    b. Right-click the OU that the GPO will be applied to and click **Link an existing**

GPO.

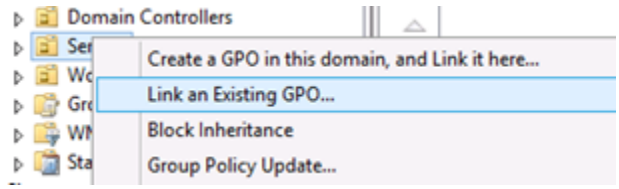

c. Select the GPO that you created and click **OK**.

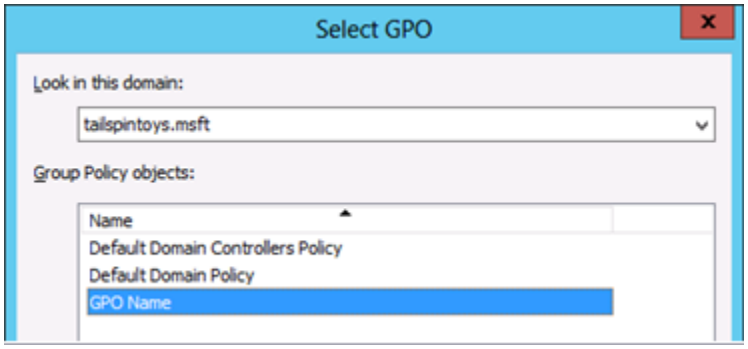

d. Create links to all other OUs that contain workstations.

e. Create links to all other OUs that contain member servers.

Verification Steps

Verify "Deny access to this computer from the network" GPO Settings

From any member server or workstation that is not affected by the GPO changes (such as a jump server), attempt to access a member server or workstation over the network that is affected by the GPO changes. To verify the GPO settings, attempt to map the system drive by using the **NET USE** command.

1. Log on locally to any member server or workstation that is not affected by the GPO changes.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **command prompt**, right-click **Command Prompt**, and then click **Run as administrator** to open an elevated command prompt.
4. When prompted to approve the elevation, click **Yes**.

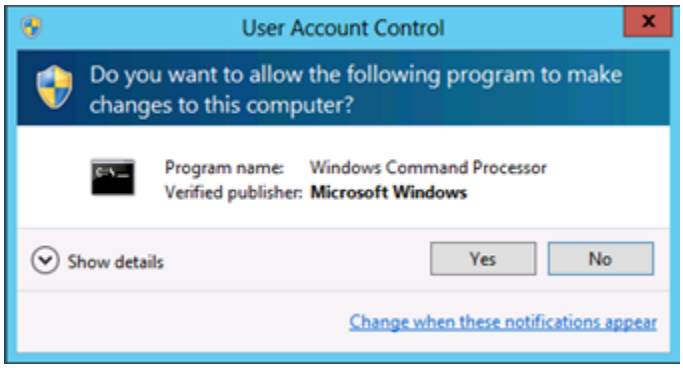

5. In the **Command Prompt** window, type **net use \\<Server Name>\c$ /user:<Server Name>\Administrator**, where <Server Name> is the name of the member server or workstation you're attempting to access over the network.

   Note
   The local Administrator credentials must be from the same system you're attempting to access over the network.

6. The following screenshot shows the error message that should appear.

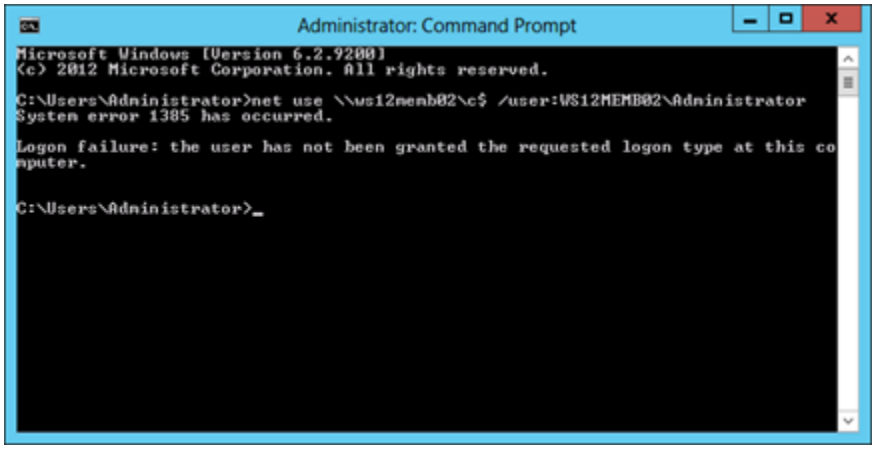

Verify "Deny log on as a batch job" GPO Settings

From any member server or workstation affected by the GPO changes, log on locally.

Create a Batch File

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type **notepad**, and click **Notepad**.
3. In **Notepad**, type **dir c:**.
4. Click **File**, and click **Save As**.

5. In the **File name** box, type **<Filename>.bat** (where <Filename> is the name of the new batch file).

Schedule a Task

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
2. In the **Search** box, type task scheduler, and click **Task Scheduler**.

   Note
   On computers running Windows 8, in the **Search** box, type **schedule tasks**, and click **Schedule tasks**.
3. Click **Action**, and click **Create Task**.
4. In the **Create Task** dialog box, type **<Task Name>** (where <Task Name> is the name of the new task).
5. Click the **Actions** tab, and click **New**.
6. In the **Action** field, click **Start a program**.
7. In the **Program/script** field, click **Browse**, locate and select the batch file created in the **Create a Batch File** section, and click **Open**.
8. Click **OK**.
9. Click the **General** tab.
10. In the **Security options** field, click **Change User or Group**.
11. Type the name of the system's local Administrator account, click **Check Names**, and click **OK**.
12. Select **Run whether the user is logged on or not** and **Do not store password**. The task will only have access to local computer resources.
13. Click **OK**.
14. A dialog box should appear, requesting user account credentials to run the task.
15. After entering the credentials, click **OK**.
16. A dialog box similar to the following should appear.

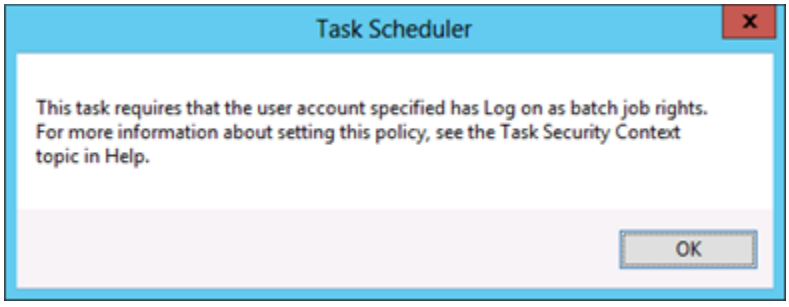

Verify "Deny log on as a service" GPO Settings

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. In **Log on as** field, click **This account**.
7. Click **Browse**, type the system's local Administrator account, click **Check Names**, and click **OK**.
8. In the **Password** and **Confirm password** fields, type the selected account's password, and click **OK**.
9. Click **OK** three more times.
10. Right-click **Print Spooler** and click **Restart**.
11. When the service is restarted, a dialog box similar to the following should appear.

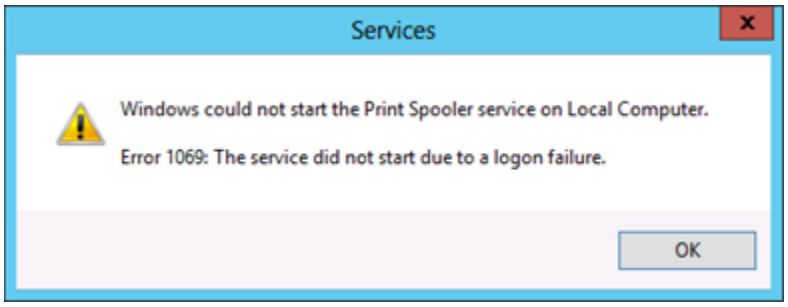

Revert Changes to the Printer Spooler Service

1. From any member server or workstation affected by the GPO changes, log on locally.
2. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.
3. In the **Search** box, type **services**, and click **Services**.
4. Locate and double-click **Print Spooler**.
5. Click the **Log On** tab.
6. In the **Log on as**: field, select **Local Systemaccount**, and click **OK**.

Verify "Deny log on through Remote Desktop Services" GPO Settings

1. With the mouse, move the pointer into the upper-right or lower-right corner of the screen. When the **Charms** bar appears, click **Search**.

2. In the **Search** box, type **remote desktop connection**, and click **Remote Desktop Connection**.
3. In the **Computer** field, type the name of the computer that you want to connect to, and click **Connect**. (You can also type the IP address instead of the computer name.)
4. When prompted, provide credentials for the system's local **Administrator** account.
5. A dialog box similar to the following should appear.

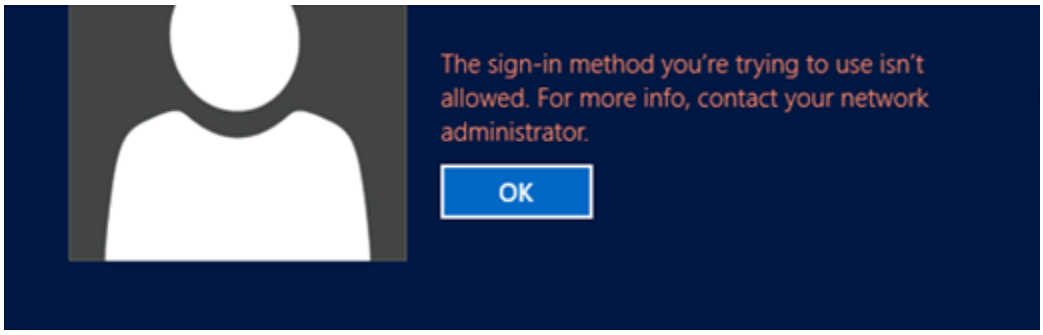

Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory

Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory

One of the challenges in implementing an Active Directory model that does not rely on permanent membership in highly privileged groups is that there must be a mechanism to populate these groups when temporary membership in the groups is required. Some privileged identity management solutions require that the software's service accounts are granted permanent membership in groups such as DA or Administrators in each domain in the forest. However, it is technically not necessary for Privileged Identity Management (PIM) solutions to run their services in such highly privileged contexts.

This appendix provides information that you can use for natively implemented or third-party PIM solutions to create accounts that have limited privileges and can be stringently controlled, but can be used to populate privileged groups in Active Directory when temporary elevation is required. If you are implementing PIM as a native solution, these accounts may be used by administrative staff to perform the temporary group population, and if you're implementing PIM via third-party software, you might be able to adapt these accounts to function as service accounts.

Note

The procedures described in this appendix provide one approach to the management of highly privileged groups in Active Directory. You can adapt these procedures to suit your needs, add additional restrictions, or omit some of the restrictions that are described here.

Creating Management Accounts for Protected Accounts and Groups in Active Directory

Creating accounts that can be used to manage the membership of privileged groups without requiring the management accounts to be granted excessive rights and permissions consists of four general activities that are described in the step-by-step instructions that follow:

1. First, you should create a group that will manage the accounts, because these accounts should be managed by a limited set of trusted users. If you do not already have an OU structure that accommodates segregating privileged and protected accounts and systems from the general population in the domain, you should create one. Although specific instructions are not provided in this appendix, screenshots show an example of such an OU hierarchy.
2. Create the management accounts. These accounts should be created as “regular” user accounts and granted no user rights beyond those that are already granted to users by default.
3. Implement restrictions on the management accounts that make them usable only for the specialized purpose for which they were created, in addition to controlling who can enable and use the accounts (the group you created in the first step).
4. Configure permissions on the AdminSDHolder object in each domain to allow the management accounts to change the membership of the privileged groups in the domain.

You should thoroughly test all of these procedures and modify them as needed for your environment before implementing them in a production environment. You should also verify that all settings work as expected (some testing procedures are provided in this appendix), and you should test a disaster recovery scenario in which the management accounts are not available to be used to populate protected groups for recovery purposes. For more information about backing up and restoring Active Directory, see the AD DS Backup and Recovery Step-by-Step Guide.

Note

By implementing the steps described in this appendix, you will create accounts that will be able to manage the membership of all protected groups in each domain, not only the highest-privilege Active Directory groups like EAs, DAs and BAs. For more information about protected groups in Active Directory, see Appendix C: Protected Accounts and Groups in Active Directory.

Step-by-Step Instructions for Creating Management Accounts for Protected Groups

Creating a Group to Enable and Disable Management Accounts

Management accounts should have their passwords reset at each use and should be disabled when activities requiring them are complete. Although you might also consider implementing smart card logon requirements for these accounts, it is an optional configuration and these instructions assume that the management accounts will be configured with a user name and long, complex password as minimum controls. In this step, you will create a group that has permissions to reset password on the management accounts and to enable and disable the accounts.

To create a group to enable and disable management accounts, perform the following steps:

1. In the OU structure where you will be housing the management accounts, right-click the OU where you want to create the group, click **New** and click **Group**.

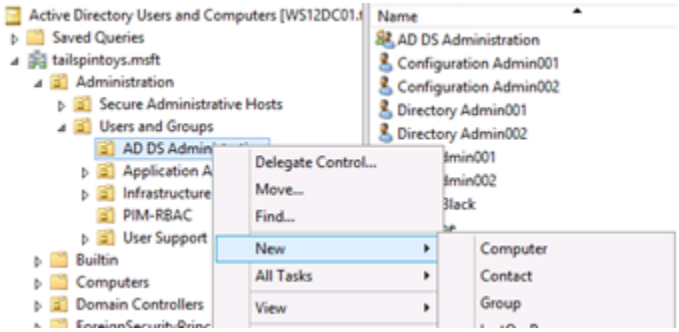

2. In the **New Object – Group** dialog box, enter a name for the group. If you plan to use this group to “activate” all management accounts in your forest, make it a universal security group. If you have a single-domain forest or if you plan to create a group in each domain, you can create a global security group. Click **OK** to create the group.

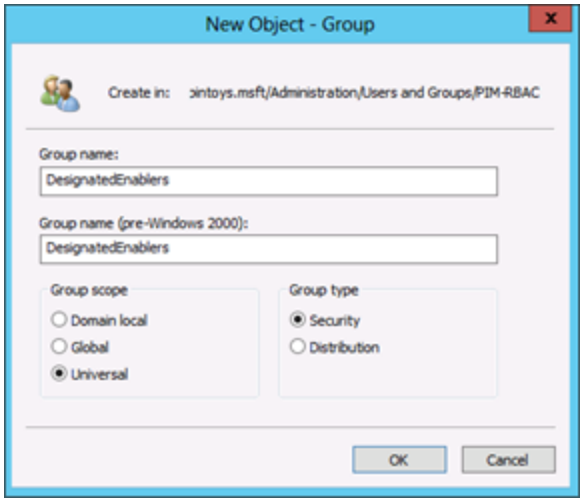

3. Right-click the group you just created, click **Properties**, and click the **Object** tab. In the group's **Object property** dialog box, select **Protect object from accidental deletion**, which will not only prevent otherwise-authorized users from deleting the group, but also from moving it to another OU unless the attribute is first deselected.

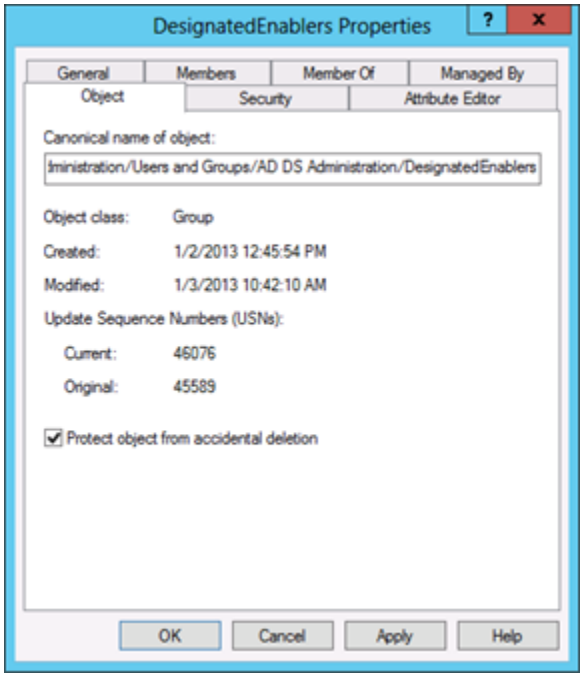

Note

If you have already configured permissions on the group's parent OUs to restrict administration to a limited set of users, you may not need to perform the following steps. They are provided here so that even if you have not yet implemented limited administrative control over the OU structure in which you've created this group, you can secure the group against modification by unauthorized users.

4. Click the **Members** tab, and add the accounts for members of your team who will be responsible for enabling management accounts or populating protected groups when necessary.

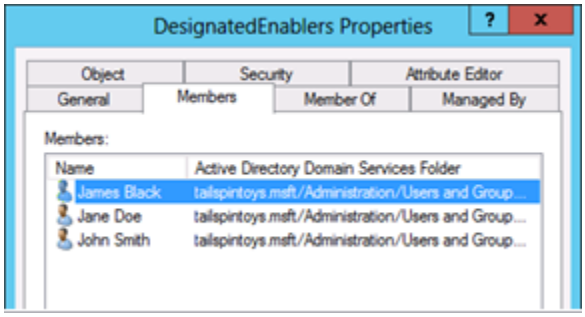

5. If you have not already done so, in the **Active Directory Users and Computers** console, click **View** and select **Advanced Features**. Right-click the group you just created, click

Properties, and click the **Security** tab. On the **Security** tab, click **Advanced**.

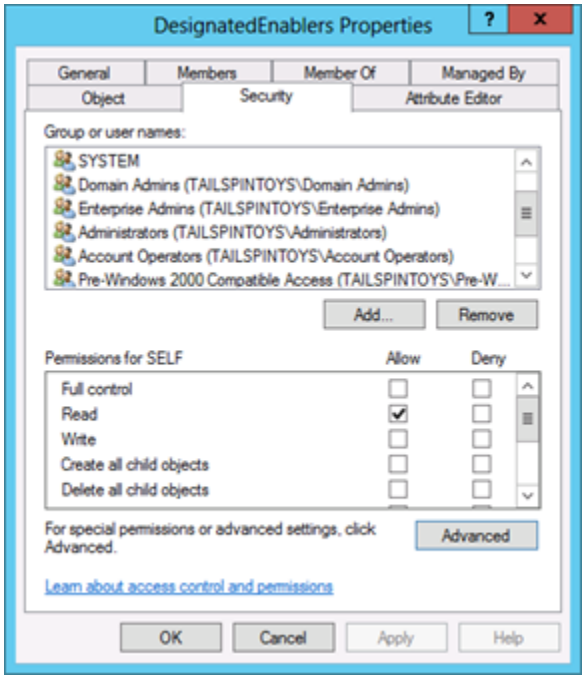

6. In the **Advanced Security Settings for [Group]** dialog box, click **Disable Inheritance**. When prompted, click **Convert inherited permissions into explicit permissions on this object**, and click **OK** to return to the group's **Security** dialog box.

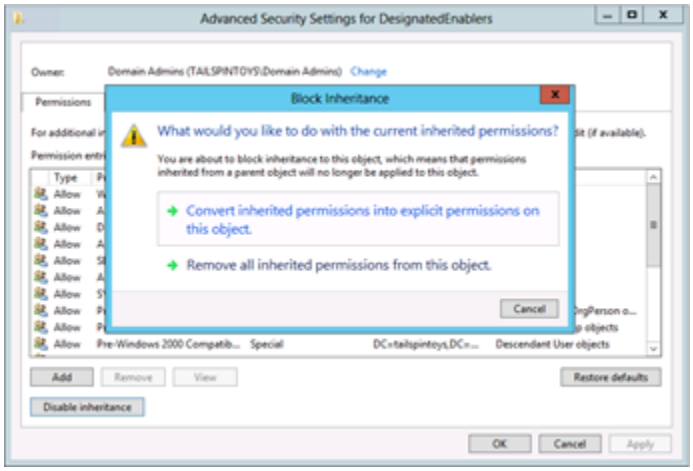

7. On the **Security** tab, remove groups that should not be permitted to access this group. For example, if you do not want Authenticated Users to be able to read the group's name and general properties, you can remove that ACE. You can also remove ACEs, such as those for account operators and pre-Windows 2000 Server compatible access. You should, however, leave a minimum set of object permissions in place. Leave the following ACEs intact:
   - SELF

- SYSTEM
- Domain Admins
- Enterprise Admins
- Administrators
- Windows Authorization Access Group (if applicable)
- ENTERPRISE DOMAIN CONTROLLERS

Although it may seem counterintuitive to allow the highest privileged groups in Active Directory to manage this group, your goal in implementing these settings is not to prevent members of those groups from making authorized changes. Rather, the goal is to ensure that when you have occasion to require very high levels of privilege, authorized changes will succeed. It is for this reason that changing default privileged group nesting, rights, and permissions are discouraged throughout this document. By leaving default structures intact and emptying the membership of the highest privilege groups in the directory, you can create a more secure environment that still functions as expected.

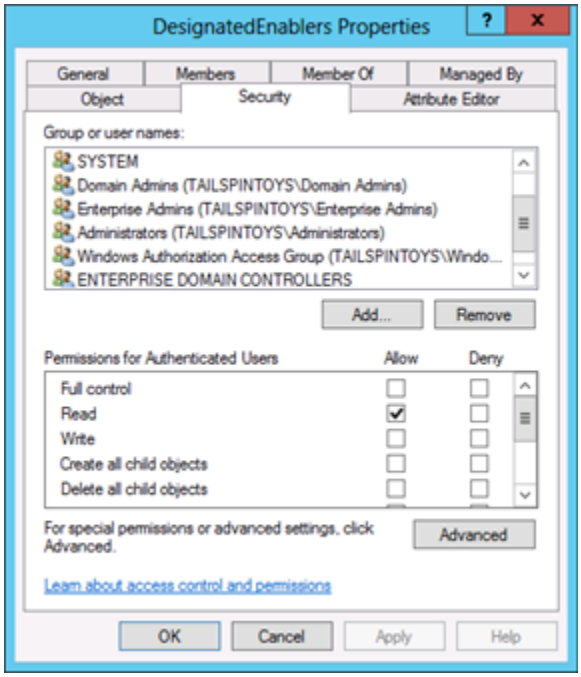

Note

If you have not already configured audit policies for the objects in the OU structure where you created this group, you should configure auditing to log changes this group.

8. You have completed configuration of the group that will be used to “check out” management accounts when they are needed and “check in” the accounts when their activities have been completed.

Creating the Management Accounts

You should create at least one account that will be used to manage the membership of privileged groups in your Active Directory installation, and preferably a second account to serve as a backup. Whether you choose to create the management accounts in a single domain in the forest and grant them management capabilities for all domains' protected groups, or whether you choose to implement management accounts in each domain in the forest, the procedures are effectively the same.

Notes

The steps in this document assume that you have not yet implemented role-based access controls and privileged identity management for Active Directory. Therefore, some procedures must be performed by a user whose account is a member of the Domain Admins group for the domain in question.

When you are using an account with DA privileges, you can log on to a domain controller to perform the configuration activities. Steps that do not require DA privileges can be performed by less-privileged accounts that are logged on to administrative workstations. Screen shots that show dialog boxes bordered in the lighter blue color represent activities that can be performed on a domain controller. Screen shots that show dialog boxes in the darker blue color represent activities that can be performed on administrative workstations with accounts that have limited privileges.

To create the management accounts, perform the following steps:

1. Log on to a domain controller with an account that is a member of the domain's DA group.
2. Launch **Active Directory Users and Computers** and navigate to the OU where you will be creating the management account.
3. Right-click the OU and click **New** and click **User**.
4. In the **New Object – User** dialog box, enter your desired naming information for the account and click **Next**.
5. Log on to a domain controller with an account that is a member of the domain's DA group.
6. Launch **Active Directory Users and Computers** and navigate to the OU where you will be creating the management account.
7. Right-click the OU and click **New** and click **User**.
8. In the **New Object – User** dialog box, enter your desired naming information for the account and click **Next**.

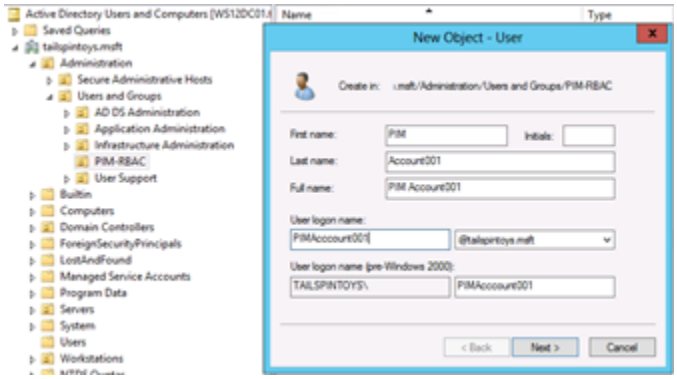

9. Provide an initial password for the user account, clear **User must change password at next logon**, select **User cannot change password** and **Account is disabled**, and click **Next**.

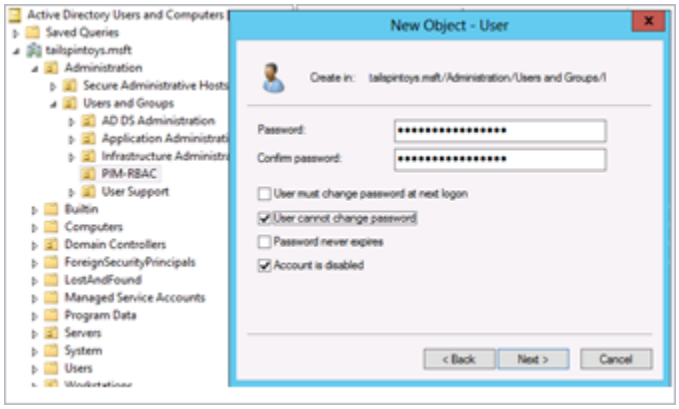

10. Verify that the account details are correct and click **Finish**.
11. Right-click the user object you just created and click **Properties**.
12. Click the **Account** tab.
13. In the **Account Options** field, select the **Account is sensitive and cannot be delegated** flag, select the **This account supports Kerberos AES 128 bit encryption** and/or the **This account supports Kerberos AES 256 encryption** flag, and click **OK**.

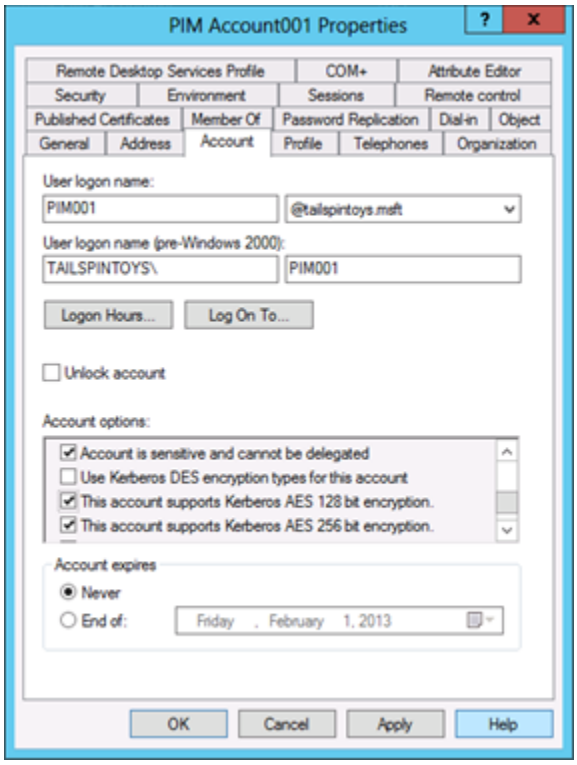

Notes

Because this account, like other accounts, will have a limited, but powerful function, the account should only be used on secure administrative hosts. For all secure administrative hosts in your environment, you should consider implementing the Group Policy setting **Network Security: Configure Encryption types allowed for Kerberos** to allow only the most secure encryption types you can implement for secure hosts.

Although implementing more secure encryption types for the hosts does not mitigate credential theft attacks, the appropriate use and configuration of the secure hosts does. Setting stronger encryption types for hosts that are only used by privileged accounts simply reduces the overall attack surface of the computers.

For more information about configuring encryption types on systems and accounts, see Windows Configurations for Kerberos Supported Encryption Type.

Note These settings are supported only on computers running

Windows Server 2012, Windows Server 2008 R2, Windows 8, or Windows 7.

14. On the **Object** tab, select **Protect object from accidental deletion**. This will not only prevent the object from being deleted (even by authorized users), but will prevent it from being moved to a different OU in your AD DS hierarchy, unless the check box is first cleared by a user with permission to change the attribute.

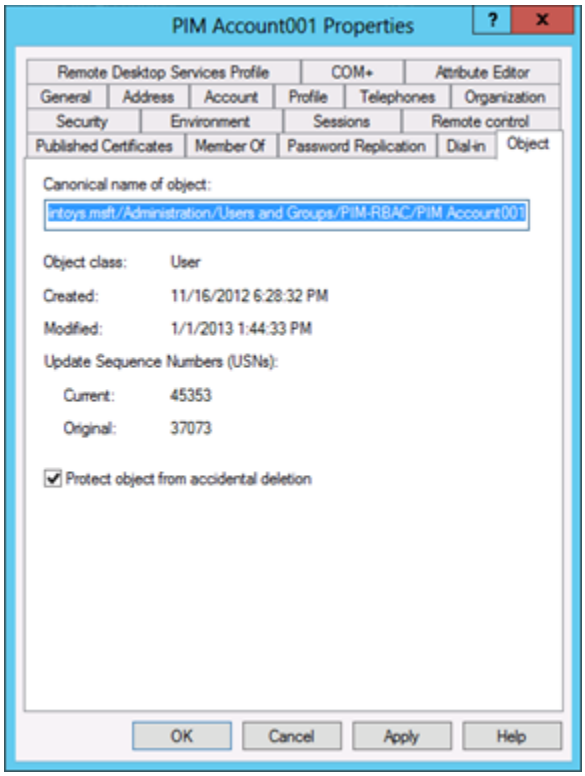

15. Click the **Remote control** tab.
16. Clear the **Enable remote control** flag. It should never be necessary for support staff to connect to this account's sessions to implement fixes.

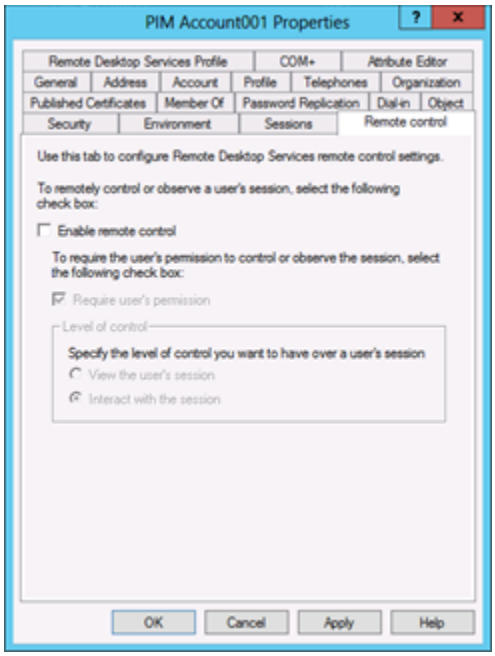

Notes

Every object in Active Directory should have a designated IT owner and a designated business owner, as described in Planning for Compromise. If you are tracking ownership of AD DS objects in Active Directory (as opposed to an external database), you should enter appropriate ownership information in this object's properties.

In this case, the business owner is most likely an IT division, andthere is no prohibition on business owners also being IT owners. The point of establishing ownership of objects is to allow you to identify contacts when changes need to be made to the objects, perhaps years from their initial creation.

17. Click on the **Organization** tab.
18. Enter any information that is required in your AD DS object standards.

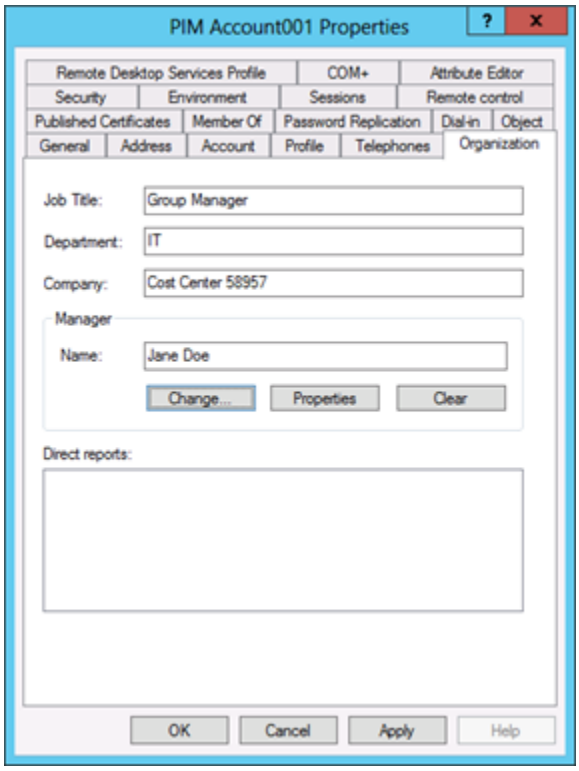

19. Click on the **Dial-in** tab.
20. In the **Network Access Permission** field, select **Deny access**.This account should never need to connect over a remote connection.

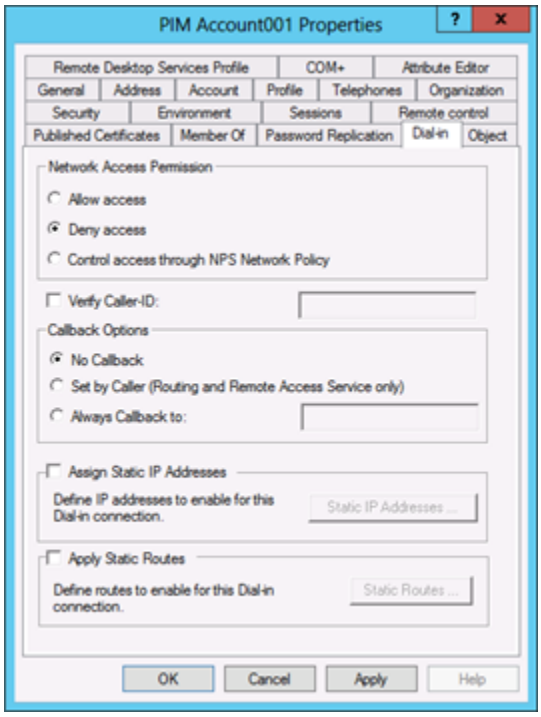

Notes

It is unlikely that this account will be used to log on to read-only domain controllers (RODCs) in your environment. However, should circumstance ever require the account to log on to an RODC, you should add this account to the Denied RODC Password Replication Group so that its password is not cached on the RODC.

Although the account's password should be reset after each use and the account should be disabled, implementing this setting does not have a deleterious effect on the account, and it might help in situations in which an administrator forgets to reset the account's password and disable it.

21. Click the **Member Of** tab.
22. Click **Add**.
23. Type **Denied RODC Password Replication Group** in the **Select Users, Contacts, Computers** dialog box and click **Check Names**. When the name of the group is underlined in the object picker, click **OK** and verify that the account is now a member of the two groups displayed in the following screenshot. Do not add the account to any protected groups.
24. Click **OK**.

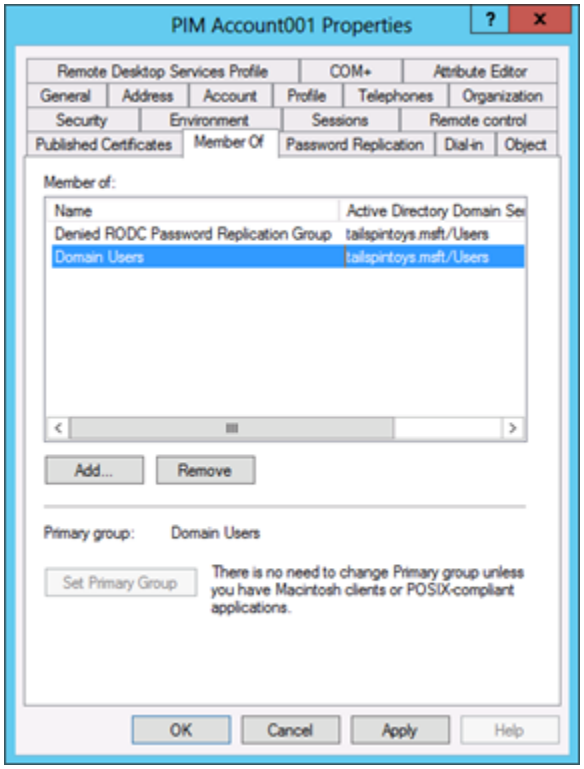

25. Click the **Security** tab and click **Advanced**.
26. In the **Advanced Security Settings** dialog box, click **Disable inheritance** and copy the inherited permissions as explicit permissions, and click **Add**.

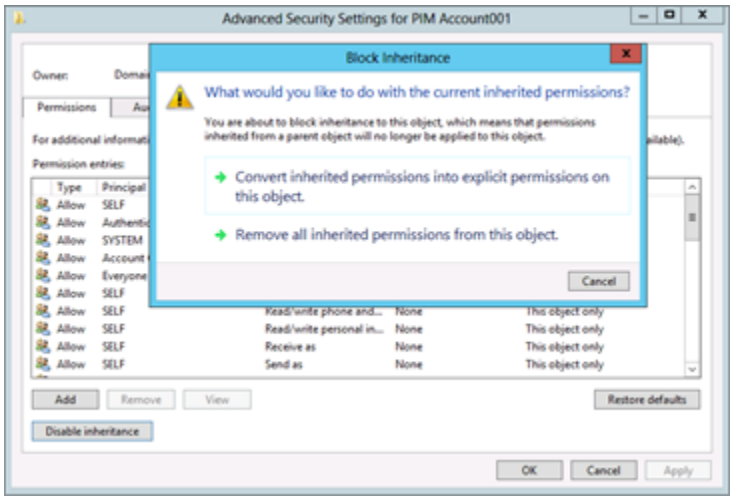

27. In the **Permission Entry for [Account]** dialog box, click **Select a principal** and add the group you created in the previous procedure. Scroll to the bottom of the dialog box and click **Clear all** to remove all default permissions.

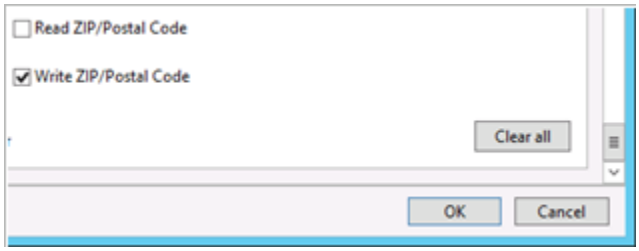

28. Scroll to the top of the **Permission Entry** dialog box. Ensure that the **Type** drop-down list is set to **Allow**, and in the **Applies to** drop-down list, select **This object only**.
29. In the **Permissions** field, select **Read all properties**, **Read permissions**, and **Reset password**.

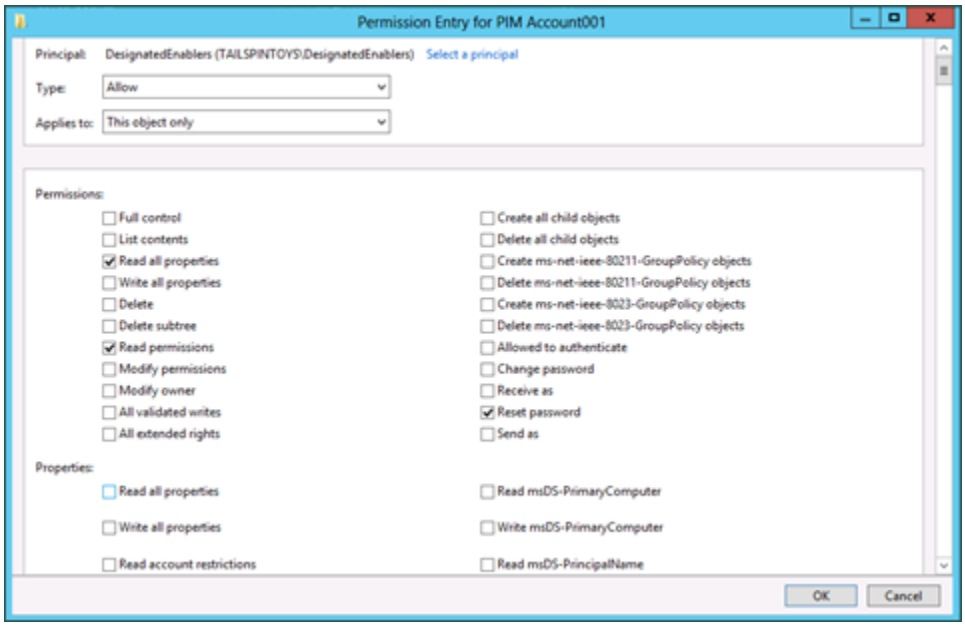

30. In the **Properties** field, select **Read userAccountControl** and **Write userAccountControl**.
31. Click **OK**, **OK** again in the **Advanced Security Settings** dialog box.

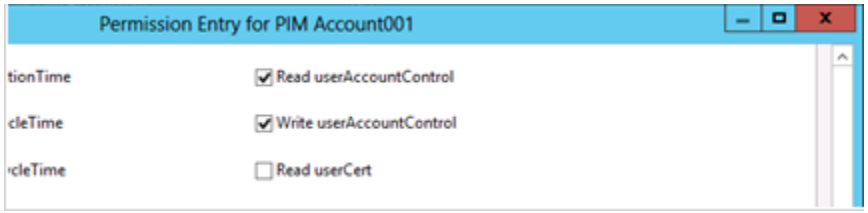

Note

The **userAccountControl** attribute controls multiple account configuration options. You cannot grant permission to change only some of the configuration options when you grant write permission to the attribute.

32. In the **Group or user names** field of the **Security** tab, remove any groups that should not be permitted to access or manage the account. Do not remove any groups that have been configured with Deny ACEs, such as the Everyone group and the SELF computed account (that ACE was set when the **user cannot change password** flag was enabled during creation of the account. Also do not remove the group you just added, the SYSTEM account, or groups such as EA, DA, BA, or the Windows Authorization Access Group.

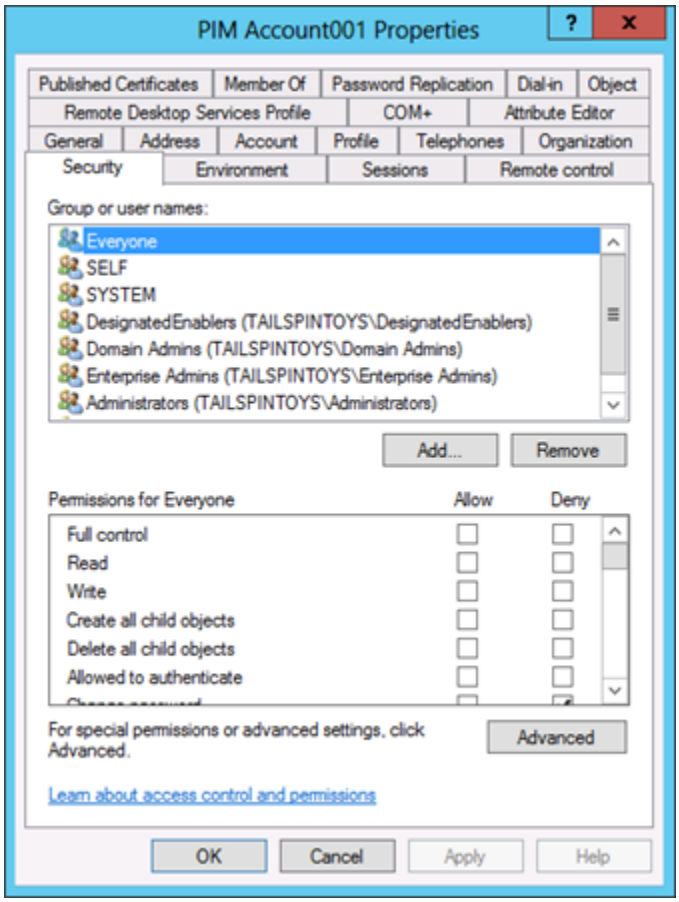

33. Click **Advanced** and verify that the Advanced Security Settings dialog box looks similar to the following screenshot.
34. Click **OK**, and **OK** again to close the account's property dialog box.

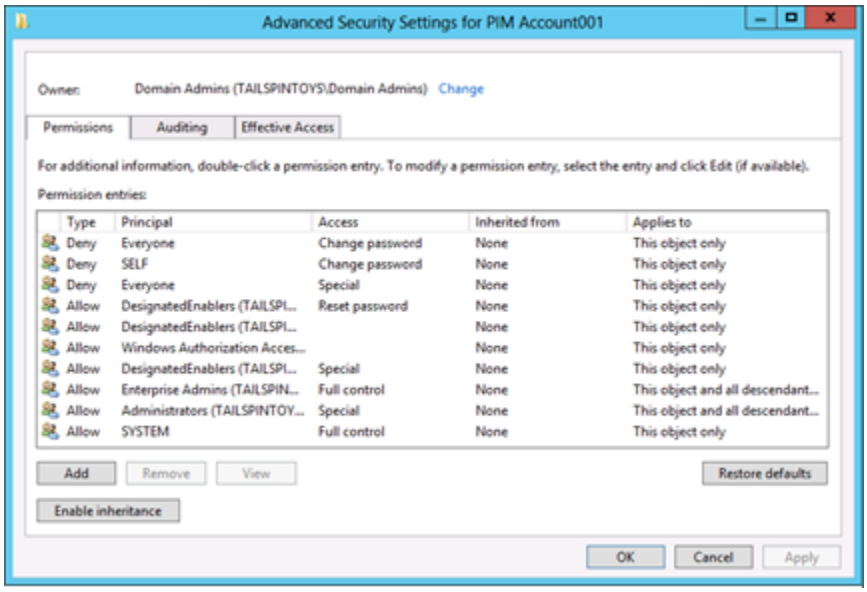

35. Setup of the first management account is now complete. You will test the account in a later procedure.

Creating Additional Management Accounts

You can create additional management accounts by repeating the previous steps, by copying the account you just created, or by creating a script to create accounts with your desired configuration settings. Note, however, that if you copy the account you just created, many of the customized settings and ACLs will not be copied to the new account and you will have to repeat most of the configuration steps.

You can instead create a group to which you delegate rights to populate and unpopulate protected groups, but you will need to secure the group and the accounts you place in it. Because there should be very few accounts in your directory that are granted the ability to manage the membership of protected groups, creating individual accounts might be the simplest approach.

Regardless of how you choose to create a group into which you place the management accounts, you should ensure that each account is secured as described earlier. You should also consider implementing GPO restrictions similar to those described in Appendix D: Securing Built-In Administrator Accounts in Active Directory.

Auditing Management Accounts

You should configure auditing on the account to log, at minimum, all writes to the account. This will allow you to not only identify successful enabling of the account and resetting of its password during authorized uses, but to also identify attempts by unauthorized users to manipulate the account. Failed writes on the account should be captured in your Security Information and Event Monitoring (SIEM) system (if applicable), and should trigger alerts that provide notification to the staff responsible for investigating potential compromises.

SIEM solutions take event information from involved security sources (for example, event logs, application data, network streams, antimalware products, and intrusion detection sources), collate

the data, and try to make intelligent views and proactive actions. There are many commercial SIEM solutions, and many enterprises create private implementations. A well designed and appropriately implemented SIEM can significantly enhance security monitoring and incident response capabilities. However, capabilities and accuracy vary tremendously between solutions. SIEMs are beyond the scope of this paper, but the specific event recommendations contained should be considered by any SIEM implementer.

For more information about recommended audit configuration settings for domain controllers, see Monitoring Active Directory for Signs of Compromise. Domain controller-specific configuration settings are provided in Monitoring Active Directory for Signs of Compromise.

Enabling Management Accounts to Modify the Membership of Protected Groups

In this procedure, you will configure permissions on the domain's AdminSDHolder object to allow the newly created management accounts to modify the membership of protected groups in the domain. This procedure cannot be performed via a graphical user interface (GUI).

As discussed in Appendix C: Protected Accounts and Groups in Active Directory, the ACL on a domain's AdminSDHolder object is effectively "copied" to protected objects when the SDProp task runs. Protected groups and accounts do not inherit their permissions from the AdminSDHolder object; their permissions are explicitly set to match those on the AdminSDHolder object. Therefore, when you modify permissions on the AdminSDHolder object, you must modify them for attributes that are appropriate to the type of the protected object you are targeting.

In this case, you will be granting the newly created management accounts to allow them to read and write the members attribute on group objects. However, the AdminSDHolder object is not a group object and group attributes are not exposed in the graphical ACL editor. It is for this reason that you will implement the permissions changes via the Dsacls command-line utility. To grant the (disabled) management accounts permissions to modify the membership of protected groups, perform the following steps:

1. Log on to a domain controller, preferably the domain controller holding the PDC Emulator (PDCE) role, with the credentials of a user account that has been made a member of the DA group in the domain.

2. Open an elevated command prompt by right-clicking **Command Prompt** and click **Run as administrator**.

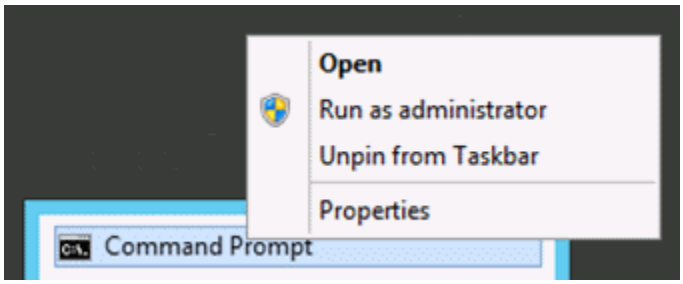

3. When prompted to approve the elevation, click **Yes**.

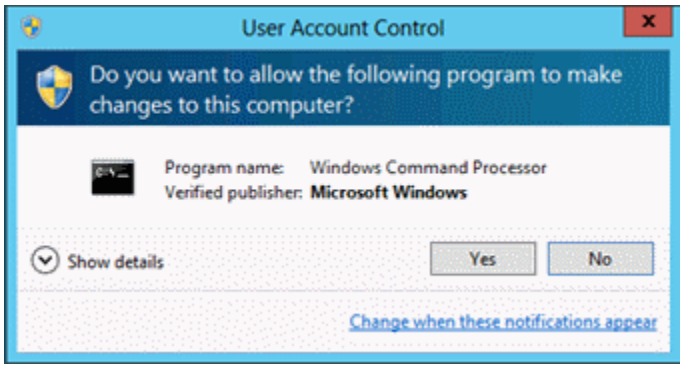

Note
For more information about elevation and user account control (UAC) in Windows, see UAC Processes and Interactions on the TechNet website.

4. At the Command Prompt, type (substituting your domain-specific information) **Dsacls [distinguished name of the AdminSDHolder object in your domain] /G [management account UPN]:RPWP;member**.

The previous command (which is not case-sensitive) works as follows:

- Dsacls sets or displays ACEs on directory objects
- CN=AdminSDHolder,CN=System,DC=TailSpinToys,DC=msft identifies the object to be modified
- /G indicates that a grant ACE is being configured
- PIM001@tailspintoys.msft is the User Principal Name (UPN) of the security principal to which the ACEs will be granted
- RPWP grants read property and write property permissions
- Member is the name of the property (attribute) on which the permissions will be set

For more information about use of **Dsacls**, type Dsacls without any parameters at a command prompt.

If you have created multiple management accounts for the domain, you should run the Dsacls command for each account. When you have completed the ACL configuration on the AdminSDHolder object, you should force SDProp to run, or wait until its scheduled run completes. For information about forcing SDProp to run, see “Running SDProp Manually” in Appendix C: Protected Accounts and Groups in Active Directory.

When SDProp has run, you can verify that the changes you made to the AdminSDHolder object have been applied to protected groups in the domain. You cannot verify this by viewing the ACL on the AdminSDHolder object for the reasons previously described, but you can verify that the permissions have been applied by viewing the ACLs on protected groups.

5. In **Active Directory Users and Computers**, verify that you have enabled **Advanced Features**. To do so, click **View**, locate the **Domain Admins** group, right-click the group and click **Properties**.
6. Click the **Security** tab and click **Advanced** to open the **Advanced Security Settings for Domain Admins** dialog box.

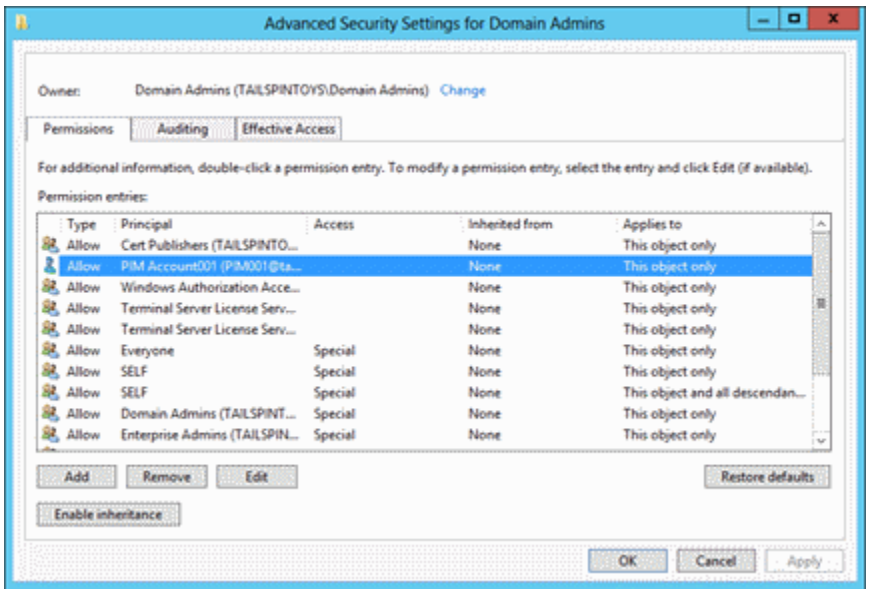

7. Select **Allow ACE for the management account** and click **Edit**. Verify that the account has been granted only **Read Members** and **Write Members** permissions on the DA group, and click **OK**.
8. Click **OK** in the **Advanced Security Settings** dialog box, and click **OK** again to close the property dialog box for the DA group.

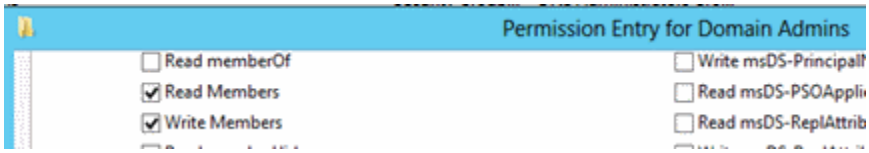

9. You can repeat the previous steps for other protected groups in the domain; the permissions should be the same for all protected groups. You have now completed

creation and configuration of the management accounts for the protected groups in this domain.

> **Note**
> Any account that has permission to write membership of a group in Active Directory can also add itself to the group. This behavior is by design and cannot be disabled. For this reason, you should always keep management accounts disabled when not in use, and should closely monitor the accounts when they're disabled and when they're in use.

Verifying Group and Account Configuration Settings

Now that you have created and configured management accounts that can modify the membership of protected groups in the domain (which includes the most highly privileged EA, DA, and BA groups), you should verify that the accounts and their management group have been created properly. Verification consists of these general tasks:

1. Test the group that can enable and disable management accounts to verify that members of the group can enable and disable the accounts and reset their passwords, but cannot perform other administrative activities on the management accounts.
2. Test the management accounts to verify that they can add and remove members to protected groups in the domain, but cannot change any other properties of protected accounts and groups.

Test the Group that Will Enable and Disable Management Accounts

1. To test enabling a management account and resetting its password, log on to a secure administrative workstation with an account that is a member of the group you created in Appendix I: Creating Management Accounts for Protected Accounts and Groups in Active Directory.

2. Open **Active Directory Users and Computers**, right-click the management account, and click **Enable Account**.

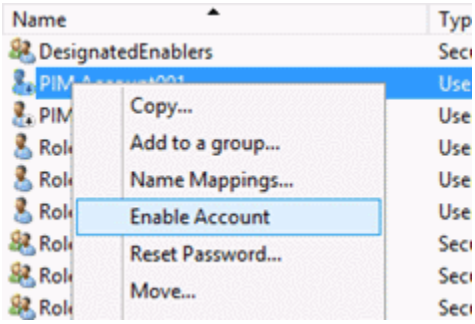

3. A dialog box should display, confirming that the account has been enabled.

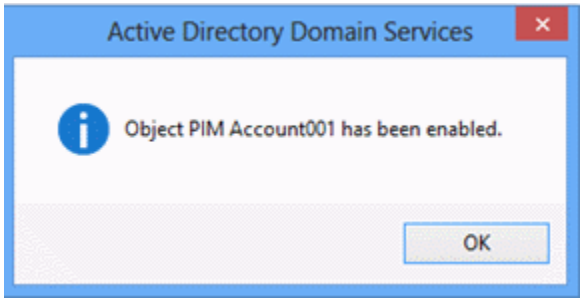

4. Next, reset the password on the management account. To do so, right-click the account again and click **Reset Password**.

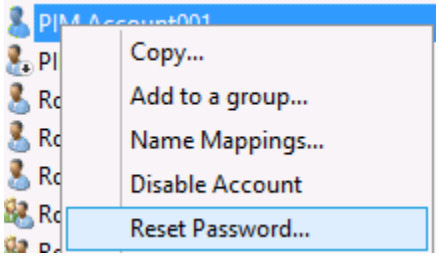

5. Type a new password for the account in the **New password** and **Confirm password** fields, and click **OK**.

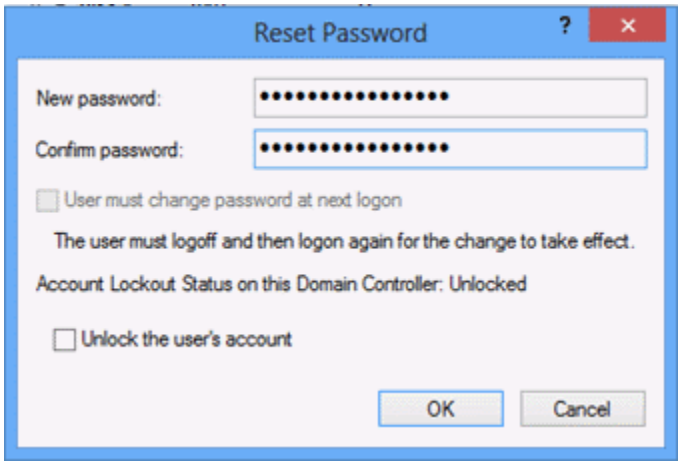

6. A dialog box should appear, confirming that the password for the account has been reset.

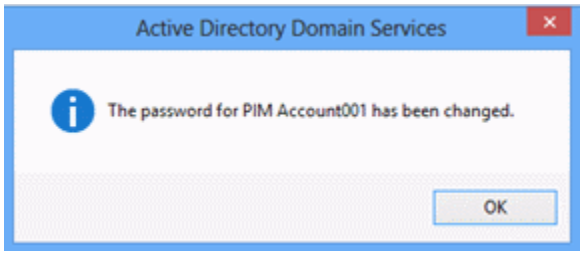

7. Now attempt to modify additional properties of the management account. Right-click the account and click **Properties**, and click the **Remote control** tab.
8. Select **Enable remote control** and click **Apply**. The operation should fail and an **Access Denied** error message should display.

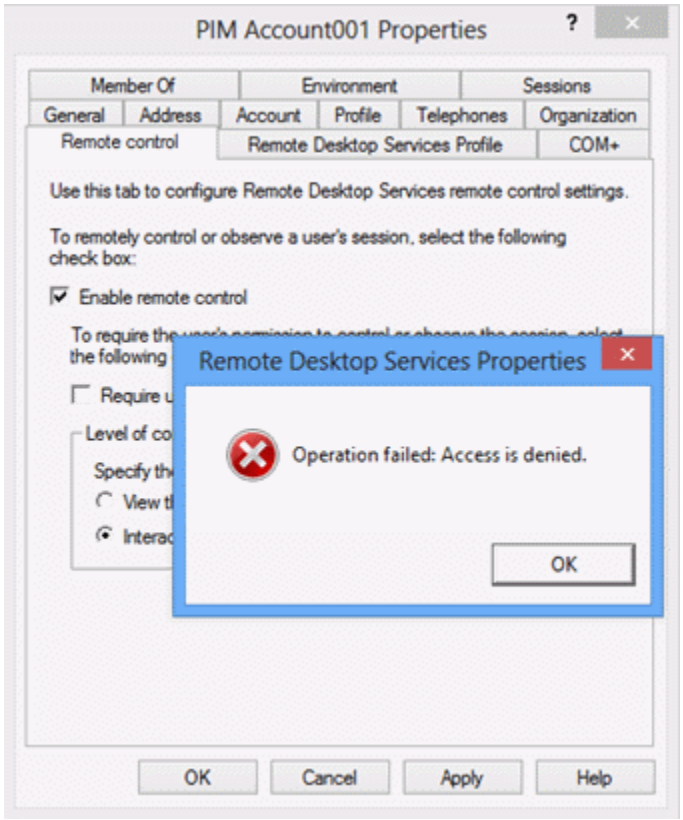

9. Click the **Account** tab for the account and attempt to change the account's name, logon hours, or logon workstations. All should fail, and account options that are not controlled by the **userAccountControl** attribute should be grayed out and unavailable for modification.

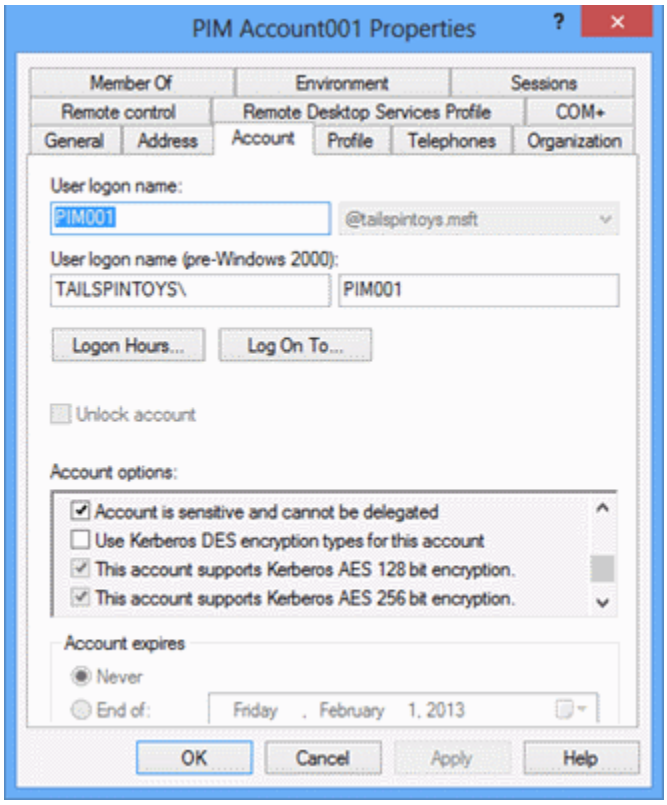

10. Attempt to add the management group to a protected group such as the DA group. When you click **OK**, a message should appear, informing you that you do not have permissions to modify the group.

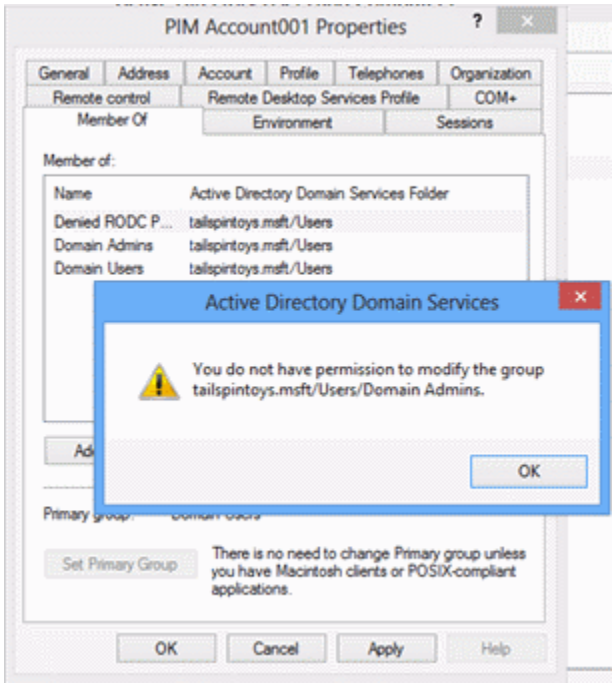

11. Perform additional tests as required to verify that you cannot configure anything on the management account except **userAccountControl** settings and password resets.

    Note
    The **userAccountControl** attribute controls multiple account configuration options. You cannot grant permission to change only some of the configuration options when you grant write permission to the attribute.

Test the Management Accounts

Now that you have enabled one or more accounts that can change the membership of protected groups, you can test the accounts to ensure that they can modify protected group membership, but cannot perform other modifications on protected accounts and groups.

1. Log on to a secure administrative host as the first management account.

2. Launch **Active Directory Users and Computers** and locate the **Domain Admins group**.
3. Right-click the **Domain Admins** group and click **Properties**.

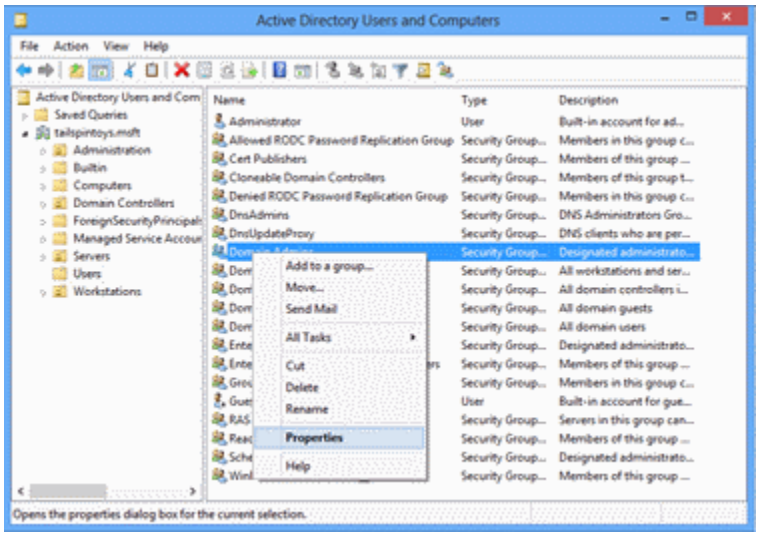

4. In the **Domain Admins Properties**, click the **Members** tab and **click** Add. Enter the name of an account that will be given temporary Domain Admins privileges and click **Check Names**. When the name of the account is underlined, click **OK** to return to the **Members** tab.

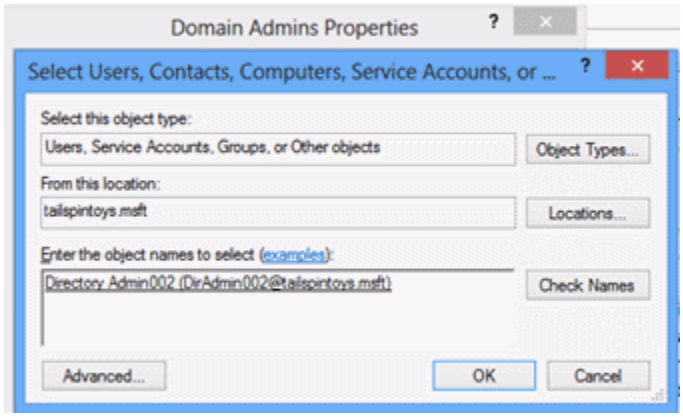

5. On the **Members** tab for the **Domain Admins Properties** dialog box, click **Apply**. After clicking **Apply**, the account should stay a member of the DA group and you should receive no error messages.

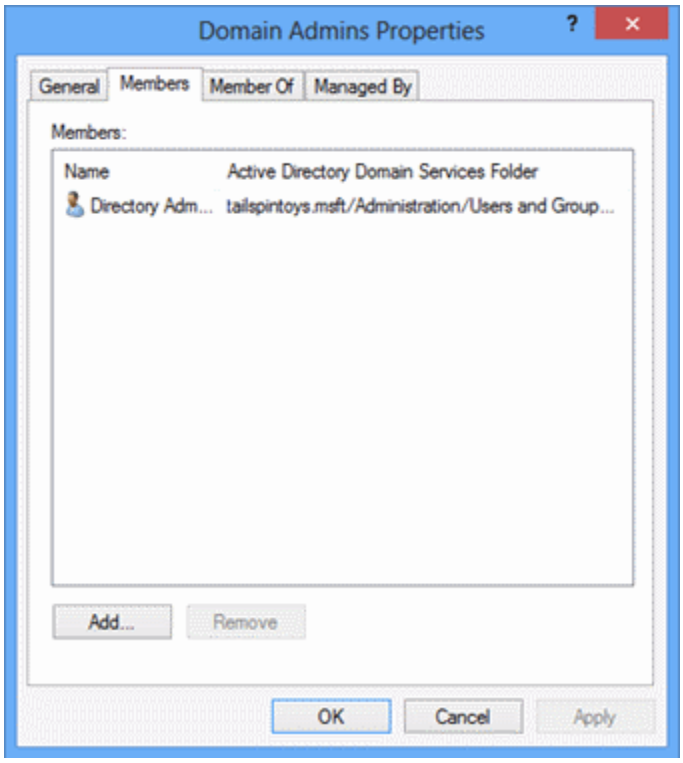

6. Click the **Managed By** tab in the **Domain Admins Properties** dialog box and verify that you cannot enter text in any fields and all buttons are grayed out.

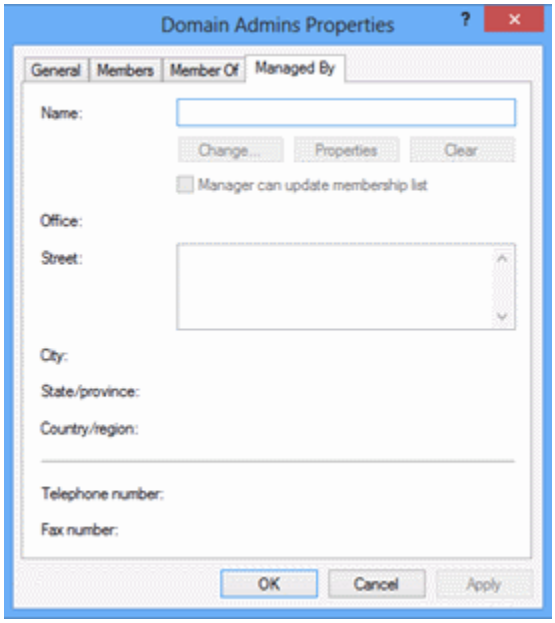

7. Click the **General** tab in the **Domain Admins Properties** dialog box and verify that you cannot modify any of the information about that tab.

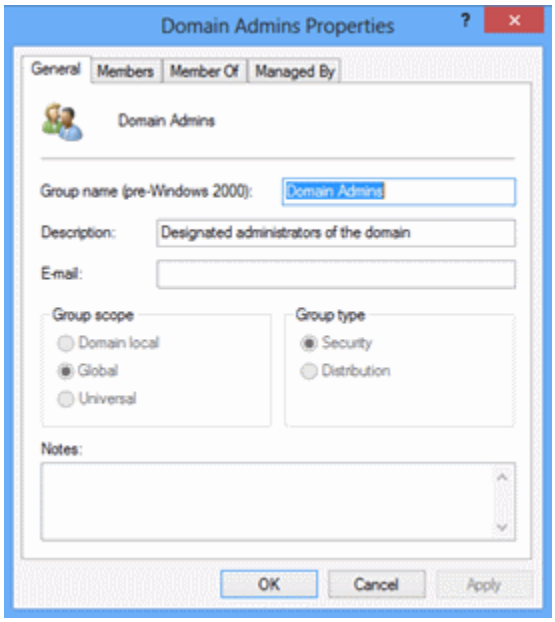

8. Repeat these steps for additional protected groups as needed. When you have finished, log on to a secure administrative host with an account that is a member of the group you created to enable and disable the management accounts. Then reset the password on the management account you just tested and disable the account. You have completed setup of the management accounts and the group that will be responsible for enabling and disabling the accounts.

Appendix J: Third-Party RBAC Vendors

Appendix J: Third-Party RBAC Vendors

Note
Descriptions of software described in this appendix were obtained from the respective vendors' websites. No endorsement of or preference for any solution is intended or implied.

The Dot Net Factory

EmpowerID includes an advanced authorization policy engine that allows organizations to define a user's access to a diverse set of corporate and cloud-hosted resources via flexible RBAC and ABAC rules. This "resultant access" information is then consumed or "pulled" by systems that support leveraging an external authorization engine to make access decisions or "pushed" down onto systems that don't.

Examples of systems supporting the “pull” model are applications that can leverage SAML or WS-Trust Identity and Claims Providers or applications supporting the Microsoft .NET Membership and Role Provider. These would include applications like Microsoft SharePoint 2010, SaaS applications, and internally developed corporate applications. EmpowerID falls into the category of a system supporting “pull” or external authorization.

Unfortunately, the majority of an enterprise’s systems do not yet support external authorization. For these systems, access is defined and controlled within each application’s security database or via ACLs. EmpowerID supports a “push” model for such cases in which the RBAC engine allows organizations to dynamically define who has access to these resources. The EmpowerID sync engine then enforces these policies by translating them into native system permissions or roles, pushing down the changes onto these systems. Additionally, the systems are monitored for permission changes so the EmpowerID engine can detect changes and roll them back when set up to do so. Examples of systems that require the “push” model are Windows Shared Folders, Group membership, Exchange mailboxes, custom database application roles and permissions, and directory ACLs.

Powerful RBAC policies leverage EmpowerID’s multitiered model to pre-calculate access to all known enterprise applications and resources based on an organization’s structure, a person’s job function, and all directly assigned access. These rules allow information from authoritative systems to drive changes in application access and provisioning policies.

ABAC policies on the other hand, provide more fine-grained on-the-fly decisions regarding a user’s access level and the actions they are authorized to perform. ABAC rules benefit from the ability to analyze contextual “in the moment” information to make decisions without the overhead and maintenance of an RBAC structure. However, ABAC rules are more limited in their use of an organization’s structural information as they must real-time decisions and cannot wait for complex analyses or pre-compilation of hierarchical information from multiple sources. The EmpowerID authorization engine is a hybrid of the RBAC and ABAC models, leveraging the best of each. It offers RBAC authorization to leverage diverse information sources for automating role-based access control and ABAC authorization, which further refines RBAC access with fine-grained controls.

Key Features and Benefits:

- Manages and enforces access control for applications that support external authorization and for enterprise systems that require permissions to be pushed down onto them
- A powerful security model supporting Role-Based Access Control (RBAC), Attribute-Based Access Control (ABAC), and Separation of Duties enforcement (SoD)
- Reduces the time to market when developing new applications by eliminating the need to write complex security code into each application
- Supports enterprise compliance initiatives by centralizing authorization into an auditable system
- Reduces risk by reducing the number of places where security logic is maintained and can be modified
- Increases agility by reducing the impact of changes in infrastructure and application providers

- Standards-based support for SAML and WS-Trust applications such as Microsoft SharePoint 2010
- Unique Rights-Based Approval Routing (RBAR) technology automatically routes requests for approval based on delegations without hard-coded logic maintained inside workflows
- Fully programmable supporting integration with custom systems via connectors or a secure web services API

IBM

Domain RBAC

Role-based access control (RBAC), introduced in AIX 6.1, provides a mechanism to split the various functions of the super user root into roles, which can be delegated to other users on the system. RBAC provides the facility to delegate duties and improves the security of the system because the auditing and tracking of activities on the system is easier. RBAC provide delegation of responsibility to another user (referred as an authorized user), but it does not provide a mechanism to limit the administrative rights of an authorized user to specific resources of the system. For example, a user that has network administrative rights can manage every network interface on the system. You cannot restrict the authorized user to modify a set of interfaces.

The domain feature for RBAC is used to restrict access to authorized users. The users and resources of the system are labeled by attaching tags called domains, and the specific access rules determine access to resources by the users.

Oracle

Oracle Solaris RBAC Elements and Basic Concepts

The RBAC model in the Oracle Solaris Operating System introduces the following elements:

Authorization – A permission that enables a user or role to perform a class of actions that require additional rights. For example, security policy at installation gives regular users the solaris.device.cdrw authorization. This authorization enables users to read and write to a CD-ROM device.

Privilege – A discrete right that can be granted to a command, a user, a role, or a system. Privileges enable a process to succeed. For example, the proc_exec privilege allows a process to call execve(). Regular users have basic privileges.

Security attributes – An attribute that enables a process to perform an operation. In a typical UNIX environment, a security attribute enables a process to perform an operation that is otherwise forbidden to regular users. For example, setuid and setgid programs have security attributes. In the RBAC model, authorizations and privileges are security attributes in addition to setuid and setgid programs. These attributes can be assigned to a user. For example, a user with the solaris.device.allocate authorization can allocate a device for exclusive use. Privileges can be

placed on a process. For example, a process with the file_flag_set privilege can set immutable, no-unlink, or append-only file attributes.

Privileged application – An application or command that can override system controls by checking for security attributes. In a typical UNIX environment and in the RBAC model, programs that use setuid and setgid are privileged applications. In the RBAC model, programs that require privileges or authorizations to succeed are also privileged applications.

Rights profile – A collection of administrative capabilities that can be assigned to a role or to a user. A rights profile can consist of authorizations, of commands with security attributes, and of other rights profiles. Rights profiles offer a convenient way to group security attributes.

Role – A special identity for running privileged applications. The special identity can be assumed by assigned users only. In a system that is run by roles, superuser is unnecessary. Superuser capabilities are distributed to different roles. For example, in a two-role system, security tasks would be handled by a security role. The second role would handle system administration tasks that are not security-related. Roles can be more fine-grained. For example, a system could include separate administrative roles for handling the cryptographic framework, printers, system time, file systems, and auditing.

Centrify

IT Security & Access Control

Strengthen IT security with Active Directory-centric access control and policy enforcement for UNIX, Linux and Mac systems and applications.

The Challenge

One of the most difficult questions asked of IT security managers in cross-platform environments is: Can you prove which users have access to a specific business-critical system or application?

For Linux and UNIX systems in particular, access controls might be stored in insecure legacy systems such as NIS or managed locally system by system. Passwords to superuser accounts may be shared among many individuals. Or a single user may have multiple identities across systems.

The Centrify Solution

Centrify addresses this challenge by giving organizations a global view of access controls and user permissions, tied to a single, centrally managed Active Directory identity. With the Centrify Suite, you can:

- Associate all access rights and permissions on audited systems to individual Active Directory accounts
- Define logical sets of systems that can have their own authorized users, administrators, and security policies, with centralized reporting of who has access to what systems

- Implement role-based access controls and limit superuser privileges to only the set of commands they need to perform their jobs
- Add additional layers of security by isolating and protecting systems holding sensitive information.
- Capture detailed logs of all user actions, and system responses, to monitor for suspicious activity
- Globally enforce consistent security and configuration policies (via Windows Group Policy) across a heterogeneous enterprise

Appendix K: Third-Party PIM Vendors

Appendix K: Third-Party PIM Vendors

Note

Descriptions of software described in this appendix were obtained from the respective vendors' websites. No endorsement of or preference for any solution is intended or implied.

Cyber-Ark

Privileged accounts and passwords are extremely powerful, allowing a privileged user to log on anonymously and have complete control of the target system with full access to all of the information about that system. This vulnerability could potentially cause tremendous financial losses and reputational damage for businesses. For enterprises, privileged accounts are especially difficult to manage:

- The average enterprise has thousands of privileged identities, accounts, and passwords. Manually managing and updating these are a time-consuming, costly and repetitive process.
- Administrative and application accounts (hard-coded, embedded credentials) are found on virtually every piece of hardware, software, and application within an organization, including virtual environments.
- Administrative or application accounts are shared, which means that the system does not track WHO logged on as an Administrator, merely that a login occurred—a significant audit challenge.
- Unlike a personal identity, such as Jdoe, administrative or application accounts are nearly impossible to disable due to high potential for disruption to business.
- Administrative and application accounts are subject to regulations such as Sarbanes Oxley, PCI, and Basel II, requiring that companies prove exactly who logs in to sensitive systems and, increasingly, what they are doing.

What is the PIM Suite?

Cyber-Ark's Privileged Identity Management (PIM) Suite is an enterprise-class, unified policy-based solution that secures, manages and logs all privileged accounts and activities associated with datacenter management whether on-premise or in the cloud:

- Control access to privileged accounts based on pre-defined security policies
- Manage application and service credentials
- Grant granular control to the commands superusers can run
- Comply with audit and regulatory requirements
- Streamline policy management of privileged accounts
- Seamlessly integrate with enterprise systems

The PIM Suite allows organizations to manage, track and audit their most privileged identities, avert internal and external threats, and prevent the loss of sensitive information. It complements the Privileged Session Management Suite designed to isolate, protect and monitor all sensitive target systems in your datacenter including servers, network devices, databases and virtual environments and records all privileged sessions on these systems for better visibility, control and smoother audit processes.

The PIM Suite: features and components

The PIM Suite offers a robust set of system features and capabilities for consistent policy definition and enforcement, automated privileged password management, and centralized reporting for compliance audits. The PIM Suite comprises three well integrated core products which can also be purchased separately as needed:

- Enterprise Password Vault
- Application Identity Manager
- On-Demand Privileges Manager

Because they share a common server platform, an initial deployment of any individual solution can quickly and easily be expanded to address any additional audit or security challenges that may arise in the future.

With Cyber-Ark's Privileged Identity Management suite you can:

- Approach Compliance with Confidence: Superior security that protects the 'keys to your kingdom' with a proven ability to meet regulatory requirements
- Minimize Internal or External Threats: Control who is accessing your most sensitive assets with out of the box best practices for defining and enforcing a unified policy for privileged identity management
- Do Business Better: Improve workforce productivity with a single access point for automatically managing privileged credentials

Quest

Privileged Account Management

Controlling and Auditing Superuser Access

Quest One helps you control and audit administrative access with privileged credentials through granular delegation and command control, keystroke logging and session audit, policy-based control, and secure and automated workflows. This approach enhances security and compliance while improving the efficiency of administering superuser access. Administrators are granted only the rights they need—nothing more, nothing less—and all activity is tracked and audited.

- Enhance security by granting administrators only the access rights required for their jobs—nothing more, nothing less—and basing those rights on established and intelligently controlled policy. In addition, when full credentials must be used, secure the process of requesting, approving and issuing access to those accounts, including the critical application-to-application (A2A) and application-to-database (A2D) passwords that pose the greatest security risk.
- Achieve compliance through access control and separation of duties for privileged access that you can track through comprehensive audit capabilities that include policy, rights and activities performed through privileged access—even down to the keystroke level on many critical systems.
- Improve efficiency through granular, policy-driven delegation of elevated access privileges and execution of specific commands across a wide range of systems and platforms, with centralized management and comprehensive audit. Through automated workflows, your administrators gain sufficient rights to do their jobs eliminating the need for manual credential management. In addition, Quest One adds significant value to Sudo by centralizing management of Sudo policy and providing visibility into Sudo-related activities.

Lieberman Software

Privileged Identity Management

Privileged identities are accounts that hold elevated permission to access files, install and run programs, and change configuration settings. These keys to your IT kingdom exist on virtually every server and desktop operating system, business application, database, Web service, and network appliance in your organization.

Risks of Unsecured Privileged Identities

Privileged identities aren't controlled by your identity access management (IAM) system, so in all likelihood:

- You do not know of all the privileged logins that exist on your network;
- You have no record of which privileged credentials are known to different individuals;
- You have no proof of who has used privileged logins to gain access to any of your IT resources, when, and for what purpose;

- There is no way to verify that each of your privileged account passwords are cryptographically strong, are sufficiently unique, and are changed often enough to be secure;
- You have no reliable list of privileged logins stored within your applications, and no way to know which in-house and vendor personnel can use these credentials to access sensitive information.

Controlling Privileged Account Access

Enterprise Random Password Manager (ERPM) from Lieberman Software can help your organization's privileged account management through a four-part I.D.E.A. process:

- Identify and document critical IT assets, their privileged accounts and their interdependencies.
- Delegate access to privileged credentials so that only appropriate personnel, using the least privilege required, can login to IT assets.
- Enforce rules for password complexity, diversity and change frequency, and synchronize changes across all dependencies.
- Audit and alert so that the requester, purpose, and duration of each privileged access request is documented.

ERPM continuously discovers, strengthens, monitors and recovers local, domain and process account passwords in the cross-platform enterprise. It identifies, secures and manages the privileged identities found throughout your IT infrastructure, including:

- Super-user login accounts utilized by individuals to change configuration settings, run programs and perform other administrative duties.
- Service accounts that require privileged login IDs and passwords to run.
- Application-to-application passwords used by web services, line-of-business applications, custom software, and other applications to connect to databases, middleware, and more.

Business Value of Privileged Identity Management

Taking control of privileged identities can help your organization:

- Reduce IT staff workloads by eliminating the manual steps required to secure privileged account credentials, access systems for maintenance, and document each access.
- Improve IT governance by automatically documenting which individuals have access to sensitive data and the ability to make changes that impact IT service delivery; at what times, and for what purpose.
- Lower cost and uncertainty of IT regulatory compliance audits by providing detailed reports that prove compliance with today's regulatory standards including SOX, PCI-DSS, HIPAA, CAG-8 and others.
- Mitigate risks whenever planned and unplanned changes happen in your IT environment or IT staff turnover occurs.

Novell

NetIQ Privileged User Manager

Secure access to UNIX, Linux and Windows systems

Do you have visibility into everything that privileged users are doing on your systems across your environment? Would you know if an unauthorized user gained access to sensitive information? The frequency and seriousness these breaches are increasing and compliance requirements for regulated information are forcing businesses to monitor of privileged user access.

NetIQ Privileged User Manager allows IT administrators to work on systems without exposing superuser (administrator or supervisor) passwords or root-account credentials to the administrator. It specifically targets managing, controlling and recording of all privileged administrator activities for UNIX, Linux and Windows environments.

Smart privileged user control features

- Secure cross-platform privileged user management Control and record “which privileged user have access to what.” You centrally define the commands that privileged users are able to run on any UNIX, Linux or Windows platform.
- Simplified policy management with web-based console. Centrally manage security policies from a single point. The intuitive drag-and-drop interface makes it easy to create rules instead of relying on manual scripting.
- Detailed analysis with color-coded risk ratings Powerful risk-analysis tools record and play back user activity—down to the keystroke level. The unique risk-profiling capability points out any collected user input that poses a risk.
- Automatic data filtering for continuous compliance Prove compliance with permanent audit records 24x7x365, not only around compliance audits. Detailed logs of user activity help maintain your compliance posture.

CA

CA IdentityMinder

CA IdentityMinder™ helps improve the operational efficiency and effectiveness of IT organizations by providing a scalable and configurable identity management foundation that can organize your identity information within the context of your unique business roles and processes. It helps streamline the on-boarding and off-boarding of users, enables the business to manage access requests, and automates identity compliance processes.

Business Challenges

Whether applications reside in the enterprise or the cloud, managing the identities and access of users to key resources is a critical function for IT organizations that are under increasing pressure to cut operating costs while demonstrating continuous compliance. They must also deal with other challenges such as:

- Mitigating risks. Protect critical systems, applications, and information from unauthorized access and use.
- Reducing costs. Increase efficiency and productivity, without sacrificing security.
- Maintaining compliance. Efficiently prove compliance with internal policies, regulations, and best practices.
- Support business initiatives. Adopt new technologies easily (such as virtualization and cloud) that support business initiatives.

Organizations are seeking solutions that automate identity-related processes throughout the enterprise—from the mainframe to the cloud, across employees, contractors, partners, and customers. The result is a smarter, more efficiently managed infrastructure that helps IT save money, reduce risk, and deliver a more reliable service.

Solution Overview

CA IdentityMinder delivers a unified approach for managing users' identities throughout their entire lifecycle and providing them with timely, appropriate access to applications and data.

CA IdentityMinder can be used to organize identity information within the context of an organization's unique business roles and processes. It helps streamline the on-boarding and off-boarding of users, enables the business to manage access requests, and automates identity compliance processes. CA IdentityMinder contains a range of features for managing identities and access rights, and meeting identity compliance requirements.

CA IdentityMinder can increase operational efficiency and user productivity while decreasing Help Desk workload and costs. In addition, the CA Technologies approach to identity management and administration helps improve your overall security posture with a consistent, auditable method for managing identity-related activities and a platform to help maintain adherence to regulations.

Key features

- **User provisioning and deprovisioning**. Automates account provisioning, removal, and approval processes throughout the user's entire lifecycle. Customizable workflows support the unique way each organization approves, alerts, and schedules these activities.
- **User self-service**. Enables users to manage attributes of their own identities, reset
- Passwords and request access to resources, easing the IT and Help Desk burden.
- **Customization without custom code**. Powerful features such as ConfigXpress, PolicyXpress, and ConnectorXpress let you customize your identity management infrastructure without custom code.
- **Securing on-premise and cloud applications**. Provides centralized control of identities, users, roles and policies across on-premise and cloud applications.

Appendix L: Events to Monitor

Appendix L: Events to Monitor

The following table lists events that you should monitor in your environment, according to the recommendations provided in Monitoring Active Directory for Signs of Compromise. In the following table, the “Current Windows Event ID” column lists the event ID as it is implemented in versions of Windows and Windows Server that are currently in mainstream support.

The “Legacy Windows Event ID” column lists the corresponding event ID in legacy versions of Windows such as client computers running Windows XP or earlier and servers running Windows Server 2003 or earlier. The “Potential Criticality” column identifies whether the event should be considered of low, medium, or high criticality in detecting attacks, and the “Event Summary” column provides a brief description of the event.

A potential criticality of High means that one occurrence of the event should be investigated. Potential criticality of Medium or Low means that these events should only be investigated if they occur unexpectedly or in numbers that significantly exceed the expected baseline in a measured period of time. All organizations should test these recommendations in their environments before creating alerts that require mandatory investigative responses. Every environment is different, and some of the events ranked with a potential criticality of High may occur due to other harmless events.

Current Windows Event ID	Legacy Windows Event ID	Potential Criticality	Event Summary
4618	N/A	High	A monitored security event pattern has occurred.
4649	N/A	High	A replay attack was detected. May be a harmless false positive due to misconfiguration error.
4719	612	High	System audit policy was changed.
4765	N/A	High	SID History was added to an account.
4766	N/A	High	An attempt to add SID History to an account failed.
4794	N/A	High	An attempt was made to set the Directory Services Restore Mode.
4897	801	High	Role separation enabled:
4964	N/A	High	Special groups have been assigned to

			a new logon.
5124	N/A	High	A security setting was updated on the OCSP Responder Service
N/A	550	Medium to High	Possible denial-of-service (DoS) attack
1102	517	Medium to High	The audit log was cleared
4621	N/A	Medium	Administrator recovered system from CrashOnAuditFail. Users who are not administrators will now be allowed to log on. Some auditable activity might not have been recorded.
4675	N/A	Medium	SIDs were filtered.
4692	N/A	Medium	Backup of data protection master key was attempted.
4693	N/A	Medium	Recovery of data protection master key was attempted.
4706	610	Medium	A new trust was created to a domain.
4713	617	Medium	Kerberos policy was changed.
4714	618	Medium	Encrypted data recovery policy was changed.
4715	N/A	Medium	The audit policy (SACL) on an object was changed.
4716	620	Medium	Trusted domain information was modified.
4724	628	Medium	An attempt was made to reset an account’s password.
4727	631	Medium	A security-enabled global group was created.
4735	639	Medium	A security-enabled local group was changed.
4737	641	Medium	A security-enabled global group was changed.
4739	643	Medium	Domain Policy was changed.
4754	658	Medium	A security-enabled universal group was created.

4755	659	Medium	A security-enabled universal group was changed.
4764	667	Medium	A security-disabled group was deleted
4764	668	Medium	A group's type was changed.
4780	684	Medium	The ACL was set on accounts which are members of administrators groups.
4816	N/A	Medium	RPC detected an integrity violation while decrypting an incoming message.
4865	N/A	Medium	A trusted forest information entry was added.
4866	N/A	Medium	A trusted forest information entry was removed.
4867	N/A	Medium	A trusted forest information entry was modified.
4868	772	Medium	The certificate manager denied a pending certificate request.
4870	774	Medium	Certificate Services revoked a certificate.
4882	786	Medium	The security permissions for Certificate Services changed.
4885	789	Medium	The audit filter for Certificate Services changed.
4890	794	Medium	The certificate manager settings for Certificate Services changed.
4892	796	Medium	A property of Certificate Services changed.
4896	800	Medium	One or more rows have been deleted from the certificate database.
4906	N/A	Medium	The CrashOnAuditFail value has changed.
4907	N/A	Medium	Auditing settings on object were changed.
4908	N/A	Medium	Special Groups Logon table modified.

4912	807	Medium	Per User Audit Policy was changed.
4960	N/A	Medium	IPsec dropped an inbound packet that failed an integrity check. If this problem persists, it could indicate a network issue or that packets are being modified in transit to this computer. Verify that the packets sent from the remote computer are the same as those received by this computer. This error might also indicate interoperability problems with other IPsec implementations.
4961	N/A	Medium	IPsec dropped an inbound packet that failed a replay check. If this problem persists, it could indicate a replay attack against this computer.
4962	N/A	Medium	IPsec dropped an inbound packet that failed a replay check. The inbound packet had too low a sequence number to ensure it was not a replay.
4963	N/A	Medium	IPsec dropped an inbound clear text packet that should have been secured. This is usually due to the remote computer changing its IPsec policy without informing this computer. This could also be a spoofing attack attempt.
4965	N/A	Medium	IPsec received a packet from a remote computer with an incorrect Security Parameter Index (SPI). This is usually caused by malfunctioning hardware that is corrupting packets. If these errors persist, verify that the packets sent from the remote computer are the same as those received by this computer. This error may also indicate interoperability problems with other IPsec implementations. In that case, if connectivity is not impeded, then these events can be ignored.

4976	N/A	Medium	During Main Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
4977	N/A	Medium	During Quick Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
4978	N/A	Medium	During Extended Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
4983	N/A	Medium	An IPsec Extended Mode negotiation failed. The corresponding Main Mode security association has been deleted.
4984	N/A	Medium	An IPsec Extended Mode negotiation failed. The corresponding Main Mode security association has been deleted.
5027	N/A	Medium	The Windows Firewall Service was unable to retrieve the security policy from the local storage. The service will continue enforcing the current policy.
5028	N/A	Medium	The Windows Firewall Service was unable to parse the new security policy. The service will continue with currently enforced policy.
5029	N/A	Medium	The Windows Firewall Service failed to initialize the driver. The service will continue to enforce the current policy.
5030	N/A	Medium	The Windows Firewall Service failed to start.
5035	N/A	Medium	The Windows Firewall Driver failed to start.
5037	N/A	Medium	The Windows Firewall Driver detected critical runtime error. Terminating.

5038	N/A	Medium	Code integrity determined that the image hash of a file is not valid. The file could be corrupt due to unauthorized modification or the invalid hash could indicate a potential disk device error.
5120	N/A	Medium	OCSP Responder Service Started
5121	N/A	Medium	OCSP Responder Service Stopped
5122	N/A	Medium	A configuration entry changed in OCSP Responder Service
5123	N/A	Medium	A configuration entry changed in OCSP Responder Service
5376	N/A	Medium	Credential Manager credentials were backed up.
5377	N/A	Medium	Credential Manager credentials were restored from a backup.
5453	N/A	Medium	An IPsec negotiation with a remote computer failed because the IKE and AuthIP IPsec Keying Modules (IKEEXT) service is not started.
5480	N/A	Medium	IPsec Services failed to get the complete list of network interfaces on the computer. This poses a potential security risk because some of the network interfaces may not get the protection provided by the applied IPsec filters. Use the IP Security Monitor snap-in to diagnose the problem.
5483	N/A	Medium	IPsec Services failed to initialize RPC server. IPsec Services could not be started.
5484	N/A	Medium	IPsec Services has experienced a critical failure and has been shut down. The shutdown of IPsec Services can put the computer at greater risk of network attack or expose the computer to potential security risks.

5485	N/A	Medium	IPsec Services failed to process some IPsec filters on a plug-and-play event for network interfaces. This poses a potential security risk because some of the network interfaces may not get the protection provided by the applied IPsec filters. Use the IP Security Monitor snap-in to diagnose the problem.
6145	N/A	Medium	One or more errors occurred while processing security policy in the Group Policy objects.
6273	N/A	Medium	Network Policy Server denied access to a user.
6274	N/A	Medium	Network Policy Server discarded the request for a user.
6275	N/A	Medium	Network Policy Server discarded the accounting request for a user.
6276	N/A	Medium	Network Policy Server quarantined a user.
6277	N/A	Medium	Network Policy Server granted access to a user but put it on probation because the host did not meet the defined health policy.
6278	N/A	Medium	Network Policy Server granted full access to a user because the host met the defined health policy.
6279	N/A	Medium	Network Policy Server locked the user account due to repeated failed authentication attempts.
6280	N/A	Medium	Network Policy Server unlocked the user account.
-	640	Medium	General account database changed
-	619	Medium	Quality of Service Policy changed
24586	N/A	Medium	An error was encountered converting volume

24592	N/A	Medium	An attempt to automatically restart conversion on volume %2 failed.
24593	N/A	Medium	Metadata write: Volume %2 returning errors while trying to modify metadata. If failures continue, decrypt volume
24594	N/A	Medium	Metadata rebuild: An attempt to write a copy of metadata on volume %2 failed and may appear as disk corruption. If failures continue, decrypt volume.
4608	512	Low	Windows is starting up.
4609	513	Low	Windows is shutting down.
4610	514	Low	An authentication package has been loaded by the Local Security Authority.
4611	515	Low	A trusted logon process has been registered with the Local Security Authority.
4612	516	Low	Internal resources allocated for the queuing of audit messages have been exhausted, leading to the loss of some audits.
4614	518	Low	A notification package has been loaded by the Security Account Manager.
4615	519	Low	Invalid use of LPC port.
4616	520	Low	The system time was changed.
4622	N/A	Low	A security package has been loaded by the Local Security Authority.
4624	528,540	Low	An account was successfully logged on.
4625	529-537,539	Low	An account failed to log on.
4634	538	Low	An account was logged off.
4646	N/A	Low	IKE DoS-prevention mode started.
4647	551	Low	User initiated logoff.
4648	552	Low	A logon was attempted using explicit

			credentials.
4650	N/A	Low	An IPsec Main Mode security association was established. Extended Mode was not enabled. Certificate authentication was not used.
4651	N/A	Low	An IPsec Main Mode security association was established. Extended Mode was not enabled. A certificate was used for authentication.
4652	N/A	Low	An IPsec Main Mode negotiation failed.
4653	N/A	Low	An IPsec Main Mode negotiation failed.
4654	N/A	Low	An IPsec Quick Mode negotiation failed.
4655	N/A	Low	An IPsec Main Mode security association ended.
4656	560	Low	A handle to an object was requested.
4657	567	Low	A registry value was modified.
4658	562	Low	The handle to an object was closed.
4659	N/A	Low	A handle to an object was requested with intent to delete.
4660	564	Low	An object was deleted.
4661	565	Low	A handle to an object was requested.
4662	566	Low	An operation was performed on an object.
4663	567	Low	An attempt was made to access an object.
4664	N/A	Low	An attempt was made to create a hard link.
4665	N/A	Low	An attempt was made to create an application client context.
4666	N/A	Low	An application attempted an operation:
4667	N/A	Low	An application client context was

			deleted.
4668	N/A	Low	An application was initialized.
4670	N/A	Low	Permissions on an object were changed.
4671	N/A	Low	An application attempted to access a blocked ordinal through the TBS.
4672	576	Low	Special privileges assigned to new logon.
4673	577	Low	A privileged service was called.
4674	578	Low	An operation was attempted on a privileged object.
4688	592	Low	A new process has been created.
4689	593	Low	A process has exited.
4690	594	Low	An attempt was made to duplicate a handle to an object.
4691	595	Low	Indirect access to an object was requested.
4694	N/A	Low	Protection of auditable protected data was attempted.
4695	N/A	Low	Unprotection of auditable protected data was attempted.
4696	600	Low	A primary token was assigned to process.
4697	601	Low	Attempt to install a service
4698	602	Low	A scheduled task was created.
4699	602	Low	A scheduled task was deleted.
4700	602	Low	A scheduled task was enabled.
4701	602	Low	A scheduled task was disabled.
4702	602	Low	A scheduled task was updated.
4704	608	Low	A user right was assigned.
4705	609	Low	A user right was removed.
4707	611	Low	A trust to a domain was removed.

4709	N/A	Low	IPsec Services was started.
4710	N/A	Low	IPsec Services was disabled.
4711	N/A	Low	May contain any one of the following: PAStore Engine applied locally cached copy of Active Directory storage IPsec policy on the computer.PAStore Engine applied Active Directory storage IPsec policy on the computer.PAStore Engine applied local registry storage IPsec policy on the computer.PAStore Engine failed to apply locally cached copy of Active Directory storage IPsec policy on the computer.PAStore Engine failed to apply Active Directory storage IPsec policy on the computer.PAStore Engine failed to apply local registry storage IPsec policy on the computer.PAStore Engine failed to apply some rules of the active IPsec policy on the computer.PAStore Engine failed to load directory storage IPsec policy on the computer.PAStore Engine loaded directory storage IPsec policy on the computer.PAStore Engine failed to load local storage IPsec policy on the computer.PAStore Engine loaded local storage IPsec policy on the computer.PAStore Engine polled for changes to the active IPsec policy and detected no changes.
4712	N/A	Low	IPsec Services encountered a potentially serious failure.
4717	621	Low	System security access was granted to an account.
4718	622	Low	System security access was removed from an account.
4720	624	Low	A user account was created.
4722	626	Low	A user account was enabled.

4723	627	Low	An attempt was made to change an account's password.
4725	629	Low	A user account was disabled.
4726	630	Low	A user account was deleted.
4728	632	Low	A member was added to a security-enabled global group.
4729	633	Low	A member was removed from a security-enabled global group.
4730	634	Low	A security-enabled global group was deleted.
4731	635	Low	A security-enabled local group was created.
4732	636	Low	A member was added to a security-enabled local group.
4733	637	Low	A member was removed from a security-enabled local group.
4734	638	Low	A security-enabled local group was deleted.
4738	642	Low	A user account was changed.
4740	644	Low	A user account was locked out.
4741	645	Low	A computer account was changed.
4742	646	Low	A computer account was changed.
4743	647	Low	A computer account was deleted.
4744	648	Low	A security-disabled local group was created.
4745	649	Low	A security-disabled local group was changed.
4746	650	Low	A member was added to a security-disabled local group.
4747	651	Low	A member was removed from a security-disabled local group.
4748	652	Low	A security-disabled local group was deleted.

4749	653	Low	A security-disabled global group was created.
4750	654	Low	A security-disabled global group was changed.
4751	655	Low	A member was added to a security-disabled global group.
4752	656	Low	A member was removed from a security-disabled global group.
4753	657	Low	A security-disabled global group was deleted.
4756	660	Low	A member was added to a security-enabled universal group.
4757	661	Low	A member was removed from a security-enabled universal group.
4758	662	Low	A security-enabled universal group was deleted.
4759	663	Low	A security-disabled universal group was created.
4760	664	Low	A security-disabled universal group was changed.
4761	665	Low	A member was added to a security-disabled universal group.
4762	666	Low	A member was removed from a security-disabled universal group.
4767	671	Low	A user account was unlocked.
4768	672,676	Low	A Kerberos authentication ticket (TGT) was requested.
4769	673	Low	A Kerberos service ticket was requested.
4770	674	Low	A Kerberos service ticket was renewed.
4771	675	Low	Kerberos pre-authentication failed.
4772	672	Low	A Kerberos authentication ticket request failed.

4774	678	Low	An account was mapped for logon.
4775	679	Low	An account could not be mapped for logon.
4776	680,681	Low	The domain controller attempted to validate the credentials for an account.
4777	N/A	Low	The domain controller failed to validate the credentials for an account.
4778	682	Low	A session was reconnected to a Window Station.
4779	683	Low	A session was disconnected from a Window Station.
4781	685	Low	The name of an account was changed:
4782	N/A	Low	The password hash an account was accessed.
4783	667	Low	A basic application group was created.
4784	N/A	Low	A basic application group was changed.
4785	689	Low	A member was added to a basic application group.
4786	690	Low	A member was removed from a basic application group.
4787	691	Low	A nonmember was added to a basic application group.
4788	692	Low	A nonmember was removed from a basic application group.
4789	693	Low	A basic application group was deleted.
4790	694	Low	An LDAP query group was created.
4793	N/A	Low	The Password Policy Checking API was called.
4800	N/A	Low	The workstation was locked.
4801	N/A	Low	The workstation was unlocked.
4802	N/A	Low	The screen saver was invoked.
4803	N/A	Low	The screen saver was dismissed.

4864	N/A	Low	A namespace collision was detected.
4869	773	Low	Certificate Services received a resubmitted certificate request.
4871	775	Low	Certificate Services received a request to publish the certificate revocation list (CRL).
4872	776	Low	Certificate Services published the certificate revocation list (CRL).
4873	777	Low	A certificate request extension changed.
4874	778	Low	One or more certificate request attributes changed.
4875	779	Low	Certificate Services received a request to shut down.
4876	780	Low	Certificate Services backup started.
4877	781	Low	Certificate Services backup completed.
4878	782	Low	Certificate Services restore started.
4879	783	Low	Certificate Services restore completed.
4880	784	Low	Certificate Services started.
4881	785	Low	Certificate Services stopped.
4883	787	Low	Certificate Services retrieved an archived key.
4884	788	Low	Certificate Services imported a certificate into its database.
4886	790	Low	Certificate Services received a certificate request.
4887	791	Low	Certificate Services approved a certificate request and issued a certificate.
4888	792	Low	Certificate Services denied a certificate request.
4889	793	Low	Certificate Services set the status of a certificate request to pending.

4891	795	Low	A configuration entry changed in Certificate Services.
4893	797	Low	Certificate Services archived a key.
4894	798	Low	Certificate Services imported and archived a key.
4895	799	Low	Certificate Services published the CA certificate to Active Directory Domain Services.
4898	802	Low	Certificate Services loaded a template.
4902	N/A	Low	The Per-user audit policy table was created.
4904	N/A	Low	An attempt was made to register a security event source.
4905	N/A	Low	An attempt was made to unregister a security event source.
4909	N/A	Low	The local policy settings for the TBS were changed.
4910	N/A	Low	The Group Policy settings for the TBS were changed.
4928	N/A	Low	An Active Directory replica source naming context was established.
4929	N/A	Low	An Active Directory replica source naming context was removed.
4930	N/A	Low	An Active Directory replica source naming context was modified.
4931	N/A	Low	An Active Directory replica destination naming context was modified.
4932	N/A	Low	Synchronization of a replica of an Active Directory naming context has begun.
4933	N/A	Low	Synchronization of a replica of an Active Directory naming context has ended.
4934	N/A	Low	Attributes of an Active Directory object were replicated.

4935	N/A	Low	Replication failure begins.
4936	N/A	Low	Replication failure ends.
4937	N/A	Low	A lingering object was removed from a replica.
4944	N/A	Low	The following policy was active when the Windows Firewall started.
4945	N/A	Low	A rule was listed when the Windows Firewall started.
4946	N/A	Low	A change has been made to Windows Firewall exception list. A rule was added.
4947	N/A	Low	A change has been made to Windows Firewall exception list. A rule was modified.
4948	N/A	Low	A change has been made to Windows Firewall exception list. A rule was deleted.
4949	N/A	Low	Windows Firewall settings were restored to the default values.
4950	N/A	Low	A Windows Firewall setting has changed.
4951	N/A	Low	A rule has been ignored because its major version number was not recognized by Windows Firewall.
4952	N/A	Low	Parts of a rule have been ignored because its minor version number was not recognized by Windows Firewall. The other parts of the rule will be enforced.
4953	N/A	Low	A rule has been ignored by Windows Firewall because it could not parse the rule.
4954	N/A	Low	Windows Firewall Group Policy settings have changed. The new settings have been applied.
4956	N/A	Low	Windows Firewall has changed the

			active profile.
4957	N/A	Low	Windows Firewall did not apply the following rule:
4958	N/A	Low	Windows Firewall did not apply the following rule because the rule referred to items not configured on this computer:
4979	N/A	Low	IPsec Main Mode and Extended Mode security associations were established.
4980	N/A	Low	IPsec Main Mode and Extended Mode security associations were established.
4981	N/A	Low	IPsec Main Mode and Extended Mode security associations were established.
4982	N/A	Low	IPsec Main Mode and Extended Mode security associations were established.
4985	N/A	Low	The state of a transaction has changed.
5024	N/A	Low	The Windows Firewall Service has started successfully.
5025	N/A	Low	The Windows Firewall Service has been stopped.
5031	N/A	Low	The Windows Firewall Service blocked an application from accepting incoming connections on the network.
5032	N/A	Low	Windows Firewall was unable to notify the user that it blocked an application from accepting incoming connections on the network.
5033	N/A	Low	The Windows Firewall Driver has started successfully.
5034	N/A	Low	The Windows Firewall Driver has been stopped.
5039	N/A	Low	A registry key was virtualized.
5040	N/A	Low	A change has been made to IPsec settings. An Authentication Set was added.

5041	N/A	Low	A change has been made to IPsec settings. An Authentication Set was modified.
5042	N/A	Low	A change has been made to IPsec settings. An Authentication Set was deleted.
5043	N/A	Low	A change has been made to IPsec settings. A Connection Security Rule was added.
5044	N/A	Low	A change has been made to IPsec settings. A Connection Security Rule was modified.
5045	N/A	Low	A change has been made to IPsec settings. A Connection Security Rule was deleted.
5046	N/A	Low	A change has been made to IPsec settings. A Crypto Set was added.
5047	N/A	Low	A change has been made to IPsec settings. A Crypto Set was modified.
5048	N/A	Low	A change has been made to IPsec settings. A Crypto Set was deleted.
5050	N/A	Low	An attempt to programmatically disable the Windows Firewall using a call to InetFwProfile.FirewallEnabled(False)
5051	N/A	Low	A file was virtualized.
5056	N/A	Low	A cryptographic self test was performed.
5057	N/A	Low	A cryptographic primitive operation failed.
5058	N/A	Low	Key file operation.
5059	N/A	Low	Key migration operation.
5060	N/A	Low	Verification operation failed.
5061	N/A	Low	Cryptographic operation.
5062	N/A	Low	A kernel-mode cryptographic self test was performed.

5063	N/A	Low	A cryptographic provider operation was attempted.
5064	N/A	Low	A cryptographic context operation was attempted.
5065	N/A	Low	A cryptographic context modification was attempted.
5066	N/A	Low	A cryptographic function operation was attempted.
5067	N/A	Low	A cryptographic function modification was attempted.
5068	N/A	Low	A cryptographic function provider operation was attempted.
5069	N/A	Low	A cryptographic function property operation was attempted.
5070	N/A	Low	A cryptographic function property modification was attempted.
5125	N/A	Low	A request was submitted to the OCSP Responder Service
5126	N/A	Low	Signing Certificate was automatically updated by the OCSP Responder Service
5127	N/A	Low	The OCSP Revocation Provider successfully updated the revocation information
5136	566	Low	A directory service object was modified.
5137	566	Low	A directory service object was created.
5138	N/A	Low	A directory service object was undeleted.
5139	N/A	Low	A directory service object was moved.
5140	N/A	Low	A network share object was accessed.
5141	N/A	Low	A directory service object was deleted.
5152	N/A	Low	The Windows Filtering Platform blocked a packet.

5153	N/A	Low	A more restrictive Windows Filtering Platform filter has blocked a packet.
5154	N/A	Low	The Windows Filtering Platform has permitted an application or service to listen on a port for incoming connections.
5155	N/A	Low	The Windows Filtering Platform has blocked an application or service from listening on a port for incoming connections.
5156	N/A	Low	The Windows Filtering Platform has allowed a connection.
5157	N/A	Low	The Windows Filtering Platform has blocked a connection.
5158	N/A	Low	The Windows Filtering Platform has permitted a bind to a local port.
5159	N/A	Low	The Windows Filtering Platform has blocked a bind to a local port.
5378	N/A	Low	The requested credentials delegation was disallowed by policy.
5440	N/A	Low	The following callout was present when the Windows Filtering Platform Base Filtering Engine started.
5441	N/A	Low	The following filter was present when the Windows Filtering Platform Base Filtering Engine started.
5442	N/A	Low	The following provider was present when the Windows Filtering Platform Base Filtering Engine started.
5443	N/A	Low	The following provider context was present when the Windows Filtering Platform Base Filtering Engine started.
5444	N/A	Low	The following sublayer was present when the Windows Filtering Platform Base Filtering Engine started.
5446	N/A	Low	A Windows Filtering Platform callout

			has been changed.
5447	N/A	Low	A Windows Filtering Platform filter has been changed.
5448	N/A	Low	A Windows Filtering Platform provider has been changed.
5449	N/A	Low	A Windows Filtering Platform provider context has been changed.
5450	N/A	Low	A Windows Filtering Platform sublayer has been changed.
5451	N/A	Low	An IPsec Quick Mode security association was established.
5452	N/A	Low	An IPsec Quick Mode security association ended.
5456	N/A	Low	PAStore Engine applied Active Directory storage IPsec policy on the computer.
5457	N/A	Low	PAStore Engine failed to apply Active Directory storage IPsec policy on the computer.
5458	N/A	Low	PAStore Engine applied locally cached copy of Active Directory storage IPsec policy on the computer.
5459	N/A	Low	PAStore Engine failed to apply locally cached copy of Active Directory storage IPsec policy on the computer.
5460	N/A	Low	PAStore Engine applied local registry storage IPsec policy on the computer.
5461	N/A	Low	PAStore Engine failed to apply local registry storage IPsec policy on the computer.
5462	N/A	Low	PAStore Engine failed to apply some rules of the active IPsec policy on the computer. Use the IP Security Monitor snap-in to diagnose the problem.
5463	N/A	Low	PAStore Engine polled for changes to the active IPsec policy and detected

			no changes.
5464	N/A	Low	PAStore Engine polled for changes to the active IPsec policy, detected changes, and applied them to IPsec Services.
5465	N/A	Low	PAStore Engine received a control for forced reloading of IPsec policy and processed the control successfully.
5466	N/A	Low	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory cannot be reached, and will use the cached copy of the Active Directory IPsec policy instead. Any changes made to the Active Directory IPsec policy since the last poll could not be applied.
5467	N/A	Low	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory can be reached, and found no changes to the policy. The cached copy of the Active Directory IPsec policy is no longer being used.
5468	N/A	Low	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory can be reached, found changes to the policy, and applied those changes. The cached copy of the Active Directory IPsec policy is no longer being used.
5471	N/A	Low	PAStore Engine loaded local storage IPsec policy on the computer.
5472	N/A	Low	PAStore Engine failed to load local storage IPsec policy on the computer.
5473	N/A	Low	PAStore Engine loaded directory storage IPsec policy on the computer.
5474	N/A	Low	PAStore Engine failed to load directory

			storage IPsec policy on the computer.
5477	N/A	Low	PAStore Engine failed to add quick mode filter.
5479	N/A	Low	IPsec Services has been shut down successfully. The shutdown of IPsec Services can put the computer at greater risk of network attack or expose the computer to potential security risks.
5632	N/A	Low	A request was made to authenticate to a wireless network.
5633	N/A	Low	A request was made to authenticate to a wired network.
5712	N/A	Low	A Remote Procedure Call (RPC) was attempted.
5888	N/A	Low	An object in the COM+ Catalog was modified.
5889	N/A	Low	An object was deleted from the COM+ Catalog.
5890	N/A	Low	An object was added to the COM+ Catalog.
6008	N/A	Low	The previous system shutdown was unexpected
6144	N/A	Low	Security policy in the Group Policy objects has been applied successfully.
6272	N/A	Low	Network Policy Server granted access to a user.
N/A	561	Low	A handle to an object was requested.
N/A	563	Low	Object open for delete
N/A	625	Low	User Account Type Changed
N/A	613	Low	IPsec policy agent started
N/A	614	Low	IPsec policy agent disabled
N/A	615	Low	IPsec policy agent
N/A	616	Low	IPsec policy agent encountered a

			potential serious failure
24577	N/A	Low	Encryption of volume started
24578	N/A	Low	Encryption of volume stopped
24579	N/A	Low	Encryption of volume completed
24580	N/A	Low	Decryption of volume started
24581	N/A	Low	Decryption of volume stopped
24582	N/A	Low	Decryption of volume completed
24583	N/A	Low	Conversion worker thread for volume started
24584	N/A	Low	Conversion worker thread for volume temporarily stopped
24588	N/A	Low	The conversion operation on volume %2 encountered a bad sector error. Please validate the data on this volume
24595	N/A	Low	Volume %2 contains bad clusters. These clusters will be skipped during conversion.
24621	N/A	Low	Initial state check: Rolling volume conversion transaction on %2.
5049	N/A	Low	An IPsec Security Association was deleted.
5478	N/A	Low	IPsec Services has started successfully.

Notes

Refer to Microsoft Support article 947226 for lists of many security event IDs and their meanings.

Run **wevtutil gp Microsoft-Windows-Security-Auditing /ge /gm:true** to get a very detailed listing of all security event IDs

For more information about Windows security event IDs and their meanings, see the Microsoft Support articles Description of security events in Windows Vista and in Windows Server 2008 and Description of security events in Windows 7 and in Windows Server 2008 R2. You can also download Security Audit Events for Windows 7 and Windows Server 2008 R2 and Windows 8

and Windows Server 2012 Security Event Details, which provide detailed event information for the referenced operating systems in spreadsheet format.

Appendix M: Document Links and Recommended Reading

Appendix M: Document Links and Recommended Reading

Document Links

The following table contains a list of links to external documents and their URLs so that readers of hard copies of this document can access this information. The links are listed in the order they appear in the document.

Links	URLs
10 Immutable Laws of Security Administration	http://technet.microsoft.com/en-us/library/cc722488.aspx
Microsoft Security Compliance Manager	http://technet.microsoft.com/en-us/library/cc677002.aspx
Gartner Symposium ITXPO	http://www.gartner.com/technology/symposium/orlando/
2012 Data Breach Investigations Report (DBIR)	http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf
Ten Immutable Laws of	http://technet.microsoft.com/en-us/security/hh278941.aspx

Security (Version 2.0)	
Using Heuristic Scanning	http://technet.microsoft.com/en-us/library/bb418939.aspx
Drive-by download	http://www.microsoft.com/security/sir/glossary/drive-by-download-sites.aspx
Microsoft Support article 2526083	http://support.microsoft.com/kb/2526083
Microsoft Support article 814777	http://support.microsoft.com/kb/814777
Open Web Application Security Project (OWASP)	https://www.owasp.org/index.php/Main_Page
Microsoft Security Development Lifecycle	http://www.microsoft.com/security/sdl/default.aspx
Mitigating Pass-the-Hash (PtH) Attacks and Other Credential Theft Techniques	http://download.microsoft.com/download/7/7/A/77ABC5BD-8320-41AF-863C-6ECFB10CB4B9/Mitigating Pass-the-Hash (PtH) Attacks and Other Credential Theft Techniques_English.pdf
Determined Adversaries and Targeted Attacks	http://www.microsoft.com/en-us/download/details.aspx?id=34793
Solution for	http://code.msdn.microsoft.com/windowsdesktop/Solution-for-management-of-

management of built-in Administrator account's password via GPO	ae44e789
Microsoft Support article 817433	http://support.microsoft.com/?id=817433
Microsoft Support article 973840	http://support.microsoft.com/kb/973840
Administrator account is disabled by default	http://technet.microsoft.com/en-us/library/cc753450.aspx
The Administrator Accounts Security Planning Guide	http://technet.microsoft.com/en-us/library/cc162797.aspx
Microsoft Windows Security Resource Kit	http://www.microsoft.com/learning/en/us/book.aspx?ID=6815&locale=en-us
Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step	http://technet.microsoft.com/en-us/library/dd378897(WS.10).aspx

Guide	
Windows Server Update Services	http://technet.microsoft.com/en-us/windowsserver/bb332157
Personal Virtual Desktops	http://technet.microsoft.com/en-us/library/dd759174.aspx
Read-Only Domain Controller Planning and Deployment Guide	http://technet.microsoft.com/en-us/library/cc771744(WS.10).aspx
Running Domain Controllers in Hyper-V	http://technet.microsoft.com/en-us/library/dd363553(v=ws.10).aspx
Hyper-V Security Guide	http://www.microsoft.com/en-us/download/details.aspx?id=16650
Ask the Directory Services Team	http://blogs.technet.com/b/askds/archive/2011/09/12/managing-rid-pool-depletion.aspx
How to configure a firewall for domains and trusts	http://support.microsoft.com/kb/179442
2009 Verizon Data Breach Report	http://www.verizonbusiness.com/resources/security/reports/2009_databreach_rp.pdf
2012 Verizon Data	http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf

Breach report	
Introducing Auditing Changes in Windows 2008	http://blogs.technet.com/b/askds/archive/2007/10/19/introducing-auditing-changes-in-windows-2008.aspx
Cool Auditing Tricks in Vista and 2008	http://blogs.technet.com/b/askds/archive/2007/11/16/cool-auditing-tricks-in-vista-and-2008.aspx
Global Object Access Auditing is Magic	http://blogs.technet.com/b/askds/archive/2011/03/10/global-object-access-auditing-is-magic.aspx
One-Stop Shop for Auditing in Windows Server 2008 and Windows Vista	http://blogs.technet.com/b/askds/archive/2008/03/27/one-stop-shop-for-auditing-in-windows-server-2008-and-windows-vista.aspx
AD DS Auditing Step-by-Step Guide	http://technet.microsoft.com/en-US/library/a9c25483-89e2-4202-881c-ea8e02b4b2a5.aspx
Getting the Effective Audit Policy in Windows 7 and 2008 R2	http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf
Sample script	http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf
Audit	http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-

Option Type	report-2012_en_xg.pdf
Advanced Security Auditing in Windows 7 and Windows Server 2008 R2	http://social.technet.microsoft.com/wiki/contents/articles/advanced-security-auditing-in-windows-7-and-windows-server-2008-r2.aspx
Auditing and Compliance in Windows Server 2008	http://technet.microsoft.com/en-us/magazine/2008.03.auditing.aspx
How to use Group Policy to configure detailed security auditing settings for Windows Vista-based and Windows Server 2008-based computers in a Windows Server 2008 domain, in a Windows Server 2003 domain, or in a Windows 2	http://support.microsoft.com/kb/921469

000 Server domain	
Advanced Security Audit Policy Step-by-Step Guide	http://technet.microsoft.com/en-us/library/dd408940(WS.10).aspx
Threats and Countermeasures Guide	http://technet.microsoft.com/en-us/library/hh125921(v=ws.10).aspx
MaxToken Size and Kerberos Token Bloat	http://blogs.technet.com/b/shanecothran/archive/2010/07/16/maxtokensize-and-kerberos-token-bloat.aspx
Authentication Mechanism Assurance	http://technet.microsoft.com/en-us/library/dd391847(v=WS.10).aspx
Microsoft Data Classification Toolkit	http://technet.microsoft.com/en-us/library/hh204743.aspx
Dynamic Access Control	http://blogs.technet.com/b/windowsserver/archive/2012/05/22/introduction-to-windows-server-2012-dynamic-access-control.aspx
Absolute Software	http://www.absolute.com/en/landing/Google/absolute-software-google/computrace-and-absolute-manage?gclid=CPPh5P6v3rMCFQtxQgodFEQAnA
Absolute Manage	http://www.absolute.com/landing/Google/absolute-manage-google/it-asset-management-software
Absolute Manage MDM	http://www.absolute.com/landing/Google/MDM-google/mobile-device-management
SolarWinds	http://www.solarwinds.com/eminentware-products.aspx
EminentWare WSUS	http://solarwinds-marketing.s3.amazonaws.com/solarwinds/Datasheets/EminentWare-WSUS-

Extension Pack	Extension-Pack-005-Datasheet2.pdf
EminentWare System Center Configuration Manager Extension Pack	http://solarwinds-marketing.s3.amazonaws.com/solarwinds/Datasheets/EminentWare-Extension-Pack-for-CM-Datasheet-006-Revised.pdf
GFI Software	http://www.gfi.com/?adv=952&loc=58&gclid=CLq9y5603rMCFal7QgodMFkAyA
GFI LanGuard	http://www.gfi.com/network-security-vulnerability-scanner/?adv=952&loc=60&gclid=CP2t-7i03rMCFQuCQgodNkAA7g
Secunia	http://secunia.com/
Secunia Corporate Software Inspector (CSI)	http://secunia.com/products/corporate/csi/
Vulnerability Intelligence Manager	http://secunia.com/vulnerability_intelligence/
eEye Digital Security	http://www.wideeyesecurity.com/?gclid=CK6b0sm13rMCFad_QgodhScAiw
Retina CS Management	http://www.wideeyesecurity.com/products.asp
Lumension	http://www.lumension.com/?rpLeadSourceId=5009&gclid=CKuai_e13rMCFal7QgodMFkAyA
Lumension Vulnerability Management	http://www.lumension.com/Solutions/Vulnerability-Management.aspx
Threats	http://technet.microsoft.com/en-us/library/hh125917(v=ws.10).aspx

and Countermeasures Guide: User Rights	
Threats and Vulnerabilities Mitigation	http://technet.microsoft.com/en-us/library/cc755181(v=ws.10).aspx
User Rights	http://technet.microsoft.com/en-us/library/dd349804(v=WS.10).aspx
Access Credential Manager as a trusted caller	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_2
Access this computer from the network	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_1
Act as part of the operating system	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_3
Add workstations to domain	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_4
Adjust memory quotas for a process	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_5
Allow log on locally	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_6
Allow log on through Terminal Services	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_7
Back up	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-

files and directories	e9845182f4b6(v=ws.10)#BKMK_8
Bypass traverse checking	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_9
Change the system time	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_10
Change the time zone	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_11
Create a pagefile	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_12
Create a token object	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_13
Create global objects	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_14
Create permanent shared objects	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_15
Create symbolic links	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_16
Debug programs	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_17
Deny access to this computer from the network	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_18
Deny log on as a batch job	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_18a
Deny log	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-

on as a service	e9845182f4b6(v=ws.10)#BKMK_19
Deny log on locally	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_20
Deny log on through Terminal Services	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_21
Enable computer and user accounts to be trusted for delegation	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_22
Force shutdown from a remote system	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_23
Generate security audits	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_24
Impersonate a client after authentication	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_25
Increase a process working set	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_26
Increase scheduling priority	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_27
Load and unload device drivers	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_28

Lock pages in memory	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_29
Log on as a batch job	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_30
Log on as a service	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_31
Manage auditing and security log	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_32
Modify an object label	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_33
Modify firmware environment values	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_34
Perform volume maintenance tasks	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_35
Profile single process	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_36
Profile system performance	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_37
Remove computer from docking station	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_38
Replace a process level token	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_39
Restore files and	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_40

directories	
Shut down the system	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_41
Synchronize directory service data	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_42
Take ownership of files or other objects	http://technet.microsoft.com/en-us/library/db585464-a2be-41b1-b781-e9845182f4b6(v=ws.10)#BKMK_43
Access Control	http://msdn.microsoft.com/en-us/library/aa374860(v=VS.85).aspx
Microsoft Support article 251343	http://support.microsoft.com/kb/251343
rootDSE Modify Operations	http://msdn.microsoft.com/en-us/library/cc223297.aspx
AD DS Backup and Recovery Step-by-Step Guide	http://technet.microsoft.com/en-us/library/cc771290(v=ws.10).aspx
Windows Configurations for Kerberos Supported Encryption Type	http://blogs.msdn.com/b/openspecification/archive/2011/05/31/windows-configurations-for-kerberos-supported-encryption-type.aspx
UAC Processes and Interactions	http://technet.microsoft.com/en-us/library/dd835561(v=WS.10).aspx#1

EmpowerID	http://www.empowerid.com/products/authorizationservices
Role-based access control (RBAC)	http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.security%2Fdoc%2Fsecurity%2Fdomain_rbac.htm
The RBAC model	http://docs.oracle.com/cd/E19082-01/819-3321/6n5i4b7ap/index.html
Active Directory-centric access control	http://www.centrify.com/solutions/it-security-access-control.asp
Cyber-Ark's Privileged Identity Management (PIM) Suite	http://www.cyber-ark.com/digital-vault-products/pim-suite/index.asp
Quest One	http://www.quest.com/landing/?id=7370&gclid=CJnNgNyr3rMCFYp_QgodXFwA3w
Enterprise Random Password Manager (ERPM)	http://www.liebsoft.com/Random_Password_Manager/
NetIQ Privileged User Manager	https://www.netiq.com/products/privileged-user-manager/
CA IdentityMinder™	http://awards.scmagazine.com/ca-technologies-ca-identity-manager
Description of security events in Windows Vista and in Windows Server 200	http://support.microsoft.com/kb/947226

8	
Description of security events in Windows 7 and in Windows Server 2008 R2	http://support.microsoft.com/kb/977519
Security Audit Events for Windows 7	http://www.microsoft.com/en-us/download/details.aspx?id=21561
Windows Server 2008 R2 and Windows 8 and Windows Server 2012 Security Event Details	http://www.microsoft.com/en-us/download/details.aspx?id=35753
Georgia Tech's Emerging Cyber Threats for 2013 report	http://www.gtsecuritysummit.com/report.html
Microsoft Security Intelligence Report	http://www.microsoft.com/security/sir/default.aspx
Australian Government Defense Signals Directory Top 35	http://www.dsd.gov.au/infosec/top35mitigationstrategies.htm

Mitigation Strategies	
Cloud Computing Security Benefits	http://www.microsoft.com/en-us/news/Press/2012/May12/05-14SMBSecuritySurveyPR.aspx
Applying the Principle of Least Privilege to User Accounts on Windows	http://www.microsoft.com/en-us/download/details.aspx?id=4868
The Administrator Accounts Security Planning Guide	http://www.microsoft.com/en-us/download/details.aspx?id=19406
Best Practice Guide for Securing Active Directory Installations for Windows Server 2003	http://www.microsoft.com/en-us/download/details.aspx?id=16755
Best Practices for Delegating Active Directory Administration for	http://www.microsoft.com/en-us/download/details.aspx?id=21678

Windows Server 2003	
Microsoft Support Lifecycle	http://support.microsoft.com/common/international.aspx?RDPATH=%2flifecycle%2fdefault.aspx
Active Directory Technical Specification	http://msdn.microsoft.com/en-us/library/cc223122(v=prot.20).aspx
Error message when nonadministrator users who have been delegated control try to join computers to a Windows Server 2003-based or a Windows Server 2008-based domain controller: “Access is denied”	http://support.microsoft.com/kb/932455
Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-	http://technet.microsoft.com/en-us/library/dd378897(WS.10).aspx

by-Step Guide	
Strict KDC Validation	http://www.microsoft.com/en-us/download/details.aspx?id=6382

Recommended Reading

The following table contains a list of recommended reading that will assist you in enhancing the security of your Active Directory systems.

Recommended Reading
Georgia Tech's Emerging Cyber Threats for 2014 Report
Microsoft Security Intelligence Report
Mitigating Pass-the-Hash (PTH) Attacks and Other Credential Theft Techniques
Australian Government Defense Signals Directory Top 35 Mitigation Strategies
2012 Data Breach Investigations Report - (Verizon, US Secret Service)
2009 Data Breach Investigations Report
Cloud Computing Security Benefits
Applying the Principle of Least Privilege to User Accounts on Windows
The Administrator Accounts Security Planning Guide
Best Practice Guide for Securing Active Directory Installations for Windows Server 2003
Best Practices for Delegating Active Directory Administration for Windows Server 2003
Microsoft Support Lifecycle
Active Directory Technical Specification - dSHeuristics information
Error message when nonadministrator users who have been delegated control try to join

computers to a Windows Server 2003-based or a Windows Server 2008-based domain controller: “Access is denied”
Best Practice Guide for Securing Active Directory Installations.doc
Hyper-V Security Guide
Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide.
Strict KDC Validation

Copyright Information

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. Microsoft makes no warranties, express or implied, in this document.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Microsoft, Active Directory, BitLocker, Hyper-V, Internet Explorer, Windows Vista, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are property of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

© 2013 Microsoft Corporation. All rights reserved.

Active Directory Domain Services Component Updates for Windows Server 2012 R2

This module introduces the components that received minor updates in the Directory Services and Identity spaces.

About the Author	
Author:	Justin Turner
Bio:	Justin is a Senior Support Escalation Engineer with the Directory Services team based in Irving, Texas, USA. He has created or contributed to many training courses and KB articles for the Microsoft Knowledgebase over the past 12 years. He teaches Microsoft employees and customers new product architecture, is a charter Microsoft Certified Master (MCM), Microsoft Certified Trainer (MCT) and holds an M.S. degree in Computer Education and Cognitive Systems.
Contributors	This training module includes contributions from *Michiko Short*, *Dean Wells*, *Alan Jowett*, *Manu Pushpendran*, *Yashar Bahman*, *Anoosh Saboori*, *Rashmi Jha*, *Justin Hall* and *Herbert Mauerer*
Reviewers	Many thanks to the following individuals that spent their own time reviewing and providing feedback: *Joey Seifert*, *Justin Hall*

Note

This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

What You Will Learn

After completing this module, you will be able to:

- Explain the component updates made within the Directory Services and Identity technology areas in Windows Server 2012 R2
  - SPN and UPN uniqueness
  - Winlogon Automatic Restart Sign-On (ARSO)
  - TPM Key Attestation
  - CA Backup and Restore Windows PowerShell cmdlets
  - Command line process auditing
  - **Credential Protection and Domain Authentication Controls**
  - Directory Services component updates
    - Domain and Forest Functional Levels
    - Deprecation of NTFRS
    - LDAP Query Optimizer changes
    - 1644 Event improvements
    - Active Directory Replication throughput improvement

Identity component updates

Lesson 1: Identity component updates

This lesson explains the Identity component updates in Windows Server 2012 R2.

What You Will Learn

After completing this lesson, you will be able to:

- Describe the following changes:
  - SPN and UPN uniqueness
  - Winlogon Automatic Restart Sign-On (ARSO)
  - TPM Key Attestation
  - CA Backup and Restore Windows PowerShell cmdlets
  - Command line process auditing
  - **Credential Protection and Domain Authentication Controls**

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

SPN and UPN uniqueness

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

Overview

Domain Controllers running Windows Server 2012 R2 block the creation of duplicate service principal names (SPN) and user principal names (UPN). This includes if the restoration or reanimation of a deleted object or the renaming of an object would result in a duplicate.

Background

Duplicate Service Principal Names (SPN) commonly occur and result in authentication failures and may lead to excessive LSASS CPU utilization. There is no in-box method to block the addition of a duplicate SPN or UPN. *

Duplicate UPN values break synchronization between on-premises AD and Office 365.

*Setspn.exe is commonly used to create new SPNs, and functionally was built into the version released with Windows Server 2008 that adds a check for duplicates.

Table 1: UPN and SPN uniqueness

Feature	Comment
UPN uniqueness	Duplicate UPNs break synchronization of on-premises AD accounts with Windows Azure AD-based services such as Office 365.
SPN uniqueness	Kerberos requires SPNs for mutual authentication. Duplicate SPNs result in authentication failures.

Symptoms

Error codes 8467 or 8468 or their hex, symbolic or string equivalents are logged in various on-screen dialogues and in event ID 2974 in the Directory Services event log. The attempt to create a duplicate UPN or SPN is blocked only under the following circumstances:

- The write is processed by a Windows Server 2012 R2 DC

Table 2: UPN and SPN uniqueness error codes

Decimal	Hex	Symbolic	String
8467	21C7	ERROR_DS_SPN_VALUE_NOT_UNIQUE_IN_FOREST	The operation failed because SPN value provided for addition/modification is not unique forest-wide.
8648	21C8	ERROR_DS_UPN_VALUE_NOT_UNIQUE_IN_FOREST	The operation failed because UPN value provided for addition/modification is not unique forest-wide.

New user creation fails if UPN is not unique

DSA.msc

The user logon name you have chosen is already in use in this enterprise. Choose another logon name, and then try again.

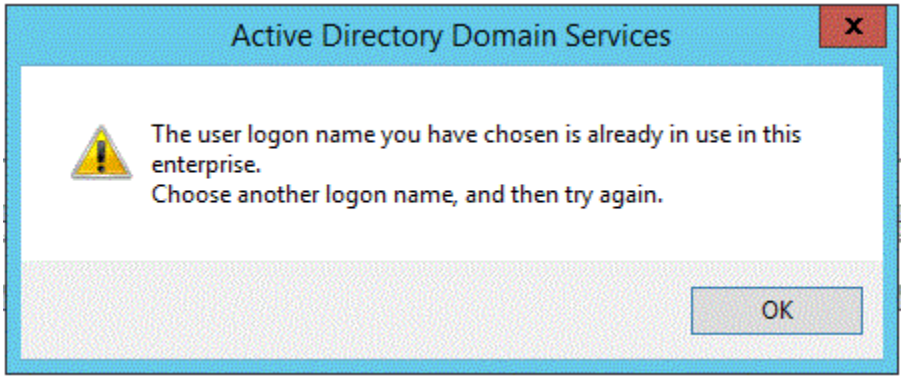

Modify an existing account:

The specified user logon name already exists in the enterprise. Specify a new one, either by changing the prefix or selecting a different suffix from the list.

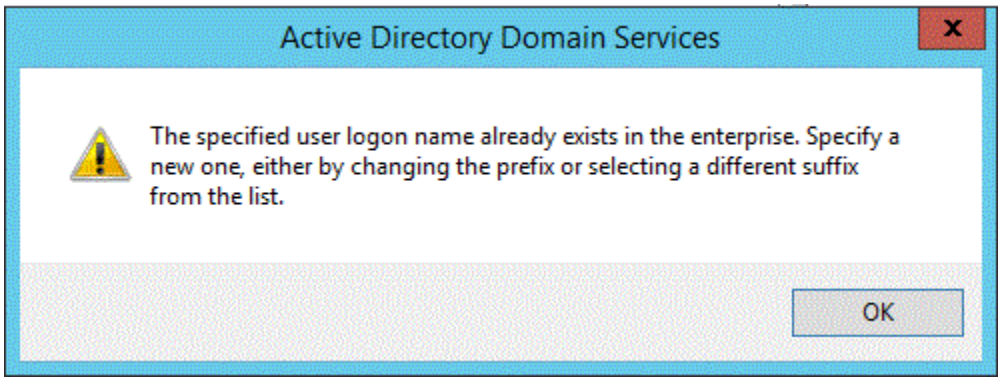

Active Directory Administrative Center (DSAC.exe)

An attempt to create a new user in Active Directory Administrative Center with a UPN that already exists will yield the following error.

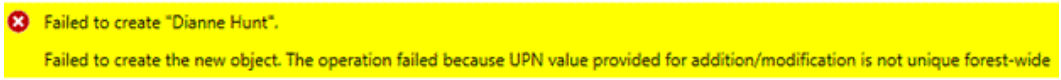

Figure 1 error displayed in AD Administrative Center when new user creation fails due to duplicate UPN

Event 2974 Source: ActiveDirectory_DomainService

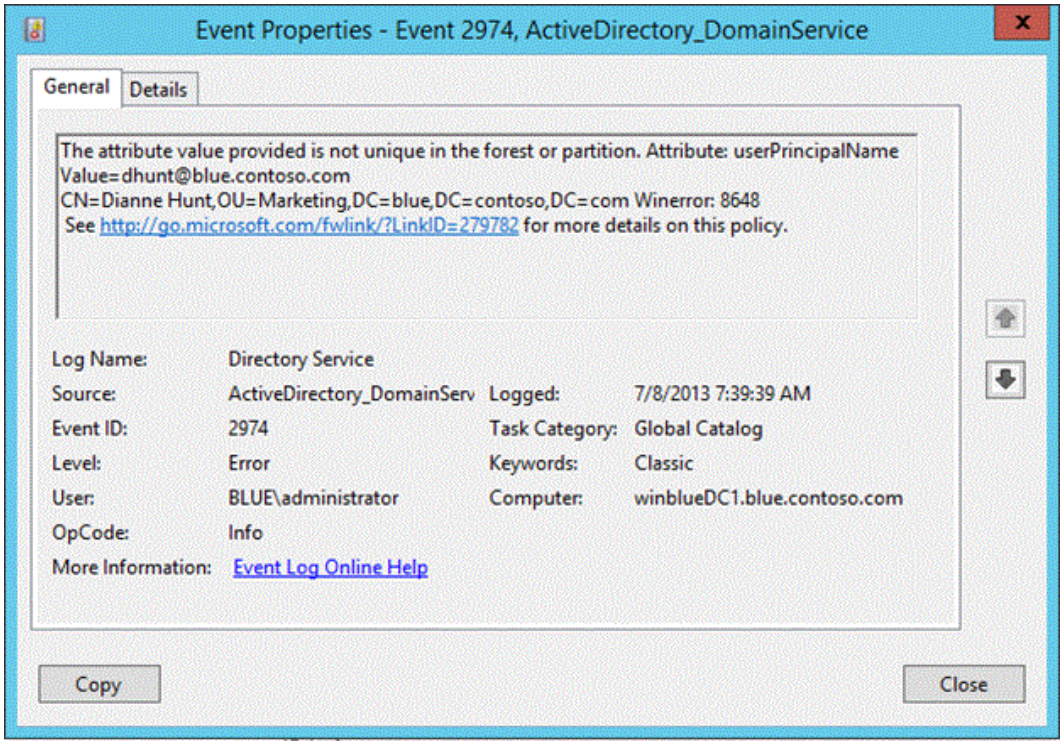

Figure 2 Event ID 2974 with error 8648

The event 2974 lists the value that was blocked and a list of one or more objects (up to 10) that already contain that value. In the following figure, you can see that UPN attribute value ***dhunt@blue.contoso.com*** already exists on four other objects. Since this is a new feature in Windows Server 2012 R2, accidental creation of duplicate UPN and SPNs in a mixed environment will still occur when down-level DCs process the write attempt.

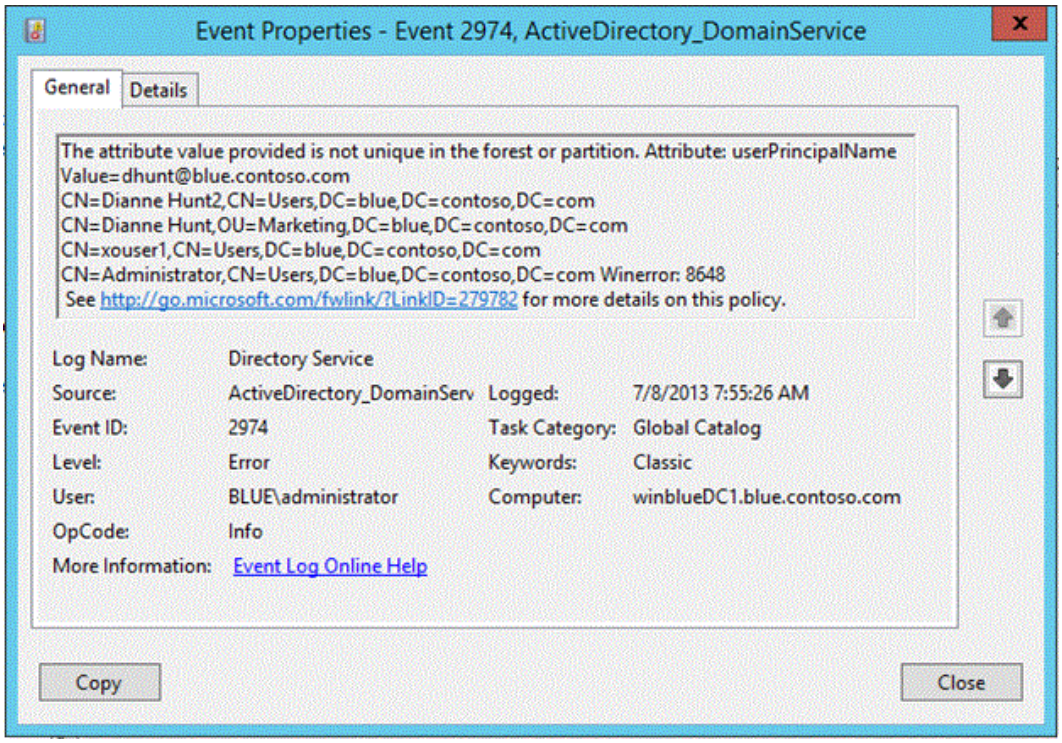

Figure 3 Event 2974 showing all objects containing the duplicate UPN

Tip

Review event ID 2974s regularly to:

- identify attempts to create duplicate UPN or SPNs
- identify objects that already contain duplicates

8648 = "The operation failed because UPN value provided for addition/modification is not unique forest-wide."

SetSPN:

Setspn.exe has had duplicate SPN detection built-in to it since the Windows Server 2008 release when using the **"-S"** option. You can bypass the duplicate SPN detection by using the **"-A"** option however. Creation of a duplicate SPN is blocked when targeting a Windows Server 2012 R2 DC using SetSPN with the -A option. The error message displayed is the same as the one displayed when using the -S option: "Duplicate SPN found, aborting operation!"

ADSIEDIT:

```
Operation failed. Error code: 0x21c8

The operation failed because UPN value provided for addition/modification is not unique
forest-wide.

000021C8: AtrErr: DSID-03200BBA, #1: 0: 000021C8: DSID-03200BBA, problem 1005
(CONSTRAINT_ATT_TYPE), data 0, Att 90290 (userPrincipalName)
```

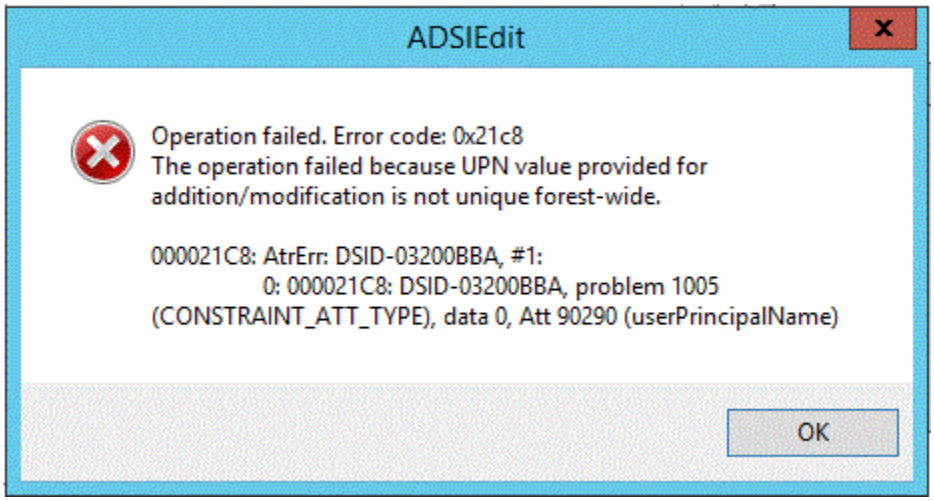

Figure 4 Error message displayed in ADSIEdit when addition of duplicate UPN is blocked

Windows PowerShell

Windows Server 2012 R2:

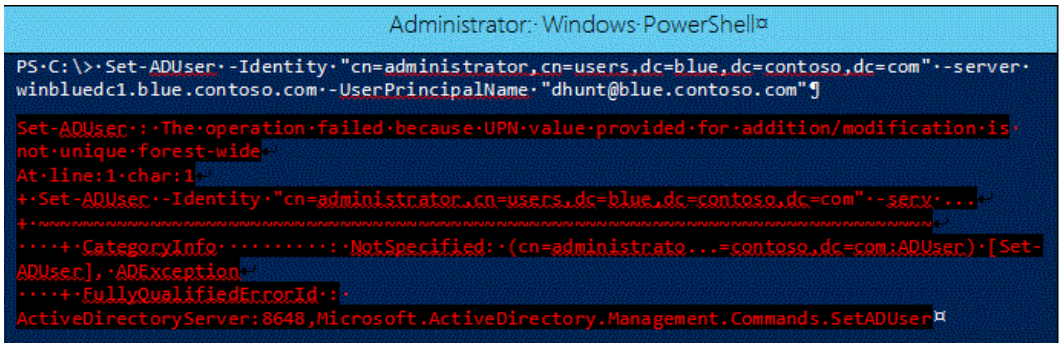

PS running from Server 2012 targeting a Windows Server 2012 R2 DC:

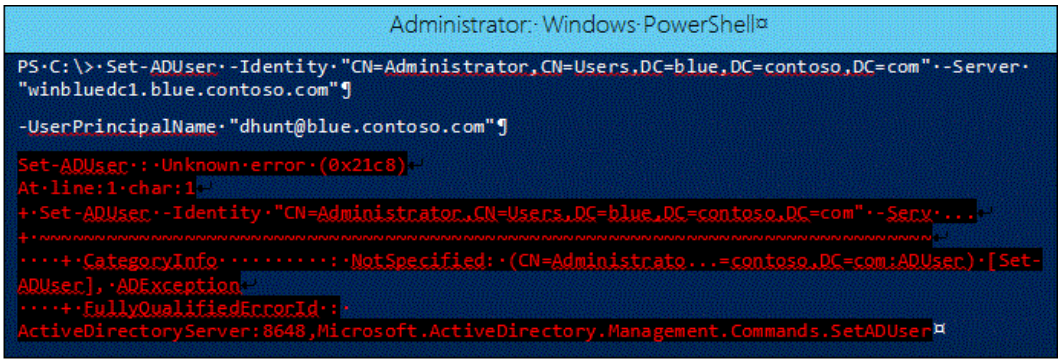

DSAC.exe running on Windows Server 2012 targeting a Windows Server 2012 R2 DC:

Figure 5 DSAC user creation error on non-Windows Server 2012 R2 while targeting Windows Server 2012 R2 DC

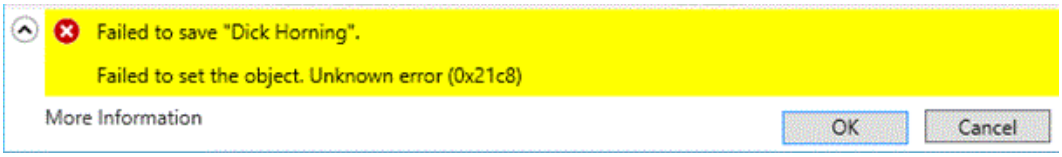

Figure 6 DSAC user modification error on non-Windows Server 2012 R2 while targeting Windows Server 2012 R2 DC

Restore of an object that would result in a duplicate UPN fails:

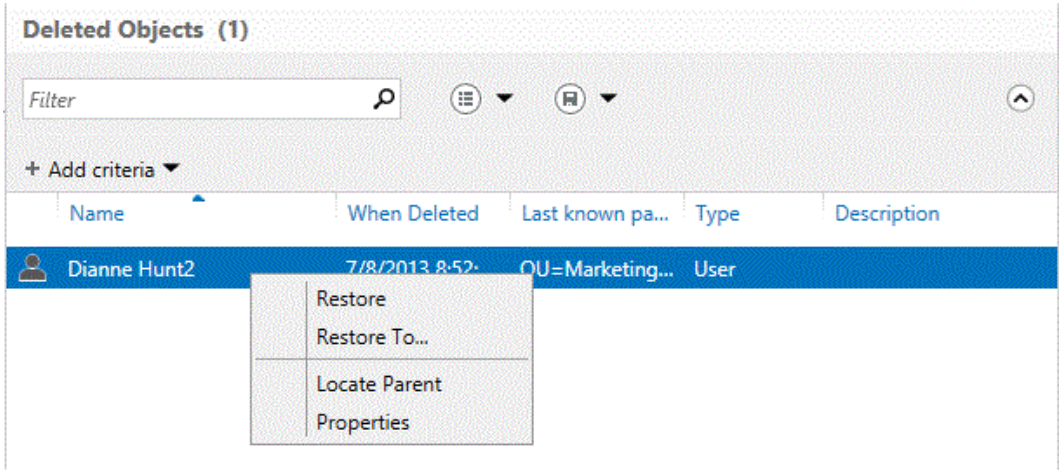

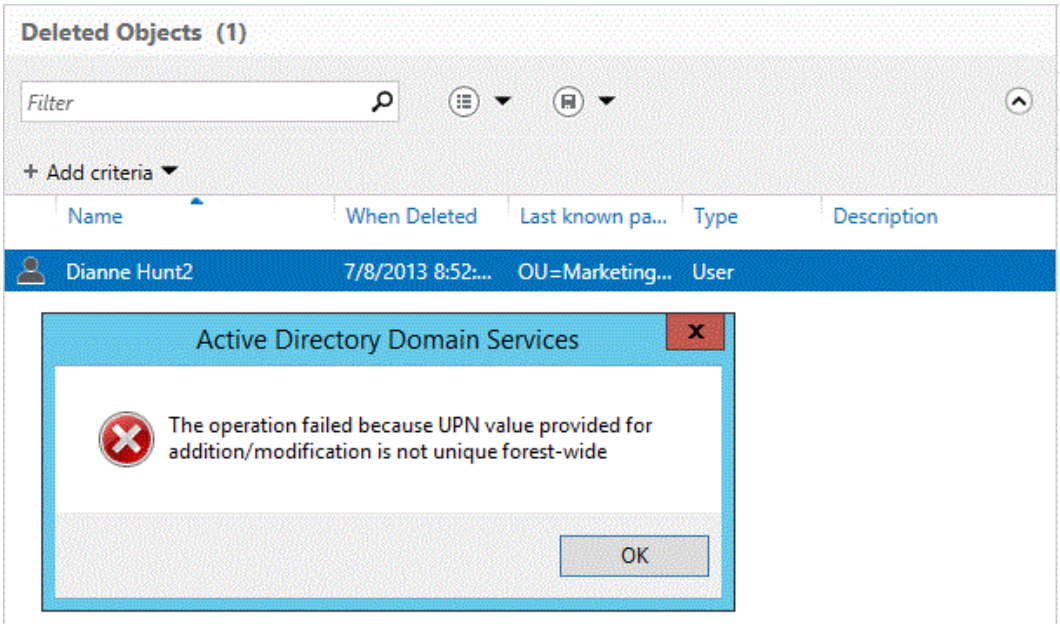

No event is logged when an object fails to restore because of a duplicate UPN / SPN.

The UPN of the object must be unique in order for it to be restored.

1. Identify the UPN that exists on the object in the Recycle Bin
2. Identify all objects that have the same value
3. Remove the duplicate UPN(s)

Identify the conflicting UPN on the deleted objectUsing repadmin.exe

```
Repadmin /showattr DCName "DN of deleted objects container" /subtree /filter:"(msDS-LastKnownRDN=<NAME>)" /deleted /atts:userprincipalname
```

```
repadmin /showattr DCName "CN=Deleted Objects,DC=blue,DC=contoso,DC=com" /subtree /filter:"(msDS-LastKnownRDN=Dianne Hunt2)" /deleted /atts:userprincipalname
```

```
C:\>repadmin /showattr winbluedc1 "cn=deleted objects,dc=blue,dc=contoso,dc=com" /subtree /filter:"(msds-lastknownrdn=Dianne Hunt2)" /deleted /atts:userprincipalname

DN: CN=Dianne Hunt2\0ADEL:dd3ab8a4-3005-4f2f-814f-d6fc54a1a1c0,CN=Deleted Objects,DC=blue,DC=contoso,DC=com

    1> userPrincipalName: dhunt@blue.contoso.com
```

To identify all objects with the same UPN:Using Repadmin.exe

```
repadmin /showattr WinBlueDC1 "DC=blue,DC=contoso,DC=com" /subtree /filter:"(userPrincipalName=dhunt@blue.contoso.com)" /deleted /atts:DN
```

```
C:\>repadmin /showattr winbluedc1 "dc=blue,dc=contoso,dc=com" /subtree /filter:"(userPrincipalName=dhunt@blue.contoso.com)" /deleted /atts:DN

DN: CN=Administrator,CN=Users,DC=blue,DC=contoso,DC=com

DN: CN=xouser1,CN=Users,DC=blue,DC=contoso,DC=com

DN: CN=xouser10,CN=Users,DC=blue,DC=contoso,DC=com

DN: CN=xouser100,CN=Users,DC=blue,DC=contoso,DC=com

DN: CN=Dianne Hunt,OU=Marketing,DC=blue,DC=contoso,DC=com

DN: CN=Dianne Hunt2\0ADEL:dd3ab8a4-3005-4f2f-814f-d6fc54a1a1c0,CN=Deleted Objects,DC=blue,DC=contoso,DC=com
```

Tip

The previously undocumented **/deleted** parameter in repadmin.exe is used to include deleted objects in the result set

Using Global Search

- Open Active Directory Administrative Center and navigate to **Global Search**
- Select the **Convert to LDAP** radio button
- Type **(userPrincipalName=***ConflictingUPN***)**
  - Replace *ConflictingUPN* with the actual UPN that is in conflict
- Select **Apply**

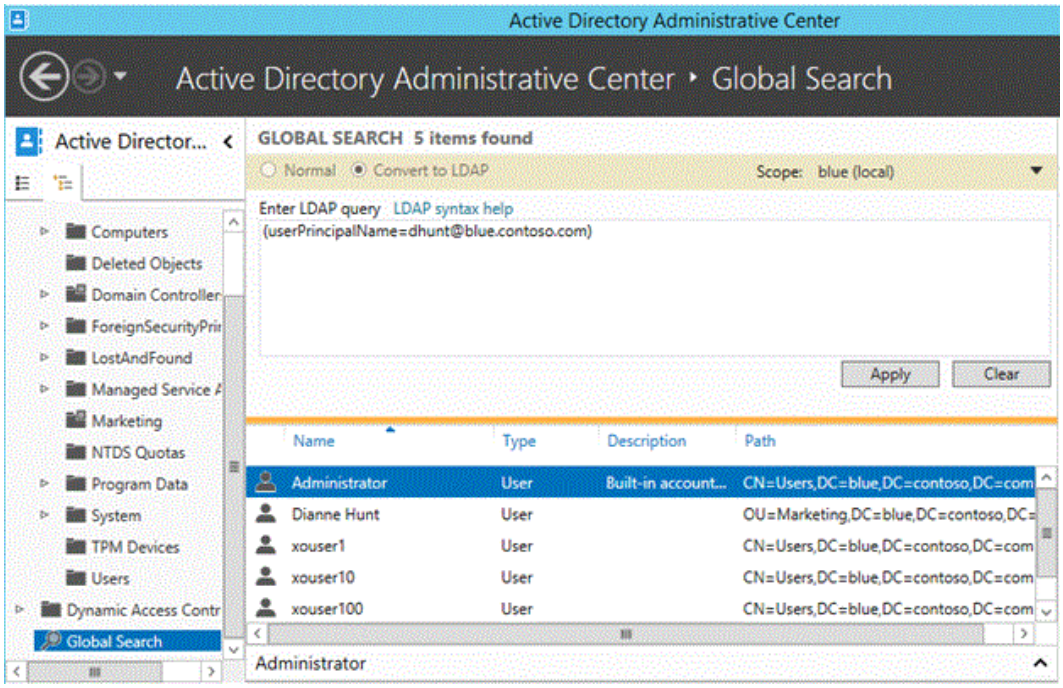

Using Windows PowerShell

```
Get-ADObject -LdapFilter "(userPrincipalName=dhunt@blue.contoso.com)" -
IncludeDeletedObjects -SearchBase "DC=blue,DC=Contoso,DC=com" -SearchScope Subtree -
Server winbluedc1.blue.contoso.com
```

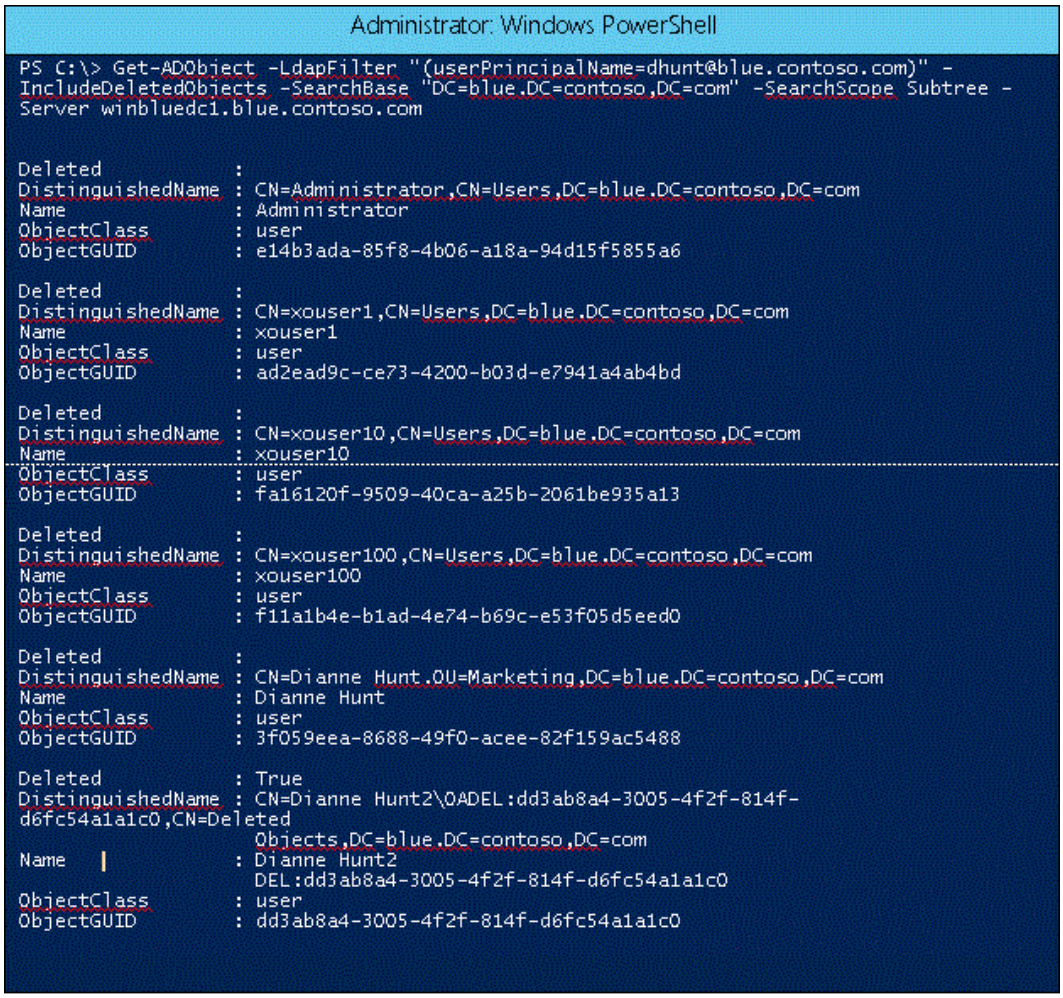

If the object needs to be restored, you will need remove the duplicate UPNs from the other objects. For only one object, it is simple enough to use ADSIEdit to remove the duplicate. If there are multiple objects with duplicates, then Windows PowerShell might be the better tool to use.

To null out the UserPrincipalName attribute using Windows PowerShell:

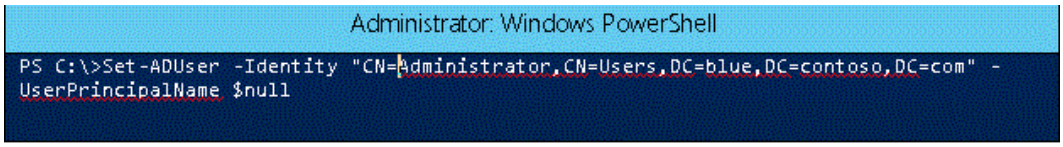

Note

The userPrincipalName attribute is single-valued attribute, so this procedure will only remove the duplicate UPN.

Duplicate SPN

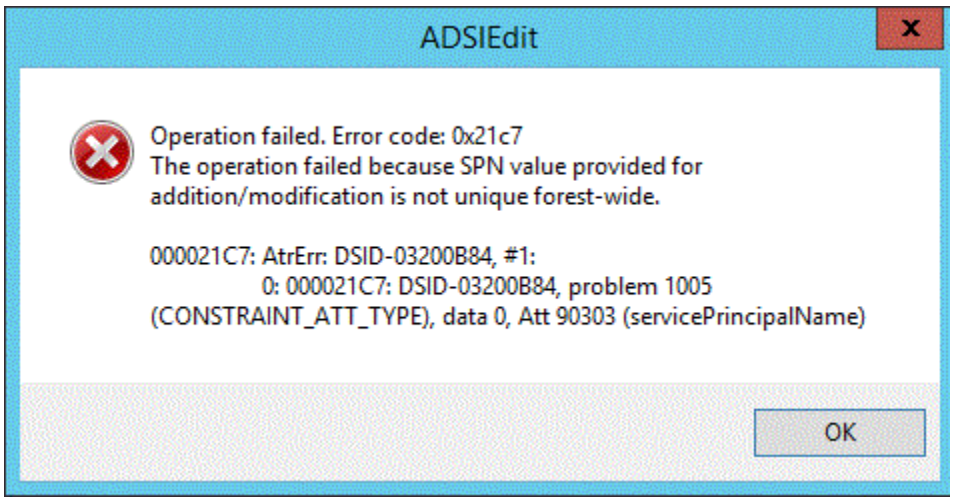

Figure 8 Error message displayed in ADSIEdit when addition of duplicate SPN is blocked

Logged in the Directory Services event log is an **ActiveDirectory_DomainService** event ID **2974**.

```
Operation failed. Error code: 0x21c7

The operation failed

The attribute value provided is not unique in the forest or partition. Attribute:

servicePrincipalName Value=<SPN>

<Object DN> Winerror: 8467
```

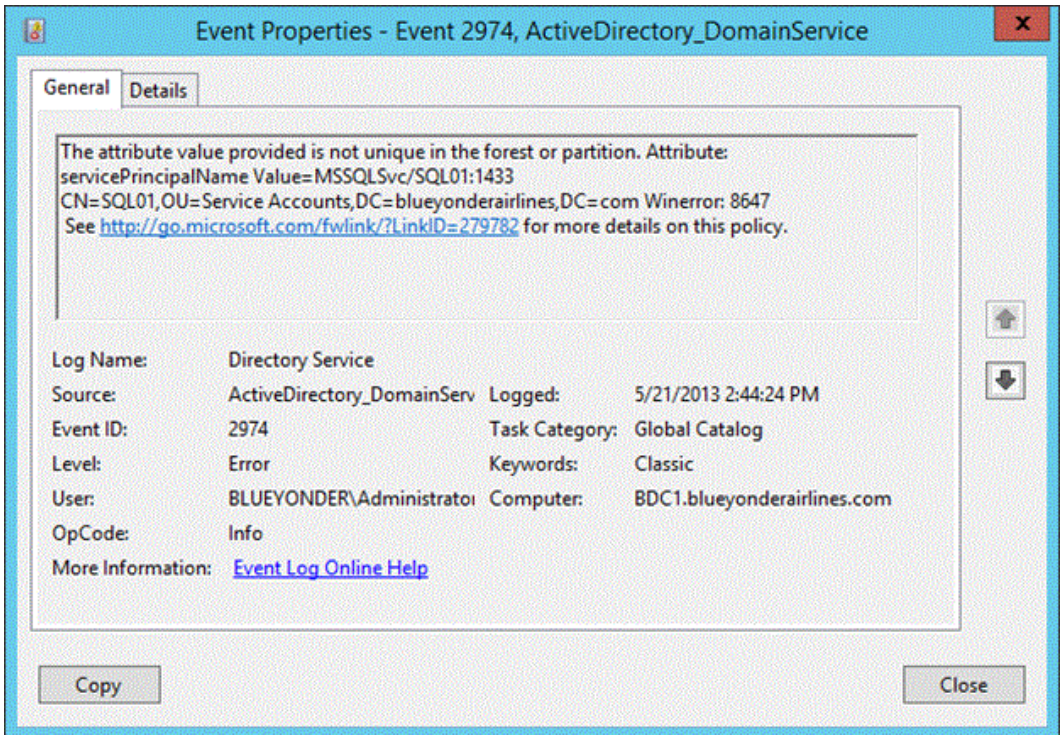

Figure 9 Error logged when creation of duplicate SPN is blocked

Workflow

- **If DC == GC**
  - No offbox call required, query can be satisfied locally
  - ***UPN case***
    - Query local forest-wide UPN index for supplied UPN (*userPrincipalName; a global index*)
      - If entries returned == 0 -> write proceeds
      - If entries returned !=0 -> write fails
        - Event logged
        - Also returns extended error:

        8648:

        ERROR_DS_UPN_VALUE_NOT_UNIQUE_IN_FOREST

  - ***SPN case***
    - Query local forest-wide SPN index for supplied SPN (*servicePrincipalName; a global index*)
      - If entries returned == 0 -> write proceeds
      - If entries returned !=0 -> write fails

- Event logged
- Also returns extended error:

8647:

ERROR_DS_SPN_VALUE_NOT_UNIQUE_IN_FOREST

- **If DC != GC**
  - Offbox call desirable but not critical, i.e. this is a best-effort uniqueness check
    - Check proceeds against local DIT only if GC cannot be located
    - Event logged to indicate such
  - ***UPN case***
    - Submit LDAP query against closest GC query GC's forest-wide UPN index for supplied UPN (*userPrincipalName; a global index*)
      - If entries returned == 0 -> write proceeds
      - If entries returned !=0 -> write fails
        - Event logged
        - Also returns extended error:

      8648:

      ERROR_DS_UPN_VALUE_NOT_UNIQUE_IN_FOREST

  - ***SPN case***
    - Submit LDAP query against closest GC query GC's forest-wide SPN index for supplied SPN (*servicePrincipalName; a global index*)
      - If entries returned == 0 -> write proceeds
      - If entries returned !=0 -> write fails
        - Event logged
        - Also returns extended error:

      8647:

      ERROR_DS_SPN_VALUE_NOT_UNIQUE_IN_FOREST

When deleted objects are re-animated, SPN or UPN values present are checked for uniqueness. If a duplicate is found, the request fails.

- For certain attribute changes like DNS Host Name, SAM Account Name etc, when the modification is made, SPNs are updated accordingly. In the process, the obsolete SPNs are deleted and new SPNs are constructed and added to the database. The requisite attribute modifications against which this path is triggered are:
  - ATT_DNS_HOST_NAME
  - ATT_MS_DS_ADDITIONAL_DNS_HOST_NAME
  - ATT_SAM_ACCOUNT_NAME

- ATT_MS_DS_ADDITIONAL_SAM_ACCOUNT_NAME
- ATT_SERVER_REFERENCE_BL
- ATT_USER_ACCOUNT_CONTROL

If any of the new SPN value is a duplicate, we fail the modification. Of the above list, the important attributes are ATT_DNS_HOST_NAME (Machine name) and ATT_SAM_ACCOUNT_NAME (SAM Account Name).

Try This: Exploring SPN and UPN uniqueness

This is the first of several "**Try This**" activities in the module. There is not a separate lab guide for this module. The **Try This** activities are essentially free-form activities that allow you explore the lesson material in the lab environment. You have the option of following the prompt or going off script and come up with your own activity.

Notes

- This is the first of several "**Try This**" activities.
- There is not a separate lab guide for this module.
- The **Try This** activities are essentially free-form activities that allow you explore the lesson material in the lab environment.
- You have the option of following the prompt or going off script and come up with your own activity.
- While not all sections have a **Try This** prompt, you are still encouraged to explore the lesson content in the lab where appropriate.

Experiment with SPN and UPN uniqueness. Follow these prompts, or complete your own.

1. Create new users with UPN
2. Create accounts with SPNs
3. Either create a new user with a UPN already previously defined or change an existing account's UPN. Do the same for a SPN on another account
    a. Populate an existing user account with a UPN already in use
        i. Using PowerShell, ADSIEDIT, or Active Directory Administrative Center (DSAC.exe)
    b. Populate an existing account with an SPN already in use
        i. Using Windows PowerShell, ADSIEDIT, or SetSPN
4. Observe the errors

Optionally

1. Verify with the classroom instructor that it is ok to enable the AD Recycle Bin in Active Directory Administrative Center. If so, move on to the next step.
2. Populate the UPN on a user account
3. Delete the account
4. Populate a different account with the same UPN as the deleted account
5. Attempt to use the Recycle Bin GUI to restore the account

6. Imagine you have just been presented with the error you see in the previous step. (and don't have a history of the steps you just performed)Your goal is to complete the restore of the account. See the workbook for example steps.

Winlogon Automatic Restart Sign-On (ARSO)

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

Overview

Windows 8 introduced lock screen apps. These are the applications that run and display notifications while the user's session is locked (calendar appointments, email and messages, etc.). Devices that are restarted due to the Windows Update process fail to display these lock screen notifications upon restart. Some users depend on these lock screen applications.

What's changed?

When a user signs in on a Windows 8.1 device, LSA will save the user credentials in encrypted memory accessible only by lsass.exe. When Windows Update initiates an automatic reboot without user presence, these credentials will be used to configure Autologon for the user. Windows Update running as system with TCB privilege will initiate the RPC call to do this.

On rebooting, the user will automatically be signed in via the Autologon mechanism and then additionally locked to protect the user’s session. The locking will be initiated via Winlogon whereas the credential management is done by LSA. By automatically signing on and locking the user on the console, the user’s lock screen applications will be restarted and available.

Note
After a Windows Update induced reboot, the last interactive user is automatically signed on and the session is locked so the user's lock screen apps can run.

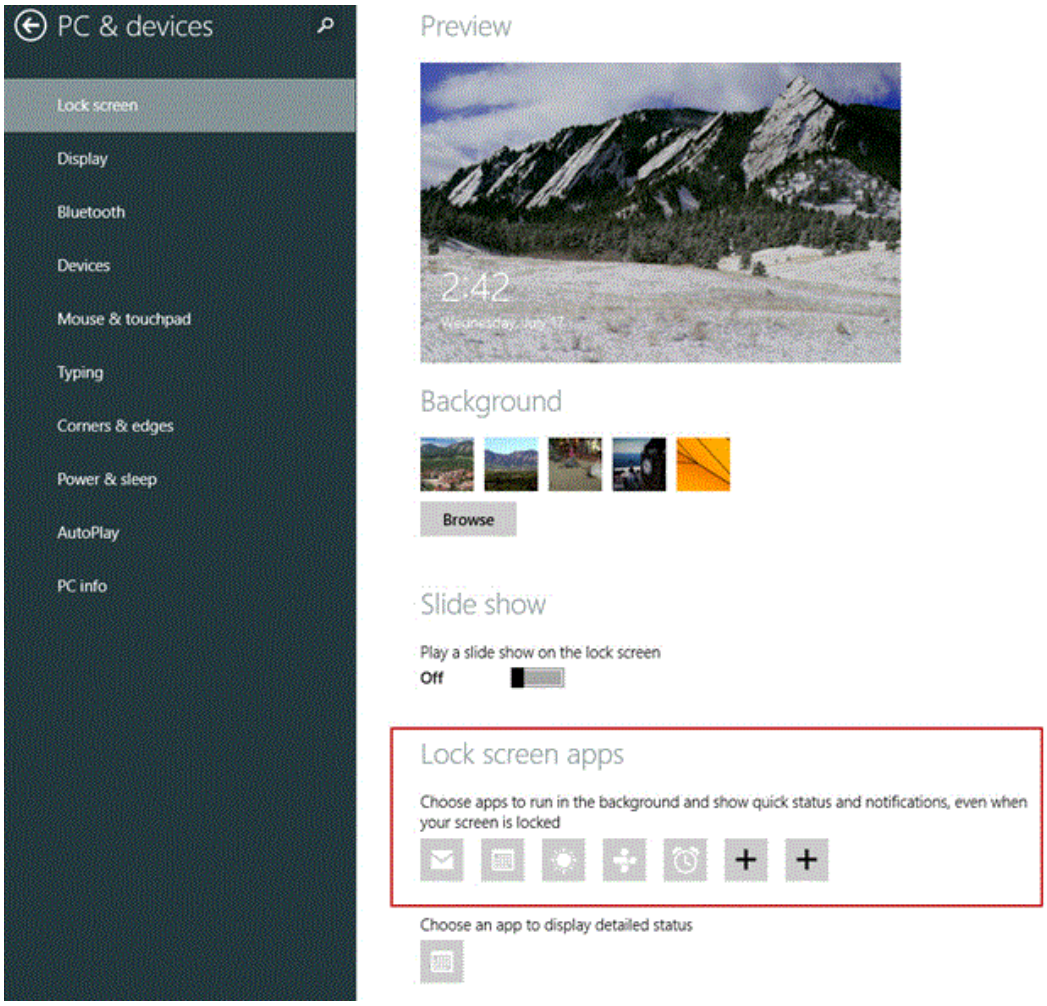
PC & devices
Lock screen
Display
Bluetooth
Devices
Mouse & touchpad
Typing
Corners & edges
Power & sleep
AutoPlay
PC info
Preview
2:42
Background
Browse
Slide show
Play a slide show on the lock screen
Off
Lock screen apps
Choose apps to run in the background and show quick status and notifications, even when your screen is locked
Choose an app to display detailed status

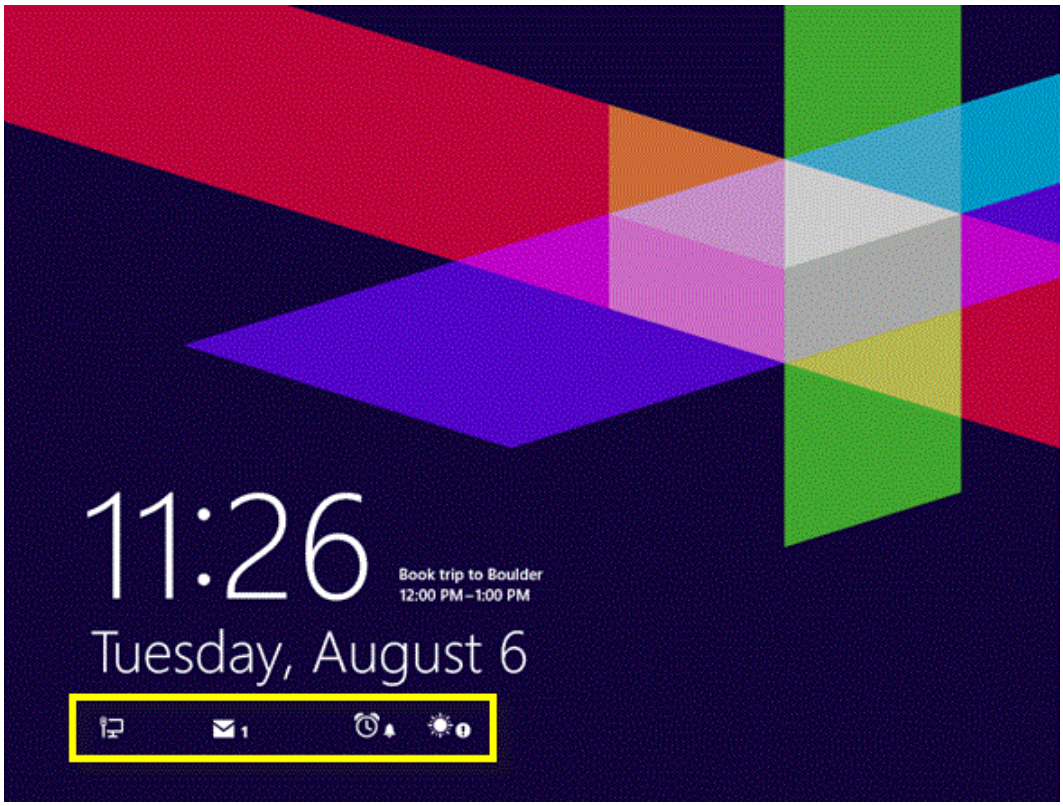

Quick Overview

- Windows Update requires restart
- Is computer able to restart (*no apps running that would lose data*)?
  - Restart for you
  - Log back in
  - Lock machine
- Enabled or disabled by Group Policy
  - Disabled by default in server SKUs
- Why?
  - Some updates cannot finish until the user logs back in.
  - Better user experience: don't have to wait 15 minutes for updates to finish installing
- How? AutoLogon
  - stores password, uses that credential to log you in
  - saves credential as an LSA secret in paged memory
  - Can only be enabled if BitLocker is enabled

Group Policy: Sign-in last interactive user automatically after a system-initiated restart

In Windows 8.1 / Windows Server 2012 R2, autologon of the lock screen user after a Windows Update restart is opt in for Server SKUs and opt out for Client SKUs.

Policy location: Computer Configuration > Policies > Administrative Templates > Windows Components > Windows Logon Option

Policy Name: Sign-in last interactive user automatically after a system-initiated restart

Supported on: At least Windows Server 2012 R2, Windows 8.1 or Windows RT 8.1

Description/Help:

This policy setting controls whether a device will automatically sign-in the last interactive user after Windows Update restarts the system.

If you enable or do not configure this policy setting, the device securely saves the user's credentials (including the user name, domain, and encrypted password) to configure automatic sign-in after a Windows Update restart. After the Windows Update restart, the user is automatically signed-in and the session is automatically locked with all the lock screen apps configured for that user.

If you disable this policy setting, the device does not store the user's credentials for automatic sign-in after a Windows Update restart. The users' lock screen apps are not restarted after the system restarts.

Registry Editor		
Value Name	Type	Data
DisableAutomaticRestartSignOn	DWORD	0 **Example:** 0 (Enabled) 1 (Disabled)
HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System		

Policy Registry Location:
HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System

Type: DWORD

Registry Name: DisableAutomaticRestartSignOn

Value: 0 or 1

0 = Enabled

1 = Disabled

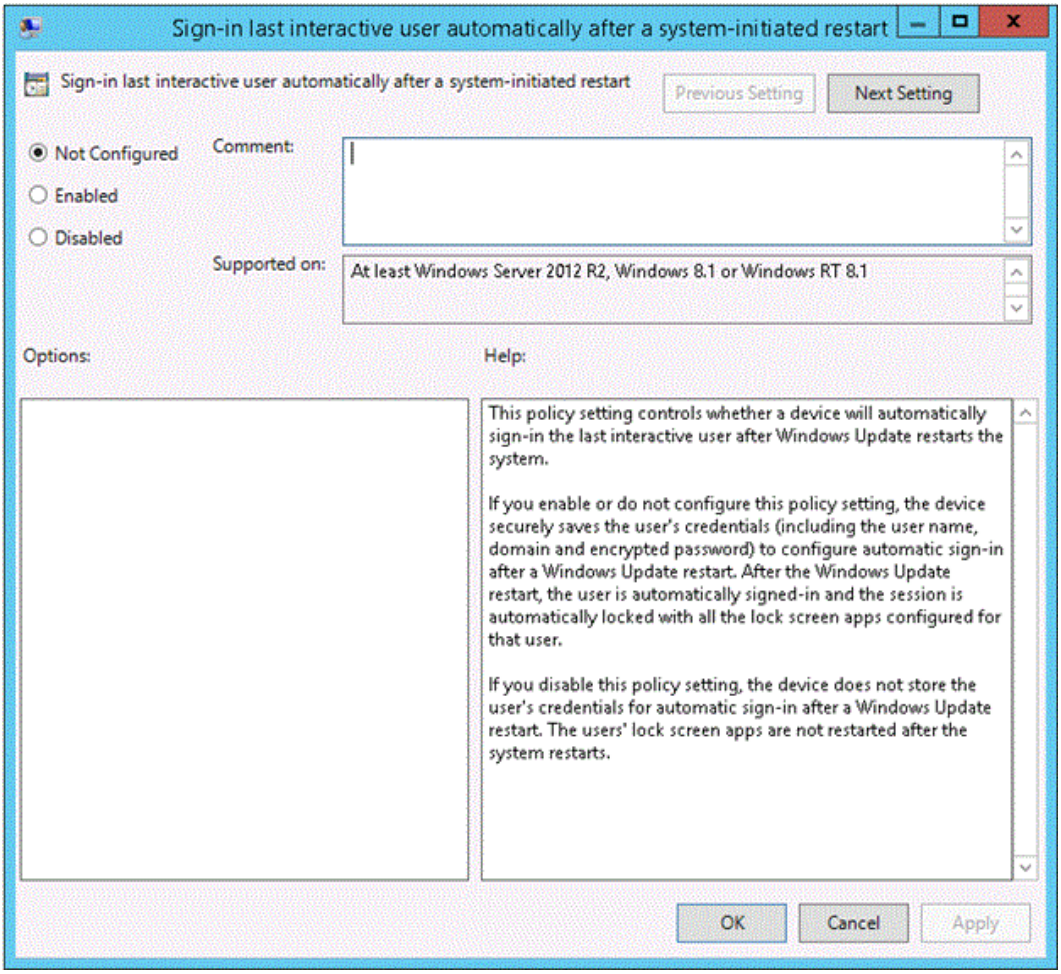

Troubleshooting

When WinLogon automatically locks, WinLogon's state trace will be stored in the WinLogon event log.

The status of an Autologon configuration attempt is logged

- If it is successful
  - records it as such
- If it is a failure:
  - records what the failure was
- When BitLocker's state changes:
  - the removal of credentials will be logged
    - These will be stored in the LSA Operational log.

Reasons why autologon might fail

There are several cases in which a user automatic login cannot be achieved. This section is intended to capture the known scenarios in which this can occur.

User Must Change Password at Next Login

User login can enter a blocked state when password change at next login is required. This can be detected prior to restart in most cases, but not all (for example, password expiry can be reached between shutdown and next login.

User Account Disabled

An existing user session can be maintained even if disabled. Restart for an account that is disabled can be detected locally in most cases in advance, depending on gp it may not be for domain accounts (some domain cached login scenarios work even if account is disabled on DC).

Logon Hours and Parental Controls

The Logon Hours and parental controls can prohibit a new user session from being created. If a restart were to occur during this window, the user would not be permitted to login. There is additional policy that causes lock or logout as a compliance action. This could be problematic for many child cases where account lockdown may occur between bed time and wake-up, particularly if the maintenance window is commonly during this time.

Additional Resources

Table 3: ARSO Glossary

Term	Definition
Autologon	Autologon is a feature that has been present in Windows for several releases. It is a documented feature of Windows that even has tools such as Autologon for Windows v3.01 http:/technet.microsoft.com/en-us/sysinternals/bb963905.aspx It allows a single user of the device to sign in automatically without entering credentials. The credentials are configured and stored in registry as an encrypted LSA secret.

TPM Key Attestation

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

Overview

While support for TPM-protected keys has existed since Windows 8, there were no mechanisms for CAs to cryptographically attest that the certificate requester private key is actually protected by a Trusted Platform Module (TPM). This update enables a CA to perform that attestation and to reflect that attestation in the issued certificate.

Note
This article assumes that the reader is familiar with certificate template concept (for reference, see Certificate Templates). It also assumes that the reader is familiar with how to configure enterprise CAs to issue certificates based on certificate templates (for reference, see Checklist: Configure CAs to Issue and Manage Certificates).

Terminology

Term	Definition
EK	Endorsement Key. This is an asymmetric key contained inside the TPM (injected at manufacturing time). The EK is unique for every TPM and can identify it. The EK cannot be changed or removed.
EKpub	Refers to public key of the EK.
EKPriv	Refers to private key of the EK.
EKCert	EK Certificate. A TPM manufacturer issued certificate for EKPub. Not all TPMs have EKCert.
TPM	Trusted Platform Module. It is designed to provide hardware-based, security-related functions. A TPM chip is a secure crypto-

Term	Definition
	processor that is designed to carry out cryptographic operations. The chip includes multiple physical security mechanisms to make it tamper resistant, and malicious software is unable to tamper with the security functions of the TPM.

Background

Beginning with Windows 8, a Trusted Platform Module (TPM) can be used to secure a certificate's private key. The Microsoft Platform Crypto Provider Key Storage Provider (KSP) enables this feature. There were two concerns with the implementation:

- There was no guarantee that the key is actually protected by a TPM (someone can easily spoof a software KSP as a TPM KSP with local administrator credentials).
- It was not possible to limit the list of TPMs that are allowed to protect enterprise issued certificates (in the event that the PKI Administrator wants to control the types of devices that can be used to obtain certificates in the environment).

TPM key attestation

TPM key attestation is the ability of the user who is requesting a certificate to cryptographically prove to a CA that the RSA key in the certificate request is protected by either “a” or “the” TPM that the CA trusts. The TPM trust model is discussed more in the Deployment overview section in this topic.

Why is TPM key attestation important?

A user certificate with a TPM-attested key provides higher security assurance backed up by non-exportability, anti-hammering, and isolation of keys provided by the TPM.

With TPM key attestation, a new management paradigm is now possible: An administrator can define the set of devices that users can use to access their corporate resources (for example, VPN or wireless access point) and have **strong** guarantees that no other devices can be used to access them. This new access control paradigm is **strong** because it is tied to a hardware-bound user identity, which is much stronger than previous software-based user identities.

How does TPM key attestation work?

In general, TPM key attestation is based on following pillars:

1. Every TPM ships with a unique asymmetric key, called the Endorsement Key (EK), burned by the manufacturer. We refer to the public portion of this key as EKPub. Some TPM chips also have an EK certificate that is issued by the manufacturer for the EKPub. We refer to this cert as EKCert.

2. A CA establishes trusts in the TPM either via EKPub or EKCert.
3. A user proves to the CA that the RSA key for which the certificate is being requested is cryptographically related to the EKPub and that the user owns the EKpriv.
4. The CA issues a certificate with a special issuance policy OID to denote that the key is now attested to be protected by a TPM.

Deployment overview

In this deployment, it is assumed that a Windows Server 2012 R2 enterprise CA is set up. Also, clients (Windows 8.1) are configured to enroll against that enterprise CA by using certificate templates. There are three steps to deploying TPM key attestation:

1. **Plan the TPM trust model:** The first step is to decide which TPM trust model to use. There are 3 supported ways for doing this:
   - **Trust based on user credential:** The enterprise CA trusts the user-provided EKPub as part of the certificate request and no validation is performed other than the user's domain credentials.
   - **Trust based on EKCert:** The enterprise CA validates the EKCert chain that is provided as part of the certificate request against an administrator-managed list of *acceptable EK cert chains*. These acceptable chains are defined per manufacturer and are expressed via two custom certificate stores (one for intermediate and one for root CA certificate) on the issuing CA. This trust mode means that **all** TPMs from a given manufacturer are trusted. Note that in this mode, TPMs in use in the environment must contain EKCerts.
   - **Trust based on EKPub:** The enterprise CA validates that the EKPub provided as part of the certificate request appears in an administrator-managed list of allowed EKPubs. This list is expressed as a directory of files where the name of each file in this directory is the SHA-2 hash of the allowed EKPub. This option offers the highest assurance level but requires more administrative effort. In this trust model, only the devices that have had their TPMs EKPub added to the allowed list of EKPubs are permitted to enroll for a TPM-attested certificate.

   Based on which method is used, the CA will use a different issuance policy OID in the issued certificate. For more details about issuance policy OIDs, see the Issuance Policy OIDs table in the Configure a certificate template section in this topic.

   Note that it is possible to choose a combination of TPM trust models. In this case, the CA will accept any of the attestation methods; however, the issuance policy OID will reflect all attestation methods that succeeded.
2. **Configure the certificate template:** Configuring the certificate template is described in the Deployment details section in this topic. This article does not cover how this certificate template is assigned to the enterprise CA or how enroll access is given to a group of users. For more information, see Checklist: Configure CAs to Issue and Manage Certificates.
3. **Configure the CA corresponding to TPM trust model**
   a. **Trust based on user credential:** No specific configuration required.
   b. **Trust based on EKCert:** The administrator must obtain the EKCert chain certificates from TPM manufacturers and import them to two new certificate stores on the CA that

perform TPM key attestation. For more information, see the CA configuration section in this topic.

c. **Trust based on EKPub:** The administrator must obtain the EKPub for each device that will need TPM-attested certificates and add them to the list of allowed EKPubs. For more information, see the CA configuration section in this topic.

Notes

- This feature requires Windows 8.1/Windows Server 2012 R2.
- TPM key attestation for third-party smart card KSPs is not supported. Microsoft Platform Crypto Provider KSP must be used.
- TPM key attestation only works for RSA keys.
- TPM key attestation is not supported for a standalone CA.

Deployment details

Configure a certificate template

To configure the certificate template for TPM key attestation, do the following configuration steps:

1. **Compatibility** tab

   In the **Compatibility Settings** section:

   - Ensure **Windows Server 2012 R2** is selected for the **Certification Authority**.
   - Ensure **Windows 8.1 / Windows Server 2012 R2** is selected for the **Certificate recipient**.

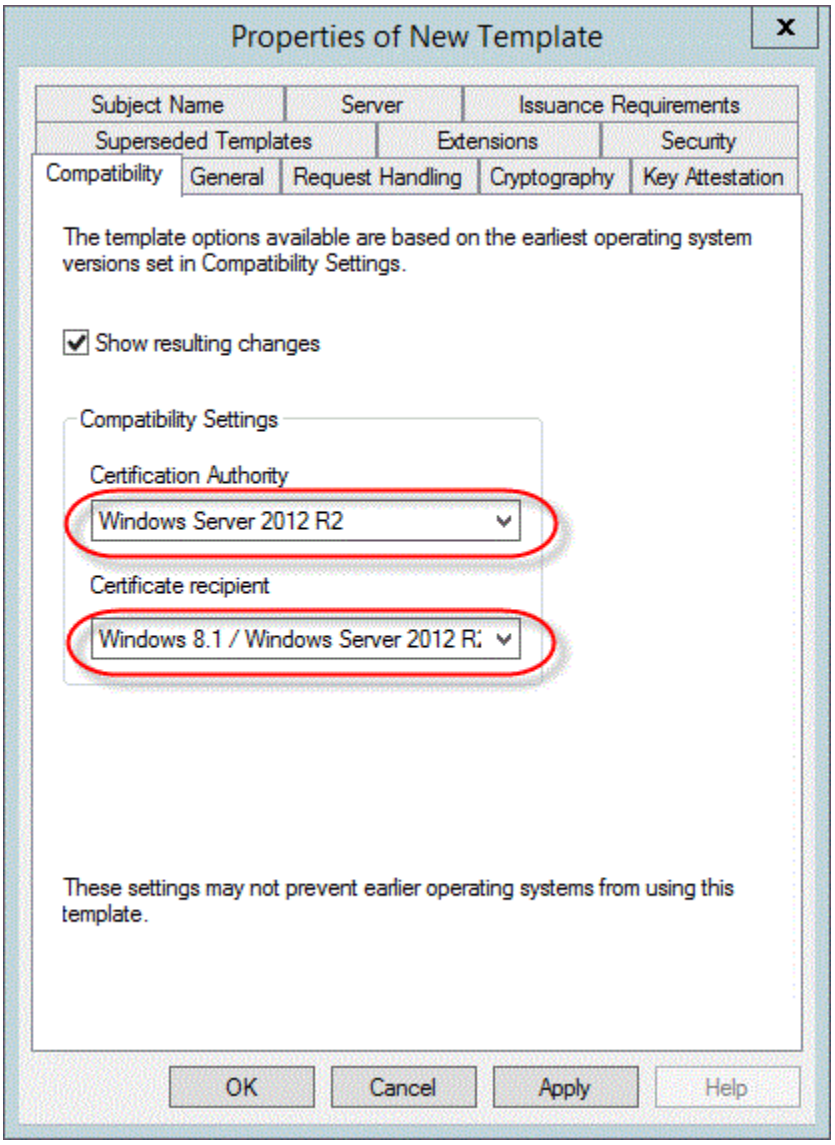

2. **Cryptography** tab

   Ensure **Key Storage Provider** is selected for the **Provider Category** and **RSA** is selected for the **Algorithm name**. Ensure **Requests must use one of the following providers** is selected and the **Microsoft Platform Crypto Provider** option is selected for the **Providers**.

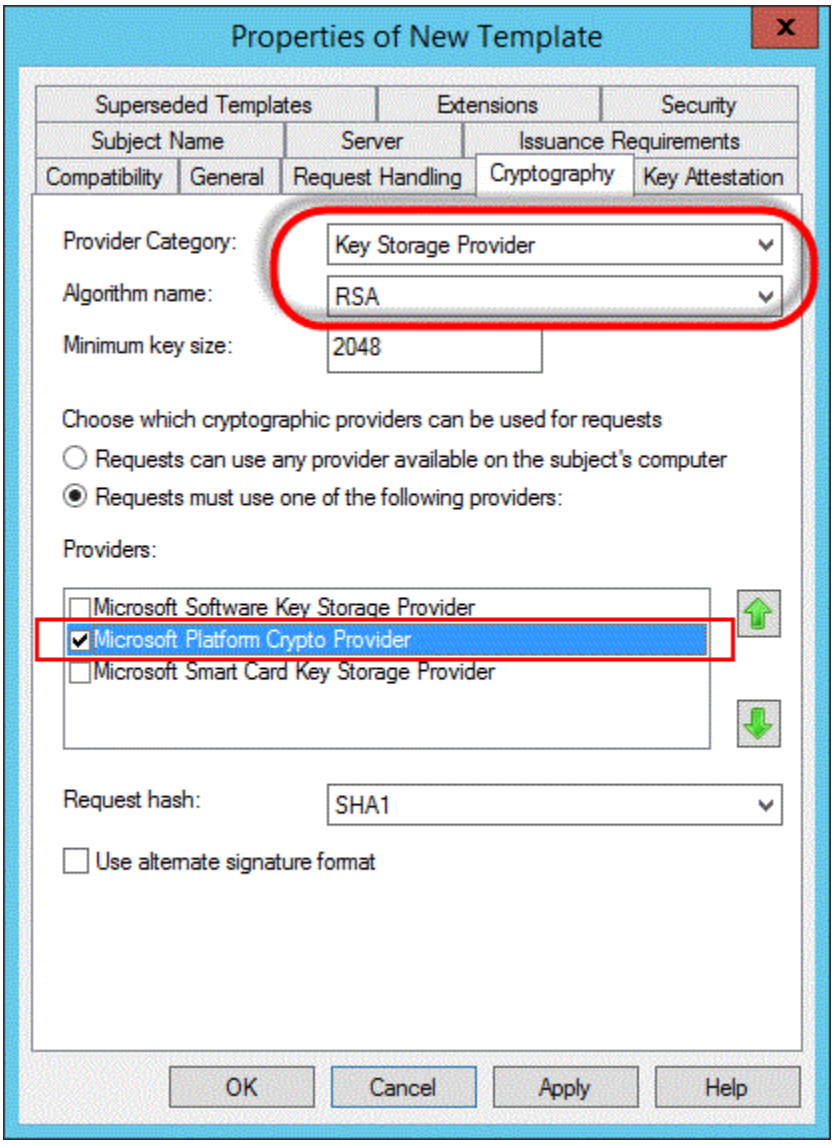

3. **Key Attestation** tab

   This is a new tab for Windows Server 2012 R2:

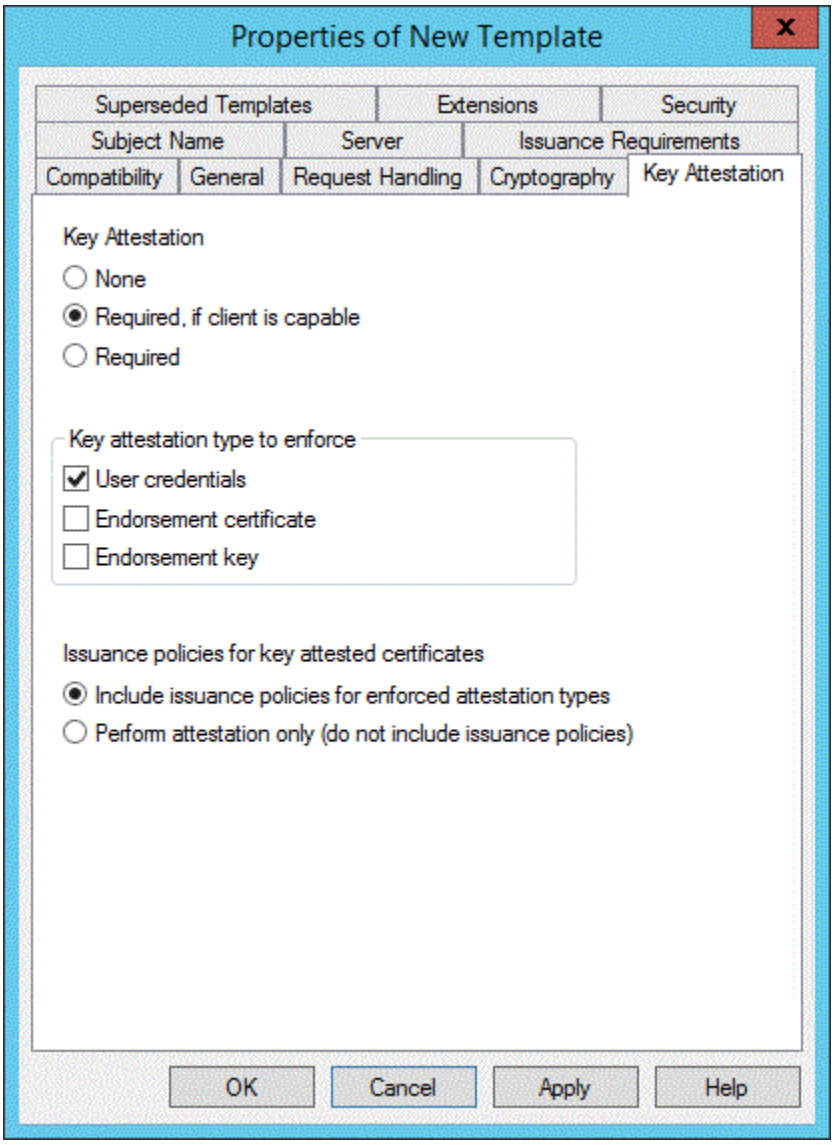

Choose an attestation mode from three possible options.

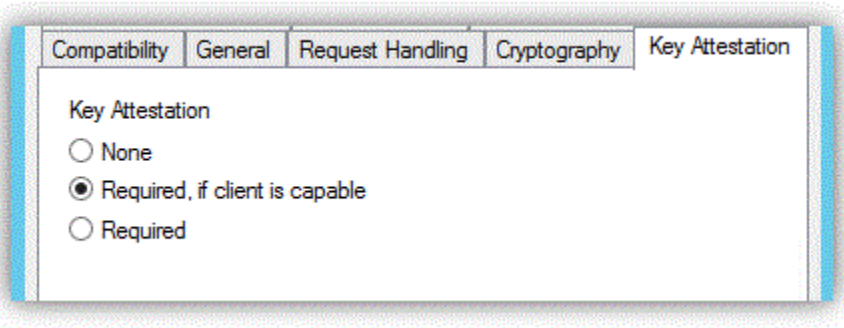

- **None:** Implies that key attestation must not be used
- **Required, if client is capable:** Allows users on a device that does not support TPM key attestation (because either the TPM is an old TPM that does not support key attestation or there is no TPM on that device) to continue enrolling for that certificate. Users who can perform attestation will be distinguished with a special issuance policy OID.
- **Required:** Client must perform TPM key attestation, otherwise the certificate request will fail.

Then choose the TPM trust model. There are again three options.

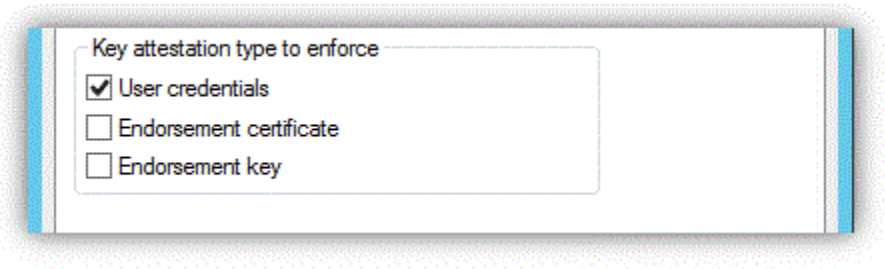

- **User credentials:** Allow an authenticating user to vouch for a valid TPM by specifying their domain credentials.
- **Endorsement certificate:** The EKCert of the device must validate through administrator-managed TPM intermediate CA certificates to an administrator-managed root CA certificate. If you choose this option, you must set up EKCA and EKRoot certificate stores on the issuing CA as described in the CA configuration section in this topic.
- **Endorsement Key:** The EKPub of the user device must appear in a PKI administrator-managed list. This option offers the highest assurance level but requires more administrative effort. If you choose this option, you must set up an EKPub list on the issuing CA as described in the CA configuration section in this topic.

Finally, decide which issuance policy to show in the issued certificate. By default, each enforcement type has an associated object identifier (also known as OID) that will be inserted into the certificate if it passes that enforcement type as described in the following table. Note that it is possible to choose a combination of enforcement methods. In this case, the CA will accept any of the attestation methods; however, the issuance policy OID will reflect all attestation methods that succeeded.

Issuance Policy OIDs

OID	Key attestation type	Description	Assurance level
1.3.6.1.4.1.311.21.30	EK	“EK Verified”: For administrator-managed list of EK	High
1.3.6.1.4.1.311.21.31	Endorsement certificate	“EK Certificate Verified”: When EK certificate chain is validated	Medium
1.3.6.1.4.1.311.21.32	User credentials	“EK Trusted on Use”: For user-attested EK	Low

The OIDs will be inserted into the issued certificate if **Include Issuance Policies** is selected (the default configuration).

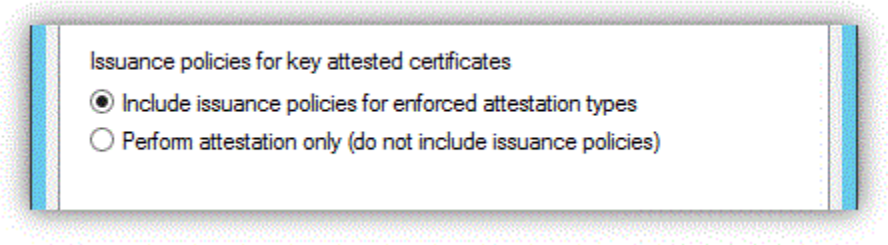

Tip

One potential use of having the OID present in the certificate is to limit access to VPN or wireless to certain devices. Example: Allow network access if OID 1.3.6.1.4.1.311.21.30 is present in the certificate. This allows you to limit access to devices whose TPM EK is present in the EKPUB list.

CA configuration

1. **Setup EKCA and EKROOT certificate stores on an issuing CA**

   If you chose **Endorsement Certificate** for the template settings, do the following configuration steps:

   a. Use Windows PowerShell to create two new certificate stores on the certification authority (CA) server that will perform TPM key attestation.

   b. Obtain the intermediate and root CA certificates(s) from manufacturer(s) that you want to allow in your enterprise environment. Those certificates must be imported into previously created certificate stores (EKCA and EKROOT) as appropriate.

   The following Windows PowerShell script performs both of these steps. In the following example, the TPM manufacturer Fabrikam has provided a root certificate *FabrikamRoot.cer* and an issuing CA certificate *Fabrikamca.cer*.

```
PS C:>\cd cert:
PS Cert:\>cd .\\LocalMachine
PS Cert:\LocalMachine> new-item EKROOT
PS Cert:\ LocalMachine> new-item EKCA
PS Cert:\EKCA\copy FabrikamCa.cer .\EKCA
PS Cert:\EKROOT\copy FabrikamRoot.cer .\EKROOT
```

2. **Setup EKPUB List if EK attestation type**

   If you chose **Endorsement Key** for the template settings, the next configuration steps are to create and configure a folder on the issuing CA where each file is 0 KB with the SHA-2 hash of the EK as the filename. This folder serves as an "allow list" of devices that are permitted to obtain TPM key attested certificates. Because you must manually add the EKPUB for each and every device that requires an attested certificate, it provides the enterprise with a guarantee of the devices that are authorized to obtain TPM key attested certificates. Configuring a CA for this mode requires two steps:

   a. **Create the EndorsementKeyListDirectories registry entry:** Use the Certutil command-line tool to configure the Windows folder locations where trusted EKpubs are defined as described in the following table.

Operation	Command syntax
Add folder locations	certutil.exe -setreg CA\EndorsementKeyListDirectories +"<folder>"
Remove folder locations	certutil.exe -setreg CA\EndorsementKeyListDirectories -"<folder>"

   The EndoresementKeyListDirectories in certutil command is a registry setting as described in the following table.

Value name	Type	Data
EndorsementKeyListDirectories	REG_MULTI_SZ	<LOCAL or UNC path to EKPUB allow list(s)> Example: *blueCA.contoso.com*\ekpub *bluecluster1.contoso.com*\ekpub D:\ekpub
HKLM\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\<CA Sanitized Name>		

EndorsementKeyListDirectories will contain a list of UNC or local file paths, each pointing to a Windows folder that the CA has Read access to. Each folder may contain zero or more allow list entries, where each entry is a file with a name that is the SHA-2 hash of a trusted EKpub (with no file extension). Creating or editing this registry key configuration will require a restart of the CA, just like existing CA registry configuration settings. However, edits to the configuration setting will take effect immediately and will not require the CA to be restarted.

Security

Secure the folders in the list from tampering and unauthorized access by configuring permissions so that only authorized administrators have Read and Write access. The computer account of the CA server requires Read access only.

b. **Populate the EKPUB list:** Use the following Windows PowerShell cmdlet to obtain the public key hash of the TPM EK by using Windows PowerShell on each device and then send this public key hash to the CA and store it on the EKPubList folder.

```
PS C:>\$a=Get-TpmEndorsementKeyInfo –hashalgorithm sha256
PS C:>$b=new-item $a.PublickKeyHash –ItemType file
PS C:>\xcopy $b \\enterpriseCA.contoso.com\EKPubList
```

Troubleshooting

Key attestation fields are unavailable on a certificate template

The Key Attestation fields are not available if the template settings do not meet the requirements for attestation. Common reasons are:

1. The compatibility settings are not configured correctly. Make sure that they are configured as follows:

a. **Certification Authority**: **Windows Server 2012 R2**
b. **Certificate Recipient**: **Windows 8.1/Windows Server 2012 R2**

2. The cryptography settings are not configured correctly. Make sure that they are configured as follows:
   a. **Provider Category**: **Key Storage Provider**
   b. **Algorithm Name**: **RSA**
   c. **Providers**: **Microsoft Platform Crypto Provider**
3. The request handling settings are not configured correctly. Make sure that they are configured as follows:
   a. The **Allow private key to be exported** option must not be selected.
   b. The **Archive subject's encryption private key** option must not be selected.

Verification of TPM device for attestation

Use the Windows PowerShell cmdlet, **Confirm-CAEndorsementKeyInfo**, to verify that a specific TPM device is trusted for attestation by CAs. There are two options: one for verifying the EKCert, and the other for verifying an EKPub. The cmdlet is either run locally on a CA, or on remote CAs by using Windows PowerShell remoting.

1. For verifying trust on an EKPub, do the following two steps:
   a. **Extract the EKPub from the client computer:** The EKPub can be extracted from a client computer via **Get-TpmEndorsementKeyInfo**. From an elevated command prompt, run the following:

      ```
      PS C:>\$a=Get-TpmEndorsementKeyInfo –hashalgorithm sha256
      ```

   b. **Verify trust on an EKCert on a CA computer:** Copy the extracted string (the SHA-2 hash of the EKPub) to the server (for example, via email) and pass it to the Confirm-CAEndorsementKeyInfo cmdlet. Note that this parameter must be 64 characters.

      ```
      Confirm-CAEndorsementKeyInfo [-PublicKeyHash] <string>
      ```

2. For verifying trust on an EKCert, do the following two steps:
   a. **Extract the EKCert from the client computer:** The EKCert can be extracted from a client computer via **Get-TpmEndorsementKeyInfo**. From an elevated command prompt, run the following:

      ```
      PS C:>\$a= Get- TpmEndorsementKeyInfo
      PS C:>\$a.manufacturerCertificates|Export-Certificate
      c:\myEkcert.cer
      ```

   b. **Verify trust on an KCert on a CA computer:** Copy the extracted EKCert (EkCert.cer) to the CA (for example, via email or xcopy). As an example, if you copy the certificate file the “c:\diagnose” folder on the CA server, run the following to finish verification:

      ```
      PS C:>new-object
      System.Security.Cryptography.X509Certificates.X509Certificate2
      “c:\diagnose\myEKcert.cer” | Confirm-CAEndorsementKeyInfo
      ```

See Also

Trusted Platform Module Technology Overview

External Resource: Trusted Platform Module

CA Backup and Restore Windows PowerShell cmdlets

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

Overview

The ADCSAdministration Windows PowerShell module was introduced in Window Server 2012. Two new cmdlets were added to this module in Window Server 2012 R2 to support the Backup and Restore of a CA.

- Backup-CARoleService
- Restore-CARoleService

Backup-CARoleService

Table 17: Backup and Restore Windows PowerShell Cmdlets

ADCSAdministration Cmdlet	Arguments - Bold arguments are required	Description
Backup-CARoleService	**-Path**	• String - location to save the backup • This is the only unnamed parameter • positional parameter **Example:** Backup-CARoleService.-Path c:\adcsbackup1 Backup-CARoleService

ADCSAdministration Cmdlet	Arguments - Bold arguments are required	Description
		c:\adcsbackup2
	-KeyOnly	• Backup the CA certificate without the database **Example:** Backup-CARoleService c:\adcsbackup3 -KeyOnly
	-Password	• Specifies the password to protect CA certificates and private keys • Must be a secure string • Not valid with the -DatabaseOnly parameter Example: Backup-CARoleService c:\adcsbackup4 -Password (Read-Host -prompt "Password:" -AsSecureString) Backup-CARoleService c:\adcsbackup5 -Password (ConvertTo-SecureString "Pa55w0rd!" -AsPlainText -Force)
	-DatabaseOnly	• Backup the database without the CA certificate Backup-CARoleService c:\adcsbackup6 -DatabaseOnly
	-Force	1. Allows you to overwrite the backup that preexists in the location specified in the -Path parameter Backup-CARoleService c:\adcsbackup1 -Force
	-Incremental	• Perform an incremental backup Backup-CARoleService c:\adcsbackup7 -Incremental

ADCSAdministration Cmdlet	Arguments - Bold arguments are required	Description
	-KeepLog	1. Instructs the command to keep log files. If the switch is not specified, log files are truncated by default except in the Incremental scenario Backup-CARoleService c:\adcsbackup7 -KeepLog

-Password <Secure String>

If the -Password parameter is used, the supplied password must be a secure string. Use the **Read-Host** cmdlet to launch an interactive prompt for secure password entry, or use the **ConvertTo-SecureString** cmdlet to specify the password in-line.

Review the following examples

Specifying a secure string for the Password parameter using Read-Host

```
Backup-CARoleService c:\adcsbackup4 -Password (Read-Host -prompt "Password:" -
AsSecureString)
```

Specifying a secure string for the Password parameter using ConvertTo-SecureString

```
Backup-CARoleService c:\adcsbackup5 -Password (ConvertTo-SecureString "Pa55w0rd!" -
AsPlainText -Force)
```

Restore-CARoleService

ADCSAdministration Cmdlet	Arguments - Bold arguments are required	Description
Restore-CARoleService	**-Path**	• String - location to restore backup from • This is the only unnamed parameter • positional parameter **Example:** Restore-CARoleService.-Path c:\adcsbackup1 -Force Restore-CARoleService c:\adcsbackup2 -Force

ADCSAdministration Cmdlet	Arguments - Bold arguments are required	Description
	-KeyOnly	• Restore the CA certificate without the database • Must be specified if the backup was taken with the -KeyOnly option **Example:** Restore-CARoleService c:\adcsbackup3 -KeyOnly -Force
	-Password	• Specifies the password of the CA certificates and private keys • Must be a secure string **Example:** Restore-CARoleService c:\adcsbackup4 -Password (read-host -prompt "Password:" -AsSecureString) -Force Restore-CARoleService c:\adcsbackup5 -Password (ConvertTo-SecureString "Pa55w0rd!" -AsPlainText -Force) -Force
	-DatabaseOnly	• Restore the database without the CA certificate Restore-CARoleService c:\adcsbackup6 -DatabaseOnly
	-Force	• Allows you to overwrite the preexisting keys • Is an optional parameter but when restoring in-place, it is likely required Restore-CARoleService c:\adcsbackup1 -Force

Issues

A non-password protected backup is taken if the ConvertTo-SecureString function fails while using the Backup-CARoleService with the -Password parameter.

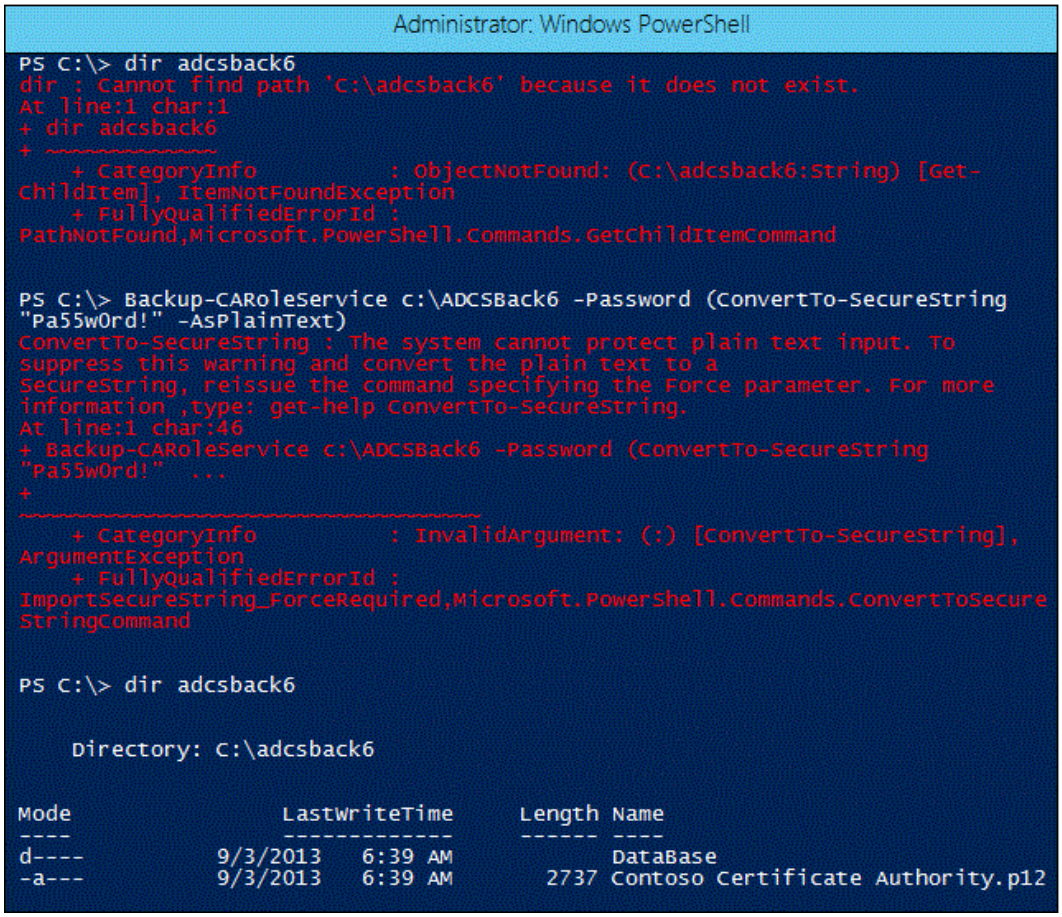

Table 18: Common Errors

Action	Error	Comment
Restore-CARoleService C:\ADCSBackup	Restore-CARoleService : The process cannot access the file because it is being used by another process. (Exception from HRESULT: 0x80070020)	Stop the Active Directory Certificate Services service prior to running the Restore-CARoleService cmdlet
Restore-CARoleService C:\ADCSBackup	Restore-CARoleService : The directory is not empty. (Exception from HRESULT: 0x80070091)	Use the -Force parameter to overwrite preexisting keys

Action	Error	Comment
Backup-CARoleService C:\ADCSBackup -Password (Read-Host -Prompt "Password:" -AsSecureString) -DatabaseOnly	Backup-CARoleService : Parameter set cannot be resolved using the specified named parameters.	The -Password parameter is only used to password protect private keys and is therefore invalid when you are not backing them up
Restore-CARoleService C:\ADCSBack15 -Password (Read-Host -Prompt "Password:" -AsSecureString) -DatabaseOnly	Restore-CARoleService : Parameter set cannot be resolved using the specified named parameters.	The -Password parameter is only used to password protect private keys and is therefore invalid when you are not restoring them
Restore-CARoleService C:\ADCSBack14 -Password (Read-Host -Prompt "Password:" -AsSecureString)	Restore-CARoleService : The system cannot find the file specified. (Exception from HRESULT: 0x80070002)	The path specified does not contain a valid database backup. Perhaps the path is invalid or the backup was taken with the -KeysOnly option?

Additional Resources

Active Directory Certificate Services Migration Guide

Backing up a CA database and private key

Restoring the CA database and configuration on the destination server

Try This: Backup the CA in your lab using Windows PowerShell

1. Use the commands in this lesson to backup the CA database and private key secured with a password.
2. Hold off on the restore of the CA at this time.

Command line process auditing

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note

This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

Overview

- The pre-existing process creation audit event ID 4688 will now include audit information for command line processes.
- It will also log SHA1/2 hash of the executable in the Applocker event log
  - Application and Services Logs\Microsoft\Windows\AppLocker
- You enable via GPO, but it is disabled by default
  - "Include command line in process creation events"

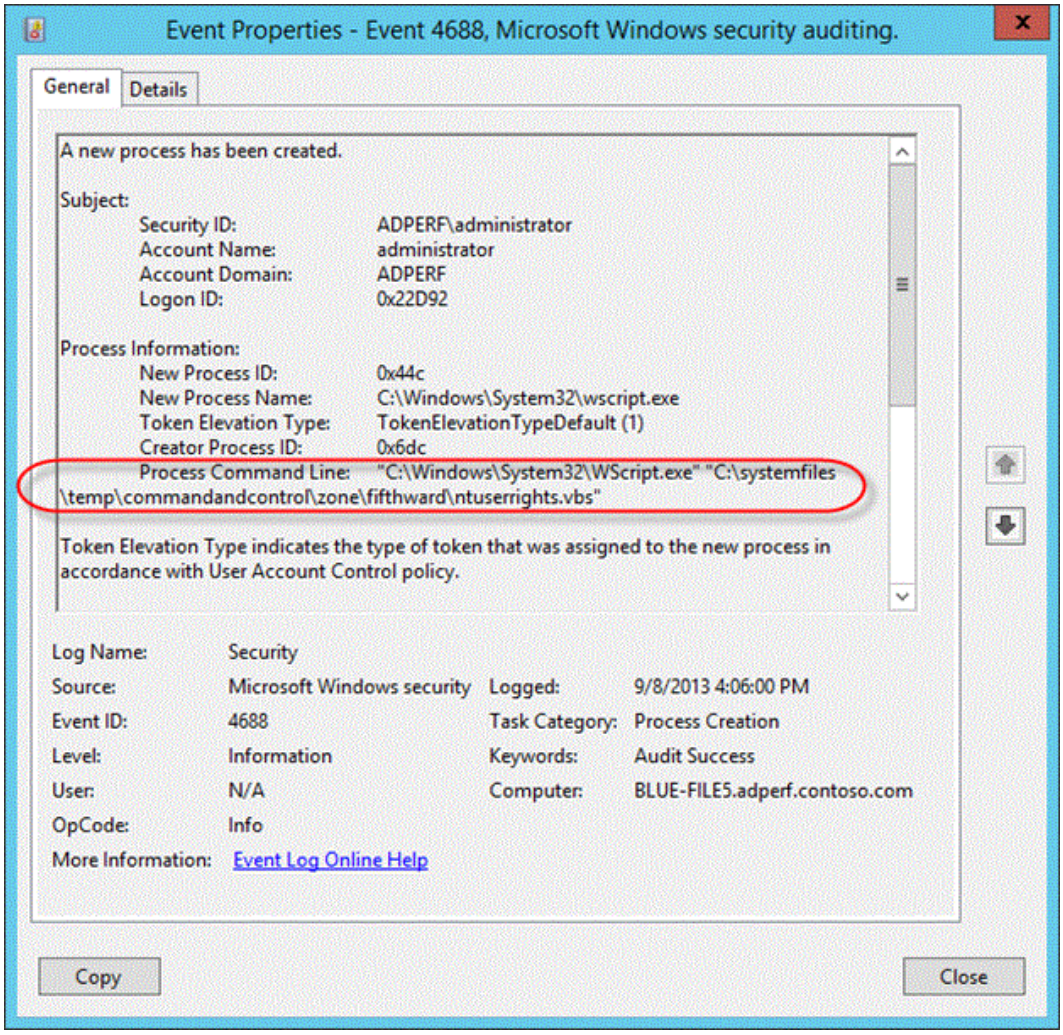

Figure 16 Event 4688

Review the updated event ID 4688 in Figure 16. Prior to this update none of the information for **Process Command Line** gets logged. Because of this additional logging we can now see that not only was the wscript.exe process started, but that it was also used to execute a VB script.

Configuration

To see the effects of this update, you will need to enable two policy settings.

You must have Audit Process Creation auditing enabled to see event ID 4688.

To enable the Audit Process Creation policy, edit the following group policy:

Policy location: Computer Configuration > Policies > Windows Settings > Security Settings > Advanced Audit Configuration > Detailed Tracking

Policy Name: Audit Process Creation

Supported on: Windows 7 and above

Description/Help:

This security policy setting determines whether the operating system generates audit events when a process is created (starts) and the name of the program or user that created it.

These audit events can help you understand how a computer is being used and to track user activity.

Event volume: Low to medium, depending on system usage

Default: Not configured

0be1c466-9e9c-4aff-8e84-cb36139db0f3

In order to see the additions to event ID 4688, you must enable the new policy setting: Include command line in process creation events

Table 19 Command line process policy setting

Policy Configuration	Details
Path	Administrative Templates\System\Audit Process Creation
Setting	**Include command line in process creation events**
Default setting	Not Configured (not enabled)
Supported on:	?

Policy Configuration	Details
Description	This policy setting determines what information is logged in security audit events when a new process has been created. This setting only applies when the Audit Process Creation policy is enabled. If you enable this policy setting the command line information for every process will be logged in plain text in the security event log as part of the Audit Process Creation event 4688, "a new process has been created," on the workstations and servers on which this policy setting is applied. If you disable or do not configure this policy setting, the process's command line information will not be included in Audit Process Creation events. Default: Not configured Note: When this policy setting is enabled, any user with access to read the security events will be able to read the command line arguments for any successfully created process. Command line arguments can contain sensitive or private information such as passwords or user data.

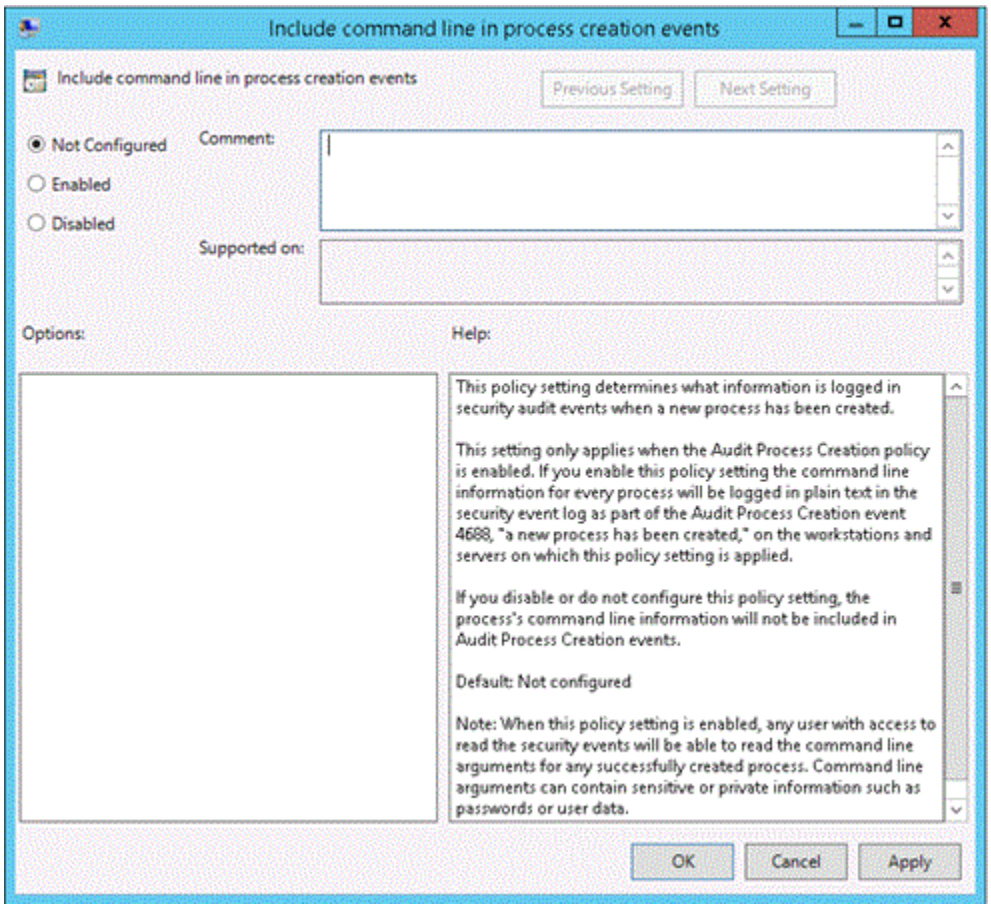

When you use Advanced Audit Policy Configuration settings, you need to confirm that these settings are not overwritten by basic audit policy settings. Event 4719 is logged when the settings are overwritten.

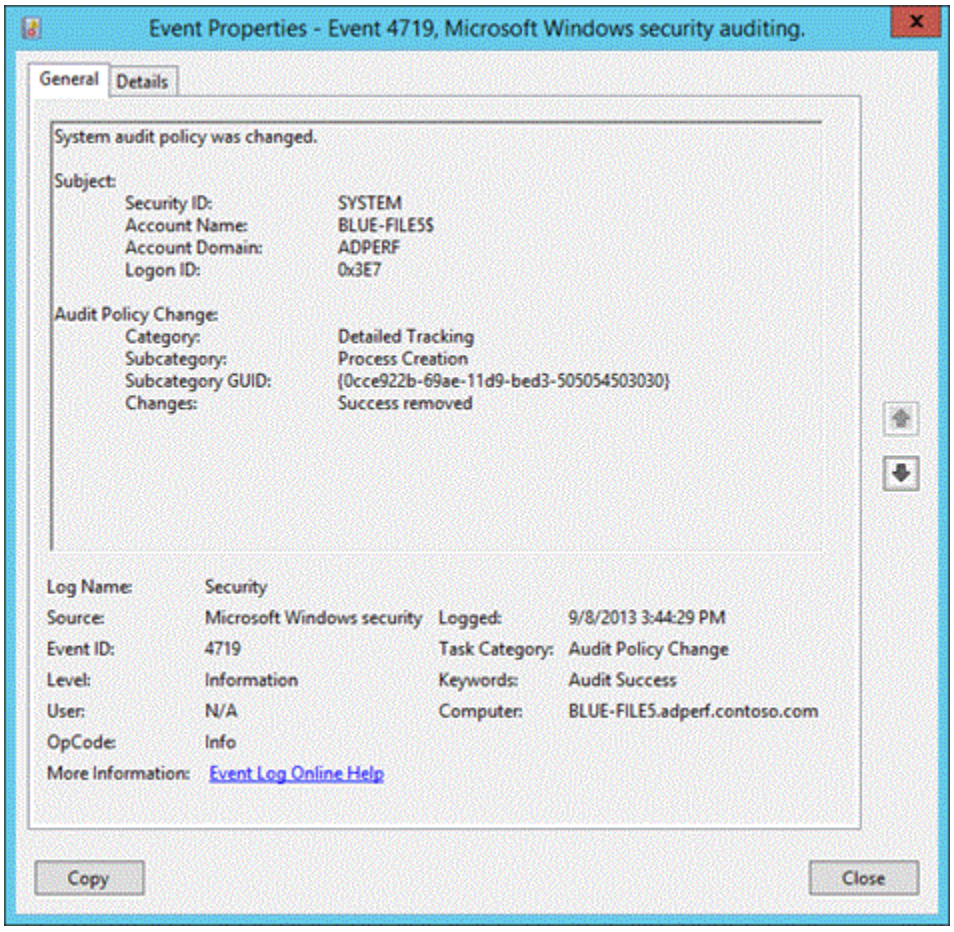

The following procedure shows how to prevent conflicts by blocking the application of any basic audit policy settings.

To ensure that Advanced Audit Policy Configuration settings are not overwritten

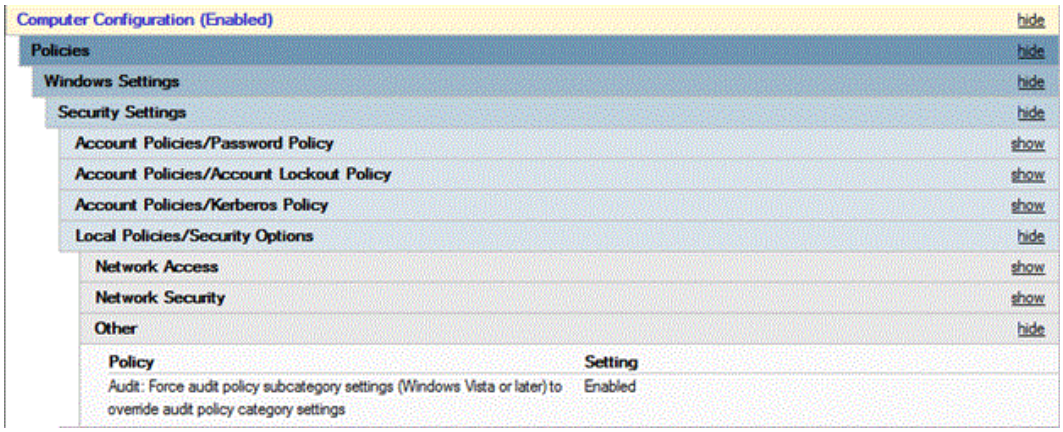

1. Open the Group Policy Management console
2. Right-click Default Domain Policy, and then click Edit.
3. Double-click Computer Configuration, double-click Policies, and then double-click Windows Settings.
4. Double-click Security Settings, and then click Security Options.
5. Double-click Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings, and then click Define this policy setting.
6. Click Enabled, and then click OK.

Additional Resources

Audit Process Creation

Advanced Security Audit Policy Step-by-Step Guide

AppLocker: Frequently Asked Questions

Try This: Explore command line process auditing

1. Enable **Audit Process Creation** events and ensure the Advance Audit Policy configuration is not overwritten
2. Create a script that will generate some events of interest and execute the script. Observe the events. The script used to generate the event in the lesson looked like this:

```
mkdir c:\systemfiles\temp\commandandcontrol\zone\fifthward
copy \\192.168.1.254\c$\hidden
c:\systemfiles\temp\hidden\commandandcontrol\zone\fifthward
start
C:\systemfiles\temp\hidden\commandandcontrol\zone\fifthward\ntus
errights.vbs
del c:\systemfiles\temp\*.* /Q
```

3. Enable the command line process auditing
4. Execute the same script as before and observe the events

Directory Services component updates

Author: Justin Turner, Senior Support Escalation Engineer with the Windows group

Note
This content is written by a Microsoft customer support engineer, and is intended for experienced administrators and systems architects who are looking for deeper technical explanations of features and solutions in Windows Server 2012 R2 than topics on

TechNet usually provide. However, it has not undergone the same editing passes, so some of the language may seem less polished than what is typically found on TechNet.

This lesson explains the Directory Services component updates in Windows Server 2012 R2.

What You Will Learn

Explain the following new Directory Services component updates:

- Explain the following new Directory Services component updates:
  - Domain and Forest Functional Levels
  - Deprecation of NTFRS
  - LDAP Query Optimizer changes
  - 1644 Event improvements
  - Active Directory Replication throughput improvement

Domain and Forest Functional Levels

Overview

The section provides a brief introduction to the domain and forest functional level changes.

New DFL and FFL

With the release, there are new domain and forest functional levels:

- Forest Functional Level: Windows Server 2012 R2
- Domain Functional Level: Windows Server 2012 R2

The Windows Server 2012 R2 Domain Functional Level enables support for the following:

1. DC-side protections for *Protected Users*

   Protected Users authenticating to a Windows Server 2012 R2 domain can no longer:

   - Authenticate with NTLM authentication
   - Use DES or RC4 cipher suites in Kerberos pre-authentication
   - Be delegated with unconstrained or constrained delegation
   - Renew user tickets (TGTs) beyond the initial 4 hour lifetime

2. Authentication Policies

   New forest-based Active Directory policies which can be applied to accounts in Windows Server 2012 R2 domains to control which hosts an account can sign-on from and apply access control conditions for authentication to services running as an account

3. Authentication Policy Silos

New forest-based Active Directory object which can create a relationship between user, managed service and computer accounts to be used to classify accounts for authentication policies or for authentication isolation.

See How to Configure Protected Accounts for more information.

In addition to the above features, the Windows Server 2012 R2 domain functional level ensures that any domain controller in the domain runs Windows Server 2012 R2. The Windows Server 2012 R2 forest functional level does not provide any new features, but it ensures that any new domain created in the forest will automatically operate at the Windows Server 2012 R2 domain functional level.

Minimum DFL enforced on new domain creation

Windows Server 2008 DFL is the minimum functional level supported on new domain creation.

Note

The deprecation of FRS is accomplished by removing the ability to install a new domain with a domain functional level lower than Windows Server 2008 with Server Manager or via Windows PowerShell.

Lowering the forest and domain functional levels

The forest and domain functional levels are set to Windows Server 2012 R2 by default on new domain and new forest creation but can be lowered using Windows PowerShell.

To raise or lower the forest functional level using Windows PowerShell, use the **Set-ADForestMode** cmdlet.

To set the contoso.com FFL to Windows Server 2008 mode:

```
Set-ADForestMode -ForestMode Windows2008Forest -Identity contoso.com
```

To raise or lower the domain functional level using Windows PowerShell, use the Set-ADDomainMode cmdlet.

To set the contoso.com DFL to Windows Server 2008 mode:

```
Set-ADDomainMode -DomainMode Windows2008Domain -Identity contoso.com
```

Promotion of a DC running Windows Server 2012 R2 as an additional replica into an existing domain running 2003 DFL works.

New domain creation in an existing forest

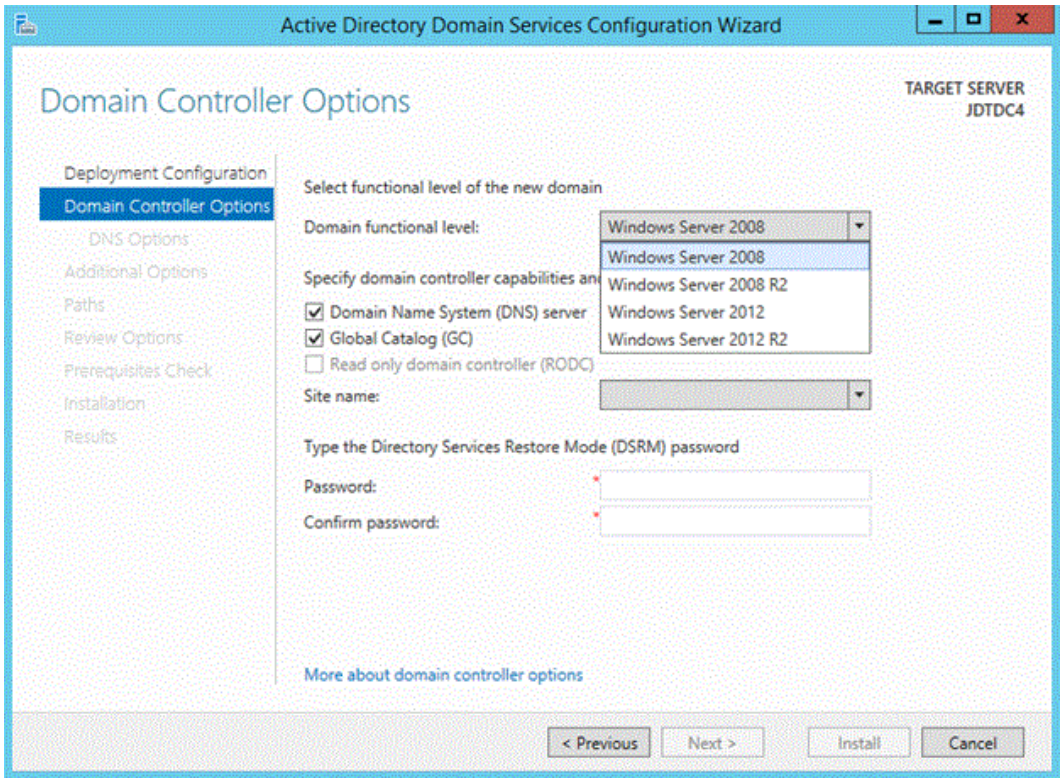

ADPREP

There are no new forest or domain operations in this release.

These .ldf files contain schema changes for the **Device Registration Service**.

1. Sch59
2. Sch61
3. Sch62
4. Sch63
5. Sch64
6. Sch65
7. Sch67

Work Folders:

1. Sch66

MSODS:

1. Sch60

Authentication Policies and Silos

1. Sch68
2. Sch69

Deprecation of NTFRS

Overview

FRS is deprecated in Windows Server 2012 R2. The deprecation of FRS is accomplished by enforcing a minimum domain functional level (DFL) of Windows Server 2008. This enforcement is present only if the new domain is created using Server Manager or Windows PowerShell.

You use the -DomainMode parameter with the Install-ADDSForest or Install-ADDSDomain cmdlets to specify the domain functional level. Supported values for this parameter can be either a valid integer or a corresponding enumerated string value. For example, to set the domain mode level to Windows Server 2008 R2, you can specify either a value of 4 or "Win2008R2". When executing these cmdlets from Server 2012 R2 valid values include those for Windows Server 2008 (3, Win2008) Windows Server 2008 R2 (4, Win2008R2) Windows Server 2012 (5, Win2012) and Windows Server 2012 R2 (6, Win2012R2). The domain functional level cannot be lower than the forest functional level, but it can be higher. Since FRS is deprecated in this release, Windows Server 2003 (2, Win2003) is not a recognized parameter with these cmdlets when executed from Windows Server 2012 R2.

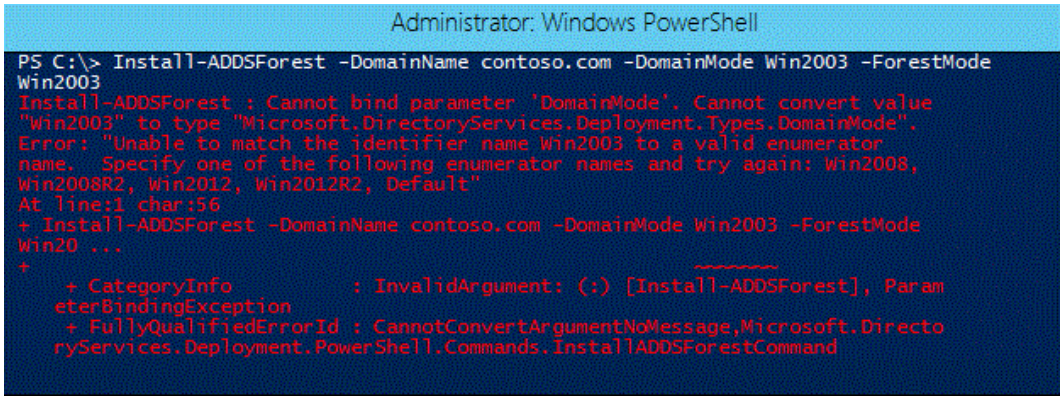

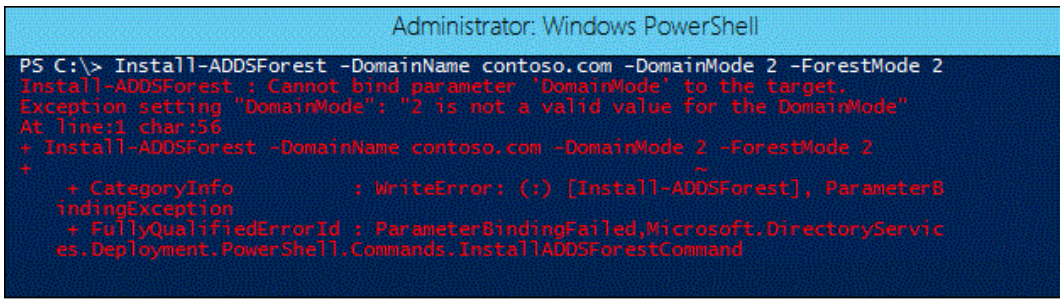

LDAP Query Optimizer changes

Overview

The LDAP query optimizer algorithm was reevaluated and further optimized. The result is the performance improvement in LDAP search efficiency and LDAP search time of complex queries.

Note

From the Developer:…improvements in the performance of searches through improvements in the mapping from LDAP query to ESE query. LDAP filters beyond a certain level of complexity prevent optimized index selection, resulting in drastically decreased performance (1000x or more). This change alters the way in which we select indices for LDAP queries to avoid this problem.

Notes

A complete overhaul of the LDAP query optimizer algorithm, resulting in:

- Faster search times
- Efficiency gains allow DCs to do more
- Less support calls regarding AD Performance issues
- Back ported to Windows Server 2008 R2 (KB 2862304)

Background

The ability to search Active Directory is a core service provided by domain controllers. Other services and line of business applications rely on Active Directory searches. Business operations can cease to a halt if this feature is not available. As a core and heavily used service, it is imperative that domain controllers handle LDAP search traffic efficiently. The LDAP query optimizer algorithm attempts to make LDAP searches efficient as possible by mapping LDAP search filters to a result set that can be satisfied via records already indexed in the database. This algorithm was reevaluated and further optimized. The result is the performance improvement in LDAP search efficiency and LDAP search time of complex queries.

Details of change

An LDAP search contains:

- A location (NC head, OU, Object) within the hierarchy to begin the search
- A search filter
- A list of attributes to return

The search process can be summarized as follows:

1. Simplify the search filter if possible.
2. Select a set of Index Keys that will return the smallest covered set.
3. Perform one or more intersections of Index Keys, to reduce the covered set.

4. For each record in the covered set, evaluate the filter expression as well as the security. If the filter evaluates to TRUE and access is granted, then return this record to the client.

The LDAP query optimization work modifies steps 2 and 3, to reduce the size of the covered set. More specifically, the current implementation selects duplicate Index Keys and performs redundant intersections.

Comparison between old and new algorithm

The target of the inefficient LDAP search in this example is a Windows Server 2012 domain controller. The search completes in approximately 44 seconds as a result of failing to find a more efficient index.

```
adfind -b dc=blue,dc=contoso,dc=com -f "(| (& (|(cn=justintu) (postalcode=80304)
(userprincipalname=justintu@blue.contoso.com)) (|(objectclass=person) (cn=justintu)) )
(&(cn=justintu)(objectclass=person)))" -stats >>adfind.txt


Using server: WINSRV-DC1.blue.contoso.com:389


…<removed search results>


Statistics
==================================
Elapsed Time: 44640 (ms)
Returned 324 entries of 553896 visited - (0.06%)


Used Filter:
 ( |  ( &  ( |  (cn=justintu)  (postalCode=80304)
(userPrincipalName=justintu@blue.contoso.com) )  ( |  (objectClass=person)  (cn=justintu)
) )  ( &  (cn=justintu)  (objectClass=person) ) )


Used Indices:
 DNT_index:516615:N


Pages Referenced          : 4619650
Pages Read From Disk      : 973
Pages Pre-read From Disk  : 180898
Pages Dirtied             : 0
Pages Re-Dirtied          : 0
```

```
Log Records Generated     : 0
Log Record Bytes Generated: 0
```

Sample results using the new algorithm

This example repeats the exact same search as above but targets a Windows Server 2012 R2 domain controller. The same search completes in less than a second due to the improvements in the LDAP query optimizer algorithm.

```
adfind -b dc=blue,dc=contoso,dc=com -f "(| (& (|(cn=justintu) (postalcode=80304)
(userprincipalname=dhunt@blue.contoso.com)) (|(objectclass=person) (cn=justintu)) )
(&(cn=justintu)(objectclass=person)))" -stats >>adfindBLUE.txt


Using server: winblueDC1.blue.contoso.com:389


….<removed search results>


Statistics
==================================
Elapsed Time: 672 (ms)
Returned 324 entries of 648 visited - (50.00%)


Used Filter:
 ( |  ( &  ( |  (cn=justintu)  (postalCode=80304)
(userPrincipalName=justintu@blue.contoso.com) )  ( |  (objectClass=person)  (cn=justintu)
) )  ( &  (cn=justintu)  (objectClass=person) ) )


Used Indices:
 idx_userPrincipalName:648:N
 idx_postalCode:323:N
 idx_cn:1:N


Pages Referenced          : 15350
Pages Read From Disk      : 176
Pages Pre-read From Disk  : 2
Pages Dirtied             : 0
Pages Re-Dirtied          : 0
```

```
Log Records Generated     : 0

Log Record Bytes Generated: 0
```

- If unable to optimize the tree:
  - For example: an expression in the tree was over a column not indexed
  - Record a list of indices that prevent optimization
  - Exposed via ETW tracing and event ID 1644

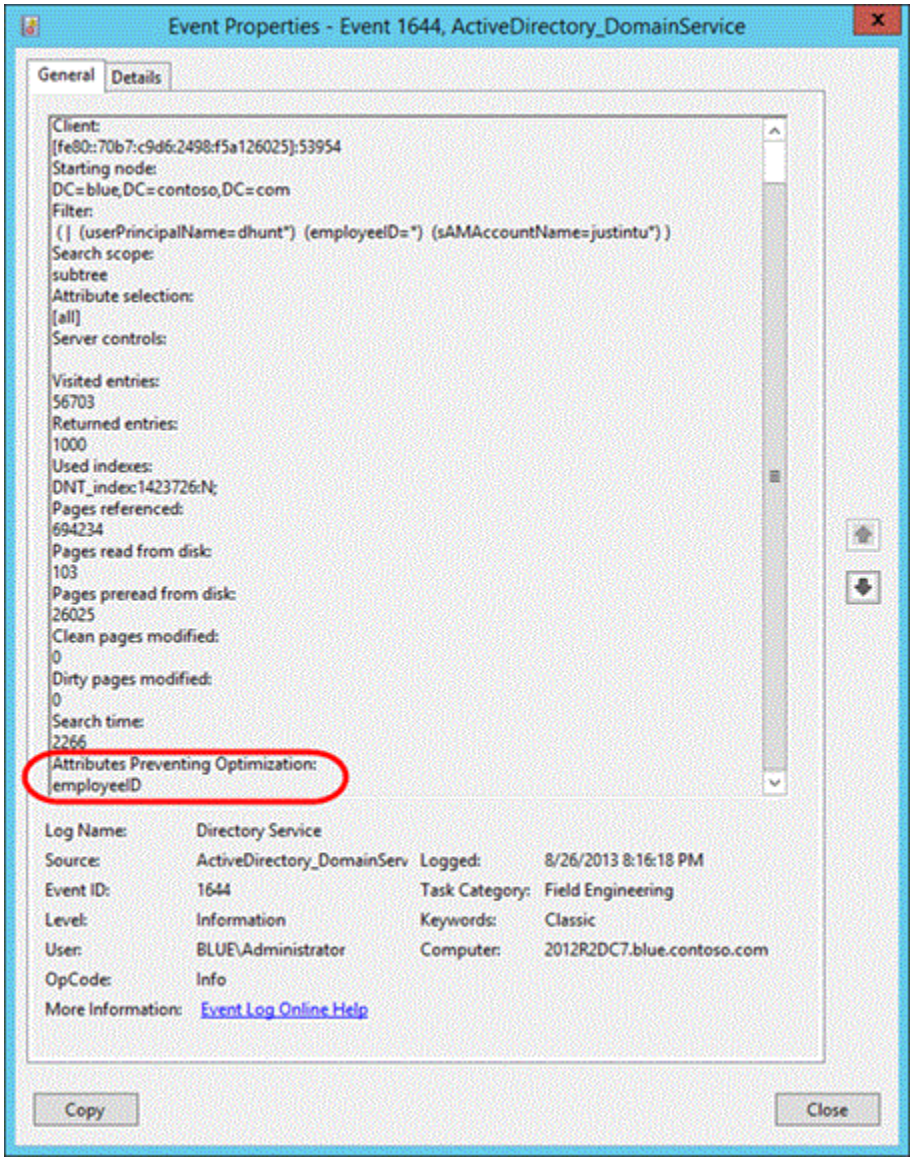

To enable the Stats control in LDP

1. Open LDP.exe, and connect and bind to a domain controller.
2. On the **Options** menu, click **Controls**.

3. On the Controls dialog box, expand the **Load Predefined** pull-down menu, click **Search Stats** and then click **OK**.

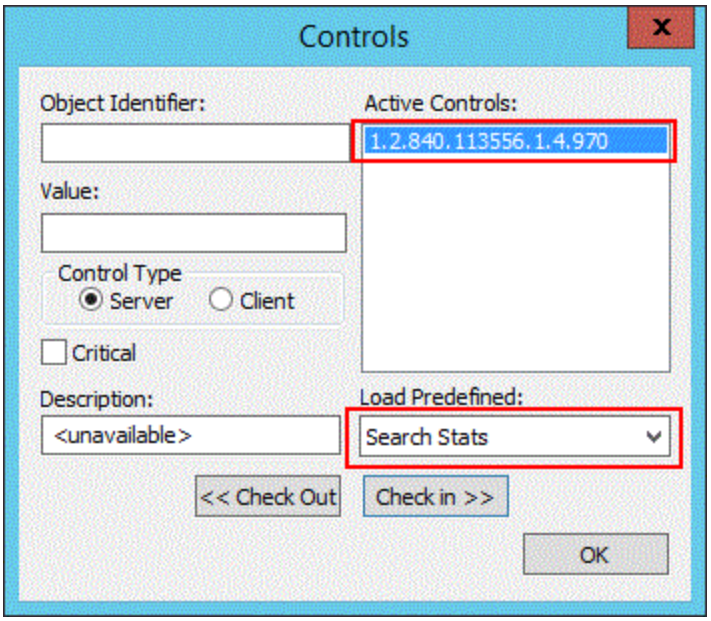

4. On the **Browse** menu, click **Search**
5. In the Search dialog box, select the **Options** button.
6. Ensure the **Extended** check box is selected on the Search Options dialog box and select **OK**.

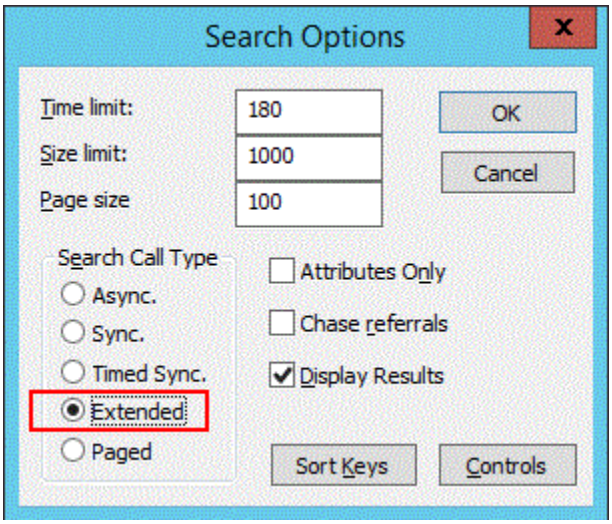

Try This: Use LDP to return query statistics

Perform the following on a domain controller, or from a domain-joined client or server that has the AD DS tools installed. Repeat the following targeting your Windows Server 2012 DC and your Windows Server 2012 R2 DC.

1. Review the "Creating More Efficient Microsoft AD Enabled Applications" article and refer back to it as needed.
2. Using LDP, enable search statistics (see To enable the Stats control in LDP)
3. Conduct several LDAP searches and observe the statistical information at the top of the results. You will repeat the same search in other activities so document them in a notepad text file.
4. Perform an LDAP search that the query optimizer should be able to optimize because of attributes indices
5. Attempt to construct a search that takes a long time to complete (you may want to increase the **Time limit** option so the search does not timeout).

Additional Resources

What Are Active Directory Searches?

How Active Directory Searches Work

Creating More Efficient Microsoft Active Directory-Enabled Applications

951581 LDAP queries are executed more slowly than expected in the AD or LDS/ADAM directory service and Event ID 1644 may be logged

1644 Event improvements

Overview

This update adds additional LDAP search result statistics to event ID 1644 to aid in troubleshooting purposes. Additionally, there is a new registry value that can be used to enable logging on a time-based threshold. These improvements were made available in Windows Server 2012 and Windows Server 2008 R2 SP1 via KB 2800945 and will be made available to Windows Server 2008 SP2.

Notes

- Additional LDAP search statistics are added to event ID 1644 to aid in troubleshooting inefficient or expensive LDAP searches
- You can now specify a Search Time Threshold (eg. Log event 1644 for searches taking longer than 100ms) instead of specifying the Expensive and Inefficient search result threshold values

Background

While troubleshooting Active Directory performance problems, it becomes apparent that LDAP search activity may be contributing to the problem. You decide to enable logging so that you can see expensive or inefficient LDAP queries processed by domain controller. In order to enable the logging, you must set the Field Engineering diagnostics value and can optionally specify the expensive / inefficient search results threshold values. Upon enabling the Field Engineering

logging level to a value of 5, any search that meets these criteria is logged in the Directory Services event log with an event ID 1644.

The event contains:

- Client IP and port
- Starting Node
- Filter
- Search scope
- Attribute selection
- Server controls
- Visited entries
- Returned entries

However, key data is missing from the event such as the amount of time spent on the search operation and what (if any) index was used.

Additional search statistics added to event 1644

- Used indexes
- Pages referenced
- Pages read from disk
- Pages preread from disk
- Clean pages modified
- Dirty pages modified
- Search time
- Attributes Preventing Optimization

New time-based threshold registry value for event 1644 logging

Instead of specifying the Expensive and Inefficient search result threshold values, you can specify Search Time Threshold. If you wanted to log all search results that took 50 ms or greater, you would specify 50 decimal / 32 hex (in addition to setting the Field Engineering value).

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters]

"Search Time Threshold (msecs)"=dword:00000032
```

Comparison of the old and new event ID 1644

OLD

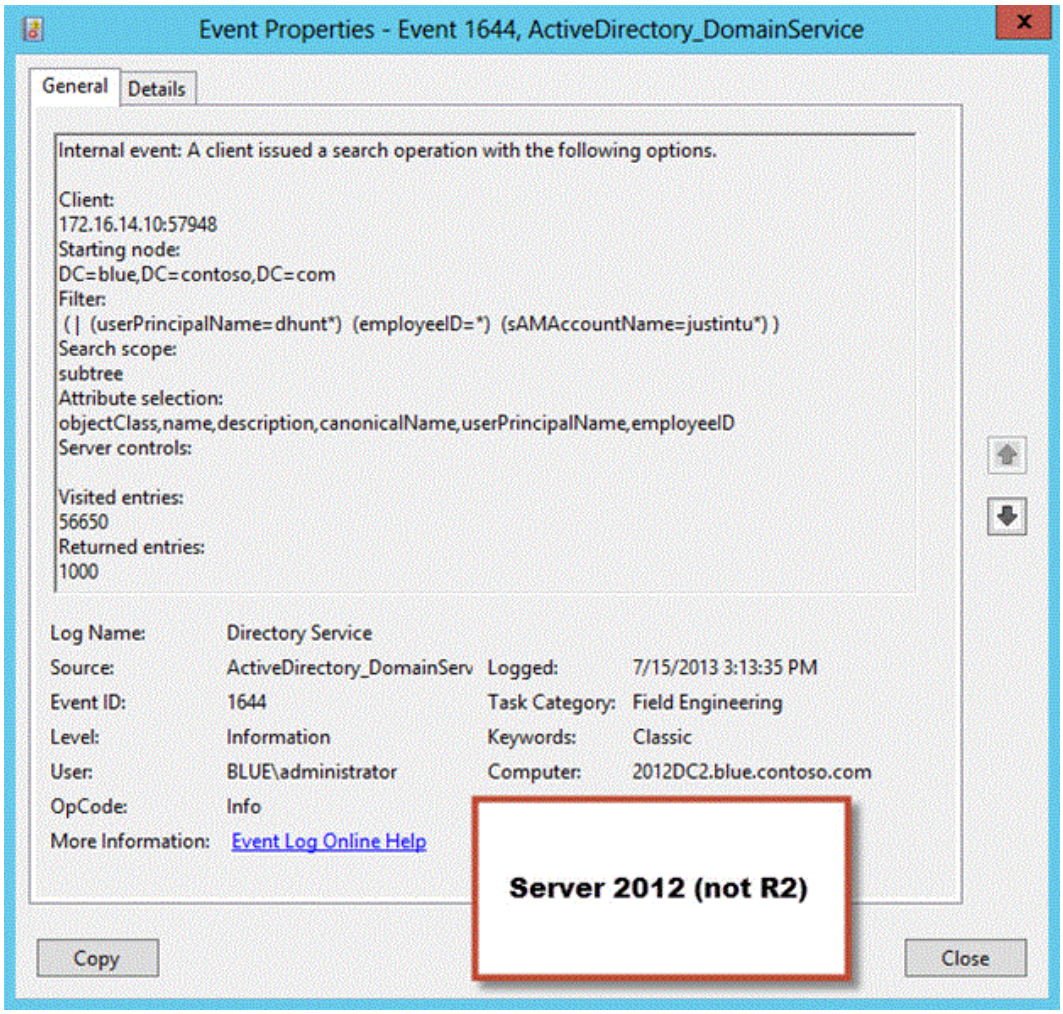
Event Properties - Event 1644, ActiveDirectory_DomainService
General
Details
Internal event: A client issued a search operation with the following options.
Client:
172.16.14.10:57948
Starting node:
DC=blue,DC=contoso,DC=com
Filter:
(| (userPrincipalName=dhunt*) (employeeID=*) (sAMAccountName=justintu*))
Search scope:
subtree
Attribute selection:
objectClass,name,description,canonicalName,userPrincipalName,employeeID
Server controls:
Visited entries:
56650
Returned entries:
1000
Log Name: Directory Service
Source: ActiveDirectory_DomainServ
Logged: 7/15/2013 3:13:35 PM
Event ID: 1644
Task Category: Field Engineering
Level: Information
Keywords: Classic
User: BLUE\administrator
Computer: 2012DC2.blue.contoso.com
OpCode: Info
More Information: Event Log Online Help
Server 2012 (not R2)
Copy
Close

NEW

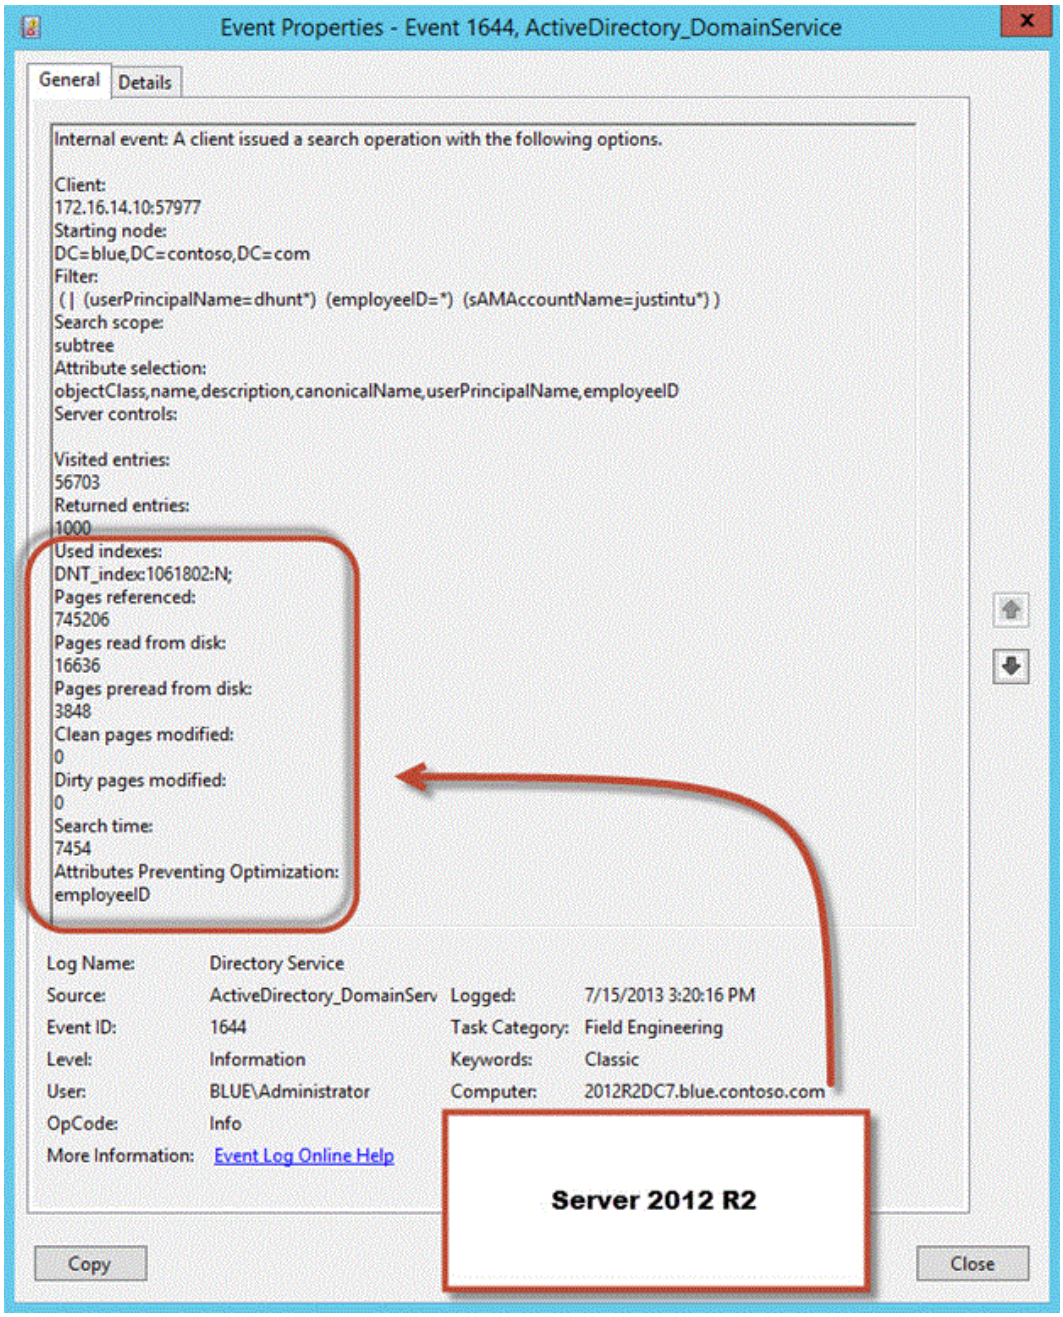

Try This: Use the event log to return query statistics

1. Repeat the following targeting your Windows Server 2012 DC and your Windows Server 2012 R2 DC. Observe the event ID 1644s on both DCs after each search.
2. Using regedit, enable event ID 1644 logging using a time-based threshold on the Windows Server 2012 R2 DC and the old method on the Windows Server 2012 DC.

3. Conduct several LDAP searches that exceed the threshold and observe the statistical information at the top of the results. Use the LDAP queries you documented earlier and repeat the same searches.
4. Perform an LDAP search that the query optimizer is not able to optimize because one or more attributes are not indexed.

Active Directory Replication throughput improvement

Overview

AD replication uses RPC for its replication transport. By default, RPC uses an 8K transmit buffer and a 5K packet size. This has the net effect where the sending instance will transmit three packets (approximately 15K worth of data) and then have to wait for a network round trip before sending more. Assuming a 3ms roundtrip time, the highest throughput would be around 40Mbps, even on 1Gbps or 10 Gbps networks.

Notes

- This update adjusts the maximum AD Replication throughput from 40Mbps to around 600 Mbps.
  - It increases the RPC send buffer size which reduces the number of network round trips
- The effect will be most noticeable on high speed, high latency network.

This updates increase the maximum throughput to around 600 Mbps by changing the RPC send buffer size from 8K to 256KB. This change allows the TCP window size to grow beyond 8K, reducing the number of network round trips.

Note

There are no configurable settings to modify this behavior.

Additional Resources

How the Active Directory Replication Model Works

How to Configure Protected Accounts

Through Pass-the-hash (PtH) attacks, an attacker can authenticate to a remote server or service by using the underlying NTLM hash of a user's password (or other credential derivatives). Microsoft has previously published guidance to mitigate pass-the-hash attacks. Windows Server 2012 R2 includes new features to help mitigate such attacks further. For more information about other security features that help protect against credential theft, see Credentials Protection and Management. This topic explains how to configure the following new features:

- Protected Users

- Authentication policies
- Authentication policy silos

There are additional mitigations built in to Windows 8.1 and Windows Server 2012 R2 to help protect against credential theft, which are covered in the following topics:

- Restricted Admin mode for Remote Desktop
- LSA Protection

Protected Users

Protected Users is a new global security group to which you can add new or existing users. Windows 8.1 devices and Windows Server 2012 R2 hosts have special behavior with members of this group to provide better protection against credential theft. For a member of the group, a Windows 8.1 device or a Windows Server 2012 R2 host does not cache credentials that are not supported for Protected Users. Members of this group have no additional protection if they are logged on to a device that runs a version of Windows earlier than Windows 8.1.

Members of the Protected Users group who are signed-on to Windows 8.1 devices and Windows Server 2012 R2 hosts can *no longer* use:

- Default credential delegation (CredSSP) - plaintext credentials are not cached even when the **Allow delegating default credentials** policy is enabled
- Windows Digest - plaintext credentials are not cached even when they are enabled
- NTLM - NTOWF is not cached
- Kerberos long term keys - Kerberos ticket-granting ticket (TGT) is acquired at logon and cannot be re-acquired automatically
- Sign-on offline - the cached logon verifier is not created

If the domain functional level is Windows Server 2012 R2, members of the group can no longer:

- Authenticate by using NTLM authentication
- Use Data Encryption Standard (DES) or RC4 cipher suites in Kerberos pre-authentication
- Be delegated by using unconstrained or constrained delegation
- Renew user tickets (TGTs) beyond the initial 4-hour lifetime

To add users to the group, you can use UI tools such as Active Directory Administrative Center (ADAC) or Active Directory Users and Computers, or a command-line tool such as Dsmod group, or the Windows PowerShell Add-ADGroupMember cmdlet. Accounts for services and computers *should not* be members of the Protected Users group. Membership for those accounts provides no local protections because the password or certificate is always available on the host.

Warning

The authentication restrictions have no workaround, which means that members of highly privileged groups such as the Enterprise Admins group or the Domain Admins group are subject to the same restrictions as other members of the Protected Users group. If all members of such groups are added to the Protected Users group, it is possible for all of

those accounts to be locked out. You should never add all highly privileged accounts to the Protected Users group until you have thoroughly tested the potential impact.

Members of the Protected Users group must be able to authenticate by using Kerberos with Advanced Encryption Standards (AES). This method requires AES keys for the account in Active Directory. The built-in Administrator does not have an AES key unless the password was changed on a domain controller that runs Windows Server 2008 or later. Additionally, any account, which has a password that was changed at a domain controller that runs an earlier version of Windows Server, is locked out. Therefore, follow these best practices:

- Do not test in domains unless **all domain controllers run Windows Server 2008 or later**.
- **Change password** for all domain accounts, for example, the domain's built-in **Administrator** account, which were created *before* the domain was created. Otherwise, these accounts cannot be authenticated.
- **Change password** for each user before adding the account to the Protected Users group or ensure that the password was changed recently on a domain controller that runs Windows Server 2008 or later.

Requirements for using protected accounts

Protected accounts have the following deployment requirements:

- To provide client-side restrictions for Protected Users, hosts must run Windows 8.1 or Windows Server 2012 R2. A user only has to sign-on with an account that is a member of a Protected Users group. In this case, the Protected Users group can be created by transferring the primary domain controller (PDC) emulator role to a domain controller that runs Windows Server 2012 R2. After that group object is replicated to other domain controllers, the PDC emulator role can be hosted on a domain controller that runs an earlier version of Windows Server.
- To provide domain controller-side restrictions for Protected Users, that is to restrict usage of NTLM authentication, and other restrictions, the domain functional level must be Windows Server 2012 R2. For more information about functional levels, see Understanding Active Directory Domain Services (AD DS) Functional Levels.

Troubleshoot events related to Protected Users

This section covers new logs to help troubleshoot events that are related to Protected Users and how Protected Users can impact changes to troubleshoot either ticket-granting tickets (TGT) expiration or delegation issues.

New logs for Protected Users

Two new operational administrative logs are available to help troubleshoot events that are related to Protected Users: Protected User – Client Log and Protected User Failures – Domain Controller Log. These new logs are located in Event Viewer and are disabled by default. To enable a log, click **Applications and Services Logs**, click **Microsoft**, click **Windows**, click **Authentication**, and then click the name of the log and click **Action** (or right-click the log) and click **Enable Log**.

For more information about events in these logs, see [Authentication Policies and Authentication Policy Silos](#).

Troubleshoot TGT expiration

Normally, the domain controller sets the TGT lifetime and renewal based on the domain policy as shown in the following Group Policy Management Editor window.

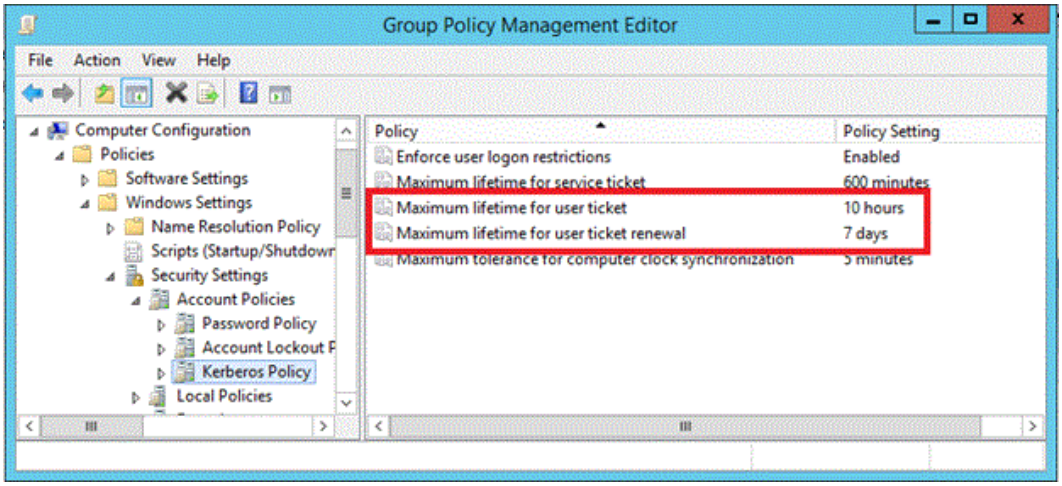

For **Protected Users**, the following settings are hard-coded:

- Maximum lifetime for user ticket: 240 minutes
- Maximum lifetime for user ticket renewal: 240 minutes

Troubleshoot delegation issues

Previously, if a technology that uses Kerberos delegation was failing, the client account was checked to see if **Account is sensitive and cannot be delegated** was set. However, if the account is a member of **Protected Users**, it might not have this setting configured in Active Directory Administrative Center (ADAC). As a result, check the setting and group membership when you troubleshoot delegation issues.

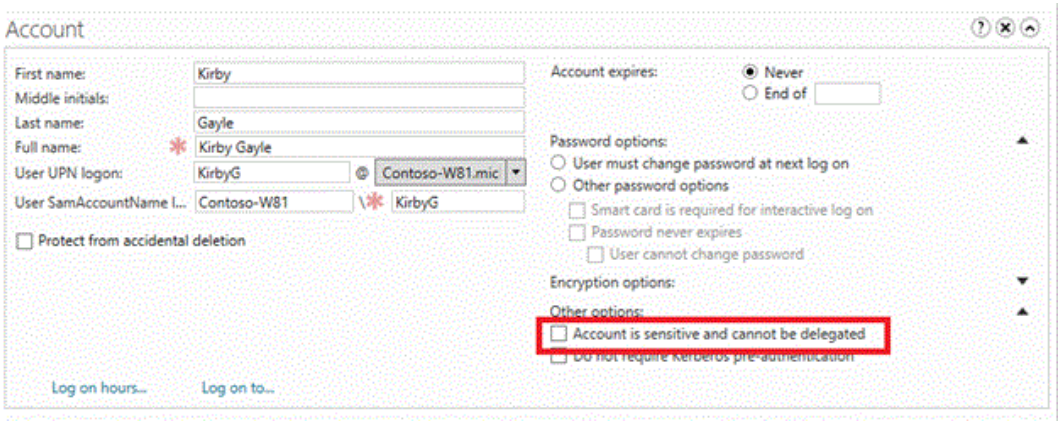

Audit authentication attempts

To audit authentication attempts explicitly for the members of the **Protected Users** group, you can continue to collect security log audit events or collect the data in the new operational administrative logs. For more information about these events, see Authentication Policies and Authentication Policy Silos

Provide DC-side protections for services and computers

Accounts for services and computers cannot be members of **Protected Users**. This section explains which domain controller-based protections can be offered for these accounts:

- Reject NTLM authentication: Only configurable via NTLM block policies
- Reject Data Encryption Standard (DES) in Kerberos pre-authentication: Windows Server 2012 R2 domain controllers do not accept DES for computer accounts unless they are configured for DES only because every version of Windows released with Kerberos also supports RC4.
- Reject RC4 in Kerberos pre-authentication: not configurable.

  Note
  Although it is possible to change the configuration of supported encryption types, it is not recommended to change those settings for computer accounts without testing in the target environment.
- Restrict user tickets (TGTs) to an initial 4-hour lifetime: Use Authentication Policies.
- Deny delegation with unconstrained or constrained delegation: To restrict an account, open Active Directory Administrative Center (ADAC) and select the **Account is sensitive and cannot be delegated** check box.

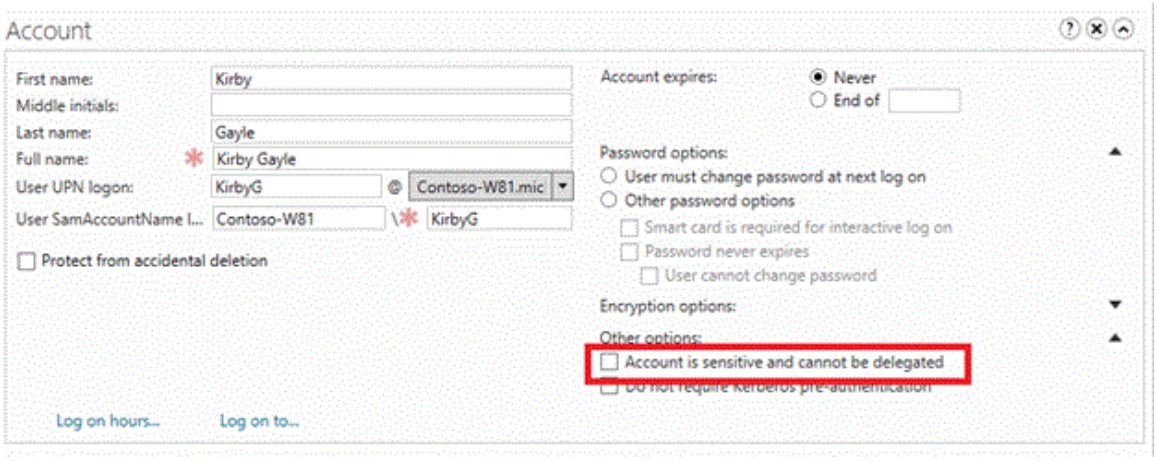

Authentication policies

Authentication Policies is a new container in AD DS that contains authentication policy objects. Authentication policies can specify settings that help mitigate exposure to credential theft, such as restricting TGT lifetime for accounts or adding other claims-related conditions.

In Windows Server 2012, Dynamic Access Control introduced an Active Directory forest-scope object class called Central Access Policy to provide an easy way to configure file servers across an organization. In Windows Server 2012 R2, a new object class called Authentication Policy (objectClass msDS-AuthNPolicies) can be used to apply authentication configuration to account classes in Windows Server 2012 R2 domains. Active Directory account classes are:

- User
- Computer
- Managed Service Account and group Managed Service Account (GMSA)

Quick Kerberos refresher

The Kerberos authentication protocol consists of three types of exchanges, also known as subprotocols:

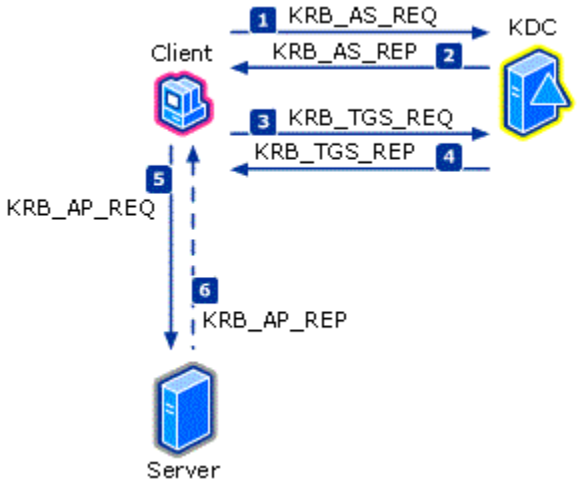

- The Authentication Service (AS) Exchange (KRB_AS_*)
- The Ticket-Granting Service (TGS) Exchange (KRB_TGS_*)
- The Client/Server (AP) Exchange (KRB_AP_*)

The AS exchange is where the client uses the account's password or private key to create a pre-authenticator to request a ticket-granting ticket (TGT). This happens at user sign-on or the first time a service ticket is needed.

The TGS exchange is where the account's TGT is used to create an authenticator to request a service ticket. This happens when an authenticated connection is needed.

The AP exchange occurs as typically as data inside the application protocol and is not impacted by authentication policies.

For more detailed information, see [How the Kerberos Version 5 Authentication Protocol Works].

Overview

Authentication policies complement Protected Users by providing a way to apply configurable restrictions to accounts and by providing restrictions for accounts for services and computers. Authentication policies are enforced during either the AS exchange or the TGS exchange.

You can restrict initial authentication or the AS exchange by configuring:

- A TGT lifetime
- Access control conditions to restrict user sign-on, which must be met by devices from which the AS exchange is coming

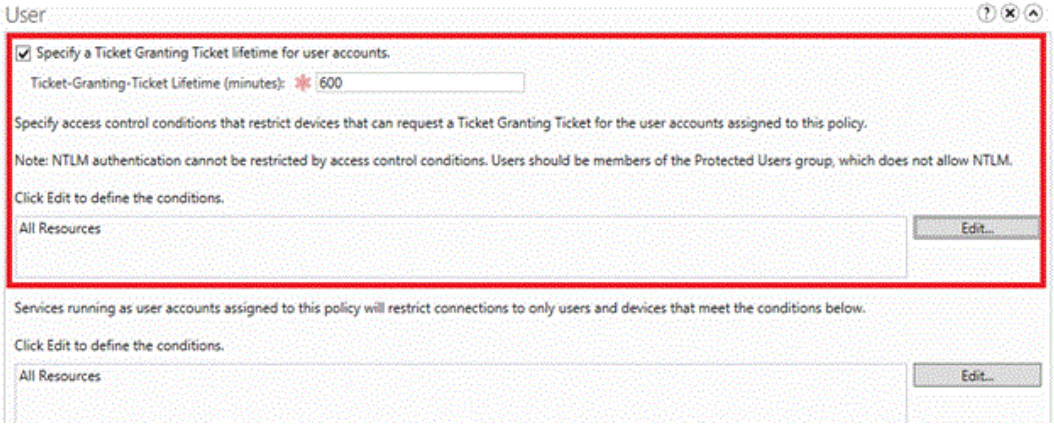

You can restrict service ticket requests through a ticket-granting service (TGS) exchange by configuring:

- Access control conditions which must be met by the client (user, service, computer) or device from which the TGS exchange is coming

Requirements for using authentication policies

Policy	Requirements
Provide custom TGT lifetimes	Windows Server 2012 R2 domain functional level account domains
Restrict user sign-on	• Windows Server 2012 R2 domain functional level account domains with Dynamic Access Control support • Windows 8, Windows 8.1, Windows Server 2012 or Windows Server 2012 R2 devices with Dynamic Access Control support

Policy	Requirements
Restrict service ticket issuance that is based on user account and security groups	Windows Server 2012 R2 domain functional level resource domains
Restrict service ticket issuance based on user claims or device account, security groups, or claims	Windows Server 2012 R2 domain functional level resource domains with Dynamic Access Control support

Restrict a user account to specific devices and hosts

A high-value account with administrative privilege should be a member of the **Protected Users** group. By default, no accounts are members of the **Protected Users** group. Before you add accounts to the group, configure domain controller support and create an audit policy to ensure that there are no blocking issues.

Configure domain controller support

The user's account domain must be at Windows Server 2012 R2 domain functional level (DFL). Ensure all the domain controllers are Windows Server 2012 R2, and then use Active Directory Domains and Trusts to raise the DFL to Windows Server 2012 R2.

To configure support for Dynamic Access Control

1. In the Default Domain Controllers Policy, click **Enabled** to enable **Key Distribution Center (KDC) client support for claims, compound authentication and Kerberos armoring** in Computer Configuration | Administrative Templates | System | KDC.

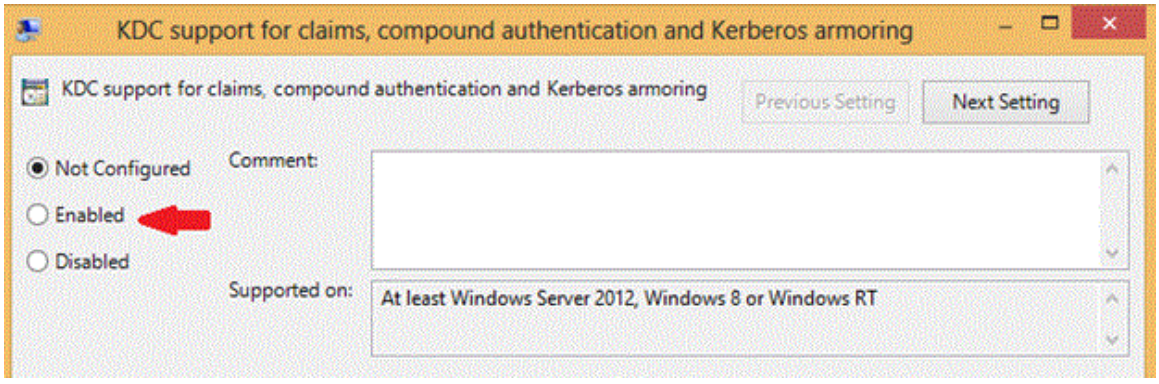

2. Under **Options**, in the drop-down list box, select **Always provide claims**.

> **Note**
> **Supported** can also be configured, but because the domain is at Windows Server 2012 R2 DFL, having the DCs always provide claims will allow user claims-based access checks to occur when using non-claims aware devices and hosts to connect to claims-aware services.

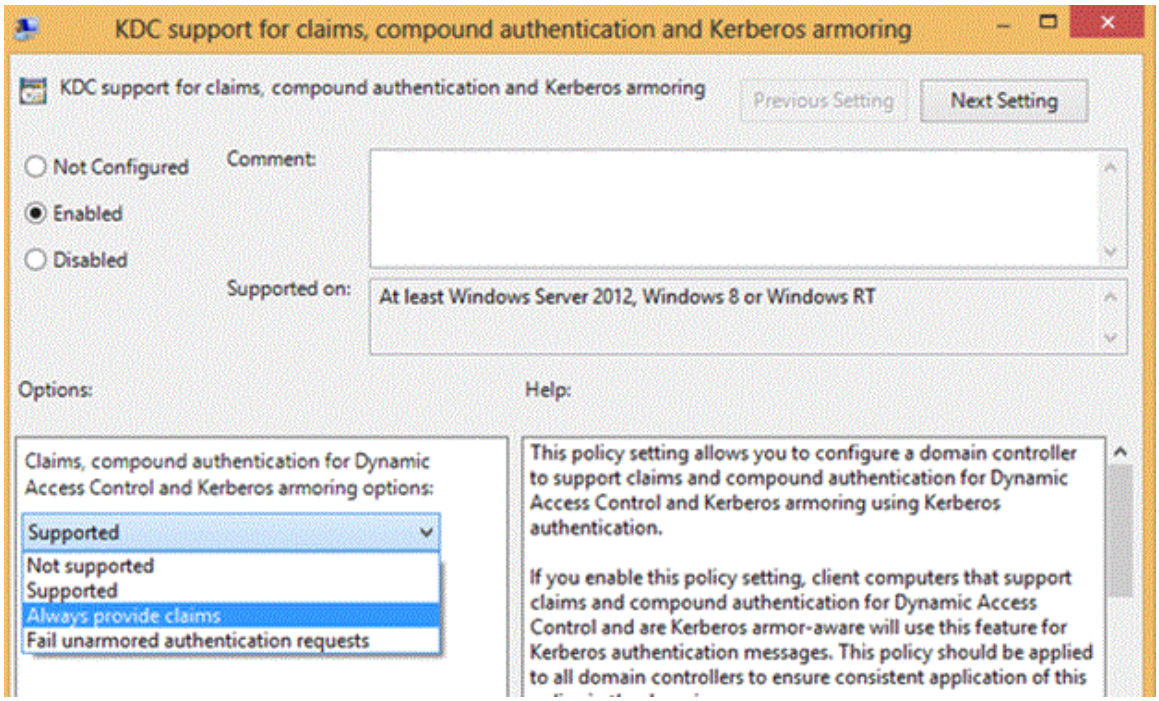

Warning

Configuring **Fail unarmored authentication requests** will result in authentication failures from any operating system which does not support Kerberos armoring, such as Windows 7 and previous operating systems, or operating systems beginning with Windows 8, which have not been explicitly configured to support it.

Create a user account audit for authentication policy with ADAC

1. Open Active Directory Administrative Center (ADAC).

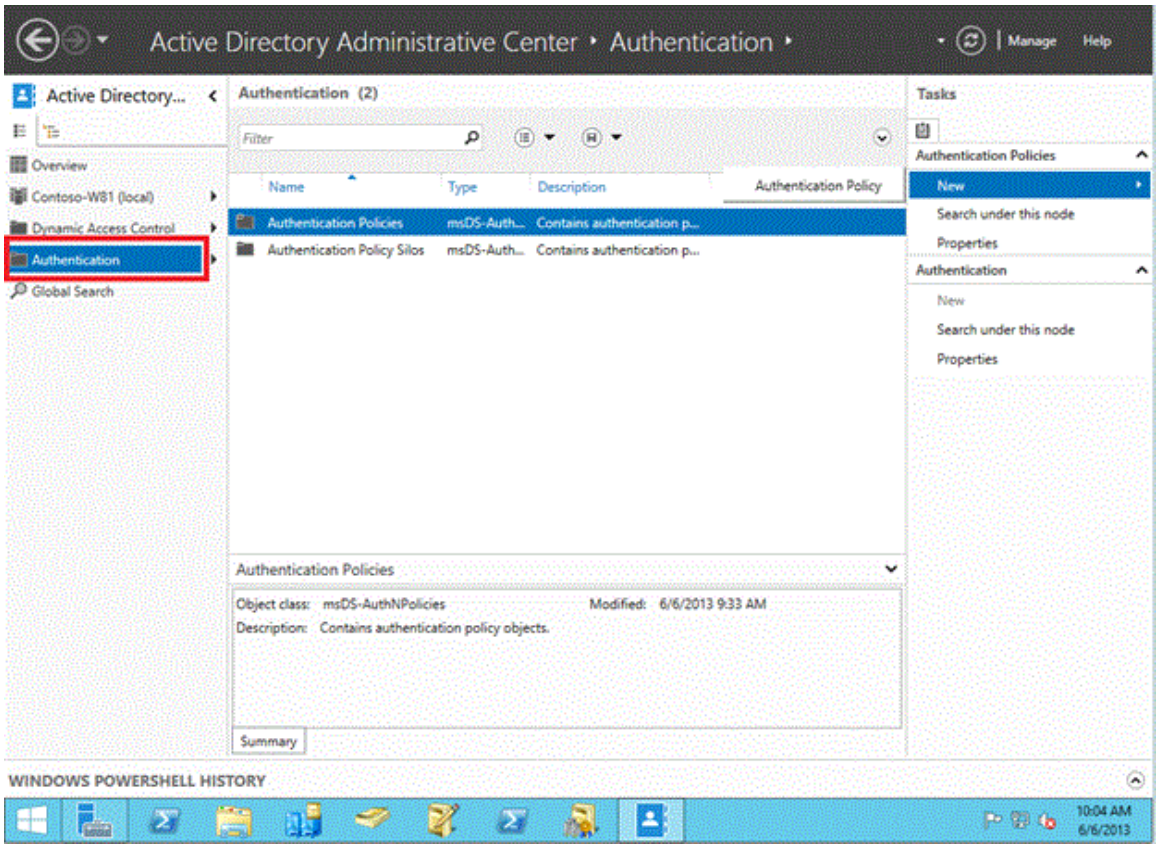

Note

The selected **Authentication** node is visible for domains which are at Windows Server 2012 R2 DFL. If the node does not appear, then try again by using a domain administrator account from a domain that is at Windows Server 2012 R2 DFL.

2. Click **Authentication Policies**, and then click **New** to create a new policy.

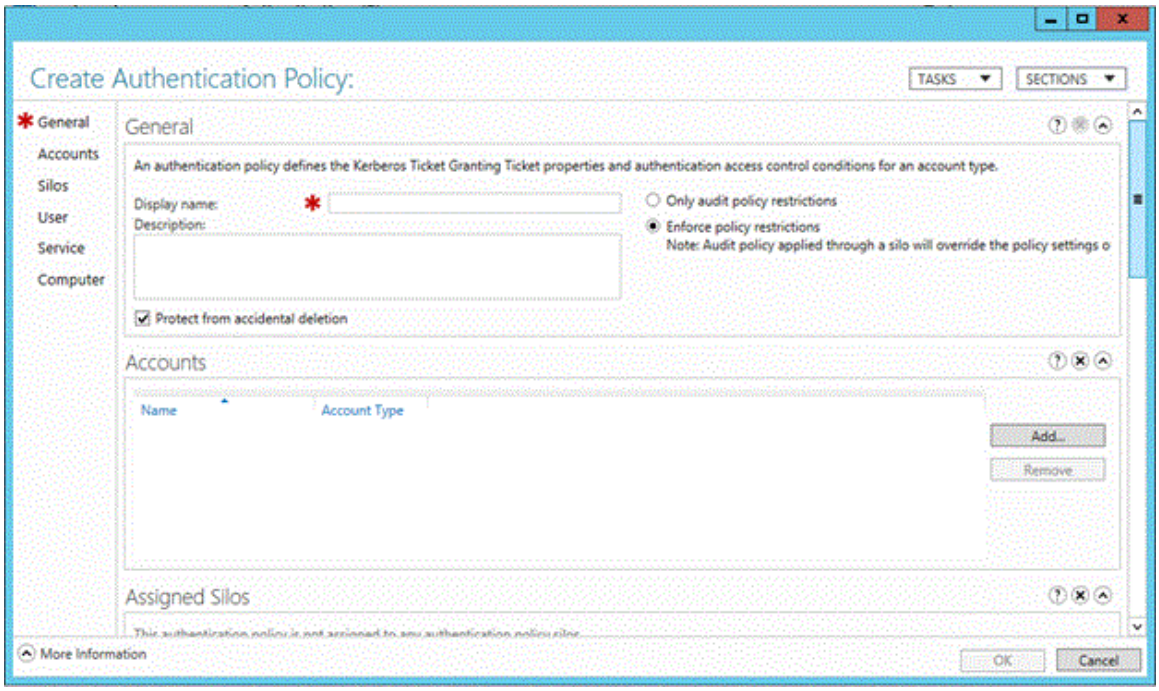

Authentications Policies must have a display name and are enforced by default.

3. To create an audit-only policy, click **Only audit policy restrictions**.

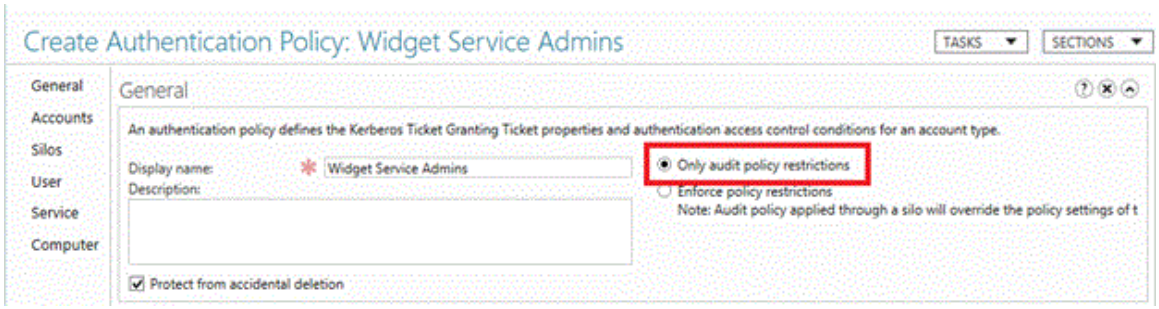

Authentication policies are applied based on the Active Directory account type. A single policy can apply to all three account types by configuring settings for each type. Account types are:

- User
- Computer
- Managed Service Account and Group Managed Service Account

If you have extended the schema with new principals that can be used by the Key Distribution Center (KDC), then the new account type is classified from the closest derived account type.

4. To configure a TGT lifetime for user accounts, select the **Specify a Ticket-Granting Ticket lifetime for user accounts** check box and enter the time in minutes.

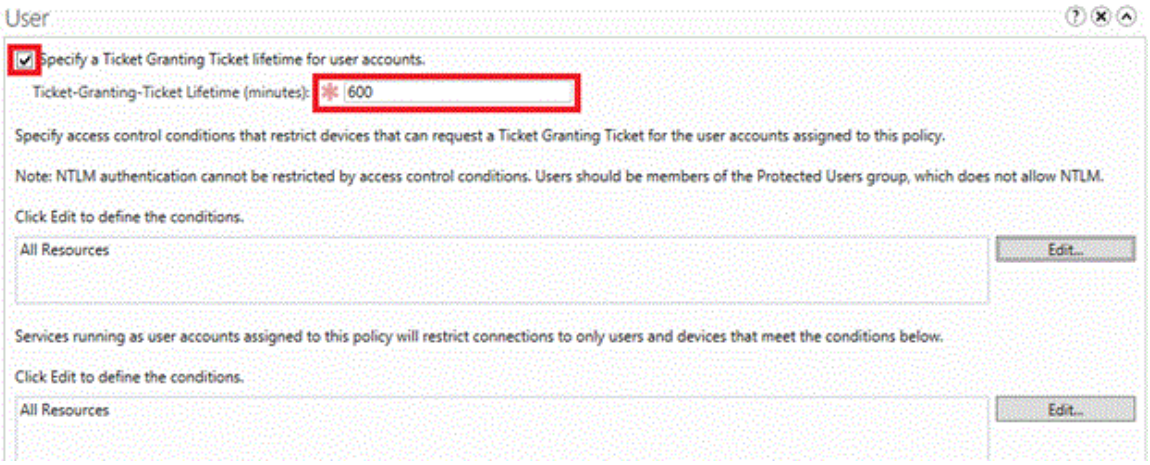

For example, if you want a 10-hour maximum TGT lifetime, enter **600** as shown. If no TGT lifetime is configured, then if the account is a member of the **Protected Users** group, the TGT lifetime and renewal is 4 hours. Otherwise, TGT lifetime and renewal are based on the domain policy as seen in the following Group Policy Management Editor window for a domain with default settings.

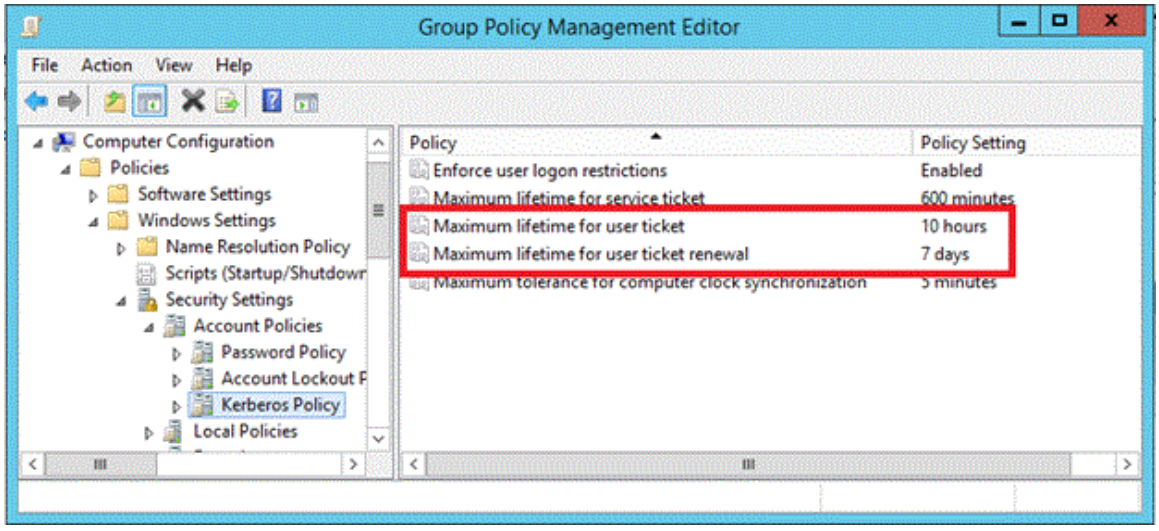

5. To restrict the user account to select devices, click **Edit** to define the conditions that are required for the device.

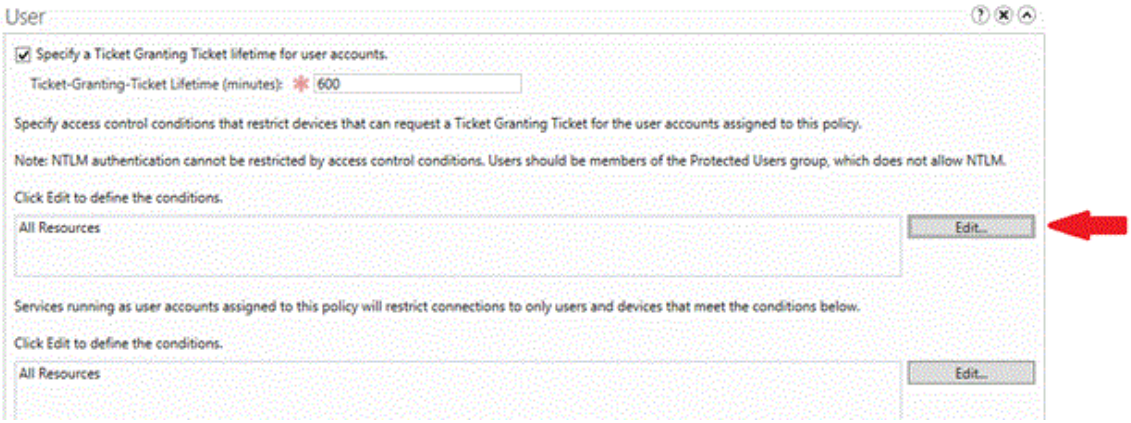

6. In the **Edit Access Control Conditions** window, click **Add a condition**.

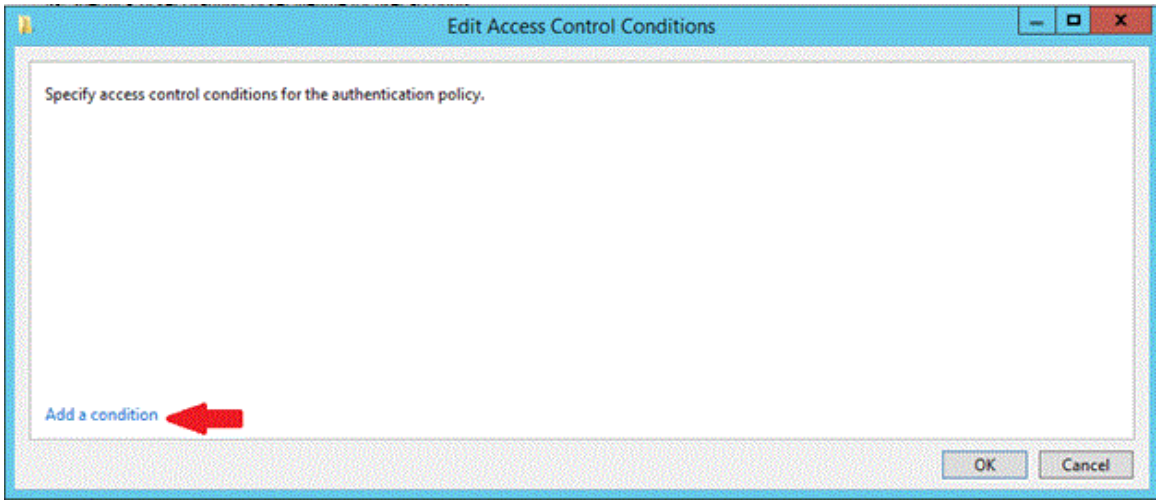

Add computer account or group conditions

1. To configure computer accounts or groups, in the drop-down list, select the drop-down list box **Member of each** and change to **Member of any**.

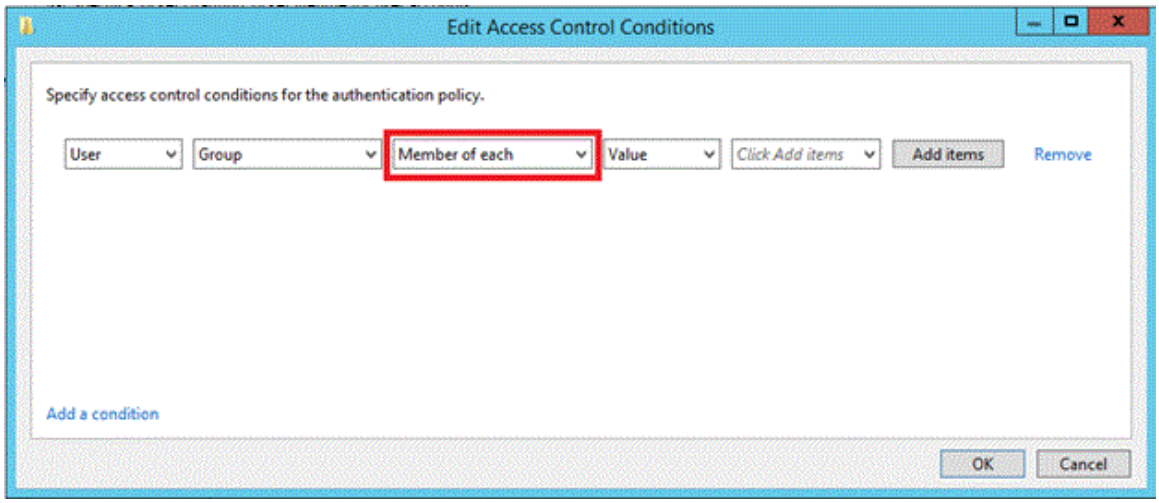

Note

This access control defines the conditions of the device or host from which the user signs on. In access control terminology, the computer account for the device or host is the user, which is why **User** is the only option.

2. Click **Add items**.

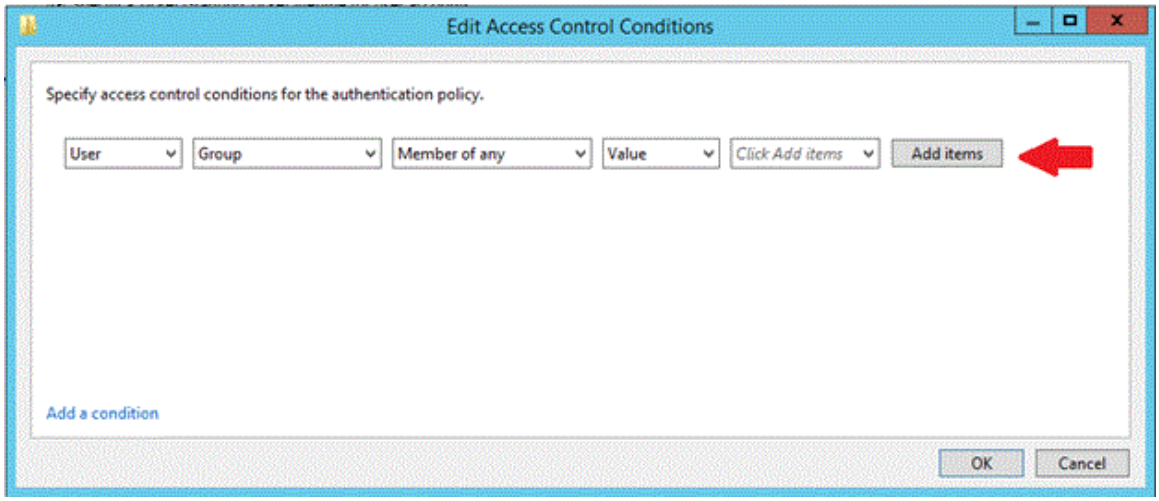

3. To change object types, click **Object Types**.

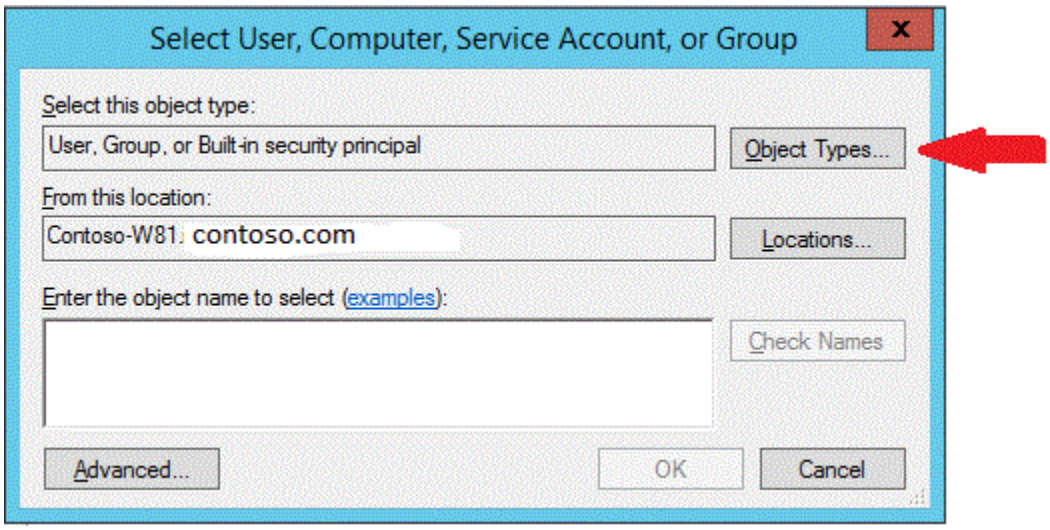

4. To select computer objects in Active Directory, click **Computers**, and then click **OK**.

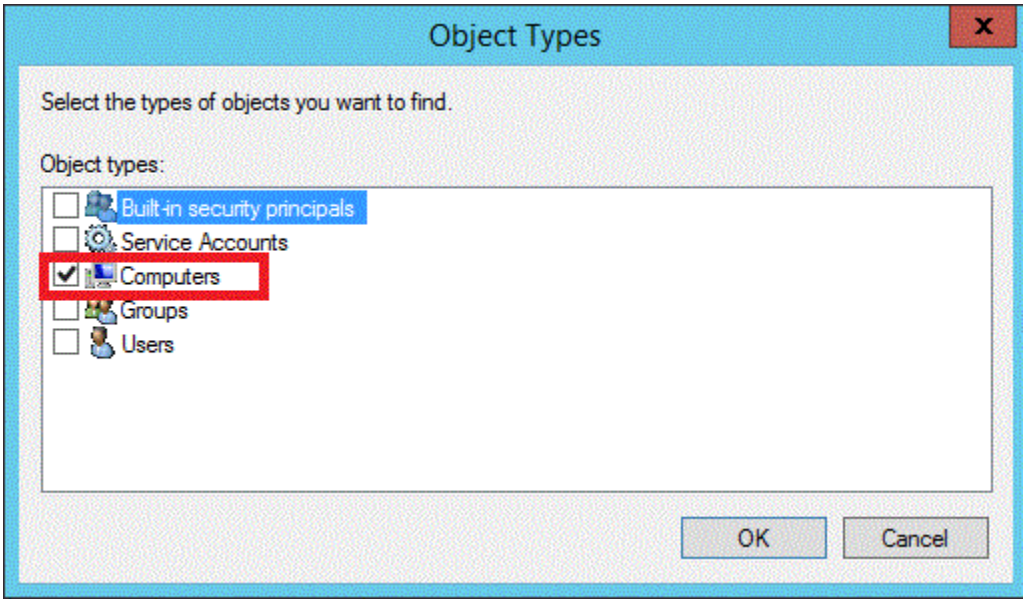

5. Type the name of the computers to restrict the user, and then click **Check Names**.

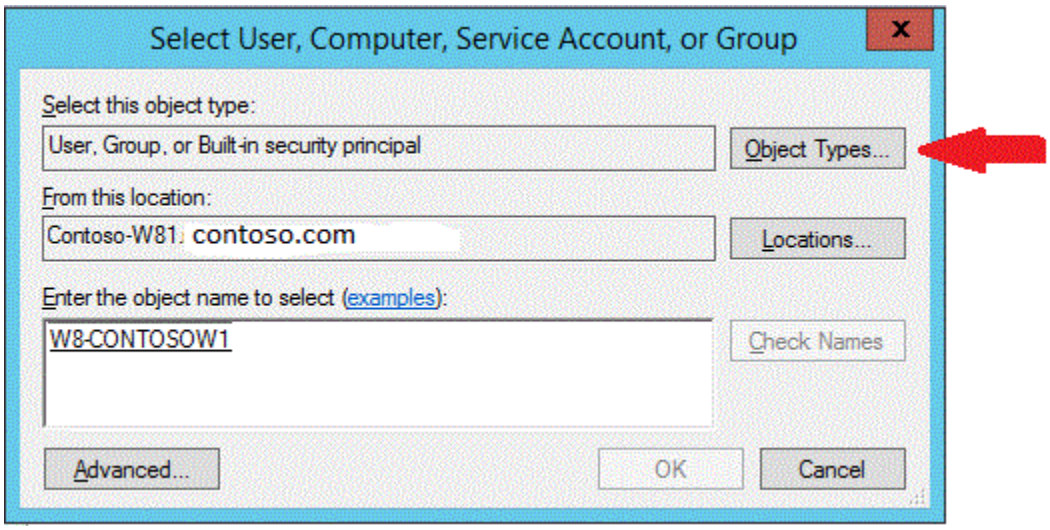

6. Click OK and create any other conditions for the computer account.

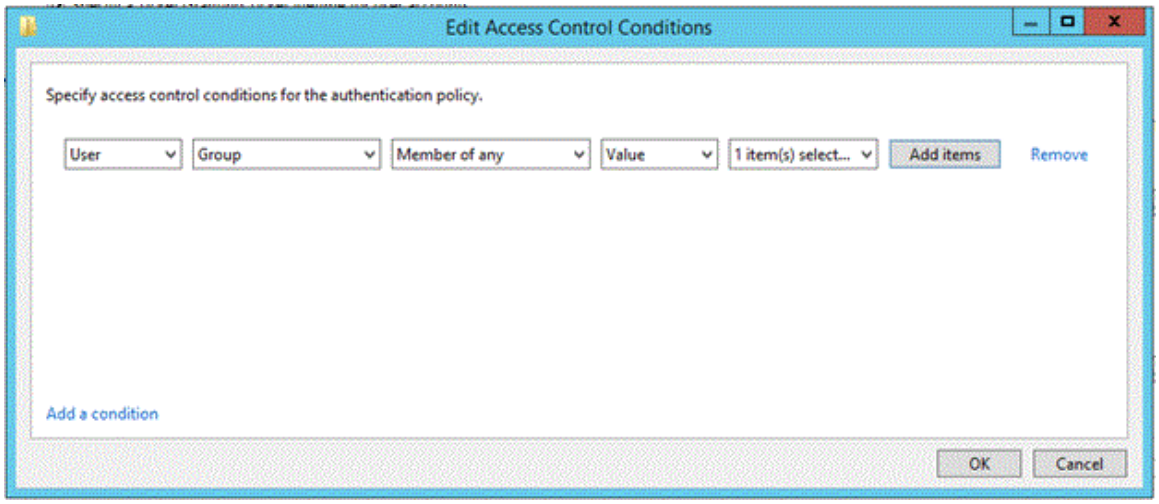

7. When done, then click **OK** and the defined conditions will appear for the computer account.

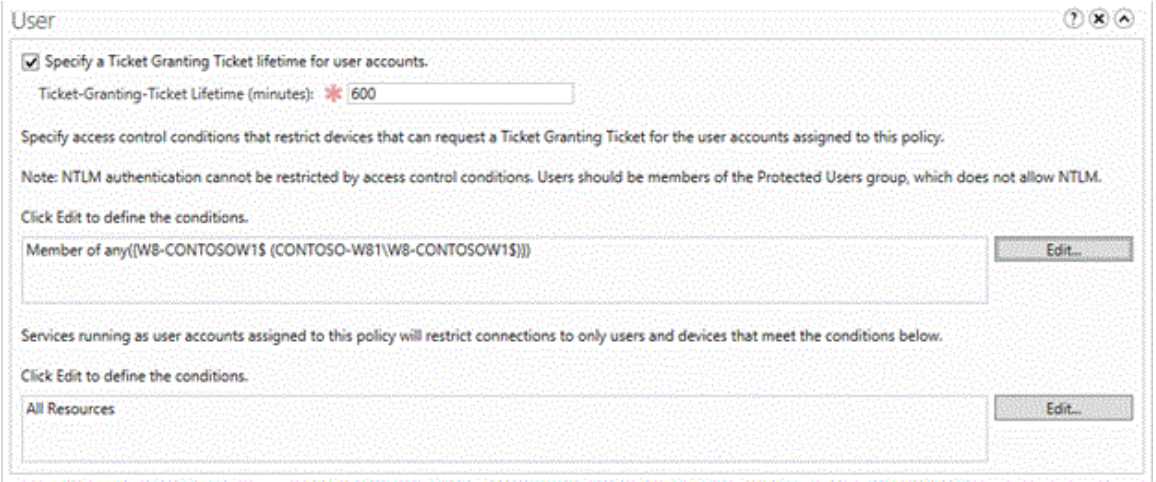

Add computer claim conditions

1. To configure computer claims, drop-down Group to select the claim.

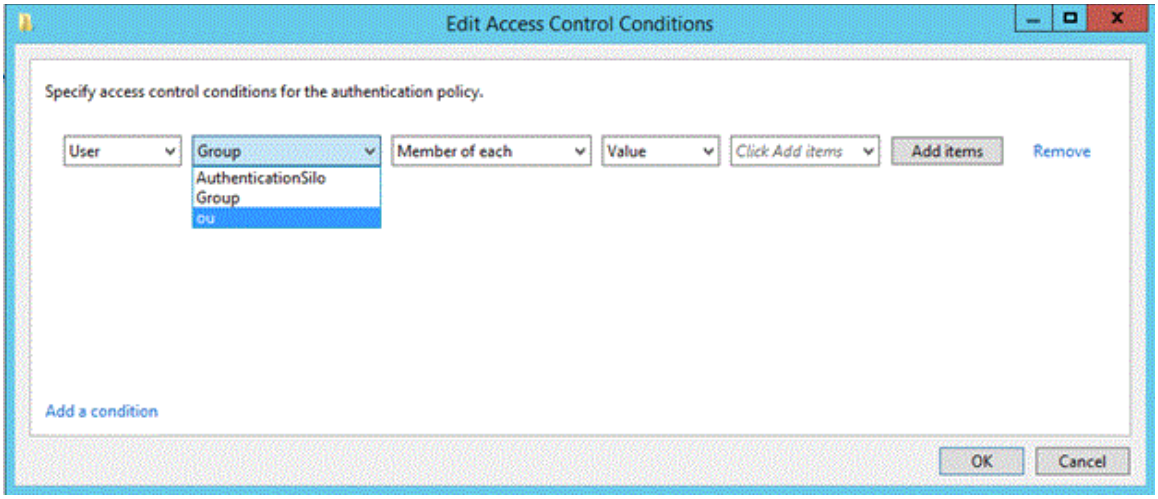

   Claims are only available if they are already provisioned in the forest.

2. Type the name of OU, the user account should be restricted to sign on.

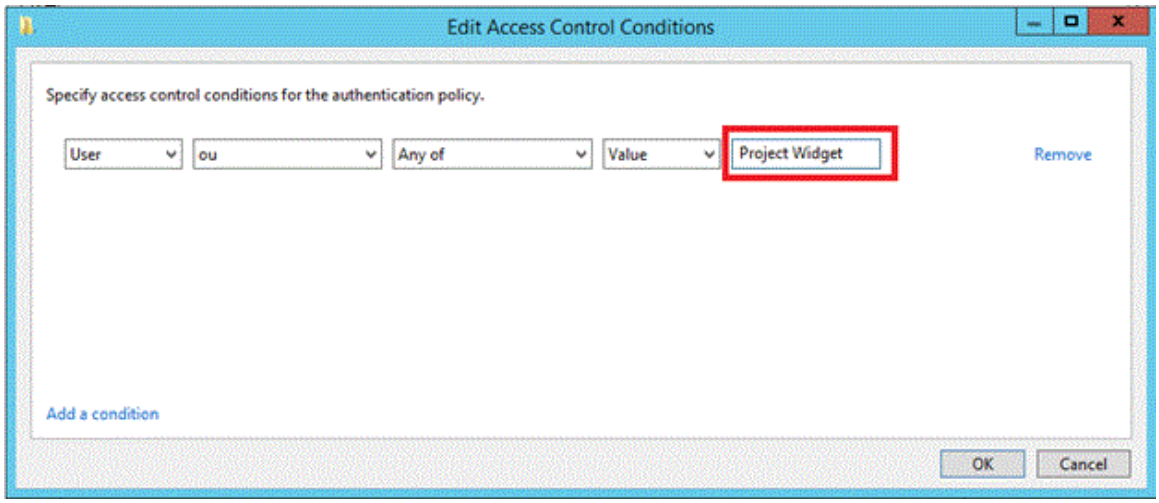

3. When done, then click OK and the box will show the conditions defined.

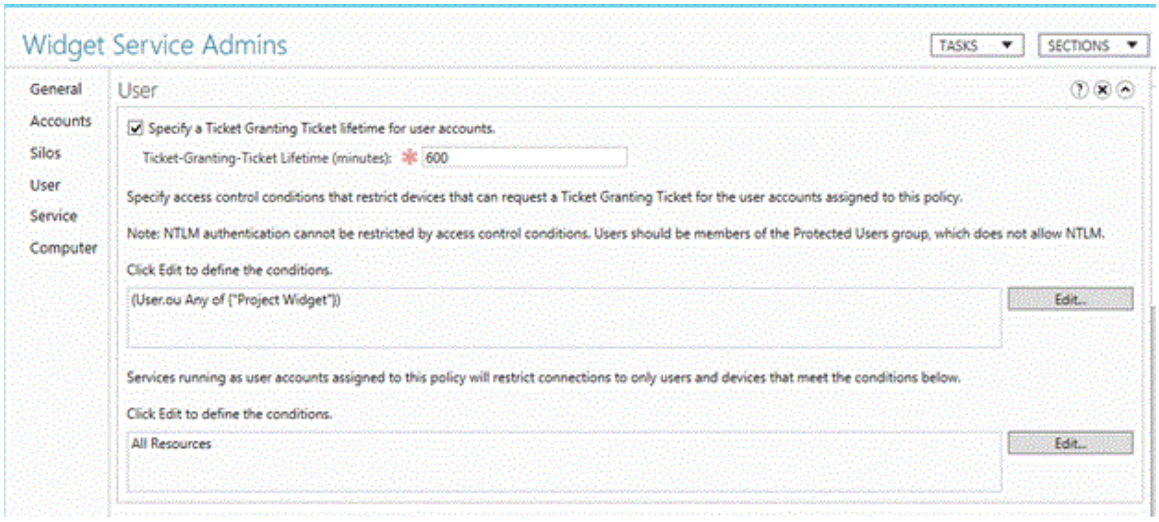

Troubleshoot missing computer claims

If the claim has been provisioned, but is not available, it might only be configured for **Computer** classes.

Let's say you wanted to restrict authentication based on the organizational unit (OU) of the computer, which was already configured, but only for **Computer** classes.

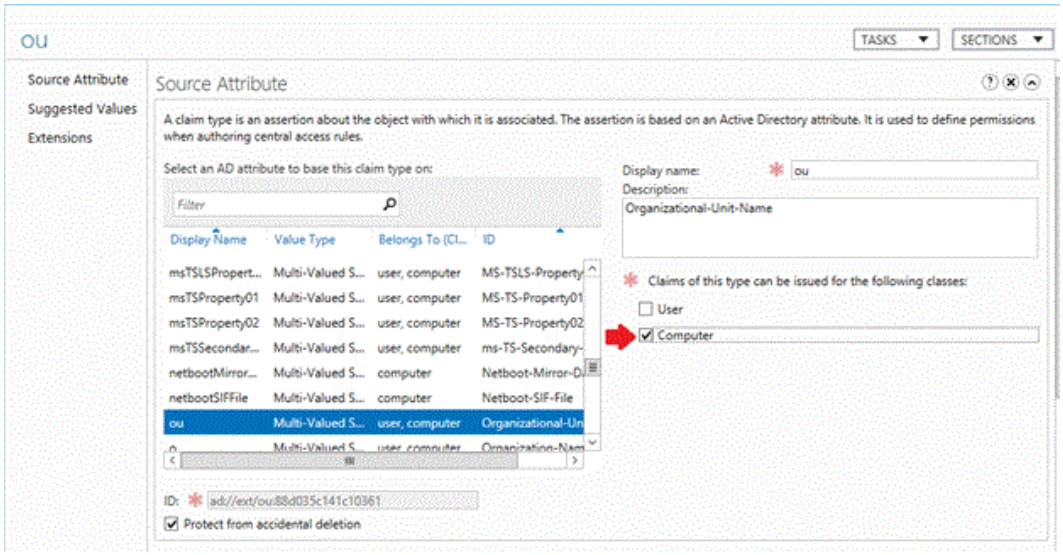

For the claim to be available to restrict User sign-on to the device, select the **User** check box.

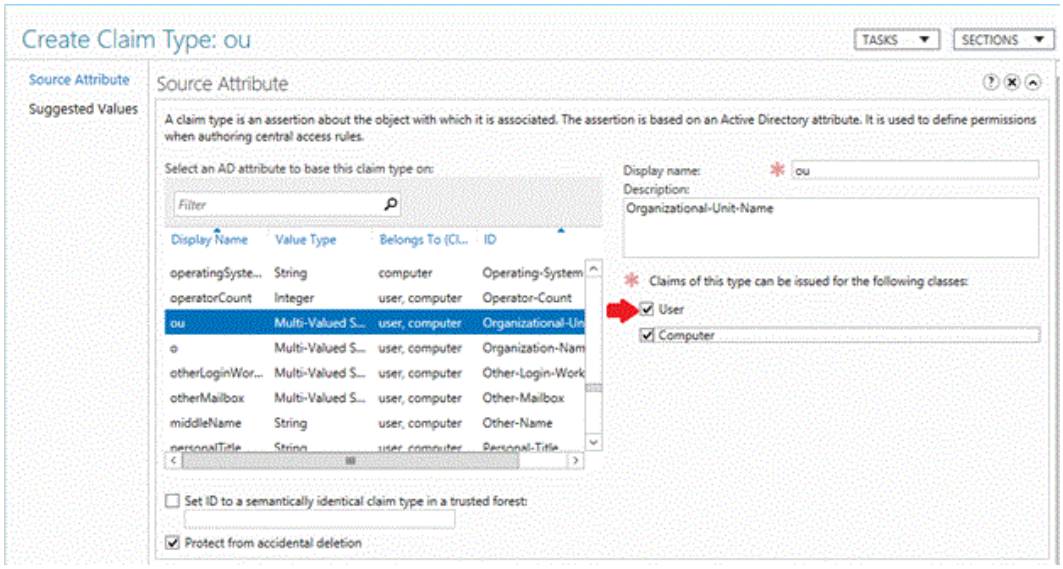

Provision a user account with an authentication policy with ADAC

1. From the **User** account, click **Policy**.

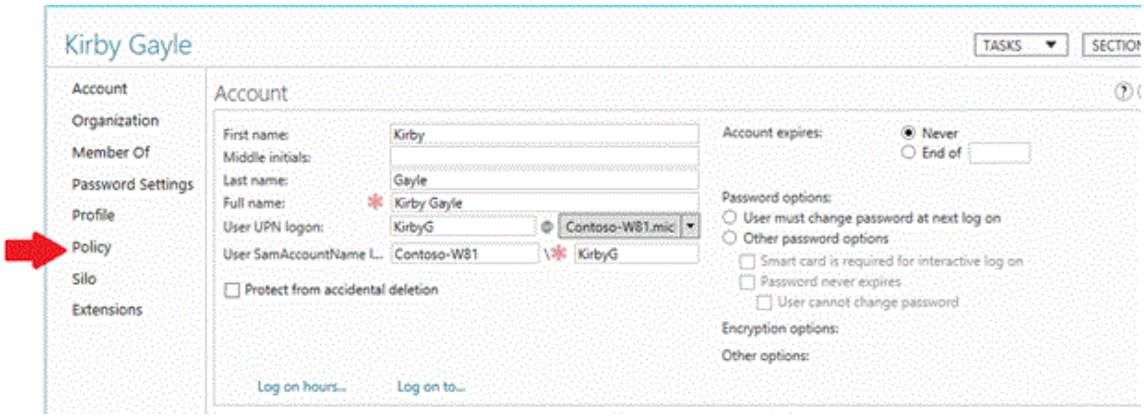

2. Select the **Assign an authentication policy to this account** check box.

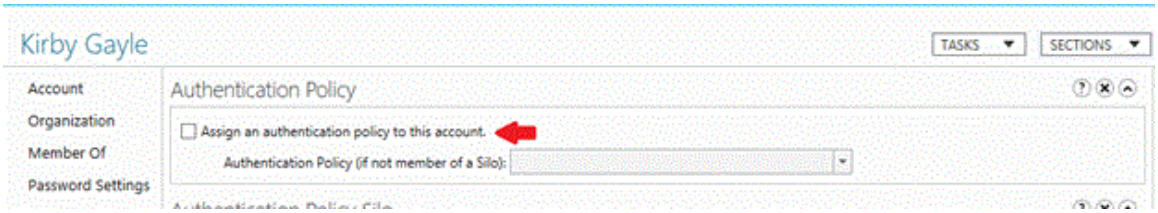

3. Then select the authentication policy to apply to the user.

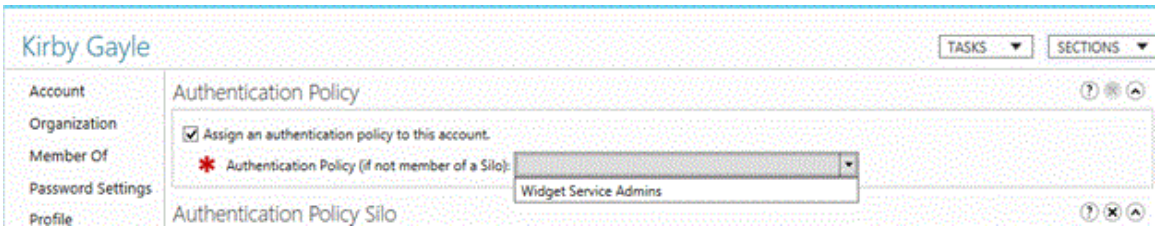

Configure Dynamic Access Control support on devices and hosts

You can configure TGT lifetimes without configuring Dynamic Access Control (DAC). DAC is only needed for checking AllowedToAuthenticateFrom and AllowedToAuthenticateTo.

Using either Group Policy or Local Group Policy Editor, enable **Kerberos client support for claims, compound authentication and Kerberos armoring** in Computer Configuration | Administrative Templates | System | Kerberos:

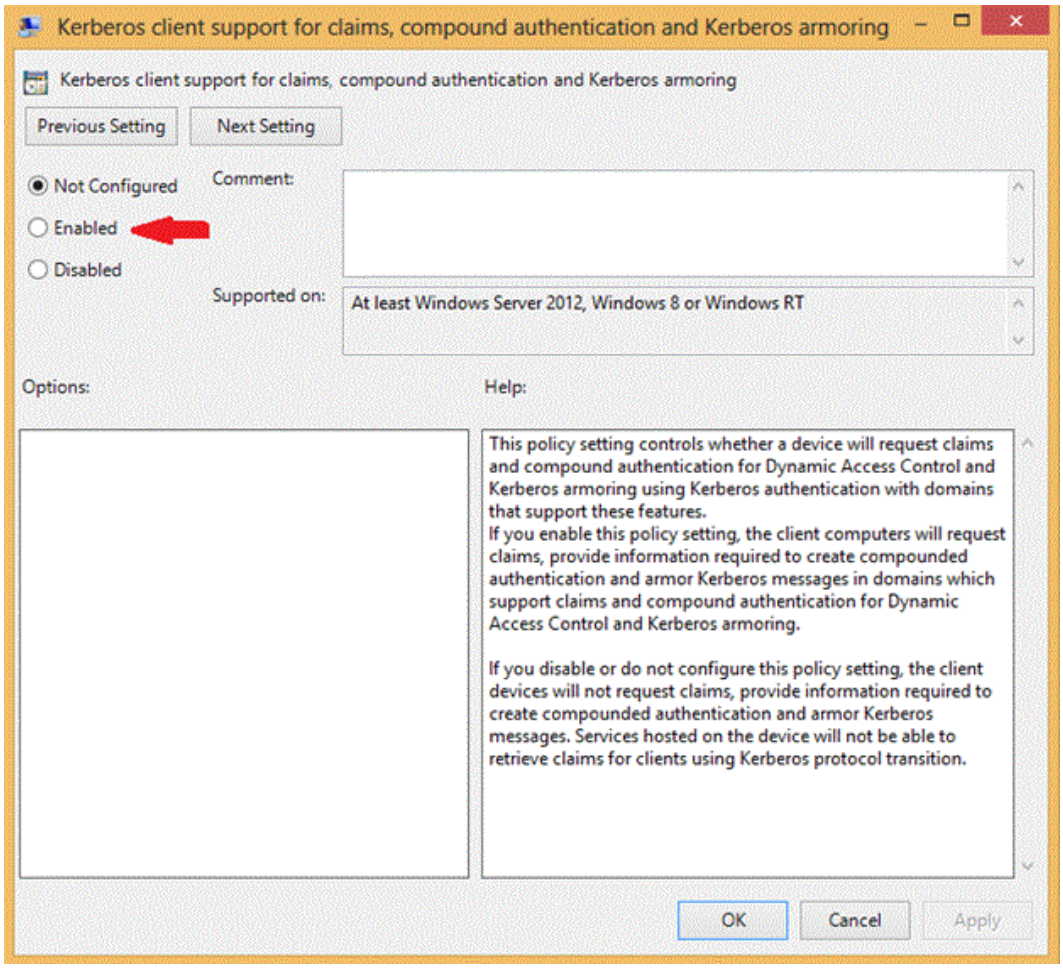

Troubleshoot Authentication Policies

Determine the accounts that are directly assigned an Authentication Policy

The accounts section in the Authentication Policy shows the accounts that have directly applied the policy.

Use the Authentication Policy Failures – Domain Controller administrative log

A new **Authentication Policy Failures – Domain Controller** administrative log under **Applications and Services Logs** > **Microsoft** > **Windows** > **Authentication** has been created to make it easier to discover failures due to Authentication Policies. The log is disabled by default. To enable it, right-click the log name and click **Enable Log**. The new events are very similar in content to the existing Kerberos TGT and service ticket auditing events. For more information about these events, see Authentication Policies and Authentication Policy Silos.

Manage authentication policies by using Windows PowerShell

This command creates an authentication policy named **TestAuthenticationPolicy**. The **UserAllowedToAuthenticateFrom** parameter specifies the devices from which users can authenticate by an SDDL string in the file named someFile.txt.

```
PS C:\> New-ADAuthenticationPolicy testAuthenticationPolicy -
UserAllowedToAuthenticateFrom (Get-Acl .\someFile.txt).sddl
```

This command gets all authentication policies that match the filter that the **Filter** parameter specifies.

```
PS C:\> Get-ADAuthenticationPolicy -Filter "Name -like 'testADAuthenticationPolicy*'" -
Server Server02.Contoso.com
```

This command modifies the description and the **UserTGTLifetimeMins** properties of the specified authentication policy.

```
PS C:\> Set-ADAuthenticationPolicy -Identity ADAuthenticationPolicy1 -Description
"Description" -UserTGTLifetimeMins 45
```

This command removes the authentication policy that the **Identity** parameter specifies.

```
PS C:\> Remove-ADAuthenticationPolicy -Identity ADAuthenticationPolicy1
```

This command uses the **Get-ADAuthenticationPolicy** cmdlet with the **Filter** parameter to get all authentication policies that are not enforced. The result set is piped to the **Remove-ADAuthenticationPolicy** cmdlet.

```
PS C:\> Get-ADAuthenticationPolicy -Filter 'Enforce -eq $false' | Remove-
ADAuthenticationPolicy
```

Authentication policy silos

Authentication Policy Silos is a new container (objectClass msDS-AuthNPolicySilos) in AD DS for user, computer, and service accounts. They help protect high-value accounts. While all organizations need to protect members of Enterprise Admins, Domain Admins and Schema Admins groups because those accounts could be used by an attacker to access anything in the forest, other accounts may also need protection.

Some organizations isolate workloads by creating accounts that are unique to them and by applying Group Policy settings to limit local and remote interactive logon and administrative privileges. Authentication policy silos complement this work by creating a way to define a relationship between User, Computer and managed Service accounts. Accounts can only belong to one silo. You can configure authentication policy for each type of account in order to control:

1. Non-renewable TGT lifetime
2. Access control conditions for returning TGT (Note: cannot apply to systems because Kerberos armoring is required)
3. Access control conditions for returning service ticket

Additionally, accounts in an authentication policy silo have a silo claim, which can be used by claims-aware resources such as file servers to control access.

A new security descriptor can be configured to control issuing service ticket based on:

- User, user's security groups, and/or user's claims
- Device, device's security group, and/or device's claims

Getting this information to the resource's DCs requires Dynamic Access Control:

- User claims:
  - Windows 8 and later clients supporting Dynamic Access Control
  - Account domain supports Dynamic Access Control and claims
- Device and/or device security group:
  - Windows 8 and later clients supporting Dynamic Access Control
  - Resource configured for compound authentication
- Device claims:
  - Windows 8 and later clients supporting Dynamic Access Control
  - Device domain supports Dynamic Access Control and claims
  - Resource configured for compound authentication

Authentication policies can be applied to all members of an authentication policy silo instead of to individual accounts, or separate authentication policies can be applied to different types of accounts within a silo. For example, one authentication policy can be applied to highly privileged user accounts, and a different policy can be applied to services accounts. At least one authentication policy must be created before an authentication policy silo can be created.

Note

An authentication policy can be applied to members of an authentication policy silo, or it can be applied independently of silos to restrict specific account scope. For example, to protect a single account or a small set of accounts, a policy can be set on those accounts without adding the accounts to a silo.

You can create an authentication policy silo by using Active Directory Administrative Center or Windows PowerShell. By default, an authentication policy silo only audits silo policies, which is equivalent to specifying the **WhatIf** parameter in Windows PowerShell cmdlets. In this case,

policy silo restrictions do not apply, but audits are generated to indicate whether failures occur if the restrictions are applied.

To create an authentication policy silo by using Active Directory Administrative Center

1. Open **Active Directory Administrative Center**, click **Authentication**, right-click **Authentication Policy Silos**, click **New**, and then click **Authentication Policy Silo**.

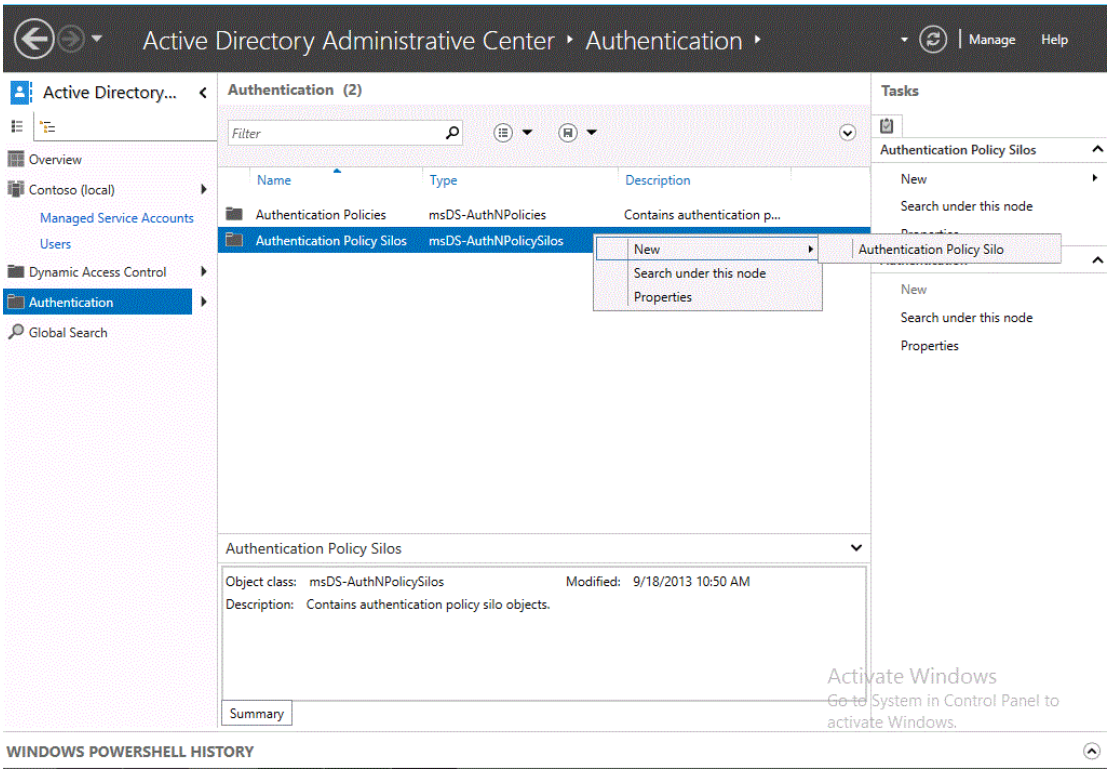

2. In **Display name**, type a name for the silo. In **Permitted Accounts**, click **Add**, type the names of the accounts, and then click **OK**. You can specify users, computers, or service accounts. Then specify whether to use a single policy for all principals or a separate policy for each type of principal, and the name of the policy or policies.

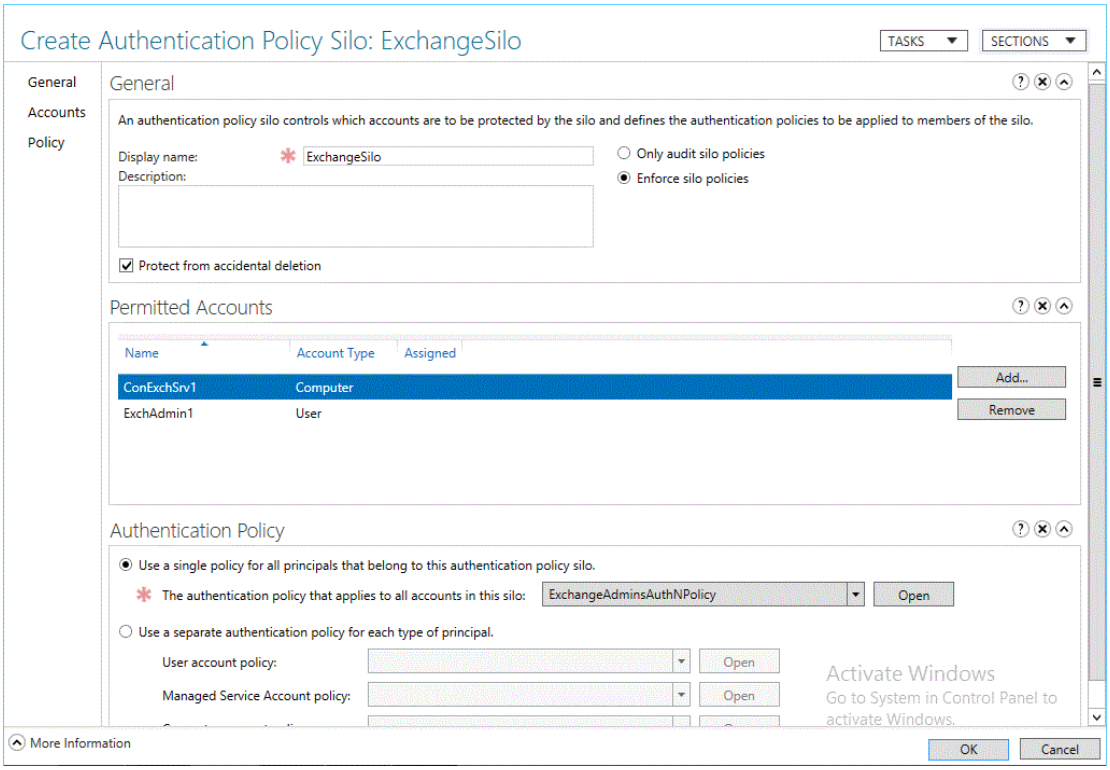

Manage authentication policy silos by using Windows PowerShell

This command creates an authentication policy silo object and enforces it.

```
PS C:\>New-ADAuthenticationPolicySilo -Name newSilo -Enforce
```

This command gets all the authentication policy silos that match the filter that is specified by the **Filter** parameter. The output is then passed to the **Format-Table** cmdlet to display the name of the policy and the value for **Enforce** on each policy.

```
PS C:\>Get-ADAuthenticationPolicySilo -Filter 'Name -like "*silo*"' | Format-Table Name,
Enforce -AutoSize

Name  Enforce
----  -------
silo     True
silos   False
```

This command uses the **Get-ADAuthenticationPolicySilo** cmdlet with the **Filter** parameter to get all authentication policy silos that are not enforced and pipe the result of the filter to the **Remove-ADAuthenticationPolicySilo** cmdlet.

```
PS C:\>Get-ADAuthenticationPolicySilo -Filter 'Enforce -eq $False' | Remove-
ADAuthenticationPolicySilo
```

This command grants access to the authentication policy silo named *Silo* to the user account named *User01*.

```
PS C:\>Grant-ADAuthenticationPolicySiloAccess -Identity Silo -Account User01
```

This command revokes access to the authentication policy silo named *Silo* for the user account named *User01*. Because the **Confirm** parameter is set to **$False**, no confirmation message appears.

```
PS C:\>Revoke-ADAuthenticationPolicySiloAccess –Identity Silo –Account User01 –
Confirm:$False
```

This example first uses the **Get-ADComputer** cmdlet to get all computer accounts that match the filter that the **Filter** parameter specifies. The output of this command is passed to **Set-ADAccountAuthenticatinPolicySilo** to assign the authentication policy silo named *Silo* and the authentication policy named *AuthenticationPolicy02* to them.

```
PS C:\>Get-ADComputer –Filter 'Name –like "newComputer*"' | Set-
ADAccountAuthenticationPolicySilo –AuthenticationPolicySilo Silo –AuthenticationPolicy
AuthenticationPolicy02
```

Active Directory Federation Services Overview

This topic provides an overview of Active Directory Federation Services (AD FS) in Windows Server® 2012 and Windows Server® 2012 R2.

Role description

AD FS provides simplified, secured identity federation and Web single sign-on (SSO) capabilities for end users who want to access applications within an AD FS-secured enterprise, in federation partner organizations, or in the cloud.

In Windows Server® 2012 R2, AD FS includes a federation service role service that acts as an identity provider (authenticates users to provide security tokens to applications that trust AD FS) or as a federation provider (consumes tokens from other identity providers and then provides security tokens to applications that trust AD FS).

The function of providing extranet access to applications and services that are secured by AD FS is now performed by a new Remote Access role service called Web Application Proxy. This is a departure from the prior versions of Windows Server in which this function was handled by an AD

FS federation server proxy. Web Application Proxy is a server role designed to provide access for the AD FS-related extranet scenario and other extranet scenarios. For more information on Web Application Proxy, see Web Application Proxy Overview.

Practical applications

AD FS simplifies access to systems and applications by using a claims-based authentication and access authorization mechanism to maintain application security. Prior versions of AD FS were used for the following:

- Providing your employees or customers with a web-based, SSO experience when accessing claims-based applications within your enterprise.
- Providing your employees or customers with a web-based, SSO experience to access resources in any federation partner organization.
- Providing your employees or customers with a Web-based, SSO experience when remote accessing internally hosted Web sites or services.
- Providing your employees or customers with a web-based, SSO experience when accessing resources or services in the cloud.

AD FS in Windows Server® 2012 R2 adds additional practical applications for AD FS, including the following:

- Device workplace join for SSO and seamless second factor authentication. This enables organizations to allow access from devices that are personally owned by users and manage the risk when providing this access. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
- Managing risk with multi-factor access control. AD FS provides a rich level of authorization that controls who has access to what applications. This can be based on user attributes (UPN, email, security group membership, authentication strength, etc.), device attributes (whether the device is workplace joined) or request attributes (network location, IP address, or user agent). For more information, see Overview: Manage Risk with Multi-Factor Access Control.
- Managing risk with additional multi-factor authentication for sensitive applications. AD FS allows you to control policies to potentially require multi-factor authentication on a per application basis. In addition, AD FS provides extensibility points for any multi-factor vendor to integrate deeply for a secure and seamless multi-factor experience for end users. For more information, see Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.
- Providing authentication and authorization capabilities for accessing web application resources from the extranet that are protected by the Web Application Proxy. For more information, see Installing and Configuring Web Application Proxy for Publishing Internal Applications, Overview: Connect to Applications and Services from Anywhere with Web Application Proxy, and Web Application Proxy Overview.

AD FS in Windows Server 2012

For Windows Server 2012, the AD FS server role includes the same functionality and feature set that is available in AD FS 2.0. It also includes the following list of new functionality that was not available in AD FS 2.0:

- **Improved installation experience using Server Manager –** With AD FS 2.0, you had to download and install the AD FS 2.0 software to deploy your AD FS server infrastructure. However, in Windows Server 2012, you install the AD FS server role using Server Manager. Server Manager provides improved AD FS configuration wizard pages that perform server validation checks before you continue with the AD FS server role installation and will automatically list and install all the services that AD FS depends on during the AD FS server role installation.
- **Additional Windows PowerShell cmdlet tools -** In addition to the Windows PowerShell based management capabilities provided in AD FS 2.0, AD FS in Windows Server 2012 and Windows Server® 2012 R2, includes new cmdlets for installing the AD FS server role and for initial configuration of the federation server and federation server proxy.

AD FS in Windows Server 2012 R2

The following sections summarize numerous changes that were made to AD FS in Windows Server® 2012 R2 in order to support newer practical applications of AD FS as well as to enhance existing functionality.

Enable users to access resources on their personal devices from anywhere

- Workplace join that enables users to join their personal devices to corporate Active Directory and as a result gain access and seamless experiences when accessing corporate resources from these devices.
- Pre-authentication of resources inside the corporate network that are protected by the Web Application proxy and accessed from the internet.
- Password change to enable users to change their password from any workplace joined device when their password has expired so that they can continue to access resources.

Enhanced access control risk management tools

Managing risk is an important aspect of governance and compliance in every IT organization. There are numerous access control risk management enhancements in AD FS in Windows Server® 2012 R2, including the following:

- Flexible controls based on network location to govern how a user authenticates to access an AD FS-secured application.
- Flexible policy to determine if a user needs to perform multi-factor authentication based on the user's data, device data, and network location.

- Per-application control to ignore SSO and force the user to provide credentials every time they access a sensitive application.
- Flexible per-application access policy based on user data, device data, or network location.
- AD FS Extranet Lockout, which enables administrators to protect Active Directory accounts from brute force attacks from the internet.
- Access revocation for any workplace joined device that is disabled or deleted in Active Directory.

Simplified deployment experience

Deploying AD FS in Windows Server® 2012 R2 is simplified by the following enhancements:

- AD FS is no longer dependent on IIS. This offers enhanced performance and reduces the foot print of services, especially when AD FS is installed on Active Directory domain controllers.
- Remote installation and configuration through Server Manager.
- UI support for installing AD FS with SQL Server
- Group Managed Service Account support. This enables AD FS to be run with service accounts without managing expiring service account passwords.
- SQL Server merge replication support when deploying AD FS across globally dispersed datacenters.

Important

Note that in Windows Server® 2012 R2, the ‘stand-alone’ mode for AD FS setup has been removed.

Enhanced sign-in with AD FS experience

The following are new AD FS capabilities in Windows Server® 2012 R2 that enable administrator to customize and enhance the sign-in experience:

- Unified customization of the AD FS service, where the changes are made once and then automatically propagated to the rest of the AD FS federation servers in a given farm.
- Updated sign-in pages that look modern and cater to different form factors automatically.
- Support for automatic fallback to forms-based authentication for devices that are not joined to the corporate domain but are still used generate access requests from within the corporate network (intranet).
- Simple controls to customize the company logo, illustration image, standard links for IT support, home page, privacy, etc.
- Customization of description messages in the sign-in pages.
- Customization of web themes.
- Home Realm Discovery (HRD) based on organizational suffix of the user for enhanced privacy of a company’s partners.
- HRD filtering on a per-application basis to automatically pick a realm based on the application.

- One-click error reporting for easier IT troubleshooting.
- Customizable error messages.
- User authentication choice when more than one authentication provider is available.

Enable developers to build modern applications

ADFS in Windows Server 2012 now supports the OAuth Authorization Grant profile with support for refresh tokens to enable modern applications that use RESTful services. ADFS also supports issuing JWT tokens that are compact tokens that are more amenable to transmit to the resources that use the REST pattern.

Other improvements

- Reduction of SSO cookie size with dynamic group SID hydration. This provides a more deterministic cookie size and reduces bloat when a user belongs to many security groups.
- Access to claims that are encoded within user certificates when using certificate authentication. This can help administrators differentiate access based on what type of certificate is used.
- Consistent client-request-id that is logged in all event logs and traces for easier troubleshooting.
- Additional request claims, for example, IP addresses, endpoint addresses, or user agents that can be used to base policy decisions on.
- Password expiry notification as claims that an administrator can configure to send to downstream applications to notify the user when their password is about to expire.

See Also

Active Directory

Getting Started with AD FS

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications
  - Walkthrough Guide: Workplace Join with an iOS Device
  - Walkthrough Guide: Workplace Join with a Windows Device
- Overview: Connect to Applications and Services from Anywhere with Web Application Proxy
  - Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy
- Overview: Manage Risk with Multi-Factor Access Control
  - Walkthrough Guide: Manage Risk with Multi-Factor Access Control
- Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

- Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

See Also

Active Directory Federation Services Overview

Set up the lab environment for AD FS in Windows Server 2012 R2

This topic outlines the steps to configure a test environment that can be used to complete the walkthroughs in the following walkthrough guides:

- Walkthrough Guide: Workplace Join with an iOS Device
- Walkthrough Guide: Workplace Join with a Windows Device
- Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy
- Walkthrough Guide: Manage Risk with Multi-Factor Access Control
- Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

Note
We do not recommend that you install the web server and the federation server on the same computer.

To set up this test environment, complete the following steps:

1. Step 1: Configure the domain controller (DC1)
2. Step 2: Configure the federation server (ADFS1) by using Device Registration Service
3. Step 3: Configure the web server (WebServ1) and a sample claims-based application
4. Step 4: Configure the client computer (Client1)

Step 1: Configure the domain controller (DC1)

For the purposes of this test environment, you can call your root Active Directory domain **contoso.com** and specify **pass@word1** as the administrator password.

- Install the AD DS role service and install Active Directory Domain Services (AD DS) to make your computer a domain controller in Windows Server 2012 R2. This action upgrades your AD DS schema as part of the domain controller creation. For more information and step-by-step instructions, see http://technet.microsoft.com/ library/hh472162.aspx.

Create test Active Directory accounts

After your domain controller is functional, you can create a test group and test user accounts in this domain and add the user account to the group account. You use these accounts to complete the walkthroughs in the walkthrough guides that are referenced earlier in this topic.

Create the following accounts:

- User: **Robert Hatley** with the following credentials: User name: **RobertH** and password: **P@ssword**
- Group: **Finance**

For information about how to create user and group accounts in Active Directory (AD), see http://technet.microsoft.com/library/cc783323%28v=ws.10%29.aspx.

Add the **Robert Hatley** account to the **Finance** group. For information on how to add a user to a group in Active Directory, see http://technet.microsoft.com/library/cc737130%28v=ws.10%29.aspx.

Create a GMSA account

The group Managed Service Account (GMSA) account is required during the Active Directory Federation Services (AD FS) installation and configuration.

To create a GMSA account

1. Open a Windows PowerShell command window and type:

```
Add-KdsRootKey –EffectiveTime (Get-Date).AddHours(-10)
New-ADServiceAccount FsGmsa -DNSHostName adfs1.contoso.com -
ServicePrincipalNames http/adfs1.contoso.com
```

Step 2: Configure the federation server (ADFS1) by using Device Registration Service

To set up another virtual machine, install Windows Server 2012 R2 and connect it to the domain **contoso.com**. Set up the computer after you have joined it to the domain, and then proceed to install and configure the AD FS role.

For a video, see Active Directory Federation Services How-To Video Series: Installing an AD FS Server Farm.

Install a server SSL certificate

You must install a server Secure Socket Layer (SSL) certificate on the ADFS1 server in the local computer store. The certificate MUST have the following attributes:

- Subject Name (CN): adfs1.contoso.com
- Subject Alternative Name (DNS): adfs1.contoso.com
- Subject Alternative Name (DNS): enterpriseregistration.contoso.com

For more information about setting up SSL certificates, see Configure SSL/TLS on a Web site in the domain with an Enterprise CA.

Active Directory Federation Services How-To Video Series: Updating Certificates.

Install the AD FS server role

To install the Federation Service role service

1. Log on to the server by using the domain administrator account administrator@contoso.com.
2. Start Server Manager. To start Server Manager, click **Server Manager** on the Windows **Start** screen, or click **Server Manager** on the Windows taskbar on the Windows desktop. On the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, verify that the target computer is selected, and then click **Next**.
6. On the **Select server roles** page, click **Active Directory Federation Services**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Federation Service (AD FS)** page, click **Next**.
9. After you verify the information on the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.
10. On the **Installation progress** page, verify that everything installed correctly, and then click **Close**.

Configure the federation server

The next step is to configure the federation server.

To configure the federation server

1. On the Server Manager **Dashboard** page, click the **Notifications** flag, and then click **Configure the federation service on the server**.

   The **Active Directory Federation Service Configuration Wizard** opens.
2. On the **Welcome** page, select **Create the first federation server in a federation server farm**, and then click **Next**.
3. On the **Connect to AD DS** page, specify an account with domain administrator rights for the **contoso.com** Active Directory domain that this computer is joined to, and then click **Next**.
4. On the **Specify Service Properties** page, do the following, and then click **Next**:
   - Import the SSL certificate that you have obtained earlier. This certificate is the required service authentication certificate. Browse to the location of your SSL certificate.
   - To provide a name for your federation service, type **adfs1.contoso.com**. This value is the same value that you provided when you enrolled an SSL certificate in Active Directory Certificate Services (AD CS).
   - To provide a display name for your federation service, type **Contoso Corporation**.
5. On the **Specify Service Account** page, select **Use an existing domain user account or group Managed Service Account**, and then specify the GMSA account **fsgmsa** that you created when you created the domain controller.
6. On the **Specify Configuration Database** page, select **Create a database on this server using Windows Internal Database**, and then click **Next**.
7. On the **Review Options** page, verify your configuration selections, and then click **Next**.
8. On the **Pre-requisite Checks** page, verify that all prerequisite checks were successfully completed, and then click **Configure**.
9. On the **Results** page, review the results, check whether the configuration has completed successfully, and then click **Next steps required for completing your federation service deployment**.

Configure Device Registration Service

The next step is to configure Device Registration Service on the ADFS1 server. For a video, see Active Directory Federation Services How-To Video Series: Enabling the Device Registration Service.

To configure Device Registration Service for Windows Server 2012 RTM

1.

   Important
   The following step applies to the Windows Server 2012 R2 RTM build.

   Open a Windows PowerShell command window and type:

   ```
   Initialize-ADDeviceRegistration
   ```

When you are prompted for a service account, type **contoso\fsgmsa$**.

Now run the Windows PowerShell cmdlet.

```
Enable-AdfsDeviceRegistration
```

2. On the ADFS1 server, in the **AD FS Management** console, navigate to **Authentication Policies**. Select **Edit Global Primary Authentication**. Select the check box next to **Enable Device Authentication**, and then click **OK**.

Add Host (A) and Alias (CNAME) Resource Records to DNS

On DC1, you must ensure that the following Domain Name System (DNS) records are created for Device Registration Service.

Entry	Type	Address
adfs1	Host (A)	IP address of the AD FS server
enterpriseregistration	Alias (CNAME)	adfs1.contoso.com

You can use the following procedure to add a host (A) resource record to corporate DNS name servers for the federation server and Device Registration Service.

Membership in the Administrators group or an equivalent is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships in the Local and Domain Default Groups (http://go.microsoft.com/fwlink/p/?LinkId=83477).

To add a host (A) and alias (CNAME) resource records to DNS for your federation server

1. On DC1, from Server Manager, on the **Tools** menu, click **DNS** to open the DNS snap-in.
2. In the console tree, expand DC1, expand **Forward Lookup Zones**, right-click **contoso.com**, and then click **New Host (A or AAAA)**.
3. In **Name,** type the name you want to use for your AD FS farm. For this walkthrough, type **adfs1**.
4. In **IP address**, type the IP address of the ADFS1 server. Click **Add Host**.
5. Right-click **contoso.com**, and then click **New Alias (CNAME)**.
6. In the **New Resource Record** dialog box, type **enterpriseregistration** in the **Alias name** box.
7. In the Fully Qualified Domain Name (FQDN) of the target host box, type **adfs1.contoso.com**, and then click **OK**.

Important

In a real-world deployment, if your company has multiple user principal name (UPN) suffixes, you must create multiple CNAME records, one for each of those UPN suffixes in DNS.

Step 3: Configure the web server (WebServ1) and a sample claims-based application

Set up a virtual machine (WebServ1) by installing the Windows Server 2012 R2 operating system and connect it to the domain **contoso.com**. After it is joined to the domain, you can proceed to install and configure the Web Server role.

To complete the walkthroughs that were referenced earlier in this topic, you must have a sample application that is secured by your federation server (ADFS1).

You can download Windows Identity Foundation SDK (http://www.microsoft.com/download/details.aspx?id=4451, which includes a sample claims-based application.

You must complete the following steps to set up a web server with this sample claims-based application.

Note
These steps have been tested on a web server that runs the Windows Server 2012 R2 operating system.

1. Install the Web Server role and Windows Identity Foundation
2. Install Windows Identity Foundation SDK
3. Configure the simple claims app in IIS
4. Create a relying party trust on your federation server

Install the Web Server role and Windows Identity Foundation

1.

   Note
   You must have access to the Windows Server 2012 R2 installation media.

   Log on to WebServ1 by using **administrator@contoso.com** and the password **pass@word1**.
2. From Server Manager, on the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, verify that the target computer is selected, and then click **Next**.
6. On the **Select server roles** page, select the check box next to **Web Server (IIS)**, click **Add Features**, and then click **Next**.
7. On the **Select features** page, select **Windows Identity Foundation 3.5**, and then click **Next**.

8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select role services** page, select and expand **Application Development**. Select **ASP.NET 3.5**, click **Add Features**, and then click **Next**.
10. On the **Confirm installation selections** page, click **Specify an alternate source path**. Enter the path to the Sxs directory that is located in the Windows Server 2012 R2 installation media. For example D:\Sources\Sxs. Click **OK**, and then click **Install**.

Install Windows Identity Foundation SDK

1. Run WindowsIdentityFoundation-SDK-3.5.msi to install Windows Identity Foundation SDK 3.5 (http://www.microsoft.com/en-us/download/details.aspx?id=4451). Choose all of the default options.

Configure the simple claims app in IIS

1. Install a valid SSL certificate in the computer certificate store. The certificate should contain the name of your web server, **webserv1.contoso.com**.
2. Copy the contents of C:\Program Files (x86)\Windows Identity Foundation SDK\v3.5\Samples\Quick Start\Web Application\PassiveRedirectBasedClaimsAwareWebApp to C:\Inetpub\Claimapp.
3. Edit the **Default.aspx.cs** file so that no claim filtering takes place. This step is performed to ensure that the sample application displays all the claims that are issued by the federation server. Do the following:
   a. Open **Default.aspx.cs** in a text editor.
   b. Search the file for the second instance of `ExpectedClaims`.
   c. Comment out the entire `IF` statement and its braces. Indicate comments by typing "//" (without the quotes) at the beginning of a line.
   d. Your `FOREACH` statement should now look like this code example.

```
Foreach (claim claim in claimsIdentity.Claims)
{
   //Before showing the claims validate that this is an
expected claim
   //If it is not in the expected claims list then don't show
it
   //if (ExpectedClaims.Contains( claim.ClaimType ) )
   // {
      writeClaim( claim, table );
   //}
}
```

   e. Save and close **Default.aspx.cs**.

f. Open **web.config** in a text editor.

g. Remove the entire `<microsoft.identityModel>` section. Remove everything starting from `including <microsoft.identityModel>` and up to and including `</microsoft.identityModel>`.

h. Save and close **web.config**.

4. **Configure IIS Manager**

   a. Open **Internet Information Services (IIS) Manager**.

   b. Go to **Application Pools**, right-click **DefaultAppPool** to select **Advanced Settings**. Set **Load User Profile** to **True**, and then click **OK**.

   c. Right-click **DefaultAppPool** to select **Basic Settings**. Change the **.NET CLR Version** to **.NET CLR Version v2.0.50727**.

   d. Right-click **Default Web Site** to select **Edit Bindings**.

   e. Add an **HTTPS** binding to port **443** with the SSL certificate that you have installed.

   f. Right-click **Default Web Site** to select **Add Application**.

   g. Set the alias to **claimapp** and the physical path to **c:\inetpub\claimapp**.

5. To configure **claimapp** to work with your federation server, do the following:

   a. Run FedUtil.exe, which is located in **C:\Program Files (x86)\Windows Identity Foundation SDK\v3.5**.

   b. Set the application configuration location to **C:\inetput\claimapp\web.config** and set the application URI to the URL for your site, **https://webserv1.contoso.com /claimapp/**. Click **Next**.

   c. Select **Use an existing STS** and browse to your AD FS server's metadata URL **https://adfs1.contoso.com/federationmetadata/2007-06/federationmetadata.xml**. Click **Next**.

   d. Select **Disable certificate chain validation**, and then click **Next**.

   e. Select **No encryption**, and then click **Next**. On the **Offered claims** page, click **Next**.

   f. Select the check box next to **Schedule a task to perform daily WS-Federation metadata updates**. Click **Finish**.

   g. Your sample application is now configured. If you test the application URL **https://webserv1.contoso.com/claimapp**, it should redirect you to your federation server. The federation server should display an error page because you have not yet configured the relying party trust. In other words, you have not secured this test application by AD FS.

You must now secure your sample application that runs on your web server with AD FS. You can do this by adding a relying party trust on your federation server (ADFS1). For a video, see Active Directory Federation Services How-To Video Series: Add a Relying Party Trust.

Create a relying party trust on your federation server

1. On you federation server (ADFS1), in the **AD FS Management console**, navigate to **Relying Party Trusts**, and then click **Add Relying Party Trust**.
2. On the **Select Data Source** page, select **Import data about the relying party**

published online or on a local network, enter the metadata URL for **claimapp**, and then click **Next**. Running FedUtil.exe created a metadata .xml file. It is located at **https://webserv1.contoso.com/claimapp/federationmetadata/2007-06/federationmetadata.xml**.

3. On the **Specify Display Name** page, specify the **display name** for your relying party trust, **claimapp**, and then click **Next**.
4. On the **Configure Multi-factor Authentication Now?** page, select **I do not want to specify multi-factor authentication setting for this relying party trust at this time**, and then click **Next**.
5. On the **Choose Issuance Authorization Rules** page, select **Permit all users to access this relying party**, and then click **Next**.
6. On the **Ready to Add Trust** page, click **Next**.
7. On the **Edit Claim Rules** dialog box, click **Add Rule**.
8. On the **Choose Rule Type** page, select **Send Claims Using a Custom Rule**, and then click **Next**.
9. On the **Configure Claim Rule** page, in the **Claim rule name** box, type **All Claims**. In the **Custom rule** box, type the following claim rule.

```
c:[ ]
=> issue(claim = c);
```

10. Click **Finish**, and then click **OK**.

Step 4: Configure the client computer (Client1)

Set up another virtual machine and install Windows 8.1. This virtual machine must be on the same virtual network as the other machines. This machine should NOT be joined to the Contoso domain.

The client MUST trust the SSL certificate that is used for the federation server (ADFS1), which you set up in Step 2: Configure the federation server (ADFS1) by using Device Registration Service . It must also be able to validate certificate revocation information for the certificate.

You also must set up and use a Microsoft account to log on to Client1.

See Also

Active Directory Federation Services How-To Video Series: Installing an AD FS Server Farm

Active Directory Federation Services How-To Video Series: Updating Certificates

Active Directory Federation Services How-To Video Series: Add a Relying Party Trust

Active Directory Federation Services How-To Video Series: Enabling the Device Registration Service

Active Directory Federation Services How-To Video Series: Installing the Web Application Proxy

Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications

The rapid increase in the number of consumer devices and ubiquitous information access is changing the way that people perceive their technology. The constant use of information technology throughout the day, along with easy access of information, is blurring traditional boundaries between work and home life. These shifting boundaries are accompanied by a belief that personal technology—selected and customized to fit users' personalities, activities, and schedules—should extend into the workplace. To accommodate the growing requirement of personal consumer devices to be connected to enterprise networks, Windows Server 2012 R2 introduces the following value propositions:

- Administrators can control who has access to company resources that are based on application, user, device, and location.
- Employees can access applications and data everywhere, on any device. Employees can use Single Sign-On in browser applications or enterprise applications.

Key concepts introduced in the solution

Workplace Join

By using Workplace Join, information workers can join their personal devices with their company's workplace computers to access company resources and services. When you join your personal device to your workplace, it becomes a known device and provides seamless second factor authentication and Single Sign-On to workplace resources and applications. When a device is joined by Workplace Join, attributes of the device can be retrieved from the directory to drive conditional access for the purpose of authorizing issuance of security tokens for applications. With Windows Server 2012 R2, Windows 8.1 and iOS devices can be joined by using Workplace Join.

Device Registration Service

Workplace Join is made possible by Device Registration Service (DRS) that is included with the Active Directory Federation Services role in Windows Server 2012 R2. When a device is joined by Workplace Join, DRS provisions a device object in Active Directory and sets a certificate on the consumer device that is used to represent the device identity. DRS is meant to face both internal and external resources. Companies that deploy both DRS and the Web Application Proxy can join devices that use Workplace Join from any Internet-connected location.

For more information about deploying Device Registration Service, see Step 5: Configure a federation server with Device Registration Service

Workplace Join as a seamless second factor authentication

Companies can manage the risk that is related to information access and drive governance and compliance while granting consumer devices access to corporate resources. Workplace Join on devices provides the following capabilities to administrators:

- Identifies known devices with device authentication. Administrators can use this information to drive conditional access and control access to resources.
- Provides a more seamless sign-in experience for users to access company resources from trusted devices.

Single Sign-On

Single Sign-On (SSO) in the context of this scenario is the functionality that reduces the number of password prompts that the end user has to enter to access company resources from known devices. This functionality implies that users are prompted only one time during the lifetime of SSO to access company applications and resource from this device. If a device uses Workplace Join, the user who is registered to use this device gets persistent SSO, by default for seven days. This user has a seamless sign-in experience in the same session or in new sessions.

Solution Overview

As part of this solution, you learn how to use Workplace Join on a Windows device, an iOS device, and experience Single Sign-On to a company resource.

This solution guides takes you through the following walkthrough steps:

1. Walkthrough Guide: Workplace Join with a Windows Device
2. Walkthrough Guide: Workplace Join with an iOS Device

See Also

Step 5: Configure a federation server with Device Registration Service

Walkthrough Guide: Workplace Join with an iOS Device

This topic demonstrates Workplace Join on an iOS device. You must complete the steps in the Set up the lab environment for AD FS in Windows Server 2012 R2 section before you can try out this walkthrough. You can use the device to access the same company web application that you accessed in Walkthrough Guide: Workplace Join with a Windows Device.

Join an iOS device with Workplace Join

Important

For Workplace Join to succeed, the iOS device must trust the Secure Socket Layer (SSL) certificate that was used to configure Active Directory Federation Services (AD FS) in Step 2: Configure the federation server (ADFS1) with Device Registration Service .

- If the AD FS SSL certificate was issued from a test certification authority (CA), you must install the certification authority certificate on your iOS device.
- If your certification authority certificate is published on a website, you can browse to the website from your iOS device and install the certificate.

In this demonstration, you join the device to the workplace.

To join an iOS device to a workplace

1. Open Apple Safari and navigate to the Device Registration Service (DRS) Over-the-Air Profile endpoint for iOS devices, **https://adf1s.contoso.com/enrollmentserver/otaprofile**.
2. Log on to the webpage by using a company domain account: **roberth@contoso.com** and password: **P@ssword**.
3. You are prompted to install a profile. On the **Install Profile** screen, click **Install**.
4. When prompted to confirm installation of the profile, click **Install Now**.
5. If your device requires a PIN to unlock the device, you are prompted to enter your PIN.
6. The profile installation is finished when you see the **Profile Installed** screen. Click **Done**.

   Return to Safari. A message informs you that you can close or leave Safari.

Tip

To view or remove the Workplace Join profile, browse to **Settings**, click **General**, and then click **Profiles** on your iOS device.

See Also

Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications

Set up the lab environment for AD FS in Windows Server 2012 R2

Walkthrough Guide: Workplace Join with a Windows Device

Walkthrough Guide: Workplace Join with a Windows Device

This topic demonstrates how to use Workplace Join to connect your Windows device with your workplace and how to access a web application by using Single Sign-On. You must complete the steps in the Set up the lab environment for AD FS in Windows Server 2012 R2 section before you can try out this walkthrough.

Access the web application before device registration

In this walkthrough, you access a company web application before you join your device to the workplace. The webpage displays the claims that were included in your security token. Notice that the list of claims does not include any information about your device. You might also observe that you do not have Single Sign-On.

To access the web application before you use Workplace Join on your device

1. Log on to Client1 with your Microsoft account.
2. Open Internet Explorer and browse to your generic claims app, **https://webserv1.contoso.com/claimapp**.
3. Log on to the webpage by using a company domain account: **roberth@contoso.com**, password: **P@ssword**.
4. The webpage lists all the claims in your security token. Only user claims are present in your security token.
5. Close Internet Explorer.
6. Open Internet Explorer and navigate to the same claims app, **https://webserv1.contoso.com/claimapp**.
7. Notice that you are prompted to enter your credentials again. You are not connected to the workplace from a device with Workplace Join and therefore do not have Single Sign-On.

Join your device with Workplace Join

Important

For Workplace Join to succeed, the client computer (Client1) must trust the SSL certificate that was used to configure Active Directory Federation Services (AD FS) in Step 2: Configure the Federation Server with Device Registration Service (ADFS1). It must also be able to validate revocation information for the certificate. If you have any issues with Workplace Join, you can view the event log on Client1.

To see the event log, open Event Viewer, expand **Applications and Services Logs**, expand **Microsoft**, expand **Windows**, and then click **Workplace Join**.

To join your device with Workplace Join

1. Log on to Client1 with your Microsoft account.
2. On the **Start** screen, open the **Charms** bar, and then select the **Settings** charm. Select **Change PC Settings**.
3. On the **PC Settings** page, select **Network**, and then click **Workplace**.
4. In the **Enter your UserID to get workplace access or turn on device management** box, type **roberth@contoso.com**, and then click **Join**.
5. When you are prompted for credentials, type **roberth@contoso.com**, and password: **P@ssword**. Click **OK**.
6. You should now see the message: "This device has joined your workplace network."

Access the web application after joining the workplace

In this part of the demonstration, you access a company web application from your device that is connected with Workplace Join. The webpage displays the claims that were included in your security token. Notice that the list of claims includes both device and user information. You might also observe that you now have Single Sign-On.

To access the web application after joining the workplace

1. Log on to **Client1** with your Microsoft account.
2. Open Internet Explorer and browse to your generic claims app, **https://webserv1.contoso.com/claimapp**.
3. Log on to the webpage by using a company domain account: **roberth@contoso.com**, password: **P@ssword**.
4. The webpage lists claims in your security token. Your token contains both user and device claims.
5. Close Internet Explorer.
6. Open Internet Explorer and navigate to the same claims app, **https://webserv1.contoso.com/claimapp**.
7. Notice that you are **not** prompted to enter your credentials again. You are connected from a device with Workplace Join and therefore have Single Sign-On.

See Also

Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications

Set up the lab environment for AD FS in Windows Server 2012 R2

Walkthrough Guide: Workplace Join with an iOS Device

Overview: Connect to Applications and Services from Anywhere with Web Application Proxy

This scenario describes how you can use Web Application Proxy – a new Remote Access role service in Windows Server® 2012 R2 – to provide access to a sample web application using claims-based authentication and a sample website using Integrated Windows authentication, both websites use Active Directory Federation Services (AD FS) preauthentication. This scenario also uses the following AD FS features:

- **Workplace Join**—Joining devices to the workplace connects these devices with Active Directory in your workplace. When you join personal devices to your workplace they become known devices and will provide seamless second factor authentication and single-sign-on to workplace resources and applications.

  This scenario configures the device registration service (DRS) to enable you to join the client device to the workplace.
- **Multifactor authentication**—This enables you to require users to provide more than one form of authentication when connecting to published applications and services. For example, using one-time passwords or smart cards. You can configure Web Application Proxy and AD FS to use multifactor authentication for all authentication requests, or per-application. In addition, configuring AD FS to allow access to only registered devices creates two-factor seamless authentication because the user must provide credentials, and the device must be registered.

  This scenario uses certificate authentication to provide the additional factor when authenticating.
- **Multifactor access control**—Access control in AD FS is implemented with authorization claim rules that are used to issue a permit or deny claims that will determine whether a user or a group of users will be allowed to access AD FS-secured resources or not. In AD FS in Windows Server 2012 R2, access control is enhanced with multiple factors, including user, device, location, and authentication data. This is made possible by a greater variety of claim types available for the authorization claim rules.

Web Application Proxy also provides built-in AD FS proxy capabilities. The following diagram shows the topology used in this scenario for Web Application Proxy to publish Microsoft applications and other line-of-business (LOB) applications.

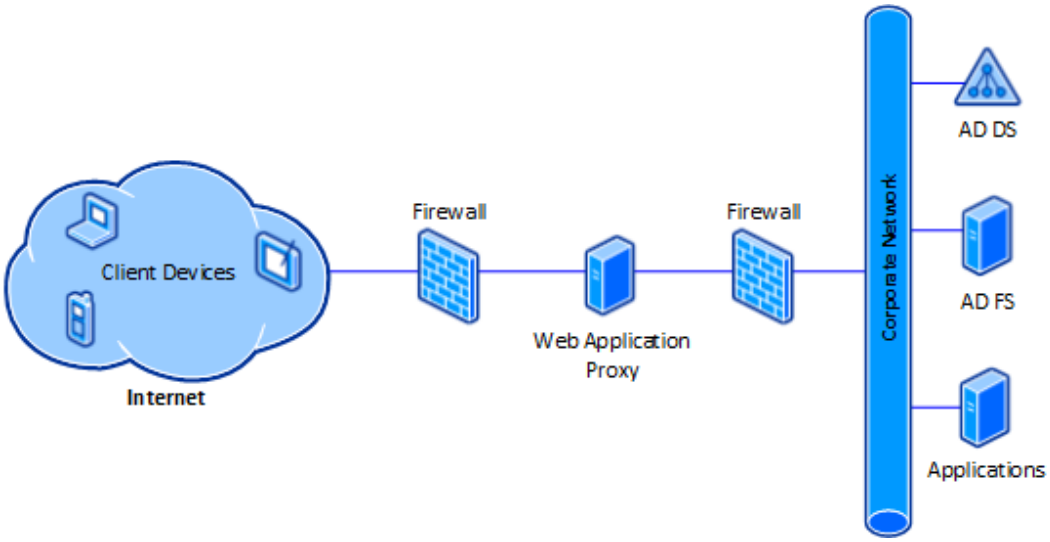

- Providing Access to Applications and Services
- Using Active Directory Federation Services
- AD FS Proxy
- Roles and Features Included in this Scenario
- Scenario Steps

Providing Access to Applications and Services

This scenario demonstrates how to plan and deploy Web Application Proxy in your organization to provide end users located outside of an organization access to applications and services running on servers inside the organization. Web Application Proxy publishing enables end users to access their organization's applications from their own devices, so that users are not limited to corporate laptops to do their work, they can use their home computer, their tablet, or their smartphone. Web Application Proxy can be used on clients with a standard browser, an Office client or a rich client using OAuth (for example Windows Store apps). Web Application Proxy serves as a reverse proxy for any application that is published through it and as such, the end user experience is the same as if the end user's device connects directly to the application.

Using Active Directory Federation Services

This scenario describes the additions and changes that you must make to your AD FS servers to provide the following functionality:

- Application publishing—For all applications and services that you want to publish through Web Application Proxy, you must configure a relying party on the AD FS server.
- Authentication—No specific configuration is required to provide authentication for published applications. However, to use workplace join, MFA, or multifactor access control in your deployment, you must perform additional configuration as described in the following guides:

- Workplace Join—Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications
- MFA—Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications
- Multifactor access control—Overview: Manage Risk with Multi-Factor Access Control.

AD FS Proxy

This scenario does not describe using Web Application Proxy as a proxy for AD FS. However, this functionality is enabled by default when you install the Web Application Proxy role service. Any AD FS endpoint that is enabled for proxy publishing is automatically published by Web Application Proxy after completing the Web Application Proxy Configuration Wizard.

Roles and Features Included in this Scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Federation Services Overview	AD FS is required to provide authentication and authorization services to Web Application Proxy and to store the Web Application Proxy configuration.
Remote Access (DirectAccess, Routing and Remote Access) Overview	Remote Access is the role containing the Web Application Proxy role service.
Active Directory Domain Services Overview	Active Directory® Domain Services is required as a prerequisite before you can deploy AD FS.
Web Server (IIS) Overview	The Web Server (IIS) role is used in this scenario to host a sample application that can be published by Web Application Proxy.

Scenario Steps

This scenario contains the following steps:

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy—This walkthrough contains the following steps:
  - Step 1: Attempt to access the web application from an Internet client
  - Step 2: Configure the Web Application Proxy server and publish the application
  - Step 3: Configure and test accessing a site using Integrated Windows authentication

- Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control

See Also

Web Application Proxy Overview

Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy

This walkthrough provides instructions for connecting to applications and services from anywhere with Web Application Proxy. For more information about connecting to applications and services from anywhere with Web Application Proxy and Windows Server® 2012 R2, see Overview: Connect to Applications and Services from Anywhere with Web Application Proxy and the video for Installing the Web Application Proxy.

Before you perform this walkthrough, make sure you have completed the lab setup guide, see Set up the lab environment for AD FS in Windows Server 2012 R2. Also make sure you have completed the following planning steps that are required in order to deploy Web Application Proxy and publish applications through it using Active Directory Federation Services (AD FS) for authentication and authorization, including workplace join, multifactor authentication (MFA), and multifactor access control.

- Plan the Web Application Proxy infrastructure—Before you deploy Web Application Proxy, you must plan the infrastructure. For this walkthrough, this includes:
  - The Web Application Proxy server should be located between the corporate network and the Internet; in this walkthrough, this is achieved by configuring two network adapters on the Web Application Proxy server.
  - For clients to reach a published web application they must be able to resolve the external address of the application that is published by Web Application Proxy. In this walkthrough, this is achieved through the hosts file on the client.
- Plan the Web Application Proxy server
  - It is recommended to join the Web Application Proxy server to the AD DS domain. Joining the Web Application Proxy server to the domain allows you to expand this walkthrough to demonstrate publishing applications that use Integrated Windows authentication.
  - Make sure that the server time of the Web Application Proxy and AD FS servers are synchronized.
  - To install the Web Application Proxy role service, you must have local administrator permissions on the server.

- To configure the Web Application Proxy server to connect to AD FS, you are required to enter the credentials of the local administrator on the AD FS server.
- Make sure that there is a certificate in the Personal certificates store on the Web Application Proxy server that covers the federation service name. This certificate is required for AD FS proxy functionality.

  To locate the federation service name value, open the AD FS Management console, click **Service**, in the **Actions** pane, click **Edit Federation Service Properties**, and then find the value in the **Federation Service name** text box.
- Make sure that there is a certificate in the Personal certificates store on the Web Application Proxy server that covers the external address of the sample web application; that is, the address that you will configure when you publish the application through Web Application Proxy. In this walkthrough, you can use a self-signed certificate, or a certificate from an Enterprise certification authority (CA).

Note
In a production environment, it is recommended to use certificates from a public CA on your Web Application Proxy server.

Task	Description
Step 1: Attempt to access the web application from an Internet client	Configure an edge server and demonstrate that before configuring Web Application Proxy and publishing the application, the web application server is protected from the Internet, and clients are unable to access the application.
Step 2: Configure the Web Application Proxy server and publish the application	Configure the edge server as a Web Application Proxy server, publish the web application, and demonstrate that with the default AD FS authentication scheme, a client device can access the published application from the Internet.
Step 3: Configure and test accessing a website using Integrated Windows authentication	Configure an additional website on WebServ1 to use Integrated Windows authentication, configure Kerberos constrained delegation in the lab, publish the application through Web Application Proxy, and demonstrate accessing this application from the Internet.
Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control	Configure Workplace Join, multifactor authentication (MFA), and multifactor access control, and demonstrate how these additional authentication factors change the user experience when client devices attempt to

Task	Description
	access the published web application.

Step 1: Attempt to access the web application from an Internet client

In this step, you will configure an edge server and attempt to access the sample application. This step does not install the Web Application Proxy role service on the server. After configuring the edge server you will modify the hosts file on the client computer and attempt to connect to the sample web application through the edge server.

To install the edge server

1. Create a Windows Server 2012 R2 server with two network adapters, one connected to the corporate network, one connected to the simulated Internet.
2. Configure the static IP addresses on the server. For example, the network adapter on the corporate network should have an IP address of 10.0.0.2, the network adapter on the simulated Internet should have an IP address of 131.107.0.2.
3. Join the server to the AD DS domain.

To access the application

1. In your lab environment, connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
2. On the client computer, open Notepad as the administrator.
3. In the Notepad window, open the hosts file located at %windir%\System32\Drivers\etc.
4. Add two new lines to the file; for example:

   ```
   131.107.0.2       webserv1.contoso.com
   131.107.0.2       adfs1.contoso.com
   131.107.0.2       enterpriseregistration.contoso.com
   ```

   - The first line contains the IP address of the Internet adapter of the edge server and the external address of the sample application that you will configure when you publish the application through Web Application Proxy.
   - The second line contains the same IP address and the address of the federation server.
   - The third line contains the same IP address and the address required for Workplace Join.
5. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.

   The edge server is not providing proxy or publishing functionality; therefore, you will not

be able to access the application.

Step 2: Configure the Web Application Proxy server and publish the application

In this step you will:

- Configure the Web Application Proxy server.
- Connect the Web Application Proxy server to the AD FS server using the Web Application Proxy Configuration Wizard.
- Publish the sample application.
- Attempt to connect to the sample application using the default AD FS authentication scheme.

Before you can publish the application, make sure you have done the following:

- Create a relying party trust for the sample application in the AD FS Management console. If you did not do this when setting up your environment, see Create a relying party trust on your federation server.
- Verify that a certificate on the Web Application Proxy server is suitable for the sample application you want to publish.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To install the Web Application Proxy role service

1. On the edge server, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. In the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. In the **Add Roles and Features Wizard**, click **Next** three times to get to the server role selection screen.
4. On the **Select server roles** dialog, select **Remote Access**, and then click **Next**.
5. Click **Next** twice.
6. On the **Select role services** dialog, select **Web Application Proxy**, click **Add Features**, and then click **Next**.
7. On the **Confirm installation selections** dialog, click **Install**.
8. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature Web-Application-Proxy -IncludeManagementTools
```

To configure Web Application Proxy

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
4. In the **Web Application Proxy Configuration Wizard**, on the **Welcome** dialog, click **Next**.
5. On the **Federation Server** dialog, do the following, and then click **Next**:
   - In the **Federation service name** box, enter the fully qualified domain name (FQDN) of the federation service; for example, adfs1.contoso.com.

     To locate this value, open the AD FS Management console, click **Service**, in the **Actions** pane, click **Edit Federation Service Properties**, and then find the value in the **Federation Service name** text box.
   - In the **User name** and **Password** boxes, enter the credentials of a local administrator account on the AD FS servers.
6. On the **AD FS Proxy Certificate** dialog, in the list of certificates currently installed on the Web Application Proxy server, select a certificate to be used by Web Application Proxy for AD FS proxy functionality, and then click **Next**.
7. On the **Confirmation** dialog, review the settings. If required, you can copy the PowerShell cmdlet to automate additional installations. Click **Configure**.
8. On the **Results** dialog, verify that the configuration was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following command will prompt you to enter credentials of a local administrator account on the AD FS servers.

```
Install-WebApplicationProxy -CertificateThumprint
'1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b' -FederationServiceName adfs1.contoso.com
```

To publish the sample claims-based application

1. On the Web Application Proxy server, in the Remote Access Management console, in the **Navigation** pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Active Directory Federation Services (AD FS)**,

and then click **Next**.

4. On the **Relying Party** page, in the list of relying parties select the relying party for the sample application, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://webserv1.contoso.com/claimapp/.
   - In the **External certificate** list, select a certificate whose subject name covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that application published successfully, and then click **Close**.

   Publishing the application may take a couple of minutes.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WebApplicationProxyApplication

    -BackendServerURL 'https://webserv1.contoso.com/claimapp/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://webserv1.contoso.com/claimapp/'

    -Name 'Claims App'

    -ExternalPreAuthentication ADFS

    -ADFSRelyingPartyName 'App_Relying_Party'
```

To access the application

1. On the client computer, open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
2. On the sign in page, enter the credentials of the test user that you created when you set up the environment.
3. The application appears in the web browser.

Step 3: Configure and test accessing a website using Integrated Windows authentication

In this step you will configure a website on WebServ1 to use Integrated Windows authentication. You can publish this website as an additional application through Web Application Proxy to demonstrate how Web Application Proxy uses Kerberos constrained delegation to authenticate users to non-claims-based applications.

Important
You must join the Web Application Proxy server to the Active Directory® Domain Services domain before you can publish applications that use Integrated Windows authentication.

Install Windows Authentication on WebServ1

When you install the Web Server (IIS) role, by default, it does not install the option to use Integrated Windows authentication (referred to by IIS as Windows Authentication). This procedure installs Windows Authentication.

To install Windows Authentication on WebServ1

1. On the edge server, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. In the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. In the **Add Roles and Features Wizard**, click **Next** three times to get to the server role selection screen.
4. On the **Select server roles** dialog, expand **Web Server (IIS)**, expand **Web Server**, expand **Security**, select **Windows Authentication**, and then click **Next**.
5. On the **Select features** dialog, click **Next**.
6. On the **Confirm installation selections** dialog, click **Install**.
7. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Create a new website using IIS

The WebServ1 server already hosts a website that uses claims-based authentication. This procedure creates a new website that uses Integrated Windows authentication.

To create a new website

1. Open an Explorer window and go to C:\inetpub\.
2. Create a new folder called NonClaims.
3. Open C:\inetpub\wwwroot\ and copy the contents to the new folder

C:\inetpub\NonClaims\.

Important
Do not copy the web.config file.

4. On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **inetmgr.exe**, and then press ENTER.
5. In the IIS console, in the **Connections** pane, expand **WebServ1**, expand **Sites**, right-click **Default Web Site**, and then click **Add Virtual Directory**.
6. On the **Add Virtual Directory** dialog box, in the **Alias** box, type **NonClaims**.
7. Next to the **Physical Path** box, type **C:\inetpub\NonClaims**, and then click **OK**.
8. In the **Connections** pane, click the **NonClaims** virtual directory.
9. In the middle pane, in the **IIS** section, double-click **Authentication**.
10. In the **Authentication** pane, right-click **Anonymous Authentication**, and then click **Disable**. Right-click **Windows Authentication**, and then click **Enable**.

Create a non-claims-aware relying party trust

To publish the application through Web Application Proxy, you must first create a non-claims-aware relying party trust on the AD FS server.

To create a non-claims-aware relying party trust

1. On the AD FS server, in the AD FS Management console, in the left pane, click **Trust Relationships**.
2. In the **Actions** pane, click **Add Non-Claims-Aware Relying Party Trust**.
3. In the **Add Non-Claims-Aware Relying Party Trust Wizard**, on the **Welcome** page, click **Start**.
4. On the **Specify Display Name** page, in **Display name**, enter a name for this trust; for example, **Non-Claims Application**, and then click **Next**.
5. On the **Configure Identifiers** page, in **Non-claims-aware relying party trust identifier**, enter an arbitrary URL as an identifier for this trust, click **Add**, and then click **Next**.
6. Click **Next** twice.
7. On the **Finish** page, click **Close**.

   The **Edit Claim Rules for <relying party>** dialog box opens.
8. On the **Edit Claim Rules for <relying party>**, click **Add Rule**.
9. In the **Add Issuance Authorization Claim Rule Wizard**, on the **Select Rule Template** page, in the **Claim rule template** list, select **Permit All Users**, and then click **Next**.
10. On the **Configure Rule** page, click **Finish**, and then on the **Edit Claim Rules for <relying party>** dialog box, click **OK**.

Configure Kerberos constrained delegation

To allow users to access applications that use Integrated Windows authentication, the Web Application Proxy server must be able to provide impersonation for users to the published application.

Tip

This procedure assumes that the Web Application Proxy server is named EDGE1.

To configure Kerberos constrained delegation

1. On the domain controller, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. Click **Tools**, and then click **ADSI Edit**.
3. On the **Action** menu, click **Connect To**, and then on the **Connection Settings** dialog box, accept the default settings to connect to the default naming context, and then click **OK**.
4. In the left pane, expand **Default naming context**, expand **DC=contoso,DC=com**, expand **CN=Computers**, right-click **CN=EDGE1**, and then click **Properties**.
5. On the **CN=EDGE1 Properties** dialog box, on the **Attribute Editor** tab, in the **Attributes** list, select **servicePrincipalName**, and then click **Edit**.
6. On the **Multi-valued String Editor** dialog box, in **Value to add**, enter **HTTP/EDGE1.contoso.com** and click **Add**. Then enter **HTTP/EDGE1** and click **Add**.

   The **Values** list now contains two new entries; for example, **HTTP/EDGE1.contoso.com** and **HTTP/EDGE1**.
7. On the **Multi-valued String Editor** dialog box, click **OK**.
8. On the **CN=EDGE1 Properties** dialog box, click **OK**.
9. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
10. In the navigation pane, under **contoso.com**, click **Computers**. In the details pane, right-click the edge server, and then click **Properties**.
11. On the **EDGE1 Properties** dialog box, on the **Delegation** tab, click **Trust this computer for delegation to specified services only**, and then click **Use any authentication protocol**.
12. Click **Add**, and on the **Add Services** dialog box, click **Users or Computers**.
13. On the **Select Users or Computers** dialog box, in **Enter the object names to select**, enter the name of the web server; for example, WebServ1, and then click **OK**.
14. On the **Add Services** dialog box, in the **Available services** list, select the **http** service type, and then click **OK**.
15. On the **EDGE1 Properties** dialog box, click **OK**.

Test accessing the application internally

To make sure that the application is working correctly, you should test accessing the application from the Web Application Proxy server before continuing.

To access the application internally

1. On the Web Application Proxy server, make sure that IE Enhanced Security Configuration is turned off.
2. Open an Internet explorer window and go to the non-claims-based web application; for example, https://webserv1.contoso.com/nonclaims/.
3. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   The default IIS website appears.

Publish the application

You must publish the application through Web Application Proxy to make it available to external users.

To publish the non-claims-based application

1. On the Edge server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER.
2. In the navigation pane, under **Configuration**, click **Web Application Proxy**.
3. In the **Tasks** pane, click **Publish**.
4. In the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
5. On the **Preauthentication** page, make sure that Active Directory Federation Services (AD FS) is selected, and then click **Next**.
6. On the **Relying Party** page, select the non-claims application relying party, and then click **Next**.
7. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application to identify it in the list of published applications; for example, **Non-Claims Application**.
   - In the **External URL** box, enter the external URL for this application; for example, **https://WebServ1.contoso.com/NonClaims/**.
   - In the **External certificate** list, select a certificate whose subject covers the external address.
   - In the **Backend server URL** box, the value is entered automatically and is the same as the external URL.
   - In the **Backend server SPN** box, enter the service principal name for this application; for example, **HTTP/WEBSERV1.contoso.com**.
8. On the **Confirmation** page, click **Publish**.

9. On the **Results** page, click **Close**.

Test accessing the application

After publishing the application through Web Application Proxy, you can test accessing the application from the simulated Internet.

To test accessing the application

1. Connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
2. Open an Internet explorer window and go to the non-claims-based web application; for example, https://webserv1.contoso.com/nonclaims/.
3. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   The default IIS website appears.

Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control

In this step you will join the client device to the workplace, and configure AD FS to enforce MFA, and multifactor access control when publishing the sample application through Web Application Proxy, and demonstrate that MFA and multifactor access control are working.

Important

For Workplace Join to succeed, the client computer (Client1) must trust the SSL certificate that was used to configure AD FS in Step 2: Configure the Federation Server with Device Registration Service (ADFS1). It must also be able to validate revocation information for the certificate. If you have any issues with Workplace Join, you can view the event log on Client1.

To see the event log, open Event Viewer, expand **Applications and Services Logs**, expand **Microsoft** , expand **Windows** and click **Workplace Join**.

To Workplace Join your device and access the application

1. On the client computer, on the **Start** screen, open the **Charms** bar and then select the **Settings** charm. Click **Change PC Settings**.
2. On the **PC Settings** page, click **Network**, and then click **Workplace**.
3. In the **Enter your UserID to get workplace access or turn on device management** box, type **RobertH@contoso.com** and then click **Join**.
4. When prompted for credentials, type **roberth@contoso.com**, and enter the password. Click **OK**.

5. You should now see the message: This device has joined your workplace network.
6. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
7. On the sign in page, enter the credentials of the test user (RobertH) that you created when you set up the environment.
8. The application appears in the web browser showing the claims in your security token. Notice that the token contains both user and device claims.

To configure MFA and attempt to access the application

1. Configure MFA as described in Step 3: Configure MFA on your federation server.
2. In your lab environment, connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
3. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
4. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   At this point, because of the MFA policy that you configured, the user will be prompted to undergo additional authentication.

   This demonstrates that the additional authentication factor that you configured is now required to access the application.

Tip

In this walkthrough, it is necessary to disable MFA before configuring multifactor access control.

To configure multifactor access control and attempt to access the application

1. Configure multifactor access control as described in Step 3: Configure multi-factor access control policy based on user data.
2. Disable MFA as follows:
   a. On the AD FS server, in the AD FS Management console, navigate to **Authentication Policies/Per Relying Party Trust**, and select the relying party trust that represents your sample application.
   b. In the **Actions** pane, click **Edit Custom Multi-factor Authentication**.
   c. In the **Edit Relying Party Trust for <relying_party_trust_name>** dialog box, in the **Users/Groups** list, select the test group that you previously added, click **Remove**, and then click **OK**.
3. On the client computer, open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
4. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   At this point, because of the access control policy that you set up in the previous steps,

an ‘access denied’ message is displayed for this user, demonstrating that the access control rule is in effect.

See also

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Overview: Connect to Applications and Services from Anywhere with Web Application Proxy

Overview: Manage Risk with Multi-Factor Access Control

- Key concepts - Multi-factor access control in AD FS
- Managing Risk with Multi-factor Access Control

Key concepts - Multi-factor access control in AD FS

The overall function of AD FS is to issue an access token that contains a set of claims. The decision regarding what claims AD FS accepts and then issues is governed by claim rules.

Access control in AD FS is implemented with issuance authorization claim rules that are used to issue a permit or deny claims that will determine whether a user or a group of users will be allowed to access AD FS-secured resources or not. Authorization rules can only be set on relying party trusts.

Rule option	Rule logic
Permit all users	If incoming claim type equals *any claim type* and value equals *any value*, then issue claim with value equals *Permit*
Permit access to users with this incoming claim	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then issue claim with value equals *Permit*
Deny access to users with this incoming claim	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then issue claim with value equals *Deny*

For more information about these rule options and logic, see When to Use an Authorization Claim Rule.

In AD FS in Windows Server® 2012 R2, access control is enhanced with multiple factors, including user, device, location, and authentication data. This is made possible by a greater variety of claim types available for the authorization claim rules. In other words, in AD FS in Windows Server® 2012 R2, you can enforce multi-factor access control based on user identity or group membership, network location, device (whether it is workplace joined, for more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications), and the authentication state (whether multifactor authentication (MFA) was performed).

Multifactor access control in AD FS in Windows Server® 2012 R2, offers the following benefits:

- Flexible and expressive per-application authorization policies, whereby you can permit or deny access based on user, device, network location, and authentication state
- Creating issuance authorization rules for relying party applications
- Rich UI experience for the common multifactor access control scenarios
- Rich claims language & Windows PowerShell support for advanced multifactor access control scenarios
- Custom (per relying party application) 'Access Denied' messages. For more information, see Customizing the AD FS Sign-in Pages. By being able to customize these messages, you can explain why a user is being denied access and also facilitate self-service remediation where it is possible, for example, prompt users to workplace join their devices. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.

The following table includes all the claim types available in AD FS in Windows Server® 2012 R2 to be used for implementing multi-factor access control.

Claim type	Description
Email Address	The email address of the user.
Given Name	The given name of the user.
Name	The unique name of the user,
UPN	The user principal name (UPN) of the user.
Common Name	The common name of the user.
AD FS 1 x E-mail Address	The email address of the user when interoperating with AD FS 1.1 or AD FS 1.0.
Group	A group that the user is a member of.
AD FS 1 x UPN	The UPN of the user when interoperating with AD FS 1.1 or AD FS 1.0.
Role	A role that the user has.
Surname	The surname of the user.

Claim type	Description
PPID	The private identifier of the user.
Name ID	The SAML name identifier of the user.
Authentication time stamp	Used to display the time and date that the user was authenticated.
Authentication method	The method used to authenticate the user.
Deny only group SID	The deny-only group SID of the user.
Deny only primary SID	The deny-only primary SID of the user.
Deny only primary group SID	The deny-only primary group SID of the user.
Group SID	The group SID of the user.
Primary group SID	The primary group SID of the user.
Primary SID	The primary SID of the user.
Windows account name	The domain account name of the user in the form of domain\user.
Is Registered User	User is registered to use this device.
Device Identifier	Identifier of the device.
Device Registration Identifier	Identifier for Device Registration.
Device Registration Display Name	Display name of Device Registration.
Device OS Type	Operating system type of the device.
Device OS Version	Operating system version of the device.
Is Managed Device	Device is managed by a management service.
Forwarded Client IP	IP address of the user.
Client Application	Type of the client application.
Client User Agent	Device type the client is using to access the application.
Client IP	IP address of the client.
Endpoint Path	Absolute Endpoint path which can be used to determine active versus passive clients.
Proxy	DNS name of the federation server proxy that passed the request.
Application Identifier	Identifier for the relying party.

Claim type	Description
Application policies	Application policies of the certificate.
Authority Key Identifier	The authority key identifier extension of the certificate that signed an issued certificate.
Basic Constraint	One of the basic constraints of the certificate.
Enhanced Key Usage	Describes one of the enhanced key usages of the certificate.
Issuer	The name of the certification authority that issued the X.509 certificate.
Issuer Name	The distinguished name of the certificate issuer.
Key Usage	One of the key usages of the certificate.
Not After	Date in local time after which a certificate is no longer valid.
Not Before	The date in local time on which a certificate becomes valid.
Certificate Policies	The policies under which the certificate has been issued.
Public Key	Public key of the certificate.
Certificate Raw Data	The raw data of the certificate.
Subject Alternative Name	One of the alternative names of the certificate.
Serial Number	The serial number of the certificate.
Signature Algorithm	The algorithm used to create the signature of a certificate.
Subject	The subject from the certificate.
Subject Key Identifier	The subject key identifier of the certificate.
Subject Name	The subject distinguished name from a certificate.
V2 Template Name	The name of the version 2 certificate template used wen issuing or renewing a certificate. This is a Microsoft-specific value.
V1 Template Name	The name of the version 1 certificate template used when issuing or renewing a certificate.

Claim type	Description
	This is a Microsoft-specific value.
Thumbprint	Thumbprint of the certificate.
X 509 Version	The X.509 format version of the certificate.
Inside Corporate Network	Used to indicate if a request originated from inside the corporate network.
Password Expiration Time	Used to display the time when the password expires.
Password Expiration Days	Used to display the number of days to password expiry.
Update Password URL	Used to display the web address of update password service.
Authentication Methods References	Used to indicate al authentication methods used to authenticate the user.

Managing Risk with Multi-factor Access Control

Using the available settings, there are many ways in which you can manage risk by implementing multifactor access control.

Common Scenarios

For example, imagine a simple scenario of implementing multi-factor access control based on the user's group membership data for a particular application (relying party trust). In other words, you can set up an issuance authorization rule on your federation server to permit users that belong to a certain group in your AD domain access to a particular application that is secured by AD FS. The detailed step by step instructions (using the UI and Windows PowerShell) for implementing this scenario are covered in Walkthrough Guide: Manage Risk with Multi-Factor Access Control. In order to complete the steps in this walkthrough, you must set up a lab environment and follow the steps in Set up the lab environment for AD FS in Windows Server 2012 R2.

Advanced Scenarios

Other examples of implementing multifactor access control in AD FS in Windows Server® 2012 R2 include the following:

- Permit access to an application secured by AD FS only if this user's identity was validated with MFA

  You can use the following code:

```
@RuleTemplate = "Authorization"
```

```
@RuleName = "PermitAccessWithMFA"
c:[Type ==
"http://schemas.microsoft.com/claims/authnmethodsreferences",
Value =~
"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$"] =>
issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "PermitUsersWithClaim");
```

- Permit access to an application secured by AD FS only if the access request is coming from a workplace joined device that is registered to the user

  You can use the following code:

```
@RuleTemplate = "Authorization"
@RuleName = "PermitAccessFromRegisteredWorkplaceJoinedDevice"
c:[Type ==
"http://schemas.microsoft.com/2012/01/devicecontext/claims/isreg
istereduser", Value =~ "^(?i)true$"] => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "PermitUsersWithClaim");
```

- Permit access to an application secured by AD FS only if the access request is coming from a workplace joined device that is registered to a user whose identity has been validated with MFA

  You can use the following code

```
@RuleTemplate = "Authorization"
@RuleName = "RequireMFAOnRegisteredWorkplaceJoinedDevice"
c1:[Type ==
"http://schemas.microsoft.com/claims/authnmethodsreferences",
Value =~
"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$"] &&
c2:[Type ==
"http://schemas.microsoft.com/2012/01/devicecontext/claims/isreg
istereduser", Value =~ "^(?i)true$"] => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "PermitUsersWithClaim");
```

- Permit extranet access to an application secured by AD FS only if the access request is coming from a user whose identity has been validated with MFA.

You can use the following code:

```
@RuleTemplate = "Authorization"
@RuleName = "RequireMFAForExtranetAccess"
c1:[Type ==
"http://schemas.microsoft.com/claims/authnmethodsreferences",
Value =~
"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$"] &&
c2:[Type ==
"http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork"
, Value =~ "^(?i)false$"] => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "PermitUsersWithClaim");
```

See Also

Walkthrough Guide: Manage Risk with Multi-Factor Access Control

Set up the lab environment for AD FS in Windows Server 2012 R2

Walkthrough Guide: Manage Risk with Multi-Factor Access Control

About This Guide

This walkthrough provides instructions for managing risk with one of the factors (user data) available through the multi-factor access control mechanism in Active Directory Federation Services (AD FS) in Windows Server® 2012 R2. For more information about multi-factor access control and authorization mechanisms in AD FS in Windows Server® 2012 R2, see Overview: Manage Risk with Multi-Factor Access Control.

This walkthrough consists of the following sections:

- Step 1: Setting up the lab environment
- Step 2: Verify the default AD FS access control mechanism
- Step 3: Configure multi-factor access control policy based on user data
- Step 4: Verify multi-factor access control mechanism

Step 1: Setting up the lab environment

In order to complete this walkthrough, you need an environment that consists of the following components:

- An Active Directory domain with a test user and group accounts, running on Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 with its schema upgraded to Windows Server® 2012 R2 or an Active Directory domain running on Windows Server® 2012 R2
- A federation server running on Windows Server® 2012 R2
- A web server that hosts your sample application
- A client computer from which you can access the sample application

Warning

It is highly recommended (both in production or test environments) that you do not use the same computer to be your federation server and your web server.

In this environment, the federation server issues the claims that are required so that users can access the sample application. The Web server hosts a sample application that will trust the users who present the claims that the federation server issues.

For instructions on how to set up this environment, see Set up the lab environment for AD FS in Windows Server 2012 R2.

Step 2: Verify the default AD FS access control mechanism

In this step you will verify the default AD FS access control mechanism, where the user is redirected to the AD FS sign-in page, provides valid credentials, and is granted access to the application. You can use the **Robert Hatley** AD account and the **claimapp** sample application that you configured in Set up the lab environment for AD FS in Windows Server 2012 R2.

To verify the default AD FS access control mechanism

1. On your client computer, open a browser window, and navigate to your sample application: **https://webserv1.contoso.com/claimapp**.

   This action automatically redirects the request to the federation server and you are prompted to sign in with a username and password.

2. Type in the credentials of the **Robert Hatley** AD account that you created in Set up the lab environment for AD FS in Windows Server 2012 R2.

   You will be granted access to the application.

Step 3: Configure multi-factor access control policy based on user data

In this step you will set up an access control policy based on the user group membership data. In other words, you will configure an **Issuance Authorization Rule** on your federation server for a relying party trust that represents your sample application - **claimapp**. By this rule's logic, **Robert Hatley** AD user will be issued claims that are required to access this application because he belongs to a **Finance** group. You have added the **Robert Hatley** account to the **Finance** group in Set up the lab environment for AD FS in Windows Server 2012 R2.

You can complete this task using either AD FS Management Console or via Windows PowerShell.

To configure multi-factor access control policy based on user data via the AD FS Management Console

1. In the AD FS Management Console, navigate to **Trust Relationships**, and then **Relying Party Trusts**.
2. Select the relying party trust that represents your sample application (**claimapp**), and then either in the **Actions** pane or by right-clicking this relying party trust, select **Edit Claim Rules**.
3. In the **Edit Claim Rules for claimapp** window, select **Issuance Authorization Rules** tab and click **Add Rule**.
4. In the **Add Issuance Authorization Claim Rule Wizard**, on the **Select Rule Template page**, select **Permit or Deny Users Based on an Incoming Claim** claim rule template and then click **Next**.
5. On the **Configure Rule** page, do all of the following and then click **Finish**:
   a. Enter a name for the claim rule, for example **TestRule**.
   b. Select **Group SID** as **Incoming claim type**.
   c. Click **Browse**, type in **Finance** for the name of your AD test group, and resolve it for the **Incoming claim value** field.
   d. Select the **Deny access to users with this incoming claim** option.
6. In the **Edit Claim Rules for claimapp** window, make sure to delete the **Permit Access to All Users** rule that was created by default when you created this relying party trust.

To configure multi-factor access control policy based on user data via Windows PowerShell

1. On your federation server, open the Windows PowerShell command window and run the following command:

```
$rp = Get-AdfsRelyingPartyTrust -Name claimapp
```

2. In the same Windows PowerShell command window, run the following command:

```
$GroupAuthzRule = '@RuleTemplate = "Authorization" @RuleName
```

```
= "Foo" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/grou
psid", Value =~ "^(?i)<group_SID>$"] =>issue(Type =
"http://schemas.microsoft.com/authorization/claims/deny",
Value = "DenyUsersWithClaim");'
Set-AdfsRelyingPartyTrust -TargetRelyingParty $rp -
IssuanceAuthorizationRules $GroupAuthzRule
```

Note
Make sure to replace <group_SID> with the value of the SID of your AD **Finance** group.

Step 4: Verify multi-factor access control mechanism

In this step you will verify the multi-factor access control policy that you set up in the previous step. You can use the following procedure to verify that **Robert Hatley** AD user can access your sample application because he belongs to the **Finance** group and AD users who do not belong to the **Finance** group cannot access the sample application.

1. On your client computer, open a browser window, and navigate to your sample application: **https://webserv1.contoso.com/claimapp**

   This action automatically redirects the request to the federation server and you are prompted to sign in with a username and password.
2. Type in the credentials of the **Robert Hatley** AD account that you created in Set up the lab environment for AD FS in Windows Server 2012 R2.

   You will be granted access to the application.
3. Type in the credentials of another AD user that does NOT belong to the **Finance** group. (For more information on how to create user accounts in AD, see http://technet.microsoft.com/en-us/library/cc783323%28v=ws.10%29.aspx).

   At this point, because of the access control policy that you set up in the previous step, an 'access denied' message is displayed for this AD user that does NOT belong to the **Finance** group. The default message text is **You are not authorized to access this site. Click here to sign out and sign in again or contact your administrator for permissions.** However, this text is fully customizable. For more information about how to customize the sign-in experience, see Customizing the AD FS Sign-in Pages.

See Also

Overview: Manage Risk with Multi-Factor Access Control

Set up the lab environment for AD FS in Windows Server 2012 R2

Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

---

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

In this guide

This guide provides the following information:

- Key Concepts - Authentication mechanisms in AD FS – description of the authentication mechanisms available in Active Directory Federation Services (AD FS) in Windows Server® 2012 R2
- Scenario Overview – a description of a scenario where you use Active Directory Federation Services (AD FS) to enable multifactor authentication (MFA) based on user's group membership.

  Note
  In AD FS in Windows Server® 2012 R2 you can enable MFA based on the network location, device identity, and user identity or group membership.

  For detailed step-by-step walkthrough instructions for configuring and verifying this scenario, see Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

Key Concepts - Authentication mechanisms in AD FS

Benefits of authentication mechanisms in AD FS

Active Directory Federation Services (AD FS) in Windows Server® 2012 R2 provides IT administrators with a richer, more flexible set of tools for authenticating users who want to access corporate resources. It empowers administrators with flexible control over the primary and the additional authentication methods, provides a rich management experience for configuring authentication polices (both through the user interface and Windows PowerShell), and enhances

the experience for the end users that access applications and services that are secured by AD FS. The following are some of the benefits of securing your application and services with AD FS in Windows Server® 2012 R2:

- Global authentication policy – a central management capability, from which an IT administrator can choose what authentication methods are used to authenticate users based on the network location from which they access protected resources. This enables administrators to do the following:
  - Mandate the use of more secure authentication methods for access requests from the extranet.
  - Enable device authentication for seamless second-factor authentication. This ties the user's identity to the registered device that is used to access the resource, thus offering more secure compound identity verification before protected resources are accessed.

    Note

    For more information about device object, Device Registration Service, Workplace Join, and the device as seamless second-factor authentication and SSO, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
  - Set MFA requirement for all extranet access or conditionally based on the user's identity, network location or a device that is used to access protected resources.
- Greater flexibility in configuring authentication policies: you can configure custom authentication policies for AD FS-secured resources with varying business values. For example, you can require MFA for application with high business impact.
- Ease of use: simple and intuitive management tools such as the GUI-based AD FS Management MMC snap-in and the Windows PowerShell cmdlets enable IT administrators to configure authentication policies with relative ease. With Windows PowerShell, you can script your solutions for use at scale and to automate mundane administrative tasks.
- Greater control over corporate assets: since as an administrator you can use AD FS to configure an authentication policy that applies to a specific resource, you have greater control over how corporate resources are secured. Applications cannot override the authentication policies specified by IT administrators. For sensitive applications and services, you can enable MFA requirement, device authentication, and optionally fresh authentication every time the resource is accessed.
- Support for custom MFA providers: for organizations that leverage third-party MFA methods, AD FS offers the ability to incorporate and use these authentication methods seamlessly.

Authentication scope

In AD FS in Windows Server® 2012 R2 you can specify an authentication policy at a global scope that is applicable to all applications and services that are secured by AD FS. You can also set authentication policies for specific applications and services (relying party trusts) that are secured by AD FS. Specifying an authentication policy for a particular application (per relying party trust) does not override the global authentication policy. If either global or per relying party trust authentication policy requires MFA, MFA will be triggered when the user tries to authenticate

to this relying party trust. The global authentication policy is a fallback for relying party trusts (applications and services) that do not have a specific authentication policy configured.

A global authentication policy applies to all relying parties that are secured by AD FS. You can configure the following settings as part of the global authentication policy:

- Authentication methods to be used for primary authentication
- Settings and methods for MFA
- Whether device authentication is enabled. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.

Per-relying party trust authentication policies apply specifically to attempts to access that relying party trust (application or service). You can configure the following settings as part of the per-relying party trust authentication policy:

- Whether users are required to provide their credentials each time at sign-in
- MFA settings based on the user/group, device registration, and access request location data

Primary and additional authentication methods

With AD FS in Windows Server® 2012 R2, in addition to the primary authentication mechanism, administrators can configure additional authentication methods. Primary authentication methods are built-in and are intended to validate users' identities. You can configure additional authentication factors to request that more information about the user's identity is provided and consequently ensure stronger authentication.

With primary authentication in AD FS in Windows Server® 2012 R2, you have the following options:

- For resources published to be accessed from outside the corporate network, Forms Authentication is selected by default. In addition, you can also enable Certificate Authentication (in other words, smart card-based authentication or user client certificate authentication that works with AD DS).
- For intranet resources, Windows Authentication is selected by default. In addition you can also enable Forms and/or Certificate Authentication.

By selecting more than one authentication method, you enable your users to have a choice of what method to authenticate with at the sign-in page for your application or service.

You can also enable device authentication for seamless second-factor authentication. This ties the user's identity to the registered device that is used to access the resource, thus offering more secure compound identity verification before protected resources are accessed.

Note

For more information about device object, Device Registration Service, Workplace Join, and the device as seamless second-factor authentication and SSO, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.

If you specify Windows Authentication method (default option) for your intranet resources, authentication requests undergo this method seamlessly on browsers that support Windows authentication.

Note

Windows authentication is not supported on all browsers. The authentication mechanism in AD FS in Windows Server® 2012 R2 detects the user's browser user agent and uses a configurable setting to determine whether that user agent supports Windows Authentication. Administrators can add to this list of user agents (via the Windows PowerShell `Set-AdfsProperties -WIASupportedUserAgents` command, in order to specify alternate user agent strings for browsers that support Windows Authentication. If the client's user agent does not support Windows Authentication, the default fallback method is Forms Authentication.

Configuring MFA

There are two parts to configure MFA in AD FS in Windows Server® 2012 R2: specifying the conditions under which MFA is required, and selecting an additional authentication method.

MFA settings

The following options are available for MFA settings (conditions under which to require MFA):

- You can require MFA for specific users and groups in the AD domain that your federation server is joined to.
- You can require MFA for either registered (workplace joined) or unregistered (not workplace joined) devices.

  Windows Server® 2012 R2 takes a user-centric approach to modern devices where device objects represent a relationship between user@device and a company. Device objects are a new class in AD in Windows Server® 2012 R2 that can be used to offer compound-identity when providing access to applications and services. A new component of AD FS - the device registration service (DRS) - provisions a device identity in Active Directory and sets a certificate on the consumer device that will be used to represent the device identity. You can then use this device identity to workplace join your device, in other words, to connect your personal device to the Active Directory of your workplace. When you join your personal device to your workplace, it becomes a known device and will provide seamless second-factor authentication to protected resources and applications. In other words, after a device is workplace joined, the user's identity is tied to this device and can be used for a seamless compound identity verification before a protected resource is accessed.

  For more information on workplace join and leave, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
- You can require MFA when the access request for the protected resources comes from either the extranet or the intranet.

Additional authentication methods

In order to enable MFA for any of these settings (conditions), you must select at least one additional authentication method. By default, in AD FS in Windows Server® 2012 R2, you can select Certificate Authentication (in other words, smart card-based authentication) as an additional authentication method.

Note
If you select Certificate Authentication, ensure that the smart card certificates have been provisioned securely and have pin requirements.

Third-party additional authentication methods

You can also configure and enable third-party authentication methods in AD FS in Windows Server® 2012 R2. In other words, you can use custom authentication methods implemented in third-party authentication providers (for example, Windows Azure Multi-Factor Authentication) to be used as MFA methods in AD FS. Once installed and registered with AD FS, you can select these MFA methods as part of the global authentication policy.

For step-by-step instructions for configuring Windows Azure Multi-Factor Authentication to function with AD FS, see the "Step 3: Configure MFA on your federation server" section in Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

Scenario Overview

In this scenario, you enable MFA based on the user's group membership data for a specific application. In other words, you will set up an authentication policy on your federation server to require MFA when users that belong to a certain group request access to a specific application that is hosted on a web server.

More specifically, in this scenario, you enable an authentication policy for a claims-based test application called **claimapp**, whereby an AD user **Robert Hatley** will be required to undergo MFA since he belongs to an AD group **Finance**.

The step-by step instructions to set up and verify this scenario are provided in Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications. In order to complete the steps in this walkthrough, you must set up a lab environment and follow the steps in Set up the lab environment for AD FS in Windows Server 2012 R2.

Other scenarios of enabling MFA in AD FS include the following:

- Enable MFA, if the access request comes from the extranet. You can modify the code presented in the "Set up MFA Policy" section of Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications with the following:

```
'c:[type ==
"http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork"
, value == "false"] =>
issue(type="http://schemas.microsoft.com/ws/2008/06/identity/cla
ims/authenticationmethod", value =
"http://schemas.microsoft.com/claims/multipleauthn" );'
```

- Enable MFA, if the access request comes from a non-workplace joined device. You can modify the code presented in the “Set up MFA Policy” section of Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications with the following:

```
'NOT
EXISTS([type=="http://schemas.microsoft.com/2012/01/devicecontex
t/claims/registrationid"]) => issue
(type="http://schemas.microsoft.com/ws/2008/06/identity/claims/a
uthenticationmethod", value =
"http://schemas.microsoft.com/claims/multipleauthn");'
```

- Enable MFA, if the access is coming from a user with a device that is workplace joined but not registered to this user. You can modify the code presented in the “Set up MFA Policy” section of Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications with the following:

```
‘c:[type=="http://schemas.microsoft.com/2012/01/devicecontext/cl
aims/isregistereduser", value == "false"] => issue
(type="http://schemas.microsoft.com/ws/2008/06/identity/claims/a
uthenticationmethod", value =
"http://schemas.microsoft.com/claims/multipleauthn");‘
```

See Also

Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

Set up the lab environment for AD FS in Windows Server 2012 R2

Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

About This Guide

This walkthrough provides instructions for configuring multifactor authentication (MFA) in Active Directory Federation Services (AD FS) in Windows Server® 2012 R2 based on the user’s group membership data.

For more information about MFA and authentication mechanisms in AD FS, see Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

This walkthrough consists of the following sections:

- Step 1: Setting up the lab environment
- Step 2: Verify the default AD FS authentication mechanism
- Step 3: Configure MFA on your federation server
- Step 4: Verify MFA mechanism

Step 1: Setting up the lab environment

In order to complete this walkthrough, you need an environment that consists of the following components:

- An Active Directory domain with a test user and group accounts, running on Windows Server® 2012 R2 or an Active Directory domain running on Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 with its schema upgraded to Windows Server® 2012 R2
- A federation server running on Windows Server® 2012 R2
- A web server that hosts your sample application
- A client computer from which you can access the sample application

Warning
It is highly recommended (both in a production and test environments) that you do not use the same computer to be your federation server and your web server.

In this environment, the federation server issues the claims that are required so that users can access the sample application. The Web server hosts a sample application that will trust the users who present the claims that the federation server issues.

For instructions on how to set up this environment, see Set up the lab environment for AD FS in Windows Server 2012 R2.

Step 2: Verify the default AD FS authentication mechanism

In this step you will verify the default AD FS access control mechanism (**Forms Authentication** for extranet and **Windows Authentication** for intranet), where the user is redirected to the AD FS sign-in page, provides valid credentials, and is granted access to the application. You can use the **Robert Hatley** AD account and the **claimapp** sample application that you configured in Set up the lab environment for AD FS in Windows Server 2012 R2.

1. On your client computer, open a browser window, and navigate to your sample application: **https://webserv1.contoso.com/claimapp**.

   This action automatically redirects the request to the federation server and you are prompted to sign in with a username and password.

2. Type in the credentials of the **Robert Hatley** AD account that you created in Set up the lab environment for AD FS in Windows Server 2012 R2.

   You will be granted access to the application.

Step 3: Configure MFA on your federation server

There are two parts to configuring MFA in AD FS in Windows Server® 2012 R2:

- Select an additional authentication method
- Set up MFA policy

Select an additional authentication method

In order to set up MFA, you must select an additional authentication method. In this walkthrough, for additional authentication method, you can choose between the following options:

- Select Certificate authentication method that is available in AD FS in Windows Server® 2012 R2 by default
- Configure and select Windows Azure Multi-Factor Authentication

Certificate authentication

Complete either of the following procedures to select Certificate authentication as the additional authentication method:

To configure Certificate authentication as an additional authentication method via the AD FS Management Console

1. On your federation server, in the AD FS Management Console, navigate to the **Authentication Policies** node, and under **Multi-factor Authentication** section, click the **Edit** link next to the **Global Settings** sub-section.
2. In the **Edit Global Authentication Policy** window, select **Certificate Authentication** as an additional authentication method, and then click **OK**.

To configure Certificate authentication as an additional authentication method via Windows PowerShell

1. On your federation server, open the Windows PowerShell command window and run the following command:

```
Set-AdfsGlobalAuthenticationPolicy –
AdditionalAuthenticationProvider CertificateAuthentication
```

Warning

To verify that this command ran successfully, you can run the `Get-AdfsGlobalAuthenticationPolicy` command.

Windows Azure Multi-Factor Authentication

Complete the following procedures in order to download and configure and select **Windows Azure Multi-Factor Authentication** as additional authentication on your federation server:

1. Create a Multi-Factor Authentication Provider via the Windows Azure Portal
2. Download the Windows Azure Multi-Factor Authentication Server
3. Install the Windows Azure Multi-Factor Authentication Server on your Federation Server
4. Configure Windows Azure Multi-Factor Authentication as an additional authentication method

Create a Multi-Factor Authentication Provider via the Windows Azure Portal

1. Log on to the Windows Azure Portal as an Administrator.
2. On the left, select Active Directory.
3. On the Active Directory page, at the top, select **Multi-Factor Auth Providers**. Then at the bottom, click **New**.
4. Under **App Services->Active Directory**, select **Multi-Factor Auth Provider**, and select **Quick Create**.
5. Under **App Services**, select **Active Auth Providers**, and select **Quick Create**.
6. Fill in the following fields and select **Create**.
   a. **Name** – The name of the Multi-Factor Auth Provider.
   b. **Usage Model** – The usage model of the Multi-Factor Authentication Provider.
      - **Per Authentication** – purchasing model that charges per authentication. Typically used for scenarios that use Windows Azure Multi-Factor Authentication in a consumer-facing application.
      - **Per Enabled User** – purchasing model that charges per enabled user. Typically used for employee-facing scenarios such as Office 365.

      For additional information on usage models, see Windows Azure pricing details.
   c. **Directory** – The Windows Azure Active Directory tenant that the Multi-Factor Authentication Provider is associated with. This is optional as the provider does not have to be linked to Windows Azure Active Directory when securing on-premises applications.
7. Once you click create, the Multi-Factor Authentication Provider will be created and you should see a message stating: Successfully created Multi-Factor Authentication Provider. Click **Ok**.

Next, you must download the Windows Azure Multi-Factor Authentication Server. You can do this by launching the Windows Azure Multi-Factor Authentication Portal through the Windows Azure portal.

Download the Windows Azure Multi-Factor Authentication Server

1. Log on to the Windows Azure Portal as an Administrator, and click on the Multi-Factor Authentication Provider you created in the procedure above. Then click the **Manage** button.

This launches the **Windows Azure Multi-Factor Authentication** portal.

2. In the **Windows Azure Multi-Factor Authentication** portal, click **Downloads**, and then click **Download** to download a copy of the Windows Azure Multi-Factor Authentication Server.

Once you have downloaded the executable for the Windows Azure Multi-Factor Authentication Server, you must install it on your federation server.

Install the Windows Azure Multi-Factor Authentication Server on your Federation Server

1. Download and double-click on the executable for the Windows Azure Multi-Factor Authentication Server. This will begin the installation.
2. On the License Agreement screen, read the agreement, select **I Agree** and click **Next**.
3. Ensure that the destination folder is correct and click **Next**.
4. Once the installation complete, click **Finish**.

You are now ready to launch the Windows Azure Multi-Factor Authentication server that you installed on your federation server and configure it as an additional authentication method.

Configure Windows Azure Multi-Factor Authentication as an additional authentication method

1. Launch **Windows Azure Multi-Factor Authentication** from where you installed it on your federation server, and on the Welcome page, check the **Skip using the Authentication Configuration Wizard** checkbox and click **Next**.
2. To activate the Multi-Factor Authentication Server, go back to the page in the Multi-Factor Authentication management portal where you downloaded the Multi-Factor Authentication Server and click the **Generate Activation Credentials** button. In the Multi-Factor Authentication Server user interface, enter the credentials that were generated and click **Activate**.
3. Next, the **Multi-Factor Authentication Server** user interface prompts you to run the **Multi-Server Configuration Wizard**. Select **No**.

   Important
   You can skip completing the **Multi-Server Configuration Wizard** given the lab environment with only one federation server that is used to complete this walkthrough. However, if your environment contains several federation servers, you must install the Multi-Factor Authentication Server and complete the **Multi-Server Configuration Wizard** on each federation server in order to enable replication between the Multi-Factor servers running on your federation servers.
4. In the **Multi-Factor Authentication Server** user interface, select the **Users** icon, click **Import from Active Directory**, select the **Robert Hatley** account to provision it in Windows Azure Multi-Factor Authentication, and then click **Import**.
5. In the **Users** list, select the **Robert Hatley** account, click **Edit**, and in the **Edit User** window, provide a cell phone number of this account, make sure the **Enabled** checkbox is checked, and then click **Apply**.

6. In the **Users** list, select the **Robert Hatley** account, and click **Test**. In the **Test User** window, provide the credentials for the **Robert Hatley** account. When the cell phone rings, press '#' to complete the account verification.
7. In the **Multi-Factor Authentication Server** user interface, select the **AD FS** icon, make sure that **Allow user enrollment**, **Allow users to select method** (including **Phone call** and **Text message**), **Use security questions for fallback** and **Enable logging** checkboxes are checked, click **Install AD FS Adapter**, and complete the **Multi-Factor Authentication AD FS Adapter** installation wizard.

   Notes

   The **Multi-Factor Authentication AD FS Adapter** installation wizard creates a security group called **PhoneFactor Admins** in your Active Directory and then adds the AD FS service account of your federation service to this group.

   It is recommended that you verify on your domain controller that the **PhoneFactor Admins** group is indeed created and that the AD FS service account is a member of this group.

   If necessary, add the AD FS service account to the **PhoneFactor Admins** group on your domain controller manually.

   For additional details on installing the AD FS Adapter, click the Help link in the top right corner of the Multi-Factor Authentication Server.
8. To register the adapter in your federation service, on your federation server, launch the Windows PowerShell command window, and run the following command: `\Program Files\Multi-Factor Authentication Server\Register-MultiFactorAuthenticationAdfsAdapter.ps1`. The adapter is now registered as **WindowsAzureMultiFactorAuthentication**. You must restart your AD FS service for the registration to take effect.
9. To configure Windows Azure Multi-Factor Authentication as the additional authentication method, in the AD FS Management Console, navigate to the **Authentication Policies** node, and under **Multi-factor Authentication** section, click the **Edit** link next to the **Global Settings** sub-section. In the **Edit Global Authentication Policy** window, select **Multi-Factor Authentication** as an additional authentication method, and then click **OK**.

   Note

   You can customize the name and description of the Windows Azure Multi-Factor Authentication method, as well as any configured third-party authentication method, as it appears in your AD FS UI, by running the **Set-AdfsAuthenticationProviderWebContent** cmdlet. For more information, see http://technet.microsoft.com/en-us/library/dn479401.aspx

Set up MFA policy

In order to enable MFA, you must set up the MFA policy on your federation server. For this walkthrough, per our MFA policy, **Robert Hatley** account is required to undergo MFA because he

belongs to the **Finance** group that you set up in Set up the lab environment for AD FS in Windows Server 2012 R2.

You can set up the MFA policy either via the AD FS Management Console or using the Windows PowerShell.

To configure the MFA policy based on user's group membership data for 'claimapp' via the AD FS Management Console

1. On your federation server, in the AD FS Management Console, navigate to **Authentication Policies\Per Relying Party Trust** node, and select the relying party trust that represents your sample application (**claimapp**).
2. Either in the **Actions** page or by right-clicking **claimapp**, select **Edit Custom Multi-factor Authentication**.
3. In the **Edit Relying Party Trust for claimapp** window, click the **Add** button next to the **Users/Groups** list. Type in **Finance** for the name of your AD group that you created in Set up the lab environment for AD FS in Windows Server 2012 R2, and click **Check Names**, and when the name is resolved, click **OK**.
4. Click **OK** in the **Edit Relying Party Trust for claimapp** window.

To configure the MFA policy based on user's group membership data for 'claimapp' via Windows PowerShell

1. On your federation server, open the Windows PowerShell command window and run the following command:

```
$rp = Get-AdfsRelyingPartyTrust -Name claimapp
```

2. In the same Windows PowerShell command window, run the following command:

```
$GroupMfaClaimTriggerRule = 'c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/grou
psid", Value =~ "^(?i) <group_SID>$"] => issue(Type =
"http://schemas.microsoft.com/ws/2008/06/identity/claims/auth
enticationmethod", Value =
"http://schemas.microsoft.com/claims/multipleauthn");'
Set-AdfsRelyingPartyTrust -TargetRelyingParty $rp -
AdditionalAuthenticationRules $GroupMfaClaimTriggerRule
```

Note
Make sure to replace <group_SID> with the value of the SID of your AD group **Finance**.

Step 4: Verify MFA mechanism

In this step you will verify the MFA functionality that you set up in the previous step. You can use the following procedure to verify that **Robert Hatley** AD user can access your sample application and this time is required to undergo MFA because he belongs to the **Finance** group.

1. On your client computer, open a browser window, and navigate to your sample application: **https://webserv1.contoso.com/claimapp**.

   This action automatically redirects the request to the federation server and you are prompted to sign in with a username and password.
2. Type in the credentials of the **Robert Hatley** AD account.

   At this point, because of the MFA policy that you configured, the user will be prompted to undergo additional authentication. The default message text is **For security reasons, we require additional information to verify your account.** However, this text is fully customizable. For more information about how to customize the sign-in experience, see Customizing the AD FS Sign-in Pages.

   If you configured Certificate authentication as the additional authentication method, the default message text is **Select a certificate that you want to use for authentication. If you cancel the operation, please close your browser and try again.**

   If you configured Windows Azure Multi-Factor Authentication as the additional authentication method, the default message text is **A call will be placed to your phone to complete your authentication.** For more information about signing in with Windows Azure Multi-Factor Authentication and using various options for the preferred method of verification, see Windows Azure Multi-Factor Authentication Overview.

See Also

Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

Set up the lab environment for AD FS in Windows Server 2012 R2

AD FS Design

- AD FS Design Guide in Windows Server 2012 R2
- AD FS Design Guide in Windows Server 2012

See Also

Active Directory Federation Services Overview

AD FS Design Guide in Windows Server 2012 R2

Active Directory Federation Services (AD FS) provides simplified, secured identity federation and Web single sign-on (SSO) capabilities for end users who want to access applications within an AD FS-secured enterprise, in federation partner organizations, or in the cloud.

In Windows Server® 2012 R2, AD FS includes a federation service role service that acts as an identity provider (authenticates users to provide security tokens to applications that trust AD FS) or as a federation provider (consumes tokens from other identity providers and then provides security tokens to applications that trust AD FS).

The function of providing extranet access to applications and services that are secured by AD FS in Windows Server 2012 R2 is now performed by a new Remote Access role service called Web Application Proxy. This is a departure from the prior versions of Windows Server in which this function was handled by an AD FS federation server proxy. Web Application Proxy is a server role designed to provide access for the AD FS-related extranet scenario and other extranet scenarios. For more information on Web Application Proxy, see Web Application Proxy Overview.

About this guide

This guide provides recommendations to help you plan a new deployment of AD FS, based on the requirements of your organization. This guide is intended for use by an infrastructure specialist or system architect. It highlights your main decision points as you plan your AD FS deployment. Before you read this guide, you should have a good understanding of how AD FS works on a functional level. For more information, see Understanding Key AD FS Concepts.

In this guide

- Identify Your AD FS Deployment Goals
- Plan Your AD FS Deployment Topology
- AD FS Requirements

Acknowledgments

Author: Gayana Bagdasaryan

Technical Reviewers: Jen Field, Samuel Devasahayam

See Also

AD FS Design

Identify Your AD FS Deployment Goals

Correctly identifying your Active Directory Federation Services (AD FS) deployment goals is essential for the success of your AD FS design project. Prioritize and, possibly, combine your deployment goals so that you can design and deploy AD FS by using an iterative approach. You can take advantage of existing, documented, and predefined AD FS deployment goals that are relevant to the AD FS designs and develop a working solution for your situation.

Prior versions of AD FS were most commonly deployed to achieve the following:

- Providing your employees or customers with a web-based, SSO experience when accessing claims-based applications within your enterprise.
- Providing your employees or customers with a web-based, SSO experience to access resources in any federation partner organization.
- Providing your employees or customers with a Web-based, SSO experience when remote accessing internally hosted Web sites or services.
- Providing your employees or customers with a web-based, SSO experience when accessing resources or services in the cloud.

In addition to these, AD FS in Windows Server® 2012 R2 adds functionality that can help you achieve the following:

- Device workplace join for SSO and seamless second factor authentication. This enables organizations to allow access from user's personal devices and manage the risk when providing this access.
- Managing risk with multi-factor access control. AD FS provides a rich level of authorization that controls who has access to what applications. This can be based on user attributes (UPN, email, security group membership, authentication strength, etc.), device attributes (whether the device is workplace joined) or request attributes (network location, IP address, or user agent).
- Managing risk with additional multi-factor authentication for sensitive applications. AD FS allows you to control policies to potentially require multi-factor authentication globally or on a per application basis. In addition, AD FS provides extensibility points for any multi-factor vendor to integrate deeply for a secure and seamless multi-factor experience for end users.
- Providing authentication and authorization capabilities for accessing web resources from the extranet that are protected by the Web Application Proxy.

To summarize, AD FS in Windows Server 2012 R2 can be deployed to achieve the following goals in your organization:

Enable your users to access resources on their personal devices from anywhere

- Workplace join that enables users to join their personal devices to corporate Active Directory and as a result gain access and seamless experiences when accessing corporate resources from these devices.
- Pre-authentication of resources inside the corporate network that are protected by the Web Application proxy and accessed from the internet.
- Password change to enable users to change their password from any workplace joined device when their password has expired so that they can continue to access resources.

Enhance your access control risk management tools

Managing risk is an important aspect of governance and compliance in every IT organization. There are numerous access control risk management enhancements in AD FS in Windows Server® 2012 R2, including the following:

- Flexible controls based on network location to govern how a user authenticates to access an AD FS-secured application.
- Flexible policy to determine if a user needs to perform multi-factor authentication based on the user's data, device data, and network location.
- Per-application control to ignore SSO and force the user to provide credentials every time they access a sensitive application.
- Flexible per-application access policy based on user data, device data, or network location.
- AD FS Extranet Lockout, which enables administrators to protect Active Directory accounts from brute force attacks from the internet.
- Access revocation for any workplace joined device that is disabled or deleted in Active Directory.

Use AD FS to enhance the sign-in experience

The following are new AD FS capabilities in Windows Server® 2012 R2 that enable administrator to customize and enhance the sign-in experience:

- Unified customization of the AD FS service, where the changes are made once and then automatically propagated to the rest of the AD FS federation servers in a given farm.
- Updated sign-in pages that look modern and cater to different form factors automatically.
- Support for automatic fallback to forms-based authentication for devices that are not joined to the corporate domain but are still used generate access requests from within the corporate network (intranet).
- Simple controls to customize the company logo, illustration image, standard links for IT support, home page, privacy, etc.
- Customization of description messages in the sign-in pages.
- Customization of web themes.

- Home Realm Discovery (HRD) based on organizational suffix of the user for enhanced privacy of a company's partners.
- HRD filtering on a per-application basis to automatically pick a realm based on the application.
- One-click error reporting for easier IT troubleshooting.
- Customizable error messages.
- User authentication choice when more than one authentication provider is available.

See Also

AD FS Design Guide in Windows Server 2012 R2

Plan Your AD FS Deployment Topology

The first step in planning a deployment of Active Directory Federation Services (AD FS) is to determine the right deployment topology to meet the needs of your organization.

Before you read this topic, review how AD FS data is stored and replicated to other federation servers in a federation server farm and make sure you understand the purpose of and the replication methods that can be used for the underlying data that is stored in the AD FS configuration database.

There are two database types that you can use to store AD FS configuration data: Windows Internal Database (WID) and Microsoft SQL Server. For more information, see The Role of the AD FS Configuration Database. Review the various benefits and limitations that are associated with using either WID or SQL Server as the AD FS configuration database, along with the various application scenarios that they support and then make your selection.

Important
To implement basic redundancy, load balancing, and the option to scale the Federation Service (if required), we recommend that you deploy at least two federation servers per federation server farm for all production environments, regardless of the type of database that you will use.

Determining which type of AD FS configuration database to use

AD FS uses a database to store configuration and—in some cases—transactional data related to the Federation Service. You can use the AD FS software to select either the built-in Windows Internal Database (WID) or Microsoft SQL Server 2008 or newer to store the data in the Federation Service.

For most purposes, the two database types are relatively equivalent. However, there are some differences to be aware of before you begin reading more about the various deployment

topologies that you can use with AD FS. The following table describes the differences in supported features between a WID database and a SQL Server database.

	Feature	Supported by WID?	Supported by SQL Server?
AD FS features	Federation server farm deployment	Yes, with a limit of five federation servers for each farm	Yes. There is no enforced limit for the number of federation servers that you can deploy in a single farm
	SAML artifact resolution **Note** This feature is not required for Microsoft Online Services, Microsoft Office 365, Microsoft Exchange, or Microsoft Office SharePoint scenarios.	No	Yes
	SAML/WS-Federation token replay detection		
Database features	Basic database redundancy using pull replication, where one or more servers hosting a read-only copy of the database request changes that are made on a source server that hosts a read/write copy of the database	Yes	No
	Database redundancy using high-availability solutions, such as failover clustering or mirroring (at the database layer only)	No	Yes

	Feature	Supported by WID?	Supported by SQL Server?
	Note All AD FS deployment topologies support clustering at the AD FS service layer.		

SQL Server considerations

You should consider the following deployment facts if you select SQL Server as the configuration database for your AD FS deployment.

- **SAML features and their effect on database size and growth**. When either the SAML artifact resolution or SAML token replay detection features are enabled, AD FS stores information in the SQL Server configuration database for each AD FS token that is issued. The growth of the SQL Server database as a result of this activity is not considered to be significant, and it depends on the configured token replay retention period. Each artifact record has a size of approximately 30 kilobytes (KB).
- **Number of servers required for your deployment**. You will need to add at least one additional server (to the total number of servers required to deploy your AD FS infrastructure) that will act as a dedicated host of the SQL Server instance. If you plan to use failover clustering or mirroring to provide fault tolerance and scalability for the SQL Server configuration database, a minimum of two SQL servers is required.

How the configuration database type you select may impact hardware resources

The impact to hardware resources on a federation server that is deployed in a farm using WID as opposed to a federation server that is deployed in a farm using the SQL Server database is not significant. However, it is important to consider that when you use WID for the farm, each federation server in that farm must store, manage, and maintain replication changes for its local copy of the AD FS configuration database while also continuing to provide the normal operations that the Federation Service requires.

In comparison, federation servers that are deployed in a farm that uses the SQL Server database do not necessarily contain a local instance of the AD FS configuration database. Therefore, they may make slightly fewer demands on hardware resources.

Where to place a federation server

As a security best practice, place AD FS federation servers in front of a firewall and connect them to your corporate network to prevent exposure from the Internet. This is important because federation servers have full authorization to grant security tokens. Therefore, they should have the same protection as a domain controller. If a federation server is compromised, a malicious user has the ability to issue full access tokens to all Web applications and to federation servers that are protected by AD FS.

Note

As a security best practice, avoid having your federation servers directly accessible on the Internet. Consider giving your federation servers direct Internet access only when you are setting up a test lab environment or when your organization does not have a perimeter network.

For typical corporate networks, an intranet-facing firewall is established between the corporate network and the perimeter network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this situation, the federation server sits inside the corporate network, and it is not directly accessible by Internet clients.

Note

Client computers that are connected to the corporate network can communicate directly with the federation server through Windows Integrated Authentication.

A federation server proxy should be placed in the perimeter network before you configure your firewall servers for use with AD FS.

Supported deployment topologies

The following topics describe the various deployment topologies that you can use with AD FS. They also describe the benefits and limitations associated with each deployment topology so that you can select the most appropriate topology for your specific business needs.

- Federation Server Farm Using WID
- Federation Server Farm Using WID and Proxies
- Federation Server Farm Using SQL Server

See Also

AD FS Design Guide in Windows Server 2012 R2

Federation Server Farm Using WID

The default topology for Active Directory Federation Services (AD FS) is a federation server farm, using the Windows Internal Database (WID), that consists of up to five federation servers hosting

your organization's Federation Service. In this topology, AD FS uses WID as the store for the AD FS configuration database for all federation servers that are joined to that farm. The farm replicates and maintains the Federation Service data in the configuration database across each server in the farm.

The act of creating the first federation server in a farm also creates a new Federation Service. When you use WID for the AD FS configuration database, the first federation server that you create in the farm is referred to as the *primary federation server*. This means that this computer is configured with a read/write copy of the AD FS configuration database.

All other federation servers that you configure for this farm are referred to as *secondary federation servers* because they must replicate any changes that are made on the primary federation server to the read-only copies of the AD FS configuration database that they store locally.

Important
We recommend the use of at least two federation servers in a load-balanced configuration.

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Organizations with 100 or fewer configured trust relationships that need to provide their internal users (logged on to computers that are physically connected to the corporate network) with single sign-on (SSO) access to federated applications or services
- Organizations that want to provide their internal users with SSO access to Microsoft Online Services or Microsoft Office 365
- Smaller organizations that require redundant, scalable services

Note
Organizations with larger databases should consider using the Federation Server Farm Using SQL Server deployment topology. Organizations with users who log in from outside the network should consider using either the Federation Server Farm Using WID and Proxies topology or the Federation Server Farm Using SQL Server topology.

What are the benefits of using this topology?

- Provides SSO access to internal users
- Data and Federation Service redundancy (each federation server replicates changes to other federation servers in the same farm)
- The farm can be scaled out by adding up to five federation servers
- WID is included with Windows; therefore, no need to purchase SQL Server

What are the limitations of using this topology?

- A WID farm has a limit of five federation servers.
- A WID farm does not support token replay detection or artifact resolution (part of the Security Assertion Markup Language (SAML) protocol).

Server placement and network layout recommendations

When you are ready to start deploying this topology in your network, you should plan on placing all of the federation servers in your corporate network behind a Network Load Balancing (NLB) host that can be configured for an NLB cluster with a dedicated cluster Domain Name System (DNS) name and cluster IP address.

Note

This cluster DNS name must match the Federation Service name, for example, fs.fabrikam.com.

The NLB host can use the settings that are defined in this NLB cluster to allocate client requests to the individual federation servers. The following illustration shows how the fictional Fabrikam, Inc., company sets up the first phase of its deployment using a two-computer federation server farm (fs1 and fs2) with WID and the positioning of a DNS server and a single NLB host that is wired to the corporate network.

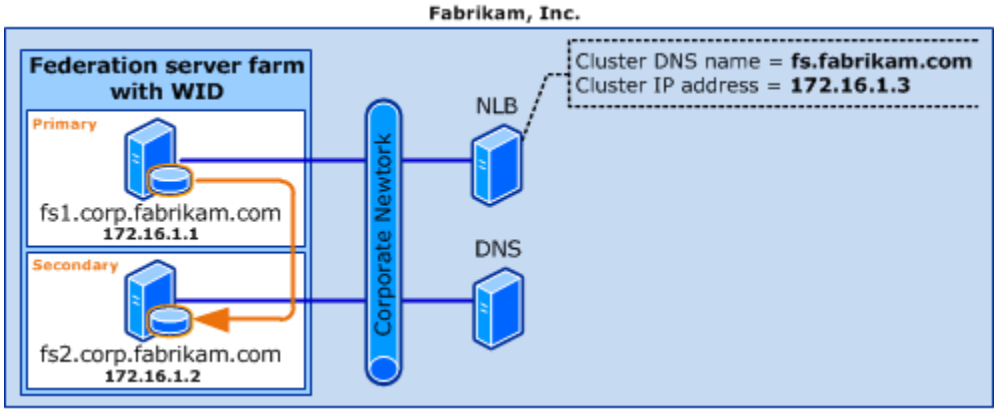

Note

If there is a failure on this single NLB host, users will not be able to access federated applications or services. Add additional NLB hosts if your business requirements do not allow having a single point of failure.

For more information about how to configure your networking environment for use with federation servers, see the Name Resolution Requirements section in AD FS Requirements.

See Also

Plan Your AD FS Deployment Topology

AD FS Design Guide in Windows Server 2012 R2

Federation Server Farm Using WID and Proxies

This deployment topology for Active Directory Federation Services (AD FS) is identical to the federation server farm with Windows Internal Database (WID) topology, but it adds proxie computers to the perimeter network to support external users. These proxies redirect client authentication requests that come from outside your corporate network to the federation server farm. In previous versions of AD FS, these proxies were called federation server proxies.

Important

In Active Directory Federation Services (AD FS) in Windows Server 2012 R2, the role of a federation server proxy is handled by a new Remote Access role service called Web Application Proxy. To enable your AD FS for accessibility from outside the corporate network, which was the purpose of deploying a federation server proxy in legacy versions of AD FS, such as AD FS 2.0 and AD FS in Windows Server 2012, you can deploy one or more web application proxies for AD FS in Windows Server 2012 R2.

In the context of AD FS, Web Application Proxy functions as an AD FS federation server proxy. In addition to this, Web Application Proxy provides reverse proxy functionality for web applications inside your corporate network to enable users on any device to access them from outside the corporate network. For more information, about the Web Application Proxy role service, see Web Application Proxy Overview.

To plan the deployment of Web Application proxy, you can review the information in the following topics:

- Step 1: Plan the Web Application Proxy Infrastructure
- Step 2: Plan the Web Application Proxy Server

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Organizations with 100 or fewer configured trust relationships that need to provide both their internal users and external users (who are logged on to computers that are physically located outside the corporate network) with single sign-on (SSO) access to federated applications or services

- Organizations that need to provide both their internal users and external users with SSO access to Microsoft Office 365
- Smaller organizations that have external users and require redundant, scalable services

What are the benefits of using this topology?

- The same benefits as listed for the Federation Server Farm Using WID topology, plus the benefit of providing additional access for external users

What are the limitations of using this topology?

- The same limitations as listed for the Federation Server Farm Using WID topology

Server placement and network layout recommendations

To deploy this topology, in addition to adding two web application proxies, you must make sure that your perimeter network can also provide access to a Domain Name System (DNS) server and to a second Network Load Balancing (NLB) host. The second NLB host must be configured with an NLB cluster that uses an Internet-accessible cluster IP address, and it must use the same cluster DNS name setting as the previous NLB cluster that you configured on the corporate network (fs.fabrikam.com). The web application proxies should also be configured with Internet-accessible IP addresses.

The following illustration shows the existing federation server farm with WID topology that was described previously and how the fictional Fabrikam, Inc., company provides access to a perimeter DNS server, adds a second NLB host with the same cluster DNS name (fs.fabrikam.com), and adds two web application proxies (wap1 and wap2) to the perimeter network.

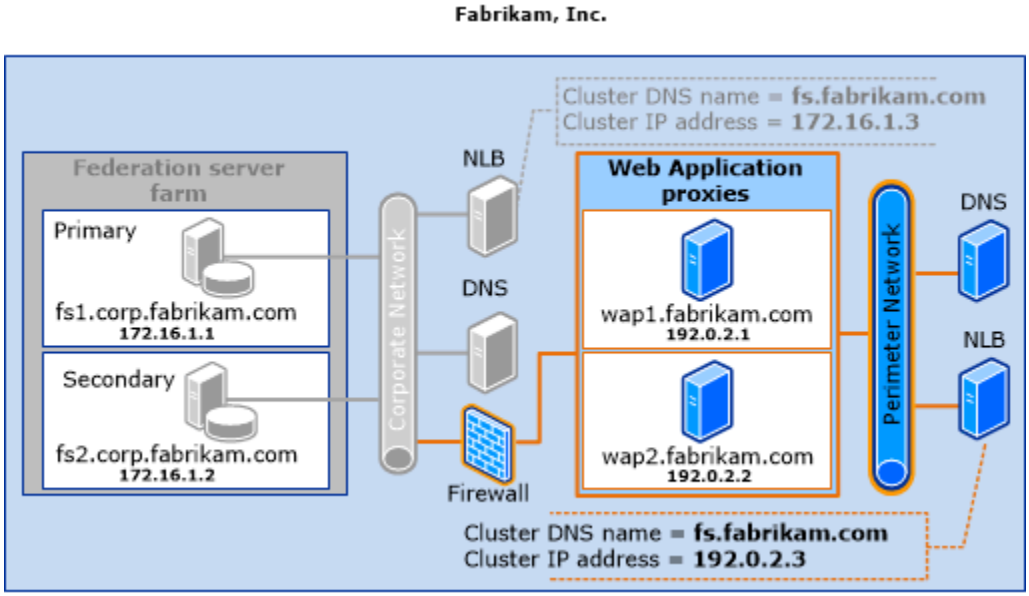

For more information about how to configure your networking environment for use with federation servers or web application proxies, see “Name Resolution Requirements” section in AD FS Requirements and Step 1: Plan the Web Application Proxy Infrastructure.

See Also

Plan Your AD FS Deployment Topology

AD FS Design Guide in Windows Server 2012 R2

Federation Server Farm Using SQL Server

This topology for Active Directory Federation Services (AD FS) differs from the federation server farm using Windows Internal Database (WID) deployment topology in that it does not replicate the data to each federation server in the farm. Instead, all federation servers in the farm can read and write data into a common database that is stored on a server running Microsoft SQL Server that is located in the corporate network.

Important

If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Large organizations with more than 100 trust relationships that need to provide both their internal users and external users with single sign-on (SSO) access to federated application or services
- Organizations that already use SQL Server and want to take advantage of their existing tools and expertise

What are the benefits of using this topology?

- Support for larger numbers of trust relationships (more than 100)
- Support for token replay detection (a security feature) and artifact resolution (part of the Security Assertion Markup Language (SAML) 2.0 protocol)
- Support for the full benefits of SQL Server, such as database mirroring, failover clustering, reporting, and management tools

What are the limitations of using this topology?

- This topology does not provide database redundancy by default. Although a federation server farm with WID topology automatically replicates the WID database on each federation server in the farm, the federation server farm with SQL Server topology contains only one copy of the database

Note
SQL Server supports many different data and application redundancy options including failover clustering, database mirroring, and several different types of SQL Server replication.

The Microsoft Information Technology (IT) department uses SQL Server database mirroring in high-safety (synchronous) mode and failover clustering to provide high-availability support for the SQL Server instance. SQL Server transactional (peer-to-peer) and merge replication have not been tested by the AD FS product team at Microsoft. For more information about SQL Server, see High Availability Solutions Overview or Selecting the Appropriate Type of Replication.

Supported SQL Server Versions

The following SQL server versions are supported with AD FS in Windows Server 2012 R2:

- SQL Server 2008 / R2
- SQL Server 2012

Server placement and network layout recommendations

Similar to the federation server farm with WID topology, all of the federation servers in the farm are configured to use one cluster Domain Name System (DNS) name (which represents the Federation Service name) and one cluster IP address as part of the Network Load Balancing

(NLB) cluster configuration. This helps the NLB host allocate client requests to the individual federation servers. Federation server proxies can be used to proxy client requests to the federation server farm.

The following illustration shows how the fictional Contoso Pharmaceuticals company deployed its federation server farm with SQL Server topology in the corporate network. It also shows how that company configured the perimeter network with access to a DNS server, an additional NLB host that uses the same cluster DNS name (fs.contoso.com) that is used on the corporate network NLB cluster, and with two web application proxies (wap1 and wap2).

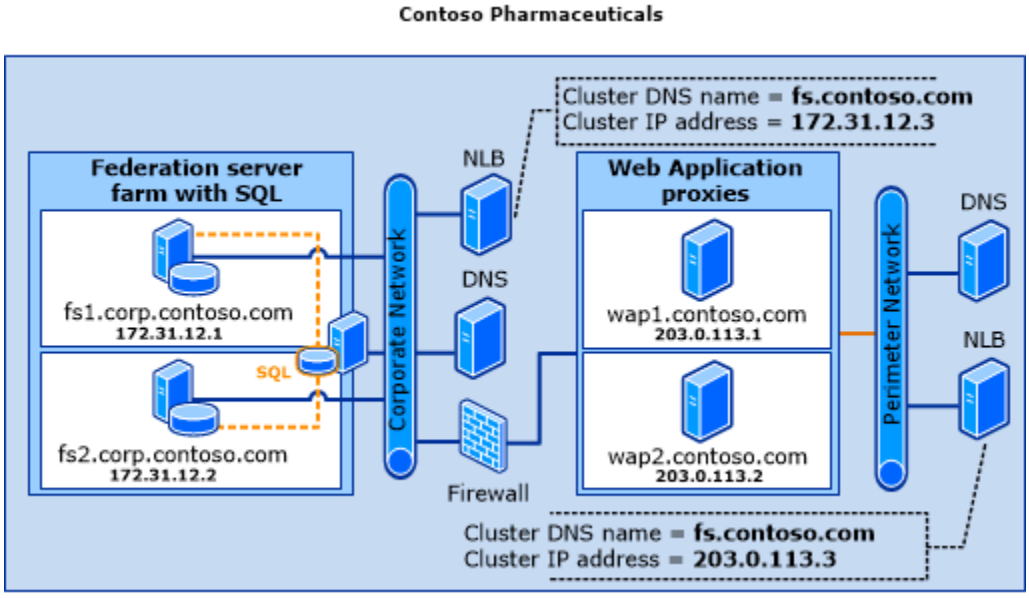

For more information about how to configure your networking environment for use with federation servers or web application proxies, see “Name Resolution Requirements” section in AD FS Requirements and Step 1: Plan the Web Application Proxy Infrastructure.

See Also

Plan Your AD FS Deployment Topology

AD FS Design Guide in Windows Server 2012 R2

AD FS Requirements

The following are the various requirements that you must conform to when deploying AD FS:

- Certificate requirements
- Hardware requirements
- Software requirements
- AD DS requirements

- Configuration database requirements
- Browser requirements
- Network requirements
- Attribute store requirements
- Application requirements
- Authentication requirements
- Workplace join requirements
- Cryptography requirements

Certificate requirements

Certificates play the most critical role in securing communications between federation servers, Web Application Proxies, claims-aware applications, and Web clients. The requirements for certificates vary, depending on whether you are setting up a federation server or federation server proxy computer, as described in this section.

Federation server certificates

Certificate type	Requirements, Support & Things to Know
Secure Sockets Layer (SSL) certificate: This is a standard SSL certificate that is used for securing communications between federation servers and clients.	• This certificate must be a publicly trusted* X509 v3 certificate. • All clients that access any AD FS endpoint must trust this certificate. It is strongly recommended to use certificates that are issued by a public (third-party) certification authority (CA). You can use a self-signed SSL certificate successfully on federation servers in a test lab environment. However, for a production environment, we recommend that you obtain the certificate from a public CA. • Supports any key size supported by Windows Server 2012 R2 for SSL certificates. • Does not support certificates that use CNG keys. • When used together with Workplace Join/Device Registration Service, the subject alternative name of the SSL certificate for the AD FS service must contain the value enterpriseregistration that is followed by the User Principal Name (UPN) suffix of your organization, for

	example, enterpriseregistration.contoso.com. • Wild card certificates are supported. When you create your AD FS farm, you will be prompted to provide the service name for the AD FS service (for example, **adfs.contoso.com**. • It is strongly recommended to use the same SSL certificate for the Web Application Proxy. This is however **required** to be the same when supporting Windows Integrated Authentication endpoints through the Web Application Proxy and when Extended Protection Authentication is turned on (default setting). • The Subject name of this certificate is used to represent the Federation Service name for each instance of AD FS that you deploy. For this reason, you may want to consider choosing a Subject name on any new CA-issued certificates that best represents the name of your company or organization to partners. The identity of the certificate must match the federation service name (for example, fs.contoso.com).The identity is either a subject alternative name extension of type dNSName or, if there are no subject alternative name entries, the subject name specified as a common name. Multiple subject alternative name entries can be present in the certificate, provided one of them matches the federation service name. • **Important:** it's strongly recommended to use the same SSL certificate across all nodes of your AD FS farm as well as all Web Application proxies in your AD FS farm.
Service communication certificate: This certificate enables WCF message security for securing communications between federation servers.	• By default, the SSL certificate is used as the service communications certificate. But you also have the option to configure another certificate as the service communication certificate. • **Important:** if you are using the SSL

	certificate as the service communication certificate, when the SSL certificate expires, make sure to configure the renewed SSL certificate as your service communication certificate. This does not happen automatically. • This certificate must be trusted by clients of AD FS that use WCF Message Security. • We recommend that you use a server authentication certificate that is issued by a public (third-party) certification authority (CA). • The service communication certificate cannot be a certificate that uses CNG keys. • This certificate can be managed using the AD FS Management console.
Token-signing certificate: This is a standard X509 certificate that is used for securely signing all tokens that the federation server issues.	• By default, AD FS creates a self-signed certificate with 2048 bit keys. • CA issued certificates are also supported and can be changed using the AD FS Management snap-in • CA issued certificates must be stored & accessed through a CSP Crypto Provider. • The token signing certificate cannot be a certificate that uses CNG keys. • AD FS does not require externally enrolled certificates for token signing. AD FS automatically renews these self-signed certificates before they expire, first configuring the new certificates as secondary certificates to allow for partners to consume them, then flipping to primary in a process called automatic certificate rollover.We recommend that you use the default, automatically generated certificates for token signing. If your organization has policies that require different certificates to be configured for token signing, you can specify the certificates at installation time using Powershell (use the –SigningCertificateThumbprint parameter of

	the Install-AdfsFarm cmdlet). After installation, you can view and manage token signing certificates using the AD FS Management console or Powershell cmdlets Set-AdfsCertificate and Get-AdfsCertificate. When externally enrolled certificates are used for token signing, AD FS does not perform automatic certificate renewal or rollover. This process must be performed by an administrator. To allow for certificate rollover when one certificate is close to expiring, a secondary token signing certificate can be configured in AD FS. By default, all token signing certificates are published in federation metadata, but only the primary token-signing certificate is used by AD FS to actually sign tokens.
Token-decryption/encryption certificate: This is a standard X509 certificate that is used to decrypt/encrypt any incoming tokens. It is also published in federation metadata.	• By default, AD FS creates a self-signed certificate with 2048 bit keys. • CA issued certificates are also supported and can be changed using the AD FS Management snap-in • CA issued certificates must be stored & accessed through a CSP Crypto Provider. • The token-decryption/encryption certificate cannot be a certificate that uses CNG keys. • By default, AD FS generates and uses its own, internally generated and self-signed certificates for token decryption. AD FS does not require externally enrolled certificates for this purpose. In addition, AD FS automatically renews these self-signed certificates before they expire. **We recommend that you use the default, automatically generated certificates for token decryption.** If your organization has policies that require different certificates to be configured for token decryption, you can specify the

	certificates at installation time using Powershell (use the –DecryptionCertificateThumbprint parameter of the Install-AdfsFarm cmdlet). After installation, you can view and manage token decryption certificates using the AD FS Management console or Powershell cmdlets Set-AdfsCertificate and Get-AdfsCertificate. **When externally enrolled certificates are used for token decryption, AD FS does not perform automatic certificate renewal. This process must be performed by an administrator**. • The AD FS service account must have access to the token-signing certificate's private key in the personal store of the local computer. This is taken care of by Setup. You can also use the AD FS Management snap-in to ensure this access if you subsequently change the token-signing certificate.

Caution

Certificates that are used for token-signing and token-decrypting/encrypting are critical to the stability of the Federation Service. Customers managing their own token-signing & token-decrypting/encrypting certificates should ensure that these certificates are backed up and are available independently during a recovery event.

Note

In AD FS you can change the Secure Hash Algorithm (SHA) level that is used for digital signatures to either SHA-1 or SHA-256 (more secure). AD FS does not support the use of certificates with other hash methods, such as MD5 (the default hash algorithm that is used with the Makecert.exe command-line tool). As a security best practice, we recommend that you use SHA-256 (which is set by default) for all signatures. SHA-1 is recommended for use only in scenarios in which you must interoperate with a product that does not support communications using SHA-256, such as a non-Microsoft product or legacy versions of AD FS.

Note

After you receive a certificate from a CA, make sure that all certificates are imported into the personal certificate store of the local computer. You can import certificates to the personal store with the Certificates MMC snap-in.

Hardware requirements

The following minimum and recommended hardware requirements apply to the AD FS federation servers in Windows Server 2012 R2:

Hardware requirement	Minimum requirement	Recommended requirement
CPU speed	1.4 GHz 64-bit processor	Quad-core, 2 GHz
RAM	512 MB	4 GB
Disk space	32 GB	100 MB

Software requirements

The following AD FS requirements are for the server functionality that is built into the Windows Server® 2012 R2 operating system:

- For extranet access, you must deploy the Web Application Proxy role service - part of the Windows Server® 2012 R2 Remote Access server role. Prior versions of a federation server proxy are not supported with AD FS in Windows Server® 2012 R2.
- A federation server and the Web Application Proxy role service cannot be installed on the same computer.

AD DS requirements

Domain controller requirements

For AD FS to operate successfully, domain controllers in all user domains and in the domain that AD FS servers are joined to must be running Windows Server 2008 or later.

Domain functional-level requirements

All user account domains and the domain that the AD FS servers are joined to must be operating at the domain functional level of Windows Server 2003 or higher.

Schema requirements

- AD FS does not require schema changes or functional-level modifications to AD DS.
- To use Workplace Join functionality, the schema of the forest that AD FS servers are joined to must be set to Windows Server 2012 R2.

Service account requirements

- Any standard service account can be used as a service account for AD FS. Group Managed Service accounts are also supported. This requires at least one domain controller (it is recommended that you deploy two or more) that is running Windows Server 2012 or higher.
- For Kerberos authentication to function between domain-joined clients and AD FS, the 'HOST/<adfs_service_name>' must be registered as a SPN on the service account. By default, AD FS will configure this when creating a new AD FS farm if it has sufficient permissions to perform this operation.

- The AD FS service account must be trusted in every user domain that contains users authenticating to the AD FS service.

Domain Requirements

- All AD FS servers must be a joined to an AD DS domain.
- All AD FS servers within a farm must be deployed in a single domain.
- The domain that the AD FS servers are joined to must trust every user account domain that contains users authenticating to the AD FS service.

Multi Forest Requirements

- The domain that the AD FS servers are joined to must trust every user account domain or forest that contains users authenticating to the AD FS service.
- The AD FS service account must be trusted in every user domain that contains users authenticating to the AD FS service.

Configuration database requirements

The following are the requirements and restrictions that apply based on the type of configuration store:

WID

- A maximum of 5 AD FS servers in a single AD FS farm is supported for AD FS in Windows Server 2012 R2.
- A maximum of 100 configuration objects (relying party trusts and claims provider trusts only) is supported for AD FS in Windows Server 2012 R2.
- Artifact resolution profile in SAML 2.0 is not supported in the WID configuration database.Token Replay Detection is not supported in the WID configuration database. This functionality is only used only in scenarios where AD FS is acting as the federation provider and consuming security tokens from external claims providers.
- Deploying AD FS servers in distinct data centers for failover or geographic load balancing is supported as long as the number of servers does not exceed 5.

SQL Server

For AD FS in Windows Server 2012 R2, you can use SQL Server 2008 and higher, including SQL Server 2012.

Browser requirements

The following are the browser requirements for the AD FS clients:

Browser	Win7	Win 8.1	WS2008R2	WS2012/ WS2012R2	iOS 6+	Android 4.3+	Mac OS-X 10.7+
IE 10.0	X	X	X	X	N/A	N/A	N/A

IE 11.0	X	X	Not Tested	X	N/A	N/A	N/A
Windows Web Authentication Broker	Not Tested	X	N/A	N/A	N/A	N/A	N/A
Firefox [v21]	X	X	Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Safari [v7]	Not Tested	Not Tested	Not Tested	Not Tested	X		X
Chrome [v27]	X	X	Not Tested	X			X

Note
AD FS supports both the 32bit and 64bit versions of Internet Explorer browsers showing in the above table.

Cookies

AD FS creates session-based and persistent cookies that must be stored on client computers to provide sign-in, sign-out, single sign-on (SSO), and other functionality. Therefore, the client browser must be configured to accept cookies. Cookies that are used for authentication are always Secure Hypertext Transfer Protocol (HTTPS) session cookies that are written for the originating server. If the client browser is not configured to allow these cookies, AD FS cannot function correctly. Persistent cookies are used to preserve user selection of the claims provider. You can disable them by using a configuration setting in the configuration file for the AD FS sign-in pages. Support for TLS/SSL is required for security reasons.

Network requirements

Configuring the following network services appropriately is critical for successful deployment of AD FS in your organization:

Configuring Corporate Firewall

Both the firewall located between the Web Application Proxy and the federation server farm and the firewall between the clients and the Web Application Proxy must have TCP port 443 enabled inbound.

In addition, if client user certificate authentication (clientTLS authentication using X509 user certificates) is required, AD FS in Windows Server 2012 R2 requires that TCP port 49443 be enabled inbound on the firewall between the clients and the Web Application Proxy. This is not required on the firewall between the Web Application Proxy and the federation servers).

Configuring DNS

- For intranet access, all clients accessing AD FS service within the internal corporate network (intranet) must be able to resolve the AD FS service name (name provided by the SSL certificate) to the load balancer for the AD FS servers or the AD FS server.

- For extranet access, all clients accessing AD FS service from outside the corporate network (extranet/internet) must be able to resolve the AD FS service name (name provided by the SSL certificate) to the load balancer for the Web Application Proxy servers or the Web Application Proxy server.
- For extranet access to function properly, each Web Application Proxy server in the DMZ must be able to resolve AD FS service name (name provided by the SSL certificate) to the load balancer for the AD FS servers or the AD FS server. This can be achieved using an alternate DNS server in the DMZ network or by changing local server resolution using HOSTS file.
- For Windows Integrated authentication to work inside the network and outside the network for a subset of endpoints exposed through the Web Application Proxy, you must use an A record (not CNAME) to point to the load balancers.

For information on configuring corporate DNS for the federation service and Device Registration Service, see Configure Corporate DNS for the Federation Service and DRS.

For information on configuring corporate DNS for Web Application proxies, see the “Configure DNS” section in Step 1: Configure the Web Application Proxy Infrastructure.

For information about how to configure a cluster IP address or cluster FQDN using NLB, see Specifying the Cluster Parameters at http://go.microsoft.com/fwlink/?LinkId=75282.

Attribute store requirements

AD FS requires at least one attribute store to be used for authenticating users and extracting security claims for those users. For a list of attribute stores that AD FS supports, see The Role of Attribute Stores.

Note
AD FS automatically creates an “Active Directory” attribute store, by default. Attribute store requirements depend on whether your organization is acting as the account partner (hosting the federated users) or the resource partner (hosting the federated application).

LDAP Attribute Stores

When you work with other Lightweight Directory Access Protocol (LDAP)-based attribute stores, you must connect to an LDAP server that supports Windows Integrated authentication. The LDAP connection string must also be written in the format of an LDAP URL, as described in RFC 2255.

It is also required that the service account for the AD FS service has the right to retrieve user information in the LDAP attribute store.

SQL Server Attribute Stores

For AD FS in Windows Server 2012 R2 to operate successfully, computers that host the SQL Server attribute store must be running either Microsoft SQL Server 2008 or higher. When you work with SQL-based attribute stores, you also must configure a connection string.

Custom Attribute Stores

You can develop custom attribute stores to enable advanced scenarios.

- The policy language that is built into AD FS can reference custom attribute stores so that any of the following scenarios can be enhanced:

    - Creating claims for a locally authenticated user
    - Supplementing claims for an externally authenticated user
    - Authorizing a user to obtain a token
    - Authorizing a service to obtain a token on behavior of a user
    - Issuing additional data in security tokens issued by AD FS to relying parties.
- All custom attribute stores must be built on top of .NET 4.0 or higher.

When you work with a custom attribute store, you might also have to configure a connection string. In that case, you can enter a custom code of your choice that enables a connection to your custom attribute store. The connection string in this situation is a set of name/value pairs that are interpreted as implemented by the developer of the custom attribute store.For more information about developing and using custom attribute stores, see Attribute Store Overview.

Application requirements

AD FS supports claims-aware applications that use the following protocols:

- WS-Federation
- WS-Trust
- SAML 2.0 protocol using IDPLite, SPLite & eGov1.5 profiles.
- OAuth 2.0 Authorization Grant Profile

AD FS also supports authentication and authorization for any non-claims-aware applications that are supported by the Web Application Proxy.

Authentication requirements

AD DS Authentication (Primary Authentication)

For intranet access, the following standard authentication mechanisms for AD DS are supported:

- Windows Integrated Authentication using Negotiate for Kerberos & NTLM
- Forms Authentication using username/passwords
- Certificate Authentication using certificates mapped to user accounts in AD DS

For extranet access, the following authentication mechanisms are supported:

- Forms Authentication using username/passwords
- Certificate Authentication using certificates that are mapped to user accounts in AD DS
- Windows Integrated Authentication using Negotiate (NTLM only) for WS-Trust endpoints that accept Windows Integrated Authentication.

For Certificate Authentication:

- Extends to smartcards that can be pin protected.
- The GUI for the user to enter their pin is not provided by AD FS and is required to be part of the client operating system that is displayed when using client TLS.
- The reader and cryptographic service provider (CSP) for the smart card must work on the computer where the browser is located.

- The smart card certificate must chain up to a trusted root on all the AD FS servers and Web Application Proxy servers.
- The certificate must map to the user account in AD DS by either of the following methods:
  - The certificate subject name corresponds to the LDAP distinguished name of a user account in AD DS.
  - The certificate subject altname extension has the user principal name (UPN) of a user account in AD DS.

For seamless Windows Integrated Authentication using Kerberos in the intranet,

- It is required for the service name to be part of the Trusted Sites or the Local Intranet sites.
- In addition, the HOST/<adfs_service_name> SPN must be set on the service account that the AD FS farm runs on.

Multi-Factor Authentication

AD FS supports additional authentication (beyond primary authentication supported by AD DS) using a provider model whereby vendors/customers can build their own multi-factor authentication adapter that an administrator can register and use during login.

Every MFA adapter must be built on top of .NET 4.5.

For more information on MFA, see Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

Device Authentication

AD FS supports device authentication using certificates provisioned by the Device Registration Service during the act of an end user workplace joining their device.

Workplace join requirements

End users can workplace join their devices to an organization using AD FS. This is supported by the Device Registration Service in AD FS. As a result, end users get the additional benefit of SSO across the applications supported by AD FS. In addition, administrators can manage risk by restricting access to applications only to devices that have been workplace joined to the organization. Below are the following requirements to enable this scenario.

- AD FS supports workplace join for Windows 8.1 and iOS 5+ devices
- To use Workplace Join functionality, the schema of the forest that AD FS servers are joined to must be Windows Server 2012 R2.
- The subject alternative name of the SSL certificate for AD FS service must contain the value enterpriseregistration that is followed by the User Principal Name (UPN) suffix of your organization, for example, enterpriseregistration.corp.contoso.com.

Cryptography requirements

The following table provides additional cryptography support information on the AD FS token signing, token encryption/decryption functionality:

Algorithm	Key lengths	Protocols/Applications/Comments
TripleDES – Default 192 (Supported 192 – 256) - http://www.w3.org/2001/04/xmlenc#tripledes-cbc	>= 192	Supported algorithm for Encrypting the security token.
AES128 - http://www.w3.org/2001/04/xmlenc#aes128-cbc	128	Supported algorithm for Encrypting the security token.
AES192 - http://www.w3.org/2001/04/xmlenc#aes192-cbc	192	Supported algorithm for Encrypting the security token.
AES256 - http://www.w3.org/2001/04/xmlenc#aes256-cbc	256	**Default**. Supported algorithm for Encrypting the security token.
TripleDESKeyWrap - http://www.w3.org/2001/04/xmlenc#kw-tripledes	All Key sizes supported by .NET 4.0+	Supported algorithm for Encrypting the symmetric key that encrypts a security token.
AES128KeyWrap - http://www.w3.org/2001/04/xmlenc#kw-aes128	128	Supported algorithm for Encrypting the symmetric key that encrypts the security token.
AES192KeyWrap - http://www.w3.org/2001/04/xmlenc#kw-aes192	192	Supported algorithm for Encrypting the symmetric key that encrypts the security token.
AES256KeyWrap - http://www.w3.org/2001/04/xmlenc#kw-aes256	256	Supported algorithm for Encrypting the symmetric key that encrypts the security token.
RsaV15KeyWrap - http://www.w3.org/2001/04/xmlenc#rsa-1_5	1024	Supported algorithm for Encrypting the symmetric key that encrypts the security token.
RsaOaepKeyWrap - http://www.w3.org/2001/04/xmlenc#rsa-oaep-mgf1p	1024	Default. Supported algorithm for Encrypting the symmetric key that encrypts the security token.
SHA1- http://www.w3.org/PICS/DSig/SHA1_1_0.html	N/A	Used by AD FS Server in artifact SourceId generation: In this scenario, the STS uses SHA1 (per

		the recommendation in the SAML 2.0 standard) to create a short 160 bit value for the artifact sourceiD. Also used by the ADFS web agent (legacy component from WS2003 timeframe) to identify changes in a “last updated” time value so that it knows when to update information from the STS.
SHA1withRSA- http://www.w3.org/PICS/DSig/RSA-SHA1_1_0.html	N/A	Used in cases when AD FS Server validates the signature of SAML AuthenticationRequest, sign the artifact resolution request or response, create token-signing certificate. In these cases, SHA256 is the default, and SHA1 is only used if the partner (relying party) cannot support SHA256 and must use SHA1.

See Also

AD FS Design Guide in Windows Server 2012 R2

AD FS Design Guide in Windows Server 2012

Note
For information about how to deploy AD FS in Windows Server 2012 R2, see Windows Server 2012 R2 AD FS Deployment Guide.

You can use Active Directory® Federation Services (AD FS) with the Windows Server® 2012 operating system in a federation services provider role to seamlessly authenticate your users to any Web-based services or applications that reside in a resource partner organization, without the need for administrators to create or maintain external trusts or forest trusts between the networks of both organizations and without the need for the users to log on a second time. The process of authenticating to one network while accessing resources in another network—without the burden of repeated logon actions by users—is known as single sign-on (SSO).

About this guide

This guide provides recommendations to help you plan a new deployment of AD FS, based on the requirements of your organization (also referred to in this guide as deployment goals) and the particular design that you want to create. This guide is intended for use by an infrastructure specialist or system architect. It highlights your main decision points as you plan your AD FS deployment. Before you read this guide, you should have a good understanding of how AD FS works on a functional level. You should also have a good understanding of the organizational requirements that will be reflected in your AD FS design.

This guide describes a set of deployment goals that are based on three primary AD FS designs, and it helps you decide the most appropriate design for your environment. You can use these deployment goals to form one of the following comprehensive AD FS designs or a custom design that meets the needs of your environment:

- Federated Web SSO to support business-to-business (B2B) scenarios and to support collaboration between business units with independent forests
- Web SSO to support customer access to applications in business-to-consumer (B2C) scenarios

For each design, you will find guidelines for gathering the required data about your environment. You can then use these guidelines to plan and design your AD FS deployment. After you read this guide and finish gathering, documenting, and mapping your organization's requirements, you will have the information necessary to begin deploying AD FS using the guidance in the Windows Server 2012 AD FS Deployment Guide.

In this guide

- Identifying Your AD FS Deployment Goals
- Mapping Your Deployment Goals to an AD FS Design
- Determine Your AD FS Deployment Topology
- Planning Your Deployment
- Planning Federation Server Placement
- Planning Federation Server Proxy Placement
- Planning for AD FS Server Capacity
- Appendix A: Reviewing AD FS Requirements

Acknowledgments

Author: Nick Pierson

Technical Reviewers: Matt Steele, Lu Zhao (Migration/Topologies/Requirements), Jen Field (Capacity Planning/Topologies/AD DS Claims), Jan Alexander (Claims Engine/Pipeline/Language), Krish Shenoy (Planning Federation Servers and Federation Server Proxies/Requirements)

Editor: Jim Becker

Identifying Your AD FS Deployment Goals

Correctly identifying your Active Directory Federation Services (AD FS) deployment goals is essential for the success of your AD FS design project. Depending on the size of your organization and the level of involvement that you want to provide for the information technology (IT) staff in any partner organizations, form a project team that can clearly articulate real-world deployment issues in a vision statement. Make sure that the members of this team understand the direction in which your deployment project must move in order to reach your AD FS deployment goals.

When you write your vision statement, identify, clarify, and refine your deployment goals. Prioritize and, possibly, combine your deployment goals so that you can design and deploy AD FS by using an iterative approach. You can take advantage of existing, documented, and predefined AD FS deployment goals that are relevant to the AD FS designs and develop a working solution for your situation.

The following table lists the tasks for articulating, refining, and documenting your AD FS deployment goals.

Task	Reference links
Evaluate predefined AD FS deployment goals that are provided in this section of the guide, and combine one or more goals to reach your organizational objectives	• Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services • Provide Your Active Directory Users Access to the Applications and Services of Other Organizations • Provide Users in Another Organization Access to Your Claims-Aware Applications and Services
Map one goal or a combination of any of the predefined AD FS deployment goals to an existing AD FS design	• Mapping Your Deployment Goals to an AD FS Design

Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services

When you are an administrator in the account partner organization in an Active Directory Federation Services (AD FS) deployment and you have a deployment goal to provide single-sign-on (SSO) access for employees on the corporate network to your hosted resources:

- Employees who are logged on to an Active Directory forest in the corporate network can use SSO to access multiple applications or services in the perimeter network in your own organization. These applications and services are secured by AD FS.

  For example, Fabrikam may want corporate network employees to have federated access to Web-based applications that are hosted in the perimeter network for Fabrikam.
- Remote employees who are logged on to an Active Directory domain can obtain AD FS tokens from the federation server in your organization to gain federated access to AD FS-secured Web-based applications or services that also reside in your organization.
- Information in the Active Directory attribute store can be populated into the employees' AD FS tokens.

The following components are required for this deployment goal:

- **Active Directory Domain Services (AD DS):** AD DS contains the employees' user accounts that are used to generate AD FS tokens. Information, such as group memberships and attributes, is populated into AD FS tokens as group claims and custom claims.

  Note
  You can also use Lightweight Directory Access Protocol (LDAP) or Structured Query Language (SQL) to contain the identities for AD FS token generation.
- **Corporate DNS:** This implementation of Domain Name System (DNS) contains a simple host (A) resource record so that intranet clients can locate the account federation server. This implementation of DNS may also host other DNS records that are required in the corporate network. For more information, see Name Resolution Requirements for Federation Servers.
- **Account partner federation server:** This federation server is joined to a domain in the account partner forest. It authenticates employee user accounts and generates AD FS tokens. The client computer for the employee performs Windows Integrated Authentication against this federation server to generate an AD FS token. For more information, see Review the Role of the Federation Server in the Account Partner.

  The account partner federation server can authenticate the following users:
  - Employees with user accounts in this domain
  - Employees with user accounts anywhere in this forest
  - Employees with user accounts anywhere in forests that are trusted by this forest (through a two-way Windows trust)
- **Employee:** An employee accesses a Web-based service (through an application) or a Web-based application (through a supported Web browser) while he or she is logged on to the corporate network. The employee's client computer on the corporate network communicates directly with the federation server for authentication.

After reviewing the information in the linked topics, you can begin deploying this goal by following the steps in Checklist: Implementing a Federated Web SSO Design.

The following illustration shows each of the required components for this AD FS deployment goal.

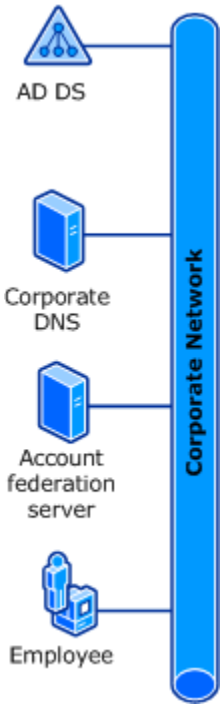

Provide Your Active Directory Users Access to the Applications and Services of Other Organizations

This Active Directory Federation Services (AD FS) deployment goal builds on the goal in Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services.

When you are an administrator in the account partner organization and you have a deployment goal to provide federated access for employees to hosted resources in another organization:

- Employees who are logged on to an Active Directory domain in the corporate network can use single-sign-on (SSO) functionality to access multiple Web-based applications or services, which are secured by AD FS, when the applications or services are in a different organization. For more information, see Federated Web SSO Design.

  For example, Fabrikam may want corporate network employees to have federated access to Web services that are hosted in Contoso.
- Remote employees who are logged on to an Active Directory domain can obtain AD FS tokens from the federation server in your organization to gain federated access to AD FS–secured Web-based applications or services that are hosted in another organization.

  For example, Fabrikam may want its remote employees to have federated access to AD FS–secured services that are hosted in Contoso, without requiring the Fabrikam employees to be on the Fabrikam corporate network.

In addition to the foundational components that are described in Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services and that are shaded in the following illustration, the following components are required for this deployment goal:

- **Account partner federation server proxy:** Employees that access the federated service or application from the Internet can use this AD FS component to perform authentication. By default, this component performs forms authentication, but it can also perform basic authentication. You can also configure this component to perform Secure Sockets Layer (SSL) client authentication if employees at your organization have certificates to present. For more information, see Where to Place a Federation Server Proxy.
- **Perimeter DNS:** This implementation of Domain Name System (DNS) provides the host names for the perimeter network. For more information about how to configure perimeter DNS for a federation server proxy, see Name Resolution Requirements for Federation Server Proxies.
- **Remote employee:** The remote employee accesses a Web-based application (through a supported Web browser) or a Web-based service (through an application), using valid credentials from the corporate network, while the employee is offsite using the Internet. The employee's client computer in the remote location communicates directly with the federation server proxy to generate a token and authenticate to the application or service.

After reviewing the information in the linked topics, you can begin deploying this goal by following the steps in Checklist: Implementing a Federated Web SSO Design.

The following illustration shows each of the required components for this AD FS deployment goal.

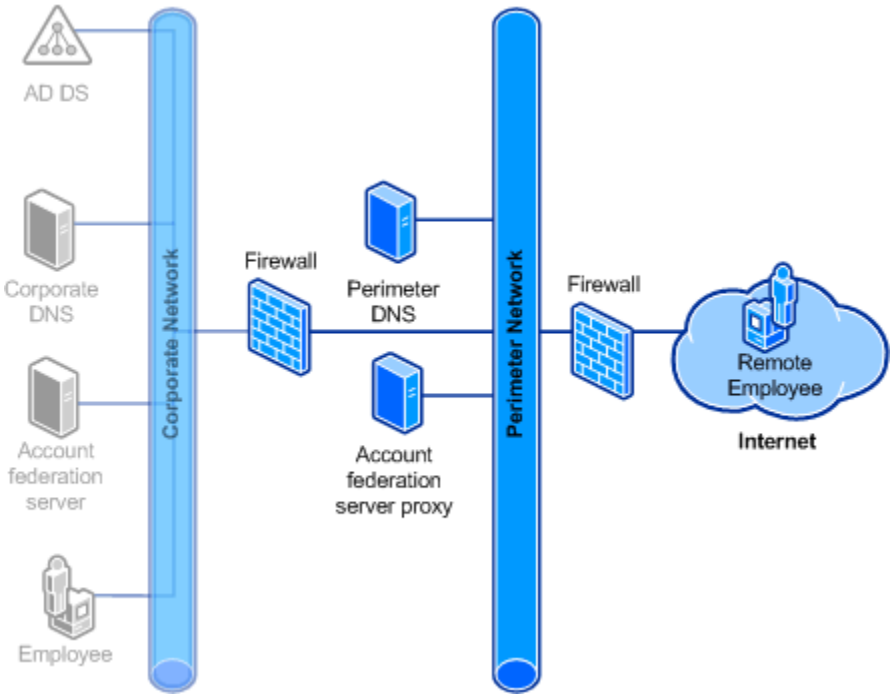

Provide Users in Another Organization Access to Your Claims-Aware Applications and Services

When you are an administrator in the resource partner organization in Active Directory Federation Services (AD FS) and you have a deployment goal to provide federated access for users in another organization (the account partner organization) to a claims-aware application or a Web-based service that is located in your organization (the resource partner organization):

- Federated users both in your organization and in organizations who have configured a federation trust to your organization (account partner organizations) can access the AD FS secured application or service that is hosted by your organization. For more information, see Federated Web SSO Design.

  For example, Fabrikam may want its corporate network employees to have federated access to Web services that are hosted in Contoso.

- Federated users who have no direct association with a trusted organization (such as individual customers), who are logged on to an attribute store that is hosted in your perimeter network, can access multiple AD FS-secured applications, which are also hosted in your perimeter network, by logging on one time from client computers that are located on the Internet. In other words, when you host customer accounts to enable access to applications or services in your perimeter network, customers that you host in an attribute store can access one or more applications or services in the perimeter network simply by logging on once. For more information, see Web SSO Design.

  For example, Fabrikam may want its customers to have single-sign-on (SSO) access to multiple applications or services that are hosted in its perimeter network.

The following components are required for this deployment goal:

- **Active Directory Domain Services (AD DS):** The resource partner federation server must be joined to an Active Directory domain.
- **Perimeter DNS:** Domain Name System (DNS) should contain a simple host (A) resource record so that client computers can locate the resource partner federation server and the Web server. The DNS server may host other DNS records that are also required in the perimeter network. For more information, see Name Resolution Requirements for Federation Servers.
- **Resource partner federation server:** The resource partner federation server validates AD FS tokens that the account partners send. Account partner discovery is performed through this federation server. For more information, see Review the Role of the Federation Server in the Resource Partner.
- **Web server:** The Web server can host either a Web application or a Web service. The Web server confirms that it receives valid AD FS tokens from federated users before it allows access to the protected Web application or Web service.

By using Windows Identity Foundation (WIF), you can develop your Web application or service so that it accepts federated user logon requests that are made with any standard logon method, such as user name and password.

After reviewing the information in the linked topics, you can begin deploying this goal by following the steps in Checklist: Implementing a Federated Web SSO Design and Checklist: Implementing a Web SSO Design.

The following illustration shows each of the required components for this AD FS deployment goal.

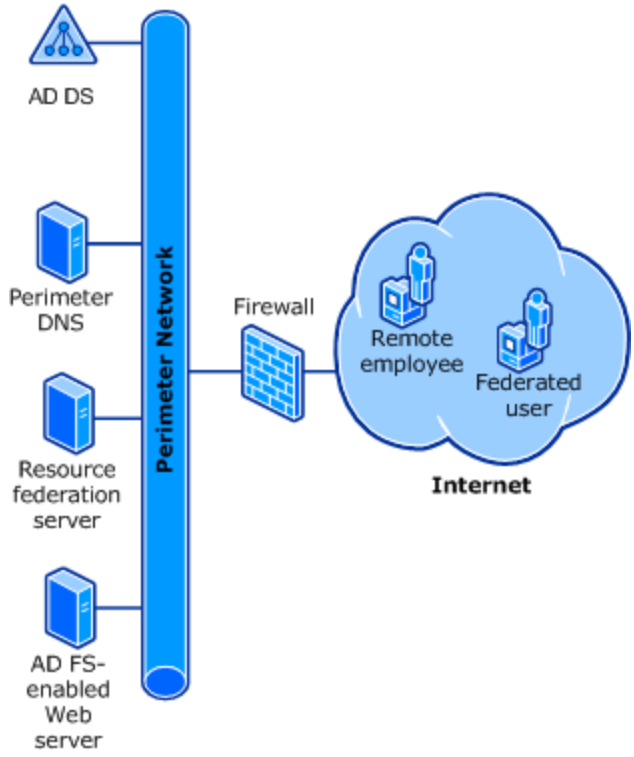

Mapping Your Deployment Goals to an AD FS Design

After you finish reviewing the existing Active Directory Federation Services (AD FS) deployment goals and you determine which goals are related to your deployment, you can map those goals to a specific AD FS design. For more information about AD FS predefined deployment goals, see Identifying Your AD FS Deployment Goals.

Use the following table to determine which AD FS design maps to the appropriate combination of AD FS deployment goals for your organization. This table refers only to the two primary AD FS designs, as described in this guide. However, you can create a hybrid or custom AD FS design by using any combination of the AD FS deployment goals to meet the needs of your organization.

AD FS deployment goal	Web SSO Design	Federated Web SSO Design
Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services	No	Yes, in the account partner
Provide Your Active Directory Users Access to the Applications and Services of Other Organizations	No	Yes, optional in the account partner
Provide Users in Another Organization Access to Your Claims-Aware Applications and Services	Yes	Yes

Web SSO Design

In the Web Single-Sign-On (SSO) design in Active Directory Federation Services (AD FS), users must authenticate only once to access multiple AD FS-secured applications or services. In this design all users are external, and no federation trust exists because there are no partner organizations. Typically, you deploy this design when you want to provide individual consumer or customer access to one or more AD FS–secured services or applications over the Internet, as shown in the following illustration.

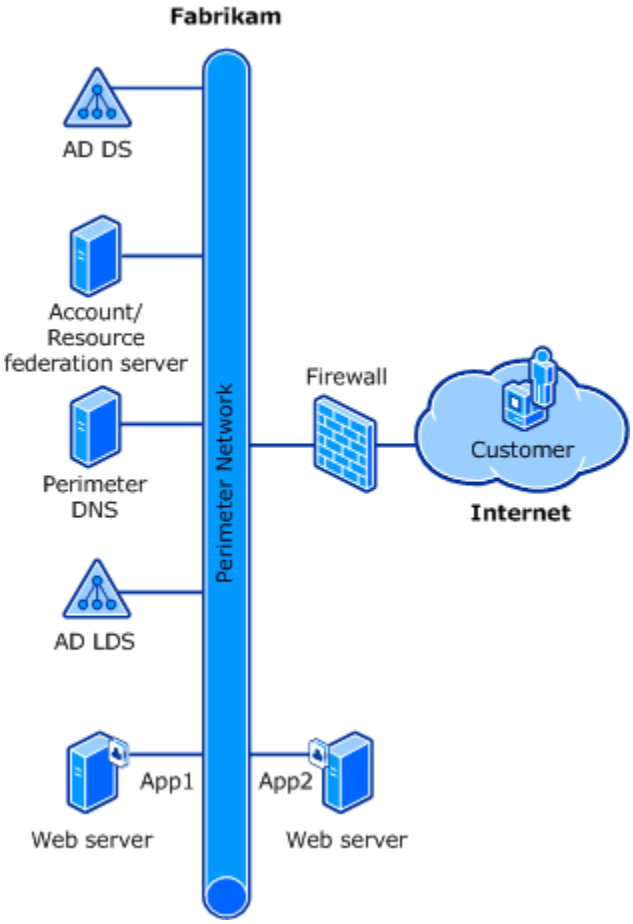

With the Web SSO design, an organization that typically hosts an AD FS-secured application or service in a perimeter network can maintain a separate store of customer accounts in the perimeter network, which makes it easier to isolate customer accounts from employee accounts.

You can manage the local accounts for customers in the perimeter network by using either Active Directory Domain Services (AD DS), SQL Server, or a custom attribute store.

This design coincides with the deployment goal in Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services.

For a list of detailed tasks that you can use to plan and deploy your Web SSO design, see Checklist: Implementing a Web SSO Design.

Federated Web SSO Design

The Federated Web Single-Sign-On (SSO) design in Active Directory Federation Services (AD FS) involves secure communication that spans multiple firewalls, perimeter networks, and name-resolution servers—in addition to the entire Internet routing infrastructure.

Typically, this design is used when two organizations agree to create a federation trust relationship to allow users in one organization (the account partner organization) to access Web-based applications or services, which are secured by AD FS, in the other organization (the resource partner organization).

In other words, a federation trust relationship is the embodiment of a business-level agreement or partnership between two organizations. As shown in the following illustration, you can establish a federation trust relationship between two businesses, which results in an end-to-end federation scenario.

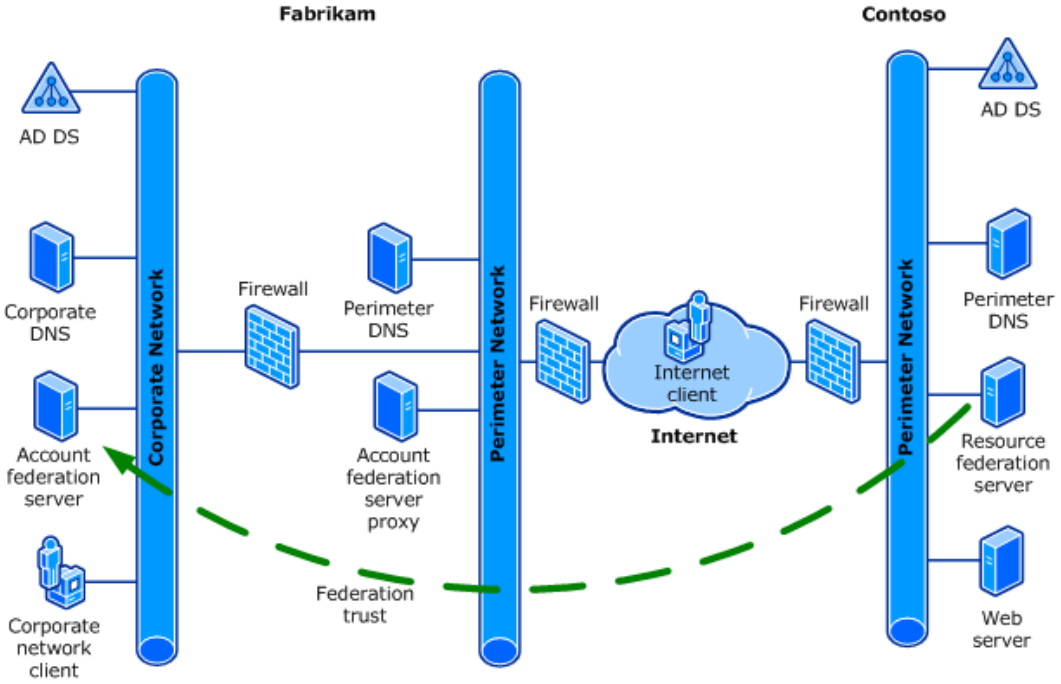

The one-way arrow in the illustration signifies the direction of the federation trust, which—like the direction of Windows trusts—always points to the account side of the forest. This means that authentication flows from the account partner organization to the resource partner organization.

In this Federated Web SSO design, two federation servers (one in Fabrikam and the other in Contoso) route authentication requests from user accounts in Fabrikam to Web-based applications or services in Contoso.

Note

For additional security, you can use federation server proxies to relay requests to federation servers that are not directly accessible from the Internet.

In this example, Fabrikam is the identity, or account, provider. The Fabrikam portion of the Federated Web SSO design uses the following AD FS deployment goal:

- Provide Your Active Directory Users Access to the Applications and Services of Other Organizations

Contoso is the resource provider. The Contoso portion of the Federated Web SSO design achieves the following AD FS deployment goals:

- Provide Users in Another Organization Access to Your Claims-Aware Applications and Services
- Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services

For a list of detailed tasks that you can use to plan and deploy the Federated Web SSO design, see Checklist: Implementing a Federated Web SSO Design.

Determine Your AD FS Deployment Topology

The first step in planning a deployment of Active Directory Federation Services (AD FS) is to determine the right deployment topology to meet the single sign-on (SSO) needs of your organization. The topics in this section describe the various deployment topologies that you can use with AD FS. They also describe the benefits and limitations associated with each deployment topology so that you can select the most appropriate topology for your specific business needs.

Before you read this deployment topology topic, we recommend that you first complete the tasks in the order shown in the following table.

Recommended task	Description	Reference
Review how AD FS data is stored and replicated to other federation servers in a federation server farm.	Understand the purpose of and the replication methods that can be used for the underlying data that is stored in the AD FS configuration database. This topic introduces the concepts of the configuration database and describes the two database types: Windows Internal Database (WID) and Microsoft SQL Server.	The Role of the AD FS Configuration Database
Select the type of AD FS configuration database that you will deploy in your organization.	Review the various benefits and limitations that are associated with using either WID or SQL Server as the AD FS configuration database, along with the various application scenarios that they support.	AD FS Deployment Topology Considerations

Note

To implement basic redundancy, load balancing, and the option to scale the Federation Service (if required), we recommend that you deploy at least two federation servers per federation server farm for all production environments, regardless of the type of database that you will use.

When you have reviewed the content in the previous table, proceed to the following topics in this section:

- Stand-Alone Federation Server Using WID
- Federation Server Farm Using WID
- Federation Server Farm Using WID and Proxies
- Federation Server Farm Using SQL Server

After you finish selecting your AD FS deployment topology, we recommend that you review the topic Planning for AD FS Server Capacity to determine the recommended number of servers that you will need to deploy to support this topology.

AD FS Deployment Topology Considerations

This topic describes important considerations to help you plan and design which Active Directory Federation Services (AD FS) deployment topology to use in your production environment. This topic is a starting point for reviewing and assessing considerations that affect what features or capabilities will be available to you after you deploy AD FS. For example, depending on which database type you choose to store the AD FS configuration database will ultimately determine whether you can implement certain Security Assertion Markup Language (SAML) features that require SQL Server.

Determining which type of AD FS configuration database to use

AD FS uses a database to store configuration and—in some cases—transactional data related to the Federation Service. You can use the AD FS software to select either the built-in Windows Internal Database (WID) or Microsoft SQL Server 2005 or newer to store the data in the Federation Service.

For most purposes, the two database types are relatively equivalent. However, there are some differences to be aware of before you begin reading more about the various deployment topologies that you can use with AD FS. The following table describes the differences in supported features between a WID database and a SQL Server database.

	Feature	Supported by WID?	Supported by SQL Server?	More information about this feature
AD FS features	Federation server farm deployment	Yes, with a limit of five federation servers for each farm	Yes. There is no enforced limit for the number of federation servers that you can deploy in a single farm	Determine Your AD FS Deployment Topology
	SAML artifact resolution **Note** This feature is not required for Microsoft Online Services, Microsoft Office 365, Microsoft Exchange, or Microsoft Office SharePoint scenarios.	No	Yes	The Role of the AD FS Configuration Database Best Practices for Secure Planning and Deployment of AD FS

	Feature	Supported by WID?	Supported by SQL Server?	More information about this feature
	SAML/WS-Federation token replay detection			
Database features	Basic database redundancy using pull replication, where one or more servers hosting a read-only copy of the database request changes that are made on a source server that hosts a read/write copy of the database	Yes	No	The Role of the AD FS Configuration Database
	Database redundancy using high-availability solutions, such as failover clustering or mirroring (at the database layer only) **Note** All AD FS deployment topologies support clustering at the AD FS service layer.	No	Yes	The Role of the AD FS Configuration Database High Availability Solutions Overview (http://go.microsoft.com/fwlink/?LinkId=179853)

SQL Server considerations

You should consider the following deployment facts if you select SQL Server as the configuration database for your AD FS deployment.

- **SAML features and their effect on database size and growth**. When either the SAML artifact resolution or SAML token replay detection features are enabled, AD FS stores information in the SQL Server configuration database for each AD FS token that is issued. The growth of the SQL Server database as a result of this activity is not considered to be significant, and it depends on the configured token replay retention period. Each artifact record has a size of approximately 30 kilobytes (KB).
- **Number of servers required for your deployment**. You will need to add at least one additional server (to the total number of servers required to deploy your AD FS infrastructure) that will act as a dedicated host of the SQL Server instance. If you plan to use failover clustering or mirroring to provide fault tolerance and scalability for the SQL Server configuration database, a minimum of two SQL servers is required.

How the configuration database type you select may impact hardware resources

The impact to hardware resources on a federation server that is deployed in a farm using WID as opposed to a federation server that is deployed in a farm using the SQL Server database is not significant. However, it is important to consider that when you use WID for the farm, each federation server in that farm must store, manage, and maintain replication changes for its local copy of the AD FS configuration database while also continuing to provide the normal operations that the Federation Service requires.

In comparison, federation servers that are deployed in a farm that uses the SQL Server database do not necessarily contain a local instance of the AD FS configuration database. Therefore, they may make slightly fewer demands on hardware resources.

Verifying that your production environment can support an AD FS deployment

In addition to the federation servers that you will deploy, and depending on how your existing production environment is set up, the following additional servers may be required to provide the necessary infrastructure to support your new AD FS deployment:

- Active Directory domain controller
- Certification authority (CA)
- Web server to host federation metadata
- Network load balancing (NLB)

Stand-Alone Federation Server Using WID

A stand-alone federation server in Active Directory Federation Services (AD FS) consists of a single server that hosts a Federation Service configured to use the Windows Internal Database (WID). This AD FS topology is for test labs. We do not recommend it for production environments because it has a limit of only one federation server, and it cannot be used to scale up to more servers.

If you want to add additional federation servers to your test lab, you must rebuild the Federation Service from scratch by deploying any of the other topologies mentioned later in this section. Therefore, we recommend that you use this topology for a test lab or a proof-of-concept environment in your private testing network in which a single federation server is adequate, as shown in the following illustration.

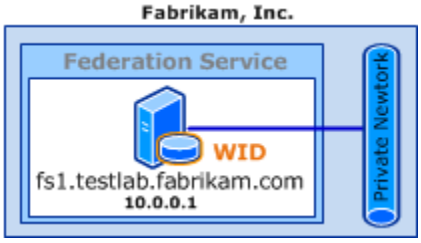

Test lab considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this topology for test lab environments.

Who should use this topology?

- Information technology (IT) professionals or IT architects who want to evaluate or develop a proof of concept for this technology

What are the benefits of using this topology?

- Easy to set up in a test lab environment

What are the limitations of using this topology?

- Only one federation server per Federation Service (no capability to scale up to a farm)
- Not redundant (only a single instance of the AD FS configuration database exists)

Federation Server Farm Using WID

The default topology for Active Directory Federation Services (AD FS) is a federation server farm, using the Windows Internal Database (WID), that consists of up to five federation servers hosting your organization's Federation Service. In this topology, AD FS uses WID as the store for the AD FS configuration database for all federation servers that are joined to that farm. The farm replicates and maintains the Federation Service data in the configuration database across each server in the farm.

The act of creating the first federation server in a farm also creates a new Federation Service. When you use WID for the AD FS configuration database, the first federation server that you create in the farm is referred to as the *primary federation server*. This means that this computer is configured with a read/write copy of the AD FS configuration database.

All other federation servers that you configure for this farm are referred to as *secondary federation servers* because they must replicate any changes that are made on the primary federation server to the read-only copies of the AD FS configuration database that they store locally.

Note

We recommend the use of at least two federation servers in a load-balanced configuration.

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Organizations with 100 or fewer configured trust relationships that need to provide their internal users (logged on to computers that are physically connected to the corporate network) with single sign-on (SSO) access to federated applications or services
- Organizations that want to provide their internal users with SSO access to Microsoft Online Services or Microsoft Office 365
- Smaller organizations that require redundant, scalable services

Note

Organizations with larger databases should consider using the Federation Server Farm Using SQL Server deployment topology, which is described later in this section. Organizations with users who log in from outside the network should consider using either the Federation Server Farm Using WID and Proxies topology or the Federation Server Farm Using SQL Server topology.

What are the benefits of using this topology?

- Provides SSO access to internal users
- Data and Federation Service redundancy (each federation server replicates changes to other federation servers in the same farm)
- The farm can be scaled out by adding up to five federation servers
- WID is included with Windows; therefore, no need to purchase SQL Server

What are the limitations of using this topology?

- A WID farm has a limit of five federation servers. For more information, see AD FS Deployment Topology Considerations.
- A WID farm does not support token replay detection or artifact resolution (part of the Security Assertion Markup Language (SAML) protocol).

Server placement and network layout recommendations

When you are ready to start deploying this topology in your network, you should plan on placing all of the federation servers in your corporate network behind a Network Load Balancing (NLB) host that can be configured for an NLB cluster with a dedicated cluster Domain Name System (DNS) name and cluster IP address.

Note
This cluster DNS name must match the Federation Service name, for example, fs.fabrikam.com.

The NLB host can use the settings that are defined in this NLB cluster to allocate client requests to the individual federation servers. The following illustration shows how the fictional Fabrikam, Inc., company sets up the first phase of its deployment using a two-computer federation server farm (fs1 and fs2) with WID and the positioning of a DNS server and a single NLB host that is wired to the corporate network.

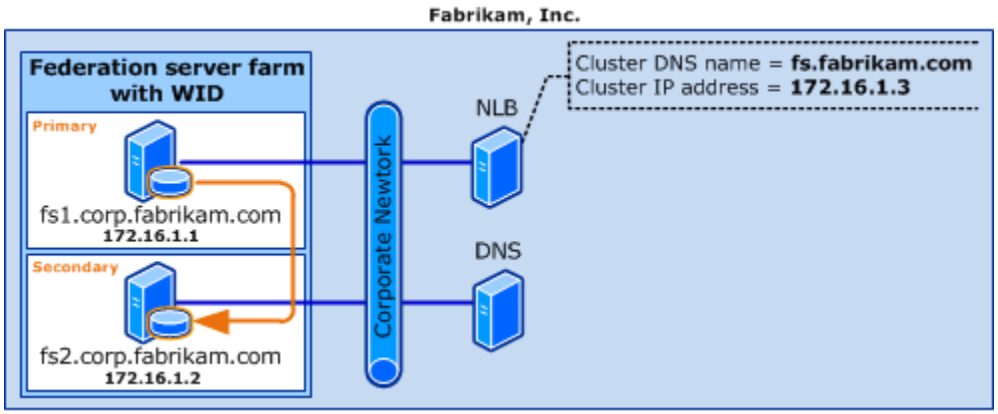

Note

If there is a failure on this single NLB host, users will not be able to access federated applications or services. Add additional NLB hosts if your business requirements do not allow having a single point of failure.

For more information about how to configure your networking environment for use with federation servers, see Name Resolution Requirements for Federation Servers in the AD FS Design Guide.

Federation Server Farm Using WID and Proxies

This deployment topology for Active Directory Federation Services (AD FS) is identical to the federation server farm with Windows Internal Database (WID) topology, but it adds federation server proxies to the perimeter network to support external users. The federation server proxies redirect client authentication requests that come from outside your corporate network to the federation server farm.

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Organizations with 100 or fewer configured trust relationships that need to provide both their internal users and external users (who are logged on to computers that are physically located outside the corporate network) with single sign-on (SSO) access to federated applications or services
- Organizations that need to provide both their internal users and external users with SSO access to Microsoft Office 365
- Smaller organizations that have external users and require redundant, scalable services

What are the benefits of using this topology?

- The same benefits as listed for the Federation Server Farm Using WID topology, plus the benefit of providing additional access for external users

What are the limitations of using this topology?

- The same limitations as listed for the Federation Server Farm Using WID topology

Server placement and network layout recommendations

To deploy this topology, in addition to adding two federation server proxies, you must make sure that your perimeter network can also provide access to a Domain Name System (DNS) server and to a second Network Load Balancing (NLB) host. The second NLB host must be configured with an NLB cluster that uses an Internet-accessible cluster IP address, and it must use the same cluster DNS name setting as the previous NLB cluster that you configured on the corporate network (fs.fabrikam.com). The federation server proxies should also be configured with Internet-accessible IP addresses.

The following illustration shows the existing federation server farm with WID topology that was described previously and how the fictional Fabrikam, Inc., company provides access to a perimeter DNS server, adds a second NLB host with the same cluster DNS name (fs.fabrikam.com), and adds two federation server proxies (fsp1 and fsp2) to the perimeter network.

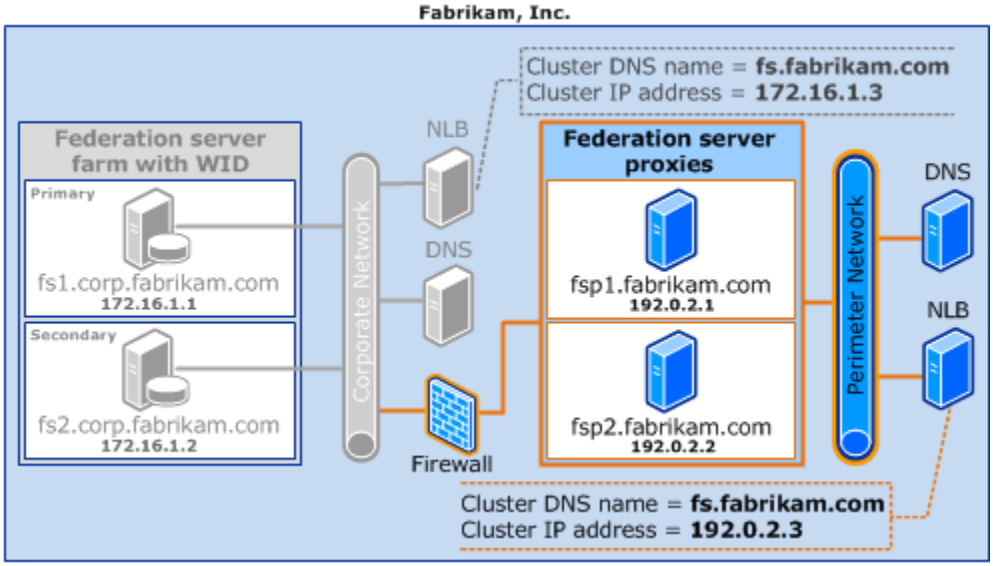

For more information about how to configure your networking environment for use with federation servers or federation server proxies, see either Name Resolution Requirements for Federation Servers or Name Resolution Requirements for Federation Server Proxies.

Federation Server Farm Using SQL Server

This topology for Active Directory Federation Services (AD FS) differs from the federation server farm using Windows Internal Database (WID) deployment topology in that it does not replicate the data to each federation server in the farm. Instead, all federation servers in the farm can read and

write data into a common database that is stored on a server running Microsoft SQL Server that is located in the corporate network.

Deployment considerations

This section describes various considerations about the intended audience, benefits, and limitations that are associated with this deployment topology.

Who should use this topology?

- Large organizations with more than 100 trust relationships that need to provide both their internal users and external users with single sign-on (SSO) access to federated application or services
- Organizations that already use SQL Server and want to take advantage of their existing tools and expertise

What are the benefits of using this topology?

- Support for larger numbers of trust relationships (more than 100)
- Support for token replay detection (a security feature) and artifact resolution (part of the Security Assertion Markup Language (SAML) 2.0 protocol)
- Support for the full benefits of SQL Server, such as database mirroring, failover clustering, reporting, and management tools

What are the limitations of using this topology?

- This topology does not provide database redundancy by default. Although a federation server farm with WID topology automatically replicates the WID database on each federation server in the farm, the federation server farm with SQL Server topology contains only one copy of the database

Note

SQL Server supports many different data and application redundancy options including failover clustering, database mirroring, and several different types of SQL Server replication.

The Microsoft Information Technology (IT) department uses SQL Server database mirroring in high-safety (synchronous) mode and failover clustering to provide high-availability support for the SQL Server instance. SQL Server transactional (peer-to-peer) and merge replication have not been tested by the AD FS product team at Microsoft. For more information about SQL Server, see High Availability Solutions Overview or Selecting the Appropriate Type of Replication.

Supported SQL Server Versions

The following SQL server versions are supported with AD FS installed with Windows Server 2012:

- SQL Server 2008 / R2
- SQL Server 2012

Server placement and network layout recommendations

Similar to the federation server farm with WID topology, all of the federation servers in the farm are configured to use one cluster Domain Name System (DNS) name (which represents the Federation Service name) and one cluster IP address as part of the Network Load Balancing (NLB) cluster configuration. This helps the NLB host allocate client requests to the individual federation servers. Federation server proxies can be used to proxy client requests to the federation server farm.

The following illustration shows how the fictional Contoso Pharmaceuticals company deployed its federation server farm with SQL Server topology in the corporate network. It also shows how that company configured the perimeter network with access to a DNS server, an additional NLB host that uses the same cluster DNS name (fs.contoso.com) that is used on the corporate network NLB cluster, and with two federation server proxies (fsp1 and fsp2).

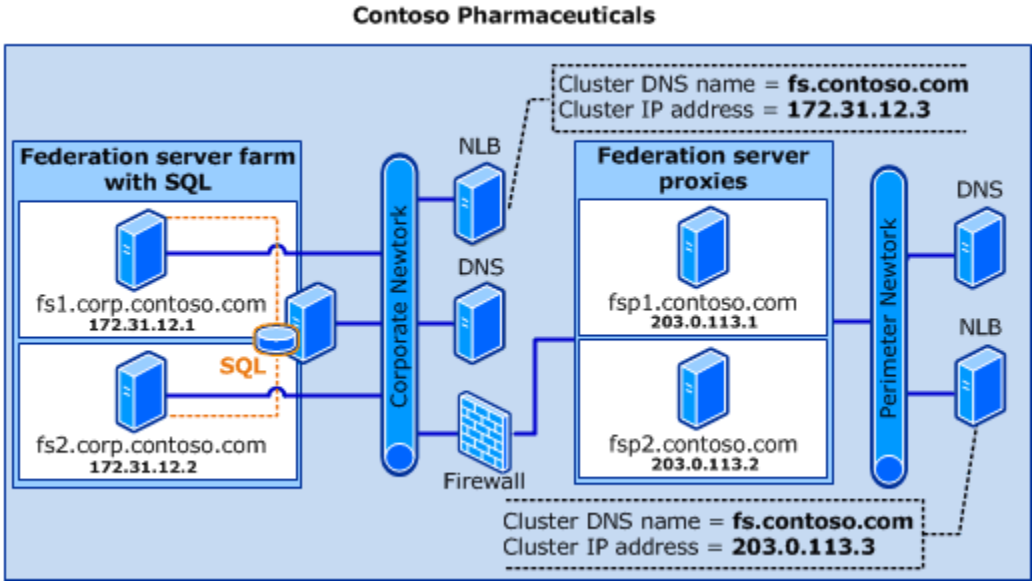

For more information about how to configure your networking environment for use with federation servers or federation server proxies, see either Name Resolution Requirements for Federation Servers or Name Resolution Requirements for Federation Server Proxies.

Planning Your Deployment

When you plan for cross-organizational (federation-based) collaboration using Active Directory Federation Services (AD FS), first determine if your organization will host a Web resource to be

accessed by other organizations across the Internet or if you will provide access to the Web resource for employees in your organization. This determination affects how you deploy AD FS, and it is fundamental in the planning of your AD FS infrastructure.

Note
Make sure that the role that organization plays in the federation agreement is clearly understood by all parties.

For the Federated Web SSO Design, AD FS uses terms such as *account partner* (also referred to as *identity provider* in the AD FS Management snap-in) and *resource partner* (also referred to as *relying party* in the AD FS Management snap-in) to help differentiate the organization that hosts the accounts (the account partner) from the organization that hosts the Web-based resources (the resource partner).

In the Web SSO Design, the organization acts in both the account partner and resource partner roles because it is providing its users with access to its applications.

The following topics explain some of the AD FS partner organization concepts. They also contain links to topics in the AD FS Deployment Guide that contain information about setting up and configuring account partner organizations and resource partner organizations based on your AD FS deployment goals.

In this section

- Best Practices for Secure Planning and Deployment of AD FS
- Planning for Interoperability with AD FS 1.x
- When to Use Identity Delegation
- Deploying AD FS in the Account Partner Organization
- Deploying AD FS in the Resource Partner Organization

Using AD DS Claims with AD FS

You can enable richer access control for federated applications by using Active Directory Domain Services (AD DS)-issued user and device claims together with Active Directory Federation Services (AD FS).

About Dynamic Access Control

In Windows Server® 2012, the Dynamic Access Control feature enables organizations to grant access to files based on user claims (which are sourced by user account attributes) and device claims (which are sourced by computer account attributes) that are issued by Active Directory Domain Services (AD DS). AD DS issued claims are integrated into Windows integrated authentication through the Kerberos authentication protocol.

For more information about Dynamic Access Control, see Dynamic Access Control Content Roadmap.

What's New in AD FS?

As an extension to the Dynamic Access Control scenario, AD FS in Windows Server 2012 can now:

- Access computer account attributes in addition to user account attributes from within AD DS. In previous versions of AD FS, the Federation Service could not access computer account attributes at all from AD DS.
- Consume AD DS issued user or device claims that reside in a Kerberos authentication ticket. In previous versions of AD FS, the claims engine was able to read user and group security IDs (SIDs) from Kerberos but was not able to read any claims information contained within a Kerberos ticket.
- Transform AD DS issued user or device claims into SAML tokens that relying applications can use to perform richer access control.

Benefits of Using AD DS Claims with AD FS

These AD DS issued claims can be inserted into Kerberos authentication tickets and used with AD FS to provide the following benefits:

- Organizations that require richer access control policies can enable claims-based access to applications and resources by using AD DS issued claims that are based on the attribute values stored in AD DS for a given user or computer account. This can help administrators to reduce additional overhead associated with creating and managing:
  - AD DS security groups that would otherwise be used for controlling access to applications and resources that are accessible via Windows Integrated authentication.
  - Forest trusts that would otherwise be used for controlling access to Business-to-Business (B2B) / Internet accessible applications and resources.
- Organizations can now prevent unauthorized access to network resources from client computers based on whether a specific computer account attribute value stored in AD DS (for example, a computer's DNS name) matches the access control policy of the resource (for example, a file server that has been ACLd with claims) or the relying party policy (for example, a claims-aware Web application). This can help administrators to set finer access control policies for resources or applications that are:
  - Only accessible via Windows Integrated authentication.
  - Internet accessible via AD FS authentication mechanisms. AD FS can be used to transform AD DS issued device claims into AD FS claims that can be encapsulated into SAML tokens which can be consumed by an Internet accessible resource or relying party application.

Differences Between AD DS and AD FS Issued Claims

There are two differentiating factors that are important to understand about claims that are issued from AD DS vs. AD FS. These differences include:

- AD DS can only issue claims that are encapsulated in Kerberos tickets, not SAML tokens. For more information about how AD DS issues claims, see Dynamic Access Control Content Roadmap.
- AD FS can only issue claims that are encapsulated in SAML tokens, not Kerberos tickets. For more information about how AD FS issues claims, see The Role of the Claims Engine.

How AD DS Issued Claims Work with AD FS

AD DS issued claims can be used with AD FS to access both user and device claims directly from the user's authentication context, rather than making a separate LDAP call to Active Directory. The following illustration and corresponding steps discusses how this process works in more detail to enable claims-based access control for the Dynamic Access Control scenario.

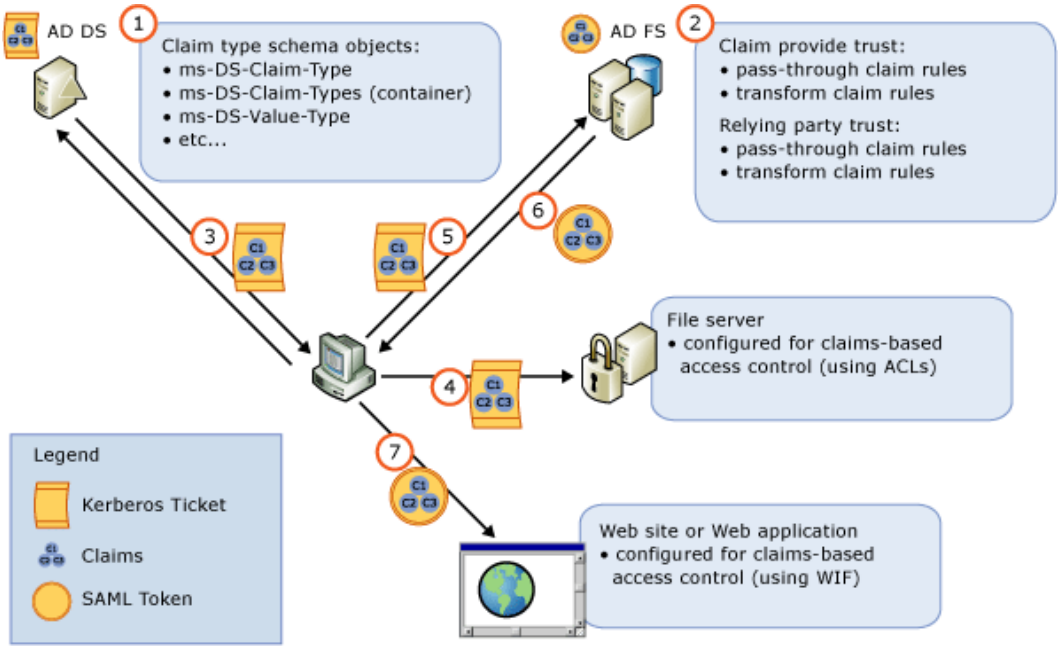

1. An AD DS administrator uses the Active Directory Administrative Center console or PowerShell cmdlets to enables specific claim type objects in the AD DS schema.
2. An AD FS administrator uses the AD FS Management console to create and configure the claims provider and relying party trusts with either pass-through or transform claim rules.
3. A Windows client attempts to access the network. As part of the Kerberos authentication process, the client presents its user and computer ticket-granting ticket (TGT) which does not yet contain any claims, to the domain controller. The domain controller then looks in AD DS for enabled claim types, and includes any resulting claims in the returned Kerberos ticket.

4. When the user/client attempts to access a file resource that is ACLd to require the claims, they can access the resource because the compound ID that was surfaced from Kerberos has these claims.
5. When the same client attempts to access a Web site or Web application that is configured for AD FS authentication, the user is redirected to an AD FS federation server that is configured for Windows integrated authentication. The client sends a request to the domain controller using Kerberos. The domain controller issues a Kerberos ticket containing the requested claims which the client can then present to the federation server.
6. Based on the way the claims rules have been configured on the claims provider and relying party trusts that the administrator configured previously, AD FS reads the claims from the Kerberos ticket and includes them in a SAML token that it issues for the client.
7. The client receives the SAML token containing the correct claims and is then redirected to the website.

For more information about how to create the claim rules required for AD DS issued claims to work with AD FS, see Create a Rule to Transform an Incoming Claim.

Best Practices for Secure Planning and Deployment of AD FS

This topic provides best-practice information to help you plan and evaluate security when you design your Active Directory Federation Services (AD FS) deployment. This topic is a starting point for reviewing and assessing considerations that affect the overall security of your use of AD FS. The information in this topic is meant to compliment and extend your existing security planning and other design best practices.

Core security best practices for AD FS

The following core best practices are common to all AD FS installations where you want to improve or extend the security of your design or deployment:

- **Use the Security Configuration Wizard to apply AD FS-specific security best practices to federation servers and federation server proxy computers**

  The Security Configuration Wizard (SCW) is a tool that comes preinstalled on all Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012 computers. You can use it to apply security best practices that can help reduce the attack surface for a server, based on the server roles that you are installing.

  When you install AD FS, the setup program creates role extension files that you can use with the SCW to create a security policy that will apply to the specific AD FS server role (either federation server or federation server proxy) that you choose during setup.

  Each role extension file that is installed represents the type of role and subrole for which each computer is configured. The following role extension files are installed in the C:\Windows\ADFS\Scw directory:

- Farm.xml
- SQLFarm.xml
- StandAlone.xml
- Proxy.xml (This file is present only if you configured the computer in the federation server proxy role.)

To apply the AD FS role extensions in the SCW, complete the following steps in order:

a. Install AD FS and choose the appropriate server role for that computer. For more information, see Install the Federation Service Proxy Role Service in the AD FS Deployment Guide.

b. Register the appropriate role extension file using the Scwcmd command-line tool. See the following table for details about using this tool in the role for which your computer is configured.

c. Verify that the command has completed successfully by examining the SCWRegister_log.xml file, which is located in the Windows\security\Msscw\Logs directory.

You must perform all these steps on each federation server or federation server proxy computer to which you want to apply AD FS–based SCW security policies.

The following table explains how to register the appropriate SCW role extension, based on the AD FS server role that you chose on the computer where you installed AD FS.

AD FS server role	AD FS configuration database used	Type the following command at a command prompt:
Stand-alone federation server	Windows Internal Database	`scwcmd register /kbname:ADFS2Standalone /kbfile:"Windows\ADFS\scw\StandAlone.xml"`
Farm-joined federation server	Windows Internal Database	`scwcmd register /kbname:ADFS2Standalone /kbfile:"Windows\ADFS\scw\Farm.xml"`
Farm-joined federation server	SQL Server	`scwcmd register /kbname:ADFS2Standalone /kbfile:"Windows\ADFS\scw\SQLFarm.xml"`
Federation server proxy	N/A	`scwcmd register /kbname:ADFS2Standalone /kbfile:"Windows\ADFS\scw\Proxy.xml"`

For more information about the databases that you can use with AD FS, see The Role of the AD FS Configuration Database.

- **Use token replay detection in situations in which security is a very important concern, for example, when kiosks are used.**
  Token replay detection is a feature of AD FS that ensures that any attempt to replay a token request that is made to the Federation Service is detected and the request is discarded. Token replay detection is enabled by default. It works for both the WS-Federation passive profile and the Security Assertion Markup Language (SAML) WebSSO profile by ensuring that the same token is never used more than once.

When the Federation Service starts, it begins to build a cache of any token requests that it fulfills. Over time, as subsequent token requests are added to the cache, the ability to detect any attempts to replay a token request multiple times increases for the Federation Service. If you disable token replay detection and later choose to enable it again, remember that the Federation Service will still accept tokens for a period of time that may have been used previously, until the replay cache has been allowed enough time to rebuild its contents. For more information, see The Role of the AD FS Configuration Database.

- **Use token encryption, especially if you are using supporting SAML artifact resolution.**

  Encryption of tokens is strongly advised to increase security and protection against potential man-in-the-middle (MITM) attacks that might be tried against your AD FS deployment. Using use encryption might have a slight impact on throughout but in general, it should not be usually noticed and in many deployments the benefits for greater security exceed any cost in terms of server performance.

  To enable token encryption, first set add an encryption certificate for your relying party trusts. You can configure an encryption certificate either when creating a relying party trust or later. To add an encryption certificate later to an existing relying party trust, you can set a certificate for use on the **Encryption** tab within trust properties while using the AD FS snap-in. To specify a certificate for an existing trust using the AD FS cmdlets, use the EncryptionCertificate parameter of either the **Set-ClaimsProviderTrust** or **Set-RelyingPartyTrust** cmdlets. To set a certificate for the Federation Service to use when decrypting tokens, use the **Set-ADFSCertificate** cmdlet and specify "`Token-Encryption`" for the *CertificateType* parameter. Enabling and disabling encryption for specific relying party trust can be done by using the *EncryptClaims* parameter of the **Set-RelyingPartyTrust** cmdlet.

- **Utilize extended protection for authentication**

  To help secure your deployments, you can set and use the extended protection for authentication feature with AD FS. This setting specifies the level of extended protection for authentication supported by a federation server.

  Extended protection for authentication helps protect against man-in-the-middle (MITM) attacks, in which an attacker intercepts client credentials and forwards them to a server. Protection against such attacks is made possible through a Channel Binding Token (CBT) which can be either required, allowed, or not required by the server when it establishes communications with clients.

  To enable the extended protection feature, use the **ExtendedProtectionTokenCheck** parameter on the **Set-ADFSProperties** cmdlet. Possible values for this setting and the level of security that the values provide are described in the following table.

Parameter Value	Security level	Protection setting
Require	Server is fully hardened.	Extended protection is enforced and always required.
Allow	Server is partially hardened.	Extended protection is enforced where systems involved have been patched to support it.
None	Server is vulnerable.	Extended protection is not enforced.

- **If you are using logging and tracing, ensure the privacy of any sensitive information.**

  AD FS does not, by default, expose or track personally identifiable information (PII) directly as part of the Federation Service or normal operations. When event logging and debug trace logging are enabled in AD FS, however, depending on the claims policy that you configure some claims types and their associated values might contain PII that might be logged in the AD FS event or tracing logs.

  Therefore, enforcing access control on the AD FS configuration and its log files is strongly advised. If you do not want this kind of information to be visible, you should disable loggin, or filter out any PII or sensitive data in your logs before you share them with others.

  The following tips can help you prevent the content of a log file from being exposed unintentionally:

  - Ensure that the AD FS event log and trace log files are protected by access control lists (ACL) that limit access to only those trusted administrators who require access to them.
  - Do not copy or archive log files using file extensions or paths that can be easily served using a Web request. For example, the .xml file name extension is not a safe choice. You can check the Internet Information Services (IIS) administration guide to see a list of extensions that can be served.
  - If you revise the path to the log file, be sure to specify an absolute path for the log file location, which should be outside of the Web host virtual root (vroot) public directory to prevent it from being accessed by an external party using a Web browser.

SQL Server–specific security best practices for AD FS

The following security best practices are specific to the use of Microsoft SQL Server® or Windows Internal Database (WID) when these database technologies are used to manage data in AD FS design and deployment.

Note

These recommendations are meant to extend, but not replace, SQL Server product security guidance. For more information about planning a secure SQL Server installation, see Security Considerations for a Secure SQL Installation (http://go.microsoft.com/fwlink/?LinkID=139831).

- **Always deploy SQL Server behind a firewall in a physically secure network environment.**

  A SQL Server installation should never be exposed directly to the Internet. Only computers that are inside your datacenter should be able to reach your SQL server installation that supports AD FS. For more information, see Security Best Practices Checklist (http://go.microsoft.com/fwlink/?LinkID=189229).

- **Run SQL Server under a service account instead of using the built-in default system service accounts.**

  By default, SQL Server is often installed and configured to use one of the supported built-in system accounts, such as the LocalSystem or NetworkService accounts. To enhance the security of your SQL Server installation for AD FS, wherever possible use a separate service account for accessing your SQL Server service and enable Kerberos authentication by registering the security principal name (SPN) of this account in your Active Directory deployment. This enables mutual authentication between client and server. Without SPN registration of a separate service account, SQL Server will use NTLM for Windows-based authentication, where only the client is authenticated.

- **Minimize the surface area of SQL Server.**

  Enable only those SQL Server endpoints that are necessary. By default, SQL Server provides a single built-in TCP endpoint that cannot be removed. For AD FS, you should enable this TCP endpoint for Kerberos authentication. To review the current TCP endpoints to see if additional user-defined TCP ports are added to a SQL installation, you can use the "SELECT * FROM sys.tcp_endpoints" query statement in a Transact-SQL (T-SQL) session. For more information about SQL Server endpoint configuration, see How To: Configure the Database Engine to Listen on Multiple TCP Ports (http://go.microsoft.com/fwlink/?LinkID=189231).

- **Avoid using SQL-based authentication.**

  To avoid having to transfer passwords as clear text over your network or storing passwords in configuration settings, use Windows authentication only with your SQL Server installation. SQL Server authentication is a legacy authentication mode. Storing Structured Query Language (SQL) login credentials (SQL user names and passwords) when you are using SQL Server authentication is not recommended. For more information, see Authentication Modes (http://go.microsoft.com/fwlink/?LinkID=189232).

- **Evaluate the need for additional channel security in your SQL installation carefully.**

  Even with Kerberos authentication in effect, the SQL Server Security Support Provider Interface (SSPI) does not provide channel-level security. However, for installations in which servers are securely located on a firewall-protected network, encrypting SQL communications may not be necessary.

Although encryption is a valuable tool to help ensure security, it should not be considered for all data or connections. When you are deciding whether to implement encryption, consider how users will access data. If users access data over a public network, data encryption might be required to increase security. However, if all access of SQL data by AD FS involves a secure intranet configuration, encryption might not be required. Any use of encryption should also include a maintenance strategy for passwords, keys, and certificates.

If there is a concern that any SQL data might be seen or tampered with over your network, use Internet Protocol security (IPsec) or Secure Sockets Layer (SSL) to help secure your SQL connections. However, this might have a negative effect on SQL Server performance, which might affect or limit AD FS performance in some situations. For example, AD FS performance in token issuance might degrade when attribute lookups from a SQL-based attribute store are critical for token issuance. You can better eliminate a SQL tampering threat by having a strong perimeter security configuration. For example, a better solution for securing your SQL Server installation is to ensure that it remains inaccessible for Internet users and computers and that it remains accessible only by users or computers within your datacenter environment.

For more information, see Encrypting Connections to SQL Server (http://go.microsoft.com/fwlink/?LinkID=189234) or SQL Server Encryption (http://go.microsoft.com/fwlink/?LinkID=189233).

- **Configure securely designed access by using stored procedures to perform all SQL-based lookups by AD FS of SQL-stored data.**

  To provide better service and data isolation, you can create stored procedures for all attribute store lookup commands. You can create a database role to which you then grant permission to run the stored procedures. Assign the service identity of the AD FS Windows service to this database role. The AD FS Windows service should not be able to run any other SQL statement, other than the appropriate stored procedures that are used for attribute lookup. Locking down access to the SQL Server database in this way reduces the risk of an elevation-of-privilege attack.

Planning for Interoperability with AD FS 1.x

Active Directory Federation Services (AD FS) federation servers running Windows Server® 2012 can interoperate with both an AD FS 1.0 (installed with Windows Server 2003 R2) Federation Service and an AD FS 1.1 (installed with Windows Server 2008 or Windows Server 2008 R2) Federation Service. Any of the following interoperability combinations are supported:

- Any AD FS 1.*x* Federation Service can send a claim that can be consumed by an AD FS Federation Service in Windows Server 2012. For more information, see Checklist: Configuring AD FS to Consume Claims from AD FS 1.x.
- Any AD FS Federation Service in Windows Server 2012 can send an AD FS 1.*x*-compatible claim that can be consumed by an AD FS 1.*x* Federation Service. For more information, see Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Federation Service.

- Any AD FS Federation Service in Windows Server 2012 can send an AD FS 1.*x*-compatible claim that can be consumed by one or more Web servers running the AD FS 1.*x* claims-aware Web agent. For more information, see Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Claims-Aware Web Agent.

Note
AD FS does not support or interoperate with the AD FS 1.*x* Windows NT token–based Web agent.

An AD FS 1.*x*-compatible claim is a claim that can be sent by an AD FS Federation Service in Windows Server 2012 and understood by an AD FS 1.*x* Federation Service. So that an AD FS 1.*x* Federation Service can consume the claims that an AD FS Federation Service sends, a Name Identifier (ID) claim type must be sent.

Understanding the Name ID claim type

The Name ID claim type is the equivalent of the identity claim type that AD FS 1.*x* uses. It must be used whenever you want to interoperate with AD FS 1.*x*. The Name ID claim type enables either an AD FS 1.*x* Federation Service or the AD FS 1.*x* claims-aware Web agent to consume claims that AD FS in Windows Server 2012 sends, as long as these claims are sent in one of the Name ID formats in the following table.

Name ID format	Corresponding URI
AD FS 1.*x* Email Address	http://schemas.xmlsoap.org/claims/EmailAddress
AD FS 1.*x* Email UPN	http://schemas.xmlsoap.org/claims/UPN
Common Name	http://schemas.xmlsoap.org/claims/CommonName
Group	http://schemas.xmlsoap.org/claims/Group

Only one Name ID claim in the appropriate format must be sent. When that criterion is satisfied, many other claims may be sent as well, assuming that they conform to the restrictions described in the table.

Note
An AD FS 1.*x* Federation Service can interpret only incoming claim types that begin with the Uniform Resource Identifier (URI) of http://schemas.xmlsoap.org/claims/.

When to Use Identity Delegation

What is identity delegation?

Identity delegation is a feature of Active Directory Federation Services (AD FS) that allows administrator-specified accounts to impersonate users. The account that impersonates the user is called the *delegate*. This delegation capability is critical for many distributed applications for which there is a series of access control checks that must be made sequentially for each application, database, or service that is in the authorization chain for the originating request. Many real-world scenarios exist in which a Web application “front end” must retrieve data from a more secure “back end”, such as a Web service that is connected to a Microsoft SQL Server database.

For example, an existing parts-ordering Web site can be enhanced programmatically so that it allows partner organizations to view their own purchase history and account status. For security reasons, all partner financial data is stored in a secure database on a dedicated Structured Query Language (SQL) server. In this situation, the code in the front-end application knows nothing about the partner organization’s financial data. Therefore, it must retrieve that data from another computer elsewhere on the network that hosts (in this case) the Web service for the parts database (the back end).

For this data-retrieval process to succeed, some succession of authorization “hand-shaking” must take place between the Web application and the Web service for the parts database, as shown in the following illustration.

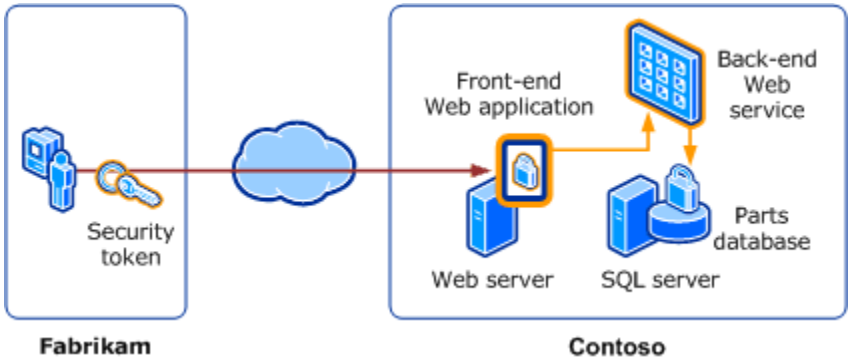

Because the original request was made to the Web server itself, which is likely to be located in a completely different organization from the organization of the user who is attempting to access the Web server, the security token that is sent along with the request does not meet the authorization criteria required to access any other computer besides the Web server. Therefore, the only way that the originating user request can be fulfilled is by placing an intermediate federation server in the resource partner organization to help with reissuing a security token that does have the appropriate access privileges.

How does identity delegation work?

Web applications in multitier application architectures often call Web services to access common data or functionality. It is important for these Web services to know the identity of the original user so that the service can make authorization decisions and facilitate auditing. In this case, the front-end Web application represents the user to the Web service as a delegate. AD FS facilitates this scenario by allowing Active Directory accounts to act as a user to another relying party. An identity delegation scenario is shown in the following illustration.

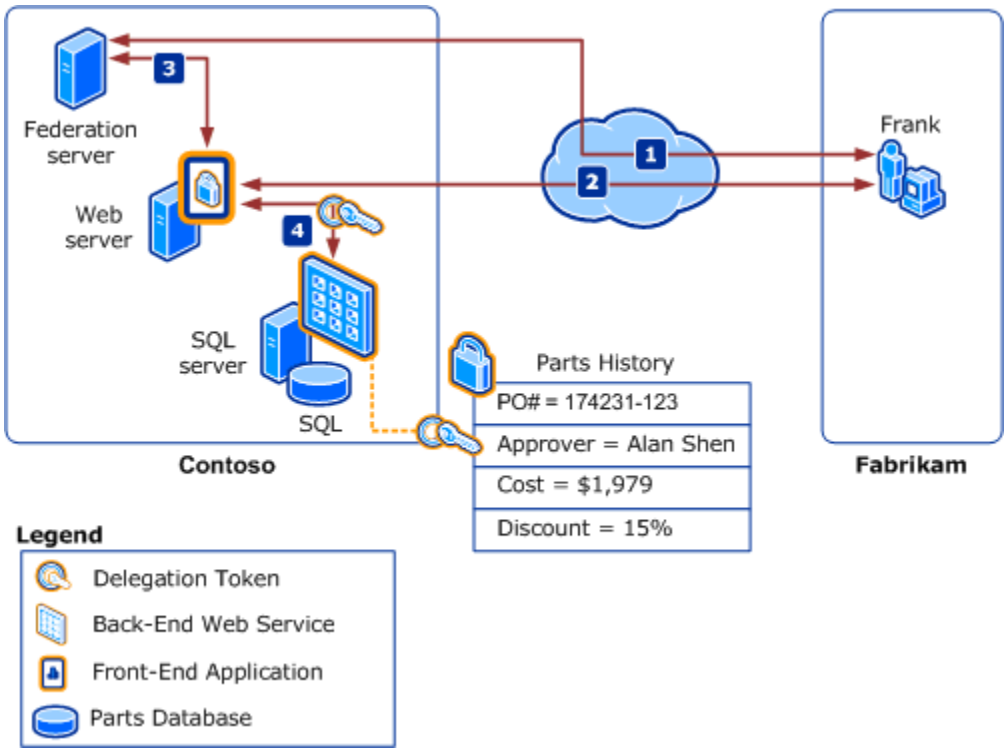

1. Frank attempts to access part-ordering history from a Web application in another organization. His client computer requests and receives a token from AD FS for the front-end part-ordering Web application.
2. The client computer sends a request to the Web application, including the token obtained in step 1, to prove the client's identity.
3. The Web application needs to communicate with the Web service to complete its transaction for the client. The Web application contacts AD FS to obtain a delegation token to interact with the Web service. Delegation tokens are security tokens that are issued to a delegate to act as a user. AD FS returns a delegation token with claims about the client, targeted for the Web service.
4. The Web application uses the token that was obtained from AD FS in step 3 to access the Web service that is acting as the client. Examining the delegation token, the Web service can determine that the Web application is acting as the client. The Web service executes its authorization policy, logs the request, and provides the needed parts history data that was originally requested by Frank to the Web application and therefore to Frank.

For a particular delegate, AD FS can limit the Web services for which the Web application may request a delegation token. The client computer does not have to have an Active Directory account for this operation to succeed. Finally, as noted previously, the Web service can easily determine the identity of the delegate that is acting as the user. This allows Web services to exhibit different behavior based on whether they are talking directly to the client computer or through a delegate.

Configuring AD FS for identity delegation

You can use the AD FS Management snap-in to configure AD FS for identity delegation whenever you need to facilitate the data retrieval process. After you configure it, AD FS can generate new security tokens that will include the authorization context that the back-end service may require before it can provide access to the protected data.

AD FS does not restrict which users can be impersonated. After you configure AD FS for identity delegation, it does the following:

- It determines which servers can be delegated the authority to request tokens to impersonate a user.
- It establishes and keeps separate both the identity context for the client account that is delegated and the server that acts as a delegate.

You can configure identity delegation by adding delegation authorization rules to a relying party trust in the AD FS Management snap-in. For more information about how to do this, see Checklist: Creating Claim Rules for a Relying Party Trust.

Configuring the front-end Web application for identity delegation

Developers have several options that they can use to appropriately program the Web front-end application or service to redirect delegation requests to an AD FS computer. For more information about how to customize a Web application to work with identity delegation, see the Windows Identity Foundation SDK.

Deploying AD FS in the Account Partner Organization

An account partner in Active Directory Federation Services (AD FS) represents the organization in the federation trust relationship that physically stores user accounts in a supported attribute store. For more information about which attribute stores are supported, see The Role of Attribute Stores.

The federation server in the account partner organization authenticates local users and creates security tokens that are used by the resource partner in making authorization decisions. Relying

parties such as Web sites and Web services are then able to easily register themselves with the federation server and consume issued tokens for authentication and access control.

In scenarios in which you need to provide your users with access to multiple federated applications or services—when each application or service is hosted by a different organization—you can configure the account partner federation server so that you can deploy multiple relying parties.

For more information about how to set up and configure an account partner organization, see Checklist: Configuring the Account Partner Organization.

In this section

- Review the Role of the Federation Server in the Account Partner
- Review the Role of the Federation Server Proxy in the Account Partner
- Prepare Client Computers in the Account Partner

Review the Role of the Federation Server in the Account Partner

A federation server in Active Directory Federation Services (AD FS) functions as a security token issuer. A federation server generates claims based on account values that reside in a local attribute store and packages them into security tokens so that users can seamlessly access Web-browser-based applications (using single sign-on (SSO)) that are hosted in a resource partner organization.

Note

When your users access federated applications by using a Web browser, a federation server automatically issues cookies to the users to maintain their logon status for that Web-browser-based application. These cookies include claims for the users. The cookies enable SSO capabilities so that the users do not have to enter credentials each time that they visit different Web-browser-based applications in the resource partner.

In the Web SSO design, organizations with a perimeter network that want Internet users to have access to applications must install a federation server proxy in the perimeter network. In the Federated Web SSO design, there must be at least one federation server installed in the corporate network of the account partner organization and at least one federation server installed in the corporate network of the resource partner organization.

Note

Before you can set up a federation server computer in the account partner organization, you must first join the computer to any domain in the Active Directory forest where the federation server will be used to authenticate users from that forest. For more information, see Checklist: Setting Up a Federation Server.

Review the Role of the Federation Server Proxy in the Account Partner

The primary role of the federation server proxy in the perimeter network of the account partner organization in Active Directory Federation Services (AD FS) is to collect authentication credentials from a client computer that logs on over the Internet and to pass those credentials to the federation server, which is located inside the corporate network of the account partner organization. The account for the client computer is stored in the account partner's attribute store.

A federation server proxy can also function in one or more of the following roles, depending on how you configure it to meet the needs of the account partner organization:

- Relay Security Tokens—The federation server issues a security token to the federation server proxy, which then relays the token to the client computer. The security token is used to provide access for that client computer to a specific relying party.
- Collect Credentials—The federation server proxy uses a default client logon Web form (clientlogon.aspx) to collect password-based credentials through forms-based authentication. However, you can customize this form to accept other supported types of authentication, such as Secure Sockets Layer (SSL) client authentication. For more information about how to customize this page, see Customizing Client Logon and Home Realm Discovery Pages (http://go.microsoft.com/fwlink/?LinkId=104275). A federation server proxy does not accept credentials through Windows Integrated Authentication.

To summarize, a federation server proxy in the account partner acts as a proxy for client logons to a federation server that is located in the corporate network. The federation server proxy also facilitates the distribution of security tokens to Internet clients that are destined for relying parties.

Caution

Exposing a federation server proxy on the account partner extranet will the client logon Web form accessible by anyone with Internet access. This can potentially leave your organization vulnerable to some password-based attacks, such as dictionary attacks or brute force attacks that can trigger account lockouts for user accounts that are stored in the corporate Active Directory Domain Services (AD DS).

Prepare Client Computers in the Account Partner

The easiest way for an administrator in an account partner organization to prepare client computers for access to Active Directory Federation Services (AD FS) federated applications is to use Group Policy. Group Policy provides a convenient way for you to push specific certificates and settings that are required for federation to all the client computers that will be used to access federated applications.

So that your client computers can seamlessly access federated applications without certificate prompts or trusted site–related prompts, we recommend that you first prepare each client computer before you deploy AD FS broadly in your organization. Consider using Group Policy to automatically:

- Configure Internet Explorer on each client computer to trust the account federation server.

  For more information, see Configure Client Computers to Trust the Account Federation Server.
- Install the appropriate account federation server, resource federation server, and Web server Secure Sockets Layer (SSL) certificates (or equivalent certificates that chain to a trusted root) on each client computer.

  For more information, see Distribute Certificates to Client Computers by Using Group Policy.

Deploying AD FS in the Resource Partner Organization

The resource partner organization in Active Directory Federation Services (AD FS) represents the organization whose Web servers may be protected by a resource-side federation server. The federation server at the resource partner uses the security tokens that are produced by the account partner to provide claims to the Web servers that are located in the resource partner.

In scenarios in which you need to provide access to federated services or applications to many different users—when some users reside in different organizations—you can configure the resource federation server so that you can deploy multiple account partners.

For more information about how to set up and configure a resource partner organization, see Checklist: Configuring the Resource Partner Organization.

In this section

- Review the Role of the Federation Server in the Resource Partner
- Review the Role of the Federation Server Proxy in the Resource Partner
- Determine Your Federated Application Strategy in the Resource Partner

Review the Role of the Federation Server in the Resource Partner

The federation server in the resource partner organization intercepts incoming security tokens that are sent by an account federation server, validates and signs them, and then issues its own security tokens that are destined for the Web-based application.

Note

When federated users use their Web browsers to access Web-based applications, the federation server in the resource partner organization builds a new authentication cookie and writes it to the browser. This cookie enables single-sign-on (SSO) capabilities so that users do not have to log on again at the federation server in the account partner when the users attempt to access different Web-based applications in the resource partner.

In the Web SSO design, at least one federation server must be installed in the perimeter network. In the Federated Web SSO design, there must be at least one federation server installed in the corporate network of the account partner organization and at least one federation server installed in the corporate network of the resource partner organization.

Note

Before you can set up a federation server computer in the resource partner organization, you must first join the computer to any Active Directory domain in the resource partner organization. For more information, see Checklist: Setting Up a Federation Server.

Review the Role of the Federation Server Proxy in the Resource Partner

A federation server proxy in Active Directory Federation Services (AD FS) can function in one or more of the following roles, depending on how you configure the server to meet the needs of the resource partner organization:

- **Account partner discovery**: An Internet client computer must identify which account partner will authenticate it. The client finds the account partner by using an account partner discovery Web form (discoverclientrealm.aspx), which is stored on the federation server proxy in the resource partner. If more than one account partner is configured in the AD FS Management snap-in, a drop-down menu appears to the client with all the available account partners that are visible to Internet client computers that access the account partner discovery Web form. You can change how the account partner discovery Web form is presented to client computers by customizing the discoverclientrealm.aspx file.
- **Security token redirection**: The federation server proxy in the account partner sends the security tokens to the resource partner. The resource federation server proxy accepts these tokens and passes them on to the federation server in the resource partner. The resource federation server then issues a security token that is bound for a specific resource Web server. The resource federation server proxy then redirects the token to the client.

To summarize, a resource federation server proxy facilitates the federated logon process by redirecting client computers to a federation server that can authenticate the clients. A resource federation server proxy also acts as a proxy for client security tokens to resource federation servers.

Note

When it is necessary to help reduce the amount of hardware and the number of required certificates, the federation server proxy can be located on the same computer as the Web server.

Determine Your Federated Application Strategy in the Resource Partner

An important part of designing a new Active Directory Federation Services (AD FS) infrastructure in the resource partner organization is determining your full set of applications and services that will be used to participate in the federation and which account partners will be the recipients of those resources. Before you design a federated application and services strategy, consider the following questions:

- Will you be enabling and deploying an ASP.NET application or a Windows Communication Foundation (WCF) service for federation?
- Will users on your corporate network require access to the federated application or service through Windows Integrated Authentication?
- Will the federated application or service be used by users in your perimeter network? If so, will Windows Integrated Authentication be required?
- Are all of the Web servers that host federated applications running a Windows Server operating system and Internet Information Services (IIS)?
- Who will the federated application or service provide resources for?

Answering these questions will help you plan a solid AD FS design. It will also assist you in creating a federated application and services strategy that is cost effective and resource efficient. For more information about designing the most appropriate federated application and services strategy for your organization, see the following topics in this guide:

- Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services
- Provide Your Active Directory Users Access to the Applications and Services of Other Organizations
- Provide Users in Another Organization Access to Your Claims-Aware Applications and Services

For more information about how to create a claims-aware ASP.NET application or WCF service, see Windows Identity Foundation SDK.

Planning Federation Server Placement

The most critical component of an Active Directory Federation Services (AD FS) deployment is the federation server. Therefore, it is important that you plan your federation server placement strategy carefully, including when and where to deploy federation servers. The information in the following topics can help you determine when and where to create a federation server or

federation server farm and whether to use that federation server in the account partner role, the resource partner role, or both:

- Review the Role of the Federation Server in the Account Partner
- Review the Role of the Federation Server in the Resource Partner
- When to Create a Federation Server
- Where to Place a Federation Server
- When to Create a Federation Server Farm
- Certificate Requirements for Federation Servers
- Name Resolution Requirements for Federation Servers

Note
Although this information might help with your placement planning for federation servers, it does not explain how to determine the proper number of federation servers and the hardware requirements for each AD FS design.

For examples of how a federation server can be placed in any of the two primary AD FS design scenarios, see Mapping Your Deployment Goals to an AD FS Design.

When to Create a Federation Server

When you create a federation serverin Active Directory Federation Services (AD FS), you provide a means by which your organization can:

- Engage in Web single-sign-on (SSO)–based communication with another organization (that also has at least one federation server) and, when necessary, with the employees in your own organization (who need access over the Internet).
- Enable front end services to impersonate users to infrastructure services using identity delegation. For more information, see When to Use Identity Delegation.

The following sections describe some of the key decisions for determining when and where to create one or more federation servers.

Determine the organizational role for the federation server

To make an informed decision regarding when to create a new federation server, you must first determine in which organization the server will reside. The role that a federation server plays in an organization depends on whether you place the federation server in the account partner organization or in the resource partner organization.

When a federation server is placed in the corporate network of the account partner, its role is to authenticate the user credentials of browser, Web service, or identity selector clients and send security tokens to the clients. For more information, see Review the Role of the Federation Server in the Account Partner.

When a federation server is placed in the corporate network of the resource partner, its role is to authenticate users, based on a security token that is issued by a federation server in the resource partner organization, or its role is to redirect token requests from configured Web applications or Web services to the account partner organization that the client belongs to. For more information, see Review the Role of the Federation Server in the Resource Partner.

Determine which AD FS design to deploy

You create federation servers in your organization whenever you want to deploy any of the following AD FS designs:

- Web SSO Design
- Federated Web SSO Design

If necessary, an organization that deploys a Federated Web SSO design can configure a single federation server so that it acts in both the account partner role and in the resource partner role. In this case, the federation server may produce Security Assertion Markup Language (SAML) tokens, based on user accounts in its own organization, or reroute token requests to the organization, based on where the users' accounts reside.

Note

For the Federated Web SSO design, there must be at least one federation server in the account partner and at least one federation server in the resource partner.

Differences between a federation server and a federation server proxy

A federation server can serve out Web pages for sign-in, policy, authentication, and discovery in the same way that a federation server proxy does. The primary differences between a federation server and a federation server proxy have to do with what operations a federation server can perform that a federation server proxy cannot perform.

The following are the operations that only a federation server can perform:

- The federation server performs the cryptographic operations that produce the token. Although federation server proxies cannot produce tokens, they can be used to route or redirect the tokens to clients and, when necessary, back to the federation server. For more information about using federation servers, see When to Create a Federation Server Proxy.
- Federation servers support the use of Windows Integrated Authentication for clients on the corporate network; federation server proxies do not. For more information about using Windows Integrated Authentication with federation server, see When to Create a Federation Server Farm.

Caution

Communication between federation servers and SQL Server configuration databases, SQL Server attribute stores, domain controllers, and AD LDS instances is not integrity or confidentiality protected by default. To mitigate this, consider protecting the

communication channel between these servers using IPSEC or using a physically secure connection between all of these servers. For communication between federation servers and SQL servers, consider using SSL protection in the connection string. For connections between federation servers and domain controllers, consider turning on Kerberos signing and encryption. For LDAP, LDAP/S is not supported for AD LDS/AD DS.

How to create a federation server

You can create a federation server using the AD FS Federation Server Configuration Wizard or the Fsconfig.exe command-line tool. When you use either of these tools, you can select any of the following options to create a federation server.

- Create a stand-alone federation server

  For more information about how to set up a stand-alone federation server, see Create a Stand-Alone Federation Server.
- Create the first federation server in a federation server farm

  For more information about how to set up the first federation server or add a federation server to a farm, see Create the First Federation Server in a Federation Server Farm.
- Add a federation server to a federation server farm

  For more information about how to add a federation server to a farm, see Add a Federation Server to a Federation Server Farm.

For more detailed information about how each of these options work, see The Role of the AD FS Configuration Database.

For more information about how to set up all the prerequisites necessary to deploy a federation server, see Checklist: Setting Up a Federation Server.

Where to Place a Federation Server

As a security best practice, place Active Directory Federation Services (AD FS) federation servers in front of a firewall and connect them to your corporate network to prevent exposure from the Internet. This is important because federation servers have full authorization to grant security tokens. Therefore, they should have the same protection as a domain controller. If a federation server is compromised, a malicious user has the ability to issue full access tokens to all Web applications and to federation servers that are protected by Active Directory Federation Services (AD FS) in all resource partner organizations.

Note

As a security best practice, avoid having your federation servers directly accessible on the Internet. Consider giving your federation servers direct Internet access only when you are setting up a test lab environment or when your organization does not have a perimeter network.

For typical corporate networks, an intranet-facing firewall is established between the corporate network and the perimeter network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this situation, the federation server sits inside the corporate network, and it is not directly accessible by Internet clients.

Note
Client computers that are connected to the corporate network can communicate directly with the federation server through Windows Integrated Authentication.

A federation server proxy should be placed in the perimeter network before you configure your firewall servers for use with AD FS. For more information, see Where to Place a Federation Server Proxy.

Configuring your firewall servers for a federation server

So that the federation servers can communicate directly with federation server proxies, the intranet firewall server must be configured to allow Secure Hypertext Transfer Protocol (HTTPS) traffic from the federation server proxy to the federation server. This is a requirement because the intranet firewall server must publish the federation server using port 443 so that the federation server proxy in the perimeter network can access the federation server.

In addition, the intranet-facing firewall server, such as a server running Internet Security and Acceleration (ISA) Server, uses a process known as server publishing to distribute Internet client requests to the appropriate corporate federation servers. This means that you must manually create a server publishing rule on the intranet server running ISA Server that publishes the clustered federation server URL, for example, http://fs.fabrikam.com.

For more information about how to configure server publishing in a perimeter network, see Where to Place a Federation Server Proxy. For information about how to configure ISA Server to publish a server, see Create a secure Web publishing rule (http://go.microsoft.com/fwlink/?LinkId=75182).

When to Create a Federation Server Farm

Consider creating a federation server farm in Active Directory Federation Services (AD FS) when you have a larger AD FS deployment and you want to provide fault tolerance, load-balancing, or scalability to your organization's Federation Service. The act of creating two or more federation servers in the same network, configuring each of them to use the same Federation Service, and adding the public key of each server's token-signing certificates to the AD FS Management snap-in creates a federation server farm.

You can create a federation server farm or install additional federation servers to an existing farm by using the AD FS Federation Server Configuration Wizard. For more information, see When to Create a Federation Server.

Note

When you choose the option to create a **New federation server farm** using the AD FS Federation Server Configuration Wizard, the wizard will attempt to create a container object (for sharing certificates) in Active Directory. Therefore, it is important that you first log on to the computer, where you are setting up the federation server role, with an account that has sufficient permissions in Active Directory to create this container object.

Before federation servers can be grouped as a farm, they must first be clustered so that requests that arrive at a single fully qualified domain name (FQDN) are routed to the various federation servers in the server farm. You can create the server cluster by deploying Network Load Balancing (NLB) inside the corporate network. This guide assumes that NLB has been configured appropriately to cluster each of the federation servers in the farm.

For more information about how to configure a cluster FQDN using Microsoft NLB technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkID=74651).

Best practices for deploying a federation server farm

We recommend the following best practices for deploying a federation server in a production environment:

- If you will be deploying multiple federation servers at the same time or you know that you will be adding more servers to the farm over time, consider creating a server image of an existing federation server in the farm and then installing from that image when you need to create additional federation servers quickly.

  Note

  If you do decide to use the server image method for deploying additional federation servers, you do not have to complete the tasks in Checklist: Setting Up a Federation Server every time that you want to add a new server to the farm.
- Use NLB or some other form of clustering to allocate a single IP address for many federation server computers.
- Reserve a static IP address for each federation server in the farm and, depending on your Domain Name System (DNS) configuration, insert an exclusion for each IP address in Dynamic Host Configuration Protocol (DHCP). Microsoft NLB technology requires that each server that participates in the NLB cluster be assigned a static IP address.
- If the AD FS configuration database will be stored in a SQL database, avoid editing the SQL database from multiple federation servers at the same time.

Configuring federation servers for a farm

The following table describes the tasks that must be completed so that each federation server can participate in a farmed environment.

Task	Description
If you are using SQL Server to store the AD FS configuration database	A federation server farm consists of two or more federation servers that share the same AD FS configuration database and token-signing certificates. The configuration database can be stored in either Windows Internal Database or in a SQL Server database. If you plan to store the configuration database in a SQL database, make sure that the configuration database is accessible so that it can be accessed by all new federation servers that participate in the farm. **Note** For farm scenarios, it is important that the configuration database be located on a computer that does not also participate as a federation server in that farm. Microsoft NLB does not allow any of the computers that participate in a farm to communicate with one another. **Note** Ensure that the identity of the AD FS AppPool in Internet Information Services (IIS)) on every federation server that participates in the farm has Read access to the configuration database.
Obtain and share certificates	You can obtain a single server authentication certificate from a public certification authority (CA)—for example, VeriSign. You can then configure the certificate so that all federation servers share the same private key portion of the certificate. For more information about how to share the same certificate, see Checklist: Setting Up a Federation Server. **Note** The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Task	Description
	For more information, see Certificate Requirements for Federation Servers.
Point to the same SQL Server instance	If the AD FS configuration database will be stored in a SQL database, the new federation server must point to the same SQL Server instance that is used by other federation servers in the farm so that the new server can participate in the farm.

Certificate Requirements for Federation Servers

In any Active Directory Federation Services (AD FS) design, various certificates must be used to secure communication and facilitate user authentications between Internet clients and federation servers. Each federation server must have a service communication certificate and a token-signing certificate before it can participate in AD FS communications. The following table describes the certificate types that are associated with federation server.

Certificate type	Description
Token-signing certificate	A token-signing certificate is an X509 certificate. Federation servers use associated public/private key pairs to digitally sign all security tokens that they produce. This includes the signing of published federation metadata and artifact resolution requests. You can have multiple token-signing certificates configured in the AD FS Management snap-in to allow for certificate rollover when one certificate is close to expiring. By default, all the certificates in the list are published, but only the primary token-signing certificate is used by AD FS to actually sign tokens. All certificates that you select must have a corresponding private key. For more information, see Token-Signing Certificates and Add a Token-Signing Certificate.

Certificate type	Description
Service communication certificate	Federation servers use a server authentication certificate, also known as a service communication for Windows Communication Foundation (WCF) Message Security. By default, this is the same certificate that a federation server uses as the Secure Sockets Layer (SSL) certificate in Internet Information Services (IIS). **Note** The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates. For more information, see Service Communications Certificates and Set a Service Communications Certificate. Because the service communication certificate must be trusted by client computers, we recommend that you use a certificate that is signed by a trusted certification authority (CA). All certificates that you select must have a corresponding private key.
Secure Sockets Layer (SSL) certificate	Federation servers use an SSL certificate to secure Web services traffic for SSL communication with Web clients and with federation server proxies. Because the SSL certificate must be trusted by client computers, we recommend that you use a certificate that is signed by a trusted CA. All certificates that you select must have a corresponding private key.
Token-decryption certificate	This certificate is used to decrypt tokens that are received by this federation server. You can have multiple decryption certificates. This makes it possible for a resource federation server to be able to decrypt tokens that are issued with an older certificate after a new certificate is set as the primary decryption certificate. All certificates can be used for

Certificate type	Description
	decryption, but only the primary token-decrypting certificate is actually published in federation metadata. All certificates that you select must have a corresponding private key. For more information, see Add a Token-Decrypting Certificate.

You can request and install an SSL certificate or service communication certificate by requesting a service communication certificate through the Microsoft Management Console (MMC) snap-in for IIS. For more general information about using SSL certificates, see IIS 7.0: Configuring Secure Sockets Layer in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkID=108544) and IIS 7.0: Configuring Server Certificates in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkID=108545).

Note

In AD FS you can change the Secure Hash Algorithm (SHA) level that is used for digital signatures to either SHA-1 or SHA-256 (more secure). AD FSdoes not support the use of certificates with other hash methods, such as MD5 (the default hash algorithm that is used with the Makecert.exe command-line tool). As a security best practice, we recommend that you use SHA-256 (which is set by default) for all signatures. SHA-1 is recommended for use only in scenarios in which you must interoperate with a product that does not support communications using SHA-256, such as a non-Microsoft product or AD FS 1. *x*.

Determining your CA strategy

AD FS does not require that certificates be issued by a CA. However, the SSL certificate (the certificate that is also used by default as the service communications certificate) must be trusted by the AD FS clients. We recommend that you not use self-signed certificates for these certificate types.

Important

Use of self-signed, SSL certificates in a production environment can allow a malicious user in an account partner organization to take control of federation servers in a resource partner organization. This security risk exists because self-signed certificates are root certificates. They must be added to the trusted root store of another federation server (for example, the resource federation server), which can leave that server vulnerable to attack.

After you receive a certificate from a CA, make sure that all certificates are imported into the personal certificate store of the local computer. You can import certificates to the personal store with the Certificates MMC snap-in.

As an alternative to using the Certificates snap-in, you can also import the SSL certificate with the IIS Manager snap-in at the time that you assign the SSL certificate to the default Web site. For more information, see Import a Server Authentication Certificate to the Default Web Site.

Note

Before you install the AD FS software on the computer that will become the federation server, make sure that both certificates are in the Local Computer personal certificate store and that the SSL certificate is assigned to the Default Web Site. For more information about the order of the tasks that are required to set up a federation server, see Checklist: Setting Up a Federation Server.

Depending on your security and budget requirements, carefully consider which of your certificates will be obtained by a public CA or a corporate CA. The following figure shows the recommended CA issuers for a given certificate type. This recommendation reflects a best-practice approach regarding security and cost.

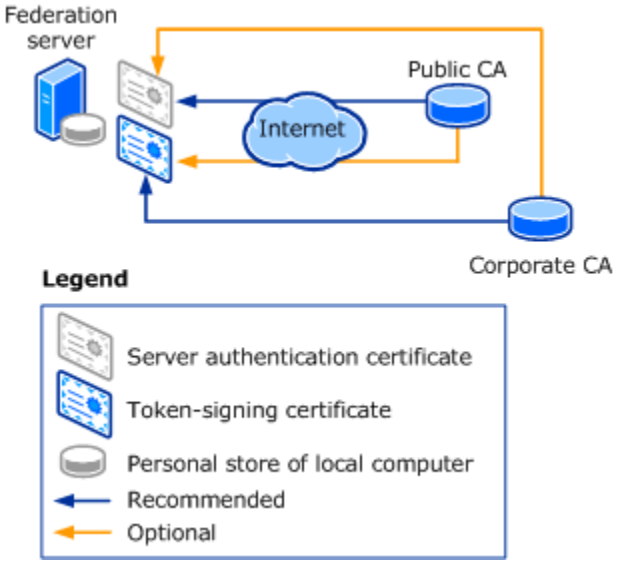

Certificate revocation lists

If any certificate that you use has CRLs, the server with the configured certificate must be able to contact the server that distributes the CRLs.

Token-Signing Certificates

Federation servers require token-signing certificates to prevent attackers from altering or counterfeiting security tokens in an attempt to gain unauthorized access to federated resources. The private/public key pairing that is used with token-signing certificates is the most important validation mechanism of any federated partnership because these keys verify that a security

token was issued by a valid partner federation server and that the token was not modified during transit.

Token-signing certificate requirements

A token-signing certificate must meet the following requirements to work with AD FS:

- For a token-signing certificate to successfully sign a security token, the token-signing certificate must contain a private key.
- The AD FS service account must have access to the token-signing certificate's private key in the personal store of the local computer. This is taken care of by Setup. You can also use the AD FS Management snap-in to ensure this access if you subsequently change the token-signing certificate.

Note

It is a public key infrastructure (PKI) best practice to not share the private key for multiple purposes. Therefore, do not use the service communication certificate that you installed on the federation server as the token-signing certificate.

How token-signing certificates are used across partners

Every token-signing certificate contains cryptographic private keys and public keys that are used to digitally sign (by means of the private key) a security token. Later, after they are received by a partner federation server, these keys validate the authenticity (by means of the public key) of the encrypted security token.

Because each security token is digitally signed by the account partner, the resource partner can verify that the security token was in fact issued by the account partner and that it was not modified. Digital signatures are verified by the public key portion of a partner's token-singing certificate. After the signature is verified, the resource federation server generates its own security token for its organization and it signs the security token with its own token-signing certificate.

For federation partner environments, when the token-signing certificate has been issued by a CA, ensure that:

1. The certificate revocation lists (CRLs) of the certificate are accessible to relying parties and Web servers that trust the federation server.
2. The root CA certificate is trusted by the relying parties and Web servers that trust the federation server.

The Web server in the resource partner uses the public key of the token-signing certificate to verify that the security token is signed by the resource federation server. The Web server then allows the appropriate access to the client.

Deployment considerations for token-signing certificates

When you deploy the first federation server in a new AD FS installation, you must obtain a token-signing certificate and install it in the local computer personal certificate store on that federation server. You can obtain a token-signing certificate by requesting one from an enterprise CA or a public CA or by creating a self-signed certificate.

When you deploy an AD FS farm, token-signing certificates are installed differently, depending on how you create the server farm.

There are two server farm options that you can consider when you obtain token-signing certificates for your deployment:

- A private key from one token-signing certificate is shared among all the federation servers in a farm.

  In a federation server farm environment, we recommend that all federation servers share (or reuse) the same token-signing certificate. You can install a single token-signing certificate from a CA on a federation server and then export the private key, as long as the issued certificate is marked as exportable.

  As shown in the following illustration, the private key from a single token-signing certificate can be shared to all the federation servers in a farm. This option—compared to the following "unique token-signing certificate" option—reduces costs if you plan to obtain a token-signing certificate from a public CA.

- There is a unique token-signing certificate for each federation server in a farm.

  When you use multiple, unique certificates throughout your farm, each server in that farm signs tokens with its own unique private key.

As shown in the following illustration, you can obtain a separate token-signing certificate for every single federation server in the farm. This option is more expensive if you plan to obtain your token-signing certificates from a public CA.

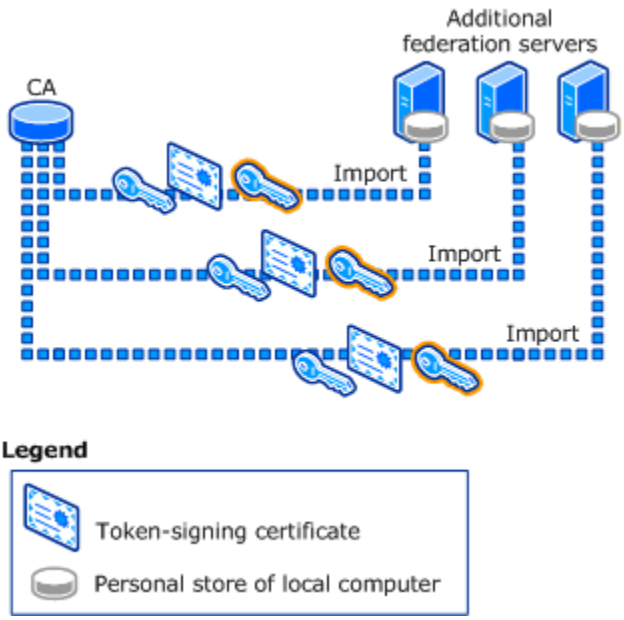

For information about installing a certificate when you use Microsoft Certificate Services as your enterprise CA, see IIS 7.0: Create a Domain Server Certificate in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkId=108548).

For information about installing a certificate from a public CA, see IIS 7.0: Request an Internet Server Certificate (http://go.microsoft.com/fwlink/?LinkId=108549).

For information about installing a self-signed certificate, see IIS 7.0: Create a Self-Signed Server Certificate in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkID=108271).

Service Communications Certificates

A federation server requires the use of service communication certificates for scenarios in which WCF message security is used.

Service communication certificate requirements

Service communication certificates must meet the following requirements to work with AD FS:

- The service communication certificate must include the server authentication enhanced key usage (EKU) extension.
- The certificate revocation lists (CRLs) must be accessible for all the certificates in the chain from the service communication certificate to the root CA certificate. The root CA must also be trusted by any federation server proxies and Web servers that trust this federation server.

- The subject name that is used in the service communication certificate must match the Federation Service name in the properties of the Federation Service.

Deployment considerations for service communication certificates

Configure service communication certificates so that all federation servers use the same certificate. If you are deploying the Federated Web Single-Sign-On (SSO) design, we recommend that your service communication certificate be issued by a public CA. You can request and install these certificates through the IIS Manager snap-in.

You can use self-signed, service communication certificates successfully on federation servers in a test lab environment. However, for a production environment, we recommend that you obtain service communication certificates from a public CA. The following are reasons why you should not use self-signed, service communication certificates for a live deployment:

- A self-signed, SSL certificate must be added to the trusted root store on each of the federation servers in the resource partner organization. While this alone does not enable an attacker to compromise a resource federation server, trusting self-signed certificates does increase the attack surface of a computer, and it can lead to security vulnerabilities if the certificate signer is not trustworthy.
- It creates a bad user experience. Clients will receive Security Alert prompts when they try to access federated resources that display the following message: "The security certificate was issued by a company you have not chosen to trust." This is expected behavior, because the self-signed certificate is not trusted.

  Note
  If necessary, you can work around this condition by using Group Policy to manually push down the self-signed certificate to the trusted root store on each client computer that will attempt to access an AD FS site.
- CAs provide additional certificate-based features, such as private key archive, renewal, and revocation, that are not provided by self-signed certificates.

Name Resolution Requirements for Federation Servers

When client computers on the corporate network attempt to access an application or Web service that is protected by Active Directory Federation Services (AD FS), they must first authenticate to a federation server. One way to authenticate is to have the corporate network clients access a local federation server through Windows Integrated Authentication.

Configure corporate DNS

So that successful name resolution through Windows Integrated Authentication on local federation servers can occur, Domain Name System (DNS) in the corporate network of the account partner must be configured for a new host (A) resource record that will resolve the fully qualified domain name (FQDN) host name of the federation server to the IP address of the federation server cluster.

In the following illustration, you can see how this task is accomplished for a given scenario. In this scenario, Microsoft Network Load Balancing (NLB) provides a single cluster FQDN name and a single cluster IP address for an existing federation server farm.

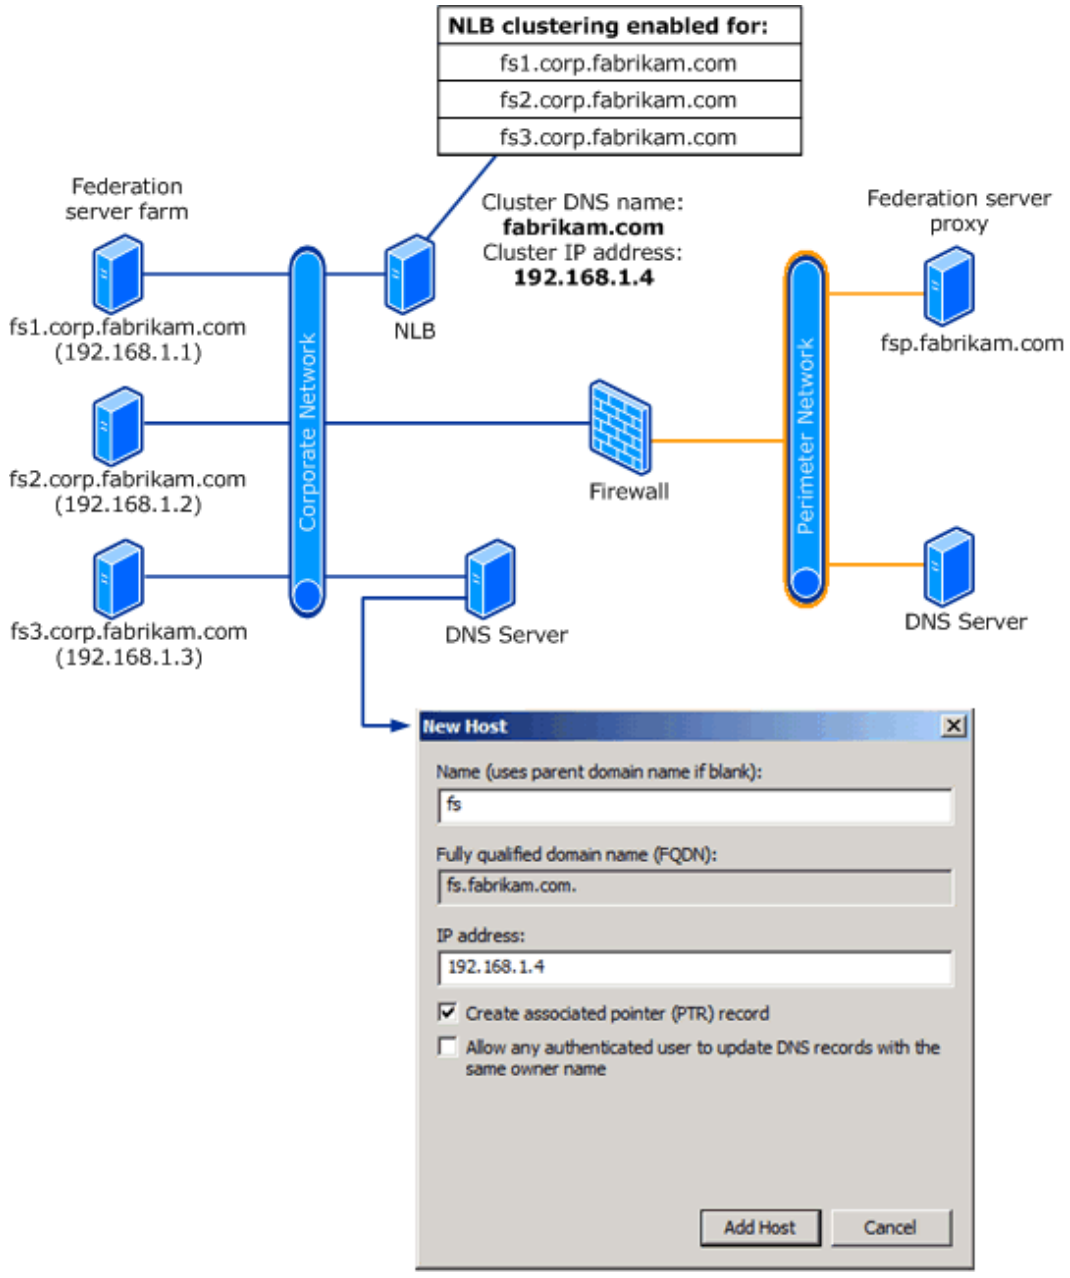

For information about how to configure a cluster IP address or cluster FQDN using NLB, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkId=75282).

For information about how to configure corporate DNS for a federation server, see Add a Host (A) Resource Record to Corporate DNS for a Federation Server.

For information about how to configure federation server proxies in the perimeter network, see Name Resolution Requirements for Federation Server Proxies.

Planning Federation Server Proxy Placement

After you gather all the information that you will use to design your Active Directory Federation Services (AD FS) infrastructure and after you plan your federation server and Web server strategy, you can plan when and where to place federation server proxies in your new design. The information in the following topics can help you determine when and where to place a federation server proxy and whether to configure it for the account partner role or the resource partner role:

- Review the Role of the Federation Server in the Account Partner
- Review the Role of the Federation Server Proxy in the Resource Partner
- When to Create a Federation Server Proxy
- Where to Place a Federation Server Proxy
- When to Create a Federation Server Proxy Farm
- Certificate Requirements for Federation Server Proxies
- Name Resolution Requirements for Federation Server Proxies

Note
Although this information might help with your placement planning for federation server proxies, it does not explain how to determine the proper number of proxies and the proxy hardware requirements for each AD FS design.

For examples of how you can place a federation server proxy in either of the two primary AD FS design scenarios, see Mapping Your Deployment Goals to an AD FS Design.

When to Create a Federation Server Proxy

Creating a federation server proxy in your organization adds additional security layers to your Active Directory Federation Services (AD FS) deployment. Consider deploying a federation server proxy in your organization's perimeter network when you want to:

- Prevent external client computers from directly accessing your federation servers. By deploying a federation server proxy in your perimeter network, you effectively isolate your federation servers so that they can be accessed only by client computers that are logged in to the corporate network through federation server proxies, which act on behalf of the external client computers. Federation server proxies do not have access to the private keys that are used to produce tokens. For more information, see Where to Place a Federation Server Proxy.
- Provide a convenient way to differentiate the sign-in experience for users who are coming from the Internet as opposed to users who are coming from your corporate network using Windows Integrated Authentication. A federation server proxy collects credentials or home realm details from Internet client computers by using the logon, logout, and identity provider discovery (homerealmdiscovery.aspx) pages that are stored on the federation server proxy.

In contrast, client computers that come from the corporate network encounter a different experience, based on the configuration of the federation server. The corporate network federation server is often configured for Windows Integrated Authentication, which provides a seamless sign-in experience for users on the corporate network.

The role that a federation server proxy plays in your organization depends on whether you place the federation server proxy in the account partner organization or in the resource partner organization. For example, when a federation server proxy is placed in the perimeter network of the account partner, its role is to collect the user credential information from browser clients. When a federation server proxy is placed in the perimeter network of the resource partner, it relays security token requests to a resource federation server and produces organizational security tokens in response to the security tokens that are provided by its account partners.

For more information, see Review the Role of the Federation Server Proxy in the Account Partner and Review the Role of the Federation Server Proxy in the Resource Partner

How to create a federation server proxy

You can create a federation server proxy using either the AD FS Federation Server Proxy Configuration Wizard or the Fsconfig.exe command-line tool. For instructions about how to do this, see Configure a Computer for the Federation Server Proxy Role.

For general information about how to set up all the prerequisites necessary to deploy a federation server proxy, see Checklist: Setting Up a Federation Server Proxy.

Where to Place a Federation Server Proxy

You can place Active Directory Federation Services (AD FS) federation server proxies in a perimeter network to provide a protection layer against malicious users that may be coming from the Internet. Federation server proxies are ideal for the perimeter network environment because they do not have access to the private keys that are used to create tokens. However, federation server proxies can efficiently route incoming requests to federation servers that are authorized to produce those tokens.

It is not necessary to place a federation server proxy inside the corporate network for either the account partner or the resource partner because client computers that are connected to the corporate network can communicate directly with the federation server. In this scenario, the federation server also provides federation server proxy functionality for client computers that are coming from the corporate network.

As is typical with perimeter networks, an intranet-facing firewall is established between the perimeter network and the corporate network, and an Internet-facing firewall is often established between the perimeter network and the Internet. In this scenario, the federation server proxy sits between both of these firewalls on the perimeter network.

Configuring your firewall servers for a federation server proxy

For the federation server proxy redirection process to be successful, all firewall servers must be configured to allow Secure Hypertext Transfer Protocol (HTTPS) traffic. The use of HTTPS is required because the firewall servers must publish the federation server proxy, using port 443, so that the federation server proxy in the perimeter network can access the federation server in the corporate network.

Note

All communications to and from client computers also occur over HTTPS.

In addition, the Internet-facing firewall server, such as a computer running Microsoft Internet Security and Acceleration (ISA) Server, uses a process known as server publishing to distribute Internet client requests to the appropriate perimeter and corporate network servers, such as federation server proxies or federation servers.

Server publishing rules determine how server publishing works—essentially, filtering all incoming and outgoing requests through the ISA Server computer. Server publishing rules map incoming client requests to the appropriate servers behind the ISA Server computer. For information about how to configure ISA Server to publish a server, see Create a Secure Web Publishing Rule (http://go.microsoft.com/fwlink/?LinkId=75182).

In the federated world of AD FS, these client requests are typically made to a specific URL, for example, a federation server identifier URL such as http://fs.fabrikam.com. Because these client requests come in from the Internet, the Internet-facing firewall server must be configured to publish the federation server identifier URL for each federation server proxy that is deployed in the perimeter network.

Configuring ISA Server to allow SSL

To facilitate secure AD FS communications, you must configure ISA Server to allow Secure Sockets Layer (SSL) communications between the following:

- **Federation servers and federation server proxies.** An SSL channel is required for all communications between federation servers and federation server proxies. Therefore, you must configure ISA Server to allow an SSL connection between the corporate network and the perimeter network.
- **Client computers, federation servers, and federation server proxies.** So that communications can occur between client computers and federation servers or between client computers and federation server proxies, you can place a computer running ISA Server in front of the federation server or federation server proxy.

  If your organization performs SSL client authentication on the federation server or federation server proxy, when you place a computer running ISA Server in front of the federation server or federation server proxy, the server must be configured for pass-through of the SSL connection because the SSL connection must terminate at the federation server or federation server proxy.

If your organization does not perform SSL client authentication on the federation server or federation server proxy, an additional option is to terminate the SSL connection at the computer running ISA Server and then re-establish an SSL connection to the federation server or federation server proxy.

Note
The federation server or federation server proxy requires that the connection be secured by SSL to protect the contents of the security token.

When to Create a Federation Server Proxy Farm

Consider installing additional federation server proxies when you have a large Active Directory Federation Services (AD FS) deployment and you want to provide fault tolerance, load-balancing, and scalability for your proxy deployment. The act of creating two or more federation server proxies in the same perimeter network and configuring each of them to protect the same AD FS Federation Service creates a federation server proxy farm.

You can create a federation server proxy farm or install additional federation server proxies to an existing farm by using the AD FS Federation Server Proxy Configuration Wizard. For more information, see When to Create a Federation Server Proxy.

Before all the federation server proxies can function together as a farm, you must first cluster them under one IP address and one Domain Name System (DNS) fully qualified domain name (FQDN). You can cluster the servers by deploying Microsoft Network Load Balancing (NLB) inside the perimeter network. The tasks in the following table require NLB to be configured appropriately to cluster the federation server proxies in the farm.

For more information about how to configure an FQDN for a cluster using Microsoft NLB technology, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?linkid=74651).

Configuring federation server proxies for a farm

The following table describes the tasks that must be completed so that each federation server proxy can participate in a farm.

Task	Description
Point all proxies in the farm to the same AD FS Federation Service name	When you create the federation server proxies, you must type the same Federation Service name in the AD FS Federation Server Proxy Configuration Wizard for all the federation server proxies that will participate in the farm. The federation server proxy uses the URL that

Task	Description
	makes up this DNS host name to determine which AD FS Federation Service instance it contacts. For more information, see Configure a Computer for the Federation Server Proxy Role.
Obtain and share certificates	You can obtain a server authentication certificate from a public certification authority (CA)—for example, VeriSign—and then configure the certificate so that all federation server proxies share the same private key portion of the same certificate on the default Web site for each federation server proxy. To share the certificate, you must install the same server authentication certificate on the default Web site for each federation server proxy. For more information, see Import a Server Authentication Certificate to the Default Web Site. For more information, see Certificate Requirements for Federation Server Proxies.

For more information about adding new federation server proxies to create a federation server proxy farm, see Checklist: Setting Up a Federation Server Proxy.

Certificate Requirements for Federation Server Proxies

Servers that are running in the federation server proxy role in Active Directory Federation Services (AD FS) are required to use Secure Sockets Layer (SSL) server authentication certificates. Federation server proxies use SSL server authentication certificates to secure Web server traffic communication with Web clients.

Federation server proxies are usually exposed to computers on the Internet that are not included in your enterprise public key infrastructure (PKI). Therefore, use a server authentication certificate that is issued by a public (third-party) certification authority (CA), for example, VeriSign.

When you have a federation server proxy farm, all federation server proxy computers must use the same server authentication certificate. For more information, see When to Create a Federation Server Proxy Farm.

It is important to verify that the subject name in the server authentication certificate matches the Federation Service name value that is specified in the AD FS Management snap-in. To locate this value, open the snap-in, right-click **Service**, click **Edit Federation Service Properties**, and then find the value in **Federation Service name** text box.

For general information about using SSL certificates, see Configuring Secure Sockets Layer in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkID=108544) and Configuring Server Certificates in IIS 7.0 (http://go.microsoft.com/fwlink/?LinkID=108545).

Note

Client authentication certificates are not required for AD FS federation server proxies.

If any certificate that you use has certificate revocation lists (CRLs), the server with the configured certificate must be able to contact the server that distributes the CRLs. The type of CRL determines what ports are used.

Name Resolution Requirements for Federation Server Proxies

When client computers on the Internet attempt to access an application that is secured by Active Directory Federation Services (AD FS), they must first authenticate to the federation server. In most cases, the federation server is usually not directly accessible from the Internet. Therefore, Internet client computers must be redirected to the federation server proxy instead. You can accomplish successful redirection by adding the appropriate Domain Name System (DNS) records to your DNS zone or zones that face the Internet.

The method that you use to redirect Internet clients to the federation server proxy depends on how you configure the DNS zone in your perimeter network or how you configure a DNS zone that you control on the Internet. Federation server proxies are intended for use in a perimeter network. They redirect Internet client requests to federation servers successfully only when DNS has been configured properly in all the Internet-facing zones that you control. Therefore, the configuration of your Internet-facing zones—whether you have a DNS zone serving only the perimeter network or a DNS zone serving both the perimeter network and Internet clients—is important.

This topic describes the steps that you can take to configure name resolution when you place a federation server proxy in your perimeter network. To determine which steps to follow, first determine which of the following DNS scenarios most closely matches the DNS infrastructure in the perimeter network of your organization. Then, follow the steps for that scenario.

DNS zone serving only the perimeter network

In this scenario, your organization has one or two DNS zones in the perimeter network, and your organization does not control any DNS zones on the Internet. Successful name resolution for a

federation server proxy in the DNS zone that serves only the perimeter network scenario depends on the following conditions:

- The federation server proxy must have a setting in the hosts file to resolve the fully qualified domain name (FQDN) of the federation server endpoint URL to an IP address of a federation server or a federation server cluster.
- DNS in the perimeter network of the account partner must be configured so that the FQDN of the federation server endpoint URL resolves to the IP address of the federation server proxy.

The following illustration and corresponding steps show how each of these conditions is achieved for a given example. In this illustration, Microsoft Network Load Balancing (NLB) technology provides a single, cluster FQDN and a single, cluster IP address for an existing federation server farm.

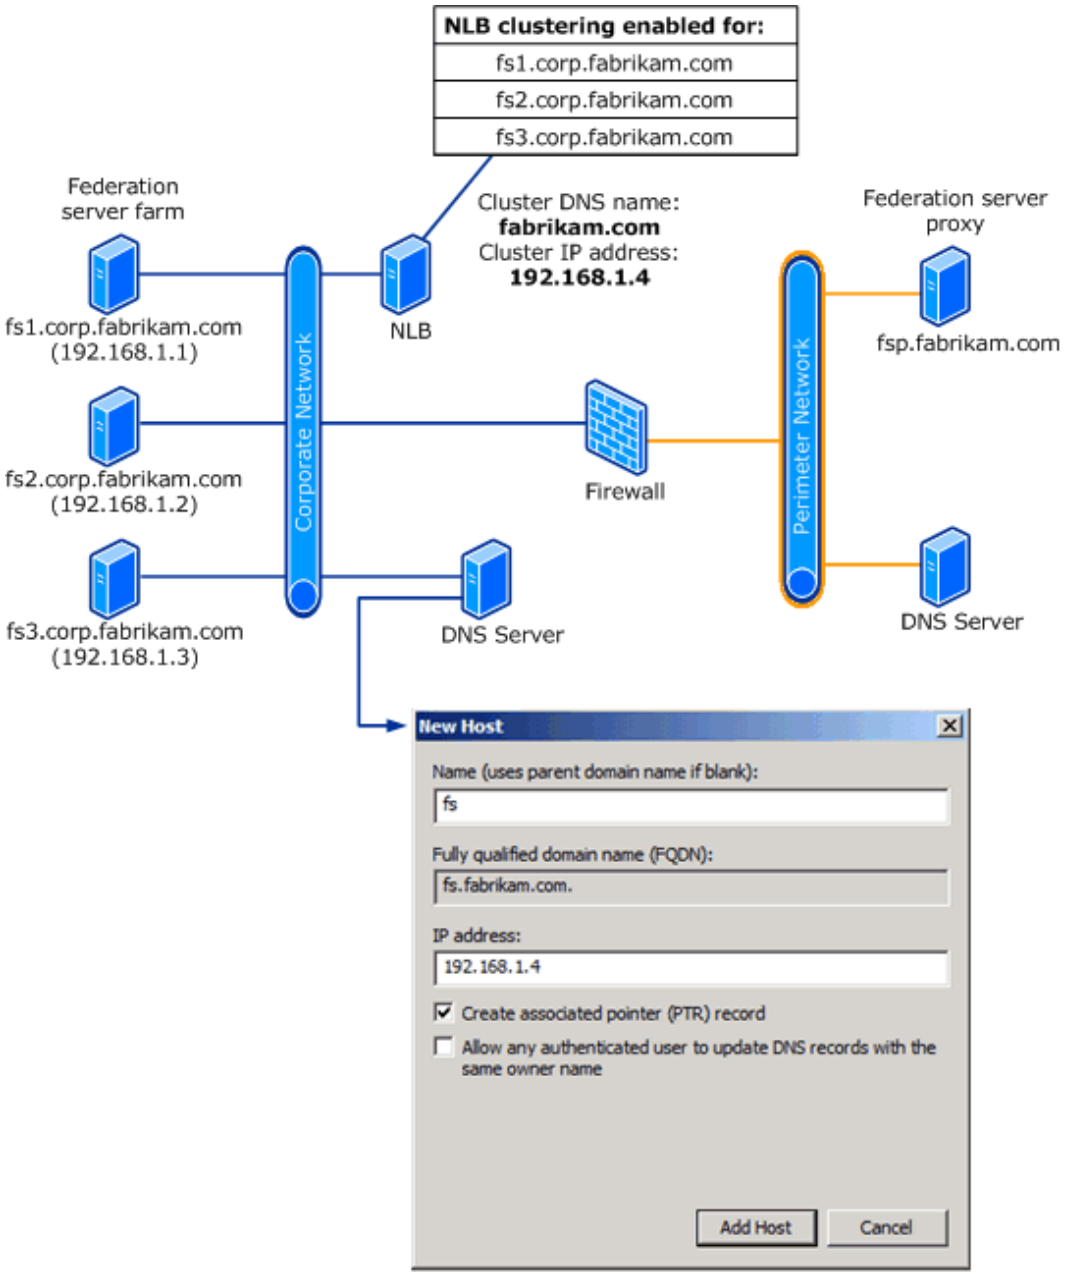

For more information about configuring a cluster IP address or a cluster FQDN using NLB, see Specifying the Cluster Parameters (http://go.microsoft.com/fwlink/?LinkId=75282).

1. Configure the hosts file on the federation server proxy

Because DNS in the perimeter network is configured to resolve all requests for fs.fabrikam.com to the account federation server proxy, the account partner federation server proxy has an entry in its local hosts file to resolve fs.fabrikam.com to the IP address of the actual account federation server (or cluster DNS name for the federation server farm) that is connected to the corporate

network. This makes it possible for the account federation server proxy to resolve the host name fs.fabrikam.com to the account federation server rather than to itself—as would occur if it attempted to look up fs.fabrikam.com using perimeter DNS—so that the federation server proxy can communicate with the federation server.

2. Configure perimeter DNS

Because there is only a single AD FS host name that client computers are directed to—whether they are on an intranet or on the Internet—client computers on the Internet that use the perimeter DNS server must resolve the FQDN for the account federation server (fs.fabrikam.com) to the IP address of the account federation server proxy on the perimeter network. So that it can forward clients on to the account federation server proxy when they attempt to resolve fs.fabrikam.com, perimeter DNS contains a limited corp.fabrikam.com DNS zone with a single host (A) resource record for fs (fs.fabrikam.com) and the IP address of the account federation server proxy on the perimeter network.

For more information about how to modify the hosts file of the federation server proxy and configure DNS in the perimeter network, see Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Only the Perimeter Network.

DNS zone serving both the perimeter network and Internet clients

In this scenario, your organization controls the DNS zone in the perimeter network and at least one DNS zone on the Internet. Successful name resolution for a federation server proxy in this scenario depends on the following conditions:

- DNS in the Internet zone of the account partner must be configured so that the FQDN of the federation server host name resolves to the IP address of the federation server proxy in the perimeter network.
- DNS in the perimeter network of the account partner must be configured so that the FQDN of the federation server host name resolves to the IP address of the federation server in the corporate network.

The following illustration and corresponding steps show how each of these conditions is achieved for a given example.

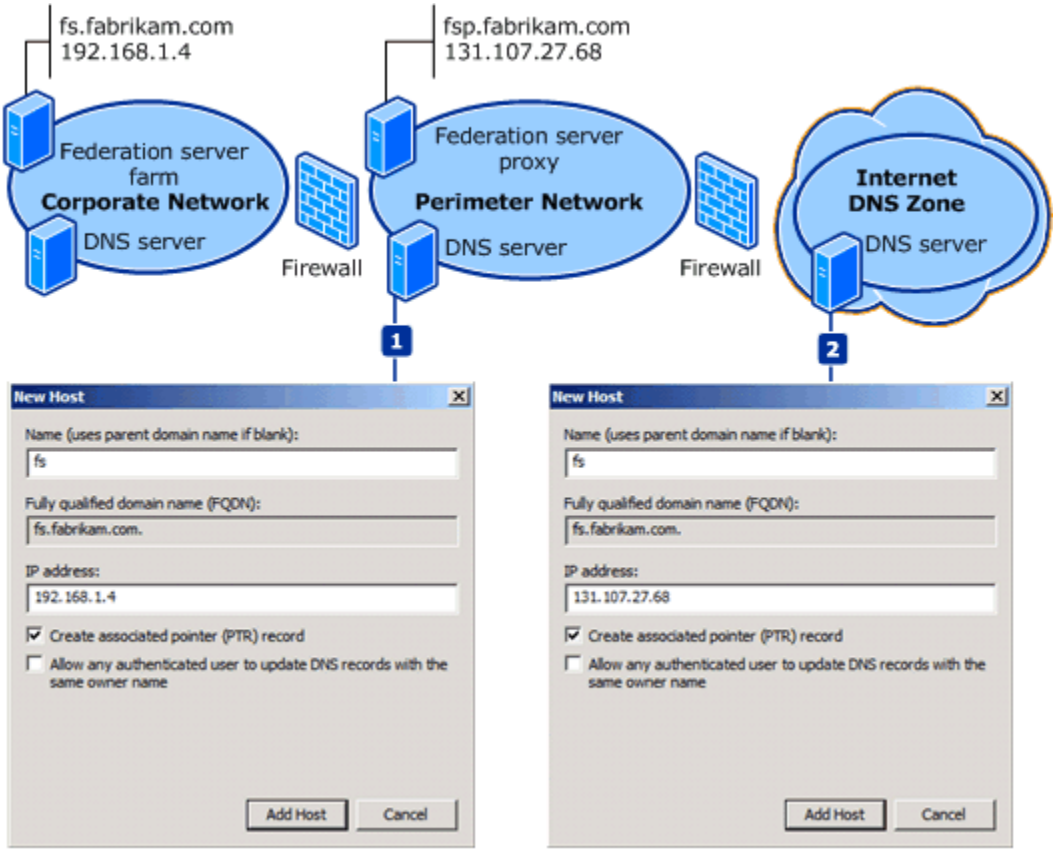

1. Configure perimeter DNS

For this scenario, because it is assumed that you will configure the Internet DNS zone that you control to resolve requests that are made for a specific endpoint URL (that is, fs.fabrikam.com) to the federation server proxy in the perimeter network, you must also configure the zone in the perimeter DNS to forward these requests to the federation server in the corporate network.

So that clients can be forwarded to the account federation server when they attempt to resolve fs.fabrikam.com, perimeter DNS is configured with a single host (A) resource record for fs (fs.fabrikam.com) and the IP address of the account federation server on the corporate network. This makes it possible for the account federation server proxy to resolve the host name fs.fabrikam.com to the account federation server rather than to itself—as would occur if it attempted to look up fs.fabrikam.com using Internet DNS—so that the federation server proxy can communicate with the federation server.

2. Configure Internet DNS

For name resolution to be successful in this scenario, all requests from client computers on the Internet to fs.fabrikam.com must be resolved by the Internet DNS zone that you control. Consequently, you must configure your Internet DNS zone to forward client requests for

fs.fabrikam.com to the IP address of the account federation server proxy in the perimeter network.

For more information about how to modify the perimeter network and Internet DNS zones, see Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Both the Perimeter Network and Internet Clients.

Planning for AD FS Server Capacity

Note

The content provided in this topic does not reflect actual testing that was performed on servers running Windows Server 2012. This topic will be updated once the required testing has been performed.

Capacity planning for Active Directory Federation Services (AD FS) is the process of forecasting peak usage periods for your Federation Service and planning or scaling-up your AD FS server deployment to meet those load requirements.

This section describes deployment guidelines for both the federation server and federation server proxy roles and is based on lab testing that was performed by the AD FS product team at Microsoft. The purpose of this content is to help you:

- Closely estimate the hardware needs for your organization's specific AD FS deployment, such as the number of AD FS servers.
- Accurately project the expected peak usage for sign-in requests, plan for growth, and ensure that your AD FS deployment is capable of handling that expected peak usage.

Before you proceed with reading this capacity planning content, we recommend that you first complete the tasks in the order shown in the following two tables. In the first table, we provide links to recommended tasks that will help provide relevant context for this capacity planning discussion.

Recommended task	Description	Reference
Understand the requirements for deploying AD FS federation servers and federation server proxies	Review important hardware and software requirements necessary for deploying federation server and federation server proxies.	Appendix A: Reviewing AD FS Requirements
Select the type of AD FS configuration database that you will deploy in your organization	Before you can begin using capacity planning data in this section, you first have to determine which AD FS configuration database type you will deploy, either	The Role of the AD FS Configuration Database; AD FS Deployment Topology Considerations

Recommended task	Description	Reference
	Windows Internal Database (WID) or a Structured Query Language (SQL) database.	
Determine the type of topology layout to use with your new AD FS configuration database selection	Once you have decided on the type of AD FS configuration database to use in your deployment, you will need to consider which deployment topology most closely matches where you will need to place federation servers and federation server proxies within your production environment.	Determine Your AD FS Deployment Topology
Understand key AD FS–related capacity planning terms	Review the definitions of common capacity planning terms that are used throughout the AD FS capacity planning discussion.	See the section titled AD FS capacity planning terms in this topic

Once you have reviewed the content in the previous table, you can now complete the prerequisite tasks in the next table.

Prerequisite task	Description	Reference
Download the AD FS Capacity Planning Sizing Spreadsheet	The AD FS Capacity Planning Sizing spreadsheet can help you to determine the number of federation servers required for an AD FS federation server farm deployment. Instructions for how to use this spreadsheet are available in the link provided below for the next task.	AD FS Capacity Planning Sizing Spreadsheet
Gather data about the number of users who will require single sign-on (SSO) access to the target claims-aware application and the expected	This user data you collect will be used for the input values required within the context of the AD FS Capacity Planning Sizing Spreadsheet.	Estimate the number of federation servers for your organization

Prerequisite task	Description	Reference
peak usage periods associated with this access		

AD FS capacity planning terms

The following table describes important terms that are used often in this capacity planning section of the AD FS Design Guide. For a more complete list of AD FS terms, see Understanding Key AD FS Concepts.

Term	Definition
Concurrent users	The estimated number of users that are expected to submit requests to the service within a given period of time, usually a peak activity period.
Active users	The approximate average number of users that are active on a system, but not necessarily submitting requests, during a given period of time.
Defined users	A theoretical maximum user count, usually based on the number of users who have defined accounts in the system.
Requests per second	The number of requests either submitted by clients (when talking about the load on a system) or processed by servers (when talking about server throughput) in a second. This metric is used in planning server processor and memory capacity.
Target server responsiveness and utilization	Success metrics that bound the acceptable server performance range. Generally, if responsiveness goes below or utilization goes above the target, the system is considered to be overloaded and more capacity is required.
Windows Internal Database (WID)	The default AD FS configuration database that can be used as an alternative to SQL Server in certain AD FS deployments.

Configuration environment used during AD FS testing

This section describes the configuration environment that the AD FS product team used to perform its tests. The team used the following computer hardware, software, and network configuration to gather performance and scalability data in tests of the federation server:

- Dual Quad Core 2.27 gigahertz (GHz) (8 cores)
- 16-GB RAM
- Windows Server 2008 R2, Enterprise Edition
- Gigabit Network

Note
Although 16 GB's of RAM was used on the federation server during testing, a more moderate memory size, such as 4 GB's of RAM per federation server can be used for most AD FS deployments. The recommendations that are provided in this AD FS Capacity Planning content along with the results provided by the AD FS Capacity Planning Spreadsheet are based on assumptions that each federation server will use approximately 4GB's of RAM for most AD FS production environments.

The product team used the following configuration to gather performance and scalability data for the federation server proxy testing:

- Dual Quad Core 2.24 GHz (4 cores)
- 4-GB RAM
- Windows Server 2008 R2, Enterprise Edition
- Gigabit Network

Note
Capacity recommendations for AD FS servers can vary considerably, depending on the specifications you choose for the hardware and network configuration to be used in a given environment. As a point of reference, the sizing guidance provided in this content is based on a utilization target of 80 percent on the computers mentioned earlier.

Measure AD FS server capacity

Typically, the hardware components that affect server performance and scalability are the CPU, memory, the disk, and network adapters. Fortunately, each of the AD FS components requires very little demand on memory and disk space. Network connectivity is an obvious requirement. Therefore, load tests that are performed on federation servers and federation server proxies concentrate on two primary areas for measuring server capacity:

- **Peak AD FS requests per second:** The number of sign-in requests that are processed per second on federation servers. This measurement can help you determine how many simultaneous users can sign in to a given server. You can use this measurement in

conjunction with the CPU consumption measurement to understand this measurement's effect on performance.

- **CPU consumption:** The percentage by which CPU capacity is measured. This measurement can help you determine the overall CPU load that occurred based on the number of incoming sign-in requests per second.

Continue reading more about AD FS capacity planning

After you have completed the prerequisite tasks and have become familiar with related terms and hardware requirements, you can use the following additional capacity planning content to help you determine the recommended number of AD FS servers required for your deployment:

- Planning for Federation Server Capacity
- Planning for Federation Server Proxy Capacity

Planning for Federation Server Capacity

Capacity planning for federation servers helps you estimate:

- Which factors grow the size of the AD FS configuration database.
- The appropriate hardware requirements for each federation server.
- The number of federation servers to place in each organization.

Federation servers issue security tokens to users. These tokens are presented to a relying party for consumption. Federation servers issue security tokens after authenticating a user or after receiving a security token that was previously issued by a partner Federation Service. A security token is requested from a Federation Service when users initially sign in to federated applications or when their security tokens expire while they are accessing federated applications.

Federation servers are designed to accommodate high-availability server farm configurations that use Microsoft Network Load Balancing (NLB) technology. Federation servers in a farm configuration can service requests independently, without accessing any common farm components for each request. Therefore, there is little overhead involved in scaling out a federation server deployment.

Recommendations:

- For mission-critical or high-availability deployments, we recommend that you create a small federation server farm in each partner organization, with at least two federation servers per farm, to provide fault tolerance.
- With the need for high availability and the ease of scaling out federation servers, scaling out is the recommended method for handling high numbers of requests per second for a particular Federation Service. Scaling up beyond the base configuration in this guide is unlikely to produce significant capacity handling gains.

AD FS configuration database size and growth

The size of the AD FS configuration database is generally considered to be small, and database size does not tend to be a major consideration in AD FS deployments. The precise size of the AD FS configuration database can depend largely on the number of trust relationships and the associated trust-related metadata—such as claims, claim rules, and monitoring settings configured for each trust. As the number of trust entries in the configuration database grows, so does the need for more disk space.

For additional deployment information about the AD FS configuration database, see AD FS Deployment Topology Considerations.

Memory, CPU and disk space requirements

Fortunately, memory, CPU and disk space requirements for federation servers are modest, and they are not likely to be a driving factor in hardware decisions. For more information about hardware requirements, see Appendix A: Reviewing AD FS Requirements.

Note

In tests that were performed by the AD FS product team using a federation server farm configured with a dedicated SQL Server to store the AD FS configuration database, the overall load on the SQL Server tended to be low. In one test using a four-federation-server farm that was configured to use a single SQL Server, CPU utilization did not exceed 10% despite testing that brought the federation servers to target utilization.

Estimate the number of federation servers for your organization

In an effort to streamline the hardware planning process for federation servers, the AD FS product team developed the AD FS Capacity Planning Sizing Spreadsheet. This Excel spreadsheet includes calculator-like functionality that will take expected usage data that you provide about users in your organization and return a recommended optimal number of federation servers for your AD FS production environment.

Note

The number of federation servers that this spreadsheet will recommend is based on the hardware and network specifications that the AD FS product team used during testing. Therefore, the number of federation servers that the spreadsheet will recommend must be understood within this context. For more information about the specifications used during testing, see the topic titled Planning for AD FS Server Capacity.

Using the AD FS Capacity Planning Sizing Spreadsheet

When you use this spreadsheet, you will need to select a value (either **40%**, **60%**, or **80%**) that best represents the percentage of total users you expect will send authentication requests to your federation servers during peak usage periods.

Then, you will need to select a value (either **1 minute**, **15 minutes**, or **1 hour**) that best represents the length of time you expect the peak usage period to last. For example, you might estimate 40% as the value for the total number of users who will login within a period of 15 minutes, or that 60% of users will login within a period of 1 hour. Together, these values define the peak load profile by which your sizing recommendation will be calculated.

Next, you will need to specify the total number of users that will require single sign-on access to the target claims-aware application, based on whether the users are:

- Logging into Active Directory from a local computer that is physically connected to your corporate network (through Windows integrated authentication)
- Logging into Active Directory remotely from a computer that is not physically connected to your corporate network (through Windows integrated authentication or Username and password)
- From another organization and are attempting to access the target claims-aware application from a trusted partner
- From a SAML 2.0 identity provider and are attempting to access the target claims-aware application

How to use this spreadsheet

You can use the following steps for each federation server farm instance you plan to deploy to determine the recommended number of federation servers.

1. Download and then open the AD FS Capacity Planning Sizing Spreadsheet
2. In the cell to the right of the **During the peak system usage period, I expect this percentage of my users to authenticate** cell, click the cell and then use the drop-down arrows to select your estimated system utilization level, either **40%**, **60%** or **80%** for the deployment.
3. In the cell to the right of the **within the following period of time** cell, click the cell and then use the drop-down arrows to select either **1 minute**, **15 minutes**, or **1 hour** to select the duration of peak load.
4. In the cell to the right of the **Enter estimated number of internal applications (such as SharePoint (2007 or 2010) or claims aware web applications)** cell, type the number of internal applications you will use in your organization.
5. In the cell to the right of the **Enter estimated number of online applications (such as Office 365 Exchange Online, SharePoint Online or Lync Online)** cell, type the number of online applications or services you will used in your organization.
6. Under the cell titled **Number of Users**, type a number on each row that applies to an example application scenario your users will need single sign-on access to. This column should contain the number of defined users, not the peak users per second. If access

attempts made to the application must first go through the home realm discovery page, type **Y**. If you are unsure of this selection, type **Y**.

7. Review the following recommended values that are provided:
   a. For the total number of recommended federation servers, see the lower right cell that is highlighted in gray.
   b. For the number of servers recommended for each example application scenario, see the cell on the row that is highlighted in gray.

Note
The value that will be automatically calculated in the cell to the right of the cell titled **Total number of federation servers recommended** at the bottom of the spreadsheet contains a formula which will add an additional 20% buffer to the sum total of all the values in each of the individual rows preceding it. The formula added to the **Total number of federation servers recommended** cell builds in this buffer to your total recommended number of deployed federation servers to make it very unlikely that the overall load on the farm will ever hit its saturation point.

Planning for Federation Server Proxy Capacity

Capacity planning for federation server proxies helps you estimate:

- The appropriate hardware requirements for each federation server proxy.
- The number of federation servers and federation server proxies to place in each organization.

Federation server proxies redirect security tokens from a protected federation server in the corporate network to federated users. The purpose of deploying a federation server proxy is to allow external users to connect to a federation server. It does not actually sign tokens or write to data in the AD FS configuration database. Therefore, the hardware requirements for the federation server proxy are usually lower than the hardware requirements for a federation server.

Because every request to a federation server proxy results in a request to a federation server or federation server farm, capacity planning for federation servers and federation server proxies must be performed in parallel.

Estimating the peak sign-ins per second for the federation server proxy requires an understanding of the usage patterns of the federated users that will be signing in through the federation server proxy. In many deployments, the federated users who sign in using the federation server proxy are located on the Internet. You can estimate the peak sign-ins per second by looking at the usage patterns of these federated users on the existing Web applications that will be protected by AD FS.

Note
For production deployments, we recommend a minimum of two federation server proxies for each federation server farm instance you deploy.

Estimate the number of federation server proxies required for your organization

Before you can estimate the number of AD FS federation server proxy machines required, you will first need to determine the total number of federation servers that you will deploy in your organization. For more information about how to do this, see Planning for Federation Server Capacity.

Once you have decided on the number of federation servers, multiply this number of servers by the percentage of incoming federated authentication requests that you expect will be made from external users (located outside of the corporate network). The value of this calculation will provide you with the estimated number of federation server proxies that will handle the incoming authentication requests for your external users.

For example, if the number of recommended federation servers is 3, and you expect that the total number of authentication requests that will be made from external users will be approximately 60% of the total number of federated authentication requests, your calculation would equal 1.8 (3 X .60) which you can round up to 2. Therefore, in this case, you would need to deploy two federation server proxy machines to accommodate the load of external user authentication requests for the three federation servers.

In tests performed by the AD FS product team, the overall CPU utilization on each federation server proxy was found to be significantly lower than the CPU utilization that was observed on the federation servers for the same farm. In one test, while one federation server CPU was indicating that it was completely saturated, the CPU for a federation server proxy providing proxy services for that same farm was observed at only 20% utilization. Therefore, our tests revealed that the load on the CPU of a federation server proxy, which uses similar hardware specifications as discussed earlier in this section, could reasonably handle the processing load for approximately three federation servers.

However, for fault tolerance purposes, we recommend a minimum of two federation server proxies for each federation server farm you deploy.

Appendix A: Reviewing AD FS Requirements

So that the organizational partners in your Active Directory Federation Services (AD FS) deployment can collaborate successfully, you must first make sure that your corporate network infrastructure is configured to support AD FS requirements for accounts, name resolution, and certificates. AD FS has the following types of requirements:

Tip

You can find additional AD FS resource links at the AD FS Content Map page on the Microsoft TechNet Wiki. This page is managed by members of the AD FS Community and is monitored on a regular basis by the AD FS Product Team.

Hardware requirements

The following minimum and recommended hardware requirements apply to the federation server and federation server proxy computers.

Hardware requirement	Minimum requirement	Recommended requirement
CPU speed	Single-core, 1 gigahertz (GHz)	Quad-core, 2 GHz
RAM	1 GB	4 GB
Disk space	50 MB	100 MB

Software requirements

AD FS relies on server functionality that is built into the Windows Server® 2012 operating system.

Note
The Federation Service and Federation Service Proxy role services cannot coexist on the same computer.

Certificate requirements

Certificates play the most critical role in securing communications between federation servers, federation server proxies, claims-aware applications, and Web clients. The requirements for certificates vary, depending on whether you are setting up a federation server or federation server proxy computer, as described in this section.

Federation server certificates

Federation servers require the certificates in the following table.

Certificate type	Description	What you need to know before deploying
Secure Sockets Layer (SSL) certificate	This is a standard Secure Sockets Layer (SSL) certificate that is used for securing communications between federation servers and clients.	This certificate must be bound to the Default Web Site in Internet Information Services (IIS) for a Federation Server or a Federation Server Proxy. For a Federation Server Proxy, the binding must be configured in IIS prior to running the Federation Server Proxy Configuration

Certificate type	Description	What you need to know before deploying
		Wizard successfully. **Recommendation:** Because this certificate must be trusted by clients of AD FS, use a server authentication certificate that is issued by a public (third-party) certification authority (CA), for example, VeriSign. **Tip** The Subject name of this certificate is used to represent the Federation Service name for each instance of AD FS that you deploy. For this reason, you may want to consider choosing a Subject name on any new CA-issued certificates that best represents the name of your company or organization to partners.
Service communication certificate	This certificate enables WCF message security for securing communications between federation servers.	By default, the SSL certificate is used as the service communications certificate. This can be changed using the AD FS Management console.
Token-signing certificate	This is a standard X509 certificate that is used for securely signing all tokens that the federation server issues.	The token-signing certificate must contain a private key, and it should chain to a trusted root in the Federation Service. By default, AD FS creates a self-signed certificate. However, you can change this later to a CA-issued certificate by using the AD FS Management snap-in, depending on the needs of your

Certificate type	Description	What you need to know before deploying
		organization.
Token-decryption certificate	This is a standard SSL certificate that is used to decrypt any incoming tokens that are encrypted by a partner federation server. It is also published in federation metadata.	By default, AD FS creates a self-signed certificate. However, you can change this later to a CA-issued certificate by using the AD FS Management snap-in, depending on the needs of your organization.

Caution
Certificates that are used for token-signing and token-decrypting are critical to the stability of the Federation Service. Because a loss or unplanned removal of any certificates that are configured for this purpose can disrupt service, you should back up any certificates that are configured for this purpose.

For more information about the certificates that federation servers use, see Certificate Requirements for Federation Servers.

Federation server proxy certificates

Federation server proxies require the certificates in the following table.

Certificate type	Description	What you need to know before deploying
Server authentication certificate	This is a standard Secure Sockets Layer (SSL) certificate that is used for securing communications between a federation server proxy and Internet client computers.	This certificate must be bound to the Default Web Site in Internet Information Services (IIS) before you can run the AD FS Federation Server Proxy Configuration Wizard successfully. **Recommendation:** Because this certificate must be trusted by clients of AD FS, use a server authentication certificate that is issued by a public (third-party) certification authority (CA), for example, VeriSign. **Tip:** The Subject name of this

Certificate type	Description	What you need to know before deploying
		certificate is used to represent the Federation Service name for each instance of AD FS that you deploy. For this reason, you may want to consider choosing a Subject name that best represents the name of your company or organization to partners.

For more information about the certificates that federation server proxies use, see Certificate Requirements for Federation Server Proxies.

Browser requirements

Although any current Web browser with JavaScript capability can be made to work as an AD FS client, the Web pages that are provided by default have been tested only against Internet Explorer versions 7.0, 8.0 and 9.0, Mozilla Firefox 3.0, and Safari 3.1 on Windows. JavaScript must be enabled, and cookies must be enabled for browser-based sign-in and sign-out to work correctly.

The AD FS product team at Microsoft successfully tested the browser and operating system configurations in the following table.

Browser	Windows 7	Windows Vista
Internet Explorer 7.0	X	X
Internet Explorer 8.0	X	X
Internet Explorer 9.0	X	Not Tested
FireFox 3.0	X	X
Safari 3.1	X	X

Note
AD FS supports both the 32bit and 64bit versions of all the browsers showing in the above table.

Cookies

AD FS creates session-based and persistent cookies that must be stored on client computers to provide sign-in, sign-out, single sign-on (SSO), and other functionality. Therefore, the client browser must be configured to accept cookies. Cookies that are used for authentication are always Secure Hypertext Transfer Protocol (HTTPS) session cookies that are written for the originating server. If the client browser is not configured to allow these cookies, AD FS cannot function correctly. Persistent cookies are used to preserve user selection of the claims provider. You can disable them by using a configuration setting in the configuration file for the AD FS sign-in pages.

Support for TLS/SSL is required for security reasons.

Network requirements

Configuring the following network services appropriately is critical for successful deployment of AD FS in your organization.

TCP/IP network connectivity

For AD FS to function, TCP/IP network connectivity must exist between the client; a domain controller; and the computers that host the Federation Service, the Federation Service Proxy (when it is used), and the AD FS Web Agent.

DNS

The primary network service that is critical to the operation of AD FS, other than Active Directory Domain Services (AD DS), is Domain Name System (DNS). When DNS is deployed, users can use friendly computer names that are easy to remember to connect to computers and other resources on IP networks.

Windows Server 2008 uses DNS for name resolution instead of the Windows Internet Name Service (WINS) NetBIOS name resolution that was used in Windows NT 4.0–based networks. It is still possible to use WINS for applications that require it. However, AD DS and AD FS require DNS name resolution.

The process of configuring DNS to support AD FS varies, depending on whether:

- Your organization already has an existing DNS infrastructure. In most scenarios, DNS is already configured throughout your network so that Web browser clients in your corporate network have access to the Internet. Because Internet access and name resolution are requirements of AD FS, this infrastructure is assumed to be in place for your AD FS deployment.
- You intend to add a federated server to your corporate network. For the purpose of authenticating users in the corporate network, internal DNS servers in the corporate network forest must be configured to return the CNAME of the internal server that is running the Federation Service. For more information, see Name Resolution Requirements for Federation Servers.

- You intend to add a federated server proxy to your perimeter network. When you want to authenticate user accounts that are located in the corporate network of your identity partner organization, the internal DNS servers in the corporate network forest must be configured to return the CNAME of the internal federation server proxy. For information about how to configure DNS to accommodate the addition of federation server proxies, see Name Resolution Requirements for Federation Server Proxies.
- You are setting up DNS for a test lab environment. If you plan to use AD FS in a test lab environment where no single root DNS server is authoritative, it is probable that you will have to set up DNS forwarders so that queries to names between two or more forests will be forwarded appropriately. For general information about how to set up an AD FS test lab environment, see AD FS Step-by-Step and How To Guides (http://go.microsoft.com/fwlink/?LinkId=180357).

Attribute store requirements

AD FS requires at least one attribute store to be used for authenticating users and extracting security claims for those users. For a list of attribute stores that AD FS supports, see The Role of Attribute Stores in the AD FS Design Guide.

Note
AD FS automatically creates an Active Directory attribute store, by default.

Attribute store requirements depend on whether your organization is acting as the account partner (hosting the federated users) or the resource partner (hosting the federated application).

AD DS

For AD FS to operate successfully, domain controllers in either the account partner organization or the resource partner organization must be running Windows Server 2003 SP1, Windows Server 2003 R2, Windows Server 2008, or Windows Server 2012.

When AD FS is installed and configured on a domain-joined computer, the Active Directory user account store for that domain is made available as a selectable attribute store.

Important
Because AD FS requires the installation of Internet Information Services (IIS), we recommend that you not install the AD FS software on a domain controller in a production environment for security purposes. However, this configuration is supported by Microsoft Customer Service Support.

Schema requirements

AD FS does not require schema changes or functional-level modifications to AD DS.

Functional-level requirements

Most AD FS features do not require AD DS functional-level modifications to operate successfully. However, Windows Server 2008 domain functional level or higher is required for client certificate authentication to operate successfully if the certificate is explicitly mapped to a user's account in AD DS.

Service account requirements

If you are creating a federation server farm, you must first create a dedicated domain-based service account in AD DS that the Federation Service can use. Later, you configure each federation server in the farm to use this account. For more information about how to do this, see Manually Configure a Service Account for a Federation Server Farm in the AD FS Deployment Guide.

LDAP

When you work with other Lightweight Directory Access Protocol (LDAP)-based attribute stores, you must connect to an LDAP server that supports Windows Integrated authentication. The LDAP connection string must also be written in the format of an LDAP URL, as described in RFC 2255.

SQL Server

For AD FS to operate successfully, computers that host the Structured Query Language (SQL) Server attribute store must be running either Microsoft SQL Server 2005 or SQL Server 2008. When you work with SQL-based attribute stores, you also must configure a connection string.

Custom attribute stores

You can develop custom attribute stores to enable advanced scenarios. The policy language that is built into AD FS can reference custom attribute stores so that any of the following scenarios can be enhanced:

- Creating claims for a locally authenticated user
- Supplementing claims for an externally authenticated user
- Authorizing a user to obtain a token
- Authorizing a service to obtain a token on behavior of a user

When you work with a custom attribute store, you may also have to configure a connection string. In this situation, you can enter any custom code you like that enables a connection to your custom attribute store. The connection string in this situation is a set of name/value pairs that are interpreted as implemented by the developer of the custom attribute store.

For more information about developing and using custom attribute stores, see Attribute Store Overview (http://go.microsoft.com/fwlink/?LinkId=190782).

Application requirements

Federation servers can communicate with and protect federation applications, such as claims-aware applications.

Authentication requirements

AD FS integrates naturally with existing Windows authentication, for example, Kerberos authentication, NTLM, smart cards, and X.509 v3 client-side certificates. Federation servers use standard Kerberos authentication to authenticate a user against a domain. Clients can authenticate by using forms-based authentication, smart card authentication, and Windows Integrated authentication, depending on how you configure authentication.

The AD FS federation server proxy role makes possible a scenario in which the user authenticates externally using SSL client authentication. You can also configure the federation server role to require SSL client authentication, although typically the most seamless user experience is achieved by configuring the account federation server for Windows Integrated authentication. In this situation, AD FS has no control over what credentials the user employs for Windows desktop logon.

Smart card logon

Although AD FS can enforce the type of credentials that it uses for authentication (passwords, SSL client authentication, or Windows Integrated authentication), it does not directly enforce authentication with smart cards. Therefore, AD FS does not provide a client-side user interface (UI) to obtain smart-card personal identification number (PIN) credentials. This is because Windows-based clients intentionally do not provide user credential details to federation servers or Web servers.

Smart card authentication

Smart card authentication uses the Kerberos protocol to authenticate to an account federation server. AD FS cannot be extended to add new authentication methods. The certificate in the smart card is not required to chain up to a trusted root on the client computer. Use of a smart-card-based certificate with AD FS requires the following conditions:

- The reader and cryptographic service provider (CSP) for the smart card must work on the computer where the browser is located.
- The smart card certificate must chain up to a trusted root on the account federation server and the account federation server proxy.
- The certificate must map to the user account in AD DS by either of the following methods:
    - The certificate subject name corresponds to the LDAP distinguished name of a user account in AD DS.
    - The certificate subject altname extension has the user principal name (UPN) of a user account in AD DS.

To support certain authentication strength requirements in some scenarios, it is also possible to configure AD FS to create a claim that indicates how the user was authenticated. A relying party can then use this claim to make an authorization decision.

AD FS Deployment

- Windows Server 2012 R2 AD FS Deployment Guide
- Windows Server 2012 AD FS Deployment Guide

See Also

Active Directory Federation Services Overview

Windows Server 2012 R2 AD FS Deployment Guide

You can use Active Directory Federation Services (AD FS) with the Windows Server 2012 R2 operating system to build a federated identity management solution that extends distributed identification, authentication, and authorization services to Web-based applications across organization and platform boundaries. By deploying AD FS, you can extend your organization's existing identity management capabilities to the Internet.

- Deploying a Federation Server Farm
- Deploying Federation Server Proxies

See Also

AD FS Deployment

Deploying a Federation Server Farm

In order to deploy a federation server farm, complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic, return to this checklist after you review the conceptual topic so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploying a Federation Server Farm

	Task	Reference
☐	Review important concepts and considerations as you prepare to	AD FS Design Guide in

	Task	Reference
	deploy Active Directory Federation Services (AD FS). **Note**	Windows Server 2012 R2 Understanding Key AD FS Concepts
	If you decide to use Microsoft SQL Server for your AD FS configuration store, ensure to deploy a functional instance of SQL Server.	What links should we point people to?
☐	Join your computer to an Active Directory domain.	Join a Computer to a Domain
☐	Enroll a Secure Socket Layer (SSL) certificate for AD FS.	Enroll an SSL Certificate for AD FS
☐	Install the AD FS role service.	Install the AD FS Role Service
☐	Configure a federation server.	Configure a Federation Server
☐	Optional step: Configure a federation server with Device Registration Service (DRS).	Configure a federation server with Device Registration Service
☐	Add a host (A) and alias (CNAME) resource record to corporate Domain Name System (DNS) for the federation service and DRS.	Configure Corporate DNS for the Federation Service and DRS
☐	Verify that a federation server is operational.	Verify That a Federation Server Is Operational

See Also

Windows Server 2012 R2 AD FS Deployment Guide

Join a Computer to a Domain

For Active Directory Federation Services (AD FS) to function, each computer that you want to configure as a federation server must be joined to a domain.

Important
For AD FS running on Windows Server 2012 R2 to function, the Active Directory domain must run one of the following operating systems:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Important
If you plan to use a group Managed Service Account (gMSA) as the service account for AD FS, you must have at least one domain controller in your environment that is running on the Windows Server 2012 or the Windows Server 2012 R2 operating system.

Important
If you plan to deploy Device Registration Service (DRS) for AD Workplace Join as a part of your AD FS deployment, the AD FS schema must be updated to the Windows Server 2012 R2 level. There are three ways to update the schema:

1. In an existing Active Directory forest, run `adprep /forestprep` from the \Support\Adprep folder of the Windows Server 2012 R2 operating system DVD on any 64-bit server that runs Windows Server 2008 or later. In this case, no additional domain controller has to be installed, and no existing domain controllers must be upgraded.

   To run `adprep /forestprep`, you must be a member of the Schema Admins group, the Enterprise Admins group, and the Domain Admins group of the domain that hosts the schema master.
2. In an existing Active Directory forest, install a domain controller that runs Windows Server 2012 R2. In this case, `adprep /forestprep` can be run automatically as part of the domain controller installation.

   During the domain controller installation, you must provide additional credentials to run `adprep /forestprep`.
3. Create a new Active Directory forest by installing AD FS on a server that runs Windows Server 2012 R2. In this case, `adprep /forestprep` does not have to be run because the schema is initially created with all the necessary containers and objects to support DRS.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group

memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To join a computer to a domain

1. On the **Start** screen, type **Control Panel**, and then press Enter.
2. Navigate to **System and Security**, and then click **System**.
3. Under **Computer name, domain, and workgroup settings**, click **Change settings**.
4. On the **Computer Name** tab, click **Change**.
5. Under **Member of**, click **Domain**, type the name of the domain that this computer joins, and then click **OK**.
6. Click **OK**, and then restart the computer.

See Also

Deploying a Federation Server Farm

Enroll an SSL Certificate for AD FS

Active Directory Federation Services (AD FS) requires a certificate for Secure Socket Layer (SSL) server authentication on each federation server in your federation server farm. The same certificate can be used on each federation server in a farm. You must have both the certificate and its private key available. For example, if you have the certificate and its private key in a .pfx file, you can import the file directly into the Active Directory Federation Services Configuration Wizard. This SSL certificate must contain the following:

1. The subject name and subject alternative name must contain your federation service name, such as fs.contoso.com.
2. The subject alternative name must contain the value **enterpriseregistration** that is followed by the User Principal Name (UPN) suffix of your organization, for example, **enterpriseregistration.corp.contoso.com**.

Warning
Specify the subject alternative name if you plan to enable the Device Registration Service (DRS) for Workplace Join.

Important
If your organization uses multiple UPN suffixes, and you plan to enable the DRS, the SSL certificate must contain a subject alternative name entry for each suffix.

See Also

Deploying a Federation Server Farm

Install the AD FS Role Service

You can use the following procedure to install the AD FS role service on a computer that is running Windows Server 2012 R2 to become the first federation server in a federation server farm or a federation server in an existing federation server farm.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To install the AD FS server role via the Add roles and features wizard

1. Open Server Manager. To open Server Manager, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop. In the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or Feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **Select a server from the server pool**, verify that the target computer is selected, and then click **Next**.
5. On the **Select server roles** page, click **Active Directory Federation Services**, and then click **Next**.
6. On the **Select features** page, click **Next**. The required prerequisites are preselected for you. You do not have to select any other features.
7. On the **Active Directory Federation Service (AD FS)** page, click **Next**.
8. After you verify the information on the **Confirm installation selections** page, click **Install**.
9. On the **Installation progress** page, verify that everything installed correctly, and then click **Close**.

To install the AD FS server role via Windows PowerShell

1. On the computer that you want to configure as a federation server, open the Windows PowerShell command window, and then run the following command: `Install-windowsfeature adfs-federation –IncludeManagementTools`.

See Also

Deploying a Federation Server Farm

Configure a Federation Server

After you install the Active Directory Federation Services (AD FS) role service on your computer, you are ready to configure this computer to become a federation server. You can do one of the following:

- Configure the first federation server in a new federation server farm
- Add a federation server to an existing federation server farm

Configure the first federation server in a new federation server farm

To configure the first federation server in a new federation server farm by using the Active Directory Federation Service Configuration Wizard

Note

Ensure that you have domain administrator permissions or have domain administrator credentials available before you perform this procedure.

1. On the Server Manager **Dashboard** page, click the **Notifications** flag, and then click **Configure the federation service on the server**.

   The **Active Directory Federation Service Configuration Wizard** opens.
2. On the **Welcome** page, select **Create the first federation server in a federation server farm**, and then click **Next**.
3. On the **Connect to AD DS** page, specify an account by using domain administrator permissions for the Active Directory (AD) domain to which this computer is joined, and then click **Next**.
4. On the **Specify Service Properties** page, do the following, and then click **Next**:
   - Import the .pfx file that contains the Secure Socket Layer (SSL) certificate and key that you have obtained earlier. In Step 2: Enroll an SSL Certificate for AD FS, you have obtained this certificate and copied it onto the computer that you want to configure as a federation server. To import the .pfx file via the wizard, click **Import**, and then browse to the file's location. Enter the password for the .pfx file when you are prompted.

- Provide a name for your federation service. For example, **fs.contoso.com**. This name must match one of the subject or subject alternative names in the certificate.
- Provide a display name for your federation service. For example, **Contoso Corporation**. Users see this name on the Active Directory Federation Services (AD FS) sign-in page.

5. On the **Specify Service Account** page, specify a service account. You can either create or use an existing group Managed Service Account (gMSA) or use an existing domain user account. If you select the option to create a new gMSA account, specify a name for the new account. If you select the option to use an existing gMSA or domain account, click **Select** to select an account.

   Note
   The benefit of using a gMSA account is its auto-negotiated password update feature.

   Warning
   If you want to use a gMSA account, you must have at least one domain controller in your environment that is running the Windows Server 2012 operating system.

   If the gMSA option is disabled, and you see an error message, such as **Group Managed Service Accounts are not available because the KDS Root Key has not been set**, you can enable gMSA in your domain by running the following Windows PowerShell command on a domain controller, which runs Windows Server 2012 or later, in your Active Directory domain: `Add-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)`. Then return to the wizard, click **Previous**, and then click **Next** to re-enter the **Specify Service Account** page. The gMSA option should now be enabled. You can select it and enter a gMSA account name that you want to use.

6. On the **Specify Configuration Database** page, specify an AD FS configuration database, and then click **Next**. You can either create a database on this computer by using Windows Internal Database (WID), or you can specify the location and the instance name of Microsoft SQL Server.

   For more information, see The Role of the AD FS Configuration Database.

   Important
   If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

7. On the **Review Options** page, verify your configuration selections, and then click **Next**.
8. On the **Pre-requisite Checks** page, verify that all prerequisite checks are successfully completed, and then click **Configure**.
9. On the **Results** page, review the results and check whether the configuration is completed successfully, and then click **Next steps required for completing your federation service deployment**. For more information, see Next steps for completing your AD FS installation. Click **Close** to exit the wizard.

To configure the first federation server in a new federation server farm via Windows PowerShell

You can create a new federation server farm by using either a new or existing gMSA account or an existing domain user account.

- **If you want to create a new federation server by using a new gMSA account, do the following:**

  Important
  You must have domain administrator permissions to create the first federation server in a new federation server farm.

  a. On the computer that you want to configure as a federation server, ensure that the required SSL certificate has been imported into the **Local Computer\My Store** directory. You can verify whether the SSL certificate has been imported by running the following command in the Windows PowerShell command window: `dir Cert:\LocalMachine\My`. The certificate is listed by its thumbprint in the **Local Computer\My Store** directory.

  b. On your domain controller, open the Windows PowerShell command window and run the following command to verify whether the KDS Root Key has been created in your domain: `Get-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)`. If it has not been created so that the output displays no information, run the following command to create the key: `Add-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)`.

  c. On the computer that you want to configure as a federation server, open the Windows PowerShell command window, and run the following command:

  ```
  Install-AdfsFarm -CertificateThumbprint <certificate_thumbprint>
  -FederationServiceName <federation_service_name> -
  GroupServiceAccountIdentifier <domain>\<GMSA_Name>$
  ```

  Warning
  The `$` sign at the end of the previous command is required.

  To obtain the value for `<certificate_thumbprint>`, run `dir Cert:\LocalMachine\My`, and then select the thumbprint of your SSL certificate. The value of `<federation_service_name>` is the name of your federation service, for example, **fs.contoso.com**.

  Note
  If this is NOT the first time that you run this command, add the `OverwriteConfiguration` parameter.

  Notes
  The previous command creates a WID farm. If you want to create a SQL Server server farm, you must have an instance of SQL Server already installed and operational.

  You can use the following command to create the first federation server in a new farm that uses an instance of SQL Server: `Install-AdfsFarm -`

```
CertificateThumbprint <certificate_thumbprint> -FederationServiceName
<federation_service_name> -GroupServiceAccountIdentifier
<domain>\<GMSA_name>$ -SQLConnectionString "Data
Source=<SQL_Host_Name?\<SQL_instance_ name>;Integrated Security=True"
```

where **<SQL_Host_Name>** is the name of the server on which SQL Server is running, and **<SQL_instance_name>** is the name of the instance of SQL Server. If you use the default instance of SQL Server, use a **SQLConnectionString** value of "**Data Source=<SQL_Host_Name>;Integrated Security=True**".

Important

If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

- **If you want to create a new federation server by using an existing domain user account, do the following:**

  a. On the computer that you want to configure as a federation server, ensure that the required SSL certificate has been imported into the **Local Computer\My Store** directory. You can verify whether the SSL certificate has been imported by running the following command in the Windows PowerShell command window: `dir Cert:\LocalMachine\My`. The certificate is listed by its thumbprint in the **Local Computer\My Store** directory.

  b. On the computer that you want to configure as a federation server, open the Windows PowerShell command window, and then run the following command: `$fscred = Get-Credential`. Enter the domain user account credentials that you want to use for the federation service account in the format domain\user name.

  c. In the same Windows PowerShell command window, run the following command:

  ```
  Install-AdfsFarm -CertificateThumbprint <certificate_thumbprint>
  -FederationServiceName <federation_service_name> -
  ServiceAccountCredential $fscred
  ```

  To obtain the value for **<certificate_thumbprint>**, run `dir Cert:\LocalMachine\My`, and then select the thumbprint of your SSL certificate. The value of **<federation_service_name>** is the name of your federation service, for example, fs.contoso.com.

  Note

  If this is NOT the first time that you run this command, add the `OverwriteConfiguration` parameter.

  Notes

  The previous command creates a WID farm. If you want to create a SQL Server farm, you must have the instance of SQL Server already installed and operational.

  You can use the following command to create the first federation server in a new farm that uses an instance of SQL Server: `Install-AdfsFarm -`

`CertificateThumbprint <certificate_thumbprint> -FederationServiceName <federation_service_name> -ServiceAccountCredential $fscredential -SQLConnectionString "Data Source=<SQL_Host_Name>\<SQL_instance_name>;Integrated Security=True"` where **SQL_Host_Name** is the name of the server on which SQL Server is running, and **SQL_instance_name** is the name of the instance of SQL Server. If you use the default instance of SQL Server, use a **SQLConnectionString** value of "**Data Source=<SQL_Host_Name>;Integrated Security=True**".

Important
If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

Add a federation server to an existing federation server farm

Important
Ensure that you have completed Step 3: Install the AD FS Role Service, before you start any of the procedures in this section.

Important
Ensure that you have obtained a valid SSL server authentication certificate before you complete this procedure.

To add a federation server to an existing federation server farm via the Active Directory Federation Service Configuration Wizard

1. On the Server Manager **Dashboard** page, click the **Notifications** flag, and then click **Configure the federation service on the server**.

   The **Active Directory Federation Service Configuration Wizard** opens.
2. On the **Welcome** page, select **Add a federation server to a federation server farm**, and then click **Next**.
3. On the **Connect to AD DS** page, specify an account by using domain administrator permissions for the AD domain to which this computer is joined, and then click **Next**.
4. On the **Specify Farm** page, provide the name of the primary federation server in a farm that uses WID or specify the database host name and the database instance name of an existing federation server farm that uses SQL Server.

   Warning
   In Windows Server® 2012 R2, there is a workaround to specify the default

instance of SQL Server. The workaround is to not use the user interface. Instead, use the steps in To configure the first federation server in a new federation server farm via Windows PowerShell.

Important
If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

5. On the **Specify SSL Certificate** page, import the .pfx file that contains the SSL certificate and key that you have obtained previously. This certificate is the required service authentication certificate. In Step 2: Enroll an SSL Certificate for AD FS, you have obtained this certificate and copied it to the computer that you want to configure as a federation server. To import the .pfx file via the wizard, click **Import** and browse to the file's location. Enter the password for the .pfx file when you are prompted.
6. On the **Specify Service Account** page, specify the same service account that you configured when you created the first federation server in the farm. You can use an existing group Managed Service Account or an existing domain user account.

   Important
   The account that you specify must be the same account as the account that was used on the primary federation server in this farm.
7. On the **Review Options** page, verify your configuration selections, and then click **Next**.
8. On the **Pre-requisite Checks** page, verify that all prerequisite checks are successfully completed, and then click **Configure**.
9. On the **Results** page, review the results and check whether the configuration is completed successfully, and then click **Next steps required for completing your federation service deployment**. For more information, see Next steps for completing your AD FS installation. Click **Close** to exit the wizard.

To add a federation server to an existing federation server farm via Windows PowerShell

You can add a federation server to an existing farm by using either an existing gMSA account or an existing domain user account.

- If you want to join a federation server to a farm by using an existing gMSA account, do the following:
  a. On the computer that you want to configure as a federation server, ensure that the required SSL certificate has been imported into the **Local Computer\My Store** directory. You can verify whether the SSL certificate has been imported by running the following command in the Windows PowerShell command window: `dir Cert:\LocalMachine\My`. The certificate is listed by its thumbprint in the **Local Computer\My Store** directory.
  b. On the computer that you want to configure as a federation server, open the Windows PowerShell command window, and run the following command.

```
Add-AdfsFarmNode -GroupServiceAccountIdentifier
<domain>\<GMSA_name>$ -PrimaryComputerName
<first_federation_server_hostname> -CertificateThumbprint
<certificate_thumbprint>
```

`<domain>\<GMSA_name>` is your AD domain and the name of your gMSA account in that domain. `<first_federation_server_hostname>` is the host name of the primary federation server in this existing farm.

You can obtain the value for `<certificate_thumbprint>` by running `dir Cert:\LocalMachine\My` in the previous step.

Note

If this is NOT the first time that you run this command, add the `OverwriteConfiguration` parameter.

Notes

The previous command creates a WID farm node. If you want to create a server farm node of computers running SQL Server, you must have the instance of SQL Server already installed and operational.

You can use the following command to add a federation server to an existing farm that is using an instance of SQL Server: `Add-AdfsFarmNode -GroupServiceAccountIdentifier <domain>\<GMSA_name>$ -SQLConnectionString "Data Source=<SQL_Host_Name>\<SQL_instance_ name>;Integrated Security=True"` where **SQL_Host_Name** is the name of the server on which SQL Server is running, and **SQL_instance_name** is the name of the instance of SQL Server. If you use the default instance of SQL Server, use a **SQLConnectionString** value of "**Data Source=<SQL_Host_Name>;Integrated Security=True**".

Important

If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

- If you want to join a federation server to a farm by using an existing domain user account, do the following:
  a. On the computer that you want to configure as a federation server, open the Windows PowerShellcommand window, and then run the following command: `$fscred = get-credential`. Enter the domain user account credentials that you want to use for the federation service account in the format domain\user name.
  b. On the computer that you want to configure as a federation server, ensure that the required SSL certificate has been imported into the **Local Computer\My Store** directory. You can verify whether the SSL certificate has been imported by running the following command in the Windows PowerShellcommand window: `dir Cert:\LocalMachine\My`. The certificate is listed by its thumbprint in the **Local Computer\My Store** directory.

c. In the same Windows PowerShell command window, run the following command.

```
Add-AdfsFarmNode -ServiceAccountCredential $fscred -
PrimaryComputerName <first_federation_server_hostname> -
CertificateThumbprint <certificate_thumbprint>
```

Note
If this is NOT the first time that you run this command, add the `OverwriteConfiguration` parameter.

Note
The previous command creates a WID farm node. If you want to create a server farm node of computers running SQL Server, you must have the instance of SQL Server already installed and operational. You can use the following command to add a federation server to an existing farm by using an instance of SQL Server: `Add-AdfsFarmNode -ServiceAccountCredential $fscred -SQLConnectionString "Data Source=<SQL_Host_Name>\<SQL_instance_ name>;Integrated Security=True"` where **SQL_Host_Name** is the name of the server on which the instance of SQL Server is running, and **SQL_instance_name** is the name of the instance of SQL Server. If you use the default instance of SQL Server, use a **SQLConnectionString** value of "**Data Source=<SQL_Host_Name>;Integrated Security=True**".

Important
If you want to create an AD FS farm and use SQL Server to store your configuration data, you can use SQL Server 2008 and newer versions, including SQL Server 2012.

See Also

Deploying a Federation Server Farm

Configure a federation server with Device Registration Service

You can enable Device Registration Service (DRS) on your federation server after you complete the procedures in Step 4: Configure a Federation Server. The Device Registration Service provides an onboarding mechanism for seamless second factor authentication, persistent single sign-on (SSO), and conditional access to consumers that require access to company resources. For more information about DRS, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications

Prepare your Active Directory forest to support devices

Note

This is a one-time operation that you must run to prepare your Active Directory forest to support devices. You must be logged on with enterprise administrator permissions and your Active Directory forest must have the Windows Server 2012 R2 schema to complete this procedure.

To prepare the Active Directory forest

1. On your federation server, open a Windows PowerShell command window and type:

```
Initialize-ADDeviceRegistration
```

2. When prompted for ServiceAccountName, enter the name of the service account you selected as the service account for AD FS. If it is a gMSA account, enter the account in the **domain\accountname$** format. For a domain account, use the format **domain\accountname**.

Enable Device Registration Service on a federation server farm node

Note

You must be logged on with domain administrator permissions to complete this procedure.

To enable Device Registration Service

1. On your federation server, open a Windows PowerShell command window and type:

```
Enable-AdfsDeviceRegistration
```

2. Repeat this step on each federation farm node in your AD FS farm..

Enable seamless second factor authentication

Seamless second factor authentication is an enhancement in AD FS that provides an added level of access protection to corporate resources and applications from external devices that are trying to access them. When a personal device is Workplace Joined, it becomes a ‘known’ device and administrators can use this information to drive conditional access and gate access to resources.

To enable seamless second factor authentication, persistent single sign-on (SSO) and conditional access for Workplace Joined devices

1. In the AD FS Management console, navigate to Authentication Policies. Select Edit Global Primary Authentication. Select the check box next to Enable Device Authentication, and then click OK.

Update the Web Application Proxy configuration

Important
You do not need to publish the Device Registration Service to the Web Application Proxy. The Device Registration Service will be available through the Web Application Proxy once it is enabled on a federation server. You may need to complete this procedure to update the Web Application Proxy configuration if it was deployed prior to enabling the Device Registration Service.

To update the Web Application Proxy Configuration

1. On your Web Application Proxy server, open a Windows PowerShell command window and type

```
Update-WebApplicationProxyDeviceRegistration
```

2. When prompted for credentials, enter the credentials of an account that has administrative rights to your federation servers.

See Also

Deploying a Federation Server Farm

Configure Corporate DNS for the Federation Service and DRS

Step 6: Add a Host (A) and Alias (CNAME) Resource Record to Corporate DNS for the Federation Service and DRS

You must add the following resource records to corporate Domain Name System (DNS) for your federation service and Device Registration Service that you configured in previous steps.

Entry	Type	Address
federation_service_name	Host (A)	IP address of the AD FS server or the IP

Entry	Type	Address
		address of the load balancer that is configured in front of your AD FS server farm
enterpriseregistration	Alias (CNAME)	federation_server_name.contoso.com

You can use the following procedure to add a host (A) and alias (CNAME) resource records to corporate DNS for the federation server and the Device Registration Service.

Membership in **Administrators**, or equivalent, is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a host (A) and alias (CNAME) resource records to DNS for your federation server

1. On you domain controller, in Server Manager, on the **Tools** menu, click **DNS** to open the DNS snap-in.
2. In the console tree, expand the **domain_controller_name** node, expand **Forward Lookup Zones**, right-click **domain_name**, and then click **New Host (A or AAAA)**.
3. In the **Name** box, type the name to use for your AD FS farm.
4. In the **IP address** box, type the IP address of your federation server. Click **Add Host**.
5. Right-click the **domain_name** node, and then click **New Alias (CNAME)**.
6. In the **New Resource Record** dialog box, type **enterpriseregistration** in the **Alias name** box.
7. In the fully qualified domain name (FQDN) of the target host box, type **federation_service_farm_name.domain_name.com**, and then click **OK**.

   Important
   In a real world deployment, if your company has multiple User Principal Name (UPN) suffixes, you must create multiple CNAME records for each of those UPN suffixes in DNS.

See Also

Deploying a Federation Server Farm

Verify That a Federation Server Is Operational

You can use the following procedure to verify that a federation server is operational; that is, that any client on the same network can reach your federation server.

Membership in **Users**, **Backup Operators**, **Power Users**, **Administrators**, or equivalent on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To verify that a federation server is operational

1. Open a browser window and in the address bar, type the federation server name, and then append it with **federationmetadata/2007-06/federationmetadata.xml** to browse to the federation service metadata endpoint. For example, **https://fs.contoso.com/federationmetadata/2007-06/federationmetadata.xml**.

   In your browser window, if you can see the federation server metadata without any Secure Socket Layer (SSL) errors or warnings, your federation server is operational.
2. You can also browse to the AD FS sign-in page where your federation service name is appended with **adfs/ls/idpinitiatedsignon.htm**, for example, **https://fs.contoso.com/adfs/ls/idpinitiatedsignon.htm**. This entry displays the AD FS sign-in page where you can sign in by using domain administrator credentials.

   Important
   Ensure to configure your browser settings to trust the federation server role by adding your federation service name, for example, **https://fs.contoso.com**, to the browser's local intranet zone.

See Also

Deploying a Federation Server Farm

Deploying Federation Server Proxies

In Active Directory Federation Services (AD FS) in Windows Server 2012 R2, the role of a federation server proxy is handled by a new Remote Access role service called Web Application Proxy. To enable your AD FS for accessibility from outside the corporate network, which was the purpose of deploying a federation server proxy in legacy versions of AD FS, such as AD FS 2.0 and AD FS in Windows Server 2012, you can deploy one or more web application proxies for AD FS in Windows Server 2012 R2.

In the context of AD FS, Web Application Proxy functions as an AD FS federation server proxy. In addition to this, Web Application Proxy provides reverse proxy functionality for web applications inside your corporate network to enable users on any device to access them from outside the corporate network. For more information, about the Web Application Proxy role service, see Web Application Proxy Overview.

To plan the deployment of Web Application proxy, you can review the information in the following topics:

- Step 1: Plan the Web Application Proxy Infrastructure
- Step 2: Plan the Web Application Proxy Server

To deploy Web Application proxy, you can follow the procedures in the following topics:

- Step 1: Configure the Web Application Proxy Infrastructure
- Step 2: Install and Configure the Web Application Proxy Server

See Also

Windows Server 2012 R2 AD FS Deployment Guide

Windows Server 2012 AD FS Deployment Guide

You can use Active Directory® Federation Services (AD FS) with the Windows Server® 2012 operating system to build a federated identity management solution that extends distributed identification, authentication, and authorization services to Web-based applications across organization and platform boundaries. By deploying AD FS, you can extend your organization's existing identity management capabilities to the Internet.

You can deploy AD FS to:

- Provide your employees or customers with a Web-based, single-sign-on (SSO) experience when they need remote access to internally hosted Web sites or services.
- Provide your employees or customers with a Web-based, SSO experience when they access cross-organizational Web sites or services from within the firewalls of your network.
- Provide your employees or customers with seamless access to Web-based resources in any federation partner organization on the Internet without requiring employees or customers to log on more than once.
- Retain complete control over your employee or customer identities without using other sign-on providers (Windows Live ID, Liberty Alliance, and others).

About this guide

This guide is intended for use by system administrators and system engineers. It provides detailed guidance for deploying an AD FS design that has been preselected by you or an infrastructure specialist or system architect in your organization.

If a design has not yet been selected, we recommend that you wait to follow the instructions in this guide until after you have reviewed the design options in the AD FS Design Guide in Windows Server 2012 and you have selected the most appropriate design for your organization. For more information about using this guide with a design that has already been selected, see Implementing Your AD FS Design Plan.

After you select your design from the design guide and gather the required information about claims, token types, attribute stores, and other items, you can use this guide to deploy your AD FS design in your production environment. This guide provides steps for deploying either of the following primary AD FS designs:

- Web SSO
- Federated Web SSO

Use the checklists in Implementing Your AD FS Design Plan to determine how best to use the instructions in this guide to deploy your particular design. For information about hardware and software requirements for deploying AD FS, see the Appendix A: Reviewing AD FS Requirements in the AD FS Design Guide.

What this guide does not provide

This guide does not provide:

- Guidance regarding when and where to place federation servers, federation server proxies, or Web servers in your existing network infrastructure. For this information, see Planning Federation Server Placement and Planning Federation Server Proxy Placement in the AD FS Design Guide.
- Guidance for using certification authorities (CAs) to set up AD FS
- Guidance for setting up or configuring specific Web-based applications
- Setup instructions that are specific to setting up a test lab environment.
- Information about how to customize federated logon screens, web.config files, or the configuration database.

In this guide

- Planning to Deploy AD FS
- Implementing Your AD FS Design Plan
- Checklist: Implementing a Web SSO Design
- Checklist: Implementing a Federated Web SSO Design
- Configuring Partner Organizations
- Configuring Claim Rules
- Deploying Federation Servers
- Deploying Federation Server Proxies
- Interoperating with AD FS 1.x

Acknowledgments

Author: Nick Pierson

Technical Reviewers: Matt Steele, Lu Zhao (Migration)

Editor: Jim Becker

Planning to Deploy AD FS

After you collect information about your environment and you decide on an Active Directory Federation Services (AD FS) design by following the guidance in the AD FS Design Guide in Windows Server 2012, you can begin to plan the deployment of your organization's AD FS design. With the completed design and the information in this topic, you can determine which tasks to perform to deploy AD FS in your organization.

Reviewing your AD FS design

If the design team that constructed the original AD FS design for your organization is different from the deployment team that will actually implement the deployment, make sure that the deployment team reviews the final design with the design team. Review the following points regarding the design:

- The design team's strategy to determine the best physical topology for the placement of federation servers in your corporate network or perimeter network. The deployment team can refer to documentation about this subject by reviewing the following topics in the AD FS Design Guide:
  - The Role of the AD FS Configuration Database
  - Planning Federation Server Placement
  - Planning Federation Server Proxy Placement

  It is possible that the design team might leave the subject of federation server or federation server proxy placement for the deployment team. The deployment team is then responsible for documenting and implementing the physical topology of the servers.
- The business reasons for your organization's designation as a claims provider, relying party, or both, within the scope of the documented AD FS design. Ensure that members of the deployment team understand the reasons why AD FS is being deployed and what other companies or organizations are involved in the federation partnership. Ensure that members of the deployment team also understand the constraints that exist for the other companies or organizations (limited hardware, no extranet environment, and so forth) that may limit the scope of the design in some way. For more information about partner organizations, see Planning Your Deployment.

After the design teams and deployment teams agree on these issues, they can proceed with the deployment of the AD FS design. For more information, see Implementing Your AD FS Design Plan.

Implementing Your AD FS Design Plan

The following environmental conditions and requirements are important factors in the implementation of your Active Directory Federation Services (AD FS) design plan:

- **Supported partners:** You usually use AD FS to work with partner organizations. To establish identity federation, determine the organizations with which you want to form a partnership. After a baseline AD FS deployment is in place, operating with partners involves adding partners, deleting partners, and updating partner information. Changes to partnerships may occur for a variety of reasons. For example, your AD FS deployment might require partnership updates if your partner changes its business significantly, your organization becomes part of a larger organization or a federation of organizations, or your organization is acquired by a different company. In any scenario in which you federate identities from multiple domains, you will need to know the domains (partners) that you are currently supporting and all the additional domains that represent potential partners.
- **Supported application and service types:** Some applications and services require access to operating system resources, while others are "claims aware." It is important to understand the types of applications and services that AD FS supports so that you can formulate administration requirements.
- **Logical and physical architectural diagrams or deployment topology:** You will need to know:
  - Whether federation servers will function in a set of farmed servers or on a single server.
  - Where your network deploys firewalls and proxies.
  - The location of resources and whether users are accessing resources from within your organization, outside the organization, or both.

How to implement your AD FS design using this guide

The next step in implementing your design is to determine in what order each deployment task must be performed. This guide uses checklists to help you walk through the various server and application deployment tasks that are required to implement your design plan. Parent and child checklists are used as necessary to represent the order in which tasks for a specific AD FS design must be processed.

Use the following parent checklists in this section of the guide to become familiar with the deployment tasks for implementing your organization's preferred AD FS design:

- Checklist: Implementing a Web SSO Design
- Checklist: Implementing a Federated Web SSO Design

Checklist: Implementing a Web SSO Design

This parent checklist includes cross-reference links to important concepts about the Web Single-Sign-On (SSO) design for Active Directory Federation Services (AD FS). It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Implementing a Web SSO Design

	Task	Reference
☐	Review important concepts about the Web SSO design and determine which AD FS deployment goals you can use to customize this design to meet the needs of your organization. **Note**	Web SSO Design Identifying Your AD FS Deployment Goals
☐	Review the hardware, software, certificate, Domain Name System (DNS), attribute store, and client requirements for deploying AD FS in your organization.	Appendix A: Reviewing AD FS Requirements
☐	According to your design plan, install one or more federation servers in the corporate network or in the perimeter network. **Note** The Web SSO design requires only one federation server to function successfully. A single federation server acts in both the claims provider role and the relying party role.	Checklist: Setting Up a Federation Server
☐	(Optional) Determine whether or not your organization needs a federation server proxy in the perimeter network.	Checklist: Setting Up a Federation Server Proxy
☐	Depending on your Web SSO design plan and how you intend to	Checklist: Configuring the Account Partner Organization

	Task	Reference
	use it, add the appropriate attribute store, relying party trusts, claims, and claim rules to the Federation Service.	
☐	If you are an administrator in the resource partner organization, claims-enable your Web browser application, Web service application, or Microsoft® Office SharePoint® Server application using WIF and the WIF SDK. **Note**	Windows Identity Foundation Windows Identity Foundation SDK

Checklist: Implementing a Federated Web SSO Design

This parent checklist includes cross-reference links to important concepts about the Federated Web Single-Sign-On (SSO) design for Active Directory Federation Services (AD FS). It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Implementing a Federated Web SSO Design

	Task	Reference
☐	Review important concepts about the Federated Web SSO design and determine which AD FS deployment goals you can use to customize this design to meet the needs of your organization.	Federated Web SSO Design Identifying Your AD FS Deployment Goals Planning Your Deployment

	Task	Reference
☐	Review the hardware, software, certificate, Domain Name System (DNS), attribute store, and client requirements for deploying AD FS in your organization.	Appendix A: Reviewing AD FS Requirements
☐	Review important concepts about claims, claim rules, attribute stores, and the AD FS configuration database before deploying AD FS in both partner organizations.	Understanding Key AD FS Concepts
☐	According to your design plan, install one or more federation servers in each partner organization. **Note** For the Federated Web SSO design, you need at least one federation server in the account partner organization and at least one federation server in the resource partner organization.	Checklist: Setting Up a Federation Server
☐	(Optional) Determine whether or not your organization needs a federation server proxy. If your design plan calls for a proxy, you can install one or more federation server proxies in each partner organization.	Checklist: Setting Up a Federation Server Proxy
☐	According to your design plan, share certificates, configure clients, and configure the trust relationships in both partner organizations so that they can communicate over a federation trust.	Checklist: Configuring the Account Partner Organization Checklist: Configuring the Resource Partner Organization
☐	If you are an administrator in the resource partner organization, claims-enable your Web browser application, Web service application, or Microsoft® Office SharePoint® Server application using WIF and the WIF SDK.	Windows Identity Foundation Windows Identity Foundation SDK

Configuring Partner Organizations

To deploy a new partner organization in Active Directory Federation Services (AD FS), complete the tasks in either Checklist: Configuring the Resource Partner Organization or Checklist: Configuring the Account Partner Organization, depending on your AD FS design.

Note

When you use either of these checklists, we strongly recommend that you first read the references to account partner or resource partner planning guidance (in the AD FS Design Guide in Windows Server 2012) before continuing to the procedures for setting up the new partner organization. Following the checklist in this way will help provide a better understanding of the full AD FS design and deployment story for the account partner or resource partner organization.

About account partner organizations

An account partner is the organization in the federation trust relationship that physically stores user accounts in an AD FS–supported attribute store. The account partner is responsible for collecting and authenticating a user's credentials, building up claims for that user, and packaging the claims into security tokens. These tokens can then be presented across a federation trust to enable access to Web-based resources that are located in the resource partner organization.

In other words, an account partner represents the organization for whose users the account-side federation server issues security tokens. The federation server in the account partner organization authenticates local users and creates security tokens that the resource partner uses in making authorization decisions.

With regard to attribute stores, the account partner in AD FS is conceptually equivalent to a single Active Directory forest whose accounts need access to resources that are physically located in another forest. Accounts in this forest can access resources in the resource forest only when an external trust or forest trust relationship exists between the two forests and the resources to which the users are trying to gain access have been set with the proper authorization permissions.

About resource partner organizations

The resource partner is the organization in an AD FS deployment where Web servers are located. The resource partner trusts the account partner to authenticate users. Therefore, to make authorization decisions, the resource partner consumes the claims that are packaged in security tokens that come from users in the account partner.

In other words, a resource partner represents the organization whose Web servers are protected by the resource-side federation server. The federation server at the resource partner uses the

security tokens that are produced by the account partner to make authorization decisions for Web servers in the resource partner.

To function as an AD FS resource, Web servers in the resource partner organization must either have Windows Identity Foundation (WIF) installed or have the Active Directory Federation Services (AD FS) 1.x Claims-Aware Web Agent role services installed. Web servers that function as an AD FS resource can host either Web-browser-based or Web-service-based applications.

Checklist: Configuring the Account Partner Organization

The account partner organization contains the users that will access Web-based applications in the resource partner. Administrators in this organization must use the AD FS Management snap-in to create relying party trusts to represent their trust relationships with resource partner organizations. In turn, the resource partner administrator must create claims provider trusts for each account partner organization that they want to trust.

This checklist includes tasks for deploying Active Directory Federation Services (AD FS) in the account partner organization. It also includes tasks for configuring the components that are required to establish one-half of a federation partnership.

If you are deploying a Web SSO Design, you do not have to follow this checklist. However, you do have to complete the tasks in this checklist to successfully deploy a Federated Web SSO Design.

Important
Make sure that the administrator in the resource partner organization follows the guidance in Checklist: Configuring the Resource Partner Organization to ensure that all necessary deployment tasks will be completed to successfully create the second half of the federation partnership.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Configuring the account partner organization

	Task	Reference
☐	If you have an existing AD FS 1.0 or 1.1 deployment in your production environment today, see the link to the right for information about how to migrate settings from	**Planning a Migration to AD FS 2.0**

	Task	Reference
	your current Federation Service to a new AD FS Federation Service. If you are deploying AD FS for the first time in your organization using AD FS, you can skip this step and continue to the next task in this checklist for information about how to set up a new account partner organization.	
☐	Based on your deployment goals, review information about the components that are required to provide users with access to the federated applications.	Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services Provide Your Active Directory Users Access to the Applications and Services of Other Organizations Provide Users in Another Organization Access to Your Claims-Aware Applications and Services
☐	Determine which AD FS design this account partner organization will be associated with.	Web SSO Design Federated Web SSO Design
☐	Before you begin deploying your AD FS servers, review the; 1.) advantages and disadvantages of choosing either Windows Internal Database (WID) or SQL Server to store the AD FS configuration database 2.) AD FS deployment topology types and their associated server placement and network layout recommendations.	Determine Your AD FS Deployment Topology AD FS Deployment Topology Considerations
☐	Review AD FS capacity planning guidance to determine the proper number of federation server and federation server proxy servers you should use in your production	Planning for AD FS Server Capacity

	Task	Reference
	environment.	
☐	To effectively plan and implement the physical topology for the account partner deployment, determine whether your AD FS design requires one or more federation servers or federation server proxies.	Checklist: Setting Up a Federation Server Checklist: Setting Up a Federation Server Proxy
☐	Determine the type of attribute store that you want to add to AD FS. Then, add the attribute store using the AD FS Management snap-in.	The Role of Attribute Stores Add an Attribute Store
☐	If you will need to send claims to or consume claims from a resource partner who is using either an AD FS 1.0 or 1.1 Federation Service, see the link to the right for information about how to configure AD FS to interoperate with previous versions of AD FS. If the resource partner organization is also using AD FS to send or consume claims to your organization, you can skip this step and continue with the next task in this checklist.	Planning for Interoperability with AD FS 1.x
☐	After you deploy the first federation server in the account partner organization, create a relying party trust relationship using the AD FS Management snap-in. You can create a relying party trust by entering data about a resource partner manually or by using a federation metadata URL that the administrator of the resource partner organization provides to you. You can use the federation metadata to retrieve the data for	Create a Relying Party Trust Manually Create a Relying Party Trust Using Federation Metadata

	Task	Reference
	the resource partner automatically. **Note** If the resource partner publishes its federation metadata or can provide a file copy of it for you to use, we recommend that you retrieve the data automatically because it can save time.	
☐	Depending on the needs of your organization, create one or more claim rule sets for each relying party trust that is specified in the AD FS Management snap-in so that claims will be issued appropriately.	Checklist: Creating Claim Rules for a Relying Party Trust
☐	A claim description must be created if one does not already exist that will fulfill the needs of your organization. AD FS ships with a default set of claim descriptions that are exposed in the AD FS Management snap-in.	Add a Claim Description
☐	Determine whether your organization will need to use identity delegation to authorize or constrain a specified account to "act as" or impersonate other users. This is often a requirement when front-end Web applications must interact with back-end Web services.	When to Use Identity Delegation
☐	Prepare client computers for federation by: • Adding the URL for the account partner federation server to the trusted sites list	Prepare Client Computers in the Account Partner Configure Client Computers to Trust the Account Federation

	Task	Reference
	for the client browser. • Using Group Policy to push the appropriate Secure Sockets Layer (SSL) certificates to client computers.	Server Distribute Certificates to Client Computers by Using Group Policy

Checklist: Configuring the Resource Partner Organization

The resource partner organization contains the Web servers hosting the Web-based applications that will be accessed by users in the account partner. Administrators in this organization must use the AD FS Management snap-in to create claims provider trusts to represent their trust relationships with account partner organizations. In turn, the account partner administrator must create relying party trusts for each account partner organization that they want to trust.

This checklist includes the tasks that are necessary for deploying Active Directory Federation Services (AD FS) in the resource partner organization. It also includes tasks for configuring the components that are required to establish one-half of a federation partnership.

If you are deploying a Web SSO Design, you do not have to follow this checklist. However, you do have to complete the tasks in this checklist to successfully deploy a Federated Web SSO Design.

Important
Make sure that the administrator of the account partner organization follows the guidance in Checklist: Configuring the Account Partner Organization to ensure that all necessary deployment tasks will be completed to successfully create the second half of the federation partnership

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Configuring the resource partner organization

	Task	Reference
☐	If you have an existing AD FS 1.0 or 1.1 deployment in your production environment today, see the link to the right for information	**Planning a Migration to AD FS**

	Task	Reference
	about how to migrate settings from your current Federation Service to a new AD FS Federation Service. If you are deploying AD FS for the first time in your organization using AD FS, you can skip this step and continue to the next task in this checklist for information about how to set up a new resource partner organization.	
☐	Based on your deployment goals, review information about the components that are required to provide users with access to the federated applications that you host in your organization.	Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services Provide Your Active Directory Users Access to the Applications and Services of Other Organizations Provide Users in Another Organization Access to Your Claims-Aware Applications and Services
☐	Determine which AD FS design this resource partner organization will be associated with.	Web SSO Design Federated Web SSO Design
☐	Review the different application types, and decide which application to deploy.	Determine Your Federated Application Strategy in the Resource Partner
☐	Before you begin deploying your AD FS servers, review the; 1.) advantages and disadvantages of choosing either Windows Internal Database (WID) or SQL Server to store the AD FS configuration database 2.) AD FS deployment topology types and their associated server placement and network layout recommendations.	Determine Your AD FS Deployment Topology AD FS Deployment Topology Considerations
☐	Review AD FS capacity planning	Planning for AD FS Server

	Task	Reference
	guidance to determine the proper number of federation server and federation server proxy servers you should use in your production environment.	Capacity
☐	To effectively plan and implement the physical topology for the resource partner deployment, determine whether your AD FS design requires one or more federation servers, or federation server proxies.	Checklist: Setting Up a Federation Server Checklist: Setting Up a Federation Server Proxy
☐	Determine whether you need an attribute store in the resource partner organization. If you need an attribute store resource partner, add the attribute store with the AD FS Management snap-in.	The Role of Attribute Stores Add an Attribute Store
☐	If you will need to send claims to or consume claims from an account partner who is using either an AD FS 1.0 or 1.1 Federation Service, see the link to the right for information about how to configure AD FS to interoperate with previous versions of AD FS. If the account partner organization is also using AD FS to send or consume claims to your organization, you can skip this step and continue with the next task in this checklist.	Planning for Interoperability with AD FS 1.x
☐	After you deploy the first federation server in the resource partner organization, create a claims provider trust relationship by using the AD FS Management snap-in. You can create a claims provider trust by entering data about an account partner manually or by	Create a Claims Provider Trust Manually Create a Claims Provider Trust Using Federation Metadata

	Task	Reference
	using a federation metadata URL that the administrator of the account partner organization provides to you. You can use the federation metadata to retrieve the data for the resource partner automatically. **Note** If the account partner publishes its federation metadata or can provide a file copy of it for you to use, we recommend that you retrieve the data automatically because it can save time.	
☐	Depending on the needs of your organization, create one or more claim rule sets for each claims provider trust that is specified in the AD FS Management snap-in so that incoming claims will be passed through, transformed, or mapped appropriately to corresponding claims in the resource partner.	Checklist: Creating Claim Rules for a Claims Provider Trust
☐	(Optional) A claim description may have to be created if one does not already exist that will fulfill the needs of your organization. AD FS includes a default set of claim descriptions that are exposed in the AD FS Management snap-in.	Add a Claim Description

Add an Attribute Store

User accounts and computer accounts that require access to a resource that is protected by Active Directory Federation Services (AD FS) are stored in an attribute store, such as

Active Directory Domain Services (AD DS). The claims issuance engine uses attribute stores to gather data that is necessary to issue claims. Data from the attribute stores is then projected as claims.

You can use the following procedure to add an attribute store to the Federation Service.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add an attribute store

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under the **AD FS\Trust Relationships** folder, right-click **Attribute Stores**, and then click **Add an attribute store**.
3. In the **Add an attribute store** dialog box, configure the following properties for the attribute store that you want to add:
   - In **Display name**, type the name that you want to use to identify the attribute store.
   - In **Attribute store type**, select a supported attribute store type, either **Active Directory**, **LDAP**, or **SQL**.
   - In **Connection string**, if you have selected either a Lightweight Directory Access Protocol (LDAP) store or a Structured Query Language (SQL) store, enter the string that you used to establish a connection to the attribute store. For Active Directory attribute stores, no connection string is necessary; therefore, this field is disabled.

     Note
     AD FS automatically creates an Active Directory attribute store, by default.
4. Click **OK**.

Additional references

Checklist: Configuring the Account Partner Organization

The Role of Attribute Stores

Create a Claims Provider Trust Using Federation Metadata

To add a new claims provider trust, using the AD FS Management snap-in, by automatically importing configuration data about the partner from federation metadata that the partner has published to a local network or to the Internet, perform the following procedure on a federation server in the resource partner organization.

Caution

Though it has long been common practice to use certificates with unqualified host names such as https://myserver, these certificates have no security value and can enable an attacker to impersonate a Federation Service that is publishing federation metadata. Therefore, when querying federation metadata, you should only use a fully qualified domain name such as https://myserver.contoso.com.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a claims provider trust using federation metadata

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under **AD FS\Trust Relationships**, right-click **Claims Provider Trusts**, and then click **Add Claims Provider Trust** to open the Add Claims Provider Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Import data about the claims provider published online or on a local network**. In **Federation metadata address (host name or URL)**, type the federation metadata URL or host name for the partner, and then click **Next**.
5. On the **Specify Display Name** page type a **Display name**, under **Notes** type a description for this claims provider trust, and then click **Next**.
6. On the **Ready to Add Trust** page, click **Next** to save your claims provider trust information.
7. On the **Finish** page, click **Close**. This will automatically display the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this claims provider trust, see the Additional references section below.

   Important

   If trusted certificate stores have been modified previously on this computer, verify that the SSL certificate that is used to secure the federation metadata retrieval is trusted by the service account that is assigned to this Federation Service. If the service account does not trust the SSL certificate of this claims provider, monitoring of the trust will fail. To prevent this failure, make sure that the issuer of the claims provider's SSL certificate is in the Local Computer Trusted Root Certification Authorities certificate store on each federation server in the farm.

Additional references

Checklist: Configuring the Resource Partner Organization

Checklist: Creating Claim Rules for a Claims Provider Trust

Create a Claims Provider Trust Manually

To add a new claims provider trust using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a resource partner federation server in the resource partner organization.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a claims provider trust manually

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under the **AD FS\Trust Relationships**, right-click **Claims Provider Trusts**, and then click **Add Claims Provider Trust** to open the Add Claims Provider Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter claims provider trust data manually**, and then click **Next**.
5. On the **Specify Display Name** page type a **Display name**, under **Notes** type a description for this claims provider trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS Profile**, click **Next**, and then go to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then move to step 8.

   If you know that you will require interoperability between this claims provider trust and other, older Active Directory Federation Services (AD FS) claims provider trusts, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile** option.
7. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 9:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Claims provider WS-Federation Passive protocol URL**, type the URL for this claims provider trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Claims provider SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this claims provider trust, and then click **Next**.

   Click the **Help** button on this page for more information about which of these options apply to the needs of your organization.
8. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this claims provider trust, and then click **Next**.
9. On the **Configure Identifier** page, under **Claims provider trust identifier**, type the appropriate identifier, and then click **Next**.
10. On the **Configure Certificates** page, click **Add** to locate a certificate file and add it to the

list of certificates, and then click **Next**.

11. On the **Ready to Add Trust** page, click **Next** to save your claims provider trust information.
12. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this claims provider trust, see the Additional references.

Additional references

Checklist: Configuring the Resource Partner Organization

Checklist: Creating Claim Rules for a Claims Provider Trust

Create a Relying Party Trust Using Federation Metadata

To add a new relying party trust, using the AD FS Management snap-in, by automatically importing configuration data about the partner from federation metadata that the partner published to a local network or to the Internet, perform the following procedure on a federation server in the account partner organization.

Caution
Though it has long been common practice to use certificates with unqualified host names such as https://myserver, these certificates have no security value and can enable an attacker to impersonate a Federation Service that is publishing federation metadata. Therefore, when querying federation metadata, you should only use a fully qualified domain name such as https://myserver.contoso.com.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a relying party trust using federation metadata

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under the **AD FS\Trust Relationships** folder, right-click **Relying Party Trusts**, and then click **Add Relying Party Trust** to open the Add Relying Party Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Import data about the relying party published online or on a local network**. In **Federation metadata address (host name or URL)**, type the federation metadata URL or host name for the partner, and then click **Next**.
5. On the **Specify Display Name** page type a name in **Display name**, under **Notes** type a

description for this relying party trust, and then click **Next**.

6. On the **Choose Issuance Authorization Rules** page, select either **Permit all users to access this relying party** or **Deny all users access to this relying party**, and then click **Next**.
7. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.
8. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this relying party trust, see the Additional references.

Important
If trusted certificate stores have been modified previously on this computer, verify that the SSL certificate that is used to secure the federation metadata retrieval is trusted by the service account that is assigned to this Federation Service. If the service account does not trust the SSL certificate of this relying party, monitoring of the trust will fail. To prevent this failure, make sure that the issuer of the relying party's SSL certificate is in the Local Computer Trusted Root Certification Authorities certificate store on each federation server in the farm.

Additional references

Checklist: Configuring the Account Partner Organization

Checklist: Creating Claim Rules for a Relying Party Trust

Create a Relying Party Trust Manually

To add a new relying party trust using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a federation server in the account partner organization.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a relying party trust manually

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under **AD FS\Trust Relationships**, right-click **Relying Party Trusts**, and then click **Add Relying Party Trust** to open the Add Relying Party Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter data about the relying party manually**, and then click **Next**.

5. On the **Specify Display Name** page type a name in **Display name**, under **Notes** type a description for this relying party trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS profile**, click **Next**, and then move to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then go to step 9.

   If you know you will require interoperability with older Active Directory Federation Services (AD FS) federation, as provided in Windows Server 2003 R2, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile**.
7. On the **Configure Certificate** page, click **Browse** to locate a certificate file, and then click **Next**.
8. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 10:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Relying party WS-Federation Passive protocol URL**, type the URL for this relying party trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Relying party SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this relying party trust, and then click **Next**.

   Click the **Help** button on this page for more information about which of these options apply to the needs of your organization.
9. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this relying party trust, and then click **Next**.
10. On the **Configure Identifiers** page, specify one or more identifiers for this relying party, click **Add** to add them to the list, and then click **Next**.
11. On the **Choose Issuance Authorization Rules** page, select either **Permit all users to access this relying party** or **Deny all users access to this relying party**, and then click **Next**.
12. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.
13. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this relying party trust, see Additional references.

Additional references

Checklist: Configuring the Account Partner Organization

Checklist: Creating Claim Rules for a Relying Party Trust

Add a Claim Description

In an account partner organization, administrators create claims to represent a user's membership in a group or role or to represent some data about a user, for example, a user's employee identification number.

In a resource partner organization, administrators create corresponding claims to represent groups and users that can be recognized as resource users. Because outgoing claims in the account partner organization map to incoming claims in the resource partner organization, the resource partner is able to accept the credentials that the account partner provides.

You can use the following procedure to add a claim.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Click **AD FS\Service**, right-click **Claim Descriptions**, and then click **Add Claim Description**.
3. On the **Add a Claim Description** dialog box, in **Display name**, type a unique name that identifies the group or role for this claim.
4. In **Claim identifier**, type a URI that is associated with the group or role of the claim that you will be using.
5. Under **Description**, type text that best describes the purpose of this claim.
6. Depending on the needs of your organization, select either of the following check boxes, as appropriate, to publish this claim into federation metadata:
   - To publish this claim to make partners aware that this server can accept this claim, click **Publish this claim in federation metadata as a claim type that this Federation Service can accept**.
   - To publish this claim to make partners aware that this server can issue this claim, click **Publish this claim in federation metadata as a claim type that this Federation Service can send**.
7. Click **OK**.

Additional references

Checklist: Configuring the Account Partner Organization

Checklist: Configuring the Resource Partner Organization

The Role of Claims

Configure Client Computers to Trust the Account Federation Server

So that client computers can successfully access federated applications using Active Directory Federation Services (AD FS), you must first configure the Internet Explorer settings on each client computer so that the browser trusts the account federation server. You can do this manually or through Group Policy, depending on your administrative preference, by completing one of the following procedures.

Configuring Internet Explorer settings manually

You can use the following procedure to manually configure each user's Internet Explorer settings to support federation through AD FS. If multiple users use a single computer, complete this procedure multiple times—once for each user profile.

To perform this procedure, log on as the user who will be accessing federated applications. This is a profile-specific setting. Therefore, it requires that you manually add this setting for each profile that exists on a specific computer.

To manually configure client computers to trust the account federation server

1. On the client computer, start Internet Explorer.
2. On the **Tools** menu, click **Internet Options**.
3. On the **Security** tab, click the **Local intranet** icon, and then click **Sites**.
4. Click **Advanced**, and in **Add this Web site to the zone**, type the full Domain Name System (DNS) name of the account federation server (for example, https://fs1.fabrikam.com), and then click **Add**.
5. Click **OK** three times.

Configuring Internet Explorer settings by using Group Policy

For most deployments, we recommend that you use Group Policy to push the appropriate Internet Explorer settings to each client computer.

Membership in **Domain Admins** or **Enterprise Admins**, or equivalent, in Active Directory Domain Services (AD DS) is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure client computers to trust the account federation server by using Group Policy

1. On a domain controller in the forest of the account partner organization, start the **Group**

Policy Management snap-in.

2. Find the appropriate Group Policy Object (GPO), right-click it, and then click **Edit**.
3. In the console tree, open **User Configuration\Preferences\Windows Settings\Internet Explorer Maintenance**, and then click **Security**.
4. In the details pane, double-click **Security Zones and Content Ratings**.
5. Under **Security Zones and Privacy**, click **Import the current security zones and privacy settings**, and then click **Modify Settings**.
6. Click **Local intranet**, and then click **Sites**.
7. In **Add this Web site to the zone**, type the full DNS name of the account federation server (for example, https://fs1.fabrikam.com), click **Add**, and then click **Close**.
8. Click **OK** two times to apply these changes to Group Policy.

Distribute Certificates to Client Computers by Using Group Policy

You can use the following procedure to push down the appropriate Secure Sockets Layer (SSL) certificates (or equivalent certificates that chain to a trusted root) for account federation servers, resource federation servers, and Web servers to each client computer in the account partner forest by using Group Policy.

Membership in **Domain Admins** or **Enterprise Admins**, or equivalent, in Active Directory Domain Services (AD DS) is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To distribute certificates to client computers by using Group Policy

1. On a domain controller in the forest of the account partner organization, start the **Group Policy Management** snap-in.
2. Find an existing Group Policy Object (GPO) or create a new GPO to contain the certificate settings. Ensure that the GPO is associated with the domain, site, or organizational unit (OU) where the appropriate user and computer accounts reside.
3. Right-click the GPO, and then click **Edit**.
4. In the console tree, open **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**, right-click **Trusted Root Certification Authorities**, and then click **Import**.
5. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
6. On the **File to Import** page, type the path to the appropriate certificate files (for example, \\fs1\c$\fs1.cer), and then click **Next**.
7. On the **Certificate Store** page, click **Place all certificates in the following store**, and

then click **Next**.

8. On the **Completing the Certificate Import Wizard** page, verify that the information you provided is accurate, and then click **Finish**.
9. Repeat steps 2 through 6 to add additional certificates for each of the federation servers in the farm.

Configuring Claim Rules

To configure claim rules in your organization, complete each of the tasks in Checklist: Creating Claim Rules for a Claims Provider Trust or Checklist: Creating Claim Rules for a Relying Party Trust, depending on the federation role that your organization will play.

Note

When you use either of these checklists, we recommend that you first read the references to understanding claims issuance and claim rules concepts in the AD FS Design Guide in Windows Server 2012 before you begin the procedures for configuring claim rules. Following the checklist in this way provides a better understanding of the design and deployment process claim rules.

About Claim Rules

In a claims-based identity model, the function of Active Directory Federation Services (AD FS) as a federation service is to issue a token that contains a set of claims. The decision regarding what claims AD FS issues is governed by claim rules. Claim rules (and all server configuration dat) are stored in the AD FS configuration database.

AD FS makes issuance decisions based on identity information that is provided to it in the form of claims and other contextual information. At a high level, AD FSoperates as a rules processor by taking one set of claims as input, performing a number of transformations, and then returning a different set of claims as output.

Checklist: Creating Claim Rules for a Claims Provider Trust

This checklist includes tasks for planning, designing, and deploying claim rules that are associated with a claims provider trust in Active Directory Federation Services (AD FS).

Note

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Creating a claim rule set for a claims provider trust

	Task	Reference
☐	Review concepts about claims, claim rules, claim rule sets, and claim rule templates and how they are associated with federated trusts.	The Role of Claims The Role of Claim Rules
☐	Review concepts about how a claim flows through all the stages in the claims issuance pipeline and how rules are processed by the claims issuance engine.	The Role of the Claims Pipeline The Role of the Claims Engine
☐	To effectively plan and implement the output claims that will be issued over this claims provider trust, determine whether one or more claim rules are needed and which claim rules you should use with this claims provider trust.	Determine the Type of Claim Rule Template to Use
☐	Review concepts about when to create one claim rule over another and how you can use the claim rule language to provide more complex logic than standard rules in order to provide a desired result in the ideal output claim set.	When to Use a Pass Through or Filter Claim Rule When to Use a Transform Claim Rule When to Use a Send LDAP Attributes as Claims Rule When to Use a Send Group Membership as a Claim Rule When to Use a Custom Claim Rule The Role of the Claim Rule Language
☐	A claim description must be created if one does not already	Add a Claim Description

	Task	Reference
	exist that will fulfill the needs of your organization. AD FS ships with a default set of claim descriptions that are exposed in the AD FS Management snap-in.	
☐	Depending on the needs of your organization, create one or more claim rules for the acceptance transform rules set that is associated with this claims provider trust so that claims will be issued appropriately.	Create a Rule to Pass Through or Filter an Incoming Claim Create a Rule to Send LDAP Attributes as Claims Create a Rule to Send Group Membership as a Claim Create a Rule to Transform an Incoming Claim Create a Rule to Send an Authentication Method Claim Create a Rule to Send an AD FS 1.x Compatible Claim Create a Rule to Send Claims Using a Custom Rule

Checklist: Creating Claim Rules for a Relying Party Trust

This checklist includes the tasks that are necessary for planning, designing, and deploying claim rules that are associated with a relying party trust in Active Directory Federation Services (AD FS).

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Creating a claim rule set for a relying party trust

	Task	Reference
☐	Review concepts about claims, claim rules, claim rule sets, and claim rule templates and how they are associated with federated trusts.	The Role of Claims The Role of Claim Rules
☐	Review concepts about how a claim flows through all the stages in the claims issuance pipeline and how rules are processed by the claims issuance engine.	The Role of the Claims Pipeline The Role of the Claims Engine
☐	To effectively plan and implement the output claims that will be issued over this relying party trust, determine whether one or more claim rules are needed and which claim rules you should use with this relying party trust.	Determine the Type of Claim Rule Template to Use
☐	Review concepts about when to create one claim rule over another and how you can use the claim rule language to provide more complex logic than standard rules in order to provide a desired result in the ideal output claim set.	When to Use a Pass Through or Filter Claim Rule When to Use a Transform Claim Rule When to Use a Send LDAP Attributes as Claims Rule When to Use a Send Group Membership as a Claim Rule When to Use an Authorization Claim Rule When to Use a Custom Claim Rule The Role of the Claim Rule Language
☐	A claim description must be created if one does not already exist that will fulfill the needs of your organization. AD FS ships with a default set of claim	Add a Claim Description

	Task	Reference
	descriptions that are exposed in the AD FS Management snap-in.	
☐	Depending on the needs of your organization, create one or more claim rules for the rule sets that are associated with this relying party trust so that claims will be issued appropriately.	Create a Rule to Pass Through or Filter an Incoming Claim Create a Rule to Send LDAP Attributes as Claims Create a Rule to Send Group Membership as a Claim Create a Rule to Transform an Incoming Claim Create a Rule to Send an Authentication Method Claim Create a Rule to Send an AD FS 1.x Compatible Claim Create a Rule to Send Claims Using a Custom Rule
☐	Depending on the needs of your organization, create one or more claim rules for either the issuance authorization rules set or the delegation authorization rules set that is associated with this relying party trust so that users will be permitted access to the relying party.	Create a Rule to Permit All Users Create a Rule to Permit or Deny Users Based on an Incoming Claim

Create a Rule to Pass Through or Filter an Incoming Claim

Using the Pass Through or Filter an Incoming Claim rule template in Active Directory Federation Services (AD FS), you can pass through all incoming claims with a selected claim type. You can also filter the values of incoming claims with a selected claim type. For example, you can use this rule template to create a rule that will send all incoming group claims. You can also use this rule to send only user principal name (UPN) claims that end with @fabrikam.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to pass through or filter an incoming claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule, in **Incoming claim type** select a claim type in the list, and then select one of the following options, depending on the needs of your organization:
   - **Pass through all claim values**
   - **Pass through only a specific claim value**
   - **Pass through only claim values that match a specific email suffix value**
   - **Pass through only claim values that start with a specific value**
7. Click the **Finish** button.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Claims Provider Trust

Checklist: Creating Claim Rules for a Relying Party Trust

When to Use a Pass Through or Filter Claim Rule

The Role of Claims

The Role of Claim Rules

Create a Rule to Permit All Users

Using the **Permit All Users** rule template in Active Directory Federation Services (AD FS), you can create an authorization rule that will give all users access to the relying party. You can use additional authorization rules to further restrict access. Users who are permitted to access the relying party from the Federation Service may still be denied service by the relying party.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to permit all users

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships\Relying Party Trusts**, click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, click the **Issuance Authorization Rules** tab or the **Delegation Authorization Rules** tab (based on the type of authorization rule you require), and then click **Add Rule** to start the **Add Authorization Claim Rule Wizard**.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Permit All Users** from the list, and then click **Next**.
6. On the **Configure Rule** page, click **Finish**.
7. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Relying Party Trust

When to Use an Authorization Claim Rule

The Role of Claims

The Role of Claim Rules

Create a Rule to Permit or Deny Users Based on an Incoming Claim

Using the **Permit or Deny Users Based on an Incoming Claim** rule template in Active Directory Federation Services (AD FS), you can create an authorization rule that will grant or deny user's access to the relying party based on the type and value of an incoming claim. For example, you can use this rule template to create a rule that will permit only users that have a group claim with

a value of Domain Admins to access the relying party. If you want to permit all users to access the relying party, use the **Permit All Users** rule template. Users who are permitted to access the relying party from the Federation Service may still be denied service by the relying party.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to permit or deny users based on an incoming claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships\Relying Party Trusts**, click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, click the **Issuance Authorization Rules** tab or the **Delegation Authorization Rules** tab (based on the type of authorization rule you require), and then click **Add Rule** to start the **Add Authorization Claim Rule Wizard**.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule, in **Incoming claim type** select a claim type in the list, under **Incoming claim value** type a value or click Browse (if it is available) and select a value, and then select one of the following options, depending on the needs of your organization:
   - **Permit access to users with this incoming claim**
   - **Deny access to users with this incoming claim**
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Relying Party Trust

When to Use an Authorization Claim Rule

The Role of Claims

The Role of Claim Rules

Create a Rule to Send LDAP Attributes as Claims

Using the Send LDAP Attributes as Claims rule template in Active Directory Federation Services (AD FS), you can create a rule that will select attributes from a Lightweight Directory Access Protocol (LDAP) attribute store, such as Active Directory, to send as claims to the relying party. For example, you can use this rule template to create a Send LDAP Attributes as Claims rule that will extract attribute values for authenticated users from the **displayName** and **telephoneNumber** Active Directory attributes and then send those values as two different outgoing claims.

You can also use this rule to send all the user's group memberships. If you want to send only individual group memberships, use the Send Group Membership as a Claim rule template. You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send LDAP attributes as claims

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims** from the list, and then click **Next**.
6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule, under **Attribute store** select **Active Directory**, and under **Mapping of LDAP attributes to outgoing claim types** select the desired **LDAP Attribute** and corresponding **Outgoing Claim Type** types from the drop-down lists.

   You have to select a new LDAP attribute and outgoing claim type pair on a different row for each Active Directory attribute that you want to issue a claim for as part of this rule.
7. Click the **Finish** button.

8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Claims Provider Trust

Checklist: Creating Claim Rules for a Relying Party Trust

The Role of Claims

The Role of Claim Rules

Create a Rule to Send Group Membership as a Claim

Using the Send Group Membership as a Claim rule template in Active Directory Federation Services (AD FS), you can create a rule that will make it possible for you to select an Active Directory security group to send as a claim. Only a single claim will be emitted from this rule, based on the group that you select. For example, you can use this rule template to create a rule that will send a group claim with a value of Admin if the user is a member of the Domain Admins security group. This rule should be used only for users in the local Active Directory domain.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send group membership as a claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group**

Membership as a Claim from the list, and then click **Next**.

6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule, in **User's group** click **Browse** and select a group, under **Outgoing claim type** select the desired claim type, and then under **Outgoing Claim Type** type a value.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Claims Provider Trust

Checklist: Creating Claim Rules for a Relying Party Trust

The Role of Claims

The Role of Claim Rules

Create a Rule to Transform an Incoming Claim

Using the Transform an Incoming Claim rule template in Active Directory Federation Services (AD FS), you can select an incoming claim, change its claim type, and change its claim value. For example, you can use this rule template to create a rule that will send a role claim with the same claim value of an incoming group claim. You can also use this rule to send a group claim with a claim value of Purchasers when there is an incoming group claim with a value of Admins, or you can send only User Principal Name (UPN) claims that end with @fabrikam.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to transform an incoming claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**

- **Issuance Transform Rules**
- **Issuance Authorization Rules**
- **Delegation Authorization Rules**

5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule, in **Incoming claim type** select a claim type in the list, in **Outgoing claim type** select a claim type in the list, and then select one of the following options, depending on the needs of your organization:
   - **Pass through all claim values**
   - **Replace an incoming claim value with a different outgoing claim value**
   - **Replace incoming e-mail suffix claims with a new e-mail suffix**

   Notes

   If you are setting up the Dynamic Access Control scenario that uses AD DS-issued claims, first create a transform rule on the claims provider trust and in **Incoming claim type**, type the name for the incoming claim or if a claim description was previously created select it from the list. Second, in **Outgoing claim type**, select the desired claim URL, and then create a transform rule on the relying party trust to issue the device claim.

   For more information about Dynamic Access Control scenarios, see Dynamic Access Control Content Roadmap or Using AD DS Claims with AD FS.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Claims Provider Trust

Checklist: Creating Claim Rules for a Relying Party Trust

When to Use a Transform Claim Rule

The Role of Claims

The Role of Claim Rules

Create a Rule to Send an Authentication Method Claim

You can use either the **Send Group Membership as Claims** rule template or the **Transform an Incoming Claim** rule template to send an authentication method claim. The relying party can use an authentication method claim to determine the logon mechanism that the user uses to authenticate and obtain claims from Active Directory Federation Services (AD FS). You can also use the Authentication Mechanism Assurance feature of Active Directory Domain Services (AD DS) in Windows Server 2008 R2 as input to generate authentication method claims for situations in which the relying party wants to determine the level of access based on smart-card logons. For example, a developer can assign different levels of access to federated users of the relying party application based on whether the users log on with their user name/password credentials, as opposed to their smart cards.

Note

If you do not write an explicit rule to generate an authentication method claim, AD FS will generate one automatically based on the logon method that the user uses to authenticate originally. For more information about the Authentication Mechanism Assurance feature, see What's New in AD DS: Authentication Mechanism Assurance (http://go.microsoft.com/fwlink/?LinkID=159947).

Depending on the needs of your organization, use one of the following procedures:

- Create this rule using the Send Group Membership as Claims rule template - You can use this rule template when you want the group that you specify in this template to ultimately determine what authentication method claim to issue.
- Create this rule using the Transform an Incoming Claim rule template - You can use this rule template in situations in which you want to change the existing authentication method to a new authentication method that will work with a product that does not understand standard AD FS authentication method claims.

To create this rule using the Send Group Membership as Claims rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**

- **Issuance Transform Rules**
- **Issuance Authorization Rules**
- **Delegation Authorization Rules**

5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. Click **Browse**, select the group whose members should receive this authentication method claim, and then click **OK**.
8. In **Outgoing claim type**, select **Authentication method** in the list.
9. In **Outgoing claim value**, type one of the default Uniform Resource Identifier (URI) values in the following table, depending on your preferred authentication method, click **Finish**, and then click **OK** to save the rule.

Actual Authentication method	Corresponding URI
Username Password authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/password
Windows authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/windows
TLS Mutual authentication using X.509 certificates	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/tlsclient
X.509-based authentication that does not use TLS	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/x509

Note
Other URI values can be used in addition to the values in the table. The URI values shown here reflect the URIs that the relying party will accept by default.

To create this rule using the Transform an Incoming Claim rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select **Authentication method** in the list.
8. In **Outgoing claim type**, select **Authentication method** in the list.
9. Select **Replace an incoming claim value with a different outgoing claim value**, and then do the following:
   a. In **Incoming claim value**, type one of the following URI values based on the actual authentication method URI that was used originally, click **Finish**, and then click **OK** to save the rule.
   b. In **Outgoing claim value**, type one of the default URI values in the following table, depending on your new preferred authentication method choice, click **Finish**, and then click **OK** to save the rule.

Actual Authentication method	Corresponding URI
Username Password authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/password
Windows authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/windows
TLS Mutual authentication using X.509	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/tlsclient

certificates	
X.509-based authentication that does not use TLS	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/x509

Note
Other URI values can be used in addition to the values in the table. The URI values shown here reflect the URIs that the relying party will accept by default.

Create a Rule to Send an AD FS 1.x Compatible Claim

In situations in which you are using Active Directory Federation Services (AD FS) to issue claims that will be received by federation servers running AD FS 1.0 (Windows Server 2003 R2) or AD FS 1.1 (Windows Server 2008 or Windows Server 2008 R2), you must do the following:

- Create a rule that will send a Name ID claim type with a format of UPN, Email, or Common Name.
- All other claims that are sent must have one of the following claim types:
  - AD FS 1.*x* Email Address
  - AD FS 1.*x* UPN
  - Common Name
  - Group
  - Any other claim type that begins with https://schemas.xmlsoap.org/claims/, such as https://schemas.xmlsoap.org/claims/EmployeeID

Depending on the needs of your organization, use one of the following procedures to create an AD FS 1.*x* compatible NameID claim:

- Create this rule to issue an AD FS 1.x Name ID claim using the Pass Through or Filter an Incoming Claim rule template
- Create this rule to issue an AD FS 1.x Name ID claim using the Transform an Incoming Claim rule template. You can use this rule template in situations in which you want to change the existing claim type to a new claim type that will work with AD FS 1. *x* claims.

Note
For this rule to work as expected, make sure that the relying party trust or claims provider trust where you are creating this rule has been configured to use the **AD FS 1.0 and 1.1 profile**.

Creating a rule to issue an AD FS 1.

You can use this rule template when you want the group that you specify in this template to ultimately determine what AD FS 1.*x*-compatible claim to issue.

To create a rule to issue an AD FS 1.x Name ID claim using the Pass Through or Filter an Incoming Claim rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select **Name ID** in the list.
8. In **Incoming name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
9. Select one of the following options, depending on the needs of your organization:
   - **Pass through all claim values**
   - **Pass through only a specific claim value**
   - **Pass through only claim values that match a specific email suffix value**
   - **Pass through only claim values that start with a specific value**
10. Click **Finish**, and then click **OK** to save the rule.

Creating a rule to issue an AD FS 1.

You can use this rule template in situations in which you want to change the incoming claim type to a new outgoing claim type of Name ID that will work with AD FS 1.*x* federation servers.

To create a rule to issue an AD FS 1.x Name ID claim using the Transform an Incoming Claim rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select the type of incoming claim that you want to transform in the list.
8. In **Outgoing claim type**, select **Name ID** in the list.
9. In **Outgoing name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
10. Select one of the following options, depending on the needs of your organization:
    - **Pass through all claim values**
    - **Replace an incoming claim value with a different outgoing claim value**
    - **Replace incoming e-mail suffix claims with a new e-mail suffix**
11. Click **Finish**, and then click **OK** to save the rule.

Create a Rule to Send Claims Using a Custom Rule

Using the Send Claims Using a Custom Rule template in Active Directory Federation Services (AD FS), you can create custom claim rules for situation in which a standard rule template is not sufficient for the needs of your organization. Custom claim rules are written in the claim rule

language and must then be copied into the **Custom rule** text box before they can be used in a rule set. For information about constructing the syntax for an advanced rule, see The Role of the Claim Rule Language.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send claims using a custom claim

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule** from the list, and then click **Next**.
6. On the **Configure Rule** page under **Claim rule name** type the display name for this rule. Under **Custom rule**, type or paste the desired claim rule language syntax for this rule.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

Additional references

Checklist: Creating Claim Rules for a Claims Provider Trust

Checklist: Creating Claim Rules for a Relying Party Trust

The Role of the Claim Rule Language

The Role of Claims

The Role of Claim Rules

Deploying Federation Servers

To deploy federation servers in Active Directory Federation Services (AD FS), complete each of the tasks in Checklist: Setting Up a Federation Server.

Note
When you use this checklist, we recommend that you first read the references to federation server planning in the AD FS Design Guide in Windows Server 2012 before you begin the procedures for configuring the servers. Following the checklist in this way provides a better understanding of the design and deployment process for federation servers.

About federation servers

Federation servers are computers running Windows Server 2008 with the AD FS software installed that have been configured to act in the federation server role. Federation servers authenticate or route requests from user accounts in other organizations and from client computers that can be located anywhere on the Internet.

The act of installing the AD FS software on a computer and using the AD FS Federation Server Configuration Wizard to configure it for the federation server role makes that computer a federation server. It also makes the AD FS Management snap-in available on that computer in the **Start\Administrative Tools** menu so that you can specify the following:

- The AD FS host name where partner organizations and applications will send token requests and responses
- The AD FS identifier that partner organizations and applications will use to identify the unique name or location of your organization
- The token-signing certificate that all federation servers in a server farm will use to issue and sign tokens
- The location of customized ASP.NET Web pages for client logon, logoff, and account partner discovery that will enhance the client experience

Note
The majority of these core user interface (UI) settings are contained in the web.config file on each federation server. The AD FS host name and AD FS identifier values are not specified in the web.config file.

Federation servers host a claims issuance engine that issues tokens based on the credentials (for example, user name and password) that are presented to it. A security token is a cryptographically signed data unit that expresses one or more claims. A claim is a statement that a server makes (for example, name, identity, key, group, privilege, or capability) about a client. After the credentials are verified on the federation server (through the user logon process), claims for the user are collected through examination of the user attributes that are stored in the specified attribute store.

In Federated Web Single-Sign-On (SSO) designs (AD FS designs in which two or more organizations are involved), claims can be modified by claim rules for a specific relying party. The claims are built into a token that is sent to a federation server in the resource partner organization. After a federation server in the resource partner receives the claims as incoming claims, it executes the claims issuance engine to run a set of claim rules to filter, pass through, or transform those claims. The claims are then built into a new token that is sent to the Web server in the resource partner.

In the Web SSO design (an AD FS design in which only one organization is involved), a single federation server can be used so that employees can log on once and still access multiple applications.

Checklist: Setting Up a Federation Server

This checklist includes the deployment tasks that are necessary to prepare a server running Windows Server® 2012 for the federation server role in Active Directory Federation Services (AD FS).

Note

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Setting up a federation server

	Task	Reference
☐	Before you begin deploying your AD FS federation servers, review the; 1.) advantages and disadvantages of choosing either Windows Internal Database (WID) or SQL Server to store the AD FS configuration database 2.) AD FS deployment topology types and their associated server placement and network layout recommendations.	Determine Your AD FS Deployment Topology AD FS Deployment Topology Considerations
☐	Review AD FS capacity planning guidance to determine the proper number of federation servers you should use in your production environment.	Planning for Federation Server Capacity
☐	Review information in the AD FS Design Guide about where to place federation servers in your organization	Planning Federation Server Placement Where to Place a

	Task	Reference
		Federation Server
☐	Determine whether a stand-alone federation server or a federation server farm is better for your deployment.	When to Create a Federation Server When to Create a Federation Server Farm
☐	Determine whether this new federation server will be created in the account partner organization or in the resource partner organization.	Review the Role of the Federation Server in the Account Partner Review the Role of the Federation Server in the Resource Partner
☐	Review information about how federation servers use service communication certificates and token-signing certificates to securely authenticate client and federation server proxy requests. **Caution** Though it has long been common practice to use certificates with unqualified host names such as https://myserver, these certificates have no security value and can enable an attacker to impersonate the AD FS Federation Service to enterprise clients. Therefore, it is recommended that you use a fully qualified domain name (FQDN) such as https://myserver.contoso.com and only use SSL certificates issued to the FQDN of your Federation Service.	Certificate Requirements for Federation Servers
☐	Review information about how to update the corporate network Domain Name System (DNS) so that successful name resolution to federation servers can occur.	Name Resolution Requirements for Federation Servers

	Task	Reference
☐	Join the computer that will become the federation server to a domain in the account partner forest or resource partner forest where it will be used to authenticate the users of that forest or from trusting forests. **Note** If you want to set up a federation server in the account partner organization, the computer must first be joined to any domain in the forest where your federation server will be used to authenticate users from that forest or from trusting forests.	Join a Computer to a Domain
☐	Create a new resource record in the corporate network DNS that points the DNS host name of the federation server to the IP address of the federation server.	Add a Host (A) Resource Record to Corporate DNS for a Federation Server
☐	(Optional) If you will be adding a federation server to a federation server farm, you might have to first export the private key of the existing token-signing certificate (on the first federation server in the farm) so that you have a file format of the certificate ready when other federation servers must import the same certificate. Exporting the private key is not required when your issued server authentication certificate can be reused by multiple computers (without the need to export) or when you will be obtaining unique server authentication certificates for each federation server in the farm. **Note** The AD FS Management snap-in refers to server authentication certificates for federation servers	Export the Private Key Portion of a Server Authentication Certificate

	Task	Reference
	as service communication certificates.	
☐	After you obtain a server authentication certificate (or private key) from a certification authority (CA), you must then import the certificate file to the default Web site for each federation server. **Note** Installing this certificate on the default Web site is a requirement before you can use the AD FS Federation Server Configuration Wizard.	Import a Server Authentication Certificate to the Default Web Site
☐	(Optional) As an alternative to obtaining a server authentication certificate from a CA, you can use Internet Information Services (IIS) to create a sample certificate for your federation server. **Caution** It is not a security best practice to deploy a federation server in a production environment by using a self-signed server authentication certificate.	IIS: Create a Self-Signed Server Certificate and then complete the procedure Import a Server Authentication Certificate to the Default Web Site
☐	If you will be configuring a federation server farm environment in an account partner organization, you must create and configure a dedicated service account in Active Directory Domain Services (AD DS) where the farm will reside and configure each federation server in the farm to use this account. By performing this procedure, you will allow clients on the corporate network to authenticate to any of the federation servers in the farm using Windows Integrated Authentication.	Manually Configure a Service Account for a Federation Server Farm
☐	Install the Federation Service role service on the computer that will become the	Install the Federation Service Role Service

	Task	Reference
	federation server.	
☐	Configure the AD FS software on the computer to act in the federation server role by using the AD FS Federation Server Configuration Wizard. Follow this procedure when you want to set up a stand-alone federation server, create the first federation server in a new farm or join a computer to an existing federation server farm. **Note** For the Federated Web Single Sign-On (SSO) design, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization.	Create a Stand-Alone Federation Server Create the First Federation Server in a Federation Server Farm Add a Federation Server to a Federation Server Farm
☐	(Optional) Use the AD FS Management snap-in to add and configure the necessary AD FS certificates required to deploy your design. For more information about when to add or change certificates using the snap-in, see Certificate Requirements for Federation Servers.	Add a Token-Signing Certificate Add a Token-Decrypting Certificate Set a Service Communications Certificate
☐	If this is the first federation server in your organization, configure the Federation Service so that it conforms to your AD FS design.	Checklist: Configuring the Account Partner Organization Checklist: Configuring the Resource Partner Organization
☐	From a client computer, verify that the federation server is operational.	Verify That a Federation Server Is Operational

Join a Computer to a Domain

For Active Directory Federation Services (AD FS) to function, each computer that functions as a federation server must be joined to a domain. federation server proxies may be joined to a domain, but this is not a requirement.

You do not have to join a Web server to a domain if the Web server is hosting claims-aware applications only.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To join a computer to a domain

1. On the **Start** screen, type **Control Panel**, and then press ENTER.
2. Navigate to **System and Security**, and then click **System**.
3. Under **Computer name, domain, and workgroup settings**, click **Change settings**.
4. On the **Computer Name** tab, click **Change**.
5. Under **Member of**, click **Domain**, type the name of the domain that this computer will join, and then click **OK**.
6. Click **OK**, and then restart the computer.

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Add a Host (A) Resource Record to Corporate DNS for a Federation Server

For clients on the corporate network to successfully access a federation server using Windows Integrated authentication, a host (A) resource record must first be created in the corporate Domain Name System (DNS) that resolves the host name of the account federation server (for example, fs.fabrikam.com) to the IP address of the federation server or federation server cluster. You can use the following procedure to add a host (A) resource record to corporate DNS for a federation server.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a host (A) resource record to corporate DNS for a federation server

1. On a DNS server for the corporate network, open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A or AAAA)**.
3. In **Name**, type only the computer name of the federation server or federation server cluster; for example, for the fully qualified domain name (FQDN) fs.fabrikam.com, type **fs**.
4. In **IP address**, type the IP address for the federation server or federation server cluster, for example, 192.168.1.4.
5. Click **Add Host**.

Additional references

Checklist: Setting Up a Federation Server

Name Resolution Requirements for Federation Servers

Export the Private Key Portion of a Server Authentication Certificate

Every federation server in an Active Directory Federation Services (AD FS) farm must have access to the private key of the server authentication certificate. If you are implementing a server farm of federation servers or Web servers, you must have a single authentication certificate. This certificate must be issued by an enterprise certification authority (CA), and it must have an exportable private key. The private key of the server authentication certificate must be exportable so that it can be made available to all the servers in the farm.

This same concept is true of federation server proxy farms in the sense that all federation server proxies in a farm must share the private key portion of the same server authentication certificate.

Note
The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Depending on which role this computer will play, use this procedure on the federation server computer or federation server proxy computer where you installed the server authentication certificate with the private key. When you finish the procedure, you can then import this certificate on the Default Web Site of each server in the farm. For more information, see Import a Server Authentication Certificate to the Default Web Site.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To export the private key portion of a server authentication certificate

1. On the **Start** screen, type **Internet Information Services (IIS) Manager**, and then press ENTER.
2. In the console tree, click **ComputerName**.
3. In the center pane, double-click **Server Certificates**.
4. In the center pane, right-click the certificate that you want to export, and then click **Export**.
5. In the **Export Certificate** dialog box, click the **…** button.
6. In **File name**, type **C:***NameofCertificate*, and then click **Open**.
7. Type a password for the certificate, confirm it, and then click **OK**.
8. Validate the success of your export by confirming that the file you specified is created at the specified location.

   Important
   So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. If this key is compromised, the security of your entire AD FS deployment (including resources within your organization and in resource partner organizations) is compromised.
9. Import the exported server authentication certificate into the certificate store on the new server before you install the Federation Service. For information about how to import the certificate, see Import a Server Certificate (http://go.microsoft.com/fwlink/?LinkId=108283).

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Certificate Requirements for Federation Servers

Certificate Requirements for Federation Server Proxies

Import a Server Authentication Certificate to the Default Web Site

After you obtain a server authentication certificate from a certification authority (CA), you must manually install that certificate on the Default Web Site for each federation server or federation server proxy in a server farm.

For Web servers, you must manually install the server authentication certificate on the appropriate Web site or virtual directory where your federated application resides.

If you are setting up a farm, be sure to perform this procedure identically—using the exact same settings—on each of the servers in your farm.

Note

The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To import a server authentication certificate to the Default Web Site

1. On the **Start** screen, type **Internet Information Services (IIS) Manager**, and then press ENTER.
2. In the console tree, click **ComputerName**.
3. In the center pane, double-click **Server Certificates**.
4. In the **Actions** pane, click **Import**.
5. In the **Import Certificate** dialog box, click the **…** button.
6. Browse to the location of the pfx certificate file, highlight it, and then click **Open**.
7. Type a password for the certificate, and then click **OK**.

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Certificate Requirements for Federation Servers

Certificate Requirements for Federation Server Proxies

Manually Configure a Service Account for a Federation Server Farm

If you intend to configure a federation server farm environment in Active Directory Federation Services (AD FS), you must create and configure a dedicated service account in Active Directory Domain Services (AD DS) where the farm will reside. You then configure each federation server in the farm to use this account. You must complete the following tasks in your organization when you want to allow client computers on the corporate network to authenticate to any of the federation servers in an AD FS farm using Windows Integrated Authentication.

Note

You have to perform the tasks in this procedure only one time for the entire federation server farm. Later, when you create a federation server by using the AD FS Federation Server Configuration Wizard, you must specify this same account on the **Service Account** wizard page on each federation server in the farm.

Create a dedicated service account

1. Create a dedicated user/service account in the Active Directory forest that is located in the identity provider organization. This account is necessary for the Kerberos authentication protocol to work in a farm scenario and to allow pass-through authentication on each of the federation servers. Use this account only for the purposes of the federation server farm.
2. Edit the user account properties, and select the **Password never expires** check box. This action ensures that this service account's function is not interrupted as a result of domain password change requirements.

   Note

   Using the Network Service account for this dedicated account will result in random failures when access is attempted through Windows Integrated Authentication, as a result of Kerberos tickets not validating from one server to another.

To set the SPN of the service account

1. Because the application pool identity for the AD FS AppPool is running as a domain user/service account, you must configure the Service Principal Name (SPN) for that account in the domain with the Setspn.exe command-line tool. Setspn.exe is installed by default on computers running Windows Server 2008. Run the following command on a computer that is joined to the same domain where the user/service account resides:

   ```
   setspn -a host/<server name> <service account>
   ```

   For example, in a scenario in which all federation servers are clustered under the Domain Name System (DNS) host name fs.fabrikam.com and the service account name that is assigned to the AD FS AppPool is named adfs2farm, type the command as follows, and then press ENTER:

   ```
   setspn -a host/fs.fabrikam.com adfs2farm
   ```

   It is necessary to complete this task only once for this account.
2. After the AD FS AppPool identity is changed to the service account, set the access control lists (ACLs) on the SQL Server database to allow Read access to this new account so that the AD FS AppPool can read the policy data.

Install the Federation Service Role Service

Now that you have properly configured a computer with the prerequisite applications and certificates, you are ready to install the Federation Service role service of Active Directory Federation Services (AD FS). When you install the Federation Service on a computer, that computer becomes a federation server.

Note

For the Federated Web Single-Sign-On (SSO) design, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. For more information, see Where to Place a Federation Server.

You can use the following procedure to install the Federation Service role service of AD FS on a computer that will become the first federation server or on a computer that will become a federation server for an existing federation server farm.

Prerequisites

Verify that an SSL certificate with the private key has already been installed or imported into the local certificate store (Personal store) before you start this procedure. If you will be using a token-signing certificate that is issued by a certification authority (CA), verify that a token-signing certificate with the private key has already been installed or imported into the local certificate store (Personal store) before you start this procedure. As an alternative, you can create a self-signed, token-signing certificate using the Add Roles Wizard, as described in this procedure. For more information about token-signing certificates, see Certificate Requirements for Federation Servers.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To install the Federation Service role service

1. On the **Start** screen, type **Server Manager**, and then press ENTER.
2. Click **Manage**, and then click **Add Roles and Features** to start the Add Roles and Features Wizard.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or Feature-based installation**, and click **Next**.
5. On the **Select destination server** page, click **Select a server from the server pool**, verify that the target computer is highlighted, and then click **Next**.
6. On the **Select server roles** page, click **Active Directory Federation Services**, and then click next.

Note
If you are prompted to install additional .NET Framework or Windows Process Activation Service features, click **Add Features** to install them.

7. On the **Select features** page, verify that the features are set, and then click **Next**.
8. On the **Active Directory Federation Service (AD FS)** page, click **Next**.
9. On the **Select role services** page, select the **Federation Service** check box, and then click **Next**.
10. On the **Web Server Role (IIS)** page, click **Next**.
11. On the **Select role services** page, click **Next**.
12. After you verify the information on the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.
13. On the **Installation progress** page, verify that everything installed correctly, and then click **Close**.

Create the First Federation Server in a Federation Server Farm

After you install the Federation Service role service and configure the required certificates on a computer, you are ready to configure the computer to become a federation server. You can use the following procedure to set up the computer to become the first federation server in a new federation server farm using the AD FS Federation Server Configuration Wizard.

The act of creating the first federation server in a farm also creates a new Federation Service and makes this computer the primary federation server. This means that this computer will be configured with a read/write copy of the AD FS configuration database. All other federation servers in this farm must replicate any changes that are made on the primary federation server to their read-only copies of the AD FS configuration database that they store locally. For more information about this replication process, see The Role of the AD FS Configuration Database.

Note
For the Federated Web Single-Sign-On (SSO) design, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. For more information, see Where to Place a Federation Server.

Membership in Domain Admins, or a delegated domain account that has been granted write access to the Program Data container in Active Directory, is the minimum required to complete this procedure.

To create the first federation server in a federation server farm

1. There are two ways to start the AD FS Federation Server Configuration Wizard. To start the wizard, do one of the following:
   - After the Federation Service role service installation is complete, open the AD FS Management snap-in and click the **AD FS Federation Server Configuration Wizard** link on the **Overview** page or in the **Actions** pane.
   - Any time after the setup wizard is complete, open Windows Explorer, navigate to the **C:\Windows\ADFS** folder, and then double-click **FsConfigWizard.exe**.
2. On the **Welcome** page, verify that **Create a new Federation Service** is selected, and then click **Next**.
3. On the **Select Stand-Alone or Farm Deployment** page, click **New federation server farm**, and then click **Next**.
4. On the **Specify the Federation Service Name** page, verify that the **SSL certificate** that is showing is correct. If this is not the correct certificate, select the appropriate certificate from the **SSL certificate** list.

   This certificate is generated from the Secure Sockets Layer (SSL) settings for the Default Web Site. If the Default Web Site has only one SSL certificate configured, that certificate is presented and automatically selected for use. If multiple SSL certificates are configured for the Default Web Site, all those certificates are listed here and you must select from among them. If there are no SSL settings configured for the Default Web Site, the list is generated from the certificates that are available in the personal certificates store on the local computer.

   Note
   The wizard will not allow you to override the certificate if an SSL certificate is configured for IIS. This ensures that any intended prior IIS configuration for SSL certificates is preserved. To work around this restriction, you can remove the certificate or reconfigure it manually with the IIS Management Console.
5. If the AD FS database that you selected already exists, the **Existing AD FS Configuration Database Detected** page appears. If that page appears, click **Delete database**, and then click **Next**.

   Caution
   Select this option only when you are sure that the data in this AD FS database is not important or that it is not used in a production federation server farm.
6. On the **Specify a Service Account** page, click **Browse**. In the **Browse** dialog box, locate the domain account that will be used as the service account in this new federation server farm, and then click **OK**. Type the password for this account, confirm it, and then click **Next**.

   Note
   See Manually Configure a Service Account for a Federation Server Farm for more information about specifying a service account for a federation server farm.

Each federation server in the federation server farm must specify the same service account for the farm to be operational. For example, if the service account that was created was contoso\ADFS2SVC, each computer that you configure for the federation server role and that will participate in the same farm must specify contoso\ADFS2SVC at this step in the Federation Server Configuration Wizard for the farm to be operational.

7. On the **Ready to Apply Settings** page, review the details. If the settings appear to be correct, click **Next** to begin configuring AD FS with these settings.
8. On the **Configuration Results** page, review the results. When all the configuration steps are finished, click **Close** to exit the wizard.

Security

For secure deployment purposes, artifact resolution and reply detection are disabled when you use the AD FS Federation Server Configuration Wizard to configure a federation server farm. This wizard automatically configures the Windows Internal Database for storing service configuration data. You might, however, mistakenly undo this change by enabling the Artifact Resolution endpoint using either the **Endpoints** node in the AD FS Management snap-in or the Enable-ADFSEndpoint cmdlet in Windows PowerShell. Be careful to not reconfigure the default setting so that this endpoint remains disabled when you use a federation server farm and the Windows Internal Database together.

Additional references

Checklist: Setting Up a Federation Server

Create a Stand-Alone Federation Server

After you install the Federation Service role service and configure the required certificates on a computer, you are ready to configure the computer to become a federation server. You can use the following procedure to set up the computer to become a stand-alone federation server. The act of creating a stand-alone federation server also creates a new Federation Service. You do create a federation server with the AD FS Federation Server Configuration Wizard.

Note

For the Federated Web Single-Sign-On (SSO) design, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. For more information, see Where to Place a Federation Server.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group

memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a stand-alone federation server

1. There are two ways to start the AD FS Federation Server Configuration Wizard. To start the wizard, do one of the following:
   - After the Federation Service role service installation is complete, open the AD FS Management snap-in and click the **AD FS Federation Server Configuration Wizard** link on the **Overview** page or in the **Actions** pane.
   - Anytime after the setup wizard is complete, open Windows Explorer, navigate to the **C:\Windows\ADFS** folder, and then double-click **FsConfigWizard.exe**.
2. On the **Welcome** page, verify that **Create a new Federation Service** is selected, and then click **Next**.
3. On the **Select Stand-Alone or Farm Deployment** page, click **Stand-alone federation server**, and then click **Next**.

   Important
   When you select the Stand-alone federation server option in the AD FS Federation Server Configuration Wizard, the service account associated with this Federation Service will automatically be assigned to the NETWORK SERVICE account. Using NETWORK SERVICE as the service account is only recommended in situations where you are evaluating AD FS in a test lab environment. If you intend to use the Stand-alone federation server option to deploy a federation server in a production environment, it is important that you change this service account to a more appropriate service account that can be dedicated to serving requests for this new Federation Service. Changing the service account to an account other than NETWORK SERVICE will mitigate possible attack vectors that would otherwise make your federation server vulnerable to malicious attacks.
4. On the **Specify the Federation Service Name** page, verify that the **SSL certificate** that is showing is correct. If not, select the appropriate certificate from the **SSL certificate** list.

   This certificate is generated from the Secure Sockets Layer (SSL) settings for the Default Web Site. If the Default Web Site has only one SSL certificate configured, that certificate is presented and automatically selected for use. If multiple SSL certificates are configured for the Default Web Site, all those certificates are listed here and you must select from among them. If there are no SSL settings configured for the Default Web Site, the list is generated from the certificates that are available in the personal certificates store on the local computer.

   Note
   The wizard will not allow you to override the certificate if an SSL certificate is configured for IIS. This ensures that any intended prior IIS configuration for SSL certificates is preserved. To work around this restriction, you can remove the

certificate or reconfigure manually it with the IIS Management Console.

5. If the AD FS database that you selected already exists, the **Existing AD FS Configuration Database Detected** page appears. If that occurs, click **Delete database**, and then click **Next**.

   Caution
   Select this option only when you are sure that the data in this AD FS database is not important or that it is not used in a production federation server farm.

6. On the **Ready to Apply Settings** page, review the details. If the settings appear to be correct, click **Next** to begin configuring AD FS with these settings.
7. On the **Configuration Results** page, review the results. When all the configuration steps are finished, click **Close** to exit the wizard.

Additional references

Checklist: Setting Up a Federation Server

Add a Federation Server to a Federation Server Farm

After you install the Federation Service role service and configure the required certificates on a computer, you are ready to configure the computer to become a federation server. You can use the following procedure to join a computer to a new federation server farm.

You join a computer to a farm with the AD FS Federation Server Configuration Wizard. When you use this wizard to join a computer to an existing farm, the computer is configured with a read-only copy of the AD FS configuration database and it must receive updates from a primary federation server.

Note
For the Federated Web Single-Sign-On (SSO) design, you must have at least one federation server in the account partner organization and at least one federation server in the resource partner organization. For more information, see Where to Place a Federation Server.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a federation server to a federation server farm

1. There are two ways to start the AD FS Federation Server Configuration Wizard. To start the wizard, do one of the following:

- After the Federation Service role service installation is complete, open the AD FS Management snap-in and click the **AD FS Federation Server Configuration Wizard** link on the **Overview** page or in the **Actions** pane.
- Anytime after the setup wizard is complete, open Windows Explorer, navigate to the **C:\Windows\ADFS** folder, and double-click **FsConfigWizard.exe**.

2. On the **Welcome** page, verify that **Add a federation server to an existing Federation Service** is selected, and then click **Next**.
3. If the AD FS database that you selected already exists, the **Existing AD FS Configuration Database Detected** page appears. If that occurs, click **Delete database**, and then click **Next**.

   Caution
   Select this option only when you are sure that the data in this AD FS database is not important or that it is not used in a production federation server farm.
4. On the **Specify the Primary Federation Server and Service Account** page, under **Primary federation server name**, type the computer name of the primary federation server in the farm, and then click **Browse**. In the **Browse** dialog box, locate the domain account that is used as the service account by all other federation servers in the existing federation server farm, and then click **OK**. Type the password and confirm it, and then click **Next**:

   Note
   For more information about specifying a service account for a federation server farm, see Manually Configure a Service Account for a Federation Server Farm. Each federation server in the federation server farm must specify the same service account for the farm to be operational. For example, if the service account that was created was contoso\ADFS2SVC, each computer you configure for the federation server role and that will participate in the same farm must specify contoso\ADFS2SVC at this step in the Federation Server Configuration Wizard for the farm to be operational.
5. On the **Ready to Apply Settings** page, review the details. If the settings appear to be correct, click **Next** to begin configuring AD FS with these settings.
6. On the **Configuration Results** page, review the results. When all the configuration steps are finished, click **Close** to exit the wizard.

Additional references

Checklist: Setting Up a Federation Server

Add a Token-Signing Certificate

Federation servers in Active Directory Federation Services (AD FS) require token-signing certificates to prevent attackers from altering or counterfeiting security tokens in an attempt to

gain unauthorized access to federated resources. Every token-signing certificate contains cryptographic private keys and public keys that are used to digitally sign (by means of the private key) a security token. Later, after these keys are received by a partner federation server, they validate the authenticity (by means of the public key) of the encrypted security token.

Caution
Certificates used for token-signing are critical to the stability of the Federation Service. Because loss or unplanned removal of any certificates configured for this purpose can disrupt service, you should backup any certificates configured for this purpose.

The token-signing certificate should chain to a trusted root in the Federation Service. You can use the following procedure to add the token-signing certificate to the AD FS Management snap-in from a file that you have exported.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a token-signing certificate

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, double-click **Service**, and then click **Certificates**.
3. In the **Actions** pane, click the **Add Token-Signing Certificate** link.
4. In the **Browse for Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.

Additional references

Checklist: Setting Up a Federation Server

Certificate Requirements for Federation Servers

Add a Token-Decrypting Certificate

Federation servers use a token-decryption certificate when a relying party federation server must decrypt tokens that are issued with an older certificate after a new certificate is set as the primary decryption certificate. Active Directory Federation Services (AD FS) uses the Secure Sockets Layer (SSL) certificate for Internet Information Services (IIS) as the default decryption certificate.

Caution
Certificates used for token-decrypting are critical to the stability of the Federation Service. Because loss or unplanned removal of any certificates configured for this purpose can disrupt service, you should backup any certificates configured for this purpose.

You can use the following procedure to add the token-decrypting certificate to the AD FS Management snap-in from a file that you have exported.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a token-decrypting certificate

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, double-click **Service**, and then click **Certificates**.
3. In the **Actions** pane, click the **Add Token-Decrypting Certificate** link.
4. In the **Browse for Certificate file** dialog box, navigate to the certificate file that you want to add, select the certificate file, and then click **Open**.

Additional references

Checklist: Setting Up a Federation Server

Certificate Requirements for Federation Servers

Set a Service Communications Certificate

Federation servers in Active Directory Federation Services (AD FS) use the service communications certificate to secure Web services traffic for Secure Sockets Layer (SSL) communication with Web clients or with federation server proxies. This is the same certificate that a federation server uses as the SSL certificate in Internet Information Services (IIS).

You can use the following procedure to change the service communications certificate with the AD FS Management snap-in.

Note
The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To set a service communications certificate

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, double-click **Service**, and then click **Certificates**.
3. In the **Actions** pane, click the **Set Service Communications Certificate** link.

4. In the **Select a service communications certificate** dialog box, navigate to the certificate file that you want to set as the service communications certificate, select the certificate file, and then click **Open**.

Additional references

Checklist: Setting Up a Federation Server

Certificate Requirements for Federation Servers

Verify That a Federation Server Is Operational

You can use the following procedures to verify that a federation server is operational; that is, that any client on the same network can reach a new federation server.

Membership in **Users**, **Backup Operators**, **Power Users**, **Administrators** or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Procedure 1: To verify that a federation server is operational

1. To verify that Internet Information Services (IIS) is configured correctly on the federation server, log on to a client computer that is located in the same forest as the federation server.
2. Open a browser window, in the address bar type the federation server's DNS host name, and then append /adfs/fs/federationserverservice.asmx to it for the new federation server, for example:

   https://fs1.fabrikam.com/adfs/fs/federationserverservice.asmx
3. Press ENTER, and then complete the next procedure on the federation server computer. If you see the message **There is a problem with this website's security certificate**, click **Continue to this website**.

   The expected output is a display of XML with the service description document. If this page appears, IIS on the federation server is operational and serving pages successfully.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Procedure 2: To verify that a federation server is operational

1. Log on to the new federation server as an administrator.

2. On the **Start** screen, type **Event Viewer**, and then press ENTER.
3. In the details pane, double-click **Applications and Services Logs**, double-click **AD FS Eventing**, and then click **Admin**.
4. In the **Event ID** column, look for event ID 100. If the federation server is configured properly, you see a new event—in the Application log of Event Viewer—with the event ID 100. This event verifies that the federation server was able to successfully communicate with the Federation Service.

Additional references

Checklist: Setting Up a Federation Server

Configure Performance Monitoring

AD FS includes its own dedicated performance counters to help you monitor the performance of both federation servers and federation server proxy computers. To use Performance Monitor to monitor the performance of your AD FS servers, it's useful to create a new data collector set and add the AD FS counters to that view. The following procedure describes how to configure performance monitoring for AD FS.

To configure performance monitoring for AD FS using Performance Monitor

1. On the **Start** screen, type **Performance Monitor**, and then press ENTER.
2. In the console tree, expand **Data Collector Sets**, right-click **User Defined**, point to **New**, and then click **Data Collector Set**.

   The Create New Data Collector Set Wizard appears.
3. In **Create New Data Collector Set**, for **Name** type a name for the new data collector set (such as "AD FS performance"), click **Create manually (Advanced)**, and then click **Next**.
4. For the type of data to include, verify that **Create data logs** is selected, and then click the check boxes for the following data types: **Performance counter**, **Event trace data**, **System configuration information**.
5. For performance counters, expand **AD FS** in the **Available counters** list, and then click **Add**.

   The AD FS performance counters should appear in the **Added counters** list.
6. When you are prompted to add event trace providers, click **Add**, select **AD FS Eventing** and **AD FS Tracing** from the list of providers.
7. When you are prompted to add registry keys to monitor, click **Next**.
8. When you are prompted to specify the location to save the performance data, you can accept the default location (**%systemdrive%\PerfLogs\Admin**<*data_collector_set*>,

and then click **Next**.

9. When you are prompted to create the data collector set, select **Save and close**, and then click **Finish**.

   The new data collector set appears in the console tree under the **User Defined** node.

10. Use the following steps to work with the AD FS performance counters:
    - To begin performance monitoring using AD FS-related counters, right-click the data collector set that you added (such as "AD FS performance"), and then click **Start**.
    - To create a report to view the performance monitoring results, right-click the data collector set that you added (such as "AD FS performance"), and then click **Latest Report**.
    - To end a capture of performance data so that you can view the latest report, right-click the data collector set that you added (such as "AD FS performance"), and then click **Stop**.

    The latest report is added and numbered automatically (starting at 000001) under the **Report\User Defined**\<*data_collector_set*> node in the console tree.

AD FS performance counters

The following table lists the AD FS performance counters and describes how they are useful for monitoring activity that relates to either a federation server or federation server proxy.

Counter	Description	Can be used on:
Token Requests	Monitors the number of token requests sent to the federation server including SSOAuth token requests.	Federation Servers
Token Requests/sec	Monitors the number of token requests sent to the federation server including SSOAuth token requests per second.	Federation Servers
Federation Metadata Requests	Monitors the number of incoming federation metadata requests sent to the federation server.	Federation Servers
Federation Metadata Requests/sec	Monitors the number of incoming federation metadata requests per second that are sent to the federation server.	Federation Servers
Artifact Resolution Requests	Monitors the number of incoming federation metadata	Federation Servers

Counter	Description	Can be used on:
	requests per second that are sent to the federation server.	
Artifact Resolution Requests/sec	Monitors the number of requests to the artifact resolution endpoint per second that are sent to the federation server.	Federation Servers
Proxy Requests	Monitors the number of incoming requests sent to the federation server proxy.	Federation Server Proxies
Proxy Requests/sec	Monitors the number of incoming requests per second that are sent to the federation server proxy.	Federation Server Proxies
Proxy MEX Requests	Monitors the number of incoming WS-Metadata Exchange (MEX) requests that are sent to the federation server proxy.	Federation Server Proxies
Proxy MEX Requests/sec	Monitors the number of incoming MEX requests per second that are sent to the federation server proxy.	Federation Server Proxies

Deploying Federation Server Proxies

To deploy federation server proxies in Active Directory Federation Services (AD FS), complete each of the tasks in Checklist: Setting Up a Federation Server Proxy.

Note

When you use this checklist, we recommend that you first read the references to federation server proxy planning guidance in the AD FS Design Guide in Windows Server 2012 before you begin the procedures for configuring the servers. Following the checklist provides a better understanding of the design and deployment process for federation server proxies.

About federation server proxies

Federation server proxies are computers that run Windows Server® 2012 and AD FS software that have been configured manually to act in the proxy role. You can use federation server proxies in your organization to provide intermediary services between an Internet client and a federation server that is behind a firewall on your corporate network.

Note

Although the federation server and the federation server proxy roles cannot be installed on the same computer, a federation server can perform federation server proxy functions. For more information, see When to Create a Federation Server.

The act of installing the AD FS software on a Windows Server® 2012 computer and configuring it to serve in the proxy role makes that computer a federation server proxy.

Checklist: Setting Up a Federation Server Proxy

This checklist includes the deployment tasks for preparing a server running Windows Server® 2012 for the federation server proxy role in Active Directory Federation Services (AD FS).

Note

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Setting Up a federation server proxy

	Task	Reference
☐	Before you begin deploying your AD FS federation server proxies, review the AD FS deployment topology types and their associated server placement and network layout recommendations.	Determine Your AD FS Deployment Topology Planning Federation Server Proxy Placement Where to Place a Federation Server Proxy
☐	Review AD FS capacity planning guidance to determine the proper number of federation server proxies you should	Planning for Federation Server Proxy Capacity

	Task	Reference
	use in your production environment.	
☐	Determine whether a single federation server proxy or a federation server proxy farm is better for your deployment. **Note** Federation servers also perform federation server proxy responsibilities.	When to Create a Federation Server Proxy When to Create a Federation Server Proxy Farm
☐	Determine whether this new federation server proxy will be created in the perimeter network of the account partner organization or the resource partner organization.	Review the Role of the Federation Server Proxy in the Account Partner Review the Role of the Federation Server Proxy in the Resource Partner
☐	Before you install AD FS on a computer that will become a federation server proxy, read about the importance of obtaining a server authentication certificate—for federation server proxy farms—adding or sharing certificates across all the servers in a farm.	Certificate Requirements for Federation Server Proxies
☐	Review information in the AD FS Design Guide about how to update Domain Name System (DNS) in the perimeter network so that successful name resolution for	Name Resolution Requirements for Federation Server Proxies

	Task	Reference
	federation servers and federation server proxies can occur.	
☐	Determine whether the federation server proxy must be joined to a domain. Although federation server proxies do not have to be joined to a domain, they are easier to manage with remote administration and Group Policy features when they are joined to a domain.	Join a Computer to a Domain
☐	Depending on how the DNS infrastructure in your perimeter network is configured, complete one of the procedures in the topics on the right before you deploy a federation server proxy in your organization. **Note** Do not perform both procedures. Read Name Resolution Requirements for Federation Server Proxies to determine which procedure best suits the requirements of your organization.	Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Only the Perimeter Network Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Both the Perimeter Network and Internet Clients
☐	After you obtain a server authentication certificate,	Import a Server Authentication Certificate to the Default Web Site

	Task	Reference
	you must install it in Internet Information Services (IIS) on the default Web site of the federation server proxy.	
☐	(Optional) As an alternative to obtaining a server authentication certificate from a certification authority (CA), you can use IIS to acquire a sample certificate for your federation server proxy. Because IIS generates a self-signed certificate that does not originate from a trusted source, use it to create a self-signed certificate only in the following scenarios: • When you have to create a Secure Sockets Layer (SSL) channel between your server and a limited, known group of users • When you have to troubleshoot third-party certificate problems **Caution** It is not a security best practice to deploy a federation server proxy in a production environment using a self-signed,	IIS: Create a Self-Signed Server Certificate(http://go.microsoft.com/fwlink/?LinkID=108271)

	Task	Reference
	server authentication certificate.	
☐	Install the Federation Service Proxy role service on the computer that will become the federation server proxy.	Install the Federation Service Proxy Role Service
☐	Configure the AD FS software on the computer to act in the federation server proxy role by using the AD FSFederation Server Proxy Configuration Wizard.	Configure a Computer for the Federation Server Proxy Role
☐	Using Event Viewer, verify that the federation server proxy service has started.	Verify That a Federation Server Proxy Is Operational

Join a Computer to a Domain

For Active Directory Federation Services (AD FS) to function, each computer that functions as a federation server must be joined to a domain. federation server proxies may be joined to a domain, but this is not a requirement.

You do not have to join a Web server to a domain if the Web server is hosting claims-aware applications only.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To join a computer to a domain

1. On the **Start** screen, type **Control Panel**, and then press ENTER.
2. Navigate to **System and Security**, and then click **System**.
3. Under **Computer name, domain, and workgroup settings**, click **Change settings**.

4. On the **Computer Name** tab, click **Change**.
5. Under **Member of**, click **Domain**, type the name of the domain that this computer will join, and then click **OK**.
6. Click **OK**, and then restart the computer.

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Only the Perimeter Network

So that name resolution can work successfully for a federation server in an Active Directory Federation Services (AD FS) scenario in which one or more Domain Name System (DNS) zones serve only the perimeter network, the following tasks must be completed:

- The hosts file on the federation server proxy must be updated to add the IP address of a federation server.
- DNS in the perimeter network must be configured to resolve all client requests for the AD FS host name to the federation server proxy. To do this, you add a host (A) resource record to perimeter DNS for the federation server proxy.

Note

These procedures assume that a host (A) resource record for the federation server has already been created in the corporate network DNS. If this record does not yet exist, create this record, and then perform these procedures. For more information about how to create the host (A) resource record for the federation server, see Add a Host (A) Resource Record to Corporate DNS for a Federation Server.

Add the IP address of a federation server to the hosts file

So that a federation server proxy can work as expected in the perimeter network of an account partner, you must add an entry to the hosts file on that federation server proxy that points to a federation server's DNS host name (for example, fs.fabrikam.com) and IP address (for example, 192.168.1.4) in the corporate network of the account partner. Adding this entry to the hosts file prevents the federation server proxy from contacting itself to resolve a client-initiated call to a federation server in the account partner.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add the IP address of a federation server to the hosts file

1. Navigate to the %systemroot%\Winnt\System32\Drivers directory folder and locate the **hosts** file.
2. Start Notepad, and then open the **hosts** file.
3. Add the IP address and the host name of a federation server in the account partner to the **hosts** file, as shown in the following example:

   192.168.1.4 fs.fabrikam.com
4. Save and close the file.

Add a host (A) resource record to perimeter DNS for a federation server proxy

So that clients on the Internet can successfully access a federation server through a newly deployed federation server proxy, you must first create a host (A) resource record in the perimeter DNS. This resource record resolves the host name of the account federation server (for example, fs.fabrikam.com) to the IP address of the account federation server proxy (for example, 131.107.27.68) in the perimeter network.

Note
It is assumed that you are using a DNS server, running Windows 2000 Server, Windows Server 2003, or Windows Server 2008 with the DNS Server service, to control the perimeter DNS zone.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a host (A) resource record to perimeter DNS for a federation server proxy

1. On a DNS server for the perimeter network, open the DNS snap-in. Click **Start**, point to **Administrative Tools**, and then click **DNS**.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A or AAAA)**.
3. In **Name**, type only the computer name of the federation server. For example, for the fully qualified domain name (FQDN) fs.fabrikam.com, type **fs**.
4. In **IP address**, type the IP address for the new federation server proxy, for example, **131.107.27.68**.
5. Click **Add Host**.

Additional references

Checklist: Setting Up a Federation Server Proxy

Name Resolution Requirements for Federation Server Proxies

Configure Name Resolution for a Federation Server Proxy in a DNS Zone That Serves Both the Perimeter Network and Internet Clients

---

So that name resolution can work successfully for a federation server proxy in an Active Directory Federation Services (AD FS) scenario in which one or more Domain Name System (DNS) zones serve both the perimeter network and Internet clients, the following tasks must be completed:

- DNS in the Internet zone that you control must be configured to resolve all Internet client requests for the AD FS host name to the federation server proxy. To accomplish this, you add a host (A) resource record to the Internet DNS zone for the federation server proxy.
- DNS in the perimeter network must be configured to resolve all incoming client requests for the AD FS host name to the federation server. To accomplish this, you add a host (A) resource record to the perimeter DNS zone for the federation server proxy.

Note

It is assumed that a host (A) resource record for the federation server has already been created in the corporate network DNS. If this record does not yet exist, create this record and then perform these procedures. For more information about how to create a host (A) resource record for the federation server, see Add a Host (A) Resource Record to Corporate DNS for a Federation Server.

Add a host (A) resource record to the Internet DNS zone for a federation server proxy

So that client computers on the Internet can successfully access a federation server through a newly deployed federation server proxy, you must first create a host (A) resource record in the Internet DNS zone that you control. This resource record resolves the host name of the account federation server (for example, fs.fabrikam.com) to the IP address of the account federation server proxy (for example, 131.107.27.68) in the perimeter network.

Note

It is assumed that you are using a DNS server running Windows 2000 Server, Windows Server 2003, or Windows Server 2008 with the DNS Server service to control the Internet DNS zone.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a host (A) resource record to the Internet DNS zone for a federation server proxy

1. On a DNS server for the Internet DNS zone, open the DNS snap-in.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A or AAAA)**.
3. In **Name**, type only the computer name of the federation server. For example, for the fully qualified domain name (FQDN) fs.fabrikam.com, type **fs**.
4. In **IP address**, type the IP address for the new federation server proxy, for example, 131.107.27.68.
5. Click **Add Host**.

Add a host (A) resource record to the perimeter DNS zone for a federation server proxy

So that Internet client requests can be processed successfully by the federation server proxy and reach the federation server after they are resolved by the Internet DNS zone, you must create a host (A) resource record in the perimeter DNS zone. This resource record resolves the host name of the account federation server (for example, fs. fabrikam.com) to the IP address of the account federation server (for example, 192.168.1.4) in the corporate network.

Note
It is assumed that you are using a DNS server running Windows 2000 Server, Windows Server 2003, Windows Server 2008, or Windows Server® 2012 with the DNS Server service to control the perimeter DNS zone.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a host (A) resource record to the perimeter DNS zone for a federation server proxy

1. On a DNS server for the perimeter network, open the **DNS snap-in**.
2. In the console tree, right-click the applicable forward lookup zone, and then click **New Host (A or AAAA)**.
3. In **Name**, type only the computer name of the federation server. For example, for the FQDN fs.fabrikam.com, type **fs**.
4. In the **IP address** text box, type the IP address for the federation server in the corporate network, for example, 192.168.1.4.
5. Click **Add Host**.

Additional references

Checklist: Setting Up a Federation Server Proxy

Name Resolution Requirements for Federation Server Proxies

Export the Private Key Portion of a Server Authentication Certificate

Every federation server in an Active Directory Federation Services (AD FS) farm must have access to the private key of the server authentication certificate. If you are implementing a server farm of federation servers or Web servers, you must have a single authentication certificate. This certificate must be issued by an enterprise certification authority (CA), and it must have an exportable private key. The private key of the server authentication certificate must be exportable so that it can be made available to all the servers in the farm.

This same concept is true of federation server proxy farms in the sense that all federation server proxies in a farm must share the private key portion of the same server authentication certificate.

Note

The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Depending on which role this computer will play, use this procedure on the federation server computer or federation server proxy computer where you installed the server authentication certificate with the private key. When you finish the procedure, you can then import this certificate on the Default Web Site of each server in the farm. For more information, see Import a Server Authentication Certificate to the Default Web Site.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To export the private key portion of a server authentication certificate

1. On the **Start** screen, type **Internet Information Services (IIS) Manager**, and then press ENTER.
2. In the console tree, click **ComputerName**.
3. In the center pane, double-click **Server Certificates**.
4. In the center pane, right-click the certificate that you want to export, and then click **Export**.
5. In the **Export Certificate** dialog box, click the **…** button.
6. In **File name**, type **C:***NameofCertificate*, and then click **Open**.
7. Type a password for the certificate, confirm it, and then click **OK**.

8. Validate the success of your export by confirming that the file you specified is created at the specified location.

   Important
   So that this certificate can be imported to the local certificate store on the new server, you must transfer the file to physical media and protect its security during transport to the new server. It is extremely important to guard the security of the private key. If this key is compromised, the security of your entire AD FS deployment (including resources within your organization and in resource partner organizations) is compromised.

9. Import the exported server authentication certificate into the certificate store on the new server before you install the Federation Service. For information about how to import the certificate, see Import a Server Certificate (http://go.microsoft.com/fwlink/?LinkId=108283).

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Certificate Requirements for Federation Servers

Certificate Requirements for Federation Server Proxies

Import a Server Authentication Certificate to the Default Web Site

After you obtain a server authentication certificate from a certification authority (CA), you must manually install that certificate on the Default Web Site for each federation server or federation server proxy in a server farm.

For Web servers, you must manually install the server authentication certificate on the appropriate Web site or virtual directory where your federated application resides.

If you are setting up a farm, be sure to perform this procedure identically—using the exact same settings—on each of the servers in your farm.

Note
The AD FS Management snap-in refers to server authentication certificates for federation servers as service communication certificates.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To import a server authentication certificate to the Default Web Site

1. On the **Start** screen, type **Internet Information Services (IIS) Manager**, and then press ENTER.
2. In the console tree, click **ComputerName**.
3. In the center pane, double-click **Server Certificates**.
4. In the **Actions** pane, click **Import**.
5. In the **Import Certificate** dialog box, click the **…** button.
6. Browse to the location of the pfx certificate file, highlight it, and then click **Open**.
7. Type a password for the certificate, and then click **OK**.

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Certificate Requirements for Federation Servers

Certificate Requirements for Federation Server Proxies

Install the Federation Service Proxy Role Service

After you configure a computer with the prerequisite applications and certificates, you are ready to install the Federation Service Proxy role service of Active Directory Federation Services (AD FS). You can use the following procedure to install the Federation Service Proxy role service. When you install the Federation Service Proxy role service on a computer, that computer becomes a federation server proxy.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To install the Federation Service Proxy role service

1. On the **Start** screen, type **Server Manager**, and then press ENTER.
2. Click **Manage**, and then click **Add Roles and Features** to start the Add Roles and Features Wizard.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or Feature-based installation**, and click **Next**.

5. On the **Select destination server** page, click **Select a server from the server pool**, verify that the target computer is highlighted, and then click **Next**.
6. On the **Select server roles** page, click **Active Directory Federation Services**, and then click next.

   Note
   If you are prompted to install additional .NET Framework or Windows Process Activation Service features, click **Add Features** to install them.
7. On the **Select features** page, verify that the features are set, and then click **Next**.
8. On the **Active Directory Federation Service (AD FS)** page, click **Next**.
9. On the **Select role services** page, select the **Federation Service Proxy** check box, and then click **Next**.
10. On the **Web Server Role (IIS)** page, click **Next**.
11. On the **Select role services** page, click **Next**.
12. After you verify the information on the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.
13. On the **Installation progress** page, verify that everything installed correctly, and then click **Close**.

Additional references

Checklist: Setting Up a Federation Server

Checklist: Setting Up a Federation Server Proxy

Configure a Computer for the Federation Server Proxy Role

After you configure a computer with the required certificates and have installed the Federation Service Proxy role service, you are ready to configure the computer to become a federation server proxy. You can use the following procedure so that the computer acts in the federation server proxy role.

Important
Before you use this procedure to configure the federation server proxy computer, make sure that you have followed all the steps in Checklist: Setting Up a Federation Server Proxy in the order that they are listed. Make sure that at least one federation server is deployed and that all the necessary credentials for authorizing a federation server proxy configuration are implemented. You must also configure Secure Sockets Layer (SSL) bindings on the Default Web Site, or this wizard will not start. All these tasks must be completed before this federation server proxy can function.

After you finish setting up the computer, verify that the federation server proxy is working as expected. For more information, see Verify That a Federation Server Proxy Is Operational.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure a computer for the federation server proxy role

1. There are two ways to start the AD FS Federation Server Configuration Wizard. To start the wizard, do one of the following:
   - On the **Start** screen, type **AD FS Federation Server Proxy Configuration Wizard**, and then press ENTER.
   - Anytime after the setup wizard is complete, open Windows Explorer, navigate to the **C:\Windows\ADFS** folder, and then double-click **FspConfigWizard.exe**.
2. Using either method, start the wizard, and on the **Welcome** page, click **Next**.
3. On the **Specify Federation Service Name** page, under **Federation Service name**, type the name that represents the Federation Service for which this computer will act in the proxy role.
4. Based on your specific network requirements, determine whether you will need to use an HTTP proxy server to forward requests to the Federation Service. If so, select the **Use an HTTP proxy server when sending requests to this Federation Service** check box, under **HTTP proxy server address** type the address of the proxy server, click **Test Connection** to verify connectivity, and then click **Next**.
5. When you are prompted, specify the credentials that are necessary to establish a trust between this federation server proxy and the Federation Service.

   By default, only the service account used by the Federation Service or a member of the local BUILTIN\Administrators group can authorize a federation server proxy.
6. On the **Ready to Apply Settings** page, review the details. If the settings appear to be correct, click **Next** to begin configuring this computer with these proxy settings.
7. On the **Configuration Results** page, review the results. When all the configuration steps are finished, click **Close** to exit the wizard.

   There is no Microsoft Management Console (MMC) snap-in to use for administering federation server proxys. To configure settings for each of the federation server proxys in your organization, use Windows PowerShell cmdlets.

Configuring an Alternate TCP/IP Port for Proxy Operations

By default, the federation server proxy service is configured to use TCP port 443 for HTTPS traffic and port 80 for HTTP traffic for communication with the federation server. To configure

different ports, such as TCP port 444 for HTTPS and port 81 for HTTP, the following tasks must be completed.

Note

If you intend to initially deploy AD FS to operate under alternate TCP/IP ports, you should first modify ports in your IIS protocol bindings for HTTP and HTTPS on both the federation server and federation server proxy computers. This should occur before you run the AD FS configuration wizards for initial configuration. If you configure Internet Information Services (IIS) first, your alternate TCP/IP port settings are discovered when wizard-based configuration occurs within AD FS, and the following procedure is not necessary. If you want to change the port settings later, update IIS protocol bindings first, and then use the following procedure to update port settings appropriately. For more information about editing IIS bindings, see article 149605 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=190275).

To configure alternate TCP/IP ports for the federation server proxy to use

1. Configure the federation server to use the nondefault ports.

   To do this, specify the nondefault port number by including it with the *HttpsPort* and *HttpPort* options as part of the **Set-ADFSProperties** cmdlet. For example, to configure these ports, use the following commands in the Windows PowerShell session on the federation server computer:

   ```
   Set-ADFSProperties -HttpsPort 444
   Set-ADFSProperties -HttpPort 81
   ```

2. Configure the federation server proxy to use the nondefault port.

   To do this, specify the nondefault port number by including it with the *HttpsPort* and *HttpPort* options as part of the **Set-ADFSProxyProperties** cmdlet. For example, to configure these ports, use the following commands in the Windows PowerShell session on the federation server computer:

   ```
   Set-ADFSProxyProperties -HttpsPort 444
   Set-ADFSProxyProperties -HttpPort 81
   ```

   Note

   Endpoint URLs are not enabled by default for the federation server proxy service. If you are configuring a new federation server installation, you must enable federation server proxy service endpoints first. For example, it is assumed that for all the endpoints that the example in this procedure refers to you have enabled them for proxy by selecting them in the AD FS Management snap-in and then selecting **Enable on proxy**.

3. Update the IIS installation at the federation server proxy so that Security Assertion Markup Language (SAML) and WS-Trust endpoints are configured to reflect the updated port number. To do this, you can use Notepad to modify the following in the Web.config file, which is located at systemdrive%\inetpub\adfs\ls\ on the federation server proxy

computer. For example, assuming that you have a federation server named sts1.contoso.com and the new port number is 444, browse to and open the Web.config file in Notepad on the federation server proxy computer, locate the following section, modify the port number as highlighted below, and then save and exit Notepad.

```
<securityTokenService
samlProtocolEndpoint="https://sts1.contoso.com:444/adfs/servi
ces/trust/samlprotocol/proxycertificatetransport"
      wsTrustEndpoint="https://sts1.contoso.com:444/adfs/serv
ices/trust/proxycertificatetransport" />
```

4. Add the federation server proxy service user account to the access control list (ACL) for the related endpoint URLs. For example, if the port number is 1234 and the user account that is used to run the AD FS federation server proxy service under is the built-in Network Service account, type the following command at a command prompt:

```
netsh http add urlacl
https://+:444/adfs/fs/federationserverservice.asmx/ user="NT
Authority\Network Service"
netsh http add urlacl https://+:444/FederationMetadata/2007-
06/ user="NT Authority\Network Service"
netsh http add urlacl https://+:444/adfs/services/ user="NT
Authority\Network Service"


netsh http add urlacl http://+:81/adfs/services/ user="NT
Authority\Network Service"
```

The previous commands must be run on both the federation server and the federation server proxy computers.

Additional references

Checklist: Setting Up a Federation Server Proxy

Verify That a Federation Server Proxy Is Operational

You can use the following procedure to verify that the federation server proxy can communicate with the Federation Service in Active Directory Federation Services (AD FS). You run this procedure after you run the **AD FS Federation Server Proxy Configuration Wizard** to configure the computer to run in the federation server proxy role. For more information about how to run this wizard, see Configure a Computer for the Federation Server Proxy Role.

Important

The result of this test is the successful generation of a specific event in Event Viewer on the federation server proxy computer.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To verify that a federation server proxy is operational

1. Log on to the federation server proxy as an administrator.
2. On the **Start** screen, type **Event Viewer**, and then press ENTER.
3. In the details pane, double-click **Applications and Services Logs**, double-click **AD FS Eventing**, and then click **Admin**.
4. In the **Event ID** column, look for event ID 198.

   If the federation server proxy is configured properly, you see a new event in the Application log of Event Viewer, with the event ID 198. This event verifies that the federation server proxy service was started successfully and now is online.

Additional references

Checklist: Setting Up a Federation Server Proxy

Configure Performance Monitoring

AD FS includes its own dedicated performance counters to help you monitor the performance of both federation servers and federation server proxy computers. To use Performance Monitor to monitor the performance of your AD FS servers, it's useful to create a new data collector set and add the AD FS counters to that view. The following procedure describes how to configure performance monitoring for AD FS.

To configure performance monitoring for AD FS using Performance Monitor

1. On the **Start** screen, type **Performance Monitor**, and then press ENTER.
2. In the console tree, expand **Data Collector Sets**, right-click **User Defined**, point to **New**, and then click **Data Collector Set**.

   The Create New Data Collector Set Wizard appears.
3. In **Create New Data Collector Set**, for **Name** type a name for the new data collector set (such as "AD FS performance"), click **Create manually (Advanced)**, and then click **Next**.

4. For the type of data to include, verify that **Create data logs** is selected, and then click the check boxes for the following data types: **Performance counter**, **Event trace data**, **System configuration information**.
5. For performance counters, expand **AD FS** in the **Available counters** list, and then click **Add**.

   The AD FS performance counters should appear in the **Added counters** list.
6. When you are prompted to add event trace providers, click **Add**, select **AD FS Eventing** and **AD FS Tracing** from the list of providers.
7. When you are prompted to add registry keys to monitor, click **Next**.
8. When you are prompted to specify the location to save the performance data, you can accept the default location (**%systemdrive%\PerfLogs\Admin**<*data_collector_set*>, and then click **Next**.
9. When you are prompted to create the data collector set, select **Save and close**, and then click **Finish**.

   The new data collector set appears in the console tree under the **User Defined** node.
10. Use the following steps to work with the AD FS performance counters:
    - To begin performance monitoring using AD FS-related counters, right-click the data collector set that you added (such as "AD FS performance"), and then click **Start**.
    - To create a report to view the performance monitoring results, right-click the data collector set that you added (such as "AD FS performance"), and then click **Latest Report**.
    - To end a capture of performance data so that you can view the latest report, right-click the data collector set that you added (such as "AD FS performance"), and then click **Stop**.

    The latest report is added and numbered automatically (starting at 000001) under the **Report\User Defined**\<*data_collector_set*> node in the console tree.

AD FS performance counters

The following table lists the AD FS performance counters and describes how they are useful for monitoring activity that relates to either a federation server or federation server proxy.

Counter	Description	Can be used on:
Token Requests	Monitors the number of token requests sent to the federation server including SSOAuth token requests.	Federation Servers
Token Requests/sec	Monitors the number of token requests sent to the federation server including SSOAuth token requests per second.	Federation Servers

Counter	Description	Can be used on:
Federation Metadata Requests	Monitors the number of incoming federation metadata requests sent to the federation server.	Federation Servers
Federation Metadata Requests/sec	Monitors the number of incoming federation metadata requests per second that are sent to the federation server.	Federation Servers
Artifact Resolution Requests	Monitors the number of incoming federation metadata requests per second that are sent to the federation server.	Federation Servers
Artifact Resolution Requests/sec	Monitors the number of requests to the artifact resolution endpoint per second that are sent to the federation server.	Federation Servers
Proxy Requests	Monitors the number of incoming requests sent to the federation server proxy.	Federation Server Proxies
Proxy Requests/sec	Monitors the number of incoming requests per second that are sent to the federation server proxy.	Federation Server Proxies
Proxy MEX Requests	Monitors the number of incoming WS-Metadata Exchange (MEX) requests that are sent to the federation server proxy.	Federation Server Proxies
Proxy MEX Requests/sec	Monitors the number of incoming MEX requests per second that are sent to the federation server proxy.	Federation Server Proxies

Interoperating with AD FS 1.x

For interoperability between Active Directory Federation Services (AD FS) in Windows Server® 2012 and AD FS 1.*x*, complete one or more of the following tasks, depending on the needs of your organization:

- Plan for interoperability between AD FS in Windows Server 2012 and previous versions of AD FS, and learn more about the Name ID claim type. For more information, see Planning for Interoperability with AD FS 1.x.
- If you will be sending claims from an AD FS Federation Service in Windows Server 2012 that can be consumed by an AD FS 1.*x* Federation Service, see Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Federation Service.
- If you will be sending claims from an AD FS Federation Service in Windows Server 2012 that can be consumed by an application that is hosted by a Web server running the AD FS 1.*x* claims-aware Web agent, see Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Claims-Aware Web Agent.
- If you will be sending claims from an AD FS 1.*x* Federation Service to be consumed by an AD FS Federation Service in Windows Server 2012, see Checklist: Configuring AD FS to Consume Claims from AD FS 1.x.

Differences between Federation Service settings

Although most of the AD FS 1.*x* Federation Service settings work in a similar manner as the AD FS Federation Service in Windows Server 2012 settings, some setting names have changed. The following table lists the names of settings for an AD FS 1.*x* Federation Service and their equivalent names for an AD FS Federation Service in Windows Server 2012.

AD FS 1.x Federation Service setting	Equivalent AD FS Federation Service in Windows Server 2012 setting
Account Partner	Claims provider trust
Resource Partner	Relying party trust
Application	Relying party trust
Application Properties	Relying party trust properties
Application URL	Relying party identifier and WS-Federation Passive Endpoint URL
Federation Service URI	Federation Service identifier
Federation Service endpoint URL	WS-Federation Passive Endpoint URL

See Also

AD FS and AD FS 1.x Interoperability

Checklist: Configuring AD FS to Consume Claims from AD FS 1.x

Checklist: Configuring AD FS to consume claims from AD FS 1.x

This checklist includes the tasks that are necessary for configuring your Active Directory Federation Services (AD FS) Federation Service in Windows Server 2012 to consume claims that are sent by an AD FS 1.*x* Federation Service.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Configuring AD FS to consume claims from AD FS 1.x

	Task	Reference
☐	Plan for interoperability between AD FS in Windows Server 2012 and previous versions of AD FS, and learn more about the Name ID claim type.	Planning for Interoperability with AD FS 1.x
☐	Before you can interoperate with a previous version of AD FS, you must first create a claims provider trust in the AD FS Federation Service. **Note** You cannot create a trust with an AD FS 1.*x* Federation Service by using federation metadata. When you set up the trust using the procedure in the link to the	Create a Claims Provider Trust Manually

	Task	Reference
	right, you must do the following in the Add Claims Provider Trust Wizard to set up this trust to interoperate with an AD FS 1.*x* Federation Service: 1. On the **Select Data Source** page, select **Enter data about the relying party trust manually**. 2. On the **Choose Profile** page, select **AD FS 1.0 and 1.1 profile**. 3. On the **Configure URL** page, under **WS-Federation Passive URL**, type the **Federation Service endpoint URL** as defined in the AD FS 1.*x* Federation Service of the partner. 4. On the **Configure Identifiers** page, under **Claims provider trust identifier**, type the **Federation Service URI** as defined in the AD FS 1.*x* Federation Service of the partner.	
☐	On the claims provider trust that you created earlier, you must create a claim rule that will take claims that are incoming from the AD FS 1.x Federation Service and pass through, filter, or transform them into a Name ID claim type. When the Name ID claim type has been passed through, filtered, or transformed, it can be used as input to another rule or rules so that it can be understood and consumed by the AD FS Federation Service in Windows Server 2012.	Create a Rule to Send an AD FS 1.x Compatible Claim

	Task	Reference
☐	Contact the administrator of the AD FS 1.*x* Federation Service and have the administrator of the AD FS 1.*x* Federation Service set up a new resource partner trust. Also, provide that administrator with the Federation Service URI (in the Federation Service properties), the Federation Service endpoint URL, and an exported token-signing certificate file (with public key only). The administrator will need these items to set up the trust.	N/A

Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Federation Service

Checklist: Configuring AD FS to send claims to an AD FS 1.x Federation Service

This checklist includes the tasks that are necessary for configuring your Active Directory Federation Services (AD FS) Federation Service in Windows Server 2012 to send claims that can be understood by an AD FS 1.*x* Federation Service.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Configuring AD FS to send claims to an AD FS 1.x Federation Service

	Task	Reference
☐	Plan for interoperability between AD FS in Windows Server 2012 and previous versions of AD FS and learn more about the Name ID claim type.	Planning for Interoperability with AD FS 1.x

	Task	Reference
☐	Before you can achieve interoperability with a previous version of AD FS, you must first create a relying party trust in the AD FS Federation Service to the AD FS 1.*x* Federation Service. **Note** You cannot create a trust with an AD FS 1.*x* Federation Service by using federation metadata. When you set up the trust using the procedure in the link to the right, you must do the following in the Add Relying Party Trust Wizard to set up this trust to interoperate with an AD FS 1.*x* Federation Service: 1. On the **Select Data Source** page, select **Enter data about the relying party trust manually**. 2. On the **Choose Profile** page, select **AD FS 1.0 and 1.1 profile**. 3. On the **Configure URL** page, under **WS-Federation Passive URL**, type the **Federation Service endpoint URL** as defined in the AD FS 1.*x* Federation Service of the partner. 4. On the **Configure Identifiers** page, under **Relying part trust identifier**, type the **Federation Service URI** as defined in the AD FS 1.*x* Federation Service of the partner.	Create a Relying Party Trust Manually
☐	On the relying party trust that you	Create a Rule to Send an AD

	Task	Reference
	created earlier, you must create claim rules that will take incoming claims that were extracted from an attribute store and pass through, filter, or transform them into a Name ID claim type that can be understood and consumed by the AD FS 1.*x* Federation Service. **Note** Before you create this rule, make sure that the claim rule set where you are creating this rule has a rule that comes before it that first extracts a Lightweight Directory Access Protocol (LDAP) attribute claim from an attribute store. This claim will be used as input to the rule that you create to send an AD FS 1.*x*-compatible claim. For more information about how to create a rule to extract an LDAP attribute, see Create a Rule to Send LDAP Attributes as Claims.	FS 1.x Compatible Claim
☐	Contact the administrator of the AD FS 1.*x* Federation Service and have the administrator of the AD FS 1.*x* Federation Service set up a new account partner trust. Also, provide that administrator with the Federation Service URI (in the Federation Service properties), the WS-Federation Passive endpoint URL (the Federation Service endpoint	N/A

	Task	Reference
	URL), and an exported token-signing certificate file (with public key only). That administrator will need these items to set up the trust.	

Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Claims-Aware Web Agent

Checklist: Configuring AD FS to send claims to an AD FS 1.x claims-aware Web agent

This checklist includes the tasks that are necessary for configuring your Active Directory Federation Services (AD FS) Federation Service in Windows Server 2012 to send claims that can be understood by an application that is hosted by a Web server running the AD FS 1.*x* claims-aware Web agent.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

Checklist: Configuring AD FS to send claims to an AD FS 1.x claims-aware Web agent

	Task	Reference
☐	Plan for interoperability between AD FS in Windows Server 2012 and previous versions of AD FS and learn more about the Name ID claim type.	Planning for Interoperability with AD FS 1.x
☐	If you have not already done so, use the link on the right to first create a relying party trust between the AD FS Federation Service in Windows Server 2012 and the AD FS 1.*x* Federation Service.	Checklist: Configuring AD FS to Send Claims to an AD FS 1.x Federation

	Task	Reference
		Service
☐	Before you can achieve interoperation with an application that is hosted by the AD FS 1.*x* claims-aware Web agent, you must first create a relying party trust in the AD FS Federation Service in Windows Server 2012 to the AD FS 1. *x* claims-aware Web agent. **Note** Creating this trust in the AD FS Federation Service is the equivalent of adding a new **Application** to the AD FS 1.x Federation Service (**Federation Service\Trust Policy\My Organization\Application**). This relying party trust is necessary because AD FS does not have an equivalent **Application** node in its own snap-in. However, it still must have a secure channel to the application. When you set up the trust using the procedure in the link to the right, you must do the following in the Add Relying Party Trust Wizard to set up this trust to interoperate with an AD FS 1.*x* claims-aware Web agent: 1. On the **Select Data Source** page, select **Enter data about the relying party trust manually**. 2. On the **Choose Profile** page, select **AD FS 1.0 and 1.1 profile**. 3. On the **Configure URL** page, under **WS-Federation Passive URL**, type the **Application URL** as defined in the AD FS 1.*x* Federation Service of the partner. 4. On the **Configure Identifiers** page, under **Relying part trust identifier**, type the **Application URL** as defined in the AD FS 1.*x* claims-aware Web agent	Create a Relying Party Trust Manually
☐	Contact the administrator of the Web server running the AD FS 1.*x* claims-aware Web agent and have that administrator edit the web.config file that is associated with the claims-aware application (under the Default Web Site in Internet Information Services (IIS)) to point the Web agent at the AD FS Federation Service. For example, replace *myresourcefederationserver* in the tag `<fs>https://myresourcefederationserver/adfs/fs/federationserverservice.asmx</fs>` of the web.config file with a valid AD FS federation server name. This is necessary for the application and AD FS 1.x claims-aware Web agent to be able to consume the claims that are sent to it from the AD FS Federation Service in Windows Server 2012.	N/A
☐	On the relying party trust that you created earlier, you have to create claim rules that will take incoming claims that were extracted from an attribute store and pass through, filter, or transform them into a Name ID claim type that can	Create a Rule to Send an AD FS 1.x

	Task	Reference
	be understood and consumed by the AD FS 1.*x* claims-aware Web agent. **Note** Before you create this rule, make sure that the claim rule set where you are creating this rule has a rule that comes before it that first extracts a Lightweight Directory Access Protocol (LDAP) attribute claim from an attribute store. This claim will be used as input to the rule that you create to send an AD FS 1.*x*-compatible claim. For more information about how to create a rule to extract an LDAP attribute, see Create a Rule to Send LDAP Attributes as Claims.	Compatible Claim

Create a Relying Party Trust Manually

To add a new relying party trust using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a federation server in the account partner organization.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a relying party trust manually

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. Under **AD FS\Trust Relationships**, right-click **Relying Party Trusts**, and then click **Add Relying Party Trust** to open the Add Relying Party Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter data about the relying party manually**, and then click **Next**.
5. On the **Specify Display Name** page type a name in **Display name**, under **Notes** type a description for this relying party trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS profile**, click **Next**, and then move to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then go to step 9.

   If you know you will require interoperability with older Active Directory Federation Services (AD FS) federation, as provided in Windows Server 2003 R2, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile**.
7. On the **Configure Certificate** page, click **Browse** to locate a certificate file, and then click **Next**.

8. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 10:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Relying party WS-Federation Passive protocol URL**, type the URL for this relying party trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Relying party SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this relying party trust, and then click **Next**.

   Click the **Help** button on this page for more information about which of these options apply to the needs of your organization.
9. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this relying party trust, and then click **Next**.
10. On the **Configure Identifiers** page, specify one or more identifiers for this relying party, click **Add** to add them to the list, and then click **Next**.
11. On the **Choose Issuance Authorization Rules** page, select either **Permit all users to access this relying party** or **Deny all users access to this relying party**, and then click **Next**.
12. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.
13. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this relying party trust, see Additional references.

Additional references

Checklist: Configuring the Account Partner Organization

Checklist: Creating Claim Rules for a Relying Party Trust

Create a Claims Provider Trust Manually

To add a new claims provider trust using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a resource partner federation server in the resource partner organization.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a claims provider trust manually

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.

2. Under the **AD FS\Trust Relationships**, right-click **Claims Provider Trusts**, and then click **Add Claims Provider Trust** to open the Add Claims Provider Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter claims provider trust data manually**, and then click **Next**.
5. On the **Specify Display Name** page type a **Display name**, under **Notes** type a description for this claims provider trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS Profile**, click **Next**, and then go to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then move to step 8.

   If you know that you will require interoperability between this claims provider trust and other, older Active Directory Federation Services (AD FS) claims provider trusts, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile** option.
7. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 9:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Claims provider WS-Federation Passive protocol URL**, type the URL for this claims provider trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Claims provider SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this claims provider trust, and then click **Next**.

   Click the **Help** button on this page for more information about which of these options apply to the needs of your organization.
8. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this claims provider trust, and then click **Next**.
9. On the **Configure Identifier** page, under **Claims provider trust identifier**, type the appropriate identifier, and then click **Next**.
10. On the **Configure Certificates** page, click **Add** to locate a certificate file and add it to the list of certificates, and then click **Next**.
11. On the **Ready to Add Trust** page, click **Next** to save your claims provider trust information.
12. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this claims provider trust, see the Additional references.

Additional references

Checklist: Configuring the Resource Partner Organization

Checklist: Creating Claim Rules for a Claims Provider Trust

Create a Rule to Send an AD FS 1.x Compatible Claim

In situations in which you are using Active Directory Federation Services (AD FS) to issue claims that will be received by federation servers running AD FS 1.0 (Windows Server 2003 R2) or AD FS 1.1 (Windows Server 2008 or Windows Server 2008 R2), you must do the following:

- Create a rule that will send a Name ID claim type with a format of UPN, Email, or Common Name.
- All other claims that are sent must have one of the following claim types:
  - AD FS 1.*x* Email Address
  - AD FS 1.*x* UPN
  - Common Name
  - Group
  - Any other claim type that begins with https://schemas.xmlsoap.org/claims/, such as https://schemas.xmlsoap.org/claims/EmployeeID

Depending on the needs of your organization, use one of the following procedures to create an AD FS 1.*x* compatible NameID claim:

- Create this rule to issue an AD FS 1.x Name ID claim using the Pass Through or Filter an Incoming Claim rule template
- Create this rule to issue an AD FS 1.x Name ID claim using the Transform an Incoming Claim rule template. You can use this rule template in situations in which you want to change the existing claim type to a new claim type that will work with AD FS 1. *x* claims.

Note

For this rule to work as expected, make sure that the relying party trust or claims provider trust where you are creating this rule has been configured to use the **AD FS 1.0 and 1.1 profile**.

Creating a rule to issue an AD FS 1.

You can use this rule template when you want the group that you specify in this template to ultimately determine what AD FS 1.*x*-compatible claim to issue.

To create a rule to issue an AD FS 1.x Name ID claim using the Pass Through or Filter an Incoming Claim rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust

that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:

- **Acceptance Transform Rules**
- **Issuance Transform Rules**
- **Issuance Authorization Rules**
- **Delegation Authorization Rules**

5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select **Name ID** in the list.
8. In **Incoming name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
9. Select one of the following options, depending on the needs of your organization:
   - **Pass through all claim values**
   - **Pass through only a specific claim value**
   - **Pass through only claim values that match a specific email suffix value**
   - **Pass through only claim values that start with a specific value**
10. Click **Finish**, and then click **OK** to save the rule.

Creating a rule to issue an AD FS 1.

You can use this rule template in situations in which you want to change the incoming claim type to a new outgoing claim type of Name ID that will work with AD FS 1.*x* federation servers.

To create a rule to issue an AD FS 1.x Name ID claim using the Transform an Incoming Claim rule template

1. On the **Start** screen, type **AD FS Management**, and then press ENTER.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, depending on the trust that you are editing and which rule set you want to create this rule in, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**

- **Delegation Authorization Rules**

5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select the type of incoming claim that you want to transform in the list.
8. In **Outgoing claim type**, select **Name ID** in the list.
9. In **Outgoing name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
10. Select one of the following options, depending on the needs of your organization:
    - **Pass through all claim values**
    - **Replace an incoming claim value with a different outgoing claim value**
    - **Replace incoming e-mail suffix claims with a new e-mail suffix**
11. Click **Finish**, and then click **OK** to save the rule.

AD FS Operations

- Customizing the AD FS Sign-in Pages
- Add an Attribute Store
- Create a Relying Party Trust
- Create a Non-Claims-Aware Relying Party Trust
- Create a Claims Provider Trust
- Configuring Claim Rules
- Configuring Authentication Policies
- Configuring AD FS Extranet Lockout
- Configuring Client Access Policies

See Also

Active Directory Federation Services Overview

Customizing the AD FS Sign-in Pages

Overview

With this release of Active Directory Federation Services (AD FS), several changes have been made to improve the experience for both administrators and users. Numerous capabilities to enable administrators to customize the sign-in experience have been added to AD FS. The new design provides Windows PowerShell commands for administrators to customize standard elements for common scenarios. In addition, all of the customizations are stored in the AD FS configuration store that is commonly shared across the AD FS service. As a result, you only have to do this task one time for the AD FS service and do not have to repeat this task for every node in the AD FS farm. The following is a list of improvements that have been added to AD FS:

- **No more dependence on IIS** - AD FS is now built directly on top of HTTP.SYS and does not require the installation of Internet Information Services (IIS).
- **Built-in webpage customization** – The appearance and behavior of webpages can be customized to provide a tailored user experience. These customizations include the following:
  - **Logo and illustration** - The logo and illustration of pages can be easily updated and changed with a simple Windows PowerShell command. It is no longer necessary to modify any of the webpage code in AD FS.
  - **Customization of standard Links** - Most organizations have some notion of standard links for IT support, home page, and a privacy link. These elements can be configured and displayed on AD FS webpages.
  - **Customization of web themes** - At the heart of the customizations is the notion of a web theme. This feature enables the administrator to customize the entire appearance and behavior of the AD FS pages. Each web theme supports all of the elements that were described previously.
- **Customization for error pages** - When something goes wrong, an error page is displayed to the end user. A few enhancements have been made to provide a better user experience. For example, AD FS now support sending error details to an IT help desk if required by using one click.
  - **Customized "Access denied" messages** – End users often get confused when they get a generic "Access denied" message. AD FS now supports customized "Access denied" messages to enable organizations to provide customized instructions for its end users. This customization supports HTML content and can be set globally or on the basis of a relying party trust.
- **Responsive design to cater to different form factors** - AD FS now supports automatic adjustment of size and positioning of elements through cascading style sheets (CSS). For example, in the smartphone scenario where the form factor is limited in width, AD FS sign-in pages would automatically adjust and also hide the illustration.
- **Improved user experience for home realm discovery** –AD FS now supports home realm discovery by looking up organizational account suffixes that a claims provider supports or by

looking up the claims provider list that is configured for a relying party trust. AD FS also supports by-passing home realm discovery for intranet access.

- **Support for updating the password** – Users can change their password without having to connect to the corporate network.

Customizing the AD FS sign-in experience

The following topics describe how to customize the AD FS sign-in experience.

General sign-in page

Out of the box, you notice the appearance and behavior of the sign-in experience to be similar to the appearance and behavior that is offered by Windows Azure Active Directory. The name of the organization, which is obtained during setup, is displayed by default. This page, as most of the other pages, can be customized and tailored to your organization's requirements. The customization includes changing the logo, the illustration, adding additional descriptions, and links. The following illustrations show an example of the default sign-in page and a customized sign-in page. Use the following Windows PowerShell cmdlets to modify and customize your AD FS sign-in page.

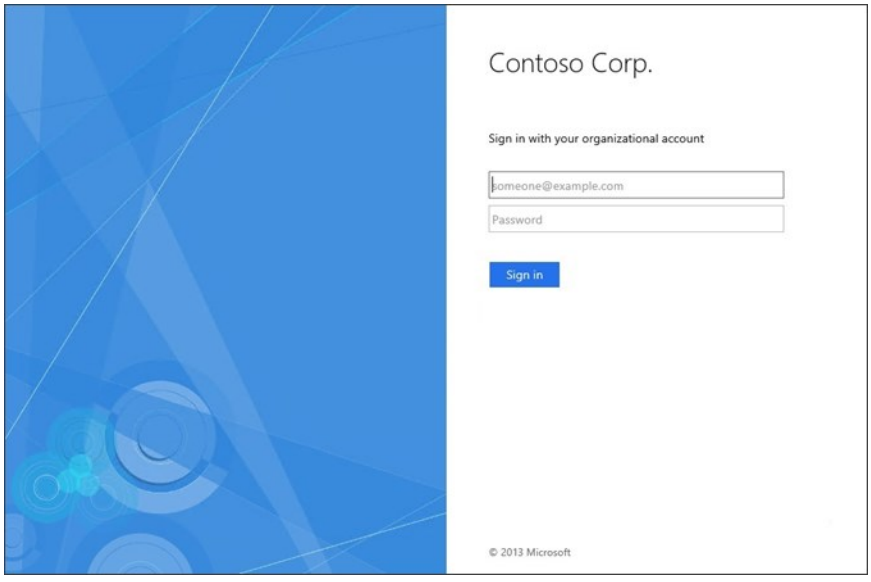

The following is an example of a customized sign-in page.

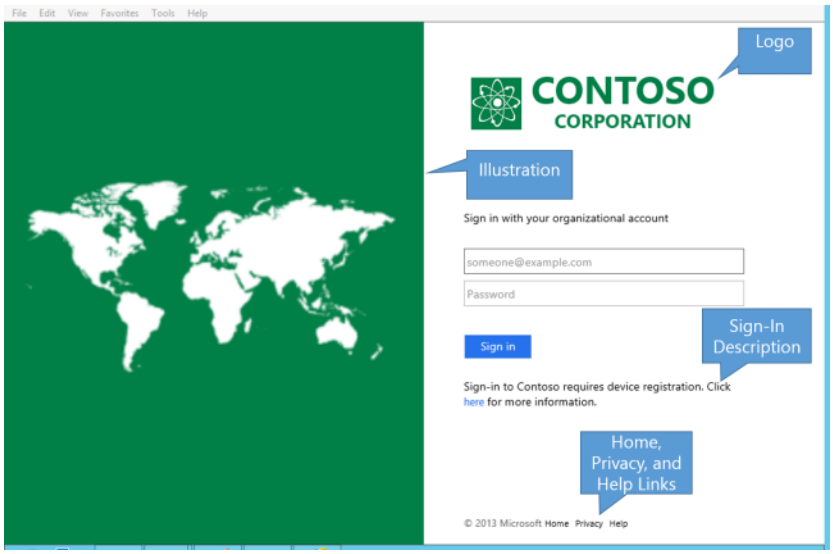

Change company name

To change the name of the company that is displayed on the sign-in page, use the following Windows PowerShell command and syntax. This value is set by default by using the value from the Federation Service display name that you entered during setup.

```
Set-AdfsGlobalWebContent –CompanyName "Contoso Corp"
```

Note

You can also use the Windows PowerShell Integrated Scripting Environment (ISE) to change the company name. By using the Windows PowerShell ISE, you can display content in a Unicode-compliant environment. For additional information, see Introducing the Windows PowerShell ISE.

Change company logo

To change the logo of the company that is displayed on the sign-in page, use the following Windows PowerShell cmdlet and syntax.

Important

We recommend the dimensions for the logo to be 260x35 @ 96 dpi with a file size of no greater than 10 KB.

```
Set-AdfsWebTheme -TargetName default -Logo @{path="c:\Contoso\logo.png"}
```

Note

The `TargetName` parameter is required. The default theme that is released with AD FS is named *default*.

Change Illustration

To change the illustration, the graphic on the left, which is displayed on the sign-in page, use the following Windows PowerShell cmdlet and syntax.

Important
We recommend the dimensions for the illustration to be 1420x1080 pixels @ 96 DPI with a file size of no greater than 200 KB.

```
Set-AdfsWebTheme -TargetName default -Illustration @{path="c:\Contoso\illustration.png"}
```

Add sign-in page description

To add a sign-in page description to the sign-in page, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -SignInPageDescriptionText "<p>Sign-in to Contoso requires
device registration. Click <A href='http://fs1.contoso.com/deviceregistration/'>here</A>
for more information.</p>"
```

Important
The string for the `SignInPageDescriptionText` parameter supports both pure HTML with the tags and without. Therefore, you can also run the following cmdlet without using the <p> tag. `Set-AdfsGlobalWebContent -SignInPageDescriptionText "Sign-in to Contoso requires device registration. Click <A href='http://fs1.contoso.com/deviceregistration/'>here</A> for more information."`

After the sign-in page is customized, the customization takes precedence; therefore, you should customize for all languages that you want to support. All customized content takes a locale parameter. When you configure localized content, it should be configured with a country-less locale first, for example, "en", before you configure country and region-specific locale such as "en-us".

Example:

Add Help Desk Link

To add the help desk link that is displayed on the sign-in page, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -HelpDeskLink https://fs1.contoso.com/help/ -HelpDeskLinkText
Help
```

Important
The `linkText` parameter in this cmdlet is not required unless you use another value than the default, which is *Help*. The advantage of using the default is that they are localized to

all client locales. After the page is customized, the customization takes precedence; therefore, you should customize for all languages that you want to support.

Add Home Link

To add the home link that is displayed on the sign-in page, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -HomeLink https://fs1.contoso.com/home/ -HomeLinkText Home
```

Important

The `linkText` parameter in this cmdlet is not required unless you use another value than the default, which is *Home*. The advantage of using the default is that they are localized to all client locales. after the sign-in page is customized, the customization takes precedence; therefore, you should customize for all languages that you want to support.

Add Privacy Link

To add the privacy link that is displayed on the sign-in page, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -PrivacyLink https://fs1.contoso.com/privacy/ -PrivacyLinkText
Privacy
```

Important

The `linkText` parameter in this cmdlet is not required unless you use another value than the default, which is *Privacy*. The advantage of using the default is that the pages are localized to all client locales. After the sign-in page is customized, the customization takes precedence; therefore, you should customize for all languages that you want to support. All customized content takes a locale parameter. When you configure localized content, you should configure it with a country-less locale first, for example, "en", before you configure country and region-specific locale, such as "en-us".

Example:

Custom Web Themes

The theme that is shipped out-of-the-box is called Default. You can export the default theme and use it so that you can start quickly. You can customize the appearance and behavior, which includes the layout by modifying the .css file, import and apply this new theme, and then you can use the customized appearance and behavior. Using the .css file also makes it easier to work with your web designers.

The following cmdlet creates a custom web theme, which duplicates the default web theme.

```
New-AdfsWebTheme –Name custom –SourceName default
```

You can modify the .css file and configure the new web theme by using the new .css file. To export a web theme, use the following cmdlet.

```
Export-AdfsWebTheme -Name default -DirectoryPath c:\theme
```

To apply the .css file to the new theme, use the following cmdlet.

```
Set-AdfsWebTheme -TargetName custom -StyleSheet @{path="c:\NewTheme.css"}
```

The following cmdlet creates a custom web theme from a new style sheet.

```
New-AdfsWebTheme -Name custom -StyleSheet @{path="c:\NewTheme.css"} -RTLStyleSheetPath
c:\NewRtlTheme.css
```

To apply the custom web theme to AD FS, use the following cmdlet.

```
Set-AdfsWebConfig -ActiveThemeName custom
```

To add JavaScript to AD FS, use the following cmdlet.

```
Set-AdfsWebTheme -TargetName custom -AdditionalFileResource @{Uri='
/adfs/portal/script/onload.js';path="D:\inetpub\adfsassets\script\onload.js"}
```

Custom Error Messages

You can configure custom error messages that can be tailored to your organization. The following illustration shows a customized error page description and a generic error message. Use the following Windows PowerShell cmdlets to customize your error messages.

Customize the error page description

To customize the error page description, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGloblaWebContent -ErrorPageDescriptionText "This is Contoso's error page description"
```

Customize a generic error message

To customize the generic error message, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -ErrorPageGenericErrorMessage "This is a generic error message. Contact Contoso IT for assistance."
```

Customize an authorization error message

To customize the authorization error message, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -ErrorPageAuthorizationErrorMessage "You have received an Authorization error.  Contact Contoso IT for assistance."
```

Customize a device authentication error message

To customize the device authentication error message, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -ErrorPageDeviceAuthenticationErrorMessage "Your device is not authorized.  Contact Contoso IT for assistance."
```

Customize a support email error message

You can configure a support email address in AD FS. If configured, AD FS automatically shows a link for end users to email the error details.

To customize the support email error message, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -ErrorPageSupportEmail  "admin@contoso.com"
```

Customize a relying party authorization message

You can configure a relying party authorization error message in AD FS.

To customize the relying party error message, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsRelyingPartyWebContent -Name fedpassive -ErrorPageAuthorizationErrorMessage "<p>
You need to be a member of Security Auditors to access this site. Click <A
href='http://accessrequest/'>here</A> for more information.</p>"
```

Home Realm Discovery

When the AD FS client first requests a resource, the resource federation server has no information about the realm of the client. The resource federation server responds to the AD FS client with a **Client Realm Discovery** page, where the user selects the home realm from a list. The list values are populated from the display name property in the Claims Provider Trusts. Use the following Windows PowerShell cmdlets to modify and customize the AD FS Home Realm Discovery experience.

Warning

Be aware that the Claims Provider name that shows up for local Active Directory is the Federation Service display name.

Configure Identity Provider to use certain email suffixes

An organization can federate with multiple claims providers. AD FS now provides the in-box capability for administrators to list the suffixes, for example, @us.contoso.com, @eu.contoso.com, that is supported by a claims provider and enable it for suffix-based discovery. With this configuration, end users can type in their organizational account, and AD FS automatically selects the corresponding claims provider.

To configure an identity provider (IDP), such as `fabrikam`, to use certain email suffixes, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsClaimsProviderTrust -TargetName fabrikam -OrganizationalAccountSuffix
@("fabrikam.com";"fabrikam2.com")
```

Configure an identity provider list per relying party

For some scenarios, an organizations might want end users to only see the claims providers that are specific to an application so that only a subset of claims provider are displayed on the home realm discovery page.

To configure an IDP list per relying party (RP), use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsRelyingPartyTrust -TargetName claimapp -ClaimsProviderName @("Fabrikam","Active
Directory")
```

Bypass Home Realm Discovery for the intranet

Most organizations only support their local Active Directory for any user who accesses from inside their firewall. In those cases, administrators can configure AD FS to bypass home realm discovery for the intranet.

To bypass HRD for the intranet, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsProperties -IntranetUseLocalClaimsProvider $true
```

Important

Please note that if an identity provider list for a relying party has been configured, even though the previous setting has been enabled and the user accesses from the intranet, AD FS still shows the home realm discovery (HRD) page. To bypass HRD in this case, you have to ensure that "Active Directory" is also added to the IDP list for this relying party.

Update password

In some instances, users might not be able to connect to the corporate network to change their account password. This factor can be problematic especially for remote employees who might live far from the nearest corporate office. For these specific cases, the update password page can be used by only connecting to the Internet.

You can customize the update password page by providing your own description for the page. This feature is only available for a workplace-joined device. See the following Windows PowerShell cmdlet on how to achieve this.

Warning

The update password page is only available for Workplace Joined devices and is disabled by default. See Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications for more information about Workplace Join and Configure a federation server with Device Registration Service for how to enable device registration service on AD FS.

To enable the password update page, go to AD FS Management under Endpoints. The endpoint for update password is located at the bottom under Other - /adfs/portal/updatepassword/. Once you have enabled the endpoint, you must restart the AD FS service. This must be done manually. You can then navigate to https://<fqdn>/adfs/portal/updatepassword/ on a workplace joined device and you should see the update password page.

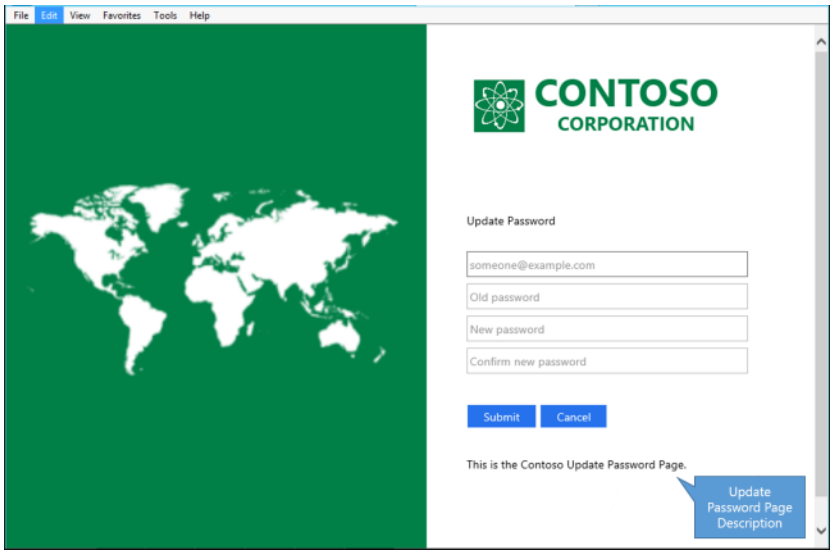

Customize the Update Password page description

To customize the update password page description, use the following Windows PowerShell cmdlet and syntax.

```
Set-AdfsGlobalWebContent -UpdatePasswordPageDescriptionText "This is the Contoso Update Password page."
```

Multi-factor authentication and external authentication providers

In AD FS, the support for multifactor authentication is provided out-of-the-box. For example, you can configure AD FS to use built-in Certificate Authentication as the second factor authentication. You can also use external authentication providers. This approach can enable AD FSto integrate with additional services, such as Windows Azure Active Authentication, or you can develop your own provider. See Solution Guide: Manage Risk with Multi-factor Access Control for more information about how to register external authentication provider by using AD FS.

We recommend that an external authentication provider use the classes that are defined in the .css file that AD FS provides to author the authentication UI. You can use the following cmdlet to export the default web theme and inspect the user interface classes and elements that are defined in the .css file. The .css file can be used in the development of the sign-in user interface of an external authentication provider.

```
Export-AdfsWebTheme -Name default -DirectoryPath C:\theme
```

The following is an example of the sign-in user interface, which is highlighted in RED, by an external authentication provider. The user interface uses the UI classes in the AD FS .css file.

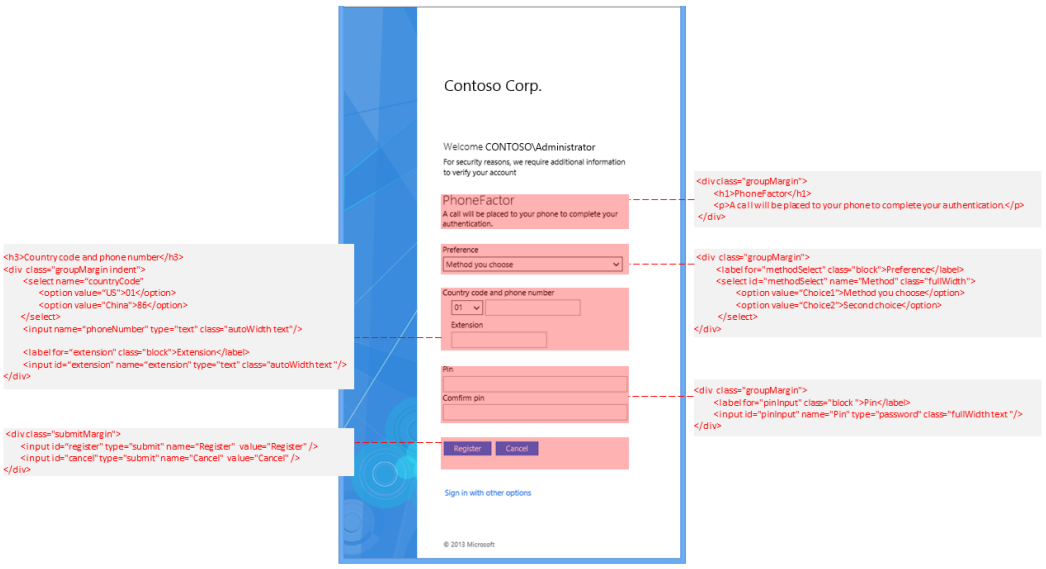

Before you write a new custom authentication method, we recommend that you study the AD FS theme and style definitions to understand the content authoring requirements.

- A custom authentication method only authors an HTML segment on the AD FS sign-in page and not the full page. You should use AD FS's style definition to get the consistent appearance and behavior.

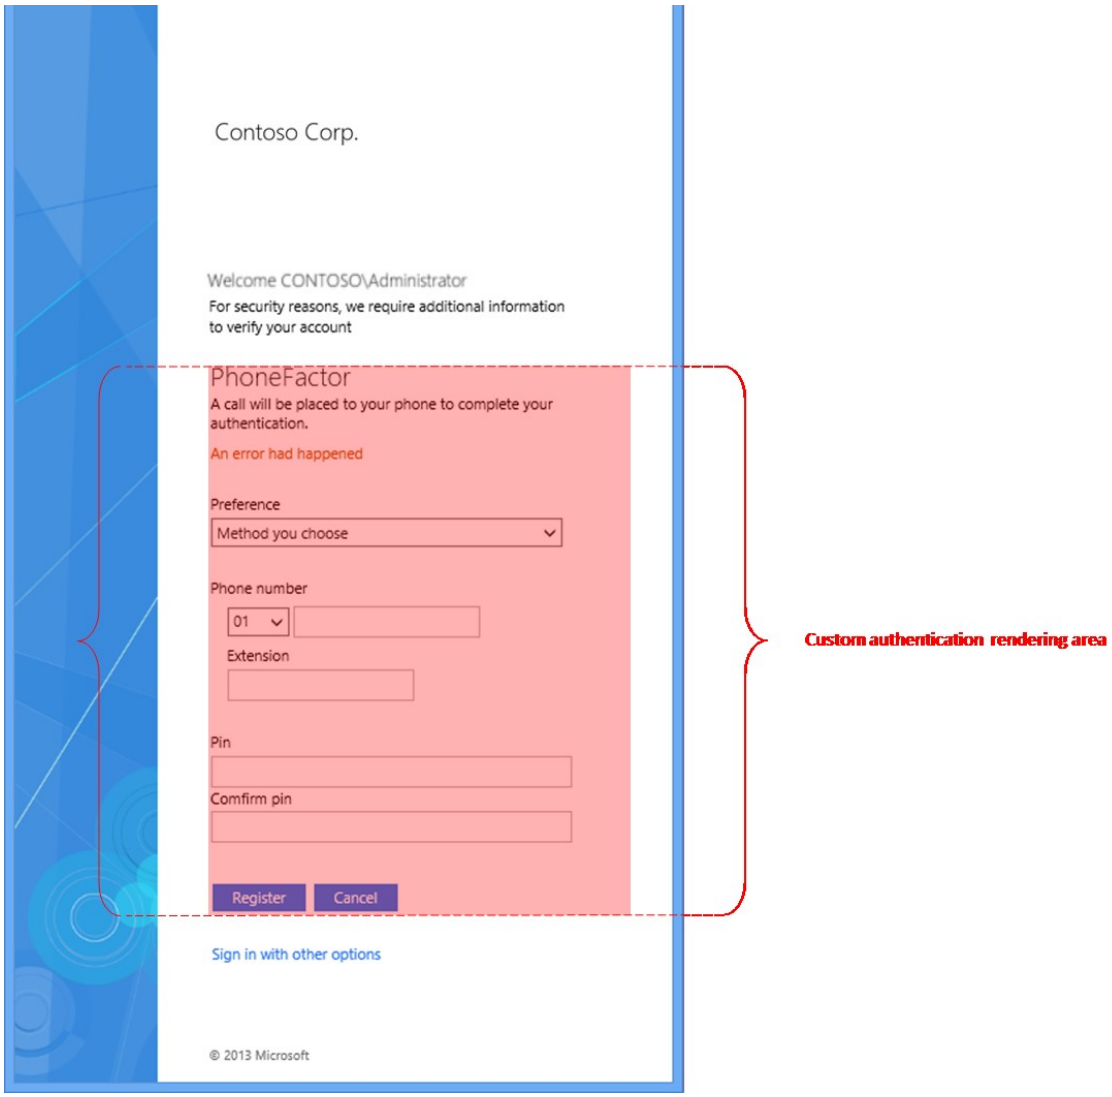

- Be aware that AD FS administrators can customize the AD FS styles. . We do not recommend to hardcode your own styles. Instead, we recommend to use AD FS styles whenever possible.
- Out-of-the box, AD FS styles are authored with one left-to-right (LTR) style and one right-to-left (RTL). Administrators can customize both, and can provide language-specific styles

through the web theme definition. Each style sheet has three sections with respective comments:

- **theme styles** - These styles should not and cannot be used. These styles are meant to define theme across all pages. They are defined by an element ID purposely so that they are not reused.
- **common styles** - These are the styles that should be used for your content.
- **form factor styles** - These are styles for different form factors. You should understand this section to ensure that your content works with different form factors, for example, phones and tablets.

For additional information, see Solution Guide: Manage Risk with Multi-factor Access Control and Solution Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

Localization

Localizing web content into languages other than English is possible. Be aware of the following considerations when you localize.

After the content is customized, the customization takes precedence; therefore, you should customize for all languages that you want to support. All customized content takes a locale parameter. When you configure localized content, configure it with a country-less locale first, for example, "en", before you configure a country and region-specific locale such as "en-us".

The following shows some additional code examples.

```
Set-AdfsWebTheme -TargetName default -Logo @{Locale="";Path="c:\contoso.png"}


Set-AdfsWebTheme -TargetName default -Illustration
@{Locale="";Path="c:\illustration.png"}
```

The following shows some additional code examples.

```
Set-AdfsGloblaWebContent -ErrorPageDescriptionText "This is Contoso's error page
description" -locale "en"


Set-AdfsGloblaWebContent -ErrorPageDescriptionText "Il s'agit de description de page
erreur de Contoso" -locale "fr"
```

If you want to customize the web content to languages other than English that requires the input of Unicode, we recommend that you use the Windows PowerShell ISE. For additional information see Introducing the Windows PowerShell ISE.

Removing the Microsoft copyright

By default, the AD FS pages contain the Microsoft copyright. To remove this copyright from your customized pages, you can use the following procedure.

To remove the Microsoft copyright

1. Create a custom theme that is based on the default.

```
New-AdfsWebTheme –Name custom –SourceName default
```

2. Export the theme by specifying the output folder.

```
Export-AdfsWebTheme -Name custom -DirectoryPath
C:\customWebTheme
```

3. Locate the Style.css file that is located in the output folder. By using the previous example, the path would be C:\CustomWebTheme\Css\Style.css.
4. Open the Style.css file with an editor, such as Notepad.
5. Locate the `#copyright` portion, and then change it to the following:

```
#copyright {color:#696969; display:none;}
```

6. Create a custom theme that is based on the new Style.css file.

```
Set-AdfsWebTheme -TargetName custom -StyleSheet
@{locale="";path="C:\customWebTheme\css\style.css"}
```

7. Activate the new theme.

```
Set-AdfsWebConfig -ActiveThemeName custom
```

See Also

AD FS Operations

Add an Attribute Store

User accounts and computer accounts that require access to a resource that is protected by Active Directory Federation Services (AD FS) are stored in an attribute store, such as Active Directory Domain Services (AD DS). The claims issuance engine uses attribute stores to gather data that is necessary to issue claims. Data from the attribute stores is then projected as claims.

You can use the following procedure to add an attribute store to the Federation Service.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add an attribute store

1. On the **Start** screen, type **AD FS Management**, and then press Enter.
2. Under the **AD FS\Trust Relationships** folder, right-click **Attribute Stores**, and then click **Add an attribute store**.
3. In the **Add an attribute store** dialog box, configure the following properties for the attribute store that you want to add:
   - In **Display name**, type the name that you want to use to identify the attribute store.
   - In **Attribute store type**, select a supported attribute store type, either **Active Directory**, **LDAP**, or **SQL**.
   - In **Connection string**, if you have selected either a Lightweight Directory Access Protocol (LDAP) store or a Structured Query Language (SQL) store, enter the string that you used to establish a connection to the attribute store. For Active Directory attribute stores, no connection string is required; therefore, this field is disabled.

     Note
     AD FS automatically creates an Active Directory attribute store, by default.
4. Click **OK**.

See Also

AD FS Operations

The Role of Attribute Stores

Create a Relying Party Trust

To add a new relying party trust by using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a federation server.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a relying party trust manually

1. In Server Manager, click **Tools**, and then select **AD FS Management**.

2. In AD FS snap-in, under **AD FS\Trust Relationships**, right-click **Relying Party Trusts**, and then click **Add Relying Party Trust** to open the Add Relying Party Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter data about the relying party manually**, and then click **Next**.
5. On the **Specify Display Name** page, type a name in **Display name**, under **Notes** type a description for this relying party trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS profile**, click **Next**, and then move to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then go to step 9.

   If you know that you will require interoperability with an older Active Directory Federation Services (AD FS) federation, as provided in Windows Server 2003 R2, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile**.
7. On the **Configure Certificate** page, click **Browse** to locate a certificate file, and then click **Next**.
8. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 10:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Relying party WS-Federation Passive protocol URL**, type the URL for this relying party trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Relying party SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this relying party trust, and then click **Next**.
9. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this relying party trust, and then click **Next**.
10. On the **Configure Identifiers** page, specify one or more identifiers for this relying party, click **Add** to add them to the list, and then click **Next**.
11. On the **Configure Multi-factor Authentication Now?** page, do one of the following:
    - Select **I do not want to configure multi-factor authentication settings for this relying party trust at this time**, and then proceed to step 13.
    - Select **Configure multi-factor authentication settings for this relying party trust**, and then proceed to step 12.

    For more information about MFA, see **Configuring authentication policies for AD FS**.
12. On the **Configure Multi-factor Authentication** page, configure multi-factor authentication (MFA) that is based on the following available settings:
    - User or group membership – In **Users/Groups**, click **Add** to specify those Active Directory (AD) users or groups for which you want to require MFA.
    - Devices – In **Devices**, select either **Unregistered** devices or **Registered devices**.
    - Locations – In **Locations**, select either **Extranet** or **Intranet**.

    For more information about MFA, see **Configuring authentication policies for AD FS**.

13. On the **Choose Issuance Authorization Rules** page, select either **Permit all users to access this relying party** or **Deny all users access to this relying party**, and then click **Next**.
14. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.
15. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box.

See Also

AD FS Operations

Create a Non-Claims-Aware Relying Party Trust

In the AD FS Management snap-in, non-claims-aware relying party trusts are objects that are created to represent the trust between the federation service and a single web-based application that is not claims-aware and that is accessed through the Web Application Proxy.

A non-claims-aware relying party trust is a relying party trust which consists of identifiers, names, and rules for authentication and authorization when the relying party trust is accessed through the Web Application Proxy. These web-based applications that do not rely on claims, in other words, these Integrated Windows Authentication-based applications, can have authorization rules that enforce access that is based on claims when the access is external to the corporate network through the Web Application Proxy.

To add a new non-claims-aware relying party trust, by using the AD FS Management snap-in, perform the following procedure.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a non-claims-aware relying party trust

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In AD FS snap-in, under **AD FS\Trust Relationships**, right-click **Relying Party Trusts**, and then click **Add Non-Claims-Aware Relying Party Trust** to open the **Add Non-Claims-Aware Relying Party Trust Wizard**.
3. On the **Welcome** page, click **Start**.
4. On the **Specify Display Name** page, type a name in **Display name**, under **Notes** type a description for this non-claims-aware relying party trust, and then click **Next**.
5. On the **Configure Identifiers** page, specify one or more identifiers for this relying party,

click **Add** to add them to the list, and then click **Next**.

6. On the **Configure Multi-factor Authentication Now?** page, do one of the following:
   - Select **I do not want to configure multi-factor authentication settings for this relying party trust at this time**, and then proceed to step 9.
   - Select **Configure multi-factor authentication settings for this relying party trust**, and then proceed to step 7.

   For more information about MFA, see **Configuring authentication policies for AD FS**.
7. On the **Configure Multi-factor Authentication** page, configure multi-factor authentication (MFA) that is based on the following available settings:
   - User or group membership – In **Users/Groups**, click **Add** to specify those Active Directory (AD) users or groups for which you want to require MFA.
   - Devices – In **Devices**, select either **Unregistered** devices or **Registered devices**.
   - Locations – In **Locations**, select either **Extranet** or **Intranet**.

   For more information about MFA, see **Configuring authentication policies for AD FS**.
8. On the **Choose Issuance Authorization Rules** page, select either **Permit all users to access this relying party** or **Deny all users access to this relying party**, and then click **Next**.
9. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.
10. On the **Finish** page, click **Close**.

Important

For the non-claims-aware relying party trusts to function there has to be a Web Application Proxy that can orchestrate the authentication and authorization and can translate to the authentication protocol of the application it represents.

See Also

AD FS Operations

Create a Claims Provider Trust

To add a new claims provider trust by using the AD FS Management snap-in and manually configure the settings, perform the following procedure on a resource partner federation server in the resource partner organization.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a claims provider trust manually

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In the AD FS snap-in, under the **AD FS\Trust Relationships**, right-click **Claims Provider Trusts**, and then click **Add Claims Provider Trust** to open the Add Claims Provider Trust Wizard.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Enter claims provider trust data manually**, and then click **Next**.
5. On the **Specify Display Name** page, type a **Display name**, under **Notes**, type a description for this claims provider trust, and then click **Next**.
6. On the **Choose Profile** page, do one of the following:
   - Click **AD FS Profile**, click **Next**, and then go to step 7.
   - Click **AD FS 1.0 and 1.1 profile**, click **Next**, and then go to step 8.

   If you know that you require interoperability between this claims provider trust and other, older Active Directory Federation Services (AD FS) claims provider trusts, click **AD FS 1.0 and 1.1 profile**. Otherwise, use the default **AD FS profile** option.
7. On the **Configure URL** page, do one or both of the following, click **Next**, and then go to step 9:
   - Select the **Enable support for the WS-Federation Passive protocol** check box. Under **Claims provider WS-Federation Passive protocol URL**, type the URL for this claims provider trust, and then click **Next**.
   - Select the **Enable support for the SAML 2.0 WebSSO protocol** check box. Under **Claims provider SAML 2.0 SSO service URL**, type the Security Assertion Markup Language (SAML) service endpoint URL for this claims provider trust, and then click **Next**.
8. On the **Configure URL** page, under **WS-Federation Passive URL**, type the URL for this claims provider trust, and then click **Next**.
9. On the **Configure Identifier** page, under **Claims provider trust identifier**, type the appropriate identifier, and then click **Next**.
10. On the **Configure Certificates** page, click **Add** to locate a certificate file and add it to the list of certificates, and then click **Next**.
11. On the **Ready to Add Trust** page, click **Next** to save your claims provider trust information.
12. On the **Finish** page, click **Close**. This action automatically displays the **Edit Claim Rules** dialog box. For more information about how to proceed with adding claim rules for this claims provider trust, see the following additional references.

Additional references

Checklist: Configuring the Resource Partner Organization

Checklist: Creating Claim Rules for a Claims Provider Trust

See Also

AD FS Operations

Configuring Claim Rules

In a claims-based identity model, the function of Active Directory Federation Services (AD FS) as federation services is to issue a token that contains a set of claims. Claims rules govern the decision in regard of claims that AD FS issues. Claim rules and all server configuration data are stored in the AD FS configuration database.

AD FS makes issuance decisions that are based on identity information that is provided to it in the form of claims and other contextual information. At a high level, AD FS operates as a rules processor by taking one set of claims as input, performs a number of transformations, and then returns a different set of claims as output.

- Create a Rule to Pass Through or Filter an Incoming Claim
- Create a Rule to Permit All Users
- Create a Rule to Permit or Deny Users Based on an Incoming Claim
- Create a Rule to Send LDAP Attributes as Claims
- Create a Rule to Send Group Membership as a Claim
- Create a Rule to Transform an Incoming Claim
- Create a Rule to Send an Authentication Method Claim
- Create a Rule to Send Claims Using a Custom Rule

See Also

AD FS Operations

Create a Rule to Pass Through or Filter an Incoming Claim

By using the **Pass Through or Filter an Incoming Claim** rule template in Active Directory Federation Services (AD FS), you can pass through all incoming claims by using a selected claim type. You can also filter the values of incoming claims by using a selected claim type. For example, you can use this rule template to create a rule that sends all incoming group claims. You can also use this rule to send only user principal name (UPN) claims that end with @fabrikam.

You can use the following procedure to create a claim rule by using the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to pass through or filter an incoming claim

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. In **Incoming claim type**, select a claim type in the list, and then select one of the following options, which depends on the requirements of your organization:
   - **Pass through all claim values**
   - **Pass through only a specific claim value**
   - **Pass through only claim values that match a specific email suffix value**
   - **Pass through only claim values that start with a specific value**
7. Click the **Finish** button.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Permit All Users

By using the **Permit All Users** rule template in Active Directory Federation Services (AD FS), you can create an authorization rule that gives all users access to the relying party. You can use

additional authorization rules to further restrict access. Users who are permitted to access the relying party from the Federation Service might still be denied service by the relying party.

You can use the following procedure to create a claim rule by using the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to permit all users

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS snap-in, in the console tree, under **AD FS\Trust Relationships\Relying Party Trusts**, click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, click the **Issuance Authorization Rules** tab or the **Delegation Authorization Rules** tab, which is based on the type of authorization rule that you require, and then click **Add Rule** to start the **Add Authorization Claim Rule Wizard**.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Permit All Users** from the list, and then click **Next**.
6. On the **Configure Rule** page, click **Finish**.
7. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Permit or Deny Users Based on an Incoming Claim

By using the **Permit or Deny Users Based on an Incoming Claim** rule template in Active Directory Federation Services (AD FS), you can create an authorization rule that grants or denies user's access permission to the relying party, which is based on the type and value of an incoming claim. For example, you can use this rule template to create a rule that grants only users that have a group claim with a value of Domain Admins access the relying party. If you want to grant all users access the relying party, use the **Permit All Users** rule template. Users who are granted access the relying party from the Federation Service might still be denied service by the relying party.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to permit or deny users based on an incoming claim

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS snap-in, in the console tree, under **AD FS\Trust Relationships\Relying Party Trusts**, click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, click the **Issuance Authorization Rules** tab or the **Delegation Authorization Rules** tab, which is based on the type of authorization rule that you require, and then click **Add Rule** to start the **Add Authorization Claim Rule Wizard**.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. In **Incoming claim type**, select a claim type in the list, under **Incoming claim value**, type a value or click **Browse**, if it is available, and select a value, and then select one of the following options, which depends on the requirements of your organization:
   - **Permit access to users with this incoming claim**
   - **Deny access to users with this incoming claim**
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Send LDAP Attributes as Claims

By using the **Send LDAP Attributes as Claims** rule template in Active Directory Federation Services (AD FS), you can create a rule that selects attributes from a Lightweight Directory Access Protocol (LDAP) attribute store, such as Active Directory, to send as claims to the relying party. For example, you can use this rule template to create a **Send LDAP Attributes as Claims** rule that extracts attribute values for authenticated users from the **displayName** and **telephoneNumber** Active Directory attributes, and then send those values as two different outgoing claims.

You can also use this rule to send all the user's group memberships. If you want to send only individual group memberships, use the **Send Group Membership as a Claim** rule template. You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send LDAP attributes as claims

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. Under **Attribute store**, select **Active Directory**, and under **Mapping of LDAP attributes to outgoing claim types**, select the **LDAP Attribute** that you want and the corresponding **Outgoing Claim Type** types from the lists.

   You have to select a new LDAP attribute and outgoing claim type pair on a different row for each Active Directory attribute for which you want to issue a claim as part of this rule.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Send Group Membership as a Claim

By using the **Send Group Membership as a Claim** rule template in Active Directory Federation Services (AD FS), you can create a rule that enables you to select an Active Directory security group to send as a claim. Only a single claim is emitted from this rule, which is based on the group that you select. For example, you can use this rule template to create a rule that sends a group claim with a value of Admin if the user is a member of the Domain Admins security group. This rule should be used only for users in the local Active Directory domain.

You can use the following procedure to create a claim rule by using the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send group membership as a claim

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. In **User's group**, click **Browse**, and then select a group. Under **Outgoing claim type**, select the claim type that you want to use, and then under **Outgoing Claim Type**, type a value.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Transform an Incoming Claim

By using the **Transform an Incoming Claim** rule template in Active Directory Federation Services (AD FS), you can select an incoming claim, change its claim type, and change its claim value. For example, you can use this rule template to create a rule that sends a role claim with the same claim value of an incoming group claim. You can also use this rule to send a group claim with a claim value of Purchasers when there is an incoming group claim with a value of Admins, or you can send only user principal name (UPN) claims that end with @fabrikam.

You can use the following procedure to create a claim rule with the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to transform an incoming claim

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. In **Incoming claim type**, select a claim type in the list. In **Outgoing claim type**, select a claim type in the list, and then select one of the following options, which depends on the requirements of your organization:
   - **Pass through all claim values**
   - **Replace an incoming claim value with a different outgoing claim value**

- **Replace incoming e-mail suffix claims with a new e-mail suffix**

Notes

If you are setting up the Dynamic Access Control scenario that uses AD FS-issued claims, first create a transform rule on the claims provider trust, and in **Incoming claim type**, type the name for the incoming claim, or, if a claim description was previously created, select it from the list. Second, in **Outgoing claim type**, select the claim URL that you want, and then create a transform rule on the relying party trust to issue the device claim.

For more information about Dynamic Access Control scenarios, see Dynamic Access Control Content Roadmap or Using AD DS Claims with AD FS.

7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Send an Authentication Method Claim

You can use either the **Send Group Membership as Claims** rule template or the **Transform an Incoming Claim** rule template to send an authentication method claim. The relying party can use an authentication method claim to determine the logon mechanism that the user uses to authenticate and obtain claims from Active Directory Federation Services (AD FS). You can also use the Authentication Mechanism Assurance feature of Active Directory Federation Services (AD FS) in Windows Server 2012 R2 as input to generate authentication method claims for situations in which the relying party wants to determine the level of access that is based on smart card logons. For example, a developer can assign different levels of access to federated users of the relying party application. The levels of access are based on whether the users log on with their user name and password credentials, as opposed to their smart cards.

Depending on the requirements of your organization, use one of the following procedures:

- Create this rule by using the **Send Group Membership as Claims** rule template - You can use this rule template when you want the group that you specify in this template to ultimately determine what authentication method claim to issue.
- Create this rule by using the **Transform an Incoming Claim** rule template - You can use this rule template when you want to change the existing authentication method to a new

authentication method that works with a product that does not recognize standard AD FS authentication method claims.

To create this rule by using the Send Group Membership as Claims rule template

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one of the following tabs, depending on the trust that you are editing, and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. Click **Browse**, select the group whose members should receive this authentication method claim, and then click **OK**.
8. In **Outgoing claim type**, select **Authentication method** in the list.
9. In **Outgoing claim value**, type one of the default uniform resource identifier (URI) values in the following table, depending on your preferred authentication method, click **Finish**, and then click **OK** to save the rule.

Actual Authentication method	Corresponding URI
User name and password authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/password
Windows authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/windows
Transport Layer Security	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/tlsclient

(TLS) Mutual authentication that uses X.509 certificates	
X.509-based authentication that does not use TLS	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/x509

Note
Other URI values can be used in addition to the values in the table. The URI values that are shown in the previous table reflect the URIs that the relying party accepts by default.

To create this rule by using the Transform an Incoming Claim rule template

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on which trust you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select **Authentication method** in the list.
8. In **Outgoing claim type**, select **Authentication method** in the list.
9. Select **Replace an incoming claim value with a different outgoing claim value**, and then do the following:
   a. In **Incoming claim value**, type one of the following URI values that are based on the actual authentication method URI that was used originally, click **Finish**, and then click **OK** to save the rule.

b. In **Outgoing claim value**, type one of the default URI values in the following table, which depends on your new preferred authentication method choice, click **Finish**, and then click **OK** to save the rule.

Actual authentication method	Corresponding URI
User name and password authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/password
Windows authentication	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/windows
TLS mutual authentication that uses X.509 certificates	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/tlsclient
X.509-based authentication that does not use TLS	http://schemas.microsoft.com/ws/2008/06/identity/authenticationmethod/x509

Note
Other URI values can be used in addition to the values in the table. The URI values that are shown ion the previous table reflect the URIs that the relying party accepts by default.

See Also

Configuring Claim Rules

Create a Rule to Send an AD FS 1.x Compatible Claim

When you use Active Directory Federation Services (AD FS) to issue claims that are to be received by federation servers that run AD FS 1.0 on the Windows Server 2003 R2 operating system or AD FS 1.1 on the Windows Server 2008 or Windows Server 2008 R2 operating system, you must do the following:

- Create a rule that sends a Name ID claim type with a format of user prinicipal name (UPN), Email, or Common Name.
- All other claims that are sent must have one of the following claim types:
  - AD FS 1.*x* Email Address
  - AD FS 1.*x* UPN
  - Common Name
  - Group
  - Any other claim type that begins with https://schemas.xmlsoap.org/claims/, such as https://schemas.xmlsoap.org/claims/EmployeeID.

Depending on the requirements of your organization, use one of the following procedures to create an AD FS 1.*x* compatible Name ID claim:

- Create this rule to issue an AD FS 1.x Name ID claim by using the **Pass Through or Filter an Incoming Claim** rule template.
- Create this rule to issue an AD FS 1.x Name ID claim by using the **Transform an Incoming Claim** rule template. You can use this rule template when you want to change the existing claim type to a new claim type that works with AD FS 1. *x* claims.

Note
For this rule to work as expected, ensure that the relying party trust or claims provider trust where you create this rule has been configured to use the **AD FS 1.0 and 1.1 profile**.

Create a rule to issue an AD FS 1.

You can use this rule template when you want the group that you specify in this template to ultimately determine what AD FS 1.*x*-compatible claim to issue.

To create a rule to issue an AD FS 1.x Name ID claim by using the Pass Through or Filter an Incoming Claim rule template

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.

4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select **Name ID** in the list.
8. In **Incoming name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
9. Select one of the following options, which depends on the requirements of your organization:
   - **Pass through all claim values**
   - **Pass through only a specific claim value**
   - **Pass through only claim values that match a specific email suffix value**
   - **Pass through only claim values that start with a specific value**
10. Click **Finish**, and then click **OK** to save the rule.

Create a rule to issue an AD FS 1.

You can use this rule template when you want to change the incoming claim type to a new outgoing claim type of Name ID that works with AD FS 1.*x* federation servers.

To create a rule to issue an AD FS 1.x Name ID claim by using the Transform an Incoming Claim rule template

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**

- **Issuance Transform Rules**
- **Issuance Authorization Rules**
- **Delegation Authorization Rules**

5. On the **Select Rule Template** page, under **Claim rule template**, select **Transform an Incoming Claim** in the list, and then click **Next**.
6. On the **Configure Rule** page, type a claim rule name.
7. In **Incoming claim type**, select the type of incoming claim that you want to transform in the list.
8. In **Outgoing claim type**, select **Name ID** in the list.
9. In **Outgoing name ID format**, select one of the following AD FS 1.*x*-compatible claim formats from the list:
   - **UPN**
   - **E-Mail**
   - **Common Name**
10. Select one of the following options, which depends on the requirements of your organization:
    - **Pass through all claim values**
    - **Replace an incoming claim value with a different outgoing claim value**
    - **Replace incoming e-mail suffix claims with a new e-mail suffix**
11. Click **Finish**, and then click **OK** to save the rule.

See Also

Configuring Claim Rules

Create a Rule to Send Claims Using a Custom Rule

By using the **Send Claims Using a Custom Rule** template in Active Directory Federation Services (AD FS), you can create custom claim rules for situation in which a standard rule template does not satisfy the requirements of your organization. Custom claim rules are written in the claim rule language and must then be copied into the **Custom rule** text box before they can be used in a rule set. For information about constructing the syntax for an advanced rule, see The Role of the Claim Rule Language.

You can use the following procedure to create a claim rule by using the AD FS Management snap-in.

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete this procedure. Review details about using the appropriate accounts and group

memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To create a rule to send claims by using a custom claim

1. In Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click either **Claims Provider Trusts** or **Relying Party Trusts**, and then click a specific trust in the list where you want to create this rule.
3. Right-click the selected trust, and then click **Edit Claim Rules**.
4. In the **Edit Claim Rules** dialog box, select one the following tabs, which depends on the trust that you are editing and in which rule set you want to create this rule, and then click **Add Rule** to start the rule wizard that is associated with that rule set:
   - **Acceptance Transform Rules**
   - **Issuance Transform Rules**
   - **Issuance Authorization Rules**
   - **Delegation Authorization Rules**
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule** from the list, and then click **Next**.
6. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule. Under **Custom rule**, type or paste the claim rule language syntax that you want for this rule.
7. Click **Finish**.
8. In the **Edit Claim Rules** dialog box, click **OK** to save the rule.

See Also

Configuring Claim Rules

Configuring Authentication Policies

In AD FS in Windows Server 2012 R2, both access control and the authentication mechanism are enhanced with multiple factors that include user, device, location, and authentication data. These enhancements enable you, either through the user interface or through Windows PowerShell, to manage the risk of granting access permissions to AD FS-secured applications via multi-factor access control and multi-factor authentication that are based on user identity or group membership, network location, device data that is workplace-joined, and the authentication state when multi-factor authentication (MFA) was performed.

For more information about MFA and multi-factor access control in Active Directory Federation Services (AD FS) in Windows Server 2012 R2, see the following topics:

- Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications
- Overview: Connect to Applications and Services from Anywhere with Web Application Proxy
- Overview: Manage Risk with Multi-Factor Access Control
- Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications

Configure authentication policies via the AD FS Management snap-in

Membership in **Administrators**, or equivalent, on the local computer is the minimum requirement to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In AD FS in Windows Server 2012 R2, you can specify an authentication policy at a global scope that is applicable to all applications and services that are secured by AD FS. You can also set authentication policies for specific applications and services that rely on party trusts and are secured by AD FS. Specifying an authentication policy for a particular application per relying party trust does not override the global authentication policy. If either global or per relying party trust authentication policy requires MFA, MFA is triggered when the user tries to authenticate to this relying party trust. The global authentication policy is a fallback for relying party trusts for applications and services that do not have a specific configured authentication policy.

To configure primary authentication globally

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In AD FS snap-in, click **Authentication Policies**.
3. In the **Primary Authentication** section, click **Edit** next to **Global Settings**. You can also right-click **Authentication Policies**, and select **Edit Global Primary Authentication**, or, under the **Actions** pane, select **Edit Global Primary Authentication**.
4. In the **Edit Global Authentication Policy** window, on the **Primary** tab, you can configure the following settings as part of the global authentication policy:
   - Authentication methods to be used for primary authentication. You can select available authentication methods under the **Extranet** and **Intranet**.
   - Device authentication via the **Enable device authentication** check box. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.

To configure primary authentication per relying party trust

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In AD FS snap-in, click **Authentication Policies\Per Relying Party Trust**, and then click the relying party trust for which you want to configure authentication policies.
3. Either right-click the relying party trust for which you want to configure authentication

policies, and then select **Edit Custom Primary Authentication**, or, under the **Actions** pane, select **Edit Custom Primary Authentication**.

4. In the **Edit Authentication Policy for <relying_party_trust_name>** window, under the **Primary** tab, you can configure the following setting as part of the **Per Relying Party Trust** authentication policy:
   - Whether users are required to provide their credentials each time at sign-in via the **Users are required to provide their credentials each time at sign-in** check box.

To configure multi-factor authentication globally

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In AD FS snap-in, click **Authentication Policies**.
3. In the **Multi-factor Authentication** section, click **Edit** next to **Global Settings**. You can also right-click **Authentication Policies**, and select **Edit Global Multi-factor Authentication**, or, under the **Actions** pane, select **Edit Global Multi-factor Authentication**.
4. In the **Edit Global Authentication Policy** window, under the **Multi-factor** tab, you can configure the following settings as part of the global multi-factor authentication policy:
   - Settings or conditions for MFA via available options under the **Users/Groups**, **Devices**, and **Locations** sections.
   - To enable MFA for any of these settings, you must select at least one additional authentication method. **Certificate Authentication** is the default available option. You can also configure other custom additional authentication methods, for example, Windows Azure Active Authentication. For more information, see Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

     Warning
     You can only configure additional authentication methods at a globally.

To configure multi-factor authentication per relying party trust

1. In Server Manager, click **Tools**, and then select **AD FS Management**.
2. In AD FS snap-in, click **Authentication Policies\Per Relying Party Trust**, and then click the relying party trust for which you want to configure MFA.
3. Either right-click the relying party trust for which you want to configure MFA, and then select **Edit Custom Multi-factor Authentication**, or, under the **Actions** pane, select **Edit Custom Multi-factor Authentication**.
4. In the **Edit Authentication Policy for <relying_party_trust_name>** window, under the **Multi-factor** tab, you can configure the following settings as part of the per-relying party trust authentication policy:
   - Settings or conditions for MFA via available options under the **Users/Groups**, **Devices**, and **Locations** sections.

Configure authentication policies via Windows PowerShell

Windows PowerShell enables greater flexibility in using various factors of access control and the authentication mechanism that are available in AD FS in Windows Server 2012 R2 to configure authentication policies and authorization rules that are necessary to implement true conditional access for your AD FS -secured resources.

Membership in Administrators, or equivalent, on the local computer is the minimum requirement to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure an additional authentication method via Windows PowerShell

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
Set-AdfsGlobalAuthenticationPolicy –
AdditionalAuthenticationProvider CertificateAuthentication
```

Warning
To verify that this command ran successfully, you can run the `Get-AdfsGlobalAuthenticationPolicy` command.

To configure MFA per-relying party trust that is based on a user's group membership data

1. On your federation server, open the Windows PowerShell command window and run the following command:

```
$rp = Get-AdfsRelyingPartyTrust –Name relying_party_trust
```

Warning
Ensure to replace *<relying_party_trust>* with the name of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$MfaClaimRule = "c:[Type ==
'"http://schemas.microsoft.com/ws/2008/06/identity/claims/gro
upsid'", Value =~ '"^(?i) <group_SID>$'"] => issue(Type =
'"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod'", Value =
'"http://schemas.microsoft.com/claims/multipleauthn'");"

Set-AdfsRelyingPartyTrust –TargetRelyingParty $rp –
```

```
AdditionalAuthenticationRules $MfaClaimRule
```

Note
Ensure to replace <group_SID> with the value of the security identifier (SID) of your Active Directory (AD) group.

To configure MFA globally based on users' group membership data

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$MfaClaimRule = "c:[Type == '"
http://schemas.microsoft.com/ws/2008/06/identity/claims/group
sid'", Value == '"group_SID'"]
 => issue(Type =
'"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod'", Value =
'"http://schemas.microsoft.com/claims/multipleauthn'");"

Set-AdfsAdditionalAuthenticationRule $MfaClaimRule
```

Note
Ensure to replace *<group_SID>* with the value of the SID of your AD group.

To configure MFA globally based on user's location

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$MfaClaimRule = "c:[Type == '"
http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwor
k'", Value == '"true_or_false'"]
 => issue(Type =
'"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod'", Value =
'"http://schemas.microsoft.com/claims/multipleauthn'");"

Set-AdfsAdditionalAuthenticationRule $MfaClaimRule
```

Note

Ensure to replace *<true_or_false>* with either `true` or `false`. The value depends on your specific rule condition that is based on whether the access request comes from the extranet or the intranet.

To configure MFA globally based on user's device data

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$MfaClaimRule = "c:[Type == '"
http://schemas.microsoft.com/2012/01/devicecontext/claims/isr
egistereduser'", Value == '"true_or_false"']
 => issue(Type =
'"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod'", Value =
'"http://schemas.microsoft.com/claims/multipleauthn'");"

Set-AdfsAdditionalAuthenticationRule $MfaClaimRule
```

Note
Ensure to replace *<true_or_false>* with either `true` or `false`. The value depends on your specific rule condition that is based on whether the device is workplace-joined or not.

To configure MFA globally if the access request comes from the extranet and from a non-workplace-joined device

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
Set-AdfsAdditionalAuthenticationRule "c:[Type ==
`"http://schemas.microsoft.com/2012/01/devicecontext/claims/i
sregistereduser`", Value == `"true_or_false`"] && c2:[Type ==
`"http://schemas.microsoft.com/ws/2012/01/insidecorporatenetw
ork`", Value == `" true_or_false `"] => issue(Type =
`"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod`", Value
=`"http://schemas.microsoft.com/claims/multipleauthn`");"
```

Note
Ensure to replace both instances of *<true_or_false>* with either `true` or `false`, which depends on your specific rule conditions. The rule conditions are based on whether the device is workplace-joined or not and whether the access request comes from the extranet or intranet.

To configure MFA globally if access comes from an extranet user that belongs to a certain group

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
Set-AdfsAdditionalAuthenticationRule "c:[Type ==
`"http://schemas.microsoft.com/ws/2008/06/identity/claims/gro
upsid`", Value == `"group_SID`"]
 && c2:[Type ==
`"http://schemas.microsoft.com/ws/2012/01/insidecorporatenetw
ork`", Value == `"true_or_false`"] => issue(Type =
`"http://schemas.microsoft.com/ws/2008/06/identity/claims/aut
henticationmethod`", Value
=`"http://schemas.microsoft.com/claims/multipleauthn`");"
```

Note
Ensure to replace *<group_SID>* with the value of the group SID and *<true_or_false>* with either `true` or `false`, which depends on your specific rule condition that is based on whether the access request comes from the extranet or intranet.

To grant access to an application based on user data via Windows PowerShell

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$rp = Get-AdfsRelyingPartyTrust -Name relying_party_trust
```

Note
Ensure to replace *<relying_party_trust>* with the value of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$GroupAuthzRule = "@RuleTemplate = `"Authorization`"
@RuleName = `"Foo`" c:[Type ==
`"http://schemas.microsoft.com/ws/2008/06/identity/claims/gro
upsid`", Value =~ `"^(?i)<group_SID>$`"] =>issue(Type =
`"http://schemas.microsoft.com/authorization/claims/deny`",
Value = `"DenyUsersWithClaim`");"
Set-AdfsRelyingPartyTrust -TargetRelyingParty $rp -
IssuanceAuthorizationRules $GroupAuthzRule
```

Note

Ensure to replace *<group_SID>* with the value of the SID of your AD group.

To grant access to an application that is secured by AD FS only if this user's identity was validated with MFA

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$rp = Get-AdfsRelyingPartyTrust -Name relying_party_trust
```

Note

Ensure to replace *<relying_party_trust>* with the value of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$GroupAuthzRule = "@RuleTemplate = `"Authorization`"
@RuleName = `"PermitAccessWithMFA`"
c:[Type ==
`"http://schemas.microsoft.com/claims/authnmethodsreferences`
", Value =~
`"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$`"
] => issue(Type =
`"http://schemas.microsoft.com/authorization/claims/permit`",
Value = '`"PermitUsersWithClaim'");"
```

To grant access to an application that is secured by AD FS only if the access request comes from a workplace-joined device that is registered to the user

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$rp = Get-AdfsRelyingPartyTrust -Name relying_party_trust
```

Note

Ensure to replace *<relying_party_trust>* with the value of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$GroupAuthzRule = "@RuleTemplate = `"Authorization`"
@RuleName =
`"PermitAccessFromRegisteredWorkplaceJoinedDevice`"
c:[Type ==
`"http://schemas.microsoft.com/2012/01/devicecontext/claims/i
sregistereduser`", Value =~ `"^(?i)true$`"] => issue(Type =
```

```
`"http://schemas.microsoft.com/authorization/claims/permit`",
Value = `"PermitUsersWithClaim`");
```

To grant access to an application that is secured by AD FS only if the access request comes from a workplace-joined device that is registered to a user whose identity has been validated with MFA

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$rp = Get-AdfsRelyingPartyTrust -Name relying_party_trust
```

Note

Ensure to replace *<relying_party_trust>* with the value of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$GroupAuthzRule = '@RuleTemplate = "Authorization"
@RuleName = "RequireMFAOnRegisteredWorkplaceJoinedDevice"
c1:[Type ==
`"http://schemas.microsoft.com/claims/authnmethodsreferences`
", Value =~
`"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$`"
] &&
c2:[Type ==
`"http://schemas.microsoft.com/2012/01/devicecontext/claims/i
sregistereduser`", Value =~ `"^(?i)true$"] => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit`",
Value = `"PermitUsersWithClaim`");"
```

To grant extranet access to an application secured by AD FS only if the access request comes from a user whose identity has been validated with MFA

1. On your federation server, open the Windows PowerShell command window and run the following command.

```
$rp = Get-AdfsRelyingPartyTrust -Name relying_party_trust
```

Note

Ensure to replace *<relying_party_trust>* with the value of your relying party trust.

2. In the same Windows PowerShell command window, run the following command.

```
$GroupAuthzRule = "@RuleTemplate = `"Authorization`"
@RuleName = `"RequireMFAForExtranetAccess`"
c1:[Type ==
`"http://schemas.microsoft.com/claims/authnmethodsreferences`
", Value =~
`"^(?i)http://schemas\.microsoft\.com/claims/multipleauthn$`"
] &&
c2:[Type ==
`"http://schemas.microsoft.com/ws/2012/01/insidecorporatenetw
ork`", Value =~ `"^(?i)false$`"] => issue(Type =
`"http://schemas.microsoft.com/authorization/claims/permit`",
Value = `"PermitUsersWithClaim`");"
```

See Also

AD FS Operations

Configuring AD FS Extranet Lockout

In the context of AD FS in Windows Server 2012 R2, Web Application Proxy functions as a federation server proxy. Web Application Proxy also serves as a barrier between the Internet and your corporate applications.

Web Application Proxy provides a number of security features to protect your corporate network, such as your users and your resources, from external threats. One of these features is AD FS extranet lockout. In case of an attack in the form of authentication requests with invalid (bad) passwords that come through the Web Application Proxy, AD FS extranet lockout enables you to protect your users from an AD FS account lockout. In addition to protecting your users from an AD FS account lockout, AD FS extranet lockout also protects against brute force password guessing attacks.

With the AD FS extranet lockout feature in Windows Server 2012 R2, an AD FS administrator can set a maximum allowed number of failed authentication requests (`ExtranetLockoutThreshold`) and an ‘observation window's time period (`ExtranetObservationWindow`). When this maximum number (`ExtranetLockoutThreshold`) of authentication requests is reached, AD FS stops trying to authenticate the supplied account credentials against AD FS for the set time period (`ExtranetObservationWindow`). This action protects this account from an AD account lockout, in other words, it protects this account from losing access to corporate resources that rely on AD FS for authentication of the user. These settings apply to all domains that the AD FS service can authenticate.

AD FS extranet lockout functions independently from the AD lockout policies. However, it can be used together with AD lockout policy. We strongly recommend that you set the `ExtranetLockoutThreshold` parameter value to a value that is less than the AD account lockout threshold.

To configure the AD FS extranet lockout, you must set three properties on the AD FS service object, in Windows PowerShell (`get-adfsproperties`/`set-adfsproperties`):

- `ExtranetLockoutEnabled`: Boolean, to enable/disable.
- `ExtranetLockoutThreshold`: Integer, that defines the maximum number of bad password attempts. The recommended value for ExtranetLockoutThreshold is 15. The AD FS lockout threshold is 20. For more information, see Strengthening Domain Policy Settings. When the AD FS lockout policies are in effect, we strongly recommend to set the value of `ExtranetLockoutThreshold` to a value below the AD FS lockout threshold value.
- `ExtranetObservationWindow`: *TimeSpan*, defines the time period where AD FS does not attempt to contact the domain controller to validate the U/P and immediately rejects the request from outside. The recommended value for `ExtranetObservationWindow` is *30 minutes*.

You can use the following Windows PowerShell command to set the AD FS extranet lockout: `Set-AdfsProperties -EnableExtranetLockout $true -ExtranetLockoutThreshold 15 -ExtranetObservationWindow ( new-timespan -Minutes 30 ).`

See Also

AD FS Operations

Configuring Client Access Policies

Some organizations may want to create policies that limit access to Microsoft ® Office 365 services, depending on where the client resides. For example, you might want to:

- Block all extranet client access to Office 365
- Block all extranet client access to Office 365, except for devices accessing Exchange Online for Exchange Active Sync
- Block all extranet client access to Office 365, except for users in one or more Active Directory group
- Block all extranet client access to Office 365, except for browser based applications

Active Directory Federation Services (AD FS) in Windows Server 2012 R2 provides a way for organizations to configure these types of policies. Office 365 customers using Single Sign-On (SSO) who require these policies can now use client access policy rules to restrict access based on the location of the computer or device that is making the request. Customers using Microsoft Online Services cloud User IDs cannot implement these restrictions at this time.

For more information about how to deploy AD FS in Windows Server 2012 R2 for SSO with Microsoft Online services, see http://go.microsoft.com/fwlink/?LInkId=321138.

Client Access Policy Claim Types

Client access policy works by identifying which authentication requests should be permitted based upon attributes of the request itself. To provide this additional request context information, AD FS populates claim values from the client request information such as the connection IP address, the AD FS endpoint, and HTTP headers sent by the client. For a detailed description of the new claim types and values, see Claim Types.

Client Access Policy Scenarios

The following table describes the scenarios supported by the client access policy feature.

Scenario	Description
Block all external access to Office 365	Office 365 access is allowed from all clients on the internal corporate network, but requests from external clients are denied based on the IP address of the external client.
Block all external access to Office 365, except Exchange ActiveSync	Office 365 access is allowed from all clients on the internal corporate network, as well as from any external client devices, such as smart phones, that make use of Exchange ActiveSync. All other external clients, such as those using Outlook, are blocked.
Block all external access to Office 365, except for browser-based applications such as Outlook Web Access or SharePoint Online	Blocks external access to Office 365, except for passive (browser-based) applications such as Outlook Web Access or SharePoint Online.
Block all external access to Office 365 for members of designated Active Directory groups	This scenario is used for testing and validating client access policy deployment. It blocks external access to Office 365 only for members of one or more Active Directory group. It can also be used to provide external access only to members of a group.

Note

All client access policy scenarios have the effect of blocking external access to Microsoft Lync Online and Office Subscription Services.

Enabling Client Access Policy

To enable client access policy in AD FS in Windows Server 2012 R2, you must update the Microsoft Office 365 Identity Platform relying party trust. Choose one of the example scenarios

below to configure the claim rules on the **Microsoft Office 365 Identity Platform** relying party trust that best meets the needs of your organization.

Scenario 1: Block all external access to Office 365

This client access policy scenario allows access from all internal clients and blocks all external clients based on the IP address of the external client. You can use the following procedures to add the correct Issuance Authorization rules to the Office 365 relying party trust for your chosen scenario.

To create rules to block all external access to Office 365

1. From **Server Manager**, click **Tools**, then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click **Relying Party Trusts**, right-click the **Microsoft Office 365 Identity Platform** trust, and then click **Edit Claim Rules**.
3. In the **Edit Claim Rules** dialog box, select the **Issuance Authorization Rules** tab, and then click **Add Rule** to start the Claim Rule Wizard.

   Note
   If you have never configured any custom Issuance Authorization Rules on the Microsoft Office 365 Identity Platform relying party trust, you will see the default **Permit Access to All Users** rule. The rules specified below are intended to replace this rule by issuing an "http://schemas.microsoft.com/authorization/claims/permit" claim based on a more specific set of conditions. Once you have finished adding the rules specified below for your scenario, be sure to come back and delete the **Permit Access to All Users** rule. Otherwise, your client access policy rules will not work.
4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
5. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example "Issue custom claim for internal client requests directly to federation server". Under **Custom rule**, type or paste the following claim rule language syntax:

   ```
   c:[Type == "http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork", Value == "true"] => issue(Type = "http://custom/allow", Value = "true");
   ```
6. Click **Finish**. Verify that the new rule appears in the Issuance Authorization Rules list.
7. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
8. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.

   On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example "Issue custom claim based on IP regex". Under **Custom rule**, type or paste the following claim rule language syntax:`c:[Type == "http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-forwarded-client-`

```
ip", Value =~ "customer provided public ip address regex"] => issue(Type =
"http://custom/allow", Value = "true");
```

You will have to replace the value above for “public ip address regex” with a valid IP expression; see Building the IP address range expression for more information.

9. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
10. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
11. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.

    On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue permit claim based on the custom claim”. Under Custom rule, type or paste the following claim rule language syntax:`c:[Type ==`

    ```
    "http://custom/allow", Value == "true"] => issue(Type =
    "http://schemas.microsoft.com/authorization/claims/permit", Value = "true");
    ```
12. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
13. If you still have the default **Permit Access to All Users** rule in the Issuance Authentication Rules list, select it and click **Remove Rule**. Click **Yes** to confirm.
14. To save the new rules, in the **Edit Claim Rules** dialog box, click **OK**.

Scenario 2: Block all external access to Office 365 except Exchange ActiveSync

The following example allows access to all Office 365 applications, including Exchange Online, from internal clients including Outlook. It blocks access from clients residing outside the corporate network, as indicated by the client IP address, except for Exchange ActiveSync clients such as smart phones.

To create rules to block all external access to Office 365 except Exchange ActiveSync

1. From **Server Manager**, click **Tools**, then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click **Relying Party Trusts**, right-click the **Microsoft Office 365 Identity Platform trust**, and then click **Edit Claim Rules**.
3. In the **Edit Claim Rules** dialog box, select the **Issuance Authorization Rules** tab, and then click **Add Rule** to start the Claim Rule Wizard.

   Note
   If you have never configured any custom Issuance Authorization Rules on the Microsoft Office 365 Identity Platform relying party trust, you will see the default “Permit Access to All Users” rule. The rules specified below are intended to replace this rule by issuing an “http://schemas.microsoft.com/authorization/claims/permit” claim based on a

more specific set of conditions. Once you have finished adding the rules specified below for your scenario, be sure to come back and delete the “Permit Access to All Users” rule. Otherwise, your client access policy rules will not work.

4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
5. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim for internal client requests directly to federation server”. Under **Custom rule**, type or paste the following claim rule language syntax:

   ```
   c:[Type == "http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork", Value == "true"] => issue(Type = "http://custom/allow", Value = "true");
   ```
6. Click **Finish**. Verify that the new rule appears in the **Issuance Authorization Rules** list.
7. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
8. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
9. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim based on IP regex”. Under **Custom rule**, type or paste the following claim rule language syntax: `c:[Type == "http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-forwarded-client-ip", Value =~ "customer provided public ip address regex"] => issue(Type = "http://custom/allow", Value = "true");`

   You will have to replace the value above for “public ip address regex” with a valid IP expression; see Building the IP address range expression for more information.
10. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
11. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
12. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.

    On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim based on the presence of the Exchange ActiveSync Header”. Under **Custom rule**, type or paste the following claim rule language syntax: `c:[type=="http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-client-application", Value=="ONE_OF_EXO_CLIENT_APP_VALUES"] => issue(Type = "http://custom/allowed", Value = "true");`

    Warning
    For more information, see X-MS-Client-Application.
13. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
14. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
15. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims**

Using a Custom Rule, and then click **Next**.

On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue permit claim based on the custom claim”. Under **Custom rule**, type or paste the following claim rule language syntax:

```
c:[Type == "http://custom/allow", Value == "true"] => issue(Type = "http://schemas.microsoft.com/authorization/claims/permit", Value = "true");
```

16. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
17. If you still have the default “Permit Access to All Users” rule in the Issuance Authentication Rules list, select it and click **Remove Rule**. Click **Yes** to confirm.
18. To save the new rules, in the **Edit Claim Rules** dialog box, click **OK**.

Scenario 3: Block all external access to Office 365 except browser-based applications

To create rules to block all external access to Office 365 except browser-based applications

1. From **Server Manager**, click **Tools**, then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click **Relying Party Trusts**, right-click the **Microsoft Office 365 Identity Platform trust**, and then click **Edit Claim Rules**.
3. In the **Edit Claim Rules** dialog box, select the **Issuance Authorization Rules** tab, and then click **Add Rule** to start the Claim Rule Wizard.

   Note

   If you have never configured any custom Issuance Authorization Rules on the Microsoft Office 365 Identity Platform relying party trust, you will see the default “Permit Access to All Users” rule. The rules specified below are intended to replace this rule by issuing an “http://schemas.microsoft.com/authorization/claims/permit” claim based on a more specific set of conditions. Once you have finished adding the rules specified below for your scenario, be sure to come back and delete the “Permit Access to All Users” rule. Otherwise, your client access policy rules will not work.
4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
5. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim for internal client requests directly to federation server”. Under **Custom rule**, type or paste the following claim rule language syntax:

   ```
   c:[Type == "http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork", Value == "true"] => issue(Type = "http://custom/allow", Value = "true");
   ```
6. Click **Finish**. Verify that the new rule appears in the Issuance Authorization Rules list.
7. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab,

click **Add Rule** to start the Claim Rule Wizard again.

8. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.

   On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim based on IP regex”. Under Custom rule, type or paste the following claim rule language syntax:
   ```
   c:[Type == "http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-forwarded-client-ip", Value =~ "customer provided public ip address regex"] => issue(Type = "http://custom/allow", Value = "true");
   ```
   You will have to replace the value above for “public ip address regex” with a valid IP expression; see Building the IP address range expression for more information.
9. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
10. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
11. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
12. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim based on the endpoint”. Under **Custom rule**, type or paste the following claim rule language syntax:
    ```
    c:[Type == "http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-endpoint-absolute-path", Value == "/adfs/ls"] => issue(Type = "http://custom/allowed", Value = "true");
    ```
13. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
14. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
15. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.

    On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue permit claim based on the custom claim”. Under Custom rule, type or paste the following claim rule language syntax:
    ```
    c:[Type == "http://custom/allow", Value == "true"] => issue(Type = "http://schemas.microsoft.com/authorization/claims/permit", Value = "true");
    ```
16. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
17. If you still have the default “Permit Access to All Users” rule in the Issuance Authentication Rules list, select it and click **Remove Rule**. Click **Yes** to confirm.
18. To save the new rules, in the **Edit Claim Rules** dialog box, click OK.

Scenario 4: Block all external access to Office 365 except for designated Active Directory groups

The following example enables access from internal clients based on IP address. It blocks access from clients residing outside the corporate network that have an external client IP address, except for those individuals in a specified Active Directory Group.Use the following steps to add the correct Issuance Authorization rules to the **Microsoft Office 365 Identity Platform** relying party trust using the Claim Rule Wizard:

To create rules to block all external access to Office 365, except for designated Active Directory groups

1. From **Server Manager**, click **Tools**, then click **AD FS Management**.
2. In the console tree, under **AD FS\Trust Relationships**, click **Relying Party Trusts**, right-click the **Microsoft Office 365 Identity Platform** trust, and then click **Edit Claim Rules**.
3. In the **Edit Claim Rules** dialog box, select the **Issuance Authorization Rules** tab, and then click **Add Rule** to start the Claim Rule Wizard.

   Note
   If you have never configured any custom Issuance Authorization Rules on the Microsoft Office 365 Identity Platform relying party trust, you will see the default "Permit Access to All Users" rule. The rules specified below are intended to replace this rule by issuing an "http://schemas.microsoft.com/authorization/claims/permit" claim based on a more specific set of conditions. Once you have finished adding the rules specified below for your scenario, be sure to come back and delete the "Permit Access to All Users" rule. Otherwise, your client access policy rules will not work.
4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
5. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example "Issue custom claim for internal client requests directly to federation server". Under **Custom rule**, type or paste the following claim rule language syntax:c:[Type == "http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork", Value == "true"]=> issue(Type = "http://custom/allow", Value = "true");
6. Click **Finish**. Verify that the new rule appears in the Issuance Authorization Rules list.
7. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
8. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
9. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example "Issue custom claim based on IP regex". Under Custom rule, type or paste the following claim rule language syntax:`c:[Type == "http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-forwarded-client-ip", Value =~ "customer provided public ip address regex"] => issue(Type =`

```
"http://custom/allow", Value = "true");
```

You will have to replace the value above for “public ip address regex” with a valid IP expression; see Building the IP address range expression for more information.

10. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
11. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
12. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
13. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue custom claim based on presence of AD group SID”. Under **Custom rule**, type or paste the following claim rule language syntax:`c:[type=="http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsid", Value=="S-1-5-21-913475896-37678304-286007319-2105"]` => issue(Type = "http://custom/allowed", Value = "true");
14. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
15. Next, in the **Edit Claim Rules** dialog box, on the **Issuance Authorization Rules** tab, click **Add Rule** to start the Claim Rule Wizard again.
16. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
17. On the **Configure Rule** page, under **Claim rule name**, type the display name for this rule, for example “Issue permit claim based on the custom claim”. Under **Custom rule**, type or paste the following claim rule language syntax:
```
c:[Type == "http://custom/allow", Value == "true"] => issue(Type = "http://schemas.microsoft.com/authorization/claims/permit", Value = "true");
```
18. Click **Finish**. Verify that the new rule appears immediately below the previous rule in the Issuance Authorization Rules list.
19. If you still have the default “Permit Access to All Users” rule in the Issuance Authentication Rules list, select it and click **Remove Rule**. Click **Yes** to confirm.
20. To save the new rules, in the **Edit Claim Rules** dialog box, click **OK**.

Building the IP address range expression

The x-ms-forwarded-client-ip claim is populated from an HTTP header that is currently set only by Exchange Online, which populates the header when passing the authentication request to AD FS. The value of the claim may be one of the following:

Note
Exchange Online currently supports only IPV4 and not IPV6 addresses.

- A single IP address: The IP address of the client that is directly connected to Exchange Online

Notes

- The IP address of a client on the corporate network will appear as the external interface IP address of the organization's outbound proxy or gateway.
- Clients that are connected to the corporate network by a VPN or by Microsoft DirectAccess (DA) may appear as internal corporate clients or as external clients depending upon the configuration of VPN or DA.

- One or more IP addresses: When Exchange Online cannot determine the IP address of the connecting client, it will set the value based on the value of the x-forwarded-for header, a non-standard header that can be included in HTTP-based requests and is supported by many clients, load balancers, and proxies on the market.

Notes

- Multiple IP addresses, indicating the client IP address and the address of each proxy that passed the request, will be separated by a comma.
- IP addresses related to Exchange Online infrastructure will not on the list.

Regular Expressions

When you have to match a range of IP addresses, it becomes necessary to construct a regular expression to perform the comparison. In the next series of steps, we will provide examples for how to construct such an expression to match the following address ranges (note that you will have to change these examples to match your public IP range):

- 192.168.1.1 – 192.168.1.25
- 10.0.0.1 – 10.0.0.14

First, the basic pattern that will match a single IP address is as follows: `\b###\.###\.###\.###\b`

Extending this, we can match two different IP addresses with an OR expression as follows:

`\b###\.###\.###\.###\b|\b###\.###\.###\.###\b`

So, an example to match just two addresses (such as 192.168.1.1 or 10.0.0.1) would be:

`\b192\.168\.1\.1\b|\b10\.0\.0\.1\b`

This gives you the technique by which you can enter any number of addresses. Where a range of address need to allowed, for example 192.168.1.1 – 192.168.1.25, the matching must be done character by character: `\b192\.168\.1\.([1-9]|1[0-9]|2[0-5])\b`

Note

The IP address is treated as string and not a number.

The rule is broken down as follows: `\b192\.168\.1\.`

This matches any value beginning with 192.168.1.

The following matches the ranges required for the portion of the address after the final decimal point:

- ([1-9] Matches addresses ending in 1-9
- |1[0-9] Matches addresses ending in 10-19
- |2[0-5]) Matches addresses ending in 20-25

Note

Note that the parentheses must be correctly positioned, so that you don't start matching other portions of IP addresses.

With the 192 block matched, we can write a similar expression for the 10 block: `\b10\.0\.0\.([1-9]|1[0-4])\b`

And putting them together, the following expression should match all the addresses for "192.168.1.1～25" and "10.0.0.1～14": `\b192\.168\.1\.([1-9]|1[0-9]|2[0-5])\b|\b10\.0\.0\.([1-9]|1[0-4])\b`

Testing the Expression

Regex expressions can become quite tricky, so we highly recommend using a regex verification tool. If you do an internet search for "online regex expression builder", you will find several good online utilities that will allow you to try out your expressions against sample data.

When testing the expression, it's important that you understand what to expect to have to match. The Exchange online system may send many IP addresses, separated by commas. The expressions provided above will work for this. However, it's important to think about this when testing your regex expressions. For example, one might use the following sample input to verify the examples above:

```
192.168.1.1, 192.168.1.2, 192.169.1.1. 192.168.12.1, 192.168.1.10, 192.168.1.25, 192.168.1.26, 192.168.1.30, 1192.168.1.20

10.0.0.1, 10.0.0.5, 10.0.0.10, 10.0.1.0, 10.0.1.1, 110.0.0.1, 10.0.0.14, 10.0.0.15, 10.0.0.10, 10,0.0.1
```

Claim Types

AD FS in Windows Server 2012 R2 provides request context information using the following claim types:

X-MS-Forwarded-Client-IP

Claim type: `http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-forwarded-client-ip`

This AD FS claim represents a "best attempt" at ascertaining the IP address of the user (for example, the Outlook client) making the request. This claim can contain multiple IP addresses, including the address of every proxy that forwarded the request. This claim is populated from an HTTP. The value of the claim can be one of the following:

- A single IP address - The IP address of the client that is directly connected to Exchange Online

  Note

The IP address of a client on the corporate network will appear as the external interface IP address of the organization's outbound proxy or gateway.

- One or more IP addresses
  - If Exchange Online cannot determine the IP address of the connecting client, it will set the value based on the value of the x-forwarded-for header, a non-standard header that can be included in HTTP based requests and is supported by many clients, load balancers, and proxies on the market.
  - Multiple IP addresses indicating the client IP address and the address of each proxy that passed the request will be separated by a comma.

Note

IP addresses related to Exchange Online infrastructure will not be present in the list.

Warning

Exchange Online currently supports only IPV4 addresses; it does not support IPV6 addresses.

X-MS-Client-Application

Claim type: `http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-client-application`

This AD FS claim represents the protocol used by the end client, which corresponds loosely to the application being used. This claim is populated from an HTTP header that is currently only set by Exchange Online, which populates the header when passing the authentication request to AD FS. Depending on the application, the value of this claim will be one of the following:

- In the case of devices that use Exchange Active Sync, the value is Microsoft.Exchange.ActiveSync.
- Use of the Microsoft Outlook client may result in any of the following values:
  - Microsoft.Exchange.Autodiscover
  - Microsoft.Exchange.OfflineAddressBook
  - Microsoft.Exchange.RPC
  - Microsoft.Exchange.WebServices
- Other possible values for this header include the following:
  - Microsoft.Exchange.Powershell
  - Microsoft.Exchange.SMTP
  - Microsoft.Exchange.PopImap

X-MS-Client-User-Agent

Claim type: `http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-client-user-agent`

This AD FS claim provides a string to represent the device type that the client is using to access the service. This can be used when customers would like to prevent access for certain devices (such as particular types of smart phones). Example values for this claim include (but are not limited to) the values below.

Note
The below are examples of what the x-ms-user-agent value might contain for a client whose x-ms-client-application is “Microsoft.Exchange.ActiveSync”

- Vortex/1.0
- Apple-iPad1C1/812.1
- Apple-iPhone3C1/811.2
- Apple-iPhone/704.11
- Moto-DROID2/4.5.1
- SAMSUNGSPHD700/100.202
- Android/0.3

Note
It is also possible that this value is empty.

X-MS-Proxy

Claim type: `http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-proxy`

This AD FS claim indicates that the request has passed through the Web Application proxy. This claim is populated by the Web Application proxy, which populates the header when passing the authentication request to the back end Federation Service. AD FS then converts it to a claim.

The value of the claim is the DNS name of the Web Application proxy that passed the request.

InsideCorporateNetwork

Claim type: `http://schemas.microsoft.com/ws/2012/01/insidecorporatenetwork`

Similar to the above x-ms-proxy claim type, this claim type indicates whether the request has passed through the web application proxy. Unlike x-ms-proxy, insidecorporatenetwork is a boolean value with True indicating a request directly to the federation service from inside the corporate network.

X-MS-Endpoint-Absolute-Path (Active vs Passive)

Claim type: `http://schemas.microsoft.com/2012/01/requestcontext/claims/x-ms-endpoint-absolute-path`

This claim type can be used for determining requests originating from “active” (rich) clients versus “passive” (web-browser-based) clients. This enables external requests from browser-based applications such as the Outlook Web Access, SharePoint Online, or the Office 365 portal to be allowed while requests originating from rich clients such as Microsoft Outlook are blocked.

The value of the claim is the name of the AD FS service that received the request.

See Also

AD FS Operations

AD FS Technical Reference

- Understanding Key AD FS Concepts

Tip
You can find additional AD FS 2.0 design content at the AD FS 2.0 Content Map page on the Microsoft TechNet Wiki. This page is managed by members of the AD FS 2.0 Community and is monitored on a regular basis by the AD FS Product Team.

See Also

Active Directory Federation Services Overview

Understanding Key AD FS Concepts

It is recommended that you learn about the important concepts for Active Directory Federation Services (AD FS) and become familiar with its feature set.

Tip
You can find additional AD FS resource links at the AD FS Content Map page on the Microsoft TechNet Wiki. This page is managed by members of the AD FS Community and is monitored on a regular basis by the AD FS Product Team.

AD FS terminology used in this guide

AD FS term	Definition
Account partner organization	A federation partner organization that is represented by a claims provider trust in the Federation Service. The account partner organization contains the users that will access Web-based applications in the resource partner.
Account federation server	The federation server in the account partner

AD FS term	Definition
	organization. The account federation server issues security tokens to users based on user authentication. The server authenticates the user, extracts the relevant attributes and group membership information out of the attribute store, packages this information into claims, and generates and signs a security token (which contains the claims) to return to the user—either to be used in its own organization or to be sent to a partner organization.
AD FS configuration database	A database used to store all configuration data that represents a single AD FS instance or Federation Service. This configuration data can be stored in either a SQL Server database or using the Windows Internal Database feature included with Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012. You can create the AD FS configuration database for SQL Server using the Fsconfig.exe command-line tool and for Windows Internal Database using the AD FS Federation Server Configuration Wizard.
Claims provider	The organization that provides claims to its users. See account partner organization.
Claims provider trust	In the AD FS Management snap-in, claims provider trusts are trust objects typically created in resource partner organizations to represent the organization in the trust relationship whose accounts will be accessing resources in the resource partner organization. A claims provider trust object consists of a variety of identifiers, names, and rules that identify this partner to the local Federation Service.
Federation metadata	The data format for communicating configuration information between a claims provider and a relying party to facilitate proper configuration of claims provider trusts and relying party trusts. The data format is defined in Security Assertion Markup Language

AD FS term	Definition
	(SAML) 2.0, and it is extended in WS-Federation.
Federation server	A computer running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012 that has been configured using the AD FS Federation Server Configuration Wizard to act in the federation server role. A federation server issues tokens and serves as part of a Federation Service.
Federation server proxy	A computer running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012 that has been configured using the AD FS Federation Server Proxy Configuration Wizard to act as an intermediary proxy service between an Internet client and a Federation Service that is located behind a firewall on a corporate network.
Primary federation server	A computer running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012 that has been configured in the federation server role using the AD FS Federation Server Configuration Wizard and has a read/write copy of the AD FS configuration database. The primary federation server is created when you use the AD FS Federation Server Configuration Wizard and select the option to create a new Federation Service and make that computer the first federation server in the farm. All other federation servers in this farm must replicate changes made on the primary federation server to a read-only copy of the AD FS configuration database that is stored locally. The term “primary federation server” does not apply when the AD FS configuration database is stored in an SQL database as all federation servers can equally read and write to a configuration database stored on a SQL Server.
Relying party	The organization that receives and processes

AD FS term	Definition
	claims. See resource partner organization.
Relying party trust	In the AD FS Management snap-in, relying party trusts are trust objects typically created in: • Account partner organizations to represent the organization in the trust relationship whose accounts will be accessing resources in the resource partner organization. • Resource partner organizations to represent the trust between the Federation Service and a single web-based application. A relying party trust object consists of a variety of identifiers, names, and rules that identify this partner or web-application to the local Federation Service.
Resource federation server	The federation server in the resource partner organization. The resource federation server typically issues security tokens to users based on a security token that is issued by an account federation server. The server receives the security token, verifies the signature, applies claim rule logic to the unpackaged claims to produce the desired outgoing claims, generates a new security token (with the outgoing claims) based on information in the incoming security token, and signs the new token to return to the user and ultimately to the Web application.
Resource partner organization	A federation partner that is represented by a relying party trust in the Federation Service. The resource partner issues claims-based security tokens that contains published Web-based applications that users in the account partner can access.

Overview of AD FS

AD FS is an identity access solution that provides client computers (internal or external to your network) with seamless SSO access to protected Internet-facing applications or services, even

when the user accounts and applications are located in completely different networks or organizations.

When an application or service is in one network and a user account is in another network, typically the user is prompted for secondary credentials when he or she attempts to access the application or service. These secondary credentials represent the user's identity in the realm where the application or service resides. They are usually required by the Web server that hosts the application or service so that it can make the most appropriate authorization decision.

With AD FS, organizations can bypass requests for secondary credentials by providing trust relationships (federation trusts) that these organizations can use to project a user's digital identity and access rights to trusted partners. In this federated environment, each organization continues to manage its own identities, but each organization can also securely project and accept identities from other organizations.

- The Role of Attribute Stores
- The Role of the AD FS Configuration Database
- The Role of Claims
- The Role of Claim Rules
- The Role of the Claims Engine
- The Role of the Claims Pipeline
- The Role of the Claim Rule Language
- Determine the Type of Claim Rule Template to Use
- How URIs Are Used in AD FS

The Role of Attribute Stores

Active Directory Federation Services (AD FS) uses the term “attribute stores” to refer to directories or databases that an organization uses to store its user accounts and their associated attribute values. After it is configured in an identity provider organization, AD FS retrieves these attribute values from the store and creates claims based on that information so that a Web application or service that is hosted in a relying party organization can make the appropriate authorization decisions whenever a federated user (a user whose account is stored in the identity provider organization) attempts to access the application or service.

For more information about how claims are generated, see The Role of Claims.

How attribute stores fit in with your AD FS deployment goals

The location of the user attribute store and the location from which users authenticate determine how you design AD FS to support the user identities. Depending on where the attribute store is located and where users will access the application (in an intranet or on the Internet), you can use one of the following deployment goals:

- Provide Your Active Directory Users Access to Your Claims-Aware Applications and Services—In this goal, users in your organization access an AD FS–secured application or service (either your own application or service or a partner's application or service) when the users are logged on to Active Directory in the corporate intranet.
- Provide Your Active Directory Users Access to the Applications and Services of Other Organizations—In this goal, users in your organization access an AD FS–secured application or service (either your own application or service or a partner's application or service) when the users are logged on to an attribute store in the corporate intranet and when they log on remotely from the Internet.
- Provide Users in Another Organization Access to Your Claims-Aware Applications and Services—In this goal, user accounts in another organization that are located in an attribute store on that organization's corporate intranet must access an AD FS–secured application in your organization. This goal also works when consumer-based user accounts that are located in an attribute store in your organization's perimeter network must be provided with access to an AD FS–secured application in your organization.

Depending on attribute store placement and other requirements of your organization, you can combine several of these deployment goals to complete the design of your AD FS deployment.

Attribute stores that are supported by AD FS

AD FS supports a wide range of directory and database stores that you can use for extracting administrator-defined attribute values and populating claims with those values. AD FS supports any of the following directories or databases as attribute stores:

- Active Directory in Windows Server 2003, Active Directory Domain Services (AD DS) in Windows Server 2008, AD DS in Windows Server® 2012
- All editions of Microsoft SQL Server 2005 and SQL Server 2008
- Custom attribute stores

The Role of the AD FS Configuration Database

The AD FS configuration database stores all the configuration data that represents a single instance of Active Directory Federation Services (AD FS) (that is, the Federation Service). The AD FS configuration database defines the set of parameters that a Federation Service requires to

identify partners, certificates, attribute stores, claims, and various data about these associated entities. You can store this configuration data in either a Microsoft SQL Server® database or the Windows Internal Database (WID) feature that is included with Windows Server® 2008, Windows Server 2008 R2 and Windows Server® 2012.

Note

The entire contents of the AD FS configuration database can be stored either in an instance of WID or in an instance of the SQL database, but not both. This means that you cannot have some federation servers using WID and others using a SQL Server database for the same instance of the AD FS configuration database.

You can use the following information in this topic along with the content provided in AD FS Deployment Topology Considerations to learn about the advantages and disadvantages of choosing either WID or SQL Server to store the AD FS configuration database:

WID uses a relational data store and does not have its own management user interface (UI). Instead, administrators can modify the contents of the AD FS configuration database by using either the AD FS Management snap-in, Fsconfig.exe, or Windows PowerShell™ cmdlets.

Using WID to store the AD FS configuration database

You can create the AD FS configuration database using WID as the store by using either the Fsconfig.exe command-line tool or the AD FS Federation Server Configuration Wizard. When you use either of these tools, you can choose any of the following options to create your federation server topology. Each of these options uses WID for storing the AD FS configuration database:

- Create a stand-alone federation server
- Create the first federation server in a federation server farm
- Add a federation server to a federation server farm

If you select the stand-alone option, WID is used to store a single instance of the AD FS configuration database. This instance cannot be shared across multiple federation servers. It is meant for test lab environments only. For more information about the stand-alone federation server option or how to set one up, see Stand-Alone Federation Server Using WID or Create a Stand-Alone Federation Server.

If you select the first federation server in a federation server farm option, WID is configured for scalability that will permit additional federation servers to be added to the farm at a later time. For more information about deploying a WID farm or how to set one up, see Federation Server Farm Using WID or Create the First Federation Server in a Federation Server Farm

If you select the add a federation server option, WID is configured to replicate configuration database changes to the new federation server at set intervals. For more information about adding a federation server to a WID farm, see Federation Server Farm Using WID or Add a Federation Server to a Federation Server Farm.

Note

When you deploy a federation server farm using WID, some features of AD FS may not be available. To have access to the full feature set when you configure your server farm, consider using Microsoft SQL Server to store the AD FS configuration database instead. For more information, see AD FS Deployment Topology Considerations.

How a WID federation server farm works

This section describes important concepts that describe how the WID federation server farm replicates data between a primary federation server and secondary federation servers. .

Primary federation server

A primary federation server is a computer running Windows Server 2008, Windows Server 2008 R2 or Windows Server® 2012 that has been configured in the federation server role with the AD FS Federation Server Configuration Wizard and that has a read/write copy of the AD FS configuration database. The primary federation server is always created when you use the AD FS Federation Server Configuration Wizard and select the option to create a new Federation Service and make that computer the first federation server in the farm. All other federation servers in this farm, also known as secondary federation servers, must synchronize changes that are made on the primary federation server to a copy of the AD FS configuration database that is stored locally.

Secondary federation servers

Secondary federation servers store a copy of the AD FS configuration database from the primary federation server, but these copies are read-only. Secondary federation servers connect to and synchronize the data with the primary federation server in the farm by polling it at regular intervals to check whether data has changed. The secondary federation servers exist to provide fault tolerance for the primary federation server while acting to load-balance access requests that are made in different sites throughout your network environment.

Note

If a primary federation server crashes and is offline, all secondary federation servers continue to process requests as normal. However, no new changes can be made to the Federation Service until the primary federation server has been brought back online. You can also nominate a secondary federation server to become the primary federation server by using Windows PowerShell. For more information, see the AD FS Administration with Windows PowerShell.

How the AD FS configuration database is synchronized

Because of the important role that the AD FS configuration database plays, it is made available on all the federation servers in the network to provide fault tolerance and load-balancing capabilities when processing requests (when network load-balancers are used). However, for

secondary federation servers to serve in this capacity, the AD FS configuration database that is stored on the primary federation server must be synchronized.

When you add a federation server to the farm, the new computer that will become a secondary federation server connects to the primary federation server to replicate the copy of the AD FS configuration database. From this point forward, the new federation server continues to pull updates from the primary federation server on a regular basis, as shown in the following illustration.

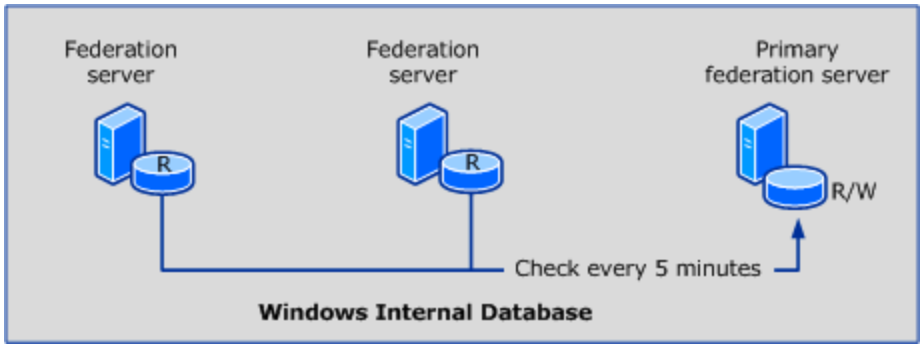

Each secondary federation server polls the primary federation server every five minutes for changes. You can adjust this default five-minute value or force an immediate synchronization anytime by using a Windows PowerShell cmdlet. For more information about how to do this, see AD FS Administration with Windows PowerShell.

The WID synchronization process also supports incremental transfers for more efficient transfers of intermediate changes. The incremental transfer process requires substantially less traffic on a network, and transfers are completed much faster.

Note
The migration of an AD FS configuration database from WID to an instance of SQL Server is supported. For more information about how to do this, see AD FS: Migrate Your AD FS Configuration Database to SQL Server on the TechNet Wiki site.

Using SQL Server to store the AD FS configuration database

You can create the AD FS configuration database using a single SQL Server database instance as the store by using the Fsconfig.exe command-line tool. Using a SQL Server database as the AD FS configuration database provides the following benefits over WID:

- Administrators can leverage the high availability features of SQL Server
- It provides additional performance increases for high traffic.
- It provides feature support of SAML artifact resolution and SAML/WS-Federation token replay detection (described below).

The term “primary federation server” does not apply when the AD FS configuration database is stored in a SQL database instance because all federation servers can equally read and write to

the AD FS configuration database that is using the same clustered SQL Server instance, as shown in the following illustration.

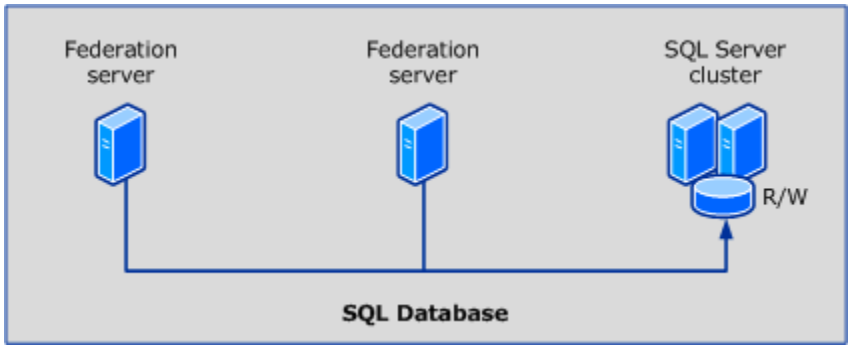

You can use SQL Server to configure two or more servers to work together as a server cluster to ensure that AD FS is made highly available to service incoming client requests. High availability provides a scale-out architecture in which you can increase server capacity by adding additional servers. Single points of failure are mitigated by automatic cluster failover.

You can achieve high availability by using the network load-balancing and failover services that SQL clustering technologies provide. For more information about how to configure SQL Server for high availability, see High Availability Solutions Overview (http://go.microsoft.com/fwlink/?LinkId=179853).

SAML artifact resolution

Security Assertion Markup Language (SAML) artifact resolution is an endpoint based on the part of the SAML 2.0 protocol that describes how a relying party can retrieve a token directly from a claims provider. In the first stage of the resolution process, a browser client contacts a resource federation server and provides it with an artifact. In the second stage, resource federation servers send the artifact to a SAML artifact endpoint URL that is hosted somewhere in an account partner organization in order to resolve the artifact message. In the final stage, the account federation server issues the token to the federation server on behalf of the browser client.

Note

If you are an administrator in an account partner organization, make sure to assign or bind an SSL certificate, which chains to a root certificate of a member of the Windows Root Certificate Program, to the federation passive Web site in IIS (<ComputerName>\Sites\Default Web Site\adfs\ls) on all the account federation servers in the farm. This is important to prevent resource federation servers from having to manually add the SSL certificate to the Local Computers Trusted People certificate store or from being unable to resolve the artifact that is published in your organization.

SAML/WS-Federation token replay detection

The term *token replay* refers to the act by which a browser client in an account partner organization attempts to send the same token it received from an account federation server

multiple times to authenticate to a resource federation server. This act occurs when a user clicks the **Back** button of their browser in an effort to resubmit the authentication page.

AD FS provides a feature referred to as *token replay detection* by which multiple token requests using the same token can be detected and then discarded. When this feature is enabled, token replay detection protects the integrity of authentication requests in both the WS-Federation passive profile and the SAML WebSSO profile by making sure that the same token is never used more than once. This feature should be enabled in situations where security is a very high concern such as when using kiosks.

In the kiosk example, a user can log off of all Web sites and later a malicious user can attempt to use the browser history in order to resubmit the federated authentication page that was loaded by the previous user. This feature mitigates this concern by storing additional information about each successful authentication made by an account partner organization in order to detect subsequent replays of the token and prevent multiple authentication attempts from succeeding.

The Role of Claims

In the claims-based identity model, claims play a pivotal role in the federation process, They are the key component by which the outcome of all Web-based authentication and authorization requests are determined. This model enables organizations to securely project digital identity and entitlement rights, or *claims*, across security and enterprise boundaries in a standardized way.

What are claims?

In its simplest form, claims are simply *statements* (for example, name, identity, group), made about users, that are used primarily for authorizing access to claims-based applications located anywhere on the Internet. Each statement corresponds to a *value* that is stored in the claim.

How claims are sourced

The Federation Service in Active Directory Federation Services (AD FS) defines which claims are exchanged between federated partners. However, before it can do this it must first populate or source the claim with either a retrieved value or a calculated value. Each claim value represents a value of a user, group, or entity and is sourced in one of two ways:

1. When the value that makes up the claim is retrieved from an attribute store, for example, when an attribute value of Sales Department is retrieved from the properties of an Active Directory user account. For more information, see The Role of Attribute Stores.
2. When the value of an incoming claim is transformed into another value based on the logic expressed in a rule. For example, when an incoming claim with the value of Domain Admins is transformed into a new value of Administrators before it is sent as an outgoing claim. For more information, see The Role of Claim Rules.

Claims can include values such as an e-mail address, User Principal Name (UPN), group membership, and other account attributes.

How claims flow

Other parties rely on the values of the claims to perform authorization tasks for Web-based applications that they host. These parties are referred to as *relying parties* in the AD FS Management snap-in. The Federation Service is responsible for brokering trust between many disparate parties. It is designed to process and flow the trusted exchange of claims from an organization that initially sources the claims, also referred to as *claims providers* in the AD FS Management snap-in, to a relying party. A relying party then uses these claims to make authorization decisions.

The flow of claims using this process is known as the *claims pipeline*. There are three steps in the flow of claims through the claims pipeline:

1. The claims that are received from the claims provider are processed by the acceptance transform rules on the claims provider trust. These rules determine which claims are accepted from the claims provider.
2. The output of the acceptance transform rules is used as input to the issuance authorization rules. These rules determine whether the user is permitted to access the relying party.
3. The output of the acceptance transform rules is used as input to the issuance transform rules. These rules determine the claims that will be sent to the relying party.

For more information, see The Role of the Claims Pipeline

How claims are issued

When you write claim rules, the source of the incoming claims for the claim rules varies based on whether you are writing rules on a claims provider trust or a relying party trust. When you write claim rules for a claims provider trust, the incoming claims are the claims sent from the trusted claims provider to the Federation Service. When you write rules for a relying party trust, the incoming claims are the claims that are output by the claim rules of the applicable claims provider trust. For more information about incoming claims and outgoing claims, see The Role of the Claims Pipeline and The Role of the Claims Engine.

What are claim types?

A claim type provides context for the claim value. It is usually expressed as a Uniform Resource Identifier (URI). AD FS can support any claim type, and it is configured with the claim types in the following table by default.

Name	Description	URI
E-Mail Address	The e-mail address of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress
Given Name	The given name of the	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname

Name	Description	URI
	user	
Name	The unique name of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name
UPN	The user principal name (UPN) of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn
Common Name	The common name of the user	http://schemas.xmlsoap.org/claims/CommonName
AD FS 1.x E-Mail Address	The e-mail address of the user when interoperating with AD FS 1.1 or ADFS 1.0	http://schemas.xmlsoap.org/claims/EmailAddress
Group	A group that the user is a member of	http://schemas.xmlsoap.org/claims/Group
AD FS 1.x UPN	The UPN of the user when interoperating with AD FS 1.1 or ADFS 1.0	http://schemas.xmlsoap.org/claims/UPN
Role	A role that the user has	http://schemas.microsoft.com/ws/2008/06/identity/claims/role
Surname	The surname of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname
PPID	The private identifier of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/privatepersonalidentifier
Name Identifier	The SAML name identifier of	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/nameidentifier

Name	Description	URI
	the user	
Authentication Method	The method used to authenticate the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/authenticationmethod
Deny Only Group SID	The deny-only group SID of the user	http://schemas.xmlsoap.org/ws/2005/05/identity/claims/denyonlysid
Deny only primary SID	The deny-only primary SID of the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/denyonlyprimarysid
Deny only primary group SID	The deny-only primary group SID of the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/denyonlyprimarygroupsid
Group SID	The group SID of the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsid
Primary group SID	The primary group SID of the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/primarygroupsid
Primary SID	The primary SID of the user	http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid
Windows account name	The domain account name of the user in the form of <domain>\<user>	http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname

What are claim descriptions?

Claim descriptions represent a list of claims types that AD FS supports and that may be published in federation metadata. The claim types mentioned in the previous table are configured as claims descriptions in the AD FS Management snap-in.

The collection of claim descriptions that will be published to federation metadata is stored in the AD FS configuration database. These claim descriptions are used by various components of the Federation Service.

Each claim description includes a claim type URI, name, publishing state, and description. You can manage the claim description collection by using the **Claim Descriptions** node in the AD FS Management snap-in. You can modify the publishing state of a claim description using the snap-in. The following settings are available:

- **Publish this claim in federation metadata as a claim type that this Federation Service can accept** (Publish as Accepted)—Indicates the claim types that will be accepted from other claims providers by this Federation Service.
- **Publish this claim in federation metadata as a claim type that this Federation Service can send** (Publish as Sent)—Indicates the claim types that are offered by this Federation Service. These are the claim types the Federation Service publishes to others as those it is willing to send. The actual claim types sent by the claims provider are often a subset of this list.

For more information about how to set the publishing state of a claim type, see Add a Claim Description in the AD FS Deployment Guide.

When generating Federation Metadata

Federation Metadata includes all the claim descriptions that are marked for publishing.

When claims rules are processed

When you keep configuration information about claims descriptions, it is easier for you to configure rules about claims. For more information about the claim rules that can be used in the claims provider organization, see The Role of Claim Rules.

The Role of Claim Rules

The overall function of the Federation Service in Active Directory Federation Services (AD FS) is to issue a token that contains a set of claims. The decision regarding what claims AD FS accepts and then issues is governed by claim rules.

What are claim rules?

A claim rule represents an instance of business logic that will take one or more incoming claims, apply conditions to them (if x then y) and produce one or more outgoing claims based on the condition parameters. For more information about incoming and outgoing claims, see The Role of Claims.

You use claim rules when you need to implement business logic that will control the flow of claims through the claims pipeline. While the claims pipeline is more a logical concept of the end-to-end

process for flowing claims, claim rules are an actual administrative element that you can use to customize the flow of claims through the claims issuance process.

For more information about the claims pipeline, see The Role of the Claims Engine.

Claim rules provide the following benefits:

- Provide a mechanism for administrators to apply run-time business logic for trusting claims from claims providers
- Provide a mechanism for administrators to define what claims are released to relying parties
- Provide rich and detailed claims-based authorization capabilities to administrators who want to permit or deny access to specific users

How claim rules are processed

Claim rules are processed through the claims pipeline using the *claims engine*. The claims engine is a logical component of the Federation Service that examines the set of incoming claims presented by a user, and will then, depending on the logic in each rule, produce an output set of claims.

Together, the claims rule engine and the set of claim rules associated with a given federated trust determine whether incoming claims should be passed through as they are, filtered to meet a specific condition's criteria or transformed into an entirely new set of claims before they are issued as outgoing claims by your Federation Service.

For more information about this process, see The Role of the Claims Engine.

What are claim rule templates?

AD FS includes a predefined set of claim rule templates that are designed to help you easily select and create the most appropriate claim rules for your particular business need. Claim rule templates are only used during the claim rule creation process.

In the AD FS Management snap-in, rules can only be created using claim rule templates. After you use the snap-in to select a claim rule template, input the necessary data for the rule logic and save it to the configuration database, it will be (from that point forward) referred to in the UI as a claim rule.

How claim rule templates work

At first glance, claim rule templates appear to be just input forms provided by the snap-in to collect data and process specific logic on incoming claims. However, at a much more detailed level, claim rule templates store the necessary claim rule language framework that make up the base logic necessary for you to quickly create a rule without needing to know the language intimately.

Each template that is provided in the user interface (UI) represents a prepopulated claim rule language syntax, based on the most commonly required administrative tasks. There is one rule template however, that is the exception. This template is referred to as the custom rule template.

With this template, no syntax is prepopulated. Instead you must directly author the claim rule language syntax in the body of the claim rule template form using the claim rule language syntax.

For more information about how to use the claim rule language syntax, see The Role of the Claim Rule Language in the AD FS Deployment Guide.

Tip

You can view the claim rule language associated with a rule at any time by clicking the **View Rule Language** button on the properties of a claim rule.

How to create a claim rule

Claim rules are created separately for each federated trust relationship within the Federation Service and are not shared across multiple trusts. You can either create a rule from a claim rule template, start from scratch by authoring the rule using the claim rule language or use Windows PowerShell to customize a rule.

All of these options coexist to provide you with the flexibility of choosing the appropriate method for a given scenario. For more information about how to create a claim rule, see Configuring Claim Rules in the AD FSDeployment Guide.

Using claim rule templates

Claim rule templates are only used during the claim rule creation process. You can use any of the following templates to create a claim rule:

- Pass Through or Filter an Incoming Claim
- Transform an Incoming Claim
- Send LDAP Attributes as Claims
- Send Group Membership as a Claim
- Send Claims Using a Custom Rule
- Permit or Deny Users Based on an Incoming Claim
- Permit All Users

For more information describing each of these claim rule templates, see Determine the Type of Claim Rule Template to Use.

Using the claim rule language

For business rules that are beyond the scope of standard claim rule templates, you can use a custom rule template to express a series of complex logic conditions using the claim rule language. For more information about using a custom rule, see When to Use a Custom Claim Rule.

Using Windows PowerShell

You can also use the ADFSClaimRuleSet cmdlet object with Windows PowerShell to create or administer rules in AD FS. For more information about how you can use Windows PowerShell with this cmdlet, see AD FS Administration with Windows PowerShell.

What is a claim rule set?

As shown in the following illustration, a claim rule set is a grouping of one or more rules for a given federated trust that will define how claims will be processed by the claims rule engine. When an incoming claim is received by the Federation Service the claim rule engine applies the logic specified by the appropriate claim rule set. It is the final sum of the logic from each rule in the set that will determine how claims will be issued for a given trust in its entirety.

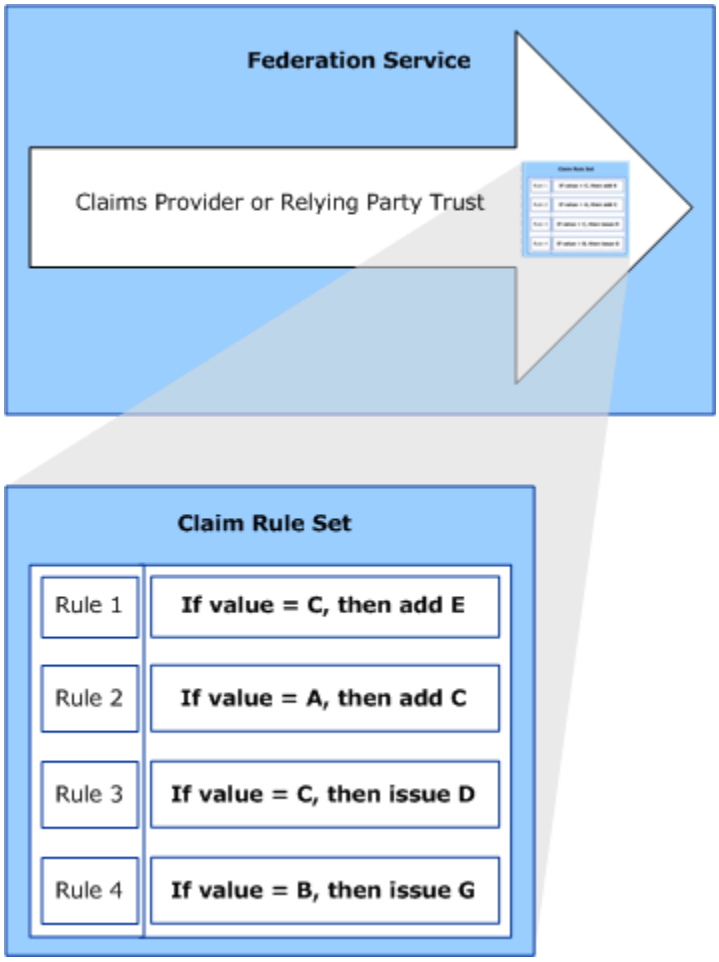

Claim rules are processed by the claims engine in chronological order within a given rule set. This order is important, because the output of one rule can be used as the input to the next rule in the set.

What are claim rule set types?

A claim rule set type is a logical segment of a federated trust that categorically identifies whether the claim rule set associated with the trust will be used for claims issuance, authorization or acceptance. Each federated trust can have one or more claim rule set types associated with it, depending on the type of trust that is used.

The following table describes the various types of claim rule sets and explains its relation with either a claims provider trust or relying party trust.

Claim rule set type	Description	Used on
Acceptance transform rule set	A set of claim rules that you use on a particular claims provider trust to specify the incoming claims that will be accepted from the claims provider organization and the outgoing claims that will be sent to the relying party trust. The incoming claims that will be used to source this rule set, will be the claims that are output by the issuance transform rule set as specified in the claims provider organization. By default, the claims provider trust node contains a claim provider trust named **Active Directory** which is used to represent the source attribute store for the acceptance transform rule set. This trust object is used to represent the connection from your Federation Service to an Active Directory database on your network. This default trust is what processes claims for users that have been authenticated by Active Directory and it cannot be deleted.	Claims provider trusts
Issuance Transform Rule Set	A set of claim rules that you use on a relying party trust to specify the claims that will be issued to the relying party.	Relying party trusts

Claim rule set type	Description	Used on
	The incoming claims that will be used to source this rule set, will initially be the claims that are output by the acceptance transform rules.	
Issuance Authorization Rule Set	A set of claim rules that you use on a relying party trust to specify the users that will be permitted to receive a token for the relying party. These rules determine whether a user can receive claims for a relying party and, therefore, access to the relying party. Unless you specify an issuance authorization rule, all users will be denied access by default.	Relying party trusts
Delegation Authorization Rule Set	A set of claim rules that you use on a relying party trust to specify the users that will be permitted to act as delegates for other users to the relying party. These rules determine whether the requester is permitted to impersonate a user while still identifying the requester in the token that is sent to the relying party. Unless you specify an issuance authorization rule, no users can act as delegates by default.	Relying party trusts
Impersonation Authorization Rule Set	A set of claim rules that you configure using Windows PowerShell to determine whether a user can fully impersonate another user to the relying party. These rules determine whether the requester is permitted to	Relying party trust

Claim rule set type	Description	Used on
	impersonate a user without identifying the requester in the token that is sent to the relying party. Impersonating another user in this way is a very powerful capability, because the relying party will not know that the user is being impersonated.	

For more information about select the appropriate claim rules to use in your organization, see Determine the Type of Claim Rule Template to Use.

The Role of the Claims Engine

At its highest level, the claims engine in Active Directory Federation Services (AD FS) is a rule-based engine that is dedicated to serving and processing claim requests for the Federation Service. The claims engine is the sole entity within the Federation Service that is responsible for running each of the rule sets across all of the federated trust relationships you have configured and handing the output result over to the claims pipeline.

While the claims pipeline is more a logical concept of the end-to-end process for flowing claims, claim rules are an actual administrative element that you can use to customize the flow of claims during the claim rules execution process. For more information about the pipeline process, see The Role of the Claims Pipeline.

As shown in the following illustration, the act of accepting incoming claims (acceptance rules), authorizing claims requesters (authorization rules) and issuing outgoing claims (issuance rules) through claim rules across all of the federated trust relationships in your organization is performed by the claims engine.

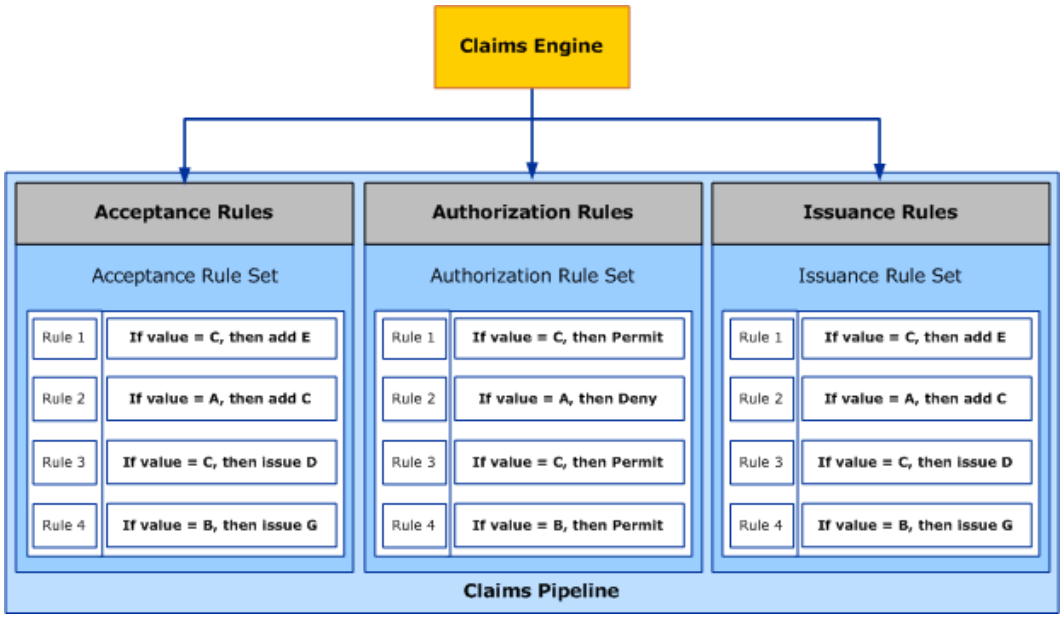

Claim rules execution process

When you configure a claims provider trust or relying party trust in your organization with claim rules, the claim rule set(s) for that trust act as a gatekeeper for incoming claims by invoking the claims engine to apply the necessary logic in the claim rules to determine whether to issue any claims and which claims to issue.

The following section outlines each of the steps that occur by the engine during the flow of claims through the claim rules execution process. Each of the steps as outlined below occurs for each stage described in the claims pipeline process. These steps include:

- Step 1 – Initialization
- Step 2 – Execution
- Step 3 – Execution result

For more information about the pipeline process, see The Role of the Claims Pipeline.

Step 1 – Initialization

In the first step of the claim rules execution process, the claims engine accepts incoming claims by first adding them to the *input claim set*. An input claim set is analogous to a cache in memory that is used to temporarily store data only as long as a required process requires that data to be made available for retrieval. The input claim set data is discarded after the rule execution finishes.

Adding a claim to the input claim set for a rule set

The input claim set is created by the claims engine when it needs to temporarily store claim data in memory while it processes the logic associated with a claim rule set. The claims engine copies

all of the incoming claims to the input claim set where they can be retrieved by the first rule in the rule set.

For example, in the illustration below, the claims engine reads the claims of A and B from the incoming claims and copies them to the input claim set. After they are in the input claim set, the claims engine retrieves and processes claims A and B as input for the logic in the first rule in the claim rule set.

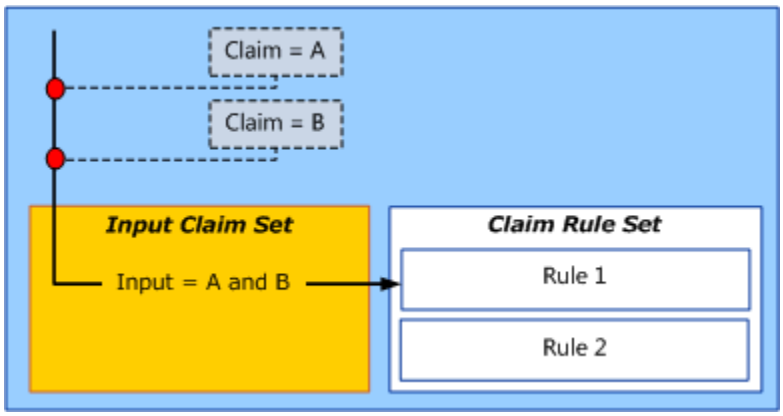

All the rules in a claim rule set share the same input claim set. Each rule in that set can add to the shared input claim set, thus affecting all subsequent rules in the set.

Step 2 – Execution

In this step of the claim rules process, claim rules are processed when the claims engine chronologically steps through all of the rules within a particular rule set one at a time. Each rule within a rule set only runs once and is executed in the order in which they appear from top to bottom as displayed in the Edit Claim Rules dialog box in the AD FS Management snap-in. The claim rule that is at the top of the rule set is processed first and then subsequent rules are processed until all of the rules have been run.

As defined in the claim rule language, a claim rule consists of two parts, condition and issuance statement. The claims engine first processes the condition part by using the data in the input claim set to determine whether the condition specified within the rule holds true for the claims contained in the input claim set (the claims that match the rule's condition are referred to as a matching claims). If any matching claims are found, the claims engine executes the issuance statement of the rule for each set of the matching claims. The issuance statement of the rule can perform either of the following tasks with matching claims:

1. Copy a matching claim into the output claim set
2. Transform the claim fields and create a new claim in either just the input claim set or in both evaluation and output claim sets.
3. Use the matching claim(s) as a key to lookup more information from an attribute store to create new claim(s) in either just the input claim set or in both input and output claim sets.

Adding a claim to the output claim set for a rule set

The *output claim set* is a location in memory that is initially empty and is important because the claims engine will only return claims that reside in the output claim set after the execution process completes. This means that any claims that reside only in the input claim set and not in the output claim set will be ignored when it comes time to calculate the final set of outgoing claims.

Adding a claim to both claim sets for a rule set

As a rule is processed, claims are either added in the input claim set or in both the input claim set and output claim set based on the statement that's used in the rule's issuance statement. The claim rule language refers to these statements as either *add* or *issue*.

If the *add* statement is used, the claims are added to just the input claim set and the claims will exist only for the purposes of the execution and will cease to exist once the execution completes. If the *issue* statement is used, the claims are added to both the input claim set and the output claim set and the claims will be returned in the output claim set once the execution completes. For more information about these statements, see The Role of the Claim Rule Language.

If the condition part of a rule within a rule set does not match any claims in the input claim set, the issuance statement part of the rule is ignored and thus no claims are added to either the output claim set or to the input claim set. The following illustration and corresponding steps show what happens when the claims engine executes a transform rule:

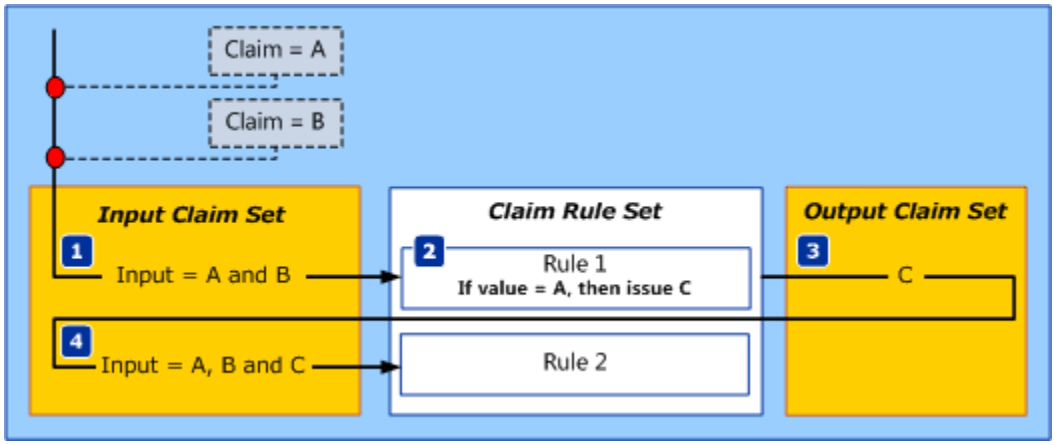

1. Incoming claims are added to the input claim set by the claims engine.
2. When the first rule executes, it sees the A and B claims, which are at that moment in time the only claims in the input claim set, and processes the conditional part of the rule logic in rule 1.
3. Since the A claim is present in the input claim set, the condition of the rule is determined to be true (matching the claim A) and a new C claim is added to both the input claim set and the output claim set.
4. Rule 2 can now use the A, B and C claims (all claims in the input claim set) as input for processing its logic.

For more information about claims transformation, see When to Use a Transform Claim Rule.

Step 3 – Execution Result

The final stage of the claim rule set execution begins once all rules have been run within a given rule set and the final set of claims is present in the output claim set. At this point, the claims engine returns the context of the output claim set as the output of the rule set execution. From this point forward it is the claims pipeline that takes over and moves this final output to the next stage in its process.

Sending the execution output to the claims pipeline

When the claims engine processes a rule set, that rule set has its own dedicated locations in memory for its input and output claim sets. This means that the input and output claim sets used by one rule set are separate from the input and output claim sets used in another rule set.

After the entire process has run for a give rule set (steps 1, 2, and 3), the newly issued outgoing claims (content of the output claim set) will be used as input to the next rule set in the claims pipeline. This allows for claims to flow from the output of one rule set to the input for another rule set, as shown in the following illustration.

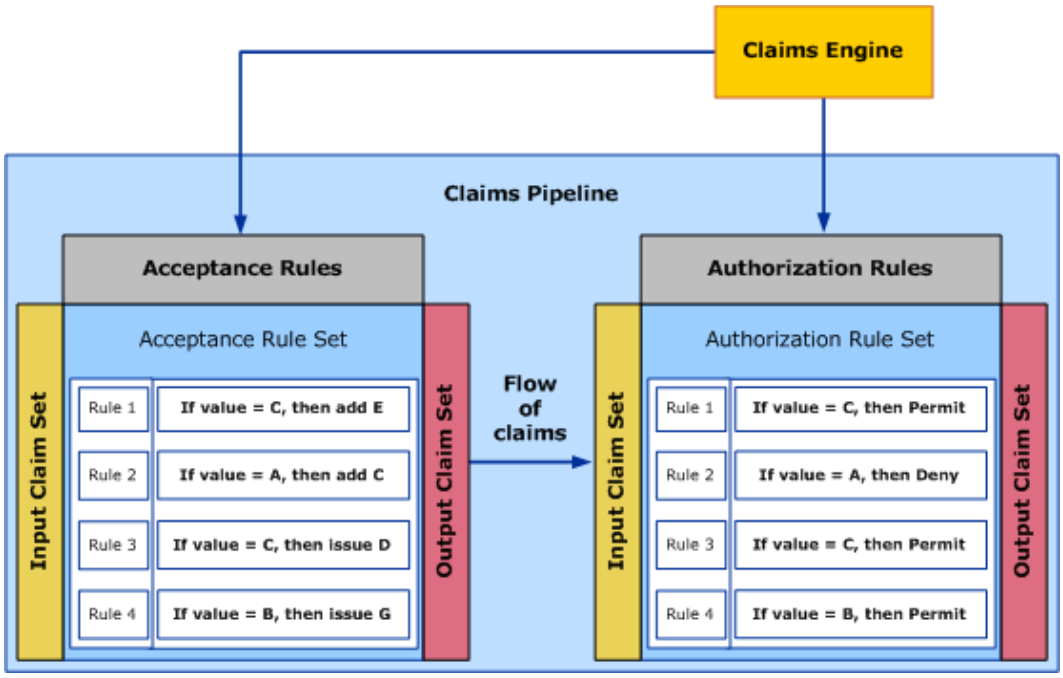

Note

Although the issuance rule set is also a critical stage in the pipeline, the illustration above does not show it only for purposes of simplifying the illustration. For an illustration that shows the issuance rule set and how it fits into the claims pipeline, see The Role of the Claims Pipeline.

In this case, the output of the acceptance rules is used by the pipeline to flow the final set of claims produced by the acceptance rules to the second stage in the pipeline, which is the processing of authorization rules. At this point, the entire claim rules execution process (steps 1, 2, and 3 above) would run again for the authorization rule set. This cycle continues until the issuance rule set (the final stage in the pipeline) has been completed.

Once the finalized outgoing claims have been returned from the engine for the issuance rule set, they will be packaged into a SAML token and the Federation Service will send the token back to the client.

Processing authorization rules

If the claim rule set that is being executed during step 2 of the claim rules execution process consists of Authorization rules (which have a different input and output claim sets than either acceptance or issuance rules), then those authorization rules will run to determine whether a token requester is authorized to obtain a security token for a given relying party from the Federation Service based on the requester's claims.

The goal of authorization rules is to issue a permit or deny claim based on whether the user is to be allowed to obtain a token for the given relying party or not. As shown in the following illustration, the output of the authorization execution is used by the pipeline to determine whether the issuance rule set is executed or not—based on presence or absence of the permit and/or deny claim—but the authorization execution output itself is not used as an input to the claim rule set.

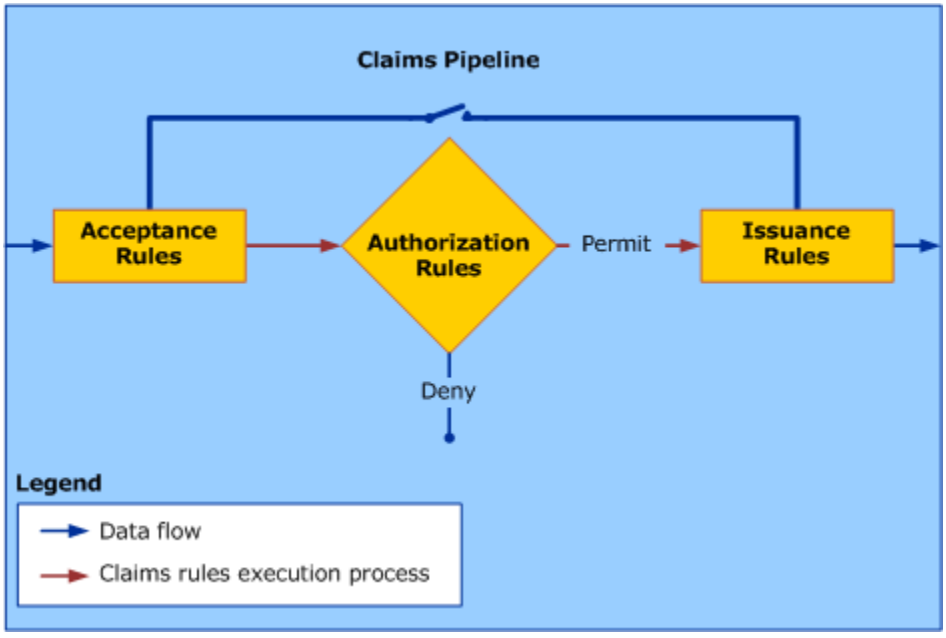

For more information about claims authorization, see When to Use an Authorization Claim Rule.

The Role of the Claims Pipeline

The claims pipeline in Active Directory Federation Services (AD FS) represents the path that claims must follow through the Federation Service before they can be issued. The Federation Service manages the entire end-to-end process of flowing claims through the various stages of the claims pipeline, which also includes the processing of claim rules by the claim rule engine.

For more information about claim rules, see The Role of Claim Rules. For more information about how the claim rule engine processes rules, see The Role of the Claims Engine.

The following section discusses the process that the Federation Service oversees in greater detail.

Claims pipeline process

The claims pipeline process consists of three high-level stages. Each stage in this process initializes the claim rule engine to process claim rules that are specific to that stage. These stages include (in the order that they occur):

1. Accepting incoming claims—This stage in the claims pipeline is used to extract the incoming claims from the token and eliminate claims that are not expected or trusted. After they are extracted, the acceptance rules that make up the acceptance transform rule set for a claims provider trust are run. These rules can be used to pass through or add new claims that can then be used in the subsequent stages of the claims pipeline. The output of this stage is used as an input to second and third stage.
2. Authorizing the claims requester—This stage is used by the claims engine to issue permit or deny claims based on whether the token requester is allowed to obtain a token for the given relying party or not. However, before this can occur the authorization rules that make up either the issuance authorization rule set or the delegation authorization rule set for a relying party trust are ran.
3. Issuing outgoing claims—This stage is used to issue outgoing claims and send them along the pipeline where they will be packaged into a security token. However, before this can occur the issuance rules that make up the issuance transform rule set for a relying party trust are ran, which will determine what claims will be issued as outgoing claims.

All three stages above perform claims rules processing but use a different set of rules. As described above, each stage has an associated set of rules based on either the issuer of the incoming claims (the acceptance rules) or the target service for which the claims/token are being issued (authorization and issuance rules).

Claims are token-agnostic but are transmitted over the network encapsulated in security tokens. The claim rules operate over claims regardless of the format of the incoming or outgoing security token.

Claims rules contain the administrator-defined logic by which the claims engine will accept the incoming claims, authorize claims based on the requester's identity and issue claims that are needed by a relying party. Ultimately, it is the claims engine that determines what claims will go

into the security token that will be issued after the claim has been flowed through the claims pipeline.

As shown in the following illustration, the claims pipeline is responsible for the entire end-to-end process of flowing a claim through the various pipeline stages in order to end up with an issued claim that will be sent over a relying party trust. The outgoing claim in the illustration represents the issued claim.

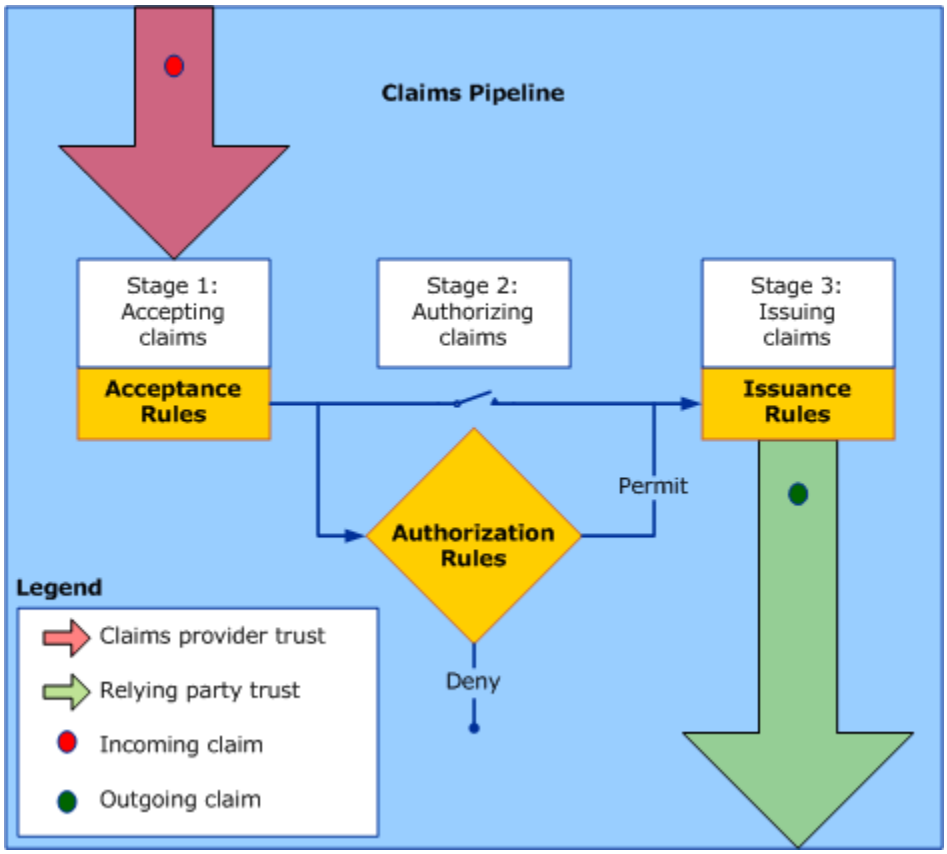

Although it is not shown in the illustration, it is the claims engine that performs the actual processing of the rules at each stage. For more information, see The Role of the Claims Engine.

The Role of the Claim Rule Language

The Active Directory Federation Services (AD FS) claim rule language acts as the administrative building block for the behavior of incoming and outgoing claims, while the claims engine acts as the processing engine for the logic in the claim rule language that defines the custom rule. For more information about how all rules are processed by the claims engine, see The Role of the Claims Engine.

Note

To watch a five-minute video of Stuart Kwan (Principal Group Program Manager for Federated Identity) explaining basic concepts of the claim rule language, click An Introduction to the Claim Rule Language (http://go.microsoft.com/fwlink/?LinkId=163262) to open the video in a new window.

Creating custom claim rules using the claim rule language

AD FS provides administrators with the option to define custom rules that they can use to determine the behavior of identity claims with the claim rule language. You can use the claim rule language syntax examples in this topic to create a custom rule that enumerates, adds, deletes, and modifies claims to meet the needs of your organization. You can build custom rules by typing in the claim rule language syntax in the **Send Claims Using a Custom Claim** rule template.

Rules are separated from each other with semicolons.

For more information about when to use custom rules, see When to Use a Custom Claim Rule.

Using claim rule templates to learn about the claim rule language syntax

AD FS also provides a set of predefined claim issuance and claim acceptance rule templates that you can use to implement common claim rules. In the **Edit Claim Rules** dialog box for a given trust, you can create a predefined rule—and view the claim rule language syntax that makes up that rule—by clicking the **View Rule Language** tab for that rule. Using the information in this section and the **View Rule Language** technique can provide insight into how to construct your own custom rules.

For more detailed information about claim rules and claim rule templates, see The Role of Claim Rules.

Understanding the components of the claim rule language

The claim rule language consists of the following components, separated by the “=>” operator:

- A condition
- An issuance statement

Conditions

You can use conditions in a rule to check input claims and determine whether the issuance statement of the rule should be executed. A condition represents a logical expression that must be evaluated to true to execute the rule body part. If this part is missing, a logical true is assumed; that is, the rule’s body is always executed. The conditions part contains a list of

conditions that are combined together with the conjunction logical operator (“&&”). All conditions in the list must be evaluated to true for the whole conditional part to be evaluated to true. The condition can be either a claims selection operator or an aggregate function call. These two are mutually exclusive, which means that claim selectors and aggregate functions cannot be combined in a single rule conditions part.

Conditions are optional in rules. For example, the following rule does not have a condition:

```
=> issue(type = "http://test/role", value = "employee");
```

There are three types of conditions:

- Single condition—This is the simplest form of a condition. Checks are made for only one expression; for example, windows account name = domain\user.
- Multiple condition—This condition requires additional checks to process multiple expressions in the rule body; for example, windows account name = domain\user and group = contosopurchasers.

Note
Another condition exists, but it is a subset of either the single condition or the multiple condition. It is referred to as a regular expression (Regex) condition. It is used to take an input expression and match the expression with a given pattern. An example of how it is can be used is shown below.

The following examples show a few of the syntax constructions, which are based on the condition types, that you can use to create custom rules.

Single-condition examples

Single-expression conditions are described in the following table. They are constructed to simply check for a claim with a specified claim type or for a claim with a specified claim type and claim value.

Condition description	Condition syntax example
This rule has a condition to check for an input claim with a specified claim type ("http://test/name"). If a matching claim is in the input claims, the rule copies the matching claim or claims to the output claims set.	`c:[type == "http://test/name"]` `  => issue(claim = c);`
This rule has a condition to check for an input claim with a specified claim type ("http://test/name") and claim value (“Terry”). If a matching claim is in the input claims, the rule copies the matching claim or claims to the output claims set.	`c:[type == "http://test/name", value == "Terry"]` `  => issue(claim = c);`

More-complex conditions are shown in the next section, including conditions to check for multiple claims, conditions to check the issuer of a claim, and conditions to check for values that match a regular expression pattern.

Multiple-condition examples

The following table provides an example of multiple-expression conditions.

Condition description	Condition syntax example

This rule has a condition to check for two input claims, each with a specified claim type ("http://test/name" and "http://test/email"). If the two matching claims are in the input claims, the rule copies the name claim to the output claims set.	`c1:[type == "http://test/name"] &&`
	`c2:[type == "http://test/email"]`
	`=> issue(claim = c1);`

Regular-condition examples

The following table provides an example of a regular, expression-based condition.

Condition description	Condition syntax example

This rule has a condition that uses a regular expression to check for an e-mail claim ending in “@fabrikam.com”. If a matching claim is found in the input claims, the rule copies the matching claim to the output claims set.	`c:[type == "http://test/email", value =~`
	`"^.+@fabrikam.com$"]`
	`=> issue(claim = c);`

Issuance statements

Custom rules are processed based on the issuance statements (*issue* or *add*) that you program into the claim rule. Depending on the desired outcome, either the issue statement or add statement can be written into the rule to populate the input claim set or the output claim set. A custom rule that uses the add statement explicitly populates claim values only to the input claim set while a custom claim rule that uses the issue statement populates claim values in both the input claim set and in the output claim set. This can be useful when a claim value is intended to be used only by future rules in the set of claim rules.

For example, in the following illustration, the incoming claim is added to the input claim set by the claims issuance engine. When the first custom claim rule executes and the criteria of domain\user is satisfied, the claims issuance engine processes the logic in the rule using the add statement,

and the value of **Editor** is added to the input claim set. Because the value of Editor is present in the input claim set, Rule 2 can successfully process the issue statement in its logic and generate a new value of **Hello**, which is added to both the output claim set and to the input claim set for use by the next rule in the rule set. Rule 3 can now use all of the values that are present in the input claim set as input for processing its logic.

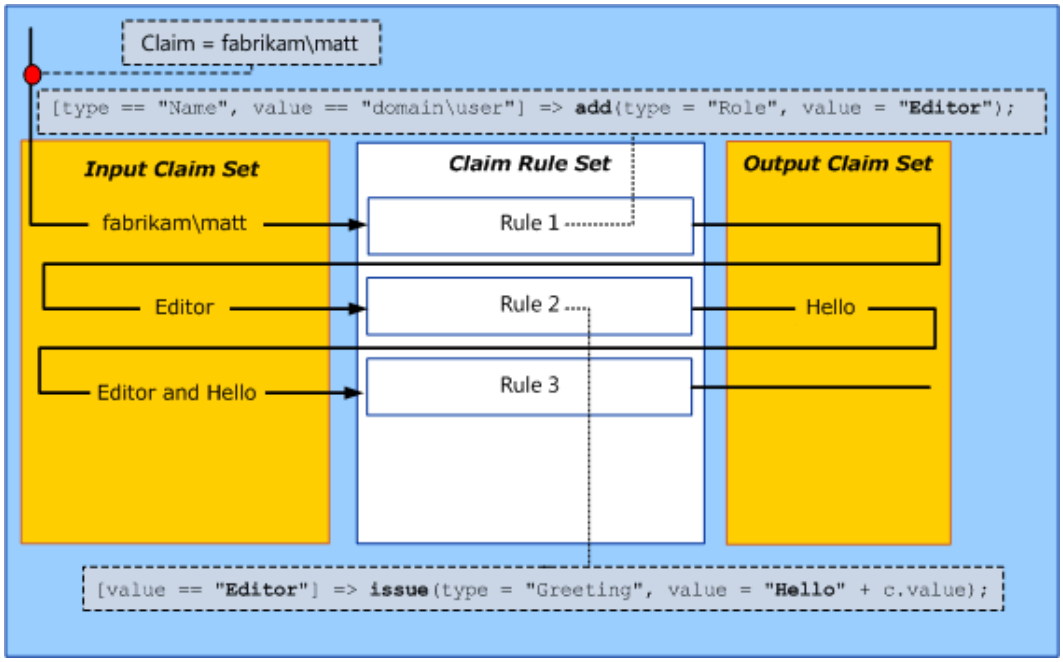

Claim issuance actions

The rule body represents a claim issuance action. There are two claim issuance actions that the language recognizes:

- **Issue statement:** The issue statement creates a claim that goes to both input and output claim sets. For example, the following statement issues a new claim based on its input claim set:

```
c:[type == "Name"]          => issue(type = "Greeting", value = "Hello " +
c.value);
```

- **Add statement:** The add statement creates a new claim that is added only to the input claim set collection. For example, the following statement adds a new claim to the input claim set:

```
[type == "Name", value == "domain\user"]          => add(type = "Role", value =
"Editor");
```

The issuance statement of a rule defines what claims will be issued by the rule when the conditions are matched. There are two forms of issuance statements regarding arguments and the statement behavior:

- **Normal**—Normal issuance statements can issue claims by using literal values in the rule or the values from claims that match the conditions. A normal issuance statement can consist of one or both of the following formats:

- *Claim copy*: The claim copy creates a copy of the existing claim in the output claim set. This issuance form only makes sense when it is combined with the “issue” issuance statement. When it is combined with the “add” issuance statement, it does not have any effect.
- *New claim*: This format reates a new claim, given the values for various claim properties. Claim.Type must be specified; all other claim properties are optional. The order of arguments for this form is ignored.

- **Attribute Store**—This form creates claims with values that are retrieved from an attribute store. It is possible to create multiple claim types by using a single issuance statement, which is important for attribute stores that make network or disk input/output (I/O) operations during the attribute retrieval. Therefore, it is desirable to limit the number of round trips between the policy engine and the attribute store. It is also legal to create multiple claims for a given claim type. When the attribute store returns multiple values for a given claim type, the issuance statement automatically creates a claim for each returned claim value. An attribute store implementation uses the param arguments to substitute the placeholders in the query argument with values that are provided in param arguments. The placeholders use the same syntax as the .NET String.Format() function (for example, {1}, {2}, and so on). The order of the arguments for this form of issuance is important, and it must be the order which is prescribed in the following grammar.

The following table describes some common syntax constructions for both types of issuance statements in claim rules.

Issuance statement type	Issuance statement description	Issuance statement syntax example

Normal	The following rule always emits the same claim whenever a user has the specified claim type and value:	`c:[type == "http://test/employee", value == "true"]`
		
		`=> issue(type = "http://test/role", value = "employee");`

Normal	The following rule converts one claim type into another. Notice that the value of the claim that matches the condition "c" is used in the issuance statement.	`c:[type == "http://test/group"]`
		
		`=> issue(type = "http://test/role", value = c.Value);`

Attribute store	The following rule uses the value of an incoming claim to	`c:[Type == "http://test/name"]`

Issuance statement type	Issuance statement description	Issuance statement syntax example

	query the Active Directory attribute store:	`=> issue(store = "Enterprise AD Attribute Store",`
		
		`types = ("http://test/email"),`
		
		`query = ";mail;{0}",`
		
		`param = c.Value)`

Attribute store	The following rule uses the value of an incoming claim to query a previously configured Structured Query Language (SQL) attribute store:	`c:[type == "http://test/name"]`
		
		`=> issue(store = "Custom SQL store",`
		
		`types = ("http://test/email","http://test/displayname"),`
		
		`query = "SELECT mail, displayname FROM users WHERE name={0}",`
		
		`param = c.value);`

Expressions

Expressions are used on the right side for both claims selector constraints and issuance statement parameters. There are various kinds of expressions that the language supports. All expressions in the language are string based, which means that they take strings as input and produce strings. Numbers or other data types, such as date/time, in expressions are not supported. The following are the types of expressions that the language supports:

- String literal: String value, delimited by the quote (“) character on both sides.

- String concatenation of expressions: The result is a string that is produced by concatenation of the left and right values.
- Function call: The function is identified by an identifier, and the parameters are passed as a comma-delimited list of expressions enclosed in brackets (“()”).
- Claim’s property access in the form of a variable name DOT property name: The result of the value of the identified claim’s property for a given variable valuation. The variable must first be bound to a claims selector before it can be used in this way. It is illegal to use the variable that is bound to a claims selector inside the constraints for that same claims selector.

The following claim properties are available for access:

- Claim.Type
- Claim.Value
- Claim.Issuer
- Claim.OriginalIssuer
- Claim.ValueType
- Claim.Properties[property_name] (This property returns an empty string if the property_name cannot be found in the claim’s properties collection.)

You can use the RegexReplace function to call inside an expression. This function takes an input expression and matches it with the given pattern. If the pattern matches, the output of the match is replaced with the replacement value.

Exists functions

The Exists function can be used in a condition to evaluate whether a claim that matches the condition exists in the input claim set. If any matching claim exists, the issuance statement is called only once. In the following example, the “origin” claim is issued exactly once—if there is at least one claim in the input claim set collection that has the issuer set to “MSFT”, no matter how many claims have the issuer set to “MSFT”. Using this function prevents duplicate claims from being issued.

```
exists([issuer == "MSFT"])
   => issue(type = "origin", value = "Microsoft");
```

Rule body

The rule body can contain only a single issuance statement. If conditions are used without using the Exists function, the rule body is executed once for each time that the conditions part is matched.

Additional references

Create a Rule to Send Claims Using a Custom Rule

How URIs Are Used in AD FS

A Uniform Resource Identifier (URI) is a string of characters that is used as a unique identifier. In AD FS, URIs are used to identify both partner network addresses and configuration objects. When used to identify partner network addresses, the URI is always a URL. When used to identify configuration objects, the URI may be a URN or a URL. For more general information about URIs, see RFC 2396 and RFC 3986.

URIs as partner network addresses

The following are the network address URLs that are most often handled by administrators in AD FS.

- The URLs of the Federation Service, including WS-Federation, SAML, WS-Trust, Federation Metadata, WS-MetadataExchange, Privacy and Organization URLs
- The URLs of a relying party trust, including WS-Federation, SAML, and Federation Metadata URLs
- The URLs of a claims provider trust, including WS-Federation, SAML, and Federation Metadata URLs

URIs as object identifiers

The following table describes the identifiers that are most often handled by administrators in AD FS.

Identifier name	Description	Comparisons
Federation Service identifier	This identifier is used to identify the Federation Service. It is used by relying parties that use claims from this Federation Service, as well as claims providers that issue claims to this Federation Service.	When a user requests claims from a claims provider for this Federation Service, the Federation Service identifier will be used to identify the target for the claims. When this Federation Service receives the claims from a claims provider, it will check to ensure the claims are scoped for it by looking for its Federation Service identifier. When a relying party is receiving claims from this Federation Service, the relying party will check that the issuer

Identifier name	Description	Comparisons
		of the claims matches the Federation Service identifier.
Relying party identifier	This identifier is used to identify the relying party to this Federation Service. It is used when issuing claims to the relying party.	When a user requests claims from this Federation Service for the relying party, the relying party identifier will be used to identify the relying party for which the claims should be targeted. This comparison is done using prefix matching (see below). When the relying party receives the claims, it will check for its identifier in the security token to ensure the claims are targeted for it.
Claims provider identifier	This identifier is used to identify the claims provider to this Federation Service. It is used when receiving claims from the claims provider.	When this Federation Service is receiving claims from the claims provider, this Federation Service will check that the issuer of the claims matches the claims provider identifier.
Claim type	This identifier is used to define the type of claim. It is used by this Federation Service, claims providers, and relying parties when sending and receiving claims.	When the Federation Service receives claims from a claims provider, the claim rules associated with the corresponding claims provider trust allow the administrator to compare claim types and process claims. The claim rules associated with a relying party trust also allow the administrator to compare claim types from the claims coming out of the claims provider trust rules, and decide which claims to issue.

URI prefix matching for relying party identifiers

The path syntax of a URI is organized hierarchically and is delimited by either all “/” characters or all “:”characters. Thus the path may be split into path sections based on the delimiting character. When prefix matching, each section must be a full match according to the matching rules (these rules govern the casing of matches). For more information about matching rules, see the RFC’s mentioned above.

When a relying party is identified in a request to the Federation Service, AD FS uses prefix matching logic to determine if there is a matching relying party trust in the AD FS configuration database.

For example, if the relying party identifier in the AD FS configuration database (URI1) is a prefix to the relying party identifier in the incoming request (URI2), then the following must be true:

- Trailing delimiters (slashes and colons) of path sections or authorities must be ignored
- The scheme and authority parts of URI1 and URI2 must be a case insensitive exact match
- Each path section of URI1 must be an exact match (based on the case sensitivity chosen) to the corresponding path section of URI2
- URI2 may have more path sections than URI1, but URI1 must not have more path sections than URI2
- URI1 cannot have more path sections than URI2
- If URI1 has a query string, it must match exactly to a URI2 query string
- If URI1 has a fragment, it must match exactly to a URI2 fragment

The following table provides additional examples.

Relying party identifier in AD FS configuration database	Relying party identifier in request message	Request identifier matches the configuration identifier?	Reason
http://contoso.com	http://contoso.com	TRUE	Exact match
http://contoso.com/	http://contoso.com	TRUE	Trailing slashes are ignored
http://contoso.com	http://contoso.com/	TRUE	Trailing slashes are ignored
http://contoso.com	http://contoso.com/hr	TRUE	URI1 has no path and matches scheme and authority to URI2

Relying party identifier in AD FS configuration database	Relying party identifier in request message	Request identifier matches the configuration identifier?	Reason
http://contoso.com/hr	http://contoso.com/hr/web	TRUE	First path sections match, URI1 has no second path section
http://contoso.com/hr	http://contoso.com/hr/web/?m=t	TRUE	Same reasons as above, query string doesn't change anything
http://contoso.com/hr/	http://contoso.com/hrw/main	FALSE	URI1 path section 1 does not match URI2 path section 1
http://contoso.com/hr	http://contoso.com	FALSE	URI1 has more path sections than URI2
http://contoso.com/hr	http://contoso.com/hrweb	FALSE	First path sections do not match
http://contoso.com/?m=t	http://contoso.com/?m=f	FALSE	Query string parts do not match
https://contoso.com	http://contoso.com	FALSE	Scheme parts do not match
http://sts.contoso.com	http://contoso.com	FALSE	Authority parts do not match

Relying party identifier in AD FS configuration database	Relying party identifier in request message	Request identifier matches the configuration identifier?	Reason
http://contoso.com	http://sts.contoso.com	FALSE	Authority parts do not match

Determine the Type of Claim Rule Template to Use

An important part of designing an Active Directory Federation Services (AD FS) infrastructure is determining the complete set of claim rules—and which corresponding claim rule templates you should use to create them—for each partner that will participate in federation with your organization. You create rules by using claim rule templates in the AD FS Management snap-in.

Each set of claim rules you configure can only be associated with one federated trust. This means that you cannot create a set of rules on one trust and use them for other trusts in your Federation Service. Instead you can easily create rules from claim rule templates to more quickly help produce a desired set of claims that are agreed upon between each federated partner and your organization.

For more information about rules and rule templates, see The Role of Claim Rules.

Before you begin to determine the types of claim rule templates you should use, consider the following questions:

- What claims will be provided by your trusted claims providers?
- What claims do you trust from each claims provider?
- What claims are required by the relying parties that trust this Federation Service?
- What claims are you willing to divulge to each relying party?
- Which users should have access to each relying party?

Answering these questions will help you plan a solid claim rule design. It will also assist you in creating a smooth authorization and access control strategy and make your deployment team more efficient during the rollout.

In this next section you can learn about the type of rule templates to select for your environment based on your business needs.

Claim rule template types

The following table describes all of the types of claim rule templates that you can use to create rules using the AD FS Management snap-in, and the benefits of using one template type over another.

Rule template type	Description	Advantages	Disadvantages
Pass Through or Filter an Incoming Claim	Used to create a rule that will pass through all claim values for a selected claim type or filter claims based on the claim values so that only certain claim values for a selected claim type will pass through. For more information, see When to Use a Pass Through or Filter Claim Rule.	• Can be used to select particular claims to be accepted or issued unchanged	• Claim type and value cannot be changed
Transform an Incoming Claim	Used to create a rule that can select an incoming claim and map it to a different claim type or map its claim value to a new claim value. For more information, see When to Use a Transform Claim Rule.	• Can be used to normalize claim types or values • Can replace an e-mail suffix of an incoming claim	• More complex string replacements require a custom rule
Send LDAP Attributes as Claims	Used to create a rule that will select attributes from an LDAP attribute store to send as claims to the relying party. For more information, see When to Use a Send LDAP Attributes	• Can source claims from any AD DS/AD LDS attribute store • Multiple claims can be issued using a single rule	• Performance – slow as a result of account lookup • Cannot use a custom LDAP filter for querying

Rule template type	Description	Advantages	Disadvantages
	as Claims Rule.		
Send Group Membership as a Claim	Used to create a rule that can send a specified claim type and value when a user is a member of an Active Directory security group. Only a single claim will be sent using this rule, based on the group that you select. For more information, see When to Use a Send Group Membership as a Claim Rule.	• Fast performance for issuing group claims – no account lookup	• User must be a member of a local Active Directory group
Send Claims Using a Custom Rule	Used to create a custom rule that will provide more advanced options than a standard rule template. You write custom rules with the AD FS claim rule language. For more information, see When to Use a Custom Claim Rule.	• Can be used to source claims from an SQL attribute store • Can be used to specify a custom LDAP filter • Can be used to issue a PPID • Can be used with a custom attribute store • Can be used to add claims only to the input claim set • Can be used to send claims based on more than one incoming claim	• More difficult to configure - Some ramp up time may be needed to initially gain knowledge of the claim rule language
Permit or Deny Users Based on an Incoming Claim	Used to create a rule that will permit or deny access by users to the relying party,	• Simplifies the authorization process	• Requires that only one claim type and one claim value be specified

Rule template type	Description	Advantages	Disadvantages
	based on the type and value of an incoming claim. For more information, see When to Use an Authorization Claim Rule.		• Does not support pattern matching for claim values
Permit All Users	Used to create a rule that will permit all users to access the relying party. For more information, see When to Use an Authorization Claim Rule.	• Simple to configure	• Less secure than using the Permit or Deny Users Based on an Incoming Claim template

When to Use an Authorization Claim Rule

You can use this rule in Active Directory Federation Services (AD FS) when you need to take an incoming claim type and then apply an action that will determine whether a user will be permitted or denied access based on the value that you specify in the rule. When you use this rule, you pass through or transform claims that match the following rule logic, based on either of the options you configure in the rule:

Rule option	Rule logic
Permit all users	If incoming claim type equals *any claim type* and value equals *any value*, then issue claim with value equals *Permit*
Permit access to users with this incoming claim	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then issue claim with value equals *Permit*
Deny access to users with this incoming claim	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then issue claim with value equals *Deny*

The following sections provide a basic introduction to claim rules and provide further details about when to use this rule.

About claim rules

A claim rule represents an instance of business logic that will take an incoming claim, apply a condition to it (if x then y) and produce an outgoing claim based on the condition parameters. The following list outlines important tips that you should know about claim rules before you read further in this topic:

- In the AD FS Management snap-in, claim rules can only be created using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on rules, you can further refine or filter claims that are generated by previous rules within a given rule set.
- Claim rule templates will always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

Permit All Users

When you use the Permit All Users rule template, all users will have access to the relying party. However, you can use additional authorization rules to further restrict access. If one rule permits a user to access the relying party, and another rule denies the user access to the relying party, the deny result overrides the permit result and the user is denied access.

Users who are permitted access to the relying party from the Federation Service may still be denied service by the relying party.

Permit access to users with this incoming claim

When you use the Permit or Deny Users Based on an Incoming Claim rule template to create a rule and set the condition to permit, you can permit specific user's access to the relying party based on the type and value of an incoming claim. For example, you can use this rule template to create a rule that will permit only those users that have a group claim with a value of Domain Admins. If one rule permits a user to access the relying party, and another rule denies the user access to the relying party, the deny result overrides the permit result and the user is denied access.

Users who are permitted to access the relying party from the Federation Service may still be denied service by the relying party. If you want to permit all users to access the relying party, use the Permit All Users rule template.

Deny access to users with this incoming claim

When you use the Permit or Deny Users Based on an Incoming Claim rule template to create a rule and set the condition to deny, you can deny user's access to the relying party based on the type and value of an incoming claim. For example, you can use this rule template to create a rule that will deny all users that have a group claim with a value of Domain Users.

If you want to use the deny condition yet also enable access to the relying party for specific users, you must later explicitly add authorization rules with the permit condition to enable those users access to the relying party.

If a user is denied access when the claims issuance engine processes the rule set, further rules processing shuts down, and AD FS returns an "Access denied" error to the user's request.

Authorizing users

In AD FS, authorization rules are used to issue a permit or deny claim that will determine whether a user or a group of users (depending on the claim type used) will be allowed to access Web-based resources in a given relying party or not. Authorization rules can only be set on relying party trusts.

Authorization rule sets

Different authorization rule sets exist depending on the type of permit or deny operations you need to configure. These rule sets include:

- **Issuance Authorization Rules**: These rules determine whether a user can receive claims for a relying party and, therefore, access to the relying party.
- **Delegation Authorization Rules**: These rules determine whether a user can act as another user to the relying party. When a user is acting as another user, claims about the requesting user are still placed in the token.
- **Impersonation Authorization Rules**: These rules determine whether a user can fully impersonate another user to the relying party. Impersonating another user is a very powerful capability, because the relying party will not know that the user is being impersonated.

For more details about how the authorization rule process fits into the claims issuance pipeline, see The Role of the Claims Issuance Engine.

Supported claim types

AD FSdefines two claim types that are used to determine whether a user is permitted or denied. These claim type Uniform Resource Identifiers (URIs) are as follows:

1. **Permit**: http://schemas.microsoft.com/authorization/claims/permit
2. **Deny**: http://schemas.microsoft.com/authorization/claims/deny

How to create this rule

You can create both authorization rules using either the claim rule language or using the **Permit All Users** rule template or the **Permit or Deny Users Based on an Incoming Claim** rule template in the AD FS Management snap-in. The Permit All Users rule template does not provide any configuration options. However, the Permit or Deny Users Based on an Incoming Claim rule template provides the following configuration options:

- Specify a claim rule name
- Specify an incoming claim type
- Type an incoming claim value
- Permit access to users with this incoming claim
- Deny access to users with this incoming claim

For more instructions on how to create this template, see Create a Rule to Permit All Users or Create a Rule to Permit or Deny Users Based on an Incoming Claim in the AD FS Deployment Guide.

Using the claim rule language

If a claim should be sent only when the claim value matches a custom pattern, you must use a custom rule. For more information, see When to Use a Custom Claim Rule.

Example of how to create an authorization rule based on multiple claims

When using the claim rule language syntax to authorize claims, a claim can also be issued based on the presence of multiple claims in the user's original claims. The following rule issues an authorization claim only if the user is a member of the group Editors and has authenticated using Windows authentication:

```
[type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/authenticationmethod",
value == "urn:federation:authentication:windows" ]
&& [type == "http://schemas.xmlsoap.org/claims/Group ", value == "editors"]
=> issue(type = "http://schemas.xmlsoap.org/claims/authZ", value = "Granted");
```

Example of how to create authorization rules that will delegate who can create or remove federation server proxy trusts

Before a Federation Service can use a federation server proxy to redirect client requests, a trust must first be established between the Federation Service and the federation server proxy computer. By default, a proxy trust is established when either of the following credentials is provided successfully in the AD FS Federation Server Proxy Configuration Wizard:

- The service account, used by the Federation Service, that the proxy will protect

- An Active Directory domain account that is a member of the local Administrators group on all federation servers in a federation server farm

When you want to specify which user or users can create a proxy trust for a given Federation Service, you can use any of the following delegation methods. This list of methods is in priority order, based on the AD FS product team's recommendations of the most secure and least problematic methods of delegation. It is necessary to use only one of these methods, depending on the needs of your organization:

1. Create a domain security group in Active Directory (for example, FSProxyTrustCreators), add this group to the local Administrators group on each of the federation servers in the farm, and then add only the user accounts to which you want to delegate this right to the new group. This is the preferred method.
2. Add the user's domain account to the administrator group on each of the federation servers in the farm.
3. If for some reason you cannot use either of these methods, you can also create an authorization rule for this purpose. Although it is not recommended—because of possible complications that may occur if this rule is not written correctly—you can use a custom authorization rule to delegate which Active Directory domain user accounts can also create or even remove the trusts between all the federation server proxies that are associated with a given Federation Service.

   If you choose method 3, you can use the following rule syntax to issue an authorization claim that will allow a specified user (in this case, contoso\frankm) to create trusts for one or more federation server proxies to the Federation Service. You must apply this rule using the Windows PowerShell command **Set-ADFSProperties AddProxyAuthorizationRules**.

```
c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windows
accountname", issuer=~"^AD AUTHORITY$" value == "contoso\frankm"
] => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "true")


exists([Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsi
d", Value == "S-1-5-32-544", Issuer =~ "^AD AUTHORITY$"])
=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "true");


c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/primary
sid", Issuer =~ "^AD AUTHORITY$" ] =>
issue(store="_ProxyCredentialStore",types=("http://schemas.micro
```

```
soft.com/authorization/claims/permit"),query="isProxyTrustManage
rSid({0})", param= c.Value );


c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/proxytr
ustid", Issuer =~ "^SELF AUTHORITY$" ] =>
issue(store="_ProxyCredentialStore",types=("http://schemas.micro
soft.com/authorization/claims/permit"),query="isProxyTrustProvis
ioned({0})", param=c.Value );
```

Later, if you want to remove the user so that the user can no longer create proxy trusts, you can revert to the default proxy trust authorization rule to remove the right for the user to create proxy trusts for the Federation Service. You must also apply this rule using the Windows PowerShell command **Set-ADFSProperties AddProxyAuthorizationRules**.

```
exists([Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsi
d", Value == "S-1-5-32-544", Issuer =~ "^AD AUTHORITY$"])
=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit",
Value = "true");


c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/primary
sid", Issuer =~ "^AD AUTHORITY$" ] =>
issue(store="_ProxyCredentialStore",types=("http://schemas.micro
soft.com/authorization/claims/permit"),query="isProxyTrustManage
rSid({0})", param= c.Value );


c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/proxytr
ustid", Issuer =~ "^SELF AUTHORITY$" ] =>
issue(store="_ProxyCredentialStore",types=("http://schemas.micro
soft.com/authorization/claims/permit"),query="isProxyTrustProvis
ioned({0})", param=c.Value );
```

For more information about how to use the claim rule language, see [The Role of the Claim Rule Language](#).

When to Use a Pass Through or Filter Claim Rule

You can use this rule in Active Directory Federation Services (AD FS) when you need to take a specific incoming claim type and then apply an action that will determine what output should occur based on the values in the incoming claim. When you use this rule, you pass through or filter any claims that match the rule logic in the following table, based on either of the options you configure in the rule.

Rule option	Rule logic
Pass through all claim values	If incoming claim type equals *specified claim type* and value equals *any value*, then pass the claim through
Pass through only a specific claim value	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then pass the claim through
Pass through only claim values that match a specific e-mail suffix value	If incoming claim type equals *specified claim type* and value equals *specified suffix value*, then pass the claim through
Pass through only claim values that start with a specific value	If incoming claim type equals *specified claim type* and value begins with *specified claim value*, then pass the claim through

The following sections provide a basic introduction to claim rules and provide further details about when to use this rule.

About claim rules

A claim rule represents an instance of business logic that will take an incoming claim, apply a condition to it (if x then y) and produce an outgoing claim based on the condition parameters. The following list outlines important tips that you should know about claim rules before you read further in this topic:

- In the AD FS Management snap-in, claim rules can only be created using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on rules, you can further refine or filter claims that are generated by previous rules within a given rule set.

- Claim rule templates will always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

Pass through all claim values

When using this action, all incoming claim values for the specified claim type are passed through as outgoing claims. For example, when the incoming claim type is specified as the Role claim type, all incoming claim values are copied individually into new outgoing claims with the outgoing claim type of Role.

Filtering a claim

In AD FS, the term *claims filtering* means to filter or restrict incoming claim values so that only certain values are passed or sent through as outgoing claims. It is the **Pass Through or Filter an Incoming Claim** rule template that makes this function possible. Within the properties of this rule, you can set conditions to filter incoming values so that only the values that meet your specified criteria are passed through.

For example, you can use this rule to only pass through claims that match the claim value of Purchaser when the incoming claim type matches the claim type of Role or you might want to issue only claims about the name of the user, but not claims containing the social security number of the user.

When you use a filter condition with this rule, all incoming claims are examined to determine which claims match the criteria set by the rule. All other claims are ignored so that only specified claim values that match a selected claim type will pass through.

For example, as shown in the following illustration, when a rule is set with the condition to filter only incoming claims that are keyed to the UPN claim type and also end with @fabrikam.com, all other incoming claims are ignored unless they meet this criteria. This includes the incoming claim with the claim type of E-Mail Address even though its claim value ends in @fabrikam.com. In this case, only the claim containing the value of Nick@fabrikam.com is sent to the relying party.

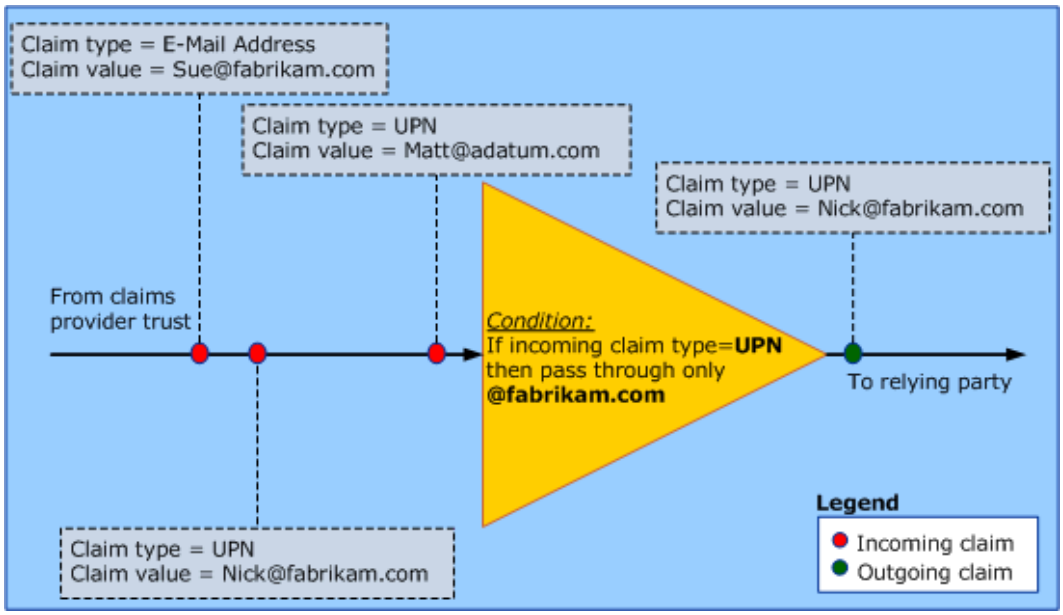

Configuring this rule on a claims provider trust

When you use a claims provider trust, this rule can be configured to pass through only incoming claims from the claims provider that match certain constraints. For example, you might want to only accept e-mail claims from the claims provider; therefore, you would use this rule template to accept e-mail claim types that end in the claims provider's Domain Name System (DNS) name.

Configuring this rule on a relying party trust

When you use a relying party trust, this rule can be configured to pass through or filter outgoing claims that will be sent to the relying party. Some relying parties might not understand certain claim types, or certain claims might contain sensitive information that should not be sent to certain relying parties. This rule template can help to enforce those policies for a particular relying party trust.

How to create this rule

You create this rule using either the claim rule language or using the Pass Through or Filter an Incoming Claim rule template in the AD FS Management snap-in. This rule template provides the following configuration options:

- Specify a claim rule name
- Specify a incoming claim type
- Pass through all claim values
- Pass through only a specific claim value
- Pass through only claim values that match a specific e-mail suffix value

- Pass through only claim values that start with a specific value

For more instructions on how to create this template, see Create a Rule to Pass Through or Filter an Incoming Claim in the AD FS Deployment Guide.

Using the claim rule language

If a claim should be sent only when the claim value matches a custom pattern, you must use a custom rule. For more information, see When to Use a Custom Rule.

Examples of how to construct a pass through or filter rule syntax

A simple filtering rule would filter claims based on one of the properties outlined above. For example, the following rule will pass through all e-mail claims:

```
c:[type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"]  =>
issue(claim  = c);
```

Filters can be logically AND-ed together. For example, the following rule will accept all e-mail claims with value johndoe@fabrikam.com:

```
c:[type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress", value ==
"johndoe@fabrikam.com "]  => issue(claim  = c);
```

In the above examples the filters always used an equality operator. The claim rule language supports the following operators:

- == - equals (case-sensitive)
- != - not equals (case-sensitive)
- =~- regular expression match
- !~ - regular expression non-match

For example, the following rule will accept all e-mail claims not issued by the local federation server that have a suffix of boeing.com:

```
c:[type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress", value =~
"^.*@boeing\.com$" , issuer != "LOCAL AUTHORITY"]  => issue(claim  = c);
```

Best practices for creating custom rules

A filter can be applied to one or more of the properties of each claim, as described in the following table.

Claim property	Description
Type	The claim type (usually represented as an Uri) reflects an implicit agreement between partners in a Federation about what kind of information is conveyed in the claim. For example, claims of type http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress

Claim property	Description
	will contain the e-mail address of the user.
Value	The value of the claim. For example, a claim of type http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress may have a value of johndoe@fabrikam.com
ValueType	The ValueType represents how the information contained in the claim's Value is to be interpreted. Typically the ValueType will be set to http://www.w3.org/2001/XMLSchema#string, but the claim value could contain Base64Binary encoded data (for example, an image) or a date, Boolean, and so on.
Issuer	The issuer represents the party that last issued the claims about the user. If the claims are obtained at a claims provider federation server the issuer of all claims is going to be set to "LOCAL AUTHORITY". If the claims were received by a Federation Provider federation server, the issuer of the claims is going to be set to the claims provider identifier of the claims provider that signed the token. Thus, when processing rules on claims received from a claims provider the issuer of all claims is going to be set to the same value. When authoring rules for a relying party, the issuer property can be used to distinguish between claims originating from different claims providers.
OriginalIssuer	This claim property is meant to convey which federation server originally issued the claim. Since the issuer property of claims is set to the last federation server that signed the token, the original issuer is useful in scenarios where a claim has flowed through more than one federation server (For example, a relying party that receives a token from a federation provider federation server might be interested which particular claims provider federation server authenticated the user)
Properties	In addition to the five properties outlined above, each claim also has a property bag where named properties can be stored. These properties are not serialized in the token and only make sense for passing information between components of the claims issuance pipeline within the scope of a single federation server. For example, setting a property during claims provider rules processing and then referring to it in relying party rules.

When to Use a Transform Claim Rule

You can use this rule in Active Directory Federation Services (AD FS) when you need to map an incoming claim type to an outgoing claim type and then apply an action that will determine what output should occur based on the values that originated in the incoming claim. When you use this rule, you pass through or transform claims that match the following rule logic, based on either of the options that you configure in the rule, as described in the following table.

Rule option	Rule logic
Pass through all incoming claims	If incoming claim type equals *specified claim type* and value equals *any value*, then pass the claim through with outgoing claim type equals *specified claim type*
Replace an incoming claim value with a different outgoing claim value	If incoming claim type equals *specified claim type* and value equals *specified claim value*, then transform the claim with new outgoing claim value *specified claim value* and with outgoing claim type *specified claim type*
Replacing incoming e-mail suffix claims with a new e-mail suffix	If incoming claim type equals *specified claim type* and value equals *any suffix value*, then transform the claim with new outgoing claim value *specified suffix value* and with outgoing claim type *specified claim type*

The following sections provide a basic introduction to claim rules and provide further details about when to use this rule.

About claim rules

A claim rule represents an instance of business logic that will take an incoming claim, apply a condition to it (if x then y) and produce an outgoing claim based on the condition parameters. The following list outlines important tips that you should know about claim rules before you read further in this topic:

- In the AD FS Management snap-in, claim rules can only be created using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on rules, you can further refine or filter claims that are generated by previous rules within a given rule set.

- Claim rule templates will always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

Pass through all claim values

When using this action, all incoming claim values that are keyed to a specified incoming claim type are mapped to a specified outgoing claim type before they are sent as outgoing claims into tokens that are signed by your Federation Service.

For example, when a rule is set with the **Pass through all claim values** option logic and the incoming claim type of Group and the outgoing claim type of Role is specified, all incoming claim values that flow in from the issuer are copied individually into new outgoing claims with claim type of Role.

Transforming a claim

In AD FS, the term *claims transformation* means to replace one incoming claim value with a different outgoing claim value. It is the Transform an Incoming Claim rule that makes this function possible. Within the properties of this rule, you can set conditions to transform incoming values with a different outgoing claim value based on the specified incoming claim type.

For example, as shown in the following illustration, when a rule is set with the condition to replace an incoming value with a different outgoing claim value, all incoming claim types of Group are mapped to new outgoing claim types of Role. In this case, the incoming claim value of Purchaser is replaced with the new outgoing claim value of Admin.

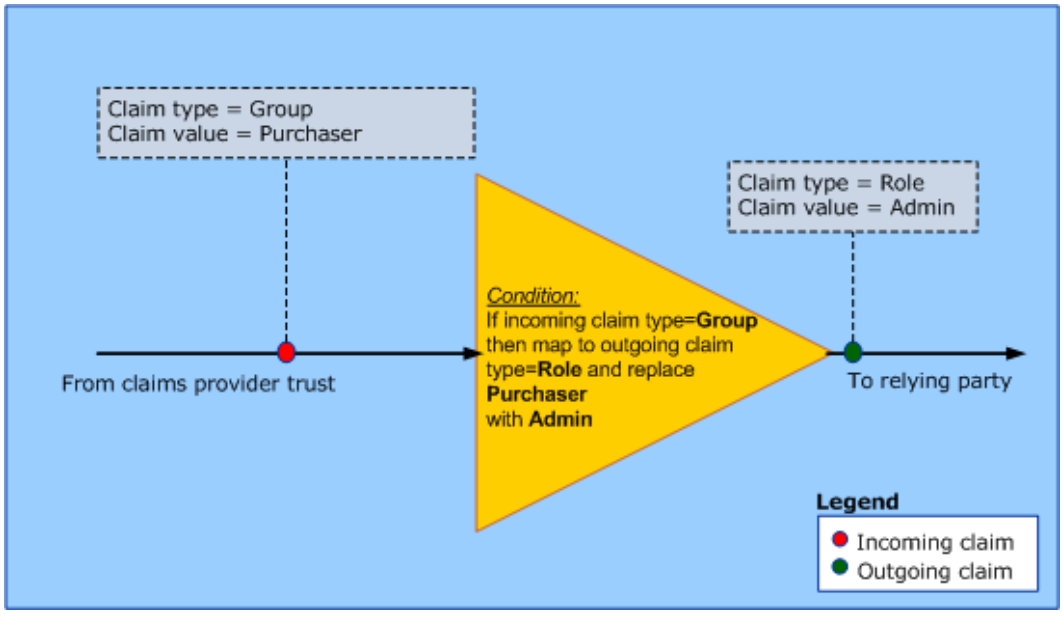

You can also use this rule to apply a condition that will replace all incoming claims with a specified e-mail suffix value with a new value. For example, you could set a condition in this rule to change all claim values with the suffix of sales.corp.fabrikam.com to fabrikam.com.

Configuring this rule on a claims provider trust

When you use a claims provider trust, this rule can be configured to transform incoming claims from the claims provider into trustworthy equivalents. Claim types or claim values can have a different meaning in your organization than in the claims provider organizations. You can use this rule to normalize the claim types and values that come from the claims provider so that their outgoing claim equivalents can be understood by the relying party.

Configuring this rule on a relying party trust

When you use a relying party trust, this rule can be configured to transform claims for the specific relying party. Claim types or claim values might have a different meaning for a specific relying party, and this rule makes it possible for you to change the outgoing claim types and values for a single relying party.

How to create this rule

You create this rule using either the claim rule language or using the **Transform an Incoming Claim** rule template in the AD FS Management snap-in. This rule template provides the following configuration options:

- Specify a claim rule name
- Transform a specific incoming claim type to a specified outgoing claim type
- Pass through all claim values
- Replace an incoming claim value with a different outgoing claim value
- Replace incoming e-mail suffix claims with a new e-mail suffix

For more instructions on how to create this template, see Create a Rule to Transform an Incoming Claim in the AD FS Deployment Guide.

Using the claim rule language

If the outgoing claim must be constructed from the content of more than one incoming claim, you must use a custom rule instead. If the claim value of the outgoing claim must be based on the value of the incoming claim—but with additional content—you must also use a custom rule in that context. For more information, see When to Use a Custom Claim Rule.

Examples of how to construct a transform rule syntax

When using the claim rule language syntax to transform claims, you can set a property of the transformed claim to a new literal value. For example, the following rule changes the value of role claims from “Administrators” to “root” while keeping the same claim type:

```
c:[type == “http://schemas.microsoft.com/ws/2008/06/identity/claims/role”, value == “Administrators”]  => issue(type = c.type, value = “root”);
```

Regular expressions can also be used for claim transformations. For example, the following rule will set the domain in windows user name claims in DOMAIN\USER format to FABRIKAM:

```
c:[type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name"] => issue(type = c.type, value = regexreplace(c.value, "(?<domain>[^\\]+)\\(?<user>.+)", "FABRIKAM\${user}"));
```

Best practices for creating custom rules

Claims transformations can be selectively applied to claims selected using basic filtering capabilities. Each of the claim properties used for filtering can be assigned values, with the following caveats:

Claim property	Description
Type, Value, ValueType	These properties will be most frequently used for assignments. At the very least type and value must be specified for the resulting transformed claim.
Issuer	While the claim rule language allows setting the Issuer of a claim, this is generally not advisable. The issuer of a claim is not serialized in the token. When a token is received the Issuer property of all claims is set to the identifier of the federation server that signed the token. Thus, setting the issuer of a claim in the rules will not have effect on the contents of the token and the setting will be lost once the claim is packaged in a token. The only scenario where setting the issuer of a claim makes sense is if it is set to a specific value in the claims provider rule set and relying party rule set is authored with rules that reference this specific value. If the Issuer property is not explicitly set to a value in a claim rule the claims issuance engine sets it to “LOCAL AUTHORITY”.

Claim property	Description
OriginalIssuer	Similarly to Issuer, OriginalIssuer should generally not be explicitly assigned a value. Unlike Issuer, the OriginalIssuer property is serialized in the token, but the expectation of token consumers is that, if set, it will contain the identifier of the federation server that originally issued a claim.
Properties	As outlined in the previous section, the Property bag of a claim is not persisted in the token, so assignments to properties should only be done if subsequent local policies are going to reference the information stored in the property.

For more information about how to use the claim rule language, see The Role of the Claim Rule Language.

When to Use a Send LDAP Attributes as Claims Rule

You can use this rule in Active Directory Federation Services (AD FS) when you want to issue outgoing claims that contain actual Lightweight Directory Access Protocol (LDAP) attribute values that exist in an attribute store and then associate a claim type with each of the LDAP attributes. For more information about attribute stores, see The Role of Attribute Stores.

When you use this rule, you issue a claim for each LDAP attribute that you specify and that matches the rule logic, as described in the following table.

Rule option	Rule logic
Mapping of LDAP attributes to outgoing claim types	If attribute store equals *specified attribute store* and LDAP attribute equals *specified value*, then map the LDAP attribute value to the *specified outgoing claim* type and issue the claim.

The following sections provide a basic introduction to claim rules. They also provide details about when to use the Send LDAP Attributes as Claims rule.

About claim rules

A claim rule represents an instance of business logic that will take an incoming claim, apply a condition to it (if x then y) and produce an outgoing claim based on the condition parameters. The following list outlines important tips that you should know about claim rules before you read further in this topic:

- In the AD FS Management snap-in, claim rules can only be created using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on rules, you can further refine or filter claims that are generated by previous rules within a given rule set.
- Claim rule templates will always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

Mapping of LDAP attributes to outgoing claim types

When you use the Send LDAP Attributes as Claims rule template, you can select attributes from an LDAP attribute store, such as Active Directory or Active Directory Domain Services (AD DS) to send their values as claims to the relying party. This essentially maps specific LDAP attributes from an attribute store that you define to a set of outgoing claims that can be used for authorization.

By using this template, you can add multiple attributes, which will be sent as multiple claims, from a single rule. For example, you can use this rule template to create a rule that will look up attribute values for authenticated users from the **company** and **department** Active Directory attributes and then send these values as two different outgoing claims.

You can also use this rule to send all the user's group memberships. If you want to send only individual group memberships, use the Send Group Membership as a Claim rule template. For more information, see When to Use a Send Group Membership as a Claim Rule.

How to create this rule

You can create this rule by using either the claim rule language or by using the Send LDAP Attributes as Claims rule template in the AD FS Management snap-in. This rule template provides the following configuration options:

- Specify a claim rule name

- Select an attribute store from which to extract LDAP attributes
- Mapping of LDAP attributes to outgoing claim types

For more information about how to create this rule, see Create a Rule to Send LDAP Attributes as Claims.

Using the claim rule language

If the query to Active Directory, AD DS, or Active Directory Lightweight Directory Services (AD LDS) must compare against an LDAP attribute other than **samAccountname**, you must use a custom rule instead. If there is no Windows Account Name claim in the input set, you must also use a custom rule to specify the claim to use for querying AD DS or AD LDS.

The following examples are provided to help you understand some of the various ways you can construct a custom rule using the claim rule language to query and extract data in an attribute store.

Example: How to query an AD LDS attribute store and return a specified value

Parameters must be separated by a semicolon. The first parameter is the LDAP filter. Subsequent parameters are the attributes to return on any matching objects.

The following example shows how to look up a user by the **sAMAccountName** attribute and issue an e-mail address claim, using the value of the user's mail attribute:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
Issuer == "AD AUTHORITY"]
=> issue(store = "AD LDS", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"), query =
"sAMAccountName={0};mail", param = regexreplace(c.Value,
"(?<domain>[^\\]+)\\(?<user>.+)", "${user}"));
```

The following example shows how to look up a user by the **mail** attribute and issue Title and Display Name claims, using the values of the user's **title** and **displayname** attributes:

```
c:[Type == " http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress ", Issuer
== "AD AUTHORITY"]
=> issue(store = "AD LDS ", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/title","http://schemas.xmlsoap.or
g/ws/2005/05/identity/claims/displayname"), query = "mail={0};title;displayname", param =
c.Value);
```

The following example shows how to look up a user by mail and title and then issue a Display Name claim using the user's **displayname** attribute:

```
c1:[Type == " http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"] &&
c2:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/title"]
=> issue(store = "AD LDS ", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/displayname"), query =
"(&(mail={0})(title={1}));displayname", param = c1.Value, param = c2.Value);
```

Example: How to query an Active Directory attribute store and return a specified value

The Active Directory query must include the user's name (with the domain name) as the final parameter so that the Active Directory attribute store can query the correct domain. Otherwise, the same syntax is supported.

The following example shows how to look up a user by the **sAMAccountName** attribute in his or her domain and then return the **mail** attribute:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
Issuer == "AD AUTHORITY"]
=> issue(store = "Active Directory", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"), query =
"sAMAccountName={0};mail;{1}", param = regexreplace(c.Value,
"(?<domain>[^\\]+)\\(?<user>.+)", "${user}"), param = c.Value);
```

Example: How to query an Active Directory attribute store based on the value of an incoming claim

```
c:[Type == "http://test/name"]

   => issue(store = "Enterprise AD Attribute Store",

         types = ("http://test/email"),

         query = ";mail;{0}",

         param = c.Value)
```

The previous query is made up of the following three parts:

- The LDAP filter—You specify this part of the query to retrieve the objects for which you want to query the attributes. For general information about valid LDAP queries, see RFC 2254. When you are querying an Active Directory attribute store and you do not specify an LDAP filter, the samAccountName={0} query is assumed and the Active Directory attribute store

expects a parameter that can feed the value for {0}. Otherwise, the query will result in an error. For an LDAP attribute store other than Active Directory, you cannot omit the LDAP filter part of the query or the query will result in an error.

- Attribute specification—In this second part of the query, you specify the attributes (which are comma-separated if you use multiple attribute values) that you want out of the filtered objects. The number of attributes that you specify must match the number of claim types that you define in the query.
- Active Directory domain—You specify the last part of the query only when the attribute store is Active Directory. (It is not necessary when you query other attribute stores.) This part of the query is used to specify the user account in the form domain\name. The Active Directory attribute store uses the domain part to determine the appropriate domain controller to connect to and run the query and request the attributes.

Example: How to use two custom rules to extract the manager e-mail from an attribute in Active Directory

The following two custom rules, when used together in the order shown below, query Active Directory for the **manager** attribute of the user account (Rule 1) and then use that attribute to query the user account of the manager for the **mail** attribute (Rule 2). Finally, the **mail** attribute is issued as a “ManagerEmail” claim. In summary, Rule 1 queries Active Directory and passes the result of the query to Rule 2, which then extracts the manager e-mail values.

For example, when these rules finish running, a claim is issued that contains the manager’s e-mail address for a user in the corp.fabrikam.com domain.

Rule 1

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname"]
=> add(store = "Active Directory", types =
("http://schemas.xmlsoap.org/claims/ManagerDistinguishedName"), query = "sAMAccountName=
{0};mail,userPrincipalName,extensionAttribute5,manager,department,extensionAttribute2,cn;
{1}", param = regexreplace(c.Value, "(?
<domain>[^\\]+)\\(?<user>.+)", "${user}"), param = c.Value);
```

Rule 2

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname"]
&& c1:[Type == "http://schemas.xmlsoap.org/claims/ManagerDistinguishedName"]
=> issue(store = "Active Directory", types =
("http://schemas.xmlsoap.org/claims/ManagerEmail"), query =
"distinguishedName={0};mail;{1}", param = c1.Value,
param = regexreplace(c1.Value, ".*DC=(?<domain>.+),DC=corp,DC=fabrikam,DC=com",
"${domain}\username"));
```

Note

These rules work only if the user's manager is in the same domain as the user (corp.fabrikam.com in this example).

Additional references

Attribute Stores

Create a Rule to Send LDAP Attributes as Claims

When to Use a Send Group Membership as a Claim Rule

You can use this rule in Active Directory Federation Services (AD FS) when you want to issue a new outgoing claim value for only those users who are members of a specified Activ Directory security group. When you use this rule, you issue a single claim for only the group that you specify and that matches the rule logic, as described in the following table.

Rule option	Rule logic
Outgoing claim value	If a user's group membership is equal to the *specified group* and outgoing claim type equals *specified claim type*, then replace the existing group name value with the *specified outgoing claim value* and issue the claim.

The following sections provide a basic introduction to claim rules. They also provide details about when to use the Send Group Membership as a Claim rule.

About claim rules

A claim rule represents an instance of business logic that will take an incoming claim, apply a condition to it (if x then y) and produce an outgoing claim based on the condition parameters. The following list outlines important tips that you should know about claim rules before you read further in this topic:

- In the AD FS Management snap-in, claim rules can only be created using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on rules, you can further refine or filter claims that are generated by previous rules within a given rule set.

- Claim rule templates will always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

Outgoing claim value

Using the Send Group Membership as a Claim rule template, you can issue a claim that is contingent on whether a user is a member of a group that you specify.

In other words, this rule template issues a claim only when the user has the group security ID (SID) that matches the Active Directory group that the administrator specifies. All users who authenticate against Active Directory Domain Services (AD DS) will have incoming group SID claims for each group that they belong to. By default, the acceptance transform rules in the Active Directory Claims Provider Trust pass through these group SID claims. Using these group SIDs as a basis for issuing claims is much faster than looking up the user's groups in AD DS.

When you use this rule, only a single claim is sent, based on the Active Directory group that you select. For example, you can use this rule template to create a rule that will send a group claim with a value of "Admin" if the user is a member of the Domain Admins security group.

Configuring this rule on a claims provider trust

Administrators should use this rule type in the acceptance transform rules of a claims provider trust only when group SIDs are being received from the claims provider, which is very uncommon for any claims providers except Active Directory or AD DS.

How to create this rule

You create this rule using either the claim rule language or by using the Send LDAP Group Membership as a Claim rule template in the AD FS Management snap-in. This rule template provides the following configuration options:

- Specify a claim rule name
- Select a user's group using object picker
- Select an outgoing claim type
- Select an outgoing name ID format (which is available only when Name ID is chosen from the outgoing claim type field)
- Specify an outgoing claim value

For more information about how to create this rule, see Create a Rule to Send Group Membership as a Claim.

Using the claim rule language

If you want to issue claims based on an incoming SID other than a group SID, use the Transform an Incoming Claim rule template. If the administrator wants to retrieve the names for all the groups that the user is a member of, use the Send LDAP Attributes as Claims rule template instead with the **tokenGroups** attribute.

Example: How to issue group claims based on the user's group membership

The following rule issues group claims for a user based on an incoming group SID:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/groupsid", Value ==
"S-1-5-21-397933417-626991126-188441444-512", Issuer == "AD AUTHORITY"]
=> issue(Type = "http://schemas.xmlsoap.org/claims/Group", Value = "administrators",
Issuer = c.Issuer, OriginalIssuer = c.OriginalIssuer, ValueType = c.ValueType);
```

Additional references

Attribute Stores

Create a Rule to Send LDAP Attributes as Claims

When to Use a Custom Claim Rule

You write a custom claim rule in Active Directory Federation Services (AD FS) using the claim rule language, which is the framework that the claims issuance engine uses to programmatically generate, transform, pass through, and filter claims. By using a custom rule, you can create rules with more complex logic than a standard rule template. Consider using a custom rule when you want to:

- Send claims based on values that are extracted from a Structured Query Language (SQL) attribute store.
- Send claims based on values that are extracted from a Lightweight Directory Access Protocol (LDAP) attribute store using a custom LDAP filter.
- Send claims based on values that are extracted from a custom attribute store.
- Send claims only when two or more incoming claims are present.
- Send claims only when an incoming claim value matches a complex pattern.
- Send claims with complex changes to an incoming claim value.
- Create claims for use only in later rules, without actually sending the claims.
- Construct an outgoing claim from the content of more than one incoming claim.

You can also use a custom rule when the claim value of the outgoing claim must be based on the value of the incoming claim, but it must also include additional content.

The claims rule language is rule based. It has a condition part and an execution part. You can use the claim rule language syntax to enumerate, add, delete, or modify claims to meet the needs of your organization. For more information about how each of these parts works, see The Role of the Claim Rule Language.

The following sections provide a basic introduction to claim rules. They also provide details about when to use a custom claim rule.

About claim rules

A claim rule represents an instance of business logic that takes an incoming claim, apply a condition to it (if x, then y) and produce an outgoing claim based on the condition parameters.

Important

- In the AD FS Management snap-in, claim rules can be created only using claim rule templates
- Claim rules process incoming claims either directly from a claims provider (such as Active Directory or another Federation Service) or from the output of the acceptance transform rules on a claims provider trust.
- Claim rules are processed by the claims issuance engine in chronological order within a given rule set. By setting precedence on the rules, you can further refine or filter claims that are generated by previous rules within a given rule set.
- Claim rule templates always require you to specify an incoming claim type. However, you can process multiple claim values with the same claim type by using a single rule.

For more detailed information about claim rules and claim rule sets, see The Role of Claim Rules. For more information about how rules are processed, see The Role of the Claims Engine. For more information how claim rule sets are processed, see The Role of the Claims Pipeline.

How to create this rule

You create this rule by first authoring the syntax that you need for your operation using the claim rule language and then pasting the result into the text box that is provided in the Send a Claims Using a Custom Rule template on the properties of either a claims provider trust or a relying party trust in the AD FS Management snap-in.

This rule template provides the following options:

- Specify a claim rule name
- Type one or more optional conditions and an issuance statement using the AD FS claim rule language

For more instructions for creating a custom rule using this template, see Create a Rule to Send Claims Using a Custom Rule in the AD FS Deployment Guide.

For a better understanding of how the claim rule language works, view the claim rule language syntax of other rules that already exist in the snap-in by clicking the **View Rule Language** tab in

the properties for that rule. Using the information in this section and the syntax information on this tab can provide insight into how to construct your own custom rules.

For more information about how to use the claim rule language, see The Role of the Claim Rule Language.

Using the claim rule language

Example: How to combine first and last names based on a user's name attribute values

The following rule syntax combines first and last names from attribute values in a given attribute store. The policy engine forms a cartesian product of the matches for each condition. For example, the output for first name {"Frank", "Alan"} and last names {"Miller", "Shen"} is {"Frank Miller", "Frank Shen", "Alan Miller", "Alan Shen"}:

```
c1:[type == "http://exampleschema/firstname" ]

&&  c2:[type == "http://exampleschema/lastname",]

=> issue(type = "http://exampleschema/name", value = c1.value + " " + c2.value);
```

Example: How to issue a manager claim based on whether users have direct reports

The following rule issues a manager claim only if the user has direct reports:

```
c:[type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name"] => add(store =
"SQL Store", types = ("http://schemas.xmlsoap.org/claims/Reports"), query = "SELECT
Reports FROM dbo.DirectReports WHERE UserName = {0}", param = c.value );

count([type == "http://schemas.xmlsoap.org/claims/Reports"] ) > 0 => issue(=
"http://schemas.xmlsoap.org/claims/ismanager", value = "true");
```

Example: How to issue a PPID claim based on an LDAP attribute

The following rule issues a Private Personal Identifier (PPID) claim based on the **windowsaccountname** and **originalissuer** attributes of users in an LDAP attribute store:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname"]

 => issue(store = "_OpaqueIdStore", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/privatepersonalidentifier"),
query = "{0};{1};{2}", param = "ppid", param = c.Value, param = c.OriginalIssuer);
```

Common attributes that can be used to uniquely identify the user for this query include the following:

- **user SID**

- **windowsaccountname**
- **samaccountname**

Active Directory Lightweight Directory Services Overview

Active Directory Lightweight Directory Services (AD LDS) is a Lightweight Directory Access Protocol (LDAP) directory service that provides flexible support for directory-enabled applications, without the dependencies and domain-related restrictions of Active Directory Domain Services (AD DS). You can run AD LDS on member servers or stand-alone servers. You can also run multiple instances of AD LDS—each with its own independently managed schema—on one server.

By using the Windows Server® 2008 Active Directory® Lightweight Directory Services (AD LDS) role, formerly known as Active Directory Application Mode (ADAM), you can provide directory services for directory-enabled applications without incurring the overhead of domains and forests and the requirements of a single schema throughout a forest.In the following sections, learn more about the AD LDS server role, the features in it, and the software and hardware considerations for installing it.

Following is the guidance available for AD LDS

- Overview of AD LDS
- Checklist: Create a new instance of AD LDS
- Checklist: Create a replica of an existing AD LDS instance
- Checklist: Import data from LDIF files
- Checklist: Create an application partition
- Checklist: Manage group memberships
- Checklist: Back up or restore an AD LDS instance
- Checklist: Synchronize data from AD DS to AD LDS
- Working with instances
- Working with Directory Partitions
- Working with Authentication and Access Control
- AD LDS Administration Tools
- User Interface: AD LDS Installation Wizard

Active Directory Rights Management Services Overview

This document provides an overview of Active Directory Rights Management Services (AD RMS) in Windows Server® 2012. AD RMS is the server role that provides you with management and development tools that work with industry security technologies—including encryption, certificates, and authentication—to help organizations create reliable information protection solutions.

Did you mean…

- Active Directory Certificate Services Overview
- Active Directory Domain Services Overview
- Active Directory Federation Services Overview
- Active Directory Lightweight Directory Services Overview

Role description

AD RMS can be used to augment the security strategy for your organization by protecting documents using information rights management (IRM).

AD RMS allows individuals and administrators through IRM policies to specify access permissions to documents, workbooks, and presentations. This helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. After permission for a file has been restricted by using IRM, the access and usage restrictions are enforced no matter where the information is, because the permission to a file is stored in the document file itself.

AD RMS and IRM help individuals enforce their personal preferences concerning the transmission of personal or private information. They also help organizations enforce corporate policy governing the control and dissemination of confidential or proprietary information.

Practical applications

IRM solutions that AD RMS enables are used to help provide the following:

- Persistent usage policies, which remain with the information, no matter where it is moved, sent or forwarded.
- An additional layer of privacy to protect sensitive information —such as financial reports, product specifications, customer data, and confidential e-mail messages—from intentionally or accidentally getting into the wrong hands.
- Prevent an authorized recipient of restricted content from forwarding, copying, modifying, printing, faxing, or pasting the content for unauthorized use
- Prevent restricted content from being copied by using the Print Screen feature in Microsoft Windows
- Support file expiration so that content in documents can no longer be viewed after a specified period of time

- Enforce corporate policies that govern the use and dissemination of content within the company

IRM-based solutions that AD RMS supports cannot prevent all types of threats to the security of sensitive documents or prevent disclosure of screen readable information under all circumstances. For example, the following are some types of document security threats that AD RMS does not address or mitigate:

- Content from being erased, stolen, or captured and transmitted by malicious programs such as Trojan horses, keystroke loggers, and certain types of spyware
- Content from being lost or corrupted because of the actions of computer viruses
- Restricted content from being hand-copied or retyped from a display on a recipient's screen
- A recipient from taking a digital photograph of the restricted content displayed on a screen
- Restricted content from being copied by using third-party screen-capture programs

For more information on how AD RMS can be used to design secure document collaboration, see AD RMS Architecture Design and Secure Collaboration Scenarios.

New and changed functionality

Several improvements have been made to the Windows Server 2012 version of AD RMS. These enhancements are covered online in the article What's New in AD RMS?

Server Manager information

The installation of AD RMS role services can be performed through the Server Manager. The following role services can be installed:

Role service	Description
Active Directory Rights Management Server	The Active Directory Rights Management Server is a required role service that installs all AD RMS features used to publish and consume rights-protected content.
Identity Federation Support	The identity federation support role service is an optional role service that allows federated identities to consume rights-protected content by using Active Directory Federation Services.

See also

The following table provides additional resources for evaluating AD CS.

Content type	References
Product evaluation	• Test Lab Guide: Deploying an AD RMS Cluster (http://go.microsoft.com/fwlink/?LinkID=252886)
Deployment	• How AD RMS Works • AD RMS Infrastructure Deployment Tips • Using AD RMS with Hardware Security Modules
Community resources	• Roadmap for AD RMS documentation in TechNet Library: AD RMS Documentation Roadmap (http://technet.microsoft.com/library/dd772711(v=WS.10).aspx) • Frequently asked questions (FAQs) list Active Directory Rights Management Services (AD RMS) Frequently Asked Questions (FAQ) (http://social.technet.microsoft.com/wiki/contents/articles/3440.ad-rms-faq.aspx) • Support forum: RMS Forum (http://go.microsoft.com/FWLink/?LinkId=262353) • Product team blog for IT professionals: AD RMS Team Blog (http://blogs.technet.com/b/rms/) • AD RMS Developer's Corner Blog: AD RMS Developer's Corner Blog (http://blogs.msdn.com/b/rms/) • Script repository: TechNet Script Center Repository (http://gallery.technet.microsoft.com/scriptcenter) search for RMS, AD RMS or rights management. • Community technology overview: Active Directory Rights Management Services (AD RMS) Overview (http://social.technet.microsoft.com/wiki/contents/articles/967.aspx)
Related technologies	Active Directory Certificate Services Active Directory Domain Services Active Directory Federation Services Active Directory Lightweight Directory Services

What's New in Active Directory Rights Management Services (AD RMS)?

This document provides details of new deployment enhancements for Active Directory Rights Management Services (AD RMS) in Windows Server® 2012. These changes should enable IT

professionals working with AD RMS to meet the needs of their business in a secure, reliable and flexible way.

Role/Feature description

AD RMS is the server role that provides you with management and development tools that work with industry security technologies—including encryption, certificates, and authentication—to help organizations create reliable information protection solutions. For more information, see Active Directory Rights Management Services Overview.

New AD RMS features covered in this topic include:

- Changes to AD RMS and Microsoft SQL Server requirements.
- Changes in deployment of AD RMS when using Server Manager and Windows PowerShell

Changes in AD RMS and SQL Server requirements

For Windows Server 2012, requirements for configuring Microsoft SQL Server to support AD RMS have been revised or changed in response to customer feedback.

What value does this change add?

The following changes to AD RMS Setup have been made to enable better support for remote deployment of AD RMS and SQL servers and to address customer feedback that requested more flexible deployment options.

What works differently?

In previous releases, AD RMS setup required that the account used to install AD RMS must have local administrator privilege on any computers hosting a SQL Server installation that was to be used to support the storage of AD RMS-related data. This was because AD RMS Setup required the ability to read SQL database settings from the Windows Registry. Because of customer feedback, this has been changed for this release.

For Windows Server 2012, AD RMS now has the following requirements for access to SQL Server.

- The AD RMS installer account must have sysadmin permissions in the SQL Server installation.
- The SQL Server Browser service must be running to locate available SQL instances.
- Firewall exceptions should be enabled on the SQL server computer for ports that will be used by AD RMS setup. The TCP port for the SQL instance that will host the AD RMS databases should be enabled. The UDP port for the SQL Server Browser service should also be enabled. For example, the default ports are usually TCP port 1433 for the SQL Server instance and UDP port 1434 for the SQL Server Browser service.

In addition to the previous access requirements, for Windows Server 2012 the following versions of Microsoft SQL Server have been tested and are supported for use with AD RMS deployment.

- SQL Server 2005 Service Pack 3
- SQL Server 2008 Service Pack 3

- SQL Server 2008 R2 Service Pack 1
- SQL Server 2012

For more information on applicable SQL Server requirements for AD RMS, see AD RMS SQL Server Requirements (http://technet.microsoft.com/library/dd772673(WS.10).aspx).

Changes in deployment of AD RMS for Server Manager and Windows PowerShell

In previous releases, AD RMS Setup supported only deployment at the same server computer where AD RMS was to be installed. Based on customer feedback, this has been changed. For Windows Server 2012, AD RMS now supports remote deployment at targeted server computers.

What value does this change add?

The following changes to AD RMS Setup have been made to enable better integration with the newly revised Server Manager in Windows Server 2012 which better supports secure and flexible remote server deployment for AD RMS and other Active Directory technologies.

What works differently?

As a result of this change, when deploying AD RMS on remote servers to ensure the security of your deployment, you will be prompted for user credentials. Deployment now also requires an additional step that can be completed using either Server Manager or Windows PowerShell.

Using Server Manager to deploy AD RMS

For Windows Server 2012, Server Manager has been redesigned to provide support for remote deployment of AD RMS as part of a two-step process that can be summarized as follows:

1. Launch the Add Roles and Features Wizard in Server Manager to add the AD RMS role. This will add and install the files necessary for AD RMS.
2. After adding the AD RMS role, launch the AD RMS Configuration wizard to select deployment options and configure the AD RMS cluster.

When the AD RMS configuration wizard first launches, if you are installing AD RMS on a remote server you will be prompted for the credentials needed to complete AD RMS configuration.

The requirements for selecting the credentials that you enter here are as follows:

- The account used to deploy AD RMS must have membership in the local Administrators group on the server computer where you are installing and configuring AD RMS.
- The account used must also have sysadmin permissions on the server that hosts the configuration database for the AD RMS cluster.

> **Note**
> Prompting for credentials only occurs in Server Manager if you are installing AD RMS to a remote server; however, the credentials requirements are the same for installing locally.

Additionally, if you wish to register the AD RMS service connection point (SCP) with Active Directory during configuration, the account also needs to be a member of the Enterprise Admins

group for the forest. Because this requires additional privilege, you might choose to complete this task later after you have completed the basic server configuration. To do so, you must use the AD RMS console. For more information, see Register a Service Connection Point (http://technet.microsoft.com/library/cc725573.aspx).

Note

If you are going to be viewing reports related to AD RMS, you must also install the .NET Framework 3.5 using Server Manager (or by typing `Install-WindowsFeature NET-Framework-Core` at the Windows PowerShell prompt) prior to installing the Microsoft Report Viewer 2008, which is used to generate troubleshooting and system health reports on AD RMS in Windows Server 2012. For more information, see the Test Lab Guide: Deploying an AD RMS Cluster.

Using Windows PowerShell to deploy AD RMS

For Windows Server 2012, you can now optionally complete the two-part AD RMS deployment process as described in the previous section using Windows PowerShell commands. In Windows Server 2008 R2, the ADRMS module was provided as part of the base operating system but this has changed to enable more flexibility and modularity in installing the product. For this release, you need to use commands provided within the new ServerManager module for Windows PowerShell.

For example, to complete the first part of the deployment process (copy and install all required files for AD RMS) you would use the following Server Manager cmdlet examples at a Windows PowerShell prompt.

Example	Description
`Add-WindowsFeature ADRMS -IncludeAllSubFeature -IncludeManagementTools`	Adds all AD RMS role services and tools. This command downloads and makes available all support files needed to work with AD RMS.
`Add-WindowsFeature ADRMS-Server`	Adds the AD RMS server role only. This command downloads and makes available those files needed to support an AD RMS server installation.
`Add-WindowsFeature ADRMS-Identity`	Adds identity federation support for AD RMS. This command downloads and makes available files needed to support AD RMS working with AD FS.

You can then install AD RMS by using the **Install-ADRMS** cmdlet which is provided as part of the AD RMS deployment module for Windows PowerShell. When using this cmdlet, use the new **Credentials** parameter as shown below. This parameter is required to support remote deployments and is recommended that you include it for all PowerShell-based deployments.

```
Install-ADRMS –Path adrmsDrive:\ -Credential
```

By adding the **-Credential** parameter you will be prompted to enter the required credentials to deploy AD RMS. If the installation is targeted for a remote computer, credentials must be passed through and validated at the remote computer before installation of AD RMS will proceed.

Using Windows PowerShell to uninstall AD RMS

For Windows Server 2012, uninstalling AD RMS can also be accomplished using Windows PowerShell as a two-step process comparable to a reversal of the process described in the previous section for AD RMS deployment. This process can be summarized as follows:

1. Uninstall the AD RMS role.

   This can be performed separately by using Windows PowerShell and running the **Uninstall-ADRMS** cmdlet in the **ADRMS** module.

2. Remove the AD RMS files and registry settings that were previously added to support AD RMS.

   If you are using Windows PowerShell, use the **Remove-WindowsFeature ADRMS** cmdlet in the ServerManager module to perform this step as a separate action.

Note

If you use the Remove Roles and Features Wizard in Server Manager to remove the AD RMS role it accomplishes as a single action both of the steps summarized above. Also, this same joint action occurs if you run **Remove-WindowsFeature ADRMS**, which calls the **Uninstall-ADRMS** cmdlet.

Changes to Windows PowerShell for deploying AD RMS

In addition to new changes to Windows PowerShell that affect deployment of the AD RMS role and its services and components, new properties have been added for Windows Server 2012 that can be used when deploying AD RMS installations. The following table provides details and comments on where these new properties can be used to support new changes such as support for strong cryptography.

Install Type	Container	Property	Default Value	Description
RootCluster	CryptoSupport	SupportCryptoMode 2	True	This container appears by default and is set to enable support for strong cryptography.
RootCluster	SSLCertificateSuppo	—	—	This container only

Install Type	Container	Property	Default Value	Description
	rt			appears when the ClusterURL property is set to HTTPS.
RootCluster	SSLCertificateSupport	SSLCertificateOption	ChooseLater	The default for this property enables you to choose a certificate later after deploying AD RMS. Other alternate values include **Create** (enables creation of a self-signed certificate) or **Existing** (Enables use of an existing certificate).
RootCluster	SSLCertificateSupport	Thumbprint	[not set]	This property is used to specify the thumbprint to identify a certificate. It only appears when the SSLCertificateOption property is set to Existing.
LicensingCluster	SSLCertificateSupport	—	—	This container only appears when the ClusterURL property is set to HTTPS.
LicensingCluster	SSLCertificateSupport	SSLCertificateOption	ChooseLater	The default for this property enables you to choose a certificate later after deploying AD RMS. Other alternate values include **Create** (enables

Install Type	Container	Property	Default Value	Description
				creation of a self-signed certificate) or **Existing** (Enables use of an existing certificate).
LicensingCluster	SSLCertificateSupport	Thumbprint	[not set]	This property is used to specify the thumbprint to identify a certificate. It only appears when the SSLCertificateOption property is set to Existing.
JoinCluster	SSLCertificateSupport	—	—	This container appears when the DatabaseName property is set to a cluster database that is using SSL.
JoinCluster	SSLCertificateSupport	SSLCertificateOption	ChooseLater	The default for this property enables you to choose a certificate later after deploying AD RMS. Other alternate values include **Create** (enables creation of a self-signed certificate) or **Existing** (Enables use of an existing certificate).
JoinCluster	SSLCertificateSupport	Thumbprint	[not set]	This property is used to specify the thumbprint to identify a certificate. It only appears when the

Install Type	Container	Property	Default Value	Description
				SSLCertificateOption property is set to Existing.
JoinCluster	ADFSSupport	—	—	This container exists by default in all AD RMS installations that are running at Windows 2008 R2 or later but only appears in Windows Server 2012 if the ADRMS-Identity role service is installed.
JoinCluster	ADFSSupport	ADFSUrl	[not set]	The value of this property must be a Web address (URL) to a valid AD FS server. For example, `"http://fs.corp.contoso.com"`

Notes

When the SSLCertificateOption property is configured to use a value of Existing and the Thumbprint property is empty or not set, it might appear incorrectly that installation of AD RMS has failed when using the **Install-ADRMS** cmdlet in a Windows PowerShell session. For example, under these conditions, you see the following error message after a PowerShell-based installation.

```
PS rmsdrive:\> install-adrms -path . -force
Install-ADRMS : Value cannot be null.
Parameter name: findValue
At line:1 char:1
```

If such an error should appear, it indicates only that AD RMS is missing an expected SSL binding. If this happens, you can assume that AD RMS is fully provisioned and operating as expected. To address the issue, you only need to configure SSL certificates on IIS before using the AD RMS server.

AD RMS fails to install if multiple installations are active simultaneously in Server Manager

In previous releases of AD RMS included with Windows Server® 2008 and Windows Server® 2008 R2 it was not possible to launch more than a single instance of the AD RMS Configuration wizard to install or update multiple AD RMS deployments from the same server computer. Because of design changes to Server Manager for Windows Server 2012, multiple instances of the Add Roles and Features Wizard can now be run simultaneously, making it possible to launch two or more instances of the AD RMS Configuration wizard.

What value does this change add?

The changes to Server Manager in Windows Server 2012 enable server roles and technologies across Windows Server 2012 to be more open, efficient and flexible when you are choosing role deployment options.

What works differently?

As a result of changes to Server Manager to enable multiple instances to be deployed and configured simultaneously in Windows Server 2012 and because this change is at odds with how AD RMS is best designed to be deployed, AD RMS setup is not able to prevent multiple AD RMS configuration wizards from being launched to create new clusters or for joining an existing cluster. Therefore, the AD RMS Configuration wizard applies the following logic and process detection to determine how to manage the occurrence of multiple instances of AD RMS setup occurring at the same time.

- If multiple instances are launched and are used to join an already existing AD RMS cluster, AD RMS setup should succeed for all active instances.
- If multiple instances are launched and are being used to try to create the same new AD RMS cluster, AD RMS setup fails for all active instances.

Deployment considerations for virtualized AD RMS servers

This section describes issues to consider when you deploy AD RMS on virtual machines. In general, AD RMS is well suited to virtualization especially for testing and evaluation purposes. There are, however, a few practices that you should observe when you are working with virtualized AD RMS servers. The following section describes these considerations.

What value does this change add?

Using virtualization platforms, such as Hyper-V to deploy AD RMS for evaluation purposes offers a number of convenience features that can make testing easier. For example, virtualization enables you to easily establish a full private virtual networked configuration separate from your production networked environment in which you can freely test and evaluate AD RMS deployment to your satisfaction before attempting a deployment of it in your production environment.

What works differently?

When working with virtualized AD RMS servers, the following are some deployment practices and features that need to be observed or followed for best results.

- To ensure consistent performance between virtualized AD RMS servers operating in test conditions to how a comparable physical AD RMS server might perform in a subsequent production deployment of AD RMS, follow the recommended guidance requirements as you create your virtual machines and then use those as well for installing AD RMS on a physical server computer. In particular, select the recommended amount of RAM for virtual AD RMS servers to the higher recommended amount. If you follow recommended hardware guidelines, there should be no perceptible difference in performance between virtualized test server deployments of AD RMS and any comparable physical server installations you might later deploy into production as a result of your virtual lab test results.
- When you install AD RMS on a virtual server you will not be able to use an internal hardware security module (HSM) for AD RMS key storage. You can, however, use all other forms of key storage: network HSM key storage, AD RMS centrally managed key storage and software CSP key storage.
- If you have previously installed AD RMS on a virtual hard drive (VHD) and then take that VHD offline to uninstall AD RMS, the AD RMS role will not be completely removed. To ensure full and proper removal of AD RMS, you should remove the AD RMS server role first before shutting down any VHD associated with the original installation of AD RMS on your virtual machine.

Server Core Support for AD RMS

For Windows Server 2012, AD RMS now joins the list of server roles such as Active Directory Domain Services (AD DS) and Active Directory Certificate Services (AD CS) that are supported for Server Core deployment. Server Core is an installation option that enables you to perform a minimal installation of the Windows Server operating system which can be useful for reducing total cost of ownership (TCO) in deploying and managing servers.

What value does this change add?

Using Server Core to deploy AD RMS can provide you with simplified management, reduced maintenance, reduced memory and disk requirements and a reduced attack surface. Servers that run under Server Core can be simpler and less costly to manage and more stable and secure in production deployments. For more information, see What is Server Core? and Why is Server Core Useful?

What works differently?

In Windows Server 2012, the Server Core installation option is no longer an irrevocable selection that is made during setup. In Windows Server 2008 R2 and Windows Server 2008, if your requirements changed, there was no way to convert between a full installation and Server Core installation without completely reinstalling the operating system. You can now administratively convert between a full Windows Server installation and running on Server Core as needed. For more information, see Windows Server Installation Options.

On Server Core installations, the optional Identity Federation Support role service for the AD RMS server role is not supported. This is because Identity Federation Support relies on a role

service of the AD FS Server role, the Claims-aware Agent, which is disabled on Server Core installations.

If you attempt to add Identity Federation Support for AD RMS, the following error will occur.

```
The request to add or remove features on the specified server failed.

Installation of one or more roles, role services, or features failed. - The source files could not be downloaded. Use the /source option to specify the location of the files that are required to restore the feature. The file location should either be the root directory of a mounted image or a component store that has the Windows Side-by-Side directory as an immediate subfolder.
```

The instructions to work around the issue by downloading the source files do not apply for this case. Identity Federation Support cannot be installed because the Claims-aware Agent is disabled on Server Core.

Recent feature updates for AD RMS

Windows Server 2012 also includes the following feature updates, which have been added recently as updates for the AD RMS role in Windows Server 2008 R2.

- Simple delegation
- Strong cryptography

Because these features are still relatively new to most customers, they are discussed more here to help increase awareness for them and provide additional context for those who are unfamiliar with their appearance in the AD RMS feature set.

Simple Delegation for AD RMS

Simple delegation for AD RMS enables you to have the same access rights to protected content that are assigned to one person delegated to other individuals within their organization.

What value does this change add?

Simple delegation provides the ability to have content rights assigned to executives and managers be easily and effectively delegated to their assistants. This enables the assistants to have the same level of access permission to Information Rights Management (IRM)-protected content as the executive. By extending the Active Directory schema to support two new attributes, delegation of rights can be easily turned on or off as needed to accomplish the appropriate level of privilege or restriction for those who support the leadership and management of your organization.

What works differently?

The requirements for enabling the use of AD RMS simple delegation in Windows Server 2012 are similar to the requirements for using it in Windows Server 2008 R2. Two attributes, **msRMSDelegator** and **msRMSDelegatorBL** must be added to the Active Directory schema. This can be done using update files executed using the Ldifde.exe tool. These extension attributes must be applied to all forests where AD RMS licensing occurs. Also, you might also

need to set additional permissions for AD RMS to read these attributes and update your SQL installation that supports AD RMS as part of preparing your deployment to support the use of simple delegation as well.

The only other change in requirements for simple delegation in Windows Server 2012 that differs from how it was enabled in Windows Server 2008 R2 is a new AD RMS PowerShell property that can be used to enable and disable it. By default, the feature is disabled, however, to enable simple delegation use the following command line at an elevated Windows PowerShell prompt at an AD RMS provided PowerShell drive:

```
PS w:\SecurityPolicy\Delegation> Set-ItemProperty -path . -Name IsEnabled -Value $true
```

For more information, see An update is available to enable simple delegation in AD RMS on Windows Server 2008 R2-based computers.

Strong cryptography for AD RMS

Strong cryptography for AD RMS enables you to increase the cryptographic strength of your AD RMS deployment by running in an advanced mode known as cryptographic mode 2.

What value does this change add?

Running AD RMS in strong cryptographic provides an updated and enhanced cryptographic implementation to support much stronger encryption as well as stronger cryptographic keys. For example, in mode 2 operation, RSA encryption is enhanced from 1024 bit encryption to 2048 bit encryption. Also, cryptographic keys used for hashing are longer (changed from 160-bit keys to 256-bit keys) and now use SHA-2 hashing algorithm (SHA-2/SHA-256) standards.

The value of strong cryptography in AD RMS is that it can be part of enabling your organization to satisfy regulatory compliance with current security standards that are set by the National Institute of Standards and Technology (NIST). Starting January 1, 2011, NIST issued Special Publication 800-57 which recommends the use of 2048-bit RSA keys. United States Federal agencies are required to comply with NIST recommendations and many private enterprises and other countries may choose to implement this recommendation. To learn more, see NIST Special Publications (http://csrc.nist.gov/publications/PubsSPs.html).

What works differently?

To enable the use of strong cryptography in your AD RMS deployment, all computers that host either AD RMS server or client software must be patched and updated.

To update AD RMS servers to cryptographic mode 2, you can use the AD RMS management console or the AD RMS cmdlets for Windows PowerShell.

Using the AD RMS management console, you can choose to update existing AD RMS servers from cryptographic mode 1 to cryptographic mode 2 by doing the following:

1. In the navigation pane, select the AD RMS server you want to upgrade.
2. From the **Action** menu (or in the Actions pane), select the **Update to Cryptographic Mode 2** option.

If you are using Windows PowerShell, you can also choose to update an existing AD RMS server from cryptographic mode 1 to cryptographic mode 2 using the ADRMS module.

To update to Cryptographic Mode 2 using Windows PowerShell, use the following syntax:

```
Update-ADRMS –UpdateCryptographicModeOnly –ServiceAccount <PSCredential> -force –
NewCSPName <”Mode2 Supported CSP”> -Regen
```

The following items apply to the parameters specified as part of the above syntax:

- **UpdateCryptographicModeOnly** is the parameter that indicates that cryptographic mode 2 should be enabled. This is a one-way operation. Once complete, you cannot return the AD RMS server to cryptographic mode 1.
- The **Force** parameter is an optional switch that can be used to override prompting for user confirmation.
- The **NewCSPName** parameter indicates the cryptographic providers that you want to use for encryption. This is an optional setting depending on your current cryptographic key settings. By default it is the Microsoft Enhanced RSA and AES Cryptographic Service Provider, but could also be another cryptographic mode 2 provider, such as one from a Hardware Security Module (HSM) that supports at least 2048 bit encryption.

> **Note**
> The CSP name cannot be specified for centrally managed keys. If an AD RMS cluster is using a centrally managed key for cryptographic mode 1, it must update to a centrally managed key for cryptographic mode 2. It cannot update to a CSP key.

As an example, if the AD RMS service account is named ADRMSSvc, you would open a Windows PowerShell prompt and run the following command to update the AD RMS server to cryptographic mode 2:

```
Update-ADRMS –UpdateCryptographicModeOnly –ServiceAccount ADRMSSvc –NewCSPName "Microsoft
Enhanced RSA and AES Cryptographic Provider"
```

For more information, see Active Directory Rights Management Service Cryptographic Modes.

See also

The following table provides additional resources for evaluating AD RMS.

Content type	References
Product evaluation	Active Directory Rights Management Services Overview
Community resources	AD RMS Blog | RMS Forum
Related technologies	Active Directory Domain Services Overview | Hyper-V Overview | Remote Desktop Services Overview

Test Lab Guide: Deploying an AD RMS Cluster

The purpose of this Test Lab Guide (TLG) is to enable you to set up a working Active Directory Rights Management Services (AD RMS) infrastructure in a test environment. During this process you create an Active Directory® domain, install a database server, install the AD RMS server role, configure the AD RMS cluster, and configure an AD RMS-enabled client computer.

Once complete, you can use the test lab environment to learn about AD RMS technology on Windows Server® 2012 and assess how it might be deployed in your organization.

As you complete the steps in this guide, you will:

- Prepare the AD RMS infrastructure.
- Install and configure AD RMS.
- Verify AD RMS functionality after you complete the configuration.

The goal of an AD RMS deployment is to be able to protect information, no matter where it goes. Once AD RMS protection is added to a digital file, the protection stays with the file. By default, only the content owner is able to remove the protection from the file. The owner grants rights to other users to perform actions on the content, such as the ability to view, copy, or print the file. For more information about the business reasons behind an AD RMS deployment, see the white paper "Windows Rights Management Services: Helping Organizations Safeguard Digital Information from Unauthorized Use" (http://go.microsoft.com/fwlink/?LinkId=64636).

Note
This guide is considered the basic AD RMS TLG. All other TLGs developed for AD RMS will assume that this guide has been completed first.

In this guide

This document contains instructions for extending the Windows Server® 2012 Base Configuration Test Lab Guide (TLG) to include an AD RMS cluster server on the APP1 server computer. In addition to extending APP1 to host the AD RMS server role, you will also need to configure the domain controller (DC1) and a desktop client computer (CLIENT1) As described in the instructions provided with the Base Configuration TLG.

In this guide you will deploy an additional SQL server computer (SQL1) which will be used to support the AD RMS configuration and logging databases, however, you will not need to configure the INET1 or EDGE1 computers from the Base Configuration TLG as they are not required for the purposes of establishing a working lab environment for testing AD RMS deployment.

Important
The configuration of the computers and network in this guide was designed to give you hands-on practice in creating an AD RMS test environment. The design decisions made

in this guide were geared toward increasing your hands-on experience and to some degree reflect AD RMS best practices configuration. For full best practices and design and planning information related to AD RMS, see AD RMS Prerequisites (http://technet.microsoft.com/library/dd772659(v=WS.10).aspx), AD RMS Performance and Logging Best Practices (http://technet.microsoft.com/library/dd941633(v=ws.10).aspx) and AD RMS Architecture Design and Secure Collaboration Scenarios (http://technet.microsoft.com/library/dd983947(v=ws.10).aspx).

What this guide does not provide

This guide does not provide the following:

- An overview of AD RMS. For more information about the advantages that AD RMS can bring to your organization, see Active Directory Rights Management Services Overview (http://go.microsoft.com/fwlink/?LinkId=84726).
- Guidance for setting up and configuring AD RMS in a production environment.
- Complete technical reference for AD RMS.

Test lab overview

The test lab configuration demonstrated in this guide extends the Windows Server 2012 Base Configuration TLG by one server computer. The additional computer will serve as a SQL server computer and be named SQL1. There are four major steps in this test lab guide to complete that contain multiple procedures as a part of completing each step.

1. Complete installation and configuration of DC1, APP1 and CLIENT 1 as described in the Base TLG Configuration.
2. Configure SQL1 as a SQL server database server.
3. Install Office 2010 trial version on CLIENT1.
4. Configure APP1 as the AD RMS root cluster server.
5. Verify AD RMS functionality at CLIENT1.

We recommend that you first use the steps provided in this guide in a test lab environment. Test lab guides are not necessarily meant to be used to deploy Windows Server features without additional deployment documentation and should be used with discretion as a stand-alone document.

Upon completion of this test lab guide, you will have a working AD RMS infrastructure. You can then test and verify AD RMS functionality as follows:

- Restrict permissions on a Microsoft Office Word 2010 document.
- Have an authorized user open and work with the document.
- Have an unauthorized user attempt to open and work with the document.

The test environment described in this guide includes four computers connected to a private network and using the operating systems, applications, and services summarized in the following table.

Computer name	Operating system	Applications and services
APP1	Windows Server 2012	AD RMS, Web Server (IIS)
DC1	Windows Server 2012	Active Directory Domain Services (AD DS), Domain Name System (DNS)
SQL1	Windows Server 2012	Microsoft SQL Server® 2012 Express Edition
CLIENT1	Windows 8	Microsoft Office Word 2010 Professional

Note

For more information about the system requirements for installing AD RMS, see Pre-installation Information for Active Directory Rights Management Services (http://go.microsoft.com/fwlink/?LinkId=84733).

The computers form a private intranet and are connected through a common hub or Layer 2 switch. This configuration can be emulated in a virtual server environment if desired. This test lab exercise uses private addresses throughout the test lab configuration. The private network ID 10.0.0.0/24 is used for the intranet. The domain controller is named DC1 for the domain named corp.contoso.com. The following figure shows the configuration of the test environment:

AD RMS test lab overview

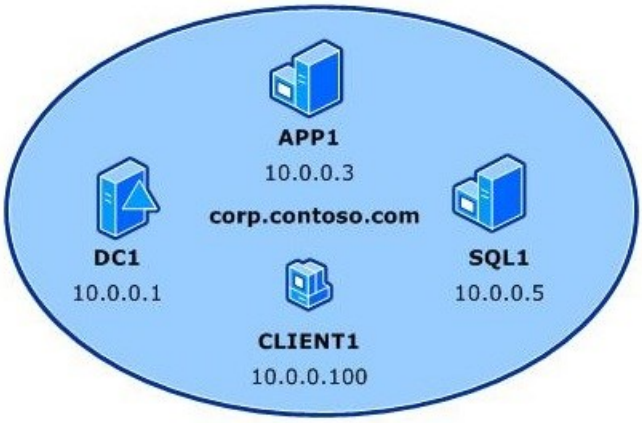

Hardware and software requirements

The following are the minimum required components of the test lab:

1. The product disc or files for Windows Server 2012.
2. The product disc or files for Windows® 8.
3. Three computers that meet the minimum hardware requirements for Windows Server 2012.
4. One computer that meets the minimum hardware requirements for Windows 8.
5. If you wish to deploy the Base Configuration test lab in a virtualized environment, your virtualization solution must support Windows Server 2012 64-bit virtual machines. The server hardware must support the amount of RAM required to run the virtual operating systems included in the Base Configuration test lab and any other virtual machines that may be required by additional TLGs.

Important

Run Windows Update on all computers or virtual machines either during the installation or immediately after installing the operating systems. After running Windows Update, you can isolate your physical or virtual test lab from your production network.

Note

If you will be installing and using the trial version of Microsoft Office 2010 Professional with the CLIENT1 computer, its best to download and complete the installation process on CLIENT1 while it is still configured with Internet access. Once Microsoft Office 2010 is installed, along with any Office updates and with activation online completed, you can finish configuring CLIENT1 by joining it to the CORP domain and then reconfiguring CLIENT1 to limit network access to only your test lab private network.

To install Office Professional Plus 2010 on CLIENT1

1. Log on to CLIENT1 with Internet access.
2. Locate your Microsoft Office Professional Plus 2010 product media or optionally, you can download Microsoft Office Professional Plus 2010 for trial installation from the Microsoft Web site. Be sure to download the Professional Plus edition as other editions might not support information rights management (IRM) using AD RMS.
3. Launch the installer for Microsoft Office 2010 to begin installation.
4. Click **Customize** as the installation type, set the installation type to **Not Available** for all applications except Word 2010, and then click **Install Now**. This might take several minutes to complete.

Step 1: Complete the Base TLG Configuration and add users and groups

Begin by completing the configuration of DC1, APP1, CLIENT1 virtual machines as described in the Windows Server 2012 Base Configuration Test Lab Guide (TLG). The Base TLG is located at http://go.microsoft.com/fwlink/p/?LinkId=236358.

Once you have completed the configuration of DC1 using the Base TLG instructions, you will want to configure the following additional user accounts for use with testing your AD RMS installation. The following table lists the user accounts that you will need to create at this time.

Account Name	User Logon Name	E-mail Address	Group
ADRMSSVC	ADRMSSVC		
ADRMSADMIN	ADRMSADMIN		Enterprise Admins
Nicole Holliday	NHOLLIDA	nhollida@contoso.com	Employees, Finance
Limor Henig	LHENIG	lhenig@contoso.com	Employees, Marketing
Stuart Railson	SRAILSON	srailson@contoso.com	Employees, Engineering

To add new user accounts

1. Log on to DC1 using the CORP\Administrator account.
2. In Server Manager, click **Local Server** in the console tree, then click **Tools**, and then select **Active Directory Users and Computers**.
3. In the console tree, expand **corp.contoso.com**.
4. Right-click **Users**, point to **New**, and then click **User**.
5. In the New Object – User dialog box, type **ADRMSSVC** in the **First name** and **User logon** name boxes, and then click **Next**.
6. In the New Object – User dialog box, type a password of your choice in the **Password** and **Confirm password** boxes. Clear the **User must change password at next logon** check box, click **Next**, and then click **Finish**.
7. Perform steps 3-6 for each of the following users: ADRMSADMIN, Nicole Holliday, Limor Henig, and Stuart Railson.

To add e-mail addresses to user accounts

1. In the **Active Directory Users and Computers** console, right-click **Nicole Holliday**, click **Properties**, type **nhollida@cpandl.com** in the **E-mail** box, and then click **OK**.
2. Repeat step 1 for Limor Henig and Stuart Railson, using the e-mail addresses for each account from the table.
3. Close the **Active Directory Users and Computers** console.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADUser -Name ADRMSSVC -GivenName ADRMSSVC -UserPrincipalName ADRMSSVC -
ChangePasswordAtLogon $false -AccountPassword (ConvertTo-SecureString -AsPlainText
"p@ssw0rd" -Force) -Enabled $true

New-ADUser -Name ADRMSADMIN -GivenName ADRMSADMIN -UserPrincipalName ADRMSADMIN -
ChangePasswordAtLogon $false -AccountPassword (ConvertTo-SecureString -AsPlainText
"p@ssw0rd" -Force) -Enabled $true

New-ADUser -Name NHOLLIDA -GivenName "Nicole Holliday" -UserPrincipalName NHOLLIDA -
EmailAddress nhollida@contoso.com -ChangePasswordAtLogon $false -AccountPassword
(ConvertTo-SecureString -AsPlainText "p@ssw0rd" -Force) -Enabled $true

New-ADUser -Name LHENIG -GivenName "Limor Henig " -UserPrincipalName LHENIG -EmailAddress
lhenig@contoso.com -ChangePasswordAtLogon $false -AccountPassword (ConvertTo-SecureString
-AsPlainText "p@ssw0rd" -Force) -Enabled $true

New-ADUser -Name SRAILSON -GivenName "Stuart Railson" -UserPrincipalName SRAILSON -
EmailAddress srailson@contoso.com -ChangePasswordAtLogon $false -AccountPassword
(ConvertTo-SecureString -AsPlainText "p@ssw0rd" -Force) -Enabled $true
```

Once you have completed the creating the additional user accounts for your AD RMS infrastructure, you will need to create some additional Active Directory groups to assign users to as well to demonstrate restricted rights and permissions in later steps as you are testing out your AD RMS installation. The following table lists the groups that you will need to create at this time.

Group Name	E-mail Address
Finance	finance@contoso.com
Marketing	marketing@contoso.com
Engineering	engineering@contoso.com
Employees	employees@contoso.com

To add new groups to Active Directory

1. In Active Directory Users and Computers console, right-click **Users**, point to **New**, and then click **Group**.
2. In the New Object – Group dialog box, type **Finance** in **Group name**, select the **Universal** option for **Group Scope**, and then click **OK**.
3. Perform the above steps 1-2 for each of the remaining groups: Marketing, Engineering, and Employees.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADGroup -Name Finance -SamAccountName Finance -GroupCategory Security -GroupScope Universal

New-ADGroup -Name Marketing -SamAccountName Marketing -GroupCategory Security -GroupScope Universal

New-ADGroup -Name Engineering -SamAccountName Engineering -GroupCategory Security -GroupScope Universal

New-ADGroup -Name Employees -SamAccountName Employees -GroupCategory Security -GroupScope Universal
```

Next, add e-mail addresses to group objects:

To add e-mail addresses to group objects

1. In Active Directory Users and Computers console, double-click **Users**, click **Finance**, and then click **Properties**.
2. Type **finance@contoso.com** in the **E-mail** box and then click **OK**.
3. Perform the above steps 1-2 for each of the remaining groups: Marketing, Engineering, and Employees.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-ADGroup Finance -Add @{mail='finance@contoso.com'}

Set-ADGroup Marketing -Add @{mail='marketing@contoso.com'}

Set-ADGroup Engineering -Add @{mail='engineering@contoso.com'}

Set-ADGroup Employees -Add @{mail='employees@contoso.com'}
```

Finally, add the user accounts to their appropriate groups. In this guide, we will add Nicole Holliday, Limor Henig, and Stuart Railson to the Employees group. Then, we will add Nicole Holliday to the Finance group, Limor Henig to the Marketing group, and finally add Stuart Railson to the Engineering group.

To add the user accounts to their respective groups, you should follow these steps:

To add user accounts to groups

1. In the **Active Directory Users and Computers** console, double-click **Users**, and then

double-click **Employees**.

2. Click **Members**, and then click **Add**.
3. Type **nhollida@contoso.com;lhenig@contoso.com;srailson@contoso.com**. and then click **OK**.
4. Perform the above steps 2 & 3 to add one member to each of the remaining groups as follows:
   - Nicole Holliday—Finance
   - Limor Henig—Marketing
   - Stuart Railson—Engineering
5. Double-click **Enterprise Admins**.
6. Click **Members**, and then click **Add**.
7. Type **adrmsadmin@contoso.com**. and then click **OK**.
8. Close the **Active Directory Users and Computers** console.
9. Log out of DC1.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-ADGroupMember -Identity Employees -Members nhollida,lhenig,srailson

Add-ADGroupMember -Identity Finance -Members nhollida

Add-ADGroupMember -Identity Marketing -Members lhenig

Add-ADGroupMember -Identity Engineering -Members srailson

Add-ADGroupMember -Identity "Enterprise Admins" -Members adrmsadmin
```

Step 2: Prepare SQL1 as a SQL Server database server for supporting AD RMS

SQL1 provides the following services:

- A SQL Server installation on a member server within corp.contoso.com domain that will host the AD RMS configuration and logging databases.

The procedures to complete the configuration of the AD RMS SQL server computer, named SQL1, include:

- Install the Operating system on SQL1 and download SQL Server 2012 trial version software.
- Install SQL Server 2012 trial version.
- Configure TCP/IP and rename computer to SQL1.
- Join the SQL1 computer to the corp.contoso.com domain.

To install the operating system and download SQL Server evaluation on to SQL1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect SQL1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

   Tip
   To access Windows Update, you can do the following: Type CTRL+ALT+DELETE and then select **Task Manager**. From Task Manager, click **More details**. From the File menu, select **Run new task**. In the Run dialog, type "**control**", select **Create this task with administrative privileges** and then click **OK**. In Control Panel, type in the search box "Windows Update" to return Windows Update. Click **Windows Update** in the search results and then update your settings to install the latest updates for Windows Server 2012.
4. Download SQL Server 2012 trial version software to the SQL1 computer.

   You can download the Microsoft SQL Server 2012 trial version software from Microsoft SQL Server 2012 Evaluation. To install the trial version later, at a minimum, you will need the following files to be downloaded at this time to a temporary directory on the SQL1 computer: SQLFULL_*architecture_language*_Lang.box, SQLFULL_*architecture_language*_Install.exe, SQLFULL_*architecture_language*_Core.box.

SQL Server 2012 will require that the .NET Framework 3.5 features are first installed. To avoid problem with installing these features later, install these features now while you have Internet access configured for the SQL1 computer.

To enable .NET Framework 3.5 Features

1. Log on to SQL1 using the local Administrator account.
2. In Server Manager, click **Local Server** in the console tree, then click **Manage**, and then select **Add Roles and Features**.
3. In Add Roles and Features wizard, click **Next**.
4. In Select Installation Type, select **Role-based or feature-based installation** and then click **Next**.
5. In Select Destination Server, verify that the local server is in the pool and that **Select a server from the server pool** is selected, and then click **Next**.
6. In Select Server Roles, click **Next**.
7. In Select Features, select **.NET Framework 3.5 Features**, and then click **Next**.
8. In Confirm Installation Options, verify that .NET Framework 3.5 Features is selected, and then click **Install**.

Next, configure the TCP/IP protocol for SQL1 with a static IP address of 10.0.0.5 and the subnet mask of 255.255.255.0.

To configure TCP/IP for SQL1

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection**.

   Note
   The link may not immediately appear. Wait for the network interfaces to be enumerated.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In IP address, type **10.0.0.5**. In **Subnet mask**, type **255.255.255.0**. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
5. Click **OK** and then close the **Wired Ethernet Properties** dialog.
6. Close the **Network Connections** window.
7. In Server Manager, click **Local Server** in the console tree. Click the link next to **Computer name** in the Properties tile.
8. On the **Computer Name** tab of the **System Properties** dialog, click **Change**.
9. In **Computer name**, type **SQL1**, click **OK** twice, and then click **Close**. When you are prompted to restart the computer, click **Restart Now**.
10. After restarting, login using the local Administrator account.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.5 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1

Rename-Computer SQL1

Restart-Computer
```

To join SQL1 to the corp.contoso.com domain

1. In **Server Manager**, click **Local Server** in the console tree. Click the link next to **Computer name** in the **Properties** tile.

2. In the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. On the Computer Name tab, under **Member of**, click **Domain**, and then type **corp.contoso.com**.
4. Click **OK**.
5. When you are prompted for a user name and password, type **User1** and its password, and then click **OK**.
6. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
7. When you are prompted that you must restart the computer, click **OK**.
8. On the **System Properties** dialog box, click **Close**.
9. When you are prompted to restart the computer, click **Restart Now**.
10. After the computer restarts, click the **Switch User arrow icon**, then click **Other User** and log on to the CORP domain with the domain **Administrator** account.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

To install SQL Server 2012 on the SQL1 computer

1. Verify that you are logged on as CORP\Administrator to the SQL1 computer.
2. Open Windows Explorer.
3. Navigate to the temporary directory (such as C:\temp) where you downloaded SQL Server 2012 installation files previously.
4. Double-click **SQLFULL_***architecture_language*_**Install.exe** to extract the SQL Server installation files.

   You will see the Unloading the Box progress dialog box as files and folder structure for setup files are being unpacked. After the files are extracted you will have a subfolder named (such as C:\temp**SQLFULL_***architecture_language*) containing the installation files for SQL Server 2012.
5. Navigate to the subfolder and double-click **Setup** to launch the SQL Server Installation Center. When prompted for administrative credentials to launch SQL Server 2012 Setup provide the current credentials for CORP\Administrator.

   In the SQL Server 2012 Installation Center, from the navigation menu on the left, click **Installation**, and then click **New SQL Server installation or add features to an existing installation**.

   SQL Server 2012 Setup will launch and run Setup Support Rules to determine that all SQL prerequisites have been met.

6. Click **OK** after the Setup Support Rules check has successfully completed.
7. In Product Key, if you are installing a trial version, under **Specify a free edition**, select **Evaluation** and then click **Next**.
8. In License Terms, read the license terms and then click the **I accept the license terms** checkbox and then click **Next**.
9. In Product Key, if you are installing a trial version, under **Specify a free edition**, select **Evaluation** and then click **Next**.
10. In Product Updates, SQL Server Setup will not be able to connect to the Windows Update service. Click **Next** to continue setup and begin installing setup files.
11. In Setup Support Rules, note the results of the rules check will typically indicate that all pre-checks have passed with warnings about Microsoft .NET Application Security and Windows Firewall. Click **Next** to continue SQL Server setup.
12. In Setup Role, select **SQL Server Feature Installation** and then click **Next**.
13. In Feature Selection, select **Database Engine Services** and **Management Tools - Basic**, accept the defaults for shared feature directory and then click **Next**.
14. In Installation Rules, verify that all rules have passed and then click **Next**.
15. In Instance Configuration, accept the **Default instance**, as well as the default values for **Instance ID** and **Instance root directory** and then click **Next**.
16. In Disk Space Requirements, review the disk space summary for the features selected is sufficient and then click **Next**.
17. In Server Configuration, accept the defaults and then click **Next**.
18. In Database Engine Configuration, accept the default authentication type (**Windows authentication**) and then for Specify SQL Server Administrators, click **Add Current User** to add CORP\Administrator to the list and then click **Next**.
19. In Error Reporting, click **Next**.
20. In Installation Configuration Rules, verify that all rules have passed and then click **Next**.
21. In Ready to Install, review installation selections and then click **Install**.
22. Click **Close** after installation has successfully completed.

Next, because ADRMSADMIN is the account you will be using as well to install and configure AD RMS in Step 2, you will want to add this account to the local Administrators group on APP1 where the AD RMS server is to be located.

To add ADRMSADMIN to the local Administrators group on APP1

1. Log on to the APP1 computer as a local administrator (APP1\Administrator).
2. In Server Manager, click **Local Server** in the console tree, then click **Tools**, and then select **Computer Management**.
3. Expand **System Tools**, expand **Local Users and Groups**, and then click **Groups**.
4. Right-click **Administrators**, click **Add to Group**, click **Add**, type **ADRMSADMIN** in **Enter the object names to select (examples)** box, and then click **OK**.
5. Click **OK**, and then close **Computer Management**.

6. Log out as APP1\Administrator.

Finally, create a shared folder on SQL1 so that other domain users can find documents saved to the network.

To create a shared network folder that can be modified by domain users

1. Verify that you are logged on as CORP\Administrator to the SQL1 computer.
2. From the server desktop, open Windows Explorer and then right-click **Local Disk (C:)**.
3. Point to **New**, and then click **Folder**.
4. Type **Public** for the new folder, and then press ENTER.
5. Right-click **Public**, point to **Share with**, and then click **Specific people**.

   The File Sharing wizard opens.
6. For **Choose people on your network to share with**, click the arrow and select **Everyone**, and then click **Add**.
7. In the list, click the arrow for **Permission Level** on the group **Everyone** and select **Read/Write**.
8. Click **Share**, then click **Done**.
9. Right-click **Public** and then click **Properties**.
10. On the Sharing tab, click **Advanced Sharing**, then click **Permissions**, and then click **Add**.
11. In Select Users, Computers, Service Accounts or Groups, type **Domain Users** and then click **Check Names**.
12. Click OK.
13. On the Share Permissions tab, verify that **Domain Users (CORP\Domain Users)** is selected in the **Group or user name** box.
14. In the Permissions for Users box select the **Full Control** check box in the **Allow** column.
15. Click **OK** twice and then **Close** to close file sharing wizard.

Next, we will want to continue working on SQL1 using SQL Server Management Studio and other administrative tools to make some configuration changes to support SQL Server access before we install AD RMS in the next step. First, the ADRMSADMIN account needs to be given SysAdmin rights on the SQL Server instance in order to be able to create the AD RMS databases during AD RMS setup.

Add the CORP\ADRMSADMIN user as a SQL login and assign it SysAdmin rights

1. While logged on to SQL1 as CORP\Administrator, press CTRL+ALT+DEL and select **Task Manager**.
2. In Task Manager, from the **File** menu, select **Run new task**, type in the following to open SQL Server Management Studio and then click **OK**:

   ```
   C:\Program Files (x86)\Microsoft SQL
   Server\110\Tools\Binn\Management Studio\Ssms.exe
   ```

3. Click **Connect** to connect to the SQL Server instance on SQL1.
4. In the console tree, expand SQL1, then expand **Security**, and then click **Logins**.
5. Right-click and select **New Login**.
6. For Login Name, click **Search**.
7. In Select User or Group, type **CORP\ADRMSADMIN** and then click **Check Names**, then click **OK**.
8. In Login - New, in the navigation pane, select **Server Roles**.
9. In the list of server roles, click the **sysadmin** checkbox.
10. Click **OK**.
11. Close SQL Server Management Studio.

Start the SQL Server Browser service on SQL1

1. While logged on to SQL1 as CORP\Administrator, in Task Manager, from the **File** menu, select **Run new task**, type in the following to open the Services console and then click **OK**:

   ```
   Services.msc
   ```

2. Scroll in the list of services and select **SQL Server Browser**.
3. Right-click and select **Properties**.
4. In SQL Server Browser Properties, for **Startup type**, select **Automatic** and then click **Apply**.
5. Click **Start** to start the SQL Server Browser service.
6. Click **OK** and then close the Services console.

Add Windows Firewall exceptions for SQL Server ports

1. While logged on to SQL1 as CORP\Administrator, in Task Manager, from the **File** menu, select **Run new task**, type the following to open the Windows Firewall with Advanced Security console and then click **OK**.

   ```
   Wf.msc
   ```

2. In the console tree, select **Inbound Rules** and then right-click and select **New Rule**.
3. For Rule Type, select **Port** and then click **Next**.
4. For Protocols and Ports, select **TCP** and then select **Specific local ports**, type in **1433** and then click **Next**.
5. For Action, select **Allow the connection** and then click **Next**.
6. For Profile, accept the defaults (**Domain**, **Private** and **Public** are all selected) and then click **Next**.
7. For Name, type **SQL_1433** and then click **Finish**.
8. Repeat steps 2-7.

   For step 4, modify Protocols and Ports to select **UDP** and for specific local ports type in **1434**. For step 7, in the Name page type **SQL_1434** to name the second rule then click

Finish.

9. Close the Windows Firewall with Advanced Security console.

Step 3: Install and configure AD RMS on APP1

To install and configure AD RMS, you must add the AD RMS server role.

Windows Server 2012 includes the option to install AD RMS as a server role through Server Manager. Both installation and configuration of AD RMS are handled through Server Manager. The first server in an AD RMS environment is the root cluster. An AD RMS root cluster is composed of one or more AD RMS servers configured in a load-balancing environment. This test lab guide will install and configure a single-server AD RMS root cluster.

To add the AD RMS server role

1. Log on to the APP1 computer as the AD RMS enterprise administrator (CORP\ADRMSADMIN).

   Note
   The ADRMSADMIN account was created for use in installing and managing the AD RMS server deployment. To ensure it has sufficient rights to accomplish its purpose, such as the ability to register the service connection point (SCP), it needs to be made a member of the Enterprise Admins group for the corp.contoso.com domain. To install the AD RMS role on APP1, the ADRMSADMIN account also needs to be added to the local Administrators group on APP1. These account and group management details are important to successfully complete the configuration of the AD RMS cluster and allow for further management of the AD RMS server.

2. In the **Dashboard** console of Server Manager, click **Add roles and features**.
3. Click **Next** three times to get to the server role selection screen.
4. In the **Select server roles** dialog, select **Active Directory Rights Management Services**, and then click **Next**.
5. When prompted to add features that are required for AD RMS, click **Add Features**.
6. In the **Select features** dialog, select **.NET Framework 3.5 Features**, and then click **Next**.

   Note
   You must install .NET Framework 3.5 prior to installing the Microsoft Report Viewer 2008 used to generate troubleshooting and system health reports on AD RMS in Windows Server 2012.

7. In **Active Directory Rights Management Services**, click **Next**.
8. In **Select role services**, verify that **Active Directory Rights Management Server** is selected, and then click **Next**.
9. Click **Install** to add the role.

10. Allow the installation to complete and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature ADRMS -IncludeManagementTools

Install-WindowsFeature NET-Framework-Core
```

In Windows Server 2012, adding the AD RMS role and configuration of a new AD RMS cluster are two separate processes. After you have completed adding the role, additional configuration is required to deploy the AD RMS role.

To configure a new AD RMS root cluster

1. In Server Manager, click the **Notifications** icon.
2. For the task event labeled **Configuration required for Active Directory Rights Management Services at APP1**, click **Perform additional configuration**.

   The AD RMS Configuration wizard opens.
3. In the configuration wizard, click **Next**.
4. For AD RMS Cluster, accept the default selection (**Create a new AD RMS root cluster**) and click **Next**.
5. For Configuration Database, accept the default selection (**Specify a database server and a database instance**) and click **Select**.
6. In Select Computer, for the object name to select, type **SQL1**, click **Check Names** and then click **OK**.
7. Click **List** to list the database instances for SQL1. After the list is populated, select **DefaultInstance** from the drop-down list options and click **Next**.
8. In Specify Service Account, click **Specify** and in the Windows Security dialog, type **ADRMSSVC** and the currently set password ("**p@ssw0rd**") and click **OK**.
9. Verify that the **Domain User Account** is set to **CORP\ADRMSSVC** and then click **Next**.
10. For Cryptographic Mode, accept the default (**Cryptographic Mode 2**) and then click **Next**.
11. For Cluster Key Storage, accept the default (**Use AD RMS centrally managed key storage**) and then click **Next**.
12. For Cluster Key Password, type and confirm a password ("p@ssw0rd") and then click **Next**.
13. For Cluster Web Site, accept the default (**Default Web Site**) and then click **Next**.
14. For Cluster Address, accept the default (**Use an SSL-encrypted connection (https://)**), for **Fully Qualified Domain Name** type in **app1.corp.contoso.com** and then click **Next**.
15. For Server Certificate, accept the default (**Create a self-signed certificate for SSL encryption**) and then click **Next**.

Tip

When using a self-signed certificate for the cluster, you can put a copy of that certificate in the Trusted Root Certification Authorities store so that it will be trusted. A copy can also be put in that same certificates store on the client computer so that the web site is trusted.

16. For Licensor Certificate, accept the default name (**APP1**) and then click **Next**.
17. For SCP Registration, accept the default (**Register the SCP now**) and then click **Next**.
18. For Confirmation, review your installation selections and then click **Install**.
19. Click **Close**.
20. Log off the server, and then log on again to update the security token of the logged-on user account.

    The user account that is logged on when the AD RMS server role is installed is automatically made a member of the AD RMS Enterprise Administrators local group. A user must be a member of that group to administer AD RMS.

Your AD RMS root cluster is now installed and configured.

Once you have completed logging in again, you can further manage AD RMS using the Active Directory Rights Management Services console.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Import-Module ADRMS

New-PSDrive -PSProvider ADRMSInstall -Name RC -Root RootCluster

# Set the AD RMS service account

$adrmssvc = Get-Credential

Set-ItemProperty -Path RC:\ -Name ServiceAccount -Value $svcacct

# Configure the AD RMS server to use the SQL1 server and the default instance

Set-ItemProperty -Path RC:\ClusterDatabase -Name ServerName -Value SQL1

Set-ItemProperty -Path RC:\ClusterDatabase -Name InstanceName -Value MSSQLSERVER

# Securely store the cluster key password string in a variable and assign it to your AD
RMS installation

$password = Read-Host -AsSecureString -Prompt "Password:"

Set-ItemProperty -Path RC:\ClusterKey -Name CentrallyManagedPassword -Value $password

# Set the AD RMS cluster address

Set-ItemProperty -Path RC:\ -Name ClusterURL -Value "http://app1.corp.contoso.com:80"

# Set the SLC name for your AD RMS installation

Set-ItemProperty -Path RC:\ -Name SLCName -Value "APP1"
```

```
# Register the SCP connection point

Set-ItemProperty -Path RC:\ -Name RegisterSCP -Value $true

# Install the AD RMS root cluster using the settings provided

Set-Location RC:\

Install-ADRMS –Path.
```

To open the Active Directory Rights Management Services console

1. In Server Manager, click **Local Server** in the console tree, then click **Tools**, and then click **Active Directory Rights Management Services**.
2. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Continue**.

From the console, you can configure trust policies, configure exclusion policies, and create rights policy templates.

Step 4: Verify AD RMS configuration after you complete the configuration

The AD RMS Client 1.0 is included in the default installation of Windows 8. Previous versions of the client are available for download for a number of earlier versions of the Windows operating system. For more information, see the AD RMS Client Requirements.

Before you can consume rights-protected content, you must add the AD RMS cluster URL to the Local Intranet security zone.

Add the AD RMS cluster URL to the Local Intranet security zone for all users who will be consuming rights-protected content.

To add AD RMS cluster to Local Intranet security zone

1. Log on to CLIENT1 as Nicole Holliday (CORP\NHOLLIDA).
2. From the Taskbar, click **Internet Explorer**.
3. Click **Tools**, and then click **Internet Options**.
4. Click the **Security** tab, click **Local intranet**, and then click **Sites**.
5. Click **Advanced**.
6. In the **Add this website to the zone**, type **https://app1.corp.contoso.com**, and then click **Add**.
7. Click **Close**.

Tip

You can now verify access to the AD RMS licensing site by typing the URL

(https://app1.corp.contoso.com) in the Address bar in Internet Explorer. You should also see a warning about the certificates for this site. That is because of the use of a self-signed certificate when AD RMS was configured. In live deployments, it is recommended that you use a signed certificate issued from a trusted Internet issuing certification authority (CA).

8. Repeat steps 1–7 for Stuart Railson and Limor Henig.

To verify the functionality of the AD RMS deployment, you will log on as Nicole Holliday and then restrict permissions on a Microsoft Word 2010 document so that members of the Engineering group are able to read the document but unable to change, print, or copy. You will then log on as Stuart Railson, verifying that the proper permission to read the document has been granted, and nothing else. Then, you will log on as Limor Henig. Since Limor is not a member of the Engineering group, he should not be able to consume the rights-protected file.

Note

In this test lab guide, when a user restricts permissions on a document or attempts to open a restricted document, a warning appears that informs you that the certificate issuer for the AD RMS Web site is unknown or untrusted. This warning results from using a self-signed certificate instead of a certificate issued by a recognized certification authority. When you receive this warning, click **Yes** to continue.

To restrict permissions on a Microsoft Word document

1. Log on to CLIENT1 as Nicole Holliday (CORP\NHOLLIDA).
2. Open Microsoft Office Word 2010.
3. Type **Engineering employees can read this document, but they cannot change, print, or copy it on the blank document page.**
4. From the File menu, click **Protect Document**, then point to **Restrict Permission by People**, and then click **Restricted Access**.

   When prompted that this page requires secure connection which includes server authentication, click **Yes**.
5. In Permissions, click the **Restrict permission to this document** check box, and then in the **Read** box, type **engineering@corp.contoso.com**.
6. Click **OK** to close the Permission dialog box.
7. From the File menu, click **Save As**, and then save the file as **\\SQL1\Public\ADRMS-TST.docx**.
8. Log off as Nicole Holliday.

Next, log on as Stuart Railson and open the document, ADRMS-TST.docx.

To view a rights-protected document

1. Log on to CLIENT1 as Stuart Railson (CORP\SRAILSON).
2. Open Windows Explorer and browse to \\SQL1\Public. Double-click ADRMS-TST.docx to

open it in Microsoft Word 2010.

When prompted that this page requires secure connection which includes server authentication, click **Yes**.

3. Note that the following message appears: "Permission to this document is currently restricted. Microsoft Office must connect to https://app1.corp.contoso.com:443/_wmcs/licensing to verify your credentials and download your permission."
4. Click **OK**.

   Note that the following message appears: "Verifying your credentials for opening content with restricted permissions…"
5. When the document opens, click the File menu. Notice that the **Print** option is not available.
6. Close Microsoft Word.
7. Log off as Stuart Railson.

Finally, log on as Limor Henig and verify that he is not able to consume the rights-protected file.

To attempt to view a rights-protected document

1. Log on to CLIENT1 as Limor Henig (CORP\LHENIG).
2. Open Windows Explorer and browse to \\SQL1\Public. Double-click ADRMS-TST.docx to open it in Microsoft Word 2010.

   When prompted that this page requires secure connection which includes server authentication, click **Yes**.
3. Note that the following message appears: "Permission to this document is currently restricted. Microsoft Office must connect to https://app1.corp.contoso.com:443/_wmcs/licensing to verify your credentials and download your permission."
4. Click **OK**.
5. The following message appears: "You do not have credentials that allow you to open this document. You can request updated permission from nhollida@contoso.com. Do you want to request updated permission?"
6. Click **No** and then close Microsoft Word.

You have successfully deployed and demonstrated the functionality of AD RMS, using the simple scenario of applying restricted permissions to a Microsoft Word 2010 document. You can also use this deployment to explore some of the additional capabilities of AD RMS through additional configuration and testing.

How AD RMS Works

Active Directory Rights Management Services (AD RMS) encompasses all of the server and client technologies that are required to support information protection through the use of rights management in an organization.

When you use an AD RMS infrastructure, you are able to able to protect the information in your organization using the following client and server components to both publish and consume rights-protected content.

- AD RMS clients, which request licenses and enforce assigned rights protection at the document level to files and messages.
- AD RMS servers, which administers account certification, licensing and publishing services that use Active Directory as well as assisting clients in locating these services.

Note

AD RMS is the name for what was originally introduced as a standalone release called Rights Management Services (RMS) 1.0 which was first released by Microsoft in 2005 for the Windows Server 2003 server platform.

How AD RMS clients work

When publishing content, AD RMS clients request and acquire new licenses for protecting content according to the usage rights and conditions that you as a publisher choose to allow for the content that you wish to protect.

When a document is authored and rights protection is chosen, the AD RMS client acquires a Client Licensor Certificate (CLC), which enables it to protect content. It then uses this CLC to encrypt the document, create and sign a Publishing License (PL) and then binds a copy of the PL to the encrypted content. This helps the content to be better protected from misuse even if it is shared to others within your organization or even to others outside of your organization.

When others receive the rights-protected content, in order to access and make use of it they will first need to use a rights-enabled application (such as Microsoft Office) to request and acquire an end-user license for the content. In order to obtain the end-user license, the AD RMS client must first determine if the recipient of the content conforms to any policies set forth in the publishing license that was used to protect the content. If the AD RMS client determines the user is eligible to access the content, the AD RMS client ensures that the user honors the conditions indicated in the end-use license, which might restrict certain actions. This ensures documents are protected as intended by authors and publishers and are only consumed by recipients according to the assigned rights policies.

How AD RMS servers work

AD RMS servers are implemented as a set of Web service components that run on Microsoft Internet Information Services (IIS) and work in connection with Microsoft SQL Server and Active Directory Domain Services (AD DS).

The various components that make up an AD RMS server are listed in the following table.

Server Component	Description
Administration web service	The AD RMS server computer hosts this web service, which is used to manage AD RMS through the use of either the AD RMS administration console or Windows PowerShell commands for AD RMS.
Account certification	AD RMS servers generate rights account certificates (RACs) that associate users with specific computers.
Licensing	AD RMS servers issue end-user licenses. An end-user license enables AD RMS client-enabled applications to access protected content within the user restrictions set by the content publisher.
Publishing	AD RMS servers also create client licensor certificates that enable content publishers to define the policies that can be enumerated in an end-user license.
Precertification	Enables a server to request a rights account certificate on behalf of a user in order for Exchange to pre-license content to Outlook users.
Service locator	Provides the URL of the account certification, licensing, and publishing services to Active Directory so that they can be discovered by AD RMS clients.

In an AD RMS cluster, all AD RMS servers are one of two types.

- **Root certification servers.** The first AD RMS server in an Active Directory forest assumes this role. There can only be one root certification server in each Active Directory forest.
- **Licensing servers.** This is the role taken on by any additional or secondary AD RMS servers added to provide independent policy options to certain groups within an Active Directory forest.

For more information on how servers work within an AD RMS infrastructure, see Server Configuration for an AD RMS Infrastructure.

Working with Templates in AD RMS

To assist in easily assigning and managing rights protection to content, rights policy templates can be used. Rights policy templates can help you to scope and create a specific set of rights that you can apply or repurpose as often as needed for protecting content. For example, a template could be created and named "Company Confidential" that where applied to protect a document would allow employees only the right to view the document, but not to forward, copy, or save it.

Rights policy templates that are used within AD RMS are stored and made available from an AD RMS server to client computers through the following example deployment workflow:

1. Bob is an administrator who chooses to make a new rights policy template for clients to use in protecting content. He uses the AD RMS console to create the template, which is stored in the AD RMS configuration database. (Through optional configuration, the AD RMS cluster also maintains a copy of all rights policy templates in a folder that Bob has specified.)
2. Bob has configured Group Policy to enable all corporate desktop computers to use AD RMS automatic rights policy template distribution. He has also modified the update frequency setting to cause AD RMS client computers to query the template distribution pipeline and update their templates within the next 5 days.

   Note
   Bob could also have used System Center Configuration Manager or manually copied the templates himself to client computers as alternative to using Group Policy.
3. Once AD RMS client computers have queried the template pipeline and have updated their local cached folder of AD RMS templates with the new policy template, users can select it and apply the template to any new documents they want to author and protect using it.

By default, new rights policy templates are distributed to all AD RMS clients but starting with Windows Server® 2008, administrators have the option to control by using the AD RMS console whether templates are archived or remain actively distributed for use. If a template is archived it is not distributed to clients but the AD RMS cluster will still be able to generate end user licenses for rights-protected content that has a publishing license generated from that template. Templates should not typically be deleted as any content protected by a template will no longer be accessible if the template is deleted from the AD RMS configuration.

For more information on templates, see AD RMS Policy Templates.

A simplified view of AD RMS in action

The following sections help to walk you through a demonstration of how AD RMS works as the process of publishing, sharing and consuming a rights-protected document occurs within first a single organization, and then as its occurs between two different organizations.

How AD RMS works within an organization

An AD RMS client and server can be used within your organization to help protect content through assigning rights protection as it is published and shared with others. This section discusses how a basic AD RMS service infrastructure (which includes an AD RMS server, a SQL Server computer hosting the AD RMS databases and Active Directory domain controllers) can be used to support rights protection.

Background

Terry Adams is a senior process engineer who has been working on site at Contoso Pharmaceuticals manufacturing facility in Redmond, Washington for the past four weeks. During that time, Terry has observed production line activities being carried out and been meeting privately with employees at the facility.

Terry has completed his process review and has been asked by his manager, Diane Margheim, to share a confidential copy of his report with the plant manager, Lola Jacobson. In order to ensure the report remains confidential, Diane has asked Terry to apply rights protection to his report document after he is finished writing it in Microsoft Word before forwarding it on to Lola to allow her preview it for comment before it is passed along to the executive leadership.

In the following graphic, we see how AD RMS works using servers and clients to support this user scenario. Terry is an information author working at his corporate desktop and using Microsoft Word to create and prepare his report.

Terry chooses the option to rights protect the document which allows him to select both the people and the level of access those people will have to his content. He grants read-only access to Lola as the information recipient for the document. This will enable Lola to view the report but deny her the ability to change, copy or print the document.

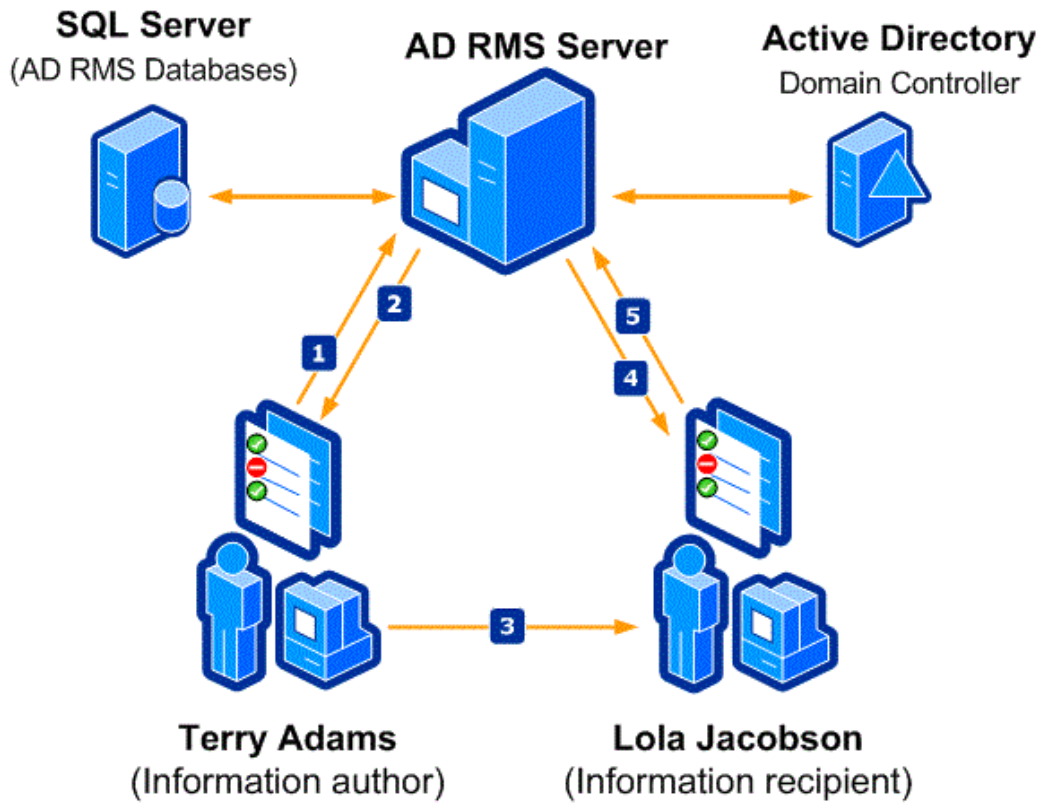

1. As Terry applies his access restrictions, the AD RMS client launches and initiates a service request on his behalf to the Contoso AD RMS server.
2. The AD RMS server for Contoso returns a Client Licensor Certificate to the AD RMS client installed at Terry's desktop, which enables him to save the document in encrypted form with the desired level of rights-protection.
3. Terry then attaches and sends the rights-protected report Word document to Lola in an email.
4. Lola receives Terry's email and saves the attached document to her local desktop and then opens it. When she does, the AD RMS client working at her desktop contacts the Contoso AD RMS server to acquire an end-user license.
5. The AD RMS client at Lola's desktop receives back the end user license, which indicates that she is permitted to view the document. The AD RMS client then decrypts the document and applies the appropriate restrictions to the enable Lola to access the content according to the access permissions that Terry assigned to it.

How AD RMS works between organizations

While AD RMS clients and servers can be used within a single organization or Active Directory forest to help protect content, you can also with some additional configurations allow for published content to be protected across organizations or forest boundaries.

In order to use AD RMS across organizational and forest boundaries, a level of trust must be established across the organizations or across forests. There are several ways in which this can occur:

- At the most basic level, Trusted User Domains (TUDs) can help allow AD RMS to process requests from users located across Active Directory forests. A simple TUD relationship between two AD RMS clusters can be established quickly and can be a lightweight and low-cost method of managing AD RMS trust between two separate organizations. For more information, see Trusted User Domains.
- Where there is a need to acquire licenses from an RMS cluster for content protected with another RMS cluster, Trusted Publishing Domains (TPDs) offer another means of enabling cross-server trust that go beyond the ability of TUDs. This option is rarely used across organizations. For more information, see Trusted Publishing Domains.
- For organizations that have invested in the federated identity capabilities of Active Directory Federation Services (AD FS), you can also leverage the power of existing federated trusts to make AD RMS work across forest boundaries. For more information, see Active Directory Federation Services with AD RMS and Federating AD RMS.

For a more detailed understanding of how AD RMS works across organizational lines, see Sharing Documents with External Organizations.

AD RMS Infrastructure Deployment Tips

Depending upon the complexity of your organization and its Active Directory deployment, deploying AD RMS can have a number of aspects to it that might require additional consideration. This topic discusses the core best practices that are common to all AD RMS deployments.

Centralize servers as much as possible

When you are setting up AD RMS within a new organization, the primary guidance to consider regarding AD RMS is that you should always try to minimize the number of AD RMS servers required. As much as possible, you should also try to centralize AD RMS servers and have them in as few locations as is necessary for your deployment.

An AD RMS infrastructure does not require a great deal of network and processor overhead to complete its day-to-day certification and licensing operations but it does need to contain and store data that must be consistently available for it to service clients properly. The more complexity in your AD RMS topology, the more complex it is to troubleshoot issues where service operations are not completing as expected.

Depending on the complexity of your organization and the scale of your Active Directory, some additional AD RMS servers might be necessary and unavoidable. This happens because as a fundamental rule of design, you must have one AD RMS root cluster for each Active Directory forest in a deployment. Therefore, if you have a need to support AD RMS in multi-forest deployment, you will need to have at least one AD RMS root cluster for each of the forests involved.

Use a single cluster if possible

If you have a single Active Directory forest for your organization, try to use only a single AD RMS cluster to support all of your deployment needs. In general, deploying fewer clusters is the better and more robust approach to use when planning and designing for AD RMS.

For example, suppose that you are seeking scalability and performance in your AD RMS deployment and have four server computers to dedicate and use to support your AD RMS deployment. In general, you will have better performance if you use all four server computers in a single root cluster that supports both certification and licensing then if you were to create two clusters, one purposed for certification and the other configured to act as a licensing-only cluster.

This happens because you can better use load balancing across a single certification cluster that is doing both licensing and certification. Therefore, a central server cluster scales better than using subordinate licensing servers because you cannot load balance a certification cluster with a licensing cluster.

It is also generally recommended that you have AD RMS licensing centralized in one cluster to improve administration. Licensing clusters are not effective to distribute load as AD RMS clients will always try to contact the licensing cluster that provided them the license originally unless reconfigured. Multiple licensing clusters also add complexity, making other administrative operations such as reporting, monitoring and troubleshooting more difficult to perform.

Use licensing-only cluster only if required

AD RMS licensing-only clusters can be useful in some circumstances for segmenting administration to the branch or department level based on the following business requirements.

- **When departments need to have independent licensing due to legal or regulatory concerns.**

  For example, in some organizations which have offices located within multiple countries, there could be legal requirements that documents encrypted and protected within one country must not be encrypted or forwarded using keys obtained from a server located or managed outside of the borders of that country. In these circumstances, a licensing-only cluster could be required to fulfill legal compliancy requirements.

- **When departments have poor connectivity and generate and consume content generally in an isolated manner.**

  In some organizations, network connectivity and content usage patterns could be decisive factors in the decision to implement a new licensing cluster. Where network connectivity is poor and 90% or more of the content that is to be published and consumed using information rights management (IRM) features occurs within the same geographic site location, deploying an additional licensing-only cluster to service local users and computers makes sense and should be considered.

Centralize licensing from multiple forests in one central cluster

When your organization has an Active Directory infrastructure that spans multiple forests, you should use a load-balanced cluster for AD RMS certification in one forest to serve publishing and licensing requests for the entire enterprise. This action simplifies administration tasks and minimizes troubleshooting work when all publication licenses come from the same source. You can then configure the licensing URLs in the different clusters so they are all the same and point to the same “central” cluster, or deploy registry keys to users in the other forests to point them to this central cluster.

Finally, you can use Trusted User Domains (TUDs) to integrate the certification clusters in the other forests with the central cluster that is doing the servicing of licensing and publishing requests. This will make the AD RMS clusters in the other forests only active for user certification and group expansion of distribution lists and ensures all the content is protected with the same keys and on the same cluster.

For more insights into how to deploy AD RMS in a multi-forest environment, see AD RMS Performance Insight from Microsoft IT's Implementation or for full details, download the technical deployment white paper, "Deploying Active Directory Rights Management Services at Microsoft" from the Microsoft Download Center.

Using AD RMS with Hardware Security Modules

Active Directory Rights Management Services (AD RMS) is an information protection platform that enables documents to be shared with protection so that only authorized users are allowed to perform specific actions on them. In order to ensure that content remains protected, AD RMS utilizes cryptographic technologies that require the generation and storage of public and private keys. Key management can affect both server performance as well as potentially raise concerns about the security of the private keys involved as any loss of such keys could cause the protection benefits offered by AD RMS to be lost. While AD RMS provides strong password-based protection for its server keys, this option relies on having good operational procedures and tight management of the servers and passwords to maintain security of the keys, constrains which in some organizations might not offer enough guarantees.

One way these implications can be addressed is by the use of hardware security modules (HSMs), which are third-party products that can be used in connection with your existing AD RMS infrastructure. By offloading key management tasks and some cryptographic processing from your AD RMS servers to an HSM, you can reduce processing overhead for each AD RMS server and also ensure that any private keys it uses are not exposed or stolen if your AD RMS deployment should ever be breached or compromised administratively.

In this guide

This guide is intended to provide additional general guidance on using HSMs with AD RMS. It also contains some instructions for extending a base setup of an AD RMS test lab to include HSMs.

To make full use of the lab setup instructions for setting up and configuring HSMs with AD RMS, you will need to have first completed the following additional test lab guides:

- Windows Server 2012 Base Configuration Test Lab Guide (TLG)
- Test Lab Guide: Deploying an AD RMS Cluster

How AD RMS Benefits From HSMs

AD RMS is an information protection platform that enables authors and publishers to define usage policies on their documents, establish who can access those documents and what operations can be performed on them. To achieve this type of document protection, the content protected by AD RMS is encrypted by supporting applications using 128-bit or 256-bit AES (Advanced Encryption Services) symmetric encryption to provide fast, strong protection.

In order to best understand how AD RMS benefits from the use of hardware security modules it helps to have a deeper understanding of how AD RMS uses AES encryption and the encryption keys required to support it.

How AD RMS Does Encryption and Uses Encryption Keys

Symmetric key encryption uses the same key to encrypt and decrypt content. But as the decryption keys for the content have to be made available to the clients that need to consume it, AD RMS uses asymmetric encryption with private/public key pairs to protect the encryption keys. All AD RMS servers, client computers, and user accounts have a public and private pair of 1,024-bit or 2,048-bit RSA–based encryption keys (depending on deployment options). AD RMS uses its private key to encrypt the content encryption symmetric key along with the rights policy data in the publishing license and the use license in each document. AD RMS also uses the private keys to digitally sign AD RMS certificates and licenses, ensuring certificates are not tampered with at the time users use them to request access to protected information.

How HSMs Help Improve Server Performance

While the symmetric encryption used by AD RMS clients to encrypt and decrypt the protected content is relatively fast and efficient, asymmetric encryption and signing are much heavier processes. Modern server CPUs provide very powerful number crunching capabilities, but in AD RMS server operations, under heavy loads, the cryptographic processing might become the factor that most limits performance.

HSMs help by offloading the cryptographic processing to a dedicated hardware unit that is capable of completing the processing of cryptographic algorithms at higher speed than most

current CPUs. This enables AD RMS servers to process licensing requests with reduced CPU effort.

Note

While offloading cryptographic operations to an HSM might improve performance, the network latency between the AD RMS servers and the HSM might lead to reduced performance. Whether an HSM affects performance positively or negatively will depend on the performance of your server CPU, the speed of your HSM, the network latency and speed between the HSM and the AD RMS server and which cryptographic mode your AD RMS deployment is operating under (with those who are running under cryptographic mode 2 seeing the maximum benefit).

How HSMs Help to Ensure the Security of Private Keys

In private/public cryptography the private key component is the most critical component. In AD RMS, the private key is used to decrypt all the content protection keys, effectively allowing whoever has access to that key to obtain access to all the documents protected by that AD RMS cluster. The private key also enables the AD RMS cluster to sign licenses and certificates that will be then trusted by all its clients and servers. So access to the private key would also grant the ability to issue licenses in the name of the AD RMS cluster. This all implies that the private key of the AD RMS cluster must be safeguarded and protected in order for the AD RMS service to provide meaningful protection.

With AD RMS in the default configuration, the private keys shared by AD RMS servers within an AD RMS cluster are stored securely through the Data Protection API (DPAPI) using the Secure Storage service. When AD RMS servers are joined in an AD RMS cluster, they all share a single root cluster key, which is stored in the configuration database in encrypted form so this key is never transmitted in a way that is accessible by external third parties. Because the security of your entire AD RMS deployment depends on the security of this key it is of central importance to protect this key as much as possible. Otherwise, if a user were to obtain a copy of this key they could access all content previously published and protected without authorization and to independently issue licenses that are valid AD RMS licenses under the name of the AD RMS cluster.

Therefore, HSMs can also provide a security advantage for AD RMS server deployments by isolating direct access to private keys, which in an HSM configuration are always securely issued and stored within the hardware device itself.

HSMs provide additional protection to the private keys by performing all the cryptographic operations internally without ever releasing the private key to any user or system, including the AD RMS servers. The private key is always stored inside the HSM in a form that is logically and physically protected from unauthorized access. The HSMs even provide physical protection against tampering, by using devices that will destroy the private key when detecting an attempt to physically breach the device or tamper with its built-in security protections.

Recommendations and best practices for using HSMs with AD RMS

While HSMs can provide significant additional protection as well as enhanced performance to AD RMS servers, they have some implications for the implementation and operation of an AD RMS infrastructure. In this section, we will discuss those concerns, and provide AD RMS deployment best practices for working with HSMs within your AD RMS design.

We will first look at best practices for deploying HSMs and then look at any concerns applicable to other processes such as decommissioning either an AD RMS server or an HSM device. After that, we will take a brief look at any recommendations for operational guidance, such as those best practices that apply to the ongoing management and maintenance of HSMs within your AD RMS deployment.

Implementation guidance

This section covers implementation guidance.

Deployment

When working with an HSM and clustered AD RMS servers, the following guidance applies:

- While it is possible to use a local (i.e. either PCI or USB-based) HSM on each physical server within a cluster, it is generally recommended to use a network HSM with a shared key that all servers can access for ease and efficiency. The following, however, are some circumstances where PCI-based HSMs are preferable:
    - A network HSM requires that the network between the AD RMS servers and the network HSM is trusted and that nobody can spoof the IP address of the AD RMS servers. If this is not possible, then a PCI-based HSM card would be a better option.
    - If you cannot assign static IP addresses to your AD RMS servers, then a network HSM will not be able to work and a PCI-based HSM would be your only choice.
    - If you have high AD RMS traffic that requires many AD RMS servers, you might need multiple network-based HSMs to accommodate the load. While it is possible to use multiple network-based HSMs to service a single AD RMS deployment, there can come a point where the ease of working with and adding network-based HSMs has to be contrasted against the optimal performance that is possible if PCI cards are used.
- The first server for an AD RMS cluster must be installed with the drivers for the HSM product and the cryptographic service provider (CSP) before that server is used to create the cluster.
- Each additional server for an AD RMS cluster must be installed with the drivers for the HSM product, the cryptographic service provider (CSP), and a copy of the cluster's key space identification before that server is joined to an AD RMS cluster. Managing the key space identification details can vary depending on the specifics for each HSM vendor, For example, in Thales nCipher products, this data is referred to as the "security world" for each device. In Safenet products, you clone all HSMs to the same "Domain". Refer to the documentation from your HSM vendor for more details.

For licensing-only AD RMS clusters, no special considerations are required as the process of working with HSMs is the same as it is for certification clusters. The licensor certificate itself is signed by the root cluster, but keys are generated in the same manner as is done for a certification cluster regardless of whether an HSM is involved in the process.

When using HSMs to support AD RMS in virtualized environments, the only factor to consider is that while physical AD RMS clusters spread over physical computers can share the same security world using either local or network HSMs, virtualized clusters can only use network HSMs as virtual machines do not work with local HSM devices installed in the host computer.

Configuration

In general, there are no special configuration settings that you need to apply when configuring AD RMS to work with an HSM device. In some cases, Microsoft support personnel have seen issues when customers using Thales nCipher products have selected to set up the security worlds for their HSM using the Operator Cards for managing a specific key setting on the HSM rather than using the Module Protection setting. If this setting is used, the HSM can prompt for the card and password during the AD RMS installation and that prompt could be suppressed. If the prompt is suppressed and the AD RMS installer does not respond, the installation might time out.

Trust

When configuring trust with AD RMS there are some differences when implementing Trusted Publishing Domains (TPDs) if an HSM is involved. A TPD is utilized to provide an AD RMS cluster with the ability to issue licenses to consume content that was protected by using another AD RMS cluster. A TPD exchange basically involves the sharing of the cluster private key, along with other supporting information such as rights policy templates.

First of all, a TPD with an HSM-based key does not actually include the cluster key, but instead references a pointer to the key in the HSM. Therefore when you export or import a TPD, you might need to export or import the HSM key separately from the source server, creating a key identification (known as a “security world” in nCipher HSMs or a "domain" in Safenet HSM products) that is equivalent to the original key so the TPD will work as expected on the target server. This has advantages in scenarios where different parties need to exchange TPDs but might want to "recall" them later if the partnership ends. By deploying the TPD with an HSM it controls to another party, taking back the HSM effectively removes access to that party to all content protected with that TPD.

Another advantage of exchanging HSM-protected TPDs is that, by decoupling the exchange of the TPD from the exchange of the private key, renewing information on rights policy templates after a TPD has been shared can be done by just re-exporting and re-importing the TPD, without exporting, storing, transporting or importing the private key, making this potentially recurrent exchange a less critical operation.

This also assumes that the cluster key can be exported in the first place from the HSM. HSM vendors, however, typically allow you to move a wrapped key between HSMs of the same type (assuming you have the right to do so).

In comparison, a trusted user domain is basically an exchange of public keys and server certificates, which do not include the private key. Therefore, the export and import of a TUD, which is normally utilized to allow users to license content from RMS clusters outside of their own forest, is done in an HSM environment in exactly the same way as in a non-HSM environment.

Decommissioning

The process for decommissioning an AD RMS server where an HSM does not differ much from decommissioning in a non-HSM environment. One additional step is that you must remove the IP address (for Thales nCipher HSMs) or the certificate (for Safenet HSMs) from all your HSMs to prevent future abuse. Also, if you are replacing an existing AD RMS cluster, the process of exporting and importing the HSM key identification settings (known as "security world" for Thales products, "domain" for Safenet products) might include additional steps such as the ones described in the previous section on trust.

Important
Be sure to erase the key identification settings (i.e. "security world" for Thales products, "Domain" for Safenet products) from any decommissioned HSMs before disposing of them and destroy (or safely archive) all administrator cards.

Operations guidance

This section covers implementation guidance for management and maintenance of HSMs to support AD RMS.

Management

The primary consideration for ensuring your HSMs are able to support your AD RMS installation well is to ensure protection of your cluster key. Refer to the documentation provided with your HSM for discussions of how to implement key protection such as by using key cards or other mechanisms as per the manufacturer recommendations.

Maintenance

As far as routine maintenance of HSMs supporting an installed AD RMS cluster, the cryptographic service provider (CSP) configuration should be backed up or documented as necessary. Also test that you have a proven backup and restore capability and a backup of your security key space (for nCipher products, this is called the "security world" which includes all settings under %nfast_kmdata%\local on each server in the cluster) and place any administrator and/or operator cards in a secure location and record their passwords).

Tip
To provide additional support for recovery of your HSM-stored keys, use k of N quorum for the administrator cards, and have two full quorums worth of administrator cards (plus spares) stored at two different secure locations.

Disaster recovery

Backing up an AD RMS cluster for disaster recovery involves backing up the AD RMS configuration database and the AD RMS cluster key.

In a software-based cluster key scenario, the AD RMS database contains the cluster key, and to recover the AD RMS cluster after a failure or outage you only need the AD RMS cluster key protection password, which you also will need to add new nodes to the cluster if the old nodes are lost or destroyed.

In an HSM-based scenario you also need to be able to restore access to the HSM. If the HSM is lost, you will need to recover it from an HSM backup, which is a process that's highly dependent on the HSM implementation. If the HSM is functional, you need to have a copy of the keys used to unlock the key space used for your AD RMS cluster in the HSM (in nCipher HSMs it's called a "security world"). The keys can take different forms depending on the HSM and the security options required, ranging from a password to a smartcard or multiple smartcards. So, as you can see, the process of recovering access to the HSM keys themselves is also completely dependent on the HSM implementation. But the logic is common: you always need to recover the configuration database, install and configure in the new node the cryptographic provider used to access the HSM, and then provide the data necessary to unlock access to the keys used by AD RMS. For more specific information on this process, see the documentation provided with your HSM.

AD RMS Best Practices Guide

A well-considered plan for designing your Active Directory Rights Management Services (AD RMS) installations makes enterprise-scale rights management deployment straightforward and easy to manage and troubleshoot.

This guide combines business and technical guidance to minimize the time and effort required to implement the AD RMS. It provides a set of checklist for specific prescriptive guidance based on best practices learned from customers that have already deployed AD RMS in their organizations and as verified by those who have in-depth experience supporting such installations.

It also provides a summary of all the main design choice tasks and decisions you need to develop a solid plan for AD RMS in your environment. The intended audience for this guide is the IT professional responsible for testing, piloting, and rolling out an AD RMS design.

AD RMS Installation Best Practices

The following are best practices to consider when preparing to install AD RMS within your organization.

Best practice rule	Notes
Use dedicated AD RMS servers.	Installing AD RMS on the same server as a

Best practice rule	Notes
	domain controller, Microsoft Exchange Server, Certification Authority, or Microsoft Office SharePoint Server is a poor security practice.
Do not install AD RMS on a domain controller.	If you do install AD RMS on a domain controller, you must add the AD RMS service account, which is normally configured with no additional permissions, to the Domain Admins group.
Do not try to install the Identity Federation Support feature until you have an Active Directory Federation Services (AD FS) server in place and be aware of case sensitivity in AD FS.	You cannot install the Identity Federation Support feature until you have an AD FS server in place. If AD FS is not yet configured in your environment at the time of AD RMS installation, you can wait until it is before attempting to add the Identity Federation Support feature. If installing Identity Federation Support, use lower case letters for the fully qualified domain name, as AD FS is case sensitive.
You should only use Windows Internal Database in a test environment.	Windows Internal Database does not support remote connections; therefore, you would be unable to add additional AD RMS servers to your cluster. In a production environment, to allow for future growth and AD RMS cluster expansion you should therefore always use Microsoft SQL Server to host AD RMS database services.
Use DNS aliases, such as CNAME records, or DNS host records, such as A Records for the fully qualified domain name of the AD RMS cluster.	This allows you to easily add additional servers to the AD RMS cluster and allows you to load balance and perform disaster recovery very easily.
Use DNS aliases, such as CNAME records, or DNS host records, such as A Records for your AD RMS database server.	This makes future migration of your AD RMS databases much easier.
Be sure that any protocol bindings and other settings that are needed by AD RMS for your Web server or its websites are configured before you install and configure AD RMS.	If you plan to deploy AD RMS on a website that is already set up, be sure that website has an HTTP binding, even if you are configuring AD RMS to use HTTPS. If you plan to deploy AD RMS on a non-default website, install the IIS 6 Management

Best practice rule	Notes
	Capability role service before you start provisioning. Using SSL protocol increases the security of the connections to the AD RMS cluster. Also, SSL is required to integrate AD RMS with AD FS. Remember that this cannot be changed once it has been specified.
Configure an extranet URL even if you do not plan to deploy one initially.	You should configure your extranet URL at the time of installation, even if it will not be initially deployed. If external access is enabled after documents are AD RMS protected you must remove the protection, remove the DRM folder on the client computers, configure extranet access, and then protect the documents again.
Use an SSL certificate issued from a trusted certification authority (CA) for the Internet.	You should use self-signed certificates only in a test environment. In a production environment you should always use an SSL certificate issued from a certification authority.
Perform additional tasks to complete configuration and ready your AD RMS installation for use.	After an installation or upgrade is complete you must log off and log back in again before you can administer AD RMS using the AD RMS console. Once installation is complete you should also back up your Server Licensor Certificate and your private key by exporting the Trusted Publishing Domain (TPD) for your cluster and storing it in a safe place. Also, be sure to keep the password that was used during the TPD export documented in and stored separately in a safe location.
Be aware of the conditions and differences when moving an existing RMS installation to AD RMS.	There are two paths to upgrading an earlier version of RMS to AD RMS: migration and in-place upgrade. Migration is the recommended process. If you choose to do an in-place upgrade, be sure to run the upgrade wizard after the operating system upgrade completes. This wizard is launched from a link in Server Manager. For more information on migrating or upgrading an AD RMS cluster see RMS to AD RMS Migration and Upgrade Guide.

Best practice rule	Notes
Always deploy AD RMS using clusters.	If necessary, you can use a single node cluster, but deploying AD RMS in cluster mode provides allows for later scalability and expansion should you ever need it. For best practices on working with AD RMS clusters, see the following section.
Use FQDNs for all URLs associated with your AD RMS deployment.	Using a single label domain can cause problems with some applications, increase the probability of spoofing attacks and cause interoperability problems with other organizations.
Use a common FQDN for both intranet and extranet access to your AD RMS installation.	Unless you have specific reasons for differentiating between the two URLs you should make the two URLs identical.

AD RMS Clustering Best Practices

The following are best practices to consider when working with AD RMS clusters within your organization.

Best practice rule	Notes
Always deploy AD RMS using clusters	If necessary, use a single node cluster and if possible, use more than one AD RMS server to support a deployment and provide additional redundancy and service availability.
Be aware of special considerations when selecting a load balancing solutions to support AD RMS, especially when working in virtualized environments.	In general, network load balancers (NLBs) or external load balancers are good alternatives with the following exception: If you are deploying AD RMS in a virtualized environment (a common or likely occurrence), an external load balancer will work better than an NLB solution. This happens because typically there will be specific operations or exceptions such as ARP registrations, port spanning that can require virtualization administrators to configure manually exceptions that override the default policy in those environments. This type of customization can

Best practice rule	Notes
	be highly undesirable to manage in large-scale virtualization deployments.
Always use virtual names for AD RMS clusters	Select a generalized name that does not refer to any specific or physical server computer. For example, something like "rms.contoso.com" is a good format to use here. Then create entries in your DNS zone that helps to direct traffic to that name such as by adding a CNAME resource record that points this name to your NLB host name and an A resource record for the NLB that directly points to its IP address. Also, be sure to use fully qualified domain names (FQDNs) and not NetBIOS names to point to your AD RMS cluster.
If possible use a single IP address for your AD RMS cluster for both intranet and extranet clients.	In general, there are no good reasons to use multiple IP addresses or to have your AD RMS server.
Always use HTTPS/SSL when configuring your AD RMS cluster URL.	Use properly issued certificates that you obtain from a recognized Internet certification authority (CA) that are matched to your cluster URL address. You can then use Internet Information Services (IIS) Manager to create certificate requests and to export your issued certificate to other nodes in your AD RMS cluster.
Only use self-signed certificates in a test environment to evaluate AD RMS.	Do not use self-signed certificates in production deployment as they can cause many additional deployment problems such as the need to manually copy and reconfigure trust for the certificate on each server node in your cluster.
If necessary, also deploy the PKI Root certificate to the trusted root containers on all AD RMS clients and servers.	If you configured SSL with a certificate issued by a private CA, store the PKI root certificate in all nodes named "Trusted Root Certification Authorities" container.
Deploy one AD RMS cluster per Active Directory forest.	Each forest where users of the solution will be hosted needs to have its own AD RMS certification cluster. Do not deploy more than one AD RMS certification cluster per forest and do not deploy Licensing Only cluster unless

Best practice rule	Notes
	there is a specific need for them.
Try to concentrate AD RMS licensing as much as possible.	If possible, point all licensing URLs in all your clusters to match the URL of one of the clusters in a forest that trusts all the other forests. That way, all licensing operations will be performed by only one cluster which simplifies logging, reporting, template configuration and troubleshooting.

AD RMS Database Server Best Practices

The following are best practices to consider when working with AD RMS database services within your organization.

Best practice rule	Notes
Consider thoughtfully whether failover clustering or log shipping will be sufficient for your AD RMS database design.	When you are designing your AD RMS database deployment, you need to consider the option to use a high availability design for your database server. If you are deploying AD RMS on Windows Server 2008 or Windows Server 2008 R2 and you are not using Exchange Prelicensing, a highly available database is not essential and a robust backup plan might be sufficient. But if deploying on Windows Server 2012 or if using Exchange Prelicensing or any other Exchange IRM integration capability, a highly available database is essential. Log shipping and failover clustering both can work well, with log shipping providing recoverability against a wider range of possible failure scenarios, and a cluster providing faster recovery times in most scenarios. Both can be combined in a single installation if you consider that your service's availability needs justify it. For more information, see **AD RMS and Database Design**.
Always use virtual names for AD RMS databases.	The reason for this is similar to the reason to use virtual names with the AD RMS cluster. It is

Best practice rule	Notes
	much more problematic or difficult to change the name of an AD RMS database once you have deployed it widely within an organizational infrastructure. Use CNAME records to support DNS name resolution for your AD RMS databases. Either CNAME or A resource records are good in case you need to move or relocate your AD RMS database servers. Also, as with AD RMS cluster virtual naming, use fully qualified domain names (FQDNs) and point the virtual names for your AD RMS database cluster to the cluster IP address. You can also use different CNAME resource records for different AD RMS databases, even if they are residing on the same server computer. This tends to make things easier if in the future you need to move or relocate databases to another server or move a server to another location on your network. You can also if you prefer use a different CNAME or A resource record for each AD RMS database.

AD RMS Client Deployment Best Practices

The following are best practices to consider when deploying AD RMS clients within your organization.

Best practice rule	Notes
With legacy AD RMS clients, always deploy the AD RMS client binaries without including additional settings in the installation package.	Do not put registry settings or other local client configuration changes in the installation package as these configuration updates can cause supportability issues. For example, you might want to change registry key settings when you re-deploy the client and in that case you would have to successfully install or re-install the client package to troubleshoot an issue. In many cases, registry settings are not updated well using this type of bundled configuration and it is more reliable to use

Best practice rule	Notes
	separate GPO updates or Windows scripting to revise such settings.
Create a group in Active Directory and use that for targeting AD RMS client deployment	It is recommended that you use GPO to deploy AD RMS client settings and that you only deploy settings as needed. Target settings using the same groups used for client deployment.
Always deploy the required settings for the AD RMS clients.	The following settings should always be configured when deploying AD RMS clients: • Root CA certificate (if needed) • AD RMS cluster URLs in the Trust Sites for Internet Explorer (note that you can use GPO to do configure this setting). • Enabling or disabling IRM functionality (you can set this either on or off depending on user group) • AD RMS template path and deployment script or settings • Text to show to non-AD RMS ready clients For more information on deploying client registry settings see either **AD RMS Client 1.0 Registry Settings** or as appropriate, see **AD RMS Client 2.0 Deployment Notes** if you are working with the AD RMS Client 2.0.

AD RMS Rights Policy Templates Best Practices

Once you have deployed your AD RMS infrastructure, you will want to ensure that you follow best practices for using rights policy templates within your organization. Rights policy templates enable content authors to quickly apply a standard level of protection for content across your organization. Furthermore, templates offer additional security options that are not available in normal protection.

For more information on how to best deploy AD RMS rights policy templates to support your deployment, see AD RMS Rights Policy Templates Best Practices.

AD RMS Performance and Logging Best Practices

Once you have deployed your AD RMS infrastructure, you will also want to ensure that it meets the necessary performance requirements for your organization and can handle expected loads as

well as monitor for errors, bottlenecks and other potential problems. To do this, adequate logging of your operations is essential to provide the necessary information to monitor performance and aid in troubleshooting issues as they arise.

For more information in how to best monitor and optimize your AD RMS deployment for performance and logging, see **AD RMS Performance and Logging Best Practices**.

Application Server Overview

Application Server provides an integrated environment for deploying and running custom, server-based business applications.

For more information, see http://technet.microsoft.com/library/cc754024(v=WS.10).aspx.

Failover Clustering Overview

This topic provides an overview of the Failover Clustering feature in Windows Server 2012 R2 and Windows Server 2012. Failover clusters provide high availability and scalability to many server workloads. These include server applications such as Microsoft Exchange Server, Hyper-V, Microsoft SQL Server, and file servers. The server applications can run on physical servers or virtual machines. This topic describes the Failover Clustering feature and provides links to additional guidance about creating, configuring, and managing failover clusters that can scale to 64 physical nodes and to 8,000 virtual machines.

Did you mean…

- Building Your Cloud Infrastructure: Scenario Overview
- Scale-Out File Server for Application Data Overview
- **Hyper-V Replica**
- Network Load Balancing Overview

Feature description

A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service.

You can manage failover clusters by using the Failover Cluster Manager snap-in and the Failover Clustering Windows PowerShell cmdlets. You can also use the tools in File and Storage Services to manage file shares on file server clusters.

Practical applications

- Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines
- Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V

New and changed functionality

New and changed functionality in Failover Clustering supports increased scalability, easier management, faster failover, and more flexible architectures for failover clusters.

For information about Failover Clustering functionality that is new or changed in Windows Server 2012 R2, see What's New in Failover Clustering in Windows Server 2012 R2.

For information about Failover Clustering functionality that is new or changed in Windows Server 2012, see What's New in Failover Clustering in Windows Server 2012.

Hardware requirements

A failover cluster solution must meet the following hardware requirements:

- Hardware components in the failover cluster solution must meet the qualifications for the Certified for Windows Server 2012 logo.
- Storage must be attached to the nodes in the cluster, if the solution is using shared storage.
- Device controllers or appropriate adapters for the storage can be Serial Attached SCSI (SAS), Fibre Channel, Fibre Channel over Ethernet (FcoE), or iSCSI.
- The complete cluster configuration (servers, network, and storage) must pass all tests in the Validate a Configuration Wizard.

Note
In the network infrastructure that connects your cluster nodes, avoid having single points of failure.

For more information about hardware compatibility, see the Windows Server Catalog.

For more information about the correct configuration of the servers, network, and storage for a failover cluster, see the following topics:

- Failover Clustering Hardware Requirements and Storage Options
- Validate Hardware for a Failover Cluster

Software requirements

- You can use the Failover Clustering feature on all editions of Windows Server 2012 R2 and Windows Server 2012. This includes Server Core installations.
- You must follow the hardware manufacturers' recommendations for firmware updates and software updates. Usually, this means that the latest firmware and software updates have been applied. Occasionally, a manufacturer might recommend specific updates other than the latest updates.

Server Manager information

In Server Manager, use the Add Roles and Features Wizard to add the Failover Clustering feature. The Failover Clustering Tools include the Failover Cluster Manager snap-in, the Failover Clustering Windows PowerShell cmdlets, the Cluster-Aware Updating (CAU) user interface and Windows PowerShell cmdlets, and related tools. For general information about installing features, see Install or Uninstall Roles, Role Services, or Features.

To open Failover Cluster Manager in Server Manager, click **Tools**, and then click **Failover Cluster Manager**.

See also

The following table provides additional resources about the Failover Clustering feature in Windows Server 2012 R2 and Windows Server 2012. Additionally, see the content on failover clusters in the Windows Server 2008 R2 Technical Library.

Content type	References
Product evaluation	What's New in Failover Clustering in Windows Server 2012 R2 | What's New in Failover Clustering in Windows Server 2012 | Microsoft Server and Cloud Platform | Cloud and Datacenter Solutions Hub
Planning	Failover Clustering Hardware Requirements and Storage Options | Validate Hardware for a Failover Cluster | Network Recommendations for a Hyper-V Cluster in Windows Server 2012
Deployment	Installing the Failover Clustering Feature and Tools | Validate Hardware for a Failover Cluster | Prestage Cluster Computer Objects in Active Directory Domain Services | Creating a Windows Server 2012 Failover Cluster | Deploy Hyper-V over SMB | Deploy Scale-Out File Server | iSCSI Target Block Storage, How

Content type	References
	To | Install and Configure High Availability Printing | Deploy an Active Directory-Detached Cluster | Deploy a Hyper-V Cluster | Using Guest Clustering for High Availability | Deploy a Guest Cluster Using a Shared Virtual Hard Disk | Building Your Cloud Infrastructure: Scenario Overview
Operations	Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster | Use Cluster Shared Volumes in a Failover Cluster | Cluster-Aware Updating Overview
Tools and settings	Failover Clustering Windows PowerShell Cmdlets
Community resources	High Availability (Clustering) Forum | Failover Clustering and Network Load Balancing Team Blog
Related technologies	Hyper-V Overview | File and Storage Services Overview | Storage Spaces Overview | Scale-Out File Server for Application Data Overview | Network Load Balancing Overview | Microsoft SQL Server | Microsoft Server and Cloud Platform | System Center Virtual Machine Manager

What's New in Failover Clustering in Windows Server 2012 R2

This topic describes the Failover Clustering functionality that is new or changed in Windows Server 2012 R2.

Feature description

Failover clusters provide high availability and scalability to many server workloads. These include server applications such as Microsoft Exchange Server, Hyper-V, Microsoft SQL Server, and file servers. The server applications can run on physical servers or virtual machines. In a failover cluster, if one or more of the clustered servers (nodes) fails, other nodes begin to provide service.

This process is known as failover. For an overview of the Failover Clustering feature, see Failover Clustering Overview.

New and changed functionality

The following table and sections summarize key Failover Clustering functionality in Windows Server 2012 R2 that is new or changed since Windows Server 2012.

Feature/Functionality	New or Improved	Description
Shared virtual hard disk (for guest clusters)	New	Enables you to use .vhdx files as shared storage in a guest cluster.
Virtual machine drain on shutdown	New	Enables a Hyper-V host to automatically live migrate running virtual machines if the computer is shut down.
Virtual machine network health detection	New	Enables a Hyper-V host to automatically live migrate virtual machines if a network disconnection occurs on a protected virtual network.
Optimized CSV placement policies	Improved	Distributes CSV ownership evenly across the failover cluster nodes.
Increased CSV resiliency	Improved	Multiple Server service instances per cluster node and CSV monitoring of the Server service provide greater resiliency.
CSV cache allocation	Improved	Increases the amount of RAM that you can allocate as CSV cache.
CSV diagnosibility	Improved	Enables you to view the state of a CSV on a per node basis, and the reason for I/O redirection.
CSV interoperability	Improved	Adds CSV support for other Windows Server 2012 R2 features.

Feature/Functionality	New or Improved	Description
Deploy an Active Directory-detached cluster	New	Enables you to deploy a failover cluster with less dependency on Active Directory Domain Services.
Dynamic witness	New	Dynamically adjusts the witness vote based on the number of voting nodes in current cluster membership.
Quorum user interface improvements	Improved	Enables you to easily view the assigned quorum vote and the current quorum vote for each node in Failover Cluster Manager.
Force quorum resiliency	New	Enables automatic recovery in the case of a partitioned failover cluster.
Tie breaker for 50% node split	New	Enables one side of a cluster to continue to run in the case of a cluster split where neither side would normally have quorum.
Configure the Global Update Manager mode	New	Helps the cluster to continue to function if there is a delay with one or more nodes.
Cluster node health detection	Improved	Increases the resiliency to temporary network failures for virtual machines that are running on a Hyper-V cluster.
Turn off IPsec encryption for inter-node cluster communication	New	Helps prevent a cluster from being affected by high latency Group Policy updates.
Cluster dashboard	New	Provides a convenient way to check the health of all managed failover clusters in Failover Cluster Manager.

Highly available virtual machine improvements

The following section provides a summary of new highly available virtual machine functionality in Windows Server 2012 R2.

Shared virtual hard disk (for guest clusters)

You can now share a virtual hard disk file (in the .vhdx file format) between multiple virtual machines. You can use these .vhdx files as shared storage for a virtual machine failover cluster, also known as a *guest cluster*. For example, you can create shared .vhdx files for data disks and for the disk witness. (You would not use a shared .vhdx file for the operating system virtual hard disk.)

What value does this change add?

In Windows Server 2012, you could deploy guest clusters using shared storage that was provided by virtual Fibre Channel or iSCSI to the guest operating system. In these configurations, the underlying storage was exposed to the user of a virtual machine. In private or public cloud deployments, there is often the need to hide the details of the underlying fabric from the user or tenant administrator. Shared .vhdx storage provides that layer of abstraction.

This change also enables easier deployment of guest cluster configurations. A shared .vhdx file configuration is easier to deploy than solutions like virtual Fibre Channel or iSCSI. When you configure a virtual machine to use a shared .vhdx file, you do not have to make storage configuration changes such as zoning and LUN masking.

What works differently?

In Windows Server 2012, shared virtual hard disks are not available. Windows Server 2012 R2 adds this functionality.

In Windows Server 2012 R2, virtual SCSI disks now appear as virtual SAS disks when you add a SCSI hard disk to a virtual machine. This includes both shared and non-shared virtual hard disk files. For example, if you view the disk in Server Manager, the bus type is listed as SAS.

For more information about shared virtual hard disks, see Virtual Hard Disk Sharing Overview and Deploy a Guest Cluster Using a Shared Virtual Hard Disk.

Virtual machine drain on shutdown

In Windows Server 2012 R2, if you shut down a Hyper-V failover cluster node without first putting the node into maintenance mode to drain any running clustered roles, the cluster now automatically live migrates all running virtual machines to another host before the computer shuts down.

Note

Make sure that there is not more than one virtual machine in a virtual machine clustered role. Starting with Windows Server 2012, we do not support this configuration. An example of this scenario is where multiple virtual machines have files on a common physical disk that is not part of Cluster Shared Volumes. A single virtual machine per

clustered role improves the management experience and the functionality of virtual machines in a clustered environment, such as virtual machine mobility.

What value does this change add?

This change provides a safety mechanism to help ensure that a server shutdown (or any action that shuts down the Cluster service) does not cause unplanned downtime for running virtual machines. This increases the availability of applications that run within the guest operating system.

Important

We still recommend that you put a node into maintenance mode or move all virtual machines to other nodes before you shut down a cluster node. This is the safest way to drain any running clustered roles.

What works differently?

In Windows Server 2012, if you shut down a cluster node without first draining the node, the virtual machines are put into a saved state, and then moved to other nodes and resumed. This means that there is an interruption to the availability of the virtual machines. If it takes too long to save state the virtual machines, they may be turned off, and then restarted on another node. In Windows Server 2012 R2, the cluster automatically live migrates all running virtual machines before shutdown.

To enable or disable this functionality, configure the **DrainOnShutdown** cluster common property. By default, this property is enabled (set to a value of "1").

To view the property value, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).DrainOnShutdown
```

Virtual machine network health detection

Network health detection and recovery is now available at the virtual machine level for a Hyper-V host cluster. If a network disconnection occurs on a protected virtual network, the cluster live migrates the affected virtual machines to a host where that external virtual network is available. For this to occur there must be multiple network paths between cluster nodes.

Note

If there are no available networks that connect to other nodes of the cluster, the cluster removes the node from cluster membership, transfers ownership of the virtual machine files, and then restarts the virtual machines on another node.

What value does this change add?

This change increases the availability of virtual machines when there is a network issue. If live migration occurs, there is no downtime because live migration maintains the session state of the virtual machine.

What works differently?

In Windows Server 2012, if there is a network disconnection at the virtual machine level, the virtual machine continues to run on that computer even though the virtual machine may not be available to users.

In Windows Server 2012 R2, there is now a **Protected network** check box in the virtual machine settings. This setting is available in the advanced features of the network adapter. By default, the setting is enabled. You can configure this setting on a per network basis for each virtual machine. Therefore, if there is a lower priority network such as one used for test or for backup, you can choose not to live migrate the virtual machine if those networks experience a network disconnection.

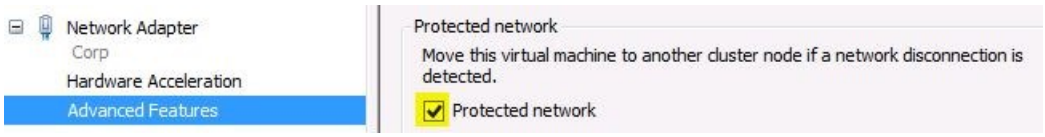

Figure 1. Protected network setting

Cluster Shared Volume (CSV) improvements

The following section provides a summary of new CSV functionality in Windows Server 2012 R2.

Optimized CSV placement policies

CSV ownership is now automatically distributed and rebalanced across the failover cluster nodes.

What value does this change add?

In a failover cluster, one node is considered the owner or "coordinator node" for a CSV. The coordinator node owns the physical disk resource that is associated with a logical unit (LUN). All I/O operations that are specific to the file system are through the coordinator node. Distributed CSV ownership increases disk performance because it helps to load balance the disk I/O.

Because CSV ownership is now balanced across the cluster nodes, one node will not own a disproportionate number of CSVs. Therefore, if a node fails, the transition of CSV ownership to another node is potentially more efficient.

This functionality is useful for a Scale-Out File Server that uses storage spaces because it ensures that storage spaces ownership is distributed.

What works differently?

In Windows Server 2012, there is no automatic rebalancing of coordinator node assignment. For example, all LUNs could be owned by the same node. In Windows Server 2012 R2, CSV ownership is evenly distributed across the failover cluster nodes based on the number of CSVs that each node owns.

Additionally in Windows Server 2012 R2, ownership is automatically rebalanced when there are conditions such as a CSV failover, a node rejoins the cluster, you add a new node to the cluster, you restart a cluster node, or you start the failover cluster after it has been shut down.

Increased CSV resiliency

Windows Server 2012 R2 includes the following improvements to increase CSV resiliency:

- Multiple Server service instances per failover cluster node. There is the default instance that handles incoming traffic from Server Message Block (SMB) clients that access regular file shares, and a second CSV instance that handles only inter-node CSV traffic. This inter-node traffic consists of metadata access and redirected I/O traffic.
- CSV health monitoring of the Server service

What value does this change add?

A CSV uses SMB as a transport for I/O forwarding between the nodes in the cluster, and for the orchestration of metadata updates. If the Server service becomes unhealthy, this can impact I/O performance and the ability to access storage. Because a cluster node now has multiple Server service instances, this provides greater resiliency for a CSV if there is an issue with the default instance. Additionally, this change improves the scalability of inter-node SMB traffic between CSV nodes.

If the Server service becomes unhealthy, it can impact the ability of the CSV coordinator node to accept I/O requests from other nodes and to perform the orchestration of metadata updates. In Windows Server 2012 R2, if the Server service becomes unhealthy on a node, CSV ownership automatically transitions to another node to ensure greater resiliency.

What works differently?

In Windows Server 2012, there was only one instance of the Server service per node. Also, there was no monitoring of the Server service.

CSV cache allocation

In Windows Server 2012 R2, you can now allocate a higher percentage of the total physical memory to the CSV cache. CSV cache enables the server to use system memory as a write-through cache.

What value does this change add?

Increasing the CSV cache limit is especially useful for Scale-Out File Server scenarios. Because Scale-Out File Servers are not typically memory constrained, you can accomplish large performance gains by using the extra memory for the CSV cache.

Tip

We recommend that you enable the CSV cache for all clustered Hyper-V and Scale-Out File Server deployments, with greater allocation for a Scale-Out File Server deployment.

What works differently?

In Windows Server 2012, you could allocate only 20% of the total physical RAM to the CSV cache. You can now allocate up to 80%.

In Windows Server 2012, the CSV cache was disabled by default. In Windows Server 2012 R2, it is enabled by default. Also, the name of the private property of the cluster Physical Disk resource has been changed from **CsvEnableBlockCache** to **EnableBlockCache**.

You must still allocate the size of the block cache to reserve. To do this, set the value of the **BlockCacheSize** cluster common property. (The name of this property was changed from **SharedVolumeBlockCacheSizeInMB** in Windows Server 2012.) For more information, see Enable the CSV cache for read-intensive workloads.

CSV diagnosibility

You can now view the state of a CSV on a per node basis. For example, you can see whether I/O is direct or redirected, or whether the CSV is unavailable. If a CSV is in I/O redirected mode, you can also view the reason.

What value does this change add?

This change enables you to optimize your cluster configuration because you can easily determine the state of a CSV.

What works differently?

You can use the new Windows PowerShell cmdlet **Get-ClusterSharedVolumeState** to view the state information (such as direct or redirected) and the redirection reason. For the state information, see the **StateInfo** property. For the I/O redirection reason, see the **FileSystemRedirectedIOReason** property and the **BlockRedirectedIOReason** property.

CSV interoperability

In Windows Server 2012 R2, CSV functionality has been enhanced to include support for the following features:

- Resilient File System (ReFS)
- Deduplication
- Parity storage spaces
- Tiered storage spaces
- Storage Spaces write-back caching

What value does this change add?

This added support expands the scenarios in which you can use CSVs, and enables you to take advantage of the efficiencies that are introduced in these features.

What works differently?

ReFS, deduplication, and parity storage spaces were not supported by CSVs in Windows Server 2012. Tiered storage spaces and Storage Spaces write-back caching are new in Windows Server 2012 R2.

Deploy an Active Directory-detached cluster

In Windows Server 2012 R2, you can deploy a failover cluster without dependencies in Active Directory Domain Services (AD DS) for network names. This is referred to as an *Active Directory-detached cluster*. When you deploy a cluster by using this method, the cluster network name (also known as the *administrative access point*) and network names for any clustered roles with

client access points are registered in Domain Name System (DNS). However, no computer objects are created for the cluster in AD DS. This includes both the computer object for the cluster itself (also known as the cluster name object or CNO), and computer objects for any clustered roles that would typically have client access points in AD DS (also known as virtual computer objects or VCOs).

Note
The cluster nodes must still be joined to an Active Directory domain.

What value does this change add?

With this deployment method, you can create a failover cluster without the previously required permissions to create computer objects in AD DS or the need to request that an Active Directory administrator pre-stages the computer objects in AD DS. Also, you do not have to manage and maintain the cluster computer objects for the cluster. For example, you can avoid the possible issue where an Active Directory administrator accidentally deletes the cluster computer object, which impacts the availability of cluster workloads.

What works differently?

The option to create an Active Directory-detached cluster is not available in Windows Server 2012. In Windows Server 2012, you can only deploy a failover cluster where the network names for the cluster are in both DNS and AD DS.

An Active Directory-detached cluster uses Kerberos authentication for intra-cluster communication. However, when authentication against the cluster network name is required, the cluster uses NTLM authentication.

Important
We do not recommend this deployment method for any scenario that requires Kerberos authentication.

To deploy this type of cluster, you must use Windows PowerShell. For deployment information and details about what is supported and not supported with this deployment method, see Deploy an Active Directory-Detached Cluster.

Quorum improvements

The following section provides a summary of improvements to cluster quorum functionality in Windows Server 2012 R2.

Dynamic witness

In Windows Server 2012 R2, if the cluster is configured to use dynamic quorum (the default), the witness vote is also dynamically adjusted based on the number of voting nodes in current cluster membership. If there are an odd number of votes, the quorum witness does not have a vote. If there is an even number of votes, the quorum witness has a vote.

The quorum witness vote is also dynamically adjusted based on the state of the witness resource. If the witness resource is offline or failed, the cluster sets the witness vote to "0."

What value does this change add?

Dynamic witness significantly reduces the risk that the cluster will go down because of witness failure. The cluster decides whether to use the witness vote based on the number of voting nodes that are available in the cluster.

This change also greatly simplifies quorum witness configuration. You no longer have to determine whether to configure a quorum witness because the recommendation in Windows Server 2012 R2 is to always configure a quorum witness. The cluster automatically determines when to use it.

Important
In Windows Server 2012 R2, we recommend that you always configure a quorum witness.

What works differently?

In Windows Server 2012, you had to determine when to configure a witness, and had to manually adjust the quorum configuration if node membership changed to keep the total number of votes at an odd number. This included every time you added or evicted cluster nodes.

Now, you no longer need to manually adjust the quorum configuration if node membership changes. By default, the cluster determines quorum management options, including the quorum witness.

Windows Server 2012 R2 also includes the new **WitnessDynamicWeight** cluster common property that you can use to view the quorum witness vote.

To view the property value, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).WitnessDynamicWeight
```

A value of "0" indicates that the witness does not have a vote. A value of "1" indicates that the witness has a vote.

Quorum user interface improvements

In Windows Server 2012 R2, you can now view the assigned quorum vote and the current quorum vote for each cluster node in the Failover Cluster Manager user interface (UI). Also, the quorum mode terminology has been simplified.

What value does this change add?

You can now easily determine in the UI which nodes have a vote, and whether that vote is active. When you click **Nodes** in Failover Cluster Manager, you can see the vote assignments.

Figure 2. Node vote assignment

What works differently?

In Windows Server 2012, you had to run the Validate Quorum Configuration validation report or use Windows PowerShell to view the vote status. You can still use these methods in Windows Server 2012 R2.

In Windows Server 2012 R2, the Validate Quorum Configuration report and the parameters for the **Set-ClusterQuorum** Windows PowerShell cmdlet are simplified to no longer use quorum mode terminology such as node majority (no witness), node majority with witness (disk or file share) or no majority (disk witness only).

Force quorum resiliency

In Windows Server 2012 R2, after there is an issue where you manually force quorum to start the cluster, for example you use the **/fq** switch when you start the Cluster service, the cluster automatically detects any partitions when connectivity is restored. The partition that you started with force quorum is now deemed authoritative. When failover cluster communication is resumed, the partitioned nodes automatically restart the Cluster service, and rejoin the cluster. The cluster is brought back into a single view of membership.

What value does this change add?

This change enables automatic recovery in the case of a partitioned failover cluster where a subset of nodes was started by forcing quorum. A partitioned failover cluster is also known as a split cluster or a "split-brain" cluster.

Note

A partitioned cluster occurs when a cluster breaks into subsets that are not aware of each other. For example, you have a multi-site cluster with three nodes in one site, and two nodes in the other. A network issue disrupts cluster communication. The site with three nodes stays running because it has quorum. The two node site without quorum shuts down. You determine that the site with three nodes does not have external connectivity, while the two node site does. Therefore, to restore service to users, you use the **/fq** switch to start the two node site. When network connectivity is restored, you have a partitioned cluster.

What works differently?

If there is a partitioned cluster in Windows Server 2012, after connectivity is restored, you must manually restart any partitioned nodes that are not part of the forced quorum subset with the **/pq** switch to prevent quorum. Ideally, you should do this as quickly as possible.

In Windows Server 2012 R2, both sides have a view of cluster membership and they will automatically reconcile when connectivity is restored. The side that you started with force quorum is deemed authoritative and the partitioned nodes automatically restart with the **/pq** switch to prevent quorum.

Tie breaker for 50% node split

As an enhancement to dynamic quorum functionality, a cluster can now dynamically adjust a running node's vote to keep the total number of votes at an odd number. This functionality works seamlessly with dynamic witness. To maintain an odd number of votes, a cluster will first adjust the quorum witness vote through dynamic witness. However, if a quorum witness is not available, the cluster can adjust a node's vote. For example:

1. You have a six node cluster with a file share witness. The cluster stretches across two sites with three nodes in each site. The cluster has a total of seven votes.
2. The file share witness fails. Because the cluster uses dynamic witness, the cluster automatically removes the witness vote. The cluster now has a total of six votes.
3. To maintain an odd number of votes, the cluster randomly picks a node to remove its quorum vote. One site now has two votes, and the other site has three.
4. A network issue disrupts communication between the two sites. Therefore, the cluster is evenly split into two sets of three nodes each. The partition in the site with two votes goes down. The partition in the site with three votes continues to function.

In addition to this automatic functionality, there is a new cluster common property that you can use to determine which site survives if there is a 50% node split where neither site has quorum. Instead of the cluster randomly picking a node to remove its quorum vote, you can set the **LowerQuorumPriorityNodeID** property to predetermine which node will have its vote removed.

What value does this change add?

With this functionality, one side of the cluster continues to run in the case of a 50% node split where neither side would normally have quorum.

By optionally setting the **LowerQuorumPriorityNodeID** property, you can control which side stays up in this scenario. For example, you can specify that the primary site stays running and that a disaster recovery site shuts down.

What works differently?

In Windows Server 2012, if there is a 50% split where neither site has quorum, both sides will go down.

In Windows Server 2012 R2, you can assign the **LowerQuorumPriorityNodeID** cluster common property to a cluster node in the secondary site so that the primary site stays running. Set this property on only one node in the site.

To set the property, start Windows PowerShell as an administrator, and then enter the following command, where "*1*" is the example node ID for a node in the site that you consider less critical:

```
(Get-Cluster).LowerQuorumPriorityNodeID = 1
```

Tip

To determine a node ID, start Windows PowerShell as an administrator, and then enter the following command, where "*Node1*" represents the name of a cluster node:

```
(Get-ClusterNode –Name "Node1").Id
```

You can also use the following command to return all the cluster node names, node IDs and the node state:

```
Get-ClusterNode | ft
```

Configure the Global Update Manager mode

When a state change occurs such as a cluster resource is taken offline, the nodes in a failover cluster must be notified of the change and acknowledge it before the cluster commits the change to the database. The Global Update Manager is responsible for managing these cluster database updates. In Windows Server 2012 R2, you can configure how the cluster manages global updates. By default, the Global Update Manager uses the following modes for failover cluster workloads in Windows Server 2012 R2:

- **All (write) and Local (read)**. In this mode, all cluster nodes must receive and process the update before the cluster considers the change committed. When a database read request occurs, the cluster reads the data from the cluster database on the local node. In this case, the local read is expected to be consistent because all nodes receive and process the updates. This is the default setting for all workloads besides Hyper-V.

  Note
  This is how global updates work for all workloads in Windows Server 2012.

- **Majority (read and write)**. In this new mode, a majority of the running cluster nodes must receive and process the update before the cluster commits the change to the database. When a database read request occurs, the cluster compares the latest timestamp from a majority of the running nodes, and uses the data with the latest timestamp. This is the default setting for Hyper-V failover clusters.

Note
There is also a new "Majority (write) and Local (read)" mode. However, this mode is not used by default for any workloads. See the "What works differently" section for more information.

What value does this change add?

The new configuration modes for Global Update Manager significantly improve cluster database performance in scenarios where there is significant network latency between the cluster nodes, for example with a stretch multi-site cluster. By association, this increases the performance of cluster workloads such as SQL Server or Exchange Server in these scenarios. Without this feature, the cluster database performs at the pace of the slowest node.

The new configuration modes can also help if there are delays that are associated with software or hardware issues. For example, a local registry update may be delayed on a node that has a hardware issue. By using a Global Update Manager mode that performs updates that are based on a majority of nodes, the cluster does not have to wait for all nodes to be notified of and acknowledge the state change before it is ready to process the next transaction.

What works differently?

In Windows Server 2012, you cannot configure the Global Update Manager mode. For all cluster workloads in Windows Server 2012, all cluster nodes must receive and process the update before the cluster considers the change committed. In Windows Server 2012 R2, you can configure the

Global Update Manager mode, with three possible values. In Windows Server 2012 R2, the majority (read and write) mode is now the default mode for Hyper-V failover clusters.

You can configure the Global Update Manager mode by using the new **DatabaseReadWriteMode** cluster common property. To view the Global Update Manager mode, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).DatabaseReadWriteMode
```

The following table shows the possible values.

Value	Description
0 = All (write) and Local (read)	• Default setting in Windows Server 2012 R2 for all workloads besides Hyper-V. • All cluster nodes must receive and process the update before the cluster commits a change to the database. • Database reads occur on the local node. Because the database is consistent on all nodes, there is no risk of out of date or "stale" data.
1 = Majority (read and write)	• Default setting in Windows Server 2012 R2 for Hyper-V failover clusters. • A majority of the cluster nodes must receive and process the update before the cluster commits the change to the database. • For a database read, the cluster compares the latest timestamp from a majority of the running nodes, and uses the data with the latest timestamp.
2 = Majority (write) and Local (read)	• A majority of the cluster nodes must receive and process the update before the cluster commits the change to the database. • Database reads occur on the local node. Because the cluster does not compare the latest timestamp on a majority of nodes, the data may be out of date or "stale."

Warning

Do not use either of the majority modes (1 or 2) for scenarios that require strong consistency guarantees from the cluster database. For example, do not use these modes

for a Microsoft SQL Server failover cluster that uses AlwaysOn availability groups, or for Microsoft Exchange Server.

Cluster node health detection

By default, cluster nodes exchange heartbeats every one second. The number of heartbeats that can be missed before failover occurs is known as the *heartbeat threshold*. In Windows Server 2012 R2, the default heartbeat threshold has been increased for Hyper-V failover clusters.

What value does this change add?

This change provides increased resiliency to temporary network failures for virtual machines that are running on a Hyper-V cluster. For example, you may not want the cluster to perform recovery actions if your network often experiences short term network interruptions. For an application that is running on a virtual machine, a short network failure is often fairly seamless because of the TCP reconnect window.

What works differently?

By default, in Windows Server 2012, a node is considered down if it does not respond within five seconds. In Windows Server 2012 R2, the default threshold value for a Hyper-V failover cluster has been increased to 10 seconds for cluster nodes in the same subnet, and 20 seconds for cluster nodes in different subnets.

The following table lists the default values in Windows Server 2012 R2.

Cluster Common Property	Default for All Clustered Roles Except Hyper-V	Default for Hyper-V Clustered Role
SameSubnetThreshold	5 seconds	10 seconds
CrossSubnetThreshold	5 seconds	20 seconds

Important

We recommend the following:

- Do not configure a heartbeat threshold value that is greater than 20 seconds because this can exceed the TCP time-out window. A value greater than 20 seconds does not increase availability and it impacts the time it takes other nodes of the cluster to detect that the node is down.
- Do not increase the heartbeat threshold from its default value for a Scale-Out File Server. A Scale-Out File Server already provides seamless recovery. If a node is considered down, the node fails over and recovers the SMB sessions on another node. This recovery action needs to occur within the SMB session time-out period. If you increase the heartbeat threshold, node detection will take longer. If it exceeds the SMB session time-out values, the recovery may no longer be seamless.

Turn off IPsec encryption for inter-node cluster communication

In Windows Server 2012 R2, you can now turn off Internet Protocol security (IPsec) encryption for inter-node cluster communication such as the cluster heartbeat.

What value does this change add?

The processing of high latency Group Policy updates can cause Active Directory Domain Services (AD DS) to become temporarily unavailable. In this situation, because IPsec encryption relies on access to AD DS, IPsec encryption is interrupted until the updates are complete. If cluster communication uses IPsec encryption, this interruption prevents inter-node cluster communication (including heartbeat messages) from being sent. If the delay exceeds the cluster heartbeat threshold for a cluster node, the cluster removes the node from cluster membership. If this occurs on multiple nodes at the same time, it could cause the cluster to lose quorum.

By turning off IPsec encryption for inter-node cluster communication, this traffic remains uninterrupted. Therefore, the ability of the cluster to provide high availability for its clustered roles is not impacted by high latency Group Policy updates.

What works differently?

You can now use the **NetFTIPSecEnabled** cluster common property to turn off IPsec encryption on port 3343 for inter-node cluster communication. By default, the **NetFTIPSecEnabled** setting is enabled (set to "1"). A value of "1" means that IPsec encryption for inter-node communication is enabled if there is an existing Group Policy setting that enforces IPsec.

To change the value to "0", which overrides any Group Policy setting and turns off IPsec encryption for inter-node cluster communication, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster). NetFTIPSecEnabled = 0
```

Warning

We recommend that you turn off IPsec encryption for inter-node cluster communication only if you experience issues because of high latency Group Policy updates. If you do turn off the setting, make sure that you thoroughly test the change because it may affect cluster performance.

Cluster dashboard

Failover Cluster Manager now includes a cluster dashboard that enables you to quickly view the health status of all managed failover clusters. You can view the name of the failover cluster together with an icon that indicates whether the cluster is running, the number and status of clustered roles, the node status, and the event status.

What value does this change add?

If you manage multiple failover clusters, this dashboard provides a convenient way for you to quickly check the health of the failover clusters.

What works differently?

In Windows Server 2012, you had to click each failover cluster name to view status information. Now, when you click **Failover Cluster Manager** in the navigation tree, there is a **Clusters** dashboard in the middle pane that shows all managed clusters.

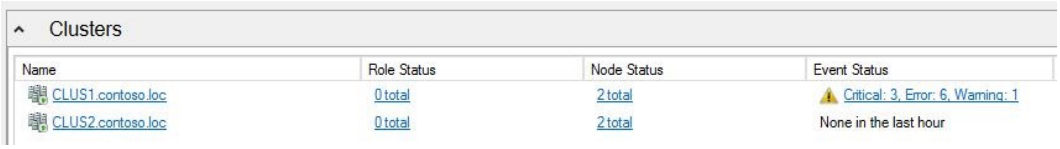

Figure 3. Cluster dashboard

See also

The following table provides additional resources that you can use to evaluate Failover Clustering.

Content type	References
Product evaluation	Failover Clustering Overview
Community resources	High Availability (Clustering) Forum | Failover Clustering and Network Load Balancing Team Blog
Related technologies	Hyper-V Overview | File and Storage Services Overview | Storage Spaces Overview | Scale-Out File Server for Application Data Overview | Network Load Balancing Overview | Microsoft SQL Server | Microsoft Server and Cloud Platform | System Center Virtual Machine Manager

What's New in Failover Clustering in Windows Server 2012

This topic describes the Failover Clustering functionality that is new or changed in Windows Server 2012. This functionality supports increased scalability, continuously available file-based server application storage, easier management, faster failover, and more flexible architectures for failover clusters.

Feature description

Failover clusters provide high availability and scalability to many server workloads. These include server applications such as Microsoft Exchange Server, Hyper-V, Microsoft SQL Server, and file servers. The server applications can run on physical servers or virtual machines. In a failover cluster, if one or more of the clustered servers (nodes) fails, other nodes begin to provide service (a process known as failover). In addition, the clustered roles (formerly called clustered services and applications) are proactively monitored to verify that they are working properly. If they are not working, they restart or move to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. For an overview of the Failover Clustering feature in Windows Server 2012, see Failover Clustering Overview.

New and changed functionality

The following table and sections summarize key Failover Clustering functionality in Windows Server 2012 that is new or changed since Windows Server 2008 R2.

Feature/functionality	New or improved	Description
Cluster scalability	Improved	Scales to 64 nodes and 8,000 virtual machines per cluster
Management of large-scale clusters by using Server Manager and Failover Cluster Manager	New	Provides GUI tools to streamline management and operation of large-scale clusters
Management and mobility of clustered virtual machines and other clustered roles	New	Helps allocate cluster resources to clustered virtual machines and other clustered roles
Cluster Shared Volumes	Improved	Improves CSV setup and enhances security, performance, and file system availability for additional cluster workloads
Support for Scale-Out File Servers	New	Provides CSV storage and integrates with File Services features to support scalable, continuously available application storage
Cluster-Aware Updating	New	Applies software updates

Feature/functionality	New or improved	Description
		across the cluster nodes while maintaining availability
Virtual machine application monitoring and management	New	Extends clustered virtual machine monitoring to the applications that run in the clustered virtual machines
Cluster validation tests	Improved	Validates Hyper-V and CSV functionality and performs faster
Active Directory Domain Services integration	Improved	Increases cluster resiliency and supports a wider range of deployments
Quorum configuration and dynamic quorum	Improved	Simplifies quorum setup and increases the availability of the cluster in failure scenarios
Cluster upgrade and migration	Improved	Allows migration of virtual machines from Windows Server 2008 R2, migration to CSVs, and reuse of existing storage
Task Scheduler integration	New	Integrates Failover Clustering with additional server functionality
Windows PowerShell support	Improved	Allows scripting of Failover Clustering functionality that was introduced in Windows Server 2012

Cluster scalability

Failover clusters in Windows Server 2012 can scale to a greater number of nodes and virtual machines than clusters in Windows Server 2008 R2, as shown in the following table:

Cluster maximum	Windows Server 2012	Windows Server 2008 R2
Nodes	64	16
Virtual machines or clustered	8,000 (up to 1,024 per node)	1,000

Cluster maximum	Windows Server 2012	Windows Server 2008 R2
roles		

Management of large-scale clusters by using Server Manager and Failover Cluster Manager

Server Manager and Failover Cluster Manager provide new capabilities in Windows Server 2012 to manage large-scale clusters.

Server Manager can discover and manage the nodes of the cluster. It enables remote multi-server management, remote role and feature installation, and the ability to start Failover Cluster Manager from the Server Manager GUI. For more information, see Manage Multiple, Remote Servers with Server Manager.

New Failover Cluster Manager features that simplify large-scale management of clustered virtual machines and other clustered roles include:

- **Search, filtering, and custom views**. Administrators can manage and navigate large numbers of clustered virtual machines or other clustered roles.
- **Multiselect**. Administrators can select a specific collection of virtual machines and then perform any needed operation (such as live migration, save, shutdown, or start).
- **Simplified live migration and quick migration of virtual machines and virtual machine storage**. Live migration and quick migration are easier to perform.
- **Simpler configuration of Cluster Shared Volumes (CSVs)**. Configuration is a right-click from the Storage pane. CSVs have additional enhancements, which are described in Cluster Shared Volumes, later in this topic.
- **Support for Hyper-V Replica**. Hyper-V Replica provides point-in-time replication of virtual machines between storage systems, clusters, and data centers for disaster recovery.

What value do these changes add?

These scalability features in Windows Server 2012 improve the configuration, management, and maintenance of large physical clusters and Hyper-V failover clusters.

Management and mobility of clustered virtual machines and other clustered roles

In Windows Server 2012, administrators can configure settings, such as prioritize starting or placing virtual machines and clustered roles on cluster nodes, to efficiently allocate resources to clustered workloads. The following table describes these settings:

Setting	Description	Scope
Priority settings: High, Medium (the default), Low, or No Auto Start	• Clustered roles with higher priority are started and are placed on nodes before	All clustered roles, including clustered virtual machines

Setting	Description	Scope
	those with lower priority. • If a No Auto Start priority is assigned, the role does not come online automatically after it fails, which keeps resources available so other roles can start.	
Preemption of virtual machines based on priority	• The Cluster service takes offline lower priority virtual machines when high-priority virtual machines do not have the necessary memory and other resources to start after a node failure. The freed-up resources can be assigned to high-priority virtual machines. • When necessary, preemption starts with the lowest priority virtual machines and continues to higher priority virtual machines. • Virtual machines that are preempted are later restarted in priority order.	Clustered virtual machines
Memory-aware virtual machine placement	• Virtual machines are placed based on the Non-Uniform Memory Access (NUMA) configuration, the workloads that are already running, and the available resources on each node. The number of failover attempts before a virtual machine is successfully started is reduced. This increases the uptime for virtual machines.	Clustered virtual machines
Virtual machine mobility features	• Multiple live migrations can be started simultaneously. The cluster carries out as	Clustered virtual machines

Setting	Description	Scope
	many as possible, and then queues the remaining migrations to complete later. Failed migrations automatically retry. • Virtual machines are migrated to nodes with sufficient memory and other resources. • A running virtual machine can be added to or removed from a failover cluster. • Virtual machine storage can be live migrated.	
Automated node draining	• The cluster automatically drains a node (moves the clustered roles that are running on the node to another node) before putting the node into maintenance mode or making other changes on the node. • Roles fail back to the original node after maintenance operations. • Administrators can drain a node with a single action in Failover Cluster Manager or by using the Windows PowerShell cmdlet, **Suspend-ClusterNode**. The target node for the moved clustered roles can be specified. • Cluster-Aware Updating uses node draining in the automated process to apply software updates to cluster nodes. For more information, see Cluster-Aware Updating later in this topic.	All clustered roles, including clustered virtual machines

What value do these changes add?

These features in Windows Server 2012 improve the allocation of cluster resources (particularly when starting or maintaining nodes) in large physical clusters and Hyper-V failover clusters.

Cluster Shared Volumes

Cluster Shared Volumes (CSVs) were introduced in Windows Server 2008 R2 to provide common storage for clustered virtual machines. In Windows Server 2012, CSVs can provide storage for additional clustered roles. CSVs allow multiple nodes in the cluster to simultaneously access the same NTFS file system without imposing hardware, file type, or directory structure restrictions. With CSVs, multiple clustered virtual machines can use the same LUN and still live migrate or quick migrate from node to node independently.

The following is a summary of new CSV functionality in Windows Server 2012.

- **Storage capabilities for a wider range of clustered roles.** Includes Scale-Out File Servers for application data, which provide continuously available and scalable file-based (SMB 3.0) server storage for Hyper-V and applications such as Microsoft SQL Server. For more information, see Support for Scale-Out File Servers later in this topic.
- **CSV proxy file system (CSVFS)**. Provides cluster shared storage with a single, consistent file namespace while still using the underlying NTFS file system.
- **Support for BitLocker Drive Encryption**. Allows decryption by using the common identity of the computer account for the cluster (also called the Cluster Name Object, or CNO). This enables physical security for deployments outside secure data centers and meets compliance requirements for volume-level encryption.
- **Ease of file backup**. Supports backup requestors that are running Windows Server 2008 R2 or Windows Server 2012 Backup. Backups can use application-consistent and crash-consistent Volume Shadow Copy Service (VSS) snapshots.
- **Direct I/O for file data access, including sparse files**. Enhances virtual machine creation and copy performance.
- **Removal of external authentication dependencies**. Improves the performance and resiliency of CSVs.
- **Integration with SMB Multichannel and SMB Direct**. Uses new SMB 3.0 features to allow CSV traffic to stream across multiple networks in the cluster and leverage network adapters that support Remote Direct Memory Access (RDMA). For more information, see Server Message Block overview.
- **Storage can be made visible to only a subset of nodes**. Enables cluster deployments that contain application and data nodes.
- **Integration with Storage Spaces**. Allows virtualization of cluster storage on groups of inexpensive disks. The Storage Spaces feature in Windows Server 2012 can integrate with CSVs to permit scale-out access to data. For more information, see Storage Spaces Overview.
- **Ability to scan and repair volumes with zero offline time**. Maintains CSV availability while the NTFS file system identifies, logs, and repairs anomalies.

What value do these changes add?

These new features provide easier CSV setup, broader workload support, enhanced security and performance in a wider variety of deployments, and greater file system availability.

What works differently?

CSVs now appear as CSV File System (CSVFS) instead of NTFS.

Support for Scale-Out File Servers

Scale-Out File Servers can host continuously available and scalable storage by using the SMB 3.0 protocol. Failover clusters in Windows Server 2012 provide the following foundational features that support this type of file server:

- A Distributed Network Name (DNN), which provides an access point for client connections to the Scale-Out File Servers.
- A Scale-out File Server resource type that supports Scale-out File Services.
- Cluster Shared Volumes (CSVs) for storage. For more information, see Cluster Shared Volumes earlier in this topic.
- Integration with File Services features to configure the clustered role for the Scale-Out File Server.

What value do these changes add?

These features support continuously available and readily scalable file services for applications and for end users. For more information, see Scale-Out File Server for Application Data Overview.

Cluster-Aware Updating

Cluster-Aware Updating (CAU) is an automated feature that allows updates to be applied automatically to the host operating system or other system components in clustered servers, while maintaining availability during the update process. This feature leverages automated draining and failback of each node during the update process. By default, it uses the Windows Update Agent infrastructure as its update source. For an overview of the CAU feature, see Cluster-Aware Updating Overview.

What value does this change add?

CAU provides increased uptime of high availability services, easier maintenance of failover clusters, and reliable and consistent IT processes.

Virtual machine application monitoring and management

In clusters running Windows Server 2012, administrators can monitor services on clustered virtual machines that are also running Windows Server 2012. This functionality extends the high-level monitoring of virtual machines that is implemented in Windows Server 2008 R2 failover clusters. If a monitored service in a virtual machine fails, the service can be restarted, or the clustered virtual machine can be restarted or moved to another node (depending on service restart settings and cluster failover settings).

What value does this change add?

This feature increases the uptime of high availability services that are running on virtual machines within a failover cluster.

Cluster validation tests

The Validate a Configuration Wizard in Failover Cluster Manager simplifies the process of validating hardware and software across servers for use in a failover cluster. The performance for large failover clusters has been improved and new tests have been added.

The following are improved features related to validation:

- **Improved performance**. Runs significantly faster, especially to test storage.
- **Targeted validation of new LUNs**. Allows specifying a new LUN (disk), rather than testing all LUNs when validating storage.
- **Integration with WMI**. Exposes cluster validation status to applications and scripts through Windows Management Instrumentation (WMI).
- **New validation tests**. Provides validation test support for CSVs, and for Hyper-V and virtual machines (when the Hyper-V role is installed).
- **Validation test awareness of replicated hardware.** Helps support multisite environments.

What value do these changes add?

The added validation tests help confirm that the servers in the cluster will support smooth failover, particularly of virtual machines from one host to another.

Active Directory Domain Services integration

Integration of failover clusters with Active Directory Domain Services (AD DS) is made more robust in Windows Server 2012 by the following features:

- **Ability to create cluster computer objects in targeted organizational units (OUs) or in the same OUs as the cluster nodes**. Aligns failover cluster dependencies on AD DS with the delegated domain administration model that is used in many IT organizations.
- **Automated repair of cluster virtual computer objects (VCOs) if they are deleted accidentally**.
- **Cluster access only to Read-only domain controllers**. Supports cluster deployments in branch office or perimeter network scenarios.
- **Ability of the cluster to start with no AD DS dependencies**. Enables certain virtualized data center scenarios.

Note

Failover clusters do not support group Managed Service Accounts.

What value do these changes add?

These features improve the configuration and resiliency of failover clusters.

Quorum configuration and dynamic quorum

The following features in Windows Server 2012 enhance the management and functionality of the cluster quorum:

- **Configure Cluster Quorum Wizard**. Simplifies quorum configuration and integrates well with new features and existing quorum functionality.
- **Vote assignment**. Allows specifying which nodes have votes in determining quorum (by default, all nodes have a vote).
- **Dynamic quorum**. Gives the administrator the ability to automatically manage the quorum vote assignment for a node, based on the state of the node. When a node shuts down or crashes, the node loses its quorum vote. When a node successfully rejoins the cluster, it regains its quorum vote. By dynamically adjusting the assignment of quorum votes, the cluster can increase or decrease the number of quorum votes that are required to keep running. This enables the cluster to maintain availability during sequential node failures or shutdowns.

What value do these changes add?

These enhancements simplify quorum setup and increase the availability of the cluster in failure scenarios.

Cluster upgrade and migration

By using the updated Migrate a Cluster Wizard in Windows Server 2012, administrators can migrate the configuration settings for clustered roles (formerly called clustered services and applications) from clusters that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. Migration enhancements in Windows Server 2012 include:

- **Export and reimport Hyper-V virtual machines**.
- **Migrate to CSV disks.**
- **Map storage and virtual networks.**
- **Reuse existing storage.**

What value does this change add?

The Migrate a Cluster Wizard provides ease and flexibility to deploy, upgrade, and migrate failover clusters.

Task Scheduler integration

In Windows Server 2012, Task Scheduler is integrated with Failover Clustering to allow the administrator to configure clustered tasks. A clustered task is a Task Scheduler task that is registered on all cluster nodes. Depending on the task, it can be enabled on all or a subset of the nodes.

The administrator can configure clustered tasks in three ways:

- **Cluster-wide.** The task is scheduled on all cluster nodes.
- **Any node**. The task is scheduled on a single, random node.

- **Resource specific**. The task is scheduled only on a node that owns a specified cluster resource.

The administrator can configure and manage clustered tasks by using the following Windows PowerShell cmdlets:

- **Register-ClusteredScheduledTask**
- **Set-ClusteredScheduledTask**
- **Get-ClusteredScheduledTask**
- **Unregister-ClusteredScheduledTask**

What value does this change add?

Clustered tasks provide a flexible mechanism to use cluster resources to run applications or processes at predefined times.

Windows PowerShell support

To use the Windows PowerShell cmdlets for Failover Clustering, you must install the Failover Cluster module for Windows PowerShell that is included with the Failover Clustering tools. For a complete list of the cmdlets, see Failover Clustering Cmdlets in Windows PowerShell.

New Windows PowerShell cmdlets support capabilities in Failover Clustering in Windows Server 2012 including the following:

- Manage cluster registry checkpoints, including cryptographic checkpoints.
- Create Scale-Out File Servers.
- Monitor virtual machine applications.
- Update the properties of a Distributed Network Name resource.
- Create and manage clustered tasks.
- Create a highly available iSCSI Target Server.

What value does this change add?

The new Windows PowerShell cmdlets provide management and scripting support for the Failover Clustering features in Windows Server 2012.

What works differently?

The **Test-ClusterResourceFailure** cmdlet replaces **Fail-ClusterResource**.

Removed or deprecated functionality

- The Cluster.exe command-line tool is deprecated, but it can be optionally installed with the Failover Clustering tools. Windows PowerShell cmdlets for Failover Clustering provide functionality that is generally equivalent to Cluster.exe commands. For more information, see Mapping Cluster.exe Commands to Windows PowerShell Cmdlets for Failover Clusters.
- The Cluster Automation Server (MSClus) COM interface is deprecated, but it can be optionally installed with the Failover Clustering tools.
- Support for 32-bit cluster resource DLLs is deprecated, but 32-bit DLLs can be optionally installed. You should update cluster resource DLLs to 64-bit.

- The Print Server role is removed from the High Availability Wizard, and it cannot be configured in Failover Cluster Manager.
- The **Add-ClusterPrintServerRole** cmdlet is deprecated, and it is not supported in Windows Server 2012.

See also

The following table provides additional resources on the web for evaluating Failover Clustering.

Content type	References
Product evaluation	Failover Clustering Overview | Microsoft Server and Cloud Platform | Private Cloud Solution Hub
Community resources	High Availability (Clustering) Forum | Failover Clustering and Network Load Balancing Team Blog
Related technologies	Hyper-V Overview | File and Storage Services Overview | Storage Spaces Overview | Scale-Out File Server for Application Data Overview | Network Load Balancing Overview | Microsoft SQL Server | Microsoft Server and Cloud Platform | System Center Virtual Machine Manager

Failover Clustering Hardware Requirements and Storage Options

You need the following hardware to create a failover cluster. To be supported by Microsoft, all hardware must be certified for the version of Windows Server that you are running, and the complete failover cluster solution must pass all tests in the Validate a Configuration Wizard. For more information about validating a failover cluster, see Validate Hardware for a Failover Cluster.

- **Servers**: We recommend that you use a set of matching computers that contain the same or similar components.
- **Network adapters and cable (for network communication)**: If you use iSCSI, each network adapter should be dedicated to either network communication or iSCSI, not both.

  In the network infrastructure that connects your cluster nodes, avoid having single points of failure. For example, you can connect your cluster nodes by multiple, distinct networks. Alternatively, you can connect your cluster nodes with one network that is constructed with

teamed network adapters, redundant switches, redundant routers, or similar hardware that removes single points of failure.

> **Note**
> If you connect cluster nodes with a single network, the network will pass the redundancy requirement in the Validate a Configuration Wizard. However, the report from the wizard will include a warning that the network should not have single points of failure.

- **Device controllers or appropriate adapters for the storage**:
  - **Serial Attached SCSI or Fibre Channel**: If you are using Serial Attached SCSI or Fibre Channel, in all clustered servers, all elements of the storage stack should be identical. It is required that the multipath I/O (MPIO) software be identical and that the Device Specific Module (DSM) software be identical. It is recommended that the mass-storage device controllers—that is, the host bus adapter (HBA), HBA drivers, and HBA firmware—that are attached to cluster storage be identical. If you use dissimilar HBAs, you should verify with the storage vendor that you are following their supported or recommended configurations.
  - **iSCSI**: If you are using iSCSI, each clustered server should have one or more network adapters or HBAs that are dedicated to the cluster storage. The network you use for iSCSI should not be used for network communication. In all clustered servers, the network adapters you use to connect to the iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or higher.
- **Storage**: You must use shared storage that is compatible with Windows Server 2012 R2 or Windows Server 2012. You can use shared storage that is attached, and you can also use SMB 3.0 file shares as shared storage for servers that are running Hyper-V that are configured in a failover cluster. For more information, see Deploy Hyper-V over SMB.

  In most cases, attached storage should contain multiple, separate disks (logical unit numbers, or LUNs) that are configured at the hardware level. For some clusters, one disk functions as the disk witness (described at the end of this subsection). Other disks contain the files required for the clustered roles (formerly called clustered services or applications). Storage requirements include the following:
  - To use the native disk support included in Failover Clustering, use basic disks, not dynamic disks.
  - We recommend that you format the partitions with NTFS. If you use Cluster Shared Volumes (CSV), the partition for each of those must be NTFS.

    > **Note**
    > If you have a disk witness for your quorum configuration, you can format the disk with either NTFS or Resilient File System (ReFS).
  - For the partition style of the disk, you can use either master boot record (MBR) or GUID partition table (GPT).

  A disk witness is a disk in the cluster storage that is designated to hold a copy of the cluster configuration database. A failover cluster has a disk witness only if this is specified as part of

the quorum configuration. For more information, see Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster.

Hardware requirements for Hyper-V

If you are creating a failover cluster that includes clustered virtual machines, the cluster servers must support the hardware requirements for the Hyper-V role. Hyper-V requires a 64-bit processor that includes the following:

- Hardware-assisted virtualization. This is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.
- Hardware-enforced Data Execution Prevention (DEP) must be available and enabled. Specifically, you must enable Intel XD bit (execute disable bit) or AMD NX bit (no execute bit).

For more information about the Hyper-V role, see Hyper-V Overview.

Deploying storage area networks with failover clusters

When deploying a storage area network (SAN) with a failover cluster, follow these guidelines:

- **Confirm compatibility of the storage**: Confirm with manufacturers and vendors that the storage, including drivers, firmware, and software used for the storage, are compatible with failover clusters in the version of Windows Server that you are running.
- **Isolate storage devices, one cluster per device**: Servers from different clusters must not be able to access the same storage devices. In most cases, a LUN used for one set of cluster servers should be isolated from all other servers through LUN masking or zoning.
- **Consider using multipath I/O software or teamed network adapters**: In a highly available storage fabric, you can deploy failover clusters with multiple host bus adapters by using multipath I/O software or network adapter teaming (also called load balancing and failover, or LBFO). This provides the highest level of redundancy and availability. For Windows Server 2012 R2 or Windows Server 2012, your multipath solution must be based on Microsoft Multipath I/O (MPIO). Your hardware vendor will typically supply an MPIO device-specific module (DSM) for your hardware, although Windows Server includes one or more DSMs as part of the operating system.

  For information about LBFO, see NIC Teaming Overview in the Windows Server Technical Library.

  Important
  Host bus adapters and multipath I/O software can be very version sensitive. If you are implementing a multipath solution for your cluster, work closely with your hardware vendor to choose the correct adapters, firmware, and software for the version of Windows Server that you are running.

- **Consider using Storage Spaces**: If you plan to deploy serial attached SCSI (SAS) clustered storage that is configured using Storage Spaces, see Deploy Clustered Storage Spaces for the requirements.

See also

- Failover Clustering Overview
- Storage Spaces Overview
- Using Guest Clustering for High Availability

Validate Hardware for a Failover Cluster

You can use the Validate a Configuration Wizard, which is integrated in Failover Cluster Manager, or the Test-Cluster Windows PowerShell cmdlet, to run a set of focused validation tests. You can run this process on a collection of servers that you intend to use as nodes in a cluster. This tests the underlying hardware and software, directly and individually, to obtain an accurate assessment of how well Failover Clustering can be supported in a given configuration.

Important
A cluster validation report is required by Microsoft as a condition of Microsoft support for a given configuration.

This topic provides steps to validate the hardware for a failover cluster.

Task	Description
Step 1: Prepare to validate hardware for a failover cluster	Learn about cluster validation and Microsoft support for a cluster configuration, and prepare your hardware for the validation tests.
Step 2: Validate a new or existing failover cluster	Run the Validate a Configuration Wizard or the Test-Cluster Windows PowerShell cmdlet.
Step 3: Analyze validation results	Review the Summary Report that is created when validation completes. If there are failures and you need support, prepare a validation report for Microsoft Customer Service and Support.
Advanced validation scenarios	Review these advanced scenarios if you need to validate an existing cluster and choose to test only certain aspects of cluster functionality.
Frequently asked questions	Get answers to your questions about the

Task	Description
	cluster validation process.

Step 1: Prepare to validate hardware for a failover cluster

What is cluster validation?

The Validate a Configuration Wizard or the Test-Cluster Windows PowerShell cmdlet enables you to run a set of focused tests on a collection of servers, networks, and associated storage that are planned for use as a failover cluster. The cluster validation process tests the underlying hardware and software to obtain an accurate assessment of how well Failover Clustering can be supported in a given configuration.

Before you create a failover cluster, we strongly recommend that you run all the cluster validation tests.

Cluster validation is intended to do the following:

- Find hardware or configuration issues before a failover cluster goes into production.
- Help ensure that the clustering solution you deploy is dependable.
- Provide a way to validate changes to the hardware of an existing cluster.
- Perform diagnostic tests on an existing cluster.

Supported cluster configurations

For Microsoft Customer Service and Support to officially support a failover cluster in Windows Server 2012 R2 or Windows Server 2012, the cluster solution must meet the following criteria:

- All hardware and software components must meet the qualifications for the appropriate logo. For Windows Server 2012 R2, this is the "Certified for Windows Server 2012 R2" logo. For Windows Server 2012, this is the "Certified for Windows Server 2012" logo. For more information, see Logo Program Requirements and Policies on the Microsoft website. The following are descriptions of the logos.
  - **Certified Windows Server 2012 R2 Systems**

    The Certified for Windows Server 2012 R2 logo demonstrates that a server system meets Microsoft's highest technical bar for security, reliability and manageability; and with other certified devices and drivers, it can support the roles, features and interfaces for Cloud and Enterprise workloads, as well as business critical applications.
  - **Certified Windows Server 2012 R2 Devices**

    The Certified for Windows Server 2012 R2 logo demonstrates that a server system meets Microsoft's highest technical bar for security, reliability and manageability; and with other certified devices and drivers, it can support the roles, features and interfaces for Cloud and Enterprise workloads, as well as business critical applications.

    - **Certified for Windows Server 2012 Devices**

      Designed for line-of-business and mission-critical applications, the Certified for Windows Server 2012 logo demonstrates that your solution meets Microsoft's highest technical bar for Windows fundamentals and platform compatibility.
    - **Certified Windows Server 2012 Systems**

      Designed for cloud and infrastructure workloads, as well as business critical applications, the Microsoft Certified for Windows Server 2012 logo demonstrates that a server system meets Microsoft's highest technical bar for security, reliability, and manageability, and with any required hardware components, can support all the roles, features, and interfaces that are supported by Windows Server 2012.
- The fully configured system (servers, network, and storage) must pass all the required tests in the Validate a Configuration Wizard, which you can run from Failover Cluster Manager. Alternatively, you can run the validation tests using the Test-Cluster Windows PowerShell cmdlet.

Microsoft support policies are also described at the Microsoft Support website.

Common validation scenarios

The following lists describe scenarios in which hardware validation is required or is useful. In general, you need to run all the validation tests (some exceptions are noted).

- **Validation before the cluster is configured**
    - **A set of servers ready to become a failover cluster**

      This is the most straightforward validation scenario. The hardware components (systems, networks, and storage) are connected, but the systems are not functioning as a cluster. Running tests in this situation has no impact on availability.
    - **Cloned or imaged systems**

      With systems that you have cloned or imaged to different hardware, you must run the Validate a Configuration Wizard as you would with any new cluster. We recommend that you run the wizard just after you connect the hardware components and install the Failover Clustering feature, before the cluster is used by clients.
    - **Virtualized servers**

      With virtualized servers in a cluster, run the Validate a Configuration Wizard as you would with any new cluster. The requirement for running the wizard is the same whether you have a "host cluster" (failover occurs between two physical computers), a "guest cluster" (failover occurs between guest operating systems on the same physical computer), or some other configuration that includes one or more virtualized servers.
- **Validation when the cluster has only one node**

  You might want to run a limited number of validation tests on a single server that you intend to use in a cluster. Some tests cannot be run in this situation; for example, tests that confirm the software and software updates match between servers, and storage tests that simulate failover between nodes. You must have at least two nodes in a cluster before you can

complete the cluster validation process. So if you bring more servers into the configuration, you must run the Cluster Validation Wizard again so that all the tests can be complete.

- **Validation after the cluster is configured and in use**
  - **To confirm validation for Microsoft support, or to rule out configuration problems**

    If you need support from Microsoft, you may be required to you provide the validation report from the wizard. If you have not already run the wizard and saved the report, you might need to take the cluster offline to run the wizard. The report shows whether your configuration is supported, and it can help Microsoft customer support troubleshoot configuration issues with hardware, drivers, and the basic system configuration.
  - **Before adding a node**

    When you add a server to a cluster, we strongly recommend that you start by connecting the server to the cluster networks and storage and then run the Validate a Configuration Wizard, specifying the existing cluster nodes and the new node.
  - **When attaching new storage**

    When you attach new storage to the cluster (which is different than exposing a new logical unit number (LUN) in existing storage), you must run the Validate a Configuration Wizard to confirm that the new storage will function correctly. To minimize the impacts to availability, we recommend that you run the wizard after you attach the storage, and before you begin to use the new LUNs in clustered services or applications.
  - **When making changes that affect firmware or drivers**

    If you want to upgrade or make other changes to the cluster that would require changing the firmware or drivers, you must run the Validate a Configuration Wizard to confirm that the new combination of hardware, firmware, drivers, and software supports your failover cluster functionality. If the change affects the firmware or drivers for storage, we recommend that you keep a small LUN available (unused by clustered roles) so that you can run the storage validation tests without taking your clustered roles offline. For more information about running storage tests on a small unused LUN, see **Considerations for validating an existing cluster**.
  - **After restoring a system from backup**

    After you restore a system from backup, run the Validate a Configuration Wizard to confirm that the system can function properly as part of a cluster. The system is not considered a supported system until the validation tests are completed.

When you validate hardware changes to an existing cluster (as an advanced scenario), you might decide to omit certain storage tests. For more information and considerations, see Advanced validation scenarios .

Categories of validation tests

The following table lists the categories of validation tests. The tests in each category are listed at the time you run the Validate a Configuration Wizard. A description of each test in each category is provided in the validation report that is saved after validation completes.

Category	Description
Cluster configuration	Lists and validates the resources that are configured for use in a cluster, including clustered roles and cluster volumes.
Hyper-V configuration	Validates the Hyper-V configuration for use in a failover cluster. **Note** Hyper-V configuration tests are only required if you use or plan to use clustered virtual machines.
Inventory	Lists host bus adapters (HBAs), devices, processes, and drivers that are in use on the computers in the cluster.
Network	Validates the configuration of the cluster networks, IP addresses, and Windows Firewall.
Storage	Validates storage disks and file systems that are available for use in a failover cluster.
System configuration	Validates operating systems, update levels, and service settings on the computers in the cluster.

Step 2: Validate a new or existing failover cluster

This step provides procedures for running the Validate a Configuration Wizard or the Test-Cluster Windows PowerShell cmdlet to validate a new or existing failover cluster.

Important
To begin the process of adding hardware (such as an additional server) to a failover cluster, connect the hardware to the failover cluster. Then run the Validate a Configuration Wizard, and specify all servers that you want to include in the cluster. The wizard tests cluster connectivity and failover, not only isolated components (such as individual servers).

To run the Validate a Configuration Wizard

1. Identify the server or servers that you want to test.
   - If the cluster does not yet exist, choose the servers that you want to include in the cluster, and make sure that you have installed the Failover Clustering feature on those servers. If the feature is not installed, see the installation instructions.
   - If the cluster already exists, make sure that you know the name of the cluster or a

node in the cluster.

2. Review the cluster requirements for the hardware for the network or the storage that you want to validate, and confirm that it is connected to the servers.
3. Decide whether you want to run all or only some of the available validation tests. In general, we recommend that you run all the tests, but the follow general guidelines can help you decide.

Type of cluster	Validation tests
New or planned cluster with all hardware connected	All tests
New or planned cluster with parts of the hardware connected	System Configuration tests, Inventory tests, and tests that apply to the hardware that is connected (that is, network tests if the network is connected or storage tests if the storage is connected)
Existing cluster to which you plan to add a server	All tests **Note** Before you run the tests, be sure to connect the networks and storage for all servers that you plan to have in the cluster.
Troubleshooting an existing cluster	All tests, although you could run only the tests that relate to the apparent issue.

Important
If a clustered role is using a disk when you start the wizard, the wizard will prompt you about whether to take that clustered role offline for the purposes of testing. If you choose to take a clustered role offline, it will remain offline until the tests finish.

4. To open the wizard, in Failover Cluster Manager, under **Actions**, click **Validate Configuration**.
5. Follow the instructions in the wizard to specify the servers (in a planned cluster) and the tests. For example, if you do not plan to use cluster features that require Hyper-V, you can omit the Hyper-V Configuration tests. The wizard then guides you to run the tests.

Note
When you run the wizard on unclustered servers, you must enter the names of all the servers that you want to test, not only one.

6. The **Summary** page appears after the tests run. On the **Summary** page, click **View Report** to view the test results.

To view the results of the tests after you close the wizard, under **Actions** in Failover Cluster Manager, click **View Validation Report**. You can see %SystemRoot%\Cluster\Reports\Validation Report *date and time*.**html** where %SystemRoot% is the folder in which the operating system is installed (for example, C:\Windows).

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example runs all cluster validation tests on the nodes named *node1* and *node2*. If *node1* or *node2* is already a member of a cluster, the tests will include all nodes in that cluster.

```
Test-Cluster -Node node1,node2
```

Step 3: Analyze validation results

After the Validate a Configuration Wizard has completed, the Failover Cluster Validation Report displays the results. All tests must pass with a green check mark, or in some cases, a yellow triangle (warning). The following table shows the symbols in the summary and explains what they mean:

Symbol	Explanation
	The corresponding validation test passed, indicating that this aspect of the cluster can be supported.
	The corresponding validation test produced a warning, indicating that this aspect of the cluster can be supported, but it might not meet the recommended best practices, and it should be reviewed. Microsoft customer support might ask you to investigate or address the issue if it appears to be directly linked to something that you are troubleshooting.
	The corresponding validation test failed, and this aspect of the cluster is not supported. You must correct the issue before you can create a failover cluster that is supported.
	The corresponding validation test was canceled. This can occur when the test depended on another test that did not complete successfully.

When you look for problem areas (warning or failures), in the test results summary, click an individual test to review the details. Also review the summary statement for information about whether the cluster is a supported configuration.

After you take action to correct the problem, you can rerun the wizard as needed to confirm that the configuration passes the tests.

What to do if a validation test fails

In most cases, if any tests in the Validate a Configuration Wizard fail, Microsoft does not consider the configuration to be supported.

If any of the Hyper-V Configuration tests fail, Hyper-V on the cluster is not configured correctly. The problem must be corrected before the virtual machines in the cluster can be supported. However, failures in this category of tests do not mean that the cluster is not supported for workloads other than clustered virtual machines.

The type of test that fails is a guideline to the corrective action to take. For example, if the **List All Disks** storage test fails, and subsequent storage tests do not run (because they would also fail), you should contact the storage vendor to troubleshoot this issue. Similarly, if a network test that is related to IP addresses fails, consult your network infrastructure team. Not all warnings or errors indicate a need to call Microsoft customer support. Most of the warnings or errors should result in working with internal teams or with a specific hardware vendor.

For information about correcting failures that are listed in a validation report, see the previous section, Step 3: Analyze validation results .

After the issues have been addressed and resolved, it is necessary to rerun the Validate a Configuration Wizard. To be considered a supported configuration, all tests are required to run and complete without failures.

Provide a validation report when you request support from Microsoft

If you need to contact Microsoft customer support about a validation problem, the support team will help you collect the validation report and other relevant configuration files by using the Microsoft Support Diagnostic Tool (MSDT). (This feature replaces the MPSReports data collection utility.) If needed, Microsoft will send instructions about how to capture the data. In some situations, Microsoft may request that the contents of the C:\Windows\Cluster\Reports folder be zipped and sent for analysis. Either method will collect the required cluster validation report.

Updates to validation tests

The Validate a Configuration Wizard provides an accurate picture of how well failover clustering can be supported for a given configuration. If an update to the Validate a Configuration Wizard becomes available, you may need to rerun the wizard and pass all tests for your configuration to continue to be supported. This may result in some solutions that previously passed to fail. The

issues that are reported in the updated tests will need to be addressed in the same manner as outlined in this guide.

Advanced validation scenarios

When you make a change to an existing cluster, you might not need to run all the cluster validation tests. The following tables in this section list the kinds of changes you might make to a cluster and the corresponding tests to run.

Key for the required validation tests shown in the following tables

- **Full**: Run the complete set of tests. This requires some cluster downtime.
- **Single LUN**: Run the complete set of tests, and run the storage tests on only one LUN. The LUN might be a small LUN that you set aside for testing purposes or the witness disk (if your cluster uses a witness disk). This validates the storage subsystem, but not specifically each individual LUN or disk. You can run these validation tests without causing downtime to your clustered services or applications.
- **Omit storage tests**: Run the system configuration, inventory, and network tests, but not the storage tests. You can run these validation tests without causing downtime to your clustered roles.
- **None**: No validation tests are needed.

Server changes

Change	Validation tests required
Physically replacing or changing a server that is used in the cluster	Full
Adding or removing CPUs	None
Adding or removing RAM on a server	None
Adding, removing, or replacing a network adapter	Omit storage tests
Updating firmware or an existing network driver	Omit storage tests
Changing the BIOS settings or firmware version	None
Adding or changing peripheral devices other than network or storage components, such as CD-ROM or DVD drives, tape drives, video cards, sound devices, and USB devices	None

Operating system changes

Change	Validation tests required
Applying operating system service pack, software updates, or hotfixes that affect the storage stack	Single LUN
Applying software updates or hotfixes that do not affect the storage stack	Omit storage tests
Installing an application that has no kernel mode or filter drivers	None
Changing or adding new kernel mode drivers	Single LUN

Cluster configuration changes

Change	Validation tests required
Adding a new node to the cluster	Full
Adding a new node that uses dissimilar hardware	Full
Removing a node from the cluster	None
Changing the quorum configuration	None

Shared storage changes

Change	Validation tests required
Changing or adding a storage array	Full
Adding another SCSI hardware RAID unit of the same type, and that unit uses an HBA that is already in the configuration	Single LUN
Making a minor (0.x) revision to the storage firmware	Single LUN
Making a major (X.0) revision to the storage firmware	Single LUN
Presenting a new disk or LUN to a cluster	Full, but test new LUNs only

SAN (switch/hub) changes

Change	Validation tests required
Adding or replacing a Fibre Channel switch or hub	Full
Changing the number of ports within a switch block	None
Making a minor (0.x) revision to the Fibre Channel switch firmware	Single LUN
Making a major (X.0) revision to the Fibre Channel switch firmware	Single LUN
Changing a switch configuration or zoning	Full, but test changed LUNs only

Host bus adapter (HBA) changes

Change	Validation tests required
Replacing an HBA (same or different type)	Full
Adding a new HBA (same or different type)	Single LUN
Changing the HBA firmware or BIOS	Single LUN
Changing the HBA driver version	Single LUN

Multipath software changes

Change	Validation tests required
Changing from single path to multipath or multipath to single path	Full
Adding a path	Single LUN
Removing a path	Single LUN
Updating the device specific module (DSM) version	Single LUN
Changing to a DSM of a different type, for example, a DSM from a different provider	Single LUN

Multisite cluster changes

Change	Validation tests required
Modifying the networks that connect the nodes	Omit storage tests
Making a minor (0.x) version change in the data replication software	Single LUN
Making a major (X.0) version change in the data replication software, or changing to a different type of replication software	Full

Networking changes

Change	Validation tests required
Modifying network firmware, software, or hardware	Omit storage tests

Including storage tests

When you perform cluster validation tests on a configured cluster, you might not always run all storage tests. This section explains what to consider if you include storage tests or do not include storage tests.

Considerations when you include storage tests

The Validate a Configuration Wizard runs all storage tests by default. All or some of the storage tests can be unselected by choosing the **Run only tests I select** option on the **Testing Options** page of the wizard. When storage tests are included, the **Review Storage Status** page of the wizard shows all of the disks and storage pools in the cluster and allows you to select the disks and storage pools to include in the storage tests. Storage tests require that a disk or storage pool that is assigned to a clustered role or Cluster Shared Volume be taken offline first. Therefore, anything using the storage will not have access to it during the storage tests. We recommend that any clustered role or other process that may be using the disk or storage pool is taken offline before the storage is included in the storage validation tests.

The Test-Cluster Windows PowerShell cmdlet runs all storage tests by default. You can specify the **–Include** parameter to only run storage tests or to run a specific storage test. You can use the **–Disk** and **-Pool** parameters to enable targeted storage validation. The **–Disk** parameter or the **–Pool** parameter allows specifying, respectively, one or more disks or one or more storage pools to include in the storage validation testing. If the **–Disk** parameter or the **–Pool** parameter is used to specify a disk or storage pool that is currently online and is assigned to a clustered role or Cluster Shared Volume, you must also specify the **–Force** parameter to validate the corresponding disk or storage pool; otherwise, you must ensure that the clustered disk or storage pool is offline before running the tests. If the **–Disk** parameter or the **–Pool** parameter is not

specified, **Test-Cluster** runs storage tests on all disks and storage pools that are available for cluster use or that are in the cluster resource offline or failed state. We recommend that any clustered role or other process that may be using the disk or storage pool is taken offline before the storage is included in the validation testing.

Storage that is not directly connected to all nodes in the cluster

There may be cases where the cluster design includes storage that is not connected to all nodes in the cluster. A common example is in multisite clusters where cluster nodes in *SiteA* are connected to one set of storage, the nodes in *SiteB* are connected to a different set of storage, and a non-Microsoft replication solution is used to ensure that both sets of storage have the same data. Failover cluster detects this asymmetric storage configuration, so the disks in *SiteA* only validate with the *SiteA* nodes, and the disks in *SiteB* only validate with nodes in *SiteB*.

One scenario where Microsoft customer support may request that you run validation tests on production clusters is when there is a cluster storage failure that could be caused by some underlying storage configuration change or problem. It may not be advisable to take a disk that is in use offline due to the availability impact it might have on the clustered roles that use it. In this situation, you can run validation tests (including storage tests) by creating or choosing a new LUN from the same shared storage device and presenting it to all nodes in the cluster. By testing this LUN, you can avoid disruption to clustered roles that are already online within the cluster and still test the underlying storage subsystem.

If a failover cluster passed the full set of validation tests and has no future hardware or software changes, it will continue to be a supported configuration. However, when you perform routine updates to software components such as drivers and firmware, it may be necessary to rerun the Configuration Wizard to ensure that the current configuration of the failover cluster is supported. The following guidelines can help you decide when this is necessary:

- All components of the storage stack should be identical across all nodes in the cluster. It is required that multipath I/O (MPIO) software and Device Specific Module (DSM) software components be identical. We recommend that the mass-storage device controllers (that is, the host bus adapter (HBA), HBA drivers, and HBA firmware) that are attached to cluster storage be identical. If you use dissimilar HBAs, you should verify with the storage vendor that you are following their supported or recommended configurations.
- A best practice is to keep a small LUN available to allow the Validate a Configuration Wizard to run tests on available storage without negatively impacting clustered roles. If Microsoft customer support requests that you run a full set of cluster validation tests, the wizard allows you to select that disk for the storage tests to verify that the storage is working properly.

Considerations when you do not include storage tests

System configuration tests, inventory tests, and network tests have very low overhead, and they can be performed without significant effect on servers in a cluster.

Microsoft customer support might request that you validate a production cluster as part of normal troubleshooting procedures (not focused on storage). In this scenario, you will use the wizard to inventory hardware and software, perform network testing, and validate system configuration. There may be certain scenarios in which only a subset of the full tests are needed. For example,

if you are troubleshooting a network issue on a production cluster, Microsoft customer support might request that you run only the hardware and software inventory and the network tests.

Frequently asked questions

If a cluster passes all tests in the Validate a Configuration Wizard, is it supported?

If all the hardware and software components in the cluster meet the qualifications for the "Certified for Windows Server 2012 R2" or "Certified for Windows Server 2012" logo, and the cluster passes the validation tests, it is considered to be supported by Microsoft Customer Service and Support for Failover Clustering. For more information, see Supported cluster configurations earlier in this guide.

Will failover cluster solutions be listed in the Windows Server Catalog?

No, Microsoft will not maintain a list of vendor solutions for failover clusters. However, many vendors list recommended failover cluster solutions and components on their websites.

How does Microsoft customer support check if the solution has been validated?

The Validate a Configuration Wizard generates a simple HTML report that clearly displays whether a solution has passed all tests. This report is collected as part of the standard diagnostics tool, MSDT.

What if I make a change to the cluster configuration, like add a node? Do I have to run the Validate a Configuration Wizard again?

Yes, the Validate a Configuration Wizard should be run any time a change is made to an existing failover cluster. For more information, see Common validation scenarios earlier in this guide.

See also

- Failover Clustering Overview
- Failover Clustering Hardware Requirements and Storage Options

Create a Failover Cluster

This topic shows how to create a failover cluster by using either the Failover Cluster Manager snap-in or Windows PowerShell. The topic covers a typical deployment, where computer objects for the cluster and its associated clustered roles are created in Active Directory Domain Services (AD DS).

Note

In Windows Server 2012 R2 you can also deploy an Active Directory-detached cluster. This deployment method enables you to create a failover cluster without permissions to create computer objects in AD DS or the need to request that computer objects are prestaged in AD DS. This option is only available through Windows PowerShell, and is only recommended for specific scenarios. For more information, see Deploy an Active Directory-Detached Cluster.

Checklist: Create a failover cluster

	Task	Reference
☐	Verify the prerequisites	Verify the prerequisites
☐	Install the Failover Clustering feature on every server that you want to add as a cluster node	Install the Failover Clustering feature
☐	Run the Cluster Validation Wizard to validate the configuration	Validate the configuration
☐	Run the Create Cluster Wizard to create the failover cluster	Create the failover cluster
☐	Create clustered roles to host cluster workloads	Create clustered roles

Create a failover cluster by using Windows PowerShell

Verify the prerequisites

Before you begin, verify the following prerequisites:

- Make sure that all servers that you want to add as cluster nodes are running the same version of Windows Server.

  Note

You can use the Failover Clustering feature on all editions of Windows Server 2012 R2 and Windows Server 2012. This includes Server Core installations.

- Review the hardware requirements to make sure that your configuration is supported. For more information, see Failover Clustering Hardware Requirements and Storage Options.
- If you want to add clustered storage during cluster creation, make sure that all servers can access the storage. (You can also add clustered storage after you create the cluster.)
- Make sure that all servers that you want to add as cluster nodes are joined to the same Active Directory domain.
- (Optional) Create an organizational unit (OU) and move the computer accounts for the servers that you want to add as cluster nodes into the OU. As a best practice, we recommend that you place failover clusters in their own OU in AD DS. This can help you better control which Group Policy settings or security template settings affect the cluster nodes. By isolating clusters in their own OU, it also helps prevent against accidental deletion of cluster computer objects.

Additionally, verify the following account requirements:

- Make sure that the account you want to use to create the cluster is a domain user who has administrator rights on all servers that you want to add as cluster nodes.
- Make sure that either of the following is true:
  - The user who creates the cluster has the **Create Computer objects** permission to the OU or the container where the servers that will form the cluster reside.
  - If the user does not have the **Create Computer objects** permission, ask a domain administrator to prestage a cluster computer object for the cluster. For more information, see Prestage Cluster Computer Objects in Active Directory Domain Services.

Note
This requirement does not apply if you want to create an Active Directory-detached cluster in Windows Server 2012 R2. For more information, see Deploy an Active Directory-Detached Cluster.

Install the Failover Clustering feature

You must install the Failover Clustering feature on every server that you want to add as a failover cluster node.

To install the Failover Clustering feature

1. Start Server Manager.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, click the server where you want to install the feature, and then click **Next**.

6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, select the **Failover Clustering** check box.
8. To install the failover cluster management tools, click **Add Features**, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**.

   Note
   A server restart is not required for the Failover Clustering feature.
10. When the installation is completed, click **Close**.
11. Repeat this procedure on every server that you want to add as a failover cluster node.

Note
After you install the Failover Clustering feature, we recommend that you apply the latest updates from Windows Update. Also, for a Windows Server 2012-based failover cluster, review the Microsoft Support article Recommended hotfixes and updates for Windows Server 2012-based Failover Clusters, and install any updates that apply.

Validate the configuration

Before you create the failover cluster, we strongly recommend that you validate the configuration to make sure that the hardware and hardware settings are compatible with failover clustering. Microsoft supports a cluster solution only if the complete configuration passes all validation tests and if all hardware is certified for the version of Windows Server that the cluster nodes are running.

Note
You must have at least two nodes to run all tests. If you have only one node, many of the critical storage tests do not run.

To run cluster validation tests

1. On a computer that has the Failover Cluster Management Tools installed from the Remote Server Administration Tools, or on a server where you installed the Failover Clustering feature, start Failover Cluster Manager. To do this on a server, start Server Manager, and then on the **Tools** menu, click **Failover Cluster Manager**.
2. In the **Failover Cluster Manager** pane, under **Management**, click **Validate Configuration**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select Servers or a Cluster** page, in the **Enter name** box, enter the NetBIOS name or the fully qualified domain name of a server that you plan to add as a failover cluster node, and then click **Add**. Repeat this step for each server that you want to add. To add multiple servers at the same time, separate the names by a comma or by a semicolon. For example, enter the names in the format *server1.contoso.com, server2.contoso.com*. When you are finished, click **Next**.

5. On the **Testing Options** page, click **Run all tests (recommended)**, and then click **Next**.
6. On the **Confirmation** page, click **Next**.

   The Validating page displays the status of the running tests.
7. On the **Summary** page, do either of the following:
   - If the results indicate that the tests completed successfully and the configuration is suited for clustering, and you want to create the cluster immediately, make sure that the **Create the cluster now using the validated nodes** check box is selected, and then click **Finish**. Then, continue to step 4 of the Create the failover cluster procedure.
   - If the results indicate that there were warnings or failures, click **View Report** to view the details and determine which issues must be corrected. Realize that a warning for a particular validation test indicates that this aspect of the failover cluster can be supported, but might not meet the recommended best practices.

     Note

     If you receive a warning for the Validate Storage Spaces Persistent Reservation test, see the blog post Windows Failover Cluster validation warning indicates your disks don't support the persistent reservations for Storage Spaces for more information.

   For more information about hardware validation tests, see Validate Hardware for a Failover Cluster.

Create the failover cluster

To complete this step, make sure that the user account that you log on as meets the requirements that are outlined in the Verify the prerequisites section of this topic.

To create the failover cluster

1. Start Server Manager.
2. On the **Tools** menu, click **Failover Cluster Manager**.
3. In the **Failover Cluster Manager** pane, under **Management**, click **Create Cluster**.

   The Create Cluster Wizard opens.
4. On the **Before You Begin** page, click **Next**.
5. If the **Select Servers** page appears, in the **Enter name** box, enter the NetBIOS name or the fully qualified domain name of a server that you plan to add as a failover cluster node, and then click **Add**. Repeat this step for each server that you want to add. To add multiple servers at the same time, separate the names by a comma or a semicolon. For example, enter the names in the format *server1.contoso.com; server2.contoso.com*. When you are finished, click **Next**.

   Note

   If you chose to create the cluster immediately after running validation in the

procedure Validate the configuration, you will not see the **Select Servers** page. The nodes that were validated are automatically added to the Create Cluster Wizard so that you do not have to enter them again.

6. If you skipped validation earlier, the **Validation Warning** page appears. We strongly recommend that you run cluster validation. Only clusters that pass all validation tests are supported by Microsoft. To run the validation tests, click **Yes**, and then click **Next**. Complete the Validate a Configuration Wizard as described in the procedure Validate the configuration.
7. On the **Access Point for Administering the Cluster** page, do the following:
   a. In the **Cluster Name** box, enter the name that you want to use to administer the cluster. Before you do, review the following information:
      - During cluster creation, this name is registered as the cluster computer object (also known as the *cluster name object* or *CNO*) in AD DS. If you specify a NetBIOS name for the cluster, the CNO is created in the same location where the computer objects for the cluster nodes reside. This can be either the default Computers container or an OU.
      - To specify a different location for the CNO, you can enter the distinguished name of an OU in the **Cluster Name** box. For example: *CN=ClusterName, OU=Clusters, DC=Contoso, DC=com*.
      - If a domain administrator has prestaged the CNO in a different OU than where the cluster nodes reside, specify the distinguished name that the domain administrator provides.
   b. If the server does not have a network adapter that is configured to use DHCP, you must configure one or more static IP addresses for the failover cluster. Select the check box next to each network that you want to use for cluster management. Click the **Address** field next to a selected network, and then enter the IP address that you want to assign to the cluster. This IP address (or addresses) will be associated with the cluster name in Domain Name System (DNS).
   c. When you are finished, click **Next**.
8. On the **Confirmation** page, review the settings. By default, the **Add all eligible storage to the cluster** check box is selected. Clear this check box if you want to do either of the following:
   - You want to configure storage later.
   - You plan to create clustered storage spaces through Failover Cluster Manager or through the Failover Clustering Windows PowerShell cmdlets, and have not yet created storage spaces in File and Storage Services. For more information, see Deploy Clustered Storage Spaces.
9. Click **Next** to create the failover cluster.
10. On the **Summary** page, confirm that the failover cluster was successfully created. If there were any warnings or errors, view the summary output or click **View Report** to view the full report. Click **Finish**.
11. To confirm that the cluster was created, verify that the cluster name is listed under **Failover Cluster Manager** in the navigation tree. You can expand the cluster name, and

then click items under **Nodes**, **Storage** or **Networks** to view the associated resources.

Realize that it may take some time for the cluster name to successfully replicate in DNS. After successful DNS registration and replication, if you click **All Servers** in Server Manager, the cluster name should be listed as a server with a **Manageability** status of **Online**.

After the cluster is created, you can do things such as verify cluster quorum configuration, and optionally, create Cluster Shared Volumes (CSV). For more information, see Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster and Use Cluster Shared Volumes in a Failover Cluster.

Create clustered roles

After you create the failover cluster, you can create clustered roles to host cluster workloads.

Note

For clustered roles that require a client access point, a virtual computer object (VCO) is created in AD DS. By default, all VCOs for the cluster are created in the same container or OU as the CNO. Realize that after you create a cluster, you can move the CNO to any OU.

To create a clustered role

1. Use Server Manager or Windows PowerShell to install the role or feature that is required for a clustered role on each failover cluster node. For example, if you want to create a clustered file server, install the File Server role on all cluster nodes.

   The following table shows the clustered roles that you can configure in the High Availability Wizard and the associated server role or feature that you must install as a prerequisite.

Clustered Role	Role or Feature Prerequisite
DFS Namespace Server	DFS Namespaces (part of File Server role)
DHCP Server	DHCP Server role
Distributed Transaction Coordinator (DTC)	None
File Server	File Server role
Generic Application	Not applicable
Generic Script	Not applicable
Generic Service	Not applicable
Hyper-V Replica Broker	Hyper-V role

iSCSI Target Server	iSCSI Target Server (part of File Server role)
iSNS Server	iSNS Server Service feature
Message Queuing	Message Queuing Services feature
Other Server	None
Virtual Machine	Hyper-V role
WINS Server	WINS Server feature

2. In Failover Cluster Manager, expand the cluster name, right-click **Roles**, and then click **Configure Role**.
3. Follow the steps in the High Availability Wizard to create the clustered role.
4. To verify that the clustered role was created, in the **Roles** pane, make sure that the role has a status of **Running**. The Roles pane also indicates the owner node. To test failover, right-click the role, point to **Move**, and then click **Select Node**. In the **Move Clustered Role** dialog box, click the desired cluster node, and then click **OK**. In the **Owner Node** column, verify that the owner node changed.

Create a failover cluster by using Windows PowerShell

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlets perform the same function as the preceding procedures in this topic. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines because of formatting constraints.

Note

You must use Windows PowerShell to create an Active Directory-detached cluster in Windows Server 2012 R2. For information about the syntax, see Deploy an Active Directory-Detached Cluster.

The following example installs the Failover Clustering feature.

```
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools
```

The following example runs all cluster validation tests on computers that are named *Server1* and *Server2*.

```
Test-Cluster -Node Server1, Server2
```

Note

The **Test-Cluster** cmdlet outputs the results to a log file in the current working directory, for example C:\Users*username*\AppData\Local\Temp.

The following example creates a failover cluster that is named *MyCluster* with nodes *Server1* and *Server2*, assigns the static IP address *192.168.1.12*, and adds all eligible storage to the failover cluster.

```
New-Cluster –Name MyCluster –Node Server1, Server2 –StaticAddress 192.168.1.12
```

The following example creates the same failover cluster as in the previous example, but it does not add eligible storage to the failover cluster.

```
New-Cluster –Name MyCluster –Node Server1, Server2 –StaticAddress 192.168.1.12 –NoStorage
```

The following example creates a cluster that is named *MyCluster* in the *Cluster* OU of the domain *Contoso.com*.

```
New-Cluster -Name CN=MyCluster,OU=Cluster,DC=Contoso,DC=com -Node Server1, Server2
```

For examples of how to add clustered roles, see topics such as Add-ClusterFileServerRole and Add-ClusterGenericApplicationRole.

See also

- Failover Clustering Overview
- Deploy a Hyper-V Cluster
- Scale-Out File Server for Application Data Overview
- Deploy an Active Directory-Detached Cluster
- Using Guest Clustering for High Availability
- Cluster-Aware Updating Overview
- New-Cluster
- Test-Cluster

Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster

This topic provides background and steps to configure and manage the quorum in a Windows Server 2012 failover cluster.

In this topic

Overview of the quorum in a failover cluster

Overview of the quorum configuration options

- Witness configuration
- Node vote assignment
- Dynamic quorum management

General recommendations for quorum configuration

Configure the cluster quorum

Recover a cluster by starting without quorum

Quorum considerations for disaster recovery configurations

Overview of the quorum in a failover cluster

The quorum for a cluster is determined by the number of voting elements that must be part of active cluster membership for that cluster to start properly or continue running. By default, every node in the cluster has a single quorum vote. In addition, a quorum witness (when configured) has an additional single quorum vote. You can configure one quorum witness for each cluster. A quorum witness can be a designated disk resource or a file share resource. Each element can cast one “vote” to determine whether the cluster can run. Whether a cluster has quorum to function properly is determined by the majority of the voting elements in the active cluster membership.

Why configure the quorum?

To increase the high availability of the cluster, and the roles that are hosted on that cluster, it is important to set the cluster quorum configuration appropriately.

The cluster quorum configuration has a direct effect on the high availability of the cluster, for the following reasons:

- It helps ensure that the failover cluster can start properly or continue running when the active cluster membership changes. Membership changes can occur because of planned or unplanned node shutdown, or when there are disruptions in connectivity between the nodes or with cluster storage.
- When a subset of nodes cannot communicate with another subset of nodes (a split cluster), the cluster quorum configuration helps ensure that only one of the subsets continues running as a cluster. The subsets that do not have enough quorum votes will stop running as a cluster. The subset that has the majority of quorum votes can continue to host clustered roles. This helps avoid partitioning the cluster, so that the same application is not hosted in more than one partition.
- Configuring a witness vote helps the cluster sustain one extra node down in certain configurations. For more information about configuring a quorum witness, see Witness configuration later in this topic.

Be aware that the full function of a cluster depends on quorum in addition to the following factors:

- Network connectivity between cluster nodes
- The capacity of each node to host the clustered roles that get placed on that node
- The priority settings that are configured for the clustered roles

For example, a cluster that has five nodes can have quorum after two nodes fail. However, each remaining node would serve clients only if it had enough capacity to support the clustered roles that failed over to it and if the role settings prioritized the most important workloads.

Overview of the quorum configuration options

The quorum model in Windows Server 2012 is flexible. If you need to modify the quorum configuration for your cluster, you can use the Configure Cluster Quorum Wizard or the Failover Clusters Windows PowerShell cmdlets. The quorum configuration options are simpler than those that are available in Windows Server 2008 R2. For steps and considerations to configure the quorum, see **Configure the cluster quorum** later in this topic.

The following table lists the three quorum configuration options that are available in the Configure Cluster Quorum Wizard.

Option	Description
Use typical settings	The cluster automatically assigns a vote to each node and dynamically manages the node votes. If it is suitable for your cluster, and there is cluster shared storage available, the cluster selects a disk witness. This option is recommended in most cases, because the cluster software automatically chooses a quorum and witness configuration that provides the highest availability for your cluster.
Add or change the quorum witness	You can add, change, or remove a witness resource. You can configure a file share or disk witness. The cluster automatically assigns a vote to each node and dynamically manages the node votes.
Advanced quorum configuration and witness selection	You should select this option only when you have application-specific or site-specific requirements for configuring the quorum. You can modify the quorum witness, add or remove node votes, and choose whether the cluster dynamically manages node votes. By default, votes are assigned to all nodes, and the node votes are dynamically managed.

Depending on the quorum configuration option that you choose and your specific settings, the cluster will be configured in one of the following quorum modes:

Mode	Description
Node majority (no witness)	Only nodes have votes. No quorum witness is configured. The cluster quorum is the majority of voting nodes in the active cluster

Mode	Description
	membership.
Node majority with witness (disk or file share)	Nodes have votes. In addition, a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated disk witness or a designated file share witness.
No majority (disk witness only)	No nodes have votes. Only a disk witness has a vote. The cluster quorum is determined by the state of the disk witness. The cluster has quorum if one node is available and communicating with a specific disk in the cluster storage. Generally, this mode is not recommended, and it should not be selected because it creates a single point of failure for the cluster.

For more information about advanced quorum configuration settings, see the following subsections:

- Witness configuration
- Node vote assignment
- Dynamic quorum management

Witness configuration

As a general rule when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster contains an even number of voting nodes, you should configure a disk witness or a file share witness. The cluster will be able to sustain one additional node down. In addition, adding a witness vote enables the cluster to continue running if half the cluster nodes simultaneously go down or are disconnected.

A disk witness is usually recommended if all nodes can see the disk. A file share witness is recommended when you need to consider multisite disaster recovery with replicated storage. Configuring a disk witness with replicated storage is possible only if the storage vendor supports read-write access from all sites to the replicated storage.

The following table provides additional information and considerations about the quorum witness types.

Witness type	Description	Requirements and recommendations
Disk witness	• Dedicated LUN that stores a copy of the cluster database • Most useful for clusters with shared (not replicated) storage	• Size of LUN must be at least 512 MB • Must be dedicated to cluster use and not assigned to a clustered role • Must be included in clustered storage and pass storage validation tests • Cannot be a disk that is a Cluster Shared Volume (CSV) • Basic disk with a single volume • Does not need to have a drive letter • Can be formatted with NTFS or ReFS • Can be optionally configured with hardware RAID for fault tolerance • Should be excluded from backups and antivirus scanning
File share witness	• SMB file share that is configured on a file server running Windows Server • Does not store a copy of the cluster database • Maintains cluster information only in a witness.log file • Most useful for multisite clusters with replicated storage	• Must have a minimum of 5 MB of free space • Must be dedicated to the single cluster and not used to store user or application data • Must have write permissions enabled for the computer object for the cluster name The following are additional considerations for a file server that hosts the file share witness: • A single file server can be configured with file share witnesses for multiple clusters. • The file server must be on a site that is separate from the

Witness type	Description	Requirements and recommendations
		cluster workload. This allows equal opportunity for any cluster site to survive if site-to-site network communication is lost. If the file server is on the same site, that site becomes the primary site, and it is the only site that can reach the file share. • The file server can run on a virtual machine if the virtual machine is not hosted on the same cluster that uses the file share witness. • For high availability, the file server can be configured on a separate failover cluster.

Node vote assignment

In Windows Server 2012, as an advanced quorum configuration option, you can choose to assign or remove quorum votes on a per-node basis. By default, all nodes are assigned votes. Regardless of vote assignment, all nodes continue to function in the cluster, receive cluster database updates, and can host applications.

You might want to remove votes from nodes in certain disaster recovery configurations. For example, in a multisite cluster, you could remove votes from the nodes in a backup site so that those nodes do not affect quorum calculations. This configuration is recommended only for manual failover across sites. For more information, see Quorum considerations for disaster recovery configurations later in this topic.

The configured vote of a node can be verified by looking up the **NodeWeight** common property of the cluster node by using the Get-ClusterNode Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote configured. A value of 1 indicates that the quorum vote of the node is assigned, and it is managed by the cluster. For more information about management of node votes, see Dynamic quorum management later in this topic.

The vote assignment for all cluster nodes can be verified by using the **Validate Cluster Quorum** validation test.

Additional considerations

- Node vote assignment is not recommended to enforce an odd number of voting nodes. Instead, you should configure a disk witness or file share witness. For more information, see Witness configuration in this topic.
- If dynamic quorum management is enabled, only the nodes that are configured to have node votes assigned can have their votes assigned or removed dynamically. For more information, see Dynamic quorum management later in this topic.

Dynamic quorum management

In Windows Server 2012, as an advanced quorum configuration option, you can choose to enable dynamic quorum management by cluster. When this option is enabled, the cluster dynamically manages the vote assignment to nodes, based on the state of each node. Votes are automatically removed from nodes that leave active cluster membership, and a vote is automatically assigned when a node rejoins the cluster. By default, dynamic quorum management is enabled.

Note

With dynamic quorum management, the cluster quorum majority is determined by the set of nodes that are active members of the cluster at any time. This is an important distinction from the cluster quorum in Windows Server 2008 R2, where the quorum majority is fixed, based on the initial cluster configuration.

With dynamic quorum management, it is also possible for a cluster to run on the last surviving cluster node. By dynamically adjusting the quorum majority requirement, the cluster can sustain sequential node shutdowns to a single node.

The cluster-assigned dynamic vote of a node can be verified with the **DynamicWeight** common property of the cluster node by using the Get-ClusterNode Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote. A value of 1 indicates that the node has a quorum vote.

The vote assignment for all cluster nodes can be verified by using the **Validate Cluster Quorum** validation test.

Additional considerations

- Dynamic quorum management does not allow the cluster to sustain a simultaneous failure of a majority of voting members. To continue running, the cluster must always have a quorum majority at the time of a node shutdown or failure.
- If you have explicitly removed the vote of a node, the cluster cannot dynamically add or remove that vote.

General recommendations for quorum configuration

The cluster software automatically configures the quorum for a new cluster, based on the number of nodes configured and the availability of shared storage. This is usually the most appropriate quorum configuration for that cluster. However, it is a good idea to review the quorum

configuration after the cluster is created, before placing the cluster into production. To view the detailed cluster quorum configuration, you can you use the Validate a Configuration Wizard, or the Test-Cluster Windows PowerShell cmdlet, to run the **Validate Quorum Configuration** test. In Failover Cluster Manager, the basic quorum configuration is displayed in the summary information for the selected cluster, or you can review the information about quorum resources that returns when you run the Get-ClusterQuorum Windows PowerShell cmdlet.

At any time, you can run the **Validate Quorum Configuration** test to validate that the quorum configuration is optimal for your cluster. The test output indicates if a change to the quorum configuration is recommended and the settings that are optimal. If a change is recommended, you can use the Configure Cluster Quorum Wizard to apply the recommended settings.

After the cluster is in production, do not change the quorum configuration unless you have determined that the change is appropriate for your cluster. You might want to consider changing the quorum configuration in the following situations:

- Adding or evicting nodes
- Adding or removing storage
- A long-term node or witness failure
- Recovering a cluster in a multisite disaster recovery scenario

For more information about validating a failover cluster, see Validate Hardware for a Failover Cluster.

Configure the cluster quorum

You can configure the cluster quorum settings by using Failover Cluster Manager or the Failover Clusters Windows PowerShell cmdlets.

Important
It is usually best to use the quorum configuration that is recommended by the Configure Cluster Quorum Wizard. We recommend customizing the quorum configuration only if you have determined that the change is appropriate for your cluster. For more information, see General recommendations for quorum configuration in this topic.

Configure the cluster quorum settings

Membership in the local **Administrators** group on each clustered server, or equivalent, is the minimum permissions required to complete this procedure. Also, the account you use must be a domain user account.

Note
You can change the cluster quorum configuration without stopping the cluster or taking cluster resources offline.

To change the quorum configuration in a failover cluster by using Failover Cluster Manager

1. In Failover Cluster Manager, select or specify the cluster that you want to change.
2. With the cluster selected, under **Actions**, click **More Actions**, and then click **Configure Cluster Quorum Settings**. The Configure Cluster Quorum Wizard appears. Click **Next**.
3. On the **Select Quorum Configuration Option** page, select one of the three configuration options and complete the steps for that option. Before you configure the quorum settings, you can review your choices. For more information about the options, see Overview of the quorum in a failover cluster, earlier in this topic.
   - To allow the cluster to automatically reset the quorum settings that are optimal for your current cluster configuration, click **Use typical settings** and then complete the wizard.
   - To add or change the quorum witness, click **Add or change the quorum witness**, and then complete the following steps. For information and considerations about configuring a quorum witness, see Witness configuration earlier in this topic.
     i. On the **Select Quorum Witness** page, select an option to configure a disk witness or a file share witness. The wizard indicates the witness selection options that are recommended for your cluster.

        Note
        You can also select **Do not configure a quorum witness** and then complete the wizard. If you have an even number of voting nodes in your cluster, this may not be a recommended configuration.

     ii. If you select the option to configure a disk witness, on the **Configure Storage Witness** page, select the storage volume that you want to assign as the disk witness, and then complete the wizard.

     iii. If you select the option to configure a file share witness, on the **Configure File Share Witness** page, type or browse to a file share that will be used as the witness resource, and then complete the wizard.
   - To configure quorum management settings and to add or change the quorum witness, click **Advanced quorum configuration and witness selection**, and then complete the following steps. For information and considerations about the advanced quorum configuration settings, see Node vote assignment and Dynamic quorum management earlier in this topic.
     i. On the **Select Voting Configuration** page, select an option to assign votes to nodes. By default, all nodes are assigned a vote. However, for certain scenarios, you can assign votes only to a subset of the nodes.

        Note
        You can also select **No Nodes**. This is generally not recommended, because it does not allow nodes to participate in quorum voting, and it requires configuring a disk witness. This disk witness becomes the single point of failure for the cluster.

     ii. On the **Configure Quorum Management** page, you can enable or disable the **Allow cluster to dynamically manage the assignment of node votes** option. Selecting this option generally increases the availability of the cluster. By default

the option is enabled, and it is strongly recommended to not disable this option. This option allows the cluster to continue running in failure scenarios that are not possible when this option is disabled.

iii. On the **Select Quorum Witness** page, select an option to configure a disk witness or a file share witness. The wizard indicates the witness selection options that are recommended for your cluster.

Note
You can also select **Do not configure a quorum witness**, and then complete the wizard. If you have an even number of voting nodes in your cluster, this may not be a recommended configuration.

iv. If you select the option to configure a disk witness, on the **Configure Storage Witness** page, select the storage volume that you want to assign as the disk witness, and then complete the wizard.

v. If you select the option to configure a file share witness, on the **Configure File Share Witness** page, type or browse to a file share that will be used as the witness resource, and then complete the wizard.

4. Click **Next**. Confirm your selections on the confirmation page that appears, and then click **Next**.

After the wizard runs and the **Summary** page appears, if you want to view a report of the tasks that the wizard performed, click **View Report**. The most recent report will remain in the *systemroot***Cluster\Reports** folder with the name **QuorumConfiguration.mht**.

Note
After you configure the cluster quorum, we recommend that you run the **Validate Quorum Configuration** test to verify the updated quorum settings.

Windows PowerShell equivalent commands

The following examples show how to use the Set-ClusterQuorum cmdlet and other Windows PowerShell cmdlets to configure the cluster quorum.

The following example changes the quorum configuration on cluster *CONTOSO-FC1* to a simple node majority configuration with no quorum witness.

```
Set-ClusterQuorum –Cluster CONTOSO-FC1 -NodeMajority
```

The following example changes the quorum configuration on the local cluster to a node majority with witness configuration. The disk resource named *Cluster Disk 2* is configured as a disk witness.

```
Set-ClusterQuorum -NodeAndDiskMajority "Cluster Disk 2"
```

The following example changes the quorum configuration on the local cluster to a node majority with witness configuration. The file share resource named *\\CONTOSO-FS\fsw* is configured as a file share witness.

```
Set-ClusterQuorum -NodeAndFileShareMajority "\\fileserver\fsw"
```

The following example removes the quorum vote from node *ContosoFCNode1* on the local cluster.

```
(Get-ClusterNode ContosoFCNode1).NodeWeight=0
```

The following example adds the quorum vote to node *ContosoFCNode1* on the local cluster.

```
(Get-ClusterNode ContosoFCNode1).NodeWeight=1
```

The following example enables the **DynamicQuorum** property of the cluster *CONTOSO-FC1* (if it was previously disabled):

```
(Get-Cluster CONTOSO-FC1).DynamicQuorum=1
```

Recover a cluster by starting without quorum

A cluster that does not have enough quorum votes will not start. As a first step, you should always confirm the cluster quorum configuration and investigate why the cluster no longer has quorum. This might happen if you have nodes that stopped responding, or if the primary site is not reachable in a multisite cluster. After you identify the root cause for the cluster failure, you can use the recovery steps described in this section.

Notes

- If the Cluster service stops because quorum is lost, Event ID 1177 appears in the system log.
- It is always necessary to investigate why the cluster quorum was lost.
- It is always preferable to bring a node or quorum witness to a healthy state (join the cluster) rather than starting the cluster without quorum.

Force start cluster nodes

After you determine that you cannot recover your cluster by bringing the nodes or quorum witness to a healthy state, forcing your cluster to start becomes necessary. Forcing the cluster to start overrides your cluster quorum configuration settings and starts the cluster in **ForceQuorum** mode.

Forcing a cluster to start when it does not have quorum may be especially useful in a multisite cluster. Consider a disaster recovery scenario with a cluster that contains separately located primary and backup sites, *SiteA* and *SiteB*. If there is a genuine disaster at *SiteA*, it could take a significant amount of time for the site to come back online. You would likely want to force *SiteB* to come online, even though it does not have quorum.

When a cluster is started in **ForceQuorum** mode, and after it regains sufficient quorum votes, the cluster automatically leaves the forced state, and it behaves normally. Hence, it is not necessary to start the cluster again normally. If the cluster loses a node and it loses quorum, it goes offline again because it is no longer in the forced state. To bring it back `online when it does not have quorum requires forcing the cluster to start without quorum.

Important

- After a cluster is force started, the administrator is in full control of the cluster.
- The cluster uses the cluster configuration on the node where the cluster is force started, and replicates it to all other nodes that are available.
- If you force the cluster to start without quorum, all quorum configuration settings are ignored while the cluster remains in **ForceQuorum** mode. This includes specific node vote assignments and dynamic quorum management settings.

Prevent quorum on remaining cluster nodes

After you have force started the cluster on a node, it is necessary to start any remaining nodes in your cluster with a setting to prevent quorum. A node started with a setting that prevents quorum indicates to the Cluster service to join an existing running cluster instead of forming a new cluster instance. This prevents the remaining nodes from forming a split cluster that contains two competing instances.

This becomes necessary when you need to recover your cluster in some multisite disaster recovery scenarios after you have force started the cluster on your backup site, *SiteB*. To join the force started cluster in *SiteB*, the nodes in your primary site, *SiteA*, need to be started with the quorum prevented.

Important
After a cluster is force started on a node, we recommend that you always start the remaining nodes with the quorum prevented.

To recover the cluster by using Failover Cluster Manager

1. In Failover Cluster Manager, select or specify the cluster you want to recover.
2. With the cluster selected, under **Actions**, click **Force Cluster Start**.

   Failover Cluster Manager force starts the cluster on all nodes that are reachable. The cluster uses the current cluster configuration when starting.

   Notes
   - To force the cluster to start on a specific node that contains a cluster configuration that you want to use, you must use the Windows PowerShell cmdlets or equivalent command-line tools as presented after this procedure.
   - If you use Failover Cluster Manager to connect to a cluster that is force started, and you use the **Start Cluster Service** action to start a node, the node is automatically started with the setting that prevents quorum.

Windows PowerShell equivalent commands

The following example shows how to use the **Start-ClusterNode** cmdlet to force start the cluster on node *ContosoFCNode1*.

```
Start-ClusterNode –Node ContosoFCNode1 –FQ
```

Alternatively, you can type the following command locally on the node:

```
Net Start ClusSvc /FQ
```

The following example shows how to use the **Start-ClusterNode** cmdlet to start the Cluster service with the quorum prevented on node *ContosoFCNode1*.

```
Start-ClusterNode –Node ContosoFCNode1 –PQ
```

Alternatively, you can type the following command locally on the node:

```
Net Start ClusSvc /PQ
```

Quorum considerations for disaster recovery configurations

This section summarizes characteristics and quorum configurations for two multisite cluster configurations in disaster recovery deployments. The quorum configuration guidelines differ depending on if you need automatic failover or manual failover for workloads between the sites. Your configuration usually depends on the service level agreements (SLAs) that are in place in your organization to provide and support clustered workloads in the event of a failure or disaster at a site.

Automatic failover

In this configuration, the cluster consists of two or more sites that can host clustered roles. If a failure occurs at any site, the clustered roles are expected to automatically fail over to the remaining sites. Therefore, the cluster quorum must be configured so that any site can sustain a complete site failure.

The following table summarizes considerations and recommendations for this configuration.

Item	Description
Number of node votes per site	Should be equal
Node vote assignment	Node votes should not be removed because all nodes are equally important
Dynamic quorum management	Should be enabled
Witness configuration	File share witness is recommended, configured in a site that is separate from the cluster sites
Workloads	Workloads can be configured on any of the sites

Additional considerations

- Configuring the file share witness in a separate site is necessary to give each site an equal opportunity to survive. For more information, see Witness configuration earlier in this topic.

Manual failover

In this configuration, the cluster consists of a primary site, *SiteA*, and a backup (recovery) site, *SiteB*. Clustered roles are hosted on *SiteA*. Because of the cluster quorum configuration, if a failure occurs at all nodes in *SiteA*, the cluster stops functioning. In this scenario the administrator must manually fail over the cluster services to *SiteB* and perform additional steps to recover the cluster.

The following table summarizes considerations and recommendations for this configuration.

Item	Description
Number of node votes per site	Can differ
Node vote assignment	• Node votes should not be removed from nodes at the primary site, *SiteA* • Node votes should be removed from nodes at the backup site, *SiteB* • If a long-term outage occurs at *SiteA*, votes must be assigned to nodes at *SiteB* to enable a quorum majority at that site as part of recovery
Dynamic quorum management	Should be enabled
Witness configuration	• Configure a witness if there is an even number of nodes at *SiteA* • If a witness is needed, configure either a file share witness or a disk witness that is accessible only to nodes in *SiteA* (sometimes called an asymmetric disk witness)
Workloads	Use preferred owners to keep workloads running on nodes at *SiteA*

Additional considerations

- Only the nodes at *SiteA* are initially configured with quorum votes. This is necessary to ensure that the state of nodes at *SiteB* does not affect the cluster quorum.
- Recovery steps can vary depending on if *SiteA* sustains a temporary failure or a long-term failure.

See Also

Failover Clustering Overview

Failover Clusters Windows PowerShell cmdlets

Use Cluster Shared Volumes in a Failover Cluster

Cluster Shared Volumes (CSV) enable multiple nodes in a failover cluster to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume. (In Windows Server 2012 R2, the disk can be provisioned as NTFS or Resilient File System (ReFS).) With CSV, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership, or dismounting and remounting a volume. CSV also help simplify the management of a potentially large number of LUNs in a failover cluster.

CSV provide a general-purpose, clustered file system, which is layered above NTFS (or ReFS in Windows Server 2012 R2). CSV applications include:

- Clustered virtual hard disk (VHD) files for clustered Hyper-V virtual machines
- Scale-out file shares to store application data for the Scale-Out File Server clustered role. Examples of the application data for this role include Hyper-V virtual machine files and Microsoft SQL Server data. (Be aware that ReFS is not supported for a Scale-Out File Server.) For more information about Scale-Out File Server, see Scale-Out File Server for Application Data Overview.

Note
CSV does not support the Microsoft SQL Server clustered workload in SQL Server 2012 and earlier versions of SQL Server.

In Windows Server 2012, CSV functionality was significantly enhanced. For example, dependencies on Active Directory Domain Services were removed. Support was added for the functional improvements in **chkdsk**, for interoperability with antivirus and backup applications, and for integration with general storage features such as BitLocker-encrypted volumes and Storage Spaces. For an overview of CSV functionality that was introduced in Windows Server 2012, see What's New in Failover Clustering in Windows Server 2012.

Windows Server 2012 R2 introduces additional functionality, such as distributed CSV ownership, increased resiliency through availability of the Server service, greater flexibility in the amount of physical memory that you can allocate to CSV cache, better diagnosibility, and enhanced interoperability that includes support for ReFS and deduplication. For more information, see What's New in Failover Clustering in Windows Server 2012 R2.

Note
For information about using data deduplication on CSV for Virtual Desktop Infrastructure (VDI) scenarios, see the blog posts Deploying Data Deduplication for VDI storage in Windows Server 2012 R2 and Extending Data Deduplication to new workloads in Windows Server 2012 R2.

In this topic

- Review requirements and considerations for using CSV in a failover cluster
  - **Network configuration considerations**

    - Storage and disk configuration requirements
    - Node requirements
- Plan to use CSV in a failover cluster
    - Arrangement of LUNs, volumes, and VHD files
    - Number and size of LUNs and volumes
- Add a disk to CSV on a failover cluster
- Enable the CSV cache for read-intensive workloads (optional)
- Back up CSV

Review requirements and considerations for using CSV in a failover cluster

Before using CSV in a failover cluster, review the network, storage, and other requirements and considerations in this section.

In this section

- **Network configuration considerations**
- Storage and disk configuration requirements
- Node requirements

Network configuration considerations

Consider the following when you configure the networks that support CSV.

- **Multiple networks and multiple network adapters**. To enable fault tolerance in the event of a network failure, we recommend that multiple cluster networks carry CSV traffic or that you configure teamed network adapters.

  If the cluster nodes are connected to networks that should not be used by the cluster, you should disable them. For example, we recommend that you disable iSCSI networks for cluster use to prevent CSV traffic on those networks. To disable a network, in Failover Cluster Manager, click **Networks**, click the network, click the **Properties** action, and then select **Do not allow cluster network communication on this network**. Alternatively, you can configure the **Role** property of the network by using the Get-ClusterNetwork Windows PowerShell cmdlet.
- **Network adapter properties**. In the properties for all adapters that carry cluster communication, make sure that the following settings are enabled:
    - **Client for Microsoft Networks** and **File and Printer Sharing for Microsoft Networks**. These settings support Server Message Block (SMB) 3.0, which is used by default to carry CSV traffic between nodes. To enable SMB, also ensure that the Server service and the Workstation service are running and that they are configured to start automatically on each cluster node.

      Note

In Windows Server 2012 R2, there are multiple Server service instances per failover cluster node. There is the default instance that handles incoming traffic from SMB clients that access regular file shares, and a second CSV instance that handles only inter-node CSV traffic. Also, if the Server service on a node becomes unhealthy, CSV ownership automatically transitions to another node.

SMB 3.0 includes the SMB Multichannel and SMB Direct features, which enable CSV traffic to stream across multiple networks in the cluster and to leverage network adapters that support Remote Direct Memory Access (RDMA). By default, SMB Multichannel is used for CSV traffic. For more information, see Server Message Block overview.

- **Microsoft Failover Cluster Virtual Adapter Performance Filter**. This setting improves the ability of nodes to perform I/O redirection when it is required to reach CSV, for example, when a connectivity failure prevents a node from connecting directly to the CSV disk. For more information, see About I/O synchronization and I/O redirection in CSV communication in this topic.

- **Cluster network prioritization**. We generally recommend that you do not change the cluster-configured preferences for the networks.
- **IP subnet configuration**. No specific subnet configuration is required for nodes in a network that use CSV. CSV can support multisubnet clusters.
- **Policy-based Quality of Service (QoS)**. We recommend that you configure a QoS priority policy and a minimum bandwidth policy for network traffic to each node when you use CSV. For more information, see Quality of Service (QoS) Overview.
- **Storage network**. For storage network recommendations, review the guidelines that are provided by your storage vendor. For additional considerations about storage for CSV, see Storage and disk configuration requirements, later in this topic.

For an overview of the hardware, network, and storage requirements for failover clusters, see Failover Clustering Hardware Requirements and Storage Options.

About I/O synchronization and I/O redirection in CSV communication

I/O synchronization CSV enable multiple nodes to have simultaneous read-write access to the same shared storage. When a node performs disk input/output (I/O) on a CSV volume, the node communicates directly with the storage, for example, through a storage area network (SAN). However, at any time, a single node (called the coordinator node) “owns” the physical disk resource that is associated with the LUN. The coordinator node for a CSV volume is displayed in Failover Cluster Manager as **Owner Node** under **Disks**. It also appears in the output of the Get-ClusterSharedVolume Windows PowerShell cmdlet.

Note

In Windows Server 2012 R2, CSV ownership is evenly distributed across the failover cluster nodes based on the number of CSV volumes that each node owns. Additionally, ownership is automatically rebalanced when there are conditions such as CSV failover, a node rejoins the cluster, you add a new node to the cluster, you restart a cluster node, or you start the failover cluster after it has been shut down.

When certain small changes occur in the file system on a CSV volume, this metadata must be synchronized on each of the physical nodes that access the LUN, not only on the single coordinator node. For example, when a virtual machine on a CSV volume is started, created, or deleted, or when a virtual machine is migrated, this information needs to be synchronized on each of the physical nodes that access the virtual machine. These metadata update operations occur in parallel across the cluster networks by using SMB 3.0. These operations do not require all the physical nodes to communicate with the shared storage.

I/O redirection Storage connectivity failures and certain storage operations can prevent a given node from communicating directly with the storage. To maintain function while the node is not communicating with the storage, the node redirects the disk I/O through a cluster network to the coordinator node where the disk is currently mounted. If the current coordinator node experiences a storage connectivity failure, all disk I/O operations are queued temporarily while a new node is established as a coordinator node.

The server uses one of the following I/O redirection modes, depending on the situation:

- **File system redirection** Redirection is per volume—for example, when CSV snapshots are taken by a backup application when a CSV volume is manually placed in redirected I/O mode.
- **Block redirection** Redirection is at the file-block level—for example, when storage connectivity is lost to a volume. Block redirection is significantly faster than file system redirection.

In Windows Server 2012 R2, you can view the state of a CSV volume on a per node basis. For example, you can see whether I/O is direct or redirected, or whether the CSV volume is unavailable. If a CSV volume is in I/O redirected mode, you can also view the reason. Use the Windows PowerShell cmdlet **Get-ClusterSharedVolumeState** to view this information.

Notes

- In Windows Server 2012, because of improvements in CSV design, CSV perform more operations in direct I/O mode than occurred in Windows Server 2008 R2.
- Because of the integration of CSV with SMB 3.0 features such as SMB Multichannel and SMB Direct, redirected I/O traffic can stream across multiple cluster networks.
- You should plan your cluster networks to allow for the potential increase in network traffic to the coordinator node during I/O redirection.

Storage and disk configuration requirements

To use CSV, your storage and disks must meet the following requirements:

- **File system format**. In Windows Server 2012 R2, a disk or storage space for a CSV volume must be a basic disk that is partitioned with NTFS or ReFS. In Windows Server 2012, a disk or storage space for a CSV volume must be a basic disk that is partitioned with NTFS.

  A CSV has the following additional requirements:

  - In Windows Server 2012 R2, you cannot use a disk for a CSV that is formatted with FAT or FAT32.

    - In Windows Server 2012, you cannot use a disk for a CSV that is formatted with FAT, FAT32, or ReFS.
    - If you want to use a storage space for a CSV, you can configure a simple space or a mirror space. In Windows Server 2012 R2, you can also configure a parity space. (In Windows Server 2012, CSV does not support parity spaces.)
    - A CSV cannot be used as a quorum witness disk. For more information about the cluster quorum, see Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster.
    - After you add a disk as a CSV, it is designated in the CSVFS format (for CSV File System). This allows the cluster and other software to differentiate the CSV storage from other NTFS or ReFS storage. Generally, CSVFS supports the same functionality as NTFS or ReFS. However, certain features are not supported. For example, in Windows Server 2012 R2, you cannot enable compression on CSV. In Windows Server 2012, you cannot enable data deduplication or compression on CSV.
- **Resource type in the cluster**. For a CSV volume, you must use the Physical Disk resource type. By default, a disk or storage space that is added to cluster storage is automatically configured in this way.
- **Choice of CSV disks or other disks in cluster storage**. When choosing one or more disks for a clustered virtual machine, consider how each disk will be used. If a disk will be used to store files that are created by Hyper-V, such as VHD files or configuration files, you can choose from the CSV disks or the other available disks in cluster storage. If a disk will be a physical disk that is directly attached to the virtual machine (also called a pass-through disk), you cannot choose a CSV disk, and you must choose from the other available disks in cluster storage.
- **Path name for identifying disks**. Disks in CSV are identified with a path name. Each path appears to be on the system drive of the node as a numbered volume under the **\ClusterStorage** folder. This path is the same when viewed from any node in the cluster. You can rename the volumes if needed.

For storage requirements for CSV, review the guidelines that are provided by your storage vendor. For additional storage planning considerations for CSV, see Plan to use CSV in a failover cluster later in this topic.

Node requirements

To use CSV, your nodes must meet the following requirements:

- **Drive letter of system disk**. On all nodes, the drive letter for the system disk must be the same.
- **Authentication protocol**. The NTLM protocol must be enabled on all nodes. This is enabled by default.

Plan to use CSV in a failover cluster

This section lists planning considerations and recommendations for using CSV in a failover cluster running Windows Server 2012 R2 or Windows Server 2012.

Important
Ask your storage vendor for recommendations about how to configure your specific storage unit for CSV. If the recommendations from the storage vendor differ from information in this topic, use the recommendations from the storage vendor.

In this section

- Arrangement of LUNs, volumes, and VHD files
- Number and size of LUNs and volumes

Arrangement of LUNs, volumes, and VHD files

To make the best use of CSV to provide storage for clustered virtual machines, it is helpful to review how you would arrange the LUNs (disks) when you configure physical servers. When you configure the corresponding virtual machines, try to arrange the VHD files in a similar way.

Consider a physical server for which you would organize the disks and files as follows:

- System files, including a page file, on one physical disk
- Data files on another physical disk

For an equivalent clustered virtual machine, you should organize the volumes and files in a similar way:

- System files, including a page file, in a VHD file on one CSV
- Data files in a VHD file on another CSV

If you add another virtual machine, where possible, you should keep the same arrangement for the VHDs on that virtual machine.

Number and size of LUNs and volumes

When you plan the storage configuration for a failover cluster that uses CSV, consider the following recommendations:

- To decide how many LUNs to configure, consult your storage vendor. For example, your storage vendor may recommend that you configure each LUN with one partition and place one CSV volume on it.
- There are no limitations for the number of virtual machines that can be supported on a single CSV volume. However, you should consider the number of virtual machines that you plan to have in the cluster and the workload (I/O operations per second) for each virtual machine. Consider the following examples:
    - One organization is deploying virtual machines that will support a virtual desktop infrastructure (VDI), which is a relatively light workload. The cluster uses high-performance storage. The cluster administrator, after consulting with the storage vendor, decides to place a relatively large number of virtual machines per CSV volume.
    - Another organization is deploying a large number of virtual machines that will support a heavily used database application, which is a heavier workload. The cluster uses lower-performing storage. The cluster administrator, after consulting with the storage vendor, decides to place a relatively small number of virtual machines per CSV volume.

- When you plan the storage configuration for a particular virtual machine, consider the disk requirements of the service, application, or role that the virtual machine will support. Understanding these requirements helps you avoid disk contention that can result in poor performance. The storage configuration for the virtual machine should closely resemble the storage configuration that you would use for a physical server that is running the same service, application, or role. For more information, see Arrangement of LUNs, volumes, and VHD files earlier in this topic.

  You can also mitigate disk contention by having storage with a large number of independent physical hard disks. Choose your storage hardware accordingly, and consult with your vendor to optimize the performance of your storage.
- Depending on your cluster workloads and their need for I/O operations, you can consider configuring only a percentage of the virtual machines to access each LUN, while other virtual machines do not have connectivity and are instead dedicated to compute operations.

Add a disk to CSV on a failover cluster

The CSV feature is enabled by default in Failover Clustering. To add a disk to CSV, you must add a disk to the **Available Storage** group of the cluster (if it is not already added), and then add the disk to CSV on the cluster. You can use Failover Cluster Manager or the Failover Clusters Windows PowerShell cmdlets to perform these procedures.

To add a disk to Available Storage

1. In Failover Cluster Manager, in the console tree, expand the name of the cluster, and then expand **Storage**.
2. Right-click **Disks**, and then click **Add Disk**. A list appears showing the disks that can be added for use in a failover cluster.
3. Select the disk or disks you want to add, and then click **OK**.

   The disks are now assigned to the **Available Storage** group.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example identifies the disks that are ready to be added to the cluster, and then adds them to the **Available Storage** group.

```
Get-ClusterAvailableDisk | Add-ClusterDisk
```

To add a disk in Available Storage to CSV

1. In Failover Cluster Manager, in the console tree, expand the name of the cluster, expand **Storage**, and then click **Disks**.
2. Select one or more disks that are assigned to **Available Storage**, right-click the selection, and then click **Add to Cluster Shared Volumes**.

The disks are now assigned to the **Cluster Shared Volume** group in the cluster. The disks are exposed to each cluster node as numbered volumes (mount points) under the %SystemDisk%ClusterStorage folder. The volumes appear in the CSVFS file system.

Note

You can rename CSV volumes in the %SystemDisk%ClusterStorage folder.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example adds *Cluster Disk 1* in **Available Storage** to CSV on the local cluster.

```
Add-ClusterSharedVolume –Name "Cluster Disk 1"
```

Enable the CSV cache for read-intensive workloads (optional)

The CSV cache provides caching at the block level of read-only unbuffered I/O operations by allocating system memory (RAM) as a write-through cache. (Unbuffered I/O operations are not cached by the cache manager.) This can improve performance for applications such as Hyper-V, which conducts unbuffered I/O operations when accessing a VHD. The CSV cache can boost the performance of read requests without caching write requests. Enabling the CSV cache is also useful for Scale-Out File Server scenarios.

Note

We recommend that you enable the CSV cache for all clustered Hyper-V and Scale-Out File Server deployments.

By default in Windows Server 2012, the CSV cache is disabled. In Windows Server 2012 R2, the CSV cache is enabled by default. However, you must still allocate the size of the block cache to reserve.

The following table describes the two configuration settings that control the CSV cache.

Property Name in Windows Server 2012 R2	Property Name in Windows Server 2012	Description
BlockCacheSize	**SharedVolumeBlockCacheSizeInMB**	This is a cluster common property that allows you to define how much memory (in megabytes) to reserve for the CSV cache on each node in the cluster. For example, if

Property Name in Windows Server 2012 R2	Property Name in Windows Server 2012	Description
		a value of 512 is defined, then 512 MB of system memory is reserved on each node. (In many clusters, 512 MB is a recommended value.) The default setting is 0 (for disabled).
EnableBlockCache	**CsvEnableBlockCache**	This is a private property of the cluster Physical Disk resource. It allows you to enable CSV cache on an individual disk that is added to CSV. In Windows Server 2012, the default setting is 0 (for disabled). To enable CSV cache on a disk, configure a value of 1. By default, in Windows Server 2012 R2, this setting is enabled.

You can monitor the CSV cache in Performance Monitor by adding the counters under **Cluster CSV Volume Cache**.

To configure the CSV cache

1. Start Windows PowerShell as an administrator.
2. To define a cache of *512* MB to be reserved on each node, type the following:
   - For Windows Server 2012 R2:

     ```
     (Get-Cluster).BlockCacheSize = 512
     ```

   - For Windows Server 2012:

     ```
     (Get-Cluster).SharedVolumeBlockCacheSizeInMB = 512
     ```

3. In Windows Server 2012, to enable the CSV cache on a CSV named *Cluster Disk 1*, type the following:

```
Get-ClusterSharedVolume "Cluster Disk 1" | Set-
ClusterParameter CsvEnableBlockCache 1
```

Notes

- In Windows Server 2012, you can allocate only 20% of the total physical RAM to the CSV cache. In Windows Server 2012 R2, you can allocate up to 80%. Because Scale-Out File Servers are not typically memory constrained, you can accomplish large performance gains by using the extra memory for the CSV cache.
- In Windows Server 2012, to avoid resource contention, you should restart each node in the cluster after you modify the memory that is allocated to the CSV cache. In Windows Server 2012 R2, a restart is no longer required.
- After you enable or disable CSV cache on an individual disk, for the setting to take effect, you must take the Physical Disk resource offline and bring it back online. (By default, in Windows Server 2012 R2, the CSV cache is enabled.)
- For more information about CSV cache that includes information about performance counters, see the blog post How to Enable CSV Cache.

Back up CSV

There are multiple methods to back up information that is stored on CSV in a failover cluster. You can use a Microsoft backup application or a non-Microsoft application. In general, CSV do not impose special backup requirements beyond those for clustered storage formatted with NTFS or ReFS. CSV backups also do not disrupt other CSV storage operations.

You should consider the following factors when you select a backup application and backup schedule for CSV:

- Volume-level backup of a CSV volume can be run from any node that connects to the CSV volume.
- Your backup application can use software snapshots or hardware snapshots. Depending on the ability of your backup application to support them, backups can use application-consistent and crash-consistent Volume Shadow Copy Service (VSS) snapshots.
- If you are backing up CSV that have multiple running virtual machines, you should generally choose a management operating system-based backup method. If your backup application supports it, multiple virtual machines can be backed up simultaneously.
- CSV support backup requestors that are running Windows Server 2012 R2 Backup, Windows Server 2012 Backup or Windows Server 2008 R2 Backup. However, Windows Server Backup generally provides only a basic backup solution that may not be suited for organizations with larger clusters. Windows Server Backup does not support application-consistent virtual machine backup on CSV. It supports crash-consistent volume-level backup only. (If you restore a crash-consistent backup, the virtual machine will be in the same state it was if the virtual machine had crashed at the exact moment that the backup was taken.) A backup of a virtual machine on a CSV volume will succeed, but an error event will be logged indicating that this is not supported.
- You may require administrative credentials when backing up a failover cluster.

Important
Be sure to carefully review what data your backup application backs up and restores, which CSV features it supports, and the resource requirements for the application on each cluster node.

Caution
If you need to restore the backup data onto a CSV volume, be aware of the capabilities and limitations of the backup application to maintain and restore application-consistent data across the cluster nodes. For example, with some applications, if the CSV is restored on a node that is different from the node where the CSV volume was backed up, you might inadvertently overwrite important data about the application state on the node where the restore is taking place.

See also

- Failover Clustering Overview
- Deploy Clustered Storage Spaces

Prestage Cluster Computer Objects in Active Directory Domain Services

This topic shows how to prestage cluster computer objects in Active Directory Domain Services (AD DS). You can use this procedure to enable a user or group to create a failover cluster when they do not have permissions to create computer objects in AD DS.

When you create a failover cluster by using the Create Cluster Wizard or by using Windows PowerShell, you must specify a name for the cluster. If you have sufficient permissions when you create the cluster, the cluster creation process automatically creates a computer object in AD DS that matches the cluster name. This object is called the *cluster name object* or CNO. Through the CNO, virtual computer objects (VCOs) are automatically created when you configure clustered roles that use client access points. For example, if you create a highly available file server with a client access point that is named *FileServer1*, the CNO will create a corresponding VCO in AD DS.

Note
In Windows Server 2012 R2, there is the option to create an Active Directory-detached cluster, where no CNO or VCOs are created in AD DS. This is targeted for specific types of cluster deployments. For more information, see Deploy an Active Directory-Detached Cluster.

To create the CNO automatically, the user who creates the failover cluster must have the **Create Computer objects** permission to the organizational unit (OU) or the container where the servers that will form the cluster reside. To enable a user or group to create a cluster without having this

permission, a user with appropriate permissions in AD DS (typically a domain administrator) can prestage the CNO in AD DS. This also provides the domain administrator more control over the naming convention that is used for the cluster, and control over which OU the cluster objects are created in.

The procedures in this topic show how to do the following:

- Step 1: Prestage the CNO in AD DS
- Step 2: Grant the user permissions to create the cluster
- Step 3: Grant the CNO permissions to the OU or prestage VCOs for clustered roles

Step 1: Prestage the CNO in AD DS

Before you begin, make sure that you know the following:

- the name that you want to assign the cluster
- the name of the user account or group to which you want to grant rights to create the cluster

As a best practice, we recommend that you create an OU for the cluster objects. If an OU already exists that you want to use, membership in the **Account Operators** group is the minimum required to complete this step. If you need to create an OU for the cluster objects, membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this step.

Note

If you create the CNO in the default Computers container instead of an OU, you do not have to complete Step 3 of this topic. In this scenario, a cluster administrator can create up to 10 VCOs without any additional configuration.

To prestage the CNO in AD DS

1. On a computer that has the AD DS Tools installed from the Remote Server Administration Tools, or on a domain controller, open **Active Directory Users and Computers**. To do this on a server, start Server Manager, and then on the **Tools** menu, click **Active Directory Users and Computers**.
2. To create an OU for the cluster computer objects, right-click the domain name or an existing OU, point to **New**, and then click **Organizational Unit**. In the **Name** box, enter the name of the OU, and then click **OK**.
3. In the console tree, right-click the OU where you want to create the CNO, point to **New**, and then click **Computer**.
4. In the **Computer name** box, enter the name that will be used for the failover cluster, and then click **OK**.

   Note

   This is the cluster name that the user who creates the cluster will specify on the **Access Point for Administering the Cluster** page in the Create Cluster wizard or as the value of the **–Name** parameter for the **New-Cluster** Windows PowerShell cmdlet.

5. As a best practice, right-click the computer account that you just created, click **Properties**, and then click the **Object** tab. On the **Object** tab, select the **Protect object from accidental deletion** check box, and then click **OK**.
6. Right-click the computer account that you just created, and then click **Disable Account**. Click **Yes** to confirm, and then click **OK**.

Note

You must disable the account so that during cluster creation, the cluster creation process can confirm that the account is not currently in use by an existing computer or cluster in the domain.

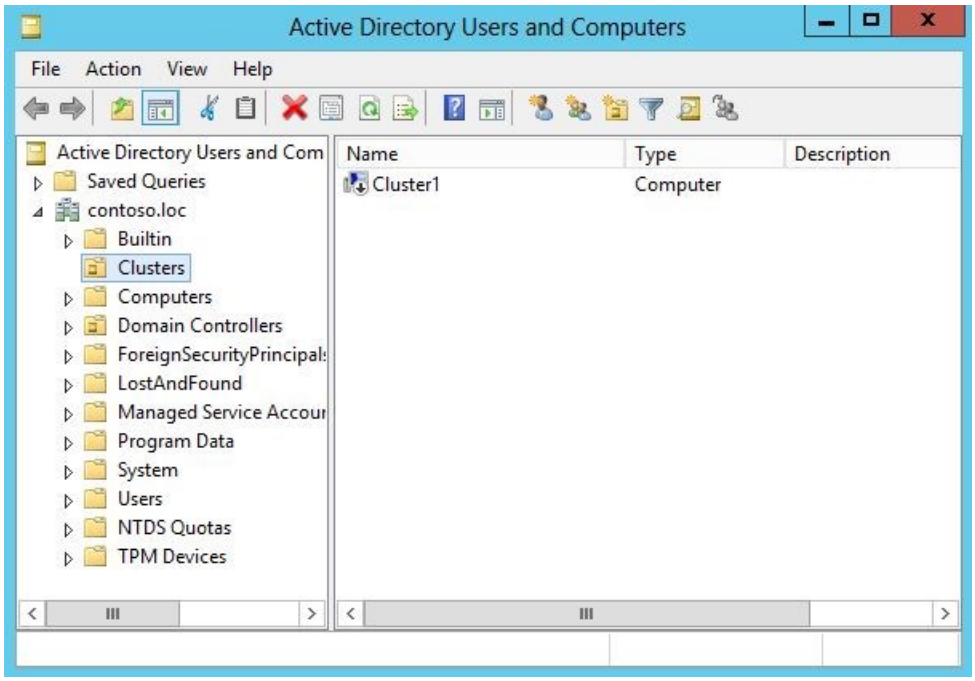

Figure 1. Disabled CNO in the example Clusters OU

Step 2: Grant the user permissions to create the cluster

You must configure permissions so that the user account that will be used to create the failover cluster has Full Control permissions to the CNO.

Membership in the **Account Operators** group is the minimum required to complete this step.

To grant the user permissions to create the cluster

1. In Active Directory Users and Computers, on the **View** menu, make sure that **Advanced Features** is selected.
2. Locate and then right-click the CNO, and then click **Properties**.

3. On the **Security** tab, click **Add**.
4. In the **Select Users, Computers, or Groups** dialog box, specify the user account or group that you want to grant permissions to, and then click **OK**.
5. Select the user account or group that you just added, and then next to **Full control**, select the **Allow** check box.

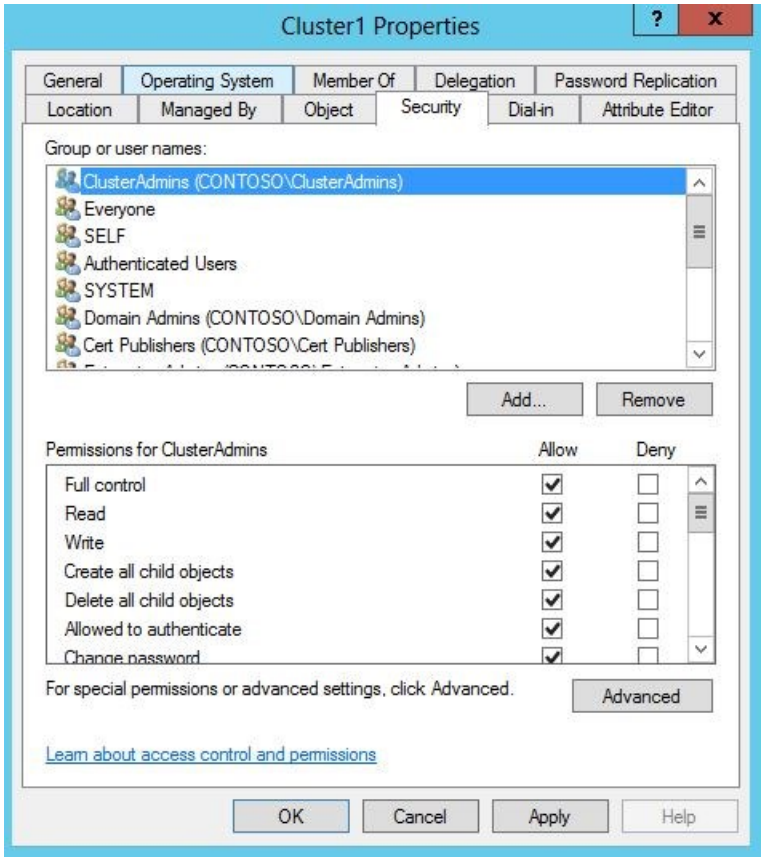

Figure 2. Granting Full Control to the user or group that will create the cluster

6. Click **OK**.

After you complete this step, the user who you granted permissions to can create the failover cluster. However, if the CNO is located in an OU, the user cannot create clustered roles that require a client access point until you complete Step 3.

Note

If the CNO is in the default Computers container, a cluster administrator can create up to 10 VCOs without any additional configuration. To add more than 10 VCOs, you must explicitly grant the **Create Computer objects** permission to the CNO for the Computers container.

Step 3: Grant the CNO permissions to the OU or prestage VCOs for clustered roles

When you create a clustered role with a client access point, the cluster creates a VCO in the same OU as the CNO. For this to occur automatically, the CNO must have permissions to create computer objects in the OU.

If you prestaged the CNO in AD DS, you can do either of the following to create VCOs:

- Option 1: Grant the CNO permissions to the OU. If you use this option, the cluster can automatically create VCOs in AD DS. Therefore, an administrator for the failover cluster can create clustered roles without having to request that you prestage VCOs in AD DS.

  Note
  Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete the steps for this option.

- Option 2: Prestage VCOs for clustered roles. Use this option if it is necessary to prestage accounts for clustered roles because of requirements in your organization. For example, you may want to control the naming convention, or control which clustered roles are created.

  Note
  Membership in the **Account Operators** group is the minimum required to complete the steps for this option.

To grant the CNO permissions to the OU

1. In Active Directory Users and Computers, on the **View** menu, make sure that **Advanced Features** is selected.
2. Right-click the OU where you created the CNO in Step 1: Prestage the CNO in AD DS, and then click **Properties**.
3. On the **Security** tab, click **Advanced**.
4. In the **Advanced Security Settings** dialog box, click **Add**.
5. Next to **Principal**, click **Select a principal**.
6. In the **Select User, Computer, Service Account, or Groups** dialog box, click **Object Types**, select the **Computers** check box, and then click **OK**.
7. Under **Enter the object names to select**, enter the name of the CNO, click **Check Names**, and then click **OK**. In response to the warning message that says that you are about to add a disabled object, click **OK**.
8. In the **Permission Entry** dialog box, make sure that the **Type** list is set to **Allow**, and the **Applies to** list is set to **This object and all descendant objects**.
9. Under **Permissions**, select the **Create Computer objects** check box.

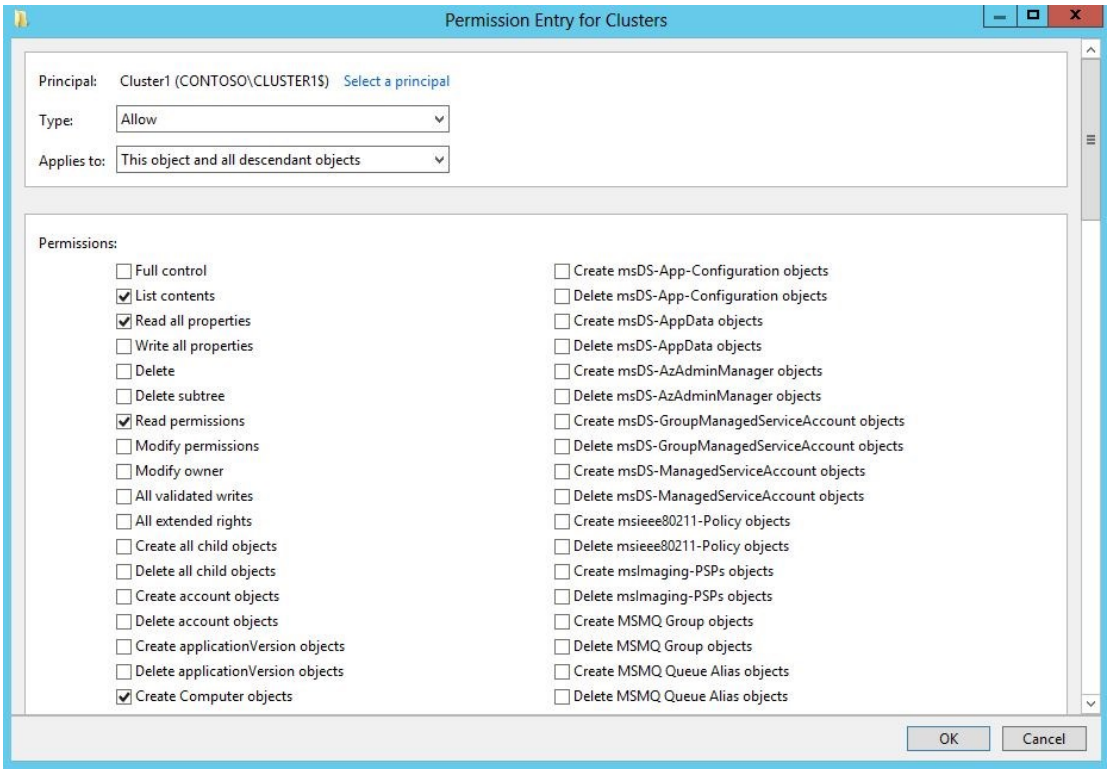

Figure 3. Granting the Create Computer objects permission to the CNO

10. Click **OK** until you return to the Active Directory Users and Computers snap-in.

An administrator on the failover cluster can now create clustered roles with client access points, and bring the resources online.

To prestage a VCO for a clustered role

1. Before you begin, make sure that you know the name of the cluster and the name that the clustered role will have.
2. In Active Directory Users and Computers, on the **View** menu, make sure that **Advanced Features** is selected.
3. In Active Directory Users and Computers, right-click the OU where the CNO for the cluster resides, point to **New**, and then click **Computer**.
4. In the **Computer name** box, enter the name that you will use for the clustered role, and then click **OK**.
5. As a best practice, right-click the computer account that you just created, click **Properties**, and then click the **Object** tab. On the **Object** tab, select the **Protect object from accidental deletion** check box, and then click **OK**.
6. Right-click the computer account that you just created, and then click **Properties**.

7. On the **Security** tab, click **Add**.
8. In the **Select User, Computer, Service Account, or Groups** dialog box, click **Object Types**, select the **Computers** check box, and then click **OK**.
9. Under **Enter the object names to select**, enter the name of the CNO, click **Check Names**, and then click **OK**. If you receive a warning message that says that you are about to add a disabled object, click **OK**.
10. Make sure that the CNO is selected, and then next to **Full control**, select the **Allow** check box.
11. Click **OK**.

An administrator on the failover cluster can now create the clustered role with a client access point that matches the prestaged VCO name, and bring the resource online.

See also

- Failover Clustering Overview

Deploy an Active Directory-Detached Cluster

In Windows Server 2012 R2, you can deploy a failover cluster without dependencies in Active Directory Domain Services (AD DS) for network names. This is referred to as an *Active Directory-detached cluster*. Using this deployment method enables you to create a failover cluster without the previously required permissions for creating computer objects in AD DS or the need to request that computer objects are prestaged in AD DS.

When you create an Active Directory-detached cluster, the cluster network name (also known as the *administrative access point*) and network names for any clustered roles with client access points are registered in Domain Name System (DNS). However, no computer objects are created for the cluster in AD DS. This includes the computer object for the cluster (also known as the cluster name object or CNO) and computer objects for any clustered roles that would typically have client access points in AD DS (also known as virtual computer objects or VCOs).

Note
This deployment method still requires that the failover cluster nodes are joined to an Active Directory domain.

Deployment considerations

An Active Directory-detached cluster uses Kerberos authentication for intracluster communication. However, when authentication against the cluster network name is required, the cluster uses NTLM authentication. Therefore, we do not recommend this deployment method for any scenario that requires Kerberos authentication.

The following table summarizes whether this deployment method is supported for a specific cluster workload.

Cluster Workload	Supported/Not Supported	More Information
SQL Server	Supported	We recommend that you use SQL Server Authentication for an Active Directory-detached cluster deployment.
File server	Supported, but not recommended	Kerberos authentication is the preferred authentication protocol for Server Message Block (SMB) traffic.
Hyper-V	Supported, but not recommended	Live migration is not supported because it has a dependency on Kerberos authentication. Quick migration is supported.
Message Queuing (also known as MSMQ)	Not supported	Message Queuing stores properties in AD DS.

In addition, be aware of the following issues for this type of cluster deployment:

- BitLocker Drive Encryption is not supported.
- Cluster-Aware Updating (CAU) in self-updating mode is not supported.

  Note
  CAU is supported in remote-updating mode.
- You cannot copy a clustered role between failover clusters that use different types of administrative access points.
- You can set the type of administrative access point only when you create the cluster. You cannot change it after the cluster is deployed.
- If you deploy a highly available file server by using this deployment method, you cannot use Server Manager to manage the file server. Instead, you must use Windows PowerShell or Failover Cluster Manager.

  To use Failover Cluster Manager, after you deploy the highly available file server, you must add the fully qualified domain name (FQDN) of the File Server clustered role to the trusted hosts list on each node of the cluster. For example, start Windows PowerShell as an administrator, and then enter the following command, where *FileServerRole1.contoso.com* and *FileServerRole2.contoso.com* represent the names of two File Server clustered roles:

  winrm set winrm/config/client '@{TrustedHosts="FileServerRole1.contoso.com,FileServerRole2.contoso.com"}'

Note
You must run this command on each node of the cluster.

How to deploy an Active Directory-detached cluster

Before you create the failover cluster, make sure that all servers that you want to add as cluster nodes meet the following prerequisites:

- All servers must be running Windows Server 2012 R2.
- All servers must be joined to the same Active Directory domain.
- All servers must have the Failover Clustering feature installed.
- All servers must use supported hardware and the collection of servers must pass all cluster validation tests. For more information, see Failover Clustering Hardware Requirements and Storage Options and Validate Hardware for a Failover Cluster.

To deploy an Active Directory-detached cluster, you must use Windows PowerShell. You cannot use Failover Cluster Manager. To create the failover cluster, start Windows PowerShell as an administrator, and then use the **New-Cluster** cmdlet with the **–AdministrativeAccessPoint** parameter set to a value of **Dns**.

The following example creates a failover cluster (*Cluster1*) from two nodes (*Node1* and *Node2*), with an administrative access point of type DNS.

```
New-Cluster Cluster1 –Node Node1,Node2 –StaticAddress 192.168.1.16 –NoStorage –
AdministrativeAccessPoint Dns
```

In this case, the cluster network name *Cluster1* will be created without a computer object in AD DS. In addition, all subsequent network names for clustered roles will be created without computer objects in AD DS.

You can run the following Windows PowerShell command to verify the type of administrative access point for a failover cluster:

```
(Get-Cluster).AdministrativeAccessPoint
```

For an Active Directory-detached cluster, the expected output value is **Dns**.

Deploy a Hyper-V Cluster

This topic provides steps to configure and deploy a simple two-node failover cluster and a clustered Hyper-V virtual machine. By using the Hyper-V role, you can use a virtualized computing environment to improve the efficiency of your computing resources and improve server availability without using as many physical computers as you would need in a failover cluster configuration that uses only physical computers. The general considerations and steps in this guide can be used to deploy clusters containing a larger number of nodes and a larger number of

virtual machines. However, this topic does not provide guidance for deploying or managing specific workloads on clustered virtual machines.

For overviews of the Hyper-V role and the Failover Clustering feature, see Hyper-V Overview and Failover Clustering Overview.

For an example of a deployment and application monitoring scenario that uses clustered virtual machines, see High Availability Printing Overview.

In this topic

- Prerequisites
- Limitations
- Step 1: Connect both physical computers to the networks and storage
- Step 2: Install Hyper-V and Failover Clustering on both physical computers
- Step 3: Create a virtual switch
- Step 4: Validate the cluster configuration
- Step 5: Create the cluster
- Step 6: Add a disk as CSV to store virtual machine data
- Step 7: Create a highly available virtual machine
- Step 8: Install the guest operating system on the virtual machine
- Step 9: Test a planned failover
- Step 10: Test an unplanned failover
- Step 11: Modify the settings of a virtual machine
- Step 12: Remove a virtual machine from a cluster

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Starting Windows PowerShell.

Prerequisites

To use the Hyper-V role on a failover cluster with two nodes, you need the hardware, software, accounts, and network infrastructure that are outlined in the sections that follow.

Hardware requirements

For general requirements for the servers, networks, network adapters, storage, and device controllers for storage for Hyper-V and for failover clusters, see "Hardware requirements" in Hyper-V Overview and Failover Clustering Hardware Requirements and Storage Options.

Note

You can use shared storage that is attached, and you can also use SMB 3.0 file shares as shared storage for servers that are running Hyper-V that are configured in a failover

cluster. For more information, see Deploy Hyper-V over SMB. The steps in this topic assume that you will be using shared storage that is attached to the cluster nodes.

Important
Microsoft supports a failover cluster solution only if all the hardware features are marked as "Certified for Windows Server 2012 R2" or "Certified for Windows Server 2012." In addition, the complete configuration (servers, network, and storage) must pass all cluster validation tests. For more information, see Step 4: Validate the cluster configuration in this topic and Validate Hardware for a Failover Cluster.

Software requirements

For the software requirements for using the Hyper-V role and the Failover Clustering feature, see Hyper-V Overview and Failover Clustering Overview.

To install the guest operating system that will run on the clustered virtual machine, you will need appropriate installation media for a supported operating system. You can install from a physical media or an image (.iso) file. You also can configure the virtual machine by using a virtual hard disk on which an operating system has already been installed. The steps in this topic assume that you will install Windows Server 2012 R2 on the virtual machine.

Network infrastructure and domain account requirements

You will need the following network infrastructure for a two-node failover cluster:

- **Networks** It is recommended to have a minimum of two networks for your failover cluster: a public network that allows clients to connect to the cluster and a separate network that is used only for communication between the clustered servers. As needed, you can configure additional networks for specific storage options or for redundancy.
- **Network settings and IP addresses** When you use identical network adapters for a network, also use identical communication settings on those adapters (for example, Speed, Duplex Mode, Flow Control, and Media Type). Also, compare the settings between the network adapter and the switch it connects to and make sure that no settings are in conflict.

  If you have private networks that are not routed to the rest of your network infrastructure, ensure that each of these private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network, and another node in a branch office that uses a separate physical network, do not specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address.
- **DNS** The servers in the cluster must be using Domain Name System (DNS) for name resolution. The DNS dynamic update protocol can be used.
- **Domain role** All servers in the cluster must be in the same Active Directory domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server. If the clustered servers are member servers, you need an additional server that acts as the domain controller in the domain that contains your cluster.

- **Clients** As needed, you can connect one or more linked clients to the failover cluster that you create, and then observe the effect on a client when you move or fail over the highly available virtual machine from one cluster node to the other.

To create a cluster or add nodes, you must be logged on to the domain with an account that has administrator rights and permissions on all servers in that cluster. The account does not need to be a Domain Admins account, but can be a Domain Users account that is in the Administrators group on each clustered server. In addition, if the account is not a Domain Admins account, the account (or the group that the account is a member of) must be given the Create Computer Objects and Read All Properties permissions in the domain.

Limitations

Hyper-V and failover clusters include the following general limitations:

- The failover cluster can have a maximum of 64 nodes.
- You can have a maximum of 8,000 virtual machines per cluster for server computer virtualization, with a maximum of 1,024 virtual machines on a single node, provided that the server hardware has the resources to support them. For example, when Hyper-V is used together with Virtual Desktop Infrastructure (VDI) for client computer virtualization, you can have a maximum of 8,000 VDI (Windows 8.1, Windows 8 or Windows 7) virtual machines per cluster, with a maximum of 1,024 on a single node.

For more information about scalability limits for Hyper-V, see Hyper-V Scalability in Windows Server 2012.

Step 1: Connect both physical computers to the networks and storage

Use the following instructions to connect your selected servers to networks and storage.

To connect the servers to the networks and storage

1. For details about the kinds of network adapters and device controllers that you can use with a failover cluster, review Prerequisites, earlier in this topic.
2. Connect and configure the networks that the servers in the cluster will use.

   Note
   If you want to include clients or a non-clustered domain controller as part of your test configuration, make sure that these computers can connect to the failover cluster servers through at least one network.
3. Follow the manufacturer's instructions for physically connecting the servers to the storage.
4. Ensure that the disks (LUNs) that you want to use in the cluster are exposed to the servers that you will cluster (and only those servers). You can use either of the following interfaces to expose disks or LUNs:

- The interface provided by the manufacturer of the storage.
- An appropriate iSCSI interface.

5. If you have purchased software that controls the format or function of the disk, follow the instructions from the vendor about how to use that software with Windows Server 2012 R2 or Windows Server 2012.
6. On one of the servers that you want to cluster, in Server Manager or in Disk Management, confirm that the cluster disks are visible.
7. If you want to have a storage volume larger than 2 terabytes, and you are using the Windows interface to control the format of the disk, convert that disk to the partition style called GUID partition table (GPT). To do this, back up any data on the disk, and delete all volumes on the disk. Then, in Disk Management, right-click the disk (not a partition), and click **Convert to GPT Disk**.

   For volumes smaller than 2 terabytes, instead of using GPT, you can use the partition style called master boot record (MBR).

   Important
   You can use either the MBR or the GPT partition style for a disk that is used by a failover cluster, but you cannot use a disk that you have converted to dynamic by using Disk Management.
8. Check the format of any exposed volume or LUN. We recommend NTFS for the format (for the quorum witness disk, you can use NTFS or ReFS).

Step 2: Install Hyper-V and Failover Clustering on both physical computers

To install the Hyper-V role and the Failover Clustering feature on each physical computer, see the following procedures:

- Add the Hyper-V role in "Install the Hyper-V Role and Configure a Virtual Machine."
- Installing the Failover Cluster Feature and Tools in Windows Server 2012

Step 3: Create a virtual switch

Perform this step on both physical computers if you did not create the virtual switch when you installed the Hyper-V role. This virtual switch provides the highly available virtual machine with access to the physical network.

Do this step using Windows PowerShell

To create a virtual switch

1. Open Hyper-V Manager.
2. From the **Actions** menu, click **Virtual Switch Manager**.
3. Under **Create virtual switch**, select **External**.

4. Click **Create Virtual Switch**. The **New Virtual Switch** page appears.
5. Type a name for the new switch. Make sure you use exactly the same name on both servers running Hyper-V.
6. Under **Connection Type**, click **External network**, and then select the physical network adapter.
7. Click **OK** to save the virtual network and close Virtual Switch Manager.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example creates the external switch *VMExternalSwitch*, which binds to the network adapter *Wired Ethernet Connection 3* and allows the management operating system to share the network adapter.

```
New-VMSwitch "VMExternalSwitch" -NetAdapterName "Wired Ethernet Connection 3" -
AllowManagementOS
```

Step 4: Validate the cluster configuration

Before you create the cluster, we strongly recommend that you run a full validation test of your cluster configuration by running the Validate a Configure Wizard in Failover Cluster Manager, or the Windows PowerShell Test-Cluster cmdlet. Specific validation tests are included for the configuration of the Hyper-V role in the failover cluster.

For detailed considerations and steps to validate the cluster, see Validate Hardware for a Failover Cluster.

Step 5: Create the cluster

To create a failover cluster by using Failover Cluster Manager or the Windows PowerShell New-Cluster cmdlet, see Creating a Windows Server 2012 Failover Cluster.

Step 6: Add a disk as CSV to store virtual machine data

To implement certain scenarios for clustered virtual machines, the virtual machine storage and virtual hard disk file should be configured as Cluster Shared Volumes (CSV). To configure a disk in clustered storage as a CSV volume, you can use Failover Cluster Manager or the Windows PowerShell Add-ClusterSharedVolume cmdlet. For detailed planning considerations and steps to create CSV, see Use Cluster Shared Volumes in a Windows Server 2012 Failover Cluster.

CSV can enhance the availability and manageability of virtual machines by enabling multiple nodes to concurrently access a single shared storage volume. For example, on a failover cluster that uses CSV, multiple clustered virtual machines that are distributed across multiple cluster

nodes can all access their virtual hard disk files at the same time, even if the files are on a single disk (LUN) in the storage. This means that the clustered virtual machines can fail over independently of one another, even if they use only a single LUN. CSV also support live migration of a Hyper-V virtual machine between nodes in a failover cluster.

Step 7: Create a highly available virtual machine

In this step, you create a virtual machine and configure it for high availability.

Note
You can run the Hyper-V New Virtual Machine Wizard directly from Failover Cluster Manager. After the virtual machine is created in this way, it is automatically configured for high availability.

Considerations for creating a virtual machine that will be highly available

- Starting in Windows Server 2012, we do not support the configuration where there is more than one virtual machine in a virtual machine clustered role. An example of this is a scenario where multiple virtual machines have files on a common physical disk that is not part of CSV. A single virtual machine per clustered role improves the management experience and the functionality of virtual machines in a clustered environment, such as virtual machine mobility.
- Choose the shared storage as the location to store the virtual machine and the virtual hard disk. Otherwise, you will not be able to make the virtual machine highly available. To make the shared storage available to the virtual machine, you must create the virtual machine on the physical computer that is the node which owns the storage.
- If you created a CSV volume in Step 6: Add a disk as CSV to store virtual machine data, in the settings for the virtual hard disk, specify the CSV volume as the location of both the virtual machine and the virtual hard disk.
- Ensure that you select a virtual hard disk option that is appropriate for the method you are using to install the guest operating system on the virtual machine (for example, from physical media or from an .iso file).

Do this step using Windows PowerShell

To create a highly available virtual machine

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. Click **Roles**.
3. In the **Actions** pane, click **Virtual Machines**, and then click **New Virtual Machine**. The New Virtual Machine Wizard appears. Click **Next**.
4. On the **Specify Name and Location** page, specify a name for the virtual machine, such as *FailoverTest*. Click **Store the virtual machine in a different location**, and then type the full path or click **Browse** and navigate to the shared storage.

5. On the **Assign Memory** page, specify the amount of memory required for the operating system that will run on this virtual machine. For example, specify 1024 MB to run Windows Server 2012 R2.
6. On the **Configure Networking** page, connect the network adapter to the virtual switch that is associated with the physical network adapter. You should specify the virtual switch that you configured in Step 3: Create a virtual switch.
7. On the **Connect Virtual Hard Disk** page, click **Create a virtual hard disk**. If you want to change the name, type a new a name for the virtual hard disk. Click **Next**.
8. On the **Installation Options** page, click **Install an operating system from a boot CD/DVD-ROM**. Under **Media**, specify the location of the media, and then click **Finish**.

   The virtual machine is created. The High Availability Wizard in Failover Cluster Manager then automatically configures the virtual machine for high availability.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example creates the virtual machine *FailoverTest*, specifying that it will be installed from a .iso file, and configures it for high availability.

```
New-VHD -Path <PathToVHDXFile> -Dynamic -SizeBytes 127GB

New-VM -Name FailoverTest -Path <PathToVMFolder> -Memory 1GB -SwitchName
"VMExternalSwitch" -BootDevice CD -VHDPath <PathToVHDXFile>

Add-VMDvdDrive -VMName FailoverTest -Path <PathtoISOFile>

Set-VM -Name FailoverTest -AutomaticStartAction Nothing

Add-ClusterVirtualMachineRole -VirtualMachine FailoverTest
```

Step 8: Install the guest operating system on the virtual machine

You then start the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine . This installs the guest operating system (which, in this topic, is assumed to be Windows Server 2012 R2). To do this, see Step 3: Install the guest operating system in “Install the Hyper-V Role and Configure a Virtual Machine.”

Note

If you are installing an operating system other than Windows Server 2012 R2 on a Windows Server 2012 R2 Hyper-V host, or an operating system other than Windows Server 2012 on a Windows Server 2012 Hyper-V host, you might also need to install Hyper-V integration services for the operating system.

Step 9: Test a planned failover

To test a planned failover, you can move the clustered virtual machine that you created in Step 7: Create a highly available virtual machine to another node.

You have the following options to move a clustered virtual machine:

- **Live migration** Move ownership of the clustered virtual machine to another node without pausing the role.
- **Quick migration** Pause the virtual machine, save state, move the role to another node, and start the virtual machine on the other node.
- **Storage migration** Move only the virtual machine data to other clustered storage.

For example, to test a planned failover by performing a live migration, you can use Failover Cluster Manager or the Windows PowerShell Move-ClusterVirtualMachineRole cmdlet.

Do this step using Windows PowerShell

To test a planned failover

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. To select the destination node for live migration of the clustered virtual machine, right-click *FailoverTest* (the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine), point to **Move**, point to **Live Migration**, and then click **Select Node**.

   As the *FailoverTest* virtual machine is moved, the status is displayed in the results pane (center pane). Optionally, you can repeat this step to move the virtual machine to an additional node or back to the original node.
3. Verify that the move succeeded by inspecting the details of each node.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example live migrates the virtual machine *FailoverTest* to the node *ContosoFCNode2*.

```
Move-ClusterVirtualMachineRole -Name "FailoverTest" -Node ContosoFCNode2
```

Step 10: Test an unplanned failover

To test an unplanned failover of the clustered virtual machine, you can stop the Cluster service on the node that owns the clustered virtual machine.

Do this step using Windows PowerShell

To test an unplanned failover

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. To minimize disruption to clients, before stopping the Cluster service on a node, move the clustered roles that are currently owned by that node (other than *FailoverTest*) to another node by doing the following:
   a. Expand the console tree under the cluster that you want to manage, and then expand **Nodes**.
   b. Select the node that owns the clustered virtual machine *FailoverTest* (the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine).
   c. Select all of the clustered roles on the node, except *FailoverTest*.
   d. To select the destination node for the selected clustered roles, right-click the roles, point to **Move**, and then click **Select Node**

   It might take several minutes to move the clustered roles to the destination node.
3. Expand the console tree under **Nodes**.
4. Right-click the node that owns *FailoverTest*, point to **More Actions**, and then click **Stop Cluster Service**.

   The virtual machine moves to the other node. There might be a short delay while this happens.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example stops the Cluster service on *ContosoFCNode2*, the node that owns the clustered virtual machine *FailoverTest*.

```
Stop-ClusterNode –Name ContosoFCNode2
```

Step 11: Modify the settings of a virtual machine

To modify the configuration of a virtual machine, you can use the Failover Clustering tools or the Hyper-V tools to access the settings. We recommend that you use the Failover Clustering tools to access the virtual machine settings. When you do this, the cluster is updated automatically after you apply the new settings, and you are able to modify all of the virtual machine settings. However, if you make changes to the virtual machine settings by using the Hyper-V tools directly, you must update the cluster manually after you make the changes, and you will be prevented from modifying certain virtual machine settings or performing actions that could conflict with settings for the failover cluster. In addition, if the configuration is not refreshed after changes are

made, a subsequent failover may not succeed, or it may succeed but the virtual machine will then be configured incorrectly.

Note
To modify certain virtual machine settings, you may be prompted to first shut down the virtual machine.

To modify the settings of a virtual machine

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. If you need to shut down the virtual machine before modifying the settings, expand **Roles**, right-click *FailoverTest*(the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine), and then click **Shut Down**.
3. Right-click *FailoverTest*, and then click **Settings**. The Settings page for the virtual machine appears.
4. Configure the settings for the virtual machine, and then click **OK**.

   The virtual machine configuration is updated in the failover cluster.
5. If you previously shut down the clustered virtual machine, right-click *FailoverTest*, point to **More Actions**, and then click **Start Role**.

Use the following procedure to manually refresh the virtual machine configuration in the failover cluster

Do this step using Windows PowerShell

To manually refresh the virtual machine configuration in the failover cluster

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. Right-click *FailoverTest* (the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine), point to **More Actions**, and then click **Refresh Virtual Machine Configuration**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example refreshes the configuration of the clustered virtual machine *FailoverTest*.

```
Update-ClusterVirtualMachineConfiguration –Name "FailoverTest"
```

Step 12: Remove a virtual machine from a cluster

When you want to remove a virtual machine from a cluster, the procedure you need to use varies depending on whether you want to keep the virtual machine and its files. This step illustrates both scenarios.

Scenario A: To remove a virtual machine from a cluster and retain the virtual machine

1. Use Failover Cluster Manager to take the virtual machine offline. Under **Roles**, right-click *FailoverTest* (the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine), point to **More Actions**, and then click **Stop Role**.
2. Optionally, export the virtual machine. Exporting a virtual machine allows you to move the virtual machine to another server running Hyper-V, such as a non-clustered server. To do this:
   a. Switch to Hyper-V Manager and verify that the *FailoverTest* virtual machine is selected.
   b. Under **Actions**, click **Export**.
   c. Type or browse to specify a location in which to export the virtual machine, and then click **Export**.

   Important
   If you plan to import the virtual machine to another cluster, use either Hyper-V Manager or Microsoft System Center Virtual Machine Manager. If you import a virtual machine using Hyper-V Manager, afterwards, configure the virtual machine by using the High Availability Wizard in Failover Cluster Manager, or the Windows PowerShell Add-ClusterVirtualMachineRole cmdlet.
3. In Hyper-V Manager, verify that the *FailoverTest* virtual machine is selected. Under **Actions**, click **Delete**.
4. In Failover Cluster Manager. expand **Roles**, right-click *FailoverTest*, and then click **Remove**.

   The virtual machine is removed from the cluster.

Scenario B: To remove a virtual machine from a cluster and delete the virtual machine

1. Use Failover Cluster Manager to take the virtual machine offline. Under **Roles**, right-click *FailoverTest* (the clustered virtual machine that you configured in Step 7: Create a highly available virtual machine), point to **More Actions**, and then click **Stop Role**.
2. Switch to Hyper-V Manager and select the *FailoverTest* virtual machine. Under **Actions**, click **Delete**.
3. In Failover Cluster Manager, expand **Roles**, right-click *FailoverTest*, and then click **Remove**.

   The virtual machine is removed from the cluster.
4. Manually delete the virtual machine and the virtual hard disk from the shared storage.

See also

- Hyper-V Overview
- Failover Clustering Overview
- System Center 2012 – Virtual Machine Manager

Network Recommendations for a Hyper-V Cluster in Windows Server 2012

There are several different types of network traffic that you must consider and plan for when you deploy a highly available Hyper-V solution. You should design your network configuration with the following goals in mind:

- To ensure network quality of service
- To provide network redundancy
- To isolate traffic to defined networks
- Where applicable, take advantage of Server Message Block (SMB) Multichannel

This topic provides network configuration recommendations that are specific to a Hyper-V cluster that is running Windows Server 2012. It includes an overview of the different network traffic types, recommendations for how to isolate traffic, recommendations for features such as NIC Teaming, Quality of Service (QoS) and Virtual Machine Queue (VMQ), and a Windows PowerShell script that shows an example of converged networking, where the network traffic on a Hyper-V cluster is routed through one external virtual switch.

Windows Server 2012 supports the concept of converged networking, where different types of network traffic share the same Ethernet network infrastructure. In previous versions of Windows Server, the typical recommendation for a failover cluster was to dedicate separate physical network adapters to different traffic types. Improvements in Windows Server 2012, such as Hyper-V QoS and the ability to add virtual network adapters to the management operating system enable you to consolidate the network traffic on fewer physical adapters. Combined with traffic isolation methods such as VLANs, you can isolate and control the network traffic.

Important
If you use System Center Virtual Machine Manager (VMM) to create or manage Hyper-V clusters, you must use VMM to configure the network settings that are described in this topic.

In this topic:

- Overview of different network traffic types
- How to isolate the network traffic on a Hyper-V cluster
- NIC Teaming (LBFO) recommendations
- Quality of Service (QoS) recommendations
- Virtual machine queue (VMQ) recommendations

- Example of converged networking: routing traffic through one Hyper-V virtual switch
- Appendix: Encryption

Overview of different network traffic types

When you deploy a Hyper-V cluster, you must plan for several types of network traffic. The following table summarizes the different traffic types.

Network Traffic Type	Description
Management	• Provides connectivity between the server that is running Hyper-V and basic infrastructure functionality. • Used to manage the Hyper-V management operating system and virtual machines.
Cluster	• Used for inter-node cluster communication such as the cluster heartbeat and Cluster Shared Volumes (CSV) redirection.
Live migration	• Used for virtual machine live migration.
Storage	• Used for SMB traffic or for iSCSI traffic.
Replica traffic	• Used for virtual machine replication through the Hyper-V Replica feature.
Virtual machine access	• Used for virtual machine connectivity. • Typically requires external network connectivity to service client requests.

The following sections provide more detailed information about each network traffic type.

Management traffic

A management network provides connectivity between the operating system of the physical Hyper-V host (also known as the *management operating system*) and basic infrastructure functionality such as Active Directory Domain Services (AD DS), Domain Name System (DNS), and Windows Server Update Services (WSUS). It is also used for management of the server that is running Hyper-V and the virtual machines.

The management network must have connectivity between all required infrastructure, and to any location from which you want to manage the server.

Cluster traffic

A failover cluster monitors and communicates the cluster state between all members of the cluster. This communication is very important to maintain cluster health. If a cluster node does not

communicate a regular health check (known as the *cluster heartbeat*), the cluster considers the node down and removes the node from cluster membership. The cluster then transfers the workload to another cluster node.

Inter-node cluster communication also includes traffic that is associated with CSV. For CSV, where all nodes of a cluster can access shared block-level storage simultaneously, the nodes in the cluster must communicate to orchestrate storage-related activities. Also, if a cluster node loses its direct connection to the underlying CSV storage, CSV has resiliency features which redirect the storage I/O over the network to another cluster node that can access the storage.

Live migration traffic

Live migration enables the transparent movement of running virtual machines from one Hyper-V host to another without a dropped network connection or perceived downtime.

We recommend that you use a dedicated network or VLAN for live migration traffic to ensure quality of service and for traffic isolation and security. Live migration traffic can saturate network links. This can cause other traffic to experience increased latency. The time it takes to fully migrate one or more virtual machines depends on the throughput of the live migration network. Therefore, you must ensure that you configure the appropriate quality of service for this traffic. To provide the best performance, live migration traffic is not encrypted.

You can designate multiple networks as live migration networks in a prioritized list. For example, you may have one migration network for cluster nodes in the same cluster that is fast (10 GB), and a second migration network for cross-cluster migrations that is slower (1 GB).

All Hyper-V hosts that can initiate or receive a live migration must have connectivity to a network that is configured to allow live migrations. Because live migration can occur between nodes in the same cluster, between nodes in different clusters, and between a cluster and a stand-alone Hyper-V host, make sure that all these servers can access a live migration-enabled network.

Storage traffic

For a virtual machine to be highly available, all members of the Hyper-V cluster must be able to access the virtual machine state. This includes the configuration state and the virtual hard disks. To meet this requirement, you must have shared storage.

In Windows Server 2012, there are two ways that you can provide shared storage:

- **Shared block storage**. Shared block storage options include Fibre Channel, Fibre Channel over Ethernet (FCoE), iSCSI, and shared Serial Attached SCSI (SAS).
- **File-based storage**. Remote file-based storage is provided through SMB 3.0.

SMB 3.0 includes new functionality known as SMB Multichannel. SMB Multichannel automatically detects and uses multiple network interfaces to deliver high performance and highly reliable storage connectivity.

By default, SMB Multichannel is enabled, and requires no additional configuration. You should use at least two network adapters of the same type and speed so that SMB Multichannel is in

effect. Network adapters that support RDMA (Remote Direct Memory Access) are recommended but not required.

SMB 3.0 also automatically discovers and takes advantage of available hardware offloads, such as RDMA. A feature known as SMB Direct supports the use of network adapters that have RDMA capability. SMB Direct provides the best performance possible while also reducing file server and client overhead.

Note
The NIC Teaming feature is incompatible with RDMA-capable network adapters. Therefore, if you intend to use the RDMA capabilities of the network adapter, do not team those adapters.

Both iSCSI and SMB use the network to connect the storage to cluster members. Because reliable storage connectivity and performance is very important for Hyper-V virtual machines, we recommend that you use multiple networks (physical or logical) to ensure that these requirements are achieved.

Note
For more information about SMB Direct and SMB Multichannel, see Improve Performance of a File Server with SMB Direct and The basics of SMB Multichannel, a feature of Windows Server 2012 and SMB 3.0.

Replica traffic

Hyper-V Replica provides asynchronous replication of Hyper-V virtual machines between two hosting servers or Hyper-V clusters. Replica traffic occurs between the primary and Replica sites.

Hyper-V Replica automatically discovers and uses available network interfaces to transmit replication traffic. To throttle and control the replica traffic bandwidth, you can define QoS policies with minimum bandwidth weight.

If you use certificate-based authentication, Hyper-V Replica encrypts the traffic. If you use Kerberos-based authentication, traffic is not encrypted.

Virtual machine access traffic

Most virtual machines require some form of network or Internet connectivity. For example, workloads that are running on virtual machines typically require external network connectivity to service client requests. This can include tenant access in a hosted cloud implementation. Because multiple subclasses of traffic may exist, such as traffic that is internal to the datacenter and traffic that is external (for example to a computer outside the datacenter or to the Internet); one or more networks are required for these virtual machines to communicate.

To separate virtual machine traffic from the management operating system, we recommend that you use VLANs which are not exposed to the management operating system.

How to isolate the network traffic on a Hyper-V cluster

To provide the most consistent performance and functionality, and to improve network security, we recommend that you isolate the different types of network traffic.

Note
Realize that if you want to have a physical or logical network that is dedicated to a specific traffic type, you must assign each physical or virtual network adapter to a unique subnet. For each cluster node, Failover Clustering recognizes only one IP address per subnet.

Isolate traffic on the management network

We recommend that you use a firewall or IPsec encryption, or both, to isolate management traffic. In addition, you can use auditing to ensure that only defined and allowed communication is transmitted through the management network.

Isolate traffic on the cluster network

To isolate inter-node cluster traffic, you can configure a network to either allow cluster network communication or not to allow cluster network communication. For a network that allows cluster network communication, you can also configure whether to allow clients to connect through the network. (This includes client and management operating system access.)

A failover cluster can use any network that allows cluster network communication for cluster monitoring, state communication, and for CSV-related communication.

To configure a network to allow or not to allow cluster network communication, you can use Failover Cluster Manager or Windows PowerShell. To use Failover Cluster Manager, click **Networks** in the navigation tree. In the **Networks** pane, right-click a network, and then click **Properties**.

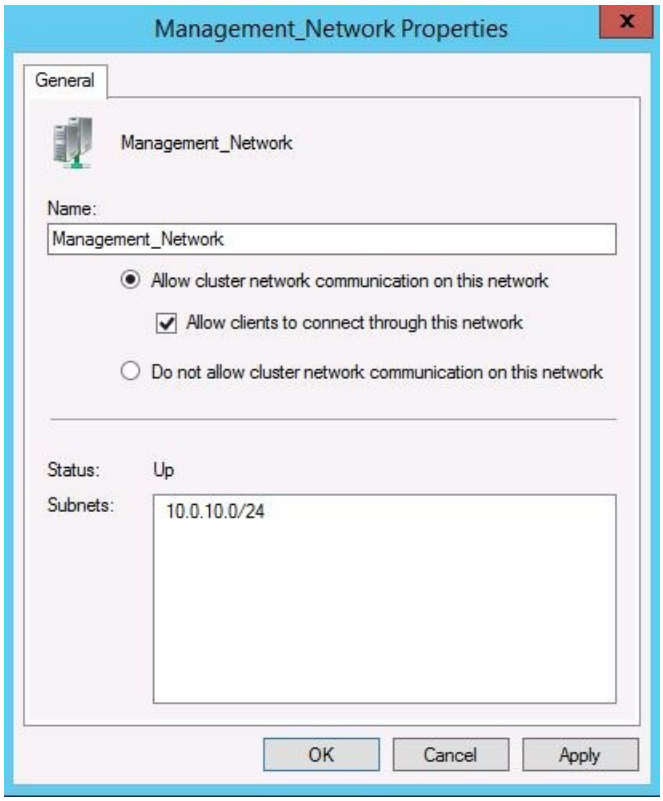

Figure 1. Failover Cluster Manager network properties

The following Windows PowerShell example configures a network named *Management Network* to allow cluster and client connectivity.

```
(Get-ClusterNetwork -Name "Management Network").Role = 3
```

The **Role** property has the following possible values.

Value	Network Setting
0	Do not allow cluster network communication
1	Allow cluster network communication only
3	Allow cluster network communication and client connectivity

The following table shows the recommended settings for each type of network traffic. Realize that virtual machine access traffic is not listed because these networks should be isolated from the management operating system by using VLANs that are not exposed to the host. Therefore, virtual machine networks should not appear in Failover Cluster Manager as cluster networks.

Network Type	Recommended Setting
Management	Both of the following: • Allow cluster network communication on this network • Allow clients to connect through this network
Cluster	Allow cluster network communication on this network **Note** Clear the **Allow clients to connect through this network** check box.
Live migration	Allow cluster network communication on this network **Note** Clear the **Allow clients to connect through this network** check box.
Storage	Do not allow cluster network communication on this network
Replica traffic	Both of the following: • Allow cluster network communication on this network • Allow clients to connect through this network

Isolate traffic on the live migration network

By default, live migration traffic uses the cluster network topology to discover available networks and to establish priority. However, you can manually configure live migration preferences to isolate live migration traffic to only the networks that you define. To do this, you can use Failover Cluster Manager or Windows PowerShell. To use Failover Cluster Manager, in the navigation tree, right-click **Networks**, and then click **Live Migration Settings**.

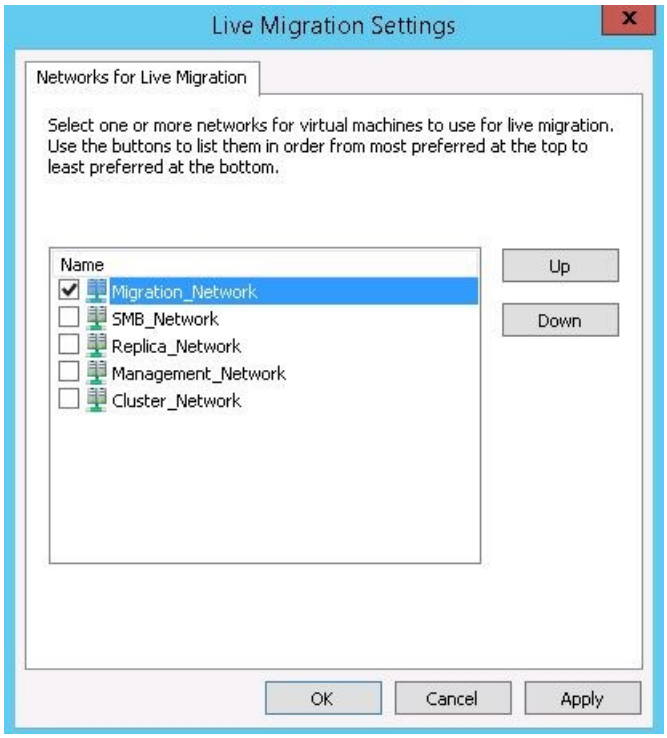

Figure 2. Live migration settings in Failover Cluster Manager

The following Windows PowerShell example enables live migration traffic only on a network that is named *Migration_Network*.

```
Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name MigrationExcludeNetworks -Value ([String]::Join(";",(Get-ClusterNetwork | Where-Object {$_.Name -ne "Migration_Network"}).ID))
```

Isolate traffic on the storage network

To isolate SMB storage traffic, you can use Windows PowerShell to set SMB Multichannel constraints. SMB Multichannel constraints restrict SMB communication between a given file server and the Hyper-V host to one or more defined network interfaces.

For example, the following Windows PowerShell command sets a constraint for SMB traffic from the file server *FileServer1* to the network interfaces *SMB1*, *SMB2*, *SMB3*, and *SMB4* on the Hyper-V host from which you run this command.

```
New-SmbMultichannelConstraint -ServerName "FileServer1" -InterfaceAlias "SMB1", "SMB2", "SMB3", "SMB4"
```

Notes

- You must run this command on each node of the Hyper-V cluster.

- To find the interface name, use the **Get-NetAdapter** cmdlet.

For more information, see New-SmbMultichannelConstraint.

To isolate iSCSI traffic, configure the iSCSI target with interfaces on a dedicated network (logical or physical). Use the corresponding interfaces on the cluster nodes when you configure the iSCSI initiator.

Isolate traffic for replication

To isolate Hyper-V Replica traffic, we recommend that you use a different subnet for the primary and Replica sites.

If you want to isolate the replica traffic to a particular network adapter, you can define a persistent static route which redirects the network traffic to the defined network adapter. To specify a static route, use the following command:

route add <destination> mask <subnet mask and gateway> if <interface> -p

For example, to add a static route to the 10.1.17.0 network (example network of the Replica site) that uses a subnet mask of 255.255.255.0, a gateway of 10.0.17.1 (example IP address of the primary site), where the interface number for the adapter that you want to dedicate to replica traffic is 8, run the following command:

route add 10.1.17.1 mask 255.255.255.0 10.0.17.1 if 8 -p

NIC Teaming (LBFO) recommendations

We recommend that you team physical network adapters in the management operating system. This provides bandwidth aggregation and network traffic failover if a network hardware failure or outage occurs.

The NIC Teaming feature, also known as load balancing and failover (LBFO), provides two basic sets of algorithms for teaming.

- **Switch-dependent modes**. Requires the switch to participate in the teaming process. Typically requires all the network adapters in the team to be connected to the same switch.
- **Switch-independent modes**. Does not require the switch to participate in the teaming process. Although not required, team network adapters can be connected to different switches.

Both modes provide for bandwidth aggregation and traffic failover if a network adapter failure or network disconnection occurs. However, in most cases only switch-independent teaming provides traffic failover for a switch failure.

NIC Teaming also provides a traffic distribution algorithm that is optimized for Hyper-V workloads. This algorithm is referred to as the *Hyper-V port* load balancing mode. This mode distributes the traffic based on the MAC address of the virtual network adapters. The algorithm uses round robin as the load-balancing mechanism. For example, on a server that has two teamed physical network adapters and four virtual network adapters, the first and third virtual network adapter will use the first physical adapter, and the second and fourth virtual network adapter will use the

second physical adapter. Hyper-V port mode also enables the use of hardware offloads such as virtual machine queue (VMQ) which reduces CPU overhead for networking operations.

Recommendations

For a clustered Hyper-V deployment, we recommend that you use the following settings when you configure the additional properties of a team.

Property Name	Recommended Setting
Teaming mode	Switch Independent (the default setting)
Load balancing mode	Hyper-V Port

Note

NIC teaming will effectively disable the RDMA capability of the network adapters. If you want to use SMB Direct and the RDMA capability of the network adapters, you should not use NIC Teaming.

For more information about the NIC Teaming modes and how to configure NIC Teaming settings, see Windows Server 2012 NIC Teaming (LBFO) Deployment and Management and NIC Teaming Overview.

Quality of Service (QoS) recommendations

You can use QoS technologies that are available in Windows Server 2012 to meet the service requirements of a workload or an application. QoS provides the following:

- Measures network bandwidth, detects changing network conditions (such as congestion or availability of bandwidth), and prioritizes - or throttles - network traffic.
- Enables you to converge multiple types of network traffic on a single adapter.
- Includes a minimum bandwidth feature which guarantees a certain amount of bandwidth to a given type of traffic.

We recommend that you configure appropriate Hyper-V QoS on the virtual switch to ensure that network requirements are met for all appropriate types of network traffic on the Hyper-V cluster.

Note

You can use QoS to control outbound traffic, but not the inbound traffic. For example, with Hyper-V Replica, you can use QoS to control outbound traffic (from the primary server), but not the inbound traffic (from the Replica server).

Recommendations

For a Hyper-V cluster, we recommend that you configure Hyper-V QoS that applies to the virtual switch. When you configure QoS, do the following:

- **Configure minimum bandwidth in weight mode instead of in bits per second.** Minimum bandwidth specified by weight is more flexible and it is compatible with other features, such

as live migration and NIC Teaming. For more information, see the **MinimumBandwidthMode** parameter in New-VMSwitch.

- **Enable and configure QoS for all virtual network adapters.** Assign a weight to all virtual adapters. For more information, see Set-VMNetworkAdapter. To make sure that all virtual adapters have a weight, configure the **DefaultFlowMinimumBandwidthWeight** parameter on the virtual switch to a reasonable value. For more information, see Set-VMSwitch.

The following table recommends some generic weight values. You can assign a value from 1 to 100. For guidelines to consider when you assign weight values, see Guidelines for using Minimum Bandwidth.

Network Classification	Weight
Default weight	0
Virtual machine access	1, 3 or 5 (low, medium and high-throughput virtual machines)
Cluster	10
Management	10
Replica traffic	10
Live migration	40
Storage	40

Virtual machine queue (VMQ) recommendations

Virtual machine queue (VMQ) is a feature that is available to computers that have VMQ-capable network hardware. VMQ uses hardware packet filtering to deliver packet data from an external virtual network directly to virtual network adapters. This reduces the overhead of routing packets. When VMQ is enabled, a dedicated queue is established on the physical network adapter for each virtual network adapter that has requested a queue.

Not all physical network adapters support VMQ. Those that do support VMQ will have a fixed number of queues available, and the number will vary. To determine whether a network adapter supports VMQ, and how many queues they support, use the Get-NetAdapterVmq cmdlet.

You can assign virtual machine queues to any virtual network adapter. This includes virtual network adapters that are exposed to the management operating system. Queues are assigned according to a weight value, in a first-come first-serve manner. By default, all virtual adapters have a weight of 100.

Recommendations

We recommend that you increase the VMQ weight for interfaces with heavy inbound traffic, such as storage and live migration networks. To do this, use the Set-VMNetworkAdapter Windows PowerShell cmdlet.

Example of converged networking: routing traffic through one Hyper-V virtual switch

The following Windows PowerShell script shows an example of how you can route traffic on a Hyper-V cluster through one Hyper-V external virtual switch. The example uses two physical 10 GB network adapters that are teamed by using the NIC Teaming feature. The script configures a Hyper-V cluster node with a management interface, a live migration interface, a cluster interface, and four SMB interfaces. After the script, there is more information about how to add an interface for Hyper-V Replica traffic. The following diagram shows the example network configuration.

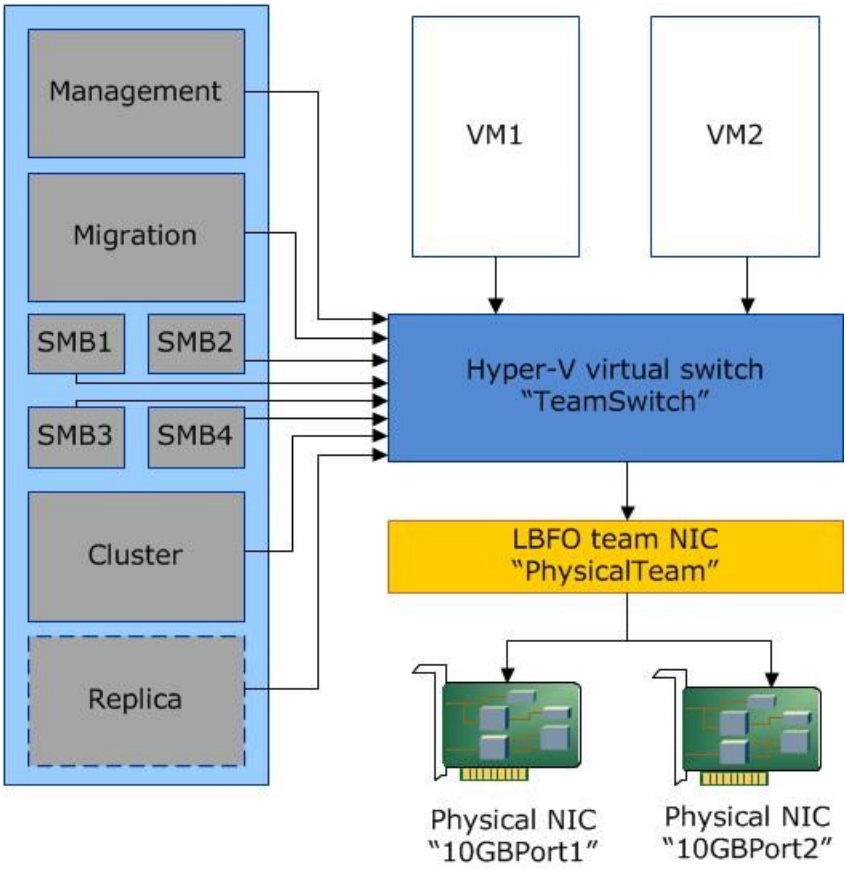

Figure 3. Example Hyper-V cluster network configuration

The example also configures network isolation which restricts cluster traffic from the management interface, restricts SMB traffic to the SMB interfaces, and restricts live migration traffic to the live migration interface.

```
# Create a network team using switch independent teaming and Hyper-V port mode

New-NetLbfoTeam "PhysicalTeam" -TeamMembers "10GBPort1", "10GBPort2" -TeamNicName
"PhysicalTeam" -TeamingMode SwitchIndependent -LoadBalancingAlgorithm HyperVPort
```

```
# Create a Hyper-V virtual switch connected to the network team
# Enable QoS in Weight mode
New-VMSwitch "TeamSwitch" –NetAdapterName "PhysicalTeam" –MinimumBandwidthMode Weight –
AllowManagementOS $false

# Configure the default bandwidth weight for the switch
# Ensures all virtual NICs have a weight
Set-VMSwitch -Name "TeamSwitch" -DefaultFlowMinimumBandwidthWeight 0

# Create virtual network adapters on the management operating system
# Connect the adapters to the virtual switch
# Set the VLAN associated with the adapter
# Configure the VMQ weight and minimum bandwidth weight
Add-VMNetworkAdapter –ManagementOS –Name "Management" –SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "Management" -Access -VlanId
10
Set-VMNetworkAdapter -ManagementOS -Name "Management" -VmqWeight 80 -
MinimumBandwidthWeight 10

Add-VMNetworkAdapter –ManagementOS –Name "Cluster" –SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "Cluster" -Access -VlanId 11
Set-VMNetworkAdapter -ManagementOS -Name "Cluster" -VmqWeight 80 -MinimumBandwidthWeight
10

Add-VMNetworkAdapter –ManagementOS –Name "Migration" –SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "Migration" -Access -VlanId
12
Set-VMNetworkAdapter -ManagementOS -Name "Migration" -VmqWeight 90 -
MinimumBandwidthWeight 40

Add-VMNetworkAdapter –ManagementOS –Name "SMB1" –SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "SMB1" -Access -VlanId 13
Set-VMNetworkAdapter -ManagementOS -Name "SMB1" -VmqWeight 100 -MinimumBandwidthWeight 40

Add-VMNetworkAdapter –ManagementOS –Name "SMB2" –SwitchName "TeamSwitch"
```

```
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "SMB2" -Access -VlanId 14
Set-VMNetworkAdapter -ManagementOS -Name "SMB2" -VmqWeight 100 -MinimumBandwidthWeight 40

Add-VMNetworkAdapter -ManagementOS -Name "SMB3" -SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "SMB3" -Access -VlanId 15
Set-VMNetworkAdapter -ManagementOS -Name "SMB3" -VmqWeight 100 -MinimumBandwidthWeight 40

Add-VMNetworkAdapter -ManagementOS -Name "SMB4" -SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "SMB4" -Access -VlanId 16
Set-VMNetworkAdapter -ManagementOS -Name "SMB4" -VmqWeight 100 -MinimumBandwidthWeight 40

# Rename the cluster networks if desired
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.10.0"}).Name =
"Management_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.11.0"}).Name = "Cluster_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.12.0"}).Name =
"Migration_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.13.0"}).Name = "SMB_Network1"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.14.0"}).Name = "SMB_Network2"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.15.0"}).Name = "SMB_Network3"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.16.0"}).Name = "SMB_Network4"

# Configure the cluster network roles
(Get-ClusterNetwork -Name "Management_Network").Role = 3
(Get-ClusterNetwork -Name "Cluster_Network").Role = 1
(Get-ClusterNetwork -Name "Migration_Network").Role = 1
(Get-ClusterNetwork -Name "SMB_Network1").Role = 0
(Get-ClusterNetwork -Name "SMB_Network2").Role = 0
(Get-ClusterNetwork -Name "SMB_Network3").Role = 0
(Get-ClusterNetwork -Name "SMB_Network4").Role = 0

# Configure an SMB Multichannel constraint
# This ensures that SMB traffic from the named server only uses SMB interfaces
```

```
New-SmbMultichannelConstraint -ServerName "FileServer1" -InterfaceAlias "vEthernet
(SMB1)", "vEthernet (SMB2)", "vEthernet (SMB3)", "vEthernet (SMB4)"


# Configure the live migration network
Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name
MigrationExcludeNetworks -Value ([String]::Join(";",(Get-ClusterNetwork | Where-Object
{$_.Name -ne "Migration_Network"}).ID))
```

Hyper-V Replica considerations

If you also use Hyper-V Replica in your environment, you can add another virtual network adapter to the management operating system for replica traffic. For example:

```
Add-VMNetworkAdapter –ManagementOS –Name "Replica" –SwitchName "TeamSwitch"
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "Replica" –Access –VlanId 17
Set-VMNetworkAdapter -ManagementOS -Name "Replica" -VmqWeight 80 -MinimumBandwidthWeight
10


# If the host is clustered – configure the cluster name and role
(Get-ClusterNetwork | Where-Object {$_.Address -eq "10.0.17.0"}).Name = "Replica"
(Get-ClusterNetwork -Name "Replica").Role = 3
```

Notes

If you are instead using policy-based QoS, where you can throttle outgoing traffic regardless of the interface on which it is sent, you can use either of the following methods to throttle Hyper-V Replica traffic:

- Create a QoS policy that is based on the destination subnet. In this example, the Replica site uses the 10.1.17.0/24 subnet.

  ```
  New-NetQosPolicy "Replica traffic to 10.1.17.0" –
  DestinationAddress 10.1.17.0/24 –MinBandwidthWeightAction 40
  ```

- Create a QoS policy that is based on the destination port. In the following example, the network listener on the Replica server or cluster has been configured to use port 8080 to receive replication traffic.

  ```
  New-NetQosPolicy "Replica traffic to 8080" –DestinationPort 8080
  –ThrottleRateActionBitsPerSecond 100000
  ```

Appendix: Encryption

Cluster traffic

By default, cluster communication is not encrypted. You can enable encryption if you want. However, realize that there is performance overhead that is associated with encryption. To enable encryption, you can use the following Windows PowerShell command to set the security level for the cluster.

```
(Get-Cluster). SecurityLevel = 2
```

The following table shows the different security level values.

Security Description	Value
Clear text	0
Signed (default)	1
Encrypted	2

Live migration traffic

Live migration traffic is not encrypted. You can enable IPsec or other network layer encryption technologies if you want. However, realize that encryption technologies typically affect performance.

SMB traffic

By default, SMB traffic is not encrypted. Therefore, we recommend that you use a dedicated network (physical or logical) or use encryption. For SMB traffic, you can use SMB encryption, layer-2 or layer-3 encryption. SMB encryption is the preferred method.

Replica traffic

If you use Kerberos-based authentication, Hyper-V Replica traffic is not encrypted. We strongly recommend that you encrypt replication traffic that transits public networks over the WAN or the Internet. We recommend Secure Sockets Layer (SSL) encryption as the encryption method. You can also use IPsec. However, realize that using IPsec may significantly affect performance.

See also

- Failover Cluster Networking Essentials (video)

Using Guest Clustering for High Availability

A failover cluster that is made up of two or more virtual machines is typically referred to as a *guest cluster*. Guest clustering supports high availability for workloads within the virtual machine. This topic includes planning and deployment information that is specific to a guest cluster scenario.

Besides the traditional benefits of virtualization, deploying a clustered application (known as a *clustered role*) inside a virtual machine provides the following additional value:

- **Proactive health monitoring of the clustered roles**. For each clustered role, there are several associated resources such as a resource for the application, the storage, and networking. The cluster continuously monitors each resource to verify that it is functioning. If the cluster detects a failure, it initiates automatic recovery steps. For example, the cluster may move the clustered role or fail it over to another node.
- **Application mobility**. You can easily move clustered roles to other nodes. In a guest cluster, you can move a clustered role to another node to service the virtual machine while the application stays in production. For example, you can apply updates or change the configuration, restart the virtual machine, and verify that it functions correctly without interrupting the availability of the application.

  Note
  Starting in Windows Server 2012, you can use the Cluster-Aware Updating feature to automate the updating of guest cluster nodes. For more information, see Cluster-Aware Updating Overview.
- **Protection from host failure**. You can place guest cluster nodes on different physical hosts to protect against physical host failure. If a guest cluster node becomes unavailable because of physical host failure, another guest cluster node automatically detects the situation. It then brings any clustered roles that were running on the unavailable node online. This preserves the availability of the clustered roles. Note that this functionality is especially valuable if the virtual machines are not running on hosts that are part of a Hyper-V failover cluster.
- **Virtual machine mobility**. If you run your clustered applications on virtual machines, you can choose to deploy them to any host that supports the virtual machine. To move a virtual machine to another host, you can live migrate the virtual machine. This provides flexibility and the ability to optimize resources on the Hyper-V hosts.

Physical host considerations

To increase the availability of a guest cluster workload, you must also consider the physical hosts on which the virtual machines will run.

Deploy a guest cluster on a physical host cluster

Although not required, to provide extra resiliency we recommend that you deploy virtual machines that are part of a guest cluster on a Hyper-V host cluster instead of on stand-alone computers that are running Hyper-V.

Consider the following scenario:

1. A two node guest cluster (Guest1 and Guest2) runs on a three node physical host cluster (Host1, Host2, and Host3). Guest1 runs on Host1. Guest2 runs on Host2.
2. Host1 suddenly fails. The guest cluster fails over any roles that were running on Guest1 to Guest2. This quickly brings the role back online and makes the workload available to clients.
3. Because Guest1 is on a physical host cluster, the Guest1 virtual machine automatically moves and starts on another host. If Guest2 has problems, Guest1 is available again for the roles to fail over to it.

In this scenario, if Guest1 was instead running on a stand-alone host, the virtual machine would be unavailable until Host1 was brought back online and the virtual machine started. If there was an issue such as a hardware failure, the virtual machine could be unavailable for a lengthy period.

Place the virtual machines on different physical hosts

We recommend that you place virtual machines that are part of the same guest cluster on different physical hosts. This enables the failover of the workload to a running guest cluster node if there is a host failure. The surviving guest cluster node detects the loss of the other node and starts the clustered roles. This shortens the recovery time and increases the availability of the workload.

If all nodes of the guest cluster are on the same host and a host failure occurs, after the host recovers you must restart at least one of the virtual machines to make the workload available. To avoid this situation, you can configure the cluster to place virtual machines on different physical nodes of the cluster. To do this, use either of the following methods:

- Use Windows PowerShell to set the **AntiAffinityClassNames** property for each virtual machine role that is part of the same guest cluster to the same text string. **AntiAffinityClassNames** is a property of a cluster group. For more information, see AntiAffinityClassNames, and the example in this section.
- If you use System Center Virtual Machine Manager (VMM), you can configure anti-affinity by using availability sets. For more information, see How to Configure Availability Sets in VMM for Virtual Machines on a Host Cluster.

> **Note**
> This functionality is available starting in System Center 2012 Service Pack 1.

When the cluster must fail over clustered roles or move a virtual machine to another node, it checks all clustered roles on the destination node to determine whether any of the roles have the same string value set for the **AntiAffinityClassNames** property. If a role on the destination node has the same string value, the cluster targets another node and does the same check until it finds a node that can host the virtual machine. If there is no available node except for one with a role that has the same string value, the cluster places the virtual machine on a node with a matching string value.

Example of how to set anti-affinity

The following Windows PowerShell example shows how to view the **AntiAffinityClassNames** property value for a clustered virtual machine that is named "*VM1*", and how to set the property

value to the string "*GuestCluster1*." You must set this property to the same string on each virtual machine that is a member of the same guest cluster.

To view the current **AntiAffinityClassNames** property value for a guest cluster node that is named "*VM1*," type the following command:

```
Get-ClusterGroup VM1 | fl anti*
```

To set the **AntiAffinityClassNames** property value to the string "*GuestCluster1*," type the following command:

```
(Get-ClusterGroup VM1).AntiAffinityClassNames = "GuestCluster1"
```

To confirm the property value for *VM1*, type the following command again:

```
Get-ClusterGroup VM1 | fl anti*
```

Guest cluster considerations for deployment and management

The deployment and management of a guest cluster is very similar to a physical host cluster. You would follow the same steps to create a guest cluster. Before you install an application on the guest cluster, check if there is a deployment guide that outlines specific requirements and recommendations for running the application in a virtual machine.

If you use System Center 2012 R2 Virtual Machine Manager, realize that it includes a new feature to deploy guest clusters. You can use service templates to create the virtual machines and to create a Windows Server 2012 R2 guest cluster that uses shared virtual hard disk (.vhdx) files as storage. For more information, see How to Create a Guest Cluster by Using a Service Template in VMM.

Support requirements

The requirements for Microsoft Customer Service and Support to officially support a guest cluster are the same as for clusters that run directly on physical hardware. This includes the requirement to pass all cluster validation tests. For more information, see Validate Hardware for a Windows Server 2012 Failover Cluster and The Microsoft support policy for Windows Server 2012 failover clusters.

To receive support for the Windows Server operating system or for server features that are running on a virtual machine, you must use a supported hypervisor. (Note that Microsoft's Life Cycle Support policy still applies to the guest operating system.) If you do not use Hyper-V, Microsoft requires that the hypervisor is validated and listed in the Microsoft Windows Server Catalog under the Server Virtualization Validation Program (SVVP). For more information, see Server Virtualization Validation Program and the Server Virtualization Validation Program FAQ.

Node limit

The same 64-node limit applies to physical host clusters and to guest clusters that are running Windows Server 2012 R2 or Windows Server 2012. The number of nodes that you deploy

depends on the clustered role. For example, if you have a database role on physical servers that is deployed on more than two nodes to spread out the instances, you can emulate this environment by using more than two nodes of a guest cluster to host the instances.

Realize that although deploying virtual machines has significant cost savings over dedicated physical servers, each virtual machine consumes CPU, memory, and other system resources that could be otherwise provisioned to other virtual machines. Therefore, do not add more nodes to a guest cluster than needed.

Active Directory domain requirements

Failover clusters require that nodes of the same cluster are members of the same Active Directory domain. However, there is no requirement that a guest cluster belongs to the same Active Directory domain as the physical hosts on which the virtual machines are running.

Storage options

The following tables lists the storage types that you can use to provide shared storage for a guest cluster.

Storage Type	Description
Shared virtual hard disk	New in Windows Server 2012 R2, you can configure multiple virtual machines to connect to and use a single virtual hard disk (.vhdx) file. Each virtual machine can access the virtual hard disk just like servers would connect to the same LUN in a storage area network (SAN). For more information, see Deploy a Guest Cluster Using a Shared Virtual Hard Disk.
Virtual Fibre Channel	Introduced in Windows Server 2012, virtual Fibre Channel enables you to connect virtual machines to LUNs on a Fibre Channel SAN. For more information, see Hyper-V Virtual Fibre Channel Overview.
iSCSI	The iSCSI initiator inside a virtual machine enables you to connect over the network to an iSCSI target. For more information, see iSCSI Target Block Storage Overview and the blog post Introduction of iSCSI Target in Windows Server 2012.

Storage requirements depend on the clustered roles that run on the cluster. Most clustered roles use clustered storage, where the storage is available on any cluster node that runs a clustered

role. Examples of clustered storage include Physical Disk resources and Cluster Shared Volumes (CSV). Some roles do not require storage that is managed by the cluster. For example, you can configure Microsoft SQL Server to use availability groups that replicate the data between nodes. Other clustered roles may use Server Message Block (SMB) shares or Network File System (NFS) shares as data stores that any cluster node can access.

Network considerations

Failover clusters, whether physical or virtual, have networking requirements for communication between the nodes and to provide client network connectivity.

Network redundancy

To achieve the goal of high availability, avoid single points of failure whenever possible. For networking, you can use NIC Teaming to protect against network adapter failures. Also, connect network adapters to different switches to provide tolerance for switch and port failures.

For guest clusters, consider the following scenarios when you plan for network redundancy:

- **Physical host NIC Teaming**. In this scenario, the host has multiple network adapters in a teamed configuration. You would configure Hyper-V with a single virtual switch that uses the team.
- **NIC Teaming inside the virtual machines**. In this scenario, the physical host provides more than one virtual network that has external connectivity. (Make sure that the virtual network adapters are connected to different external switches.) You then use NIC Teaming inside the virtual machine to provide failover between the networks if there is a connection failure.

  Important
  To use NIC Teaming inside a virtual machine, you must also configure each Hyper-V switch port that is associated with the virtual machine to allow teaming. To do this, open the virtual machine settings in Hyper-V Manager. In the **Advanced Features** of each network adapter, under **NIC Teaming**, select the **Enable this network adapter to be part of a team in the guest operating system** check box. Or, you can use the following Windows PowerShell command:

  ```
  Set-VMNetworkAdapter –VMName VMName –AllowTeaming On
  ```

Because most failures occur at the infrastructure level, it is best to provide resiliency for the physical host. You can use NIC Teaming inside the guest virtual machines when you do not have NIC Teaming at the physical host level. For more information about NIC Teaming, see NIC Teaming Overview.

Number of networks

When you deploy a guest cluster, you must consider the following general categories of network communication. Each has slightly different requirements.

Traffic Type	Description
Client access	Most cluster deployments have clustered roles that provide services to users or to other

Traffic Type	Description
	systems. For instance, most database applications have client applications that store and retrieve records or other information. This type of communication typically requires routing to other subnets and a default gateway.
Cluster communication	The Cluster service and other cluster components such as Cluster Shared Volumes (CSV) use one or more networks to transfer information between nodes. For this communication, the cluster can use private networks that have no gateway or routed networks.
Storage	Clustered roles may require storage that can be accessed from any node of the cluster. If you use SMB file shares or iSCSI targets for the guest cluster, both use Ethernet networks to connect to the storage.

Realize that there are different ways to provide these requirements. These include the following:

- **Separate networks for each category**. Some guest cluster deployments may have one or more virtual networks for each type of network traffic. This emulates the common physical network configuration for failover clusters. There is a network for client access (preferably teamed) that has a default gateway and is routable, a network that is on a non-routed subnet for communications between cluster nodes, and if storage uses a network, one or more network adapters that use either NIC Teaming or Multipath I/O (MPIO) for resiliency.
- **Converged networking**. In this scenario you use one virtual network for all networking I/O on the guest cluster virtual machines. Because the network is shared for all types of network traffic, you should configure QoS policies on the physical host to make sure that the cluster and CSV traffic has reserved bandwidth to meet I/O demands. For more information, see the "Hyper-V QoS" section of Policy-based QoS and Hyper-V QoS, and the "4 NICs in two NIC teams" example in QoS Common Configurations.

  Note
  When you run cluster validation tests on a cluster with a single network adapter, the networking test reports a warning about lack of redundancy. This warning does not prevent you from deploying the cluster or from receiving support. It is provided to make sure that you understand the importance of redundancy.

Protect against short-term network interruptions

Failover cluster nodes use the network to send heartbeat packets to other nodes of the cluster. If a node does not receive a response from another node for a specified period of time, the cluster removes the node from cluster membership. By default, a guest cluster node is considered down

if it does not respond within 5 seconds. Other nodes that are members of the cluster will take over any clustered roles that were running on the removed node.

Typically, during the live migration of a virtual machine there is a fast final transition when the virtual machine is stopped on the source node and is running on the destination node. However, if something causes the final transition to take longer than the configured heartbeat threshold settings, the guest cluster considers the node to be down even though the live migration eventually succeeds. If the live migration final transition is completed within the TCP time-out interval (typically around 20 seconds), clients that are connected through the network to the virtual machine seamlessly reconnect.

To make the cluster heartbeat time-out more consistent with the TCP time-out interval, you can change the **SameSubnetThreshold** and **CrossSubnetThreshold** cluster properties from the default of 5 seconds to 20 seconds. By default, the cluster sends a heartbeat every 1 second. The threshold specifies how many heartbeats to miss in succession before the cluster considers the cluster node to be down.

Important

Do not increase the heartbeat threshold from its default value for a Scale-Out File Server. A Scale-Out File Server already provides seamless recovery. If a node is considered down, the node fails over and recovers the SMB sessions on another node. This recovery action needs to occur within the SMB session time-out period. If you increase the heartbeat threshold, node detection will take longer. If it exceeds the SMB session time-out values, the recovery may no longer be seamless.

The following example shows how to use Windows PowerShell to view the **CrossSubnetThreshold** and **SameSubnetThreshold** cluster properties, and how to change the values for each to 20.

To view the current property values, type the following command:

```
Get-Cluster | fl *subnet*
```

To set the values of the **CrossSubnetThreshold** and **SameSubnetThreshold** properties to 20, type the following commands:

```
(Get-Cluster).CrossSubnetThreshold = 20
```

```
(Get-Cluster).SameSubnetThreshold = 20
```

To confirm the property values, type the following command again:

```
Get-Cluster | fl *subnet*
```

See also

- High Availability Solutions (SQL Server)
- Failover Clustering Overview

Deploy a Guest Cluster Using a Shared Virtual Hard Disk

The procedures in this topic show how you can designate a virtual hard disk (.vhdx file) as shared storage for multiple virtual machines that are configured as a guest failover cluster. You can use shared .vhdx files for data disks and for the disk witness. Do not use a shared .vhdx file for the operating system disk.

For an overview of shared .vhdx files that includes deployment diagrams, see Virtual Hard Disk Sharing Overview.

Deployment scenarios

You can deploy a guest failover cluster that uses shared .vhdx files on a Hyper-V failover cluster in either of the following configurations:

Scenario 1: Hyper-V failover cluster using Cluster Shared Volumes (CSV) on block-level storage

In this scenario, all virtual machine files, including the shared .vhdx files are stored on a CSV that is configured as shared storage for a Hyper-V failover cluster.

This scenario has the following physical host prerequisites:

- At least two servers that are running Windows Server 2012 R2.
- The servers must be members of the same Active Directory domain.
- The servers must support Hyper-V. For more information, see the "Hardware requirements" section of Hyper-V Overview.
- The servers must meet the requirements for failover clustering. For more information, see Failover Clustering Hardware Requirements and Storage Options and Validate Hardware for a Failover Cluster.
- The servers must have access to block-level storage, which you can add as shared storage to the physical cluster. This storage can be iSCSI, Fibre Channel, Serial Attached SCSI (SAS), or clustered storage spaces that use a set of shared SAS just-a-bunch-of-disks (JBOD) enclosures.

Scenario 2: Hyper-V failover cluster using file-based storage in a separate Scale-Out File Server

This scenario uses Server Message Block (SMB) file-based storage as the location of the shared .vhdx files. You must deploy a Scale-Out File Server and create an SMB file share as the storage location. You also need a separate Hyper-V failover cluster.

The following table describes the physical host prerequisites.

Cluster Type	Requirements
Scale-Out File Server	• At least two servers that are running

Cluster Type	Requirements
	Windows Server 2012 R2. • The servers must be members of the same Active Directory domain. • The servers must meet the requirements for failover clustering. For more information, see Failover Clustering Hardware Requirements and Storage Options and Validate Hardware for a Failover Cluster. • The servers must have access to block-level storage, which you can add as shared storage to the physical cluster. This storage can be iSCSI, Fibre Channel, SAS, or clustered storage spaces that use a set of shared SAS JBOD enclosures.
Hyper-V failover cluster	• At least two servers that are running Windows Server 2012 R2. • The servers must be members of the same Active Directory domain. • The servers must support Hyper-V. For more information, see the “Hardware requirements” section of Hyper-V Overview. • The servers must meet the requirements for failover clustering. For more information, see Failover Clustering Hardware Requirements and Storage Options and Validate Hardware for a Failover Cluster.

Step 1: Configure the physical servers

In this step, you create the failover clusters that are required to host the virtual machines, and create the storage locations for the virtual machine files.

Use the appropriate procedure for your deployment scenario. Before you begin, make sure that the servers meet the prerequisites that are outlined in the Deployment scenarios section of this topic.

Scenario 1: Hyper-V failover cluster using CSV

1. Create at least a two-node Hyper-V failover cluster by following Steps 1 through 5 in the topic Deploy a Hyper-V Cluster.

2. In Step 6 of the Deploy a Hyper-V Cluster topic, add the storage where you want to store the virtual machine files to CSV. For more information, see the "Add a disk as a CSV to a failover cluster" section of Use Cluster Shared Volumes in a Failover Cluster.

Scenario 2: Hyper-V cluster using file-based storage in a separate Scale-Out File Server

1. Create at least a two-node Hyper-V failover cluster by following Steps 1 through 5 in the topic Deploy a Hyper-V Cluster.

   Note
   If you plan to store all virtual machine files on the Scale-Out File Server, you do not need shared block-level storage on the Hyper-V cluster.

2. Deploy a Scale-Out File Server by following Steps 1 and 2, and Task 1.1 of Step 3 in the guide Deploy Scale-Out File Server.

Step 2: Create the virtual machines

This step shows how to create the virtual machines that you can use to create the guest failover cluster. In this step, you create the virtual machine configuration files and the virtual hard disk that is used for the operating system.

Guest operating system requirements

You must use Windows Server 2012 or Windows Server 2012 R2 for the guest operating system. If you use Windows Server 2012, you must upgrade the Hyper-V integration services to the Windows Server 2012 R2 version.

Note
As a best practice, we recommend that if you create a virtual machine on a failover cluster node, you create it as a highly available virtual machine.

To create the virtual machines

1. From one of the Hyper-V failover cluster nodes, create two or more highly available virtual machines. To do this, follow Steps 7 and 8 in the topic Deploy a Hyper-V Cluster. When you create the virtual machines, note the following:
   a. On the **Specify Name and Location** page, select the **Store the virtual machine in a different location** check box, click **Browse**, and then navigate to the shared storage location.
      - If the shared storage is a CSV disk, navigate to: *C:*\ClusterStorage\Volume*X*, where *C:*\ represents the system drive, and *X* represents the volume number.
      - If the shared storage is an SMB file share, navigate to: *ServerName**ShareName*, where *ServerName* represents the client access point for the Scale-Out File Server, and *ShareName* represents the name of the SMB file share.
   b. On the **Specify Generation** page, select **Generation 1** or **Generation 2**. Both types of virtual machines are supported. However, if you want to use an existing virtual hard disk that is in the .vhd file format, you must use a Generation 1 virtual machine.

Note
The virtual hard disk for the operating system can use the .vhd or the .vhdx file format.

c. On the **Connect Virtual Hard Disk** page, create a virtual hard disk or use an existing virtual hard disk that resides on shared storage. (If you use an existing virtual hard disk, the guest operating system must be Windows Server 2012 or Windows Server 2012 R2.) In the **Location** box, specify a CSV path or the SMB file share path.

2. If you chose to create a virtual hard disk, install an operating system on each virtual machine. You must use Windows Server 2012 or Windows Server 2012 R2.
3. Ensure that both virtual machines are running and that they are joined to the same Active Directory domain.
4. If you installed Windows Server 2012 as the guest operating system, upgrade the Hyper-V integration services to the Windows Server 2012 R2 version. For information about how to upgrade integration services, see Step 4 of the topic Install the Hyper-V Role and Configure a Virtual Machine.

Step 3: Create and enable a shared virtual hard disk

This step shows how to create and then share a virtual hard disk that is in the .vhdx file format. Repeat this step for each shared .vhdx file that you want to add. For example, you may want to add one or more shared disks that will act as data disks, and a separate shared disk that you can designate as the disk witness for the guest failover cluster.

To create a virtual hard disk to share

1. In Failover Cluster Manager, expand the cluster name, and then click **Roles**.
2. In the **Roles** pane, right-click the virtual machine on which you want to add a shared virtual hard disk, and then click **Settings**.
3. In the virtual machine settings, under **Hardware**, click **SCSI Controller**.
4. In the details pane, click **Hard Drive**, and then click **Add**.
5. In the **Hard Drive** details pane, under **Virtual hard disk**, click **New**.

   The New Virtual Hard Disk Wizard opens.
6. On the **Before You Begin** page, click **Next**.
7. On the **Choose Disk Format** page, accept the default format of **VHDX**, and then click **Next**.

   Note
   To share the virtual hard disk, the format must be .vhdx.
8. On the **Choose Disk Type** page, select **Fixed size** or **Dynamically expanding**, and then click **Next**.

Note

A differencing disk is not supported for a shared virtual hard disk.

9. On the **Specify Name and Location** page, do the following:
   a. In the **Name** box, enter the name of the shared virtual hard disk.
   b. In the **Location** box, enter the path of the shared storage location.

      For Scenario 1, where the shared storage is a CSV disk, enter the path:

      C:\ClusterStorage\Volume*X*, where *C:*\ represents the system drive, and *X* represents the desired CSV volume number.

      For Scenario 2, where the shared storage is an SMB file share, specify the path:

      *ServerName**ShareName*, where *ServerName* represents the client access point for the Scale-Out File Server, and *ShareName* represents the name of the SMB file share.
   c. Click **Next**.
10. On the **Configure Disk** page, accept the default option of **Create a new blank virtual hard disk**, specify the desired size, and then click **Next**.
11. On the **Completing the New Virtual Hard Disk Wizard** page, review the configuration, and then click **Finish**.

Important

If the virtual machine is running, do not click **Apply** in the virtual machine settings before you continue to the next procedure.

If you do click **Apply** on a running virtual machine, you will need to stop the virtual machine or remove and then add the virtual hard disk without clicking **Apply**. For more information, see the Note in Step 3 of the following procedure.

To share the virtual hard disk

1. In the virtual machine settings, under **SCSI Controller**, expand the hard drive that you created in the previous procedure.
2. Click **Advanced Features**.
3. In the details pane, select the **Enable virtual hard disk sharing** check box.

Notes

If the check box appears dimmed and is unavailable, you can do either of the following:

- Remove and then add the virtual hard disk to the running virtual machine. When you do, ensure that you do not click **Apply** when the New Virtual Hard Disk Wizard completes. Instead, immediately configure sharing in **Advanced Features**.
- Stop the virtual machine, and then select the **Enable virtual hard disk sharing** check box.

4. Click **Apply**, and then click **OK**.
5. Add the virtual hard disk to each virtual machine that will use the shared .vhdx file. When

you do, repeat this procedure to enable virtual hard disk sharing for each virtual machine that will use the disk.

Tip

To share a virtual hard disk by using Windows PowerShell, use the **Add-VMHardDiskDrive** cmdlet with the **–ShareVirtualDisk** parameter. You must run this command as an administrator on the Hyper-V host for each virtual machine that will use the shared .vhdx file.

For example, the following command adds a shared virtual hard disk (*Data1.vhdx*) on volume 1 of CSV to a virtual machine that is named *VM1*.

```
Add-VMHardDiskDrive -VMName VM1 -Path C:\ClusterStorage\Volume1\Data1.vhdx -
ShareVirtualDisk
```

The following command adds a shared virtual hard disk (*Witness.vhdx*) that is stored on an SMB file share (*Server1**Share1*) to a virtual machine that is named *VM2*.

```
Add-VMHardDiskDrive -VMName VM2 -Path \\Server1\Share1\Witness.vhdx -
ShareVirtualDisk
```

Step 4: Create the guest cluster

This step shows how to verify that the shared .vhdx files show up as shared storage when you create the guest cluster.

To create the guest cluster and verify the shared storage

1. On one of the virtual machines, in Server Manager or in Disk Management, confirm that the shared .vhdx file is visible as a disk. Bring the disk online, initialize the disk, and then create and format the volume.
2. Create the guest cluster by adding the virtual machines as cluster nodes. When you create the cluster, add all eligible storage. (The shared .vhdx files will be added as available storage.)
3. After the failover cluster is created, verify that the storage was added. To do this, open Failover Cluster Manager, expand the cluster name, expand **Storage**, and then click **Disks**.

   In the **Disks** pane, verify that the disks appear.
4. Optionally, add the available storage to CSV. For more information, see the "Add a disk as a CSV to a failover cluster" section of Use Cluster Shared Volumes in a Failover Cluster.

Cluster-Aware Updating Overview

This topic provides an overview of Cluster-Aware Updating (CAU), a feature for failover clusters that was introduced in Windows Server 2012. CAU automates the software updating process on clustered servers while maintaining availability. This topic describes scenarios and applications for using CAU, and provides links to content that describes how to integrate CAU into other IT automation and management processes.

Did you mean…

CAU is related to but is distinct from the following foundational technologies:

- Failover Clustering Overview
- Windows Update/Microsoft Update
- Windows Server Update Services
- Windows Management Instrumentation
- Windows Remote Management

Feature description

CAU is an automated feature that enables you to update clustered servers with little or no loss of availability during the update process. During an Updating Run, CAU transparently performs the following tasks:

- Puts each node of the cluster into node maintenance mode
- Moves the clustered roles off the node
- Installs the updates and any dependent updates
- Performs a restart if necessary
- Brings the node out of maintenance mode
- Restores the clustered roles on the node
- Moves to update the next node

For many clustered roles (formerly called clustered applications and services) in the cluster, the automatic update process triggers a planned failover, and it can cause a transient service interruption for connected clients. However, in the case of continuously available workloads such as Hyper-V with live migration or file server with SMB Transparent Failover, CAU can coordinate cluster updates with no impact to the service availability.

Note

The CAU feature is only compatible with Windows Server 2012 R2 and Windows Server 2012 failover clusters and the clustered roles that are supported on those versions.

Practical applications

- CAU reduces service outages in clustered services, reduces the need for manual updating workarounds, and makes the end-to-end cluster updating process more reliable for the

administrator. When the CAU feature is used together with continuously available cluster workloads, such as continuously available file servers (file server workload with SMB Transparent Failover) or Hyper-V, the cluster updates can be performed with zero impact to service availability for clients.

- CAU facilitates the adoption of consistent IT processes across the enterprise. You can create Updating Run Profiles for different classes of failover clusters and then manage them centrally on a file share to ensure that CAU deployments throughout the IT organization apply updates consistently, even if the clusters are managed by different lines-of-business or administrators.
- CAU can schedule Updating Runs on regular daily, weekly, or monthly intervals to help coordinate cluster updates with other IT management processes.
- CAU provides an extensible architecture to update the cluster software inventory in a cluster-aware fashion. This can be used by publishers to coordinate the installation of software updates that are not published to Windows Update or Microsoft Update or that are not available from Microsoft, for example, updates for non-Microsoft device drivers.
- CAU self-updating mode enables a “cluster in a box” appliance (a set of clustered physical machines running Windows Server 2012, typically packaged in one chassis) to update itself. Typically, such appliances are deployed in branch offices with minimal local IT support to manage the clusters. Self-updating mode offers great value in these deployment scenarios.

Important functionality

Following is a description of important CAU functionality:

- A user interface (UI) and a set of Windows PowerShell cmdlets that you can use to preview, apply, monitor, and report on the updates
- An end-to-end automation of the cluster-updating operation (an Updating Run), orchestrated by one or more Update Coordinator computers
- A default plug-in that integrates with the existing Windows Update Agent (WUA) and Windows Server Update Services (WSUS) infrastructure in Windows Server 2012 R2 or Windows Server 2012 to apply important Microsoft updates
- A second plug-in that can be used to apply Microsoft hotfixes, and that can be customized to apply non-Microsoft updates
- Updating Run Profiles that you configure with settings for Updating Run options, such as the maximum number of times that the update will be retried per node. Updating Run Profiles enable you to rapidly reuse the same settings across Updating Runs and easily share the update settings with other failover clusters.
- An extensible architecture that supports new plug-in development to coordinate other node-updating tools across the cluster, such as custom software installers, BIOS updating tools, and network adapter or host bus adapter (HBA) updating tools.

CAU can coordinate the complete cluster updating operation in two modes:

- **Self-updating mode** For this mode, the CAU clustered role is configured as a workload on the failover cluster that is to be updated, and an associated update schedule is defined. The cluster updates itself at scheduled times by using a default or custom Updating Run profile. During the Updating Run, the CAU Update Coordinator process starts on the node that

currently owns the CAU clustered role, and the process sequentially performs updates on each cluster node. To update the current cluster node, the CAU clustered role fails over to another cluster node, and a new Update Coordinator process on that node assumes control of the Updating Run. In self-updating mode, CAU can update the failover cluster by using a fully automated, end-to-end updating process. An administrator can also trigger updates on-demand in this mode, or simply use the remote-updating approach if desired. In self-updating mode, an administrator can get summary information about an Updating Run in progress by connecting to the cluster and running the **Get-CauRun** Windows PowerShell cmdlet.

- **Remote-updating mode** For this mode, a remote computer that is running Windows Server 2012 R2, Windows Server 2012, Windows 8.1 or Windows 8, which is called an Update Coordinator, is configured with the CAU tools. The Update Coordinator is not a member of the cluster that is updated during the Updating Run. From the remote computer, the administrator triggers an on-demand Updating Run by using a default or custom Updating Run profile. Remote-updating mode is useful for monitoring real-time progress during the Updating Run, and for clusters that are running on Server Core installations of Windows Server 2012 R2 or Windows Server 2012.

Hardware and software requirements

CAU can be used on all editions of Windows Server 2012 R2 and Windows Server 2012, including Server Core installations. For detailed requirements information, see Requirements and Best Practices for Cluster-Aware Updating.

To use CAU, you must install the Failover Clustering feature in Windows Server 2012 R2 or Windows Server 2012 and create a failover cluster. The components that support CAU functionality are automatically installed on each cluster node.

To install the Failover Clustering feature, you can use the following tools:

- Add Roles and Features Wizard in Server Manager
- **Add-WindowsFeature** Windows PowerShell cmdlet
- Deployment Image Servicing and Management (DISM) command-line tool

For more information, see Install or Uninstall Roles, Role Services, or Features.

You must also install the CAU tools, which are included in the Failover Clustering Tools (which are also part of the Remote Server Administration Tools, or RSAT). The CAU tools consist of the CAU UI and the CAU Windows PowerShell cmdlets. You must install the Failover Clustering Tools as follows to support the different CAU updating modes:

- To use CAU in self-updating mode, the Failover Clustering Tools must be installed on each cluster node. (This is the default installation.)
- To enable remote-updating mode, you must install the Failover Clustering Tools from the RSAT on a local or a remote computer that is running Windows Server 2012 R2, Windows Server 2012, Windows 8.1 or Windows 8 and that has network connectivity to the failover cluster.

Notes

- You must use the Failover Clustering Tools from the Windows Server 2012 R2 RSAT to remotely manage updates for a Windows Server 2012 R2 failover cluster. You can also use the Windows Server 2012 R2 RSAT to remotely manage updates on a Windows Server 2012 failover cluster.
- To use CAU only in remote-updating mode, installation of the Failover Clustering Tools on the cluster nodes is not required. However, certain CAU features will not be available. For more information, see Requirements and Best Practices for Cluster-Aware Updating.
- Unless you are using CAU only in self-updating mode, the computer on which the CAU tools are installed and that coordinates the updates cannot be a member of the failover cluster.

For more information about installing the Failover Clustering feature, see Installing the Failover Clustering Feature and Tools.

For more information about deploying RSAT, see Deploy Remote Server Administration Tools.

To enable self-updating mode, the CAU clustered role must also be added to the failover cluster. To do this by using the CAU UI, under **Cluster Actions**, use the **Configure Self-Updating Options** action. Alternatively, run the Add-CauClusterRole Windows PowerShell cmdlet.

To uninstall CAU, uninstall the Failover Clustering feature or Failover Clustering Tools by using Server Manager, Windows PowerShell cmdlets, or the DISM command-line tools.

Additional requirements and best practices

To ensure that CAU can update the cluster nodes successfully, and for additional guidance for configuring your failover cluster environment to use CAU, you can run the CAU Best Practices Analyzer model.

For detailed requirements and best practices for using CAU, and information about running the CAU Best Practices Analyzer model, see Requirements and Best Practices for Cluster-Aware Updating.

Starting CAU

You can start the CAU UI from the following locations:

- Server Manager
- ClusterUpdateUI.exe located in the %systemroot%\system32 folder
- Failover Cluster Manager

To start the CAU UI from Server Manager

1. Start Server Manager.
2. Do one of the following:
   - On the **Tools** menu, click **Cluster-Aware Updating**.
   - If one or more cluster nodes, or the cluster, is added to Server Manager, on the **All Servers** page, right-click the name of a node (or the name of the cluster), and then click **Update Cluster**.

For more information about managing servers by using Server Manager, see Manage Multiple, Remote Servers with Server Manager.

See also

The following table provides links to additional information about using CAU that is available on the web.

Content type	References
Product evaluation	Cluster-Aware Updating: Frequently Asked Questions
Planning	Requirements and Best Practices for Cluster-Aware Updating
Deployment	Requirements and Best Practices for Cluster-Aware Updating
Operations	Advanced Options and Updating Run Profiles for CAU | How CAU Plug-ins Work
Troubleshooting	Troubleshoot Cluster-Aware Updating (CAU) in Windows Server 2012 Failover Clusters
Tools and settings	Windows PowerShell® Cmdlets for Cluster-Aware Updating | Cluster-Aware Updating Plug-in Reference
Community resources	Storage Team at Microsoft: File Cabinet Blog
Related technologies	Failover Clustering Overview | File and Storage Services Overview | Hyper-V Overview | Scale-Out File Server for Application Data Overview | Windows Server Update Services | System Center Configuration Manager | System Center Virtual Machine Manager

Cluster-Aware Updating: Frequently Asked Questions

What is CAU?

Cluster-Aware Updating (CAU) is a feature that coordinates software updates on all servers in a failover cluster in a way that does not impact the service availability any more than a planned failover of a cluster node. For some applications with continuous availability features (such as Hyper-V with live migration, or an SMB 3.0 file server with SMB Transparent Failover), CAU can coordinate automated cluster updating with no impact on service availability.

Does CAU work with Windows Server 2008 R2 or Windows 7?

No. CAU coordinates the cluster updating operation only from computers running the Windows Server 2012 R2, Windows Server 2012, Windows 8.1 or Windows 8 operating system.

Is CAU functionality limited to only specific clustered applications?

No. CAU is agnostic to the type of the clustered application. CAU is an external cluster-updating solution that is layered on top of clustering APIs and Windows PowerShell cmdlets. As such, CAU can coordinate updating for any clustered application that is configured in a Windows Server 2012 R2 or Windows Server 2012 failover cluster.

Note

Currently, the following Windows Server 2012 clustered workloads are tested and certified for CAU: SMB, Hyper-V, DFS Replication, DFS Namespaces, iSCSI, and NFS.

Does CAU support updates from Microsoft Update and Windows Update?

Yes. By default, CAU is configured with a plug-in that uses the Windows Update Agent (WUA) utility APIs on the cluster nodes. The WUA infrastructure can be configured to point to Microsoft Update and Windows Update or to Windows Server Update Services (WSUS) as its source of updates.

Does CAU support WSUS updates?

Yes. By default, CAU is configured with a plug-in that uses the Windows Update Agent (WUA) utility APIs on the cluster nodes. The WUA infrastructure can be configured to point to Microsoft Update and Windows Update or to Windows Server Update Services (WSUS) as its source of updates.

Can CAU apply limited distribution release updates?

Yes. Limited distribution release (LDR) updates, also called hotfixes, are not published through Microsoft Update or Windows Update, so they cannot be downloaded by the Windows Update Agent (WUA) plug-in that CAU uses by default.

However, CAU includes a second plug-in that you can select to apply hotfix updates. This hotfix plug-in can also be customized to apply non-Microsoft driver, firmware, and BIOS updates.

Can I use CAU to apply cumulative updates?

Yes. If the cumulative updates are general distribution release updates or LDR updates, CAU can apply them.

Can I schedule CAU to apply updates?

Yes. CAU supports the following updating modes, both of which allow updates to be scheduled:

Self-updating Enables the cluster to update itself according to a defined profile and a regular schedule, such as during a monthly maintenance window. You can also start a Self-Updating Run on demand at any time. To enable self-updating mode, you must add the CAU clustered role to the cluster. The CAU self-updating feature performs like any other clustered workload, and it can work seamlessly with the planned and unplanned failovers of an update coordinator computer.

Remote-updating Enables you to start an Updating Run at any time from a computer running Windows Server 2012 R2, Windows Server 2012, Windows 8.1 or Windows 8. You can start an Updating run through the UI or by using the **Invoke-CauRun** Windows PowerShell cmdlet. Remote-updating is the default updating mode for CAU. You can use Task Scheduler to run the **Invoke-CauRun** cmdlet on a desired schedule from a remote computer that is not one of the cluster nodes.

Can CAU updates be scheduled when a backup is in progress?

Yes. CAU does not impose any constraints in this regard. However, performing software updates on a server (with the associated potential restarts) while a server backup is in progress is not an IT best practice. Be aware that CAU relies only on clustering APIs to determine resource failovers and failbacks; thus, CAU is unaware of the server backup status.

Can CAU work well with System Center Configuration Manager?

Yes. A server administrator should consider the following levels of interoperability between System Center Configuration Manager and CAU:

Coexistence CAU is a tool that coordinates software updates on a cluster node, and Configuration Manager also performs server software updates. It is important to configure these tools so that they do not have overlapping coverage of the same servers in any data-center deployment. This ensures that the objective behind using CAU is not inadvertently defeated, because Configuration Manager-driven updating does not incorporate cluster awareness.

Leverage CAU and Configuration Manager can work together to deliver synergistic value. By using the public plug-in interface architecture in CAU, Configuration Manager can leverage the cluster awareness of CAU. This allows a customer who already has Configuration Manager deployed to use the cluster awareness capabilities of CAU while taking advantage of the

Configuration Manager infrastructure, such as distribution points, approvals, and the Configuration console. For more information, see Cluster-Aware Updating Plug-in Reference.

Note
A Configuration Manager plug-in for CAU is not available in Windows Server 2012.

Do I need administrative credentials to run CAU?

Yes. For running the CAU tools, CAU needs administrative credentials on the local server, or it needs the **Impersonate a Client after Authentication** user right on the local server or the client computer on which it is running. However, to coordinate software updates on the cluster nodes, CAU requires cluster administrative credentials on every node. Although the CAU UI will start without the credentials, it prompts for the cluster administrative credentials when it connects to a cluster instance to preview or apply updates.

Can I script CAU functionality to automate it further?

Yes. CAU comes with Windows PowerShell cmdlets that offer a rich set of scripting options. These are the same cmdlets that the CAU UI calls to perform CAU actions.

When CAU coordinates updates, what happens to the clustered roles that are active on each cluster node?

Clustered roles (formerly called applications and services) that are active on a node, fail over to other nodes before software updating can commence. CAU orchestrates these failovers by using the maintenance mode, which pauses and drains the node of all active clustered roles. When the software updates are complete, CAU resumes the node and the clustered roles fail back to the updated node. This ensures that the distribution of clustered roles relative to nodes stays the same across the CAU Updating Runs of a cluster.

How does CAU select the target nodes that each clustered role should fail over to?

CAU relies on clustering APIs to coordinate the failovers. The clustering API implementation selects the target nodes by relying on internal metrics and intelligent placement heuristics (such as workload levels) across the target nodes.

Does CAU load balance the clustered roles at the end of the updating process?

CAU does not load balance the clustered nodes, but it attempts to preserve the distribution of clustered roles. When CAU finishes updating a cluster node, it attempts to fail back previously hosted clustered roles to that node. CAU relies on clustering APIs to fail back the resources to the beginning of the pause process. Thus in the absence of unplanned failovers and preferred owner settings, the distribution of clustered roles should remain unchanged.

How does CAU select the order of nodes to update?

By default, CAU selects the order of nodes to update based on the level of activity. The nodes that are hosting the fewest clustered roles are updated first. However, an administrator can specify a particular order for updating the nodes by specifying a parameter for the Updating Run in the CAU UI or by using the Windows PowerShell cmdlets.

When a CAU Updating Run is initiated on the cluster, what happens if all the cluster nodes are not online?

The administrator who initiates an Updating Run can specify the acceptable threshold for the number of nodes that can be offline. Therefore, an Updating Run can proceed on a cluster even if all the cluster nodes are not online.

Can I use CAU to select and update only a single node?

No. CAU is a cluster-scoped updating tool, so it only allows you to select clusters to update. If you want to update a single node, you can use existing server updating tools independently of CAU.

Can CAU report cluster node updates that are initiated from outside CAU?

No. CAU can only report Updating Runs that are initiated from within CAU. However, when a subsequent CAU Updating Run is launched, updates that were installed through non-CAU methods are appropriately considered to determine the additional updates that might be applicable to each cluster node.

Is CAU flexible enough to support my unique IT process needs?

Yes. CAU offers the following dimensions of flexibility to suit enterprise customers' unique IT process needs:

Scripts An Updating Run can specify a pre-update Windows PowerShell script and a post-update Windows PowerShell script. The pre-update script runs on each cluster node before the node is paused. The post-update script runs on each cluster node after the node updates are installed.

Note

.NET Framework 4.5 and Windows PowerShell must be installed on each cluster node on which you want to run the pre-update and post-update scripts. You must also enable Windows PowerShell remoting on the cluster nodes. For detailed system requirements, see Requirements and Best Practices for Cluster-Aware Updating.

Advanced Updating Run options The administrator can additionally specify from a large set of advanced Updating Run options such as the maximum number of times that the update process will be retried on each node. These options can be specified using either the CAU UI or the CAU

Windows PowerShell cmdlets. These custom settings can be saved in an Updating Run Profile and reused for later Updating Runs.

Public plug-in architecture CAU includes features to Register, Unregister, and Select plug-ins. CAU ships with two default plug-ins: one coordinates the Windows Update Agent (WUA) APIs on each cluster node; the second applies hotfixes that are manually copied to a file share that is accessible to the cluster nodes. If an enterprise has unique needs that cannot be met with these two plug-ins, the enterprise can build a new CAU plug-in according to the public API specification. For more information, see Cluster-Aware Updating Plug-in Reference.

For information about configuring and customizing CAU plug-ins to support different updating scenarios, see How CAU Plug-ins Work.

How can I export the CAU preview and update results?

CAU offers export options through the command-line interface and through the UI.

Command-line interface options:

- Preview results by using the Windows PowerShell cmdlet **Invoke-CauScan | ConvertTo-Xml**. Output: XML
- Report results by using the Windows PowerShell cmdlet **Invoke-CauRun | ConvertTo-Xml**. Output: XML
- Report results by using the Windows PowerShell cmdlet **Get-CauReport | Export-CauReport**. Output: HTML, CSV

UI options:

- Copy the report results from the **Preview updates** screen. Output: CSV
- Copy the report results from the **Generate report** screen. Output: CSV
- Export the report results from the **Generate report** screen. Output: HTML

How do I install CAU?

CAU is included with all Windows Server 2012 R2 and Windows Server 2012 editions. A CAU installation is seamlessly integrated into the Failover Clustering feature. CAU is installed as follows:

- When Failover Clustering is installed on a cluster node, the CAU Windows Management Instrumentation (WMI) provider is automatically installed.
- When Failover Clustering tools are installed on a server or client computer, the CAU Windows PowerShell cmdlets are automatically installed.
- When the Failover Clustering UI is installed on a server or client computer, the CAU Windows PowerShell cmdlets and the CAU UI are automatically installed.

Does CAU need any service or component running on the cluster nodes that are being updated?

CAU does not need a service running on the cluster nodes. However, CAU needs a software component (the WMI provider) installed on the cluster nodes. This component is installed with the Failover Clustering feature.

To enable self-updating mode, the CAU clustered role must also be added to the cluster.

What is the difference between using CAU and using System Center 2012 Virtual Machine Manager to update Hyper-V clusters?

- System Center 2012 Virtual Machine Manager (VMM) is focused on updating only Hyper-V clusters, whereas CAU can update any type of supported failover cluster, including Hyper-V clusters.
- VMM requires additional licensing, whereas CAU is licensed for all Windows Server 2012 R2 and Windows Server 2012 editions. The CAU features, tools, and UI are installed with Failover Clustering components.
- If you already own a System Center license, you can continue to use VMM to update Hyper-V clusters because it offers an integrated management and software updating experience.
- CAU is supported only on clusters that are running Windows Server 2012 R2 and Windows Server 2012. VMM also supports Hyper-V clusters on computers running Windows Server 2008 and Windows Server 2008 R2.

Can I use remote-updating on a cluster that is configured for self-updating?

Yes. A failover cluster in a self-updating configuration can be updated through remote-updating on-demand, just as you can force a Windows Update scan at any time on your computer, even if Windows Update is configured to install updates automatically. However, you need to make sure that an Updating Run is not already in progress.

Can I reuse my cluster update settings across clusters?

Yes. CAU supports a number of Updating Run options that determine how the Updating Run behaves when it updates the cluster. These options can be saved as an Updating Run Profile, and they can be reused across any cluster. We recommend that you save and reuse your settings across failover clusters that have similar updating needs. For example, you might create a “Business-Critical SQL Server Cluster Updating Run Profile” for all Microsoft SQL Server clusters that support business-critical services.

Where can I get more information about the public CAU plug-in specification?

- Cluster-Aware Updating Plug-in Reference

- Cluster Aware Updating plug-in sample

See also

- Cluster-Aware Updating Overview

Requirements and Best Practices for Cluster-Aware Updating

This section describes the requirements and dependencies that are needed to use Cluster-Aware Updating (CAU) to apply updates to a Windows Server 2012 R2 or Windows Server 2012 failover cluster. For additional information about best practices and tests for cluster updating readiness, see the following sections later in this topic:

- Best practices recommendations for using CAU
- Test cluster updating readiness

Note
You may need to independently validate that your cluster environment is ready to apply updates if you use a plug-in other than **Microsoft.WindowsUpdatePlugin**. If you are using a non-Microsoft plug-in, contact the publisher for more information. For more information about plug-ins, see How CAU Plug-ins Work.

For an overview of the CAU feature and CAU updating scenarios, see Cluster-Aware Updating Overview.

Install the Failover Clustering feature and the Failover Clustering Tools

CAU requires an installation of the Failover Clustering feature and the Failover Clustering Tools. The Failover Clustering Tools include the CAU tools (clusterawareupdating.dll), the Failover Clustering Windows PowerShell cmdlets, and other components needed for CAU operations. For steps to install the Failover Clustering feature, see Installing the Failover Clustering Feature and Tools.

The exact installation requirements for the Failover Clustering Tools depend on whether CAU coordinates updates as a clustered role on the failover cluster (by using self-updating mode) or from a remote computer running Windows Server 2012 R2, Windows Server 2012, Windows 8.1 or Windows 8 (by using remote-updating mode). The self-updating mode of CAU additionally requires the installation of the CAU clustered role on the failover cluster by using the CAU tools.

Note

You must use the Failover Clustering Tools from the Windows Server 2012 R2 RSAT to remotely manage updates for a Windows Server 2012 R2 failover cluster. You can also use the Windows Server 2012 R2 RSAT to remotely manage updates on a Windows Server 2012 failover cluster.

The following table summarizes the CAU feature installation requirements for the two CAU updating modes.

Installed component	Self-updating mode	Remote-updating mode
Failover Clustering feature	Required on all cluster nodes	Required on all cluster nodes
Failover Clustering Tools	Required on all cluster nodes	• Required on remote-updating computer • Required on all cluster nodes to run the Save-CauDebugTrace Windows PowerShell cmdlet
CAU clustered role	Required	Not required

Obtain an administrator account

The following administrator requirements are necessary to use CAU features.

- To preview or apply update actions by using the CAU user interface (UI) or the CAU Windows PowerShell cmdlets, you must use a domain account that has local administrator rights and permissions on all the cluster nodes. If the account does not have sufficient privileges on every node, you are prompted in the CAU UI to supply the necessary credentials when you perform these actions. To use the CAU Windows PowerShell cmdlets, you can supply the necessary credentials as a cmdlet parameter.
- If you use CAU in remote-updating mode when you are logged on with an account that does not have local administrator rights and permissions on the cluster nodes, you must run the CAU tools as an administrator by using a local administrator account on the Update Coordinator computer, or by using an account that has the **Impersonate a client after authentication** user right.
- To run the CAU Best Practices Analyzer, you must use an account that has administrative privileges on the cluster nodes and local administrative privileges on the computer that is used to run the **Test-CauSetup** cmdlet or to analyze cluster updating readiness using the CAU UI. For more information, see Test cluster updating readiness.

Verify the cluster configuration

The following are general requirements for a failover cluster to support updates by using CAU. Additional configuration requirements for remote management on the nodes are listed in Configure the nodes for remote management later in this topic.

- Sufficient cluster nodes must be online so that the cluster has quorum.
- All cluster nodes must be in the same Active Directory domain.
- The cluster name must be resolved on the network using DNS.
- If CAU is used in remote-updating mode, the Update Coordinator computer must have network connectivity to the failover cluster nodes, and it must be in the same Active Directory domain as the failover cluster.
- The Cluster service should be running on all cluster nodes. By default this service is installed on all cluster nodes and is configured to start automatically.
- To use Windows PowerShell pre-update or post-update scripts during a CAU Updating Run, ensure that the scripts are installed on all cluster nodes or that they are accessible to all nodes, for example, on a highly available network file share. If scripts are saved to a network file share, configure the folder for Read permission for the Everyone group.

Configure the nodes for remote management

The following table and the remaining sections in this topic summarize additional configuration or software that must be installed on all cluster nodes to remotely support the two CAU updating modes. These requirements are in addition to the installation requirements for the Install the Failover Clustering feature and the Failover Clustering Tools and the general clustering requirements that are described in previous sections in this topic.

Requirement	Self-updating mode	Remote-updating mode
Enable Windows Management Instrumentation	Required on all cluster nodes	Required on all cluster nodes
Enable Windows PowerShell and Windows PowerShell remoting	Required on all cluster nodes	Required on all cluster nodes to run the following: • The Save-CauDebugTrace Windows PowerShell cmdlet • Windows PowerShell pre-update and post-update scripts during an Updating Run • Tests of cluster updating readiness using the CAU UI or the Test-CauSetup Windows PowerShell cmdlet
Install .NET Framework 4.5	Required on all cluster nodes	Required on all cluster nodes to run the following: • The Save-CauDebugTrace Windows PowerShell cmdlet • Windows PowerShell pre-

Requirement	Self-updating mode	Remote-updating mode
		update and post-update scripts during an Updating Run • Tests of cluster updating readiness using the CAU UI or the Test-CauSetup Windows PowerShell cmdlet
Enable a firewall rule to allow automatic restarts	Required on all cluster nodes if a firewall is in use	Required on all cluster nodes if a firewall is in use

Enable Windows Management Instrumentation

All cluster nodes must be configured for remote management using Windows Management Instrumentation (WMI). This is enabled by default.

To manually enable remote management, do the following:

1. In the Services console, start the **Windows Remote Management** service and set the startup type to **Automatic**.
2. Run the **Set-WSManQuickConfig** Windows PowerShell cmdlet, or run the following command from an elevated command prompt:

```
winrm quickconfig -q
```

To support WMI remoting, if Windows Firewall is in use on the cluster nodes, the inbound firewall rule for **Windows Remote Management (HTTP-In)** must be enabled on each node. By default, this rule is enabled.

Enable Windows PowerShell and Windows PowerShell remoting

To enable self-updating mode and certain CAU features in remote-updating mode, Windows PowerShell 3.0 or 4.0 must be installed and enabled to run remote commands on all cluster nodes. By default, Windows PowerShell is installed and is enabled for remoting.

To install the Windows PowerShell feature, use the Add Roles and Features Wizard in Server Manager.

To enable Windows PowerShell remoting, you can use one of the following methods:

- Run the **Enable-PSRemoting** Windows PowerShell cmdlet.
- Configure a domain-level Group Policy setting for Windows Remote Management (WinRM).

For more information about enabling Windows PowerShell remoting, see about_Remote_Requirements.

Install .NET Framework 4.5

To enable self-updating mode and certain CAU features in remote-updating mode, .NET Framework 4.5 must be installed on all cluster nodes. By default, NET Framework 4.5 is installed.

To install the .NET Framework 4.5 Features, use the Add Roles and Features Wizard in Server Manager. For more information, see Install or Uninstall Roles, Role Services, or Features.

Enable a firewall rule to allow automatic restarts

To allow automatic restarts after updates are applied (if the installation of an update requires a restart), if Windows Firewall or a non-Microsoft firewall is in use on the cluster nodes, a firewall rule must be enabled on each node that allows the following traffic:

- Protocol: TCP
- Direction: inbound
- Program: wininit.exe
- Ports: RPC Dynamic Ports
- Profile: Domain

If Windows Firewall is used on the cluster nodes, you can do this by enabling the **Remote Shutdown** Windows Firewall rule group on each cluster node. When you use the CAU UI to apply updates and to configure self-updating options, the **Remote Shutdown** Windows Firewall rule group is automatically enabled on each cluster node.

Note
The **Remote Shutdown** Windows Firewall rule group cannot be enabled when it will conflict with Group Policy settings that are configured for Windows Firewall.

Windows PowerShell equivalent commands

The **Remote Shutdown** firewall rule group is also enabled by specifying the **–EnableFirewallRules** parameter when running the following CAU Windows PowerShell cmdlets: Add-CauClusterRole, Invoke-CauRun, and SetCauClusterRole.

The following example shows an additional method to enable automatic restarts on a cluster node.

```
Set-NetFirewallRule -Group "@firewallapi.dll,-36751" -Profile Domain -Enabled true
```

Best practices recommendations for using CAU

Recommendations for applying Microsoft updates

We recommend that when you begin to use CAU to apply updates with the default **Microsoft.WindowsUpdatePlugin** plug-in on a cluster, you stop using other methods to install software updates from Microsoft on the cluster nodes.

Caution

Combining CAU with methods that update individual nodes automatically (on a fixed time schedule) can cause unpredictable results, including interruptions in service and unplanned downtime.

We recommend that you follow these guidelines:

- For optimal results, we recommend that you disable settings on the cluster nodes for automatic updating, for example, through the Automatic Updates settings in Control Panel, or in settings that are configured using Group Policy.

  Caution

  Automatic installation of updates on the cluster nodes can interfere with installation of updates by CAU and can cause CAU failures.

  If they are needed, the following Automatic Updates settings are compatible with CAU, because the administrator can control the timing of update installation:

  - Settings to notify before downloading updates and to notify before installation
  - Settings to automatically download updates and to notify before installation

  However, if Automatic Updates is downloading updates at the same time as a CAU Updating Run, the Updating Run might take longer to complete.
- Do not configure an update system such as Windows Server Update Services (WSUS) to apply updates automatically (on a fixed time schedule) to cluster nodes.
- All cluster nodes should be uniformly configured to use the same update source, for example, a WSUS server, Windows Update, or Microsoft Update.
- If you use a configuration management system to apply software updates to computers on the network, exclude cluster nodes from all required or automatic updates. Examples of configuration management systems include Microsoft System Center Configuration Manager 2007 and Microsoft System Center Virtual Machine Manager 2008.
- If internal software distribution servers (for example, WSUS servers) are used to contain and deploy the updates, ensure that those servers correctly identify the approved updates for the cluster nodes.

Apply Microsoft updates in branch office scenarios

To download Microsoft updates from Microsoft Update or Windows Update to cluster nodes in certain branch office scenarios, you may need to configure proxy settings for the Local System account on each node. For example, you might need to do this if your branch office clusters access Microsoft Update or Windows Update to download updates by using a local proxy server.

If necessary, configure WinHTTP proxy settings on each node to specify a local proxy server and configure local address exceptions (that is, a bypass list for local addresses). To do this, you can run the following command on each cluster node from an elevated command prompt:

```
netsh winhttp set proxy <ProxyServerFQDN >:<port> "<local>"
```

where *<ProxyServerFQDN>* is the fully qualified domain name for the proxy server.

For example, to configure WinHTTP proxy settings for the Local System account specifying the proxy server *MyProxy.CONTOSO.com*, with local address exceptions, type the following command:

```
netsh winhttp set proxy MyProxy.CONTOSO.com:443 "<local>"
```

Recommendations for using the Microsoft.HotfixPlugin

- We recommend that you configure permissions in the hotfix root folder and hotfix configuration file to restrict Write access to only local administrators on the computers that are used to store these files. This helps prevent tampering with these files by unauthorized users that could compromise the functionality of the failover cluster when hotfixes are applied.
- To help ensure data integrity for the server message block (SMB) connections that are used to access the hotfix root folder, you should configure SMB Encryption in the SMB shared folder, if it is possible to configure it. The **Microsoft.HotfixPlugin** requires that SMB signing or SMB Encryption is configured to help ensure data integrity for the SMB connections.

  Note
  SMB Encryption, which provides enhanced security and better performance in many environments, is supported starting in Windows Server 2012.

  For more information, see Restrict access to the hotfix root folder and hotfix configuration file in How CAU Plug-ins Work.

Additional recommendations

- To avoid interfering with a CAU Updating Run that may be scheduled at the same time, do not schedule password changes for cluster name objects and virtual computer objects during scheduled maintenance windows.
- You should set appropriate permissions on pre-update and post-update scripts that are saved on network shared folders to prevent potential tampering with these files by unauthorized users.
- To configure CAU in self-updating mode, a virtual computer object (VCO) for the CAU clustered role must be created in Active Directory. CAU can create this object automatically at the time that the CAU clustered role is added, if the failover cluster has sufficient permissions. However, because of the security policies in certain organizations, it may be necessary to prestage the object in Active Directory. For a procedure to do this, see Steps for prestaging an account for a clustered role.
- To save and reuse Updating Run settings across failover clusters with similar updating needs in the IT organization, you can create Updating Run Profiles. Additionally, depending on the updating mode, you can save and manage the Updating Run Profiles on a file share that is accessible to all remote Update Coordinator computers or failover clusters. For more information, see Advanced Options and Updating Run Profiles for CAU.

Test cluster updating readiness

You can run the CAU Best Practices Analyzer (BPA) model to test whether a specified failover cluster and the network environment meet many of the requirements to have software updates applied by CAU. Many of the tests check the environment for readiness to apply Microsoft updates by using the default plug-in, **Microsoft.WindowsUpdatePlugin**.

Note

You may need to independently validate that your cluster environment is ready to apply software updates by using a plug-in other than **Microsoft.WindowsUpdatePlugin**. If you are using a non-Microsoft plug-in, contact the publisher for more information.

You can run the BPA in the following two ways:

1. Select **Analyze cluster updating readiness** in the CAU console. After the BPA completes the readiness tests, a test report appears in the CAU UI. If issues are detected on cluster nodes, the specific issues and the nodes where the issues appear are identified so that you can take corrective action. The tests can take several minutes to complete.
2. Run the **Test-CauSetup** Windows PowerShell cmdlet. For more information, see Test-CauSetup. You can run the cmdlet on a local or remote computer on which the CAU Windows PowerShell cmdlets are installed. You can also run the cmdlet on a node of the failover cluster.

Notes

- You must use an account that has administrative privileges on the cluster nodes and local administrative privileges on the computer that is used to run the **Test-CauSetup** cmdlet or to analyze cluster updating readiness using the CAU UI. To run the tests using the CAU UI, you must be logged on to the computer with the necessary credentials.
- The tests assume that the CAU tools that are used to preview and apply software updates run from the same computer and with the same user credentials as are used to test cluster updating readiness.

Important

We highly recommend that you test the cluster for updating readiness in the following situations:

- Before you use CAU for the first time to apply software updates.
- After you add a node to the cluster or perform other hardware changes in the cluster that require running the Validate a Cluster Wizard.
- After you change an update source, or change update settings or configurations (other than CAU) that can affect the application of updates on the nodes.

Tests for cluster updating readiness

The following table lists the cluster updating readiness tests, some common issues, and resolution steps.

Test	Possible issues and impacts	Resolution steps
The failover cluster must be available	Cannot resolve the failover cluster name, or one or more cluster nodes cannot be accessed. The BPA cannot run the cluster readiness tests.	• Check the spelling of the name of the cluster specified during the BPA run. • Ensure that all nodes of the cluster are online and running. • Check that the Validate a Configuration Wizard can successfully run on the failover cluster.
The failover cluster nodes must be enabled for remote management via WMI	One or more failover cluster nodes are not enabled for remote management by using Windows Management Instrumentation (WMI). CAU cannot update the cluster nodes if the nodes are not configured for remote management.	Ensure that all failover cluster nodes are enabled for remote management through WMI. For more information, see Configure the nodes for remote management in this topic.
Windows PowerShell remoting should be enabled on each failover cluster node	Windows PowerShell 3.0 or 4.0 is not installed or is not enabled for remoting on one or more failover cluster nodes. CAU cannot be configured for self-updating mode or use certain features in remote-updating mode.	Ensure that Windows PowerShell is installed on all cluster nodes and is enabled for remoting. For more information, see Configure the nodes for remote management in this topic.
Failover cluster version	One or more nodes in the failover cluster do not run the Windows Server 2012 R2 or Windows Server 2012 operating system. CAU cannot update the failover cluster.	Verify that the failover cluster that is specified during the BPA run is running Windows Server 2012 R2 or Windows Server 2012. For more information, see Install the Failover Clustering feature and the

Test	Possible issues and impacts	Resolution steps
		Failover Clustering Tools in this topic.
The required versions of .NET Framework and Windows PowerShell must be installed on all failover cluster nodes	.NET Framework 4.5 or Windows PowerShell 3.0 or 4.0 is not installed on one or more cluster nodes. Some CAU features might not work.	Ensure that .NET Framework 4.5 and Windows PowerShell are installed on all cluster nodes, if they are required. For more information, see Configure the nodes for remote management in this topic.
The Cluster service should be running on all cluster nodes	The Cluster service is not running on one or more nodes. CAU cannot update the failover cluster.	• Ensure that the Cluster service (clussvc) is started on all nodes in the cluster, and it is configured to start automatically. • Check that the Validate a Configuration Wizard can successfully run on the failover cluster. For more information, see Install the Failover Clustering feature and the Failover Clustering Tools in this topic.
Automatic Updates must not be configured to automatically install updates on any failover cluster node	On at least one failover cluster node, Automatic Updates is configured to automatically install Microsoft updates on that node. Combining CAU with other update methods can result in unplanned downtime or unpredictable results.	If Windows Update functionality is configured for Automatic Updates on one or more cluster nodes, ensure that Automatic Updates is not configured to automatically install updates. For more information, see Recommendations for applying Microsoft updates.
The failover cluster nodes should use the	One or more failover cluster nodes are	Ensure that every cluster node is configured to use

Test	Possible issues and impacts	Resolution steps
same update source	configured to use an update source for Microsoft updates that is different from the rest of the nodes. Updates might not be applied uniformly on the cluster nodes by CAU.	the same update source, for example, a WSUS server, Windows Update, or Microsoft Update. For more information, see Recommendations for applying Microsoft updates.
A firewall rule that allows remote shutdown should be enabled on each node in the failover cluster	One or more failover cluster nodes do not have a firewall rule enabled that allows remote shutdown, or a Group Policy setting prevents this rule from being enabled. An Updating Run that applies updates that require restarting the nodes automatically might not complete properly.	If Windows Firewall or a non-Microsoft firewall is in use on the cluster nodes, configure a firewall rule that allows remote shutdown. For more information, see Enable a firewall rule to allow automatic restarts in this topic.
The proxy server setting on each failover cluster node should be set to a local proxy server	One or more failover cluster nodes have an incorrect proxy server configuration. If a local proxy server is in use, the proxy server setting on each node must be configured properly for the cluster to access Microsoft Update or Windows Update.	Ensure that the WinHTTP proxy settings on each cluster node are set to a local proxy server if it is needed. If a proxy server is not in use in your environment, this warning can be ignored. For more information, see Apply Microsoft updates in branch office scenarios in this topic.
The CAU clustered role should be installed on the failover cluster to enable self-updating mode	The CAU clustered role is not installed on this failover cluster. This role is required for cluster self-updating.	To use CAU in self-updating mode, add the CAU clustered role on the failover cluster in one of the following ways:

Test	Possible issues and impacts	Resolution steps
		• Run the Add-CauClusterRole Windows PowerShell cmdlet. • Select the **Configure cluster self-updating options** action in the CAU UI.
The CAU clustered role should be enabled on the failover cluster to enable self-updating mode	The CAU clustered role is disabled. For example, the CAU clustered role is not installed, or it has been disabled by using the Disable-CauClusterRole Windows PowerShell cmdlet. This role is required for cluster self-updating.	To use CAU in self-updating mode, enable the CAU clustered role on this failover cluster in one of the following ways: • Run the Enable-CauClusterRole Windows PowerShell cmdlet. • Select the **Configure cluster self-updating options** action in the CAU UI.
The configured CAU plug-in for self-updating mode must be registered on all failover cluster nodes	The CAU clustered role on one or more nodes of this failover cluster cannot access the CAU plug-in module that is configured in the self-updating options. A self-updating run might fail.	• Ensure that the configured CAU plug-in is installed on all cluster nodes by following the installation procedure for the product that supplies the CAU plug-in. • Run the Register-CauPlugin Windows PowerShell cmdlet to register the plug-in on the required cluster nodes.
All failover cluster nodes should have the same set of registered CAU plug-ins	A self-updating run might fail if the plug-in that is configured to be used in an Updating Run is changed to one that is	• Ensure that the configured CAU plug-in is installed on all cluster nodes by following the

Test	Possible issues and impacts	Resolution steps
	not available on all cluster nodes.	installation procedure for the product that supplies the CAU plug-in. • Run the Register-CauPlugin Windows PowerShell cmdlet to register the plug-in on the required cluster nodes.
The configured Updating Run options must be valid	The self-updating schedule and Updating Run options that are configured for this failover cluster are incomplete or are not valid. A self-updating run might fail.	Configure a valid self-updating schedule and set of Updating Run options. For example, you can use the Set-CauClusterRole Windows PowerShell cmdlet to configure the CAU clustered role.
At least two failover cluster nodes must be owners of the CAU clustered role	An Updating Run launched in self-updating mode will fail because the CAU clustered role does not have a possible owner node to move to.	Use the Failover Clustering Tools to ensure that all cluster nodes are configured as possible owners of the CAU clustered role. This is the default configuration.
All failover cluster nodes must be able to access Windows PowerShell scripts	Not all possible owner nodes of the CAU clustered role can access the configured Windows PowerShell pre-update and post-update scripts. A self-updating run will fail.	Ensure that all possible owner nodes of the CAU clustered role have permissions to access the configured Windows PowerShell pre-update and post-update scripts.
All failover cluster nodes should use identical Windows PowerShell scripts	Not all possible owner nodes of the CAU clustered role use the same copy of the specified Windows PowerShell pre-update	Ensure that all possible owner nodes of the CAU clustered role use the same Windows PowerShell pre-update and post-update scripts.

Test	Possible issues and impacts	Resolution steps
	and post-update scripts. A self-updating run might fail or show unexpected behavior.	
The WarnAfter setting specified for the Updating Run should be less than the StopAfter setting	The specified CAU Updating Run timeout values make the warning timeout ineffective. An Updating Run might be canceled before a warning event log can be generated.	In the Updating Run options, configure a **WarnAfter** option value that is less than the **StopAfter** option value.

See also

- Cluster-Aware Updating Overview

How CAU Plug-ins Work

To preview or apply updates by using the Cluster-Aware Updating (CAU) user interface (UI) or the CAU Windows PowerShell cmdlets, you must specify one or more CAU plug-ins. A plug-in installs updates across all the nodes in a failover cluster by using a specific updating tool and update source. This topic provides information about using the built-in CAU plug-ins or other plug-ins that you install for CAU.

For detailed guidance about the CAU built-in plug-ins, see the following sections later in this topic:

- Use the Microsoft.WindowsUpdatePlugin
- Use the Microsoft.HotfixPlugin

Install a plug-in

A plug-in other than the default plug-ins that are installed with CAU (**Microsoft.WindowsUpdatePlugin** and **Microsoft.HotfixPlugin**) must be installed separately. If CAU is used in self-updating mode, the plug-in must be installed on all cluster nodes. If CAU is used in remote-updating mode, the plug-in must be installed on the remote Update Coordinator computer. A plug-in that you install may have additional installation requirements on each node.

To install a plug-in, follow the instructions from the plug-in publisher. To manually register a plug-in with CAU, run the Register-CauPlugin cmdlet on each computer where the plug-in is installed.

Specify a plug-in and plug-in arguments

Specify a CAU plug-in

Tip

In the CAU UI, you can only specify a single plug-in for CAU to use to preview or to apply updates during an Updating Run. By using the CAU Windows PowerShell cmdlets, you can specify one or more plug-ins. If you need to install multiple types of updates on the cluster, it is usually more efficient to specify multiple plug-ins in one Updating Run, rather than using a separate Updating Run for each plug-in. For example, fewer node restarts will typically occur.

In the CAU UI, you select a plug-in from a drop-down list of available plug-ins when you use CAU to perform the following actions:

- Apply updates to the cluster
- Preview updates for the cluster
- Configure cluster self-updating options

By default, CAU selects the plug-in **Microsoft.WindowsUpdatePlugin**. However, you can specify any plug-in that is installed and registered with CAU.

By using the CAU Windows PowerShell cmdlets that are listed in the following table, you can specify one or more plug-ins for an Updating Run or scan by passing the **–CauPluginName** parameter. You can specify multiple plug-in names by separating them with commas. If you specify multiple plug-ins, you can also control how the plug-ins influence each other during an Updating Run by specifying the **-RunPluginsSerially**, **-StopOnPluginFailure**, and **–SeparateReboots** parameters. For more information about using multiple plug-ins, use the links provided to the cmdlet documentation in the following table.

Cmdlet	Description
Add-CauClusterRole	Adds the CAU clustered role that provides the self-updating functionality to the specified cluster.
Invoke-CauRun	Performs a scan of cluster nodes for applicable updates and installs those updates through an Updating Run on the specified cluster.
Invoke-CauScan	Performs a scan of cluster nodes for applicable updates and returns a list of the initial set of updates that would be applied to each node in the specified cluster.
Set-CauClusterRole	Sets configuration properties for the CAU clustered role on the specified cluster.

If you do not specify a CAU plug-in parameter by using these cmdlets, the default is the plug-in **Microsoft.WindowsUpdatePlugin**.

Specify CAU plug-in arguments

When you configure the Updating Run options, you can specify one or more *name=value* pairs (arguments) for the selected plug-in to use. For example, in the CAU UI, you can specify multiple arguments as follows:

Name1=Value1;Name2=Value2;Name3=Value3

These *name=value* pairs must be meaningful to the plug-in that you specify. For some plug-ins the arguments are optional.

The syntax of the CAU plug-in arguments follows these general rules:

- Multiple *name=value* pairs are separated by semicolons.
- A value that contains spaces is surrounded by quotation marks, for example: **Name1="Value with Spaces"**.
- The exact syntax of *value* depends on the plug-in.

To specify plug-in arguments by using the CAU Windows PowerShell cmdlets that support the **–CauPluginParameters** parameter, pass a parameter of the form:

-CauPluginArguments @{Name1=Value1;Name2=Value2;Name3=Value3}

You can also use a predefined Windows PowerShell hash table. To specify plug-in arguments for more than one plug-in, pass multiple hash tables of arguments, separated with commas. Pass the plug-in arguments in the plug-in order that is specified in **CauPluginName**.

Specify optional plug-in arguments

The plug-ins that CAU installs (**Microsoft.WindowsUpdatePlugin** and **Microsoft.HotfixPlugin**) provide additional options that you can select. In the CAU UI, these appear on an **Additional Options** page after you configure Updating Run options for the plug-in. If you are using the CAU Windows PowerShell cmdlets, these options are configured as optional plug-in arguments. For more information, see Use the Microsoft.WindowsUpdatePlugin and Use the Microsoft.HotfixPlugin later in this topic.

Manage plug-ins using Windows PowerShell cmdlets

Cmdlet	Description
Get-CauPlugin	Retrieves information about one or more software updating plug-ins that are registered on the local computer.

Cmdlet	Description
Register-CauPlugin	Registers a CAU software updating plug-in on the local computer.
Unregister-CauPlugin	Removes a software updating plug-in from the list of plug-ins that can be used by CAU. **Note** The plug-ins that are installed with CAU (**Microsoft.WindowsUpdatePlugin** and the **Microsoft.HotfixPlugin**) cannot be unregistered.

Use the Microsoft.WindowsUpdatePlugin

The following table summarizes information about the default plug-in for CAU: **Microsoft.WindowsUpdatePlugin**.

Item	Details
Description	• Communicates with the Windows Update Agent on each failover cluster node to apply updates that are needed for the Microsoft products that are running on each node. • Installs cluster updates directly from Windows Update or Microsoft Update, or from an on-premises Windows Server Update Services (WSUS) server. • Installs only selected, general distribution release (GDR) updates. By default, the plug-in applies only important software updates. **Note** To apply updates other than the important software updates that are selected by default (for example, driver updates), you can configure an optional plug-in parameter. For more information, see Configure the Windows Update Agent query string.
Requirements	• The failover cluster and remote Update

Item	Details
	Coordinator computer (if used) must meet the requirements for CAU and the configuration that is required for remote management listed in Requirements and Best Practices for Cluster-Aware Updating. • Review Recommendations for applying Microsoft updates, and then make any necessary changes to your Microsoft Update configuration for the failover cluster nodes. • For best results, we recommend that you run the CAU Best Practices Analyzer (BPA) to ensure that the cluster and update environment are configured properly to apply updates by using CAU. For more information, see Test CAU updating readiness.
Configuration	No configuration is required. The default configuration downloads and installs important GDR updates on each node. **Note** Updates that require the acceptance of Microsoft license terms or require user interaction are excluded, and they must be installed manually.
Additional options	Optionally, you can specify either of the following plug-in arguments to augment or restrict the set of updates that are applied by the plug-in: • To configure the plug-in to apply recommended updates in addition to important updates on each node, in the CAU UI, on the **Additional Options** page, select the **Give me recommended updates the same way that I receive important updates** check box. Alternatively, configure the **'IncludeRecommendedUpdates'='True'** plug-in argument. • To configure the plug-in to filter the types of GDR updates that are applied to each

Item	Details
	cluster node, specify a Windows Update Agent query string using a **QueryString** plug-in argument. For more information, see Configure the Windows Update Agent query string.

Configure the Windows Update Agent query string

You can configure a plug-in argument for the default plug-in, **Microsoft.WindowsUpdatePlugin**, that consists of a Windows Update Agent (WUA) query string. This instruction uses the WUA API to identify one or more groups of Microsoft updates to apply to each node, based on specific selection criteria. You can combine multiple criteria by using a logical AND or a logical OR. The WUA query string is specified in a plug-in argument as follows:

QueryString="Criterion1=Value1 and/or Criterion2=Value2 and/or…"

For example, **Microsoft.WindowsUpdatePlugin** automatically selects important updates by using a default **QueryString** argument that is constructed using the **IsInstalled**, **Type**, **IsHidden**, and **IsAssigned** criteria:

QueryString="IsInstalled=0 and Type='Software' and IsHidden=0 and IsAssigned=1"

If you specify a **QueryString** argument, it is used in place of the default **QueryString** that is configured for the plug-in.

Example 1

To configure a **QueryString** argument that installs a specific update as identified by ID *f6ce46c1-971c-43f9-a2aa-783df125f003*:

QueryString="UpdateID='f6ce46c1-971c-43f9-a2aa-783df125f003' and IsInstalled=0"

Notes

The preceding example is valid for applying updates by using the Cluster-Aware Updating Wizard. If you want to install a specific update by configuring self-updating options with the CAU UI or by using the **Add-CauClusterRole** or **Set-CauClusterRole** Windows PowerShell cmdlet, you must format the UpdateID value with two single-quote characters:

QueryString="UpdateID=''f6ce46c1-971c-43f9-a2aa-783df125f003'' and IsInstalled=0"

Example 2

To configure a **QueryString** argument that installs only drivers:

QueryString="IsInstalled=0 and Type='Driver' and IsHidden=0"

For more information about query strings for the default plug-in, **Microsoft.WindowsUpdatePlugin**, the search criteria (such as **IsInstalled**), and the syntax that

you can include in the query strings, see the section about search criteria in the Windows Update Agent (WUA) API Reference.

Use the Microsoft.HotfixPlugin

The plug-in **Microsoft.HotfixPlugin** can be used to apply Microsoft limited distribution release (LDR) updates (also called hotfixes, and formerly called QFEs) that you download independently to address specific Microsoft software issues.

Item	Details
Description	• Installs LDR Microsoft updates (hotfixes) from a root folder on an SMB file share. **Note** Hotfixes are sometimes available for download from Microsoft in Knowledge Base articles, but they are also provided to customers on an as-needed basis. • Can also be customized to apply non-Microsoft driver, firmware, and BIOS updates.
Requirements	• The failover cluster and remote Update Coordinator computer (if used) must meet the requirements for CAU and the configuration that is required for remote management listed in Requirements and Best Practices for Cluster-Aware Updating. • Review Recommendations for using the Microsoft.HotfixPlugin. • For best results, we recommend that you run the CAU Best Practices Analyzer (BPA) model to ensure that the cluster and update environment are configured properly to apply updates by using CAU. For more information, see Test CAU updating readiness. • Obtain the updates from the publisher, and copy them or extract them to a Server Message Block (SMB) file share (hotfix root folder) that supports at least SMB 2.0 and that is accessible by all of the cluster nodes and the remote Update Coordinator computer (if CAU is used in remote-updating mode). For more information, see Configure a hotfix root folder structure later in this topic. **Note** By default, this plug-in only Installs hotfixes with the following file name extensions: .msu, .msi, and .msp. • Copy the DefaultHotfixConfig.xml file (which is provided in the **%systemroot%\System32\WindowsPowerShell\v1.0\Modules\ClusterAwareUpdating** folder on a computer where the CAU tools are installed) to the hotfix root folder that you created and under which you extracted the hotfixes. For example, copy the configuration file to *\\MyFileServer\Hotfixes\Root*\.

Item	Details
	Note To install most hotfixes provided by Microsoft and other updates, the default hotfix configuration file can be used without modification. If your scenario requires it, you can customize the configuration file as an advanced task. The configuration file can include custom rules, for example, to handle hotfix files that have specific extensions, or to define behaviors for specific exit conditions. For more information, see Customize the hotfix configuration file later in this topic.
Configuration	Configure the following settings. For more information, see the links to sections later in this topic. • The path to the shared hotfix root folder that contains the updates to apply and that contains the hotfix configuration file. You can type this path in the CAU UI or configure the **HotfixRootFolderPath=<Path>** Windows PowerShell plug-in argument. **Note** You can specify the hotfix root folder as a local folder path or as a UNC path of the form *ServerName**Share**RootFolderName*. A domain-based or standalone DFS Namespace path can be used. However, the plug-in features that check access permissions in the hotfix configuration file are incompatible with a DFS Namespace path, so if you configure one, you must disable the check for access permissions by using the CAU UI or by configuring the **DisableAclChecks='True'** plug-in argument. For more information, see Configure permissions to access the hotfix root folder later in this topic. • Settings on the server that hosts the hotfix root folder to check for appropriate permissions to access the folder and ensure the integrity of the data accessed from the SMB shared folder (SMB signing or SMB Encryption). For more information, see Restrict access to the hotfix root folder.
Additional options	• Optionally, configure the plug-in so that SMB Encryption is enforced when accessing data from the hotfix file share. In the CAU UI, on the **Additional Options** page, select the **Require SMB Encryption in accessing the hotfix root folder** option, or configure the **RequireSMBEncryption='True'** Windows PowerShell plug-in argument. **Important** • You must perform additional configuration steps on the SMB server to enable SMB data integrity with SMB signing or SMB Encryption. For more information, see Step 4 in Restrict access to the hotfix root folder. • If you select the option to enforce the use of SMB Encryption, and the

Item	Details
	hotfix root folder is not configured for access by using SMB Encryption, the Updating Run will fail. • Optionally, disable the default checks for sufficient permissions for the hotfix root folder and the hotfix configuration file. In the CAU UI, select **Disable check for administrator access to the hotfix root folder and configuration file**, or configure the **DisableAclChecks='True'** plug-in argument. • Optionally, configure the **HotfixInstallerTimeoutMinutes=<Integer>** argument to specify how long the hotfix plug-in waits for the hotfix installer process to return. (The default is 30 minutes.) For example, to specify a timeout period of two hours, set **HotfixInstallerTimeoutMinutes=120**. • Optionally, configure the **HotfixConfigFileName = <name>** plug-in argument to specify a name for the hotfix configuration file that is located in the hotfix root folder. If not specified, the default name DefaultHotfixConfig.xml is used.

Configure a hotfix root folder structure

For the hotfix plug-in to work, hotfixes must be stored in a well-defined structure in an SMB file share (hotfix root folder), and you must configure the hotfix plug-in with the path to the hotfix root folder by using the CAU UI or the CAU Windows PowerShell cmdlets. This path is passed to the plug-in as the **HotfixRootFolderPath** argument. You can choose one of several structures for the hotfix root folder, according to your updating needs, as shown in the following examples. Files or folders that do not adhere to the structure are ignored.

Example 1. Folder structure used to apply hotfixes to all cluster nodes

To specify that hotfixes apply to all cluster nodes, copy them to a folder named **CAUHotfix_All** under the hotfix root folder. In this example, the **HotfixRootFolderPath** plug-in argument is set to *\\MyFileServer\Hotfixes\Root*\. The **CAUHotfix_All** folder contains three updates with the extensions .msu, .msi, and .msp that will be applied to all cluster nodes. The update file names are only for illustration purposes.

Note

In this and the following examples, the hotfix configuration file with its default name DefaultHotfixConfig.xml is shown in its required location in the hotfix root folder.

```
\\MyFileServer\Hotfixes\Root\

   DefaultHotfixConfig.xml

   CAUHotfix_All\

      Update1.msu

      Update2.msi

      Update3.msp
```

```
...
```

Example 2. Folder structure used to apply certain updates only to a specific node

To specify hotfixes that apply only to a specific node, use a subfolder under the hotfix root folder with the name of the node. Use the NetBIOS name of the cluster node, for example, *ContosoNode1*. Then, move the updates that apply only to this node to this subfolder. In the following example, the **HotfixRootFolderPath** plug-in argument is set to *\\MyFileServer\Hotfixes\Root*. Updates in the **CAUHotfix_All** folder will be applied to all cluster nodes, and *Node1_Specific_Update.msu* will be applied only to *ContosoNode1*.

```
\\MyFileServer\Hotfixes\Root\

   DefaultHotfixConfig.xml

   CAUHotfix_All\

      Update1.msu

      Update2.msi

      Update3.msp

      ...

   ContosoNode1\

      Node1_Specific_Update.msu

      ...
```

Example 3. Folder structure used to apply updates other than .msu, .msi, and .msp files

By default, **Microsoft.HotfixPlugin** only applies updates with the .msu, .msi, or .msp extension. However, certain updates might have different extensions and require different installation commands. For example, you might need to apply a firmware update with the extension .exe to a node in a cluster. You can configure the hotfix root folder with a subfolder that indicates a specific, non-default update type should be installed. You must also configure a corresponding folder installation rule that specifies the installation command in the `<FolderRules>` element in the hotfix configuration XML file.

In the following example, the **HotfixRootFolderPath** plug-in argument is set to *\\MyFileServer\Hotfixes\Root*. Several updates will be applied to all cluster nodes, and a firmware update *SpecialHotfix1.exe* will be applied to *ContosoNode1* by using *FolderRule1*. For information about configuring *FolderRule1* in the hotfix configuration file, see Customize the hotfix configuration file later in this topic.

```
\\MyFileServer\Hotfixes\Root\

   DefaultHotfixConfig.xml

   CAUHotfix_All\

      Update1.msu

      Update2.msi
```

```
    Update3.msp
    ...

ContosoNode1\
    FolderRule1\
        SpecialHotfix1.exe
    ...
```

Customize the hotfix configuration file

The hotfix configuration file controls how **Microsoft.HotfixPlugin** installs specific hotfix file types in a failover cluster. The XML schema for the configuration file is defined in HotfixConfigSchema.xsd, which is located in the following folder on a computer where the CAU tools are installed:

%systemroot%\System32\WindowsPowerShell\v1.0\Modules\ClusterAwareUpdating folder

To customize the hotfix configuration file, copy the sample configuration file DefaultHotfixConfig.xml from this location to the hotfix root folder and make appropriate modifications for your scenario.

Important
To apply most hotfixes provided by Microsoft and other updates, the default hotfix configuration file can be used without modification. Customization of the hotfix configuration file is a task only in advanced usage scenarios.

By default, the hotfix configuration XML file defines installation rules and exit conditions for the following two categories of hotfixes:

- Hotfix files with extensions that the plug-in can install by default (.msu, .msi, and .msp files). These are defined as `<ExtensionRules>` elements in the `<DefaultRules>` element. There is one `<Extension>` element for each of the default supported file types. The general XML structure is as follows:

```
<DefaultRules>
    <ExtensionRules>
      <Extension name="MSI">
        <!-- Template and ExitConditions elements for
installation of .msi files follow -->
          ...
      </Extension>
      <Extension name="MSU">
```

```
        <!-- Template and ExitConditions elements for
installation of .msu files follow -->
          ...
      </Extension>
      <Extension name="MSP">
        <!-- Template and ExitConditions elements for
installation of .msp files follow -->
          ...
      </Extension>
          ...
   </ExtensionRules>
</DefaultRules>
```

If you need to apply certain update types to all cluster nodes in your environment, you can define additional `<Extension>` elements.

- Hotfix or other update files that are not .msi, .msu, or .msp files, for example, non-Microsoft drivers, firmware, and BIOS updates.

  Each non-default file type is configured as a `<Folder>` element in the `<FolderRules>` element. The name attribute of the `<Folder>` element must be identical to the name of a folder in the hotfix root folder that will contain updates of the corresponding type. The folder can be in the **CAUHotfix_All** folder or in a node-specific folder. For example, if *FolderRule1* is configured in the hotfix root folder, configure the following element in the XML file to define an installation template and exit conditions for the updates in that folder:

```
<FolderRules>
      <Folder name="FolderRule1">
        <!-- Template and ExitConditions elements for
installation of updates in FolderRule1 follow -->
          ...
      </Folder>
      ...
</FolderRules>
```

The following tables describe the `<Template>` attributes and the possible `<ExitConditions>` subelements.

<Template> attribute	Description
path	The full path to the installation program for the file type that is defined in the <Extension name> attribute. To specify the path to an update file in the hotfix root folder structure, use $update$.
parameters	A string of required and optional parameters for the program that is specified in path. To specify a parameter that is the path to an update file in the hotfix root folder structure, use $update$.

<ExitConditions> subelement	Description
<Success>	Defines one or more exit codes that indicate the specified update succeeded. This is a required subelement.
<Success_RebootRequired>	Optionally defines one or more exit codes that indicate the specified update succeeded and the node must restart. **Note** Optionally, the <Folder> element can contain the alwaysReboot attribute. If this attribute is set, it indicates that if a hotfix installed by this rule returns one of the exit codes that is defined in <Success>, it is interpreted as a <Success_RebootRequired> exit condition.
<Fail_RebootRequired>	Optionally defines one or more exit codes that indicate the specified update failed and the node must restart.
<AlreadyInstalled>	Optionally defines one or more exit codes that indicate the specified update was not applied because it is already installed.
<NotApplicable>	Optionally defines one or more exit codes that indicate the specified update was not applied because it does not apply to the cluster node.

Important
Any exit code that is not explicitly defined in `<ExitConditions>` is interpreted as the update failed, and the node does not restart.

Restrict access to the hotfix root folder

You must perform several steps to configure the SMB file server and file share to help secure the hotfix root folder files and hofix configuration file for access only in the context of **Microsoft.HotfixPlugin**. These steps enable several features that help prevent possible tampering with the hotfix files in a way that might compromise the failover cluster.

The general steps are as follows:

1. Identify the user account that is used for Updating Runs by using the plug-in
2. Configure this user account in the necessary groups on an SMB file server
3. Configure permissions to access the hotfix root folder
4. Configure settings for SMB data integrity
5. Enable a Windows Firewall rule on the SMB server

Step 1. Identify the user account that is used for Updating Runs by using the hotfix plug-in

The account that is used in CAU to check security settings while performing an Updating Run using **Microsoft.HotfixPlugin** depends on whether CAU is used in remote-updating mode or self-updating mode, as follows:

- **Remote-updating mode** The account that has administrative privileges on the cluster to preview and apply updates.
- **Self-updating mode** The name of the virtual computer object that is configured in Active Directory for the CAU clustered role. This is either the name of a prestaged virtual computer object in Active Directory for the CAU clustered role or the name that is generated by CAU for the clustered role. To obtain the name if it is generated by CAU, run the **Get-CauClusterRole** CAU Windows PowerShell cmdlet. In the output, **ResourceGroupName** is the name of the generated virtual computer object account.

Step 2. Configure this user account in the necessary groups on an SMB file server

Important
You must add the account that is used for Updating Runs as a local administrator account on the SMB server. If this is not permitted because of the security policies in your organization, configure this account with the necessary permissions on the SMB server by using the following procedure.

To configure a user account on the SMB server

1. Add the account that is used for Updating Runs to the Distributed COM Users group and to one of the following groups: Power User, Server Operation, or Print Operator.
2. To enable the necessary WMI permissions for the account, start the WMI Management Console on the SMB server. Start Windows PowerShell and then type the following

command:

```
wmimgmt.msc
```

3. In the console tree, right-click **WMI Control (Local)**, and then click **Properties**.
4. Click **Security**, and then expand **Root**.
5. Click **CIMV2**, and then click **Security**.
6. Add the account that is used for Updating Runs to the **Group or user names** list.
7. Grant the **Execute Methods** and **Remote Enable** permissions to the account that is used for Updating Runs.

Step 3. Configure permissions to access the hotfix root folder

By default, when you attempt to apply updates, the hotfix plug-in checks the configuration of the NTFS file system permissions for access to the hotfix root folder. If the folder access permissions are not configured properly, an Updating Run using the hotfix plug-in might fail.

If you use the default configuration of the hotfix plug-in, ensure that the folder access permissions meet the following requirements.

- The Users group has Read permission.
- If the plug-in will apply updates with the .exe extension, the Users group has Execute permission.
- Only certain security principals are permitted (but are not required) to have Write or Modify permission. The allowed principals are the local Administrators group, SYSTEM, CREATOR OWNER, and TrustedInstaller. Other accounts or groups are not permitted to have Write or Modify permission on the hotfix root folder.

Optionally, you can disable the preceding checks that the plug-in performs by default. You can do this in one of two ways:

- If you are using the CAU Windows PowerShell cmdlets, configure the **DisableAclChecks='True'** argument in the **CauPluginArguments** parameter for the hotfix plug-in.
- If you are using the CAU UI, select the **Disable check for administrator access to the hotfix root folder and configuration file** option on the **Additional Update Options** page of the wizard that is used to configure Updating Run options.

However, as a best practice in many environments, we recommend that you use the default configuration to enforce these checks.

Step 4. Configure settings for SMB data integrity

To check for data integrity in the connections between the cluster nodes and the SMB file share, the hotfix plug-in requires that you enable settings on the SMB file share for SMB signing or SMB Encryption. SMB Encryption, which provides enhanced security and better performance in many environments, is supported starting in Windows Server 2012. You can enable either or both of these settings, as follows:

- To enable SMB signing, see the procedure in the article 887429 in the Microsoft Knowledge Base.

- To enable SMB Encryption for the SMB shared folder, run the following Windows PowerShell cmdlet on the SMB server:

  ```
  Set-SmbShare <ShareName> -EncryptData $true
  ```

  Where <*ShareName*> is the name of the SMB shared folder.

Optionally, to enforce the use of SMB Encryption in the connections to the SMB server, select the **Require SMB Encryption in accessing the hotfix root folder** option in the CAU UI, or configure the **RequireSMBEncryption='True'** plug-in argument by using the CAU Windows PowerShell cmdlets.

Important
If you select the option to enforce the use of SMB Encryption, and the hotfix root folder is not configured for connections that use SMB Encryption, the Updating Run will fail.

Step 5. Enable a Windows Firewall rule on the SMB server

You must enable the **File Server Remote Management (SMB-in)** rule in Windows Firewall on the SMB file server. This is enabled by default in Windows Server 2012 R2 and Windows Server 2012.

See also

- Cluster-Aware Updating Overview
- Cluster-Aware Updating Windows PowerShell Cmdlets
- Cluster-Aware Updating Plug-in Reference

Advanced Options and Updating Run Profiles for CAU

This topic describes Updating Run options that can be configured for a CAU Updating Run. These advanced options can be configured when you use either the CAU UI or the CAU Windows PowerShell cmdlets to apply updates or to configure self-updating options.

Most configuration settings can be saved as an XML file called an Updating Run Profile and reused for later Updating Runs. The default values for the Updating Run options that are provided by CAU can also be used in many cluster environments.

For information about additional options that you can specify for each Updating Run and about Updating Run Profiles, see the following sections later in this topic:

- Options that you specify when you request an Updating Run
- Use Updating Run Profiles

Options that can be set in an Updating Run Profile

The following table lists options that you can set in a CAU Updating Run Profile.

Important
To set the **PreUpdateScript** or **PostUpdateScript** option, ensure that Windows PowerShell 3.0 or 4.0 and .NET Framework 4.5 are installed and that Windows PowerShell remoting is enabled on each node in the cluster. For more information, see Configure nodes for remote management in Requirements and Best Practices for Cluster-Aware Updating.

Option	Default value	Details
StopAfter	Unlimited time	Time in minutes after which the Updating Run will be stopped if it has not completed. **Note** If you specify a pre-update or a post-update Windows PowerShell script, the entire process of running scripts and performing updates must be complete within the **StopAfter** time limit.
WarnAfter	By default, no warning appears	Time in minutes after which a warning will appear if the Updating Run (including a pre-update script and a post-update script, if they are configured) has not completed.
MaxRetriesPerNode	3	Maximum number of times that the update process (including a pre-update script and a post-update script, if they are configured) will be retried per node. The maximum is 64.
MaxFailedNodes	For most clusters, an integer that is approximately one-third of the number of cluster nodes	Maximum number of nodes on which updating can fail, either because the nodes fail or the Cluster service stops running. If one more node fails, the Updating

Option	Default value	Details
		Run is stopped. The valid range of values is 0 to 1 less than the number of cluster nodes.
RequireAllNodesOnline	None	Specifies that all nodes must be online and reachable before updating begins.
RebootTimeoutMinutes	15	Time in minutes that CAU will allow for restarting a node (if a restart is necessary) and starting all auto-start services. If the restart process does not complete within this time, the Updating Run on that node is marked as failed.
PreUpdateScript	None	The path and file name for a Windows PowerShell script to run on each node before updating begins, and before the node is put into maintenance mode. The file name extension must be **.ps1**, and the total length of the path plus file name must not exceed 260 characters. As a best practice, the script should be located on a disk in cluster storage, or at a highly available network file share, to ensure that it is always accessible to all of the cluster nodes. If the script is located on a network file share, ensure that you configure the file share for Read permission for the Everyone group, and restrict write access to prevent tampering with the files by unauthorized users. If you specify a pre-update script, be sure that settings such as the time limits (for example, **StopAfter**) are configured to allow the script to run successfully.

Option	Default value	Details
		These limits span the entire process of running scripts and installing updates, not just the process of installing updates.
PostUpdateScript	None	The path and file name for a Windows PowerShell script to run after updating completes (after the node leaves maintenance mode). The file name extension must be **.ps1** and the total length of the path plus file name must not exceed 260 characters. As a best practice, the script should be located on a disk in cluster storage, or at a highly available network file share, to ensure that it is always accessible to all of the cluster nodes. If the script is located on a network file share, ensure that you configure the file share for Read permission for the Everyone group, and restrict write access to prevent tampering with the files by unauthorized users. If you specify a post-update script, be sure that settings such as the time limits (for example, **StopAfter**) are configured to allow the script to run successfully. These limits span the entire process of running scripts and installing updates, not just the process of installing updates.
ConfigurationName	This setting only has an effect if you run scripts. If you specify a pre-update script or a post-update script, but you do not specify a **ConfigurationName**, the default session configuration for Windows	Specifies the Windows PowerShell session configuration that defines the session in which scripts (specified by **PreUpdateScript** and **PostUpdateScript**) are run, and can limit the commands that can

Option	Default value	Details
	PowerShell (Microsoft.PowerShell) is used.	be run.
CauPluginName	**Microsoft.WindowsUpdatePlugin**	Plug-in that you configure Cluster-Aware Updating to use to preview updates or perform an Updating Run. For more information, see How CAU Plug-ins Work.
CauPluginArguments	None	A set of *name=value* pairs (arguments) for the updating plug-in to use, for example: **Domain=Domain.local** These *name=value* pairs must be meaningful to the plug-in that you specify in **CauPluginName**. To specify an argument using the CAU UI, type the *name*, press the Tab key, and then type the corresponding *value*. Press the Tab key again to provide the next argument. Each *name* and *value* are automatically separated with an equal (=) sign. Multiple pairs are automatically separated with semicolons. For the default **Microsoft.WindowsUpdatePlugin** plug-in, no arguments are needed. However, you can specify an optional argument, for example to specify a standard Windows Update Agent query string to filter the set of updates that are applied by the plug-in. For a *name*, use **QueryString**, and for a *value*, enclose the full query in quotation marks. For more information, see How CAU Plug-ins Work.

Options that you specify when you request an Updating Run

The following table lists options (other than those in an Updating Run Profile) that you can specify when you request an Updating Run. For information about options that you can set in an Updating Run Profile, see the preceding table.

Option	Default value	Details
ClusterName	None **Note** This option must be set only when the CAU UI is not run on a failover cluster node, or you want to reference a failover cluster different from where the CAU UI is run.	NetBIOS name of the cluster on which to perform the Updating Run.
Credential	Current account credentials	Administrative credentials for the target cluster on which the Updating Run will be performed. You may already have the necessary credentials if you start the CAU UI (or open a Windows PowerShell session, if you are using the CAU Windows PowerShell cmdlets) from an account that has administrator rights and permissions on the cluster.
NodeOrder	By default, CAU starts with the node that owns the smallest number of clustered roles, then progresses to the node that has the second smallest number, and so on.	Names of the cluster nodes in the order that they should be updated (if possible).

Use Updating Run Profiles

Each Updating Run can be associated with a specific Updating Run Profile. The default Updating Run Profile is stored in the %windir%\cluster folder. If you are using the CAU UI in remote-updating mode, you can specify an Updating Run Profile at the time that you apply updates, or you can use the default Updating Run profile. If you are using CAU in self-updating mode, you can import the settings from a specified Updating Run Profile when you configure the self-updating options. In both cases, you can override the displayed values for the Updating Run options according to your needs. If you want, you can save the Updating Run options as an Updating Run Profile with the same file name or a different file name. The next time that you apply updates or configure self-updating options, CAU automatically selects the Updating Run Profile that was previously selected.

An existing Updating Run Profile can also be modified, or a new Updating Run Profile can be created, by selecting **Create or modify Updating Run Profile** in the CAU UI.

Windows PowerShell equivalent commands

You can import the settings from an Updating Run Profile when you run the **Invoke-CauRun**, **Add-CauClusterRole**, or **Set-CauClusterRole** cmdlet.

The following example performs a scan and a full Updating Run on the cluster named *CONTOSO-FC1*, using the Updating Run options that are specified in *C:\Windows\Cluster\DefaultParameters.xml*. Default values are used for the remaining cmdlet parameters.

```
$MyRunProfile = Import-Clixml C:\Windows\Cluster\DefaultParameters.xml

Invoke-CauRun –ClusterName CONTOSO-FC1 @MyRunProfile
```

By using an Updating Run Profile, you can update a failover cluster in a repeatable fashion with consistent settings for exception management, time bounds, and other operational parameters. Because these settings are typically specific to a class of failover clusters—such as “All Microsoft SQL Server clusters”, or “My business-critical clusters”—you might want to name each Updating Run Profile according to the class of Failover Clusters it will be used with. In addition, you might want to manage the Updating Run Profile on a file share that is accessible to all of the failover clusters of a specific class in your IT organization.

Notes

- An Updating Run Profile does not store cluster-specific information such as administrative credentials. If you are using CAU in self-updating mode, the Updating Run Profile also does not store the self-updating schedule information. This makes it possible to share an Updating Run Profile across all failover clusters in a specified class.
- If you configure self-updating options using an Updating Run Profile and later modify the profile with different values for the Updating Run options, the self-updating configuration does not change automatically. To apply the new Updating Run settings, you must configure the self-updating options again.

See also

- Cluster-Aware Updating Overview
- **CAU cmdlets**

File and Storage Services Overview

This topic discusses the File and Storage Services server role in Windows Server 2012 R2 and Windows Server 2012, including what's new, a list of role services, and where to find evaluation and deployment information.

Did you mean…

- File Services in Windows Server 2008 R2
- File Services Windows Server Protocols
- Windows Storage Server 2008 R2
- Step 4: Configure iSCSI storage provider identity credentials

Role description

File and Storage Services includes technologies that help you set up and manage one or more file servers, which are servers that provide central locations on your network where you can store files and share them with users. If your users need access to the same files and applications, or if centralized backup and file management are important to your organization, you should set up one or more servers as a file server by installing the File and Storage Services role and the appropriate role services.

The File and Storage Services role and the Storage Services role service are installed by default, but without any additional role services. This basic functionality enables you to use Server Manager or Windows PowerShell to manage the storage functionality of your servers. However, to set up or manage a file server, you should use the Add Roles and Features Wizard in Server Manager or the `Install-WindowsFeature` Windows PowerShell cmdlet to install additional File and Storage Services role services, such as the role services discussed in this topic.

Practical applications

Administrators can use the File and Storage Services role to set up and manage multiple file servers and their storage capabilities by using Server Manager or Windows PowerShell. Some of the specific applications include the following:

- **Work Folders** Use to enable users to store and access work files on personal PCs and devices, in addition to corporate PCs. Users gain a convenient location to store work files and access them from anywhere. Organizations maintain control over corporate data by storing the files on centrally managed file servers and optionally specifying user device policies (such as encryption and lock screen passwords). Work Folders is a new role service in Windows Server 2012 R2.

- **Data Deduplication** Use to reduce the disk space requirements of your files, saving money on storage.
- **iSCSI Target Server** Use to create centralized, software-based, and hardware-independent iSCSI disk subsystems in storage area networks (SANs).
- **Storage Spaces** Use to deploy high availability storage that is resilient and scalable by using cost-effective industry-standard disks.
- **Server Manager** Use to remotely manage multiple file servers from a single window.
- **Windows PowerShell** Use to automate the management of the majority of administration tasks for file servers.

New and changed functionality

The following table describes some of the major changes in File and Storage Services functionality that is available in Windows Server 2012 R2.

Feature/functionality	New or updated?	Description
Work Folders	New	Provides a consistent way for users to access their work files from their personal computers and devices. See Work Folders Overview for more information.
Server Message Block	Updated	Enhancements include automatic rebalancing of Scale-Out File Server clients, improved performance of SMB Direct, and improved SMB event messages. See What's New in SMB in Windows Server 2012 R2 for more information.
Storage Spaces	Updated	Enhancements include SSD and HDD storage tiers, an SSD-based write-back cache, parity space support for failover clusters, dual parity support, and greatly decreased storage space rebuild times. See What's New in Storage Spaces in Windows Server 2012 R2 for more information.
DFS Replication	Updated	Enhancements include

Feature/functionality	New or updated?	Description
		database cloning for large performance gains during initial sync, a Windows PowerShell module for DFS Replication, a new DFS Replication WMI provider, faster replication on high bandwidth connections, conflict and preexisting data recovery, and support for rebuilding corrupt databases without unexpected data loss. See What's New in DFS Replication in Windows Server 2012 R2 for more information.
iSCSI Target Server	Updated	Updates include virtual disk enhancements, manageability enhancements in a hosted or private cloud, and improved optimization to allow disk-level caching. See What's New in iSCSI Target Server in Windows Server 2012 R2 for more information.

The following table describes some of the major changes in File and Storage Services functionality that are available in Windows Server 2012.

Feature/functionality	New or updated?	Description
Data Deduplication	New	Saves disk space by storing a single copy of identical data on the volume.
iSCSI Target Server	New	Provides block storage to other servers and applications on the network by using the Internet SCSI (iSCSI) standard.
Storage Spaces and storage pools	New	Enables you to virtualize storage by grouping industry-standard disks into storage

Feature/functionality	New or updated?	Description
		pools and then creating storage spaces from the available capacity in the storage pools.
Unified remote management of File and Storage Services in Server Manager	New	Enables you to remotely manage multiple file servers, including their role services and storage, from a single window.
Windows PowerShell cmdlets for File and Storage Services	New	Provides Windows PowerShell cmdlets to perform the majority of administration tasks for file and storage servers.

For more information about what else is new in File and Storage Services and related technologies, see the following topics.

- DFS Namespaces and DFS Replication Overview
- Encrypted Hard Drive
- NTFS Health and Chkdsk
- What's New in File Server Resource Manager in Windows Server 2012
- Server Message Block overview
- Thin Provisioning and Trim Storage Overview
- Windows Offloaded Data Transfers Overview

Data Deduplication

By using the Data Deduplication role service to reduce the number of duplicate blocks of data in storage, you can store much more data in a given amount of storage capacity than was possible in previous releases that used Single Instance Storage (SIS) or NTFS file system compression. General purpose file servers can typically reduce storage capacity utilization by a 2:1 ratio (for example, files that previously used 1 TB would use 500 GB after data deduplication). Servers that host virtualization data (such as VHD files) often reduce storage capacity utilization by a 20:1 ratio, which reduces 1 TB of data to 50 GB.

Data deduplication is highly scalable, resource efficient, and nonintrusive. It can run on dozens of large volumes of primary data simultaneously without affecting other workloads on the server. Low impact on the server workloads is maintained by throttling the CPU and memory resources that are consumed. By using data deduplication jobs, you can schedule when data deduplication should run, specify the resources to duplicate, and tune file selection. Data integrity and recoverability are maximized through checksum and other algorithms by using selective redundancy.

When combined with BranchCache, the same optimization techniques are applied to data that is transferred over the WAN to a branch office. This results in faster file download times and reduced bandwidth consumption.

What value does this change add?

Data deduplication uses variable-size chunking and compression, which together deliver storage optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data.

What works differently?

Windows Server 2012 includes Data Deduplication as a role service that you can install and manage by using Server Manager or Windows PowerShell cmdlets. The default settings can quickly reduce the amount of storage capacity that is used by your data. You can fine-tune the settings to see more gains or use Windows PowerShell cmdlets to create scripts that will trigger storage optimization when and where you want it.

For more detailed information, see Data Deduplication Overview.

iSCSI Target Server

iSCSI Target Server provides block storage to other servers and applications on the network by using the Internet SCSI (iSCSI) standard. When combined with other continuously available technologies in Windows Server 2012, iSCSI Target Server provides continuously available storage that was previously available only on expensive, high-end storage area network (SAN) devices.

What value does this change add?

- **Network or diskless boot** By using boot-capable network adapters or a software loader, you can deploy hundreds of diskless servers. With the iSCSI Target Server, the deployment is fast (in Microsoft testing, 256 computers deployed in 34 minutes). By using differencing virtual hard disks, you can save up to 90% of the storage space for the operating system images. This is ideal for large deployments of identical operating system images, such as a server farm that is running Hyper-V or high-performance computing (HPC) clusters.
- **Server application storage** Some applications (for example, Hyper-V and Exchange Server) require block storage, which is raw storage that appears to applications as an unformatted disk, which is ready for direct management by the application. iSCSI Target Server can provide these applications with continuously available block storage. Because the storage is remotely accessible, iSCSI Target Server can also consolidate block storage for central or branch office locations.
- **Heterogeneous storage** iSCSI Target Server supports iSCSI initiators on operating systems other than Windows, making it easy to share storage in a heterogeneous environment.
- **Dev, test, and demo lab environments** With iSCSI Target Server, any computer running Windows Server 2012 can be a network-accessible block storage device. This is useful for testing applications prior to deployment on a SAN storage device.

What works differently?

In Windows Server 2012, iSCSI Target Server management functionality moves from a separate download to part of the Windows Server operating system. You can use Server Manager or Windows PowerShell cmdlets to set up, configure, and manage iSCSI Target Server. Windows Server 2012 also includes changes to the clustering resource model that improve scalability so that more initiators can connect to the target servers.

For more detailed information, see iSCSI Target Block Storage Overview.

Storage Spaces and storage pools

Storage Spaces is a storage subsystem that is included in Windows that enables you to group industry-standard disks (such as Serial ATA or Serial Attached SCSI disks) into one or more storage pools, and then create virtual disks known as “storage spaces” from the available capacity in the storage pools. Storage Spaces provides resilient storage virtualization capabilities for business-critical virtual or physical deployments, including deployments on scalable multinode servers.

After you group physical disks into storage pools, you can create virtual disks from the available capacity without individually managing each physical disk. This aggregation of disks enables you to make more efficient use of disk capacity, easily add storage without impacting users, and delegate administration of storage. In Windows Server, you can use storage pools with Storage Spaces or with non-Microsoft storage subsystems, including subsystems that use the SMI-S standard.

What value does this change add?

Storage Spaces and storage pools reduce administration costs by decreasing the time administrators spend provisioning storage. They also simplify administration tasks, enabling administrators who are not storage professionals to set up and manage resilient and high availability storage. Storage Spaces also saves hardware costs by using industry-standard disks for resilient storage with high availability.

What works differently?

With storage pools, instead of managing each disk individually, you add physical disks to one or more pools and then create virtual disks from the available capacity. You then create volumes on the virtual disks, as if they were physical disks. When you run low on the available capacity in the pool, simply create more virtual disks.

Virtual disks also provide thinly provisioning, which means that pool capacity is used only according to the size of the files that you copy to the virtual disks, not the size of the virtual disk you create. For example, a 10 TB virtual disk might consume 100 MB of pool capacity if you haven’t copied many files to the disk yet.

For more information, see Storage Spaces Overview and Storage Management Overview.

Unified remote management of File and Storage Services in Server Manager

The File and Storage Services role in Server Manager enables you to remotely manage multiple file servers from a single window in Windows Server 2012, including role services and storage. The File and Storage Services page in Server Manager provides the following sections to manage all servers running Windows Server 2012 that have been added to Server Manager:

- **Servers** Manage basic server functionality on servers running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. You can use the **Servers** page to perform such tasks as restarting servers and launching administrative tools.
- **Storage Pools** Manage storage pools, including the physical disks that make up the pools and virtual disks that are created from available capacity in the pools.
- **Volumes** Manage volumes, including scanning for file system errors, extending volumes, and configuring Data Deduplication.
- **Shares** Manage SMB and NFS shares, including creating new shares and setting up quotas.
- **iSCSI Virtual Disks** Manage iSCSI virtual disks, including creating new iSCSI virtual disks and targets.

What value does this change add?

Managing multiple file servers and file server technologies from a single Server Manager window enables administrators to work more efficiently and get a better overall view of the servers they manage, so managing multiple servers becomes as easy as managing one.

What works differently?

Prior to Windows Server 2012, managing multiple file servers meant using Remote Desktop to connect to each server, or opening multiple instances of the administration console (one per server). In Windows Server 2012, you can use Server Manager to perform many of the functions that are provided by the following standalone management consoles:

- Disk Management
- Share and Storage Management
- File Server Resource Manager (quotas and assigning folder management properties to file shares)
- Microsoft iSCSI Software Target (not available in Windows Server 2012)
- Storage Manager for SANs (not available in Windows Server 2012)

In addition to the integrated functionality, you can use the **Tools** menu in Server Manager to launch administrative tools on any of the managed servers, including DFS Management, File Server Resource Manager, and Services for Network File System (NFS).

Note

Although the File and Storage Services role in Server Manager does not support fully managing servers running Windows Server 2008 R2 or Windows Server 2008, you can add these servers to Server Manager and use the **Servers**, and **All Servers** pages to view details about the servers and launch administrative tools. For more information, see

Managing Down-Level Windows-based Servers from Server Manager in Windows Server 2012.

Windows PowerShell cmdlets for File and Storage Services

Windows Server 2012 includes Windows PowerShell cmdlets to perform the majority of administration tasks for file and storage servers.

What value does this change add?

Extensive Windows PowerShell cmdlets allow administrators to automate common administration tasks by using Windows PowerShell scripts.

What works differently?

Instead of using a variety of snap-ins or disparate command-line utilities, administrators can manage their servers by using Windows PowerShell cmdlets and scripts. Windows Server 2012 R2 and Windows Server 2012 include Windows PowerShell cmdlets to manage the following file and storage technologies.

- BranchCache
- Data Deduplication
- DFS Namespaces
- DFS Replication (New in Windows Server 2012 R2)
- File Server Resource Manager
- iSCSI
- iSCSI Target
- MultiPath I/O (MPIO)
- Primary Computer
- Server for NFS
- SMB Shares
- SMB Witness
- SMI-S
- Storage
- Storage Spaces
- Sync Share (Work Folders) (New in Windows Server 2012 R2)

There are numerous related Windows PowerShell cmdlets that are also useful for file and storage workloads. For example, you can use the Best Practices Analyzer (BPA) cmdlets to compare servers to a known set of best practices for the File and Storage Services role.

For a reference sheet that contains some of the File and Storage Services cmdlets that are most commonly used, download the Windows PowerShell Reference Sheet for File and Storage Services.

Removed or deprecated functionality

The following features are included in Windows Server 2012, but they are being phased out, and they likely will be removed entirely from future versions of Windows Server.

Deprecated feature	Replacement
DFS Namespaces command-line tool: Dfscmd	Windows PowerShell cmdlets for DFS Namespaces
File Replication Service (FRS)	DFS Replication
File Server Resource Manager command-line tools: Dirquota, Filescrn, and Storrept	Windows PowerShell cmdlets for File Server Resource Manager
Share and Storage Management snap-in	File and Storage Services role in Server Manager
Shared Folders snap-in	File and Storage Services role in Server Manager
Virtual Disk Service (VDS) provider	Storage Management APIs and storage provider or the SMI-S standard and a compliant storage provider

The following features are not present in Windows Server 2012.

Deprecated feature	Replacement
Storage Manager for SANs snap-in	File and Storage Services role in Server Manager and Storage Management Windows PowerShell cmdlets
SCSIport host-bus adapter drivers	Storport drivers or a different host-bus adapter

For a list of deprecated features in the Windows Server 2012, see Features Removed or Deprecated in Windows Server 2012.

For more information about File Replication Service support in Windows Server 2012 and Windows Server 2008 R2, see File Replication Service (FRS) Is Deprecated in Windows Server 2008 R2.

Requirements for running File and Storage Services

There are no special hardware or software requirements for running File and Storage Services. However, File Server Resource Manager and DFS Replication support only volumes that are

formatted with the NTFS file system—the Resilient File System (ReFS) and the FAT file system are not supported.

The following sections explain common questions about File and Storage Services requirements.

How do I deploy and configure this role in a multiserver environment?

You can split the functionality of File and Storage Services across multiple servers by installing the role and appropriate role services on each appropriate server. You can then add all the servers to Server Manager for centralized management, or use Windows PowerShell scripts to simultaneously manage multiple servers.

Can I run this role on virtual machines?

Yes. You can run File and Storage Services and all its role services on Hyper-V virtual machines.

Important
Using Hyper-V snapshots to restore a server running DFS Replication for replication of anything other than the SYSVOL folder causes DFS Replication to fail, which requires special database recovery steps. For more information, see article 2517913 in the Microsoft Knowledge Base.

Can I run this role in a clustered environment?

Yes. You can run all File and Storage Services role services in a clustered environment. However, DFS Replication does not support replicating content that is stored in Cluster Shared Volumes.

Are there considerations for managing this role remotely?

You can manage File and Storage Services remotely by using the following tools:

- Server Manager
- Windows PowerShell cmdlets
- DFS Management
- File Server Resource Manager
- Services for Network File System
- DFS command-line utilities

No special considerations are necessary.

Are there considerations for managing the role on the Server Core installation option?

You can install and run File and Storage Services on the Server Core installation option or the Minimal Server Interface of Windows Server 2012. Server Manager and the administrative

consoles do not work on the Server Core installation option, although they can be used with the Minimal Server Interface. They can also be used to remotely administer servers that use the Server Core option.

Server Manager information

The following role services can be installed with the File and Storage Services server role.

Note
Server Manager cannot be used to add roles and features to servers that are running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. For more information, see **Remote, Multiserver Management**. Also, Server Manager can receive only an online or offline status from servers that are running Windows Server 2003.

Role service	Description
File Services	Provides a grouping for role services that are related to file servers—it doesn't install anything.
File Server	Manages file shares and enables users to access files on this computer from the network. For more information about file shares that use the Server Message Block (SMB) protocol, see Server Message Block overview.
BranchCache for Network Files	Enables computers in branch offices to cache commonly downloaded files from file shares on which BranchCache is enabled, and then provide those files to other computers in the branch office. This reduces network bandwidth usage and provides faster access to the files. For more information about BranchCache, see BranchCache Overview.
Data Deduplication	Saves disk space by storing a single copy of identical data on the volume. For more information about Data Deduplication, see Data Deduplication Overview.
DFS Namespaces	Enables you to group file shares that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single file share with a series of subfolders. However, the underlying structure of the namespace can

Role service	Description
	consist of numerous file shares located on different servers in multiple sites. Because the underlying structure of file shares is hidden from users, a single folder in a DFS namespace can correspond to multiple file shares on multiple servers. This structure provides fault tolerance and the ability to automatically connect users to local file shares, when they are available, instead of routing users over wide area network (WAN) connections. For more information, see DFS Namespaces and DFS Replication Overview.
DFS Replication	Replicates data between multiple servers over limited bandwidth network connections and local area network connections. It is a multiple master replication engine that uses the Remote Differential Compression (RDC) protocol to update only the portions of files that have changed since the last replication. DFS Replication can be used in conjunction with DFS Namespaces or by itself. For more information, see DFS Namespaces and DFS Replication Overview.
File Server Resource Manager	Helps you manage and understand the files and folders on a file server by scheduling file management tasks and storage reports, classifying files and folders, configuring folder quotas, and defining file screening policies. For more information, see File Server Resource Manager Overview.
File Server VSS Agent Service	Enables you to create volume shadow copies of applications that store data files on this file server.
iSCSI Target Server	Provides management tools for iSCSI targets. For more information, see iSCSI Target Block Storage Overview.
Server for Network File System (NFS)	Enables this computer to share files with UNIX-based computers and other computers that use the network file system (NFS) protocol.

Role service	Description
Storage Services	Provides storage management functionality that is always installed, including Storage Spaces and storage pools.

See also

For related information, see the following resources.

Content type	References
Product evaluation	Data Deduplication Overview | DFS Namespaces and DFS Replication Overview | **High-Performance, Continuously Available File Share Storage for Server Applications Technical Preview** | iSCSI Target Block Storage Overview | Network File System Overview | Windows Offloaded Data Transfers Overview | Resilient File System Overview | Increasing Server, Storage, and Network Availability: Scenario Overview | Scenario: Get Insight into Your Data by Using Classification | Scenario: Implement Retention of Information on File Servers | Server Message Block overview | Introduction to SMI-S | Storage Management Overview | Storage Spaces Overview | Supporting Information Workers with Reliable File Services and Storage | Thin Provisioning and Trim Storage Overview
Deployment	Plan to Deploy Data Deduplication | Deploy Network File System | Scale-Out File Server for Application Data Overview | Understand and Troubleshoot Storage Spaces and Pools in Windows Server "8" Beta
Operations	iSCSI Target Block Storage, How To | Windows PowerShell Reference Sheet for File and Storage Services
Tools and settings	**N:BranchCache** | **N:Deduplication** | **N:iSCSI** | **N:IscsiTarget** | **N:MPIO** | **N:Storage** | Storage

Content type	References
	Management API Classes
Troubleshooting	Known Issues: File and Storage Services in Windows Server 2012
Community resources	The File Services and Storage TechNet Forum | Storage Team at Microsoft File Cabinet Blog | Ask the Directory Services Team Blog | Jose Barreto's Blog | Windows Storage Team Blog | DFS Replication Survival Guide (TechNet Wiki) | Windows Server Information Experience on Twitter
Related technologies	Folder Redirection, Offline Files, and Roaming User Profiles overview | Failover Clustering Overview

Data Deduplication Overview

This topic describes the data deduplication feature in Windows Server 2012 R2 and Windows Server 2012, and it explains practical applications for the feature.

Feature description

Data deduplication involves finding and removing duplication within data without compromising its fidelity or integrity. The goal is to store more data in less space by segmenting files into small variable-sized chunks (32–128 KB), identifying duplicate chunks, and maintaining a single copy of each chunk. Redundant copies of the chunk are replaced by a reference to the single copy. The chunks are compressed and then organized into special container files in the System Volume Information folder.

The result is an on-disk transformation of each file as shown in Figure 1. After deduplication, files are no longer stored as independent streams of data, and they are replaced with stubs that point to data blocks that are stored within a common chunk store. Because these files share blocks, those blocks are only stored once, which reduces the disk space needed to store all files. During file access, the correct blocks are transparently assembled to serve the data without calling the application or the user having any knowledge of the on-disk transformation to the file. This enables administrators to apply deduplication to files without having to worry about any change in behavior to the applications or impact to users who are accessing those files.

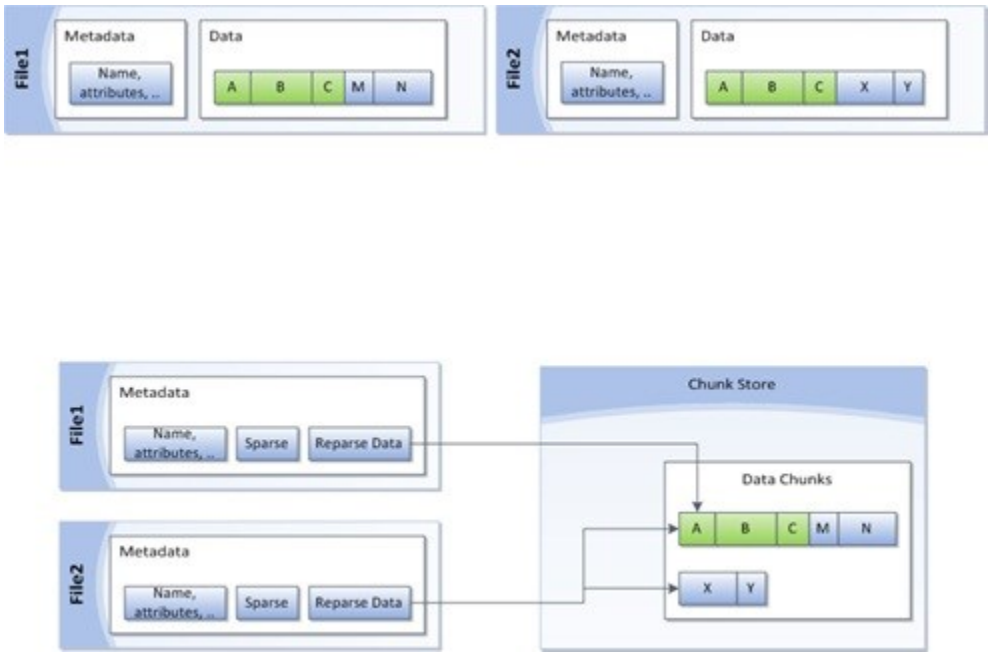

Figure 1 On-disk transformation of files during data deduplication

After a volume is enabled for deduplication and the data is optimized, the volume contains the following:

- **Unoptimized files**. For example, unoptimized files could include files that do not meet the selected file-age policy setting, system state files, alternate data streams, encrypted files, files with extended attributes, files smaller than 32 KB, other reparse point files, or files in use by other applications (the “in use” limit is removed in Windows Server 2012 R2).
- **Optimized files**. Files that are stored as reparse points that contain pointers to a map of the respective chunks in the chunk store that are needed to restore the file when it is requested.
- **Chunk store**. Location for the optimized file data.
- **Additional free space**. The optimized files and chunk store occupy much less space than they did prior to optimization.

Practical applications

To cope with data storage growth in the enterprise, administrators are consolidating servers and making capacity scaling and data optimization key goals. Data deduplication provides practical ways to achieve these goals, including:

- **Capacity optimization.** Data deduplication stores more data in less physical space. It achieves greater storage efficiency than was possible by using features such as Single Instance Storage (SIS) or NTFS compression. Data deduplication uses subfile variable-size chunking and compression, which deliver optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data.
- **Scale and performance.** Data deduplication is highly scalable, resource efficient, and nonintrusive. It can process up to 50 MB per second in Windows Server 2012 R2, and about

20 MB of data per second in Windows Server 2012. It can run on multiple volumes simultaneously without affecting other workloads on the server. Low impact on the server workloads is maintained by throttling the CPU and memory resources that are consumed. If the server gets very busy, deduplication can stop completely. In addition, administrators have the flexibility to run data deduplication jobs at any time, set schedules for when data deduplication should run, and establish file selection policies.

- **Reliability and data integrity.** When data deduplication is applied, the integrity of the data is maintained. Data Deduplication uses checksum, consistency, and identity validation to ensure data integrity. For all metadata and the most frequently referenced data, data deduplication maintains redundancy to ensure that the data is recoverable in the event of data corruption.
- **Bandwidth efficiency with BranchCache.** Through integration with BranchCache, the same optimization techniques are applied to data transferred over the WAN to a branch office. The result is faster file download times and reduced bandwidth consumption.
- **Optimization management with familiar tools.** Data deduplication has optimization functionality built into Server Manager and Windows PowerShell. Default settings can provide savings immediately, or administrators can fine-tune the settings to see more gains. One can easily use Windows PowerShell cmdlets to start an optimization job or schedule one to run in the future. Installing the Data Deduplication feature and enabling deduplication on selected volumes can also be accomplished by using an Unattend.xml file that calls a Windows PowerShell script and can be used with Sysprep to deploy deduplication when a system first boots.

New and changed functionality

The following table describes the changes in Data Deduplication functionality. For more information, see What's New in Data Deduplication in Windows Server 2012 R2.

Feature/Functionality	New or updated?	Description
Data deduplication for remote storage of Virtual Desktop Infrastructure (VDI) workloads	New in Windows Server 2012 R2	Optimize active virtual hard disks (VHDs) for Virtual Desktop Infrastructure (VDI) workloads by implementing Data Deduplication on Cluster Shared Volumes (CSVs).
Expand an optimized file on its original path.	New in Windows Server 2012 R2	Use the new **Expand-DedupFile** cmdlet in Windows PowerShell to expand optimized files on a specified path on the original path if needed for compatibility with applications, performance, or other requirements. For more information about the cmdlet, see

Feature/Functionality	New or updated?	Description
		T:Deduplication.Expand-DedupFile.

Requirements

To take advantage of data deduplication, the environment must meet the following requirements:

- Server: One computer or virtual machine running Windows Server 2012 R2 or Windows Server 2012 with at least one data volume
- (Optional) Another computer or virtual machine running Windows Server 2012 R2 or Windows Server 2012 that is connected to the server over a network

> **Important**
> If Data Deduplication is performed on VDI workloads, all VHD files must be stored on a file server running Windows Server 2012 R2, and the storage node and computation node must run on different servers.

Interoperability with Windows Azure virtual machines

You can run this Windows Server role service on a virtual machine in Windows Azure. This scenario has been tested with Windows Server 2012 R2. We recommend using Data Deduplication with Windows Azure virtual machines that have volumes with frequent reads but infrequent writes. In these circumstances, Data Deduplication can provide an efficient way to store more data in Azure VMs.

The following workloads can be good candidates for use with Data Deduplication on Azure VMs:

- General file servers with relatively static content
- Microsoft SharePoint sites with relatively static content
- Web sites with relatively static content

The following workloads aren't recommended for use with Data Deduplication on Azure VMs due to the frequency of changes to the large files used by the workloads:

- Messaging servers such as Microsoft Exchange Server
- Database servers such as Microsoft SQL Server

To learn about how to get started with Windows Azure virtual machines, visit the Windows Azure web site.

Architecture overview

The Data Deduplication feature consists of a filter driver that monitors local or remote I/O and a deduplication service that controls the three types of jobs that are available (Optimization, Garbage Collection, and Scrubbing).

Inherent in the deduplication architecture is resiliency during hardware failures—with full checksum validation on data and metadata, including redundancy for metadata and the most accessed data chunks.

Data Deduplication can potentially process all of the data on a selected volume (except a file size less than 32 KB, files in folders that are excluded, or files that have age settings applied). You should carefully determine if a server and attached volumes are suitable candidates for deduplication prior to enabling the feature. We strongly recommend that during deduplication, you regularly back up important data.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	• Extending Data Deduplication to new workloads in Windows Server 2012 R2
Deployment	• Deploying Data Deduplication for VDI storage in Windows Server 2012 R2 • Plan to Deploy Data Deduplication • Backup and Restore Considerations for Deduplicated Volumes • Data Deduplication Interoperability • Install and Configure Data Deduplication
Operations	• Monitor and Report for Data Deduplication
Community resources	• The File Services and Storage TechNet Forum • Storage Team at Microsoft File Cabinet Blog • Ask the Directory Services Team Blog • Jose Barreto's Blog • Windows Server Information Experience on Twitter
Related technologies	• File and Storage Services Overview • Failover Clustering Overview • Storage Spaces Overview

What's New in Data Deduplication in Windows Server 2012 R2

This topic describes features that were added to the Data Deduplication feature in Windows Server 2012 R2, including support for the optimization of live virtual hard disks (VHDs) for Virtual Desktop Infrastructure (VDI) workloads.

Feature description

Introduced in Windows Server 2012, Data Deduplication involves finding and removing duplication within data without compromising its fidelity or integrity. The goal is to store more data in less space by segmenting files into small variable-sized chunks (32–128 KB), identifying duplicate chunks, and maintaining a single copy of each chunk. Redundant copies of the chunk are replaced by a reference to the single copy. The chunks are compressed and then organized into special container files in the System Volume Information folder.

Note
For more information, see Data Deduplication Overview. For recommended usage scenarios, see Plan to Deploy Data Deduplication.

New and changed functionality

The following table describes the changes in Data Deduplication functionality in Windows Server 2012 R2.

Functionality	New or updated?	Description
Data deduplication for remote storage of VDI workloads	New	Optimize active VHDs for VDI workloads by implementing Data Deduplication on Cluster Shared Volumes (CSVs).
Expand an optimized file on its original path.	New	Use the **Expand-DedupFile** Windows PowerShell cmdlet to expand optimized files on a specified path if needed for compatibility with applications, performance, or other requirements.

Data deduplication for remote storage of VDIworkloads

In Windows Server 2012 R2, Data Deduplication can be installed on a scale-out file server and used to optimize live VHDs for VDI workloads.

What value does this change add?

By optimizing CSV volumes for your VDI workloads, you can stretch the virtual machine capacity of your existing storage subsystem. Storage savings as great as 95 percent can be achieved by implementing Data Deduplication on live VHDs for VDI deployments.

Important

In Windows Server 2012 R2, the performance of VHDs that are optimized through Data Deduplication is fully tested and supported only on VDI workloads. The same performance gains are not guaranteed for non-VDI workloads on virtual machines running Hyper-V; nor does Microsoft offer support for these scenarios in Windows Server 2012 R2.

The space savings from data deduplication make it feasible to deploy solid-state drives (SSDs) (with their improved I/O) for VDI, and to simplify supporting infrastructure such as “just a bunch of disks” (JBOD) enclosures, cooling, and power.

By consolidating files, data deduplication can improve caching efficiency, and as a result, I/O, on the storage subsystem for some types of operation.

For more information about the benefits of using data deduplication with VDI workloads, see Extending Data Deduplication to new workloads in Windows Server 2012 R2.

What works differently?

This feature is implemented through specifying **HyperV** with the new **UsageType** parameter for the **Enable-DedupVolume** Windows PowerShell cmdlet. This command can perform optimization on active VHD files. To enable the use of scale-out file servers for VDI workloads, data deduplication of CSVs is supported.

Data deduplication in Windows Server 2012 R2 implements Improvements in write efficiency and faster optimization speeds , which makes optimization on active VHD files feasible. However, when Data Deduplication involves virtualization, the computer that enables data deduplication cannot be the same server that is running Hyper-V. This ensures that optimization does not compete with the virtual machines for resources on the Hyper-V management operating system. For more information, see **T:Deduplication.Enable-DedupVolume**.

Your Hyper-V and VDI infrastructure can remain the same, with one exception: all VHD files for the virtual machines must be stored on a file server running Windows Server 2012 R2. The storage on that file server can be provided by directly attached disks (such as JBOD enclosures used with Storage Spaces), or it can be provided by a storage area network (SAN) or an iSCSI storage device. To help ensure that the storage is highly available, we recommend that you use a clustered file server with CSVs that provide storage for the VHDs. For procedures that describe how to set up Data Deduplication for a VHD workload, see Deploying Data Deduplication for VDI storage in Windows Server 2012 R2.

Expand an optimized file on its original path

The new **Expand-DataDedupFile** Windows PowerShell cmdlet enables you to expand optimized files on a specified path if needed for application compatibility, performance, or other requirements. The files are expanded on the original path. For more information, see **T:Deduplication.Expand-DedupFile**.

What value does this change add?

This gives you a way to expand individual files within an optimized volume if for any reason the optimized files are resulting in compatibility or performance issues.

What works differently?

This capability is new in Windows Server 2012 R2.

See also

The following table provides additional resources for evaluating the Data Deduplication feature in Windows Server 2012 R2.

Content type	References
Product evaluation	• Extending Data Deduplication to new workloads in Windows Server 2012 R2
Deployment	• Deploying Data Deduplication for VDI storage in Windows Server 2012 R2 • Plan to Deploy Data Deduplication • Backup and Restore Considerations for Deduplicated Volumes • Data Deduplication Interoperability • Install and Configure Data Deduplication
Operations	• Monitor and Report for Data Deduplication
Tools and settings	• Deduplication Cmdlets in Windows PowerShell
Community resources	• The File Services and Storage TechNet Forum • Storage Team at Microsoft File Cabinet Blog • Ask the Directory Services Team Blog • Jose Barreto's Blog • Windows Server Information Experience on Twitter
Related technologies	• File and Storage Services Overview

Content type	References
	• Failover Clustering Overview • Storage Spaces Overview

Plan to Deploy Data Deduplication

This document explains how to plan for deploying Data Deduplication. To be safe, backup and restore a deduplicated volume to ensure that your deployment is covering the entire dataset.

In this document

1. Step 1: Target deployments
2. Step 2: Determine which volumes are candidates for deduplication
3. Step 3: Evaluate savings with the Deduplication Evaluation Tool
4. Step 4: Plan the rollout, scalability, and deduplication policies

Step 1: Target deployments

Data Deduplication is designed to be installed on primary data volumes without adding additional dedicated hardware. This means that you can install and use the feature without impacting the primary workload on the server. The default settings are nonintrusive because they allow data to age for five days before processing a particular file, and has a default minimum file size of 32 KB. The implementation is designed for low memory and CPU usage, and if memory utilization becomes high, deduplication will wait for available resources. Administrators can schedule more aggressive deduplication based on the type of data that is involved and the frequency and volume of changes that occur to the volume or particular file types.

The ideal workloads for deduplication include:

- **General file shares**: Group content publication and sharing, user home folders, and Folder Redirection/Offline Files
- **Software deployment shares**: Software binaries, images, and updates
- **VHD libraries**: Virtual hard disk (VHD) file storage for provisioning to hypervisors
- **VDI Deployments (Windows Server 2012 R2 only)**: Virtual Desktop Infrastructure (VDI) deployments using Hyper-V

Step 2: Determine which volumes are candidates for deduplication

Deduplication can be very effective to optimize storage and reduce the amount of disk space consumed—50% to 90% when applied to the right data. Evaluate the following considerations when you select data for deduplication:

5. 0.1. **Does the data have duplication within it?**

   File shares or servers that host user documents, software deployment binaries, or virtual hard disk files tend to have plenty of duplication, and they will yield high savings from deduplication. The following table highlights typical deduplication savings for various content types. Results will vary by data type, mix, and size. It is best to run an evaluation on a data sample before you enable deduplication. For Details about using the DDPEval tool to evaluate deduplication savings see Step 3: Evaluate savings with the Deduplication Evaluation Tool later in this topic.

Scenario	Content	Typical Space Savings
User documents	Documents, photos, music, videos	30-50%
Deployment shares	Software binaries, cab files, symbols files	70-80%
Virtualization libraries	Virtual hard disk files	80-95%
General file share	All of the above	50-60%

2. **Does the data access pattern allow for sufficient time for deduplication?**

   Files that change often and are constantly accessed by users or applications are not good candidates for deduplication. The constant access and change to the data are likely to cancel any optimization gains made by deduplication, and deduplication may not be able to process the files.

   - A good candidate for deduplication is a file share that hosts user documents, virtual files, or software deployment files that contain data that is modified infrequently and read frequently.
   - Poor candidates for deduplication are a constantly-mounted SQL Server database that is running virtual machines, and live Exchange Server databases.

   Good candidates allow time to deduplicate the files. File age policies can be applied to control when files are deduplicated to prevent early or frequent deduplication of files that are still likely to be modified significantly.

3. **Does the server have sufficient resources and time to run deduplication?**

   Deduplication requires reading, processing, and writing large amounts of data. This consumes server resources and must be considered when planning a deployment. Servers typically have periods of high activity and times when there is low resource utilization. Deduplication can get the most work done when resources are available. If a server is always running at maximum capacity, it is probably not a good candidate for deduplication, even though the deduplication processes can optimize some files by using background optimization jobs.

   Based on observed savings and typical resource usage, deployment candidates for deduplication are ranked as follows:

Great candidates for deduplication:

- Folder redirection servers
- Virtualization depot or provisioning library
- Software deployment shares
- SQL Server and Exchange Server backup volumes
- VDI VHDs (supported only on Windows Server 2012 R2)

Should be evaluated based on content:

- Line-of-business servers
- Static content providers
- Web servers
- High-performance computing (HPC)

Not good candidates for deduplication:

- Hyper-V hosts
- WSUS
- Servers running SQL Server or Exchange Server
- Files approaching or larger than 1 TB in size
- VDI VHDs on Windows Server 2012

Server and volume requirements for deduplication

Servers

See the following list for server requirements for deduplication:

- Server hardware should meet the minimum requirements for running Windows Server 2012. The deduplication feature was designed to support minimal configurations such as a single processor system with 4 GB of RAM and one SATA hard disk drive.
- If you plan to support deduplication on multiple volumes on the same server, you need to plan an appropriately size for the system to ensure that it can process the data. The general rule is that the server needs 1 CPU-core and 350 MB of free memory to run a deduplication job on a single volume, and that job can process about 100 GB per hour or around 2 TB per day. Deduplication scales with additional CPU core processors and available memory to enable parallel processing of multiple volumes.

  For example: If you have a server with 16 CPU core processors and 16 GB of memory, deduplication uses 25% of the system memory in the default Background Processing mode, and in this case, that would be 4 GB. If you divide by 350 MB, you can calculate that the server could process about 11 volumes at a time. If you add 8 GB of memory, the system could process 17 volumes at a time. If you set an optimization job to run in Throughput mode, the system will use up to 50% of the system's memory for the optimization job.
- Data Deduplication supports up to 90 volumes at a time; however, deduplication can simultaneously process one volume per physical CPU core processor plus one. Hyper-

threading does not impact this because only physical core processors can be used to process a volume. A system with 16 CPU core processors and 90 volumes will process 17 volumes at a time until all 90 volumes are done, if there is sufficient memory.

- Virtual server instances should follow the same guidance as physical hardware regarding server resources.

Volumes

Volumes that are candidates for deduplication must conform to the following requirements:

- Must not be a system or boot volume. Deduplication is not supported on operating system volumes.
- Can be partitioned as a master boot record (MBR) or a GUID Partition Table (GPT), and must be formatted using the NTFS file system.
- Can reside on shared storage, such as storage that uses a Fibre Channel or an SAS array, or when an iSCSI SAN and Windows Failover Clustering is fully supported.
- Do not rely on Cluster Shared Volumes (CSVs). You can access data if a deduplication-enabled volume is converted to a CSV, but you cannot continue to process files for deduplication.
- Do not rely on the Microsoft Resilient File System (ReFS).
- Must be exposed to the operating system as non-removable drives. Remotely-mapped drives are not supported.

Note

Files with extended attributes, encrypted files, files smaller than 32 KB, and reparse point files are not processed by deduplication.

Unsupported configurations

Constantly open or changing files

In Windows Server 2012, deduplication is not supported for files that are open and constantly changing for extended periods of time or that have high I/O requirements, for example, running virtual machines on a Hyper-V host, live SQL Server databases, or active VDI sessions.

Beginning with Windows Server 2012 R2, deduplication adds support for running virtual machines in a VDI configuration. The VHDs used must be stored on a file server separate from the Hyper-V host computer.

Deduplication optimization of open files is supported only for VDI scenarios with Hyper-V using remote storage over SMB 3.0. Deduplication of open files has not been fully validated for general virtualization or other applications, and additional issues may be encountered.

Deduplication can be set to process files that are 0 days old and the system will continue to function as expected, but it will not process files that are exclusively open. It is not a good use of server resources to deduplicate a file that is constantly being written to, or will be written to in the near future. If you adjust the default minimum file age setting to 0, test that deduplication is not constantly being undone by changes to the data.

In Windows Server 2012, deduplication will not process files that are constantly and exclusively open for write operations. This means that you will not get any deduplication savings unless the file is closed when an optimization job attempts to process a file that meets your selected deduplication policy settings.

Beginning with Windows Server 2012 R2, the VHDs of running virtual machines deployed in a VDI configuration can be optimized. The deduplication process will optimize the parts of the VHD file that are older than the selected minimum days defined by the deduplication policy. Running virtual machine VHDs can realize considerable storage savings.

FSRM hard quotas

Setting a hard quota on a volume root folder that also has deduplication enabled is not supported. When a hard quota is present on a volume root, the actual free space on the volume and the quota-restricted space on the volume are not the same. This may cause deduplication jobs to fail. Creating a soft quota on a volume root that has deduplication enabled is supported; however, when FSRM quotas encounter a deduplicated file, it accounts for it based on the file's logical size. Quota usage (including any quota thresholds) does not change when a file is processed by deduplication. All other FSRM quota functionality, including volume-root soft quotas and quotas on subfolders, will work normally when using deduplication.

Deduplication considerations

Working with initial datasets

Volumes brought under deduplication control may be empty or they may already contain data. If the volume already has data, deduplication may take an extended amount of time to process the data for the first time. This is a normal condition and should be expected. Data Deduplication processes files when the server has available resources and will automatically resume processing. The rate of processing depends on the system specifications. Volume processing tests by Microsoft show about 20 MB per second or 2 TB per day. Administrators can monitor the progress of deduplication by using Windows PowerShell or in Server Manager. After the initial dataset is processed, deduplication has minimal impact on the server, and it does not require manual maintenance.

Write path for new files

Data deduplication allows new writing directly to an NTFS volume. Files on the volume are evaluated periodically and processed if they meet the deduplication policy criteria. Deduplication processes files inside an enabled volume once per hour in the background or on additional schedules that you specify by using Throughput mode optimization jobs.

Deduplication throughput

The deduplication feature can process roughly 2 TB of data per volume in a 24-hour period (100 GB per hour) when running a throughput optimization job on a single volume. Multiple volumes can run in parallel if CPU core processors and memory are available. The actual speed of deduplication is dependent on CPU, disk I/O, memory and available server resources at the time that deduplication is running. Deduplication is designed to pause if the server starts to run out of resources so it does not impact the server's regular workload. The busier the server is, the

longer it will take to process the data. If you have a data volume that is very busy, you might consider scheduling a regular throughput optimization job during off-peak hours. For more information, see **Enable data deduplication on the volume**.

Running low on disk space

Running very low on disk space on a deduplicated volume is not advisable because certain types of data access operations can fail unexpectedly. For example, memory-mapped operations on deduplicated files can fail due to lack of disk space. This is because NTFS may reserve disk space in response to paging I/O that is triggered through memory-mapped or buffered operations. Notepad is an example of an application that accesses files by using memory-mapped operations. If the volume is at or near zero free space, opening deduplicated files in Notepad may return an error message that says disk access has failed. Similar data-access failures can occur on sparse or compressed files on a full volume without deduplication installed. We recommend keeping at least one gigabyte of free space on deduplicated volumes to avoid failures.

If the volume is full and you encounter file access failures, consider one of the following options to work around the issue:

- Run a Garbage Collection job to reclaim disk space by using an elevated Windows PowerShell command prompt:

  ```
  PS C:\> Start-DedupJob E: -Type GarbageCollection -full
  ```

- Copy the files elsewhere by using a tool that supports unbuffered I/O mode, such as XCOPY with the /J option or Robocopy with the /J option.

Step 3: Evaluate savings with the Deduplication Evaluation Tool

When the deduplication feature is installed, the Deduplication Evaluation Tool (DDPEval.exe) is automatically installed to the \Windows\System32\ directory. This tool can be copied from any server running Windows Server 2012 R2 or Windows Server 2012 to systems running Windows Server 2012, Windows Server 2008 R2, or Windows 7. You can use it to determine the expected savings that you would get if deduplication is enabled on a particular volume.

DDPEval.exe supports local drives, and it supports evaluating mapped or unmapped remote shares. The following example shows running DDPEval.exe by using a command syntax:

```
DDPEval <VolumePath:>


C:\> DDPEVAL.EXE \\Server\Testshare

C:\> DDPEVAL.EXE E:\Test\

C:\> DDPEVAL.EXE F:
```

This will produce output similar to the following:

```
Data Deduplication Savings Evaluation Tool
```

```
Copyright 2011-2012 Microsoft Corporation.  All Rights Reserved.


Evaluated folder: E:\Test

Processed files: 34

Processed files size: 12.03MB

Optimized files size: 4.02MB

Space savings: 8.01MB

Space savings percent: 66

Optimized files size (no compression): 11.47MB

Space savings (no compression): 571.53KB

Space savings percent (no compression): 4

Files with duplication: 2

Files excluded by policy: 20

Files excluded by error: 0
```

When the space savings percent is calculated, you can make an informed decision to install deduplication or not.

Determine potential space reclamation with the Measure-DedupFileMetadata cmdlet

By running the Windows PowerShell **Measure-DedupFileMetadata** cmdlet, you can determine how much potential disk space can be reclaimed on a volume if you delete a group of folders, a single folder, or a single file, and then run a Garbage Collection job. Specifically, the DedupDistinctSize value will tell you how much space you would get back if you deleted those files. Files often have chunks that are shared across other folders, so the deduplication engine will calculate which chunks are unique and would be deleted after the Garbage Collection job.

Here is an example:

```
PS C:\Windows\system32> Measure-DedupFileMetadata -path X:\A_Data,X:\Archive1


Path                    : {X:\A_Data, X:\Archive1}

Volume                  : X:

VolumeId                : \\?\Volume{88f3b478-72e9-11e1-b0f1-806e6f6e6963}

FilesCount              : 575

OptimizedFilesCount     : 541

Size                    : 713.46 MB

SizeOnDisk              : 2.44 MB
```

```
DedupSize               : 239.81 MB

DedupChunkCount         : 5136

DedupDistinctSize       : 700 KB

DedupDistinctChunkCount : 11
```

In this example, 700 KB of space will be reclaimed if Garbage Collection is run with the **–full** parameter.

Step 4: Plan the rollout, scalability, and deduplication policies

When appropriate servers and volumes have been identified, you can enable the feature on the server and assign volumes. For more information, see Install and Configure Data Deduplication.

Deduplication can be enabled locally or remotely by using Windows PowerShell or Server Manager. When it is enabled, the default policies and schedules are created.

The default policy settings for data deduplication include:

- Process files that have a minimum age of five days according to the Last Modified Time. If Last Access Time is enabled on the server (this is not the default setting), deduplication will use the Last Access Time.
- Process files in background mode every hour. In background mode, the system uses up to 25% of the system memory during optimization jobs, whereas manual Throughput jobs use up to 50% of the system memory.
- Do not exclude any directories or file types. The default setting is to process the entire volume.
- Run a garbage collection job every Saturday at 1:45 AM. Garbage collection reclaims space on a volume by deleting chunks from the chunk store that are no longer referenced. Garbage collection compacts a container only if approximately 50 MB of chunks exist that have no references. Every fourth run of garbage collection incorporates the **-full** parameter, which instructs the job to reclaim all available space and maximize all container compaction.
- Run a data scrubbing job every Saturday at 2:45 AM. Scrubbing jobs verify data integrity and automatically attempt to repair corruptions that are found.

Note

Do not enable **NtfsDisableLastAccessUpdate** when using System Center Virtual Machine Manager (VMM) VHD library server, and deduplication on a hosting volume. The VMM library server will toggle the last-access date often, preventing deduplication from optimizing files correctly.

The default deduplication policy settings are usually sufficient to get excellent savings without impacting regular server workload. If a particular deployment has any of the following conditions, you might consider altering the default settings.

Potential deduplication policy settings

Condition	Action to consider
Incoming data is static or expected to be read-only, and you want to process files on the volume sooner.	Change the MinimumFileAgeDays setting to a smaller number of days to process files earlier.
You have directories that you do not want to deduplicate.	Add a directory to the exclusion list.
You have file types that you do not want to deduplicate.	Add a file type to the exclusion list.
The server has different off-peak hours than the default and you want to change the Garbage Collection and Scrubbing schedules.	Update the schedules using Windows PowerShell.

Backup and Restore Considerations for Deduplicated Volumes

One of the benefits of deduplication is that by reducing the used space on a data volume, it makes backing up that volume faster, which enables longer retention times, greater efficiency to the backup storage target, and quicker restoration operations.

Block-based backup applications should work without modification, and they maintain the optimization on the backup media. File-based backup operations that do not use deduplication usually copy the files in their original form. If an administrator is using XCOPY, for instance, the restoration happens in memory during the backup copy operation. The backup target must be large enough to hold the logical size of the entire dataset because none of the deduplication is retained.

Caution

Do not attempt to manually back up the reparse points and the deduplication chunk-store files that are located in the hidden Sysvol folder by using Robocopy. Data Deduplication may be modifying data, and this could cause data loss.

Windows Server Backup

Windows Server Backup has the ability to back up an optimized volume “as-is” (that is, without removing deduplicated data). The following examples show how to back up a volume and how to restore a volume or selected files from a volume.

1. Install Windows Server Backup by running the following Windows PowerShell command:

```
PS C:\> Add-WindowsFeature -name Windows-Server-Backup
```

2. To back up the E: volume to another volume, run the following from a command prompt, substituting the correct volume names for your situation:

   ```
   C:\> wbadmin start backup –include:E: -backuptarget:F: -quiet
   ```

3. Get the version ID of the backup you just created:

   ```
   C:\> wbadmin get versions
   ```

   This output version ID will be a date and time string, for example: 02/16/2012-06:22.

4. Restore the entire volume:

   ```
   C:\> Wbadmin start recovery –version:02/16/2012-06:22 -
   itemtype:Volume  -items:E: -recoveryTarget:E:
   ```

5. Restore a particular folder (in this case, the E:\Docs folder):

   ```
   C:\> Wbadmin start recovery –version:02/16/2012-06:22 -
   itemtype:File  -items:E:\Docs  -recursive
   ```

Optimized backup considerations

Performing an optimized backup results in a smaller, faster backup. The backup is smaller because the total size of the optimized files, non-optimized files (files that are not included in policies), and data deduplication chunk store files are significantly smaller than the full logical size of the volume (that is, the size of all the files before deduplication). Optimized backups are faster because there is less I/O and because the optimized files are not restoredrehydrated during the file copy operations. Selective backup (in which most of the volume is being backed up) may also benefit from using the optimized backup approach if the logical size of the selected files is significantly greater than the physical size of the optimized files plus the chunk store container files.

Incremental optimized backup considerations

A backup application can perform an incremental optimized backup as follows:

- Back up only the changed user files, including reparse points that were created, modified, or deleted since the previous backup.
- Back up the changed chunk store container files. (Because new chunks are appended to the current chunk store container, when its size reaches about 1 GB, that container file is sealed and a new container file is created.)
- Perform an incremental backup at the sub-file level. (You can back up file ranges instead of entire files for large files that rarely change).

Full volume restore from optimized backup

File-level full volume restore should be performed in an optimized manner by essentially reversing the procedure described earlier. The complete set of data deduplication metadata and

container files are restored, then the complete set of data deduplication reparse points are restored, followed by restoration of all non-deduplicated files.

Restoring volumes in this manner provides faster recovery time (compared to full-volume restore from a non-optimized backup store) because the amount of copied data is much smaller. Also the size of the restored data in an optimized full-volume restore fits on the original size volume. Restoring to the original or an equivalently sized volume from a non-optimized backup will likely fill the disk because the files will be replaced at their full logical (non-deduplicated) size.

Block-level restore from an optimized backup is automatically an optimized restore because the restore process occurs under Data Deduplication, which works at the file level.

Note
We recommend that you always perform a full volume restore to the original reformatted volume or to a freshly formatted volume of equivalent or greater size.

Selective restore from an optimized backup

Selective restore from an optimized backup can be done by using one of the following approaches:

- Partial volume restore to a newly formatted volume
- Restore from a -backup store that allows data deduplication by using a backup application that supports the selective restore API provided by Data Deduplication

Selective partial volume restore to a newly formatted volume

This approach is similar to the file-level, full-volume restore described earlier (only a subset of the data deduplication reparse points are restored to a freshly formatted volume).

Note
With this approach, the entire set of data deduplication metadata and container files are copied to the new volume.

This approach is appropriate if most of the volume is being restored, and only when there are portions of the volume that are not required on the restored volume. The application can make a calculated decision about whether this is an appropriate technique by comparing the size of the data deduplication store files plus the physical size of all of the restored files to the total logical size of the target restore file set.

Selective restore from a backup store that allows data deduplication

If the backup store is located on a volume that is enabled for data deduplication, the backup store and the original volume are optimized. In this case, the backup application can copy the files as normal files back to the target volume. The files are restored in memory from the backup store, so the application is subjected to increased I/O latency while it is reading the backup store.

If the file on the target volume is missing, or if it has been fully recalled from the data deduplication store, there will be no increased latency on the target volume. However, to be sure, it is a good practice to first delete the file (if it exists) from the target volume before you restore the file with this approach.

Non-optimized backup and restore

In non-optimized backup and restore, the backup application does not use the Data Deduplication backup and restore API. Instead, the backup application opens the files and copies them without specifying the reparse point flag.

The optimized files are coped to the backup volume as normal files, not as optimized files. The conversion from optimized files to normal files is performed transparently in memory by Data Deduplication when the backup application copies the files. Restoring from such a backup store is a normal file-copy operation.

The size of the data in a non-optimized backup is normally much larger than the original optimized volume because of the space savings that is provided by deduplication. A full volume restore from a non-optimized backup will usually not fit on the original or an equivalently sized volume.

Non-optimized backup is normally used only for the following situations:

- To selectively back up a small percentage of the files on the volume, where the logical size of the selected files is significantly smaller than the physical size of the optimized files plus the chunk store container files.
- To support restoring a Windows Server 2012 backup to an earlier version of the Windows Server operating system.
- To support restoring a Windows Server 2012 backup to a computer that is not running Windows Server.

Incremental non-optimized backup

A backup application can perform an incremental optimized backup as follows:

- To back up changed user files
- To skip reparse point changes

Data Deduplication Interoperability

When using the Data Deduplication feature for the first time or migrating from a previous version of Windows Server, be sure to consider the following related technologies and issues:

- BranchCache
- Failover Clusters
- DFS Replication

- FSRM quotas
- Single Instance Storage

BranchCache

Access to the data over the network can be optimized by enabling BranchCache on servers and clients. When a BranchCache-enabled system communicates over a WAN with a remote file server that is running data deduplication, all of the deduplicated files are already indexed and hashed, so requests for data from a branch office are quickly computed. This is similar to preindexing or prehashing a BranchCache-enabled server.

About BranchCache BranchCache is a feature that can reduce wide area network (WAN) utilization and enhance network application responsiveness when users access content in a central office from branch office locations. When you enable BranchCache, a copy of the content that is retrieved from the web server or file server is cached within the branch office. If another client in the branch requests the same content, the client can download it directly from the local branch network without needing to retrieve the content by using the WAN.

Failover clusters

Failover clusters are fully supported, and deduplicated volumes will failover gracefully. Each node that accesses deduplicated volumes must be running the Data Deduplication feature. A deduplicated volume is a self-contained and portable unit—all of the data and configuration information is contained on the volume. Deduplication schedules are managed by using the task scheduler and when a cluster is formed, the schedule information is put into the schedule for the cluster. If a deduplicated volume is taken over by another node, the scheduled job will be applied on the next scheduled interval.

DFS Replication

The Data Deduplication feature has interoperability with Distributed File System (DFS) Replication. Optimizing or unoptimizing a file will not trigger a replication because the file does not change. DFS Replication uses Remote Differential Compression (RDC), not the chunks in the chunk store, for over-the-wire savings. The files on the replica can also be optimized by using deduplication if the replica is running Windows Server 2012.

File Server Resource Manager quotas

Creating a hard quota on a volume root folder that also has deduplication enabled is not supported. When a hard quota is present on a volume root, the actual free space on the volume and the quota restricted space on the volume are not the same. This may cause deduplication optimization jobs to fail.

Creating a soft quota on a volume root that has deduplication enabled is supported. When FSRM quotas encounter a deduplicated file, File Server Resource Manager accounts for it based on the

file's logical size. Quota usage (including any quota thresholds) does not change when a file is processed by deduplication. All other FSRM quota functionality, including volume-root soft quotas and quotas on subfolders, work normally when using deduplication.

Single Instance Storage

Volumes that have been created and optimized by using the Windows Storage Server version of deduplication, Single Instance Storage (SIS), must not also run Data Deduplication. We recommend that prior to migrating a volume to Windows Server 2012 R2 or Windows Server 2012, you remove SIS from the volume by using SISAdmin.exe within Windows Storage Server or by copying the data to a different volume that is not running SIS. Windows Server 2012 R2 and Windows Server 2012 support reading and writing to SIS-controlled volumes, but you cannot continue to use SIS by using Windows Server 2012 R2 or Windows Server 2012. You can use the following command syntax to install the SIS-Limited feature (which includes the SIS filter driver) on Windows Server 2012 R2 or Windows Server 2012:

You can install the SIS filter driver to read SIS-controlled volumes. The data can be copied to a non-SIS controlled volume so that data deduplication can be installed on the volume.

You can use the following command syntax to install the SIS-Limited feature (which includes the SIS filter driver) on Windows Server 2012 R2 or Windows Server 2012:

```
C:\> dism /online /enable-feature:SIS-Limited
```

Note

Windows Server 2012 R2 and Windows Server 2012 do not support sisadmin.exe and cannot be used to remove SIS from a volume.

Before you install the SIS filter driver, consider the following:

1. You should remove SIS from your volumes before you install Data Deduplication.
2. Do not restore SIS links from a backup to a deduplicated volume.
3. Restoring SIS volumes to Windows Server 2012 R2 or Windows Server 2012 is supported if you load the SIS-Limited feature.

Migrating SIS volumes

You have several options when it comes to migrating Windows Storage Server 2008 volumes to Windows Server 2012 R2 or Windows Server 2012 to take advantage of the Data Deduplication feature.

You can migrate your existing SIS-installed Windows Storage Server 2008 volumes. However, migration is not automatic because Single Instance Storage (SIS) and data deduplication are mutually-exclusive technologies.

You need to first open the volumes in Windows Storage Server 2008, remove SIS from the volumes, and then uninstall SIS before migrating to Windows Server 2012.

To remove SIS

- From a command prompt in Windows Storage Server 2008 R2 or Windows Storage Server 2008, type:

  ```
  sisadmin.exe [/m <server>] [/u <volumes>
  ```

  where:

  /m <server> shifts the focus of the command line to a remote server. If the **/m** option is not specified, the command line is applied to the local server. **<server>** can be expressed as a host name, fully qualified domain name (FQDN), or an IP address.

  /u <volumes> is used to remove SIS from a volume. This restores all file copies and removes reparse points.

  For each command option that uses **<volumes>** as a parameter, **<volumes>** represents a space-delimited list of volume names (for example, D: E: F: G:).

  For example:

  To remove SIS entirely from volume F: of a remote server by using the IP address of the server, you might use the following command:

  ```
  sisadmin.exe /m 192.168.1.50 /u F:
  ```

Install and Configure Data Deduplication

This document explains how to set up a server, enable data deduplication, optimize a volume, and carry out advanced deduplication operations.

In this document

- Prerequisites
- Step 1: Set up the server
- **Step 2: Enable data deduplication**
- **Step 3: Set data deduplication jobs**
- **Step 4: Set data deduplication schedules**

Prerequisites

Before installing and configuring Data Deduplication, review the following resources:

- Plan to Deploy Data Deduplication
- Backup and Restore Considerations for Deduplicated Volumes
- Data Deduplication Interoperability

Step 1: Set up the server

To install deduplication components on the server by using Server Manager

1. From the Add Roles and Features Wizard, under Server Roles, select **File and Storage**

Services (if it has not already been installed).

2. Select the **File Services** check box, and then select the **Data Deduplication** check box.
3. Click **Next** until the **Install** button is active, and then click **Install**.

To install deduplication components on the server by using Windows PowerShell

1. Start Windows PowerShell. Right-click the Windows PowerShell icon on the taskbar, and then click **Run as Administrator**.
2. Run the following Windows PowerShell commands:

```
Import-Module ServerManager

Add-WindowsFeature -name FS-Data-Deduplication

Import-Module Deduplication
```

Step 2: Enable data deduplication

To enable data deduplication by using Server Manager

1. From the Server Manager dashboard, right-click a data volume and choose **Configure Data Deduplication**. The Deduplication Settings page appears.
2. Enable data Deduplication.

   On Windows Server 2012 R2: In the **Data deduplication** box, select the workload you want to host on the volume. Select **General purpose file server** for general data files or **Virtual Desktop Infrastructure (VDI) server** when configuring storage for running virtual machines.

   On Windows Server 2012: Select the **Enable data deduplication** check box
3. Enter the number of days that should elapse from the date of file creation until files are deduplicated, enter the extensions of any file types that should not be deduplicated, and then click **Add** to browse to any folders with files that should not be deduplicated.
4. Click **Apply** to apply these settings and return to the Server Manager dashboard, or click the **Set Deduplication Schedule** button to continue to set up a schedule for deduplication.

To enable data deduplication by using Windows PowerShell

1. To enable deduplication on a volume, run the following Windows PowerShell command on the server. In this example deduplication is enabled on volume E.

   On Windows Server 2012 R2:

```
Enable-DedupVolume E: -UsageType HyperV

Enable-DedupVolume E: -UsageType Default
```

   On Windows Server 2012:

```
Enable-DedupVolume E:
```

2. Optionally, set the minimum number of days that must pass before a file is deduplicated by using the following command.

   ```
   Set-Dedupvolume E: -MinimumFileAgeDays 20
   ```

   If you set MinimumFileAgeDays to 0, deduplication will process all files, regardless of their age. This is suitable for a test environment, where you want to exercise maximum deduplication. In a production environment, however, it is preferable to wait for a number of days (the default is three days in Windows Server 2012 R2 and five days in Windows Server 2012), because files tend to change a lot for a brief period of time before the change rate slows. This allows for the most efficient use of your server resources.

To return a list of the volumes that have been enabled for data deduplication by using Windows PowerShell

Run the following Windows PowerShell commands on the server.

```
Get-DedupVolume
```

```
Get-DedupVolume | format-list
```

The first command returns summary information and the second returns details about the volume data deduplication settings.

Step 3: Set data deduplication jobs

Data deduplication jobs can be run on demand (manually) or scheduled. There are three types of jobs that you can perform on a volume: Optimization, Data Scrubbing, and Garbage Collection.

Optimization jobs

The Data Deduplication feature comes with built-in jobs that will automatically launch and optimize the specified volume(s) on a regular basis. Optimization jobs deduplicate data and compress file chunks on a volume per the policy settings. After the initial optimization is complete, optimization jobs run on the files that are included in the policies, according to the job schedules that you have configured or the default job schedules that ship with the product.

You can trigger an optimization job on demand in Windows PowerShell by using the **Start-DedupJob** cmdlet. For example:

```
Start-DedupJob –Volume E: –Type Optimization
```

This command returns immediately and the job is launched asynchronously. If you want the job to complete at a later time , add the –wait parameter, like this:

```
Start-DedupJob E: –Type Optimization -Wait
```

You can query the progress of the job on the volume by using the **Get-DedupJob** cmdlet:

```
Get-DedupJob
```

The **Get-DedupJob** command show current jobs that are running or are queued to run.

You can query the key status statistics including the achieved savings on the volume by using the **Get-DedupStatus** cmdlet:

```
Get-DedupStatus | fl
```

The **Get-DedupStatus** command shows the free space, space saved, optimized files, InPolicyfiles (the number of files that fall within the volume deduplication policy, based on the defined file age, size, type, and location criteria), and the associated drive identifier.

Note

You can also view the deduplication savings in Server Manager on the Volumes page. From Server Manager, click **File Services**, and then click **Volumes**. Right-click the column heading to add **Deduplication Savings**.

Optimization job queuing

Optimization jobs are started in the following order:

1. **Preemptive (manually run jobs that are not scheduled)**

   Any manual jobs that include the **–Preempt** option will terminate any jobs that are currently running, and start immediately. (Note that the **–Preempt** option is ignored in scheduled jobs.)

2. **StopWhenSystemBusy parameter**

   Jobs that contain this parameter will stop if resources are not available to run the job without interfering with the server's workload.

3. **Priority**

   Among jobs that do not have the same **StopWhenSystemBusy** setting, high priority jobs are queued first, normal jobs are queued second, and low priority job are queued last.

4. **Manual or scheduled**

   Manual jobs are queued before scheduled jobs.

Memory settings are not considered as part of the optimization job queue algorithm.

Optimization metadata

Metadata provides you with evidence about savings that you gleaned from using optimization. There are three cmdlets that output this metadata: **Update-DedupStatus**, **Get-DedupMetadata**, and **Measure-DedupFileMetadata**. This metadata can help you assess the impact of some optimization configuration options.

Update-DedupStatus returns the following metadata:

Metadata	What it indicates
DedupSavedSpace	Difference between the logical size of the optimized files and the logical size of the store (the deduplicated user data plus deduplication metadata). This number changes continually.
DedupRate	Ratio of DedupSavedSpace to the logical size of all of the files on the volume, and it is

Metadata	What it indicates
	expressed as a percentage. This number changes continually.
OptimizedFilesCount	Number of optimized files on the specified volume. Note that this number will remain steady (instead of decrease) as users delete files from or add files to the volume, until you run a Garbage Collection job. This count is most accurate after a full garbage collection job runs.
OptimizedFilesSize	Aggregate size of all optimized files on the specified volume. Note that this number remains steady (instead of decreasing) as users delete files from or add new files to the volume, until you run a garbage collection job. This number is most accurate after a full garbage collection job runs.
InPolicyFilesCount	Number of files that currently qualify for optimization. This number stays relatively constant between optimization jobs.
InPolicyFilesSize	Aggregate size of all files that currently qualify for optimization. This number stays relatively constant between optimization jobs.
LastOptimizationTime	Date and time when an optimization job was last run on the specified volume. This date and time stays constant between optimization jobs.
LastGarbageCollectionTime	Date and time when a garbage collection job was run last on the specified volume. This date and time stays constant between optimization jobs.
LastScrubbingTime	Date and time when a scrubbing job was run last on the specified volume. This date and time stays constant between optimization jobs.

Get-DedupMetadata returns the following metadata:

Metadata	What it indicates
DataChunkCount	Number of data chunks on the volume.

Metadata	What it indicates
DataContainerCount	Number of containers in the data store.
DataChunkAverageSize	Data store size (not including chunk metadata) divided by the total number of data chunks in the data store.
StreamMapCount	Number of data streams on the volume.
StreamMapContainerCount	Number of containers in the stream map store.
StreamMapAverageChunkCount	Stream map store size divided by the total number of streams in the store.
HotspotCount	Number of “hotspot” chunks on the volume. A hotspot is a chunk that is referenced over 100 times. All hotspot chunks are duplicated on the volume to provide automatic data corruption recovery in the event that corruption occurs on the disk and impacts one of these popular chunks.
HotspotContainerCount	Number of hotspot containers.
CorruptionLogEntryCount	Number of corrupted items on the volume.

Data Scrubbing jobs

Data Deduplication has built-in data integrity features such as checksum validation and metadata consistency checking. It also has built-in redundancy for critical metadata and the most popular data chunks. As data is accessed or jobs process data, these features may encounter corruption, and they will record the corruption in a log file. Scrubbing jobs use these features to analyze the chunk store corruption logs, and when possible, to make repairs. Possible repair operations include using three sources of redundant data:

1. Deduplication keeps backup copies of popular chunks when they are referenced over 100 times in an area called the hotspot. If the working copy is corrupted, deduplication will use the backup.
2. When using Storage Spaces in a mirrored configuration, deduplication can use the mirror image of the redundant chunk to serve the I/O and fix the corruption.
3. If a file is processed with a chunk that is corrupted, the corrupted chunk is eliminated, and the new incoming chunk is used to fix the corruption.

Scrubbing jobs output a summary report in the Windows event log located here:

Event Viewer\Applications and Services Logs\Microsoft\Windows\Deduplication\Scrubbing

Data Deduplication defaults create a data integrity scrubbing job on a weekly basis, but you can also trigger one on demand by using the following command:

```
Start-DedupJob E: -Type Scrubbing
```

This initiates a job that attempts to repair all corruptions that were logged in to the deduplication internal corruption log, during the I/O to deduplication files. To check the data integrity of all the deduplicated data on the volume, use the **-full** parameter:

```
Start-DedupJob E: -Type Scrubbing -full
```

Also known as Deep Scrubbing, the **-full** parameter will scrub the entire set of deduplicated data and look for all corruptions that are causing data access failures.

Garbage Collection jobs

Data Deduplication includes garbage collection jobs to process deleted or modified data on the volume so that any data chunks no longer referenced are cleaned up. Garbage collection jobs process previously deleted or logically overwritten optimized content to create usable volume free space. When an optimized file is deleted or overwritten by new data, the old data in the chunk store is not immediately deleted.

Note

Garbage collection is a processing-intensive operation, so you should allow the deletion load to reach a threshold and then manually run this job type, or schedule it for off hours.

Garbage collection can also be triggered on demand. For example:

```
Start-DedupJob E: -Type GarbageCollection
```

This command removes unreferenced chunks and compacts containers that have more than 5% unreferenced data. By adding the **–full** parameter, the job will compact all containers to the maximum extent possible. For example:

```
Start-DedupJob E: -Type GarbageCollection -full
```

Step 4: Set data deduplication schedules

Data Deduplication comes with three schedules that are set up immediately. Optimization runs every hour, and Garbage Collection and Scrubbing are set for once a week. You can view the schedules by using this Windows PowerShell command:

```
Get-DedupSchedule

Enabled    Type               StartTime             Days              Name
-------    ----               ---------             ----              ----
True       Optimization
BackgroundOptimization
True       GarbageCollection  2:45 AM               Saturday
WeeklyGarbageCollection
True       Scrubbing          3:45 AM               Saturday          WeeklyScrubbing
```

Two additional schedules can be used immediately to add jobs. These job schedules run on all volumes on the server. If you want to run a job only on a particular volume, you must create a new job. You can create, modify, or view job schedules from the **Deduplication Settings** page in Server Manager, or by using the following Windows PowerShell commands:

```
Set-DedupSchedule <ScheduleName> <properties>

Remove-DedupSchedule <ScheduleName>
```

Note
Deduplication only supports weekly job scheduling. If you want to create a schedule for a monthly job or any other time period, use Windows Task Scheduler. However, you will be unable to view custom job schedules that are created or modified with Task Scheduler by using the **Get-DedupSchedule** cmdlet. Such schedules are not migrated with server upgrades.

The built-in job schedules that Data Deduplication handles are:

- **BackgroundModeOptimization** Use this job schedule to run an Optimization job with the following parameters:

Parameter	Value
Enabled	True
Priority	Low
Memory	25
ContinueWhenSystemBusy	False
ScheduledTask	Microsoft\Windows\Deduplication\BackgroundModeOptimization
Start	00:00:00
Days	{Mon,Tues,Wed,Thurs,Fri,Sat,Sun}
Duration	0
Repeat	1

- **ThroughputModeOptimization** Use this job schedule to run an Optimization job with the following parameters:

Parameter	Value
Enabled	False
Priority	Normal
Memory	50
ContinueWhenSystemBusy	False
ScheduledTask	Microsoft\Windows\Deduplication\ThroughputModeOptimization
Start	00:00:00
Days	{Mon,Tues,Wed,Thurs,Fri,Sat,Sun}
Duration	4
Repeat	0

- **ThroughputModeOptimization2** Use this job schedule to run an Optimization job with the following parameters:

Parameter	Value
Enabled	False
Priority	Normal
Memory	50
ContinueWhenSystemBusy	False
ScheduledTask	Microsoft\Windows\Deduplication\ThroughputModeOptimization
Start	00:00:00
Days	{Mon,Tues,Wed,Thurs,Fri,Sat,Sun}
Duration	4
Repeat	0

- **WeeklyGarbageCollection** This default setting is scheduled to run a Garbage Collection job with the following parameters:

Parameter	Value
Enabled	True
Priority	Normal
Memory	50
ContinueWhenSystemBusy	False
ScheduledTask	Microsoft\Windows\Deduplication\WeeklyGarbageCollection
Start	01:45:00
Days	{Sat}
Duration	0
Repeat	0

- **WeeklyScrubbing** Use this job schedule to run a Scrubbing job with the following parameters:

Parameter	Value
Enabled	True
Priority	Normal
Memory	50
ContinueWhenSystemBusy	False
ScheduledTask	Microsoft\Windows\Deduplication\WeeklyScrubbing
Start	02:45:00
Days	{Sat}
Duration	0
Repeat	0

Operational considerations

Some files cannot be read when the free disk space on a deduplicated volume approaches zero. To resolve this issue, do one of the following:

- Run a Garbage Collection task to reclaim disk space.
- Copy files elsewhere (if there is not a recent memory map of the files).
- Run Robocopy.exe in non-cached Read mode to copy files elsewhere (if there is a recent memory map of the files). For more information about using Robocopy, see Robocopy.

Advanced deduplication policy considerations

With some server configurations, you may want to speed up deduplication. Here are some scenarios that might warrant additional job scheduling:

Condition	Action to consider
Significant incoming data	Add additional throughput optimization jobs.
More volumes than CPU core processors with significant incoming data	Add additional throughput optimization jobs.
Data deletions exceed 50 GB per hour, and you want to get the free space back as quickly as possible.	Add additional garbage collection jobs to reclaim the free space.

Monitor and Report for Data Deduplication

Administrators can query the progress of a deduplication job, view the achieved space savings on the volume, and view the status of the deduplication process by using the **Get-DedupStatus** and **Get-DedupVolume** Windows PowerShell cmdlets.

Here is an example of **Get-DedupStatus** output:

```
PS C:\> Get-DedupStatus

FreeSpace         SavedSpace   OptimizedFiles     InPolicyFiles      Volume

--------------    ----------   --------------     -------------      ------

280.26 GB         529.94 GB    36124              36125              X:

151.26 GB         84.19 GB     43017              43017              Z:
```

Here is an example of **Get-DedupStatus** verbose output for drive X: in the previous example by using the **|fl** pipe:

```
PS C:\> Get-DedupStatus |fl

Volume                              : X:

VolumeId                            : \\?\Volume{1234567890}\

Capacity                            : 465.76 GB

FreeSpace                           : 280.26 GB
```

```
UsedSpace                           : 185.5 GB
UnoptimizedSize                     : 715.44 GB
SavedSpace                          : 529.94 GB
SavingsRate                         : 74 %
OptimizedFilesCount                 : 36124
OptimizedFilesSize                  : 696.46 GB
OptimizedFilesSavingsRate           : 76 %
InPolicyFilesCount                  : 36125
InPolicyFilesSize                   : 704.44 GB
LastOptimizationTime                : 5/23/2012 2:45:09 PM
LastOptimizationResult              : 0x00000000
LastOptimizationResultMessage       : The operation completed successfully.
LastGarbageCollectionTime           : 5/19/2012 2:45:51 AM
LastGarbageCollectionResult         : 0x00000000
LastGarbageCollectionResultMessage : The operation completed successfully.
LastScrubbingTime                   : 5/19/2012 3:45:21 AM
LastScrubbingResult                 : 0x00000000
LastScrubbingResultMessage          : The operation completed successfully.
```

Here is an example of **Get-DedupVolume** output:

```
PS C:\> Get-DedupVolume

Enabled        SavedSpace          SavingsRate         Volume
-------       ----------           -----------         ------
True          529.94 GB            74 %                X:
```

Here is an example of **Get-DedupVolume** verbose output for drive X: in the previous example by using the **|fl** pipe:

```
PS C:\> Get-DedupVolume |fl

Volume                  : X:
VolumeId                : \\?\Volume{1234567890}\
Enabled                 : True
DataAccessEnabled       : True
Capacity                : 465.76 GB
FreeSpace               : 280.26 GB
```

```
UsedSpace                : 185.5 GB
UnoptimizedSize          : 715.44 GB
SavedSpace               : 529.94 GB
SavingsRate              : 74 %
MinimumFileAgeDays       : 0
MinimumFileSize          : 32768
NoCompress               : False
ExcludeFolder            :
ExcludeFileType          :
NoCompressionFileType    : {aac, aif, aiff, asf...}
ChunkRedundancyThreshold : 100
Verify                   : False
```

In addition to the status of a drive or volume, you can get detailed data about how deduplication processed the data by running the **Get-DedupMetadata** cmdlet:

```
PS C:\> Get-DedupMetadata
Volume                          : X:
VolumeId                        : \\?\Volume{1234567890}\
StoreId                         : {FDF037AE-4AAC-440F-84DF-A94A3D82DF7C}
DataChunkCount                  : 3314369
DataContainerCount              : 173
DataChunkAverageSize            : 52.48 KB
DataChunkMedianSize             : 0 B
DataStoreUncompactedFreespace   : 0 B
StreamMapChunkCount             : 36575
StreamMapContainerCount         : 16
StreamMapAverageDataChunkCount  :
StreamMapMedianDataChunkCount   :
StreamMapMaxDataChunkCount      :
HotspotChunkCount               : 10561
HotspotContainerCount           : 1
HotspotMedianReferenceCount     :
CorruptionLogEntryCount         : 0
TotalChunkStoreSize             : 167.2 G
```

The **Get-DedupMetadata** cmdlet shows the number of chunks in the chunk store, the number of StreamMaps, and the hotspot counts. (Hotspots are the most frequently-referenced chunks.)

Here is an example of a status report generated by **Get-DedupSchedule**:

```
PS C:\> Get-DedupSchedule

Enabled    Type                StartTime             Days              Name

-------    ----                ---------             ----              ----

True       Optimization
BackgroundOptimization

True       GarbageCollection   2:45 AM               Saturday
WeeklyGarbageCollection

True       Scrubbing           3:45 AM               Saturday          WeeklyScrubbing
```

Here is an example of **Get-DedupSchedule** verbose output in the previous example by using the `|fl` pipe:

```
PS C:\> Get-DedupSchedule |fl

Name               : BackgroundOptimization

Enabled            : True

Type               : Optimization

Days               :

Start              :

DurationHours      :

StopWhenSystemBusy : True

Memory             : 25 %

Priority           : Low

ScheduledTask      : \Microsoft\Windows\Deduplication\BackgroundOptimization

Full               :

ReadOnly           :


Name               : WeeklyGarbageCollection

Enabled            : True

Type               : GarbageCollection

Days               : Saturday

Start              : 3/22/2012 2:45:00 AM

DurationHours      : 48

StopWhenSystemBusy : True
```

```
Memory              : 50 %
Priority            : Normal
ScheduledTask       : \Microsoft\Windows\Deduplication\WeeklyGarbageCollection
Full                : False
ReadOnly            : False

Name                : WeeklyScrubbing
Enabled             : True
Type                : Scrubbing
Days                : Saturday
Start               : 3/22/2012 3:45:00 AM
DurationHours       : 0
StopWhenSystemBusy  : True
Memory              : 50 %
Priority            : Normal
ScheduledTask       : \Microsoft\Windows\Deduplication\WeeklyScrubbing
Full                : False
ReadOnly            : False
```

A note about File Explorer properties

When you use File Explorer to view the properties of a deduplicated file, you will see that **Size** shows the logical size of the file, and **Size on Disk** shows the true physical allocation to that file. **Size on Disk** is less than the actual file size because deduplication has moved the contents of the file to a common chunk store and replaced the original file with an NTFS reparse point stub and metadata. Use the **Measure-DedupFileMetadata** Windows PowerShell cmdlet to determine how much space could be freed if you deleted particular files or folders.

How do I tell if deduplication is keeping up with the rate of incoming data?

For most deployments, the default settings will be appropriate to process existing data and keep up with the incoming data. In some environments, you may want to monitor deduplication if additional jobs need to be scheduled. The easiest way is to use the **Get-DedupStatus** Windows PowerShell cmdlet to monitor the number of optimized files compared to the number of in-policy files. This enables you to see if all in-policy files are processed, for example:

```
PS X:\> Get-DedupStatus

FreeSpace    SavedSpace   OptimizedFiles     InPolicyFiles      Volume
```

```
---------    ----------    --------------    -------------      ------
222.14 GB    528.93 GB     36130             36126              X:
151.26 GB    84.19 GB      43017             43017              Z:
```

If the number of in-policy files is continuously rising faster than the number of optimized files, you might consider running additional jobs.

Monitoring the event log can also be helpful to understand deduplication events and status. To view deduplication events, in File Explorer, navigate to **Applications and Services Logs**, click **Microsoft**, click **Windows**, and then click **Deduplication**.

If the value **LastOptimizationResult = 0x00000000** appears in the **Get-DedupStatus |fl** Windows PowerShell results, the entire dataset was processed by the previous optimization job.

DFS Namespaces and DFS Replication Overview

This topic discusses the DFS Namespaces (DFSN or DFS-N) and DFS Replication (DFSR or DFS-R) role services, which together comprise Distributed File System (DFS) in Windows Server 2012 R2 and Windows Server 2012. This topic discusses how to install DFS, what's new, and where to find evaluation and deployment information.

Did you mean…

- What's New in DFS Replication in Windows Server 2012 R2
- DFS Management (online Help)
- DFS Step-by-Step Guide for Windows Server 2008
- File and Storage Services Overview
- File Services Windows Server Protocols

Role service descriptions

DFS Namespaces and DFS Replication are role services in the File and Storage Services role.

- **DFS Namespaces** Enables you to group shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites.
- **DFS Replication** Enables you to efficiently replicate folders (including those referred to by a DFS namespace path) across multiple servers and sites. DFS Replication uses a compression algorithm known as remote differential compression (RDC). RDC detects changes to the data in a file, and it enables DFS Replication to replicate only the changed file blocks instead of the entire file.

Practical applications

You can use DFS Namespaces and DFS Replication to achieve the following goals:

- **DFS Namespaces** Provide users with a centralized folder namespace through which to access and store files. You can host the underlying file shares on different servers and in different sites to improve availability and performance.
- **DFS Replication** Synchronize folder contents efficiently between servers, across LAN or WAN network connections.

Note

DFS Replication can also be used to replicate the SYSVOL folder in Active Directory Domain Services (AD DS) in domains that use the Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 domain functional levels.

New and changed functionality

The following table describes the changes in DFS Namespaces and DFS Replication functionality.

For more information about what's new in DFS Replication in Windows Server 2012 R2, see What's New in DFS Replication in Windows Server 2012 R2; for Windows Server 2012, see What's New in DFS Namespaces and DFS Replication in Windows Server 2012.

Feature/functionality	New or updated?	Description
Windows PowerShell module for DFS Replication	New in Windows Server 2012 R2	Provides Windows PowerShell cmdlets for performing the majority of administration tasks for DFS Replication, as well as new functionality.
DFS Replication: Windows Management Infrastructure provider	New in Windows Server 2012 R2	Provides the latest Windows Management Infrastructure (WMI)-based methods to manage DFS Replication.
DFS Replication: Database cloning for initial sync	New in Windows Server 2012 R2	Provides support for bypassing initial replication when creating new replicated folders, replacing servers, or recovering from a disaster.
DFS Replication: Database corruption recovery	New in Windows Server 2012 R2	Provides support for rebuilding corrupt databases without unexpected data loss caused by

Feature/functionality	New or updated?	Description
		non-authoritative initial sync.
DFS Replication: Cross-file RDC disable	New in Windows Server 2012 R2	Provides the option to disable cross-file remote differential compression (RDC) between servers.
DFS Replication: File staging tuning	New in Windows Server 2012 R2	Provides the option to configure variable file staging sizes on individual servers.
DFS Replication: Preserved file restoration	New in Windows Server 2012 R2	Provides the capability to restore files from the ConflictAndDeleted and PreExisting folders.
DFS Replication: Database dirty shutdown recovery	Updated in Windows Server 2012 R2	Enables automatic recovery after a loss of power or an unexpected stopping of the DFS Replication service.
DFS Replication: Membership disable	Updated in Windows Server 2012 R2	Stops DFSR private folder cleanup when disabling a server's membership in a replicated folder.
Windows PowerShell module for DFS Namespaces	New in Windows Server 2012	Provides Windows PowerShell cmdlets for performing the majority of administration tasks for DFS Namespaces.
DFS Namespaces: Site awareness for DirectAccess clients	New in Windows Server 2012	Provides remote client computers with optimal referrals to corporate resources when the computers connect by using DirectAccess to the corporate network.
DFS Namespaces: Windows Management Infrastructure provider	New in Windows Server 2012	Provides the latest Windows Management Infrastructure-based management methods to manage DFS Namespaces.
DFS Replication: Support for Data Deduplication volumes	New in Windows Server 2012	Provides support for replicated folders stored on volumes that use Data Deduplication.

Removed or deprecated functionality

The following features are included in Windows Server 2012 R2 and Windows Server 2012, but they are being phased out, and they likely will be removed from future versions of the Windows Server operating system.

Deprecated feature	Replacement
DFS Namespaces command-line tool, Dfscmd	DFS Namespaces module for Windows PowerShell
File Replication Service (FRS)	DFS Replication In Windows Server 2012 R2, it is no longer possible to use Windows PowerShell or Server Manager to create new domains with a Windows Server 2003 domain functional level. This means that new FRS deployments are now blocked and DFS Replication will always be used for SYSVOL replication in new domains.

For a complete list of deprecated features in Windows Server 2012, see Features Removed or Deprecated in Windows Server 2012.

For more information about FRS support in Windows Server 2012 and Windows Server 2008 R2, see File Replication Service (FRS) Is Deprecated in Windows Server 2008 R2.

Requirements for running DFS

There are no additional hardware or software requirements for running DFS Management or using DFS Namespaces.

Before you can deploy DFS Replication, you must configure your servers as follows:

- Update the Active Directory Domain Services (AD DS) schema to include Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 R2 schema additions. (If you install a domain controller running Windows Server 2012 the schema is automatically updated.) You cannot use read-only replicated folders with the Windows Server 2003 R2 or older schema additions.
- Ensure that all servers in a replication group are located in the same forest. You cannot enable replication across servers in different forests.
- Install DFS Replication on all servers that will act as members of a replication group.
- Contact your antivirus software vendor to check that your antivirus software is compatible with DFS Replication.

- Locate any folders that you want to replicate on volumes formatted with the NTFS file system. DFS Replication does not support the Resilient File System (ReFS) or the FAT file system. DFS Replication also does not support replicating content stored on Cluster Shared Volumes.

Server Manager information

DFS Namespaces and DFS Replication are a part of the File and Storage Services role. The management tools for DFS (DFS Management, the DFS Namespaces module for Windows PowerShell, and command-line tools) are installed separately as part of the Remote Server Administration Tools.

Installing DFS Namespaces, DFS Replication, and DFS Management

To install the role services and the DFS Management Tools, use one of the following methods.

- Use Server Manager
- Use Windows PowerShell

To install DFS by using Server Manager

1. Open Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard appears.
2. On the **Server Selection** page, select the server or virtual hard disk (VHD) of an offline virtual machine on which you want to install DFS.
3. Select the role services and features that you want to install.
   - To install the DFS Namespaces and DFS Replication services, on the **Server Roles** page, select **DFS Namespaces** and **DFS Replication**.
   - To install only the DFS Management Tools, on the **Features** page, expand **Remote Server Administration Tools**, **Role Administration Tools**, expand **File Services Tools**, and then select **DFS Management Tools**.

     DFS Management Tools installs the DFS Management snap-in, the DFS Namespaces module for Windows PowerShell, and command-line tools, but it does not install any DFS services on the server.

To install DFS by using Windows PowerShell

- Open a Windows PowerShell session with elevated user rights, and then type the following command, where <name> is the role service or feature that you want to install (see the following table for a list of relevant role service or feature names):

```
Install-WindowsFeature <name>
```

Role service or feature	Name
DFS Namespaces	FS-DFS-Namespace

DFS Replication	`FS-DFS-Replication`
DFS Management Tools	`RSAT-DFS-Mgmt-Con`

For example, to install the Distributed File System Tools portion of the Remote Server Administration Tools feature, type:

```
Install-WindowsFeature RSAT-DFS-Mgmt-Con
```

To install the DFS Namespaces, DFS Replication, and the Distributed File System Tools portions of the Remote Server Administration Tools feature, type:

```
Install-WindowsFeature FS-DFS-Namespace, FS-DFS-Replication,
RSAT-DFS-Mgmt-Con
```

Interoperability with Windows Azure virtual machines

Using DFS Namespaces and DFS Replication on a virtual machine in Windows Azure has been tested with Windows Server 2012 R2; however, there are some limitations and requirements that you must follow.

For DFS Namespaces:

- You can't cluster stand-alone namespaces in Windows Azure virtual machines.
- You can host domain-based namespaces in Windows Azure virtual machines, including environments with Windows Azure AD, though a single namespace can't encompass both on premise namespace servers and namespace servers hosted in Azure VMs, even when using Active Directory Federation Services.

For DFS Replication:

- Using snapshots or saved states to restore a server running DFS Replication for replication of anything other than the SYSVOL folder causes DFS Replication to fail, which requires special database recovery steps. Similarly, don't export, clone, or copy the virtual machines. For more information, see article 2517913 in the Microsoft Knowledge Base, as well as Safely Virtualizing DFSR.
- When backing up data in a replicated folder hosed in a virtual machine, you must use backup software from within the guest virtual machine.
- DFS Replication requires access to physical or virtualized domain controllers – it can't communicate directly with Windows Azure Active Directory (Windows Azure AD).
- DFS Replication requires a VPN connection between your on premise replication group members and any members hosted in Azure VMs. You also need to configure the on premise router (such as Forefront Threat Management Gateway) to allow the RPC Endpoint Mapper (port 135) and a randomly assigned port between 49152 and 65535 to pass over the VPN connection. You can use the **Set-DfsrMachineConfiguration** cmdlet or the **Dfsrdiag** command line tool to specify a static port instead of the random port. For more information about how to specify a static port for DFS Replication, see Set-DfsrServiceConfiguration. For

information about related ports to open for managing Windows Server, see article 832017 in the Microsoft Knowledge Base.

To learn about how to get started with Windows Azure virtual machines, visit the Windows Azure web site.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	What's New in DFS Replication in Windows Server 2012 R2 (Windows Server 2012 R2) | What's New in DFS Namespaces and DFS Replication in Windows Server 2012 (Windows Server 2012) | What's New in Distributed File System in Windows Server 2008 R2
Deployment	SYSVOL Replication Migration Guide: FRS to DFS Replication | DFS Operations Guide: Migrating from FRS to DFS Replication | DFS Namespace Scalability Considerations
Operations	DFS Namespaces: Frequently Asked Questions | DFS Replication: Frequently Asked Questions | Administering DFS-Replicated SYSVOL | Windows Server 2008 Command Reference
Troubleshooting	Understanding DFSR Dirty (Unexpected) Shutdown Recovery | Types of DFS Replication Errors | DFSR Event 1008 (DFS Replication) | DFSR Event 1202 (DFS Replication) | DFSR Event 5002 (DFS Replication) | DFSR Event 5012 (DFS Replication) | DFSR Event 6002 (DFS Replication)
Community resources	The File Services and Storage TechNet Forum | Storage Team at Microsoft File Cabinet Blog | Ask the Directory Services Team Blog | Jose Barreto's Blog |DFS Replication Survival Guide (TechNet Wiki) | Windows Server Information Experience on Twitter
Related technologies	Folder Redirection, Offline Files, and Roaming User Profiles overview | Failover Clustering

Content type	References
	Overview

What's New in DFS Replication in Windows Server 2012 R2

This topic describes the features that were added to Distributed File System (DFS) Replication (DFSR or DFS-R) in Windows Server 2012 R2.

For information about the features that were added in Windows Server 2012, see What's New in DFS Namespaces and DFS Replication in Windows Server 2012.

Role service description

DFS Replication is a role service in the File and Storage Services role. It enables you to efficiently replicate folders (including those referred to by a DFS namespace path) across multiple servers and sites. DFS Replication uses a compression algorithm known as remote differential compression (RDC). RDC detects changes to the data in a file, and it enables DFS Replication to replicate only the changed file blocks instead of the entire file.

New and changed functionality

The following table describes the changes in DFS Replication functionality that are available in this release.

Feature/functionality	New or updated?	Description
Windows PowerShell module for DFS Replication	New	Provides Windows PowerShell cmdlets for performing the majority of administrative tasks for DFS Replication, in addition to new functionality.
DFS Replication Windows Management Infrastructure provider	New	Provides the latest Windows Management Infrastructure (WMI)-based methods to manage DFS Replication.
Database cloning for initial sync	New	Provides support for bypassing initial replication when creating new replicated folders, replacing servers, or recovering from a

Feature/functionality	New or updated?	Description
		disaster.
Database corruption recovery	New	Provides support for rebuilding corrupt databases without unexpected data loss caused by nonauthoritative initial sync.
Cross-file RDC disable	New	Provides the option to disable cross-file remote differential compression (RDC) between servers.
File staging tuning	New	Provides the option to configure variable file staging sizes on individual servers.
Preserved file restoration	New	Provides the capability to restore files from the **ConflictAndDeleted** and **PreExisting** folders.
Unexpected shutdown database recovery improvements	Updated	Enables automatic recovery after a loss of power or an unexpected stoppage of the DFS Replication service.
Membership disabling improvements	Updated	Stops DFS Replication private folder cleanup when disabling a server's membership in a replicated folder.

Windows PowerShell module for DFS Replication

The DFS Replication module provides Windows PowerShell cmdlets for performing the majority of administrative tasks for DFS Replication, in addition to the new functionality described in this topic.

What value does this change add?

Administrators can use the extensive Windows PowerShell cmdlets to perform common administrative tasks, and optionally automate them by using Windows PowerShell scripts. These tasks include operational actions such as creating, modifying, and removing replication settings. New functionality is also included with the cmdlets, such as the ability to clone DFS Replication databases and restore preserved files.

What works differently?

Instead of using DFS Management or DFS Replication command-line tools, administrators can perform all common tasks by using the Windows PowerShell cmdlets.

The Windows PowerShell cmdlets are available on computers that run Windows Server 2012 R2 or Windows 8.1 and that have the DFS Management Tools feature (part of the Remote Server Administration Tools) installed.

Tip

To use the Windows PowerShell module for DFS Replication from a computer that doesn't have the module installed, use the **Enter-PSSession** or **Invoke-Command** cmdlets to establish a session with a computer that has the DFS Management Tools feature installed.

DFS Replication Windows Management Infrastructure provider

Windows Server 2012 R2 includes new Windows Management Infrastructure (sometimes referred to as WMI v2) provider functionality, which provides programmatic access to manage DFS Replication.

What value does this change add?

Management programs can use the latest Windows Management Infrastructure-based methods to manage DFS Replication.

What works differently?

Windows Management Infrastructure-based management access occurs over a firewall friendly Windows Remote Management (WinRM) transport protocol. The WMI v1 namespace is still available for backwards compatibility.

Database cloning for initial sync

Windows Server 2012 R2 includes new database cloning functionality, which can accelerate initial replication when you create new replicated folders, replace servers, or recover from disaster.

What value does this change add?

You can now export a DFS Replication database from a volume on one server, and then preseed replicated files and import the database on multiple servers. The replicated folders begin replicating with a greatly reduced initial setup time (the setup time is usually reduced approximately 99%).

What works differently?

In Windows Server 2012 and earlier operating systems, initial sync replication required that a destination server populate its database over the network through a resource-expensive version-vectoring process with the source server. For larger datasets, this can take considerable time (days or weeks), even when you preseed data on the destination server.

Now the **Export-DfsrClone** cmdlet allows you to export the DFS Replication database and volume configuration .xml file settings for a given volume from the local computer to clone that

database on another computer. Running the cmdlet triggers the export in the DFS Replication service and then waits for the service to complete the operation. During the export, DFS Replication stores file metadata in the database as part of the validation process. After you preseed the data and copy the exported database and .xml file to the destination server, you use **Import-DfsrClone** to import the database to a volume and validate the files in the file system. Any files that perfectly match don't require expensive interserver metadata and file synchronization, which leads to dramatic performance improvements during the initial sync.

For more information, see Export a Clone of the DFS Replication Database.

Database corruption recovery

Windows Server 2012 R2 provides support for rebuilding corrupt databases without unexpected data loss caused by a nonauthoritative initial sync.

What value does this change add?

When DFS Replication detects database corruption, it rebuilds the database and then resumes replication normally, with no files arbitrarily losing conflicts. When replicating with a Read-only partner, DFS Replication resumes replication without waiting indefinitely for an administrator to manually set the primary flag.

What works differently?

Previously, a corrupt database would trigger DFS Replication to delete the database and start the nonauthoritative initial sync process again, as if replication was being set up for the first time. Any files on the recovering server would lose all conflicts automatically, even if they were the latest version of the file. Those conflicts would move into the **ConflictAndDeleted** or **PreExisting** folders, leading to perceived or real data loss.

In Windows Server 2012 R2, when DFS Replication detects database corruption, it rebuilds the database by using local file and update sequence number (USN) change journal information, and then marks each file with a Normal replicated state. DFS Replication then contacts its partner servers and merges the changes, which allows the last writer to save the most recent changes as if this was normal ongoing replication.

Cross-file RDC disable

Windows Server 2012 R2 provides the option to disable cross-file remote differential compression between servers.

What value does this change add?

You can now specifically choose to use the cross-file remote differential compression (RDC) capability, depending on your data and network topologies. For servers on LANs, turning off cross-file RDC may reduce server resource overhead and increase replication performance.

What works differently?

In Windows Server 2012, DFS Replication always enables cross-file RDC. Cross-file RDC uses up to five existing previously replicated files on a volume to seed a new replicating file. Applying cross-file RDC over low-bandwidth network connections with files that are similar results in large

bandwidth savings and potentially large time savings. However, when you use cross-file RDC on high-bandwidth network connections, cross-file RDC might increase local processing time and negatively affect performance. In extremely large datasets (millions of files on a volume with a great deal of similarity), cross-file RDC might also negatively affect CPU and disk utilization.

In Windows Server 2012 R2, DFS Replication allows you to choose whether to use cross-file RDC on a per-connection basis between partners. Disabling cross-file RDC might increase performance at the cost of higher bandwidth usage.

File staging tuning

Windows Server 2012 R2 provides the option to configure variable file staging sizes on individual servers.

What value does this change add?

You can now choose a minimum file size for a file to stage if you have not configured RDC for a smaller size. For servers on LANs with larger files, increasing the minimum staging size for files can increase replication performance.

What works differently?

In Windows Server 2012 and earlier operating systems, DFS Replication uses a hard-coded 256 KB file size to determine staging requirements. If RDC is enabled and the RDC minimum size (by default, this is 64 KB) is larger than 256 KB, a file will be staged before it replicates. Usually this means that files smaller than 64 KB don't get staged on servers running Windows Server 2012 and earlier operating systems. File staging adds replication time to allow for performing RDC operations and to lower the chance of files not replicating because they are in use by applications that created file locks on them.

DFS Replication now allows you to configure the staging minimum size from as little as 256 KB to as large as 512 TB. When you are not using RDC or staging, files are no longer compressed or copied to the staging folder, which can increase performance at the cost of much higher bandwidth usage.

Preserved file restoration

Windows Server 2012 R2 provides the ability to restore files from the **ConflictAndDeleted** and **PreExisting** folders.

What value does this change add?

This allows recovery of obfuscated user data from hidden DFS Replication private folders.

What works differently?

Windows Server 2012 and earlier operating systems don't provide any tools to recover these files.

DFS Replication now enables you to inventory and retrieve the conflicted, deleted, and preexisting files by using the **Get-DfsrPreservedFiles** and **Restore-DfsrPreservedFiles** cmdlets. You can restore these files and folders to their previous location or to a new location.

You can choose to move or copy the files, and you can keep all versions of a file or only the latest version.

Unexpected shutdown database recovery improvements

Windows Server 2012 R2 enables automatic recovery after a loss of power or an unexpected stoppage of the DFS Replication service.

What value does this change add?

When DFS Replication detects an unexpected database shutdown (for example, after a power outage or service termination), it automatically validates the database against the file system and then resumes replication normally, settling any file conflicts normally.

What works differently?

In Windows Server 2012 and Windows Server 2008 R2, an unexpected shutdown required you to re-enable replication manually by using a WMI method. You could disable this default behavior by using a per-computer registry value.

When DFS Replication detects an unexpected shutdown in Windows Server 2012 R2, it defaults to triggering the automatic recovery process. You must opt out of this behavior by using the registry value. In addition, if the only replicated folder on a volume is the built-in SYSVOL folder of a domain controller, it automatically triggers recovery regardless of the registry setting.

Membership disabling improvements

Windows Server 2012 R2 stops DFS Replication private folder cleanup when disabling a server's membership in a replicated folder.

What value does this change add?

Files that had previously moved to the **ConflictAndDeleted** or **PreExisting** folders are no longer deleted when you disable a server's replication group membership. Additionally, the message that is displayed by the management tools now states the processing that occurs after memberships are disabled. The message also explains that re-enabling a membership starts nonauthoritative synchronization.

What works differently?

In Windows Server 2012 and earlier operating systems, disabling a membership immediately deleted the **DfsrPrivate** folder for that membership, including the **Staging**, **ConflictAndDeleted**, and **PreExisting** folders. After these folders are deleted, you can't easily recover data from them without reverting to a backup.

DFS Replication now leaves the **DfsrPrivate** folder untouched when you disable a membership. You can recover conflicted, deleted, and preexisting files from that location if the membership is not re-enabled. (Enabling the membership deletes the content of all private folders.)

Removed or deprecated functionality

The following features are included in Windows Server 2012 R2, but they are being phased out, and they likely will be removed from future versions of the Windows Server operating system.

Deprecated feature	Replacement
DFS Namespaces command-line tool, Dfscmd	DFS Namespaces module for Windows PowerShell
File Replication Service (FRS)	DFS Replication In Windows Server 2012 R2, it is no longer possible to use Windows PowerShell or Server Manager to create new domains with a Windows Server 2003 domain functional level. This means that new FRS deployments are blocked, and DFS Replication is always used for SYSVOL replication in new domains. To use FRS to replicate SYSVOL on domain controllers running Windows Server 2012 R2, the domain controllers must belong to an existing domain that uses the Windows Server 2003 domain functional level.

See also

For additional related information, see the following resources.

Content Type	References
Product evaluation	• What's New in DFS Replication in Windows Server 2012 R2 (Windows Server 2012 R2) • What's New in DFS Namespaces and DFS Replication in Windows Server 2012 (Windows Server 2012) • What's New in Distributed File System in Windows Server 2008 R2
Related technologies	• Folder Redirection, Offline Files, and Roaming User Profiles overview • Failover Clustering Overview

What's New in DFS Namespaces and DFS Replication in Windows Server 2012

This topic describes the features that were added to DFS Namespaces (DFS or DFS-N) and DFS Replication (DFSR or DFS-R) in Windows Server 2012.

For information about the features that were added in Windows Server 2012 R2, see What's New in DFS Namespaces and DFS Replication in Windows Server 2012.

Role service descriptions

DFS Namespaces and DFS Replication are role services in the File and Storage Services role. DFS Namespaces enables you to group shared folders that are located on different servers into one or more logically structured namespaces. DFS Replication enables you to efficiently replicate folders (including those referred to by a DFS namespace path) across multiple servers and sites.

New and changed functionality

The following table describes the changes in DFS Namespaces and DFS Replication functionality.

Feature/functionality	Description
Windows PowerShell module for DFS Namespaces	Provides Windows PowerShell cmdlets for performing the majority of administration tasks for DFS Namespaces.
DFS Namespaces: Site awareness for DirectAccess clients	Provides remote client computers with optimal referrals to corporate resources when the computers connect by using DirectAccess to the corporate network.
DFS Namespaces: Windows Management Infrastructure provider	Provides the latest Windows Management Infrastructure-based management methods to manage DFS Namespaces.
DFS Replication: Support for Data Deduplication volumes	Provides support for replicated folders stored on volumes that use Data Deduplication.

Windows PowerShell module for DFS Namespaces

Windows Server 2012 includes Windows PowerShell cmdlets for performing the majority of administration tasks for DFS Namespaces.

What value does this change add?

Administrators can use the extensive Windows PowerShell cmdlets to perform common administration tasks, and optionally automate them by using Windows PowerShell scripts.

What works differently?

Instead of using DFS Management or DFS command-line utilities, administrators can perform all common tasks by using the Windows PowerShell cmdlets.

The Windows PowerShell cmdlets are available on computers that run Windows Server 2012 or Windows 8 and that have the DFS Management Tools feature (part of the Remote Server Administration Tools) installed.

For more information, see DFS Namespace (DFSN) Cmdlets in Windows PowerShell and Introducing DFS Namespaces Windows PowerShell Cmdlets.

Tip

To use the Windows PowerShell module for DFS Namespaces from a computer that does not have the module installed, use the **Enter-PSSession** cmdlet to establish a session with a computer that has the DFS Management Tools feature installed.

DFS Namespaces: Site awareness for DirectAccess clients

Windows Server 2012 includes site awareness for DirectAccess functionality to provide optimal site-aware referrals to remote users who access a corporate network by using DirectAccess.

What value does this change add?

Remote users who connect to DFS namespaces over a DirectAccess connection receive referrals to the namespace servers and folder targets that are closest to their location. This reduces network latency when accessing files on the corporate network, making more efficient use of bandwidth.

What works differently?

When accessing a DFS namespace path by using DirectAccess in Windows 7 or Windows Server 2008 R2, remote computers that have IP addresses outside of the sites specified in Active Directory Domain Services receive a randomly ordered referral that could include servers in distant sites, even when servers in a nearby site are available.

When a remote computer accesses a DFS namespace path by using DirectAccess on computers running Windows 8 or Windows Server 2012, the computer provides a site name in the referral request to the namespace server running Windows Server 2012. The namespace server uses the site name to provide a referral to the closest site available.

For site awareness to work for clients that are using DirectAccess, the client computer must be running Windows 8 or Windows Server 2012 and the server must be running Windows Server 2012.

DFS Namespaces: Windows Management Infrastructure provider

Windows Server 2012 includes new Windows Management Infrastructure (the new version of WMI, sometimes known as MI or WMIv2) provider functionality, which provides programmatic access to manage DFS Namespaces.

What value does this change add?

Management programs can use the latest Windows Management Infrastructure-based methods to manage DFS Namespaces.

What works differently?

Windows Management Infrastructure-based management access occurs over a firewall-friendly Windows Remote Management (WinRM) transport protocol.

DFS Replication: Support for Data Deduplication volumes

Windows Server 2012 includes support for replicated folders stored on volumes on which Data Deduplication is enabled.

What value does this change add?

You can use the new Data Deduplication functionality to reduce storage requirements without impacting DFS Replication.

What works differently?

DFS Replication replicates the new reparse point tag IO_REPARSE_TAG_DEDUP.

Removed or deprecated functionality

The following features are included in Windows Server 2012, but they are being phased out, and they likely will be removed from future versions of the Windows Server operating system.

Deprecated feature	Replacement
DFS Namespaces command-line tool, Dfscmd	DFS Namespaces module for Windows PowerShell
File Replication Service (FRS)	DFS Replication

For a complete list of deprecated features in Windows Server 2012, see Features Removed or Deprecated in Windows Server 2012.

For more information about FRS support in Windows Server 2012 and Windows Server 2008 R2, see File Replication Service (FRS) Is Deprecated in Windows Server 2008 R2.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	What's New in DFS Replication in Windows Server 2012 R2 | What's New in Distributed File System in Windows Server 2008 R2
Deployment	SYSVOL Replication Migration Guide: FRS to DFS Replication | DFS Operations Guide: Migrating from FRS to DFS Replication | DFS Namespace Scalability Considerations
Operations	DFS Namespaces: Frequently Asked Questions | DFS Replication: Frequently Asked Questions | Administering DFS-Replicated SYSVOL | Windows Server 2008 Command Reference
Troubleshooting	Understanding DFSR Dirty (Unexpected) Shutdown Recovery | Types of DFS Replication Errors | DFSR Event 1008 (DFS Replication) | DFSR Event 1202 (DFS Replication) | DFSR Event 5002 (DFS Replication) | DFSR Event 5012 (DFS Replication) | DFSR Event 6002 (DFS Replication)
Community resources	The File Services and Storage TechNet Forum | Storage Team at Microsoft File Cabinet Blog | Ask the Directory Services Team Blog | Jose Barreto's Blog |DFS Replication Survival Guide (TechNet Wiki) | Windows Server Information Experience on Twitter
Related technologies	Folder Redirection, Offline Files, and Roaming User Profiles overview | Failover Clustering Overview

Export a Clone of the DFS Replication Database

This topic describes how to export a clone of the Distributed File System Replication (DFS Replication) database, which is a new capability introduced in Windows Server 2012 R2.

You can remove the need for members of a replication group to exchange files or metadata during replication setup by preseeding a replicated folder and exporting a clone of the database.

This can reduce the initial synchronization time by up to 99% under ideal circumstances, and it's useful when you configure replication for the first time, add new replication partners, replace server hardware, or recover from loss or corruption of the DFS Replication database.

This document assumes that you have a basic knowledge of DFS Replication. For more information, see Overview of DFS Replication.

Tip
To maximize cloning success, test the procedures in your lab prior to attempting production changes.

Prerequisites

The following requirements are necessary on the source server and the destination server:

- The DFS Replication role service must be installed and running.
- You can perform only one cloning operation at a time per server.
- You must be signed in as a member of the Domain Administrators security group or the Enterprise Administrators security group, or you must have delegated permissions to create replication groups. For more information, see Delegate the Ability to Manage DFS Replication.

The following requirements are necessary on the source server from which you want to export a database clone:

- The DFS Replication database must be on a volume that contains at least one Read-write replicated folder. This is because SYSVOL and Read-only replicated folders are not included in the database export process.
- The volume you are exporting must contain replicated folders that are in the **Normal** state. The replicated folders in that database cannot be in initial sync, initial building, or recovery states.
- Use only approved preseeding techniques to copy files in replicated folders to the destination server.
- Confirm that the preseeding worked properly by using the **Get-DfsrFileHash** cmdlet. For more information on preseeding replicated folders, see DFS Replication: Copying Files to Preseed or Stage Initial Synchronization.
- To maximize cloning performance and data fidelity, prevent all user changes to files on the source server during the cloning operation. The more changes that occur during cloning, the more DFS Replication must update the destination server.
- Export a database clone only when it is acceptable that replication will pause. Cloning takes the DFS Replication volume manager offline, which prevents any replication from running on that volume during cloning. When cloning completes, DFS Replication automatically resumes replication.

The following requirements are necessary on the destination server on which you want to import a database clone:

- The destination volume cannot have a replicated folder on it.

- If the destination volume previously contained a replicated folder, delete the following hidden folder on the volume: <volume>:\system volume information\dfsr\database_<string>.
- To export database clones to a hub server with numerous replication partners, you might have to use NTFS volume mount points on the hub server to create a separate volume for each imported database. This is because each volume can only host one DFS Replication database, and cloning can't merge an existing database with an imported database.
- Prevent all Write access to existing file shares on the destination replication folders by changing the share permissions to Read-only for all users. Don't create any new file shares for the folders until cloning and the post-cloning initial sync processing finishes. As an alternative to changing file share permissions, you can note and remove the file shares. As with any initial sync, users shouldn't use a destination server until all the files are available because the appearance of missing or out-of-date files could lead them to contact the Help Desk, or they could attempt to fix the apparent issues themselves.

Note
You can clone a clone. That is, you can create a clone, import the cloned database on another server, and then clone that database. You can also import a cloned database on many servers.

Note
The supported limit for files per volume is 70 million, with 120 million files per server supported in Windows Server 2012 R2.

Step 1: Create an example replication environment

This procedure shows you how to create an example replication environment to test DFS Replication database cloning. To clone an existing server that already performs replication, skip this section. Cloning supports multiple destination servers, but this example includes only a single destination server.

This example uses SRV01 as the source (upstream or authoritative) server, SRV02 as the destination (downstream or nonauthoritative) server, and drive H on both servers. It also uses a replication group named RG01 and a replicated folder named RF01. All servers must run Windows Server 2012 R2.

To create an example replication environment for testing

1. On the source server *SRV01*, create a folder on drive H named *RF01*.
2. Create or copy in some test files to the H:\RF01 folder, such as by copying the 1,000 largest files in the C:\Windows\sysWOW64 folder.
3. Create the *RG01* test replication group, the *RF01* test replicated folder, and add only *SRV01* as a member by running the following commands from an elevated Windows PowerShell session:

```
$DfsrComputerName = "SRV01"
$DfsrReplicatedFolderPath = "H:\RF01"
$DfsrReplicationGroupName = "RG01"
$DfsrReplicatedFolderName = "RF01"

New-DfsReplicationGroup -GroupName $DfsrReplicationGroupName
New-DfsReplicatedFolder -GroupName $DfsrReplicationGroupName
-FolderName $DfsrReplicatedFolderName
Add-DfsrMember -GroupName $DfsrReplicationGroupName -
ComputerName $DFSRComputername
Set-DfsrMembership -GroupName $DfsrReplicationGroupName -
FolderName $DfsrReplicatedFolderName -ContentPath
$DfsrReplicatedFolderPath -ComputerName $DFSRComputername -
PrimaryMember $True
```

Important
Do not add SRV02 as a member or create a connection between the servers in this new replication group yet. However, you can add multiple replicated folders on a volume and clone them.

4. Wait for DFS Replication event 4112 in the DFS Replication event log, which indicates that the replication folder initialized successfully as the primary member (there is one event per replicated folder). To view this event, in Event Viewer, navigate to **Applications and Services Logs**, then **DFS Replication**.

   To reduce the wait time, you can force Active Directory Domain Services (AD DS) replication, and then use the DFS Replication cmdlet **Update-DfsrConfigurationFromAD**. For more information, see Force Replication Between Domain Controllers or Synchronize Replication with All Partners.

Step 2: Export a clone of a DFS Replication database

When you have a DFS Replication database to export (from an example replication environment or from a production environment), use the following procedure to export a clone of the database, which is used for all replicated folders on the volume you specify.

This example uses SRV01 as the source server, SRV02 as the destination server, and drive H on both servers. It also uses a replication group named RG01 and a replicated folder named RF01. All servers must run Windows Server 2012 R2.

If you're performing an upgrade installation (instead of a clean install) of Windows Server, you don't need to use cloning or the procedures in this topic.

Tip

To maximize cloning performance and data fidelity, prevent all user changes to files on the source server during the cloning operation. The more changes that occur during cloning, the more DFS Replication must update the destination server.

To export a clone of a DFS Replication database

1. Validate that all the existing replicated folders on the volume that stores the replicated folder that you want to preseed are in the **Normal**, noninitial sync state. Replicated folders in other states are skipped during cloning.

   To validate the state, examine the DFS Replication event log to ensure that all replicated folders have a 4112 or 4104 event. You can also use the following sample command to determine which replicated folders are in the **Normal**, noninitial sync state:

   ```
   Get-WmiObject -Namespace "root\Microsoft\Windows\DFSR" -Class
   msft_dfsrreplicatedfolderinfo -ComputerName srv01
   | ft replicatedfoldername,state -auto -wrap
   ```

   All replicated folders listed for that computer will show state 4 (that is, **Normal**) if they have completed the initial build or initial sync, and they are ready to clone.

2. Export the cloned database and volume configuration XML by running the following sample command from an elevated Windows PowerShell session:

   ```
   $DfsrCloneVolume = "H:"
   $DfsrCloneDir = "\DfsrClone"

   New-Item -Path $DfsrCloneVolume\$DfsrCloneDir -Type Directory
   Export-DfsrClone -Volume $DfsrCloneVolume -Path
   $DfsrCloneVolume\$DfsrCloneDir
   ```

   Note

   Although the **Export-DfsrClone** cmdlet does not return any output until cloning completes, you can safely close the Windows PowerShell session after running the command. The DFS Replication service performs the processing while the command synchronously waits for the result. If you close the Windows PowerShell console or exit the command, you can continue to see progress by examining the DFS Replication event log, the DFS Replication debug logs, or by using the **Get-DfsrCloneState** cmdlet.

3. After completion, make a note of the Robocopy sample commands that are displayed by the **Export-DfsrClone** cmdlet. You can later use these commands to copy the data to the destination server. You can also use **Export-DfsrClone** with the **-WhatIf** parameter to see the Robocopy samples without performing cloning, for example:

   ```
   Robocopy.exe "H:\DfsrClone" "<destination path>" /B
   Robocopy.exe "H:\RF01" "<destination path>" /E /B /COPYALL
   ```

```
/R:6 /W:5 /MT:64 /XD DfsrPrivate /TEE /LOG+:preseed.log
```

4. Wait for a DFS Replication Event 2402 in the DFS Replication event log, which indicates that the export completed successfully. You can also use the **Get-DfsrCloneState** cmdlet to see the export status. The cmdlet returns **Ready** when the export process completes.

Step 3: Import a clone of a DFS Replication database

After you export a clone of a DFS Replication database, use the following procedure to import it on the destination server.

Import the cloned database on the destination server

1. Use an approved method to preseed the destination server with the contents of the replicated folders from the source server. If the destination server is the same as the source server (for example, if you are replacing hardware or installing a new operating system), or you are recovering from a corrupted DFS Replication database, skip this step. In these instances, there is no need to remove the previously replicated data.

   For example, you could use the following Robocopy sample command with the instructions in the following article: Replacing DFSR Member Hardware or OS (Part 2: Preseeding):

   ```
   Robocopy.exe "H:\RF01" "\\Srv02\H$\Rf01" /E /B /COPYALL /R:6
   /W:5 /MT:64 /XD DfsrPrivate /TEE /LOG+:preseed.log
   ```

2. Copy to the destination server the exported dfsr.db.clone database and config.xml files that you created with the **Export-DfsrClone** cmdlet in Step 2 of To export a clone of a DFS Replication database. For example, you could use the following Robocopy sample command:

   ```
   Robocopy.exe "H:\DfsrClone" "\\Srv02\H$\DfsrClone" /B
   ```

   Important
   If you are using database cloning to recover from a corrupted database, remove the destination server from the affected replication group membership prior to copying. For more information, see Use cloning to replace a corrupted DFS Replication database.

3. On the destination server *SRV02*, ensure that a DFS Replication database does not already exist on the volume into which you will import data. To do so, use the following procedure:

   a. Launch an elevated Windows PowerShell session and then type the following, where *H* is the drive letter of the appropriate replicated folder:

      ```
      Get-ChildItem -path "H:\System Volume Information\dfsr" -
      hidden
      ```

   b. If there aren't any DFS Replication databases on the volume, there won't be any

output. If a directory is listed, use DFS Management or the following Windows PowerShell command to confirm that there aren't any replicated folders on drive *H:*\ (you cannot clone into an existing DFS Replication database).

```
Get-DfsrMembership | Where-Object -Property ContentPath -Like H:\* | Format-Table
```

c. If the destination volume doesn't currently contain any replicated folders, but it did previously (such as reusing storage when replacing a server), in the elevated Windows PowerShell session, type the following command:

```
Stop-Service DFSR
Remove-Item -path "H:\system volume information\dfsr" -recurse -force
Start-Service DFSR
```

4. After all preseeding operations are complete, import the cloned database and volume configuration XML by using a validation level of basic (1). This provides a balance of speed and data fidelity for the initial sync. Run the following sample command from an elevated Windows PowerShell session on the destination server:

Important

Do not start the import until you complete all preseeding operations.

```
Import-DfsrClone -Volume h: -Path "h:\Dfsrclone"
```

Note

Although the **Import-DfsrClone** cmdlet does not return any output until cloning completes, you can safely close the Windows PowerShell session after running the command. The DFS Replication service performs the processing while the command synchronously waits for the result. If you close the Windows PowerShell console or exit the command, you can continue to see progress by examining the DFS Replication event log, the DFS Replication debug logs, or by using the **Get-DfsrCloneState** cmdlet.

5. Wait for a DFS Replication Event 2404 in the DFS Replication event log, which indicates that the import completed successfully. Event 2416 also contains cloning progress events. You can also use the **Get-DfsrCloneState** cmdlet to see the export status. The cmdlet returns **Ready** when the export process completes.
6. Add *SRV02* to the replication group with a connection to the source server by running the following sample commands from an elevated Windows PowerShell session.

   If the destination server is a production server that is already a member of the same replication group as the source server (and you did not remove all of its memberships while replacing a corrupted database), skip the **Add-DfsrMember** and **Add-DfsrConnection** cmdlets, and use only the **Set-DfsrMembership** cmdlet in the following command.

```
$DfsrSourceComputerName = "Srv01"
```

```
$DfsrDestinationComputerName = "Srv02"
$DfsrReplicationGroupName = "RG01"
$DfsrReplicatedFolderName = "RF01"
$DfsrReplicatedFolderPath = "H:\RF01"

Add-DfsrMember -GroupName $DfsrReplicationGroupName -
ComputerName $DfsrDestinationComputerName
Add-DfsrConnection -GroupName $DfsrReplicationGroupName -
SourceComputerName $DfsrSourceComputerName –
DestinationComputerName $DfsrDestinationComputerName
Set-DfsrMembership -GroupName $DfsrReplicationGroupName -
FolderName $DfsrReplicatedFolderName -ContentPath
$DfsrReplicatedFolderPath -ComputerName
$DfsrDestinationComputerName
```

7. Wait for DFS Replication event 4104 in the DFS Replication event log, which indicates that the replication folder initialized successfully as a nonauthoritative member (there is one event per replicated folder).
8. Open an elevated Windows PowerShell session, and use the **Get-DfsrPreservedFiles** cmdlet to note any files that were conflicted or preexisting on the destination server. If there aren't any, you should see the result **Could not find the specified manifest file**. There should not be any preexisting files, and the only conflicts would be files that were modified after cloning started and before it completed, for example:

```
Get-DfsrPreservedFiles -path
"h:\rf01\dfsrprivate\conflictanddeletedmanifest.xml"
Get-DfsrPreservedFiles -path
"h:\rf01\dfsrprivate\preexistingmanifest.xml"
```

9. Use the **Start-DfsrPropagationTest** and **Write-DfsrPropagationReport** cmdlets to validate replication, for example:

```
Start-DfsrPropagationTest –GroupName "RG01" –FolderName * -
ReferenceComputer SRV01
Write-DfsrPropagationReport –GroupName "RG01" –FolderName * -
ReferenceComputer SRV01 –Path "c:\reports"
```

Use cloning to replace a corrupted DFS Replication database

In the event of DFS Replication database corruption that is caused by hardware issues (such as abrupt power loss), you can save time by cloning the database from another server rather than waiting for the automatic nonauthoritative recovery to complete.

To perform a database recovery by using cloning, follow the procedures to export a database clone and import a database clone, but skip preseeding the replicated folder. Also, prior to performing "Step 3: Import a clone of a DFS Replication database," remove the server from the replicated folder memberships that are affected by the corruption. This prevents DFS Replication from attempting to rebuild the database.

To do so, from an elevated Windows PowerShell session, type the following command to remove only the member. (This removes all of that computer's memberships in that replication group because Windows PowerShell does not support removing a single membership.)

```
Remove-DfsrMember -GroupName RG01 -ComputerName SRV02
```

Important
If you remove a member from a replication group, all of its memberships are affected. Only perform this operation if all the server's memberships exist only on the volume with the corrupt database.

If there are volumes without corrupt databases, you should use the **DFSRADMIN** command instead. Open an elevated command prompt window, and type the following command to remove only the single affected membership:

```
Dfsradmin membership delete /rgname:rg01 /rfname:rf01 /memname:srv02
```

See also

Content Type	References
Product evaluation	• What's New in DFS Replication in Windows Server 2012 R2
Operations	• DFS Management Help • DFS Replication: Frequently Asked Questions • DFS Replication: Copying Files to Preseed or Stage Initial Synchronization • DFS Replication Cmdlets
Community resources	• The File Services and Storage TechNet Forum • Storage Team at Microsoft File Cabinet Blog • Ask the Directory Services Team Blog • DFS Replication Survival Guide (TechNet Wiki) • Windows Server Information Experience on Twitter

DFS Replication: Copying Files to Preseed or Stage Initial Synchronization

The topics in this guide describe how to preseed files for Distributed File System (DFS) Replication (also known as DFSR or DFS-R) to speed up initial synchronization when you set up DFS Replication, add a new replication partner, or replace a server. Preseeding is required when you set up the database cloning feature, which was introduced in Windows Server 2012 R2.

Distributed File System Replication synchronizes files between two or more servers. When you set up replication, the initial synchronization time drops significantly if DFS Replication can use files that are already in place on the destination volume.

DFS Replication was introduced in Windows Server 2003 R2. Methods described in this guide can be used to preseed files for DFS Replication from servers running Windows Server 2003 R2 and later operating systems to any server that supports DFS Replication, including servers running Windows Server 2012 R2 or Windows Server 2012. Not all methods support all of these operating systems; for guidance, see Comparison of tools, later in this topic. The methods can be used to preseed files from folders that are already replicated through DFS Replication or folders that will be replicated for the first time.

Note
This guide assumes that you have a basic knowledge of DFS Replication. For more information, see DFS Namespaces and DFS Replication Overview.

Overview of preseeding files for DFS Replication

DFS Replication is a state-based, multi-master file replication service that uses a database and the NTFS USN journal to track file changes. In its simplest configuration, two servers communicate with each other bidirectionally (writable) or unidirectionally (Read-only). For an individual file update operation, each server can operate in an upstream or downstream fashion.

When configuring replication on the second – and subsequent – servers, DFS Replication completes a two-phase process:

- **Phase 1: Initial build** – DFS Replication scans all existing files and folders in the replicated folder and adds a record for each one to the database.
- **Phase 2: Initial sync** – DFS Replication exchanges replicated file metadata - such as each file hash, unique identifier, and global version sequence number – with its partner server. Inbound, the downstream (destination) server replicates any missing files or files that have a different hash from the file on the upstream (source) server. If DFS Replication finds any files that have a matching hash on the same relative path on both servers, the file is not replicated downstream.

When you preseed file data on the downstream server, DFS Replication does not have to replicate those files inbound during initial synchronization. For larger data sets, this can save significant setup time and allow DFS Replication faster entry into the normal replication state, in which only new, changed, or modified files are replicated between servers.

Note

To speed up initial synchronization when you add a new replication partner or replace a server in an existing DFS Replication setup, you can preseed the replicated folder on the destination server and then export a clone of the DFS Replication database. If you both preseed the replicated folder and export a database clone, members of a replication group do not have to exchange files or metadata during replication setup. For more information about this feature, introduced in Windows Server 2012 R2, see Export a Clone of the DFS Replication Database.

For preseeding to work, however, you must copy files from the upstream servers to the downstream servers without changing their file hashes. A file hash created by DFS Replication incorporates the data stream, the alternate data stream, and the access control list (ACL) applied directly to that file. This ensures that any aspect of file changes replicates to all other nodes. Incorrectly preseeding files on a downstream server leads to slower replication than not preseeding, because DFS Replication takes the additional step of verifying every existing file, but then must discard and replicate each file.

Comparison of tools

The tool that is best suited for your environment depends on the environment, the operating systems on the source and destination servers, and the nature and size of the data sets that you want to preseed for DFS Replication. The following table compares three tools that can be used for preseeding: Robocopy, Windows Server Backup, and NTBackup. This guide provides step-by-step procedures for preseeding files with each of these tools.

Preseeding tool	Advantages	Limitations
Robocopy	• Copies files directly from the source directory to the destination directory. • Can be used to preseed files from Windows Server 2003 R2 and later, through Windows Server 2012 R2.	• Cannot copy files with exclusive locks. To effectively preseed, you must stabilize files on the source computer before copying begins. • Requires manual installation of the correct Robocopy version for the source or destination operating system.
Windows Server Backup	• Easily installed feature of Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. • Copies files with exclusive	• Requires backing up to an interim storage location. • Cannot be used to preseed files from Windows Server 2003 R2. (Must be installed on both the source

Preseeding tool	Advantages	Limitations
	locks.	and destination server.) • On a source server running Windows Server 2008, Windows Server Backup can only back up volumes, not individual folders. However, individual folders can be restored.
Windows NTBackup	• If you are accustomed to using NTBackup, provides a familiar tool for preseeding files from Windows Server 2003 R2 for DFS Replication in Windows Server 2008 R2 or Windows Server 2008. • Copies files with exclusive locks.	• Requires backing up to an interim storage location. • Cannot be used to preseed files for DFS Replication in Windows Server 2012 R2 or Windows Server 2012.

Additional preseeding methods might be available for your storage devices and accompanying storage software:

- **Hardware options** - Some SAN and NAS devices allow synchronous or snapshot-based data duplication mirroring onto another volume. For example, you might be able to copy a volume currently in use on the source server to a new volume attached to the destination server. Because the new volume is a block-level image of the previous volume, the data and files should be perfectly preseeded.

  Note
  Be aware that Storage Spaces in Windows Server 2012 R2 and Windows Server 2012 does not allow creation of new volumes using mirror disks that were evicted from a previous storage pool.

- **Third-party software options** - Your backup or storage software vendors might offer additional options for preseeding files. Review documentation from your vendor for more information.

Important
If you plan to preseed data by using any non-Microsoft solution, you should test the solution and then validate the preseeded data on the destination server to ensure that the

file hashes are identical on the source and destination servers. For more information, see Step 2: Validate Preseeded Files for DFS Replication.

In this guide

This guide provides detailed instructions for preseeding files for DFS Replication, validating the preseeded files, and troubleshooting any issues that occur during preseeding:

- Step 1: Preseed Files for DFS Replication
  - Use Robocopy to Preseed Files for DFS Replication
  - Use Windows Server Backup to Preseed Files for DFS Replication
  - Use Windows NTBackup to Preseed Files for DFS Replication
- Step 2: Validate Preseeded Files for DFS Replication
- Troubleshoot Preseeding for Initial Synchronization in DFS Replication

Additional resources

Content type	References
Product evaluation	What's New in DFS Replication in Windows Server 2012 R2 | DFS Namespaces and DFS Replication Overview
Operations	DFS Replication help and basic operational guidance | DFS Replication: Frequently Asked Questions (FAQ) | **Robocopy** (command-line reference) | Windows Server Backup Step-by-Step Guide for Windows Server 2008 R2 | Windows Server Backup Step-by-Step Guide for Windows Server 2008 | Export a Clone of the DFS Replication Database
Community resources	The File Services and Storage TechNet Forum | Storage Team at Microsoft - File Cabinet Blog | Ask the Directory Services Team blog | DFS Replication: Survival Guide (TechNet Wiki) | Windows Server Information Experience on Twitter

Step 1: Preseed Files for DFS Replication

The topics in this section provide detailed walkthroughs of three methods for preseeding files on the destination server before you set up Distributed File System (DFS) Replication (also known as DFSR or DFS-R), add a new replication partner, or replace a server. By preseeding files, you can speed up initial synchronization for DFS Replication significantly and enable cloning of the DFS Replication database in Windows Server 2012 R2.

Step-by-step procedures are provided for the following methods:

- Use Robocopy to Preseed Files for DFS Replication
- Use Windows Server Backup to Preseed Files for DFS Replication
- Use Windows NTBackup to Preseed Files for DFS Replication

The best method to use depends on your environment, the operating systems on your source and destination computers, and the data set that you plan to replicate. For a comparison of these and other tools and an overview of the preseeding process, see DFS Replication: Copying Files to Preseed or Stage Initial Synchronization.

Important
To maximize the chances of preseeding successfully, we recommend that you test the procedures in your lab before you attempt production changes. Testing can also demonstrate which method of preseeding is best for your environment and data set.

Use Robocopy to Preseed Files for DFS Replication

This topic explains how to use the command-line tool Robocopy.exe to preseed files when setting up replication for Distributed File System (DFS) Replication (also known as DFSR or DFS-R) in Windows Server. By preseeding files before you set up DFS Replication, add a new replication partner, or replace a server, you can speed up initial synchronization and enable cloning of the DFS Replication database in Windows Server 2012 R2. The Robocopy method is one of several preseeding methods; for an overview, see Step 1: Preseed Files for DFS Replication.

The Robocopy (Robust File Copy) command-line utility is included with Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. The utility provides extensive options that include copying security, backup API support, retry capabilities, and logging. Later versions include multi-threading and un-buffered I/O support.

Important
Robocopy does not copy exclusively locked files. If users tend to lock many files for long periods on your file servers, consider using a different preseeding method. Preseeding does not require a perfect match between file lists on the source and destination servers, but the more files that do not exist when initial synchronization is performed for

DFS Replication, the less effective preseeding is. To minimize lock conflicts, use Robocopy during non-peak hours for your organization. Always examine the Robocopy logs after preseeding to ensure that you understand which files were skipped because of exclusive locks.

Preseeding files for DFS Replication by using Robocopy involves the following steps:

- Step 1: Download and install the latest version of Robocopy
- Step 2: Stabilize files that will be replicated
- Step 3: Copy the replicated files to the destination server

Prerequisites

Because preseeding does not directly involve DFS Replication, you need only meet the requirements for performing a file copy by using Robocopy:

- You will need to use an account that is a member of the local Administrators group on both the source server and the destination server.
- Install the most recent version of Robocopy on the server that you will use to copy the files – either the source server or the destination server; you will need to install the most recent version for the operating system version. For instructions, see Step 2: Stabilize files that will be replicated. Unless you are preseeding files from a server running Windows Server 2003 R2, you can run Robocopy on either the source or destination server. The destination server, which typically has the more recent operating system version, gives you access to the most recent version of Robocopy.
- Ensure that sufficient storage space is available on the destination drive. Do not create a folder on the path that you plan to copy to: Robocopy must create the root folder.

  Note
  When you decide how much space to allocate for the preseeded files, consider expected data growth over time and storage requirements for DFS Replication. For planning help, see Edit the Quota Size of the Staging Folder and Conflict and Deleted Folder in Managing DFS Replication.
- On the source server, optionally install Process Monitor or Process Explorer, which you can use to check for applications that are locking files. For download information, see Process Monitor v3.05 and Process Explorer v15.40.

Step 1: Download and install the latest version of Robocopy

Before you use Robocopy to preseed files, you should download and install the latest version of Robocopy.exe. This ensures that DFS Replication does not skip files because of issues within the shipping versions of Robocopy.

The source for the latest compatible Robocopy version depends on the version of Windows Server that is running on the server. For information about downloading the hotfix with the most recent version of Robocopy for Windows Server 2008 R2 or Windows Server 2008, see article 968429 in the Microsoft Knowledge Base.

Alternatively, you can locate and install the latest hotfix for an operating system by taking the following steps.

To locate and install the latest version of Robocopy for a specific version of Windows Server

1. In a web browser, open http://support.microsoft.com.
2. In **Search Support**, type the following string, replacing <*operating system version*> with the appropriate operating system, and press ENTER:

   ```
   robocopy.exe kbqfe "<operating system version>"
   ```

   For example, enter **robocopy.exe kbqfe "Windows Server 2008 R2"**.
3. Locate and download the hotfix with the highest ID number (that is, the latest version).
4. Install the hotfix on the server.

Step 2: Stabilize files that will be replicated

After you install the latest version of Robocopy on the server, you should minimize the chances that locked files will block copying by using methods in the following table. Most applications do not exclusively lock files. However, during normal operations, a small percentage of files might be locked on file servers.

Source of the lock	Explanation	Mitigation
Users remotely open files on shares.	Employees connect to a standard file server and edit documents, multimedia content, or other files. Sometimes referred to as the traditional home folder or shared data workloads.	Only perform Robocopy operations during off-peak, non-business hours. This minimizes the number of files that Robocopy must skip during preseeding. Consider temporarily setting Read-only access on the file shares that will be replicated by using the Windows PowerShell **Grant-SmbShareAccess** and **Close-SmbSession** cmdlets. If you set permissions for a common group such as Everyone or Authenticated Users to READ, standard users might

Source of the lock	Explanation	Mitigation
		be less likely to open files with exclusive locks (if their applications detect the Read-only access when files are opened). You might also consider setting a temporary firewall rule for SMB port 445 inbound to that server to block access to files or use the **Block-SmbShareAccess** cmdlet. However, both of these methods are very disruptive to user operations.
Applications open files local.	Application workloads running on a file server sometimes lock files.	Temporarily disable or uninstall the applications that are locking files. You can use Process Monitor or Process Explorer to determine which applications are locking files. To download Process Monitor, open Process Monitor v3.05 and Process Explorer v15.40.

Step 3: Copy the replicated files to the destination server

After you minimize locks on the files that will be replicated, you are ready to preseed the files from the source server to the destination server.

Note

You can run Robocopy on either the source computer or the destination computer. The following procedure describes running Robocopy on the destination server, which typically is running a more recent operating system, to take advantage of any additional Robocopy capabilities that the more recent operating system might provide.

To preseed the replicated files onto the destination server by using Robocopy

1. Log on to the destination server with an account that is a member of the local Administrators group on both the source server and the destination server.
2. Open an elevated command prompt.
3. To preseed the files from the source to destination server, run the following command,

substituting your own source, destination, and log file paths:

```
robocopy "<source replicated folder path>" "<destination
replicated folder path>" /e /b /copyall /r:6 /w:5 /MT:64 /xd
DfsrPrivate /tee /log:<log file path> /v
```

This command copies all contents of the source folder to the destination folder, with the following parameters:

Parameter	Description
"<source replicated folder path>"	Specifies the source folder to preseed on the destination server.
"<destination replicated folder path>"	Specifies the path to the folder that will store the preseeded files. **Important** The destination folder *must not already exist* on the destination server. To get matching file hashes, Robocopy must create the root folder when it preseeds the files.
/e	Copies subdirectories and their files, as well as empty subdirectories.
/b	Copies files in Backup mode.
/copyall	Copies all file information, including data, attributes, time stamps, the NTFS access control list (ACL), owner information, and auditing information.
/r:6	Retries the operation 6 times when an error occurs.
/w:5	Waits 5 seconds between retries.
MT:64	Copies 64 files simultaneously.
/xd DfsrPrivate	Excludes the DfsrPrivate folder.
/tee	Writes status output to the console window, as well as to the log file.
/log <log file path>	Specifies the log file to write. Overwrites the file's existing contents. (To append the

	entries to the existing log file, use `/log+ <log file path>`.)
`/v`	Produces verbose output that includes skipped files.

For example, the following command replicates files from the source replicated folder E:\RF01 to data drive D: on the destination server:

```
robocopy.exe "\\srv01\e$\rf01" "d:\rf01" /e /b /copyall /r:6
/w:5 /MT:64 /xd DfsrPrivate /tee
/log:c:\temp\preseedsrv02.log
```

Note

We recommend that you use the parameters described above when you use Robocopy to preseed files for DFS Replication. However, you can change some of their values or add additional parameters. For example, you might find out through testing that you have the capacity to set a higher value (thread count) for the **/MT** parameter. Also, if you will replicate larger files primarily, you might be able to increase copy performance by adding the **/j** option for unbuffered I/O. For more information about Robocopy parameters, see the Robocopy command-line reference.

Caution

To avoid potential data loss when you use Robocopy to preseed files for DFS Replication, do not make the following changes to the recommended parameters:

- Do not use the **/mir** parameter (which mirrors a directory tree) or the **/mov** parameter (which moves the files, and then deletes them from the source).
- Do not remove the **/e**, **/b**, and **/copyall** options.

4. After copying completes, examine the log for any errors or skipped files. Use Robocopy to copy any skipped files individually; do not re-copy the entire set of files. If files were skipped because of exclusive locks, retry copying individual files with Robocopy later, or accept that those files will require over-the-wire replication by DFS Replication during initial synchronization.

Next step

After you complete the initial copy, and use Robocopy to resolve issues with as many skipped files as possible, you will use the **Get-DfsrFileHash** cmdlet in Windows PowerShell or the **Dfsrdiag** command to validate the preseeded files by comparing file hashes on the source and destination servers. For instructions, see Step 2: Validate Preseeded Files for DFS Replication.

Use Windows Server Backup to Preseed Files for DFS Replication

This topic explains how to use the Windows Server Backup feature to preseed data when setting up replication for Distributed File System (DFS) Replication (also known as DFSR or DFS-R). By preseeding data before you set up DFS Replication, add a new replication partner, or replace a server, you can speed up initial synchronization and enable cloning of the DFS Replication database in Windows Server 2012 R2.

The Windows Server Backup graphical utility (wbadmin.msc) and command-line utility (wbadmin.exe) are included with Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. You can use Windows Server Backup to back up and restore data from one server to another using a transitional storage location. Windows Server Backup can back up exclusively locked files.

Note

Windows Server Backup is one of several tools that can be used to preseed data for DFS Replication. For a comparison of tools and an overview of the preseeding process, see Step 1: Preseed Files for DFS Replication.

Preseeding data for DFSR by using Windows Server Backup involves the following steps:

- Step 1: Install the Windows Server Backup feature
- Step 2: Back up the replicated files from the source server
- Step 3: Restore the replicated files to the destination server

Prerequisites

Because preseeding does not directly involve DFS Replication, you need only meet the requirements for performing a file copy by using Windows Server Backup:

- You will need to use an account that is a member of the local Administrators group on both the source server and the destination server.
- The source and destination computers must have an operating system that includes the Windows Server Backup feature. Windows Server Backup is included with the following operating systems: Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008.
- Install the Windows Server Backup feature on the source and destination servers. For more information, see Step 1: Install the Windows Server Backup feature.
- A separate volume must be available with ample storage for backing up the source folders. The backup volume can be on the source computer, the destination computer, or on network storage that is accessible to both computers. The backup volume can be stored on fixed or removable storage.

  Note

When you decide how much space to allocate for the preseeded files, consider expected data growth over time and storage requirements for DFS Replication. For planning help, see Edit the Quota Size of the Staging Folder and Conflict and Deleted Folder in Managing DFS Replication.

- On the destination server, the path that you plan to restore the data to should not contain any files. Any files restored as copies of existing files will be moved to the PreExisting folder when initial synchronization for DFS Replication starts.

Step 1: Install the Windows Server Backup feature

The Windows Server Backup feature is not installed by default. You must add it to both the source server and the destination server. The following procedures explain how to install Windows Server Backup.

To install Windows Server Backup on Windows Server 2012 R2 or Windows Server 2012 by using Server Manager

1. Start **Server Manager** (servermanager.exe).
2. On the Dashboard, click **Add roles and features**.
3. Proceed to the **Features** page, and select **Windows Server Backup**.
4. Follow the wizard instructions to install the feature.

To install Windows Server Backup on Windows Server 2012 R2 or Windows Server 2012 by using Windows PowerShell

1. Open Windows PowerShell with elevated rights.
2. Run the following command:

```
Install-WindowsFeature -Name Windows-Server-Backup -
IncludeManagementTools
```

To install Windows Server Backup on Windows Server 2008 R2 or Windows Server 2008

- To find out how to install Windows Server Backup on Windows Server 2008 R2 by using Server Manager or from the command line, see Windows Server Backup Step-by-Step Guide for Windows Server 2008 R2.
- To find out how to install Windows Server Backup on Windows Server 2008 by using Server Manager or from the command line, see Windows Server Backup Step-by-Step Guide for Windows Server 2008.

Step 2: Back up the replicated files from the source server

To back up files from the source server, you must have a different volume that has ample storage for the backed up data. That volume can be located on the source server, the destination server,

or on network storage that both servers can access. The volume can be on fixed or removable storage.

Note

The procedures in this topic describe how to back up to and restore selected folders from a shared folder. You can also back up to a local drive on the destination server. In that case, for Windows Server Backup to discover the backup on the destination server, you must back up to the drive's root folder.

To back up files using the wbadmin.msc graphical tool

1. Sign in on the source server with an account that is a member of the local Administrators group.
2. Start Windows Server Backup by using either of the following methods:
   - In **Server Manager**, on the **Tools** menu, click **Windows Server Backup**.
   - At an elevated command prompt, enter **wbadmin.msc**.
3. In the console tree, click **Local Backup**.
4. On the **Actions** pane, click **Backup Once**.

   The Backup Once Wizard opens.
5. On the **Backup Options** page, click **Next** to begin specifying an unscheduled backup.
6. On the **Select Backup Configuration** page, select **Custom**, and then click **Next**.
7. On the **Select Items for Backup** page, click **Add Items**, and then select the folders that you want to preseed on the destination server. Then click **Next**.

   Note

   The option to select specific folders to back up is not available in Windows Server 2008. That version of Windows Server Backup only supports full volume backups. However, it does allow you to restore specific folders.
8. On the **Specify Destination Type** page, select **Remote shared folder**, and then click **Next**.
9. On the **Specify Remote Folder** page, enter or select the shared folder where you want to store the backup, and then click **Next**. A **WindowsImageBackup** folder will be created in the shared folder to store the backup.
10. On the **Confirmation** page, review the options that you specified, and then click **Backup**.
11. In Windows Server 2008 only, on the **Specify advanced option** page, click **Next** to the default selection.
12. After the backup operations complete, verify that no errors occurred. In the console tree, click **Local Backup**, and then click the backup that just completed. If the status is not **Successful**, click **View details** in the **Status** pane below to review any errors that occurred.

To back up the replicated files by using the wbadmin.exe command-line tool

1. Sign in on the source server with an account that is a member of the local Administrators

group.

2. Open an elevated command prompt.
3. Run the following command, substituting your own path and folders:

```
wbadmin start backup -backuptarget:<storage location for the
backup>: -include:<folder1, folder2, foldern>
```

Note
In Windows Server 2008, you cannot use the **–include** parameter to specify individual folders to back up. You can only back up full volumes.

Example 1: The following example backs up the replicated folders RF01 and RF02 on F: to a local fixed disk on the G: volume:

```
wbadmin start backup -backuptarget:g: -include:f:\rf01,f:\rf02
```

Example 2: The following example backs up the same replicated folders to a remote file share on the SRV01 server:

```
wbadmin start backup -backuptarget:\\SRV01\BackDfsr$ -include:f:\rf01,f:\rf02
```

Step 3: Restore the replicated files to the destination server

You can now restore the backup that you just created to preseed the data from the source server to the destination server. The following procedures describe how to do this by using the graphical tool (wbadmin.msc) or the command-line utility (wbadmin.exe).

The following procedure describes how to use Windows Server Backup to preseed files from a backup stored on a shared folder.

To restore the replicated files by using the wbadmin.msc graphical tool

1. Sign in on the destination server with an account that is a member of the local Administrators group on both the source and destination servers.
2. Start **Windows Server Backup** from **Server Manager**, or by running Start Wbadmin.msc.
3. In the console tree, click **Local Backup**.
4. On the **Actions** pane, click **Recover**.

   The Recovery Wizard opens.
5. On the **Getting Started** page, select **A backup stored on another location** (to indicate that you saved the backup from a different server), and then click **Next**.
6. On the **Specify Location Type** page, click **Remote shared folder**. (Select this option even if the share that stores the backup is on the destination server.)
7. On the **Specify Remote Folder** page, type the UNC path to the shared folder where you stored your backup. For example, type **\\MyFileServer\Backups**. After you type the server name and a backslash, you can select from shared folders that store backups on

the server.

8. On the **Select Backup Date** page, select the date and time of the backup that you want to preseed on the destination server, and then click **Next**.
9. On the **Select Recovery Type** page, select **File and Folders**, and then click **Next**.
10. On the **Select Items to Recover** page, under **Available Items**, expand the backup, select the folder that stores the files you want to preseed, and then click **Next**.
11. On the **Specify Recovery Options** page, enter the following options, and then click **Next**:
    a. Select **Another location**, and then specify or select the location on the destination server to which you want to preseed the files.

    Important
    The target on the destination server should not contain any files. Any files restored as copies of existing files will be moved to the PreExisting folder when initial synchronization for DFS Replication starts.

    b. Leave the **Create copies so that you have both versions** and **Leave Restore access control list (ACL) permissions to the file or folder being recovered** options selected.
12. On the **Confirmation** page, review the options that you selected, and then click **Recover**.
13. After the restore operations complete, verify that no errors occurred. If any files are missing, check the backup to make sure the files were backed up. Then check the restore log to ensure that no files were skipped during the restore.

To restore the replicated files using wbadmin.exe command-line tool

1. On the destination server, open an elevated command prompt.
2. To retrieve a list of backups, run the following command:

   ```
   wwbadmin get versions –backuptarget:<path to saved backup>
   ```

   Example 1: Lists backups that were saved to the H: drive on the (local) destination server:

   ```
   wbadmin get versions -backupTarget:h:
   ```

   Example 2: Lists backups that were saved to a transitional file server share:

   ```
   wbadmin get versions -backupTarget:\\SRV01\BackDfsr
   ```

3. To restore the backup to the destination server, run the following command:

   ```
   wbadmin start recovery -version:<version> -itemtype:File -
   backuptarget:<path to saved backup> -items:<path to files to
   restore> -recursive -overwrite:createcopy -
   recoverytarget:<path to restore to>
   ```

   Example 1: Restores files from two folders - RF01 and RF02, which were backed up to the Q: drive on the destination server. The files are restored to the F: drive on the

destination (local) server:

```
wbadmin start recovery -version:06/07/2013-16:23 -itemtype:File -backuptarget:q: -items:f:\rf01,f:\rf02 -recursive -overwrite:createcopy -recoverytarget:f:\
```

Example 2: Restores files that were backed up to a transitional file server share named \\srv01\q$:

```
wbadmin start recover -version:06/07/2013-16:23 -itemtype:File -backuptarget:\\srv01\q$ -items:f:\rf01,f:\rf02 -recursive -overwrite:createcopy -recoverytarget:g:\
```

4. Review the list of restored files to ensure that all files from the backup were restored successfully.

Next step

After you use Windows Server Backup to back up files to an interim storage location and then restore them to the desired location on the destination server, and you resolve any issues with the backup and restore, use the **Get-DfsrFileHash** cmdlet in Windows PowerShell or the **Dfsrdiag** command to validate the preseeded files by comparing file hashes on the source and destination servers. For instructions, see Step 2: Validate Preseeded Files for DFS Replication.

Use Windows NTBackup to Preseed Files for DFS Replication

This topic provides step-by-step procedures for using the NTBackup command-line utility to preseed files for Distributed File System (DFS) Replication (also known as DFSR or DFS-R) from a computer running Windows Server 2003 R2 to a computer running Windows Server 2008 R2 or Windows Server 2008. By preseeding data before you set up DFS Replication, add a new replication partner, or replace a server, you can speed up initial synchronization.

The Windows NTBackup graphical utility (ntbackup.exe) is a tool for restoring backups made in Windows Server 2003 R2 to a Windows Server 2008 R2 or Windows Server 2008 computer. NTBackup provides the capability to back up and restore data from one server to another by using a transitional storage location. NTBackup can back up exclusively locked files. The tool is included with the Windows Server 2003 R2 operating system, but not with more recent operating systems. Downloadable versions of NTBackup are available for restoring files to a computer running Windows Server 2008 R2 or Windows Server 2008.

If you are accustomed to using Windows NTBackup, this tool provides a familiar method for preseeding files for DFS Replication on a Windows Server 2008 R2 or Windows Server 2008 computer. However, the use of NTBackup is limited:

- In Windows Server 2008 R2 and Windows Server 2008, NTBackup can be used to restore backups but not to create backups.

- Microsoft does not provide NTBackup tools for Windows Server 2012 R2 and Windows Server 2012. To preseed files from Windows Server 2003 R2 or Windows Server 2003 to Windows Server 2012 or later, you must use the Robocopy tool or non-Microsoft methods.

Note
NTBackup is one of several tools that can be used to preseeding data for DFS Replication. For a comparison of tools and an overview of the preseeding process, see DFS Replication: Copying Files to Preseed or Stage Initial Synchronization.

Preseeding files for DFS Replication by using NTBackup involves the following steps:

- Step 1: Install NTBackup on the destination server
- Step 2: Back up the files from the source server
- Step 3: Restore the backup to the destination server

Prerequisites

Because preseeding does not directly involve DFS Replication, you need only meet the requirements for backing up and restoring files by using Windows NTBackup:

- You will need to use an account that is a member of the local Administrators group on both the source server and the destination server.
- The destination computer must have the Windows Server 2008 R2 or Windows Server 2008 operating system installed.
- You must install NTBackup on the destination computer. (NTBackup will be installed already on your source computer, which is running Windows Server 2003 R2.) For more information, see Step 1: Install NTBackup on the destination server.
- To back up the files from the source computer, ample storage must be available on a different volume. The backup volume can exist on the source server, the destination server, or on network storage that can be accessed by both servers. The backup volume can be stored on fixed or removable storage.

  Note
  When you decide how much space to allocate for the preseeded files, consider expected data growth over time and storage requirements for DFS Replication. For planning help, see Edit the Quota Size of the Staging Folder and Conflict and Deleted Folder in Managing DFS Replication.

Step 1: Install NTBackup on the destination server

On a Windows Server 2008 R2 or Windows Server 2008 destination server, you must add the NTBackup feature. On the source computer, running Windows Server 2003 R2, NTBackup is installed automatically during operating system installation.

To install NTBackup on Windows Server 2008 R2

1. Log on to the Windows Server 2008 R2 computer as an administrator.

2. Download and install NTBackup. For download instructions, see article 974674 in the Microsoft Knowledge Base.

To install NTBackup on Windows Server 2008

1. Log on to the Windows Server 2008 computer as an administrator.
2. Open **Server Manager**, navigate to the **Features** section, and click **Add features**.
3. Select **Removable Storage Manager**, click **Next**, and then click **Install**.
4. Download and install the appropriate NTBackup MSI file from Windows NT Backup - Restore Utility.

Step 2: Back up the files from the source server

Your next step is to back up the files to interim storage on the source server, the destination server, or on network storage that can be accessed by both servers. You must have ample storage for the backed up data on a different volume. The backup can be stored on fixed or removable storage.

To back up files by using NTBackup

1. On the source server, start **Ntbackup.exe**.
2. If the **Backup or Restore Wizard** dialog box appears, clear the **Always start in wizard mode** check box, click **Cancel**, and then restart NTBackup.exe.
3. Click the **Backup** tab.
4. Select replicated folders to back up:
   a. Select each replicated folder that you want to back up.
   b. Clear the selection of the DfsrPrivate folder under each replicated folder so that those folders are not backed up. The DfsrPrivate folder increases the size of the backup but does not contain data useful to the destination server.
5. Under **Backup media or file name**, browse or type a path to a backup BKF file. This path can be a local path on the server or a remote share; the path can specify the same volume or a different volume. The target can be the destination server for the restore.
6. Click **Start Backup**. Select any desired options, entering a custom description and label if desired, and then click **Start Backup**.
7. After the backup operations complete, verify that no errors occurred during the backup.

Step 3: Restore the backup to the destination server

You are now ready to preseed the files from the source server to the destination server by restoring the backup.

To restore the files by using NtbackupRestore or NtBackup

1. On the destination server, start the Windows NTBackup restore utility, **NtBackupRestoreUtility.exe**.
2. Click **Restore Wizard (Advanced)**.
3. On the **Welcome to the Restore Wizard** page, click **Next**.
4. On the **What to restore** page, select the backup that you want to preseed on the destination server. To do this, use the **Browse** button to select the BKF file that was created during the backup.
5. Select each replicated folder that you want to restore, and then click **Next**.
6. On the **Completing the Restore Wizard** page, click **Advanced**.
7. On the **Where to Restore** page, specify **Alternate location**, select the correct location, and then click **Next**.
8. On the **How to Restore** page, select the **Replace existing files if they are older than the backup files** option, then click **Next**.
9. On the **Advanced Restore Options** page, select the **Restore security settings** option, clear all other options, and then click **Next**.
10. On the **Completing the Restore Wizard** page, click **Finish**.
11. After the restore operations complete, verify that that no errors occurred.

Next step

After you restore your files to the destination server and you resolve any issues with the backup and restore, use the **Get-DfsrFileHash** cmdlet in Windows PowerShell or the **Dfsrdiag** command to validate the preseeded files by comparing file hashes on the source and destination servers. For instructions, see Step 2: Validate Preseeded Files for DFS Replication.

Step 2: Validate Preseeded Files for DFS Replication

After you preseed files on the destination server in preparation for setting up replication for Distributed File System (DFS) Replication (also known as DFSR or DFS-R), you should confirm that the preseeding succeeded by using the tools included with Windows Server or the Remote Server Administration Tools (RSAT) client.

These utilities perform a simulated file marshaling using the same hash computation function that the DFS Replication service uses, and they can guarantee that a file is identical on two servers for the purpose of replication.

Caution

Do not use checksum comparison tools designed to calculate hashes for downloaded files. This includes Microsoft tools such as the Microsoft File Checksum Integrity Verifier and LogParser 2.2. Those tools do not calculate the security and alternate data stream portions of files as part of the hash.

Best practices for validating preseeded files

By following these guidelines, you can get an accurate picture of the validity of your preseeded files efficiently and without affecting performance:

- When you compare files on the source and destination servers, do not attempt to validate every file hash. It is better to pick several locations within a preseeded data set and spot check the file hashes. If you preseeded millions of files, validating every file hash might take too long.

  By selecting a handful of subfolders after preseeding, as well as the direct contents of the root folder, you can get a good picture of the validity of the file hashes without degrading performance. Typically, an incorrectly preseeded data set yields mismatched hashes for every file and folder; it is difficult to preseed with only partial success.
- If users are accessing files on the source server while you preseed data, some files probably will have differing hashes simply from user actions. If you see individual file hash mismatches, but most of the files have matching hashes, check the file size, last modified date and time, and permissions on each copy of the file to confirm that the difference is expected.
- Consider counting the replicated files and folders on the source and destination server to ensure that the destination server does not have a gross mismatch of fewer objects. This is not a precise test if you allowed users to create files on the source server during preseeding, but it provides another data point to consider.

  For instance, you can use the following Windows PowerShell command on both the source server and the destination server, replacing the **–Path** value with the path that you want to check, to ensure that the file counts are similar:

  ```
  (Get-ChildItem -Path h:\rf01 -Recurse -Force).count
  ```

Two methods for validating preseeded files

The following sections provide a detailed walkthrough of two validation techniques available to administrators who manage DFS Replication. You must decide which one works best for your environment.

- Get-DfsrFileHash method
- Dfsrdiag method

Get-DfsrFileHash method

The DFSR module of Windows PowerShell, which is provided with Windows Server 2012 R2, includes the **Get-DfsrFileHash** cmdlet, which you can use to calculate file hashes on local or

remote computers. The cmdlet has only one parameter, **-Path**, which accepts UNC paths, mapped drives, and local volumes. The cmdlet does not require that the DFS Replication service be installed or that replication be configured on the server. You can run the cmdlet against any target operating system. If you have a single computer running Windows Server 2012 R2, you can retrieve any file hashes on any Windows operating system.

Note

The DFSR module for Windows PowerShell also is provided with the Windows 8.1 Remote Server Administration Tools (RSAT). For download information, see Remote Server Administration Tools for Windows 8.1.

Compare a file hash for a single file

The following example code uses **Get-DfsrFileHash** to retrieve the hash of the file rascfg.dll on both the source server and the destination server. In this example, the file hash on the local server is not a match for the file hash on the remote server (named SRV02), which indicates that preseeding was not successful.

```
PS C:\> Get-DfsrFileHash -Path 'C:\rf03\rascfg.dll' | Format-Table -Auto -Wrap

Path               FileHash

----               --------

C:\rf03\rascfg.dll 04D164A0-ADEB6583-7DB21B5B-22AF572C


PS C:\> Get-DfsrFileHash -Path '\\srv02\c$\rf03\rascfg.dll' | Format-Table -Auto -Wrap

Path                          FileHash

----                          --------

\\srv02\c$\rf03\rascfg.dll    D7F1EF42-5FEF2DDE-2F3ECDF0-AD26F343
```

Compare file hashes for a folder's immediate contents

The following sample code uses **Get-DfsrFileHash** retrieve the files hashes for all files in direct child folders on both the source server and the destination server for the preseeded data set. In this example, all files hatches match for files on the local server and the remote server (named SRV02).

```
PS C:\> Get-DfsrFileHash -Path 'C:\rf03\*' | Format-Table -Auto -Wrap


Path                FileHash

----                --------

C:\rf03\app1        DF896CD6-23100551-A2D4F4D7-7BBCA290

C:\rf03\canary      DF896CD6-23100551-A2D4F4D7-7BBCA290

C:\rf03\doc2        DF896CD6-23100551-A2D4F4D7-7BBCA290
```

```
C:\rf03\rasauto.dll  E725F26C-9A62CDD9-9A4E47D3-93F61FD9
C:\rf03\rasautou.exe 1C42D753-48E566B2-7FCFDC9D-2F2C1FF2
C:\rf03\rascfg.dll   04D164A0-ADEB6583-7DB21B5B-22AF572C

PS C:\> Get-DfsrFileHash -Path '\\srv02\c$\rf03\*' | Format-Table -Auto -Wrap

Path                          FileHash
----                          --------
\\srv02\c$\rf03\app1          DF896CD6-23100551-A2D4F4D7-7BBCA290
\\srv02\c$\rf03\canary        DF896CD6-23100551-A2D4F4D7-7BBCA290
\\srv02\c$\rf03\doc2          DF896CD6-23100551-A2D4F4D7-7BBCA290
\\srv02\c$\rf03\rasauto.dll   E725F26C-9A62CDD9-9A4E47D3-93F61FD9
\\srv02\c$\rf03\rasautou.exe 1C42D753-48E566B2-7FCFDC9D-2F2C1FF2
\\srv02\c$\rf03\rascfg.dll    04D164A0-ADEB6583-7DB21B5B-22AF572C
```

Note

A simple method for comparing hashes between servers is to use a Windows PowerShell pipeline to send the output to a file (for example, use **Export-CSV**), and then open the file in Microsoft Excel. You can then compare the results of each server's output.

Get-DfsrFileHash (like **Get-Item**) does not return hidden files when wildcards are used. **Get-DfsrFileHash** also does not retrieve file hashes recursively for deeper folders and files. Both of these behaviors are by design. Searching recursively and for every hidden file typically is not necessary and can be inefficient when your goal is to spot check preseeding results.

If you want to retrieve nested hashes and hashes for hidden files and folders, you can use the **Get-ChildItem** cmdlet with the following parameters, substituting your own paths:

```
Get-DfsrFileHash -Path (Get-ChildItem -Path c:\Rf01\subfolder1\subfolder2 –Recurse -Force).fullname
```

Important

Use care with this approach to hash retrieval: on large data sets, returning the data takes significant time. If you specify the root of the replicated folder, **Get-ChildItem** by default also returns the hidden DfsrPrivate folder. If you do not want to retrieve hash data for the DfsrPrivate folder, use the **–Exclude** parameter to exclude it.

Compare file hashes between two servers by using Compare-Object

To retrieve the hash of all direct child files and folders in a replicated folder on both the source server and the destination server, and then compare those hashes automatically, you can leverage mapped drives and the **Compare-Object** cmdlet.

The following code sample includes a series of commands that perform the following tasks:

1. Map a drive to the source server.
2. Save the file hashes to a file.
3. Delete the drive on the source server.
4. Map a drive to a relative path on the destination server.
5. Save file hashes for that drive to a file.
6. Compare the two files.

The file comparison that is returned shows both equal and unequal results. Each difference is identified by a => or <= marker. In this case, most files and folders match, but the hashes for rascfg.dll differ.

```
PS C:\> net use x: \\Srv01\c$\Rf03

PS C:\> Get-DfsrFileHash x:\* | Out-File c:\Srv01.txt

PS C:\> net use x: /d

PS C:\> net use x: \\Srv02\c$\rf03

PS C:\> Get-DfsrFileHash x:\* | Out-File c:\Srv02.txt

PS C:\> net use x: /d

PS C:\> Compare-Object -ReferenceObject (Get-Content C:\Srv01.txt) -DifferenceObject
(Get-Content C:\Srv02.txt) -IncludeEqual


InputObject                     SideIndicator

-----------                     -------------

                                ==

Path                            ... ==

----                            ... ==

x:\app1                         ... ==

x:\canary                       ... ==

x:\doc2                         ... ==

x:\rasauto.dll                  ... ==

                                ==

                                ==

x:\rascfg.dll                   ... =>

x:\rascfg.dll                   ... <=
```

Dfsrdiag method

As an alternative to the **Get-DfsrFileHash** cmdlet, you can use the older **DFSRdiag.exe** tool with the FILEHASH option to calculate file hashes on local or remote computers. This updated tool is included with Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2.

Like **Get-DfsrFileHash**, this tool accepts UNC paths, mapped drives, and local volumes. It does not require that the DFS Replication service be installed or that replication be configured. And **Dfsrdiag.exe** can run against any target operating system. If you have a single computer running Windows Server 2008 R2 or later, you can retrieve any file hashes on any Windows operating system.

However, this older tool can only return one file or folder at a time. For this reason, the **Get-DfsrFileHash** cmdlet is the recommended tool for comparing file hashes after preseeding files.

Compare a file hash for a single file

The following example code uses **Dfsrdiag.exe** to retrieve the hash of the file rascfg.dll on both the source and destination server. In this example, the file on the local server matches the file on SRV02.

```
C:\> dfsrdiag filehash /path:f:\rf01\rascfg.dll

   File Hash: A35BA042-F0DC3008-A5863965-F5010326

Operation Succeeded


C:\> dfsrdiag filehash /path:\\srv02\g$\rf01\rascfg.dll

   File Hash: A35BA042-F0DC3008-A5863965-F5010326

Operation Succeeded
```

Note

TechNet offers an unsupported sample script that automates **Dfsrdiag**, provides source and destination comparison, and allows you to select random files. For more information, see PowerShell Script to Calculate and compare filehashes for DFSR Pre-Staging.

Related resources

- Troubleshoot Preseeding for Initial Synchronization in DFS Replication
- DFS Replication: Copying Files to Preseed or Stage Initial Synchronization
- Step 1: Preseed Files for DFS Replication

Troubleshoot Preseeding for Initial Synchronization in DFS Replication

If you encounter problems while you are preseeding files for Distributed File System (DFS) Replication (also known as DFSR or DFS-R), or you find later that all of the files and folders have mismatched hashes, you will need to troubleshoot. This topic explains how to resolve issues that are most commonly reported in forums and Microsoft Customer Support Services.

Guidance for following issues is provided:

- Troubleshoot mismatched hashes
- Troubleshoot missing files and folders
- Troubleshoot poor copy or restore performance

Troubleshoot mismatched hashes

If you find that files or folders on the source and destination servers have mismatched hashes, validate the following:

- If you used Robocopy to preseed files for DFS Replication, check for the following issues:
  - Ensure that the downstream server did not already contain the destination folder before you copied the folder and its contents. If the destination folder already exists on the destination server, the permissions are not copied correctly to the root folder.
  - Ensure that you included the **/copyall** option when you ran Robocopy.exe.
- If you used a third-party backup solution, ensure that your backup utility restores files without altering the security or alternate data streams. Contact the backup vendor to confirm this capability.

Troubleshoot missing files and folders

If you find that files are missing after preseeding, validate the following:

- If you used Robocopy to preseed files for DFS Replication, check for the following issues:
  - Ensure that you provided the **/E /B** options during the copy.
  - Examine the Robocopy log file for files that were skipped because of exclusive locks.
- If you restored from a backup:
  - Verify that the backup you restored included the files.
  - Examine the restore log to ensure that no files were skipped during the restore.
- If you are using security software such as anti-malware protection, ensure that files were not deleted or quarantined during copy or restore from backup because of virus heuristics.

Troubleshoot poor copy or restore performance

If copy or restore performance is slower than expected, validate the following:

- If you are using security software such as anti-malware protection, ensure that file scanning that the security software performs is not affecting performance. Consider setting exclusions for the replicated folders.
- If you are copying large amounts of data over slower, lower-bandwidth, or high-latency networks, consider copying during off-peak hours, when more throughput is available. Alternatively, ship a copy of the data using a physical mail or package service.
- Ensure that you have installed the latest networking drivers from your vendors.
- If you are using Robocopy, check for the following issues:
  - Ensure that you specified an appropriate **/MT** (multithread) value to ensure maximum throughput (an **/MT** value of 64 is recommended). If you are copying mostly large files

(multi-Gigabyte or larger), consider using the **/j** option if that option is supported by your version of Robocopy.

- Ensure that you are not using the **/zb** options, as this might be slower than copying with just **/b** (Backup mode) in some versions of Robocopy. The **/zb** options together copy files in Uses Restart mode; however, if access is denied, the files are copied Backup mode.
- Ensure that **/R:6 /W:5** or similar settings are used, to ensure that copying is not blocked by long-running retries for a file. If **/R** and **/W** are not set, the default behavior is for Robocopy to retry each file 1 million times, waiting 30 seconds between each retry (that is, Robocopy might wait one full year for a file).

Related resources

- DFS Replication: Copying Files to Preseed or Stage Initial Synchronization
- Step 1: Preseed Files for DFS Replication
- Step 2: Validate Preseeded Files for DFS Replication

File Server Resource Manager Overview

This topic describes the File Server Resource Manager features that are included with the File and Storage Services server role in Windows Server 2012.

Did you mean…

- What's New in File Server Resource Manager in Windows Server 2012
- File and Storage Services Overview
- Dynamic Access Control: Scenario Overview

Feature descriptions

File Server Resource Manager is a set of features that allow you to manage and classify data that is stored on file servers. File Server Resource Manager includes the following features:

- **File Classification Infrastructure** File Classification Infrastructure provides insight into your data by automating classification processes so that you can manage your data more effectively. You can classify files and apply policies based on this classification. Example policies include dynamic access control for restricting access to files, file encryption, and file expiration. Files can be classified automatically by using file classification rules or manually by modifying the properties of a selected file or folder.
- **File Management Tasks** File Management Tasks enables you to apply a conditional policy or action to files based on their classification. The conditions of a file management task include the file location, the classification properties, the date the file was created, the last modified date of the file, or the last time the file was accessed. The actions that a file management task can take include the ability to expire files, encrypt files, or run a custom command.

- **Quota management** Quotas allow you to limit the space that is allowed for a volume or folder, and they can be automatically applied to new folders that are created on a volume. You can also define quota templates that can be applied to new volumes or folders.
- **File screening management** File screens help you control the types of files that user can store on a file server. You can limit the extension that can be stored on your shared files. For example, you can create a file screen that does not allow files with an MP3 extension to be stored in personal shared folders on a file server.
- **Storage reports** Storage reports are used to help you identify trends in disk usage and how your data is classified. You can also monitor a selected group of users for attempts to save unauthorized files.

The features included with File Server Resource Manager can be configured and managed by using the File Server Resource Manager Microsoft Management Console (MMC) or by using Windows PowerShell.

Important
File Server Resource Manager supports volumes formatted with the NTFS file system only. The Resilient File System is not supported.

Practical applications

Some practical applications for File Server Resource Manager include:

- Use File Classification Infrastructure with the Dynamic Access Control scenario to create a policy that grants access to files and folders based on the way files are classified on the file server.
- Create a file classification rule that tags any file that contains at least 10 social security numbers as having personally identifiable information.
- Expire any file that has not been modified in the last 10 years.
- Create a 200 megabyte quota for each user's home directory and notify them when they are using 180 megabytes.
- Do not allow any music files to be stored in personal shared folders.
- Schedule a report that runs every Sunday night at midnight that generates a list of the most recently accessed files from the previous two days. This can help you determine the weekend storage activity and plan your server downtime accordingly.

See also

Use the links in the following table to learn more about File Server Resource Manager.

Content type	References
Product evaluation	What's New in File Server Resource Manager in Windows Server 2012 | What's New in File Server Resource Manager in Windows Server

Content type	References
	2012 R2 | Dynamic Access Control: Scenario Overview | File and Storage Services Overview

What's New in File Server Resource Manager in Windows Server 2012 R2

This topic describes the File Server Resource Manager functionality that is new or changed in Windows Server 2012 R2.

Feature description

File Server Resource Manager provides a set of features that enable you to manage and classify data that is stored on file servers. For an overview of the File Server Resource Manager feature, see File Server Resource Manager Overview.

You can configure the features that are included with File Server Resource Manager by using the File Server Resource Manager snap-in or by using Windows PowerShell.

New and changed functionality

The following table summarizes the File Server Resource Manager functionality in Windows Server 2012 R2 that is new or changed since Windows Server 2012.

Feature/Functionality	New or Improved	Description
Clear classification property values that no longer apply to an updated file	New	Provides the ability to clear property values that no longer apply to an updated file during reevaluation of existing classification property values.
Set maximum values for storage reports	New	Enables you to configure the maximum number of files per storage report, and to configure maximum values in the default parameters for specific storage reports.

Clear classification property values that no longer apply to an updated file

The **Evaluation Type** tab in the properties of a classification rule now includes options to clear automatically-classified and user-classified property values during reevaluation of existing property values if the property value no longer applies to the updated file.

What value does this change add?

This change enables File Server Resource Manager to dynamically remove classification values that no longer apply to a file. For example, you may have a file that was classified as sensitive because it contained a specific string such as "Private." If the string is removed, the file is no longer considered sensitive after reevaluation occurs.

What works differently?

In Windows Server 2012, if there is a conflict during the reevaluation of property values, you can overwrite the existing value or aggregate the values. Although these options are still available in Windows Server 2012 R2, if you choose to overwrite the existing value, you can also specify whether you want to clear a property value that no longer applies.

Set maximum values for storage reports

On the **Storage Reports** tab when you configure File Server Resource Manager options, you can now specify the maximum number of files per report. When you configure default parameters for a report, you can also configure the following maximum values:

- For the **Duplicate Files** report, you can specify the maximum number of files in a duplicate group per report, and the maximum number of groups of duplicate files per report.
- For the **Files by File Group** report, you can specify the maximum number of duplicate file groups per report, and the maximum number of files in any file group per report.
- For the **Files by Owner** report, you can specify the maximum number of owners per report, and the maximum number of files per owner per report.
- For the **Files by Property** and the **Folders by Property** reports, you can specify the maximum number of property values per report, and the maximum number of files for each property value. Realize that if you set the values for one of these report types, it affects both report types.

What value does this change add?

If you have more files or groups than the report default limits, you can increase the maximum values so that you can see all files and groups in the report.

What works differently?

In Windows Server 2012, the default limit for the number of files per storage report was 1000, and the default limit for the number of groups was 10. You could not change these values through the user interface or through Windows PowerShell. Options to configure maximum values in the default parameters of a report type were not available in Windows Server 2012.

See also

The following table provides additional resources for evaluating File Server Resource Manager.

Content type	References
Product evaluation	What's New in File Server Resource Manager in Windows Server 2012
Related technologies	File and Storage Services Overview | Dynamic Access Control: Scenario Overview

What's New in File Server Resource Manager in Windows Server 2012

This topic describes the features that were added to File Server Resource Manager in Windows Server 2012.

Did you mean…

- File Server Resource Manager Overview
- Dynamic Access Control: Scenario Overview
- File and Storage Services Overview

Feature description

File Server Resource Manager provides a set of features that allow you to manage and classify data that is stored on file servers. File Server Resource Manager includes the following features:

- **File Classification Infrastructure** Provides insight into your data by automating classification processes so that you can manage your data more effectively. You can classify files and apply policies based on this classification. Example policies include Dynamic Access Control for restricting access to files, file encryption, and file expiration. Files can be classified automatically by using file classification rules, or manually by modifying the properties of a selected file or folder.
- **File Management Tasks** Apply a conditional policy or action to files, based on their classification. The conditions of a file management task include the file location, the classification properties, the date the file was created, the last modified date of the file, or the last time the file was accessed. The actions that a file management task can take include the ability to expire files, encrypt files, or run a custom command.
- **Quota Management** Limit the space allowed for a volume or folder and quotas can be automatically applied to new folders that are created on a volume. You can also define quota templates that can be applied to new volumes or folders.

- **File Screening Management** Control the types of files that users can store on a file server. You can limit the extension that can be stored on your file shares. For example, you can create a file screen that does not allow files with an MP3 extension to be stored in personal shared folders on a file server.
- **Storage Reports** Identify trends in disk usage and how your data is classified, and monitor attempts by a selected group of users to save unauthorized files.

The features that are included with File Server Resource Manager can be configured and managed by using the File Server Resource Manager MMC console or by using Windows PowerShell.

New and changed functionality

The following table lists the changes that are new to File Server Resource Manager in Windows Server 2012.

Feature/functionality	New or updated?	Description
Dynamic Access Control	New	Dynamic Access Control uses File Classification Infrastructure to help you centrally control access and audit access to files on your file servers.
Automatic classification	Updated	The updates to automatic classification enable you to get more precise control on how data is classified on your file servers, including continuous classification, using Windows PowerShell for custom classification, updates to the existing content classifier, and dynamic name space for classification rules.
Manual classification	New	Manual classification enables users to classify files and folders manually without the need to create automatic classification rules.
File management tasks	Updated	The updates to file management tasks include an Active Directory Rights

Feature/functionality	New or updated?	Description
		Management Services (AD RMS) file management task, continuous file management tasks, and dynamic name space for file management tasks.
Access-denied assistance	New	Access-denied assistance allows you to customize the access denied error message in Windows 8 that users see when they do not have access to a file or folder.

Dynamic Access Control

In Windows Server 2012, you can apply data governance across your file servers to control who can access information and to audit who has accessed information. Dynamic Access Control lets you:

- Identify data by automatically or manually classifying files.
- Control access to files by applying central access policies.
- Audit access to files by using central audit policies for compliance reporting and forensic analysis.

What value does this change add?

Organizations can define central access and audit policies in Active Directory and use them to control who can access information and to track who accessed information in file servers.

Automatic classification

Automatic classification by using File Classification Infrastructure has been enhanced in Windows Server 2012 in the following ways:

- **Continuous classification** Configure File Classification Infrastructure to classify files a few seconds after they are created or modified on the file server, without having to wait for the next scheduled time for the classification to occur.
- **Windows PowerShell classifier** Classify a file automatically by running a Windows PowerShell script that determines the file classification. The Windows PowerShell classifier makes it easier to implement custom classification logic that is specific to your organization. For example, you can classify a file based on who last created or modified the file.
- **Enhanced content classifier** Specify minimum and maximum occurrences of a string or regular expression. For example, you could classify a file that contains more than ten social security numbers as having personally identifiable information.

- **Dynamic name space for classification rules** Specify the type of information that a folder contains, such as application data, group data, or user data, and then configure classification rules based on the type of information in which you want these to operate.

Manual classification

On the **Classification** tab of the file properties in Windows Server 2012, File Classification Infrastructure adds the ability to manually classify files. You can also classify folders so that any file added to the classified folder will inherit the classifications of the parent folder.

What value does this change add?

Manual classification gives users and content owners the ability to classify their files and folders by using the properties sheet of that file or folder.

File management tasks

In Windows Server 2012, file management tasks have been updated in the following ways:

- **Active Directory Rights Management Services file management task** Encrypt automatically any file that has an AD RMS protector when a specified condition is met. You can select an existing AD RMS rights policy template or specify the policy manually.
- **Continuous file management tasks** Configure file management tasks to run a few seconds after files are created or modified on a file server when classification properties are defined as a condition in the file management task.

  Note
  You cannot configure the file management tasks to be continuous if you have configured a notification or if a fixed schedule is assigned.
- **Dynamic name space for file management tasks** Specify the type of information that a folder contains, such as application data, group data, or user data, and then configure file management tasks based on the type of information in which you want these to operate.

Access-denied assistance

Access-denied assistance enables you to customize the access-denied error message that is displayed when a user running Windows 8 does not have access to a file or folder on a file server. You can configure the error message so that the user can request access to the file directly from the dialog box. You can also specify the user group that is sent the access request by using File Server Resource Manager.

Access-denied assistance can be configured by using Group Policy or by using the File Server Resource Manager console on each file server. You can also customize the error message per file server or you can have a separate error message for each file share on the file server. For more information about configuring access-denied assistance see Scenario: Access-Denied Assistance.

What value does this change add?

Access-denied assistance makes it easier for users to troubleshoot access issues for users so they can get access to the file and folders they need in a more efficient way.

Removed or deprecated functionality

In Windows Server 2012, the File Server Resource Manager command-line tools (dirquota.exe, filescrn.exe, and storrept.exe) are deprecated and all functionality has been replicated in Windows PowerShell cmdlets. The command-line tools still exist in the product, but all management tasks can be accomplished by using Windows PowerShell cmdlets.

Note
Any tasks that require a schedule cannot be configured by using the command line. You must use the Windows PowerShell cmdlets in File Server Resource Manager.

See also

The following table provides additional resources for evaluating File Server Resource Manager.

Content type	References
Product evaluation	What's New in File Server Resource Manager in Windows Server 2012
Related technologies	File and Storage Services Overview

Folder Redirection, Offline Files, and Roaming User Profiles overview

This topic discusses the Folder Redirection, Offline Files (client-side caching or CSC), and Roaming User Profiles (sometimes known as RUP) technologies, including what's new in Windows 8 and Windows Server 2012 and where to find additional information.

Did you mean…

- Offline Files and Folder Redirection (Windows 7 and Windows Server 2008 R2)
- Work Folders Overview
- BranchCache Overview

Technology description

Folder Redirection and Offline Files are used together to redirect the path of local folders (such as the Documents folder) to a network location, while caching the contents locally for increased speed and availability. Roaming User Profiles is used to redirect a user profile to a network location. These features used to be referred to as Intellimirror.

- **Folder Redirection** enables users and administrators to redirect the path of a known folder to a new location, manually or by using Group Policy. The new location can be a folder on the local computer or a directory on a file share. Users interact with files in the redirected folder as if it still existed on the local drive. For example, you can redirect the Documents folder, which is usually stored on a local drive, to a network location. The files in the folder are then available to the user from any computer on the network.
- **Offline Files** makes network files available to a user, even if the network connection to the server is unavailable or slow. When working online, file access performance is at the speed of the network and server. When working offline, files are retrieved from the Offline Files folder at local access speeds. A computer switches to Offline Mode when:
  - The new *Always Offline* mode has been enabled
  - The server is unavailable
  - The network connection is slower than a configurable threshold
  - The user manually switches to Offline Mode by using the **Work offline** button in Windows Explorer
- **Roaming User Profiles** redirects user profiles to a file share so that users receive the same operating system and application settings on multiple computers. When a user signs in to a computer by using an account that is set up with a file share as the profile path, the user's profile is downloaded to the local computer and merged with the local profile (if present). When the user signs out of the computer, the local copy of their profile, including any changes, is merged with the server copy of the profile. Roaming User Profiles is typically enabled on domain accounts by a network administrator.

Practical applications

Administrators can use Folder Redirection, Offline Files, and Roaming User Profiles to centralize storage for user data and settings and to provide users with the ability to access their data while offline or in the event of a network or server outage. Some specific applications include:

- Centralize data from client computers for administrative tasks, such as using a server-based backup tool to back up user folders and settings.
- Enable users to continue accessing network files, even if there is a network or server outage.
- Optimize bandwidth usage and enhance the experience of users in branch offices who access files and folders that are hosted by corporate servers located offsite.
- Enable mobile users to access network files while working offline or over slow networks.

New and changed functionality

The following table describes some of the major changes in Folder Redirection, Offline Files, and Roaming User Profiles that are available in this release.

Feature/functionality	New or updated?	Description
Always Offline mode	New	Provides faster access to files and lower bandwidth usage by always working offline, even when connected through a high-speed network connection.
Cost-aware synchronization	New	Helps users avoid high data usage costs from synchronization while using metered connections that have usage limits, or while roaming on another provider's network.
Primary Computer support	New	Enables you to limit the use of Folder Redirection, Roaming User Profiles, or both to only a user's primary computers.

Always Offline mode

In Windows 8 and Windows Server 2012, administrators can configure the experience for users of Offline Files to always work offline, even when they are connected through a high-speed network connection. Windows updates files in the Offline Files cache by synchronizing hourly in the background, by default.

What value does this change add?

The Always Offline mode provides the following benefits:

- Users experience faster access to files in redirected folders, such as the Documents folder.
- Network bandwidth is reduced, decreasing costs on expensive WAN connections or metered connections such as a 4G mobile network.

What works differently?

Prior to Windows 8 and Windows Server 2012, users would transition between the Online and Offline modes, depending on network availability and conditions, even when the Slow-Link mode (also known as the Slow Connection mode) was enabled and set to a 1 millisecond latency threshold.

With Always Offline mode, computers running Windows 8 and Windows Server 2012 never transition to Online mode when the **Configure slow-link mode** Group Policy setting is configured and the **Latency** threshold parameter is set to 1 millisecond. Changes are synced in the background every 120 minutes, by default, but synchronization is configurable by using the **Configure Background Sync** Group Policy setting.

For more information, see Enable the Always Offline Mode to Provide Faster Access to Files.

Cost-aware synchronization

With cost-aware synchronization, Windows disables background synchronization when the user is using a metered network connection, such as a 4G mobile network, and the subscriber is near or over their bandwidth limit, or roaming on another provider's network.

Note
Metered network connections usually have round-trip network latencies that are slower than the default 35 millisecond latency value for transitioning to Offline (Slow Connection) mode in Windows 8 and Windows Server 2012. Therefore, these connections usually transition to Offline (Slow Connection) mode automatically.

What value does this change add?

Cost-aware synchronization helps users avoid unexpectedly high data usage costs while using metered connections that have usage limits, or while roaming on another provider's network.

What works differently?

Prior to Windows 8 and Windows Server 2012, users who wanted to minimize fees while using Offline Files on metered network connections had to track their data usage by using tools from the mobile network provider. The users could then manually switch to Offline mode when they were roaming, near their bandwidth limit, or over their limit.

In Windows 8 and Windows Server 2012, Windows automatically tracks roaming and bandwidth usage limits while on metered connections. When the user is roaming, near their bandwidth limit, or over their limit, Windows switches to Offline mode and prevents all synchronization. Users can still manually initiate synchronization, and administrators can override cost-aware synchronization for specific users, such as executives.

For more information, see Enable Background File Synchronization on Metered Networks.

Primary computers for Folder Redirection and Roaming User Profiles

In Windows Server 2012, you can designate a set of computers, known as primary computers, for each domain user, which enables you to control which computers use Folder Redirection, Roaming User Profiles, or both. Designating primary computers is a simple and powerful method to associate user data and settings with particular computers or devices, simplify administrator oversight, improve data security, and help protect user profiles from corruption.

What value does this change add?

There are four major benefits to designating primary computers for users:

- The administrator can specify which computers users can use to access their redirected data and settings. For example, the administrator can choose to roam user data and settings between a user's desktop and laptop, and to not roam the information when that user logs on to any other computer, such as a conference room computer.
- Designating primary computers reduces the security and privacy risk of leaving residual personal or corporate data on computers where the user has logged on. For example, a general manager who logs on to an employee's computer for temporary access does not leave behind any personal or corporate data.
- Primary computers enable the administrator to mitigate the risk of an improperly configured or otherwise corrupt profile, which could result from roaming between differently configured systems, such as between x86-based and x64-based computers.
- The amount of time required for a user's first sign-in on a non-primary computer, such as a server, is faster because the user's roaming user profile and/or redirected folders are not downloaded. Sign-out times are also reduced, because changes to the user profile do not need to be uploaded to the file share.

What works differently?

To limit downloading private user data to primary computers, the Folder Redirection and Roaming User Profiles technologies perform the following logic checks when a user signs in to a computer:

1. The Windows operating system checks the new Group Policy settings (**Download roaming profiles on primary computers only** and **Redirect folders on primary computers only**) to determine if the **msDS-Primary-Computer** attribute in Active Directory Domain Services (AD DS) should influence the decision to roam the user's profile or apply Folder Redirection.
2. If the policy setting enables primary computer support, Windows verifies that the AD DS schema supports the **msDS-Primary-Computer** attribute. If it does, Windows determines if the computer that the user is logging on to is designated as a primary computer for the user as follows:
    a. If the computer is one of the user's primary computers, Windows applies the Roaming User Profiles and Folder Redirection settings.
    b. If the computer is not one of the user's primary computers, Windows loads the user's cached local profile, if present, or it creates a new local profile. Windows also removes any existing redirected folders according to the removal action that was specified by the previously applied Group Policy setting, which is retained in the local Folder Redirection configuration.

For more information, see Deploy Primary Computers for Folder Redirection and Roaming User Profiles

Hardware requirements

Folder Redirection, Offline Files, and Roaming User Profiles require an x64-based or x86-based computer, and they are not supported by Windows on ARM (WOA)-based computers.

Software requirements

To designate primary computers, your environment must meet the following requirements:

- The Active Directory Domain Services (AD DS) schema must be updated to include Windows Server 2012 schema additions (installing a Windows Server 2012 domain controller automatically updated the schema). For information about updating the AD DS schema, see What's new for Adprep.exe? and Running Adprep.exe.
- Client computers must run Windows 8 or Windows Server 2012 and be joined to the Active Directory domain that you are managing.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	Supporting Information Workers with Reliable File Services and Storage | What's New in Offline Files (Windows 7 and Windows Server 2008 R2) | What's New in Offline Files for Windows Vista | Changes to Offline Files in Windows Vista (TechNet Magazine)
Deployment	Deploy Folder Redirection, Offline Files, and Roaming User Profiles | Implementing an End-User Data Centralization Solution: Folder Redirection and Offline Files Technology Validation and Deployment | Managing Roaming User Data Deployment Guide | Configuring New Offline Files Features for Windows 7 Computers Step-by-Step Guide | Using Folder Redirection | Implementing Folder Redirection (Windows Server 2003)
Tools and settings	Primary Computer Windows PowerShell cmdlets | Offline Files on MSDN | Offline Files Group Policy Reference (Windows 2000)
Community resources	File Services and Storage Forum | Hey, Scripting Guy! How Can I Work with the Offline Files Feature in Windows? | Hey, Scripting Guy! How Can I Enable and Disable Offline Files?
Related technologies	Active Directory Domain Services Overview | File and Storage Services Overview | Remote

Content type	References
	Desktop Services Overview

Deploy Folder Redirection, Offline Files, and Roaming User Profiles

This series of topics discuss how to use Windows Server to deploy Folder Redirection and Roaming User Profiles (sometimes known as RUP) to Windows client computers. The topics also discuss how to enable new Offline Files (client-side caching or CSC) functionality included with Windows Server 2012 and Windows 8.

- Deploy Folder Redirection with Offline Files
- Deploy Roaming User Profiles
- Deploy Primary Computers for Folder Redirection and Roaming User Profiles
- Enable Optimized Moving of Redirected Folders
- Enable the Always Offline Mode to Provide Faster Access to Files
- Enable Background File Synchronization on Metered Networks
- Disable Offline Files on Individual Redirected Folders
- Troubleshoot User Profiles with Events

See Also

Folder Redirection, Offline Files, and Roaming User Profiles overview

File and Storage Services Overview

Deploy Folder Redirection with Offline Files

This topic describes how to use Windows Server to deploy Folder Redirection with Offline Files to Windows client computers.

In this document

- Prerequisites
- Step 1: Create a folder redirection security group
- Step 2: Create a file share for redirected folders
- Step 3: Create a GPO for Folder Redirection
- Step 4: Configure folder redirection with Offline Files
- Step 5: Enable the Folder Redirection GPO
- Step 6: Test Folder Redirection

- Appendix A: Checklist for deploying Folder Redirection

Prerequisites

Hardware requirements

Folder Redirection requires an x64-based or x86-based computer; it is not supported by Windows® RT.

Software requirements

Folder Redirection has the following software requirements:

- To administer Folder Redirection, you must be signed in as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- Client computers must run Windows 8.1, Windows 8, Windows 7, Windows Vista, Windows XP, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.
- Client computers must be joined to the Active Directory Domain Services (AD DS) that you are managing.
- A computer must be available with Group Policy Management and Active Directory Administration Center installed.
- A file server must be available to host redirected folders.
  - If the file share uses DFS Namespaces, the DFS folders (links) must have a single target to prevent users from making conflicting edits on different servers.
  - If the file share uses DFS Replication to replicate the contents with another server, users must be able to access only the source server to prevent users from making conflicting edits on different servers.

Note

To use new features in Folder Redirection, there are additional client computer and Active Directory schema requirements. For more information, see Folder Redirection, Offline Files, and Roaming User Profiles overview.

Step 1: Create a folder redirection security group

If your environment is not already set up with Folder Redirection, the first step is to create a security group that contains all users to which you want to apply Folder Redirection policy settings.

To create a security group for Folder Redirection

1. Open Server Manager on a computer with Active Directory Administration Center installed.

2. On the **Tools** menu, click **Active Directory Administration Center**. Active Directory Administration Center appears.
3. Right-click the appropriate domain or OU, click **New**, and then click **Group**.
4. In the **Create Group** window, in the **Group** section, specify the following settings:
   - In **Group name**, type the name of the security group, for example: **Folder Redirection Users**.
   - In **Group scope**, click **Security**, and then click **Global**.
5. In the **Members** section, click **Add**. The Select Users, Contacts, Computers, Service Accounts or Groups dialog box appears.
6. Type the names of the users or groups to which you want to deploy Folder Redirection, click **OK**, and then click **OK** again.

Step 2: Create a file share for redirected folders

If you do not already have a file share for redirected folders, use the following procedure to create a file share on a server running Windows Server 2012.

Note

Some functionality might differ or be unavailable if you create the file share on a server running another version of Windows Server.

To create a file share on Windows Server 2012

1. In the Server Manager navigation pane, click **File and Storage Services**, and then click **Shares** to display the Shares page.
2. In the **Shares** tile, click **Tasks**, and then click **New Share**. The New Share Wizard appears.
3. On the **Select Profile** page, click **SMB Share – Quick**. If you have File Server Resource Manager installed and are using folder management properties, instead click **SMB Share - Advanced**.
4. On the **Share Location** page, select the server and volume on which you want to create the share.
5. On the **Share Name** page, type a name for the share (for example, **Users$**) in the **Share name** box.

   Tip

   When creating the share, hide the share by putting a **$** after the share name. This will hide the share from casual browsers.
6. On the **Other Settings** page, optionally select the **Enable access-based enumeration** and **Encrypt data access** checkboxes.
7. On the **Permissions** page, click **Customize permissions…**. The Advanced Security Settings dialog box appears.
8. Click **Disable inheritance**, and then click **Convert inherited permissions into explicit**

permission on this object.

9. Set the permissions as described Table 1 and shown in Figure 1, removing permissions for unlisted groups and accounts, and adding special permissions to the Folder Redirection Users group that you created in Step 1.

Figure 1 Setting the permissions for the redirected folders share

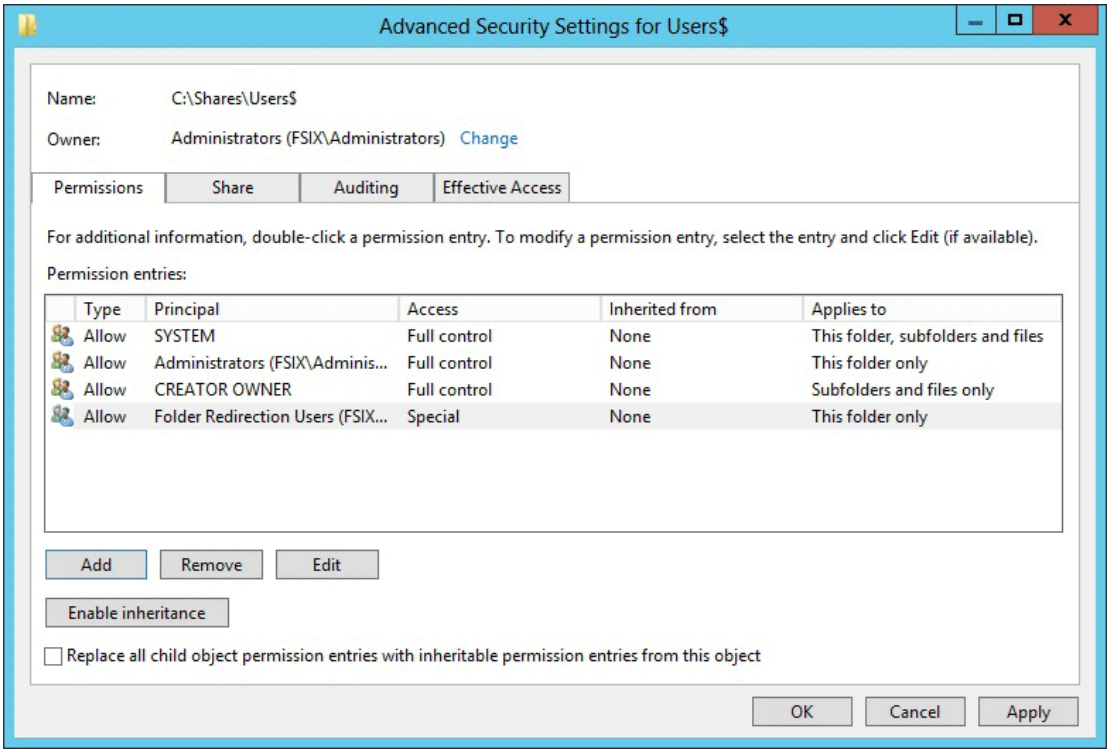

Figure 1 Setting the permissions for the redirected folders share

10. If you chose the **SMB Share - Advanced** profile, on the **Management Properties** page, select the **User Files** Folder Usage value.
11. If you chose the **SMB Share - Advanced** profile, on the **Quota** page, optionally select a quota to apply to users of the share.
12. On the **Confirmation** page, click **Create.**

Table 1 Required permissions for the file share hosting redirected folders

User Account	Access	Applies to
System	Full control	This folder, subfolders and files
Administrators	Full Control	This folder only
Creator/Owner	Full Control	Subfolders and files only
Security group of users needing	List folder / read data[1]	This folder only

to put data on share (Folder Redirection Users)	Create folders / append data[1]	
Other groups and accounts	None (remove)	

1 Advanced permissions

Step 3: Create a GPO for Folder Redirection

If you do not already have a GPO created for Folder Redirection settings, use the following procedure to create one.

To create a GPO for Folder Redirection

1. Open Server Manager on a computer with Group Policy Management installed.
2. From the **Tools** menu click **Group Policy Management**. Group Policy Management appears.
3. Right-click the domain or OU in which you want to setup Folder Redirection and then click **Create a GPO in this domain, and Link it here**.
4. In the **New GPO** dialog box, type a name for the GPO (for example, **Folder Redirection Settings**), and then click **OK**.
5. Right-click the newly created GPO and then clear the **Link Enabled** checkbox. This prevents the GPO from being applied until you finish configuring it.
6. Select the GPO. In the **Security Filtering** section of the **Scope** tab, select **Authenticated Users**, and then click **Remove**.
7. In the **Security Filtering** section, click **Add**.
8. In the **Select User, Computer, or Group** dialog box, type the name of the security group you created in Step 1 (for example, **Folder Redirection Users**), and then click **OK**.

Step 4: Configure folder redirection with Offline Files

After creating a GPO for Folder Redirection settings, edit the Group Policy settings to enable and configure Folder Redirection, as discussed in the following procedure.

Note
Offline Files is enabled by default for redirected folders on Windows client computers, and disabled on computers running Windows Server, unless changed by the user. To use Group Policy to control whether Offline Files is enabled, use the **Allow or disallow use of the Offline Files feature** policy setting.

To configure Folder Redirection in Group Policy

1. In Group Policy Management, right-click the GPO you created (for example, **Folder**

Redirection Settings), and then click **Edit**.

2. In the Group Policy Management Editor window, navigate to **User Configuration**, then **Policies**, then **Windows Settings**, and then **Folder Redirection**.
3. Right-click a folder that you want to redirect (for example, **Documents**), and then click **Properties**.
4. In the **Properties** dialog box, from the **Setting** box click **Basic - Redirect everyone's folder to the same location**.

   Note
   To apply Folder Redirection to client computers running Windows XP or Windows Server 2003, click the **Settings** tab and select the **Also apply redirection policy to Windows 2000, Windows 2000 Server, Windows XP, and Windows Server 2003 operating systems** checkbox.
5. In the **Target folder location** section, click **Create a folder for each user under the root path** and then in the **Root Path** box, type the path to the file share storing redirected folders, for example: **\\fs1.corp.contoso.com\users$**
6. Click the **Settings** tab, and in the **Policy Removal** section, optionally click **Redirect the folder back to the local userprofile location when the policy is removed** (this setting can help make Folder Redirection behave more predictably for adminisitrators and users).
7. Click **OK**, and then click **Yes** in the Warning dialog box.

Step 5: Enable the Folder Redirection GPO

Once you have completed configuring the Folder Redirection Group Policy settings, the next step is to enable the GPO, permitting it to be applied to affected users.

Tip
If you plan to implement primary computer support or other policy settings, do so now, before you enable the GPO. This prevents user data from being copied to non-primary computers before primary computer support is enabled.

To enable the Folder Redirection GPO

1. Open Group Policy Management.
2. Right-click the GPO that you created, and then click **Link Enabled**. A checkbox appears next to the menu item.

Step 6: Test Folder Redirection

To test Folder Redirection, sign in to a computer with a user account configured for Folder Redirection. Then confirm that the folders and profiles are redirected.

To test Folder Redirection

1. Sign in to a primary computer (if you enabled primary computer support) with a user account for which you have enabled Folder Redirection.
2. If the user has previously signed in to the computer, open an elevated command prompt, and then type the following command to ensure that the latest Group Policy settings are applied to the client computer:

   ```
   gpupdate /force
   ```
3. Open File Explorer.
4. Right-click a redirected folder (for example, the My Documents folder in the Documents library), and then click **Properties**.
5. Click the **Location** tab, and confirm that the path displays the file share you specified instead of a local path.

Appendix A: Checklist for deploying Folder Redirection

	1. Prepare domain
	• Join computers to domain
	• Create user accounts
	2. Create security group for Folder Redirection
	• Group name:
	• Members:
	3. Create a file share for redirected folders
	• File share name:
	4. Create a GPO for Folder Redirection
	• GPO name:
	5. Configure Folder Redirection and Offline Files policy settings
	• Redirected folders:
	• Windows 2000, Windows XP, and Windows Server 2003 support enabled?
	• Offline Files enabled? (enabled by default on Windows client computers)

	• Always Offline Mode enabled?
	• Background file synchronization enabled?
	• Optimized Move of redirected folders enabled?
	6. (Optional) Enable primary computer support
	• Designate primary computers for users • Location of user and primary computer mappings:
	• (Optional) Enable primary computer support for Folder Redirection • Computer-based or User-based?
	• (Optional) Enable primary computer support for Roaming User Profiles
	7. Enable the Folder Redirection GPO
	8. Test Folder Redirection

See Also

Deploy Folder Redirection, Offline Files, and Roaming User Profiles

Deploy Primary Computers for Folder Redirection and Roaming User Profiles

Microsoft's Support Statement Around Replicated User Profile Data

How to Add and Remove Apps

Troubleshooting packaging, deployment, and query of Windows Runtime-based apps

Deploy Roaming User Profiles

This topic describes how to use Windows Server to deploy Roaming User Profiles to Windows client computers.

For a list of recent changes to this topic, see the Change History section of this topic.

In this document

- Prerequisites
- Step 1: Enable the use of separate profile versions
- Step 2: Create a Roaming User Profiles security group
- Step 3: Create a file share for roaming user profiles
- Step 4: Optionally create a GPO for Roaming User Profiles

- Step 5: Optionally set up Roaming User Profiles on user accounts
- Step 6: Optionally set up Roaming User Profiles on computers
- Step 7: Enable the Roaming User Profiles GPO
- Step 8: Test Roaming User Profiles
- Appendix A: Checklist for deploying Roaming User Profiles
- Appendix B: Profile version reference information

Prerequisites

Hardware requirements

Roaming User Profiles require an x64-based or x86-based computer; it is not supported by Windows® RT.

Software requirements

Roaming User Profiles has the following software requirements:

- If you are deploying Roaming User Profiles with Folder Redirection in an environment with existing local user profiles, deploy Folder Redirection before Roaming User Profiles to minimize the size of roaming profiles. After the existing user folders have been successfully redirected, you can deploy Roaming User Profiles.
- To administer Roaming User Profiles, you must be signed in as a member of the Domain Administrators security group, the Enterprise Administrators security group, or the Group Policy Creator Owners security group.
- Client computers must run Windows 8.1, Windows 8, Windows 7, Windows Vista, Windows XP, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. Windows XP and Windows Server 2003 support Roaming User Profiles only when set up on user accounts by using AD DS - they don’t support enabling Roaming User Profiles on a per-computer basis by using Group Policy.
- Client computers must be joined to the Active Directory Domain Services (AD DS) that you are managing.
- A computer must be available with Group Policy Management and Active Directory Administration Center installed.
- A file server must be available to host roaming user profiles.
    - If the file share uses DFS Namespaces, the DFS folders (links) must have a single target to prevent users from making conflicting edits on different servers.
    - If the file share uses DFS Replication to replicate the contents with another server, users must be able to access only the source server to prevent users from making conflicting edits on different servers.

Note

To use new features in Roaming User Profiles, there are additional client computer and Active Directory schema requirements. For more information, see Folder Redirection, Offline Files, and Roaming User Profiles overview.

Considerations when using Roaming User Profiles on multiple versions of Windows

If you decide to use Roaming User Profiles across multiple versions of Windows, we recommend taking the following actions:

- Configure Windows to maintain separate profile versions for each operating system version. This helps prevent undesirable and unpredictable issues such as profile corruption.
- Use Folder Redirection to store user files such as documents and pictures outside of user profiles. This enables the same files to be available to users across operating system versions. It also keeps profiles small and sign-ins quick.
- Allocate sufficient storage for Roaming User Profiles. If you support two operating system versions, profiles will double in number (and thus total space consumed) because a separate profile is maintained for each operating system version.
- Don't use Roaming User Profiles across computers running Windows Vista/Windows Server 2008 and Windows 7/Windows Server 2008 R2. Roaming between these operating system versions isn't supported due to incompatibilities in their profile versions.
- Inform your users that changes made on one operating system version won't roam to another operating system version.
- When moving your environment to a new version of Windows users will receive a new, empty profile. There isn't a supported method of migrating user profiles from one operating system version to another.

Step 1: Enable the use of separate profile versions

If you are deploying Roaming User Profiles on computers running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012, we recommend making a couple of changes to your Windows environment prior to deploying. These changes help ensure that future operating system upgrades go smoothly, and facilitate the ability to simultaneously run multiple versions of Windows with Roaming User Profiles.

To make these changes, use the following procedure.

Enable the use of separate profiles for each version of Windows

1. Download and install the appropriate software update on all computers on which you're going to use roaming, mandatory, super-mandatory, or domain default profiles:
   - Windows 8.1, or Windows Server 2012 R2: install the software update described in article 2887595 in the Microsoft Knowledge Base (when released).
   - Windows 8 or Windows Server 2012: install the software update described in article

2887239 in the Microsoft Knowledge Base.

2. On all computers running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012 on which you will use Roaming User Profiles, use Registry Editor or Group Policy to create the following registry key DWORD Value and set it to **1**. For information about creating registry keys by using Group Policy, see Configure a Registry Item.

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ProfSvc\
   Parameters\UseProfilePathExtensionVersion
   ```

   Caution
   Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

3. Restart the computers.

Step 2: Create a Roaming User Profiles security group

If your environment is not already set up with Roaming User Profiles, the first step is to create a security group that contains all users and/or computers to which you want to apply Roaming User Profiles policy settings.

- Administrators of general-purpose roaming user profiles deployments typically create a security group for users.
- Administrators of Remote Desktop Services or virtualized desktop deployments typically use a security group for users and the shared computers.

To create a security group for Roaming User Profiles

1. Open Server Manager on a Windows Server 2012 R2 or Windows Server 2012 computer with Active Directory Administration Center installed.
2. On the **Tools** menu, click **Active Directory Administration Center**. Active Directory Administration Center appears.
3. Right-click the appropriate domain or OU, click **New**, and then click **Group**.
4. In the **Create Group** window, in the **Group** section, specify the following settings:
   - In **Group name**, type the name of the security group, for example: **Roaming User Profiles Users and Computers**.
   - In **Group scope**, click **Security**, and then click **Global**.
5. In the **Members** section, click **Add**. The Select Users, Contacts, Computers, Service Accounts or Groups dialog box appears.
6. If you want to include computer accounts in the security group, click **Object Types**, select the **Computers** check box and then click **OK**.
7. Type the names of the users, groups, and/or computers to which you want to deploy

Roaming User Profiles, click **OK**, and then click **OK** again.

Step 3: Create a file share for roaming user profiles

If you do not already have a separate file share for roaming user profiles (independent from any shares for redirected folders to prevent inadvertent caching of the roaming profile folder), use the following procedure to create a file share on a server running Windows Server 2012.

Note
Some functionality might differ or be unavailable if you create the file share on a server running another version of Windows Server.

To create a file share on Windows Server 2012

1. In the Server Manager navigation pane, click **File and Storage Services**, and then click **Shares** to display the Shares page.
2. In the Shares tile, click **Tasks**, and then click **New Share**. The New Share Wizard appears.
3. On the **Select Profile** page, click **SMB Share – Quick**. If you have File Server Resource Manager installed and are using folder management properties, instead click **SMB Share - Advanced**.
4. On the **Share Location** page, select the server and volume on which you want to create the share.
5. On the **Share Name** page, type a name for the share (for example, **User Profiles$**) in the **Share name** box.

   Tip
   When creating the share, hide the share by putting a **$** after the share name. This hides the share from casual browsers.
6. On the **Other Settings** page, optionally select the **Enable access-based enumeration** and **Encrypt data access** checkboxes.
7. On the **Permissions** page, click **Customize permissions…**. The Advanced Security Settings dialog box appears.
8. Click **Disable inheritance**, and then click **Convert inherited permissions into explicit permission on this object**.
9. Set the permissions as described Table 1 and shown in Figure 1, removing permissions for unlisted groups and accounts, and adding special permissions to the Roaming User Profiles Users and Computers group that you created in Step 1.

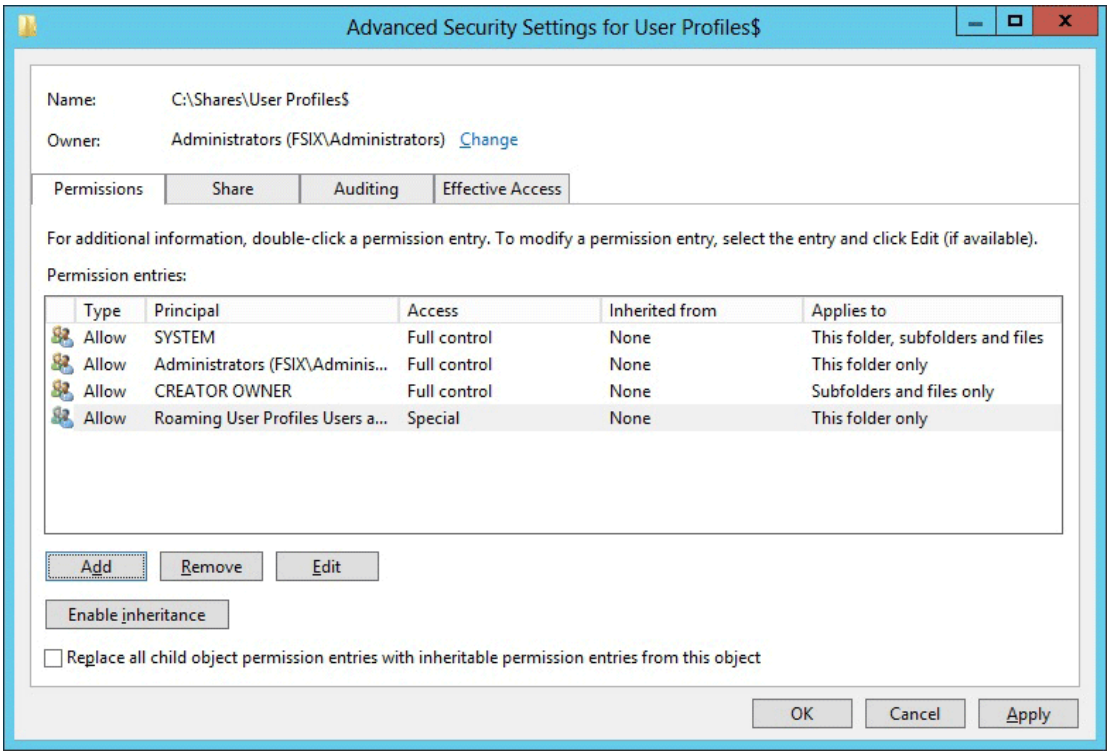

Figure 1 Setting the permissions for the roaming user profiles share

10. If you chose the **SMB Share - Advanced** profile, on the **Management Properties** page, select the **User Files** Folder Usage value.
11. If you chose the **SMB Share - Advanced** profile, on the **Quota** page, optionally select a quota to apply to users of the share.
12. On the **Confirmation** page, click **Create.**

Table 1 Required permissions for the file share hosting roaming user profiles

User Account	Access	Applies to
System	Full control	This folder, subfolders and files
Administrators	Full Control	This folder only
Creator/Owner	Full Control	Subfolders and files only
Security group of users needing to put data on share (Roaming User Profiles Users and Computers)	List folder / read data[1] Create folders / append data[1]	This folder only
Other groups and accounts	None (remove)	

1 Advanced permissions

Step 4: Optionally create a GPO for Roaming User Profiles

If you do not already have a GPO created for Roaming User Profiles settings, use the following procedure to create an empty GPO for use with Roaming User Profiles. This GPO allows you to configure Roaming User Profiles settings (such as primary computer support, which is discussed separately), and can also be used to enable Roaming User Profiles on computers, as is typically done when deploying in virtualized desktop environments or with Remote Desktop Services.

To create a GPO for Roaming User Profiles

1. Open Server Manager on a computer with Group Policy Management installed.
2. From the **Tools** menu click **Group Policy Management**. Group Policy Management appears.
3. Right-click the domain or OU in which you want to setup Roaming User Profiles and then click **Create a GPO in this domain, and Link it here**.
4. In the **New GPO** dialog box, type a name for the GPO (for example, **Roaming User Profile Settings**), and then click **OK**.
5. Right-click the newly created GPO and then clear the **Link Enabled** checkbox. This prevents the GPO from being applied until you finish configuring it.
6. Select the GPO. In the **Security Filtering** section of the **Scope** tab, select **Authenticated Users**, and then click **Remove**.
7. In the **Security Filtering** section, click **Add**.
8. In the **Select User, Computer, or Group** dialog box, type the name of the security group you created in Step 1 (for example, **Roaming User Profiles Users and Computers**), and then click **OK**.

Step 5: Optionally set up Roaming User Profiles on user accounts

If you are deploying Roaming User Profiles to user accounts, use the following procedure to specify roaming user profiles for user accounts in Active Directory Domain Services. If you are deploying Roaming User Profiles to computers, as is typically done for Remote Desktop Services or virtualized desktop deployments, instead use the procedure documented in Step 5 of this topic.

Note

If you set up Roaming User Profiles on user accounts by using Active Directory and on computers by using Group Policy, the computer-based policy setting takes precedence.

To set up Roaming User Profiles on user accounts

1. In Active Directory Administration Center, navigate to the **Users** container (or OU) in the appropriate domain.
2. Select all users to which you want to assign a roaming user profile, right-click the users and then click **Properties**.
3. In the **Profile** section, select the **Profile path:** checkbox and then enter the path to the file share where you want to store the user's roaming user profile, followed by **%username%** (which is automatically replaced with the user name the first time the user signs in). For example:

   \\fs1.corp.contoso.com\User Profiles$\%username%

   To specify a mandatory roaming user profile, specify the path to the NTuser.man file that you created previously, for example, **\\fs1.corp.contoso.com\User Profiles$\default**. For more information, see Creating a Mandatory User Profile.
4. Click **OK**.

Notes

By default, deployment of all Windows® Runtime-based (Windows Store) apps is allowed when using Roaming User Profiles. However, when using a special profile, apps are not deployed by default. Special profiles are user profiles where changes are discarded after the user signs out:

- Roaming user profiles to which the **Delete cached copies of roaming profiles** Group Policy setting applies (located in Computer Configuration\Policies\Administrative Templates\System\User Profiles)
- Mandatory user profiles and super-mandatory profiles, created by an administrator
- Temporary user profiles, created when an error prevents the correct profile from loading
- User profiles for the Guest account and members of the Guests group

To remove restrictions on app deployment for special profiles, enable the **Allow deployment operations in special profiles** policy setting (located in Computer Configuration\Policies\Administrative Templates\Windows Components\App Package Deployment). However, deployed apps in this scenario will leave some data stored on the computer, which could accumulate, for example, if there are hundreds of users of a single computer. To cleanup apps, locate or develop a tool that uses the CleanupPackageForUserAsync API to cleanup app packages for users who no longer have a profile on the computer.

For additional background information about Windows Store apps, see **Managing Client Access to the Windows Store**.

Step 6: Optionally set up Roaming User Profiles on computers

If you are deploying Roaming User Profiles to computers, as is typically done for Remote Desktop Services or virtualized desktop deployments, use the following procedure. If you are deploying Roaming User Profiles to user accounts, use the procedure described in Step 4 of this topic.

You can use Group Policy to apply Roaming User Profiles to computers running Windows 8.1, Windows 8, Windows 7, Windows Vista, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008.

Note

If you set up Roaming User Profiles on computers by using Group Policy and on user accounts by using Active Directory, the computer-based policy setting takes precedence.

To set up Roaming User Profiles on computers

1. Open Server Manager on a computer with Group Policy Management installed.
2. From the **Tools** menu click **Group Policy Management**. Group Policy Management appears.
3. In Group Policy Management, right-click the GPO you created in Step 3 (for example, **Roaming User Profiles Settings**), and then click **Edit**.
4. In the Group Policy Management Editor window, navigate to **Computer Configuration**, then **Policies**, then **Administrative Templates**, then **System**, and then **User Profiles**.
5. Right-click **Set roaming profile path for all users logging onto this computer** and then click **Edit**.

   Tip

   A user's home folder, if configured, is the default folder used by some programs such as Windows PowerShell. You can configure an alternative local or network location on a per-user basis by using the **Home folder** section of the user account properties in AD DS. To configure the home folder location for all users of a computer running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012 in a virtual desktop environment, enable the **Set user home folder** policy setting, and then specify the file share and drive letter to map (or specify a local folder). Do not use environment variables or ellipses. The user's alias is appended to the end of the path specified during user sign on.
6. In the **Properties** dialog box, click **Enabled**
7. In the **Users logging onto this computer should use this roaming profile path** box, enter the path to the file share where you want to store the user's roaming user profile, followed by **%username%** (which is automatically replaced with the user name the first time the user signs in). For example:

   \\fs1.corp.contoso.com\User Profiles$\%username%

   To specify a mandatory roaming user profile, which is a preconfigured profile to which

users cannot make permanent changes (changes are reset when the user signs out), specify the path to the NTuser.man file that you created previously, for example, **\\fs1.corp.contoso.com\User Profiles$\default**. For more information, see Creating a Mandatory User Profile.

8. Click **OK**.

Step 7: Enable the Roaming User Profiles GPO

If you set up Roaming User Profiles on computers by using Group Policy, or if you customized other Roaming User Profiles settings by using Group Policy, the next step is to enable the GPO, permitting it to be applied to affected users.

Tip

If you plan to implement primary computer support or other policy settings, do so now, before you enable the GPO. This prevents user data from being copied to non-primary computers before primary computer support is enabled.

To enable the Roaming User Profile GPO

1. Open Group Policy Management.
2. Right-click the GPO that you created and then click **Link Enabled**. A checkbox appears next to the menu item.

Step 8: Test Roaming User Profiles

To test Roaming User Profiles, sign in to a computer with a user account configured for Roaming User Profiles, or sign in to a computer configured for Roaming User Profiles. Then confirm that the profile is redirected.

To test Roaming User Profiles

1. Sign in to a primary computer (if you enabled primary computer support) with a user account for which you have enabled Roaming User Profiles enabled. If you enabled Roaming User Profiles on specific computers, sign in to one of these computers.
2. If the user has previously signed in to the computer, open an elevated command prompt, and then type the following command to ensure that the latest Group Policy settings are applied to the client computer:

   ```
   GpUpdate /Force
   ```

3. To confirm that the user profile is roaming, open **Control Panel**, click **System and Security**, click **System**, click **Advanced System Settings**, click **Settings** in the User Profiles section and then look for **Roaming** in the **Type** column.

Appendix A: Checklist for deploying Roaming User Profiles

	1. Prepare domain
	• Join computers to domain
	• Enable the use of separate profile versions
	• Create user accounts
	• (Optional) Deploy Folder Redirection
	2. Create security group for Roaming User Profiles
	• Group name:
	• Members:
	3. Create a file share for Roaming User Profiles
	• File share name:
	4. Create a GPO for Roaming User Profiles
	• GPO name:
	5. Configure Roaming User Profiles policy settings
	6. Enable Roaming User Profiles:
	• Enabled in AD DS on user accounts?
	• Enabled in Group Policy on computer accounts?
	7. Enable primary computer support
	• Designate primary computers for users • Location of user and primary computer mappings:
	• (Optional) Enable primary computer support for Folder Redirection • Computer-based or User-based?
	• (Optional) Enable primary computer support for Roaming User Profiles
	8. Enable the Roaming User Profiles GPO

	9. Test Roaming User Profiles

Appendix B: Profile version reference information

The following table lists the location of Roaming User Profiles on various versions of Windows.

Operating system version	Roaming User Profile location
Windows XP and Windows Server 2003	\\<*servername*>\<*fileshare*>\<*username*>
Windows Vista and Windows Server 2008	\\<*servername*>\<*fileshare*>\<*username*>.v2
Windows 7 and Windows Server 2008 R2	\\<*servername*>\<*fileshare*>\<*username*>.v2
Windows 8 and Windows Server 2012	\\<*servername*>\<*fileshare*>\<*username*>.v3 (after the software update and registry key are applied) \\<*servername*>\<*fileshare*>\<*username*>.v2 (before the software update and registry key are applied)
Windows 8.1 and Windows Server 2012 R2	\\<servername>\<fileshare>\<username>.v4 (after the software update and registry key are applied) \\<*servername*>\<*fileshare*>\<*username*>.v2 (before the software update and registry key are applied)

Change History

The following table summarizes some of the most important changes to this topic.

Date	Description	Reason
October 9th, 2013	Revised for Windows Server 2012 R2 and Windows 8.1, clarified a few things, and added the Considerations when using Roaming User Profiles on multiple versions of Windows and Appendix B: Profile version reference information sections.	Updates for new version; customer feedback

See Also

Deploy Folder Redirection, Offline Files, and Roaming User Profiles

Deploy Primary Computers for Folder Redirection and Roaming User Profiles

Implementing User State Management

Microsoft's Support Statement Around Replicated User Profile Data

How to Add and Remove Apps

Troubleshooting packaging, deployment, and query of Windows Runtime-based apps

Deploy Primary Computers for Folder Redirection and Roaming User Profiles

This topic describes how to enable primary computer support and designate primary computers for users. Doing so enables you to control which computers use Folder Redirection and Roaming User Profiles.

Important
When enabling primary computer support for Roaming User Profiles, always enable primary computer support for Folder Redirection as well. This keeps documents and other user files out of the user profiles, which helps profiles remain small and sign on times stay fast.

In this document

- Prerequisites
- Step 1: Designate primary computers for users
- Step 2: Optionally enable primary computers for Folder Redirection in Group Policy
- Step 3: Optionally enable primary computers for Roaming User Profiles in Group Policy
- Step 4: Enable the GPO
- Step 5: Test primary computer function

Prerequisites

Software requirements

Primary computer support has the following requirements:

- The Active Directory Domain Services (AD DS) schema must be updated to include Windows Server 2012 schema additions (installing a Windows Server 2012 domain controller automatically updates the schema). For information about updating the AD DS schema, see What's new for Adprep.exe? and Running Adprep.exe.

- Client computers must run Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012.

Tip

Although primary computer support requires Folder Redirection and/or Roaming User Profiles, if you are deploying these technologies for the first time, it is best to set up primary computer support before enabling the GPOs that configure Folder Redirection and Roaming User Profiles. This prevents user data from being copied to non-primary computers before primary computer support is enabled. For configuration information, see Deploy Folder Redirection, Offline Files, and Roaming User Profiles and Deploy Roaming User Profiles.

Step 1: Designate primary computers for users

The first step in deploying primary computers support is designating the primary computers for each user. To do so, use Active Directory Administration Center to obtain the distinguished name of the relevant computers and then set the **msDs-PrimaryComputer** attribute.

Tip

To use Windows PowerShell to accomplish this step, download and install the Primary Computer Windows PowerShell cmdlets and then use the Add-XADPrimaryComputer cmdlet.

To specify the primary computers for users

1. Open Server Manager on a computer with Active Directory Administration Tools installed.
2. On the **Tools** menu, click **Active Directory Administration Center**. Active Directory Administration Center appears.
3. Navigate to the **Computers** container in the appropriate domain.
4. Right-click a computer that you want to designate as a primary computer and then click **Properties**.
5. In the Navigation pane, click **Extensions**.
6. Click the **Attribute Editor** tab, scroll to **distinguishedName**, click **View**, right-click the value listed, click **Copy**, click **OK**, and then click **Cancel**.
7. Navigate to the **Users** container in the appropriate domain, right-click the user to which you want to assign the computer, and then click **Properties**.
8. In the Navigation pane, click **Extensions**.
9. Click the **Attribute Editor** tab, select **msDs-PrimaryComputer** and then click **Edit**. The Multi-valued String Editor dialog box appears.
10. Right-click the text box, click **Paste**, click **Add**, click **OK**, and then click **OK** again.

Step 2: Optionally enable primary computers for Folder Redirection in Group Policy

The next step is to optionally configure Group Policy to enable primary computer support for Folder Redirection. Doing so enables a user's folders to be redirected on computers designated as the user's primary computers, but not on any other computers. You can control primary computers for Folder Redirection on a per-computer basis, or a per-user basis.

To enable primary computers for Folder Redirection

1. In Group Policy Management, right-click the GPO you created when doing the initial configuration of Folder Redirection and/or Roaming User Profiles (for example, **Folder Redirection Settings** or **Roaming User Profiles Settings**), and then click **Edit**.
2. To enable primary computers support using computer-based Group Policy, navigate to **Computer Configuration**. For user-based Group Policy, navigate to **User Configuration**.
   - Computer-based Group Policy applies primary computer processing to all computers to which the GPO applies, affecting all users of the computers.
   - User-based Group Policy to applies primary computer processing to all user accounts to which the GPO applies, affecting all computers to which the users sign on.
3. Under **Computer Configuration** or **User Configuration**, navigate to **Policies**, then **Administrative Templates**, then **System**, then **Folder Redirection**.
4. Right-click **Redirect folders on primary computers only**, and then click **Edit**.
5. Click **Enabled**, and then click **OK**.

Step 3: Optionally enable primary computers for Roaming User Profiles in Group Policy

The next step is to optionally configure Group Policy to enable primary computer support for Roaming User Profiles. Doing so enables a user's profile to roam on computers designated as the user's primary computers, but not on any other computers.

To enable primary computers for Roaming User Profiles

1. Enable primary computer support for Folder Redirection, if you have not already.

   This keeps documents and other user files out of the user profiles, which helps profiles remain small and sign on times stay fast.
2. In Group Policy Management, right-click the GPO you created (for example, **Folder Redirection and Roaming User Profiles Settings**), and then click **Edit**.
3. Navigate to **Computer Configuration**, then **Policies**, then **Administrative Templates**, then **System**, and then **User Profiles**.
4. Right-click **Download roaming profiles on primary computers only,** and then click **Edit**.

5. Click **Enabled**, and then click **OK**.

Step 4: Enable the GPO

Once you have completed configuring Folder Redirection and Roaming User Profiles, enable the GPO if you have not already. Doing so permits it to be applied to affected users and computers.

To enable the Folder Redirection and/or Roaming User Profiles GPOs

1. Open Group Policy Management
2. Right-click the GPOs that you created, and then click **Link Enabled**. A checkbox should appear next to the menu item.

Step 5: Test primary computer function

To test primary computer support, sign in to a primary computer, confirm that the folders and profiles are redirected, then sign in to a non-primary computer and confirm that the folders and profiles are not redirected.

To test primary computer functionality

1. Sign in to a designated primary computer with a user account for which you have enabled Folder Redirection and/or Roaming User Profiles.
2. If the user account has signed on to the computer previously, open a Windows PowerShell session or Command Prompt window as an administrator, type the following command and then sign off when prompted to ensure that the latest Group Policy settings are applied to the client computer:

   ```
   Gpupdate /force
   ```
3. Open File Explorer.
4. Right-click a redirected folder (for example, the My Documents folder in the Documents library), and then click **Properties**.
5. Click the **Location** tab, and confirm that the path displays the file share you specified instead of a local path. To confirm that the user profile is roaming, open **Control Panel**, click **System and Security**, click **System**, click **Advanced System Settings**, click **Settings** in the User Profiles section and then look for **Roaming** in the **Type** column.
6. Sign in with the same user account to a computer that is not designated as the user's primary computer.
7. Repeat steps 2-5, instead looking for local paths and a **Local** profile type.

Note

If folders were redirected on a computer before you enabled primary computer support, the folders will remain redirected unless the following setting is configured in each folder's folder redirection policy setting: **Redirect the folder back to the local userprofile location when the policy is removed**. Similarly, profiles that were

previously roaming on a particular computer will show **Roaming** in the **Type** columns; however, the **Status** column will show **Local**.

See Also

Deploy Folder Redirection with Offline Files

Deploy Roaming User Profiles

Folder Redirection, Offline Files, and Roaming User Profiles overview

Primary Computer Windows PowerShell cmdlets

Enable Advanced Offline Files Functionality

This series of topics discusses how to enable new Offline Files (client-side caching or CSC) functionality included with Windows Server 2012 and Windows 8.

- Enable Optimized Moving of Redirected Folders
- Enable the Always Offline Mode to Provide Faster Access to Files
- Enable Background File Synchronization on Metered Networks
- Disable Offline Files on Individual Redirected Folders

See Also

Deploy Folder Redirection, Offline Files, and Roaming User Profiles

Deploy Primary Computers for Folder Redirection and Roaming User Profiles

Folder Redirection, Offline Files, and Roaming User Profiles overview

Enable Optimized Moving of Redirected Folders

This topic describes how to perform an optimized move of redirected folders (Folder Redirection) to a new file share. If you enable this policy setting, when an administrator moves the file share hosting redirected folders and updates the target path of the redirected folders in Group Policy, the cached content is simply renamed in the local Offline Files cache without any delays or potential data loss for the user.

Previously, administrators could change the target path of the redirected folders in Group Policy and let the client computers copy the files at the affected user's next sign in, causing a delayed sign in. Alternatively, administrators could move the file share and update the target path of the redirected folders in Group Policy. However, any changes made locally on the client computers between the start of the move and the first sync after the move would be lost.

In this topic

- Prerequisites
- Step 1: Enable optimized move in Group Policy
- Step 2: Relocate the file share for redirected folders

Prerequisites

Optimized move has the following requirements:

- Folder Redirection must be setup. For more information see Deploy Folder Redirection with Offline Files.
- Client computers must run Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012

Step 1: Enable optimized move in Group Policy

To optimize the relocation of Folder Redirection data, use Group Policy to enable the **Enable optimized move of contents in Offline Files cache on Folder Redirection server path change** policy setting for the appropriate Group Policy Object (GPO). Configuring this policy setting to **Disabled** or **Not Configured** causes the client to copy all the Folder Redirection content to the new location and then delete the content from the old location if the server path changes.

To enable optimized moving of redirected folders

1. In Group Policy Management, right-click the GPO you created for Folder Redirection settings (for example, **Folder Redirection and Roaming User Profiles Settings**), and then click **Edit**.
2. Under **User Configuration**, navigate to **Policies**, then **Administrative Templates**, then **System**, then **Folder Redirection**.
3. Right-click **Enable optimized move of contents in Offline Files cache on Folder Redirection server path change**, and then click **Edit**.
4. Click **Enabled**, and then click **OK**.

Step 2: Relocate the file share for redirected folders

When moving the file share that contains users' redirected folders, it is import to take precautions to ensure that the folders are relocated properly.

Important

If the users' files are in use or if the full file state is not preserved in the move, users might experience poor performance as the files are copied over the network, synchronization conflicts generated by Offline Files, or even data loss.

1. Notify users in advance that the server hosting their redirected folders will change and recommend that they perform the following actions:
   - Synchronize the contents of their Offline Files cache and resolve any conflicts.
   - Open an elevated command prompt, type **GpUpdate /Target:User /Force**, and then sign out and sign back in to ensure that the latest Group Policy settings are applied to the client computer

     Note
     By default, client computers update Group Policy every 90 minutes, so if you allow sufficient time for client computers to receive updated policy, you do not need to ask users to use GpUpdate.
2. Remove the file share from the server to ensure that no files in the file share are in use. To do so in Server Manager, on the **Shares** page of File and Storage Services, right-click the appropriate file share and then click **Remove**.

   Users will work offline using Offline Files until the move is complete and they receive the updated Folder Redirection settings via Group Policy.
3. Using an account with backup privileges, move the contents of the file share to the new location using a method that preserves file timestamps, such as a backup and restore utility. To use the **Robocopy** command, open an elevated command prompt, and then type the following command, where *<Source>* is the current location of the file share, and *<Destination>* is the new location:

   ```
   Robocopy /B <Source> <Destination> /Copyall /MIR /EFSRAW
   ```

   Note
   The **Xcopy** command does not preserve all of the file state.
4. Edit the Folder Redirection policy settings, updating the target folder location for each redirected folder that you want to relocate. For more information, see Step 4 of Deploy Folder Redirection with Offline Files.
5. Notify users that the server hosting their redirected folders has changed, and that they should use the **GpUpdate /Target:User /Force** command, and then sign out and sign back in to get the updated configuration and resume proper file synchronization.

   Users should sign on to all machines at least once to ensure that the data gets properly relocated in each Offline Files cache.

See Also

Deploy Folder Redirection with Offline Files

Deploy Roaming User Profiles

Folder Redirection, Offline Files, and Roaming User Profiles overview

Enable the Always Offline Mode to Provide Faster Access to Files

This document describes how to use the Always Offline mode of Offline Files to provide faster access to cached files and redirected folders. Always Offline also provides lower bandwidth usage because users are always working offline, even when they are connected through a high-speed network connection.

In this document

- Prerequisites
- Enabling the Always Offline mode

Prerequisites

To enable the Always Offline mode, your environment must meet the following prerequisites.

- An Active Directory Domain Services (AD DS) domain, with client computers joined to the domain. There are no forest or domain functional-level requirements or schema requirements.
- Client computers running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012. (Client computers running earlier versions of Windows might continue to transition to Online mode on very high-speed network connections.)
- A computer with Group Policy Management installed.

Enabling the Always Offline mode

To enable the Always Offline mode, use Group Policy to enable the **Configure slow-link mode** policy setting and set the latency to **1** (millisecond). Doing so causes client computers running Windows 8 or Windows Server 2012 to automatically use the Always Offline mode.

Note

Computers running Windows 7, Windows Vista, Windows Server 2008 R2, or Windows Server 2008 might continue to transition to the Online mode if the latency of the network connection drops below one millisecond.

To enable the Always Offline Mode

1. Open **Group Policy Management**.
2. To optionally create a new Group Policy Object (GPO) for Offline Files settings, right-click the appropriate domain or organizational unit (OU), and then click **Create a GPO in this domain, and link it here**.
3. In the console tree, right-click the GPO for which you want to configure the Offline Files settings and then click **Edit**. The **Group Policy Management Editor** appears.
4. In the console tree, under **Computer Configuration**, expand **Policies**, expand

Administrative Templates, expand **Network**, and expand **Offline Files**.

5. Right-click **Configure slow-link mode**, and then click **Edit**. The **Configure slow-link mode** window appears.
6. Click **Enabled**.
7. In the **Options** box, click **Show**. The **Show Contents window** appears.
8. In the **Value name** box, specify the file share for which you want to enable Always Offline mode.
9. To enable Always Offline mode on all file shares, type *.
10. In the **Value** box, type **Latency=1** to set the latency threshold to one millisecond, and then click **OK**.

Note

By default, when in Always Offline mode, Windows synchronizes files in the Offline Files cache in the background every two hours. To change this value, use the **Configure Background Sync** policy setting.

See Also

Folder Redirection, Offline Files, and Roaming User Profiles overview

Deploy Folder Redirection with Offline Files

Enable Background File Synchronization on Metered Networks

This document describes how to enable background file synchronization of Offline Files while using metered connections that have usage limits, and while roaming on another provider's network.

Windows 8.1, Windows 8, Windows Server 2012 R2, and Windows Server 2012 support cost-aware synchronization by automatically tracking roaming and bandwidth usage limits while on metered connections. By default, when the user is using a metered connection (such as a 4G mobile network) and is near or over their bandwidth limit or roaming on another provider's network, Windows switches to Offline mode and disables background synchronization. Users can still manually initiate synchronization, and administrators can override cost-aware synchronization for specific users, such as executives.

In this document

- Prerequisites
- Enable background file synchronization of Offline Files on metered networks

Prerequisites

To enable the background file synchronization of Offline Files for groups of users when they are using metered connections, your environment must meet the following prerequisites.

- An Active Directory Domain Services (AD DS) domain, with client computers joined to the domain. There are no forest or domain functional-level requirements or schema requirements.
- Client computers running Windows 8 or Windows Server 2012.
- A computer with Group Policy Management installed.

Tip

To ensure client computers that are using metered network connections properly transition to the Offline (Slow Connection) mode instead of the more bandwidth-intensive Online mode, test the latencies of your metered network connections. Metered network connections typically have a round-trip network latency of 50 milliseconds or more. However, if the latency is less (faster) than 35 milliseconds on computers running Windows 8 or Windows Server 2012, you can adjust the threshold by using the **Configure slow-link mode** Group Policy setting, or you can enable the Always Offline mode.

Enable background file synchronization of Offline Files on metered networks

To enable background file synchronization of Offline Files for a group of users while using metered connections, use Group Policy to enable the **Enable file synchronization on costed networks** policy setting for the appropriate Group Policy Object (GPO).

To enable background file synchronization of Offline Files on metered networks

1. Open **Group Policy Management**.
2. To optionally create a new Group Policy Object (GPO) that specifies which users should perform background synchronization on metered networks, right-click the appropriate domain or organizational unit, and then click **Create a GPO in this domain, and link it here**.
3. In the console tree, right-click the GPO for which you want to configure the Offline Files settings, and then click **Edit**. The **Group Policy Management Editor** appears.
4. In the console tree, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Network**, and expand **Offline Files**.
5. Right-click **Enable file synchronization on costed networks**, and then click **Edit**. The **Enable file synchronization on costed networks** window appears.
6. Click **Enabled**, and then click **OK**.

See Also

Folder Redirection, Offline Files, and Roaming User Profiles overview

Enable the Always Offline Mode to Provide Faster Access to Files

Deploy Folder Redirection with Offline Files

Disable Offline Files on Individual Redirected Folders

This topic describes how to disable Offline Files caching on individual folders that are redirected to network shares by using Folder Redirection. This provides the ability to specify which folders to exclude from caching locally, reducing the Offline Files cache size and time required to synchronize Offline Files.

In this document

- Prerequisites
- Disabling Offline Files on individual redirected folders

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

To disable Offline Files caching of specific redirected folders, your environment must meet the following prerequisites.

- An Active Directory Domain Services (AD DS) domain, with client computers joined to the domain. There are no forest or domain functional-level requirements or schema requirements.
- Client computers running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012.
- A computer with Group Policy Management installed.

Disabling Offline Files on individual redirected folders

To disable Offline Files caching of specific redirected folders, use Group Policy to enable the **Do not automatically make specific redirected folders available offline** policy setting for the appropriate Group Policy Object (GPO). Configuring this policy setting to **Disabled** or **Not Configured** makes all redirected folders available offline.

Note

Only domain administrators, enterprise administrators, and members of the Group Policy creator owners group can create GPOs.

Do this step using Windows PowerShell

To disable Offline Files on specific redirected folders

1. Open **Group Policy Management**.
2. To optionally create a new GPO that specifies which users should have redirected folders excluded from being made available offline, right-click the appropriate domain or organizational unit (OU) and then click **Create a GPO in this domain, and Link it here**.
3. In the console tree, right-click the GPO for which you want to configure the Folder Redirection settings and then click **Edit**. The Group Policy Management Editor appears.
4. In the console tree, under **User Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and expand **Folder Redirection**.
5. Right-click **Do not automatically make specific redirected folders available offline** and then click **Edit**. The **Do not automatically make specific redirected folders available offline** window appears.
6. Click **Enabled**. In the **Options** pane select the folders that should not be made available offline by selecting the appropriate check boxes. Click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This example creates a new GPO named *Offline Files Settings* in the *MyOu* organizational unit in the *contoso.com* domain (the LDAP distinguished name is `"ou=MyOU,dc=contoso,dc=com"`). It then disables Offline Files for the Videos redirected folder.

```
New-GPO -Name "Offline Files Settings" | New-Gplink -Target "ou=MyOu,dc=contoso,dc=com" -
LinkEnabled Yes


Set-GPRegistryValue -Name "Offline Files Settings" -Key
"HKCU\Software\Policies\Microsoft\Windows\NetCache\{18989B1D-99B5-455B-841C-
AB7C74E4DDFC}" -ValueName DisableFRAdminPinByFolder -Type DWORD -Value 1
```

See the following table for a listing of registry key names (folder GUIDs) to use for each redirected folder.

Redirected folder	Registry key name (folder GUID)
AppData(Roaming)	{3EB685DB-65F9-4CF6-A03A-E3EF65729F3D}
Desktop	{B4BFCC3A-DB2C-424C-B029-7FE99A87C641}

Redirected folder	Registry key name (folder GUID)
Start Menu	{625B53C3-AB48-4EC1-BA1F-A1EF4146FC19}
Documents	{FDD39AD0-238F-46AF-ADB4-6C85480369C7}
Pictures	{33E28130-4E1E-4676-835A-98395C3BC3BB}
Music	{4BD8D571-6D19-48D3-BE97-422220080E43}
Videos	{18989B1D-99B5-455B-841C-AB7C74E4DDFC}
Favorites	{1777F761-68AD-4D8A-87BD-30B759FA33DD}
Contacts	{56784854-C6CB-462b-8169-88E350ACB882}
Downloads	{374DE290-123F-4565-9164-39C4925E467B}
Links	{BFB9D5E0-C6A9-404C-B2B2-AE6DB6AF4968}
Searches	{7D1D3A04-DEBB-4115-95CF-2F29DA2920DA}
Saved Games	{4C5C32FF-BB9D-43B0-B5B4-2D72E54EAAA4}

See Also

Folder Redirection, Offline Files, and Roaming User Profiles overview

Deploy Folder Redirection with Offline Files

Troubleshoot User Profiles with Events

This topic discusses how to troubleshoot problems loading and unloading user profiles by using events and trace logs. The following sections describe how to use the three event logs that record information about user profiles.

In this document

- Step 1: Checking events in the Application log
- Step 2: Viewing the Operational log for the User Profile Service
- Step 3: Enabling and viewing analytic and debug logs

- Step 4: Creating and decoding a trace

Step 1: Checking events in the Application log

The first step in troubleshooting issues loading and unloading user profiles (including roaming user profiles) is to use Event Viewer to examine any Warning and Error events that User Profile Service records in the Application log.

To view User Profile Services events in the Application log

1. Start Event Viewer. To do so, open **Control Panel**, click **System and Security**, and then, in the **Administrative Tools** section, click **View event logs**. The Event Viewer window opens.
2. In the console tree, navigate to **Windows Logs**, and then **Application**.
3. In the Actions pane, click **Filter Current Log**. The Filter Current Log dialog box opens.
4. In the **Event sources** box, select the **User Profiles Service** checkbox, and then click **OK**.
5. Review the listing of events, paying particular attention to Error events.
6. When you find noteworthy events, click the Event Log Online Help link to display additional information and troubleshooting procedures.
7. To perform further troubleshooting, note the date and time of noteworthy events and then examine the Operational log (as described in Step 2) to view details about what the User Profile Service was doing around the time of the Error or Warning events.

Note

You can safely ignore User Profile Service event 1530 "Windows detected your registry file is still in use by other applications or services."

Step 2: Viewing the Operational log for the User Profile Service

If you cannot resolve the issue using the Application log alone, use the following procedure to view User Profile Service events in the Operational log. This log shows some of the inner workings of the service, and can help pinpoint where in the profile load or unload process the problem is occurring.

Both the Windows Application log and the User Profile Service Operational log are enabled by default in all Windows installations.

To view the Operational log for the User Profile Service

1. In the Event Viewer console tree, navigate to **Applications and Services Logs**, then **Microsoft**, then **Windows**, then **User Profile Service**, and then **Operational**.
2. Examine the events that occurred around the time of the Error or Warning events that

you noted in the Application log.

3. If you need further details, use Step 3 to enable and view analytic and debug logs.

Step 3: Enabling and viewing analytic and debug logs

If you require more detail than the Operational log provides, you can enable analytic and debug logs on the affected computer. This level of logging is much more detailed and should be disabled except when trying to troubleshoot an issue.

To enable and view analytic and debug logs

1. In the **Actions** pane of Event Viewer, click **View**, and then click **Show Analytic and Debug Logs**.
2. Navigate to **Applications and Services Logs**, then **Microsoft**, then **Windows**, then **User Profile Service**, and then **Diagnostic**.
3. Click **Enable Log** and then click **Yes**. This enables the Diagnostic log, which will start logging
4. If you require even more detailed information, see Step 3 for information about how to create a trace log.
5. When you are finished troubleshooting the issue, navigate to the **Diagnostic** log, click **Disable Log**, click **View** and then clear the **Show Analytic and Debug Logs** checkbox to hide analytic and debug logging.

Step 4: Creating and decoding a trace

If you cannot resolve the issue using events, you can create a trace log (an ETL file) while reproducing the issue, and then decode it using public symbols from the Microsoft symbol server. Trace logs provide very specific information about what the User Profile Service is doing, and can help pinpoint where the failure occurred.

The best strategy when using ETL tracing is to first capture the smallest log possible. Once the log is decoded, search the log for failures.

To create and decode a trace for the User Profile Service

1. Sign on to the computer where the user is experiencing problems, using an account that is a member of the local Administrators group.
2. From an elevated command prompt enter the following commands, where *<Path>* is the path to a local folder that you have previously created, for example `C:\logs`:

```
logman create trace -n RUP -o <Path>\RUP.etl -ets
logman update RUP -p {eb7428f5-ab1f-4322-a4cc-1f1a9b2c5e98}
0x7FFFFFFF 0x7 -ets
```

3. From the Start screen, click the user name, and then click **Switch account**, being careful not to log off the administrator. If you are using Remote Desktop, close the Administrator session to establish the user session.
4. Reproduce the problem. The procedure to reproduce the problem is typically to sign on as the user experiencing the issue, sign the user off, or both.
5. After reproducing the problem, sign on as the local administrator again.
6. From an elevated command prompt run the following command to save the log into an ETL file:

```
logman stop -n RUP -ets
```

7. Type the following command to export the ETL file into a human-readable file in the current directory (likely your home folder or the %WINDIR%\System32 folder):

```
Tracerpt <path>\RUP.etl
```

8. Open the Summary.txt file and Dumpfile.xml file (which you can open in Microsoft Excel to more easily view the complete details of the log). Look for events that include **fail** or **failed**; you can safely ignore lines that include the **Unknown** event name.

See Also

Deploy Roaming User Profiles

iSCSI Target Block Storage Overview

In Windows Server 2012, iSCSI Target Server is available as a role service under the File and Storage Services role in Server Manager. If you are looking for ways to automate management tasks, a complete set of Windows PowerShell cmdlets are also available.

iSCSI Target Server is ideal for the following:

1. **Network and diskless boot** By using boot-capable network adapters or a software loader, you can deploy hundreds of diskless servers. With iSCSI Target Server, the deployment is fast. In Microsoft internal testing, 256 computers deployed in 34 minutes. By using differencing virtual hard disks, you can save up to 90% of the storage space that was used for operating system images. This is ideal for large deployments of identical operating system images, such as on virtual machines running Hyper-V or in high-performance computing (HPC) clusters.
2. **Server application storage** Some applications require block storage. iSCSI Target Server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it can also consolidate block storage for central or branch office locations.
3. **Heterogeneous storage** iSCSI Target Server supports non-Microsoft iSCSI initiators, making it easy to share storage on servers in a mixed software environment.
4. **Development, test, demonstration, and lab environments** When iSCSI Target Server is enabled, a computer running the Windows Server operating system becomes a network-

accessible block storage device. This is useful for testing applications prior to deployment in a storage area network (SAN).

Block storage requirements

Enabling iSCSI Target Server to provide block storage leverages your existing Ethernet network. No additional hardware is needed. If high availability is an important criterion, consider setting up a high-availability cluster. You need shared storage for a high-availability cluster—either hardware for Fibre Channel storage or a serial attached SCSI (SAS) storage array.

If you enable guest clustering, you need to provide block storage. Any servers running Windows Server software with iSCSI Target Server can provide block storage.

See Also

iSCSI Target Block Storage, How To

What's New in iSCSI Target Server in Windows Server 2012 R2

What's New in iSCSI Target Server in Windows Server 2012 R2

This topic describes the new and changed functionality of iSCSI Target Server in Windows Server 2012 R2.

Feature description

iSCSI Target Server enables you to network boot multiple computers from a single operating system image that is stored in a centralized location. This improves efficiency, manageability, availability, and security. iSCSI Target Server can boot hundreds of computers by using a single operating system image. Other scenarios in which you can use iSCSI Target Server include:

- Some server applications require block storage, and iSCSI Target Server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it can also consolidate block storage for central and branch office locations.
- iSCSI Target Server supports non-Microsoft iSCSI initiators, making it easy to share storage on servers in a mixed, heterogeneous software environment.
- When iSCSI Target Server is enabled, a computer running the Windows Server operating system becomes a network-accessible block storage device. This is useful for testing applications prior to deployment in a storage area network (SAN).

New and changed functionality

The following table lists functionality in iSCSI Target Server that is new or has been changed for this release.

Feature/functionality	New/updated	Summary
Virtual disks enhancements	New	Includes a redesigned data persistence layer that is based on a new version of the VHD format called VHDX (VHD 2.0). VHDX has a much larger storage capacity than the older VHD format. iSCSI Target Server also provides data corruption protection during power failures and optimizes structural alignments of dynamic and differencing disks to prevent performance degradation on new, large-sector physical disks. You can still import VHD 1.0 disks by using Windows Server 2012 R2.
Manageability enhancements	Updated	Uses the SMI-S provider in Windows Server 2012 R2 with System Center Virtual Machine Manager (VMM) to manage iSCSI Target Server in a hosted, and/or private cloud. Additionally, new Windows PowerShell cmdlets for iSCSI Target Server enable you to export and import configuration files, and they provide the ability to disable remote management when iSCSI Target Server is deployed in a dedicated Windows-based appliance scenario (for example, Windows Storage Server).

Feature/functionality	New/updated	Summary
Improved optimization to allow disk-level caching	Updated	iSCSI Target Server now sets the disk cache bypass flag on a hosting disk I/O, through Force Unit Access (FUA), only when the issuing initiator explicitly requests it. This change can potentially improve performance. Previously, iSCSI Target Server would always set the disk cache bypass flag on all I/O's. System cache bypass functionality remains unchanged in iSCSI Target Server; for instance, the file system cache on the target server is always bypassed.
Scalability limits	Updated	Increases the maximum number of sessions per target server to 544, and increases the maximum number of logical units per target server to 256.
Local mount functionality	Updated	Deprecates local mount functionality for snapshots. As a workaround, you can use the local iSCSI initiator on the target server computer (this is also called the loopback initiator) to access the exported snapshots.

See also

The following list provides additional resources about iSCSI Target Server and related technologies.

- iSCSI Target Block Storage Overview
- iSCSI Target Boot Overview
- File and Storage Services

iSCSI High-Availability Block Storage Technical Preview

In a virtualized environment, building a host cluster with Windows Server 2012 can limit downtime due to hardware failures. However, this scenario does not help if you need to take down the virtual machine for maintenance. If downtime of any kind is not an option, you should consider creating a virtual machine guest cluster on top of your host cluster—a feature that has become more relevant in Windows Server 2012. This enables you to service the workload by taking one virtual machine node down for maintenance while the others are still running.

In a host cluster only scenario, the cluster service runs inside the parent partition of the Hyper-V-enabled physical computer. The cluster manages the virtual machines, which reduces downtime when you need to make hardware changes or software updates to the parent partition. The cluster also provides the ability to migrate virtual machines to other servers to load balance.

Guest clustering runs inside a Hyper-V guest computer, and it supports high availability for workloads within the virtual machine. If you add guest clustering to the previous scenario, you can manage cluster-aware applications and services by moving applications between virtual machines within the cluster, without downtime. You can use this process for operating system, service, or application updates. In addition, shared block storage is directly presented to the virtual machine by using the Microsoft iSCSI Software Initiator.

Requirements

Enabling Microsoft iSCSI Software Target to provide block storage leverages your existing Ethernet network. No additional hardware is needed. If high availability is an important criteria, consider setting up a high availability server. With a high availability server, you will need shared storage for the cluster—either hardware Fibre Channel storage or a serial attached SCSI (SAS) storage array.

If you enable guest clustering, you need to provide block storage. Any servers running Windows Server software with Microsoft iSCSI Software Target can provide block storage.

Technical overview

Microsoft iSCSI Software Target is a feature under the File Server role in Windows Server 2012, and it is in full compliance with the iSCSI protocol that is outlined by the Internet Engineering Task Force. The Microsoft iSCSI Initiator provides a way to enable an iSCSI initiator in virtual machines so that they can connect to target servers.

Because iSCSI is an industrial standard, Microsoft iSCSI Software Target interoperates with non-Microsoft implementations of an iSCSI initiator. In a heterogeneous environment with other operating systems, Microsoft iSCSI Software Target can also provide storage access, giving you more choices to access storage.

Windows Server 2012 includes a change in Microsoft iSCSI Software Target when it is used in a clustering configuration. This change improves scalability so that more initiators can connect to the target servers. This supports the goal to provide continuous availability and support workloads with less than 60 seconds of brown time.

iSCSI Target Block Storage, How To

With iSCSI Target Server, any computer running Windows Server 2012 can become a storage device.

- **Prerequisites**
- Step 1: Install iSCSI Target Server and failover clustering
- Step 2: Configure the iSCSI initiator
- Step 3: Create the high availability iSCSI Target Server role service
- Step 4: Configure iSCSI storage provider identity credentials

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

Before you configure block storage, you need to set up two or more servers that are running Windows Server 2012.

Step 1: Install iSCSI Target Server and failover clustering

Install the required role service and features by using Windows PowerShell commands:

1. Start Windows PowerShell if it is not already open. From the Start menu, click **Administrative Tools**, and then click **Windows PowerShell Modules**. The Windows PowerShell Modules command-line console appears. The first time you open the tool and click ENTER, Windows PowerShell imports all the cmdlets.
2. Add the iSCSI Target Server role service by running the following command:

```
PS C:\> Add-WindowsFeature fs-iscsitarget-server
```

3. Verify that the iSCSI Target Server role service is installed by running the following command:

```
PS C:\> Get-WindowsFeature fs-iscsitarget-server
```

Note
"[X]" indicates that the role service is installed.

4. Install failover clustering by running the following command on the cluster node:

```
PS C:\> Add-WindowsFeature failover-clustering
```

5. Verify that the feature is installed on both nodes by running the following command:

```
PS C:\> Get-WindowsFeatures failover-clustering
```

Note
"[X]" indicates that the feature is installed.

Step 2: Configure the iSCSI initiator

1. On the computer that you set up as the iSCSI initiator, start the iSCSI Initiator Properties. From Server Manager, click the **Tools** menu, and then choose **iSCSI Initiator**.
2. Click the **Targets** tab, and then type the iSCSI target IP address.
3. Click **Quick Connect**, and then select the **Connect** button.

After a connection has been established, you can use Windows PowerShell cmdlets to bring the disks online.

To use PowerShell cmdlets to bring the disks online

1. From Windows PowerShell, type:

```
New-IscsiTargetPortal –TargetPortalAddress <IP_address>
$t = Get-IscsiDiscoveredTarget
Connect-IscsiDiscoveredTarget –NodeAddress $t.NodeAddress
Get-IscsiPersistentTarget
```

Note
If there is no record returned, you can register the iSCSI target by typing:

```
Register-IscsiPersistentTarget –TargetName $t.NodeAddress
```

2. Bring the disk online and create a volume on it:

```
Get-disk
```

Note
To manage iSCSI Target Server on a cluster, you will need to use Server Manager or Windows PowerShell cmdlets.

Step 3: Create the high availability iSCSI Target Server role service

To ensure transparent failover, create the iSCSI Target Server role service on a cluster.

To create the iSCSI Target Server role service on a cluster

1. Open Failover Cluster Manager.

2. Right-click **Roles**, and then click **Configure Role**. The High Availability Wizard, Select Service or Application page appears.
3. Select **iSCSI Target Server**. The **Client Access Point** page appears.
4. Type the service or application name and the IP address for the iSCSI Target Server high-availability instance.
5. Select the **Networks Address** check box, and then click **Next**. The **Select Storage** page appears.

   Note

   Each time that you assign storage to the cluster nodes, by default, the disks are offline. You need to manually bring them online. To bring the disks online, leave the High Availability Wizard open (you will return to it shortly) and open the Disk Management tool.
6. In the Disk Management tool, right-click the disks to bring them online.
7. After bringing the disks online, you need to put them in the storage. Open the Failover Cluster Manager. From the Actions pane, click **Storage**, and then click **Add Disk**.
8. Return to the High Availability Wizard, select the check box for this storage volume, and then click **Next**. The **Confirmation** page appears.
9. Return to the Failover Cluster Manager and verify that the storage volume was created successfully.

Note

To manage iSCSI Target Server on a cluster, you need to use Server Manager or Windows PowerShell cmdlets.

Step 4: Configure iSCSI storage provider identity credentials

The iSCSI VDS and VSS storage providers must be configured to run under the administrative credentials of the iSCSI Target Server. You can configure these credentials by using Windows PowerShell or by using Component Services (DcomCnfg). For servers that are not part of a domain, create a local user account with the same user name and password for each server.

To use Windows PowerShell to configure credentials

1. Open Windows PowerShell.
2. Run the following command:

```
$PsCred = Get-Credential
$PrvdIdentityPath = New-Object
System.Management.ManagementPath("root\wmi:WT_iSCSIStoragePro
viderIdentity")
$PrvdIdentityClass = New-Object
```

```
System.Management.ManagementClass($PrvdIdentityPath)
$PrvdIdentityClass.SetProviderIdentity("{88155B26-CE61-42FB-AF31-
E024897ADEBF}",$PsCred.UserName,$PsCred.GetNetworkCredential(
).Password)
$PrvdIdentityClass.SetProviderIdentity("{9D884A48-0FB0-4833-AB70-
A19405D58616}",$PsCred.UserName,$PsCred.GetNetworkCredential(
).Password)
```

To use Component Services to configure credentials

1. Open **Component Services**.
2. In the left pane, click **Computers**, click **My Computer**, and then click **DCOM Config**..
3. Right-click **WTVdsProv**, click **Properties**, click the **Identity** tab, click **This user**, enter the user name and password for the iSCSI Target Server account, and then click **OK**.
4. Right-click **WTSnapsthoProvider**, click **Properties**, click the **Identity** tab, click **This user**, enter the user name and password of the iSCSI Target Server account, and then click **OK**.
5. Click **Run**, and then type **regedit**. In the left pane, click **HKey_Local_Machine**, click **Software**, click **Microsoft**, click **Windows NT**, click **CurrentVersion**, click **WTVDSProv**, and then click **WtServers**. . Right-click **WtServers**, point to **New**, click **String Value**, and then type the name of the iSCSI Target Server. (The **Data** value can remain blank.)

See also

- iSCSI Target Block Storage Overview

iSCSI Target Implementation Notes

iSCSI Target Server was implemented in compliance with the industry standards that are specified by RFC 3720 and RFC 5048. To verify this compliance, iSCSI Target Server was tested by a compliance testing organization.

This document lists the issues that were found during testing of iSCSI Target Server on Windows Server 2012. For the previous version – Microsoft iSCSI Software Target version 3.3 - known issues are listed at Microsoft iSCSI Software Target 3.3 Implementation Notes (http://technet.microsoft.com/en-us/library/gg983494(v=ws.10)).

Known Issues

The following are the known issues with Windows® Server 2012 iSCSI Target Server.

SCSI Response Excess Immediate Data

RFC 3720, section 3.2.4.2, specifies the following:

It is also an error for an initiator to send more unsolicited data, whether immediate or as separate PDUs, than FirstBurstLength.

The behavior exhibited by ISCSI Target Server:

When responding to a **Write** command with more immediate data than specified by the FirstBurstLength, iSCSI Target Server sends a SCSI Response PDU with a status of **Good**.

SCSI Response to Incorrect Amount of Data Condition

RFC 3720, section 10.4.7.2, specifies the following:

The target reports the 'Incorrect amount of data' condition if during data output the total data length to output is greater than FirstBurstLength and the initiator sent unsolicited non-immediate data but the total amount of unsolicited data is different than FirstBurstLength. The target reports the same error when the amount of data sent as a reply to an R2T does not match the amount requested.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server responds with status of 0x00 (good) after receiving more than 512 bytes of data when FirstBurstLength=512.

The iSCSI Target Server responds with status of 0x00 (good) after receiving data less than specified ExpectedDataTranserLength and F bit set to 1.

SCSI Response to Unexpected Unsolicited Data

RFC 3720, section 12.11, specifies the following:

If ImmediateData is set to No and InitialR2T is set to Yes, then the initiator MUST NOT send unsolicited data and the target MUST reject unsolicited data with the corresponding response code.

The behavior exhibited by ISCSI Target Server:

When the ImmediateData is set to No, and InitialR2T is set to Yes, iSCSI Target Server was observed to respond with the status Good when it received a Write command with unexpected unsolicited data.

Handling Text Response Text Field in Discovery Session

RFC 3720, section 10.11.1, specifies the following:

A Text Response with the F bit set to 1 in response to a Text Request with the F bit set to 0 is a protocol error.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to send a Text Response with:TargetTransferTag = 0xFFFFFFFF while the F bit was not set.

Handling Text Request During Normal Session

RFC 3720, section 3.5.3.1, specifies the behavior of Text Request and Text Response as follows:

Text requests and responses are designed as a parameter negotiation vehicle and as a vehicle for future extension.

In the data segment, Text Requests/Responses carry text information using a simple "key=value" syntax.

Text Request/Responses may form extended sequences using the same Initiator Task Tag. The initiator uses the F (Final) flag bit in the text request header to indicate its readiness to terminate a sequence. The target uses the F (Final) flag bit in the text response header to indicate its consent to sequence termination.

Text Request and Responses also use the Target Transfer Tag to indicate continuation of an operation or a new beginning. A target that wishes to continue an operation will set the Target Transfer Tag in a Text Response to a value different from the default 0xffffffff.

An initiator willing to continue will copy this value into the Target Transfer Tag of the next Text Request. If the initiator wants to restart the current target negotiation (start fresh) will set the Target Transfer Tag to 0xffffffff.

Although a complete exchange is always started by the initiator, specific parameter negotiations may be initiated by the initiator or target.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to disconnect the session after it received a Text Request during a normal session, when the following conditions existed:

F bit was set to 0

There was an attached MaxRecvDataSegmentLengthKey

This happens with or without a valid ITT (Initiator Task tag).

Handling Text Request During Discovery Session

RFC 3720, section 3.3, specifies the following:

Discovery-session - a session only opened for target discovery. The target MUST ONLY accept text requests with the SendTargets key and a logout request with the reason "close the session". All other requests MUST be rejected.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to disconnect after receiving a text request with an attached MaxRecvDataSegmentLength key during a Discovery session.

Handling Text Request with SendTargets=

RFC 3720, Appendix D, specifies the following:

If the SendTargets value is nothing in the Text Request, “The session should only respond with addresses for the target to which the session is logged in. This MUST be supported on operational sessions, and MUST NOT return targets other than the one to which the session is logged in.”

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to disconnect after receiving a text request with an attached “SendTargets=” key during a discovery session and during a normal session.

Handling Text Request with SendTargets=All during Normal Session

RFC 3720, Appendix D, specifies the following:

If the SendTargets value is “All” in the Text Request, “The initiator is requesting that information on all relevant targets known to the implementation be returned. This value MUST be supported on a discovery session, and MUST NOT be supported on an operational session.”

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to disconnect after receiving a text request with an attached “SendTargets=All” key during a normal session.

Text Response Negotiation Failure

RFC 3720, Section 6.10, specifies the following:

A negotiation failure is considered to be one or more of the following:

None of the choices, or the stated value, is acceptable to one of the sides in the negotiation.

The text request timed out and possibly terminated.

The text request was answered with a Reject PDU.

The behavior implemented by the ISCSI Software Target in Windows Server 2012:

The iSCSI Target Server was observed to disconnect the session after receiving a text request with the F bit set to 0, and an attached MaxRecvDataSegmentLengthKey.

Handling Text Request with C bit

RFC 3720, Section 5.2, specifies the following:

A target receiving a Text or Login Request with the C bit set to 1 MUST answer with a Text or Login Response with no data segment (DataSegmentLength 0).

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to disconnect after receiving a text request with C bit set to 1.

Standard Login Key Negotiation

RFC 3720, section 10.13.3, specifies the following:

For a new session, the target MUST generate a non-zero TSIH and ONLY return it in the Login Final-Response.

The behavior exhibited by ISCSI Target Server:

When the iSCSI initiator performed a standard login and negotiated the login parameters, the iSCSI Target Server was observed to set the TSIH field in the first Login Response PDU.

Duplicate or Out of Order CmdSN

RFC 3720, section 3.2.2.1 specifies the following:

The target MUST silently ignore any non-immediate command outside of this range or non-immediate duplicates within the range.

For non-immediate commands, the CmdSN field can take any value from ExpCmdSN to MaxCmdSN inclusive.The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server does not send an R2T in response to a WRITE command which was followed by a duplicate WRITE. It silently ignores any PDU which CmdSN is not the next sequential CmdSN.

Note
Since the iSCSI Target Server supports one connection per session, it currently silently ignores any PDU which CmdSN is not the next sequential CmdSN.

Task Management Response for Non-Existent Task

RFC 3720, section 10.6.1, specifies the following:

If the Referenced Task Tag does not identify an existing task and if the CmdSN indicated by the RefCmdSN field in the Task Management function request is outside the valid CmdSN window, then targets must return the "Task does not exist" response.

The target provides a Response, which may take on the following values:

a) 0 - Function complete.

b) 1 - Task does not exist.

c) 2 - LUN does not exist.

d) 3 - Task still allegiant.

e) 4 - Task allegiance reassignment not supported.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to respond with 0 (function complete) when receiving a non-existent RefrencedTaskTag and a RefCmdSN outside the valid CmdSN window.

Incorrect Logout Reason Code Being Accepted During Discovery

RFC 3720, section 3.3, specifies the following:

Discovery-session - a session only opened for target discovery. The target MUST ONLY accept text requests with the SendTargets key and a logout request with the reason "close the session". All other requests MUST be rejected.

The behavior exhibited by ISCSI Target Server:

During the discovery session, the iSCSI Target Server sends a response code of 0x00 (success) in response to a logout request with reason code 2 (remove connection with recovery).

Logout Response Non-existent Connection

RFC 3720, section 10.14, specifies the following:

When receiving a Logout Request with the reason code of "close the connection" or "close the session", the target MUST terminate all pending commands, whether acknowledged via ExpCmdSN or not, on that connection or session respectively.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server sends a response code of 0x00 (success) in response to a logout request with reason code 1 (closes the connection) for a non-existent CID.

Note
The iSCSI Target Server supports one session per connection and currently ignores CID value.

Handling Not Understood for Required Keys

RFC 3720, section 5.2, specifies the following:

If the acceptor sends "Reject" as an answer the negotiated key is left at its current value (or default if no value was set). If the current value is not acceptable to the proposer on the connection or to the session it is sent, the proposer MAY choose to terminate the connection or session.

All keys in this document, except for X extension formats, MUST be supported by iSCSI initiators and targets when used as specified here. If used as specified, these keys MUST NOT be answered with NotUnderstood.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to send a login success message when it received a request with the TargetPortalGroupTag=NotUnderstood key.

Handling Re-negotiation During Login

RFC 3720, section 5.3, specifies the following:

Neither the initiator nor the target should attempt to declare or negotiate a parameter more than once during login except for responses to specific keys that explicitly allow repeated key declarations (e.g., TargetAddress). An attempt to renegotiate or redeclare parameters not specifically allowed MUST be detected by the initiator and target. If such an attempt is detected by the target, the target MUST respond with Login Reject (initiator error).

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to allow renegotiation of the following parameters during the login phase:

ImmediateData key=value pair to yes

MaxBurstLength key=value pair to 32524

DataDigest key=value pair to CRC32C

Double negotiation of DataDigest key=value pair in the same PDU

When it receives the previous parameters twice in the negotiation phase, the iSCSI Target Server proceeds with the login request.

Reject SNACK

RFC 3720, section 10.17.3, specifies the following:

StatSN, ExpCmdSN and MaxCmdSN carry their usual values and are not related to the rejected command. StatSN is advanced after a Reject.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to send an incorrect StatSN when rejecting SNACK request.

Text mode negotiation

RFC 3720, section 5.2, specifies the following:

Any key not understood by the acceptor may be ignored by the acceptor without affecting the basic function. However, the answer for a key not understood MUST be key=NotUnderstood.

The behavior exhibited by ISCSI Target Server:

The iSCSI Target Server was observed to ignore the vendor specific keys and continue with the Login phase when C bit is set. If C bit is not set, iSCSI Target Server will comply with the standard.

iSCSI Target Boot Overview

iSCSI Target Server in Windows Server 2012 can boot hundreds of computers from a single operating system image that is stored in a centralized location. This improves efficiency, manageability, availability, and security.

Feature description

By using differencing virtual hard disks (VHDs), you can use a single operating system image (the “master image”) to boot up to 256 computers. As an example, let’s assume that you deployed Windows Server 2008 R2 HPC with an operating system image of approximately 20 GB, and you used two mirrored disk drives to act as the boot volume. You would have needed approximately 10 TB of storage for only the operating system image to boot 256 computers. With iSCSI Target Server, you will use 40 GB for the operating system base image, and 2 GB for differencing virtual hard disks per server instance, totaling 552 GB for the operating system images. This provides a savings of over 90% on storage for the operating system images alone.

Practical applications

Using a controlled operating system image provides the following benefits:

More secure and easier to manage. Some enterprises require that data be secured by physically locking storage in a centralized location. In this scenario, servers access the data remotely, including the operating system image. With iSCSI Target Server, administrators can centrally manage the operating system boot images, and control which applications to use for the master image.

Rapid deployment. Because the master image is prepared by using Sysprep, when computers boot from the master image, they skip the file copying and installation phase that occurs during Windows Setup, and they go straight to the customization phase. In Microsoft internal testing, 256 computers were deployed in 34 minutes.

Fast recovery. Because the operating system images are hosted on the computer running iSCSI Target Server, if a diskless client needs to be replaced, the new computer can point to the operating system image, and boot up immediately.

Note

Various vendors provide a storage area network (SAN) boot solution, which can be used by the iSCSI Target Server in Windows Server 2012 on commodity hardware.

Hardware requirements

iSCSI Target Server does not require special hardware for functional verification. In data centers with large-scale deployments, the design should be validated against specific hardware. For reference, Microsoft internal testing indicated that a 256 computer deployment required 24x15k-RPM disks in a RAID 10 configuration for storage. A network bandwidth of 10 GB is optimal. A general estimate is 60 iSCSI boot servers per 1 GB network adapter.

A network adapter is not required for this scenario, and a software boot loader can be used (such as iPXE open source boot firmware).

Software requirements

iSCSI Target Server can be installed as part of the File and iSCSI Services role service in Server Manager.

For automation, you can use the corresponding Windows PowerShell cmdlets.

Datacenter Diskless Boot Technical Preview

The Microsoft iSCSI Software Target feature is built into Windows Server 2012, and it lets you create a Storage Area Network (SAN) device on hardware running the Windows operating system. iSCSI Software Target also enables you to network boot multiple computers from a single operating system image that is stored in a centralized location. This improves efficiency, manageability, availability, and security. iSCSI Software Target in Windows Server 2012 can boot hundreds of computers by using a single operating system image and it provides the following benefits:

Cost savings on operating system storage. By using differencing virtual disks, you can use a single operating system image (the “golden image”) to boot up to 256 computers. As an example, in a deployment of Windows Server 2008 R2 HPC edition, the operating system image is approximately 20 GB. A common deployment is to have two mirrored disk drives that act as the boot volume. Rounding the operating system storage to 40 GB per instance, you would need approximately 10 TB of storage — for only the operating system image—to boot 256 computers. With iSCSI Software Target boot, however, you will use 40 GB for the operating system base image, and 2 GB for differencing virtual hard disks (VHDs) per server instance, totaling 552 GB for the operating system images. This provides a savings of over 90% on storage for the operating system images alone.

Controlled operating system images make it more secure and easy to manage. Some enterprises require that data be secured by physically locking storage in a centralized location. In this scenario, servers access the data remotely, including the operating system image. With iSCSI Software Target boot, administrators can centrally manage the operating system boot images, and control which applications to put in the golden image.

Rapid deployment. Because the golden image is a sysprepped operating system image, when the computers boot from the golden image, they skip the file copying and installation phase that occurs during Windows Setup, and they go straight to the customization phase. In our testing, we deployed 256 computers in 34 minutes.

Fast recovery. Because the operating system images are hosted on the iSCSI Software Target server, if the diskless client needs to be replaced, the new computer can point to the operating system image, and boot up immediately.

A SAN boot is a solution that has been offered from various vendors. In Windows Server 2012, the new iSCSI Software Target feature provides this network boot capability on commodity hardware.

Requirements

iSCSI Software Target can be installed as part of the File Server role.

The easy management experience is provided through Server Manager. For automation, you can use the corresponding Windows PowerShell cmdlets.

This feature does not require special hardware for functional verification. In data centers with large-scale deployments, the design should be validated against specific hardware. For reference, Microsoft internal testing indicated that for a 256-iSCSI boot deployment, 24x15k-RPM (revolution per minute) disks in a RAID 10 configuration were required for storage. A network bandwidth of 10 GB is optimal. A general estimate is 60 iSCSI boot servers per 1 GB network adapter. However, an iSCSI boot-capable network adapter is not required for this scenario. If the network adapter does not support it, a software boot loader can be used (such as iPXE open source boot firmware).

Technical overview

The iSCSI Software Target feature in Windows Server 2012 supports diskless network boot without the need for special hardware or additional software. iSCSI Software Target fully complies with the iSCSI protocol as specified by the Internet Engineering Task Force. A key component to this feature is that it supports differencing of virtual hard disks. This is critical in the boot scenario because multiple servers running Windows Server 2012 can boot by using only one base image.

iSCSI Software Target accomplishes diskless boot as follows:

- iSCSI Software Target supports creating differencing virtual disks based on a golden image.
- Each diskless client boots from its own differencing virtual disk.
- The diskless client reads from the golden image, and writes to its own differencing VHD.

Because the image is hosted in a centralized place, storage space savings are realized on operating system deployment. That can translate into freed-up storage space that was previously allocated to operating system boot images or budget savings because you no longer need to purchase hard disks for this purpose and the power to cool them.

In Windows Server 2012, the scale of iSCSI boot has increased to 256. Therefore, the clustered target can support 256 boot clients, and it does not experience operating system errors with iSCSI Software Target service failover from one clustered node to another.

Network File System Overview

This topic describes the Network File System role service and features that are included with the File and Storage Services server role in Windows Server 2012 R2 and Windows Server 2012. Network File System (NFS) provides a file sharing solution for enterprises that have heterogeneous environments that consist of both Windows and non-Windows computers.

Did you mean…

- Deploy Network File System

- File and Storage Services Overview
- Services for NFS Step-by-Step Guide for Windows Server 2008 R2

Feature description

Using the NFS protocol, you can transfer files between computers running Windows and other non-Windows operating systems, such as Linux or UNIX.

In Windows Server 2012, NFS includes the components, Server for NFS and Client for NFS. Server for NFS enables a computer running Windows Server 2012 to act as a NFS file server for other non-Windows client computers. Client for NFS enables a Windows-based computer that is running Windows Server 2012 to access files that are stored on a non-Windows NFS server.

Practical applications

Some practical applications for NFS include:

- Use a Windows NFS file server to provide multi-protocol access to the same file share over both the SMB and NFS protocols from multi-platform clients.
- Deploy a Windows NFS file server in a predominantly non-Windows operating system environment to provide access to NFS file shares for non-Windows client computers.
- Migrate applications from one operating system to another by storing the data on file shares that are accessible by both SMB and NFS protocols.

New and changed functionality

New and changed functionality in Network File System includes support for the NFS version 4.1 and easier deployment and manageability. For information about functionality that is new or changed in Windows Server 2012, review the following table:

Feature/functionality	New or updated	Description
NFS version 4.1	New	Increases security, performance, and interoperability over NFS version 3.
NFS infrastructure	Updated	Provides easier deployment and manageability, and increases security.
NFS version 3 continuous availability	Updated	Improves continuous availability on NFS version 3 clients.
Deployment and manageability	Updated	Enables you to easily deploy and manage NFS using new

Feature/functionality	New or updated	Description
improvements		Windows PowerShell cmdlets and a new WMI provider.

NFS version 4.1

NFS version 4.1 implements all of the required aspects, in addition to some of the optional aspects, of RFC 5661:

- **Pseudo file system.** This file system includes a separation of the physical and logical namespace and is compatible with NFS version 3 and NFS version 2. An alias is provided for the exported file system, which is part of the pseudo file system.
- **Compound RPCs.** Combines relevant operations and reduces chattiness.
- **Sessions and session trunking.** Enables just one semantic and allows continuous availability and better performance while utilizing multiple networks between NFS 4.1 clients and the NFS Server.

NFS infrastructure

Improvements to the overall NFS infrastructure in Windows Server 2012 are detailed below:

- **Remote Procedure Call (RPC)/External Data Representation (XDR).** The RPC/XDR transport infrastructure is powered by the WinSock network protocol and is available for both Server for NFS and Client for NFS. This replaces Transport Device Interface (TDI), offers better support, and provides better scalability and Receive Side Scaling (RSS).
- **RPC port multiplexer.** This feature is firewall-friendly (less ports to manage) and simplifies deployment of NFS.
- **Auto-tuned caches and thread pools.** The resource management capability of the new RPC/XDR infrastructure is dynamic and automatically tunes the cache and thread pools based on the workload. This completely removes the guesswork involved when tuning parameters, and therefore, provides optimal performance as soon as NFS is deployed.
- **New Kerberos privacy implementation and authentication options.** NFS deployments are more secure with the addition of Kerberos privacy (Krb5p) support along with the existing krb5 and krb5i authentication options.
- **Identity Mapping Windows PowerShell module.** These new cmdlets make it easier to manage identity mapping, configure Active Directory Lightweight Directory Services (AD LDS), and set up UNIX and Linux passwd and flat files.
- **Volume mount point.** Volumes mounted under an NFS share can be accessed using NFS version 4.1.
- **Port Multiplexing.** This feature supports the RPC port multiplexer (port 2049), which is firewall-friendly and simplifies deployment of NFS.

NFS version 3 continuous availability

NFS version 3 clients can have fast and transparent failovers with more availability and reduced downtime. The failover process is faster for NFS version 3 clients because:

- The clustering infrastructure now allows one resource per network name instead of one resource per share. This significantly improves the failover time of resources.
- Failover paths within an NFS server are tuned for better performance.
- Wildcard registration in an NFS server is no longer required, and the failovers are more fine-tuned.
- Network Status Monitor (NSM) notifications are sent out after a failover process and clients no longer need to wait for TCP timeouts to reconnect to the server that failed over.

Deployment and manageability improvements

Deploying and managing NFS has improved in the following ways:

- Over forty new Windows PowerShell cmdlets make it easier to configure and manage NFS file shares. For more information, see NFS Cmdlets in Windows PowerShell.
- Identity mapping is improved with a local flat file mapping store and with new Windows PowerShell cmdlets available to configure identity mapping.
- The Server Manager graphical user interface is easier to use.
- The new WMI version 2 provider is available for easier management.
- The RPC port multiplexer (port 2049) is firewall-friendly and simplifies deployment of NFS.

Server Manager information

In Server Manager, use the Add Roles and Features Wizard to add the Server for NFS role service (under the File and iSCSI Services role). For general information about installing features, see Install or Uninstall Roles, Role Services, or Features. Tools for Server for NFS include the Services for Network File System MMC snap-in to manage the Server for NFS and Client for NFS components. Using the snap-in, you can manage the Server for NFS components that are installed on the computer. Server for NFS also contains several Windows command-line administration tools:

- **Mount.** Mounts a remote NFS share (also known as an export) locally and maps it to a local drive letter on the Windows client computer.
- **Nfsadmin.** Manages configuration settings of the Server for NFS and Client for NFS components.
- **Nfsshare.** Configures NFS share settings for folders that are shared using Server for NFS.
- **Nfsstat.** Displays or resets statistics of calls received by Server for NFS.
- **Showmount.** Displays mounted file systems exported by Server for NFS.
- **Umount.** Removes NFS-mounted drives.

NFS in Windows Server 2012 introduces the NFS module for Windows PowerShell with several new cmdlets specifically for NFS. These cmdlets provide an easy way to automate NFS management tasks. For more information, see NFS Cmdlets in Windows PowerShell.

See also

The following table provides additional resources for evaluating NFS.

Content type	References
Deployment	Deploy Network File System
Operations	NFS Cmdlets in Windows PowerShell
Related technologies	File and Storage Services Overview

Deploy Network File System

Network File System (NFS) provides a file sharing solution that enables you to transfer files between computers running Windows Server 2012 and UNIX operating systems using the NFS protocol. This topic describe the steps you should follow to deploy NFS.

In this topic

What's new in Network File System

Scenarios for using Network File System

System requirements

Deploy NFS infrastructure

Install Network File System

Configure NFS authentication

Create an NFS file share

What's new in Network File System

The following improvements are available for NFS in Windows Server 2012:

- **Support for NFS version 4.1**. This protocol version includes the following enhancements.

  Navigating firewalls is easier which improves accessibility.

  Supports the RPCSEC_GSS protocol which provides stronger security, and the ability for clients and servers to negotiate security.

  Supports UNIX and Windows file semantics.

  Takes advantage of clustered file server deployments.

  Supports WAN friendly compound procedures.

- **NFS module for Windows PowerShell**. The availability of built-in NFS cmdlets makes it easier to automate various operations. The cmdlet names are consistent with other Windows PowerShell cmdlets (using verbs such as ‘Get’ and ‘Set’) which makes it easier for users familiar with Windows PowerShell to learn to use new cmdlets.
- **NFS management improvements**. A new centralized UI-based management console simplifies configuration and management of SMB and NFS shares, quotas, file screens and classification, in addition to managing clustered file servers.
- **Identity Mapping improvements**. New UI support and task-based Windows PowerShell cmdlets for configuring identity mapping, which allows administrators to quickly configure an identity mapping source, and then create individual mapped identities for users. Improvements make it easy for administrators to set up a share for multi-protocol access over both NFS and SMB.
- **Cluster resource model restructure**. This improvement brings consistency between the cluster resource model for the Windows NFS and SMB protocol servers and simplifies administration. For NFS servers that have many shares, the resource network and the number of WMI calls required fail over a volume containing a large number of NFS shares are reduced.
- **Integration with Resume Key Manager**. The Resume Key Manager is a component that tracks file server and file system state and enables the Windows SMB and NFS protocol servers to fail over without disrupting clients or server applications that store their data on the file server. This improvement is a key component of the continuous availability capability of the file server running Windows Server 2012.

Scenarios for using Network File System

NFS supports a mixed environment of Windows-based and UNIX-based operating systems. The following deployment scenarios are examples of how you can deploy a continuously available Windows Server 2012 file server using NFS.

Provision file shares in heterogeneous environments

This scenario applies to organizations with heterogeneous environments that consist of both Windows and other operating systems, such as UNIX or Linux-based client computers. With this scenario, you can provide multi-protocol access to the same file share over both the SMB and NFS protocols. Typically, when you deploy a Windows file server in this scenario, you want to facilitate collaboration between users on Windows and UNIX-based computers. When a file share is configured, it is shared with both the SMB and NFS protocols, with Windows users accessing their files over the SMB protocol, and users on UNIX-based computers typically access their files over the NFS protocol.

For this scenario, you must have a valid identity mapping source configuration. Windows Server 2012 supports the following identity mapping stores:

- Mapping File
- Active Directory Domain Services (AD DS)
- RFC 2307 compliant LDAP stores such as Active Directory Lightweight Directory Services (AD LDS)

- User Name Mapping (UNM) server

Provision file shares in UNIX-based environments

In this scenario, Windows file servers are deployed in a predominantly UNIX-based environment to provide access to NFS file shares for UNIX-based client computers. An Unmapped UNIX User Access (UUUA) option was initially implemented for NFS shares in Windows Server 2008 R2 so that Windows servers can be used for storing NFS data without creating UNIX-to-Windows account mapping. UUUA allows administrators to quickly provision and deploy NFS without having to configure account mapping. When enabled for NFS, UUUA creates custom security identifiers (SIDs) to represent unmapped users. Mapped user accounts use standard Windows security identifiers (SIDs) and unmapped users use custom NFS SIDs.

System requirements

Server for NFS can be installed on any version of Windows Server 2012. You can use NFS with UNIX-based computers that are running an NFS server or NFS client if these NFS server and client implementations comply with one of the following protocol specifications:

1. NFS Version 4.1 Protocol Specification (as defined in RFC 5661)
2. NFS Version 3 Protocol Specification (as defined in RFC 1813)
3. NFS Version 2 Protocol Specification (as defined in RFC 1094)

Deploy NFS infrastructure

You need to deploy the following computers and connect them on a local area network (LAN):

- One or more computers running Windows Server 2012 on which you will install the two main Services for NFS components: Server for NFS and Client for NFS. You can install these components on the same computer or on different computers.
- One or more UNIX-based computers that are running NFS server and NFS client software. The UNIX-based computer that is running NFS server hosts an NFS file share or export, which is accessed by a computer that is running Windows Server 2012 as a client using Client for NFS. You can install NFS server and client software either in the same UNIX-based computer or on different UNIX-based computers, as desired.
- A domain controller running at the Windows Server 2008 R2 functional level. The domain controller provides user authentication information and mapping for the Windows environment.
- When a domain controller is not deployed, you can use a Network Information Service (NIS) server to provide user authentication information for the UNIX environment. Or, if you prefer, you can use Password and Group files that are stored on the computer that is running the User Name Mapping service.

Install Network File System

To install Network File System on the server using Server Manager:

1. From the Add Roles and Features Wizard, under Server Roles, select **File and Storage Services** if it has not already been installed.
2. Under **File and iSCSI Services**, select **File Server** and **Server for NFS**. Click **Add Features** to include selected NFS features.
3. Click **Install** to install the NFS components on the server.

To install Network File System on the server using Windows PowerShell:

1. Start Windows PowerShell. Right-click the PowerShell icon on the taskbar, and select **Run as Administrator**.
2. Run the following Windows PowerShell commands:

```
PS C:\> Import-Module ServerManager

PS C:\> Add-WindowsFeature FS-NFS-Services

PS C:\> Import-Module NFS
```

Configure NFS authentication

When using the NFS version 4.1 and NFS version 3.0 protocols, you have the following authentication and security options.

RPCSEC_GSS:

- **Krb5**. Uses the Kerberos version 5 protocol to authenticate users before granting access to the file share.
- **Krb5i**. Uses Kerberos version 5 protocol to authenticate with integrity checking (checksums), which verifies that the data has not been altered.
- **Krb5p** Uses Kerberos version 5 protocol, which authenticates NFS traffic with encryption for privacy.

AUTH_SYS:

You can also choose not to use server authorization (AUTH_SYS), which gives you the option to enable unmapped user access. When using unmapped user access, you can specify to allow unmapped user access by UID / GID, which is the default, or allow anonymous access.

Instructions for configuring NFS authentication on discussed in the following section.

Create an NFS file share

You can create an NFS file share using either Server Manager or Windows PowerShell NFS cmdlets.

To create an NFS file share by using Server Manager

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.

3. On the left, click **File and Storage Services**, and then click **Shares**.
4. Click **To create a file share, start the New Share Wizard**.
5. On the **Select Profile** page, select either **NFS Share – Quick** or **NFS Share - Advanced**, and click **Next**.
6. On the **Share Location** page, select a server and a volume, and click **Next**.
7. On the **Share Name** page, specify a name for the new share, and click **Next**.
8. On the **Authentication** page, specify the authentication method you want to use for this share.
9. On the **Share Permissions** page, click **Add**, and then specify the host, client group or netgroup you want to grant permission to the share.
10. In **Permissions**, configure the type of access control you want the users to have, and click **OK**.
11. On the **Confirmation** page, review your configuration, and click **Create** to create the NFS file share.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

On the server, type the following to configure an NFS file share (where `nfs1` is the name of the share and `C:\shares\nfsfolder` is the path):

```
New-NfsShare -name nfs1 -Path C:\shares\nfsfolder
```

NTFS

This topic provides overview and background information for NTFS file system.

Feature description

NTFS provides performance and reliability, and includes built-in security features, such as file and folder permissions. NTFS also provides support for volumes up to 256 terabytes in size, support for disk quotas and compression, and support for mounted drives. For information on new and changed functionality in NTFS, see What's New in NTFS for Windows Server 2012 R2.

Practical applications

The list below describes some of the practical applications in which NTFS should be used as the file system.

- **Increasing reliability**. NTFS uses its log file and checkpoint information to restore the consistency of the file system when the computer is restarted in the event of a system failure. In the event of a bad-sector error, NTFS dynamically remaps the cluster containing the bad sector and allocates a new cluster for the data, as well as marking the cluster as bad and no longer using it. For example, by formatting a POP3 mail server with NTFS, the mail store can offer logging and recovery. In the event of a server crash, NTFS can recover data by replaying its log files.
- **Increasing security**. NTFS allows you to set permissions on a file or folder, and specify the groups and users whose access you want to restrict or allow, and then select the type of access.
- **Limited space on a volume**. If your organization has limited space on a volume, NTFS provides support for increasing storage on a server with limited disk space.
  - Disk quotas allow you to track and control user disk space usage for NTFS volumes.
  - NTFS supports compression as well as adding unallocated space from the same disk or from another disk to increase the size of an NTFS volume.
  - Mounted volumes allow you to mount a volume at any empty folder on a local NTFS volume if you run out of drive letters or need to create additional space that is accessible from an existing folder.
- **Supporting large volumes**. NTFS allows you to create an NTFS volume up to 16 terabytes using the default cluster size (4 KB) for large volumes. You can create NTFS volumes up to 256 terabytes using the maximum cluster size of 64 KB.

See also

The following list provides additional resources on the web about related technologies in Windows Server 2012.

- NTFS Health and Chkdsk
- File and Storage Services Overview
- Resilient File System Overview
-

What's New in NTFS for Windows Server 2012 R2

This topic introduces the new features and functionality for NTFS in Windows Server 2012 R2.

Feature description

A file system is a part of the operating system that determines how files are named, stored, and organized on a volume. A file system manages files and folders, and the information needed to

locate and access these items. NTFS allows you to maximize security, availability, and reliability as well as a design for greater storage growth.

New and changed functionality

The following table describes the new features and improvements for NTFS in Windows Server 2012 R2.

Feature/functionality	New or updated	Summary
Sparse file performance	Updated	Sparse file performance is improved to enable increased deduplication performance.
Offline repair for volumes (spot fixing)	Updated	A larger number of errors and issues are repaired using the spot-fix corruption handling stages. This results in the file system being offline only for a brief period of time to repair errors. Integration with Cluster Shared Volumes allows online spot-fixing.
Memory changes	Updated	The file system cache in the management operating system scales up and down as virtual machine clients consume and release memory. In addition, the file system cache in the guest operating system scales up and down as virtual machines consume and release memory.

See also

The following list provides additional resources about NTFS and related technologies.

- NTFS

- File and Storage Services Overview
- Resilient File System Overview
- Understanding the Impact of Large Sector Media for IT Pros

NTFS Health and Chkdsk

Windows Server 2012 and Windows 8 enables customers to confidently deploy multiterabyte NTFS file system volumes, which supports consolidation scenarios and maximizes storage utilization. Historically, customers have not deployed large NTFS volumes. The average deployment sizes were approximately 500 gigabyte (GB) because in the event of file system corruption, the entire volume would be offline for an indeterminate period of time. Chkdsk in Windows Server 2012 and Windows 8 introduces a new approach that prioritizes volume availability and allows for the detection of corruption while the volume remains online with data available.

The new model of chkdsk has the following benefits:

- Customers can confidently deploy large volumes. Corruption-related downtime is now proportional to only the number of corruptions on the volume.
- Customers who are using clustered shared volumes do not see any downtime, even for correcting corruption events that would normally require a remount.
- Windows Server 2012 actively monitors the health state of the file system volume, and it always provides the health state to the administrator.
- Customers do not see any downtime for transient corruption events.
- Customers experience significantly fewer corruption events.

Requirements

- Chkdsk is enabled by default on Windows Server 2012.
- No special hardware is needed.

Technical overview

In the event of file system corruption in previous versions of Windows Server, the file system volume would be taken offline, and its data would be unavailable for an indeterminate period of time while chkdsk ran and attempted to repair the corruption. The time for chkdsk to complete was proportional to the number of files on the volume, and this time could not be accurately predicted by customers. For a single corruption issue, the entire volume and all of its files would still be analyzed offline. This model was problematic for customers because it potentially imposed significant downtimes, which made it challenging for customers to meet necessary business objectives. In addition, analysis indicated that a significant portion of file system corruptions were due to transient errors—chkdsk would run, but it would not find any actual corruption.

Windows Server 2012 addresses this problem by prioritizing file system availability, even when corruption is present, and developing a new model for managing file system corruption. The model includes:

- Improved self-healing: NTFS instantaneously self-heals more issues online without requiring chkdsk to run offline. This reduces the number of times that chkdsk is required to run.
- Online analysis: In the previous model, almost all of the offline time required to run chkdsk was spent scanning and analyzing the drive. In Windows Server 2012, the analysis phase of chkdsk, which was responsible for the majority of the offline time, becomes an online, background task. This allows the volume to remain online and available while the system determines whether there is corruption. There is also added logic in the model that verifies the type of corruption is not transient, preventing unnecessary analysis.
- Corruption correction: When the scan is completed, Windows Server 2012 informs the administrator (by using events and the management consoles), that the volumes need to be repaired and suggests a solution, such as performing a remount or reboot. Because the analysis phase has already completed, no additional scanning or detection is required. Chkdsk directly fixes the identified corruption, and the offline time is minimized to seconds. Therefore, the offline time for a volume is no longer proportional to the number of files on the volume, but rather to the number of corruptions on the volume.

Windows Server 2012 always provides the current health state of the file system volume by using standard events that indicate the current state of the volume. The following figure illustrates how significantly the new chkdsk model changes downtime in the event of corruption.

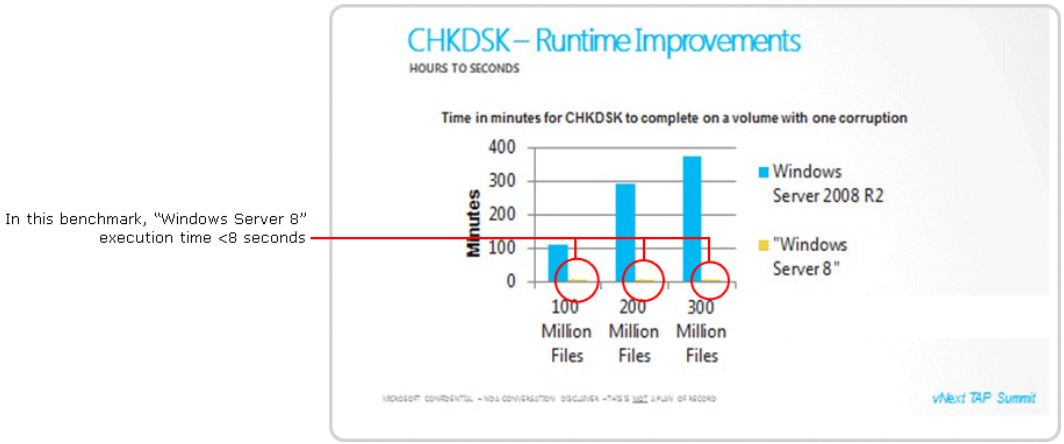

Figure 1 Chkdsk downtimes

Resilient File System Overview

This topic describes Resilient File System (ReFS), a new file system in Windows Server 2012, the deployment scenarios for the new file system, and the new and changed functionality for ReFS in Windows Server 2012 R2.

Feature description

Windows customers want a cost-effective platform that maximizes data availability, scales efficiently to very large data sets across diverse workloads, and guarantees data integrity by means of resiliency to corruption (regardless of software or hardware failures). ReFS is a new file system that targets these needs while providing a foundation for significant future innovations. By utilizing an integrated storage stack comprising ReFS and the new Storage Spaces feature in Windows Server 2012, customers can now deploy the most cost-effective platform for available and scalable data access using commodity storage.

Storage Spaces protects data from partial and complete disk failures by maintaining copies on multiple disks. ReFS interfaces with Storage Spaces to automatically repair the corruption. For more information about Storage Spaces, see Storage Spaces Overview and the Storage Spaces Microsoft TechNet blog.

The key attributes of ReFS include:

- Maintaining a high level of data availability and reliability, even when the individual underlying storage devices experience failures.
- Providing a full, end-to-end resilient architecture when used in conjunction with Storage Spaces. When used together, ReFS and Storage Spaces provide enhanced resiliency to storage device failures.

For information about new and updated functionality in Windows Server 2012 R2, see New and updated functionality later in this topic.

Deployment scenarios

ReFS provides functionality that helps customers store and protect data, regardless of the reliability of the underlying hardware and software stack. This minimizes the cost of storage and reduces capital expenditures for businesses. Customers can deploy a Windows Server 2012 R2 file server attached to a just a bunch of disks (JBOD) storage configuration with Serial ATA (SATA) or Serially Attached SCSI (SAS) drives. Additionally, the deployment could include failover clustering where the customer deploys a scale-out, two-node file server cluster with Storage Spaces, where the cluster uses a shared JBOD storage configuration with SAS drives.

Important functionality

The significant functionality included with ReFS is described as follows:

- **Integrity**. ReFS stores data in a way that protects it from many of the common errors that can normally cause data loss. When ReFS is used in conjunction with a mirror space or a parity space, detected corruption—both metadata and user data, when integrity streams are enabled—can be automatically repaired using the alternate copy provided by Storage Spaces. In addition, there are Windows PowerShell cmdlets (**Get-FileIntegrity** and **Set-FileIntegrity**) that you can use to manage the integrity and disk scrubbing policies.
- **Availability**. ReFS prioritizes the availability of data. Historically, file systems were often susceptible to data corruption that would require the system to be taken offline for repair.

With ReFS, if corruption occurs, the repair process is both localized to the area of corruption and performed online, requiring no volume downtime. Although rare, if a volume does become corrupted or you choose not to use it with a mirror space or a parity space, ReFS implements *salvage*, a feature that removes the corrupt data from the namespace on a live volume and ensures that good data is not adversely affected by nonrepairable corrupt data. Because ReFS performs all repair operations online, it does not have an offline **chkdsk** command.

- **Scalability**. As the amount and size of data that is stored on computers continues to rapidly increase, ReFS is designed to work well with extremely large data sets—petabytes and larger—without performance impact. ReFS is not only designed to support volume sizes of 2^64 bytes (allowed by Windows stack addresses), but ReFS is also designed to support even larger volume sizes of up to 2^78 bytes using 16 KB cluster sizes. This format also supports 2^64-1 byte file sizes, 2^64 files in a directory, and the same number of directories in a volume.
- **Proactive Error Correction**. The integrity capabilities of ReFS are leveraged by a data integrity scanner, which is also known as a *scrubber*. The integrity scanner periodically scans the volume, identifying latent corruptions and proactively triggering a repair of that corrupt data.

New and updated functionality

The following table lists the features and functionality that are updated or new to ReFS in Windows Server 2012 R2.

Feature/functionality	New or updated	Description
Corruptions on parity spaces	Updated	When a corruption occurs on a parity space, ReFS automatically corrects the corruption.
Subfolder recovery from ReFS metadata corruption	Updated	When corruption of ReFS metadata occurs, subfolders and their associated files are automatically recovered.
ReFS is available on client operating systems	New	ReFS is now available to use on Windows 8.1.
ReFS registry entry	New	The new **RefsDisableLastAccessUpdate** registry entry is equivalent to **NtfsDisableLastAccessUpdate**.
Storage cmdlets in Windows	New	New storage cmdlets, **Get-FileIntegrity** and **Set-FileIntegrity**, are available to manage integrity and

Feature/functionality	New or updated	Description
		disk scrubbing policies.

Corruptions on parity spaces

ReFS can automatically correct corruption on a parity space when integrity streams are enabled to detect corrupt data and because ReFS examines the second (and third) copies that the data parity spaces contain. ReFS then uses the correct version to correct the corruption.

Note
ReFS can already detect corruption on mirrored spaces and automatically repair those corruptions.

Subfolder recovery from ReFS metadata corruption

When the metadata for a ReFS directory is corrupted, subfolders and their associated files are automatically recovered. ReFS identifies and recovers the files while ReFS remains online. Unrecoverable corruption of the ReFS directory metadata affects only those files that are in the directory in which the corruption has occurred.

ReFS is available on client operating systems

In Windows 8.1, you have the option to format a storage space as ReFS using mirrored Storage Spaces in Control Panel. ReFS also supports alternate data streams (up to 128K for both Windows 8.1 and Windows Server 2012 R2), which enables Office document thumbnails, Mark of the Web (MOTW), and similar features available on Windows 8.1.

ReFS registry entry

ReFS includes a new registry entry, **RefsDisableLastAccessUpdate**, which is the equivalent of the previous **NtfsDisableLastAccessUpdate** registry entry. For more information, see NtfsDisableLastAccessUpdate.

Storage cmdlets in Windows

New storage cmdlets in Windows PowerShell are available (Get-FileIntegrity and Set-FileIntegrity) for you to manage the integrity and disk scrubbing policies.

See also

The following list provides additional resources on the web about related technologies in Windows Server 2012.

- File and Storage Services Overview
- Storage Spaces Overview

- Scale-Out File Server for Application Data Overview

Server for NFS Data Store

In the past several years, organizations have been increasingly pressured to reduce their IT spending. Reduced IT budgets drive a business need for IT to find new ways to decrease costs. Because the cost of storage constitutes a sizeable portion of IT budgets, this trend has resulted in the shift towards converged storage fabric. Virtualization is a large driver of the increase in demand for storage capacity. File-based storage has gained credence as a viable alternative to more expensive storage area network (SAN) block-based storage because it is simple to provision and manage. An example of this trend is the recent popularity of deploying and running VMware ESX and VMware ESXi on virtual machines from file-based storage accessed over the Network File System (NFS) protocol.

The need to store and run critical virtualized enterprise workloads from file-based storage demands a guarantee for increased availability and reliability. Windows Server 2012 includes several key improvements to the file system, storage, and networking stack. Server for NFS is the NFS server implementation that is included with Windows Server operating systems. For Windows Server 2012, Server for NFS has been updated to support continuous availability. These improvements make it possible to reliably store and run VMware ESX on a virtual machine that is running Windows Server 2012 by using the NFS protocol to share files. Additional improvements in failover clustering make it possible to deploy Server for NFS in a clustered configuration and take advantage of better resilience to hardware and software outages that may afflict individual cluster nodes. At the same time, failover clustering isolates VMware ESX hypervisor hosts from intermittent glitches during failovers.

Figure 1 NFS sharing files

Requirements

The following items are required to deploy and run VMware ESX on virtual machines that are using Server for NFS as a data store:

- A computer running Windows Server 2012 with the Failover Clustering feature installed and shared storage

- A server running VMware ESX 4.1
- A computer that is used for management (such as a computer running Windows 7) with VMware vSphere Client version 4.1 installed on it
- Installable media, such as a CD or ISO image of the guest operating system that you want to run on the VMware ESX hypervisor host

On the computer running Windows Server 2012, you need to use Server Manager to install the following features:

- The File Services role
- The Server for NFS role service
- The Failover Clustering feature

This feature set includes Windows PowerShell cmdlets that you can use to set NFS to share files. You can also use these tools to quickly provision the shared files as a data store for servers running VMware ESX and to run virtual machines on the shared files by using the NFS protocol.

Technical overview

Windows Server 2012 includes a new open network computing remote procedure call (ONCRPC), which consists of an External Data Representation (XDR) runtime built on the Winsock Kernel (WSK) interface. This configuration replaces the previous runtime that was built on the legacy Transport Driver Interface (TDI). The new XDR runtime is equipped with a dynamic autotuned thread pool that does not require performance tuning using registry keys, and it features a more efficient XDR encoding and decoding implementation. In addition, it implements dynamic endpoints and consumes Plug and Play (PnP) notifications to support the addition and removal of network interfaces at runtime. With this new design, the XDR runtime infrastructure is capable of performing faster and more targeted failovers to achieve continuous availability. In Windows Server 2012, Server for NFS and Client for NFS have been ported to run on this new scalable, high-performance XDR runtime infrastructure.

The failover cluster resource DLL is completely rewritten in Windows Server 2012, which enables Server for NFS to fail over much faster than on previous versions of Windows Server. The previous per-share cluster resource for NFS has been replaced by a consolidated per-virtual disk cluster resource, which results in fewer cluster resources on the failover cluster and faster end-to-end failovers. The failover logic within the NFS cluster resource has also been reworked for performance. The failover clustering infrastructure has undergone several key performance enhancements, which has resulted in faster failovers of the disk and network resources. This helps reduce the total amount of time for a failover to occur, and It makes it easy to deploy shared files as a reliable file-based storage backend for VMware ESX hypervisor hosts.

In addition, Windows Server 2012 includes a comprehensive set of task-oriented Windows PowerShell cmdlets that make it easy to provision and manage the NFS shares.

Server Message Block overview

This topic describes the SMB 3.0 feature in Windows Server® 2012 -- practical uses for the feature, the most significant new or updated functionality in this version compared to previous versions, and the hardware requirements.

Feature description

The Server Message Block (SMB) protocol is a network file sharing protocol that allows applications on a computer to read and write to files and to request services from server programs in a computer network. The SMB protocol can be used on top of its TCP/IP protocol or other network protocols. Using the SMB protocol, an application (or the user of an application) can access files or other resources at a remote server. This allows applications to read, create, and update files on the remote server. It can also communicate with any server program that is set up to receive an SMB client request. Windows Server 2012 introduces the new 3.0 version of the SMB protocol.

Practical applications

This section discusses some new practical ways to use the new SMB 3.0 protocol.

File storage for virtualization (Hyper-V™ over SMB). Hyper-V can store virtual machine files, such as configuration, Virtual hard disk (VHD) files, and snapshots, in file shares over the SMB 3.0 protocol. This can be used for both stand-alone file servers and clustered file servers that use Hyper-V together with shared file storage for the cluster.

Microsoft SQL Server over SMB. SQL Server can store user database files on SMB file shares. Currently, this is supported with SQL Server 2008 R2 for stand-alone SQL servers. Upcoming versions of SQL Server will add support for clustered SQL servers and system databases.

Traditional storage for end-user data. The SMB 3.0 protocol provides enhancements to the Information Worker (or client) workloads. These enhancements include reducing the application latencies experienced by branch office users when accessing data over wide area networks (WAN) and protecting data from eavesdropping attacks.

New and changed functionality

For information on new and changed functionality in Windows Server 2012 R2, see What's New in SMB in Windows Server 2012 R2.

SMB in Windows Server 2012 includes the new SMB 3.0 protocol and many new improvements which are described in the following table.

Feature/functionality	New or updated	Summary
SMB Transparent Failover	New	Enables administrators to

Feature/functionality	New or updated	Summary
		perform hardware or software maintenance of nodes in a clustered file server without interrupting server applications storing data on these file shares. Also, if a hardware or software failure occurs on a cluster node, SMB clients transparently reconnect to another cluster node without interrupting server applications that are storing data on these file shares.
SMB Scale Out	New	Using Cluster Shared Volumes (CSV) version 2, administrators can create file shares that provide simultaneous access to data files, with direct I/O, through all nodes in a file server cluster. This provides better utilization of network bandwidth and load balancing of the file server clients, and optimizes performance for server applications.
SMB Multichannel	New	Enables aggregation of network bandwidth and network fault tolerance if multiple paths are available between the SMB 3.0 client and the SMB 3.0 server. This enables server applications to take full advantage of all available network bandwidth and be resilient to a network failure.
SMB Direct	New	Supports the use of network adapters that have RDMA capability and can function at full speed with very low

Feature/functionality	New or updated	Summary
		latency, while using very little CPU. For workloads such as Hyper-V or Microsoft SQL Server, this enables a remote file server to resemble local storage.
Performance Counters for server applications	New	The new SMB performance counters provide detailed, per-share information about throughput, latency, and I/O per second (IOPS), allowing administrators to analyze the performance of SMB 3.0 file shares where their data is stored. These counters are specifically designed for server applications, such as Hyper-V and SQL Server, which store files on remote file shares.
Performance optimizations	Updated	Both the SMB 3.0 client and SMB 3.0 server have been optimized for small random read/write I/O, which is common in server applications such as SQL Server OLTP. In addition, large Maximum Transmission Unit (MTU) is turned on by default, which significantly enhances performance in large sequential transfers, such as SQL Server data warehouse, database backup or restore, deploying or copying virtual hard disks.
SMB-specific Windows PowerShell cmdlets	New	With Windows PowerShell cmdlets for SMB, an administrator can manage file shares on the file server, end to end, from the command line.

Feature/functionality	New or updated	Summary
SMB Encryption	New	Provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. Requires no new deployment costs, and no need for Internet Protocol security (IPsec), specialized hardware, or WAN accelerators. It may be configured on a per share basis, or for the entire file server, and may be enabled for a variety of scenarios where data traverses untrusted networks.
SMB Directory Leasing	New	Improves application response times in branch offices. With the use of directory leases, roundtrips from client to server are reduced since metadata is retrieved from a longer living directory cache. Cache coherency is maintained because clients are notified when directory information on the server changes. Works with scenarios for *HomeFolder* (read/write with no sharing) and *Publication* (read-only with sharing).

Hardware requirements

SMB Transparent Failover has the following requirements:

- A failover cluster running Windows Server 2012 with at least two nodes configured. The cluster must pass the cluster validation tests included in the validation wizard.
- File shares must be created with the Continuous Availability (CA) property, which is the default.
- File shares must be created on CSV volume paths to attain SMB Scale-Out.

- Client computers must be running Windows® 8 or Windows Server 2012, both of which include the updated SMB client that supports continuous availability.

Note
Note that down-level clients can connect to file shares that have the CA property, but transparent failover will not be supported for these clients.

SMB Multichannel has the following requirements:

- At least two computers running Windows Server 2012 are required. No extra features need to be installed—the technology is on by default.
- For information on recommended network configurations, see the See Also section at the end of this overview topic.

SMB Direct has the following requirements:

- At least two computers running Windows Server 2012 are required. No extra features need to be installed—the technology is on by default.
- Network adapters with RDMA capability are required. Currently, these adapters are available in three different types: iWARP, Infiniband, or RoCE (RDMA over Converged Ethernet).

See also

The following list provides additional resources on the web about SMB and related technologies in Windows Server 2012.

- File and Storage Services Overview
- Scale-Out File Server for Application Data Overview
- Deploying Fast and Efficient File Servers for Server Applications

Deploy Hyper-V over SMB

SMB 3.0 file shares can be used as shared storage for Hyper-V in Windows Server 2012 R2 and Windows Server 2012. With this capability, Hyper-V can store virtual machine files, which includes configuration, virtual hard disk (VHD) files, and snapshots, on SMB file shares. Listed below are the main advantages of storing application data for Hyper-V on SMB file shares:

- **Ease of provisioning and management.** You can manage file shares instead of storage fabric and logical unit numbers (LUNs).
- **Increased flexibility.** You can dynamically migrate virtual machines or databases in the data center.
- **Ability to take advantage of existing investment in a converged network.** You can use your existing converged network with no specialized storage networking hardware.
- **Reduced capital expenditures.** Capital expenses (acquisition costs) are reduced.
- **Reduced operating expenditures.** You can reduce operating costs because there is no need for specialized storage expertise.

Note

Some Hyper-V features in Windows Server 2012 R2 and Windows Server 2012 now rely on using SMB file shares, such as some types of live migration.

In this guide

- Prerequisites
- Step 1: Configuring file server clusters
- Step 2: Install Hyper-V
- Step 3: Create an SMB file share
- Step 4: Create a virtual machine and virtual hard disk file on the file share
- Step 5: Migrate virtual machine storage to an SMB file share
- Step 6: Initiate a live migration of a virtual machine to another cluster node
- Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage
- Troubleshooting

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

Using Hyper-V with SMB has the following requirements:

- One or more computers running Windows Server 2012 R2 or Windows Server 2012 with the Hyper-V role installed. You can also use non-Microsoft file servers that implement the SMB 3.0 protocol.
- One or more computers running Windows Server 2012 R2 or Windows Server 2012 with the File and Storage Services role installed.
- A common Active Directory infrastructure. The servers running Active Directory Domain Services (AD DS) do not need to run Windows Server 2012 R2 or Windows Server 2012.

The three most common file server configurations for Hyper-V over SMB are a single-node file server, a dual-node file server, and a multi-node file server, as shown in the following figure.

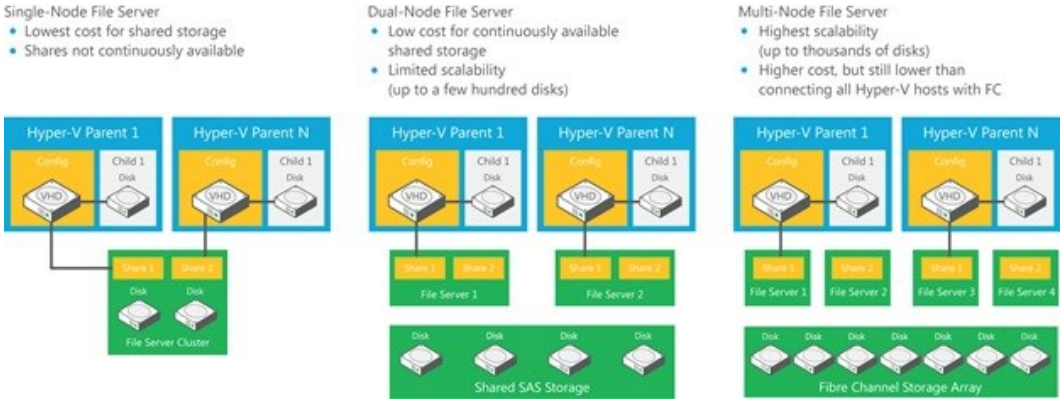

Figure 1 Common configurations for Hyper-V over SMB2

The two supported Hyper-V configurations for Hyper-V over SMB are:

- Standalone Hyper-V servers (not a high availability solution)
- Hyper-V servers configured in a failover cluster

Note
The Hyper-V host must have Windows Server 2012 R2 or Windows Server 2012 installed.

Considerations when using Hyper-V with SMB

- An Active Directory infrastructure is required, so you can grant permissions to the computer account of the Hyper-V hosts.
- The file server must have Windows Server 2012 R2 or Windows Server 2012 installed, so the new SMB 3.0 protocol is available. You can also use non-Microsoft file servers that implement the SMB 3.0 protocol. Hyper-V does not block older versions of SMB, however, the Hyper-V Best Practice Analyzer issues an alert when an older version of SMB is detected.
- Loopback configurations (where the computer that is running Hyper-V is used as the file server for virtual machine storage) are not supported.
- You must have separate failover clusters for Hyper-V and for the file server.

Step 1: Configuring file server clusters

To deploy Hyper-V over SMB, use one of the following procedures for your file server configuration. All servers in your file server configuration must have Windows Server 2012 R2 or Windows Server 2012 installed.

To configure a standalone file server

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, select the appropriate server, and then click **Next**. The local server is selected by default.
6. On the **Select server roles** page, click **File and Storage Services**, and then click **Next**.
7. On the **Confirm installation selections** page, click **Install**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To add the File and Storage Services role, type:

```
Install-WindowsFeature File-Services, FS-FileServer
```

If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available. To view the list of network adapters, type:

```
Get-NetAdapter
```

```
Get-SmbServerNetworkInterface
```

To configure a clustered file server

1. With Windows Server 2012 R2 or Windows Server 2012 installed on two servers, add the File and Storage Services role and the Failover Clustering feature on each server by typing:

   ```
   Install-WindowsFeature File-Services, FS-FileServer,
   Failover-Clustering
   ```

   ```
   Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature
   ```

2. If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available and that they are configured on different subnets. To view the list of network adapters, type:

   ```
   Get-NetAdapter
   ```

   ```
   Get-SmbServerNetworkInterface
   ```

3. To create a failover cluster using the two servers, type:

   ```
   New-Cluster –Name ClusterName –Node FileServer1, FileServer2
   ```

4. To create a file server cluster to host continuously available SMB file shares, where `FST` is the name of the file server cluster and `Cluster Disk 1` is the storage, type:

   ```
   Add-ClusterFileServerRole -Name FST -Storage "Cluster Disk 1"
   –StaticAddress 192.168.101.22/24, 192.168.102.22/24
   ```

   The above example assumes you are using two networks with addresses, 192.168.101.22/24 and 192.168.102.22/24, for SMB network traffic. Two networks are recommended for network fault tolerance.

To configure a failover cluster with a scale-out file server

1. Follow steps 1-3 in the previous procedure, To configure a clustered file server.
2. To create a scale-out file server on the failover cluster to host continuously available SMB file shares, where `FSO` is the name of the scale-out file server cluster and `Cluster Disk 2` is the name of the storage, type:

   ```
   Add-ClusterSharedVolume "Cluster Disk 2"
   ```

```
Add-ClusterScaleOutFileServerRole -Name FSO
```

Step 2: Install Hyper-V

To continue deploying Hyper-V over SMB, install the Hyper-V role on a separate server.

To install the Hyper-V role

1. To install the Hyper-V role on a separate server, follow steps 1-5 from To configure a standalone file server.
2. On the **Select server roles** page, click **Hyper-V**, and then click **Next**.
3. On the **Confirm installation selections** page, click **Install**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To install the Hyper-V role and Hyper-V Windows PowerShell cmdlets and tools, type:

```
Install-WindowsFeature Hyper-V, Hyper-V-PowerShell, Hyper-V-Tools
```

If you are using SMB Multichannel, ensure there are two network adapters with identical type and speed available, and they are not connected to the virtual switch. To view the list of network adapters, type:

```
Get-NetAdapter
```

```
Get-SmbClientNetworkInterface
```

Step 3: Create an SMB file share

The folder used by Hyper-V to store virtual machine data requires specific permissions to access the SMB file share. You need to make sure that the Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators have full control permissions.

To create an SMB file share by using Server Manager

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. On the left, click **File and Storage Services**.
4. Click **Tasks**, and then click **New Share** to open the **New Share Wizard**.
5. On the **Select Profile** page, select **SMB Share – Applications**, and click **Next**.
6. On the **Share Location** page, select a server and a volume, and click **Next**.
7. On the **Share Name** page, specify a name for the new share, and click **Next**.
8. On the **Permissions** page, click **Customize Permissions**.

9. Click **Add**, click **Select a Principal**, and then click **Object Types**.
10. In **Object Types**, click to select **Computers**, and click **OK**.
11. Enter the name of the computer, and click **OK**.
12. In **Permissions Entry**, select **Full Control**, and click **OK**.
13. Repeat the previous three steps for the second Hyper-V server. Click OK when finished.
14. On the **Permissions** page, click **Next**.
15. Click **Create** to create the SMB file share.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

For a standalone file server or a clustered file server, type the following to configure an SMB file share (where `HV1` and `HV2` are servers running Hyper-V, `HVC` is the Hyper-V cluster account, and `HVadmin` is the Hyper-V administrator account):

```
# Create folder

MD X:\VMS

# Create file share

New-SmbShare -Name VMS1 -Path X:\VMS -FullAccess Domain\HVAdmin, Domain\HV1$,
Domain\HV2$, Domain\HVC$

# Set NTFS permissions from the file share permissions

Set-SmbPathAcl –Name VMS1
```

For a scale-out file server cluster, type the following to configure an SMB file share (where `HV1` and `HV2` are servers running Hyper-V and `HVadmin` is the Hyper-V administrator account):

```
# Create folder

MD X:\VMS

# Create file share

New-SmbShare -Name VMS1 -Path X:\VMS -FullAccess Domain\HVAdmin, Domain\HV1$,
Domain\HV2$, Domain\HVC$

apply permissions

# Set NTFS permissions from the file share permissions

Set-SmbPathAcl –Name VMS1
```

Step 4: Create a virtual machine and virtual hard disk file on the file share

To create a virtual hard disk (VHD) and virtual machine on an SMB file share, use Hyper-V Manager or Hyper-V Windows PowerShell cmdlets. You will also need to specify a Universal Naming Convention (UNC) path (for example, \\servername\sharename).

To create a virtual machine

1. Open Hyper-V Manager. From the Server Manager Tools menu, click **Hyper-V Manager**.
2. From the navigation pane of Hyper-V Manager, select the computer running Hyper-V.
3. From the Actions pane, click **New** and then click **Virtual Machine**.
4. The New Virtual Machine wizard opens. Click **Next**.
5. On the **Specify Name and Location** page, type an appropriate name and location (using a UNC path). To configure Hyper-V over SMB, you must use a UNC path.
6. On the **Assign Memory** page, specify enough memory to start the guest operating system.
7. On the **Configure Networking** page, connect the virtual machine to the switch you created when you installed Hyper-V.
8. On the **Connect Virtual Hard Disk** and **Installation Options** pages, choose **Create a virtual hard disk**. Click **Next**, and then click the option that describes the type of media you will use. For example, to use an .iso file, click **Install an operating system from a boot CD/DVD** and then specify the path to the .iso file.
9. On the Summary page, verify your selections and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

For a file server named `FS1` and a file share named `VMS`, to create a virtual machine named `VM1`, type:

```
New-VHD -Path \\FS1\VMS\VM1.VHDX -VHDType Dynamic -SizeBytes 127GB

New-VM -Name VM1 -Path \\FS1\VMS -Memory 1GB -VHDPath \\FS1\VMS\VM1.VHDX
```

Step 5: Migrate virtual machine storage to an SMB file share

You can migrate virtual machine storage from both direct-attached storage (DAS) to an SMB file share and from a SMB file share to another SMB file share.

To migrate virtual machine storage from local storage to an SMB file share

1. To confirm that there is a virtual machine using local storage, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

   To migrate virtual machine storage from an SMB file share to another SMB file share, specify the SMB file share in this step.
2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To migrate the virtual machine storage to an SMB file share, type:

   ```
   Move-VMStorage –VMName VM1 –DestinationStoragePath \\FST\VMS
   ```
4. To confirm that a virtual machine is using an SMB file share and the workload is not interrupted, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

Step 6: Initiate a live migration of a virtual machine to another cluster node

You can transparently move running virtual machines from one cluster node to another node in the same cluster without a dropped network connection or perceived downtime.

To initiate a live migration of a virtual machine to another cluster node

1. To confirm that a virtual machine is running in a cluster node and is using an SMB file share, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   Get-ClusterGroup VM1 | FT Name, OwnerNode, State
   ```
2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To perform a live migration of a virtual machine to another cluster node, type:

   ```
   Move-ClusterVirtualMachineRole -Name VM1 -Node HV2 –
   VmMigrationType Live
   ```
4. To confirm that the virtual machine moved to another cluster node and the workload is not interrupted, type:

   ```
   Get-ClusterGroup VM1 | FT Name, OwnerNode, State
   ```

Step 7: Move virtual machines to another Hyper-V host and migrate virtual machine storage

To move running virtual machines from one cluster node to another node in the same cluster and migrate virtual machine storage, use one of the following procedures:

To move a virtual machine to another Hyper-V host

1. To confirm that a virtual machine is running on the Hyper-V host, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   ```

2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To perform a live migration of a virtual machine to another Hyper-V host, type:

   ```
   Move-VM –Name VM1 –DestinationHost HV2
   ```

4. To confirm that the virtual machine moved to another Hyper-V host and the workload is not interrupted, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   ```

To move a virtual machine with direct-attached storage and migrate its virtual machine storage to an SMB file share

1. To confirm that a virtual machine with local storage is running on the Hyper-V host, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   ```

   ```
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

   To migrate virtual machine storage from an SMB file share to another SMB file share, specify the SMB file share in this step.
2. Start a prolonged process, such as running a file copy workload, inside of the guest operating system.
3. To move the virtual machine to another Hyper-V host and the storage to an SMB file share, type:

   ```
   Move-VM –Name VM1 -DestinationHost HV2 –
   DestinationStoragePath \\FST\VMS
   ```

4. To confirm that the virtual machine moved to another Hyper-V host using the SMB file share, and the workload is not interrupted, type:

   ```
   Get-VM VM1 | FT Name, Path, State
   ```

   ```
   Get-VMHardDiskDrive VM1 | FT VMName, Path
   ```

Troubleshooting

This section covers some common issues that you might encounter when using Hyper-V over SMB.

Check and fix issues with permissions

You may experience issues with permissions that are related to accessing an SMB file share or the NTFS folder on which the share was created. To check permissions on an SMB file share (where `VMS1` is the share and `X:\VMS` is the NTFS folder), type the following:

```
Get-SmbShareAccess -Name VMS1 -Path X:\VMS
```

If you determine that a specific computer account is missing from permissions, you can add the account to both the file share and the folder to fix the issue. To fix the permissions, type the following:

```
Grant-SmbShareAccess –Name VMS1 –AccountName Domain\HV3$ -AccessRight Full
```

```
Set-SmbPathAcl –Name VMS1
```

Using constrained delegation

When using Hyper-V Manager from a computer running Windows Server 2012 R2 or Windows Server 2012 to manage virtual machines on another computer running Windows Server 2012 R2 or Windows Server 2012, you may experience an error that says access to an SMB file share is denied. Typically, this is because you need delegation rights to use your credentials to access the remote share on another computer. This is a security feature that prevents a user from gaining access to a computer in your network for the purpose of performing actions on other computers in your network. To address this issue, you have two choices:

Option 1: Use Remote Desktop. Use Remote Desktop to access the computer and run Hyper-V Manager directly on that computer.

Option 2: Configure constrained delegation. You can change the properties of the computer account in Active Directory Users and Computers to allow delegation. When enabled, constrained delegation gives you the ability to use a specific SMB remote file share without requiring you to perform an action on any computer. Constrained delegation tells Active Directory Users and Computers that between two computers, (in this case, the Hyper-V server and the SMB file server), and for specific services, (in this case, SMB), it is allowed to re-issue access to the resources.

To configure constrained delegation, for each server running Hyper-V, perform one of the following procedures.

Note

In Windows Server 2012 R2, you can perform this procedure using a new set of Windows PowerShell SMB cmdlets that simplifies the configuration of constrained delegation.

To configure constrained delegation

1. In Active Directory Users and Computers, click to open **Properties** for the computer account, and then click to open the **Delegation** tab.

2. Select both **Trust this computer for delegation to the specified services only** and **Use Kerberos only**.
3. Click **Add**, and provide the name of the SMB file server (or the **Cluster Access Point** for a Scale-Out File Server).
4. Select the **CIFS** service. Note that Common Internet File System (CIFS) is the previous name for SMB.
5. On the SMB file share created for virtual machines, add Full Control permissions for the Hyper-V Administrators.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Before you perform this procedure using the Windows PowerShell SMB cmdlets, the Active Directory module for Windows PowerShell must be available. To install the Active Directory cmdlets, type:

```
Install-WindowsFeature RSAT-AD-PowerShell
```

To configure constrained delegation (where `FileServer1` and `FileServer2` are the servers running Hyper-V and `HV1` and `HV2` are the SMB clients), type:

```
Enable-SmbDelegation –SmbServer FileServer1 –SmbClient HV1

Enable-SmbDelegation –SmbServer FileServer1 –SmbClient HV2

Enable-SmbDelegation –SmbServer FileServer2 –SmbClient HV1

Enable-SmbDelegation –SmbServer FileServer2 –SmbClient HV2
```

Note

This procedure works only with the resource-based constrained delegation available starting in Windows Server 2012, therefore, the Active Directory forest must be at the Windows Server 2012 functional level.

See also

- SMB Share Cmdlets in Windows PowerShell
- File and Storage Services Overview
- Increasing Server, Storage, and Network Availability: Scenario Overview

Improve Performance of a File Server with SMB Direct

Windows Server 2012 R2 and Windows Server 2012 include a feature called SMB Direct, which supports the use of network adapters that have Remote Direct Memory Access (RDMA) capability. Network adapters that have RDMA can function at full speed with very low latency, while using very little CPU. For workloads such as Hyper-V or Microsoft SQL Server, this enables a remote file server to resemble local storage. SMB Direct includes:

- Increased throughput: Leverages the full throughput of high speed networks where the network adapters coordinate the transfer of large amounts of data at line speed.
- Low latency: Provides extremely fast responses to network requests, and, as a result, makes remote file storage feel as if it is directly attached block storage.
- Low CPU utilization: Uses fewer CPU cycles when transferring data over the network, which leaves more power available to server applications.

SMB Direct is automatically configured by Windows Server 2012 R2 and Windows Server 2012.

SMB Multichannel and SMB Direct

SMB Multichannel is the feature responsible for detecting the RDMA capabilities of network adapters to enable SMB Direct. Without SMB Multichannel, SMB uses regular TCP/IP with the RDMA-capable network adapters (all network adapters provide a TCP/IP stack along with the new RDMA stack).

With SMB Multichannel, SMB detects whether a network adapter has the RDMA capability, and then creates multiple RDMA connections for that single session (two per interface). This allows SMB to use the high throughput, low latency, and low CPU utilization offered by RDMA-capable network adapters. It also offers fault tolerance if you are using multiple RDMA interfaces.

Notes

- You should not team RDMA-capable network adapters if you intend to use the RDMA capability of the network adapters. When teamed, the network adapters will not support RDMA.
- After at least one RDMA network connection is created, the TCP/IP connection used for the original protocol negotiation is no longer used. However, the TCP/IP connection is retained in case the RDMA network connections fail.

In this guide

- Requirements
  - Deploy SMB Direct with Ethernet (iWARP) Network Adapters
  - Deploy SMB Direct with InfiniBand Network Adapters
  - Deploy SMB Direct with RoCE Network Adapters
- Enabling and disabling SMB Direct
- Test performance of SMB Direct

Requirements

SMB Direct requires the following:

- At least two computers running Windows Server 2012 R2 or Windows Server 2012
- One or more network adapters with RDMA capability. Currently, these network adapters are available in three different types: iWARP, Infiniband, or RoCE (RDMA over Converged Ethernet). Use the links below for additional information on each of these types of RDMA-capable network adapters:

  Deploy SMB Direct with Ethernet (iWARP) Network Adapters

  Deploy SMB Direct with InfiniBand Network Adapters

  Deploy SMB Direct with RoCE Network Adapters

Considerations when using SMB Direct

- You can use SMB Direct in a failover cluster; however, you need to make sure that the cluster networks used for client access are adequate for SMB Direct. Failover clustering supports using multiple networks for client access, along with network adapters that are RSS (Receive Side Scaling)-capable and RDMA-capable.
- You can use SMB Direct on the Hyper-V management operating system to support using Hyper-V over SMB, and to provide storage to a virtual machine that uses the Hyper-V storage stack. However, RDMA-capable network adapters are not directly exposed to a Hyper-V client. If you connect an RDMA-capable network adapter to a virtual switch, the virtual network adapters from the switch will not be RDMA-capable.
- If you disable SMB Multichannel, SMB Direct is also disabled. Since SMB Multichannel detects network adapter capabilities and determines whether a network adapter is RDMA-capable, SMB Direct cannot be used by the client if SMB Multichannel is disabled.
- SMB Direct is not supported on Windows RT. SMB Direct requires support for RDMA-capable network adapters, which is available only on Windows Server 2012 R2 and Windows Server 2012.
- SMB Direct is not supported on down-level versions of Windows Server. It is supported only on Windows Server 2012 R2 and Windows Server 2012.

Enabling and disabling SMB Direct

SMB Direct is enabled by default when Windows Server 2012 R2 or Windows Server 2012 is installed. The SMB client automatically detects and uses multiple network connections if an appropriate configuration is identified.

Disabling SMB Direct

Typically, you will not need to disable SMB Direct, however, you can disable it by running one of the following Windows PowerShell scripts.

To disable RDMA for a specific interface, type:

```
Disable-NetAdapterRdma <name>
```

To disable RDMA for all interfaces, type:

```
Set-NetOffloadGlobalSetting -NetworkDirect Disabled
```

When you disable RDMA on either the client or the server, the systems cannot use it. *Network Direct* is the internal name for Windows Server 2012 R2 and Windows Server 2012 basic networking support for RDMA interfaces.

Re-enabling SMB Direct

After disabling RDMA, you can re-enable it by running one of the following Windows PowerShell scripts.

To re-enable RDMA for a specific interface, type:

```
Enable-NetAdapterRDMA <name>
```

To re-enable RDMA for all interfaces, type:

```
Set-NetOffloadGlobalSetting -NetworkDirect Enabled
```

You need to enable RDMA on both the client and the server to start using it again.

Test performance of SMB Direct

You can test how the performance is working by using one of the following procedures.

Compare a file copy with and without using SMB Direct

To measure the increased throughput of SMB Direct

1. Configure SMB Direct
2. Measure the amount of time to run a large file copy using SMB Direct.
3. Disable RDMA on the network adapter, see Enabling and disabling SMB Direct.
4. Measure the amount of time to run a large file copy without using SMB Direct.
5. Re-enable RDMA on the network adapter, and then compare the two results.
6. To avoid the impact of caching, you should do the following:
   a. Copy a large amount of data (more data than memory is capable of handling).
   b. Copy the data twice, with the first copy as practice and then timing the second copy.
   c. Restart both the server and the client before each test to make sure they operate under similar conditions.

Fail one of multiple network adapters during a file copy with SMB Direct

To confirm the failover capability of SMB Direct

1. Ensure that SMB Direct is functioning in a multiple network adapter configuration.
2. Run a large file copy. While the copying is run, simulate a failure of one of the network paths by disconnecting one of the cables (or by disabling one of the network adapters).
3. Confirm that the file copying continues using one of the remaining network adapters, and that there are no file copy errors.

Note
To avoid failures of a workload that does not use SMB Direct, make sure there are no other workloads using the disconnected network path.

See also

- Deploy SMB Direct with InfiniBand Network Adapters
- Deploy SMB Direct with RoCE Network Adapters
- Deploy SMB Direct with Ethernet (iWARP) Network Adapters
- Server Message Block overview
- Increasing Server, Storage, and Network Availability: Scenario Overview
- Deploy Hyper-V over SMB

Deploy SMB Direct with InfiniBand Network Adapters

This topic explains how to deploy SMB Direct with InfiniBand network adapters.

Overview

When you deploy SMB Direct with an RDMA-capable network adapter, the network adapter functions at full speed with very low latency, while using very little CPU. You can use the InfiniBand series of networks adapters to take full advantage of the capabilities of SMB Direct.

Hardware and software requirements

Use the following requirements to implement and test this deployment:

- Two or more computers running Windows Server 2012 R2 or Windows Server 2012
- One or more InfiniBand network adapters for each server
- One or more InfiniBand switches
- Two or more network cables as required for the InfiniBand network adapters that you select. These are typically enhanced small form-factor pluggable (SFP+), QSFP, or QSFP+ connectors.

Example InfiniBand configurations

There are many options available for InfiniBand network adapters, cables, and switches. The following are some examples of configurations that you can use with Windows Server 2012 R2 or Windows Server 2012:

1. **Two computers using FDR (Fourteen Data Rate) InfiniBand network adapters**. FDR InfiniBand has a 54 Gbps data rate. In this deployment, the minimum configuration is two network adapters and a network cable. You will need a system with PCIe Gen3 slots to achieve the 54 Gbps data rate. PCIe Gen3 slots are available on newer computer systems. If

you use an older system, the network adapter will be limited by the speed of the older PCIe Gen2 host bus adapter.

2. **Ten computers using dual FDR InfiniBand network adapters**. To increase throughput in a private cloud deployment, you can configure a two-node file server cluster, in addition to an eight-node Hyper-V cluster. For added performance and fault tolerance, use two InfiniBand network adapters for each system. In this configuration, you need 20 FDR InfiniBand network adapters and a 20-port FDR switch.

For additional information about the InfiniBand network adapters (and part numbers) that support SMB Direct, see the following blog post: Deploying Windows Server 2012 with SMB Direct (SMB over RDMA) and the Mellanox ConnectX-2/ConnectX-3 using InfiniBand – Step by Step.

Download and update the latest drivers

Although Windows Server 2012 R2 and Windows Server 2012 might include drivers for the network adapter, you should always verify that you have the latest version of the driver from the manufacturer of the network adapter. Follow the manufacturer's installation instructions to install the driver.

Configure a subnet manager

When you use an InfiniBand switch, you are required to have a subnet manager running. The best option is to use a managed InfiniBand switch (which runs a subnet manager), but you can also install a subnet manager on a computer that is connected to an unmanaged switch. This section discusses the options that are available for configuring a subnet manager.

Use a managed switch with a built-in subnet manager.

Many InfiniBand switches have an internal subnet manager that runs by default when the switch is powered on or runs when it is enabled by using the switch's web interface.

Use OpenSM with an unmanaged switch (recommended only for test environments)

If a managed switch is not available, you can use one of the computers running Windows Server 2012 R2 or Windows Server 2012 to run your subnet manager. You can use the OpenSM application that is provided by the OpenFabrics Alliance, which is available with some InfiniBand device driver distributions. The executable file for the OpenSM application is opensm.exe. Make sure that you install the tool as a service that starts automatically.

Important

We recommend that you do not use OpenSM with an unmanaged switch in production environments.

For additional fault tolerance, make sure that you have two computers on your network configured to run OpenSM. Do not run OpenSM on more than two computers that are connected to an InfiniBand switch.

Configure the IP address

After you have installed the latest version of the drivers, you should configure the IP address for the network adapters. If you are using a DHCP server, the network adapters are automatically configured, and you can move ahead to the following section.

To assign an IP address to your network adapter, you can use **IP Address Management** or the Net TCP/IP cmdlets in Windows PowerShell. For more information, see Net TCP/IP Cmdlets in Windows PowerShell.

The following Windows PowerShell example shows how to assign an IP address to a network adapter called `RDMA1`, with an IP address of `192.168.1.10`, and a DNS server at `192.168.1.2`.

```
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled

Remove-NetIPAddress -InterfaceAlias RDMA1 -AddressFamily IPv4 -Confirm:$false

New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.1.10 -PrefixLength 24 -Type
Unicast

Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.1.2
```

Verify the configuration

To verify that the configuration is working correctly, you can review configuration information about the network adapter and the SMB connection. For more information, see the Verify the configuration section.

Review the performance counters

There are several performance counters that you can use to verify that the RDMA-capable network adapters are being used and that the SMB Direct connections are established. For more information, see the Review the performance counters section.

Review the event logs for RDMA network adapters

You can use several events in the Windows event log to verify that the RDMA-capable network adapters are operational, and the SMB Direct connections are accessible. For more information, see the Review the event logs for RDMA network adapters section.

See also

- Improve Performance of a File Server with SMB Direct
- Deploy SMB Direct with Ethernet (iWARP) Network Adapters

- Server Message Block overview
- Increasing Server, Storage, and Network Availability: Scenario Overview
- Deploy Hyper-V over SMB

Deploy SMB Direct with RoCE Network Adapters

This topic explains how to deploy SMB Direct with RoCE network adapters.

Overview

When you deploy SMB Direct with an RDMA-capable network adapter, the network adapter functions at full speed with very low latency, while using very little CPU. You can use the RoCE series of networks adapters to take full advantage of the capabilities of SMB Direct.

Hardware and software requirements

Use the following requirements to implement and test this deployment:

- Two or more computers running Windows Server 2012 R2 or Windows Server 2012
- One or more RoCE network adapters for each server
- One or more 10 gigabit or 40 gigabit Ethernet switches with the Priority Flow Control (PFC) capability
- Two or more network cables as required for the RoCE network adapters that you select. These are typically enhanced small form-factor pluggable (SFP+) connectors for 10 gigabit switches or QSFP connectors 40 gigabit switches.

Example RoCE configurations

There are many options available for RoCE network adapters, cables, and switches. The following are some examples of configurations that you can use with Windows Server 2012 R2 or Windows Server 2012:

- **Two computers using 10 gigabit RoCE network adapters**. In this deployment, the minimum configuration is two network adapters and a network cable. You can use this deployment for testing a single file server with a server running Hyper-V. For 10 gigabit Ethernet, use network adapters with SFP+ connectors.
- **Eight computers using dual 10 gigabit RoCE network adapters**. In this deployment, you can configure a two-node file server cluster, in addition to a six-node Hyper-V cluster. In this configuration, you need eight computers with a dual port RoCE network adapter on each computer. In addition, you need a 10 gigabit switch with at least 16 ports (using 10 gigabit and SFP+ connectors). The network adapters you select must support Priority Flow Control (PFC).
- **Two computers using 40 gigabit RoCE network adapters**. FDR RoCE has a 54 Gbps data rate. In this deployment, the minimum configuration is two network adapters and a network cable (you should use a QSFP connector for 40 gigabit). You will need a system with

PCIe Gen3 slots to achieve the 54 Gbps data rate. PCIe Gen3 slots are available on newer computer systems. If you use an older system, the network adapter will be limited by the speed of the older PCIe Gen2 host bus adapter.

- **Ten computers using dual 40 gigabit RoCE network adapters**. To increase throughput in a private cloud deployment, you can configure a two-node file server cluster, in addition to an eight-node Hyper-V cluster. For added performance and fault tolerance, use two 40 gigabit RoCE network adapters for each system. In this configuration, you need 20 RoCE network adapters and a 20-port switch. You will need a system with PCIe Gen3 slots to achieve the rate speed of the network adapter. PCIe Gen3 slots are available on newer computer systems. If you use an older system, the network adapter will be limited by the speed of the older PCIe Gen2 host bus adapter.

For detailed information on the RoCE network adapters that support SMB Direct, see this blog post.

Download and update the latest drivers

Although Windows Server 2012 R2 and Windows Server 2012 might include drivers for the network adapter, you should always verify that you have the latest version of the driver from the manufacturer of the network adapter. Follow the manufacturer's installation instructions to install the driver.

Configuring Priority Flow Control

For RoCE to function reliably, you need to configure Priority Flow Control (PFC) on all cluster nodes and on all switches in the flow path.

Configure PFC on cluster nodes

To configure PFC on the cluster nodes, perform the following steps:

- If applicable, clear all previous configurations
- Enable Data Center Bridging (DCB) on both client and server
- Create a Quality of Service (QoS) policy to identify RoCE traffic on both client and server
- Enable PFC on a specific priority (the example shown below uses **Priority 3**)
- Enable the DCB settings for the network adapters (the example shown below assumes the network adapter is named **Ethernet 4**)
- Optionally, you can limit the bandwidth used by the SMB traffic (the example shown elow has a bandwidth limit of 60%)

The following example shows how to configure PFC using Windows PowerShell. To perform the steps listed above, type the following in Windows PowerShell:

```
# Clear previous configurations

Remove-NetQosTrafficClass

Remove-NetQosPolicy -Confirm:$False
```

```
# Enable DCB
Install-WindowsFeature Data-Center-Bridging

# Disable the DCBx setting:
Set-NetQosDcbxSetting -Willing 0

# Create QoS policies and tag each type of traffic with the relevant priority
New-NetQosPolicy "SMB" -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3
New-NetQosPolicy "DEFAULT" -Default -PriorityValue8021Action 3
New-NetQosPolicy "TCP" -IPProtocolMatchCondition TCP -PriorityValue8021Action 1
New-NetQosPolicy "UDP" -IPProtocolMatchCondition UDP -PriorityValue8021Action 1

# If VLANs are used, mark the egress traffic with the relevant VlanID:
Set-NetAdapterAdvancedProperty -Name <Network Adapter Name> -RegistryKeyword "VlanID" -
RegistryValue <ID>

# Enable Priority Flow Control (PFC) on a specific priority. Disable for others
Enable-NetQosFlowControl -Priority 3
Disable-NetQosFlowControl 0,1,2,4,5,6,7

# Enable QoS on the relevant interface
Enable-NetAdapterQos -InterfaceAlias "Ethernet 4"

# Optionally, limit the bandwidth used by the SMB traffic to 60%
New-NetQoSTrafficClass "SMB" -Priority 3 -Bandwidth 60 -Algorithm ETS
```

If you have a Kernel Debugger attached to the computer (which is typically applicable just for developers), flow control is always disabled. In this case, you need to run the following script in Windows PowerShell to disable this behavior:

```
Set-ItemProperty HKLM:"\SYSTEM\CurrentControlSet\Services\NDIS\Parameters"
AllowFlowControlUnderDebugger -Type DWORD -Value 1 –Force
```

Configure PFC on a switch

For RoCE to perform reliably, you must enable PFC on the switch. Configuring PFC on switches varies according to the type of switch you select. Refer to the hardware documentation for the switch for details on how to perform this step.

Important
This is not an optional step. If you do not configure PFC on the switch, your RoCE-capable network adapters will perform poorly.

Configure the IP address

After you have installed the latest version of the drivers, you should configure the IP address for the network adapters. If you are using a DHCP server, the network adapters are automatically configured, and you can move ahead to the following section.

To assign an IP address to your network adapter, you can use **IP Address Management** or the Net TCP/IP cmdlets in Windows PowerShell. For more information, see Net TCP/IP Cmdlets in Windows PowerShell.

The following Windows PowerShell example shows how to assign an IP address to a network adapter called `RDMA1`, with an IP address of `192.168.1.10`, and a DNS server at `192.168.1.2`.

```
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled

Remove-NetIPAddress -InterfaceAlias RDMA1 -AddressFamily IPv4 -Confirm:$false

New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.1.10 -PrefixLength 24 -Type
Unicast

Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.1.2
```

Verify the configuration

To verify that the configuration is working correctly, you can review configuration information about the network adapter and the SMB connection. For more information, see the Verify the configuration section.

Review the performance counters

There are several performance counters that you can use to verify that the RDMA-capable network adapters are being used and that the SMB Direct connections are established. For more information, see the Review the performance counters section.

Review the event logs for RDMA network adapters

You can use several events in the Windows event log to verify that the RDMA-capable network adapters are operational, and the SMB Direct connections are accessible. For more information, see the Review the event logs for RDMA network adapters section.

See also

- Improve Performance of a File Server with SMB Direct

- Deploy SMB Direct with InfiniBand Network Adapters
- Deploy SMB Direct with Ethernet (iWARP) Network Adapters
- Server Message Block overview
- Increasing Server, Storage, and Network Availability: Scenario Overview
- Deploy Hyper-V over SMB

Deploy SMB Direct with Ethernet (iWARP) Network Adapters

This topic explains how to deploy SMB Direct with an Ethernet (iWARP) network adapter.

Overview

When you deploy SMB Direct with an RDMA-capable network adapter, the network adapter functions at full speed with very low latency, while using very little CPU. You can use the Ethernet (iWARP) series of networks adapters to take full advantage of the capabilities of SMB Direct.

Hardware and software requirements

Use the following requirements to implement and test this deployment:

- Two or more computers running Windows Server 2012 R2 or Windows Server 2012
- For each server running Windows Server 2012 R2 or Windows Server 2012, use one or more iWARP network adapters
- One or more 10 gigabit or 40 gigabit Ethernet switches
- Two or more network cables as required for the iWARP network adapters that you select. These are typically enhanced small form-factor pluggable (SFP+) connectors for 10 gigabit Ethernet switches, and QSPF connectors for 40 gigabit Ethernet switches.

For detailed information about the iWARP network adapters that support SMB Direct, see the following blog post: Deploying Windows Server 2012 with SMB Direct (SMB over RDMA) and the Chelsio T4 Cards Using iWARP – Step-by-Step.

Download and update the latest drivers

Although Windows Server 2012 R2 and Windows Server 2012 might include drivers for the network adapter, you should always verify that you have the latest version of the driver from the manufacturer of the network adapter. Follow the manufacturer's installation instructions to install the driver.

Configure the IP addresses

After you have the latest version of the drivers installed, you should configure the IP address for the network adapters. If you are using DHCP, network adapters are automatically configured, and you can move ahead to the following section.

To assign an IP address to your network adapter, you can use either **IP Address Management** or the Net TCP/IP cmdlets in Windows PowerShell. For more information, see Net TCP/IP Cmdlets in Windows PowerShell.

The following Windows PowerShell example shows how to assign an IP address to a network adapter called `RDMA1`, with an IP address of `192.168.1.10`, a default gateway at `192.168.1.1`, and a DNS server at `192.168.1.2`.

```
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled

Remove-NetIPAddress -InterfaceAlias RDMA1 -AddressFamily IPv4 -Confirm:$false

New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.1.10 -PrefixLength 24 -Type
Unicast -DefaultGateway 192.168.1.1

Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.1.2
```

Configure Windows Firewall

Because iWARP uses TCP/IP to communicate, you need to configure Windows Firewall to allow TCP/IP traffic. Basically, you need to configure a firewall rule for the SMB server that allows incoming traffic from the SMB Direct clients. In Windows Server 2012, when SMB Direct is used with iWARP, it uses TCP port 5445, in addition to TCP port 445, which is used for SMB.

The following example shows how to configure the firewall rule by using Windows PowerShell. To allow access by the clients on the SMB server, type the following in Windows PowerShell:

```
Enable-NetFirewallRule FPSSMBD-iWARP-In-TCP
```

You should note that FPSSMBD-iWARP-In-TCP is preconfigured for Windows Server 2012 R2 and Windows Server 2012, and it allows incoming traffic on port 5445. You should enable it only with iWARP network adapters that have the appropriate drivers for use with SMB Direct. If you have multiple SMB servers, you need to enable this firewall rule on every server that uses SMB Direct with iWARP.

Allow access across multiple subnets

One key advantage of iWARP RDMA-capable network adapters is their ability to be routed across different subnets. Although the most common configuration is a single subnet deployment, you can use iWARP network adapters to connect computers across subnets. In Windows Server 2012 R2 and Windows Server 2012, this capability is disabled by default because it is not always required.

To enable **Network Direct** support, and therefore SMB Direct, you need to configure all SMB servers and SMB clients to allow RDMA to be routed across subnets. To do this, type the following in Windows PowerShell:

```
Set-NetOffloadGlobalSetting -NetworkDirectAcrossIPSubnets Allow

Disable-NetAdapter -InterfaceAlias RDMA1

Enable-NetAdapter -InterfaceAlias RDMA1
```

Note

If you disable and then re-enable the network adapter, the settings take effect without requiring you to restart the server.

We recommend that you apply the previous settings before you create any file shares. If you apply the settings (or make any other major network configuration changes) after you create a file share, the SMB client re-evaluates its connections when the new network adapters are detected (or every 10 minutes). You can also configure SMB to immediately update its connections by typing the following in Windows PowerShell on the SMB clients:

```
Update-SmbMultichannelConnection
```

Verify the configuration

You should use the following information to verify that the configuration is working correctly.

Verify the network adapter configuration

You can use the Net TCP/IP cmdlets in Windows PowerShell to verify that **Network Direct** is globally enabled and that you have RDMA-capable network adapters. In Windows PowerShell, type the following on both the SMB server and the SMB client:

```
Get-NetOffloadGlobalSetting | Select NetworkDirect

Get-NetAdapterRDMA

Get-NetAdapterHardwareInfo
```

Verify the SMB configuration

You can use the SMB Share cmdlets in Windows PowerShell to verify that SMB Multichannel is enabled, that the network adapters are properly recognized by SMB, and that the RDMA capability of the network adapters is properly identified. For more information, see SMB Share Cmdlets in Windows PowerShell.

On the SMB client, type the following in Windows PowerShell:

```
Get-SmbClientConfiguration | Select EnableMultichannel

Get-SmbClientNetworkInterface
```

On the SMB server, type the following in Windows PowerShell:

```
Get-SmbServerConfiguration | Select EnableMultichannel

Get-SmbServerNetworkInterface

netstat.exe -xan | ? {$_ -match "445"}
```

Note

The **Netstat** command confirms if the file server is listening on the RDMA-capable network adapters.

Verify the SMB connection

To verify that the SMB connection is working, on the SMB client, start an ongoing file copy to create a lasting session with the SMB server. While the file copy is ongoing, use the SMB Share cmdlets in Windows PowerShell to verify that the connection is using the correct SMB dialect and that SMB Direct is working. To do this, type the following in Windows PowerShell:

```
Get-SmbConnection

Get-SmbMultichannelConnection

netstat.exe -xan | ? {$_ -match "445"}
```

Note

You must run the script above while a workload is running on that specific client. If you have no activity while the script is running, the script might return nothing. This is because your SMB session has expired, and there are no current connections.

Review the performance counters

There are several performance counters that you can use to verify that the RDMA-capable network adapters are being used and that the SMB Direct connections are being established. You can also use the specific performance counters for the SMB server and the SMB client to verify the performance of SMB, including:

- I/O processors (to verify data requests per second)
- Latency (to verify average seconds per request)
- Throughput (to verify data bytes per second)

Following is a list of the relevant performance counters.

For SMB client:

- **RDMA activity**: One instance for each RDMA network adapter.
- **SMB Direct connection**: One instance for each SMB Direct connection.
- **SMB client shares**: One instance for each SMB file share that the client is currently using.

For SMB server:

- **RDMA activity**: One instance for each RDMA network adapter.
- **SMB Direct connection**: One instance for each SMB Direct connection.
- **SMB server shares**: One instance for each SMB file share that the server is currently sharing.
- **SMB server session**: One instance for each client SMB session that is established with the server.

Review the event logs for RDMA network adapters

You can use several events in the Windows event log to verify that the RDMA-capable network adapters are operational, and the SMB Direct connections are accessible.

To view all events associated with the RDMA-capable network adapters, type the following in Windows PowerShell:

```
Get-WinEvent -ListLog "*SMB*" | Get-WinEvent | ? Message -like "*RDMA*" | FL LogName, Id, TimeCreated, Level, Message
```

To view only the critical events associated with the RDMA-capable network adapters, type the following in Windows PowerShell:

```
Get-WinEvent -ListLog "*SMB*" | Get-WinEvent | ? { $_.Level -lt 4 -and $_. Message -like "*RDMA*" } | FL LogName, Id, TimeCreated, Level, Message
```

See also

- Improve Performance of a File Server with SMB Direct
- Deploy SMB Direct with InfiniBand Network Adapters
- Deploy SMB Direct with RoCE Network Adapters
- Server Message Block overview
- Increasing Server, Storage, and Network Availability: Scenario Overview
- Deploy Hyper-V over SMB

What's New in SMB in Windows Server 2012 R2

This topic introduces the new features and functionality for Server Message Block (SMB) in Windows Server 2012 R2.

Feature description

Server Message Block (SMB) is a network file sharing protocol that allows applications on a computer to read and write to files and to request services from server programs in a computer network.

For an overview description of SMB, see Server Message Block Overview.

New and changed functionality

The following table describes the new features and improvements for SMB 3.0 in Windows Server 2012 R2.

Feature/functionality	New or updated	Summary
Automatic rebalancing of Scale-Out File Server clients	New	This functionality improves scalability and manageability

Feature/functionality	New or updated	Summary
		for Scale-Out File Servers. SMB client connections are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share. This improves efficiency by reducing redirection traffic between file server nodes. Clients are redirected following an initial connection and when cluster storage is reconfigured.
Improved performance of SMB Direct (SMB over RDMA)	Updated	Improves performance for small I/O workloads by increasing efficiency when hosting workloads with small I/Os (such as an online transaction processing (OLTP) database in a virtual machine). These improvements are evident when using higher speed network interfaces, such as 40 Gbps Ethernet and 56 Gbps InfiniBand.
Improved SMB event messages	Updated	SMB events now contain more detailed and helpful information. This makes troubleshooting easier and reduces the need to capture network traces or enable more detailed diagnostic event logging. By default, the most relevant event channels are turned on, so you instantly capture all of the essential information. In addition, some events now include details on configuration and troubleshooting solutions.

Feature/functionality	New or updated	Summary
VHDX files as shared storage for guest clustering	New	Simplifies the creation of guest clusters by using shared VHDX files for shared storage inside the virtual machines. You can use this feature with VHDX files that are stored in Cluster Shared Volumes (CSV) or SMB Scale-Out file shares. For more information, see Deploy a Guest Cluster Using a Shared Virtual Hard Disk.
Hyper-V Live Migration over SMB	New	Enables you to perform a live migration of virtual machines by using SMB 3.0 as a transport. This allows you to take advantage of key SMB features, such as SMB Direct and SMB Multichannel, by providing high speed migration with low CPU utilization.
Improved SMB bandwidth management	New	Enables you to configure SMB bandwidth limits to control different SMB traffic types. There are three SMB traffic types: default, live migration, and virtual machine.
Support for multiple SMB instances on a Scale-Out File Server	New	Provides an additional instance on each cluster node in Scale-Out File Servers specifically for CSV traffic. A default instance can handle incoming traffic from SMB clients that are accessing regular file shares, while another instance only handles inter-node CSV traffic. This feature improves scalability and reliability of traffic between CSV nodes.
SMB 1.0 is now an optional	Updated	The SMB 1.0 features,

Feature/functionality	New or updated	Summary
feature		including the legacy computer browser service and Remote Administration Protocol (RAP), are now separate and can be eliminated. These features are still enabled by default, but if you have no older SMB clients, such as Windows XP or Windows Server 2003, you can remove the SMB 1.0 features to increase security and potentially reduce patching.

See also

The following list provides additional resources about SMB and related technologies.

- Server Message Block overview
- File and Storage Services Overview
- Scale-Out File Server for Application Data Overview
- Deploying Fast and Efficient File Servers for Server Applications

SMB Security Enhancements

This topic explains the SMB security enhancements in Windows Server 2012 R2 and Windows Server 2012, including:

- SMB Encryption
- Secure dialect negotiation
- New signing algorithm
- Disabling SMB 1.0

SMB Encryption

SMB Encryption provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. You can deploy SMB Encryption with minimal effort, but it may require small additional costs for specialized hardware or software. It has no requirements for Internet Protocol security (IPsec) or WAN accelerators. SMB Encryption can be configured on a per share basis or for the entire file server, and it can be enabled for a variety of scenarios where data traverses untrusted networks.

Note

SMB Encryption does not cover security at rest, which is typically handled by BitLocker Drive Encryption.

SMB Encryption should be considered for any scenario in which sensitive data needs to be protected from man-in-the-middle attacks. Possible scenarios include:

- An information worker's sensitive data is moved by using the SMB protocol. SMB Encryption offers an end-to-end privacy and integrity assurance between the file server and the client, regardless of the networks traversed, such as wide area network (WAN) connections that are maintained by non-Microsoft providers.
- SMB 3.0 enables file servers to provide continuously available storage for server applications, such as SQL Server or Hyper-V. Enabling SMB Encryption provides an opportunity to protect that information from snooping attacks. SMB Encryption is simpler to use than the dedicated hardware solutions that are required for most storage area networks (SANs).

Important

You should note that there is a notable performance operating cost with any end-to-end encryption protection when compared to non-encrypted.

Enable SMB Encryption

You can enable SMB Encryption for the entire file server or only for specific file shares. Use one of the following procedures to enable SMB Encryption:

To enable SMB Encryption by using Windows PowerShell

1. To enable SMB Encryption for an individual file share, type the following script on the server:

```
Set-SmbShare –Name <sharename> -EncryptData $true
```

2. To enable SMB Encryption for the entire file server, type the following script on the server:

```
Set-SmbServerConfiguration –EncryptData $true
```

3. To create a new SMB file share with SMB Encryption enabled, type the following script:

```
New-SmbShare –Name <sharename> -Path <pathname> –
EncryptData $true
```

To enable SMB Encryption by using Server Manager

1. In Server Manager, open **File and Storage Services**.
2. Click **Shares** to open the Shares management page.
3. Right-click the share on which you want to enable SMB Encryption, and then click **Properties**.
4. On the **Settings** page of the share, click **Encrypt data access**. Remote file access to

this share is encrypted.

Considerations for deploying SMB Encryption

By default, when SMB Encryption is enabled for a file share or server, only SMB 3.0 clients are allowed to access the specified file shares. This enforces the administrator's intent of safeguarding the data for all clients that access the shares. However, in some circumstances, an administrator may want to allow unencrypted access for clients that do not support SMB 3.0 (for example, during a transition period when mixed client operating system versions are being used). To allow unencrypted access for clients that do not support SMB 3.0, type the following script in Windows PowerShell:

```
Set-SmbServerConfiguration –RejectUnencryptedAccess $false
```

The secure dialect negotiation capability described in the next section prevents a man-in-the-middle attack from downgrading a connection from SMB 3.0 to SMB 2.0 (which would use unencrypted access). However, it does not prevent a downgrade to SMB 1.0, which would also result in unencrypted access. To guarantee that SMB 3.0 clients always use SMB Encryption to access encrypted shares, you must disable the SMB 1.0 server. (For instructions, see the section Disabling SMB 1.0.) If the **–RejectUnencryptedAccess** setting is left at its default setting of **$true**, only encryption-capable SMB 3.0 clients are allowed to access the file shares (SMB 1.0 clients will also be rejected).

Notes

- SMB Encryption uses the Advanced Encryption Standard (AES)-CCM algorithm to encrypt and decrypt the data. AES-CCM also provides data integrity validation (signing) for encrypted file shares, regardless of the SMB signing settings. If you want to enable SMB signing without encryption, you can continue to do this. For more information, see the following blog post: The Basics of SMB Signing.
- You may encounter issues when you attempt to access the file share or server if your organization uses wide area network (WAN) acceleration appliances.
- With a default configuration (where there is no unencrypted access allowed to encrypted file shares), if clients that do not support SMB 3.0 attempt to access an encrypted file share, Event ID 1003 is logged to the Microsoft-Windows-SmbServer/Operational event log, and the client will receive an **Access denied** error message.
- SMB Encryption and the Encrypting File System (EFS) in the NTFS file system are unrelated, and SMB Encryption does not require or depend on using EFS.
- SMB Encryption and the BitLocker Drive Encryption are unrelated, and SMB Encryption does not require or depend on using BitLocker Drive Encryption.

Secure dialect negotiation

SMB 3.0 is capable of detecting man-in-the-middle attacks that attempt to downgrade the SMB 2.0 or SMB 3.0 protocol or the capabilities that the client and server negotiate. When such an attack is detected by the client or the server, the connection is disconnected and event ID 1005 is logged in the Microsoft-Windows-SmbServer/Operational event log. Secure dialect negotiation cannot detect or prevent downgrades from SMB 2.0 or 3.0 to SMB 1.0. Because of

this, and to take advantage of the full capabilities of SMB Encryption, we strongly recommend that you disable the SMB 1.0 server. For more information, see Disabling SMB 1.0.

The secure dialect negotiation capability that is described in the next section prevents a man-in-the-middle attack from downgrading a connection from SMB 3 to SMB 2 (which would use unencrypted access); however, it does not prevent downgrades to SMB 1, which would also result in unencrypted access. For more information on potential issues with earlier non-Windows implementations of SMB, see the Microsoft Knowledge Base.

New signing algorithm

SMB 3.0 uses a more recent encryption algorithm for signing: Advanced Encryption Standard (AES)-cipher-based message authentication code (CMAC). SMB 2.0 used the older HMAC-SHA256 encryption algorithm. AES-CMAC and AES-CCM can significantly accelerate data encryption on most modern CPUs that have AES instruction support. For more information, see the following blog post: The Basics of SMB Signing.

Disabling SMB 1.0

The legacy computer browser service and Remote Administration Protocol features in SMB 1.0 are now separate, and they can be eliminated. These features are still enabled by default, but if you do not have older SMB clients, such as computers running Windows Server 2003 or Windows XP, you can remove the SMB 1.0 features to increase security and potentially reduce patching.

Note

SMB 2.0 was introduced in Windows Server 2008 and Windows Vista. Older clients, such as computers running Windows Server 2003 or Windows XP, do not support SMB 2.0; and therefore, they will not be able to access file shares or print shares if the SMB 1.0 server is disabled. In addition, some non-Microsoft SMB clients may not be able to access SMB 2.0 file shares or print shares (for example, printers with “scan-to-share” functionality).

Determine whether SMB clients use SMB 1.0

To determine whether any SMB clients are currently connected to the server running SMB 1.0, type the following script in Windows PowerShell:

```
Get-SmbSession | Select Dialect,ClientComputerName,ClientUserName | ? Dialect -lt 2
```

Note

You should run this script repeatedly over the course of a week (multiple times each day) to build an audit trail. You could also run this as a scheduled task.

Disable SMB 1.0

To disable SMB 1.0, type the following script in Windows PowerShell:

```
Set-SmbServerConfiguration –EnableSMB1Protocol $false
```

Note

If an SMB client connection is denied because the server running SMB 1.0 has been disabled, event ID 1001 will be logged in the Microsoft-Windows-SmbServer/Operational event log. You can find the name and IP address of the denied client in the event log details.

See also

The following list provides additional resources on the web about SMB and related technologies in Windows Server 2012.

- Server Message Block overview
- File and Storage Services Overview
- Scale-Out File Server for Application Data Overview

Storage Management Overview

This topic discusses the unified interface for storage management in Windows Server 2012, including requirements and an overview of the technology.

Did you mean…

- File and Storage Services Overview
- Storage Spaces Overview
- File Services Windows Server Protocols

Technology description

Windows Server 2012 enables storage management that is comprehensive and fully scriptable, and administrators can manage it remotely. A WMI-based interface provides a single mechanism through which to manage all storage, including non-Microsoft intelligent storage subsystems and virtualized local storage (known as Storage Spaces). Additionally, management applications can use a single Windows API to manage different storage types by using standards-based protocols such as Storage Management Initiative Specification (SMI-S).

Requirements

Your storage infrastructure should support the Storage Management Provider interface or support an appropriate version of SMI-S.

Technical overview

The unified interface for storage management in Windows Server 2012 provides not only a core set of defined WMI and Windows PowerShell interfaces, but also features for more advanced management, as shown in the following figure.

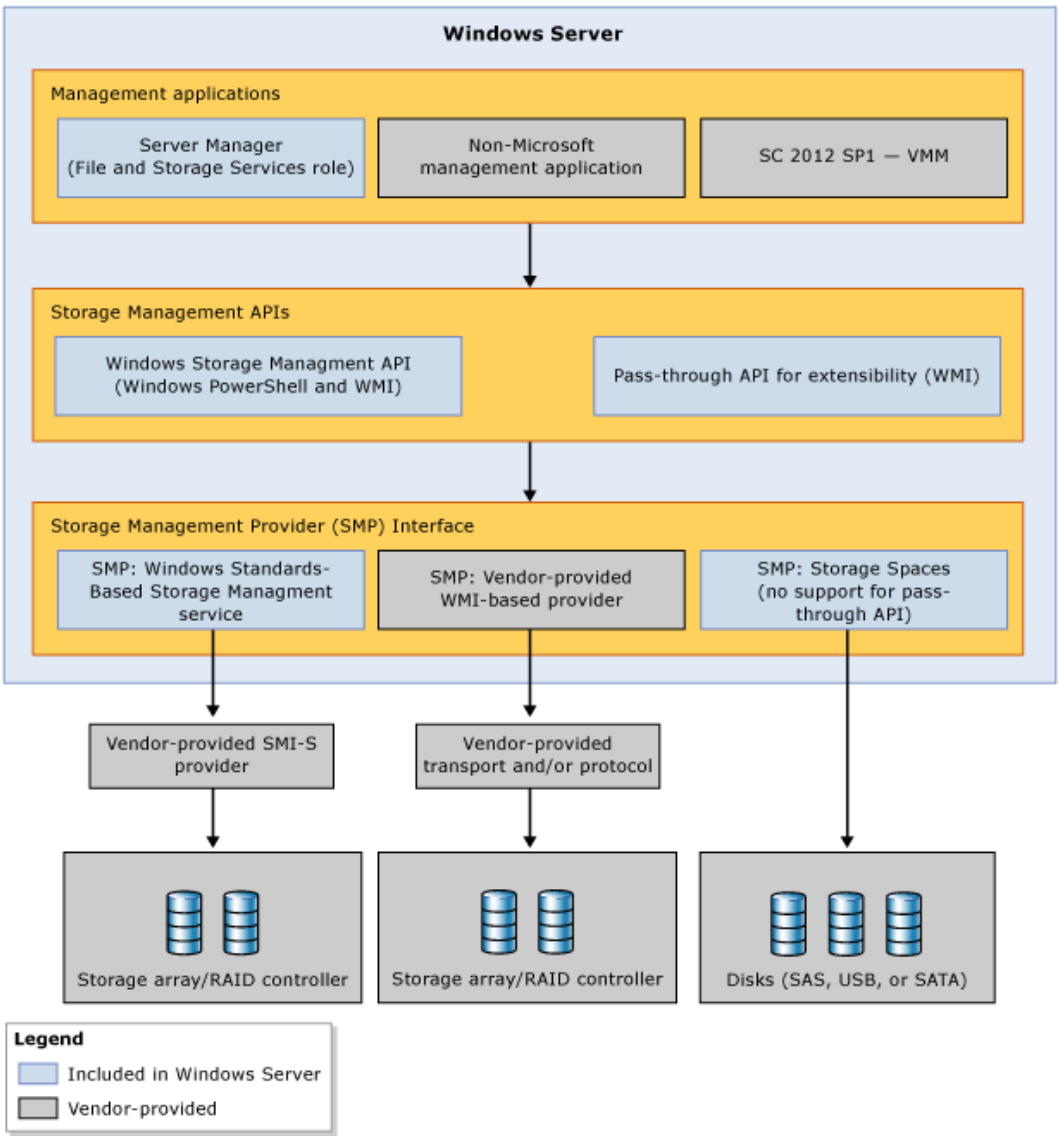

Figure 1 Unified storage management architecture

This design offers specific advantages for the following users:

- **Enterprise system administrators**. Uniform, scriptable management by using Windows PowerShell with a comprehensive set of cmdlets that can be used for discovery; thin provisioning support; snapshot management; replication; masking and unmasking; enumerating HBA ports; and creating pools, logical units, and volumes. As an example, an

administrator can use a single script to configure host resources and to configure and present the storage to Windows.

- **ISVs**. Flexibility to administer any type of storage that is connected to a Windows system.
- **Storage manufacturers**. Seamless integration of devices with any storage management client, which ensures consistent experience for users.

The unified storage management interface requires the use of a storage management provider that is based on SMI-S or WMI. If you install the Windows Standards-Based Storage Management feature, the interface integrates seamlessly with services offered by SMI-S providers, which enables administrators to manage non-Microsoft intelligent storage subsystems that have SMI-S providers. For WMI-based services, the interface includes a rich set of built-in storage management features that ISVs will find particularly helpful.

For more information about SMI-S and Windows Server 2012, see Introduction to SMI-S.

See also

For additional related information, see the following resources.

Content type	References
Evaluation	Introduction to SMI-S
Deployment	Understand and Troubleshoot Storage Spaces and Pools in Windows Server "8" Beta
Tools and settings	**N:Storage** (Windows PowerShell cmdlets) | Storage Management API Classes
Community resources	The File Services and Storage TechNet Forum | Windows Storage Team Blog | Storage Team at Microsoft File Cabinet Blog | Jose Barreto's Blog | Windows Server Information Experience on Twitter
Related technologies	Data Deduplication Overview | **High-Performance, Continuously Available File Share Storage for Server Applications Technical Preview** | iSCSI Target Block Storage Overview | Windows Offloaded Data Transfers Overview | Resilient File System Overview | Storage Spaces Overview | Supporting Information Workers with Reliable File Services and Storage | Thin Provisioning and Trim Storage Overview

Storage Spaces Overview

This topic discusses Storage Spaces, a technology in Windows and Windows Server that enables you to virtualize storage by grouping industry-standard disks into storage pools, and then create virtual disks called *storage spaces* from the available capacity in the storage pools. This topic provides an overview of the technology and new Storage Spaces functionality in Windows Server 2012 R2.

Did you mean…

- Storage Spaces: Frequently Asked Questions (Windows 8.1)
- High Performance Storage Solutions with Windows 8
- Storage Management Overview
- File and Storage Services Overview
- Storage Spaces Frequently Asked Questions (FAQ) (TechNet Wiki)

Technology description

Storage Spaces enables cost-effective, highly available, scalable, and flexible storage solutions for business-critical (virtual or physical) deployments. Storage Spaces delivers sophisticated storage virtualization capabilities, which empower customers to use industry-standard storage for single computer and scalable multi-node deployments. It is appropriate for a wide range of customers, including enterprise and cloud hosting companies, which use Windows Server for highly available storage that can cost-effectively grow with demand.

With Storage Spaces the Windows storage stack has been fundamentally enhanced to incorporate two new abstractions:

- **Storage pools**. A collection of physical disks that enable you to aggregate disks, expand capacity in a flexible manner, and delegate administration.
- **Storage spaces**. Virtual disks created from free space in a storage pool. Storage spaces have such attributes as resiliency level, storage tiers, fixed provisioning, and precise administrative control.

Storage Spaces is manageable through the Windows Storage Management API in Windows Management Instrumentation (WMI) and Windows PowerShell, and through the File and Storage Services role in Server Manager. Storage Spaces is completely integrated with failover clustering for high availability, and it is integrated with CSV for scale-out deployments.

Important functionality

Storage Spaces includes the following features:

- **Storage pools**. Storage pools are the fundamental building blocks for Storage Spaces. Storage administrators are already familiar with this concept, obviating the need to learn a new model. They can flexibly create storage pools based on the needs of the deployment. For example, given a set of physical disks, an administrator can create one pool (by using all the available physical disks) or multiple pools (by dividing the physical disks as required).

Furthermore, to maximize the value from storage hardware, the administrator can combine hard disks and solid-state drives (SSDs) in the same pool, using storage tiers to move frequently accessed portions of files to SSD storage, and using write-back caches to buffer small random writes to SSD storage. Pools can be expanded dynamically by simply adding additional drives, thereby seamlessly scaling to cope with unceasing data growth.

- **Resilient storage**. Storage Spaces provides three storage layouts (also known as resiliency types):
  - **Mirror**. Data is duplicated on two or three physical disks, increasing reliability, but reducing capacity. This storage layout requires at least two disks to protect you from a single disk failure, or at least five disks to protect you from two simultaneous disk failures.
  - **Parity**. Data and parity information are striped across physical disks, increasing reliability, but somewhat reducing capacity. This storage layout requires at least three disks to protect you from a single disk failure and at least seven disks to protect you from two disk failures.
  - **Simple (no resiliency)**. Data is striped across physical disks, maximizing capacity and increasing throughput, but decreasing reliability. This storage layout requires at least one disk and does not protect you from a disk failure.

  Additionally, Storage Spaces can automatically rebuild mirror and parity spaces in which a disk fails by using dedicated disks that are reserved for replacing failed disks (hot spares), or more rapidly by using spare capacity on other drives in the pool. Storage Spaces also includes background scrubbing and intelligent error correction to allow continuous service availability despite storage component failures. In the event of a power failure or cluster failover, the integrity of data is preserved so that recovery happens quickly and does not result in data loss.
- **Continuous availability**. Storage Spaces is fully integrated with failover clustering, which allows it to deliver continuously available service deployments. One or more pools can be clustered across multiple nodes within a single cluster. Storage spaces can then be instantiated on individual nodes, and the storage will seamlessly fail over to a different node when necessary (in response to failure conditions or due to load balancing). Integration with CSVs permits scale-out access to data.
- **Storage tiers**. Storage Spaces in Windows Server 2012 R2 combines the best attributes of SSDs and hard disk drives (HDDs) by enabling the creation of virtual disks comprised of two tiers of storage – an SSD tier for frequently accessed data, and a HDD tier for less-frequently accessed data. Storage Spaces transparently moves data at a sub-file level between the two tiers based on how frequently data is accessed. As a result, storage tiers can dramatically increase performance for the most used (“hot”) data by moving it to SSD storage, without sacrificing the ability to store large quantities of data on inexpensive HDDs.
- **Write-back cache**. Storage Spaces in Windows Server 2012 R2 supports creating a write-back cache that uses a small amount of space on existing SSDs in the pool to buffer small random writes. Random writes, which often dominate common enterprise workloads, are directed to SSDs and later are written to HDDs.
- **Operational simplicity**. The Windows Storage Management API, WMI, and Windows PowerShell permit full scripting and remote management. Storage Spaces can also be easily managed through the File and Storage Services role in Server Manager. Storage Spaces

also provides notifications when the amount of available capacity in a storage pool hits a configurable threshold.

- **Multitenancy**. Administration of storage pools can be controlled through access control lists (ACLs) and delegated on a per-pool basis, thereby supporting hosting scenarios that require tenant isolation. Storage Spaces follows the familiar Windows security model; therefore, it can be fully integrated with Active Directory Domain Services.

New and changed functionality

The following table describes some of the major changes in Storage Spaces. For more detailed information, see What's New in Storage Spaces in Windows Server 2012 R2.

Feature/functionality	New or updated?	Description
Storage tiers	New	Automatically moves frequently accessed data to faster (SSD) storage and infrequently accessed data to slower (HDD) storage.
Write-back cache	New	Buffers small random writes to SSD storage, reducing the latency of writes.
Parity space support for failover clusters	New	Enables you to create parity spaces on failover clusters.
Dual parity	New	Stores two copies of the parity information on a parity space, helping protect you from two simultaneous disk failures while optimizing for storage efficiency.
Automatically rebuild storage spaces from storage pool free space	New	Greatly decreases how long it takes to rebuild a storage space after a disk failure by using spare capacity in the pool instead of a single hot spare.

Requirements

Storage Spaces has the following requirements:

- Windows Server 2012 R2, Windows Server 2012, Windows 8.1, or Windows 8

- Serial ATA (SATA) or Serial Attached SCSI (SAS) connected disks, optionally in a just-a-bunch-of-disks (JBOD) enclosure

  RAID adapters, if used, must have all RAID functionality disabled and must not obscure any attached devices, including enclosure services provided by an attached JBOD.

  Note

  Consumers can use USB drives with Storage Spaces, though USB 3 drives are recommended to ensure a high level of performance. USB 2 drives will decrease performance – a single USB 2 hard drive can saturate the bandwidth available on the shared USB bus, limiting performance when multiple drives are attached to the same USB 2 controller. When using USB 2 drives, plug them directly into different USB controllers on your computer, do not use USB hubs, and add USB 2 drives to a separate storage pool used only for storage spaces that do not require a high level of performance.
- For shared-storage deployments on failover clusters:
  - Two or more servers running Windows Server 2012 R2 or Windows Server 2012
  - Requirements as specified for failover clustering and Cluster Shared Volumes (CSV)
  - SAS connected JBODs that comply with Windows Certification requirements

    For a list of certified JBODs, see the Windows Server Catalog.
  - Identical SAS HBAs that are certified for use with Storage Spaces

    These HBAs are connected to all JBOD enclosures in the file server cluster, and can't have built-in RAID functionality. For a list of certified HBAs, see the Windows Server Catalog.

See also

For additional related information, see the following resources.

Content type	References
Evaluation	What's New in Storage Spaces in Windows Server 2012 R2 | Virtualizing storage for scale, resiliency, and efficiency (Building Windows 8 Blog) | Storage Spaces Overview (TechNet Wiki)
Deployment	Deploy Storage Spaces on a Stand-Alone Server | Deploy Clustered Storage Spaces | Storage Spaces - Designing for Performance
Operations	Managing Storage with Windows PowerShell on Windows Server 2012 | Deploy and Manage Storage Spaces with Windows PowerShell | Storage Spaces Cmdlets | How Storage

Content type	References
	Spaces Makes Use of Hot Spares | How Storage Spaces responds to errors on physical disks
Community resources	Storage Spaces Frequently Asked Questions (FAQ) | Windows Server Storage Spaces: What is it and why should I use it? (CuratedAnswers) | Step-by-step for Storage Spaces Tiering in Windows Server 2012 R2 | The File Services and Storage TechNet Forum | Windows Storage Team Blog |Storage Team at Microsoft File Cabinet Blog | Jose Barreto's Blog | Windows Server Information Experience on Twitter
Related technologies	Data Deduplication Overview | **High-Performance, Continuously Available File Share Storage for Server Applications Technical Preview** | iSCSI Target Block Storage Overview | Windows Offloaded Data Transfers Overview | Resilient File System Overview | Storage Spaces Overview | Supporting Information Workers with Reliable File Services and Storage | Thin Provisioning and Trim Storage Overview

Deploy Storage Spaces on a Stand-Alone Server

This topic describes how to deploy Storage Spaces on a stand-alone Windows Server 2012−based server. For information about how to create a clustered storage space, see Deploy Clustered Storage Spaces.

To create a storage space, you must first create one or more storage pools. A storage pool is a collection of physical disks. A storage pool enables storage aggregation, elastic capacity expansion, and delegated administration.

From a storage pool, you can create one or more virtual disks. These virtual disks are also referred to as *storage spaces*. A storage space appears to the Windows operating system as a regular disk from which you can create formatted volumes. When you create a virtual disk through the File and Storage Services user interface, you can configure the resiliency type

(simple, mirror, or parity), the provisioning type (thin or fixed), and the size. Through Windows PowerShell, you can set additional parameters such as the number of columns, the interleave value, and which physical disks in the pool to use. For information about these additional parameters, see New-VirtualDisk and What are columns and how does Storage Spaces decide how many to use? in Storage Spaces Frequently Asked Questions (FAQ).

Note

You cannot use a storage space to host the Windows Server 2012 operating system.

From a virtual disk, you can create one or more volumes. When you create a volume, you can configure the size, the drive letter or folder, the file system (NTFS file system or Resilient File System (ReFS)), the allocation unit size, and an optional volume label.

The following figure illustrates the Storage Spaces workflow.

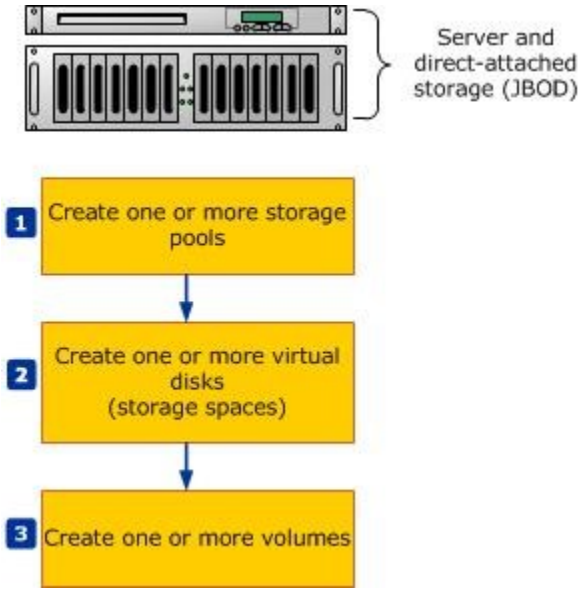

Figure 1: Storage Spaces workflow

In this topic

- Prerequisites
- Step 1: Create a storage pool
- Step 2: Create a virtual disk
- Step 3: Create a volume

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

To use Storage Spaces on a stand-alone Windows Server 2012−based server, make sure that the physical disks that you want to use meet the following prerequisites.

Important
If you want to deploy Storage Spaces on a failover cluster, see the topic Deploy Clustered Storage Spaces. Realize that a failover cluster deployment has different prerequisites, such as the supported disk bus types, the supported resiliency types, and the minimum number of disks that are required.

Area	Requirement
Disk bus types	• Serial Attached SCSI (SAS) • Serial Advanced Technology Attachment (SATA) **Notes** • You can also use USB drives. However, we do not recommend that you use USB drives in a server environment. • Storage Spaces does not support iSCSI and Fibre Channel controllers.
Disk configuration	• Physical disks must be at least 4 GB. • Disks must be blank and not formatted. Do not create volumes.
HBA considerations	We recommend that you use simple host bus adapters (HBAs) that do not support RAID functionality. If RAID capable, HBAs must be in non-RAID mode with all RAID functionality disabled. Adapters must not abstract the physical disks, cache data, or obscure any attached devices. This includes enclosure services that are provided by attached just-a-bunch-of-disks (JBOD) devices. Storage Spaces is compatible only with HBAs where you can completely disable all RAID functionality.
JBOD enclosures	A JBOD enclosure is optional. For full Storage Spaces functionality if you are using a JBOD enclosure, verify with your storage vendor that the JBOD enclosure supports Storage Spaces. To determine whether the JBOD enclosure

Area	Requirement
	supports enclosure and slot identification, run the following Windows PowerShell cmdlet: `Get-PhysicalDisk | ? {$_.BusType –eq "SAS"} | fc` If the **EnclosureNumber** and **SlotNumber** fields contain values, this indicates that the enclosure supports these features.

To plan for the number of physical disks and the desired resiliency type for a stand-alone server deployment, use the following guidelines.

Resiliency Type	Disk Requirements	When To Use
Simple • Stripes data across physical disks. • Maximizes disk capacity and increases throughput. • Does not provide resiliency.	Requires at least one physical disk. **Warning** A simple space does not protect from disk failure.	• Do not use to host irreplaceable data. • Use to host temporary or easily recreated data at a reduced cost. • Suited for high performance workloads where resiliency is not required or is provided by the application.
Mirror • Stores two or three copies of the data across the set of physical disks. • Increases reliability, but reduces capacity. Duplication occurs with every write. A mirror space also stripes the data across multiple physical drives. • Greater data throughput than parity, and lower access latency. • Uses dirty region tracking (DRT) to track modifications to the disks in the pool. When the	• Requires at least two physical disks to protect from single disk failure. • Requires at least five physical disks to protect from two simultaneous disk failures.	Use for most deployments. For example, mirror spaces are suited for a general-purpose file share or a virtual hard disk (VHD) library.

Resiliency Type	Disk Requirements	When To Use
system resumes from an unplanned shutdown and the spaces are brought back online, DRT makes disks in the pool consistent with each other.		
Parity • Stripes data and parity information across physical disks. • Increases reliability when it is compared to a simple space, but somewhat reduces capacity. • Increases resiliency through journaling. This helps prevent data corruption if an unplanned shutdown occurs.	Requires at least three physical disks to protect from single disk failure.	Use for workloads that are highly sequential, such as archive or backup.

Step 1: Create a storage pool

You must first group available physical disks into one or more storage pools.

Do this step using Windows PowerShell

To create a storage pool

1. In the Server Manager navigation pane, click **File and Storage Services**.
2. In the navigation pane, click the **Storage Pools** page.

   By default, available disks are included in a pool that is named the *primordial* pool. If no primordial pool is listed under **STORAGE POOLS**, this indicates that the storage does not meet the requirements for Storage Spaces. Make sure that the disks meet the requirements that are outlined in the Prerequisites section.

   Tip
   If you select the **Primordial** storage pool, the available physical disks are listed under **PHYSICAL DISKS**.
3. Under **STORAGE POOLS**, click the **TASKS** list, and then click **New Storage Pool**.

   The New Storage Pool Wizard opens.
4. On the **Before you begin** page, click **Next**.
5. On the **Specify a storage pool name and subsystem** page, enter a name and optional

description for the storage pool, select the group of available physical disks that you want to use, and then click **Next**.

6. On the **Select physical disks for the storage pool** page, do the following, and then click **Next**:
   a. Select the check box next to each physical disk that you want to include in the storage pool.
   b. If you want to designate one or more disks as hot spares, under **Allocation**, click the drop-down arrow, and then click **Hot Spare**.
7. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.
8. On the **View results** page, verify that all tasks completed, and then click **Close**.

   Note
   Optionally, to continue directly to the next step, you can select the **Create a virtual disk when this wizard closes** check box.
9. Under **STORAGE POOLS**, verify that the new storage pool is listed.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example shows which physical disks are available in the primordial pool.

```
Get-StoragePool -IsPrimordial $true | Get-PhysicalDisk | Where-Object CanPool -eq $True
```

The following example creates a new storage pool that is named *StoragePool1*. It uses all available disks.

```
New-StoragePool –FriendlyName StoragePool1 –StorageSubsystemFriendlyName “Storage Spaces*” –PhysicalDisks (Get-PhysicalDisk –CanPool $True)
```

The following example creates a new storage pool *StoragePool1* that uses four of the available disks.

```
New-StoragePool –FriendlyName StoragePool1 –StorageSubsystemFriendlyName “Storage Spaces*” –PhysicalDisks (Get-PhysicalDisk PhysicalDisk1, PhysicalDisk2, PhysicalDisk3, PhysicalDisk4)
```

The following example sequence of cmdlets shows how to add an available physical disk *PhysicalDisk5* as a hot spare to the storage pool *StoragePool1*.

```
$PDToAdd = Get-PhysicalDisk –FriendlyName PhysicalDisk5

Add-PhysicalDisk –StoragePoolFriendlyName StoragePool1 –PhysicalDisks $PDToAdd –Usage HotSpare
```

Step 2: Create a virtual disk

Next, you must create one or more virtual disks from the storage pool. When you create a virtual disk, you can select how the data is laid out across the physical disks. This affects both reliability and performance. You can also select whether to create thin- or fixed-provisioned disks.

Do this step using Windows PowerShell

To create a virtual disk

1. If the New Virtual Disk Wizard is not already open, on the **Storage Pools** page in Server Manager, under **STORAGE POOLS**, make sure that the desired storage pool is selected.
2. Under **VIRTUAL DISKS**, click the **TASKS** list, and then click **New Virtual Disk**.

   The New Virtual Disk Wizard opens.
3. On the **Before you begin** page, click **Next**.
4. On the **Select the storage pool** page, click the desired storage pool, and then click **Next**.
5. On the **Specify the virtual disk name** page, enter a name and optional description, and then click **Next**.
6. On the **Select the storage layout** page, click the desired layout, and then click **Next**.

   Note
   If you select a layout where you do not have enough physical disks, you will receive an error message when you click **Next**. For information about which layout to use and the disk requirements, see the Prerequisites section of this topic.
7. If you selected **Mirror** as the storage layout, and you have five or more disks in the pool, the **Configure the resiliency settings** page appears. Select one of the following options:
   - **Two-way mirror**
   - **Three-way mirror**
8. On the **Specify the provisioning type** page, click one of the following options, and then click **Next**.
   - **Thin**

     With thin provisioning, space is allocated on an as-needed basis. This optimizes the usage of available storage. However, because this enables you to over-allocate storage, you must carefully monitor how much disk space is available.
   - **Fixed**

     With fixed provisioning, the storage capacity is allocated immediately, at the time a virtual disk is created. Therefore, fixed provisioning uses space from the storage pool that is equal to the virtual disk size.

     Tip
     With Storage Spaces, you can create both thin- and fixed-provisioned virtual disks in the same storage pool. For example, you could use a thin-provisioned

virtual disk to host a database and a fixed-provisioned virtual disk to host the associated log files.

9. On the **Specify the size of the virtual disk** page, do the following:

   If you selected thin provisioning in the previous step, in the **Virtual disk size** box, enter a virtual disk size, select the units (**MB**, **GB**, or **TB**), and then click **Next**.

   If you selected fixed provisioning in the previous step, click one of the following:

   - **Specify size**

     To specify a size, enter a value in the **Virtual disk size** box, and then select the units (**MB**, **GB**, or **TB**).

     If you use a storage layout other than simple, the virtual disk uses more free space than the size that you specify. To avoid a potential error where the size of the volume exceeds the storage pool free space, you can select the **Create the largest virtual disk possible, up to the specified size** check box.

   - **Maximum size**

     Select this option to create a virtual disk that uses the maximum capacity of the storage pool.

10. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.
11. On the **View results** page, verify that all tasks completed, and then click **Close**.

    Tip
    By default, the **Create a volume when this wizard closes** check box is selected. This takes you directly to the next step.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example creates a 50 GB virtual disk that is named *VirtualDisk1* on a storage pool that is named *StoragePool1*.

```
New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName VirtualDisk1 -Size (50GB)
```

The following example creates a mirrored virtual disk that is named *VirtualDisk1* on a storage pool that is named *StoragePool1*. The disk uses the maximum storage capacity of the storage pool.

```
New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName VirtualDisk1 -ResiliencySettingName Mirror -UseMaximumSize
```

The following example creates a 50 GB virtual disk that is named *VirtualDisk1* on a storage pool that is named *StoragePool1*. The disk uses the thin provisioning type.

```
New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName VirtualDisk1 -Size (50GB) -ProvisioningType Thin
```

The following example creates a virtual disk that is named *VirtualDisk1* on a storage pool that is named *StoragePool1*. The virtual disk uses three-way mirroring and is a fixed size of 20 GB.

Note
You must have at least five physical disks in the storage pool for this cmdlet to work. (This does not include any disks that are allocated as hot spares.)

```
New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName VirtualDisk1 -
ResiliencySettingName Mirror -NumberOfDataCopies 3 -Size 20GB -ProvisioningType Fixed
```

Step 3: Create a volume

Next, you must create a volume from the virtual disk. You can assign an optional drive letter or folder, and then format the volume with a file system.

Do this step using Windows PowerShell

To create a volume

1. If the New Volume Wizard is not already open, on the **Storage Pools** page in Server Manager, under **VIRTUAL DISKS**, right-click the desired virtual disk, and then click **New Volume**.

   The New Volume Wizard opens.
2. On the **Before you begin** page, click **Next**.
3. On the **Select the server and disk** page, do the following, and then click **Next**.
   a. In the **Server** area, click the server on which you want to provision the volume.
   b. In the **Disk** area, click the virtual disk on which you want to create the volume.
4. On the **Specify the size of the volume** page, enter a volume size, specify the units (**MB**, **GB**, or **TB**), and then click **Next**.
5. On the **Assign to a drive letter or folder** page, configure the desired option, and then click **Next**.
6. On the **Select file system settings** page, do the following, and then click **Next**.
   a. In the **File system** list, click **NTFS** or **ReFS**.
   b. In the **Allocation unit size** list, either leave the setting at **Default** or set the allocation unit size.

      Note
      For more information about allocation unit size, see Default cluster size for NTFS, FAT, and exFAT.
   c. Optionally, in the **Volume label** box, enter a volume label name, for example **HR Data**.
7. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.
8. On the **View results** page, verify that all tasks completed, and then click **Close**.

9. To verify that the volume was created, in Server Manager, click the **Volumes** page.

   The volume is listed under the server where it was created. You can also verify that the volume is in Windows Explorer.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the previous procedure. Enter the command on a single line.

The following example initializes the disks for virtual disk *VirtualDisk1*, creates a partition with an assigned drive letter, and then formats the volume with the default NTFS file system.

```
Get-VirtualDisk -FriendlyName VirtualDisk1 | Get-Disk | Initialize-Disk -Passthru | New-Partition -AssignDriveLetter -UseMaximumSize | Format-Volume
```

See also

- Storage Spaces Overview
- Storage Cmdlets in Windows PowerShell
- Deploy Clustered Storage Spaces
- Storage Spaces Frequently Asked Questions (FAQ)

Deploy Clustered Storage Spaces

This topic describes how to deploy clustered storage spaces by using Storage Spaces and Failover Clustering in Windows Server. Together, these technologies provide a resilient, highly available, and cost-efficient solution that you can scale from simple deployments to the needs of a large datacenter. You can build a failover cluster for your physical workloads or for virtual workloads that are available through the Hyper-V role.

The basic building block of a clustered storage spaces deployment is a small collection of servers, typically two to four, and a set of shared Serial Attached SCSI (SAS) just-a-bunch-of-disks (JBOD) enclosures. The JBOD enclosures should be connected to all the servers, where each server has redundant paths to all the disks in each JBOD enclosure. The following figure shows an example of the basic building block.

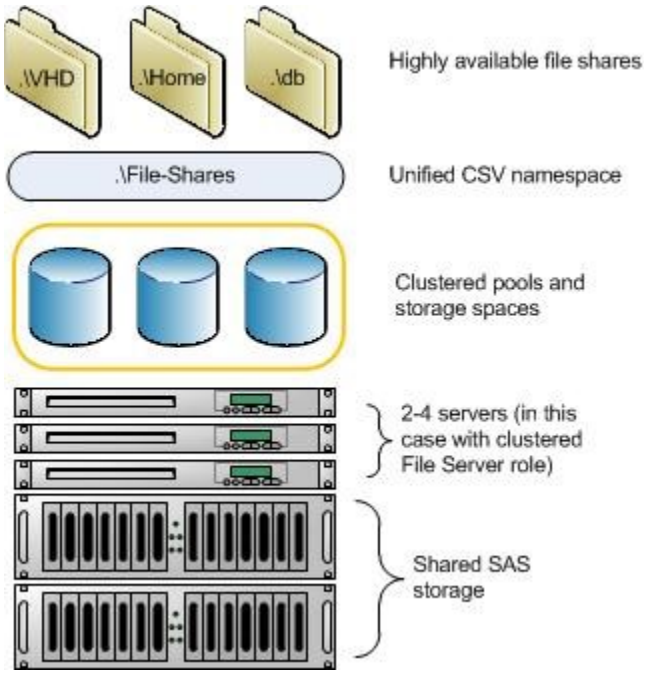

Figure 1: Example of clustered storage spaces

By using Cluster Shared Volumes (CSVs), you can unify storage access into a single namespace for ease of management. A common namespace folder is created at the path C:\ClusterStorage\ with all the CSVs in the failover cluster. All cluster nodes can access a CSV at the same time, regardless of the number of servers, the number of JBOD enclosures, or the number of provisioned virtual disks. This unified namespace enables highly available workloads to transparently fail over to another server if a server failure occurs. It also enables you to easily take a server offline for maintenance.

Clustered storage spaces can help protect against the following risks:

- **Physical disk failures** When you deploy a clustered storage space, protection against physical disk failures is provided by creating storage spaces with the mirror or parity resiliency types. Additionally, mirror spaces use dirty region tracking (DRT) to track modifications to the disks in the pool. When the system resumes from a power fault or a hard reset event and the spaces are brought back online, DRT makes disks in the pool consistent with each other.
- **Data access failures** If you have redundancy at all levels, you can protect against failed components, such as a failed cable from the enclosure to the server, a failed SAS adapter, power faults or failure of a JBOD backplane. For example, in an enterprise deployment, you should have redundant SAS adapters, SAS I/O modules, and power supplies. To protect against complete disk enclosure failure, you can use redundant JBOD enclosures.
- **Data corruptions and volume unavailability** Both the NTFS file system and the Resilient File System (ReFS) help protect against corruption. For NTFS, improvements to the Chkdsk tool in Windows Server 2012 can greatly improve availability. For more information, see NTFS Health and Chkdsk. For more information about ReFS, see Resilient File System Overview.

- **Server node failures** Through the Failover Clustering feature, you can provide high availability for the underlying storage and workloads. This helps protect against server failure and enables you to take a server offline for maintenance without service interruption.

Note
You cannot use a storage space to host the Windows system drive.

In this topic

- Prerequisites
- Step 1: Enable MPIO on each server
- Step 2: Verify that all servers can see the shared disks
- Step 3: (optional) Create storage spaces through File and Storage Services
- Step 4: Create a failover cluster
- Step 5: (optional) Create clustered storage spaces in Failover Cluster Manager
- Step 6: (optional) Add a cluster disk to a CSV

This topic includes two methods that you can use to create a clustered storage space. You can do either of the following:

- Configure storage spaces in File and Storage Services before you create the failover cluster, and then add the storage spaces as eligible storage during failover cluster creation. If you want to use this method, follow steps 1 through 4 in this topic. Optionally, you can complete step 6 to add the volume to a CSV.
- Create a clustered storage space by using Failover Cluster Manager after you create a failover cluster. If you want to use this method, skip step 3 in this topic. Optionally, you can complete step 6 to add the volume to a CSV.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

To support clustered storage spaces, your configuration must meet the prerequisites that are outlined in the following table.

Area	Requirement
Operating system	Storage Spaces and the Failover Clustering feature are supported in all editions of Windows Server 2012 R2 and Windows Server 2012. This includes Server Core installations.
Disk bus type	The disk bus type must be SAS. **Notes** • We recommend dual-port SAS drives for

Area	Requirement
	redundancy. • Storage Spaces does not support iSCSI and Fibre Channel controllers.
Resiliency type	The simple, mirror, and parity resiliency types are supported for a failover cluster though the parity resiliency type is not supported on failover clusters running Windows Server 2012. **Warning** Do not use a simple space to host irreplaceable data.
Disk configuration	For a failover cluster, you must have at least three physical disks in addition to any hot spare disks. Also, realize the following: • For two-way mirroring on failover clusters, you must have at least three physical disks. • For three-way mirroring on failover clusters, you must have at least five physical disks. • The disks must be blank and not formatted. Do not create volumes. • All disks must be at least 4 GB. • The physical disks that are used for a clustered storage pool must be dedicated to that pool. You cannot add startup disks.
HBA considerations	We recommend that you use simple host bus adapters (HBAs) that do not support RAID functionality. If RAID capable, HBAs must be in non-RAID mode with all RAID functionality disabled. Adapters must not abstract the physical disks, cache data, or obscure any attached devices. This includes enclosure services that are provided by attached JBOD devices. Storage Spaces is compatible only with HBAs where you can completely disable all RAID functionality. For a list of certified HBAs, see the Windows Server Catalog.
JBOD enclosures	For full Storage Spaces functionality, use a certified JBOD. For a list of certified devices,

Area	Requirement
	see the Windows Server Catalog. Disks must provide persistent reservations and pass cluster validation. To determine whether the JBOD enclosure supports enclosure and slot identification, run the following Windows PowerShell cmdlet: `Get-PhysicalDisk | ? {$_.BusType -eq "SAS"} | fc` If the **EnclosureNumber** and **SlotNumber** fields have values, this indicates that the enclosure supports these features.

Step 1: Enable MPIO on each server

With highly available storage, where there are multiple paths from the physical disks in the storage enclosure to the server, you must use multipath I/O (MPIO) software. The Windows Multipath I/O feature combines the multiple physical paths to physical disks as a single logical path that is used for data access. Also, it provides resiliency to a failure of a physical path and load balancing across physical paths.

For multipath support, you must enable MPIO on each server that can access the shared storage.

Do this step using Windows PowerShell

To enable MPIO

1. In Server Manager, click the server on which you want to install MPIO.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Select features** page of the Add Roles and Features Wizard, select the **Multipath I/O** check box.
4. Complete the Add Roles and Features Wizard.
5. Repeat this procedure for each server that you want to add as a failover cluster node.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example installs the MPIO feature on a server that is named *Server1*.

```
Install-WindowsFeature –ComputerName Server1 –Name Multipath-IO
```

Step 2: Verify that all servers can see the shared disks

Before you create the failover cluster, verify that all servers that you want to add as cluster nodes can see the shared disks. You can do this by using Server Manager, the Disk Management tool, or through Windows PowerShell.

Do this step using Windows PowerShell

To verify that a server can see the shared disks

1. To verify from one console that all servers can see the shared disks, make sure that you add all computers that you want to add as cluster nodes to Server Manager.
2. In Server Manager, click **File and Storage Services**, and then under **Volumes**, click **Disks**.
3. Under each server, verify that the shared disks are listed. Also, verify the following:
   - If you want to create storage spaces before you create the failover cluster, the disks do not have to be online or initialized. The disks must not be partitioned.
   - If you want to create a clustered storage space by using Failover Cluster Manager after you create a failover cluster, the disks must be online and initialized without partitions.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example lists the disks that are visible to the server. This includes their online or offline status and excludes the system disks.

```
Get-Disk | Where-Object IsSystem –eq $False
```

The following example lists the disks that are available for the storage pool, where the bus type is SAS.

```
Get-PhysicalDisk –CanPool $true | Where-Object BusType –eq “SAS”
```

The following example initializes the disks that you specify. You can obtain the disk numbers from the output of the **Get-Disk** cmdlet.

```
Initialize-Disk 1, 2, 3
```

Step 3: (optional) Create storage spaces through File and Storage Services

This step shows how to create storage spaces before you create the failover cluster by using the shared physical disks. When you create the failover cluster, if you designate that you want to add all eligible storage, the storage pools are automatically added as clustered storage pools with the associated virtual disks and volumes.

Skip this step if you want to create a clustered storage space by using Failover Cluster Manager after you create a failover cluster.

Do this step using Windows PowerShell

To create storage spaces through File and Storage Services

1. On one of the servers that you want to add as a cluster node, configure storage spaces as described in Deploy Storage Spaces on a Stand-Alone Server. However, to use with failover clustering, make sure that you do the following:
   - Verify that the disks meet the requirements that are outlined in the Prerequisites section of this topic. Only SAS disks are supported in a clustered storage space.
   - When you create a storage pool, you must select at least three physical disks. If you want to use the mirror resiliency type to protect against disk failure, make sure that you select at least three physical disks for a two-way mirror or at least five physical disks for a three-way mirror.
   - When you create virtual disks, you must do the following:
     - Specify the storage layout.
     - Specify **Fixed** as the provisioning type.
     - If you have multiple JBOD enclosures, and you want Storage Spaces to try to construct a storage space across them to help protect against complete JBOD failure, you must use Windows PowerShell to create virtual disks. For example, the following cmdlet shows how to create a virtual disk with the **IsEnclosureAware $True** parameter.

   New-VirtualDisk –StoragePoolFriendlyName *StoragePool* –FriendlyName *VirtualDisk1* –ResiliencySettingName Mirror –ProvisioningType Fixed –Size *10TB* –IsEnclosureAware $True

     - If you want to create storage spaces with storage tiers, first use the New-StorageTier cmdlet to create the appropriate tiers, then use the **New-VirtualDisk** cmdlet with the **–Storageliers** parameter.
   - When you create volumes, you must format the volumes with NTFS if you want to add the volumes to CSVs on servers running Windows Server 2012.

Windows PowerShell equivalent commands

The Deploy Storage Spaces on a Stand-Alone Server topic includes the associated Windows PowerShell examples.

Step 4: Create a failover cluster

This step shows how to install a failover cluster that you can use to host clustered storage spaces.

Note

For more information about failover clustering, see Failover Clustering Overview.

Do this step using Windows PowerShell

To create the failover cluster

1. Install the Failover Clustering feature on each node. For more information, see Install or Uninstall Roles, Role Services, or Features.
2. On one of the nodes, open Failover Cluster Manager, and then under **Management**, click **Validate Configuration** to validate the cluster. When you run the validation tests, make sure that you include the storage tests.
3. When you create the cluster, do either of the following:
   - If you have already created storage spaces in File and Storage Services, on the **Confirmation** page of the Create Cluster Wizard, verify that the **Add all eligible storage to the cluster** check box is selected.
   - If you have not yet created any storage spaces, make sure that you do not add the physical disks as cluster disks. Therefore, on the **Confirmation** page of the Create Cluster Wizard, clear the **Add all eligible storage to the cluster** check box.
4. If you had existing storage pools that you added to the failover cluster as eligible storage, after you create the failover cluster, verify that the storage appears in Failover Cluster Manager. To do this, follow these steps:
   a. In Failover Cluster Manager, expand *ClusterName*, expand **Storage**, and then click **Pools**.
   b. In the **Pools** pane, verify that any storage pools that you created in File and Storage Services are listed.

      Tip
      For detailed information, click a storage pool. In the results pane, you can view information on the **Summary**, **Virtual Disks**, or the **Physical Disks** tabs.

   c. To verify that virtual disks were added as cluster disks, under the **Storage** node, click **Disks**. To add a cluster disk to a CSV, continue to Step 6: (optional) Add a cluster disk to a CSV.

Note
To create clustered storage spaces through Failover Cluster Manager, continue to step 5.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example installs the Failover Clustering feature.

```
Install-WindowsFeature –Name Failover-Clustering –IncludeManagementTools
```

The following example runs all cluster validation tests on computers that are named *Server1* and *Server2*.

```
Test-Cluster –Node Server1, Server2
```

Note

The **Test-Cluster** cmdlet outputs the results to a log file in the current working directory, for example C:\Users*username*\AppData\Local\Temp.

The following example creates a failover cluster that is named *MyCluster* with nodes *Server1* and *Server2*, assigns the static IP address *192.168.1.12*, and adds all eligible storage to the failover cluster. If you have already created storage spaces in File and Storage Services, add all eligible storage.

```
New-Cluster –Name MyCluster –Node Server1, Server2 –StaticAddress 192.168.1.12
```

The following example creates the same failover cluster as in the previous example, but it does not add eligible storage to the failover cluster. If you have not yet created any storage spaces, make sure that you do not add eligible storage.

```
New-Cluster –Name MyCluster –Node Server1, Server2 –StaticAddress 192.168.1.12 –NoStorage
```

Step 5: (optional) Create clustered storage spaces in Failover Cluster Manager

If you have an existing failover cluster, you can create clustered storage spaces through Failover Cluster Manager.

Tip

For an existing failover cluster, you can also create storage spaces in File and Storage Services, and then add the storage space in Failover Cluster Manager. To do this, in Failover Cluster Manager, expand *ClusterName*, expand **Storage**, right-click **Pools**, and then click **Add Storage Pool**. In the **Add Pools to a Cluster** dialog box, select the check box next to the pool, and then click **OK**. To add virtual disks and associated volumes, right-click **Disks**, and then click **Add Disk**. In the **Add Disks to a Cluster** dialog box, select the check box next to each cluster disk that you want to add, and then click **OK**.

Important

To have Storage Spaces try to construct a mirror access across multiple supported JBOD enclosures, you must use Windows PowerShell to create the virtual disk by using the **IsEnclosure $True** parameter. For more information, see the second example in the Create a virtual disk section of this step.

Do this step using Windows PowerShell

To create clustered storage spaces in Failover Cluster Manager

1. In Failover Cluster Manager, expand *ClusterName*, and then expand **Storage**.
2. Right-click **Pools**, and then click **New Storage Pool**.

   The New Storage Pool Wizard opens.
3. On the **Before you begin** page, click **Next**.
4. On the **Specify a storage pool name and subsystem** page, enter a name and optional

description for the storage pool, select the group of available physical disks that you want to use, and then click **Next**.

Note
In the **Available to** column, verify that the cluster name is listed.

5. On the **Select physical disks for the storage pool** page, do the following, and then click **Next**:
   a. Select the check box next to each physical disk that you want to include in the storage pool. For a failover cluster, you must select at least three physical disks.
   b. If you want to designate one or more disks as hot spares, under **Allocation**, click the drop-down arrow, and then click **Hot Spare**.
6. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.
7. On the **View results** page, verify that all tasks are completed, select the **Create a virtual disk when this wizard closes** check box, and then click **Close**.

   The New Virtual Disk Wizard opens.
8. On the **Before you begin** page of the New Virtual Disk Wizard, click **Next**.
9. On the **Select the storage pool** page, click the desired storage pool, and then click **Next**.
10. On the **Specify the virtual disk name** page, enter a name and optional description for the virtual disk, and then click **Next**.
11. On the **Select the storage layout** page, click the desired layout, and then click **Next**.

    Note
    For a failover cluster, you can select either **Simple** or **Mirror**. For resiliency to recover from disk failure, you must select **Mirror**.
12. If you selected **Mirror** in the previous step, and you have at least five physical disks, the **Configure the resiliency settings page** appears. Select one of the following options:
    - **Two-way mirror**
    - **Three-way mirror**
13. On the **Specify the size of the virtual disk** page, do either of the following, and then click **Next**:
    - To specify a size, click **Specify size**, enter a value in the **Virtual disk size** box, and then select the units (**MB**, **GB** or **TB**).

      If you use a mirror storage layout, the virtual disk uses more free space than the size that you specify. To avoid a potential error where the size of the volume exceeds the storage pool free space, you can select the **Create the largest virtual disk possible, up to the specified size** check box.
    - To create a virtual disk that uses the maximum capacity of the storage pool, click **Maximum size**.
14. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.

15. On the **View results** page, verify that all tasks completed, make sure that the **Create a volume when this wizard closes** check box is selected, and then click **Close**.

    Important
    By default, the **Create a volume when this wizard closes** check box is selected. If you clear the check box, you must create a volume through File and Storage Services in Server Manager or through Windows PowerShell. You cannot create the volume through Failover Cluster Manager.

16. On the **Before you begin** page of the New Volume Wizard, click **Next**.
17. On the **Select the server and disk** page, do the following, and then click **Next**.
    a. In the **Server** area, click the failover cluster on which you want to provision the volume.
    b. In the **Disk** area, click the virtual disk on which you want to create the volume.
18. On the **Specify the size of the volume** page, enter a volume size, specify the units (**MB**, **GB**, or **TB**), and then click **Next**.
19. On the **Assign to a drive letter or folder** page, configure the desired option, and then click **Next**.
20. On the **Select file system settings** page, do the following, and then click **Next**.
    a. In the **File system** list, click **NTFS** or **ReFS**.

       Important
       If you want to add the virtual disk to a CSV, you must use NTFS.

       In Windows Server 2012, ReFS does not support CSVs, data deduplication, or File Server Resource Manager (FSRM).

    b. In the **Allocation unit size** list, either leave the setting at **Default** or set the allocation unit size.

       Note
       For more information about allocation unit size, see Default cluster size for NTFS, FAT, and exFAT.

    c. Optionally, in the **Volume label** box, enter a volume label name, for example **HR Data**.
21. On the **Confirm selections** page, verify that the settings are correct, and then click **Create**.
22. On the **View results** page, verify that all tasks completed, and then click **Close**.
23. To verify that the storage pool with associated virtual disks and volumes was added to the cluster, follow these steps:
    a. In Failover Cluster Manager, expand *ClusterName*, expand **Storage**, and then click **Pools**.
    b. In the **Pools** pane, verify that the storage pool is listed.

       Tip
       For more detailed information, click the storage pool. In the results pane, you

can view information on the **Summary**, **Virtual Disks**, or the **Physical Disks** tabs.

c. To verify that virtual disks were added as cluster disks, under the **Storage** node, click **Disks**. Make sure that cluster disks are listed, for example, *Cluster Disk 1*. Click a cluster disk, and then view the information in the results pane. The results pane shows information such as the virtual disk name, the health, the resiliency type, and volume information.

> **Tip**
> You can also view the volumes on the **Volumes** page of File and Storage Services in Server Manager. Or, in Windows Explorer, you can view the volumes on the node that owns the cluster disk.

24. To add a cluster disk to a CSV, continue to Step 6: (optional) Add a cluster disk to a CSV.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

1. Create a storage pool

The following example creates a storage pool *StoragePool1* that uses all available SAS disks. It also sets the fixed provisioning type as the default.

```
$physdisk = Get-PhysicalDisk –CanPool $true | Where-Object BusType –eq “SAS”

$stsubsys = Get-StorageSubsystem -FriendlyName "Storage Spaces*"

$stsubsys | New-StoragePool –FriendlyName StoragePool1 –PhysicalDisks $physdisk –
ProvisioningTypeDefault Fixed
```

The following example sequence of cmdlets creates a storage pool *StoragePool1* that uses three physical disks: *PhysicalDisk1*, *PhysicalDisk2*, and *PhysicalDisk3*. It also sets the fixed provisioning type as the default.

```
$physdisk = Get-StorageSubSystem -FriendlyName "Storage Spaces*" | Get-PhysicalDisk
PhysicalDisk1, PhysicalDisk2, PhysicalDisk3

New-StoragePool –FriendlyName StoragePool1 –StorageSubsystemFriendlyName "Storage
Spaces*" –PhysicalDisks $physdisk –ProvisioningTypeDefault Fixed
```

The following example adds the hot spare disk that you select to *StoragePool1*.

```
$hotsparedisk = Get-PhysicalDisk –CanPool $true | Out-GridView –PassThru

Add-PhysicalDisk –StoragePoolFriendlyName StoragePool1 –PhysicalDisks $hotsparedisk –
Usage HotSpare
```

2. Create a virtual disk

The following example creates a 10 GB virtual disk *VirtualDisk1* with the mirror resiliency type in the storage pool *StoragePool1*.

```
$newspace = New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName
VirtualDisk1 -ResiliencySettingName Mirror -Size (10GB)
```

The following example is the same as the previous example, but it adds the **IsEnclosureAware $True** parameter. Specify this parameter if you want Storage Spaces to try to construct a mirror across multiple supported JBOD enclosures.

```
$newspace = New-VirtualDisk -StoragePoolFriendlyName StoragePool1 -FriendlyName
VirtualDisk1 -ResiliencySettingName Mirror -Size (10GB) -IsEnclosureAware $True
```

3. Create a volume

The following example sequence of cmdlets initializes the virtual disk that you created in the previous step, creates a partition with the default partition style of GUID partition table (GPT), assigns a drive letter, and then formats the volume with the default of NTFS.

```
$newvol = $newspace | Get-Disk

Initialize-Disk -Number $newvol.Number

$partition = New-Partition -DiskNumber $newvol.Number -AssignDriveLetter -UseMaximumSize

Format-Volume -Partition $partition
```

4. Add the virtual disk to the cluster

The following example sequence of cmdlets adds the virtual disk *VirtualDisk1* to the failover cluster.

```
$space = Get-VirtualDisk -FriendlyName VirtualDisk1

Add-ClusterDisk $space
```

Step 6: (optional) Add a cluster disk to a CSV

This step shows how to add a cluster disk to a CSV. By doing this, the clustered storage is available across a uniform namespace on all nodes. By default, this namespace is C:\ClusterStorage, where C:\ represents the system disk.

Do this step using Windows PowerShell

To add a cluster disk to a CSV

1. In Failover Cluster Manager, expand *ClusterName*, expand **Storage**, and then click **Disks**.
2. Right-click a cluster disk, and then click **Add to Cluster Shared Volumes**.

   The **Assigned To** column changes to **Cluster Shared Volume**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example adds the cluster disk *Cluster Disk 1* to a CSV on the local failover cluster.

```
Add-ClusterSharedVolume –Name "Cluster Disk 1"
```

See also

- Storage Spaces Overview
- Deploy Storage Spaces on a Stand-Alone Server
- Storage Cmdlets in Windows PowerShell
- Use Cluster Shared Volumes in a Windows Server 2012 Failover Cluster
- Storage Spaces Frequently Asked Questions (FAQ)

How Storage Spaces Makes Use of Hot Spares

This topic describes how to designate disks as hot spares for use in automatically replacing failed disks in a storage pool, and how Storage Spaces monitors disk health and automatically recovers from failed disks.

How to designate hot spares

You can designate a disk as a hot spare when creating a storage pool or when adding disks to a pool by using the File and Storage Services role in Server Manager, or by using Windows PowerShell.

When creating a storage pool

To designate one or more disks as hot spares when creating a storage pool, use the following procedure:

Create a storage pool with hot spares

1. In **Server Manager**, click **File and Storage Services**, and then **Storage Pools**.
2. In the **STORAGE POOLS** tile, click **Tasks**, and then **New Storage Pool**.
3. On the **Select physical disks for the storage pool** page, select the disks you want to designate as hot spares, and then choose **Hot Spare** in the **Allocation** column.

When adding a physical disk

To designate one or more disks as hot spares when adding physical disks to a storage pool, use the following procedure:

Do this step using Windows PowerShell

Add a hot spare to a storage pool

1. In **Server Manager**, click **File and Storage Services**, and then **Storage Pools**.
2. Select the appropriate storage pool, and then in the **PHYSICAL DISKS** tile, click **Tasks**, and then **Add Physical Disk**.
3. Select the disks you want to designate as hot spares, and then choose **Hot Spare** in the **Allocation** column.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following example adds the physical disk named *PhysicalDisk4* to the *Company Data* storage pool and designates the disk as a hot spare.

```
Add-PhysicalDisk -StoragePoolFriendlyName "Company Data" -PhysicalDisks (Get-PhysicalDisk -FriendlyName PhysicalDisk4) -Usage HotSpare
```

How Storage Spaces monitors disk health

Write errors

If a disk reports a media error during a write operation, Storage Spaces performs the following actions:

- Marks the disk as Failed
- Activates the failure LED for the disk, if the disk is in a JBOD that supports SCSI Enclosure Services
- Replaces the failed disk with a hot spare and retries the write operation, if there is a hot spare available in the storage pool

Note

Modern hard disks are designed to silently handle a limited number failures without reporting an error to the operating system. If the disk exceeds the number of allowed failures during a write operation, it reports a media error, which is typically an indicator that the disk is approaching a catastrophic failure.

Read errors

If an error occurs while reading data from a disk, Storage Spaces performs the following actions:

1. Attempts to retrieve a copy of the data from another disk in the mirror or parity space. If this attempt succeeds, the read error is cleared.
2. Attempts to write a fresh copy of the data that was retrieved to the location that produced the read error. If this attempt fails, Storage Spaces produces a write error and takes action as described in the Write Errors section.

Disconnected disks

If a disk is unexpectedly disconnected or silently fails read or write operations and a hot spare is available in the storage pool Storage Spaces performs the following actions by default:

- Marks the disk as Failed
- Replaces the failed disk with a hot spare and retries the write operation
- If there is not a hot spare available, Storage Spaces waits for the disk to reconnect or for administrator to take action.

This behavior is configurable by using the `Set-StoragePool -RetireMissingPhysicalDisks <value>` cmdlet from a Windows PowerShell session. There are three settings:

Setting	Behavior
Auto	If a hot spare is available in the storage pool follow the **Enabled** logic. Otherwise, follow the **Disabled** logic.
Enabled	Mark the disk as Failed, replace it with a hot spare, and then retry the operation
Disabled	Wait for the disk to reconnect or for an administrator to take action (do not use any available hot spares)

Note

If an entire enclosure is disconnected, no action is taken even when set to **Auto** or **Enabled** because the system has likely lost more disks than can be replaced by hot spares.

What's New in Storage Spaces in Windows Server 2012 R2

This topic describes the features that were added to Storage Spaces in Windows Server 2012 R2, including storage tiers, write-back cache, and dual parity.

Technology description

Storage Spaces enables cost-effective, high availability, scalable, and flexible storage solutions for virtual or physical deployments. Storage Spaces delivers sophisticated storage virtualization capabilities that empower customers to use industry-standard storage for scalable multinode deployments or single server deployments. It is appropriate for a wide range of customers,

including enterprise and cloud hosting companies that use Windows Server for high availability and scalable storage that is cost effective. For more information, see Storage Spaces Overview.

New and changed functionality

The following table describes the changes in Storage Spaces functionality that are available in this release.

Feature/functionality	New or updated?	Description
Storage tiers	New	Automatically moves frequently accessed data to faster (solid-state drive) storage and infrequently accessed data to slower (hard disk) storage.
Write-back cache	New	Buffers small random writes to solid-state drives, reducing the latency of writes.
Parity space support for failover clusters	New	Enables you to create parity spaces on failover clusters.
Dual parity	New	Stores two copies of the parity information on a parity space, which helps protect you from two simultaneous physical disk failures and optimizes storage efficiency.
Automatically rebuild storage spaces from storage pool free space	New	Decreases how long it takes to rebuild a storage space after a physical disk failure by using spare capacity in the pool instead of a single hot spare.

Storage tiers

With **Storage tiers**, you can create virtual disks that are comprised of two tiers of storage:

- Solid-state drive tier for frequently accessed data
- Hard disk drive tier for less-frequently accessed data

Storage Spaces transparently moves data at a subfile level between the two tiers, based on how frequently data is accessed.

What value does this change add?

Storage tiers combine the best attributes of solid-state drives and hard disk drives. They increase the performance of the most used (“hot”) data by moving it to solid-state drives, without sacrificing the ability to store large quantities of data on inexpensive hard disks.

What works differently?

The following aspects work differently with storage tiers:

- To create a storage space with storage tiers, the storage pool must have a sufficient number of hard disks and solid-state drives to support the selected storage layout, and the hard disks must contain enough free space.
- When you create a storage space by using the New Virtual Disk Wizard or the **New-VirtualDisk** cmdlet, you can specify to create the virtual disk with storage tiers.
- To create a storage space with storage tiers, the virtual disk must use fixed provisioning, and the number of columns will be identical on both tiers (a four-column, two-way mirror with storage tiers would require eight solid-state drives and eight hard disk drives).
- Volumes that are created on virtual disks that use storage tiers should be the same size as the virtual disk.
- Administrators can *pin* (assign) files to the standard hard disk drive tier or to the faster solid-state drive tier by using the **Set-FileStorageTier** cmdlet. This ensures that the files are always accessed from the appropriate tier.
- There are the following new Windows PowerShell cmdlets for working with storage tiers, and the **New-VirtualDisk** cmdlet includes the **–StorageTiers** and **-StorageTierSizes** parameters.
    - Set-FileStorageTier
    - Get-FileStorageTier
    - Clear-FileStorageTier
    - New-StorageTier
    - Get-StorageTier
    - Set-StorageTier
    - Resize-StorageTier
    - Remove-StorageTier
    - Get-StorageTierSupportedSize

For information about using storage tiers, see: Step-by-step for Storage Spaces Tiering in Windows Server 2012 R2.

Write-back cache

Storage Spaces can use existing solid-state drives in the storage pool to create a write-back cache that is tolerant of power failures and that buffers small random writes to solid-state drives before later writing them to hard disk drives.

What value does this change add?

Small random writes often dominate common enterprise workloads, and they can impact the performance of other data transfers that are taking place. By using solid-state drives (which excel

at random access) for a write-back cache, Storage Spaces can reduce the latency of the random writes and also greatly reduce any impact on the performance of other data transfers.

What works differently?

The write-back cache is transparent to administrators and users, and it is created on all new virtual disks if there are a sufficient number of solid-state drives in the storage pool, as determined by the following requirements for the associated storage space:

- Simple spaces require one solid-state drive
- Two-way mirror spaces and single-parity spaces require two solid-state drives
- Three-way mirror spaces and dual parity spaces require three solid-state drives

The write-back cache works with all types of storage spaces, including storage spaces with storage tiers.

Newly created storage spaces automatically use a 1 GB write-back cache by default when the storage pool contains enough physical disks with **MediaType** set to **SSD** or **Usage** set to **Journal** to support the specified resiliency setting. If there aren't enough physical disks with these settings, the write-back cache size is set to 0, except for parity spaces, when it's set to 32 MB.

Parity space support for failover clusters

With Storage Spaces, you can create parity spaces on failover clusters.

What value does this change add?

Workloads that require greater capacity utilization than mirror spaces can now use parity spaces with one or two copies of parity information (single or dual parity) to maximize capacity and resiliency, while still offering the ability to fail over to another cluster node.

What works differently?

The parity storage layout is an option when you create virtual disks in Server Manager or with the **New-VirtualDisk** cmdlet.

Dual parity

Dual parity stores two copies of parity information in a parity space, helping protect you from two simultaneous disk failures.

What value does this change add?

Dual parity enables you to keep a high level of resiliency when you use a parity space with a large number of disks or any time you need to help protect against two simultaneous disk failures.

What works differently?

The dual parity resiliency type is an option when you create virtual disks in Server Manager or with the **New-VirtualDisk** cmdlet.

Automatically rebuild storage spaces from storage pool free space

Storage Spaces now includes the ability to automatically rebuild storage spaces from storage pool free space instead of using hot spares.

What value does this change add?

Rebuild times are accelerated because multiple disks in the pool can accept the data that was stored on the failed disk instead of waiting for a single hot spare to write all of the data. Additionally, hot spare drives are no longer needed, and storage pool free space can provide additional capacity and performance to the pool.

What works differently?

When a physical disk fails, instead of writing a copy of the data that was on the failed disk to a single hot spare, the data is copied to multiple physical disks in the pool so that the previous level of resiliency is achieved. Administrators no longer need to allocate physical disks as hot spares in the storage pool.

See also

The following table provides additional resources for evaluating Storage Spaces.

Content type	References
Product evaluation	• Storage Spaces: What's New in Windows Server 2012 R2(TechEd 2013 Video)
Community resources	• Windows Server 2012 R2 Storage Space Overview (YouTube video) • Windows Server Storage Spaces: What is it and why should I use it? • Step-by-step for Storage Spaces Tiering in Windows Server 2012 R2
Related technologies	• File and Storage Services Overview

Supporting Information Workers with Reliable File Services and Storage

One of the major goals of IT professionals is to support information workers who use client applications such as a word processor to work on files that are stored on file shares. Information workers require high levels of availability and performance from the servers that store their files. IT professionals also require those servers to be cost-effective, manageable, and built with security in mind.

A collection of features have been improved or added to Windows Server 2012 in response to the needs of information workers and the IT professionals who support them. At a high level, most of these features fall within the following three categories.

- Availability and Performance
- Cost Effectiveness
- Security and Manageability

Availability and Performance

In response to customer requests, we made availability a top priority as we were planning for Windows Server 2012.

Always Offline mode

In Windows 8 and Windows Server 2012, administrators can configure the experience for users of Offline Files to always work offline, even when they are connected through a high-speed network connection. Windows updates files in the Offline Files cache by synchronizing hourly in the background, by default. This addresses a customer concern with Offline Files where users would transition between the Online and Offline modes, depending on network availability and conditions, even when the Slow-Link mode (also known as the Slow Connection mode) was enabled and set to a 1 millisecond latency threshold.

For more information, see Folder Redirection, Offline Files, and Roaming User Profiles overview.

Cluster-Aware Updating for minimal downtime

Existing server updating tools are oriented toward standalone servers. Using these tools on clustered servers can be laborious and error-prone, and can result in unintended downtime during necessary update maintenance. The Cluster-Aware Updating (CAU) feature provides you with options for performing planned updates on server nodes in a failover cluster with minimal downtime and administrative overhead.

For more information, see Cluster-Aware Updating Overview.

SMB and NFS v3 Transparent Failover

Transparent Failover is a technology that allows file servers to fail over to another node in the cluster in a way that is completely transparent to applications that have files open on SMB or NFS v3 file shares on the failed server node. The delay for application file-requests might be only a few seconds, while the operations are transferred to the next node in the cluster. When combined with CAU, Transparent Failover takes care of both planned maintenance and unplanned failures with minimal impact.

For more information, see Server Message Block overview.

Chkdsk improvements

In the past, file system corruptions could result in unexpected downtime while Chkdsk scanned the volume and corrected these corruptions, with the downtime proportional to the number of files on the volume.

Windows Server 2012 introduces a new NTFS health model and approach to Chkdsk, reducing surprises, putting the user in control, and making computers more available and scalable. Detection of file system corruptions occurs online, while the correction of these errors happens in a small number of seconds, while offline. The design goal for Windows Server 2012 is to enable administrators to confidently deploy large volumes, including 64 TB volumes, without concern for availability or scalability.

Storage Spaces and Storage Pools

Storage Spaces enables you to virtualize storage by grouping cost-effective industry-standard disks into storage pools, and then create storage spaces (referred to as *virtual disks* in Server Manager) from the available capacity in the storage pools.

Storage Spaces helps you increase the availability of business-critical data in a cost-efficient way, regardless of whether you work in a large enterprise, a hosting company, or a small business. Storage Spaces reduces administration costs by decreasing the time administrators spend provisioning storage. It also simplifies administration tasks, enabling administrators who are not storage professionals to set up and manage resilient and high availability storage. Storage Spaces also saves hardware costs by enabling the use of industry-standard disks for resilient storage up to multiple petabytes in size, with high availability.

For more information, see File and Storage Services Overview.

SMB Directory Leasing

For branch office users connecting over WAN networks with high latencies, reducing the number of round-trips between a client and a server dramatically improves application response times. Directory Leasing is an enhancement to the existing SMB protocol that reduces the number of round-trips between client and server.

With Directory Leasing, client computers request a lease lock on directory information (which contains the file metadata), enabling Windows to cache the information for longer periods. Because many end-user applications access file metadata frequently, caching it for longer periods reduces the round-trips that a client has to wait for, improving the application response time.

For more information, see Server Message Block overview.

Cost Effectiveness

We recognize that cost effectiveness is one of your top priorities, closely rivaling availability.

Thin provisioning of storage

Windows Server 2012 identifies thinly-provisioned *virtual disks* (not to be confused with virtual hard disks or VHDs) created in a storage pool and enables transparent reclamation of capacity when files are deleted. This improves capacity utilization when used with the Storage Spaces subsystem or a third-party storage subsystem

Sophisticated applications can identify thinly provisioned disks by using Windows Management Instrumentation (WMI) or Windows PowerShell, and can invoke commands to reclaim capacity from deleted files as appropriate. Applications are notified via standard Windows events whenever the available capacity in the underlying storage pool for a thinly-provisioned virtual disk falls below the specified threshold, enabling the applications to take action.

To reclaim unused capacity on thinly provisioned virtual disks and perform other optimization tasks, the new Storage Optimizer runs automatically on a periodic basis. Together, these enhancements enable administrators to separate deployment requirements from purchasing decisions and to offload the burden of accurately predicting future storage capacity requirements.

For more information, see Thin Provisioning and Trim Storage Overview.

Data Deduplication

By using the new Data Deduplication role service to reduce the number of duplicate blocks of data in storage, you can store much more data in a given amount of storage capacity than was previously possible using Single Instance Storage (SIS) or NTFS file system compression. General purpose file servers can typically reduce storage capacity utilization by a 2:1 ratio (for example, files that previously used 1 TB would use 500 GB after deduplication). Servers that host virtualization data (such as VHD files) often reduce storage capacity utilization by a 20:1 ratio, which reduces 1 TB of data to 50 GB. Data deduplication has minimal CPU and memory overhead and maintains a high level of data integrity.

For more information, see Data Deduplication Overview.

BranchCache for Network Files version 2

The BranchCache for Network Files role service in Windows Server 2012 uses the same technology as the Data Deduplication role service to reduce duplicate data replicated over WAN connections. As a result, the BranchCache in Windows Server 2012 can further reduce WAN costs, and improve application response times for branch office users.

Enhanced support for low cost SATA disks

Windows Server 2012 features enhanced support for low-cost SATA disks. This enhancement decreases the possibility of metadata inconsistency due to unexpected power loss on commodity hard disks while enabling the disk to cache associated data as long as safely possible. This enables customers to more reliably utilize cost-effective SATA drives in enterprise deployments.

Access-denied assistance

In the past, when users are denied access to a file or folder, typically on a file share, users received little help with troubleshooting the issue. Often times, the solution was to determine the owner of the file share and then send the owner an email asking for access. In other instances, users would call the help desk or give up.

With Access-denied assistance, file server administrators can provide users with a customized access-denied message, including troubleshooting information and the ability to request assistance via email. Upon receiving the email requesting assistance, the designated owner of the folder can look at the details provided by Windows to help determine why the user was denied access and then take the appropriate action.

Access-denied assistance can reduce support costs by enabling server administrators to provide users with the means to help themselves solve access-denied issues and obtain assistance directly from the folder owners when they need additional help.

For more information, see Access-denied assistance.

Security and Manageability

Primary Computers support for Folder Redirection and Roaming User Profiles

In Windows Server 2012, you can designate a set of computers, known as primary computers, for each domain user, which enables you to control which computers use Folder Redirection, Roaming User Profiles, or both. For example, the administrator can choose to roam user data and settings between a user's desktop and laptop, and to not roam the information when that user logs on to any other computer, such as a conference room computer. Designating primary computers is a simple and powerful method to associate user data and settings with particular computers or devices, simplify administrator oversight, improve data security, and help protect user profiles from corruption.

For more information, see Folder Redirection, Offline Files, and Roaming User Profiles overview.

SMB Encryption

A security concern that comes up with data traversing untrusted networks is that it is prone to eavesdropping attacks. Existing solutions to the problem typically use IPsec, WAN accelerators, or other dedicated hardware solutions. However, these solutions are expensive to setup and maintain.

Windows Server 2012 includes encryption built-in to the SMB protocol. This allows end to end data protection from snooping attacks with no additional deployment costs. You have the flexibility to decide whether the entire server or only specific shares should be enabled for encryption. SMB encryption is also relevant to server application workloads if the application data is on a file server and traverses untrusted networks. With this feature, data security is maintained while it is on the wire.

For more information, see Server Message Block overview.

File Classification Infrastructure and Dynamic Access Control

Data governance, data retention, and regulatory compliance are responsible for a significant portion of IT budgets. Dynamic access control and the File Classification Infrastructure technologies included in Windows Server 2012 provide new ways to classify data, control access to data, and achieve regulatory compliance. It provides sophisticated authorization and auditing controls, along with data classification capabilities that enable organizations to apply information governance to unstructured data on their file servers. You now have greater control over file server data, a conditional expression ACL model, centralized authorization and auditing policy management, and business-driven data.

For more information, see What's New in File Server Resource Manager in Windows Server 2012.

Unified remote management of File and Storage Services in Server Manager

The File and Storage Services role in Server Manager enables you to remotely manage multiple file servers, including their role services and storage, all from a single window.

This functionality is built on top of a Windows Management Instrumentation (WMI) and Windows PowerShell infrastructure. As a result, Windows Server 2012 includes Windows PowerShell cmdlets for performing the majority of administration tasks for file and storage servers, simplifying the task of automating common tasks.

For more information, see File and Storage Services Overview.

Improved management of Folder Redirection, Offline Files, and Roaming User Profiles

Windows Server 2012 includes improved WMI interfaces for Folder Redirection, Offline Files, and Roaming User Profiles, making it easier for independent software vendors (ISVs) that deliver management suites to integrate managing, monitoring, and troubleshooting these technologies into their existing management tools.

See Also

File and Storage Services Overview

Scale-Out File Server for Application Data Overview

Thin Provisioning and Trim Storage Overview

This topic discusses just-in-time allocations (also known as thin provisioning) and the ability to reclaim storage that is no longer needed (also known as trim), including benefits and requirements and an overview of the technologies.

Did you mean…

- Storage Management Overview
- Storage Spaces Overview
- File and Storage Services Overview
- File Services Windows Server Protocols
- Windows Storage Server 2008 R2

Technology description

Sophisticated storage solutions offer just-in-time allocations (also known as thin provisioning) and the ability to reclaim storage that is no longer needed (also known as trim). Windows Server 2012 integrates with these sophisticated storage solutions to enable organizations to get the most out of their storage infrastructures.

Requirements

Thin provisioning and trim are enabled by default in Windows Server 2012, and have the following requirements:

- Storage infrastructure that complies with the certification that is required for Windows Server 2012
- Standards-compliant hardware for identification

Technical overview

To confidently deploy sophisticated storage solutions that support just-in-time allocation, you need to know that you can provision additional capacity as needed. Windows Server 2012 identifies thinly provisioned virtual disks, provides standardized notifications when use thresholds are crossed, and provides a platform that enables applications to give up storage when it is no longer needed, thereby ensuring maximal use.

The following capabilities are included:

- **Identification**. Windows Server 2012 uses a standardized method to detect and identify thinly provisioned virtual disks, thereby enabling additional capabilities delivered by the storage stack that is provided in the operating system and through storage management applications.

- **Notification**. When configured physical storage use thresholds are reached, Windows Server 2012 notifies the administrator through events, which enables the administrator to take appropriate action as soon as possible. These events can also be used for automated actions by sophisticated management applications, such as Microsoft System Center.
- **Optimization**. Windows Server 2012 provides a new API that lets applications return storage when it is no longer needed. NTFS issues trim notifications in real time when appropriate. Additionally, trim notifications are issued as part of storage consolidation (optimization), which is performed regularly on a scheduled basis.

See also

For additional related information, see the following resources.

Content type	References
Deployment	Understand and Troubleshoot Storage Spaces and Pools in Windows Server "8" Beta
Community resources	The File Services and Storage TechNet Forum | Storage Team at Microsoft File Cabinet Blog | Jose Barreto's Blog | Windows Server Information Experience on Twitter
Related technologies	Data Deduplication Overview | **High-Performance, Continuously Available File Share Storage for Server Applications Technical Preview** | iSCSI Target Block Storage Overview | Windows Offloaded Data Transfers Overview | Resilient File System Overview | Storage Spaces Overview | Supporting Information Workers with Reliable File Services and Storage | Thin Provisioning and Trim Storage Overview

Plan and Deploy Thin Provisioning

Before you use the thin provisioning feature in Windows Server 2012, you should determine if thin provisioning is appropriate for your environment, and if it is, plan for its use. As a system administrator, you should work closely with your storage administrator and other IT staff to ensure that you make accurate predictions about disk capacity needs, that you agree on threshold notification settings that give you enough time to react to disk space issues and to procure new disks if necessary, and that the system administrator and the storage administrator develop a coordinated response to events and alerts.

In this topic

- Determine if thin provisioning is appropriate for your environment
- Perform capacity forecasting
- Determine threshold notification and resource exhaustion settings
- Monitor and respond to events
- Consider space reclamation and potential performance impact

Note

To fully support the thin provisioning feature in Windows Server 2012, including space reclamation (also known as *file-level trim*), your storage array must comply with the certification requirements for Windows Server 2012.

Determine if thin provisioning is appropriate for your environment

Before you use thin provisioning, ensure that it is an appropriate storage solution for your environment. Thin provisioning is not designed for all storage usage models. Consider using thin provisioning if the following conditions are true:

- There is predictable storage consumption on a storage volume or a low-volatile usage pattern.
- The storage volume can tolerate a brief outage. For example, do not use thin provisioning if a mission critical cluster disk cannot tolerate any downtime or delays that are caused by thin provisioning. For example, there may be delays that are associated with temporary resource exhaustion or with space reclamation that occurs after large files are deleted.
- There is an adequate storage monitoring processes that detects when critical thresholds are crossed. You should have well-defined policies for monitoring and for response.
- You understand the time requirement to acquire new storage. If you do not understand the storage resource procurement process, thin-provisioned logical units (LUNs) may be exposed to the risk of permanent resource exhaustion.

Perform capacity forecasting

Predictable storage consumption is extremely important for thin-provisioned LUN deployment. A storage array that can support thin provisioning enables storage administrators to create large size LUNs with less physical disk space. A thin-provisioned LUN acts like a virtual storage device. Physical disk resources are allocated only when they are needed. This design reduces the amount of unused physical disk capacity and enhances the efficiency of storage use. However, this also introduces the risk that you will run out of physical disk space unless you carefully plan for and predict capacity requirements.

There are two basic forecasting approaches: quantitative and qualitative. For storage resource forecasting, it is a best practice to use both methods.

- **Quantitative forecasting**. Quantitative forecasting uses statistical analysis to analyze historical resource consumption. This provides a basis to determine future resource needs. Quantitative analysis provides a predicted growth rate for storage resource consumption, and it provides some insight into seasonal peaks that may occur, such as yearly application rollouts or semiannual migrations to data warehouses. The main drawback of using only quantitative forecasting is that history is not always an accurate predictor of future needs.
- **Qualitative forecasting**. Qualitative forecasting uses the expertise of the IT staff and other stakeholders to forecast needs based on their real-world knowledge and experience. We recommend that you work as a team, leveraging the insight of the system administrators, storage administrators, line-of-business (LOB) application owners, and senior managers who can help provide a more accurate forecast.

We recommend that you do all the following as part of capacity planning:

- Gather forecasting data from the IT department. When you do, communicate the value of forecasting to senior IT managers and to the business units.
- Develop accurate metrics. It is crucial to develop the appropriate forecasting metrics so that they are meaningful to all interested parties. Meaningful metrics will help to solidify the value of the initiative for stakeholders such as senior managers.
- Discuss capacity-on-demand with your storage vendors. A capacity-on-demand buying process is a low-risk way to build flexibility into a resource plan.

Use the storage resource forecasting results to create a plan to deploy thin-provisioned LUNs and to procure physical storage. If storage capacity prediction is not available or reliable, you should not use the thin provisioning feature.

Determine threshold notification and resource exhaustion settings

A storage administrator can configure threshold notification and resource exhausting settings.

Threshold notification settings

A thin provisioning threshold notification is an alert setting that warns system administrators and storage administrators when a predetermined capacity threshold is reached. There are two types of thresholds: LUN threshold and pool threshold. A storage administrator must configure these threshold values by using the vendor's storage management tools. When a threshold is reached, the Windows Server operating system logs an event to the System log in Event Viewer to alert the system administrator that the resource use or availability threshold was reached for the LUN or pool.

As a system administrator, you should work with your storage administrator to confirm the threshold notification settings. As a best practice, we recommend the following:

- When a system administrator requests thin-provisioned LUNs from the storage administrator, the system administrator should provide the desired threshold notification settings.
- If a system administrator does not provide the desired threshold notification settings, the storage administrator should determine a rational threshold setting and inform the system administrator.

Resource exhaustion settings

There are two kinds of resource exhaustion for a thin-provisioned LUN: temporary and permanent.

- **Temporary resource exhaustion setting**. Temporary resource exhaustion occurs when a LUN reaches the configured size limit, but automatic growth (or *auto-grow*) is enabled. When auto-grow or over-subscription is supported for a LUN, the storage array can use the temporary resource exhaustion setting to interact with the host server while it allocates more resources to the LUN. If temporary resource exhaustion occurs, the application that is using the disk is not interrupted or stopped.
- **Permanent resource exhaustion setting**. Permanent resource exhaustion represents a hard limit on the LUN's capacity consumption. This prevents the LUN from unexpectedly using capacity when other LUNs need it. If resource exhaustion is permanent, the storage array cannot allocate more resources to the LUN. The Windows partition manager sets the LUN to offline when the permanent resource exhaustion threshold limit is reached.

  Note

  It is important that the system administrator and the storage administer confirm the permanent resource exhaustion setting.

Monitor and respond to events

After you deploy thin-provisioned LUNs to application and data servers, and threshold notification and resource exhaustion settings are configured by the storage administrator. The system administrator and the storage administrator should monitor and respond to related events.

You can monitor thin provisioning threshold notifications from the System log in Event Viewer or from the vendor's storage management tools. The following table lists the threshold notification warning events.

Event ID	General Information
144	Threshold notification without additional information
145	Threshold notification without specific information
146	Threshold notification with used LUN capacity and available LUN capacity information
147	Threshold notification with used LUN capacity and available pool capacity information
148	Threshold notification with used pool capacity and available LUN capacity information
149	Threshold notification with used pool capacity

Event ID	General Information
	and available pool capacity information

When a threshold notification event is detected, the system administrator should immediately notify the storage administrator that an event has occurred and provide the event details. Working together, you should do the following:

1. Review the System log information, LUN use status, and pool consumption level.
2. Project LUN and pool capacity requirements.
3. Expand the LUN size or increase the pool size as needed.
4. In the storage management tools, reset the threshold notification setting.

How to respond to a permanent resource exhaustion event

If a permanent resource exhaustion threshold is reached, the following occurs:

- Event 150 is logged in the System log.
- The thin-provisioned LUN is set to offline, and is therefore unavailable.

If this occurs, you can use the following steps to recover the offline LUN:

1. The system administrator reviews the information that is associated with event ID 150 in the System log.
2. The system administrator informs the storage administrator about the issue.
3. The storage administrator either extends the LUN capacity or adds more physical disk resources to resolve the permanent resource exhaustion.
4. The storage administrator notifies the system administrator about the change in LUN status.
5. The system administrator uses File and Storage Services in Server Manager or the Disk Management tool to rescan the disk, bring the disk online, and extend the disk size.

Consider space reclamation and potential performance impact

In Windows Server 2012, when a large file is deleted, the file system performs space reclamation according to the storage array setting. Large file deletions can potentially affect the performance of the regular I/O. To protect against this, you can do either of the following:

- Configure a registry setting to disable the file-delete notification that triggers real-time space reclamation. Realize that if you configure this setting, this disables the file-delete notification for all LUNs that are assigned to the host. If file-delete notification is disabled, you can still use the Defragment and Optimize Drives tool to perform space reclamation on-demand or on a scheduled basis.
- If you want to use real-time space reclamation, but you have a thin-provisioned LUN that is used by a highly volatile application where large file deletions are common, we recommend that you convert the thin-provisioned LUN to a thick (or full) provisioned LUN. A storage administrator can do this by using the vendor's storage management tools.

To disable the file-delete notification in the registry, follow these steps:

Caution

Incorrectly editing the registry may severely damage your system. Before you make changes to the registry, back up valued data on the computer.

1. Start Registry Editor.
2. Locate the following registry subkey:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem
3. Double-click **DisableDeleteNotification**.
4. In the **Value data** box, enter a value of **1**, and then click **OK**.

If file-delete notification is disabled, you can use the Defragment and Optimize Drives tool to reclaim free storage space. You can initiate a manual request or configure optimization to occur automatically on a daily, weekly, or monthly basis. If you initiate a manual request, you can try to target periods of low disk I/O, for example an evening or a weekend.

To start the Defragment and Optimize Drives tool, on the Server Manager menu bar, click **Tools**, and then click **Defragment and Optimize Drives**. If the disk drive is a thin-provisioned LUN, the media type of the disk drive appears as “Thin Provisioned Drive.”

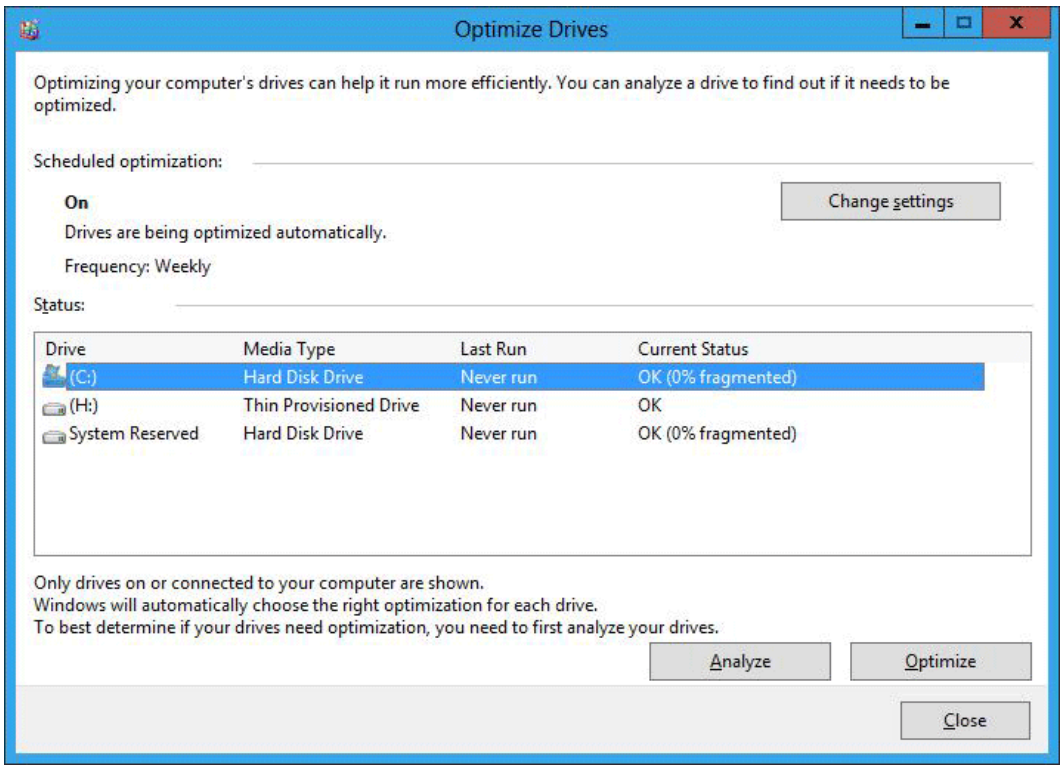

Figure 1: Defragment and Optimize Drives tool

If you have concerns or questions regarding the status of space reclamation, we recommend that you contact your storage administrator or your storage array vendor for more detailed information or for troubleshooting recommendations.

See also

For additional information, see the following resources.

Content type	References
Product evaluation	Thin Provisioning and Trim Storage Overview | Thin Provisioning in Windows Server 2012
Hardware certification requirements	Windows 8 Hardware Certification Requirements **Note** In the "Windows 8 Device Requirements" document that is available on the website, see the "Device.Storage.Hd.ThinProvisioning" section.

Windows Offloaded Data Transfers Overview

This topic provides an overview of Windows Offloaded Data Transfer (ODX, also known as copy offload) in Windows Server 2012. ODX enables direct data transfers within or between compatible storage devices without transferring the data through the host computer.

Did you mean…

- Storage Management Overview
- Deploy Windows Offloaded Data Transfers

Feature description

Windows Offloaded Data Transfer (ODX) functionality in Windows Server 2012 maximizes an enterprise's investment in intelligent storage arrays by enabling the arrays to directly transfer data within or between compatible storage devices, bypassing the host computer.

By offloading the file transfer to the storage array, ODX minimizes latencies, maximizes array throughput, and reduces resource usage such as CPU and network consumption on the host computer. Windows Server 2012 offloads file transfers transparently and automatically when you move or copy files, irrespective of whether you drag-and-drop files through File Explorer or use command-line file copy commands.

Practical applications

Some of the applications of ODX include:

- Rapidly import and export Hyper-V virtual machines that are stored on an ODX-capable storage array and accessed via iSCSI, Fibre Channel, or SMB file shares
- Transfer large files such as database files or video files with increased speed and decreased CPU and network resource consumption on the host server

Important functionality

In traditional host-based file transfers, the data to be transferred must be:

1. Read from the storage through the source server
2. Transferred across the network to the destination server
3. Written back to the storage through the destination server

To eliminate this inefficiency, ODX uses a token-based mechanism for reading and writing data within or between intelligent storage arrays. Instead of routing the data through the host, a small token is copied between the source server and destination server. The token serves as a point-in-time representation of the data. As an example, when you copy a file or migrate a virtual machine between storage locations (within or between storage arrays), a token representing the virtual machine file is copied, thereby removing the need to copy the underlying data through the servers.

The following figure explains the steps that are involved with a token-based copy operation.

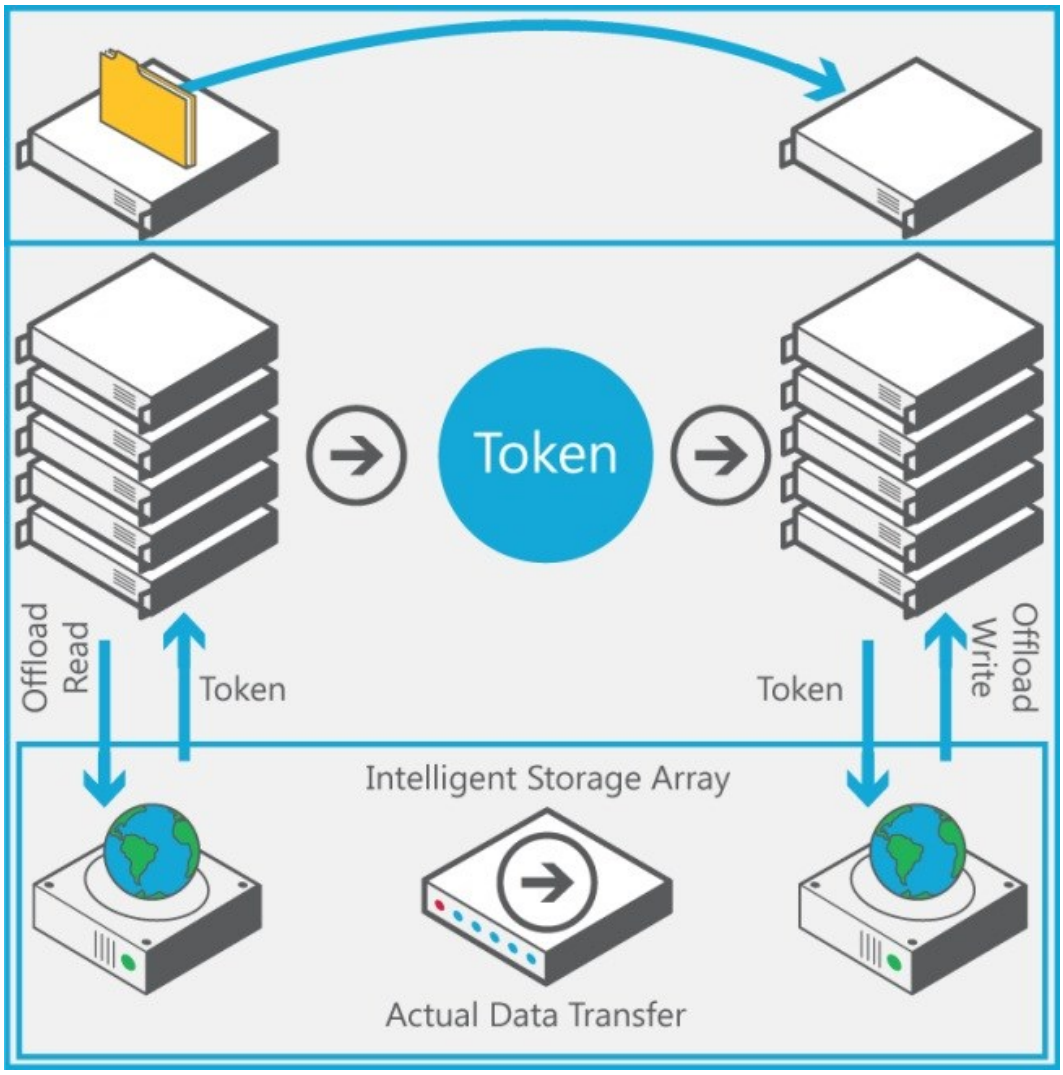

Figure 1 Token-based copy operation

This procedure is described in the following steps:

1. A user copies or moves a file by using Windows Explorer, a command line interface, or as part of a virtual machine migration.
2. Windows Server 2012 automatically translates this transfer request into an ODX (if supported by the storage array), and it receives a token that represents the data.
3. The token is copied between the source server and destination server.
4. The token is delivered to the storage array.
5. The storage array internally performs the copy or move and provides status information to the user.

Notes

- In the event of an MPIO path failover, Windows retries the ODX transfer. If this fails, Windows initiates a cluster failover (when part of a failover cluster).
- In the event of a cluster failover, if the application is cluster aware Windows resumes the ODX transfer after the failover.
- If Windows cannot resume or restart an ODX transfer after an MPIO path or cluster failover, Windows issues a LUN reset to the storage device, ending all outstanding operations on the LUN. It then returns an IO failure back to the application.

Hardware requirements

To use ODX, your storage arrays must meet the following requirements:

- Must be certified compatible with Windows Offloaded Data Transfer (ODX) on Windows Server 2012
- To support ODX between storage arrays, the copy manager for the storage arrays must support cross-storage array ODX, and the storage arrays must be from the same vendor
- Must be connected by using one of the following protocols:
  - iSCSI
  - Fibre Channel
  - Fibre Channel over Ethernet
  - Serial Attached SCSI (SAS)
- Must use one of the following configurations:
  - One server with one storage array
  - One server with two storage arrays
  - Two servers with one storage array
  - Two servers with two storage arrays

Software requirements

To use ODX, your environment must support the following:

- The computer initiating the data transfer must be running Windows 8 or Windows Server 2012.
- File system filter drivers such as antivirus and encryption programs need to opt-in to ODX. ODX is not supported by the following file system filter drivers:
  - Data Deduplication
  - BitLocker Drive Encryption
- Files must be on an unencrypted basic partition. Storage Spaces and dynamic volumes are not supported.
- Files must be on a volume formatted using NTFS. ReFS and FAT are not supported. Files can be directly transferred to or from this volume, or from one of the following containers:
  - A Virtual Hard Disk (VHD) that uses the VHD or VHDX formats
  - A file share that uses the SMB protocol

- The files must be 256 KB or larger – smaller files are transferred using a traditional (non-ODX) file transfer.
- The application that performs the data transfer must be written to support ODX. The following currently support ODX:
  - Hyper-V management operations that transfer large amounts of data at a time, such as creating a fixed size virtual hard disk (VHD), merging snapshot or converting virtual hard disks.
  - File Explorer
  - Copy commands in Windows PowerShell
  - Copy commands in Windows command prompt (including Robocopy)
- Files should not be highly fragmented. Transfers of highly fragmented files will have reduced performance.

Hyper-V Requirements

To use ODX with virtual machines hosted by Hyper-V, the virtual machines need to access storage from an ODX-capable storage array. You can achieve this by using any of the following approaches.

- Store the VHD on an ODX-capable iSCSI LUN
- Assign ODX-capable iSCSI LUNs to the virtual machine's iSCSI initiator
- Assign ODX-capable Fibre Channel LUNs to the virtual machine's virtual Fibre Channel adapter
- Connect the host or virtual machine to an SMB file share on another computer that is hosted on an ODX-capable storage array

See also

For more information, see the following resources.

Content type	References
Deployment	Deploy Windows Offloaded Data Transfers
API Reference	FSCTL_OFFLOAD_READ control code | FSCTL_OFFLOAD_WRITE control code | T10 SCSI Primary Command Spec (SPC4r36) | Windows 8 Devices Requirements

Deploy Windows Offloaded Data Transfers

This topic discusses how to deploy Windows Offloaded Data Transfer (ODX) in Windows Server 2012 to directly transfer data within or between compatible storage devices, bypassing the host computer. This topic also discusses hardware requirements, software requirements, and how to verify the performance of ODX after you implement it.

Although no action is required to use ODX with a compatible storage array and compatible applications, there are a number of tasks that you should perform to confirm that your environment is compatible with ODX and to verify that you are receiving the performance benefits of ODX, as discussed in this topic.

- **In this document**

- Prerequisites
- Step 1: Gather storage array information
- Step 2: Validate file system filter drivers
- Step 3: Establish a performance baseline
- Step 4: Test ODX performance
- Appendix: Deployment checklist

Prerequisites

To use ODX, your environment must meet the following hardware and software requirements.

Hardware requirements

To use ODX, your storage arrays must meet the following requirements:

- Be certified as compatible with Windows Offloaded Data Transfer (ODX) on Windows Server 2012
- Support cross-storage array ODX. To support ODX between storage arrays, the copy manager for the storage arrays must support cross-storage array ODX, and the storage arrays must be from the same vendor
- Be connected by using one of the following protocols:
  - iSCSI
  - Fibre Channel
  - Fibre Channel over Ethernet
  - Serial Attached SCSI (SAS)
- Use one of the following configurations:
  - One server with one storage array
  - One server with two storage arrays
  - Two servers with one storage array
  - Two servers with two storage arrays

Software requirements

To use ODX, your environment must support the following:

- The computer initiating the data transfer must be running Windows 8 or Windows Server 2012.
- File system filter drivers such as antivirus and encryption programs need to opt-in to ODX. ODX is not supported by the following file system filter drivers:
  - Data Deduplication
  - BitLocker Drive Encryption
- Files must be on an unencrypted basic partition. Storage Spaces and dynamic volumes are not supported.
- Files must be on a volume that is formatted by using NTFS. ReFS and FAT are not supported. Files can be directly transferred to or from this volume, or from one of the following containers:
  - A virtual hard disk (VHD) that uses the .vhd or .vhdx format
  - A file share that uses the SMB protocol
- The files must be 256 KB or larger. Smaller files are transferred by using a traditional (non-ODX) file transfer operation.
- The application that performs the data transfer must be written to support ODX. The following currently support ODX:
  - Hyper-V management operations that transfer large amounts of data at one time, such as creating a fixed size VHD, merging a snapshot, or converting VHDs.
  - File Explorer
  - Copy commands in Windows PowerShell
  - Copy commands in the Windows command prompt (including Robocopy)
- Files should not be highly fragmented. Transfers of highly fragmented files will have reduced performance.

Hyper-V Requirements

To use ODX with virtual machines on a server running Hyper-V, the virtual machines need to access storage from an ODX-capable storage array. You can achieve this by using any of the following approaches.

- Store the VHD on an ODX-capable iSCSI LUN.
- Assign ODX-capable iSCSI LUNs to the virtual machine's iSCSI initiator.
- Assign ODX-capable Fibre Channel LUNs to the virtual machine's virtual Fibre Channel adapter.
- Connect the host or virtual machine to an SMB file share on another computer that is hosted on an ODX-capable storage array.

Step 1: Gather storage array information

Before deploying ODX, gather the following information about the copy manager (operating system) of the storage array:

- What is the name and version of the copy manager?
- Does the copy manager support ODX?
- Does the copy manager support an ODX operation across multiple storage arrays from the same vendor?
- What is the default inactive timer value? This specifies how long the copy manager waits to invalidate the idle token after the timer expiration.
- What is the maximum token capacity of the copy manager?
- What is the optimal transfer size? This tells Windows how to send Read and Write commands that are optimally sized for the storage array.

Step 2: Validate file system filter drivers

To use ODX, validate all the file system filter drivers on all servers that are hosting the storage support ODX.

To validate the opt-in status of file system filter drivers, use the following procedure:

To validate file system filter driver opt-in status

1. On each server on which you want to use ODX, list all of the file system filter drivers that are attached to the volume on which you want to enable ODX. To do so, open a Windows PowerShell session as an administrator, and then type the following command, where *<volume>* is the drive letter of the volume:

   ```
   Fltmc instances -v <volume>
   ```

2. For each filter driver listed, query the registry to determine whether the filter driver has opted-in to ODX support. To do so, type the following command for each filter previously listed, and replace *<FilterName>* with the name of the filter.

   ```
   Get-ItemProperty
   hklm:\system\currentcontrolset\services\<FilterName> -Name
   "SupportedFeatures"
   ```

3. If the **SupportedFeatures** registry value equals **3**, the filter driver supports ODX. If the value is not **3**, contact the file system filter driver vendor for an ODX-compatible version.

Step 3: Establish a performance baseline

To establish a performance baseline, use the following procedures to disable ODX on the server and create a System Performance Report during a representative data transfer.

Disable ODX

To establish a baseline of non-offloaded data transfer performance, first disable ODX on the server by following these steps:

To disable ODX on the server

1. Open a Windows PowerShell session as an administrator.
2. Check whether ODX is currently enabled (it is by default) by verifying that the **FilterSupportedFeaturesMode** value in the registry equals **0**. To do so, type the following command:

   ```
   Get-ItemProperty
   hklm:\system\currentcontrolset\control\filesystem -Name
   "FilterSupportedFeaturesMode"
   ```

3. Disable ODX support. To do so, type the following command:

   ```
   Set-ItemProperty
   hklm:\system\currentcontrolset\control\filesystem -Name
   "FilterSupportedFeaturesMode" -Value 1
   ```

Create a System Performance Report during a data transfer

To record the baseline performance of data transfers, use Performance Monitor to record system performance during a represenative data transfer. To do so, follow these steps:

To create a System Performance Report

1. In Server Manager, on the **Tools** menu, click **Performance Monitor**.
2. Initiate a large data transfer that is representative of the workload you want to accelerate and that is within or between the storage arrays that support ODX.
3. Start the System Performance data collector set. To do so, expand **Data Collector Sets**, expand **System**, right-click **System Performance**, and then click **Start**. Performance Monitor will collect data for 60 seconds.
4. Expand **Reports**, expand **System**, expand **System Performance**, and then click the most recent report.
5. Review the System Performance Report, and take note of the following counters:
   - CPU Utilization (in the Resource Overview section)
   - Network Utilization (in the Resource Overview section)
   - Disk Bytes/sec (in the Disk section, under Physical Disk)

Step 4: Test ODX performance

After you establish a baseline of system performance during traditional data transfers, use the following procedures to enable ODX on the server and test offloaded data transfers:

Enable ODX

To enable ODX on the server, follow these steps:

To enable ODX

1. Open a Windows PowerShell session as an administrator.
2. Type the following command:

```
Set-ItemProperty
hklm:\system\currentcontrolset\control\filesystem -Name
"FilterSupportedFeaturesMode" -Value 0
```

Verify ODX performance

After ODX is enabled, create a System Performance Report during a large offloaded data transfer (see the Create a System Performance Report during a data transfer section earlier in this topic for the procedure).

When you evaluate the performance of offloaded data transfers, you should see the following differences from the baseline that you created when ODX was disabled:

- CPU utilization should be much lower (only slightly higher than prior to the data transfer). This shows that the server did not need to manage the data transfer.
- Network utilization should be much lower (only slightly higher than prior to the data transfer). This shows that the data transfer bypassed the server.
- Disk Bytes/sec should be much higher. This reflects increased performance from direct transfers within an array or within the SAN.

After you verify ODX performance, periodically create another System Performance Report during offloaded data transfers to confirm that ODX is still operating as expected. If any performance degradation is detected, contact Microsoft Customer Support and the storage array vendor.

Tip

You can use the following command in a Windows PowerShell session to display a list of storage subsystems that support ODX and use a storage management provider. This command does not display storage subsystems that use the Storage Management Initiative Specification (SMI-S) protocol.

```
Get-OffloadDataTransferSetting | Get-StorageSubSystem
```

Appendix: Deployment checklist

Use the following checklist to confirm that you completed all the steps for the deployment.

	Deploying Windows Offloaded Data Transfers Checklist

	Check the Windows Offloaded Data Transfers prerequisites.
	Gather storage array information.
	Validate the file system filter drivers.
	Establish a performance baseline.
	Test the ODX performance.

See Also

T10 SCSI Primary Command Spec (SPC4r36)

Windows 8 Devices Requirements

Work Folders Overview

This topic discusses Work Folders, which is a role service for file servers running Windows Server that provides a consistent way for users to access their work files from their PCs and devices.

Did you mean…

- Folder Redirection, Offline Files, and Roaming User Profiles overview
- BranchCache Overview

Role description

With Work Folders users can store and access work files on personal computers and devices, often referred to as bring-your-own device (BYOD), in addition to corporate PCs. Users gain a convenient location to store work files, and they can access them from anywhere. Organizations maintain control over corporate data by storing the files on centrally managed file servers, and optionally specifying user device policies such as encryption and lock-screen passwords.

Work Folders can be deployed with existing deployments of Folder Redirection, Offline Files, and home folders. Work Folders stores user files in a folder on the server called a *sync share*. You can specify a folder that already contains user data, which enables you to adopt Work Folders without migrating servers and data or immediately phasing out your existing solution.

Practical applications

Administrators can use Work Folders to provide users with access to their work files while keeping centralized storage and control over the organization's data. Some specific applications for Work Folders include:

- Provide a single point of access to work files from a user's work and personal computers and devices
- Access work files while offline, and then sync with the central file server when the PC or device next has Internet or intranet connectivity
- Deploy with existing deployments of Folder Redirection, Offline Files, and home folders
- Use existing file server management technologies, such as file classification and folder quotas, to manage user data
- Specify security policies to instruct user's PCs and devices to encrypt Work Folders and use a lock screen password
- Use Failover Clustering with Work Folders to provide a high-availability solution

Important functionality

Work Folders includes the following functionality.

Functionality	Availability	Description
Work Folders role service in Server Manager	Windows Server 2012 R2	File and Storage Services provides a way to set up sync shares (folders that store user's work files), monitors Work Folders, and manages sync shares and user access
Work Folders cmdlets	Windows Server 2012 R2	A Windows PowerShell module that contains comprehensive cmdlets for managing Work Folders servers
Work Folders integration with Windows 8.1	Windows 8.1, Windows RT 8.1	Work Folders provides the following functionality in Windows computers: • A Control Panel item that sets up and monitors Work Folders • File Explorer integration that enables easy access to files in Work Folders • A sync engine that transfers files to and from a central file server while maximizing battery life and system performance
Work Folders app for devices	Not yet announced	An app that allows popular

Functionality	Availability	Description
		devices to access files in Work Folders

Software requirements

Work Folders has the following software requirements for file servers and your network infrastructure:

- A server running Windows Server 2012 R2 for hosting sync shares with user files
- A volume formatted with the NTFS file system for storing user files

To enable users to sync across the Internet, there are additional requirements:

- A server certificate for each file server that will host Work Folders (plus a server certificate for your reverse proxy server). These certificates should be from a certification authority (CA) that is trusted by your users—ideally a public CA
- The ability to make a server accessible from the Internet by creating publishing rules in your organization's reverse proxy or network gateway
- A publicly registered domain name and the ability to create additional public DNS records for the domain
- (Optional) An Active Directory Domain Services forest with schema extensions in Windows Server 2012 R2 to support automatically referring PCs and devices to the correct file server when using multiple file servers
- (Optional) Active Directory Federation Services (AD FS) infrastructure when using AD FS authentication

Work Folders has the following software requirements for client computers:

- Computers must be running one of the following operating systems:
  - Windows 8.1
  - Windows RT 8.1
- Enough free space on a local, NTFS-formatted drive to store all the user's files in Work Folders, plus an additional 6 GB of free space if Work Folders is located on the system drive, as it is by default. Work Folders uses the following location by default:
  %USERPROFILE%\Work Folders

  However, users can change the location during setup (microSD cards and USB drives formatted with the NTFS file system are supported locations, though sync will stop if the drives are removed).

  The maximum size for individual files is 10 GB by default. There is no per-user storage limit, although administrators can use the quotas functionality of File Server Resource Manager to implement quotas.
- Work Folders doesn't support rolling back the virtual machine state of client virtual machines. Instead perform backup and restore operations from inside the client virtual machine by using System Image Backup or another backup app.

Work Folders compared to other sync technologies

The following table discusses how various Microsoft sync technologies are positioned and when to use each.

	Work Folders	Offline Files	SkyDrive Pro	SkyDrive
Technology summary	Syncs files that are stored on a file server with PCs and devices	Syncs files that are stored on a file server with PCs that have access to the corporate network (can be replaced by Work Files)	Syncs files that are stored in Office 365 or in SharePoint with PCs and devices inside or outside a corporate network, and provides document collaboration functionality	Syncs personal files that are stored in SkyDrive with PCs, Mac computers, and devices
Intended to provide user access to work files	Yes	Yes	Yes	No
Cloud service	None	None	Office 365	Microsoft SkyDrive
Internal network servers	File servers running Windows Server 2012 R2	File servers	SharePoint server (optional)	None
Supported clients	PCs and devices* inside or outside a corporate network	PCs in a corporate network or connected through DirectAccess, VPNs, or other remote access technologies	PCs, iOS, Windows Phone	PCs, Mac computers, Windows Phone, iOS, Android

*Work Folders apps not yet announced.

Note

In addition to the sync technologies listed in the previous table, Microsoft offers other replication technologies, including DFS Replication, which is designed for server-to-server replication, and BranchCache, which is designed as a branch office WAN acceleration technology. For more information, see DFS Namespaces and DFS Replication Overview and **BranchCache Overview**

Server Manager information

Work Folders is part of the File and Storage Services role. You can install Work Folders by using the Add Roles and Features Wizard or the `Install-WindowsFeature` cmdlet. Both methods accomplish the following:

- Adds the **Work Folders** page to **File and Storage Services** in Server Manager
- Installs the Windows Sync Shares service, which is used by Windows Server to host sync shares
- Installs the SyncShare Windows PowerShell module to manage Work Folders on the server

Interoperability with Windows Azure virtual machines

You can run this Windows Server role service on a virtual machine in Windows Azure. This scenario has been tested with Windows Server 2012 R2.

To learn about how to get started with Windows Azure virtual machines, visit the Windows Azure web site.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	• Introducing Work Folders on Windows Server 2012 R2 (blog post) • Introduction to Work Folders (Channel 9 Video) • Work Folders Test Lab Deployment (blog post)
Deployment	• Designing a Work Folders Implementation • Deploying Work Folders • Performance Considerations for Work Folders Deployments
Operations	• Work Folders Certificate Management (blog

Content type	References
	post) • Monitoring Windows Server 2012 R2 Work Folders Deployments (blog post) • SyncShare (Work Folders) Cmdlets in Windows PowerShell • Storage and File Services PowerShell Cmdlets Quick Reference Card For Windows Server 2012 R2 Preview Edition
Troubleshooting	• Windows Server 2012 R2 – Resolving Port Conflict with IIS Websites and Work Folders (blog post)
Community resources	• File Services and Storage Forum • The Storage Team at Microsoft - File Cabinet Blog • Ask the Directory Services Team Blog
Related technologies	• File and Storage Services Overview • File Server Resource Manager Overview • Folder Redirection, Offline Files, and Roaming User Profiles overview • BranchCache Overview • DFS Namespaces and DFS Replication Overview

Designing a Work Folders Implementation

This topic explains the design process for a Work Folders implementation, and assumes that you have the following background:

- Have a basic understanding of Work Folders (as described in Work Folders Overview)
- Have a basic understanding of Active Directory Domain Services (AD DS) concepts
- Have a basic understanding of Windows file sharing and related technologies
- Have a basic understanding of SSL certificate usage
- Have a basic understanding of enabling web access to internal resources via a web reverse proxy

The following sections will help you design your Work Folders implementation. Deploying Work Folders is discussed in the next topic, Deploying Work Folders.

Software requirements

Work Folders has the following software requirements for file servers and your network infrastructure:

- A server running Windows Server 2012 R2 for hosting sync shares with user files
- A volume formatted with the NTFS file system for storing user files

To enable users to sync across the Internet, there are additional requirements:

- A server certificate for each file server that will host Work Folders (plus a server certificate for your reverse proxy server). These certificates should be from a certification authority (CA) that is trusted by your users—ideally a public CA
- The ability to make a server accessible from the Internet by creating publishing rules in your organization's reverse proxy or network gateway
- A publicly registered domain name and the ability to create additional public DNS records for the domain
- (Optional) An Active Directory Domain Services forest with schema extensions in Windows Server 2012 R2 to support automatically referring PCs and devices to the correct file server when using multiple file servers
- (Optional) Active Directory Federation Services (AD FS) infrastructure when using AD FS authentication

Work Folders has the following software requirements for client computers:

- Computers must be running one of the following operating systems:
  - Windows 8.1
  - Windows RT 8.1
- Enough free space on a local, NTFS-formatted drive to store all the user's files in Work Folders, plus an additional 6 GB of free space if Work Folders is located on the system drive, as it is by default. Work Folders uses the following location by default: **%USERPROFILE%\Work Folders**

  However, users can change the location during setup (microSD cards and USB drives formatted with the NTFS file system are supported locations, though sync will stop if the drives are removed).

  The maximum size for individual files is 10 GB by default. There is no per-user storage limit, although administrators can use the quotas functionality of File Server Resource Manager to implement quotas.
- Work Folders doesn't support rolling back the virtual machine state of client virtual machines. Instead perform backup and restore operations from inside the client virtual machine by using System Image Backup or another backup app.

Note

Make sure to install the Windows 8.1 and Windows Server 2012 R2 General Availability update rollup on all Work Folders servers and any client computers running Windows 8.1. For more information, see article 2883200 in the Microsoft Knowledge Base.

Deployment scenarios

Work Folders can be implemented on any number of file servers within a customer environment. This allows Work Folders implementations to scale based on customer needs and can result in highly individualized deployments. However, most deployments will fall into one of the following three basic scenarios.

Single-Site Deployment

In a single-site deployment, file servers are hosted within a central site in the customer infrastructure. This deployment type is seen most often in customers with a highly centralized infrastructure or with smaller branch offices that do not maintain local file servers. This deployment model can be easier for IT staff to administer, since all server assets are local, and internet ingress/egress is likely centralized at this location as well. However, this deployment model also relies on good WAN connectivity between the central site and any branch offices, and users in branch offices are vulnerable to an interruption of service due to network conditions.

Multiple-Site Deployment

In a multiple-site deployment, file servers are hosted in multiple locations within the customer's infrastructure. This could mean multiple datacenters or it could mean that branch offices maintain individual file servers. This deployment type is seen most often in larger customer environments or in customers that have several larger branch offices that maintain local server assets. This deployment model is more complex for IT personnel to administer, and relies on careful coordination of data storage and maintenance of Active Directory Domain Services (AD DS) to ensure that users are using the correct sync server for Work Folders.

Hosted Deployment

In a hosted deployment, sync servers are deployed in an IAAS (Infrastructure-as-a-Service) solution such as Windows Azure VM. This deployment method has the advantage of making the availability of file servers less dependent on WAN connectivity within a customer's business. If a device is able to connect to the Internet, it can get to its sync server. However, the servers deployed in the hosted environment still need to be able to reach the organization's Active Directory domain to authenticate users, and the customer trades infrastructure requirements on premise for additional complexity in maintaining that connection.

Deployment technologies

Work Folders deployments consist of a number of technologies that work together to provide service to devices on both the internal and external networks. Before designing a Work Folders deployment, customers should be familiar with the requirements of each of the following technologies.

Active Directory Domain Services

AD DS provides two important services in a Work Folders deployment. First, as the back-end for Windows authentication, AD DS provides the security and authentication services that are used to grant access to user data. If a domain controller cannot be reached, a file server will be unable to authenticate an incoming request and the device will not be able to access any data stored in that file server's sync share.

Second, AD DS (with the Windows Server 2012 R2 schema update) maintains the msDS-SyncServerURL attribute on each user, which is used to automatically direct users to the appropriate sync server.

File Servers

File servers running Windows Server 2012 R2 host the Work Folders role service, and host the sync shares that store users' Work Folders data. File servers can also host data stored by other technologies operating on the internal network (such as file shares), and can be clustered to provide fault tolerance for user data.

Group Policy

You can use Group Policy to specify a Work Folders server to domain-joined PCs. This simplifies Work Folders setup a little bit– users would otherwise need to enter their work email address to lookup the settings (assuming that Work Folders is set up properly), or enter the Work Folders URL that you explicitly provided them via email or another means of communication.

You can also use Group Policy to forcibly set up Work Folders on a per-user or per-computer basis, though doing so causes Work Folders to sync on every PC a user signs in to (when using the per-user policy setting), and prevents users from specifying an alternate location for Work Folders on their PC (such as on a microSD card to conserve space on the primary drive). We suggest carefully evaluating your user's needs before forcing automatic setup.

Windows Intune

Windows Intune also provides a layer of security and manageability for non-domain-joined devices that would not otherwise be present. You can use Windows Intune to configure and manage users' personal devices such as tablets that connect to Work Folders from across the Internet. Windows Intune can provide devices with the sync server URL to use – otherwise users must enter their work email address to lookup the settings (if you publish a public Work Folders URL in the form of https://workfolders.*contoso.com*), or enter the sync server URL directly.

Without a Windows Intune deployment, users must configure external devices manually, which can result in increased demands on a customer's help desk staff.

You can also use Windows Intune to selectively wipe the data from Work Folders on a user's device without affecting the rest of their data – handy for if a user leaves your organization or has their device stolen.

Web Application Proxy/Device Registration Service

Work Folders is built around the concept of allowing Internet-connected devices to retrieve business data securely from the internal network, which allows users to "take their data with them" on their tablets and devices that would not normally be able to access work files. To do this, a reverse proxy must be used to publish sync server URLs and make them available to Internet clients.

While customers can leverage any of several different technologies to provide this functionality, Active Directory Federation Services (AD FS) in Windows Server 2012 R2 offers the following functionality that can be useful with Work Folders deployments:

- Web Application Proxy – a reverse proxy implementation that pre-authenticates access to web applications using AD FS. For more information, see Web Application Proxy Overview.
- Workplace Join and Device Registration Service– optionally allows users to join devices to the workplace that would not normally be domain-joined; for example, personal laptops, tablets, and smartphones. When these technologies are enabled, the AD FS administrator can configure applications and policies for corporate resources that require devices to be registered before they can gain access to published applications and resources. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
- Multifactor authentication (MFA) – optionally allows administrators to set up a policy based on user, device, or location data that requires users to authenticate with more than one authentication method; for example, a one-time password or a smart card. See Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.
- Multi-domain authentication – with AD FS you can support users from multiple domains syncing with a single sync server across the Internet. Without this capability, users from domains other than the one the sync server is joined to wouldn't be able to sync across the Internet.

Additional design considerations

In addition to understanding each of the components noted above, customers need to spend time in their design thinking about the number of sync servers and shares to operate, and whether or not to leverage failover clustering to provide fault tolerance on those sync servers

Number of Sync Servers

It is possible for a customer to operate multiple sync servers in an environment. This can be a desirable configuration for several reasons:

- Geographic distribution of users – for example, branch office files servers or regional datacenters
- Data storage requirements – certain business departments might have specific data storage or handling requirements that are easier with a dedicated server
- Load balancing – in large environments, storing user data on multiple servers can increase server performance and uptime.

For information on Work Folders server scaling and performance, see Performance Considerations for Work Folders Deployments.

Note

When using multiple sync servers, we recommend setting up automatic server discovery for users. This process relies upon the configuration of an attribute on each user account in AD DS. The attribute is named **msDS-SyncServerURL** and becomes available on user accounts after a Windows Server 2012 R2 domain controller is added to the domain or the Active Directory schema updates are applied. This attribute should be set for each user to ensure that users connect to the appropriate sync server. By using automatic server discovery, organizations can publish Work Folders behind a "friendly" URL such as *https://workfolders.contoso.com*, regardless of the number of sync servers in operation.

Number of Sync Shares

Individual sync servers can maintain multiple sync shares. This can be useful for the following reasons:

- Auditing and security requirements – If data used by a certain department must be more heavily audited or retained for a longer period of time, separate sync shares can help administrators keep user folders with differing audit levels separated.
- Differing quotas or file screens – If you want to set different storage quotas or limits on which file types are allowed in Work Folders (file screens) for different groups of users, separate sync shares can help.
- Departmental control – If administrative duties are distributed by department, utilizing separate shares for different departments can aid administrators in enforcing quotas or other policies.
- Differing device policies –If an organization needs to maintain multiple device policies (such as encrypting Work Folders) for different groups of users, using multiple shares enables this.
- Storage capacity –If a file server has multiple volumes, additional shares can be used to take advantage of these additional volumes. An individual share has access to only the volume that it is hosted on, and is unable to take advantage of additional volumes on a file server.

Access to Sync Shares

While the sync server that a user accesses is determined by the URL entered at their client (or the URL published for that user in AD DS when using server automatic discovery), access to individual sync shares is determined by the permissions present on the share.

As a result, if a customer is hosting multiple sync shares on the same server, care must be taken to ensure that individual users have permissions to access only one of those shares. Otherwise, when users connect to the server, their client may connect to the wrong share. This can be accomplished by creating a separate security group for each sync share.

Further, access to an individual user's folder inside a sync share is determined by ownership rights on the folder. When creating a sync share, Work Folders by default grants users exclusive access to their files (disabling inheritance and making them the owner of their individual folders).

Design checklist

The following set of design questions is intended to aid customers in designing a Work Folders implementation that best serves their environment. Customers should work through this checklist prior to attempting to deploy servers.

- Intended Users
  - Which users will use Work Folders?
  - How are users organized? (Geographically, by office, by department, etc)
  - Do any users have special requirements for data storage, security, or retention?
  - Do any users have specific device policy requirements, such as encryption?
  - Do you need to interoperate with or migrate from other user data management solutions such as Folder Redirection?
  - Do users from multiple domains need to sync across the Internet with a single server?
  - Do you need to support users who aren't members of the Local Administrators group on their domain-joined PCs? (If so, you'll need to exclude the relevant domains from Work Folders device policies such as encryption and password policies)
- Infrastructure and Capacity Planning
  - In what sites should sync servers be located on the network?
  - Will any sync servers be hosted by an Infrastructure as a Service (IaaS) provider such as in an Azure VM?
  - Will dedicated servers be needed for any specific user groups, and if so, how many users for each dedicated server?
  - Where are the Internet ingress/egress points on the network?
  - Will sync servers be clustered for fault-tolerance?
  - Will sync servers need to maintain multiple data volumes to host user data?
- Data Security
  - Will multiple sync shares need to be created on any sync servers?
  - What groups will be used to provide access to sync shares?
  - If you're using multiple sync servers, what security group will you create for the delegated ability to modify the msDS-SyncServerURL property of user objects?
  - Are there any special security or auditing requirements for individual sync shares?
  - Is multi-factor authentication (MFA) required?
  - Do you need the ability to remotely wipe Work Folders data from PCs and devices?
- Device Access
  - What URL will be used to provide access for Internet-based devices (*the default URL that is required for email-based automatic server discovery is workfolders.domainname*)?

- How will the URL be published to the Internet?
- Will automatic server discovery be used?
- Will Group Policy be used to configure domain-joined PCs?
- Will Windows Intune be used to configure external devices?
- Will Device Registration be required for devices to connect?

Next steps

After designing your Work Folders implementation, it's time to deploy Work Folders. For more information, see Deploying Work Folders.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	• Work Folders Overview
Deployment	• Deploying Work Folders • Performance Considerations for Work Folders Deployments • Work Folders Test Lab Deployment (blog post)

Deploying Work Folders

This topic discusses the steps needed to implement Work Folders. It assumes that you've already read Designing a Work Folders Implementation.

To deploy Work Folders, a process that can involve multiple servers and technologies, use the following steps.

Tip
The simplest Work Folders deployment is a single file server (often called a sync server) without support for syncing over the Internet, which can be a useful deployment for a test lab or as a sync solution for domain-joined client computers. To create a simple deployment, these are minimum steps to follow:

- Step 3: Install Work Folders on file servers
- Step 5: Create security groups for Work Folders
- Step 7: Create sync shares for user data

Step 1: Obtain SSL certificates

To allow users to sync across the Internet, the URL published by Work Folders must be protected by an SSL certificate. The requirements for SSL certificates used by Work Folders are as follows:

- The certificate must be issued by a trusted certification authority. For most Work Folders implementations, a publicly trusted CA is recommended, since certificates will be used by non-domain-joined, Internet-based devices.
- The certificate must be valid.
- The private key of the certificate must be exportable (as you will need to install the certificate on multiple servers).
- The subject name of the certificate must contain the public Work Folders URL used for discovering the Work Folders service from across the Internet – this must be in the format of **workfolders.**<*domain_name*>.
- Subject alternative names (SANs) must be present on the certificate listing the server name for each sync server in use.

Step 2: Create DNS records

To allow users to sync across the Internet, you must create a Host (A) record in public DNS to allow Internet clients to resolve your Work Folders URL. This DNS record should resolve to the external interface of the reverse proxy server.

On your internal network, you should create a DNS alias record for the Work Folders URL that resolves to the server names of all sync servers on the network.

Step 3: Install Work Folders on file servers

You can install Work Folders on a domain-joined server by using Server Manager or by using Windows PowerShell, locally or remotely across a network. This is useful if you are configuring multiple sync servers across your network.

To deploy the role in Server Manager, do the following:

1. Start the **Add Roles and Features Wizard**.
2. On the **Select installation type** page, choose **Role-based or feature-based deployment**.
3. On the **Select destination server** page, select the server on which you want to install Work Folders.
4. On the **Select server roles** page, expand **File and Storage Services**, expand **File and iSCSI Services**, and then select **Work Folders**.
5. When asked if you want to install **IIS Hostable Web Core**, click **Ok** to install the minimal version of Internet Information Services (IIS) required by Work Folders.
6. Click **Next** until you have completed the wizard.

To deploy the role by using Windows PowerShell, use the following cmdlet:

```
Add-WindowsFeature FS-SyncShareService
```

Step 4: Binding the SSL certificate on the sync servers

Work Folders installs the IIS Hostable Web Core, which is an IIS component designed to enable web services without requiring a full installation of IIS. After installing the IIS Hostable Web Core, you should bind the SSL certificate for the server to the Default Web Site on the file server. However, the IIS Hostable Web Core does not install the IIS Management console.

There are two options for binding the certificate to the Default Web Interface. To use either option you must have installed the private key for the certificate into the computer's personal store.

- Utilize the IIS management console on a server that has it installed. From within the console, connect to the file server you want to manage, and then select the Default Web Site for that server. The Default Web Site will appear disabled, but you can still edit the bindings for the site and select the certificate to bind it to that web site.
- Use the netsh command to bind the certificate to the Default Web Site https interface. The command is as follows:

```
netsh http add sslcert ipport=<IP address>:443 certhash=<Cert
thumbprint> appid={CE66697B-3AA0-49D1-BDBD-A25C8359FD5D}
certstorename=MY
```

Step 5: Create security groups for Work Folders

Before creating sync shares, a member of the Domain Admins or Enterprise Admins groups needs to create some security groups in Active Directory Domain Services (AD DS) for Work Folders (they might also want to delegate some control as described in Step 6). Here are the groups you need:

- One group per sync share to specify which users are allowed to sync with the sync share
- One group for all Work Folders administrators so that they can edit an attribute on each user object that links the user to the correct sync server (if you're going to use multiple sync servers)

Groups should follow a standard naming convention and should be used only for Work Folders to avoid potential conflicts with other security requirements.

To create the appropriate security groups, use the following procedure multiple times – once for each sync share, and once to optionally create a group for file server administrators.

To create security groups for Work Folders

1. Open Server Manager on a Windows Server 2012 R2 or Windows Server 2012 computer with Active Directory Administration Center installed.
2. On the **Tools** menu, click **Active Directory Administration Center**. Active Directory Administration Center appears.
3. Right-click the container where you want to create the new group (for example, the Users container of the appropriate domain or OU), click **New**, and then click **Group**.

4. In the **Create Group** window, in the **Group** section, specify the following settings:
   - In **Group name**, type the name of the security group, for example: **HR Sync Share Users**, or **Work Folders Administrators**.
   - In **Group scope**, click **Security**, and then click **Global**.
5. In the **Members** section, click **Add**. The Select Users, Contacts, Computers, Service Accounts or Groups dialog box appears.
6. Type the names of the users or groups to which you grant access to a particular sync share (if you're creating a group to control access to a sync share), or type the names of the Work Folders administrators (if you're going to configure user accounts to automatically discover the appropriate sync server), click **OK**, and then click **OK** again.

To create a security group by using Windows PowerShell, use the following cmdlets:

```
$GroupName = "Work Folders Administrators"

$DC = "DC1.contoso.com"

$ADGroupPath = "CN=Users,DC=contoso,DC=com"

$Members = "CN=Maya Bender,CN=Users,DC=contoso,DC=com","CN=Irwin
Hume,CN=Users,DC=contoso,DC=com"


New-ADGroup -GroupCategory:"Security" -GroupScope:"Global" -Name:$GroupName -
Path:$ADGroupPath -SamAccountName:$GroupName -Server:$DC

Set-ADGroup -Add:@{'Member'=$Members} -Identity:$GroupName -Server:$DC
```

Step 6: Optionally delegate user attribute control to Work Folders administrators

If you are deploying multiple sync servers and want to automatically direct users to the correct sync server, you'll need to update an attribute on each user account in AD DS. However, this normally requires getting a member of the Domain Admins or Enterprise Admins groups to update the attributes, which can quickly become tiresome if you need to frequently add users or move them between sync servers.

For this reason, a member of the Domain Admins or Enterprise Admins groups might want to delegate the ability to modify the msDS-SyncServerURL property of user objects to the Work Folders Administrators group you created in Step 5, as described in the following procedure.

Delegate the ability to edit the msDS-SyncServerURL property on user objects in AD DS

1. Open Server Manager on a Windows Server 2012 R2 or Windows Server 2012 computer with Active Directory Users and Computers installed.
2. On the **Tools** menu, click **Active Directory Users and Computers**. Active Directory Users and Computers appears.

3. Right-click the OU under which all user objects exist for Work Folders (if users are stored in multiple OUs or domains, right-click the container that is common to all of the users), and then click **Delegate Control…**. The Delegation of Control Wizard appears.
4. On the **Users or Groups** page, click **Add…** and then specify the group you created for Work Folders administrators (for example, **Work Folders Administrators**).
5. On the **Tasks to Delegate** page, click **Create a custom task to delegate**.
6. On the **Active Directory Object Type** page, click **Only the following objects in the folder**, and then select the **User objects** checkbox.
7. On the **Permissions** page, clear the **General** checkbox, select the **Property-specific** checkbox, and then select the **Read msDS-SyncServerUrl**, and **Write msDS-SyncServerUrl** checkboxes.

To delegate the ability to edit the msDS-SyncServerURL property on user objects by using Windows PowerShell, use the following example script that makes use of the DsAcls command.

```
$GroupName = "Contoso\Work Folders Administrators"

$ADGroupPath = "CN=Users,dc=contoso,dc=com"


DsAcls $ADGroupPath /I:S /G ""$GroupName":RPWP;msDS-SyncServerUrl;user"
```

Note
The delegation operation might take a while to run in domains with a large number of users.

Step 7: Create sync shares for user data

At this point, you're ready to designate a folder on the sync server to store your user's files. This folder is called a sync share, and you can use the following procedure to create one.

1. If you don't already have an NTFS volume with free space for the sync share and the user files it will contain, create a new volume and format it with the NTFS file system.
2. In Server Manager, click **File and Storage Services**, and then click **Work Folders**.
3. A list of any existing sync shares is visible at the top of the details pane. To create a new sync share, from the **Tasks** menu choose **New Sync Share…**. The New Sync Share Wizard appears.
4. On the **Select the server and path** page, specify where to store the sync share. If you already have a file share created for this user data, you can choose that share. Alternatively you can create a new folder.

   Note
   By default, sync shares aren't directly accessible via a file share (unless you pick an existing file share). If you want to make a sync share accessible via a file share, use the **Shares** tile of Server Manager or the New-SmbShare cmdlet to create a file share, preferably with access-based enumeration enabled.

5. On the **Specify the structure for user folders** page, choose a naming convention for user folders within the sync share. There are two options available:
   - **User alias** creates user folders that don't include a domain name. If you are using a file share that is already in use with Folder Redirection or another user data solution, select this naming convention. You can optionally select the **Sync only the following subfolder** checkbox to sync only a specific subfolder, such as the Documents folder.
   - **User alias@domain** creates user folders that include a domain name. If you aren't using a file share already in use with Folder Redirection or another user data solution, select this naming convention to eliminate folder naming conflicts when multiple users of the share have identical aliases (which can happen if the users belong to different domains).
6. On the **Enter the sync share name** page, specify a name and a description for the sync share. This is not advertised on the network but is visible in Server Manager and Windows Powershell to help distinguish sync shares from each other.
7. On the **Grant sync access to groups** page, specify the group that you created that lists the users allowed to use this sync share.

   Important

   To improve performance and security, grant access to groups instead of individual users and be as specific as possible, avoiding generic groups such as Authenticated Users and Domain Users. Granting access to groups with large numbers of users increases the time it takes Work Folders to query AD DS. If you have a large number of users, create multiple sync shares to help disperse the load.
8. On the **Specify device policies** page, specify whether to request any security restrictions on client PCs and devices. There are two device policies that can be individually selected:
   - **Encrypt Work Folders** Requests that Work Folders be encrypted on client PCs and devices
   - **Automatically lock screen, and require a password** Requests that client PCs and devices automatically lock their screens after 15 minutes, require a six-character or longer password to unlock the screen, and activate a device lockout mode after 10 failed retries

     Security

     To apply Work Folders password policies that are more restrictive than policies currently in place, users must belong to the Administrators group on their computers. While users usually are administrators of their personal devices and computers, they aren't always administrators on domain-joined computers. To enforce password policies for non-administrators on domain-joined computers, you should instead use Group Policy password policies for the user and/or computer domains and exclude these domains from the Work Folders password policies. You can exclude domains by using the Set-Syncshare - PasswordAutoExcludeDomain cmdlet after creating the sync share.
9. Review your selections and complete the wizard to create the sync share.

You can create sync shares using Windows PowerShell by using the New-SyncShare cmdlet. Below is an example of this method:

```
New-SyncShare "HR Sync Share" K:\Share-1 -User "HR Sync Share Users"
```

The example above creates a new sync share named *Share01* with the path *K:\Share-1*, and access granted to the group named *HR Sync Share Users*

Tip

After you create sync shares you can use File Server Resource Manager functionality to manage the data in the shares. For example, you can use the **Quota** tile inside the Work Folders page in Server Manager to set quotas on the user folders. You can also use File Screening Management to control the types of files that Work Folders will sync, or you can use the scenarios described in Dynamic Access Control: Scenario Overview for more sophisticated file classification tasks.

Step 8: Optionally specify a tech support email address and Active Directory Federation Services authentication

After installing Work Folders on a file server, you probably want to specify an administrative contact email address for the server, and you might want to enable Active Directory Federation Services (AD FS) authentication. To do either of these tasks, use the following procedure:

Specifying an administrative contact email or AD FS authentication

1. In Server Manager, click **File and Storage Services**, and then click **Servers**.
2. Right-click the sync server, and then click **Work Folders Settings**. The Work Folders Settings window appears.
3. On the **Authentication** page, optionally choose **Active Directory Federation Services** and specify a Federation Service URL. For more information about AD FS, see Active Directory Federation Services Overview.

   Note

   If the sync server isn't in the same Active Directory site as the AD FS server and the network traffic must go through a proxy server, you need to configure the sync server to use the correct proxy configuration. For more information, see the following topic: How to configure a proxy server for the Work Folders service.
4. In the navigation pane, click **Support Email** and then type the email address or addresses that users should use when emailing for help with Work Folders. Click **OK** when you're finished.

   Work Folders users can click a link in the Work Folders Control Panel item that sends an email containing diagnostic information about the client PC to the address(es) you specify here.

Step 9: Optionally set up server automatic discovery

If you are hosting multiple sync servers in your environment, you should configure server automatic discovery by populating the **msDS-SyncServerURL** property on user accounts in AD DS.

Before you can do so, you must install a Windows Server 2012 R2 domain controller or update the forest and domain schemas by using the `Adprep /forestprep` and `Adprep /domainprep` commands. For information on how to safely run these commands, see Running Adprep.

You probably also want to create a security group for file server administrators and give them delegated permissions to modify this particular user attribute, as described in Step 5 and Step 6. Without these steps you would need to get a member of the Domain Admins or Enterprise Admins group to configure automatic discovery for each user.

To specify the sync server for users

1. Open Server Manager on a computer with Active Directory Administration Tools installed.
2. On the **Tools** menu, click **Active Directory Administration Center**. Active Directory Administration Center appears.
3. Navigate to the **Users** container in the appropriate domain, right-click the user you want to assign to a sync share, and then click **Properties**.
4. In the Navigation pane, click **Extensions**.
5. Click the **Attribute Editor** tab, select **msDS-SyncServerUrl** and then click **Edit**. The Multi-valued String Editor dialog box appears.
6. In the **Value to add** box, type the URL of the sync server with which you want this user to sync, click **Add**, click **OK**, and then click **OK** again.

   Note
   The sync server URL is simply **https://** or **http://** (depending on whether you want to require a secure connection) followed by the fully qualified domain name of the sync server. For example, **https://sync1.contoso.com**.

To populate the attribute for multiple users, use Active Directory PowerShell. Below is an example that populates the attribute for all members of the *HR Sync Share Users* group, discussed in Step 5.

```
$SyncServerURL = "https://sync1.contoso.com"

$GroupName = "HR Sync Share Users"


Get-ADGroupMember -Identity $GroupName |

Set-ADUser -Add @{"msDS-SyncServerURL"=$SyncServerURL}
```

Step 10: Configure Web Application Proxy or another reverse proxy

To enable users to sync their Work Folders across the Internet, you need to publish Work Folders through a reverse proxy, making Work Folders available externally on the Internet. You can use Web Application Proxy, which is included in Active Directory Federation Services (AD FS), to publish Work Folders on the Internet, or you can use another reverse proxy solution.

For background information about Web Application Proxy, see Web Application Proxy Overview. For details on publishing applications such as Work Folders on the Internet using Web Application Proxy, see Overview: Connect to Applications and Services from Anywhere with Web Application Proxy.

Step 11: Optionally use Group Policy to configure domain-joined PCs

If you have a large number of domain-joined PCs to which you want to deploy Work Folders, you can use Group Policy to do the following client PC configuration tasks:

- Specify which sync server users should sync with
- Force Work Folders to be set up automatically, using default settings (review the Group Policy discussion in Designing a Work Folders Implementation before doing this)

To control these settings, create a new Group Policy object (GPO) for Work Folders and then configure the following Group Policy settings as appropriate:

- "Specify Work Folders settings" policy setting in User Configuration\Policies\Administrative Templates\Windows Components\WorkFolders
- “Force automatic setup for all users” policy setting in Computer Configuration\Policies\Administrative Templates\Windows Components\WorkFolders

Note
These policy settings are available only when editing Group Policy from a computer running Group Policy Management on Windows 8.1 or Windows Server 2012 R2. Versions of Group Policy Management from earlier operating systems do not have this setting available.

See also

For additional related information, see the following resources.

Content type	References
Product evaluation	• Work Folders Overview
Deployment	• Designing a Work Folders Implementation • Work Folders Test Lab Deployment (blog

Content type	References
	post) • A new user attribute for Work Folders server Url (blog post)
Technical Reference	• Interactive logon: Machine account lockout threshold • Sync Share Cmdlets

Group Policy Overview

Use this topic to find the documentation resources and other technical information you need to accomplish key Group Policy tasks. You can learn about new and updated functionality for Group Policy in Windows Server 2012 R2 and Windows 8.1and ways to automate common Group Policy tasks by using Windows PowerShell.

Did you mean…

- Group Policy Settings Reference for Windows Server 2012 and Windows 8
- Overview of Group Policy (Windows Server 2008 R2)
- Overview of Preferences
- Refresh Group Policy settings with GPUpdate.exe
- Security Policy Settings Overview
- **Remote Server Administration Tools for Windows Server 2012**
- **Windows Firewall with Advanced Security Overview**

Feature description

Group Policy is an infrastructure that allows you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences. To configure Group Policy settings that affect only a local computer or user, you can use the Local Group Policy Editor. You can manage Group Policy settings and Group Policy Preferences in an Active Directory Domain Services (AD DS) environment through the Group Policy Management Console (GPMC). Group Policy management tools also are included in the Remote Server Administration Tools pack to provide a way for you to administer Group Policy settings from your desktop.

Windows PowerShell When the GPMC is installed on servers or client computers, the Windows PowerShell module is also installed. You have full Windows PowerShell functionality. If you install the Remote Server Administration Tools pack, the latest Windows PowerShell cmdlets for Group Policy are also installed. For more information about Windows PowerShell cmdlets and scripts that you can use to manage Group Policy, see Group Policy Cmdlets.

You can configure the Group Policy feature by using Windows PowerShell cmdlets. For more information about using Server Manager cmdlets to install the Group Policy Management Console, see Install or Uninstall Roles, Role Services, or Features and Server Manager deployment cmdlet module.

Practical applications

By using Group Policy, you can significantly reduce your organization's total cost of ownership. Various factors, such as the large number of policy settings available, the interaction between multiple policies, and inheritance options, can make Group Policy design complex. By carefully planning, designing, testing, and deploying a solution based on your organization's business requirements, you can provide the standardized functionality, security, and management control that your organization needs.

Here are some Windows Server 2012 scenarios that use Group Policy to implement a solution:

- Configuring Policy-based Quality of Service (QoS)
- Managing the New Wired Network (IEEE 802.3) Policies Settings
- Managing the Wireless Network (IEEE 802.11) Policies
- Securing End-to-End IPsec Connections by Using IKEv2 in Windows Server 2012

New and changed functionality

Group Policy designs can become very complex. Various factors, such as the large number of policy settings and preference items available, the interaction between multiple policies, and inheritance options, can make it difficult to determine if Group Policy is functioning correctly on each computer.

In Windows Server 2012, Group Policy focused on improving the Group Policy troubleshooting experience. Windows Server 2012 R2 expands the support for IPv6 networking, adds policy caching to reduce sign-in times in synchronous mode, and provides more detailed event logging. For more details about these changes and more information about the additional changes in functionality that are not listed here, see What's New in Group Policy in Windows Server 2012 and What's New in Group Policy in Windows Server 2012 R2.

See also

Content type	References
Planning	Not yet available.
Deployment	Not yet available.
Operations	Not yet available.
Troubleshooting	Group Policy Analysis and Troubleshooting

Content type	References
	Overview
Security	Not yet available.
Tools and settings	• Group Policy Settings Reference for Windows and Windows Server • Group Policy Cmdlets
Community resources	• Group Policy Team Blog • Group Policy forum
Related technologies	Not yet available.

What's New in Group Policy in Windows Server 2012 R2

This topic describes the new and changed functionality of the Group Policy feature in Windows Server 2012 R2.

Feature description

Group Policy is an infrastructure that enables you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences. For Group Policy settings that affect only a local computer or user, you can use the Local Group Policy Editor. You can manage Group Policy settings and Group Policy Preferences in an Active Directory Domain Services (AD DS) environment through the Group Policy Management Console (GPMC). Group Policy management tools also are included in the Remote Server Administration Tools pack, which provides a way for you to administer Group Policy settings from your desktop. For more information about Group Policy, see the Group Policy Overview.

New and changed functionality

The following table lists functionality in Group Policy that is new for this release or has been changed.

Feature/functionality	New or Updated
IPv6 support	Updated
Policy caching	New
Event logging	Updated

IPv6 support

Windows Server 2012 R2 expands support for IPv6 in Group Policy. This expanded support encompasses printers, item-level targeting, and VPN networks. The following procedures describe how you can use this expanded support.

Printers

1. Under **User Configuration**, expand **Preferences**, and expand **Control Panel Settings**.
2. Click **Printers**, right-click, click **New**, and then click **TIP/IP Printer**.
3. Select the **Use IPv6 Address** check box and enter an IPv6 address in the corresponding field.

Item-level targeting

1. In any Group Policy preference, click **Properties**, and click the **Common** tab.
2. Select the **Item-level targeting** check box, and click the **Targeting…** button.
3. In the **Targeting Editor**, click the **New Item** menu, and click **IP Address Range**.
4. Select the **Use IPv6** check box and enter an IPv6 address in the corresponding field. You can filter by **submpa**.

VPN connections

1. Under **User Configuration**, expand **Preferences**, and expand **Control Panel Settings**.
2. Click **Network Option**, right-click, click **New**, and click **VPN Connection**.
3. In the **New VPN Connection** window, select the **Use IPv6** check box.

Policy caching

Windows Server 2012 R2 , when Group Policy gets the latest version of a policy from the domain controller, it writes that policy to a local store. Then if Group Policy is running in synchronous mode the next time the computer reboots, it reads the most recently downloaded version of the policy from the local store, instead of downloading it from the network. This reduces the time it takes to process the policy. Consequently, the boot time is shorter in synchronous mode. This is especially important if you have a latent connection to the domain controller, for example, with DirectAccess or for computers that are off premises. This behavior is controllable by a new policy called Configure Group Policy Caching.

Event logging

Windows Server 2012 R2 provides more detailed events for Group Policy settings in the operational event log. The logged events include more information about how long it takes to download and process policies. There are also details about WMI processing. This can help you analyze and diagnose long sign-in times.

See also

The following table provides additional resources for evaluating Group Policy.

Content type	References
Product evaluation	Group Policy Overview | Group Policy Techcenter
Community resources	Group Policy Team Blog | TechNet Wiki
Related technologies	Active Directory Domain Services Overview

What's New in Group Policy in Windows Server 2012

This topic describes the new and changed functionality of the Group Policy feature in Windows Server 2012

Feature description

Group Policy is an infrastructure that enables you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences. For Group Policy settings that affect only a local computer or user, you can use the Local Group Policy Editor. You can manage Group Policy settings and Group Policy Preferences in an Active Directory Domain Services (AD DS) environment through the Group Policy Management Console (GPMC). Group Policy management tools also are included in the Remote Server Administration Tools pack, which provides a way for you to administer Group Policy settings from your desktop. For more information about Group Policy, see the Group Policy Overview.

New and changed functionality

The following table lists functionality in Group Policy that is new for this release or has been changed.

Feature/functionality	New or Updated
Remote Group Policy Update	New
Group Policy Results report improvements	Updated
Group Policy infrastructure status	New
Local Group Policy support for Windows RT	New

Feature/functionality	New or Updated
Sign-in optimizations	Updated
Fast Startup	New
New Group Policy starter GPOs	New
Group Policy cmdlet changes	Updated
Registry.pol changes	Updated
Group Policy Client service idle state	Updated
Group Policy settings in Internet Explorer 10	New
Group Policy Preferences for Internet Explorer 10	New

Remote Group Policy Update

In Windows Server 2012, you can refresh Group Policy settings, including security settings that are set on a group of remote computers, by using the functionality that is added to the context menu for an organizational unit (OU) in the Group Policy Management Console (GPMC). This functionality schedules a task on all computers in a selected OU, which refreshes the computer and user Group Policy settings.

What value does this change add?

When you troubleshoot Group Policy issues for a specific computer or user, you can run **gpupdate.exe** to verify that the most current Group Policy settings have been applied. This command-line utility needs to be run on a specific computer. In Windows Server 2012, you can schedule **gpupdate.exe** to run on multiple computers from the GPMC or from a Windows PowerShell session by using the new **Invoke-GPUpdate** cmdlet.

What works differently?

Prior to Windows Server 2012, you would have to remote to a specific computer and run **gpupdate.exe** from the command-line. In Windows Server 2012 you can update Group Policy for all computers in a specific OU and the OUs that it contains. For more information about the new remote Group Policy update feature, see Force a Remote Group Policy Refresh (GPUpdate)

Group Policy Results report improvements

Group Policy Results in Windows Server 2012 includes more information to help determine if a Group Policy setting was applied to a computer or user. If the results do not match the expected results, there is information about why this happened.

What value does this change add?

It is sometimes hard to determine why Group Policy applied the specific policy settings and Preferences. The **Group Policy Results** report includes the following new information to help you understand why a particular Group Policy result was achieved:

- Whether the connection was determined to be a slow link or fast link
- Whether block inheritance has been set
- Whether loopback has been set
- The processing time for each client-side extension
- The GPO name is now displayed with each Group Policy setting and preference item. This identifies which is the winning GPO for a particular policy setting or preference item.

What works differently?

The following applicable conditions are displayed on the Group Policy Results **Summary** tab:

- If a slow link or fast link is detected
- If block inheritance is set
- If loopback is enabled

The Group Policy Results **Details** tab displays:

1. The OU that contains the computer or user.
2. The **Component Status** section displays the amount of time each client-side extension took to process and the last time each client-side extension processed.
3. The **Component Status** section provides a link in the **Event Log** column that displays the event log messages from the last Group Policy refresh. This functionality is equivalent to the information that is returned from the GPLogview.exe utility.
4. The **Winning GPO** name is displayed in a table with each **Policy** setting name and the value that is set for each policy setting and preference item.

Important

To view the **Group Policy Results** for a specific computer, the following firewall rules must be set on each client computer to allow the following connections:

- Remote Event Log Management (NP-IN)
- Remote Event Log Management (RPC)
- Remote Event Log Management (RPC-EPMAP)
- Windows Management Instrumentation (WMI-IN)

If you do not want to allow the connections on computers, you can also run **Gpresult.exe /h** *<filename.html>* from the command line on the each local computer, where filename.html is the name of a file to which **Gpresult** writes the output.

Group Policy infrastructure status

Display the status of Active Directory and SYSVOL replication as it relates to all Group Policy Objects or a single Group Policy Object.

What value does this change add?

Group Policy relies on being stored and replicated to all domain controllers in a domain. There can be a lag time after a change is made on one domain controller before the change is replicated to all other domain controllers. Until changes to a GPO are replicated to the domain controller that a client computer is accessing, that computer will receive the earlier version of the GPO during Group Policy refresh. In earlier versions of the Windows operating system, administrators had to download **GPOtool.exe** to diagnose these issues.

What works differently?

In Windows Server 2012, you no longer need to download and run a separate tool for monitoring and diagnosing replication issues related to Group Policy at the domain level. Potential differences that can be viewed by using the Group Policy infrastructure status are:

- Active Directory and SYSVOL security descriptor (ACL details)
- Active Directory and SYSVOL GPO version details
- Number of GPOs listed in Active Directory and SYSVOL for each domain controller

For more information about the new Group Policy infrastructure status, see Check Group Policy Infrastructure Status.

Local Group Policy support for Windows RT

You can configure Group Policy on devices that run Windows RT.

What value does this change add?

Local Group Policy is available for Windows RT. It is off by default, but can be turned on by the local administrator.

What works differently?

For Windows RT devices, the Group Policy Client service is disabled by default. The Group Policy Client service must be set to Automatic and started by the administrator before Group Policy is processed on the device.

To turn on the Group Policy Client service

1. From the start screen, type **Services.msc**.
2. Double-click **Group Policy Client** to open the **Group Policy Client Properties (Local Computer)** dialog box.
3. Set the **Startup type** to **Automatic**, and then click the **Start** button.

Sign-in optimizations

Slow link processing during sign-in is improved.

What value does this change add?

More control to determine if the network connection should be processed as a slow link improves the sign-in experience for users by allowing users to sign-in faster.

What works differently?

For DirectAccess connections, when the network connection speed cannot be determined, Group Policy processing defaults to slow-link mode. During sign-in, if a slow link is detected, Group Policy automatically switches to asynchronous processing. A new policy setting enables administrators to configure all 3G connections so that they are treated as a slow link. To disable 3G slow-link connections, select the **Always treat WWAN connections as a slow link** check box after you have enabled the **Configure Group Policy slow link detection** policy setting.

Note
The **Configure Group Policy slow link detection** policy setting is located under Computer Configuration\Policies\Administrative Templates\System\Group Policy in the Group Policy Management Editor.

Fast Startup

Improvement in start times affects Group Policy processing.

What value does this change add?

Fast Startup decreases the time it takes to shut down and start a computer by causing the computer to enter a hibernate state instead of performing a full shutdown.

What works differently?

Group Policy settings or scripts that are applied during start or shutdown might not be applied when computers are configured to use Fast Startup. For more information about the impact of this change, see Understand the Effect of Fast Logon Optimization and Fast Startup on Group Policy.

New Group Policy starter GPOs

You can configure firewall port requirements to allow Group Policy remote RSoP reporting and remote Group Policy update.

What value does this change add?

The two starter GPOs make configuring Group Policy firewall port requirements easier. You can import the starter GPOs at the same time that you create a new GPO for this purpose.

What works differently?

For more information about the new Starter GPOs, see Configure Firewall Port Requirements for Group Policy.

Group Policy cmdlet changes

Automate the remote Group Policy update by using the new **Invoke-GPUpdate** cmdlet.

Note
Group Policy cmdlets can also run on the Server Core installation option.

What value does this change add?

The new **Invoke-GPUpdate** cmdlet provides more functionality than applying remote Group Policy update through the GPMC interface. For example, the **Invoke-GPUpdate** cmdlet enables

you to refresh computers located in the default computer container, while the remote Group Policy update functionality in the GPMC enables you to remotely refresh only computers that are located in an OU.

On a Server Core installation, you can manage GPMC functionality by using Group Policy cmdlets. This provides you with more flexibility for managing Group Policy.

The **Get-GPPermissions** cmdlet and the **Set-GPPermissions** cmdlet are renamed to the singular form: **Get-GPPermission** and **Set-GPPermission**. Both cmdlets have aliases for their previous names to support backward compatibility.

What works differently?

For more information about all Group Policy cmdlets, see Group Policy Cmdlets in Windows PowerShell.

Registry.pol changes

An increase in the maximum size allowed for registry.pol enables faster downloads of registry.pol files from domain controllers.

What value does this change add?

With this change, there should be very few situations where the maximum size of the registry.pol file restrains administrators from adding new Administrative Template settings to a GPO. With the faster download of registry.pol files, Group Policy processing should increase.

What works differently?

The registry.pol file maximum size is increased to 100 MB. The Group Policy processing has been changed to read larger amounts of data from a registry.pol file when processing the Administrative Template settings. This change results in less network access for reading the registry.pol file from the domain controller, which speeds up Group Policy processing.

Group Policy Client service idle state

The Group Policy Client service will sleep when the Group Policy service is idle for more than 10 minutes.

What value does this change add?

Group Policy processes approximately every 90 minutes, by default. Setting the Group Policy Client service to sleep in between processing helps create better performance for client computers.

What works differently?

Group Policy background refresh starts as a scheduled task, not as a service that continuously checks to determine when it is time to run the background refresh. The scheduled task model requires less overhead processing, which creates better performance for client computers.

Group Policy settings in Internet Explorer 10

Group Policy Administrative Template settings that support Internet Explorer 10 are added.

What value does this change add?

New Group Policy settings in the Internet Explorer 10 Administrative Template support new features.

What works differently?

For more information about the new and changed policy settings that can be used to manage and control your Internet Explorer 10 configuration, see Group Policy Settings in Internet Explorer 10.

Group Policy Preferences for Internet Explorer 10

Windows Server 2012 and Windows 8 include Group Policy Preferences support for Internet Explorer 10.

What value does this change add?

Group Policy Preferences consolidates multiple ways to configure Internet Explorer preference settings.

What works differently?

The Internet Explorer Maintenance (IEM) snap-in is replaced by the Internet Explorer 10 preference extension. Administrators can use the Internet Explorer 10 preference extension or the Internet Explorer Administration Kit (IEAK) to configure Group Policy settings. Information about the Internet Explorer 10 preference extension can be found at:

- Internet Explorer Maintenance Replacements
- Configure an Internet Explorer 10 Item

Removed or deprecated functionality

The following features and functionalities have been removed from this release of Group Policy. Applications, code, or usage that depends on these features will not function in this release unless you employ an alternate method. For more information about removed or deprecated functionality in this release, see Features Removed or Deprecated in the Windows Server 2012.

- The Immediate Task preference item no longer supports the following actions:
  - Send an email
  - Display a message

  For more information, see Configure an Immediate Task Item (Windows Vista and later versions)
- NetMeeting Administrative Template files are removed

See also

The following table provides additional resources for evaluating Group Policy.

Content type	References
Product evaluation	Group Policy Overview | Group Policy Techcenter
Community resources	Group Policy Team Blog | TechNet Wiki
Related technologies	Active Directory Domain Services Overview

Configure Firewall Port Requirements for Group Policy

You can use the information in this topic to configure the firewall port requirements for Group Policy. Group Policy requires that firewall ports are opened on client computers for an administrator to perform these two remote operations:

- Remote Resultant Set of Policy (RSoP) Group Policy results: ports that require firewall rules
- Remote Group Policy refresh: ports that require firewall rules

By default, Windows Firewall enables all outbound network traffic,and it allows only inbound traffic that is enabled by firewall rules. This topic identifies the TCP and UDP ports for which you must have active firewall rules to allow the inbound traffic. This allows Group Policy to perform remote Group Policy Results reporting from client computers and to perform remote Group Policy refresh to client-based computers. You can use the information in this topic to configure non-Microsoft firewall products and to create a GPO to configure a client computer with the required firewall rules. This topic also presents two new starter Group Policy Objects (GPOs) that configure the proper firewall rules on client computers.

Notes

If you have configured client computers by using Group Policy, the Group Policy settings override any manual configuration of client computers to which the policies are applied. If you want to review these rules, from the Group Policy Management Console (GPMC), you can run a Group Policy Results report or Group Policy Planning report. Or from a client computer, open the Windows Firewall with Advanced Security MMC snap-in and click **Inbound Rules**.

Membership in the Administrators group or equivalent is the minimum permissions required to make these configuration changes.

If you use a non-Microsoft firewall product, check your firewall product documentation for instructions about how to open these ports to allow network traffic as required by Group Policy.

Remote Resultant Set of Policy (RSoP) Group Policy results: ports that require firewall rules

You can use the RSoP feature of the GPMC to create detailed reports about the policy settings that are applied to computers or users who have signed in. When RSoP reporting targets remote computers, all connections are direct to each remote client from the computer running the GPMC (that is, not transitively through a domain controller).

To use RSoP reporting for remotely targeted computers through the firewall, you must have firewall rules that allow inbound network traffic on the ports listed in the following table. This allows remote WMI and event log traffic to flow between the computer running the GPMC and the remotely targeted computer.

Server port	Type of network traffic
TCP SMB 445, all services and programs	Remote Event Log Management (NP-in)
TCP RPC dynamic ports, EventLog (Windows Event Log service)	Remote Event Log Management (RPC)
TCP port 135, RPCSS (Remote Procedure Call service)	Remote Event Log Management (RPC-EPMAP)
TCP all ports, Winmgmt (Windows Management Instrumentation service)	Windows Management Instrumentation (WMI-in)

Configure firewall rules by creating a GPO from the Group Policy Reporting Firewall Ports Starter GPO and linking to the domain

In Windows Server 2012, Group Policy adds a new Starter GPO called, **Group Policy Reporting Firewall Ports**. This Starter GPO includes policy settings to configure the firewall rules that are specified in the previous table. This enables inbound network traffic on the ports, which is necessary to allow the GPMC to gather the Group Policy results RSoP information from a remote computer. It is a best practice to create a new GPO from this Starter GPO, and then link the new GPO to your domain with a higher precedence than the Default Domain GPO, so that you can configure all computers in the domain for remote Group Policy results reporting.

Do this step using Windows PowerShell

To create a GPO from the Group Policy Reporting Firewall Ports Starter GPO and link to the domain

1. In the GPMC console tree, right-click the domain for which you want to configure all computers to enable a remote Group Policy refresh, and then click **Create a GPO in this**

domain, and Link it here…

2. In the **New GPO** dialog box, type the name of the new Group Policy Object in the **Name** box.
3. Select the **Group Policy Reporting Firewall Ports** Starter GPO from the **Source Starter GPO** list that you want to use to create a new Group Policy Object.

   Notes

   If you do not see any Starter GPOs listed, cancel creating a GPO and do the following before you return to Step 1:

   a. Navigate to **Starter GPOs**.

   b. In the results pane, click **Create Starter GPOs Folder**.
4. Click **OK**.
5. In the results pane, click the **Linked Group Policy Objects** tab.
6. Select the GPO that you just created. Click the Up arrow until the GPO you just created is located above the Default Domain Policy. The new GPO will then have a smaller link-order value than the Default Domain Policy.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

You can use the **New-GPO** cmdlet with the **–StarterGpoName** parameter to create a new GPO. You can then pipe the output from the **New-GPO** cmdlet to the **New-GPLink** cmdlet.

For example, to create a new GPO, called *Configure firewall rules for remote reporting*, based on the **Group Policy Reporting Firewall Ports** Starter GPO, and link the GPO to the Contoso.com domain, type the following:

```
New-GPO –Name "Configure firewall rules for remote reporting" –StarterGpoName "Group Policy Reporting Firewall Ports" | New-GPLink –target "dc=Contoso,dc=com" –LinkEnabled yes
```

For more information about the **New-GPO** cmdlet and the **New-GPLink** cmdlet, see:

- New-GPO
- New-GPLink

Remote Group Policy refresh: ports that require firewall rules

To schedule a remote Group Policy refresh for domain-joined computers you must have firewall rules that enable inbound network traffic on the ports listed in the following table.

Server port	Type of network traffic
TCP RPC dynamic ports, Schedule (Task Scheduler service)	Remote Scheduled Tasks Management (RPC)
TCP port 135, RPCSS (Remote Procedure Call service)	Remote Scheduled Tasks Management (RPC-EPMAP)
TCP all ports, Winmgmt (Windows Management Instrumentation service)	Windows Management Instrumentation (WMI-in)

Configure firewall rules by creating a GPO from the Group Policy Remote Update Firewall Ports Starter GPO and linking to the domain

In Windows Server 2012, Group Policy adds a new Starter GPO called **Group Policy Remote Update Firewall Ports**. This Starter GPO includes policy settings to configure the firewall rules that are specified in the previous table. This enables inbound network traffic on the ports, which is necessary to allow the remote Group Policy refresh to run. It is a best practice to create a new GPO from this Starter GPO, and then link the new GPO to your domain with a higher precedence than the Default Domain GPO, so that you can configure all computers in the domain to enable a remote Group Policy refresh.

Do this step using Windows PowerShell

To create a GPO from the Group Policy Remote Update Firewall Ports Starter GPO and link to the domain

1. In the GPMC console tree, right-click the domain for which you want to configure all computers to enable a remote Group Policy refresh, and then click **Create a GPO in this domain, and Link it here…**
2. In the **New GPO** dialog box, type the name of the new Group Policy Object in the **Name** box.
3. Select the **Group Policy Remote Update Firewall Ports** Starter GPO from the **Source Starter GPO** list.

   Notes

   If you do not see any Starter GPOs listed, cancel creating a GPO and do the following before you return to the Step 1:

   a. Navigate to **Starter GPOs**

   b. In the results pane, click **Create Starter GPOs Folder**
4. Click **OK**.
5. In the results pane, click the **Linked Group Policy Objects** tab.
6. Select the GPO that you just created. Click the up arrow until the GPO you just created is

above the Default Domain Policy in link order. The new GPO will then have a smaller link order value than the Default Domain Policy.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

You can use the **New-GPO** cmdlet with the **–StarterGpoName** parameter to create a new GPO. You can then pipe the output from the **New-GPO** cmdlet to the **New-GPLink** cmdlet.

For example, to create a new GPO, called *Configure firewall rules for remote gpupdate*, which is based on the **Group Policy Remote Update Firewall Ports** Starter GPO, and link the GPO to the Contoso.com domain, type the following:

```
New-GPO –Name "Configure firewall rules for remote gpupdate" –StarterGpoName "Group Policy Remote Update Firewall Ports" | New-GPLink –target "dc=Contoso,dc=com" –LinkEnabled yes
```

For more information about the **New-GPO** cmdlet and the **New-GPLink** cmdlet, see:

- New-GPO
- New-GPLink

Group Policy Management Console

Install Group Policy Management Console (GPMC)

Group Policy Management Console (GPMC) is a comprehensive administrative tool for Group Policy management. Administrators use GPMC to perform all Group Policy management tasks, with the exception of configuring individual policy settings in Group Policy Objects themselves, which is done with the Local Group Policy Object Editor. You can install GPMC through Server Manager on the server as follows.

1. On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type Server Manager. Click **Server Manager**.
2. In the main pane in **Server Manager**, click **Add roles and features**.
3. Follow the **Add Roles and Features Wizard** until you get to the **Features** menu.
4. Select **Group Policy Management** from the list of available features. Click **Install** and follow the steps in the wizard.

Open GPMC

To start GPMC, do the following:

On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **gpmc.msc**, and then click **OK** or press **ENTER**.

Resultant Set of Policy

You can use Resultant Set of Policy (RSoP) to simulate and test policy settings that are applied to computers or users using Group Policy.

Open RSoP as an MMC snap-in

1. Open Microsoft Management Console.
2. On the **File** menu, click **Add/Remove Snap-in**.
3. In the **Available snap-ins** box, click **Resultant Set of Policy**, and then click **Add**.
4. Click **OK**.

Note

You must add a Resultant Set of Policy snap-in to Microsoft Management Console (MMC) for each new RSoP query.

RSoP overview

Resultant Set of Policy (RSoP) is an addition to Group Policy to assist in policy implementation and troubleshooting. RSoP is a query engine that polls existing policies and planned policies, and then reports the results of those queries. It polls existing policies based on site, domain, domain controller, and organizational unit. RSoP gathers this information from the Common Information Management Object Model (CIMOM) database (otherwise known as CIM-compliant object repository) through Windows Management Instrumentation (WMI).

RSoP provides details about all policy settings that are configured by an Administrator, including Administrative Templates, Folder Redirection, Internet Explorer Maintenance, Security Settings, Scripts, and Group Policy Software Installation.

When policies are applied on multiple levels (for example, site, domain, domain controller, and organizational unit), the results can conflict. RSoP can help you determine a set of applied policies and their precedence (the order in which policies are applied).

RSoP consists of two modes: planning mode and logging mode. With planning mode, you can simulate the effect of policy settings that you want to apply to a computer and user. Logging mode reports the existing policy settings for a computer and user that is currently logged on.

The Resultant Set of Policy Wizard helps you create an RSoP query. You can open the wizard from Microsoft Management Console (MMC), Active Directory Users and Computers, or Active Directory Sites and Services. You must run the wizard at least once to create an RSoP query. When complete, the wizard displays the query results in the RSoP snap-in in MMC. From here,

you can save, change, and refresh your queries. You can create many RSoP queries by adding multiple Resultant Set of Policy snap-ins to MMC, one RSoP snap-in per query.

Local Group Policy Editor

Local Group Policy Editor is a Microsoft Management Console (MMC) snap-in that is used to configure and modify Group Policy settings within Group Policy Objects (GPOs).

Administrators need to be able to quickly modify Group Policy settings for multiple users and computers throughout a network environment. The Local Group Policy Editor provides administrators with a hierarchical tree structure for configuring Group Policy settings in GPOs. These GPOs can then be linked to sites, domains, and organizational units (OU) that contain computer or user objects.

Local Group Policy Editor consists of two main sections:

User Configuration Holds settings that are applied to users (at sign-in and during periodic background refresh),

Computer Configuration Holds settings that are applied to computers (at startup and during periodic background refresh).

These sections are further divided into the different types of policies that can be set, such as Administrative Templates, Security, and Folder Redirection.

To work efficiently, administrators need to have immediate access to information about the function and purpose of individual policy settings. For Administrative Templates policy settings, Local Group Policy Editor provides information about each policy setting directly in the web view of the console. This information shows operating system requirements, defines the policy setting, and includes any specific details about the effect of enabling or disabling the policy setting.

In addition, developers should be able to quickly and easily add Group Policy support to their software products. The Local Group Policy Editor is designed to be extensible. The easiest way for developers to extend Local Group Policy Editor for their applications is to write custom Administrative Template files that they can add to Local Group Policy Editor.

Use the Local Group Policy Editor

The following procedures explain how to open the Local Group Policy Editor from the command line or as an MMC snap-in.

To open the Local Group Policy Editor from the command line

- On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **gpedit.msc**, and then press **ENTER**.

To open the Local Group Policy Editor as a snap-in

1. On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc**, and then

press **ENTER**.

2. On the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Local Group Policy Editor**, and then click **Add**.
4. In the **Select Group Policy Object** dialog box, click **Browse**.
5. Click **This computer** to edit the Local Group Policy Object, or click **Another computer** to edit the Local Group Policy Object for a different computer. Click **Users** to edit Administrator, Non-Administrator, or per-user Local Group Policy Objects.
6. Click **Finish**.

Notes

If you want to save a Local Group Policy Editor console and choose which GPO opens in it from the command line:

In the **Select Group Policy Object** dialog box, select the **Allow the focus of the Group Policy Snap-in to be changed when launching from the command line** check box.

Group Policy Preferences

Group Policy Preferences is a collection of Group Policy client-side extensions that deliver preference settings to domain-joined computers running Microsoft Windows desktop and server operating systems. Preference settings are administrative configuration choices deployed to desktops and servers. Preference settings differ from policy settings because users have a choice to alter the administrative configuration. Policy settings administratively enforce setting, which restricts user choice.

Group Policy Preferences are distributed to domain-joined computers using the Group Policy. The flexibility of Group Policy enables it to deliver opaque configuration data to a domain-joined computer running Windows. The opaque data is then transferred to a Group Policy client side extension at which point the opaque data becomes relevant because the client-side extension understands the data.

This document describes how the Group Policy Drive Maps and Printers client-side extensions process their configuration data. With this knowledge, administrators can more effective design and deploy Group Policy Drive Map and Printer items in their environment. And, the information presented in this technical reference enables IT Professionals to troubleshoot Group Policy Drive Map and Printer processing.

Prerequisite Fundamentals

Group Policy

Group Policy is a management technology included in Windows Server that enables you to secure computer and user settings. Securing these settings ensures a common computing environment for users and lowers the total cost of ownership by restricting accidental or deliberate configurations that adversely affect the operating system.

A Group Policy object (GPO) is a logical object composed of two components, a Group Policy container and a Group Policy template. Windows stores both of these objects on domain controllers in the domain. The Group Policy container object is stored in the domain partition of Active Directory. The Group Policy template is a collection of files and folders stored on the system volume (SYSVOL) of each domain controller in the domain. Windows copies the container and template to all domain controllers in a domain. Active Directory replication copies the Group Policy container while the File Replication Service (FRS) or the Distributed File System Replication (DFSR) service copies the data on SYSVOL.

The Group Policy container and template together; make the logical object called a Group Policy object. Each Group Policy object contains two classes of configuration: user and computer. Computer configuration settings affect the computer as whole, regardless of the logged on user. User configuration settings affect the currently logged on user, and may vary with each user. Some examples of computers settings are power management, user rights, and firewall settings. Examples of user settings include Internet Explorer, display settings, and Folder Redirection.

Group Policy objects and their settings apply to computers and user to which they are linked. You can link GPOs to an Active Directory site, domain, organizational unit, or nested organizational unit. Group Policy objects separate from the containers to which they are linked. This separation enables you to link a single GPO to multiple containers. Linking GPOs to many containers enables a single GPO to apply to users or computer within multiple container. This defines the scope of the GPO. Computer configurations apply to computers within the container or nested containers. User configurations apply to users in the same fashion.

Policy settings apply to computers at computer startup and to users during user logon. Windows Server 2012 and Windows 8 includes a Group Policy service. During computer startup, the Group Policy service queries Active Directory for the list of GPOs that are within scope (linked) of the computer object. Again, this includes:

- The site in which the computer resides
- The domain in which the computer is a member
- The parent organizational unit to which the computer is a direct member and any other organizational units above the parent OU.

The Group Policy service decides which GPOs apply to computers (there are many ways to filer GPOs from applying, which is beyond the scope of this introduction) and applies those policy settings. Client-side extensions (CSEs) are responsible for applying policy settings contained in the GPOs. A Group Policy client-side extension is a separate component from the Group Policy service that is responsible for reading specific policy setting data from the GPO and applying it to

the computer or user. For example, the Group Policy registry client-side extension reads registry policy setting data from each GPO and then applies that information into the registry. The security CSE reads and applies security policy settings. The Folder Redirection CSE reads and applies Folder Redirection policy settings.

Group Policy processing repeats when the user logs on the computer. The Group Policy service decides the GPOs that apply to the user and then applies user policy settings.

It's important that you have a firm understanding of how to create, modify, and link Group Policy objects to containers in Active Directory. Group Policy Preferences use the same concepts as Group Policy. In fact, you manage Group Policy Preferences the same way that you manage Group Policy. This is a review of Group Policy; it's not complete. If you are unfamiliar with how to manage Group Policy or you need a thorough refresher, then you can read the Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista (Microsoft Press 2008).

Client-side Extensions

A Group Policy client-side extensions is an isolated component that is responsible for processing specific policy settings delivered by the Group Policy infrastructure. The format in which each Group Policy client-side extension saves data can be unique to each extension. And, the Group Policy infrastructure is unaware of this format, nor does it care. Group Policy's purpose is to deliver settings to the computer where each client-side extension applies their portion of the policy settings from multiple Group Policy objects.

To help understand the relationship between the Group Policy infrastructure and the Group Policy client-side extensions-- consider a postal carrier. The postal carrier collects information from various sources and delivers that information to you. The postal carrier has no idea what information they are delivering. The information could be a letter, a DVD, or a CD with photos. The postal carrier only knows they are to deliver the information to a specific address.

In this analogy, the Group Policy service is the postal carrier-- it delivers the information without out any knowledge about the information. The information delivered by the postal carrier represents the different policy settings. The Group Policy client-side extension represents the person receiving the information. Addresses can have many recipients. Each recipient receives their own mail in an expected format. The Group Policy client side extension reads its respective policy setting information and performs actions based on information contains in the policy settings.

Group Policy Processing

Group Policy application is the process of deciding which Group Policy objects that Windows applies to a user or computer and then applying those settings. Understanding Group Policy processing is key to planning and deploying Group Policy settings. Misunderstanding Group Policy processing is the most common cause of unwanted and unexplainable policy settings.

The key to understanding Group Policy processing is Scope. **Scope** is simply a collection of all Group Policy objects that should apply to a user or computer based on their object's location in

Active Directory. You create scope by **linking** Group Policy objects to specific locations within Active Directory.

The key to understanding Group Policy processing is Scope. **Scope** is simply a collection of all Group Policy objects that should apply to a user or computer based on their object's location in Active Directory. You create scope by **linking** Group Policy objects to specific locations within Active Directory.

Group Policy provides options that can change the scope of Group Policy object. Changing the scope of Group Policy objects affects which policy settings apply and those that do not. You change the scope of Group Policy using **processing order**, **filtering**, and **link options**.

Scope

Group Policy processing must identity the scope to which it is applying policy settings. Scope is simply states as where the user or computer object resides within the Active Directory hierarchy. The easiest way to discover the scope of a user or computer object is to lookup the respective user or computer's distinguished name in Active Directory. An object's distinguished name in a directory provides the objects identity and the objects location within the directory. Consider the following distinguished name.

```
CN=Kim Akers,OU=Human Resources, DC=corp,DC=contoso,DC=com
```

From this, the Group Policy service determines the name of the user object, the organizational unit that contains the user object, and the domain in which the user object resides.

```
CN=Jeff Low,OU=Managers,OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com
```

Linking

Understanding Group Policy scope requires knowing where to link Group Policy objects so they apply to users or computer. To enable a Group Policy object to apply to a user or computer, you associate it with a specific location within Active Directory. Associating a Group Policy object with an object in Active Directory is called linking.

Active Directory has rules that govern where you can link Group Policy objects. Active Directory objects to which you can link Group Policy objects include:

- Site objects
- Domain objects
- Organizational Unit objects

Linking Group Policy objects to these Active Directory objects is strategic in deploying Group Policy. These are container objects. Container objects, as the name implies, means they can include other objects within them-- they representing hierarchical grouping of objects in a directory. Site objects can contain computer objects from multiple domains. Domain objects can contain multiple Organizational Units, computers and user objects. Organizational Unit objects can contain other Organizational Unit objects, computers, and users. Let's look at the distinguished name again.

```
CN=Jeff Low,OU=Managers,OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com
```

Close examination of the distinguished name reveals each container object that could potentially apply Group Policy settings to the user. The CN=Jeff Low is the user object name. You cannot link Group Policy directly to a user object. However, the remaining portion of the name shows the object's location. Working left to right, you can discover each container object that is capable of apply Group Policy to the user.

```
OU=Managers,OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com

OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com

OU=RandD,DC=corp,DC=contoso,DC=com

DC=corp,DC=contoso,DC=com
```

Each of these locations represent the scope of Group Policy. The Group Policy service collects linked Group Policy objects from each of these locations in the directory. This represents the scope of Group Policy for the user or computer.

Notice the order in which Windows collects the list of Group Policy objects? It begins with the OU closest to the user and traverses up the directory to the object furthest away from the user, which is typically the domain object. Through linking, you have a list of Group Policy objects that are in scope with the user or computer. However, not every GPO in the list should apply to the user or computer.

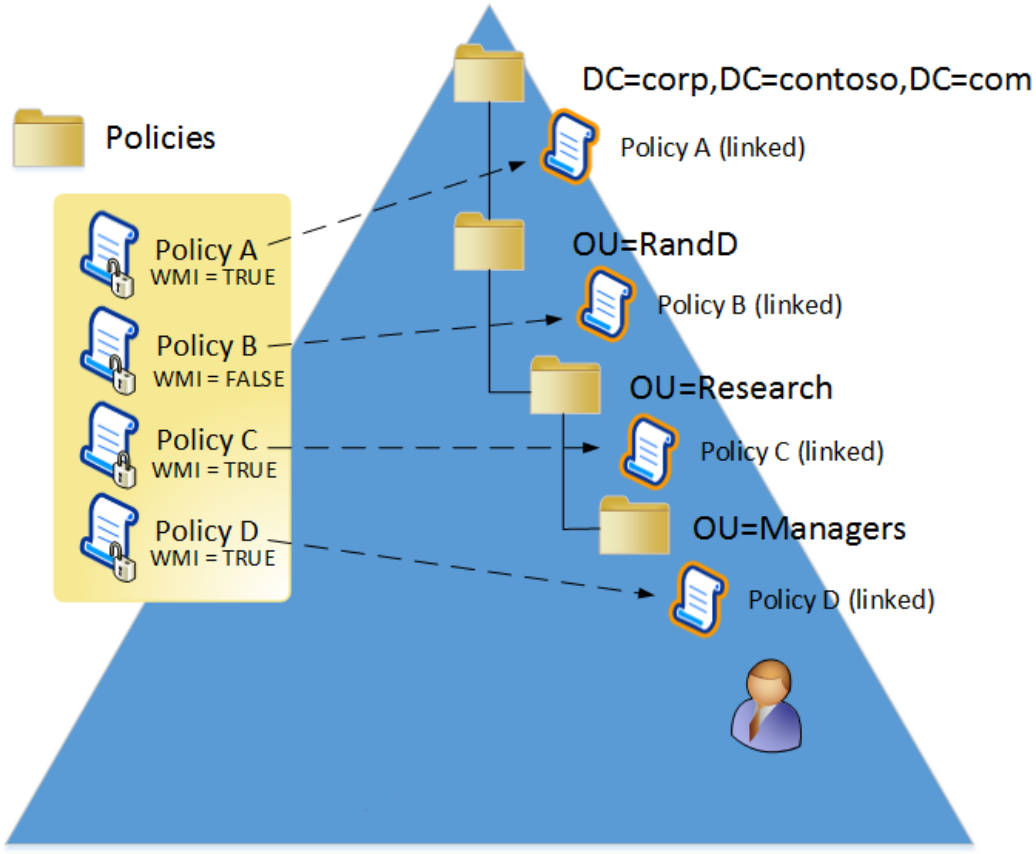

Security Filtering

Group Policy scope is the list of all Group Policy objects that may be applicable to the user or computer because of their object's location within Active Directory. Security Filtering determines if the respective user or computer has the proper permissions to apply the Group Policy object. A user or computer must have the **Read** and **Apply Group Policy** permissions for the Group Policy service to consider the Group Policy object applicable to the user.

The Group Policy services iterates through the entire list of Group Policy objects determining if the user or computer has the proper permissions to the GPO. If the user or computer has the permissions to apply the GPO, then the Group Policy service moves that GPO into a filtered list of GPOs. It continues to filter each Group Policy object based on permissions until it reaches the end of the list. The filtered list of Group Policy objects contains all GPOs within scope of the user or computer and are applicable to the user or computer based on permissions.

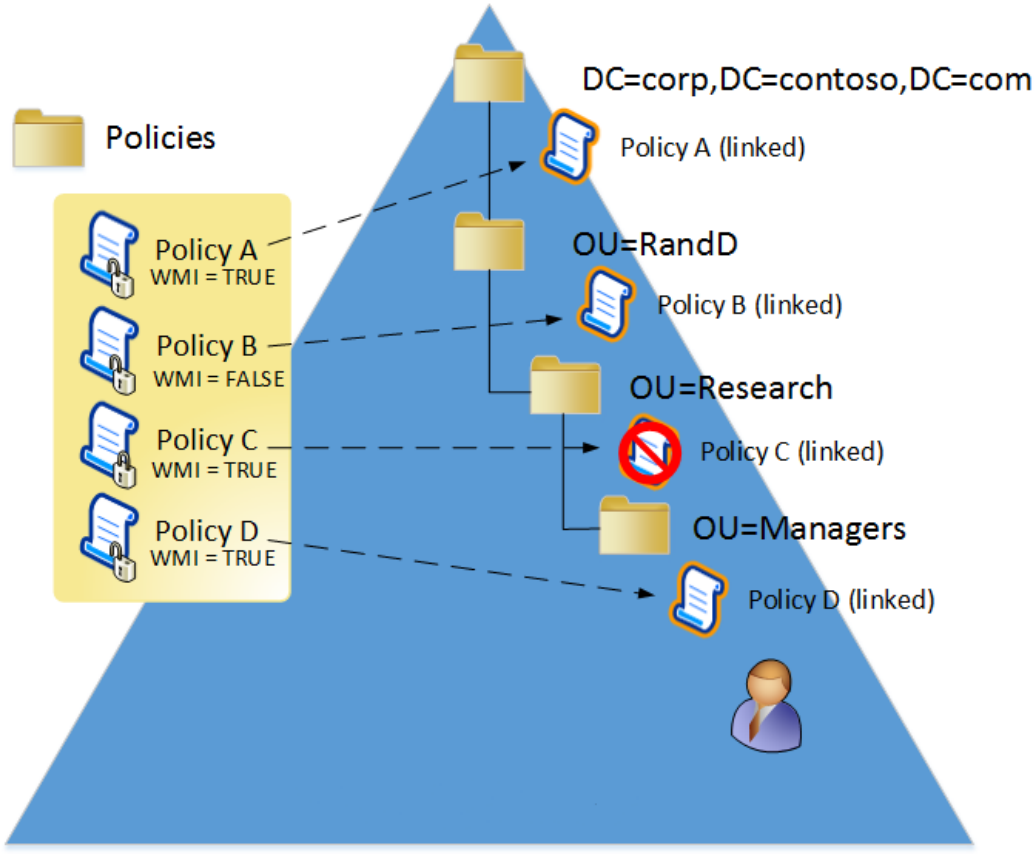

WMI Filtering

WMI filtering is the final phase of determining the scope of Group Policy objects that apply to a user or computer.

Windows Management Instrumentation (WMI) is the Microsoft implementation of Web-Based Enterprise Management (WBEM). WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components.

Group Policy provides more filters to control the scope of applicable Group Policy objects. WMI enables you to create queries to interrogate specific features of the computer, operating system, and other managed components. In the form of queries, you create criteria that behave like logical expressions-- where the result equates to true or false. You associated, or link these criteria to a Group Policy object. If the criteria evaluates to true, the Group Policy object remains applicable to the user and is kept in the filtered list. If the criteria evaluates to false, the Group Policy service removes the Group Policy object from the filtered list.

Once WMI filtering completes, the Group Policy service has a list of filter Group Policy objects. This final list represents all applicable Group Policy objects for the user or computer. Internally, Security and WMI filtering occur in one cycle.

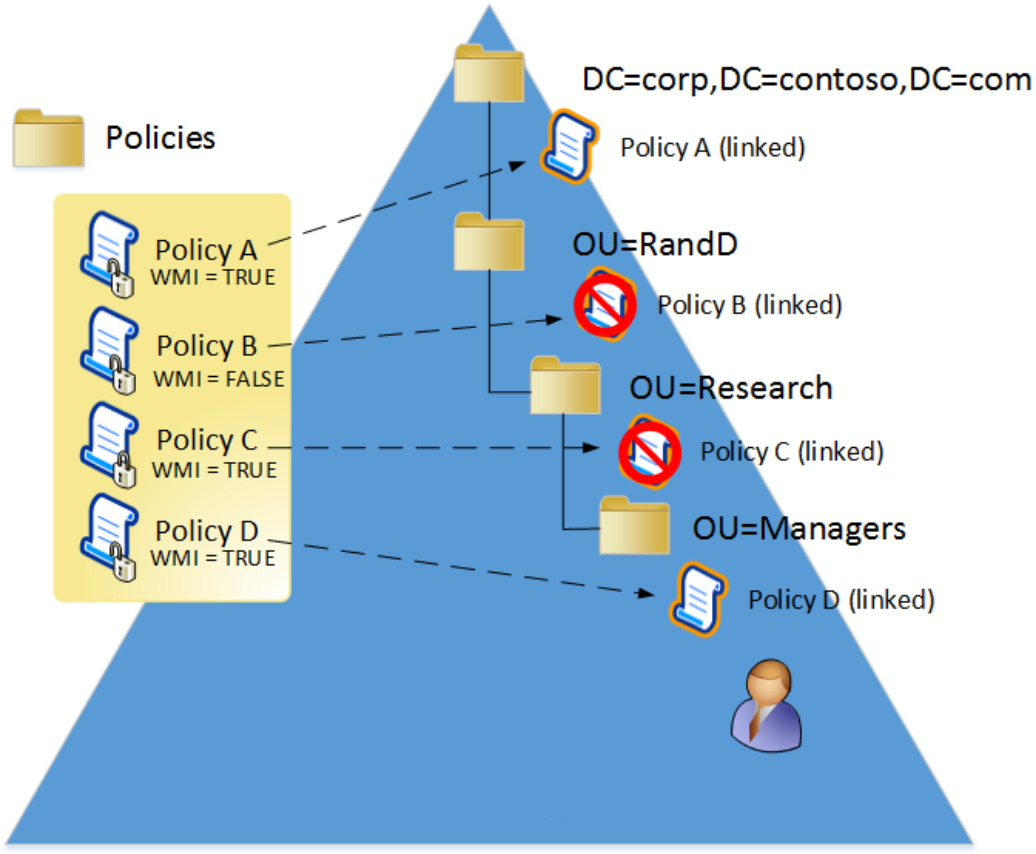

Processing Order

Group Policy has a specific order in which it applies Group Policy objects. Understanding the order in which Group Policy objects apply is important because Group Policy uses the order of application to resolve conflicting policy settings among different Group Policy objects linked to different locations within Active Directory.

Local, Site, Domain, and OU

The Group Policy service applies the Local Group Policy first, then Group Policy objects from the Site, followed by Group Policy objects from the domain, and Group Policy objects from organization units. If the targeted user or computer to receive Group Policy settings, then the Group Policy service applies Group Policy objects from OUs furthest in lineage from the user to closest in lineage to the user. Consider the filtered list of applicable Group Policy objects.

```
DC=corp,DC=contoso,DC=com

OU=RandD,DC=corp,DC=contoso,DC=com

OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com

OU=Managers,OU=Research,OU=RandD,DC=corp,DC=contoso,DC=com
```

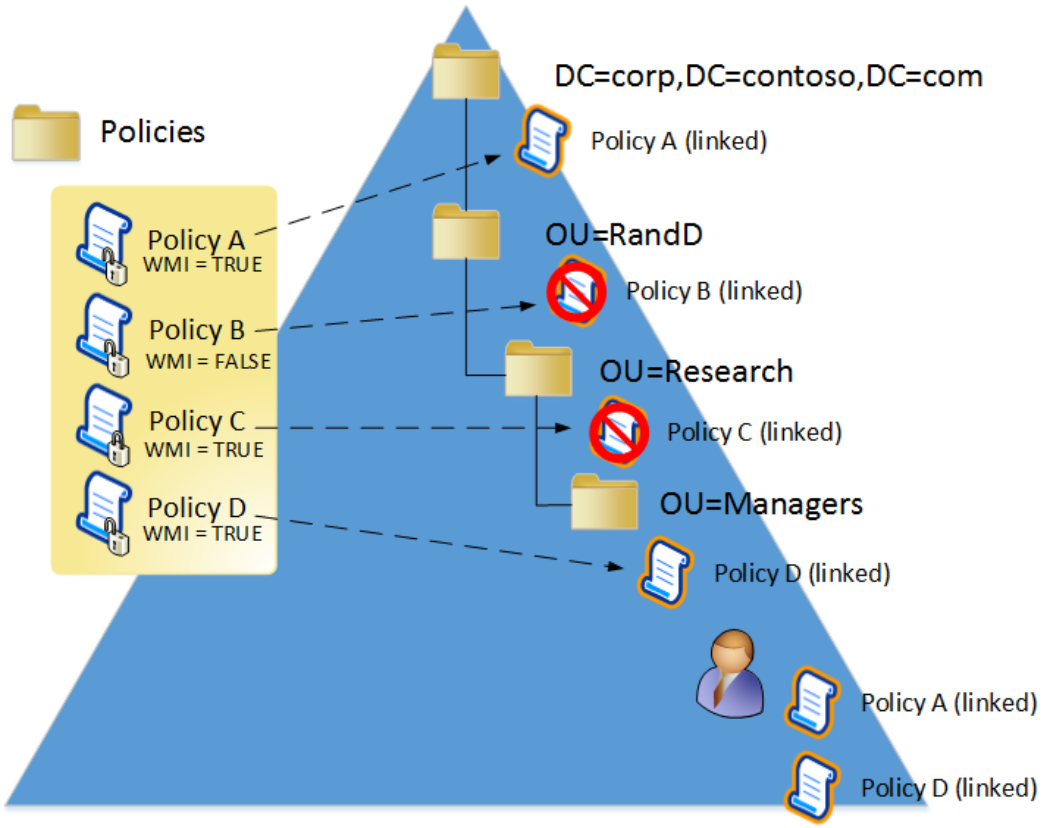

Notice the order of Group Policy objects has changed from the first list. This reordering of Group Policy occurs during the Security and WMI filter processing. The Group Policy service builds the first list of GPOs by finding the user or computer object and then collecting all linked GPOs as it walks up the directory tree. The GPOs are listed backwards from the order they apply because as the Group Policy service adds the newly discovered link location to the bottom of the list. This explains why the domain location is at the bottom of the list.

However, when filtering the list for security and WMI filters, the Group Policy service starts at the top of the list, which is the OU closest in lineage to the user or computer object. The service builds a new list (the filtered list) by placing the GPOs that pass through the filter into the filtered list. The service inverts the order of the original list, making the domain location at the top of the list. The location closest to the user is at the bottom of the list —the order Group Policy applies GPOs to users and computers.

Conflict Resolution

Each Group Policy object contains the same number of potential policy settings. Therefore, it is possible to have the same policy setting defined in multiple Group Policy objects. Conflicts occurs when the same policy setting is configured in multiple Group Policy objects. Like two cars competing for the same space on the road—one wins and the other loses. Group Policy handles conflicts by using a method known as last-writer-wins. Last-writer-wins resolves conflicts by declaring the prevailing setting as the setting that Group Policy writes last. Therefore, the Group

Policy object containing the conflicting policy setting that applies last is the setting that wins over all other settings.

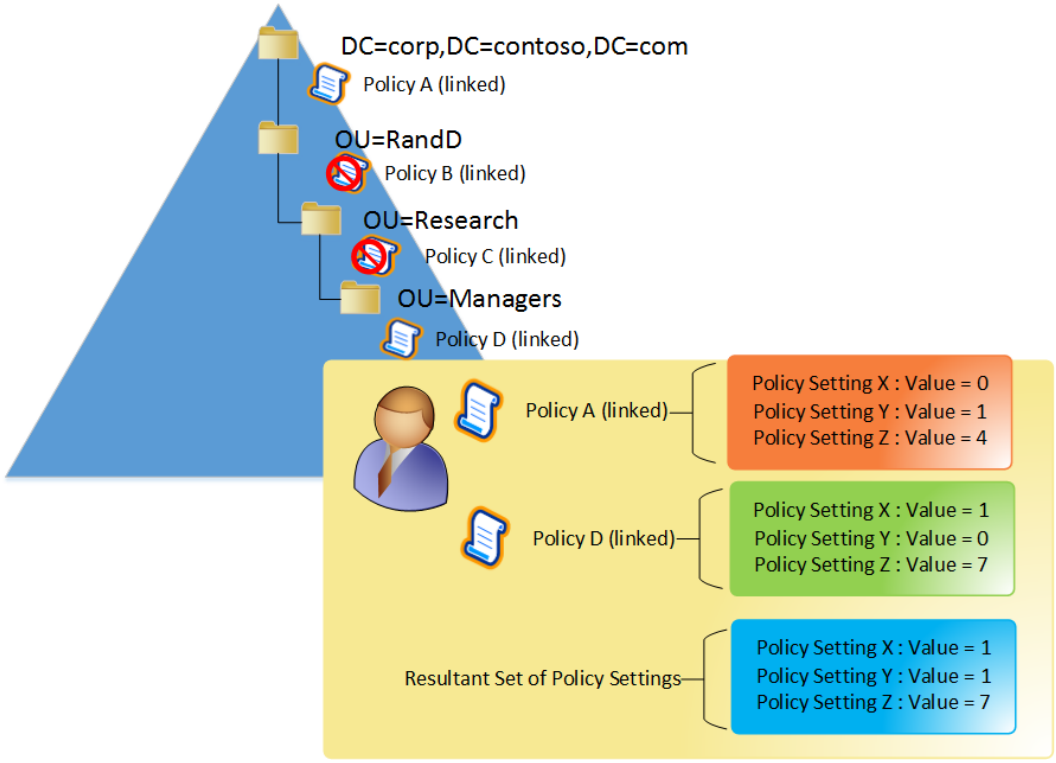

The Processing Order section of this document describes that Group Policy objects apply in Local, Site, Domain, and Organizational Unit order. Based on this processing hierarchy:

- Policy settings in Group Policy objects linked to the Active Directory site resolve policy setting conflicts between the Local Group Policy object and Group Policy objects linked to the Active Directory site.
- Policy settings in GPOs linked to the domain resolve policy setting conflicts between Group Policy objects linked to the Active Directory site and GPOs linked to the Active Directory domain.
- Policy settings in GPOs linked to an organizational unit resolve policy setting conflicts between Group Policy objects linked to the Active Directory domain and GPOs linked to an organizational Unit.
- Policy settings in GPOs linked to a child organizational unit resolve policy settings conflicts between Group Policy objects linked to the child organizational unit and GPOs linked to the parent organizational unit.

Conflict Resolution among GPOs linked at the same Location

Group Policy enables you to link multiple Group Policy objects at each site, domain, and organization unit locations in the directory. Until now, conflict resolution only identified resolutions between conflicting policy settings linked at two different locations in Active Directory. What about conflicting policy settings in Group Policy objects that are linked at the same location?

Group Policy continues to use the last-writer-wins method for resolving policy setting conflicts among Group Policy objects linked as the same location in Active Directory. Understanding how the Group Policy Management Console (GPMC) links Group Policy objects to locations in Active Directory explains the processing order of Group Policy objects link at the same location in Active Directory.

GPLink Attribute

The locations that support Group Policy linking, Active Directory sites, domains, and organizational units, do so because each of these objects have a GPLink attribute. The GPLink attribute is a single-valued attribute that accepts a value of a string data type. While the Active Directory Schema enforces the single-valued nature of the GPLink attribute, Group Policy uses the attribute as a multivalued attribute. The GPMC writes the value of the GPLink attribute using the following format.

```
[distinguishedNameOfGroupPolicyContainer;linkOPtions][…][…]
```

The distingushedNameOfGroupPolicyContainer token represents the distinguished name of the Group Policy Container. A Group Policy object is a single logical object composed of two components of information. The component of information stored on the file system is the Group Policy template. The remaining component, the Group Policy Container is an object in Active Directory object that lives in the domain partition of Active Directory. As previously covered, the distinguished name of a directory object provides the object's name and location in the directory.

The linkOptions token is an integer value that defines the link options associated with the Group Policy object. Currently, you can enable or disable linked of Group Policy objects. Also, you can configure the link as enforced. The linkOptions value is a bit value where combining values varies the configurations.

```
Enabled0x0

Disabled 0x1

Enforced0x2
```

Disabling the link of a Group Policy objects prevents the Group Policy service from including that GPO in the list of GPOs within scope of the targeted user or computer. The distinguishedNameOfGroupPolicyContainer and the linkOptions token are enclosed in square brackets ([ ]) and separated by a semicolon (;). This represents a singly linked Group Policy object. Linking another Group Policy object to the location inserts a new distingushedNameOfGroupPolicyContainer and linkOptions combination before the existing combination; it does not add the new combination to the end. The linking pattern continues to insert newly linked GPOs at the beginning of the value; by moving existing values to the right.

The Group Policy service reads this long string as a list of values from left to right. The first GPO link entry in the value is the first to apply at this location. The next entry in the value applies afterwards. The process continues until the last GPO in the value applies.

Group Policy inherently assigns each GPO precedence based on the order it reads the list—left to right. Therefore, the first GPO in the value has the lowest precedence in the list of linked Group Policy objects. The next GPO in the value has a higher precedence than the previous GPO because it applies its policy settings after the previous GPO; by winning any policy setting

conflicts between the two GPOs. Each GPO that follows has a higher precedence than the Group Policy object before it in the link order. The last GPO in the value has the highest precedence because it is the last Group Policy object the Group Policy service applies.

The best way to understand this is to think of the long string as a list of GPOs. Take the first GPO (the left most GPO) in the value and place it the list. Take the next links GPO listed and place on top of the list (causing all others to move down in the list by one). Continue this process until the last GPO is on top of the list. This final GPO linked entries list is in precedence order, which means the list is processed from the bottom to the top.

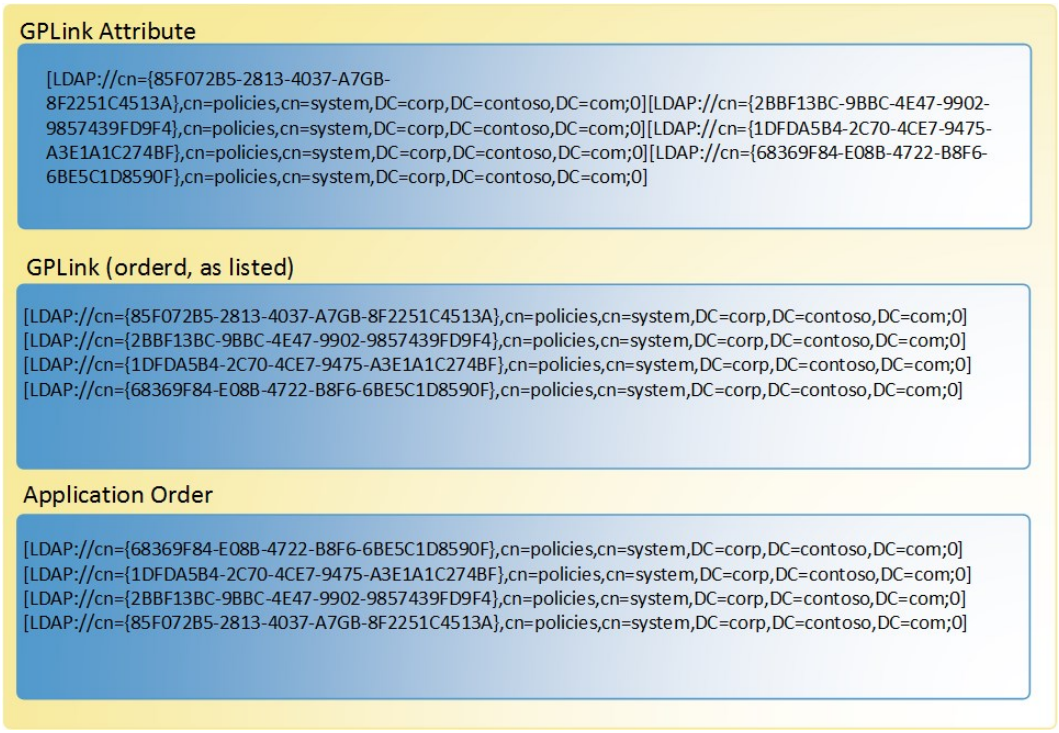

When viewed in a list in precedence order, it's easy to discover that GPOs higher in the list have more precedence than GPOs lower in the list. As a result, GPOs lower in the list lose policy setting conflicts and GPOs higher in the list win policy setting conflicts.

Link Options

As previously stated, a Group Policy linked as options of enabled, disabled, and enforced. The enabled and disabled options are intuitive to understand. When an enabled link is considered in the scope of Group Policy for the targeted user or computer. A disabled linked behaves as if the Group Policy object was never linked.

Enforced

The Enforced link option is the exception to all rules. The Enforced option ensures the settings from the linked GPO always win conflicts regardless of any other Group Policy object that contains policy settings that may conflict with those of the linked GPO. The GPMC visually represents an enforced Group Policy link by adding a padlock to the existing linked policy icon.

Group Policy settings from an enforced link always apply, even if the organizational unit has block policy inheritance enabled

Block Policy Inheritance

The last item about Group Policy processing order is Block Policy Inheritance, or simply known as Block Inheritance in the Group Policy Management Console. Each domain and organizational unit in Active Directory object contains a **GPOptions** attribute. This setting enables you to block Group Policy settings linked higher in the processing order from applying to users and computers that are typically in containers lower in the processing order.

For example, policy settings linked to the domain apply to computers and users within the entire domain, regardless of their parent organizational unit. However, you can use GPMC to block inheritance on the domain or an organizational unit to prevent normal Group Policy setting from applying to users and computers within that container. Blocking policy inheritance on the domain prevents Group Policy settings from GPOs linked to the Active Directory site from applying to the domain. Blocking policy inheritance on organizational units prevents normal Group Policy settings from GPOs linked to sites and domains from applying to the organizational units.

Block policy inheritance does not prevent Group Policy settings from enforced linked Group Policy objects from applying to users and computers. Group Policy settings from enforced links apply regardless of the block policy inheritance status on domain and organizational unit objects.

Group Policy Preferences

Group Policy Preferences extends Group Policy. Preferences are not Group Policy settings. Windows stores both settings in the registry; however; policy settings have an advantage over preferences—they typically override a preference.

You can configure Windows using the user interface. The user interface presents you with choices; you choose the options you like; and click OK or close the dialog box. Windows then saves your choices to the registry so it can recall those settings later. Settings configurable by the user are known as preferences (notice the lowercase “p”). Mapping a shared folder or choosing a default home page is an example of preferences. When you set the home page using Internet Explorer, you can close the web browser and open it up again and it remembers your home page. Policy settings differ from preferences because policy settings are enforced on the user or computer. Policy prevents the user from changing their settings. Typically, users configure preferences.

Group Policy Preferences enables you to deploy desired configurations to computers and users without limiting the user from choosing a different configuration. It is important to remember that while the user can change the configuration, Group Policy Preferences are Group Policy client-side extensions. Group Policy Preferences refresh with Group Policy; therefore, Group Policy overwrites any preference settings altered by the user with the value configured in a Group Policy Preference. Replacing a user configured preference setting with one configured using Group Policy Preferences is not the same as Group Policy. A true Group Policy setting enforces the setting and restricts the user from changing the setting. Users can easily change preference

values enabled by Group Policy Preferences until the next refresh of Group Policy (which returns the preference settings back to the value configured in the Group Policy Preference item).

Client-side Extensions

Group Policy Preferences are Group Policy client-side extensions. There are 20 extensions that makes up Group Policy Preferences. These extensions include

Client Side Extension	Description
Group Policy Environment	Create, modify, or delete environment variables.
Group Policy Local Users and Groups	Create, modify, or delete local users and groups.
Group Policy Device Settings	Enable or disable hardware devices or classes of devices.
Group Policy Network Options	Create, modify, or delete virtual private networking (VPN) or dial-up networking (DUN) connections.
Group Policy Drive Maps	Create, modify, or delete mapped drives, and configure the visibility of all drives.
Group Policy Folders	Create, modify, or delete folders.
Group Policy Network Shares	Create, modify, or delete network shares
Group Policy Files	Copy, modify the attributes of, replace, or delete files.
Group Policy Data Sources	Create, modify, or delete Open Database Connectivity (ODBC) data source names.
Group Policy INI Files	Add, replace, or delete sections or properties in configuration settings (.ini) or setup information (.inf) files.
Group Policy Folder Options	Create, modify, or delete folders.
Group Policy Schedule Tasks	Create, modify, or delete scheduled or immediate tasks.
Group Policy Registry	Copy registry settings and apply them to other computers. Create, replace, or delete registry settings.
Group Policy Printers	Create, modify, or delete TCP/IP, shared, and local printer connections.

Group Policy Shortcuts	Create, modify, or delete shortcuts.
Group Policy Internet Settings	Modify user-configurable Internet settings
Group Policy Start Menu Settings	Modify Start menu options.(Not applicable for Windows 8 and Windows Server 2012)
Group Policy Regional Options	Modify regional options.
Group Policy Power Options	Modify power options and create, modify, or delete power schemes.
Group Policy Applications	Configure settings for applications.

Common Configurations

Most Group Policy Preference items share a common configuration that enable you to control the scope of Group Policy Preference processing for each configured preference item.

Stop processing items in this extension if an error occurs on this item

Each preference extension can contain one or more preference items. By default, a failing preference item does not prevent other preference items in the same extension from processing.

If the **Stop processing items in this extension if an error occurs on this item** option is selected, a failing preference item prevents remaining preference items within the extension from processing. This change in behavior is limited to the hosting Group Policy object (GPO) and client-side extension. It does not extend to other GPOs.

It's important to understand that Group Policy Preference extensions process preference items from the top of the list and work their way to the bottom. The preference extension only stops processing preference items that follow the failing preference item (items appearing below the failing preference items as they appear in the list).

Run in logged-on user's security context (user policy option)

There are two security contexts in which Group Policy applies user preferences: the SYSTEM account and the logged-on user.

By default, Group Policy processes user preference items using the security context of the SYSTEM account. In this security context, the preference extension is limited to environment variables and system resources available only to the computer.

If the **Run in logged-on user's security context** option is selected, it changes the security context under which the preference item is processed. The preference extension processes preference items in the security context of the logged-on user. This allows the preference extension to access resources as the user rather than the computer. This can be important when using drive maps or other preferences in which the computer may not have permissions to resources or when using environment variables. The value of many environment variables differ when evaluated in a security context other than the logged-on user.

Group Policy Preference extensions that need to process in the user's security context, such as Drive Maps and Printers **automatically** switch to the user's context and do not need you to adjust this setting.

Remove this item when it is no longer applied

Group Policy applies policy settings and preference items to users and computers. You decide which users and computers receive these items by linking one or more Group Policy objects (GPOs) to Active Directory sites, domains, or organizational units. User and computer objects in these containers receive policy settings and preference items defined in the linked GPOs because they are within the scope of the GPO.

Unlike policy settings, the Group Policy service does not remove preference settings when the hosting GPO becomes out of scope for the user or computer.

If the **Remove this item when it is no longer applied** option is selected, it changes this behavior. After selecting this option, the preference extension decides if the preference item should not apply to targeted users or computers (out of scope). If the preference extension decides the preference item is out of scope, it removes the settings associated with the preference item.

Selecting this setting changes the preference item's action to **Replace**. During Group Policy application, the preference extension recreates (deletes and creates) the results of the preference item. When the preference item is out of scope for the user or computer, the results of the preference item are deleted, but not created. Preference items can become out of scope by using item-level targeting or by higher-level Group Policy filters such as WMI and security group filters.

The **Remove this item when it is no longer applied** option is not available when you set the preference item action to Delete.

Apply once and do not reapply

Preference items apply when Group Policy refreshes.

By default, the results of preference items are rewritten each time Group Policy refreshes. This ensures the preference item results are consistent with what you configured in the Group Policy object.

If the **Apply once and do not reapply** option is selected, it changes this behavior, so the preference extension applies the results of the preference item to the user or computer only once. This option is useful when you do not want the results of a preference item to reapply.

Item-level Targeting

Group Policy provides filters to control which policy settings and preference items apply to users and computers. Preferences provide an added layers of filtering called targeting. Item-level targeting enables you to control if a preference item applies to a group of users or computers.

Use item-level targeting to change the scope of individual preference items, so they apply only to selected users or computers. Within a single Group Policy object (GPO), you can include multiple preference items—each customized for selected users or computers and each targeted to apply settings only to the relevant users or computers.

Each targeting item results in a value of either true or false. You can apply multiple targeting items to a preference item and select the logical operation (AND or OR) by which to combine each targeting item with the preceding one. If the combined result of all targeting items for a preference item is false, then the settings in the preference item are not applied to the user or computer. Using targeting collections, you can also create parenthetical expressions.

Battery Present

A Battery Present targeting item allows a preference item to be applied to computers or users only if one or more batteries are present in the processing computer. If Is Not is selected, it allows the preference item to be applied only if the processing computer does not have one or more batteries present.

If an uninterruptible power supply (UPS) is connected to the processing computer, a Battery Present targeting item may detect the UPS and identify it as a battery.

Computer Name

A Computer Name targeting item allows a preference item to be applied to computers or users only if the computer's name matches the specified computer name in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the computer's name does not match the specified computer name in the targeting item.

CPU Speed

A CPU Speed targeting item allows a preference item to be applied to computers or users only if the processing computer's CPU speed is greater than or equal to the value specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the processing computer's CPU speed is less than or equal to the value specified in the targeting item.

Date Match

A Date Match targeting item allows a preference item to be applied to computers or users only if the day or date matches that specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the day or date does not match that specified in the targeting item.

Dial-up Connection

A Dial-Up Connection targeting item allows a preference item to be applied to users only if a network connection of the type specified in the targeting item is connected. If Is Not is selected, it allows the preference item to be applied only if no network connection of the type specified in the targeting item is connected.

Dial-Up Connection targeting items detect whether a type of network connection exists, not whether the user is logged on through a connection of that type.

Disk Space

A Disk Space targeting item allows a preference item to be applied to computers or users only if the processing computer's available disk space is greater than or equal to the amount specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the

processing computer's available disk space is less than or equal to the amount specified in the targeting item.

Domain

A Domain targeting item allows a preference item to be applied to computers or users only if the user is logged on to or the computer is a member of the domain or workgroup specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the user is not logged on to or the computer is not a member of the domain or workgroup specified in the targeting item.

Environment Variables

An Environment Variable targeting item allows a preference item to be applied to computers or users only if the environment variable and value specified in the targeting item are equal. If Is Not is selected, it allows the preference item to be applied only if the environment variable and value specified in the targeting item are not equal or if the environment variable does not exist.

If you want to restrict the scope of multiple preference items with a complex set of targeting items, you can simplify configuration by using an environment variable. For example, create an Environment Variable preference item that generates a new environment variable with a value of 1, and apply the targeting items to it. To apply the same targeting to other preference items, add an Environment Variable targeting item to those preference items, and configure it to require a value of 1 for the variable that you created using an Environment Variable preference item.

File Match

A File Match targeting item allows a preference item to be applied to computers or users only if the file or folder specified in the targeting item exists, or only if the file exists and is a version within the range specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the file or folder specified in the targeting item does not exist, or only if the version of the file is not within the range specified in the targeting item.

IP Address Match

An IP Address Range targeting item allows a preference item to be applied to computers or users only if the processing computer's IP address is within the range specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the processing computer's IP address is not within the range specified in the targeting item.

Language

A Language targeting item allows a preference item to be applied to computers or users only if the locale specified in the targeting item is installed on the processing computer. Additional options allow you to restrict the targeting to the user's or computer's locale. If Is Not is selected, it allows the preference item to be applied only if the processing computer's locale does not match the specified locale in the targeting item.

A locale is composed of a language and, in some cases, a geographic area in which the language is spoken or the alphabet used. For example, French (Canada) is a locale composed of the language French and the geographic area Canada.

LDAP Query

An LDAP Query targeting item allows a preference item to be applied to computers or users only if the LDAP query returns a value for the attribute specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the LDAP query does not return a value for the attribute specified in the targeting item.

MAC Address Range

A MAC Address Range targeting item allows a preference item to be applied to computers or users only if any of the processing computer's MAC addresses are within the range specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if none of the processing computer's MAC addresses are not within the range specified in the targeting item.

Range starting points and ending points are inclusive. You can specify a single address by typing the same value in both boxes.

MSI Query

An MSI Query targeting item allows a preference item to be applied to computers or users only if certain aspects of an MSI installed product, update, or component on the processing computer match the specified criteria in the targeting item. If Is Not is selected, it allows the preference item to be applied only if certain aspects of an MSI installed product, update, or component on the processing computer do not match the specified the specified criteria in the targeting item.

Operating System

An Operating System targeting item allows a preference item to be applied to computers or users only if the processing computer's operating system's product name, release, edition, or computer role matches those specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the operating system's product name, release, edition, or computer role does not match those specified in the targeting item.

Organizational Unit

An Organizational Unit targeting item allows a preference item to be applied to computers or users only if the user or computer is a member of the organizational unit (OU) specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the user or computer is a not member of the OU specified in the targeting item.

PCMCIA Present

A PCMCIA Present targeting item allows a preference item to be applied to computers or users only if the processing computer has at least one PCMCIA slot present. If Is Not is selected, it allows the preference item to be applied only if the processing computer does not have any PCMCIA slots present.

A PCMCIA slot is considered present when the drivers for the slot are installed and the slot is functioning correctly.

Portable Computer

A Portable Computer targeting item allows a preference item to be applied to computers or users only if the processing computer is identified as a portable computer in the current hardware profile on the processing computer or if the processing computer is identified as a portable computer

with the docking state specified in the targeting item. When Is Not is selected, it allows the preference item to be applied only if the processing computer is not identified as a portable computer in the current hardware profile on the processing computer or if the docking state of the processing computer differs from the docking state specified in the targeting item.

Processing Mode

A Processing Mode targeting item allows a preference item to be applied to computers or users only if the Group Policy processing mode or conditions on the processing computer match at least one of those specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the Group Policy processing mode or conditions on the processing computer do not match any of those specified in the targeting item.

RAM

A RAM targeting item allows a preference item to be applied to computers or users only if total amount of physical memory in the processing computer is greater than or equal to the amount specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the total amount of physical memory in the processing computer is less than the amount specified in the targeting item. Provide the total amount of physical memory in megabytes (MB). One gigabyte (GB) of physical memory is entered as 1024. Four gigabytes of physical memory are entered as 4096.

Registry Match

A Registry Match targeting item allows a preference item to be applied to computers or users only if the registry key or value specified in the targeting item exists, if the registry value contains the data specified in the targeting item, or if the version number in the registry value is within the range specified in the targeting item. If the targeting item allows the preference item and if Get value data is selected in the targeting item, then the targeting item saves the value data of the specified registry value to the environment variable specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the registry key or value specified in the targeting item does not exist, if the registry value does not contains the data specified in the targeting item, or if the version number in the registry value is not within the range specified in the targeting item.

Security Group

A Security Group targeting item allows a preference item to be applied to computers or users only if the processing computer or user is a member of the group specified in the targeting item and optionally only if the specified group is the primary group for the processing computer or user. If Is Not is selected, it allows the preference item to be applied only if the processing computer or user is not a member of the group specified in the targeting item and optionally only if the specified group is not the primary group for the processing computer or user.

Security Group

- Domain groups
    - Domain local
    - Global groups
    - Universal groups

- Local groups
  - Local groups (including built-in groups)
  - Well-known

Site

A Site targeting item allows a preference item to be applied to computers or users only if the processing computer is in the site in Active Directory specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the processing computer is not in the site in Active Directory specified in the targeting item.

Targeting Collection

The targeting items applied to a preference item are evaluated as a logical expression. A targeting collection allows you create a parenthetical grouping within that expression. You can nest one targeting collection within another to create more complex logical expressions.

A targeting collection allows a preference item to be applied to computers or users only if the collection of targeting items specified results in a value of true. If Is Not is selected, it allows the preference item to be applied only if the collection of targeting items specified results in a value of false.

Terminal Session

A Terminal Session targeting item allows a preference item to be applied to users only if the processing user is logged on to a terminal services session with the settings specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the user is not logged on to a terminal services session or the user is logged on to a terminal services session without the settings specified in the targeting item.

Time Range

A Time Range targeting item allows a preference item to be applied to computers or users only if the current time on the end user's computer is within the time range specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the current time on the end user's computer is not within the range specified in the targeting item.

User

A User targeting item allows a preference item to be applied to users only if the processing user is the user specified in the targeting item. If Is Not is selected, it allows the preference item to be applied only if the processing user is not the user specified in the targeting item.

WMI Query

A WMI Query targeting item allows a preference item to be applied to computers or users only if the processing computer evaluates the WMI query as true. If Is Not is selected, it allows the preference item to be applied only if the processing computer evaluates the WMI query as false.

Processing

Earlier, this document explained Group Policy processing. Group Policy Preference client-side extensions adhere to these same rules. Therefore, linked hierarchy, security and WMI filtering can change the scope of Group Policy object configured with Group Policy Preferences. By

changing the scope, users and computers may or may not receive settings or preference items configured in these Group Policy objects.

However, Group Policy Preference client-side extensions have their own internal processing. You can configure one or more preference items for a single Group Policy Preference extension to process within a single Group Policy object. For example, you can configure a single GPO to contain 10 Drive Map Preference items within a single GPO.

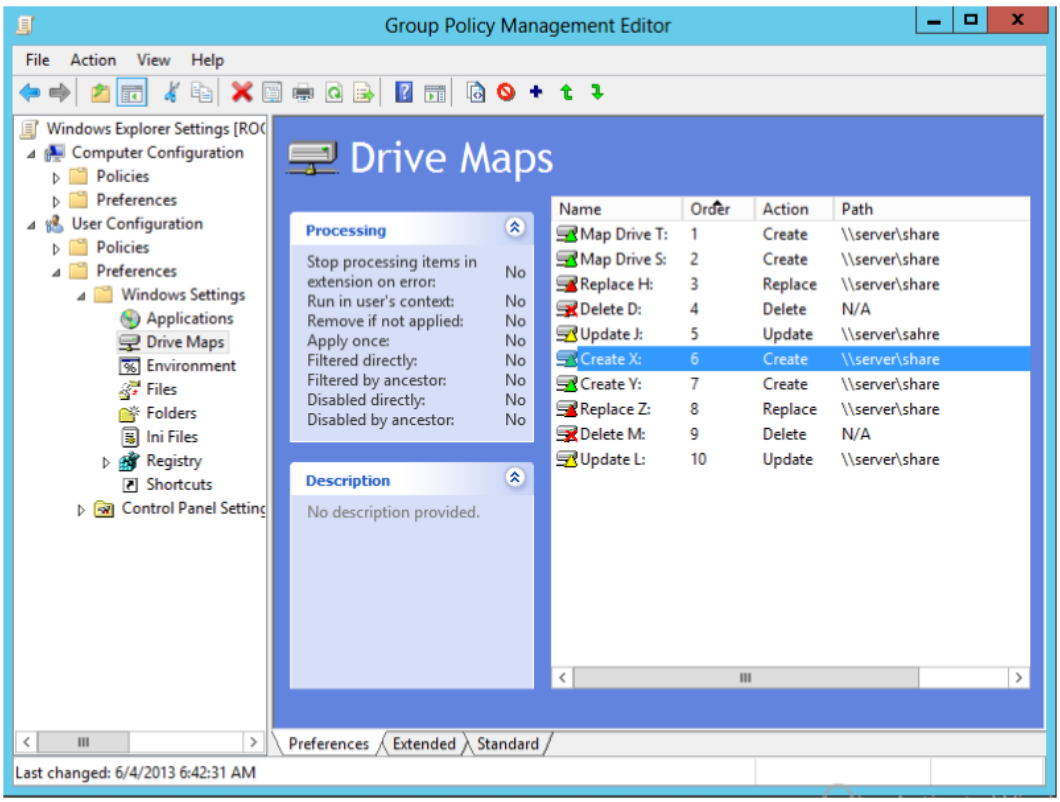

During Group Policy processing, the Group Policy infrastructure cycles through a list of Group Policy extensions. As it moves to each extension, it shares information relevant for the extension to process its portion of Group Policy. Critical components of the information shared with the extensions include a list of Group Policy objects that included changes, a list of Group Policy objects that are no longer in scope with the user or computer. Also, the Group Policy infrastructure provides information specific to this instance of Group Policy processing such as if the network connection is considered a slow link.

The Group Policy Preference extension uses the information about the changed and out-of-scope Group Policy objects to process its policy settings. Group Policy Preference client-side extensions process preference items in order from the top of the list to the bottom of the list.

The results of processing each preference item vary depending on the action configured in the preference item. Also, item-level targeting can prevent the preference item from applying to the user or computer. The Group Policy Preference client-side extension applies each item in the list until it reaches the end of the list, or exits because of a common configuration settings such as

Stop processing items in this extension if an error occurs on this item or Apply once and do not reapply. Once the preference extensions applies all preference items in the list, it returns control to the Group Policy service.

Drive Map

File servers revolutionized centralized storage of user data. However, it also introduced usability problems for end users attempting to remember long and complex paths to access their files and folders. IT Professionals resolved that problem by introducing users to Drive Mappings. Drive maps enable users to simply associate a single drive letter to files and folders that reside on file servers. Drive maps reduced returned usability back to the end user because it removed the requirement of users remembering long and complex paths to access centralized data. Now, when a user needs access to their data, they can simply select the H drive from Windows Explorer, which represents the users HOME folder. Or, they can access the S drive for data that is shared among the department or organization.

While mapped drives returned usability to end users, it created complexity for IT professionals. To facilitate mapped drives for users, IT Professionals created a large and complicated matrix of users who belongs to groups that correlated to specific mapped drives needed by the users. And, this mess of user-to-group-to-mapped-drives complexity was implemented in an elaborate logon script. This limited the number of IT Professionals that could support and update the script to only those that had scripting knowledge. Additionally, complex scripts are difficult to troubleshoot and support. Group Policy Preferences Drive Maps resolve this problem.

Overview

The Group Policy Preference Drive Map extension enables you to configure mapped drives for an enterprise environment without the need of complicated and cumbersome logon scripts. The Drive Map preference extensions is configured using the same tools to configure Group Policy; thereby simplifying management to using the Group Policy Management Editor.

Capabilities

The Group Policy Drive Map extensions provides you with basic drive map needs such as creating, deleting, or replacing drive mappings. In addition, the drive map extension enables updating labels, and the visibility status of existing drive mappings. Other capabilities provided by the Group Policy Drive Map extension includes enabling persistent drive mappings and creating drive mappings using alternate credentials.

You can combine Drive Map preference items with Group Policy Preferences Item-level targeting to achieve granular and complex drive map scenarios that remain simple, easy to manage, and do not use or require knowledge of scripting languages.

Configurable Options

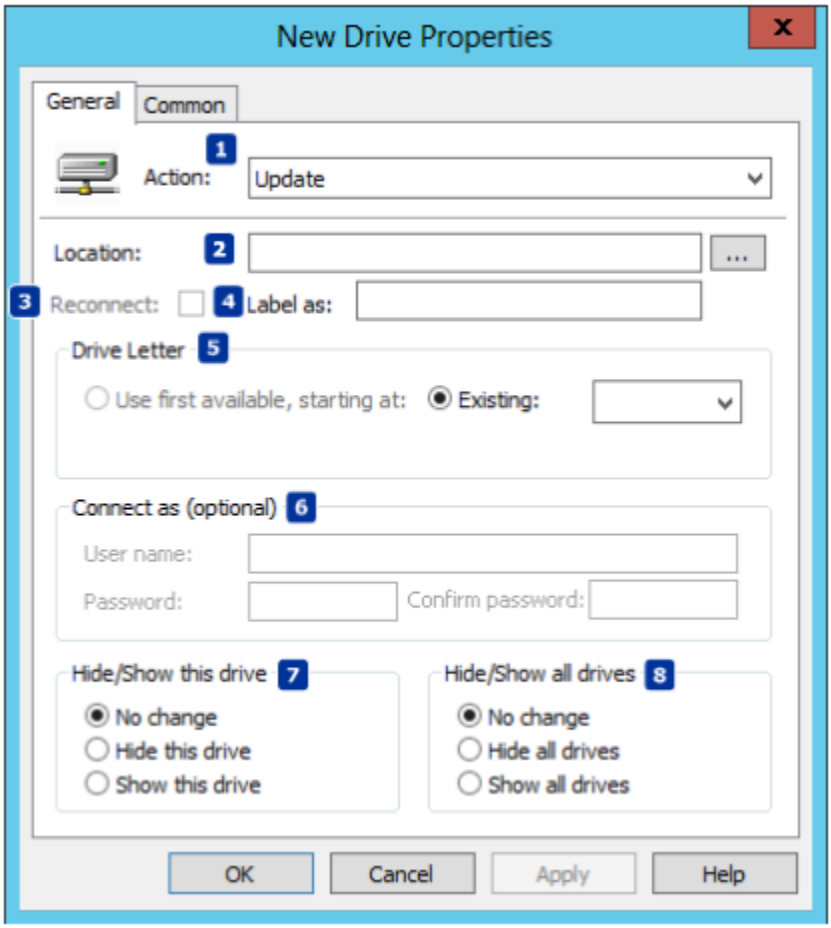

The Group Policy Drive Map preference item user interface provides several choices to help achieve most drive map scenarios that exist in enterprise environments.

Action (1)

The preference item provides a choice of four actions: **Create**, **Replace**, **Update**, and **Delete**. The behavior of the preference item varies with the action selected and whether the drive letter exists.

1. **Create**

   Create a new mapped drive for users.

2. **Delete**

   Remove a mapped drive for users.

3. **Replace**

Delete and **Create** mapped drives for users. The result of the **Replace** action is to overwrite all existing settings associated with the mapped drive. If the drive mapping does not exist, then the **Replace** action creates a new drive mapping.

4. **Update**

   Modify settings of an existing mapped drive for users. This action differs from **Replace** in that it only updates settings defined within the preference item. All other settings remain as configured on the mapped drive. If the drive mapping does not exist, then the **Update** action creates a new drive mapping.

Location (2)

To configure a new drive mapping or recreate a drive mapping, type a fully qualified UNC path for the network share (such as \\server\sharename, \\server\hiddenshare$, or \\server\sharename\foldername). Alternatively, you can use built-in Group Policy Preference variables by pressing F3 and selecting a variable from the Variable browser.

To modify an existing drive mapping, which are matched by drive letter, leave this field blank.

This option is available only if the action selected is **Create**, **Replace**, or **Update**.

Reconnect (3)

To save this mapped drive in the user's settings and attempt to restore it at each subsequent logon, select this check box. Otherwise, the drive is mapped, but not saved in the user's settings.

This option is available only if the action selected is **Create**, **Replace**, or **Update**.

Label as (4)

To provide a descriptive label that appears next to the drive letter, type the label in this field. Alternatively, you can use built-in Group Policy Preference variables by pressing F3 and selecting a variable from the Variable browser.

This option is available only if the action selected is **Create**, **Replace**, or **Update**.

Drive Letter (5)

Select the mapped drives, identified by drive letter, to configure:

To assign the first available drive letter to the mapped drive, select Use first available, starting at, and then select a drive letter at which to begin checking for availability. This option is available only if the action selected is **Create**, **Replace**, or **Update**.

To assign a specific drive letter to the mapped drive, select Use, and then select a drive letter. This option is available only if the action selected is **Create**, **Replace**, or **Update**, and if you have typed a location.

To modify an existing drive mapping (identified by drive letter), select **Existing**, and then select a drive letter. This option is available only if the **Location** field is blank, and the action selected is **Update**.

To delete all drive mappings from a particular drive letter onward, select **Delete all, starting at**, and then select a drive letter at which to begin deleting drive mappings. Physical drives are skipped without error. This option is available only if the action selected is **Delete**.

To delete a specific mapped drive, select **Delete**, and then select the drive letter. This option is available only if the action selected is **Delete**.

Connect as (6)

To implement a drive mapping using credentials other than those of the currently logged on user, type the credentials to be used. This option is available only if the action selected is **Create**, **Replace**, or **Update**.

The Group Policy Drive Map management extension stores alternative credentials in the configuration XML file, which is designed to be read by all domain users. To mitigate storing of passwords in clear text, the management extensions obfuscates the password using a derived Advanced Encryption Standards (AES) key.

Hide/Show this drive (7)

Configure the visibility of the mapped drive:

To keep mapped drive visibility as it is currently configured, select **No change**. This does not take precedence over the **Hide/Show all drives** setting.

To prevent the drive from being displayed in Windows Explorer, select **Hide this drive**. This takes precedence over the **Hide/Show all drives** setting.

To allow this drive to be displayed in Windows Explorer, select **Show this drive**. This takes precedence over the **Hide/Show all drives** setting.

This option is available only if the action selected is **Create**, **Replace**, or **Update**.

You can use a Drive Map preference item to configure the visibility of a physical drive rather than a mapped drive. To do so, select the **Update** action, leave the **Location** field blank, select the drive letter of the physical drive, and then configure the **Hide/Show this drive** and **Hide/Show all drives** options.

Hide/Show all drives (8)

Configure the visibility of all mapped and physical drives in Windows Explorer. The options are comparable to those for **Hide/Show this drive**, but apply globally to all drives.

Hide/Show this drive options have precedence over **Hide/Show all drives**. For example, if a Drive Map preference item has the **Hide/Show this drive** option set to **Hide this drive** and the

Hide/Show all drives option set to **Show all drives**, then all drives are visible except the drive designated as hidden.

How does it work?

The Group Policy Drive Map preference extension is a Group Policy client side extension that is hosted in the gpprefcl.dll dynamic linked library. As part of the Group Policy specification, each component of data stored in a Group Policy object must be have two 128-bit unique identifiers in string format.

Client-side and Tool identifiers

The first identifier associates the portion of data stored in the Group Policy object with the entity responsible for consuming, or processing the data. This identifier is known as the client-side extension identifier.

The second identifier associates the portion of the data stored in the Group Policy object with the entity responsible for authoring, or managing it. This identifier is known as the snap-in identifier, or tool identifier.

The client-side extension unique identifier for the Group Policy Drive Map preference extension is {5794DAFD-BE60-433f-88A2-1A31939AC01F}. The snap-in unique identifier for the Group Policy Drive Map preference extension is {2EA1A81B-48E5-45E9-8BB7-A6E3AC170006}.

Group Policy uses the client-side identifier in many locations. First, the client side identifier is stored in the registry of the local computer. The Group Policy infrastructure requires registration of all client-side extensions. The registry location is

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions
```

The Group Policy infrastructure reads the preceding registry location to discover all Group Policy client-side extensions on the computer. Client-side extensions not listed at this location do not process during the application of Group Policy. The registration information includes a sub-key with the client-side extensions unique identifier as the name. Several registry values reside beneath the identifier key. These values are used by the Group Policy infrastructure, but also provide valuable information such as the name of the client side extension (in the **Default** registry value), the binary file that hosts the client-side extension code (in the **DLLName** registry value), and the event source and event log the extension uses to report information, warnings, and errors (in the **EventSources** registry value, separated by a comma, respectively).

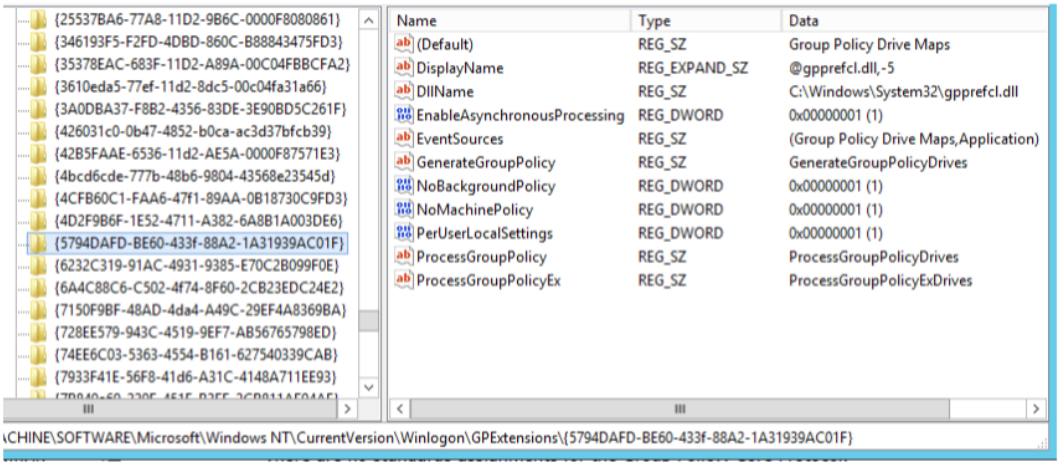

When a Group Policy object is configured to include data from a specific client side extension, the editor must include the client-side extension and snap-in unique identifiers as part of the data saved in the Group Policy object.

Extensions specific data configured to apply to computers is saved to the **gPCMachineExtensionNames** attribute of the Group Policy container object. Extension specific data configured to apply to users is saved to the **gPCUserExtensionNames** attribute. If extension specific data is configured to apply to both user and computer, then the editor saves the extension identifier and the snap-in identifier to both attributes of the Group Policy container object

The **gPCMachineExtensionNames** and the **gPCUserExtensionNames** attribute can hold one or more groups of unique identifiers. Each identifier group is comprised of an extension identifier followed by one or more snap-in identifiers. The group of identifiers are bracketed accordingly by surrounding both identifiers within square brackets ([ ]).

```
extensionID1snapinID1snapinID2…… [ extensionIDnsnapinIDn…]
```

The unique identifier grouping must be listed in ascending alphanumeric order. Saving these in any other order prevents Group Policy from correctly applying or prevents you from editing existing extensions specific data within the Group Policy object.

Storing the extension identifier and the snap-in identifiers in the Group Policy object provides hints that optimize Group Policy processing and management. Processing is optimal because the Group Policy infrastructure knows prior to processing which extensions are required to apply the data hosted in the Group Policy object. Authoring is optimal because the editor has hints to determine what specific snap-ins are needed to edit the data.

XML File Configuration

The Group Policy Drive Map preference extension stores all configuration data associated with the Group Policy object in an XML-formatted file. The **Drives.xml** file resides in the **Drives** folder in the Group Policy template that is hosted and shared on all domain controllers' SYSVOL share.

The Group Policy template location varies for each Group Policy object and is the string value stored in the **gPCFileSysPath** attribute on the Group Policy container object.

```
\\dns_domain_name\Sysvol\dns_domain_name\Policies\group_policy_id
```

Elements

Group Policy Preference configuration files use two XML elements to describe the configuration data: an outer element and an inner element. The outer element serves as a logical container of the inner XML elements. The inner elements represent the configuration of one or more preference items belonging to the container (outer XML element).

Outer Element

The Group Policy Drive Map configuration file uses an outer XML element named **Drives**. This element serves as a collection of inner elements where each inner element represents a drive map preference item. The **Drives** outer element supports two attributes: **clsid** and **disabled**.

Attributes

The **Drives** outer element of the Group Policy Drive Map configuration file contains one required and one optional attribute.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of outer element. The Group Policy Drive Map client-side extension expects the outer element's **clsid** value to equal {8FDDCC1A-0C3C-43cd-A6B4-71A6DF20DA8C}. This is how the client-side extension identifies and validates the outer-element is of type **Drives**.

Disabled

The disabled attribute is an optional attribute for the Drives outer element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A Drives outer element with a disabled attribute equaling true indicates the configuration for all inner elements is disabled. This prevents the Group Policy Drive Map extension from processing any of the inner elements of the configuration file. A disabled attribute equaling false indicates the configuration for all inner elements is enabled and that processing continues as normal.

A **Drives** outer element without a **disabled** attribute indicates that the outer element has never been disabled. In this configuration, the default value for the **disabled** attribute is false, which means the outer element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the outer element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

Inner Element

The Drive outer element accepts one type of inner element. The **Drive** (non-plural) inner element represents a single Drive Map preference item. There can be one or more **Drive** inner elements within the **Drives** outer element.

Each **Drive** inner element contains attributes that describe the inner element. These attributes are jointly used by management editors and client side extensions as a way to identify one **Drive** inner element from the next. **Drive** inner element attributes do not contain configuration data.

Attributes

ByPassErrors

The **byPassErrors** attribute is an optional attribute for the **Drive** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

One or more Drive inner element are contained within a single **Drives** outer element, which enables the client-side extension to process one or more preference items. A Drive inner element with a byPassErrors attribute equally true indicates the Group Policy Drive Map extension should continue processing the next Drive inner element regardless if the result of the current **Drives** inner element is a failure. A byPassErrors attribute equaling false indicates the extension should stop processing subsequent **Drive** inner elements if the result of the current **Drive** inner element is a failure.

A **Drive** inner element without a *byPassErrors* attribute indicates that the inner element has never been configured to halt on errors. In this configuration, the default value for the **byPassErrors** attribute is true, which means the any inner element failure does not halt the processing of subsequent inner elements. The **byPassErrors** attribute appears in the configuration the first time the inner element is configured to stop on errors. Once in this configuration, the value of the attribute changes with each respective configuration, but never removed.

Changed

The **changed** attribute is an optional attribute for the **Drives** inner element. The attribute uses a string value to represent the date and time of when the inner element was last modified in UTC. The string value uses the format of YYYY-MM-DD HH:MM:SS.

Changed

The **changed** attribute is an optional attribute for the **Drives** inner element. The attribute uses a string value to represent the date and time of when the inner element was last modified in UTC. The string value uses the format of YYYY-MM-DD HH:MM:SS.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of inner element. The Group Policy Drive Map client-side extension expects each inner element's **clsid** value to equal {935D1B74-9CB8-4e3c-9914-7DD559B7A417}. This is how the client-side extension identifies and validates the inner element is of type **Drive**

Disabled

The **disabled** attribute is an optional attribute for the **Drive** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A **Drive** inner element with a **disabled** attribute equaling true indicates the configuration for that inner element is disabled. This prevents the Group Policy Drive Map extension from processing that inner element in the configuration file. A **disabled** attribute equaling false indicates the configuration for that inner element is enabled and that processing continues as normal.

A **Drive** inner element without a **disabled** attribute indicates that the inner element has never been disabled. In this configuration, the default value for the **disabled** attribute is false, which means the inner element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the inner element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

Image

The **image** attribute is a numeric string value that represents the index of a bitmap resource used by the Group Policy Management editor as the icon image used to display the preference item in the editor. The value of the numeric string typically corresponds to the value of the **action** attribute in the **properties** element.

Image Value	Action Value
0	C
1	R
2	U
3	D

Name

The name attribute is a string value that represents the display name of the preference item. The name attribute has no impact on the processing of the preference item and is strictly used for management and reporting.

removePolicy

The **removePolicy** attribute is an optional attribute for the **Drive** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **Drive** inner element with a **removePolicy** attribute equaling true indicates the Group Policy Drive Map extension should remove the drive map if the Group Policy object hosting the Drive Map preference item is not within scope of the user. Once the Group Policy object hosting the Drive map preference item no longer applies to the user, the Group Policy Drive map extension deletes the drive map based on the drive letter criterion.

The **removePolicy** attribute correlates to the value of the **action** attribute in the **Properties** element. A **Drive** inner element with a **removePolicy** attribute equals true requires the value of the **action** attribute of **Properties** element to equal an uppercase R.

A **Drive** inner element with a **removePolicy** attribute equaling false indicates the extension should leave existing mapped drives resulting from a Drive Map preference item configured.

A **Drive** inner element without a **removePolicy** attribute indicates the inner element has been configured to remove when it no longer applies. In this configuration, the default value for the **removePolicy** attribute is false, which means all map drives resulting from a Drive Map preference item remain configured. The **removePolicy** attribute appears in the configuration the first time the inner element is configured to remove items when it no longer applies. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

Status

The **status** attribute is string value that represents a message that the Group Policy Preference editor can display in the status area of the Microsoft Management Console. The **status** attribute remains as part of the configuration; however, is no longer being actively implemented.

Uid

The **uid** attribute is a 128-bit identifier represented as global unique identifier (GUID). The **uid** value uniquely identifies each inner element. This is how the client-side extension identifies and differentiates each **Drive** inner element.

UserContext

The **userContext** attribute is an optional attribute for the **Drive** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A Drive inner element with a **userContext** attribute equaling true indicates the Group Policy Drive Map extension should create the mapped drive using the security context of the current user. A **Drive** inner element with a **removePolicy** attribute equaling false indicates the extension should create the mapped drive using the security context of the local system.

A **Drive** inner element without a **userContext** attribute indicates the inner element has been manually configured to apply in the security context of the current user. In this configuration, the default value varies on a per-policy basis using the simple guidelines of computer preference items run in the security context of the local system and user preference items run in the security context of the current user.

The **userContext** attribute appears in the configuration the first time the inner element is manually configured to apply using the security context of the current user. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

The **userContext** attribute remains as part of the configuration; however, is no longer being actively implemented. Most of the Group Policy Preference extensions programmatically switch between the current user and system security contexts as needed, irrespective of the **userContext** attribute value present in the inner element configuration. The recommended

configuration is to leave each preference item to the defaults and to not manually configure the security context of any preference item.

```
<Drive clsid="{935D1B74-9CB8-4e3c-9914-7DD559B7A417}" name="Map Drive F:" status="F:"
image="2" changed="2013-03-26 14:56:38" uid="{CABDD26A-E1E7-412F-A6B5-CE20F9A8E1D4}">
```

Properties Element

Each inner element requires a Properties element in the configuration file. Group Policy Preference extensions read the Properties element for the preference item configuration.

Attributes

The **Properties** element shares unique and common attributes among all the Group Policy Preferences inner element classes. The common attributes characterize abstract functionality provided by the inner element classes (preference items) such as create, delete, or update. However, the implementation of these actions varies across each inner element.

Action

The **action** attribute is an optional, string attribute in the **Properties** element used to describe the action the Group Policy Drive Map extension performs for the associated inner element. That **action** attribute value is a single, uppercase letter of C, R, U, or D. If the **action** attribute's is not present in the **Properties** element then the default value is U.

Create

An **action** attribute value equaling an upper case C instructs the Group Policy Drive Map extension to create a new drive mapping using the settings in the current **Drive** inner element.

The create action has two behaviors based on the value in the **useLetter** attribute. When the **useLetter** value equals 1, the create action creates a new drive mapping only when the current user does not have an active drive letter mapped with a value that matches the value configured in the **letter** attribute. The newly created drive mapping uses the drive letter value configured in the **letter** attribute.

If the current user has an existing drive mapping using the value configured in the letter attribute then the extension ignores the instruction and does not report error.

When the **useLetter** value equals 0, the create action creates a new drive mapping only when the current user does not have an active drive letter mapped to a file share that matches the value of the **path** attribute. The newly created drive mapping uses the first available drive letter starting with letter value configured in the **letter** attribute.

If the current user has an existing drive mapping to a file share that matches the value in the path attribute then the extension ignores the instruction and does not report an error.

The create action, after successfully creating the newly mapped drive letter, also sets the drive letters visibility based on the value configured in the **thisDrive** attribute.

Delete

An **action** attribute value equaling an upper case D instructs the Group Policy Drive Map extension to remove an existing drive mapping using the settings in the current **Drive** inner element.

The delete action has two behaviors based on the value in the **useLetter** attribute. When the **useLetter** attribute value equals 1, the delete action removes the mapped drive only when the current user has an active drive letter mapped with a value that matches the value configured in the **letter** attribute.

If the current user does not have an existing drive mapping using the value configured in the letter attribute then the extension ignores the instruction and does not report an error.

When the **useLetter** attribute value equals 0, the delete action removes all existing mapped drives starting with the drive letter that equals the value in the **letter** attribute (inclusive) all the way through the letter Z.

The extension skips any drive letter that is not currently mapped to a shared location and does not report errors for these drive letters. The extension skips drive letters mapped to physical devices, but reports an error when attempting to remove drive letters associated with physical devices.

Replace

An **action** attribute value equaling an upper case R instructs the Group Policy Drive Map extension to replace an existing drive mapping with the settings in the current **Drive** inner element. The replace action instructs the extension to perform the delete action first, and then the create action. The delete instruction removes an existing mapped drive using the criteria identified in the Delete section of this document. The create instruction creates a new mapped drive using the criteria identified in the Create section of this document. The results of the replace action give the perception that the extension overwrote all settings associated with the mapped drive.

Update

The **action** attribute value equaling an upper case U instructs the Group Policy Drive Map extension to update an existing drive mapping with the settings configured in the current **Drive** inner element.

The update action has two behaviors based on the value in the **useLetter** attribute. When the **useLetter** value equals 1, the update action updates the an existing drive mapping only when the current user has an active drive letter mapped with a value that matches the value configured in the **letter** attribute.

If the current user does not have an existing drive mapping using the value configured in the letter attribute then the extension invokes the create action to create a new mapped drive using the settings configured in the current **Drive** inner element.

An empty **path** attribute value causes the extension to report an error when attempting to create a new mapped drive.

A **useLetter** attribute value equaling 0 causes the extension to perform one of two behaviors. An empty **path** attribute value instructs the extension to perform the update on all existing drive

letters starting with the drive letter that equals the value in the **letter** attribute (inclusive) and ending with the drive letter Z.

This aspect of the update portion does not skip physical drive letters. As a result, the update action can change visibly and label options for physical as well as mapped drive letters.

The second possible action from this configuration occurs when the **useLetter** attribute value equals 0 and the **path** attribute value contains a UNC path of remote file share. This configuration instructs the extension to search through the existing drive letters starting with the drive letter that equals the value in the **letter** attribute (inclusive) and perform the update on the drive letter where the drive letter's mapped network location matches the value configured in the **path** attribute.

It is important to understand that the Group Policy Drive Map extension cannot update the Location, Reconnect, or Connect as (optional) settings. Configure the preference item action to Replace to modify these drive map settings.

```
<Properties action="C" path="\\server\share" useLetter="1" letter="F"/>
```

allDrives

The **allDrives** attribute is an optional attribute for the **Properties** element used to configure the visibility of all physical and mapped drive letters as they appear in Windows Explorer through the Group Policy Drive Map extension. The **allDrives** attribute value is string value represented using all upper case characters. Acceptable values for the **allDrives** attribute include NOCHANGE, HIDE, and SHOW.

The NOCHANGE value instructs the extension not to change drive letter visibility for all the drives of the current user. Therefore, drive letters visible before the extensions processes remain visible after the extension complete. The same behavior is observed for hidden drive letters.

The HIDE value instructs the extension to hide all the drive letters for all the drives of the current user. The results of this configuration hides all drive letters, physical and mapped, from view in Windows Explorer.

The SHOW value instructs the extension to show all the drive letters for all the drivers of the current user. The results of this configuration displays all drive letters, physical and mapped, in Windows Explorer.

Remember, drive visibility and labels extend beyond mapped drive letters and include physical drive letters.

The **allDrives** attribute when combined with the **thisDrive** attribute by provides varying configuration permutations. It is important to remember that the **thisDrive** attribute value has precedence over the **allDrives** attribute value, but only for the drive letter configured in the **letter** attribute's value of the **Properties** element.

The default value for the **allDrives** attribute is NOCHANGE.

Cpassword

The cpassword attribute is an optional string attribute in the Properties element used to store a password needed to successfully connect the configured mapped drive. The attribute value is a

string representation of the password that is obfuscated using an AES-derived encryption key. The default value for the cpassword attribute is an empty string, or blank.

Label

The label attribute is an optional string attribute in the Properties element used to configure a description on the designated mapped drive, which is viewable in Windows Explorer and Windows utilities such as NET.EXE. The attribute value accepts any alphanumeric or printable symbol characters in either upper or lower case.

Letter

The letter attribute is a required attribute in the Properties element that accepts a single upper case alphabetic character value. The extension uses the letter attribute as a reference to a specific drive letter or the beginning point of a range of drive letters. How the extension interprets the letter attribute value depends on the value of the useLetter attribute.

A useLetter attribute value equaling 0 causes the extension to treat the letter attribute value as the beginning letter in a range of upper case letters that stops with the letter Z. A useLetter attribute value equaling 1 causes the extension to treat the letter attribute value as a single, specific drive letter.

Path

The path attribute is a required attribute in the Properties element that accepts a string value Universal Naming Convention (UNC) path to a shared network folder.

The create, replace, and update actions use this value as the target network location to which the extension should map the drive letter.

Also, the extension uses the path attribute value as an alternative criterion for the update action when the value of the useLetter attribute equals 0. Additionally, the value of the path attribute can equal and empty string (blank) when the action value equals U. An empty path attribute value during the update action causes the extension to report an error when the value of useLetter attribute does not equal 1.

Persistent

The path attribute is a required attribute in the Properties element that accepts a string value Universal Naming Convention (UNC) path to a shared network folder.

The create, replace, and update actions use this value as the target network location to which the extension should map the drive letter.

Also, the extension uses the path attribute value as an alternative criterion for the update action when the value of the useLetter attribute equals 0. Additionally, the value of the path attribute can equal and empty string (blank) when the action value equals U. An empty path attribute value during the update action causes the extension to report an error when the value of useLetter attribute does not equal 1.

Persistent

The persistent attribute is an optional attribute in the Properties element. The attribute uses a numeric string to represent a Boolean value. A value equaling 1 indicates true and a value of 0 indicates false.

A Properties element within a Drive inner element with a persistent attribute equaling true indicates the Group Policy Drive Map extension creates the drive mapping to remain persistent between user logons and computer reboots. Windows remembers the mapped drive and reestablishes the drive mapping on subsequent user logons.

A Properties element within a Drive inner element with a persistent attribute equaling false indicates the Group Policy Drive Map extension creates the drive mapping not to persist between user logons and computer reboots. Windows does not remember the mapped drive and does not reestablish the drive mapping on subsequent reboots.

It's important to understand that Windows configures drive persistency at the time it creates the mapped drive. Therefore, if you want need to use the Replace action if you need to change a mapped drive persistent state.

Useletter

The useLetter attribute value is a required attribute for the **Properties** element. The attribute uses a numeric string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 represents false.

A **useLetter** attribute value equaling false causes the extension to treat the **letter** attribute value as the beginning letter in a range of upper case letters that stops with the letter Z. A **useLetter** attribute value equaling true causes the extension to treat the letter attribute value as a single, specific drive letter.

Username

The username attribute is an optional string attribute in the **Properties** element used to store the domain and username needed to successfully connect the configured mapped drive.

thisDrive

The **thisdrive** attribute is an optional attribute for the **Properties** element used to configure the Windows Explorer visibility for the **letter** attribute's value in the **Properties** element. The **thisDrives** attribute value is a string value represented using all upper case characters. Acceptable values for the **thisDrives** attribute included NOCHANGE, HIDE, and SHOW.

The NOCHANGE value instructs the extension not to change configured drive letter's visibility for the current user. Therefore, a drive letter visible before the extensions processes and remains visible after the extension completes. The same behavior is observed for hidden drive letters.

The HIDE value instructs the extension to hide the configured drive letter for the current user. The results of this configuration hides the configured drive letter, physical or mapped, from view in Windows Explorer.

The SHOW value instructs the extension to show the configured drive letter for current user. The results of this configuration displays the configured drive letter, physical or mapped, in Windows Explorer.

Remember, drive visibility and labels extend beyond mapped drive letters and include physical drive letters.

The **allDrives** attribute when combined with the **thisDrive** attribute by provides varying configuration permutations. It is important to remember that the **thisDrive** attribute value has

precedence over the **allDrives** attribute value, but only for the drive letter configured in the **letter** attribute's value of the **Properties** element.

The default value for the **thisDrives** attribute is NOCHANGE.

```
<Drives clsid="{8FDDCC1A-0C3C-43cd-A6B4-71A6DF20DA8C}" disabled="0">
<Drive clsid="{935D1B74-9CB8-4e3c-9914-7DD559B7A417}" name="F:" status="F:" image="0"
changed="2013-04-04 16:48:16" uid="{B225E3B9-C639-4AA3-AB79-CB65D626CF9A}">
 <Properties action="C" thisDrive="NOCHANGE" allDrives="NOCHANGE" userName=""
path="\\server\share" label="" persistent="0" useLetter="1" letter="F"/>
 </Drive>
<Drive clsid="{935D1B74-9CB8-4e3c-9914-7DD559B7A417}" name="G:" status="G:" image="2"
changed="2013-03-26 14:56:38" uid="{CABDD26A-E1E7-412F-A6B5-CE20F9A8E1D4}">
 <Properties action="U" thisDrive="NOCHANGE" allDrives="NOCHANGE" userName=""
path="\\server\share" label="" persistent="0" useLetter="1" letter="G"/>
 </Drive>
<Drive clsid="{935D1B74-9CB8-4e3c-9914-7DD559B7A417}" name="H:" status="H:" image="1"
changed="2013-04-04 16:48:40" uid="{94BCBDAE-E11C-4C3A-9E05-CE1B416C1759}" disabled="0">
 <Properties action="R" thisDrive="NOCHANGE" allDrives="NOCHANGE" userName=""
path="\\server\share" label="" persistent="0" useLetter="1" letter="H"/>
 </Drives>
```

Evaluation of XML

The extension reads the entire XML file into memory subsequently processes each element in the XML file starting with the outer element. The outer element defines the overarching class of preference items that appear in the inner elements of the file. Additionally, the outer element may contain the disabled attribute. A disabled attribute value equaling 1 in the outer element stop any further evaluation of the XML, effectively preventing all of the Drive inner elements from processing (applying to the user).

After reading the Drives element, the extension focuses its attention to the next element in the XML file—one of possibly many Drive inner element. As noted in the attribute section of this document, the Drive inner element contains metadata and processing control information used be the Group Policy management snap-ins and the client-side extensions.

The extension reads the attributes and values from the entire Drive inner element, which includes the attributes and values of the Drive element, the attributes and values of the Properties element and the attributes and values of an optional Filters element along with all its child elements and their attributes and values.

A Drive inner element containing a disabled attribute equaling 1 stops the extension from processing this element (applying to the user) and the extension moves its focus to the next Drive inner element.

The extension continues processing the current Drive inner element when it does not contain a disabled attribute or when it contains a disabled attribute with a value equaling 0.

Next, the extension moves its focus on the Filter element of the Drive inner element.

The Filters element holds Item-level targeting information that the extension uses as additional criteria to determine if the associated preference items (Mapped Drives) should apply to the user or computer.

The Filters element contains attributes, values and child elements that represent of sophisticated set of Boolean rules. The details of the Filters element are beyond the scope of this portion of the document. In this context, all that is needed to know is a Filter element, when evaluated by an extension, only returns true or false.

When present, the extension reads the entire Filters element and evaluates each Boolean expression included in the section. The evaluation of the Filters element can only result in a true or false response. A filter evaluation result equaling false means the one or more Boolean expressions in the Filters element evaluated to false. The false evaluation forces the extension to stop any further processing of the current Drive inner element and moves focus to the next Drive inner element.

A Drive inner element that does not contains a Filters inner element or an inner element containing a Filters inner element that evaluates to true moves its focus to the Properties element, which is the next element within the Drive inner element.

The Properties element within a Drive inner element contains all the Drive Map configuration information the extension needs to create, replace, update, or delete a mapped drive. The Group Policy Drive Map extension reads the attributes and values and performs the requested instructions and determines if the actions performed were successful.

If the actions performed were unsuccessful (an error occurred), the extension checks the bypassErrors attribute value read from the current Drive inner element. A false value (zero [0]) instructs the extension to discontinue further processing of all Drive inner elements while a true value (one [1]) instructs the extension to ignore the unsuccessful result and continue processing additional Drive inner elements.

The extension then moves its focus to the next Drive inner element and the process repeats until the XML file reaches the Drives closing outer element.

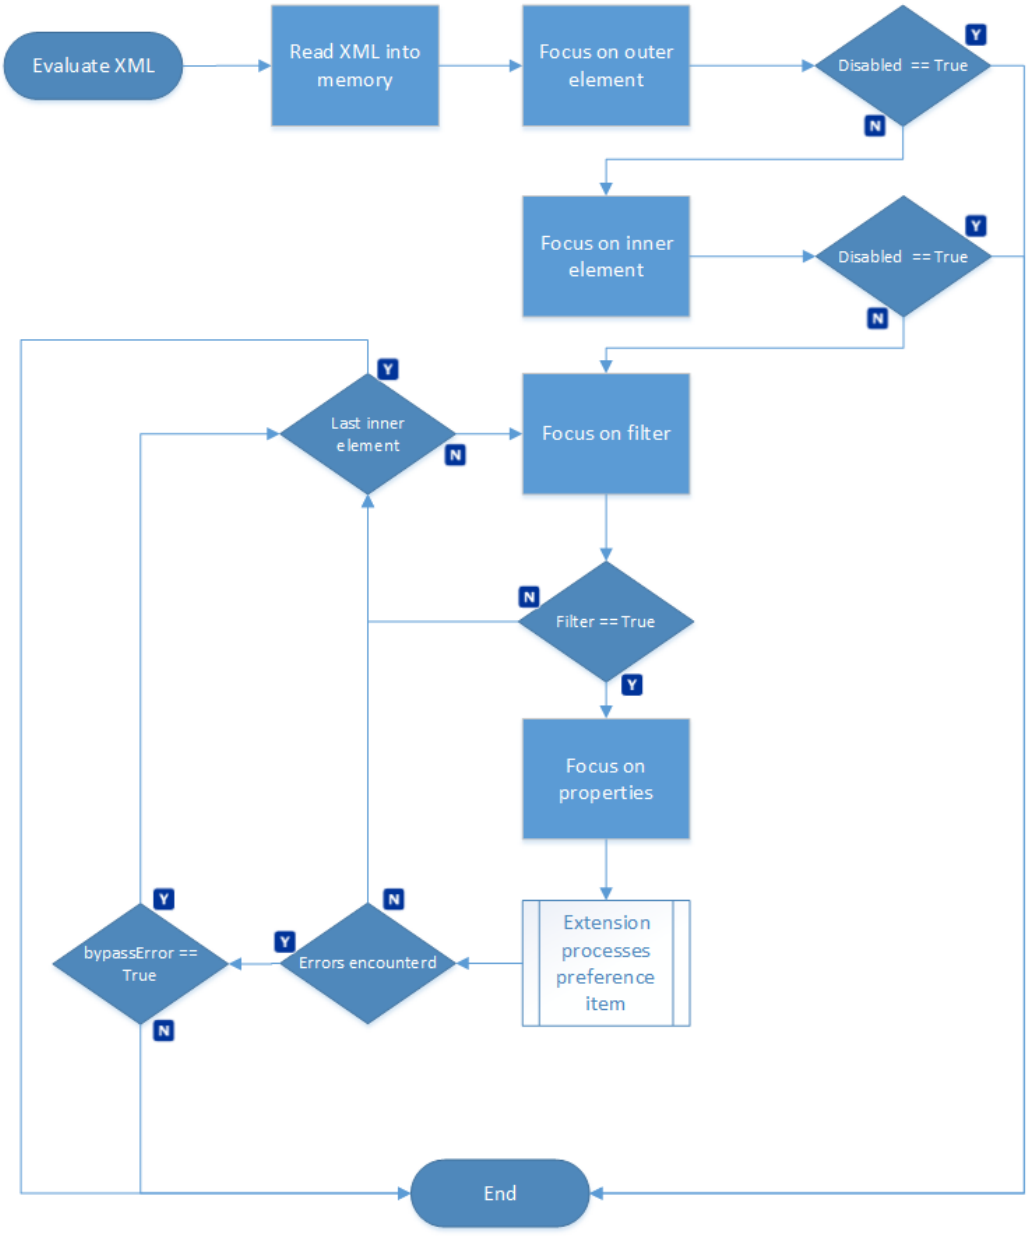

Xml to User Interface mapping

The Group Policy Drive Map management snap-in determines the values and attributes stored in the XML configuration. The following figures correlate which components of the user interface control or shows the various elements and attributes in the XML file.

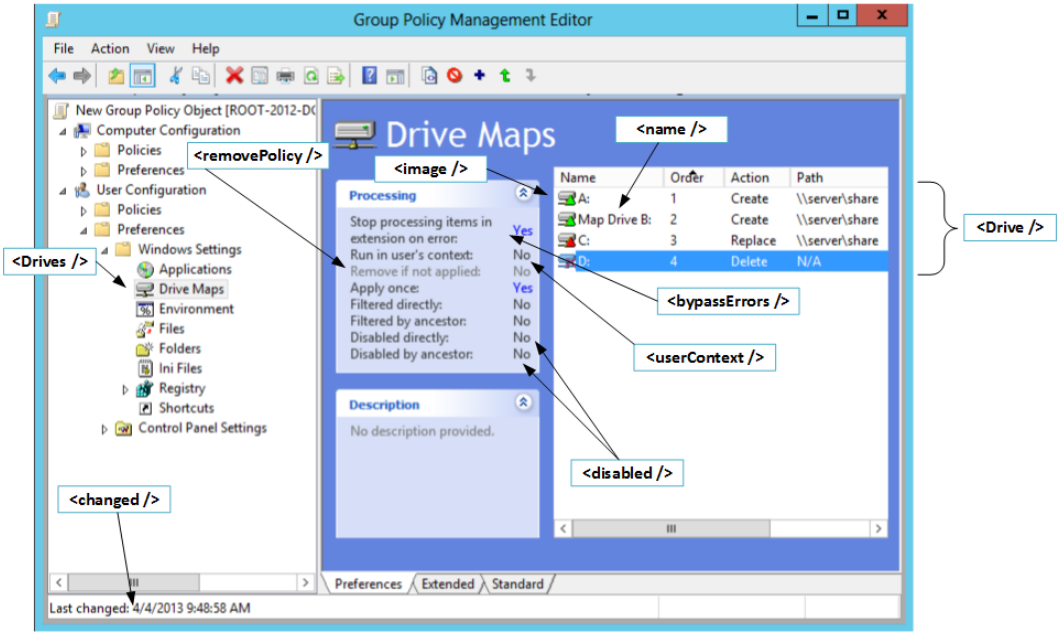
Group Policy Management Editor
File Action View Help
New Group Policy Object [ROOT-2012-DC
Computer Configuration
Policies
Preferences
User Configuration
Policies
Preferences
Windows Settings
Applications
Drive Maps
Environment
Files
Folders
Ini Files
Registry
Shortcuts
Control Panel Settings
Drive Maps
Processing
Stop processing items in extension on error: Yes
Run in user's context: No
Remove if not applied: No
Apply once: Yes
Filtered directly: No
Filtered by ancestor: No
Disabled directly: No
Disabled by ancestor: No
Description
No description provided.
Name Order Action Path
A: 1 Create \\server\share
Map Drive B: 2 Create \\server\share
C: 3 Replace \\server\share
D: 4 Delete N/A
Preferences Extended Standard
Last changed: 4/4/2013 9:48:58 AM
<name />
<removePolicy />
<image />
<Drive />
<Drives />
<bypassErrors />
<userContext />
<disabled />
<changed />

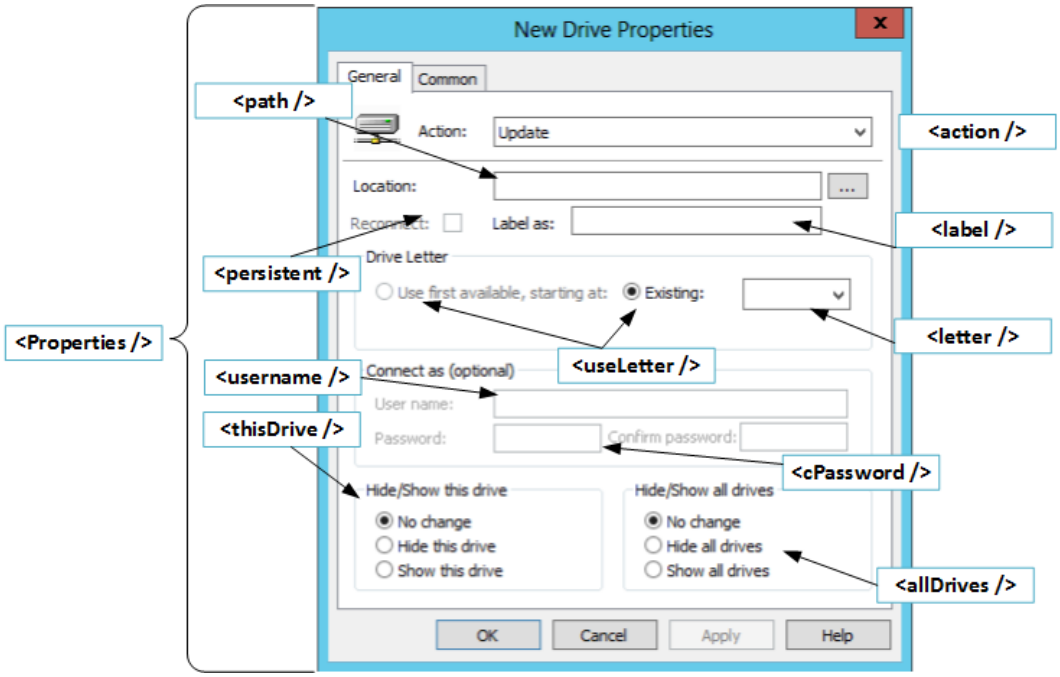
New Drive Properties
General
Common
Action: Update
Location:
Reconnect:
Label as:
Drive Letter
Use first available, starting at:
Existing:
Connect as (optional)
User name:
Password:
Confirm password:
Hide/Show this drive
No change
Hide this drive
Show this drive
Hide/Show all drives
No change
Hide all drives
Show all drives
OK
Cancel
Apply
Help
<path />
<action />
<label />
<persistent />
<Properties />
<letter />
<useLetter />
<username />
<thisDrive />
<cPassword />
<allDrives />

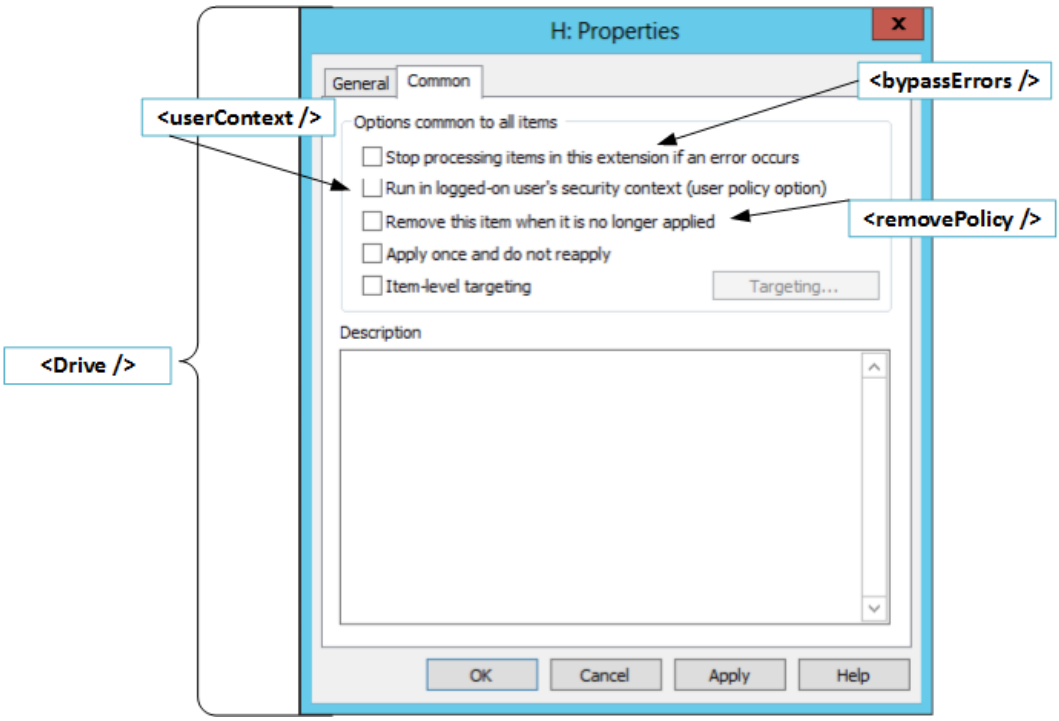

CSE Processing

Group Policy client side extensions enable Group Policy to process a variety of different tasks because each client side extension is responsible for processing its portion of data stored in a Group Policy object.

The Group Policy Drive Maps client side extension is one of many extensions included in the gpprefcl.dll file that is included in the Windows operating system starting with Windows Server 2008 and Windows Vista Service Pack 1.

Group Policy's modular design enables each extension to process the data that is relevant to them. Therefore, the implementation of each Group Policy client side extension may be different.

The Group Policy Drive Map extension is one CSE in a family of Group Policy client side extensions known as Group Policy Preferences. The goal of this extension is to manage drive letters associated with network specific shared folders and visibility properties of driver letters (remotely mapped drives or physical). The extension comes registered on the Windows operating system by virtue of CSE specific metadata stored in the following registry location

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows
NT\CurrentVersion\Winlogon\GPExtensions\{5794DAFD-BE60-433f-88A2-1A31939AC01F}
```

The information stored in this location is consistent on all computers. Therefore, restoring the information on a computer where it is missing is relatively easy using the export and import feature of the Registry editor.

Startup

The Group Policy service invokes this client side extension when it encounters a Group Policy object configured with Group Policy Drive maps. When activated, the extension must locate its configuration file.

Prerequisite information

When Group Policy invokes the client side extension, it provides to the extension information that enables it to process its portion of Group Policy. This information includes operational information that describes the current processing context such as if the processing is background or foreground, is the processing occurring over a slow link, or is the processing for the computer rather than the user. Also included in this information is a list of the Group Policy objects that were removed from the scope of application and the list of Group Policy objects that were changed.

Each entry in the list of Group Policy objects provides to the client side extension provides information about the Group Policy object in question. The entry includes the information in Table 2.

Field	Description
Options	Client side extensions use the options field to determine if the Group Policy object is disabled or if the Group Policy object is
Version	The Group Policy object version number increments with saved changes to the computer and user portions of a GPO, respectively. Client side extensions can use the version number to determine if any changes were made to the computer or user portion of a GPO.
DSPath	The DS path provides the LDAP path, which provides the location and common name of the Group Policy object in Active Directory. Remember, the common name of a GPO is its unique identifier; not its display name.
File system path	The file system path provides the root location of the files associated with the Group Policy object. These collection of files typically are used by one or more client side extensions as the data store for group policy settings.
Link Path	The link path field provides an LDAP path of

	the container to which the Group Policy is linked
Link Type	The link type field provides information about the container to which the GPO is linked such as computer, site, domain, or organizational unit.
Extensions	The extension field provides a list of the extensions that are configured in the Group Policy object.
Display name	The display name is the name of the GPO that is displayed in user interfaces and reporting tools.

The information in Table 2 is present for each GPO in both lists. The out-of-scope list contains a list of GPOs that applied to the user or computer on the last application of Group Policy but no longer apply on this application of Group Policy. The list of changed list contains a list of GPOs that include changes for one or more aspects of the Group Policy object.

The extension reads the operational information to understand what type of Group Policy processing Windows is performing: user or computer. How the processing is occurring: foreground or background. This enables the client side extensions to make operational decisions based on this processing characteristics.

The extension uses the lists of Group Policy objects to complete its portion of Group Policy processing. The Group Policy Drive Maps extension starts by processing each of the GPOs in the out-of-scope list.

Out-of-Scope Group Policy objects

The Group Policy Drives Maps processes the entire list of out-of-scope GPOs sequentially. The extension logs any errors reading the Group Policy objects and continues processing until it reaches the end of the list.

The extension cycles through the list of Group Policy objects in the list starting with the first entry in the list. The extension collects the prerequisite information about the current GPO. From this information and its own information, the extension builds file path location to the configuration file. The configuration file location path begins with the extensions history path.

Group Policy Preference History

The history path is a local data store of the last successfully applied data configuration files for each Group Policy object. The extension saves the last applied configuration files in the user's local application folder, which holds a value equivalent to the environment variable

%LOCALAPPDATA% on Windows 8 and Windows Server 2012. The root of the history folder looks similar to

```
C:\users\[username]\AppData\Local\Microsoft\Group Policy\History
```

Windows protects the history folder with the hidden attribute; therefore, by default, it is not visible from Windows Explorer or from the command line. The history folder contains sub folders, where each folder is named to the globally unique identifier (GUID) of a corresponding Group Policy object.

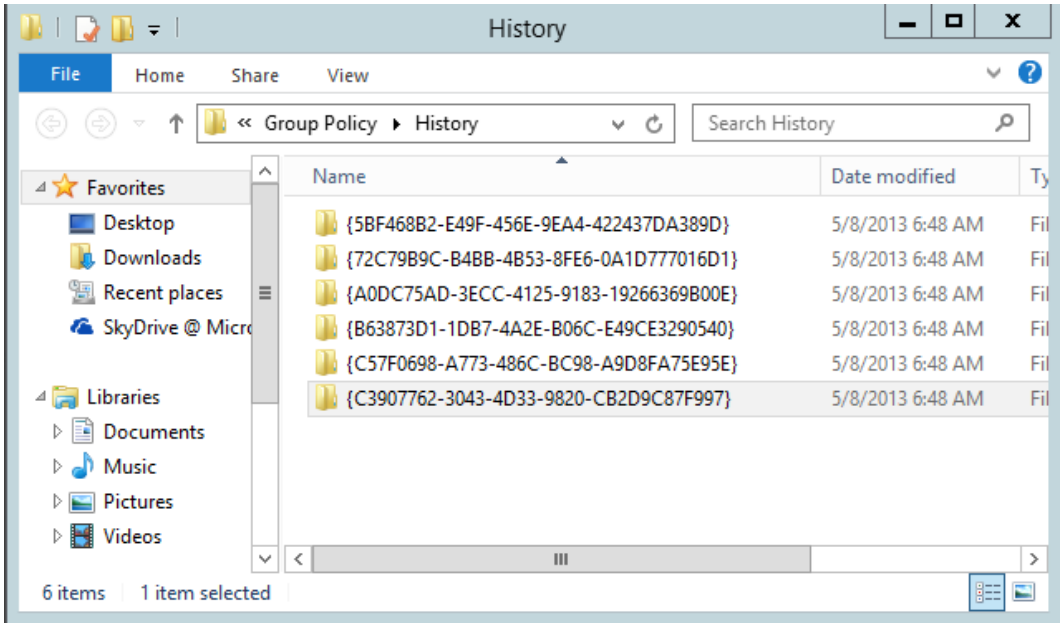

Each Group Policy object folder within the history folder contains sub folders where each folder is named to a corresponding security identifier (SID) of a user who has logged on to the current computer. The complete path to the history file for the Group Policy Drive Maps extension is

```
C:\users\[username]\AppData\Local\Microsoft\Group
Policy\History\[GUID]\SID\Preferences\Drives\Drives.xml
```

The extension then concatenates the string \User\Preferences\Drives\Drives.xml to the end of the GPO's file system path. This creates a full path, including file name, to the configuration file.

Note

Windows only allows the configuration and application of Group Policy Drives Map to users.

The extensions opens this file and reads the entire contents of the file into memory and then closes the file. The CSE logs any errors occurring that may occur with the file operations and moves to the next GPO in the list.

Configuration File

The configuration file contains an outer element with multiple inner elements where each inner element represents a separate drive map preference item. Each inner element includes a Properties element that holds the respective elements configuration information.

The extension begins by settings its focus on the first inner element. The extension scans the current inner element for the removePolicy attribute. An inner element without a removePolicy attribute or an inner element with a removePolicy attribute value equaling false (zero [0]) causes the extension to move to the next inner element.

An inner element with a removePolicy attribute equaling true (one [1]) indicates to the extension that it must remove the settings in the current inner element. The extension uses the settings from the current inner element and its associated Properties element to undo, or remove the settings. The extension accomplish this by forcing the processing to use the Remove action regardless of the action configured in the current inner element and continues this cycle until it has processed all inner elements in the configuration file. Once processing of all inner elements is complete, the extension deletes the \Drives\Drives.xml folder and file for the specific user and GPO from the local history.

The extension then moves to the next Group Policy object in the list and processes it in the same fashion until it has processed all GPOs in the out-of-scope list. The extension begins processing the new and changed GPO list.

New and Changed Group Policy objects

The Group Policy Drives Maps processes the entire list of new and changed GPOs sequentially. The extension logs any errors reading the Group Policy objects and continues processing until it reaches the end of the list.

The extension cycles through the list of Group Policy objects in the list starting with the first entry in the list. The extension collects the prerequisite information about the current GPO. From this information, the extension builds file path location to the configuration file. The configuration file location path begins with the GPO's file system path. The extension then concatenates the string \User\Preferences\Drives\Drives.xml to the end of the GPO's file system path. This creates a full path, including file name, to the configuration file.

Note

Windows only allows the configuration and application of Group Policy Drives Map to users.

Configuration File

The extension reads the entire configuration file into memory. The configuration file contains an outer element with multiple inner elements where each inner element represents a separate drive map preference item. Each inner element includes a Properties element that holds the respective elements configuration information.

Outer Element Processing

The extension evaluates the outer element for the disabled attribute. An outer element with a disabled attribute equaling true (one [1]) indicates the all inner elements are considered disabled. The extension discontinues processing any part of the current configuration and moves to the next Group Policy object in the list. An outer element without a disabled attribute or with a disabled attribute value equaling false (zero [0]) continues processing by evaluating the inner elements,

Inner Element Processing

The extension begins by settings its focus on the first inner element. The extension evaluates the current inner element for the disabled attribute. An inner element with a disabled attribute equaling true (one [1]) indicates the inner element is disabled. The extension discontinues processing the current inner element and move to the next inner element. An inner element without a disabled attribute or inner element with a disabled attribute value equaling false (zero [0]) continues evaluating the inner element for item-level targeting information.

Item-level Targeting Processing

The extension scans the current inner element for a Filters element. The Filters element is an optional element that holds the item-level targeting configuration information for each inner element. A Filters element existing within the inner element indicates the extension must evaluate one or more Boolean criteria before processing the inner element configuration. A false result from a Filters evaluation forces the extension to ignore the current inner element and move to the next inner element. A true result from a Filters evaluation indicates the extension should continue processing the inner element and action configured within the current element.

Action Processing

Action processing occurs once the client side extension has evaluated all other criteria that would instruct it to bypass processing outer or the current inner element. Action processing correlates to the action configured in each inner element (reminder an inner element represents a single preference item). The four actions the Group Policy Drive Map extension can process are create, replace, update, and delete.

Create

The create actions represents a configuration where the extension creates a new mapped drive based on the configuration in the current inner element.

Prologue

The extension extracts the values for the <path>, <letter>, <persistent>, <username>, and <cPassword> attributes from the configuration information. If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extension continues extracting the values for the <allDrives>, <thisDrive>, and <useLetter> attributes. The extension completes the prologue portion of the processing by impersonating the

user logging on the computer using the user's full authentication token. Any failure attempting the impersonation results in the extension aborting its processing and to return the underlying error code that resulted in the impersonation failure.

Core

The core phase of create action processing is to use the configuration information read during the prologue phase to create a mapped network drive for the current user.

Specific Drive Letter

The extension checks if the value of the <useLetter> attribute equals true (one [1]). A true value determines the desired configuration is for a specific drive letter.

The extensions starts by ensure the drive letter is normalized in all lower case and that its syntax matches [driveLetter:]. Any failures encountered during normalization cause the extension to abort and to return an invalid argument error.

The extension then uses the Windows API WNetGetConnection to determine if the value <letter> is currently used by an existing mapped network drive or by a physical device such as a hard drive, CD-ROM/DVD-ROM, ram disk or removable drive.

A <letter> value that currently exists for the user but is not a mapped drive or not considered a physical device causes the extension to abort its process and to return unexpected error. A <letter> value that currently exists for the user that is either mapped to a shared network folder or associated with a physical device causes the extension to abort its processing and to return an object name exists error.

Create Mapped Network Drive

The extension determines if the current user's token is not a restricted token and the <persistent> attribute value equals false (zero [0]). Upon validating these factors, the extension uses the WNetUseConnection API to create a newly mapped drive using the value of the <letter> attribute as the drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and to return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write "Connected with access name…" to the trace log.

If the extension determines the current user's token is a restricted token (User Account Control is enabled) then the extension must create the mapped drive under the restricted token rather than the full token. This explains why mapped drives for administrators and some users are not always visible in Windows Explorer.

The extension stops impersonating the current user with the user's full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write "Failed to revert from the user's full token" to the trace log, and return the underlying error code that is the result of failure to cease

impersonation. Any failure in requesting a restricted token for the current user instructs the extension to abort its processing, write “Failed to impersonate with user’s restricted token” to the trace log, and return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write “Failed to connect drive with restricted token” along with the underlying error code it received from the Windows API to the trace log. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token causes the extension to write “Failed to revert user from restricted context” to the trace log. A failure impersonating the user with their full token causes the extension to write “Failed to impersonate with user’s full token” to the trace log. The extension continues processing by applying the configured label to the newly created mapped drive.

First Available Drive Letter, Starting at

A <useletter> attribute value equaling false (zero [0]) determines the desired configuration is to create a map drive using the first available letter starting with the <letter> attribute value.

Determine the First Available Drive Letter

The extension needs to determine the first available drive letter before it can map any network location. To accomplish this, the extensions creates a list of lower case letters starting with the <letter> attribute’s value and ending with the letter ‘z’. The extension then iterates through each letter in the list of letter to determine the first available drive letter, starting with the <letter> attribute’s value.

The extension determines a letter qualifies as available first by using the WNetGetConnection API to determine if the current letter in the list of letters is mapped to a network drive or associated with a physical device.

Drive Letter not in use

A drive letter not mapped to a shared folder instructs the extension to check if a previously available drive letter was already saved. If the extension already knows of the first available drive letter, than it performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. If the extension is unaware of any previously saved available drive letter, then ensures the current letter is not associated with a physical device. A drive letter associated with a physical device is not eligible as an available drive letter. The

extension performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. A drive letter making to this part of the process has passed all the criteria needed to become the first available drive letter.

Drive Letter in use by a Mapped Drive

A drive letter mapped to a shared folder instructs the extension to evaluate the shared folder location of the existing mapped drive to the <path> attribute's value in the inner element configuration.

If the shared folder path matches the <path> attribute value, then the extension considers the mapped drive already successfully created. The extension writes "Path is connected already on letter" along with the current drive letter to the trace log, discontinues iterating the drive letters, and proceeds to the Drive Label portion of the core processing phase.

If the shared folder path does not match the <path> attribute value, then the extension considers the current drive letter unavailable. The extension performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list.

The extension completes its iteration through the drive letter list and validates a free drive letter exits. If the extension remains unaware of an available drive letter and it did not locate an existing drive letter mapped to the network path identical to the <path> attribute value, then the extension aborts its processing, and returns a not connected error. Otherwise, the extension is aware of an available drive letter and must now connect the shared network folder to the available drive letter.

Connecting to First Available Drive

The extension must create a new mapped drive using the discovered first available drive letter and the <path> attribute's value. To do this, the extension determines the type of token used by the current user and the value of the <persistent> attribute.

If the user has a full token or if the <persistent> attribute equals false (zero [1]), then the extension use the WNetUseConnection API to create a new mapped drive where the drive letter used is the discovered first available drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value, which is false, is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

If the user has a full token or if the <persistent> attribute equals false (zero [1]), then the extension use the WNetUseConnection API to create a new mapped drive where the drive letter used is the discovered first available drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value, which is false, is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log.

A user with a restricted token (User Account Control enabled) indicates to the extension that it must also create the new drive mapping while impersonating the user using the user’s restricted token.

The extension stops impersonating the current user with the user’s full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write “Failed to revert from the user’s full token” to the trace log, and return the underlying error code that is the result of failure to cease impersonation.

Any failure in requesting a restricted token for the current user instructs the extension to abort its processing and to return the underlying error code that is the result of the failure to when requesting the restricted token.

A failure occurring when extension impersonates the user using the restricted token cause the extension to abort its processing, write “Failed to impersonate with user’s restricted token” to the trace log, and return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write “Failed to connect drive with restricted token” to the trace log along with the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token cases the extension to write “Failed to revert user from restricted context” to the trace log. A failure impersonating the user with their full token causes the extension to write “Failed to impersonate with user’s full token” to the trace log. The extension continues processing by applying the configured label to the newly created mapped drive.

Drive Label

To this point in the process, the extension has successfully created a mapped network drive using either a specific drive letter or a first available drive letter. And, the extension has mapped the drive using the user’s full token and restricted token, if needed.

The extension needs to apply the <label> attribute value to the newly created drive letter. The extension ensures a value is present for the <label> attribute and that the mapped drive was newly created and not an existing local or remote drive letter. Once the extensions confirms these conditions, it replaces any backslash (\) characters in the <path> attribute's value with pounds (#) characters. Next, the extension creates a registry key using the modified <path> value under

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Mountpoints2
```

The extension creates a new registry value named _LabelFromReg. The value assigned to _LabelFromReg is original, unmodified <path> attribute value. The extension ignores any failures that may occurs during drive label processing.

Hide or Show Drives

The Group Policy Drive Map extension can show or hide drives using two configurations. It can hide or show a single, specified drive letter and it can hide or show all drive letters. The extension processes all drives first and specified drive last. Thus, the outcome of configuring both results in the specified drive configuration have precedence over the all drives configuration.

The extension determines if it needs to show or hide any drive letters by checking the <thisDrive> and <AllDrives> attribute values. If both attribute values equal "NOCHANGE' then the extension performs no other actions and processes the epilog portion of core processing. Otherwise, the extension begins the hide or show drives process by changing to the security context to System. Any failure in changing to System security context causes the extension to abort its processing and return the underlying error responsible for the failure.

All Drives

The extension determines if it needs to hide or show all drive letters by evaluating the value of the <AllDrives> attribute. If the attribute's value equals "SHOW" then the extension deletes the NoDrives registry value from the following registry hive of the current user.

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

If the attribute's value equals "HIDE" then the extension writes the hexadecimal value 0x03FFFFFF to the NoDrives registry value of the previously mentioned registry hive of the current user (the extension creates the NoDrives registry value if it does not exist). The extension then processes hide or show settings for a specific drive letter.

Specified Drive

The extension validates that the <thisDrive> attribute equals "SHOW" or "HIDE". Next, the extension validates the <letter> attribute value is not blank. If the <letter> attribute value is blank, then the extension aborts its processing and returns an invalid argument error.

The extension retrieves the NoDrives registry value from the following registry hive of the current user

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

The extension uses an algorithm based on information in Microsoft Knowledgebase article 231289, Using Group Policy Objects to hide specified drives, to determine equivalent bit value of the <letter> attribute value.

If the <thisDrive> attribute value equals “HIDE” then the extension enables that bit in the NoDrives registry value. If the <thisDrive> attribute value equals “SHOW” then the extension disabled the bit in the NoDrives registry value.

After modifying the NoDrives registry value, the extension checks its value to ensure it does not equal zero (0). If the newly calculated NoDrives registry value equals zero, then the extension deletes the NoDrives registry value. Otherwise, the extension saves the newly calculated NoDrives registry value to the registry.

Epilogue

The extension completes the processing by updating its %LastDriveMapped% environment variable and by securely deleting the clear text instance of the <cPassword> attribute value. Additionally, if the <persistent> attribute value equals true (one [1]), then the extension signals to the Group Policy infrastructure that the next foreground processing must be synchronous.

Create action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Delete

The delete action represents a configuration where the extension deletes all mapped drives or a specific mapped drive based on the configuration in the current inner element.

Prologue

The extension extracts the values for the <path>, <letter>, <persistent>, <username>, and <cPassword> attributes from the configuration information. If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extension continues extracting the values for the <allDrives>, <thisDrive>, and <useLetter> attributes. The extension completes the prologue portion of the processing by impersonating the user logging on the computer using the user’s full authentication token. Any failure attempting the impersonation results in the extension aborting its processing and to return the underlying error code that resulted in the impersonation failure.

Core

The core phase of delete action processing is to use the configuration information read during the prologue phase to identify and delete a specific mapped drive, or to delete all mapped drives starting with a specific drive letter for the current user.

Specific Drive Letter

The extension checks if the value of the <useLetter> attribute equals true (one [1]). A true value determines the desired configuration is for a specific drive letter.

The extensions starts by ensure the drive letter is normalized in all lower case and that its syntax matches [driveLetter:]. Any failures encountered during normalization cause the extension to abort and to return an invalid argument error.

A <letter> value that currently exists for the user but is not a mapped drive or not considered a physical device causes the extension to abort its process and to return unexpected error.

A <letter> value that currently exists for the user and is associated with a physical device cause the extension to abort its processing and return an already assigned error.

Delete Mapped Drive

A <letter> value that currently exists for the user and is remotely connected to a network share causes the extension to use the Windows API WNetCancelConnection2 to remove the connection and its persistent state from the user's session.

If the Windows API failed to delete the mapped drive, then the extension aborts its processing and returns the underlying error it received from the Windows API. A successful response from the Windows API causes the extension to proceed to the epilog phase.

Delete All, Starting at

A <useletter> attribute value equaling false (zero [0]) determines the desired configuration is to delete all mapped drives starting with starting with the <letter> attribute value.

The extensions creates a list of lower case letters starting with the <letter> attribute value and ending with the letter 'z'. The extension iterates through each letter in the list to determine if the drive is mapped to a shared network folder using the WNetGetConnection Windows API.

If the extension determines the current drive letter is not mapped to a shared network folder, then it completes the current iteration and begins a new integration using the next drive letter in the list as the current drive letter.

If the extension determines the current drive letter is mapped to a shared network folder, then it uses the Windows API WNetCancelConnection2 to delete the connection and its persistent state from the user session. The extension then completes the current iteration and begins a new iteration using the next drive letter in the list as the current drive letter.

This cycles completes until the extension has evaluated all drive letters starting with the <letter> attribute value through the letter 'z'. The extension then proceeds to the epilogue phase.

Epilogue

The extension completes the processing by securely deleting the clear text instance of the <cPassword> attribute value. Additionally, if the <persistent> attribute value equals true (one [1]), then the extension signals to the Group Policy infrastructure that the next foreground processing must be synchronous.

Delete action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Replace

The replace action represents a configuration where the extension replaces a new mapped drive with a new mapped drive based on the configuration in the current inner element.

It is important to understand the difference between replace and update. The replace action typically deletes the existing mapped drive and creates a new one—essentially changing all configurations elements of the mapped drive. The update action only updates a portion of configuration elements in an existing mapped drive.

Prologue

The extension extracts the values for the <path>, <letter>, <persistent>, <username>, and <cPassword> attributes from the configuration information. If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extension continues extracting the values for the <allDrives>, <thisDrive>, and <useLetter> attributes. The extension completes the prologue portion of the processing by impersonating the user logging on the computer using the user's full authentication token. Any failure attempting the impersonation results in the extension aborting its processing and to return the underlying error code that resulted in the impersonation failure.

Core

The core phase of replace action processing is to use the configuration information read during the prologue phase to identify an existing mapped network drive, delete the specified network drive, if found, and to create a newly mapped network drive for the current user.

Specific Drive Letter

The extension checks if the value of the <useLetter> attribute equals true (one [1]). A true value determines the desired configuration is for a specific drive letter.

The extensions starts by ensure the drive letter is normalized in all lower case and that its syntax matches [driveLetter:]. Any failures encountered during normalization cause the extension to abort and to return an invalid argument error.

The extension then uses the Windows API WNetGetConnection to determine if the value <letter> is currently used by an existing mapped network drive or by a physical device such as a hard drive, CD-ROM/DVD-ROM, ram disk or removable drive.

A <letter> value that currently exists for the user but is not a mapped drive or not considered a physical device causes the extension to abort its process and to return unexpected error.

Delete Mapped Drive

A <letter> value that currently exists for the user and is remotely connected to a network share causes the extension to use the Windows API WNetCancelConnection2 to remove the connection and its persistent state from the user's session. A <letter> value that currently exists for the user and is associated with a physical device cause the extension to abort its processing and return an already assigned error.

Create Mapped Network Drive

The extension determines if the current user's token is not a restricted token and the <persistent> attribute value equals false (zero [0]). Upon validating these factors, the extension uses the WNetUseConnection API to create a newly mapped drive using the value of the <letter> attribute as the drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder.

Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and to return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log.

If the extension determines the current user’s token is a restricted token (User Account Control is enabled) then the extension must create the mapped drive under the restricted token rather than the full token. This explains why mapped drives for administrators and some users are not always visible in Windows Explorer.

The extension stops impersonating the current user with the user’s full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write “Failed to revert from the user’s full token” to the trace log and to return the underlying error code that is the result of failure to cease impersonation. Any failure in requesting a restricted token for the current user instructs the extension to abort its processing, write “Failed to impersonate with user’s restricted token”, and to return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write “Failed to connect drive with restricted token” along with the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token causes the extension to write “Failed to revert user from restricted context” to the trace log. A failure impersonating the user with their full token causes the extension to write “Failed to impersonate with user’s full token” to the trace log. The extension continues processing by apply the configured label to the newly created mapped drive.

First Available Drive Letter, Starting at

A <useletter> attribute value equaling false (zero [0]) determines the desired configuration is to replace a mapped drive using the first available letter starting with the <letter> attribute value.

In this configuration, the extension uses the <letter> attribute value as the starting point of the drive letters through which it searches for an existing mapped drive. The <path> attribute value becomes the matching criterion that determines if the mapped drive exists.

Determine the First Available Drive Letter

The extension needs to determine the first available drive letter before it can map any network location. To accomplish this, the extensions creates a list of lower case letters starting with the <letter> attribute value and ending with the letter 'z'. The extension iterates through each letter in the list to determine the first available drive letter, starting with the <letter> attribute value.

The extension determines a letter qualifies as available by using the WNetGetConnection API to determine if the current letter in the list of letters is mapped to a network drive or associated with a physical device.

Drive Letter not in use

A drive letter not mapped to a shared folder instructs the extension to check if it previously saved an available drive letter. If the extension already knows of the first available drive letter, then it performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. The extension purposely cycles through the entire list to ensure no other drive letter within the range of drive letters is mapped to the specified location.

If the extension is unaware of any previously saved available drive letter, then it ensures the current letter is not associated with a physical device. A drive letter associated with a physical device is not eligible as an available drive letter. The extension performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. A drive letter making it to this part of the process has passed all the criteria needed to become the first available drive letter. The extension saves this drive letter and continues the evaluation over again using the next letter in the list.

Drive Letter in use by a Mapped Drive

A drive letter mapped to a shared folder instructs the extension to evaluate the shared folder location of the existing mapped drive to the <path> attribute value in the inner element configuration.

If the shared folder path matches the <path> attribute value, then the extension considers the mapped drive already connected. The extension writes "Path is connected already on letter" along with the current drive letter to the trace log, discontinues iterating the drive letters, and proceeds to the Drive Label portion of the core processing phase.

The extension completes its iteration through the drive letter list and validates a free drive letter exits. If the extension remains unaware of an available drive letter and it did not locate an existing drive letter mapped to the network path identical to the <path> attribute value, then the extension aborts its processing, and returns a not connected error. Otherwise, the extension is aware of an available drive letter and must now connect the shared network folder to the available drive letter.

Connecting to First Available Drive

The extension must create a new mapped drive using the discovered first available drive letter and the <path> attribute value. To do this, the extension determines the type of token used by the current user and the value of the <persistent> attribute.

If the user has a full token or if the <persistent> attribute equals false (zero [1]), then the extension use the WNetUseConnection API to create a new mapped drive where the drive letter

used is the discovered first available drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value, which is false, is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write "Connected with access name…" to the trace log.

A user with a restricted token (User Account Control enabled) indicates to the extension that it must also create the new drive mapping while impersonating the user using the user's restricted token.

The extension stops impersonating the current user with the user's full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write "Failed to revert from the user's full token" to the trace log, and to return the underlying error code that is the result of failure to cease impersonation.

Any failure in requesting a restricted token for the current user instructs the extension to abort its processing and to return the underlying error code that is the result of the failure to when requesting the restricted token.

A failure occurring when extension impersonates the user using the restricted token cause the extension to abort its processing, write "Failed to impersonate with user's restricted token" to the trace log, and return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write "Failed to connect drive with restricted token" to the trace log along with the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write "Connected with access name…" to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token cases the extension to write "Failed to revert user from restricted context" to the trace log. A failure impersonating the user with their full token causes the extension to write "Failed to impersonate with user's full token" to

the trace log. The extension continues processing by apply the configured label to the newly created mapped drive.

Drive Label

To this point in the process, the extension has successfully deleted a mapped drive, if discovered and created a mapped network drive using either a specific drive letter or a first available drive letter. Also, the extension has mapped the drive using the user's full token and restricted token, if needed.

The extension needs to apply the <label> attribute value to the newly created drive letter. The extension ensures a value is present for the <label> attribute and that the mapped drive was newly created and not an existing local or remote drive letter. Once the extensions confirms these conditions, it replaces any backslash (\) characters in the <path> attribute's value with pound (#) characters. Next, the extension creates a registry key using the modified <path> value under

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Mountpoints2
```

The extension creates a new registry value named _LabelFromReg. The value assigned to _LabelFromReg is the original, unmodified <path> attribute value. The extension ignores any failures that may occurs during drive label processing.

Hide or Show Drives

The Group Policy Drive Map extension can show or hide drives using two configurations during the replace, create, and update actions. It can hide or show a single, specified drive letter and it can hide or show all drive letters. The extension processes all drives first and specified drive last. Thus, the outcome of configuring both results in the specified drive configuration have precedence over the all drives configuration.

The extension determines if it needs to show or hide any drive letters by checking the <thisDrive> and <AllDrives> attribute values. If both attribute values equal "NOCHANGE' then the extension performs no other actions and processes the epilog portion of core processing. Otherwise, the extension begins the hide or show drives process by changing to the security context to System. Any failure in changing to System security context causes the extension to abort its processing and return the underlying error responsible for the failure.

All Drives

The extension determines if it needs to hide or show all drive letters by evaluating the value of the <AllDrives> attribute. If the attribute's value equals "SHOW" then the extension deletes the NoDrives registry value from the following registry hive of the current user.

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

If the attribute's value equals "HIDE" then the extension writes the hexadecimal value 0x03FFFFFF to the NoDrives registry value of the previously mentioned registry hive of the current user (the extension creates the NoDrives registry value if it does not exist). The extension then processes hide or show settings for a specific drive letter.

Specified Drive

The extension validates that the <thisDrive> attribute equals “SHOW” or “HIDE”. Next, the extension validates the <letter> attribute value is not blank. If the <letter> attribute value is blank, then the extension aborts its processing and returns an invalid argument error.

The extension retrieves the NoDrives registry value from the following registry hive of the current user

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

The extension uses an algorithm based on information in Microsoft Knowledgebase article 231289, Using Group Policy Objects to hide specified drives, to determine equivalent bit value of the <letter> attribute value.

If the <thisDrive> attribute value equals “HIDE” then the extension enables that bit in the NoDrives registry value. If the <thisDrive> attribute value equals “SHOW” then the extension disabled the bit in the NoDrives registry value.

After modifying the NoDrives registry value, the extension checks its value to ensure it does not equal zero (0). If the newly calculated NoDrives registry value equals zero, then the extension deletes the NoDrives registry value. Otherwise, the extension saves the newly calculated NoDrives registry value to the registry.

Epilogue

The extension completes the processing by updating its %LastDriveMapped% environment variable and by securely deleting the clear text instance of the <cPassword> attribute value. Additionally, if the <persistent> attribute value equals true (one [1]), then the extension signals to the Group Policy infrastructure that the next foreground processing must be synchronous.

Replace action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Update

The update action represents a configuration where the extension refreshes an existing mapped drive based on the configuration in the current inner element. It is important to understand the difference between replace and update.

The replace action typically deletes the existing mapped drive and creates a new one—essentially changing all configurations elements of the mapped drive. The update action only updates a portion of configuration elements in an existing mapped drive.

Prologue

The extension extracts the values for the <path>, <letter>, <persistent>, <username>, and <cPassword> attributes from the configuration information. If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extension continues extracting the values for the <allDrives>, <thisDrive>, and <useLetter> attributes. The extension completes the prologue portion of the processing by impersonating the

user logging on the computer using the user's full authentication token. Any failure attempting the impersonation results in the extension aborting its processing and to return the underlying error code that resulted in the impersonation failure.

Core

The core phase of update action processing is to use the configuration information read during the prologue phase to identify an existing mapped network drive and, if found, update drive label and visibility settings, or create a newly mapped network drive for the current user.

Specified Drive Exists

A <letter> value that currently exists for the user and is mapped to a shared network folder or a physical device causes the extension to continue processing by applying the configured label portion of the preference item to the existing drive letter.

Non-existent Specified Drive

A <letter> value that does not exists for the user causes the extension to create a new mapped network drive.

Create Mapped Network Drive

The extension determines if the current user's token is not a restricted token and the <persistent> attribute value equals false (zero [0]). Upon validating these factors, the extension uses the WNetUseConnection API to create a newly mapped drive using the value of the <letter> attribute as the drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and to return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write "Connected with access name…" to the trace log.

If the extension determines the current user's token is a restricted token (User Account Control is enabled) then the extension must create the mapped drive under the restricted token rather than the full token. This explains why mapped drives for administrators and some users are not always visible in Windows Explorer.

The extension stops impersonating the current user with the user's full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write "Failed to revert from the user's full token" to the trace log, and return the underlying error code that is the result of failure to cease impersonation. Any failure in requesting a restricted token for the current user instructs the extension to abort its processing, write "Failed to impersonate with user's restricted token" to the trace log, and return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value

of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write “Failed to connect drive with restricted token” along with the underlying error code it received from the Windows API to the trace log. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token causes the extension to write “Failed to revert user from restricted context” to the trace log. A failure impersonating the user with their full token causes the extension to write “Failed to impersonate with user’s full token” to the trace log. The extension continues processing by applying the configured label to the newly created mapped drive.

First Available Drive Letter, Starting at

A <useletter> attribute value equaling false (zero [0]) determines the desired configuration is to update an existing mapped drive using the first available letter starting with the <letter> attribute value.

In this configuration, the extension uses the <letter> attribute value as the starting point of the drive letters through which it searches for an existing mapped drive. The <path> attribute value becomes the matching criterion that determines if the mapped drive exists.

Determine the First Available Drive Letter

The extension needs to determine the first available drive letter before it can map any network location. To accomplish this, the extension creates a list of lower case letters starting with the <letter> attribute value and ending with the letter ‘z’. The extension iterates through each letter in the list to determine the first available drive letter, starting with the <letter> attribute value.

The extension determines a letter qualifies as available by using the WNetGetConnection API to determine if the current letter in the list of letters is mapped to a network drive or associated with a physical device.

Drive Letter not in use

A drive letter not mapped to a shared folder instructs the extension to check if it previously saved an available drive letter. If the extension already knows of the first available drive letter, then it performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. The extension purposely cycles through the entire list to ensure no other drive letter within the range of drive letters is mapped to the specified location.

If the extension is unaware of any previously saved available drive letter, then it ensures the current letter is not associated with a physical device. A drive letter associated with a physical device is not eligible as an available drive letter. The extension performs no other actions on the current drive letter and starts the evaluation over again using the next letter in the list. A drive

letter making it to this part of the process has passed all the criteria needed to become the first available drive letter. The extension saves this drive letter and continues the evaluation over again using the next letter in the list.

Drive Letter in use by a Mapped Drive

A drive letter mapped to a shared folder instructs the extension to evaluate the shared folder location of the existing mapped drive to the <path> attribute value in the inner element configuration.

If the shared folder path matches the <path> attribute value, then the extension considers the mapped drive already connected. The extension writes “Path is connected already on letter” along with the current drive letter to the trace log, discontinues iterating the drive letters, and proceeds to the Drive Label portion of the core processing phase.

The extension completes its iteration through the drive letter list and validates a free drive letter exits. If the extension remains unaware of an available drive letter and it did not locate an existing drive letter mapped to the network path identical to the <path> attribute value, then the extension aborts its processing, and returns a not connected error. Otherwise, the extension is aware of an available drive letter and must now connect the shared network folder to the available drive letter.

Connecting to First Available Drive

The extension must create a new mapped drive using the discovered first available drive letter and the <path> attribute value. To do this, the extension determines the type of token used by the current user and the value of the <persistent> attribute.

If the user has a full token or if the <persistent> attribute equals false (zero [1]), then the extension use the WNetUseConnection API to create a new mapped drive where the drive letter used is the discovered first available drive letter. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value, which is false, is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response from the Windows API instructs the extension to abort its processing and return the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log.

A user with a restricted token (User Account Control enabled) indicates to the extension that it must also create the new drive mapping while impersonating the user using the user’s restricted token.

The extension stops impersonating the current user with the user’s full token and requests Windows to provide a restricted token for the user. Any failure in the cessation of impersonation instructs the extension to abort its processing, write “Failed to revert from the user’s full token” to the trace log, and to return the underlying error code that is the result of failure to cease impersonation.

Any failure in requesting a restricted token for the current user instructs the extension to abort its processing and to return the underlying error code that is the result of the failure to when requesting the restricted token.

A failure occurring when extension impersonates the user using the restricted token cause the extension to abort its processing, write “Failed to impersonate with user’s restricted token” to the trace log, and return the underlying error code that is the result of the failure to impersonate using the restricted token.

While impersonating the user using the restricted token, the extension use the WNetUseConnnection API to create a new mapped drive where the drive letter used is the value of the <letter> attribute. The shared folder used for the mapped drive is the value of the <path> attribute. If the <username> and <cPassword> attributes contains values, then their respective values are the credentials used to authenticate and authorize the connection to the shared folder. Additionally, the <persistent> attribute value is passed to the Windows API to indicate if the newly created mapped drive should persist between logons and reboots.

An unsuccessful response instructs the extension to abort its processing, write “Failed to connect drive with restricted token” to the trace log along with the underlying error code it received from the Windows API. A successful response from Windows API instructs the extension to write “Connected with access name…” to the trace log. Then, the extension discontinues impersonating the current user under the restricted token and returns impersonating the user using the full token.

A failure reverting from impersonating the user with a restricted token cases the extension to write “Failed to revert user from restricted context” to the trace log. A failure impersonating the user with their full token causes the extension to write “Failed to impersonate with user’s full token” to the trace log. The extension continues processing by apply the configured label to the newly created mapped drive.

Drive Label

To this point in the process, the extension has successfully deleted a mapped drive, if discovered and created a mapped network drive using either a specific drive letter or a first available drive letter. Also, the extension has mapped the drive using the user’s full token and restricted token, if needed.

The extension needs to apply the <label> attribute value to the newly created drive letter. The extension ensures a value is present for the <label> attribute and that the mapped drive was newly created and not an existing local or remote drive letter. Once the extensions confirms these conditions, it replaces any backslash (\) characters in the <path> attribute’s value with pound (#) characters. Next, the extension creates a registry key using the modified <path> value under

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Explorer\Mountpoints2
```

The extension creates a new registry value named _LabelFromReg. The value assigned to _LabelFromReg is the original, unmodified <path> attribute value. The extension ignores any failures that may occurs during drive label processing.

Hide or Show Drives

The Group Policy Drive Map extension can show or hide drives using two configurations during the replace, create, and update actions. It can hide or show a single, specified drive letter and it can hide or show all drive letters. The extension processes all drives first and specified drive last. Thus, the outcome of configuring both results in the specified drive configuration have precedence over the all drives configuration.

The extension determines if it needs to show or hide any drive letters by checking the <thisDrive> and <AllDrives> attribute values. If both attribute values equal "NOCHANGE' then the extension performs no other actions and processes the epilog portion of core processing. Otherwise, the extension begins the hide or show drives process by changing to the security context to System. Any failure in changing to System security context causes the extension to abort its processing and return the underlying error responsible for the failure.

All Drives

The extension determines if it needs to hide or show all drive letters by evaluating the value of the <AllDrives> attribute. If the attribute's value equals "SHOW" then the extension deletes the NoDrives registry value from the following registry hive of the current user.

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

If the attribute's value equals "HIDE" then the extension writes the hexadecimal value 0x03FFFFFF to the NoDrives registry value of the previously mentioned registry hive of the current user (the extension creates the NoDrives registry value if it does not exist). The extension then processes hide or show settings for a specific drive letter.

Specified Drive

The extension validates that the <thisDrive> attribute equals "SHOW" or "HIDE". Next, the extension validates the <letter> attribute value is not blank. If the <letter> attribute value is blank, then the extension aborts its processing and returns an invalid argument error.

The extension retrieves the NoDrives registry value from the following registry hive of the current user

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer
```

The extension uses an algorithm based on information in Microsoft Knowledgebase article 231289, Using Group Policy Objects to hide specified drives, to determine equivalent bit value of the <letter> attribute value.

If the <thisDrive> attribute value equals "HIDE" then the extension enables that bit in the NoDrives registry value. If the <thisDrive> attribute value equals "SHOW" then the extension disabled the bit in the NoDrives registry value.

After modifying the NoDrives registry value, the extension checks its value to ensure it does not equal zero (0). If the newly calculated NoDrives registry value equals zero, then the extension deletes the NoDrives registry value. Otherwise, the extension saves the newly calculated NoDrives registry value to the registry.

Epilogue

The extension completes the processing by updating its %LastDriveMapped% environment variable and by securely deleting the clear text instance of the <cPassword> attribute value. Additionally, if the <persistent> attribute value equals true (one [1]), then the extension signals to the Group Policy infrastructure that the next foreground processing must be synchronous.

Update action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Dependencies

The Group Policy Drive Maps extension is one of many extensions provided in a single dynamic linked library, gpprecl.dll. This extension has one or more dependencies on the following dynamic linked libraries that are included in the Windows operating system.

1. ACTIVEDS.DLL
2. ADVAPI32.DLL
3. IPHLPAPI.DLL
4. KERNEL32.DLL
5. MPR.DLL
6. MSI.DLL
7. MSVCRT.DLL
8. NETAPI32.DLL
9. NTDLL.DLL
10. OLE32.DLL
11. OLEAUT32.DLL
12. POWRPROF.DLL
13. RPCRT4.DLL
14. SECUR32.DLL
15. SETUPAPI.DLL
16. SHELL32.DLL
17. SHLWAPI.DLL
18. USER32.DLL
19. USERENV.DLL
20. VERSION.DLL
21. WINSPOOL.DRV
22. WINSTA.DLL
23. WLDAP32.DLL
24. WS2_32.DLL
25. WTSAPI32.DLL
26. XMLLITE.DLL

Interaction with other components

The Windows operating system enables many different configurations to accomplish a single task. Mapping shared network folders to local drive letters is one of those tasks.

The most common configuration for mapping network drives occurs during the user logon. Historically, IT departments have accomplished this configuration using network logon scripts—scripts that process as part of the user logon. Another configuration for mapping Home drives is to configure the shared network folder on the user's object in Active Directory. Windows clients looking for this configuration and automatically connect the shared network folder to the configured drive letter. Users can map network drives using Windows Explorer and make them persistent between computer reboots. Lastly, you can use the Group Policy Drive Map extension to map drives during logon.

Flexible configurations are important; however, understanding that incorporating the one or more of the configurations into one solution needs to be thoroughly tested for each specific environment as many of these technologies can encroach on the other thereby creating unpredictable results. A suggested best practice is to pick one configuration and use it exclusively in your environment, especially when mapping network drives during logon. Each configuration maps the drive during different stages of the logon process. Incorporating multiple configurations could be counterproductive and cause conflict between the two configurations. This conflict is likely to produce unwanted results, increases the complexity of the solution, increases administrative overhead of the solution, and increase the difficulty of troubleshooting the cause of the problem.

Logging

In addition to the event messages recorded in the application log, the Group Policy Drive Map extension has trace logging to help diagnose preference items that are not applying or undesired or unexpected results from preference items that are applying.

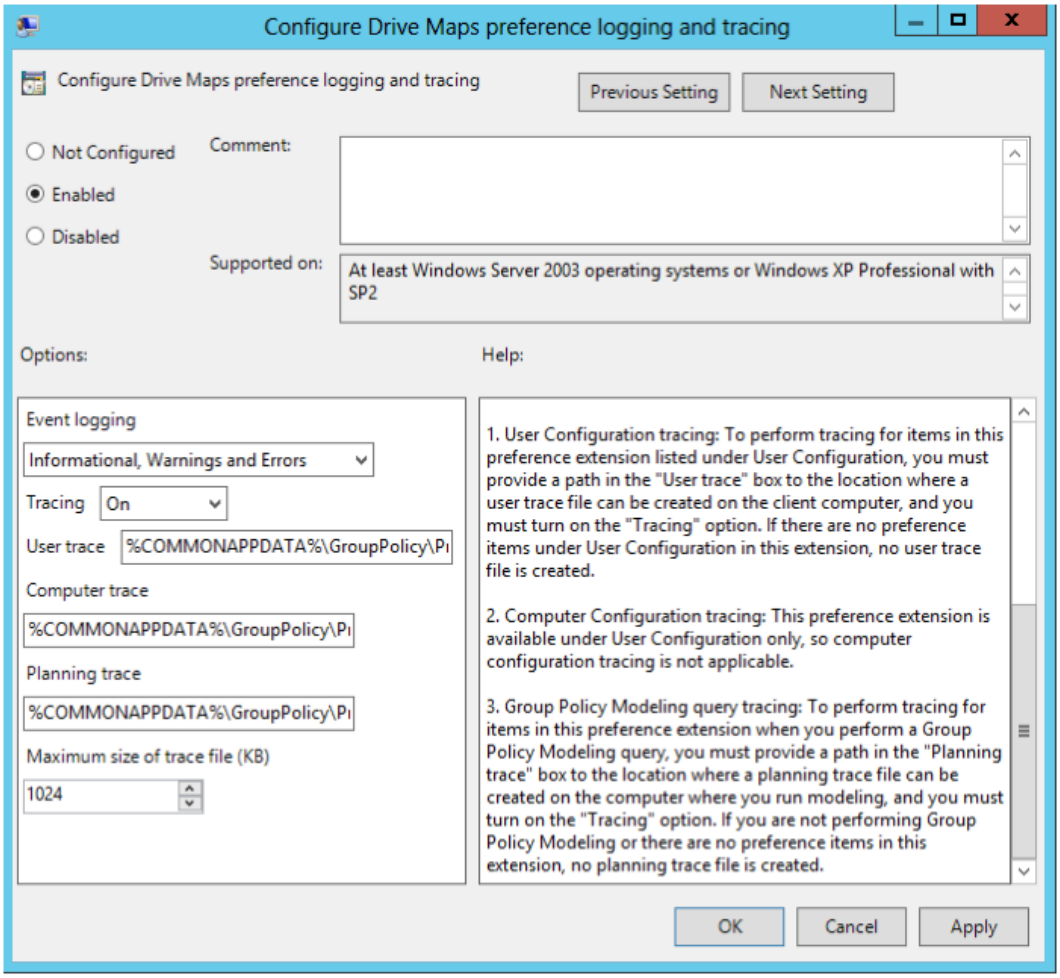

You enable Group Policy Drive Map logging using Computer\Policies namespace of the Group Policy Management Editor. The **Configure Drive Maps preference logging and tracing** policy setting resides under the **Administrative Templates\System\Group Policy\Logging and tracing node**. The policy setting enables a variety of configuration options that affect the types of events the extension reports to the event log, trace logging, and the filename and location of the trace logs.

The best configuration for capturing Group Policy Drive Map information for troubleshooting is to configure **Event Logging** using the **Informational, Warnings, and Errors** option and configure **Tracing** to **On**.

Trace logging file name and locations are configurable to any name and location to where the user or computer has read and write permissions. The default paths for these files use a Group Policy Preference **variable %COMMONAPPDATA%\GroupPolicy\Preference\Trace\<filename>**. The Group Policy Preference %COMMONAPPDATA% variable equates to the All Users' application data folder.

Starting with Windows Vista, the All Users' folder was deprecated with the introduction of the version 2 user profile structure. To ensure compatibility among applications, Windows creates a

symbolic link named All Users in the file system under the %SYSTEMDRIVE%\Users folder. The symbolic link redirects file system requests to the All Users folder to the **C:\ProgramData** folder. Therefore, the default location for Group Policy Preference trace logs is **C:\ProgramData\GroupPolicy\Preference\Trace\<filename>**.

Printers

File and print services has been a long requirement of the traditional network operating system. Over the years, networks have grown from a few computers in a single room to campuses of buildings hosting hundreds to thousands of computers around the globe, all of which need to print.

Regardless of how green an organization is, the need to print remains and has only become more complicated as networks span floors and buildings and mobile users are in the Mobile, Alabama branch today and the Charlotte, North Carolina branch tomorrow. IT networks added more printers; however, users were challenged to remember what printer to use depending on their building, floor or geographical location. To resolve these challenges, IT professionals developed mapped printer connections.

Mapped printer connections attacked the problem of users needing to remember what printer to use by reconciling users to printers based on group memberships or other criteria and then implemented that logic in complicated user logon scripts. User burden was reduced, but at the expense of the administrative overhead. Now IT administrators need to learn a complex logic scripting and programming experience. Group Policy Printers resolve aims to problem a simple resolution for both user and administrators.

Overview

The Group Policy Printers extension enables you to configure local, shared, or TCP/IP printers for an enterprise environment without the complicated mess of login scripts. You configure the Group Policy Printers extension using the same tools to configure Group Policy. As a result, managing printer connections for users and computers is consistent and does not require programming or scripting knowledge.

Capabilities

The Group Policy Printers extension provides the basic enterprise printer deployment and management by enabling you to create, replace, and delete local, shared, and TCP/IP based printers. In addition, the Group Policy Printers extension enables updating settings such as description and default printer status on these types of printers.

You can combine Printer preference items with Group Policy Preferences Item-level targeting to achieve granular and complex connected printers scenarios that remain simple, easy to manage, and do not use or require knowledge of scripting languages.

Pushed Printers

The Group Policy Printers extension should not be confused with the Pushed Printer Group Policy extension. Pushed printers is an alternative technology used to deploy printers in the enterprise that stores data in Active Directory as well as on SYSVOL share of a domain controllers. Group Policy Printers exclusively stores its configuration information on the SYSVOL share of domain controllers.

Processing Complete

The extension has completed processing the all inner and outer elements in the configuration file. To finalize the processing, the extension must updated the Group Policy history based on its actions during processing. Understand, the actions are not based on individual preference items, but rather the entire configuration file.

The extension caches the configuration files, if configured, for all Group Policy objects applied from the new or changed GPO list in the Group Policy Preference history. A configuration file hosted in a GPO appearing in the Out-of-scope GPO list is removed from the Group Policy Preference history.

Local Printer

The Group Policy Printers preference extension allows you to create, configure, and delete local printers by using the Local Printer preference item.

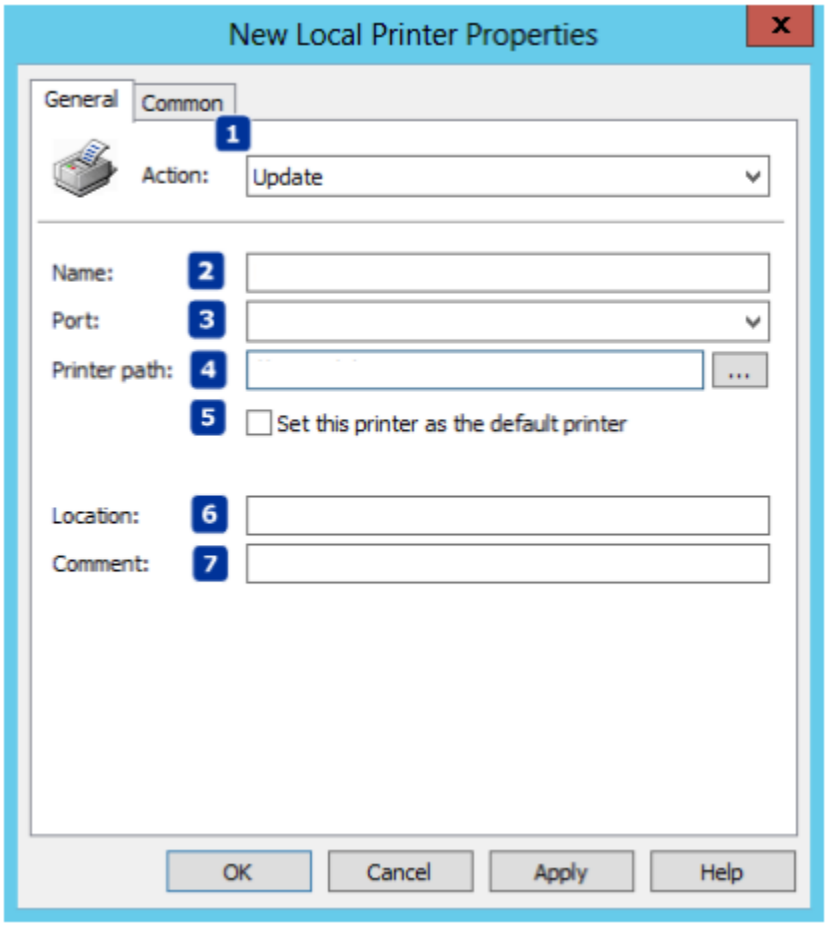

The Local Printer preference item provides several configuration points to help achieve most local printer scenarios that exist in enterprise environments.

Action (1)

This type of preference item provides a choice of four actions: **Create**, **Replace**, **Update**, and **Delete**. The behavior of the preference item varies with the action selected and whether the printer connection already exists.

Create

Create a new local printer. If a local printer with the same name exists, then it does not modify it.

Delete

Remove a local printer with the same name. The extension performs no action if the local printer does not exist. If a printer with the same name exists, the extension only removes the connection. It does not remove the printer driver.

Replace

Delete and **Create** the local printer. The net result of the **Replace** action overwrites all existing settings associated with the local printer. If the local printer does not exist, then the **Replace** action creates a new local printer.

Update

Rename or modify a local printer. The action differs from **Replace** in that it updates the settings defined within the preference item. All other settings remain as they were previously configured. If the local printer does not exist, then the **Update** action creates a new local printer.

Name (2)

The name of the targeted local printer. The preference extension creates a new local printer with this name if the local printer does not exist. If the printer exists, the preference extension uses the local printer with this name as the target of the requested action. Press F3 to display a list of variables from which you can select.

Port (3)

Choose a local port from the supplied list or type the name of the local port.

Printer Path (4)

The fully qualified UNC path of a shared printer connection. The preference extension uses this shared connection as an installation source for the printer driver. The actual printer should be physically connected to the workstation.

Set this printer as the default printer (5)

Makes the local printer the default Windows printer for the current user.

Location (6)

A description to where the printer is located. This information appears in the **Location** text box on the printer properties dialog box. Press F3 to display a list of variables from which you can select.

Comment (7)

Text that provides additional comments about the printer. This information appears in the Comments text box on the printer properties dialog box. Press F3 to display a list of variables from which you can select.

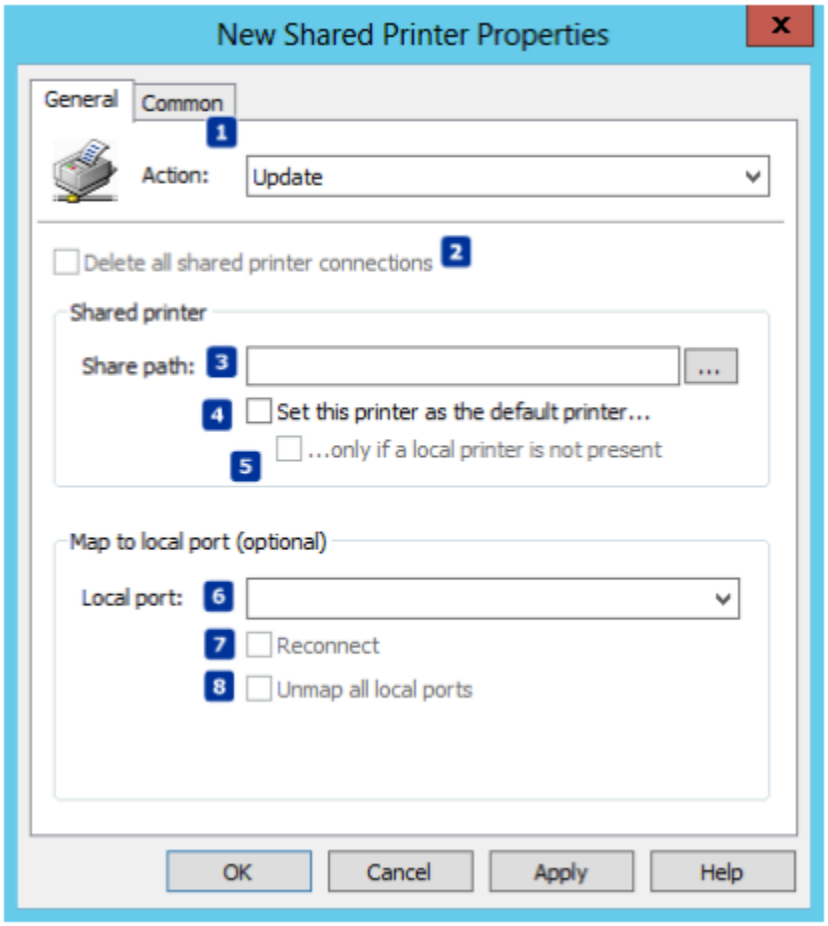

Shared Printer

The Printers preference extension allows you to create, configure and delete shared printer connections by using the Shared Printers preference item.

Action (1)

This type of preference item provides a choice of four actions: **Create**, **Replace**, **Update**, and **Delete**. The behavior of the preference item varies with the action selected and whether the printer connection already exists.

Create

Create a new local printer. If a local printer with the same name exists, then it does not modify it.

Delete

Remove a local printer with the same name. The extension performs no action if the local printer does not exist. If a printer with the same name exists, the extension only removes the connection. It does not remove the printer driver.

Replace

Delete and **Create** the local printer. The net result of the **Replace** action overwrites all existing settings associated with the local printer. If the local printer does not exist, then the **Replace** action creates a new local printer.

Update

Rename or modify a local printer. The action differs from **Replace** in that it updates the settings defined within the preference item. All other settings remain as they were previously configured. If the local printer does not exist, then the **Update** action creates a new local printer.

Action (1)

This type of preference item provides a choice of four actions: **Create**, **Replace**, **Update**, and **Delete**. The behavior of the preference item varies with the action selected and whether the printer connection already exists.

Create

Create a new shared printer connection. If a local printer with the same name exists, then it does not modify it.

Delete

Remove a shared printer connection with the same share path. The extension performs no action if the shared printer connection does not exist. If a printer with the same share path exists, the extension only removes the connection. It does not remove the printer driver.

Replace

Delete and **Create** the shared printer connection. The net result of the Replace action overwrites all existing settings associated with the shared printer connection. If the shared printer connection does not exist, then the **Replace** action creates a new shared printer connection.

Update

Modify a shared printer connection. The action differs from **Replace** in that it updates the settings defined within the preference item. All other settings remain as they were previously configured. If the shared printer connection does not exist, then the **Update** action creates a new shared printer connection.

Share Path (3)

A fully qualified UNC path of a shared printer.

Set this printer as the default printer (4)

Make the shared printer connection the default Windows printer for the current user.

Only if local printer is not present (5)

Bypasses changing the default printer if there is a local printer configured on the computer. This setting is unavailable until you select the Set this printer as the default printer check box.

A local printer is any printer that is not connected to a shared network printer. This includes physical printers connected to parallel, serial, and USB ports, TCP/IP printers, and virtual printers installed through software.

Local Port (6)

Choose a local port to which you want the shared connection mapped (optional).

When **Local Port** contains a value and the preference item action is set to **Delete**, the preference extension deletes the shared printer connection associated with that local port.

Reconnect (7)

Select this check box if you want the shared printer connection persistent. This option is unavailable until you choose a value from the Local Port list.

Unmap all local ports (8)

Select this check box if you want to delete shared printer connections from all local ports. This setting is available only when the preference item's action is set to Delete.

TCP/IP Printer (Port printer)

The Printers preference extension allows you to create, configure, and delete TCP/IP printers by using the TCP/IP Printer preference item.

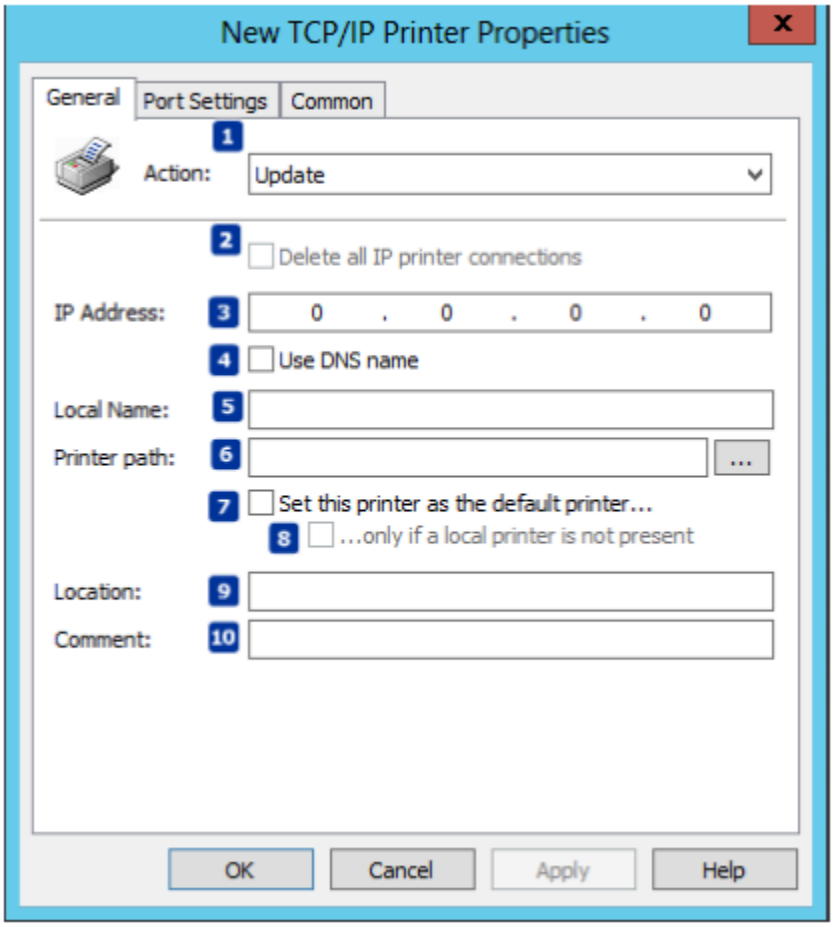

General Tab

Action (1)

Create

Create a new TCP/IP printer connection. If a TCP/IP printer connection with the same IP address exists, then it does not modify it.

Delete

Remove a TCP/IP printer connection with the same IP address. The extension performs no action if the TCP/IP printer connection does not exist. The Delete action does not remove the printer driver or port. It only removes the TCP/IP printer connection.

Replace

Delete and Create the TCP/IP printer connection. The net result of the Replace action overwrites all existing settings associated with the TCP/IP printer connection. If the TCP/IP printer connection does not exist, then the Replace action creates a new TCP/IP printer connection.

Update

Modify a TCP/IP printer connection. The action differs from **Replace** in that it updates the settings defined within the preference item. All other settings remain as they were previously configured. If the TCP/IP printer connection does not exist, then the **Update** action creates a new TCP/IP printer connection.

Delete all IP printer connections (2)

Deletes all TCP/IP printer connections for the current user. This setting is available only when the preference item's action is set to **Delete**.

IP Address/DNS name (3)

The IP address of the remote printer. Or select the **Use DNS name** check box and type the fully qualified domain name of the remote printer.

Use DNS Name (4)

Toggles the formatting of the IP Address/ DNS name text box. When Use DNS Name cleared, the preference item configuration formats IP for Address text box for an IPv4 address. When Use DNS Name is selected, the preference item configuration formats the DNS name text box for a fully qualified domain name. The Group Policy Printers extension does directly support IPv6 addresses.

Local Name (5)

The local name of the targeted TCP/IP printer connection. The preference extension creates a new TCP/IP printer connection with this name if one does not exist. If a TCP/IP printer connection with this name exists, the preference extension uses the TCP/IP printer with this name as the target of the requested action. Press F3 to display a list of variables from which you can select.

The Group Policy Printers preference extension uses Local Name to determine if a TCP/IP printer exists when a Local Name is provided. Otherwise, the preference extension uses the TCP/IP address or the DNS Name to determine if the TCP/IP connection exits.

Printer path (6)

Type a fully qualified UNC path of a shared printer connection. The preference extension uses this shared connection as an installation source for the printer driver.

Set this printer as the default printer (7)

Make the TCP/IP printer connection the default Windows printer for the current user.

Only if local printer is not present (8)

Bypasses changing the default printer if there is a local printer configured on the computer. This setting is unavailable until you select the **Set this printer as the default printer** check box.

A local printer is any printer that is not connected to a shared network printer. This includes physical printers connected to parallel, serial, and USB ports, TCP/IP printers, and virtual printers installed through software.

Location (9)

A description to where the printer is located. This information appears in the **Location** text box on the printer properties dialog box. Press F3 to display a list of variables from which you can select.

Comment (10)

Text that provides additional comments about the printer. This information appears in the **Comments** text box on the printer properties dialog box. Press F3 to display a list of variables from which you can select.

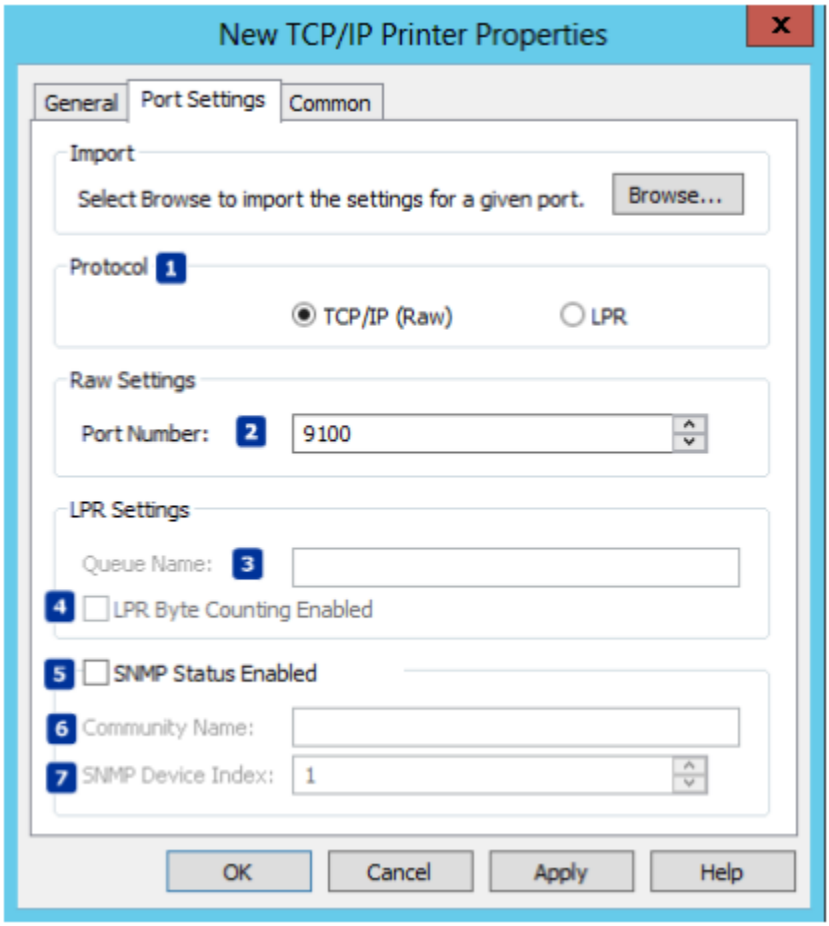

Port Settings Tab

Protocol (1)

Specifies either a TCP/IP (Raw) printer or an RFC1179 compliant printer.

Port Number (2)

The port number used when communicating to the printer. The default port number is 9100.

LPR Settings (3)

The name of the LPR Queue to use when communicating with the printer. The default is "LPR".

LPR Byte Counting Enabled (4)

When selected, document sizes are calculated locally prior to being sent to the printer. Some LPR ports require this setting to be enabled to prevent data loss.

SNMP Status Enabled (5)

Specifies whether or not the printer port will support RFC1759

Community Name (6)

Specifies the SNMP community name used by the printer. The default is "public".

SNMP Device Index (7)

Specifies the SNMP device index used by the printer. The default is 1.

How does it work?

The Group Policy Printer preference extension is a Group Policy client side extension that is hosted in the gpprefcl.dll dynamic linked library. As part of the Group Policy specification, each component of data stored in a Group Policy object must have two 128-bit unique identifiers in a string format.

Client-side and Tool Identifiers

The first identifier associates the portion of data stored in the Group Policy object with the entity responsible for consuming, or processing the data. This identifier is known as the client-side extension identifier.

The second identifier associates the portion of data stored in the Group Policy object with the entity responsible for authoring, or managing it. This identifier is known as the snap-in identifier, or tool identifier.

The client –side extension unique for the Group Policy Printers preference extension is {BC75B1ED-5833-4858-9BB8-CBF0B166DF9D}. The snap-in unique identifier for the Group Policy Printers preference extension is {A8C42CEA-CDB8-4388-97F4-5831F933DA84}.

Group Policy uses the client-side identifier in many locations. First, the client side identifier is stored in the registry of the local computer. The registry location is

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions
```

The Group Policy infrastructure reads the preceding registry location to discover all Group Policy client-side extensions on the computer. Client-side extensions not listed at this location do not process during the application of Group Policy. The registration information includes a sub-key with the client-side extensions unique identifier as the name. Several registry values reside beneath the identifier key. These values are used by the Group Policy infrastructure, but also provide valuable information such as the name of the client side extension (in the **Default** registry value), the binary file that hosts the client-side extension code (in the **DLLName** registry value), and the event source and event log the extension uses to report information, warnings, and errors (in the **EventSources** registry value, separated by a comma, respectively).

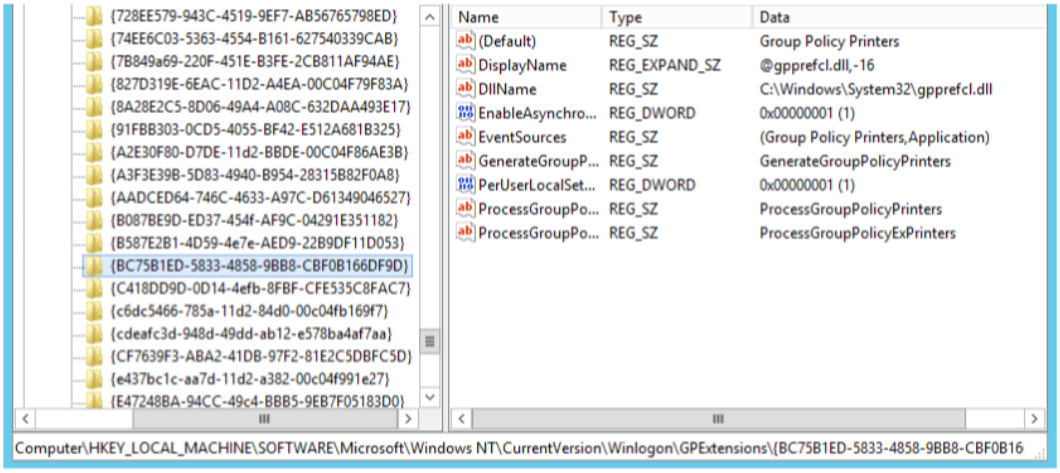

When a Group Policy object is configured to include data from a specific client side extension, the editor must include the client-side extension and snap-in unique identifiers as part of the data saved in the Group Policy object.

Extensions specific data configured to apply to computers is saved to the **gPCMachineExtensionNames** attribute of the Group Policy container object. Extension specific data configured to apply to users is saved to the gPCUserExtensionNames attribute. If extension specific data is configured to apply to both user and computer, then the editor saves the extension identifier and the snap-in identifier to both attributes of the Group Policy container object

The **gPCMachineExtensionNames** and the **gPCUserExtensionNames** attribute can hold one or more groups of unique identifiers. Each identifier group is comprised of an extension identifier followed by one or more snap-in identifiers. The group of identifiers are bracketed accordingly by surrounding both identifiers within square brackets ([ ]).

```
extensionID1snapinID1snapinID2…… [ extensionIDnsnapinIDn…]
```

The unique identifier grouping must be listed in ascending alphanumeric order. Saving these in any other order prevents Group Policy from correctly applying or prevents you from editing existing extensions specific data within the Group Policy object.

Storing the extension identifier and the snap-in identifiers in the Group Policy object provides hints that optimize Group Policy processing and management. Processing is optimal because the Group Policy infrastructure knows prior to processing which extensions are required to apply the data hosted in the Group Policy object. Authoring is optimal because the editor has hints to determine what specific snap-ins are needed to edit the data.

XML Configuration

The Group Policy Printers preference extension stores all configuration data associated with the Group Policy object in an XML-formatted file. The Printers.xml file resides in the Printers folder in the Group Policy template that is hosted and shared on all domain controllers' SYSVOL share.

The Group Policy template location varies for each Group Policy object and is the string value stored in the gPCFileSysPath attribute on the Group Policy container object.

```
\\dns_domain_name\Sysvol\dns_domain_name\Policies\group_policy_id
```

XML Declaration

The Group Policy Printers configuration file starts with the XML declaration. The declaration describes two characteristics of the XML content: version and encoding.

The version declaration identifies the version of XML syntax the document uses. The encoding declaration identifies the characters used throughout the document.

The XML declaration used in the Group Policy Printers configuration file uses a value of 1 for the XML version and the encoding value is UTF-8

```
<?xml version="1.0" encoding="UTF-8"?>
```

The Group Policy Printers configuration file does not include the optional standalone declaration.

Elements

Group Policy Printers configuration files use four XML elements to describe the configuration data: an outer element and an inner element for each printer preference type. The outer element serves as a logical container of the inner XML elements. The inner elements represent the configuration of one or more preference items belonging to the container (outer XML element).

Outer Elements

The Group Policy Printers configuration file uses an outer XML element named **Printers**. This element serves as a collection of inner elements where each inner element represents a printer preference item. The **Printers** outer element supports two attributes: **clsid** and **disabled**.

Attributes

The **Printers** outer element of the Group Policy Printers configuration file contains one required and one optional attribute.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of outer element. The Group Policy Printers client-side extension expects the outer element's clsid value to equal {1F577D12-3D1B-471e-A1B7-060317597B9C}. This is how the client-side extension identifies and validates the outer-element is of type **Printers**

Disabled

The **disabled** attribute is an optional attribute for the **Drives** outer element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A **Printers** outer element with a **disabled** attribute equaling true indicates the configuration for all inner elements is disabled. This prevents the Group Policy Printers extension from processing any of the inner elements of the configuration file. A **disabled** attribute equaling false indicates the configuration for all inner elements is enabled and that processing continues as normal.

A **Printers** outer element without a **disabled** attribute indicates that the outer element has never been disabled. In this configuration, the default value for the **disabled** attribute is false, which means the outer element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the outer element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

```
<Printers clsid=" {1F577D12-3D1B-471e-A1B7-060317597B9C}" disabled="0">
```

Inner Elements

The **Printers** outer element accepts three types of inner elements. The **LocalPrinter** inner element represents a single local printer preference item. The **SharedPrinter** inner element represents a single shared printer preference item. And the **PortPrinter** inner element represents a single TCP/IP printer preference item. There can be any combination of one or more **LocalPrinter**, **SharedPrinter**, or **PortPrinter** inner elements within the **Printers** outer element.

Each inner element contains attributes that describe the inner element. These attributes are jointly used by management editors and client side extensions as a way to identify one inner element from the next. The inner element attributes do not contain configuration data.

Local Printer Element

Attributes

ByPassErrors

The **byPassErrors** attribute is an optional attribute for the **LocalPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

One or more printer inner elements are contained within a single **Printers** outer element, which enables the client-side extension to process one or more preference items. A **LocalPrinter** inner element with a **byPassErrors** attribute equaling true indicates the Group Policy Printers extension should continue processing the next printer inner element regardless if the result of the current printer inner element is a failure. A **byPassErrors** attribute equaling false indicates the extension should stop processing subsequent printer inner elements if the result of the current inner element is a failure.

A **LocalPrinter** inner element without a **byPassErrors** attribute indicates that the inner element has never been configured to halt on errors. In this configuration, the default value for the

byPassErrors attribute is true, which means the any inner element failure does not halt the processing of subsequent inner elements. The **byPassErrors** attribute appears in the configuration the first time the inner element is configured to stop on errors. Once in this configuration, the value of the attribute changes with each respective configuration, but never removed.

Changed

The **changed** attribute is an optional attribute for the **LocalPrinter** inner element. The attribute uses a string value to represent the date and time of when the inner element was last modified in UTC. The string value uses the format of YYYY-MM-DD HH:MM:SS.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of inner element. The Group Policy Printers client-side extension expects the **LocalPrinter** inner element's **clsid** value to equal {F08996D5-568B-45f5-BB7A-D3FB1E370B0A}. This is how the client-side extension identifies and validates the inner element is of type **LocalPrinter**.

Disabled

The **disabled** attribute is an optional attribute for the **LocalPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A **LocalPrinter** inner element with a **disabled** attribute equaling true indicates the configuration for that inner element is disabled. This prevents the Group Policy Printers extension from processing that inner element in the configuration file. A **disabled** attribute equaling false indicates the configuration for that inner element is enabled and that processing continues as normal.

A **LocalPrinter** inner element without a **disabled** attribute indicates that the inner element has never been disabled. In this configuration, the default value for the **disabled** attribute is false, which means the inner element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the inner element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

Image

The **image** attribute is a numeric string value that represents the index of a bitmap resource used by the Group Policy Management editor as the icon image used to display the preference item in the editor. The value of the numeric string typically corresponds to the value of the **action** attribute in the **properties** element.

Image Value	Action Value
0	C
1	R

2	U
3	D

Name

The name attribute is a string value that represents the display name of the preference item. The **name** attribute has no impact on the processing of the preference item and is strictly used for management and reporting.

removePolicy

The **removePolicy** attribute is an optional attribute for the **LocalPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **LocalPrinter** inner element with a **removePolicy** attribute equaling true indicates the Group Policy Printers extension should remove the configuration included in the inner element if the Group Policy object hosting the preference item is not within scope of the user or computer. Once the Group Policy object hosting the preference item no longer applies to the user, the Group Policy Printers extension removes the effective results of the inner element's configuration from the user or computer.

The **removePolicy** attribute correlates to the value of the **action** attribute in the **Properties** element. A **LocalPrinter** inner element with a **removePolicy** attribute equal to true requires the value of the **action** attribute of **Properties** element to equal an uppercase R.

A **LocalPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item as configured.

A **LocalPrinter** inner element without a **removePolicy** attribute indicates the inner element has never been configured to remove when it no longer applies. In this configuration, the default value for the **removePolicy** attribute is false, which means the extension processes the preference item as configured. The **removePolicy** attribute appears in the configuration the first time the inner element is configured to remove items when it no longer applies. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

Status

The **status** attribute is string value that represents a message that the Group Policy Preference editor can display in the status area of the Microsoft Management Console. The **status** attribute remains as part of the configuration; however, is no longer being actively implemented.

Uid

The **uid** attribute is a 128-bit identifier represented as global unique identifier (GUID). The **uid** value uniquely identifies each inner element. This is how the client-side extension identifies and differentiates each LocalPrinter inner element.

UserContext

The **userContext** attribute is an optional attribute for the **LocalPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **LocalPrinter** inner element with a **userContext** attribute equaling true indicates the Group Policy Printers extension should create the mapped drive using the security context of the current user. A **LocalPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item using the security context of the local system.

A **LocalPrinter** inner element without a **userContext** attribute indicates the inner element has been manually configured to apply in the security context of the current user. In this configuration, the default value varies on a per-policy basis using the simple guidelines of computer preference items run in the security context of the local system and user preference items run in the security context of the current user.

The **userContext** attribute appears in the configuration the first time the inner element is manually configured to apply using the security context of the current user. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

The **userContext** attribute remains as part of the configuration; however, is no longer actively implemented. Most of the Group Policy Preference extensions programmatically switch between the current user and system security contexts as needed, irrespective of the **userContext** attribute value present in the inner element configuration. The recommended configuration is to leave each preference item to the defaults and to not manually configure the security context of any preference item.

```
<LocalPrinter clsid="{F08996D5-568B-45f5-BB7A-D3FB1E370B0A}" name="LocalPrinter"
status="Location" image="2" changed="2013-05-16 17:33:56" uid="{A10F181B-8938-44F6-9818-
84A28B3A37E6}" bypassErrors="1">
```

Properties Element

Each inner element requires a Properties element in the configuration file. Group Policy Preference extensions read the Properties element for the preference item configuration.

Attributes

The Properties element shares unique and common attributes among all the Group Policy Preferences inner element classes. The common attributes characterize abstract functionality provided by the inner element classes (preference items) such as create, delete, or update. However, the implementation of these actions varies across each inner element.

Action

The **action** attribute is an optional string attribute in the **Properties** element used to describe the action the Group Policy Printers extension performs for the associated inner element. That **action** attribute value is a single, uppercase letter of C, R, U, or D. If the **action** attribute is not present in the **Properties** element then the default value is U.

Create

An action attribute value equaling an upper case C instructs the extension to create a new local printer using the settings in the current **LocalPrinter** inner element.

The create action instructs the extension to create a new local printer only when the computer or current user does not have a local printer that shares the same name with the value in the **name** attribute.

If the computer or current user has an existing local printer name that matches the value in the **name** attribute then the extension ignores the instruction and does not report an error.

The create action, after successfully creating the new local printer, may also set the newly created local printer as the default local printer based on the value configured in the **default** attribute.

Delete

An action attribute value equaling an upper case D instructs the Group Policy Printers extension to remove an existing local printer using the settings in the current **LocalPrinter** inner element.

The delete action has two behaviors based on the value in the **deleteAll** attribute. When the **deleteAll** attribute value equals 1, the delete action removes all local printers for the computer or current user. A **deleteAll** attribute value equaling 0 instructs the extension to remove a local printer only when the computer or current user has a local printer with a name that matches the value of the **name** attribute.

If the computer or current user does not have an existing local printer that shares the same name with the value in the name attribute then the extension ignores the instruction and does not report an error.

Replace

An **action** attribute value equaling an upper case R instructs the extension to replace the existing local printer with the settings in the current **LocalPrinter** inner element. The replace action instructs the extension to perform the delete action first, and then the create action. The delete instruction removes an existing local printer using the criteria identified in the **Delete** action. The create instruction creates a new local printer using the criteria identified in the **Create** action. The results of the replace action gives the perception that the extension overwrote all settings associated with the local printer.

Update

The **action** attribute value equaling an upper case U instructs the Group Policy Printers extension to update an existing local printer with the settings configured in the current **LocalPrinter** inner element.

The update action instructs the extension to update settings on the local printer only when the computer or current user has a local printer that shares the same name with the value in the **name** attribute.

If the computer or current user does not have an existing local printer name that matches the value in the **name** attribute then the extension creates a new local printer using the settings configured in the **LocalPrinter** inner element.

Comment

The **comment** attribute is an optional string attribute in the **Properties** element used to include additional remarks about the printer that are visible in the printer's property page.

Default

The **default** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **LocalPrinter** inner element with a **default** attribute equaling true indicates the extension should make the printer targeted in the current inner element the default printer after it performs the configured action.

A **Properties** element within a **LocalPrinter** inner element with a **default** attribute equaling false indicates the extension should not make the printer in the current inner element the default printer.

Location

The **location** attribute is an optional string attribute in the **Properties** element used to include a description of the printer's location that appears on the printer's property page.

Name

The **name** attribute is a required attribute in the **Properties** element that accepts a string value. The extension uses the value of the **name** attribute as the name of newly created printer when the **action** attribute value equals an upper case C, R, or U and a local printer does not exist for the computer or local user.

A computer or current user with a local printer sharing the same name as the value of the **name** attribute causes the extension to use that printer as the target of the configured action when the **action** attribute value equals an upper case D, R, or U.

Path

The **path** attribute is a required attribute in the **Properties** element that accepts a string value Universal Naming Convention (UNC) path to a shared printer. The Group Policy Printers extension uses this path as the installation source for the printer driver used to configure the local (physically attached) printer.

Port

The **port** attribute is a required attribute in the **Properties** element that accepts a string value of the local computer port the extension uses when creating the new local printer.

```
<LocalPrinter clsid="{F08996D5-568B-45f5-BB7A-D3FB1E370B0A}" name="printer1"
status="Location" image="2" changed="2013-05-16 17:33:56" uid="{A10F181B-8938-44F6-9818-
84A28B3A37E6}" bypassErrors="1"><Properties action="U" name="printer1" port="LPT4:"
path="\\server\somePrinter" default="0" deleteAll="0" location="Location"
comment="Comments"/></LocalPrinter>
```

XML to User Interface mapping

The Group Policy Printers management snap-in determines the values and attributes stored in the XML configuration. The following figures correlate which components of the user interface control or show various elements and attributes in the XML file.

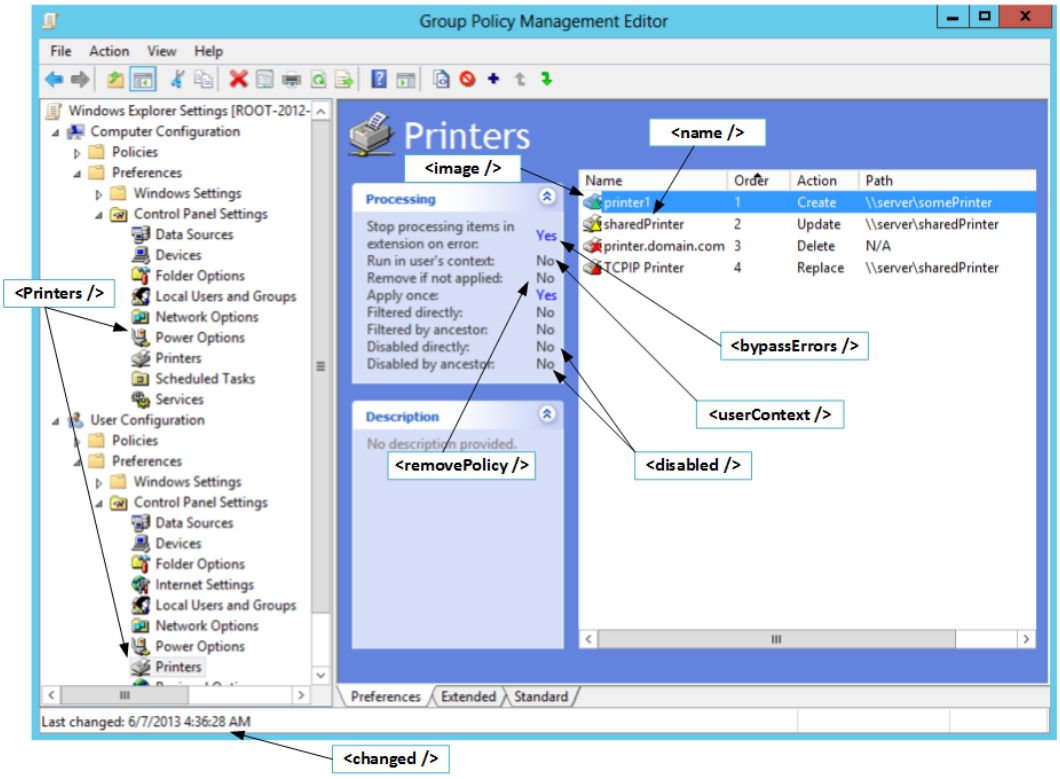
Group Policy Management Editor
File Action View Help
Windows Explorer Settings [ROOT-2012-
Computer Configuration
Policies
Preferences
Windows Settings
Control Panel Settings
Data Sources
Devices
Folder Options
Local Users and Groups
Network Options
Power Options
Printers
Scheduled Tasks
Services
User Configuration
Policies
Preferences
Windows Settings
Control Panel Settings
Data Sources
Devices
Folder Options
Internet Settings
Local Users and Groups
Network Options
Power Options
Printers
Printers
<image />
<name />
Processing
Stop processing items in extension on error: Yes
Run in user's context: No
Remove if not applied: No
Apply once: Yes
Filtered directly: No
Filtered by ancestor: No
Disabled directly: No
Disabled by ancestor: No
Description
No description provided.
Name Order Action Path
printer1 1 Create \\server\somePrinter
sharedPrinter 2 Update \\server\sharedPrinter
printer.domain.com 3 Delete N/A
TCPIP Printer 4 Replace \\server\sharedPrinter
Preferences Extended Standard
Last changed: 6/7/2013 4:36:28 AM
<Printers />
<bypassErrors />
<userContext />
<removePolicy />
<disabled />
<changed />

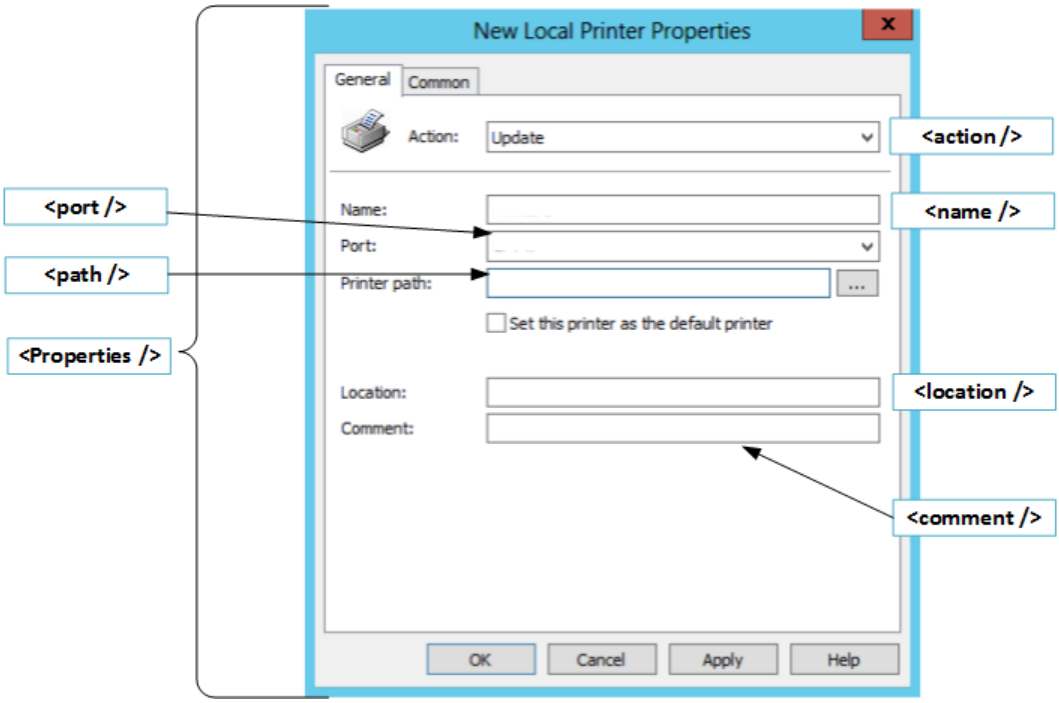
New Local Printer Properties
General
Common
Action:
Update
Name:
Port:
Printer path:
...
Set this printer as the default printer
Location:
Comment:
OK
Cancel
Apply
Help
<action />
<name />
<port />
<path />
<Properties />
<location />
<comment />

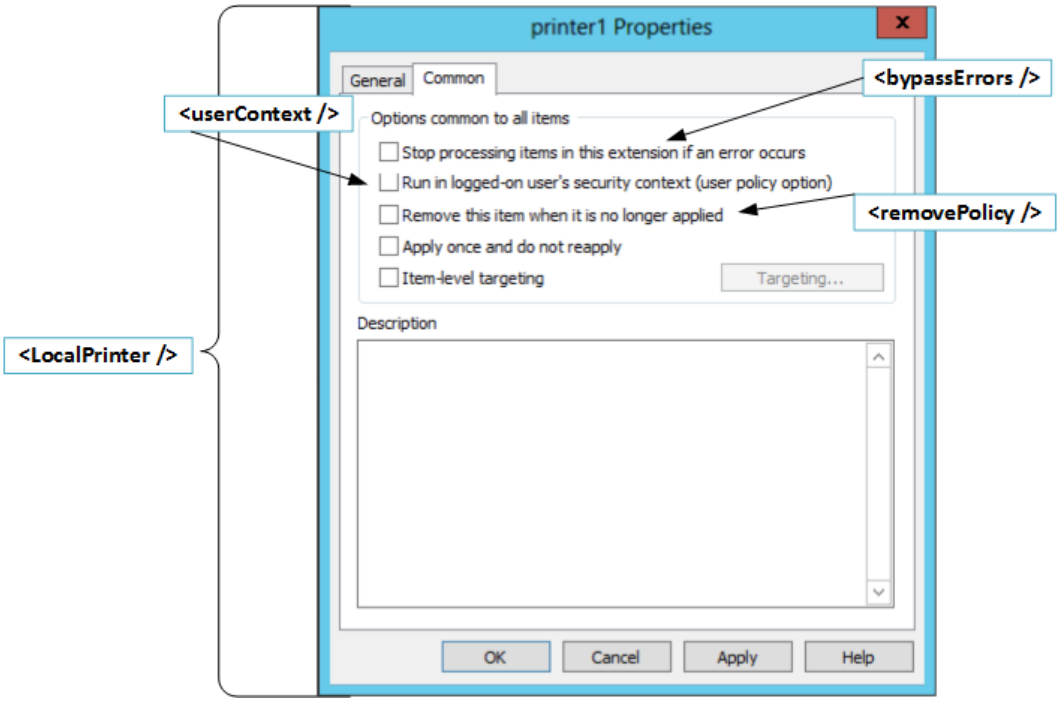

Port Printer Element

Attributes

ByPassErrors

The **byPassErrors** attribute is an optional attribute for the **PortPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

One or more printer inner elements are contained within a single **Printers** outer element, which enables the client-side extension to process one or more preference items. A **PortPrinter** inner element with a **byPassErrors** attribute equaling true indicates the Group Policy Printers extension should continue processing the next printer inner element regardless if the result of the current printer inner element is a failure. A **byPassErrors** attribute equaling false indicates the extension should stop processing subsequent printer inner elements if the result of the current inner element is a failure.

A **LocalPrinter** inner element without a **byPassErrors** attribute indicates that the inner element has never been configured to halt on errors. In this configuration, the default value for the **byPassErrors** attribute is true, which means the any inner element failure does not halt the processing of subsequent inner elements. The **byPassErrors** attribute appears in the configuration the first time the inner element is configured to stop on errors. Once in this configuration, the value of the attribute changes with each respective configuration, but never removed.

Changed

The **changed** attribute is an optional attribute for the **PortPrinter** inner element. The attribute uses a string value to represent the date and time of when the inner element was last modified in UTC. The string value uses the format of YYYY-MM-DD HH:MM:SS.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of inner element. The Group Policy Printers client-side extension expects the **PortPrinter** inner element's **clsid** value to equal {C3A739D2-4A44-401e-9F9D-88E5E77DFB3E}. This is how the client-side extension identifies and validates the inner element is of type **PortPrinter**.

Disabled

The **disabled** attribute is an optional attribute for the **PortPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A **PortPrinter** inner element with a **disabled** attribute equaling true indicates the configuration for that inner element is disabled. This prevents the Group Policy Printers extension from processing that inner element in the configuration file. A **disabled** attribute equaling false indicates the configuration for that inner element is enabled and that processing continues as normal.

A PortPrinter inner element without a **disabled** attribute indicates that the inner element has never been disabled. In this configuration, the default value for the **disabled** attribute is false, which means the inner element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the inner element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

Image

The **image** attribute is a numeric string value that represents the index of a bitmap resource used by the Group Policy Management editor as the icon image used to display the preference item in the editor. The value of the numeric string typically corresponds to the value of the **action** attribute in the **properties** element.

Image Value	Action Value
0	C
1	R
2	U
3	D

Name

The **name** attribute is a string value that represents the display name of the preference item. The **name** attribute has no impact on the processing of the preference item and is strictly used for management and reporting.

RemovePolicy

The **removePolicy** attribute is an optional attribute for the **PortPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **PortPrinter** inner element with a **removePolicy** attribute equaling true indicates the Group Policy Printers extension should remove the configuration included in the inner element if the Group Policy object hosting the preference item is not within scope of the user or computer. Once the Group Policy object hosting the preference item no longer applies to the user, the Group Policy Printers extension removes the effective results of the inner element's configuration from the user or computer.

The **removePolicy** attribute correlates to the value of the **action** attribute in the **Properties** element. A **PortPrinter** inner element with a **removePolicy** attribute equal to true requires the value of the **action** attribute of **Properties** element to equal an uppercase R.

A **PortPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item as configured.

A **PortPrinter** inner element without a **removePolicy** attribute indicates the inner element has never been configured to remove when it no longer applies. In this configuration, the default value for the **removePolicy** attribute is false, which means the extension processes the preference item as configured. The **removePolicy** attribute appears in the configuration the first time the inner element is configured to remove items when it no longer applies. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

Status

The **status** attribute is string value that represents a message that the Group Policy Preference editor can display in the status area of the Microsoft Management Console. The **status** attribute remains as part of the configuration; however, is no longer being actively implemented.

Uid

The **uid** attribute is a 128-bit identifier represented as global unique identifier (GUID). The **uid** value uniquely identifies each inner element. This is how the client-side extension identifies and differentiates each **PortPrinter** inner element.

UserContext

The **userContext** attribute is an optional attribute for the **PortPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **PortPrinter** inner element with a **userContext** attribute equaling true indicates the Group Policy Printers extension should create the mapped drive using the security context of the current user. A **LocalPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item using the security context of the local system.

A **LocalPrinter** inner element without a **userContext** attribute indicates the inner element has been manually configured to apply in the security context of the current user. In this configuration, the default value varies on a per-policy basis using the simple guidelines of

computer preference items run in the security context of the local system and user preference items run in the security context of the current user.

The **userContext** attribute appears in the configuration the first time the inner element is manually configured to apply using the security context of the current user. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

The **userContext** attribute remains as part of the configuration; however, is no longer actively implemented. Most of the Group Policy Preference extensions programmatically switch between the current user and system security contexts as needed, irrespective of the **userContext** attribute value present in the inner element configuration. The recommended configuration is to leave each preference item to the defaults and to not manually configure the security context of any preference item.

```
<PortPrinter clsid="{C3A739D2-4A44-401e-9F9D-88E5E77DFB3E}" name="TCP/IP Printer"
status="TCP/IP Printer" image="2" changed="2013-06-05 22:11:27" uid="{3840AAFA-BFFE-46DD-
8FDE-8DFE33F34908}">
```

Properties Element

Each inner element requires a Properties element in the configuration file. Group Policy Preference extensions read the Properties element for the preference item configuration.

Attributes

The **Properties** element shares unique and common attributes among all the Group Policy Preferences inner element classes. The common attributes characterize abstract functionality provided by the inner element classes (preference items) such as create, delete, or update. However, the implementation of these actions varies across each inner element.

Action

The **action** attribute is an optional string attribute in the **Properties** element used to describe the action the Group Policy Printers extension performs for the associated inner element. That **action** attribute value is a single, uppercase letter of C, R, U, or D. If the action attribute is not present in the **Properties** element then the default value is U.

Create

An action attribute value equaling an upper case C instructs the extension to create a new TCP/IP printer connection using the settings in the current **PortPrinter** inner element.

The create action instructs the extension to create a new TCP/IP printer connection only when the computer or current user does not have an existing TCP/IP printer connection. The extension determines the matching criterion by the **localName** attribute value.

A string value in the **localName** attribute causes the extension to use the **localName** attribute value as the matching criterion to determine if the printer connection exists. An empty **localName** attribute value causes the extension to use the value in the **ipAddress** attribute as the matching criterion to determine if the printer connection exists. An **ipAddress** attribute containing the fully qualified domain name of the remote printer causes the extension to resolve

the name to an IPv4 address and use the resolved IPv4 address as the value to determine if the printer connection exists.

An existing printer connection found for the computer or current user causes the extension to ignore the action and does not report an error. A computer or current user without an existing printer connection causes the extension to create a new TCP/IP printer connection using the settings from the current **PortPrinter** inner element.

The create action, after successfully created the new TCP/IP printer connection, may also set the newly created TCP/IP printer connection as the default printer based on the value configured in the **default** attribute.

Delete

An action attribute value equaling an upper case D instructs the Group Policy Printers extension to remove an existing TCP/IP printer connection using the settings in the current **PortPrinter** element.

The delete action has two behaviors based on the value in the **deleteAll** attribute. When the **deleteAll** attribute value equals true (one [1]), the delete action removes all TCP/IP printer connections for the computer or current user. A **deleteAll** attribute value equaling false (zero [0]) instructs the extension to remove a TCP/IP printer connection only when the computer or current user has a TCP/IP printer connection matching the settings in the current **PortPrinter** inner element.

The extension determines the matching criterion by the **localName** attribute value. A string value in the **localName** attribute causes the extension to use the **localName** attribute value as the matching criterion to determine if the printer connection exists. An empty **localName** attribute value causes the extension to use the value in the **ipAddress** attribute as the matching criterion to determine if the printer connection exists. An **ipAddress** attribute containing the fully qualified domain name of the remote printer causes the extension to resolve the name to an IPv4 address and use the resolved IPv4 address as the value to determine if the printer connection exists.

Replace

An **action** attribute value equaling an upper case R instructs the extension to replace the existing TCP/IP printer connection with the settings in the current **PortPrinter** inner element. The replace action instructs the extension to perform the delete action first, and then the create action. The delete instruction removes an existing TCP/IP printer connection using the criteria identified in the **Delete** action. The create instruction creates a new TCP/IP printer connection using the criteria identified in the **Create** action. The results of the replace action gives the perception that the extension overwrote all settings associated with the local printer.

Update

The action attribute value equaling an upper case U instructs the Group Policy Printers extension to update an existing TCP/IP printer connection with the settings configured in the current **PortPrinter** inner element.

The update action instructs the extension to update settings on the TCP/IP printer connection only when the computer or current user has a matching TCP/IP printer.

The extension determines the matching criterion by the **localName** attribute value. A string value in the **localName** attribute causes the extension to use the **localName** attribute value as the matching criterion to determine if the printer connection exists. An empty **localName** attribute value causes the extension to use the value in the **ipAddress** attribute as the matching criterion to determine if the printer connection exists. An **ipAddress** attribute containing the fully qualified domain name of the remote printer causes the extension to resolve the name to an IPv4 address and use the resolved IPv4 address as the value to determine if the printer connection exists.

If the computer or current user does not have a matching TCP/IP printer connection then the extension creates a TCP/IP printer connection using the settings configured in the **PortPrinter** inner element.

Comment

The **comment** attribute is an optional string attribute in the **Properties** element used to include additional remarks about the printer that are visible in the printer's property page.

Default

The **default** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **PortIPrinter** inner element with a **default** attribute equaling true indicates the extension makes the printer targeted in the current inner element the default printer after it performs the configured action.

A **Properties** element within a **PortPrinter** inner element with a **default** attribute equaling false indicates the extension does not make the printer in the current inner element the default printer.

DeleteAll

The deleteAll attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **PortPrinter** inner element with a **deleteAll** attribute equaling true indicates the extension deletes all TCP/IP printer connections for the computer or current user when the action attribute value equals an upper case D.

A **Properties** element within a **PortPrinter** inner element with a **deleteAll** attribute equaling false indicates the extension does not delete all TCP/IP printer connections for the computer or current user when the **action** attribute value equals an upper case D.

DoubleSpool

The **doubleSpool** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **PortPrinter** inner element with a **doubleSpool** attribute equaling true indicates the extension enables double spooling for the targeted TCP/IP printer connection for the computer or current user.

A **Properties** element within a **PortPrinter** inner element with a **doubleSpool** attribute equaling false indicates the extension disabled or does not enabled double spooling for the targeted TCP/IP printer connections for the computer or current user.

The **doubleSpool** attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of 0.

IPAddress

The **ipAddress** attribute is a required attribute in the **Properties** element that accepts a string value. The extension uses the value of the **ipAddress** to locate the remote printer it uses to process the configured action.

Also, the extension uses the **ipAddress** attribute value as the criterion used to determine if an existing TCP/IP printer connection exists when the **localName** attribute is not configured. An **ipAddress** attribute containing the fully qualified domain name of the remote printer causes the extension to resolve the name to an IPv4 address and use the resolved IPv4 address as the value to determine if the printer connection exists.

The extension expects the value of the **ipAddress** to be a numerical string of an IPv4 address of the remote printer when the **useDNS** attribute value equals false (zero [0]). When the **useDNS** attribute value equals true (one [1]), the extension expects the value of the **ipAddress** to be fully qualified domain name of the remote printer.

UseDNS

The **useDNS** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **PortPrinter** inner element with a **useDNS** attribute equaling true indicates the extension expects the **ipAddress** attribute value to be a fully qualified domain name of the remote printer. Additionally, the extension resolves the fully qualified domain name to an IPv4 address before performing the configured action.

A **Properties** element within a **PortPrinter** inner element with a **useDNS** attribute equaling false indicates the extension expects the **ipAddress** attribute value to be numerical string IPv4 address.

LocalName

The **localName** attribute is an optional string attribute. The extension uses this attribute as a name for newly created TCP/IP printer connections and as the matching criterion that determines if a TCP/IP printer connection exists for the computer or current user.

An unconfigured **localName** attribute instructs the extension to use the configured IPv4 address as the matching criterion that determines if a TCP/IP printer connection exists. Otherwise, the extension uses the **localName** attribute value as the matching criterion to determine if a TCP/IP printer connection exists. A TCP/IP printer connection sharing the same name as the **localName** attribute value becomes the targeted TCP/IP printer connection for the configured action.

Location

The **location** attribute is an optional string attribute in the **Properties** element used to include a description of the printer's location that appears on the printer's property page.

LPRQueue

The **lprQueue** attribute is an optional string attribute in the **Properties** element. A **protocol** attribute value equaling "PROTOCOL_LPR_TYPE" requires the **lprQueue** attribute to contain a string value that represents the name of the LPR queue. The extension ignored the **lprQueue** attribute when the **protocol** attribute value equals "PROTOCOL_RAWTCP_TYPE"

Path

The **path** attribute is a required attribute in the **Properties** element that accepts a string value Universal Naming Convention (UNC) path to a shared printer. The Group Policy Printers extension uses this path as the installation source for the printer driver used to configure the TCP/IP printer connection.

PortNumber

The **portNumber** attribute is an optional, numerical-string attribute in the **Properties** element when configuring a raw TCP/IP printer and a required numerical-string attribute when configuring a LPR printer. The attribute's value must be the port number used to communicate with the targeted printer. A **protocol** attribute value equaling PROTOCOL_LPR_TYPE requires the **portNumber** to contain the numerical-string value 515.

The **portNumber** attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of 9100.

Protocol

The **protocol** attribute is an optional, string attribute in the **Properties** element when configuring raw TCP/IP printer and a required string attribute when configuring a LPR printer. A **protocol** attribute value equaling "PROTOCOL_LPR_TYPE" indicates the targeted printer is configured as an LPR printer. A **protocol** attribute value equaling "PROTOCOL_RAWTCP_TYPE" or a missing **protocol** attribute indicates the targeted printer is configured as a raw TCP/IP printer.

The protocol attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of "PROTOCOL_RAWTCP_TYPE".

SkipLocal

The **skipLocal** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false. The extension ignores attribute if the **default** attribute value equals false.

A **Properties** element within a **PortPrinter** inner element with a **skipLocal** attribute equaling true indicates the extension must bypass changing the default printer if a local printer is configured for the computer or current user.

A **Properties** element within a **PortPrinter** inner element with a **skipLocal** attribute equaling false indicates the extension does not bypass changing the default printer if a local printer is configured for the computer or current user.

SNMPCommunity

The **snmpCommunity** attribute is an optional string attribute in the **Properties** element. When configured, the **snmpCommunity** value represents the community named in which the printer participates. The attribute's value defaults to "public" and applies to both raw and LPR TCP/IP printers.

The **snmpCommunity** attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of "public."

SNMPDevIndex

The **snmpDevIndex** is an optional attribute in the **Properties** element. The attribute uses a numerical string value to indicate the SNMP device index used for the targeted printer.

The **snmpDevIndex** attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of 0.

SNMPEnabled

The **snmpEnabled** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **PortPrinter** inner element with a **snmpEnabled** attribute equaling true indicates the extension enables SNMP and its settings on the printer targeted in the current inner element.

A **Properties** element within a **PortPrinter** inner element with a **snmpEnabled** attribute equaling false indicates the extension does not enable SNMP and its settings on the printer targeted in the current inner element.

The **snmpEnabled** attribute only appears in the inner element when you configure one or more port settings for the targeted printer. Otherwise, the extension uses the default value of 0.

```
<PortPrinter clsid="{C3A739D2-4A44-401e-9F9D-88E5E77DFB3E}" name="dnsName"
status="dnsName" image="2" changed="2013-06-05 22:11:27" uid="{3840AAFA-BFFE-46DD-8FDE-
8DFE33F34908}"><Properties ipAddress="dnsName" action="U" location="Location"
localName="Local Name" comment="Comment" default="1" skipLocal="1" useDNS="1"
path="printer path" deleteAll="0"/></PortPrinter>
```

XML to User Interface mapping

The Group Policy Printers management snap-in determines the values and attributes stored in the XML configuration. The following figures correlate which components of the user interface control or show various elements and attributes in the XML file.

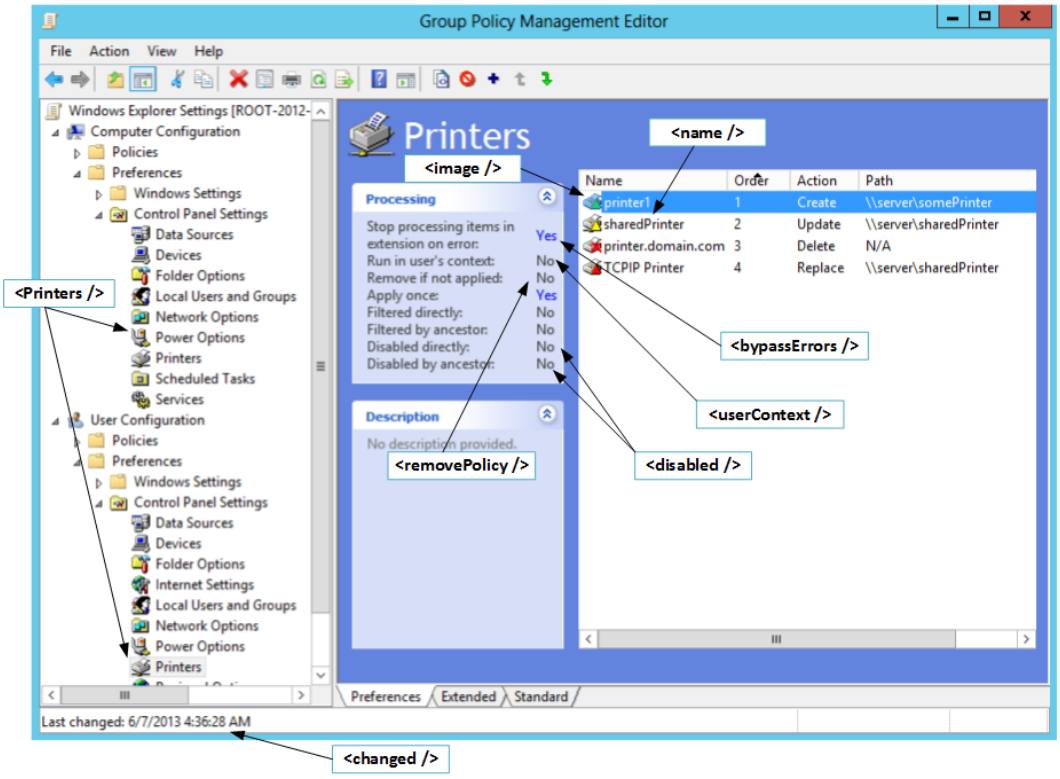
Group Policy Management Editor
File Action View Help
Windows Explorer Settings [ROOT-2012-
Computer Configuration
Policies
Preferences
Windows Settings
Control Panel Settings
Data Sources
Devices
Folder Options
Local Users and Groups
Network Options
Power Options
Printers
Scheduled Tasks
Services
User Configuration
Policies
Preferences
Windows Settings
Control Panel Settings
Data Sources
Devices
Folder Options
Internet Settings
Local Users and Groups
Network Options
Power Options
Printers
Printers
<image />
<name />
Processing
Stop processing items in extension on error: Yes
Run in user's context: No
Remove if not applied: No
Apply once: Yes
Filtered directly: No
Filtered by ancestor: No
Disabled directly: No
Disabled by ancestor: No
Description
No description provided.
Name Order Action Path
printer1 1 Create \\server\somePrinter
sharedPrinter 2 Update \\server\sharedPrinter
printer.domain.com 3 Delete N/A
TCPIP Printer 4 Replace \\server\sharedPrinter
<Printers />
<bypassErrors />
<userContext />
<removePolicy />
<disabled />
Preferences Extended Standard
Last changed: 6/7/2013 4:36:28 AM
<changed />

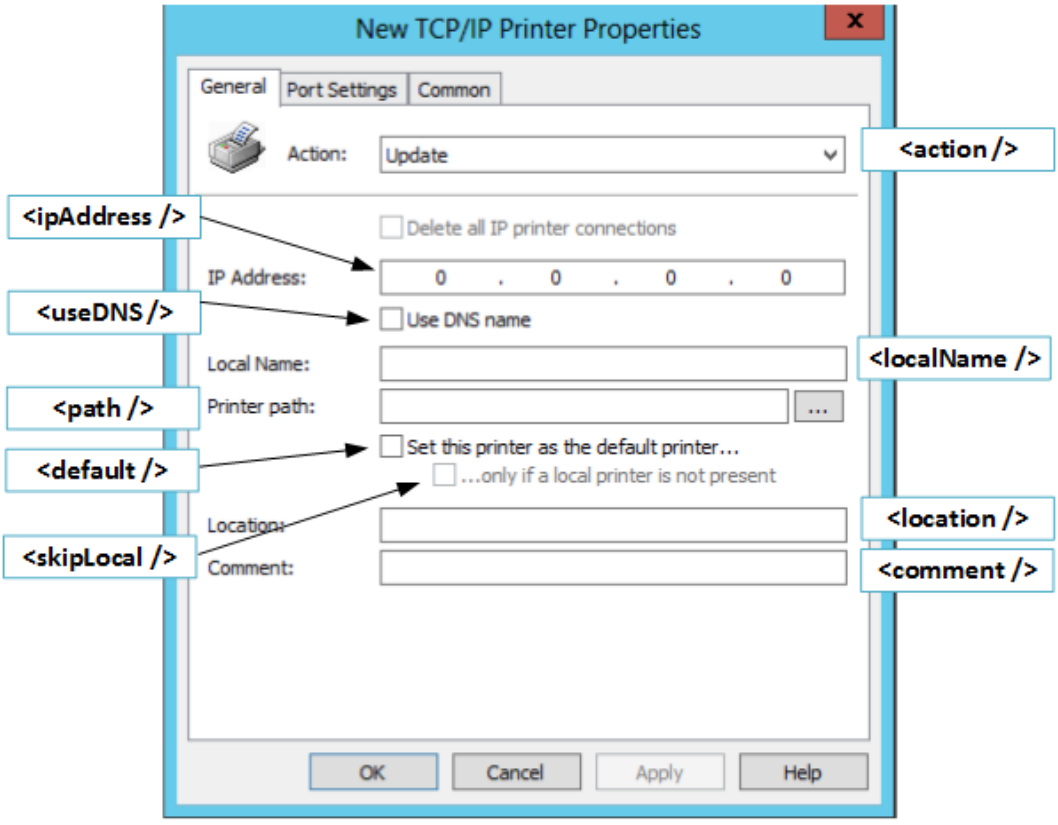
New TCP/IP Printer Properties
General Port Settings Common
Action: Update
<action />
<ipAddress />
Delete all IP printer connections
IP Address: 0 . 0 . 0 . 0
<useDNS />
Use DNS name
Local Name:
<localName />
<path />
Printer path:
...
<default />
Set this printer as the default printer...
...only if a local printer is not present
Location:
<location />
<skipLocal />
Comment:
<comment />
OK Cancel Apply Help

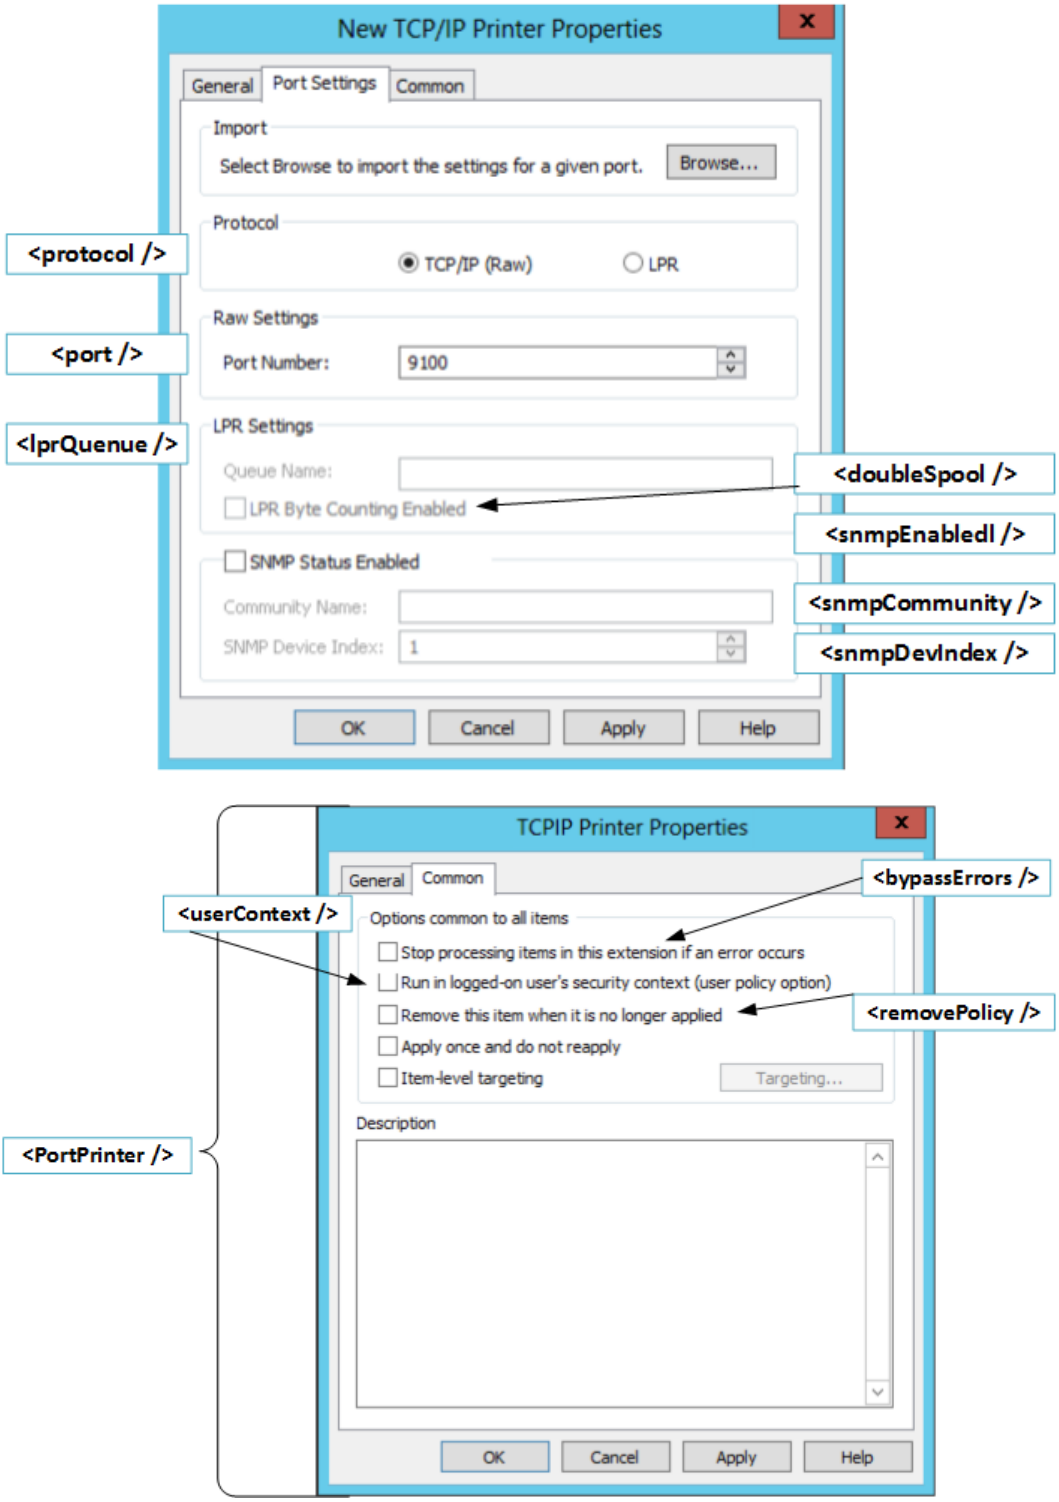
New TCP/IP Printer Properties
General
Port Settings
Common
Import
Select Browse to import the settings for a given port.
Browse...
Protocol
TCP/IP (Raw)
LPR
Raw Settings
Port Number:
9100
LPR Settings
Queue Name:
LPR Byte Counting Enabled
SNMP Status Enabled
Community Name:
SNMP Device Index:
1
OK
Cancel
Apply
Help
<protocol />
<port />
<lprQuenue />
<doubleSpool />
<snmpEnabledl />
<snmpCommunity />
<snmpDevIndex />
TCPIP Printer Properties
General
Common
Options common to all items
Stop processing items in this extension if an error occurs
Run in logged-on user's security context (user policy option)
Remove this item when it is no longer applied
Apply once and do not reapply
Item-level targeting
Targeting...
Description
OK
Cancel
Apply
Help
<userContext />
<bypassErrors />
<removePolicy />
<PortPrinter />

Shared Printer Element

Attributes

ByPassErrors

The **byPassErrors** attribute is an optional attribute for the **SharedPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

One or more printer inner elements are contained within a single **Printers** outer element, which enables the client-side extension to process one or more preference items. A **SharedPrinter** inner element with a **byPassErrors** attribute equaling true indicates the Group Policy Printers extension should continue processing the next printer inner element regardless if the result of the current printer inner element is a failure. A **byPassErrors** attribute equaling false indicates the extension should stop processing subsequent printer inner elements if the result of the current inner element is a failure.

A **SharedPrinter** inner element without a **byPassErrors** attribute indicates that the inner element has never been configured to halt on errors. In this configuration, the default value for the **byPassErrors** attribute is true, which means the any inner element failure does not halt the processing of subsequent inner elements. The **byPassErrors** attribute appears in the configuration the first time the inner element is configured to stop on errors. Once in this configuration, the value of the attribute changes with each respective configuration, but never removed.

Changed

The **changed** attribute is an optional attribute for the **SharedPrinter** inner element. The attribute uses a string value to represent the date and time of when the inner element was last modified in UTC. The string value uses the format of YYYY-MM-DD HH:MM:SS.

Clsid

The **clsid** attribute is a 128-bit unique identifier represented as global unique identifier (GUID). The **clsid** value identifies the type of inner element. The Group Policy Printers client-side extension expects the **SharedPrinter** inner element's **clsid** value to equal {9A5E9697-9095-436d-A0EE-4D128FDFBCE5}. This is how the client-side extension identifies and validates the inner element is of type **SharedPrinter**.

Disabled

The **disabled** attribute is an optional attribute for the **SharedPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicated false.

A **SharedPrinter** inner element with a **disabled** attribute equaling true indicates the configuration for that inner element is disabled. This prevents the Group Policy Printers extension from processing that inner element in the configuration file. A **disabled** attribute equaling false indicates the configuration for that inner element is enabled and that processing continues as normal.

A **SharedPrinter** inner element without a **disabled** attribute indicates that the inner element has never been disabled. In this configuration, the default value for the **disabled** attribute is false,

which means the inner element is enabled and processing continues normally. The **disabled** attribute appears in the configuration the first time the inner element is disabled and remains in the configuration for the remainder configuration files lifetime. Once in the configuration, the value of the attribute is changed with the respective configuration, but never removed.

Image

The **image** attribute is a numeric string value that represents the index of a bitmap resource used by the Group Policy Management editor as the icon image used to display the preference item in the editor. The value of the numeric string typically corresponds to the value of the **action** attribute in the **properties** element.

Image Value	Action Value
0	C
1	R
2	U
3	D

Name

The **name** attribute is a string value that represents the display name of the preference item. The **name** attribute has no impact on the processing of the preference item and is strictly used for management and reporting.

RemovePolicy

The **removePolicy** attribute is an optional attribute for the **SharedPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **SharedPrinter** inner element with a **removePolicy** attribute equaling true indicates the Group Policy Printers extension should remove the configuration included in the inner element if the Group Policy object hosting the preference item is not within scope of the user or computer. Once the Group Policy object hosting the preference item no longer applies to the user, the Group Policy Printers extension removes the effective results of the inner element's configuration from the user or computer.

The **removePolicy** attribute correlates to the value of the **action** attribute in the **Properties** element. A **SharedPrinter** inner element with a **removePolicy** attribute equal to true requires the value of the **action** attribute of **Properties** element to equal an uppercase R.

A **SharedPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item as configured.

A **SharedPrinter** inner element without a **removePolicy** attribute indicates the inner element has never been configured to remove when it no longer applies. In this configuration, the default value for the **removePolicy** attribute is false, which means the extension processes the preference item as configured. The **removePolicy** attribute appears in the configuration the first

time the inner element is configured to remove items when it no longer applies. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

Status

The **status** attribute is string value that represents a message that the Group Policy Preference editor can display in the status area of the Microsoft Management Console. The **status** attribute remains as part of the configuration; however, is no longer being actively implemented.

Uid

The **uid** attribute is a 128-bit identifier represented as global unique identifier (GUID). The **uid** value uniquely identifies each inner element. This is how the client-side extension identifies and differentiates each **SharedPrinter** inner element.

UserContext

The **userContext** attribute is an optional attribute for the **SharedPrinter** inner element. The attribute uses a numeric string to represent a Boolean value. A value of 1 indicates true and a value of 0 indicates false.

A **SharedPrinter** inner element with a **userContext** attribute equaling true indicates the Group Policy Printers extension should create the mapped drive using the security context of the current user. A **SharedPrinter** inner element with a **removePolicy** attribute equaling false indicates the extension should process the preference item using the security context of the local system.

A **SharedPrinter** inner element without a **userContext** attribute indicates the inner element has been manually configured to apply in the security context of the current user. In this configuration, the default value varies on a per-policy basis using the simple guidelines of computer preference items run in the security context of the local system and user preference items run in the security context of the current user.

The **userContext** attribute appears in the configuration the first time the inner element is manually configured to apply using the security context of the current user. Once in this configuration, the value of the attribute changes with each respective configuration, but is never removed.

The **userContext** attribute remains as part of the configuration; however, is no longer actively implemented. Most of the Group Policy Preference extensions programmatically switch between the current user and system security contexts as needed, irrespective of the **userContext** attribute value present in the inner element configuration. The recommended configuration is to leave each preference item to the defaults and to not manually configure the security context of any preference item.

```
<SharedPrinter clsid="{9A5E9697-9095-436d-A0EE-4D128FDFBCE5}" name="sharePath"
status="sharePath" image="2" changed="2013-06-05 22:06:43" uid="{7C851C09-C0D2-4EFC-8FD6-
6A5093EE644E}">
```

Properties Element

Action

The Properties element shares unique and common attributes among all the Group Policy Preferences inner element classes. The common attributes characterize abstract functionality provided by the inner element classes (preference items) such as create, delete, or update. However, the implementation of these actions varies across each inner element.

Create

An action attribute value equaling an upper case C instructs the extension to create a new shared printer connection using the settings in the current **SharedPrinter** inner element.

The create action instructs the extension to create a new shared printer connection only when the computer or current user does not have a shared printer connection that shares the same path with the value in the **path** attribute.

If the computer or current user has a shard printer connection with a share path that matches the value in the **path** attribute then the extension ignores the instruction and does not report an error.

The create action, after successfully creating the new shared printer connection, may also set the newly created shared printer connection as the default printer based on the value configured in the **default** attribute.

Delete

An action attribute value equaling an upper case D instructs the Group Policy Printers extension to remove an existing shared printer connection using the settings in the current **SharedPrinter** inner element.

The delete action has two behaviors based on the value in the **deleteAll** attribute and two behaviors based on the value in the **deleteMaps** attribute.

When the **deleteAll** attribute value equals true (one [1]), the delete action removes all shared printer connections from the computer or current user. A **deleteAll** attribute equaling false (zero [0]) instructs the extension to remove the shared printer connection only when the computer or current user has a shared printer connection with a share path that matches the value in the **path** attribute of the current inner element.

A computer or current user not having an existing shared printer connection with a share path matching the value in the **path** attribute causes the extension to ignore the instruction and not report an error.

When the **deleteMaps** attribute value equals true (one [1]), the delete action removes all shared printer connections from associated local ports for the computer or current user. A **deleteMaps** attribute value equaling false (zero [0]) case the extension to remove the shared printer connection associated with the port matching the value in the **port** attribute.

Replace

An action attribute value equaling an upper case R instructs the extension to replace the existing shared printer connecting with the settings in the current **SharedPrinter** inner element. The replace action instructs the extension to perform the delete action first, and then the create action. The delete instruction removes an existing shared printer connection using the criteria identified in the **Delete** action. The create instruction creates a new shared printer connection using the criteria identified in the **Create** action. The results of the replace action gives the perception that the extension overwrote all settings associated with the shared printer connection.

Update

The action attribute value equaling an upper case U instructs the Group Policy Printers extension to update an existing local printer with the settings configured in the current **SharedPrinter** inner element.

The update action instructs the extension to update the settings on the shared printer connections only when the computer or current user has a shared printer connection with a share path matching the **path** attribute value in the current inner element.

A computer or current user not having a shared printer connection with a share path matching the value in the **path** attribute cause the extension to create a new shared printer connection using the settings configured in the **SharedPrinter** inner element.

Comment

The **comment** attribute is an optional string attribute in the **Properties** element used to include additional remarks about the printer that are visible in the printer's property page.

Cpassword

The **cpassword** attribute is an optional string attribute in the **Properties** element used to store a password needed to successfully connect the configured shared printer. The attribute value is a string representation of the password that is obfuscated using an AES-derived encryption key. The default value for the **cpassword** attribute is an empty string, or blank.

Default

The **default** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **SharedPrinter** inner element with a default attribute equaling true indicates the extension makes the printer targeted in the current inner element the default printer after it performs the configured action.

A **Properties** element within a **SharedPrinter** inner element with a **default** attribute equaling false indicates the extension does not make the printer in the current inner element the default printer.

DeleteAll

The **deleteAll** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **SharedPrinter** inner element with a **deleteAll** attribute equaling true indicates the extension deletes all shared printer connections for the computer or current user when the **action** attribute value equals an upper case D.

A **Properties** element within a **SharedPrinter** inner element with a **deleteAll** attribute equaling false indicates the extension does not delete all shared printer connections for the computer or current user when the **action** attribute value equals an upper case D.

DeleteMaps

The **deleteMaps** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value equaling 0 indicates false.

A **Properties** element within a **SharedPrinter** inner element with a **deleteMaps** attribute equaling true indicates the extension removes all shared printer connections from associated local ports for the computer or current user when the action attribute value equals an upper case D.

A **Properties** element within a **SharedPrinter** inner element with a **deleteAll** attribute equaling false indicates the extension removes the shared printer connection associated with the port matching the value in the **port** attribute for the computer of current user when the **action** attribute value equals an upper case D.

Location

The **location** attribute is an optional string attribute in the **Properties** element used to include a description of the printer's location that appears on the printer's property page.

Path

The **path** attribute is a required attribute in the **Properties** element that accepts a string value Universal Naming Convention (UNC) path to a shared printer. The Group Policy Printers extension uses this path as the installation source for the printer driver used to configure the shared printer connection.

Persistent

The **persistent** attribute is an optional attribute in the **Properties** element. The attribute uses a numeric string to represent a Boolean value. A value equaling 1 indicates true and a value of 0 indicates false.

A **Properties** element within a **SharedPrinter** inner element with a **persistent** attribute equaling true indicates the Group Policy Printers extension creates the shared printer connection to remain persistent between user logons and computer reboots. Windows remembers the shared printer connection and reestablishes the connection on subsequent user logons.

A **Properties** element within a **SharedPrinter** inner element with a **persistent** attribute equaling false indicates the Group Policy Printers extension creates the shared printer connection not to persist between user logons and computer reboots. Windows does not remember the shared printer connection and does not reestablish the connection on subsequent reboots.

It's important to understand that Windows configures persistency at the time it creates the shared printer connection. Therefore, if you want need to use the Replace action if you need to change a shared printer connection's persistent state.

Port

The **port** attribute is a required attribute in the **Properties** element that accepts a string value of the local computer port the extension uses when creating the new shared printer connection.

SkipLocal

The **skipLocal** attribute is an optional attribute in the **Properties** element. The attribute uses a numerical string to represent a Boolean value. A value equaling 1 indicates true and a value

equaling 0 indicates false. The extension ignores attribute if the default attribute value equals false.

A **Properties** element within a **SharedPrinter** inner element with a **skipLocal** attribute equaling true indicates the extension must bypass changing the default printer if a local printer is configured for the computer or current user.

A **Properties** element within a **SharedPrinter** inner element with a **skipLocal** attribute equaling false indicates the extension does not bypass changing the default printer if a local printer is configured for the computer or current user.

Username

The **username** attribute is an optional string attribute in the **Properties** element used to store the domain and username needed to successfully connect to the configured shared printer.

```
<SharedPrinter clsid="{9A5E9697-9095-436d-A0EE-4D128FDFBCE5}" name="sharePath"
status="sharePath" image="2" changed="2013-06-05 22:06:43" uid="{7C851C09-C0D2-4EFC-8FD6-
6A5093EE644E}"><Properties action="U" comment="" path="sharePath" location="" default="1"
skipLocal="1" deleteAll="0" persistent="1" deleteMaps="0" port="Local
Port"/></SharedPrinter>
```

XML to User Interface mapping

The Group Policy Printers management snap-in determines the values and attributes stored in the XML configuration. The following figures correlate which components of the user interface control or show various elements and attributes in the XML file.

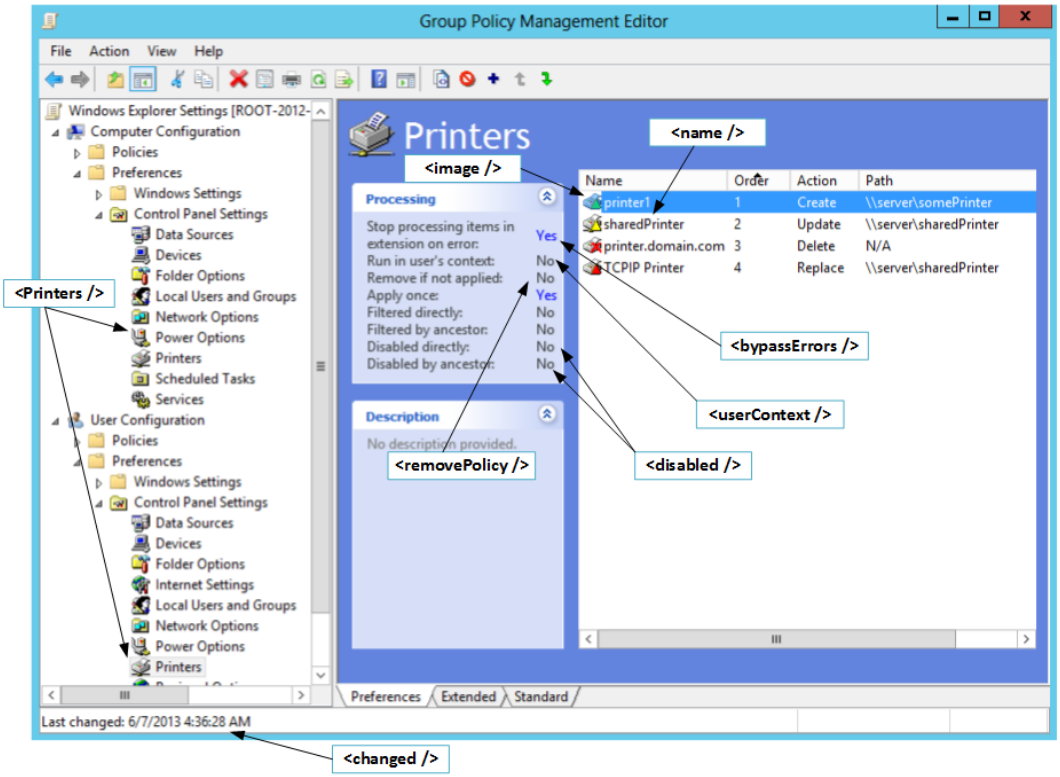

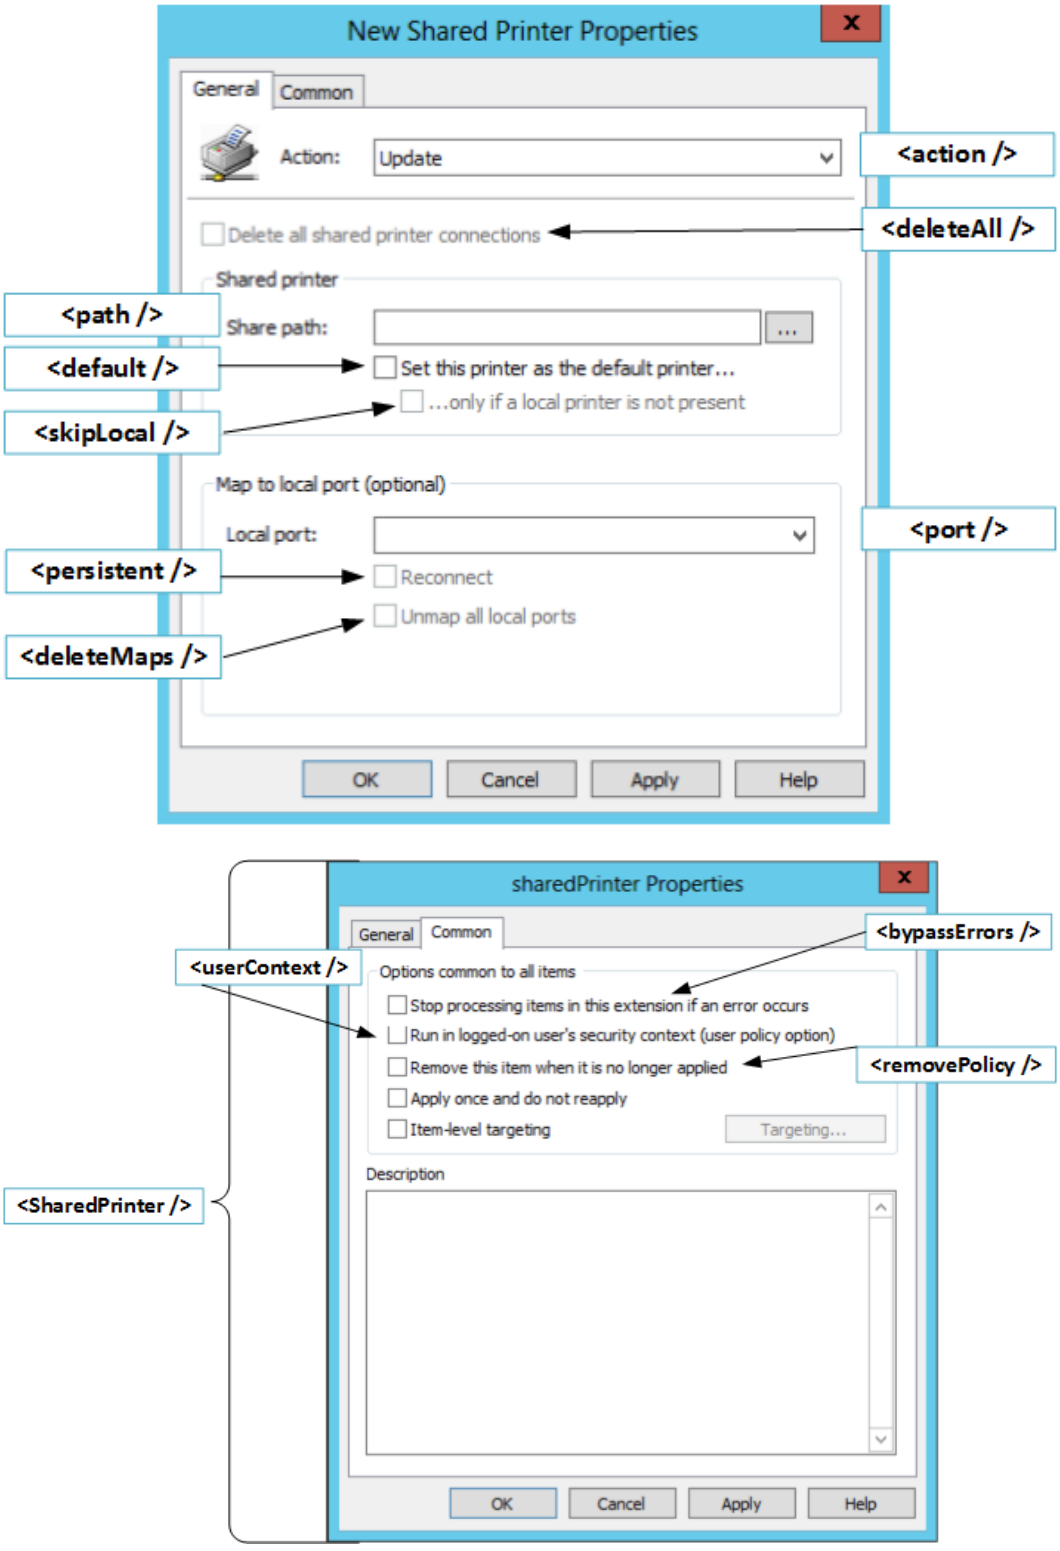
New Shared Printer Properties
General
Common
Action:
Update
<action />
Delete all shared printer connections
<deleteAll />
Shared printer
<path />
Share path:
...
<default />
Set this printer as the default printer...
...only if a local printer is not present
<skipLocal />
Map to local port (optional)
Local port:
<port />
<persistent />
Reconnect
Unmap all local ports
<deleteMaps />
OK
Cancel
Apply
Help
sharedPrinter Properties
General
Common
<bypassErrors />
<userContext />
Options common to all items
Stop processing items in this extension if an error occurs
Run in logged-on user's security context (user policy option)
Remove this item when it is no longer applied
<removePolicy />
Apply once and do not reapply
Item-level targeting
Targeting...
Description
<SharedPrinter />
OK
Cancel
Apply
Help

Evaluation of XML

The extension reads the entire XML file into memory subsequently processes each element in the XML file starting with the outer element. The outer element defines the overarching class of preference items that appear in the inner elements of the file. Additionally, the outer element may contain the **disabled** attribute. A **disabled** attribute value equaling true (one [1]) in the outer element stops any further evaluation of the XML, effectively preventing all **LocalPrinter**, **SharedPrinter**, or **PortPrinter** inner elements from processing (applying to the computer or user).

After reading the **Printers** element, the extension focuses its attention to the next element in the XML file—one of possibly many printer inner element. As noted in the attribute section, printer inner elements contain metadata and processing control information used be the Group Policy management snap-ins and the client-side extensions.

The extension reads the attributes and values from the entire printer inner element, which includes the attributes and values of the **LocalPrinter**, **SharedPrinter**, or **PortPrinter** elements, the attributes and values of the Properties element and the attributes and values of an optional **Filters** element along with all its child elements and their attributes and values.

A printer inner element containing a **disabled** attribute equaling true (one [1]) stops the extension from processing this element (applying to the computer or user) and the extension moves its focus to the next inner element.

The extension continues processing the current inner element when it does not contain a **disabled** attribute or when it contains a **disabled** attribute with a value equaling false (zero [0]).

Next, the extension moves its focus on the **Filter** element of the inner element.

The **Filters** element holds Item-level targeting information that the extension uses as additional criteria to determine if the associated preference items (Printers) should apply to the user or computer.

The Filters element contains attributes, values and child elements that represent a sophisticated set of Boolean rules. The details of the Filters element are beyond the scope of this section. In this context, all that is needed to know is a Filter element, when evaluated by an extension, only returns true or false.

When present, the extension reads the entire **Filters** element and evaluates each Boolean expression included in the section. The evaluation of the **Filters** element can only result in a true or false response. A filter evaluation result equaling false means the one or more Boolean expressions in the **Filters** element evaluated to false. The false evaluation forces the extension to stop any further processing of the current inner element and moves focus to the next inner element.

A **LocalPrinter**, **SharedPrinter**, or **PortPrinter** inner element that does not contain a **Filters** inner element or an inner element containing a **Filters** inner element that evaluates to true moves its focus to the **Properties** element, which is the next element within the current inner element.

The **Properties** element within a **LocalPrinter** inner element contains all the local printer configuration information the extension needs to create, replace, update, or delete a local printer.

The Group Policy Printer extension reads the attributes and values and performs the requested instructions and determines if the actions performed were successful.

If the actions performed were unsuccessful (an error occurred), the extension checks the **bypassErrors** attribute value read from the current **LocalPrinter** inner element. A false value (zero [0]) instructs the extension to discontinue further processing of all printer inner elements while a true value (one [1]) instructs the extension to ignore the unsuccessful result and continue processing additional inner elements.

The extension then moves its focus to the next printer inner element and the process repeats until the XML file reaches the Printers closing outer element.

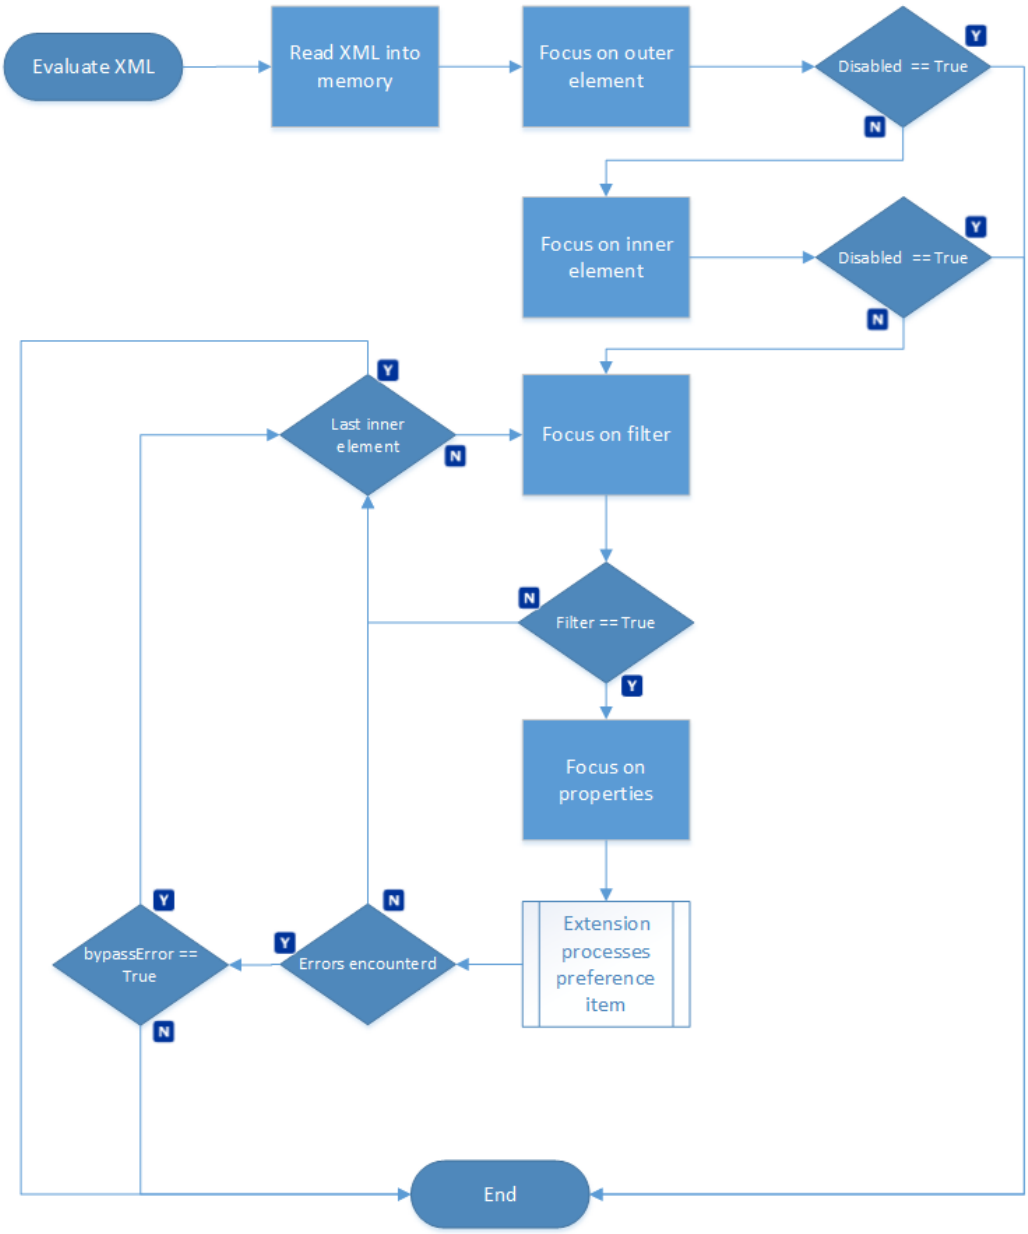

CSE Processing

Group Policy client side extensions enable Group Policy to process a variety of different tasks because each client side extension is responsible for processing its portion of data stored in a Group Policy object.

The Group Policy Printers client side extension is one of many extensions included in the gpprefcl.dll file that is included in the Windows operating system starting with Windows Server 2008 and Windows Vista Service Pack 1.

Group Policy's modular design enables each extension to process the data that is relevant to them. Therefore, the implementation of each Group Policy client side extension may be different.

The Group Policy Printers extension is one CSE in a family of Group Policy client side extensions known as Group Policy Preferences. The goal of this extension is to manage printer deployment scenarios. The extension comes registered on the Windows operating system by virtue of CSE specific metadata stored in the following registry location

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows
NT\CurrentVersion\Winlogon\GPExtensions\{BC75B1ED-5833-4858-9BB8-CBF0B166DF9D}
```

The information stored in this location is consistent on all computers. Therefore, restoring the information on a computer where it is missing is relatively easy using the export and import feature of the Registry editor.

Startup

The Group Policy service invokes this client side extension when it encounters a Group Policy object configured with Group Policy Printers. When activated, the extension must locate its configuration file.

Prerequisite information

When Group Policy invokes the client side extension, it provides to the extension information that enables it to process its portion of Group Policy. This information includes operational information that describes the current processing context such as if the processing is background or foreground, is the processing occurring over a slow link, or is the processing for the computer rather than the user. Also included in this information is a list of the Group Policy objects that were removed from the scope of application and the list of Group Policy objects that were changed.

Each entry in the list of Group Policy objects provides to the client side extension provides information about the Group Policy object in question. The entry includes the following table.

Field	Description
Options	Client side extensions use the options field to determine if the Group Policy object is disabled or if the Group Policy object is

Version	The Group Policy object version number increments with saved changes to the computer and user portions of a GPO, respectively. Client side extensions can use the version number to determine if any changes were made to the computer or user portion of a GPO.
DSPath	The DS path provides the LDAP path, which provides the location and common name of the Group Policy object in Active Directory. Remember, the common name of a GPO is its unique identifier; not its display name.
File system path	The file system path provides the root location of the files associated with the Group Policy object. These collection of files typically are used by one or more client side extensions as the data store for group policy settings.
Link Path	The link path field provides an LDAP path of the container to which the Group Policy is linked
Link Type	The link type field provides information about the container to which the GPO is linked such as computer, site, domain, or organizational unit.
Extensions	The extension field provides a list of the extensions that are configured in the Group Policy object.
Display name	The display name is the name of the GPO that is displayed in user interfaces and reporting tools.

The information in the table above is present for each GPO in both lists. The out-of-scope list contains a list of GPOs that applied to the user or computer on the last application of Group Policy but no longer apply on this application of Group Policy. The changed list contains a list of GPOs that include changes for one or more aspects of the Group Policy object.

The extension reads the operational information to understand what type of Group Policy processing Windows is performing: user or computer. How the processing is occurring: foreground or background. This enables the client side extensions to make operational decisions based on this processing characteristics.

The extension uses the lists of Group Policy objects to complete its portion of Group Policy processing. The Group Policy Printers extension starts by processing each of the GPOs in the out-of-scope list.

Out-of-Scope Group Policy objects

The Group Policy Printers extension processes the entire list of out-of-scope GPOs sequentially. The extension logs any errors reading the Group Policy objects and continues processing until it reaches the end of the list.

The extension cycles through the list of Group Policy objects in the list starting with the first entry in the list. The extension collects the prerequisite information about the current GPO. From this information and its own information, the extension builds file path location to the configuration file. The configuration file location path begins with the extensions history path.

Group Policy Preference History

The history path is a local data store of the last successfully applied data configuration files for each Group Policy object. The extension saves the last applied configuration files in the user's local application folder, which holds a value equivalent to the Group Policy Preference environment variable %LOCALAPPDATA% on Windows 8 and Windows Server 2012. The root of the history folder looks similar to

```
C:\users\[username]\AppData\Local\Microsoft\Group Policy\History
```

Windows protects the history folder with the hidden attribute; therefore, by default, it is not visible from Windows Explorer or from the command line. The history folder contains sub folders, where each folder is named to the globally unique identifier (GUID) of a corresponding Group Policy object.

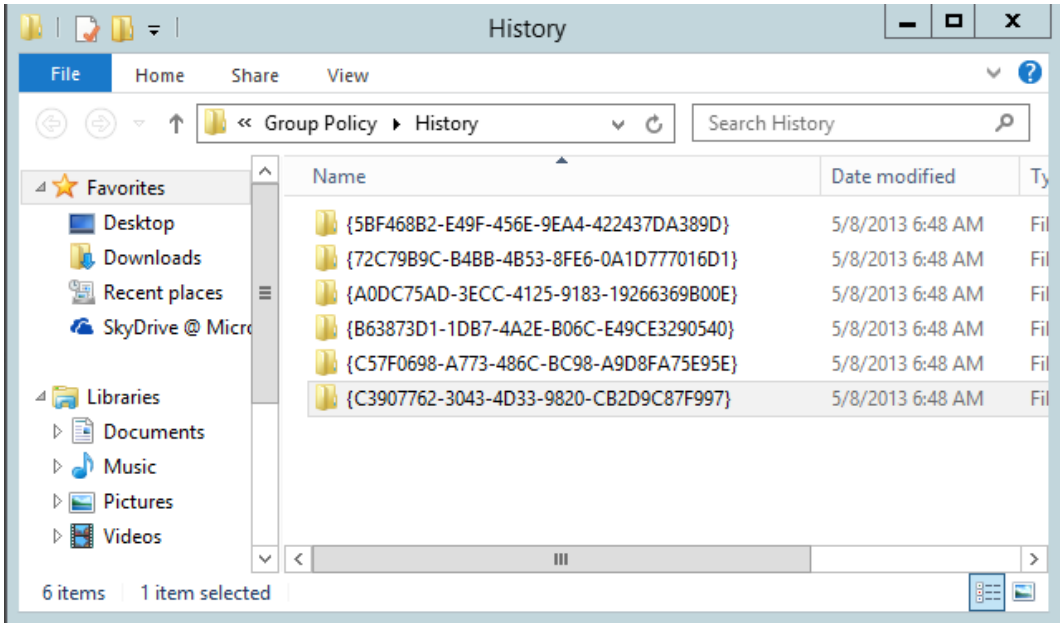

Each Group Policy object folder within the history folder contains sub folders where each folder is named to a corresponding security identifier (SID) of a user who has logged on to the current computer. The complete path to the history file for the Group Policy Printers extension is

```
C:\users\[username]\AppData\Local\Microsoft\Group
Policy\History\[GUID]\SID\Preferences\Printers\Printers.xml
```

The extension then concatenates the string **\User\Preferences\Printers\Printers.xml** to the end of the GPO's file system path. This creates a full path, including file name, to the configuration file.

Note
Windows allows the configuration and application of Group Policy Printers to computers and users.

The extensions opens this file and reads the entire contents of the file into memory and then closes the file. The CSE logs any errors occurring that may occur with the file operations and moves to the next GPO in the list.

Configuration File

The configuration file contains an outer element with multiple inner elements where each inner element represents a separate drive map preference item. Each inner element includes a Properties element that holds the respective elements configuration information.

The extension begins by settings its focus on the first inner element. The extension scans the current inner element for the **removePolicy** attribute. An inner element without a **removePolicy** attribute or an inner element with a **removePolicy** attribute value equaling false (zero [0]) causes the extension to move to the next inner element.

An inner element with a **removePolicy** attribute equaling true (one [1]) indicates to the extension that it must remove the settings in the current inner element. The extension uses the settings from the current inner element and its associated **Properties** element to undo, or remove the settings. The extension accomplishes this by forcing the processing to use the Remove action regardless of the action configured in the current inner element and continues this cycle until it has processed all inner elements in the configuration file. Once processing of all inner elements is complete, the extension deletes the **\Printers\Printers.xml** folder and file for the specific computer or user and GPO from the local history.

The extension then moves to the next Group Policy object in the list and processes it in the same fashion until it has processed all GPOs in the out-of-scope list. The extension begins processing the new and changed GPO list.

New and Changed Group Policy objects

The Group Policy Printers extension processes the entire list of new and changed GPOs sequentially. The extension logs any errors reading the Group Policy objects and continues processing until it reaches the end of the list.

The extension cycles through the list of Group Policy objects in the list starting with the first entry in the list. The extension collects the prerequisite information about the current GPO. From this information, the extension builds file path location to the configuration file. The configuration file location path begins with the GPO's file system path. The extension then concatenates the string **\User\Preferences\Printers\Printers.xml** to the end of the GPO's file system path. This creates a full path, including file name, to the configuration file.

Note

Windows allows the configuration and application of Group Policy Printers to computers and users.

Configuration File

The extension reads the entire configuration file into memory. The configuration file contains an outer element with multiple inner elements where each inner element represents a separate drive map preference item. Each inner element includes a **Properties** element that holds the respective elements configuration information.

Outer Element Processing

The extension evaluates the outer element for the **disabled** attribute. An outer element with a **disabled** attribute equaling true (one [1]) indicates the all inner elements are considered disabled. The extension discontinues processing any part of the current configuration and moves to the next Group Policy object in the list. An outer element without a **disabled** attribute or with a **disabled** attribute value equaling false (zero [0]) continues processing by evaluating the inner elements,

Inner Element Processing

The extension begins by settings its focus on the first inner element. The extension evaluates the current inner element for the **disabled** attribute. An inner element with a **disabled** attribute equaling true (one [1]) indicates the inner element is disabled. The extension discontinues processing the current inner element and move to the next inner element. An inner element without a **disabled** attribute or inner element with a **disabled** attribute value equaling false (zero [0]) continues evaluating the inner element for item-level targeting information.

Item-level Targeting Processing

The extension scans the current inner element for a **Filters** element. The **Filters** element is an optional element that holds the item-level targeting configuration information for each inner element. A **Filters** element existing within the inner element indicates the extension must evaluate one or more Boolean criteria before processing the inner element configuration. A false result from a **Filters** evaluation forces the extension to ignore the current inner element and move to the next inner element. A true result from a **Filters** evaluation indicates the extension should continue processing the inner element and action configured within the current element.

Action Processing

Action processing occurs once the client side extension has evaluated all other criteria that would instruct it to bypass processing outer or the current inner element. Action processing correlates to the action configured in each inner element (reminder an inner element represents a single preference item). The four actions the Group Policy Printers extension can process are create, replace, update, and delete. The tree preferences items the Group Policy Printers extension recognizes are local printers, shared printers, and TCP/IP printers.

Local Printer

Create

The create action represents a configuration where the extension creates a local printer based on the configuration in the current inner element.

Prologue

The extension sets the default values for <default> and <deleteAll> to false and then extracts the values for <port>, <name>, <location>, <comment>, <path>, <processor>, and <dataType> from the configuration information.

Core

The core phase of create action processing uses the configuration information read during the prologue phase to create a local printer for the computer or current user.

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to abort its process and return an object already exists error. A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to proceed with the driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing and to return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the driver needed for the printer installation is not located on the remote location specified by the <path> attribute value, then the extension closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed with adding the local printer to the computer or current user.

Add the printer

The extension continues processing by creating a new local printer based on the configuration in the current inner element. Adding a printer begins with the extension reacquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API.

A successful response from the API enables the extension to gather supplemental configuration information from the remote printer using the EnumPrinters Windows API. An unsuccessful response from the EnumPrinters API causes the extension to close the printer connection using the ClosePrinter Windows API, abort its processing, and return the underlying error it received from the EnumPrinters API. A successful response from the EnumPrinters API causes the extension to continue processing by closing the printer connection using the ClosePrinter API.

The extensions modifies a portion of the supplemental data received from the EnumPrinters API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension ensures the <port> and <name> attribute values are not empty. If either of these attribute values are empty, then the extension aborts its processing and returns an invalid printer name error.

Once the <port> and <name> attributes are validated, the extension then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered. The extension then determines if the <default> attribute value equals true. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental data. A false (zero [0]) causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the <port>, <name>, <location>,<path>, and <comment> attributes to the AddPrinter Windows API. An unsuccessful result from the API causes the extension to abort its processing and return the underlying error it received from the API while a successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

To this point, the extension verified a local printer using the same name as the value configured in <name> attribute does not exist. The extension connected to the remote printer, downloaded and installed the printer driver, and subsequently collected supplemental data in which it combined with its configuration data to create the newly created local printer. The create action processing concludes by setting the default printer.

Set the default printer

The last successful result received from Windows was from using the AddPrinter API. Therefore, the extension now determines whether to set the default printer by checking the <default> attribute value. A false (zero [0]) attribute value instructs the extension that the printer is not configured to be the default printer and exits without an error.

A true (one [1]) value instructs the extension to make the newly created printer the default printer for the computer or user. To do this, the extension switches to the security context of the local user (when applying to the current user). Any failure resulting from the security context switch causes the extension to abort its processing and return the underlying error it received when attempting the security context switch.

The extension proceeds by passing the <name> attribute value to the SetDefaultPrinter Windows API. The extension exits without an error when it receives a successful response from API. Otherwise, the extension exits with the underlying error it received from the SetDefaultPrinter API.

Create action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Delete

The delete action represents a configuration where the extension deletes a specific local printer based on the configuration in the current inner element.

Prologue

The extension sets the default values for <default> and <deleteAll> to false and then extracts the values for <port>, <name>, <location>, <comment>, <path>, <processor>, and <dataType> from the configuration information.

Core

The core phase of delete action processing uses the configuration information read during the prologue phase to delete a local printer for the computer or current user.

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to gracefully exit without an error. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to proceed with the deletion phase.

Epilogue

The extension has inventoried all local printers for the computer or current user and has found a local printer with name matching the <name> attribute value.

Deletion

The extension begins the deletion processing by passing the <name> attribute value to the OpenPrinter Windows API requesting PRINTER_ALL_ACCESS. An unsuccessful response from the API with an underlying invalid printer name error causes the extension to abort its processing and return an unexpected error. An unsuccessful result from the API causes the extension to abort its processing and return the underlying error it received from the API.

A successful response enables the extension to pass the printer handle it received from the OpenPrinter API to the DeletePrinter Windows API. An unsuccessful response from the API causes the extension to close the printer connection using the ClosePrinter API and returning the underlying it received from the API. A successful response from the API cause the extension to close the printer connection using the ClosePrinter API and exiting the function without an error.

Delete action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Replace

The replace action represents a configuration where the extension replaces an existing local printer with a newly created local printer based on the configuration in the current inner element.

It is important to understand the differences between replace and update. The replace action typically deletes an existing local printer and creates a new one—essentially changing all configuration elements of the local printer. The update action only updates a portion of the configuration elements for an existing local printer.

Prologue

The extension sets the default values for <default> and <deleteAll> to false and then extracts the values for <port>, <name>, <location>, <comment>, <path>, <processor>, and <dataType> from the configuration information.

Core

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to proceed to the Driver installation phase. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to proceed with the deletion phase.

Deletion

The extension begins deletion processing by passing the <name> attribute value to the OpenPrinter Windows API requesting PRINTER_ALL_ACCESS permissions. An unsuccessful response from the API with an underlying invalid printer name error causes the extension to abort its processing and return an unexpected error. An unsuccessful result from the API causes the extension to abort its processing and return the underlying error it received from the API.

A successful response enables the extension to pass the printer handle it received from the OpenPrinter API to the DeletePrinter Windows API. An unsuccessful response from the DeletePrinter API causes the extension to close the printer connection using the ClosePrinter API and return the underlying it received from the DeletePrinter API. A successful response from the API causes the extension to close the printer connection using the ClosePrinter API and proceed with Driver installation.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing and to return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the extension does not locate the driver needed for the printer installation on the remote location specified by the <path> attribute value, then it closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed with adding the local printer to the computer or current user.

Add the printer

Adding a printer begins with the extension reacquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API.

A successful response from the API enables the extension to gather supplemental configuration information from the remote printer using the EnumPrinters Windows API. An unsuccessful response from the EnumPrinters API causes the extension to close the printer connection using the ClosePrinter Windows API, abort its processing, and return the underlying error it received from the EnumPrinters API. A successful response from the API causes the extension to continue processing by closing the printer connection using the ClosePrinter API.

The extensions modifies a portion of the supplemental data received from the EnumPrinters API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension ensures the <port> and <name> attribute values are not empty. If either of these attribute values are empty, then the extension aborts its processing and returns an invalid printer name error.

Once the <port> and <name> attributes are validated, the extension then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered. The extension then determines if the <default> attribute value equals true. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental data. A false (zero [0]) causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the <port>, <name>, <location>, <path>, and <comment> attributes to the AddPrinter Windows API. An unsuccessful result from the API causes the extension to abort its processing and return the underlying error it received from the API while a successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

Set the default printer

The last successful result received from Windows was from using the AddPrinter API. Therefore, the extension now determines whether to set the default printer by checking the <default> attribute value. A false (zero [0]) attribute value instructs the extension that the printer is not configured to be the default printer and exits without an error.

A true (one [1]) value instructs the extension to make the newly created printer the default printer for the computer or user. To do this, the extension switches to the security context of the local user (when applying to the current user). Any failure resulting from the security context switch causes the extension to abort its processing and return the underlying error it received when attempting the security context switch.

The extension proceeds by passing the <name> attribute value to the SetDefaultPrinter Windows API. The extension exits without an error when it receives a successful response from API. Otherwise, the extension exits with the underlying error it received from the SetDefaultPrinter API.

Replace action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Update

The update action represents a configuration where the extension refreshes an existing local printer based on the configuration in the current inner element. It is important to understand the difference between replace and update.

The replace action typically deletes the existing local printer and creates a new one—essentially changing all configurations elements of the local printer. The update action only updates a portion of configuration elements in an existing local printer.

Prologue

The extension sets the default values for <default> and <deleteAll> to false and then extracts the values for <port>, <name>, <location>, <comment>, <path>, <processor>, and <dataType> from the configuration information.

Core

The core phase of create action processing uses the configuration information read during the prologue phase to create a local printer for the computer or current user.

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute.

A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to proceed to the Set the default printer phase. A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to proceed with the driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing and to return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the driver needed for the printer installation is not located on the remote location specified by the <path> attribute value, then the extension closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed with adding the local printer to the computer or current user.

Add the printer

The extension continues processing by creating a new local printer based on the configuration in the current inner element. Adding a printer begins with the extension reacquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API.

A successful response from the API enables the extension to gather supplemental configuration information from the remote printer using the EnumPrinters Windows API. An unsuccessful response from the EnumPrinters API causes the extension to close the printer connection using the ClosePrinter Windows API, abort its processing, and return the underlying error it received from the EnumPrinters API. A successful response from the EnumPrinters API causes the extension to continue processing by closing the printer connection using the ClosePrinter API.

The extensions modifies a portion of the supplemental data received from the EnumPrinters API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension ensures the <port> and <name> attribute values are not empty. If either of these attribute values are empty, then the extension aborts its processing and returns an invalid printer name error.

Once the <port> and <name> attributes are validated, the extension then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered. The extension then determines if the <default> attribute value equals true. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental data. A false (zero [0]) causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the <port>, <name>, <location>,<path>, and <comment> attributes to the AddPrinter Windows API. An unsuccessful result from the API causes the extension to abort its processing and return the underlying error it received from the API while a successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

To this point, the extension verified a local printer using the same name as the value configured in <name> attribute does not exist. The extension connected to the remote printer, downloaded and installed the printer driver, and subsequently collected supplemental data in which it combined with its configuration data to create the newly created local printer. The create action processing concludes by setting the default printer.

Set the default printer

The last successful result received from Windows was from using the AddPrinter API. Therefore, the extension now determines whether to set the default printer by checking the <default>

attribute value. A false (zero [0]) attribute value instructs the extension that the printer is not configured to be the default printer and exits without an error.

A true (one [1]) value instructs the extension to make the newly created printer the default printer for the computer or user. To do this, the extension switches to the security context of the local user (when applying to the current user). Any failure resulting from the security context switch causes the extension to abort its processing and return the underlying error it received when attempting the security context switch.

The extension proceeds by passing the <name> attribute value to the SetDefaultPrinter Windows API. The extension exits without an error when it receives a successful response from API. Otherwise, the extension exits with the underlying error it received from the SetDefaultPrinter API.

Update action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Port Printer

Create

The create action represents a configuration where the extension creates a TCP/IP printer connection based on the configuration in the current inner element.

Prologue

The extension begins processing **portPrinter** inner elements by extracting the values of the <comment>, <location>, <path>, <localName>, <ipAddress>, <deleteAll>, <default>, <skipLocal>, <useDNS>, and <lprQueue> attributes.

The extension reads the next set of attributes from the current inner element and assigns each a default value if a value is not present in the inner element. These attributes include

The <snmpCommunity> with a default value “public”

The <protocol> attribute with a default value of “PROTOCOL_RAWTCP_TYPE”

The <port> attribute with a default value of “9100”

The <doubleSpool>, <snmpEnabled>, and <snmpDevIndex> with a default value of “0”

Lastly, the extension reads the <processor> attribute and provides a default value of “winprint”, and the <dataType> attribute with a default value of “RAW”.

Core

Determine IP address

The extension begins by determining the IP address it uses for the remote printer. It determines this decision based on the <useDNS> attribute value. A <useDNS> attribute value equaling false (zero [0]) causes the extension to use the value of the <ipAddress> attribute as the IPv4 address of the remote printer and proceed with the Check for printer process.

A <useDNS> attribute value equaling true (one [1]) causes the extension to read the fully qualified domain name from the <ipAddress> attribute value and proceed with the Name resolution phase.

Name resolution

The extension begins the name resolution process by using the LoadLibrary Windows API to load the wsock32.dll dynamic linked library into memory. An unsuccessful response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API causes the extension to resolve the name into an IPv4 address.

The extension resolves the name into an IPv4 address by passing the <ipAddress> value, which is a fully qualified domain name of a remote printer, to the GetAddrInfoW Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API returns a list of network addresses for the name.

The extension enters a critical section, iterates through the list of address, stores all IPv4 addresses in memory, and exits the critical section. Any errors encountered due to memory allocations causes the extension to abort its processing and return an out of memory error. Otherwise the extension continues by calling the WsaCleanup API and then evaluates the list of return IPv4 addresses.

The extension uses the first IPv4 address in the list of addresses returned from the GetAddrInfoW API. The extension converts the IPv4 address into a dotted, numerical string format and proceeds to the Check for printer process.

Any error encounter encountered during the name resolution phase causes the extension to write “Error resolving IP address” and the underlying error code to the trace log and proceed to the Check for printer phase.

Check for printer

The extension continues processing by determining if a TCP/IP printer exists for the computer or current user. To do this, the extension prefixes the IPv4 address discovered in the Determine IP address phase with “IP_”. The new variation of the IPv4 address looks like “IP_0.0.0.0”; however, with the IPv4 address of the remote printer.

The extension determines how it searches for an existing TCP/IP printer connection by checking the value of the <localName> attribute. If an attribute value exists, then the extension proceeds to the Matching the printer by name phase. If an attribute value does not exist, then the extension proceeds with the Matching by IPv4 address phase.

Matching by IPv4 address

Detect port and TCP/IP printer

The extension detects an existing TCP/IP Printer connection by port name using the varied instance of the IPv4 address created in the Check for printer process as the matching criterion.

The extension uses the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write “Error finding TCP/IP printer…” to the trace log, and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers. With each iteration, the extension verifies the current printer in the list of printers is not a network printer connection by validating the PRINTER_ATTRIBUTE_NETWORK bit is disabled

on the current printers attributes. Next, the extension performs a case insensitive comparison of the current printer's port name with the varied instance of the IPv4 port name.

The extensions encounters a matched TCP/IP printer connection when the current printer is not a shared printer connection and the current printer's port name matches the varied instance of the IPv4 port name.

A complete cycle through the list of printers without finding a matched printer (port or printer name) causes the extension to proceed with the Get remote printer settings phase.

Upon matching a TCP/IP printer connection (port and printer), the extension saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Get remote printer settings phase.

Matching by name

Detect port

The matching by name phase uses the varied instance of the IPv4 address created in the Check for printer phase and passes it to the OpenPrinter Windows API using the XcvPort and SERVER_ACCESS_ADMINISTRATOR permissions. Any error encountered when attempting to connect to the printer causes the extension to internally report that it could not locate the specified remote port and proceeds with the Detect local printers phase.

A successful response from the API instructs the extension to use the GetConfigInfo API from the Windows Device Driver TCPMON XcvData category of functions. An unsuccessful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it could not locate the specified remote port, and proceed with the Inventory local printers phase. A successful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it located the specified remote port, and proceed with the Detect local printers phase.

Detect local printers

The extension continues by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to internally report it found a matching printer and to proceed to the Get remote printer settings phase.

A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to internally report it did not find a matching printer and to proceed to the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by evaluating the <localName> attribute value. The extension uses this value to determine if it needs retrieve the printer name for the remote printer. The

extension already has the printer name if it matched the printer by name. In this scenario, the extension proceeds to the Detect default printer override phase.

The extension needs to retrieve the printer name if it matched the printer by IPv4 address. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing, write "Error getting TCP/IP printer settings …" to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Detect default printer override phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create port phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write "Error finding local printers", and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates that the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the

previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to the Create port phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]) and proceed to the Create port phase.

Create port

The extension determines if it needs to create the port for the computer or current user by internally referencing if it discovered the port through either the Detect port and TCP/IP printer phase or the Detect port phase. The extension proceeds to the Create printer phase if it successfully discovered the port.

Before creating the port, the extension ensures the computer has the necessary components installed to create a TCP/IP printer port. The extension accomplishes this using the EnumMonitors Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful result from the API also returns a list of print monitors available on the computer. The extension iterates through the list of print monitors until it locates a monitor name that matches the string “Standard TCP/IP Port”. If the extension does not discover the Standard TCP/IP Port monitor, then the extension aborts its processing, writes “Error adding TCP/IP printer port to the trace log, and returns an invalid printer monitor error.

The extension continues with port creation by passing the varied instance of the IPv4 address (port name) to the OpenPrinter Windows API using the XcvPort moniker and SERVER_ACCESS_ADMINISTRATOR permissions. A failed result from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful response from the API causes the extension use the AddPort Windows Device Driver TCPMON XcvData API to create the port using the following values:

1. The varied instance of the IPv4 address (port name)
2. The DNS host name or IPv4 address (whichever is the value of the <ipAddress> attribute)
3. The numerical value of the <protocol> attribute

4. The <portNumber> attribute value
5. The <doubleSpool> attribute value
6. The <snmpEnabled> attribute value
7. The <snmpDevIndex> attribute value
8. The <lprQueue> attribute value
9. The <snmpCommunity> attribute value
10.

The extension ignores a printer already existed error, uses the ClosePrinter API, and proceeds with the Create printer phase. All other unsuccessful results returned from the API cause the extension to use the ClosePrinter API, abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the AddPort API. A successful response from the AddPort API causes the extension to use the ClosePrinter API and proceed to the Create Printer phase.

Create printer

The extension determines if it needs to create the printer for the computer or current user by internally referencing if it discovered the printer through either the Detect port and TCP/IP printer phase or the Detect port phase. The extension gracefully exits with an object name exists error if it previously discovered the printer. Otherwise the extension proceeds to the Driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the new TCP/IP printer connection. An unsuccessful response from the Windows API causes the extension to abort its processing, write “Error installing printer driver” to the trace log, and return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the extension cannot locate the driver needed for the printer installation on the remote location specified by the <path> attribute value, then it closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an out of memory error.

All other failed responses from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed to the Add TCP/IP printer phase.

Add TCP/IP printer

The extension continues processing by creating a new TCP/IP printer connection based on the configuration in the current inner element. Adding a printer begins with the extension acquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error.

A successful response from the API causes the extension to retrieve additional information about the remote printer. Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error. The additional information the extension retrieves from the remote printer includes the following information: printer name, driver name, server name, share name, printer processor, location, and comments. The extension continues by releasing the printer handle using the ClosePrinter API.

The extension modifies a portion of the supplemental data received from the GetPrinter API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension continues by ensuring it has previously discovered the port name and printer name. If either of these internal values remain empty, then the extension aborts its processing, writes “Error creating TCP/IP printer connection” to the trace log, and returns an invalid printer name error.

The extension assigns the port name and printer name values to newly created printer and then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered.

The remaining data modified by the extension include <location>, <comment>, <processor>, and <dataType> values. The extension overwrites the information retrieved from the GetPrinter API with information from the configuration file, when the values in the configuration file are not empty.

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental

data. A false (zero [0]) value causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the port name, printer name, <location>, <path>, and <comment> attributes to the AddPrinter Windows API.

The extension ignores an unsuccessful result from the API with a printer already exists error code. Otherwise, all other unsuccessful results from the API cause the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API. A successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

Set default printer

The extension concludes its processing by configuring the printer as the default printer. The extension has already gracefully exited if it previously discovered the printer during the Detect port and TCP/IP printers or Detect local printers phases. Therefore, the extension processes the default printer phase only when it has not previously discovered the printer (just created the new printer).

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. An internal <default> value equaling false (zero [0]) causes the extension to exit gracefully without reporting an error.

An internal <default> value equaling true (one [1]) causes the extension to change its security context to the local system. Any error occurring when switching security contexts causes the extension to abort its processing, write “Error settings default printer” to the trace log, and return the underlying error it encountered during the security context switch.

The extension sets the default printer by passing the internal printer name value to the SetDefaultPrinter Windows API. An unsuccessful result from the API causes the extension write “Error setting default printer” to the trace log, exit gracefully, and return the underlying error it received from the API. Otherwise, the extension exits gracefully without any error.

Create action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Delete

The delete action represents a configuration where the extension deletes a specific TCP/IP printer connection based on the configuration in the current inner element for the computer or current user or all TCP/IP printer connections for the computer or current user.

Prologue

The extension begins processing portPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <localName>, <ipAddress>, <deleteAll>, <default>, <skipLocal>, <useDNS>, and <lprQueue> attributes.

The extension reads the next set of attributes from the current inner element and assigns each a default value if a value is not present in the inner element. These attributes include

The <snmpCommunity> with a default value “public”

The <protocol> attribute with a default value of “PROTOCOL_RAWTCP_TYPE”

The <port> attribute with a default value of “9100”

The <doubleSpool>, <snmpEnabled>, and <snmpDevIndex> with a default value of “0”

Lastly, the extension reads the <processor> attribute and provides a default value of “winprint”, and the <dataType> attribute with a default value of “RAW”.

Core

Delete all TCP/IP printer connections

The extension beings the core processing by checking the value of the <deleteAll> attribute. A <deleteAll> attribute value equaling false causes the extension to proceed to the Delete a specific TCP/IP printer connection phase.

Inventory local printers

The extension inventories local printers by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API. A successful result from the API causes the extension to cycle through the list of returned local printers to determine which local printers are TCP/IP printer connections.

Identify and delete TCP/IP printer connections

The extensions begins the process of deleting TCP/IP printer connections by integrating through the list of local printers. The extension determines if the current printer’s port name of the list of local printers begins with the string “IP_”.

If the current printer’s port name begins with the string “IP_”, then the extension passes the current printer’s name and PRINTER_ALL_ACCESS to the OpenPrinter Windows API.

An unsuccessful response from the API with an invalid printer name error causes the extension to abort its processing and return an unexpected error. All other unsuccessful response from the API cause the extension to abort its processing, and return the underlying error it received from the API. A successful response from API causes the extension to pass the printer handle received from the OpenPrinter API to the DeletePrinter Windows API.

An unsuccessful response from the API causes the extension to close the printer handle using the ClosePrinter API, abort its processing, and return the underlying error it received from the API. A successful response from the API causes the extension to continue its iteration of printers and subsequent deletion.

The extension continues the cycle through printers until it has examined all printers in the list of printers or it encounters an error attempting to delete a printer. The extension closes the printer handle it received from OpenPrinter API and gracefully exits without an error after successfully examining all printers in the list.

Delete a specific TCP/IP printer connection

The extension needs to delete a specific TCP/IP printer connection rather than all TCP/IP printer connections based on the <deleteAll> attribute value equaling false (zero [0]).

Determine IP address

The extension begins by determining the IP address it uses for the remote printer. It determines this decision based on the <useDNS> attribute value. A <useDNS> attribute value equaling false (zero [0]) causes the extension to use the value of the <ipAddress> attribute as the IPv4 address of the remote printer and proceed with the Check for printer process.

A <useDNS> attribute value equaling true (one [1]) causes the extension to read the fully qualified domain name from the <ipAddress> attribute value and proceed with the Name resolution phase.

Name resolution

The extension begins the name resolution process by using the LoadLibrary Windows API to load the wsock32.dll dynamic linked library into memory. An unsuccessful response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API causes the extension to resolve the name into an IPv4 address.

The extension resolves the name into an IPv4 address by passing the <ipAddress> value, which is a fully qualified domain name of a remote printer, to the GetAddrInfoW Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API returns a list of network addresses for the name.

The extension enters a critical section, iterates through the list of address, stores all IPv4 addresses in memory, and exits the critical section. Any errors encountered due to memory allocations causes the extension to abort its processing and return an out of memory error. Otherwise the extension continues by calling the WsaCleanup API and then evaluates the list of return IPv4 addresses.

The extension uses the first IPv4 address in the list of addresses returned from the GetAddrInfoW API. The extension converts the IPv4 address into a dotted, numerical string format and proceeds to the Check for printer phase.

Any error encounter encountered during the name resolution phase causes the extension to write “Error resolving IP address” and the underlying error code to the trace log and proceed to the Check for printer phase.

Check for printer

The extension continues processing by determining if a TCP/IP printer exists for the computer or current user. To do this, the extension prefixes the IPv4 address discovered in the Determine IP address phase with “IP_”. The new variation of the IPv4 address looks like “IP_0.0.0.0”; however, with the IPv4 address of the remote printer.

The extension determines how it searches for an existing TCP/IP printer connection by checking the value of the <localName> attribute. If an attribute value exists, then the extension proceeds to the Matching the printer by name phase. If an attribute value does not exist, then the extension proceeds with the Matching by IPv4 address phase.

Matching by IPv4 address

Detect port and TCP/IP printer

The extension detects an existing TCP/IP Printer connection by port name using the varied instance of the IPv4 address created in the Check for printer process as the matching criterion.

The extension uses the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write “Error finding TCP/IP printer…” to the trace log, and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers. With each iteration, the extension verifies the current printer in the list of printers is not a network printer connection by validating the PRINTER_ATTRIBUTE_NETWORK bit is disabled on the current printers attributes. Next, the extension performs a case insensitive comparison of the current printer’s port name with the varied instance of the IPv4 port name.

The extensions encounters a matched TCP/IP printer connection when the current printer is not a shared printer connection and the current printer’s port name matches the varied instance of the IPv4 port name.

A complete cycle through the list of printers without finding a matched printer (port or printer name) causes the extension to proceed with the Get remote printer settings phase.

Upon matching a TCP/IP printer connection (port and printer), the extension saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Delete TCP/IP printer connection phase.

Matching by name

Detect port

The matching by name phase uses the varied instance of the IPv4 address created in the Check for printer phase and passes it to the OpenPrinter Windows API using the XcvPort and SERVER_ACCESS_ADMINISTRATOR permissions. Any error encountered when attempting to connect to the printer causes the extension to internally report that it could not locate the specified remote port and proceeds with the Detect local printers phase.

A successful response from the API instructs the extension to use the GetConfigInfo API from the Windows Device Driver TCPMON XcvData category of functions. An unsuccessful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it could not locate the specified remote port, and proceed with the Inventory local printers phase. A successful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it located the specified remote port, and proceed with the Detect local printers phase.

Detect local printers

The extension continues by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to internally report it found a matching printer and to proceed to the Delete TCP/IP printer connection phase.

A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to internally report it did not find a matching printer and to proceed to the Delete TCP/IP printer connection phase.

Delete TCP/IP printer connection

The extension internally references if it successfully discovered the printer during the Check for printer phase. If the extension did not successfully discover the printer at that time, then the extension proceeds to the Get remote printer settings phase.

Successfully discovering the printer during the Check for printer phase causes the extension to pass the current printer's name and PRINTER_ALL_ACCESS to the OpenPrinter Windows API.

An unsuccessful response from the API with an invalid printer name error causes the extension to abort its processing and return an unexpected error. All other unsuccessful response from the API cause the extension to abort its processing, and return the underlying error it received from the API. A successful response from API causes the extension to pass the printer handle received from the OpenPrinter API to the DeletePrinter Windows API.

An unsuccessful response from the API causes the extension to close the printer handle using the ClosePrinter API, abort its processing, and return the underlying error it received from the API. A successful response from the API causes the extension to proceed to the Get remote printer settings phase.

Epilogue

Get remote printer settings

The extension begins this phase by evaluating the <localName> attribute value. The extension uses this value to determine if it needs retrieve the printer name for the remote printer. The extension already has the printer name if it matched the printer by name. In this scenario, the extension proceeds to the Detect default printer override phase.

The extension needs to retrieve the printer name if it matched the printer by IPv4 address. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing, write "Error getting TCP/IP printer settings …" to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error. A successful response from the API causes the extension to

retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and gracefully exits without reporting an error.

Delete action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Replace

The replace action represents a configuration where the extension replaces an existing TCP/IP printer connection with a newly created TCP/IP printer connection based on the configuration in the current inner element.

It is important to understand the differences between replace and update. The replace action typically deletes an existing TCP/IP printer connection and creates a new one—essentially changing all configuration elements of the printer connection. The update action only updates a portion of the configuration elements for an existing TCP/IP printer connection.

Prologue

The extension begins processing portPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <localName>, <ipAddress>, <deleteAll>, <default>, <skipLocal>, <useDNS>, and <lprQueue> attributes.

The extension reads the next set of attributes from the current inner element and assigns each a default value if a value is not present in the inner element. These attributes include

The <snmpCommunity> with a default value “public”

The <protocol> attribute with a default value of “PROTOCOL_RAWTCP_TYPE”

The <port> attribute with a default value of “9100”

The <doubleSpool>, <snmpEnabled>, and <snmpDevIndex> with a default value of “0”

Lastly, the extension reads the <processor> attribute and provides a default value of “winprint”, and the <dataType> attribute with a default value of “RAW”.

Core

Delete a specific TCP/IP printer connection

The extension needs to delete a specific TCP/IP printer connection based on the <deleteAll> attribute value equaling false (zero [0]).

Determine IP address

The extension begins by determining the IP address it uses for the remote printer. It determines this decision based on the <useDNS> attribute value. A <useDNS> attribute value equaling false (zero [0]) causes the extension to use the value of the <ipAddress> attribute as the IPv4 address of the remote printer and proceed with the Check for printer process.

A <useDNS> attribute value equaling true (one [1]) causes the extension to read the fully qualified domain name from the <ipAddress> attribute value and proceed with the Name resolution phase.

Name resolution

The extension begins the name resolution process by using the LoadLibrary Windows API to load the wsock32.dll dynamic linked library into memory. An unsuccessful response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API causes the extension to resolve the name into an IPv4 address.

The extension resolves the name into an IPv4 address by passing the <ipAddress> value, which is a fully qualified domain name of a remote printer, to the GetAddrInfoW Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API returns a list of network addresses for the name.

The extension enters a critical section, iterates through the list of address, stores all IPv4 addresses in memory, and exits the critical section. Any errors encountered due to memory allocations causes the extension to abort its processing and return an out of memory error. Otherwise the extension continues by calling the WsaCleanup API and then evaluates the list of return IPv4 addresses.

The extension uses the first IPv4 address in the list of addresses returned from the GetAddrInfoW API. The extension converts the IPv4 address into a dotted, numerical string format and proceeds to the Check for printer phase.

Any error encounter encountered during the name resolution phase causes the extension to write “Error resolving IP address” and the underlying error code to the trace log and proceed to the Check for printer phase.

Check for printer

The extension continues processing by determining if a TCP/IP printer exists for the computer or current user. To do this, the extension prefixes the IPv4 address discovered in the Determine IP address phase with “IP_”. The new variation of the IPv4 address looks like “IP_0.0.0.0”; however, with the IPv4 address of the remote printer.

The extension determines how it searches for an existing TCP/IP printer connection by checking the value of the <localName> attribute. If an attribute value exists, then the extension proceeds to the Matching the printer by name phase. If an attribute value does not exist, then the extension proceeds with the Matching by IPv4 address phase.

Matching by IPv4 address

Detect port and TCP/IP printer

The extension detects an existing TCP/IP Printer connection by port name using the varied instance of the IPv4 address created in the Check for printer process as the matching criterion.

The extension uses the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write “Error finding TCP/IP printer…” to the trace log, and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers. With each iteration, the extension verifies the current printer in the list of printers is not a network printer connection by validating the PRINTER_ATTRIBUTE_NETWORK bit is disabled on the current printers attributes. Next, the extension performs a case insensitive comparison of the current printer’s port name with the varied instance of the IPv4 port name.

The extensions encounters a matched TCP/IP printer connection when the current printer is not a shared printer connection and the current printer’s port name matches the varied instance of the IPv4 port name.

A complete cycle through the list of printers without finding a matched printer (port or printer name) causes the extension to proceed with the Get remote printer settings phase.

Upon matching a TCP/IP printer connection (port and printer), the extension saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Delete TCP/IP printer connection phase.

Matching by name

Detect port

The matching by name phase uses the varied instance of the IPv4 address created in the Check for printer phase and passes it to the OpenPrinter Windows API using the XcvPort and SERVER_ACCESS_ADMINISTRATOR permissions. Any error encountered when attempting to connect to the printer causes the extension to internally report that it could not locate the specified remote port and proceeds with the Detect local printers phase.

A successful response from the API instructs the extension to use the GetConfigInfo API from the Windows Device Driver TCPMON XcvData category of functions. An unsuccessful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it could not locate the specified remote port, and proceed with the Inventory local printers phase. A successful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it located the specified remote port, and proceed with the Detect local printers phase.

Detect local printers

The extension continues by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer’s name with the value of the <name> attribute. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to internally report it found a matching printer and to proceed to the Delete TCP/IP printer connection phase.

A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to internally report it did not find a matching printer and to proceed to the Delete TCP/IP printer connection phase.

Delete TCP/IP printer connection

The extension internally references if it successfully discovered the printer during the Check for printer phase. If the extension did not successfully discover the printer at that time, then the extension proceeds to the Get remote printer settings phase.

Successfully discovering the printer during the Check for printer phase causes the extension to pass the current printer's name and PRINTER_ALL_ACCESS to the OpenPrinter Windows API.

An unsuccessful response from the API with an invalid printer name error causes the extension to abort its processing and return an unexpected error. All other unsuccessful response from the API cause the extension to abort its processing, and return the underlying error it received from the API. A successful response from API causes the extension to pass the printer handle received from the OpenPrinter API to the DeletePrinter Windows API.

An unsuccessful response from the API causes the extension to close the printer handle using the ClosePrinter API, abort its processing, and return the underlying error it received from the API. A successful response from the API causes the extension to proceed to the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by evaluating the <localName> attribute value. The extension uses this value to determine if it needs retrieve the printer name for the remote printer. The extension already has the printer name if it matched the printer by name. In this scenario, the extension proceeds to the Detect default printer override phase.

The extension needs to retrieve the printer name if it matched the printer by IPv4 address. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing, write "Error getting TCP/IP printer settings …" to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, saves the printer name to memory, alters a copy

of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Detect default printer override phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create port phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write “Error finding local printers”, and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates that the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to the Create port phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]) and proceed to the Create port phase.

Create port

The extension determines if it needs to create the port for the computer or current user by internally referencing if it discovered the port through either the Detect port and TCP/IP printer

phase or the Detect port phase. The extension proceeds to the Create printer phase if it successfully discovered the port.

Before creating the port, the extension ensures the computer has the necessary components installed to create a TCP/IP printer port. The extension accomplishes this using the EnumMonitors Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful result from the API also returns a list of print monitors available on the computer. The extension iterates through the list of print monitors until it locates a monitor name that matches the string “Standard TCP/IP Port”. If the extension does not discover the Standard TCP/IP Port monitor, then the extension aborts its processing, writes “Error adding TCP/IP printer port to the trace log, and returns an invalid printer monitor error.

The extension continues with port creation by passing the varied instance of the IPv4 address (port name) to the OpenPrinter Windows API using the XcvPort moniker and SERVER_ACCESS_ADMINISTRATOR permissions. A failed result from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful response from the API causes the extension use the AddPort Windows Device Driver TCPMON XcvData API to create the the port using the following values:

1. The varied instance of the IPv4 address (port name)
2. The DNS host name or IPv4 address (whichever is the value of the <ipAddress> attribute)
3. The numerical value of the <protocol> attribute
4. The <portNumber> attribute value
5. The <doubleSpool> attribute value
6. The <snmpEnabled> attribute value
7. The <snmpDevIndex> attribute value
8. The <lprQueue> attribute value
9. The <snmpCommunity> attribute value

The extension ignores a printer already existed error, uses the ClosePrinter API, and proceeds with the Create printer phase. All other unsuccessful results returned from the API cause the extension to use the ClosePrinter API, abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the AddPort API. A successful response from the AddPort API causes the extension to use the ClosePrinter API and proceed to the Create Printer phase.

Create printer

The extension determines if it needs to create the printer for the computer or current user by internally referencing if it discovered the printer through either the Detect port and TCP/IP printer phase or the Detect port phase. The extension gracefully exits with an object name exists error if it previously discovered the printer. Otherwise the extension proceeds to the Driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the new TCP/IP printer connection. An unsuccessful response from the Windows API causes the extension to abort its processing, write “Error installing printer driver” to the trace log, and return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the extension cannot locate the driver needed for the printer installation on the remote location specified by the <path> attribute value, then it closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an out of memory error.

All other failed responses from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed to the Add TCP/IP printer phase.

Add TCP/IP printer

The extension continues processing by creating a new TCP/IP printer connection based on the configuration in the current inner element. Adding a printer begins with the extension acquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error.

A successful response from the API causes the extension to retrieve additional information about the remote printer. Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error. The additional information the extension retrieves from the remote printer includes the following information: printer name, driver name, server name, share name, printer processor, location, and comments. The extension continues by releasing the printer handle using the ClosePrinter API.

The extension modifies a portion of the supplemental data received from the GetPrinter API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension continues by ensuring it has previously discovered the port name and printer name. If either of these internal values remain empty, then the extension aborts its processing, writes “Error creating TCP/IP printer connection” to the trace log, and returns an invalid printer name error.

The extension assigns the port name and printer name values to newly created printer and then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered.

The remaining data modified by the extension include <location>, <comment>, <processor>, and <dataType> values. The extension overwrites the information retrieved from the GetPrinter API with information from the configuration file, when the values in the configuration file are not empty.

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental data. A false (zero [0]) value causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the port name, printer name, <location>, <path>, and <comment> attributes to the AddPrinter Windows API.

The extension ignores an unsuccessful result from the API with a printer already exists error code. Otherwise, all other unsuccessful results from the API cause the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API. A successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

Set default printer

The extension concludes its processing by configuring the printer as the default printer. The extension has already gracefully exited if it previously discovered the printer during the Detect port and TCP/IP printers or Detect local printers phases. Therefore, the extension processes the default printer phase only when it has not previously discovered the printer (just created the new printer).

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. An internal <default> value equaling false (zero [0]) causes the extension to exit gracefully without reporting an error.

An internal <default> value equaling true (one [1]) causes the extension to change its security context to the local system. Any error occurring when switching security contexts causes the extension to abort its processing, write “Error settings default printer” to the trace log, and return the underlying error it encountered during the security context switch.

The extension sets the default printer by passing the internal printer name value to the SetDefaultPrinter Windows API. An unsuccessful result from the API causes the extension write “Error setting default printer” to the trace log, exit gracefully, and return the underlying error it received from the API. Otherwise, the extension exits gracefully without any error.

Update

The update action represents a configuration where the extension refreshes an existing TCP/IP printer connection based on the configuration in the current inner element. It is important to understand the difference between replace and update.

The replace action typically deletes the existing TCP/IP printer connection and creates a new one—essentially changing all configurations elements of the local printer. The update action only updates a portion of configuration elements in an existing TCP/IP printer connection.

Prologue

The extension begins processing portPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <localName>, <ipAddress>, <deleteAll>, <default>, <skipLocal>, <useDNS>, and <lprQueue> attributes.

The extension reads the next set of attributes from the current inner element and assigns each a default value if a value is not present in the inner element. These attributes include

The <snmpCommunity> with a default value “public”

The <protocol> attribute with a default value of “PROTOCOL_RAWTCP_TYPE”

The <port> attribute with a default value of “9100”

The <doubleSpool>, <snmpEnabled>, and <snmpDevIndex> with a default value of “0”

Lastly, the extension reads the <processor> attribute and provides a default value of “winprint”, and the <dataType> attribute with a default value of “RAW”.

Core

Determine IP address

The extension begins by determining the IP address it uses for the remote printer. It determines this decision based on the <useDNS> attribute value. A <useDNS> attribute value equaling false (zero [0]) causes the extension to use the value of the <ipAddress> attribute as the IPv4 address of the remote printer and proceed with the Check for printer process.

A <useDNS> attribute value equaling true (one [1]) causes the extension to read the fully qualified domain name from the <ipAddress> attribute value and proceed with the Name resolution phase.

Name resolution

The extension begins the name resolution process by using the LoadLibrary Windows API to load the wsock32.dll dynamic linked library into memory. An unsuccessful response from the API

causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API causes the extension to resolve the name into an IPv4 address.

The extension resolves the name into an IPv4 address by passing the <ipAddress> value, which is a fully qualified domain name of a remote printer, to the GetAddrInfoW Windows API. A failed response from the API causes the extension to abort its processing and return the underlying error it received from the API. A successful response from the API returns a list of network addresses for the name.

The extension enters a critical section, iterates through the list of address, stores all IPv4 addresses in memory, and exits the critical section. Any errors encountered due to memory allocations causes the extension to abort its processing and return an out of memory error. Otherwise the extension continues by calling the WsaCleanup API and then evaluates the list of return IPv4 addresses.

The extension uses the first IPv4 address in the list of addresses returned from the GetAddrInfoW API. The extension converts the IPv4 address into a dotted, numerical string format and proceeds to the Check for printer process.

Any error encounter encountered during the name resolution phase causes the extension to write “Error resolving IP address” and the underlying error code to the trace log and proceed to the Check for printer phase.

Check for printer

The extension continues processing by determining if a TCP/IP printer exists for the computer or current user. To do this, the extension prefixes the IPv4 address discovered in the Determine IP address phase with “IP_”. The new variation of the IPv4 address looks like “IP_0.0.0.0”; however, with the IPv4 address of the remote printer.

The extension determines how it searches for an existing TCP/IP printer connection by checking the value of the <localName> attribute. If an attribute value exists, then the extension proceeds to the Matching the printer by name phase. If an attribute value does not exist, then the extension proceeds with the Matching by IPv4 address phase.

Matching by IPv4 address

Detect port and TCP/IP printer

The extension detects an existing TCP/IP Printer connection by port name using the varied instance of the IPv4 address created in the Check for printer process as the matching criterion.

The extension uses the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write “Error finding TCP/IP printer…” to the trace log, and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers. With each iteration, the extension verifies the current printer in the list of printers is not a network printer connection by validating the PRINTER_ATTRIBUTE_NETWORK bit is disabled on the current printers attributes. Next, the extension performs a case insensitive comparison of the current printer’s port name with the varied instance of the IPv4 port name.

The extensions encounters a matched TCP/IP printer connection when the current printer is not a shared printer connection and the current printer's port name matches the varied instance of the IPv4 port name.

A complete cycle through the list of printers without finding a matched printer (port or printer name) causes the extension to proceed with the Get remote printer settings phase.

Upon matching a TCP/IP printer connection (port and printer), the extension saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Get remote printer settings phase.

Matching by name

Detect port

The matching by name phase uses the varied instance of the IPv4 address created in the Check for printer phase and passes it to the OpenPrinter Windows API using the XcvPort and SERVER_ACCESS_ADMINISTRATOR permissions. Any error encountered when attempting to connect to the printer causes the extension to internally report that it could not locate the specified remote port and proceeds with the Detect local printers phase.

A successful response from the API instructs the extension to use the GetConfigInfo API from the Windows Device Driver TCPMON XcvData category of functions. An unsuccessful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it could not locate the specified remote port, and proceed with the Inventory local printers phase. A successful response from the API causes the extension to close the connection to the remote printer using the ClosePrinter API, internally report that it located the specified remote port, and proceed with the Detect local printers phase.

Detect local printers

The extension continues by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs a case insensitive comparison of the current printer's name with the value of the <name> attribute. A computer or current user with a local printer name matching the value of the <name> attribute causes the extension to internally report it found a matching printer and to proceed to the Get remote printer settings phase.

A computer or current user without a local printer name matching the value of the <name> attribute causes the extension to internally report it did not find a matching printer and to proceed to the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by evaluating the <localName> attribute value. The extension uses this value to determine if it needs retrieve the printer name for the remote printer. The extension already has the printer name if it matched the printer by name. In this scenario, the extension proceeds to the Detect default printer override phase.

The extension needs to retrieve the printer name if it matched the printer by IPv4 address. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing, write "Error getting TCP/IP printer settings …" to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write "Error getting TCP/IP printer settings" to the trace log, and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, saves the printer name to memory, alters a copy of the printer name by removing any indication of a UNC path from the name—if any—leaving just the printer name, and proceeds with the Detect default printer override phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create port phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to abort its processing, write "Error finding local printers", and return the underlying error received from the Windows API.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates that the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to the Create/Update port phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]) and proceed to the Create port phase.

Create/Update port

Before creating the port, the extension ensures the computer has the necessary components installed to create a TCP/IP printer port. The extension accomplishes this using the EnumMonitors Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful result from the API also returns a list of print monitors available on the computer. The extension iterates through the list of print monitors until it locates a monitor name that matches the string “Standard TCP/IP Port”. If the extension does not discover the Standard TCP/IP Port monitor and the extension aborts its processing. If the extension did not previously discover the port, then it writes “Error adding TCP/IP printer port” to the trace log. The extension writes “Error configuring TCP/IP printer port” to the trace log if it did not previously discover the port. Regardless of port discovery status, the extension returns an invalid printer monitor error.

If the extension did not previously discover the port, then it proceeds with the Create port sub-phase. Otherwise, the extension proceeds with the Update port sub-phase.

Create port

The extension continues with port creation by passing the varied instance of the IPv4 address (port name) to the OpenPrinter Windows API using the XcvPort moniker and SERVER_ACCESS_ADMINISTRATOR permissions. A failed result from the API causes the extension to abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful response from the API causes the extension use the AddPort Windows Device Driver TCPMON XcvData API to create the port using the following values:

1. The varied instance of the IPv4 address (port name)
2. The DNS host name or IPv4 address (whichever is the value of the <ipAddress> attribute)
3. The numerical value of the <protocol> attribute

4. The <portNumber> attribute value
5. The <doubleSpool> attribute value
6. The <snmpEnabled> attribute value
7. The <snmpDevIndex> attribute value
8. The <lprQueue> attribute value
9. The <snmpCommunity> attribute value

The extension ignores a printer already existed error, uses the ClosePrinter API, and proceeds with the Create printer phase. All other unsuccessful results returned from the API cause the extension to use the ClosePrinter API, abort its processing, write “Error adding TCP/IP printer port” to the trace log, and return the underlying error it received from the AddPort API. A successful response from the AddPort API causes the extension to use the ClosePrinter API and proceed to the Create Printer phase.

Update port

The extension continues with updating the port by passing the varied instance of the IPv4 address (port name) to the OpenPrinter Windows API using the XcvPort moniker and SERVER_ACCESS_ADMINISTRATOR permissions. A failed result from the API causes the extension to abort its processing, write “Error configuring TCP/IP printer port” to the trace log, and return the underlying error it received from the API.

A successful response from the API causes the extension use the ConfigurePort Windows Device Driver TCPMON XcvData API to configure the port using the following values:

1. The varied instance of the IPv4 address (port name)
2. The DNS host name or IPv4 address (whichever is the value of the <ipAddress> attribute)
3. The numerical value of the <protocol> attribute
4. The <portNumber> attribute value
5. The <doubleSpool> attribute value
6. The <snmpEnabled> attribute value
7. The <snmpDevIndex> attribute value
8. The <lprQueue> attribute value
9. The <snmpCommunity> attribute value

The extension ignores a printer already existed error, uses the ClosePrinter API, and proceeds with the Create printer phase. All other unsuccessful results returned from the API cause the extension to use the ClosePrinter API, abort its processing, write “Error configuring TCP/IP printer port” to the trace log, and return the underlying error it received from the ConfigurePort API. A successful response from the ConfigurePort API causes the extension to use the ClosePrinter API and proceed to the Create Printer phase.

Create printer

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the new TCP/IP printer connection. An unsuccessful response from the Windows API causes the extension to

abort its processing, write “Error installing printer driver” to the trace log, and return the underlying error provided by the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install driver, provided the driver is not already installed.

If the extension cannot locate the driver needed for the printer installation on the remote location specified by the <path> attribute value, then it closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns an out of memory error.

All other failed responses from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “Error installing printer driver” to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection and proceed to the Add TCP/IP printer phase.

Add TCP/IP printer

The extension continues processing by creating a new TCP/IP printer connection based on the configuration in the current inner element. Adding a printer begins with the extension acquiring a handle to the remote printer using the <path> attribute value and the OpenPrinter Windows API. A failed response from the API causes the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the GetPrinter API causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error.

A successful response from the API causes the extension to retrieve additional information about the remote printer. Any errors regarding memory allocation occurring while gather remote printer information causes the extension to abort its processing, write “Error creating TCP/IP printer settings” to the trace log, and return an unexpected error. The additional information the extension retrieves from the remote printer includes the following information: printer name, driver name, server name, share name, printer processor, location, and comments. The extension continues by releasing the printer handle using the ClosePrinter API.

The extension modifies a portion of the supplemental data received from the GetPrinter API. The first of the data modified by the extension is the security descriptor. The extensions changes the security descriptor data returned by the API to a null value. The Windows authorization

subsystem interprets a null security descriptor as if every user on the computer has full access to the printer.

The extension continues by ensuring it has previously discovered the port name and printer name. If either of these internal values remain empty, then the extension aborts its processing, writes “Error creating TCP/IP printer connection” to the trace log, and returns an invalid printer name error.

The extension assigns the port name and printer name values to newly created printer and then removes the PRINTER_ATTRIBUTE_SHARED bit from the attributes of the supplemental data previously gathered.

The remaining data modified by the extension include <location>, <comment>, <processor>, and <dataType> values. The extension overwrites the information retrieved from the GetPrinter API with information from the configuration file, when the values in the configuration file are not empty.

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. A <default> attribute value equaling true (one [1]) causes the extension to add the PRINTER_ATTRIBUTE_DEFAULT bit to the attributes of the supplemental data. A false (zero [0]) value causes the extension to remove the PRINTER_ATTRIBUTE_DEFAULT bit from the attributes of the supplemental data.

The extension concludes this phase by passing the supplemental data along with the values of the port name, printer name, <location>, <path>, and <comment> attributes to the AddPrinter Windows API.

The extension ignores an unsuccessful result from the API with a printer already exists error code. Otherwise, all other unsuccessful results from the API cause the extension to abort its processing, write “Error creating TCP/IP printer connection” to the trace log, and return the underlying error it received from the API. A successful result from the API enables the extension to continue processing, and set the default printer.

Epilogue

`Set default printer`

The extension concludes its processing by configuring the printer as the default printer. The extension has already gracefully exited if it previously discovered the printer during the Detect port and TCP/IP printers or Detect local printers phases. Therefore, the extension processes the default printer phase only when it has not previously discovered the printer (just created the new printer).

The extension references the internal <default> value as it may have modified the attribute value during the Detect default printer override phase to determine if the new TCP/IP printer connection should be the default printer. An internal <default> value equaling false (zero [0]) causes the extension to exit gracefully without reporting an error.

An internal <default> value equaling true (one [1]) causes the extension to change its security context to the local system. Any error occurring when switching security contexts causes the

extension to abort its processing, write “Error settings default printer” to the trace log, and return the underlying error it encountered during the security context switch.

The extension sets the default printer by passing the internal printer name value to the SetDefaultPrinter Windows API. An unsuccessful result from the API causes the extension write “Error setting default printer” to the trace log, exit gracefully, and return the underlying error it received from the API. Otherwise, the extension exits gracefully without any error.

Update action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Shared Printer

Create

The create action represents a configuration where the extension creates a local printer based on the configuration in the current inner element.

Prologue

The extension begins processing SharedPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <port>, <username>, and <cPassword> attributes from the current inner element.

The extension reads the next set of attributes from the current inner element and assign each attribute a default value of false (zero [0]) if a value is not present in the inner element. These attributes include <deleteAll>, <default>, <skipLocal>, <persistent>, and <deleteMaps>.

If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extensions prepares for the core processing phase by setting the security context to the computer or current user (depending how the preference item is deployed).

Core

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension indicate a matching printer does not exists and proceeds with the Detect default printer override phase.

A successful result from the API causes the extension to cycle through the list of returned local printers. The extension begins the cycle with the current printer and determines if the printer is a network printer by the PRINTER_ATTRIBUTE_NETWORK bit in the current printer’s attributes. If the PRINTER_ATTRIBUTE_NETWORK bit is not enabled, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the PRINTER_ATTRIBUTE_NETWORK bit is disabled, then the extension attempts to discover a matching printer.

Match printer by share name

The extension detects a matched printer when it identifies a local printer sharing the same share name as the <path> attribute value. To begin, the extension uses the GetComputerName Windows API to retrieve and store the computers name. The extension begins a series of comparisons starting with the current printer’ server name.

If the current printer's server name is empty, then the extension creates a temporary path value represented as UNC path using the recently discover computer name for the server portion and the current printer's share name value as the share portion (\\computerName\printerShareName). Additionally, the extension creates a temporay name value represented as a UNC path using the recently discovered computer name for the server portion and the current printer's name as the share portion (\\computerName\printerName).

If the current printer's server name contains a value, then the extension creates a temporary path value represented as a UNC path using the current printer's server name as the server portion and the current printer's share name as the share portion (\\currentPrinterServerName\currentPrinterShareName). Additionally, the extension creates a temporary name value and assigns the current printer's name to the temporary value.

The extension determines if the current printer matches the printer specified in the inner element configuration by comparing the temporary path value to the current printer's path value, comparing the temporary name value to the current printer's path value, and comparing the current printer's port name to the current printer's path value.

If none of these comparisons indicate a successful match, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the extension completes the iteration of printers without finding a matching printer, then the extension indicates a matching printer does not exist and proceeds to the Detect default printer override phase.

If at least one of the comparisons indicates a successful match, then the extension stops iterating through the list of printers and indicates a matching printer exists. The extension writes "FindSharedPrinterEx: " and the <path> attribute value to the trace log, and proceeds to the Detect default printer override phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create printer phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to proceed to the Create printer phase.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the

previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]), write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

Create printer

The extension internally references if it previously detected a local printer. If the extension discovered a printer existed, then the extension exits gracefully with an object name exists error. If the extension did not discover a printer during the Match printer by share name phase, then it proceeds with the Driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing, write “InstallDriversEx:” and the <path> attribute value to the trace log, and return the underlying error it received from the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install the driver, provided the driver is not already installed.

If the driver needed for the printer installation is not located on the remote location specified by the <path> attribute value, then the extension closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx:” and the <path> attribute value to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection, write “InstallDriversEx: “and the <path> attribute value to the trace log, and proceed with the Add shared printer phase.

Add shared printer

The extension adds the shared printer to the computer or current user by passing the <path> attribute value to the AddPrinterConnection Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error code I received from the API.

A successful response from the API causes the extension write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Set default printer phase.

Set default printer

To set the default printer, the extension internally references if it previously discovered an existing printer. If the extension previously discovered a printer, then the extension proceeds to the Map printer port phase. If the extension did not previously discover a printer, then it checks the internal <default> attribute to determine if it was modified during the Detect default printer override phase.

If the extension determined the internal <default> attribute value equals false (zero [0]), then it proceeds to the map printer port phase. If the extension determines the internal <default> attribute value equals true (one [1]), then it passes the <path> attribute value to the SetDefaultPrinter Windows API.

An unsuccessful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error provided by the API. A successful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Map printer port phase.

Map printer port

The extension begins by verifying that the <port> attribute value is not empty. If the <port> attribute value is empty, then the extension proceeds to the Epilogue phase. Otherwise, the extension proceeds with the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by retrieving information about the remote printer. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the API causes

the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and returning the underlying response code it received from the GetPrinter API. The extension creates a temporary internal string value by combining the printers share path and share name in a UNC syntax.

Detect mapped port

Next, the extension determines if the port matching the <port> attribute value is currently mapped to another printer. The extension determines this by passing the <port> attribute value to the WNetGetConnection Windows API.

An unsuccessful response from the API also returns an empty internal remote path value, which indicates the port is not mapped. The extension writes “IsMapped: “along with the <port> attribute value, and the blank internal remote path value to the trace log. The extension proceeds to the Map port phase.

A successful response from the API also returns an internal remote path value in the form of a UNC path to which the port is mapped. The extension writes “IsMapped: “along with the <port> attribute value and the internal remote path value to which the port is mapped to the trace log.

The extension then performs a case insensitive comparison between the <path> attribute value and the internal remote path value. If the comparison results in a match, then the extension proceeds to the Epilogue phase.

If the comparisons results conclude that the two string values do not match, then the extension writes the <port> attribute value “already mapped to “ the remote path value to the trace log, sets the error code to already exists, and proceeds to the Epilogue phase.

Map port

The extension maps the local port matching the <port> attribute value to the <path> attribute value by passing the <port>, <path>, <persistent> attribute values to the WNetUseConnection Windows API. The extension also passes the <username> and <cPassword> attribute values to the API provided these values are not empty.

Regardless of the response, the extension writes “MapPort: “along with the <port> and <shared> attribute values to the trace log. The extension also sets the error code to the underlying error code it received from the WNetUseConnection API. The extension proceeds to the Epilogue phase.

Epilogue

The extension completes the processing by securely deleting the clear text instance of the <cPassword> attribute value.

Create action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Delete

The delete action represents a configuration where the extension deletes a specific shared printer connection based on the configuration in the current inner element for the computer or current user or all shared printer connections for the computer or current user. Also the extension deletes all LPT printer mappings or a specific printer mapping based on the current inner element.

Prologue

The extension begins processing SharedPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <port>, <username>, and <cPassword> attributes from the current inner element.

The extension reads the next set of attributes from the current inner element and assign each attribute a default value of false (zero [0]) if a value is not present in the inner element. These attributes include <deleteAll>, <default>, <skipLocal>, <persistent>, and <deleteMaps>.

If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extensions prepares for the core processing phase by setting the security context to the computer or current user (depending how the preference item is deployed).

Core

The extension begins this processing by determining if it is configured to delete all shared printer connections or a specific printer connection by evaluating the <deleteAll> attribute. If the attribute value equals false (zero [0]), then the extension proceed to the Delete a specific shared printer connection phase. If the attribute value equals true, then the extension proceeds to the Delete all shared printer connections phases

Delete all shared printer connections

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension abort its processing and proceed to the Epilogue phase.

A successful result from the API causes the extension to cycle through the list of returned local printers. The extension begins the cycle with the current printer and determines if the printer is a network printer by the PRINTER_ATTRIBUTE_NETWORK bit in the current printer's attributes. If the PRINTER_ATTRIBUTE_NETWORK bit is not enabled, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the PRINTER_ATTRIBUTE_NETWORK bit is disabled, then the extension attempts to discover a matching printer.

Delete shared printer connection

The extension deletes the shared printer connection by passing the current printer name to the DeletePrinterConnection Windows API. An unsuccessful response indicating an invalid printer name causes the extension to abort its processing and return an unexpected error. All other unsuccessful responses cause the extension to abort its processing and return the underlying error it received from the API. A successful response from the API cause the extension to proceed to the Epilogue phase.

Delete a specific shared printer connection

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension indicate a matching printer does not exists and proceeds with the Detect default printer override phase.

A successful result from the API causes the extension to cycle through the list of returned local printers. The extension begins the cycle with the current printer and determines if the printer is a network printer by the PRINTER_ATTRIBUTE_NETWORK bit in the current printer's attributes. If the PRINTER_ATTRIBUTE_NETWORK bit is not enabled, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the PRINTER_ATTRIBUTE_NETWORK bit is disabled, then the extension attempts to discover a matching printer.

Match printer by share name

The extension detects a matched printer when it identifies a local printer sharing the same share name as the <path> attribute value. To begin, the extension uses the GetComputerName Windows API to retrieve and store the computers name. The extension begins a series of comparisons starting with the current printer' server name.

If the current printer's server name is empty, then the extension creates a temporary path value represented as UNC path using the recently discover computer name for the server portion and the current printer's share name value as the share portion (\\computerName\printerShareName). Additionally, the extension creates a temporary name value represented as a UNC path using the recently discovered computer name for the server portion and the current printer's name as the share portion (\\computerName\printerName).

If the current printer's server name contains a value, then the extension creates a temporary path value represented as a UNC path using the current printer's server name as the server portion and the current printer's share name as the share portion (\\currentPrinterServerName\currentPrinterShareName). Additionally, the extension creates a temporary name value and assigns the current printer's name to the temporary value.

The extension determines if the current printer matches the printer specified in the inner element configuration by comparing the temporary path value to the current printer's path value, comparing the temporary name value to the current printer's path value, and comparing the current printer's port name to the current printer's path value.

If none of these comparisons indicate a successful match, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the extension

completes the iteration of printers without finding a matching printer, then the extension indicates a matching printer does not exist and proceeds to the Detect default printer override phase.

If at least one of the comparisons indicates a successful match, then the extension stops iterating through the list of printers and indicates a matching printer exists. The extension writes “FindSharedPrinterEx: “ and the <path> attribute value to the trace log, and proceeds to the Detect default printer override phase.

Delete shared printer connection

The extension determines if it previously discovered a shared printer matching the <path> attribute value. If the extension did not discover a shared printer with a matching path name then it proceeds to the Delete printer mappings phase.

The extension deletes the shared printer connection by passing the <path> attribute value to the DeletePrinterConnection Windows API. An unsuccessful response indicating an invalid printer name causes the extension to abort its processing and return an unexpected error. All other unsuccessful responses cause the extension to abort its processing and return the underlying error it received from the API. A successful response from the API cause the extension to proceed to the Delete printer mapping phase.

Delete printer mappings

The extension determines if it should delete all printer mappings or a specific mapped printer by evaluating the <deleteMaps> attribute value. An attribute value equaling true (one [1]) causes the extension to proceed to the Delete all printer mappings phase. An attribute value equaling false causes the extension to proceed to the Delete a specific mapped printer.

Delete all printer mappings

The extension begins by using the EnumPorts Windows API to gather a list of ports available for printer for the computer or current user. An unsuccessful response from the API causes proceed to the Epilogue phase. A successful response from the API also returns a list of ports for the computer or current user. The extension begins to iterate through the list of ports by making the first port in the list the current port.

Detect LPT port

The extension looks for the string “LPT’ in the current port’s name. If the current port does not include the string “LPT” in its name, then the extension moves to the next port in the list and repeats the evaluation with the new current port. If the current port does include the string “LPT” in its name, then the extension proceeds to the Detect mapped port phase.

Detect and delete mapped port

Next, the extension determines if the port matching the <port> attribute value is currently mapped. The extension determines this by passing the <port> attribute value to the WNetGetConnection Windows API.

An unsuccessful response from the API also returns an empty internal remote path value, which indicates the port is not mapped. This result causes the extension to makes the next port in the list ports the current port and return to the Detect LPT port phases.

A successful response from the API also returns an internal remote path value in the form of a UNC path to which the port is mapped, which indicates the port is mapped. This result causes the extension to pass the current port to the WNetCancelConnection2 Windows API, which removes the mapping from the current port.

All responses from the API cause the extension to make the next port in the list ports the current port and return to the Detect LPT port phase. The extension proceeds to the Epilogue phases if it completely iterates through the list of ports in the list.

Delete a specific mapped printer

The extension determines confirms the <port> attribute value is present in the configuration. If the <port> attribute value is empty, then the extension proceeds to the Epilogue phase; otherwise it proceeds to the Detect and delete mapped port phase.

Detect and delete mapped port

Next, the extension determines if the port matching the <port> attribute value is currently mapped. The extension determines this by passing the <port> attribute value to the WNetGetConnection Windows API.

An unsuccessful response from the API also returns an empty internal remote path value, which indicates the port is not mapped. This result causes the extension to write “IsMapped: “along with the <port> attribute value to the trace log and proceed to the Epilogue phase.

A successful response from the API also returns an internal remote path value in the form of a UNC path to which the port is mapped, which indicates the port is mapped. This result causes the extension to pass the current port to the WNetCancelConnection2 Windows API, which removes the mapping from the current port.

All responses from the API cause the extension to write “UnMapPort: “along with the <port> attribute value to the trace log and proceed to the Epilogue phase.

Epilogue

The extension completes the processing by securely deleting the clear text instance of the <cPassword> attribute value.

Delete action processing is complete. The extension moves to the next inner element (preference item

Replace

The replace action represents a configuration where the extension replaces an existing shared printer connection with a newly created shared printer connection based on the configuration in the current inner element.

It is important to understand the differences between replace and update. The replace action typically deletes an existing shared printer connection and creates a new one—essentially changing all configuration elements of the printer connection. The update action only updates a portion of the configuration elements for an existing shared printer connection.

Prologue

The extension begins processing **SharedPrinter** inner elements by extracting the values of the <comment>, <location>, <path>, <port>, <username>, and <cPassword> attributes from the current inner element.

The extension reads the next set of attributes from the current inner element and assign each attribute a default value of false (zero [0]) if a value is not present in the inner element. These attributes include <deleteAll>, <default>, <skipLocal>, <persistent>, and <deleteMaps>.

If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extensions prepares for the core processing phase by setting the security context to the computer or current user (depending how the preference item is deployed).

Core

Delete a specific shared printer connection

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension indicate a matching printer does not exists and proceeds with the Detect default printer override phase.

A successful result from the API causes the extension to cycle through the list of returned local printers. The extension begins the cycle with the current printer and determines if the printer is a network printer by the PRINTER_ATTRIBUTE_NETWORK bit in the current printer's attributes. If the PRINTER_ATTRIBUTE_NETWORK bit is not enabled, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the PRINTER_ATTRIBUTE_NETWORK bit is disabled, then the extension attempts to discover a matching printer.

Match printer by share name

The extension detects a matched printer when it identifies a local printer sharing the same share name as the <path> attribute value. To begin, the extension uses the GetComputerName Windows API to retrieve and store the computers name. The extension begins a series of comparisons starting with the current printer' server name.

If the current printer's server name is empty, then the extension creates a temporary path value represented as UNC path using the recently discover computer name for the server portion and the current printer's share name value as the share portion (\\computerName\printerShareName). Additionally, the extension creates a temporary name value represented as a UNC path using the recently discovered computer name for the server portion and the current printer's name as the share portion (\\computerName\printerName).

If the current printer's server name contains a value, then the extension creates a temporary path value represented as a UNC path using the current printer's server name as the server portion and the current printer's share name as the share portion (\\currentPrinterServerName\currentPrinterShareName). Additionally, the extension creates a temporary name value and assigns the current printer's name to the temporary value.

The extension determines if the current printer matches the printer specified in the inner element configuration by comparing the temporary path value to the current printer's path value,

comparing the temporary name value to the current printer's path value, and comparing the current printer's port name to the current printer's path value.

If none of these comparisons indicate a successful match, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the extension completes the iteration of printers without finding a matching printer, then the extension indicates a matching printer does not exist and proceeds to the Detect default printer override phase.

If at least one of the comparisons indicates a successful match, then the extension stops iterating through the list of printers and indicates a matching printer exists. The extension writes "FindSharedPrinterEx: " and the <path> attribute value to the trace log, and proceeds to the Detect default printer override phase.

Delete shared printer connection

The extension determines if it previously discovered a shared printer matching the <path> attribute value. If the extension did not discover a shared printer with a matching path name then it proceeds to the Delete printer mappings phase.

The extension deletes the shared printer connection by passing the <path> attribute value to the DeletePrinterConnection Windows API. An unsuccessful response indicating an invalid printer name causes the extension to abort its processing and return an unexpected error. All other unsuccessful responses cause the extension to abort its processing and return the underlying error it received from the API. A successful response from the API cause the extension to proceed to the Delete printer mapping phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create printer phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to proceed to the Create printer phase.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]), write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

Create printer

The extension internally references if it previously detected the local printer or if it previously deleted the local printer. If the extension did not previously delete or previously detect a local printer, then the extension proceeds to the Set default printer phase. If the extension discovers it previously deleted or detected a local printer then the extension proceeds with the Driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing, write “InstallDriversEx:” and the <path> attribute value to the trace log, and return the underlying error it received from the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install the driver, provided the driver is not already installed.

If the driver needed for the printer installation is not located on the remote location specified by the <path> attribute value, then the extension closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx:” and the <path> attribute value to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection, write “InstallDriversEx: “and the <path> attribute value to the trace log, and proceed with the Add shared printer phase.

Add shared printer

The extension adds the shared printer to the computer or current user by passing the <path> attribute value to the AddPrinterConnection Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error code I received from the API.

A successful response from the API causes the extension write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Set default printer phase.

Set default printer

To set the default printer, the extension internally references if it previously discovered an existing printer. If the extension previously discovered a printer, then the extension proceeds to the Map printer port phase. If the extension did not previously discover a printer, then it checks the internal <default> attribute to determine if it was modified during the Detect default printer override phase.

If the extension determined the internal <default> attribute value equals false (zero [0]), then it proceeds to the map printer port phase. If the extension determines the internal <default> attribute value equals true (one [1]), then it passes the <path> attribute value to the SetDefaultPrinter Windows API.

An unsuccessful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error provided by the API. A successful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Map printer port phase.

Map printer port

The extension begins by verifying that the <port> attribute value is not empty. If the <port> attribute value is empty, then the extension proceeds to the Epilogue phase. Otherwise, the extension proceeds with the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by retrieving information about the remote printer. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the API causes

the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and returning the underlying response code it received from the GetPrinter API. The extension creates a temporary internal string value by combining the printers share path and share name in a UNC syntax.

Detect mapped port

Next, the extension determines if the port matching the <port> attribute value is currently mapped to another printer. The extension determines this by passing the <port> attribute value to the WNetGetConnection Windows API.

An unsuccessful response from the API also returns an empty internal remote path value, which indicates the port is not mapped. The extension writes “IsMapped: “along with the <port> attribute value, and the blank internal remote path value to the trace log. The extension proceeds to the Map port phase.

A successful response from the API also returns an internal remote path value in the form of a UNC path to which the port is mapped. The extension writes “IsMapped: “along with the <port> attribute value and the internal remote path value to which the port is mapped to the trace log.

The extension then performs a case insensitive comparison between the <path> attribute value and the internal remote path value. If the comparison results in a match, then the extension proceeds to the Epilogue phase.

If the comparisons results conclude that the two string values do not match, then the extension writes the <port> attribute value “already mapped to “ the remote path value to the trace log, sets the error code to already exists, and proceeds to the Epilogue phase.

Map port

The extension maps the local port matching the <port> attribute value to the <path> attribute value by passing the <port>, <path>, <persistent> attribute values to the WNetUseConnection Windows API. The extension also passes the <username> and <cPassword> attribute values to the API provided these values are not empty.

Regardless of the response, the extension writes “MapPort: “along with the <port> and <shared> attribute values to the trace log. The extension also sets the error code to the underlying error code it received from the WNetUseConnection API. The extension proceeds to the Epilogue phase.

Epilogue

The extension completes the processing by securely deleting the clear text instance of the <cPassword> attribute value.

Replace action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Update

The update action represents a configuration where the extension refreshes an existing shared printer connection based on the configuration in the current inner element. It is important to understand the difference between replace and update.

It is important to understand the differences between replace and update. The replace action typically deletes an existing shared printer connection and creates a new one—essentially changing all configuration elements of the printer connection. The update action only updates a portion of the configuration elements for an existing shared printer connection.

Prologue

The extension begins processing SharedPrinter inner elements by extracting the values of the <comment>, <location>, <path>, <port>, <username>, and <cPassword> attributes from the current inner element.

The extension reads the next set of attributes from the current inner element and assign each attribute a default value of false (zero [0]) if a value is not present in the inner element. These attributes include <deleteAll>, <default>, <skipLocal>, <persistent>, and <deleteMaps>.

If the <cPassword> attribute contains a value, then the extension reverse-obfuscates the password into its clear text equivalent. The extensions prepares for the core processing phase by setting the security context to the computer or current user (depending how the preference item is deployed).

Core

Delete a specific shared printer connection

Inventory local printers

The extension begins core processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension indicate a matching printer does not exists and proceeds with the Detect default printer override phase.

A successful result from the API causes the extension to cycle through the list of returned local printers. The extension begins the cycle with the current printer and determines if the printer is a network printer by the PRINTER_ATTRIBUTE_NETWORK bit in the current printer's attributes. If the PRINTER_ATTRIBUTE_NETWORK bit is not enabled, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the PRINTER_ATTRIBUTE_NETWORK bit is disabled, then the extension attempts to discover a matching printer.

Match printer by share name

The extension detects a matched printer when it identifies a local printer sharing the same share name as the <path> attribute value. To begin, the extension uses the GetComputerName

Windows API to retrieve and store the computers name. The extension begins a series of comparisons starting with the current printer' server name.

If the current printer's server name is empty, then the extension creates a temporary path value represented as UNC path using the recently discover computer name for the server portion and the current printer's share name value as the share portion (\\computerName\printerShareName). Additionally, the extension creates a temporary name value represented as a UNC path using the recently discovered computer name for the server portion and the current printer's name as the share portion (\\computerName\printerName).

If the current printer's server name contains a value, then the extension creates a temporary path value represented as a UNC path using the current printer's server name as the server portion and the current printer's share name as the share portion (\\currentPrinterServerName\currentPrinterShareName). Additionally, the extension creates a temporary name value and assigns the current printer's name to the temporary value.

The extension determines if the current printer matches the printer specified in the inner element configuration by comparing the temporary path value to the current printer's path value, comparing the temporary name value to the current printer's path value, and comparing the current printer's port name to the current printer's path value.

If none of these comparisons indicate a successful match, then the extension makes the next printer in the list of printers the current printer and starts a new iteration. If the extension completes the iteration of printers without finding a matching printer, then the extension indicates a matching printer does not exist and proceeds to the Detect default printer override phase.

If at least one of the comparisons indicates a successful match, then the extension stops iterating through the list of printers and indicates a matching printer exists. The extension writes "FindSharedPrinterEx: "and the <path> attribute value to the trace log, and proceeds to the Detect default printer override phase.

Detect default printer override

The extension detects the default printer override option by validating both the <default> and <skipLocal> attribute values equal true (one [1]). If either attribute equals false, the extension proceeds to the Create printer phase.

When both attributes equal true, the extension continues processing by using the Windows API EnumPrinters to retrieve a list of locally installed printers. A failed result from the API causes the extension to proceed to the Create printer phase.

A successful result from the API causes the extension to cycle through the list of returned local printers where it performs multiple evaluations. The extension first validates the current printer from the list of local printers is not a network connected printer by determining if the PRINTER_ATTRIBUTE_NETWORK bit is enabled in the current printers attributes. If the current printer is a network printer, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. If the current printer is not a network printer, then the extension proceeds with validating the printer name.

The second evaluation performed on the current printer is to guarantee the current printer does not match the previously discovered printer name. If the current printer name matches the

previously discovered printer name, then the extension skips the remaining evaluations and cycles to the next printer in the list of local printers. Otherwise, the current printer name and the previously discovered printer name do not match and the extension continues with its final evaluation.

The final evaluation performed by the extension is to determine if the current printer port is configured with a local port. The extension checks if the port name of the current printer from the list of local printers contains “LPT”, “COM”, or “USB” in the name. If the current printer’s port name does not contain these strings, then the extension cycles to the next printer in the list of local printers.

If the extensions iterates through all the printers and does not discover a printer in the list of local printers that passes these evaluations, then it internally reports that a local printer is not present. If the extension identifies a printer from the list of local printers passing all three evaluations, then it will cease the iteration and internally report a local printer exists.

A local printer not existing for the computer or current user with both the <default> and <skipLocal> attributes equaling true causes the extension to proceed to write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

A local printer existing for the computer or current user with both the <default>< and <skipLocal> attributes equaling true causes the extension to internally change the <default> attribute to equal false (zero [0]), write “IsOtherLocalPrinterEx” and the last observed error code to the trace log and proceed to the Create printer phase.

Create printer

The extension internally references if it previously detected the local printer. If the extension did previously detect a local printer, then the extension proceeds to the Set default printer phase. If the extension discovers it did not previously detected a local printer then the extension proceeds with the Driver installation phase.

Driver installation

The extension begins the driver installation phase by using the Windows API OpenPrinter using the <path> attribute value as the shared printer that hosts the printer driver for the newly created local printer. An unsuccessful response from the Windows API causes the extension to abort its processing, write “InstallDriversEx:” and the <path> attribute value to the trace log, and return the underlying error it received from the API.

A successful response from the OpenPrinter API causes the extension to use the GetPrinterDriver Windows API to retrieve the printer driver from the <path> attribute value and install the driver, provided the driver is not already installed.

If the driver needed for the printer installation is not located on the remote location specified by the <path> attribute value, then the extension closes the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx:” and the <path> attribute value to the trace log, and returns an unknown printer driver error.

A failed response from the GetPrinterDriver API indicating a problem with allocating memory causes the extension to close the printer connection using the ClosePrinter API. The extension

aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns an out of memory error.

A failed response from the GetPrinterDriver API causes the extension to close the printer connection using the ClosePrinter Windows API. The extension aborts its processing, writes “InstallDriversEx: “and the <path> attribute value to the trace log, and returns the underlying error it received from the Windows API.

A successful response from the API causes the extension to close the printer connection, write “InstallDriversEx: “and the <path> attribute value to the trace log, and proceed with the Add shared printer phase.

Add shared printer

The extension adds the shared printer to the computer or current user by passing the <path> attribute value to the AddPrinterConnection Windows API. An unsuccessful response from the API causes the extension to abort its processing, write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error code I received from the API.

A successful response from the API causes the extension write “AddSharedPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Set default printer phase.

Set default printer

To set the default printer, the extension internally references if it previously discovered an existing printer. If the extension previously discovered a printer, then the extension proceeds to the Map printer port phase. If the extension did not previously discover a printer, then it checks the internal <default> attribute to determine if it was modified during the Detect default printer override phase.

If the extension determined the internal <default> attribute value equals false (zero [0]), then it proceeds to the map printer port phase. If the extension determines the internal <default> attribute value equals true (one [1]), then it passes the <path> attribute value to the SetDefaultPrinter Windows API.

An unsuccessful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and return the underlying error provided by the API. A successful response from the API causes the extension to abort its processing, write “SetDefaultPrinterEx: “and the <path> attribute value to the trace log, and proceed to the Map printer port phase.

Map printer port

The extension begins by verifying that the <port> attribute value is not empty. If the <port> attribute value is empty, then the extension proceeds to the Epilogue phase. Otherwise, the extension proceeds with the Get remote printer settings phase.

Get remote printer settings

The extension begins this phase by retrieving information about the remote printer. The extension gets this information by passing the <path> attribute value to the OpenPrinter Windows API. An unsuccessful response from the API causes the extension to abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return the underlying error it received from the API.

A successful response causes the extension to pass the printer handle it received from the OpenPrinter API to the GetPrinter Windows API. An unsuccessful response from the API causes the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error.

Any errors regarding memory allocation occurring while gather remote printer information causes the extension to close the printer handle using the ClosePrinter API; abort its processing; write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and return an unexpected error. A successful response from the API causes the extension to retrieve additional information about the remote printer. The additional information retrieved from the remote printer includes: printer name, driver name, server name, share name, printer processor, location, and comments.

Upon successfully retrieving the remote printer information, the extension closes the connection to the remote printer using the ClosePrinter API, write “GetRemotePrinterSettingsEx: ”along with printer path, printer share path, and printer share name to the trace log; and returning the underlying response code it received from the GetPrinter API. The extension creates a temporary internal string value by combining the printers share path and share name in a UNC syntax.

Detect mapped port

Next, the extension determines if the port matching the <port> attribute value is currently mapped to another printer. The extension determines this by passing the <port> attribute value to the WNetGetConnection Windows API.

An unsuccessful response from the API also returns an empty internal remote path value, which indicates the port is not mapped. The extension writes “IsMapped: “along with the <port> attribute value, and the blank internal remote path value to the trace log. The extension proceeds to the Map port phase.

A successful response from the API also returns an internal remote path value in the form of a UNC path to which the port is mapped. The extension writes “IsMapped: “along with the <port> attribute value and the internal remote path value to which the port is mapped to the trace log.

The extension then performs a case insensitive comparison between the <path> attribute value and the internal remote path value. If the comparison results in a match, then the extension proceeds to the Epilogue phase.

If the comparisons results conclude that the two string values do not match, then the extension writes the <port> attribute value “already mapped to “ the remote path value to the trace log, sets the error code to already exists, and proceeds to the Epilogue phase.

Map port

The extension maps the local port matching the <port> attribute value to the <path> attribute value by passing the <port>, <path>, <persistent> attribute values to the WNetUseConnection Windows API. The extension also passes the <username> and <cPassword> attribute values to the API provided these values are not empty.

Regardless of the response, the extension writes “MapPort: “along with the <port> and <shared> attribute values to the trace log. The extension also sets the error code to the underlying error

code it received from the WNetUseConnection API. The extension proceeds to the Epilogue phase.

Epilogue

The extension completes the processing by securely deleting the clear text instance of the <cPassword> attribute value.

Update action processing is complete. The extension moves to the next inner element (preference item) and begins inner element processing on the new, current inner element.

Dependencies

The Group Policy Printers extension is one of many extensions provided in a single dynamic linked library, gpprecl.dll. This extension has one or more dependencies on the following dynamic linked libraries that are included in the Windows operating system.

1. ACTIVEDS.DLL
2. ADVAPI32.DLL
3. IPHLPAPI.DLL
4. KERNEL32.DLL
5. MPR.DLL
6. MSI.DLL
7. MSVCRT.DLL
8. NETAPI32.DLL
9. NTDLL.DLL
10. OLE32.DLL
11. OLEAUT32.DLL
12. POWRPROF.DLL
13. RPCRT4.DLL
14. SECUR32.DLL
15. SETUPAPI.DLL
16. SHELL32.DLL
17. SHLWAPI.DLL
18. USER32.DLL
19. USERENV.DLL
20. VERSION.DLL
21. WINSPOOL.DRV
22. WINSTA.DLL
23. WLDAP32.DLL
24. WS2_32.DLL
25. WTSAPI32.DLL
26. XMLLITE.DLL

Interaction with other components

The Windows operating system enables many different configurations to accomplish a single task. Creating and managing printers is one of those tasks.

The most common configuration for mapping printers occurs during the user logon. Historically, IT departments have accomplished this configuration using network logon scripts—scripts that process as part of the user logon. Users can map printers using Windows Explorer and make them persistent between computer reboots. You can use Deployed Printers, which is an alternative Group Policy-based technology that stores configuration information in Active Directory. Lastly, you can use the Group Policy Drive Map extension to map drives during logon.

Flexible configurations are important; however, understanding that incorporating the one or more of the configurations into one solution needs to be thoroughly tested for each specific environment as many of these technologies can encroach on the other thereby creating unpredictable results. A suggested best practice is to pick one configuration and use it exclusively in your environment, especially when managing printers during logon. Incorporating multiple configurations could be counterproductive and cause conflict between the configurations. This conflict is likely to produce unwanted results, increases the complexity of the solution, increases administrative overhead of the solution, and increase the difficulty of troubleshooting the cause of the problem.

Logging

In addition to the event messages recorded in the application log, the Group Policy Printers extension has trace logging to help diagnose preference items that are not applying or undesired or unexpected results from preference items that are applying.

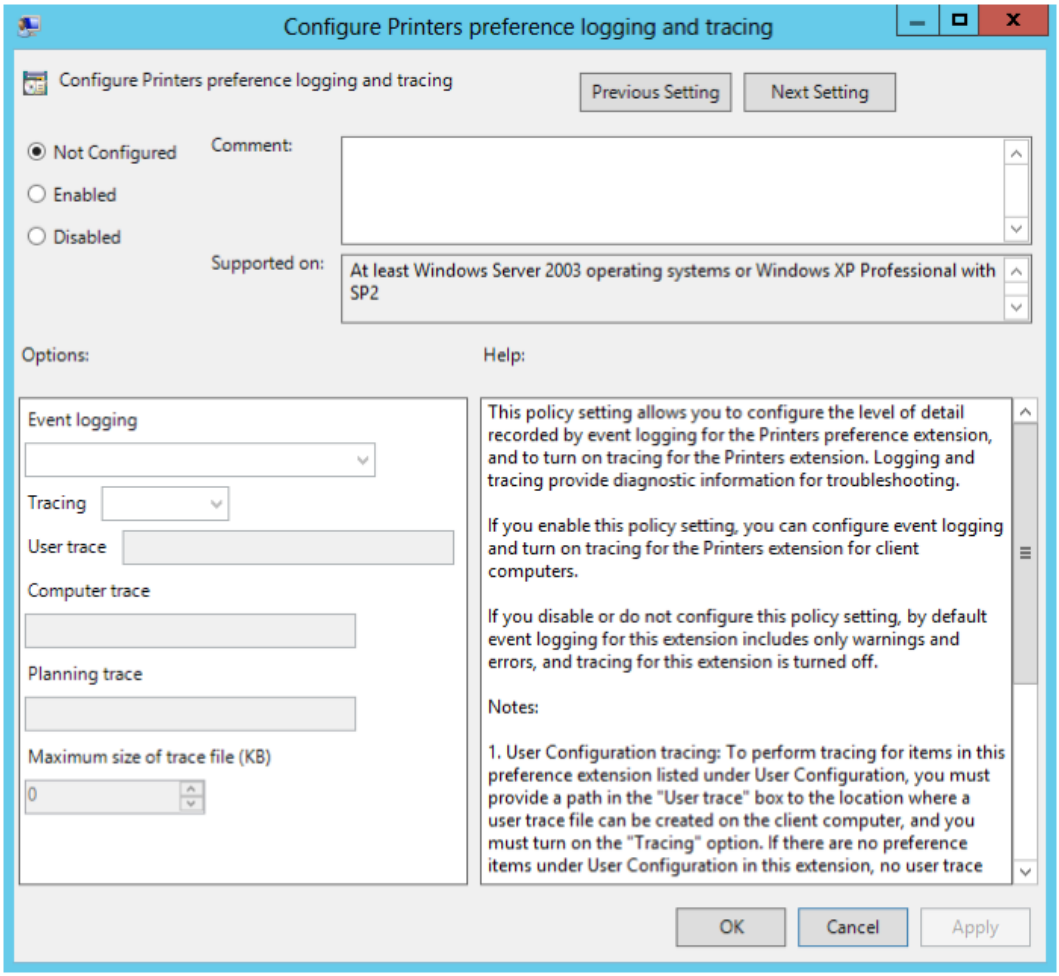

You enable Group Policy Preference logging using Computer\Policies namespace of the Group Policy Management Editor. The **Configure Printers preference logging and tracing** policy setting resides under the **Administrative Templates\System\Group Policy\Logging and tracing node**. The policy setting enables a variety of configuration options that affect the types of events the extension reports to the event log, trace logging, and the filename and location of the trace logs.

The best configuration for capturing Group Policy Printers information for troubleshooting is to configure **Event Logging** using the **Informational, Warnings, and Errors** option and configure **Tracing** to **On**.

Trace logging file name and locations are configurable to any name and location to where the user or computer has read and write permissions. The default paths for these files use a Group Policy Preference variable **%COMMONAPPDATA%\GroupPolicy\Preference\Trace\<filename>**. The Group Policy Preference %COMMONAPPDATA% variable equates to the All Users' application data folder.

Starting with Windows Vista, the All Users' folder was deprecated with the introduction of the version 2 user profile structure. To ensure compatibility among applications, Windows creates a

symbolic link named All Users in the file system under the %SYSTEMDRIVE%\Users folder. The symbolic link redirects file system requests to the All Users folder to the **C:\ProgramData** folder. Therefore, the default location for Group Policy Preference trace logs is **C:\ProgramData\GroupPolicy\Preference\Trace\<filename>.**

Group Policy Analysis and Troubleshooting Overview

This topic provides high-level information about concepts and techniques that are used to troubleshoot Group Policy settings, in addition to links to detailed reference topics, procedures, and troubleshooting scenario guides.

Did you mean…

- Troubleshooting Group Policy by Using Event Logs for Vista and Windows Server 2008
- Group Policy Infrastructure Events and Errors for Windows Server 2008

Group Policy troubleshooting description

Understanding how to troubleshoot Group Policy is important for maintaining corporate standards. IT professionals depend on the reliability of Group Policy to keep networks secure and managed, and to lower operating costs. When an IT pro is trying to understand why a Group Policy setting did not apply to a computer or user in an expected manner, understanding where to locate and how to read and analyze information can make the difference between a small network issue and hours of downtime.

Troubleshooting Group Policy becomes complicated because Group Policy is not just a single component or application that resides locally on a single computer. Group Policy is a processing infrastructure that is used to deliver and apply one or more desired configurations or policy settings to a set of targeted users and computers within an Active Directory environment. This processing infrastructure consists of a Group Policy engine and multiple client-side extensions (CSEs) that are responsible for reading specific policy settings on target client computers. During Group Policy processing, Group Policy must access domain-based information about Group Policy Objects (GPOs) that is stored in Active Directory, the Group Policy container, and SYSVOL (which provides a standard location to store the Group Policy template).

Practical considerations

The Group Policy Management Console (GPMC) provides a starting point to approach troubleshooting Group Policy. Functionality from number of separate tools, such as GPUpdate.exe and GPOTool.exe, were rolled into the GPMC in Windows Server 2012.

A methodology for using the GPMC to troubleshoot Group Policy considers the following factors:

- Ensure Group Policy has processed the most current set of computer and user settings.

    - Force a Remote Group Policy Refresh (GPUpdate).
    - Check that Group Policy Object information has replicated to all domain controllers. See Check Group Policy Infrastructure Status.
    - Check for Group Policy errors and events.
- Investigate configurations that affect the Resultant Set of Policy.
    - Slow link
    - Loopback
    - Security filtering
    - Block inheritance
    - Enforced GPO
    - Disable the entire GPO, or the computer or user portion of the GPO.
- Understand core Group Policy processing.
    - Understand how changes to the Windows Server 2012 and Windows 8 startup process impact Group Policy processing.
    - Understand how migrating Group Policy to a separate service affects core Group Policy processing.
    - Check for Group Policy errors and events by using the GPMC reports.
- Understand Group Policy client-side extension processing.
    - Understand the expected behavior of the Group Policy preference extensions.
    - Determine which GPO contains a setting that has been applied to a computer or user.

See also

The following references provide troubleshooting information for components that affect Group Policy functionality.

Troubleshooting considerations	References
Active Directory replication information	• Introduction to Active Directory Replication and Topology Management Using Windows PowerShell (Level 100) • Troubleshooting replication with repadmin
DFRS and FRS replication information	1. DFS Replication: Frequently Asked Questions (FAQ) 2. What's New in Distributed File System 3. DFS Replication

Check Group Policy Infrastructure Status

Group Policy is a complicated infrastructure that allows you to apply policy settings to remotely configure computer and user experiences within a domain. Most of the troubleshooting advice comes to IT administrators when Group Policy fails to perform in the expected manner. But this advice does not scale well when you have thousands of Group Policy Objects (GPOs) that apply to thousands of computers and users.

In some cases, when you have a large enterprise spread over multiple countries and time zones, there could be a significant replication delay between domain controllers. GPO mismatches in version numbers between the Group Policy container and the Group Policy template, or GPO mismatches between different domain controllers, can be expected due to a delay in replication. They can also indicate an issue with Group Policy. In previous versions of Windows, there were tools (such as **GPOTool.exe**) to provide a view of the GPO replication, but they provided inconsistent information.

In Windows Server 2012, the Group Policy Management Console (GPMC) was enhanced to provide a report for the overall health state of the Group Policy infrastructure for a domain, or to scope the health view to a single GPO.

The GPMC domain status tab (added in Windows Server 2012) displays individual pieces of information that indicate the health of the Group Policy infrastructure with regards to domain controllers, GPO replication, and GPO versioning. This Group Policy infrastructure health status can help you find inconsistencies and anticipate issues.

In this document

We will look at what makes up a GPO and where the parts of the GPO are stored in a domain-based environment. Then we will look at how to use the Group Policy infrastructure health status information to monitor for issues with GPO replication and what to do for further investigation of replication issues.

- Prerequisites
- Step 1: Check Group Policy infrastructure health
- Step 2: Check the results of the Group Policy infrastructure status report
- Step 3: Check Active Directory replication issues
- Step 4: Check file services replication issues

Prerequisites

You can run Group Policy infrastructure health checks only by using the GPMC from domain-joined computers running:

- Windows Server 2012 or Windows Server 2012 R2
- Windows 8 or Windows 8.1 with Remote Server Administration Tools for Windows 8

An understanding of how GPOs are stored in a domain will help you make full use of the data that is displayed on the infrastructure status page.

In a domain that contains more than one domain controller, Group Policy information takes time to propagate or replicate from one domain controller to another. Low-bandwidth network connections between domain controllers slow replication. The Group Policy infrastructure has mechanisms to manage these issues.

Each GPO is stored partly in Active Directory and partly in SYSVOL on the domain controller. The portion of the GPO that is stored in Active Directory is called the Group Policy container, and the portion of the GPO that is stored in SYSVOL is called the Group Policy template. The GPMC and Group Policy Management Editor manage the GPO as a single unit. For example, when you set permissions on a GPO in the GPMC, the GPMC is actually setting permissions for objects in Active Directory and in SYSVOL. It is not recommended that you manipulate these separate objects independently outside of the GPMC and the Group Policy Management Editor. It is important to understand that these two separate features of a GPO rely on different replication mechanisms. The file system portion, Group Policy template, is replicated through Distributed File Service Replication (DFS-R) or File Replication Service (FRS), independently of the replication that is handled by Active Directory in the Group Policy container.

Temporary lack of synchronization can occur between the GPO data that is stored in Active Directory (Group Policy container) and the GPO data that is stored on SYSVOL (Group Policy template). This is due to the differences in the replication schemes used by Active Directory and DFS-R or FRS.

Group Policy provides a framework that extends Group Policy management to more features. These components are referred to as Group Policy snap-ins or extensions. For those Group Policy extensions that store data in only one data store (Active Directory or SYSVOL), this is not an issue, and Group Policy is applied as it is read. Such extensions include administrative templates, scripts, folder redirection, and most of the security settings.

For any Group Policy extension that stores data in Active Directory and SYSVOL, the extension must properly handle the possibility that the data is unsynchronized. This is also true for extensions that need multiple objects in a single store to be atomic in nature, because neither storage location handles transactions.

An example of an extension that stores data in Active Directory and SYSVOL is Software Installation. The script files are stored in SYSVOL, and the Windows Installer package definition is stored in Active Directory. If a script exists, but the corresponding Active Directory features are not present, nothing happens. If a script file is missing, but the package is known in Active Directory, application installation fails gracefully and retries on the next processing of Group Policy.

The tools used to manage Active Directory and Group Policy (such as GPMC, the Group Policy Object Editor, and Active Directory Users and Computers) communicate through domain controllers. If there are several domain controllers available, changes to objects (including users, computers, organizational units, and GPOs) may take time to appear on other domain controllers. The administrator may see different data depending on the last domain controller on which the changes were made and on which domain controller they are currently viewing the data.

Ideally, you would like to proactively check to make sure that all your GPOs within a domain are consistent. In the past, **GPOTool.exe** was used for this check. However, **GPOTool.exe** is limited

to checking the values in the GPT.ini file, which only indicates if the versions in the Group Policy container and the Group Policy template on each domain controller are in sync.

In Windows Server 2012 and Windows Server 2012 R2, the GPMC contains a more comprehensive replication check than what was previously done with **GPOTool.exe**.

Step 1: Check Group Policy infrastructure health

Windows Server 2012 includes a graphical reporting feature in the GPMC that enables you to choose a baseline domain controller for comparison. You can see the current Group Policy replication status with synchronization details when a comparison finds a differential from the baseline domain controller.

Important

You can only run the Group Policy infrastructure status report from the GPMC. You cannot schedule this report, and there is no Windows PowerShell or other command-line equivalent.

The report does not persist between GPMC sessions. Do not close the GPMC until the report gathering process has completed to avoid losing the report information.

To create and analyze an infrastructure status report

1. 0.1. Run an infrastructure status report for a domain or for a GPO:
   a. .a. **For an entire domain** In the GPMC console tree, click the name of the domain for which you want to check the replication status of all the GPOs.
   b. **For a single GPO** In the GPMC console tree, navigate to the Group Policy Objects container. Expand the Group Policy Objects container and click the GPO for which you want to check the replication status.
2. Click the **Status** tab in the results pane.
3. Click the **Detect Now** button to gather infrastructure status from all of the domain controllers in this domain.

Step 2: Check the results of the Group Policy infrastructure status report

To validate consistency between the baseline and each domain controller, the GPMC compares the Group Policy information that is stored in Active Directory (Group Policy container), and it separately compares the Group Policy information that is stored in SYSVOL (Group Policy template) on all the domain controllers in the selected domain.

Here is a list of comparisons that are performed:

- ACL for each Group Policy container
- Version number attribute on each Group Policy template
- Count of Group Policy container objects

- ACL on each Group Policy template
- Version number stored in the gpt.ini file in each Group Policy container
- Count of Group Policy template folders and files
- File hash for each file in Group Policy template folders

During the Group Policy infrastructure status report gathering process, if the GPMC cannot contact a domain controller, or if a domain controller is found to not be consistent with the baseline domain controller, the domain controller in question is added to the **Domain controller(s) with replication in progress** list.

To verify GPO replication consistency for all domain controllers

1. When the infrastructure status report gathering process is complete, click the **Status** tab, and in the **Status Details** section, verify that there are no domain controllers listed in the **Domain controller(s) with replication in progress** list.
2. Group Policy application requires clients to access specified servers, including domain controllers and other servers, such as those used for sharing documents and as installation points. Group Policy management also requires access to domain controllers.

   If a domain controller is listed as **Inaccessible** in the **Domain controller(s) with replication in progress** section of the **Status Details**, check the following:

   a. The domain controller is working and accessible on the network.

   b. The event logs and operation logs of the domain controller do not report any errors that would cause the server to not respond to network requests.

   c. SYSVOL and Active Directory are available from the domain controller that is listed as inaccessible.
3. If a domain controller is listed in **Domain controller(s) with replication in progress** due to inconsistencies in the Group Policy container portion of the GPO, see Step 3: Check Active Directory replication issues
4. If a domain controller is listed in **Domain controller(s) with replication in progress** due to inconsistencies in the Group Policy template portion of the GPO, see Step 4: Check file services replication issues

Step 3: Check Active Directory replication issues

Inbound or outbound replication failure can cause Active Directory objects that represent the replication topology, replication schedule, domain controllers, users, computers, passwords, security groups, group memberships, and Group Policy to be inconsistent between domain controllers.

If you suspect that you have an issue with Active Directory replication, see the following resources to understand how these mechanisms work and how you can troubleshoot issues:

- Troubleshooting replication with repadmin
- Troubleshooting Active Directory Replication Problems
- Active Directory Replication and Topology Management Using Windows PowerShell

Step 4: Check file services replication issues

The Group Policy template portion of Group Policy resides in SYSVOL. Group Policy depends on SYSVOL replication working correctly for the Group Policy template to remain consistent between domain controllers within a domain. SYSVOL replication depends on the Distributed File System Replication (DFS-R) or the File Replication Service (FSR), depending on how you have configured your domain controllers.

If you suspect that you have an issue with SYSVOL replication, see the following resources to understand how these mechanisms work and how you can troubleshoot issues:

- Administering DFS-Replicated SYSVOL
- DFS Replication: Frequently Asked Questions (FAQ)
- What's New in Distributed File System
- DFS Management Help
- How To Troubleshoot the File Replication Service in Windows Server 2003
- Troubleshooting File Replication Service
- Using the BurFlags registry key to reinitialize File Replication Service replica sets
- Ask the Directory Services Team blog contains extensive troubleshooting information for Active Directory replication, DFSR and Group Policy

See also

- Group Policy Overview
- Group Policy Analysis and Troubleshooting Overview

Force a Remote Group Policy Refresh (GPUpdate)

Group Policy is a complicated infrastructure that enables you to apply policy settings to remotely configure a computer and user experience within a domain. When the Resultant Set of Policy settings does not conform to your expectations, a best practice is to first verify that the computer or user has received the latest policy settings. In previous versions of Windows, this was accomplished by having the user run **GPUpdate.exe** on their computer.

With Windows Server 2012 and Windows 8, you can remotely refresh Group Policy settings for all computers in an organizational unit (OU) from one central location by using the Group Policy Management Console (GPMC). Or you can use the **Invoke-GPUpdate** Windows PowerShell cmdlet to refresh Group Policy for a set of computers, including computers that are not within the OU structure—for example, if the computers are located in the default computers container.

The remote Group Policy refresh updates all Group Policy settings, including security settings that are set on a group of remote computers, by using the functionality that is added to the context menu for an OU in the Group Policy Management Console (GPMC). When you select an

OU to remotely refresh the Group Policy settings on all the computers in that OU, the following operations happen:

1. An Active Directory query returns a list of all computers that belong to that OU.
2. For each computer that belongs to the selected OU, a WMI call retrieves the list of signed in users.
3. A remote scheduled task is created to run **GPUpdate.exe /force** for each signed in user and once for the computer Group Policy refresh. The task is scheduled to run with a random delay of up to 10 minutes to decrease the load on the network traffic. This random delay cannot be configured when you use the GPMC, but you can configure the random delay for the scheduled task or set the scheduled task to run immediately when you use the **Invoke-GPUpdate** cmdlet.

This document describes a method to force a remote Group Policy refresh to all computers in an OU and all OUs that are contained within the selected OU by using the GPMC. An equivalent Windows PowerShell method is also presented for each procedure.

In this document

- Prerequisites
- Step 1: Configure firewall rules on each client that will be managed with remote Group Policy refresh
- Step 2: Schedule a remote Group Policy refresh

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

You can only schedule to force a remote Group Policy update by using the GPMC from domain-joined computers that are running:

- Windows Server 2012 or Windows Server 2012 R2
- Windows 8 or Windows 8.1 with Remote Server Administration Tools for Windows 8

You can schedule a remote Group Policy refresh for any computer running:

- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2
- Windows Server 2008
- Windows 8.1
- Windows 8
- Windows 7
- Windows Vista

Step 1: Configure firewall rules on each client that will be managed with remote Group Policy refresh

To schedule a Group Policy refresh for domain-joined computers by using the GPMC or the **Invoke-GPUpdate** cmdlet, you must have firewall rules that enable inbound network traffic on the ports listed in the following table.

Server port	Type of network traffic
TCP RPC dynamic ports, Schedule (Task Scheduler service)	Remote Scheduled Tasks Management (RPC)
TCP port 135, RPCSS (Remote Procedure Call service)	Remote Scheduled Tasks Management (RPC-EPMAP)
TCP all ports, Winmgmt (Windows Management Instrumentation service)	Windows Management Instrumentation (WMI-in)

In Windows Server 2012, Group Policy added a Starter GPO called, **Group Policy Remote Update Firewall Ports**. This Starter GPO includes policy settings to configure the firewall rules that are specified in the previous table. It is a best practice to create a new GPO from this Starter GPO. Link the GPO to your domain at a higher precedence than the Default Domain GPO, and then use it to configure all the computers in the domain to enable a remote Group Policy refresh.

Do this step using Windows PowerShell

To create a GPO from the Group Policy Remote Update Firewall Ports Starter GPO and link to the domain

1. In the GPMC console tree, locate the domain for which you want to configure all the computers to enable a remote Group Policy refresh.
2. Right-click the selected domain, and click **Create a GPO in this domain, and link it here…**
3. In the **New GPO** dialog box, type the name of the new Group Policy object in the **Name** box.
4. In the **Source Starter GPO** list, select the **Group Policy Remote Update Firewall Ports** Starter GPO that you want to use to create a new Group Policy object, and click **OK**.
5. In the results pane, click the **Linked Group Policy Objects** tab.
6. Select the GPO that you just created, and click the Up arrow until the GPO is listed above the Default Domain Policy. The new GPO will have a smaller link order value than the Default Domain Policy.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Use the **New-GPO** cmdlet with the **–StarterGpoName** parameter, and then pipe the output to the **New-GPLink** cmdlet.

For example, to create a new GPO called *Configure firewall rules for remote gpupdate* by using the **Group Policy Remote Update Firewall Ports** Starter GPO, then link the new GPO to the Contoso.com domain, use the following script:

```
New-GPO –Name "Configure firewall rules for remote gpupdate" –StarterGpoName "Group Policy Remote Update Firewall Ports" | New-GPLink –target "dc=Contoso,dc=com" –LinkEnabled yes
```

For more information about the **New-GPO** cmdlet and the **New-GPLink** cmdlet, see:

- New-GPO
- New-GPLink

Step 2: Schedule a remote Group Policy refresh

You can schedule **gpupdate.exe** to run on multiple computers from the GPMC or from a Windows PowerShell session using the **Invoke-GPUpdate** cmdlet.

Do this step using Windows PowerShell

To schedule a Group Policy refresh to run on all computers in an OU by using the GPMC

1. In the GPMC console tree, locate the OU for which you want to refresh Group Policy for all computers.

   Note
   Group Policy will also be refreshed for all computers that are located in the OUs contained in the selected OU.

2. Right-click the selected OU, and click **Group Policy Update…**
3. Click **Yes** in the **Force Group Policy update** dialog box. This is the equivalent to running **GPUpdate.exe /force** from the command line.
4. The **Remote Group Policy update results** window displays only the status of scheduling a Group Policy refresh for each computer located in the selected OU and any OUs contained within the selected OU. This display does not show the success or failure of the actual Group Policy refresh for each computer.
5. Use Resultant Set of Policy to determine the success of the scheduled Group Policy refresh, Determine Resultant Set of Policy.

   Note
   You should plan a delay of up to 10 minutes to start a Group Policy refresh when you are verifying the results for each computer.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The **Invoke-GPUpdate** cmdlet allows you to schedule a remote Group Policy update for a specified computer with all the options that the **GPUpdate.exe** command-line utility provides. This allows more freedom to determine which set of computers is to be refreshed than if you schedule the refresh through the GPMC. Additionally, you have the freedom to configure the interval of time to wait before a Group Policy refresh is performed by using the **–RandomDelayInMinutes** parameter. If set to a zero (0) value, the scheduled task for the Group Policy refresh is configured to start immediately. For more information, see Invoke-GPUpdate.

You can refresh the changed Group Policy settings for the computer that you are signed in to by running the **Invoke-GPUpdate** cmdlet without including any parameters, for example:

```
Invoke-GPUpdate
```

You cannot schedule a Group Policy refresh for the **Computers** container by using the GPMC **Group Policy Update…** functionality. The **Computers** container is a default location for computer accounts. It is not implemented as an OU that can be managed by the GPMC. However, by combining the use of the Windows PowerShell cmdlet, **Get-ADComputer**, with the **Invoke-GPUpdate** cmdlet, you can schedule a remote refresh for all computers in the **Computers** container. For more information about available Windows PowerShell cmdlets for Active Directory, see AD DS Administration Cmdlets in Windows PowerShell.

First obtain the list of computers in the **Computers** container by using the **Get-ADComputer** cmdlet. Then supply the name of each computer that is returned to the **Invoke-GPUpdate** cmdlet. For example, to force a refresh of all Group Policy settings for all computers in the **Computers** container for the Contoso.com domain, use the following script:

```
Get-ADComputer –filter * -Searchbase "cn=computers, dc=Contoso,dc=com" | foreach{ Invoke-GPUpdate –computer $_.name -force}
```

You can force a Group Policy refresh for all Group Policy settings for all computers in a single OU when you combine the **Get-ADComputer** with the **Invoke-GPUpdate** cmdlet. For example, to force a refresh of all Group Policy settings for all computers in the *Accounting* OU of the Contoso.com domain, use the following script:

```
Get-ADComputer –filter * -Searchbase "ou=Accounting, dc=Contoso,dc=com" | foreach{ Invoke-GPUpdate –computer $_.name -force}
```

You can force an immediate Group Policy refresh for all Group Policy settings for all computers in a single OU when you combine the **Get-ADComputer** with the **Invoke-GPUpdate** cmdlet and set the **–-RandomDelayInMinutes** to 0. For example, to force a refresh of all Group Policy settings for all computers in the *Accounting* OU of the Contoso.com domain, use the following script:

```
Get-ADComputer –filter * -Searchbase "ou=Accounting, dc=Contoso,dc=com" | foreach{ Invoke-GPUpdate –computer $_.name –force –-RandomDelayInMinutes 0}
```

See also

- Group Policy Overview
- Group Policy Analysis and Troubleshooting Overview

Understand the Effect of Fast Logon Optimization and Fast Startup on Group Policy

This troubleshooting topic explains the interaction between the required synchronous processing for certain Group Policy client-side extensions (CSEs) and Fast Logon Optimization, which was introduced in Windows XP. Fast Logon Optimization may cause you to sign in or restart the computer multiple times before certain Group Policy settings are applied to the client computer. This topic also discusses the effect that the fast startup feature introduced in Windows 8 has on Group Policy settings that are applied during startup or shutdown.

Prerequisites

This topic applies to computers that are running Windows 8.1, Windows 8, Windows 7, Windows Vista, or Windows XP.

Group Policy settings and CSEs

Policy settings are grouped into categories, such as Administrative Templates, Security Settings, Folder Redirection, Disk Quota, Software Installation, and the Group Policy preference extensions. The settings in each category require a specific CSE to process them, and each CSE has its own rules for processing settings. For this discussion, it's important to understand that Group Policy preference extensions represent a set of client-side extensions, not a single CSE. Each Group Policy preference extension has rules to process settings.

Asynchronous and synchronous processing

Asynchronous processing refers to processes that do not depend on the outcome of other processes. Therefore, they can occur on different threads simultaneously.

Synchronous processing refers to processes that depend on each other's outcome. Therefore, synchronous processes must wait for the previous process to finish before the next process can start.

Asynchronous processing affects Group Policy in the following ways:

- **At startup** The client computer does not wait for the network to be fully initialized before sign-in is available to the user. Group Policy for the client computer processes when the network becomes available in parallel to startup and sign-in activities.

- **At sign-in** A user does not have to wait for Group Policy to finish processing before signing in.

For CSEs such as Folder Redirection, Software Installation, and Drive Maps preference extension, the outcome of Group Policy processing might adversely affect the user's experience, for example, a program being uninstalled while the user is working with it. To keep this from happening, the CSE is designed to require synchronous processing to apply the new settings. During asynchronous processing, the CSE signals the system to indicate that a synchronous application of Group Policy is required. A synchronous application of Group Policy occurs at the next startup (if it is signaled during the computer Group Policy refresh) or at the user's next sign-in (if it is signaled during the user Group Policy refresh).

Foreground and background processing

Group Policy foreground processing applies when the computer starts or shuts down and when the user signs in or signs out. During foreground processing, policy settings can be applied asynchronously or synchronously.

Group Policy background processing applies during periodic refreshes after the computer has started or a user has signed in. All requested Group Policy refreshes that are performed by using **GPUpdate.exe** also run as background processing. During background processing, policy settings are only applied asynchronously.

CSE processing requirements

When Group Policy processes on a Windows-based computer, client-side extensions interpret the stored policy settings and make the appropriate changes to the environment. When troubleshooting a given client-side extension's application of Group Policy, the administrator can view the configuration parameters that affect that extension's operation. These parameters are in the form of registry values.

Important

This section provides an overview about the information that is stored in the registry. This is for informational purposes, and it is not recommended that you adjust the CSE processing properties by using the registry. Serious problems might occur if you modify the registry incorrectly. If you do make any changes to the registry, it is recommended that you back up the registry before you modify it. Then, you can restore the registry if a problem occurs. For more information about how to back up and restore the registry, see article 322756 in the Microsoft Knowledge Base.

Each client-side extension is identified by a subkey under the following path:

HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Winlogon\GPExtensions

The subkey name takes the form of a Globally Unique Identifier (GUID). For more information about identifying registered Group Policy client-side extensions, see article 216357 in the Microsoft Knowledge Base.

The **NoBackgroundPolicy** registry key determines if a CSE is called during Group Policy background policy processing. If this value is set to 1, the CSE will not be called during background processing. A CSE that requires synchronous processing can still be called during background processing. However, because background processing is always performed in asynchronous mode, a CSE that requires synchronous processing exits after signaling a request for the next sign-in to be run in synchronous mode.

Each CSE determines if it requires synchronous processing to apply changes to the policy settings.

The following table provides information about the synchronous processing requirements and the default background processing behavior of CSEs in Group Policy.

CSE name	Requires synchronous processing	Called during background processing
Disk Quota	No	Yes
EFS Recovery	No	Yes
Folder Redirection	Yes	Yes
Internet Explorer Maintenance	No	Yes
IP security	No	Yes
Registry	No	Yes
Scripts	• Startup scripts run synchronously by default. • Sign-in scripts run asynchronously by default.	Yes, but scripts are run only in foreground processing.
Security	No	Yes
Software Installation	Yes	Yes
Wired Networking	No	Yes
Wireless Networking	No	Yes
Applications preference extension	No	Yes
Data Sources preference extension	No	Yes
Devices preference extension	No	Yes
Drive Maps preference extension	Yes	No
Environment preference	No	Yes

CSE name	Requires synchronous processing	Called during background processing
extension		
Files preference extension	No	Yes
Folder Options preference extension	No	Yes
Folders preference extension	No	Yes
Ini Files preference extension	No	Yes
Internet Settings preference extension	No	Yes
Local Users and Groups preference extension	No	Yes
Network Options preference extension	No	Yes
Network Shares preference extension	No	Yes
Power Options preference extension	No	Yes
Printers preference extension	No	Yes
Regional Options preference extension	No	Yes
Registry preference extension	No	Yes
Scheduled Tasks preference extension	No	Yes
Services preference extension	No	Yes
Shortcuts preference extension	No	Yes
Start Menu preference extension	No	Yes

Although the requirement to run a client-side extension in synchronous mode is not configurable, other default behavior for each CSE can be configured through policy settings. To configure the properties that are associated with a CSE, apply the Administrative Templates computer configuration policy settings that are created for this purpose. The following table provides the

names of the policy settings and the CSE properties that can be modified. If the CSE property includes an X, it can be modified for the specific CSE.

Policy Setting Name	Allows processing across a slow network connection	Do not apply during periodic background processing	Process even if the Group Policy Objects have not changed	Allows background processing priority
Configure disk quota policy processing	X	X	X	
Configure EFS recovery policy processing	X	X	X	
Configure Folder Redirection policy processing	X		X	
Configure Internet Explorer Maintenance policy processing	X	X	X	
Configure IP security policy processing	X	X	X	
Configure Registry policy processing		X	X	
Configure Scripts policy processing	X	X	X	
Configure Security policy processing		X	X	
Configure Software Installation policy processing	X		X	
Configure Wired Networking policy processing	X	X	X	
Configure Wireless Networking policy	X	X	X	

Policy Setting Name	Allows processing across a slow network connection	Do not apply during periodic background processing	Process even if the Group Policy Objects have not changed	Allows background processing priority
processing				
Configure Applications preference extension policy processing	X	X	X	X
Configure Data Sources preference extension policy processing	X	X	X	X
Configure Devices preference extension policy processing	X	X	X	X
Configure Drive Maps preference extension policy processing	X		X	X
Configure Environment preference extension policy processing	X	X	X	X
Configure Files preference extension policy processing	X	X	X	X
Configure Folder Options preference extension policy processing	X	X	X	X
Configure Folders preference	X	X	X	X

Policy Setting Name	Allows processing across a slow network connection	Do not apply during periodic background processing	Process even if the Group Policy Objects have not changed	Allows background processing priority
extension policy processing				
Configure Ini Files preference extension policy processing	X	X	X	X
Configure Internet Settings preference extension policy processing	X	X	X	X
Configure Local Users and Groups preference extension policy processing	X	X	X	X
Configure Network Options preference extension policy processing	X	X	X	X
Configure Network Shares preference extension policy processing	X	X	X	X
Configure Power Options preference extension policy processing	X	X	X	X
Configure Printers preference extension policy processing	X	X	X	X
Configure Regional Options preference	X	X	X	X

Policy Setting Name	Allows processing across a slow network connection	Do not apply during periodic background processing	Process even if the Group Policy Objects have not changed	Allows background processing priority
extension policy processing				
Configure Registry preference extension policy processing	X	X	X	X
Configure Scheduled Tasks preference extension policy processing	X	X	X	X
Configure Services preference extension policy processing	X	X	X	X
Configure Shortcuts preference extension policy processing	X	X	X	X
Configure Start Menu preference extension policy processing	X	X	X	X

Fast Logon Optimization and Group Policy processing

By default in Windows 8.1, Windows 8, Windows 7, Windows Vista, and Windows XP, the Fast Logon Optimization feature is set for domain and workgroup members. Policy settings apply asynchronously when the computer starts and when the user signs in. As a result, these operating systems do not wait for the network to be fully initialized at startup and sign-in. Existing users are signed in by using cached credentials. This results in shorter sign-in times. Group Policy is applied after the network becomes available.

Fast Logon Optimization is always off during sign-in when a user:

- First signs in to a computer.
- Has a roaming user profile or a home directory for sign-in purposes.
- Has synchronous sign-in scripts.

Note
Under the preceding conditions, computer startup can still be asynchronous. However, because sign-in is synchronous under these conditions, sign-in does not exhibit optimization.

How Fast Logon Optimization affects CSE processing

With Fast Logon Optimization, enabled Group Policy foreground processing runs asynchronously. When a CSE requires synchronous processing:

- Policy settings do not apply during the asynchronous processing cycle.
- The CSE that requires synchronous policy application can still be called during asynchronous policy application. Instead of applying policy settings, the CSE signals for synchronous processing to be applied at the next startup or sign-in.

Fast Logon Optimization and required synchronous policy application

When synchronous policy application is required and Fast Logon Optimization is enabled, it can take up to two reboots or two sign-ins for the policy settings to apply. This is demonstrated in the following examples.

- Example 1: Fast Logon Optimization with synchronous processing when the user is not signed in. A user is not signed in to the client computer when a change is made to user policy settings, and synchronous processing is not requested. In this case, the policy settings are applied after two sign-ins.
- Example 2: Fast Logon Optimization with synchronous processing when the user is signed in. A user is signed in to the client computer when a change is made to user policy settings, and the user stays signed in to the computer so that background processing takes place prior to sign-out. In this case, the policy settings are applied after one sign-in.

Notes
These examples assume that only policy settings for a single CSE that requires synchronous processing have changed.

When you troubleshoot Group Policy, consider that there may be interactions with multiple CSEs. You may find that the expected behavior for a single CSE varies when combined with the processing requirements of other CSEs. For more specific examples, see the Group Policy topics in Wiki: Group Policy Troubleshooting Portal.

Example 1: Fast Logon Optimization with synchronous processing when the user is not signed in

In this example, the user is not signed in to the computer when there is a change to the user policy setting for a CSE that requires synchronous processing. Synchronous processing was not requested when the user policy setting was last processed. In this example, it takes two sign-ins for the CSE to apply the policy settings.

When the user signs in, Fast Logon Optimization is in effect and Group Policy processes asynchronously. The Group Policy engine determines that there is a change to the policy settings for the CSE, and the Group Policy engine calls that CSE.

The CSE determines that it requires synchronous processing to apply the new settings. With the current processing session in asynchronous mode, the CSE exits without applying the policy settings, and the CSE signals to the Group Policy engine that synchronous processing is needed for the next sign-in.

At the second sign-in, Group Policy processes synchronously. The Group Policy engine determines that the CSE did not completely process the policy settings during the last sign-in session. The Group Policy engine calls the CSE. The CSE determines that Group Policy is processing synchronously, and the CSE applies the policy settings.

Example 2: Fast Logon Optimization with synchronous processing when the user is signed in

In this example, the user is already signed in to the computer when there is a change to the user policy setting for a CSE that requires synchronous processing. The user stays signed in to the computer until Group Policy background processing occurs, and the CSE applies the policy settings after a single sign-in.

When the user stays signed in to a client computer, by default, background processing occurs approximately every ninety minutes. Background processing is always run asynchronously. During this background processing, the Group Policy engine determines that there is a change to the policy settings for a CSE. The Group Policy engine calls that CSE.

The CSE determines that it requires synchronous processing to apply the new settings. With the background processing session in asynchronous mode, the CSE exits without applying the policy settings. The CSE signals to the Group Policy engine that synchronous processing is needed for the next sign-in. During subsequent background refreshes, the CSE is called, but it exits without applying the policy settings because it still requires synchronous processing.

At the next sign-in, Group Policy processes synchronously. The Group Policy engine determines that the CSE did not complete processing during the last sign-in session. The Group Policy engine calls the CSE. The CSE determines that Group Policy is processing synchronously, and the CSE applies the policy settings with a single sign-in.

Fast Logon Optimization, required synchronous processing, and required foreground processing for a CSE

A CSE can require synchronous processing and foreground processing to apply settings. The foreground processing requirement is set when the **NoBackgroundPolicy** registry key is set to a value of 1. In this case, the CSE applies the policy settings with two sign-ins.

For example, the user is signed in to the computer when there is a change to a user policy setting that requires synchronous processing and can only be applied during foreground processing. The user stays signed in to the client computer until Group Policy background processing occurs (by default, approximately every ninety minutes). During background processing, the Group Policy engine determines that there is a change to the policy settings for a CSE, but the **NoBackgroundPolicy** registry key is set to 1, and the Group Policy engine does not call this CSE during background processing. Additionally, during subsequent background refreshes, the Group Policy engine does not call the CSE.

At the next sign-in, Group Policy processes asynchronously because Fast Logon Optimization is in effect. This is foreground processing, so the Group Policy engine calls the CSE, which was set to only process in the foreground.

The CSE determines that it requires synchronous processing to apply the new policy settings. With the current processing session in asynchronous mode, the CSE exits without applying the policy settings. The CSE signals to the Group Policy engine that synchronous processing is needed for the next sign-in.

At the next sign-in, Group Policy processes synchronously. The Group Policy engine determines that the CSE did not completely process the policy settings during the last sign-in session. The Group Policy engine calls the CSE. The CSE determines that Group Policy is processing synchronously and applies the policy settings. In this case, it takes two sign-ins for the CSE to apply the settings.

Notes

This example assumes that only policy settings for a single CSE that requires synchronous processing and foreground processing, have changed.

When you troubleshoot Group Policy, consider that there may be interactions with multiple CSEs, and you may find that the expected behavior for a single CSE varies when combined with the processing requirements of other CSEs. For more specific examples, see the Group Policy topics in Wiki: Group Policy Troubleshooting Portal.

Fast startup and Group Policy processing

Group Policy settings or scripts that are applied during startup or shutdown might not be applied on computers that are running Windows 8.1 or Windows 8 because, by default, these computers are not fully shut down by the **Shut down** command. Instead, the computer enters a hibernate state. Full shutdown only occurs when the system is set to reboot or when a pending software update or other event causes the computer to process a full shutdown. You can use the **Restart**

Computer command in Windows 8.1 and Windows 8 to achieve startup or shutdown on client computers.

If you do not want fast startup to apply to computers in your enterprise, you can apply a GPO that disables the **Computer Configuration/Policies/Administrative Templates/System/Shutdown/Require use of fast startup** policy setting. Be aware that disabling this policy setting results in full shutdowns and longer startup times for client computers. The enabled state of the **Require use of fast startup** policy setting does not take precedence over any Group Policy setting that disables hibernation.

Related resources

- Troubleshooting Group Policy Client-Side Extension Behavior
- Description of the Windows XP Professional Fast Logon Optimization feature

Hyper-V Overview

Tip
Hyper-V is enhanced in Windows Server 2012 and Windows Server 2012 R2.

- Explore the Evaluation Guide.
- Download Windows Server 2012 R2.
- Download the Microsoft Hyper-V Server 2012 R2 software.

This topic describes the Hyper-V role, practical uses for the role, the most significant new or updated functionality in this version compared to previous versions of Hyper-V, hardware requirements, and a list of operating systems (known as *guest operating systems*) supported for use in a Hyper-V virtual machine.

Did you mean…

- Client Hyper-V
- Virtual Machine Manager
- What's New in Hyper-V for Windows Server 2012
- Overview of Hyper-V (Windows Server 2008 R2)

Role and technology description

The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built in to Windows Server. Installing the Hyper-V role installs the required components and optionally installs management tools. The required components include Windows hypervisor, Hyper-V Virtual Machine Management Service, the virtualization WMI provider, and other virtualization components such as the virtual machine bus (VMbus), virtualization service provider (VSP) and virtual infrastructure driver (VID).

The management tools for the Hyper-V role consist of:

- GUI-based management tools: Hyper-V Manager, a Microsoft Management Console (MMC) snap-in, and Virtual Machine Connection, which provides access to the video output of a virtual machine so you can interact with the virtual machine.
- Hyper-V-specific cmdlets for Windows PowerShell. Windows Server 2012 includes a Hyper-V module, which provides command-line access to all the functionality available in the GUI, as well functionality not available through the GUI. For more information about the Hyper-V module, see Hyper-V Module for Windows PowerShell.

If you use Server Manager to install the Hyper-V role, the management tools are included unless you specifically exclude them. If you use Windows PowerShell to install the Hyper-V role, the management tools are not included by default. To install the tools, use the parameter –IncludeManagementTools. For instructions about installing the Hyper-V role, see Install Hyper-V and Configure a Virtual Machine.

Notes

Hyper-V management tools in Windows Server 2012 are designed to manage this version of Hyper-V only. The tools cannot be used to manage earlier versions of Hyper-V.

The Hyper-V management tools in Windows Server 2012 R2 can be used to manage Hyper-V on Windows Server 2012 R2 and Windows Server 2012.

The Hyper-V technology virtualizes hardware to provide an environment in which you can run multiple operating systems at the same time on one physical computer. Hyper-V enables you to create and manage virtual machines and their resources. Each virtual machine is an isolated, virtualized computer system that can run its own operating system. The operating system that runs within a virtual machine is called a *guest operating system*.

For more information about the architecture of Hyper-V, see the Windows Server 2012 Hyper-V Component Architecture Poster.

Practical applications

Hyper-V provides infrastructure so you can virtualize applications and workloads to support a variety of business goals aimed at improving efficiency and reducing costs, such as:

- Establish or expand a private cloud environment. Hyper-V can help you move to or expand use of shared resources and adjust utilization as demand changes, to provide more flexible, on-demand IT services.
- Increase hardware utilization. By consolidating servers and workloads onto fewer, more powerful physical computers, you can reduce consumption of resources such as power and physical space.
- Improve business continuity. Hyper-V can help you minimize the impact of both scheduled and unscheduled downtime of your workloads.
- Establish or expand a virtual desktop infrastructure (VDI). A centralized desktop strategy with VDI can help you increase business agility and data security, as well as simplify regulatory compliance and management of the desktop operating system and applications. Deploy Hyper-V and Remote Desktop Virtualization Host (RD Virtualization Host) on the same

physical computer to make personal virtual desktops or virtual desktop pools available to your users.

- Increase efficiency in development and test activities. You can use virtual machines to reproduce different computing environments without the need for acquiring or maintaining all the hardware you would otherwise need.

New and changed functionality

Hyper-V in Windows Server 2012 and Windows Server 2012 R2 includes improvements in many areas.

For information about changes in Hyper-V in Windows Server 2012, see What's New in Hyper-V for Windows Server 2012.

For information about changes in Hyper-V in Windows Server 2012 R2, see What's New in Hyper-V for Windows Server 2012 R2.

Hardware requirements

Hyper-V requires a 64-bit processor that includes the following:

- Hardware-assisted virtualization. This is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.
- Hardware-enforced Data Execution Prevention (DEP) must be available and enabled. Specifically, you must enable Intel XD bit (execute disable bit) or AMD NX bit (no execute bit).

Important

The Hyper-V role is not supported on Windows Azure Virtual Machines. For a list of supported server roles on Windows Azure Virtual Machines, see Microsoft server software support for Windows Azure Virtual Machines.

Software requirements (for supported guest operating systems)

Hyper-V includes a software package for supported guest operating systems that improves integration between the physical computer and the virtual machine. This package is referred to as integration services. In general, you install this package in the guest operating system as a separate procedure after you set up the operating system in the virtual machine. However, some operating systems have the integration systems built-in and do not require a separate installation. For instruction on installing integration services, see Install Hyper-V and Configure a Virtual Machine.

The following table lists the operating systems supported in Windows Server 2012 and Windows Server 2012 R2 for use as guest operating systems in Hyper-V virtual machines, as well as provides information about integration services. This includes information about certain Linux

distributions for which integration services are available. Microsoft provides support for issues found in the integration services.

Important
For more information about running Linux as a guest operating system on Hyper-V, see Linux Virtual Machines on Hyper-V.

Microsoft provides support for the operating systems listed in the tables below in the following manner:

- Issues found in Microsoft operating systems and in integration services are supported by Microsoft support.
- For issues found in other operating systems that have been certified by the operating system vendor to run on Hyper-V, support is provided by the vendor. For example, see the Red Hat Enterprise Linux 5.9 entry in the table below.
- For issues found in other operating systems, Microsoft submits the issue to the multi-vendor support community, TSANet.

Guest operating system (server)	Maximum number of virtual processors	Notes
Windows Server 2012 R2	64	Integration services do not require a separate installation because they are built-in.
Windows Server 2012	64	Integration services do not require a separate installation because they are built-in. **Note** Upgrade the integration services after you set up the operating system in the virtual machine in Hyper-V running on Windows Server 2012 R2.
Windows Server 2008 R2 with Service Pack 1 (SP 1)	64	Datacenter, Enterprise, Standard and Web editions. Install the integration services after you set up the operating system in the virtual machine.
Windows Server 2008 R2 **Important** Only supported on	64	Datacenter, Enterprise, Standard and Web editions. Upgrade the integration services

Guest operating system (server)	Maximum number of virtual processors	Notes
Hyper-V running on Windows Server 2012.		after you set up the operating system in the virtual machine.
Windows Server 2008 with Service Pack 2 (SP 2)	4	Datacenter, Enterprise, Standard and Web editions (32-bit and 64-bit). Install the integration services after you set up the operating system in the virtual machine.
Windows Home Server 2011	4	Edition information is not applicable. Install the integration services after you set up the operating system in the virtual machine.
Windows Small Business Server 2011	Essentials edition - 2 Standard edition - 4	Essentials and Standard editions. Install the integration services after you set up the operating system in the virtual machine.
Windows Server 2003 R2 with Service Pack 2 (SP2)	2	Standard, Web, Enterprise, and Datacenter editions (32-bit and 64-bit). Install the integration services after you set up the operating system in the virtual machine.
Windows Server 2003 with Service Pack 2	2	Standard, Web, Enterprise, and Datacenter editions (32-bit and 64-bit). Install the integration services after you set up the operating system in the virtual machine.

Guest operating system (client)	Maximum number of virtual processors	Notes
Windows 8.1	32	Integration services do not require a separate installation

Guest operating system (client)	Maximum number of virtual processors	Notes
		because they are built-in.
Windows 8	32	Integration services do not require a separate installation because they are built-in. **Note** Upgrade the integration services after you set up the operating system in the virtual machine in Hyper-V running on Windows Server 2012 R2.
Windows 7 with Service Pack 1 (SP 1)	4	Ultimate, Enterprise, and Professional editions (32-bit and 64-bit). Upgrade the integration services after you set up the operating system in the virtual machine.
Windows 7	4	Ultimate, Enterprise, and Professional editions (32-bit and 64-bit). Upgrade the integration services after you set up the operating system in the virtual machine.
Windows Vista with Service Pack 2 (SP2)	2	Business, Enterprise, and Ultimate, including N and KN editions. Install the integration services after you set up the operating system in the virtual machine.
Windows XP with Service Pack 3 (SP3)	2	Professional. Install the integration services after you set up the operating system in the virtual machine.
Windows XP x64 Edition with Service Pack 2 (SP 2)	2	Professional. Install the integration services after you set up the operating system in the

Guest operating system (client)	Maximum number of virtual processors	Notes
		virtual machine.

See also

The following table provides links to related articles and technologies.

Content type	References
Product evaluation	• Building Your Cloud Infrastructure: Scenario Overview • Linux Virtual Machines on Hyper-V • Increasing Server, Storage, and Network Availability: Scenario Overview • Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100) • Hyper-V Comparison Guide • Understand and Troubleshoot Hyper-V Replica in Windows Server "8"
Troubleshooting	• Understand and Troubleshoot Hyper-V Replica in Windows Server "8"
Tools and settings	• Hyper-V Module for Windows PowerShell • Hyper-V Cmdlets in Windows PowerShell
Community resources	• Ben Armstrong's Virtual PC Guy's WebLog • Hyper-V Portal on the TechNet Wiki • Hyper-V in the Windows Server forum • General discussion in the Windows Server 2012 forum
Related technologies	• **Microsoft Hyper-V Server 2012** • Failover Clustering Overview • File and Storage Services Overview • Remote Desktop Services Overview

Feature Overviews

Feature overviews for Hyper-V running on Windows Server® 2012 R2:

- Virtual Machine Connection - Enhanced Session Mode Overview
- Automatic Virtual Machine Activation
- Virtual Hard Disk Sharing Overview
- Online Virtual Hard Disk Resizing Overview
- Storage Quality of Service for Hyper-V
- Generation 2 Virtual Machine Overview

Feature overviews for Hyper-V running on Windows Server 2012 R2 and Windows Server® 2012

- Hyper-V Automation Overview
- Hyper-V Dynamic Memory Overview
- Hyper-V Network Virtualization Overview
- Hyper-V Offloaded Data Transfer Overview
- Hyper-V Replica Feature Overview
- Hyper-V Resource Metering Overview
- Hyper-V Support for Large Sector Disks Overview
- Hyper-V Virtual Fibre Channel Overview
- Hyper-V Virtual Hard Disk Format Overview
- **Hyper-V Virtual Switch Technical Preview**
- Hyper-V Support for Scaling Up and Scaling Out Overview
- Virtual Machine Storage Migration Overview
- Virtual Machine Live Migration Overview
- Simplified Import Overview
- Hyper-V Network Virtualization Overview
- Simplified Import Overview
- Hyper-V Virtual NUMA Overview

Hyper-V Automation Overview

Windows PowerShell is the scripting solution for automating tasks in Windows Server. However, in current versions of Hyper-V, writing scripts for Hyper-V with in-box tools requires users to learn Windows Management Instrumentation (WMI). WMI provides a very flexible set of interfaces, but they are designed for developers, not IT pros. Hyper-V in Windows Server 2012 addresses this issue by introducing more than 140 built-in Hyper-V cmdlets for Windows PowerShell. Administrators can use these new cmdlets to more easily enable automation of datacenter tasks that range from basic to complex.

Key benefits

The new Hyper-V cmdlets for Windows PowerShell provide IT pros with an easy way to enable automation of management tasks within Windows Server 2012. With the extensive number of Hyper-V cmdlets and the close integration with other parts of the operating system, administrators can now easily enable automation of Hyper-V-related tasks in their environment.

Requirements

To try out the Hyper-V cmdlets, you will need the following:

- A computer running Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer that has processor support for hardware virtualization.
- A user account that is a member of the Administrators group or Hyper-V Administrators group.

Technical overview

Hyper-V cmdlets enable administrators to perform the tasks that are available in the graphical user interface of Hyper-V Manager—as well as a number of tasks that can be done exclusively through the cmdlets in Windows PowerShell.

Designed for IT pros

Windows PowerShell is designed intentionally for IT pros. This design decision is reflected in several ways:

- **Task-oriented interface.** Hyper-V cmdlets are designed so that it is easy for IT pros to go from thinking about the task to actually performing the task. The following table shows the task and the associated cmdlet syntax:

Task	**Windows PowerShell command to perform the task**
Create a new virtual machine named “test”	New-VM –Name Test
Get a list of all virtual machines	Get-VM
Create a new virtual hard disk at d:\VHDs\test.vhd	New-VHD –Path D:\VHDs\test.vhd
Start all virtual machines whose name begins with “web”	Start-VM –Name web*
Connect the virtual network adapter on the “test” virtual machine to the “QA” switch.	Connect-VMNetworkAdapter –VMName test –SwitchName QA

- **Use of standard cmdlet verbs.** Hyper-V administrators often need to manage more than just Hyper-V. By using the same verbs as other Windows cmdlets, the Hyper-V cmdlets make easier for administrators to extend their existing knowledge of Windows PowerShell. For example, administrators who are familiar with managing services through Windows PowerShell can reuse the same verbs to perform the corresponding tasks on a virtual machine, as shown in the following table:

Task	**Cmdlet for performing task on a service**	**Cmdlet for performing task on a virtual machine**
Get	Get-Service	Get-VM
Configure	Set-Service	Set-VM
Create	New-Service	New-VM
Start	Start-Service	Start-VM
Stop	Stop-Service	Stop-VM
Restart	Restart-Service	Restart-VM
Suspend	Suspend-Service	Suspend-VM
Resume	Resume-Service	Resume-VM

There are similar examples with other core Windows PowerShell cmdlets as well:

Core Windows PowerShell cmdlet	**Hyper-V cmdlet**
Import-Csv	Import-VM
Export-Csv	Export-VM
Enable-PSRemoting	Enable-VMMigration
Checkpoint-Computer	Checkpoint-VM
Measure-Command	Measure-VM

- **Consistent cmdlet nouns to simplify discoverability.** The nouns of the Hyper-V cmdlets are designed to make it easier for administrators to discover the cmdlets they need when they need them. All cmdlets in the Hyper-V module use one of three following noun prefixes:

Prefix	**Purpose**
VM	Cmdlets for managing virtual machines
VHD	Cmdlets for managing virtual hard disk files
VFD	Cmdlets for managing virtual floppy disk files

Hyper-V Dynamic Memory Overview

If you have idle or low-load virtual machines, as in pooled Virtual Desktop Infrastructure (VDI) environments, Dynamic Memory enables you to increase consolidation and improve reliability for restart operations. You also gain agility in responding to requirement changes with these new capabilities.

Key benefits

With the Dynamic Memory improvements for Hyper-V in Windows Server 2012, you can attain higher consolidation numbers with improved reliability for restart operations. This can lead to lower costs, especially in environments that have many idle or low-load virtual machines, such as pooled VDI environments. Dynamic Memory run-time configuration changes can reduce downtime and provide increased agility to respond to requirement changes.

Technical overview

Dynamic Memory, introduced in Windows Server 2008 R2 Service Pack 1 (SP1), defined startup memory as the minimum amount of memory that a virtual machine can have. However, Windows requires more memory during startup than the steady state. As a result, administrators sometimes assign extra memory to a virtual machine because Hyper-V cannot reclaim memory from these virtual machines after startup. In Windows Server 2012, Dynamic Memory introduces a minimum memory setting, which allows Hyper-V to reclaim the unused memory from the virtual machines. This is reflected as increased virtual machine consolidation numbers, especially in Virtual Desktop Infrastructure (VDI) environments.

Windows Server 2012 also introduces Smart Paging for reliable virtual machine restart operations. Although minimum memory increases virtual machine consolidation numbers, it also brings a challenge. If a virtual machine has a smaller amount of memory than its startup memory and if it is restarted, Hyper-V needs additional memory to restart the virtual machine. Due to host memory pressure or virtual machine states, Hyper-V may not always have additional memory available. This can cause sporadic virtual machine restart failures. Smart Paging is used to bridge the memory gap between minimum memory and startup memory, and allow virtual machines to restart reliably.

Minimum memory configuration with reliable restart operation

As in the previous version of Dynamic Memory, you can configure a minimum memory amount for virtual machines, and Hyper-V continues to ensure that this amount is always assigned to running virtual machines.

To provide a reliable restart experience for the virtual machines configured with less minimum memory than startup memory, Hyper-V uses Smart Paging. This memory management method

uses disk resources as additional, temporary memory when more memory is required to restart a virtual machine. This approach has advantages and drawbacks. It provides a reliable way to keep the virtual machines running when there is no available physical memory. However, it can degrade virtual machine performance because disk access speeds are much slower than memory access speeds.

Important

When installing or upgrading the operating system of a virtual machine, the amount of memory that is available to the virtual machine during the installation and upgrade process is the value specified as **Startup RAM**. Even if Dynamic Memory has been configured for the virtual machine, the virtual machine only uses the amount of memory as configured in the **Startup RAM** setting. Ensure the **Startup-RAM** value meets the minimum memory requirements of the operating system during the install or upgrade procedure.

To minimize the performance impact of Smart Paging, Hyper-V uses it only when all of the following occurs:

- The virtual machine is being restarted.
- There is no available physical memory.
- No memory can be reclaimed from other virtual machines running on the host.

Smart Paging is not used when:

- A virtual machine is being started from an “off state” (instead of a restart).
- Oversubscribing memory for a running virtual machine is required.
- A virtual machine is failing over in Hyper-V clusters.

When host memory is oversubscribed, Hyper-V continues to rely on the paging operation in the guest operating system because it is more effective than Smart Paging. The paging operation in the guest operating system is performed by Windows Memory Manager. Windows Memory Manager has more information than the Hyper-V host about memory usage within the virtual machine, which means it can provide Hyper-V with better information to use when choosing the memory to be paged. Because of this, less overhead to the system is incurred compared to Smart Paging.

To further reduce the impact of Smart Paging, Hyper-V removes memory from the virtual machine after it completes the start process. It accomplishes this by coordinating with Dynamic Memory components inside the guest operating system (a process sometimes referred to as “ballooning”), so the virtual machine stops using Smart Paging. With this technique, the use of Smart Paging is temporary and is not expected to be longer than 10 minutes.

Also note the following about how Smart Paging is used:

- Smart Paging files are created only when needed for a virtual machine.
- After the additional amount of memory is removed, Smart Paging files are deleted.
- Smart Paging is not used for this virtual machine again until another restart occurs and there is not enough physical memory.

Run-time Dynamic Memory configuration changes

Hyper-V enables users to make the following configuration changes to Dynamic Memory when the virtual machine is running:

- Increase the maximum memory.
- Decrease the minimum memory.

About the Dynamic Memory Settings

The following virtual machine settings in Hyper-V Manager allow you to configure Dynamic Memory. Review the following descriptions to understand how the features work and how they are related.

Setting	Description
Startup RAM	Specifies the amount of memory required to start the virtual machine. The value needs to be high enough to allow the guest operating system to start, but should be as low as possible to allow for optimal memory utilization and potentially higher consolidation ratios.
Minimum RAM	Specifies the minimum amount of memory that should be allocated to the virtual machine after the virtual machine has started. The value can be set as low as 32 MB to a maximum value equal to the **Startup RAM** value.
Maximum RAM	Specifies the maximum amount of memory that this virtual machine is allowed to use. The value can be set from as low as the value for Startup RAM to as high as 1 TB. However, a virtual machine can use only as much memory as the maximum amount supported by the guest operating system. For example, if you specify 64 GB for a virtual machine running a guest operating system that supports a maximum of 32 GB, the virtual machine cannot use more than 32 GB.
Memory buffer	Specifies how much memory Hyper-V will attempt to assign to the virtual machine compared to the amount of memory actually needed by the applications and services running inside the virtual machine. Memory

Setting	Description
	buffer is specified as a percentage because the actual amount of memory that represents the buffer changes in response to changes in memory usage while the virtual machine is running. Hyper-V uses performance counters in the virtual machine that identify committed memory to determine the current memory requirements of the virtual machine and then calculates the amount of memory to add as a buffer. The buffer is determined using the following formula: Amount of memory buffer = how much memory the virtual machine actually needs / (memory buffer value / 100). For example, if the memory committed to the guest operating system is 1000 MB and the memory buffer is 20%, Hyper-V will attempt to allocate an additional 20% (200 MB) for a total of 1200 MB of physical memory allocated to the virtual machine. **Note** The buffer is not maintained when there is not enough physical memory available in the computer to give every virtual machine its requested memory buffer.
Memory weight	Provides Hyper-V with a way to determine how to distribute memory among virtual machines if there is not enough physical memory available in the computer to give every virtual machine its requested amount of memory.

Guest operating systems that support Dynamic Memory

The following table lists the guest operating systems that support Dynamic Memory on Hyper-V in Windows Server 2012 and Windows Server 2012 R2.

For a list of the operating systems supported for use as guest operating systems in Hyper-V virtual machines see Software requirements (for supported guest operating systems).

Guest operating system	Notes
Linux distributions	For details on Linux distributions and versions that support dynamic memory, and the requirements for Linux Integration Services, see Linux Virtual Machines on Hyper-V.
Windows Server 2012 and Windows Server 2012 R2	Datacenter, Standard, and Essentials editions. Integration services are built-in and do not require a separate download and installation.
Windows 8 and Windows 8.1	Windows 8, Windows 8 Pro, and Windows 8 Enterprise editions. Windows 8.1, Windows 8.1 Pro, and Windows 8.1 Enterprise editions. Integration services are built-in and do not require a separate download and installation.
Windows Server 2008 R2 with Service Pack 1 (SP1)	Datacenter, Enterprise, Standard, and Web editions.
Windows Server 2008 R2	Datacenter and Enterprise editions.
Windows 7	Enterprise and Ultimate editions (32-bit and 64-bit).
Windows Server 2008 with Service Pack 2 (SP2)	Standard and Web editions (32-bit and 64-bit). Apply the following hotfix as described in article 2230887(http://go.microsoft.com/fwlink/?LinkId=206472)
Windows Server 2008 with Service Pack 2 (SP2)	Datacenter and Enterprise editions (32-bit and 64-bit).
Windows Vista with Service Pack 1 (SP1)	Ultimate and Enterprise editions (32-bit and 64-bit).
Windows Server 2003 R2 with Service Pack 2 (SP2)	Datacenter, Enterprise, Standard, and Web editions (32-bit and 64-bit).
Windows Server 2003 with Service Pack 2 (SP2)	Datacenter, Enterprise, Standard, and Web editions (32-bit and 64-bit).

Note
Install the latest integration services after you set up the operating system in the virtual machine.

Note

Hyper-V running on Windows Server 2012 and Windows Server 2012 R2 automatically calculates an amount of memory to reserve for exclusive use by the management host operating system. This memory is used to run virtualization services. If the computer is part of a failover cluster, Hyper-V also reserves enough memory to run failover cluster services. In previous versions of Hyper-V you might have needed to modify the **MemoryReserve** registry setting to ensure adequate memory for these functions. This is no longer necessary for Hyper-V running on Windows Server 2012 and Windows Server 2012 R2.

Dynamic Memory / Hyper-V Memory Management FAQs

The following are some frequently asked questions regarding Dynamic Memory.

Do I need to manually configure the size of the page file after installing the Hyper-V role?

No, it is recommended to allow the system to manage the page file size when running Hyper-V.

How is the memory divided up between the host operating system and running virtual machines on Windows Server 2012 R2?

Hyper-V running on Windows Server 2012 R2 monitors the runtime health of the host operating system to determine how much memory can be safely allocated to virtual machines.

In previous versions, Hyper-V performed calculations at boot-time to estimate how much memory could be consumed by virtual machines and how much should be reserved for the host operating system.

The move to runtime monitoring in Windows Server 2012 R2 improves overall system reliability and performance.

My virtual machine failed to start with the following error message “Not enough memory in the system to start the virtual machine” or “Could not initialize memory: Ran out of memory (0x8007000E)”. How can I tell how much memory is available for virtual machines?

As described above, Hyper-V running on Windows Server 2012 R2 monitors the runtime health of the host operating system to determine how much memory can be safely allocated to virtual machines.

The current amount of memory available to virtual machines can be viewed in the following Performance Monitor counter, **Hyper-V Dynamic Memory Balancer – Available Memory**.

Should I configure the MemoryReserve registry setting?

When running Hyper-V on Windows Server 2012 R2 it is recommend that the **MemoryReserve** registry setting no longer be created.

With Hyper-V running on Windows Server 2008 R2 or Windows Server 2008 R2 Service Pack 1, certain cases necessitated modifying the **MemoryReserve** registry setting to ensure adequate memory for virtualization and failover clustering services in the host operating system. Improved memory reserve calculations in Windows Server 2012 and the addition of runtime monitoring of the host operating system health in Windows Server 2012 R2 makes use of the **MemoryReserve** registry setting unnecessary.

Hyper-V Offloaded Data Transfer Overview

To take advantage of innovations in storage hardware that provide near-instantaneous copying of large amounts of data, Hyper-V in Windows Server 2012 introduces Offloaded Data Transfer. With this new feature, Hyper-V workloads use the offload semantics of the host hardware, as well as the virtual storage stack, to perform certain internal operations on virtual hard disks that require large amounts of data to be copied. Hyper-V performs these operations faster than was previously possible.

Key benefits

Hyper-V Offloaded Data Transfer is a new hardware feature that allows for copying large amounts of data from one location to another. By providing support for offloaded data transfer in the Hyper-V storage stack, this feature makes it possible to complete these operations in a fraction of the time it would have taken without the support. Allowing the virtualized workload to use the semantics for offloaded data transfer by passing from the workload to the host hardware helps the virtualized workload operate as efficiently as it would in a non-virtualized environment.

Requirements

Hyper-V Offloaded Data Transfer requires the following:

- Offloaded Data Transfer–capable hardware to host the virtual hard disk files. The hardware needs to be connected to the virtual machine as virtual SCSI devices or directly attached physical disks (sometimes referred to as pass-through disks).
- This optimization is also supported for natively attached, VHDX-based virtual disks.
- VHD-based or VHDX-based virtual disks attached to a virtual IDE controller do not support this optimization because integrated development environment (IDE) devices lack support for Offloaded Data Transfer.

Technical overview

Storage area network (SAN) vendors are working to provide near-instantaneous copy operations of large amounts of data. This storage is designed to allow the system above the disks to specify the move of a specific data set from one location to another (an “offloaded data transfer”).

Crucial maintenance tasks for virtual hard disks—such as merge, move, and compact—depend on copying large amounts of data. The current method of copying data requires data to be read in and written to different locations, which can be a time-consuming process.

The storage stack of Hyper-V in Windows Server 2012 supports offloaded data transfer operations so that these operations can be passed from the guest operating system to the host hardware. This ensures that the workload can use storage enabled for offloaded data transfer as it would if it were running in a non-virtualized environment. The Hyper-V storage stack also issues offloaded data transfer operations during maintenance operations for virtual hard disks, such as merging disks and storage migration meta-operations where large amounts of data are moved.

Hyper-V Replica Feature Overview

With Hyper-V Replica, administrators can replicate their Hyper-V virtual machines from one Hyper-V host at a primary site to another Hyper-V host at the Replica site. This feature lowers the total cost-of-ownership for an organization by providing a storage-agnostic and workload-agnostic solution that replicates efficiently, periodically, and asynchronously over IP-based networks across different storage subsystems and across sites. This scenario does not rely on shared storage, storage arrays, or other software replication technologies. The following figure demonstrates how Hyper-V Replica lets administrators easily replicate virtual machines to a remote site over a WAN link.

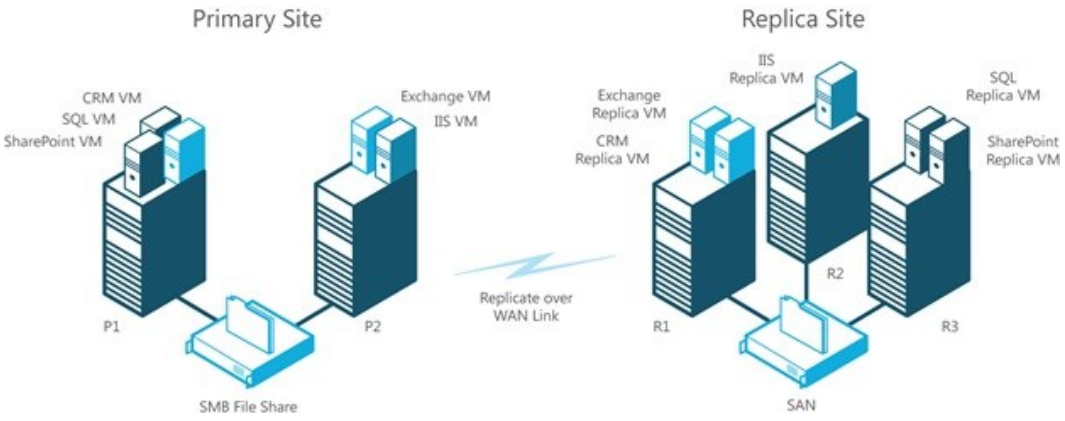

Figure 1 Replicated virtual machines

Additionally, administrators can use Hyper-V Replica to test the Replica virtual machine without disrupting the ongoing replication. If a failure occurs at the primary site, administrators can quickly restore their business operations by bringing up the replicated virtual machine at the Replica site.

Key benefits

Virtual machines can be easily replicated to different locations for greater protection and availability.

Key features

Hyper-V Replica tracks the write operations on the primary virtual machine and then replicates these changes to the Replica server over a WAN. The network connection between the two servers uses the HTTP protocol and supports Kerberos authentication and certificate-based authentication, with optional support for encryption.

Hyper-V Replica is closely integrated with failover clustering in Windows Server 2012, and it provides nearly seamless replication across different migration scenarios in the primary and Replica servers. This allows virtual hard disks to be stored in a different location to enable recovery in case the data center goes down due to natural disaster or other causes.

Hyper-V Resource Metering Overview

IT organizations need tools to charge back business units that they support while providing the business units with the right amount of resources to match their needs. For hosting providers, it is equally important to issue chargebacks based on the amount of usage by each customer.

To implement advanced billing strategies that measure both the assigned capacity of a resource and its actual usage, earlier versions of Hyper-V required users to develop their own chargeback solutions that polled and aggregated performance counters. These solutions could be expensive to develop and sometimes led to loss of historical data.

To assist with more accurate, streamlined chargebacks while protecting historical information, Hyper-V in Windows Server 2012 introduces Resource Metering, a feature that allows customers to create cost-effective, usage-based billing solutions. With this feature, service providers can choose the best billing strategy for their business model, and independent software vendors can develop more reliable, end-to-end chargeback solutions on top of Hyper-V.

Key benefits

Hyper-V Resource Metering in Windows Server 2012 allows organizations to avoid the expense and complexity associated with building in-house metering solutions to track usage within specific business units. It enables hosting providers to quickly and cost-efficiently create a more advanced, reliable, usage-based billing solution that adjusts to the provider's business model and strategy.

Use of network metering port ACLs

Enterprises pay for the Internet traffic in and out of their data centers, but not for the network traffic within their data centers. For this reason, providers generally consider Internet and intranet traffic separately for the purposes of billing. To differentiate between Internet and intranet traffic, providers can measure incoming and outgoing network traffic for any IP address range, by using network metering port ACLs.

Virtual machine metrics

Windows Server 2012 provides two options for administrators to obtain historical data on a client's use of virtual machine resources: Hyper-V cmdlets in Windows PowerShell and the new APIs in the Virtualization WMI provider. These tools expose the metrics for the following resources used by a virtual machine during a specific period of time:

- Average CPU usage, measured in megahertz over a period of time.
- Average physical memory usage, measured in megabytes.
- Minimum memory usage (lowest amount of physical memory).
- Maximum memory usage (highest amount of physical memory).
- Maximum amount of disk space allocated to a virtual machine.
- Total incoming network traffic, measured in megabytes, for a virtual network adapter.
- Total outgoing network traffic, measured in megabytes, for a virtual network adapter.

Movement of virtual machines between Hyper-V hosts—for example, through live, offline, or storage migrations—does not affect the collected data.

Hyper-V Support for Large Sector Disks Overview

Hyper-V in Windows Server 2012 introduces support for large sector size disks to help ensure compatibility with emerging innovations in storage hardware.

The data storage industry is transitioning the physical format of hard disk drives from 512-byte sectors to 4096-byte sectors (also known as 4K or 4 KB sectors). This transition is driven by several factors, including increases in storage density and reliability.

However, most of the software industry has depended on disk sectors of 512 bytes in length. A change in sector size introduces compatibility issues in many applications. To minimize the impact on the industry, hard drive vendors are introducing transitional 512-byte emulation (512e) drives. These drives offer some of the advantages offered by 4 KB native drives (such as improved format efficiency and an improved scheme for error correction codes (ECC)), but with fewer compatibility issues than by exposing a 4 KB sector size at the disk interface.

Key benefits

The storage industry is introducing 4 KB physical format drives to provide increased capacity. This new format is expected to be widely adopted by customers. Updates to the virtualization stack in Hyper-V in Windows Server 2012 ensure compatibility for storage configurations where the underlying physical hard disk uses the new 4 KB format.

Requirements

This feature requires physical disk drives that use either the 512e or native 4 KB format.

Technical overview

Hyper-V in Windows Server 2012 provides the following:

Support for improved performance of VHDs on 512e disks

512e disks can perform write operations only in terms of a physical sector—that is, it cannot directly write a 512-byte sector write issued to it. The internal process in the disk that makes this write possible consists of the following steps:

1. The disk reads the 4 KB physical sector into its internal cache, which contains the 512-byte logical sector referred to in the write.
2. Data in the 4 KB buffer is modified to include the updated 512-byte sector.
3. The disk performs a write of the updated 4 KB buffer back to its physical sector on the disk.

This process is called a Read-Modify-Write (RMW). This RMW process causes performance degradation in virtual hard disks (VHDs) because to the following reasons:

1. Dynamic and differencing VHDs have a 512-byte sector bitmap in front of their data payload. In addition, footer/header/Parent Locators all align to a 512-byte sector. Therefore, it is common for the VHD driver to issue 512-byte writes to update these structures, which results in the RMW behavior described above.
2. It is common for applications to issue reads and writes in multiples of 4 KB sizes (the default cluster size of NTFS). Because there is a 512-byte sector bitmap in front of the data payload block of dynamic and differencing VHDs, the 4 KB blocks are not aligned to the physical 4 KB boundary. Each 4 KB write issued by the current parser to update the payload data results in two reads for two blocks on the disk, which are then updated and subsequently written back to the two disk blocks.

The overall performance impact to the workloads was in the range of 30% to 80% and was even higher at times.

Hyper-V mitigates the performance impact of 512e disks on the VHD stack by padding the previously mentioned structures to be aligned to 4 KB boundaries in the VHD format. This mitigates the RMW impact when accessing the data within the VHD file, as well as when updating the VHD metadata structures.

Support for hosting VHDs on native 4 KB disks

The current VHD driver assumes a physical sector size of 512 bytes and issues 512-byte I/Os, which makes it incompatible with these disks. As a result, the current VHD driver cannot open VHD files on physical 4 KB sector disks. Hyper-V makes it possible to store VHDs on 4 KB disks by implementing a software RMW algorithm in the VHD layer to convert the 512-byte access and update request to the VHD file to corresponding 4 KB accesses and updates.

Hyper-V Virtual Fibre Channel Overview

You need your virtualized workloads to connect easily and reliably to your existing storage arrays. Hyper-V provides Fibre Channel ports within the guest operating system, which allows you to connect to Fibre Channel directly from within virtual machines. This feature protects your investments in Fibre Channel, enables you to virtualize workloads that use direct access to Fibre Channel storage, allows you to cluster guest operating systems over Fibre Channel, and provides an important new storage option for servers hosted in your virtualization infrastructure.

Key benefits

With this Hyper-V virtual Fibre Channel feature, you can connect to Fibre Channel storage from within a virtual machine. This allows you to use your existing Fibre Channel investments to support virtualized workloads. Support for Fibre Channel in Hyper-V guests also includes support for many related features, such as virtual SANs, live migration, and MPIO.

Requirements

The virtual Fibre Channel feature in Hyper-V requires the following:

- One or more installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer with processor support for hardware virtualization.
- A computer with one or more Fibre Channel host bus adapters (HBAs) that have an updated HBA driver that supports virtual Fibre Channel. Updated HBA drivers are included with the in-box HBA drivers for some models. The HBA ports to be used with virtual Fibre Channel should be set up in a Fibre Channel topology that supports NPIV. To determine whether your hardware supports virtual Fibre Channel, contact your hardware vendor or OEM.
- An NPIV-enabled SAN.
- Virtual machines configured to use a virtual Fibre Channel adapter, which must use Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 as the guest operating system.
- Storage accessed through a virtual Fibre Channel supports devices that present logical units. Virtual Fibre Channel logical units cannot be used as boot media.

Technical overview

Virtual Fibre Channel for Hyper-V provides the guest operating system with unmediated access to a SAN by using a standard World Wide Name (WWN) associated with a virtual machine. Hyper-V users can now use Fibre Channel SANs to virtualize workloads that require direct access to SAN logical unit numbers (LUNs). Fibre Channel SANs also allow you to operate in new scenarios, such as running the Failover Clustering feature inside the guest operating system of a virtual machine connected to shared Fibre Channel storage.

Mid-range and high-end storage arrays are capable of advanced storage functionality that helps offload certain management tasks from the hosts to the SANs. Virtual Fibre Channel presents an alternate hardware-based I/O path to the Windows software virtual hard disk stack. This allows you to use the advanced functionality offered by your SANs directly from Hyper-V virtual machines. For example, you can use Hyper-V to offload storage functionality (for example, taking a snapshot of a LUN) on the SAN hardware by using a hardware Volume Shadow Copy Service (VSS) provider from within a Hyper-V virtual machine.

NPIV support

Virtual Fibre Channel for Hyper-V guests uses the existing N_Port ID Virtualization (NPIV) T11 standard to map multiple virtual N_Port IDs to a single physical Fibre Channel N_port. A new NPIV port is created on the host each time you start a virtual machine that is configured with a virtual HBA. When the virtual machine stops running on the host, the NPIV port is removed. Due to the use of NPIV, the HBA ports used for virtual Fibre Channel should be set up in a Fibre Channel topology that supports NPIV, and the SAN should support NPIV ports.

Virtual SAN support

Hyper-V allows you to define virtual SANs on the host to accommodate scenarios where a single Hyper-V host is connected to different SANs through multiple Fibre Channel ports. A virtual SAN defines a named group of physical Fibre Channel ports that are connected to the same physical SAN. For example, assume that a Hyper-V host is connected to two SANs—a production SAN and a test SAN. The host is connected to each SAN through two physical Fibre Channel ports. In this example, you might configure two virtual SANs—one named “Production SAN” that has the two physical Fibre Channel ports connected to the production SAN and one named “Test SAN” that has two physical Fibre Channel ports connected to the test SAN. You can use the same technique to name two separate paths to a single storage target.

You can configure as many as four virtual Fibre Channel adapters on a virtual machine and associate each one with a virtual SAN. Each virtual Fibre Channel adapter connects with one WWN address or two WWN addresses to support live migration. You can set each WWN address automatically or manually.

Tape library support

Tape libraries are not supported on virtual machines configured with a Virtual Fibre Channel adapter.

Live migration

To support live migration of virtual machines across Hyper-V hosts while maintaining Fibre Channel connectivity, two WWNs are configured for each virtual Fibre Channel adapter, as shown in Figure 1: Set A and Set B. Hyper-V automatically alternates between the Set A and Set B WWN addresses during a live migration. This ensures that all LUNs are available on the destination host before the migration and that no downtime occurs during the migration.

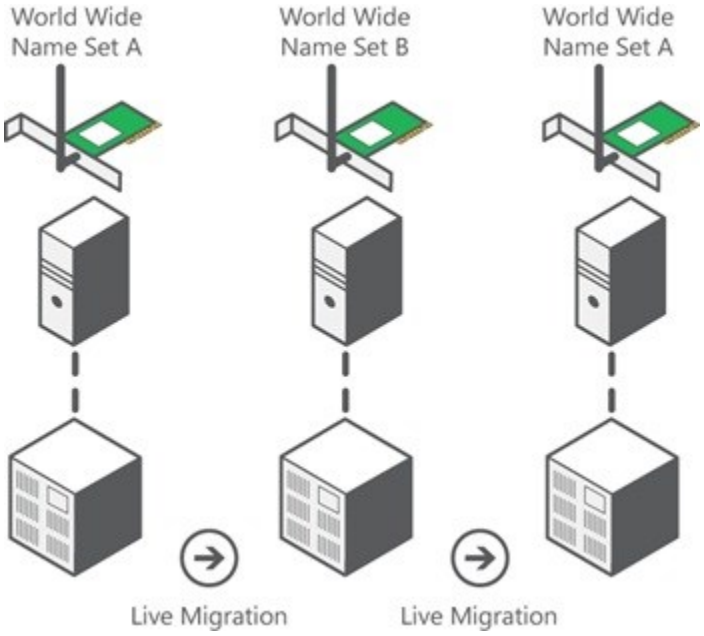

Figure 1 Alternating WWN addresses during a live migration

MPIO functionality

Hyper-V in Windows Server 2012 can use the multipath I/O (MPIO) functionality to ensure continuous connectivity to Fibre Channel storage from within a virtual machine. You can use MPIO functionality with Fibre Channel in the following ways:

- Use MPIO for host access. Install multiple Fibre Channel ports on the host, and use MPIO to provide highly available connectivity to the LUNs accessible by the host.
- Virtualize workloads that use MPIO. Configure multiple virtual Fibre Channel adapters inside a virtual machine, and use a separate copy of MPIO within the guest operating system of the virtual machine to connect to the LUNs that the virtual machine can access. This configuration can coexist with a host MPIO setup.
- Use different device specific modules (DSMs) for the host or each virtual machine. This approach allows live migration of the virtual machine configuration, including the configuration

of DSM and connectivity between hosts and compatibility with existing server configurations and DSMs.

See also

Implement Hyper-V Virtual Fibre Channel

Implement Hyper-V Virtual Fibre Channel

Until recently, some variants of physical servers have proven challenging to virtualize—specifically those that are dependent on virtual Fibre Channel storage. In Windows Server 2012 R2 and Windows Server 2012, this issue has been resolved. Hyper-V provides virtual Fibre Channel ports within the guest operating system, which allows you to connect to virtual Fibre Channel directly from within virtual machines.

The virtual Fibre Channel feature in Windows Server 2012 R2 and Windows Server 2012 makes it possible for you to virtualize workloads that previously depended on Fibre Channel storage. Your virtual machines can now directly access Fibre Channel storage as if your operating system was running on a physical server.

Support for virtual Fibre Channel in Hyper-V guests also includes support for many related features, such as virtual SANs (vSAN), live migration, quick migration, MPIO, Import and Export, Save and Restore, Pause and Resume, and host initiated backups.

For more information, refer to the Hyper-V Virtual Fibre Channel Overview.

To implement virtual Fibre Channel, complete the following steps.

In this document

- Prerequisites
- Step 1: Review considerations for live migration
- Step 2: Review considerations for MPIO connectivity to virtual Fibre Channel storage
- Step 3: Configure a virtual Fibre Channel adapter
- Terms and definitions

Prerequisites

The virtual Fibre Channel feature has the following prerequisites:

- A computer with one or more Fibre Channel host bus adapters (HBAs) or Fibre Channel over Ethernet (FCoE) converged network adapters. These devices must use updated drivers that support virtual Fibre Channel. The HBA ports should be set up in a Fibre Channel topology that supports N_Port ID Virtualization (NPIV). To determine if your hardware supports virtual Fibre Channel, contact your hardware vendor or OEM.
- An NPIV-enabled storage area network (SAN).

- In Windows Server 2012 R2 and Windows Server 2012, you can have up to four virtual ports per virtual machine. Depending on your configuration, your hardware may restrict the number of virtual ports per HBA and the logical unit numbers (LUNs) per port.
- Virtual machines that are configured to use a virtual Fibre Channel adapter must use Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 as the guest operating system. No other guest operating systems are currently supported for use with virtual Fibre Channel.
- Hardware can limit the number of virtual ports per physical host bus adapter. This restricts the number of virtual machines that can be associated with each physical host bus adapter that is installed on the server running Hyper-V.
- Often NPIV is disabled by default on your host bus adapter, so refer to your manufacturer's recommendations to enable this functionality on your host bus adapter.
- Storage that is accessed through a virtual Fibre Channel adapter supports devices that present logical unit numbers. Virtual Fibre Channel LUNs cannot be used as boot media. Fibre Channel tape devices are not supported.
- Always refer to the documentation that is provided by your storage vendor when working with storage hardware configurations. There may be additional vendor-specific steps required.

The basic virtual Fibre Channel architecture is shown in Figure 1.

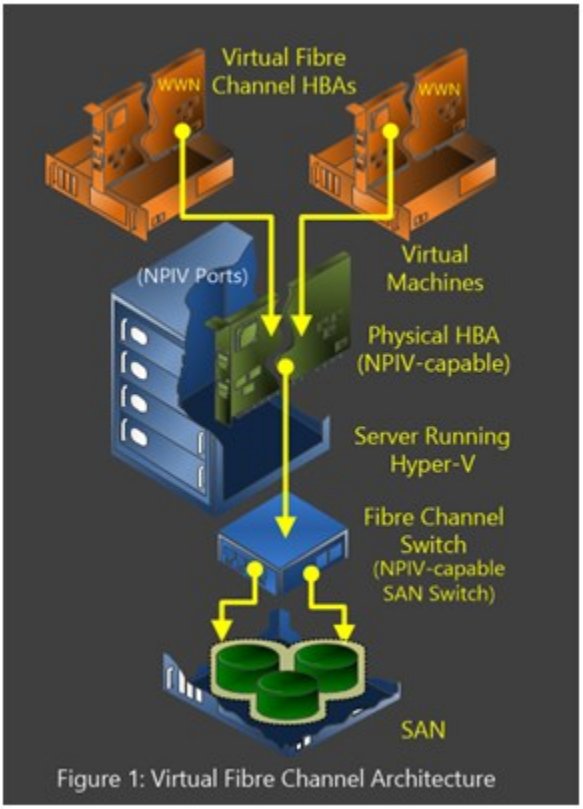

Figure 1: Virtual Fibre Channel Architecture

Step 1: Review considerations for live migration

Hyper-V in Windows Server 2012 R2 and Windows Server 2012 supports live migration of virtual machines across computers running Hyper-V, while maintaining Fibre Channel connectivity. Live migration is accomplished with failover clusters, and each virtual Fibre Channel adapter is configured with two World Wide Names (WWNs).

As shown in Figure 2, Hyper-V automatically changes between Set A and Set B WWN addresses during a live migration. Hyper-V ensures that all LUNs are available on the destination computer, and then it performs the live migration.

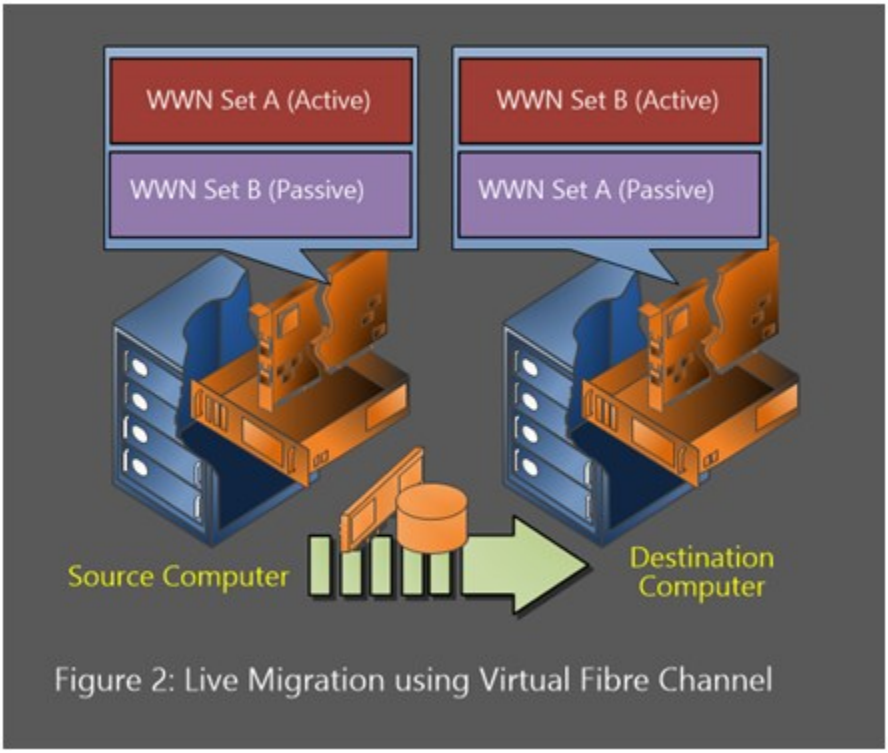

Figure 2: Live Migration using Virtual Fibre Channel

If you want to use live migration with virtual Fibre Channel, each server running Hyper-V must have access to any virtual SANs that are being used by the virtual machine. In addition, the virtual SAN connectivity must have the same number of ports on the SAN to expose the LUNs.

Step 2: Review considerations for MPIO connectivity to virtual Fibre Channel storage

MPIO is used for load balancing and failover scenarios. Aggregating multiple active links provides an increase in throughput, reduces latency, and increases availability. Hyper-V in Windows Server 2012 R2 and Windows Server 2012 can use MPIO functionality to ensure continuous connectivity to Fibre Channel storage from within a virtual machine. As shown in Figure 3, MPIO routes network traffic through multiple network adapters.

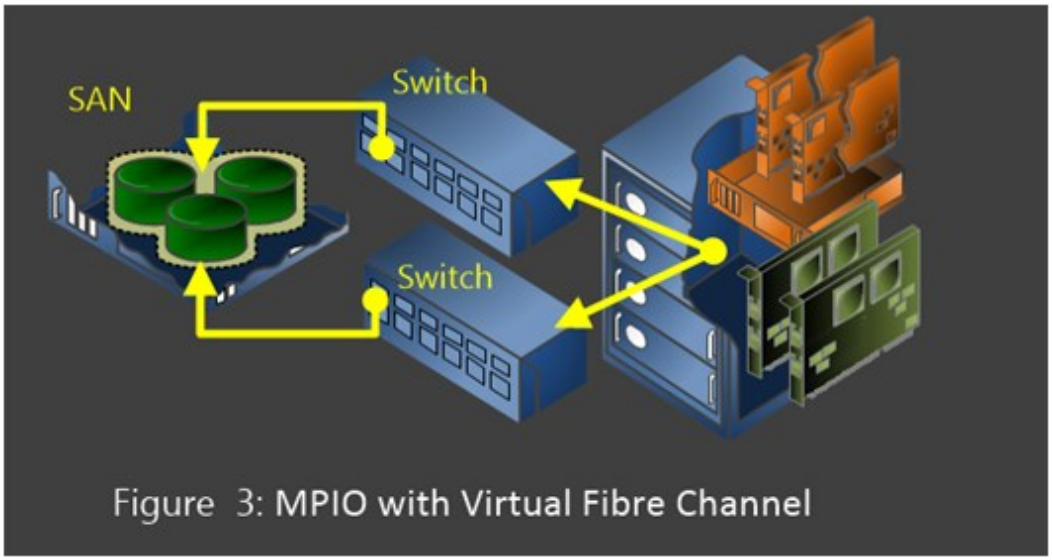

Figure 3: MPIO with Virtual Fibre Channel

MPIO on the virtual machine It is possible to configure multiple (up to four) virtual Fibre Channel adapters inside a virtual machine. You can then use the SAN vendor's recommended MPIO policy within the guest operating system of the virtual machine.

MPIO on the server running Hyper-V You can install multiple virtual Fibre Channel ports on the server running Hyper-V, and then use MPIO to provide high availability to the LUNs that are accessible by the server.

Step 3: Configure a virtual Fibre Channel adapter

The virtual Fibre Channel feature in Hyper-V makes it possible for a virtual machine to connect directly to a SAN through a virtualized Fibre Channel adapter. Establishing connectivity to your SAN by using a virtual Fibre Channel adapter is a two-step process. You create a virtual SAN, and then you add a virtual Fibre Channel adapter to your virtual machine and connect it to your vSAN.

A vSAN groups physical HBA ports together. You can create a virtual Fibre Channel SAN by using the following procedure.

To create a virtual SAN

1. On the Windows **Start** screen, click **Server Manager**.
2. Click **Tools**, and then click **Hyper-V Manager**.
3. In the **Action** pane, click **Virtual SAN Manager**.
4. To create a new vSAN, click **Create**.
5. In the **Name** text box, type the name for your new vSAN.
6. In the **Notes** text box, type any notes that you deem important for the vSAN.
7. Ensure you select at least one available WWN for the vSan to be used by the virtual Fibre Channel adapter.

8. Click **OK** to create the vSAN. You will be returned to the **Hyper-V Manager** screen.Step 2
9. Click **Virtual SAN Manager** to see your new vSAN.

Note

You can remove an existing vSAN by using the same basic procedure, but click **Remove** on the **Virtual SAN Manager** page.

A virtual Fibre Channel adapter is required in each virtual machine that requires Fibre Channel connectivity to the vSAN. Prior to implementing the following procedure, the virtual machine must be shut down.

To add a virtual Fibre Channel adapter and connect it to a vSAN

1. On the Windows **Start** screen, click **Server Manager**.
2. Click **Tools**, and then click **Hyper-V Manager**.
3. On the **Hyper-V Manager** screen, click the name of the desired virtual machine under **Virtual Machines**.
4. In the **Action** pane, click **Settings**.
5. Under **Add Hardware**, you will see a list of hardware options that you can add to your virtual machine. Click **Fibre Channel Adapter**, and then click **Add**.
6. On the **Settings** page, you must connect the virtual Fibre Channel adapter to a vSAN that you previously created. Click the **Virtual SAN** drop-down list that shows **Not connected**, and then click the name of the required vSan.
7. Configure your WWNs (Set A and Set B) for your virtual Fibre Channel adapter. You can use the default automatically generated WWN or manually assign them. You need to configure both WWN sets to ensure a successful live migration of your virtual machines.

   Note

   When your WWNs are assigned to the virtual Fibre Channel adapter, ensure that your LUNs and physical SAN are configured to recognize and work correctly with your virtual Fibre Channel adapter.
8. Click **OK**.

A virtual Fibre Channel adapter is now added to your virtual machine and connected to a virtual SAN. You can now turn on the virtual machine and connect to your SAN storage.

Terms and definitions

The following terms are commonly used when discussing virtual Fibre Channel:

- **MPIO** Multipath I/O (MPIO) software is used to provide load balancing and fault tolerance for multiple virtual Fibre Channel connections in the same server running Hyper-V. Native MPIO is available in Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008.

- **NPIV** Virtual Fibre Channel for Hyper-V guest operating systems uses the existing N_Port ID Virtualization (NPIV) T11 standard to map multiple virtual N_Port IDs to a single physical Fibre Channel N_port. A new NPIV port is created on the server running Hyper-V each time you start a virtual machine that is configured with a virtual HBA. When the virtual machine stops on the server running Hyper-V, the NPIV port is removed.
- **Virtual SAN** A virtual SAN (vSAN) defines a named group of physical Fibre Channel ports that are connected to the same physical SAN.
- **World Wide Names** A World Wide Port Name (WWPN) is a unique number that is provided to a virtual Fibre Channel HBA, similar to a MAC address. This unique key is used to allow the storage fabric to recognize a specific HBA. A World Wide Node Name (WWNN) is assigned to a node (an endpoint, a device) in a virtual Fibre Channel fabric.

See also

- Hyper-V Virtual Fibre Channel Troubleshooting Guide

The following posters provide a visual reference for understanding key Hyper-V technologies. They focus on Hyper-V Replica, networking, virtual machine mobility (live migration), storage, failover clustering, and scalability.

- Windows Server 2012 R2: Hyper-V Component Architecture
- Windows Server 2012: Hyper-V Component Architecture

Hyper-V Virtual Hard Disk Format Overview

As enterprise workloads for virtual environments grow in size and in performance demands, virtual hard disk (VHD) formats need to accommodate them. Hyper-V in Windows Server 2012 introduces a new version of the VHD format called VHDX, which is designed to handle current and future workloads.

VHDX has a much larger storage capacity than the older VHD format. It also provides data corruption protection during power failures and optimizes structural alignments of dynamic and differencing disks to prevent performance degradation on new, large-sector physical disks.

Key benefits

The new VHDX format in Windows Server 2012 addresses the technological demands of an evolving enterprise by increasing storage capacity, protecting data, and ensuring quality performance on large-sector disks.

Technical overview

The main new features of the VHDX format are:

- Support for virtual hard disk storage capacity of up to 64 TB.

- Protection against data corruption during power failures by logging updates to the VHDX metadata structures.
- Improved alignment of the virtual hard disk format to work well on large sector disks.

The VHDX format also provides the following features:

- Larger block sizes for dynamic and differencing disks, which allows these disks to attune to the needs of the workload.
- A 4-KB logical sector virtual disk that allows for increased performance when used by applications and workloads that are designed for 4-KB sectors.
- The ability to store custom metadata about the file that the user might want to record, such as operating system version or patches applied.
- Efficiency in representing data (also known as “trim”), which results in smaller file size and allows the underlying physical storage device to reclaim unused space. (Trim requires physical disks directly attached to a virtual machine or SCSI disks, and trim-compatible hardware.)

Hyper-V Support for Scaling Up and Scaling Out Overview

The use of virtualization has gained widespread acceptance as a means for lowering costs through consolidation of multiple server roles, typically hosting utility and productivity applications. These workloads generally consume comparatively little in the way of CPU and memory resources, and generate only moderate amounts of I/O. To support these workloads, configuring a virtual machine to use 2 or 4 virtual processors and relatively modest amounts of memory is sufficient. Hyper-V in Windows Server 2008 R2 supports configuring virtual machines with a maximum of 4 virtual processors, and up to 64 GB of memory.

However, IT organizations are increasingly looking to leverage virtualization to deploy mission-critical, tier-1 business applications. These larger and more demanding workloads—including high-end, online transaction processing (OLTP) databases and online transaction analysis (OLTA) business intelligence solutions—are typically run on systems with 16 or more processors, and demand large amounts of memory. For example, for SQL workloads of this type, a general practice is to allocate 8 GB of memory per logical processor. For this class of workloads, fewer virtual machines would typically be run on each virtualization host, but each virtual machine would require more virtual processors and greater amounts of virtual machine memory than current less demanding workloads.

As workloads demand greater and greater system resources to run these business-critical applications, highly scalable servers continue to expand the limits of processor core counts, and offer increased system memory capacity. With current processors providing up to 10 cores and 20 threads per socket, 8 socket systems are available that support up to 160 logical processors and up to 2 TB or more of memory. As multicore processors continue to evolve with increasing core counts, these high-end servers push the boundaries of scale support for operating systems and virtualization hosts.

Hyper-V in Windows Server® 2012 supports running on large host systems through expanded support for host processors and memory, and it enables the virtualization of high-performance, scale-up workloads by supporting the configuration of large, high-performance virtual machines.

Key benefits

Hyper-V in Windows Server 2012 provides support for 320 logical processors and 4 TB of system memory, which enables customers to run Hyper-V on the largest scale-up server systems currently available. Hyper-V enables customers to virtualize their most-demanding, mission-critical, tier-1 workloads by supporting large, high-performance virtual machines with up to 64 virtual processors and 1 TB of memory in a virtual machine. By projecting a virtual NUMA topology into large virtual machines, the guest operating system and applications such as SQL Server can leverage their existing thread scheduler and memory allocation optimizations, which ensures maximum performance and scalability of demanding workloads in a virtual machine. Support for SR-IOV–capable systems and network devices allows SR-IOV–capable network adapters to be assigned directly to a virtual machine, which maximizes network throughput while minimizing network latency as well as the CPU overhead required for processing network traffic.

Requirements

- One or more installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a server that is capable of running Hyper-V. Specifically, it must have processor support for hardware virtualization.
- The number of virtual processors that may be configured in a virtual machine depends on the number of processors on the physical computer. You must have at least as many logical processors in the virtualization host as the number of virtual processors required in the virtual machine. For example, to configure a virtual machine with the maximum of 64 virtual processors, you must be running Hyper-V on a virtualization host that has 64 or more logical processors.
- SR-IOV networking requires:
    - A host system which supports SR-IOV (for example, Intel VT-d), including chipset support for interrupt and DMA remapping, and proper firmware support to enable and describe the platform’s SR-IOV capabilities to the operating system.
    - An SR-IOV–capable network adapter and driver in both the management operating system (which runs the Hyper-V role) and each virtual machine where a virtual function is assigned.

Technical overview

Hyper-V in the Release Candidate of Windows Server 2012 supports running on large server systems and enables the virtualization of high-performance, scale-up workloads by including the following changes and features:

Increased hardware support for the virtualization host – Hyper-V now supports running on a host system with up to 320 logical processors and 4 TB of memory, which ensures compatibility with the largest scale-up server systems.

Support for large virtual machines – Hyper-V now supports configuring a virtual machine with up to 64 virtual processors and up to 1 TB of memory, which is a significant increase from previous versions.

NUMA support in a virtual machine – Non-Uniform Memory Architecture, or (NUMA), is a computer architecture used in multiprocessor systems in which the time required for a processor to access memory depends on the memory's location relative to the processor. With NUMA, a processor can access local memory (memory attached directly to the processor) faster than it can access remote memory (local to another processor in the system). Modern operating systems and high-performance applications such as Microsoft SQL Server have developed optimizations that recognize the system's NUMA topology and consider NUMA when scheduling threads or allocating memory to increase performance.

Projecting a virtual NUMA topology into a virtual machine enables optimal performance and workload scalability in large virtual machine configurations by allowing the guest operating system and applications such as SQL to leverage their inherent NUMA performance optimizations. The default virtual NUMA topology projected into a Hyper-V virtual machine is optimized to match the host's NUMA topology.

Note

If the virtual machine is configured to use Dynamic Memory, only one virtual NUMA node (that is, a flat NUMA topology) will be projected into the guest, which effectively disables virtual NUMA support.

Support for SR-IOV networking devices – Single Root I/O Virtualization (SR-IOV) is a standard introduced by the PCI-SIG. SR-IOV works in conjunction with system chipset support for virtualization technologies. This provides remapping of interrupts and DMA and allows SR-IOV capable devices to be assigned directly to a virtual machine. Hyper-V in Windows Server 2012 enables support for SR-IOV–capable network devices and allows an SR-IOV virtual function of a physical network adapter to be assigned directly to a virtual machine. This increases network throughput and reduces network latency, while also reducing the host CPU overhead required for processing network traffic.

Virtual Machine Storage Migration Overview

In Windows Server 2008 R2, you can move a running instance of a virtual machine using live migration, but you are not able to move the virtual machine's storage while the virtual machine is running.

Hyper-V in Windows Server 2012 introduces support for moving virtual machine storage without downtime by making it possible to move the storage while the virtual machine remains running.

You can perform this task by using a new wizard in Hyper-V Manager or by using new Hyper-V cmdlets for Windows PowerShell.

You can add storage to either a stand-alone computer or to a Hyper-V cluster, and then move virtual machines to the new storage while the virtual machines continue to run.

The most common reason for moving a virtual machine's storage is to update the physical storage that is available to Hyper-V. You can also move virtual machine storage between physical storage devices, at run time, to respond to reduced performance that results from bottlenecks in the storage throughput.

Key benefits

Hyper-V in Windows Server 2012 makes it possible to move virtual machine storage while a virtual machine is running.

Requirements

You need the following to use the Hyper-V functionality of moving virtual machine storage:

- One or more installations of Windows Server 2012 with the Hyper-V role installed.
- A server that is capable of running Hyper-V. Specifically, it must have processor support for hardware virtualization.
- Virtual machines that are configured to use only virtual hard disks for storage.

Note
You cannot move the storage of a virtual machine when any of its storage is directly attached to a physical disk.

Technical overview

This new feature allows you to move the virtual hard disks of a virtual machine while those virtual hard disks remain available for use by the running virtual machine. When you move a running virtual machine's virtual hard disks, Hyper-V performs the following steps, as shown in Figure 1:

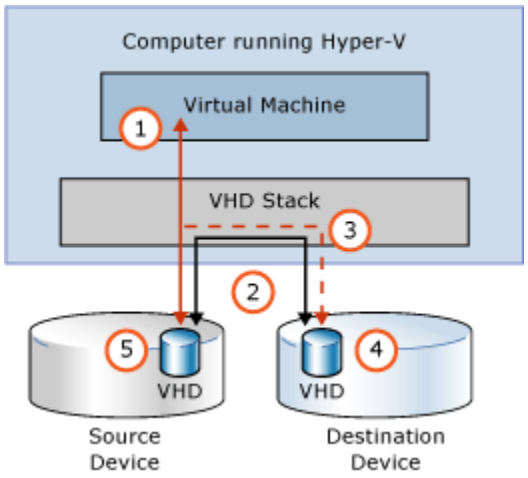

Figure 1 Moving virtual hard disks

1. Throughout most of the move operation, disk reads and writes go to the source virtual hard disk.
2. While reads and writes occur on the source virtual hard disk, the disk contents are copied to the new destination virtual hard disk.
3. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.
4. After the source and destination virtual hard disks are completely synchronized, the virtual machine switches over to using the destination virtual hard disk.
5. The source virtual hard disk is deleted.

Virtual Machine Live Migration Overview

Hyper-V live migration moves running virtual machines from one physical server to another with no impact on virtual machine availability to users. Hyper-V in Windows Server 2012 and Windows Server 2012 R2 introduce the following live migration improvements:

Faster and simultaneous migration.

In Windows Server 2012 live migrations are now able to utilize higher network bandwidths (up to 10 Gigabit) to complete migrations faster. You can also perform multiple simultaneous live migrations to enable you to move many virtual machines quickly. These changes allow you to implement high levels of mobility and flexibility in private cloud solutions.

In Windows Server 2012 R2 Hyper-V live migration has been updated to allow the administrator to select the optimal performance options when moving virtual machines to a different server. In larger scale deployments, such as private cloud deployments and hosting providers, this update can reduce the amount of time it takes to move a virtual machine to another server running Hyper-V. The reduction in time to move a virtual machine can also reduce overhead on the

network and CPU usage. For more information about live migration performance options, see Performance options.

Live migration outside of a clustered environment.

In Windows Server 2012, you can configure a virtual machine so that it is stored on an SMB file share. You can then perform a live migration on this running virtual machine between non-clustered servers running Hyper-V, while the virtual machine's storage remains on the central SMB share. This allows users to gain the benefits of virtual machine mobility without having to invest in the clustering infrastructure if they do not need guarantees of availability in their environment. (Hyper-V with SMB storage can also be configured with Failover Clustering if you do require high availability. For information about different configurations of Hyper-V with SMB storage, see Deploy Hyper-V over SMB.)

You can also perform a live migration of a virtual machine between two non-clustered servers running Hyper-V when you are only using local storage for the virtual machine. (This is sometimes referred to as a "shared nothing" live migration. In this case, the virtual machines storage is mirrored to the destination server over the network, and then the virtual machine is migrated, while it continues to run and provide network services. For instructions, see Configure Live Migration and Migrating Virtual Machines without Failover Clustering.

This functionality allows live migration in the most basic deployments and in more advanced scenarios, such as performing a live migration for a virtual machine between multiple, separate clusters to load balance across an entire data center.

Cross version live migration

Upgrading to a new version of Windows Server no longer requires downtime to the virtual machines. In Windows Server 2012 R2 Hyper-V live migration has been updated to support migrating Hyper-V virtual machines in Windows Server 2012 to Hyper-V in Windows Server 2012 R2.

When moving a virtual machine on Windows Server 2012, the specified destination server can now be a computer running Windows Server 2012 R2. This applies to a move that is initiated in Hyper-V Manager or when using the **Move-VM** Windows PowerShell cmdlet.

Note
Moving a Hyper-V virtual machine in Windows Server 2012 R2 to Hyper-V in Windows Server 2012 is not supported.

Key benefits

Live migration of virtual machines in Windows Server 2012 delivers improved performance and flexibility. It is also now available inside and outside of clustered environments—both with and without shared storage.

Requirements

Common requirements for any form of live migration:

- Two (or more) servers running Hyper-V that:
  - Support hardware virtualization.
  - Use processors from the same manufacturer. For example, all AMD or all Intel.
  - Belong to either the same Active Directory domain, or to domains that trust each other.
- Virtual machines must be configured to use virtual hard disks or virtual Fibre Channel disks (no physical disks). For more information on virtual Fibre Channel, see Hyper-V Virtual Fibre Channel Overview.
- Use of an isolated network, physically or through another networking technology such as VLANs, is recommended for live migration network traffic.

Requirements for live migration in a cluster:

- Windows Failover Clustering is enabled and configured.
- Cluster Shared Volume (CSV) storage in the cluster is enabled.

Important
Physical disks that are directly attached to a virtual machine (pass-through disks) are supported when all of the following conditions are met:

- The virtual machine configured with one or more physical disks is running on a Hyper-V failover cluster.
- The virtual machine configuration file is hosted on a Cluster Shared Volume (CSV).
- The physical disks are configured as a storage disk resource under control of the failover cluster and are properly configured as a dependent resource for the virtual machine.

Requirements for live migration using shared storage:

- All files that comprise a virtual machine (for example, virtual hard disks, snapshots, and configuration) are stored on an SMB share.
- Permissions on the SMB share have been configured to grant access to the computer accounts of all servers running Hyper-V.

Requirements for live migration with no shared infrastructure:

- No extra requirements exist.

Note
Physical disks that are directly attached to a virtual machine (pass-through disks) are not supported in live migration without shared storage (also known as “shared nothing” live migration).

Performance options

In larger scale deployments, such as private cloud deployments or cloud hosting providers, this update can reduce overhead on the network and CPU usage in addition to reducing the amount of time for a live migration. Hyper-V administrators can configure the appropriate live migration performance options based on their environment and requirements. Live migration performance options are configured on the host settings in the Hyper-V Manager console or via the **Set-VMHost** Windows PowerShell cmdlet and applies to all live migrations initiated from the host.

The following live migrations performance options are available.

Option	Description
TCP/IP	The memory of the virtual machine is copied to the destination server over a TCP/IP connection. This is the same method that is used in Hyper-V in Windows Server 2012.
Compression	The memory content of the virtual machine that is being migrated is compressed and then copied to the destination server over a TCP/IP connection. **Note** This is the **default** setting in Hyper-V in Windows Server 2012 R2.
SMB	The memory content of the virtual machine is copied to the destination server over a SMB 3.0 connection. • SMB Direct is used when the network adapters on the source and destination servers have Remote Direct Memory Access (RDMA) capabilities enabled. • SMB Multichannel automatically detects and uses multiple connections when a proper SMB Multichannel configuration is identified. For more information, see Improve Performance of a File Server with SMB Direct.

For information about how to modify the live migration performance options, see Step3: [optional] Configure performance options for live migration

Technical overview

Hyper-V live migration moves running virtual machines from one physical server to another with no impact on virtual machine availability to users. By pre-copying the memory of the migrating virtual machine to the destination server, live migration minimizes the transfer time of the virtual machine. A live migration is deterministic, which means that the administrator, or script, that initiates the live migration determines which computer is used as the destination for the live migration. The guest operating system of the migrating virtual machine is not aware that the migration is happening, so no special configuration for the guest operating system is needed.

After initiating a live migration, the following process occurs:

1. Live migration setup occurs.

During the live migration setup stage, the source server creates a connection with the destination server. This connection transfers the virtual machine configuration data to the destination server. A skeleton virtual machine is set up on the destination server and memory is allocated to the destination virtual machine.

2. Memory pages are transferred from the source node to the destination node.

In the second stage of a live migration, the memory assigned to the migrating virtual machine is copied over the network to the destination server. This memory is referred to as the “working set” of the migrating virtual machine. A page of memory is 4 KB.

For example, suppose that a virtual machine named “test virtual machine” configured with 1024 MB of RAM is migrating to another server running Hyper-V. The entire 1024 MB of RAM assigned to this virtual machine is the working set of “test virtual machine.” The utilized pages within the “test virtual machine” working set are copied to the destination server.

In addition to copying the working set of “test virtual machine” to the destination server, Hyper-V monitors the pages in the working set for “test virtual machine” on the source server. As memory pages are modified by “test virtual machine,” they are tracked and marked as being modified. The list of modified pages is simply the list of memory pages “test virtual machine” has modified after the copy of its working set has begun.

During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, with each iteration requiring a smaller number of modified pages to be copied. After the working set is copied to the destination server, the next stage of the live migration begins.

3. Modified pages are transferred.

The third stage of a live migration is a memory copy process that duplicates the remaining modified memory pages for “test virtual machine” to the destination server. The source server transfers the CPU and device state of the virtual machine to the destination server.

During this stage, the network bandwidth available between the source and destination servers is critical to the speed of the live migration. Using a 1 Gigabit Ethernet or faster is important. The faster the source server transfers the modified pages from the migrating virtual machines working set, the more quickly the live migration is completed.

The number of pages transferred in this stage is determined by how actively the virtual machine accesses and modifies the memory pages. The more modified pages there are, the longer it takes to transfer all pages to the destination server.

After the modified memory pages are copied completely to the destination server, the destination server has an up-to-date working set for “test virtual machine.” The working set for “test virtual machine” is present on the destination server in the exact state it was in when “test virtual machine” began the migration process.

Note

You can cancel the live migration process at any point before this stage of the migration.

4. The storage handle is moved from the source server to the destination server.

During the fourth stage of a live migration, control of the storage associated with “test virtual machine,” such as any virtual hard disk files or physical storage attached through a virtual Fibre Channel adapter, is transferred to the destination server. (Virtual Fibre Channel is also a new Hyper-V feature in Windows Server 2012.)

5. The virtual machine is brought online on the destination server.

In the fifth stage of a live migration, the destination server now has the up-to-date working set for “test virtual machine,” as well as access to any storage used by “test virtual machine.” At this point “test virtual machine” is resumed.

6. Network cleanup occurs.

In the final stage of a live migration, the migrated virtual machine is running on the destination server. At this point, a message is sent to the network switch. This message causes the network switch to obtain the new the MAC addresses of the migrated virtual machine so that network traffic to and from “test virtual machine” can use the correct switch port.

The live migration process completes in less time than the TCP time-out interval for the virtual machine being migrated. TCP time-out intervals vary based on network topology and other factors. The following variables may affect live migration speed:

- The number of modified pages on the virtual machine to be migrated—the larger the number of modified pages, the longer the virtual machine will remain in a migrating state.
- Available network bandwidth between source and destination servers.
- Hardware configuration of source and destination servers.
- Load on source and destination servers.
- Available bandwidth (network or Fibre Channel) between servers running Hyper-V and shared storage.

The live migration process for a virtual machine inside a cluster (when the virtual machine is stored on a CSV volume) and for a virtual machine outside of a cluster (when the virtual machine is stored on an SMB share) is practically identical.

When performing a live migration of a virtual machine between two computers with no shared infrastructure, the first thing that Hyper-V does is perform a partial migration of the virtual machines storage, as follows:

1. Throughout most of the move operation, disk reads and writes go to the source virtual hard disk.
2. While reads and writes occur on the source virtual hard disk, the disk contents are copied over the network to the new destination virtual hard disk.
3. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.
4. After the source and destination virtual hard disks are completely synchronized, the virtual machine live migration is initiated, following the same process that is used for live migration with shared storage.
5. Once the live migration is complete and the virtual machine is successfully running on the destination server, the files on the source server are deleted.

See Also

Configure Live Migration and Migrating Virtual Machines without Failover Clustering

Simplified Import Overview

Administrators often think of a virtual machine as a single, stand-alone entity that they can move around to address their operational needs. However, a virtual machine consists of several parts, which administrators do not normally need to think about:

- Virtual hard disks, stored as files on the physical storage.
- Virtual machine snapshots, stored as a special type of virtual hard disk file.
- The saved state of the different, host-specific devices.
- The memory file for the virtual machine or its snapshot.
- The virtual machine configuration file, which organizes all of those parts and arranges them into a working virtual machine.

Each virtual machine and every snapshot associated with it must be unique, so globally unique identifiers are used. Additionally, virtual machines store and use some host-specific information, such as the path information for virtual hard disk files. When Hyper-V tries to start a virtual machine, it goes through a series of validation checks before being started. Problems such as hardware differences that might exist when a virtual machine is moved to another host can cause these validation checks to fail. That, in turn, prevents the virtual machine from starting. The administrator is left with files on the disk that take up space and are not useful.

Hyper-V in Windows Server 2012 introduces a new Import wizard that detects and fixes more than 40 different types of incompatibilities. The Import wizard walks you through the steps of addressing incompatibilities when you import the virtual machine to the new host—so this wizard can help with configuration that is associated with physical hardware, such as memory, virtual switches, and virtual processors.

Also, you no longer need to export a virtual machine to be able to import it. You can simply copy a virtual machine and its associated files to the new host, and then use the Import wizard to specify the location of the files. This “registers” the virtual machine with Hyper-V and makes it available for use. You can copy a virtual machine to an NTFS-formatted USB drive, and you can recover virtual machines in cases where the system drive fails but the data drive that stores the virtual machines is intact.

In addition to the new wizard, automation support is available. The new Hyper-V module for Windows PowerShell includes cmdlets for importing virtual machines.

Key benefits

The new Import wizard in Windows Server 2012 provides a simpler, improved way to import or copy virtual machines.

Requirements

To try out the Import wizard, you will need the following:

- Two installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer that has processor support for hardware virtualization.
- A virtual machine.
- A user account that belongs to the local Hyper-V Administrators group.

Technical overview

To import a virtual machine, the wizard does the following:

1. **Creates a copy of the virtual machine configuration file.** This is done as a precaution in case an unexpected restart occurs on the host, such as from a power outage.
2. **Validates hardware.** Information in the virtual machine configuration file is compared to hardware on the new host.
3. **Compiles a list of errors.** This list indentifies what needs to be reconfigured and determines which pages appear next in the wizard.
4. **Displays the relevant pages, one category at a time.** The wizard explains each incompatibility to help you reconfigure the virtual machine so it is compatible with the new host.
5. **Removes the copy of the configuration file.** After the wizard does this, the virtual machine is ready to be started.

Hyper-V Network Virtualization Overview

In Windows Server 2012 R2 and System Center 2012 R2 Virtual Machine Manager, Microsoft provides an end-to-end network virtualization solution. There are four major components that comprise Microsoft's network virtualization solution:

- **Windows Azure Pack for Windows Server** provides a tenant facing portal to create virtual networks.
- **System Center 2012 R2 Virtual Machine Manager** (VMM) provides centralized management of the virtual networks.
- **Hyper-V Network Virtualization** provides the infrastructure needed to virtualize network traffic.
- **Hyper-V Network Virtualization gateways** provides connections between virtual and physical networks.

For more information about a network virtualization solution that uses all of these components, see **Hybrid Cloud Multi-Tenant Networking Solution Guide**.

This topic introduces concepts and explains the key benefits and capabilities of Hyper-V Network Virtualization (one part of the overall network virtualization solution) in Windows Server 2012 R2.

It explains how network virtualization benefits both private clouds looking for enterprise workload consolidation and public cloud service providers of Infrastructure as a Service (IaaS).

For more technical details about networking virtualization, see Hyper-V Network Virtualization technical details.

Did you mean…

- **Hyper-V FAQ**
- Hyper-V Virtual Switch Overview

Feature description

Hyper-V Network Virtualization provides “virtual networks” (called a VM network) to virtual machines similar to how server virtualization (hypervisor) provides “virtual machines” to the operating system. Network virtualization decouples virtual networks from the physical network infrastructure and removes the constraints of VLAN and hierarchical IP address assignment from virtual machine provisioning. This flexibility makes it easy for customers to move to IaaS clouds and efficient for hosters and datacenter administrators to manage their infrastructure, while maintaining the necessary multi-tenant isolation, security requirements, and supporting overlapping Virtual Machine IP addresses.

Customers want to seamlessly extend their datacenters to the cloud. Today there are technical challenges in making such seamless hybrid cloud architectures. One of the biggest hurdles customers face is reusing their existing network topologies (subnets, IP addresses, network services, and so on.) in the cloud and bridging between their on-premise resources and their cloud resources. Hyper-V Network Virtualization provides the concept of a VM Network that is independent of the underlying physical network. With this concept of a VM Network, composed of one or more Virtual Subnets, the exact location in the physical network of virtual machines attached to a virtual network is decoupled from the virtual network topology. As a result, customers can easily move their virtual subnets to the cloud while preserving their existing IP addresses and topology in the cloud so that existing services continue to work unaware of the physical location of the subnets. That is, Hyper-V Network Virtualization enables a seamless hybrid cloud.

In addition to hybrid cloud, many organizations are consolidating their datacenters and creating private clouds to internally get the efficiency and scalability benefit of cloud architectures. Hyper-V Network Virtualization allows better flexibility and efficiency for private clouds by decoupling a business unit’s network topology (by making it virtual) from the actual physical network topology. In this way, the business units can easily share an internal private cloud while being isolated from each other and continue to keep existing network topologies. The datacenter operations team has flexibility to deploy and dynamically move workloads anywhere in the datacenter without server interruptions providing better operational efficiencies and an overall more effective datacenter.

For workload owners, the key benefit is that they can now move their workload “topologies” to the cloud without changing their IP addresses or re-writing their applications. For example, the typical three-tier LOB application is composed of a front end tier, a business logic tier, and a database tier. Through policy, Hyper-V Network Virtualization allows customer onboarding all or parts of the

three tiers to the cloud, while keeping the routing topology and the IP addresses of the services (i.e. virtual machine IP addresses), without requiring the applications to be changed.

For infrastructure owners, the additional flexibility in virtual machine placement makes it possible to move workloads anywhere in the datacenters without changing the virtual machines or reconfiguring the networks. For example Hyper-V Network Virtualization enables cross subnet live migration so that a virtual machine can live migrate anywhere in the datacenter without a service disruption. Previously live migration was limited to the same subnet restricting where virtual machines could be located. Cross subnet live migration allows administrators to consolidate workloads based on dynamic resource requirements, energy efficiency, and can also accommodate infrastructure maintenance without disrupting customer workload up time.

Practical applications

With the success of virtualized datacenters, IT organizations and hosting providers (providers who offer colocation or physical server rentals) have begun offering more flexible virtualized infrastructures that make it easier to offer on-demand server instances to their customers. This new class of service is referred to as Infrastructure as a Service (IaaS). Windows Server 2012 R2 provides all the required platform capabilities to enable enterprise customers to build private clouds and transition to an IT as a service operational model. Windows Server 2012 R2 also enables hosters to build public clouds and offer IaaS solutions to their customers. When combined with Virtual Machine Manager to manage Hyper-V Network Virtualization policy, Microsoft provides a powerful cloud solution.

Windows Server 2012 R2 Hyper-V Network Virtualization provides policy-based, software-controlled network virtualization that reduces the management overhead faced by enterprises when they expand dedicated IaaS clouds, and it provides cloud hosters better flexibility and scalability for managing virtual machines to achieve higher resource utilization.

An IaaS scenario that has virtual machines from different organizational divisions (dedicated cloud) or different customers (hosted cloud) requires secure isolation. Today's solution, virtual local area networks (VLANs), can present significant disadvantages in this scenario.

VLANs Currently, VLANs are the mechanism that most organizations use to support address space reuse and tenant isolation. A VLAN uses explicit tagging (VLAN ID) in the Ethernet frame headers, and it relies on Ethernet switches to enforce isolation and restrict traffic to network nodes with the same VLAN ID. The main disadvantages with VLANs are as follows:

- Increased risk of an inadvertent outage due to cumbersome reconfiguration of production switches whenever virtual machines or isolation boundaries move in the dynamic datacenter.
- Limited in scalability because there is a maximum of 4094 VLANs and typical switches support no more than 1000 VLAN IDs.
- Constrained within a single IP subnet, which limits the number of nodes within a single VLAN and restricts the placement of virtual machines based on physical locations. Even though VLANs can be expanded across sites, the entire VLAN must be on the same subnet.

IP address assignment In addition to the disadvantages that are presented by VLANs, virtual machine IP address assignment presents issues, which include:

- Physical locations in datacenter network infrastructure determine virtual machine IP addresses. As a result, moving to the cloud typically requires changing IP addresses of the service workloads.
- Policies are tied to IP addresses, such as firewall rules, resource discovery and directory services, and so on. Changing IP addresses requires updating all the associated policies.
- Virtual machine deployment and traffic isolation are dependent on the topology.

When datacenter network administrators plan the physical layout of the datacenter, they must make decisions about where subnets will be physically placed and routed. These decisions are based on IP and Ethernet technology that influence the potential IP addresses that are allowed for virtual machines running on a given server or a blade that is connected to a particular rack in the datacenter. When a virtual machine is provisioned and placed in the datacenter, it must adhere to these choices and restrictions regarding the IP address. Therefore, the typical result is that the datacenter administrators assign new IP addresses to the virtual machines.

The problem with this requirement is that in addition to being an address, there is semantic information associated with an IP address. For instance, one subnet may contain given services or be in a distinct physical location. Firewall rules, access control policies, and IPsec security associations are commonly associated with IP addresses. Changing IP addresses forces the virtual machine owners to adjust all their policies that were based on the original IP address. This renumbering overhead is so high that many enterprises choose to deploy only new services to the cloud, leaving legacy applications alone.

Hyper-V Network Virtualization decouples virtual networks for customer virtual machines from the physical network infrastructure. As a result, it enables customer virtual machines to maintain their original IP addresses, while allowing datacenter administrators to provision customer virtual machines anywhere in the datacenter without reconfiguring physical IP addresses or VLAN IDs. The next section summarizes the key functionality.

Important functionality

The following is a list of the key functionality, benefits, and capabilities of Hyper-V Network Virtualization in Windows Server 2012 R2:

- **Enables flexible workload placement – Network isolation and IP address re-use without VLANs**

  Hyper-V Network Virtualization decouples the customer's virtual networks from the physical network infrastructure of the hosters, providing freedom for workload placements inside the datacenters. Virtual machine workload placement is no longer limited by the IP address assignment or VLAN isolation requirements of the physical network because it is enforced within Hyper-V hosts based on software-defined, multitenant virtualization policies.

  Virtual machines from different customers with overlapping IP addresses can now be deployed on the same host server without requiring cumbersome VLAN configuration or violating the IP address hierarchy. This can streamline the migration of customer workloads into shared IaaS hosting providers, allowing customers to move those workloads without modification, which includes leaving the virtual machine IP addresses unchanged. For the

hosting provider, supporting numerous customers who want to extend their existing network address space to the shared IaaS datacenter is a complex exercise of configuring and maintaining isolated VLANs for each customer to ensure the coexistence of potentially overlapping address spaces. With Hyper-V Network Virtualization, supporting overlapping addresses is made easier and requires less network reconfiguration by the hosting provider.

In addition, physical infrastructure maintenance and upgrades can be done without causing a down time of customer workloads. With Hyper-V Network Virtualization, virtual machines on a specific host, rack, subnet, VLAN, or entire cluster can be migrated without requiring a physical IP address change or major reconfiguration.

- **Enables easier moves for workloads to a shared IaaS cloud**

  With Hyper-V Network Virtualization, IP addresses and virtual machine configurations remain unchanged. This enables IT organizations to more easily move workloads from their datacenters to a shared IaaS hosting provider with minimal reconfiguration of the workload or their infrastructure tools and policies. In cases where there is connectivity between two datacenters, IT administrators can continue to use their tools without reconfiguring them.

- **Enables live migration across subnets**

  Live migration of virtual machine workloads traditionally has been limited to the same IP subnet or VLAN because crossing subnets required the virtual machine's guest operating system to change its IP address. This address change breaks existing communication and disrupts the services running on the virtual machine. With Hyper-V Network Virtualization, workloads can be live migrated from servers running Windows Server 2012 in one subnet to servers running Windows Server 2012 in a different subnet without changing the workload IP addresses. Hyper-V Network Virtualization ensures that virtual machine location changes due to live migration are updated and synchronized among hosts that have ongoing communication with the migrated virtual machine.

- **Enables easier management of decoupled server and network administration**

  Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, and network administrators can focus on overall network infrastructure and traffic management. This enables datacenter server administrators to deploy and migrate virtual machines without changing the IP addresses of the virtual machines. There is reduced overhead because Hyper-V Network Virtualization allows virtual machine placement to occur independently of network topology, reducing the need for network administrators to be involved with placements that might change the isolation boundaries.

- **Simplifies the network and improves server/network resource utilization**

  The rigidity of VLANs and the dependency of virtual machine placement on a physical network infrastructure results in overprovisioning and underutilization. By breaking the dependency, the increased flexibility of virtual machine workload placement can simplify the network management and improve server and network resource utilization. Note that Hyper-V Network Virtualization supports VLANs in the context of the physical datacenter. For

example, a datacenter may want all Hyper-V Network Virtualization traffic to be on a specific VLAN.

- **Is compatible with existing infrastructure and emerging technology**

  Hyper-V Network Virtualization can be deployed in today's datacenter, yet it is compatible with emerging datacenter "flat network" technologies.

- **Provides for interoperability and ecosystem readiness**

  Hyper-V Network Virtualization supports multiple configurations for communication with existing resources, such as cross premise connectivity, storage area network (SAN), non-virtualized resource access, and so on. Microsoft is committed to working with ecosystem partners to support and enhance the experience of Hyper-V Network Virtualization in terms of performance, scalability, and manageability.

- **Uses Windows PowerShell and WMI**

  Hyper-V Network Virtualization supports Windows PowerShell and Windows Management Instrumentation (WMI) for configuring the network virtualization and isolation policies. The Windows PowerShell cmdlets for Hyper-V Network Virtualization enable administrators to build command-line tools or automated scripts to configure, monitor, and troubleshoot network isolation policies.

Software requirements

Hyper-V Network Virtualization requires Windows Server 2012 or Windows Server 2012 R2 and the Hyper-V role.

See also

To learn more about Hyper-V Network Virtualization see the following links:

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2 • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2)
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**

Content type	References
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide • Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	• Windows Server Gateway • Hyper-V Virtual Switch Overview

What's New in Hyper-V Network Virtualization in Windows Server 2012 R2

This topic describes the new or changed features and functionality for Hyper-V Network Virtualization (HNV) in Windows Server 2012 R2.

Feature description

HNV provides a virtual network abstraction on top of a physical network. This abstraction provides virtual networks to virtual machines which operate the same as a physical network. This is similar to the abstraction hypervisors provide to the operating system running in virtual machines. HNV provides this abstraction through an overlay network on top of the physical network for each VM network. For a HNV overview, see Hyper-V Network Virtualization Overview.

New and changed functionality

The following HNV features have been added or updated for Windows Server 2012 R2.

Feature/functionality	Windows Server 2012	Windows Server 2012 R2
Inbox HNV Gateway		X
HNV interoperability with Hyper-V Virtual Switch Extensions	X	X
HNV VM Network Diagnostics	X	X

Feature/functionality	Windows Server 2012	Windows Server 2012 R2
HNV Architecture	X	X
Dynamic IP Address Learning		X
HNV + Windows NIC Teaming	X	X
NVGRE Encapsulated Task Offload	X	X

Inbox HNV Gateway

The inbox HNV gateway is a multi-tenant gateway that performs Site-to-Site (VPN), NAT, and Forwarding functions.

What value does this change add?

It is now easier to set up a gateway which can connect multiple tenant VPN connections in a hybrid cloud scenario, supports multiple tenants connecting to the Internet, and forwards network traffic from a datacenter network to virtual networks in a private cloud scenario.

What works differently?

- You can use System Center 2012 R2 Virtual Machine Manager to fully manage the HNV gateway.
- Supports guest clustering for high availability
- Includes BGP for dynamic routes update

HNV Architecture

The HNV filter moved from being an NDIS lightweight filter (LWF) to being part of the Hyper-V virtual switch.

What value does this change add?

Forwarding switch extensions can co-exist with HNV allowing multiple network virtualization solutions (one provided by HNV and another provided by the forwarding switch extension) to co-exist on the same host running Hyper-V.

What works differently?

- Improved interoperability with switch extensions
- The HNV NDIS LWF does not have to be bound to network adapters anymore. Once you attach a network adapter to the virtual switch you can enable HNV simply by assigning a Virtual Subnet ID to a particular virtual network adapter. For those using Virtual Machine Manager to manage VM networks this is transparent, but for anyone using Windows PowerShell this will save an often-missed step.

HNV interoperability with Hyper-V Virtual Switch Extensions

Switch extensions work in both the Customer Address (CA) and Provider Address (PA) space.

What value does this change add?

Third-party switch extensions can now work with HNV traffic, because the extensions have visibility into both the provider address (PA) space, and the customer address (CA) space.

What works differently?

- The HNV module was moved to inside the virtual switch so that extensions can see both the provider (PA) and virtual (CA) IP address spaces. This allows forwarding and other types of extensions to make decisions with knowledge of both address spaces.
- Hybrid forwarding is implemented. Hybrid forwarding directs packets to different forwarding agents, based upon the packet type. In the Windows Server 2012 R2 implementation, an NVGRE packet is forwarded by the HNV module. A packet that is not NVGRE is forwarded normally by the forwarding extension.

HNV VM Network Diagnostics

Some new diagnostic tools have been included.

What value does this change add?

This enhances your ability to diagnose HNV networks.

What works differently?

- Enhanced ping.exe (ping –p) to allow pinging to and from provider addresses
- Two new Windows PowerShell cmdlets (**Test-VMNetworkAdapter** and **Select-NetVirtualizationNextHop**) that enables diagnostics of HNV policy and the Customer Address space.
- Added the ability for Message Analyzer to decode NVGRE packets

For more details, check out the New Networking Diagnostics with PowerShell in Windows Server R2 blog post.

Dynamic IP Address Learning

HNV learns about the IP address of a virtual machine that has been assigned manually or set via DHCP on the virtual network.

What value does this change add?

- Enables high availability scenarios for both virtual machines on a VM network and the HNV gateway.
- Allows you to run DHCP, DNS, and Active Directory in your VM networks.

What works differently?

- For broadcast or multicast packets in a VM network, a PA multicast IP address is used if configured. If a PA multicast address is not available an intelligent PA unicast replication is used.

Packets are unicasted only to PA addresses that are configured for the particular virtual subnet the packet is on. In addition, only one unicast packet is sent per host no matter how many relevant virtual machines are on the host.

- Once a host learns a new IP address it will then notify Virtual Machine Manager. At this point, the learned IP address becomes part of the centralized policy that Virtual Machine Manager pushes out. This allows for both rapid dissemination of HNV routing policy and limits the network overhead for disseminating this HNV routing policy.
- Includes support for Duplicate Address Detection (DAD), network unreachability detection (NUD) and Address Resolution Protocol (ARP) packets in the CA address space for both IPv4 and IPv6. The HNV filter also provides a reliable ARP proxy for any known routing policies once again reducing the amount of control traffic that goes out on the physical network.
- Fully supported in Windows PowerShell.

HNV + Windows NIC Teaming

Integrates HNV and Windows NIC Teaming to allow multiple network adapters to be placed into a team for the purposes of bandwidth aggregation, and/or traffic failover to maintain connectivity in the event of a network component failure.

What value does this change add?

Integrating HNV with Windows NIC Teaming increases HNV network throughput and reliability.

What works differently?

Both inbound and outbound spread of virtualized traffic on a NIC Team is enabled. This means that both traffic leaving a host or coming into a host can utilize all the network adapters in the NIC team.

NVGRE Encapsulated Task Offload

In NDIS 6.30 and later (Windows Server 2012 and later), Network Virtualization using Generic Routing Encapsulation (NVGRE) task offload makes it possible to use Generic Routing Encapsulation (GRE)-encapsulated packets with:

- Large Send Offload (LSO)
- Receive Side Scaling (RSS)
- Virtual Machine Queue (VMQ)

What value does this change add?

To increase its performance, NVGRE can offload tasks to a network adapter that has the appropriate task offload capabilities.

What works differently?

Two partners have announced that their next generation network adapters will support NVGRE Encapsulated Task Offload. You can read the press releases from Mellanox and Emulex for more details.

Microsoft is continuing to work with additional network vendors to enable NVGRE Task Offload. More announcements will be made in the future.

See also

The following table provides additional resources for evaluating HNV.

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2 • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2) • Test Lab Guide: Windows Server 2012 R2 Hyper-V Network Virtualization with System Center 2012 R2 VMM
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide • Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	Hyper-V Virtual Switch Overview

Hyper-V Network Virtualization technical details

Server virtualization enables multiple server instances to run concurrently on a single physical host; yet server instances are isolated from each other. Each virtual machine essentially operates

as if it is the only server running on the physical computer. Network virtualization provides a similar capability, in which multiple virtual network infrastructures run on the same physical network (potentially with overlapping IP addresses), and each virtual network infrastructure operates as if it is the only virtual network running on the shared network infrastructure. Figure 1 shows this relationship.

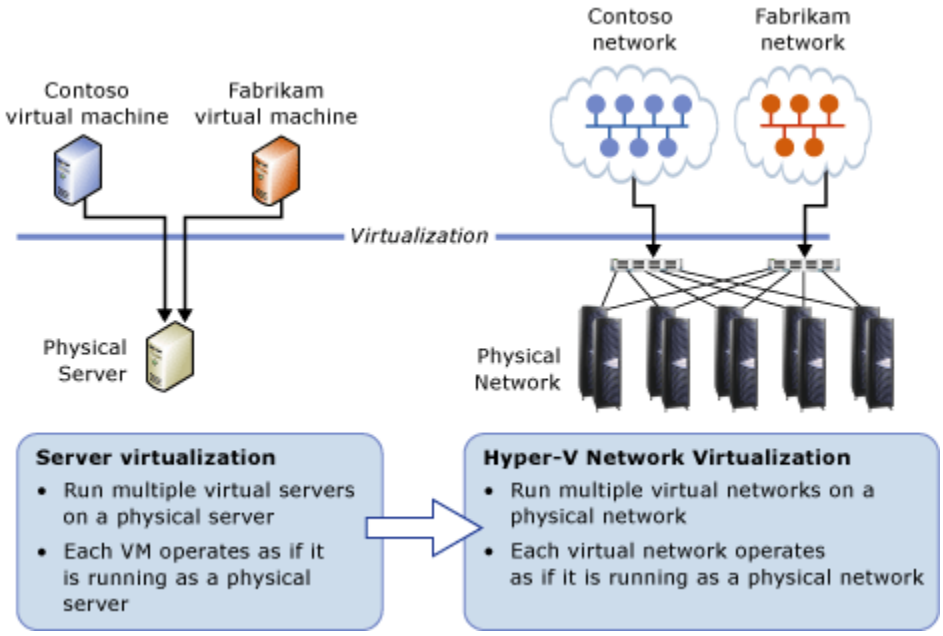

Figure 1: Server virtualization versus network virtualization

Hyper-V Network Virtualization Concepts

In Hyper-V Network Virtualization (HNV), a customer is defined as the “owner” of a group of virtual machines that are deployed in a datacenter. A customer can be a corporation or enterprise in a multitenant public datacenter, or a division or business unit within a private datacenter. Each customer can have one or more VM networks in the datacenter, and each VM network consists of one or more virtual subnets.

VM network

- Each VM network consists of one or more virtual subnets. A VM network forms an isolation boundary where the virtual machines within a VM network can communicate with each other. As a result, virtual subnets in the same VM network must not use overlapping IP address prefixes.
- Each VM network has a Routing Domain which identifies the VM network. The Routing Domain ID (RDID), which identifies the VM network, is assigned by datacenter administrators or datacenter management software, such as System Center 2012 R2 Virtual Machine Manager (VMM). The RDID is a Windows GUID — for example, “{11111111-2222-3333-4444-000000000000}”.

Virtual subnets

- A virtual subnet implements the Layer 3 IP subnet semantics for the virtual machines in the same virtual subnet. The virtual subnet is a broadcast domain (similar to a VLAN). Virtual machines in the same virtual subnet must use the same IP prefix.
- Each virtual subnet belongs to a single VM network (RDID), and it is assigned a unique Virtual Subnet ID (VSID). The VSID must be unique within the datacenter and is in the range 4096 to 2^24-2).

A key advantage of the VM network and routing domain is that it allows customers to bring their network topologies to the cloud. Figure 2 shows an example where the Contoso Corp has two separate networks, the R&D Net and the Sales Net. Because these networks have different routing domain IDs, they cannot interact with each other. That is, Contoso R&D Net is isolated from Contoso Sales Net even though both are owned by Contoso Corp. Contoso R&D Net contains three virtual subnets. Note that both the RDID and VSID are unique within a datacenter.

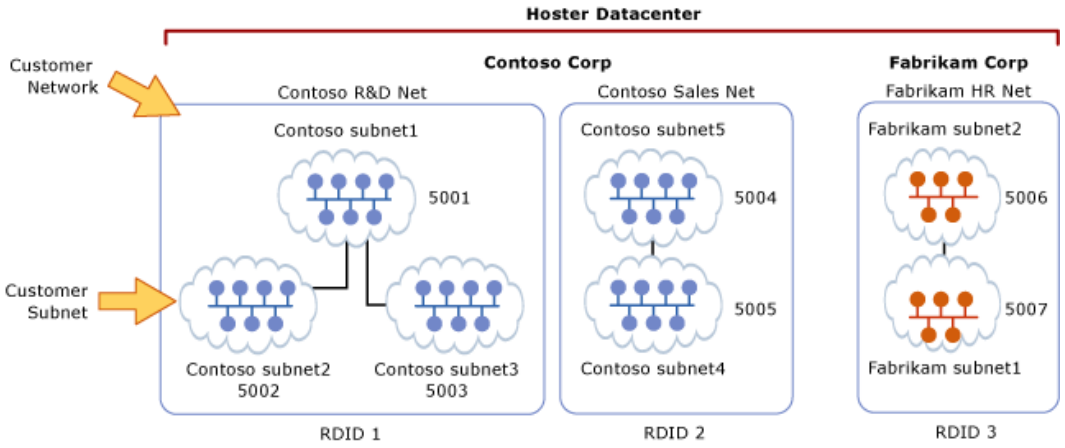

Figure 2: Customer networks and virtual subnets

In Figure 2, the virtual machines with VSID 5001 can have their packets routed or forwarded by HNV to virtual machines with VSID 5002 or VSID 5003. Before delivering the packet to the Hyper-V switch, HNV will update the VSID of the incoming packet to the VSID of the destination virtual machine. This will only happen if both VSIDs are in the same RDID. If the VSID that is associated with the packet does not match the VSID of the destination virtual machine, the packet will be dropped. Therefore, virtual network adapters with RDID1 cannot send packets to virtual network adapters with RDID2.

Note

In the packet flow description above, the term “virtual machine” actually means the “virtual network adapter” on the virtual machine. The common case is that a virtual machine only has a single virtual network adapter. In this case, the words virtual machine and virtual network adapter can conceptually mean the same thing. Because a virtual machine can have multiple virtual network adapters, and these virtual network adapters can have different VirtualSubnetIDs (VSIDs) or RoutingDomainIDs (RDIDs), HNV specifically focuses on the packets sent and received between virtual network adapters.

Each virtual subnet defines a Layer 3 IP subnet and a Layer 2 (L2) broadcast domain boundary similar to a VLAN. When a virtual machine broadcasts a packet, this broadcast is limited to the virtual machines that are attached to switch ports with the same VSID. Each VSID can be associated with a multicast address in the PA. All broadcast traffic for a VSID is sent on this multicast address.

Note

HNV does NOT depend on broadcast or multicast. For broadcast or multicast packets in a VM network, a PA multicast IP address is used if configured. However, the many datacenter operators do not enable multicast in their environments. As a result, when a PA multicast address is not available an intelligent PA unicast replication is used. This means that packets are unicasted only to PA addresses that are configured for the particular virtual subnet the packet is on. In addition, only one unicast packet per host is sent no matter how many relevant virtual machines are on the host.

In addition to being a broadcast domain, the VSID provides isolation. A virtual network adapter in HNV is connected to a Hyper-V switch port that has a VSID ACL. If a packet arrives on this Hyper-V switch port with a different VSID the packet is dropped. Packets will only be delivered on a Hyper-V switch port if the VSID of the packet matches the VSID of the switch port. This is the reason, in the above example of Figure 2, that packets flowing from VSID 5001 to 5003 must have the VSID in the packet modified before delivery to the destination virtual machine.

If the Hyper-V switch port does not have a VSID ACL, the virtual network adapter that is attached to that switch port is not part of a HNV virtual subnet. Packets sent from a virtual network adapter that does not have a VSID ACL will pass unmodified through the Hyper-V switch.

When a virtual machine sends a packet, the VSID of the Hyper-V switch port is associated with this packet. On the receiving side, HNV delivers to the Hyper-V switch the VSID in the OOB along with the decapsulated packet. On the receiving end, HNV performs a policy lookup and adds the VSID to the OOB data before the packet is passed to the Hyper-V switch.

Note

Hyper-V Switch Extensions can operate in both the Provider Address (PA) space and the Customer Address (CA) space. This means the VSID is available to the switch extensions. This allows the switch extension to become multitenant aware. For example, a firewall switch extension can differentiate CA IP address 10.1.1.5 with OOB containing VSID 5001 from the same CA IP address with VSID 6001.

Routing in Hyper-V Network Virtualization

As in physical networks, routing is an important part of HNV. There are two key aspects to understand: how packets are routed between virtual subnets and how packets are routed outside a virtual network.

Routing Between Virtual Subnets

In a physical network, a subnet is the Layer 2 (L2) domain where computers (virtual and physical) can directly communicate with each other without having to be routed. In Windows, if you statically configure a network adapter you can set a “default gateway” which is the IP address to send all traffic that is going out of the particular subnet so that it can be routed appropriately. This is typically the router for your physical network. HNV uses a built in router that is part of every host to form a distributed router for a virtual network. This means that every host, in particular the Hyper-V Virtual Switch, acts as the default gateway for all traffic that is going between Virtual Subnets that are part of the same VM network. In Windows Server 2012 and Windows Server 2012 R2 the address used as the default gateway is the lowest entry for the subnet (as an example, it is the “.1” address for a /24 subnet prefix). This address is reserved in each virtual subnet for the default gateway and cannot be used by virtual machines in the virtual subnet.

HNV acting as a distributed router allows for a very efficient way for all traffic inside a VM Network to be routed appropriately because each host can directly route the traffic to the appropriate host without needing an intermediary. This is particularly true when two virtual machines in the same VM Network but different Virtual Subnets are on the same physical host. As you will see later in this section, the packet never has to leave the physical host.

Routing Outside a Virtual Network

Most customer deployments will require communication from the HNV environment to resources that are not part of the HNV environment. Network Virtualization gateways are required to allow communication between the two environments. Scenarios requiring a HNV Gateway include Private Cloud and Hybrid Cloud. Basically, HNV gateways are required for VPNs and routing.

Gateways can come in different physical form factors. They can be built on Windows Server 2012 R2, incorporated into a Top of Rack (TOR) switch, a load balancer, put into other existing network appliances, or can be a new stand-alone network appliance.

Note

For more information about gateways, see Hyper-V Network Virtualization Gateway Architectural Guide and Simple demo with gateway script.

Private Cloud (Routing)

Large enterprises may be either hesitant, or for compliance reasons, unable to move some of their services and data to a public cloud hoster. However, enterprises still want to obtain the benefits of the cloud provided by HNV by consolidating their datacenter resources into a private cloud. In a private cloud deployment, overlapping IP addresses may not be needed because corporations typically have sufficient internal non-routable address (e.g. 10.x.x.x or 192.x.x.x) space. Consider the example shown in Figure 3.

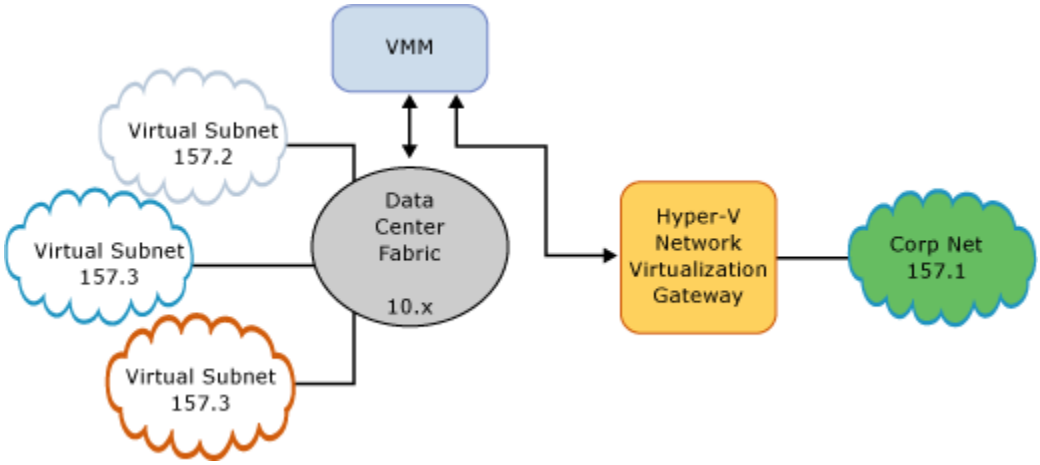

Figure 3: Private Cloud Deployment

Notice in this example that the Customer Addresses in the virtual subnets are 157.x addresses while the IP addresses in the non-network virtualized part of the network (Corp Net) are also 157.x addresses. In this case the PA addresses for the virtual subnets in the datacenter are 10.x IP addresses. This deployment allows the enterprise to take advantage of HNV's ability to offer flexibility in both virtual machine placement and cross-subnet live migration in the datacenter fabric. This increases datacenter efficiency thereby reducing both Operational Expenses (OpEx) and Capital Expenses (CapEx). In this scenario the HNV gateway provides routing between the 10.x and 157.1 IP addresses.

Hybrid Cloud (Site to site VPN)

A key advantage of HNV is that it can seamlessly extend an on-premise datacenter to a Windows Server 2012 based cloud datacenter. This is called a Hybrid Cloud model as shown in Figure 4.

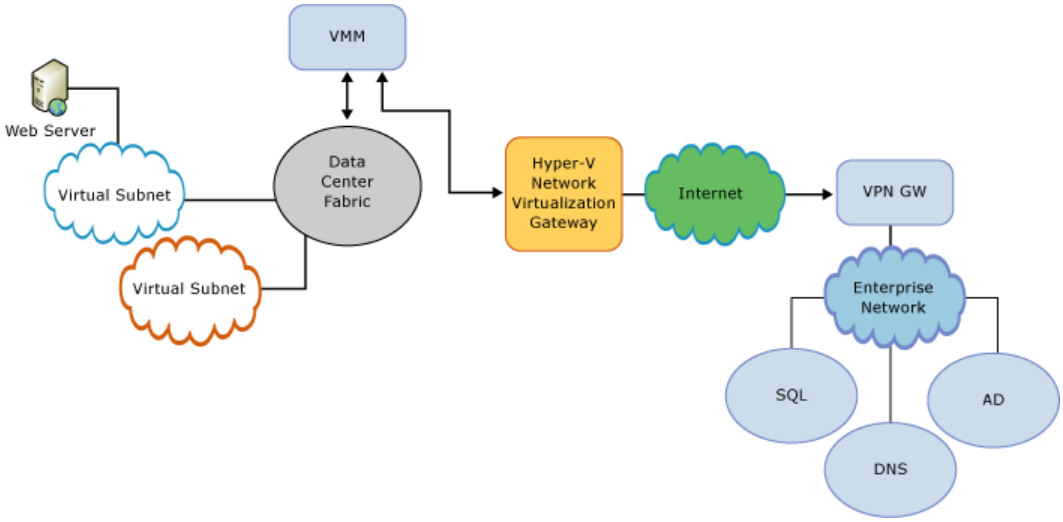

Figure 4: Hybrid Cloud Deployment

In this scenario an internal subnet, such as the subnet containing web servers is moved from the Enterprise Network into a Cloud Hoster's datacenter. Taking advantage of *Bring Your Own IP Address* offered by the Hoster, the Enterprise does not need to change the network configuration of the Web Server virtual machine or any other network endpoint that references that Web Server. The Hoster provides a secure link via a HNV Gateway Appliance. The Enterprise admins need only configure their on-premise VPN with the appropriate IP address. The Web Server virtual machine is unaware that it has been moved to the cloud. It remains domain-joined with Active Directory (AD) and uses the Enterprise's DNS server. The Web Server virtual machine also continues to interact with other servers in the enterprise such as a SQL Server.

The HNV Gateway can support multiple site to site (S2S) VPN tunnels as shown in Figure 5. Note that a VMM is not pictured in the diagram but is required for HNV deployments.

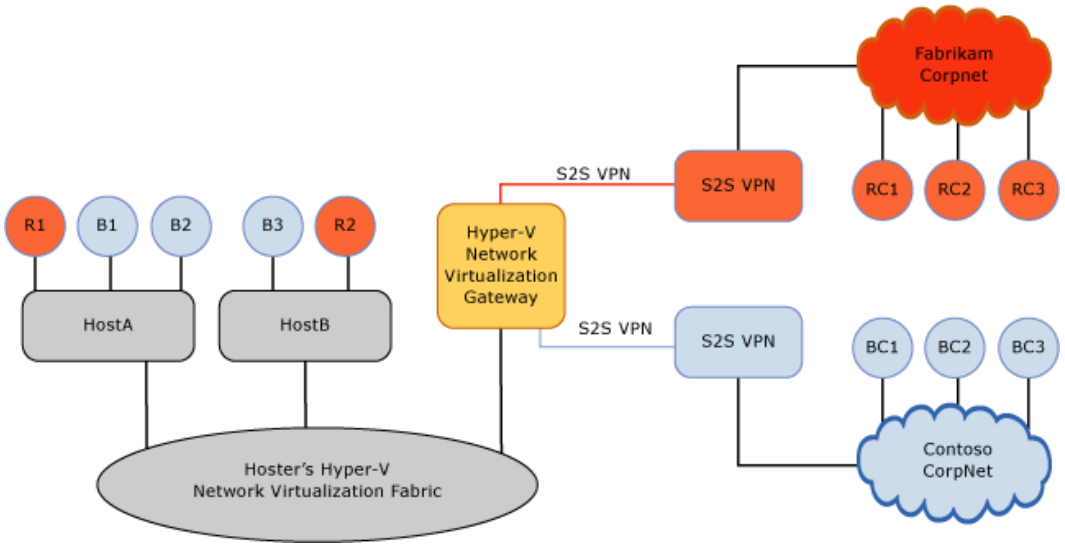

Figure 5: HNV Gateway

Packet Encapsulation

Each virtual network adapter in HNV is associated with two IP addresses:

- **Customer Address** (CA) The IP address that is assigned by the customer, based on their intranet infrastructure. This address enables the customer to exchange network traffic with the virtual machine as if it had not been moved to a public or private cloud. The CA is visible to the virtual machine and reachable by the customer.
- **Provider Address** (PA) The IP address that is assigned by the hoster or the datacenter administrators based on their physical network infrastructure. The PA appears in the packets on the network that are exchanged with the server running Hyper-V that is hosting the virtual machine. The PA is visible on the physical network, but not to the virtual machine.

The CAs maintain the customer's network topology, which is virtualized and decoupled from the actual underlying physical network topology and addresses, as implemented by the PAs. The

following diagram shows the conceptual relationship between virtual machine CAs and network infrastructure PAs as a result of network virtualization.

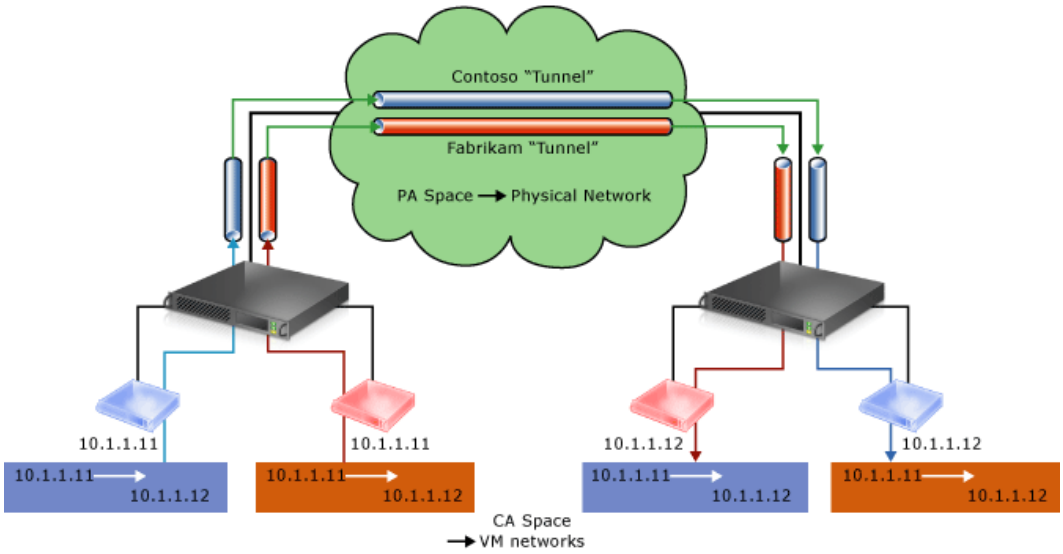

Figure 6: Conceptual diagram of network virtualization over physical infrastructure

In the diagram, customer virtual machines are sending data packets in the CA space, which traverse the physical network infrastructure through their own virtual networks, or "tunnels". In the example above, the tunnels can be thought of as "envelopes" around the Contoso and Fabrikam data packets with green shipping labels (PA addresses) to be delivered from the source host on the left to the destination host on the right. The key is how the hosts determine the "shipping addresses" (PA's) corresponding to the Contoso and the Fabrikam CA's, how the "envelope" is put around the packets, and how the destination hosts can unwrap the packets and deliver to the Contoso and Fabrikam destination virtual machines correctly.

This simple analogy highlighted the key aspects of network virtualization:

- Each virtual machine CA is mapped to a physical host PA. There can be multiple CAs associated with the same PA.
- Virtual machines send data packets in the CA spaces, which are put into an "envelope" with a PA source and destination pair based on the mapping.
- The CA-PA mappings must allow the hosts to differentiate packets for different customer virtual machines.

As a result, the mechanism to virtualize the network is to virtualize the network addresses used by the virtual machines. The next section describes the actual mechanism of address virtualization.

Network virtualization through address virtualization

HNV supports Network Virtualization for Generic Routing Encapsulation (NVGRE) as the mechanism to virtualize the IP Address:

Generic Routing Encapsulation This network virtualization mechanism uses the Generic Routing Encapsulation (NVGRE) as part of the tunnel header. In NVGRE, the virtual machine's packet is encapsulated inside another packet. The header of this new packet has the appropriate source and destination PA IP addresses in addition to the Virtual Subnet ID, which is stored in the Key field of the GRE header, as shown in Figure 7.

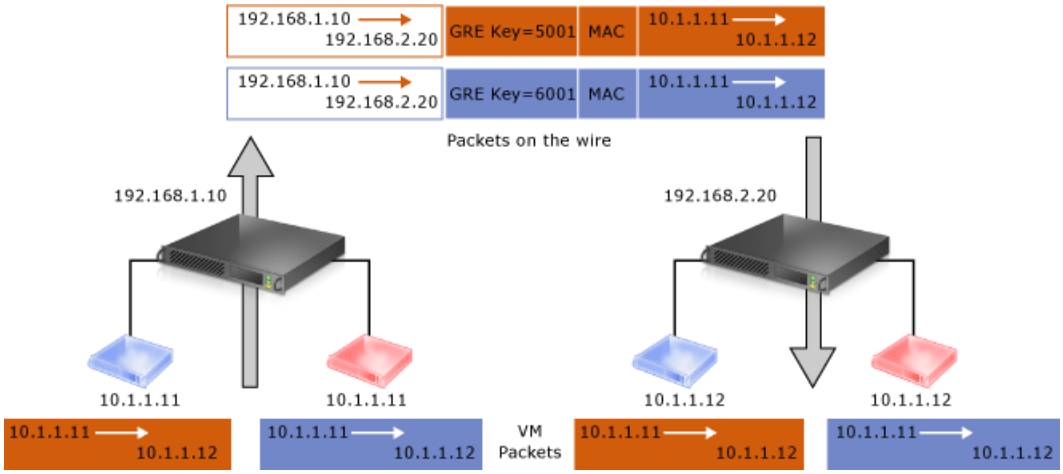

Figure 7: Network virtualization - NVGRE encapsulation

The Virtual Subnet ID allows hosts to identify the customer virtual machine for any given packet, even though the PA's and the CA's on the packets may overlap. This allows all virtual machines on the same host to share a single PA, as shown in Figure 7.

Sharing the PA has a big impact on network scalability. The number of IP and MAC addresses that need to be learned by the network infrastructure can be substantially reduced. For instance, if every end host has an average of 30 virtual machines, the number of IP and MAC addresses that need to be learned by the networking infrastructure is reduced by a factor of 30.The embedded Virtual Subnet IDs in the packets also enable easy correlation of packets to the actual customers.

With Windows Server 2012 and later, HNV fully supports NVGRE out of the box; it does NOT require upgrading or purchasing new network hardware such as NICs (Network Adapters), switches, or routers. This is because the NVGRE packet on the wire is a regular IP packet in the PA space, which is compatible with today's network infrastructure.

Windows Server 2012 made working with standards a high priority. Along with key industry partners (Arista, Broadcom, Dell, Emulex, Hewlett Packard, and Intel) Microsoft published a draft RFC that describes the use of Generic Routing Encapsulation (GRE), which is an existing IETF standard, as an encapsulation protocol for network virtualization. For more information, see the

following Internet Draft: [Network Virtualization using Generic Routing Encapsulation](). As NVGRE-aware becomes commercially available the benefits of NVGRE will become even greater.

Multitenant deployment example

The following diagram shows an example deployment of two customers moving in a cloud datacenter with the CA-PA relationship defined by the HNV policies.

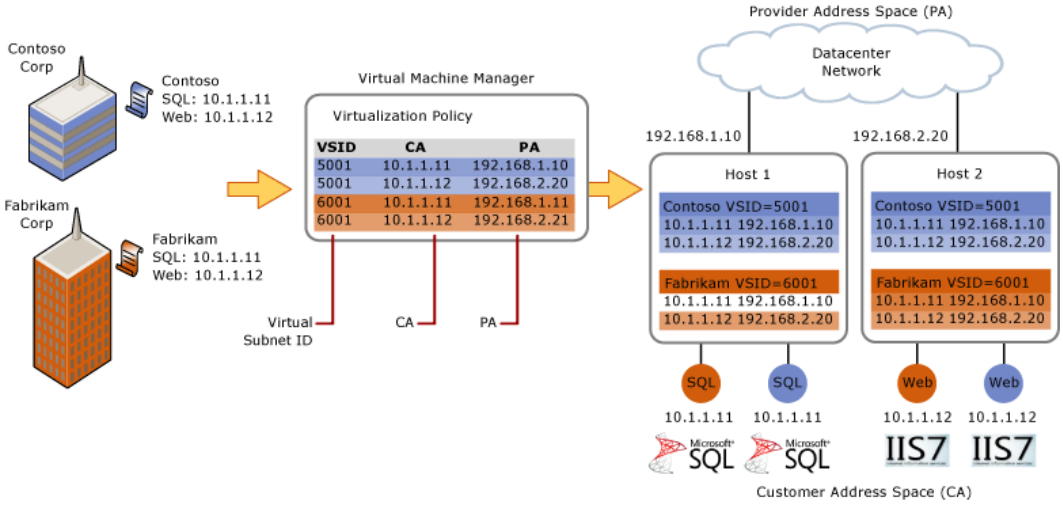

Figure 8: Multi-tenant deployment example

Consider the example in Figure 8. Prior to moving to the hosting provider's shared IaaS service:

- Contoso Corp ran a SQL Server (named **SQL**) at the IP address 10.1.1.11 and a web server (named **Web**) at the IP address 10.1.1.12, which uses its SQL Server for database transactions.
- Fabrikam Corp ran a SQL Server, also named **SQL** and assigned the IP address 10.1.1.11, and a web server, also named **Web** and also at the IP address 10.1.1.12, that uses its SQL Server for database transactions.

Contoso Corp and Fabrikam Corp move their respective SQL Servers and web servers to the same hosting provider's shared IaaS service where, coincidentally, they run the **SQL** virtual machines on Hyper-V Host 1 and the **Web** (IIS7) virtual machines on Hyper-V Host 2. All virtual machines maintain their original intranet IP addresses (their CAs).

Both companies are assigned the following Virtual Subnet ID (VSID) and PAs by their hosting provider when the virtual machines are provisioned:

- PAs of Contoso Corp's virtual machines: **VSID** is 5001, **SQL** is 192.168.1.10, **Web** is 192.168.2.20
- PAs of Fabrikam Corp's virtual machines: **VSID** is 6001, **SQL** is 192.168.1.10, **Web** is 192.168.2.20

The hosting provider creates policy settings, consisting of a customer virtual subnet for Fabrikam Corp that maps the CAs of the Fabrikam Corp virtual machines to their assigned PAs and VSID, and a separate customer virtual subnet for Contoso Corp that maps the CAs of the Contoso Corp

virtual machines to their assigned PAs and VSID. The provider applies these policy settings to Hyper-V Host 1 and Hyper-V Host 2.

When the Contoso Corp Web virtual machine on Hyper-V Host 2 queries its SQL Server at 10.1.1.11, the following happens:

Hyper-V Host 2, based on its policy settings, translates the addresses in the packet from:

- Source: 10.1.1.12 (the CA of Contoso Corp **Web**)
- Destination: 10.1.1.11 (the CA of Contoso Corp **SQL**)

The encapsulated packet contains:

- GRE header with VSID: 5001
- Outer source: 192.168.2.20 (the PA for Contoso Corp Web)
- Outer destination: 192.168.1.10 (the PA for Contoso Corp SQL)

When the packet is received at Hyper-V Host 1, based on its policy settings, it will decapsulate the NVGRE packet with:

- Outer source: 192.168.2.20 (the PA for Contoso Corp **Web**)
- Outer destination: 192.168.1.10 (the PA for Contoso Corp **SQL**)
- GRE header with VSID: 5001

The decapsulated packet (the original packet sent from the Contoso Corp **Web** virtual machine) is delivered to the Contoso Corp **SQL** virtual machine:

- Source: 10.1.1.12 (the CA of Contoso Corp **Web**)
- Destination: 10.1.1.11 (the CA of Contoso Corp **SQL**)

When the Contoso Corp **SQL** virtual machine on Hyper-V Host 1 responds to the query, the following happens:

Hyper-V Host 1, based on its policy settings, translates the addresses in the packet from:

- Source: 10.1.1.11 (the CA of Contoso Corp **SQL**)
- Destination 10.1.1.12 (the CA of Contoso Corp **Web**)

The packet is encapsulated with:

- GRE header with VSID: 5001
- Outer source: 192.168.1.10 (the PA for Contoso Corp SQL)
- Outer destination: 192.168.2.20 (the PA for Contoso Corp Web)

When it is received at Hyper-V Host 2, it will, based on its policy settings, decapsulate the packet:

- Source: 192.168.1.10 (the PA for Contoso Corp **SQL**)
- Destination: 192.168.2.20 (the PA for Contoso Corp **Web**)
- GRE header with VSID: 5001

The decapsulated packet is delivered to the Contoso Corp **Web** virtual machine with:

- Source: 10.1.1.11 (the CA of Contoso Corp **SQL**)
- Destination: 10.1.1.12 (the CA of Contoso Corp **Web**)

A similar process for traffic between the Fabrikam Corp **Web** and **SQL** virtual machines uses the HNV policy settings for the Fabrikam Corp. As a result, with HNV, Fabrikam Corp and Contoso

Corp virtual machines interact as if they were on their original intranets. They can never interact with each other, even though they are using the same IP addresses.

The separate addresses (CAs and PAs), the policy settings of the Hyper-V hosts, and the address translation between the CA and the PA for inbound and outbound virtual machine traffic isolate these sets of servers. Furthermore, the virtualization mappings and transformation decouples the virtual network architecture from the physical network infrastructure. Although Contoso **SQL** and **Web** and Fabrikam **SQL** and **Web** reside in their own CA IP subnets (10.1.1/24), their physical deployment happens on two hosts in different PA subnets, 192.168.1/24 and 192.168.2/24, respectively. The implication is that cross-subnet virtual machine provisioning and live migration become possible with HNV.

Note

For more information about packet flow, download the Hyper-V Network Virtualization Packet Flow PowerPoint presentation at http://www.microsoft.com/en-us/download/details.aspx?id=34782.

Hyper-V Network Virtualization architecture

In Windows Server 2012, HNV policy enforcement and IP virtualization was performed by the Network Driver Interface Specification (NDIS) Lightweight Filter (LWF) called Windows Network Virtualization (WNV). The WNV filter was located below the Hyper-V switch as shown in Figure 9. A side effect of this architecture was that switch extensions could only see the CA address space traffic, but not for the PA space.

In Windows Server 2012 R2, HNV is now a part of the virtual switch, enabling extensions to gain visibility into both the CA and PA space addresses.

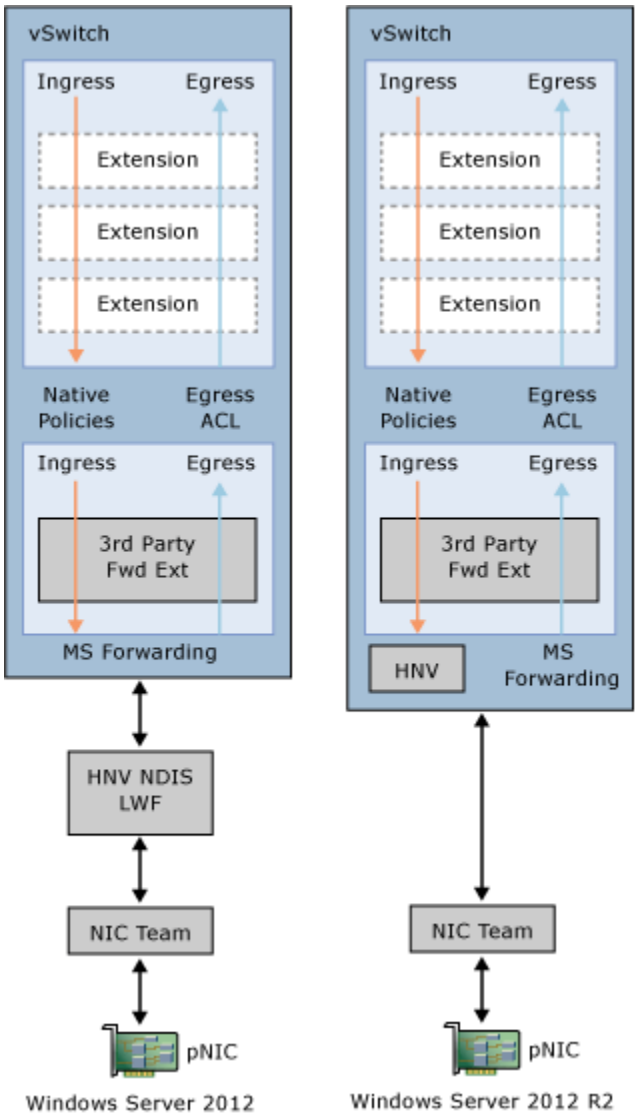

Figure 9: HNV Architecture

Each virtual machine network adapter is configured with an IPv4, and/or an IPv6 address. These are the CAs that will be used by the virtual machines to communicate with each other, and they are carried in the IP packets from the virtual machines. HNV virtualizes the CAs to PAs based on the network virtualization policies.

A virtual machine sends a packet with source address CA1, which is virtualized based on HNV policy in the Hyper-V switch. A special network virtualization access control list based on VSID isolates the virtual machine from other virtual machines that are not part of the same virtual subnet or part of the same routing domain.

Hyper-V Network Virtualization Policy Management

The Windows platform provides public APIs for datacenter management software to manage HNV. Virtual Machine Manager is one such datacenter management product. The management software contains all of the HNV policies. Because the virtual machine manager must be aware of virtual machines and more importantly provisions virtual machines and complete customer virtual networks in the datacenter and must be multi-tenant aware, managing HNV policy is a natural extension for policy-based networking.

Summary

Cloud-based datacenters can provide many benefits such as improved scalability and better resource utilization. To realize these potential benefits requires a technology that fundamentally addresses the issues of multi-tenant scalability in a dynamic environment. HNV was designed to address these issues and also improve the operational efficiency of the datacenter by decoupling the virtual network topology for the physical network topology. Building on an existing standard, HNV runs in today's datacenter and as NVGRE-aware hardware becomes available the benefits will continue to increase. Customers, with HNV can now consolidate their datacenters into a private cloud or seamlessly extend their datacenters to a hoster's environment with a hybrid cloud.

See also

To learn more about HNV see the following links:

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2 • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2)
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide

Content type	References
	• Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	Hyper-V Virtual Switch Overview

Hyper-V Network Virtualization Gateway Architectural Guide

Windows Server 2012 introduces a scalable, secure multi-tenant solution called Hyper-V Network Virtualization. This technology can be used to build cloud datacenters and makes it easier for customers to incrementally move their network infrastructure to private, hybrid, and public clouds.

For an overview of Hyper-V Network Virtualization, see Hyper-V Network Virtualization Overview, Hyper-V Network Virtualization Technical Details, and Windows Server® 2012 Hyper-V Survival Guide at the Microsoft TechNet web site.

Hyper-V Network Virtualization is deployable with existing server and switch hardware. However, gateway functionality, in addition to a datacenter management solution, is required to complete key scenarios including private cloud, cross-premise, and hybrid cloud deployments.

This article provides requirements for the different types of Hyper-V Network Virtualization gateways and gateway functionality. Note that gateway functionality is specifically called out because existing network appliances may incorporate Hyper-V Network Virtualization gateway functionality. Additional technical details regarding the packet encapsulation format used by Hyper-V Network Virtualization may be found in the NVGRE draft RFC.

Gateway Scenarios

Most customer deployments require communication from the network virtualized environment to the non-network virtualized environment. Therefore Hyper-V Network Virtualization gateways are required to bridge the two environments.

Gateways can come in different form factors. They can be built upon Windows Server 2012, incorporated into a Top of Rack (TOR) switch, put into an existing network appliance, or can be a stand-alone network appliance. The routing functionality is required for the private cloud scenario described in the Private Cloud section. In addition to routing, VPN functionality is required for the hybrid cloud scenario (described in the Hybrid Cloud section). A third scenario incorporates the gateway functionality into a load balancer (described in the Load Balancer section).

Private Cloud (routing)

Large enterprises may either be hesitant, or for compliance reasons, unable to move some of their services and data to a public cloud hoster. However, enterprises still want to obtain the benefits of the cloud and network virtualization by consolidating their datacenter resources into a private cloud. In a private cloud deployment, overlapping IP addresses may not be needed because corporations have sufficient internal non-routable address space (for example, 10.x.x.x or 192.x.x.x). Consider the example shown in Figure 1.

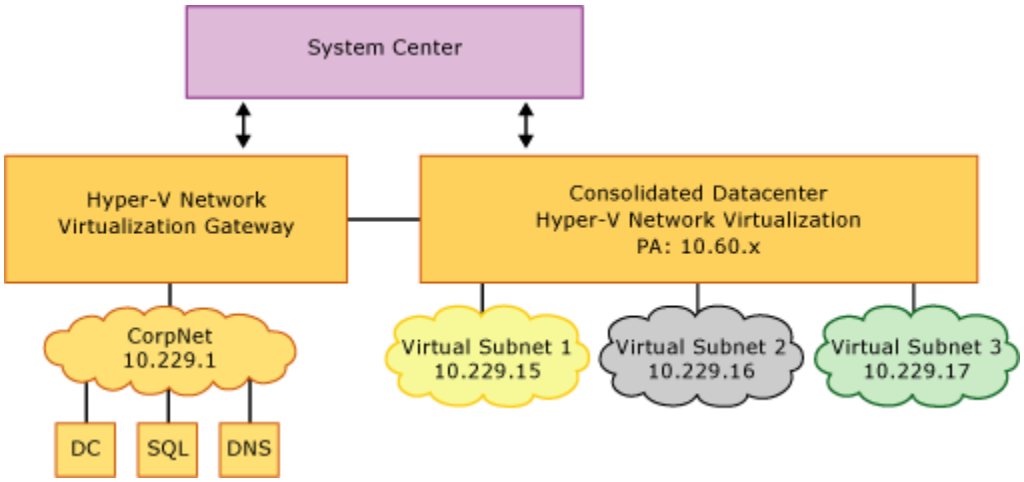

Figure 1: Private cloud deployment

Notice in this example that the customer addresses in the virtual subnets are 10.229.x IP addresses while the IP addresses in the non-network virtualized part of the network (CorpNet) are 10.229.1 addresses. In this case the PA addresses for the virtual subnets in the datacenter are 10.60.x IP addresses. This deployment allows the enterprise to take advantage of Hyper-V Network Virtualization's ability to offer flexibility in both virtual machine placement and cross-subnet live migration in the datacenter fabric. This increases datacenter efficiency thereby reducing both Operational Expenses (OpEx) and Capital Expenses (CapEx).

The Private Cloud scenario shown in Figure 1 requires a private cloud gateway appliance or gateway functionality incorporated into an existing network appliance. The purpose of the private cloud gateway is to provide routing and transition between physical and virtual addresses.

Hybrid Cloud (S2S VPN)

A key advantage of Hyper-V Network Virtualization is that it can seamlessly extend an on-premise datacenter to a Windows Server 2012 based cloud datacenter. This is called a Hybrid Cloud model as shown in Figure 2.

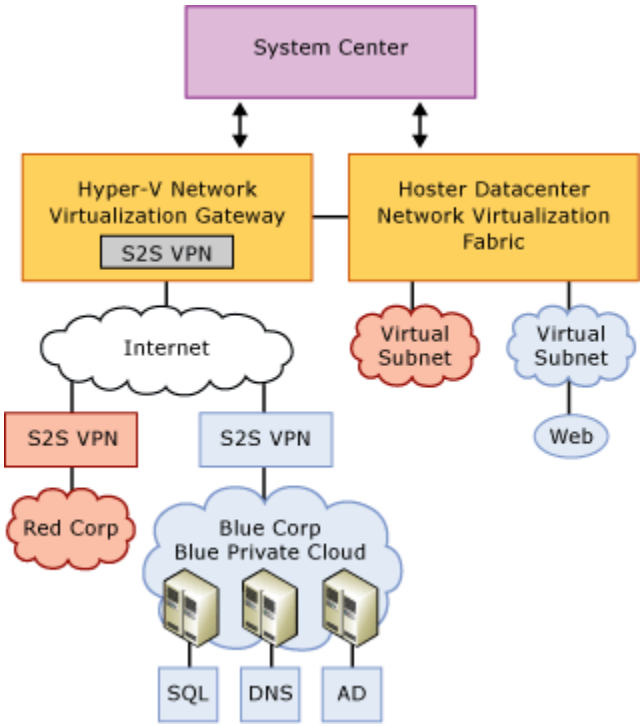

Figure 2: Hybrid cloud deployment

In this scenario an internal server, such as a web server, is moved from the enterprise network into a cloud hoster's datacenter. Taking advantage of *Bring Your Own IP Address* offered by the hoster, the enterprise does not need to change the network configuration of the web server virtual machine or any other network endpoint that references that web server. The hoster provides a secure link via a VPN gateway appliance. The enterprise administrators need to only configure their on-premise VPN with the appropriate IP address. The web server virtual machine is unaware that it has been moved to the cloud. It remains domain-joined with Active Directory (AD) and uses the enterprise's DNS server. The web server virtual machine also continues to interact with other servers in the enterprise such as a SQL Server. In this example, all three of these services (AD, DNS, SQL) remain on-premise. Network diagnostic tools such as tracert, which counts the number of network hops between a source and destination, show that the network traffic between the web server virtual machine and SQL is no longer local but routed across the Internet.

The Hybrid Cloud scenario shown in Figure 2 requires a multi-tenant VPN appliance such as a network virtualization gateway appliance. The VPN appliance gateway must interact with the datacenter orchestrator (for example, System Center Virtual Machine Manager) to obtain the appropriate network virtualization policies.

Load Balancer

A load balancer provides the illusion of a single IP address (Virtual IP address or VIP) to clients requesting a service. The service is implemented on multiple different IP addresses (Direct IP

address or DIP). A traditional load balancer maps network traffic from VIP to DIP via Network Address Translation (NAT). The load balancer re-writes the destination IP address (VIP) of the incoming packet with one of the IP addresses (DIP) behind the load balancer that performs the service. In the case where the load balancer is acting as a bi-directional NAT, the DIP sends a return packet to the load balancer and the load balancer rewrites the source IP (DIP) with the VIP. In this case the client only knows about the VIP and is unaware of the DIP.

In the context of Hyper-V Network Virtualization the load balancer emits NVGRE packets to get the packets to their destination virtual machines.

Consider the following scenario where the Blue Company and the Red Company have their virtual machines in the same multi-tenant datacenter that provides load balancing for the tenants. Blue Company has four virtual machines offering a service behind a load balancer. The Red Company has two virtual machines offering a different service that they want load balanced.

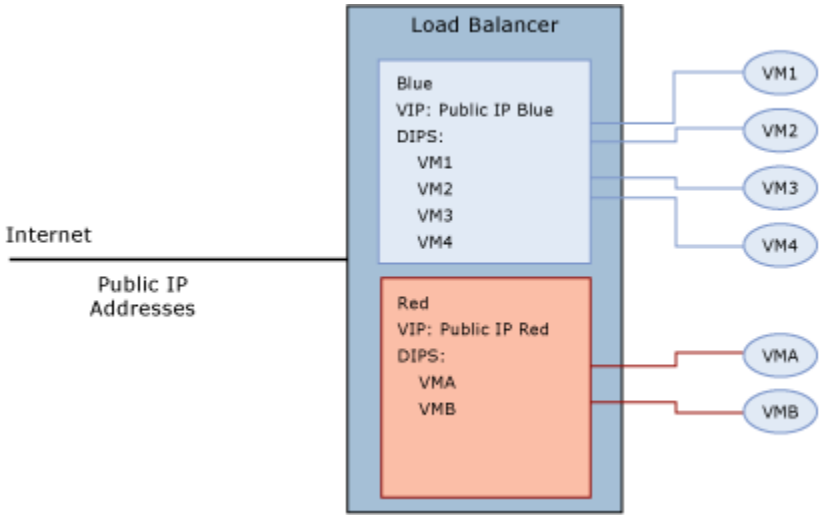

Figure 3: Multi-tenant Load Balancer Example

Figure 3 shows a virtual topology where publically routable IP addresses are load balanced across multiple destination virtual machines. For example, clients from the Internet send traffic to the Blue public IP (VIP). The load balancer, based on Network Virtualization policy, emits an NVGRE packet to get the packet to the appropriate Hyper-V host. They Hyper-V host will decapsulate the packet and then deliver it to the destination virtual machine.

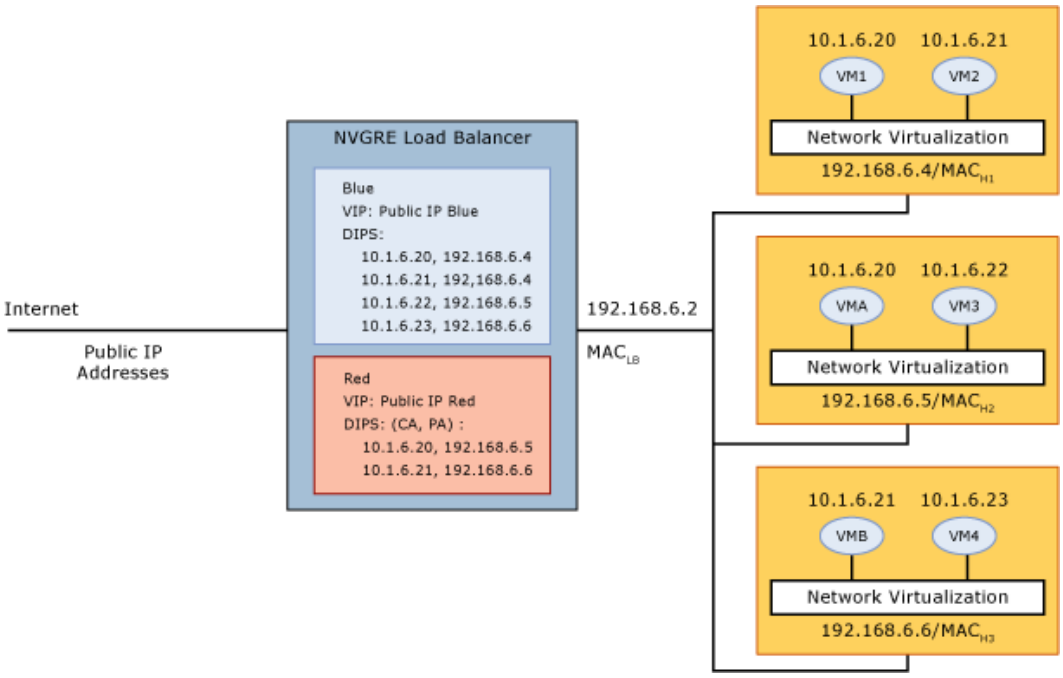

Figure 4: Multi-tenant Load Balancer deployment

Figure 4 show a possible deployment of Figure 3. In this case the PA space of the load balanced virtual machines is 192.168.6.x. Blue Customer has four virtual machines that are running on three physical hosts. Red Customer has two virtual machines running on two different hosts. Both Blue and Red are hosting web services. The outside world accesses these services via the appropriate publically routable VIP (Blue VIP and Red VIP). The load balancer, using network virtualization policy, generates the appropriate NVGRE packet.

Consider the case where a connection initiated from ClientIP and sent to the Blue VIP is load balanced to Blue VM2. The load balancer needs to send a packet to Blue VM2 which has a VSID of BlueVSID1, CA of 10.1.6.21 and a PA of 192.168.6.4. The load balancer's PA is 192.168.6.2. Therefore the load balancer generates the NVGRE packet header:

MAC_{LB} --> MAC_{H1}	192.168.6.2 --> 192.168.6.4	BlueVSID1	MAC_{ext} --> MAC_{VM2}	ClientIP --> 10.1.6.21

Now consider the scenario shown in Figure 5 where Blue and Red customers have multi-tier applications that require load balancing. From the perspective of the virtual machines they communicate to the next tier via VIP1. For example virtual machines in Tier B1 communicate to virtual machines in Tier B2 via Blue VIP1. Blue VIP1 is a Blue CA IP address which will also have a corresponding PA address on the load balancer. A possible deployment topology is shown in Figure 6.

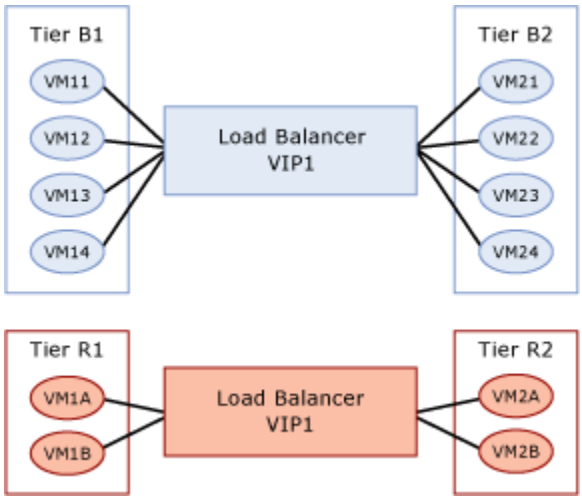

Figure 5: Load balancing between tiers (logical perspective)

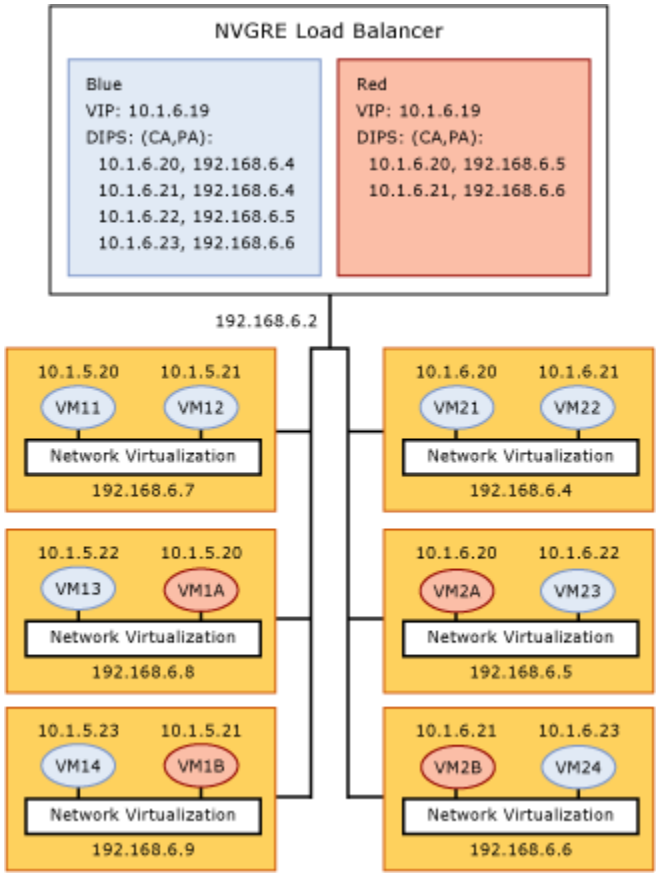

Figure 6: Load Balancer placed internally in the datacenter

The Blue CA VIP and the Red CA VIP on the load balancer both have the PA of 192.168.6.2. The load balancer needs to look at the VSID in the NVGRE packet and then determine the appropriate VIP to DIP mapping. For example Blue virtual machine 11 will send a packet to Blue

CA VIP. The NVGRE packet header will have Source IP: 192.168.6.7, Destination IP: 192.168.6.2, GRE Key Field Blue VSID1, Inner Source IP: 10.1.5.20, Inner Destination IP: Blue CA VIP, and the rest of the original packet. The load balancer's policy must match on both the VSID and the VIP because with overlapping CA IP addresses the VIP by itself may not be unique.

External Router Gateway Configuration

Gateways must contain Hyper-V Network Virtualization policy and perform actions similar to a Windows Server 2012 host. In the typical deployment the Windows Server 2012 host is aware of all the policy information for every routing domain for all virtual machines located on that host. For example, if a host has a virtual machine with routing domain X and another virtual machine from routing domain Y the host has all the policy information for all virtual subnets contained in both routing domain X and routing domain Y. Even in the case of a single routing domain there could be a substantial amount of policy that needs to be on every Hyper-V host.

Consider the case of a private cloud shown in Figure 1. In this scenario a corporation has an internal cloud with a single routing domain. In this case every Hyper-V host must have policy for all network virtualized virtual machines. This can place a burden on the datacenter orchestrator to maintain this policy on all hosts. Note that it is rare for a virtual machine to actually communicate with every other virtual subnet in the corporation so much of this policy on a given host is in practice not needed.

To reduce the amount of policy needed on each host the Windows Server 2012 host can be put into a special mode where it routes traffic within a given Virtual Subnet based on its own host policy. Any cross virtual-subnet traffic is sent to an external router for a network virtualization policy check and potentially forwarding the packet along the path to its eventual destination. In such a deployment the host only needs to be aware of policy for the virtual subnets of virtual machines running on that host and not for every virtual subnet of the routing domains of the virtual machines running on the host. In this scenario the Hyper-V Network Virtualization Gateway makes all cross-subnet routing decisions. The Hyper-V Network Virtualization Gateway must have all policy information for all virtual machines in the virtual subnets to which it routes packets.

Management

For any Hyper-V Network Virtualization deployment a datacenter orchestrator is required. One such datacenter orchestrator is System Center Virtual Machine Manager (VMM). The datacenter orchestrator is responsible for providing a mechanism to get Hyper-V Network Virtualization policies to the appropriate gateway appliance. Note that PA IP routing is not managed by the datacenter orchestrator.

Gateway Console

The Hyper-V Network Virtualization Gateway partner should provide a management console for basic configuration of the gateway appliance or appliance functionality. This management console should handle operations such as PA assignment, high availability, monitoring, and

authentication. The management console will be responsible for configuring the aspects of the gateway not managed by a datacenter orchestrator such as VMM.

Multiple Gateways

High Availability by using multiple physical Hyper-V Network Virtualization gateways is outside the scope of this article. The datacenter orchestrator (for example VMM) will need to configure and manage the gateway. Scenarios in which the datacenter deploys multiple gateways include scalability and resiliency. For scalability a single physical gateway may not be able to handle the load required for a datacenter. Additionally a single gateway introduces a single point of failure. The partner's management console should handle deployments consisting of multiple gateways.

System Center Virtual Machine Manager

In the VMM model the Hyper-V Network Virtualization Gateway is managed via a PowerShell plug-in module. Partners building Hyper-V Network Virtualization gateways need to create a PowerShell plug-in module which physically runs on the VMM server. This plug-in module will communicate policy to the gateway. Figure 7 shows a block diagram of VMM managing a Hyper-V Network Virtualization deployment. Note that a partner plug-in runs inside the VMM server. This plug-in communicates to the gateway appliances. The protocol used for this communication is not specified here. The partner may determine the appropriate protocol. Note that VMM uses the Microsoft implementation of WS-Management Protocol called Windows Remote Management (WinRM) and Windows Management Instrumentation (WMI) to manage the Windows Server 2012 hosts and update network virtualization policies.

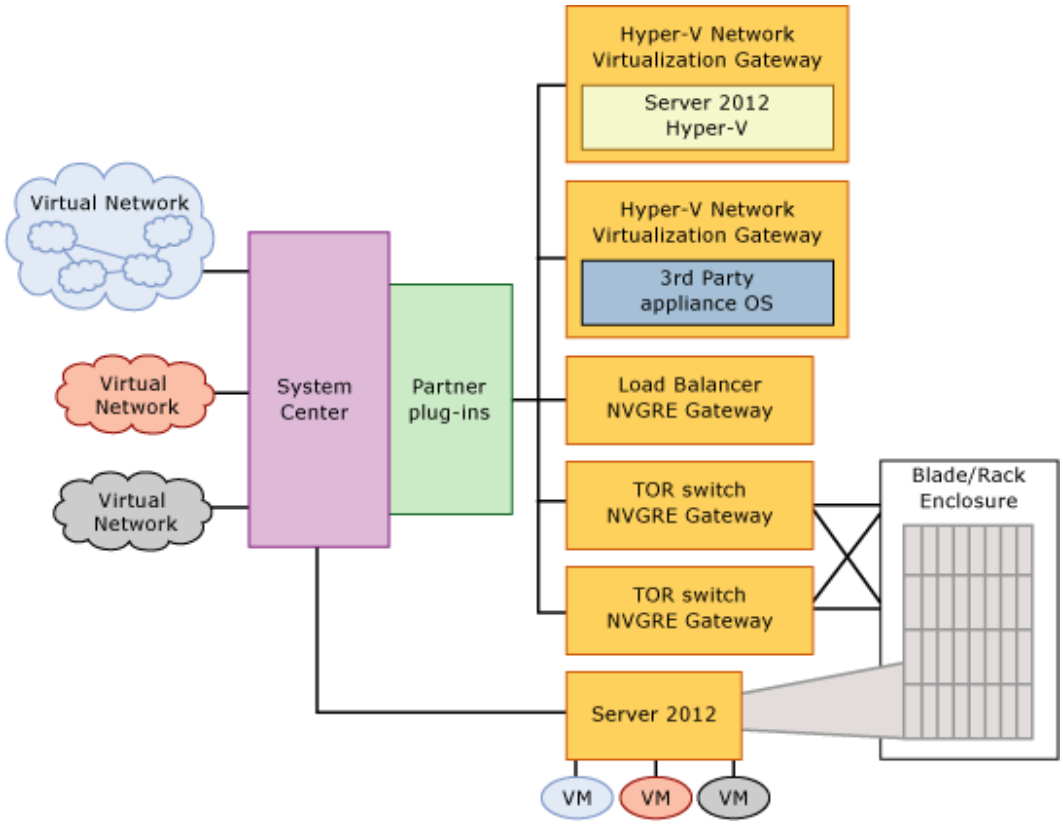

Figure 7: Using System Center to manage Hyper-V Network Virtualization

Windows-Based Gateway Appliances

Windows Server 2012 can be used as the base platform for a Hyper-V Network Virtualization Gateway appliance.

Private Cloud Router Architecture

Figure 8 shows the architecture for the Private Cloud Router using Windows Server 2012 as the base platform. In this scenario all virtual machines are in the same routing domain. The red, blue, and purple virtual machines have a CA IP prefix of 157.4, 157.3, 157.2, respectively. The services on CorpNet have IP addresses with a prefix of 157.1.

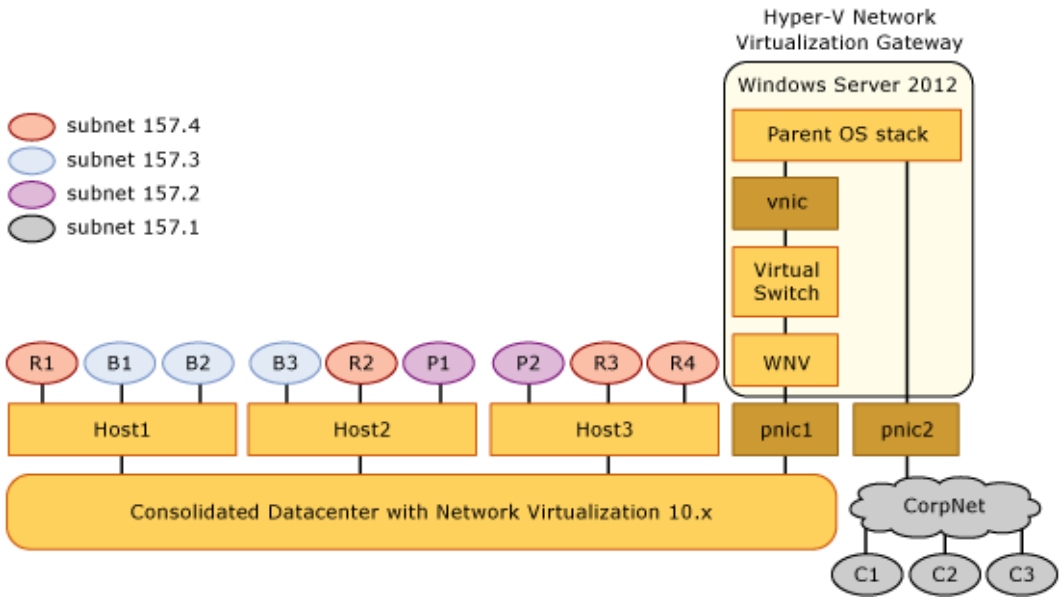

Figure 8: Private Cloud Gateway using Windows Server 2012

The WNV filter in the gateway receives traffic from the red, blue, and purple virtual subnets and sends it to CorpNet via the parent(Management) operating system stack. The virtual switch has a parent (host) VNIC for communication to the host operating system stack from the network virtualized world. The other network interface is bound to the physical network interface (pnic2) connected to CorpNet. Note that there is no requirement that there be two physical network interfaces in the Hyper-V Network Virtualization Gateway shown in Figure 8. A single physical network interface could be used along with two parent virtual network interfaces. However, datacenter administrators may prefer a physical separation of traffic. Two physical network interfaces in the gateway allow the consolidated datacenter and CorpNet traffic to interact with the gateway via different physical switches. Two physical network interfaces, one for the network virtualized world and one for the non-network virtualized world provides more flexibility in customer deployments.

Cross Premise Gateway

The Hybrid Cloud scenario enables an enterprise to seamlessly expand their on-premises datacenter into the cloud. This requires a site to site VPN tunnel. This can be accomplished with Windows Server 2012 as the host platform and a per tenant Windows Server 2012 guest virtual machine running a Site To Site (S2S) VPN tunnel connecting the cloud datacenter with various on-premise datacenters. Windows Server 2012 S2S VPN supports IKEv2 and configuration of remote policy can be accomplished via PowerShell/WMI. In addition Windows Server 2012 guest virtual machines support new network interface offload capabilities that enhance the performance and scalability of the gateway appliance. These offload capabilities are discussed below in the Hardware Considerations section.

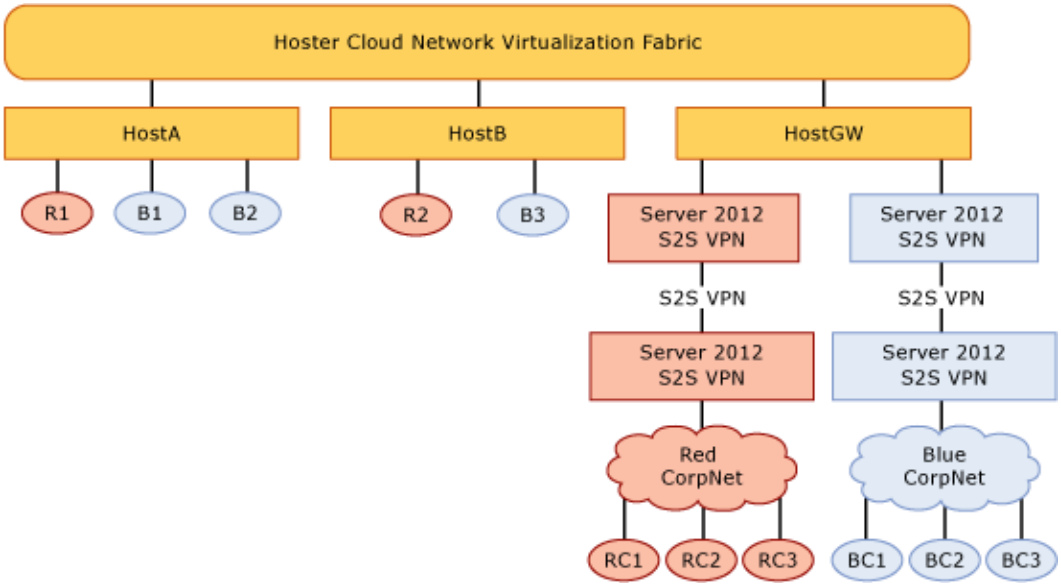

Figure 9: Hybrid Cloud with Windows Server 2012 based Gateway

Figure 9 shows a scenario where Red Corp and Blue Corp are customers of Hoster Cloud. Red Corp and Blue Corp seamlessly extend their datacenter into Hoster Cloud has deployed Windows Server 2012 based per tenant virtual machine gateways allowing Red Corp and Blue Corp to seamlessly extend their on-premise datacenters. In Figure 10 there is no requirement that Red Corp or Blue Corp run Windows Server 2012 S2S VPN, only that the customer's on premise S2S VPN support IKEv2 to interact with corresponding Windows Server 2012 S2S virtual machines running on HostGW.

Figure 9 shows the internal architecture for HostGW. Each Routing Domain requires its own virtual machine. The technical reason for this is that a vmnic can only be associated with a single Virtual Subnet (VSID) and a VSID can only be part of a single routing domain. The VSID switch port ACL does not support trunking of VSIDs. Therefore the simplest way to provide isolation is with a per tenant (Routing Domain) gateway virtual machine.

Each of the virtual machines is dual homed which means they have two virtual network interfaces. One of the virtual network interfaces has the appropriate VSID associated with it. The other virtual network interface has a VSID of 0 which means traffic is not modified by the WNV filter. The Windows Server 2012 virtual machine is running RRAS and using IKEv2 to create a secure tunnel between Hoster Cloud and the customer's on premise gateway.

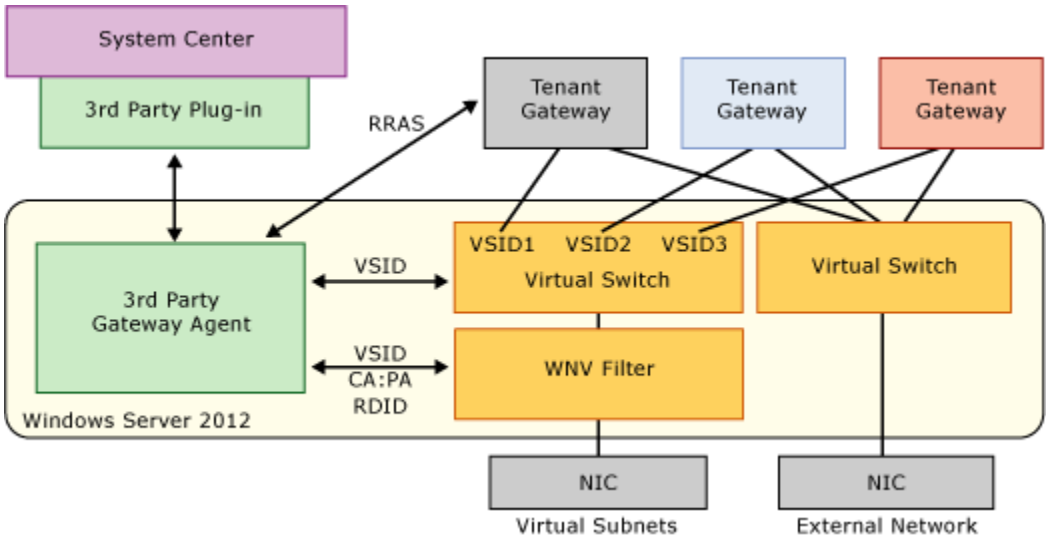

Figure 10: Hybrid Cloud with Windows Server 2012 based Per Tenant VM Gateways

Figure 10 shows the architecture where VMM is managing a Hyper-V Network Virtualization deployment. The partner has a plug-in that runs in the VMM server. When using Windows Server 2012 as a Hyper-V Network Virtualization gateway appliance a local management process running in Windows is required as the end point for this communication from the plug-in running in the VMM server. This is how the plug-in is able to communicate network virtualization policy to the WNV filter running on HostGW.

Hardware Considerations

The scalability of the Windows Server 2012 based gateway appliance will be determined by the server hardware including:

- NUMA capabilities
- number of CPU cores
- amount of RAM
- offload capabilities and performance of the network interfaces

In general, a dual socket server class platform provided by major OEMs should be a sufficient base platform. The number of virtual machines that can effectively run on the server appliance is typically determined by the amount of RAM in the server.

The choice of network interfaces can also have a great impact on the performance and scalability of the gateway appliance. Figure 10 shows a two network interface configuration. Note that ALL traffic sent on the virtual subnet path will be NVGRE traffic. For encapsulated packets traditional offloads such as Large Send Offload (LSO) on the send path and Virtual Machine Queues (VMQ) on the receive path will not provide the expected benefit because they operate on the outer packet header. The outer header for NVGRE traffic makes it appear as though all traffic is generated using the same IP address and destined for a single IP address. Network interface offloads provide substantial benefit to the Hyper-V switch performance. A new feature in Windows

Server 2012 is GRE Task Offload where the network interface will operate on the inner packet header for standard offloads such that LSO and VMQ provide their expected performance and scalability benefits. Therefore the network interface on the path for Virtual Subnets in Figure 10 should support GRE Task Offload. Such network interfaces are expected to be commercially available by the end of 2012.

Another new feature in Windows Server 2012 is IPsec Task Offload for guest virtual machines. IPsec is used for securing site to site VPN tunnels. Both sides of the S2S tunnel must support IKEv2. If the Cross-Premise gateway is NOT behind a NAT then IPsec Task Offload (IPsecTO) can provide better performance by having the External Network network interface in Figure 9 perform the per packet encryption operations on behalf of the per Tenant RRAS virtual machine. Both sides of the S2S tunnel would need to use AES-GCM128 to take advantage of IPsecTO. Therefore, the following conditions must be met for the RRAS virtual machine to take advantage of IPsecTO:

- Both sides of the S2S tunnel must use AES-GCM128
- The gateway appliance must not be behind a NAT

Because IPsec operations can consume a large amount of CPU, IPsecTO can provide a significant CPU reduction and also an increase in network throughput.

SR-IOV is a new Windows Server 2012 feature that directly exposes a physical network interface to the virtual machine. SR-IOV bypasses the virtual switch. Unfortunately, current IPsecTO network interfaces on the market do not support IPsecTO on the SR-IOV path. The Windows Server 2012 Hyper-V Switch introduces new policy-based features such as ACLs and QoS. If any of these Hyper-V switch policies are set on a virtual network interface , the virtual network interface traffic will not take the SR-IOV path and instead will need to go through the Hyper-V Switch path for policy enforcement. When traffic does take the SR-IOV path instead of the Hyper-V Switch path there will be a significant reduction in CPU utilization.

For the Private Cloud gateway similar considerations should be evaluated regarding hardware requirements.

Additional Resources

Additional information regarding Hyper-V Network Virtualization and System Center Virtual Machine Manager can be found in the following articles (articles without links are still under development):

- Windows Server 2012 Hyper-V Network Virtualization Survival Guide
- Hyper-V Network Virtualization Overview
- Network Virtualization Technical Details
- //BUILD talk on Hyper-V Network Virtualization
- //BUILD Keynote Demo (starts 13.5 minutes into video)
- NVGRE Draft RFC

Virtual Hard Disk Sharing Overview

Starting in Windows Server 2012 R2, Hyper-V makes it possible to share a virtual hard disk file between multiple virtual machines. Sharing a virtual hard disk file (.vhdx) provides the shared storage that is necessary for a Hyper-V guest failover cluster. This is also referred to as a virtual machine failover cluster. For information on creating and enabling a shared virtual hard disk, see Configure a Shared Virtual Hard Disk.

Key benefits

Sharing a virtual hard disk file (.vhdx) means that you can create and manage a guest failover cluster to protect the application services running inside your virtual machines. Before Windows Server 2012 R2, if you wanted to create a Hyper-V guest failover cluster, you needed to expose your storage topology to the virtual machine.

Starting in Windows Server 2012 R2, you can deploy a Hyper-V guest failover cluster that is no longer bound to your storage topology. You can implement a guest failover cluster by using a shared virtual hard disk, Fibre Channel, Server Message Block (SMB), Storage Spaces, or iSCSI storage options. Shared virtual hard disks are only available in Windows Server 2012 R2. Hyper-V makes it possible to share a virtual hard disk file between multiple virtual machines. Sharing a virtual hard disk file (.vhdx) provides the shared storage that is necessary for a Hyper-V guest failover cluster.

Using a shared virtual hard disk is ideal for the following situations:

- SQL Server database files.
- File server services running within a virtual machine.
- Database files that reside on shared disks.

Guest Failover Clusters

The following information will help you understand and deploy a guest failover cluster by using a shared virtual hard disk.

Shared virtual hard disk format

Shared virtual hard disk functionality in guest failover clusters exclusively uses the .vhdx file format. Although the shared virtual hard disk must use the .vhdx file format for the data drive, the operating system disk for a virtual machine can use either the .vhd or the .vhdx file format.

Supported virtual machines

Generation 1 and Generation 2 virtual machines are supported in a guest failover cluster using a shared virtual hard disk.

Supported guest operating systems

Supported guest operating systems for guest clusters that use a shared virtual hard disk include Windows Server 2012 R2 and Windows Server 2012. To support Windows Server 2012 as a guest operating system, Windows Server 2012 R2 Integration Services must be installed within the guest (virtual machine).

Deploying a shared virtual hard disk in a guest failover cluster

A Hyper-V guest failover cluster that uses a shared virtual hard disk has two preferred deployment models. The shared virtual hard disk for the guest failover cluster can be deployed on:

- Cluster Shared Volumes (CSVs) on block storage (including clustered Storage Spaces).
- Scale-Out File Server with SMB 3.0 on file-based storage.

A Hyper-V guest failover cluster can be configured and deployed by using Hyper-V Manager and Failover Cluster Manager. You can also use Windows PowerShell.

For deployment steps, see Deploy a Guest Cluster Using a Shared Virtual Hard Disk.

Using Cluster Shared Volume

One option for using virtual hard disk sharing is to use Cluster Shared Volumes. The following graphic provides an example configuration. You can also use other types of shared storage if required.

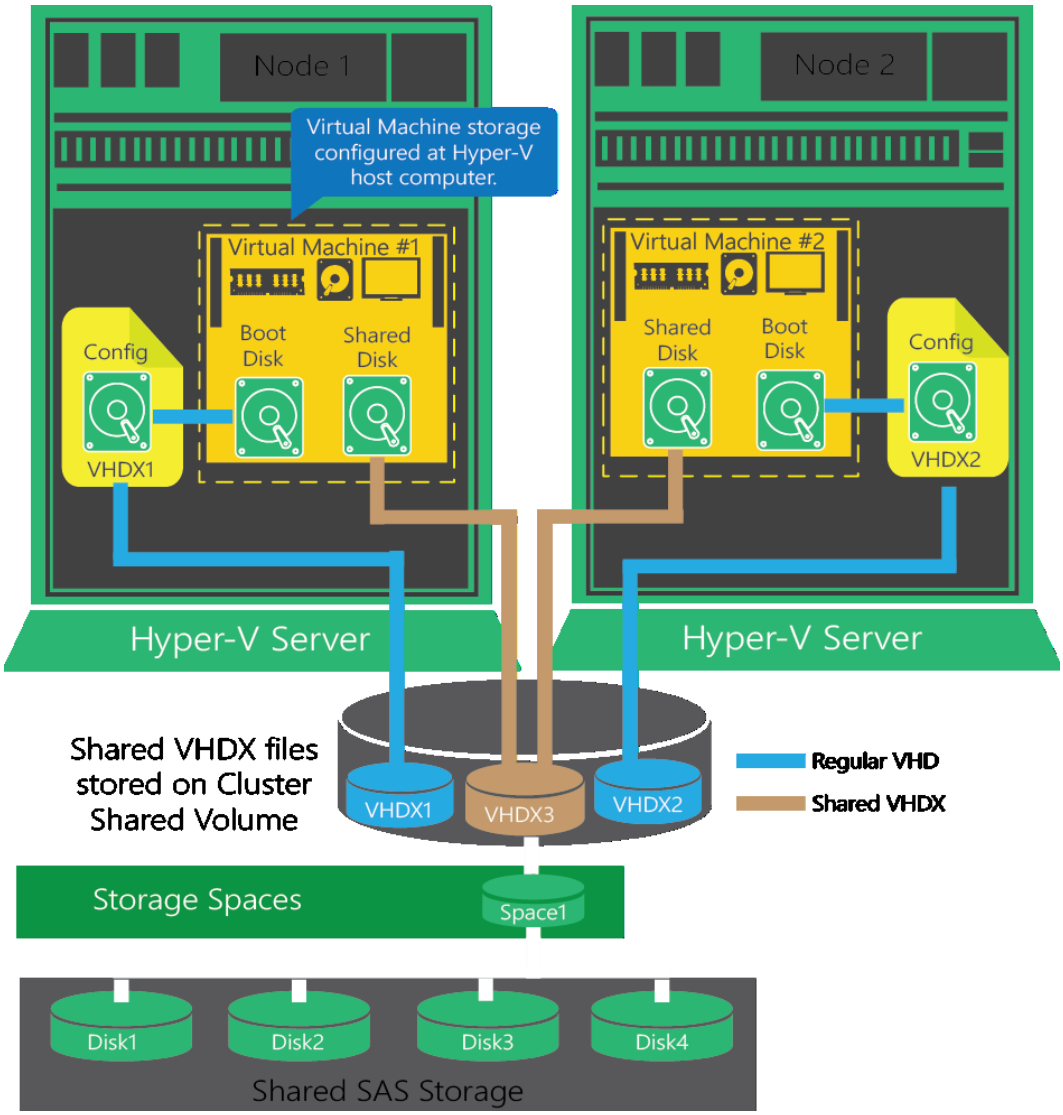

Using Scale-Out File Server with SMB 3.0

A second option for virtual hard disk sharing is to use a Scale-Out File Server. The following graphic provides an example configuration. You can also use other types of shared storage if required.

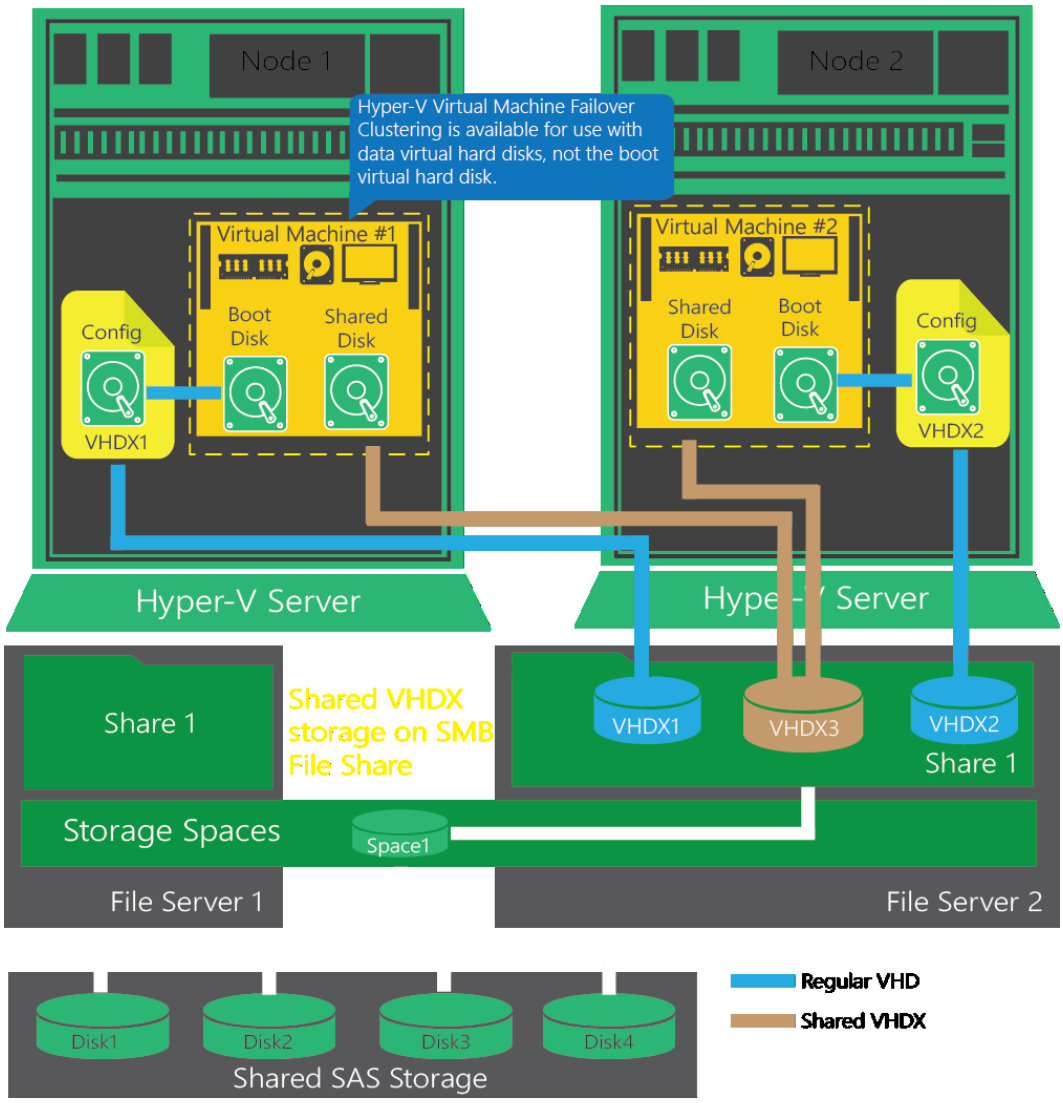

Requirements

To configure a guest failover cluster that uses shared virtual hard disks, you require the following:

- At least a two-node Hyper-V failover cluster.
- Servers must be running Windows Server 2012 R2.
- Servers must belong to the same Active Directory domain.
- Availability of configured shared storage resources—for example, CSVs on block storage (such as clustered storage spaces) or a Scale-Out File Server cluster (running Windows Server 2012 R2) with SMB 3.0 (for file-based storage).
- Sufficient memory, disk, and processor capacity within the failover cluster to support multiple virtual machines that are implemented as guest failover clusters.

Hyper-V Virtual NUMA Overview

Windows Server 2012 introduced support for projecting a virtual NUMA topology into Hyper-V virtual machines. This capability can help improve the performance of workloads running on virtual machines that are configured with large amounts of memory. This document briefly defines what NUMA is, explains how applications can take advantage of NUMA to optimize their performance, describes how virtual NUMA can be configured, and provides guidance about understanding virtual NUMA and its impact on virtual machine workloads.

This document is intended for Hyper-V administrators who are running high-performance NUMA-aware applications (such as database servers and web servers), and for system or software architects looking for information about virtual NUMA.

This topic includes the following sections:

- Introduction to NUMA
- NUMA topology in Hyper-V virtual machines
- Configuring Virtual NUMA
- Virtual NUMA and Dynamic Memory
- Virtual NUMA FAQs

Introduction to NUMA

Non-Uniform Memory Access (NUMA) is a computer system architecture that is used with multiprocessor designs in which some regions of memory have greater access latencies. This is due to how the system memory and processors are interconnected. Some memory regions are connected directly to one or more processors, with all processors connected to each other through various types of interconnection fabric. For large multiprocessor systems, this arrangement results in less contention for memory and increased system performance.

A NUMA architecture divides memory and processors into groups, called *NUMA nodes*. From the perspective of any single processor in the system, memory that is in the same NUMA node as that processor is referred to as *local*, and memory that is contained in another NUMA node is referred to as *remote*. Processors can access local memory faster.

Most modern operating systems and many high-performance applications that typically scale to utilize many processors and large amounts of memory, such as Microsoft SQL Server, include optimizations that recognize and adapt to a computer's NUMA topology. To avoid remote access penalties, a NUMA-aware application attempts to allocate storage for data and schedule processor threads to access that data in the same NUMA node. These optimizations minimize memory access latencies and reduce memory interconnect traffic.

NUMA topology in Hyper-V virtual machines

When a virtual machine is started, Hyper-V attempts to allocate all the memory for that virtual machine from a single physical NUMA node, if sufficient memory is available. If the memory

requirements for the virtual machine cannot be satisfied from a single node, Hyper-V allocates memory from another physical NUMA node. This is known as NUMA spanning. For more information about NUMA Spanning, see the Host settings – NUMA Spanning section later in this document.

Virtual NUMA

To enable virtualizing large scale-up workloads, Hyper-V in Windows Server 2012 offers expanded virtual machine scale limits. A single virtual machine can be assigned up to 64 virtual processors and 1 TB of memory. When creating much larger virtual machines, memory from multiple NUMA nodes on the host system will likely be utilized. In such a virtual machine configuration, if the virtual processors and memory continue to be allocated from a single NUMA node, workloads would not be taking advantage of NUMA optimizations, and performance will suffer.

In Windows Server 2012, Hyper-V presents a virtual NUMA topology to virtual machines. By default, this virtual NUMA topology is optimized to match the NUMA topology of the underlying host computer. Exposing a virtual NUMA topology into a virtual machine allows the guest operating system and any NUMA-aware applications running within it to take advantage of the NUMA performance optimizations, just as they would when running on a physical computer.

There is no distinction between a virtual and a physical NUMA from the workload's perspective. Inside a virtual machine, when a workload allocates local memory for data, and accesses that data in the same NUMA node, fast local memory access results on the underlying physical system. Performance penalties due to remote memory access are successfully avoided.

Configuring Virtual NUMA

Virtual NUMA is configured in the following locations:

- Host settings – NUMA Spanning
- Virtual machine settings - NUMA

Host settings – NUMA Spanning

This setting affects all virtual machines on the server running Hyper-V. The NUMA spanning setting determines virtual machine behavior when there is not enough memory on the host to perform an optimal mapping between the virtual NUMA nodes and the host's physical NUMA nodes. This setting affects the following virtual machine operations:

- Startup
- Restore
- Migration

When NUMA spanning is enabled and one or more virtual NUMA nodes must be split between two or more physical NUMA nodes due to memory constraints, the virtual machines are still allowed to start, or be restored or migrated to another host. This splitting, however, results in a suboptimal mapping between virtual and physical NUMA nodes. Even NUMA-aware workloads

may perform sub optimally; for example, memory access within the virtual machine that appears to be local can actually become remote on the host's physical NUMA topology. Additionally, because mapping virtual NUMA nodes to physical NUMA nodes can vary based on memory availability at the time the virtual machine is started, performance may be inconsistent across virtual machine restarts.

This virtual NUMA to physical NUMA split can occur in the following scenarios:

- A virtual machine has been configured with virtual NUMA nodes larger than the host's physical NUMA nodes. This might be caused by manual configuration of the virtual machine's NUMA settings, or because the virtual machine was initially created on a server running Hyper-V and then moved to a server that has a different physical NUMA topology.
- The current distribution of available memory across the physical NUMA nodes requires that at least one virtual NUMA node be backed by memory on multiple physical NUMA nodes. Uneven distribution of available memory could be caused by the memory usage of other virtual machines or by the host computer system.

Important

NUMA spanning is enabled by default.

When NUMA Spanning is disabled, Hyper-V will not start, restore, or accept a live migration of a virtual machine if it would be forced to split a virtual NUMA node between two or more physical NUMA nodes. This ensures that virtual machines always run with optimal performance. Each virtual NUMA node is backed by one physical NUMA node, and the virtual NUMA topology aligns exactly with the physical server's NUMA topology. This setting is recommended for NUMA-aware workloads that require optimal performance at all times, but it might cause a virtual machine to not start, restore, or migrate.

Note

The NUMA spanning setting applies to all the virtual machines on the server running Hyper-V.

The following table lists the benefits and disadvantages of enabling or disabling NUMA spanning.

NUMA Spanning	Benefits	Disadvantages
Enabled (default setting)	• Virtual machines can utilize whatever memory is available on the host, regardless of the distribution across physical NUMA nodes. • Virtual machines will start regardless of their virtual NUMA topology, even if it does not align well with the host's physical NUMA	• Virtual machine performance may vary between virtual machine restarts. • Virtual NUMA alignment may vary between boots. • NUMA-aware workloads may not perform optimally because virtual NUMA topologies may not align well with physical NUMA

NUMA Spanning	Benefits	Disadvantages
	topology. • Virtual machines with Dynamic Memory enabled can use all available system resources, even if those resources come from multiple physical NUMA nodes.	topologies.
Disabled	• NUMA-aware workloads perform optimally because virtual NUMA topologies will align with physical NUMA topologies. • Virtual machine performance is consistent between reboots.	• Virtual machines will fail to start if any virtual NUMA node cannot be placed entirely within a single physical NUMA node. • Virtual machines may fail to migrate to a given host if their virtual NUMA topologies aren't configured to align with the target host's physical NUMA topology. • Virtual machines with Dynamic Memory enabled cannot use more processors or memory than available in a single physical NUMA node.

Virtual machine settings - NUMA

The virtual NUMA topology can be configured for each individual virtual machine. The maximum amount of memory and the maximum number of virtual processors in each virtual NUMA node can be configured, in addition to the maximum number of virtual NUMA nodes presented in each virtualized socket. By default, these values are set to align with the host's physical NUMA topology.

Maximum values are set to the highest that allow a single virtual NUMA node to be contained entirely within a physical NUMA node on the host. These default values account for virtualization overheads, such as memory that is required by virtual devices. If changes have been made to these settings or the virtual machine has been moved to host with a different physical NUMA topology, you can restore the default virtual NUMA topology by clicking **Use Hardware Topology** on the **NUMA Configuration** page in the settings of a virtual machine.

If a virtual machine will be moved or migrated between multiple hosts with different physical NUMA topologies, and NUMA spanning has been disabled, configure the virtual NUMA setting for each virtual machine to the lowest default value that is available across the potential hosts. This ensures the virtual NUMA topology for the virtual machine will fit within the physical NUMA topology on each host that it could be migrated to—which is required when NUMA spanning is disabled.

Virtual NUMA and Dynamic Memory

Virtual NUMA and Dynamic Memory features cannot be used at the same time. A virtual machine that has Dynamic Memory enabled effectively has only one virtual NUMA node, and no NUMA topology is presented to the virtual machine regardless of the virtual NUMA settings. As a result, the NUMA spanning setting must be enabled when you have virtual machines that are large enough to be split between two or more physical NUMA nodes. If the NUMA spanning setting is disabled, the virtual machine must fit entirely within a single physical NUMA node, or the virtual machine will not start, or be restored or migrated.

Workloads that are not NUMA-aware will not take advantage of virtual NUMA. However, the guest operating system may perform some NUMA optimization. Enabling Dynamic Memory (therefore presenting only a single virtual NUMA) should not cause performance degradation.

Similarly, virtual machines that are not large enough to consume resources from multiple physical NUMA nodes receive no benefit from enabling virtual NUMA. Enabling Dynamic Memory, and so presenting only a single virtual NUMA node, will not cause any performance-related degradation.

When trying to decide which feature to use, you should take the following questions into consideration. If the answer to both is yes, enable virtual NUMA and do not enable Dynamic Memory.

1. Is the workload running in the virtual machine NUMA-aware?
2. Will the virtual machine consume more resources, processors, or memory than are available on a single physical NUMA node?

Important

- If a virtual machine is configured through Hyper-V Manager, Windows PowerShell, or the virtualization v1 namespace for WMI (root\virtualization\v1), enabling Dynamic Memory automatically disables virtual NUMA. Similarly, disabling Dynamic Memory automatically enables virtual NUMA when the virtual machine has enough virtual processors to warrant it.
- If a virtual machine is configured through the virtualization v2 namespace for WMI (root\virtualization\v2), the two features are configured independently, and an error will result if both features are enabled simultaneously.

For more information about Dynamic Memory, see Hyper-V Dynamic Memory Overview topic.

Virtual NUMA FAQs

This section explains some questions that are commonly asked about virtual NUMA.

Why is NUMA virtualized in Windows Server 2012?

Prior to Windows Server 2012, all virtual processors and memory were presented to a virtual machine as if they belonged to the same NUMA node, irrespective of the host's physical hardware topology. At the time, this was reasonable given that virtual machines could be configured with a maximum of four virtual processors.

In Windows Server 2012, virtual machines can have up to 64 virtual processors and 1 TB of memory. At this scale, virtual machines must utilize memory and processors from multiple physical NUMA nodes simultaneously. If the virtual processors and memory continued to be presented to the virtual machine as belonging to the same NUMA node, performance (especially for NUMA-aware scale up workloads) would suffer.

Do all virtual machines take advantage of virtual NUMA?

No, only virtual machines running NUMA-aware workloads that are large enough to require resources from more than one physical NUMA node, and are not configured to use Dynamic Memory, will benefit from virtual NUMA.

How much of a benefit does virtual NUMA provide?

A number of factors contribute, such as the architecture of the host's memory bus and processors, the topology to which the memory and processors are connected, how the memory is physically installed on the motherboards, the memory allocation techniques in use, and the workload's memory access patterns. The simplest way to find out is to benchmark your workload with virtual NUMA disabled, and then again with virtual NUMA enabled.

Can virtual NUMA and Dynamic Memory be used together?

No. Virtual NUMA and Dynamic Memory cannot be used together by a virtual machine. When Dynamic Memory is enabled, a single virtual NUMA node will be presented to the virtual machine regardless of the virtual NUMA settings, and virtual NUMA is disabled. When trying to decide which feature to use, you should consider the following questions. If the answer to both is yes, enable virtual NUMA.

1. Is the workload running in the virtual machine NUMA-aware?
2. Will the virtual machine consume more resources, processors, or memory than are available on a single physical NUMA node?

How can I tell if my system in NUMA-based?

The **Get-VMHostNumaNode** Windows PowerShell cmdlet gets the NUMA topology of a Hyper-V host and returns an object for each of the host's NUMA nodes. If the results return more than one NUMA node, your host is NUMA based. However, if the results return one node, your host is not NUMA based.

Example of a NUMA based host

```
PS C:\> Get-VMHostNumaNode



NodeId                 : 0
ProcessorsAvailability : {94, 99, 100, 100}
MemoryAvailable        : 3196
MemoryTotal            : 4093
ComputerName           : ServerName_01


NodeId                 : 1
ProcessorsAvailability : {66, 73, 98, 99}
MemoryAvailable        : 2533
MemoryTotal            : 4096
ComputerName           : ServerName_01
```

How can I determine if a virtual machine's performance is being impacted by NUMA spanning?

Hyper-V provides performance counters that enable you to judge how well a virtual machine's virtual NUMA nodes align with the host's physical NUMA nodes. You can look at the following counters:

- **Hyper-V VM Vid Partition : Remote Physical Pages** Indicates how many of a virtual machine's memory pages were allocated on a non-aligned physical NUMA node.

- **Hyper-V VM Vid Partition : Physical Pages Allocated** Indicates how many memory pages a virtual machine has in total.

If the **Remote Physical Pages** counter is zero, the virtual machine is perfectly aligned. If the **Remote Physical Pages** counter is not zero, comparing it with the total number of memory pages. The **Physical Pages Allocated** counter enables you to determine the percentage of the virtual machine's memory that is split between more than one physical NUMA node.

Performance counters information can be retrieved through Performance monitor or the **Get-Counter** cmdlet.

Example

In this example, the **Get-Counter** cmdlet reveals that a virtual machine's virtual NUMA topology exactly matches the host's physical NUMA topology.

```
PS C:\> Get-Counter "\Hyper-V VM VID Partition(*)\*"
```

```
Timestamp                 CounterSamples
---------                 --------------
5/30/2013 2:42:04 PM      \\ServerName_01\hyper-v vm vid partition(_total)\remote
physical pages :
                          0

                          \\ServerName_01\hyper-v vm vid partition(_total)\preferred numa
node index :
                          1

                          \\ServerName_01\hyper-v vm vid partition(_total)\physical pages
allocated :
                          526344

                          \\ServerName_01\hyper-v vm vid partition(Windows® server 2012 -
server 1)\remote physical pages :
                          0

                          \\ServerName_01\hyper-v vm vid partition(Windows® server 2012 -
server 1)\preferred numa node index :
                          1

                          \\ServerName_01\hyper-v vm vid partition(Windows® server 2012 -
server 1)\physical pages allocated :
                          526344
```

What can I do if a performance counters indicate that a virtual machine's virtual NUMA nodes do not align well with the Hyper-V host's physical NUMA topology?

First, determine if the virtual machine's workload is performing well. If it's already performing well, the workload may not depend heavily on memory access latency, or your host may provide remote memory access latencies that are low enough that the misalignment does not cause any impact.

If the workload is not performing optimally, and performance counters indicate a large proportion of the virtual machine's memory is allocated on misaligned physical NUMA nodes, consider the following actions:

- Ensure that the configured NUMA topology of the virtual machine aligns well with the Hyper-V host. Try resetting the virtual machine's NUMA settings to their default values. For more information about resetting a virtual machine's NUMA settings, see the Virtual machine settings - NUMA section.
- If you have multiple Hyper-V hosts, consider moving the virtual machine to another host. The distribution of available memory across NUMA nodes on the current host may not be ideal with the memory that was allocated for the virtual machine.
- Consider disabling the Hyper-V NUMA spanning setting. For more information about the Hyper-V NUMA spanning setting, see the Host settings – NUMA Spanning section.

How can I see which physical NUMA nodes a virtual machine is using?

To see which physical NUMA nodes a virtual machine is using, copy and save the following script and run it on your server running Hyper-V. See the example for the required script syntax.

Example

In this example, the virtual NUMA topology matches the physical NUMA topology. Each virtual NUMA node fits in a physical NUMA node without the need to split the virtual NUMA nodes between more than one physical NUMA node. Virtual NUMA node number 0 resides in physical NUMA node 3.

```
PS C:\> .\Users\vmuser\Desktop\GetBackingPhyNumaNode.ps1 -VmName NumaTest


VMName                                  : NumaTest
vNUMANode                               : 0,          1,           2
VMMem                                   : 6828,       6826,        6826
PhyNUMANode                             : 3,          1,           2
PhyNUMANodeMemRemaining                 : 1836,       1840,        1840
```

Sample code

```
param ($VmName=$(throw "Please provide the name of the virtual machine whose virtual to
physical NUMA mapping is to be found"))


$Namespace = 'root\virtualization\v2'


$RUNNING=2
$STOPPED=3
```

```
$VNUMA = New-Object System.Object
[Array]$vNodes = @()

function EditMember
{
    param ($Object=$(throw "Please provide the name of the class to be modified"), `
            $MemberName=$(throw "Please provide the name of the member to be modified"),
`
            $MemberValue=$(throw "Please provide the value with which the member is to be
modified"))

    if ($MemberValue -isnot [String])
    {
        $MemberValue = [System.Convert]::ToString($MemberValue)
    }

    $Error.Clear()
    $currValue = Get-Member -InputObject $Object -Name $MemberName
    if ($currValue -eq $null)
    {
        $Object | Add-Member -MemberType NoteProperty -Name $MemberName -Value
$MemberValue
    }
    else
    {
        $Object.$MemberName += ", " + $MemberValue
    }
}

$vms = $null

if ($VmName -eq "*")
{
```

```
    $vms = GWMI -n $Namespace Msvm_ComputerSystem | ? { $_.ElementName -ne $_.__SERVER }
}
else
{
    $vms = GWMI -n $Namespace Msvm_ComputerSystem | ? { $_.ElementName -eq $VmName }
}

if ($vms -eq $null)
{
    Throw "$VmName does not seem to exist on $($ENV:COMPUTERNAME)"
}

$vmIndex = 0
$vms | % {
    $VNUMA = New-Object System.Object

    $vm = $_
    $VNUMA | Add-Member -MemberType NoteProperty -Name VMName -Value $($vm.ElementName)
    $vmMem = GWMI -N $Namespace Msvm_Memory | ? { $_.Name.Contains($vm.Name) }
    $vmMemStg = GWMI -n $Namespace Msvm_MemorySettingData | ? {
$_.__PATH.Contains($vm.Name) }

    if ($($vm.EnabledState) -ne $RUNNING)
    {
        if ($vmMemStg -is [Array])
        {
            $vmMemStg = $vmMemStg[0]
        }
        EditMember -Object $VNUMA -MemberName vNUMANode -MemberValue "NA (Off)"
        EditMember -Object $VNUMA -MemberName VMMem -MemberValue
$($vmMemStg.VirtualQuantity)
        EditMember -Object $VNUMA -MemberName PhyNUMANode -MemberValue "NA"
        EditMember -Object $VNUMA -MemberName PhyNUMANodeMemRemaining -MemberValue "NA"
    }
```

```
    else
    {
        if ($vmMem -is [Array])
        {
            $vmMem | % {
                if (-not $_.DeviceID.EndsWith("Aggregate"))
                {
                    $memSrcNode = $_.GetRelated("Msvm_NumaNode")
                    EditMember -Object $VNUMA -MemberName vNUMANode -MemberValue
$($_.DeviceID.Split("\")[1])
                    EditMember -Object $VNUMA -MemberName VMMem -MemberValue
$($_.NumberOfBlocks)
                    if ($memSrcNode -ne $null)
                    {
                        EditMember -Object $VNUMA -MemberName PhyNUMANode -MemberValue
$($($($_.GetRelated("Msvm_NumaNode")).NodeId).Split("\")[1])
                        EditMember -Object $VNUMA -MemberName PhyNUMANodeMemRemaining -
MemberValue $($($_.GetRelated("Msvm_NumaNode")).CurrentlyConsumableMemoryBlocks)
                    }
                    else
                    {
                        EditMember -Object $VNUMA -MemberName PhyNUMANode -MemberValue
"Multi"
                        EditMember -Object $VNUMA -MemberName PhyNUMANodeMemRemaining -
MemberValue "NA"
                    }
                }
            }
        }
        else
        {
            $memSrcNode = $vmMem.GetRelated("Msvm_NumaNode")
            if ($memSrcNode -ne $null)
            {
```

```
                EditMember -Object $VNUMA -MemberName vNUMANode -MemberValue
$($vmMem.DeviceID.Split("\")[1])
                EditMember -Object $VNUMA -MemberName VMMem -MemberValue
$($vmMem.NumberOfBlocks)
                EditMember -Object $VNUMA -MemberName PhyNUMANode -MemberValue
$($($($vmMem.GetRelated("Msvm_NumaNode")).NodeId).Split("\")[1])
                EditMember -Object $VNUMA -MemberName PhyNUMANodeMemRemaining -
MemberValue $($($vmMem.GetRelated("Msvm_NumaNode")).CurrentlyConsumableMemoryBlocks)
            }
            else
            {
                EditMember -Object $VNUMA -MemberName vNUMANode -MemberValue "NA"
                if ($vmMemStg.DynamicMemoryEnabled -eq $true)
                {
                    EditMember -Object $VNUMA -MemberName VMMem -MemberValue "Dynamic"
                }
                else
                {
                    EditMember -Object $VNUMA -MemberName VMMem -MemberValue "Multi"
                }
                EditMember -Object $VNUMA -MemberName PhyNUMANode -MemberValue "NA"
                EditMember -Object $VNUMA -MemberName PhyNUMANodeMemRemaining -
MemberValue "NA"
            }
        }
    }

    $vNodes += $VNUMA
}

$vNodes
```

Why do I see that a virtual machine is assigned memory from different physical NUMA nodes than expected?

Other virtual machines running on the host will dictate which physical NUMA nodes can best be used to provide memory. When the NUMA spanning setting is enabled, Hyper-V may choose NUMA distributions that don't seem to match the physical host well, but this configuration makes sense in the context of the other running virtual machines. Disabling the NUMA spanning setting can make the NUMA distribution of memory more predictable, but virtual machines may fail to start if this predictable distribution cannot be supported with the currently available resources. For more information about NUMA spanning, see the Host settings – NUMA Spanning section.

Automatic Virtual Machine Activation

Automatic Virtual Machine Activation (AVMA) acts as a proof-of-purchase mechanism, helping to ensure that Windows products are used in accordance with the Product Use Rights and Microsoft Software License Terms.

AVMA lets you install virtual machines on a properly activated Windows server without having to manage product keys for each individual virtual machine, even in disconnected environments. AVMA binds the virtual machine activation to the licensed virtualization server and activates the virtual machine when it starts up. AVMA also provides real-time reporting on usage and historical data on the license state of the virtual machine. Reporting and tracking data is available on the virtualization server.

Practical Applications

On virtualization servers that are activated using Volume Licensing or OEM licensing, AVMA offers several benefits.

Server datacenter managers can use AVMA to do the following:

- Activate virtual machines in remote locations
- Activate virtual machines with or without an internet connection
- Track virtual machine usage and licenses from the virtualization server, without requiring any access rights on the virtualized systems

There are no product keys to manage and no stickers on the servers to read. The virtual machine is activated and continues to work even when it is migrated across an array of virtualization servers.

Service Provider License Agreement (SPLA) partners and other hosting providers do not have to share product keys with tenants or access a tenant's virtual machine to activate it. Virtual machine activation is transparent to the tenant when AVMA is used. Hosting providers can use the server logs to verify license compliance and to track client usage history.

System Requirements

AVMA only requires a virtualization server running Windows Server 2012 R2 Datacenter. The guest virtual machine operating system must be Windows Server 2012 R2 Datacenter, Windows Server 2012 R2 Standard, or Windows Server 2012 R2 Essentials.

How to implement AVMA

1. On a Windows Server 2012 R2 Datacenter virtualization server, install and configure the Microsoft Hyper-V Server role. For more information, see Install Hyper-V Server 2012.
2. Create a virtual machine and install a supported server operating system on it.
3. Install the AVMA key in the virtual machine. From an elevated command prompt, run the following command:

   ```
   slmgr /ipk <AVMA_key>
   ```

The virtual machine will automatically activate the license against the virtualization server.

Tip

You can also employ the AVMA keys in any Unattend.exe setup file.

AVMA keys

The following AVMA keys can be used for Windows Server 2012 R2.

Edition	AVMA key
Datacenter	Y4TGP-NPTV9-HTC2H-7MGQ3-DV4TW
Standard	DBGBW-NPF86-BJVTX-K3WKJ-MTB6V
Essentials	K2XGM-NMBT3-2R6Q8-WF2FK-P36R2

Reporting and Tracking

The registry (KVP) on the virtualization server provides real-time tracking data for the guest operating systems. Because the registry key moves with the virtual machine, you can get license information as well. By default the KVP returns information about the virtual machine, including the following:

- Fully qualified domain name
- Operating system and service packs installed
- Processor architecture
- IPv4 an IPv6 network addresses
- RDP addresses

For more information about how to get this information, see Hyper-V Script: Looking at KVP GuestIntrinsicExchangeItems.

Note
KVP data is not secured. It can be modified and is not monitored for changes.

Important
KVP data should be removed if the AVMA key is replaced with another product key (retail, OEM, or volume licensing key).

Historical data about AVMA requests is available in a log file on the virtualization server (EventID 12310).

Since the AVMA activation process is transparent, error messages are not displayed. However, the following events are captured in a log file on the virtual machines (EventID 12309).

Notification	Description
AVMA Success	The virtual machine was activated.
Invalid Host	The virtualization server is unresponsive. This can happen when the server is not running a supported version of Windows.
Invalid Data	This usually results from a failure in communication between the virtualization server and the virtual machine, often caused by corruption, encryption, or data mismatch.
Activation Denied	The virtualization server could not activate the guest operating system because the AVMA ID did not match.

2.0

Online Virtual Hard Disk Resizing Overview

Starting in Windows Server® 2012 R2, Hyper-V delivers the ability to expand or shrink the size of a virtual hard disk while the virtual machine is still running. For more information about resizing virtual hard disks, see Configure Online Virtual Hard Disk Resize.

Key benefits

Storage administrators can avoid costly downtime by performing maintenance operations on running virtual hard disks. Shutting down virtual machines is no longer required and this eliminates interruption to users accessing those virtual machines and helps reduce maintenance

costs. Prior to Windows Server 2012 R2, if a virtual machine was online or a virtual hard disk was in use, it was not possible to perform maintenance on the virtual hard disk without temporarily shutting down the virtual machine.

Requirements

The following functionality is required for resizing a virtual hard disk:

- A server capable of running Hyper-V. The server must have processor support for hardware virtualization. The Hyper-V role must be installed.
- A user account that is a member of the local Hyper-V Administrators group or the Administrators group.

The following functionality is required for resizing a virtual hard disk:

- **VHDX** - the ability to expand and shrink virtual hard disks is exclusive to virtual hard disks that are using the .vhdx file format. Online resizing is supported for VHDX disk types, including fixed, differencing, and dynamic disks. Virtual hard disks that use the .vhd file format are not supported for resizing operations.
- **SCSI controller** - the ability to expand or shrink the capacity of a virtual hard disk is exclusive to .vhdx files that are attached to a SCSI controller. VHDX files that are attached to an IDE controller are not supported.

Understanding online resizing functionality

There are important concepts to understand when increasing or decreasing the capacity of a virtual hard disk.

Expanding a virtual hard disk

Expanding a virtual hard disk increases the disk capacity of the virtual hard disk. However, to make the additional disk space available to the virtual machine requires some extra configuration. From the perspective of the virtual machine, the virtual hard disk expansion is reflected under **Disk Manager** as an unallocated disk volume. The size of this unallocated volume is the difference between the original virtual hard disk and the nominated size of the expanded virtual hard disk.

To make the full virtual hard disk capacity available to the virtual machine, you need to use **Disk Manager** to expand the volume within the virtual machine. You can do this by using the Extend Volume Wizard within **Disk Manager**. After this is complete, you will be able to view the expanded disk capacity in the operating system of the virtual machine.

Shrinking a virtual hard disk

As expected, when you shrink a virtual hard disk, the virtual hard disk capacity is decreased. However, there is a limit to amount of disk space that can be adjusted when you shrink a virtual hard disk. That limit is defined by the size of the volume that is currently in use by the virtual

machine. The **Hyper-V Manager** user interface shows the minimum disk size available for shrinking a virtual hard disk.

If you want to shrink the size of the virtual hard disk below the size of the currently used volume within the virtual machine, you must first use **Disk Manager** in the operating system of the virtual machine to shrink the volume. After that is complete, you can reduce the size of the virtual hard disk.

Performing online resizing operations

In Windows Server 2012 R2, the option to resize a virtual hard disk is available through the Edit Virtual Hard Disk Wizard. However, the user interface option to shrink a virtual hard disk is only visible for virtual hard disks that have been previously expanded.

Administrators can use the Windows PowerShell interface for online virtual hard disk resizing operations.

Storage Quality of Service for Hyper-V

Starting in Windows Server® 2012 R2, Hyper-V includes the ability to set certain quality-of-service (QoS) parameters for storage on the virtual machines. For more information about configuring storage QoS for Hyper-V, see Configure Storage Quality of Service.

Storage QoS provides storage performance isolation in a multitenant environment and mechanisms to notify you when the storage I/O performance does not meet the defined threshold to efficiently run your virtual machine workloads.

Key benefits

Storage QoS provides the ability to specify a maximum input/output operations per second (IOPS) value for your virtual hard disk. An administrator can throttle the storage I/O to stop a tenant from consuming excessive storage resources that may impact another tenant.

An administrator can also set a minimum IOPS value. They will be notified when the IOPS to a specified virtual hard disk is below a threshold that is needed for its optimal performance.

The virtual machine metrics infrastructure is also updated, with storage related parameters to allow the administrator to monitor the performance and chargeback related parameters.

Maximum and minimum values are specified in terms of normalized IOPS where every 8 K of data is counted as an I/O.

Key features

Storage QoS allows administrators to plan for and gain acceptable performance from their investment in storage resources Administrators can:

- Specify the maximum IOPS allowed for a virtual hard disk that is associated with a virtual machine.
- Receive a notification when the specified minimum IOPS for a virtual hard disk is not met.
- Monitor storage-related metrics through the virtual machine metrics interface.

Requirements

Storage QoS requires that the Hyper-V role is installed. The Storage QoS feature cannot be installed separately. When you install Hyper-V, the infrastructure is enabled for defining QoS parameters associated with your virtual hard disks.

Note

Storage QoS is not available if you are using shared virtual hard disks.

Technical overview

Virtual hard disk maximum IOPS

Storage QoS provides the following features for setting maximum IOPS values (or limits) on virtual hard disks for virtual machines:

- You can specify a maximum setting that is enforced on the virtual hard disks of your virtual machines. You can define a maximum setting for each virtual hard disk.
- Virtual disk maximum IOPS settings are specified in terms of normalized IOPS. IOPS are measured in 8 KB increments.
- You can use the WMI interface to control and query the maximum IOPS value you set on your virtual hard disks for each virtual machine.
- Windows PowerShell enables you to control and query the maximum IOPS values you set for the virtual hard disks in your virtual machines.
- Any virtual hard disk that does not have a maximum IOPS limit defined defaults to 0.
- The Hyper-V Manager user interface is available to configure maximum IOPS values for Storage QoS.

Virtual hard disk minimum IOPS threshold notifications

Storage QoS provides the following features for setting minimum values (or reserves) on virtual hard disks for virtual machines:

- You can define a minimum IOPS value for each virtual hard disk, and an event-based notification is generated when the minimum IOPS value is not met.
- Virtual hard disk minimum values are specified in terms of normalized IOPS. IOPS are measured in 8 KB increments.

- You can use the WMI interface to query the minimum IOPS value you set on your virtual hard disks for each virtual machine.
- Windows PowerShell enables you to control and query the minimum IOPS values you set for the virtual hard disks in your virtual machines.
- Any virtual hard disk that does not have a minimum IOPS value defined will default to 0.
- The Hyper-V Manager user interface is available to configure minimum IOPS settings for Storage QoS.

Generation 2 Virtual Machine Overview

The generation of a virtual machine determines the virtual hardware and functionality that is presented to the virtual machine. In Hyper-V there are two supported virtual machine generations, generation 1 and generation 2. Generation 2 virtual machines have a simplified virtual hardware model, and supports Unified Extensible Firmware Interface (UEFI) firmware instead of BIOS-based firmware. The majority of legacy devices have also been removed from generation 2 virtual machines.

This topic describes the various benefits to a generation 2 virtual machine, and it lists the requirements for generation 2 virtual machines.

In previous versions of Hyper-V, there was only one generation of virtual machine. In Windows Server 2012 R2 there are two generations of virtual machines to choose from when you create a new virtual machine:

- **Generation 1** Provides the same virtual hardware to the virtual machine as in previous versions of Hyper-V.
- **Generation 2** Provides the following new functionality on a virtual machine:
  - PXE boot by using a standard network adapter
  - Boot from a SCSI virtual hard disk
  - Boot from a SCSI virtual DVD
  - Secure Boot (enabled by default)
  - UEFI firmware support

Note
IDE drives and legacy network adapter support has been removed.

Requirements

The following guest operating systems are supported as generation 2 virtual machines:

- Windows Server 2012 R2
- Windows Server 2012

- 64-bit versions of Windows 8.1
- 64-bit versions of Windows 8

Caution
After a virtual machine is created, you cannot change its generation.

Generation 2 virtual machine features

PXE boot by using a standard network adapter

In previous versions of Hyper-V, if you wanted to perform a remote installation of the guest operating system by using PXE boot, you were required to install a legacy network adapter. Legacy network adapters are not available in generation 2 virtual machines. Generation 2 virtual machines support PXE boot by using a standard network adapter.

Boot from SCSI controller

In previous versions of Hyper-V, you could not boot a virtual machine from a SCSI-attached virtual hard disk or from a DVD.

Generation 2 virtual machines can boot from a virtual hard disk or DVD that is attached to the SCSI controller. The virtual Integrated Device Electronics (IDE) controller is not available in generation 2 virtual machines.

Secure Boot

Secure Boot is a feature that helps prevent unauthorized firmware, operating systems, or UEFI drivers (also known as option ROMs) from running at boot time.

Note
Secure Boot is enabled by default for generation 2 virtual machines. This can be modified after you create the virtual machine. For more information, see **Secure Boot**.

Device support comparison

Following is a device comparison between generation 1 and generation 2 virtual machines.

Generation 1 Device	Generation 2 Replacement	Generation 2 Enhancements
IDE controller	Virtual SCSI controller	Boot from .vhdx (64 TB maximum size, and online resize capability)
IDE CD-ROM	Virtual SCSI CD-ROM	Support for up to 64 SCSI

Generation 1 Device	Generation 2 Replacement	Generation 2 Enhancements
		DVD devices per SCSI controller.
Legacy BIOS	UEFI firmware	Secure Boot
Legacy network adapter	Synthetic network adapter	Network boot with IPv4 and IPv6
Floppy controller and DMA controller	No floppy controller support	
Universal asynchronous receiver/transmitter (UART) for COM ports	Optional UART for debugging	Faster and more reliable
i8042 keyboard controller	Software-based input	Uses fewer resources because there is no emulation. Also reduces the attack surface from the guest operating system.
PS/2 keyboard	Software-based keyboard	Uses fewer resources because there is no emulation. Also reduces the attack surface from the guest operating system.
PS/2 mouse	Software-based mouse	Uses fewer resources because there is no emulation. Also reduces the attack surface from the guest operating system.
S3 video	Software-based video	Uses fewer resources because there is no emulation. Also reduces the attack surface from the guest operating system.
PCI bus	No longer required	
Programmable interrupt controller (PIC)	No longer required	
Programmable interval timer	No longer required	
Super I/O device	No longer required	

Generation 2 FAQ

The following are some frequently asked questions regarding generation 2 virtual machines.

Can I run generation 1 and generation 2 virtual machines together?

Yes. This is expected because generation 2 does not support all of the guest operating systems that support Hyper-V. Guest operating systems that are not supported by generation 2 virtual machines need to be installed in a generation 1 virtual machine.

Is there a performance gain to generation 2 virtual machines?

In the daily running of a virtual machine, there is no performance difference between generation 1 and generation 2 virtual machines. You might see a slight improvement in virtual machine boot and installation times with generation 2 virtual machines.

What are the benefits of using a generation 2 virtual machine?

Following are the benefits of using a generation 2 virtual machine:

- Can boot the virtual machine from a SCSI device or a standard network adapter.
- Helps prevent unauthorized firmware, operating systems, or UEFI drivers from running at boot time when Secure Boot is enabled.

How can I enable kernel debugging by using a COM port on a generation 2 virtual machine?

By default, COM ports are not visible in generation 2 virtual machines. You can configure COM ports by using Windows PowerShell or Windows Management Instrumentation (WMI).

To enable kernel debugging by using a COM port, follow these steps:

1. Disable Secure Boot. Kernel debugging is not compatible with Secure Boot, so Secure Boot must be disabled in the virtual machine. Use the **Set-VMFirmware** Windows PowerShell cmdlet to disable Secure Boot when the virtual machine is in an Off state. For example, the following command will disable Secure Boot on virtual machine TestVM:

   ```
   Set-VMFirmware –Vmname TestVM –EnableSecureBoot Off
   ```

2. Configure a COM port. Use the **Set-VMComPort** Windows PowerShell cmdlet to add a COM port to the virtual machine. For example, the following command will configure the first COM port on virtual machine, TestVM, to connect to the named pipe, TestPipe, on the local computer:

   ```
   Set-VMComPort –VMName TestVM 1 \\.\pipe\TestPipe
   ```

Note
Configured COM ports are not visible when looking at the settings of a virtual machine in the Hyper-V Manager console.

Is RemoteFX supported by generation 2 virtual machines?

No. RemoteFX is not supported by generation 2 virtual machines.

Can I attach a physical CD or DVD to a generation 2 virtual machine?

No. The virtual DVD drive in generation 2 virtual machines only supports ISO image files.

Can I attach a virtual hard disk in VHD format to a generation 2 virtual machine?

No. Generation 2 virtual machines only support VHDX format virtual hard drives. You can convert VHD files to VHDX format using the Edit Virtual Hard Disk Wizard in the Hyper-V Manager console or use the **Convert-VHD** Windows PowerShell cmdlet.

Can a VHDX file that was converted from a VHD file be used to boot a generation 2 virtual machine?

No. A clean installation of the operating system is required.

Can I resize a VHDX file that contains the boot volume for a generation 2 virtual machine while the virtual machine is running?

Yes.

What is the maximum size of a VHDX file that is supported by a generation 2 virtual machine?

64 TB

Can I create a VHDX file that can be used to boot both generation 1 and generation 2 virtual machines?

Although this is technically possible, it is possible that unexpected errors can occur. Therefore, a VHDX file that is bootable to both generation 1 and generation 2 virtual machines is not supported or recommended.

Can I mount a VHDX file that is used by a generation 2 virtual machine?

Yes. A VHDX file is a container for a disk, and as such, it can be mounted regardless of the virtual machine generation.

Can I tell whether a VHDX file was created by a generation 1 or a generation 2 virtual machine?

No. The virtual machine generation is a property of the virtual machine, not a property of the virtual hard disk.

Can a generation 2 VHDX file be attached to a generation 1 virtual machine?

Yes. The VHDX file can be attached to the IDE controller or the SCSI controller of a generation 1 virtual machine. However, if this is a bootable VHDX file, the generation 1 virtual machine will not boot.

Can I import a generation 2 virtual machine to Hyper-V in Windows Server 2012?

No. Generation 2 virtual machine is a new feature in Windows Server 2012 R2.

Is Secure Boot or UEFI firmware required on the physical host?

No. The virtual firmware and its configuration are independent of the physical platform.

Does the UEFI firmware in a generation 2 virtual machine support setup mode for Secure Boot?

No.

Can I run a UEFI shell or other UEFI application in a generation 2 virtual machine?

Microsoft does not provide or support a UEFI shell or other UEFI applications. Using a non-Microsoft UEFI shell or UEFI applications is technically possible if they are compiled directly from the sources. If these applications are not appropriately digitally signed, you will be required to disable Secure Boot for the virtual machine.

Where are the boot entries for generation 2 virtual machines stored?

The boot entries for a generation 2 virtual machine are stored in the configuration file for the virtual machine.

What is the default generation of a new virtual machine?

By default, generation 1 virtual machines are created by using the New Virtual Machine Wizard or the **New-VM** Windows PowerShell cmdlet . In the New Virtual Machine Wizard, you can select **Generation 2** on the **Specify Generation** page. In Windows PowerShell, you can use the **–Generation** parameter to specify the generation of a virtual machine.

Can I change the generation of a virtual machine after it has been created?

No.

Why are 64-bit versions of Windows Server 2008 R2 and Windows 7 not supported as a guest operating system for generation 2 virtual machines?

Although Windows Server 2008 R2 and Windows 7 support UEFI, they depend on a programmable interrupt controller (PIC), which is not present in generation 2 virtual machine hardware.

Why are 32-bit versions of Windows 8.1 and Windows 8 not supported as guest operation systems for generation 2 virtual machines?

Generation 2 virtual machines are not supported on 32-bit guest operating systems because the 64-bit-only UEFI firmware does not have a compatibility support module.

What is the maximum number of network adapters that can be supported by a generation 2 virtual machine?

Eight. In addition, generation 2 virtual machines do not support legacy network adapters.

What is the maximum number of storage devices (VHDX files or DVD drives), that can be supported by generation 2 virtual machines?

256. Generation 2 virtual machines support up to four SCSI controllers, with each controller supporting 64 devices. Generation 2 virtual machines can have more than four DVD drives.

When I create a new virtual machine with the New-VM Windows PowerShell cmdlet, why does the generation 2 virtual machine not have a DVD drive?

On a generation 2 virtual machine, you can add a DVD drive while the virtual machine is running.

Does a generation 2 virtual machine support iSCSI or Fibre Channel boot?

No.

Can I perform a network installation that uses IPv6?

Yes. By default, generation 2 virtual machines use IPv4. You can set a virtual machine to use IPv6 instead of IPv4 by using the **Set-VMFirmware** Windows PowerShell cmdlet. For example, the following command will set the preferred protocol to IPv6 for a virtual machine named TestVM:

```
Set-VMFirmware –VMName TestVM –IPProtocolPreference IPv6
```

Can I use a virtual floppy disk (.vfd) as an answer file for an unattended installation of a generation 2 virtual machine?

No. Generation 2 virtual machines do not support virtual floppy disks. As an alternate, use an ISO image file in a DVD drive.

Virtual Machine Specifications for Hyper-V in Windows Server 2012 R2

Hyper-V in Windows Server 2012 R2 and Client Hyper-V in Windows 8.1 introduce generation 2 virtual machines that have a revised set of virtual hardware. The virtual hardware that was exposed in Windows Server 2012 and Windows 8 is preserved as generation 1 virtual machines.

Generation 1 virtual machine components

The following table lists basic virtual machine components and their corresponding emulated hardware. Collapse this tableExpand this table

Component	Purpose
Basic input/output system (BIOS)	Implements American Megatrends (AMI) BIOS on an emulated Intel 440BX chipset with PIIX4 on-board components:

	• Battery-backed CMOS • Real-Time clock • Two 8327 DMA controllers • I/O APIC • PCI-to-ISA bridge (also known as PIIX or PCI-to-ISA / IDE Xcelerator) • PCI bus as found in the Intel 440BX chipset • Host-to-PCI bridge as found in the Intel 440BX chipset • Two cascaded 8259 programmable interrupt controllers (PICs) • Programmable Interval Timer (PIT) • Power management hardware as provided in the Intel 440BX chipset • SMBus • Two 16550 UART controllers that support up to two serial ports • Winbond 83977 Super I/O chipset • Keyboard Controller Device (i8042) • Standard Microsoft® 101-Key PS-2 Keyboard or PS/2 Mouse • ISA Bus • PIC Device: Generic VESA 2.0 Video Device • DEC 21140 Ethernet Adapter • Floppy disk controller We recommend that the guest operating system implements the Advanced Configuration and Power Interface (ACPI) specification for best performance.
Floppy disk drive	Supports a single 1.44 MB floppy disk drive and mapping floppy drive images.
Serial (COM) port	Emulates up to two serial ports that can be mapped to local named pipes and files.
Mouse	Emulates a standard PS/2 Microsoft IntelliMouse pointing device.
Keyboard	Emulates a standard PS/2 101-key Microsoft keyboard.

Network adapter (multifunction)	Emulates the multiport DEC 21140 10/100TX 100 MB Ethernet network adapter with one to four network connections. In some cases, the DEC 21140 network adapter might be displayed on the virtual machine as an Intel 21140 network adapter. They are equivalent network adapters. The virtual network adapters and the network driver that controls them do not support the virtual local area network (VLAN) identifier (ID) in a tag header. Up to four such adapters can be added to a virtual machine.
Memory	Supports 1 terabyte (TB) of RAM per virtual machine.
Video card	Emulates a generic graphics adapter with 4 MB of Video RAM (VRAM), VGA, and SVGA support that is compliant with VESA 2.0, a hardware cursor, and support for Microsoft DirectX.
IDE/ATAPI storage	Emulates up to four IDE devices, hard drives, or CD-ROM or DVD-ROM drives (or ISO images), and virtual hard disks up to 2040 MB in size per IDE channel.

The following table lists software components that are designed for optimal performance in virtualized environments.

Component	**Purpose**
Microsoft Virtual Machine Bus	Provides a high-speed Hyper-V-specific communication channel between a virtual machine and a host that is installed as a part of integration services.
Microsoft Virtual Machine Bus Network Adapter	Provides a network adapter that can be added to the virtual machine when the integration services are installed in the operating system on the virtual machine.
Hyper-V-specific SCSI Controller	Provides a SCSI adapter that can be added to the virtual machine when the integration

	services are installed in the operating system on the virtual machine. It supports up to 64 devices per controller with a maximum of four SCSI controllers per virtual machine.
Physical disk attached to the virtual machine (also known as pass-through storage)	Provides mechanisms to enable virtual machines to read and write directly to a storage device accessible to the virtual machine. The virtual machine treats it as a disk.
Hyper-V-specific Video Card	Provides a Hyper-V-specific video card that is added when the integration services are installed in the operating system on the virtual machine.
Hyper-V-specific Mouse	Provides a Hyper-V-specific mouse that is added when the integration services are installed in the operating system on the virtual machine.
Virtual Fibre Channel Adapter	Provides a Hyper-V-specific virtual Fibre Channel adapter that is added to the guest operating system when the integration services are installed in the operating system on the virtual machine and a virtual Host Bus Adapter (HBA) is added to the virtual machine.
Virtual PCI Bus	Provides a Hyper-V-specific virtual PCI bus that is added when the integration services are installed in the operating system on the virtual machine and Single Root I/O Virtualization (SR-IOV) networking is being used.
Hyper-V-specific Generation Counter device	Provides a Hyper-V-specific virtual machine Generation Counter device that is added when the integration services are installed in the operating system on the virtual machine. The device provides a 128-bit virtual machine Generation Identifier for a virtual machine that stays the same unless the virtual machine is reverted to a snapshot.
Automatic Virtual Machine Activation	Enables the Windows Server 2012 R2 guest operating systems to automatically activate themselves when running on a Windows Server 2012 R2 Datacenter host, which makes it easier and faster to deploy Microsoft cloud

	solutions that are based on Windows Server 2012 R2 Datacenter.
Enhanced Session Mode	Enables Hyper-V to redirect local resources to a virtual machine session through the Virtual Machine Connection tool. The enhanced session mode connection uses a Remote Desktop Connection session via the virtual machine bus (VMBus), so that no network connection to the virtual machine is required.

Generation 2 virtual machine components

The following table lists basic virtual machine components in a generation 2 virtual machine.

Component	Purpose
UEFI firmware	Provides system firmware.
Serial (COM) ports	Provides optional emulated serial ports that can be mapped to local named pipes and files.
Real-Time clock	Provides emulated time support.
I/O APIC	Provides an emulated Advanced Programmable Interrupt Controller (I/O APIC).
Memory	Supports 1 TB of RAM per virtual machine.

The following table lists software components that are designed for optimal performance in virtualized environments.

Component	Purpose
Microsoft Hyper-V Virtual Machine Bus	Provides a communication channel between a virtual machine and the host.
Microsoft Hyper-V Network Adapter	Supports network input and output.
Microsoft Hyper-V SCSI Controller	Supports storage input and output.
Physical disk attached to the virtual machine (also known as pass-through storage)	Provides mechanisms to enable virtual machines to read and write directly to a storage device that is accessible to the virtual machine. The virtual machine treats it as a disk.

Microsoft Hyper-V Video	Supports video output.
Microsoft Hyper-V Input	Supports human interface device (HID)-compliant mouse input.
Microsoft Hyper-V Virtual Keyboard	Supports keyboard input.
Microsoft Hyper-V Virtual Fibre Channel Adapter	Supports fibre-channel-based input and output.
Microsoft Hyper-V Virtual PCI Bus	Supports SR-IOV network input and output.
Microsoft Hyper-V Generation Counter device	Provides a 128-bit Generation Identifier that stays the same unless the virtual machine is reverted to a checkpoint.
Microsoft Virtual DVD-ROM	Supports CD/DVD-ROM input.
Automatic Virtual Machine Activation	Enables Windows Server 2012 R2 guest operating systems to automatically activate themselves when running on a Windows Server 2012 R2 Datacenter host, which makes it easier and faster to deploy Microsoft cloud solutions that are based on Windows Server 2012 R2 Datacenter.
Enhanced Session Mode	Enables Hyper-V to redirect local resources to a virtual machine session through the Virtual Machine Connection tool. The enhanced session mode connection uses a Remote Desktop Connection session via the virtual machine bus (VMBus), so that no network connection to the virtual machine is required.

Hypervisor top-level functional specification (TLFS)

The Hypervisor top-level functional specification (TLFS) for the third-generation hypervisor by Microsoft specifies the externally visible behavior of the hypervisor, which is a component of the Hyper-V role in Windows Server 2012 R2.

The specification provides details on the functionalities of the hypervisor. It covers topics such as Feature and Interface Discovery, Hypercall Interfaces, Partition Management, Physical Hardware Management, Resource Management, Guest Physical Address Spaces, Intercepts, Virtual Processor Management, Virtual Processor Execution, Virtual MMU and Caching, Virtual Interrupt Control, Inter-Partition Communication, Timers, Message Formats, Partition Save and Restore, Scheduler, Event Logging, Guest Debugging Support, Statistics, Booting, and System Properties.

Virtual Machine Connection - Enhanced Session Mode Overview

Virtual Machine Connection in Hyper-V now supports redirecting local resources to a virtual machine session. This feature provides a functionality that is similar to a remote desktop connection when you are interacting with a virtual machine.

Enhanced session mode

In previous versions of Hyper-V, Virtual Machine Connection provided redirection of only the virtual machine screen, keyboard, and mouse with limited copy functionality. To get additional redirection abilities, a remote desktop connection to the virtual machine could be initiated, but this required a network path to the virtual machine.

Starting with Hyper-V in Windows Server 2012 R2, Hyper-V can now redirect local resources to a virtual machine session through Virtual Machine Connection tool. The enhanced session mode connection uses a Remote Desktop Connection session via the virtual machine bus (VMBus), so no network connection to the virtual machine is required.

The following local resources can be redirected when using Virtual Machine Connection.

- Display configuration
- Audio
- Printers
- Clipboard
- Smart cards
- USB devices
- Drives
- Supported Plug and Play devices

This feature is enabled by default in Client Hyper-V, and it is disabled by default on Hyper-V in Windows Server.

The following guest operating systems support enhanced session mode connections:

- Windows Server 2012 R2
- Windows 8.1

Practical applications

In previous versions of Hyper-V, the way to copy files to and from a virtual machine was by using a Remote Desktop Connection session or a network file copy, but these solutions require a working network connection. When connecting to a virtual machine by using the Virtual Machine Connection tool, you could copy only a limited amount of ASCII text to the virtual machine.

Enhanced session mode allows you to redirect local resources to a virtual machine in the same way that you can with a remote desktop connection, but you do not need a network connection to the virtual machine. You need to connect only to the server where the virtual machine is running.

Enhanced session mode can be useful in the following scenarios:

- Troubleshooting a virtual machine without a network connection to the virtual machine
- Signing in to the virtual machine by using a smart card
- Printing from a virtual machine to a local printer
- Testing and troubleshooting developer applications running in a virtual machine that require USB and sound redirection without using Remote Desktop Connection.

Configuring enhanced session mode

Configuring enhanced session mode is done in the following areas:

- Server settings:**Enhanced Session Mode Policy**
- User setting:**Enhanced Session Mode**
- Guest operating system

Server settings

The **Enhanced Session Mode Policy** setting affects all virtual machines on the server running Hyper-V.

The **Allow enhanced session mode** setting determines whether to allow enhanced session mode connections to virtual machines on the server running Hyper-V. If this setting is enabled, Hyper-V allows enhanced session mode connections to a virtual machine when the following conditions are met:

- An enhanced session mode connection has been requested by Virtual Machine Connection.
- The virtual machine is running an operating system that supports enhanced session mode.
- Remote Desktop Services is running on the virtual machine.

When the **Allow enhanced session mode** setting is disabled, connections from Virtual Machine Connection use a basic session.

Important
The default setting for the **Allow enhanced session mode** is:

- Hyper-V in Windows Server 2012 R2: Disabled
- Client Hyper-V: Enabled

User settings

The **Enhanced Session Mode** setting determines if Virtual Machine Connection attempts to use enhanced session mode when available in a guest operating system. When **Use enhanced session mode** is enabled, device redirection takes place when the following conditions are met:

- **Allow enhanced session mode** is enabled on the server running Hyper-V.

- The virtual machine is running an operating system that supports enhanced session mode.
- Remote Desktop Services is running on the virtual machine.

Guest operating system

The following guest operating systems support enhanced session mode connections:

- Windows Server 2012 R2
- Windows 8.1

Additionally, Remote Desktop Services needs to be running, and the user account you use to sign in to the virtual machine needs to be a member of the **Remote Desktop Users** local group or the local **Administrators** group.

Important

The Remote Desktop Services Group Policy setting and the local Group Policy settings are enforced for enhanced session mode connections. So if there is a policy in effect that does not allow the redirection of printers, printer redirection in a Virtual Machine Connection session is blocked as it would be for a remote desktop connection.

FAQs

The following are some frequently asked questions regarding enhanced session mode.

Does enhanced session mode require Terminal Services CALs?

No. Virtual Machine Connection creates an administrative remote desktop session from the client to the virtual machine. These sessions do not require Terminal Services CALs.

Can I save the settings for a connection to a virtual machine?

Yes. When you connect to a virtual machine that supports enhanced session mode, you will see the **Connect to** dialog box. In this dialog box, you can set which local resources will be available in the session with the virtual machine. To save your settings, click **Show Option**, select **Save my settings for future connections to this virtual machine**, and then click **Connect**. The settings will be saved in a file in the following folder:

%appdata%\roaming\Microsoft\Windows\Hyper-V\Client\1.0

A file will be created for each virtual machine with the following naming convention:

vmconnect.rdp.{virtual machine GUID}.config

What's New in Hyper-V for Windows Server 2012 R2

This topic explains the new and changed functionality of the Hyper-V role in Windows Server 2012 R2.

Role description

The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built in to Windows Server 2012 R2. Hyper-V virtualizes hardware to provide an environment in which you can run multiple operating systems at the same time on one physical computer, by running each operating system on its own virtual machine. For more information about Hyper-V, see the Hyper-V Overview.

New and changed functionality

The following table lists functionality in Hyper-V that is new for this release or has been changed.

Feature or functionality	New or updated
Shared virtual hard disk	New
Resize virtual hard disk	Updated
Storage Quality of Service	New
Live migrations	Updated
Virtual machine generation	New
Integration services	Updated
Export	Updated
Failover Clustering and Hyper-V	Updated
Enhanced session mode	New
Hyper-V Replica	Updated
Linux support	Updated
Management	Updated
Automatic Virtual Machine Activation	New
Hyper-V Networking	Updated

Shared virtual hard disk

Hyper-V in Windows Server 2012 R2 enables clustering virtual machines by using shared virtual hard disk (VHDX) files.

What value does this change add?

This feature is used to build a high availability infrastructure, and it is especially important for private cloud deployments and cloud-hosted environments that manage large workloads. Shared virtual hard disks enable multiple virtual machines to access the same virtual hard disk (VHDX) file, which provides shared storage for use by Windows Failover Clustering.. The shared virtual hard disk files can be hosted on Cluster Shared Volumes (CSV) or on Server Message Block (SMB)-based Scale-Out File Server file shares.

What works differently?

This feature is new in Windows Server 2012 R2. It was not possible to cluster virtual machines by using a shared virtual hard disk in previous releases of Windows Server.

For more information, see Virtual Hard Disk Sharing Overview,

Resize virtual hard disk

Hyper-V storage has been updated to support resizing virtual hard disks while the virtual machine is running.

What value does this change add?

Resizing virtual hard disks while the virtual machine is running enables an administrator to perform configuration and maintenance operations on the virtual hard disks while the associated virtual machine is online or the virtual hard disk data disk is in use.

What works differently?

Online virtual hard disk resizing is only available for VHDX files that are attached to a SCSI controller. The virtual hard disk size can be increased or decreased through the user interface while virtual hard disk is in use.

For more information, see Online Virtual Hard Disk Resizing Overview.

Storage Quality of Service

Hyper-V in Windows Server 2012 R2 includes storage Quality of Service (QoS). Storage QoS enables you to manage storage throughput for virtual hard disks that are accessed by your virtual machines.

What value does this change add?

Storage QoS enables you to specify the maximum and minimum I/O loads in terms of I/O operations per second (IOPS) for each virtual disk in your virtual machines. Storage QoS ensures that the storage throughput of one virtual hard disk does not impact the performance of another virtual hard disk on the same host.

What works differently?

This feature is new in Windows Server 2012 R2. It was not possible to configure storage QoS parameters for your virtual hard disks in previous releases of Windows Server.

For more information, see Storage Quality of Service for Hyper-V.

Live migrations

Hyper-V live migration has been updated with the following capabilities.

Improved performance

Hyper-V live migration has been updated to allow the administrator to select the optimal performance options when moving virtual machines to a different server.

What value does this change add?

In larger scale deployments, such as private cloud deployments or cloud hosting providers, this update can reduce overhead on the network and CPU usage in addition to reducing the amount of time for a live migration. Hyper-V administrators can configure the appropriate live migration performance options based on their environment and requirements.The following live migrations options are now available.

Option	Description
TCP/IP	The memory of the virtual machine is copied to the destination server over a TCP/IP connection. This is the same method that is used in Hyper-V in Windows Server 2012.
Compression	The memory content of the virtual machine that is being migrated is compressed and then copied to the destination server over a TCP/IP connection. This is the default setting in Hyper-V in Windows Server 2012 R2.
SMB 3.0 protocol	The memory content of the virtual machine is copied to the destination server over a SMB 3.0 connection. • SMB Direct is used when the network adapters on the source and destination servers have Remote Direct Memory Access (RDMA) capabilities enabled. • SMB Multichannel automatically detects and uses multiple connections when a proper SMB Multichannel configuration is identified. For more information, see Improve

Option	Description
	Performance of a File Server with SMB Direct.

Cross-version live migrations

Hyper-V live migration has been updated to support migrating Hyper-V virtual machines in Windows Server 2012 to Hyper-V in Windows Server 2012 R2.

What value does this change add?

Upgrading to a new version of Windows Server no longer requires downtime to the virtual machines.

Hyper-V administrators can move Hyper-V virtual machines in Windows Server 2012 to Hyper-V in Windows Server 2012 R2. Moving a virtual machine to a down-level server running Hyper-V is not supported.

What works differently?

When moving a virtual machine, the specified destination server can now be a computer running Windows Server 2012 R2. This applies to a move that is initiated in Hyper-V Manager or when using the **Move-VM** Windows PowerShell cmdlet.

Virtual machine generation

Virtual machine generation determines the virtual hardware and functionality that is presented to the virtual machine.

What value does this change add?

Hyper-V in Windows Server 2012 R2 includes two supported virtual machine generations.

- **Generation 1** Provides the same virtual hardware to the virtual machine as in previous versions of Hyper-V.
- **Generation 2** Provides the following new functionality on a virtual machine:
  - Secure Boot (enabled by default)
  - Boot from a SCSI virtual hard disk
  - Boot from a SCSI virtual DVD
  - PXE boot by using a standard network adapter
  - UEFI firmware support

Note
IDE drives and legacy network adapter support has been removed.

The following guest operating systems are supported as a generation 2 virtual machine.

- Windows Server 2012
- Windows Server 2012 R2
- 64-bit versions of Windows 8
- 64-bit versions of Windows 8.1

What works differently?

When creating a new virtual machine in Hyper-V Manager or by using the **New-VM** Windows PowerShell cmdlet, you need to specify a virtual machine generation.

Note

After a virtual machine has been created, you cannot change its generation.

For more information, see Generation 2 Virtual Machine Overview.

Integration services

Hyper-V integration services are updated with a new service that allows Hyper-V administrators to copy files to the virtual machine while the virtual machine is running without using a network connection.

What value does this change add?

In previous versions of Hyper-V, a Hyper-V administrator may have needed to shut down a virtual machine to copy files to it. A new Hyper-V integration service has been added that allows the Hyper-V administrator to copy files to a running virtual machine without using a network connection.

What works differently?

A Windows PowerShell cmdlet, **Copy-VMFile**, also has been added for this new feature. The following services must be enabled for this feature to work.

- **Guest services** on the **Integration Services** property page of the virtual machine needs to be selected. By default this setting is not selected.

  Or you can enable the **Guest services** by using the **Enable-VMIntegrationService** Windows PowerShell cmdlet.
- The **Hyper-V Guest Service Interface** service in the guest operating system must be running.

Note

The **Hyper-V Guest Service Interface** service enters a running state when the **Guest services** service is selected on the **Integration Services** property page of the virtual machine. To disable this feature in the guest operating system, the guest operating system administrator can set the **Hyper-V Guest Service Interface** service startup type to **Disabled**.

Export

Hyper-V is updated to support exporting a virtual machine or a virtual machine checkpoint while the virtual machine is running. You no longer need to shut down a virtual machine before exporting.

What value does this change add?

Exporting a virtual machine while the virtual machine is running allows the administrator to export the virtual machine without incurring any downtime.

This assists in the following scenarios:

- Duplicating an existing production environment or part of an environment to a test lab.
- Testing a planned move to a cloud hosting provider or to a private cloud.
- Troubleshooting an application issue.

What works differently?

The **Export** option is now available as an action for a running virtual machine from Hyper-V Manager. The following Windows PowerShell cmdlets can be used on a running virtual machine, **Export-VM** and **Export-VMSnapshot**.

Failover Clustering and Hyper-V

Using Windows Failover Clustering with Hyper-V enables virtual network adapter protection and virtual machine storage protection.

What value does this change add?

Hyper-V has been enhanced to detect physical storage failures on storage devices that are not managed by Windows Failover Clustering (SMB 3.0 file shares). Storage failure detection can detect the failure of a virtual machine boot disk or any additional data disks associated with the virtual machine. If such an event occurs, Windows Failover Clustering ensures that the virtual machine is relocated and restarted on another node in the cluster. This eliminates situations where unmanaged storage failures would not be detected and where virtual machine resources may become unavailable.

Hyper-V and Windows Failover Clustering are enhanced to detect network connectivity issues for virtual machines. If the physical network assigned to the virtual machine suffers a failure (such as a faulty switch port or network adapter, or a disconnected network cable), the Windows Failover Cluster will move the virtual machine to another node in the cluster to restore network connectivity.

Enhanced session mode

Virtual Machine Connection in Hyper-V now allows redirection of local resources in a Virtual Machine Connection session.

What value does this change add?

Virtual Machine Connection enhances the interactive session experience provided for Hyper-V administrators who want to connect to their virtual machines. It provides functionality that is similar to a remote desktop connection when you are interacting with a virtual machine.

In previous versions of Hyper-V, Virtual Machine Connection provided redirection of only the virtual machine screen, keyboard, and mouse with limited copy functionality. To get additional redirection abilities, a remote desktop connection to the virtual machine could be initiated, but this required a network path to the virtual machine.

The following local resources can be redirected when using Virtual Machine Connection.

- Display configuration
- Audio
- Printers
- Clipboard
- Smart cards
- Drives
- USB devices
- Supported Plug and Play devices

What works differently?

This feature is enabled by default in Client Hyper-V, and it is disabled by default on Hyper-V in Windows Server.

The following guest operating systems support enhanced session mode connections:

- Windows Server 2012 R2
- Windows 8.1

For additional information, see Virtual Machine Connection - Enhanced Session Mode Overview.

Hyper-V Replica

Hyper-V Replica adds the following new features in Windows Server 2012 R2:

- You can configure extended replication. In extended replication, your Replica server forwards information about changes that occur on the primary virtual machines to a third server (the extended Replica server). After a planned or unplanned failover from the primary server to the Replica server, the extended Replica server provides further business continuity protection. As with ordinary replication, you configure extended replication by using Hyper-V Manager, Windows PowerShell, or WMI.
- The frequency of replication, which previously was a fixed value, is now configurable. You can also access recovery points for 24 hours. Previous versions had access to recovery points for only 15 hours.

Linux support

As part of Microsoft's continuing commitment to making Hyper-V the best all-around virtual platform for hosting providers, there are now more built-in Linux Integration Services for newer distributions and more Hyper-V features are supported for Linux virtual machines.

What value does this change add?

Linux support for Hyper-V in Windows Server 2012 R2 has now been enhanced in the following ways:

- **Improved video** - a Hyper-V-specific video driver is now included for Linux virtual machines to provide an enhanced video experience with better mouse support.

- **Dynamic Memory** - Dynamic Memory is now fully supported for Linux virtual machines, including both hot-add and remove functionality. This means you can now run Windows and Linux virtual machines side-by-side on the same host machine while using Dynamic Memory to ensure fair allocation of memory resources to each virtual machine on the host.
- **Online VHDX resize** - virtual hard disks attached to Linux virtual machines can be resized while the virtual machine is running.
- **Online backup** - you can now back up running Linux virtual machines to Windows Azure using the Windows Azure Online Backup capabilities of the in-box Windows Server Backup utility, System Center Data Protection Manager, or any third-party backup solution that supports backing up Hyper-V virtual machines.

What works differently?

The Linux Integration Services are built into many distributions now, so you do not have to download and install LIS separately. For more information, see: Linux Virtual Machines on Hyper-V.

Management

You can manage Hyper-V in Windows Server 2012 from a computer running Windows Server 2012 R2 or Windows 8.1. In previous releases, you could not connect to and manage a down-level version of Hyper-V. A solution was to create a remote desktop session to a down-level server running Hyper-V and run the Hyper-V management operating system from within the remote desktop session. This solution required that Remote Desktop Services was running and properly configured, and the solution was not viable when Hyper-V was installed on Server Core installation.

What value does this change add?

You can manage Hyper-V in Windows Server 2012 from Hyper-V Manager in Windows Server 2012 R2 or Windows 8.1. This enables you to upgrade your management workstation to the latest version of the operating system and to connect and manage Hyper-V in Windows Server 2012.

You can deploy the latest version of Hyper-V without upgrading the management workstation immediately.

Note
When connecting to Hyper-V in Windows Server 2012 R2 from a computer running Windows Server 2012 or Windows 8, you can only perform actions that are supported by Hyper-V in Windows Server 2012.

Automatic Virtual Machine Activation

Automatic Virtual Machine Activation (AVMA) lets you install virtual machines on a computer where Windows Server 2012 R2 is properly activated without having to manage product keys for each individual virtual machine, even in disconnected environments. AVMA binds the virtual machine activation to the licensed virtualization server and activates the virtual machine when it

starts. AVMA also provides real-time reporting on usage, and historical data on the license state of the virtual machine. Reporting and tracking data is available on the virtualization server.

What value does this change add?

AVMA requires a virtualization server running Windows Server 2012 R2 Datacenter. The operating system on the guest virtual machine must be Windows Server 2012 R2 Datacenter, Windows Server 2012 R2 Standard, or Windows Server 2012 R2 Essentials.

Datacenter managers can use AVMA to do the following:

- Activate virtual machines in remote locations
- Activate virtual machines with or without an Internet connection
- Track virtual machine usage and licenses from the virtualization server, without requiring any access rights on the virtual machines

What works differently?

There are no product keys to manage and no stickers to read on the servers. The virtual machine is activated and continues to work even when it is migrated across an array of virtualization servers.

Service Provider License Agreement (SPLA) partners and other hosting providers do not have to share product keys with tenants or access a tenant's virtual machine to activate it. Virtual machine activation is transparent to the tenant when AVMA is used. Hosting providers can use the server logs to verify license compliance and to track client usage history.

For more information, see Automatic Virtual Machine Activation.

Hyper-V Networking

- What's New in Hyper-V Network Virtualization in Windows Server 2012 R2
- What's New in Hyper-V Virtual Switch in Windows Server 2012 R2

What's New in Hyper-V for Windows Server 2012

This topic describes the new and changed functionality of the Hyper-V role in Windows Server® 2012.

If you are looking for information about Windows Server® 2012 R2, see What's New in Hyper-V for Windows Server 2012 R2.

Role description

The Hyper-V role enables you to create and manage a virtualized computing environment by using virtualization technology that is built in to Windows Server 2012. Hyper-V virtualizes hardware to provide an environment in which you can run multiple operating systems at the same

time on one physical computer, by running each operating system in its own virtual machine. For more information about Hyper-V, see the Hyper-V Overview.

New and changed functionality

The following table lists functionality in Hyper-V that is new for this release or has been changed.

Feature/functionality	New or Updated
Client Hyper-V	New
Dynamic Memory	Updated
Hyper-V module for Windows PowerShell	New
Hyper-V Replica	New
Importing of virtual machines	Updated
Live migration	Updated
Resource metering	New
Significantly increased scale and improved resiliency	Updated
Simplified authorization	New
SR-IOV	New
Storage migration	New
Storage on SMB 3.0 file shares	New
Virtual Fibre Channel	New
Virtual hard disk format	Updated
Virtual machine snapshots	Updated
Virtual NUMA	New
Virtual switch	Updated
Sleep support	New

Client Hyper-V

The Hyper-V technology is now available in a desktop operating system version of Windows.

What value does this change add?

This change enables you to use Hyper-V without installing a server operating system. For more information, see **Client Hyper-V**.

Dynamic Memory

Dynamic Memory improvements include support for configuring minimum memory, and Smart Paging, which is a memory management technique to provide a reliable restart experience for virtual machines configured with less minimum memory than startup memory.

What value does this change add?

These changes can help you improve your hardware utilization by making it more practical to run more virtual machines on a single physical computer at the same time. Memory resources can be allocated to virtual machines more efficiently, because you can configure a virtual machine with less minimum memory than startup memory without having to plan for enough memory to be available to restart the virtual machine.

What works differently?

You can now configure a virtual machine with minimum memory, using either Hyper-V Manager or the Set-VMMemory cmdlet. Smart Paging reduces the risk that a virtual machine might not be able to restart if a restart is required when there is not enough available physical memory. Smart Paging uses disk resources as additional, temporary memory when more memory is required to restart a virtual machine than the amount of memory currently allocated to a virtual machine. To minimize the potential performance impact of Smart Paging, Hyper-V uses it only when all of the following occurs:

1. The virtual machine is being restarted.
2. There is no available physical memory.
3. No memory can be reclaimed from other virtual machines running on the host.

For more information about improvements to Dynamic Memory, see the Hyper-V Dynamic Memory Overview.

Hyper-V module for Windows PowerShell

Hyper-V module for Windows PowerShell includes more than 160 cmdlets to manage Hyper-V, virtual machines, and virtual hard disks.

What value does this change add?

The availability of built-in Hyper-V cmdlets makes it easier to automate various operations than was possible in previous versions of Hyper-V. The cmdlet names are consistent with other Windows PowerShell cmdlets (using verbs such as ‘Get’ and ‘Set’) which makes it easier for users familiar with Windows PowerShell to learn to use new cmdlets.

What works differently?

Hyper-V cmdlets are available after being installed. The cmdlets can be installed when you install the role, or as a feature available through the Remote Server Administration Tools. For more information, see Hyper-V Automation Overview.

Hyper-V Replica

This feature enables you to replicate virtual machines between storage systems, clusters, and data centers in two sites to provide business continuity and disaster recovery.

What value does this change add?

This feature makes it easier and less expensive to plan and implement business continuity and disaster recovery solutions for your virtual machines.

What works differently?

You can use a built-in solution to configure virtual machines to be replicated to a second server running Hyper-V. For more information, see **Maintaining Business Continuity of Virtualized Environments with Hyper-V Replica: scenario overview**.

Importing of virtual machines

The import process has been updated to help you resolve configuration problems that would otherwise prevent a virtual machine from being imported. You also can import a virtual machine after copying the files manually, rather than exporting the virtual machine first.

What value does this change add?

The import process has been simplified and is more reliable.

What works differently?

Before this update was introduced, the import process copied a virtual machine and its files but did not check for configuration problems. Now, the import process has been redesigned to detect and help you correct certain configuration problems.

Live migration

It is now possible to perform a live migration in a non-clustered environment, as well as perform more than one live migration at the same time and use higher networks bandwidths.

What value does this change add?

These improvements provide more flexibility in moving a running virtual machine by making it easier and faster to perform live migrations.

What works differently?

You can now perform a live migration without setting up failover clustering and Cluster Shared Volumes. Live migration in a non-clustered environment is supported both when storage is kept on a central SMB share, and when the storage is local. You can complete migrations faster by using higher network bandwidths (up to 10 gigabits). You can also perform more than one live migration at the same time to quickly move many virtual machines. For more information, see Virtual Machine Live Migration Overview. For instructions on performing a live migration outside a cluster, see Configure Live Migration and Migrating Virtual Machines without Failover Clustering.

Resource metering

You can now track and gather data about physical processor, memory, storage, and network usage by specific virtual machines.

What value does this change add?

You can determine the resource usage of specific virtual machines. You can use this data to perform capacity planning, to monitor consumption by different business units or customers, or to capture data needed to help redistribute the costs of running a workload.

What works differently?

You can configure a virtual machine to measure the physical processor, memory, storage, and network usage, and then measure the usage periodically. For more information, see Hyper-V Resource Metering Overview.

Significantly increased scale and improved resiliency

This version of Hyper-V offers significantly larger compute and storage resources than was previously possible and improves the handling of hardware errors.

What value does this change add?

These features help ensure that you can configure large, high-performance virtual machines to support workloads that might need to scale up significantly.

What works differently?

The computer that runs Hyper-V in Windows Server 2012 can be configured with up to 320 logical processors and 4 terabytes (TB) of memory. Virtual machines can be configured with 64 virtual processors and 1 TB of memory. Improved handling of hardware errors increases the resiliency and stability of the virtualization environment. For more information, see Hyper-V Scalability in Windows Server 2012and Hyper-V Support for Scaling Up and Scaling Out Overview.

Simplified authorization

The Hyper-V Administrators group is introduced and is implemented as a local security group.

What value does this change add?

This group can reduce the number of users that belong to the local Administrators group while providing users with access to Hyper-V.

What works differently?

The Hyper-V Administrators group is a new local security group. Add users to this group instead of the local Administrators group to provide them with access to Hyper-V. Members of the Hyper-V Administrators have complete and unrestricted access to all features of Hyper-V.

SR-IOV

This new feature allows you to assign a network adapter that supports single-root I/O virtualization (SR-IOV) directly to a virtual machine.

What value does this change add?

Use of SR-IOV maximizes network throughput while minimizing network latency as well as the CPU overhead required for processing network traffic.

What works differently?

You can provide a virtual machine with direct connectivity to a physical network adapter. For more information, see Hyper-V Support for Scaling Up and Scaling Out Overview.

Storage migration

Move the virtual hard disks used by a virtual machine to different physical storage while the virtual machine remains running.

What value does this change add?

This new feature makes it easier to manage the physical storage subsystems used by virtual machines.

What works differently?

It is no longer necessary to take a virtual machine offline to move, maintain, or upgrade the underlying physical storage.

Storage on SMB 3.0 file shares

Hyper-V now supports use of SMB 3.0 file shares to provide storage for virtual machines.

What value does this change add?

This new feature provides a new way to provide virtual machines with shared storage, without use of a storage area network (SAN).

What works differently?

You can specify a file share as the storage location for virtual machines. For more information, see Deploy Hyper-V over SMB.

Virtual Fibre Channel

This new feature allows you to connect directly to Fibre Channel storage from within the guest operating system that runs in a virtual machine.

What value does this change add?

This feature makes it possible to virtualize workloads and application that require direct access to Fibre Channel-based storage. It also makes it possible to configure clustering directly within the guest operating system (sometimes referred to as guest clustering).

What works differently?

This feature makes host bus adapter (HBA) ports available within the guest operating system. For more information, see Hyper-V Virtual Fibre Channel Overview.

Virtual hard disk format

The new format has been introduced to meet evolving requirements and take advantage of innovations in storage hardware.

What value does this change add?

The new format increases the maximum storage size per virtual hard disk, as well as improves the stability and efficiency of those disks.

What works differently?

A new virtual hard disk format (VHDX) supports up to 64 terabytes of storage. It also provides built-in protection from corruption stemming from power failures and prevents performance degradation on some large-sector physical disks. For more information, see Hyper-V Virtual Hard Disk Format Overview.

Virtual machine snapshots

After a virtual machine snapshot is deleted, the storage space that the snapshot consumed before being deleted is now made available while the virtual machine is running.

What value does this change add?

It is no longer necessary to shut down, turn off, or put the virtual machine into a saved state to recover the storage space.

What works differently?

Operations that keep the virtual machine and any remaining snapshots synchronized after deleting a snapshot now occur while the virtual machine remains running. After those operations are complete, the files that contained the snapshots (.avhd files) are deleted. Previously, those operations did not occur while the virtual machine was running or paused.

Virtual NUMA

A virtual NUMA topology is made available to the guest operating system in a virtual machine.

What value does this change add?

Newer operating systems and high-performance applications such as SQL Server include optimizations that recognize a computer's NUMA topology to increase performance by considering NUMA when scheduling threads or allocating memory. The virtual NUMA feature makes it possible for the guest operating system and NUMA-aware applications running in the virtual machine (such as SQL Server) to take advantage of these performance optimizations.

What works differently?

A virtual NUMA topology is made available to the guest operating system. The default virtual NUMA topology is optimized to match the NUMA topology of physical computer.

Virtual switch

The architecture of the virtual switch has been updated to provide an open framework that allows third parties to add new functionality to the virtual switch.

What value does this change add?

This change enables Microsoft partners to extend the virtual switch by providing new functionality such as monitoring, forwarding, and filtering into the virtual switch.

What works differently?

The open architecture makes it possible to use extensions, which are implemented by using Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers. These two public Windows platforms for extending Windows networking functionality are used as follows:

- NDIS filter drivers are used to monitor or modify network packets in Windows.
- WFP callout drivers allow independent software vendors (ISVs) to create drivers to filter and modify TCP/IP packets, monitor or authorize connections, filter IP security (IPsec)-protected traffic, and filter remote procedure calls (RPCs).

For more information, see the Hyper-V Virtual Switch Overview.

Sleep support

Sleep and hibernate are no longer disabled when the Hyper-V role is installed.

What value does this change add?

Computers running Hyper-V can be put into sleep or hibernate mode. This is particularly useful to users who are running Hyper-V on laptops computers.

What works differently?

When installing the Hyper-V role sleep and hibernate are no longer disabled. When entering sleep or hibernate Hyper-V will save running virtual machines. When coming out of sleep or hibernate Hyper-V will attempt to restart the virtual machines that were running before entering sleep or hibernate.

Sleep support

Sleep and hibernate are no longer disabled when the Hyper-V role is installed.

What value does this change add?

Computers running Hyper-V can be put into sleep or hibernate mode. This is particularly useful to users who are running Hyper-V on laptops computers.

What works differently?

When installing the Hyper-V role sleep and hibernate are no longer disabled. When entering sleep or hibernate Hyper-V will save running virtual machines. When coming out of sleep or hibernate Hyper-V will attempt to restart the virtual machines that were running before entering sleep or hibernate.

Removed or deprecated functionality

The following features and functionalities have been removed from this release of Hyper-V. Applications, code, or usage that depends on these features will not function in this release unless you employ an alternate method. For more information about removed or deprecated functionality in this release, see Features Removed or Deprecated in Windows Server 2012.

- Virtual machine (VM) Chimney, also called TCP Offload, has been removed. The TCP chimney will not be available to guest operating systems.

The following features and functionalities are deprecated starting with this release of Hyper-V. Eventually, they will be completely removed from the product. We recommend that you begin planning now to employ alternate methods for any applications, code, or usage that depend on these features.

- The WMI root\virtualization namespace is deprecated. The new namespace is root\virtualization\v2.
- Windows Authorization Manager (AzMan) is deprecated. You may need to use new management tools for virtual machines or redesign the authorization model you use.

See also

The following table provides additional resources for evaluating Hyper-V.

Content type	References
Product evaluation	Install Hyper-V and Configure a Virtual Machine | Understand and Troubleshoot Hyper-V Replica in Windows Server "8" | Building Your Cloud Infrastructure: Scenario Overview | Increasing Server, Storage, and Network Availability: Scenario Overview
Community resources	Ben Armstrong's Virtual PC Guy's WebLog | Hyper-V Portal on the TechNet Wiki
Related technologies	Failover Clustering Overview | File and Storage Services Overview | Remote Desktop Services Overview

Hyper-V Scalability in Windows Server 2012

Hyper-V in Windows Server® 2012 supports significantly larger configurations of virtual and physical components than in previous releases of Hyper-V. This increased capacity enables you to run Hyper-V on large physical computers and to virtualize high-performance, scale-up

workloads. This topic lists the supported maximum configuration for the various components. As you plan your deployment of Hyper-V, consider the maximums that apply to each virtual machine as well as those that apply to the physical computer that runs the Hyper-V role.

Note
For information about System Center Virtual Machine Manager (VMM), see Virtual Machine Manager. VMM is a Microsoft product for managing a virtualized data center that is sold separately—it is not a component of Windows Server 2012.

Virtual machines

The following table lists the maximums that apply to each virtual machine.

Component	Maximum	Notes
Virtual processors	64	The number of virtual processors supported by a guest operating system might be lower. For more information, see the Hyper-V Overview.
Memory	1 TB	Review the requirements for the specific operating system to determine the minimum and recommended amounts.
Virtual hard disk capacity	64 TB supported by the VHDX format introduced in Windows Server 2012 and Windows® 8; 2040 GB supported by the VHD format.	Each virtual hard disk is stored on physical media as either a .vhdx or a .vhd file, depending on the format used by the virtual hard disk.
Virtual IDE disks	4	The startup disk (sometimes referred to as the boot disk) must be attached to one of the IDE devices. The startup disk can be either a virtual hard disk or a physical disk attached directly to a virtual machine.
Virtual SCSI controllers	4	Use of virtual SCSI devices requires integration services to be installed in the guest operating system. For a list of the guest operating systems for which integration services are

Component	Maximum	Notes
		available, see the Hyper-V Overview.
Virtual SCSI disks	256	Each SCSI controller supports up to 64 disks, which means that each virtual machine can be configured with as many as 256 virtual SCSI disks. (4 controllers x 64 disks per controller)
Virtual Fibre Channel adapters	4	As a best practice, we recommended that you connect each virtual Fibre Channel Adapter to a different virtual SAN.
Size of physical disks attached directly to a virtual machine	Varies	Maximum size is determined by the guest operating system.
Snapshots	50	The actual number may be lower, depending on the available storage. Each snapshot is stored as an .avhd file that consumes physical storage.
Virtual network adapters	12	• 8 can be the “network adapter” type. This type provides better performance and requires a virtual machine driver that is included in the integration services packages. • 4 can be the “legacy network adapter” type. This type emulates a specific physical network adapter and supports the Pre-execution Boot Environment (PXE) to perform network-based installation of an operating system.

Component	Maximum	Notes
Virtual floppy devices	1 virtual floppy drive	None.
Serial (COM) ports	2	None.

Server running Hyper-V

The following table lists the requirements and maximums that apply to the server running Hyper-V.

Component	Maximum	Notes
Logical processors	320	Both of the following must be available and enabled in the BIOS: • Hardware-assisted virtualization • Hardware-enforced Data Execution Prevention (DEP)
Virtual processors per logical processor	No ratio imposed by Hyper-V.	None.
Running virtual machines per server	1024	None.
Virtual processors per server	2048	None.
Memory	4 TB	None.
Storage	Limited by what is supported by the management operating system. No limits imposed by Hyper-V.	**Note** Microsoft supports network-attached storage (NAS) for Hyper-V in Windows Server 2012 when using SMB 3.0. NFS-based storage is not supported.
Virtual storage area networks (SANs)	No limits imposed by Hyper-V	None.
Physical network adapters	No limits imposed by Hyper-V.	None.
Network adapter teams (NIC Teaming)	No limits imposed by Hyper-V.	For more information about NIC Teaming in Windows Server

Component	Maximum	Notes
		2012, see NIC Teaming Overview.
Virtual switches	Varies; no limits imposed by Hyper-V.	The practical limit depends on the available computing resources.
Virtual network switch ports per server	Varies; no limits imposed by Hyper-V.	The practical limit depends on the available computing resources.

Failover Clusters and Hyper-V

The following table lists the maximums that apply to highly available servers running Hyper-V. It is important to do capacity planning to ensure that there will be enough hardware resources to run all the virtual machines in a clustered environment. For more information about requirements for failover clusters, see Failover Clustering Hardware Requirements and Storage Options.

Component	Maximum	Notes
Nodes per cluster	64	Consider the number of nodes you want to reserve for failover, as well as maintenance tasks such as applying updates. We recommend that you plan for enough resources to allow for 1 node to be reserved for failover, which means it remains idle until another node is failed over to it. (This is sometimes referred to as a passive node.) You can increase this number if you want to reserve additional nodes. There is no recommended ratio or multiplier of reserved nodes to active nodes; the only specific requirement is that the total number of nodes in a cluster cannot exceed the maximum of 64.

Component	Maximum	Notes
Running virtual machines per cluster and per node	8,000 per cluster	Several factors can affect the real number of virtual machines that can be run at the same time on one node, such as: • Amount of physical memory being used by each virtual machine. • Networking and storage bandwidth. • Number of disk spindles, which affects disk I/O performance.

Hyper-V Component Architecture Posters

Architecture posters provide a visual reference for understanding key Hyper-V technologies.

- Download the **Windows Server 2012 R2 Hyper-V Component Architecture Poster and Hyper-V Mini-Posters** here: http://www.microsoft.com/download/details.aspx?id=40732.
- Download the **Windows Server 2012 R2 Private Cloud Virtualization and Storage Poster and Mini-Posters** here: http://www.microsoft.com/download/details.aspx?id=41665
- Download the **Windows Server 2012 Hyper-V Component Architecture** poster as a .pdf file here: http://go.microsoft.com/fwlink/p/?LinkId=242689.

Tip

The mini-posters were designed to be printed on a standard printer.

Other posters in this series include the previously published:

- Windows Server 2008 R2: Hyper-V Component Architecture - http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=3501
- Windows Server 2008 Component Posters - http://go.microsoft.com/fwlink/?LinkId=193499
- Windows Server 2008 R2 Feature Components Poster - http://go.microsoft.com/fwlink/?LinkId=179116

Install Hyper-V and Configure a Virtual Machine

This content shows you how to install the Hyper-V role on a computer running Windows Server® 2012. It also provides basic instructions for creating and configuring a virtual machine after Hyper-V is installed. It doesn't provide configuration instructions for specific usage scenarios, such as configuring Hyper-V to use scale-out file servers. Instead, it provides links to other documents that provide those instructions, when available.

In this document

- Prerequisites
- Step 1: Install Hyper-V
- Step 2: Create a virtual machine
- Step 3: Install the guest operating system
- Step 4: Install or upgrade integration services

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

Before you attempt to install Hyper-V, make sure that you have the following:

- A user account with the appropriate administrator to the computer.
- Enough memory to run all the virtual machines that you plan to run at the same time.
- Software for the virtual machine. For example, to test a particular workload, you will need installation media for the operating system and the workload. If you want to test certain features of Windows Server 2012, you can use the same installation media in the virtual machine that you used to install Windows Server 2012 on the physical computer.

Step 1: Install Hyper-V

Install Hyper-V so you can create and run virtual machines on this computer. Select one of the installation methods:

- To add the Hyper-V role in Windows Server
- To enable Client Hyper-V
- To install the Hyper-V role using the Install-WindowsFeature cmdlet on Windows Server
- To install Client Hyper-V using the Get-WindowsOptionalFeature cmdlet

To add the Hyper-V role in Windows Server

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following:
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** page, type any part of the name **Server Manager**. Click the shortcut for Server Manager when it is displayed on the **Start** page in the **Apps** results. To pin the Server Manager shortcut to the **Start** page, right-click the shortcut, and then click **Pin to Start**.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page, verify that your destination server and network environment are prepared for the role and feature you want to install. Click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation** and then click **Next**.
5. On the **Select destination server** page, select a server from the server pool and then click **Next**.
6. On the **Select server roles** page, select **Hyper-V**.
7. To add the tools that you use to create and manage virtual machines, click **Add Features**. On the Features page, click **Next**.
8. On the **Create Virtual Switches** page, **Virtual Machine Migration** page, and **Default Stores** page, select the appropriate options.
9. On the **Confirm installation selections** page, select **Restart the destination server automatically if required**, and then click **Install**.
10. When installation is finished, verify the installation by opening the **All Servers** page in Server Manager, selecting a server on which you installed Hyper-V, and viewing the **Roles and Features** tile on the page for the selected server.

To enable Client Hyper-V

1.

   Important
   Because Hyper-V is an optional feature, the files required for installation may or may not be present on your computer, depending on your organization's IT policy. If you are connected to the Internet, enabling the feature will automatically download the required files. If you are not connected to the Internet, you can download the required files and manually copy them to your computer, or otherwise provide the installation media.

In the Control Panel, click **Programs**, and then click **Programs and Features**.

2. Click **Turn Windows features on or off**.
3. Click **Hyper-V**, click **OK**, and then click **Close**.

To install the Hyper-V role using the Install-WindowsFeature cmdlet on Windows Server

1.

   Warning

   In Windows PowerShell, unlike in the Add Roles and Features Wizard, management tools and snap-ins for a role are not included by default. To include management tools as part of a role installation, add the `-IncludeManagementTools` parameter to the cmdlet. If you are installing roles and features on a server that is running the Server Core installation option of Windows Server 2012, and you add a role's management tools to an installation, you are prompted to change the installation option to a minimal-shell option that allows the management tools to run. Otherwise, management tools and snap-ins cannot be installed on servers that are running the Server Core installation option of Windows Server.

   For more information, see Install-WindowsFeature.

   Do one of the following to open a Windows PowerShell session with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** page, type any part of the name **Windows PowerShell**. Right-click the shortcut for Windows PowerShell when it is displayed on the **Start** page in the **Apps** results, click **Advanced**, and then click **Run as Administrator**. To pin the Windows PowerShell shortcut to the **Start** page, right-click the shortcut, and then click **Pin to Start**.

2. Type the following, and then press **Enter**, where *computer_name* represents a remote computer on which you want to install Hyper-V. To install Hyper-V directly from a console session, do not include `-ComputerName <computer_name>` in the command.

   ```
   Install-WindowsFeature –Name Hyper-V -ComputerName
   <computer_name> -IncludeManagementTools -Restart
   ```

3. To view a list of available and installed roles and features on the local server, type **Get-WindowsFeature** and then press **Enter**. The results of the cmdlet contain the command names of roles and features that have been added to this computer.

   Note

   In Windows PowerShell 3.0, there is no need to import the Server Manager cmdlet module into the Windows PowerShell session before running cmdlets that are part of the module. A module is automatically imported the first time you run

a cmdlet that is part of the module. Also, Windows PowerShell cmdlets are not case-sensitive.

4. When the installation is finished, verify installation by running the `Get-WindowsFeature`. If you installed Hyper-V remotely, include the ComputerName parameter (Get-WindowsFeature -ComputerName <*computer_name*>) to view a list of roles and features that are installed on the server.

To install Client Hyper-V using the Get-WindowsOptionalFeature cmdlet

1.

   Note
   For more information, see Get-WindowsOptionalFeature and Using PowerShell to Set Up Hyper-V

   Do one of the following to open a Windows PowerShell session with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** page, type any part of the name **Windows PowerShell**. Right-click the shortcut for Windows PowerShell when it is displayed on the **Start** page in the **Apps** results, click **Advanced**, and then click **Run as Administrator**. To pin the Windows PowerShell shortcut to the **Start** page, right-click the shortcut, and then click **Pin to Start**.

2. Type the following, and then press **Enter**.

   ```
   Get-WindowsOptionalFeature -Online -FeatureName *hyper-v*all
   | Enable-WindowsOptionalFeature -Online
   ```

3. When the installation is finished, verify installation by running the `Get-WindowsOptionalFeature -Online` and then press **Enter**. The results of the cmdlet contain the command names features that have been added to this computer.

Step 2: Create a virtual machine

Before beginning, review the following considerations. These are not required, but can help you make the most appropriate choices. Consider:

- What to name the virtual machine. The name provides you with a way to identify the virtual machine. For example, you might use a name that identifies the type of server you want to configure on this virtual machine, such as a web server. Or, you might use a name that identifies the guest operating system.
- How you will install the operating system. You can install from a physical media or an image (.iso) file. You also can use a virtual hard disk on which an operating system has already been installed.

Do this step using Windows PowerShell

To create a virtual machine

1. Open Hyper-V Manager.
2. From the navigation pane of Hyper-V Manager, select the computer running Hyper-V.
3. From the Actions pane, click **New** and then click **Virtual Machine**.
4. The New Virtual Machine wizard opens. Click **Next**.
5. On the **Specify Name and Location** page, type an appropriate name.
6. On the **Assign Memory** page, specify enough memory to start the guest operating system.
7. On the **Configure Networking** page, connect the virtual machine to the switch you created when you installed Hyper-V.
8. On the **Connect Virtual Hard Disk** and **Installation Options** pages, choose the option that is appropriate for how you plan to install the guest operating system:
   - If you will install the guest operating system from a DVD or an image file (an .ISO file), choose **Create a virtual hard disk**. Click **Next**, and then click the option that describes the type of media you will use. For example, to use an .iso file, click **Install an operating system from a boot CD/DVD** and then specify the path to the .iso file.
   - If the guest operating system is already installed in a virtual hard disk, choose **Use an existing virtual hard disk** and click **Next**. Then, choose **Install an operating system later**.
9. On the Summary page, verify your selections and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Run the following command to create a virtual machine named web server with 1 GB of startup memory and use an existing virtual hard disk in which a guest operating system has already been installed.

```
New-VM –Name “web server” –MemoryStartupBytes 1GB –VHDPath d:\vhd\BaseImage.vhdx
```

Step 3: Install the guest operating system

This step assumes that you configured the boot media for the virtual machine when you created the virtual machine.

Note
This step must be completed through the GUI. It cannot be automated or performed within a Windows PowerShell session.

To install the guest operating system

1. From Hyper-V Manager, in the **Virtual Machines** section of the results pane, right-click the name of the virtual machine and click **Connect**.
2. The Virtual Machine Connection tool opens.
3. From the **Action** menu in the Virtual Machine Connection window, click **Start**.
4. The virtual machine starts, searches the startup devices, and loads the installation package.
5. Proceed through the installation.

Step 4: Install or upgrade integration services

Hyper-V includes a software package for supported guest operating systems that improves integration between the physical computer and the virtual machine. This package is referred to as integration services.

Note
This step must be completed through the GUI. It cannot be automated or performed within a Windows PowerShell session.

To install or upgrade integration services

1. Open Hyper-V Manager. From the Server Manager Tools menu, click **Hyper-V Manager**.
2. Connect to the virtual machine. From the Virtual Machines section of the results pane, using one of the following methods:
   - Right-click the name of the virtual machine and click **Connect**.
   - Select the name of the virtual machine. In the **Action** pane, click **Connect**.
3. The Virtual Machine Connection tool opens. From the Action menu of Virtual Machine Connection, click **Insert Integration Services Setup Disk**. This action loads the setup disk in the virtual DVD drive.
4. Depending on the operating system being installed, you may need to start the installation manually. Click anywhere in the guest operating system window and navigate to the CD drive. Use the method that is appropriate for the guest operating system to start the installation package from the CD drive.
5. After the installation finishes, all integration services are available for use.

See also

- Hyper-V Overview
- What's New in Windows Server
- Hyper-V Module for Windows PowerShell
- **Client Hyper-V**
- Linux Virtual Machines on Hyper-V

Hyper-V Replica Overview

This topic describes how to deploy Hyper-V Replica, a new feature of Hyper-V that replicates all changes on a virtual machine to a counterpart virtual machine hosted by a different server. It covers basic planning decisions and includes steps for preparing a virtualized environment for Replica, configuring and enabling replication, testing the deployment, performing planned failovers, and responding to unplanned failovers.

Did you mean…

- Increasing Server, Storage, and Network Availability: Scenario Overview
- Building Your Cloud Infrastructure: Scenario Overview
- Windows Azure Backup Overview
- Hyper-V Overview
- Failover Clustering Overview

Role/Feature description

Hyper-V Replica provides asynchronous replication of Hyper-V virtual machines between two hosting servers. It is simple to configure and does not require either shared storage or any particular storage hardware. Any server workload that can be virtualized in Hyper-V can be replicated. Replication works over any ordinary IP-based network, and the replicated data can be encrypted during transmission. Hyper-V Replica works with standalone servers, failover clusters, or a mixture of both. The servers can be physically co-located or widely separated geographically. The physical servers do not need to be in the same domain, or even joined to any domain at all.

Key concepts

In this scenario, we define two “sites”: the “primary site,” which is the location where the virtualized environment normally operates; and the “Replica site,” which is the location of the server that will receive the replicated data. At the primary site, the primary server is the physical server that hosts one or more primary virtual machines. At the Replica site, the Replica server similarly hosts the Replica virtual machines.

Once replication is configured and enabled, an initial copy of data from the primary virtual machines must be sent to the Replica virtual machines. We call this “initial replication” and you can choose to accomplish it directly over the network or by copying the data to a physical device and transporting that to the Replica site.

When replication is underway, changes in the primary virtual machines are transmitted over the network periodically to the Replica virtual machines. The exact frequency varies depending on how long a replication cycle takes to finish (depending in turn on the network throughput, among other things), but generally, replication data is sent to the Replica server every 5 minutes in Windows Server 2012. In Windows Server 2012 R2, you can configure the replication frequency, so that the changes are sent every 30 seconds, every 5 minutes, or every 15 minutes.

You can choose to move operations on any primary virtual machine to its corresponding Replica virtual machine at any time, an action we call “planned failover.” In a planned failover, any un-replicated changes are first copied over to the Replica virtual machine and the primary virtual machine is shut down, so no loss of data occurs. After the planned failover, the Replica virtual machine takes over the workload; to provide similar protection for the virtual machine that is now servicing the workload, you configure “reverse replication” to send changes back to the primary virtual machine (once that comes back online).

If the primary server should fail unexpectedly, perhaps as a result of a major hardware failure or a natural disaster, you can bring up the Replica virtual machines to take over the workload—this is “unplanned failover.” In unplanned failover, there is the possibility of data loss, since there was no opportunity to copy over changes that might not have been replicated yet.

New features starting with Windows Server 2012 R2

Starting in Windows Server 2012 R2, the frequency of replication, which previously was a fixed value, is configurable. You can set this value in the **Configure Replication Frequency** page of the Enable Replication wizard or in the **Replication** section for the virtual machine in Hyper-V Manager. You can configure the replication frequency so that the changes are sent every 30 seconds, every 5 minutes, or every 15 minutes.

You can also access recovery points up to 24 hours old (previously, recovery points up to 15 hours old were available).

You can configure extended replication. In extended replication, your Replica server forwards changes that occur on the primary virtual machines to a third server (the extended Replica server). After a planned or unplanned failover from the primary server to the Replica server, the extended Replica server provides further business continuity protection. As with ordinary replication, you configure extended replication by using Hyper-V Manager, Windows PowerShell (using the –Extended option), or WMI.

In practice, you configure the extended Replica server with a subset of the usual protection parameters. The extended Replica server offers most of the same features as the Replica server, though it does not support application-consistent replication and must use the same VHDs that the Replica server is using. You can fail over to the extended Replica server if both the primary and Replica servers go down; you can also perform a test failover to either the Replica or extended Replica server without disrupting workloads. You can move the target of replication from the extended Replica server to a different extended Replica server.

Practical applications

Because Hyper-V Replica is so simple and flexible, it can be used in a wide variety of potential scenarios of varying complexity. Some examples include:

Head office and branch office

In this scenario, there are two sites: a main head office and one or more branch offices in different physical locations. Taking advantages of virtualized workloads, Hyper-V Replica can be used to provide disaster recovery support for the branch offices. The servers at any of the sites can be clustered or standalone.

For this situation, day-to-day operations would run on the virtual machines running on primary servers at the various branch offices. Each branch office would have a Replica server standing by at the head office to take over the workload in the event that the primary server must go offline for any reason.

In Windows Server 2012 R2, you can extend this scenario further using extended replication. Each branch office would have its Replica server located nearby, which then uses extended replication to send changes to an extended Replica server back at the head office.

This scenario can be scaled up to involve large datacenters with many servers without requiring any different management activities with respect to Hyper-V Replica.

Cloud service provider

In this scenario, the hosting provider sets up a Replica server at their datacenter which receives replication data from a number of primary servers running virtualized workloads on the premises of their various customers. The hosting provider's Replica server thereby provides disaster recovery capability for the customers who subscribe to it.

To assure security for the customers, this scenario would involve certificate-based authentication using certificates probably serviced by a separate certificate server owned by the hosting provider. In addition, the Trusted Group feature of Replica allows the hosting provider to segregate the replicated data from each customer, using separate storage locations and tagging to prevent data from various customers from being mixed.

Hardware requirements

You can set up replication of Hyper-V virtual machines as long as you have any two physical Windows Server 2012 or Windows Server 2012 R2 servers which support the Hyper-V role. The two (or three, in the case of extended replication) servers can be physically co-located or in completely separate geographical locations. Either (or both of) the primary and Replica servers can be part of a failover cluster and mixed standalone and clustered environments are supported.

Software requirements

No special software, other than Windows Server 2012 or Windows Server 2012 R2, is required. If you plan to use certificate-based authentication (required for the replicated data to be encrypted during transmission), you will need an appropriate certificate, which can either be local and self-signed, or supplied by a certificate server in your deployment.

See also

This table contains links to related resources that explore other aspects of this feature, including blogs and forums where you can find the most up-to-date information.

Content type	References
Product evaluation	Windows Server 2012 R2 and Windows Server 2012| Understand and Troubleshoot Hyper-V Replica
Planning	Understand and Troubleshoot Hyper-V Replica| Feature Overviews
Deployment	Understand and Troubleshoot Hyper-V Replica
Operations	Understand and Troubleshoot Hyper-V Replica
Troubleshooting	Understand and Troubleshoot Hyper-V Replica
Security	Understand and Troubleshoot Hyper-V Replica
Tools and settings	Hyper-V Module for Windows PowerShell
Community resources	Ben Armstrong's Virtual PC Guy's WebLog| John Howard's Hyper-V and virtualization blog| Martin McClean's ThePosterGuy blog| Hyper-V Portal on the TechNet Wiki| Virtualization Blog
Related technologies	Hyper-V Overview | Failover Clustering Overview

Demonstrate Planned Failover in Hyper-V Replica

This document shows you how to demonstrate using Hyper-V Replica to move virtual machines from one physical server to another without interrupting their operations, a means of maintaining business continuity in virtualized environments called "planned failover." It will show you how to enable Hyper-V, configure the physical servers for Hyper-V Replica, start and monitor replication, perform a test failover, and perform an actual failover.

In this document

- Prerequisites
- Step 1: Enable the Hyper-V Role and create a virtual machine
- Step 2: Configure the Replica server to accept replication

- Step 3: Enable and begin replication of virtual machines
- Step 4: Monitor replication
- Step 5: Demonstrate test failover
- Step 6: Demonstrate planned failover

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

To demonstrate or evaluate the planned failover of working virtual machines using Hyper-V Replica, ensure that you have two physical servers, each with the characteristics below. These two servers can be physically next to one another for ease of access or in completely separate geographic locations for a more realistic demonstration.

Note
If you are doing this test or demonstration with servers located in different sites, ensure that the firewall at each site is properly configured to allow communications between the servers running Hyper-V. Hyper-V Replica includes pre-configured inbound firewall rules that can be enabled depending on the authentication mechanism chosen (HTTP or HTTPS).

- Hardware that supports the Hyper-V role
- Sufficient storage to host the files used by whatever virtualized workloads you are using
- Sufficient network bandwidth between the locations of the primary and Replica sites (if not co-located)

Step 1: Enable the Hyper-V Role and create a virtual machine

To enable the Hyper-V server role and create virtual machines, see Install Hyper-V and Configure a Virtual Machine.

Step 2: Configure the Replica server to accept replication

The server that you configure as a Replica server can receive replicated virtual machine data from any number of primary servers. You can configure the Replica server to accept replication from only specific primary servers, from any primary server, or from a collection of primary servers defined by fully qualified domain names with wildcard characters.

In this example, the steps will enable the Replica server to accept replication traffic using Kerberos authentication over port 8080.

Do this step using Windows PowerShell

To configure the Replica server to accept replication

1. Ensure that an appropriate inbound firewall rule exists that will allow incoming replication traffic. As an example, you can use the following Windows PowerShell command to create a rule with the name HVRAllowReplicaTraffic over port 8080. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

   ```
   Enable-Netfirewallrule -displayname "Hyper-V Replica HTTP
   Listener (TCP-In)”
   ```

2. Start Hyper-V Manager and in the **Actions** pane for the server that will be the Replica server, click **Hyper-V Settings**.
3. In the **Hyper-V Settings** dialog box, click **Replication Configuration**.
4. In the **Details** pane, select **Enable this computer as a Replica server**.
5. In the **Authentication and ports** section, choose an authentication method and the port to be used (for this example, choose Kerberos authentication and port 8080).

   Note
   When using Kerberos authentication, the replicated data is not encrypted over the network. Only certificate-based authentication encrypts the replicated data during transmission.

6. In the **Authorization and storage** section, specify the location to store the Replica virtual machine files (or accept the default location). By default, you specify which primary servers will be allowed to send replication data to this Replica server, along with the storage location to use for each virtual machine. To specify a server, use its fully qualified domain name (FQDN) or the international equivalent (FQIDN). You can use “*” as a wildcard character in the first octet (for example, “*.contoso.com”). You can also choose not to specify allowed primary servers and instead allow any server that successfully authenticates to replicate data to this Replica server. Click **Apply** or **OK** to finish.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sequence of cmdlets will enable the server for replication using Kerberos authentication over port 8080 and storing the Replica data on the volume D:\Example. It will allow replication from any server that successfully authenticates. All steps must be completed by a user with administrative privileges.

```
Enable-Netfirewallrule -displayname "Hyper-V Replica HTTP Listener (TCP-In)”

Import-Module Hyper-V


$RecoveryPort = 8080
```

```
$ReplicaStorageLocation = "D:\Example"


Set-VMReplicationServer -ReplicationEnabled $true -AllowedAuthenticationType Integrated -
IntegratedAuthenticationPort $RecoveryPort -DefaultStorageLocation
$ReplicaStorageLocation -ReplicationAllowedFromAnyServer $true
```

Step 3: Enable and begin replication of virtual machines

You can enable replication (or not) for each virtual machine, regardless of the workload the virtual machine is running.

In this example, the steps will enable a virtual machine called "CRMVM" to replicate to the previously enabled Replica server. In this example, only the latest recovery point will be maintained and the initial replication occurs over the network.

Do this step using Windows PowerShell

To enable and begin replication of a virtual machine

1. In Hyper-V Manager, right-click the virtual machine you want to configure and point to **Enable Replication…**. The Enable Replication wizard opens.
2. On the **Specify Replica Server** page, provide the name of the Replica server which will host the virtual machine that will replicate the current virtual machine.
3. On the **Specify Connection Parameters** page, enter the port that was configured to receive replication traffic in the port on the Replica server (in this example, 8080). If the system has been able to contact the Replica server already, these values will be pre-populated.
4. Choose either Kerberos or certificate-based mutual authentication to match the selection you made on the Replica server. For this example, choose **Kerberos authentication (HTTP)**.
5. On the **Choose Replication VHDs** page, select those disks which you do *not* want to replicate for the virtual machine, and then click **Next**. Disks which contain data that is not necessary for the virtual machine to run once operations have switched to the Replica server can be excluded from replication. VHDs that contain data that is rapidly changing and not used by the Replica server after failover, such as page file disks, should be excluded from replication to conserve network bandwidth.
6. Use the **Configure Recovery History** page to specify the number of recovery points to maintain on the Replica server. If you choose not to store any additional recovery points, the Replica server maintains only the latest received replica data. These points are updated approximately every 5-15 minutes, depending on the time required for each replication operation. When you choose to maintain an additional one or more recovery points, they are created once every hour. In addition, you can also select to store

application-consistent snapshots at a specified interval. These snapshots use the Volume Shadow Copy Service (VSS) and preserve the state of applications that might be running in the virtual machine that is being replicated.

7. Use the **Choose Initial Replication Method** page to specify the method you will use to transfer the initial copy of the virtual machine data to the Replica server. You have three options:
   - Initial replication over the network, which transfers the set of selected VHDs over the network to the Replica server either immediately or at a later time that you specify.
   - Using a pre-existing restored virtual machine on the Replica server (for example, if you have restored an earlier backup of the virtual machine on the Replica server) as the initial copy. With this option, only the changes that have occurred on the production virtual machine since it was stored are sent over the network.
   - Using external media. You can save network bandwidth by copying the initial copy to external media and then physically delivering the media to the Replica site.

   For this example, choose **Send initial copy over the network**.
8. On the **Completing the Enable Replication Relationship** page, review the information in the **Summary** and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sequence of cmdlets will enable and start replication for two virtual machines (“CRMVM” and “IISVM”). All steps must be completed by a user with administrative privileges.

```
Import-Module Hyper-V


$ReplicaServer = “Recovery1.contoso.com”

$RecoveryPort = 8080

$PrimaryVM1 = “CRMVM”

$PrimaryServer = “Primary1.contoso.com”


Set-VMReplication -VMName $PrimaryVM1 -ReplicaServerName $ReplicaServer -
ReplicaServerPort $RecoveryPort -AuthenticationType Integrated -CompressionEnabled $true
-RecoveryHistory 0


Start-VMInitialReplication –VMName $PrimaryVM1
```

Step 4: Monitor replication

You can monitor replication status in Hyper-V Manager by adding the appropriate column. In the **View** menu, use the **Add/Remove Columns** option and add **Replication Health** to the list of **Available Columns** in Hyper-V Manager.You can also view replication health for a given monitoring interval. Select the virtual machine from Hyper-V Manager, right-click the replicating virtual machine, and choose **View Replication Health…** from the **Replication** option.

The health report shows the current replication state and the health of the connection. You can also view replication statistics.

Step 5: Demonstrate test failover

To ensure that the replicated virtual machines (and applications running within them) will function properly on the Replica server just as they do on the primary server, you can conduct a test failover at any time. When you conduct a test failover, a temporary virtual machine is created on the Replica server. You can test any applications in that test virtual machine without interrupting the ongoing replication. When you conclude the test, the temporary virtual machine is deleted.

Note

This test virtual machine will not by default be connected to any network. If you need to conduct tests that require a network, you can modify the settings of this test virtual machine in the same way you would modify settings of any ordinary virtual machine.

Do this step using Windows PowerShell

To demonstrate test failover

1. Access the Replica server, and in Hyper-V Manager, right-click the virtual machine you want to test failover for, point to **Replication…**, and then point to **Test Failover…**.
2. Choose a recovery point to use. This creates and starts a virtual machine with the name “<virtual machine name>-Test” (for example, “CRMVM-Test”).
3. Conduct your tests on the test virtual machine. For example, you might verify that the virtual machine starts, pauses, and stops, and that any applications in the virtual machine run properly.

   Tip

   After you have concluded your testing, discard the test virtual machine by choosing **Stop Test Failover** under the **Replication** option or simply delete the test virtual machine directly.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sequence of cmdlets will conduct a test failover of the virtual machine “CRMVM.”

```
$ReplicaVM1 = "CRMVM"


$TestReplicaVM1 = Start-VMFailover -AsTest -VMName $ReplicaVM1


Start-VM $TestReplicaVM1
```

To stop the test failover operation and discard the test virtual machine:

```
Stop-VMFailover –VMName $ReplicaVM1
```

Step 6: Demonstrate planned failover

In a planned failover, the latest changes on the primary virtual machine are replicated to its counterpart Replica virtual machine, then the Replica virtual machine starts, effectively transferring the load from the primary server to the Replica server with no loss of data. You would perform a planned failover, for example, to demonstrate and confirm readiness for a disaster recovery scenario or anytime you need to do any maintenance or upgrades which would require shutting down the physical primary server.

Do this step using Windows PowerShell

To demonstrate planned failover

1. If you have not already done so, follow the procedure in Step 2: Configure the Replica server to accept replication to configure the *primary* server to receive replicated data. You may have already used the procedure to configure the Replica server, but for a planned failover, the primary server must also be enabled for replication, so you might have to repeat Step 2 for the primary server.
2. Start Hyper-V Manager on the primary server and choose a virtual machine to fail over. Turn off the virtual machine that you want to fail over.
3. Right-click the virtual machine, point to **Replication**, and then point to **Planned Failover**.
4. Click **Fail Over** to actually transfer operations to the virtual machine on the Replica server. Failover will not occur if the prerequisites have not been met, but this should not happen after following the instructions in this demonstration.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sequence of cmdlets will conduct a planned failover of the virtual machine “CRMVM,” provided you have already configured the *primary* server (in addition to the Replica server) to accept replication traffic using the procedure in Step 2. Run these cmdlets on the primary server:

```
Stop-VM $PrimaryVM1

Start-VMFailover -VMName $PrimaryVM1 –prepare
```

Then run these cmdlets on the Replica server:

```
$ReplicaVM1 = “CRMVM”

Start-VMFailover -VMName $ReplicaVM1

Set-VMReplication -reverse -VMName $ReplicaVM1

Start-VM $ReplicaVM1
```

See also

- For an overview of this scenario, see **Maintaining Business Continuity of Virtualized Environments with Hyper-V Replica**
- For additional information about the Hyper-V, see Hyper-V Overview.

Deploy Hyper-V Replica

This topic contains information that will guide you through making basic decisions about setting up Hyper-V Replica, configuring Replica, starting replication, and dealing with both planned and unplanned failover activities.

- Step 1: Prepare to Deploy Hyper-V Replica
- Step 2: Enable Replication
- Step 3: Test the Replication Deployment
- Step 4: Perform a Planned Failover
- Step 5: Respond to an Unplanned Failover
- Step 6: Configure and Use Extended Replication

For more information, see Feature Overviews.

Step 1: Prepare to Deploy Hyper-V Replica

This topic helps you make the basic decisions required so that you can get Hyper-V Replica working for your particular situation as smoothly and quickly as possible. Because Replica is so

simple and flexible, you can use it in a wide variety of deployment scenarios. Deciding on a few key factors ahead of time will allow you to quickly move through the Enable Replication wizard.

For further assistance in making detailed planning decisions for your unique deployment (such as what replication frequency would be best, predicting the network load generated by replication, and more, download and use the Capacity Planner for Hyper-V Replica from http://www.microsoft.com/en-us/download/details.aspx?id=39057.

Task	Description
1.1. Make basic planning decisions	Make basic planning decisions.
1.2. Install the Hyper-V server role	Install the Hyper-V server role, if it is not already installed.
1.3. Configure the firewall	Configure the firewall, if the replication data will be crossing it.
1.4. Configure Hyper-V Replica Broker	Configure the Hyper-V Replica Broker, if the servers are part of a failover cluster.
1.5. Provide a certificate for encrypted data	Create a certificate for Replica to use, if the replication data should be encrypted.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Make basic planning decisions

By deciding on a few basic parameters ahead of time, you can save some time when you run the Enable Replication wizard and ensure that your replication deployment will work perfectly from the start.

To plan configuration

1. **Will both the primary and Replica servers be behind the same firewall?** Because the Replica site is likely remote from the primary site, you should configure your firewall to allow replication data through. **If the primary and Replica servers will be separated by a firewall, add Step 1.3 to your to-do list.**
2. **Will either the primary or Replica server be part of a failover cluster?** If so, Replica needs a NetBIOS name and IP address to use as a connection point to the cluster. To enable this, install Hyper-V on each node of the cluster and configure a Hyper-V Replica Broker role for the cluster. **If a failover cluster is involved in replication, add Step 1.4 to your to-do list.**
3. **Does the replication data sent between the primary and Replica servers need to be encrypted?** To provide encrypted replication data transfer, Replica uses certificate-

based authentication, so you must have or create an appropriate security certificate. For example, in any remote office or other cross-premises scenario, the data should probably be encrypted for security. However, if both the primary and Replica servers are physically co-located and are behind the same firewall, the built-in Kerberos authentication may be adequate.

> **Note**
> Kerberos authentication is only available when the primary and Replica servers are members of the same domain or in mutually trusted domains. Other scenarios require certificate-based authentication.

If the transmitted replication data must be encrypted, add Step 1.5 to your to-do list.

4. **Which VHDs need to be replicated?** VHDs that contain data that is rapidly changing and not used by the Replica server after failover, such as page file disks, should be excluded from replication to conserve network bandwidth. **Make a note of which VHDs can be excluded from replication.**
5. **How many recovery points do you need?** You can configure Replica to store only the latest received replication data; the data on the Replica server is updated according to the replication frequency you configure. You can also configure Replica to maintain an additional one or more recovery points, which are created once approximately every hour. Having additional recovery points also allows you to recover virtual machine operations to an earlier point in time when you perform a failover. A maximum of 15 (24 in Windows Server 2012 R2) such recovery points is possible. **Make a note of the number of recovery points you need.**
6. **Do you need applications to remain consistent?** The standard replication configuration maintains consistency in the state of the operating system of the replicated virtual machines after a failover, but not the state of applications that may be running in the virtual machine. To preserve the state of applications, you can choose to have Replica create additional application-consistent recovery points that use VSS on a set schedule. **Make a note whether you need application-consistent recovery points.**
7. **Decide on the initial replication method**. To start replicating data, Replica needs to transfer the current state of the virtual machines. This initial state can be transmitted directly over the existing network, either immediately or at a later time that you configure. You can also use a pre-existing restored virtual machine on the Replica server (for example, if you have restored an earlier backup of the virtual machine on the Replica server) as the initial copy. Or, you can save network bandwidth by copying the initial copy to external media and then physically delivering the media to the Replica site. **Make a note of the initial replication method you plan to use.**

   > **Note**
   > If you plan to use a pre-existing virtual machine restored from backup as the source for the initial replication, delete all previous snapshots associated with the virtual machine before you enable replication and start the initial replication.
8. **(For Windows Server 2012 R2) Will you use extended replication?** In extended replication, your Replica server forwards changes that occur on the primary virtual

machines to a third server (the extended Replica server). After a planned or unplanned failover from the primary server to the Replica server, the extended Replica server provides further business continuity protection. **Once you have completed Steps 1 through 5 to enable, configure, and test replication between the primary and Replica servers, see Step 6 for information about configuring and working with extended replication.**

1.2. Install the Hyper-V server role

Hyper-V Replica is a feature of Hyper-V, so you must enable the Hyper-V server role if you aren't already using it. Replica also requires virtual machines, so if you haven't already done so, you need to create one or more. To enable the Hyper-V server role and create virtual machines, see Install Hyper-V and Configure a Virtual Machine.

1.3. Configure the firewall

Hyper-V Replica already includes firewall rules that will allow the replication data to pass through, but you do have to enable the appropriate exceptions. Follow the steps below depending on whether you plan to use Kerberos or certificate-based authentication on standalone (non-clustered) servers. If your servers are part of a failover cluster, run the Windows PowerShell cmdlets that follow the procedures on any node in the cluster to enable the firewall rules on all nodes in the cluster.

Note
The details of this step apply only to Windows Firewall. If you are using a non-Microsoft firewall, different steps may be required.

To enable the firewall rules for Kerberos authentication

1. Open **Windows Firewall with Advance Security** and click **Inbound Rules**.
2. Right-click **Hyper-V Replica HTTP Listener (TCP-In)** and click **Enable Rule**.

To enable the firewall rules for certificate-based authentication

1. Open **Windows Firewall with Advance Security** and click **Inbound Rules**.
2. Right-click **Hyper-V Replica HTTPS Listener (TCP-In)** and click **Enable Rule**.

For servers that are part of a failover cluster, run this Windows PowerShell cmdlet on any node in the cluster if you will be using **Kerberos authentication** for Replica. The cmdlet must be run by a user with administrative privileges.

```
get-clusternode | ForEach-Object  {Invoke-command -computername $_.name -scriptblock
{Enable-Netfirewallrule -displayname "Hyper-V Replica HTTP Listener (TCP-In)"}}
```

For servers that are part of a failover cluster, run this Windows PowerShell cmdlet on any node in the cluster if you will be using **certificate-based authentication** for Replica. The cmdlet must be run by a user with administrative privileges.

```
get-clusternode | ForEach-Object  {Invoke-command -computername $_.name -scriptblock
{Enable-Netfirewallrule -displayname "Hyper-V Replica HTTPS Listener (TCP-In)"}}
```

1.4. Configure Hyper-V Replica Broker

If you plan to use a failover cluster with Hyper-V Replica and haven't already set up the cluster, see http://go.microsoft.com/fwlink/?LinkId=253006

Each node of the failover cluster that is involved in Replica must have the Hyper-V server role installed. If you haven't already done so, refer to Step 1.2 in this topic and install the Hyper-V role on each cluster node.

Finally, configure the Hyper-V Replica Broker for the cluster with the following steps.

To configure the Hyper-V Replica Broker

1. In Server Manager, open Failover Cluster Manager.
2. In the left pane, connect to the cluster, and while the cluster name is highlighted, click **Configure Role** in the **Actions** pane. The High Availability wizard opens
3. In the **Select Role** screen, select **Hyper-V Replica Broker**.
4. Complete the wizard, providing a NetBIOS name and IP address to be used as the connection point to the cluster (called a "client access point"). The Hyper-V Replica Broker is configured, resulting in a client access point name. **Make a note of the client access point name for configuring Replica later on.**
5. Verify that the Hyper-V Replica Broker role comes online successfully and can fail over between all nodes of the cluster. To do this, right-click the role, point to **Move**, and then click **Select Node**. Then, select a node, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This example sequence of cmdlets will create a Hyper-V Replica Broker names "HVR-Broker" that uses the static IP address 192.168.1.5. All steps must be completed by a user with administrative privileges.

```
$BrokerName = "HVR-Broker"

Add-ClusterServerRole -Name $BrokerName –StaticAddress 192.168.1.5

Add-ClusterResource -Name "Virtual Machine Replication Broker" -Type "Virtual Machine
Replication Broker" -Group $BrokerName

Add-ClusterResourceDependency "Virtual Machine Replication Broker" $BrokerName

Start-ClusterGroup $BrokerName
```

1.5. Provide a certificate for encrypted data

If you configure Replica to use Kerberos authentication, the data transmitted from the primary to the Replica server is not encrypted. For the data to be encrypted, you should use certificate-based authentication. You can either provide an existing X.509v3 certificate or create a self-signed certificate.

If you are already using a certificate server or public key infrastructure (PKI) system, you can use an existing certificate, provided that it meets the following criteria:

- The certificate must not be expired or revoked.
- The certificate must include both client and server authentication extensions for enhanced key usage (EKU) and an associated private key.
- The certificate must terminate at a valid root certificate in the Trusted Root Certification Authorities store on the Replica server.
- If the virtual machine is hosted by a standalone server, the subject common name (CN) must be equal to (or the subject alternative name—DNS name) should contain the FQDN of the host. If the virtual machine is hosted by a failover cluster, the subject common name (CN) must be equal to (or the subject alternative name—DNS name) should contain the FQDN of the Hyper-V Replica Broker.

For additional information about using certificates with Hyper-V Replica, see http://blogs.technet.com/b/virtualization/archive/2012/03/13/hyper-v-replica-certificate-requirements.aspx

You can validate any candidate certificate by running **certutil –store my** on both the primary and Replica servers. In both cases, the output of the command should indicate that the certificate passed the encryption test.

To create a self-signed certificate with Makecert.exe

1. **On the primary server**, copy the Makecert.exe utility locally.
2. Create a self-signed test root authority certificate by running the following command from an elevated command prompt:

   makecert -pe -n "CN=PrimaryTestRootCA" -ss root -sr LocalMachine -sky signature -r "PrimaryTestRootCA.cer"
3. Create a new certificate signed by the test root authority certificate by running the following command from an elevated command prompt, supplying the FQDN of the primary server:

   makecert -pe -n "CN=<FQDN>" -ss my -sr LocalMachine -sky exchange -eku 1.3.6.1.5.5.7.3.1,1.3.6.1.5.5.7.3.2 -in "PrimaryTestRootCA" -is root -ir LocalMachine -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 PrimaryTestCert.cer
4. **On the Replica server**, copy the Makecert.exe utility locally.
5. **Create a self-signed test root authority certificate by running the following command from an elevated command prompt:**

   makecert -pe -n "CN=ReplicaTestRootCA" -ss root -sr LocalMachine -sky signature

-r "ReplicaTestRootCA.cer"

6. Create a new certificate signed by the test root authority certificate by running the following command from an elevated command prompt, supplying the FQDN of the **Replica** server:

   makecert -pe -n "CN=<FQDN>" -ss my -sr LocalMachine -sky exchange -eku 1.3.6.1.5.5.7.3.1,1.3.6.1.5.5.7.3.2 -in "ReplicaTestRootCA" -is root -ir LocalMachine -sp "Microsoft RSA SChannel Cryptographic Provider" -sy 12 ReplicaTest.cer

7. Copy the file ReplicaTestRootCA.cer from the Replica server to the primary server, and then import it with the following command:

   certutil -addstore -f Root "ReplicaTestRootCA.cer"

8. Copy the file PrimaryTestRootCA.cer from the primary server to the Replica server, and then import it with the following command:

   certutil -addstore -f Root "PrimaryTestRootCA.cer"

9. By default, a certificate revocation check is required; however, self-signed certificates don't support revocation checks. Disable the check by editing the registry on both the primary and Replica servers with the following command:

   reg add "HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Virtualization\Replication" /v DisableCertRevocationCheck /d 1 /t REG_DWORD /f

Note

A slightly different version of this procedure that uses the Certificates snap-in to Microsoft Management Console is available at the Virtualizaition Blog.

See also

- Step 2: Enable Replication
- Deploy Hyper-V Replica

Step 2: Enable Replication

This topic discusses the steps to actually enable replication, including both standalone servers and servers that are part of a failover cluster. You enable replication in both cases both for each virtual machine that is to be replicated and the physical server that is hosting the virtual machines. The activities in Step 1: Prepare to Deploy Hyper-V Replica will enable you to work through this step very quickly.

Although Hyper-V Replica will work if you configure only the Replica server (to receive replication from primary servers), it is a good practice to also configure the primary servers to receive replication. This allows you to reverse the replication direction after a failover, which provides replication protection for the newly active Replica server.

Task	Description
2.1 Configure the Replica server	Configure server that receives the replication data and hosts virtual machines after a failover.
2.2. Configure a Replica server that is part of a failover cluster (optional)	Replica servers that are part of failover clusters are configured for Replica slightly differently.
2.3 Enable replication for virtual machines	Each virtual machine that is to be replicated must be enabled for replication.
2.4 Configure primary server to receive replication	Configure a primary server to receive replication so that it can participate in reverse replication.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

2.1 Configure the Replica server

You configure a server to receive replication with Hyper-V Manager, which is available locally on any server that has Hyper-V installed or installable separately as part of the Remote Server Administration Tools (RSAT). Servers that are part of failover clusters use Failover Cluster Manager to also configure Hyper-V replication. With any of these tools, you can connect to the Replica server remotely to enable replication for it.

Note
The Replica server need not have exactly the same resources (such as memory and CPU) as the primary server. If the Replica server has fewer resources than the primary server, Hyper-V Replica will compensate automatically to the extent possible. However, if the Replica server is completely missing a resource that is present on the primary server (for example, a network switch with a particular name), enabling replication will succeed, but the Replica server will not be able to use the missing resource in the event of a failover.

To configure the Replica server

1. In Hyper-V Manager, click **Hyper-V Settings** in the **Actions** pane.
2. In the **Hyper-V Settings** dialog, click **Replication Configuration**.
3. In the Details pane, select **Enable this computer as a Replica server**.
4. In the **Authentication and ports** section, select the authentication method you decided on in Step 1: Prepare to Deploy Hyper-V Replica. For either authentication method, specify the port to be used (the default ports are 80 for Kerberos over HTTP and 443 for

certificate-based authentication over HTTPS).

5. If you are using certificate-based authentication, click **Select Certificate** and provide the request certificate information.
6. In the **Authorization and storage** section, use the radio buttons to specify whether to allow **any** authenticated (primary) server to send replication data to this Replica server or to limit acceptance to data from specific primary servers. You can use wildcard characters to limit acceptance to servers from a particular domain without having to specify them all individually (for example, *.contoso.com). If you specify individual primary servers, you can designate a separate storage location for Replica data for each one, as well grouping them with the **Trust Group** tag.
7. Click **OK** or **Apply** when you are finished.

```
Set-VMReplicationServer -ReplicationEnabled $true -AllowedAuthenticationType Certificate
-ReplicationAllowedFromAnyServer $true -CertificateThumbprint "<CertThumbprint>" -
DefaultStorageLocation "<Storage Location>" -CertificateAuthenticationPort <Listenerport>
```

Verify that the https listener is running with **netsh http show servicestate**.

Ensure that Windows Firewall is allowing traffic through the port with the following cmdlet:

```
Enable-Netfirewallrule -displayname "Hyper-V Replica HTTPS Listener (TCP-In)"
```

If the Replica server is part of a failover cluster, ensure that Windows Firewall is allowing traffic with this cmdlet instead:

```
Get-ClusterNode | ForEach-Object  {Invoke-command -computername $_.name -scriptblock
{Enable-Netfirewallrule -displayname "Hyper-V Replica HTTPS Listener (TCP-In)"}}
```

2.2. Configure a Replica server that is part of a failover cluster (optional)

Having configured the Hyper-V Replica Broker (see Step 1.4), you enable replication directly from Failover Cluster Manager. If you haven't configured the Hyper-V Replica Broker, follow the directions in Step 1.4. When you complete these steps, Replica is configured for all nodes of the cluster.

Note
If you use certificate-based authentication, ensure that each cluster node and the Hyper-V Replica Broker's client access point all have appropriate machine certificates.

Note
If you use network address translation (NAT), ensure that the inbound and outbound ports are configured to use the same port number. Replica only listens on one port.

To configure a Replica server that is part of a failover cluster

1. In Server Manager, open Failover Cluster Manager.
2. In the left pane, connect to the cluster, and while the cluster name is highlighted, click

Roles in the **Navigate** category of the **Details** pane

3. Right-click the role and choose **Replication Settings**.
4. In the Details pane, select **Enable this cluster as a Replica server**.
5. In the **Authentication and ports** section, select the authentication method you decided on in Step 1: Prepare to Deploy Hyper-V Replica. For either authentication method, specify the port to be used (the default ports are 80 for Kerberos over HTTP and 443 for certificate-based authentication over HTTPS).
6. If you are using certificate-based authentication, click **Select Certificate** and provide the request certificate information.
7. In the **Authorization and storage** section, use the radio buttons to specify whether to allow **any** authenticated (primary) server to send replication data to this Replica server or to limit acceptance to data from specific primary servers. You can use wildcard characters to limit acceptance to servers from a particular domain without having to specify them all individually (for example, *.contoso.com). If you specify individual primary servers, you can designate a separate storage location for Replica data for each one, as well grouping them with the **Trust Group** tag.

2.3 Enable replication for virtual machines

Once the hosting server is configured for Replica, you can enable replication for each virtual machine that you want to be replicated. The details of replication are configured on a per-virtual-machine basis, so each one can have slightly different settings.

Do this step using Windows PowerShell

To enable replication for a virtual machine

1. In the Details pane of Hyper-V Manager, select a virtual machine by clicking it.
2. Right-click the selected virtual machine and point to **Enable Replication**. The Enable Replication wizard opens.
3. On the **Specify Replica Server** page, in the **Replica Server** box, enter either the NetBIOS or fully qualified international domain name (FQIDN) of the Replica server that you configured in Step 2.1. If the Replica server is part of a failover cluster, enter the name of the Hyper-V Replica Broker that you configured in Step 1.4. Click **Next**.
4. On the **Specify Connection Parameters** page, the authentication and port settings you configured for the Replica server in Step 2.1 will automatically be populated, provided that Remote WMI is enabled. If it is not enabled, you will have to provide the values. Click **Next**.
5. On the **Choose Replication VHDs** page, clear the checkboxes for any VHDs that you want to **exclude** from replication, per your decision from 1.1. Make basic planning decisions, and then click **Next**.
6. On the **Configure Recovery History** page, select the number and types of recovery points to be created on the Replica server, per your decisions from 1.1. Make basic planning decisions, and then click **Next**.

Note
In Windows Server 2012 R2, this page of the wizard has been replaced by two pages: **Configure Replication Frequency** and **Configure Additional Recovery Points**.

7. On the **Choose Initial Replication** page, select the initial replication method to use per your decisions from 1.1. Make basic planning decisions. If you choose to have initial replication over the network occur at a later time, configure that here as well, and then click **Next**.
8. On the **Completing the Enable Replication Relationship Wizard** page, review the information in the Summary and then click **Finish**.
9. A dialog box appears indicating that replication was successfully enabled. In this dialog, click the Settings button and provide settings to configure the network that the virtual machine will connect to on the Replica server. The Replica virtual machine does not automatically connect to any network on the Replica server (after a failover) by default, so these settings are important. You can configure the network settings so that the virtual machine will connect to a different network after a failover to the Replica server than it used when it was on the primary server.

   Note
   The configuration data for a virtual machine is replicated only once, at the time of initial replication. After that point, you manage the primary and Replica virtual machines independently.
10. A Replica virtual machine is created on the Replica server. If you elected to send the initial copy over the network, the transmission begins either immediately or at the time you configured. If you elected to use external media for the initial copy, the necessary files are copied to a local location. Copy the files to the media you will use to transfer the initial copy and then send the media to the Replica site.
11. Using Hyper-V Manager or Windows PowerShell on the Replica server, perform initial replication for each virtual machine, using the files you sent to the Replica site.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This example sequence of cmdlets will enable and start replication for a virtual machine (“CRMVM”) that is hosted by a Replica server named “Recovery1.contoso.com”. In this example, Kerberos authentication is configured. All steps must be completed by a user with administrative privileges.

```
Import-Module Hyper-V


$ReplicaServer = “Recovery1.contoso.com”

$RecoveryPort = 8080
```

```
$PrimaryVM1 = "CRMVM"

$PrimaryServer = "Primary1.contoso.com"


Enable-VMReplication -VMName $PrimaryVM1 -ReplicaServerName $ReplicaServer -
ReplicaServerPort $RecoveryPort -AuthenticationType Integrated -CompressionEnabled $true
-RecoveryHistory 0


Start-VMInitialReplication -VMName $PrimaryVM1
```

2.4 Configure primary server to receive replication

Hyper-V Replica normally sends changes that occur on the primary virtual machines to the Replica virtual machines, but after a failover, it can send the data the other direction. By doing this, when you fail over operations from the current primary server to the Replica server, the replication direction can be reversed back to the primary server from the Replica server once the primary has come back online. In this way, you provide replication protection for the Replica server which is now servicing the virtual machine load.

To do this, simply repeat 2.1 Configure the Replica server (or 2.2. Configure a Replica server that is part of a failover cluster (optional)) on one or more of the primary servers, using the same settings you used for the Replica server.

See also

- Step 3: Test the Replication Deployment
- Step 1: Prepare to Deploy Hyper-V Replica
- Deploy Hyper-V Replica

Step 3: Test the Replication Deployment

This topic includes steps to verify the proper functioning of your Hyper-V Replica deployment, primarily by means of the test failover feature.

To ensure that the replicated virtual machines (and applications running within them) will function properly on the Replica server just as they do on the primary server, you can conduct a test failover at any time. When you conduct a test failover, a temporary virtual machine is created on the Replica server. You can test any applications in that test virtual machine without interrupting the ongoing replication. When you conclude the test, the temporary virtual machine is deleted.

Note

This test virtual machine will not by default be connected to any network. If you need to conduct tests that require a network, you can modify the settings of this test virtual machine in the same way you would modify settings of any ordinary virtual machine.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

3.1. Perform test failover

For a test failover to succeed, you must have enabled replication for at least one virtual machine and have completed the initial replication by any of the available methods. To verify failover with a recovery point other than the most recent one, replication must have been running long enough for at least one additional recovery point to have been created.

Do this step using Windows PowerShell

To perform a test failover

1. Access the Replica server, and in Hyper-V Manager, right-click the virtual machine you want to test failover for, point to **Replication…**, and then point to **Test Failover…**.
2. Choose a recovery point to use. This creates and starts a virtual machine with the name "<virtual machine name>-Test" (for example, "CRMVM-Test").
3. Conduct your tests on the test virtual machine. For example, you might verify that the virtual machine starts, pauses, and stops, and that any applications in the virtual machine run properly.

   After you have concluded your testing, discard the test virtual machine by choosing **Stop Test Failover** under the **Replication** option.

   Important
   Operations that move the test virtual machine, such as quick migration, are not supported and could result in the test virtual machine becoming orphaned.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This example sequence of cmdlets will conduct a test failover of a virtual machine named "CRMVM".

```
$ReplicaVM1 = "CRMVM"


$TestReplicaVM1 = Start-VMFailover -AsTest -VMName $ReplicaVM1


Start-VM $TestReplicaVM1
```

To stop the test failover operation and discard the test virtual machine:

```
Stop-VMFailover –VMName $ReplicaVM1
```

See also

- Step 4: Perform a Planned Failover
- Step 2: Enable Replication
- Deploy Hyper-V Replica
- Hyper-V Replica Overview

Step 4: Perform a Planned Failover

This topic includes the steps to perform a planned failover. A planned failover is an operation in which you move operations on any primary virtual machine to its corresponding Replica virtual machine. For example, you might perform a planned failover in order to perform maintenance or upgrades on a primary virtual machine or when you have advance notice of an impending event that might require taking the primary server offline. You might also perform a planned failover to demonstrate and confirm readiness for a disaster recovery scenario.

In a planned failover, any un-replicated changes are first copied over to the Replica virtual machine, so no loss of data occurs. After the planned failover, the Replica virtual machine takes over the workload; to provide similar protection for the virtual machine that is now servicing the workload, you configure “reverse replication” to send changes back to the primary virtual machine (once that comes back online).

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

4.1. Perform a planned failover

Do this step using Windows PowerShell

To perform a planned failover

1. If you have not already done so, follow the procedure in 2.4 Configure primary server to receive replication (optional) to configure the **primary** server to receive replicated data. You might have already used the procedure to configure the Replica server, but for a planned failover, the primary server must also be enabled for replication, so you might

have to repeat the step for the primary server.

2. Start Hyper-V Manager on the primary server and choose a virtual machine to fail over. Turn off the virtual machine that you want to fail over.
3. Right-click the virtual machine, point to **Replication**, and then point to **Planned Failover**.
4. Click **Fail Over** to actually transfer operations to the virtual machine on the Replica server. Failover will not occur if the prerequisites have not been met.

Note

If you remove the recovery points on the Replica server after failover has begun, you commit to the operation and cannot cancel the failover.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sequence of cmdlets will conduct a planned failover of the virtual machine “CRMVM,” provided you have already configured the *primary* server (in addition to the Replica server) to accept replication traffic using the procedure in Step 2. Run these cmdlets on the primary server:

```
$PrimaryVM1 = "CRMVM"

Stop-VM $PrimaryVM1

Start-VMFailover -VMName $PrimaryVM1 -prepare
```

Then run these cmdlets on the Replica server:

```
$ReplicaVM1 = "CRMVM"

Start-VMFailover -VMName $ReplicaVM1
```

Once the primary virtual machine has come online, you can optionally run these commands to reverse the replication direction:

```
Start-VM $ReplicaVM1

Set-VMReplication -reverse -VMName $ReplicaVM1
```

See also

- Step 3: Test the Replication Deployment
- Step 5: Respond to an Unplanned Failover
- Deploy Hyper-V Replica

Step 5: Respond to an Unplanned Failover

This document includes steps to follow to get your virtual machines back online after an unplanned failover has occurred. In an unplanned failover, the primary server has failed unexpectedly, perhaps as a result of a major hardware failure or natural disaster. There is the possibility of data loss, since there was no opportunity to transmit changes on the primary virtual machines that might not have been replicated yet.

In contrast to failover clustering, the workload does not automatically transfer to the virtual machines on the Replica server if the primary server goes offline.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

5.1. Start the Replica virtual machines

Do this step using Windows PowerShell

To start the Replica virtual machines

1. Open Hyper-V Manager and connect to the Replica server.
2. Right-click the name of the virtual machine you want to use, point to **Replication**, and then point to **Failover…**.
3. In the dialog that opens, choose the recovery snapshot you want the virtual machine to recover to, and then click **Failover…**. The Replication Status will change to **Failed over – Waiting completion** and the virtual machine will start using the network parameters you previously configured for it (see 2.3 Enable replication for virtual machines). If you have not previously configured the network parameters or if you want to use different parameters at this point, you can use the Windows PowerShell cmdlet below to do this.
4. Use the Complete-VMFailover Windows PowerShell cmdlet below to complete failover.
5. Repeat these steps for each virtual machine you want to bring back online.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

This sample script will complete an unplanned failover:

```
$Primary = "host1"

$Replica = "host2"

$VMName = "vm01"
```

```
Start-VMFailover -VMName $VMName -ComputerName $Replica

Complete-VMFailover -VMName $VMName -ComputerName $Replica

Start-VM -VMName $VMName -ComputerName $Replica

#
# Assuming that Primary server is recovered.
#

Set-VMReplication -VMName $VMName -AsReplica -ComputerName $Primary

Set-VMReplication -VMName $VMName -Reverse -ReplicaServerName $Primary -ComputerName
$Replica

Start-VMInitialReplication -VMName $VMName -ComputerName $Replica
```

To provide new or different network parameters for the Replica virtual machine to use when it comes online, use the following example cmdlet, substituting appropriate values:

```
$NetworkAdapter = Get-VMNetworkAdapter -VMName $Replica -Name $VMNetworkAdapterName
$NetworkAdapter | Set-VMNetworkAdapterFailoverConfiguration -IPV4Address 172.20.31.100 -
IPV4SubnetMask 255.255.255.0 -IPV4DefaultGateway 172.20.31.1 -IPV4PreferredDNSServer
172.20.31.101 -IPV4AlternateDNSServer 172.20.31.102
```

5.2. Start reverse replication

Once the Replica virtual machine(s) are back online running on the Replica server, they are unprotected against further failure. You can provide this protection by enabling reverse replication back to the primary server once the server comes back online.

Note

The Replica virtual machines must be replicated back to their original primary server. Do not attempt to configure replication from the Replica server to a third server.

Do this step using Windows PowerShell

To start reverse replication

1. Open Hyper-V Manager and connect to the Replica server.
2. Right-click the name of the virtual machine you want to reverse replicate, point to **Replication**, and then point to **Reverse replication…**. The Reverse Replication wizard opens.
3. Complete the Reverse Replication wizard. You will find the requested information to be very similar if not identical to the information you provided in the Enable Replication wizard (see 2.3 Enable replication for virtual machines.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To reverse replication of the virtual machine, use the following example cmdlet:

```
Set-VMReplication -reverse -VMName $ReplicaVM2
```

See also

- Step 4: Perform a Planned Failover
- Deploy Hyper-V Replica

Step 6: Configure and Use Extended Replication

This topic explains how to configure and work with extended replication. In extended replication (available in Windows Server 2012 R2), your Replica server forwards changes that occur on the primary virtual machines to a third server (the extended Replica server). After a planned or unplanned failover from the primary server to the Replica server, the extended Replica server provides further business continuity protection. As with ordinary replication, you configure extended replication by using Hyper-V Manager, Windows PowerShell (using the `-Extended` option), or WMI.

Having the third server allows for new options for planned, unplanned, and test failovers, as well as different ways of moving the virtualized workload among the three servers. Once extended replication is established, you can do essentially all of the same things with the Replica and extended Replica servers just as you would with the primary server and the Replica server.

Note

Because the roles of primary, Replica, and extended Replica can be exchanged among the various servers, to keep the explanations as clear as possible, in this topic we will assume a starting configuration of three servers: Server A, Server B, and Server C. This topic assumes that you have already configured standard replication along the lines of Steps 1-5 of this topic collection, where Server A is the primary server and Server B is the Replica server. The extended Replica server created in these steps will initially be Server C.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Enable extended replication

Do this step using Windows PowerShell

To enable extended replication

1. In Hyper-V Manager, connect to the Replica server (in the initial state of this example, Server B).
2. Right-click the virtual machine and point to **Extend Replication**. The Extend Replication wizard opens.
3. On the **Specify Replica Server** page, enter the name of the server you intend to use as the extended Replica (in this example, Server C; not the current Replica server, Server B).
4. Complete the wizard, providing values as you did when you originally configure replication between the primary and the Replica server. You can use different values for replication frequency, authentication method, and initial replication method than those you used for the basic replication configuration.

   Notes
   There are certain differences in the Replica-extended Replica relationship as compared to the primary-Replica relationship:

   - You must use a longer replication frequency for extended replication than was configured for the primary-Replica relationship. Only frequencies of 5 or 15 minutes are supported for extended replication.
   - If you have configured application-consistent snapshots in the primary-Replica relationship, these snapshots are forwarded to the extended Replica. However, you cannot set a different frequency for the snapshots or enable application-consistent snapshots for extended replication independent of the primary-Replica relationship.
   - The collection of VHDs used in the primary-Replica relationship cannot be different in extended replication.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

You extend replication with the same cmdlet used to enable ordinary replication. However, you run it from the current Replica server and use `-ReplicaServerName` to define the *extended* Replica server.

```
Enable-VMReplication –VMName <vmname> -ReplicaServerName <extended_server_name> -
ReplicaServerPort <Auth_port> -AuthenticationType <Certificate/Kerberos> -
ReplicationFrequencySec <300/900> [--other optional parameters if needed--]
```

Monitor extended replication

To monitor the status and statistics of extended replication, start Hyper-V Manager, right-click the Replica virtual machine, and then point to **View Replication Health**.

Note

To see the data for the extended replica, you must view the health report from the Replica server. From the primary server, you can only access the report for the primary replication relationship.

With Windows PowerShell, you can use the `Measure-VMReplication` cmdlet with the extra parameter `-ReplicationRelationshipType`. If you use the parameter without specifying a value, statistics for both the primary and extended replication are returned. To specify the extended replication relationship, use the parameter with the value `Extended`; for the primary relationship, use the value `Simple`.

Examples:

To view the statistics of the extended replication relationship from the Replica server:

```
Measure-VMReplication –VMName <name> -ReplicationRelationshipType Extended | select *
```

To view the statistics of the primary replication relationship from the Replica server:

```
Measure-VMReplication –VMName <name> -ReplicationRelationshipType Simple | select *
```

To view replication relationship statistics for both the primary and extended replication relationships:

```
Measure-VMReplication –VMName <name> | select *
```

Pause and resume extended replication

Pausing and resuming extended replication works the same way as for primary replication. You can control primary and extended replication separately. If you access the Replica virtual machine with Hyper-V Manager, you will find separate options for pausing or resuming either the

primary replication or the extended replication. If you access the extended Replica virtual machine, you will find pause and resume options for only the extended replication.

With Windows PowerShell, you can use the `Suspend-VMReplication` and `Resume-VMReplication` cmdlets with the extra parameter `-ReplicationRelationshipType`. To specify pausing or resuming the extended replication relationship, use the parameter with the value `Extended`; for the primary relationship, use the value `Simple` (which is the default if you don't specify the replication relationship type.

Examples:

To pause or resume extended replication from the Replica server:

```
Suspend-VMReplication/ Resume-VMReplication –VMName <name> -ReplicationRelationshipType Extended
```

To pause or resume primary replication from the Replica server:

```
Suspend-VMReplication/ Resume-VMReplication –VMName <name> -ReplicationRelationshipType Simple
```

To pause or resume primary replication (if performed on the primary server) or extended replication (if performed on the extended Replica server):

```
Suspend-VMReplication/ Resume-VMReplication –VMName <name>
```

Resynchronization

To recover from certain errors, it might be necessary to resynchronize replication. If an error has occurred, you can resynchronize either primary or extended Replication with the following commands.

Examples:

To resynchronize replication from the primary server in the event of an error:

```
Suspend-VMReplication –VMName <name>//puts the virtual machine state to suspended

Resume-VMReplication –VMName <name> -ReplicationRelationshipType Extended -Resync
```

To resynchronize extended replication from the Replica server:

```
Suspend-VMReplication –VMName <name> -ReplicationRelationshipType Extended //puts the extended virtual machine state to suspended

Resume-VMReplication –VMName <name> -ReplicationRelationshipType Extended -Resync
```

In some circumstances, it is possible that replication could be in the "Resynchronization Required" state. If this occurs, you can use Hyper-V Manager to perform resynchronization immediately or schedule it for a later time. If the primary replication relationship is resynchronizing, extended replication will move into the Resynchronization Required state. When primary resynchronization is complete, extended resynchronization commences automatically.

You can also start resynchronization with Windows PowerShell.

Examples:

To resynchronize replication from the primary server when in Resynchronization Required state:

```
Resume-VMReplication –VMName <name> -Resync
```

To resynchronize extended replication from the Replica server when in Resynchronization Required state:

```
Resume-VMReplication –VMName <name> -ReplicationRelationshipType Extended -Resync
```

Test failover with an extended Replica server

Test failovers from the Replica server to the extended Replica server work just as they do in a basic replication relationship (described in Step 3: Test the Replication Deployment). During the test failover, the relationship between the Replica server and the extended Replica server remains intact. During the test, you can also pause or resume operations between Replica and extended Replica.

New options for planned failover

Having an extended Replica server affords more options when you perform a planned failover. In a basic replication configuration (without an extended Replica server), once you perform a planned failover from Server A to Server B, you typically reverse replication from B to A to provide replication protection for B, which is now hosting the active virtualized workload.

With an extended Replica server, once you have completed a planned failover from A to B, B acts as the primary server, but you can either use C as its Replica server or reverse replication back to A.

To fail over to Server B and use Server C as the new Replica

1. Use Hyper-V Manager or Windows PowerShell to shut down the virtual machine on the primary server (in this example, Server A).
2. Start a planned failover from Server A to Server B, as described in Step 4: Perform a Planned Failover.
3. In the Planned Failover dialog box, clear **Reverse the replication direction after failover** checkbox and select the **Start the Replica virtual machine after failover** checkbox.
4. Click **Fail Over**.
5. After failover is complete, in Hyper-V Manager, right-click the virtual machine that is now on the Replica server (Server B) and point to **Resume Extended Replication**. This removes the replication relationship between Servers B and A; as result, B is now the primary server, and C is its Replica server.

To do this with Windows PowerShell, run `Start-VMFailover -VMName <name> -Prepare` on the primary server (Server A), then run `Start-VMFailover -VMName <name>` on Server B. The virtual machine is now failed over to Server B, which is the new primary server.

To use Server C as the new Replica server to protect Server B, run the following cmdlets on Server B:

```
Complete-VMFailover -VMName <name> // to commit the failover

Start-VMFailover -Prepare

Resume-VMReplication -VMName <name> -Continue // to resume extended replication
```

To fail over to Server B and use Server A as the new Replica

1. Use Hyper-V Manager or Windows PowerShell to shut down the virtual machine on the primary server (in this example, Server A).
2. Start a planned failover from Server A to Server B, as described in Step 4: Perform a Planned Failover.
3. In the Planned Failover dialog box, select the **Reverse the replication direction after failover** checkbox and select the **Start the Replica virtual machine after failover** checkbox.
4. After failover is complete, in Hyper-V Manager, right-click the virtual machine that is now on the Replica server (Server B) and point to **Reverse Replication**. This removes the extended replication relationship between Servers B and C; as result, B is now the primary server, and A is its Replica server.

To do this with Windows PowerShell, run `Start-VMFailover -VMName <name> -Prepare` on the primary server (Server A), then run `Start-VMFailover -VMName <name>` on Server B. The virtual machine is now failed over to Server B, which is the new primary server.

To use Server A as the new Replica server to protect Server B, run the following cmdlets on Server B:

```
Remove-VMReplication -VMName <name> -ReplicationRelationshipType Extended

Set-VMReplication -VMName <name> -Reverse
```

In some circumstances, you might want to take both the primary and Replica servers (Servers A and B) offline, and have the extended Replica (Server C) maintain the virtualized workload.

To fail over completely to the extended Replica (Server C)

1. Shut down the virtual machine on the primary server (Server A), and then complete a planned failover to the Replica server (Server B) (see Step 4: Perform a Planned Failover).

2. Once failover to the Replica server is complete, complete a planned failover to the extended Replica server (Server C).

To do this with Windows PowerShell, run `Start-VMFailover –VMName <name> -Prepare` on the primary server (Server A), then run `Start-VMFailover –VMName <name>` and `Complete-VMFailover –VMName <name>`on Server B. The virtual machine is now failed over to Server B, which is the new primary server.

Fail over from Server B to Server C by running `Start-VMFailover –VMName <name> -Prepare` on Server B.

Finally, complete the failover to Server C by running the following on Server C:

```
Start-VMFailover –VMName <name>

Complete-VMFailover –VMName <name>
```

Unplanned failover from primary to Replica with an extended Replica server

If the primary server (Server A) goes down, you can complete an unplanned failover to the Replica server (Server B) with Hyper-V Manager by right-clicking the virtual machine on Server B and pointing to **Failover**. In the dialog that opens, select a recovery point and click **Fail Over**.

To do this with Windows PowerShell, run `Start-VMFailover –VMName <name> -VMRecoverySnapshot <VMSnapshot>`.

Once failover to Server B is complete, you can choose to continue replication to Server C (to provide protection to Server B), or if Server A comes back online, reverse replication from Server B to Server A.

To continue replication to Server C, on Server B, right-click the virtual machine in Hyper-V Manager and point to **Resume Extended Replication**.

Note

This action removes replication between Servers A and B. Server B becomes the primary server and Server C becomes the Replica server. Replication from Server B to Server C goes into the resynchronization state and replication continues automatically once resynchronization is complete. The virtual machine on Server A becomes orphaned.

To do this with Windows PowerShell, run the following cmdlets on Server B:

```
Complete-VMFailover –VMName <name> // to commit the failover

Resume-VMReplication –VMName <name> -Continue // to resume extended replication
```

At a later point, when all three servers are back online, you can restore each server to its original function by following these steps:

To restore the servers to their original roles

1. Remove replication from all three servers with Hyper-V Manager or Windows PowerShell.
2. Run the Enable Replication wizard on Server B, choosing **Use an existing virtual machine on the Replica server as the initial copy**.
3. Perform a planned failover from Server B to Server A.
4. Extend replication from Server B to Server C.

To reverse replication to Server A after the unplanned failover, on Server B, right-click the virtual machine in Hyper-V Manager and point to **Reverse Replication**. The extended replication relationship is removed; Server B becomes the primary server and Server A becomes the Replica server. The virtual machine on Server C is orphaned.

To do this with Windows PowerShell, on either Server A or Server B, run `Set-VMReplication -VMName <name> -AsReplica`. Then on Server B, run:

```
Set-VMReplication –VMName <name> -Reverse -ReplicaServerName <new_recovery_server_name> -ReplicaServerPort <Auth_port> -AuthenticationType <Certificate/Kerberos> -ReplicationFrequencySec <30/300/900> [--other optional parameters if needed--]
```

At a later point, when all three servers are back online, you can restore each server to its original function by following these steps:

To restore the servers to their original roles

1. Remove replication from Server C with Hyper-V Manager or Windows PowerShell.
2. Perform a planned failover from Server B to Server A; choose the **Reverse Replication** option in the wizard. At this point, Server A is the primary server and Server B is the Replica server.
3. On Server B, run the Extend Replication wizard, choosing to **Use an existing virtual machine on the Replica server as the initial copy**.

Unplanned failover when both primary and Replica servers go offline

If a failover has shifted the virtualized workload to the extended Replica server (Server C in these examples), it is now functioning as the primary server. If both Server A and Server B go down, complete the unplanned failover to Server C with Hyper-V Manager by right-clicking the virtual machine on Server C and pointing to **Failover**. In the dialog that opens, select a recovery point and click **Fail Over**.

Once failover is complete, you have three options for putting the other servers back into service:

Repurpose the Replica server (Server B, which previously replicated Server A) to be a Replica server replicating Server C:

In Hyper-V Manager, access Server C, right-click the virtual machine, and then point to **Reverse Replication**.

If a primary replication relationship is detected on Server B, you will receive an error. If this occurs, remove the replication relationship on B, and then repeat the reverse replication step from Server C.

When reverse replication from Server C is complete, Server C will be the primary server, Server B is the Replica server, and the virtual machine on Server A is orphaned.

To restore the servers to their original roles

1. Perform a planned failover of the virtual machine from Server C to Server B.
2. Remove replication on all three servers.
3. Enable replication from Server B to Server A, choosing to **Use an existing virtual machine on the Replica server as the initial copy**
4. Perform a planned failover from Server B to Server A, selecting the check box for enabling reverse replication.
5. On Server B, run the Extend Replication wizard, to extend replication to Server C.

Repurpose the original primary server (Server A) to be a Replica server replicating Server C:

With Hyper-V Monitor, access Server A, right-click the virtual machine, and point to **Remove Replication**.

On Server B, run the Enable Replication wizard to set up replication to Server A, choosing to **Use an existing virtual machine on the Replica server as the initial copy**. If an orphaned virtual machine is still present on Server A, the wizard can use it as a data source to complete the initial replication.

To restore the servers to their original roles

1. Perform a planned failover of the virtual machine from Server C to Server A.
2. Remove replication on all three servers.
3. Enable replication from Server A to Server B, choosing to **Use an existing virtual machine on the Replica server as the initial copy**
4. Perform a planned failover from Server A to Server B, selecting the check box for enabling reverse replication.
5. On Server B, run the Extend Replication wizard, to extend replication to Server C.

Use a new server (Server D) as the Replica server

Prepare a new server for Hyper-V and replication as described in Step 2: Enable Replication.

Access Server C and remove replication for the virtual machine.

Enable replication to Server D for the virtual machine.

Moving primary or Replica virtual machines to different servers or clusters

You can move either primary or Replica virtual machines around to different servers either with quick migration or export/import, but in many cases resynchronization will be required.

Removing replication or extended replication

You can remove extended replication with Hyper-V Manager by right-clicking the Replica virtual machine and then pointing to **Replication**. Point to Remove Replication and then select either **Remove extended replication** or **Remove all replication**. If you remove the replication relationship between the Replica and the extended Replica, changes in the virtual machine are no longer forwarded to the extended Replica, but the virtual machine on the extended Replica is not deleted.

Note

If you want to remove replication between the primary and the Replica, remove replication between the Replica and the extended Replica first.

With Windows PowerShell, use the `Remove-VMReplication` cmdlet with the parameter `-ReplicationRelationshipType`. Use the value `Extended` to remove extended replication or `Simple` to remove primary replication.

Examples:

To remove extended replication from the Replica server:

```
Remove-VMReplication –VMName <name> -ReplicationRelationshipType Extended
```

To remove primary replication from the Replica server:

```
Remove-VMReplication –VMName <name> -ReplicationRelationshipType Simple
```

To remove all replication (run from primary or extended Replica server):

```
Remove-VMReplication –VMName <name>
```

See also

- Step 3: Test the Replication Deployment
- Step 5: Respond to an Unplanned Failover
- Deploy Hyper-V Replica

Hyper-V Module for Windows PowerShell

Hyper-V in Windows Server® 2012 introduces the Hyper-V module for Windows PowerShell, which includes more than 160 Hyper-V cmdlets. These cmdlets provide an easy way to automate Hyper-V management tasks.

For an alphabetical list of all Hyper-V cmdlets Windows Server 2012, as well as links to reference content for each cmdlet (including examples), see Hyper-V Cmdlets in Windows PowerShell.

Get updated Help

The Hyper-V module supports the Update-Help cmdlet, which allows you to download the newest help files for Windows PowerShell modules and install the files on your computer. Your account must be a member of the Administrators group on the computer you want to update. To update Help for the Hyper-V module, start Windows PowerShell with the "Run as administrator" option (`Start-Process PowerShell -Verb RunAs`) and run the following command:

```
PS C:\> Update-Help -Module Hyper-V
```

For more information, see Update-Help.

See Also

Hyper-V Overview

Hyper-V Configuration

Hyper-V configuration:

Configure Live Migration and Migrating Virtual Machines without Failover Clustering

Configure Storage Quality of Service

Configure a Shared Virtual Hard Disk

Configure Online Virtual Hard Disk Resize

Configure Live Migration and Migrating Virtual Machines without Failover Clustering

This article shows you how to configure and perform a live migration without using failover clustering. A live migration moves running virtual machines between servers running Hyper-V without any perceived downtime. This topic also includes instructions for configuring constrained delegation if you want to use remote management tools to perform live migrations.

In this document

- Prerequisites
- Step 1: [Optional] Configure constrained delegation
- Step 2: Configure the source and destination computers for live migration
- Step3: [optional] Configure performance options for live migration
- Step 4: Move a running virtual machine
- Step 5: [Optional] Move a running virtual machine again (back to the original host or to another host)

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

Before you perform the steps in this document, make sure that your environment meets the following prerequisites:

- The Hyper-V role in Windows Server 2012 or Windows Server 2012 R2 is installed on the source and destination servers, and at least one running virtual machine exists on the source computer. For instructions, see Install Hyper-V and Configure a Virtual Machine.

  Warning
  In Windows Server 2012 R2 Hyper-V live migration has been updated to support migrating Hyper-V virtual machines in Windows Server 2012 to Hyper-V in Windows Server 2012 R2. Moving a virtual machine to a down-level server running Hyper-V is not supported.
- The source and destination computers either belong to the same Active Directory domain, or belong to domains that trust each other.
- The user account has the appropriate permission to perform the various steps:
  - The account you use to configure constrained delegation must be a member of the Domain Administrators group.
  - The account you use to configure live migration and to perform the live migrations must be a member of the local Hyper-V Administrators group or the Administrators group on both the source and destination computers.

- You have a computer running Windows® 8 with the Hyper-V management tools installed. For instructions, see http://technet.microsoft.com/library/hh857623.aspx.

In addition to ensuring all prerequisites are met, consider the following planning decisions. These decisions determine how you configure and perform live migrations.

- Do you plan to sign on to each server to perform the given task (either through a local console session, a Remote Desktop session, or a remote Windows PowerShell session), or do you want perform the tasks with remote management tools? The answer determines whether you should select Kerberos or Credential Security Support Provider (CredSSP) to authenticate live migration traffic. To manage the tasks with remote management tools, configure constrained delegation and select Kerberos as the authentication protocol. Otherwise, you must sign on to the source computer to perform a live migration, and CredSSP is used to authenticate the live migration.

  Notes

  The requirement of signing in to the source computer has implications that might not be obvious. For example, if you sign in to TestServer01 to move a virtual machine to TestServer02, and then want to move the virtual machine back to TestServer01, the operation will fail unless you sign in to TestServer02 before you try to move the virtual machine back to TestServer01.

  If the connection between the source and destination computers cannot be authenticated, an error occurs and the following message is displayed:

  Virtual machine migration operation failed at migration Source.

  Failed to establish a connection with host <*computer name*>**: No credentials are available in the security package (0x8009030E).**

- Live migration performance options. Applies to Hyper-V in Windows Server 2012 R2.

  The performance options can reduce overhead on the network and CPU usage in addition to the reducing the amount of time for a live migration. Hyper-V administrators can configure the appropriate live migration performance options based on their environment and requirements.

  The following live migrations performance options are available.

Option	Description
TCP/IP	The memory of the virtual machine is copied to the destination server over a TCP/IP connection. This is the same method that is used in Hyper-V in Windows Server 2012.
Compression	The memory content of the virtual machine that is being migrated is compressed and then copied to the destination server over a TCP/IP connection. **Note** This is the **default** setting in Hyper-V in Windows Server 2012 R2.
SMB	The memory content of the virtual machine is copied to the destination server over a SMB 3.0 connection. • SMB Direct is used when the network adapters on the source and destination servers have Remote Direct Memory Access (RDMA) capabilities enabled. • SMB Multichannel automatically detects and uses multiple connections when a proper SMB Multichannel configuration is identified. For more information, see Improve Performance of a File Server with SMB Direct.

- Will you allow live migration traffic through any available network, or isolate the traffic to specific networks? As a security best practice, we recommend that you isolate the traffic onto trusted, private networks because live migration traffic is not encrypted when it is sent over the network. Network isolation can be achieved through a physically isolated network or through another trusted networking technology such as VLANs.

Step 1: [Optional] Configure constrained delegation

If you have decided to use Kerberos to authenticate live migration traffic, configure constrained delegation before you proceed to the rest of the steps. As noted in the “Prerequisites” section,

you must use an account that is a member of the Domain Administrators group to configure constrained delegation.

To configure constrained delegation

1. Open the Active Directory Users and Computers snap-in. For example, to do this from Server Manager, select the server if it not already selected. After the server is selected, click **Tools**, and then click **Active Directory Users and Computers**. This opens the Active Directory Users and Computers snap-in.
2. From the navigation pane, select the domain and double-click the **Computers** folder.
3. From the Computers folder, right-click the computer account of the source server and then click **Properties**.
4. In the Properties dialog box, click the **Delegation** tab.
5. On the delegation tab, select **Trust this computer for delegation to the specified services only**. Under that option, select **Use Kerberos only**.
6. Click **Add**.
7. In the Add Services dialog box, click **Users or Computers**.
8. In the **Select Users or Computers** dialog box, type the name of the destination server. Click **Check Names** to verify that you typed the name correctly, and then click **OK**.
9. In the Add Services dialog box, in the list of available services, do the following and then click **OK**.
   - To move virtual machine storage, select **cifs**. This is required if you want to move the storage along with the virtual machine, as well as if you want to move only a virtual machine's storage. If the server is configured to use SMB storage for Hyper-V, this should already be selected.
   - To move virtual machines, select **Microsoft Virtual System Migration Service**.
10. On the **Delegation** tab of the Properties dialog box, verify that the services you selected in the previous step are listed as the services to which the destination computer can present delegated credentials. Click **OK**.
11. From the **Computers** folder, select the computer account of the destination server and repeat the process. In the **Select Users or Computers** dialog box, be sure to specify the name of the source server.

Notes

The configuration changes do not take effect until the following has occurred:

- The changes have replicated to the domain controllers that the servers running Hyper-V are logged into.
- A new kerboros ticket has been issued.

Step 2: Configure the source and destination computers for live migration

In this step, you will enable live migrations, configure the source and destination servers so they will send and receive live migrations, optionally configure the authentication protocol. When you configure the servers, you choose whether to allow live migration traffic on any available network, or only on specified networks. As a security best practice, we recommend that you select specific networks to use for live migration traffic, as discussed in the “Prerequisites” section.

Do this step using Windows PowerShell

To configure the source and destination computers for live migration

1. Open Hyper-V Manager if it is not already open. (From Server Manager, click **Tools** and then click **Hyper-V Manager**.)
2. In the navigation pane, select one of the servers that you want to configure for live migrations. (If none of the servers are listed, click **Hyper-V Manager** in the navigation pane. Then, in the **Action** pane, click **Connect to Server** and specify each server name. After that, select one of the servers in navigation pane.)
3. In the **Action** pane, click **Hyper-V Settings**.
4. In Hyper-V Settings dialog box, click **Live Migrations**.
5. In the **Live Migrations** pane, check **Enable incoming and outgoing live migrations**.
6. Under **Simultaneous live migrations**, specify a different number if you don’t want to use the default of 2.
7. Under **Incoming live migrations**, if you want to use specific network connections to accept live migration traffic, click **Add** to type the IP address information. Otherwise, click **Use any available network for live migration**. Click **OK**.
8. If you have configured constrained delegation in Step 1: [Optional] Configure constrained delegation, expand **Live Migrations** and then select **Advanced Features**.

   In the **Advanced Features** pane, under **Authentication protocol**, select Kerberos.
9. Click **OK**.
10. Select the other server in Hyper-V Manager and repeat the remaining steps.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To use Windows PowerShell to configure a server running Hyper-V for live migration, you can use the Enable-VMMigration, Set-VMMigrationNetwork, and Set-VMHost cmdlets, depending on how you want to configure the host. The following example commands configure live migration on the local host, allow incoming migration traffic only on the specified network, and specify Kerberos as the authentication protocol. Each line represents a separate command.

```
PS C:\> Enable-VMMigration
```

```
PS C:\> Set-VMMigrationNetwork 192.168.10.1
```

```
PS C:\> Set-VMHost –VirtualMachineMigrationAuthenticationType Kerberos
```

Step3: [optional] Configure performance options for live migration

Note

This applies to Hyper-V in Windows Server 2012 R2.

In this step, you configure the live migration performance options. The performance options can reduce overhead on the network and CPU usage in addition to the reducing the amount of time for a live migration. Hyper-V administrators can configure the appropriate live migration performance options based on their environment and requirements.

Do this step using Windows PowerShell

To configure live migration performance options

1. From Hyper-V Manager, in the navigation pane, select one of the servers you want to configure for live migrations.
2. In the **Action** pane, click **Hyper-V Settings**.
3. In the Hyper-V Settings dialog box, expand **Live Migrations** and then select **Advanced Features**.
4. In the **Advanced Features** pane, under **Performance options**, select the appropriate option for your environment, and then click **OK**.

   Note

   When you enable live migration the first time the default setting for **Performance options** is **Compression**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To use Windows PowerShell to configure a server running Hyper-V for live migration, you can use the Enable-VMMigration, Set-VMMigrationNetwork, and Set-VMHost cmdlets, depending on how you want to configure the host. The following example command sets the live migration performance option to SMB.

```
PS C:\> Set-VMHost -VirtualMachineMigrationPerformanceOption SMB
```

Step 4: Move a running virtual machine

In this step, you move a running virtual machine from one server running Hyper-V in Windows Server 2012 to another server running Windows Server 2012 or Windows Server 2012 R2.

Do this step using Windows PowerShell

To move a running virtual machine

1. Connect to the source server using one of the following methods (unless you are signed in to the source server):
   - From Server Manager, click **All Servers** and then click the name of the source server. (For more information about managing multiple servers, see Manage Multiple, Remote Servers with Server Manager.)
   - From Hyper-V Manager, in the navigation pane, click the name of the source server.
2. From the **Virtual Machines** section of Hyper-V Manager, right-click the virtual machine and then click **Move**.
3. On the **Choose Move Type** page of the Move Wizard, choose **Move the virtual machine**.
4. On the **Specify Destination** page, type the name or browse to the destination computer.
5. On the **Choose Move Options** page, select whether you want to move various items stored for a virtual machine, such as virtual hard disks, snapshots, and paging file.
6. On the **Summary** page, review your choices and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To use Windows PowerShell to move a running virtual machine, you use the Move-VM cmdlet. The following example moves a virtual machine named *LMTest* to a destination server named *TestServer02* and moves the virtual machine's virtual hard disks and other files (such as any snapshots and Smart Paging files) to the *D:\LMTest* directory on the destination server.

```
PS C:\> Move-VM LMTest TestServer02 –IncludeStorage –DestinationStoragePath D:\LMTest
```

For reference information about the Move-VM cmdlet, see http://go.microsoft.com/fwlink/?LinkID=217784.

Step 5: [Optional] Move a running virtual machine again (back to the original host or to another host)

This step is optional because you might not need to move the virtual machine back again. However, it's included to mention the fact that, if you haven't enabled constrained delegation, you must sign in to source server to complete the step. Then, complete this step using the same instructions as given in the preceding section.

Important
If the original source server was Windows Server 2012 and original destination server was Windows Server 2012 R2 you cannot migrate back to the original source server.

See also

- Hyper-V Overview
- What's New in Hyper-V for Windows Server 2012
- Virtual Machine Live Migration Overview
- Deploy Hyper-V over SMB

Configure Storage Quality of Service

Hyper-V enables an administrator to configure your virtual machine hard disks for Storage Quality of Service (QoS).

Note
Storage QoS cannot be installed as a separate feature. It is only available when a user installs the Hyper-V role. For more information about storage QoS, see Storage Quality of Service for Hyper-V.

To configure a virtual hard disk for Storage Quality of Service

You can configure the Storage QoS parameters for the bandwidth on the virtual hard disk of your virtual machine using the following procedure.

1. In **Server Manager**, point to **Tools**, and then click **Tools**.
2. Point to **Hyper-V Manager** and then click **Hyper-V Manager**.
3. Click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**
4. In the **Action** pane, click **Settings**.
5. Double click **SCSI Controller**.
6. Point to **SCSI Controller** and then click **SCSI Controller**.
7. Click **Hard Drive**. You will see **Hard Drive** settings displayed in the right pane.
8. To the left of the Hard Drive icon, click "**+**".
9. Click **Advanced Features**.
10. In the **Advanced Features** pane, click **Enable Quality of Service management**.
11. Enter **Maximum** IOPS values. The default value is 0..
12. Enter your **Minimum** IOPS values. The default value is 0.
13. Click **OK**. .

Configure a Shared Virtual Hard Disk

Shared VHDX files should be deployed on Cluster Shared Volumes (CSV) or on a Scale-Out File Server Cluster with SMB 3.0. For more information about shared virtual hard disks, see Virtual Hard Disk Sharing Overview.

To create and enable a shared virtual hard disk on a virtual machine

1. Open Hyper-V Manager if it is not already open. (From Server Manager, click **Tools** and then click **Hyper-V Manager**.)
2. Under **Virtual Machines**, select the virtual machine that you want to configure with a shared virtual hard disk..
3. In the **Actions** pane, click **Settings**.
4. Click **SCSI Controller**.
5. On the left pane, under **SCSI Controller**, click **Add**.
6. On the left pane, under **Hard Drive**, click **New** to start the **New Virtual Hard Disk Wizard.**.
7. On the **Before You Begin** page, click **Next**
8. On the **Choose Disk Format**, accept the default **VHDX**. Click **Next**.
9. On the **Choose Disk Type** Page, accept the default **Dynamic Expanding**. Click **Next**.
10. On the **Specify Name and Location** page, enter the name of your shared virtual hard disk..
11. On the **Specify Name and Location** page, click **Browse**. Navigate to your shared storage location. (Cluster Shared Volumes (CSV) or on a Scale-Out File Server Cluster with SMB 3.0).
12. Click **Select Folder**. Click **Next**.
13. On the **Configure Disk** page, verify the size of the new virtual hard disk and alter to desired size or accept default. Click **Next**.
14. On the Completing the New Virtual Hard Disk Wizard, review your configuration and Click **Finish**
15. The new virtual hard disk is now created and needs to be enabled as a shared virtual hard disk. Click the Hard Drive you just created under **SCSI Controller**.
16. Click the "**+**" symbol next to the hard disk.
17. Click **Advanced Features**.
18. On the right hand size, under **Virtual hard disk sharing**, click **Enable virtual hard disk sharing**.
19. Click **OK**. The shared virtual hard disk is now created and enabled.

Configure Online Virtual Hard Disk Resize

Starting in Windows Server® 2012 R2, Hyper-V provides you with the ability to expand or shrink decrease the size of a virtual hard disk while the virtual machine is still running. This topic creates a virtual hard disk (attached to a SCSI Controller), adds the virtual hard disk to the virtual machine and details the process for expanding and shrinking a virtual hard disk.

Note

The user interface option to shrink a virtual hard disk is only visible for virtual hard disks that have been expanded previously. Online resize is exclusive to SCSI attached VHDX files. For more information about resizing virtual hard disks, see Online Virtual Hard Disk Resizing Overview.

Prerequisites

Before you attempt to edit the size of your virtual hard disk, make sure that you have the following:

- A computer running Windows Server® 2012 R2 or later with the Hyper-V role installed.
- A user account that is a member of the local Hyper-V Administrators group, or the Administrators group.

To create a virtual hard disk for a virtual machine

Create a virtual hard disk (attached to a SCSI controller) for a virtual machine using the following procedure:

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Tools** menu, click **Hyper-V Manager**.
3. Under **Virtual Machines**, select the virtual machine..
4. In the **Actions** pane, click **Settings**.
5. Click **SCSI Controller**. On the right pane, under **SCSI Controller**, click **Hard Drive**..
6. Click **Add**
7. On the right pane, under **Hard Drive**, click **New**.
8. On the **Before You Begin** page, click **Next**
9. On the **Choose Disk Format** page, choose **VHDX** and click **Next**.
10. On the **Specify Name and Location** page, type the name of the new virtual hard disk. If required, type the location of the virtual hard disk. Click **Next**.
11. On the **Configure Disk** page, type disk size and click **Next**
12. On the **Completing the New Virtual Hard Disk Wizard** page, click **Finish**.

Adding a virtual hard disk to the virtual machine

After adding a virtual hard disk, it is necessary to configure your virtual machine so that the operating system within the virtual machine can see the additional disk capacity. The disk must be brought online, initialized and formatted.

These instructions assume a virtual hard disk has been created for a Windows Server® 2012 R2virtual machine.

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Tools** menu, click **Hyper-V Manager**.
3. Under **Virtual Machines**, double click to connect to the virtual machine.
4. Enter credentials and login to the virtual machine.
5. To start **Disk Manager**, start the **Server Manager**. .
6. On the **Tools** menu, click **Computer Management**. **Add**
7. Under **Storage**, click **Disk Management**.
8. On the middle pane, locate the new hard disk. Right click the disk and click **Online**.
9. On the middle pane, locate the new hard disk. Right click the disk and click **Initialize**.
10. On the **Initialize Disk** page, click **OK**.
11. On the middle pane, right click the unallocated volume and click **New Simple Volume**.
12. In the **New Simple Volume Wizard**, click **Next**.
13. On the **Specify Volume Size** page, click **Next**.
14. On the **Assign Drive Letter or Path** page, click **Next**.

    On the **Format Partition** page, click **Next**

    On the **Completing the New Simple Volume** page, click **Finish**.

The new virtual hard disk will now be visible to the virtual machines operating system.

To expand the size of a virtual hard disk

You can expand the storage capacity of the virtual hard disk using the following procedure:

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Tools** menu, click **Hyper-V Manager**.
3. Under **Virtual Machines**, select the virtual machine with the virtual hard disk that you want to expand.
4. In the **Actions** pane, click **Edit Disk** to start the **Edit Virtual Hard Disk Wizard**.
5. On the Before you Begin page, click Next.
6. On the **Locate Virtual Hard Disk** page, click **Browse**.
7. Navigate to your virtual hard disk. Select your virtual hard disk. Click **Open**.
8. Click **Next**.

9. On the **Expand Virtual Disk** page, enter the new virtual hard disk size. The size is specified in gigabytes with a maximum size of 64TB for any virtual hard disk.
10. Click **Next**.
11. On the **Completing the Edit Virtual Hard Disk Wizard** page, click **Finish**.

Expand the volume of a virtual hard disk connected to a virtual machine

After expanding a virtual hard disk, it is necessary to configure the disk management within your virtual machine operating system so that it can see the new disk space. This space is currently residing as an unallocated volume created in the virtual disk.

These instructions assume a virtual hard disk has been expanded and that it is connected to a Windows Server® 2012 R2virtual machine.

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Tools** menu, click **Hyper-V Manager**.
3. Under **Virtual Machines**, double click the virtual machine.
4. Enter credentials and login to the virtual machine.
5. To start **Disk Manager**, start the **Server Manager**.
6. On the **Tools** menu, click **Computer Management.**
7. Under **Storage**, click **Disk Management**.
8. On the middle pane, locate the existing volume in use by the virtual machine.
9. Right click the disk and click **Extend Volume**.
10. On the **Welcome to the Extended Volume Wizard**, click **Next**.
11. On the **Select Disk** page, click **Next**.
12. On the **Completing the Extended Volume** page, click **Finish**.

The operating system within the virtual machine will now see the expanded disk capacity.

To shrink the size of a virtual hard disk

You can shrink the size of a virtual hard disk. However, you cannot shrink the size of the virtual hard disk below the size of the disk volume that is in use within the operating system of the virtual machine.

Reduce the storage capacity of the virtual hard disk using the following procedure:

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Tools** menu, click **Hyper-V Manager**.
3. Under **Virtual Machines**, select the virtual machine with the virtual hard disk that you want to shrink.
4. In the **Actions** pane, click **Edit Disk** to start the **Edit Virtual Hard Disk Wizard**.
5. On the **Before you Begin** page, click **Next**.
6. On the **Locate Virtual Hard Disk** page, click **Browse**.

7. Navigate to your virtual hard disk. Select your virtual hard disk. Click **Open**.
8. Click **Next**
9. On the **Choose Action** page, click **Shrink**. Click **Next**
10. On the **Expand Virtual Disk** page, enter the new virtual hard disk size. The size is specified in gigabytes.
11. Click **Next**.
12. On the Completing the **Edit Virtual Hard Disk Wizard** page, click **Finish**.

Linux Virtual Machines on Hyper-V

Hyper-V supports both emulated and Hyper-V-specific devices for Linux virtual machines. When running with emulated devices, no additional software is required to be installed on a Linux virtual machine. However emulated devices do not provide high performance and cannot leverage the rich virtual machine management infrastructure that the Hyper-V technology offers. In order to make full use of all benefits that Hyper-V provides, it is best to use Hyper-V-specific devices for Linux. The collection of drivers that are required to run Hyper-V-specific devices are known as Linux Integration Services (LIS).

For older Linux distributions, Microsoft provides an ISO containing installable LIS driver for Linux virtual machines. For newer Linux distributions, LIS is built in to the Linux operating system, and no separate download or installation is required. Note that the LIS collection from Microsoft cannot be installed into a Linux distribution that already has LIS built-in.

The goal of this content is to provide information that helps facilitate your Linux deployment on Hyper-V. Specific details include:

- Linux distributions that require the download and installation of LIS drivers.
- Linux distributions that contain pre-built LIS drivers.
- Feature distribution maps that indicate the features in major Linux distributions.
- Known issues and workarounds for each distribution.
- Feature description for each LIS feature.

In this section:

- CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map
- SUSE virtual machines feature distribution map
- Ubuntu virtual machines feature distribution map
- Feature Descriptions

CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map

The following feature distribution map indicates the features that are present in each version. The known issues and workarounds for each distribution are listed after the table.

Table legend

√ - Feature available

(*blank*) - Feature not available

Feature	**Windows Server operating system version**	**6.4**	**6.0 – 6.3**		**5.9**	**5.7 – 5.8**		**5.5 – 5.6**
Availability		Built in	LIS 3.5	LIS 3.4	Built in	LIS 3.5	LIS 3.4	LIS 3.5
Core	2012 R2, 2012, 2008 R2	√	√	√	√	√	√	√
Networking								
Jumbo frames	2012 R2, 2012, 2008 R2	√	√	√	√	√	√	√
VLAN tagging and trunking	2012 R2, 2012, 2008 R2	√	√	√	√	√	√	√
Live Migration	2012 R2, 2012, 2008 R2	√	√	√	√	√	√	√
Static IP Injection	2012 R2, 2012	√	√ Note 1	√		√ Note 1		√ Note 1
Storage								
VHDX resize	2012 R2		√			√		√
Virtual Fibre	2012 R2		√ Note			√ Note		√ Note

Channel			2			2		2
Live virtual machine backup	2012 R2	√ Note 5	√ Note 3, 4		√ Note 5	√ Note 3, 4		√ Note 3, 4
TRIM support	2012 R2							
Memory								
Configuration of MMIO gap	2012 R2	√	√	√	√	√	√	√
Dynamic Memory – Hot Add	2012 R2, 2012, 2008 R2							
Dynamic Memory – Ballooning	2012 R2, 2012	√ Note 6	√ Note 6			√ Note 6		√ Note 6
Video								
Hyper-V-specific video device	2012 R2, 2012, 2008 R2		√			√		√
Miscellaneous								
Key-Value Pair	2012 R2, 2012, 2008 R2	√ Note 7	√	√ Note 7		√		√
Non-Maskable Interrupt	2012 R2	√	√	√	√	√	√	√
PAE Kernel Support		√	√	√	√	√		√

Notes

1. Static IP injection may not work if Network Manager has been configured for a given synthetic network adapter on the virtual machine. To ensure smooth functioning of static IP injection please ensure that either Network Manager is either turned off completely or has been turned off for a specific network adapter through its ifcfg-ethX file.
2. While using virtual fibre channel devices, ensure that logical unit number 0 (LUN 0) has been populated. If LUN 0 has not been populated, a Linux virtual machine might not be able to mount fibre channel devices natively.

3. If there are open file handles during a live virtual machine backup operation, then in some corner cases, the backed-up VHDs might have to undergo a file system consistency check (fsck) on restore.
4. Live backup operations can fail silently if the virtual machine has an attached iSCSI device or a passthrough disk.
5. Live backup support for certain distributions such as RHEL/CentOS 5.9/6.4/6.5 is available through Linux Backup Essentials hosted at Hyper-V Backup Essentials.
6. LIS 3.5 only provides Ballooning support and does not provide Hot-Add support. In such a scenario, the dynamic memory feature can be used by setting the Startup memory parameter to a value which is equal to the Maximum memory parameter. This results in all the requisite memory being Hot-Added to the virtual machine at boot time and then later depending upon the memory requirements of the host, Hyper-V can freely allocate or deallocate memory from the guest using Ballooning. Please ensure that **Startup Memory** and **Minimum Memory** are not set below the recommended value for the distribution.
7. In Windows Server 2012 R2, the key/value pair (KVP) infrastructure might not function correctly without a Linux software update. Contact your distribution vendor to obtain the software update in case you see problems with this feature.

See Also

- SUSE virtual machines feature distribution map
- Ubuntu virtual machines feature distribution map
- Feature Descriptions

SUSE virtual machines feature distribution map

The following is a feature distribution map that indicates the features in each version. The known issues and workarounds for each distribution are listed after the table.

Table legend

√ - Feature available

(*blank*) - Feature not available

Feature	Windows Server operating system version	SLES 11 SP3	SLES 11 SP2	Open SUSE 12.3
Availability		Built-in	Built-in	Built-in
Core	2012 R2, 2012, 2008 R2	√	√	√
Networking				

Jumbo frames	2012 R2, 2012, 2008 R2	√	√	√
VLAN tagging and trunking	2012 R2, 2012, 2008 R2	√	√	√
Live migration	2012 R2, 2012, 2008 R2	√	√	√
Static IP Injection	2012 R2, 2012, 2008 R2	√ Note 1	√ Note 1	√ Note 1
Storage				
VHDX resize	2012 R2	√		
Virtual Fibre Channel	2012 R2			
Live virtual machine backup	2012 R2	√ Note 2, 3		
TRIM support	2012 R2			
Memory				
Configuration of MMIO gap	2012 R2	√	√	√
Dynamic Memory – Hot Add	2012 R2, 2012, 2008 R2	√ Note 4, 5		
Dynamic Memory – Ballooning	2012 R2, 2012, 2008 R2	√ Note 4, 5		√ Note 4, 5
Video				
Hyper-V-specific video device	2012 R2, 2012, 2008 R2	√		
Miscellaneous				
Key/value pair	2012 R2, 2012, 2008 R2	√ Note 6	√ Note 6	√ Note 6
Non-Maskable Interrupt	2012 R2	√	√	√
PAE Kernel Support	2012 R2, 2012, 2008 R2	√	√	√

Notes

1. Static IP injection may not work if **Network Manager** has been configured for a given Hyper-V-specific network adapter on the virtual machine. To ensure smooth functioning of static IP injection please ensure that Network Manager is turned off completely or has been turned off for a specific network adapter through its **ifcfg-ethX** file.
2. If there are open file handles during a live virtual machine backup operation, the backed-up virtual hard disks (VHDs) might have to undergo a file system consistency check (fsck) on restore.
3. Backup operations can fail silently if the virtual machine has an attached iSCSI device or direct-attached storage (also known as a pass-through disk).
4. Dynamic memory operations can fail if the guest operating system is running too low on memory. The following are some best practices:
    - Startup memory and minimal memory should be equal to or greater than the amount of memory that the distribution vendor recommends.
    - Applications that tend to consume the entire available memory on a system are limited to consuming up to 80 percent of available RAM.
5. If you are using Dynamic Memory on Windows Server 2008 R2 or Windows Server 2012 operating systems, specify **Startup memory**, **Minimum memory**, and **Maximum memory** parameters in multiples of 128 megabytes (MB). Failure to do so can lead to Hot-Add failures, and you may not see any memory increase in a guest operating system.
6. In Windows Server 2012 R2, the key/value pair infrastructure might not function correctly without a Linux software update. Contact your distribution vendor to obtain the software update in case you see problems with this feature.

See Also

- CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map
- Ubuntu virtual machines feature distribution map
- Feature Descriptions

Ubuntu virtual machines feature distribution map

The following feature distribution map indicates the features in each version. The known issues and workarounds for each distribution are listed after the table.

Table legend

√ - Feature available

(*blank*) - Feature not available

Feature	Windows Server	13.10	13.04	12.10	12.04

	operating system version				
Availability		Built-in	Built-in	Built-in	Built-in
Core	2012 R2, 2012, 2008 R2	√	√	√	√
Networking					
Jumbo frames	2012 R2, 2012, 2008 R2	√	√	√	√
VLAN tagging and trunking	2012 R2, 2012, 2008 R2	√	√	√	√
Live migration	2012 R2, 2012, 2008 R2	√	√	√	√
Static IP Injection	2012 R2, 2012, 2008 R2	√ Note 1	√ Note 1	√ Note 1	√ Note 1
Storage					
VHDX resize	2012 R2	√			
Virtual Fibre Channel	2012 R2	√ Note 2			
Live virtual machine backup	2012 R2	√ Note 3, 4, 5			
TRIM support	2012 R2	√			
Memory					
Configuration of MMIO gap	2012 R2	√	√	√	√
Dynamic Memory – Hot Add	2012 R2, 2012, 2008 R2	√ Note 6, 7			
Dynamic Memory – Ballooning	2012 R2, 2012, 2008 R2	√ Note 6, 7			
Video					
Hyper-V specific video device	2012 R2, 2012, 2008 R2	√			
Miscellaneous					
Key/value pair	2012 R2, 2012,	√ Note 5, 8	√ Note 5, 8	√ Note 5, 8	√ Note 5, 8

	2008 R2				
Non-Maskable Interrupt	2012 R2	√	√	√	√
PAE Kernel Support	2012 R2, 2012, 2008 R2	√	√	√	√

Notes

1. Static IP injection may not work if **Network Manager** has been configured for a given Hyper-V-specific network adapter on the virtual machine. To ensure smooth functioning of static IP injection please ensure that Network Manager is turned off completely or has been turned off for a specific network adapter through its **ifcfg-ethX** file.
2. While using virtual fiber channel devices, ensure that logical unit number 0 (LUN 0) has been populated. If LUN 0 has not been populated, a Linux virtual machine might not be able to mount fiber channel devices natively.
3. If there are open file handles during a live virtual machine backup operation, the backed-up VHDs might have to undergo a file system consistency check (fsck) on restore.
4. Backup operations can fail silently if the virtual machine has an attached iSCSI device or direct-attached storage (also known as a pass-through disk).
5. Certain daemons must be installed manually for Ubuntu distributions:
   - **VSS Snapshot daemon** – Without this daemon, live virtual machine backup does not function correctly. To install the VSS snapshot daemon, use the following command:

     ```
     # sudo apt-get install hv-kvp-daemon-init
     ```
   - **KVP daemon** – To install the KVP daemon, use the following command:

     ```
     # sudo apt-get install hv-kvp-daemon-init
     ```
6. Dynamic Memory operations can fail if the guest operating system is running too low on memory. The following are some best practices:
   - Startup memory and minimal memory should be equal to or greater than the amount of memory that the distribution vendor recommends.
   - Applications that tend to consume the entire available memory on a system are limited to consuming up to 80 percent of available RAM.
7. If you are using Dynamic Memory on Windows Server 2008 R2 or Windows Server 2012 operating systems, specify **Startup memory**, **Minimum memory**, and **Maximum memory** parameters in multiples of 128 megabytes (MB). Failure to do so can lead to Hot-Add failures, and you might not see any memory increase on a guest operating system.
8. In Windows Server 2012 R2, the key/value pair infrastructure might not function correctly without a Linux software update. Contact your distribution vendor to obtain the software update in case you see problems with this feature.

See Also

- CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map
- SUSE virtual machines feature distribution map
- Feature Descriptions

Feature Descriptions

The following features are described in this section:

- Core
- Networking
- Storage
- Memory
- Video
- Miscellaneous

Core

Feature	Description
Integrated shutdown	With this feature, an administrator can shut down virtual machines from the Hyper-V Manager.
Time synchronization	This feature ensures that the maintained time inside a virtual machine is kept synchronized with the maintained time on the host.
Multiprocessing support	With this feature, a virtual machine can use multiple processors on the host by configuring multiple virtual CPUs.
Heartbeat	With this feature, the host to can track the state of the virtual machine.
Integrated mouse support	With this feature, you can use a mouse on a virtual machine's desktop and also use the mouse seamlessly between the Windows Server desktop and the Hyper-V console for the virtual machine.
Hyper-V-specific Storage device	This feature grants high-performance access to storage devices that are attached to a virtual machine.

Hyper-v specific Network device	This feature grants high-performance access to network adapters that are attached to a virtual machine.

Networking

Feature	Description
Jumbo frames	With this feature, an administrator can increase the size of network frames beyond 1500 bytes, which leads to a significant increase in network performance.
VLAN tagging and trunking	With this feature, you can configure virtual LAN (VLAN) traffic for virtual machines.
Live Migration	With this feature, you can migrate a virtual machine from one host to another host.
Static IP Injection	With this feature, you can replicate the static IP address of a virtual machine after it has been failed over to its replica on a different host. Such IP replication ensures that network workloads continue to work seamlessly after a failover event.

Storage

Feature	Description
VHDX resize	With this feature, an administrator can resize a fixed-size .vhdx file that is attached to a virtual machine.
Virtual Fibre Channel	With this feature, virtual machines can recognize a fiber channel device and mount it natively.

Live virtual machine backup	This feature facilitates zero down time backup of live virtual machines. Note that the backup operation does not succeed if the virtual machine has virtual hard disks (VHDs) that are hosted on remote storage such as a Server Message Block (SMB) share or an iSCSI volume. Additionally, ensure that the backup target does not reside on the same volume as the volume that you back up.
TRIM support	TRIM hints notify the drive that certain sectors that were previously allocated are no longer required by the app and can be purged. This process is typically used when an app makes large space allocations via a file and then self-manages the allocations to the file, for example, to virtual hard disk files.

Memory

Feature	Description
Configuration of MMIO gap	With this feature, appliance manufacturers can configure the location of the Memory Mapped I/O (MMIO) gap. The MMIO gap is typically used to divide the available physical memory between an appliance's Just Enough Operating Systems (JeOS) and the actual software infrastructure that powers the appliance.
Dynamic Memory – Hot Add	With this feature, you can dynamically increase the amount of memory that is available to a live virtual machine. To provision a virtual machine with dynamic memory, the administrator must specify three parameters: Startup memory, Minimum memory, and Maximum memory.

	When the virtual machine is first started, the available memory is equal to the Startup memory. As the memory requirements within the virtual machine increase due to application workloads, Hyper-V dynamically provides more memory via its Hot-Add mechanism. The maximum amount of memory that can be made available to the virtual machine is capped by the value of the **Maximum memory** parameter.
Dynamic Memory – Ballooning	This feature dynamically provisions and deprovisions the amount of memory that is allocated to a live virtual machine. If a virtual machine is configured with dynamic memory, Hyper-V continuously monitors memory requirements within the virtual machine. If the required amount of memory falls below the amount of memory that is allocated, then Hyper-V can automatically deprovision memory from the virtual machine. Similarly, if memory requirements increase, Hyper-V dynamically provisions more memory to the virtual machine. Note that if the amount of allocated memory falls to the **Minimum memory** parameter that is specified for the virtual machine, Hyper-V abstains from further deprovisioning. Also note that this feature is different from Hot Add. Hot Add increases the amount of memory **available** (or visible) to a virtual machine. Ballooning adjusts the actual amount of memory **allocated** to a virtual machine. The allocated memory can take up any value between the specified Minimum memory and the total available memory.

Video

Feature	Description
Hyper-V-specific video device	This feature provides high-performance graphics and superior resolution for virtual machines.

Miscellaneous

Feature	Description
KVP (Key-Value pair) exchange	This feature provides a key/value pair (KVP) exchange service for virtual machines. Typically administrators use the KVP mechanism to perform read and write custom data operations on a virtual machine.
Non-Maskable Interrupt	With this feature, an administrator can issue Non-Maskable Interrupts (NMI) to a virtual machine. NMIs are useful in obtaining crash dumps of operating systems that have become unresponsive due to application bugs. These crash dumps can be diagnosed after you restart.
PAE Kernel Support	Physical Address Extension (PAE) technology allows a 32-bit kernel to access a physical address space that is larger than 4GB.Older Linux distributions such as RHEL 5.x used to ship a separate kernel that was PAE enabled. Newer distributions such as RHEL 6.x have prebuilt PAE support.

See Also

- CentOS, Oracle Linux, Red Hat Enterprise Linux virtual machines feature distribution map
- SUSE virtual machines feature distribution map
- Ubuntu virtual machines feature distribution map

Networking Overview

This collection contains detailed information about networking products and features for the IT professional to design, deploy, and maintain Windows Server.

Important
Networking resources for operating systems earlier than Windows Server 2012 are available at the following locations.

- Windows Server 2008 and Windows Server 2008 R2 Networking
- Windows Server 2003 Network Services

Networking in Windows Server 2012 and Windows Server 2012 R2

Following are the networking technologies and documentation for Windows Server 2012 and Windows Server 2012 R2.

What's New in Networking in Windows Server 2012 R2	Discover new networking technologies and new features for existing technologies in Windows Server 2012 R2.
What's New in Networking in Windows Server 2012	Discover new networking technologies and new features for existing technologies in Windows Server 2012.
Core Network Guide and Companion Guides Overview	Learn how to deploy a Windows Server network with the Core Network Guide and Companion Guides for Windows Server 2012.
802.1X Authenticated Wired Access Overview	Windows Server® 2012 and Windows® 8 provide features that you can use to deploy IEEE 802.1X authenticated wired service for Ethernet network clients. In combination with the 802.1X-capable Ethernet switches and other Windows Server® 2012 services that you deploy on your network, you can control who can access your network.
802.1X Authenticated Wireless Access Overview	Windows Server® 2012 and Windows® 8 provides features that you can use to deploy IEEE 802.1X authenticated wireless service for IEEE 802.11 wireless network clients. In combination with the 802.1X-capable wireless access points and other Windows Server®

	2012 services that you deploy on your network, you can control who can access your wireless network.
BranchCache Overview	BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server 2012 and Windows 8 operating systems. To optimize WAN bandwidth when users access content on remote servers, BranchCache copies content from your main office or hosted cloud content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.
Data Center Bridging (DCB) Overview	Data Center Bridging is a suite of technologies that aim to enable converged network fabrics – where storage, data networking, cluster inter-process communication (IPC), and management traffic all share the same Ethernet network infrastructure. DCB can provide cost savings and simplified management for enterprise and hoster data centers.
Domain Name System (DNS) Overview	DNS is a hierarchical, distributed database that contains mappings of domain names to various types of data, such as IP addresses. DNS enables you to use friendly names, such as www.microsoft.com, to easily locate computers and other resources on a TCP/IP-based network. The DNS service in Windows Server® 2012 adds enhanced support for DNS Security Extensions (DNSSEC), including online signing and automated key management.
Dynamic Host Configuration Protocol (DHCP) Overview	DHCP is a client-server technology that enables DHCP servers to assign, or lease, IP addresses to computers and other devices that are enabled as DHCP clients. When you deploy DHCP servers on your network, you can automatically provide client computers and other IPv4- and IPv6-based network devices with valid IP addresses and the additional

	configuration parameters these clients and devices need. The DHCP Server service in Windows Server® 2012 includes support for policy based assignment and DHCP failover.
Extensible Authentication Protocol (EAP) for Network Access Overview	Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by allowing arbitrary authentication methods that use credential and information exchanges of arbitrary lengths. EAP was developed in response to demand for authentication methods that use security devices, such as smart cards, token cards, and crypto calculators. EAP provides an industry-standard architecture for supporting additional authentication methods within PPP.
High-Speed Networking	High-speed networking features include TCP Chimney Offload, Virtual Machine Chimney, Virtual Machine Queue, receive-side scaling (RSS), and Network Direct Memory Access (NetDMA). These features can improve network and operating system performance.
Hyper-V Network Virtualization Overview	Hyper-V Network Virtualization provides the concept of a Virtual Network that is independent of the underlying physical network. With this concept of Virtual Networks, which are composed of one or more Virtual Subnets, the exact physical location of an IP subnet is decoupled from the virtual network topology. As a result, customers can easily move their subnets to the cloud while preserving their existing IP addresses and topology in the cloud so that existing services continue to work unaware of the physical location of the subnets. That is, Hyper-V Network Virtualization enables a seamless hybrid cloud.
Hyper-V Virtual Switch Overview	The Virtual Switch is part of the Hyper-V infrastructure in Windows Server® 2012 that enables virtual machines to connect to each other and to the host computer's physical network. In Windows Server® 2012, the switch provides services that can be enhanced and

	extended by adding plug-ins that supply additional functionality.
IP Address Management (IPAM) Overview	IPAM is a built-in framework in Windows Server® 2012 for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. IPAM enables easier administration and monitoring of DHCP and DNS.
Internet Protocol version 6 (IPv6)	IPv6 is the replacement for IPv4, the Internet layer protocol of the TCP/IP protocol stack in prevalent use around the world today. IPv6 is fully supported in both Windows Server® 2012 and Windows® 8. IPv6 solves many of the problems and shortcomings of IPv4, providing an Internet layer protocol that can scale to the future needs of devices that will connect to the Internet. IPv6 uses 128-bit addresses to enable 3.4×10^{38} possible addresses, more than enough to handle today's needs and those of the foreseeable future.
Low Latency Workloads Technologies	Discover this group of technologies that are designed for or which were improved in Windows Server 2012 to address low latency computing scenarios.
Network Load Balancing Overview	Network Load Balancing in Windows Server 2012 distributes traffic across several servers by using the TCP/IP networking protocol. It enhances the availability and scalability of Internet server applications.
Network Policy and Access Services Overview	The Network Policy and Access Services server role in Windows Server 2012 provides the Network Policy Server (NPS), Health Registration Authority (HRA), and Host Credential Authorization Protocol (HCAP) role services. You can use this server role to deploy and configure Network Access Protection (NAP) and Remote Authentication Dial-In User Service (RADIUS) servers and proxies. RADIUS servers perform network access authorization and authentication for network connections through the network access

	servers that you deploy, such as 802.1X-capable wireless access points, 802.1X-capable Ethernet switches, and virtual private network (VPN) servers.
Network Shell (Netsh) Overview	Netsh is a command-line scripting utility that allows you to display or modify the network configuration of a computer. You can use Netsh commands to configure the local computer and remote computers. Netsh also provides a scripting feature that allows you to run a group of commands in batch mode against a specified computer. With netsh, you can save a configuration script in a text file for archival purposes or to help you configure other computers.
Network Subsystem Performance Tuning	This topic provides information about choosing the right network adapter for your server workload, network related performance counters, and performance tuning network adapters and related networking technologies, such as Receive Side Scaling (RSS), Receive Side Coalescing (RSC), and others.
NIC Teaming Overview	Network Interface Card (NIC) Teaming, also called load balancing and failover (LBFO), allows multiple network adapters on a computer that is running Windows Server® 2012 to be placed into a team to help prevent connectivity loss in the event of a network component failure and to enable bandwidth aggregation to help improve network performance.
Quality of Service (QoS) Overview	Policy-based QoS in Windows Server® 2012 and Windows® 8 enables you to define the priority of traffic, negotiate finer service levels with bandwidth providers, and control bandwidth costs.
Remote Access (DirectAccess, Routing and Remote Access) Overview	Remote Access in Windows Server® 2012 enables you to support remote network users by providing them with connectivity to your organization’s private network from anywhere.
Telephony Server and the Telephony API	TAPI integrates telecommunications with the

Overview	computer. TAPI supports both traditional and IP telephony to provide voice, data, and video communication.
TCP/IP	Transmission Control Protocol/Internet Protocol (TCP/IP) is an industry standard suite of protocols that is designed for large networks consisting of network segments that are connected by routers. TCP/IP is the core protocol suite that is used on the Internet. In Windows Server 2012, TCP/IP – including both Internet Protocol version 4 (IPv4) and IPv6 – is unchanged from TCP/IP in Windows Server 2008 R2. For more information, see TCP/IP in the Windows Server 2008 R2 Technical Library.
Virtual Receive-side Scaling in Windows Server 2012 R2	New in Windows Server 2012 R2, virtual Receive-side Scaling (vRSS) enables network adapters to distribute network processing load across multiple virtual processors in multi-core virtual machines (VMs).
Windows Firewall with Advanced Security Overview	Windows Firewall with Advanced Security is an advanced interface for IT professionals to use to configure both Windows Firewall and Internet Protocol security (IPsec) settings for the computers that are running Windows Server 2012 and Windows 8 on their networks. Windows Firewall with Advanced Security in Windows Server 2012 adds support for IKEv2, Windows Store apps, and Windows PowerShell.
Windows Internet Name Service (WINS) Overview	WINS provides a distributed database that maps network basic input/output system (NetBIOS) names to IP addresses and was designed to solve the problems arising from NetBIOS name resolution in routed environments.
Windows PowerShell Cmdlets for Networking	Use this list of links to connect to command references for Windows PowerShell cmdlets for networking technologies.
Windows Server Gateway	Windows Server Gateway is new in Windows

	Server 2012 R2 and is not included in Windows Server 2012. Windows Server Gateway is a virtual machine (VM)-based virtual software router that allows Cloud Service Providers (CSPs) and Enterprises to enable datacenter and cloud network traffic routing between virtual and physical networks, including the Internet.

What's New in Networking in Windows Server 2012 R2

New and improved technologies for Windows Server Networking

The following networking technologies are new or improved in Windows Server® 2012 R2.

802.1X Authenticated Wired Access

802.1X Authenticated Wired Access in Windows 8.1 and Windows Server 2012 R2 provides new features and capabilities over previous versions.

For more information, see What's New in 802.1X Authenticated Wired Access in Windows Server 2012 R2.

802.1X Authenticated Wireless Access

802.1X Authenticated Wireless Access in Windows 8.1 and Windows Server 2012 R2 provides new features and capabilities over previous versions.

For more information, see What's New in 802.1X Authenticated Wireless Access in Windows Server 2012 R2.

Domain Name System

Domain Name System (DNS) in Windows Server 2012 R2 provides new features and capabilities over previous versions.

For more information, see What's New in DNS Server in Windows Server 2012 R2.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) in Windows Server 2012 R2 provides new features and capabilities over previous versions.

For more information, see What's New in DHCP in Windows Server 2012 R2.

Hyper-V Network Virtualization

Hyper-V Network Virtualization (HNV) has many important updates that enable hybrid cloud and private cloud solutions.

For more information, see What's New in Hyper-V Network Virtualization in Windows Server 2012 R2.

Hyper-V Virtual Switch

Hyper-V Virtual Switch provides new features and capabilities over previous versions.

For more information, see What's New in Hyper-V Virtual Switch in Windows Server 2012 R2.

Internet Protocol Address Management

IP Address Management (IPAM) is a feature that was first introduced in Windows Server 2012 that provides highly customizable administrative and monitoring capabilities for the IP address infrastructure on a corporate network. IPAM in Windows Server 2012 R2 includes many enhancements.

For more information, see What's New in IPAM in Windows Server 2012 R2.

Remote Access (DirectAccess, Routing and Remote Access)

Remote Access provides new features and capabilities over previous versions.

For more information, see What's New in Remote Access in Windows Server 2012 R2.

Virtual Receive-side Scaling

New in Windows Server 2012 R2, virtual Receive-side Scaling (vRSS) enables network adapters to distribute network processing load across multiple virtual processors in multi-core virtual machines (VMs).

For more information, see Virtual Receive-side Scaling in Windows Server 2012 R2.

Windows Server Gateway

New in Windows Server 2012 R2, Windows Server Gateway is a virtual machine (VM)-based software router that allows Cloud Service Providers (CSPs) and Enterprises to enable datacenter and cloud network traffic routing between virtual and physical networks, including the Internet.

Windows Server Gateway routes network traffic between the physical network and VM network resources, regardless of where the resources are located. You can use Windows Server Gateway to route network traffic between physical and virtual networks at the same physical location or at many different physical locations, providing network traffic flow in private and hybrid cloud scenarios.

For more information, see Windows Server Gateway.

What's New in Networking in Windows Server 2012

New and improved technologies for Windows Server Networking

The following networking technologies are new or improved in Windows Server® 2012.

802.1X Authenticated Wired and Wireless Access

For Windows Server 2012 and Windows® 8, 802.1X Authenticated Wired and Wireless Access provide the Extensible Authentication Protocol (EAP) authentication type Tunneled Transport Layer Security, or EAP-TTLS. EAP-TTLS is new in Windows Server 2012 and Windows 8 and is not available in other versions of Windows Server.

EAP-TTLS is a standards-based EAP tunneling method that supports mutual authentication. EAP-TTLS provides a secure tunnel for client authentication using EAP methods and other legacy protocols. EAP-TTLS also provides you with the ability to configure EAP-TTLS on client computers for network access solutions in which non-Microsoft Remote Authentication Dial In User Service (RADIUS) servers that support EAP-TTLS are used for authentication.

For more information about 802.1X Authenticated Wired Access, see 802.1X Authenticated Wired Access Overview.

For more information about 802.1X Authenticated Wireless Access, see 802.1X Authenticated Wireless Access Overview.

BranchCache

BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server 2012 and Windows 8 operating systems. To optimize WAN bandwidth when users access content on remote servers, BranchCache copies content from your main office or hosted cloud content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.

BranchCache improvements in Windows Server 2012 include automatic BranchCache client computer configuration, deep integration with the Windows file server, the ability to cache small updates to files, resulting in more bandwidth savings, security improvements, simplification of hosted cache server deployment, and much more. For more information, see What's New in BranchCache.

Data Center Bridging (DCB)

DCB is introduced as a new technology in Windows Server 2012. DCB is a suite of Institute of Electrical and Electronics Engineers (IEEE) standards that enable Converged Fabrics in the data center, where storage, data networking, cluster IPC and management traffic all share the same Ethernet network infrastructure. DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control.

Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support Internet Small Computer System Interface (iSCSI), Remote Direct Memory Access (RDMA) over Converged Ethernet, or Fiber Channel over Ethernet (FCoE). Priority-based flow control is essential if the upper layer protocol, such as Fiber Channel, assumes a lossless underlying transport. For more information, see Data Center Bridging (DCB) Overview.

Domain Name System (DNS)

The Domain Name System (DNS) is used in TCP/IP networks for naming computers and network services. Using DNS, computers are able to locate devices and services on a network through user-friendly names.

DNS in Windows includes the DNS Client service and the DNS Server service. Windows Server 2012 and Windows 8 include several enhancements to DNS. For more information see What's New in DNS.

Dynamic Host Configuration Protocol (DHCP)

DHCP is an Internet Engineering Task Force (IETF) standard designed to reduce the administration burden and complexity of configuring hosts on a TCP/IP-based network, such as a private intranet. Using the DHCP Server service, the process of configuring TCP/IP on DHCP clients is automatic. Windows Server 2012 provides several enhancements to the DHCP Server service, including DHCP failover and DHCP policy based assignment. For more information see What's New in DHCP in Windows Server 2012.

Hyper-V Network Virtualization

With the success of virtualized datacenters, IT organizations and hosting providers (providers who offer colocation or physical server rentals) have begun offering more flexible virtualized infrastructures that make it easier to offer on-demand server instances to their customers. This new class of service is referred to as Infrastructure as a Service (IaaS). Windows Server 2012 provides all the required platform capabilities to enable enterprise customers to build private clouds and transition to an IT as a service operational model. Windows Server 2012 also enables hosters to build public clouds and offer IaaS solutions to their customers. When combined with System Center to manage Hyper-V Network Virtualization policy, Microsoft provides a powerful cloud solution.

Windows Server 2012 Hyper-V Network Virtualization provides policy-based, software-controlled network virtualization that reduces the management overhead that is faced by enterprises when they expand dedicated IaaS clouds, and it provides cloud hosters better flexibility and scalability for managing virtual machines to achieve higher resource utilization.

For more information, see Hyper-V Network Virtualization Overview.

Hyper-V Virtual Switch

The Hyper-V virtual switch is a virtual layer-2 network switch that provides programmatically managed and extensible capabilities to connect virtual machines to the physical network. Hyper-V virtual switch provides policy enforcement for security, isolation, and service levels. Additionally, the Hyper-V virtual switch in Windows Server® 2012 introduces a number of new and enhanced capabilities for tenant isolation, traffic shaping, protection against malicious virtual machines, and simplified troubleshooting. For more information, see What's New in Hyper-V Virtual Switch in Windows Server 2012.

IP Address Management (IPAM)

IPAM in Windows Server 2012 is an entirely new built-in framework for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. IPAM provides for administration and monitoring of servers running Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS). For more information, see IP Address Management (IPAM) Overview.

Low Latency Workloads Technologies

Low latency computing environments typically contain applications that require very fast inter-process communication (IPC) and inter-computer communications, a high degree of predictability regarding latency and transaction response times, and the ability to handle very high message rates.

In Windows Server 2012, low latency workloads technologies include Data Center Bridging (DCB), Kernel Mode Remote Direct Memory Access (kRDMA), NIC Teaming, NetworkDirect, TCP Loopback Optimization, and more. For more information, see Low Latency Workloads Technologies.

Network Load Balancing (NLB)

The Network Load Balancing (NLB) feature distributes traffic across several servers by using the TCP/IP networking protocol. By combining two or more computers that are running applications in Windows Server 2012 into a single virtual cluster, NLB provides reliability and performance for web servers and other mission-critical servers.

The servers in an NLB cluster are called *hosts*, and each host runs a separate copy of the server applications. NLB distributes incoming client requests across the hosts in the cluster. You can configure the load that is to be handled by each host, and you can also add hosts dynamically to

the cluster to handle increased load. NLB can also direct all traffic to a designated single host, which is called the *default host*.

Note
The NLB functionality in Windows Server 2012 is generally the same as in Windows Server® 2008 R2, however some task details are changed in Windows Server 2012. For information on new ways to do tasks, see Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012.

For more information on NLB, see **Network Load Balancing Overview**.

Network Policy and Access Services

Network Policy and Access Services in Windows Server 2012 includes the Network Policy Server (NPS), Health Registration Authority (HRA), and Host Credential Authorization Protocol (HCAP) role services. You can use the Network Policy and Access Services server role to deploy NPS as a Remote Authentication Dial-In User Service (RADIUS) server and RADIUS proxy that performs authentication and authorization for connection requests from RADIUS clients, such as 802.1X capable Ethernet switches and wireless access points. In addition you can configure Network Access Protection (NAP) health policies in NPS.

You can now use Windows PowerShell to automate the installation of the Network Policy and Access Services server role. You can also deploy and configure some aspects of NPS by using Windows PowerShell.

For more information, see Network Policy and Access Services Overview.

NIC Teaming

Network Interface Card (NIC) Teaming is introduced as a new technology in Windows Server 2012. NIC Teaming, also known as load balancing and failover (LBFO), allows multiple network adapters on a computer to be placed into a team for the following purposes:

- Bandwidth aggregation
- Traffic failover to prevent connectivity loss in the event of a network component failure

For more information, see NIC Teaming Overview.

Quality of Service (QoS)

QoS is a set of technologies for managing network traffic in a cost effective manner, to enhance user experiences in enterprise environments, as also in home and small offices. QoS technologies allow you to measure bandwidth, detect changing network conditions (such as congestion or availability of bandwidth), and prioritize or throttle traffic. For example, you can use QoS to prioritize traffic for latency-sensitive applications (such as voice or video), and to control the impact of latency-insensitive traffic (such as bulk data transfers).

In Windows Server 2012, QoS includes new bandwidth management features that enable cloud hosting providers and enterprises to provide services that deliver predictable network

performance to virtual machines on a server running Hyper-V. In hosted environments, Hyper-V QoS enables hosting providers to guarantee specific performance levels based on service level agreements (SLAs). Hyper-V QoS helps ensure that no customer is impacted or compromised by other customers on their shared infrastructure, which includes computing, storage, and network resources.

Some additional new features for QoS include providing you with the ability to enforce minimum bandwidth for a traffic flow, to configure rate limiting on a Hyper-V virtual switch per switch port by using either Windows PowerShell or Windows Management Instrumentation, and to enforce QoS policies on Single Root I/O Virtualization (SR-IOV)-capable network adapters that support bandwidth reservation per Virtual Port.

For more information, see Quality of Service (QoS) Overview.

Remote Access

Remote Access in Windows Server 2012 combines two networking services into one unified server role:

DirectAccess and RRAS unified server role

Windows Server® 2008 R2 introduced DirectAccess, a new remote access feature that allows connectivity to corporate network resources without the need for traditional Virtual Private Network (VPN) connections. DirectAccess provides support only for domain-joined Windows 7 Enterprise and Ultimate edition clients. The Windows Routing and Remote Access Server (RRAS) provides traditional VPN connectivity for legacy clients, non-domain joined clients, and third party VPN clients. RRAS also provides site-to-site connections between servers. RRAS in Windows Server 2008 R2 cannot coexist on the same edge server with DirectAccess, and must be deployed and managed separately from DirectAccess.

Windows Server 2012 combines the DirectAccess feature and the RRAS role service into a new unified server role. This new Remote Access server role allows for centralized administration, configuration, and monitoring of both DirectAccess and VPN-based remote access services. Additionally, Windows Server 2012 DirectAccess provides multiple updates and improvements to address deployment blockers and provide simplified management. For more information, see Remote Access (DirectAccess, Routing and Remote Access) Overview.

Windows Firewall with Advanced Security

The following are some of the new capabilities in Windows Firewall with Advanced Security:

- IKEv2 for IPsec transport mode is added to Windows Server 2012, which provides interoperability for Windows with other operating systems using IKEv2 for end-to-end security. This adds support for Suite B (RFC 4869) requirements.
- Administrators can custom configure Windows Firewall to fine tune network access if they desire more control of their Windows Store applications.
- Windows PowerShell has extensive cmdlets to allow Windows Firewall configuration and management.

For more information, see **Windows Firewall with Advanced Security Overview**.

Core Network Guide and Companion Guides Overview

This topic contains the following sections.

- Introduction to the Windows Server Core Network
- Windows Server 2012 Core Network Guide
- Windows Server 2012 Core Network Companion Guides
  - Core Network Companion Guide: Server Certificate Deployment
  - Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access
  - Core Network Companion Guide: Computer and User Certificate Deployment
  - Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode
  - Core Network Companion Guide: Group Policy Deployment

Introduction to the Windows Server Core Network

A core network is a collection of network hardware, devices, and software that provides the fundamental services for your organization's information technology (IT) needs.

A Windows Server core network provides you with many benefits, including the following.

- Core protocols for network connectivity between computers and other Transmission Control Protocol/Internet Protocol (TCP/IP) compatible devices. TCP/IP is a suite of standard protocols for connecting computers and building networks. TCP/IP is network protocol software provided with Microsoft® Windows® operating systems that implements and supports the TCP/IP protocol suite.
- Dynamic Host Configuration Protocol (DHCP) server automatic IP addressing. Manual configuration of IP addresses on all computers on your network is time-consuming and less flexible than dynamically providing computers and other devices with IP address leases from a DHCP server.
- Domain Name System (DNS) name resolution service. DNS allows users, computers, applications, and services to find the IP addresses of computers and devices on the network by using the Fully Qualified Domain Name of the computer or device.
- A forest, which is one or more Active Directory domains that share the same class and attribute definitions (schema), site and replication information (configuration), and forest-wide search capabilities (global catalog).
- A forest root domain, which is the first domain created in a new forest. The Enterprise Admins and Schema Admins groups, which are forest-wide administrative groups, are located in the forest root domain. In addition, a forest root domain, as with other domains, is a collection of computer, user, and group objects that are defined by the administrator in Active Directory

Domain Services (AD DS). These objects share a common directory database and security policies. They can also share security relationships with other domains if you add domains as your organization grows. The directory service also stores directory data and allows authorized computers, applications, and users to access the data.

- A user and computer account database. The directory service provides a centralized user accounts database that allows you to create user and computer accounts for people and computers that are authorized to connect to your network and access network resources, such as applications, databases, shared files and folders, and printers.

A core network also allows you to scale your network as your organization grows and IT requirements change. For example, with a core network you can add domains, IP subnets, remote access services, wireless services, and other features and server roles provided by Windows Server® 2012.

Windows Server 2012 Core Network Guide

The Windows Server 2012 Core Network Guide provides instructions on how to plan and deploy the core components required for a fully functioning network and a new Active Directory® domain in a new forest. Using this guide, you can deploy computers configured with the following Windows server components:

- The Active Directory Domain Services (AD DS) server role
- The Domain Name System (DNS) server role
- The Dynamic Host Configuration Protocol (DHCP) server role
- The Network Policy Server (NPS) role service of the Network Policy and Access Services server role
- The Web Server (IIS) server role
- Transmission Control Protocol/Internet Protocol version 4 (TCP/IP) connections on individual servers

This guide also serves as a foundation for companion guides that show you how to deploy additional technologies in Windows Server 2012.

This guide is available at the following locations.

- The Windows Server 2012 Core Network Guide in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/p/?LinkId=255199.
- The Windows Server 2012 Core Network Guide in HTML format in the Technical Library.

Windows Server 2012 Core Network Companion Guides

While the Core Network Guide provides instructions on how to deploy a new Active Directory forest with a new root domain and the supporting networking infrastructure, Companion Guides provide you with the ability to add features to your network. Each companion guide allows you to accomplish a specific goal after you have deployed your core network. In some cases, there are

multiple companion guides that, when deployed together and in the correct order, allow you to accomplish very complex goals in a measured, cost-effective, reasonable manner.

If you deployed your Active Directory domain and core network before encountering the Core Network Guide, you can still use the Companion Guides to add features to your network. Simply use the Core Network Guide as a list of prerequisites, and know that to deploy additional features with the Companion Guides, your network must meet the prerequisites that are provided by the Core Network Guide.

Core Network Companion Guide: Server Certificate Deployment

This guide explains how to build upon the foundation network by deploying server certificates for computers that are running Network Policy Server (NPS), Routing and Remote Access Service (RRAS), or both. Server certificates are required when you deploy certificate-based authentication methods with Extensible Authentication Protocol (EAP) and Protected EAP (PEAP) for network access authentication. Deploying server certificates with Active Directory Certificate Services (AD CS) for EAP and PEAP certificate-based authentication methods provides the following benefits:

- Binding the identity of the NPS or RRAS server to a private key
- A cost-efficient and secure method for automatically enrolling certificates to domain member NPS and RRAS servers
- An efficient method for managing certificates and certification authorities
- Security provided by certificate-based authentication
- The ability to expand the use of certificates for additional purposes

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Server Certificate Deployment in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/p/?LinkID=255200.
- The Windows Server 2012 Core Network Companion Guide: Server Certificate Deployment in HTML format in the Technical Library.

Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access

This guide explains how to build upon a core network by providing instructions about how to deploy Institute of Electrical and Electronics Engineers (IEEE) 802.1X-authenticated IEEE 802.11 wireless access using Protected Extensible Authentication Protocol – Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAP v2).

This guide provides instructions about how to deploy a WiFi access infrastructure that uses the following components:

- One or more 802.1X-capable 802.11 wireless access points (APs)
- AD DS Users and Computers
- Group Policy Management.

- One or more Network Policy Server (NPS) servers
- Server certificates for computers running NPS.
- Wireless client computers running Windows® 8, Windows® 7, Windows Vista® or Windows XP with Service Pack

This guide is available at the following locations.

1. The Windows Server 2012 Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access in Word format in the Microsoft Download Center at http://www.microsoft.com/en-us/download/details.aspx?id=35555.
2. The Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access in HTML format in the Technical Library.

Core Network Companion Guide: Computer and User Certificate Deployment

This companion guide explains how to build on the core network by providing instructions for deploying client computer and user certificates with Active Directory Certificate Services (AD CS).

Certificates are used for network access authentication because they provide strong security for authenticating users and computers and they eliminate the need for less secure password-based authentication methods.

When you deploy Extensible Authentication Protocol with Transport Layer Security (EAP-TLS) or Protected EAP with TLS (PEAP-TLS), certificates are required for the authentication of servers, clients, and users during network connection attempts through network access servers such as 802.1X-capable switches and wireless access points and virtual private network (VPN) servers.

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Computer and User Certificates Deployment in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/?LinkId=255201.
- The Windows Server 2012 Core Network Companion Guide: Computer and User Certificates Deployment in HTML format in the Technical Library at http://technet.microsoft.com/library/jj129712.

Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode

This guide builds upon the core network by showing you how to deploy BranchCache in hosted cache mode in your branch offices. This guide shows how to use new BranchCache features for Windows Server 2012 and Windows 8 to deploy your hosted cache server by Service Connection Point, and how to prehash content on your Web and file content servers so that you can preload the content on your hosted cache servers.

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode in Word format in the Microsoft Download Center.

- The Windows Server 2012 Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode in HTML format in the Windows Server 2012 Technical Library.

Core Network Companion Guide: Group Policy Deployment

This guide explains how to build on the core network by providing instructions for deploying Group Policy Objects (GPOs) by using membership groups instead of the organizational units (OUs) that form the hierarchy of an Active Directory domain.

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Group Policy Deployment in Word format in the Microsoft Download Center.
- The Windows Server 2012 Core Network Companion Guide: Group Policy Deployment in HTML format in the Windows Server 2012 Technical Library.

Core Network Guide

This guide contains the following sections.

1. About this guide
2. Core Network Overview
3. Core Network Planning
4. Core Network Deployment
5. Additional Technical Resources
6. Appendices A through E

About this guide

This guide is available at the following locations.

- The Windows Server 2012 Core Network Guide in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/p/?LinkId=255199.
- The Windows Server 2012 Core Network Guide in HTML format in the Windows Server 2012 Technical Library.

This guide is designed for network and system administrators who are installing a new network or who want to create a domain-based network to replace a network that consists of workgroups. The deployment scenario provided in this guide is particularly useful if you foresee the need to add more services and features to your network in the future.

It is recommended that you review design and deployment guides for each of the technologies used in this deployment scenario to assist you in determining whether this guide provides the services and configuration that you need.

A core network is a collection of network hardware, devices, and software that provides the fundamental services for your organization's information technology (IT) needs.

A Windows Server core network provides you with many benefits, including the following.

- Core protocols for network connectivity between computers and other Transmission Control Protocol/Internet Protocol (TCP/IP) compatible devices. TCP/IP is a suite of standard protocols for connecting computers and building networks. TCP/IP is network protocol software provided with Microsoft® Windows® operating systems that implements and supports the TCP/IP protocol suite.
- Dynamic Host Configuration Protocol (DHCP) automatic IP address assignment to computers and other devices that are configured as DHCP clients. Manual configuration of IP addresses on all computers on your network is time-consuming and less flexible than dynamically providing computers and other devices with IP address configurations using a DHCP server.
- Domain Name System (DNS) name resolution service. DNS allows users, computers, applications, and services to find the IP addresses of computers and devices on the network by using the Fully Qualified Domain Name of the computer or device.
- A forest, which is one or more Active Directory domains that share the same class and attribute definitions (schema), site and replication information (configuration), and forest-wide search capabilities (global catalog).
- A forest root domain, which is the first domain created in a new forest. The Enterprise Admins and Schema Admins groups, which are forest-wide administrative groups, are located in the forest root domain. In addition, a forest root domain, as with other domains, is a collection of computer, user, and group objects that are defined by the administrator in Active Directory Domain Services (AD DS). These objects share a common directory database and security policies. They can also share security relationships with other domains if you add domains as your organization grows. The directory service also stores directory data and allows authorized computers, applications, and users to access the data.
- A user and computer account database. The directory service provides a centralized user accounts database that allows you to create user and computer accounts for people and computers that are authorized to connect to your network and access network resources, such as applications, databases, shared files and folders, and printers.

A core network also allows you to scale your network as your organization grows and IT requirements change. For example, with a core network you can add domains, IP subnets, remote access services, wireless services, and other features and server roles provided by Windows Server® 2012.

Network hardware requirements

To successfully deploy a core network, you must deploy network hardware, including the following:

- Ethernet, Fast Ethernet, or Gigabyte Ethernet cabling
- A hub, Layer 2 or 3 switch, router, or other device that performs the function of relaying network traffic between computers and devices.
- Computers that meet the minimum hardware requirements for their respective client and server operating systems.

What this guide does not provide

This guide does not provide instructions for deploying the following:

- Network hardware, such as cabling, routers, switches, and hubs
- Additional network resources, such as printers and file servers
- Internet connectivity
- Remote access
- Wireless access
- Client computer deployment

Note
Client computers running Windows® 8, Windows® 7, and Windows Vista® are configured by default to receive IP address leases from the DHCP server. Therefore, no additional DHCP or Internet Protocol version 4 (IPv4) configuration of client computers is required.

Technology Overviews

The following sections provide brief overviews of the required technologies that are deployed to create a core network.

Active Directory Domain Services

A directory is a hierarchical structure that stores information about objects on the network, such as users and computers. A directory service, such as AD DS, provides the methods for storing directory data and making this data available to network users and administrators. For example, AD DS stores information about user accounts, including names, email addresses, passwords, and phone numbers, and enables other authorized users on the same network to access this information.

DNS

DNS is a name resolution protocol for TCP/IP networks, such as the Internet or an organization network. A DNS server hosts the information that enables client computers and services to resolve easily recognized, alphanumeric DNS names to the IP addresses that computers use to communicate with each other.

DHCP

DHCP is an IP standard for simplifying the management of host IP configuration. The DHCP standard provides for the use of DHCP servers as a way to manage dynamic allocation of IP addresses and other related configuration details for DHCP-enabled clients on your network.

DHCP allows you to use a DHCP server to dynamically assign an IP address to a computer or other device, such as a printer, on your local network. Every computer on a TCP/IP network must

have a unique IP address, because the IP address and its related subnet mask identify both the host computer and the subnet to which the computer is attached. By using DHCP, you can ensure that all computers that are configured as DHCP clients receive an IP address that is appropriate for their network location and subnet, and by using DHCP options, such as default gateway and DNS servers, you can automatically provide DHCP clients with the information that they need to function correctly on your network.

For TCP/IP-based networks, DHCP reduces the complexity and amount of administrative work involved in reconfiguring computers.

TCP/IP

TCP/IP in Windows Server 2012 is the following:

- Networking software based on industry-standard networking protocols.
- A routable enterprise networking protocol that supports the connection of your Windows-based computer to both local area network (LAN) and wide area network (WAN) environments.
- Core technologies and utilities for connecting your Windows-based computer with dissimilar systems for the purpose of sharing information.
- A foundation for gaining access to global Internet services, such as the World Wide Web and File Transfer Protocol (FTP) servers.
- A robust, scalable, cross-platform, client/server framework.

TCP/IP provides basic TCP/IP utilities that enable Windows-based computers to connect and share information with other Microsoft and non-Microsoft systems, including:

- Windows Server 2012
- Windows 8
- Windows Server 2008 R2
- Windows 7
- Windows Server 2008
- Windows Vista
- Windows Server 2003 operating systems
- Windows XP
- Internet hosts
- Apple Macintosh systems
- IBM mainframes
- UNIX systems
- Open VMS systems
- Network-ready printers, such as HP LaserJet series printers that use HP JetDirect cards

Core Network Overview

The following illustration shows the Windows Server Core Network topology.

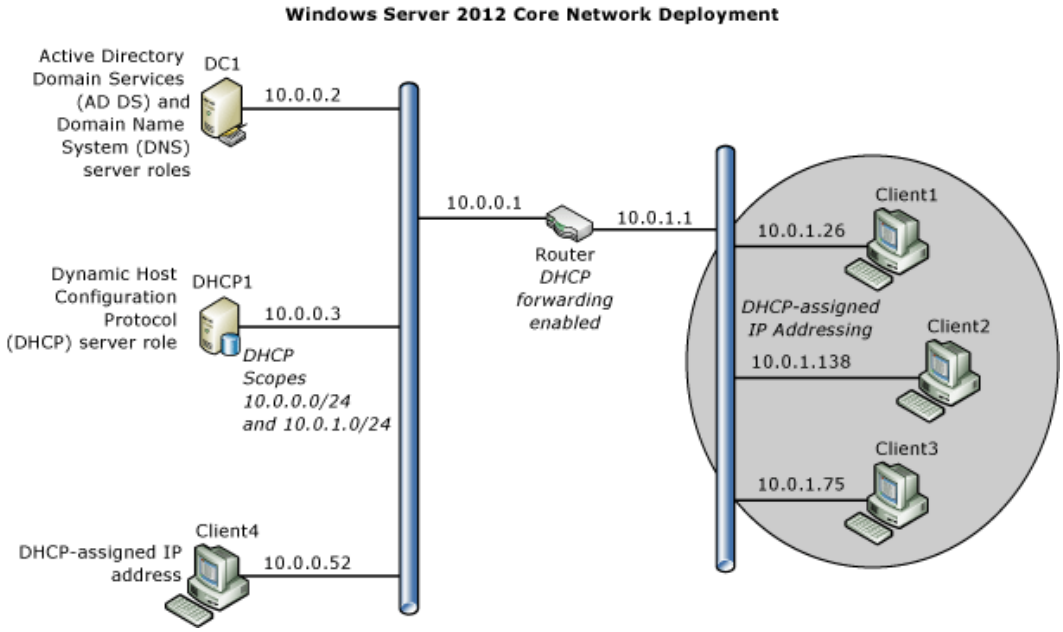

Note

This guide also includes instructions for adding optional Network Policy Server (NPS) and Web Server (IIS) servers to your network topology to provide the foundation for secure network access solutions, such as 802.1X wired and wireless deployments that you can implement using Core Network Companion guides. For more information, see Deploying optional features for network access authentication and Web services.

Core Network Components

Following are the components of a core network.

Router

This deployment guide provides instructions for deploying a core network with two subnets separated by a router that has DHCP forwarding enabled. You can, however, deploy a Layer 2 switch, a Layer 3 switch, or a hub, depending on your requirements and resources. If you deploy a switch, the switch must be capable of DHCP forwarding or you must place a DHCP server on each subnet. If you deploy a hub, you are deploying a single subnet and do not need DHCP forwarding or a second scope on your DHCP server.

Static TCP/IP configurations

The servers in this deployment are configured with static IPv4 addresses. Client computers are configured by default to receive IP address leases from the DHCP server.

Active Directory Domain Services global catalog and DNS server DC1

Both Active Directory Domain Services (AD DS) and Domain Name System (DNS) are installed on this server, named DC1, which provides directory and name resolution services to all computers and devices on the network.

DHCP server DHCP1

The DHCP server, named DHCP1, is configured with a scope that provides Internet Protocol (IP) address leases to computers on the local subnet. The DHCP server can also be configured with additional scopes to provide IP address leases to computers on other subnets if DHCP forwarding is configured on routers.

Client computers

Client computers running Windows 8, Windows 7, and Windows Vista are configured by default as DHCP clients, which obtain IP addresses and DHCP options automatically from the DHCP server.

Core Network Planning

Before you deploy a core network, you must plan the following items.

- Planning subnets
- Planning basic configuration of all servers
- Planning the deployment of DC1
- Planning domain access
- Planning the deployment of DHCP1

The following sections provide more detail on each of these items.

Note

For assistance with planning your deployment, also see Appendix E - Core Network Planning Preparation Sheet.

Planning subnets

In Transmission Control Protocol/Internet Protocol (TCP/IP) networking, routers are used to interconnect the hardware and software used on different physical network segments called subnets. Routers are also used to forward IP packets between each of the subnets. Determine the physical layout of your network, including the number of routers and subnets you need, before proceeding with the instructions in this guide.

In addition, to configure the servers on your network with static IP addresses, you must determine the IP address range that you want to use for the subnet where your core network servers are located. In this guide, the private IP address ranges 10.0.0.1 – 10.0.0.254 and 10.0.1.1 – 10.0.1.254 are used as examples, but you can use any private IP address range that you prefer.

Important

After you select the IP address ranges that you want to use for each subnet, ensure that you configure your routers with an IP address from the same IP address range as that used on the subnet where the router is installed. For example, if your router is configured by default with an IP address of 192.168.1.1, but you are installing the router on a subnet with an IP address range of 10.0.0.0/24, you must reconfigure the router to use an IP address from the 10.0.0.0/24 IP address range.

The following recognized private IP address ranges are specified by Internet Request for Comments (RFC) 1918:

- 10.0.0.0 – 10.255.255.255
- 172.16.0.0 – 172.31.255.255
- 192.168.0.0 – 192.168.255.255

When you use the private IP address ranges as specified in RFC 1918, you cannot connect directly to the Internet using a private IP address because requests going to or from these addresses are automatically discarded by Internet service provider (ISP) routers. To add Internet connectivity to your core network later, you must contract with an ISP to obtain a public IP address.

Important

When using private IP addresses, you must use some type of proxy or network address translation (NAT) server to convert the private IP address ranges on your local network to a public IP address that can be routed on the Internet. Most routers provide NAT services, so selecting a router that is NAT-capable should be fairly simple.

For more information, see Planning the deployment of DHCP1.

Planning basic configuration of all servers

For each server in the core network, you must rename the computer and assign and configure a static IPv4 address and other TCP/IP properties for the computer.

Planning naming conventions for computers and devices

For consistency across your network, it is a good idea to use consistent names for servers, printers, and other devices. Computer names can be used to help users and administrators easily identify the purpose and location of the server, printer, or other device. For example, if you have three DNS servers, one in San Francisco, one in Los Angeles, and one in Chicago, you might use the naming convention *server function-location-number*:

- DNS-DEN-01. This name represents the DNS server in Denver, Colorado. If additional DNS servers are added in Denver, the numeric value in the name can be incremented, as in DNS-DEN-02 and DNS-DEN-03.
- DNS-SPAS-01. This name represents the DNS server in South Pasadena, California.
- DNS-ORL-01. This name represents the DNS server in Orlando, Florida.

For this guide, the server naming convention is very simple, and consists of the primary server function and a number. For example, the domain controller is named DC1 and the DHCP server is named DHCP1.

It is recommended that you choose a naming convention before you install your core network using this guide.

Planning static IP addresses

Before configuring each computer with a static IP address, you must plan your subnets and IP address ranges. In addition, you must determine the IP addresses of your DNS servers. If you plan to install a router that provides access to other networks, such as additional subnets or the Internet, you must know the IP address of the router, also called a default gateway, for static IP address configuration.

The following table provides example values for static IP address configuration.

Configuration items	Example values
IP address	10.0.0.2
Subnet mask	255.255.255.0
Default gateway (Router IP address)	10.0.0.1
Preferred DNS server	10.0.0.2

Note
If you plan on deploying more than one DNS server, you can also plan the Alternate DNS Server IP address.

Planning the deployment of DC1

Following are key planning steps before installing Active Directory Domain Services (AD DS) and DNS on DC1.

Planning the name of the forest root domain

A first step in the AD DS design process is to determine how many forests your organization requires. A forest is the top-level AD DS container, and consists of one or more domains that share a common schema and global catalog. An organization can have multiple forests, but for most organizations, a single forest design is the preferred model and the simplest to administer.

When you create the first domain controller in your organization, you are creating the first domain (also called the forest root domain) and the first forest. Before you take this action using this guide, however, you must determine the best domain name for your organization. In most cases, the organization name is used as the domain name, and in many cases this domain name is registered. If you are planning to deploy external-facing Internet based Web servers to provide

information and services for your customers or partners, choose a domain name that is not already in use, and then register the domain name so that your organization owns it.

Planning the forest functional level

While installing AD DS, you must choose the forest functional level that you want to use. Domain and forest functionality, introduced in Windows Server 2003 Active Directory, provides a way to enable domain- or forest-wide Active Directory features within your network environment. Different levels of domain functionality and forest functionality are available, depending on your environment.

Forest functionality enables features across all the domains in your forest. The following forest functional levels are available:

- Windows Server 2003. This forest functional level supports only Windows Server 2003 domain controllers and domain controllers that are running later versions of the Windows Server operating system.
- Windows Server 2008. This forest functional level supports only domain controllers that are running Windows Server 2008 and later versions of the Windows Server operating system.
- Windows Server 2008 R2. This forest functional level supports Windows Server 2008 R2 domain controllers and domain controllers that are running later versions of the Windows Server operating system.
- Windows Server 2012. This forest functional level supports Windows Server 2012 domain controllers and domain controllers that are running later versions of the Windows Server operating system.

If you are deploying a new domain in a new forest and all of your domain controllers will be running Windows Server 2012, it is recommended that you configure AD DS with the Windows Server 2012 forest functional level during AD DS installation.

Important

After the forest functional level is raised, domain controllers that are running earlier operating systems cannot be introduced into the forest. For example, if you raise the forest functional level to Windows Server 2008 R2, domain controllers running Windows Server 2003 or Windows Server 2008 cannot be added to the forest.

Example configuration items for AD DS are provided in the following table.

Configuration items:	Example values:
Full DNS name	Examples: • corp.contoso.com • example.com
Forest functional level	• Windows Server 2003 • Windows Server 2008 • Windows Server 2008 R2

Configuration items:	Example values:
	• Windows Server 2012
Active Directory Domain Services Database folder location	E:\Configuration\ Or accept the default location.
Active Directory Domain Services Log files folder location	E:\Configuration\ Or accept the default location.
Active Directory Domain Services SYSVOL folder location	E:\Configuration\ Or accept the default location
Directory Restore Mode Administrator Password	**J*p2leO4$F**
Answer file name (optional)	**AD DS_AnswerFile**

Planning DNS zones

On primary, Active Directory-integrated DNS servers, a forward lookup zone is created by default during installation of the DNS Server role. A forward lookup zone allows computers and devices to query for another computer's or device's IP address based on its DNS name. In addition to a forward lookup zone, it is recommended that you create a DNS reverse lookup zone. With a DNS reverse lookup query, a computer or device can discover the name of another computer or device using its IP address. Deploying a reverse lookup zone typically improves DNS performance and greatly increases the success of DNS queries.

When you create a reverse lookup zone, the in-addr.arpa domain, which is defined in the DNS standards and reserved in the Internet DNS namespace to provide a practical and reliable way to perform reverse queries, is configured in DNS. To create the reverse namespace, subdomains within the in-addr.arpa domain are formed, using the reverse ordering of the numbers in the dotted-decimal notation of IP addresses.

The in-addr.arpa domain applies to all TCP/IP networks that are based on Internet Protocol version 4 (IPv4) addressing. The New Zone Wizard automatically assumes that you are using this domain when you create a new reverse lookup zone.

While you are running the New Zone Wizard, the following selections are recommended:

Configuration Items	Example values
Zone type	**Primary zone**, and **Store the zone in Active Directory** is selected
Active Directory Zone Replication Scope	**To all DNS servers in this domain**
First Reverse Lookup Zone Name wizard page	**IPv4 Reverse Lookup Zone**

Configuration Items	Example values
Second Reverse Lookup Zone Name wizard page	Network ID = 10.0.0.
Dynamic Updates	**Allow only secure dynamic updates**

Planning domain access

To log on to the domain, the computer must be a domain member computer and the user account must be created in AD DS before the logon attempt.

Note
Individual computers that are running Windows have a local users and groups user account database that is called the Security Accounts Manager (SAM) user accounts database. When you create a user account on the local computer in the SAM database, you can log onto the local computer, but you cannot log on to a domain. Domain user accounts are created with the Active Directory Users and Computers Microsoft Management Console (MMC) on a domain controller, not with local users and groups on the local computer.

After the first successful logon with domain logon credentials, the logon settings persist unless the computer is removed from the domain or the logon settings are manually changed.

Before you log on to the domain:

- Create user accounts in Active Directory Users and Computers. Each user must have an Active Directory Domain Services user account in Active Directory Users and Computers. For more information, see Create a User Account in Active Directory Users and Computers.
- Ensure the correct IP address configuration. To join a computer to the domain, the computer must have an IP address. In this guide, servers are configured with static IP addresses and client computers receive IP address leases from the DHCP server. For this reason, the DHCP server must be deployed before you join clients to the domain. For more information, see Deploying DHCP1.
- Join the computer to the domain. Any computer that provides or accesses network resources must be joined to the domain. For more information, see Joining Server Computers to the Domain and Logging On and Joining Client Computers to the Domain and Logging On.

Planning the deployment of DHCP1

Following are key planning steps before installing the DHCP server role on DHCP1.

Planning DHCP servers and DHCP forwarding

Because DHCP messages are broadcast messages, they are not forwarded between subnets by routers. If you have multiple subnets and want to provide DHCP service for each subnet, you must do one of the following:

- Install a DHCP server on each subnet
- Configure routers to forward DHCP broadcast messages across subnets and configure multiple scopes on the DHCP server, one scope per subnet.

In most cases, configuring routers to forward DHCP broadcast messages is more cost effective than deploying a DHCP server on each physical segment of the network.

Planning IP address ranges

Each subnet must have its own unique IP address range. These ranges are represented on a DHCP server with scopes.

A scope is an administrative grouping of IP addresses for computers on a subnet that use the DHCP service. The administrator first creates a scope for each physical subnet and then uses the scope to define the parameters used by clients.

A scope has the following properties:

- A range of IP addresses from which to include or exclude addresses used for DHCP service lease offerings.
- A subnet mask, which determines the subnet prefix for a given IP address.
- A scope name assigned when it is created.
- Lease duration values, which are assigned to DHCP clients that receive dynamically allocated IP addresses.
- Any DHCP scope options configured for assignment to DHCP clients, such as DNS server IP address and router/default gateway IP address.
- Reservations are optionally used to ensure that a DHCP client always receives the same IP address.

Before deploying your servers, list your subnets and the IP address range you want to use for each subnet.

Planning subnet masks

Network IDs and host IDs within an IP address are distinguished by using a subnet mask. Each subnet mask is a 32-bit number that uses consecutive bit groups of all ones (1) to identify the network ID and all zeroes (0) to identify the host ID portions of an IP address.

For example, the subnet mask normally used with the IP address 131.107.16.200 is the following 32-bit binary number:

```
11111111 11111111 00000000 00000000
```

This subnet mask number is 16 one-bits followed by 16 zero-bits, indicating that the network ID and host ID sections of this IP address are both 16 bits in length. Normally, this subnet mask is displayed in dotted decimal notation as 255.255.0.0.

The following table displays subnet masks for the Internet address classes.

Address class	Bits for subnet mask	Subnet mask
Class A	11111111 00000000 00000000 00000000	255.0.0.0
Class B	11111111 11111111 00000000 00000000	255.255.0.0
Class C	11111111 11111111 11111111 00000000	255.255.255.0

When you create a scope in DHCP and you enter the IP address range for the scope, DHCP provides these default subnet mask values. Typically, default subnet mask values are acceptable for most networks with no special requirements and where each IP network segment corresponds to a single physical network.

In some cases, you can use customized subnet masks to implement IP subnetting. With IP subnetting, you can subdivide the default host ID portion of an IP address to specify subnets, which are subdivisions of the original class-based network ID.

By customizing the subnet mask length, you can reduce the number of bits that are used for the actual host ID.

To prevent addressing and routing problems, you should make sure that all TCP/IP computers on a network segment use the same subnet mask and that each computer or device has an unique IP address.

Planning exclusion ranges

When you create a scope on a DHCP server, you specify an IP address range that includes all of the IP addresses that the DHCP server is allowed to lease to DHCP clients, such as computers and other devices. If you then go and manually configure some servers and other devices with static IP addresses from the same IP address range that the DHCP server is using, you can accidentally create an IP address conflict, where you and the DHCP server have both assigned the same IP address to different devices.

To solve this problem, you can create an exclusion range for the DHCP scope. An exclusion range is a contiguous range of IP addresses within the scope's IP address range that the DHCP server is not allowed to use. If you create an exclusion range, the DHCP server does not assign the addresses in that range, allowing you to manually assign these addresses without creating an IP address conflict.

You can exclude IP addresses from distribution by the DHCP server by creating an exclusion range for each scope. You should use exclusions for all devices that are configured with a static IP address. The excluded addresses should include all IP addresses that you assigned manually to other servers, non-DHCP clients, diskless workstations, or Routing and Remote Access and PPP clients.

It is recommended that you configure your exclusion range with extra addresses to accommodate future network growth. The following table provides an example exclusion range for a scope with an IP address range of 10.0.0.1 - 10.0.0.254 and a subnet mask of 255.255.255.0.

Configuration items	Example values
Exclusion range Start IP Address	10.0.0.1
Exclusion range End IP Address	10.0.0.25

Planning TCP/IP static configuration

Certain devices, such as routers, DHCP servers, and DNS servers, must be configured with a static IP address. In addition, you might have additional devices, such as printers, that you want to ensure always have the same IP address. List the devices that you want to configure statically for each subnet, and then plan the exclusion range you want to use on the DHCP server to ensure that the DHCP server does not lease the IP address of a statically configured device. An exclusion range is a limited sequence of IP addresses within a scope, excluded from DHCP service offerings. Exclusion ranges assure that any addresses in these ranges are not offered by the server to DHCP clients on your network.

For example, if the IP address range for a subnet is 192.168.0.1 through 192.168.0.254 and you have ten devices that you want to configure with a static IP address, you can create an exclusion range for the 192.168.0.*x* scope that includes ten or more IP addresses: 192.168.0.1 through 192.168.0.15.

In this example, you use ten of the excluded IP addresses to configure servers and other devices with static IP addresses and five additional IP addresses are left available for static configuration of new devices that you might want to add in the future. With this exclusion range, the DHCP server is left with an address pool of 192.168.0.16 through 192.168.0.254.

Additional example configuration items for AD DS and DNS are provided in the following table.

Configuration items	Example values
Network Connect Bindings	Ethernet
DNS Server Settings	DC1.corp.contoso.com
Preferred DNS server IP address	10.0.0.2
Add Scope dialog box values 1. Scope Name 2. Starting IP Address 3. Ending IP Address 4. Subnet Mask 5. Default Gateway (optional)	 1. Primary Subnet 2. 10.0.0.1 3. 10.0.0.254 4. 255.255.255.0 5. 10.0.0.1

Configuration items	Example values
6. Lease duration	6. 8 days
IPv6 DHCP Server Operation Mode	Not enabled

Core Network Deployment

To deploy a core network, the basic steps are as follows:

1. Configuring All Servers
2. Deploying DC1
3. Joining Server Computers to the Domain and Logging On
4. Deploying DHCP1
5. Joining Client Computers to the Domain and Logging On
6. Deploying optional features for network access authentication and Web services

Notes

- Equivalent Windows PowerShell commands are provided for most procedures in this guide. Before running these cmdlets in Windows PowerShell, replace example values with values that are appropriate for your network deployment. In addition, you must enter each cmdlet on a single line in Windows PowerShell. In this guide, individual cmdlets might appear on several lines due to formatting constraints and the display of the document by your browser or other application.
- The procedures in this guide do not include instructions for those cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Configuring All Servers

Before installing other technologies, such as Active Directory Domain Services or DHCP, it is important to configure the following items.

- Rename the computer
- Configure a static IP address

You can use the following sections to perform these actions for each server.

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

Rename the computer

You can use the procedure in this section to provide computers running Windows Server 2012 with a different computer name. Renaming the computer is useful for circumstances in which the operating system has automatically created a computer name that you do not want to use.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlets on separate lines, and then press ENTER. You must also replace *ComputerName* with the name that you want to use.

```
Rename-Computer ComputerName

Restart-Computer
```

To rename computers running Windows Server 2012

1. In Server Manager, click **Local Server**. The computer **Properties** are displayed in the details pane.
2. In **Properties**, in **Computer name**, click the existing computer name. The **System Properties** dialog box opens. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
3. In the **Computer Name/Domain Changes** dialog box, in **Computer name**, type a new name for your computer. For example, if you want to name the computer DC1, type **DC1**.
4. Click **OK** twice, and then click **Close**. If you want to restart the computer immediately to complete the name change, click **Restart Now**. Otherwise, click **Restart Later**.

Note

For information on how to rename computers that are running other Microsoft operating systems, see Appendix A - Renaming computers.

Configure a static IP address

You can use the procedures in this topic to configure the Internet Protocol version 4 (IPv4) properties of a network connection with a static IP address for computers running Windows Server 2012.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlets on separate lines, and then press ENTER. You must also replace interface names and IP addresses in this example with the values that you want to use to configure your computer.

```
New-NetIPAddress -IPAddress 10.0.0.2 -InterfaceAlias "Ethernet" -DefaultGateway
10.0.0.1 -AddressFamily IPv4 -PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Ethernet" -ServerAddresses 127.0.0.1
```

To configure a static IP address on a computer running Windows Server 2012

1. In the task bar, right-click the Network icon, and then click **Open Network and Sharing Center**.
2. In **Network and Sharing Center**, click **Change adapter settings**. The **Network Connections** folder opens and displays the available network connections.

3. In **Network Connections**, right-click the connection that you want to configure, and then click **Properties**. The network connection **Properties** dialog box opens.
4. In the network connection **Properties** dialog box, in **This connection uses the following items**, select **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**. The **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box opens.
5. In **Internet Protocol Version 4 (TCP/IPv4) Properties**, on the **General** tab, click **Use the following IP address**. In **IP address**, type the IP address that you want to use.
6. Press tab to place the cursor in **Subnet mask**. A default value for subnet mask is entered automatically. Either accept the default subnet mask, or type the subnet mask that you want to use.
7. In **Default gateway**, type the IP address of your default gateway.

   Note

   You must configure **Default gateway** with the same IP address that you use on the local area network (LAN) interface of your router. For example, if you have a router that is connected to a wide area network (WAN) such as the Internet as well as to your LAN, configure the LAN interface with the same IP address that you will then specify as the **Default gateway**. In another example, if you have a router that is connected to two LANs, where LAN A uses the address range 10.0.0.0/24 and LAN B uses the address range 192.168.0.0/24, configure the LAN A router IP address with an address from that address range, such as 10.0.0.1. In addition, in the DHCP scope for this address range, configure **Default gateway** with the IP address 10.0.0.1. For the LAN B, configure the LAN B router interface with an address from that address range, such as 192.168.0.1, and then configure the LAN B scope 192.168.0.0/24 with a **Default gateway** value of 192.168.0.1.
8. In **Preferred DNS server**, type the IP address of your DNS server. If you plan to use the local computer as the preferred DNS server, type the IP address of the local computer.
9. In **Alternate DNS Server**, type the IP address of your alternate DNS server, if any. If you plan to use the local computer as an alternate DNS server, type the IP address of the local computer.
10. Click **OK**, and then click **Close**.

Note

For information on how to configure a static IP address on computers that are running other Microsoft operating systems, see Appendix B - Configuring static IP addresses.

Deploying DC1

To deploy DC1, which is the computer running Active Directory Domain Services (AD DS) and DNS, you must complete these steps in the following order:

- Perform the steps in the section Configuring All Servers.
- Install AD DS and DNS for a New Forest

- Create a User Account in Active Directory Users and Computers
- Add a Group
- Assign Group Membership
- Configure a DNS Reverse Lookup Zone

Administrative privileges

If you are installing a small network and are the only administrator for the network, it is recommended that you create a user account for yourself, and then add your user account as a member of both Enterprise Admins and Domain Admins. Doing so will make it easier for you to act as the administrator for all network resources. It is also recommended that you log on with this account only when you need to perform administrative tasks, and that you create a separate user account for performing non-IT related tasks.

If you have a larger organization with multiple administrators, refer to AD DS documentation to determine the best group membership for organization employees.

Differences between domain user accounts and user accounts on the local computer

One of the advantages of a domain-based infrastructure is that you do not need to create user accounts on each computer in the domain. This is true whether the computer is a client computer or a server.

Because of this, you should not create user accounts on each computer in the domain. Create all user accounts in Active Directory Users and Computers and use the preceding procedures to assign group membership. By default, all user accounts are members of the Domain Users group.

All members of the Domain Users group can log on to any client computer after it is joined to the domain.

You can configure user accounts to designate the days and times that the user is allowed to log on to the computer. You can also designate which computers each user is allowed to use. To configure these settings, open Active Directory Users and Computers, locate the user account that you want to configure, and double-click the account. In the user account **Properties**, click the **Account** tab, and then click either **Logon Hours** or **Log On To**.

Install AD DS and DNS for a New Forest

You can use this procedure to install Active Directory Domain Services (AD DS) and DNS and to create a new domain in a new forest.

Membership in **Administrators** is the minimum required to perform this procedure.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlets on separate lines, and then press ENTER. You must also replace the domain name in this example with the value that you want to use.

```
Install-WindowsFeature AD-Domain-Services -IncludeManagementTools

Install-ADDSForest -DomainName corp.contoso.com
```

To install Active Directory Domain Services and DNS

1. On DC1, in **Server Manager**, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
2. In **Before You Begin**, click **Next**.

   Note
   The **Before You Begin** page of the Add Roles and Features Wizard is not displayed if you have previously selected **Skip this page by default** when the Add Roles and Features Wizard was run.
3. In **Select Installation Type**, ensure that **Role-Based or feature-based installation** is selected, and then click **Next**.
4. In **Select destination server**, ensure that **Select a server from the server pool** is selected. In **Server Pool**, ensure that the local computer is selected. Click **Next**.
5. In **Select server roles**, in **Roles**, click **Active Directory Domain Services**. In **Add features that are required for Active Directory Domain Services**, click **Add Features**. Click **Next**.
6. In **Select features**, click **Next**, and in **Active Directory Domain Services**, review the information that is provided, and then click **Next**.
7. In **Confirm installation selections**, click **Install**. The Installation progress page displays status during the installation process. When the process completes, in the message details, click **Promote this server to a domain controller**. The Active Directory Domain Services Configuration Wizard opens.
8. In **Deployment Configuration**, select **Add a new forest**. In **Root domain name**, type the fully qualified domain name (FQDN) for your domain. For example, if your FQDN is corp.contoso.com, type **corp.contoso.com**. Click **Next**.
9. In **Domain Controller Options**, in **Select functional level of the new forest and root domain**, select the forest functional level and domain functional level that you want to use. In **Specify domain controller capabilities**, ensure that **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected. In **Password** and **Confirm password**, type the Directory Services Restore Mode (DSRM) password that you want to use. Click **Next**.
10. In **DNS Options**, click **Next**.
11. In **Additional Options**, verify the NetBIOS name that is assigned to the domain, and change it only if necessary. Click **Next**.
12. In **Paths**, in **Specify the location of the AD DS database, log files, and SYSVOL**, do one of the following:
    - Accept the default values.
    - Type folder locations that you want to use for **Database folder**, **Log files folder**, and **SYSVOL folder**.
13. Click **Next**.
14. In **Review Options**, review your selections.

15. If you want to export settings to a Windows PowerShell script, click **View script**. The script opens in Notepad, and you can save it to the folder location that you want. Click **Next**. In **Prerequisites Check**, your selections are validated. When the check completes, click **Install**. When prompted by Windows, click **Close**.

Create a User Account in Active Directory Users and Computers

You can use this procedure to create a new domain user account in Active Directory Users and Computers Microsoft Management Console (MMC).

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlet on one line, and then press ENTER. You must also replace the user account name in this example with the value that you want to use.

```
New-ADUser -SamAccountName User1 -AccountPassword (read-host "Set user password" -
assecurestring) -name "User1" -enabled $true -PasswordNeverExpires $true -
ChangePasswordAtLogon $false
```

After you press ENTER, type the password for the user account. The account is created and, by default, is granted membership to the Domain Users group.

With the following cmdlet, you can assign additional group memberships for the new user account. The example below adds User1 to the Domain Admins and Enterprise Admins groups. Ensure before running this command that you change the user account name, domain name, and groups to match your requirements.

```
Add-ADPrincipalGroupMembership -Identity
"CN=User1,CN=Users,DC=corp,DC=contoso,DC=com" -MemberOf "CN=Enterprise
Admins,CN=Users,DC=corp,DC=contoso,DC=com","CN=Domain
Admins,CN=Users,DC=corp,DC=contoso,DC=com"
```

To create a user account

1. On DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click **corp.contoso.com**.
2. In the details pane, right-click the folder in which you want to add a user account.

   Where?

   - Active Directory Users and Computers/*domain node*/*folder*
3. Point to **New**, and then click **User**. The **New Object – User** dialog box opens.
4. In **First name**, type the user's first name.
5. In **Initials**, type the user's initials.
6. In **Last name**, type the user's last name.

7. Modify **Full name** to add initials or reverse the order of first and last names.
8. In **User logon name**, type the user logon name. Click **Next**.
9. In **New Object - User**, in **Password** and **Confirm password**, type the user's password, and then select the appropriate password options.
10. Click **Next**, review the new user account settings, and then click **Finish**.

Add a Group

You can use this procedure to create a new group in Active Directory Users and Computers Microsoft Management Console (MMC).

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To add a group

1. On DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click **corp.contoso.com**.
2. In the details pane, right-click the folder in which you want to add a new group.

   Where?

   - Active Directory Users and Computers/*domain node*/*folder*
3. Point to **New**, and then click **Group**.
4. In **New Object – Group**, in **Group name**, type the name of the new group.

   By default, the name you type is also entered as the pre-Windows 2000 name of the new group.
5. In **Group scope**, select **Global**, unless you have a specific reason for chosing another option.
6. In **Group type**, select **Security**.
7. Click **OK**.

Assign Group Membership

You can use this procedure to add a user, computer, or group to a group in Active Directory Users and Computers Microsoft Management Console (MMC).

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

To assign group membership

1. On DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click **corp.contoso.com**.

2. In the details pane, double-click the folder that contains the group to which you want to add a member.

   Where?

   - **Active Directory Users and Computers**/*domain node*/*folder that contains the group*

3. In the details pane, right-click the object that you want to add to a group, such as a user or computer, and then click **Properties**. The object's **Properties** dialog box opens. Click the **Member of** tab.
4. On the **Member of** tab, click **Add**.
5. In **Enter the object names to select**, type the name of the group to which you want to add the object, and then click **OK**.
6. To assign group membership to other users, groups or computers, repeat steps 4 and 5 of this procedure.

Configure a DNS Reverse Lookup Zone

You can use this procedure to configure a reverse lookup zone in Domain Name System (DNS).

Membership in **Domain Admins** is the minimum required to perform this procedure.

Notes

- For medium and large organizations, it's recommended that you configure and use the DNSAdmins group in Active Directory Users and Computers. For more information, see Additional Technical Resources
- To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlet on one line, and then press ENTER. You must also replace the DNS reverse lookup zone and zonefile names in this example with the values that you want to use. Ensure that you reverse the network ID for the reverse zone name. For example, if the network ID is 192.168.0, create the reverse lookup zone name **0.168.192.in-addr.arpa**.

```
Add-DnsServerPrimaryZone 0.0.10.in-addr.arpa -ZoneFile 0.0.10.in-addr.arpa.dns
```

To configure a DNS reverse lookup zone

1. On DC1, in Server Manager, click **Tools**, and then click **DNS**. The DNS MMC opens.
2. In DNS, if it is not already expanded, double-click the server name to expand the tree. For example, if the DNS server name is DC1, double-click **DC1**.
3. Select **Reverse Lookup Zones**, right-click **Reverse Lookup Zones**, and then click **New Zone**. The New Zone Wizard opens.
4. In **Welcome to the New Zone Wizard**, click **Next**.
5. In **Zone Type**, select **Primary zone**.
6. If your DNS server is a writeable domain controller, ensure that **Store the zone in Active Directory** is selected. Click **Next**.
7. In **Active Directory Zone Replication Scope**, select **To all DNS servers running on domain controllers in this domain**, unless you have a specific reason to choose a different option. Click **Next**.
8. In the first **Reverse Lookup Zone Name** page, select **IPv4 Reverse Lookup Zone**. Click

Next.

9. In the second **Reverse Lookup Zone Name** page, do one of the following:
   - In **Network ID**, type the network ID of your IP address range. For example, if your IP address range is 10.0.0.1 through 10.0.0.254, type **10.0.0**.
   - In **Reverse lookup zone name**, your IPv4 reverse lookup zone name is automatically added. Click **Next**.
10. In **Dynamic Update**, select the type of dynamic updates that you want to allow. Click **Next**.
11. In **Completing the New Zone Wizard**, review your choices, and then click **Finish**.

Joining Server Computers to the Domain and Logging On

After you have installed Active Directory Domain Services (AD DS) and created one or more user accounts that have permissions to join a computer to the domain, you can join core network servers to the domain and log on to the servers in order to install additional technologies, such as Dynamic Host Configuration Protocol (DHCP).

On all servers that you are deploying, except for the server running AD DS, do the following:

1. Complete the procedures provided in Configuring All Servers.
2. Use the instructions in the following two procedures to join your servers to the domain and to log on to the servers to perform additional deployment tasks:

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlet, and then press ENTER. You must also replace the domain name with the name that you want to use.

```
Add-Computer -DomainName corp.contoso.com
```

When you are prompted to do so, type the user name and password for an account that has permission to join a computer to the domain. To restart the computer, type the following command and press ENTER.

```
Restart-Computer
```

To join computers running Windows Server 2012 to the domain

1. In Server Manager, click **Local Server**. In the details pane, click **WORKGROUP**. The **System Properties** dialog box opens.
2. In the **System Properties** dialog box, click **Change**. The **Computer Name/Domain Changes** dialog box opens.
3. In **Computer Name**, in **Member of**, click **Domain**, and then type the name of the domain that you want to join. For example, if the domain name is corp.contoso.com, type **corp.contoso.com**.
4. Click **OK**. The **Windows Security** dialog box opens.
5. In **Computer Name/Domain Changes**, in **User name**, type the user name, and in **Password**, type the password, and then click **OK**. The **Computer Name/Domain**

Changes dialog box opens, welcoming you to the domain. Click **OK**.

6. The **Computer Name/Domain Changes** dialog box displays a message indicating that you must restart the computer to apply the changes. Click **OK**.
7. On the **System Properties** dialog box, on the **Computer Name** tab, click **Close**. The **Microsoft Windows** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes. Click **Restart Now**.

Note

For information on how to join computers that are running other Microsoft operating systems to the domain, see Appendix C – Joining computers to the domain.

To log on to the domain using computers running Windows Server 2012

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click the left arrow, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain corp.contoso.com with an account named **User-01**, type **CORP\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Note

For information on how to log on to the domain using computers that are running other Microsoft operating systems, see Appendix D – Log on to the domain.

Deploying DHCP1

Before deploying this component of the core network, you must do the following:

- Perform the steps in the section Configuring All Servers.
- Perform the steps in the section Joining Server Computers to the Domain and Logging On.

To deploy DHCP1, which is the computer running the Dynamic Host Configuration Protocol (DHCP) server role, you must complete these steps in the following order:

- Install Dynamic Host Configuration Protocol (DHCP)
- Create and Activate a New DHCP Scope

Notes

To perform these procedures by using Windows PowerShell, open PowerShell and type the following cmdlets on separate lines, and then press ENTER. You must also replace the scope name, IP address start and end ranges, subnet mask, and other values in this example with the values that you want to use.

```
Install-WindowsFeature DHCP -IncludeManagementTools

Add-DhcpServerv4Scope -name "Corpnet" -StartRange 10.0.0.1 -EndRange 10.0.0.254 -
SubnetMask 255.255.255.0 -State Active
```

```
Add-DhcpServerv4ExclusionRange –ScopeID 10.0.0.0 –StartRange 10.0.0.1 –EndRange
10.0.0.15
Set-DhcpServerv4OptionValue -OptionID 3 -Value 10.0.0.1 -ScopeID 10.0.0.0 -
ComputerName DHCP1.corp.contoso.com
Add-DhcpServerv4Scope -name "Corpnet2" -StartRange 10.0.1.1 -EndRange 10.0.1.254 -
SubnetMask 255.255.255.0 –State Active
Add-DhcpServerv4ExclusionRange –ScopeID 10.0.1.0 –StartRange 10.0.1.1 –EndRange
10.0.1.15
Set-DhcpServerv4OptionValue -OptionID 3 -Value 10.0.1.1 -ScopeID 10.0.1.0 -
ComputerName DHCP1.corp.contoso.com
Set-DhcpServerv4OptionValue -DnsDomain corp.contoso.com -DnsServer 10.0.0.2
Add-DhcpServerInDC -DnsName DHCP1.corp.contoso.com
```

Install Dynamic Host Configuration Protocol (DHCP)

You can use this procedure to install and configure the DHCP Server role using the Add Roles and Features Wizard.

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To install DHCP

1. On DHCP1, in Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
2. In **Before You Begin**, click **Next**.

   Note
   The **Before You Begin** page of the Add Roles and Features Wizard is not displayed if you have previously selected **Skip this page by default** when the Add Roles and Features Wizard was run.
3. In **Select Installation Type**, ensure that **Role-Based or feature-based installation** is selected, and then click **Next**.
4. In **Select destination server**, ensure that **Select a server from the server pool** is selected. In **Server Pool**, ensure that the local computer is selected. Click **Next**.
5. In **Select Server Roles**, in **Roles**, select **DHCP Server**, and then click **Next**. In **Add features that are required for DHCP Server**, click **Add Features**. Click **Next**.
6. In **Select features**, click **Next**, and in **DHCP Server**, review the information that is provided, and then click **Next**.
7. In **Confirm installation selections**, click **Restart the destination server automatically if required**. When you are prompted to confirm this selection, click **Yes**, and then click **Install**. The Installation progress page displays status during the installation process. When the process completes, the message “Configuration required. Installation succeeded on *ComputerName*” is displayed, where *ComputerName* is the name of the computer upon which you installed DHCP Server. In the message window, click

Complete DHCP configuration. The DHCP Post-Install configuration wizard opens. Click **Next**.

8. In **Authorization**, specify the credentials that you want to use to authorize the DHCP server in Active Directory Domain Services, and then click **Commit**. After authorization is complete, click **Close**.

Create and Activate a New DHCP Scope

You can use this procedure to create a new DHCP scope using the DHCP Microsoft Management Console (MMC). When you complete the procedure, the scope is activated and the exclusion range that you create prevents the DHCP server from leasing the IP addresses that you use to statically configure your servers and other devices that require a static IP address.

Membership in **DHCP Administrators**, or equivalent, is the minimum required to perform this procedure.

To create and activate a new DHCP Scope

1. On DHCP1, in Server Manager, click **Tools**, and then click **DHCP**. The DHCP MMC opens.
2. In **DHCP**, double-click the server name. For example, if the DHCP server name is DHCP1.corp.contoso.com, double-click **DHCP1.corp.contoso.com**.
3. Right-click **IPv4**, and then click **New Scope**. The New Scope Wizard opens.
4. In **Welcome to the New Scope Wizard**, click **Next**.
5. In **Scope Name**, in **Name**, type a name for the scope. For example, type **Subnet 1**.
6. In **Description**, type a description for the new scope, and then click **Next**.
7. In **IP Address Range**, do the following:
   a. In **Start IP address**, type the IP address that is the first IP address in the range. For example, type **10.0.0.1**.
   b. In **End IP address**, type the IP address that is the last IP address in the range. For example, type **10.0.0.254**. Values for **Length** and **Subnet mask** are entered automatically, based on the IP address you entered for **Start IP address**.
   c. If necessary, modify the values in **Length** or **Subnet mask**, as appropriate for your addressing scheme.
   d. Click **Next**.
8. In **Add Exclusions**, do the following:
   a. In **Start IP address**, type the IP address that is the first IP address in the exclusion range. For example, type **10.0.0.1**.
   b. In **End IP address**, type the IP address that is the last IP address in the exclusion range, For example, type **10.0.0.15**.
9. Click **Add**, and then click **Next**.
10. In **Lease Duration**, modify the default values for **Days**, **Hours**, and **Minutes**, as appropriate for your network, and then click **Next**.
11. In **Configure DHCP Options**, select **Yes, I want to configure these options now**, and

then click **Next**.

12. In **Router (Default Gateway)**, do one of the following:
    - If you do not have routers on your network, click **Next**.
    - In **IP address**, type the IP address of your router or default gateway. For example, type **10.0.0.1**. Click **Add**, and then click **Next**.
13. In **Domain Name and DNS Servers**, do the following:
    a. In **Parent domain**, type the name of the DNS domain that clients use for name resolution. For example, type **corp.contoso.com**.
    b. In **Server name**, type the name of the DNS computer that clients use for name resolution. For example, type **DC1**.
    c. Click **Resolve**. The IP address of the DNS server is added in **IP address**. Click **Add**, wait for DNS server IP address validation to complete, and then click **Next**.
14. In **WINS Servers**, because you do not have WINS servers on your network, click **Next**.
15. In **Activate Scope**, select **Yes, I want to activate this scope now**.
16. Click **Next**, and then click **Finish**.

Important

To create new scopes for additional subnets, repeat this procedure. Use a different IP address range for each subnet that you plan to deploy, and ensure that DHCP message forwarding is enabled on all routers that lead to other subnets.

Joining Client Computers to the Domain and Logging On

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following cmdlet, and then press ENTER. You must also replace the domain name with the name that you want to use.

```
Add-Computer -DomainName corp.contoso.com
```

When you are prompted to do so, type the user name and password for an account that has permission to join a computer to the domain. To restart the computer, type the following command and press ENTER.

```
Restart-Computer
```

To join computers running Windows 8 to the domain

1. Log on to the computer with the local Administrator account.
2. Right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
3. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.
4. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
5. In **Computer Name**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is corp.contoso.com, type

corp.contoso.com.

6. Click **OK**. The **Windows Security** dialog box opens.
7. In **Computer Name/Domain Changes**, in **User name**, type the user name, and in **Password**, type the password, and then click **OK**. The **Computer Name/Domain Changes** dialog box opens, welcoming you to the domain. Click **OK**.
8. The **Computer Name/Domain Changes** dialog box displays a message indicating that you must restart the computer to apply the changes. Click **OK**.
9. On the **System Properties** dialog box, on the **Computer Name** tab, click **Close**. The **Microsoft Windows** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes. Click **Restart Now**.

To log on to the domain using computers running Windows 8

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click the left arrow, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain corp.contoso.com with an account named **User-01**, type **CORP\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Deploying optional features for network access authentication and Web services

If you intend to deploy network access servers, such as wireless access points or VPN servers, after installing your core network, it is recommended that you deploy both an NPS server and a Web server. For network access deployments, the use of secure certificate-based authentication methods is recommended. You can use NPS to manage network access policies and to deploy secure authentication methods. You can use a Web server to publish the certificate revocation list (CRL) of your certification authority (CA) that provides certificates for secure authentication.

Note

You can deploy server certificates and other additional features by using Core Network Companion Guides. For more information, see Additional Technical Resources.

The following illustration shows the Windows Server Core Network topology.

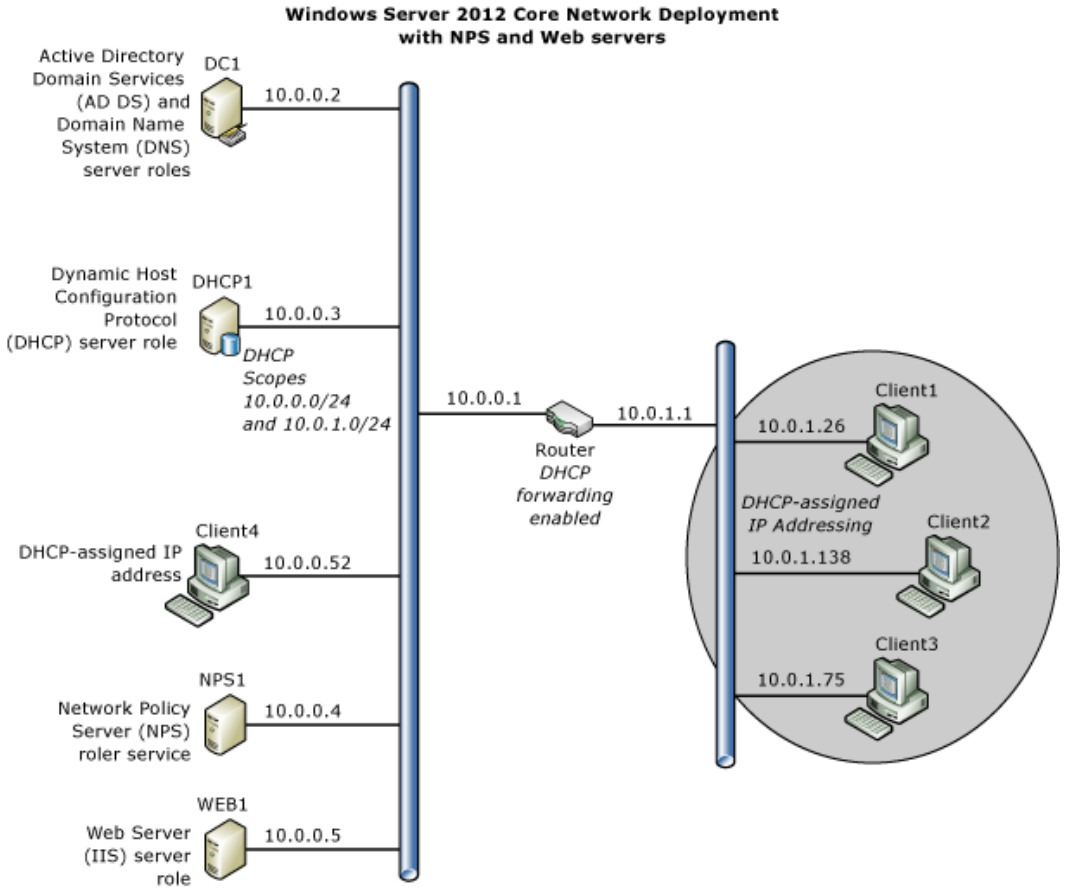

The following sections provide information on adding NPS and Web servers to your network.

- Deploying NPS1
- Deploying WEB1

Deploying NPS1

The Network Policy Server (NPS) server is installed as a preparatory step for deploying other network access technologies, such as virtual private network (VPN) servers, wireless access points, and 802.1X authenticating switches. In addition, installing NPS prepares your network for the deployment of Network Access Protection (NAP).

Network Policy Server (NPS) allows you to centrally configure and manage network policies with the following three features: Remote Authentication Dial-In User Service (RADIUS) server, RADIUS proxy, and Network Access Protection (NAP) policy server.

NPS is an optional component of a core network, but you should install NPS if any of the following are true:

- You are planning to expand your network to include remote access servers that are compatible with the RADIUS protocol, such as a computer running Windows Server® 2012, Windows Server 2008 R2 or Windows Server 2008 and Routing and Remote Access service, Terminal Services Gateway, or Remote Desktop Gateway.

- You plan to deploy NAP.
- You plan to deploy 802.1X wired or wireless access.

Before deploying this role service, you must do the following:

- Perform the steps in the section Configuring All Servers.
- Perform the steps in the section Joining Server Computers to the Domain and Logging On

To deploy NPS1, which is the computer running the Network Policy Server (NPS) role service of the Network Policy and Access Services server role, you must complete this step:

- Planning the deployment of NPS1
- Install Network Policy Server (NPS)
- Register the NPS Server in the Default Domain

Planning the deployment of NPS1

If you intend to deploy network access servers, such as wireless access points or VPN servers, after deploying your core network, it is recommended that you deploy NPS.

When you use NPS as a Remote Authentication Dial-In User Service (RADIUS) server, NPS performs authentication and authorization for connection requests through your network access servers. NPS also allows you to centrally configure and manage network policies that determine who can access the network, how they can access the network, and when they can access the network.

Following are key planning steps before installing NPS.

- Plan the user accounts database. By default, if you join the server running NPS to an Active Directory domain, NPS performs authentication and authorization using the AD DS user accounts database. In some cases, such as with large networks that use NPS as a RADIUS proxy to forward connection requests to other RADIUS servers, you might want to install NPS on a non-domain member computer.
- Plan the use of Network Access Protection (NAP). With some NAP enforcement methods, it is required that you install NPS on a specific server. For example, if you deploy NAP with DHCP, NPS must be installed on the DHCP server.
- Plan RADIUS accounting. NPS allows you to log accounting data to a SQL Server database or to a text file on the local computer. If you want to use SQL Server logging, plan the installation and configuration of your server running SQL Server.

Install Network Policy Server (NPS)

You can use this procedure to install Network Policy Server (NPS) by using the Add Roles Wizard. NPS is a role service of the Network Policy and Access Services server role.

Note

By default, NPS listens for RADIUS traffic on ports 1812, 1813, 1645, and 1646 on all installed network adapters. If Windows Firewall with Advanced Security is enabled when you install NPS, firewall exceptions for these ports are automatically created during the installation process for both Internet Protocol version 6 (IPv6) and IPv4 traffic. If your network access servers are configured to send RADIUS traffic over ports other than

these defaults, remove the exceptions created in Windows Firewall with Advanced Security during NPS installation, and create exceptions for the ports that you do use for RADIUS traffic.

Administrative Credentials

To complete this procedure, you must be a member of the **Domain Admins** group.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following, and then press ENTER.

```
Install-WindowsFeature NPAS-Policy-Server -IncludeManagementTools
```

To install NPS

1. On NPS1, in Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
2. In **Before You Begin**, click **Next**.

   Note

   The **Before You Begin** page of the Add Roles and Features Wizard is not displayed if you have previously selected **Skip this page by default** when the Add Roles and Features Wizard was run.
3. In **Select Installation Type**, ensure that **Role-Based or feature-based installation** is selected, and then click **Next**.
4. In **Select destination server**, ensure that **Select a server from the server pool** is selected. In **Server Pool**, ensure that the local computer is selected. Click **Next**.
5. In **Select Server Roles**, in **Roles**, select **Network Policy and Access Services**, and then click **Next**
6. In **Select features**, click **Next**, and in **Network Policy and Access Services**, review the information that is provided, and then click **Next**.
7. In **Select role services**, click **Network Policy Server**. In **Add features that are required for Network Policy Server**, click **Add Features**. Click **Next**.
8. In **Confirm installation selections**, click **Restart the destination server automatically if required**. When you are prompted to confirm this selection, click **Yes**, and then click **Install**. The Installation progress page displays status during the installation process. When the process completes, the message “Installation succeeded on *ComputerName*” is displayed, where *ComputerName* is the name of the computer upon which you installed Network Policy Server. Click **Close**.

Register the NPS Server in the Default Domain

You can use this procedure to register an NPS server in the domain where the server is a domain member.

NPS servers must be registered in Active Directory so that they have permission to read the dial-in properties of user accounts during the authorization process. Registering an NPS server adds the server to the **RAS and IAS Servers** group in Active Directory.

Administrative credentials

To complete this procedure, you must be a member of the **Domain Admins** group.

To register an NPS server in its default domain

1. On NPS1, in Server Manager, click Tools, and then click **Network Policy Server**. The Network Policy Server MMC opens.
2. Right-click **NPS (Local)**, and then click **Register server in Active Directory**. The **Network Policy Server** dialog box opens.
3. In **Network Policy Server**, click **OK**, and then click **OK** again.

Deploying WEB1

The Web Server (IIS) role in Windows Server 2012 provides a secure, easy-to-manage, modular and extensible platform for reliably hosting websites, services, and applications. With IIS 8 you can share information with users on the Internet, an intranet, or an extranet. IIS 8 is a unified web platform that integrates IIS, ASP.NET, FTP services, PHP, and Windows Communication Foundation (WCF).

In addition to allowing you to publish a CRL for access by domain member computers, the Web Server (IIS) server role allows you to use the Web Server (IIS) role to set up and manage multiple websites, web applications, and FTP sites. IIS8 also provides the following benefits:

- Maximize web security through a reduced server foot print and automatic application isolation.
- Easily deploy and run ASP.NET, classic ASP, and PHP web applications on the same server.
- Achieve application isolation by giving worker processes a unique identity and sandboxed configuration by default, further reducing security risks.
- Easily add, remove, and even replace built-in IIS components with custom modules, suited for customer needs.
- Speed up your website through built-in dynamic caching and enhanced compression.

To deploy WEB1, which is the computer that is running the Web Server (IIS) server role, you must do the following:

- Perform the steps in the section Configuring All Servers.
- Perform the steps in the section Joining Server Computers to the Domain and Logging On
- Install the Web Server (IIS) server role

Install the Web Server (IIS) server role

To complete this procedure, you must be a member of the **Administrators** group.

Notes

To perform this procedure by using Windows PowerShell, open PowerShell and type the following, and then press ENTER.

```
Install-WindowsFeature Web-Server -IncludeManagementTools
```

1. In **Server Manager**, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
2. In **Before You Begin**, click **Next**.

   Note
   The **Before You Begin** page of the Add Roles and Features Wizard is not displayed if you have previously selected **Skip this page by default** when the Add Roles and Features Wizard was run.
3. On the **Installation Type** page, click **Next**.
4. On the **Server Selection** page, click **Next**.
5. On the **Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
6. Click **Next** until you have accepted all of the default web server settings, and then click **Install**.
7. Verify that all installations were successful, and then click **Close**.

Additional Technical Resources

For more information about the technologies in this guide, see the following resources:

Windows Server 2012 Technical Library Resources

- Active Directory Domain Services overview at http://technet.microsoft.com/library/hh831484.aspx.
- Domain Name System (DNS) overview at http://technet.microsoft.com/library/hh831667.aspx.
- Implementing the DNS Admins Role at http://technet.microsoft.com/library/cc756152(WS.10).aspx.
- Dynamic Host Configuration Protocol (DHCP) overview at http://technet.microsoft.com/library/hh831825.aspx.
- Network Policy and Access Services overview at http://technet.microsoft.com/library/hh831683.aspx.
- Web Server (IIS) overview at http://technet.microsoft.com/library/hh831725.aspx.

Windows Server® 2008 and Windows Server® 2008 R2 Technical Library Resources

- Active Directory Domain Services at http://go.microsoft.com/fwlink/?LinkId=96418
- Domain Name System (DNS) at http://go.microsoft.com/fwlink/?LinkId=110949
- Dynamic Host Configuration Protocol (DHCP) at http://go.microsoft.com/fwlink/?LinkId=96419
- Network Policy Server (NPS) at http://go.microsoft.com/fwlink/?LinkId=104545
- TCP/IP at http://go.microsoft.com/fwlink/?LinkId=103329

Appendices A through E

The following sections contain additional configuration information for computers that are running operating systems other than Windows Server 2012 and Windows 8, as well as a network preparation worksheet to assist you with your deployment.

1. Appendix A - Renaming computers
2. Appendix B - Configuring static IP addresses
3. Appendix C – Joining computers to the domain
4. Appendix D – Log on to the domain
5. Appendix E - Core Network Planning Preparation Sheet

Appendix A - Renaming computers

You can use the procedures in this section to provide computers running Windows Server® 2008 R2, Windows® 7, Windows Server® 2008, Windows Vista®, Windows Server 2003, and Windows XP with a different computer name.

- Windows Server 2008 R2 and Windows 7
- Windows Server 2008 and Windows Vista
- Windows Server 2003 and Windows XP

Windows Server 2008 R2 and Windows 7

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

To rename computers running Windows Server 2008 R2 and Windows 7

1. Click **Start**, right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
2. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.

   Note
   On computers running Windows 7, before the **System Properties** dialog box opens, the **User Account Control** dialog box opens, requesting permission to continue. Click **Continue** to proceed.
3. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
4. In **Computer Name**, type the name for your computer. For example, if you want to name the computer DC1, type **DC1**.
5. Click **OK** twice, click **Close**, and then click **Restart Now** to restart the computer.

Windows Server 2008 and Windows Vista

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

To rename computers running Windows Server 2008 and Windows Vista

1. Click **Start**, right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
2. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.

   Note
   On computers running Windows Vista, before the **System Properties** dialog box opens, the **User Account Control** dialog box opens, requesting permission to continue. Click **Continue** to proceed.
3. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
4. In **Computer Name**, type the name for your computer. For example, if you want to name the computer DC1, type **DC1**.
5. Click **OK** twice, click **Close**, and then click **Restart Now** to restart the computer.

Windows Server 2003 and Windows XP

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

To rename computers running Windows Server 2003 and Windows XP

1. Click **Start**, right-click **My Computer**, and then click **Properties**. The **System Properties** dialog box opens.
2. Click **Computer Name**, and then click **Change**. The **Computer Name Changes** dialog box opens.
3. In **Computer name**, type the name for your computer. For example, if you want the computer named Client-01, type **Client-01**.
4. Click **OK**. The **System Setting Changes** dialog box opens, indicating that you must restart the computer before the changes take effect.
5. Click **OK**, click **OK** again to close the dialog box, and then click **Yes** to restart the computer.

Appendix B - Configuring static IP addresses

This topic provides procedures for configuring static IP addresses on computers running the following operating systems:

- Windows Server 2008 R2

- Windows Server 2008
- Windows Server 2003

Windows Server 2008 R2

Membership in **Administrators**, or equivalent, is the minimum required to perform this procedure.

To configure a static IP address on a computer running Windows Server 2008 R2

1. Click **Start**, and then click **Control Panel**.
2. In **Control Panel**, click **Network and Internet**. **Network and Internet** opens.

   In **Network and Internet**, click **Network and Sharing Center**. **Network and Sharing Center** opens.
3. In **Network and Sharing Center**, click **Change adapter settings**. **Network Connections** opens.
4. In **Network Connections**, right-click the network connection that you want to configure, and then click **Properties**.
5. In **Local Area Connection Properties**, in **This connection uses the following items**, select **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**. The **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box opens.
6. In **Internet Protocol Version 4 (TCP/IPv4) Properties**, on the **General** tab, click **Use the following IP address**. In **IP address**, type the IP address that you want to use.
7. Press tab to place the cursor in **Subnet mask**. A default value for subnet mask is entered automatically. Either accept the default subnet mask, or type the subnet mask that you want to use.
8. In **Default gateway**, type the IP address of your default gateway.
9. In **Preferred DNS server**, type the IP address of your DNS server. If you plan to use the local computer as the preferred DNS server, type the IP address of the local computer.
10. In **Alternate DNS Server**, type the IP address of your alternate DNS server, if any. If you plan to use the local computer as an alternate DNS server, type the IP address of the local computer.
11. Click **OK**, and then click **Close**.

Windows Server 2008

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

To configure a static IP address on a computer running Windows Server 2008

1. Click **Start**, and then click **Control Panel**.
2. In **Control Panel**, verify that **Classic View** is selected, and then double-click **Network and Sharing Center**.

3. In **Network and Sharing Center**, in **Tasks**, click **Manage Network Connections**.
4. In **Network Connections**, right-click the network connection that you want to configure, and then click **Properties**.
5. In **Local Area Connection Properties**, in **This connection uses the following items**, select **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**. The **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box opens.
6. In **Internet Protocol Version 4 (TCP/IPv4) Properties**, on the **General** tab, click **Use the following IP address**. In **IP address**, type the IP address that you want to use.
7. Press tab to place the cursor in **Subnet mask**. A default value for subnet mask is entered automatically. Either accept the default subnet mask, or type the subnet mask that you want to use.
8. In **Default gateway**, type the IP address of your default gateway.
9. In **Preferred DNS server**, type the IP address of your DNS server. If you plan to use the local computer as the preferred DNS server, type the IP address of the local computer.
10. In **Alternate DNS Server**, type the IP address of your alternate DNS server, if any. If you plan to use the local computer as an alternate DNS server, type the IP address of the local computer.
11. Click **OK**, and then click **Close**.

Windows Server 2003

Membership in **Administrators**, or equivalent, is the minimum required to perform these procedures.

To configure a static IP address on a computer running Windows Server 2003

1. Click **Start**, click **Control Panel**, right-click **Network Connections**, and then click **Open**.
2. In **Network Connections**, right-click the network connection that you want to configure, and then click **Properties**.
3. In **Local Area Connection Properties**, in **This Connection uses the following Items**, select **Internet Protocol (TCP/IP)**, and then click **Properties**. The **Internet Protocol (TCP) Properties** dialog box opens.
4. In **Internet Protocol Version 4 (TCP/IPv4) Properties**, on the **General** tab, click **Use the following IP address**. In **IP address**, type the IP address that you want to use.
5. In **Subnet mask**, either accept the default subnet mask, or type the subnet mask that you want to use.
6. In **Default gateway**, type the IP address of your default gateway.
7. In **Preferred DNS server**, type the IP address of your DNS server.
8. In **Alternate DNS Server**, type the IP address of your alternate DNS server, if any.
9. Click **OK**, and then click **Close**.

Appendix C – Joining computers to the domain

You can use these procedures to join computers running Windows Server® 2008 R2, Windows® 7 Windows Server® 2008, Windows Vista®, and Windows Server 2003 to the domain.

- Windows Server 2008 R2 and Windows 7
- Windows Server 2008 and Windows Vista
- Windows Server 2003 and Windows XP

Important

To join a computer to a domain, you must be logged on to the computer with the local Administrator account or, if you are logged on to the computer with a user account that does not have local computer administrative credentials, you must provide the credentials for the local Administrator account during the process of joining the computer to the domain. In addition, you must have a user account in the domain to which you want to join the computer. During the process of joining the computer to the domain, you will be prompted for your domain account credentials (user name and password).

Windows Server 2008 R2 and Windows 7

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

To join computers running Windows Server 2008 R2 and Windows 7 to the domain

1. Log on to the computer with the local Administrator account.
2. Click **Start**, right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
3. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.

   Note

   On computers running Windows® 7, before the **System Properties** dialog box opens, the **User Account Control** dialog box opens, requesting permission to continue. Click **Continue** to proceed.
4. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
5. In **Computer Name**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is corp.contoso.com, type **corp.contoso.com**.
6. Click **OK**. The **Windows Security** dialog box opens.
7. In **Computer Name/Domain Changes**, in **User name**, type the user name, and in **Password**, type the password, and then click **OK**. The **Computer Name/Domain Changes** dialog box opens, welcoming you to the domain. Click **OK**.
8. The **Computer Name/Domain Changes** dialog box displays a message indicating that you must restart the computer to apply the changes. Click **OK**.
9. On the **System Properties** dialog box, on the **Computer Name** tab, click **Close**. The

Microsoft Windows dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes. Click **Restart Now**.

Windows Server 2008 and Windows Vista

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

To join computers running Windows Server 2008 and Windows Vista to the domain

1. Log on to the computer with the local Administrator account.
2. Click **Start**, right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
3. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.

   Note
   On computers running Windows® 7, before the **System Properties** dialog box opens, the **User Account Control** dialog box opens, requesting permission to continue. Click **Continue** to proceed.
4. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
5. In **Computer Name**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is corp.contoso.com, type **corp.contoso.com**.
6. Click **OK**. The **Windows Security** dialog box opens.
7. In **Computer Name/Domain Changes**, in **User name**, type the user name, and in **Password**, type the password, and then click **OK**. The **Computer Name/Domain Changes** dialog box opens, welcoming you to the domain. Click **OK**.
8. The **Computer Name/Domain Changes** dialog box displays a message indicating that you must restart the computer to apply the changes. Click **OK**.
9. On the **System Properties** dialog box, on the **Computer Name** tab, click **Close**. The **Microsoft Windows** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes. Click **Restart Now**.

Windows Server 2003 and Windows XP

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

To join computers running Windows Server 2003 and Windows XP to the domain

1. Click **Start**, right-click **My Computer**, and then click **Properties**. The **System Properties** dialog box opens.
2. Click **Change**. The **Computer Name Changes** dialog box opens.
3. In **Computer Name Changes**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is corp.contoso.com, type

corp.contoso.com.

4. Click **OK**. The **Computer Name Changes** dialog box opens. In **User name**, type the domain administrator account name, and in **Password**, type the administrator password, and then click **OK**.
5. The **Computer Name Changes** dialog box opens, welcoming you to the domain.
6. Click **OK**. The **Computer Name Changes** dialog box displays a message indicating that you must restart the computer to apply the changes.
7. Click **OK**.
8. On the **System Properties** dialog box, on the **Computer Name** tab, click **OK**, to close the **System Properties** dialog box. The **System Settings Change** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes.
9. Click **Yes**.

Appendix D – Log on to the domain

You can use these procedures to log on to the domain using computers running Windows Server 2008 R2, Windows 7, Windows Server® 2008, Windows Vista®, and Windows Server 2003.

- Windows Server 2008 R2 and Windows 7
- Windows Server 2008 and Windows Vista
- Windows Server 2003 and Windows XP

Windows Server 2008 R2 and Windows 7

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

Log on to the domain using computers running Windows Server 2008 R2 and Windows 7

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click **Switch User**, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain corp.contoso.com with an account named **User-01**, type **CORP\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Windows Server 2008 and Windows Vista

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

Log on to the domain using computers running Windows Server 2008 and Windows Vista

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click **Switch User**, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain corp.contoso.com with an account named **User-01**, type **CORP\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Windows Server 2003 and Windows XP

Membership in **Domain Users**, or equivalent, is the minimum required to perform this procedure.

Log on to the domain using computers running Windows Server 2003 and Windows XP

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The **Log On to Windows** dialog box appears.
3. If **Log on to** is not displayed, click **Options**.
4. In **Log on to**, in the drop down list, select your domain. For example, in the corp.contoso.com domain, select EXAMPLE.
5. Type your domain and user name in the format *domain\user*. For example, to log on to the corp.contoso.com domain with an account named **User-01**, type **CORP\User-01**.
6. In **Password**, type your domain password, and then press ENTER.

Appendix E - Core Network Planning Preparation Sheet

You can use this Network Planning Preparation Sheet to gather the information required to install a core network. This topic provides tables that contain the individual configuration items for each server computer for which you must supply information or specific values during the installation or configuration process. Example values are provided for each configuration item.

For planning and tracking purposes, spaces are provided in each table for you to enter the values used for your deployment. If you log security-related values in these tables, you should store the information in a secure location.

The following links lead to the sections in this topic that provide configuration items and example values that are associated with the deployment procedures presented in this guide.

1. Installing Active Directory Domain Services and DNS
   - Configuring a DNS Reverse Lookup Zone
2. Installing DHCP

- Creating an exclusion range in DHCP
- Creating a new DHCP scope

3. Installing Network Policy Server (optional)

Installing Active Directory Domain Services and DNS

The tables in this section list configuration items for pre-installation and installation of Active Directory Domain Services (AD DS) and DNS.

Pre-installation configuration items for AD DS and DNS

The following tables list pre-installation configuration items as described in Configuring All Servers:

- Configure a static IP address

Configuration items	Example values	Values
IP address	10.0.0.2	
Subnet mask	255.255.255.0	
Default gateway	10.0.0.1	
Preferred DNS server	127.0.0.1	
Alternate DNS server	10.0.0.15	

- Rename the computer

Configuration item	Example value	Value
Computer name	DC1	

AD DS and DNS installation configuration items

Configuration items for the Windows Server Core Network deployment procedure Install AD DS and DNS for a New Forest:

Configuration items	Example values	Values
Full DNS name	corp.contoso.com	
Forest functional level	Windows Server 2003	
Active Directory Domain Services database folder location	E:\Configuration\ Or accept the default location.	
Active Directory Domain	E:\Configuration\	

Configuration items	Example values	Values
Services log files folder location	Or accept the default location.	
Active Directory Domain Services SYSVOL folder location	E:\Configuration\ Or accept the default location	
Directory Restore Mode Administrator password	J*p2leO4$F	
Answer file name (optional)	AD DS_AnswerFile	

Configuring a DNS Reverse Lookup Zone

Configuration items	Example values	Values
Zone type:	• Primary zone • Secondary zone • Stub zone	
Zone type **Store the zone in Active Directory**	• Selected • Not selected	
Active Directory zone replication scope	• To all DNS servers in this forest • To all DNS servers in this domain • To all domain controllers in this domain • To all domain controllers specified in the scope of this directory partition	
Reverse lookup zone name (IP type)	• IPv4 Reverse Lookup Zone • IPv6 Reverse Lookup Zone	
Reverse lookup zone name (network ID)	10.0.0	

Installing DHCP

The tables in this section list configuration items for pre-installation and installation of DHCP.

Pre-installation configuration items for DHCP

The following tables list pre-installation configuration items as described in Configuring All Servers:

- Configure a static IP address

Configuration items	Example values	Values
IP address	10.0.0.3	
Subnet mask	255.255.255.0	
Default gateway	10.0.0.1	
Preferred DNS server	10.0.0.2	
Alternate DNS server	10.0.0.15	

- Rename the computer

Configuration item	Example value	Value
Computer name	DHCP1	

DHCP installation configuration items

Configuration items for the Windows Server Core Network deployment procedure Install Dynamic Host Configuration Protocol (DHCP):

Configuration items	Example values	Values
Network connect bindings	Ethernet	
DNS server settings	DC1	
Preferred DNS server IP address	10.0.0.2	
Alternate DNS server IP address	10.0.0.15	
Scope name	Corp1	
Starting IP address	10.0.0.1	
Ending IP address	10.0.0.254	
Subnet mask	255.255.255.0	
Default gateway (optional)	10.0.0.1	

Configuration items	Example values	Values
Lease duration	8 days	
IPv6 DHCP server operation mode	Not enabled	

Creating an exclusion range in DHCP

Configuration items to create an exclusion range while creating a scope in DHCP.

Configuration items	Example values	Values
Scope name	Corp1	
Scope description	Main office subnet 1	
Exclusion range start IP address	10.0.0.1	
Exclusion range end IP address	10.0.0.15	

Creating a new DHCP scope

Configuration items for the Windows Server Core Network deployment procedure Create and Activate a New DHCP Scope:

Configuration items	Example values	Values
New scope name	Corp2	
Scope description	Main office subnet 2	
(IP address range) Start IP address	10.0.1.1	
(IP address range) End IP address	10.0.1.254	
Length	8	
Subnet mask	255.255.255.0	
(Exclusion range) Start IP address	10.0.1.1	
Exclusion range end IP address	10.0.1.15	
Lease duration	• 8	

Configuration items	Example values	Values
Days Hours Minutes	• 0 • 0	
Router (default gateway) IP address	10.0.1.1	
DNS parent domain	corp.contoso.com	
DNS server IP address	10.0.0.2	

Installing Network Policy Server (optional)

The tables in this section list configuration items for pre-installation and installation of NPS.

Pre-installation configuration items

The following three tables list pre-installation configuration items as described in Configuring All Servers:

- Configure a static IP address

Configuration items	Example values	Values
IP address	10.0.0.4	
Subnet mask	255.255.255.0	
Default gateway	10.0.0.1	
Preferred DNS server	10.0.0.2	
Alternate DNS server	10.0.0.15	

- Rename the computer

Configuration item	Example value	Value
Computer name	NPS1	

Network Policy Server installation configuration items

Configuration items for the Windows Server Core Network NPS deployment procedures Install Network Policy Server (NPS) and Register the NPS Server in the Default Domain.

- No additional configuration items are required to install and register NPS.

Core Network Companion Guide: Server Certificate Deployment

The Windows Server 2012 Core Network Guide provides instructions for planning and deploying the core components required for a fully functioning network and a new Active Directory® domain in a new forest.

This guide explains how to build on the core network by providing instructions for deploying server certificates for computers that are running Network Policy Server (NPS), Routing and Remote Access Service (RRAS), or both.

Tip

This guide is available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/p/?linkid=255200).

This guide contains the following sections.

- Prerequisites for using this guide
- About this guide
- What this guide does not provide
- Technology overviews
- Server Certificate Deployment Overview
- Server Certificate Deployment Planning
- Server Certificate Deployment
- Additional Resources

Prerequisites for using this guide

This is a companion guide to the Windows Server 2012 Core Network Guide. To deploy server certificates with this guide, you must first do the following.

- Deploy a core network using the Core Network Guide, or already have the technologies provided in the Core Network Guide installed and functioning correctly on your network. These technologies include TCP/IP v4, DHCP, Active Directory Domain Services (AD DS), DNS, NPS, and Web Server (IIS).

  Notes

  The Windows Server 2012 Core Network Guide is available in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154884).

  The Core Network Guide is also available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=157742).

About this guide

This guide provides instructions for deploying server certificates to servers running NPS, RRAS, or both, by using AD CS in Windows Server 2012.

Server certificates are required when you deploy certificate-based authentication methods with Extensible Authentication Protocol (EAP) and Protected EAP (PEAP) for network access authentication.

Deploying server certificates with Active Directory Certificate Services (AD CS) for EAP and PEAP certificate-based authentication methods provides the following benefits:

- Binding the identity of the server running NPS or the RRAS server to a private key
- A cost-effective and secure method for automatically enrolling certificates to domain member NPS and RRAS servers
- An efficient method for managing certificates and certification authorities (CAs)
- Security provided by certificate-based authentication
- The ability to expand the use of certificates for additional purposes

This guide is designed for network and system administrators who have followed the instructions in the Windows Server 2012 Core Network Guide to deploy a core network, or for those who have previously deployed the technologies included in the Core Network Guide, including Active Directory Domain Services (AD DS), Domain Name Service (DNS), Dynamic Host Configuration Protocol (DHCP), TCP/IP, Web Server (IIS), and Network Policy Server (NPS).

Important

This guide, which provides instructions for deploying server certificates using an online Enterprise Root certification authority (CA), is designed for small organizations that have limited computing resources. For security reasons - if your organization has the computing resources - it is recommended that you deploy an offline Enterprise Root CA in a two tier public key infrastructure (PKI). For more information, see Additional Resources.

It is recommended that you review the design and deployment guides for each of the technologies that are used in this deployment scenario. These guides can help you determine whether this deployment scenario provides the services and configuration that you need for your organization's network.

Requirements for deploying server certificates

Following are the requirements for using certificates:

- To deploy server certificates by using autoenrollment, AD CS requires the Windows Server 2012 Standard, Enterprise, or Datacenter operating systems. AD DS must be installed before AD CS is installed. Although AD CS can be deployed on a single server, many deployments involve multiple servers configured as CAs.
- To provide computers with access to the Authority Information Access (AIA) and certificate revocation list (CRL) that is generated by your certification authority, you must have a Web server that is properly configured according to the instructions in this guide.

- To deploy PEAP or EAP for virtual private networks (VPNs), you must deploy RRAS configured as a VPN server. The use of NPS is optional; however, if you have multiple VPN servers, using NPS is recommended for ease of administration and for the RADIUS accounting services that NPS provides.
- To deploy PEAP or EAP for Remote Desktop Gateway (RD Gateway), you must deploy RD Gateway and NPS.

  Note
  In previous versions of Windows Server, Remote Desktop Services was named Terminal Services.
- To deploy PEAP or EAP for 802.1X secure wired or wireless, you must deploy NPS and additional hardware, such as 802.1X-capable switches and wireless access points.
- To deploy certificate-based authentication methods that require certificates for user and computer authentication in addition to requiring certificates for server authentication, such as EAP with Transport Layer Security (EAP-TLS) or PEAP-TLS, you must also deploy user or computer certificates through autoenrollment or by using smart cards.

What this guide does not provide

This guide does not provide comprehensive instructions for designing and deploying a public key infrastructure (PKI) by using AD CS. It is recommended that you review AD CS documentation and PKI design documentation before deploying the technologies in this guide. For more information, see the Additional Resources section later in this document.

This guide does not provide instructions on how to install Web Server (IIS) or Network Policy Server technologies on server computers; those instructions are provided in the Core Network Guide.

This guide also does not provide detailed instructions for deploying the network access technologies for which server certificates can be used.

Technology overviews

Following are technology overviews for EAP, PEAP, and AD CS.

EAP

Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by allowing arbitrary authentication methods that use credential and information exchanges of arbitrary lengths. EAP was developed in response to an increasing demand for authentication methods that use security devices such as smart cards, token cards, and crypto calculators. EAP provides an industry-standard architecture for supporting additional authentication methods within PPP.

With EAP, an arbitrary authentication mechanism is used to verify the identities of the client and server that are establishing a network access connection. The exact authentication scheme to be used is negotiated by the access client and the authenticator - the network access server or the Remote Authentication Dial-In User Service (RADIUS) server.

With EAP authentication, both the network access client and the authenticator (such as the server running NPS) must support the same EAP type for successful authentication to occur.

Important
Strong EAP types, such as those that are based on certificates, offer better security against brute-force attacks, dictionary attacks, and password-guessing attacks than password-based authentication protocols, such as CHAP or MS-CHAP, version 1.

EAP in Windows Server 2012

Windows Server 2012 includes an EAP infrastructure, EAP types, and the ability to pass EAP messages to a RADIUS server (EAP-RADIUS) such as NPS.

By using EAP, you can support additional authentication schemes, known as EAP types. The EAP types that are supported by Windows Server 2012 are:

- **Transport Layer Security (TLS)**. EAP-TLS requires the use of computer certificates or user certificates, in addition to server certificates that are enrolled to computers running NPS.
- **Microsoft Challenge-Handshake Authentication Protocol, version 2 (MS-CHAP v2)**. This EAP type is a password-based authentication protocol. When used within EAP as the authentication method EAP-MS-CHAP v2, NPS and RRAS servers provide a server certificate as proof of identity to client computers, while users prove their identity with a user name and password.
- **Tunneled Transport Layer Security (TTLS)**. EAP-TTLS is new in Windows Server 2012 and is not available in other versions of Windows Server. EAP-TTLS is a standards-based EAP tunneling method that supports mutual authentication. EAP-TTLS provides a secure tunnel for client authentication using EAP methods and other legacy protocols. EAP-TTLS also provides you with the ability to configure EAP-TTLS on client computers for network access solutions in which non-Microsoft Remote Authentication Dial In User Service (RADIUS) servers that support EAP-TTLS are used for authentication.

In addition, you can install other non-Microsoft EAP modules on the server running NPS or Routing and Remote Access to provide other EAP authentication types. In most cases, if you install additional EAP types on servers, you must also install matching EAP client authentication components on client computers so that the client and server can successfully negotiate an authentication method to use for connection requests.

PEAP

PEAP uses TLS to create an encrypted channel between an authenticating PEAP client, such as a wireless computer, and a PEAP authenticator, such as a server running NPS or other RADIUS server.

PEAP does not specify an authentication method, but it provides additional security for other EAP authentication protocols (such as EAP-MSCHAP v2) that can operate through the TLS-encrypted channel provided by PEAP. PEAP is used as an authentication method for access clients that are connecting to your organization's network through the following types of network access servers:

- 802.1X-capable wireless access points

- 802.1X-capable authenticating switches
- Computers running Windows Server 2012 or Windows Server 2008 R2 and RRAS that are configured as VPN servers
- Computers running Windows Server 2012 or Windows Server 2008 R2 and RD Gateway

Features of PEAP

To enhance the EAP protocols and network security, PEAP provides:

- A TLS channel that provides protection for the EAP method negotiation that occurs between the client and server. This TLS channel helps prevent an attacker from injecting packets between the client and the network access server to cause the negotiation of a less secure EAP type. The encrypted TLS channel also helps prevent denial of service attacks against the server running NPS.
- Support for the fragmentation and reassembly of messages, which allows the use of EAP types that do not provide this functionality.
- Clients with the ability to authenticate the NPS or other RADIUS server. Because the server also authenticates the client, mutual authentication occurs.
- Protection against the deployment of an unauthorized wireless access point at the moment when the EAP client authenticates the certificate provided by the server running NPS. In addition, the TLS master secret that is created by the PEAP authenticator and the client is not shared with the access point. Because of this, the access point cannot decrypt the messages that are protected by PEAP.
- PEAP fast reconnect, which reduces the delay between an authentication request by a client and the response by the NPS or other RADIUS server. Fast reconnect also allows wireless clients to move between access points that are configured as RADIUS clients to the same RADIUS server without repeated requests for authentication. This reduces resource requirements for the client and the server, and it minimizes the number of times that users are prompted for credentials.

Active Directory Certificate Services

AD CS in Windows Server 2012 provides customizable services for creating and managing the X.509 certificates that are used in software security systems that employ public key technologies. Organizations can use AD CS to enhance security by binding the identity of a person, device, or service to a corresponding public key. AD CS also includes features that allow you to manage certificate enrollment and revocation in a variety of scalable environments.

Server Certificate Deployment Overview

This topic contains the following sections.

- Server certificate deployment components
- Server certificate deployment process overview

Server certificate deployment components

You can use this guide to install Active Directory® Certificate Services (AD CS) as an Enterprise root certification authority (CA) and to enroll a server certificate to servers running Network Policy Server (NPS), Routing and Remote Access service (RRAS), or both NPS and RRAS.

If you deploy certificate-based authentication, servers running NPS and RRAS are required to use a server certificate to prove their identities to client computers that are attempting to connect to the network.

The following illustration shows the components that are required to deploy server certificates to your NPS server.

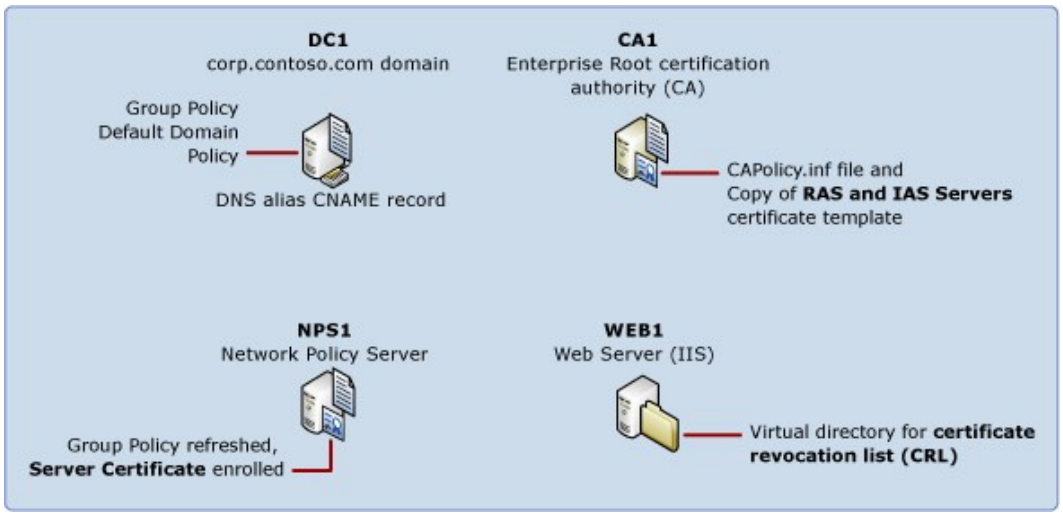

Note

In the illustration above, four servers are depicted: DC1, NPS1, WEB1, and CA1. This guide provides instructions for deploying and configuring CA1, and for configuring the other three servers, which this guide assumes you have already installed on your network by using the Core Network Guide.

For more information on each item depicted in the illustration above, see the following:

- CA1 running the AD CS server role
- WEB1 running the Web Services (IIS) server role
- DC1 running the AD DS and DNS server roles
- NPS1 running the Network Policy Server role service of the Network Policy and Access Services server role

CA1 running the AD CS server role

The enterprise root certification authority (CA) is also an issuing CA. The CA issues certificates to computers and users that have the correct security permissions to enroll a certificate. Active Directory Certificate Services (AD CS) is installed on CA1.

In this scenario, the enterprise root CA is also an issuing CA. For larger networks or where security concerns provide justification, you can separate the roles of root CA and issuing CA, and deploy subordinate CAs that are issuing CAs.

In the most secure deployments, the enterprise root CA is taken offline and physically secured. For more information, see Additional Resources.

CAPolicy.inf

Before you install AD CS, you configure the CAPolicy.inf file with specific settings for your deployment.

Copy of the RAS and IAS servers certificate template

When you deploy server certificates, you make a copy of the **RAS and IAS servers** certificate template and then configure the template according to your requirements and the instructions in this guide. You will be using a copy rather than the original so that the configuration of the original template is preserved for possible future use. The CA uses the copy of the RAS and IAS servers template to create server certificates that it issues to NPS servers that are members of the RAS and IAS servers group in Active Directory Users and Computers.

Note

NPS servers that you have registered in Active Directory are automatically added to the RAS and IAS servers group.

Additional CA1 configuration

The CA publishes a certificate revocation list (CRL) that computers must check to ensure that certificates that are presented to them as proof of identity are valid certificates and have not been revoked. You must configure your CA with the correct location of the CRL so that computers know where to look for the CRL during the authentication process.

WEB1 running the Web Services (IIS) server role

The Web server is installed when you perform the tasks in the Windows Server 2012 Core Network Guide, so before you perform the tasks in this guide, you should already have a Web server installed on your network. On the computer that is running the Web Server (IIS) server role, WEB1, you must create a folder in Windows Explorer for use as the location for the CRL and AIA.

Virtual directory for the CRL and AIA

After you create a folder in Windows Explorer, you must configure the folder as a virtual directory in Internet Information Services (IIS) Manager, as well as configuring the access control list for the virtual directory to allow computers to access the AIA and CRL after they are published there.

DC1 running the AD DS and DNS server roles

The domain controller and DNS server is installed when you perform the tasks in the Windows Server 2012 Core Network Guide, so before you perform the tasks in this guide, you should already have a domain controller and DNS server installed on your network.

Group Policy default domain policy

After you configure the certificate template on the CA, you can configure the default domain policy in Group Policy so that certificates are autoenrolled to NPS servers. Group Policy is configured in AD DS on the server DC1.

DNS alias (CNAME) resource record

You must create an alias (CNAME) resource record for the Web server to ensure that other computers can find the server, as well as the AIA and the CRL that are stored on the server. In addition, using an alias CNAME resource record provides flexibility so that you can use the Web server for other purposes, such as hosting Web and FTP sites.

NPS1 running the Network Policy Server role service of the Network Policy and Access Services server role

The NPS server is installed when you perform the tasks in the Windows Server 2012 Core Network Guide, so before you perform the tasks in this guide, you should already have an NPS server installed on your network.

Group Policy applied and certificate enrolled to NPS1

After you have configured the certificate template and autoenrollment, you can refresh Group Policy on the NPS server. At this time, the NPS server enrolls the server certificate from CA1.

Server certificate deployment process overview

Note
The details of how to perform these steps are provided in the section Server Certificate Deployment.

The process of configuring NPS and RRAS server certificate enrollment occurs in these stages:

1. On DC1, create an alias (CNAME) record for your Web server, WEB1.
2. Configure your Web server to host the CRL from the CA, then publish the CRL and copy the Enterprise Root CA certificate into the new virtual directory.
3. On the computer where you are planning to install AD CS, assign the computer a static IP address, rename the computer, join the computer to the domain, and then log on to the computer with a user account that is a member of the Domain Admins and Enterprise Admins groups.

4. On the computer where you are planning to install AD CS, configure the CAPolicy.inf file with settings that are specific to your deployment.
5. Install the AD CS server role and perform additional configuration of the CA.
6. Copy the CRL and CA certificate to the share on the Web server.
7. On the CA, configure a server certificate template. The CA issues certificates based on a certificate template, so you must configure the template for the server certificate before the CA can issue a certificate.
8. Configure server certificate autoenrollment in Group Policy. When you configure autoenrollment, all domain member servers that are running NPS, RRAS, or both on your network will automatically receive a server certificate when Group Policy on each server is refreshed. If you add more NPS or RRAS servers later, they will automatically receive a server certificate, too.
9. Refresh Group Policy on servers running NPS and RRAS. When Group Policy is refreshed, the servers receive the server certificate, which is based on the template that you configured in the previous step. This certificate is used by the server to prove its identity to client computers that attempt to connect to your network.

   Note
   All domain member computers automatically receive the Enterprise Root CA's certificate without the configuration of autoenrollment. This certificate is different than the server certificate that you configure and distribute by using autoenrollment. The CA's certificate is automatically installed in the Trusted Root Certification Authorities certificate store for all domain member computers so that they will trust certificates that are issued by this CA. For example, if you deploy EAP-TLS, client computers use a certificate to prove their identities to the NPS server. When the NPS server receives a certificate from a client computer as proof of the client computer's identity, trust for the certificate is established because NPS has the issuing CA certificate in its own Trusted Root Certification Authorities certificate store.
10. Verify that the NPS server enrolled a valid server certificate.

Server Certificate Deployment Planning

Before you deploy server certificates, you must plan the following items:

- Plan basic server configuration
- Plan domain access
- Plan the location and name of the virtual directory on your Web server
- Plan a DNS alias (CNAME) record for your Web server
- Plan configuration of CAPolicy.inf
- Plan configuration of the CDP and AIA extensions on CA1
- Plan the copy operation between the CA and the Web server
- Plan the configuration of the server certificate template on the CA

Plan basic server configuration

After you install Windows Server 2012 on your certification authority, you must rename the computer and assign and configure a static IP address for the local computer.

For more information, see the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library.

Plan domain access

To log on to the domain, the computer must be a domain member computer and the user account must be created in AD DS before the logon attempt. In addition, most procedures in this guide require that the user account is a member of the Enterprise Admins or Domain Admins groups in Active Directory Users and Computers, so you must log on to the CA with an account that has the appropriate group membership.

For more information, see the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library.

Plan the location and name of the virtual directory on your Web server

To provide access to the CRL and the CA certificate to other computers, you must store these items in a virtual directory on your Web server. In this guide, the virtual directory is located on the Web server WEB1. This folder is on the “C:” drive and is named “pki.” You can locate your virtual directory on your Web server at a folder location that works for your deployment.

Plan a DNS alias (CNAME) record for your Web server

Alias (CNAME) resource records are also sometimes called canonical name resource records. With these records, you can use more than one name to point to a single host, making it easy to do such things as host both a File Transfer Protocol (FTP) server and a Web server on the same computer. For example, the well-known server names (ftp, www) are registered using alias (CNAME) resource records that map to the Domain Name System (DNS) host name, such as WEB1, for the server computer that hosts these services.

This guide provides instructions for configuring your Web server to host the certificate revocation list (CRL) for your certification authority (CA). Because you might also want to use your Web server for other purposes, such as to host an FTP or Web site, it’s a good idea to create an alias resource record in DNS for your Web server. In this guide, the CNAME record is named “pki,” but you can choose a name that is appropriate for your deployment.

Plan configuration of CAPolicy.inf

Before you install AD CS, you must configure CAPolicy.inf on the CA with information that is correct for your deployment. A CAPolicy.inf file contains the following information:

```
[Version]

Signature="$Windows NT$"

[PolicyStatementExtension]

Policies=InternalPolicy

[InternalPolicy]

OID=1.2.3.4.1455.67.89.5

Notice="Legal Policy Statement"

URL=http://pki.corp.contoso.com/pki/cps.txt

[Certsrv_Server]

RenewalKeyLength=2048

RenewalValidityPeriod=Years

RenewalValidityPeriodUnits=5

CRLPeriod=weeks

CRLPeriodUnits=1

LoadDefaultTemplates=0

AlternateSignatureAlgorithm=1
```

You must plan the following items for this file:

- **URL**. The example CAPolicy.inf file has a URL value of **http://pki.corp.contoso.com/pki/cps.txt**. This is because the Web server in this guide is named WEB1 and has a DNS CNAME resource record of pki. The Web server is also joined to the corp.contoso.com domain. In addition, there is a virtual directory on the Web server named “pki” where the certificate revocation list is stored. Ensure that the value that you provide for URL in your CAPolicy.inf file points to a virtual directory on your Web server in your domain.
- **RenewalKeyLength**. The default renewal key length for AD CS in Windows Server 2012 is 2048. The key length that you select should be as long as possible while still providing compatibility with the applications that you intend to use.
- **RenewalValidityPeriodUnits**. The example CAPolicy.inf file has a RenewalValidityPeriodUnits value of 5 years. This is because the expected lifespan of the CA is around ten years. The value of RenewalValidityPeriodUnits should reflect the overall validity period of the CA or the highest number of years for which you want to provide enrollment.
- **CRLPeriodUnits**. The example CAPolicy.inf file has a CRLPeriodUnits value of 1. This is because the example refresh interval for the certificate revocation list in this guide is 1 week. At the interval value that you specify with this setting, you must publish the CRL on the CA to

the Web server virtual directory where you store the CRL and provide access to it for computers that are in the authentication process.

- **AlternateSignatureAlgorithm**. This CAPolicy.inf implements an improved security mechanism by implementing alternate signature formats. You should not implement this setting if you still have Windows XP clients that require certificates from this CA. For more information, see Additional Resources.

If you do not plan on adding any subordinate CAs to your public key infrastructure at a later time, and if you want to prevent the addition of any subordinate CAs, you can add the PathLength key to your CAPolicy.inf file with a value of 0. To add this key, copy and paste the following code into your file:

```
[BasicConstraintsExtension]

PathLength=0

Critical=Yes
```

Important
It is not recommended that you change any other settings in the CAPolicy.inf file unless you have a specific reason for doing so.

Plan configuration of the CDP and AIA extensions on CA1

When you configure the Certificate Revocation List (CRL) Distribution Point (CDP) and the Authority Information Access (AIA) settings on CA1, you need the name of your Web server and your domain name. You also need the name of the virtual directory that you create on your Web server where the certificate revocation list (CRL) and the certification authority certificate are stored.

The CDP location that you must enter during this deployment step has the format:

http://*DNSAlias(CNAME)RecordName*.*Domain*.com/*VirtualDirectoryName*/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl.

For example, if your Web server is named WEB1 and your DNS alias CNAME record for the Web server is “pki,” your domain is corp.contoso.com, and your virtual directory is named pki, the CDP location is:

http://pki.corp.contoso.com/pki/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl

The AIA location that you must enter has the format:

http://*DNSAlias(CNAME)RecordName*.*Domain*.com/*VirtualDirectoryName*/<ServerDNSName>_<CaName><CertificateName>.crt.

For example, if your Web server is named WEB1 and your DNS alias CNAME record for the Web server is “pki,” your domain is corp.contoso.com, and your virtual directory is named pki, the AIA location is:

http://pki.corp.contoso.com/pki/<ServerDNSName>_<CaName><CertificateName>.crt

Plan the copy operation between the CA and the Web server

To publish the CRL from the CA to the Web server virtual directory, you can run the certutil –crl command after you configure the CDP and AIA locations on the CA. Ensure that you configure the correct paths on the CA Properties **Extensions** tab before you run this command using the instructions in this guide. In addition, to copy the Enterprise CA certificate to the Web server, you must have already created the virtual directory on the Web server and configured the folder as a shared folder.

Plan the configuration of the server certificate template on the CA

To deploy autoenrolled server certificates, you must copy the certificate template named **RAS and IAS Server**. By default, this copy is named **Copy of RAS and IAS Server**. If you want to rename this template copy, plan the name that you want to use during this deployment step.

Note
The last three deployment sections in this guide – which allow you to configure server certificate autoenrollment, refresh Group Policy on the NPS server, and verify that the NPS server has received a valid server certificate from the CA – do not require additional planning steps.

Server Certificate Deployment

Follow these steps to install an enterprise root certification authority (CA) and to deploy server certificates for use with PEAP and EAP.

Important

- To perform the steps in this guide, you must already have a server that is running Network Policy Server (NPS) and a server that is running the Web Server (IIS) server role. If you do not have an NPS server and a Web server, you cannot complete all of the steps in this guide. You must install these servers before performing the steps in this guide. For more information on how to accomplish these tasks, see the Windows Server® 2012 Core Network Guide at http://technet.microsoft.com/en-us/library/hh911995.aspx.
- Before you install Active Directory Certificate Services, you must name the computer, configure the computer with a static IP address, and join the computer to the domain. After you install AD CS, you cannot change the computer name or the domain membership of the computer, however you can change the IP address if needed. For more information on how

to accomplish these tasks, see the Windows Server® 2012 Core Network Guide at http://technet.microsoft.com/en-us/library/hh911995.aspx.

- Create an Alias (CNAME) Record in DNS for WEB1
- Configure WEB1 to Distribute Certificate Revocation Lists (CRLs)
- Prepare the CAPolicy.inf File
- Install the Certification Authority
- Configure the CDP and AIA Extensions on CA1
- Copy the CA Certificate and CRL to the Virtual Directory
- Configure the Server Certificate Template
- Configure Server Certificate Autoenrollment
- Refresh Group Policy
- Verify NPS Server Enrollment of a Server Certificate

Note

The procedures in this guide do not include instructions for cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Create an Alias (CNAME) Record in DNS for WEB1

You can use this procedure to add a CNAME resource record for your Web server to a zone in DNS on your domain controller. When you perform this procedure, replace Alias name and other variables with values that are appropriate for your deployment.

To perform this procedure, you must be a member of Domain Admins.

To add an alias (CNAME) resource record to a zone

1. On DC1, in Server Manager, click **Tools** and then click **DNS**. The DNS Manager Microsoft Management Console (MMC) opens.
2. In the console tree, double-click **Forward Lookup Zones**, right-click the forward lookup zone where you want to add the Alias resource record, and then click **New Alias (CNAME)**. The **New Resource Record** dialog box opens.
3. In **Alias name**, type the alias name **pki**.
4. When you typed the Alias name, the Fully qualified domain name (FQDN) auto-fills in the dialog box. For example, if your alias name is “pki” and your domain is corp.contoso.com, the value **pki.corp.contoso.com** is auto-filled for you.

5. In **Fully qualified domain name (FQDN) for target host**, type the FQDN of the DNS host computer for which this alias is to be used.

   Note
   As an option, you can click **Browse** to search the DNS namespace for hosts in this domain that have host (A) resource records already defined.

6. Click **OK** to add the new record to the zone.

Configure WEB1 to Distribute Certificate Revocation Lists (CRLs)

You can use this procedure to configure the web server WEB1 to distribute CRLs.

In the extensions of the root CA, it was stated that the CRL from the root CA would be available via http://pki.corp.contoso.com/pki. Currently, there is not a PKI virtual directory on WEB1, so one must be created.

To perform this procedure, you must be a member of Domain Admins.

Note
In the procedure below, replace the user account name, the Web server name, folder names and locations, and other values with those that are appropriate for your deployment.

To configure WEB1 to distribute certificates and CRLs

1. On WEB1, run Windows PowerShell as an administrator, type `explorer c:\`, and then press ENTER. Windows Explorer opens.
2. Create a new folder named PKI on the C: drive. To do so, click **Home**, and then click **New Folder**. Type **pki** and then press ENTER.
3. In Windows Explorer, right-click the folder you just created, hover the mouse cursor over **Share with**, and then click **Specific people**. The **File Sharing** dialog box opens.
4. In **File Sharing**, type **Cert Publishers**, and then click **Add**. The Cert Publishers group is added to the list. In the list, in **Permission Level**, click the arrow next to **Cert Publishers**, and then click **Read/Write**. Click **Share**, and then click **Done**.
5. Close Windows Explorer.
6. Open the IIS console. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
7. In the Internet Information Services (IIS) Manager console tree, expand **WEB1**. If you are

invited to get started with Microsoft Web Platform, click **Cancel**.

8. Expand **Sites** and then right-click the **Default Web Site** and then click **Add Virtual Directory**.
9. In **Alias**, type **pki**. In **Physical path** type **C:\pki**, then click **OK**.
10. Enable Anonymous access to the pki virtual directory, so that any client can check the validity of the CA certificates and CRLs. To do so:
    a. In the **Connections** pane, ensure that **pki** is selected.
    b. On **pki Home** click **Authentication**.
    c. In the **Actions** pane, click **Edit Permissions**.
    d. On the **Security** tab, click **Edit**
    e. On the **Permissions for pki** dialog box, click **Add**.
    f. In the **Select Users, Computers, Service Accounts, or Groups**, type **ANONYMOUS LOGON; Everyone** and then click **Check Names**. Click **OK**.
    g. Click **OK** on the **Select Users, Computers, Service Accounts or Groups** dialog box.
    h. Click **OK** on the **Permissions for pki** dialog box.
11. Click **OK** on the **pki Properties** dialog box.
12. In the **pki Home** pane, double-click **Request Filtering**.
13. The **File Name Extensions** tab is selected by default in the **Request Filtering** pane. In the **Actions** pane, click **Edit Feature Settings**.
14. In **Edit Request Filtering Settings**, select **Allow double escaping** and then click **OK**.
15. In the Internet Information Services (IIS) Manager MMC, click your Web server name. For example, if your Web server is named WEB1, click **WEB1**.
16. In **Actions**, click **Restart**. Internet services are stopped and then restarted.

Prepare the CAPolicy.inf File

On CA1, you must prepare the CAPolicy.inf file before installing Active Directory Certificate Services.

To perform this procedure, you must be a member of the Administrators group.

To prepare the CAPolicy.inf file

1. Open Windows PowerShell, type **notepad c:\Windows\CAPolicy.inf** and press ENTER.
2. When prompted to create a new file, click **Yes**.
3. Enter the following as the contents of the file:

```
[Version]
Signature="$Windows NT$"
[PolicyStatementExtension]
Policies=InternalPolicy
[InternalPolicy]
OID=1.2.3.4.1455.67.89.5
Notice="Legal Policy Statement"
URL=http://pki.corp.contoso.com/pki/cps.txt
[Certsrv_Server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=5
CRLPeriod=weeks
CRLPeriodUnits=1
LoadDefaultTemplates=0
AlternateSignatureAlgorithm=1
[CRLDistributionPoint]
[AuthorityInformationAccess]
```

4. Click **File**, and then click **Save As**. Ensure the following:
   - **File name** is set to **CAPolicy.inf**
   - **Save as type** is set to **All Files**
   - **Encoding** is **ANSI**
5. When you are prompted to overwrite the file, click **Yes**.

Ensure CAPolicy.inf file has appropriate settings

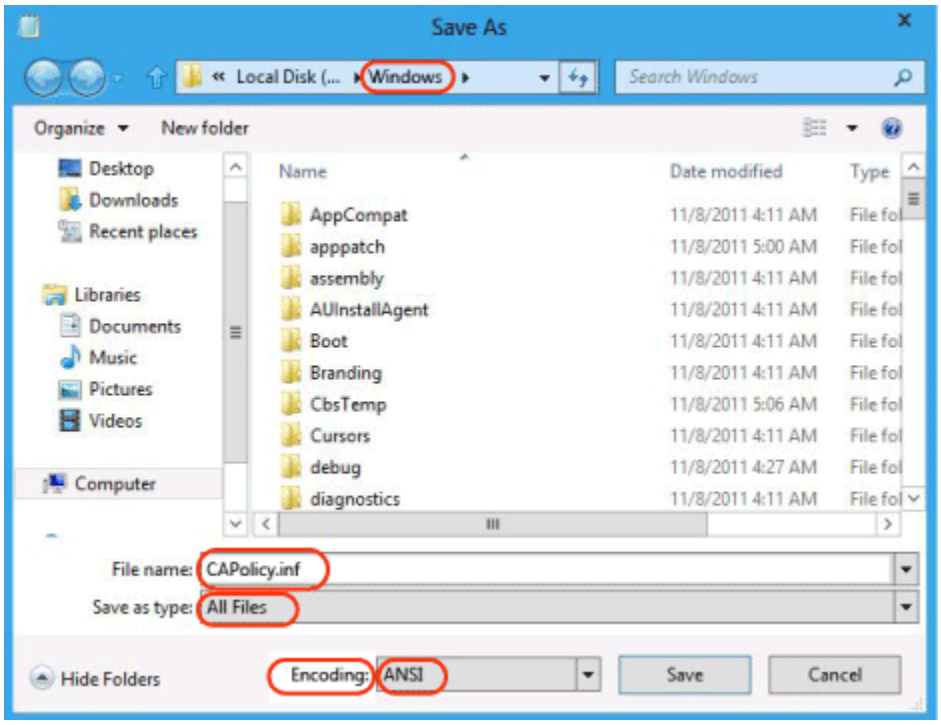

Caution

Be sure to save the CAPolicy.inf with the inf extension. If you do not specifically type **.inf** at the end of the file name and select the options as described, the file will be saved as a text file and will not be used during CA installation.

6. Close Notepad.

Important

In the CAPolicy.inf, you can see there is a line specifying the URL http://pki.corp.contoso.com/pki/cps.txt. The Internal Policy section of the CAPolicy.inf is just shown as an example of how you would specify the location of a certificate practice statement (CPS). In this guide, you are not instructed to create the certificate practice statement (CPS). To learn more about policy statements including CPS, see Additional Resources.

Install the Certification Authority

You can use this procedure to install Active Directory® Certificate Services (AD CS) so that you can enroll a server certificate to servers running Network Policy Server (NPS), Routing and Remote Access Service (RRAS), or both.

Important

- Before you install Active Directory Certificate Services, you must name the computer, configure the computer with a static IP address, and join the computer to the domain. For more information on how to accomplish these tasks, see the Windows Server® 2012 Core Network Guide at http://technet.microsoft.com/en-us/library/hh911995.aspx.
- To perform this procedure, the computer on which you are installing AD CS must be joined to a domain where Active Directory Domain Services (AD DS) is installed. To autoenroll server certificates to computers running NPS and RRAS, you must install AD CS on either the Windows Server 2012 Enterprise or Datacenter operating systems.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

Notes

To perform this procedure by using Windows PowerShell, open Windows PowerShell and type the following command, and then press ENTER. You must also replace the domain name with the name that you want to use.

```
Add-WindowsFeature Adcs-Cert-Authority -IncludeManagementTools
```

After AD CS is installed, type the following command and press ENTER.

```
Install-AdcsCertificationAuthority –CAType EnterpriseRootCA
```

To install Active Directory Certificate Services

1. Log on as a member of both the Enterprise Admins group and the root domain's Domain Admins group.
2. In Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
3. In **Before You Begin**, click **Next**.

   Note

   The **Before You Begin** page of the Add Roles and Features Wizard is not displayed if you have previously selected **Skip this page by default** when the Add Roles and Features Wizard was run.
4. In **Select Installation Type**, ensure that **Role-Based or feature-based installation** is selected, and then click **Next**.
5. In **Select destination server**, ensure that **Select a server from the server pool** is selected. In **Server Pool**, ensure that the local computer is selected. Click **Next**.
6. In **Select Server Roles**, in **Roles**, select **Active Directory Certificate Services**. When you are prompted to add required features, click **Add Features**, and then click **Next**.
7. In **Select features**, click **Next**.
8. In **Active Directory Certificate Services**, read the provided information, and then click **Next**.
9. In **Confirm installation selections**, click **Install**. Do not close the wizard during the installation process. When installation is complete, click **Configure Active Directory Certificate Services on the destination server**. The AD CS Configuration wizard

opens. Read the credentials information and, if needed, provide the credentials for an account that is a member of the Enterprise Admins group. Click **Next**.

10. In **Role Services**, click **Certification Authority**, and then click **Next**.
11. On the **Setup Type** page, verify that **Enterprise CA** is selected, and then click **Next**.
12. On the **Specify the type of the CA** page, verify that **Root CA** is selected, and then click **Next**.
13. On the **Specify the type of the private key** page, verify that **Create a new private key** is selected, and then click **Next**.
14. On the **Cryptography for CA** page, keep the default settings for CSP (**RSA#Microsoft Software Key Storage Provider**) and hash algorithm (**SHA1**), and determine the best key character length for your deployment. Large key character lengths provide optimal security; however, they can impact server performance and might not be compatible with legacy applications. It is recommended that you keep the default setting of 2048. Click **Next**.
15. On the **CA Name** page, keep the suggested common name for the CA or change the name according to your requirements. Ensure that you are certain the CA name is compatible with your naming conventions and purposes, because you cannot change the CA name after you have installed AD CS. Click **Next**.
16. On the **Validity Period** page, in **Specify the validity period**, type the number and select a time value (Years, Months, Weeks, or Days). The default setting of five years is recommended. Click **Next**.
17. On the **CA Database** page, in **Specify the database locations**, specify the folder location for the certificate database and the certificate database log. If you specify locations other than the default locations, ensure that the folders are secured with access control lists (ACLs) that prevent unauthorized users or computers from accessing the CA database and log files. Click **Next**.
18. In **Confirmation**, click **Configure** to apply your selections, and then click **Close**.

Configure the CDP and AIA Extensions on CA1

You can use this procedure to configure the Certificate Revocation List (CRL) Distribution Point (CDP) and the Authority Information Access (AIA) settings on CA1.

To perform this procedure, you must be a member of Domain Admins.

To configure the CDP and AIA extensions on CA1

1. In Server Manager, click **Tools** and then click **Certification Authority**.
2. In the Certification Authority console tree, right-click **corp-CA1-CA**, and then click **Properties**.

   Note
   The name of your CA is different if you did not name the computer CA1 and your domain name is different than the one in this example. The CA name is in the format *domain-CAComputerName*-CA.
3. Click the **Extensions** tab. Ensure that **Select extension** is set to **CRL Distribution Point (CDP)**, and in the **Specify locations from which users can obtain a certificate revocation list (CRL)**, do the following:
   a. Select the entry **file://\\<ServerDNSName>/CertEnroll/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**, and then click **Remove**. In **Confirm removal**, click **Yes**.
   b. Select the entry **http://<ServerDNSName>/CertEnroll/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**, and then click **Remove**. In **Confirm removal**, click **Yes**.
   c. Select the entry that starts with the path **ldap://CN=<CATruncatedName><CRLNameSuffix>,CN=<ServerShortName>**, and then click **Remove**. In **Confirm removal**, click **Yes**.
4. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, click **Add**. The **Add Location** dialog box opens.
5. In **Add Location**, in **Location**, type **http://pki.corp.contoso.com/pki/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**, and then click **OK**. This returns you to the CA properties dialog box.
6. On the **Extensions** tab, select the following checkboxes:
   - **Include in CRLs. Clients use this to find the Delta CRL locations**
   - **Include in the CDP extension of issued certificates**
7. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, click **Add**. The **Add Location** dialog box opens.
8. In **Add Location**, in **Location**, type **file://\\pki.corp.contoso.com/pki/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**, and then click **OK**. This returns you to the CA properties dialog box.
9. On the **Extensions** tab, select the following checkboxes:
   - **Publish CRLs to this location**
   - **Publish Delta CRLs to this location**
10. Change **Select extension** to **Authority Information Access (AIA)**, and in the **Specify locations from which users can obtain a certificate revocation list (CRL)**, do the following:
    a. Select the entry that starts with the path **ldap://CN=<CATruncatedName>,CN=AIA,CN=Public Key Services**, and then click **Remove**. In **Confirm removal**, click **Yes**.

b. Select the entry **http://<ServerDNSName>/CertEnroll/<ServerDNSName>_<CaName><CertificateName>.crt**, and then click **Remove**. In **Confirm removal**, click **Yes**.

c. Select the entry **file://\\<ServerDNSName>/CertEnroll/<ServerDNSName>_<CaName><CertificateName>.crt**, and then click **Remove**. In **Confirm removal**, click **Yes**.

11. In **Specify locations from which users can obtain a certificate revocation list (CRL)**, click **Add**. The **Add Location** dialog box opens.
12. In **Add Location**, in **Location**, type **http://pki.corp.contoso.com/pki/<ServerDNSName>_<CaName><CertificateName>.crt**, and then click **OK**. This returns you to the CA properties dialog box.
13. On the **Extensions** tab, select **Include in the AIA of issued certificates**.
14. In **Add Location**, in **Location**, type **file://\\pki.corp.contoso.com/pki/<ServerDNSName>_<CaName><CertificateName>.crt**, and then click **OK**. This returns you to the CA properties dialog box.

    Important
    Ensure that **Include in the AIA extension of issued certificates** is not selected.
15. When prompted to restart Active Directory Certificate Services, click **No**. You will restart the service later.

Copy the CA Certificate and CRL to the Virtual Directory

You can use this procedure to copy the Certificate Revocation List and Enterprise root CA certificate from your certification authority to a virtual directory on your Web server, and to ensure that AD CS is configured correctly. Before running the commands below, ensure that you replace directory and server names with those that are appropriate for your deployment.

To perform this procedure you must be a member of Domain Admins.

To copy the certificate revocation list from CA1 to WEB1

1. On CA1, run Windows PowerShell as an Administrator, and then publish the CRL with the following command:

   Type `certutil -crl`, and then press ENTER.

   To copy the CA certificate to the file share on your Web server, type `copy C:\Windows\system32\certsrv\certenroll\*.crt \\WEB1\pki`, and then press ENTER.

To restart AD CS, type `Restart-Service certsvc`, and then press ENTER.

2. To verify that your CDP and AIA extension locations are correctly configured, type `pkiview.msc`, and then press ENTER. The pkiview Enterprise PKI MMC opens.
3. Click your CA name. For example, if your CA name is corp-CA1-CA, click **corp-CA1-CA**. In the details pane, verify that the **Status** value for the **CA Certificate**, **AIA Location #1**, and **CDP Location #1** are all **OK**.

Important

If **Status** for any item is not **OK**, do the following:

- Verify that you have entered the correct locations for the CDP and AIA on the CA Extensions tab. Ensure that there are no extra spaces or other characters in the locations that you have provided.
- Verify that you copied the CRL and CA certificate to the correct location on your Web server, and that the location matches the location you provided for the CDP and AIA locations on the CA.
- Verify that you correctly configured permissions for the virtual folder where the CA certificate and CRL are stored.

Configure the Server Certificate Template

You can use this procedure to configure the certificate template that Active Directory® Certificate Services (AD CS) uses as the basis for server certificates enrolled to servers running Network Policy Server (NPS), Routing and Remote Access Service (RRAS), or both.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure the certificate template

1. On CA1, in Server Manager, click **Tools**, and then click **Certification Authority**. The Certification Authority Microsoft Management Console (MMC) opens.
2. In the MMC, double-click the CA name, right-click **Certificate Templates**, and then click **Manage**.
3. The Certificate Templates console opens. All of the certificate templates are displayed in the details pane.
4. In the details pane, click the **RAS and IAS Server** template.
5. Click the **Action** menu, and then click **Duplicate Template**. The template Properties dialog box opens.
6. Click the **Security** tab. In **Group or user names**, click **RAS and IAS servers**.
7. In **Permissions for RAS and IAS servers**, under **Allow**, ensure that **Enroll** is selected, and then select the **Autoenroll** check box. Click **OK**, and close the Certificate Templates

MMC.

8. In the Certification Authority MMC, click **Certificate Templates**. On the **Action** menu, point to **New**, and then click **Certificate Template to Issue**. The **Enable Certificate Templates** dialog box opens.
9. In **Enable Certificate Templates**, click the name of the certificate template that you just configured, and then click **OK**. For example, if you did not change the default certificate template name, click **Copy of RAS and IAS Server**, and then click **OK**.

Configure Server Certificate Autoenrollment

Note
Before you perform this procedure, you must configure a server certificate template by using the Certificate Templates Microsoft Management Console snap-in on a CA that is running AD CS.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure server certificate autoenrollment

1. On the computer where AD DS is installed, open Windows PowerShell®, type **mmc**, and then press ENTER. The Microsoft Management Console opens.
2. On the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Available snap-ins**, scroll down to and double-click **Group Policy Management Editor**. The **Select Group Policy Object** dialog box opens.

   Important
   Ensure that you select **Group Policy Management Editor** and not **Group Policy Management**. If you select **Group Policy Management**, your configuration using these instructions will fail and a server certificate will not be autoenrolled to your NPS servers.
4. In **Group Policy Object**, click **Browse**. The **Browse for a Group Policy Object** dialog box opens.
5. In **Domains, OUs, and linked Group Policy Objects,** click **Default Domain Policy**, and then click **OK**.
6. Click **Finish**, and then click **OK**.
7. Double-click **Default Domain Policy**. In the console, expand the following path: **Computer Configuration**, **Policies**, **Windows Settings**, **Security Settings**, and then **Public Key Policies**.
8. Click **Public Key Policies**. In the details pane, double-click **Certificate Services Client -**

Auto-Enrollment. The **Properties** dialog box opens. Configure the following items, and then click **OK**:

a. In **Configuration Model**, select **Enabled**.

b. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.

c. Select the **Update certificates that use certificate templates** check box.

9. Click **OK**.

Refresh Group Policy

You can use this procedure to manually refresh Group Policy on the local computer. When Group Policy is refreshed, if certificate autoenrollment is configured and functioning correctly, the local computer is autoenrolled a certificate by the certification authority (CA).

Note
Group Policy is automatically refreshed when you restart the domain member computer, or when a user logs on to a domain member computer. In addition, Group Policy is periodically refreshed. By default, this periodic refresh is performed every 90 minutes with a randomized offset of up to 30 minutes.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure.

To refresh Group Policy on the local computer

1. On the computer where NPS is installed, open Windows PowerShell® by using the icon on the taskbar.
2. At the Windows PowerShell prompt, type **gpupdate**, and then press ENTER.

Verify NPS Server Enrollment of a Server Certificate

You can use this procedure to verify that your NPS server has enrolled a server certificate from the certification authority. To verify that a server certificate is correctly configured and is enrolled to the NPS server, you must configure a test network policy and allow NPS to verify that NPS can use the certificate for authentication.

Membership in the **Domain Admins** group is the minimum required to complete this procedure.

To verify NPS server enrollment of a server certificate

1. In Server Manager, click **Tools**, and then click **Network Policy Server**. The Network Policy Server Microsoft Management Console (MMC) opens.
2. Double-click **Policies**, right-click **Network Policies**, and click **New**. The New Network Policy wizard opens.
3. In **Specify Network Policy Name and Connection Type**, in **Policy name**, type **Test policy**. Ensure that **Type of network access server** has the value **Unspecified**, and then click **Next**.
4. In **Specify Conditions**, click **Add**. In **Select condition**, click **Windows Groups**, and then click **Add**.
5. In **Groups**, click **Add Groups**. In **Select Group**, type **Domain Users**, and then press ENTER. Click **OK**, and then click **Next**.
6. In **Specify Access Permission**, ensure that **Access granted** is selected, and then click **Next**.
7. In **Configure Authentication Methods**, click **Add**. In **Add EAP**, click **Microsoft: Protected EAP (PEAP)**, and then click **OK**. In **EAP Types**, select **Microsoft: Protected EAP (PEAP)**, and then click **Edit**. The **Edit Protected EAP Properties** dialog box opens.
8. In the **Edit Protected EAP Properties** dialog box, in **Certificate issued to**, NPS displays the name of your server certificate in the format *ComputerName*.*Domain*. For example, if your NPS server is named NPS-01 and your domain is example.com, NPS displays the certificate **NPS-01.example.com**. In addition, in **Issuer**, the name of your certification authority is displayed, and in **Expiration date**, the date of expiration of the server certificate is shown. This demonstrates that your NPS server has enrolled a valid server certificate that it can use to prove its identity to client computers that are trying to access the network through your network access servers, such as virtual private network (VPN) servers, 802.1X-capable wireless access points, Remote Desktop Gateway servers, and 802.1X-capable Ethernet switches.

   Important
   If NPS does not display a valid server certificate and if it provides the message that such a certificate cannot be found on the local computer, there are two possible reasons for this problem. It is possible that Group Policy did not refresh properly, and the NPS server has not enrolled a certificate from the CA. In this circumstance, restart the NPS server. When the computer restarts, Group Policy is refreshed, and you can perform this procedure again to verify that the server certificate is enrolled. If refreshing Group Policy does not resolve this issue, either the certificate template, certificate autoenrollment, or both are not configured correctly. To resolve these issues, start at the beginning of this guide and perform all steps again to ensure that the settings that you have provided are accurate.
9. When you have verified the presence of a valid server certificate, you can click **OK** and **Cancel** to exit the New Network Policy wizard.

Note
Because you are not completing the wizard, the test network policy is not created in NPS.

Additional Resources

For more information about the technologies that are discussed in this guide, see the following resources:

- The Windows Server 2012 Core Network Guide in Word format in the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=30132) and in HTML format in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh911995.aspx).
- PKI Design Guidance in the Microsoft TechNet Wiki (http://social.technet.microsoft.com/wiki/contents/articles/2901.pki-design-guidance.aspx)
- Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Design Guide in the Microsoft TechNet Wiki http://social.technet.microsoft.com/wiki/contents/articles/7421.ad-cs-pki-design.aspx
- Guidelines for Using Alternate Signature Formats, which is available at the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://technet.microsoft.com/library/cc753169.aspx)
- Creating Certificate Policies and Certificate Practice Statements in the Technical Library (http://technet.microsoft.com/library/cc780454.aspx)
- Request for Comments 2527, Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, at the IETF Web site (http://www.ietf.org/rfc/rfc2527.txt)
- CA Policy.inf Syntax in the Technical Library (http://technet.microsoft.com/library/cc728279.aspx)
- Deployment Planning (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure) in Windows Server TechCenter (http://go.microsoft.com/fwlink/?LinkId=106049)
- Active Directory Domain Services in the Windows Server® 2012 Technical Library (http://technet.microsoft.com/library/hh831484.aspx)
- Group Policy in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh831683.aspx)
- Network Policy and Access Services in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154883)
- Windows Server 2008 R2 Core Network Guide, which is available at the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154884)

- "Configuring All Servers" topic in the Windows Server 2008 R2 Core Network Guide in the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154885)
- "Joining Computers to the Domain and Logging On" topic in the Windows Server 2008 R2 Core Network Guide in the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154886)

Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access

This Guide is a companion guide to the Windows Server® 2012 Core Network Guide. The Core Network Guide provides instructions for planning and deploying the components required for a fully functioning network and a new Active Directory® domain in a new forest.

This guide explains how to build upon a core network by providing instructions about how to deploy Institute of Electrical and Electronics Engineers (IEEE) 802.1X-authenticated IEEE 802.11 wireless access using Protected Extensible Authentication Protocol – Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAP v2).

Because PEAP-MS-CHAP v2 requires that users provide password-based credentials rather than a certificate during the authentication process, it is typically easier and less expensive to deploy than EAP-TLS or PEAP-TLS.

Note
In this guide, IEEE 802.1X Authenticated Wireless Access with PEAP-MS-CHAP v2 is abbreviated to “wireless access” and “WiFi access.”

About this guide

This guide provides instructions about how to deploy a WiFi access infrastructure that uses the following components:

- One or more 802.1X-capable 802.11 wireless access points (APs).
- AD DS Users and Computers.
- Group Policy Management.
- One or more Network Policy Server (NPS) servers.
- Server certificates for computers running NPS.
- Wireless client computers running Windows® 8, Windows® 7, Windows Vista® or Windows XP with Service Pack 2.

This guide is designed for network and system administrators who have:

1. Followed the instructions in the Windows Server 2012 **Core Network Guide** to deploy a core network, or for those who have previously deployed the core technologies included in the core

network, including AD DS, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), TCP/IP, NPS, and Windows Internet Name Service (WINS).

The Core Network Guide is available at the following locations:

- The Windows Server 2012 Core Network Guide is available in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154884).

  The Core Network Guide is also available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=157742).

2. Either followed the instructions in the Windows Server 2012 **Core Network Companion Guide: Server Certificate Deployment** to deploy and use Active Directory Certificate Services (AD CS) to autoenroll server certificates to computers running NPS, or who have purchased a server certificate from a public CA, such as VeriSign, that client computers already trust. A client computer trusts a CA if that CA cert is already in the Trusted Root Certification Authorities certificate store on Windows-based computers. By default, computers running Windows have multiple public CA certificates installed in their Trusted Root Certification Authorities certificate store.

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Server Certificate Deployment in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/p/?LinkID=255200.
- The Windows Server 2012 Core Network Companion Guide: Server Certificate Deployment in HTML format in the Technical Library.

It is recommended that you review the design and deployment guides for each of the technologies that are used in this deployment scenario. These guides can help you determine whether this deployment scenario provides the services and configuration that you need for your organization's network.

Requirements

Following are the requirements for deploying a wireless access infrastructure by using the scenario documented in this guide:

- Before deploying this scenario, you must first purchase and install 802.1X-capable wireless access points to provide wireless coverage in the desired locations at your site.
- Active Directory Domain Services (AD DS) is installed, as are the other network technologies, according to the instructions in the Windows Server 2012 Core Network Guide.
- Server certificates are required when you deploy the PEAP-MS-CHAP v2 certificate-based authentication method.
- You or someone else in your organization is familiar with the IEEE 802.11 standards that are supported by your wireless APs and the wireless network adapters installed in the client computers on your network; for example, radio frequency types, 802.11 wireless authentication (WPA2 or WPA), and ciphers (AES or TKIP).

What this guide does not provide

Following are some items this guide does not provide:

Comprehensive guidance for selecting 802.1X-capable wireless access points

Because many differences exist between brands and models of 802.1X-capable wireless APs, this guide does not provide detailed information about:

- Determining which brand or model of wireless AP is best suited to your needs.
- The physical deployment of wireless APs on your network.
- Advanced wireless AP configuration, such as for wireless VLAN.
- Instructions on how to configure wireless AP vendor-specific attributes in NPS.

Additionally, terminology and names for settings vary between wireless AP brands and models, and might not match the generic setting names referenced in this guide. For wireless AP configuration details, you must review the product documentation provided by the manufacturer of your wireless APs.

Instructions for deploying NPS server certificates

There are two alternatives for deploying NPS server certificates. This guide does not provide comprehensive guidance to help you determine which alternative will best meet your needs. In general, however, the choices you face are:

Purchasing certificates from a public CA, such as VeriSign, that are already trusted by Windows-based clients. This option is typically recommended for smaller networks.

Deploying a Public Key Infrastructure (PKI) on your network by using AD CS.

The following table describes some of the main considerations for deciding whether to deploy a PKI or purchase a certificate from a public CA.

Features	Purchased Public CA certificate	PKI
Scales well	No. Certificates must be purchased and installed on a per-server basis.	Yes. If autoenrollment is used, NPS servers automatically enroll certificates. New NPS servers you add later will also automatically enroll certificates.
Has recurring costs over time	Yes. Public CA certificates must be renewed.	No. When certificates expire, the CA automatically issues new ones.
Requires extensive planning and knowledge	No. On smaller networks, certificates can be purchased and installed on a per-server basis more easily than	Yes. Deploying a PKI requires knowledge of AD CS.

Features	Purchased Public CA certificate	PKI
	deploying a PKI.	
Requires additional hardware	No.	Maybe. To deploy one or more CAs, you must have additional servers or virtual machines.
Is easily extensible	No.	Yes. You can use the PKI to deploy additional authentication methods and certificates used for other purposes.

NPS network policies and other NPS settings

Except for the configuration settings made when you run the **Configure 802.1X** wizard, as documented in this guide, this guide does not provide detailed information for manually configuring NPS conditions, constraints or other NPS settings.

For more information about NPS, see **Additional Resources** in this guide.

DHCP

This deployment guide does not provide information about designing or deploying DHCP subnets for wireless LANs.

For more information about DHCP, see the **Additional Resources** in this guide.

Technology overviews

Following are technology overviews for deploying wireless access:

IEEE 802.1X

The IEEE 802.1X standard defines the port-based network access control that is used to provide authenticated network access to Ethernet networks. This port-based network access control uses the physical characteristics of the switched LAN infrastructure to authenticate devices attached to a LAN port. Access to the port can be denied if the authentication process fails. Although this standard was designed for wired Ethernet networks, it has been adapted for use on 802.11 wireless LANs.

802.1X-capable wireless access points (APs)

This scenario requires the deployment of one or more 802.1X-capable wireless APs that are compatible with the Remote Authentication Dial-In User Service (RADIUS) protocol.

802.1X and RADIUS-compliant APs, when deployed in a RADIUS infrastructure with a RADIUS server such as an NPS server, are called *RADIUS clients*.

Wireless clients

This guide provides comprehensive configuration details to supply 802.1X authenticated access for domain-member users who connect to the network with wireless client computers running Windows 8, Windows 7, Windows Vista or Windows XP with Service Pack 2 or later. Computers must be joined to the domain in order to successfully establish authenticated access.

Wireless computers that are running Windows Server 2012 configured by the same configuration settings as for Windows 8, Windows 7 and Windows Vista. Wireless computers running Windows Server 2003 are configured by the same wireless security and connectivity settings as for computers running Windows XP.

Note

On your domain wireless client computers running Windows 8, Windows 7, and Windows Vista, users can view the profiles you configure in the Windows Vista Wireless Network Policy by opening **Network and Sharing Center** and then clicking **Manage wireless networks**

The Windows Vista Wireless Policies in Windows Server 2012 Group Policy also provide settings that you can use to manage specific features and enhancements that are found only in wireless client computers running Windows 7.

Support for IEEE 802.11 Standards

Windows 8, Windows Server 2012, Windows 7, Windows Vista Windows Server 2003, and, Windows XP, provide built-in support for 802.11 wireless networking. An installed 802.11 wireless network adapter appears as a wireless network connection in the Network Connections folder. Although there is built-in support for 802.11 wireless networking, the wireless components of Windows are dependent upon the following:

- The capabilities of the wireless network adapter. The installed wireless network adapter must support the wireless LAN or wireless security standards that you require. For example, if the wireless network adapter does not support Wi-Fi Protected Access (WPA), you cannot enable or configure WPA security options.
- The capabilities of the wireless network adapter driver. To allow you to configure wireless network options, the driver for the wireless network adapter must support the reporting of all of its capabilities to Windows. Verify that the driver for your wireless network adapter was written for the capabilities of Windows Vista or Windows XP and is the most current version by checking Microsoft Update or the Web site of the wireless network adapter vendor.

The following table shows the transmission rates and frequencies for common IEEE 802.11 wireless standards.

Standards	Frequencies	Bit Transmission Rates	Usage
802.11	S-Band Industrial, Scientific, and Medical (ISM) frequency range	2 megabits per second (Mbps)	Obsolete. Not commonly used.

Standards	Frequencies	Bit Transmission Rates	Usage
	(2.4 to 2.5 GHz)		
802.11b	S-Band ISM	11 Mbps	Commonly used.
802.11a	C-Band ISM (5.725 to 5.875 GHz)	54 Mbps	Not commonly used due to expense and limited range.
802.11g	S-Band ISM	54 Mbps	Widely used. 802.11g devices are compatible with 802.11b devices.
802.11n (IEEE standards development are in progress)	C-Band and S-Band ISM	250 Mbps	Devices based on the pre-ratification IEEE 802.11n standard became available in August 2007. Many 802.11n devices are compatible with 802.11a, b, and g devices.

Wireless network security methods

Wireless network security methods is an informal grouping of wireless authentication (sometimes referred to as wireless security) and wireless security encryption. Wireless authentication and encryption are used in pairs to prevent unauthorized users from accessing the wireless network, and to protect wireless transmissions. When configuring wireless security settings in the Wireless Network Policies of Group Policy there are multiple combinations to choose from. However, only the WPA2-Enterprise, WPA-Enterprise, and Open with 802.1X authentication standards are supported for 802.1X Authenticated wireless deployments. You must select WPA2-Enterprise, WPA-Enterprise, or Open with 802.1X in order to gain access the EAP settings in the Wireless Network Policies that are required for 802.1X authenticated wireless deployments.

Wireless authentication

This guide recommends the use of two wireless authentication standards for 802.1X authenticated wireless deployments:

Wi-Fi Protected Access – Enterprise (WPA-Enterprise) WPA is an interim standard developed by the WiFi Alliance to comply with the 802.11 wireless security protocol. The WPA

protocol was developed in response to a number of severe flaws that were discovered in the preceding Wired Equivalent Privacy (WEP) protocol.

WPA-Enterprise provides improved security over WEP by:

1. Requiring authentication that uses the 802.1X EAP framework as part of the infrastructure that ensures centralized mutual authentication and dynamic key management
2. Enhancing the Integrity Check Value (ICV) with a Message Integrity Check (MIC), to protect the header and payload
3. Implementing a frame counter to discourage replay attacks

Wi-Fi Protected Access 2 – Enterprise (WPA2-Enterprise) Like the WPA-Enterprise standard, WPA2-Enterprise uses the 802.1X and EAP framework. WPA2-Enterprise provides stronger data protection for multiple users and large managed networks. WPA2-Enterprise is a robust protocol that is designed to prevent unauthorized network access by verifying network users through an authentication server.

Wireless security encryption

Wireless security encryption is used to protect the wireless transmissions that are sent between the wireless client and the wireless AP. Wireless security encryption is used in conjunction with the selected network security authentication method. By default, computers running Windows 8, Windows 7 and Windows Vista support two encryption standards:

1. **Temporal Key Integrity Protocol** (TKIP) is an older encryption protocol that was originally designed to provide more secure wireless encryption than what was provided by the inherently weak Wired Equivalent Privacy (WEP) protocol. TKIP was designed by the IEEE 802.11i task group and the Wi-Fi Alliance to replace WEP without requiring the replacement of legacy hardware. TKIP is a suite of algorithms that encapsulates the WEP payload, and allows users of legacy WiFi equipment to upgrade to TKIP without replacing hardware. Like WEP, TKIP uses the RC4 stream encryption algorithm as its basis. The new protocol, however, encrypts each data packet with a unique encryption key, and the keys are much stronger than those by WEP. Although TKIP is useful for upgrading security on older devices that were designed to use only WEP, it does not address all of the security issues facing wireless LANs, and in most cases is not sufficiently robust to protect sensitive government or corporate data transmissions.
2. **Advanced Encryption Standard** (AES) is the preferred encryption protocol for the encryption of commercial and government data. AES offers a higher level of wireless transmission security than either TKIP or WEP. Unlike TKIP and WEP, AES requires wireless hardware that supports the AES standard. AES is a symmetric-key encryption standard that uses three block ciphers, AES-128, AES-192 and AES-256.

Important

Wired Equivalency Privacy (WEP) was the original wireless security standard that was used to encrypt network traffic. You should not deploy WEP on your network as there are well-known vulnerabilities in this outdated form of security.

Enhanced encryption

AES encryption is strengthened by enabling the Federal Information Processing Standard (FIPS) 140-2 standard. FIPS 140-2 is a U.S. government computer security standard that is used to certify cryptographic modules, and specify that wireless transmissions adhere to the FIPS 140-2 standard for cryptography. The option to enable FIPS 140-2 is only available when WPA2-Enterprise with AES, or WPA2-Personal with AES are selected. You must select WPA2-Enterprise with AES to deploy FIPS-140-2 in 802.1X authenticated wireless deployments.

Wireless authentication and encryption pairs

The following table shows wireless security standards (as listed in the Wireless Network Policies extensions of Group Policy) and their corresponding authentication and encryption methods that can be used in 802.1X authenticated deployments.

Wireless network authentication type:	**Wireless encryption:**	**Encryption key bit size:**	**Comments:**
WPA2-Enterprise	Advanced Encryption Standard (AES)	128	Strongest 802.1X-based wireless network authentication with very strong AES encryption.
WPA2-Enterprise	Temporal Key Integrity Protocol (TKIP)	128	Strongest 802.1X-based wireless network authentication with less strong TKIP encryption.
WPA-Enterprise	Advanced Encryption Standard (AES)	128	Mid-strength 802.1X-based wireless network authentication with very strong AES encryption.
WPA-Enterprise	Temporal Key Integrity Protocol (TKIP)	128	Mid-strength 802.1X-based wireless network authentication with less strong TKIP encryption.
IEEE 802.1X	Open	N/A	Not recommended for production environments.
IEEE 802.11	WEP	40 or 104	Use in production environments is strongly discouraged due to weak Wi-Fi authentication and

Wireless network authentication type:	Wireless encryption:	Encryption key bit size:	Comments:
			encryption.

Active Directory Doman Services (AD DS)

AD DS provides a distributed database that stores and manages information about network resources and application-specific data from directory-enabled applications. Administrators can use AD DS to organize elements of a network, such as users, computers, and other devices, into a hierarchical containment structure. The hierarchical containment structure includes the Active Directory forest, domains in the forest, and organizational units (OUs) in each domain. A server that is running AD DS is called a *domain controller*.

AD DS contains the user accounts, computer accounts, and account properties that are required by IEEE 802.1X and PEAP-MS-CHAP v2 to authenticate user credentials and to evaluate authorization for wireless connections.

Active Directory Users and Computers

Active Directory Users and Computers is a component of AD DS that contains accounts that represent physical entities, such as a computer, a person, or a security group. A *security group* is a collection of user or computer accounts that administrators can manage as a single unit. User and computer accounts that belong to a particular group are referred to as *group members*.

Group Policy Management

Group Policy Management is a Windows Server 2012 feature that enables directory-based change and configuration management of user and computer settings, including security and user information. You use Group Policy to define configurations for groups of users and computers. With Group Policy, you can specify settings for registry entries, security, software installation, scripts, folder redirection, remote installation services, and Internet Explorer maintenance. The Group Policy settings that you create are contained in a Group Policy object (GPO). By associating a GPO with selected Active Directory system containers—sites, domains, and OUs—you can apply the GPO's settings to the users and computers in those Active Directory containers. To manage Group Policy objects across an enterprise, you can use the Group Policy Management Editor Microsoft Management Console (MMC).

This guide provides detailed instructions about how to specify settings in the Wireless Network (IEEE 802.11) Policies extension of Group Policy Management. The Wireless Network (IEEE 802.11) Policies configure domain-member wireless client computers with the necessary connectivity and wireless settings for 802.1X authenticated wireless access.

Server certificates

This deployment scenario requires server certificates for each NPS server that performs 802.1X authentication.

A server certificate is a digital document that is commonly used for authentication and to secure information on open networks. A certificate securely binds a public key to the entity that holds the corresponding private key. Certificates are digitally signed by the issuing CA, and they can be issued for a user, a computer, or a service.

A certification authority (CA) is an entity responsible for establishing and vouching for the authenticity of public keys belonging to subjects (usually users or computers) or other CAs. Activities of a certification authority can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and revoking certificates.

Active Directory Certificate Services (AD CS) is a Windows Server 2012 server role that issues certificates as a network CA. An AD CS certificate infrastructure, also known as a *public key infrastructure (PKI)*, provides customizable services for issuing and managing certificates for the enterprise.

EAP, PEAP, and PEAP-MS-CHAP v2

Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by allowing additional authentication methods that use credential and information exchanges of arbitrary lengths. With EAP authentication, both the network access client and the authenticator (such as the NPS server) must support the same EAP type for successful authentication to occur. Windows Server 2012 includes an EAP infrastructure, supports two EAP types, and the ability to pass EAP messages to NPS servers. By using EAP, you can support additional authentication schemes, known as *EAP types*. The EAP types that are supported by Windows Server 2012 are:

- Transport Layer Security (TLS)
- Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)

Security

Strong EAP types (such as those that are based on certificates) offer better security against brute-force attacks, dictionary attacks, and password guessing attacks than password-based authentication protocols (such as CHAP or MS-CHAP version 1).

Protected EAP (PEAP) uses TLS to create an encrypted channel between an authenticating PEAP client, such as a wireless computer, and a PEAP authenticator, such as an NPS server or other RADIUS servers. PEAP does not specify an authentication method, but it provides additional security for other EAP authentication protocols (such as EAP-MS-CHAP v2) that can operate through the TLS encrypted channel provided by PEAP. PEAP is used as an authentication method for access clients that are connecting to your organization's network through the following types of network access servers (NASs):

- 802.1X-capable wireless access points
- 802.1X-capable authenticating switches

- Computers running Windows Server 2012 and the Routing and Remote Access service (RRAS) that are configured as virtual private network (VPN) servers
- Computers running Windows Server 2012 and Terminal Services Gateway

PEAP-MS-CHAP v2 is easier to deploy than EAP-TLS because user authentication is performed by using password-based credentials (user name and password), instead of certificates or smart cards. Only NPS or other RADIUS servers are required to have a certificate. The NPS server certificate is used by the NPS server during the authentication process to prove its identity to PEAP clients.

This guide provides instructions to configure your wireless clients and your NPS server(s) to use PEAP-MS-CHAP v2 for 802.1X authenticated access.

Network Policy Server

Network Policy Server (NPS) is included in Windows Server 2012, and allows you to centrally configure and manage network policies by using the following three components: Remote Authentication Dial-In User Service (RADIUS) server, RADIUS proxy, and Network Access Protection (NAP) policy server. NPS is an optional service of a core network, but it is required to deploy 802.1X wireless access.

When you configure your 802.1X wireless access points as RADIUS clients in NPS, NPS processes the connection requests sent by the APs. During connection request processing, NPS performs authentication and authorization. Authentication determines whether the client has presented valid credentials. If NPS successfully authenticates the requesting client, then NPS determines whether the client is authorized to make the requested connection, and either allows or denies the connection. This is explained in more detail as follows:

Authentication

Successful mutual PEAP-MS-CHAP v2 authentication has two main parts:

1. The client authenticates the NPS server. During this phase of mutual authentication, the NPS server sends its server certificate to the client computer so that the client can verify the NPS server's identity with the certificate. To successfully authenticate the NPS server, the client computer must trust the CA that issued the NPS server certificate. The client trusts this CA when the CA's certificate is present in the Trusted Root Certification Authorities certificate store on the client computer.

   If you deploy your own private CA, the CA certificate is automatically installed in the Trusted Root Certification Authorities certificate store for the Current User and for the Local Computer when Group Policy is refreshed on the domain member client computer. If you decide to deploy server certificates from a public CA, ensure that the public CA certificate is already in the Trusted Root Certification Authorities certificate store.
2. The NPS server authenticates the user. After the client successfully authenticates the NPS server, the client sends the user's password-based credentials to the NPS server, which verifies the user's credentials against the user accounts database in Active Directory Doman Services (AD DS).

If the credentials are valid and authentication succeeds, the NPS server begins the authorization phase of processing the connection request. If the credentials are valid and authentication succeeds, the NPS server begins the authorization phase of processing the connection request. If the credentials are not valid and authentication fails, NPS sends an Access Reject message and the connection request is denied.

Authorization

The server running NPS performs authorization as follows:

1. NPS checks for restrictions in the user or computer account dial-in properties in AD DS.
2. NPS then processes its network policies to find a policy that matches the connection request. If a matching policy is found, NPS either grants or denies the connection based on that policy's configuration.

If both authentication and authorization are successful, and if the matching network policy grants access, NPS grants access to the network, and the user and computer can connect to network resources for which they have permissions.

Note

To deploy wireless access, you must configure NPS policies. This guide provides instructions to use the **Configure 802.1X wizard** in NPS to create NPS policies for 802.1X authenticated wireless access.

Bootstrap profiles

In 802.1X-authenticated wireless networks, wireless clients must provide security credentials that are authenticated by a RADIUS server in order to connect to the network. For Protected EAP [PEAP]-Microsoft Challenge Handshake Authentication Protocol version 2 [MS-CHAP v2], the security credentials are a username and password. For EAP-Transport Layer Security [TLS], the security credentials are certificates, such as client user and computer certificates or smart cards.

When connecting to a network that is configured to perform either PEAP-MS-CHAP v2 or EAP-TLS authentication, by default, Windows wireless clients must also validate a computer certificate that is sent by the RADIUS server. The computer certificate that is sent by the RADIUS server for every authentication session is commonly referred to as a server certificate.

As mentioned previously, RADIUS servers can be issued server certificate in one of two ways: from a commercial CA (such as VeriSign, Inc.,), or from a private CA that you deploy on your network. If the RADIUS server sends a computer certificate that was issued by a commercial CA that already has a root certificate installed in the client's Trusted Root Certifications Authorities certificate store, then the wireless client can validate the RADIUS server's computer certificate, regardless of whether the wireless client has joined the Active Directory domain. In this case the wireless client can connect to the wireless network, and then join the computer to the domain.

Note

The behavior requiring the client to validate the server certificate can be disabled, but disabling server certificate validation is not recommended in production environments.

Wireless bootstrap profiles are temporary profiles that are configured in such a way as to enable wireless client users to connect to the 802.1X-authenticated wireless network before the computer is joined to the domain, and/or before the user has successfully logged on to the domain by using a given wireless computer for the first time. This section summarizes what problem is encountered when trying to join a wireless computer to the domain, or for a user to use a domain-joined wireless computer for the first time to log on to the domain.

For deployments in which the user or IT administrator cannot physically connect a computer to the wired Ethernet network to join the computer to the domain, and the computer does not have the necessary issuing root CA certificate installed in its **Trusted Root Certification Authorities** certificate store, you can configure wireless clients running Windows 8, Windows 7 and Windows Vista with a temporary wireless connection profile, called a *bootstrap profile*, to connect to the wireless network. A *bootstrap profile* removes the requirement to validate the RADIUS server's computer certificate. This temporary configuration enables the wireless user to join the computer to the domain, at which time the Wireless Network (IEEE 802.11) Policies are applied and the appropriate root CA certificate is then installed on the computer. When Group Policy is applied, one or more wireless connection profiles that enforce the requirement for mutual authentication are applied on the computer; the bootstrap profile is no longer removed. After joining the computer to the domain and restarting the computer, the user can use a wireless connection to log on to the domain.

Wireless Access Deployment Overview

The following illustration shows the components that are required to deploy 802.1X authenticated wireless access with PEAP-MS-CHAP v2.

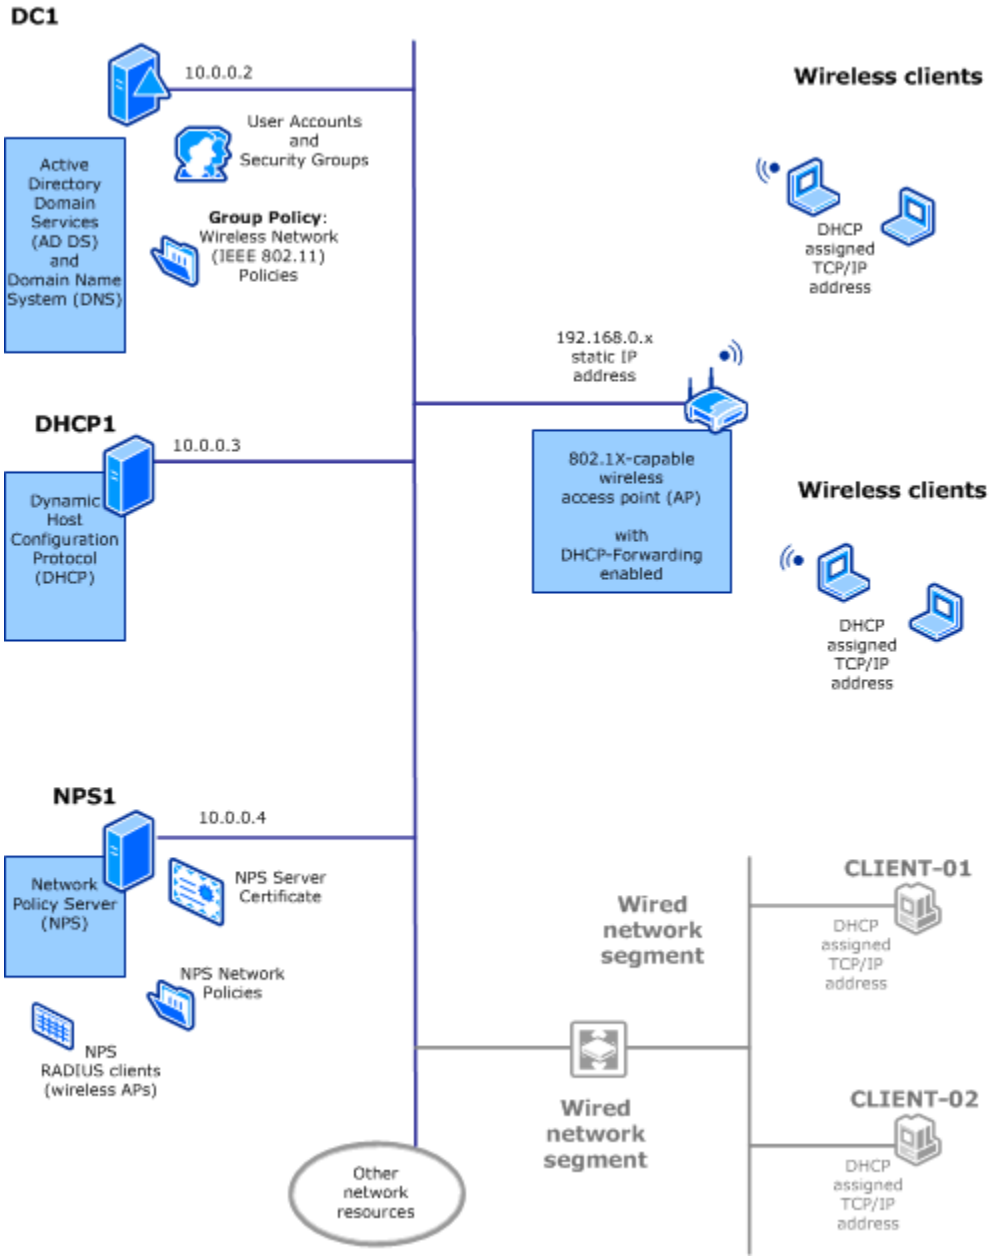

Wireless access deployment components

The following components are required for this wireless access deployment:

802.1X-capable Wireless access points

After the required network infrastructure services supporting your wireless local area network are in place, you can begin the design process for the location of the wireless APs. The wireless AP deployment design process involves these steps:

- Identify the areas of coverage for wireless users. While identifying the areas of coverage, be sure to identify whether you want to provide wireless service outside the building, and if so, determine specifically where those external areas are.
- Determine how many wireless APs to deploy to ensure adequate coverage.
- Determine where to place wireless APs.
- Select the channel frequencies for wireless APs.

Active Directory Domain Services

Users and Computers

Use the Active Directory Users and Computers snap-in to create and manage user accounts, and to create a wireless security group for each domain member to whom you want to grant wireless access.

Wireless Network (IEEE 802.11) Policies

You can use the Wireless Network (IEEE 802.11) Policies extension of Group Policy Management to configure policies for computers that are running Windows 8, Windows® 7 and Windows Vista®, and Windows XP. As is the case with Group Policy Management found in Windows Server 2012, there are two separate wireless policy nodes in the Windows Server 2012 Wireless Network (IEEE 802.11) Policies extension of Group Policy. For the purpose of this discussion, a wireless policy node is a collection of Group Policy settings that can be applied to computers running specific operating systems. By default, the two wireless policy nodes in the Windows Server 2012 Wireless Network (IEEE 802.11) Policies are named:

New XP Wireless Network Policy

New Wireless Network Policy

Tip

When configuring either the New XP Wireless Network Policy or the New Wireless Network Policy you are provided the option to change the name and description of the policy. If you change the name of either policy, that change is reflected in the Details pane of Group Policy Management Editor, and on the title bar of the wireless network policy dialog. Regardless of how you rename your policies, the New XP Wireless Policy will always be listed in Group Policy Management Editor with the **Type** displaying **XP**. The New Wireless Network Policy will always be listed with the **Type** showing **Vista and later Releases**.

The following table describes which of the operating systems can be configured by each wireless policy.

Policy	Can be applied to computers running Windows XP, Windows Server 2003	Can be applied to computers running Windows Vista, Windows Server 2008	Can be applied to Windows 7 Windows Server 2008 R2, Windows 8, Windows Server 2012
New XP Wireless Network Policy	Yes	Yes But this policy cannot configure new wireless features in Windows Vista	Yes But this policy cannot configure new wireless features in Windows 7, Windows Server 2008 R2, Windows 8, and Windows Server 2012.
New Wireless Network Policy	No	Yes Includes settings for all of the wireless features in Windows Vista.	Yes Includes settings for all of the wireless features in Windows Vista, and settings for wireless feature enhancements in Windows 7, Windows Server 2008 R2, Windows 8, and Windows Server 2012. Computers running Windows Vista will ignore Windows 7-specific settings.

Note
To enable wireless service on computers running Windows Server 2008 R2 and Windows Server 2012, you must enable the wireless LAN Service. For more information, see Wireless LAN Service Overview.

The Windows Vista Wireless Network Policy enables you to configure, prioritize, and manage multiple wireless profiles. A wireless profile is a collection of connectivity and security settings that are used to connect to a specific wireless network. When Group Policy is updated on your wireless client computers, the profiles you create in the Windows Vista Wireless Network Policy are automatically added to your wireless client computers that are running Windows 7 and Windows Vista to which the Wireless Network Policy applies.

If you have wireless clients that you want to be able to connect to more than one wireless network, you can configure a wireless profile that contains the specific connectivity and security settings for each network. For example, assume your company has one wireless network for the main corporate office, with a service set identifier (SSID) WlanCorp. Your branch office also has a wireless network to which you also want to connect. The branch office has the SSID configured as WlanBranch. In this scenario, you can configure a profile for each network, and laptops that are used at both the corporate office and branch office will be able to connect to the wireless networks when they are at either location.

Alternately, assume your network has a mixture of wireless computers that support different security standards. Perhaps some older computers have wireless adapters that can only use WPA-Enterprise, while newer devices can use the stronger WPA2-Enterprise standard. You can create two different profiles that use the same SSID and nearly identical connectivity and security settings. However in one profile, you can set the wireless authentication to WPA2-Enterprise with AES, and in the other profile you can specify WPA-Enterprise with TKIP. This is commonly known as a mixed-mode deployment. The ability to configure mixed-mode deployments using a common SSID is one of the enhancements in the Wireless Network (IEEE 802.11) Policies for Windows Vista.

NPS

Network Policy Server (NPS) enables you to create and enforce network access policies for client health, connection request authentication, and connection request authorization. When you use NPS as a RADIUS server, you configure network access servers, such as wireless access points, as RADIUS clients in NPS. You also configure the network policies that NPS uses to authenticate access clients and authorize their connection requests.

Wireless client computers

For the purpose of this guide, wireless client computers are computers that are equipped with IEEE 802.11 wireless network adapters and that are most typically running Windows 8, Windows 7, Windows Vista or Windows XP. Within the context of this deployment scenario, however, wireless client computers can also be computers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2003.

By default, the functionality for 802.11 wireless is disabled on computers that are running Windows Server 2012 and Windows Server 2008 R2. To use 802.11 wireless on computers running Windows Server 2012 and Windows Server 2008 R2 you must first install and enable the Wireless LAN Service feature on the server computer. You can install the Wireless LAN Service feature by using the Add Features Wizard in Server Manager.

Wireless access deployment process

The process of configuring and deploying wireless access occurs in these stages:

Stage 1 – AP Deployment

Plan, deploy, and configure your APs for wireless client connectivity and for use with NPS. Depending on your preference and network dependencies, you can either pre-configure settings on your wireless APs prior to installing them on your network, or you can configure them remotely after installation.

Stage 2 – AD DS Group Configuration

You must create one or more wireless users security groups. Then, you must add each user for whom you want to allow wireless access to the wireless network to the appropriate wireless users security group.

Stage 3 – Group Policy Configuration

Configure the Wireless Network (IEEE 802.11) Policies extension of Group Policy by using the Group Policy Management Editor Microsoft Management Console (MMC).

To configure domain-member computers using the settings in the wireless network policies, you must apply Group Policy. When a computer is first joined to the domain Group Policy is automatically applied. If changes are made to Group Policy, the new settings are automatically applied:

- by Group Policy at pre-determined intervals
- if a domain user logs off and then back on to the network
- by restarting the client computer and logging on to the domain

You can also force Group Policy to refresh by running **gpupdate** at the command prompt.

Stage 4 – NPS server configuration

Use a configuration wizard in NPS to add wireless access points as RADIUS clients, and to create the network policies that NPS uses when processing connection requests. When using the wizard to create the network policies, specify PEAP as the EAP type, and the wireless users security group that was created in the second stage.

Stage 5 – Deploy wireless clients

Use client computers to connect to the network.

For domain member computers that can log on to the wired LAN, the necessary wireless configuration settings are automatically applied when Group Policy is refreshed. If you have enabled the setting in Wireless Network (IEEE 802.11) Policies to connect automatically when the computer is within broadcast range of the wireless network, your wireless, domain-joined computers will then automatically attempt to connect to the wireless LAN. To connect to the wireless network, users need only supply their domain user name and password credentials when prompted by Windows.

Wireless Access Deployment Planning

Before you deploy wireless access, you must plan the following items:

- Installation of wireless access points (APs) on your network
- Wireless client configuration and access

Planning wireless AP installations

When you design your wireless network access solution, you must determine what standards your wireless APs must support, the areas where you want to provide wireless service, and where to locate wireless APs. Additionally, you must plan an IP address scheme for your wireless AP's and wireless clients. See the section **Plan the configuration of wireless AP's in NPS** below for related information.

Verify wireless AP support for standards

For the purposes of consistency and ease of deployment and AP management, it is recommended that you deploy wireless APs of the same brand and model.

The wireless APs that you deploy must support the following:

IEEE 802.1X

RADIUS authentication

Wireless Authentication and Cipher. Listed in order of most to least preferred:

a. WPA2-Enterprise with AES
b. WPA2-Enterprise with TKIP
c. WPA-Enterprise with AES
d. WPA-Enterprise with TKIP

Note

To deploy WPA2, use wireless network adapters and wireless APs that also support WPA2. Otherwise, use WPA-Enterprise.

In addition, to provide enhanced security for the network, the wireless APs must support the following options:

- **DHCP filtering.** The wireless AP must filter on IP ports to prevent the transmission of DHCP broadcast messages in those cases in which the client is a DHCP server. The wireless AP must block the client from sending IP packets from UDP port 68 to the network.
- **DNS filtering.** The wireless AP must filter on IP ports to prevent a client from performing as a DNS server. The wireless AP must block the client from sending IP packets from TCP or UDP port 53 to the network.
- **Client isolation** If your wireless access point provides client isolation capabilities, it should be enabled to prevent possible Address Resolution Protocol (ARP) spoofing, exploits.

Identify areas of coverage for wireless users

Use architectural drawings of each floor for each building to identify the areas where you want to provide wireless coverage. For example, identify the appropriate offices, conferences rooms, lobbies, cafeterias, or courtyards. On the drawings, indicate any devices that interfere with the wireless signals, such as medical equipment, wireless video cameras, cordless telephones that operate in the 2.4 through 2.5 GHz Industrial, Scientific and Medical (ISM) range, and Bluetooth-enabled devices. On the drawing, mark aspects of the building that might interfere with wireless

signals; metal objects used in the construction of a building can affect the wireless signal. For example, the following common objects can interfere with signal propagation: Elevators, heating and air-conditioning ducts, and concrete support girders.

Refer to your AP manufacturer for information about sources that might cause wireless AP radio frequency attenuation. Most APs provide testing software that you can use to check for signal strength, error rate, and data throughput.

Determine where to install wireless APs

On the architectural drawings, locate your wireless APs close enough together to provide ample wireless coverage but far enough apart that they do not interfere with each other. The necessary distance between APs depends upon the type of AP and AP antenna, aspects of the building that block wireless signals, as well as other sources of interference. Typically, mark wireless APs placements so that each wireless AP is not more than 300 feet from any adjacent wireless AP. See the wireless AP manufacturer's documentation for AP specifications and guidelines for placement.

Temporarily install wireless APs in the locations specified on your architectural drawings. Then using a laptop equipped with an 802.11 wireless adapter and the site survey software that is commonly supplied with wireless adapters, determine the signal strength within each coverage area. In coverage areas where signal strength is low, reposition the AP to improve signal strength for the coverage area, install additional wireless APs to provide the necessary coverage, relocate or remove sources of signal interference.

Update your architectural drawings to indicate the final placement of all wireless APs. Having an accurate AP placement map will assist later during troubleshooting operations or when you want to upgrade or replace APs.

Plan the configuration of wireless APs in NPS

NPS has the ability to configure wireless APs individually or in groups. If the following conditions are met, you can configure the wireless APs that you deploy in groups:

1. If you are running NPS servers on computers running Windows Server 2012 Enterprise or Windows Server 2012 Datacenter
2. If you have deployed wireless APs within the same IP address range
3. And if the wireless APs are all configured with the same shared secret

The advantage of configuring your wireless APs in NPS by groups is that it cuts down on your administrative overhead; which is very useful feature if you have deployed a large number of wireless APs. If you have not deployed your NPS servers on computers running either Windows Server 2012 Enterprise or Windows Server 2012 Datacenter, then you must configure each wireless AP individually in NPS.

Plan the use of PEAP Fast Reconnect

In an 802.1X infrastructure, wireless access points are configured as RADIUS clients to RADIUS servers. When PEAP fast reconnect is deployed, a wireless client that roams between two or more access points is not required to be authenticated with each new association. PEAP fast reconnect reduces the response time for authentication between client and authenticator because the authentication request is forwarded from the new access point to the NPS server that originally performed authentication and authorization for the client connection request. Because both the PEAP client and NPS server both use previously cached Transport Layer Security (TLS) connection properties (the collection of which is named the TLS handle), the NPS server can quickly determine that the client is authorized for a reconnect.

Important

For fast reconnect to work, the APs must be configured as RADIUS clients of the same NPS server

If the original NPS server becomes unavailable, or if the client moves to an access point that is configured as a RADIUS client to a different RADIUS server, full authentication must occur between the client and the new authenticator.

Wireless AP configuration

The following list summarizes items commonly configured on 802.1X-capable wireless APs:

Note

The item names can vary by brand and model and might be different from those in the following list. See your wireless AP documentation for configuration-specific details.

- **Service set identifier (SSID).** This is the name of the wireless network (for example, ExampleWlan), and the name that is advertised to wireless clients. To reduce confusion, the SSID that you choose to advertise should not match the SSID that is broadcast by any wireless networks that are within reception range of your wireless network.

  In cases in which multiple wireless APs are deployed as part of the same wireless network, configure each wireless AP with the same SSID. In cases in which multiple wireless APs are deployed as part of the same wireless network, configure each wireless AP with the same SSID.

  In cases where you have a need to deploy different wireless networks to meet specific business needs, your wireless AP's on one network should broadcast a different SSID than the SSID your other network(s). For example, if you need a separate wireless network for your employees and guests, you could configure your wireless APs for the business network with the SSID set to broadcast **ExampleWLAN**. For your guest network, you could then set each wireless AP's SSID to broadcast **GuestWLAN**. In this way your employees and guests can connect to the intended network without unnecessary confusion.

  Tip

Some wireless AP's have the ability to broadcast multiple SSID's to accommodate multi-network deployments. Wireless AP's that can broadcast multiple SSID's can reduce deployment and operational maintenance costs.

- **Wireless authentication and encryption.**

  Wireless authentication is the security authentication that is used when the wireless client associates with a wireless access point.

  Wireless encryption is the security encryption cipher that is used with wireless authentication to protect the communications that are sent between the wireless AP and the wireless client.

- **Wireless AP IP address (static).** Each wireless AP, configure a unique static IP address that falls within the DHCP exclusion range as documented in the Windows Server 2012 Core Network Guide procedure "Creating a new DHCP Scope."
- **DNS name.** Some wireless APs can be configured with a DNS name. Configure each wireless AP with a unique name. For example, if you have a deployed wireless APs in a multi-story building, you might name the first three wireless APs that are deployed on the third floor AP3-01, AP3-02, and AP3-03.
- **Wireless AP subnet mask.** Configure the mask to designate which portion of the IP address that is the network ID and which portion of the IP address is the host.
- **AP DHCP service.** If your wireless AP has a built-in DHCP service, disable it.
- **RADIUS shared secret.** Use a unique RADIUS shared secret for each wireless AP. Each shared secret should be a random sequence at least 22 characters long of uppercase and lowercase letters, numbers, and punctuation. To ensure randomness, you can use a random character generation program to create your shared secrets. It is recommended that you record the shared secret for each wireless AP and store it in a secure location, such as an office safe. When you configure RADIUS clients in the NPS console you will create a virtual version of each AP. The shared secret that you configure on each virtual AP in NPS must match the shared secret on the actual, physical AP.
- **RADIUS server IP address.** Type the IP address of the NPS server that you want to use to authenticate and authorize connection requests to this access point...
- **UDP port(s).** By default, NPS uses UDP ports 1812 and 1645 for RADIUS authentication messages and UDP ports 1813 and 1646 for RADIUS accounting messages. It is recommended that you do not change the default RADIUS UDP ports settings.
- **VSAs.** Some wireless APs require vendor-specific attributes (VSAs) to provide full wireless AP functionality.
- **DHCP filtering.** Configure wireless APs to block wireless clients from sending IP packets from UDP port 68 to the network. See the documentation for your wireless AP to configure DHCP filtering.
- **DNS filtering.** Configure wireless APs to block wireless clients from sending IP packets from TCP or UDP port 53 to the network. See the documentation for your wireless AP to configure DNS filtering.

Planning wireless client configuration and access

When planning the deployment of 802.1X-authenticated wireless access, you must consider several client-specific factors:

- **Planning support for multiple standards.**

  Determine whether your wireless computers are all using the same version of Windows. For example, determine whether all of your wireless client computers all running Windows 8, or whether your wireless deployment include have a mixture of computers running Windows 8, Windows 7, Windows Vista and Windows XP.

  Determine whether all of the wireless network adapters on all of the wireless client computers support the same wireless standards, or whether you need to support varying standards. For example, determine whether some network adapter hardware drivers support WPA2-Enterprise and AES, while others support only WPA-Enterprise and TKIP.

- **Planning client authentication mode.** Authentication modes define how Windows clients process domain credentials. You can select from the following three network authentication modes in the wireless network policies.

  a. **User re-authentication** This mode specifies that authentication is always performed by using security credentials based on the computer's current state. When no users are logged on to the computer, authentication is performed by using the computer credentials. When a user is logged on to the computer, authentication is always performed by using the user credentials.

  b. **Computer only** Computer only mode specifies that authentication is always performed by using only the computer credentials.

  c. **User authentication** User authentication mode specifies that authentication is only performed when the user is logged on to the computer. When there are no users logged on to the computer, authentication attempts are not performed.

- **Planning wireless restrictions.** Determine whether you want to provide all of your wireless users with the same level of access to your wireless network, or whether you want to restrict access for some of your wireless users. You can apply restrictions in NPS against specific groups of wireless users. For example, you can define specific days and hours that certain groups are permitted access the wireless network.

- **Planning methods for adding new wireless computers.** For wireless-capable computers that are joined to your domain before you deploy your wireless network, if the computer is connected to a segment of the wired network that is not protected by 802.1X, the wireless configuration settings are automatically applied on those computers after you configure Wireless Network (IEEE 802.11) Policies and Group Policy is refreshed.

  For computers that are not already joined to your domain, however, you must plan a method to apply the settings that are required for 802.1X-authenticated access. For example, determine whether you want to join the computer to the domain by:

  a. Connecting the computer to a segment of the wired network that is not protected by 802.1X, the joining the computer to the domain

  b. Provide your wireless users with the steps and settings that they require to add their own wireless bootstrap profile, and the join the computer to the domain

c. Having the configuration performed by members of your IT staff

Planning support for multiple standards

In Windows Server 2012, the Wireless Network (IEEE 802.11) Policies extension in Group Policy provides a wide range of configuration options to support a variety of deployment options. You can deploy wireless APs that are configured with the standards that you want to support, and then configuring multiple wireless profiles in Wireless Network (IEEE 802.11) Policies, with each profile specifying one of the required set of standards.

For example, if your network has wireless computers that support WPA2-Enterprise and AES, other computers that support WPA-Enterprise and AES, while other computers support only WPA-Enterprise and TKIP, you must determine whether you want to:

- Configure a single profile to support all of the wireless computers using the weakest encryption but which all of your computer can use; in this case, WPA-Enterprise and TKIP.
- Configure two profiles to provide the best possible security that is supported by each wireless computer. In this instance you would configure one profile that specifies the strongest encryption (WPA2-Enterprise and AES), and one profile that uses the weaker WPA-Enterprise and TKIP encryption. In this example, it is essential that you place the profile that uses WPA2-Enterprise and AES highest in the preference order. Computers that are not capable of using WPA2-Enterprise and AES will automatically skip to the next profile in the preference order and process the profile that specifies WPA-Enterprise and TKIP.

Important

The profile with the most secure standards should be placed higher in the list because connecting computers will use the first profile that they are capable of using.

Planning restricted access to the wireless network

In many cases, you might want to provide wireless users with varying levels of access to the wireless network. For example, you might want to allow some users unrestricted access, any hour of the day, every day of the week. For other users, you might only want to allow access during core hours, Monday through Friday, and deny access on Saturday and Sunday.

This guide provides instructions to create an access environment that places all of your wireless users in a group with common access to wireless resources. You create one wireless users security group in the Active Directory Users and Computers snap-in, and then make every user for whom you want to grant wireless access – a member of that group. When you configure NPS network policies, you specify the wireless users security group as the object that NPS processes when determining authorization.

However, if your deployment requires support for varying levels of access you need only do the following:

1. Follow the procedure **Create a Wireless Users Security Group** in this guide, to create additional wireless security groups in Active Directory Users and Computers, each security group specifying a unique name.

2. Follow the procedure **Add Users to the Wireless Security Group** to make each user a member of the appropriate security group.
3. Follow the procedure **Create NPS Policies for 802.1X Wireless Using a Wizard** to configure an additional set of NPS policies for each additional wireless security group. In step 9 of the procedure, in **Specify User Groups**, click **Add**, and then type the name of the security group that you configured in Active Directory Users and Computers.

Planning methods for adding new wireless computers

The preferred method to join new wireless computers to the domain and then log on to the domain is by using a wired connection to a segment of the LAN that has access to domain controllers, and is not protected by an 802.1X authenticating Ethernet switch.

In some cases, however, it might not be practical to use a wired connection to join computers to the domain, or, for a user to use a wired connection for their first logon attempt by using computers that are already joined to the domain. To join a computer to the domain by using a wireless connection or for users to logon to the domain the first time by using a domain-joined computer and a wireless connection, wireless clients must first establish a connection to the wireless network on a segment that has access to the network domain controllers.

For more information about the steps to join computers to the domain by using a wired connection, and to log on to the domain by using a wired connection, see the Windows Server 2012 Core Network Guide, in the section titled **Joining computers to the Domain and Logging On**. The Windows Server 2012 Core Network Guide is available for download in Word formant at the Microsoft Download Center (http://www.microsoft.com/download/details.aspx?id=29248) and in HTML format in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh911995.aspx).

This guide provides the following methods to configure wireless computers running Windows Vista with wireless profiles that users can use to either join the computer to the domain by using a wireless connection, or to log on to the domain by using a wireless connection and a computer that is already joined to the domain:

1. **A member of the IT staff joins a wireless computer to the domain, and then configures a Single Sign On bootstrap wireless profile.** In this method, an IT administrator connects the wireless computer to the wired Ethernet network, and then joins the computer to the domain. Then the administrator distributes the computer to the user. When the user starts the computer, the domain credentials that they manually specify for the user logon process are used to both establish a connection to the wireless network and log on to the domain.
2. **The user manually configures wireless computer with bootstrap wireless profile, and then joins the domain.** In this method, users manually configure their wireless computers with a bootstrap wireless profile based on instructions from an IT administrator. The bootstrap wireless profile allows users to establish a wireless connection, and then join the computer to the domain. After joining the computer to the domain and restarting the computer, the user can log on to the domain by using a wireless connection and their domain account credentials.

Wireless Access Deployment

Follow these steps to deploy wireless access:

- **Deploying and Configuring Wireless APs**
- **Creating Security Groups for Wireless Users**
- **Configuring Wireless Network (IEEE 802.11) Policies**
- **Configuring your NPS Server**
- **Joining New Wireless Computers to the Domain**

Deploying and Configuring Wireless APs

Follow these steps to deploy and configure your wireless APs:

- **Specify Channel Frequencies for Wireless APs [preliminary]**
- **Configure Wireless APs [Preliminary]**

Note

The procedures in this guide do not include instructions for cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Specify Wireless AP Channel Frequencies

When you deploy multiple wireless APs at a single geographical site, you must configure wireless APs that have overlapping signals to use unique channel frequencies to reduce interference between wireless APs.

To specify unique channel frequencies for wireless APs

1. If there are other organizations that have offices in close proximity or in the same building as your organization, identify whether there are any wireless networks owned those organizations. Find out both the placement and the assigned channel frequencies of their wireless AP's, because you need to assign different channel frequencies to your AP's and you need to determine the best location to install your AP's.
2. Identify overlapping wireless signals on adjacent floors within your own organization. After identifying overlapping coverage areas outside and within your organization, assign channel frequencies for your wireless APs, ensuring that any two wireless APs with overlapping coverage are assigned different channel frequencies.

Configure Wireless APs

Use this information with the product documentation provided by the wireless AP manufacturer to configure your wireless APs.

This procedure enumerates items commonly configured on a wireless AP. The item names can vary by brand and model and might be different from those listed in the following list. For configuration-specific details, see your wireless AP documentation.

To configure your wireless APs

- **SSID.** Specify the name of the wireless network(s) (for example, ExampleWLAN). This is the name that is advertised to wireless clients.
- **Encryption.** Specify WPA2-Enterprise (preferred) or WPA-Enterprise, and either AES (preferred) or TKIP encryption cipher, depending on which versions are supported by your wireless client computer network adapters.
- **Wireless AP IP address (static).** On each AP, configure a unique static IP address that falls within the exclusion range of the DHCP scope. Using an address that is excluded from assignment by DHCP prevents the DHCP server from assigning the same IP address to a computer or other device.
- **Subnet mask.** Configure this to match the subnet mask settings of the LAN to which you have connected the wireless AP.
- **DNS name.** Some wireless APs can be configured with a DNS name. The DNS service on the network can resolve DNS names to an IP address. On each wireless AP that supports this feature, enter a unique name for DNS resolution.
- **DHCP service.** If your wireless AP has a built-in DHCP service, disable it.
- **RADIUS shared secret.** Use a unique RADIUS shared secret for each wireless AP. Each shared secret should be a random sequence at least 22 characters long of uppercase and lowercase letters, numbers, and punctuation. To ensure randomness, you can use a random character generation, such as the random character generator found in the NPS **Configure 802.1X** wizard, to create the shared secrets.

  Tip
  Record the shared secret for each wireless AP and store it in a secure location, such as an office safe. You must know the shared secret for each wireless AP when you configure RADIUS clients in the NPS.
- **RADIUS server IP address.** Type the IP address of the server running NPS.
- **UDP port(s).** By default, NPS uses UDP ports 1812 and 1645 for authentication messages and UDP ports 1813 and 1646 for accounting messages.

  Tip
  Do not change the default RADIUS UDP port settings.
- **VSAs.** Some wireless APs require vendor-specific attributes (VSAs) to provide full wireless AP functionality. VSAs are added in NPS network policy.
- **DHCP filtering.** Configure wireless APs to block wireless clients from sending IP packets from UDP port 68 to the network, as documented by the wireless AP manufacturer.
- **DNS filtering.** Configure wireless APs to block wireless clients from sending IP packets from TCP or UDP port 53 to the network, as documented by the wireless AP manufacturer.

Creating Security Groups for Wireless Users

Follow these steps to create one or more wireless users security groups, and then add users to the appropriate wireless users security group:

- **Create a Wireless Users Security Group**
- **Add Users to the Wireless Security Group**

Create a Wireless Users Security Group

You can use this procedure to create a wireless security group in the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in.

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To create a wireless users security group

1. Click **Start**, click **Administrative Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers snap-in opens. If it is not already selected, click the node for your domain. For example, if your domain is example.com, click **example.com**.
2. In the details pane, right-click the folder in which you want to add a new group (for example, right-click **Users**), point to **New**, and then click **Group**.
3. In **New Object – Group**, in **Group name**, type the name of the new group. For example, type **Wireless Group**.
4. In **Group scope**, select one of the following options:
   a. **Domain local**
   b. **Global**
   c. **Universal**
5. In **Group type**, select **Security**.
6. Click **OK**.

Add Users to the Wireless Users Security Group

You can use this procedure to add a user, computer, or group to your wireless security group in the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in.

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

To add users to the wireless security group

1. Click **Start**, click **Administrative Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is example.com,

click **example.com**.

2. In the details pane, double-click the folder that contains your wireless security group.
3. In the details pane, right-click the wireless security group, and then click **Properties**. The **Properties** dialog box for the security group opens.
4. On the **Members** tab, click **Add**, and then complete one of the following procedures:

To add a user or group

1. In **Enter the object names to select**, type the name of the user or group that you want to add, and then click **OK**.
2. To assign group membership to other users or groups, repeat step 1 of this procedure.

To add a computer

1. Click **Object Types**. The **Object Types** dialog box opens.
2. In **Object types**, select **Computers**, and then click **OK**.
3. In **Enter the object names to select**, type the name of the computer that you want to add, and then click **OK**.
4. To assign group membership to other computers, repeat steps 1-3 of this procedure.

Configuring Wireless Network (IEEE 802.11) Policies

Follow these steps to configure Wireless Network (IEEE 802.11) Policies Group Policy extension:

- **Open or Add and Open a Group Policy Object**
- **Activate Default Wireless Network (IEEE 802.11) Policies**
- **Open Wireless Network (IEEE 802.11) Policies for Editing**
- **Configure the New Wireless Network Policy**
- **Configure the New Windows XP Wireless Network Policy**

Open or Add and Open a Group Policy Object

By default, the Group Policy Management feature is installed on computers running Windows Server 2012 when the Active Directory Domain Services (AD DS) server role is installed. The procedure that follows describes how to open the Group Policy Management Console (GPMC) on your domain controller running Windows Server 2012. The procedure then describes how to either open an existing domain-level Group Policy object (GPO) for editing, or create a new domain GPO and open it for editing.

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To open or add and open a Group Policy object

1. On your Windows Server 2012 domain controller, click **Start**, point to **Administrative Tools**, and then double-click **Group Policy Management**. The Group Policy Management Console opens.
2. In the left pane, double-click your forest. For example, double-click **Forest: example.com**.
3. In the left pane, double-click **Domains**, and then double-click the domain for which you want to manage a Group Policy object. For example, double-click **example.com**.
4. Do one of the following:
   a. **To open an existing domain-level GPO for editing**, double click the domain that contains the Group Policy object that you want to manage, right-click the domain policy you want to manage, and then click **Edit**.
   b. **To create a new Group Policy object and open for editing**, right-click the domain for which you want to create a new Group Policy object, and then click **Create a GPO in this domain, and Link it here**.

      In **New GPO**, in **Name**, type a name for the new Group Policy object, and then click **OK**.

      Right-click your new Group Policy object, and then click **Edit**. Group Policy Management Editor opens.

Activate Default Wireless Network (IEEE 802.11) Policies

This procedure describes how to activate the default Wireless Network (IEEE 802.11) Policies by using the Group Policy Management Editor (GPME).

Note
After you activate either the Windows Vista version of the Wireless Network (IEEE 802.11) Policies or the Windows XP version, it is removed from the list of options when you right-click **Wireless Network (IEEE 802.11) Policies**, and it is added in the details pane of the GPME when you select the **Wireless Network (IEEE 802.11) Policies** node. This state remains until the wireless policy is deleted, at which time the wireless policy returns to the menu when you right-click **Wireless Network (IEEE 802.11) Policies** in the GPME. Additionally, the Windows Vista and Windows XP wireless policies are only listed in the GPME details pane when the **Wireless Network (IEEE 802.11) Policies** node is selected.

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To activate default Wireless Network (IEEE 802.11) Policies

1. On your Windows Server 2012 domain controller, do one of the following:
   a. If Group Policy Management Editor is already open, proceed to step 2.
   b. If GPME is not already open, do the following:
      i. Click **Start**, point to **Administrative Tools**, and then double-click **Group Policy**

Management. The Group Policy Management Microsoft Management Console (MMC) snap-in opens.

ii. In the left pane, double-click your forest. For example, double-click **Forest: example.com**.

iii. In the left pane, double-click **Domains**, and then double-click the domain in which you want to manage a Group Policy object (GPO). For example, double-click **example.com**.

iv. Right-click the domain-level GPO you want to manage, and then click **Edit**. The Group Policy Management Editor MMC opens. Proceed to step 2.

2. In the GPME, in the left pane, double-click **Computer Configuration**, double-click **Policies**, double-click **Windows Settings**, and then double-click **Security Settings**.
3. In **Security Settings**, right-click **Wireless Network (IEEE 802.11) Policies**, and then click either **Create a new Wireless Policy for Windows Vista and Later Releases** or **Create a new Windows XP Policy**. The properties dialog box opens for the policy you selected.
4. Click **OK**. The default policy is activated and listed in the details pane of the GPME.

To access the properties of a policy you have already created, select **Wireless Network (IEEE 802.11) Policies**. In the details pane, right-click either the Windows Vista or Windows XP policy that you want to modify, and then click **Properties**.

Open Wireless Network (IEEE 802.11) Policies for Editing

You can use this procedure to open activated Wireless Network (IEEE 802.11) Policies for editing. If a policy was previously activated, you do not need to perform this step.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To open activated Wireless Network (IEEE 802.11) Policies for editing

1. On your Windows Server 2012 domain controller, if Group Policy Management Editor is not already open, do the following: click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**. The Group Policy Management Microsoft Management Console (MMC) snap-in opens.
2. In the left pane, double-click your forest. For example, double-click **Forest: example.com**.
3. In the left pane, double-click **Domains**, and then double-click the domain in which you want to manage a Group Policy object. For example, double-click **example.com**.
4. Right-click the Group Policy object you want to manage, and then click **Edit**. For example, right-click **Default Domain Policy**, and then click **Edit**. The Group Policy Management Editor opens.

Note

The Group Policy object that you select must be the same object that you specified when you activated the Wireless Network (IEEE 802.11) Policies.

5. In Group Policy Management Editor, in the right pane, if it is not already expanded, double-click **Computer Configuration**, double-click **Policies**, double-click **Windows Settings**, double-click **Security Settings**, and then select **Wireless Network (IEEE 802.11) Policies**.
6. In the details pane: right-click either **New Wireless Network Policy** or **New XP Wireless Network Policy**, and then click **Properties**. The properties dialog box for the policy you selected opens.
   a. To open the Windows Vista policy, right-click **New Wireless Network Policy**, and then click **Properties**. The **New Wireless Network Policy Properties** dialog opens.
   b. To open the Windows XP policy, right click **New XP Wireless Network Policy**, and then click **Properties**. The **New XP Wireless Network Policy Properties** opens.

Note
Wireless network policies are not necessarily listed as **New Wireless Network Policy** or **New XP Wireless Network Policy** in the details pane of the Group Policy Management Editor. If the default policy name was previously changed from to another name, the name change is reflected in the Group Policy Management Editor details pane, but with the **Type** specified as either **Vista** or **XP**.

Configure the New Wireless Network Policy

Use the procedures in this section to configure Wireless Network (IEEE 802.11) Policies for client computers running Windows 8, Windows 7 and Windows Vista that connect to your wireless network by using 802.1X-capable wireless access points (APs). This policy enables you to configure security and authentication settings, manage wireless profiles, and specify permissions for wireless networks that are not configured as preferred networks.

- **Configure a Wireless Connection Profile for PEAP-MS-CHAP v2**
- **Set the Preference Order for Wireless Connection Profiles**
- **Define Network Permissions**

Configure a Wireless Connection Profile for PEAP-MS-CHAP v2

This procedure provides the steps required to configure a PEAP-MS-CHAP v2 wireless profile.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To configure a Windows Vista wireless connection profile for PEAP-MS-CHAP v2

1. If you have not already done so, use the steps in the topic **Open Wireless Network (IEEE 802.11) Policies for Editing** to open the Windows Vista **New Wireless Network Policy Properties** dialog.
2. On your domain controller running Windows Server 2012, in **New Wireless Network**

Policy Properties, on the **General** tab, in the **Policy Name** field, type a name for the policy, or accept the default name.

Tip

If you change the default policy name, **New Wireless Network Policy**, then both the **New Wireless Network Policy Properties** dialog and the activated policy module in Group Policy Management Editor will change to match the new policy name.

3. In **Description**, type a brief description for the policy.
4. To specify that WLAN AutoConfig is used to configure wireless network adapter settings, select **Use Windows WLAN AutoConfig service for clients**.
5. On the **General** tab, do one of the following:
   a. To add and configure a new profile, click **Add**, and then select **Infrastructure**.
   b. To edit an existing profile, select the profile you want to modify, and then click **Edit**. The **New Profile properties** dialog opens.
6. On the **Connection** tab, in the **Profile Name** field:
   a. If you are adding a new profile, it is recommended that you type a new name for the profile. For example, type **Example.com WLAN Profile for Windows Vista**.
   b. If you are editing a profile that is already added, use the existing profile name, or modify the name as needed.
7. In **Network Name(s) (SSID)**, type the SSID that corresponds to the SSID configured on your wireless APs, and then click **Add**.

   If your deployment uses multiple SSIDs and each wireless AP uses the same wireless security settings, repeat this step to add the SSID for each wireless AP to which you want this profile to apply.

   If your deployment uses multiple SSIDs and the security settings for each SSID do not match, configure a separate profile for each group of SSIDs that use the same security settings. For example, if you have one group of wireless APs configured to use WPA2-Enterprise and AES, and another group of wireless APs to use WPA-Enterprise and TKIP, configure a profile for each group of wireless APs.
8. If **NEWSSID** is present, select it, and then click **Remove**.
9. If you deployed wireless access points that are configured to suppress the broadcast beacon, select **Connect even if the network is not broadcasting**.

   Note

   Enabling this option can create a security risk because wireless clients will probe for and attempt connections to any wireless network. By default, this setting is not enabled.
10. Click the **Security** tab, click **Advanced**, and then configure the following:
    a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

       When the advanced 802.1X settings are enforced, the default values for **Max Eapol-**

Start Msgs, **Held Period**, **Start Period**, and **Auth Period** are sufficient for typical wireless deployments.

b. To enable Single Sign On, select **Enable Single Sign On for this network**.

c. The remaining default values in **Single Sign On** are sufficient for typical wireless deployments.

d. In **Fast Roaming**, select **This network uses pre-authentication**, if your wireless AP is configured for pre-authentication.

11. To specify that wireless communications meet FIPS 140-2 standards, select **Perform cryptography in FIPS 140-2 certified mode**.
12. Click **OK** to return to the **Security** tab. In **Select the security methods for this network**, in **Authentication**, select **WPA2-Enterprise** if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **WPA-Enterprise**.
13. In **Encryption**, select **AES**, if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **TKIP**.

    Note

    The settings for both **Authentication** and **Encryption** must match the settings configured on your wireless AP. The default settings for **Authentication Mode**, **Max Authentication Failures**, and **Cache user information for subsequent connections to this network** are sufficient for typical wireless deployments.

14. In **Select a network authentication method**, select **Protected EAP (PEAP)**, and then click **Properties**. The **Protected EAP Properties** page opens.
15. In **Protected EAP Properties**, confirm that **Verify the server's identity by validating the certificate** is selected.
16. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your NPS server.

    Note

    This setting limits the root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, then clients will trust all root CAs listed in their trusted root certification authority store.

17. In the **Select Authentication Method** list, select **Secured password (EAP-MS-CHAP v2)**.
18. To enable PEAP Fast Reconnect, select **Enable Fast Reconnect**.
19. If Network Access Protection (NAP) is configured on your network, select **Network Access Protection**. Otherwise, clear this check box.
20. To require server cryptobinding TLV on connection attempts, select **Disconnect if server does not present cryptobinding TLV**.
21. To specify that user identity is masked in phase one of authentication, select **Enable Identity Privacy**, and in the textbox, type an anonymous identity name, or leave the textbox blank.

    Tip

The NPS policy for 802.1X Wireless must be created by using NPS **Connection Request Policy**. If the NPS policy is created in by using NPS **Network Policy**, then identity privacy will not work.

Note
EAP identity privacy is provided by certain EAP methods where an empty or an anonymous identity (different from the actual identity) is sent in response to the EAP identity request. PEAP sends the identity twice during the authentication. In the first phase, the identity is sent in plain text and this identity is used for routing purposes, not for client authentication. The real identity—used for authentication—is sent during the second phase of the authentication, within the secure tunnel that is established in the first phase. If **Enable Identity Privacy** checkbox is selected, the username is replaced with the entry specified in the textbox. For example, assume **Enable Identity Privacy** is selected and the identity privacy alias **anonymous** is specified in the textbox. For a user with a real identity alias **jdoe@example.com**, the identity sent in first phase of authentication will be changed to **anonymous@example.com**. The realm portion of the 1st phase identity is not modified as it is used for routing purposes.

22. Click **Configure**. In the **EAP MSCHAPv2 Properties** dialog box, verify **Automatically use my Windows logon name and password (and domain if any)** is selected, click **OK**, and then click **OK** to close **Protected EAP Properties**.
23. Click **OK** to close the **Security** tab, and then click **OK** again to close the Vista Wireless Network Policy.

Set the Preference Order for Wireless Connection Profiles

This procedure provides the steps to specify the order in which wireless connection profiles are used to connect domain member wireless clients to wireless networks.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To set the preference order for wireless connection profiles

1. On your domain controller running Windows Server 2012, open **Vista Wireless Network (IEEE 802.11) Policy Properties**.
2. On the **General** tab, in **Connect to available networks in the order of profiles listed below**, select the profile that you want to reposition, and then click either the "up arrow" button or “down arrow” button to move the profile to the desired location in the list.
3. Repeat step 2 for each profile that you want to reposition in the list.
4. Click **OK** to save all changes.

Define Network Permissions

You can configure settings on the **Network Permissions** tab for your domain members running Windows 8, Windows 7 and Windows Vista to which Wireless Network (IEEE 802.11) Policies apply. You can only apply these setting for wireless networks that are not configured on the **General** tab in the **Vista Wireless Network Policy Properties** page:

- Allow or deny connections to specific wireless networks that you specify by network type and Service Set Identifier (SSID)
- Allow or deny connections to ad hoc networks
- Allow or deny connections to infrastructure networks
- Allow or deny users to view network types (ad hoc or infrastructure) to which they are denied access
- Allow or deny users to create a profile that applies to all users
- Users can only connect to allowed networks by using Group Policy profiles

Membership in **Domain Admins**, or equivalent, is the minimum required to complete these procedures.

To allow or deny connections to specific wireless networks

1. On your domain controller running Windows Server 2012, open **Vista Wireless Network (IEEE 802.11) Policy Properties**, and then click **Network Permissions**.
2. On the **Network Permissions** tab, click **Add**. The **New Permissions Entry** dialog box opens.
3. In the **New Permission Entry** dialog box, in the **Network Name (SSID)** field, type the network SSID of the network for which you want to define permissions.
4. In **Network Type**, select **Infrastructure** or **Ad hoc**.

   Note
   If you are uncertain whether the broadcasting network is an infrastructure or ad hoc network, you can configure two network permission entries, one for each network type.
5. In **Permission**, select **Allow** or **Deny**.
6. Click **OK**, to return to the **Network Permissions** tab.

To specify additional network permissions

1. On the **Network Permissions** tab, configure any or all of the following:
   a. To deny your domain members running Windows 8, Windows 7 and Windows Vista access to ad hoc networks, select **Prevent connections to ad-hoc networks**.
   b. To deny your domain members running Windows 8, Windows 7 and Windows Vista access to infrastructure networks, select **Prevent connections to infrastructure networks**.
   c. To allow your domain members running Windows 8, Windows 7 and Windows Vista to view network types (ad hoc or infrastructure) to which they are denied access,

select **Allow user to view denied networks**.

d. To allow users of computers that are running Windows 8, Windows 7 or Windows Vista to create profiles that apply to all users, select **Allow everyone to create all user profiles**.

e. To specify that your users can only connect to allowed networks by using Group Policy profiles, select **Only use Group Policy profiles for allowed networks**.

Windows 7-specific settings

1. To block users from hosting a wireless network on computers running Windows 7 that are equipped with wireless network adapters that support the Soft Access Point and Virtual Wi-Fi capability, select **Don't allow hosted networks**.

   Note
   Computers running Windows Vista are not affected by these Windows 7 settings.

2. To deny users with computers running Windows 7 to enter and store their domain credentials (username and password) which the computer can then use to authenticate to the network (even though the user is not actively logged on), in **Windows 7 Policy Settings**, select **Don't allow shared user credentials for network authentication**.

   Note
   Shared user credentials can be allowed to enable the computer to reconnect to the network after the computer is restarted. This enables the computer to continue to receive updates, such as those made through Group Policy and Windows Updates, during extended periods when a user is not actively logged on to the computer.

3. To specify the duration for which computers running Windows 7 are prohibited from making auto connection attempts to the network, select **Enable Block Period**, and in **Block Period (minutes)**, specify the number of minutes for which you want the block period to apply. The valid range of minutes is 1-60.

   Note
   For more information about the settings on any tab, press F1 while viewing that tab.

4. Click **OK** to save the settings, and close the **Network Permissions** tab.

Configure the New Windows XP Wireless Network Policy

Use the procedures in this section to configure Wireless Network (IEEE 802.11) Policies for client computers running Windows XP that connect to your wireless network by using 802.1X-capable wireless access points (APs).

To configure a Windows XP wireless connection profile for PEAP-MS-CHAP v2

This procedure provides the steps required to configure a PEAP-MS-CHAP v2 wireless profile for computers running Windows XP.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To configure a Windows XP wireless connection profile for PEAP-MS-CHAP v2

1. If you have not already done so, use the steps in the topic **Open Wireless Network (IEEE 802.11) Policies for Editing** to open the **XP Wireless Network (IEEE 802.11) Policies Properties** page.
2. On the **General** tab of the policy properties, in **XP Policy Name**, type a name for the policy, or leave the default name. In **Description**, type a brief description of the policy.

   Note
   For conceptual information about the settings on any tab of Wireless Network (IEEE 802.11) Policies, press F1 while viewing that tab.
3. In **Networks to access**, select either **Any available network (wireless AP preferred)** or **Access Point (infrastructure) network only**.
4. To specify that WLAN AutoConfig is used to configure wireless network adapter settings, select **Use Windows WLAN AutoConfig service for clients**.
5. To allow clients to automatically connect to networks that are not specifically defined on the **Preferred Networks** tab, select **Automatically connect to non-preferred networks**.
6. On the **Preferred Networks** tab, in **Networks**, click **Add**, and then select **Infrastructure**. The **Network Properties** dialog box opens.
7. On the **Network Properties** dialog box, in **Network Name (SSID)**, type the Service Set Identifier (SSID) that corresponds with the SSID configured on the wireless access point (AP).
8. In **Description**, enter a description for the wireless network.
9. If you deployed wireless APs that are configured to suppress the broadcast beacon frames, select **Connect even if network is not broadcasting**.

   Note
   Enabling this option can create a security risk because wireless clients will probe for and attempt connections to any wireless network. By default, this setting is not enabled.
10. In **Select the security methods for this network**, in **Authentication**, select either **WPA2** or **WPA**, and then in **Encryption**, specify either **AES** or **TKIP**.

    Additional considerations for these settings:

    a. In the XP Wireless Network (IEEE 802.11) Policy, WPA2 and WPA correspond to the Windows Vista Wireless Network (IEEE 802.11) Policy settings, WPA2-Enterprise

and WPA-Enterprise settings, respectively.

b. WPA2 is preferred over WPA; AES is preferred over TKIP. However, not all wireless network adapter drivers in Windows XP and Windows Vista support WPA2 or AES.

c. Selecting WPA2 exposes additional settings for Fast Roaming that are not provided by WPA. The default settings for Fast Roaming are sufficient for typical deployments.

d. Although available, do not select either WPA2-PSK or WPA-PSK. WPA2-PSK and WPA-PSK are intended for small office and home office networks, and cannot be used in this scenario.

11. Click the **IEEE 802.1X** tab. In **EAP type**, by default, **Protected EAP (PEAP)** is selected.
12. Click **Settings**. The **Protected EAP Properties** page opens.
13. On the **Protected EAP Properties** page, in **When Connecting**, do the following:

a. To specify that wireless clients must verify the authenticity of the NPS server certificate, select **Verify the server's identity by validating the certificate** (recommended).

b. To specify which RADIUS servers wireless clients must use to provide network authentication and authorization, type the name of each NPS server exactly as it appears in the **Subject** field of each RADIUS server's certificate.

c. In **Trusted Root Certification Authorities**, select the trusted root certification authority corresponding to your NPS server certificate. For example, if your domain CA in example.com is named CA-01, select **example-CA-01-CA**.

d. Click **OK** to close the **Protected EAP Properties** page, and then click **OK** to close the **XP Wireless Network (IEEE 802.11) Policies Properties** page.

Configuring your NPS Server

Follow these steps to configure NPS to perform 802.1X authentication for wireless access:

- **Register NPS in Active Directory Domain Services**
- **Configure a Wireless AP as an NPS RADIUS Client**
- **Create NPS Policies for 802.1X Wireless using a Wizard**

Register NPS in Active Directory Domain Services

You can use this procedure to register a server running Network Policy Server (NPS) in Active Directory Domain Services (AD DS) in the domain where the NPS server is a member. For NPS servers to be granted permission to read the dial-in properties of user accounts during the authorization process, each NPS server must be registered in AD DS. Registering an NPS server adds the server to the **RAS and IAS Servers** security group in AD DS.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To register an NPS server in its default domain

1. On your NPS server, click **Start**, click **Administrative Tools**, and then click **Network**

Policy Server. The NPS snap-in opens.

2. Right-click **NPS (Local)**, and then click **Register Server in Active Directory**. The **Network Policy Server** dialog box opens.
3. In **Network Policy Server**, click **OK**, and then click **OK** again.

Configure a Wireless AP as an NPS RADIUS Client

Use this procedure to configure a wireless access point (AP), also known as a *network access server (NAS)*, as a Remote Authentication Dial-In User Service (RADIUS) client by using the NPS snap-in. Unless your NPS servers are running Windows Server 2012 Enterprise or Windows Server 2012 Datacenter, you must repeat this procedure for every wireless AP that you deploy on your network.

Important

Client computers, such as wireless portable computers and other computers running client operating systems, are not RADIUS clients. RADIUS clients are network access servers—such as wireless access points, 802.1X-capable switches, virtual private network (VPN) servers, and dial-up servers—because they use the RADIUS protocol to communicate with RADIUS servers such as Network Policy Server (NPS) servers.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

To add a network access server as a RADIUS client in NPS

1. On the NPS server, click **Start**, click **Administrative Tools**, and then click **Network Policy Server**. The NPS snap-in opens.
2. In the NPS snap-in, double-click **RADIUS Clients and Servers**. Right-click **RADIUS Clients**, and then click **New**.
3. In **New RADIUS Client**, verify that the **Enable this RADIUS client** check box is selected.
4. In **New RADIUS Client**, in **Friendly name**, type a display name for the wireless access point.

   Tip

   In NPS a wireless access point is one type of device that falls within a group called network access server (NAS).

   For example, if you want to add a wireless access point (AP) named AP-01, type **AP-01**.
5. In **Address (IP or DNS)**, type the IP address or fully qualified domain name (FQDN) for the NAS.

   If you enter the FQDN, to verify that the name is correct and maps to a valid IP address, click **Verify**, and then in **Verify Address**, in the **Address** field, click **Resolve**. If the FQDN name maps to a valid IP address, the IP address of that NAS will automatically appear in **IP address**. If the FQDN does not resolve to an IP address you will receive a message indicating that no such host is known. If this occurs, verify that you have the

correct AP name and that the AP is powered on and connected to the network.

Click **OK** to close **Verify Address**.

6. In **New RADIUS Client**, in **Shared Secret**, do one of the following:
   a. To manually configure a RADIUS shared secret, select **Manual**, and then in **Shared secret**, type the strong password that is also entered on the NAS. Retype the shared secret in **Confirm shared secret**.
   b. To automatically generate a shared secret, select the **Generate** check box, and then click the **Generate** button. Save the generated shared secret, and then use that value to configure the NAS so that it can communicate with the NPS server.

      Important
      The RADIUS shared secret that you enter for your virtual AP's in NPS must exactly match the RADIUS shared secret that is configured on your actual wireless AP's. If you use the NPS option to generate a RADIUS shared secret, then you must configure the matching actual wireless AP with the RADIUS shared secret that was generated by NPS.
7. In **New RADIUS Client**, on the **Advanced** tab, in **Vendor name**, specify the NAS manufacturer name. If you are not sure of the NAS manufacturer name, select **RADIUS standard**.
8. In **Additional Options**, if you are using any authentication methods other than EAP and PEAP, and if your NAS supports the use of the message authenticator attribute, select **Access Request messages must contain the Message-Authenticator attribute**.
9. If you plan on deploying Network Access Protection (NAP) and your NAS supports NAP, select **RADIUS client is NAP-capable**.
10. Click **OK**. Your NAS appears in the list of RADIUS clients configured on the NPS server.

Create NPS Policies for 802.1X Wireless Using a Wizard

You can use this procedure to create the connection request policies and network policies required to deploy either 802.1X-capable wireless access points as Remote Authentication Dial-In User Service (RADIUS) clients to the RADIUS server running Network Policy Server (NPS).

Important
Client computers, such as wireless portable computers and other computers running client operating systems, are not RADIUS clients. RADIUS clients are network access servers—such as wireless access points, 802.1X-capable switches, virtual private network (VPN) servers, and dial-up servers—because they use the RADIUS protocol to communicate with RADIUS servers such as Network Policy Server (NPS) servers.

After you run the wizard, the following policies are created:

- One connection request policy
- One network policy

Note

You can run the New IEEE 802.1X Secure Wired and Wireless Connections wizard every time you need to create new policies for 802.1X authenticated access.

Membership in **Domain Admins**, or equivalent, is the minimum required to complete this procedure.

Create policies for 802.1X authenticated wireless by using a wizard

1. Open the NPS snap-in. If it is not already selected, click **NPS (Local)**. If you are running the NPS MMC snap-in and want to create policies on a remote NPS server, select the server.
2. In **Getting Started**, in **Standard Configuration**, select **RADIUS server for 802.1X Wireless or Wired Connections**. The text and links below the text change to reflect your selection.
3. Click **Configure 802.1X**. The Configure 802.1X wizard opens.
4. On the **Select 802.1X Connections Type** wizard page, in **Type of 802.1X connections**, select **Secure Wireless Connections**, and in **Name**, type a name for your policy, or leave the default name **Secure Wireless Connections**. Click **Next**.
5. On the **Specify 802.1X Switches** wizard page, in **RADIUS clients**, all 802.1X switches and wireless access points that you have added as RADIUS Clients in the NPS snap-in are shown. Do any of the following:
   a. To add additional network access servers (NASs), such as wireless APs, in **RADIUS clients**, click **Add**, and then in **New RADIUS client**, enter the information for: **Friendly name**, **Address (IP or DNS)**, and **Shared Secret**.
   b. To modify the settings for any NAS, in **RADIUS clients**, select the AP for which you want to modify the settings, and then click **Edit**. Modify the settings as required.
   c. To remove a NAS from the list, in **RADIUS clients**, select the NAS, and then click **Remove**.

      Warning
      Removing a RADIUS client from within the **Configure 802.1X** wizard deletes the client from the NPS server configuration. All additions, modifications, and deletions that you make within the **Configure 802.1X** wizard to RADIUS clients are reflected in the NPS snap-in, in the **RADIUS Clients** node under **NPS** / **RADIUS Clients and Servers**. For example, if you use the wizard to remove an 802.1X switch, the switch is also removed from the NPS snap-in.
6. Click **Next**. On the **Configure an Authentication Method** wizard page, in **Type (based on method of access and network configuration)**, select **Microsoft: Protected EAP (PEAP)**, and then click **Configure**.

   Tip
   If you receive an error message indicating that a certificate cannot be found for use with the authentication method, and you have configured Active Directory Certificate Services to automatically issue certificates to RAS and IAS servers on your network, first ensure that you have followed the steps to

Register NPS in Active Directory Domain Services, then use the following steps to update Group Policy: Click **Start**, click **Run**, and in **Open**, type **gpupdate**, and the press ENTER. When the command returns results indicating that both user and computer Group Policy have updated successfully, select **Microsoft: Protected EAP (PEAP)** again, and then click **Configure**.

If after refreshing Group Policy you continue to receive the error message indicating that a certificate cannot be found for use with the authentication method, the certificate is not being displayed because it does not meet the minimum server certificate requirements as documented in the Core Network Companion Guide: Server Certificate Deployment. If this happens, you must discontinue NPS configuration, revoke the certificate issued to your NPS server(s), and then follow the instructions to configure a new certificate by using the version of deployment guide that corresponds to the operating system installed on your CA.

a. For Windows Server 2012, the **Core Network Companion Guide: Server Certificate Deployment**, available for download in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=251761) and in HTML format in the Windows 8 Technical Library (http://technet.microsoft.com/library/jj125379.aspx).

b. For Windows 8, the **Core Network Companion Guide: Server Certificate Deployment**, available for download in Word format at the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=29691) and in HTML format in the Windows 8 Technical Library (http://technet.microsoft.com/en-us/library/jj125379.aspx).

7. On the **Edit Protected EAP Properties** wizard page, in **Certificate issued**, ensure that the correct NPS server certificate is selected, and then do the following:

   Note

   Verify that the value in **Issuer** is correct for the certificate selected in **Certificate issued**. For example, the expected issuer for a certificate issued by a CA running Windows Server 2012 Active Directory Certificate Services (AD CS) named corp-DC1, in the domain contoso.com, is **corp-DC1-CA**.

   a. To allow users to roam with their wireless computers between access points without requiring them to reauthenticate each time they associate with a new AP, select **Enable Fast Reconnect**.

   b. To specify that connecting wireless clients will end the network authentication process if the RADIUS server does not present cryptobinding Type-Length-Value (TLV), select **Disconnect Clients without Cryptobinding**.

   c. To modify the policy settings for the EAP type, in **EAP Types**, click **Edit**, in **EAP MSCHAPv2 Properties**, modify the settings as needed, and then click **OK**.

8. Click **OK**. The Edit Protected EAP Properties dialog box closes, returning you to the **Configure 802.1X** wizard. Click **Next**.

9. In **Specify User Groups**, click **Add**, and then type the name of the security group that

you configured for your wireless clients in the Active Directory Users and Computers snap-in. For example, if you named your wireless security group Wireless Group, type **Wireless Group**. Click **Next**.

10. Click **Configure** to configure RADIUS standard attributes and vendor-specific attributes for virtual LAN (VLAN) as needed, and as specified by the documentation provided by your wireless AP hardware vendor. Click **Next**.
11. Review the configuration summary details, and then click **Finish**.

Joining New Wireless Computers to the Domain

The easiest method to join new wireless computers that are running Windows 8, Windows 7, Windows Vista, or Windows XP to the domain is to physically attach the computer to a segment of the wired LAN (a segment not controlled by an 802.1X switch) before joining the computer to the domain. This is easiest because wireless group policy settings are automatically and immediately applied and, if you have deployed your own PKI, the computer receives the CA certificate and stores it in the Trusted Root Certification Authorities certificate store, allowing the wireless client to trust NPS servers with server certs issued by your CA.

Likewise, after a new wireless computer is joined to the domain, the preferred method for users to log on to the domain is to perform log on by using a wired connection to the network.

Computers running Windows Vista

In cases where it is not practical to join computers running Windows Vista to the domain by using a wired Ethernet connection, or in cases where the user cannot log on to the domain for the first time by using a wired connection, you must use an alternative method. This guide provides the following alternative methods to configure profiles that allow users to join computers to the domain and then log on, or log on to the domain by using a wireless connection:

Method 1. A member of the IT staff joins a wireless computer running Windows Vista to the domain and configures a Single Sign On bootstrap wireless profile. In this method, the IT administrator connects the wireless computer to the wired Ethernet network and joins the computer to the domain. Then the administrator distributes the computer to the user. When the user starts the computer without using a wired connection, the domain credentials that they manually specify for the user logon are used to both establish a connection to the wireless network and to log on to the domain.

For more information, see **Join the Domain and Log On by using Wireless Method 1**.

Method 2. The user manually configures the wireless computer running Windows Vista with bootstrap wireless profile and joins the domain based on instructions from an IT administrator. The bootstrap wireless profile allows the user to establish a wireless connection and then join the domain. After joining the computer to the domain and restarting the computer, the user can log on to the domain by using a wireless connection and their domain account credentials.

For more information, see **Join the Domain and Log On by using Wireless Method 2**.

Computers running Windows XP

In cases where it is not practical to either join computers running Windows XP to the domain by using a wired Ethernet connection or the user cannot log on to the domain for the first time by using a wired connection, you must manually configure a connection profile. This guide provides the following alternative methods to configure a profile that allows users to join computers running Windows XP to the domain by using a wireless connection, and then log on to the domain by using a wireless connection:

Method 3. The user manually configures the properties of the wireless connection in **Network Connections** on the wireless computer running Windows XP, based on instructions from an IT administrator. The configuration allows the user to establish a wireless connection and then join the domain. After the computer is joined to the domain and restarted, the user can log on to the domain by using a wireless connection and their domain account credentials. For more information, see **Join the Domain and Log On by using Wireless Method 3** .

Join the Domain and Log On by using Wireless Method 1

Domain member users with domain-joined wireless client computers running Windows Vista can use a temporary wireless profile to connect to an 802.1X-authenticated wireless network without first connecting to the wired LAN. This temporary wireless profile, known as a *bootstrap wireless profile*, requires the user to manually specify their domain user account credentials, and does not validate the certificate of the Remote Authentication Dial-In User Service (RADIUS) server running Network Policy Server (NPS). After establish wireless connectivity, Group Policy is applied on the wireless client computer, and a new wireless profile is issued. The new policy automatically uses the computer and user account credentials for client authentication. Additionally, as part of the PEAP-MS-CHAP v2 mutual authentication, the client validates the credentials of the RADIUS server.

After you join the computer running Windows Vista to the domain, use this procedure to configure a Single Sign On bootstrap wireless profile, before distributing the wireless computer to the domain-member user.

To configure a Single Sign On bootstrap wireless profile

1. Configure a bootstrap profile by using the procedure **Configure a Wireless Connection Profile for PEAP-MS-CHAP v2** with the following settings specified:
   a. PEAP-MS-CHAP v2 authentication
   b. Validate RADIUS server certificate disabled
   c. Single Sign On enabled
2. In Windows Vista Wireless Network (IEEE 802.11) Policies, on the **General** tab, click **Export** to export the profile to a network share, USB flash drive, or other easily accessible location.
3. Join the new wireless computer to the domain (for example, through an Ethernet connection that does not require IEEE 802.1X authentication) and add the bootstrap wireless profile to the computer by using the **netsh wlan add profile** command.

   Note

For more information, see Netsh Commands for Wireless Local Area Network (WLAN) at http://technet.microsoft.com/library/dd744890.aspx.

4. Distribute the new wireless computer to the user with the procedure to “Log on to the domain using computers running Windows Vista.”

When the user starts the computer, Windows Vista prompts the user to enter their domain user account name and password. Because Single Sign On is enabled, the computer uses the domain user account credentials to first establish a connection with the wireless network and then log on to the domain.

Log on to the domain using computers running Windows Vista

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click **Switch User**, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain example.com with an account named **User-01**, type **example\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Join the Domain and Log On by using Wireless Method 2

In this method, you complete the steps in the General steps section, then you provide your domain-member users with the instructions about how to manually configure a wireless computer that is running Windows 8, Windows 7 or Windows Vista® with a bootstrap wireless profile. The bootstrap wireless profile allows the user to establish a wireless connection and then join the domain. After the computer is joined to the domain and restarted, the user can log on to the domain through a wireless connection.

General steps

1. Configure a local computer administrator account, in **Control Panel**, for the user.

   Important
   To join a computer to a domain, the user must be logged on to the computer with the local Administrator account. Alternatively, the user must provide the credentials for the local Administrator account during the process of joining the computer to the domain. In addition, the user must have a user account in the domain to which the user wants to join the computer. During the process of joining the computer to the domain, the user will be prompted for domain account credentials (user name and password).

2. Provide your domain users with the instructions for configuring a bootstrap wireless profile, as documented in the following procedure: To configure a Bootstrap Wireless Profile. Additionally, provide the user with both the local computer credentials (user name and password), and domain credentials (domain user account name and password) in

the form *DomainName**UserName*, as well as the procedures to “Join the computer to the domain,” and to “Log on to the domain,” as documented in the Windows 8 Core Network Guide.

Note

After completing the general steps, provide the following procedures to users at your organization who will connect to your wireless network with computers running Windows 8, Windows 7 or Windows Vista.

To configure a bootstrap wireless profile

1. Use the credentials provide to you by your network administrator or IT support professional to log on to the computer with the local computers administrator account.
2. Click **Start**, click **Connect to**, and then click **Set up a connection or network**. The **Connect to a Network** dialog box opens.
3. Click **Manually connect to a wireless network**, and then click **Next**.
4. In **Manually connect to a wireless network**, in **Network name**, type the SSID name of the AP.
5. In **Security type**, select the setting provided by your administrator.
6. In **Encryption type**, select the setting provided by your administrator.
7. Select **Start this connection automatically**, and then click **Next**.
8. In **Successfully added** *Your Network SSID*, click **Change connection settings**.
9. Click **Change connection settings**. The *Your Network SSID* Wireless Network property dialog box opens.
10. Click the **Security** tab, and then in **Choose a network authentication method**, select **Protected EAP (PEAP)**.
11. Click **Settings**. The **Protected EAP (PEAP) Properties** page opens.
12. In the **Protected EAP (PEAP) Properties** page, clear **Validate server certificate**, click **OK** twice, and then click **Close**.
13. Windows Vista attempts to connect to the wireless network. The settings of the bootstrap wireless profile specify that you must provide your domain credentials. When Windows Vista prompts you for an account name and password, type your domain account credentials as follows: *Domain Name**User Name*, *Domain Password*.

To join a computer running to the domain

1. Log on to the computer with the local Administrator account.
2. Click **Start**, right-click **Computer**, and then click **Properties**. The **System** dialog box opens.
3. In **Computer name, domain, and workgroup settings**, click **Change settings**. The **System Properties** dialog box opens.

   Note

   On computers running Windows Vista, before the **System Properties** dialog box

opens, the **User Account Control** dialog box opens, requesting permission to continue. Click **Continue** to proceed.

4. Click **Change**. The **Computer Name/Domain Changes** dialog box opens.
5. In **Computer Name**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is example.com, type **example.com**.
6. Click **OK**. The **Windows Security** dialog box opens.
7. In **Computer Name/Domain Changes**, in **User name**, type the user name, and in **Password**, type the password, and then click **OK**. The **Computer Name/Domain Changes** dialog box opens, welcoming you to the domain. Click **OK**.
8. The **Computer Name/Domain Changes** dialog box displays a message indicating that you must restart the computer to apply the changes. Click **OK**.
9. On the **Computer Name** tab of the **System Properties** page, click **Close**. The **Microsoft Windows** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes. Click **Restart Now**.

Log on to the domain

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The logon screen appears.
3. Click **Switch User**, and then click **Other User**.
4. In **User name**, type your domain and user name in the format *domain\user*. For example, to log on to the domain example.com with an account named **User-01**, type **example\User-01**.
5. In **Password**, type your domain password, and then click the arrow, or press ENTER.

Join the Domain and Log On by using Wireless Method 3

In this method, users manually configure the wireless connection settings on computers that are running Windows XP, based on instructions that you provide them. The configuration allows users to establish a wireless connection and then join the domain. After the computer is joined to the domain and restarted, users can log on to the domain by using a wireless connection and their domain account credentials.

General steps

1. The IT administrator configures a local computer administrator account, in **Control Panel**, for the user.

   Important
   To configure the necessary wireless connection settings and join a computer to a domain, the user must be logged on to the computer with a local Administrator account. In addition, the user must have a user account in the domain to which the user wants to join the computer. During the process of joining the computer

to the domain, the user will be prompted for domain account credentials (user name and password).

2. The IT administrator provides the user with the following items:
   - A wireless computer running Windows XP with Service Pack 2.
   - The instructions to manually configure the wireless connection settings in the properties of the wireless connection in **Network Connections**, as documented in the procedure that follows, “Manually configure wireless connection settings for Windows XP.”

     When you (the IT administrator) give your users the instructions to manually configure their wireless connection settings, you must provide the following values for your wireless network:
     - **Network name (SSID)**: Specify the service set identifier for your wireless network.
     - **Connect even if this network is not broadcasting**: Specify whether to clear this check box if your wireless APs are not configured to suppress the broadcast beacon or to select this check box if your wireless APs are configured to suppress the broadcast beacon.
     - **Network Authentication**: Specify Wi-Fi Protected Access (WPA) or WPA2.
     - **Data encryption**: Specify Temporal Key Integrity Protocol (TKIP) or Advanced Encryption Standard (AES).
     - **Validate server certificate**: Specify whether to clear this check box if you have deployed a private certification authority (CA) on your network to issue server certificates to your NPS servers or whether to select this check box if you have purchased server certificates for your servers running Network Policy Server (NPS) from a public CA which are already trusted by your network clients.
     - **Trusted Root Certification Authorities**: Specify the name of the CA that issued the server certificate to your NPS server.

       Tip

       The procedure “Manually configure wireless connection settings for Windows XP” provides a step to specify the CA that issued server certificates to your NPS servers. To simplify the instructions given to your users for configuring wireless connection settings, retain this step regardless of whether **Validate server certificate** is selected or cleared.
   - Both the local computer credentials (user name and password), and domain credentials (domain user account name and password) in the form *DomainName**UserName*.
   - The procedures about joining the computer to the domain, and about logging on to the domain, which are provided later in this topic and documented in the Windows 8 Foundation Network Guide.

Manually configure wireless connection settings for Windows XP

1. Log on to the computer with your local computer Administrator account.
2. Click **Start**, point to **Connect To**, and then click **Show all connections**. Network Connections opens.
3. In **Network Connections**, in **LAN or High-Speed Internet**, right-click your wireless network connection, and then click **Properties**. The Wireless Network Connection Properties dialog box opens.
4. In **Wireless Network Connection Properties**, click the **Wireless Networks** tab.
5. On the **Wireless Networks** tab, click **Add**. The Wireless Network properties dialog box opens.
6. On the **Association** tab, in **Network name (SSID)**, type the SSID value specified by your administrator.
7. In **Connect even if this network is not broadcasting**, select or clear the check box as specified by your administrator. By default, this check box is cleared.
8. In **Network Authentication**, select **TKIP** or **AES** as specified by your administrator.
9. Ensure that **This is a computer-to-computer (ad-hoc) network: wireless access points are not used** is cleared, and then click the **Authentication** tab.
10. On the **Authentication** tab, ensure that **Enable IEEE 802.1x authentication for this network** is selected.
11. In **EAP type**, select **Protected EAP (PEAP)**, and then click **Properties**. The Protected EAP Properties dialog box opens.
12. On the **Protected EAP Properties** dialog box, in **Validate server certificate**, select or clear the check box as specified by your administrator.
13. In **Trusted Root Certification Authorities**, select the CA specified by your administrator.
14. In **Select Authentication Method**, ensure **Secured password (EAP-MSCHAP v2)** is selected, and then click **OK** two times. The Protected EAP Properties dialog box closes, and the Authentication tab closes, returning you to the Wireless Networks tab.
15. On the **Wireless Network** tab, in **Preferred networks**, select the SSID that you specified in step 6 of this procedure, and then click **Move up** until the wireless network is positioned at the top of the list in **Preferred networks**.
16. Click **OK**. The Wireless Network Connection Properties closes, returning you to Network Connections.
17. Close **Network Connections**.

To join computers running Windows XP to the domain

1. Log on to the computer with your local computer Administrator account.
2. Click **Start**, right-click **My Computer**, and then click **Properties**. The **System Properties** dialog box opens.
3. Click **Change**. The **Computer Name Changes** dialog box opens.
4. In **Computer Name Changes**, in **Member of**, select **Domain**, and then type the name of the domain you want to join. For example, if the domain name is example.com, type

example.com.

5. Click **OK**. The **Computer Name Changes** dialog box opens. In **User name**, type the domain user account name, and in **Password**, type the domain user password, and then click **OK**.
6. The **Computer Name Changes** dialog box opens, welcoming you to the domain.
7. Click **OK**. The **Computer Name Changes** dialog box displays a message indicating that you must restart the computer to apply the changes.
8. Click **OK**.
9. In the **System Properties** dialog box, on the **Computer Name** tab, click **OK**, to close the **System Properties** dialog box. The **System Settings Change** dialog box opens, and displays a message, again indicating that you must restart the computer to apply the changes.
10. Click **Yes**.

Log on to the domain using computers running Windows XP

1. Log off the computer, or restart the computer.
2. Press CTRL + ALT + DELETE. The **Log On to Windows** dialog box appears.
3. If **Log on to** is not displayed, click **Options**.
4. In **Log on to**, in the drop down list, select your domain. For example, in the example.com domain, select EXAMPLE.
5. Type your domain and user name in the format *domain\user*. For example, to log on to the example.com domain with an account named **User-01**, type **example\User-01**.
6. In **Password**, type your domain password, and then press ENTER.

Additional Resources

For more information about the technologies in this guide, see the following resources:

- 802.1X Authenticated Wireless Access Overview in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/hh994700.aspx
- Wireless LAN Service Overview, for 802.11 wireless-capable computers running Windows Server 2012, in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/hh994698.aspx
- New Wireless Connection Processes, for computers running Windows 8, in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/jj200213.aspx
- Managing the Wireless Network (IEEE 802.11) Policies in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/hh994695.aspx
- Improvements to Certificate-based Authentication, for wireless computers running Server 8, in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/jj200227.aspx
- Extensible Authentication Protocol (EAP) Settings for Network Access in the Windows Server 2012 Technical Library at http://technet.microsoft.com/en-us/library/hh945104.aspx

- Active Directory Certificate Services in the Windows Server 2012 Technical Library at http://go.microsoft.com/fwlink/?LinkId=218045
- Active Directory Domain Services in the Windows Server 2012 Technical Library at http://go.microsoft.com/fwlink/?LinkId=96418
- Domain Name System (DNS) in the Windows Server 2012 Technical Library at http://technet.microsoft.com/library/hh831667.aspx
- Group Policy in the Windows Server 2012 Technical Library at http://technet.microsoft.com/library/hh831791.aspx
- Netsh Commands for Wireless Local Area Network (WLAN) in the Windows 8 Technical Library at http://go.microsoft.com/fwlink/?LinkID=81752
- Network Policy Server (NPS) in the Windows Server 2012 Technical Library at http://go.microsoft.com/fwlink/?LinkId=104545 and Network Policy Server at http://go.microsoft.com/fwlink/?LinkId=93758

Core Network Companion Guide: Computer and User Certificates Deployment

The Windows Server® 2012 Core Network Guide provides instructions for planning and deploying the components required for a fully functioning network and a new Active Directory® domain in a new forest.

This companion guide explains how to build on the core network by providing instructions for deploying client computer and user certificates with Active Directory Certificate Services (AD CS).

Note
This guide is available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/p/?LinkId=255201).

This guide contains the following sections.

- Prerequisites for using this guide
- About this guide
- What this guide does not provide
- Technology overviews
- Certificate Deployment Overview
- Certificate Deployment Planning
- Certificate Deployment
- Additional Resources

Prerequisites for using this guide

This is a companion guide to the Windows Server 2012 Core Network Guide. To deploy computer and user certificates with this guide, you must first do the following.

1. Deploy a core network using the Core Network Guide, or already have the technologies provided in the Core Network Guide installed and functioning correctly on your network. These technologies include TCP/IP v4, DHCP, Active Directory Domain Services (AD DS), DNS, NPS, and Web Server (IIS).

   Notes

   The Windows Server 2012 Core Network Guide is available in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154884).

   The Core Network Guide is also available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=157742).

2. If you are using computer certificates or user certificates for network access authentication with certificate-based authentication methods, you must deploy server certificates to NPS servers, RRAS servers, or both by using the Core Network Companion Guide: Server Certificate Deployment; or you must have already deployed a public key infrastructure (PKI) and server certificates that meet the requirements for network access authentication.

   Notes

   The Windows Server 2012 Core Network Companion Guide: Server Certificate Deployment is available in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/p/?LinkId=251948).

   The Core Network Companion Guide: Server Certificate Deployment is also available in Word format at the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=30131).

About this guide

This guide provides instructions for deploying client computer and user certificates to domain member computers and domain users by using AD CS.

Certificates are used for network access authentication because they provide strong security for authenticating users and computers and they eliminate the need for less secure password-based authentication methods.

When you deploy Extensible Authentication Protocol with Transport Layer Security (EAP-TLS) or Protected EAP with TLS (PEAP-TLS), certificates are required for the authentication of servers, and for either client computers or users, during network connection attempts through network access servers such as 802.1X-capable switches and wireless access points, virtual private network (VPN) servers, and computers running Windows Server 2012 and Remote Desktop Gateway (RD Gateway) or Windows Server 2008 and Terminal Services Gateway (TS Gateway).

Note

All of these network access servers are also called Remote Authentication Dial-In User Service (RADIUS) clients, because they use the RADIUS protocol to send connection requests and other RADIUS messages to RADIUS servers. RADIUS servers process the connection requests and perform authentication and authorization. The RADIUS server and proxy in Windows Server 2012 is Network Policy Server (NPS).

Deploying certificates with AD CS for EAP and PEAP certificate-based authentication methods provides the following benefits:

- Security provided with certificate-based authentication, which binds the identity of the server running NPS, RRAS server, user, or client computer to a private key
- A cost-effective and secure method for managing certificates, allowing you to automatically enroll, renew, and revoke certificates to domain member computers and domain users
- An efficient method for managing certification authorities (CAs)
- The ability to deploy other types of certificates that are used for purposes other than computer, user, or server authentication. For example, you can deploy certificates that provide users with the ability to digitally sign e-mail, or you can issue certificates used for software code signing.

This guide is designed for network and system administrators who have followed the instructions in the Windows Server 2012 Core Network Guide to deploy a network, or for those who have previously deployed the technologies included in the core network, including Active Directory Domain Services (AD DS), Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), TCP/IP, Network Policy Server (NPS), and Web Server (IIS).

It is recommended that you review the design and deployment guides for each of the technologies that are used in this deployment scenario. These guides can help you determine whether this deployment scenario provides the services and configuration that you need for your organization network.

Requirements for deploying computer and user certificates

Following are the requirements for deploying client computer and user certificates using autoenrollment:

- AD DS is installed, as are other network technologies, according to the instructions in the Windows Server 2012 Core Network Guide.
- To perform autoenrollment of client computer and user certificates, your CA must be running the Windows Server 2008 or Windows Server® 2008 R2 Enterprise or Datacenter operating system and must be an issuing CA; or your CA must be running the Windows Server 2012 Standard, Enterprise, or Datacenter operating system and must be an issuing CA. Although AD CS can be deployed on a single server, many deployments use a two tier public key infrastructure with multiple servers configured as subordinate CAs.
- To deploy EAP-TLS or PEAP-TLS, you must enroll server certificates to NPS servers and, if you are using RRAS servers as virtual private network (VPN) servers, to computers running Windows Server 2008, Windows Server® 2008 R2, or Windows Server 2012 and Routing and Remote Access service (RRAS). This guide assumes that you have autoenrolled server certificates in accordance with the Windows Server 2012 Core Network Companion Guide: Deploying Server Certificates, which is available in HTML format in the Windows Server 2012 Technical Library (http://technet.microsoft.com/en-us/library/jj125379). This paper is also available in Word format at the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=19224).

Note

If you deploy one or more RRAS servers as VPN servers and you do not have NPS installed, network policies and the authentication methods included in these policies are configured individually per RRAS server, which can be time-consuming and can create opportunities for configuration errors. When you install NPS, you can configure RRAS servers as RADIUS clients in NPS, and then use NPS to centrally manage all policies and authentication methods used per policy.

- To deploy PEAP or EAP for VPN, you must deploy Routing and Remote Access configured as a VPN server. The use of NPS is optional; however, if you have multiple VPN servers, using NPS is recommended for ease of administration and for the RADIUS accounting services that NPS provides.
- To deploy PEAP or EAP for TS Gateway in Windows Server 2008 or RD Gateway in Windows Server 2012, you must deploy TS Gateway and NPS or RD Gateway and NPS, respectively.
- To deploy PEAP or EAP for 802.1X secure wired or wireless, you must deploy NPS and additional hardware, such as 802.1X-capable switches or wireless access points.

What this guide does not provide

This guide does not provide information about the following:

- How to deploy client computer and user certificates with smart cards.
- How to deploy server certificates with autoenrollment.
- How to design and deploy a public key infrastructure (PKI) by using AD CS. It is recommended that you review AD CS design and deployment documentation before deploying the technologies in this guide. For more information, see **Additional Resources**.
- How to deploy the network access technologies for which server certificates can be used. There might be other companion guides available that provide instructions for deploying these network access solutions. You might also want to review the NPS documentation for this information.

Technology overviews

Following are technology overviews for client computer and user certificates, EAP, PEAP, and AD CS.

AD CS

AD CS in Windows Server 2012 provides customizable services for creating and managing the X.509 certificates that are used in software security systems that employ public key technologies. Organizations can use AD CS to enhance security by binding the identity of a person, device, or service to a corresponding public key. AD CS also includes features that allow you to manage certificate enrollment and revocation in a variety of scalable environments.

Client computer and user certificates

When you deploy EAP-TLS or PEAP-TLS, you can deploy computer certificates for client computer authentication, user certificates for user authentication, or both.

Note

EAP does not provide mechanisms that perform dual authentication — that is, the authentication of both the computer being used to access the network and the user who is attempting to connect. For this reason, you are not required to issue both computer and user certificates when you deploy EAP and PEAP with certificate-based authentication types.

There are two methods for deploying client computer and user certificates:

- **Using smart cards**. When you deploy certificates using smart cards, you must purchase additional hardware to imprint certificates on user identification or other cards that your employees use to log on to the network. In addition, users must be supplied with smart card readers, which are used to read the certificate that is imprinted on the smart card when they log on.

  Important

  This guide does not provide information about how to deploy client computer and user certificates with smart cards.

- **Using autoenrollment**. When you deploy certificates using autoenrollment, you configure the CA to automatically enroll certificates to computers that are members of the Domain Computers group and to users who are members of the Domain Users group. No additional hardware is required to autoenroll certificates, because the certificates are stored on the computer that is connecting to the network. When a computer receives a computer or user certificate from the CA, the certificate is stored locally in a data store named the certificate store.

Important

You should enroll certificates only to the computers and users to whom you want to grant network access through RADIUS clients. You do not have to autoenroll certificates to all members of the Domain Users and Domain Computers groups. Instead, you can issue certificates to subsets of the Domain Users and Domain Computers groups, such as to the Sales team or the Accounting department. To enroll certificates to other groups, create the groups and then add members to the groups in Active Directory Users and Computers. In the Certificate Templates snap-in, remove the Domain Users or Domain Computers groups from the access control list (ACL) on the certificate templates (the User template and the Workstation Authentication template, respectively), and then add the groups that you created to the template.

Certificate store

On computers that are running the Windows operating system, certificates that are installed on the computer are kept in a storage area called the *certificate store*. The certificate store is accessible using the Certificates Microsoft Management Console (MMC) snap-in.

This store contains multiple folders, where certificates of different types are stored. For example, the certificate store contains a Trusted Root Certification Authorities folder where the certificates from all trusted root CAs are kept.

When your organization deploys a PKI and installs a private trusted root CA using AD CS, the CA automatically sends its certificate to all domain member computers in the organization. The domain member client and server computers store the CA certificate in the Trusted Root Certification Authorities folder in the Current User and the Local Computer certificate stores. After this occurs, the domain member computers trust certificates that are issued by the trusted root CA.

Similarly, when you autoenroll computer certificates to domain member client computers, the certificate is kept in the Personal certificate store for the Local Computer. When you autoenroll certificates to users, the user certificate is kept in the Personal certificate store for the Current User.

EAP

Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by allowing arbitrary authentication methods that use credential and information exchanges of arbitrary lengths. EAP was developed in response to the demand for authentication methods that use security devices such as smart cards, token cards, and crypto calculators. EAP provides an industry-standard architecture for supporting additional authentication methods within PPP.

With EAP, an arbitrary authentication mechanism is used to verify the identities of the client and server that are establishing a network access connection. The authentication scheme to be used is negotiated by the access client and the authenticator (the network access server or the RADIUS server).

For successful authentication to occur, both the network access client and the authenticator (such as the server running NPS) must support the same EAP type.

Important

Strong EAP types (such as those that are based on certificates) offer better protection against brute-force attacks, dictionary attacks, and password guessing attacks than password-based authentication protocols (such as CHAP or MS-CHAP, version 1).

EAP in Windows Server 2012

Windows Server 2012 includes an EAP infrastructure, two EAP types, and the ability to pass EAP messages to a RADIUS server (EAP-RADIUS) such as NPS.

By using EAP, you can support additional authentication schemes, known as EAP types. The EAP types that are supported by Windows Server 2012 are:

- **Transport Layer Security (TLS)**. EAP-TLS requires the use of computer certificates or user certificates, in addition to server certificates that are enrolled to computers running NPS.
- **Microsoft Challenge-Handshake Authentication Protocol, version 2 (MS-CHAP v2)**. This EAP type is a password-based authentication protocol. When used within EAP as the authentication method EAP-MS-CHAP v2, NPS and RRAS servers provide a server certificate as proof of identity to client computers, while users prove their identity with a user name and password.
- **Tunneled Transport Layer Security (TTLS)**. EAP-TTLS is new in Windows Server 2012 and is not available in other versions of Windows Server. EAP-TTLS is a standards-based EAP tunneling method that supports mutual authentication. EAP-TTLS provides a secure tunnel for client authentication using EAP methods and other legacy protocols. EAP-TTLS also provides you with the ability to configure EAP-TTLS on client computers for network access solutions in which non-Microsoft Remote Authentication Dial In User Service (RADIUS) servers that support EAP-TTLS are used for authentication.

In addition, you can install other non-Microsoft EAP modules on the server running NPS or Routing and Remote Access to provide other EAP authentication types. In most cases, if you install additional EAP types on servers, you must also install matching EAP client authentication components on client computers so that the client and server can successfully negotiate an authentication method to use for connection requests.

PEAP

PEAP uses TLS to create an encrypted channel between an authenticating PEAP client, such as a wireless computer, and a PEAP authenticator, such as a server running NPS or other RADIUS server.

PEAP does not specify an authentication method, but it provides additional security for other EAP authentication protocols (such as EAP-MSCHAP v2) that can operate through the TLS-encrypted channel provided by PEAP. PEAP is used as an authentication method for access clients that are connecting to your organization's network through the following types of network access servers:

- 802.1X-capable wireless access points
- 802.1X-capable authenticating switches
- Computers running Windows Server 2012 or Windows Server 2008 R2 and RRAS that are configured as VPN servers
- Computers running Windows Server 2012 or Windows Server 2008 R2 and RD Gateway

Features of PEAP

To enhance the EAP protocols and network security, PEAP provides:

- A TLS channel that provides protection for the EAP method negotiation that occurs between the client and server. This TLS channel helps prevent an attacker from injecting packets between the client and the network access server to cause the negotiation of a less secure EAP type. The encrypted TLS channel also helps prevent denial of service attacks against the server running NPS.
- Support for the fragmentation and reassembly of messages, which allows the use of EAP types that do not provide this functionality.

- Clients with the ability to authenticate the NPS or other RADIUS server. Because the server also authenticates the client, mutual authentication occurs.
- Protection against the deployment of an unauthorized wireless access point at the moment when the EAP client authenticates the certificate provided by the server running NPS. In addition, the TLS master secret that is created by the PEAP authenticator and the client is not shared with the access point. Because of this, the access point cannot decrypt the messages that are protected by PEAP.
- PEAP fast reconnect, which reduces the delay between an authentication request by a client and the response by the NPS or other RADIUS server. Fast reconnect also allows wireless clients to move between access points that are configured as RADIUS clients to the same RADIUS server without repeated requests for authentication. This reduces resource requirements for the client and the server, and it minimizes the number of times that users are prompted for credentials.

EAP-TLS and PEAP-TLS deployment overview

The following are the general steps for deploying EAP-TLS or PEAP-TLS:

- Deploy network access servers (RADIUS clients) that are both EAP and RADIUS compliant.
- Autoenroll server certificates to servers running NPS and, if applicable, Routing and Remote Access VPN servers.
- Autoenroll computer certificates, user certificates, or both, to domain member computers and users, respectively, or to other groups that you have created.
- Configure network access servers as RADIUS clients in NPS.
- Configure EAP authentication in NPS or RRAS network policy.
- Ensure that client computers support EAP. By default, Windows® 8, Windows® 7, and Windows Vista® support EAP.

Group Policy

Group Policy in Windows Server 2012 is an infrastructure used to deliver and apply one or more desired configurations or policy settings to a set of targeted users and computers within an Active Directory environment. This infrastructure consists of a Group Policy engine and multiple client-side extensions (CSEs) responsible for reading policy settings on target client computers. Group Policy is used in this scenario to enroll and distribute certificates to users, computers, or both.

Certificate Deployment Overview

This topic contains the following sections.

- Computer and user certificate deployment components
- Computer and user certificate deployment process

Computer and user certificate deployment components

The following illustration depicts the components that are required to deploy client computer and user certificates.

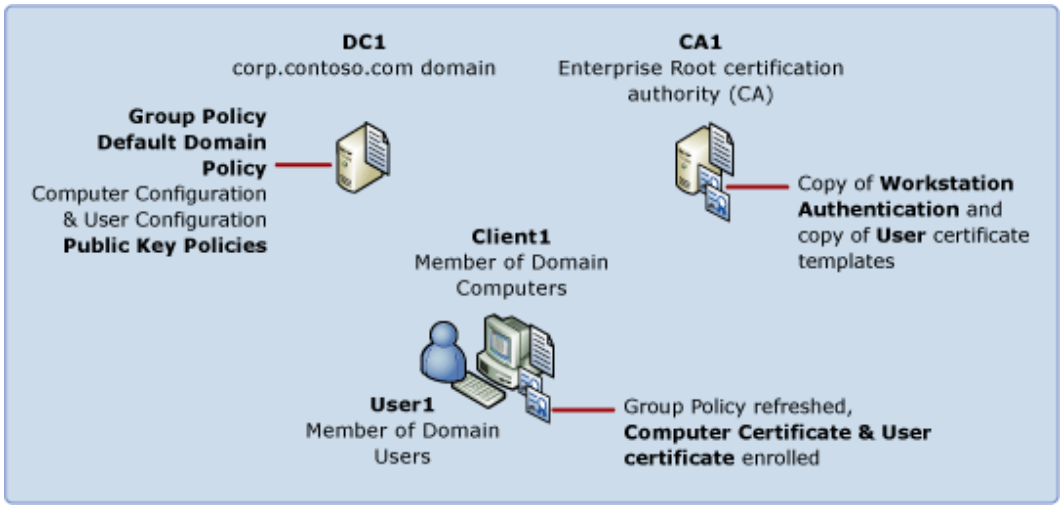

For more information on each item depicted in the illustration above, see the following:

- CA1 running the AD CS server role
- DC1 running the AD DS and DNS server roles
- Client1

The following components are required to deploy client computer and user certificates.

CA1 running the AD CS server role

The enterprise root certification authority (CA) is also an issuing CA. The CA issues certificates to computers and users that have the correct security permissions to enroll a certificate. Active Directory Certificate Services (AD CS) is installed on CA1.

In this scenario, the enterprise root CA is also an issuing CA. For larger networks or where security concerns justify it, you can separate the roles of root CA and issuing CA, and deploy subordinate CAs that are issuing CAs.

In the most secure deployments, the enterprise root CA is taken offline and physically secured. For more information, see Additional Resources.

Copy of the User certificate template

When you deploy user certificates, you make a copy of the User certificate template and then configure the template according to your requirements and the instructions in this guide. You will be using a copy rather than the original so that the configuration of the original template is preserved for possible future use. The CA uses the copy of the user template to create user

certificates that it issues to users that are members of the Domain Users group in Active Directory Users and Computers.

Copy of the Workstation Authentication certificate template

When you deploy computer certificates, you make a copy of the Workstation Authentication certificate template and then configure the template according to your requirements and the instructions in this guide. You will be using a copy rather than the original so that the configuration of the original template is preserved for possible future use. The CA uses the copy of the certificate template to create computer certificates that it issues to computers that are members of the Domain Computers group in Active Directory Users and Computers.

DC1 running the AD DS and DNS server roles

After you configure the certificate templates on the CA, you can configure the default domain policy in Group Policy so that certificates are autoenrolled to users, computers, or both. Group Policy is configured in AD DS on the server DC1.

Client1

After you have configured certificate templates and Group Policy, you can refresh Group Policy on client computers to enroll certificates.

Computer and user certificate deployment process

The process of configuring computer and user certificate enrollment occurs in these stages:

- Deploy a public key infrastructure (PKI). This step is required only if you have not already deployed a CA on your network. Instructions for installing AD CS as an enterprise root CA are included in the Core Network Companion Guide: Deploying Server Certificates, which is available in the Windows Server 2012 Technical Library. For more information, see Additional Resources.
- On CA1, which is the computer running the AD CS server role, configure copies of the Workstation Authentication and User certificate templates and add them to the CA. The CA issues certificates based on these certificate templates, so you must configure the templates for the computer and user certificates, and then add them to the CA, before the CA can issue them.
- On DC1, which is the computer running both the Active Directory Domain Services server role and the DNS server role, configure computer and user certificate autoenrollment in Group Policy. When you configure the autoenrollment of computer certificates, user certificates, or both, all domain member computers, domain users, or both will automatically receive a certificate when Group Policy on the user's computer is refreshed. If you add more computers or users later, they will automatically receive a certificate.

Note

If you removed the Domain Users or Domain Computers groups from the certificate template ACL and replaced these groups with custom groups that you created in Active Directory Users and Computers, certificates will be enrolled only to the members of your custom groups.

- Refresh Group Policy on domain member computers. When Group Policy is refreshed:
  - If you have deployed computer certificates, the domain member computer enrolls a computer certificate that is based on the template that you configured in the previous step.
  - If you have deployed user certificates, the domain user enrolls a user certificate that is based on the template that you configured in the previous step.
  - If you have deployed both user and computer certificates, the domain member computer enrolls a computer certificate and the user enrolls a user certificate that are based on the templates that you configured in the previous steps.
  - Because you have already deployed a PKI, the CA certificate already exists in the Trusted Root Certification Authorities folder in the certificate store for both the Current User and the Local Computer.

Note

Group Policy is automatically refreshed if you restart the domain member computer, or if a user logs on to a domain member computer. In addition, Group Policy is periodically refreshed. By default, this periodic refresh is performed every 90 minutes with a randomized offset of up to 30 minutes.

Certificate Deployment Planning

Follow these steps for computer and user certificates deployment planning.

- Planning certificate configuration
- Planning User certificate enrollment

Planning certificate configuration

The procedures for configuring certificate templates allow you to autoenroll computer and user certificates that meet X.509, Secure Sockets Layer/Transport Layer Security (SSL/TLS), Extensible Authentication Protocol (EAP), and Protected Extensible Authentication Protocol (PEAP) requirements.

Computer and user certificates have the following minimum requirements:

- The Subject name must contain a value. If you issue a certificate that has a blank Subject, the certificate cannot be used for authentication.
- The certificate chains to a trusted root CA and does not fail any of the checks that are performed by CryptoAPI and that are specified in NPS network policy. Because you are issuing certificates from your own enterprise root CA, computer and user certificates issued by the CA automatically chain to the CA. Domain member computers have the CA certificate

in the Trusted Root Certification Authorities folder in both the Local Computer and Current User certificate stores, which means that they trust the CA.

- The user or computer certificate is configured with the Client Authentication purpose in Application Policies extensions (also called Enhanced Key Usage or EKU extensions). The object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2. By default, the User and Workstation Authentication certificate templates contain this purpose in Application Policies extensions.
- For user certificates, the Subject Alternative Name (SubjectAltName) extension, if used, must contain the user principal name (UPN). By default, the User certificate template is configured with the UPN.
- For computer certificates, the SubjectAltName extension, if used, must contain the computer's fully qualified domain name (FQDN), which is also called the DNS name. By default, the Workstation Authentication certificate template is not configured with this value and must be reconfigured to meet this requirement according to the instructions in Configure the Workstation Authentication Certificate Template.

Client computer display of certificates

With PEAP and EAP-TLS, client computers display a list of all installed certificates in the computer certificate store, with the following exceptions:

- Wireless clients do not display registry-based and smart card-logon certificates.
- Wireless clients and VPN clients do not display password-protected certificates.
- Certificates that do not contain the Client Authentication purpose in Application Policies extensions are not displayed.

In addition, 802.1X client computers do not use registry-based certificates that are either smart card-logon or password-protected.

Planning User certificate enrollment

Only domain member computers can refresh and be configured through Group Policy. For this reason, you can successfully autoenroll user certificates only if the user is connecting to the network using a domain member computer.

To enroll user certificates to users connecting to the network on non-domain member computers, you must provide users with the certificate using one of the following enrollment methods:

- Install from removable media.
- Use the Web enrollment pages provided with AD CS. For more information, see "Request a Certificate Over the Web" in the AD CS Help in the Windows Server 2008 Technical Library (http://go.microsoft.com/fwlink/?LinkId=110404).
- Deploy smart cards, with the user certificate stored on the smart card.

If users connect to the organization network with both a domain member computer and a non-domain member computer, you can autoenroll user certificates to the user when the user logs on to the domain member computer, and then provide the user with instructions for using the

Certificates Microsoft Management Console (MMC) snap-in to export the certificate to a file for import using the MMC on the non-domain member computer.

Important
Always use the Certificates MMC snap-in to import a certificate from file; do not attempt to drag and drop the certificate into the MMC. In addition, if a certificate needs to be moved in the certificate store, do not drag and drop the certificate to another location. To move a certificate, you must export and then import it.

Certificate Deployment

Follow these steps to deploy client computer certificates, user certificates, or both.

- Certificate Deployment for Computers
- Certificate Deployment for Users
- Refresh Group Policy
- Verify Computer or User Enrollment of a Certificate

Note
The procedures in this guide do not include instructions for cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Certificate Deployment for Computers

You can use the following topics to configure the Workstation Authentication certificate template, configure autoenrollment in Group Policy, and refresh Group Policy on the computer where you want to enroll a certificate.

- Configure the Workstation Authentication Certificate Template
- Configure Computer Certificate Autoenrollment
- Refresh Group Policy

Configure the Workstation Authentication Certificate Template

You can use this procedure to configure the certificate template that Active Directory® Certificate Services (AD CS) uses as the basis for computer certificates that are enrolled to domain member client computers, or to members of other groups that you specify.

Note

If you want to enroll computer certificates to members of groups other than the Domain Computers group, remove the Domain Computers group from the template's access control list (ACL) while performing this procedure, and then add the groups you prefer to the ACL. After you add new groups to the ACL, ensure that you allow Enroll and Autoenroll permissions.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure the certificate template and autoenrollment

1. On CA1, in Server Manager, click **Tools**, and then click **Certification Authority**. The Certification Authority Microsoft Management Console (MMC) opens.
2. In the MMC, double-click the CA name, right-click **Certificate Templates**, and then click **Manage**.
3. The **Certificate Templates** console opens. All of the certificate templates are displayed in the details pane.
4. In the details pane, click the **Workstation Authentication** template.
5. On the **Action** menu, click **Duplicate Template**. The **Properties** dialog box for the certificate template opens.
6. On the **General** tab, in **Display Name**, type a new name for the certificate template or keep the default name.
7. Click the **Subject Name** tab. Ensure that **Build from this Active Directory information** is selected. In **Subject name format**, select **Fully distinguished name**.
8. Click the **Security** tab. In **Group or user names**, click **Domain Computers**.
9. In **Permissions for Domain Computers**, under **Allow**, ensure that **Enroll** is selected, and then select the **Autoenroll** check box. Click **OK**, and close the Certificate Templates MMC.
10. In the Certification Authority MMC, click **Certificate Templates**. On the **Action** menu, point to **New**, and then click **Certificate Template to Issue**. The **Enable Certificate Templates** dialog box opens.
11. Click the name of the certificate template you just configured, and then click **OK**. For example, if you did not change the default certificate template name, click **Copy of Workstation Authentication**, and then click **OK**.

Configure Computer Certificate Autoenrollment

You can use this procedure to automatically enroll, or autoenroll, client computer certificates to domain member computers.

Note

In this procedure, you are instructed to enable the **Certificate Services Client - Auto-Enrollment** Group Policy setting. If you have previously deployed server certificates and configured autoenrollment of server certificates, you do not need to perform this procedure again; however, you can use this procedure to verify that Group Policy is configured correctly to autoenroll certificates.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure client computer certificate autoenrollment

1. On the computer where AD DS is installed, open Windows PowerShell®, type **mmc**, and then press ENTER. The Microsoft Management Console opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In the **Add or Remove Snap-ins** dialog box, in **Available snap-ins**, scroll down to and double-click **Group Policy Management Editor**. The Group Policy Wizard opens.

   Important

   Ensure that you select **Group Policy Management Editor** and not **Group Policy Management**. If you select **Group Policy Management**, your configuration using these instructions will fail and a server certificate will not be autoenrolled to your computers.
4. In **Select Group Policy Object**, click **Browse**. The **Browse for a Group Policy Object** dialog box opens.
5. In **Domains, OUs, and linked Group Policy Objects**, click **Default Domain Policy**, and then click **OK**.
6. Click **Finish**, and then click **OK**.
7. In the MMC, expand the following path: **Default Domain Policy**, **Computer Configuration**, **Policies**, **Windows Settings**, **Security Settings**, **Public Key Policies**.
8. Click **Public Key Policies**, and then in the details pane double-click **Certificate Services Client - Auto-Enrollment**. The **Certificate Services Client - Auto-Enrollment Properties** dialog box opens. Configure the following items, and then click **OK**:
   a. In **Configuration Model**, select **Enabled**.
   b. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.

c. Select the **Update certificates that use certificate templates** check box.

Refresh Group Policy

You can use this procedure to manually refresh Group Policy on the local computer. When Group Policy is refreshed, if certificate autoenrollment is configured and functioning correctly, the local computer is autoenrolled a certificate by the certification authority (CA).

Note
Group Policy is automatically refreshed when you restart the domain member computer, or when a user logs on to a domain member computer. In addition, Group Policy is periodically refreshed. By default, this periodic refresh is performed every 90 minutes with a randomized offset of up to 30 minutes.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure.

To refresh Group Policy on the local computer

1. Open Windows PowerShell®.
2. Type the following command, and then press ENTER: `gpupdate /force`

Certificate Deployment for Users

You can use the following topics to configure the User certificate template, configure autoenrollment in Group Policy, and refresh Group Policy on the computer where you want to enroll a User certificate.

- Configure the User Certificate Template
- Configure User Certificate Autoenrollment
- Refresh Group Policy

Configure the User Certificate Template

You can use this procedure to configure the certificate template that Active Directory® Certificate Services (AD CS) uses as the basis for user certificates that are enrolled to domain users or members of other groups that you specify.

Note

If you want to enroll user certificates to members of groups other than the Domain Users group, remove the Domain Users group from the template's access control list (ACL) while performing this procedure, and then add the groups you prefer to the ACL. After you add new groups to the ACL, ensure that you allow Enroll and Autoenroll permissions.

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure the certificate template and autoenrollment

1. On CA1, in Server Manager, click **Tools**, and then click **Certification Authority**. The Certification Authority Microsoft Management Console (MMC) opens.
2. In the MMC, double-click the CA name, right-click **Certificate Templates**, and then click **Manage**.
3. The Certificate Templates console opens. All of the certificate templates are displayed in the details pane.
4. In the details pane, click the **User** template.
5. On the **Action** menu, click **Duplicate Template**. The **Properties of New Template** dialog box opens.
6. In **Properties of New Template**, on the **General** tab, in **Display Name**, type a new name for the certificate template or keep the default name.
7. Click the **Security** tab. In **Group or user names**, click **Domain Users**.
8. In **Permissions for Domain Users**, under **Allow**, ensure that **Enroll** is selected, and then select the **Read** and **Autoenroll** check boxes.
9. Click the **Subject Name** tab. Ensure that **Build from this Active Directory** information is selected. Also ensure that **Subject name format** has the value of **Fully distinguished name**. In **Include this information in alternate subject name**, ensure that **User principal name (UPN)** is selected.

   Important
   On the **Subject Name** tab, ensure that the following items are not selected:

   - In **Subject name format**, ensure that **Include e-mail name in subject name** is not selected.
   - In **Include this information in alternate subject name**, ensure that **E-mail name** is not selected.
10. Click **OK**, and close the Certificate Templates MMC.
11. In the Certification Authority MMC, click **Certificate Templates**. On the **Action** menu, point to **New**, and then click **Certificate Template to Issue**. The **Enable Certificate Templates** dialog box opens.
12. Click the name of the certificate template you just configured, and then click **OK**. For example, if you did not change the default certificate template name, click **Copy of User**, and then click **OK**.

Configure User Certificate Autoenrollment

You can use this procedure to automatically enroll, or autoenroll, user certificates to members of the Domain Users group in Active Directory Domain Services (AD DS).

Membership in both the **Enterprise Admins** and the root domain's **Domain Admins** group is the minimum required to complete this procedure.

To configure user certificate autoenrollment

1. On the computer where AD DS is installed, open Windows PowerShell®, type **mmc**, and then press ENTER. The Microsoft Management Console opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Available snap-ins**, scroll down to and double-click **Group Policy Management Editor**. The Group Policy Wizard opens.

   Important
   Ensure that you select **Group Policy Management Editor** and not **Group Policy Management**. If you select **Group Policy Management**, your configuration using these instructions will fail and a server certificate will not be autoenrolled to your users.
4. In **Select Group Policy Object**, click **Browse**. The **Browse for a Group Policy Object** dialog box opens.
5. In **Domains, OUs, and linked Group Policy Objects**, click **Default Domain Policy**, and then click **OK**.
6. Click **Finish**, and then click **OK**.
7. In the MMC, expand the following path: **Default Domain Policy**, **User Configuration**, **Policies**, **Windows Settings**, **Security Settings**, **Public Key Policies**.
8. In the MMC, click **Public Key Policies**. In the details pane, double-click **Certificate Services Client - Auto-Enrollment**. The **Certificate Services Client - Auto-Enrollment Properties** dialog box opens. Configure the following items, and then click **OK**:
   a. In **Configuration Model**, select **Enabled**.
   b. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.
   c. Select the **Update certificates that use certificate templates** check box.

Refresh Group Policy

You can use this procedure to manually refresh Group Policy on the local computer. When Group Policy is refreshed, if certificate autoenrollment is configured and functioning correctly, the local computer is autoenrolled a certificate by the certification authority (CA).

Note
Group Policy is automatically refreshed when you restart the domain member computer, or when a user logs on to a domain member computer. In addition, Group Policy is periodically refreshed. By default, this periodic refresh is performed every 90 minutes with a randomized offset of up to 30 minutes.

Membership in **Administrators**, or equivalent, is the minimum required to complete this procedure.

To refresh Group Policy on the local computer

1. Open Windows PowerShell®.
2. Type the following command, and then press ENTER: `gpupdate /force`

Verify Computer or User Enrollment of a Certificate

You can use this procedure to view the certificates that are in the certificate store to verify that a computer certificate, a user certificate, or both are enrolled on the local computer.

To perform this procedure to verify enrollment of a computer certificate, the computer must be a member of **Domain Computers** or of the custom group that you specify, and you must be logged on to the computer with an account that has Administrator privileges. To perform this procedure to verify enrollment of a user certificate, you must be logged on to the computer with an account that is a member of the **Domain Users** group, or of the custom group that you specify.

To verify computer or user enrollment of a certificate

1. Open Windows PowerShell®, type **mmc**, and then press ENTER. The Microsoft Management Console opens.

   Note
   If you are logged on to the computer with an account that does not have Administrator privileges, and if you want to view the Local Computer certificate store, you must open Windows PowerShell with the parameter **Run as administrator**, and then provide the credentials for an account that has

Administrator privileges.

2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Available snap-ins**, click **Certificates**, and then click **Add**. Select one of the following.
   - To create a snap-in for your user account: In **Certificates snap-in**, ensure that **My user account** is selected, and then click **Finish**.
   - To create a snap-in for your computer account: In **Certificates snap-in**, ensure that **Computer account** is selected, and then click **Next**. In **Select Computer**, ensure that **Local computer** is selected, and then click **Finish**.

   Tip
   To create a snap-in that displays both the user account and computer account certificate stores, repeat the previous step and add the account snap-in that you did not select during the first selection process.
4. To close the **Add or Remove Snap-ins** dialog box, click **OK**. In the MMC, the certificate store for the **Current User**, the certificate store for the **Local Computer**, or both are displayed, depending on your selections in the previous step.
5. Expand the certificate store that you want to view, click **Personal**, and in **Object Type**, double-click **Certificates**. The certificate enrolled to the computer or user is displayed in this store. For example, if you are viewing the **Certificates (Local Computer)** certificate store and your computer is named Client1 and your domain is named corp.contoso.com, the Personal Certificates store contains a certificate named **Client1.corp.contoso.com**. In another example, if you are viewing the **Certificates – Current User** certificate store, your user name is User1 and your domain name is corp.contoso.com, the Personal Certificates store contains a certificate named **User1.corp.contoso.com**.

Additional Resources

For more information about the technologies that are discussed in this guide, see the following resources:

- The Windows Server 2012 Core Network Guide in Word format in the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=30132) and in HTML format in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh911995.aspx).
- PKI Design Guidance in the Microsoft TechNet Wiki (http://social.technet.microsoft.com/wiki/contents/articles/2901.pki-design-guidance.aspx)
- Active Directory Certificate Services (AD CS) Public Key Infrastructure (PKI) Design Guide in the Microsoft TechNet Wiki http://social.technet.microsoft.com/wiki/contents/articles/7421.ad-cs-pki-design.aspx

- Guidelines for Using Alternate Signature Formats, which is available at the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://technet.microsoft.com/library/cc753169.aspx)
- Creating Certificate Policies and Certificate Practice Statements in the Technical Library (http://technet.microsoft.com/library/cc780454.aspx)
- Request for Comments 2527, Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, at the IETF Web site (http://www.ietf.org/rfc/rfc2527.txt)
- CA Policy.inf Syntax in the Technical Library (http://technet.microsoft.com/library/cc728279.aspx)
- Deployment Planning (Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure) in Windows Server TechCenter (http://go.microsoft.com/fwlink/?LinkId=106049)
- Active Directory Domain Services in the Windows Server® 2012 Technical Library (http://technet.microsoft.com/library/hh831484.aspx)
- Group Policy in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh831683.aspx)
- Network Policy and Access Services in the Windows Server 2012 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154883)
- Windows Server 2008 R2 Core Network Guide, which is available at the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154884)
- "Configuring All Servers" topic in the Windows Server 2008 R2 Core Network Guide in the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154885)
- "Joining Computers to the Domain and Logging On" topic in the Windows Server 2008 R2 Core Network Guide in the Windows Server 2008 and Windows Server 2008 R2 Technical Library (http://go.microsoft.com/fwlink/?LinkId=154886)

Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode

The Windows Server 2012 Core Network Guide provides instructions for planning and deploying the core components required for a fully functioning network and a new Active Directory® domain in a new forest.

This guide explains how to build on the core network by providing instructions for deploying BranchCache in hosted cache mode in a branch office with a Read-Only Domain Controller where client computers are running Windows® 8 and are joined to the domain.

Important

Do not use this guide if you are planning to deploy or have already deployed a BranchCache hosted cache server that is running Windows Server® 2008 R2. This guide

provides instructions for deploying hosted cache mode with a hosted cache server that is running Windows Server 2012.

This guide contains the following sections.

- Prerequisites for using this guide
- About this guide
- What this guide does not provide
- Technology overviews
- BranchCache Hosted Cache Mode Deployment Overview
- BranchCache Hosted Cache Mode Deployment Planning
- BranchCache Hosted Cache Mode Deployment
- Additional Resources

Notes

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode in Word format in the Microsoft Download Center.
- The Windows Server 2012 Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode in HTML format in the Windows Server 2012 Technical Library.

Prerequisites for using this guide

This is a companion guide to the Windows Server 2012 Core Network Guide. To deploy BranchCache in hosted cache mode with this guide, you must first do the following.

- Deploy a core network in your main office by using the Core Network Guide, or already have the technologies provided in the Core Network Guide installed and functioning correctly on your network. These technologies include TCP/IP v4, DHCP, Active Directory Domain Services (AD DS), DNS, NPS, and Web Server (IIS).

  Notes

  The Windows Server 2012 Core Network Guide is available in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh911995.aspx).

  The Core Network Guide is also available in Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/p/?LinkId=255199).
- Deploy BranchCache content servers that are running Windows Server 2012 in your main office or in a cloud data center. For information on how to deploy BranchCache content servers, see Additional Resources.
- Establish wide area network (WAN) connections between your branch office, your main office and, if appropriate, your Cloud resources, by using a virtual private network (VPN), DirectAccess, or other connection method.
- Deploy client computers in your branch office that are running Windows 8.

About this guide

This guide is designed for network and system administrators who have followed the instructions in the Windows Server 2012 Core Network Guide to deploy a core network, or for those who have previously deployed the technologies included in the Core Network Guide, including Active Directory Domain Services (AD DS), Domain Name Service (DNS), Dynamic Host Configuration Protocol (DHCP), TCP/IP, Web Server (IIS), and Network Policy Server (NPS).

It is recommended that you review the design and deployment guides for each of the technologies that are used in this deployment scenario. These guides can help you determine whether this deployment scenario provides the services and configuration that you need for your organization's network.

What this guide does not provide

This guide does not provide conceptual information about BranchCache, including information about BranchCache modes and capabilities.

This guide does not provide information about how to deploy WAN connections or other technologies in your branch office, such as DHCP, a RODC, or a VPN server.

In addition, this guide does not provide guidance on the hardware you should use when you deploy a hosted cache server. It is possible to run other services and applications on your hosted cache server, however you must make the determination, based on workload, hardware capabilities, and branch office size, whether to install BranchCache hosted cache server on a particular computer, and how much disk space to allocate for the cache.

This guide does not provide instructions for configuring computers that are running Windows 7. If you have client computers that are running Windows® 7 in your branch offices, you must configure them using procedures that are different than those provided in this guide for client computers that are running Windows® 8.

In addition, if you have computers running Windows 7, you must configure your hosted cache server with a server certificate that is issued by a certification authority that client computers trust. (If all of your client computers are running Windows 8, you do not need to configure the hosted cache server with a server certificate.)

Important

If your hosted cache servers are running Windows Server 2008 R2, use the Windows Server 2008 R2 BranchCache Deployment Guide rather than this guide to deploy BranchCache in hosted cache mode. Apply the Group Policy settings that are described in that guide to both Windows 8 and Windows 7 BranchCache clients. Computers that are running Windows Server 2008 R2 cannot be configured by using the steps in this guide. For more information, see the Windows Server® 2008 R2 BranchCache Deployment Guide

Technology overviews

For this companion guide, BranchCache is the only technology that you need to install and configure. You must run Windows PowerShell BranchCache commands on your content servers, such as Web and file servers, however you do not need to change or reconfigure the content servers in any other way. In addition, you must configure client computers by using Group Policy.

BranchCache

BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server 2012 and Windows 8 operating systems, as well as in some editions of Windows Server® 2008 R2 and Windows® 7. To optimize WAN bandwidth when users access content on remote servers, BranchCache copies content from your main office or hosted cloud content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.

When you deploy BranchCache in hosted cache mode, you must configure client computers in the branch office as hosted cache mode clients, and then you must deploy a hosted cache server in the branch office. This guide demonstrates how to deploy your hosted cache server with prehashed and preloaded content from your Web and file server-based content servers.

Group Policy

Group Policy in Windows Server 2012 is an infrastructure used to deliver and apply one or more desired configurations or policy settings to a set of targeted users and computers within an Active Directory environment. This infrastructure consists of a Group Policy engine and multiple client-side extensions (CSEs) responsible for reading policy settings on target client computers.

Group Policy is used in this scenario to configure domain member client computers with BranchCache hosted cache mode.

BranchCache Hosted Cache Mode Deployment Overview

This topic contains the following sections.

- Hosted Cache Server deployment components
- Hosted Cache Server deployment process overview

Hosted Cache Server deployment components

You can use this guide to deploy a BranchCache hosted cache server in a branch office where computers are joined to a domain. In this deployment, the hosted cache server is deployed by using service connection points in Active Directory Domain Services (AD DS), and you have the

option with BranchCache in Windows Server 2012 to prehash the shared content on Web and file based content servers, then preload the content on the hosted cache server.

The following illustration shows the components that are required to deploy a BranchCache hosted cache server.

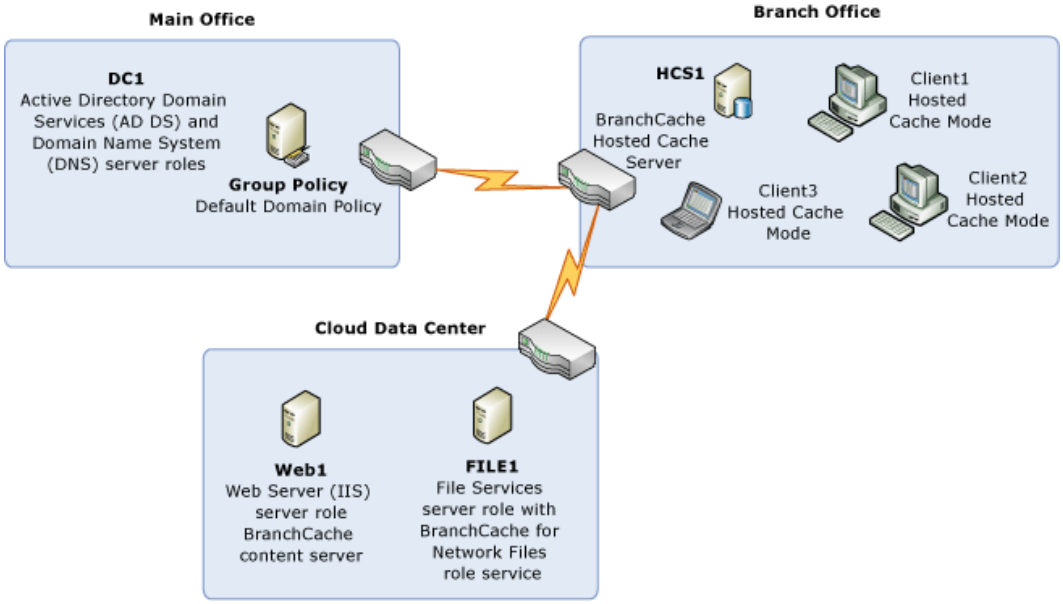

Important

Although this deployment depicts content servers in a cloud data center, you can use this guide to deploy a BranchCache hosted cache server regardless of where you deploy your content servers – in your main office or in a cloud location.

HCS1 in the branch office

You must configure this computer as a hosted cache server. If you choose to prehash content server data so that you can preload the content on your hosted cache servers, you can import data packages that contain the content from your Web and file servers.

WEB1 in the cloud data center

WEB1 is a BranchCache-enabled content server. If you choose to prehash content server data so that you can preload the content on your hosted cache servers, you can prehash the shared content on WEB1, then create a data package that you copy to HCS1.

FILE1 in the cloud data center

FILE1 is a BranchCache-enabled content server. If you choose to prehash content server data so that you can preload the content on your hosted cache servers, you can prehash the shared content on FILE1, then create a data package that you copy to HCS1.

DC1 in the main office

DC1 is a domain controller, and you must configure the Default Domain Policy with BranchCache Group Policy settings to enable Automatic Hosted Cache Discovery by Service Connection Point. When client computers in the branch have Group Policy refreshed and this policy setting is applied, they automatically locate and begin to use the hosted cache server in the branch office.

Client computers in the branch office

You must refresh Group Policy on client computers to apply new BranchCache Group Policy settings and to allow clients to locate and use the hosted cache server.

Hosted Cache Server deployment process overview

Note
The details of how to perform these steps are provided in the section BranchCache Hosted Cache Mode Deployment.

The process of deploying a BranchCache Hosted Cache Server occurs in these stages:

Note
Some of the steps below are optional, such as those steps that demonstrate how to prehash and preload content on hosted cache servers. When you deploy BranchCache in hosted cache mode, you are not required to prehash content on your Web and file content servers, to create a data package, and to import the data package in order to preload your hosted cache servers with content. The steps are noted as optional in this section and in the section BranchCache Hosted Cache Mode Deployment so that you can skip them if you prefer.

1. On HCS1, use Windows PowerShell commands to configure the computer as a hosted cache server and to register a Service Connection Point in Active Directory.
2. (Optional) On HCS1, if the BranchCache default values do not match your deployment goals for the server and the hosted cache, configure the amount of disk space that you want to allocate for the hosted cache. Also configure the disk location that you prefer for the hosted cache.
3. (Optional) Prehash content on content servers, create data packages, and preload content on the hosted cache server.

   Note
   Prehashing and preloading content on your hosted cache server is optional, however if you choose to do so, you must perform all of the steps that are applicable to your deployment. (For example, if you do not have Web servers, you do not need to perform any of the steps related to prehashing and preloading Web server content.)

   a. On WEB1, prehash Web server content and create a data package.

b. On FILE1, prehash file server content and create a data package.
c. From WEB1 and FILE1, copy the data packages to the hosted cache server HCS1.
d. On HCS1, import the data packages to preload the data cache.

4. On DC1, configure domain joined branch office client computers for hosted cache mode by configuring Group Policy with BranchCache policy settings.
5. On client computers, refresh Group Policy.

BranchCache Hosted Cache Mode Deployment Planning

Before you deploy your hosted cache server, you must plan the following items:

- Plan basic server configuration
- Plan domain access
- Plan the location and size of the hosted cache
- Plan the share to which the content server packages are to be copied
- Plan prehashing and data package creation on content servers

Plan basic server configuration

Note
If you are planning on using an existing server in your branch office as your hosted cache server, you do not need to perform this planning step, because the computer is already named and has an IP address configuration.

After you install Windows Server 2012 on your hosted cache server, you must rename the computer and assign and configure a static IP address for the local computer.

Note
In this guide, the hosted cache server is named HCS1, however you should use a server name that is appropriate for your deployment.

For more information, see the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library.

Plan domain access

Note
If you are planning on using an existing server in your branch office as your hosted cache server, you do not need to perform this planning step, unless the computer is not currently joined to the domain.

To log on to the domain, the computer must be a domain member computer and the user account must be created in AD DS before the logon attempt. In addition, you must join the computer to the domain with an account that has the appropriate group membership.

For more information, see the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library.

Plan the location and size of the hosted cache

On HCS1, determine where on your hosted cache server you want to locate the hosted cache. For example, decide the hard disk and volume where you plan to store the cache.

In addition, decide what percentage of disk space you want to allocate for the hosted cache.

Plan the share to which the content server packages are to be copied

After you create data packages on your content servers, you must copy them over the network to a share on your hosted cache server. Plan the folder location and sharing permissions for the shared folder. In addition, if your content servers host a large amount of data and the packages that you create will be large files, plan to perform the copy operation during off –peak hours so that WAN bandwidth is not consumed by the copy operation during a time when others need to use the bandwidth for business purposes.

Plan prehashing and data package creation on content servers

Before you prehash content on your content servers, you must identify the folders and files that contain content that you want to add to the data package. In addition, you must plan on the local folder location where you can store the data packages before copying them to the hosted cache server.

BranchCache Hosted Cache Mode Deployment

Follow these steps to deploy BranchCache hosted cache mode.

- Install the BranchCache Feature and Configure the Hosted Cache Server by Service Connection Point
- Move and Resize the Hosted Cache (Optional)
- Prehash and Preload Content on the Hosted Cache Server (Optional)
- Configure Client Automatic Hosted Cache Discovery by Service Connection Point

Note

The procedures in this guide do not include instructions for cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Install the BranchCache Feature and Configure the Hosted Cache Server by Service Connection Point

You can use this procedure to install the BranchCache feature on your hosted cache server, HCS1, and to configure the server to register a Service Connection Point (SCP) in Active Directory Domain Services (AD DS).

When you register hosted cache servers with an SCP in AD DS, the SCP allows client computers that are configured correctly to automatically discover hosted cache servers by querying AD DS for the SCP. Instructions on how to configure client computers to perform this action are provided later in this guide.

Important

Before you perform this procedure, you must join the computer to the domain and configure the computer with a static IP address.

To perform this procedure, you must be a member of the Administrators group.

To install the BranchCache feature and configure the hosted cache server

1. On the server computer desktop, in the Taskbar, right-click the Windows PowerShell icon, right-click the words **Windows PowerShell**, and then click **Run as Administrator**.

   Note

   If you have previously run Windows PowerShell as an Administrator on this computer, you do not need to right-click the words **Windows PowerShell**, because **Run as Administrator** is present as a right-click menu option.

2. Windows PowerShell opens. Type the following command, and then press ENTER.

   ```
   Install-WindowsFeature BranchCache
   ```

3. To configure the computer as a hosted cache server after the BranchCache feature is installed, and to register a Service Connection Point in AD DS, type the following command in Windows PowerShell, and then press ENTER.

   ```
   Enable-BCHostedServer -RegisterSCP
   ```

4. To verify the hosted cache server configuration, type the following command and press ENTER.

```
Get-BCStatus
```

The results of the command display status for all aspects of your BranchCache installation. Following are a few of the BranchCache settings and the correct value for each item:

- BranchCacheIsEnabled: True
- HostedCacheServerIsEnabled: True
- HostedCacheScpRegistrationEnabled: True

5. To prepare for the step of copying your data packages from your content servers to your hosted cache servers, either identify an existing share on the hosted cache server or create a new folder and share the folder so that it is accessible from your content servers. After you create your data packages on your content servers, you will copy the data packages to this shared folder on the hosted cache server.
6. If you are deploying more than one hosted cache server, repeat this procedure on each server.

Move and Resize the Hosted Cache (Optional)

You can use this procedure to move the hosted cache to the drive and folder that you prefer, and to specify the amount of disk space that the hosted cache server can use for the hosted cache.

This procedure is optional. If the default cache location (%windir%\ServiceProfiles\NetworkService\AppData\Local\PeerDistPub) and size – which is 5% of the total hard disk space – are appropriate for your deployment, you do not need to change them.

You must be a member of the Administrators group to perform this procedure.

To move and resize the hosted cache

1. On the server computer desktop, in the Taskbar, right-click the Windows PowerShell icon, right-click the words **Windows PowerShell**, and then click **Run as Administrator**.

   Note
   If you have previously run Windows PowerShell as an Administrator on this computer, you do not need to right-click the words **Windows PowerShell**, because **Run as Administrator** is present as a right-click menu option.
2. Windows PowerShell opens. Type the following command to move the hosted cache to another location on the local computer, and then press ENTER.

   Important
   Before running the following command, replace parameter values, such as –Path

and –MoveTo, with values that are appropriate for your deployment.

```
Set-BCCache -Path C:\datacache –MoveTo D:\datacache
```

3. Type the following command to resize the hosted cache –specifically the datacache - on the local computer. Press ENTER.

   Important
   Before running the following command, replace parameter values, such as -Percentage, with values that are appropriate for your deployment.

   ```
   Set-BCCache -Percentage 20
   ```

4. To verify the hosted cache server configuration, type the following command and press ENTER.

   ```
   Get-BCStatus
   ```

   The results of the command display status for all aspects of your BranchCache installation. Following are a few of the BranchCache settings and the correct value for each item:

   - DataCache | CacheFileDirectoryPath: Displays the hard disk location that matches the value you provided with the –MoveTo parameter of the SetBCCache command. For example, if you provided the value D:\datacache, that value is displayed in the command output.
   - DataCache | MaxCacheSizeAsPercentageOfDiskVolume: Displays the number that matches the value you provided with the –Percentage parameter of the SetBCCache command. For example, if you provided the value 20, that value is displayed in the command output.

Prehash and Preload Content on the Hosted Cache Server (Optional)

You can use the procedures in this section to prehash content on your content servers, add the content to data packages, and then preload the content on your hosted cache servers. These procedures are optional because you do not have to prehash and preload content on your hosted cache servers. If you do not preload content, data is added to the hosted cache automatically as clients download it over the WAN connection.

Important
Although these procedures are collectively optional, if you decide to prehash and preload content on your hosted cache servers, performing both procedures is required.

- Create Content Server Data Packages for Web and File Content (Optional)
- Import Data Packages on the Hosted Cache Server (Optional)

Create Content Server Data Packages for Web and File Content (Optional)

You can use this procedure to prehash content on Web and file servers, and then create data packages to import on your hosted cache server. This procedure is optional because you do not have to prehash and preload content on your hosted cache servers. If you do not pre-load content, data is added to the hosted cache automatically as clients download it over the WAN connection.

This procedure provides instructions for prehashing content on both file servers and Web servers. If you do not have one of those types of content servers, you do not have to perform the instructions for that content server type.

Important
Before you perform this procedure, you must install and configure BranchCache on your content servers. In addition, if you plan on changing the server secret on a content server, do so before pre-hashing content – modifying the server secret invalidates previously-generated hashes. For more information about configuring content servers, see Additional Resources.

To perform this procedure, you must be a member of the Administrators group.

To create content server data packages

1. On each content server, locate the folders and files that you want to prehash and add to a data package. Identify or create a folder where you want to save your data package later in this procedure.
2. On the server computer desktop, in the Taskbar, right-click the Windows PowerShell icon, right-click the words **Windows PowerShell**, and then click **Run as Administrator**.

   Note
   If you have previously run Windows PowerShell as an Administrator on this computer, you do not need to right-click the words **Windows PowerShell**, because **Run as Administrator** is present as a right-click menu option.
3. Windows PowerShell opens. Do one or both of the following, depending on the types of content servers that you have:

   Note
   The value for the –Path parameter is the folder where your content is located. You must replace the example values in the commands below with a valid folder location on your content server that contains data that you want to prehash and add to a package.

   - If the content that you want to prehash is on a file server, type the following command, and then press ENTER.

     ```
     Publish-BCFileContent -Path D:\share -StageData
     ```

- If the content that you want to prehash is on a Web server, type the following command, and then press ENTER.

  ```
  Publish-BCWebContent –Path D:\inetpub\wwwroot -StageData
  ```

4. Create the data package by running the following command on each of your content servers. Replace the example value (D:\temp) for the –Destination parameter with the location that you identified or created at the beginning of this procedure.

   ```
   Export-BCDataPackage –Destination D:\temp
   ```

5. From the content server, access the share on your hosted cache servers where you want to preload content, and copy the data packages to the shares on the hosted cache servers.

Import Data Packages on the Hosted Cache Server (Optional)

You can use this procedure to import data packages and preload content on your hosted cache servers. This procedure is optional because you do not have to prehash and preload content on your hosted cache servers. If you do not pre-load content, data is added to the hosted cache automatically as clients download it over the WAN connection.

You must be a member of the Administrators group to perform this procedure.

To import data packages on the hosted cache server

1. On the server computer desktop, in the Taskbar, right-click the Windows PowerShell icon, right-click the words **Windows PowerShell**, and then click **Run as Administrator**.

   Note
   If you have previously run Windows PowerShell as an Administrator on this computer, you do not need to right-click the words **Windows PowerShell**, because **Run as Administrator** is present as a right-click menu option.

2. Windows PowerShell opens. Type the following command, replacing the value for the –Path parameter with the folder location where you have stored your data packages, and then press ENTER.

   ```
   Import-BCCachePackage –Path D:\temp\PeerDistPackage.zip
   ```

3. If you have more than one hosted cache server where you want to preload content, perform this procedure on each hosted cache server.

Configure Client Automatic Hosted Cache Discovery by Service Connection Point

With this procedure you can use Group Policy to enable and configure BranchCache hosted cache mode on domain-joined computers that are running Windows Server 2012 or Windows 8.

Note

To configure domain-joined computers that are running Windows Server 2008 R2 or Windows 7, see the Windows Server 2008 R2 BranchCache Deployment Guide.

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

To use Group Policy to configure clients for hosted cache mode

1. On a computer upon which the Active Directory Domain Services server role is installed, open Server Manager, select the Local Server, click **Tools**, and then click **Group Policy Management**. The Group Policy Management console opens.
2. In the Group Policy Management console, expand the following path: **Forest:** *corp.contoso.com*, **Domains**, *corp.contoso.com*, **Group Policy Objects**, where *corp.contoso.com* is the name of the domain where the BranchCache client computer accounts that you want to configure are located.
3. Right-click **Group Policy Objects**, and then click **New**. The **New GPO** dialog box opens. In **Name**, type a name for the new Group Policy object (GPO). For example, if you want to name the object BranchCache Client Computers, type **BranchCache Client Computers**. Click **OK**.
4. In the Group Policy Management console, ensure that **Group Policy Objects** is selected, and in the details pane right-click the GPO that you just created. For example, if you named your GPO BranchCache Client Computers, right-click **BranchCache Client Computers**. Click **Edit**. The Group Policy Management Editor console opens.
5. In the Group Policy Management Editor console, expand the following path: **Computer Configuration**, **Policies**, **Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer**, **Network**, **BranchCache**.
6. Click **BranchCache**, and then in the details pane, double-click **Turn on BranchCache**. The **Turn on BranchCache** dialog box opens.
7. In the **Turn on BranchCache** dialog box, click **Enabled**, and then click **OK**.
8. In the Group Policy Management Editor console, ensure that **BranchCache** is still selected, and then in the details pane double-click **Enable Automatic Hosted Cache Discovery by Service Connection Point**. The policy setting dialog box opens.
9. In the **Enable Automatic Hosted Cache Discovery by Service Connection Point** dialog box, click **Enabled**, and then click **OK**.
10. To enable client computers to download and cache content from BranchCache file server-based content servers: In the Group Policy Management Editor console, ensure that **BranchCache** is still selected, and then in the details pane double-click

BranchCache for network files. The **Configure BranchCache for network files** dialog box opens.

11. In the **Configure BranchCache for network files** dialog box, click **Enabled**. In **Options**, type a numeric value, in milliseconds, for the maximum round trip network latency time, and then click **OK**.

    Note
    By default, client computers cache content from file servers if the round trip network latency is longer than 80 milliseconds.

12. To configure the amount of hard disk space allocated on each client computer for the BranchCache cache: In the Group Policy Management Editor console, ensure that **BranchCache** is still selected, and then in the details pane double-click **Set percentage of disk space used for client computer cache**. The **Set percentage of disk space used for client computer cache** dialog box opens. Click **Enabled**, and then in **Options** type a numeric value that represents the percentage of hard disk space used on each client computer for the BranchCache cache. Click **OK**.
13. To specify the default age, in days, for which segments are valid in the BranchCache data cache on client computers: In the Group Policy Management Editor console, ensure that **BranchCache** is still selected, and then in the details pane double-click **Set age for segments in the data cache**. The **Set age for segments in the data cache** dialog box opens. Click **Enabled**, and then in the details pane type the number of days that you prefer. Click **OK**.
14. Configure additional BranchCache policy settings for client computers as appropriate for your deployment.
15. Refresh Group Policy on branch office client computers by running the command **gpupdate /force**, or by rebooting the client computers.

Additional Resources

For more information about the technologies that are discussed in this guide, see the following resources:

- BranchCache
- Install and Configure Content Servers
- BranchCache Cmdlets in Windows PowerShell
- Group Policy
- Building Your Cloud Infrastructure: Scenario Overview
- Windows Server 2008 R2 BranchCache Deployment Guide

Core Network Companion Guide: Group Policy Deployment

This is a companion guide to the Windows Server® 2012 Core Network Guide, which is available for download in Microsoft Office Word format at the Microsoft Download Center (http://go.microsoft.com/fwlink/?LinkId=255199) and in HTML format in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh911995.aspx).

The Windows Server 2012 Core Network Guide provides instructions for planning and deploying the core components required for a fully functioning network and a new Active Directory domain in a new forest.

This guide explains how to build on the core network by providing instructions for deploying Group Policy Objects (GPOs) by using membership groups instead of the organizational units (OUs) that form the hierarchy of an Active Directory domain.

This guide contains the following sections.

- About this guide
- Requirements
- What this guide does not provide
- Technology overview of Group Policy
- Group Policy Deployment Overview
- Group Policy Deployment Planning
- Group Policy Deployment
- Additional Resources

Notes

This guide is available at the following locations.

- The Windows Server 2012 Core Network Companion Guide: Group Policy Deployment in Word format in the Microsoft Download Center.
- The Windows Server 2012 Core Network Companion Guide: Group Policy Deployment in HTML format in the Windows Server 2012 Technical Library.

About this guide

This guide provides instructions for deploying Group Policy settings to a set of client computers or users by using membership groups rather than account location in the OU hierarchy of a domain.

The method described in this guide shows you how to create a single membership group into which you can add the user or computer accounts that are to receive a configuration through the use of GPOs. Membership in the group, rather than the account location in the OU hierarchy, determines whether the computer receives one of the GPOs associated with the membership group. In addition, Windows Management Instrumentation (WMI) filters are used to ensure that

only the GPO with the settings that correspond to the version of Windows running on the computer is applied.

There are two main benefits to using this method to deploy GPOs:

- It is totally independent of the OU structure of your domain. To apply a GPO to one computer no longer means moving computers to another OU or restructuring your OU hierarchy.
- It is very easy to apply or stop applying the settings in a GPO. You simply remove the user or computer account from the membership group. This removes the user or computer from the scope of the GPO without affecting any other GPOs that apply to the user or computer.

This guide is designed for network and system administrators who have followed the instructions in the Windows Server 2012 Core Network Guide to deploy a core network, or for those who have previously deployed the core technologies included in the core network, including Active Directory Domain Services (AD DS), Domain Name Service (DNS), Dynamic Host Configuration Protocol (DHCP), TCP/IP, and Windows Internet Name Service (WINS) (optional).

It is recommended that you review the design and deployment guides for each of the technologies that are used in this deployment scenario. These guides can help you determine whether this deployment scenario provides the services and configuration that you need for your organization's network.

Requirements

Caution

We recommend that you use the methods documented in this guide only for GPOs that must be deployed to the majority of the computers in your organization, and only when the OU hierarchy in your Active Directory domain does not match up well with the deployment needs of these GPOs. When the OU hierarchy supports it, deploy a GPO by linking it to the lowest level OU that contains all of the accounts to which the GPO will apply.

In a large enterprise environment with hundreds or thousands of GPOs, the use of this method can result in user or computer accounts that are made members of an excessive number of groups; this can result in network connectivity problems when network protocol limits have been exceeded. For more information about the problems associated with excessive group membership, see the following articles in the Microsoft Knowledge Base:

- What's New for Kerberos in Windows Server 2012 provides information about improvements to reduce authentication failures due to large service tickets, see http://technet.microsoft.com/library/hh831747.aspx
- Article 327825 "New resolution for problems with Kerberos authentication when users belong to many groups" (http://go.microsoft.com/fwlink/?LinkId=23044)
- Article 263693 "Group Policy may not be applied to users belonging to many groups" (http://go.microsoft.com/fwlink/?LinkId=126293)
- Article 328889 "Users who are members of more than 1,015 groups may fail logon authentication" (http://go.microsoft.com/fwlink/?LinkId=115213)

Following are the requirements for using Group Policy:

- To deploy Group Policy, you must have an Active Directory domain controller hosting a domain and computers that are joined to the domain.
- To configure GPOs, create membership groups, and assign members to them, you must be logged on as a member of the Domain Admins group.

What this guide does not provide

This guide does not provide comprehensive instructions for designing and deploying a Group Policy infrastructure by using AD DS. It is recommended that you review AD DS and Group Policy documentation before you deploy the technologies in this guide. For more information, see Additional Resources.

Technology overview of Group Policy

Group Policy is a management technology that provides users with consistent access to their applications, application settings, roaming user profiles, and user data, from any managed computer—even when they are disconnected from the network. Group Policy settings are contained in GPOs, which are linked to sites, domains, or OUs within an Active Directory domain. The settings within GPOs are then evaluated and applied by the targeted computers and users. Group Policy is one of the top reasons to deploy Active Directory Domain Services (AD DS) because it allows you to manage user and computer objects.

Most administrators associate GPO deployment with the OU hierarchy of an Active Directory domain. You can link a GPO to an OU, and any computers or users in that OU or one of its descendants receive and apply the policy. However, an Active Directory domain can contain only a single hierarchy of OUs, and computer and user accounts can be placed into only a single OU. For this reason, there are times when an OU hierarchy that is appropriate for solving one problem is inappropriate for another. For example, many organizations design the OU hierarchy to support delegated administration. Computer and user accounts are placed into OUs for which an IT team has been assigned responsibility. By granting the IT team administrative permissions on the OU container, they can manage the computers and users whose accounts are in the OU. This same hierarchy might be ineffective for deploying Group Policy settings that affect computers across the entire organization, for example, when deploying Internet Protocol security (IPsec) settings for server and domain isolation scenarios.

In addition, configuring one version of Windows might require you to use a policy setting that is different from the policy setting used in another version of Windows. For example, IPsec rules in Windows Server 2012, Windows Server 2008, Windows 8, Windows 7, and Windows Vista are managed by a different part of the GPO than the IPsec rules for Windows Server 2003 and earlier versions of Windows. This means that you might have six separate GPOs that all perform the same function, one each for Windows Server 2012, Windows Server 2008, Windows Server 2003, Windows 8, Windows 7, and Windows Vista. The Windows Server 2003 GPO also applies to earlier versions of Windows.

Group Policy Deployment Overview

The following illustration shows the components that are required to deploy Group Policy by using a membership group.

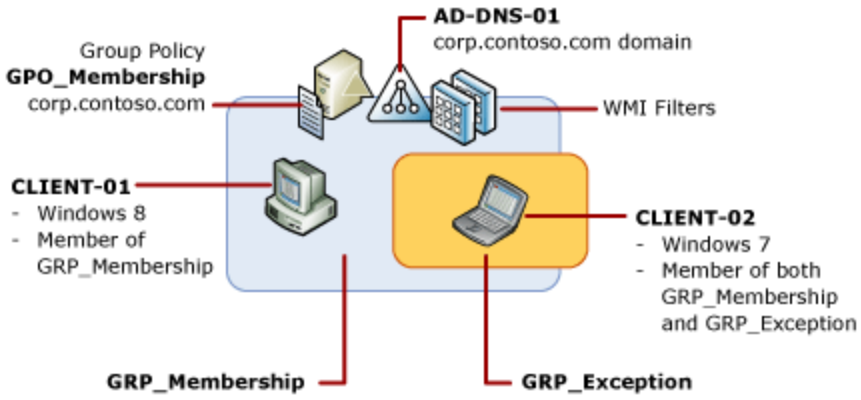

Group Policy deployment components

The components include:

Domain controller: AD-DNS-01

The domain controller is a computer running Windows Server 2012 and Active Directory Domain Services (AD DS). The illustration shows a domain controller named AD-DNS-01 configured in the corp.contoso.com forest and domain.

Group Policy: GPO_Membership

Group Policy is configured in AD DS on the domain controller. You create one or more GPOs that are associated with the membership group and configure the settings required by the set of computers that must receive each GPO. For example, you can create a GPO named GPO_Membership, as shown in the illustration.

WMI Filters

Windows Management Instrumentation (WMI) filters allow Group Policy to query the computer for conditions that must be true for the GPO to apply. In this guide, WMI filters are used to query for the version of Windows to ensure that only a GPO designed for that version will apply.

Membership group: GRP_Membership

The membership group contains the user or computer accounts that will receive one of the GPOs associated with the group. The choice of GPO depends on the WMI filters used for checking the

version of Windows and any membership in an exception group. For example, you can create a group named GRP_Membership, as shown in the above illustration.

Exception group: GRP_Exception

The exception group contains the user or computer accounts that might be in the membership group, but must not be allowed to apply a particular membership GPO. Exception groups are assigned deny permissions on the membership GPOs. For example, you can create a group named GRP_Exception, as shown in the above illustration.

Exception groups are used for the following reasons:

- When a computer or user is in two membership groups but must apply only one of the two GPOs. You use one of the membership groups as if it were an exception group for the GPOs that pertain to the other membership group. For example, consider membership groups Group A and Group B, each with a GPO, GPO A and GPO B. Only one of the GPOs must apply to any computer. Members of Group B must receive only the GPO B, even if they are also members of Group A. To do this, treat the Group B membership group as an exception group for Group A. On GPO A, deny Apply Group Policy permissions to Group B. Because deny permissions override allow permissions, members of Group B will only be allowed to apply GPO B.
- When a WMI filter cannot be used.

Group Policy deployment process

The process for configuring Group Policy based on a membership group occurs in these stages:

1. Create the membership group in AD DS.
2. Create the exception group in AD DS.
3. Create the GPOs, one for each version of Windows that requires different settings to achieve the desired configuration. Remember that often you can copy a GPO for one version of Windows to serve as a starting point for another version. This can save a lot of time if your GPOs contain a large number of settings. For example, Windows 8 and Windows 7 support almost the same collection of settings. You could create the GPO for Windows 7, configure it, and then make a copy of it for Windows 8. Then you only need to change the few settings that are different between Windows 7 and Windows 8.
4. Create the WMI filters that allow you to distinguish between different versions of Windows. In theory, you could have six GPOs for a single membership group, one each for Windows Server 2012, Windows Server 2008, Windows Server 2003, Windows 8, Windows 7, Windows Vista, and Windows XP. Each GPO must have a WMI filter that prevents it from being applied to any version of Windows except the one for which it is designed.
5. Grant Read and Apply Group Policy permissions on the GPOs to the membership group. Remove the default entry for Authenticated Users. Deny Apply Group Policy permissions to the exception group. Assign the corresponding WMI filter to each GPO.

Important

If you have membership groups in which a computer might be a member of more than one membership group, and only one of the GPOs must be applied, then treat one of the membership groups as an exception group for the other.

6. Link the GPOs to the domain container to make it visible to all of the computers in the domain. The security group and WMI filters limit their application to only the specified set of computers.
7. Add a small number of test computers to the membership and exception groups. Refresh Group Policy on the test computers and make sure that each receives the correct GPO and applies the correct settings.
8. After testing is complete, add the production computers to the membership and exception groups.

Group Policy Deployment Planning

Before you deploy Group Policy, you must plan the membership and exception groups.

Planning the membership and exception groups

There is one membership group for each set of GPOs that contain configuration data for your client computers and users. Adding a user or computer account to the group enables that user or computer to read and apply all of the GPOs associated with the group.

To limit the user or computer to only one GPO of the several that might be associated with the membership group, create and assign WMI filters to each GPO. A WMI filter is evaluated to determine if a GPO should be applied to the user or computer. For example, the WMI filters described in this guide contain information about the version of the Windows operating system. For more information about creating WMI filters, see WMI Filtering Using GPMC (http://go.microsoft.com/fwlink/?linkid=93188).

If there are some computers in the membership group that should not apply the GPO, then you can create an exception group that is denied permission to apply the GPO. Because deny permissions override allow permissions, a user or computer that is a member of both groups will not apply the GPO.

Planning domain access

To log on to the domain, the computer must be a domain member computer and the user account must be created in AD DS before the logon attempt.

For more information, see the "Joining Computers to the Domain and Logging On" topic in the Core Network Guide.

Group Policy Deployment

Follow these steps to create Group Policy by using a membership group.

1. Create the Membership Group
2. Create the Exception Group
3. Create the WMI Filters
4. Create the Group Policy Objects for Each Version of Windows
5. Apply Security Group Filters and WMI Filters to the GPOs
6. Link the GPOs to the Domain
7. Add Computers to the Exception Group
8. Add Test Computers to the Membership Group
9. Add Production Computers to the Membership Group

Note

The procedures in this guide do not explicitly mention the display of the **User Account Control** dialog box. If this dialog box appears in response to your actions while you are performing the procedures in this guide, click **Continue**.

Create the Membership Group

You can use this procedure to create a new group in Active Directory Users and Computers Microsoft Management Console (MMC).

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To create a membership group

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click **corp.contoso.com**.
2. In the details pane, right-click the folder in which you want to add a new group.

   Where?

   - Active Directory Users and Computers/*domain node*/*folder*
3. Point to **New**, and then click **Group**.
4. In **New Object – Group**, in **Group name**, type the name of the new group.

   By default, the name you type is also entered as the pre-Windows 2000 name of the new group.
5. In **Group scope**, select **Global**, unless you have a specific reason for chosing another option.

6. In **Group type**, select **Security**.
7. Click **OK**.

Create the Exception Group

Use this procedure to create an exception group. Members of this group are prevented from applying a GPO.

Membership in **Domain Admins**, or equivalent, is the minimum required to perform this procedure.

To create an exception group

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click **corp.contoso.com**.
2. In the details pane, right-click the folder in which you want to add a new group.

   Where?

   - Active Directory Users and Computers/*domain node*/*folder*
3. Point to **New**, and then click **Group**.
4. In **New Object – Group**, in **Group name**, type the name of the new group.

   By default, the name you type is also entered as the pre-Windows 2000 name of the new group.
5. In **Group scope**, select **Global**, unless you have a specific reason for chosing another option.
6. In **Group type**, select **Security**.
7. Click **OK**.

Create the WMI Filters

Use this procedure to create the WMI filters that are used to restrict the application of a GPO to the computers in the membership group that are running the version of Windows for which the GPO is intended.

Membership in **Domain Admins** is the minimum required to complete this procedure.

Create WMI Filters

To create the WMI filters

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Group Policy Management**. The Group Policy Management Console (GPMC) opens.
2. In the GPMC navigation pane, expand **Features**, expand **Group Policy Management**, expand **Forest:** *your forest name*, expand **Domains**, expand your domain, and then click **WMI Filters**.
3. Click **Action**, and then click **New**.
4. In the **New WMI Filter** dialog box, in **Name**, type the name of the filter. Use a name that indicates the target group (for example, **Windows Server 2012**).
5. In **Description**, type the purpose of the filter.
6. Click **Add**.

   The **WMI Query** dialog box appears.
7. In **Query**, type the WMI Query Language (WQL) string that returns **TRUE** when applied to the correct version of Windows. For example, for Windows Server 2012 type:

   ```
   select * from Win32_OperatingSystem where Version like "6.2%"
   and ProductType = "3"
   ```

   Both Windows Server 2012 and Windows 8 return version numbers that begin with `6.2`. To differentiate between the client and server versions, include the clause to check the **ProductType** field. This value returns **1** for client versions of Windows such as Windows 8, **2** for server versions of Windows operating as domain controllers, and **3** for server versions of Windows that are not operating as domain controllers.
8. Repeat step 8 for each of the other versions of Windows. The following table shows sample query statements for each version.

WMI Filter Name	WQL Query Statement
Windows Server 2012	`select * from Win32_OperatingSystem where Version like "6.2%" and ProductType = "3"`
Windows 8	`select * from Win32_OperatingSystem where Version like "6.2%" and ProductType = "1"`
Windows Server 2008 R2	`select * from Win32_OperatingSystem where Version like "6.1%" and ProductType = "3"`
Windows 7	`select * from Win32_OperatingSystem where Version like "6.1%" and ProductType = "1"`

Windows Server 2008	`select * from Win32_OperatingSystem where Version like "6.0%" and ProductType = "3"`
Windows Vista	`select * from Win32_OperatingSystem where Version like "6.0%" and ProductType = "1"`
Windows Server 2003	`select * from Win32_OperatingSystem where Version like "5.2%" and ProductType = "3"`
Windows XP	`select * from Win32_OperatingSystem where (Version like "5.1%" or Version like "5.2%") and ProductType = "1"`

Note

The filter for Windows 2000 is used to prevent computers that are running later versions of Windows from applying the GPO. You cannot use a WMI filter to apply a GPO to computers that are running Windows 2000 because that version of the operating system does not support WMI filters.

You can also create combination filters when required by your design. The following table shows query statements for common operating system combinations.

WMI Filter Name	WQL Query Statement
Windows Vista and Windows Server 2008	`select * from Win32_OperatingSystem where Version like "6.0%" and ProductType<>"2"`
Windows Server 2003 and Windows Server 2008	`select * from Win32_OperatingSystem where (Version like "5.2%" or Version like "6.0%") and ProductType="3"`
Windows 2000, XP and 2003	`select * from Win32_OperatingSystem where Version like "5.%" and ProductType<>"2"`

Create the Group Policy Objects for Each Version of Windows

Use this procedure to create the GPOs. Create one GPO for each version of Windows that requires different settings to achieve the desired configuration.

Membership in **Domain Admins** is the minimum required to complete this procedure.

To create the GPOs

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Group Policy Management**. The Group Policy Management Console (GPMC) opens.
2. In the GPMC navigation pane, expand **Forest:** *your forest name*, expand **Domains**, expand your domain, and then click **Group Policy Objects**.
3. Click **Action**, and then click **New**.
4. In the **New GPO** dialog box, in **Name**, type the name of the GPO. Use a name that indicates both the function and the target group (for example, **GPO_FunctionA_W2K12**). Click **OK** to save your GPO.
5. Repeat step 5 for each of the other versions of Windows supported by the membership group (**GPO_FunctionA_W2K8 R2**, **GPO_FunctionA_W2K8**, **GPO_FunctionA_W2K3**, **GPO_FunctionA_Windows_8**, **GPO_FunctionA_Windows_7**, **GPO_FunctionA_Vista**, and **GPO_FunctionA_XP**).

Apply Security Group Filters and WMI Filters to the GPOs

Use this procedure to assign the WMI and security group filters that you created earlier to restrict each GPO to the computers in the membership group that are running the version of Windows for which the GPO is intended.

Membership in **Domain Admins** is the minimum required to complete this procedure.

To assign the security group and WMI filters to your GPOs

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Group Policy Management**. The Group Policy Management Console (GPMC) opens.
2. In the GPMC navigation pane, expand **Forest:** *your forest name*, expand **Domains**, and then expand your domain.
3. For each of the GPOs that contains settings for a different zone and version of Windows, click the GPO and perform the following steps on the **Scope** tab in the results pane:
   a. Remove the default security group filter that allows any computer to apply the GPO.

In **Security Filtering**, click **Authenticated Users**, and then click **Remove**.

b. Add a security group filter that permits only accounts in the membership group to apply the GPO. In **Security Filtering**, click **Add**, enter the membership group account name that you created for the zone, and then click **OK**.

Note
We recommend that you use the GPMC to remove **Authenticated Users** from **Security Filtering** in order to automatically remove both the Apply Group Policy permission and read access permission. If the Apply Group Policy permission is not set, but the Read permission is, the GPO is still inspected (although not applied) by any user or computer that is in the OU hierarchy where the GPO is linked. This inspection process increases logon time slightly for each GPO with read access still set after removal of the Apply Group Policy permission.

c. Add a security group filter that prevents members of an exception group from applying the GPO. Click the **Delegation** tab, click **Advanced**, click **Add**, enter the exception group account name, and then click **OK**. In the **Group or user names** list, select the group you just added, in the **Permissions for** *Exception Group Name* list, clear all of the **Allow** check boxes, select the **Deny** check boxes for **Read** and **Apply Group Policy**, click **OK** and then **Yes** in **Windows Security**.

d. Assign the WMI filter that limits the GPO to only computers running the specified versions of Windows. Click the **Scope** tab, and then in the **WMI Filtering** list, select the WMI filter you created earlier. In the confirmation dialog box, click **Yes**.

Note
We recommend that you use WMI filters primarily for exception management. They can provide a powerful solution for targeting GPOs to specific users and computers, but because WMI filters are evaluated every time Group Policy is processed, they increase startup and logon time. Also, there is no time-out for WMI filters. Therefore, you should use them only when necessary.

Link the GPOs to the Domain

Use this procedure to link the GPOs that you created earlier to the domain container in the Active Directory hierarchy. Because the assignment of the GPOs is controlled by the combination of the security group and WMI filters, you do not have to assign the GPOs to OU containers.

Important
Linking GPOs to the domain container is appropriate for GPOs that must be used by computers throughout the Active Directory hierarchy. Link GPOs that apply only to a

subset of the computers in the network to the OU container that is lowest in the hierarchy but still includes all of the targeted computers in the containers below it.

Membership in **Domain Admins** is the minimum required to complete this procedure.

To link the GPOs to the domain container

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Group Policy Management**. The Group Policy Management Console (GPMC) opens.
2. In the GPMC navigation pane, expand **Forest:** *your forest name*, expand **Domains**, expand your domain, and then click **Group Policy Objects**.
3. Drag each GPO in the **Group Policy Objects** node to the node for your domain container. In the confirmation dialog box, click **OK**.

Add Computers to the Exception Group

Before you add any computers to the membership groups, use this procedure to add computers to the appropriate exception groups. This will prevent the computer accounts in the exception group from applying GPOs that are not intended for them.

Membership in **Domain Admins** is the minimum required to complete this procedure.

To add computers to the exception groups

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click *corp.contoso.com*.
2. In the navigation pane, expand **Active Directory Users and Computers**, expand your domain, and then select **Computers**. (If you created the exception groups somewhere other than in the Computers container, select the appropriate container.)
3. In the details pane, double-click the exception group to which you want to add computers.
4. Select the **Members** tab, and then click **Add**.
5. Type the name of the computer in the text box, and then click **OK**.
6. Repeat steps 5 and 6 for each additional computer account or group that you want to add.
7. Click **OK** to close the group properties dialog box.
8. Restart the client computers that were added to the exception groups. Changes in group membership do not take effect until after the computer is restarted.

Add Test Computers to the Membership Group

Use this procedure to add test computers to membership groups, force GPO updates, and confirm the results. Be sure to add computers that represent the overall composition of your network, at least one for every GPO that you deploy.

Membership in **Domain Admins** is the minimum required to complete this procedure.

To add the test computers to the membership groups

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click *corp.contoso.com*.
2. In the navigation pane, expand **Active Directory Users and Computers**, expand your domain, and then select **Computers**. (If you created the exception groups somewhere other than in the Computers container, select the appropriate container.)
3. In the details pane, double-click the GPO membership group to which you want to add computers.
4. Select the **Members** tab, and then click **Add**.
5. Type the name of the computer in the text box, and then click **OK**.
6. Repeat steps 5 and 6 for each additional computer account or group that you want to add.
7. Click **OK** to close the group properties dialog box.
8. Restart the client computers that were added to the membership groups. Changes in group membership do not take effect until after the computer is restarted.

To force a GPO update on the test computers

1. For a computer that is running Windows Server 2012, Windows Server 2008, Windows 8, Windows 7, Windows Vista, open an elevated command prompt, and then type the following command:

   ```
   gpupdate /target:computer /force
   ```

2. For a computer that is running Windows Server 2003 or Windows XP, open a command prompt, and then type the following command:

   ```
   gpupdate /target:computer /force
   ```

To verify the GPO application on the test computers

1. For a computer that is running Windows Server 2012, Windows Server 2008, Windows 8, Windows 7, Windows Vista, open an elevated command prompt, and then type the following command:

```
gpresult /r /scope:computer
```

2. For a computer that is running Windows Server 2003 or Windows XP, open a command prompt, and then type the following command:

```
gpresult /scope:computer
```

Add Production Computers to the Membership Group

After you have thoroughly tested your GPO deployment, use this procedure to add the production computers to the membership groups.

To add production computers to the membership groups

1. On AD-DNS-01, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers MMC opens. If it is not already selected, click the node for your domain. For example, if your domain is corp.contoso.com, click *corp.contoso.com*.
2. In the navigation pane, expand **Active Directory Users and Computers**, expand your domain, and then select **Computers**. (If you created the exception groups somewhere other than in the Computers container, select the appropriate container.)
3. In the details pane, double-click the GPO membership group to which you want to add computers.

   Caution
   If you choose to add a large group like Domain Computers to the membership group, make sure that your exception groups are fully populated and tested before adding the large group.
4. Select the **Members** tab, and then click **Add**.
5. Type the name of the computer in the text box, and then click **OK**.
6. Repeat steps 5 and 6 for each additional computer account or group that you want to add.
7. Click **OK** to close the group properties dialog box.
8. Restart the client computers that were added to the membership groups. Changes in group membership do not take effect until after the computer is restarted.

Additional Resources

For more information about the technologies that are discussed in this guide, see the following resources:

- Active Directory Domain Services in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh831484.aspx)
- Group Policy in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh831791.aspx)
- Core Network Guide
- "Configuring All Servers" section of the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library, see Configuring All Servers
- "Joining Computers to the Domain and Logging On" sections of the Windows Server 2012 Core Network Guide in the Windows Server 2012 Technical Library, see Joining Client Computers to the Domain and Logging On and Joining Server Computers to the Domain and Logging On

802.1X Authenticated Wired Access Overview

This document provides introductory information about Institute of Electrical and Electronics Engineers (IEEE) 802.1X authenticated access for IEEE 802.3 wired Ethernet connections. Links to resources with information about technologies that are closely related to 802.1X Authenticated Wired Access, or otherwise relevant to wired access are also provided.

Notes

In addition to this topic, the following 802.1X Authenticated Wired Access documentation is available.

- What's New in 802.1X Authenticated Wired Access in Windows Server 2012 R2

Did you mean…

- 802.1X Authenticated Wireless Access
- Netsh Commands for Wired Local Area Network (LAN) in the Windows Server 2008 R2 and Windows Server 2008 technical library on TechNet.
- 802.1X Authenticated Wired Access in the Windows Server 2008 R2 and Windows Server 2008 technical library on TechNet.

Feature description

IEEE 802.1X authentication provides an additional security barrier for your intranet that you can use to prevent guest, rogue, or unmanaged computers that cannot perform a successful authentication from connecting to your intranet.

For the same reason that administrators deploy IEEE 802.1X authentication for IEEE 802.11 wireless networks—enhanced security—network administrators want to implement the IEEE

802.1X standard to help protect their wired network connections. Just as an authenticated wireless client must submit a set of credentials to be validated before being allowed to send wireless frames to the intranet, an IEEE 802.1X wired client must also perform authentication prior to being able to send traffic over its switch port.

Important terminology and technology overviews

Following are overviews that will help you to understand the various technologies that are required to deploy 802.1X authenticated wired access.

Note
In this document, 802.1X authenticated wired access is referred to as wired access.

IEEE 802.1X

The IEEE 802.1X standard defines the port-based network access control that is used to provide authenticated wired access to Ethernet networks. This port-based network access control uses the physical characteristics of the switched Local Area Network (LAN) infrastructure to authenticate devices attached to a LAN port. Access to the port can be denied if the authentication process fails. Although this standard was designed for wired Ethernet networks, it has also been adapted for use on 802.11 wireless LANs.

IEEE 802.1X-capable wired Ethernet switches

To deploy 802.1X wired access you must install and configure one or more 802.1X-capable Ethernet switches on your network. The switches must be compatible with the Remote Authentication Dial-In User Service (RADIUS) protocol.

When 802.1X and RADIUS-compliant switches are deployed in a RADIUS infrastructure, with a RADIUS server such as an NPS server, they are called *RADIUS clients*.

IEEE 802.3 Ethernet

IEEE 802.3 is a collection of standards that defines the Layer-1 (physical layer) and Layer-2 (data-link layer media access control (MAC)) of wired Ethernet. 802.3 Ethernet is typically implemented on a LAN, and also in some wide area network (WAN) applications.

Network Policy Server

Network Policy Server (NPS) lets you centrally configure and manage network policies by using the following three components: RADIUS server, RADIUS proxy, and Network Access Protection (NAP) policy server. NPS is required to deploy 802.1X wired access.

Server certificates

Wired access deployment requires server certificates for each NPS server that performs 802.1X authentication.

A *server certificate* is a digital document that is commonly used for authentication and to help secure information on open networks. A certificate securely binds a public key to the entity that holds the corresponding private key. Certificates are digitally signed by the issuing certification authority (CA), and they can be issued for a user, a computer, or a service.

A CA is an entity responsible for establishing and vouching for the authenticity of public keys that belong to subjects (usually users or computers) or other CAs. Activities of a CA can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and revoking certificates.

Active Directory Certificate Services (AD CS) is a Windows Server 2012 server role that issues certificates as a network CA. An AD CS certificate infrastructure, also called a *public key infrastructure (PKI)*, provides customizable services for issuing and managing certificates for the enterprise.

EAP

Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by enabling additional authentication methods that use credential and information exchanges of arbitrary lengths. With EAP authentication, both the network access client and the authenticator (such as an NPS server) must support the same EAP type for successful authentication to occur.

New and changed functionality

In Windows Server 2012, wired access includes only minimal changes to the wired access solution provided in Windows Server 2008 R2. That change is summarized as follows:

Feature/functionality	Previous operating system	New operating system
The addition of EAP-Tunneled Transport Layer Security (EAP-TTLS) to the list of network authentication methods that are included by default	Not included	Included by default

See also

Following is a table of resources related to 802.1X authenticated wired access.

Content type	References
Product evaluation	IEEE 802.1X Authenticated Wired Access – Cable Guy article |
Planning	Windows Server 2008 802.1X Authenticated Wired Access Design Guide |
Deployment	Windows Server 2008 802.1X Authenticated Wired Access Deployment Guide | Windows Server 2008 R2 Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access
Operations	Windows Server 2008 R2 Netsh Commands for Wired Local Area Network (LAN) |
Troubleshooting	Windows Server 2008 R2 Network Diagnostics Framework (NDF) and Network Tracing |
Security	Not applicable
Tools and settings	Content not available
Community resources	Content not available
Related technologies	Windows Server 2008 R2 802.1X Authenticated Wireless Access | Windows Server 2008 R2 Network Policy and Access Services

What's New in 802.1X Authenticated Wired Access in Windows Server 2012 R2

This topic provides information about the new features for 802.1X Authenticated Wired Access in Windows Server® 2012 R2 and Windows® 8.1.

Extending the use of passwords for Enterprise wired Ethernet access

If you have deployed password-based 802.1X authentication methods for wired and wireless connections through Ethernet switches and wireless access points, users with non-domain joined computers and devices that are running Windows 8.1 and Windows Server 2012 R2 can bring

their own devices to your organization and enjoy the advantages of password-based credential reuse.

When password-based Extensible Authentication Protocol (EAP) and Protected EAP (PEAP) authentication methods are deployed, end users can provide their credentials the first time they connect to your organization's network, then connect to all the resources they want to without being prompted repeatedly for their credentials, because the credentials are stored on the local computer for reuse.

This is especially useful for users who are connecting to multiple network resources, such as Enterprise intranet Web sites, Enterprise printers, and line of business applications.

For security reasons, when the user's computer or device disconnects from the network, the stored credentials are discarded.

This feature is available for non-domain joined computers and devices that are running Windows 8.1 and Windows Server 2012 R2 when you have deployed the following authentication methods on your network.

- EAP with Microsoft Challenge Handshake Protocol version 2 (EAP-MS-CHAP v2)
- PEAP-EAP-MS-CHAP v2
- EAP-TTLS with EAP-MS-CHAP v2

In Windows 8.1 and Windows Server 2012 R2, this feature is enabled by default. You can use the following registry key to disable or re-enable user password storing if you have previously disabled it:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eaphost\Configuration

Name	Type	Default	Exists by default
DisableUsrPwdStoring	DWORD (1 or 0)	0 (false)	No

When this registry key is set to 1 (true) EAP methods do not store any credential information in Credential Manager.

Managing the New Wired Network (IEEE 802.3) Policies Settings

This topic provides details about configuring 802.1X authenticated wired Ethernet access using the Wired Network (IEEE 802.3) Policies extension of Group Policy. Information is presented in two main sections:

Configuring Wired Network (IEEE 802.3) Policies provides step-by-step details for configuring 802.1X authenticated Wired Access based on one of the authentication methods you have selected:

- Configure a wired connection profile for PEAP-MS-CHAP v2

- Configure a wired connection profile for PEAP-TLS
- Configure a wired connection profile for EAP-TLS

Per-setting details provides information about the individual settings in the Wired Network (IEEE 802.3) Policies extension of Group Policy:

- General – settings
- Security - settings
- Advanced security settings for Wired and Wireless Network Policies

Configuring Wired Network (IEEE 802.3) Policies

This section provides step-by-step details to configure Wired Network (IEEE 802.3) Policies extension of Group Policy for 802.1X authenticated wired network access.

- Configure a wired connection profile for PEAP-MS-CHAP v2
- Configure a wired connection profile for PEAP-TLS
- Configure a wired connection profile for EAP-TLS

Configure a wired connection profile for PEAP-MS-CHAP v2

This procedure provides the steps that are required to configure a wired access connection profile for Protected Extensible Authentication Protocol–Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAP v2) for authentication by using secure passwords.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

Tip

For more information about individual controls on any active dialog box in Wired Network (IEEE 802.3) Policies, press **F1** while viewing that dialog box.

To configure a wired connection profile for PEAP-MS-CHAP v2

1. If you have not already done so, open the New Wired Network (IEEE 802.3) Policies Properties dialog.
2. In **New Wired Network (IEEE 802.3) Policies Properties**, on the **General** tab, in **Policy Name**, type a name for your network policy, or leave the default name **New Wired Network Policy**.
3. In **Description**, type a description for your network policy.
4. Select **Use Windows Wired Auto Config service for clients** to specify that Wired AutoConfig is used to configure wired network adapter settings.
5. To prevent the use of shared user credentials for authentication for computers running Windows 8 and Windows 7, select **don’t allow shared user credentials for network authentication**.
6. To specify the duration for which automatic authentication attempts will be blocked from occurring after a failed authentication attempt, select **Enable block period (minutes)**.

7. Click the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

      When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, **Auth Period**, and **Eapol-Start Message** are sufficient for common wired access deployments.
   b. To enable Single Sign On, select **Enable Single Sign On for this network**.
   c. To specify whether Single Sign On is performed before or after user logon, select either **Perform immediately before User Logon** or **Perform immediately after User Logon**.
   d. To specify a maximum connectivity delay, select **Max delay for connectivity (seconds)**, and then specify a value, or leave the default value of 10.
8. To specify whether additional dialogs may be displayed during Single Sign On, select **Allow additional dialogs to be displayed during Single Sign On**.
9. If your network is configured to use different VLANs for authentication that uses computer credentials and authentication that uses user credentials, select **This network uses different VLAN for authentication with machine and user credentials**.
10. Click **OK** to return to the **Security** tab. In **Select a network authentication method**, select **Protected EAP (PEAP)**, and then click **Properties**. The **Protected EAP Properties** page opens.
11. In **Protected EAP Properties**, verify that **Validate server certificate** is selected.
12. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computer running Network Policy Server (NPS).

    Note
    This setting limits which root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, clients will trust all root CAs listed in their trusted root certification authority store.
13. Select **Do not prompt user to authorize new servers or trusted certification authorities**. Selecting this setting provides an enhanced user experience and better security.
14. In the **Select Authentication Method** list, select **Secured password (EAP-MS-CHAP v2)**.
15. To enable PEAP fast reconnect, select **Enable Fast Reconnect**.
16. If Network Access Protection (NAP) is configured on your network, select **Enable Quarantine checks**. Otherwise, clear this check box.
17. Click **Configure**. In the **EAP MSCHAPv2 Properties** dialog box, verify **Automatically use my Windows logon name and password (and domain if any)** is selected, click **OK**, and then click **OK** to close **Protected EAP Properties**.
18. Click **OK** to close the **Security** tab, and then click **OK** again to close the Windows Vista Wired (IEEE 802.3) Network Policy.

Back to top

Configure a wired connection profile for PEAP-TLS

This procedure provides the steps that are required to configure a wired access connection profile for Protected Extensible Authentication Protocol – Transport Layer Security (PEAP-TLS) for authentication by using smart cards or user and computer digital certificates.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

Tip

For more information about individual controls on any active dialog box in Wired Network (IEEE 802.3) Policies, press **F1** while viewing that dialog box.

To configure a wired connection profile for PEAP-TLS

1. If you have not already done so, open the New Wired Network (IEEE 802.3) Policies Properties dialog.
2. In **New Wired Network (IEEE 802.3) Policies Properties**, on the **General** tab, in **Policy Name**, type a name for your network policy, or leave the default name **New Wired Network Policy**.
3. In **Description**, type a description for your network policy.
4. Select **Use Windows Wired Auto Config service for clients** to specify that Wired AutoConfig is used to configure wired network adapter settings.
5. To prevent the use of shared user credentials for authentication for computers running Windows 8 and Windows 7, select **don't allow shared user credentials for network authentication**.
6. To specify the duration for which automatic authentication attempts will be blocked from occurring after a failed authentication attempt, select **Enable block period (minutes)**.
7. Click the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

      When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, **Auth Period**, and **Eapol-Start Message** are sufficient for common wired access deployments.

   b. To enable Single Sign On, select **Enable Single Sign On for this network**.
   c. To specify whether Single Sign On is performed before or after user logon, select either **Perform immediately before User Logon** or **Perform immediately after User Logon**.
   d. To specify a maximum connectivity delay, select **Max delay for connectivity (seconds)**, and then specify a value, or leave the default value of 10.
8. To specify whether additional dialogs may be displayed during Single Sign On, select **Allow additional dialogs to be displayed during Single Sign On**.

9. If your network is configured to use different VLANs for authentication that uses computer credentials and authentication that uses user credentials, select **This network uses different VLAN for authentication with machine and user credentials**.
10. Click **OK** to return to the **Security** tab. In **Select a network authentication method**, select **Protected EAP (PEAP)**, and then click **Properties**. The **Protected EAP Properties** page opens.
11. In **Protected EAP Properties**, verify that **Validate server certificate** is selected.
12. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computer running Network Policy Server (NPS).

    Note
    This setting limits which root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, clients will trust all root CAs listed in their trusted root certification authority store.
13. Select **Do not prompt user to authorize new servers or trusted certification authorities**. Selecting this setting provides an enhanced user experience and better security.
14. In the **Select Authentication Method** list, select **Smart Card or other certificate**, and then click **Configure**. The **Smart Card or other Certificate Properties** dialog box opens.
15. In the **Smart Card or other Certificate Properties** dialog box, in **When connecting**, for smart card deployments, select **Use my smart card**. Otherwise, for computer and user digital certificate deployments, select **Use a certificate on this computer**.
16. Verify that **Validate server certificate** is selected.
17. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computer running Network Policy Server (NPS).
18. Select **Do not prompt user to authorize new servers or trusted certification authorities**.
19. Click **OK**. The **Smart Card or other Certificate Properties** dialog box closes, returning you to **Protected EAP Properties**.
20. On the **Protected EAP Properties** dialog box, to enable PEAP fast reconnect, select **Enable Fast Reconnect**.
21. If Network Access Protection (NAP) is configured on your network, select **Enable Quarantine checks**. Otherwise, clear this check box.
22. Click **OK**, to close the **Protected EAP Properties** dialog box, click **OK** to close the **Security** tab, and then click **OK** again to close the Windows Vista Wired (IEEE 802.3) Network Policy.

Back to top

Configure a wired connection profile for EAP-TLS

This procedure provides the steps that are required to configure a wired access connection profile for Extensible Authentication Protocol – Transport Layer Security (PEAP-TLS) for authentication by using smart cards or user and computer digital certificates.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

Tip

For more information about individual controls on any active dialog box in Wired Network (IEEE 802.3) Policies, press **F1** while viewing that dialog box.

To configure a wired connection profile for EAP-TLS

1. If you have not already done so, open the New Wired Network (IEEE 802.3) Policies Properties dialog.
2. In **New Wired Network (IEEE 802.3) Policies Properties**, on the **General** tab, in **Policy Name**, type a name for your network policy, or leave the default name **New Wired Network Policy**.
3. In **Description**, type a description for your network policy.
4. Select **Use Windows Wired Auto Config service for clients** to specify that Wired AutoConfig is used to configure wired network adapter settings.
5. To prevent the use of shared user credentials for authentication for computers running Windows 8 and Windows 7, select **don't allow shared user credentials for network authentication**.
6. To specify the duration for which automatic authentication attempts will be blocked from occurring after a failed authentication attempt, select **Enable block period (minutes)**.
7. Click the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

      When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, **Auth Period**, and **Eapol-Start Message** are sufficient for common wired access deployments.

   b. To enable Single Sign On, select **Enable Single Sign On for this network**.
   c. To specify whether Single Sign On is performed before or after user logon, select either **Perform immediately before User Logon** or **Perform immediately after User Logon**.
   d. To specify a maximum connectivity delay, select **Max delay for connectivity (seconds)**, and then specify a value, or leave the default value of 10.
8. To specify whether additional dialogs may be displayed during Single Sign On, select **Allow additional dialogs to be displayed during Single Sign On**.
9. If your network is configured to use different VLANs for authentication that uses computer credentials and authentication that uses user credentials, select **This network uses**

different VLAN for authentication with machine and user credentials.

10. Click **OK** to return to the **Security** tab. In the **Select Authentication Method** list, select **Smart Card or other certificate**, and then click **Properties**. The **Smart Card or other Certificate Properties** dialog box opens.
11. In the **Smart Card or other Certificate Properties** dialog box, in **When connecting**, for smart card deployments, select **Use my smart card**. Otherwise, for computer and user digital certificate deployments, select **Use a certificate on this computer**.
12. Verify that **Validate server certificate** is selected.
13. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computer running Network Policy Server (NPS).
14. Select **Do not prompt user to authorize new servers or trusted certification authorities**.
15. Click **OK**. The **Smart Card or other Certificate Properties** dialog box closes, returning you to the **Security** tab.
16. Click **OK** to close the **Security** tab, and then click **OK** again to close the Windows Vista Wired (IEEE 802.3) Network Policy.

Back to top

Per-setting details

This section provides details about settings in the Wired Network (IEEE 802.3) Policies extension of Group Policy:

1. General - settings
2. Security - settings

Note
Information about the Advanced security settings is provided in the topic, Advanced Security Settings for Wired and Wireless Network Policies. Clicking this link will open a new Web page. A link is provided in the Additional Resources section of that topic to return you to this page.

General - settings

Use the Wired Network (IEEE 802.3) Policies General tab to specify whether the Wired AutoConfig Service is used to configure local area network (LAN) adapters to connect to the wired network. You can also specify the policy name and description.

Item	Details
Policy Name	Provides a location for you to type a name for the Network Policy that is applied to your wired clients running Windows Vista.

Item	Details
	Note When you name the Wired Network (IEEE 802.3) Policy, the name is displayed as the title of its property pages, and under **Name** in the details pane in the Group Policy Management Console (GPMC) for that policy.
Description	Provides a location for you to type a description for your Wired Network (IEEE 802.3) Policy. **Note** The description you type is displayed under **Description** in the details pane of the Group Policy Management Console (GPMC) for that policy.
Use Windows wired network service for clients	Specifies that Wired AutoConfig Service is used to configure and connect clients running Windows Vista to the 802.3 wired Ethernet network. Default = enabled
Don’t allow shared user credentials for network authentication	Specifies that users with computers running Windows 7 are not allowed to store their user credentials (such as user name and password), which the computer can then use to log on to the network (even though the user is not actively logged on to the computer). Default = not selected. Users are allowed to enter and store their user credentials in profiles that they configure. The following statements summarize the behavior of this setting: • This setting is provided in both the Wired Network (IEEE 802.3) Policies and the version of the Wireless Network (IEEE 802.11) Policies that you can configure for client computers running Windows Vista and subsequent releases of Windows. • This setting is per-Group Policy object (GPO), and within each GPO it is specific to the Wired Network (IEEE 802.3) Policies

Item	Details
	and Wireless Network (IEEE 802.11) Policies. • For as long as the default setting remains unchanged, the local administrator can use the **netsh LAN set allowexplicitcreds** command on client computers running Windows 7 to allow or disallow users to store their user credentials. • After the first time that the default is changed the administrator can use this setting to enable or disable the capability for users to configure non-Group Policy based user profiles to store their user credentials. User credentials cannot be stored with Group Policy profiles.
Enable block period (minutes)	Specifies whether to prohibit computers running Windows 7 from making auto connection attempts to the network for a specified amount of time, following a failed authentication attempt. The blocked state is reset upon a manual connection attempt, a session change, or a media connect. Default = not enabled. If enabled, the default is 20 minutes. The valid range of minutes is 1-60.

Back to top

Security - settings

Security configuration items include all of the settings on the **Security** tab. These settings specify whether to perform 802.1X authentication for connecting clients and which network authentication method to use.

Item	Details
Enable use of IEEE 802.1X authentication for network access	Specifies that 802.1X authentication is performed for access requests to the wired network. Default = enabled
Select a network authentication	• Specifies the network authentication method that connecting clients use:

Item	Details
	• **Smart Card or other certificate** (EAP-TLS) • **Protected EAP (PEAP)** Default = Protected EAP (PEAP)
Properties	Opens the properties page of the selected network authentication method. For setting information specific to network authentication methods, see: Extensible Authentication Protocol (EAP) Settings for Network Access
Authentication Mode	Specifies how network authentication is performed: • **User re-authentication**. An 802.1X-compliant device always uses security credentials based on the computer's current state. Authentication is performed by using the computer credentials when no users are logged on to the computer. When a user logs on to the computer, authentication is always performed using the user credentials. This is the recommended setting. • **Computer only**. Authentication is always performed using only the computer credentials. • **User authentication**. Specifies that when users are not logged on to the computer, authentication is performed by using the computer credentials. After a user logs on to the computer, authentication is still based on the computer credentials. Authentication is performed by using the user credentials if the user travels to a new wireless access point. • **Guest authentication**. Allows connections to the network that are regulated by the restrictions and permissions set for the Guest user account. Default = User re-authentication
Max Authentication Failures	Specifies the maximum number of failed

Item	Details
	authentication attempts that can occur with a specific set of credentials before notification is displayed to indicate that authentication has failed. Default = 3
Advanced	Opens the Advances security dialog. For information about advanced security settings, see Advanced Security Settings for Wired and Wireless Network Policies

Back to top

See Also

802.1X Authenticated Wired Access Overview

802.1X Authenticated Wireless Access Overview

This document provides introductory information about Institute of Electrical and Electronics Engineers (IEEE) 802.1X authenticated access for IEEE 802.11 wireless access. Links to resources with information about technologies that are closely related to 802.1X authenticated wireless access, or otherwise relevant to wireless access are also provided.

Notes

In addition to this topic, the following 802.1X Authenticated Wireless Access documentation for Windows Server 2012 is also available.

- What's New in 802.1X Authenticated Wireless Access in Windows Server 2012 R2
- Improvements to Certificate-based Authentication
- Wireless LAN Service Overview
- Managing the Wireless Network (IEEE 802.11) Policies
- Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access
- New Wireless Connection Processes

Did you mean…

- 802.1X Authenticated Wired Access
- Netsh Commands for Wireless Local Area Network (WLAN) in the Windows Server 2008 R2 and Windows Server 2008 technical library on TechNet.

- 802.1X Authenticated Wireless Access in the Windows Server 2008 R2 and Windows Server 2008 technical library on TechNet.

Feature description

IEEE 802.1X authentication provides an additional security barrier for your intranet that you can use to prevent guest, rogue, or unmanaged computers that cannot perform a successful authentication from connecting to your intranet.

For the same reason that administrators deploy IEEE 802.1X authentication for IEEE 802.3 wired networks—enhanced security—network administrators want to implement the IEEE 802.1X standard to help protect their wireless network connections. Just as an authenticated wired client must submit a set of credentials to be validated before being allowed to send frames over the wired Ethernet intranet, an IEEE 802.1X wireless client must also perform authentication prior to being able to send traffic over its wireless access point (AP) port, and over the network.

Important terminology and technology overviews

Following are overviews that will help you to understand the various technologies that are required to deploy 802.1X authenticated wireless access.

Note
In this document, 802.1X authenticated wireless access is referred to as WiFi access.

IEEE 802.1X

The IEEE 802.1X standard defines the port-based network access control that is used to provide authenticated WiFi access to corporate networks. This port-based network access control uses the physical characteristics of the 802.1X capable wireless APs infrastructure to authenticate devices attached to a LAN port. Access to the port can be denied if the authentication process fails. Although this standard was first designed for wired Ethernet networks, it has also been adapted for use on 802.11 wireless LANs.

IEEE 802.1X-capable wired Ethernet switches

To deploy 802.1X wireless access you must install and configure one or more 802.1X-capable wireless APs on your network. The wireless APs must be compatible with the Remote Authentication Dial-In User Service (RADIUS) protocol.

When 802.1X and RADIUS-compliant wireless APs are deployed in a RADIUS infrastructure, with a RADIUS server such as an NPS server, they are called *RADIUS clients*.

IEEE 802.11 wireless

IEEE 802.11 is a collection of standards that defines the Layer-1 (physical layer) and Layer-2 (data-link layer media access control (MAC)) of WiFi access.

Network Policy Server

Network Policy Server (NPS) lets you centrally configure and manage network policies by using the following three components: RADIUS server, RADIUS proxy, and Network Access Protection (NAP) policy server. NPS is required to deploy 802.1X wireless access.

Server certificates

WiFi access deployment requires server certificates for each NPS server that performs 802.1X authentication.

A *server certificate* is a digital document that is commonly used for authentication and to help secure information on open networks. A certificate securely binds a public key to the entity that holds the corresponding private key. Certificates are digitally signed by the issuing certification authority (CA), and they can be issued for a user, a computer, or a service.

A CA is an entity responsible for establishing and vouching for the authenticity of public keys that belong to subjects (usually users or computers) or other CAs. Activities of a CA can include binding public keys to distinguished names through signed certificates, managing certificate serial numbers, and revoking certificates.

Active Directory Certificate Services (AD CS) is a Windows Server 2012 server role that issues certificates as a network CA. An AD CS certificate infrastructure, also called a *public key infrastructure (PKI)*, provides customizable services for issuing and managing certificates for the enterprise.

EAP

Extensible Authentication Protocol (EAP) extends Point-to-Point Protocol (PPP) by enabling additional authentication methods that use credential and information exchanges of arbitrary lengths. With EAP authentication, both the network access client and the authenticator (such as an NPS server) must support the same EAP type for successful authentication to occur.

New and changed functionality

In Windows Server 2012, WiFi access includes only minimal changes to the wired access solution provided in Windows Server 2008 R2. That change is summarized as follows:

Feature/functionality	Previous operating system	New operating system
The addition of EAP-Tunneled Transport Layer Security (EAP-TTLS) to the list of network authentication methods that are included by default	Not included	Included by default

See also

Following are additional resources that pertain to 802.1X authenticated wireless access.

Content type	References
Product evaluation	Connecting to Wireless Networks with Windows 7 The Cable Guy - July 2010 |
Planning	Windows Server 2008 802.1X Authenticated Wireless Access Design Guide |
Deployment	Windows Server 2008 802.1X Authenticated Wireless Access Deployment Guide | Windows Server 2012 Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access | Windows Server 2008 R2 Core Network Companion Guide: Deploying Password-based 802.1X Authenticated Wireless Access
Operations	Windows Server 2012 Managing the Wireless Network (IEEE 802.11) Policies extension of Group Policy| Extensible Authentication Protocol (EAP) Settings for Network Access | Windows Server 2012 Wireless LAN Service Overview| Windows Server 2008 R2 Netsh Commands for Wireless Local Area Network (WLAN) |
Troubleshooting	Windows Server 2008 R2 Network Diagnostics Framework (NDF) and Network Tracing |
Security	Content not available
Tools and settings	Windows Server 2012 Managing the Wireless Network (IEEE 802.11) Policies extension of Group Policy
Community resources	Content not available
Related technologies	Windows Server 2008 R2 802.1X Authenticated Wired Access | Windows Server 2008 R2 Network Policy and Access Services

What's New in 802.1X Authenticated Wireless Access in Windows Server 2012 R2

This topic provides information about the new features for 802.1X Authenticated Wireless Access in Windows Server® 2012 R2 and Windows® 8.1.

Role and technology description

IEEE 802.1X authentication provides an additional security barrier for your intranet that you can use to prevent guest, rogue, or unmanaged computers that cannot perform a successful authentication from connecting to your intranet.

For the same reason that administrators deploy IEEE 802.1X authentication for IEEE 802.3 wired networks—enhanced security—network administrators want to implement the IEEE 802.1X standard to help protect their wireless network connections. Just as an authenticated wired client must submit a set of credentials to be validated before being allowed to send frames over the wired Ethernet intranet, an IEEE 802.1X wireless client must also perform authentication prior to being able to send traffic over its wireless access point (AP) port, and over the network.

New and changed functionality

The following table provides a list of the new features for Wireless Access.

Feature/functionality	New or updated?	Description
Support for 802.11ac	New	802.11ac provides better bandwidth and faster connections than previous standards.
Wireless Display	New	Wirelessly project your laptop or tablet screen for viewing on larger monitors and televisions.
Extending the use of passwords for Enterprise wireless access	New	Your organization employees can bring their own non-domain joined wireless devices running Windows Server 2012 R2 and Windows 8.1 to work and enter passwords only once during their wireless session.

Support for 802.11ac

Wireless connections in Windows® 8.1 and Windows Server® 2012 R2 now support Institute of Electrical and Electronics Engineers (IEEE) 802.11ac wireless access.

802.11ac provides higher speeds, throughput, reliability, and quality of wireless communications than previous standards, such as 802.11n, 802.11g, and 802.11b. This is possible because 802.11ac uses wider RF bandwidths, more spatial streams, multi user multiple in, multiple out (MIMO), and enhanced high density modulation.

802.11ac allows better performance for applications that require higher bandwidth, such as wireless display, streaming high fidelity videos, distribution of HDTV and other content between devices, rapid upload and download of files, and playing online games.

Note

To see the benefits of using 802.11ac connections, your other wireless network hardware, such as wireless network adapters and wireless access points, must also support 802.11ac. In addition, 802.11ac functions regardless of whether you have deployed 802.1X authentication; the two technologies are not related.

The following table provides a comparison between 802.11n and 802.11ac.

Feature	802.11n	802.11ac
Operation frequency	2.4 GHz and 5 GHz	5 GHz Only
Channel width	20 MHz, 40 MHz	20 MHz, 40 MHz, 80 MHz; 160 MHz optional
Modulation schemes	64 QAM	BPSK, QPSK, 16QAM, 64 QAM, 256 QAM (optional)
Spatial streams	4 spatial MIMO	Up to 8 Spatial MIMO
Multi User MIMO	No	Yes

In Windows® 8.1 and Windows Server® 2012 R2, the following Network Shell (netsh) commands in the **netsh wlan** context are modified to include information about 802.11ac:

- show drivers

  When you run the **netsh wlan show drivers** command, if the drivers support 802.11ac, the command output item **Radio types supported** now includes 802.11ac.

- show networks

  When you run the **netsh wlan show networks** command with the **mode** parameter set to **bssid**, if the network driver supports 802.11ac, the command output item **Radio types supported** now includes 802.11ac. Following is an example of this command with the mode parameter.

```
netsh wlan show networks mode=bssid
```

- show profiles

  When you run the **netsh wlan show profiles** command, if the profile includes the radio type 802.11ac, the command output item **Radio type** now includes 802.11ac.

Wireless Display

You can use Wireless Display to wirelessly project the screen from your computer or device that is running Windows 8.1 or Windows® RT 8.1 to any Miracast-compatible display, such as a television, monitor, or projector. This capability is especially useful in circumstances where your source device has a small screen and you are projecting onto a device with a larger screen, allowing you to easily share experiences and information with others.

Note
Wireless Display is available in Windows 8.1 and Windows® RT 8.1. In addition, Wireless Display functions regardless of whether you have deployed 802.1X authentication; the two technologies are not related.

For example, at home you can use Wireless Display to project your laptop or tablet screen to your large screen television, allowing your family to engage in planning a family vacation, view photos, watch videos, or surf the Internet together.

In the office you can use Wireless Display to connect to a projector from your laptop or tablet, then run Microsoft PowerPoint and other applications for your presentation.

Wireless Display is compatible with the WFA Miracast specification and works with Miracast-certified receivers. For more information, see Certified Products.

To use Wireless Display, you must move your computer or device running Windows 8.1 or Windows RT 8.1 within range of a Miracast-compatible display, and then access the Project menu by pressing the Windows key and K simultaneously (Win+K). The **Project** screen opens and provides a list of compatible devices to which you can connect. If you have not previously connected to a device, click **Add a Wireless Display** to add a device.

Note
If you have upgraded your computer or device from Windows 8 to Windows 8.1, it is recommended that you obtain and install the most recent drivers for your devices from Windows Update or your product manufacturer before you use Wireless Display.

For more information, see Project to a wireless display with Miracast.

Extending the use of passwords for Enterprise wireless access

If you have deployed password-based 802.1X authentication methods for wired and wireless connections through Ethernet switches and wireless access points, users with non-domain joined computers and devices that are running Windows 8.1 and Windows Server 2012 R2 can bring

their own devices to your organization and enjoy the advantages of password-based credential reuse.

When password-based Extensible Authentication Protocol (EAP) and Protected EAP (PEAP) authentication methods are deployed, end users can provide their credentials the first time they connect to your organization's network, then connect to all the resources they want to without being prompted repeatedly for their credentials, because the credentials are stored on the local computer for reuse.

This is especially useful for users who are connecting to multiple network resources, such as Enterprise intranet Web sites, Enterprise printers, and line of business applications.

For security reasons, when the user's computer or device disconnects from the network, the stored credentials are discarded.

This feature is available for non-domain joined computers and devices that are running Windows 8.1 and Windows Server 2012 R2 when you have deployed the following authentication methods on your network.

- EAP with Microsoft Challenge Handshake Protocol version 2 (EAP-MS-CHAP v2)
- PEAP-EAP-MS-CHAP v2
- EAP-TTLS with EAP-MS-CHAP v2

In Windows 8.1 and Windows Server 2012 R2, this feature is enabled by default. You can use the following registry key to disable or re-enable user password storing if you have previously disabled it:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Eaphost\Configuration

Name	Type	Default	Exists by default
DisableUsrPwdStoring	DWORD (1 or 0)	0 (false)	No

When this registry key is set to 1 (true) EAP methods do not store any credential information in Credential Manager.

Improvements to Certificate-based Authentication

In Windows 8, there are many improvements in how you can manage the use of certificates for authentication and in how Windows selects certificates during the authentication process.

In Windows 7 and earlier operating systems, using certificates for user or computer authentication for wireless and virtual private network (VPN) connections is a complex and non-intuitive experience due to the following issues:

- **Manual certificate selection is frequently required**. In many cases, you must manually select a certificate. Often it is unclear which certificate to select, causing a circumstance where you must employ a trial and error method to make the right selection.
- **Windows can have difficulty filtering the list of certificates**. Windows uses some basic rules to filter out some of the certificates, such as ignoring expired or revoked certificates, however if there are multiple certificates that have the same subject name and issuer name, Windows takes a different action depending on whether the certificate is a user certificate or a computer certificate. For user certificates, Windows prompts the user to make a manual selection of which certificate to use. For computer certificates, the certificate with the highest weight is selected. If the selected certificate is the incorrect certificate for the connection, authentication fails. These filtering mechanisms are very rudimentary and user intervention is still required in most cases.
- **Problems selecting certificates on smart cards**. If there are multiple certificates present on a smart card, Windows enumerates a list of the certificates and then always picks the last certificate from the list. In many cases, this is not the correct certificate to use for the authentication attempt. This is problematic when the same smart card is used for multiple purposes and contains a separate certificate for every case. There is no workaround available to select the right certificate.
- **Registry certificates cannot be used for two factor authentication**. Although certificates are ideal candidates for two-factor authentication, registry certificates – which are protected by a strong private key and are the most appropriate certificates for two-factor authentication - cannot be used. The reason for this is that Windows does not support registry certificates and completely ignores them. As a result, organizations must deploy and manage complex and expensive smart card solutions rather than using registry based certificates.

In Windows Server 2012 and Windows 8, certificate filtering provides Windows with the ability to select the correct certificate for use with certificate-based authentication during the network connection process.

Certificate filtering

Certificate filtering allows you to configure rules during the process of creating either Wireless or VPN profiles on Windows Server 2012. When a Wireless or VPN profile is applied to a computer running Windows 8, these rules limit the number of certificates that Windows 8 can choose during the authentication process.

You can configure certificate filtering in Windows Server 2012 so that computers that are running Windows 8 that are selecting a certificate for an authentication attempt can automatically eliminate all certificates from consideration except for the certificate that you want them to choose. This certificate filtering therefore assists client computers in achieving a higher number of successful authentication attempts, because Windows 8 selects the correct certificate for the attempt.

The certificate filtering criteria that you select can apply to registry, smart card, and client computer certificates.

While you are configuring a Wireless or VPN profile in Windows Server 2012, you can use the following new filters:

- **Certificate Issuer**. When you configure the certificate issuer filter, you can specify that computers that are running Windows 8 must select a certificate that is issued by a specific CA. When you configure this filter, the filter enumerates the root and intermediate root certification authorities (CAs) for all certificates that are present in the certificate store on the local computer, and allows you to select the correct CA from the list.
- **Extended Key Usage (EKU)**. The EKU filter allows you to choose which EKUs must be present in the certificate. When you configure the EKU filter, you can specify that computers that are running Windows 8 must enumerate all EKU properties for all certificates that are present in the certificate store on the local computer. You can also add a custom EKU to the filtering rule by specifying the object identifier (OID) of the EKU.

Note

Extended Key Usage extensions are also commonly called Enhanced Key Usage extensions, and the acronym for both terms is EKU.

The illustration below depicts the **Configure Certificate Selection** dialog box where you can apply certificate filtering.

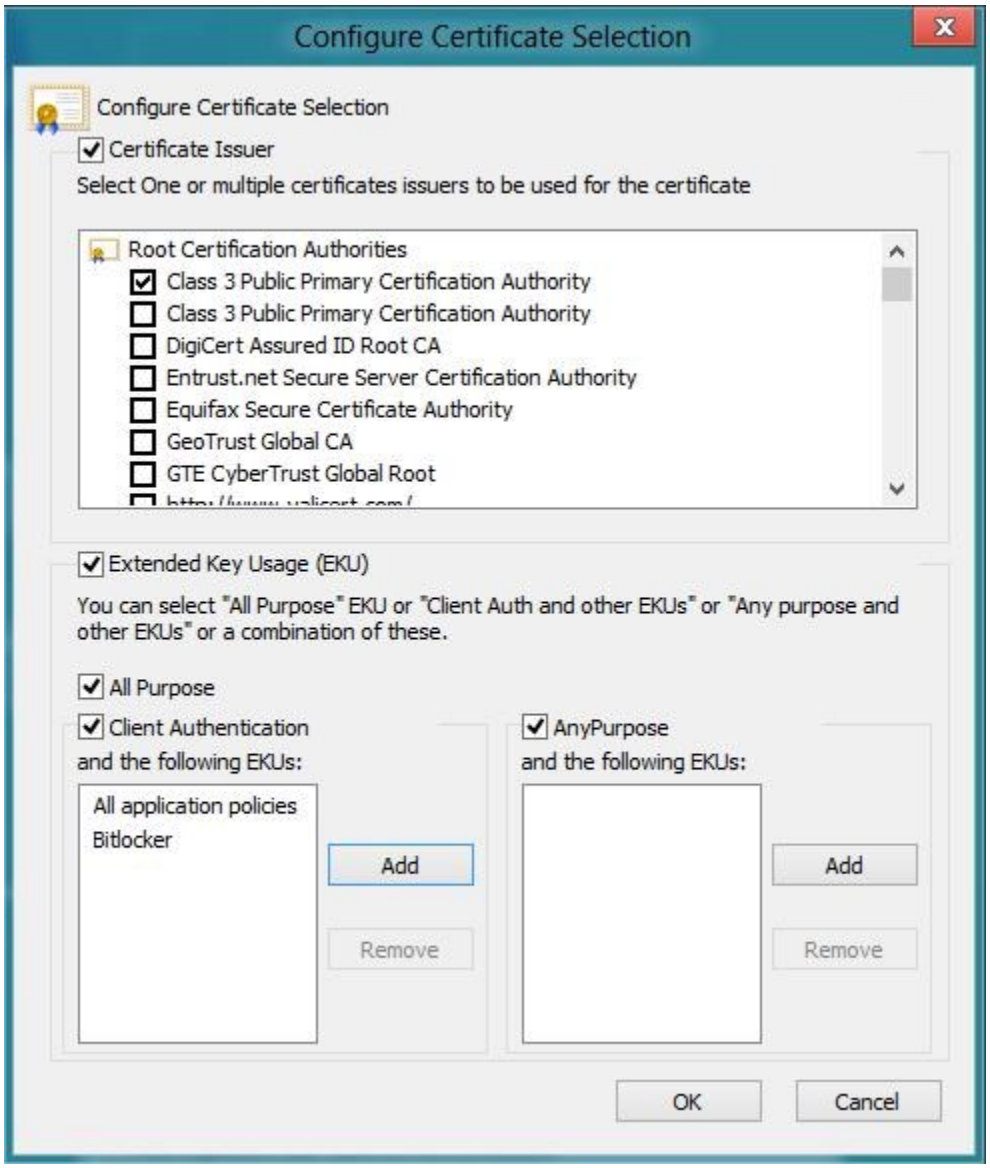

When a Wireless profile where these rules are configured is applied to a computer running Windows 8 and the client attempts to connect to a network, Windows applies the filtering rules and selects one or more certificates that match all of the rules. If multiple compliant certificates are found, Windows displays a dialog box that lists the certificates so that you can choose one. This dialog box is displayed for registry, smart card, and client computer certificates.

Certificate weight as a filtering mechanism

When a Smart Card certificate is used for Pre-Logon-Access Provider (PLAP) scenarios, the weight of the certificate is also used for filtering. The weight of a certificate is determined by the certificate revocation list Distribution Point (CDP) and by the Authority Information Access (AIA)

properties that are present in the certificate. AIA has a weight of 2 and CDP has a weight of 1. If both properties are present then Windows adds their weights together to determine the certificate weight. After this process, Windows selects and uses the certificate that has the highest weight value.

Private Key Protected Registry Certificates

Windows 8 recognizes private key protected registry certificates that are present in the Current User certificate store on the local computer, and these certificates can be used for two-factor authentication. Windows does not look for private key protected certificates in any other certificate store.

Certificate filtering compatibility information

Wireless and VPN profiles created on Windows Server 2012 that contain certificate filters can be applied to computers that are running operating systems earlier than Windows 8, such as Windows 7 and Windows Vista, however the filters cannot be parsed by these earlier operating systems and only the legacy rudimentary filtering is applied during the authentication process.

Wireless LAN Service Overview

The Wireless LAN service is a feature in Windows Server® 2012 that you can use to enable the wireless WLAN AutoConfig service, and to configure the WLAN AutoConfig service for automatic startup. Once enabled, the WLAN AutoConfig service dynamically selects which wireless network the computer automatically connects to, and configures the necessary settings on the wireless network adapter. This includes automatically selecting and connecting to a more preferred wireless network when one becomes available.

Enabling the Wireless LAN Service

In Windows Server® 2008, installing the Wireless LAN Service server feature did not require restarting the computer. The restart was added in Windows Server 2012 to support the Windows Connection Manager service. The Windows Connection Manager service runs in an svchost.exe process which is already running when the Wireless LAN Service feature is installed, which is why a restart is required. Additionally, to start the WAL AutoConfig service requires an additional step, which is explained in the following procedure.

To enable the Wireless LAN Service

1. In **Server Manager Dashboard**, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
2. Click **Next**. In **Select installation type**, select **Role-based or feature-based installation**, and then click **Next**.

3. In **Select destination server**, enable **Select a server from the server pool**, and in **Server Pool**, select the server for which you want to enable the Wireless LAN Service, and then click **Next**.
4. In **Select server roles**, click **Next**.
5. In **Select Server features**, in **Features**, select **Wireless LAN Service**, and then click **Next**.

Warning

The computer must be restarted to complete the installation of the Wireless LAN Service. Regardless of whether you select the option to automatically restart the computer after adding the Wireless LAN Service, or you manually restart the computer after adding the wireless LAN service, after the computer restarts, the Wireless LAN Service (WLAN AutoConfig) does not automatically start. You can verify this behavior as follows:

1. In Network Connections, the Wireless Adapter (named Wi-Fi by default), is disabled.
2. In the Services console [Control Panel\All Control Panel Items\Administrative Tools\Services], the WLAN AutoConfig service startup type is Automatic, however the status does not indicate the service is running.
3. In the Notification Area, Clicking on the Network Connections icon opens a list of networks, but the heading **Wi-Fi** is not listed, and no wireless network are detected.

   You can start the WLAN AutoConfig service, and resolve the service startup configuration in one of two ways:

- Manually restart the computer again.
- In the Services console [Control Panel\All Control Panel Items\Administrative Tools\Services], double click WLAN AutoConfig, click Start, and then click OK.

Managing the Wireless Network (IEEE 802.11) Policies

In Windows Server® 2012, the Wireless Network (IEEE 802.11) Policies extension of Group Policy contains two policy nodes. You can use the New Wireless Network Policy to configure wireless settings on computers running Windows® 8, Windows® 7, and Windows Vista®. You can use the Wireless Network Policy to configure wireless settings on computers running Windows®°XP°with°SP2, Windows°XP°with°SP1, or Windows°Server®°2003.

The following resources provide details about configuring 802.1X authenticated wireless access using the Wireless Network (IEEE 802.11) Policies extension of Group Policy.

- Information about the New Wireless Network Policy is provided in the topic, Managing the New Wireless Network (IEEE 802.11) Policies Settings. This topic provides prescriptive guidance for configuring 802.1X authenticated wireless access based on one of the authentication methods you have selected when configuring the New Wireless Network Policy, as well as configuration details about the individual settings.

The Advanced tab in the New Wireless Network Policies contains many of the same settings as are found on the Advanced tab in the Wired Network (IEEE 802.3) Policies. The settings for both the New Wireless Network Policies and the Wired Network (IEEE 802.3) Policies are all documented in the topic: Advanced Security Settings for Wired and Wireless Network Policies.

- Information about the Wireless Network Policy to configure wireless settings on computers running Windows°XP is provided in the topic, Managing the Windows XP-based Wireless Network Policies. This topic provides prescriptive guidance for configuring 802.1X authenticated wireless access based on one of the authentication methods you have selected when configuring the New Wireless Network Policy, as well as configuration details about the individual settings.
- Within the Wireless Network (IEEE 802.11) Policies there are settings to configure Extensible Authentication Protocol (EAP).

  The EAP is an architectural framework that provides extensibility for authentication methods for commonly used protected network access technologies, such as IEEE 802.1X-based wireless access, IEEE 802.1X-based wired access, and Point-to-Point Protocol (PPP) connections such as Virtual Private Networking (VPN). EAP is not an authentication method like MS-CHAP v2, but rather a framework on the access client and authentication server that allows networking vendors to develop and easily install new authentication methods known as EAP methods. For EAP setting details, see Extensible Authentication Protocol (EAP) Settings for Network Access

Managing the New Wireless Network (IEEE 802.11) Policies Settings

Wireless network settings in the Wireless Network (IEEE 802.11) Policies Group Policy extension include global wireless settings, the list of preferred networks, Wired Equivalent Privacy (WEP) and Wi-Fi Protected Access (WPA) settings, and Institute of Electrical and Electronics Engineers (IEEE) 802.1X settings.. You can use the New Wireless Network Policy to configure wireless settings on computers running Windows® 8, Windows® 7, and Windows Vista®.

There are two main sections in this topic:

1. Configuring Wireless Network (IEEE 802.11) Policies, provides prescriptive guidance for configuring 802.1X authenticated wireless access based on one of the authentication methods you have selected when configuring this policy.
2. Wireless Network (IEEE 802.11) Policies settings, provides configuration details about the individual settings in this policy.

Configuring Wireless Network (IEEE 802.11) Policies

This section provides step-by-step details to configure Wireless Network (IEEE 802.11) Policies extension of Group Policy for 802.1X authenticated wireless network access.

In this section:

1. Configure a wireless connection profile for PEAP-MS-CHAP v2
2. Configure a wireless connection profile for PEAP-TLS
3. Configure a wireless connection profile for EAP-TLS

Configure a wireless connection profile for PEAP-MS-CHAP v2

This procedure provides the steps that are required to configure a wireless connection profile for Protected Extensible Authentication Protocol–Microsoft-Challenge Handshake Authentication Protocol (PEAP-MS-CHAP v2) for authentication by using secure passwords.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

To configure a Windows Vista wireless connection profile for PEAP-MS-CHAP v2

1. If you have not already done so, open the Windows Vista Wireless Network (IEEE 802.11) Policies properties page.
2. In **Windows Vista Wireless Network (IEEE 802.11) Policies Properties**, on the **General** tab, select **Use Windows to configure wireless network settings for clients** to specify that WLAN AutoConfig is used to configure wireless network adapter settings.
3. In **Windows Vista Wireless Network (IEEE 802.11) Policies Properties**, on the **General** tab, do one of the following:
   - To add and configure a new profile, click **Add**, and then select **Infrastructure**.
   - To edit an existing profile, select the profile that you want to modify, and then click **Edit**.
4. On the **Connection** tab, in **Profile Name**, if you are adding a new profile, type a name for the profile. If you are editing a profile that is already added, use the existing profile name, or modify the name as needed.
5. In **Network Name(s) (SSID)**, type the service set identifier (SSID) that corresponds to the SSID configured on your wireless APs, and then click **Add**.

   If your deployment uses multiple SSIDs and each wireless AP uses the same wireless security settings, repeat this step to add the SSID for each wireless AP to which you want this profile to apply.

   If your deployment uses multiple SSIDs and the security settings for each SSID do not match, configure a separate profile for each group of SSIDs that use the same security settings. For example, if you have one group of wireless APs configured to use WPA2-Enterprise and AES, and another group of wireless APs to use WPA-Enterprise and

TKIP, configure a profile for each group of wireless APs.

6. If **NEWSSID** is present, select it, and then click **Remove**.
7. If you deployed wireless access points that are configured to suppress the broadcast beacon, select **Connect even if the network is not broadcasting**.

   Note
   Enabling this option can create a security risk because wireless clients will probe for and try to connect to any wireless network. By default, this setting is not enabled.

8. Click the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

      When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, and **Auth Period** are sufficient for most wireless deployments.

   b. To enable Single Sign On, select **Enable Single Sign On for this network**.
   c. The remaining default values in **Single Sign On** are sufficient for typical wireless deployments.
   d. In **Fast Roaming**, if your wireless AP is configured for pre-authentication, select **This network uses pre-authentication**.
9. Click **OK** to return to the **Security** tab. In **Select the security methods for this network**, in **Authentication**, if it is supported by your wireless AP and wireless client network adapters, select **WPA2-Enterprise**. Otherwise, select **WPA-Enterprise**.
10. In **Encryption**, select **AES**, if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **TKIP**.

    Note
    The settings for both **Authentication** and **Encryption** must match the settings configured on your wireless AP. On the **Security** tab, the default settings for **Authentication Mode**, **Max Authentication Failures**, and **Cache user information for subsequent connections to this network** are sufficient for typical wireless deployments.

11. In **Select a network authentication method**, select **Protected EAP (PEAP)**, and then click **Properties**. The **Protected EAP Properties** page opens.
12. In **Protected EAP Properties**, verify that **Validate server certificate** is selected.
13. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computer running Network Policy Server (NPS).

    Note
    This setting limits the root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, clients will trust all root CAs listed in their trusted root certification authority store.

14. Select **Do not prompt user to authorize new servers or trusted certification authorities**. Selecting this setting provides an enhanced user experience and better security.
15. In the **Select Authentication Method** list, select **Secured password (EAP-MS-CHAP v2)**.
16. To enable PEAP Fast Reconnect, select **Enable Fast Reconnect**.
17. If Network Access Protection (NAP) is configured on your network, select **Enable Quarantine checks**. Otherwise, clear this check box.
18. Click **Configure**. In the **EAP MSCHAPv2 Properties** dialog box, verify **Automatically use my Windows logon name and password (and domain if any)** is selected, click **OK**, and then click **OK** to close **Protected EAP Properties**.
19. Click **OK** to close the **Security** tab, and then click **OK** again to close the Windows Vista Wireless Network Policy.

Configure a wireless connection profile for PEAP-TLS

This procedure provides the steps that are required to configure a Protected Extensible Authentication Protocol–Transport Level Security (PEAP-TLS) wireless connection profile for authentication using smart cards or other certificates.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

To configure a Windows Vista PEAP-TLS wireless connection profile

1. If you have not already done so, open the Windows Vista Wireless Network (IEEE 802.11) Policies properties page.
2. In **Windows Vista Wireless Network (IEEE 802.11) Policies Properties**, on the **General** tab, select **Use Windows to configure wireless network settings for clients** to specify that WLAN AutoConfig is used to configure wireless network adapter settings.
3. In **Windows Vista Wireless Network (IEEE 802.11) Policies Properties**, on the **General** tab, do one of the following:
   - To add and configure a new profile, click **Add**, and then select **Infrastructure**.
   - To edit an existing profile, select the profile that you want to modify, and then click **Edit**.
4. On the **Connection** tab, in **Profile Name**, if you are adding a new profile, type a name for the profile. If you are editing a profile that is already added, use the existing profile name, or modify the name as needed.
5. In **Network Name(s) (SSID)**, type the service set identifier (SSID) that corresponds to the SSID configured on your wireless APs, and then click **Add**.

   If your deployment uses multiple SSIDs and each wireless AP uses the same wireless security settings, repeat this step to add the SSID for each wireless AP to which you want this profile to apply.

   If your deployment uses multiple SSIDs and the security settings for each SSID do not

match, configure a separate profile for each group of SSIDs that use the same security settings. For example, if you have one group of wireless APs configured to use WPA2-Enterprise and AES, and another group of wireless APs to use WPA-Enterprise and TKIP, configure a profile for each group of wireless APs.

6. If **NEWSSID** is present, select it, and then click **Remove**.
7. If you deployed wireless access points that are configured to suppress the broadcast beacon, select **Connect even if the network is not broadcasting**.

   Note
   Enabling this option can create a security risk because wireless clients will probe for and try to connect to any wireless network. By default, this setting is not enabled.

8. Click the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

      When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, and **Auth Period** are sufficient for most wireless deployments.

   b. To enable Single Sign On, select **Enable Single Sign On for this network**.
   c. The remaining default values in **Single Sign On** are sufficient for typical wireless deployments.
   d. In **Fast Roaming**, if your wireless AP is configured for pre-authentication, select **This network uses pre-authentication**.
9. Click **OK** to return to the **Security** tab. In **Select the security methods for this network**, in **Authentication**, select **WPA2-Enterprise** if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **WPA-Enterprise**.
10. In **Encryption**, if it is supported by your wireless AP and wireless client network adapters, select **AES**. Otherwise, select **TKIP**.

    Note
    The settings for both **Authentication** and **Encryption** must match the settings configured on your wireless AP. On the **Security** tab, the default settings for **Authentication Mode**, **Max Authentication Failures**, and **Cache user information for subsequent connections to this network** are sufficient for typical wireless deployments.

11. In **Select a network authentication method**, select **Protected EAP (PEAP)**, and then click **Properties**. The **Protected EAP Properties** page opens.
12. In **Protected EAP Properties**, verify that **Validate server certificate** is selected.
13. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued server certificates to your computers running Network Policy Server (NPS).

    Note

This setting limits the root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, clients will trust all root CAs listed in their trusted root certification authority store.

14. Select **Do not prompt user to authorize new servers or trusted certification authorities**. Selecting this setting provides an enhanced user experience and better security.
15. In the **Select Authentication Method** list, select **Smart Card or other certificate**.
16. To enable PEAP Fast Reconnect, select **Enable Fast Reconnect**.
17. If Network Access Protection (NAP) is configured on your network, select **Enable Quarantine checks**. Otherwise, clear this check box.
18. Click **Configure**. In the **Smart Card or other Certificate Properties** dialog box, in **When connecting**, specify one of the following.
    - For deployments that use smart cards, select **Use my smart card**.
    - For other certificate deployments, select **Use a certificate on this computer**.
19. Select **Validate server certificate**.
20. To specify which NPS servers your client computers can use for authentication and authorization, select **connect to these servers**, and then for each NPS server, type the name of each server, exactly as it appears in the **Subject** filed of each NPS server's certificate, separated by a semicolon.
21. In **Trusted Root Certification Authorities**, select the CA that issued certificates to your NPS servers.
22. Click **OK** two times, to close the Protected EAP properties, returning you to the **Security** tab.
23. Click **OK** to close the **Security** tab, and then click **OK** again to close the Windows Vista Wireless Network Policy.

Configure a wireless connection profile for EAP-TLS

This procedure provides the steps that are required to configure an Extensible Authentication Protocol–Transport Level Security (EAP-TLS) wireless profile for authentication by using secure passwords.

Membership in Domain Admins, or equivalent, is the minimum requirement to complete this procedure.

To configure an EAP-TLS wireless profile

1. In **Windows Vista Wireless Network (IEEE 802.11) Policies Properties**, on the **General** tab, click **Add**, and then select **Infrastructure**.

   Note
   For more information about the settings on any tab, press F1 while you are viewing that tab.

2. On the **Connection** tab, do the following:

a. In **Profile Name**, type a name for the EAP-based profile.
b. In **Network Name(s) (SSID)**, type the service set identifier (SSID) that corresponds to the SSID configured on your wireless APs, and then click **Add.**
c. If present, select **NEWSSID**, and then click **Remove**.
d. If your wireless access point is configured to suppress its broadcast beacon, select **Connect even if the network is not broadcasting**.

Note
Enabling this option can create a security risk because wireless clients will probe for and try to connect to any wireless network. By default, this setting is not enabled.

3. Select the **Security** tab, click **Advanced**, and then configure the following:
   a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**.

Note
When the advanced 802.1X settings are enforced, the default values for **Max Eapol-Start Msgs**, **Held period**, **Start Period**, and **Auth Period** are sufficient for most wireless deployments.

   b. In **Single Sign On**, select **Enable Single Sign On for this network**.

Note
The remaining default values in **Single Sign On** are sufficient for most wireless deployments.

   c. In **Fast Roaming**, if your wireless AP is configured for pre-authentication, select **This network uses pre-authentication**.
4. Click **OK** to return to the **Security** tab, and then configure the following:
   a. In **Select the security methods for this network**, for **Authentication**, if it is supported by your wireless AP and wireless client network adapters, select **WPA2-Enterprise**. Otherwise, select **WPA-Enterprise**.
   b. In **Encryption**, if it is supported by your wireless AP and wireless client network adapters, select **AES** (preferred). Otherwise, select **TKIP**.

Note
The settings for both **Authentication** and **Encryption** must match the settings configured on your wireless AP. On the **Security** tab, the default settings for **Authentication Mode**, **Max Authentication Failures**, and **Cache user information for subsequent connections to this network** are sufficient for most wireless deployments.

5. In **Select a network authentication method**, select **Smart Card or other certificate (EAP-TLS)**. On the **Security** tab, click **Properties**, and then configure the following:
   a. In **When connecting**, verify that **Use a certificate on this computer** and **Use simple certificate selection** are selected.

b. Verify that **Validate server certificate** is selected.

In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your computers running Network Policy Server (NPS).

Note
This setting limits the trusted root CAs that clients trust to the selected values. If no trusted root CAs are selected, clients will trust all trusted root CAs listed in their trusted root certification authority store.

6. Click **OK** to close **Smart Card or other Certificate Properties**, and then click **OK** again to close the **EAP Profile**.

Wireless Network (IEEE 802.11) Policies settings

This section provides configuration setting details about the New Wireless Network Policy. You can use the New Wireless Network Policy to configure wireless settings on computers running Windows® 8, Windows® 7, and Windows Vista®.

Wireless Policies for computers running Windows Vista and subsequent releases of Windows

In this section:

1. General - settings
2. Import and Export Wireless Network Profiles
3. Connection - settings
4. New Permissions Entry - settings
5. Security – configuration items
   a. Select the security methods for this network - settings
   b. Select a network authentication method - settings

Note
Information about the network authentication methods is provided in the topic, Advanced Security Settings for Wired and Wireless Network Policies. Clicking this link will open a new Web page. A link is provided in the Additional Resources section of that topic to return you to this page.

Note
Information about the Advanced Security settings is provided in the topic, Advanced Security Settings for Wired and Wireless Network Policies. Clicking this link will open a new Web page. A link is provided in the Additional Resources section of that topic to return you to this page.

6. Additional resources

General - settings

You can use the **General** tab to create and manage wireless network profiles and to define a list of preferred wireless networks, which prioritizes the order in which your domain-member clients attempt to connect. You can also specify whether the WLAN AutoConfig Service is used to configure 802.11 wireless adapters to connect to wireless networks.

Item	Details
Vista Policy Name	Provides a location for you to type a friendly name for the Wireless Network Policy. **Note** When you name the Wireless Network (IEEE 802.11) Policy, the name you specify is displayed as the title of its properties page, and under **Name** in the details pane in the Group Policy Management Console (GPMC) for that policy.
Description	Provides a location for you to type a description for the Wireless Policy. **Note** The description you type is displayed under **Description** in the details pane of the Group Policy Management Console (GPMC) for that policy.
Use Windows WLAN Auto Config service for clients	Specifies that the WLAN AutoConfig Service is used to configure and connect clients running Windows Vista® to the wireless network. Default = enabled
Connect to available networks in the order of profiles listed below	Click the desired profile, and then use the **Move Up** and **Move down** buttons to specify the preferred order for clients to attempt connections. **Note** Profiles for ad-hoc networks cannot be prioritized higher than infrastructure profiles.
Add	Allows you to add a wireless network profile and specify the profile properties.

Item	Details
	When you click **Add**, you will be presented with two options: **Ad-hoc**. Specifies a computer-to-computer network. Access points are not used. **Infrastructure**. Specifies a network that uses one or more wireless access points.
Edit	Select a network under **Profile name** and then click **Edit** to modify the properties of the wireless network profile.
Remove	Select a network under **Profile name** and then click **Remove** to delete the wireless network profile.
Import	Allows you to incorporate an existing wireless network profile, from a location you specify, into the list of available networks.
Export	Allows you to save a configured profile as an Extensible Markup Langage (XML) file.

Import and Export Wireless Network Profiles

Profile import and export are managed by using the following two interfaces. You can use **Import a Profile** to add a wireless network profile, from a location you specify, into the list of available wireless networks. You can use **Save Export Profile** to export any profile listed under **Connect to available networks in the order of profiles listed below** on the **General** tab, and save it to a location you specify.

Open for import a profile (Import Profiles)

Item	Details
Save in	Specifies the location to save a profile.
Name	Lists saved wireless profiles.
File name	Provides a location for you to type a name for the profile.
Save as type	Specifies the file type used to save the profile.

Save export profile as (Export Profiles)

Item	Details
Look in	Specifies the location of a saved profile.
Name	Lists saved profiles. Select the profile you want to export, and then click **Open**.
File name	Provides a location for you to type a new name or modify the existing profile name.
Files of type	Specifies the file type for files that are displayed under **Name**.

Connection - settings

The **Connection** tab for the Wireless Network (IEEE 802.11) Policy allows you to create wireless network profiles and specify the name of each wireless network to which domain-member wireless clients can connect. Each network is saved as a profile. A *profile* is the collection of configuration settings for a wireless networks, saved as an Extensible Markup Language (XML) file.

Item	Details
Profile name	Specifies the friendly name for the wireless network profile.
Network Name (SSID)	Specifies the name of the wireless network. This must match the Service Set Identifier (SSID) configured on the wireless access points for this network.
Connect automatically when this network is in range	Specifies that Windows will automatically attempt to connect to the wireless network when the computer is within broadcast range of any wireless access points that are broadcasting the SSID. **Note** Windows attempts connections to wireless networks in the order they are listed in preferred networks.
Connect to a more preferred network if available	Specifies that clients will attempt to connect to wireless networks in the order specified under **Connect to available networks in the order**

Item	Details
	of profiles listed below on the **General** tab.
Connect even if the network is not broadcasting	Specifies that wireless clients will actively probe for wireless access points with the specified SSID, in case wireless access points are configured to suppress beacon broadcasts.

Network Permissions - settings

You can use Network Permissions to list and configure wireless networks that are not defined on the **General** tab in the **Connect to available networks in the order of profiles listed below** preferred list. You can use these settings to define additional wireless networks and specify whether you want to allow or deny connections by your domain-member wireless clients. Alternatively, you can block the additional wireless networks from being displayed to your domain-member wireless clients. These settings are specific to the wireless networks listed on the **Network Permissions** tab under **Network Name (SSID)**.

Connections to the wireless networks that are listed under **Network Name (SSID)** on the **Network Permissions** tab are only possible if the permission is set to **Allow**. If the permission is set to **Allow**, your domain-member wireless clients first attempt to connect to a preferred network, before attempting to connect to non-preferred networks. However, domain members can actively launch a connection attempt to listed networks that have permissions set to **Allow**.

Item	Details
Network Name (SSID)	Lists wireless networks, for which you want to allow or deny permissions, but that are not defined on the **General** tab in **Connect to available networks in the order of profiles listed below**. Default = disabled
Add	Opens the **New Permission Entry** dialog box, for adding a new wireless network to the permission list. When you click **Add**, you will be presented with two options: • **Ad-hoc**. Specifies a computer-to-computer network. Access points are not used. • **Infrastructure**. Specifies a network that uses one or more wireless access points.
Remove	Deletes the selected wireless network from the

Item	Details
	permission list.
Prevent connections to ad-hoc networks	Specifies that domain-member wireless clients cannot form a new ad-hoc network or connect to any ad-hoc networks in the permission list. Default = not enabled
Prevent connections to infrastructure networks	Specifies that domain-member wireless clients cannot connect to any infrastructure networks in the permission list. Default = not enabled
Allow user to view denied networks	Specifies whether domain-member wireless clients can view wireless networks in the permission list that have permissions set to **Deny**. Default = not enabled
Allow everyone to create all user profiles	Specifies that any user to which Wireless Network (IEEE 802.11) Policy applies can create all user wireless profiles on the computer. When this setting is not selected, only users in the Network Administrators group or Network Operators group can create all user wireless profiles on the computer. *All user profiles* are wireless profiles that are used to connect to a specific wireless network and that can be started by any user with an account on the computer. Every user with an account on the computer is contained within the computer-wide access control list (ACL) named **Everyone**. Based on the Everyone ACL group, network administrators can use this setting to limit access to the controls that enable users to create all user profiles. Default = enabled

New Permissions Entry - settings

You can use the Wireless Network (IEEE 802.11) Policy New Permission page to add new wireless networks to the permission list on the **Networks Permissions** tab. You can use New

Permissions to specify, by Service Set Identifier (SSID), which wireless networks your wireless domain members are allowed to connect to, and which are denied.

Items	Details
Network Name (SSID)	Provides a location for you to type the name for the wireless network for which you want to set permissions.
Network Type	Specifies whether the network is infrastructure (uses a wireless access point) or ad-hoc (computer-to-computer).
Permission	Specifies whether to permit or deny connections to the selected network.

Security – configuration items

Security configuration items include all of the settings on the **Security** tab. These settings specify whether to perform 802.1X authentication for connecting clients, which network security authentication method to use, and which network authentication method to use.

Select the security methods for this network settings pertain to the authentication and security that the client uses to exchange data with the wireless access point.

Select a network authentication method settings pertain to the EAP method that is used for 802.1X authenticated network access.

Select the security methods for this network - settings

Item	Details
Authentication	Specifies the security authentication method to use when the wireless client associates with the wireless access point. • **Open** • **Shared** • **WPA-Enterprise** • **WPA-Personal** • **WPA2-Enterprise** • **WPA2-Personal** • **Open with 802.1X** **Note** Network Authentication Methods and

Item	Details
	advanced settings are only exposed by selecting **WPA-Enterprise**, **WPA2-Enterprise**, and **Open with 802.1X** for the network security authentication method. Default = Defaults to the most secure setting supported by the wireless hardware and drivers.
Encryption	Specifies the security encryption to use for the selected network security authentication method. If the **Authentication** is set to **WPA-Enterprise**, **WPA-Personal**, **WPA2-Enterprise**, or **WPA2-Personal**, the encryption options are **AES** or **TKIP**. If **Authentication** is set to **Open**, **Shared**, or **Open with 802.1X**, then encryption options are either **WEP** or **Disabled**.

Select a network authentication method - settings

Item	Details
Select a network authentication method	Specifies the network authentication method that connecting wireless clients use. • **Smart Card or other certificate** • **Protected EAP (PEAP)** Default = Protected EAP (PEAP)
Properties	Opens the properties page of the selected network authentication method. For setting information for network authentication methods, see: Extensible Authentication Protocol (EAP) Settings for Network Access. Clicking this link will open a new Web page. A link is provided in the Additional Resources section of that topic to return you to this page.
Authentication Mode	Specifies how credentials are used for network

Item	Details
	authentication. The following options are available. • **User re-authentication**. An 802.1X-compliant device always uses security credentials based on the computer's current state. Authentication is performed by using the computer credentials when no users are logged on to the computer. When a user logs on to the computer, authentication is always performed by using the user credentials. This is the recommended setting. • **Computer only**. Authentication is always performed using only the computer credentials. • **User authentication**. Specifies that when users are not logged on to the computer, authentication is performed by using the computer credentials. After a user logs on to the computer, authentication is still based on the computer credentials. Authentication is performed by using the user credentials if the user travels to a new wireless access point. • **Guest authentication**. Allows connections to the network which are regulated by the restrictions and permissions set for the Guest user account. Default = User re-authentication
Max Authentication Failures	Specifies the maximum number of failed authentication attempts that can occur with a specific set of credentials before notification is displayed to indicate that authentication has failed. Default = 3
Clear credentials after logoff	Specifies that when the user logs off, the user credential data is removed from the registry. The result is that at the next user log on, they will be prompted for their credentials (such as user name and password).

Item	Details
	Default = Not enabled
Authenticate as guest when user or computer information is unavailable	Specifies that client connection requests that cannot meet computer or user authentication requirements can connect to the network by using permissions configured for the Guest account. Default = Not enabled
Advanced	Provides access to advanced security settings. **Note** Information about the Advanced Security settings is provided in the topic, Advanced Security Settings for Wired and Wireless Network Policies. Clicking this link will open a new Web page. A link is provided in the Additional Resources section of that topic to return you to this page.

Additional resources

For additional information about the Advanced Security settings, see Advanced Security Settings for Wired and Wireless Network Policies, in the Microsoft technical library on TechNet.

For additional information about network authentication methods, see Extensible Authentication Protocol (EAP) Settings for Network Access, in the Microsoft technical library on TechNet.

Advanced Security Settings for Wired and Wireless Network Policies

The Wireless Network (IEEE 802.11) Policies Advanced settings and Wired Network (IEEE 802.11) Policies specify settings associated with 802.1X authentication requests. For the Wireless Network (IEEE 802.11) Policies, the Advanced settings are only exposed by enabling Wi-Fi Protected Access 2 (WPA2)-Enterprise, WPA-Enterprise, or Open with 802.1X as the network authentication setting on the **Security** tab in the Wireless Network (IEEE 802.11) Policy.

Advanced security settings for the Wireless Network Policies and Wired Network Policies

IEEE 802.1X - settings

IEEE 802.1X settings specify the behavior of wireless client 802.1X authentication requests.

Item	Details
Enforce advanced 802.1x settings	Specifies that the settings listed are enforced as configured: • **Max Eapol-Start Msgs** • **Held Period (seconds)** • **Start Period (seconds)** • **Auth Period (seconds)**
Max Eapol-Start Msgs	EAPOL is the Extensible Authentication Protocol (EAP) over local area network (LAN) protocol. If no response is received to the original EAPOL-Start message, this setting specifies the maximum number of subsequent EAPOL-Start messages sent. Default = 3
Held Period (seconds)	After a client has received notification of authentication failure, this setting specifies the number of seconds an authenticating client waits before it performs another 802.1X authentication request. Default = 60
Start Period	If no response is received to the original EAPOL-Start message, this setting specifies the number of seconds between the retransmission of subsequent EAPOL-Start messages. Default = 5
Auth Period	After end-to-end 802.1X authentication is initiated, this setting specifies the number of seconds authenticating clients must wait before retransmitting any 802.1X requests.

Item	Details
	Default = 30

Single Sign On - settings

In Windows Server® 2008 and Windows Vista, Single Sign On performs 802.1X authentication based on the network security configuration during the user logon process. This feature enables scenarios — such Group Policy updates, implementation login scripts, and joining of wireless clients to domains — that require network connectivity prior to user logon.

You can use Group Policy settings to configure Single Sign On profiles for your wireless client computers. When a Single Sign On profile is configured, 802.1X authentication is conducted prior to computer logon to the domain; users are only prompted for credential information if needed.

Item	Details
Enable Single Sign On for this network	Specifies that Single Sign On is activated for the wireless network profile for this network. Default = Not enabled
Perform immediately before User Logon	Specifies that Single Sign On performs 802.1X authentication before user logon completes. Default = Not enabled
Perform immediately after User Logon	Specifies that Single Sign On performs 802.1x authentication immediately upon successful user logon. Default = Not enabled
Max delay for connectivity (Seconds)	Specifies the maximum time, in seconds, in which 802.1X authentication must complete and authorize network access. This setting allows the network administrator to define the maximum length of time a user waits at the logon screen. Default = 10
This network uses different VLAN for authentication with machine and user credentials	Specifies that wireless computers are placed on one virtual local area network (VLAN) at startup, and then — based on user permissions — transitions to a different VLAN network after the user logs on to the computer. This setting is used in scenarios where it is desirable to separate traffic by using VLANs.

Item	Details
	For example, one VLAN, "VLAN-a," allows access only to authenticated computers, typically with a restricted set of assets. A second VLAN, "VLAN-b," provides authenticated and authorized users with access to a broader set of assets, such as e-mail, build servers, or the intranet. Default = Not enabled

Fast Roaming settings

Fast Roaming is a feature of WPA2 that uses pre-authentication and Pairwise Master Key (PMK) Caching to enable wireless clients to roam more quickly among wireless access points (APs).

Fast Roaming settings are only exposed by enabling **WPA2-Enterprises** on the Security tab in the Windows Vista Wireless Network (IEEE 802.11) Policy.

Item	Details
Enable Pairwise Master Key (PMK) Caching	Specifies that Pairwise Master Key (PMK) Caching is used for WPA2 Fast Roaming. Default = Enabled
PMK time to Live (Minutes)	Specifies the duration, in minutes, the PMK is held in cache. Default = 720
Number of Entries in PMK Cache	Specifies the maximum number of PMK entries that are stored in cache. Default = 128
This network uses pre-authentication	Specifies that pre-authentication is used for WPA2 Fast Roaming. Pre-authentication enables WPA2 wireless clients that are connected to one wireless AP to perform 802.1X authentication with other wireless APs within its range. Pre-authentication stores the PMK and its associated information in the PMK cache. When the wireless client connects to a wireless AP with which it has pre-authenticated, it uses the cached PMK information to reduce the time

Item	Details
	required to authenticate and connect. **Note** WPA2 client pre-authentication is only possible with wireless access points that broadcast pre-authentication capability in Beacon and Probe Response messages. Default = Enabled
Maximum pre-authentication attempts	Specifies the maximum allowed pre-authentication attempts. This setting is available only when **This network uses pre-authentication** is selected. Default = 3

Item	Details
Perform cryptography in FIPS 140-2-certified mode	Specifies that wireless transmissions adhere to the Federal Information Processing Standard (FIPS) 140-2 mode for cryptography. FIPS 140-2 is a U.S. government computer security standard that is used to certify cryptographic modules. Default = not enabled

Additional resources

For additional information about wireless settings in Group Policy, see Managing the New Wireless Network (IEEE 802.11) Policies Settings

For additional information about wireless settings in Group Policy, see Managing the New Wired Network (IEEE 802.3) Policies Settings

Managing the Windows XP-based Wireless Network Policies

Wireless network settings in the Wireless Network (IEEE 802.11) Policies Group Policy extension include global wireless settings, the list of preferred networks, Wired Equivalent Privacy (WEP)

and Wi-Fi Protected Access (WPA) settings, and Institute of Electrical and Electronics Engineers (IEEE) 802.1X settings. These settings encompass all of the items on the Association and Authentication tabs in the Wireless Network (IEEE 802.11) Policies to configure wireless network access for client running Windows®°XP with Service°Pack°(SP)°2, Windows°XP with SP1, or Windows°Server®°2003.

This topic contains the following sections:

1. General - settings
2. 802.1X - settings
3. Network Properties - settings
4. Preferred Networks - settings

General - settings

Item	Details
XP Policy Name	Provides a location for you to type a friendly name for the wireless network policy. **Note** When you name the Wireless Network (IEEE 802.11) Policy, the name you specify is displayed as the name of the associated properties page, and under **XP Policy Name** in the details pane in the Group Policy Management Console (GPMC) for that policy.
Description	Provides a location for you to type a description for the wireless policy. **Note** The description you type for is displayed under **Description** in the details pane of the Group Policy Management Console (GPMC) for that policy.
Networks to access	Specifies the types of wireless networks with which the wireless client is allowed to create connections: • **Any available network (access point preferred)** • **Access point (infrastructure) networks**

Item	Details
	only • **Computer-to-computer (ad hoc) networks only**
Use Windows WLAN Auto Config service for clients	Enables the Wireless Zero Configuration service. When enabled, the Wireless Zero Configuration service is used to configure and connect clients running Windows°XP to wireless networks. **Note** Clients running Windows°XP use Wireless Zero Configuration while Clients running Windows Vista use WLAN AutoConfig Service.
Automatically connect to non-preferred networks	Specifies that wireless clients can attempt to automatically connect to wireless networks that are not configured as preferred networks. If selected, your domain-member wireless clients can connect to a non-preferred network after first attempting to connect to preferred networks.

802.1X - settings

IEEE 802.1X settings specify the behavior of wireless clients when connecting to an infrastructure wireless network that is configured for 802.1X authentication. Settings include authentication methods, (such as Extensible Authentication Protocol (EAP), authentication mode, and 802.1X-specific parameters.

Item	Details
Enable network access control using IEEE 802.1X	Specifies whether you want to use Institute of Electrical and Electronics Engineers (IEEE) 802.1X to perform authentication for your wireless network. If you clear this check box, all of the other settings on this tab become unavailable.
EAP Type	Specifies the network authentication method that connecting wireless clients use.

Item	Details
	• **Smart Card or other certificate** (EAP-TLS) • **Protected EAP (PEAP)** Default = Protected EAP (PEAP)
Settings	Opens the properties page of the selected network authentication method. For setting information for network authentication methods, see: **Authentication Methods Properties [Preliminary]**
Eapol-Start message	EAPOL is the Extensible Authentication Protocol (EAP) over local area network (LAN) protocol. This setting specifies the transmission behavior of the EAPOL-Start message when authenticating. You can select from the following: • **Do not transmit**. Specifies that EAPOL-Start messages are not sent. • **Transmit Determines**. when to send EAPOL-Start messages and, if needed, sends an EAPOL-Start message. • **Transmit per IEEE 802.1x**. Sends an EAPOL-Start message upon association to initiate the 802.1X authentication process.
Authentication Mode	Specifies how credentials are used for network authentication. The following options are available. • **User re-authentication**. An 802.1X-compliant device always uses security credentials based on the computer's current state. Authentication is performed using the computer credentials when no users are logged on to the computer. When a user logs on to the computer, authentication is always performed using the user credentials. This is the recommended setting. • **Computer only**. Authentication is always

Item	Details
	performed using only the computer credentials. • **User authentication**. Specifies that when users are not logged on to the computer, authentication is performed by using the computer credentials. After a user logs on to the computer, authentication is still based on the computer credentials. Authentication is preformed by using the user credentials if the user travels to a new wireless access point. • **Guest authentication**. Allows connections to the network which are regulated by the restrictions and permissions set for the Guest user account. Default = User re-authentication
Authenticate as computer when computer information is available	Specifies whether the computer will attempt to authenticate using computer credentials when the user is not logged on. Computer credentials are typically a computer certificate. Default = Enabled
Authenticate as guest when user or computer information is unavailable	Specifies that client connection requests that cannot meet computer or user authentication requirements can connect to the network by using permissions configured for the Guest account. Default = Not enabled
Max Eapol-Start Msgs	If no response is received to the original EAPOL-Start message, this setting specifies the maximum number of subsequent EAPOL-Start messages sent. Default = 3
Held Period	After a client has received notification of authentication failure, this setting specifies the number of seconds an authenticating client waits before it performs another 802.1X authentication request. Default = 60

Item	Details
Start Period	If no response is received to the original EAPOL-Start message, this setting specifies the number of seconds between the retransmission of subsequent EAPOL-Start messages. Default = 5
Auth Period	After end-to-end 802.1X authentication is initiated, this setting specifies the number of seconds authenticating clients must wait before retransmitting any 802.1X requests. Default = 30

Network Properties - settings

You can use the Wireless Network (IEEE 802.11) Policy for Windows°XP clients to specify the name of the wireless network, network security authentication, such as Wi-Fi Protected Access 2 (**WPA2**), and Fast Roaming features.

Item	Details
Network Name (SSID)	Provides a location for you to type the wireless LAN network name, also known as the Service Set Identifier (SSID).
Description	Provides a location for you to type a description of the wireless LAN network.
Connect even if network is not broadcasting	Specifies that wireless clients will actively probe for wireless access points with the specified SSID, if wireless access points are configured to suppress beacon broadcasts.
Authentication	Specifies the security authentication method that is used between the wireless access point and the wireless client. • **Open** • **Shared** • **WPA-Enterprise** • **WPA-Personal** • **WPA2-Enterprise**

Item	Details
	• **WPA2-Personal** • **Open with 802.1X** **Note** Network Authentication Methods and advanced settings are only exposed by selecting **WPA-Enterprise**, **WPA2-Enterprise**, and **Open with 802.1X** for the network security authentication method. Default = Defaults to the most secure setting supported by the wireless hardware and drivers.
Encryption	Specifies the security encryption to use for the selected network security authentication method. If Authentication is set to **WPA-Enterprise**, **WPA-Personal**, **WPA2-Enterprise**, or **WPA2-Personal**, the encryption options are **AES** or **TKIP**. If Authentication is set to **Open**, **Shared**, or **Open with 802.1X**, then encryption options are **WEP** or **Disabled**.
Enable Pairwise Master Key (PMK) Caching	Specifies that Pairwise Master Key (PMK) Caching is used for WPA2 Fast Roaming. **Note** Fast Roaming is a feature of Wi-Fi Protected Access 2 (WPA2) that uses pre-authentication and PMK Caching to enable wireless clients to roam more quickly among wireless access points (APs). Fast Roaming settings are only exposed if WPA2 is selected as the network security method.
PMK Time to Live (Minutes)	Specifies the duration, in minutes, the PMK is held in cache. Default = 720
Number of Entries in PMK Cache	Specifies the maximum number of PMK entries

Item	Details
	that are stored in cache. Default = 128
This network uses pre-authentication	Specifies that pre-authentication is used for WPA2 Fast Roaming. Pre-authentication enables WPA2 wireless clients that are connected to one wireless AP to perform 802.1X authentication with other wireless APs within its range. Pre-authentication stores the PMK and its associated information in the PMK cache. When the wireless client connects to a wireless AP with which it has pre-authenticated, it uses the cached PMK information to reduce the time required to authenticate and connect. **Note** WPA2 client pre-authentication is only possible with wireless access points that broadcast pre-authentication capability in Beacon and Probe Response messages. Default = Enabled
Maximum Pre-authentication attempts	Specifies the maximum allowed pre-authentication attempts. This setting is available only when **This network uses pre-authentication** is selected. Default = 3

Preferred Networks - settings

You can use the settings on the Preferred Networks tab to add wireless networks, by name, to a prioritized list. That list specifies the order in which your domain-member wireless clients will attempt to connect to wireless networks.

If a client cannot connect to a preferred wireless network, and the wireless network is configured to prevent automatic connection to wireless networks that are not in the preferred list (the default), then the Wireless Zero Configuration service places the wireless network adapter in infrastructure mode using a random wireless network name. At this point, the wireless adapter is not connected to any wireless networks and continues to scan for preferred wireless networks.

Item	Details
Profile name	Displays the prioritized list of preferred wireless networks. General profile information, such as Service Set Identifier (SSID), IEEE 802.1X status **Enabled** or **Disabled**, and encryption type, is displayed.
Add	Allows you to add a wireless network profile, and specify the profile properties. When you click **Add**, you will be presented with two options: • **Ad-hoc**. Specifies a computer-to-computer network. Access points are not used. • **Infrastructure**. Specifies a network that uses one or more wireless access points.
Edit	Allows you to change the configuration settings for the selected preferred wireless network profile.
Remove	Deletes the selected network from the list of preferred networks.
Up arrow	Moves a selected network to a higher connection priority in the list of preferred networks. **Note** Profiles for ad-hoc networks cannot be prioritized higher than infrastructure profiles.
Down arrow	Moves a selected network to a lower connection priority in the list of preferred networks.

New Wireless Connection Processes

This section provides information about how computers running Windows 8 connect to different types of networks based on the authentication methods that are allowed for the network connection.

- Connections to Organization Networks that Require RADIUS Server Authentication by Wireless Clients

- Connections to Organization Networks with Multiple RADIUS Servers
- Connections to Networks Whose RADIUS Server Certificate is Issued by a Public Certification Authority
- Connections to Public Wi-Fi Hotspots

Connections to Organization Networks that Require RADIUS Server Authentication by Wireless Clients

When you use a laptop or other wireless device to connect to an organization network and the authentication method deployed for the connection requires that the wireless client authenticates the RADIUS server by using the server certificate, events differ based upon whether the wireless device is joined to the domain.

Domain member computers and other devices automatically receive the private certification authority (CA) certificate when they log on to the network. This certificate is kept in the Trusted Root Certification Authorities certificate store on the domain member wireless client. During the authentication process for subsequent network connection attempts, the wireless client receives the server certificate from the Remote Authentication Dial-In User Service (RADIUS) server, and then checks the certificates in the Trusted Root Certification Authorities certificate store to determine whether the CA is trusted. When the wireless client finds that the server certificate is issued by a trusted CA, the wireless client successfully authenticates the RADIUS server, and the rest of the authentication process can continue.

When your wireless client is not a domain member, however, the CA certificate is not present in the Trusted Root Certification Authorities certificate store, so the wireless client cannot automatically trust the CA that issued the server certificate.

For non-domain joined wireless clients that are running Windows® 8, when you attempt to connect to an organization network whose authentication methods require RADIUS server authentication by the client, the dialog box in the illustration below automatically opens to allow you to determine whether you trust the server certificate and want to connect to the network.

If you trust the server certificate and choose to connect to the network, the certificate information is stored on the local computer so that subsequent connection attempts to the same wireless network and RADIUS server happen seamlessly and automatically.

For information about connecting to networks with multiple RADIUS servers, see Connections to Organization Networks with Multiple RADIUS Servers.

Connections to Organization Networks with Multiple RADIUS Servers

Wireless devices that are used on a network where more than one RADIUS server is deployed can encounter authentication issues, such as failure to automatically authenticate the RADIUS server, in circumstances where the server certificates on the RADIUS servers are not all issued by the same private certification authority (CA). Even if the wireless device successfully verifies

the identity of a RADIUS server and connects to the network, moving to a new location or encountering a situation where a different RADIUS server processes a reconnection attempt can result in failure to automatically authenticate the server and establish a wireless connection. In this circumstance, you must manually connect to the network and accept the server certificate for the connection.

To understand how to deploy multiple RADIUS servers without causing these authentication issues, it is important to understand the server certificate checks that are performed by Windows® 8 devices during the authentication process. The client performs the following checks on the RADIUS server certificate:

- The client compares the hashes of the issuing CA in the RADIUS server certificate with the hashes of the issuing CA of a certificate that you have previously chosen to trust. The issuing CA hashes in both certificates must match in order for authentication to proceed.
- The client verifies that the server certificate is issued by a private CA, which is determined by the presence of the RADIUS Enhanced Key Usage (EKU) extension in the certificate.

If the server certificate passes both of these checks, then the authentication process continues; if the RADIUS server authenticates and authorizes the connection, the wireless device connects to the network seamlessly.

If the server certificate does not pass both of these checks, Windows silently drops the connection and you must manually connect to the network.

Deployment requirements for multiple RADIUS servers

In order to ensure that your RADIUS server deployment leverages the server certificate checks that are performed by wireless clients that are running Windows 8, and to provide wireless users with a seamless, automatic wireless connection experience, you must deploy server certificates with the following requirements.

- You must deploy a private CA rather than obtain server certificates from a third party public CA. In addition, the certificate template that you use to issue the certificates must contain the RADIUS EKU extension. This extension is **id-kp-eapOverLAN** and the object identifier (OID) for this EKU is 1.3.6.1.5.5.7.3.14. This EKU extension can only be configured on a private CA and is used by Windows 8 to determine whether a private CA issued the certificate.
- All server certificates that are issued to all RADIUS servers must chain to the same private trusted root CA.

Important

If you deploy server certificates using Active Directory Certificate Services (AD CS) using the RAS and IAS Servers certificate template, you must add the **id-kp-eapOverLAN** EKU extension to the RAS and IAS Servers certificate template; it is not included in the template by default.

If your server certificate deployment does not meet these requirements and you attempt to connect to the network, you must manually accept the server certificate when prompted by Windows. After you have accepted a server certificate, it is stored on the local computer and Windows attempts to use the certificate during subsequent connection attempts. If you have

multiple RADIUS servers with server certificates that do not match the requirements, you must manually accept each certificate as you move through the network and encounter different authenticating servers.

Connections to Networks Whose RADIUS Server Certificate is Issued by a Public Certification Authority

When you encounter a wireless network whose RADIUS server certificate is issued by a public certification authority (CA), and whose configuration requires use of an authentication method in which the client computer must verify the identity of the RADIUS server, the client computer checks the certificates that are installed in the Trusted Root Certification Authorities certificate store. If the public CA certificate exists in the certificate store, the client computer trusts the RADIUS server and authentication of the server is successful.

When connecting to this type of network for the first time, if the public CA certificate is not in the Trusted Root Certification Authorities certificate store, you must manually accept the certificate.

During subsequent auto-connects, Windows compares the new certificate with the server certificate that was received during the first connection attempt. If the certificate is the same, then trust is established and Windows seamlessly authenticates the RADIUS server. If the certificate is not the same, the connection is dropped silently and you must manually connect to the network.

Connections to Public Wi-Fi Hotspots

In Windows® 8, there are many improvements to 802.1X authenticated wireless access, especially for circumstances where the computer or other device is connecting for the first time to a secure network that does not use captive portals.

Making 802.1X wireless connections on computers that are running Windows® 7 and earlier operating systems is difficult because the default settings provide no automatic way for authentication to occur when Extensible Authentication Protocol (EAP)-based authentication methods are required for the connection.

For example, when you try to connect an organization-managed, domain joined Windows 7 computer to a public wireless hotspot, Windows 7 automatically provides domain credentials as proof of identity to the public hotspot's wireless access point (WAP). Because the public hotspot is not part of the organization’s domain, it cannot authenticate the client computer, and authentication fails. In this circumstance, you must determine the authentication protocol that the

WAP requires, and then follow a complex set of steps to select and configure the protocol in the Wireless profile on your computer.

Connections to public Wi-Fi hotspots

When you encounter a wireless access signal using a computer or other device that is running Windows 8, Windows informs you that a wireless network is available. If you select a network and click Connect, the following dialog box opens.

This dialog box provides the ability to log on to the network using different credential types. Windows presents credential options that are based on the types of credentials that it finds on the computer.

The following types of credentials are supported:

- User name and password. In addition to standard credentials that you type in the fields provided, the following types of user name and password are supported:
  - One Time Password (OTP) user name and token codes are accepted in the form of user name and password. In this circumstance, the password consists of your personal identification number (PIN) and token.
  - **Use Windows Logon Credentials**. When you select this check-box, Windows 8 automatically fills **UserName** and **Password** with your Windows logon credentials.
- Registry certificates that are stored in the certificate store on the local computer. Smart Card certificates are not supported.
- Global System for Mobile Communications (GSM) subscriber identity module (SIM) cards.

By default, the ability to type user name and password is provided because password-based authentication is the most popular authentication method that is used by Wi-Fi hotspots and other organizations.

If Windows 8 finds other forms of credentials on the computer, Windows provides a link that you can click to use these credentials; however this link is not displayed if the operating system does not find other types of credentials on the computer or device, and you must type a user name and password or Windows logon credentials for authentication.

In the illustration above, Windows 8 has discovered alternate credential types on the local computer - both a certificate and SIM card are found - so the **Other ways to connect** link is displayed.

When you click the link, a new dialog box opens, and you are provided options to authenticate the wireless connection by using a certificate or a SIM card.

Notes

- If Windows 8 finds only a registry certificate on the computer, the link is named **Connect using a certificate**.
- If Windows 8 finds only a SIM card on the computer, the link is named **Connect using a SIM card**.

If you choose one of these alternate credential types, you are prompted to enter additional credential information through an appropriate dialog box, if required, such as a certificate password or a SIM card PIN.

The credentials that you provide are then used for authentication to the wireless network.

Selecting the Authentication Protocol

When you select the type of credential that you want to use to authenticate to the wireless network, Windows 8 automatically selects the appropriate authentication protocol to use for the connection. The following steps provide details on how this process works.

1. Windows checks all installed EAP methods in the operating system, including the built-in EAP methods and third party methods, and creates a list, called a *short list*, of authentication methods that match the credential type that you selected.

- PEAP with Tunneled Transport Layer Security (TTLS), or PEAP-TTLS, is mapped to password-based methods, as are third party OTP-based methods. The authentication methods EAP with Message-Digest Algorithm (MD5), or EAP-MD5, as well as EAP with Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), or EAP-MS-CHAP v2, are always ignored because, for security purposes, the Windows 8 wireless client does not use them.
- Microsoft built-in methods are always preferred over third party implementations. For example, if a third party implementation of PEAP is found Windows will still pick only the built-in Microsoft implementation when short-listing the protocols.
- In case of SIM-based authentication Windows short lists only those protocols that are supported by your SIM card.

2. Windows uses this protocol list and negotiates the final protocol with the Remote Authentication Dial In User Service (RADIUS) server that is processing the connection request and performing authentication and authorization, such as a server that is running Windows Server 2012 with Network Policy Server (NPS) installed and configured as a RADIUS server.
   - Based on credential type, Windows proposes the protocols to the RADIUS server in the order provided in the table below.
   - For PEAP and TTLS only EAP-MS-CHAP v2 and Password Authentication Protocol (PAP) are negotiated as inner methods.

Protocol order

Credential type	Order
Password	<EAP-PEAP> - <EAP-TTLS> - <other password-based and OTP-based protocols third party methods in the order in which they were short-listed>
Certificate	<EAP-TLS> - <other certificate-based third party methods in the order in which they were short-listed>
SIM	<EAP-AKA'> - <EAP-AKA> - <EAP-SIM>

For more information, see Extensible Authentication Protocol (EAP) for Network Access Overview at http://technet.microsoft.com/en-us/library/hh945105.

Notes

AKA is an acronym for UMTS Authentication and Key Agreement, which is described in the Internet Engineering Task Force (IETF) request for comments 4187, Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA).

AKA' is the EAP AKA prime protocol, which is described in the IETF request for comments 5448, Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA').

Caveats for third party authentication methods

A third party method is short-listed by Windows for authentication protocol negotiation only if it provides a mechanism to generate the method configuration blob. Following are the options that are available to generate this blob.

- The third party method can implement a new application programming interface (API) exposed by EAPHost that generates the user blob and the configuration blob of the method when passed the credential type and the actual credentials. Details of the API can be found on MSDN at http://msdn.microsoft.com/en-us/library/windows/desktop/hh706635(v=vs.85).aspx.
- If the method is password-based then Windows can still use the method's X2B API to generate the configuration blob. If the method has implemented the QueryCredential API, then a user blob can also be generated.

After an EAP method is negotiated with the server, Windows uses the method's blobs for authentication. The configuration blob is needed at that time so that the method doesn't have to launch its own dialog box to collect user credentials again.

BranchCache Overview

This topic, which is intended for Information Technology (IT) professionals, provides overview information about BranchCache, including BranchCache modes, features, capabilities, and the BranchCache functionality that is available in different operating systems.

Notes

In addition to this topic, the following BranchCache documentation is available.

- BranchCache Learning Roadmap
- What's New in BranchCache
- BranchCache Network Shell and Windows PowerShell Commands
- BranchCache Deployment Guide
- Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode
- BranchCache Frequently Asked Questions
- BranchCache for Windows Server 2008 R2

Who will be interested in BranchCache?

If you are a system administrator, network or storage solution architect, or other IT professional, BranchCache might interest you under the following circumstances:

- You design or support IT infrastructure for an organization that has two or more physical locations and a wide area network (WAN) connection from the branch offices to the main office.
- You design or support IT infrastructure for an organization that has deployed cloud technologies, and a WAN connection is used by workers to access data and applications at remote locations.
- You want to optimize WAN bandwidth usage by reducing the amount of network traffic between branch offices and the main office.
- You have deployed or are planning on deploying content servers at your main office that match the configurations that are described in this topic.
- The client computers in your branch offices are running either Windows® 8 or Windows® 7.

This topic includes the following sections:

- What is BranchCache?
- BranchCache modes
- BranchCache-enabled content servers
- BranchCache and the cloud
- Content information versions
- BranchCache installation guide
- Operating system versions for BranchCache
- BranchCache Security
- Content flow and processes
- Cache Security

What is BranchCache?

BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server® 2012 and Windows® 8 operating systems, as well as in some editions of Windows Server® 2008 R2 and Windows® 7. To optimize WAN bandwidth when users access content on remote servers, BranchCache copies content from your main office or hosted cloud content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.

At branch offices, content is stored either on servers that are configured to host the cache or, when no server is available in the branch office, on client computers that are running Windows 8 or Windows 7. After a client computer requests and receives content from the main office and the content is cached at the branch office, other computers at the same branch office can obtain the content locally rather than downloading the content from the content server over the WAN link.

When subsequent requests for the same content are made by client computers, the clients download *content information* from the server instead of the actual content. Content information consists of hashes that are calculated using chunks of the original content, and are extremely small compared to the content in the original data. Client computers then use the content information to locate the content from a cache in the branch office, whether the cache is located

on a client computer or on a server. Client computers and servers also use content information to secure cached content so that it cannot be accessed by unauthorized users.

BranchCache increases end user productivity by improving content query response times for clients and servers in branch offices, and can also help improve network performance by reducing traffic over WAN links.

BranchCache modes

BranchCache has two modes of operation: distributed cache mode and hosted cache mode.

When you deploy BranchCache in distributed cache mode, the content cache at a branch office is distributed among client computers.

When you deploy BranchCache in hosted cache mode, the content cache at a branch office is hosted on one or more server computers, which are called hosted cache servers.

Note

You can deploy BranchCache using both modes, however only one mode can be used per branch office. For example, if you have two branch offices, one which has a server and one which does not, you can deploy BranchCache in hosted cache mode in the office that contains a server, while deploying BranchCache in distributed cache mode in the office that contains only client computers.

In the following illustration, BranchCache is deployed in both modes.

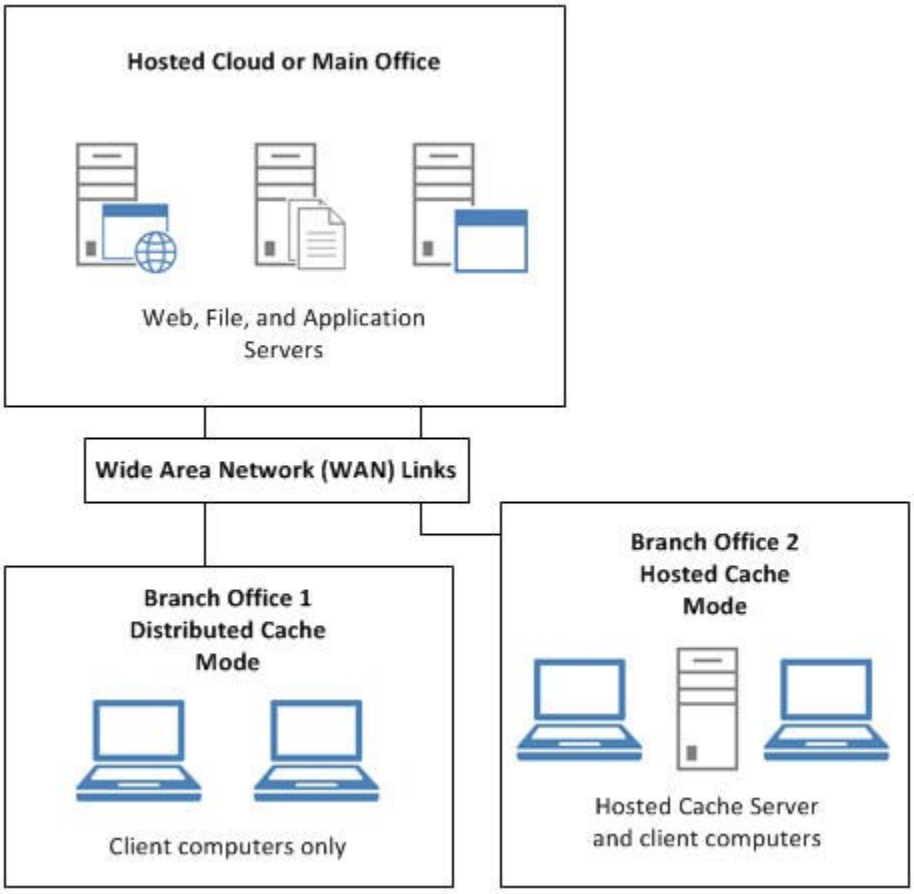

Distributed cache mode is best suited for small branch offices that do not contain a local server for use as a hosted cache server. Distributed cache mode allows you to deploy BranchCache with no additional hardware in branch offices.

If the branch office where you want to deploy BranchCache contains additional infrastructure, such as one or more servers that are running other workloads, deploying BranchCache in hosted cache mode is beneficial for the following reasons:

Increased cache availability

Hosted cache mode increases the cache efficiency because content is available even if the client that originally requested and cached the data is offline. Because the hosted cache server is always available, more content is cached, providing greater WAN bandwidth savings, and BranchCache efficiency is improved.

Centralized caching for multiple-subnet branch offices

Distributed cache mode operates on a single subnet. At a multiple-subnet branch office that is configured for distributed cache mode, a file downloaded to one subnet cannot be shared with client computers on other subnets. Because of this, clients on other subnets, unable to discover

that the file has already been downloaded, get the file from the main office content server, using WAN bandwidth in the process. When you deploy hosted cache mode, however, this is not the case - all clients in a multiple-subnet branch office can access a single cache, which is stored on the hosted cache server, even if the clients are on different subnets. In addition, BranchCache in Windows Server 2012 provides the ability to deploy more than one hosted cache server per branch office.

Caution

If you use BranchCache for SMB caching of files and folders, do not disable Offline Files. If you disable Offline Files, BranchCache SMB caching does not function correctly.

BranchCache-enabled content servers

When you deploy BranchCache, the source content is stored on BranchCache-enabled content servers in your main office. The following types of content servers are supported by BranchCache:

Note

Only source content – that is, content that client computers initially obtain from a BranchCache-enabled content server – is accelerated by BranchCache. Content that client computers obtain directly from other sources, such as Web servers on the Internet or Windows Update, is not cached by client computers or hosted cache servers and then shared with other computers in the branch office. If you want to accelerate Windows Update content, however, you can install a Windows Server Update Services (WSUS) application server at your main office and configure it as a BranchCache content server.

Web servers

Supported Web servers include computers that are running Windows Server 2012 or Windows Server® 2008 R2 that have the Web Server (IIS) server role installed and that use Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS). In addition, the Web server must have the BranchCache feature installed. For more information about Web Server (IIS) in Windows Server 2012, see Web Server (IIS) Overview.

File servers

Supported file servers include computers that are running Windows Server 2012 or Windows Server 2008 R2 that have the File Services server role and the BranchCache for Network Files role service installed. These file servers use Server Message Block (SMB) to exchange information between computers. After you complete installation of your file server, you must also share folders and enable hash generation for shared folders by using Group Policy or Local Computer Policy to enable BranchCache. For more information about File and Storage Services in Windows Server 2012, see File and Storage Services Overview.

Application servers

Supported application servers include computers that are running Windows Server 2012 or Windows Server 2008 R2 with Background Intelligent Transfer Service (BITS) installed and enabled. In addition, the application server must have the BranchCache feature installed. For example, you can deploy Microsoft Windows Server Update Services (WSUS) and Microsoft System Center Configuration Manager Branch Distribution Point servers as BranchCache content servers.

BranchCache and the cloud

The cloud has enormous potential to reduce operational expenses and achieve new levels of scale, but moving workloads away from the people who depend on them can increase networking costs and hurt productivity. Users expect high performance and don't care where their applications and data are hosted. BranchCache can improve the performance of networked applications and reduce bandwidth consumption with a shared cache of data. It improves productivity in branch offices and in headquarters, where workers are using servers that are deployed in the cloud.

Because BranchCache does not require new hardware or network topology changes, it is an excellent solution for improving communication between office locations and both public and private clouds.

For more information about cloud technologies in Windows Server 2012, see Building Your Cloud Infrastructure: Scenario Overview.

Content information versions

There are two versions of content information:

- Content information that is compatible with computers running Windows Server 2008 R2 and Windows 7 is called version 1, or V1. With V1 BranchCache file segmentation, file segments are larger than in V2 and are of fixed size. Because of large fixed segment sizes, when a user makes a change that modifies the file length, not only is the segment with the change invalidated, but all of the segments to the end of the file are invalidated. The next call for the changed file by another user in the branch office therefore results in reduced WAN bandwidth savings because the changed content and all content after the change are sent over the WAN link.
- Content information that is compatible with computers running Windows Server 2012 and Windows 8 is called version 2, or V2. V2 content information uses smaller, variable-sized segments that are more tolerant to changes within a file. This increases the probability that segments from an older version of the file can be reused when users access an updated version, causing them to retrieve only the changed portion of the file from the content server, and using less WAN bandwidth.

The following table provides information on the content information version that is used depending upon which client, content server, and hosted cache server operating systems you are using in your BranchCache deployment.

Note
In the table below, the acronym “OS” means operating system.

Client OS	Content Server OS	Hosted Cache Server OS	Content Information Version
Windows Server 2008 R2 and Windows 7	Windows Server 2012 or Windows Server 2008 R2	Windows Server 2012 or Windows Server 2008 R2; none for distributed cache mode	V1
Windows Server 2012 and Windows 8	Windows Server 2008 R2	Windows Server 2012 or Windows Server 2008 R2; none for distributed cache mode	V1
Windows Server 2012 and Windows 8	Windows Server 2012	Windows Server 2008 R2; none for distributed cache mode	V1
Windows Server 2012 and Windows 8	Windows Server 2012	Windows Server 2012; none for distributed cache mode	V2

When you have content servers and hosted cache servers that are running Windows Server 2012, they use the content information version that is appropriate based on the operating system of the BranchCache client that requests information. When computers running Windows Server 2012 and Windows 8 request content, the content and hosted cache servers use V2 content information; when computers running Windows Server 2008 R2 and Windows 7 request content, the content and hosted cache servers use V1 content information.

When you deploy BranchCache in distributed cache mode, clients that use different content information versions do not share content with each other.

BranchCache installation guide

You can use Server Manager in Windows Server 2012 to install either the BranchCache feature or the BranchCache for Network Files role service of the File Services server role. You can use the following table to determine whether to install the role service or the feature.

Functionality	Computer location	Install this BranchCache element
Content server (BITS-based application server)	Main office or cloud data center	BranchCache feature
Content server (Web server)	Main office or cloud data center	BranchCache feature
Content server (file server using the SMB protocol)	Main office or cloud data center	BranchCache for Network Files role service of the File Services server role
Hosted cache server	Branch office	BranchCache feature with hosted cache server mode enabled
BranchCache-enabled client computer	Branch office	No installation needed; just enable BranchCache and a BranchCache mode (distributed or hosted) on the client

To install either the role service or the feature, open Server Manager and select the computers where you want to enable BranchCache functionality. In Server Manager, click **Manage**, and then click **Add Roles and Features**. The **Add Roles and Features** wizard opens. As you run the wizard, make the following selections:

- On the wizard page **Select Installation Type**, select **Role-based or Feature-based Installation**.
- On the wizard page **Select Server Roles**, if you are installing a BranchCache-enabled file server, select **File Services**. Later in the wizard configuration process, also select **BranchCache for Network Files**. If you do not want to install a BranchCache-enabled file server, do not install the File Services role with the BranchCache for Network Files role service.
- On the wizard page **Select features**, if you are installing a content server that is not a file server or you are installing a hosted cache server, select **BranchCache**. If you do not want to install a content server other than a file server or a hosted cache server, do not install the BranchCache feature.

Operating system versions for BranchCache

Following is a list of operating systems that support different types of BranchCache functionality.

Operating systems for BranchCache client computer functionality

- Windows® 8 Enterprise
- Windows® 7 Enterprise
- Windows® 7 Ultimate

Operating systems for BranchCache content server functionality

The Windows Server 2012 family of operating systems can be used as BranchCache content servers.

In addition, the Windows Server® 2008 R2 family of operating systems can be used as BranchCache content servers, with the following exceptions:

- BranchCache is not supported in Server Core installations of Windows Server® 2008 R2 Enterprise with Hyper-V.
- BranchCache is not supported in Server Core installations of Windows Server® 2008 R2 Datacenter with Hyper-V.

Operating systems for BranchCache hosted cache server functionality

The Windows Server 2012 family of operating systems can be used as BranchCache hosted cache servers.

In addition, the following Windows Server® 2008 R2 operating systems can be used as BranchCache hosted cache servers:

- Windows Server® 2008 R2 Enterprise
- Windows Server 2008 R2 Enterprise with Hyper-V
- Windows Server 2008 R2 Enterprise Server Core Installation
- Windows Server 2008 R2 Enterprise Server Core Installation with Hyper-V
- Windows Server 2008 R2 for Itanium-Based Systems
- Windows Server® 2008 R2 Datacenter
- Windows Server® 2008 R2 Datacenter with Hyper-V
- Windows Server 2008 R2 Datacenter Server Core Installation with Hyper-V

BranchCache Security

BranchCache implements a secure-by-design approach that works seamlessly alongside your existing network security architectures, without the requirement for additional equipment or complex additional security configuration.

BranchCache is non-invasive and does not alter any Windows authentication or authorization processes. After you deploy BranchCache, authentication is still performed using domain credentials, and the way in which authorization with Access Control Lists (ACLs) functions is

unchanged. In addition, other configurations continue to function just as they did before BranchCache deployment.

The BranchCache security model is based on the creation of metadata, which takes the form of a series of hashes. These hashes are also called content information.

After content information is created, it is used in BranchCache message exchanges rather than the actual data, and it is exchanged using the supported protocols (HTTP, HTTPS, and SMB).

Cached data is kept encrypted and cannot be accessed by clients that do not have permission to access content from the original source. Clients must be authenticated and authorized by the original content source before they can retrieve content metadata, and must possess content metadata to access the cache in the local office.

How BranchCache generates content information

Because content information is created from multiple elements, the value of the content information is always unique. These elements are:

- The actual content (such as Web pages or shared files) from which the hashes are derived.
- Configuration parameters, such as the hashing algorithm and block size. To generate content information, the content server divides the content into segments and then subdivides those segments into blocks. BranchCache uses secure cryptographic hashes to identify and verify each block and segment, supporting the SHA256 hash algorithm.
- A server secret. All content servers must be configured with a server secret, which is a binary value of arbitrary length.

Note

The use of a server secret ensures that client computers are not able to generate the content information themselves. This prevents malicious users from using brute force attacks with BranchCache-enabled client computers to guess minor changes in content across versions in situations in which the client had access to a previous version but does not have access to the current version.

Content information details

BranchCache uses the server secret as a key in order to derive a content-specific hash that is sent to authorized clients. Applying a hashing algorithm to the combined server secret and the Hash of Data generates this hash.

This hash is called the segment secret. BranchCache uses segment secrets to secure communications. In addition, BranchCache creates a Block Hash List, which is list of hashed data blocks, and the Hash of Data, which is generated by hashing the Block Hash List.

The content information includes the following:

- The Block Hash List:

```
BlockHashi = Hash(dataBlocki)   1<=i<=n
```

- The Hash of Data (HoD):

```
HoD = Hash(BlockHashList)
```

- Segment Secret (Kp):

```
Kp = HMAC(Ks, HoD)
```

BranchCache uses the Peer Content Caching protocol and the Retrieval Framework protocol to implement the processes that are required to ensure the secure caching and retrieval of data between content caches.

In addition, BranchCache handles content information with the same degree of security that it uses when handling and transmitting the actual content itself.

Content flow and processes

The flow of content information and actual content is divided into four phases:

1. BranchCache processes: Request content
2. BranchCache processes: Locate content
3. BranchCache processes: Retrieve content
4. BranchCache processes: Cache content

The following sections describe these phases.

BranchCache processes: Request content

In the first phase, the client computer in the branch office requests content, such as a file or a Web page, from a content server in a remote location, such as a main office. The content server verifies that the client computer is authorized to receive the requested content. If the client computer is authorized and both content server and client are BranchCache-enabled, the content server generates content information.

The content server then sends the content information to the client computer using the same protocol as would have been used for the actual content. For example, if the client computer requested a Web page over HTTP, the content server sends the content information using HTTP. Because of this, the wire-level security guarantees of the content and the content information are identical.

After the initial portion of content information (Hash of Data + Segment Secret) is received, the client computer performs the following actions:

- Uses the Segment Secret (Kp) as the encryption key (Ke).
- Generates the Segment ID (HoHoDk) from the HoD and Kp:

```
HoHoDk = HMAC(Kp, HoD + C), where C is the ASCII string "MS_P2P_CACHING" with NUL
terminator.
```

The primary threat at this layer is the risk to the Segment Secret, however BranchCache encrypts the content data blocks to protect the Segment Secret. BranchCache does this by using the encryption key that is derived from the Segment Secret of the content segment within which the

content blocks are located. This approach ensures that an entity that is not in possession of the server secret cannot discover the actual content in a data block. The Segment Secret is treated with the same degree of security as the plaintext segment itself, because knowledge of the Segment Secret for a given segment enables an entity to obtain the segment from peers and then decrypt it. Knowledge of the Server Secret does not immediately yield any particular plaintext but can be used to derive certain types of data from the cipher text and then to possibly expose some partially known data to a brute-force guessing attack. The server secret, therefore, should be kept confidential.

BranchCache processes: Locate content

After the content information is received by the client computer, the client uses the Segment ID to locate the requested content in the local branch office cache, whether that cache is distributed between client computers or is located on a hosted cache server.

If the client computer is configured for hosted cache mode, it is configured with the computer name of the hosted cache server and contacts that server to retrieve the content.

If the client computer is configured for distributed cache mode, however, the content might be stored across multiple caches on multiple computers in the branch office. The client computer must discover where the content is located before the content is retrieved.

When they are configured for distributed cache mode, client computers locate content by using a discovery protocol that is based on the Web Services Dynamic Discovery (WS-Discovery) protocol. Clients send WS-Discovery multicast Probe messages to discover cached content over the network. Probe messages include the Segment ID, which enables clients to check whether the requested content matches the content stored in their cache. Clients that receive the initial Probe message reply to the querying client with unicast Probe-Match messages if the Segment ID matches content that is cached locally.

The success of the WS-Discovery process depends on the fact that the client that is performing the discovery has the correct content information, which was provided by the content server, for the content that it is requesting.

The main threat to data during the Request content phase is information disclosure, because access to the content information implies authorized access to content. To mitigate this risk, the discovery process does not reveal the content information, other than the Segment ID, which does not reveal anything about the plaintext segment that contains the content.

In addition, another client computer run by a malicious user on the same network subnet can see the BranchCache discovery traffic to the original content source going through the router.

If the requested content is not found in the branch office, the client requests the content directly from the content server across the WAN link.

After the content is received, it is added to the local cache, either on the client computer or on a hosted cache server. In this case, the content information prevents a client or hosted cache server from adding to the local cache any content that does not match the hashes. The process of verifying content by matching hashes ensures that only valid content is added to the cache, and the integrity of the local cache is protected.

BranchCache processes: Retrieve content

After a client computer locates the desired content on the content host, which is either a hosted cache server or a distributed cache mode client computer, the client computer begins the process of retrieving the content.

First the client computer sends a request to the content host for the first block that it requires. The request contains the Segment ID and block range that identify the desired content. Because only one block is returned, the block range contains only a single block. (Requests for multiple blocks are currently not supported.) The client also stores the request in its local Outstanding Request List.

Upon receiving a valid request message from a client, the content host checks whether the block specified in the request exists in the content host's content cache.

If the content host is in possession of the content block, then the content host sends a response that contains the Segment ID, the Block ID, the encrypted data block, and the initialization vector that is used for encrypting the block.

If the content host is not in possession of the content block, the content host sends an empty response message. This informs the client computer that the content host does not have the requested block. An empty response message contains the Segment ID and Block ID of the requested block, along with a zero-sized data block.

When the client computer receives the response from the content host, the client verifies that the message corresponds to a request message in its Outstanding Request List. (The Segment ID and block index must match that of an outstanding request.)

If this verification process is unsuccessful and the client computer does not have a corresponding request message in its Outstanding Request List, the client computer discards the message.

If this verification process is successful and the client computer has a corresponding request message in its Outstanding Request List, the client computer decrypts the block. The client then validates the decrypted block against the appropriate block hash from the content information that the client initially obtained from the original content server.

If the block validation is successful, the decrypted block is stored in the cache.

This process is repeated until the client has all of the required blocks.

Note

If the complete segments of content do not exist on one computer, the retrieval protocol retrieves and assembles content from a combination of sources: a set of distributed cache mode client computers, a hosted cache server, and - if the branch office caches do not contain the complete content - the original content server in the main office.

Before BranchCache sends content information or content, the data is encrypted. BranchCache encrypts the block in the response message. In Windows 7, the default encryption algorithm that BranchCache uses is AES-128, the encryption key is Ke, and the key size is 128 bits, as dictated by the encryption algorithm. BranchCache generates an initialization vector that is suitable for the encryption algorithm and uses the encryption key to encrypt the block. BranchCache then records the encryption algorithm and the initialization vector in the message. Servers and clients never

exchange, share, or send each other the encryption key. The client receives the encryption key from the content server that hosts the source content. Then, using the encryption algorithm and initialization vector it received from the server, it decrypts the block. There is no other explicit authentication or authorization built into the download protocol.

Security threats

The primary security threats at this layer include:

- Tampering with data:

  A client serving data to a requester tampers with the data. The BranchCache security model uses hashes to confirm that neither the client nor the server has altered the data.

- Information disclosure:

  BranchCache sends encrypted content to any client that specifies the appropriate Segment ID. Segment IDs are public, so any client can receive encrypted content. However, if a malicious user obtains encrypted content, they must know the encryption key to decrypt the content. The upper layer protocol performs authentication and then gives the content information to the authenticated and authorized client. The security of the content information is equivalent to the security provided to the content itself, and BranchCache never exposes the content information.

  An attacker sniffs the wire to obtain the content. BranchCache encrypts all transfers between clients by using AES128 where the secret key is Ke, preventing data from being sniffed from the wire. Content information that is downloaded from the content server is protected in exactly the same way as the data itself would have been and is hence no more or less protected from information disclosure than if BranchCache had not been used at all.

- Denial of Service:

  A client is overwhelmed by requests for data. BranchCache protocols incorporate queue management counters and timers to prevent clients from being overloaded.

BranchCache processes: Cache content

On distributed cache mode client computers and hosted cache servers that are located in branch offices, content caches are built up over time as content is retrieved over WAN links.

When client computers are configured with hosted cache mode, they add content to their own local cache and also offer data to the hosted cache server. The Hosted Cache Protocol provides a mechanism for clients to inform the hosted cache server about content and segment availability. To upload content to the hosted cache server, the client informs the server that it has a segment that is available. The hosted cache server then retrieves all of the content information that is associated with the offered segment, and downloads the blocks within the segment that it actually needs. This process is repeated until the client has no more segments to offer the hosted cache server.

To update the hosted cache server by using the Hosted Cache Protocol, the following requirements must be met:

- The client computer is required to have a set of blocks within a segment that it can offer to the hosted cache server. The client must supply content information for the offered segment; this is comprised of the Segment ID, the segment Hash of Data, the Segment Secret, and a list of all block hashes that are contained within the segment.
- For hosted cache servers that are running Windows Server 2008 R2, a hosted cache server certificate and associated private key are required, and the certification authority (CA) that issued the certificate must be trusted by client computers in the branch office. This allows the client and server to participate successfully in HTTPS Server authentication.

  Important
  Hosted cache servers that are running Windows Server 2012 do not require a hosted cache server certificate and associated private key.
- The client computer is configured with the computer name of the hosted cache server and the Transmission Control Protocol (TCP) port number upon which the hosted cache server is listening for BranchCache traffic. The hosted cache server's certificate is bound to this port. The computer name of the hosted cache server can be a fully qualified domain name (FQDN), if the hosted cache server is a domain member computer; or it can be the NetBIOS name of the computer if the hosted cache server is not a domain member.
- The client computer actively listens for incoming block requests. The port on which it is listening is passed as part of the offer messages from the client to the hosted cache server. This enables the hosted cache server to use BranchCache protocols to connect to the client computer to retrieve data blocks in the segment.
- The hosted cache server starts to listen for incoming HTTP requests when it is initialized.
- If the hosted cache server is configured to require client computer authentication, both the client and the hosted cache server are required to support HTTPS authentication.

Hosted cache mode cache population

The process of adding content to the hosted cache server's cache in a branch office begins when the client sends an INITIAL_OFFER_MESSAGE, which includes the Segment ID. The Segment ID in the INITIAL_OFFER_MESSAGE request is used to retrieve the corresponding segment Hash of Data, list of block hashes, and the Segment Secret from the hosted cache server's block cache. If the hosted cache server already has all the content information for a particular segment, the response to the INITIAL_OFFER_MESSAGE will be OK, and no request to download blocks occurs.

If the hosted cache server does not have all of the offered data blocks that are associated with the block hashes in the segment, the response to the INITIAL_OFFER_MESSAGE is INTERESTED. The client then sends a SEGMENT_INFO_MESSAGE that describes the single segment that is being offered. The hosted cache server responds with an OK message and initiates the download of the missing blocks from the offering client computer.

The segment Hash of Data, list of block hashes, and the segment secret are used to ensure that the content that is being downloaded has not been tampered with or otherwise altered. The downloaded blocks are then added to the hosted cache server's block cache.

Cache Security

This section provides information on how BranchCache secures cached data on client computers and on hosted cache servers.

Client computer cache security

The greatest threat to data stored in the BranchCache is tampering. If an attacker can tamper with content and content information that is stored in the cache, then it might be possible to use this to try and launch an attack against the computers that are using BranchCache. Attackers can initiate an attack by inserting malicious software in place of other data. BranchCache mitigates this threat by validating all content using block hashes found in the content information. If an attacker attempts to tamper with this data, it is discarded and is replaced with valid data from the original source.

A secondary threat to data stored in the BranchCache is information disclosure. In distributed cache mode, the client caches only the content that it has requested itself; however, that data is stored in clear text, and might be at risk. To help restrict cache access to the BranchCache Service only, the local cache is protected by file system permissions that are specified in an ACL. Although the ACL is effective in preventing unauthorized users from accessing the cache, it is possible for a user with administrative privileges to gain access to the cache by manually changing the permissions that are specified in the ACL. BranchCache does not protect against the malicious use of an administrative account.

Data that is stored in the content cache is not encrypted, so if data leakage is a concern, you can use encryption technologies such as BitLocker or the Encrypting File System (EFS). The local cache that is used by BranchCache does not increase the information disclosure threat borne by a computer in the branch office; the cache contains only copies of files that reside unencrypted elsewhere on the disk. Encrypting the entire disk is particularly important in environments in which the physical security of the clients is difficult to ensure. For example, encrypting the entire disk helps to secure sensitive data on mobile computers that might be removed from the branch office environment.

Hosted cache server cache security

In hosted cache mode, the greatest threat to the security of the hosted cache server is information disclosure. BranchCache in a hosted cache environment behaves in a similar manner to distributed cache mode, with file system permission protecting the cached data. The difference is that the hosted cache server stores all of the content that any BranchCache-enabled computer in the branch office requests, rather than just the data that a single client requests. The consequences of unauthorized intrusion into this cache could be much more serious, because much more data is at risk.

In a hosted cache environment where the hosted cache server is running Windows Server 2008 R2, the use of encryption technologies such as BitLocker or EFS is advisable if any of the clients in the branch office can access sensitive data across the WAN link. It is also necessary to prevent physical access to the hosted cache, because disk encryption works only

when the computer is turned off when the attacker gains physical access. If the computer is turned on or is in sleep mode, then disk encryption offers little protection.

Note
Hosted cache servers that are running Windows Server 2012 encrypt all data in the cache by default, so the use of additional encryption technologies is not required.

Even if a client is configured in hosted cache mode, it will still cache data locally, and you might want to take steps to protect the local cache in addition to the cache on the hosted cache server.

BranchCache Learning Roadmap

BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server® 2012 and Windows® 8 operating systems.

To optimize WAN bandwidth, BranchCache copies content from your main office content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.

At branch offices, content is cached either on servers that are running the BranchCache feature of Windows Server 2012 or, when no server is available in the branch office, on computers running Windows 8. After a client computer requests and receives content from the main office and the content is cached at the branch office, other computers at the same branch office can obtain the content locally rather than contacting the main office over the WAN link.

BranchCache helps improve content query response times for clients and servers in branch offices, and can also help improve network performance by reducing traffic over WAN links.

If you are new to BranchCache, this topic can help you identify what you need to learn to fully understand how to deploy and troubleshoot BranchCache. It includes prerequisite topics that cover a variety of networking and IT infrastructure fundamentals. You must understand the prerequisite technologies first, because BranchCache builds upon them and assumes an understanding of them. Afterwards, you can begin learning about BranchCache through the resources in the Level 100, 200, and 300 sections.

We recommend that you read the topics in the order listed.

- Prerequisite information
- Level 100
- Level 200
- Level 300
- Optional information

Note
Some of the hypertext links in the following sections provide access to Windows Server® 2008 R2 documentation.

Prerequisite information

The following documents contain the information that you need to fully understand the tools that are used to deploy and manage BranchCache.

- **Step 1: Learn about the Microsoft Management Console (MMC)**.

  See Microsoft Management Console 3.0 (http://go.microsoft.com/fwlink/?LinkID=70036).

  This document describes the Microsoft Management Console (MMC), which hosts administrative tools that you can use to administer networks, computers, services, and other system components.

  Your goal is to understand how to use the MMC so that you can add and remove MMC snap-ins to manage Group or Local Computer Policy and configure BranchCache policies.

- **Step 2: Learn about Group and Local Computer Policy**.

  See the Group Policy Overview (http://technet.microsoft.com/library/hh831791) and Local Group Policy Editor (http://go.microsoft.com/fwlink/?LinkID=67684).

  These documents describe domain-based Group Policy, Local Computer Policy, and tools that you can use to create and modify policies.

  Your goal is to understand the basics of how to use either Group Policy with Active Directory Domain Services or Local Computer Policy, because you use these policies to deploy BranchCache.

- **Step 3: Learn about Windows PowerShell**.

  See the Windows PowerShell User's Guide (http://technet.microsoft.com/library/cc196356).

  This document describes Windows PowerShell, which is a task-based command-line shell and scripting language that allows you to configure and display the status of various network communications server roles and components after they are installed on computers.

  Your goal is to understand how to run commands at the Windows PowerShell prompt, how to get information about commands, and working with objects so that you can use the BranchCache cmdlets to manage BranchCache.

- **Step 4: Learn how to install server roles and server features by using Server Manager**.

  See Manage Multiple, Remote Servers with Server Manager (http://technet.microsoft.com/library/hh831456).

  This document describes Server Manager, which is an expanded Microsoft Management Console (MMC) that allows you to view and manage virtually all of the information and tools that affect your server's productivity.

  Your goal is to understand how to install the BranchCache feature and how to install server roles for BranchCache content servers that you want to deploy, such as file servers, Web servers, and Background Intelligent Transfer Service (BITS)-based application servers.

Level 100

The following documents and videos contain introductory information about BranchCache.

- **Step 1: Learn the benefits of BranchCache**.

  See the BranchCache Executive Overview and view the video Branch Cache in Windows 7.

  This document and video describe BranchCache and provide an overview of BranchCache functionality.

  Your goal is to understand the business benefits of BranchCache.
- **Step 2: Learn about the components of BranchCache**.

  See the BranchCache Overview.

  This document describes the technologies that are used to design and deploy BranchCache.

  Your goal is to understand BranchCache modes, hardware requirements, and operating system versions required for each BranchCache component.
- **Step 3: Learn about BranchCache security**.

  See the section “BranchCache Security” in the topic BranchCache Overview.

  This document provides guidance to help organizations understand and manage the security of the Microsoft® BranchCache™ feature that is provided in some editions of Windows 8 and Windows Server 2012.

  Your goal is to understand how BranchCache security works.
- **Step 4: Learn about new BranchCache features**.

  See the topic What's New in BranchCache.

  This document provides an overview of new BranchCache features for Windows Server 2012 and Windows 8.

  Your goal is to familiarize yourself with BranchCache settings before deploying the technology.
- **Step 5: Find out more about BranchCache**.

  See the TechNet Wiki topic BranchCache Frequently Asked Questions.

  This topic provides information about different aspects of BranchCache, from preloading content in branch offices to information about deploying hosted cache mode.

  Your goal is to have a more detailed understanding of how BranchCache works.

Level 200

The following documents contain intermediate information about BranchCache in Windows Server 2012.

- **Step 1: Learn how to create an effective design for a BranchCache deployment**.

  See the section Choosing a BranchCache Design in the BranchCache Deployment Guide.

  This section of the BranchCache Deployment Guide describes the design considerations for BranchCache to meet your organization’s requirements for deploying BranchCache in branch offices that have different configurations.

  Your goal is to understand the different BranchCache modes, the infrastructure elements and requirements, and how to map your deployment goals to a BranchCache design.

- **Step 2: Learn how to deploy BranchCache content servers**.

  See the section Install and Configure Content Servers in the BranchCache Deployment Guide.

  This section of the BranchCache Deployment Guide provides instructions on which BranchCache software to install on your content server to enable BranchCache functionality.

  Your goal is to install the correct BranchCache software for the types of content servers that you are using at your main office.
- **Step 3: Deploy a BranchCache distributed cache mode design**.

  See the section Configure BranchCache Client Computers in the BranchCache Deployment Guide.

  This section of the BranchCache Deployment Guide provides instructions on how to deploy BranchCache in distributed cache mode using Windows 8.

  Your goal is to deploy BranchCache in distributed cache mode.
- **Step 4: Learn basic troubleshooting for BranchCache**.

  See the BranchCache Troubleshooting.

  This document describes common troubleshooting situations and the tools that you can use to help diagnose and resolve BranchCache deployment problems.

  Your goal is to understand the kinds of problems that can occur when deploying BranchCache.

Level 300

The following documents contain advanced information about BranchCache.

- **Step 1: Learn the Windows PowerShell cmdlets for BranchCache**

  See the BranchCache Cmdlets in Windows PowerShell (http://technet.microsoft.com/library/hh848392.aspx).

  This topic, which is part of the Windows PowerShell cmdlet reference, provides a comprehensive list of BranchCache commands, including syntax and examples.

  Your goal is to use these commands either manually or in scripts to manage and configure BranchCache.
- **Step 2: Deploy BranchCache in hosted cache mode**.

  Deploy hosted cache mode by using the Core Network Companion Guide: Deploying BranchCache Hosted Cache Mode.

  This guide includes instructions on how to deploy BranchCache in hosted cache mode and, optionally, how to prehash the content on your Web and file servers, and how to preload the content on your hosted cache server.

  Your goal is to deploy BranchCache in hosted cache mode.

Optional information

The following documents and videos contain information about the BranchCache Netsh commands and the technologies that you can use to deploy BranchCache content servers, including Web servers, file servers, and BITS based application servers.

- **Learn the network shell (netsh) commands for BranchCache**.

  See the Network Shell (Netsh) Commands for BranchCache (http://go.microsoft.com/fwlink/?LinkId=156640)

  This topic, which is part of the Netsh command reference, provides a comprehensive list of netsh branchcache commands, including syntax and examples.

  Your goal is to use the netsh branchcache commands either manually or in scripts to manage and configure BranchCache.

- **Deploy one or more file servers for use as BranchCache content servers**.

  See File and Storage Services Overview.

  This topic contains detailed information about File Services technologies in Windows Server 2012. File Services provides technologies that help manage storage, enable file replication, manage shared folders, ensure fast file searching, and enable access for UNIX client computers.

  Your goal is to learn how to deploy file servers that can be used as BranchCache content servers, and to manage shares that you can enable for use with BranchCache.

- **Deploy one or more application servers for use as BranchCache content servers**.

  See Application Server Overview.

  This document describes the Application Server role, which provides an integrated environment for deploying and running custom, server-based business applications that are built with the Microsoft .NET Framework 3.0. The Application Server role supports applications that use COM+, Message Queuing, Web services, and distributed transactions.

  Your goal is to learn how to deploy BITS-based application servers that can be used as BranchCache content servers.

- **Deploy one or more Web servers for use as BranchCache content servers**.

  See Web Server (IIS) Overview.

  This document describes the Web Server (IIS) role, which includes Internet Information Services (IIS) 7, which is a unified Web platform that integrates IIS, ASP.NET, Windows Communication Foundation, and Windows SharePoint Services.

  Your goal is to learn how to deploy Web servers that can be used as BranchCache content servers.

What's New in BranchCache

BranchCache in Windows Server 2012 and Windows 8 provides substantial performance, manageability, scalability, and availability improvements.

- **Office sizes and the number of branch offices are not limited**. BranchCache supports any number of offices of any size when you deploy hosted cache mode with multiple hosted cache servers.
- **There is no requirement for a Group Policy Object (GPO) for each office location, streamlining deployment**. A single GPO that contains a small group of settings is all that is required to deploy BranchCache in any size organization, from a large enterprise to a small business.
- **Client computer configuration is automatic**. Clients can be configured through Group Policy as distributed cache mode clients by default, however they will search for a hosted cache server – and if one is discovered, clients automatically self-configure as hosted cache mode clients.
- **BranchCache is deeply integrated with the Windows file server**. BranchCache uses Windows file server's state-of-the-art technology to divide files into small pieces and eliminate duplicates. This greatly increases the chance of finding duplicate pieces in independent files, resulting in greater bandwidth savings. BranchCache is also more tolerant of small changes in large files. In addition, file division calculations are performed only one time. When a client computer that is running Windows 8 downloads content from a file server or web server that is running Windows Server 2012 and is using new Data Deduplication technology, there is no need for BranchCache to spend CPU cycles calculating how to divide the content – because the file server and web server have already made these calculations.

  For more information about Data Deduplication in File and Storage Services, see Data Deduplication Overview.
- **Duplicate content is stored once and downloaded once**. When identical content exists in a file – or across many files on the content server or hosted cache server – BranchCache stores only one instance of the content, which provides disk storage savings. In addition, client computers at office locations download only one instance of duplicate content, saving additional WAN bandwidth.
- **Small changes to large files produce bandwidth savings**. BranchCache now divides files and web pages into smaller pieces. BranchCache uses a highly optimized file server chunking system to choose intelligent split points inside of files, providing the ability for client computers to download only the part of the file that is changed.
- **Offline creation of content information**. When you deploy a BranchCache enabled file server or Web server as a content server, content information is calculated offline, well before a BranchCache client requests a file. This provides faster performance and more bandwidth savings, because content information is ready for the first client that requests the content and calculations have already been performed.
- **Cache Encryption**. Cached data is now stored encrypted by default. This allows you to ensure data security without using drive encryption technologies.

- **Cache preloading**. New tools are available that allow you to load cacheable content onto hosted cache servers before the content is requested by client computers. You can preload content from media, such as DVDs or hard disks, or you can transfer the content over the network to the hosted cache server.
- **BranchCache is now manageable with PowerShell and Windows Management Instrumentation (WMI)**. This enables scripting and remote management of BranchCache content servers, hosted cache servers, and client computers.

In addition to these improvements, the following improvements to hosted cache mode and hosted cache servers are included in Windows Server 2012.

- **Deployment of multiple hosted cache servers**. In the previous version of BranchCache, you were able to deploy only one hosted cache server per office location. Windows Server 2012 provides the ability to scale hosted cache mode deployments for offices of any size by allowing you to deploy as many hosted cache servers as are needed at a location.
- **Improved database performance**. BranchCache now uses the Extensible Storage Engine (ESE) database technology that powers Microsoft Exchange Server. This allows a single hosted cache server to keep up with the demands of more people while using the same hardware. It also allows a hosted cache server to store significantly more data (on the order of terabytes), which is necessary to provide high optimization for large organizations.
- **No server certificate requirement**. Previously, hosted cache servers were required to have a server certificate that was issued by a certification authority (CA) that client computers at the office location trusted. Deploying a public key infrastructure with one or more CAs is complex and expensive, and this requirement is now removed, because BranchCache security is improved with data encryption and other technologies.

For more information, see BranchCache overview at http://technet.microsoft.com/library/hh831696.aspx.

New BranchCache Group Policy settings

BranchCache includes the following new Group Policy and Local Computer Policy settings.

- **Hash Version Support for BranchCache**. This policy setting specifies whether the BranchCache hash generation service supports version 1 (V1) hashes, version 2 (V2) hashes, or both V1 and V2 hashes. Hashes, also called content information, are created based on the data in shared folders where BranchCache is enabled. V2 content information describes smaller, variable sized chunks of data and enables greater savings of WAN bandwidth. V1 hashes are compatible with Windows 7 and Windows Server 2008 R2; V2 hashes are compatible with Windows 8 and Windows Server 2012.If you specify only one version that is supported, content information for that version is the only type that is generated by BranchCache, and it is the only type of content information that can be retrieved by client computers. For example, if you enable support for V1 hashes, BranchCache generates only V1 hashes and client computers – even those that are running Windows 8 - can retrieve only V1 hashes.
- **Set age for segments in the data cache**. This policy setting specifies the age, in days, for which segments are valid in the BranchCache data cache on client computers. If you enable

this policy setting, you can configure the maximum age for segments in the data cache. If you disable or do not configure this policy setting, the default age is set to 28 days.

- **Configure Client BranchCache Version Support**. This policy setting specifies whether BranchCache-capable client computers operate in a downgraded mode in order to maintain compatibility with previous versions of BranchCache. If client computers do not use the same BranchCache version, cache efficiency might be reduced because client computers that are using different versions of BranchCache might store cache data in incompatible formats. If you enable this policy setting, all clients use the version of BranchCache that you specify in the policy options. The versions that you can select in the policy options are:
  - Windows Vista with BITS 4.0 installed, Windows 7, or Windows Server 2008 R2. If you select this version, later versions of Windows run the version of BranchCache that is included in these operating systems rather than later versions of BranchCache.
  - Windows 8 or Windows Server 2012. If you select this version, later versions of Windows run the version of BranchCache that is included in these operating systems.

  If you do not configure this setting, all clients will use the version of BranchCache that matches their operating system.
- **Configure Hosted Cache Servers**. This policy setting specifies whether client computers are configured to use hosted cache mode and provides the computer names of the hosted cache servers that are available to the client computers. Hosted cache mode enables client computers in branch offices to retrieve content from one or more hosted cache servers that are installed in the same office location. You can use this setting to automatically configure client computers that are configured for hosted cache mode with the computer names of the hosted cache servers in the branch office.

  If you enable this policy setting and specify valid computer names of hosted cache servers, hosted cache mode is enabled for all client computers to which the policy setting is applied. For this policy setting to take effect, you must also enable the "Turn on BranchCache" policy setting.

  This policy setting can only be applied to client computers that are running Windows 8. This policy has no effect on computers that are running Windows 7 or Windows Vista. Client computers to which this policy setting is applied, in addition to the "Set BranchCache Hosted cache mode" policy setting, use the hosted cache servers that are specified in this policy setting and do not use the hosted cache server that is configured in the policy setting “Set BranchCache Hosted cache mode.”

  If you do not configure this policy setting, or if you disable this policy setting, client computers that are configured with hosted cache mode still function correctly.
- **Enable Automatic Hosted Cache Discovery by Service Connection Point**. This policy setting specifies whether client computers should attempt the automatic configuration of hosted cache mode by searching for hosted cache servers publishing service connection points that are associated with the client's current Active Directory site. If you enable this policy setting, client computers to which the policy setting is applied search for hosted cache servers using Active Directory, and will prefer these servers and hosted cache mode rather than manual BranchCache configuration or BranchCache configuration by other group policies.

If you enable this policy setting in addition to the "Turn on BranchCache" policy setting, BranchCache clients attempt to discover hosted cache servers in the local branch office. If client computers detect hosted cache servers, hosted cache mode is turned on. If they do not detect hosted cache servers, hosted cache mode is not turned on, and the client uses any other configuration that is specified manually or by Group Policy.

When this policy setting is applied, the client computer performs or does not perform automatic hosted cache server discovery under the following circumstances:

- If no other BranchCache mode-based policy settings are applied, the client computer performs automatic hosted cache server discovery. If one or more hosted cache servers is found, the client computer self-configures for hosted cache mode.
- If the policy setting "Set BranchCache Distributed cache mode" is applied in addition to this policy, the client computer performs automatic hosted cache server discovery. If one or more hosted cache servers are found, the client computer self-configures for hosted cache mode only.
- If the policy setting "Set BranchCache Hosted cache mode" is applied, the client computer does not perform automatic hosted cache discovery. This is also true in cases where the policy setting "Configure Hosted Cache Servers" is applied.

This policy setting can only be applied to client computers that are running Windows 8. This policy has no effect on computers that are running Windows 7 or Windows Vista. If you disable, or do not configure this setting, a client will not attempt to discover hosted cache servers by service connection point.

Additional BranchCache Group Policy settings

The following BranchCache Group Policy settings were also present in the previous version of BranchCache. These policy settings are supported on computers that are running Windows 8, Windows Server 2012, Windows 7, Windows Server 2008 R2, and Windows Vista Business, Enterprise, and Ultimate editions with Background Intelligent Transfer Service (BITS) 4.0 installed.

- **Turn on BranchCache**. This policy setting specifies whether BranchCache is enabled on client computers to which this policy is applied. In addition to this policy setting, you must specify whether the client computers are hosted cache mode or distributed cache mode clients. To do so, configure one of the following the policy settings:
  - Set BranchCache Distributed Cache mode
  - Set BranchCache Hosted Cache mode
  - Configure Hosted Cache Servers
- **Set BranchCache Distributed Cache Mode**. This policy setting specifies whether BranchCache distributed cache mode is enabled on client computers to which this policy is applied. In addition to this policy, you must use the policy “Turn on BranchCache” to enable BranchCache on client computers. In distributed cache mode, client computers download content from BranchCache-enabled main office content servers, cache the content locally, and serve the content to other BranchCache distributed cache mode clients in the branch office.

- **Set BranchCache Hosted Cache mode**. This policy setting specifies whether BranchCache hosted cache mode is enabled on client computers to which this policy is applied. In addition to this policy, you must use the policy “Turn on BranchCache” to enable BranchCache on client computers. When a client computer is configured as a hosted cache mode client, it is able to download cached content from a hosted cache server that is located at the branch office. In addition, when the hosted cache client obtains content from a content server, the client can upload the content to the hosted cache server for access by other hosted cache clients at the branch office.
- **Configure BranchCache for network files**. This policy setting is used only when you have deployed one or more BranchCache-enabled file servers at your main office. This policy setting specifies when client computers in branch offices start caching content from file servers based on the network latency – or delay – that occurs when the clients download content from the main office over a Wide Area Network (WAN) link. When you configure a value for this setting, which is the maximum round trip network latency allowed before caching begins, clients do not cache content until the network latency reaches the specified value; when network latency is greater than the value, clients begin caching content after they receive it from the file servers.
- **Set percentage of disk space used for client computer cache**. This policy setting specifies the default percentage of total disk space that is allocated for the BranchCache disk cache on client computers. If you enable this policy setting, you can configure the percentage of total disk space to allocate for the cache. If you disable or do not configure this policy setting, the cache is set to 5 percent of the total disk space on the client computer.
- **Hash Publication for BranchCache**. This policy setting specifies whether a hash generation service generates hashes, also called content information, for data that is stored in shared folders. This policy setting must be applied to server computers that have the File Services role and both the File Server and the BranchCache for network files role services installed.

For more information, see BranchCache Settings for Windows Server 2008 R2, Windows 7 Enterprise, and Windows 7 Ultimate.

BranchCache Network Shell and Windows PowerShell Commands

In Windows Server® 2012 you can configure and manage BranchCache by using either Windows PowerShell or the Network Shell (Netsh) commands for BranchCache.

Important
The netsh commands for BranchCache have not changed since the previous operating system version. The netsh commands for BranchCache in Windows Server 2012 are identical to the commands that are provided in Windows Server® 2008 R2.

In future versions of Windows, Microsoft might remove the netsh functionality for BranchCache. Microsoft recommends that you transition to Windows PowerShell if you currently use netsh to configure and manage BranchCache.

Windows PowerShell and netsh command references are at the following locations.

- Netsh Commands for BranchCache in Windows Server 2008 R2
- BranchCache Cmdlets in Windows PowerShell for Windows Server 2012.

Tip

To view a list of Windows PowerShell commands for BranchCache at the Windows PowerShell prompt, type `Get-Command -Module BranchCache` at the Windows PowerShell prompt, and then press ENTER.

BranchCache Deployment Guide

This guide contains the following sections.

- Choosing a BranchCache Design
- Deploy BranchCache
- Additional Resources

Notes

This guide is available at the following locations.

- The Windows Server 2012 BranchCache Deployment Guide in Word format in the Microsoft Download Center at http://go.microsoft.com/fwlink/p/?LinkId=259572.
- The Windows Server 2012 BranchCache Deployment Guide in HTML format in the Windows Server 2012 Technical Library at http://go.microsoft.com/fwlink/?LinkId=259777.

BranchCache is a wide area network (WAN) bandwidth optimization technology that is included in some editions of the Windows Server® 2012 and Windows® 8 operating systems.

To optimize WAN bandwidth, BranchCache copies content from your main office content servers and caches the content at branch office locations, allowing client computers at branch offices to access the content locally rather than over the WAN.

At branch offices, content is cached either on servers that are running the BranchCache feature of Windows Server 2012 or, when no server is available in the branch office, on computers running Windows 8. After a client computer requests and receives content from the main office or cloud data center and the content is cached at the branch office, other computers at the same branch office can obtain the content locally rather than contacting the content server over the WAN link.

Benefits of deploying BranchCache

BranchCache caches file, web, and application content at branch office locations, allowing client computers to access data using the local area network (LAN) rather than accessing the content over slow WAN connections.

BranchCache reduces both WAN traffic and the time that is required for branch office users to open files on the network. BranchCache always provides users with the most recent data, and it protects the security of your content by encrypting the caches on the hosted cache server and on client computers.

What this guide provides

This deployment guide allows you to deploy BranchCache in the following modes:

- Distributed cache mode. In this mode, branch office client computers download content from the content servers in the main office and then cache the content for other computers in the same branch office. Distributed cache mode does not require a server computer in the branch office.
- Hosted cache mode. In this mode, branch office client computers download content from the content servers in the main office, and a hosted cache server retrieves the content from the clients. The hosted cache server then caches the content for other client computers.

This guide also provides instructions on how to deploy three types of content servers. Content servers contain the source content that is downloaded by branch office client computers, and one or more content server is required to deploy BranchCache in either mode. The content server types are:

- **Web server-based content servers**. These content servers send content to BranchCache client computers using the HTTP and HTTPS protocols. These content servers must be running Windows Server 2008 R2 or Windows Server 2012 versions that support BranchCache and upon which the BranchCache feature is installed.
- **BITS-based application servers**. These content servers send content to BranchCache client computers using the Background Intelligent Transfer Service (BITS). These content servers must be running Windows Server 2008 R2 or Windows Server 2012 versions that support BranchCache and upon which the BranchCache feature is installed.
- **File server-based content servers**. These content servers must be running Windows Server 2008 R2 or Windows Server 2012 versions that support BranchCache and upon which the File Services server role is installed. In addition, the **BranchCache for network files** role service of the File Services server role must be installed and configured. These content servers send content to BranchCache client computers using the Server Message Block (SMB) protocol.

BranchCache deployment requirements

Following are the requirements for deploying BranchCache by using this guide.

- **File and Web content servers** must be running either Windows Server 2008 R2 or Windows Server 2012 to provide BranchCache functionality. Windows 8 clients continue to see benefits from BranchCache when accessing content servers that are running Windows Server 2008 R2, however they are unable to make use of the new chunking and hashing technologies in Windows Server 2012.
- **Client computers** must be running Windows 8 to make use of the new deployment model and the chunking and hashing improvements.

- **Hosted cache servers** must be running Windows Server 2012 to make use of the deployment improvements and scale features described in this document. A computer that is running Windows Server 2012 that is configured as a hosted cache server can continue to serve client computers that are running Windows 7, but to do so, it must be equipped with a certificate that is suitable for Transport Layer Security (TLS), as described in the BranchCache Deployment Guide for Windows Server 2008 R2 and Windows 7.
- **An Active Directory domain** is required to take advantage of Group Policy and hosted cache automatic discovery, but a domain is not required to use BranchCache. You can configure individual computers by using Windows PowerShell. In addition, it is not required that your domain controllers are running Windows Server 2012 to utilize new BranchCache Group Policy settings; you can import the BranchCache administrative templates onto domain controllers that are running other operating systems, or you can author the group policy objects remotely on other computers that are running Windows 8 or Windows Server 2012.
- **Active Directory sites** are used to limit the scope of hosted cache servers that are automatically discovered. To automatically discover a hosted cache server, both the client and server computers must belong to the same site. BranchCache is designed to have a minimal impact on clients and servers and does not impose additional hardware requirements beyond those needed to run their respective operating systems.

BranchCache history and documentation

BranchCache was first introduced in Windows® 7 and Windows Server® 2008 R2, and is improved in Windows Server 2012 and Windows 8. For more information about BranchCache in Windows Server 2012 and Windows 8, including security information and a comprehensive list of new features, see BranchCache overview.

For more information about new BranchCache features, see What's New in BranchCache.

Note
For information about BranchCache in Windows 7 and Windows Server 2008 R2, see http://branchcache.com and BranchCache for Windows Server 2008 R2.

Choosing a BranchCache Design

You can use this guide to deploy BranchCache in the following modes and mode combinations.

- All branch offices are configured for distributed cache mode.
- All branch offices are configured for hosted cache mode and have a hosted cache server on site.
- Some branch offices are configured for distributed cache mode and some branch offices have a hosted cache server on site and are configured for hosted cache mode.

The following illustration depicts a dual mode installation, with one branch office configured for distributed cache mode and one branch office configured for hosted cache mode.

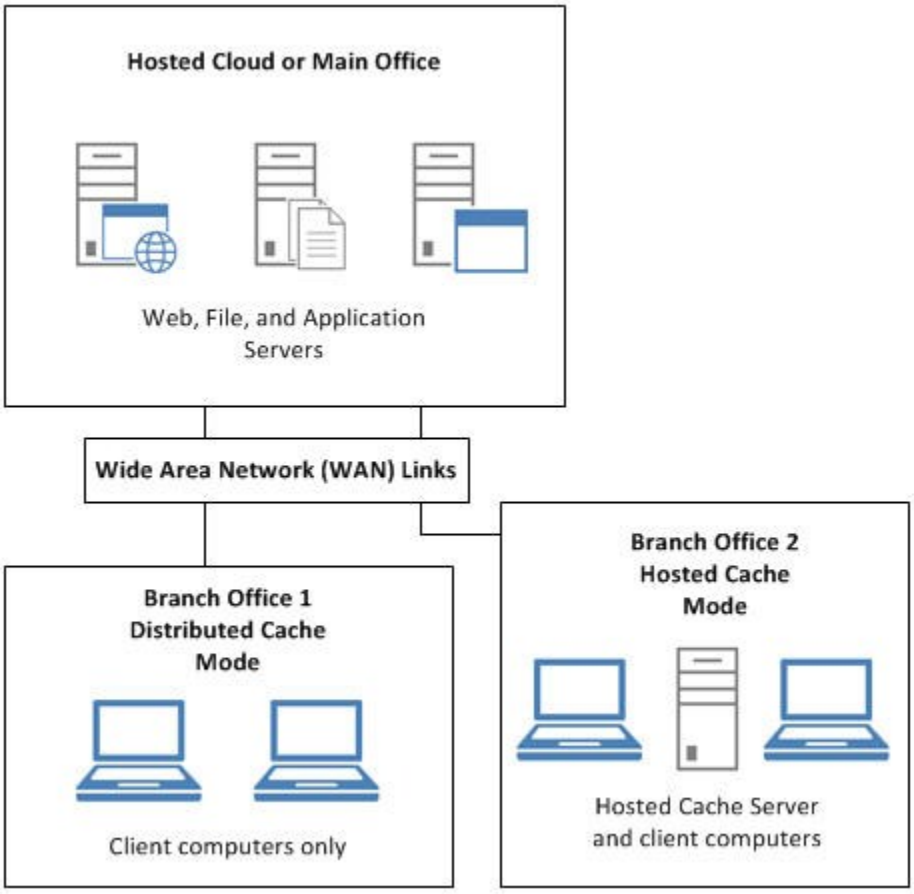

Before you deploy BranchCache, select the mode you prefer for each branch office in your organization.

Deploy BranchCache

The following sections provide information about deploying BranchCache in distributed and hosted cache modes.

- Install and Configure Content Servers
- Deploy Hosted Cache Servers (Optional)
- Prehashing and Preloading Content on Hosted Cache Servers (Optional)
- Configure BranchCache Client Computers

Note

The procedures in this guide do not include instructions for those cases in which the **User Account Control** dialog box opens to request your permission to continue. If this dialog

box opens while you are performing the procedures in this guide, and if the dialog box was opened in response to your actions, click **Continue**.

Install and Configure Content Servers

When you deploy BranchCache in distributed cache mode or hosted cache mode, you must deploy one or more content servers at your main office. Content servers that are Web servers or application servers use the BranchCache feature. Content servers that are file servers use the **BranchCache for network files** role service of the File Services server role in Windows Server 2012.

See the following topics to deploy content servers.

- Install Content Servers that Use the BranchCache Feature
- Install File Services Content Servers

Install Content Servers that Use the BranchCache Feature

To deploy content servers that are Secure Hypertext Transfer Protocol (HTTPS) 1.1 Web servers, Hypertext Transfer Protocol (HTTP) 1.1 Web servers, and Background Intelligent Transfer service (BITS)-based application servers, such as Windows Server Update Services (WSUS) and System Center Configuration Manager branch distribution site system servers, you must install the BranchCache feature, start the BranchCache service, and (for WSUS servers only) perform additional configuration steps.

See the following topics to deploy content servers.

- Install the BranchCache Feature
- Configure Windows Server Update Services (WSUS) Content Servers

Install the BranchCache Feature

You can use this procedure to install the BranchCache feature and start the BranchCache service on a computer running Windows Server 2012.

Membership in **Administrators** or equivalent is the minimum required to perform this procedure.

Before you perform this procedure, it is recommended that you install and configure your BITS-based application or Web server.

Notes

To perform this procedure by using Windows PowerShell, run Windows PowerShell as an Administrator, type the following commands at the Windows PowerShell prompt, and then press ENTER.

```
Install-WindowsFeature BranchCache

Restart-Computer
```

To install and enable the BranchCache feature

1. In Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features wizard opens. Click **Next**.
2. In **Select Installation Type**, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
3. In **Select destination server**, ensure that the correct server is selected, and then click **Next**.
4. In **Select server roles**, click **Next**.
5. In **Select features**, click **BranchCache**, and then click **Next**.
6. In **Confirm installation selections**, click **Install**. In **Installation progress**, the BranchCache feature installation proceeds. When installation is complete, click **Close**.

After you install the BranchCache feature, the BranchCache service – also called the PeerDistSvc - is enabled, and the start type is Automatic.

Configure Windows Server Update Services (WSUS) Content Servers

After installing the BranchCache feature and starting the BranchCache service, WSUS servers must be configured to store update files on the local computer. When you configure WSUS servers to store update files on the local computer, both the update metadata and the update files are downloaded by and stored directly upon the WSUS server. This ensures that BranchCache client computers receive Microsoft product update files from the WSUS server rather than directly from the Microsoft Update Web site.

Note

To learn more about WSUS server configuration, see “Advanced Synchronization Options for WSUS” on Microsoft TechNet at http://go.microsoft.com/fwlink/?LinkId=150597.

Install File Services Content Servers

To deploy content servers that are running the File Services server role, you must install the BranchCache for network files role service of the File Services server role. In addition, you must enable BranchCache on file shares according to your requirements.

During the configuration of the content server, you can allow BranchCache publication of content for all file shares or you can select a subset of file shares to publish.

Notes

When you deploy a BranchCache enabled file server or Web server as a content server, content information is now calculated offline, well before a BranchCache client requests a file. Because of this improvement, you do not need to configure hash publication for content servers, as you did in the previous version of BranchCache.

This automatic hash generation provides faster performance and more bandwidth savings, because content information is ready for the first client that requests the content, and calculations have already been performed.

See the following topics to deploy content servers.

- Configure the File Services server role
- Enable Hash Publication for File Servers
- Enable BranchCache on a File Share (Optional)

Configure the File Services server role

You can deploy BranchCache file server-based content servers on computers running Windows Server 2012 and the File Services server role with the BranchCache for network files role service installed.

- To install a BranchCache content server on a computer that does not already have File Services installed, see Install a New File Server as a Content Server.
- To install a BranchCache content server on a computer that is already configured with the File Services server role, see Configure an Existing File Server as a Content Server.

Install a New File Server as a Content Server

You can use this procedure to install the File Services server role and the BranchCache for network files role service on a computer running Windows Server® 2012.

Membership in **Administrators**, or equivalent is the minimum required to perform this procedure.

Notes

To perform this procedure by using Windows PowerShell, run Windows PowerShell as an Administrator, type the following commands at the Windows PowerShell prompt, and then press ENTER.

```
Install-WindowsFeature FS-BranchCache -IncludeManagementTools

Restart-Computer
```

To install the Data Deduplication role service, type the following command, and then press ENTER.

```
Install-WindowsFeature FS-Data-Deduplication -IncludeManagementTools
```

To install File Services and the BranchCache for network files role service

1. In Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens. In **Before You Begin**, click **Next**.
2.
3. In **Select Installation Type**, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
4. In **Select destination server**, ensure that the correct server is selected, and then click **Next**.
5. In **Select Server Roles**, in **Roles**, note that the **File And Storage Services** role is already installed; click the arrow to the left of the role name to expand the selection of role services, and then click the arrow to the left of **File and iSCSI Services**.
6. Select the check boxes for **File Server** and **BranchCache for Network Files**.

   Tip
   It is recommended that you also select the check box for **Data Deduplication**. For more information about Data Deduplication, see Additional Resources.

   Click **Next**.
7. In **Select features**, click **Next**.
8. In **Confirm Installation Selections**, review your selections, and then click **Install**. The **Installation progress** pane is displayed during installation. When installation is complete, click **Close**.

Configure an Existing File Server as a Content Server

You can use this procedure to install the BranchCache for network files role service of the File Services server role on a computer running Windows Server 2012.

Important

If the File Services server role is not already installed, do not follow this procedure. Instead, see Install a New File Server as a Content Server

Membership in **Administrators**, or equivalent is the minimum required to perform this procedure.

Notes

To perform this procedure by using Windows PowerShell, run Windows PowerShell as an Administrator, type the following commands at the Windows PowerShell prompt, and then press ENTER.

```
Install-WindowsFeature FS-BranchCache -IncludeManagementTools
```

To install the Data Deduplication role service, type the following command, and then press ENTER.

```
Install-WindowsFeature FS-Data-Deduplication -IncludeManagementTools
```

To install the BranchCache for network files role service

1. In Server Manager, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features wizard opens. Click **Next**.
2. In **Select Installation Type**, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
3. In **Select destination server**, ensure that the correct server is selected, and then click **Next**.
4. In **Select Server Roles**, in **Roles**, note that the **File And Storage Services** role is already installed; click the arrow to the left of the role name to expand the selection of role services, and then click the arrow to the left of **File and iSCSI Services**.
5. Select the check box for **BranchCache for Network Files**.

   Tip

   If you have not already done so, it is recommended that you also select the check box for **Data Deduplication**. For more information about Data Deduplication, see Additional Resources.

   Click **Next**.
6. In **Select features**, click **Next**.
7. In **Confirm Installation Selections**, review your selections, and then click **Install**. The **Installation progress** pane is displayed during installation. When installation is complete, click **Close**.

Enable Hash Publication for File Servers

You can enable BranchCache hash publication on one file server or on multiple file servers.

- To enable hash publication on one file server using local computer Group Policy, see Enable Hash Publication for Non-Domain Member File Servers.
- To enable hash publication on multiple file servers using domain Group Policy, see Enable Hash Publication for Domain Member File Servers.

Note
If you have multiple file servers and you want to enable hash publication per share, rather than enabling hash publication for all shares, you can use the instructions in the topic Enable Hash Publication for Non-Domain Member File Servers.

Enable Hash Publication for Non-Domain Member File Servers

You can use this procedure to configure hash publication for BranchCache using local computer Group Policy on a file server that is running Windows Server® 2012 with the BranchCache for network files role service of the File Services server role installed. This procedure is intended for use on a non-domain member file server. If you perform this procedure on a domain member file server and you also configure BranchCache using domain Group Policy, domain Group Policy settings override local Group Policy settings.

Membership in **Administrators**, or equivalent is the minimum required to perform this procedure.

Note
If you have one or more domain member file servers, you can add them to an organizational unit (OU) in Active Directory Domain Services and then use Group Policy to configure hash publication for all of the file servers at one time, rather than individually configuring each file server. For more information, see Enable Hash Publication for Domain Member File Servers.

To enable hash publication for one file server

1. Open Windows PowerShell, type **mmc**, and then press ENTER. The Microsoft Management Console (MMC) opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Add or Remove Snap-ins**, in **Available snap-ins**, double-click **Group Policy Object Editor**. The Group Policy Wizard opens with the Local Computer object selected. Click **Finish**, and then click **OK**.
4. In the Local Group Policy Editor MMC, expand the following path: **Computer Configuration**, **Administrative Templates**, **Network**, **Lanman Server**. Click **Lanman Server**.
5. In the details pane, double-click **Hash Publication for BranchCache**. The **Hash Publication for BranchCache** dialog box opens.

6. In the **Hash Publication for BranchCache** dialog box, click **Enabled**.
7. In **Options**, click **Allow hash publication for all shared folders**, and then click one of the following:
    a. To enable hash publication for all shared folders on this computer, click **Allow hash publication for all shared folders**.
    b. To enable hash publication only for shared folders for which BranchCache is enabled, click **Allow hash publication only for shared folders on which BranchCache is enabled**.
    c. To disallow hash publication for all shared folders on the computer even if BranchCache is enabled on the file shares, click **Disallow hash publication on all shared folders**.
8. Click **OK**.

Enable Hash Publication for Domain Member File Servers

When you're using Active Directory Domain Services (AD DS), you can use domain Group Policy to enable BranchCache hash publication for multiple file servers. To do so, you must create an organizational unit (OU), add file servers to the OU, create a BranchCache hash publication Group Policy Object (GPO), and then configure the GPO.

See the following topics to enable hash publication for multiple file servers.

- Create the BranchCache File Servers Organizational Unit
- Move File Servers to the BranchCache File Servers Organizational Unit
- Create the BranchCache Hash Publication Group Policy Object
- Configure the BranchCache Hash Publication Group Policy Object

Create the BranchCache File Servers Organizational Unit

You can use this procedure to create an organizational unit (OU) in Active Directory Domain Services (AD DS) for BranchCache file servers.

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

To create the BranchCache file servers organizational unit

1. On a computer where AD DS is installed, in Server Manager, click **Tools**, and then click

Active Directory Users and Computers. The Active Directory Users and Computers console opens.

2. In the Active Directory Users and Computers console, right-click the domain to which you want to add an OU. For example, if your domain is named example.com, right click **example.com**. Point to **New**, and then click **Organizational Unit**. The **New Object – Organizational Unit** dialog box opens.
3. In the **New Object – Organizational Unit** dialog box, in **Name**, type a name for the new OU. For example, if you want to name the OU BranchCache file servers, type **BranchCache file servers**, and then click **OK**.

Move File Servers to the BranchCache File Servers Organizational Unit

You can use this procedure to add BranchCache file servers to an organizational unit (OU) in Active Directory Domain Services (AD DS).

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

Note
You must create a BranchCache file servers OU in the Active Directory Users and Computers console before you add computer accounts to the OU with this procedure. For more information, see Create the BranchCache File Servers Organizational Unit.

To move file servers to the BranchCache file servers organizational unit

1. On a computer where AD DS is installed, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**. The Active Directory Users and Computers console opens.
2. In the Active Directory Users and Computers console, locate the computer account for a BranchCache file server, left-click to select the account, and then drag and drop the computer account on the BranchCache file servers OU that you previously created. For example, if you previously created an OU named **BranchCache file servers**, drag and drop the computer account on the **BranchCache file servers** OU.
3. Repeat the previous step for each BranchCache file server in the domain that you want to move to the OU.

Create the BranchCache Hash Publication Group Policy Object

You can use this procedure to create the BranchCache hash publication Group Policy Object (GPO).

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

Note
Before performing this procedure, you must create the BranchCache file servers organizational unit and move file servers into the OU. For more information, see Enable Hash Publication for Domain Member File Servers.

To create the BranchCache hash publication Group Policy Object

1. Open Windows PowerShell, type **mmc**, and then press ENTER. The Microsoft Management Console (MMC) opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Add or Remove Snap-ins**, in **Available snap-ins**, double-click **Group Policy Management**, and then click **OK**.
4. In the Group Policy Management MMC, expand the path to the BranchCache file servers OU that you previously created. For example, if your forest is named example.com, your domain is named example1.com, and your OU is named BranchCache file servers, expand the following path: **Group Policy Management**, **Forest: example.com**, **Domains**, **example1.com**, **BranchCache file servers**.
5. Right-click **BranchCache file servers**, and then click **Create a GPO in this domain, and Link it here**. The **New GPO** dialog box opens. In **Name**, type a name for the new GPO. For example, if you want to name the object BranchCache Hash Publication, type **BranchCache Hash Publication**. Click **OK**.

Configure the BranchCache Hash Publication Group Policy Object

You can use this procedure to configure the BranchCache hash publication Group Policy Object (GPO) so that all file servers that you added to your OU have the same hash publication policy setting applied to them.

Membership in **Domain Admins**, or equivalent is the minimum required to perform this procedure.

Note

Before performing this procedure, you must create the BranchCache file servers organizational unit, move file servers into the OU, and create the BranchCache hash publication GPO. For more information, see Enable Hash Publication for Domain Member File Servers.

To configure the BranchCache hash publication Group Policy Object

1. Run Windows PowerShell as an Administrator, type **mmc**, and then press ENTER. The Microsoft Management Console (MMC) opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Add or Remove Snap-ins**, in **Available snap-ins**, double-click **Group Policy Management**, and then click **OK**.
4. In the Group Policy Management MMC, expand the path to the BranchCache hash publication GPO that you previously created. For example, if your forest is named example.com, your domain is named example1.com, and your GPO is named **BranchCache Hash Publication**, expand the following path: **Group Policy Management**, **Forest: example.com**, **Domains**, **example1.com**, **Group Policy Objects**, **BranchCache Hash Publication**.
5. Right-click the **BranchCache Hash Publication** GPO and click **Edit**. The Group Policy Management Editor console opens.
6. In the Group Policy Management Editor console, expand the following path: **Computer Configuration**, **Policies**, **Administrative Templates**, **Network**, **Lanman Server**.
7. In the Group Policy Management Editor console, click **Lanman Server**. In the details pane, double-click **Hash Publication for BranchCache**. The **Hash Publication for BranchCache** dialog box opens.
8. In the **Hash Publication for BranchCache** dialog box, click **Enabled**.
9. In **Options**, click **Allow hash publication for all shared folders**, and then click one of the following:
    a. To enable hash publication for all shared folders for all file servers that you added to the OU, click **Allow hash publication for all shared folders**.
    b. To enable hash publication only for shared folders for which BranchCache is enabled, click **Allow hash publication only for shared folders on which BranchCache is enabled**.
    c. To disallow hash publication for all shared folders on the computer even if BranchCache is enabled on the file shares, click **Disallow hash publication on all shared folders**.
10. Click **OK**.

Note

In most cases, you must save the MMC console and refresh the view to display the configuration changes you have made.

Enable BranchCache on a File Share (Optional)

You can use this procedure to enable BranchCache on a file share.

Important
You do not need to perform this procedure if you configure the hash publication setting with the value **Allow hash publication for all shared folders**.

Membership in **Administrators**, or equivalent is the minimum required to perform this procedure.

To enable BranchCache on a file share

1. Open Windows PowerShell, type **mmc**, and then press ENTER. The Microsoft Management Console (MMC) opens.
2. In the MMC, on the **File** menu, click **Add/Remove Snap-in**. The **Add or Remove Snap-ins** dialog box opens.
3. In **Add or Remove Snap-ins**, in **Available snap-ins**, double-click **Shared Folders**. The Shared Folders Wizard opens with the Local Computer object selected. Configure the View that you prefer, click **Finish**, and then click **OK**.
4. Double-click **Shared Folders (Local)**, and then click **Shares**.
5. In the details pane, right-click a share, and then click **Properties**. The share's **Properties** dialog box opens.
6. In the **Properties** dialog box, on the **General** tab, click **Offline Settings**. The **Offline Settings** dialog box opens.
7. Ensure that **Only the files and programs that users specify are available offline** is selected, and then click **Enable BranchCache**.
8. Click **OK** twice.

Deploy Hosted Cache Servers (Optional)

You can use this procedure to install and configure BranchCache hosted cache servers that are located in branch offices where you want to deploy BranchCache hosted cache mode. With BranchCache in Windows Server 2012, you can deploy multiple hosted cache servers in one branch office.

Important
This step is optional because distributed cache mode does not require a hosted cache server computer in branch offices. If you are not planning on deploying hosted cache mode in any branch offices, you do not need to deploy a hosted cache server, and you do not need to perform the steps in this procedure.

You must be a member of **Administrators**, or equivalent to perform this procedure.

To install and configure a hosted cache server

1. On the computer that you want to configure as a hosted cache server, run the following command at a Windows PowerShell prompt to install the BranchCache feature.

   ```
   Install-WindowsFeature BranchCache -IncludeManagementTools
   ```

2. Configure the computer as a hosted cache server by using one of the following commands:
   - To configure a non-domain joined computer as a hosted cache server, type the following command at the Windows PowerShell prompt, and then press ENTER.

     ```
     Enable-BCHostedServer
     ```

   - To configure a domain joined computer as a hosted cache server, and to register a service connection point in Active Directory for automatic hosted cache server discovery by client computers, type the following command at the Windows PowerShell prompt, and then press ENTER.

     ```
     Enable-BCHostedServer -RegisterSCP
     ```

3. To verify the correct configuration of the hosted cache server, type the following command at the Windows PowerShell prompt, and then press ENTER.

   ```
   Get-BCStatus
   ```

   Note
   After you run this command, in the section **HostedCacheServerConfiguration**, the value for **HostedCacheServerIsEnabled** is **True**. If you configured a domain joined hosted cache server to register a service connection point (SCP) in Active Directory, the value for **HostedCacheScpRegistrationEnabled** is **True**.

Prehashing and Preloading Content on Hosted Cache Servers (Optional)

You can use this procedure to force the creation of content information – also called hashes - on BranchCache-enabled Web and file servers. You can also gather the data on file and web servers into packages that can be transferred to remote hosted cache servers. This provides you with the ability to preload content on remote hosted cache servers so that data is available for the first client access.

You must be a member of **Administrators**, or equivalent to perform this procedure.

To prehash content and preload the content on hosted cache servers

1. Log on to the file or Web server that contains the data that you wish to preload, and

identify the folders and files that you wish to load on one or more remote hosted cache servers.

2. Run Windows PowerShell as an Administrator. For each folder and file, run either the `Publish-BCFileContent` command or the `Publish-BCWebContent` command, depending on the type of content server, to trigger hash generation and to add data to a data package.
3. After all the data has been added to the data package, export it by using the `Export-BCCachePackage` command to produce a data package file.
4. Move the data package file to the remote hosted cache servers by using your choice of file transfer technology. FTP, SMB, HTTP, DVD and portable hard disks are all viable transports.
5. Import the data package file on the remote hosted cache servers by using the `Import-BCCachePackage` command.

Configure BranchCache Client Computers

You can use the following topics to configure domain member and non-domain member client computers as BranchCache distributed cache or hosted cache mode clients.

- Use Group Policy to Configure Domain Member Client Computers
- Use Windows PowerShell to Configure Non-Domain Member Client Computers
- Configure Firewall Rules for Non-Domain Members to Allow BranchCache Traffic
- Verify Client Computer Settings

Use Group Policy to Configure Domain Member Client Computers

You can use these procedures to create a Group Policy Object for all of the computers in your organization, to configure domain member client computers with distributed cache mode or hosted cache mode, and to configure Windows Firewall with Advanced Security to allow BranchCache traffic.

This section contains the following procedures.

1. To create a Group Policy Object and configure BranchCache modes
2. To configure Windows Firewall with Advanced Security Inbound Traffic Rules
3. To configure Windows Firewall with Advanced Security Outbound Traffic Rules

Tip

In the following procedure, you are instructed to create a Group Policy Object in the Default Domain Policy, however you can create the object in an organizational unit (OU) or other container that is appropriate for your deployment.

You must be a member of **Domain Admins**, or equivalent to perform these procedures.

To create a Group Policy Object and configure BranchCache modes

1. On a computer upon which the Active Directory Domain Services server role is installed, in Server Manager, click **Tools**, and then click **Group Policy Management**. The Group Policy Management console opens.
2. In the Group Policy Management console, expand the following path: **Forest:** *example.com*, **Domains**, *example.com*, **Group Policy Objects**, where *example.com* is the name of the domain where the BranchCache client computer accounts that you want to configure are located.
3. Right-click **Group Policy Objects**, and then click **New**. The **New GPO** dialog box opens. In **Name**, type a name for the new Group Policy Object (GPO). For example, if you want to name the object BranchCache Client Computers, type **BranchCache Client Computers**. Click **OK**.
4. In the Group Policy Management console, ensure that **Group Policy Objects** is selected, and in the details pane right-click the GPO that you just created. For example, if you named your GPO BranchCache Client Computers, right-click **BranchCache Client Computers**. Click **Edit**. The Group Policy Management Editor console opens.
5. In the Group Policy Management Editor console, expand the following path: **Computer Configuration**, **Policies**, **Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer**, **Network**, **BranchCache**.
6. Click **BranchCache**, and then in the details pane, double-click **Turn on BranchCache**. The policy setting dialog box opens.
7. In the **Turn on BranchCache** dialog box, click **Enabled**, and then click **OK**.
8. To enable BranchCache distributed cache mode, in the details pane, double-click **Set BranchCache Distributed Cache mode**. The policy setting dialog box opens.
9. In the **Set BranchCache Distributed Cache mode** dialog box, click **Enabled**, and then click **OK**.
10. If you have one or more branch offices where you are deploying BranchCache in hosted cache mode, and you have deployed hosted cache servers in those offices, double-click **Enable Automatic Hosted Cache Discovery by Service Connection Point**. The policy setting dialog box opens.
11. In the **Enable Automatic Hosted Cache Discovery by Service Connection Point** dialog box, click **Enabled**, and then click **OK**.

    Note
    When you enable both the **Set BranchCache Distributed Cache mode** and the **Enable Automatic Hosted Cache Discovery by Service Connection Point** policy settings, client computers operate in BranchCache distributed cache mode unless they find a hosted cache server in the branch office, at which point they operate in hosted cache mode.
12. Use the procedures below to configure firewall settings on client computers by using Group Policy.

To configure Windows Firewall with Advanced Security Inbound Traffic Rules

1. In the Group Policy Management console, expand the following path: **Forest:** *example.com*, **Domains**, *example.com*, **Group Policy Objects**, where *example.com* is the name of the domain where the BranchCache client computer accounts that you want to configure are located.
2. In the Group Policy Management console, ensure that **Group Policy Objects** is selected, and in the details pane right-click the BranchCache client computers GPO that you created previously. For example, if you named your GPO BranchCache Client Computers, right-click **BranchCache Client Computers**. Click **Edit**. The Group Policy Management Editor console opens.
3. In the Group Policy Management Editor console, expand the following path: **Computer Configuration**, **Policies**, **Windows Settings**, **Security Settings**, **Windows Firewall with Advanced Security**, **Windows Firewall with Advanced Security – LDAP…**, **Inbound Rules**.
4. Right-click **Inbound Rules**, and then click **New Rule**. The New Inbound Rule Wizard opens.
5. In **Rule Type**, click **Predefined**, expand the list of choices, and then click **BranchCache – Content Retrieval (Uses HTTP)**. Click **Next**.
6. In **Predefined Rules**, click **Next**.
7. In **Action**, ensure that **Allow the connection** is selected, and then click **Finish**.

   Important
   You must select **Allow the connection** for the BranchCache client to be able to receive traffic on this port.
8. To create the WS-Discovery firewall exception, again right-click **Inbound Rules**, and then click **New Rule**. The New Inbound Rule Wizard opens.
9. In **Rule Type**, click **Predefined**, expand the list of choices, and then click **BranchCache – Peer Discovery (Uses WSD)**. Click **Next**.
10. In **Predefined Rules**, click **Next**.
11. In **Action**, ensure that **Allow the connection** is selected, and then click **Finish**.

    Important
    You must select **Allow the connection** for the BranchCache client to be able to receive traffic on this port.

To configure Windows Firewall with Advanced Security Outbound Traffic Rules

1. In the Group Policy Management Editor console, right-click **Outbound Rules**, and then click **New Rule**. The New Outbound Rule Wizard opens.
2. In **Rule Type**, click **Predefined**, expand the list of choices, and then click **BranchCache – Content Retrieval (Uses HTTP)**. Click **Next**.
3. In **Predefined Rules**, click **Next**.

4. In **Action**, ensure that **Allow the connection** is selected, and then click **Finish**.

   Important
   You must select **Allow the connection** for the BranchCache client to be able to send traffic on this port.

5. To create the WS-Discovery firewall exception, again right-click **Outbound Rules**, and then click **New Rule**. The New Outbound Rule Wizard opens.
6. In **Rule Type**, click **Predefined**, expand the list of choices, and then click **BranchCache – Peer Discovery (Uses WSD)**. Click **Next**.
7. In **Predefined Rules**, click **Next**.
8. In **Action**, ensure that **Allow the connection** is selected, and then click **Finish**.

   Important
   You must select **Allow the connection** for the BranchCache client to be able to send traffic on this port.

Use Windows PowerShell to Configure Non-Domain Member Client Computers

You can use this procedure to manually configure a BranchCache client computer for distributed cache mode or hosted cache mode.

Note
If you have configured BranchCache client computers using Group Policy, the Group Policy settings override any manual configuration of client computers to which the policies are applied.

Membership in **Administrators**, or equivalent is the minimum required to perform this procedure.

To enable BranchCache distributed or hosted cache mode

1. On the BranchCache client computer that you want to configure, run Windows PowerShell as an Administrator, and then do one of the following.
   - To configure the client computer for BranchCache distributed cache mode, type the following command, and then press ENTER.

     ```
     Enable-BCDistributed
     ```

   - To configure the client computer for BranchCache hosted cache mode, type the following command, and then press ENTER.

     ```
     Enable-BCHostedClient
     ```

     Tip
     If you want to specify the available hosted cache servers, use the -

`ServerNames` parameter with a comma separated list of your hosted cache servers as the parameter value. For example, if you have two hosted cache servers named HCS1 and HCS2, configure the client computer for hosted cache mode with the following command.

```
Enable-BCHostedClient –ServerNames HCS1,HCS2
```

Configure Firewall Rules for Non-Domain Members to Allow BranchCache Traffic

You can use the information in this topic to configure third party firewall products and to manually configure a client computer with firewall rules that allow BranchCache to run in distributed cache mode.

Notes

- If you have configured BranchCache client computers using Group Policy, the Group Policy settings override any manual configuration of client computers to which the policies are applied.
- If you have deployed BranchCache with DirectAccess, you can use the settings in this topic to configure IPsec rules to allow BranchCache traffic.

Membership in **Administrators**, or equivalent is the minimum required to make these configuration changes.

[MS-PCCRD]: Peer Content Caching and Retrieval Discovery Protocol

Distributed cache clients must allow inbound and outbound MS-PCCRD traffic, which is carried in the Web Services Dynamic Discovery (WS-Discovery) protocol.

Firewall settings must allow multicast traffic in addition to inbound and outbound traffic. You can use the following settings to configure firewall exceptions for distributed cache mode.

IPv4 multicast: 239.255.255.250

IPv6 multicast: FF02::C

Inbound traffic: Local port: 3702, Remote port: ephemeral

Outbound traffic: Local port: ephemeral, Remote port: 3702

Program: %systemroot%\system32\svchost.exe (BranchCache Service [PeerDistSvc])

[MS-PCCRR]: Peer Content Caching and Retrieval: Retrieval Protocol

Distributed cache clients must allow inbound and outbound MS-PCCRR traffic, which is carried in the HTTP 1.1 protocol as documented in request for comments (RFC) 2616.

Firewall settings must allow inbound and outbound traffic. You can use the following settings to configure firewall exceptions for distributed cache mode.

Inbound traffic: Local port: 80, Remote port: ephemeral

Outbound traffic: Local port: ephemeral, Remote port: 80

Verify Client Computer Settings

You can use this procedure to verify that the client computer is correctly configured for BranchCache.

Note

This procedure includes steps for manually updating Group Policy and for restarting the BranchCache service. You do not need to perform these actions if you reboot the computer, as they will occur automatically in this circumstance.

You must be a member of **Administrators**, or equivalent to perform this procedure.

To verify BranchCache client computer settings

1. To refresh Group Policy on the client computer whose BranchCache configuration you want to verify, run Windows PowerShell as an Administrator, type the following command, and then press ENTER.

   ```
   gpupdate /force
   ```

2. For client computers that are configured in hosted cache mode and are configured to automatically discover hosted cache servers by service connection point, run the following commands to stop and restart the BranchCache service.

   ```
   net stop peerdistsvc
   ```

   ```
   net start peerdistsvc
   ```

3. Inspect the current BranchCache operational mode by running the following command.

   ```
   Get-BCStatus
   ```

4. In Windows PowerShell, review the output of the **Get-BCStatus** command.

   The value for **BranchCacheIsEnabled** should be **True**.

   In **ClientSettings**, the value for **CurrentClientMode** should be **DistributedClient** or **HostedCacheClient**, depending on the mode that you configured using this guide.

   In **ClientSettings**, if you configured hosted cache mode and provided the names of your hosted cache servers during configuration, or if the client has automatically located

hosted cache servers using service connection points, **HostedCacheServerList** should have a value that is the same as the name or names of your hosted cache servers. For example, if your hosted cache server is named HCS1 and your domain is corp.contoso.com, the value for **HostedCacheServerList** is **HCS1.corp.contoso.com**.

5. If any of the BranchCache settings listed above do not have the correct values, use the steps in this guide to verify the Group Policy or Local Computer Policy settings, as well as the firewall exceptions, that you configured, and ensure that they are correct. In addition, either restart the computer or follow the steps in this procedure to refresh Group Policy and restart the BranchCache service.

Additional Resources

For more information about the technologies mentioned in this guide, see the following resources in the Windows Server® 2012 Technical Library.

- Active Directory Domain Services Overview (http://technet.microsoft.com/library/hh831484)
- BranchCache Overview (http://technet.microsoft.com/library/hh831696)
- What's New in BranchCache (http://technet.microsoft.com/library/jj127252)
- BranchCache Cmdlets in Windows PowerShell (http://technet.microsoft.com/library/hh848392.aspx)
- Data Deduplication overview (http://technet.microsoft.com/library/hh831602)
- File and Storage Services Overview (http://technet.microsoft.com/library/hh831487)
- Group Policy Overview (http://technet.microsoft.com/library/hh831791)
- Web Server (IIS) Overview (http://technet.microsoft.com/library/hh831725)
- Windows Server Update Services Overview (http://technet.microsoft.com/library/hh852345)

Data Center Bridging (DCB) Overview

This document provides introductory information about Data Center Bridging (DCB). Links to resources with information about technologies that are closely related to DCB, or otherwise relevant to DCB are also provided.

Did you mean…

- Hyper-V Virtual Fiber Channel, see Hyper-V Virtual Fibre Channel Overview
- Hyper-V Virtual Switch, see **Hyper-V Virtual Switch Technical Preview**
- iSCSI High-Availability Block Storage, see iSCSI High-Availability Block Storage Technical Preview
- Latency mitigation, see Low Latency Workloads Technologies
- Network Bridge, see How the Windows XP Network Bridge Works

Feature description

Data Center Bridging DCB is a suite of Institute of Electrical and Electronics Engineers (IEEE) standards that enable Converged Fabrics in the data center, where storage, data networking, cluster IPC and management traffic all share the same Ethernet network infrastructure. DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control. Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support Internet Small Computer System Interface (iSCSI), Remote Direct Memory Access (RDMA) over Converged Ethernet, or Fiber Channel over Ethernet (FCoE). Priority-based flow control is essential if the upper layer protocol, such as Fiber Channel, assumes a lossless underlying transport.

Practical applications

Many enterprises have large Fiber Channel (FC) storage area network (SAN) installations for storage service. FC SAN requires special network adapters on servers and FC switches in the network. In general, FC hardware is significantly more expensive to deploy than the Ethernet hardware, which results in large capital expenditures. Additionally, having separate adapter and switch hardware to support Ethernet network traffic and FC SAN services requires additional space, power and cooling capacity in a datacenter, which results in additional, ongoing operational expenditures. From a cost perspective, it is advantageous for many enterprises to merge their FC technology with their Ethernet-based hardware solution to provide both storage and data networking services.

For the enterprises that already have a large FC SAN but want to migrate away from additional investment in the FC technology, DCB will enable them to build an Ethernet based converged fabric for both storage and data networking. A DCB converged fabric can reduce the future total cost of ownership (TCO) and simplify the management.

For hosters who have already adopted, or who plan to adopt iSCSI as their storage solution, DCB can provide hardware-assisted bandwidth reservation for iSCSI traffic to ensure performance isolation. And unlike other proprietary solutions, DCB is standard-based and therefore relatively easy to deploy and manage in a heterogeneous network.

A Windows Server® 2012-based implementation of DCB alleviates many of the issues that can occur when converged fabric solutions are provided by multiple original equipment manufacturers (OEMs). Implementations of proprietary solutions provided by multiple OEM’s might not interoperate with one another, might be difficult to manage, and will typically have different software maintenance schedules. By contrast, Windows Server® 2012 DCB is standard-based and therefore relatively easy to deploy and manage in a heterogeneous network.

Important functionality

Following is a list that summarizes the functionality that is provided by DCB.

1. Provides interoperability between DCB-capable network adapters and DCB-capable switches.
2. Provides a lossless Ethernet transport between a computer running Windows Server® 2012 and its neighbor switch by turning on priority-based flow control on the network adapter.
3. Provides the ability to allocate bandwidth to a Traffic Control (TC) by percentage, where the TC might consist of one or more classes of traffic that are differentiated by 802.1p.
4. Enables server administrators or network administrators to assign an application to a particular traffic class or priority based on well-known protocols, well-known TCP/UDP port, or NetworkDirect port used by that application.
5. Provides DCB management through Windows Server® 2012 Windows Management Instrumentation (WMI) and PowerShell.
6. Provides DCB management through Windows Server® 2012 Group Policy.
7. Supports the co-existence of Windows Server® 2012 Quality of Service (QoS) solutions.

Hardware requirements

Following is a list of hardware requirements for DCB:

1. Requires that DCB-capable Ethernet network adapter(s) are installed in computers that are providing Windows Server® 2012 DCB.
2. Requires that DCB-capable hardware switches are deployed on your network.

See also

Following is a table of resources for technologies related to DCB.

Content type	References
Product evaluation	DCB Windows PowerShell User Scripting Guide
Planning	Content not available
Deployment	Content not available
Operations	Content not available
Troubleshooting	Content not available
Security	Content not available
Tools and settings	Content not available
Community resources	Content not available
Related technologies	Content not available

DCB Windows PowerShell User Scripting Guide

The purpose of this document is to guide IT pros or datacenter administrator to configure Data Center Bridging (DCB) on a capable network adapter on Windows Server 2012 using the Windows PowerShell management tool.

DCB Overview

This guide contains the following sections:

Prerequisites

- Windows Server 2012
- Enable Data Center Bridging
- Import DCB PowerShell modules

Windows PowerShell tips

- Find Windows PowerShell commands
- Learn command syntax
- Learn more about command parameters

DCB configurations

- Traffic Class management
- Priority Flow Control settings:
- Application Priority assignment

DCBX settings

- Set the Willing bit
- Display the Willing bit

DCB configuration on network adapters

- Display DCB setting on network adapters
- Disable DCB on network adapters
- Enable DCB on network adapters

Prerequisites

Following are the prerequisites for configuring and managing DCB.

Windows Server 2012

DCB is a new feature introduced in Windows Server 2012. It is not available on Windows client operating systems. The PowerShell command output samples and the DCB configuration

examples provided in this document are from an interim Windows Server 2012 build. The PowerShell commands, known as PowerShell cmdlets, as well as the user interface, are subject to change before the official release of Windows Server 2012.

Enable Data Center Bridging

DCB includes a software control protocol called DCB Exchange Protocol (DCBX). This protocol allows network admin to configure DCB on a switch and propagates the configurations to an end device. If you prefer to manage DCB from a switch you don't need to install DCB as a feature in Windows Server 2012. You can still view the configurations in Windows Server 2012 using PowerShell though. See DCB Configurations.

Note

If you want to manage DCB on a server either because DCBX is not used or supported in your network, or because the traffic classification rules are too complicated to express in DCBX, you can use either of the following two procedures to install DCB as a feature in Windows Server 2012.

To install DCB using Windows PowerShell

1. In the Task Bar, right-click the Windows PowerShell icon, and then click **Run as Administrator**. Windows PowerShell opens in elevated mode.
2. In the Windows PowerShell console, type: **Install-WindowsFeature "data-center-bridging"**, and then press Enter.

In Windows PowerShell the command and results appear as follows:

```
PS C:\> Install-WindowsFeature "data-center-bridging"

Success  Restart  Needed Exit Code      Feature Result
---------  --------  -------------------      ----------------
True       No         Success               {Data Center Bridging}

WARNING: Windows automatic updating is not enabled. To ensure that your newly installed
role or feature is automatically updated, turn on Windows Update.

PS C:\>
```

To install DCB using the Add Roles and Features Wizard in Server Manager

1. Open Server Manager.

2. In the Server Manager Dashboard menu, click **Manage**, and then click **Add Roles and Features**. The Add Roles and Features Wizard opens.
3. In the Add Roles and Features Wizard, in **Before You Begin**, click **Next**.
4. In **Select installation type** ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
5. In **Select destination server**, ensure that **Select a server from the server pool** is selected, and in Server Pool, select the server on which you want to install DCB, and then click **Next**.
6. On **Select server roles**, click **Next**.
7. In **Select features** select **Data Center Bridging**, and then click **Install**.

Import DCB PowerShell modules

The commands to configure DCB are included in three PowerShell modules: netqos, dcbqos, and netadapter. You can import these modules as follows.

```
PS C:\> import-module netqos

PS C:\> import-module dcbqos

PS C:\> import-module netadapter

PS C:\> get-module


ModuleType       Name                                       ExportedCommands

----------       ----                                       ----------------

Manifest         dcbqos                                     {Disable-NetQosFlowControl,
Enable-NetQosFlowControl, Get-NetQosDcbxS...

Manifest         Microsoft.PowerShell.Management            {Add-Computer, Add-Content,
Checkpoint-Computer, Clear-Content...}

Manifest         Microsoft.PowerShell.Utility               {Add-Member, Add-Type, Clear-
Variable, Compare-Object...}

Manifest         netadapter                                 {Disable-NetAdapter, Disable-
NetAdapterBinding, Disable-NetAdapterChe...

Manifest         netqos                                     {Get-NetQosPolicy, New-
NetQosPolicy, Remove-NetQosPolicy, Set-NetQosP...

Script           ServerManager                              {Get-WindowsFeature, Install-
WindowsFeature, Uninstall-WindowsFeature...


PS C:\>
```

Importing the modules before any DCB configuration is optional. The only consideration is if you do not import them beforehand, you might experience a few seconds delay when invoking any command from a module for the first time.

Windows PowerShell tips

If you are familiar with Windows PowerShell you can proceed to the next section, DCB configuration. If you are not already familiar with Windows PowerShell, this following section provides several useful tips.

Find Windows PowerShell commands

There are a couple ways to find available PowerShell commands. The first one is to use the Help command.

```
PS C:\> help *qos*


Set-NetQosPolicy                Function  netqos                ...
Disable-NetQosFlowControl       Function  dcbqos                ...
Enable-NetQosFlowControl        Function  dcbqos                ...
Get-NetQosDcbxSetting           Function  dcbqos                ...
Get-NetQosFlowControl           Function  dcbqos                ...
Get-NetQosTrafficClass          Function  dcbqos                ...
New-NetQosTrafficClass          Function  dcbqos                ...
Remove-NetQosTrafficClass       Function  dcbqos                ...
Set-NetQosDcbxSetting           Function  dcbqos                ...
Set-NetQosFlowControl           Function  dcbqos                ...
Set-NetQosTrafficClass          Function  dcbqos                ...
Disable-NetAdapterQos           Function  netadapter            ...
Enable-NetAdapterQos            Function  netadapter            ...
Get-NetAdapterQos               Function  netadapter            ...
Set-NetAdapterQos               Function  netadapter            ...


PS C:\>
```

The use of keyword “QoS” and the wildcard “*” returns all of the QoS commands from the netqos, dcbqos, and netadapter modules that were previously imported.

You can use the Get-Command to find all of the commands from a specific PowerShell module.

```
PS C:\> get-command -module netqos

CommandType  Name                  ModuleName
-----------  -----                 ----------
Function     Get-NetQosPolicy      netqos
Function     New-NetQosPolicy      netqos
Function     Remove-NetQosPolicy   netqos
Function     Set-NetQosPolicy      netqos

PS C:\>
```

Learn command syntax

You can learn the syntax of a command by using the Man command.

```
PS C:\> man New-NetQoSTrafficClass

NAME
    New-NetQosTrafficClass

SYNTAX
    New-NetQosTrafficClass [-Name] <string> [-Algorithm] <Algorithm> {Strict | ETS} [-
Priority] <byte[]>
    [-BandwidthPercentage <byte>] [-CimSession <CimSession[]>] [-ThrottleLimit <int>] [-
AsJob] [-WhatIf] [-Confirm]
    [<CommonParameters>]

ALIASES
    None

REMARKS
```

```
    Get-Help cannot find the Help files for this cmdlet on this computer. It is
displaying only partial help.
        -- To download and install Help files for the module that includes this cmdlet,
use Update-Help.


PS C:\>
```

Alternately, you can add “-?” after the command. For example, **New-NetQosTrafficClass -?** which will render the same command synopsis as the Man command.

Learn more about command parameters

You can use “get-help” coupled with “-full” to learn more about the command parameters:

```
PS C:\> get-help New-NetQoSTrafficClass -full | more


NAME
    New-NetQosTrafficClass


SYNTAX
    New-NetQosTrafficClass [-Name] <string> [-Algorithm] <Algorithm> {Strict | ETS} [-
Priority] <byte[]>
    [-BandwidthPercentage <byte>] [-CimSession <CimSession[]>] [-ThrottleLimit <int>] [-
AsJob] [-WhatIf] [-Confirm]
    [<CommonParameters>]



PARAMETERS
    -Algorithm <Algorithm>

        Required?                    true
        Position?                    1
        Accept pipeline input?       false
        Parameter set name           cim:CreateInstance0
        Aliases                      tsa
        Dynamic?                     false
```

```
    -AsJob

        Required?                    false
        Position?                    Named
        Accept pipeline input?       false
        Parameter set name           cim:CreateInstance0
        Aliases                      None
        Dynamic?                     false

    -BandwidthPercentage <byte>

        Required?                    false
        Position?                    Named
        Accept pipeline input?       false
        Parameter set name           cim:CreateInstance0
        Aliases                      Bandwidth, bw
        Dynamic?                     false

    -CimSession <CimSession[]>

        Required?                    false
        Position?                    Named
        Accept pipeline input?       false
        Parameter set name           cim:CreateInstance0
        Aliases                      Session
        Dynamic?                     false

    -Confirm

        Required?                    false
        Position?                    Named
        Accept pipeline input?       false
-- More  --
```

DCB configurations

The configurations of traffic class, PFC and application priority assignment at the system level will not be applied on network adapters until you:

1. Turn the DCBX Willing bit to false (zero)
2. Enable DCB on the network adapters. See **DCB Configuration on Network Adapters**.

Note

If you want to configure DCB from a switch through DCBX, see **DCBX Settings**

The DCBX Willing bit is described in the DCB specification. If the Willing bit on a device is set to true, the device is willing to accept configurations from a remote device through DCBX. If the Willing bit on a device is set to false, the device will reject all configuration attempts from remote devices and enforce only the local configurations.

If DCB is not installed in Windows Server 2012 the value of the Willing bit is irrelevant as far as the operating system is concerned because the operating system has no local settings apply to network adapters. After DCB is installed, the default value of the Willing bit is true. This design allows network adapters to keep whatever configurations they may have received from their remote peers.

If a network adapter doesn't support DCBX it will never receive configurations from a remote device. It will receive configurations from the operating system, and only after the DCBX Willing bit is set to false.

To enforce operating system configurations of traffic class, PFC and application priority assignment on network adapters, or to simply override the configurations from remote devices—if there are any—you can run the following command:

```
PS C:\> set-netQoSdcbxSetting -Willing 0


Confirm
Are you sure you want to perform this action?
Set-NetQosDcbxSetting -Willing $false
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"): y
PS C:\>
```

Set the Willing bit @@

To display the state of the Willing bit setting, you can use the following command:

```
PS C:\> get-NetQosDcbxSetting
```

```
Willing
-------
False


PS C:\>
```

DCB Configuration on Network Adapter: Enabling DCB on a network adapter will allow you to see the configuration propagated from a switch to the network adapter.

DCB configurations in Windows Server 2012 involve a few steps:

1. Configure DCB settings at the system level, which includes:
    a. Traffic Class Management
    b. Priority Flow Control (PFC) Settings
    c. Application Priority Assignment
    d. DCBX settings
2. Configure DCB on the Network Adapter

DCB PowerShell commands have “QoS” as opposed to “DCB” in the name string. This is by design because QoS and DCB are integrated in Windows Server 2012 to provide a seamless QoS management experience.

Traffic Class management

Following are command examples for Traffic Class management:

Create Traffic Class

```
PS C:\> New-NetQosTrafficClass -name "SMB class" -priority 4 -bandwidthPercentage 30 -
Algorithm ETS


Name                      Algorithm Bandwidth(%) Priority
------                    --------- ------------ --------
SMB class                 ETS       30           4


PS C:\>
```

By default, all 802.1p values are mapped to a default traffic class, which has 100% of the bandwidth of the physical link. The above command just creates a new traffic class, to which any packet tagged with 8 Microsoft Confidential 802.1p value 4 will be mapped, and its Transmission Selection Algorithm (TSA) is ETS and has 30% of the bandwidth.

You can create up to 7 new traffic classes. Plus the default traffic class, there will be 8 traffic classes in the system at most. However, a DCB capable network adapter may not support that many traffic classes in the hardware. If you create more traffic classes than can be accommodated on a network adapter and you enable DCB on that network adapter, the miniport driver should report to theoperating systeman error, which will be event-logged.

Display traffic class

```
PS C:\> Get-NetQoSTrafficClass

Name                    Algorithm  Bandwidth(%)  Priority
------                  ---------  ------------  --------
[Default]               ETS        70            0-3,5-7
SMB class               ETS        30            4

PS C:\>
```

Modify Traffic Class

```
PS C:\> Set-NetQoSTrafficClass -Name "SMB class" -BandwidthPercentage 40
PS C:\> get-NetQosTrafficClass

Name            Algorithm  Bandwidth(%)  Priority
----            ---------  ------------  --------
[Default]       ETS        60            0-3,5-7
SMB class       ETS        40            4

PS C:\>
```

After you have a traffic class you can change its settings independently. The settings you can change include:

1. bandwidth allocation (-BandwidthPercentage)
2. TSA (-Algorithm)
3. priority mapping (-Priority)

Remove Traffic Class

```
PS C:\> Remove-NetQosTrafficClass -Name "SMB class"
PS C:\> Get-NetQosTrafficClass

Name             Algorithm  Bandwidth(%)  Priority
----             ---------  ------------  --------
[Default]        ETS        100           0-7

PS C:\>
```

After you remove a traffic class, the 802.1p value mapped to that traffic class is remapped to the default traffic class. Bandwidth is also returned back to the default traffic class.

You cannot remove the default traffic class

Priority Flow Control settings:

Following are command examples for Priority Flow Control settings:

Enable Priority Flow Control

```
PS C:\> Enable-NetQosFlowControl -priority 4
PS C:\> Get-NetQosFlowControl -priority 4

Priority   Enabled
--------   -------
4           True

PS C:\>
```

Display Priority Flow Control…

```
PS C:\> Get-NetQosFlowControl

Priority   Enabled
--------   -------
0          False
1          False
2          False
3          False
4          True
5          False
6          False
7          False

PS C:\>
```

Disable Priority Flow Control

```
PS C:\> disable-NetQosflowControl -priority 4
PS C:\> get-NetQosFlowControl -priority 4

Priority    Enabled
--------    -------
4           False

PS C:\>
```

Application Priority assignment

Following are examples of Priority assignment

Create QoS Policy

```
PS C:\> New-NetQosPolicy -Name "SMB policy" -PriorityValue8021Action 4


Name              : SMB policy
Owner             : Group Policy (Machine)
NetworkProfile    : All
Precedence        : 127
PriorityValue     : 4


PS C:\>
```

The previous command creates a new policy for SMB. –SMB is an inbox filter that matches TCP port 445 (reserved for SMB). If a packet is sent to TCP port 445 it will be tagged by theoperating systemwith 802.1p value of 4 before the packet is passed to a network miniport driver.

In addition to –SMB, other default filters include –iSCSI (matching TCP port 3260), -NFS (matching TCP port 2049), -LiveMigration (matching TCP port 6600), -FCOE (matching EtherType 0x8906) and –NetworkDirect.

NetworkDirect is an abstract layer we create on top of any RDMA implementation on a network adapter. –NetworkDirect must be followed by a Network Direct port.

In addition to the default filters, you can classify traffic by application's executable name (as in the first example below), or by IP address, port, or protocol (as shown in the second example):

By executable name

```
PS C:\> New-NetQosPolicy "background" -AppPathNameMatchcondition "c:\Program files
(x86)\backup.exe" -PriorityValue8021Action 1


Name           : background
Owner          : Group Policy (Machine)
NetworkProfile : All
Precedence     : 127
AppPathName    : c:\Program files (x86)\backup.exe
PriorityValue  : 1


PS C:\>
```

By IP address port or protocol

```
PS C:\> New-NetQosPolicy "Network Management" -IPDstPrefixMatchCondition 10.240.1.0/24 -
IPProtocolMatchCondition both -NetworkProfile all-PriorityValue8021Action 7

Name :  Network Management
Owner            :  Group Policy (Machine)
Network Profile  :  All
Precedence       :  127
IPProtocol       :  Both
IPDstPrefix      :  10.240.1.0/24
PriorityValue    :  7
```

Display QoS Policy

```
Display QoS policy
PS C:\> Get-NetQosPolicy

Name            : Network Management
Owner           : (382ACFAD-1E73-46BD-A0A-6-4EE0E587B95)
NetworkProfile  : All
Precedence      : 127
IPProtocol      : Both
IPDstPrefix     : 10.240.1.0/24
PriorityValue   : 7

Name            : SMB policy
Owner           : (382AFAD-1E73-46BD-A0A-6-4EE0E587B95)
NetworkProfile  : All
Precedence      : 127
Template        : SMB
PriorityValue   : 4
```

Modify QoS Policy

You can either modify QoS policies inline (first screen shot below) or in an object oriented fashion (second screen shot below).

Inline

```
PS C:\> Set-NetqosPolicy -Name "Network Management" -IPSrcPrefixMatchCondition
10.235.2.0/24 -IPProtocolMatchCondition both -PriorityValue 802.1Action 7
PS C:\> Get-NetQosPolicy -name "network management"

Name             : Network Management
Owner            : {382ACFD-1E73-46BD-A0A0-4EE0E587B95}
NetworkProfile   : All
Precedence       : 127
IPProtocol       : Both
IPSrcPrefix      : 10.235.2.0/24
IPDstPrefix      : 10.240.1.0/24
PriorityValue    : 7
```

Object oriented

```
PS C:\> $obj = Get-NetQosPolicy "Network Management"
PS C:\> $obj

Name            : Network Management
Owner           : (382ACFAD-1E73-46BD-A0A0-64EE0E587B95)
NetworkProfile : 127
IPProtocol      : Both
IPSrcPrefix     : 10.235.2.0/24
IPDstPrefix     : 10.240.1.0/24
PriorityValue   : 7

PS C:\> $obj.NetworkProfile = "Domain"
PS C:\> $obj.PriorityValue = 6
PS C:\> $obj | Set-NetQosPolicy
```

```
PS C:\> Get-NetQosPolicy "Network Management"


Name            : Network Management
Owner           : (382ACFAD-1E73-46BD-A0A0-64EE0E587B95)
NetworkProfile  : Domain
Precedence      : 127
IPProtocol      : Both
IPSrcPrefix     : 10.235.2.0/24
IPDstPrefix     : 10.240.1.0/24
PriorityValue   : 6
```

Remove QoS Policy

```
PS C:\> Remove-NetQosPolicy -Name "Network Management"
PS C:\>
```

DCBX settings

You can either modify QoS policies inline (first screen shot below) or in an object oriented fashion (second screen shot below).

- Turn the DCBX Willing bit to false (zero)
- Enable DCB on the network adapters (see DCB Configuration on Network Adapter)

The DCBX Willing bit is specified in the DCB specification. If the Willing bit on a device is true it means the device is willing to accept configurations from a remote device through DCBX. If the Willing bit on a device is false it means the device will reject any configuration from a remote device and enforce only the local configurations.

If DCB is not installed in Windows Server 2012 the value of the Willing bit is irrelevant as far as theoperating systemis concerned because theoperating systemhas no local settings to push down to network adapters. After DCB is installed, the default value of this Willing bit is true – this design allows network adapters to keep whatever configurations they may have received from their remote peers.

If a network adapter doesn’t support DCBX it will never receive configurations from a remote device. It will receive configurations from theoperating systemonly after the DCBX Willing bit is set to false in the operating system.

Set the Willing bit

To enforce theoperating systemconfigurations of traffic class, PFC and application priority assignment on network adapters – or simply override the configurations from remote devices, if there's any, you should run the following command:

```
PS C:\Windows\system32> Set-NetQosDcbxSetting -Willing 0


Confirm

Are you sure you want to perform this action?

Set-NetQosDcbxSetting -Willing $false

[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"): y

PS C:\Windows\system32>cls
```

Display the Willing bit

```
PS C:\> Get-NetQoSDcbxSetting


Willing

-------

False


PS C:\>
```

DCB configuration on network adapters

DCB configuration on network adapters is independent of DCB configuration at the system level described above. Regardless of whether DCB is installed in Windows Server 2012 or not, you can always run the following commands. If you configure DCB from a switch and rely on DCBX to propagate the configurations to network adapters, you can examine what configurations are received and enforced on the network adapters from theoperating systemside after you enable DCB on the network adapters.

Display DCB setting on network adapters

```
PS C:\> Get-NetAdapterQos
```

```
Network Adapter Name     : Wired Ethernet Connection 4
QoS Enabled              : True
MACsec Bypass Supported  : False
Traffic Classes (TCs)    : 8
ETS-Capable TCs          : 8
PFC-Enabled TCs          : 8
Operational TC Mappings  : TC   TSA  Bandwidth  Priorities
         --   ---  ---------  ----------
          0   ETS       100%  0-7
Operational PFC Enabled  : Priority 4
Remote TC Mappings       : TCTSA  Bandwidth  Priorities
         -----  --------   ----------
          0ETS       12%   1
          1ETS       13%   0
          2ETS       12%   2
          3ETS       13%   3
          4ETS       12%   4
          5ETS       13%   5
          6ETS       12%   6
          7ETS       13%   7
Remote PFC Enabled       : All Priorities Disabled
Remote Classifications   : Protocol   Port/Type  Priority
        --------   ---------  --------
                           Ethertype     0x8906  1
                                 TCP       3260  2
```

Disable DCB on network adapters

```
PS C:\> Disable-NetAdapterQos "Wired Ethernet Connection 4"
PS C:\> Get-NetAdapterQos "Wired Ethernet Connection 4"

Network Adapter Name    : Wired Ethernet Connection 4
```

```
QoS Enabled             : False
MACsec Bypass Supported : False
Traffic Classes (TCs)   : 0
ETS-Capable TCs         : 0
PFC-Enabled TCs         : 0
```

Enable DCB on network adapters

```
PS C:\> Enable-NetAdapterQos "Wired Ethernet Connection 4"
PS C:\>
```

More Information

For more information about DCB, see: Data Center Bridging (DCB) Overview.

Domain Name System (DNS) Overview

This topic provides overview information about the Domain Name System (DNS) Server role in Windows Server® 2012 and Windows Server® 2012 R2, and the DNS Client service in Windows Server 2012, Windows® 8, Windows Server 2012 R2, and Windows 8.1. Also see the following topics:

1. What's New in DNS Server in Windows Server 2012 R2: Provides information about new DNS Server features and functionality in Windows Server 2012 R2.
2. What's New in DNS Client in Windows 8.1: Provides information about new DNS Client features and functionality in Windows 8.1 and Windows Server 2012 R2.
3. What's New in DNS Server in Windows Server 2012: Provides information about new DNS Server features and functionality in Windows Server 2012.
4. What's New in DNS Client in Windows 8: Provides information about new DNS Client features and functionality in Windows 8 and Windows Server 2012.
5. DNSSEC in Windows Server 2012: Provides information about DNS Security Extensions support in Windows Server 2012 and Windows Server 2012 R2, including concepts and detailed deployment procedures.
6. Step-by-Step: Demonstrate DNSSEC in a Test Lab: Provides step by step instructions for deploying DNSSEC in a test lab.

7. Windows Server 2012 Virtual Labs combines IP address management (IPAM) in a virtual lab with DNSSEC and DHCP failover in Windows Server 2012. See Building a Resilient Network Infrastructure to start the virtual lab. Note: The virtual lab will take some time to start.

Role description

DNS is a system that is used in TCP/IP networks for naming computers and network services. DNS naming locates computers and services through user-friendly names. When a user enters a DNS name in an application, DNS services can resolve the name to other information that is associated with the name, such as an IP address.

Windows Server 2012 provides several enhancements to the DNS Server service. For more information about what's new in DNS Server and DNS Client for Windows Server 2012, Windows Server 2012 R2, Windows 8, and Windows 8.1, see What's New in DNS.

Practical applications

DNS in Windows Server 2012 and Windows Server 2012 R2 provides the following:

1. The DNS Server service in Windows Server 2012 has greatly enhanced support for DNS Security Extensions (DNSSEC), a suite of extensions that adds security to the DNS protocol. For more information about DNSSEC, see Overview of DNSSEC.
2. DNS integration with Active Directory is the same in Windows Server 2012 and Windows Server 2012 R2 as in previous operating systems. When DNS Server is installed on a domain controller, DNS is integrated Active Directory Domain Services (AD DS) to store and replicate DNS zones. This makes multimaster replication possible, along with more secure transmission of DNS data. In turn, AD DS requires DNS so that clients can locate domain controllers.
3. DNS and DHCP integration is the same in Windows Server 2012 and Windows Server 2012 R2 as in previous operating systems. DNS and DHCP integration enables DNS resource records to be dynamically updated for new computers and devices, or when IP addresses of existing devices change on the network.

New and changed functionality

See What's New in DNS.

Server Manager information

Installation of the DNS Server role can be performed using Server Manager. The following features and tools are installed automatically when you install DNS Server:

Feature or Tool	Description
Remote Server Administration Tools	DNS Server Tools are required to manage the DNS Server role, but do not have to be

Feature or Tool	Description
	installed on the same server. The DNS Manager console is installed automatically when you install DNS Server unless you choose to cancel installation of Remote Server Administration Tools.

Important

The DNS Server role is frequently installed with Active Directory Domain Services (AD DS). DNS Server and can be installed on a domain controller using Server Manager's Add Roles and Features Wizard, or it can be installed as an available option during promotion of a server to a domain controller. If you install AD DS and promote the server to a domain controller using Windows PowerShell, you can install DNS by specifying the **–InstallDns** parameter when using the **Install-ADDSDomainController** cmdlet.

DNS Server and AD DS role services can be installed independently or together. However, the AD DS role service is required if the DNS server will host Active Directory-integrated DNS zones.

What's New in DNS

Information about new features and capabilities for Domain Name System (DNS) services is provided in this topic. Links are provided below to information about changes in functionality in the DNS Server service and DNS Client service in Windows Server 2012, Windows 8, Windows Server 2012 R2, and Windows 8.1.

For information about changes in functionality in other versions of Windows, see What's New in DNS in Windows Server 2008 R2 and What's New in DNS in Windows Server 2008.

Role description

DNS is a system that is used in TCP/IP networks for naming computers and network services. DNS naming locates computers and services through user-friendly names. When a user enters a DNS name in an application, DNS services can resolve the name to other information that is associated with the name, such as an IP address.

The following topics describe changes to the DNS Server service and the DNS Client service.

1. What's New in DNS Server in Windows Server 2012
2. What's New in DNS Client in Windows 8
3. What's New in DNS Server in Windows Server 2012 R2
4. What's New in DNS Client in Windows 8.1

See Also

Domain Name System (DNS) Overview

What's New in DNS Server in Windows Server 2012

This topic provides information about new and changed functionality in the DNS Server service in Windows Server 2012. For information about the DNS Client service or the DNS Server service in other versions of Windows Server, see What's New in DNS.

New and changed functionality

New functionality in DNS Server for Windows Server 2012 includes the following.

Feature/functionality	New or improved	Description
Enhanced support for DNSSEC	Improved	DNS Security Extensions (DNSSEC) support is extended to include online signing and automated key management.
Windows PowerShell support	Improved	DNS configuration and management is greatly enhanced with Windows PowerShell.

Enhanced support for DNSSEC

DNSSEC support is significantly enhanced in Windows Server 2012 compared to previous versions of Windows Server. For more information about DNSSEC, see Overview of DNSSEC. Also see Step-by-Step: Demonstrate DNSSEC in a Test Lab.

Specific enhancements to DNSSEC include:

1. Support for Active Directory-integrated DNS scenarios including DNS dynamic updates in DNSSEC signed zones.
2. Support for updated DNSSEC standards, including NSEC3 and RSA/SHA-2.
3. Automated trust anchor distribution through Active Directory.
4. Automated trust anchor rollover support per RFC 5011.
5. Updated user interface with deployment and management wizards.
6. Validation of records signed with updated DNSSEC standards (NSEC3, RSA/SHA-2).

7. Easy extraction of the root trust anchor.

Windows PowerShell support

DNS configuration and management is greatly enhanced with Windows PowerShell, including:

1. Parity with the user interface and dnscmd.exe.
2. DNS Server role installation/removal using Windows PowerShell.
3. Windows PowerShell client query with DNSSEC validation results.
4. Server configuration is enabled for computers running older operating systems.

What's New in DNS Client in Windows 8

This topic provides information about new and changed functionality in the DNS Client service in Windows 8. For an overview of DNS Server service and the DNS Client service, and links to review changes in other versions of Windows, see What's New in DNS.

Note
Changes to the DNS Client service in Windows 8 are also present in computers running Windows Server 2012.

New and changed functionality

New functionality in DNS Client for Windows 8 includes the following.

Feature/functionality	New or improved	Description
Link-local multicast name resolution (LLMNR)	Improved	Outbound LLMNR queries are not sent to mobile broadband and VPN interfaces.
Network basic input/output system (NETBIOS)	Improved	Outbound NETBIOS queries are not sent to mobile broadband interfaces.
LLMNR timeout	Improved	The LLMNR query timeout has been increased to 410 msec for the 1st retry and 410 msec for the 2nd retry. The total timeout value is now 820 msec instead of 300 msec. This change is to solve a problem with computers in power saving mode. LLMNR and

Feature/functionality	New or improved	Description
		NETBIOS queries are also issued in parallel, improving response times for all queries.
Parallel queries	Improved	LLMNR and NETBIOS are issued in parallel and optimized for IPv4 and IPv6 queries.
Binding order optimization	Improved	Interfaces are divided into networks to send parallel DNS queries and prefer binding order responses.
Protocol reordering	Improved	If a specific interface is hijacking DNS names, then for flat names on those networks LLMNR and NETBIOS queries are sent in parallel with DNS queries and the LLMNR or NETBIOS response is preferred.
Asynchronous DNS cache	Improved	All the queries in DNS cache service are asynchronous and response timing is optimized.

What's New in DNS Server in Windows Server 2012 R2

This topic provides information about new and changed functionality in the DNS Server service in Windows Server 2012 R2. For information about the DNS Client service or the DNS Server service in other versions of Windows Server, see What's New in DNS.

New and changed functionality

New functionality in DNS Server for Windows Server 2012 R2 includes the following.

Feature/functionality	New or improved	Description
Enhanced zone level statistics	Improved	Zone level statistics are

Feature/functionality	New or improved	Description
		available for different resource record types, zone transfers, and dynamic updates.
Enhanced DNSSEC support	Improved	DNSSSEC key management and support for signed file-backed zones is improved.
Enhanced Windows PowerShell support	Improved	New Windows PowerShell parameters are available for DNS Server.

Enhanced zone level statistics

DNS server statistics available in Windows Server® 2012 using the **Get-DnsServerStatistics** Windows PowerShell cmdlet are: CacheStatistics, DatabaseStatistics, DnssecStatistics, DsStatistics, ErrorStatistics, MasterStatistics, MemoryStatistics, NetBiosStatistics, PacketStatistics, PrivateStatistics, Query2Statistics, QueryStatistics, RecordStatistics, RecursionStatistics, SecondaryStatistics, SecurityStatistics, TimeoutStatistics, TimeStatistics, UpdateStatistics, and WinsStatistics.

In Windows Server 2012 R2, the following additional statistics are available:

ZoneQueryStatistics: Zone query statistics provide the information about:

- QueriesFailure: The number of queries that did not result in a successful response, for example when the response is DNS SERVER FAILURE.
- QueriesNameError: The number of queries that resulted in an NXDOMAIN or EMPTY AUTH response.
- QueriesReceived: The total number of queries received for the specified record type.
- QueriesResponded: The total number of queries that resulted in a valid DNS response.

Query information is provided for the following resource record types: A, AAAA, PTR, CNAME, MX, AFSDB, ATMA, DHCID, DNAME, HINFO, ISDN, MG (Mail Group), MB (Mail Box), MINFO (Mail Box Information), NAPTR (Naming Authority Pointer), NXT (Next Domain), KEY (Public Key), MR (Renamed Mailbox), RP (Responsible Person), RT (Route Through), SRV (Service Location), SIG (Signature), TXT (Text), WKS (Well Known Services), X.25, DNSKEY, DS, NS, SOA.

ZoneTransferStatistics: Zone transfer statistics provide the information about AXFR and IXFR transations, including:

- RequestReceived: The total number of zone transfer requests received by the DNS Serverservice when operating as a primary server for a specific zone.
- RequestSent: The total number of zone transfer requests sent by the DNS Serverservice when operating as a secondary server for a specific zone.

- ResponseReceived: The total number of zone transfer requests received by the DNS Serverservice when operating as a secondary server for a specific zone.
- SuccessReceived: The total number of zone transfers received by the DNS Server servicewhen operating as a secondary server for a specific zone.
- SuccessSent: The total number of zone transfers successfully sent by the DNS Serverservice when operating as a primary server for a specific zone.

ZoneUpdateStatistics: Zone update statistics provide the information about:

- DynamicUpdateReceived: The total number of dynamic update requests received by theDNS server.
- DynamicUpdateRejected: The total number of dynamic updates rejected by the DNSserver.

To get zone level statistics, type the following commands at an elevated Windows PowerShell prompt.

```
PS C:\> $statistics = Get-DnsServerStatistics –ZoneName contoso.com

$statistics.ZoneQueryStatistics

$statistics.ZoneTransferStatistics

$statistics.ZoneUpdateStatistics
```

Enhanced DNSSEC support

Enhanced support for DNSSEC includes changes to online signing for file-backed zones, and enhanced signing key management support:

1. In Windows Server 2012 R2, the Key Master role is introduced for file-backed multi-master zones.

   Note
   The Key Master is an authoritative DNS server that generates and manages signing keys for a zone that is protected with DNSSEC. The Key Master role was introduced in Windows Server 2012 for Active Directory-integrated zones. For more information, see Overview of DNSSEC.

2. DNSSEC is enhanced to enable isolation of the key management process from primary DNS servers which are not the key masters of a zone. The entire process of signing key generation, storage, rollover, retirement, and deletion can be initiated only from the Key Master while the other primary servers can continue the zone signing by accessing these keys.

   DNSSEC key separation is accomplished by enabling generation and storage of keys on a cryptographic next-generation (CNG) compliant offline storage module.

Enhanced Windows PowerShell support

The following new Windows PowerShell cmdlets and parameters are introduced in Windows Server 2012 R2:

1. **Step-DnsServerSigningKeyRollover**: This cmdlet forces a KSK rollover when waiting for a parent delegation signer (DS) update. If a server hosting a securely delegated zone is unable to check if the DS record in parent is updated, this parameter enables you to force a rollover. It is expected that the DS record has been manually updated in the parent.
2. **Add-DnsServerTrustAnchor -Root**: The Root parameter set enables you to retrieve trust anchors from the URL specified in **RootTrustAnchorsURL** property of the DNS server. This cmdlet has the following alias: **Retrieve-DnsServerRootTrustAnchor**.
3. **RootTrustAnchorsURL**: The Get-DnsServerSetting and Set-DnsServerSetting cmdlets are extended to add a new output string of **RootTrustAnchorURL**.

What's New in DNS Client in Windows 8.1

This topic provides information about new and changed functionality in the DNS Client service in Windows 8.1. For an overview of DNS Server service and the DNS Client service, and links to review change in other versions of Windows, see What's New in DNS.

Note
Changes to the DNS Client service in Windows 8.1 are also present in computers running Windows Server 2012 R2.

New and changed functionality

New functionality in DNS Client for Windows 8.1 includes the following.

Feature/functionality	New or improved	Description
AutoDial trigger for VPN	New	This DNS feature enables Windows Store apps to use VPN auto-dialers. To accomplish this, the DNS client notifies the VPN platform that a VPN connection is needed if receives a name resolution request and the appropriate configuration parameter is present on the client. The **VPN Trigger** parameter is a setting that can be configured under the generic server settings in the name resolution policy table (NRPT).
DNS query adaptive timeout	New	This feature enables the

Feature/functionality	New or improved	Description
		timeout for DNS queries to adapt based on the time required for previous queries, reducing the timeout for most queries. Timeouts can also be increased for high-latency links, such as satellite links. Configuration of DNS timeouts is enabled on a per network interface basis, and can be optimized by Windows Store apps. Instead of waiting for 1000ms before timing out a DNS query, the first timeout is adjusted to be between 25ms and 1000ms, based on past performance of the network.
DNS server non-responsive cache	New	Nonresponsive DNS servers are cached and periodically retried. This enables the DNS client to use the best available server consistently, and to spend less time waiting for unresponsive DNS servers.
DNS cache improvements	New	The DNS cache is improved to consolidate cache entries, enable more cache entries, and to permit caching of additional records.
DNS query coalescing	New	Multiple DNS queries for the same name are combined, resulting in only one DNS query. This optimizes client, network, and server resources.
DNS SQM	New	Improvements are made to Software Quality Metrics (SQM) reporting for the DNS client. This information can be used to

Feature/functionality	New or improved	Description
		improve performance and reliability.
Events for name resolution	New	Event tracing for Windows (ETW) events are added to DNS logging. This feature will assist with troubleshooting DNS issues.
Parallelize A and AAAA queries	New	A and AAAA DNS queries are issued in parallel, saving time for interfaces that have both IPv6 and IPv4 addresses.
Per interface Winsock name resolution	New	This feature enables the **GetAddrInfoEx()** application programming interface (API) to issue a name query on a specific network interface.
Asynchronous Winsock name resolution	New	This feature enables the **GetAddrInfoEx()** API to issue asynchronous name resolution queries.
Persistent cache	New	The DNS cache is now persistent across changes that occur on the same network, including address change notifications and sleep-resume-standby state transitions.

DNSSEC in Windows Server 2012

This guide provides information to help you understand and deploy Domain Name System Security Extensions (DNSSEC) on Windows Server® 2012 and Windows Server® 2012 R2.

In this guide

The following sections are available in this guide:

Topic	Details
Overview of DNSSEC	Information and examples to help you understand how DNSSEC works.
DNSSEC in Windows	A summary of DNSSEC support in Windows.
DNSSEC Deployment Planning	Things to consider before you deploy DNSSEC.
Deploy DNSSEC with Windows Server 2012	Detailed steps to deploy DNSSEC.
Appendix A: DNSSEC Terminology	DNSSEC terminology with definitions.
Appendix B: Windows PowerShell for DNS Server	A summary of Windows PowerShell cmdlets for DNS server.

See Also

Step-by-Step: Demonstrate DNSSEC in a Test Lab

Secure DNS Deployment Guide

Domain Name System (DNS) Overview

Overview of DNSSEC

DNSSEC is a suite of extensions that add security to the DNS protocol by enabling DNS responses to be validated. Specifically, DNSSEC provides origin authority, data integrity, and authenticated denial of existence. With DNSSEC, the DNS protocol is much less susceptible to certain types of attacks, particularly DNS spoofing attacks.

The core DNSSEC extensions are specified in the following RFCs. Additional RFCs provide supporting information.

RFC 4033: DNS Security Introduction and Requirements

RFC 4034: Resource Records for the DNS Security Extensions

RFC 4035: Protocol Modifications for the DNS Security Extensions

In this section

How DNSSEC works

- DNSSEC-related resource records
- Validation of DNS responses
- Example DNS queries
- DNSSEC scenarios

How DNSSEC works

If supported by an authoritative DNS server, a DNS zone can be secured with DNSSEC using a process called zone signing. Signing a zone with DNSSEC adds validation support to a zone without changing the basic mechanism of a DNS query and response.

Validation of DNS responses occurs through the use of digital signatures that are included with DNS responses. These digital signatures are contained in new, DNSSEC-related resource records that are generated and added to the zone during zone signing. See the following figure:

Figure 1: DNS resource records in the zone secure.contoso.com before and after zone signing:

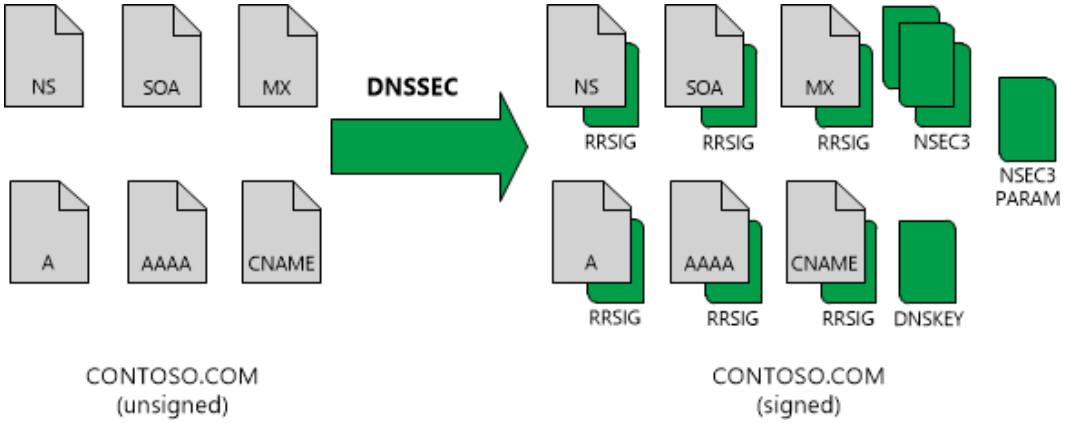

See the DNSSEC-related resource records section below for more information about each of these resource records.

When a DNSSEC-aware recursive or forwarding DNS server receives a query from a DNS client for a DNSSEC-signed zone, it will request that the authoritative DNS server also send DNSSEC records, and then attempt to validate the DNS response using these records. A recursive or forwarding DNS server knows that the zone supports DNSSEC if it has a DNSKEY (also called a trust anchor) for that zone.

Note

A non-authoritative DNS server might use recursion or forwarding to resolve a DNS query. This topic refers to the non-authoritative server as a recursive DNS server, however if the server uses forwarding then the process used for DNSSEC validation of DNS responses is the same.

DNSSEC validation

A recursive DNS server uses the DNSKEY resource record to validate responses from the authoritative DNS server by decrypting digital signatures contained in DNSSEC-related resource records and then computing and comparing hash values. If hash values are the same, it provides a reply to the DNS client with the DNS data it requested (such as an A record). If hash values are not the same, it replies with a SERVFAIL message. In this way, a DNSSEC-capable resolving

DNS server with a valid trust anchor installed protects against DNS spoofing attacks whether or not DNS clients are DNSSEC-aware.

Additionally, if the DNS client is DNSSEC-aware, it can be configured to require that the DNS server perform DNSSEC validation. See the following figure:

Figure 2: The validation process:

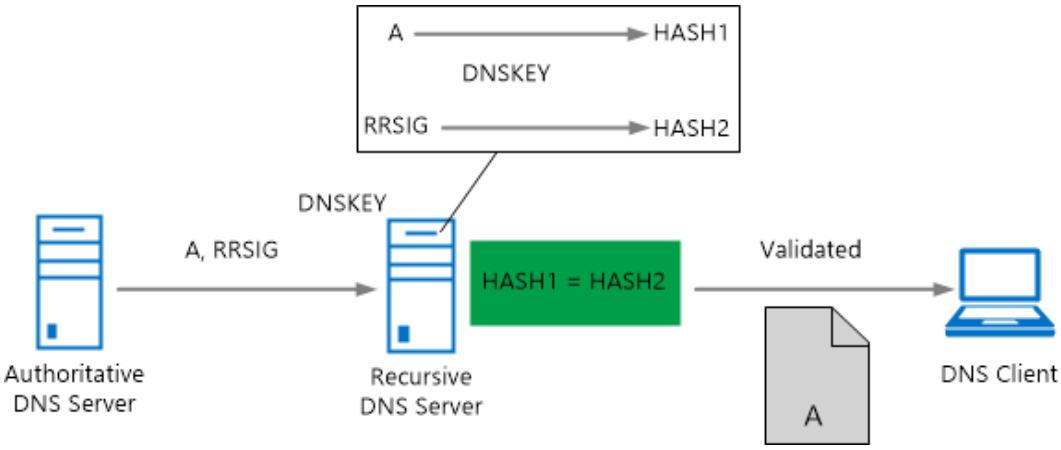

DNSKEYs are used to compute hash values and decrypt RRSIG records. The figure does not display all validation processes that are performed. Additional validation is also carried out to ensure the DNSKEYs are valid, and that DS records are valid, if they exist (not shown above).

See the Validation of DNS responses section below for additional information about how DNSSEC data can be added to the DNS query and response process.

DNSSEC-related resource records

Table 1: New resource record types used with DNSSEC:

Record Type	Description
Resource Record Signature (RRSIG)	Signatures that are generated with DNSSEC are contained in RRSIG records. Each RRSIG record is matched to another record in the zone for which it provides a digital signature. When a resolver issues a query for a name, one or more RRSIG records are returned in the response.
Next Secure (NSEC)	An NSEC record is used to prove non-existence of a DNS name. NSEC records prevent spoofing attacks intended to fool a

Record Type	Description
	DNS client into believing that a DNS name does not exist.
Next Secure 3 (NSEC3)	NSEC3 is a replacement or alternative to NSEC that has the additional benefit of preventing “zone walking” which is the process of repeating NSEC queries in order to retrieve all the names in a zone. A DNS server running Windows Server 2012 or a later operating system supports both NSEC and NSEC3. A zone can be signed with either NSEC or NSEC3, but not both.
Next Secure 3 Parameter (NSEC3PARAM)	The NSEC3PARAM record is used to determine which NSEC3 records to include in responses for non-existing DNS names.
DNS Key (DNSKEY)	A DNSKEY resource record stores a public cryptographic key that is used to verify a signature. The DNSKEY record is used by a DNS server during the validation process. DNSKEY records can store public keys for a zone signing key (ZSK) or a key signing key (KSK).
Delegation Signer (DS)	A DS record is a DNSSEC record type used to secure a delegation. DS records are used to build authentication chains to child zones.

Addition of DNSSEC-related resource records

With the exception of the DS record, all of these records are added to a zone automatically when it is signed with DNSSEC. The DS record is a special record that can be manually added to a parent zone to create a secure delegation for a child zone. For example, the contoso.com zone can contain a DS record for secure.contoso.com; however this record must either be created in the parent zone, or created in a child zone and then propagated to the parent zone. The DS record is not automatically created when you sign a zone.

NSEC or NSEC3 records are automatically added to a zone during zone signing. However, a signed zone cannot have both NSEC and NSEC3 records. The type of record (NSEC or NSEC3) added to the zone depends on how zone signing is configured. In the previous example, the zone is signed using NSEC3.

Trust anchors

DNSKEY and DS record are also called **trust anchors** or **trust points**. A trust anchor must be distributed to all non-authoritative DNS servers that will perform DNSSEC validation of DNS responses for a signed zone. If the DNS server is running on a domain controller, trust anchors are stored in the forest directory partition in Active Directory Domain Services (AD DS) and can be replicated to all domain controllers in the forest. On standalone DNS servers, trust anchors are stored in a file named **TrustAnchors.dns**. A DNS server running Windows Server 2012 or a later operating system also displays configured trust anchors in the DNS Manager console tree in the **Trust Points** container. You can also use Windows PowerShell or dnscmd.exe to view trust anchors (note: dnscmd.exe is deprecated and might be removed in a future version of Windows Server).

An example of the Windows PowerShell cmdlet **Get-DnsServerTrustAnchor** is shown below:

```
PS C:\> Get-DnsServerTrustAnchor -Name secure.contoso.com


TrustAnchorName                TrustAnchorType      TrustAnchorState     TrustAnchorData
---------------                ---------------      ----------------     ---------------
secure.contoso.com.            DNSKEY               Valid
[15952][DnsSec][RsaSha256][AwEAAe3JOsLYe17k...
secure.contoso.com.            DNSKEY               Valid
[38431][DnsSec][RsaSha256][AwEAAdsXYyqxjwBc...
```

To view all the current trust points on a server, you can use the **Get-DnsServerTrustPoint** cmdlet:

```
PS C:\> Get-DnsServerTrustPoint


TrustPointName                          TrustPointState      LastActiveRefreshTime
NextActiveRefreshTime
--------------                          ---------------      ---------------------     -
--------------------
secure.contoso.com.                     Active               10/11/2013 4:44:21 PM
10/15/2013 11:22:27 AM
```

How DNSSEC uses these new resource records to secure DNS responses is discussed below.

Validation of DNS responses

The following figure and table provides a simplified illustration of how DNSSEC can be incorporated into the DNS query and response process to provide validation.

In the example, a DNS client computer queries a recursive (caching) DNS server, which in turn queries authoritative DNS servers before returning a response. This example assumes that DNS data is not yet cached on the client or server. If a zone is signed with DNSSEC, and if DNS

servers and clients are DNSSEC-aware, then DNSSEC data can be used to validate that DNS responses are genuine.

Figure 3: The recursive DNS query process

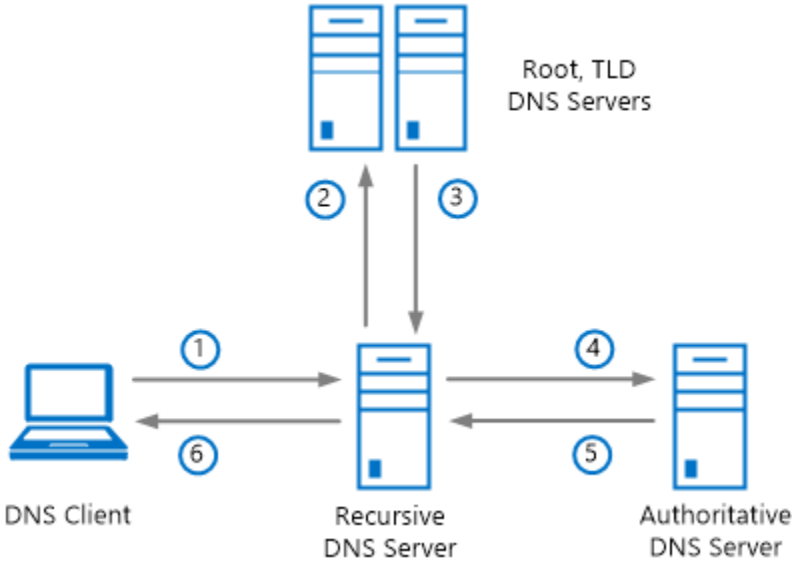

Table 2: Steps in a DNS query and response with optional DNSSEC data:

Step	Query-Response	Optional DNSSEC Data
1	A DNS client sends a DNS query to a recursive DNS server.	The DNS client can indicate it is DNSSEC-aware (DO=1).
2	The recursive DNS server sends a DNS query to the root and TLD DNS servers.	The recursive DNS server can indicate it is DNSSEC-aware (DO=1).
3	The root and TLD servers return a DNS response to the recursive DNS server providing the IP address of the authoritative DNS server for the zone.	Authoritative servers for the parent zone can indicate that the child zone is signed using DNSSEC and include a secure delegation (DS record).
4	The recursive DNS server sends a DNS query to the authoritative DNS server for the zone.	The recursive DNS server can indicate it is DNSSEC-aware (DO=1) and capable of validating signed resource records (CD=1) to be sent in the

Step	Query-Response	Optional DNSSEC Data
		response.
5	The authoritative DNS server returns a DNS response to the recursive DNS server, providing the resource record data.	The authoritative DNS server can include DNSSEC signed resource record (RRSIG record) in the DNS response, for use in validation.
6	The recursive DNS server returns a DNS response to the DNS client, providing the resource record data.	The recursive DNS server can indicate whether or not the DNS response was validated (AD=1) using DNSSEC.

Including DNSSEC data:

Three important DNSSEC-related flags (bits) mentioned in the previous table and described below are used in a DNS query and response to determine whether or not DNSSEC data is included and whether or not validation was performed. These flags are set by turning on or turning off extended data bits in the DNS packet header. When these flags are turned on, this is referred to as “setting” the bit which corresponds to a value of one (1). Turning a flag off is referred to as “clearing” the bit and corresponds to a value of zero (0).

- **DO**: The DO bit is included in a DNS query and is an abbreviation for “DNSSEC OK.” If the DO bit is set (DO=1), then the client is DNSSEC-aware and it is OK for the DNS server to return DNSSEC data in a response. If the DO bit is not set (DO=0), then the client is not DNSSEC-aware and the DNS server must not include any DNSSEC data in a DNS response. DNS clients can still be protected by DNSSEC even if they are not DNSSEC-aware. In this context, a DNS client is any computer that sends a DNS query. When a recursive DNS server sends a query to the authoritative DNS server, the recursive DNS server (computer) must indicate that it is DNSSEC aware so that the authoritative DNS server will send DNSSEC data in the response.
- **AD**: The AD bit is included in a DNS response and is an abbreviation for “authenticated data.” If the AD bit is set (AD=1), then it means the DNS response is authentic because it was validated using DNSSEC. A non-validating DNSSEC-aware computer, such as one running Windows 8, does not perform DNSSEC validation but can be configured to require that DNS responses are authentic. If the AD bit is not set (AD=0), then the DNS response was not validated, either because validation was not attempted, or because validation failed.
- **CD**: The CD bit is included in a DNS query and is an abbreviation for “checking disabled.” If the CD bit is set (CD=1) in a query, then it means a DNS response should be sent whether or not validation was successfully performed. If the CD bit is not set (CD=0), then a DNS response will not be sent if validation was required and failed. If the CD bit is clear (CD=0), this essentially means “checking enabled” and DNSSEC validation can occur. The CD bit might be set (CD=1) in a query because the host is capable of performing DNSSEC

validation, such as a recursive DNS server running Windows Server 2012. A non-validating stub resolver, such as a computer running Windows 8 will always set CD=0.

A fourth important flag (bit) that can be present in a DNS packet header is the AA bit. This flag is not new with DNSSEC, but it can be used when DNSSEC is deployed:

- **AA**: The AA bit is included in a DNS response and is an abbreviation for “authoritative answer.” If the AA bit is set, it means that the DNS response is authentic because it came directly from an authoritative DNS server. A client that requires DNSSEC validation (AD=1) will accept the AA bit instead (AA=1, AD=0) if the client directly queries an authoritative server because authoritative responses do not need to be validated. It would be redundant for an authoritative server to validate its own response.

Example DNS queries

The following examples display DNS query results performed from a DNS client computer running Windows 8.1 using the **Resolve-DnsName** cmdlet. The Resolve-DnsName cmdlet was introduced in Windows Server 2012 and Windows 8 and can be used to display DNS queries that include DNSSEC data.

Important
Do not use the **nslookup** command line tool to test DNSSEC support for a zone. The **nslookup** tool uses an internal DNS client that is not DNSSEC-aware.

Example 1: In the following example, a query is sent to a recursive DNS server for an address (A) record in the signed zone **secure.contoso.com** with DO=0.

```
PS C:\> Resolve-DnsName finance.secure.contoso.com -type A -server dns1.contoso.com

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
finance.secure.contoso.com                     A      2848  Answer     192.168.0.200
```

In this example, the DO bit was not set because the dnssekok parameter was not included.

Example 2: In the following example, the DO=1 flag is set by adding –dnssecok parameter:

```
PS C:\> resolve-dnsname -name finance.secure.contoso.com -type A -server dns1.contoso.com
-dnssecok

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
finance.secure.contoso.com                     A      2848  Answer     192.168.0.200


Name        : finance.secure.contoso.com
QueryType   : RRSIG
```

```
TTL         : 2848

Section     : Answer

TypeCovered : A

Algorithm   : 8

LabelCount  : 4

OriginalTtl : 3600

Expiration  : 10/18/2013 8:23:41 PM

Signed      : 10/8/2013 7:23:41 PM

Signer      : secure.contoso.com

Signature   : {86, 229, 217, 39...}


Name      : .

QueryType : OPT

TTL       : 32768

Section   : Additional

Data      : {}
```

When DO=0, the DNS server will not send DNSSEC data in the DNS reply. When DO=1, the client indicates that it wants to receive DNSSEC data if available. Because the secure.contoso.com zone is signed, an RRSIG resource record was included with the DNS response when DO=1.

In both example 1 and example 2, validation is not required for the secure.contoso.com zone because the Name Resolution Policy Table (NRPT) is not configured to require validation.

You can use the **Get-DnsClientNrptPolicy** cmdlet to view current NRPT rules. See Get-DnsClientNrptPolicy for more information.

Example 3: In the following example, an NRPT rule is displayed for secure.contoso.com:

```
PS C:\> Get-DnsClientNrptPolicy


Namespace                       : .secure.contoso.com

QueryPolicy                     :

SecureNameQueryFallback         :

DirectAccessIPsecCARestriction  :

DirectAccessProxyName           :

DirectAccessDnsServers          :

DirectAccessEnabled             :

DirectAccessProxyType           : NoProxy
```

```
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired   : False
NameServers                      :
DnsSecIPsecCARestriction         :
DnsSecQueryIPsecEncryption       :
DnsSecQueryIPsecRequired         : False
DnsSecValidationRequired         : False
NameEncoding                     : Utf8WithoutMapping
```

In this example, the value for **DnsSecValidationRequired** is **False**. This means that DNSSEC validation is not required.

Example 4: After enabling DNSSEC validation for secure.contoso.com, the NRPT displays True for DnsSecValidationRequired. This example only displays the secure.contoso.com namespace, and the DnsSecValidationRequired parameter:

```
PS C:\> (Get-DnsClientNrptPolicy -NameSpace secure.contoso.com).DnsSecValidationRequired
True
```

If the value of **DnsSecValidationRequired** is **True**, then DNSSEC-aware client computers will always send queries with DO=1, even if the dnssekok parameter is not included.

Example 5: When DNSSEC validation is required by the NRPT, the DNSSEC OK bit is automatically set (DO=1) for DNSSEC-aware clients:

```
PS C:\> resolve-dnsname -name finance.secure.contoso.com -type A -server dns1.contoso.com

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
finance.secure.contoso.com                     A      2848  Answer     192.168.0.200

Name        : finance.secure.contoso.com
QueryType   : RRSIG
TTL         : 2848
Section     : Answer
TypeCovered : A
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 10/18/2013 8:23:41 PM
Signed      : 10/8/2013 7:23:41 PM
```

```
Signer      : secure.contoso.com
Signature   : {86, 229, 217, 39...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

In this example, an RRSIG record is sent in the DNS response in order to fulfill the validation requirements for secure.contoso.com. A valid trust anchor is also configured on the recursive DNS server (dns1.contoso.com).

If a DNS client is not DNSSEC-aware, the NRPT rule does not apply and queries are sent with DO=0 even if an NRPT rule exists that requires DNSSEC validation.

Example 6: In the following example the same query is performed as in example 5, but without a valid trust anchor configured on dns1.contoso.com.

```
PS C:\> resolve-dnsname -name finance.secure.contoso.com -type A -server dns1.contoso.com
resolve-dnsname : finance.secure.contoso.com : Unsecured DNS packet
At line:1 char:1
+ resolve-dnsname -name finance.secure.contoso.com -type A -server dns1.contoso.co ...
+ ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    + CategoryInfo          : ResourceUnavailable: (finance.secure.contoso.com:String)
[Resolve-DnsName], Win32Excepti
   on
    + FullyQualifiedErrorId :
DNS_ERROR_UNSECURE_PACKET,Microsoft.DnsClient.Commands.ResolveDnsName
```

In this example, DNS resolution fails with the message DNS_ERROR_UNSECURE_PACKET because validation is required by the NRPT but cannot be performed due to the lack of a valid trust anchor for the secure.contoso.com zone. The Resolve-DnsName cmdlet reports detailed results for the type of failure encountered. If the DNS client tried to resolve finance.secure.contoso.com in order to connect to this host (for example to share a file), this would fail.

DNSSEC scenarios

Because DNSSEC can be deployed in many different environments with unique server and client settings, it is important to understand how DNS queries and responses are affected.

Consider the following four DNSSEC-related statements. Each statement affects how DNSSEC works in a given scenario:

1. The finance.secure.contoso.com resource record is correctly signed with DNSSEC.
2. A recursive DNS server is capable of validating responses to a query for finance.secure.contoso.com.
3. A DNS client is DNSSEC-aware.
4. A DNS client is configured to require validation for all queries in the secure.contoso.com namespace.

Let's examine each of these conditions, and their consequences, in more detail:

1. **DNSSEC signing status**: Because DNSSEC signs all records in the zone, this condition refers to the state of the secure.contoso.com zone, and not just the finance.secure.contoso.com resource record. You cannot sign some records and not sign other records; therefore the DNSSEC status of finance.secure.contoso.com depends on the DNSSEC status of secure.contoso.com.

   Signed correctly: The secure.contoso.com zone can be signed in a valid manner, thus permitting finance.secure.contoso.com to be validated as genuine. To be valid, the zone must be signed with valid, unexpired keys and all required DNSSEC-related resource records must be present.

   Not signed: The secure.contoso.com zone might not be signed, in which case there will be no RRSIG record associated with finance.secure.contoso.com, and DNS responses to queries for finance.secure.contoso.com cannot be validated. If a client requires validation (condition #4 below) in this scenario, then a DNS query sent to a recursive DNS server will fail because the DNS client will not accept a non-validated response. Note: If a client directly queries an authoritative server this will not fail validation because the response is already considered authentic.

   Not signed correctly: The secure.contoso.com zone might be signed, but not in a valid manner. For example, one or more DNSSEC signing keys might be expired. After initially signing a zone, the zone must periodically be re-signed with new keys before the signing key validity period expires, in order to maintain a secure DNSSEC operational status. The validity period for DNSSEC signing keys should be short enough to maintain security but long enough to enable easy administration. DNSSEC in Windows Server 2012 and Windows Server 2012 R2 supports automatic key rollover, providing both security and ease of administration for your DNSSEC-signed zones.

2. **Validation status**: A recursive DNS server with a valid trust anchor (public cryptographic key) for the secure.contoso.com zone will be capable of validating a DNS response for the finance.secure.contoso.com resource record. A recursive DNS server must also support the DNSSEC standards and algorithms used to sign the zone.

   Can validate: If the recursive DNS server supports all cryptographic algorithms used to sign the secure.contoso.com zone, and it has a valid trust anchor that it can use to decrypt the DNSSEC signature associated with the signed resource record, then it can validate the finance.secure.contoso.com resource record as genuine.

   Cannot validate: A non-DNSSEC aware DNS server is not capable of validation. Similarly, a DNS server that does not currently have a valid trust anchor for the secure.contoso.com zone will not be capable of validating a DNS response for finance.secure.contoso.com.

Trust anchors must be updated when a zone is re-signed, for example during key rollover. DNSSEC in Windows Server 2012 and Windows Server 2012 R2 supports automatic update of trust anchors on key rollover per RFC 5011.

3. **DNSSEC-aware status**: The DNSSEC-aware status of a DNS client depends on the operating system that it is running. Note: The Windows DNS Client in Windows 7, Windows Server 2008 R2, and later operating systems are DNSSEC-aware, non-validating stub resolvers. Previous Windows operating systems are non-DNSSEC-aware. The DNS client will inform a DNS server whether or not it is DNSSEC-aware when it sends a DNS query.

   Both client and server are DNSSEC-aware: If both the client and the server are DNSSEC-aware, then the DNS response from the server to the client will include the DNSSEC authenticated data (AD) flag. If the DNS response is validated with DNSSEC, then AD=1 will be sent. If the DNS response was not validated, then AD=0 will be sent.

   The DNS server is not DNSSEC-aware: If the DNS server is not DNSSEC-aware, then no validation is performed and the AD flag is not set (AD=0) regardless of whether or not the DNS client is DNSSEC-aware.

   The client is not DNSSEC-aware: If the client is not DNSSEC-aware, the DNS server will not set the AD flag in its response to the client even if it understands DNSSEC. However, if the DNS server is DNSSEC-aware and it has a trust anchor for the zone, the server will attempt to validate the response from the authoritative server. If the validation fails, it will return a DNS server failure to the DNS client. If the validation succeeds, it will return the query results to the client. In this way, a DNSSEC-aware recursive DNS server can protect non-DNSSEC-aware DNS clients. In this scenario, if the zone is not signed, no validation is attempted and the response is returned normally to the client.

4. **Validation requirement**: This setting determines the requirement of a DNSSEC-aware DNS client for DNSSEC data (the AD flag) in a response from a DNS server. Note: In order for the DNS client to require validation, it must be DNSSEC-aware. Non-DNSSEC-aware DNS clients cannot be forced to require DNSSEC validation. If the DNS client is directly querying an authoritative DNS server, the response will always appear to be validated, even if the zone is not signed. This is because authoritative DNS servers always return authentic responses.

   Validation is required: If validation is required, the client must receive AD=1 (validation successful) or it will reject the DNS response. If validation was unsuccessful or not attempted (AD=0) then DNS resolution will fail. This is true even if the zone is not signed. Note: This only applies to queries against a recursive, non-authoritative DNS server.

   Validation is not required: If validation is not required, the client will accept a response from a non-DNSSEC-aware recursive DNS server. However, if the recursive DNS server is DNSSEC-aware and validation fails, it will return a server failover to the client even if the client does not require validation.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Deploy DNSSEC with Windows Server 2012

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

DNS Overview

Step-by-Step: Demonstrate DNSSEC in a Test Lab

DNSSEC in Windows

Support for DNSSEC in Windows Server® 2012 and Windows Server® 2012 R2 is significantly enhanced compared to previous versions of Windows. See the following topics for more information:

DNS Servers: Support for DNSSEC-signed zones on primary, authoritative DNS servers was added with Windows Server 2008 R2. However, support in Windows Server 2008 R2 was limited to offline signing of static zones and standards such as NSEC3 and RSA/SHA-2 were not supported. DNSSEC support is enhanced significantly in Windows Server 2012 and later operating systems.

DNS Clients: The DNS Client service in Windows 7 and later operating systems is DNSSEC-aware. Previous operating systems were not DNSSEC-aware. Operating systems that are DNSSEC aware can be configured to require DNSSEC validation.

DNS Zones: Signing a zone with DNSSEC protects it from spoofing attacks. Before you can sign a zone with DNSSEC, you must specify several DNSSEC options and parameters. You can specify zone signing parameters and sign a zone with Windows PowerShell, or you can use the Zone Signing Wizard that is provided in the DNS Manager console.

Trust Anchors: A trust anchor is a public cryptographic key that enables a DNS server to validate DNS responses in a namespace as genuine. If a DNS server has a trust anchor, it will automatically attempt to validate DNS responses for that namespace. Trust anchors must be updated each time a zone is signed.

The NRPT: The Name Resolution Policy Table (NRPT) is a table of namespaces and corresponding settings stored in the Windows Registry that determines the DNS client's behavior when issuing queries and processing responses. You can use the NRPT to configure security-aware DNS clients to require validation of DNS responses.

See also

Overview of DNSSEC

DNSSEC Deployment Planning

Deploy DNSSEC with Windows Server 2012

DNS Servers

If a zone is signed with DNSSEC and validation is required, DNSSEC must be supported by each DNS server that is involved in the process of a DNS query and response. This support creates a chain of trust that enables DNSSEC data to be passed from the authoritative DNS server that hosts a signed zone to other DNS servers that can verify this data and thus validate that DNS responses are genuine. DNS clients are protected against spoofing attacks when they receive this validated response.

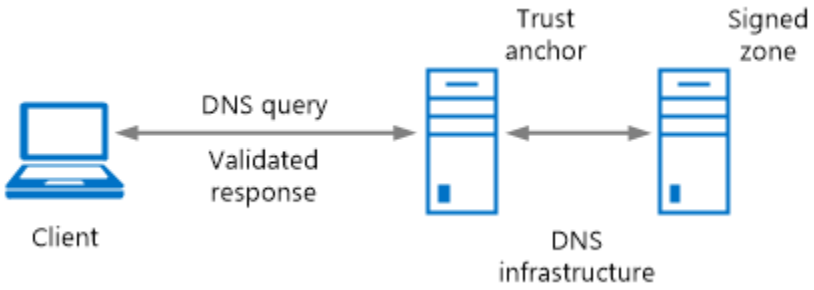

To enable DNSSEC and the chain of trust:

- Primary authoritative DNS servers must support zone signing.
- Primary and secondary authoritative DNS servers must support hosting of signed zones.
- Recursive or forwarding DNS servers must support DNSSEC validation of DNS responses.

DNS clients that are security-aware can also be configured to require that DNSSEC validation be performed. Clients that are not security-aware or do not require validation can still be protected with DNSSEC. For more information, see DNS Clients.

Note
DNSSEC requires Extension Mechanisms for DNS (EDNS0), a DNS extension that permits large UDP packets to be used in DNS responses. Some network infrastructure devices such as routers and firewalls can interfere with transmission of large UDP packets and cause DNS queries to fail when EDNS0 is enabled. For more information, see Microsoft knowledge base article 2028240 (http://support.microsoft.com/kb/2028240).

In this section

DNSSEC support in Windows Server
- DNSSEC mixed-mode deployment
- WINS lookup for signed zones

DNSSEC in Windows Server 2012
- Online signing and dynamic updates
- Read-only domain controllers (RODCs)

DNSSEC in DNS Manager
The Key Master
Transferring the Key Master role
Seizing the Key Master role
DNSSEC standards
Zone transfers

DNSSEC support in Windows Server

DNSSEC support in Windows Server is greatly improved in Windows Server 2012 and Windows Server 2012 R2. See the following table.

Table 1: Summary of DNSSEC support in Windows Server operating systems

Operating System Version	DNSSEC Support
Windows Server 2003 and Windows Server 2008	DNSSEC is implemented on secondary zones as described in RFC 2535. RFC 2535 has been made obsolete by RFCs 4033, 4034, and 4035 therefore this DNSSEC implementation is not interoperable with Windows Server 2008 R2 or any later operating systems.
Windows Server 2008 R2	Limited DNSSEC support was added in Windows Server 2008 R2 to enable offline signing for static zones. DNSSEC in Windows Server 2008 R2 is not intended to be used with dynamic, Active Directory-integrated DNS zones. If a zone is DNSSEC-signed on a DNS server running Windows Server 2008 R2, all types of dynamic updates (secure and non-secure) are disabled on that zone. To support staged migration, you can deploy DNSSEC in a mixed environment with Windows Server 2008 R2 and Windows Server 2012 DNS servers. For more information, see the DNSSEC mixed-mode deployment topic in this section.
Windows Server 2012 and Windows Server 2012 R2	DNSSEC support is significantly enhanced in Windows Server 2012 and Windows Server

Operating System Version	DNSSEC Support
	2012 R2. Enhancements include online signing, DNS dynamic update support for signed zones, Windows PowerShell support, automated trust anchor distribution and rollover, support for NSEC3 and RSA/SHA-2 standards, and updated user interface deployment and management wizards.

Note: The DNS client in Windows 7, Windows Server 2008 R2 and later versions of Windows is a non-validating DNSSEC-aware stub resolver. Earlier versions of the Windows DNS Client were not DNSSEC-aware. For more information about DNSSEC and DNS clients, see the DNS Clients topic in this guide.

For detailed information about DNSSEC in Windows Server 2008 R2, see Deploying DNS Security Extensions (DNSSEC) in the Secure DNS Deployment Guide for Windows Server 2008 R2.

DNSSEC mixed-mode deployment

To support staged migration, you can deploy DNSSEC in a mixed environment with some DNS servers running Windows Server 2008 R2 and other DNS servers running Windows Server 2012 or a later operating system.

The following specifications apply to mixed-mode deployments:

- Updated DNSSEC standards (ex: NSEC3) that are available in Windows Server 2012 and Windows Server 2012 R2 will be incompatible with DNS servers running Windows Server 2008 R2. As a result, servers running Windows Server 2008 R2 will load zones that are signed on Windows Server 2012 as unsigned zones.
- Zones that are signed on (legacy) DNS servers running Windows Server 2008 R2 are compatible with some DNSSEC operations on (current) DNS servers running Windows Server 2012 or a later operating system. DNSSEC-related resource records (ex: RRSIG, DNSKEY) in these zones will be loaded on current DNS servers and the zones will be capable of DNSSEC validation. However, the zone will not appear to be signed in DNS Manager or using current Windows PowerShell cmdlets.
- A zone that is signed on Windows Server 2008 R2 cannot be unsigned on a server running Windows Server 2012 or a later operating system using the **Invoke-DnsServerZoneUnsign** cmdlet or the DNS Manager console. To unsign a zone that was signed on a legacy DNS server, you must use legacy procedures. These procedures can be performed on a current DNS server or a legacy DNS server.

- Trust anchors that are created on legacy DNS servers can be imported and will be displayed on current DNS servers using the DNS Manager console or Windows PowerShell. The format of trust anchors is the same in both legacy and current operating systems.

Note
DNSSEC-related NRPT rules are forward and backward compatible. Rules that are created on computers running Windows Server 2008 R2 are displayed and can be modified on computers running Windows Server 2012 or Windows Server 2012 R2, and vice-versa.

Warning
An issue has been identified where DNSSEC validation can fail on computers running Windows 7 when you distribute trust anchors in Active Directory. To repair this issue, import trust anchors manually to validating DNS servers.

WINS lookup for signed zones

You can configure a DNS server to use Windows Internet Name Service (WINS) to look up names that are not found in DNS by checking the NetBIOS namespace.

To use WINS lookup integration, two special resource record types—the WINS and WINS-R resource records—are enabled and added to a zone. When the WINS resource record is used, DNS queries that fail to find a matched host (A) resource record in the zone are forwarded to WINS servers that are configured in the WINS resource record. For reverse lookup zones, the WINS-R resource record can be enabled and used to provide a similar benefit for further resolving a reverse DNS query that is not found.

In Windows Server 2008 R2, WINS lookup was disabled for DNSSEC-signed zones. For unsigned zones, the WINS tab is displayed in DNS Manager when viewing properties of the zone. See the following example:

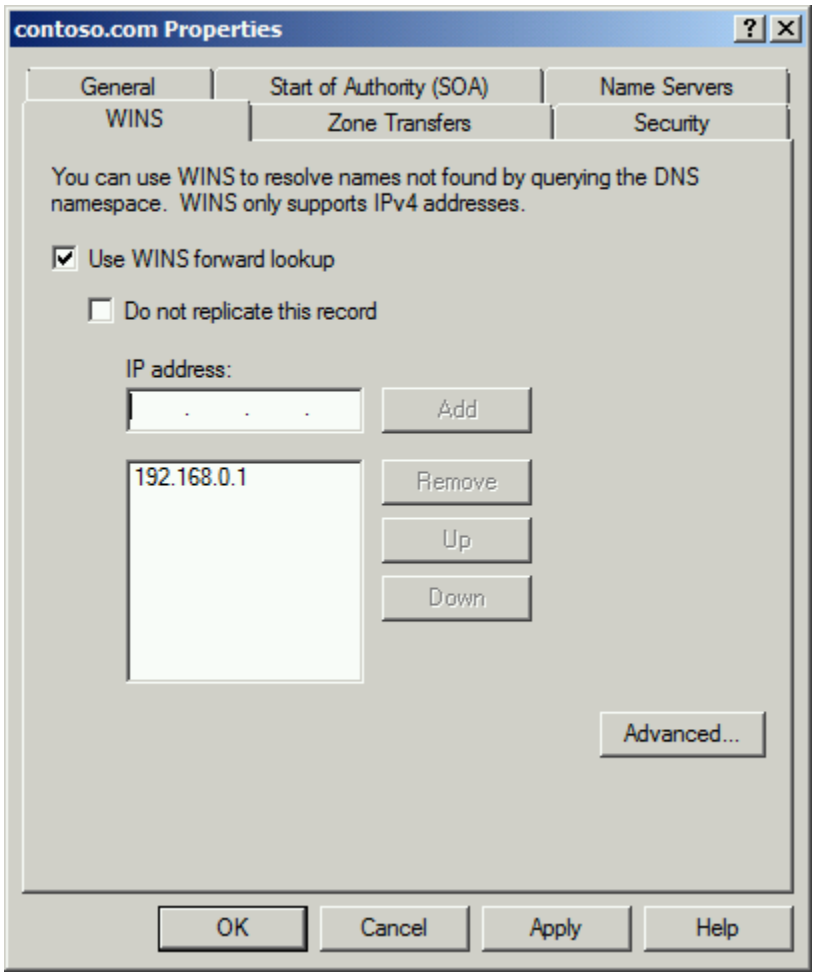

If a zone was signed, WINS lookup was disabled and the WINS tab was not displayed.

In Windows Server 2012 and later operating systems, WINS forwarding is enabled for signed zones. If you choose to enable WINS lookup on a signed zone, a warning will appear. See the following example:

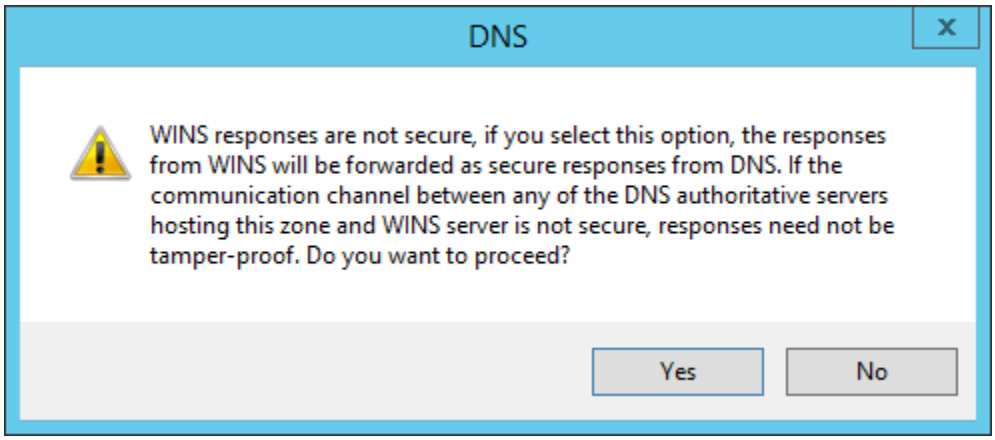

A similar message is displayed if you attempt to sign a zone that previously has had WINS lookup enabled. Warning messages are also displayed when you attempt to sign a zone with Windows PowerShell that has WINS lookup enabled.

In summary, WINS lookups are available for both signed and unsigned zones on Windows Server 2012 or later, but WINS lookup is not tamper-proof and constitutes a security risk if it is enabled for a DNSSEC-signed zone.

DNSSEC in Windows Server 2012

Windows Server 2012 and Windows Server 2012 R2 provide many enhancements over Windows Server 2008 R2 that simplify DNSSEC deployment and administration, including online signing of dynamic, Active Directory-integrated DNS zones, Windows PowerShell support, and wizards integrated into the DNS console. These key DNSSEC-related features are discussed in this topic and summarized in the table below.

Table 2: Summary of new DNSSEC features in Windows Server 2012

Feature	Description	More Information
Dynamic online signing	Online signing of dynamic, Active Directory-integrated zones is supported.	Online signing and dynamic updates
DNSSEC and read-only domain controllers (RODCs)	RODCs load signed zones as file-backed secondary zones.	Read-only domain controllers (RODCs)
DNSSEC deployment wizards	Zone signing and unsigning is greatly simplified.	DNSSEC in DNS Manager
DNSSEC administration	The Key Master is a DNS server that generates and manages signing keys for a	The Key Master

Feature	Description	More Information
	zone. DNSSEC key management is greatly enhanced with automatic key rollover and trust anchor distribution.	
Standards	Updated standards like NSEC3 and RSA/SHA-2 are supported.	DNSSEC standards
Windows PowerShell support	Windows PowerShell has parity with **dnscmd** for management of DNSSEC signed zones.	Appendix B: Windows PowerShell for DNS Server

Online signing and dynamic updates

Support for DNSSEC zone signing was added with Windows Server® 2008 R2, and subsequently extended in Windows Server 2012 to support online signing.

DNSSEC in Windows Server 2008 R2

In Windows Server 2008 R2, DNSSEC-signed zones could only be signed offline, over a file-based copy of the zone. It was not possible to generate or update signatures on a zone while the zone was online. Nearly all configuration required manual administration, including:

- You could only generate the keys required to sign a zone manually using the dnscmd tool. Although it is recommended that a zone be signed with both a KSK and a ZSK, the keys had to be generated one at a time.
- There was no built-in provision to generate additional keys automatically for rollovers. You were required to plan the type of rollover desired in advance, and then generate keys manually.
- Dnscmd required several inputs to generate keys and assumed no defaults. In some cases, you had to specify the option even though there was only one available – such as algorithm (the only supported algorithm was RSA/SHA-1).
- Zone signing was also a manual operation through dnscmd and was done over a file copy of the zone offline. You then had to manually import this signed file copy of the zone onto the server.
- Windows Server 2008 R2 supported signing Active Directory-integrated zones, but a signed zone could not have dynamic updates enabled. You were required to manually re-sign a zone whenever an update was made to the zone.
- There was no provision to replicate or distribute private keys. Keys were stored in the machine certificate store on the computer where they were generated, and could not be easily exported to other DNS servers.

DNSSEC in Windows Server 2012 and Windows Server 2012 R2

With Windows Server 2012 and later operating systems, DNSSEC support is enhanced to include online signing of dynamic zones. In addition, multiple zone management tasks are automated:

- Signing keys are generated automatically by DNS Manager or Windows PowerShell when parameters have been specified. You can also choose a default parameter set.
- For each signing key, the administrator is given the choice of enabling or disabling automatic key rollover at a specified frequency.
- DNS Manager or Windows PowerShell can be used to sign a zone using default values or using custom values. If the zone was previously signed, you can also choose to re-use these parameter values. DNS Manager also provides the option of signing a zone using the same values that were used to sign another zone.
- Zone signing is now performed on active, online zones. There is no need to take a zone offline for signing, and then import a signed copy of the zone after signing is complete.
- If a zone is updated, it is automatically re-signed. Updates to signed zones can be performed manually or dynamically. Dynamic updates are no longer disabled for DNSSEC-signed zones.
- For Active Directory-integrated zones, signing keys are automatically replicated to all primary authoritative DNS servers.

DNS in Windows Server 2012 and Windows Server 2012 R2 also includes a DNSSEC wizard in DNS Manager, which walks an administrator through the signing and unsigning process. The wizard generates all keys necessary to sign a zone automatically. For more information, see DNSSEC in DNS Manager.

For Active Directory-integrated zones, private zone signing keys replicate automatically to all primary DNS servers through Active Directory replication. Each authoritative server signs its own copy of the zone when it receives the key. For optimal performance, and to prevent increasing the size of the Active Directory database file, the signed copy of the zone remains in memory for Active Directory-integrated zones. A DNSSEC signed zone is only committed to disk for file-backed zones. Secondary DNS servers pull a full copy of the zone, including signatures, from the primary DNS server.

In general, cryptographic operations are computationally expensive. For large zones, the DNS server may take several minutes to sign the zone depending on the key length and size of the zone. To prevent performance degradation from occurring when all DNS servers start to sign the zone at the same time, signing is staggered. When a replica domain controller sees the DNSSEC keys and configuration, it will wait for a random period between 5 minutes and 30 minutes before it begins signing the zone.

Read-only domain controllers (RODCs)

In Windows Server 2008 and Windows Server 2008 R2, DNS servers running on read-only domain controllers (RODCs) host Active Directory-integrated copies of all zones. However, since the zone is read-only, the DNS server cannot make any updates to the zones it hosts. Instead, updates occur on other DNS servers and are transferred to the RODC through Active Directory replication.

When an Active Directory-integrated zone is signed with DNSSEC, private keys are also replicated to all DNS servers running on domain controllers, with an exception: Private keys are not replicated to an RODC because RODCs are intended to operate in insecure environments.

In Windows Server 2012 and Windows Server 2012 R2, an RODC will load unsigned zones from Active Directory with no change in functionality from Windows Server 2008 R2. However, if the RODC finds a DNSSEC-signed zone in Active Directory, it will not load the zone as Active Directory-integrated. Instead, it will create a secondary copy of the zone, and then configure the closest writeable domain controller for the domain as the primary server. The RODC will then attempt to perform a zone transfer. Zone transfers must be enabled on the primary DNS server for this transfer to succeed. If zone transfers are not enabled, the RODC will log an error event and will take no further action. In this scenario, you must manually enable zone transfers on the primary selected by the RODC. Alternately, you can choose to reconfigure the RODC to point to a different primary DNS server that has zone transfers enabled.

DNSSEC in DNS Manager

DNS Manager in Windows Server 2012 and Windows Server 2012 R2 includes wizards to help you sign zones with DNSSEC, to view and edit DNSSEC parameters, and to easily unsign a zone.

The DNSSEC status of a zone is displayed in the console tree with a lock icon displayed next to the name of the zone if it is signed with DNSSEC. A signed zone will also display a status of **Signed** in the right-hand pane in the **DNSSEC Status** column and the name of the The Key Master is shown in the corresponding column. See the following example:

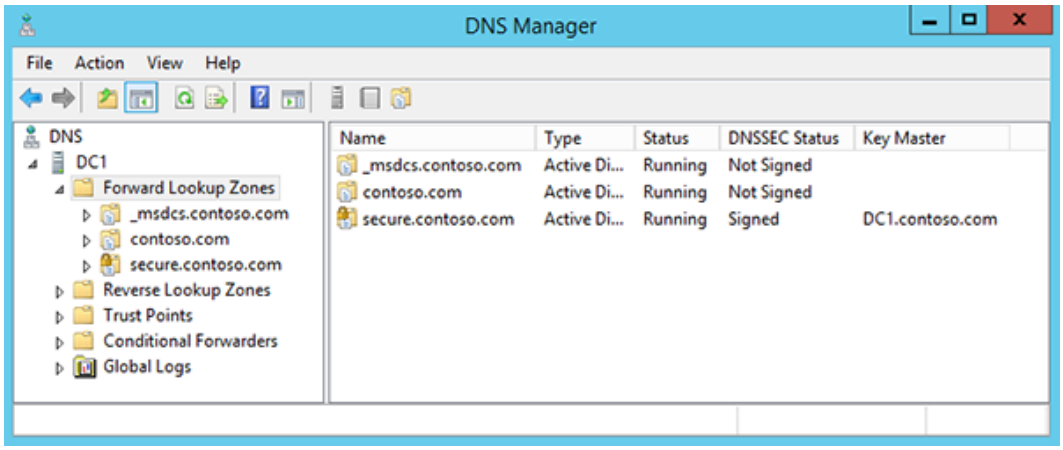

If you recently signed the zone, you might need to refresh the console view to see the DNSSEC-signed status of a zone.

You can also right-click the zone in DNS Manager and point to **DNSSEC**. If the zone is not yet signed, the only choice available is **Sign the Zone**. See the following example:

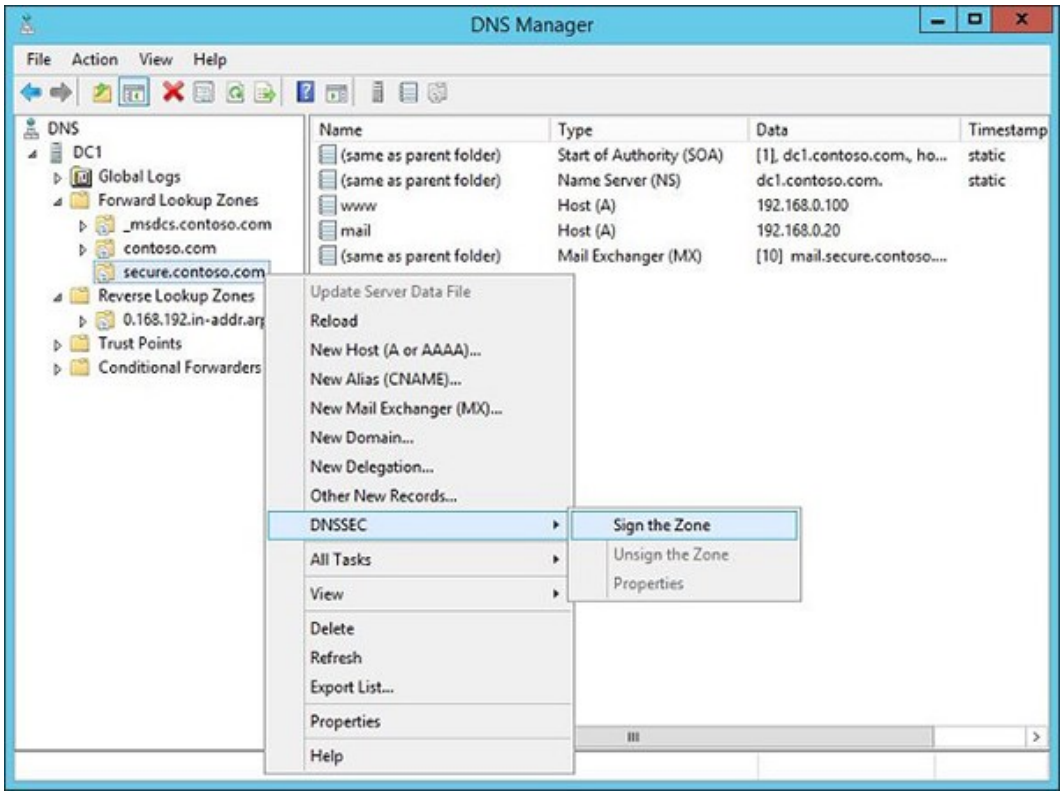

For information about signing and unsigning a zone, see DNS Zones.

The Key Master

The DNSSEC Key Master is a new concept and component of a Windows DNSSEC deployment introduced in Windows Server 2012.

In Windows DNSSEC, the Key Master is the DNS server that is responsible for key generation and key management for a DNSSEC-signed zone. When you use default settings to sign a zone, the local server will be selected as the Key Master. You also have the option of choosing a different DNS server from a list of servers that support online DNSSEC signing. Only one DNS server can be the Key Master for a given zone at a given time.

The Key Master must be a primary, authoritative server for the zone and must be capable of online zone signing. However, in a Microsoft multi-master DNS deployment environment, the following is possible:

- The Key Master role can be moved to a different authoritative name server after zone signing. This can be done gracefully if the current Key Master is online, or it can be done as part of a disaster recovery scenario if the current Key Master is offline. For more information, see Transferring the Key Master role and Seizing the Key Master role below.
- One authoritative server can be the Key Master for multiple zones.
- Different authoritative servers can be Key Masters for different zones.

You cannot move the Key Master role if a zone is file-backed, because these zones have only one primary, authoritative DNS server.

Important

If a zone is Active Directory-integrated, the Key Master will be a domain controller and can benefit by enhanced security considerations used with domain controllers. If the zone is file-backed, the Key Master might not be also a domain controller. In this scenario, it is recommended to take additional security precautions to protect the Key Master from attack and private key material from becoming compromised. For security reasons, a Key Master that is not a domain controller should only have the DNS Server role installed in order to limit its attack surface.

The name of the Key Master is displayed in DNS manager when you click on **Forward Lookup Zones** or **Reverse Lookup Zones**, and it is displayed on the **Key Master** tab of the DNSSEC properties page. You can also use the **Get-DnsServerDnsSecZoneSetting** Windows PowerShell cmdlet to view the Key Master. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com | Select
KeyMasterServer


KeyMasterServer
---------------
DC1.contoso.com
```

The Key Master generates all keys for the zone, and is responsible for distribution of private keys and zone signing information. The Key Master is also responsible for performing all zone signing key (ZSK) and key signing key (KSK) rollovers, and for polling child zones to keep signed delegations up to date.

Note

An unsigned zone can also be assigned a Key Master. All zones that have been signed will have a Key Master setting, whether or not they are currently signed. A zone that has never been signed typically does not have a Key Master, but can be configured with a Key Master in preparation for zone signing with Windows PowerShell.

Transferring the Key Master role

If the Key Master is online, you can perform a graceful transfer the Key Master role to another DNS server. Another qualifying DNS server must be available on the network.

To transfer the Key Master role using DNS Manager, view the DNSSEC properties of the zone, click the **Key Master** tab, and then choose **Use the following DNS server as the Key Master**. When you click the drop-down list, a popup will ask if you wish for the local server to build a list of available, qualifying DNS servers that can be the Key Master. Click **OK**, choose a server from the list, and then click **OK**. See the following examples:

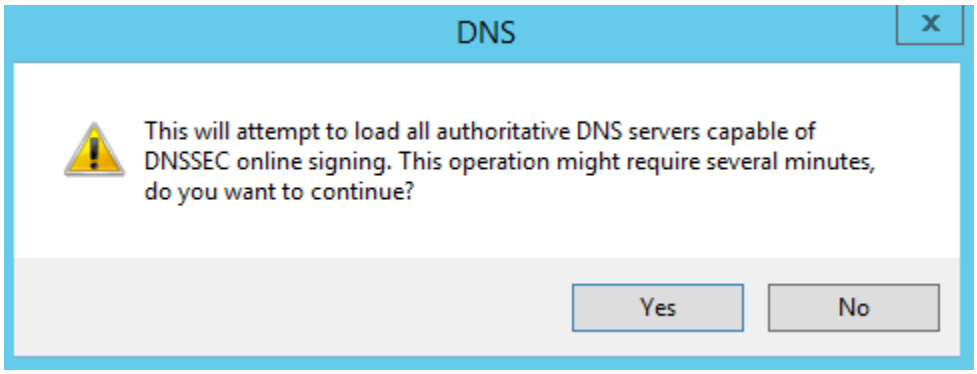
DNS
This will attempt to load all authoritative DNS servers capable of DNSSEC online signing. This operation might require several minutes, do you want to continue?
Yes
No

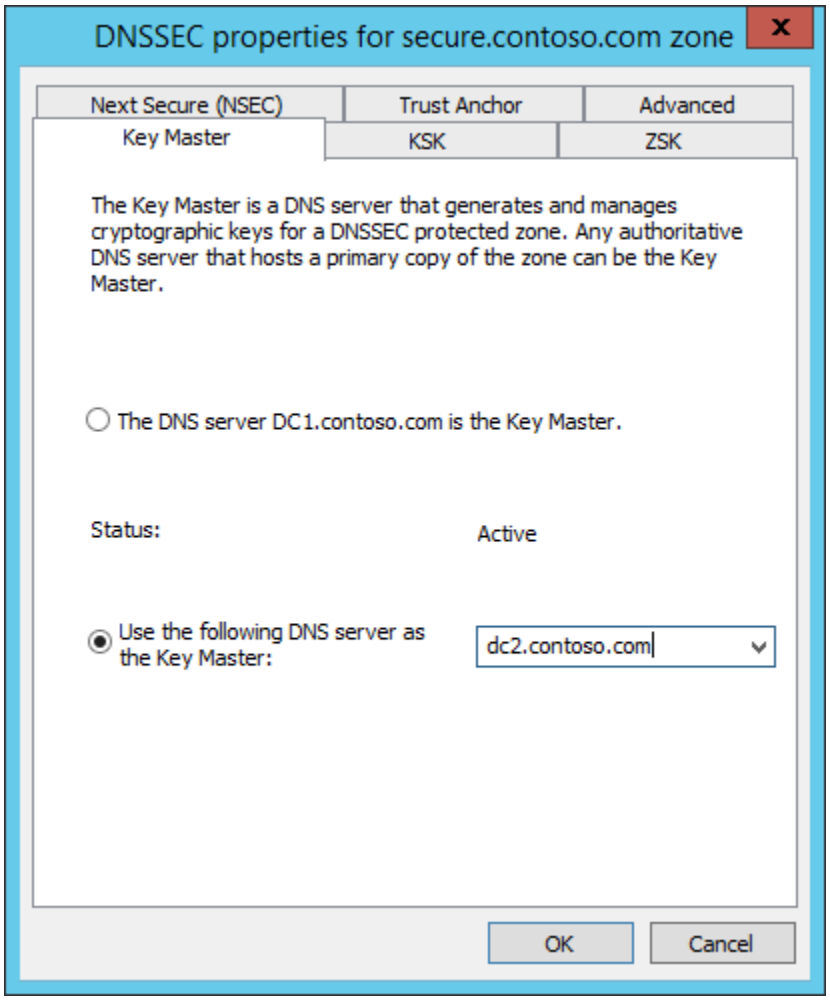
DNSSEC properties for secure.contoso.com zone
Next Secure (NSEC)
Trust Anchor
Advanced
Key Master
KSK
ZSK
The Key Master is a DNS server that generates and manages cryptographic keys for a DNSSEC protected zone. Any authoritative DNS server that hosts a primary copy of the zone can be the Key Master.
The DNS server DC1.contoso.com is the Key Master.
Status:
Active
Use the following DNS server as the Key Master:
dc2.contoso.com
OK
Cancel

Important

The signed-in user must have Domain Admin rights, or equivalent, on a DNS server in order for it to be displayed in the list.

You can also perform this operation using Windows PowerShell with the **Reset-DnsServerZoneKeyMasterRole** cmdlet. See the following example:

```
PS C:\> Reset-DnsServerZoneKeyMasterRole -ZoneName secure.contoso.com -KeyMasterServer
dc2.contoso.com -force
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com | Select
KeyMasterServer


KeyMasterServer
---------------
DC2.contoso.com
```

In this example, the **Get-DnsServerDnsSecZoneSetting cmdlet** is also used to verify that the Key Master role was transferred successfully.

Seizing the Key Master role

The server designated as the Key Master should be online and highly available to ensure uninterrupted service for key signing operations. However, if the Key Master becomes offline and cannot easily be restored to online status, it is possible to forcibly transfer the Key Master role to another qualifying DNS server, using a process called seizing. Do not seize the Key Master role unless it is unavoidable. A graceful online transfer of the Key Master role is always preferred.

Private key material

Important
The Key Master must have access to private key material for a DNSSEC-signed zone. If the current Key Master is offline, other DNS servers might have access to private key material if it is stored in a shared location such as Active Directory. If private key material is not stored in Active Directory, and the new Key Master cannot access the private keys for a zone through other means, then new keys must be generated and the zone must be re-signed with these new keys. After re-signing with new keys, all trust anchors that exist on other DNS servers will be invalid and must be updated.

To store private key material in Active Directory, select the **Replicate this private key to all DNS servers authoritative for this zone** checkbox in KSK settings for all KSKs that are used. See the following example:

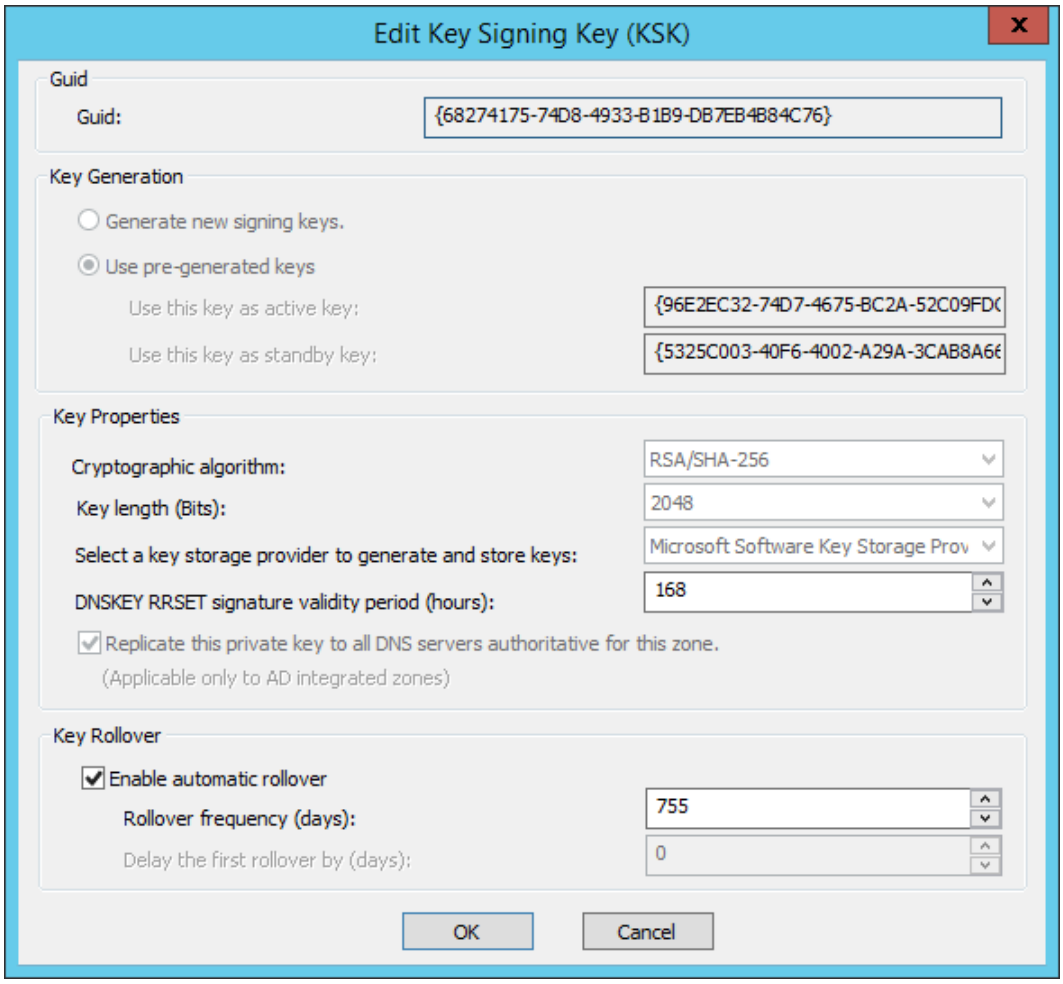

If you do not wish to store private key material in Active Directory, you might also be able to provide access to private key material using a certificate or hardware storage module (HSM) device. This storage media must be accessible to the DNS server that is selected as the new Key Master.

If private key material is not stored in Active Directory or an external device, you can use the following command to store this material in a certificate on the local computer:

```
MakeCert -ss MS-DNSSEC -sr LocalMachine
```

To use the MakeCert command, you must first download and install the Windows SDK from http://go.microsoft.com/fwlink/p/?linkid=84091.

If the new Key Master cannot access private key material for the zone, a notification is displayed indicating that private key material is not accessible and that new keys must be generated. See the following example:

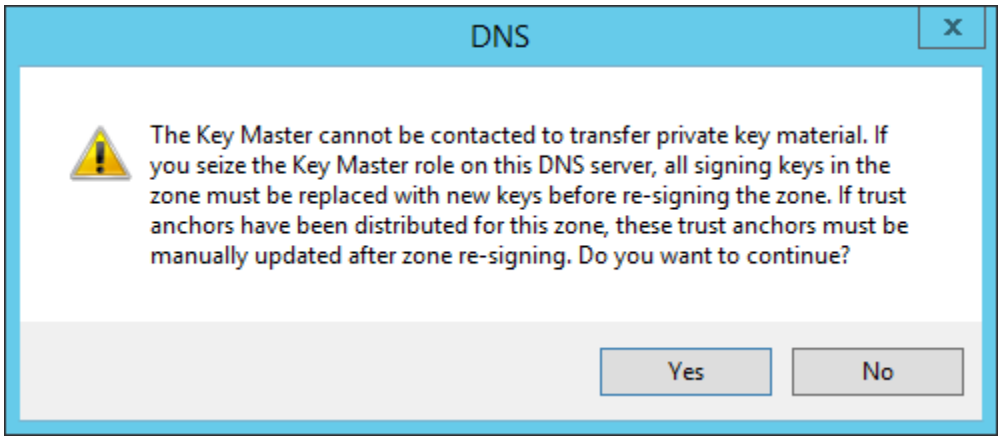

Important

If you click **Yes** and continue seizing the Key Master role, new ZSKs and KSKs are not automatically generated. You must re-sign the zone so that private key material is available. If trust anchors were distributed for the zone, these must also be replaced. If the original Key Master becomes available before the zone is re-signed, you can move the Key Master role back to this server without the requirement to re-sign the zone and redistribute trust anchors.

If private key material is stored in Active Directory, you can seize the Key Master role on another primary, authoritative Active Directory-integrated DNS server and have full access to private key material. In this case, signing keys (ZSKs and KSKs) do not need to be replaced.

If you use DNS Manager to access an Active Directory-integrated authoritative DNS server with a primary copy of the zone when the Key Master is offline, a notification will be displayed when you view DNSSEC properties of the zone indicating that DNSSEC settings could not be loaded. See the following example:

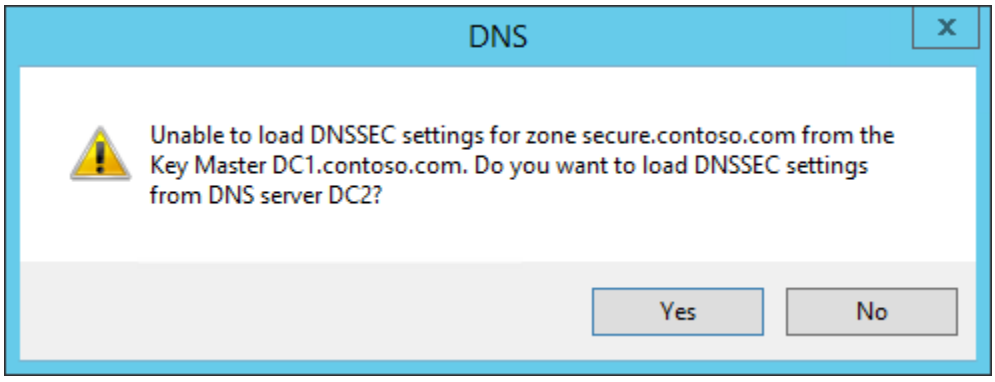

Click **OK** to load DNSSEC settings using the local server. These settings will be the settings that were configured during the last successful replication.

On the **Key Master** tab, you can choose a new Key Master from a list of available DNS servers, similar to the procedure used for a graceful transfer of the Key Master role, described earlier in this topic.

After choosing a new Key Master and clicking **OK**, a notification is displayed with information about the changes that will be made. Clicking **OK** again will proceed with the seizing operation and another notification will be displayed with the status of the role transfer. See the following examples:

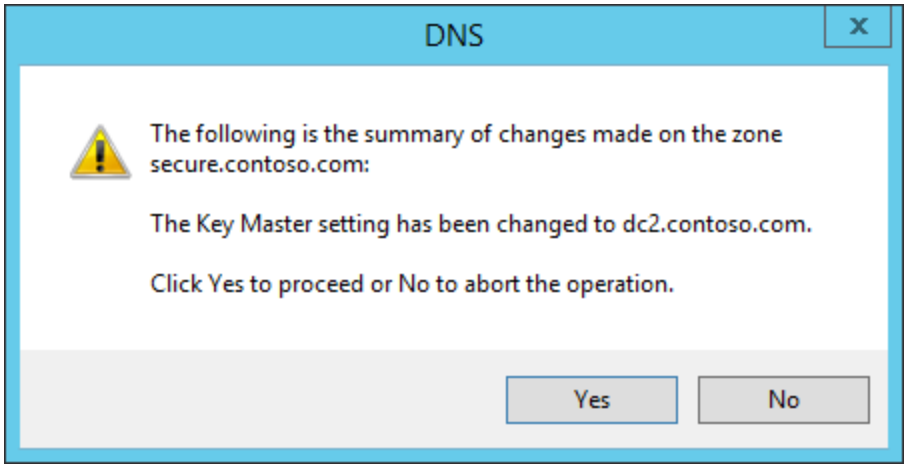

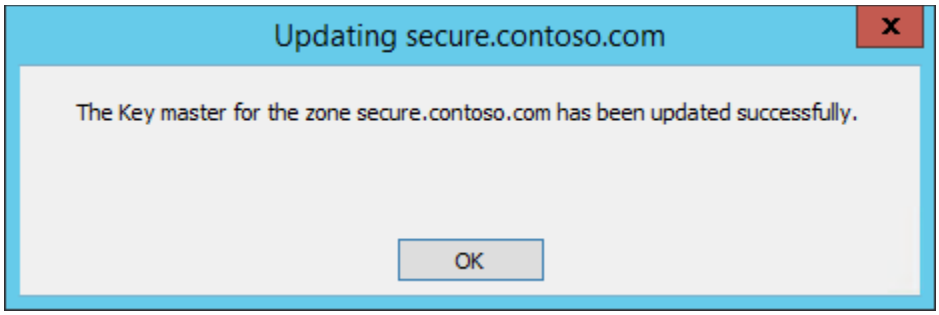

Information about the new Key Master replicates in Active Directory to all primary DNS servers. After seizing the Key Master role on another server, if the old Key Master comes online it will recognize that it is no longer the Key Master. You do not need to modify settings further.

DNSSEC standards

In addition to online signing and unsigning of dynamic Active-Directory integrated zones, DNSSEC in Windows Server 2012 and Windows Server 2012 R2 includes support for:

- Automated key rollover
- Automated trust anchor updates (RFC 5011)
- NSEC3 zone signing and validation

- NSEC3 random and user defined salt
- RSA/SHA-2 zone signing and validation

For more information about DNSSEC settings used in zone signing, see Parameter values.

Zone transfers

Zone transfers can be secured using IPsec. Windows Server 2012 and Windows Server 2012 R2 do not support Transaction Signature (TSIG) for zone transfers.

Certificate-based authentication can be used to establish an IPsec session between DNS servers. Each endpoint must present a certificate to prove its identity. This method requires that certificates are created and configured on all DNS servers that participate in zone transfers for DNSSEC-signed zones.

To configure certificate-based IPsec policies on DNS servers see the following topics:

Procedure: Deploy Certificates for DNS Server Authentication

Procedure: Deploy IPsec Policy to DNS Servers

Warning

A bug has been identified in Windows Server 2012 and Windows Server 2012 R2 whereby multiple RRSIG records in a DNSSEC-signed zone can accumulate on a secondary DNS server, particularly for resource records at the zone root, such as DNSKEY signatures. These additional resource records do not cause DNS resolution problems. These extra RRSIG records accumulate because incremental zone transfers (IXFR) from the primary DNS server will add new RRSIG records but do not automatically remove older RRSIG records. To remove these records, initiate a full zone transfer. It is also possible for RRSIG records on a secondary DNS server to expire. To prevent this, you can periodically initiate a full zone transfer from the primary to the secondary DNS server. DNS-Server-Service event ID 1525 (warning) and event ID 1524 (error) will be recorded in Event Viewer when signatures are about to expire, and have expired, respectively. This bug is pending a hotfix.

See also

Overview of DNSSEC

DNS Clients

DNS Zones

DNS Clients

In Windows® 7 and later operating systems, the Windows DNS client is security-aware, which means it can determine whether or not a DNS response that it receives was validated as genuine. The DNS client itself is non-validating, and it depends on a DNS server to provide DNSSEC

validation. The client can require that this validation be performed if rules are configured in the Name Resolution Policy Table (NRPT) and applied by Group Policy. See the following figure.

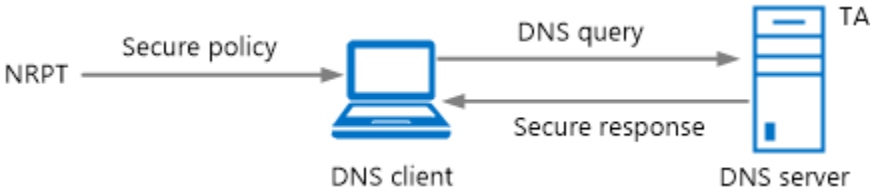

For information about configuring name resolution policy, see the The NRPT.

DNS clients are much less susceptible to DNS spoofing attacks when name resolution policies are applied and DNS responses are validated. For more information about DNS spoofing attacks, see Why DNSSEC.

In this section

Security-aware client

DNS queries and responses

Managing validation

DNS client resolver behavior

Security-aware client

With the implementation of DNSSEC in Windows Server 2008 R2 and Windows 7, the client is security-aware, which means that it understands DNSSEC and can require validation of DNS responses. Validation is not possible unless a zone is signed and the validating DNS server supports DNSSEC. For more information about how validation works, see Overview of DNSSEC.

The DNS client service in previous operating systems was not security-aware. This information is summarized in the following table:

Table 1: DNS client security-aware status

Operating system	Security aware
Windows XP, Windows Server 2003, and previous operating systems	No
Windows Vista and Windows Server 2008	No
Windows 7 and Windows Server 2008 R2	Yes
Windows 8, Windows Server 2012, and later operating systems	Yes

Note: A Windows DNS client is any computer that is running the Windows DNS Client service. The DNS Client service is included in every version of the Windows operating system, even if the computer is also a DNS server.

Security-aware DNS clients do not automatically require that DNS responses are validated. This requirement must be specifically enabled for a given namespace by name resolution policy. Policies can be configured to require validation for some DNS zones (i.e. namespaces) and not for other zones.

The following section summarizes how DNS queries and responses are affected when a zone is signed with DNSSEC.

DNS queries and responses

When a zone has been signed with DNSSEC, the success or failure of a DNS query depends on the following settings:

- **Name resolution policy**: Name resolution policy can only be applied to security-aware DNS clients. Non-security-aware clients will ignore policy settings. Policy can be configured to either require DNSSEC validation or not require DNSSEC validation for a given namespace. Rules for specific namespaces are created in Group Policy using a special policy table called the Name Resolution Policy Table, or NRPT. For more information, see The NRPT.
- **Trust anchor status**: Trust anchors are installed on DNS servers. A trust anchor is a special resource record that holds a public cryptographic key used to validate DNSSEC-signed DNS responses for a specific namespace. A trust anchor must be installed on any DNS server that will perform DNSSEC validation for DNS clients, unless the DNS server is already authoritative for the namespace (i.e. it hosts the zone). A trust anchor can be 1) available and valid, 2) available and invalid, or 3) missing. A security-aware DNS server will automatically attempt to validate a DNS response if the server has a trust anchor installed for the specified zone. You cannot install a trust anchor on a DNS server that does not support DNSSEC. For more information about trust anchors, see How DNSSEC works and Trust Anchors.
- **Security-aware status**: If a DNS client is not security-aware, it can still be protected by DNSSEC because a DNS server with an installed trust anchor for a given namespace will automatically attempt to validate all responses in that namespace, and will discard invalid (spoofed) responses. However, a non-security-aware client will ignore name resolution policy requirements for DNSSEC validation. Therefore, if a trust anchor is not installed on a DNS server used by the client, and NRPT policy requires that validation be performed, only a security-aware DNS client will fail DNS resolution. The behavior of a non-security-aware client is the same as a security-aware client that does not have any NRPT rules applied.

The following table summarizes how DNSSEC validation works for security-aware and non-security aware DNS clients under different conditions:

Table 2: Responses to DNS queries from security-aware and non-security aware clients

DNSSEC validation required	Trust anchor (TA) status	Non-security-aware client	Security-aware client
No	TA: None	Success	Success

DNSSEC validation required	Trust anchor (TA) status	Non-security-aware client	Security-aware client
No	TA: Valid	Success	Success
No	TA: Invalid	Failure	Failure
Yes	TA: None	Success	Failure
Yes	TA: Valid	Success	Success+RRSIG
Yes	TA: Invalid	Failure	Failure

Important

If the DNS server used by a DNS client has a trust anchor for the zone that is queried, DNSSEC validation will be attempted even if it is not required by NRPT policy. Consequently, if a zone is unsigned it is critical that you also remove all trust anchors that were installed on DNS servers.

The results displayed for **TA: Invalid** apply to all cases where validation fails, such as during a DNS spoofing attack. In a spoofing attack, the trust anchor itself is valid, but the resource records to be validated are invalid. The DNS server returns a failure in this case to both the security-aware and non-security-aware client.

If no TA is present, the DNS server does not attempt validation, and the DNS client will fail to resolve the name only if: 1) it is security-aware, and 2) the NRPT policy that is applied to the client requires validation for that namespace.

Managing validation

In some cases, DNSSEC validation might be broken for a specific namespace. To temporarily disable DNSSEC validation for this zone, but continue requiring DNSSEC validation for other zones, consider a forward-around design. See the following example:

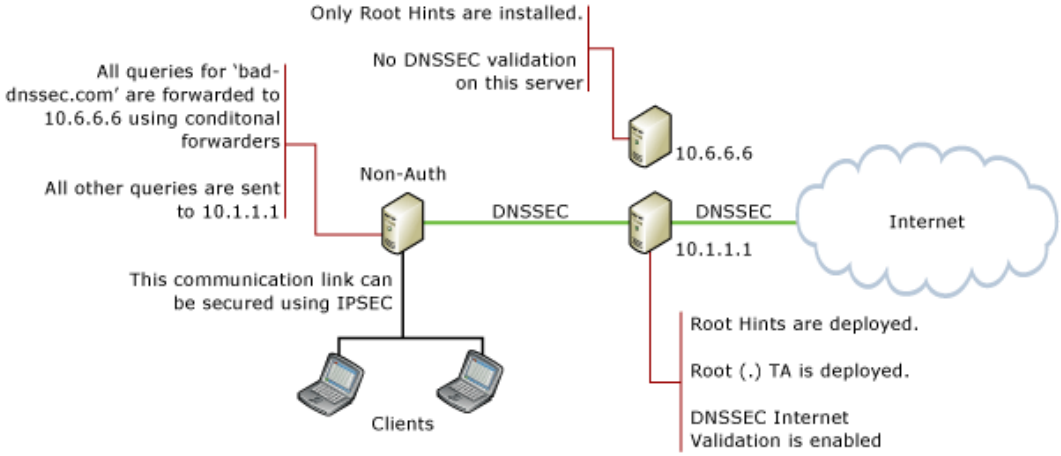

In this example, a conditional forwarder is added to non-authoritative resolving DNS servers to temporarily disable DNSSEC validation for the zone *bad-dnssec.com*. DNS queries for this zone are sent to the DNS server at *10.6.6.6* that does not have a trust anchor installed and does not perform DNSSEC validation. Other queries are sent to the DNS server at *10.1.1.1* where DNSSEC validation is performed.

DNS client resolver behavior

It is important to remember that a DNS client can be configured to query multiple DNS servers to obtain a response. Queries will continue until a response is received, therefore it is important to enable DNSSEC validation on secondary DNS servers in case the primary DNS server fails to respond. A response of SERVFAIL from a DNS server is considered a response; therefore a Windows DNS client will not attempt to obtain additional responses after DNSSEC validation has failed. The DNS query process is summarized below, including queries that can be issued from a DNS client with multiple network adapters installed.

DNS query process

The Windows DNS client is a stub resolver, which means that when it needs to resolve a DNS name, it issues a single recursive query to the primary DNS server configured on its network interface.

As long as that DNS server responds, no other DNS queries are issued. The DNS server is responsible for handling all additional tasks needed to resolve the DNS query, and the final result is returned to the DNS client.

If a DNS client does not receive a response to a DNS query and additional DNS servers are configured on one or more network interfaces, it can send additional queries to these DNS servers until it receives a response. A negative response, such as name not found is also considered a valid response.

In Windows 7, Windows Server 2008 R2 and later operating systems, the DNS client will query DNS servers that are configured on its network interfaces using the following process.

- **Query 1**: The DNS Client service sends the query to the first DNS server on the preferred network adapter's list of DNS servers and waits one (1) second for a response. The preferred network adapter on a computer with multiple network interfaces is the one that is first in the binding order.
- **Query 2**: If the DNS Client service does not receive a response from the first DNS server on the preferred network adapter within one second, it sends the query to the next DNS server on the preferred interface and first DNS servers on all other network adapters and waits one (1) second for a response.
- **Query 3**: If the DNS Client service does not receive a response from any of the DNS servers, it sends the query to the next DNS servers on each of the network adapters that are still under consideration and waits two (2) seconds for a response. A network interface is considered "under consideration" if no response has been received from any DNS server configured on that interface, and the list of DNS servers has not been exhausted.

- **Query 4**: If the DNS Client service still does not receive a response from any DNS server, it sends the name query to all DNS servers on all network adapters that are still under consideration and waits four (4) seconds for a response.
- **Query 5**: If the DNS Client service still does not receive a response from any DNS server, it repeats the previous query, sending the name query to all DNS servers on all adapters that are still under consideration and waiting four (4) seconds for a response.

The total wait time is 12 seconds (1, 1, 2, 4, 4) if there is no response. If a query is sent to multiple DNS servers and more than one DNS server responds, the first response that is received is used.

If it is necessary for a DNS client to use another DNS server to resolve a query, this server is temporarily set to be first on the priority list of DNS servers used. After 15 minutes, the DNS client will reset DNS server priority list back to default settings.

See also

Overview of DNSSEC

DNS Servers

DNS Zones

The NRPT

DNS Zones

Signing a zone with DNSSEC adds several new DNSSEC-related resource records to the zone. These new resource records can be sent in response to DNS queries for the zone along with the traditional, non-DNSSEC-related resource records. See the following figure:

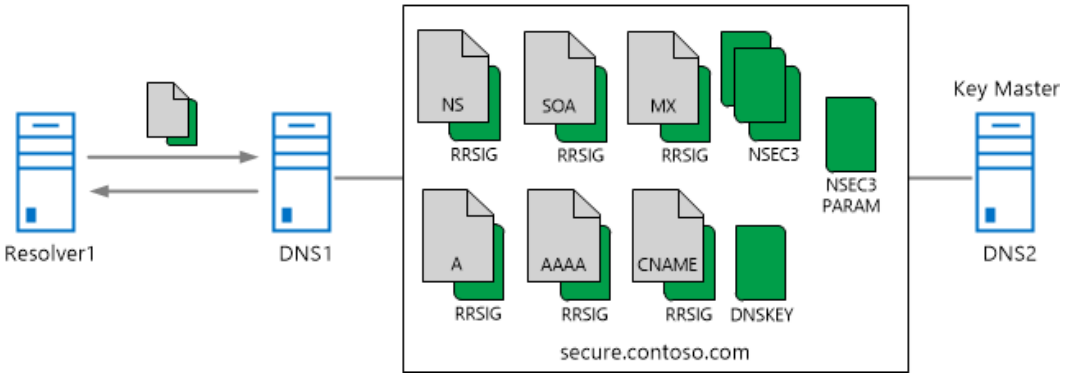

In the previous figure, both DNS1 and DNS2 are authoritative for the signed zone *secure.contoso.com*. DNS2 is the Key Master for the zone and DNS1 can be either an Active Directory-integrated primary DNS server, or a file-backed secondary authoritative DNS server. The DNS server Resolver1 will resolve DNS names through recursion or forwarding on behalf of a DNS client, and is capable of validating DNS responses for the *secure.contoso.com* zone.

DNS servers must specifically request that DNSSEC records be included in DNS responses. Therefore, servers that do not support DNSSEC will receive only traditional resource records in a DNS response.

A signed zone can be a forward or reverse lookup zone, a static or dynamic zone, and Active Directory-integrated or file-backed. If the zone is Active Directory-integrated, private zone signing keys replicate automatically to all primary, master DNS servers through Active Directory replication. Each master server signs its own copy of the zone when it receives the key. For optimal performance, and to prevent increasing the size of the Active Directory database file, the signed copy of the zone remains in memory for Active Directory-integrated zones. A DNSSEC signed zone is only committed to disk for file-backed zones.

In an Active Directory multi-master configuration, zone settings, including DNSSEC parameters, can be modified on any primary authoritative DNS server provided that the Key Master is online. You cannot perform zone signing and unsigning or modify DNSSEC parameters of a zone on a secondary DNS server.

In this section

Signing a zone

DNS Manager

- Zone Signing Wizard
  - Signing options
    - Custom parameters
    - Existing parameters
    - Default parameters
  - Parameter values
- Unsign zone wizard
- DNSSEC properties page
  - Modifying DNSSEC properties

Windows PowerShell

- Sign a zone
- Unsign a zone
- Review DNSSEC properties
- Review the status of signing keys
- Modify DNSSEC properties

Signing a zone

In Windows Server 2012 and Windows Server 2012 R2, the process of zone signing is greatly simplified over previous operating systems.

On a DNS server running Windows Server 2012 or Windows Server 2012 R2 with role administration tools including **DNS Server Tools** installed, you can use DNS Manager or Windows PowerShell to sign a zone. These methods are summarized below. The dnscmd.exe command line tool can also be used, but is deprecated and not recommended.

- DNS Manager - The zone signing wizard provides a step-by-step interface to input DNSSEC parameters and initiate zone signing. If the zone is already signed, you can also view or edit settings and re-sign a zone using the DNSSEC properties page. An unsign zone wizard is also provided to remove DNSSEC-signed records from a zone.
- Windows PowerShell - Many cmdlets are available to sign and unsign DNS zones, and to view or edit DNSSEC parameters.

Both methods for zone signing are compatible with each other. You can sign a zone with DNS Manager and view or edit settings with Windows PowerShell, and vice-versa.

After selecting a zone signing method, you must also define the DNSSEC Parameter values to use. Three options are available to specify DNSSEC parameter values:

- Custom parameters - You can provide custom values for each parameter.
- Existing parameters - You can use values that were used previously.
- Default parameters - You can use a default parameter value set.

See Signing options in this topic for more information about specifying DNSSEC parameter values using the zone signing wizard in DNS Manager. For information about specifying DNSSEC parameters in Windows PowerShell, see Checklist: Sign a Zone and Appendix B: Windows PowerShell for DNS Server.

DNS Manager

DNS Manager in Windows Server 2012 and Windows Server 2012 R2 provides a Zone Signing Wizard and Unsign zone wizard to greatly simplify the process of applying or removing DNSSEC from a zone. After a zone has been signed, you can also view and edit several properties and values using the DNSSEC properties page.

Zone Signing Wizard

To start the **Zone Signing Wizard** in DNS Manager, right-click a forward or reverse DNS zone on a primary, authoritative DNS server, point to **DNSSEC**, and click **Sign the Zone**. See the following example:

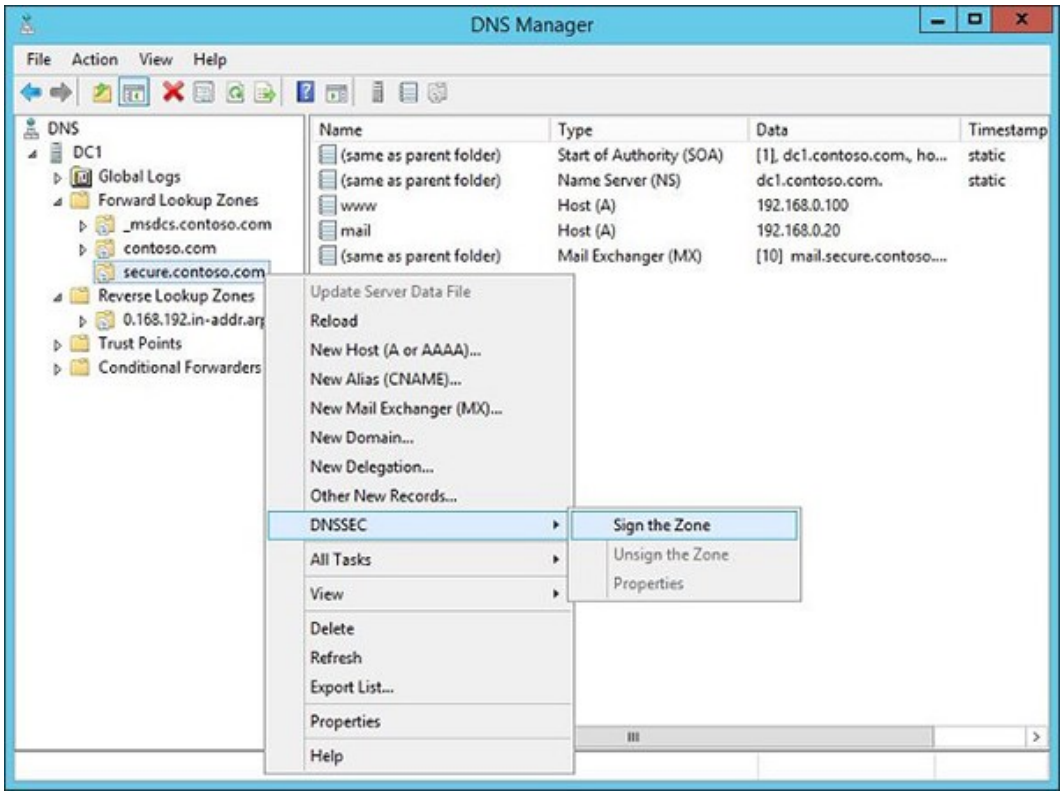

Important

You cannot sign or unsign a zone, or view DNSSEC properties of a zone on a secondary DNS server. The server must be a primary authoritative DNS server. DNSSEC related resource records are transferred to secondary servers, but you cannot view or edit the DNSSEC properties of a secondary zone.

After you click **Sign the Zone**, the wizard will display an introductory page with a summary of DNSSEC (not shown). Click **Next** to view the signing options page.

Signing options

The signing options page requires you to choose a method for specifying DNSSEC parameter values. See the following example:

The method that you choose on this page determines which pages the zone signing wizard will display. Choosing to use default settings or parameters of an existing zone will skip custom parameter pages. Some parameter values can also be modified after the zone is signed if desired.

The following options are available to specify the DNSSEC parameter values that will be used for zone signing:

- **Customize zone signing parameters**. This is the default option. If you choose this option, the wizard will display several pages that you can use to specific different parameter values, and provide you with default values for each parameter. These values can also be modified to any value within an acceptable range. If the zone was previously signed and then un-signed, the parameters that were most recently used are provided as the default values. For more information about this option, see Custom parameters.
- **Sign the zone with parameters of an existing zone**. To use this option, the DNS server must already be authoritative for one or more DNSSEC-signed zones. If you choose this option, you must provide the name of a signed zone to use as a template for signing the current zone. For more information about this option, see Existing parameters.
- **Use default settings to sign the zone**. If you choose this option, the zone will be signed with default parameters. These parameters are the same as those that are displayed by default if you choose to customize zone signing parameters, provided the zone has not been previously signed with a different set of parameters. If you choose to sign the zone with default settings, the wizard will skip to the summary page. For more information about this option, see Default parameters.

For a list of the default parameter values used in the wizard, and the allowed range of custom parameter values, see Parameter values.

The length of time required to sign a zone will depend on the parameters that you specify and the size of the zone.

Custom parameters

The following table lists pages that are displayed in the zone signing wizard when you choose to customize zone signing parameters. Click the link in the “More information” column for an explanation and example of each page.

Table 1: Custom parameter pages

Action	Description	More information
Select the Key Master	Only one server can be the Key Master at a given time. By default, the Key Master is the local server, but a different DNS server can be selected if it meets requirements.	Choose the Key Master
Add signing keys	You must choose at least one zone signing key (ZSK) and one key signing key (KSK).	Signing keys
Specify KSK parameter values	For each KSK that is added to a zone, you must specify individual parameter values.	KSK configuration
Specify ZSK parameter values	For each ZSK that is added to a zone, you must specify individual parameter values.	ZSK configuration
Specify Next secure (NSEC) values	You can choose NSEC or NSEC3 to provide authenticated denial of existence.	NSEC
Choose trust anchor distribution options	Select how trust anchors will be distributed and whether or not automatic updates are enabled.	Trust anchors
Specify signing and polling parameters	Specify DS and DNSKEY parameters.	Signing and polling
Summary	Before zone signing starts, you are provided with a summary of all DNSSEC parameter values	Summary

Action	Description	More information
	that will be used to sign the zone. Clicking **Next** on this page will initiate zone signing.	

Choose the Key Master

In Windows DNS, signing a zone requires that you designate one DNS server as the Key Master. You can change the Key Master, but only one DNS server can be the Key Master at a given time. The Key Master is the DNS server where zone signing occurs, and it is responsible for DNSSEC-related cryptographic processes, such as signing key creation and management. The Key Master must be a primary authoritative DNS server for the zone.

The following options are available when selecting the Key Master:

- The DNS server <local computer> is the Key Master.
- Select another primary server as the Key Master: <allowed list>

See the following example:

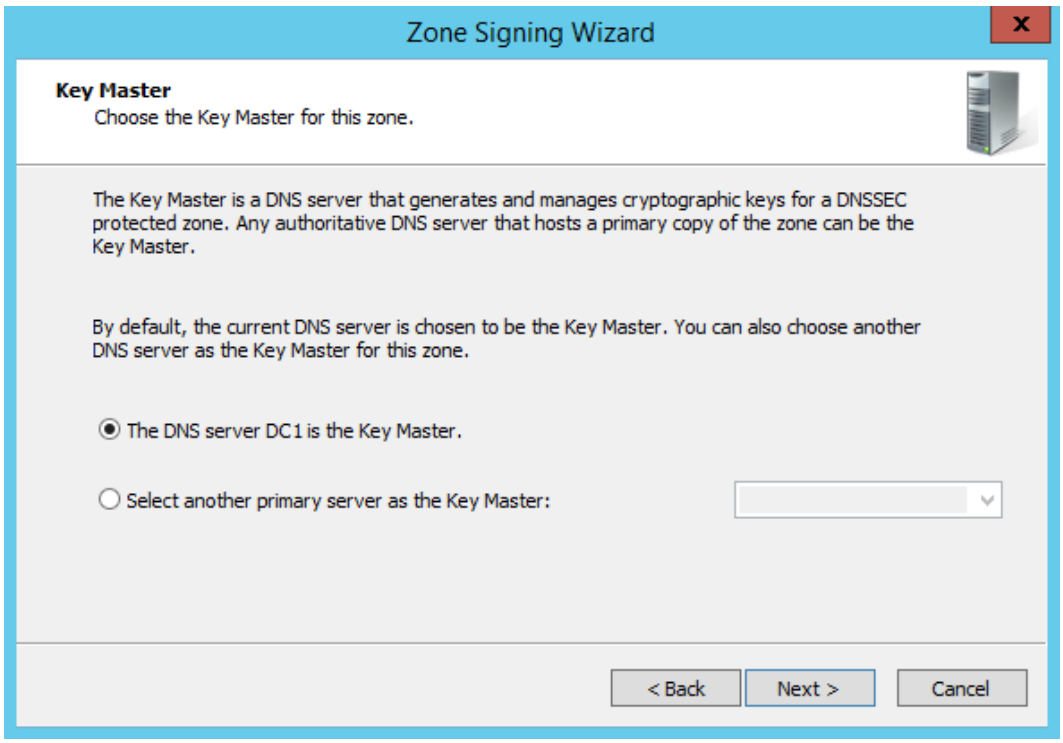

By default, the local server is selected as the Key Master. You also have the option of choosing a different DNS server from a list of online, primary, authoritative DNS servers that support online DNSSEC signing.

Tip

The signed-in user must have administrative privileges on a DNS server in order for it to be displayed in the allowed list of Key Master servers.

Only one DNS server can be the Key Master for a given zone at a given time. For more information about the Key Master, including options to transfer the Key Master role to a different DNS server after a zone is signed, see The Key Master.

Signing keys

The zone signing wizard displays six pages that provide information and enable you to add signing keys to the zone:

1. **Key signing key (KSK)**: *A KSK is an authentication key used to sign other keys*.

   This introductory page provides a short description of a KSK and a link to more information. Click **Next** on this page to add or edit KSKs.
2. **Key Signing Key (KSK)**: *Configure one or more KSKs*.

   This page lists the current KSKs. Click **Add**, **Edit**, or **Remove** to modify KSKs and then click **Next** when finished.
3. **New/Edit Key Signing Key (KSK)**

   Use this page to edit individual KSK parameter values.
4. **Zone signing key (ZSK)**: *A ZSK is an authentication key used to sign the zone data*.

   This page provides a short description of a ZSK and a link to more information. Click **Next** on this page to add or edit ZSKs.
5. **Zone Signing Key (ZSK)**: *Configure one or more ZSKs*.

   This page lists the current ZSKs. Click **Add**, **Edit**, or **Remove** to modify ZSKs and then click **Next** when finished.
6. **New/Edit Zone Signing Key (ZSK)**

   Use this page to edit individual ZSK parameter values.

For more information, including examples of these pages see the KSK configuration and ZSK configuration topics in this section.

To sign a zone with DNSSEC, you must configure at least one KSK and one ZSK. Multiple keys are also supported: up to three (3) for each cryptographic algorithm. A list of available Cryptographic algorithms is provided below.

Typically, a zone is signed with a single ZSK and a single KSK. You might use more than one key to sign a zone under the following circumstances:

1. All DNS validators do not support a single cryptographic algorithm
2. You are transitioning the DNS infrastructure to a use new signing key properties
3. You are merging DNS zones or network elements that use different signing key properties

Each additional signing key and cryptographic algorithm that you add will increase the time and computational cost of zone signing. For more information about signing keys and DNS server performance, see DNSSEC Performance Considerations.

Important

To support automatic key rollover, an additional signing key is generated for each ZSK and KSK that is added to the zone. For automatic KSK rollover, only the Double Signature rollover method is supported. For automatic ZSK rollover, only the Pre-publish rollover method is supported. For more information, see **Manage Signing Keys**.

Cryptographic algorithms

The following cryptographic algorithms are available for use in either a KSK or ZSK. You do not need to use the same algorithm for both types of signing keys, but you cannot use RSA/SHA-1 and RSA-SHA-1 (NSEC3) together.

Table 2: Cryptographic algorithms and compatible NSEC methods

Algorithm	Compatible NSEC methods
ECDSAP256/SHA-256	NSEC, NSEC3
ECDSAP384/SHA-384	NSEC, NSEC3
RSA/SHA-1	NSEC
RSA/SHA-1 (NSEC3)	NSEC3
RSA/SHA-256	NSEC, NSEC3
RSA/SHA-512	NSEC, NSEC3

You must have at least one (1) KSK and one (1) ZSK. You can have up to three (3) KSKs and three (3) ZSKs that use the same cryptographic algorithm, and you have multiple cryptographic algorithms.

KSK configuration

If the zone has not been previously signed, the KSK configuration page requires that you add at least one KSK. See the following example:

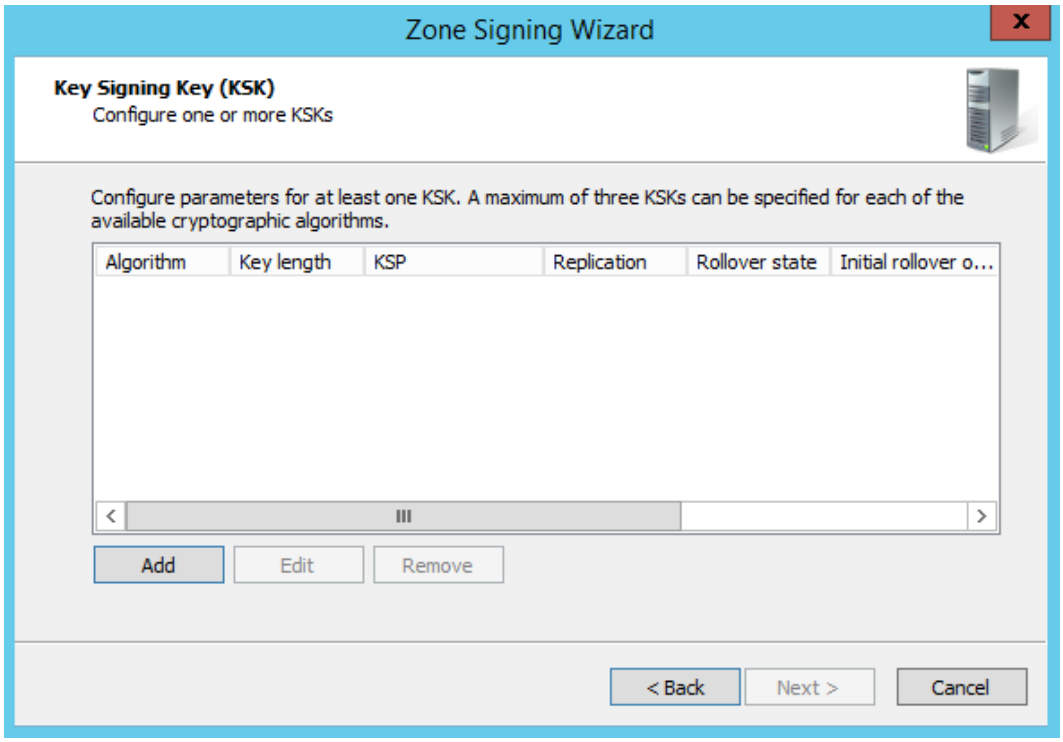

Click **Add** to display the New Key Signing Key (KSK) page:

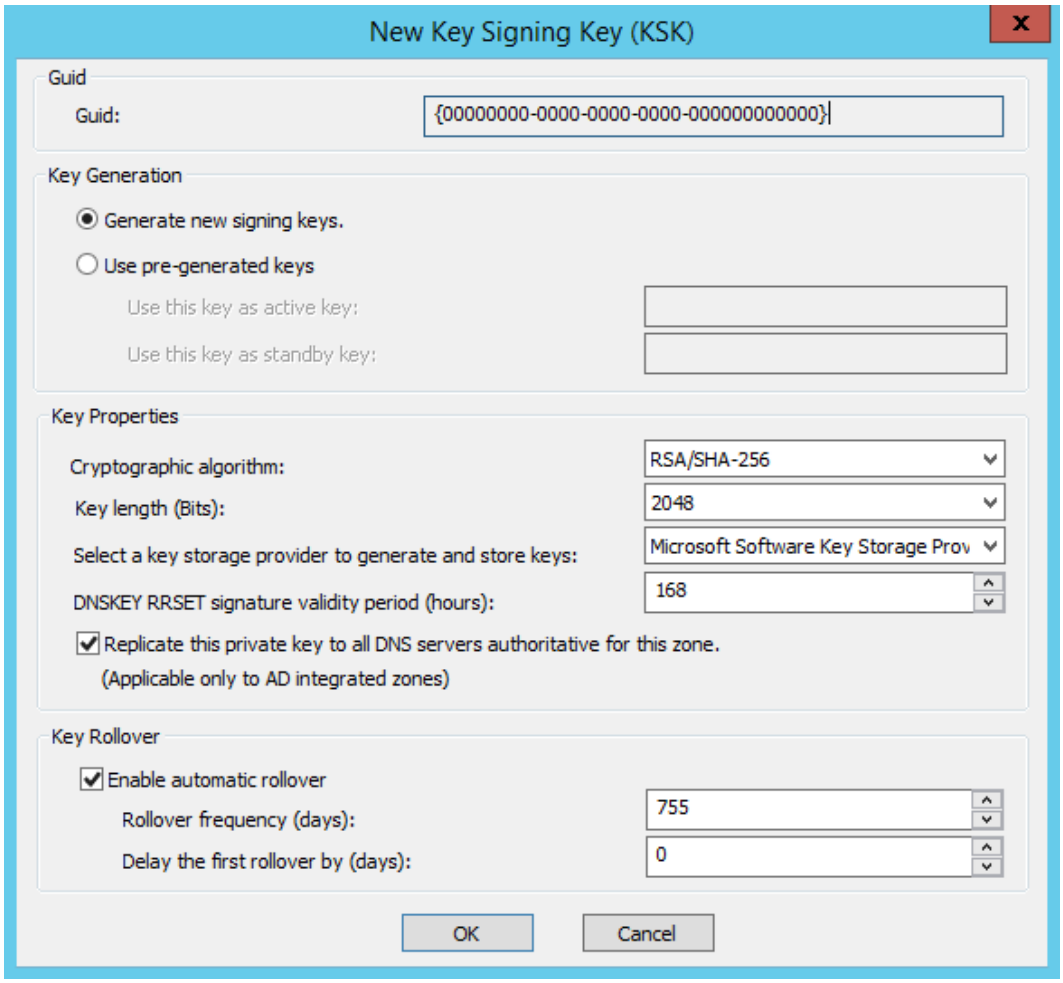

Note

If the zone has been signed and then unsigned, the KSKs and settings that were used to previously sign the zone are automatically provided again. You have the option to use the same keys and settings as before, to edit parameter values for these keys, or to remove these keys and add new ones.

Click **OK** to accept default parameter values for the KSK, or enter new parameter values and then click **OK**. Information about KSK parameters is provided in the following section.

Add additional KSKs if desired, and then click **Next** to continue.

KSK parameters

The key signing key, or KSK, is an authentication key that signs all of the DNSKEY records at the root of the zone, and it is part of the chain of trust.

When you generate a new KSK, the Key Master can create an active and a standby key. Properties for the KSK include the following:

- **Cryptographic algorithm**: By default, the RSA/SHA-1 (NSEC3) algorithm is chosen. To sign a zone with NSEC3, you cannot choose the RSA/SHA-1 algorithm (RSA/SHA-1 (NSEC3) and RSA/SHA-1 are two different available algorithms).
- **Key length**: By default, this is 2048 bits for keys that use the RSA/SHA-1 (NSEC3) algorithm. Typically, the KSK will have a longer key length than a ZSK for greater security. The KSK key length affects the time that is required for zone signing less than the ZSK key length.
- **Key storage provider**: If keys will be distributed by using Active Directory Domain Services, you must choose Microsoft Software Key Storage Provider.
- **DNSKEY RRSET signature validity period**: By default, this is set to 72 hours. Signatures that are generated using the KSK might have a longer validity period than signatures that are generated by the ZSK to provide a more stable secure entry point into the zone.
- **Key rollover**: You can choose whether to use automatic rollover and supply a rollover frequency. Automatic key rollover is highly recommended. You can also delay the first rollover by a specified number of days.

Also see: Parameter values

ZSK configuration

If the zone has not been previously signed, the ZSK configuration page requires that you add at least one ZSK. See the following example:

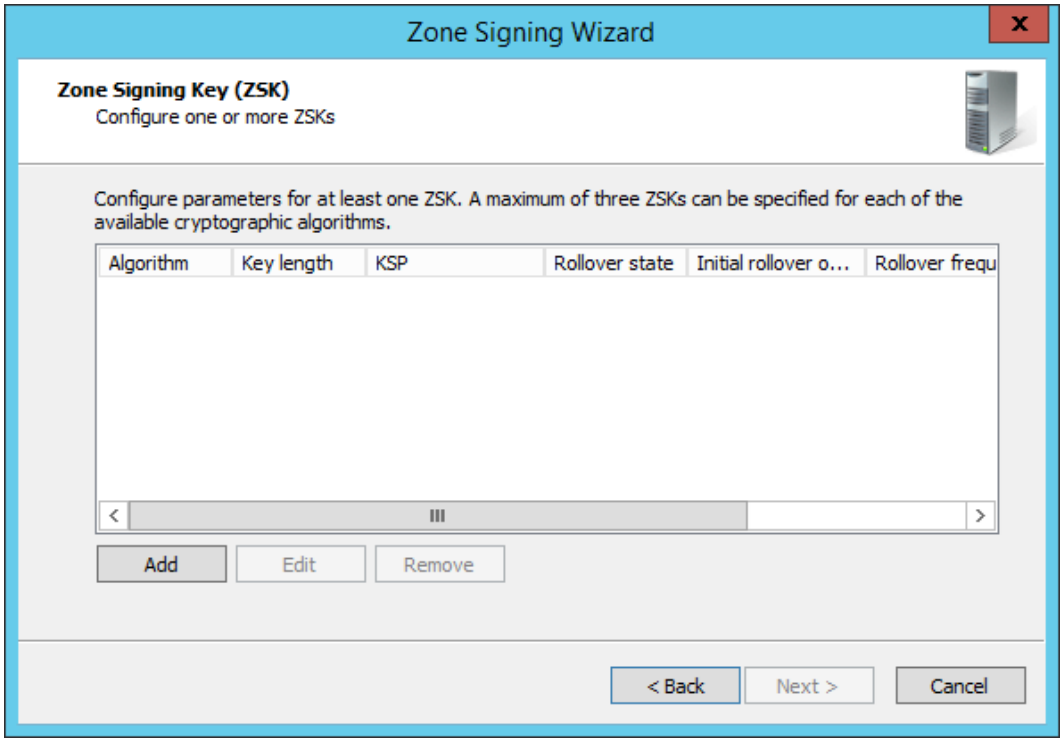

Click **Add** to display the New Zone Signing Key (ZSK) page:

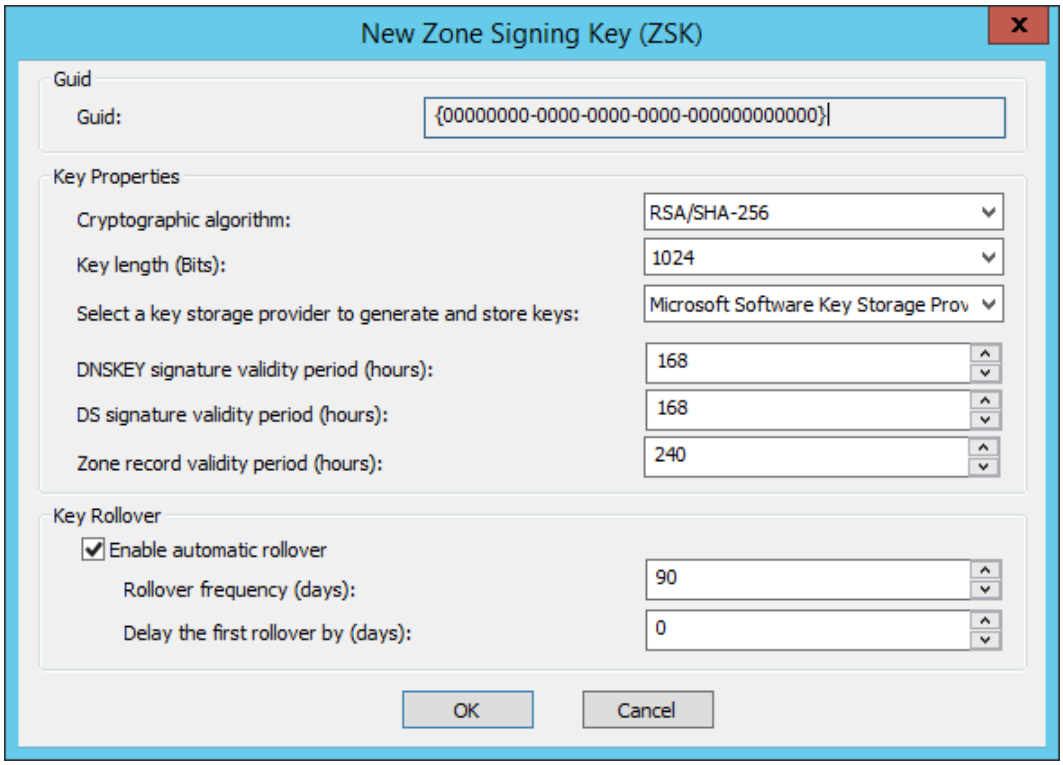

Note

If the zone has been signed and then unsigned, the ZSKs and settings that were used to previously sign the zone are automatically provided again. You have the option to use the same keys and settings as before, to edit parameter values for these keys, or to remove these keys and add new ones.

Click **OK** to accept default parameter values for the ZSK, or enter new parameter values and then click **OK**.

Add additional ZSKs if desired, and then click **Next** to continue.

ZSK parameters

The zone signing key, or ZSK, is used to sign zone data. A ZSK is typically rolled over more frequently than a KSK. Properties for the ZSK include the following:

- **Cryptographic algorithm**: By default, the RSA/SHA-1 (NSEC3) algorithm is chosen. To sign a zone with NSEC3, you cannot choose the RSA/SHA-1 algorithm (RSA/SHA-1 (NSEC3) and RSA/SHA-1 are two different available algorithms). The algorithm you choose can affect the time required for zone signing.
- **Key length**: By default, this is 1024 bits for keys that use the RSA/SHA-1 (NSEC3) algorithm. Typically, the ZSK will have a shorter key length than a KSK. The ZSK key length can affect the time required for zone signing.

- **Key storage provider**: If keys will be distributed by using Active Directory Domain Services, you must choose Microsoft Software Key Storage Provider.
- **DNSKEY signature validity period**: By default, this is set to 72 hours.
- **DS signature validity period**: By default this is set to 72 hours.
- **Zone record validity period**: By default this is set to 240 hours.

Also see: Parameter values

NSEC

The next page in the zone signing wizard is the Next Secure (NSEC) page. On this page, you are asked to choose whether to use NSEC3 or NSEC for authenticated denial of existence. See the following example:

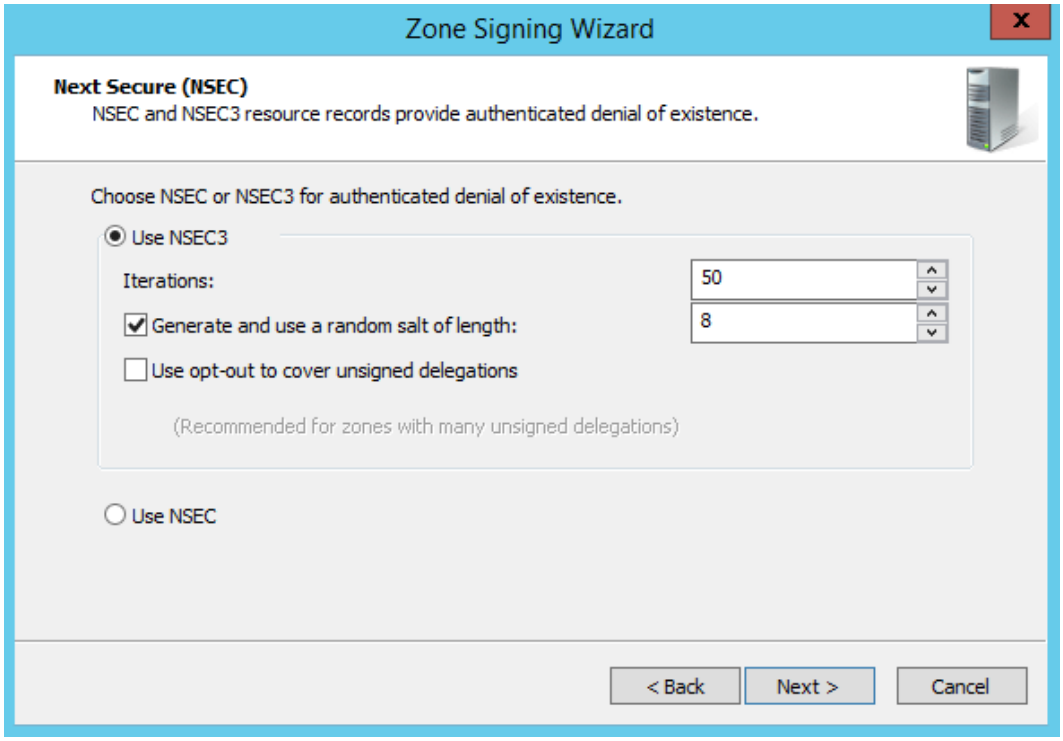

You should choose NSEC3 on this page unless you have a good reason to use NSEC, such as a requirement to use the RSA/SHA-1 algorithm, or some other incompatibility with NSEC3. For more information about NSEC and NSEC3, see DNSSEC-related resource records.

Customize settings on this page and then click **Next** to continue.

Trust anchors

This page enables you to configure distribution and update settings for trust anchors (TAs). See the following example:

Customize settings on this page and then click **Next** to continue.

Signing and polling

The signing and polling page enables you to configure values for DS and DNSKEY records in the zone. See the following example:

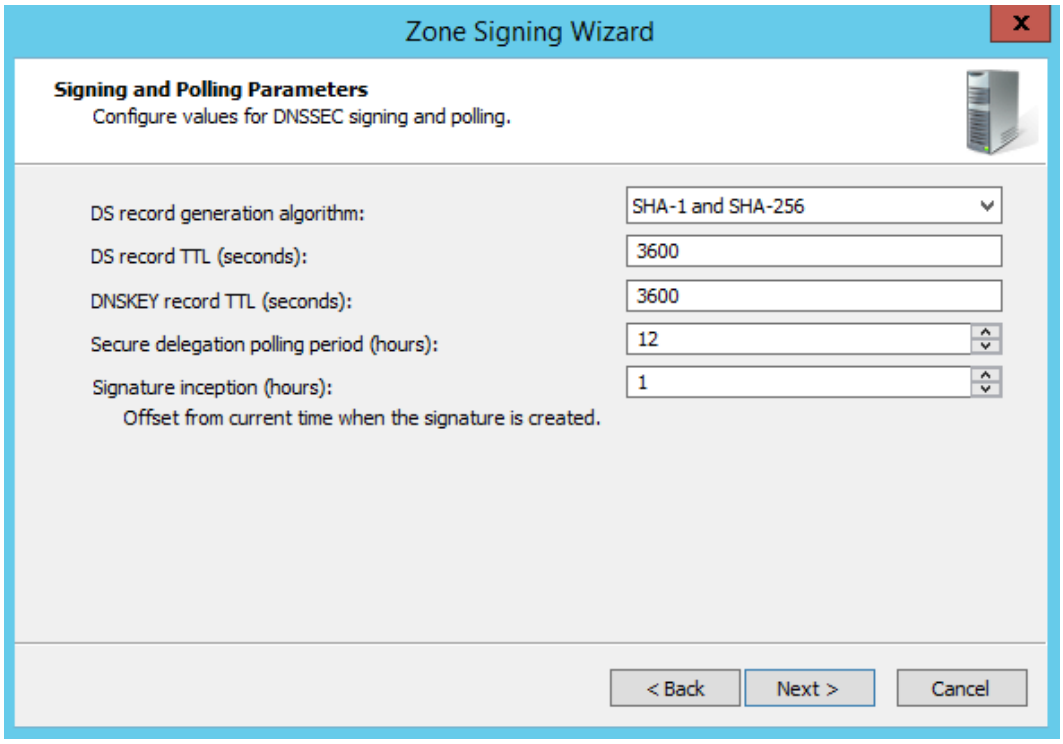

Customize settings on this page and then click **Next** to continue.

Summary

The summary page is entitled **DNS Security Extensions (DNSSEC)**.

Caution

Clicking **Next** on the summary page will initiate zone signing.

Do not click **Next** on this page unless you are ready to begin signing the zone.

You can also click **Back** to change settings, or click **Cancel** to skip signing the zone. See the following example:

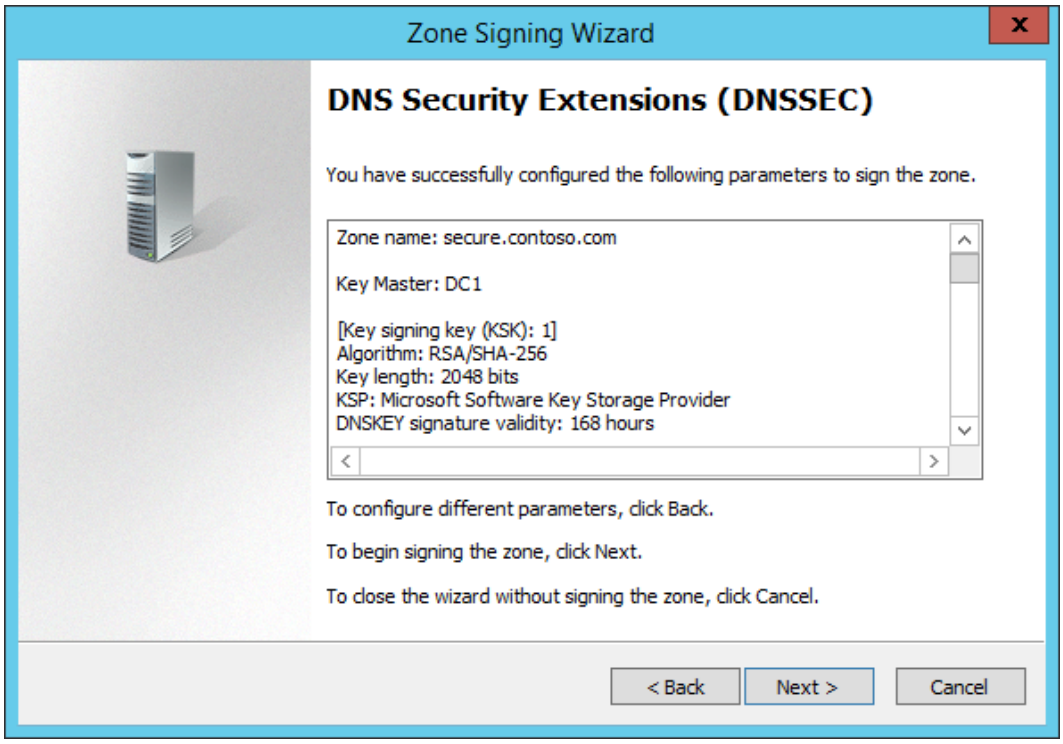

If you click **Next**, and zone signing is successful, a final page will be displayed indicating that the zone has been successfully signed. Error details are displayed if zone signing was not successful. Click **Finish** to close the wizard.

Existing parameters

If you choose **Sign the zone with parameters of an existing zone**, you must provide the name of another zone hosted by the local DNS server that is signed with DNSSEC. See the following example:

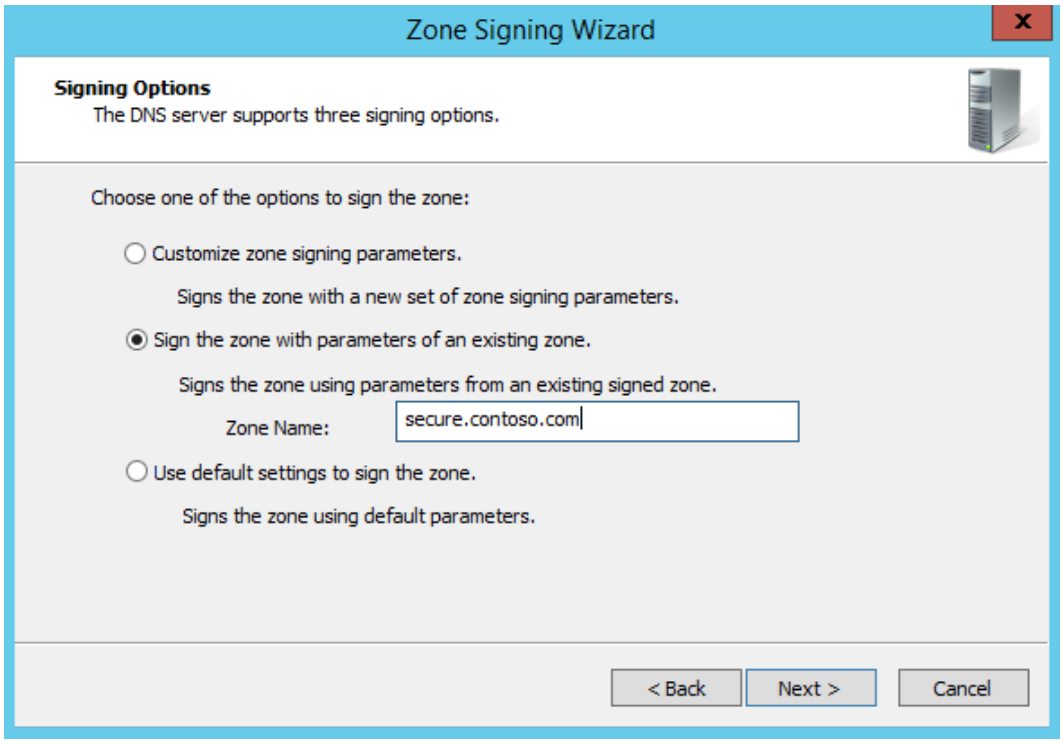

Note

A zone that was signed offline on a DNS server running Windows Server 2008 R2 is not considered to be signed by DNS servers running Windows Server 2012 or Windows Server 2012 R2. If the local DNS server is a secondary DNS server for the specified zone, it will attempt to retrieve DNSSEC settings from the primary DNS server.

If DNSSEC settings for the specified zone are successfully retrieved, the zone signing wizard will skip to the Choose the Key Master page. Clicking **Next** on this page will display the Summary page, and clicking **Next** on this page will initiate zone signing.

Default parameters

Choosing **Use default settings to sign the zone** and clicking **Next** will select the local DNS server as the Key Master and skip to the Summary page. Clicking **Next** on the summary page will initiate zone signing. Default settings used to sign a zone are provided in the table below.

Parameter values

The following table lists default and allowed DNSSEC parameter values used in the zone signing wizard.

Table 3: DNSSEC parameters

Action	Default	Allowed
Key Master	The local DNS server is the Key Master	Any DNS server found in Active Directory that is primary and authoritative for the zone can be the Key Master. The signed in user must have Domain Admin privileges, or equivalent, to choose the Key Master.
KSK	One KSK	A minimum of one KSK is required. Up to three KSKs can be added for each of the available cryptographic algorithms.
KSK: GUID	By default, all zeros are displayed. These values are a placeholder. The key's GUID is generated automatically during zone signing.	No values are allowed in this field. You cannot manually input a GUID. The key GUID is also called the KeyId.
KSK: Key Generation	Generate new signing keys	Use pre-generated keys: If this option is chosen, you must specify an available active and standby key. Use Get-DnsServerSigningKey to view available keys.
KSK: Key Properties*	Cryptographic algorithm: RSA/SHA-256 Key length (bits): 2048 Select a key storage provider to generate and store keys: Microsoft Software Key Storage Provider DNSKEY RRSET signature validity period (hours): 168 Replicate this private key to all DNS servers authoritative for this zone: Enabled	Cryptographic algorithm: ECDSAP256/SHA-256, ECDSAP256/SHA-384, RSA/SHA-1, RSA/SHA-1 (NSEC3), RSA/SHA-256, RSA/SHA-512. Key length (bits): 1024 – 4096 in increments of 64 Select a key storage provider to generate and store keys: Microsoft Software Key Storage Provider, Microsoft Platform Crypto Provider, Microsoft Smart Card Key Storage Provider DNSKEY RRSET signature

Action	Default	Allowed
		validity period (hours): 6 - 720 Replicate this private key to all DNS servers authoritative for this zone: Enabled or disabled
KSK: Key Rollover	Enable automatic rollover: Enabled Rollover frequency (days): 755 Delay the first rollover by (days): 0	Enable automatic rollover: Enabled or disabled Rollover frequency (days): 7 - 7300 Delay the first rollover by (days): 0 - 30
ZSK	One ZSK	A minimum of one ZSK is required. Up to three ZSKs can be added for each of the available cryptographic algorithms.
ZSK: GUID	By default, all zeros are displayed. These values are a placeholder. The key's GUID is generated automatically during zone signing.	No values are allowed in this field. You cannot manually input a GUID. The key GUID is also called the KeyId.
ZSK: Key Properties*	Cryptographic algorithm: RSA/SHA-256 Key length (bits): 1024 Select a key storage provider to generate and store keys: Microsoft Software Key Storage Provider DNSKEY signature validity period (hours): 168 DS signature validity period (hours): 168 Zone record validity period (hours): 240	Cryptographic algorithm: ECDSAP256/SHA-256, ECDSAP256/SHA-384, RSA/SHA-1, RSA/SHA-1 (NSEC3), RSA/SHA-256, RSA/SHA-512. Key length (bits): 1024 – 4096 in increments of 64 Select a key storage provider to generate and store keys: Microsoft Software Key Storage Provider, Microsoft Platform Crypto Provider, Microsoft Smart Card Key Storage Provider DNSKEY signature validity period (hours): 6 - 720 DS signature validity period (hours): 6 - 720 Zone record validity period

Action	Default	Allowed
		(hours): 6 - 720
ZSK: Key Rollover	Enable automatic rollover: Enabled Rollover frequency (days): 90 Delay the first rollover by (days): 0	Enable automatic rollover: Enabled or disabled Rollover frequency (days): 7 - 1825 Delay the first rollover by (days): 0 - 30
Next Secure (NSEC) **	Choose NSEC or NSEC3 for authenticated denial of existence: Use NSEC3	Choose NSEC or NSEC3 for authenticated denial of existence: Use NSEC3, Use NSEC
Next Secure (NSEC): NSEC3 options	Iterations: 50 Generate and use a random salt of length: 8 Use opt-out to cover unsigned delegations: Disabled	Iterations: For key length values of 256-1024, 1088-2048, and 2112-4096, iterations must be in the range of 0-150, 0-500, and 0-2500, respectively. Generate and use a random salt of length: 1 - 255 Use opt-out to cover unsigned delegations: Enabled or disabled
Trust Anchors	Enable the distribution of trust anchors for this zone: Disabled Enable automatic update of trust anchors on key rollover (RFC 5011): Enabled	Enable the distribution of trust anchors for this zone: Enabled or disabled Enable automatic update of trust anchors on key rollover (RFC 5011): Enabled or disabled.
Signing and polling parameters	DS record generation algorithm: SHA-1 and SHA-256 DS record TTL (seconds): 3600 DNSKEY record TTL (seconds): 3600 Secure delegation polling period (hours): 12 Signature inception (hours): 1	DS record generation algorithm: All, None, SHA-1, SHA-1 and SHA-256, SHA-1 and SHA-384, SHA-256, SHA-256 and SHA-384, SHA-384 DS record TTL (seconds): 0 - 604800 DNSKEY record TTL (seconds): 0 - 604800 Secure delegation polling period (hours): 1 - 168 Signature inception (hours): 0 -

Action	Default	Allowed
		168

**You cannot sign a zone with a key that uses the RSA/SHA-1 (NSEC3) algorithm if another signing key uses the RSA/SHA-1 algorithm. These two cryptographic algorithms are incompatible.

**You cannot use NSEC for authenticated denial of existence if you have chosen RSA/SHA-1 (NSEC3) as a cryptographic algorithm in any KSK or ZSK. You also cannot use NSEC3 if you are using RSA/SHA-1 with any signing key.

Unsign zone wizard

DNS Manager also provides a wizard to unsign a DNS zone. To use DNS Manager to unsign a zone that was previously signed on a server running Windows Server 2012 or Windows Server 2012 R2, right-click the zone in the console tree, point to **DNSSEC**, and then click **Unsign the Zone** to display the **Unsign zone** wizard. See the following example:

To unsign the zone, click **Next**. The zone will immediately be unsigned. You can unsign a zone on any primary, authoritative DNS server.

After unsigning a zone using the wizard, you cannot click **Back** to re-sign the zone. Similarly, you cannot unsign a zone by clicking **Back** after you have completed the zone signing process.

However, when you unsign a zone, the parameters that were used for zone signing are saved so that you can use them again. To re-sign a zone using the same parameters that were used previously, choose **Customize zone signing parameters** and accept the default values that are provided for each parameter. The KSK and ZSK's that were used to sign the zone previously, along with values for all custom DNSSEC parameters will be automatically populated in the **Zone Signing Wizard**. If you are using Windows PowerShell, you can re-use the same parameters by specifying the **DoResign** parameter with the **Invoke-DnsServerZoneSign** cmdlet.

If do not wish to use the same parameters, you also have the option of changing them. Re-signing the zone again replaces this default parameter set with the current one.

DNSSEC properties page

To view the DNSSEC properties of a signed zone in DNS Manager, right-click the signed zone in the DNS Manager console tree, point to **DNSSEC**, and then click **Properties**.

Tip

To view DNSSEC properties, you must view the zone on a primary authoritative DNS server. You cannot view the DNSSEC properties page on a secondary authoritative DNS server. The DNS server does not need to be the Key Master.

Click the **Key Master**, **KSK**, **ZSK**, **Next Secure (NSEC)**, **Trust Anchor**, or **Advanced** tab to see information about the zone. See the following example:

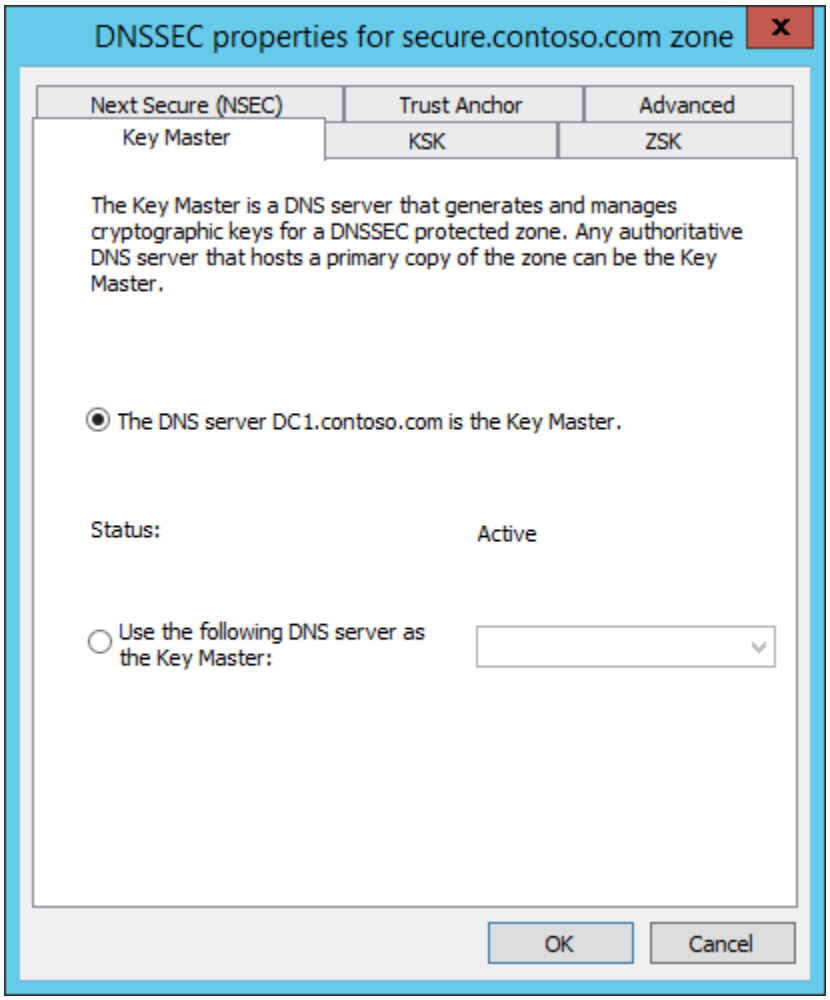

Tip

The DNSSEC node in DNS Manager introduces a second properties page for each DNS zone. You can right-click a zone and then click **Properties**, or right-click a zone, point to **DNSSEC**, and click **Properties**. If you do not see the DNSSEC properties dialog box, verify that you pointed to **DNSSEC** first.

You can modify DNSSEC parameters on the properties page. For example, you can click the **ZSK** tab, click a ZSK in the list, click **Edit**, and modify one or more of the parameters.

If you edit DNSSEC parameters, the zone is re-signed using these new parameters when you click **OK** in the DNSSEC properties dialog, and you are provided with zone signing status. See the following example.

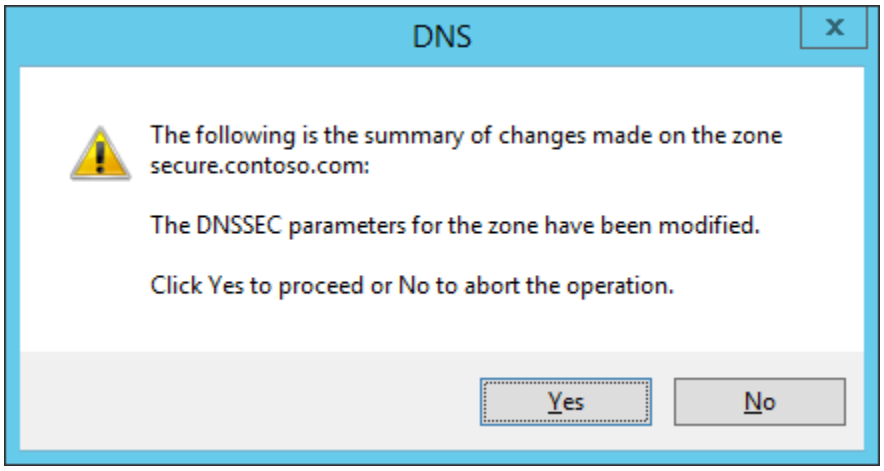

If you do not wish to commit the changes that you made to DNSSEC parameters, click **No** and parameters will not be changed. Clicking **Yes** will re-sign the zone with new parameters.

Modifying DNSSEC properties

After a zone is signed, certain DNSSEC properties can be modified. For example, keys may be added after the zone is signed. If you add a new signing key, the zone is immediately signed again with this key and any other keys that are still present in the zone. Active signing keys cannot be deleted, but you can retire a signing key provided that it is not the last KSK or ZSK in the zone. For more information, see Procedure: Retire a Signing Key.

Some properties can be changed without un-signing and resigning the zone. However, settings will not be used immediately. For example, changes to KSK settings will not take effect until the next KSK rollover. This includes signature and polling settings on the **Advanced** tab of DNSSEC properties. Similarly, changes to ZSK settings do not take effect until the next ZSK rollover.

Certain DNSSEC settings cannot be modified when a zone is signed. Changing these settings requires that the zone first be unsigned, and then signed again. In this scenario, any trust anchors that have been distributed to resolving DNS servers must be manually updated.

See Procedure: Modify Zone Signing Parameters for more information about how to modify DNSSEC properties of a zone that is currently signed.

Windows PowerShell

See Appendix B: Windows PowerShell for DNS Server for a summary of Windows PowerShell cmdlets you can use to manage a DNSSEC deployment. Some basic examples are also provided below.

Sign a zone

You can use the Invoke-DnsServerZoneSign cmdlet to sign a DNS zone. See the following example:

```
PS C:\> Invoke-DnsServerZoneSign -ZoneName secure.contoso.com -ComputerName DC1 -
SignWithDefault -PassThru -Verbose -Force
VERBOSE: This will reset the DNSSEC settings for the zone to defaults and sign it.
The following parameters will be configured:
One KeySigningKey with:
 CryptoAlgorithm: RsaSha256
 Key Length: 2048
 Rollover Period: 755.00:00:00
One ZoneSigningKey with:
 CryptoAlgorithm: RsaSha256
 Key Length: 1024
 Rollover Period: 90.00:00:00
Authenticated Denial of Existence: NSec3
Random Salt Length: 8 octets
NSec3 Hash Algorithm: RsaSha1
NSec3 Iterations: 50
NSec3 OptOut: No.

ZoneName                            ZoneType        IsAutoCreated   IsDsIntegrated
IsReverseLookupZone  IsSigned
--------                            --------        -------------   --------------  -----
--------------  --------
secure.contoso.com                  Primary         False           True            False
True
```

In this example, the secure.contoso.com zone is signed using default parameters. The computer DC1 is specified as the Key Master because it is where zone signing is done. If the ComputerName parameter is not specified, signing will be attempted using the local computer as the Key Master.

Unsign a zone

You can use the Invoke-DnsServerZoneUnSign cmdlet to sign a DNS zone. See the following example:

```
PS C:\> Invoke-DnsServerZoneUnSign -ZoneName secure.contoso.com -ComputerName DC2 -
PassThru -Verbose -Force

VERBOSE: This will initiate unsigning of the zone secure.contoso.com on server DC2.


ZoneName                            ZoneType        IsAutoCreated   IsDsIntegrated
IsReverseLookupZone  IsSigned
--------                            --------        -------------   --------------  -----
--------------  --------
secure.contoso.com                  Primary         False           True            False
False
```

In this example, the secure.contoso.com zone is unsigned from the computer DC2. You can unsign a zone from any primary, authoritative DNS server that has hosts a signed version of the zone; it does not need to be The Key Master. If the **ComputerName** parameter is not specified, unsigning will be attempted using the local computer.

Review DNSSEC properties

You can use the Get-DnsServerDnsSecZoneSetting cmdlet to view settings for a signed zone. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com -ComputerName DC1



ZoneName                      : secure.contoso.com
IsKeyMasterServer             : True
KeyMasterServer               : DC1.contoso.com
KeyMasterStatus               : Online
DenialOfExistence             : NSec3
NSec3HashAlgorithm            : RsaSha1
NSec3Iterations               : 50
NSec3OptOut                   : False
IsNSec3SaltConfigured         : True
NSec3RandomSaltLength         : 8
NSec3UserSalt                 : -
DnsKeyRecordSetTTL            : 01:00:00
```

```
DSRecordSetTTL                : 01:00:00
DSRecordGenerationAlgorithm   : {Sha1, Sha256}
DistributeTrustAnchor         : {None}
EnableRfc5011KeyRollover      : True
ParentHasSecureDelegation     : False
SecureDelegationPollingPeriod : 12:00:00
PropagationTime               : 2.00:00:00
SignatureInceptionOffset      : 01:00:00
```

In this example, the secure.contoso.com zone is viewed on DC1. Because DC1 is The Key Master, the value of **IsKeyMasterServer** is **True**. If a different primary authoritative DNS server is specified on the command line, this value would be **False** to indicate that DNSSEC information is being viewed on a server that is not the Key Master. If can be important if zone settings have recently changed and are not yet replicated to other primary DNS servers. You cannot view DNSSEC properties of a signed zone on a secondary DNS server.

Several other Windows PowerShell cmdlets are also available to view more detailed information for individual DNSSEC settings. For more information, see Appendix B: Windows PowerShell for DNS Server.

Review the status of signing keys

A DNSSEC key management strategy includes planning for key generation, key storage, key expiration, and key replacement. Together, key expiration and replacement in DNSSEC is called key rollover. In Windows Server 2012, key management is made easier with simple and flexible key generation, Active Directory storage and replication, an automated key rollover.

You can use the Get-DnsServerSigningKey cmdlet to view the status of KSKs and ZSKs for a signed zone. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com | fl *


KeyId                         : 7479d44e-70ff-48d3-8f64-95622692f647
IsRolloverEnabled             : True
ActiveKey                     : {68550441-8E4D-4D99-921C-BC1807FFF7ED}
CryptoAlgorithm               : RsaSha256
CurrentRolloverStatus         : NotRolling
CurrentState                  : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
```

```
KeyLength                     : 2048
KeyStorageProvider            : Microsoft Software Key Storage Provider
KeyType                       : KeySigningKey
LastRolloverTime              :
NextKey                       :
NextRolloverAction            : Normal
NextRolloverTime              : 12/3/2015 2:05:12 PM
RolloverPeriod                : 755.00:00:00
RolloverType                  : DoubleSignature
StandbyKey                    : {287C4DF6-E3A4-40F0-B09A-F7287343F2A9}
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      : root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm, CurrentRolloverStatus,
CurrentState...}
CimSystemProperties           : Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                         : 00be2f93-ae8b-4381-8165-74f9028210d6
IsRolloverEnabled             : True
ActiveKey                     : {38C23E3E-6A6C-4C40-A792-208962C8EF21}
CryptoAlgorithm               : RsaSha256
CurrentRolloverStatus         : NotRolling
CurrentState                  : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                     : 1024
KeyStorageProvider            : Microsoft Software Key Storage Provider
KeyType                       : ZoneSigningKey
LastRolloverTime              :
NextKey                       : {ABF62EF6-3F6D-43D5-A09B-CECEF22E99C8}
NextRolloverAction            : Normal
```

```
NextRolloverTime              : 2/6/2014 2:05:12 PM
RolloverPeriod                : 90.00:00:00
RolloverType                  : PrePublish
StandbyKey                    :
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      : root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm, CurrentRolloverStatus,
CurrentState...}
CimSystemProperties           : Microsoft.Management.Infrastructure.CimSystemProperties
```

To limit the output of this command to specified signing keys and options, use the **KeyId** parameter with the **Select** option. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com -KeyId 925795c2-4772-48d8-
b0f0-0d79c8a7e5c7 | Select CurrentState, CurrentRolloverStatus, NextRolloverAction,
NextRolloverTime


CurrentState                  CurrentRolloverStatus             NextRolloverAction
NextRolloverTime
------------                  ---------------------             ------------------
----------------
Active                        NotRolling                        Normal
4/16/2014 6:09:37 PM
```

Since any cryptographic security protocol like DNSSEC can be subject to brute force attacks, it is advisable to roll over (replace) zone signing keys on a regular basis. Several properties relating to rollover status of a signing key are displayed by the **Get-DnsServerSigningKey** cmdlet, including **CurrentRolloverStatus**, **LastRolloverTime**, **NextRolloverAction**, **NextRolloverTime**, and others. Additionally, DNS events are logged each time DNSSEC keys are rolled over, or when keys do not roll over because of a problem.

Important

It is a best practice to monitor key rollovers and ensure that they are successful. If a key rollover does not occur, an administrator must identify and correct the problem, then initiate a manual key rollover. The DNS server will not attempt to repeat a signing key rollover attempt that has failed.

Modify DNSSEC properties

You can use the Set-DnsServerDnsSecZoneSetting cmdlet to configure settings for a signed zone. See the following example:

```
PS C:\> Set-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com -ComputerName DC1 -
PropagationTime 1:00:00 -Verbose
VERBOSE: Modifies the DNSSEC properties for the zone secure.contoso.com on server DC1.
VERBOSE: PropagationTime successfully set on server DC1.
```

In this example, the propagation time for secure.contoso.com zone is modified on DC1. This change must be made on DC1 because it is The Key Master for the secure.contoso.com zone. You cannot specify a non-Key Master DNS server on the command line to modify DNSSEC properties.

As previously stated, some DNSSEC settings cannot be modified when a zone is signed. Changing these settings requires that the zone first be unsigned, and then signed again.

For information about additional Windows PowerShell cmdlets that you can use to modify DNSSEC settings, see Appendix B: Windows PowerShell for DNS Server.

See also

DNS Servers

DNS Clients

Trust Anchors

The NRPT

Trust Anchors

A trust anchor (or trust “point”) is a public cryptographic key for a signed zone. Trust anchors must be configured on every non-authoritative DNS server that will attempt to validate DNS data. This section discusses some key concepts about trust anchors.

In this section

Types of trust anchors

Delegations and the chain of trust

- The delegation signer (DS) record

Working with trust anchors

Trust anchor status

Types of trust anchors

Trust anchors can be of two types: DNSKEY or DS. Only DNSKEY trust anchors were supported in Windows Server 2008 R2. Both DNSKEY and DS trust anchors are supported in Windows Server 2012 and Windows Server 2012 R2. For more information about DNSKEY, DS, other DNSSEC-related resource records, and how to view them, see DNSSEC-related resource records.

A DNSKEY trust anchor is a public key from a DNSKEY resource record. It is typically larger in size than a DS trust anchor. A DS trust anchor is a hash of a public key. Because it is only a hash, DS trust anchors are typically smaller. The following are examples of DS and DNSKEY trust anchors, respectively:

```
secure.contoso.com      DS   32801 8 1 (
                        4E1FD2F988BC942E666132250395123B0A54
                        D246 )
                        DS   32801 8 2 (
                        92838CE772067BF04FA942C44C952384267C
                        8D8091A700278D9DC5244F973D88 )
                        DS   53750 8 1 (
                        1D423EBD2C45D7CD7B957C7D7E7795C1C726
                        3275 )
                        DS   53750 8 2 (
                        1073224EBC9ECFEF5C335A4AD10721EA2894
                        63A08CE5E820C511A48220067A2F )


secure.contoso.com   1546847232   DNSKEY   257   3   7   (
                        AwEAAb4v+rZFLosf4qV1/0KZrRcPdj/ljk+Y
                        56XQtuIjGzZG+HH3ys8zhGwJ2WRH3AyRlzLQ
                        5ebMVP12Fmacx5UjhLAydT3rrCbqmHvrrLq9
                        7yJa9Sgeakb/0bwUIE2K2n/g73RFAu8eM0Ut
                        geLZAGkRxdvS8tTAHgoeVCmwbZitxwjVSsZ1
                        cEK7yQJKspwghszXBfonC202GvgCp6F30J4E
                        1LjVJsqPo9ZrccFpFZS4kSEfEz45YXPP/VBx
                        IFRgejdtBL9j6yqf4qo/GynjPSWJY3KOLn/r
                        2wWzujxTLIdpXUdowXGPa8jJU/tDi9avo7+M
                        XIF3LBpvF9derokq/3L6AmE=
```

```
    ) ; key tag = 37458
1546847232   DNSKEY   257   3   7   (
    AwEAAdwL+faFmv/JToQFA94PnqfXI/ZCYvEq
    S5jxmLeZKjl5/xMA2hhaGeKtmS8yTMqeYWm5
    p0MkzDKFMyisjF03N7aBsgLy33GVAea2ipxt
    hjec14q/M37KTgAFh2OyjM4z+Rf+osip0dsv
    KVl5vXgYYFXqf52aEoYiYe2PfWM0mJtl1NlA
    uYAu4h3140AHynIIJYJkCV+ptHuSECEFYNBf
    aO1Qg9KKVvH7sH9Crfz2sHb0+sVsi2+IgBjD
    ILXCW8af7bowl0yhL1qB+njegKxzI8IpnOA+
    MDWZvaJCtTQGPUjTwG0qjWBzO4bONHTLqkyR
    /jvPTciJ+Qah3h5apITtwL8=
    ) ; key tag = 9468
```

For information about the components of a trust anchor, see Working with trust anchors in this topic.

A validating DNS server that uses a DS trust anchor must query the authoritative DNS server to obtain the full DNSKEY resource record set (RRSet). This is called “priming” of the trust anchor, and can be an advantage because it forces the validating server to acquire up-to-date information.

You can install one or both types of trust anchors on a validating DNS server. Active Directory-integrated DNS servers store trust anchors in Active Directory. Non Active Directory-integrated DNS servers store trust anchors in the text file: **%windir%\system32\dns\TrustAnchors.dns**.

DNSSEC validation is enabled by default on the **Advanced** tab of DNS server properties. You can modify and check the status of this setting with Dnscmd.exe. See the following example:

```
PS C:\> dnscmd /info /enablednssec

Query result:

Dword:  0 (00000000)

Command completed successfully.
```

In the previous example, DNSSEC validation is disabled. To enable DNSSEC validation, you must configure the value to 1:

```
PS C:\> dnscmd /config /enablednssec 1

Registry property enablednssec successfully reset.

Command completed successfully.


PS C:\> dnscmd /info /enablednssec

Query result:
```

```
Dword:  1 (00000001)

Command completed successfully.
```

You can modify this setting in DNS Manager or using the Dnscmd.exe tool.

Delegations and the chain of trust

A delegation is a referral of authority from one DNS server to a different DNS server for a specific namespace. Delegations are commonly used to assign authority for child zones.

For example, the Internet root DNS servers delegate DNS authority for .com to generic top-level domain (gTLD) servers. See the following example using nslookup.exe:

```
PS C:\> nslookup -q=NS microsoft.com. a.root-servers.net

com   nameserver = m.gtld-servers.net

com   nameserver = l.gtld-servers.net

com   nameserver = k.gtld-servers.net

com   nameserver = j.gtld-servers.net

com   nameserver = i.gtld-servers.net

com   nameserver = h.gtld-servers.net

com   nameserver = g.gtld-servers.net

com   nameserver = f.gtld-servers.net

com   nameserver = e.gtld-servers.net

com   nameserver = d.gtld-servers.net

com   nameserver = c.gtld-servers.net

com   nameserver = b.gtld-servers.net

com   nameserver = a.gtld-servers.net

…
```

In turn, the gTLD servers delegate authority for microsoft.com to Microsoft's authoritative DNS servers. See the following example:

```
PS C:\> nslookup -q=NS microsoft.com. a.gtld-servers.net

microsoft.com   nameserver = ns3.msft.net

microsoft.com   nameserver = ns1.msft.net

microsoft.com   nameserver = ns5.msft.net

microsoft.com   nameserver = ns2.msft.net

microsoft.com   nameserver = ns4.msft.net

…
```

The following example uses the **Get-DnsServerZoneDelegation** cmdlet to verify delegation for *corp.adatum.com* and *finance.corp.adatum.com*:

```
PS C:\> Get-DnsServerZoneDelegation -Name adatum.com -ChildZoneName corp -ComputerName
dc2.contoso.com

ZoneName                            ChildZoneName                       NameServer
IPAddress
--------                            -------------                       ----------
---------
adatum.com                          corp.adatum.com.                    dc2.contoso.com.
192.168.0.2



PS C:\> Get-DnsServerZoneDelegation -Name corp.adatum.com -ChildZoneName finance -
ComputerName dc2.contoso.com

ZoneName                            ChildZoneName                       NameServer
IPAddress
--------                            -------------                       ----------
---------
corp.adatum.com                     finance.corp.adatum.com.            dc2.contoso.com.
192.168.0.2
```

Delegations are important because they establish a chain of authentication for child zones. If all zones in the chain are signed with DNSSEC, resolving DNS servers can have only a single delegation signer (DS) trust anchor installed, provided that proper DS records are available in the parent zone.

With delegations, it is not always necessary to install a trust anchor for every zone that a caching DNS server will validate. Under certain conditions, you can build a chain of authentication (a chain of trust) where it is sufficient for the DNS server to have only a DS trust anchor for the parent zone installed.

However if the chain of trust is broken, then a child zone must have its own trust anchor installed. For example, DNS responses for the signed zone A.B.C.com can be validated using a trust anchor for the signed zone C.com, provided that B.C.com is also signed. If B.C.com is not signed, then C.com becomes an island of trust and requires its own trust anchor. The following diagram illustrates this scenario:

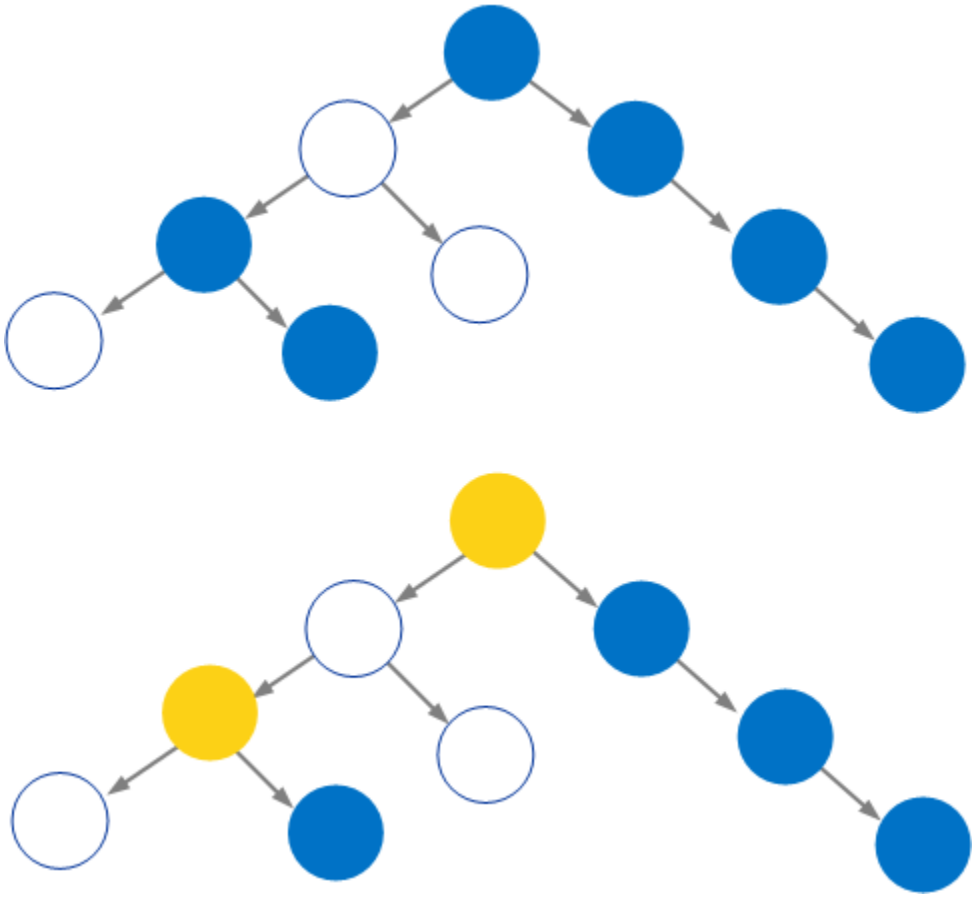

In scenarios where islands of trust are present, a trust anchor must be configured for the root of each island. In the previous diagram, zones shown in blue are signed while those in white are not. The right side of the tree is completely signed; the left side of the tree is intermittently signed. In order for a DNS server to be able to perform DNSSEC validation for all the signed zones, you must configure trust anchors for the zones in yellow.

Additional trust anchors for the other signed zones can be added to a DNS server, but the trust anchors for the zones in yellow are sufficient for full validation of all signed zones.

After a zone is DNSSEC-signed, and if the parent of the zone is also DNSSEC-signed, the signed delegation records must be added to the parent zone and the parent zone must be re-signed. This process must be performed every time a new child zone is signed for the first time, or if a child zone is re-signed with a new key signing key. If you own a signed zone but do not own the children of the zone, and if the children zones are in the process of being DNSSEC-signed one at a time, you must re-sign your zone after adding the delegation records each time a new child zone is signed. However, the process of signing multiple zones can be optimized if you own the parent as well as the children zones that are to be signed.

The delegation signer (DS) record

Like the DNSKEY resource record, the DS record can be used to create a trust anchor for a signed zone. The DS record is smaller in size than a DNSKEY because it contains only a hash of the public key.

The DS resource record is not added to a zone during the signing process like some DNSSEC-related resource records, even if a delegation already exists in the zone. To add a DS record, you must manually add or import it. Fortunately, the DS resource record set (DSSET) is automatically added as a file to the Key Master when a zone is signed. You can use this file with the **Import-DnsServerResourceRecordDS** cmdlet to import DS records to the parent zone. See the following example:

```
PS C:\> Import-DnsServerResourceRecordDS -ZoneName adatum.com -DSSetFile
"c:\windows\system32\dns\dsset-corp.adatum.com"
PS C:\> Get-DnsServerResourceRecord -ZoneName adatum.com -RRType DS


HostName                  RecordType Timestamp            TimeToLive      RecordData
--------                  ---------- ---------            ----------      ----------
corp                      DS         0                    01:00:00
[58555][Sha1][RsaSha1NSec3]
corp                      DS         0                    01:00:00
[58555][Sha256][RsaSha1NSec3]
corp                      DS         0                    01:00:00
[63513][Sha1][RsaSha1NSec3]
corp                      DS         0                    01:00:00
[63513][Sha256][RsaSha1NSec3]
```

In the previous example, DS records for the child zone, corp.adatum.com were imported into the parent zone, adatum.com, using the DSSET file located in the **c:\windows\system32\dns** directory. The DSSET file was located in this directory because the local DNS server is the Key Master for the child zone. If the Key Master DNS server for a child zone is not the same computer as the primary authoritative DNS server for the parent zone where you are adding the DS record, you must obtain the DSSET file for the child zone and make it available to the primary authoritative server for the parent zone. Alternatively, you can manually add DS records.

DS records must be added to each parent zone in a chain of authentication. For example, DS records for finance.corp.adatum.com must be added to corp.adatum.com to complete a chain of authentication from finance.corp.adatum.com to adatum.com. See the following example:

```
PS C:\> Import-DnsServerResourceRecordDS -ZoneName corp.adatum.com -DSSetFile
"c:\windows\system32\dns\dsset-finance.corp.adatum.com"
PS C:\> Get-DnsServerResourceRecord -ZoneName corp.adatum.com -RRType DS
```

```
HostName                  RecordType Timestamp            TimeToLive      RecordData
--------                  ---------- ---------            ----------      ----------
finance                   DS         0                    01:00:00
[58990][Sha1][RsaSha256]
finance                   DS         0                    01:00:00
[58990][Sha256][RsaSha256]
finance                   DS         0                    01:00:00
[40231][Sha1][RsaSha256]
finance                   DS         0                    01:00:00
[40231][Sha256][RsaSha256]
```

If the chain of authentication is completed with proper delegations in each zone and DS records added to the parent, a recursive DNS server can have a trust anchor present for only the parent zone and provide validation for all child zones. In the following example, a DNS client requires validation for the *adatum.com* namespace:

```
PS C:\> Get-DnsClientNrptPolicy -Effective



Namespace                       : .adatum.com
QueryPolicy                     : QueryIPv6Only
SecureNameQueryFallback         : FallbackPrivate
DirectAccessIPsecCARestriction  :
DirectAccessProxyName           :
DirectAccessDnsServers          :
DirectAccessEnabled             : False
DirectAccessProxyType           :
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired  :
NameServers                     :
DnsSecIPsecCARestriction        :
DnsSecQueryIPsecEncryption      :
DnsSecQueryIPsecRequired        : False
DnsSecValidationRequired        : True
NameEncoding                    :
```

A resolving DNS server with IP address of 192.168.0.4 has a trust anchor for *adatum.com*:

```
PS C:\> Get-DnsServerTrustPoint -ComputerName 192.168.0.4
```

```
TrustPointName                          TrustPointState      LastActiveRefreshTime
NextActiveRefreshTime
--------------                          ---------------      ---------------------     -
--------------------
adatum.com.                             Active               1/22/2014 2:05:50 PM
1/22/2014 2:06:06 PM
```

DNSSEC validation is successful for DNS queries in the *finance.corp.adatum.com* zone because this zone and *corp.adatum.com* have the proper delegations and DS records:

```
PS C:\> Resolve-DnsName -Name mail.finance.corp.adatum.com -Type A -Server 192.168.0.4 -
dnssecok

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
mail.finance.corp.adatum.com                   A      3599  Answer     192.168.0.150

Name        : mail.finance.corp.adatum.com
QueryType   : RRSIG
TTL         : 3599
Section     : Answer
TypeCovered : A
Algorithm   : 8
LabelCount  : 5
OriginalTtl : 3600
Expiration  : 1/31/2014 10:03:18 PM
Signed      : 1/21/2014 9:03:18 PM
Signer      : finance.corp.adatum.com
Signature   : {80, 88, 162, 208...}

…
```

Working with trust anchors

Because trust anchors represent public keys for corresponding zones, trust anchors must be updated if the public keys for a zone change, such as happens when a zone is re-signed. Additionally, if trust anchors are distributed after a zone is signed and the zone becomes

unsigned at a later date, trust anchors for the zone must be removed. If trust anchors are not removed when a zone is unsigned, DNS servers will continue to attempt validation of DNS responses for the zone, but this validation will fail and therefore DNS resolution for the zone will fail.

Trust anchors can be updated manually, or you can use the automated update mechanism described in RFC 5011 (http://tools.ietf.org/search/rfc5011).

Important
Windows Server 2012 and Windows Server 2012 R2 fully support RFC 5011. Support for RFC 5011 is enabled on the authoritative DNS server on a zone-by-zone basis. However, trust anchor updates are only supported for automated key rollover. If you manually re-sign or unsign a zone, or manually roll zone signing keys, trust anchors on validating DNS servers are not automatically updated.

Trust anchors that have been enabled for RFC 5011 automatic updates have information stored in the file: **%windir%\system32\dns\RFC5011.csv** on the DNS server where the trust anchor is installed.

When you sign a zone on a DNS server running Windows Server 2012 or later, the public DNSKEY and DS resource record sets (the set of all resource records of that type, also known as the RRSET) are stored as text files in the **%windir%\system32\dns** directory on the Key Master. The DNSKEY RRSET (or KEYSET) is stored as a file with the name *keyset-<zone name>* and the DS RRSET (or DSSET) is stored as a file named *dsset-<zone name>*.

For example, after signing the secure.contoso.com zone, the following files are created on the Key Master:

Key type	Filename
DNSKEY	*keyset-secure.contoso.com*
DS	*dsset-secure.contoso.com*

During zone signing, DNSKEY resource records are also added to the zone along with RRSIG, NSEC, and other records. DS records are not automatically added to a zone.

These KEYSET and DSSET files can be used to export trust anchors for the zone, and to manually add a trust anchor. For more information, see Procedure: Export a Trust Point and Procedure: Add a Trust Point.

The KEYSET (DNSKEY) file has the following format: `<zone name> <TTL> IN DNSKEY <flags> <protocol> <algorithm> <Base64Data>; key tag = <key tag>`. See the following example:

```
secure.contoso.com.   3600   IN   DNSKEY   257 3 8 (
                        AwEAAbMK6im/lrLB546v9jl41p/g7/30X9Cn
                        k5anwOnQu8JBOYt0CUiXci1lxO/280NHB/Eu
```

```
                1r7Nk0AQPwvZ7Dqbdd7eqvKX/WJgK85UtONR
                PjhbCNe4vj7Qr8GmUW/aQdDIZqqdstLJINfE
                k3GgOFslSUkoWpP+z9BwnuEnv6jFkguLikj6
                X9aM1N0bjLMDvi34PlRKGxOA6WsDyN96abzz
                Iuj00LjKBL4+0JBwuRB+tIOAMLLsXtLI4Dby
                Zz6/ewZHI0FELbDnUDj5hqCWU+J2IBTbs4mW
                WKRHpTdZtNYda8Oy9+SMQJh/pIepKAjrJkzM
                wZWFgg0qqn5QXVIms3T/pZk=
                ) ; key tag = 62111
              3600   IN   DNSKEY   257 3 8 (
                AwEAAY/kENeHkSmtdBzPhiHiDGPcbwA7nkau
                MsASiXOMIM+2btAYq6GczO45kuemvXmJU1o6
                eyeJmphDGm649UOjOiebIokJ31m3K1mymrpF
                R6DfY2WTwLMtnPuiiyvif9qcF5yKoVthownG
                kG3/vJbDyRGONteTtJFOZZDtDHLOh624jz73
                KUds4uq6VWZ+UhCZFneNNtz2e4NyjxGSf5ZK
                kKOpmj/Q2veAbV8Ck5fieydOMvOi2HM6oop9
                YrAGRUvmwH2Bg/enDiWr7+FDCEzem4bP4QDT
                ng6qjSvNGfRTtq6x4zySPcd7BMv1Q5nBpdjU
                UCNXRU37nTXB7/sSKQCmFic=
                ) ; key tag = 8000
```

Two keys will be present in the file. If the flag field is an odd number, the key represents a KSK (secure entry point). If the flag field is an even number, the key is a ZSK. In the previous example, both keys are KSKs because the flag field is 257. The protocol is always DNSSEC, represented by the number 3, and the algorithm is RSA/SHA-256, corresponding to the number 8.

The DSSET (DS) file has the following format: `<zone name> <TTL> IN DS <key tag> <algorithm> <type> ( digest )`. See the following example:

```
secure.contoso.com.   3600   IN   DS   16945 8 1 (
                        86A50FB9A986CB32BF878FFFEF3DC00D5732
                        3ED8 )
                      3600   IN    DS   16945 8 2 (
                        59FFBE7F9C06FC939C772DDAA8D7A5F3B316
                        4082059FBC5D689DFD6B0290310E )
                      3600   IN    DS   39301 8 1 (
```

```
                565B5B1FC1616CF84514D28814D34A1C2397
                8008 )
            3600   IN    DS   39301 8 2 (
                6A634D28ABF17E723EDBCD981120DC736A08
                CF7170E849303B9357B52D18F9D3 )
```

If you add a DS trust anchor to a DNS server and the DS record is present in the DNS zone - recall that DS records are not automatically added during zone signing – the DNS server will use the digest in the DS record to acquire a DNSKEY for the zone.

Trust anchor status

To view the status of a trust anchor on a DNS server, you can use the DNS Manager console or Windows PowerShell. Trust anchors are organized and displayed in the DNS Manager console tree under Trust Points. The Windows PowerShell cmdlets **Get-DnsServerTrustAnchor** and **Get-DnsServerTrustPoint** can also provide status of trust anchors. The Get-DnsServerTrustAnchor cmdlet requires that you specify a namespace, whereas Get-DnsServerTrustPoint will display all configured trust anchors on the DNS server if a name is not specified. See the following examples:

```
PS C:\> Get-DnsServerTrustAnchor -Name adatum.com


TrustAnchorName                TrustAnchorType      TrustAnchorState     TrustAnchorData
---------------                ---------------      ----------------     ---------------
adatum.com.                    DNSKEY               Valid
[48955][DnsSec][RsaSha256][AwEAAcTxpB/tf4DA...
adatum.com.                    DNSKEY               Valid
[2504][DnsSec][RsaSha256][AwEAAYkCDUmGSkh+y...


PS C:\> Get-DnsServerTrustPoint


TrustPointName                           TrustPointState      LastActiveRefreshTime
NextActiveRefreshTime
--------------                           ---------------      ---------------------     -
--------------------
.                                        DSPending            1/24/2014 1:02:28 PM
1/24/2014 1:02:44 PM
adatum.com.                              Active               1/24/2014 1:02:28 PM
1/24/2014 1:02:44 PM
```

```
corp.adatum.com.                    Active              1/24/2014 1:02:28 PM
1/24/2014 1:02:44 PM
```

In the previous example, the root (.) trust anchor is not active, instead displaying a status of **DSPending**. The following trust anchor states are possible:

Status	Description
Valid or Active	Trusted for DNSSEC validation
Add Pending	Not yet trusted for DNSSEC validation, but will become valid after an add hold-down period.
Missing	Trusted for DNSSEC validation, but was not found during the last active refresh cycle.
Revoked	Revoked by an administrator, and no longer valid for DNSSEC validation.
DS Pending	Waiting for a matching DNSKEY, which will become a valid trust anchor.
DS Invalid	Matched with a DNSKEY that cannot be a valid trust anchor.

In the previous example, the DS trust anchor for the root (.) zone is in a DS Pending state because access to the root DNSKEY record is currently blocked.

See also

DNS Servers

DNS Clients

DNS Zones

The NRPT

The NRPT

The Name Resolution Policy Table (NRPT) in Windows Server 2012 and Windows Server 2012 R2 enables you to enforce name resolution policies on security-aware DNS clients. For information about configuring the NRPT, see Procedure: Configure the NRPT.

In this section

Introduction to the NRPT

[NRPT rule processing]
[Filter NRPT policy]
[View NRPT policy]

Introduction to the NRPT

The NRPT is a table that contains rules that you can configure to specify DNS settings or special behavior for names or namespaces. The NRPT can be configured using the Group Policy Management Editor under **Computer Configuration\Policies\Windows Settings\Name Resolution Policy**, or with Windows PowerShell. See the following example:

Figure 1: The Name Resolution Policy Table in Windows Server 2012 R2

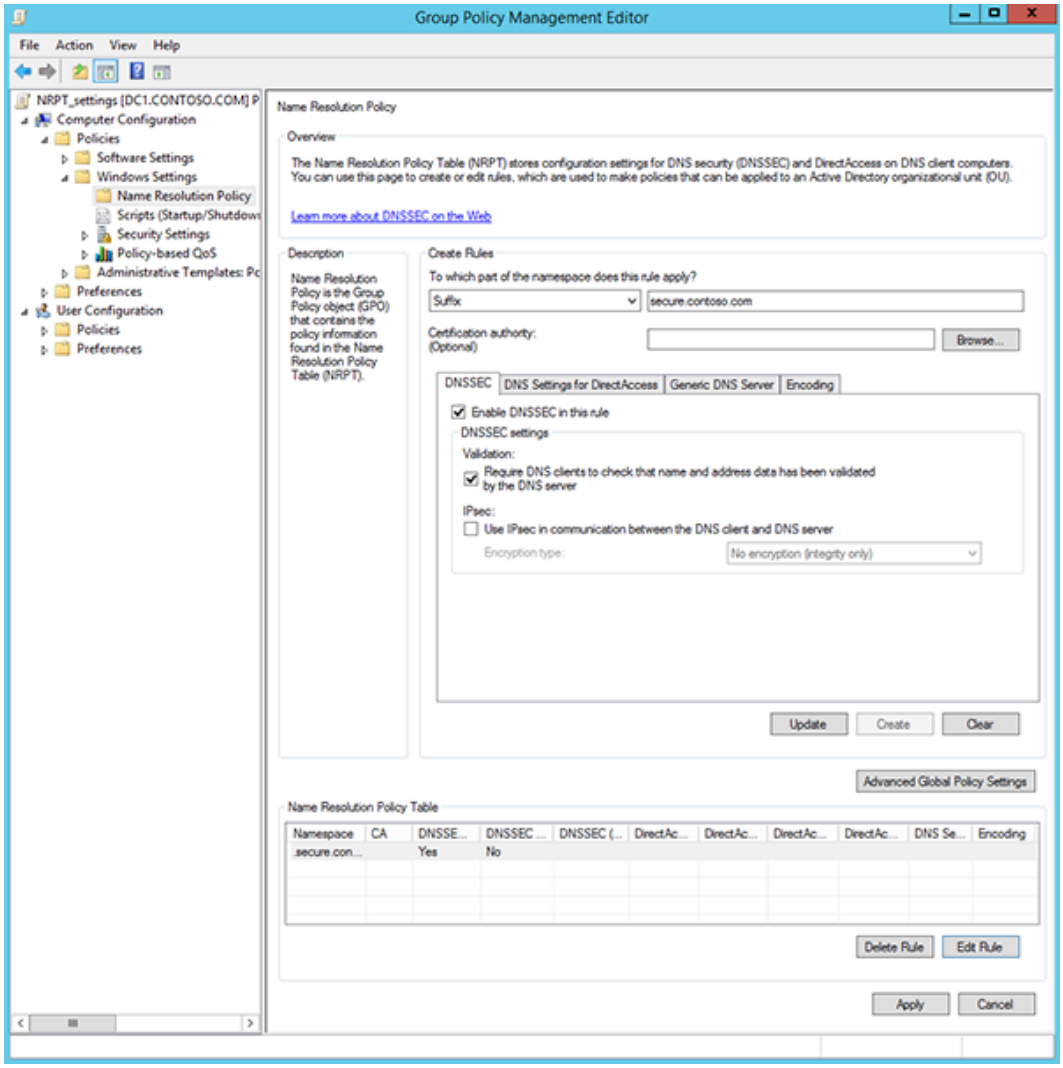

When performing DNS name resolution, the DNS Client service checks the NRPT before sending a DNS query. If a DNS query matches an entry in the NRPT, it is handled according to settings in the policy. Queries that do not match an NRPT entry are processed normally. You can use the NRPT to require that DNSSEC validation is performed on DNS responses for queries in the namespaces that you specify.

The namespace to which a policy applies can be configured using the following options:

Option	Usage
Suffix	The policy applies to anything ending in *.contoso.com*. This includes child domains.
Prefix	The policy applies only to a host name. This policy will be triggered only if the hostname portion of a DNS query matches the flat name configured here. For example, an entry of *server1* will apply to both *server1.contoso.com* and *server1.nwtraders.com*.
FQDN	The policy applies only to the specified host. This is not the FQDN of a domain, but an FQDN of a host computer.
Subnet (IPv4)	This configures a policy which applies to reverse IPv4 lookup queries.
Subnet (IPv6)	This configures a policy which applies to reverse IPv6 lookup queries.
Any	This configures the default policy.

Note
The NRPT can be used to create different types of rules by choosing one of four available tabs. This topic only discusses rules configured using the **DNSSEC** tab, that apply to a DNSSEC deployment.

Client computers that are not security-aware, such as computers running Windows XP, will ignore NRPT settings. For more information, see DNS Clients.

NRPT rule processing

NRPT rules can be configured in a local, site, domain, or OU-linked Group Policy Object (GPO). Normal Group Policy processing rules apply. See Group Policy processing and precedence (http://technet.microsoft.com/library/cc785665.aspx) for more information. This means that NRPT rules are processed in the following order:

1. Local GPO

2. Site-linked GPOs
3. Domain-linked GPOs
4. Organizational unit-linked GPOs

If NRPT rules are present in multiple GPOs that apply to the same user or computer, the rules are merged, with more discrete rules taking precedence over more general rules. For example, a rule applying to a child namespace will take precedence over a rule for the parent namespace. See the following example:

GPO name	Namespace	DNSSEC enabled	Validation required
Default Domain Policy	contoso.com	√	X
NRPT_settings	secure.contoso.com	√	√

In this example, DNSSEC validation will be required for DNS queries that have a suffix of secure.contoso.com (ex: www.secure.contoso.com). This includes queries in a child domain (for example: www.corp.secure.contoso.com). However, validation is not required for *.contoso.com or *.corp.contoso.com. Therefore, if parent and child domains require different name resolution policy, you must explicitly create policies for the child namespace.

Caution

NRPT rules do not overwrite one another. If two rules are created in two different GPOs that apply to the same namespace for the same user or computer, a conflict will occur and neither rule will be applied. If the rules apply to the same user or computer for different namespaces, they are merged. This rule does not apply to local Group Policy, however. If any NRPT settings are configured in domain Group Policy, then all local Group Policy NRPT settings are ignored.

Filter NRPT policy

By default, new domain-linked Group Policy Objects (GPOs) apply to the authenticated users group. You can apply NRPT policy to selected groups, users, and computers by removing the authenticated users group from Security Filtering and then adding a custom user, group, or computer. See the following example:

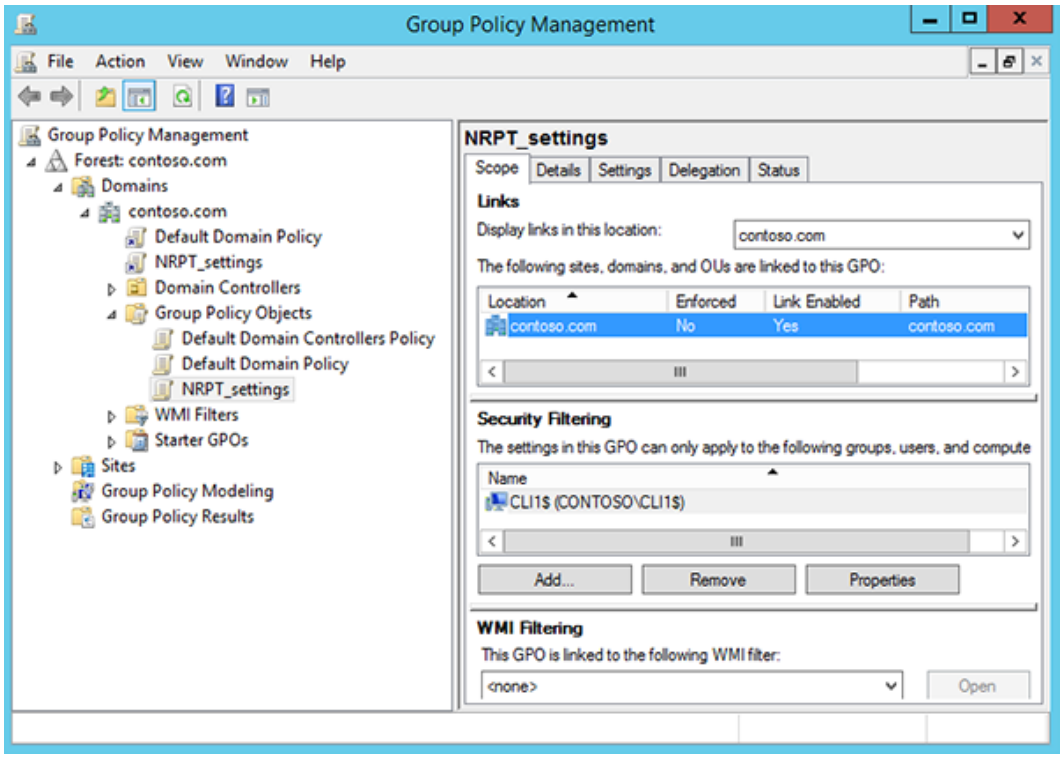

In this example, a new GPO is created named **NRPT_settings**. This GPO only applies to the computer **CLI1** using security filtering.

View NRPT policy

You can use the Get-DnsClientNrptPolicy cmdlet to view NRPT policies. See the following example:

```
PS C:\> Get-DnsClientNrptPolicy -Effective


Namespace                       : .contoso.com
QueryPolicy                     : QueryIPv6Only
SecureNameQueryFallback         : FallbackPrivate
DirectAccessIPsecCARestriction  :
DirectAccessProxyName           :
DirectAccessDnsServers          :
DirectAccessEnabled             : False
DirectAccessProxyType           :
```

```
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired   :
NameServers                      :
DnsSecIPsecCARestriction         :
DnsSecQueryIPsecEncryption       :
DnsSecQueryIPsecRequired         : False
DnsSecValidationRequired         : False
NameEncoding                     :

Namespace                        : .secure.contoso.com
QueryPolicy                      : QueryIPv6Only
SecureNameQueryFallback          : FallbackPrivate
DirectAccessIPsecCARestriction   :
DirectAccessProxyName            :
DirectAccessDnsServers           :
DirectAccessEnabled              : False
DirectAccessProxyType            :
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired   :
NameServers                      :
DnsSecIPsecCARestriction         :
DnsSecQueryIPsecEncryption       :
DnsSecQueryIPsecRequired         : False
DnsSecValidationRequired         : True
NameEncoding                     :
```

The cmdlet gets policy that applies to the local computer. In this example, the output indicates that DNSSEC validation is required for *.secure.contoso.com, and is not required for *.contoso.com.

See also

DNS Servers

DNS Clients

DNS Zones

Trust Anchors

DNSSEC Deployment Planning

The following topics provide planning guidance. Review the information in these topics before you sign a DNS zone with DNSSEC:

Why DNSSEC: Consider the reasons you are deploying DNSSEC. Possible reasons include protection against spoofing attacks and compliance with government corporate policies.

Stage a DNSSEC Deployment: Before fully deploying DNSSEC in a production environment, try signing a test zone.

DNSSEC Performance Considerations: DNSSEC signing and validation requires processor resources on DNS servers. DNSSEC-signed zones and DNS responses are also larger in size, requiring additional disk and network resources.

DNSSEC Requirements: Before you deploy DNSSEC, verify that your DNS infrastructure meets specifications.

See also

Overview of DNSSEC

DNSSEC in Windows

Why DNSSEC

A DNS zone that is signed with DNSSEC is more secure, but is also associated with more administrative tasks. Before you sign a zone with DNSSEC, consider both the potential advantages and disadvantages.

In this section

Why to deploy DNSSEC

- DNS spoofing

Signing domain zones

Signed zones

Why to deploy DNSSEC

In environments that do not employ security technologies such as IPsec or HTTPS, the DNS protocol can be vulnerable to attack due to an inherent lack of authentication and integrity checking of data that is exchanged between DNS servers or sent to DNS clients. As originally designed, DNS itself does not offer any form of security and is vulnerable to spoofing and man-in-the-middle attacks. An attacker that has compromised a DNS server can gain access to all

network communications sent by a targeted host. If DNS servers are vulnerable to attack, securing them with DNSSEC can be critical.

DNSSEC includes changes to client and server DNS components, enabling DNS data to be cryptographically signed and to enforce name validation policies that protect DNS communications. With DNSSEC, a DNS server can validate responses that it receives as genuine. By validating DNS responses, DNS servers and clients are protected against the single greatest vulnerability in DNS: DNS spoofing.

DNS spoofing

DNS spoofing is the practice of assuming the DNS name of another system either by corrupting a name service cache or by compromising a DNS server for a valid domain. When a DNS resolver sends a remote query, it tags the query with a 16-bit XID (transaction ID) value in the DNS packet header and expects that the remote DNS server will respond on the same port with the same XID value. The query is typically sent over UDP, which is more vulnerable to attacks than TCP because of TCP's three step "handshake." Generally, TCP is only used after a UDP response has been truncated. When the resolver receives a UDP DNS response it can only weakly verify that the response is authentic by matching the following parameters:

- **Remote DNS server address**: This check is often disabled by default since many network devices make it appear that valid responses come from a different address than the query was sent to, making it even easier to spoof a DNS response.
- **Port on which the packet was received**: The resolver will typically send from an ephemeral port to port 53. DNS servers respond from port 53 to whatever source port was used by the requestor.
- **Packet XID value**: The XID value is set in the request by the resolver and must match the value contained in the response. A strongly random value can and should be used for the XID but it is only 16 bits long. The XID value, like the rest of the DNS packet, is sent in the clear.
- **Query name and type**: The DNS server echoes the query name and type in the question section of the DNS response.

If an attacker has access to a DNS client or server's network traffic, they may be able to guess that a DNS client or server has sent a DNS query and is waiting for a DNS response. When they have determined this to be true, the attacker can send one or more spoofed DNS response packets and attempt to beat the authentic response back to the requestor. An attacker (or a botnet under their control) can attack DNS clients and DNS servers that send remote DNS queries. Since the attacker controls the time-to-live (TTL) assigned to the spoofed response, once the attack is successful it can persist for days, weeks, or even longer.

The following figure shows how DNSSEC validation can be used to prevent spoofing attacks on DNS clients and recursive DNS servers that are meant to redirect client computers to malicious sites.

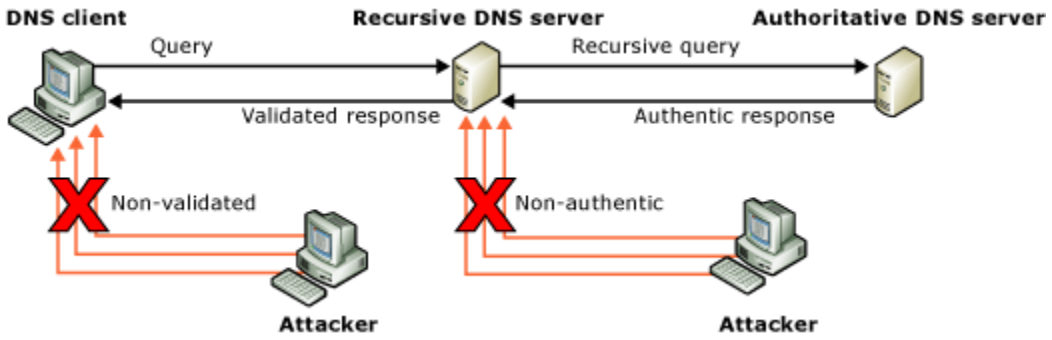

With DNSSEC, a DNS client computer can ignore DNS responses that are not validated as genuine by a recursive DNS server. Recursive DNS servers will also fail to validate non-authentic DNS responses after checking digital signatures. In this way, DNSSEC can be an effective method to prevent DNS spoofing attacks.

Signing domain zones

Public DNS zones that are connected to the Internet and must be available to customers and vendors are particularly vulnerable to attack. These kinds of zones can often benefit from DNSSEC protection. Usually these zones are not domain zones, or otherwise part of an Active Directory Domain Services (AD DS) infrastructure.

Internal zones are typically less vulnerable to attack, either because they are not exposed to the Internet or because other security protocols have been implemented to protect them. These zones might not benefit as much from signing with DNSSEC, and be associated with more administrative burden because of multiple domain-joined client computers that require access to domain services such as LDAP, Net Logon, Kerberos, and others. However, some domain zones are required to be signed, such as those controlled by government agencies. Because internal domain zones often provide support for additional services that are not implemented by public DNS zones, it is even more important to properly manage your DNSSEC deployment when signing these types of zones.

In general, if an internal domain zone is not at risk for attack it is not advisable to sign it with DNSSEC unless required by government or corporate policy.

Signed zones

Because DNSSEC adds additional records to a zone and requires encryption and decryption of DNS messages, signing a zone with DNSSEC increases its administrative burden and is associated with a performance impact on DNS infrastructure. In addition, since any cryptographic security protocol like DNSSEC can be subject to brute force attacks, it is advisable to roll over (replace) cryptographic keys on a regular basis, adding further to management requirements. For more information about requirements for signed zones, see DNS Zones.

The length of a cryptographic key (in DNSSEC, a KSK or ZSK) also affects its ability to be compromised using brute force methods. Longer keys are more difficult to break, but are also associated with a greater performance impact. For more information about the performance impact of different cryptographic algorithms and key lengths, see DNSSEC Performance Considerations.

It is therefore necessary to balance security with manageability and performance when configuring zone signing parameters and deciding whether or not to sign a zone. Do not sign a zone that will not benefit from enhanced security, and if you decide to sign a zone, be sure to stage your DNSSEC deployment carefully so that you fully understand the performance and administrative impacts. For more information about staging a DNSSEC deployment, see Stage a DNSSEC Deployment.

If you decide to sign a zone, Windows Server 2012 and Windows Server 2012 R2 provide excellent DNSSEC support to simplify management tasks and integrate with other systems, including:

- Support for Active Directory-integrated DNS scenarios including DNS dynamic updates in DNSSEC signed zones.
- Support for updated DNSSEC standards, including NSEC3 and RSA/SHA-2.
- Automated trust anchor distribution through Active Directory.
- Automated trust anchor rollover support per RFC 5011.
- Updated user interface with deployment and management wizards.
- Validation of records signed with updated DNSSEC standards (NSEC3, RSA/SHA-2).
- Easy extraction of the root trust anchor.

In Windows Server 2012 R2, DNSSEC support is enhanced with additional Windows PowerShell and key management support. For more information, see What's New in DNS Server in Windows Server 2012 R2.

See also

Stage a DNSSEC Deployment

DNSSEC Performance Considerations

DNSSEC Requirements

Stage a DNSSEC Deployment

You can use the following guidance to gradually deploy DNSSEC on your network.

In this section

Identify goals

DNSSEC pilot

DNSSEC staging

Identify goals

The first step in a DNSSEC staged deployment is to evaluate your goals. Based on these goals, decide if there are zones will benefit from zone signing and other zones that do not need to be signed with DNSSEC.

Include preliminary goals if you are beginning with a pilot deployment, and any additional goals you hope to achieve when DNSSEC is deployed in a production environment. Some possible goals include:

- **Improved security**: DNSSEC can protect you from DNS spoofing attacks.
- **Compliance with corporate policies**: Signing some or all zones might be required by your organization.
- **Compliance with government mandates**: Signing some or all zones might be required by your government.
- **Proof of concept**: You might be deploying DNSSEC only in a test environment in order to learn about the technology and evaluate it for a production deployment in the future.

DNSSEC pilot

Start with a pilot project. A DNSSEC pilot project will enable you to become more familiar with DNSSEC zone signing and administration before you deploy signed zones and configure secure name resolution policies in a larger environment.

Notes

In a Microsoft Active Directory-integrated multi-master DNS environment, all primary authoritative DNS servers will distribute and share DNS information based on the Active Directory zone replication scope. Consequently, if you sign an Active Directory-integrated zone on a DNS server that is also a domain controller, Active Directory will attempt to replicate the zone to other DNS servers in the zone replication scope. If other DNS servers are running Windows Server 2012 or a later operating system, the zone will be also be signed on these servers. To limit the size of your pilot project, sign a zone that is file-backed or has a relatively small zone replication scope.

Once a zone is signed for DNSSEC, the DNS server will explicitly block attempts to change the zone replication scope or zone type. This is primarily to avoid complexities related to key storage. To change the zone replication scope, you must first unsign the zone.

Consider the following factors when testing DNSSEC in your environment:

Scope: Begin with a small number of zones and resource records, or a small number of DNS servers and clients. This enables you to easily test the zone signing parameters that you plan to use and make necessary adjustments. You might wish to only use a lab environment to test DNSSEC before deploying it in production. Before deploying DNSSEC in production, test a zone

that is similar in size to one that will be used in production so that you can evaluate the performance impact.

Validation testing: To test DNSSEC validation, you must deploy trust anchors and configure the NRPT. You should verify that clients can resolve names in the signed zone when DNSSEC validation is successful, and that DNS resolution fails when DNSSEC validation is not successful. For more information about conditions that can cause DNS resolution to succeed or fail for a signed zone, see DNS Clients.

Administration: A critical part of managing a DNSSEC deployment is performing a key rollover and deploying trust anchors. You should test these procedures as part of your DNSSEC pilot. You can also transfer the Key Master role to a different DNS server as part of a disaster recovery simulation.

Performance testing: Review network, disk, and processor usage on the DNS server before and after zone signing. Evaluate performance of the DNS server during zone signing and key rollover. Reload the DNS zone or restart the DNS Server service to review zone loading time.

Education: It is important to provide your support personnel with information about how DNSSEC works, additional administrative tasks associated with signed zones, and procedures to help troubleshoot DNSSEC related issues. New tools and procedures might be developed to help you manage a DNSSEC deployment. DNSSEC related events are recorded in the DNS Server log in Event Viewer (Custom Views\Server Roles\DNS Server).

Best practices: A DNSSEC pilot is ideal for starting a list of best practices for deployment, administration, and troubleshooting of DNSSEC-signed zones.

DNSSEC staging

To allow for staged migration of server operating systems, DNSSEC deployment does not require that all DNS servers are running Windows Server® 2012 or a later operating system to host a signed zone. Although it is recommended to have all DNS servers hosting the signed zone upgraded to Windows Server 2012 or Windows Server 2012 R2 before validation is enabled on resolvers, only one authoritative DNS server running Windows Server 2012 or Windows Server 2012 R2 is required.

Only DNS servers running Windows Server 2012 or a later operating system will host a signed copy of a zone that was signed using Windows Server 2012 or Windows Server 2012 R2. Servers running earlier versions of Windows will host an unsigned version of the zone.

In a mixed-mode Active Directory environment, with some DNS servers running Windows Server 2012 or a more recent operating system and some DNS servers running Windows Server 2008 R2 or a previous operating system, the Active Directory schema must be upgraded before signing a DNS zone. Schema update will occur automatically if you promote a computer running Windows Server 2012 or Windows Server 2012 R2 to be a domain controller. See

For more information about updating the Active Directory schema, see Active Directory Schema Update (http://technet.microsoft.com/library/cc755834.aspx).

You can use the following steps to stage a DNSSEC deployment:

1. **Choose a zone**: Start by signing a single zone or only a few zones. The zones should be small enough to monitor easily, but of sufficient size to enable thorough testing.
2. **Choose a Key Master**: Select a primary authoritative DNS server to be the Key Master for your zone. This is the server where you will initiate zone signing, and should be online at all times.
3. **Choose methods and settings**: Decide if you will use DNS Manager or Windows PowerShell to sign the zone, and document the DNSSEC parameter values that you will use.
4. **Sign the zone**: Using the methods and settings you have chosen, initiate zone signing on the Key Master and monitor progress.
5. **Verify zone signing**: Verify that zone signing was successful on the Key Master. If the zone is Active Directory-integrated, verify that replication of the signed zone is successful. If secondary servers are configured, verify that zone transfers are successful. Also verify that the zone is functioning as expected. Signing a zone with DNSSEC does not affect your ability to add, delete, or update resource records.
6. **Choose DNS clients**: Select which DNS clients will require that DNSSEC validation be performed for DNS queries in this zone. The choice of DNS clients also affects which DNS servers will require trust anchors. DNS clients that require validation must use DNS servers that have valid trust anchors for the zone.
7. **Deploy trust anchors**: Deploy trust anchors to all non-authoritative DNS servers that will perform validation of DNS queries for the zone. If applicable, this includes deploying DS records to DNS registrars that manage parent DNS zones.
8. **Configure and deploy name resolution policy**: If desired, deploy name resolution policy in your environment. Configure the NRPT to require validation of DNS queries for the signed zone and deploy this policy only to DNS clients that use DNS servers that are capable of DNSSEC validation. If some authoritative DNS servers in your environment are not capable of hosting a signed version of the zone, ensure that these DNS servers are not queried when DNSSEC validation is required.
9. **Verify name resolution policy**: Verify that NRPT settings are applied to DNS clients and that all DNS queries are successful.
10. **Monitor and manage**: Review the impact that zone signing has on your DNS server and support infrastructure. Initiate the key rollover process or wait for automated key rollover to occur and verify that the zone is successfully re-signed. Redistribute trust anchors as needed following key rollover.
11. **Scale out**: When you are ready to deploy DNSSEC on a larger scale, add additional zones, servers, and clients. Repeat the previous steps in this section to verify that new zones are signed correctly and that your DNS infrastructure is operating as expected.

See also

DNSSEC Deployment Planning

Why DNSSEC

DNSSEC Performance Considerations

DNSSEC Requirements

DNSSEC Performance Considerations

Prior to signing a DNS zone, consider the impact that DNSSEC can have on DNS server performance and resource utilization. This section provides information to help you understand how server performance is affected by DNSSEC.

In this section

DNS performance

Cryptographic algorithms

DNS performance

Performance considerations for DNSSEC-signed zones include the following:

- Signed zones can require up to five times more memory than an equivalent unsigned zone. Ensure that sufficient memory is available on all the DNS servers that will host signed zones.
- Zone loading time is increased for DNSSEC-signed zones. A DNS server with several large DNSSEC-signed zones will require additional time to restart the DNS Server service, or to reload individual signed zones.
- When responding to queries, the DNS server will respond with additional DNSSEC resource records. This will increase the number of packets on the wire and can decrease the maximum query throughput of the DNS server by a factor of two to three times compared to an unsigned zone.
- A DNS server that is performing validation of DNSSEC data can experience an increase in CPU usage. A server that is authoritative for signed zones will typically experience only a small increase in CPU usage associated with its zone handling. However, the increase in CPU load can be more significant if the DNS server is also performing validation on data it obtains from other remote signed zones, when servicing queries from clients or when configured to act as a forwarder for other DNS servers, depending on how many queries it needs to service.
- A signed zone can have four or more times the number of records as the unsigned version of the zone depending on the number of signing keys and algorithms used. Depending on the signature size, the actual disk space consumed by a file-backed signed zone can be more than ten times that of the unsigned zone. This is not a concern for AD-integrated zones, since the signatures are only stored in memory and are never committed to disk.

The actual performance of a DNSSEC-signed zone depends on the following factors:

- **Zone type**: The DNS server must generate signatures when loading Active Directory-integrated zones, but will have signatures stored on disk for a file-backed zone. This will significantly affect the server performance with respect to zone loading.
- **Cryptographic algorithms used**: RSA based algorithms are generally faster from a signature generation and validation standpoint as compared to Elliptical Curve based algorithms, for example.

- **Key length**: Longer key lengths will create larger signatures and increased signature computation and verification overhead.
- **Denial of Existence (NSEC vs. NSEC3)**: NSEC3 requires the calculation of hashes in the authenticating and non-authenticating functions. The number of hash iterations can also affect performance by adding overhead to zone signing time and validation.

Cryptographic algorithms

In general, cryptographic operations are computationally expensive. For large zones, the DNS server may take several minutes to sign the zone depending on the key length and size of the zone. To prevent performance degradation from occurring when all DNS servers start to sign the zone at the same time, signing is staggered. When a replica DC sees the DNSSEC keys and configuration, it will wait for a random period between 5 minutes and 30 minutes before it begins signing the zone. Different cryptographic algorithms also have a different impact on DNS server resources. See the following table.

Table 1: Available cryptographic algorithms on Windows Server 2012 R2:

Algorithm	Key Length (bits)
ECDSAP256/SHA-256	256
ECDSAP384/SHA-384	384
RSA/SHA-1	1024 - 4096
RSA/SHA-1 (NSEC3)	1024 - 4096
RSA/SHA-256	1024 - 4096
RSA/SHA-512	1024 - 4096

The following figure displays results from lab testing of Active Directory-integrated zone signing performance and memory usage for different choices of ZSK cryptographic algorithms, key lengths, and NSEC choices. The results are intended to show a general trend for different choices of signing key parameters. Precise results will vary based on the DNS environment.

Figure 1: Zone signing performance and memory usage for different cryptographic algorithms and key lengths

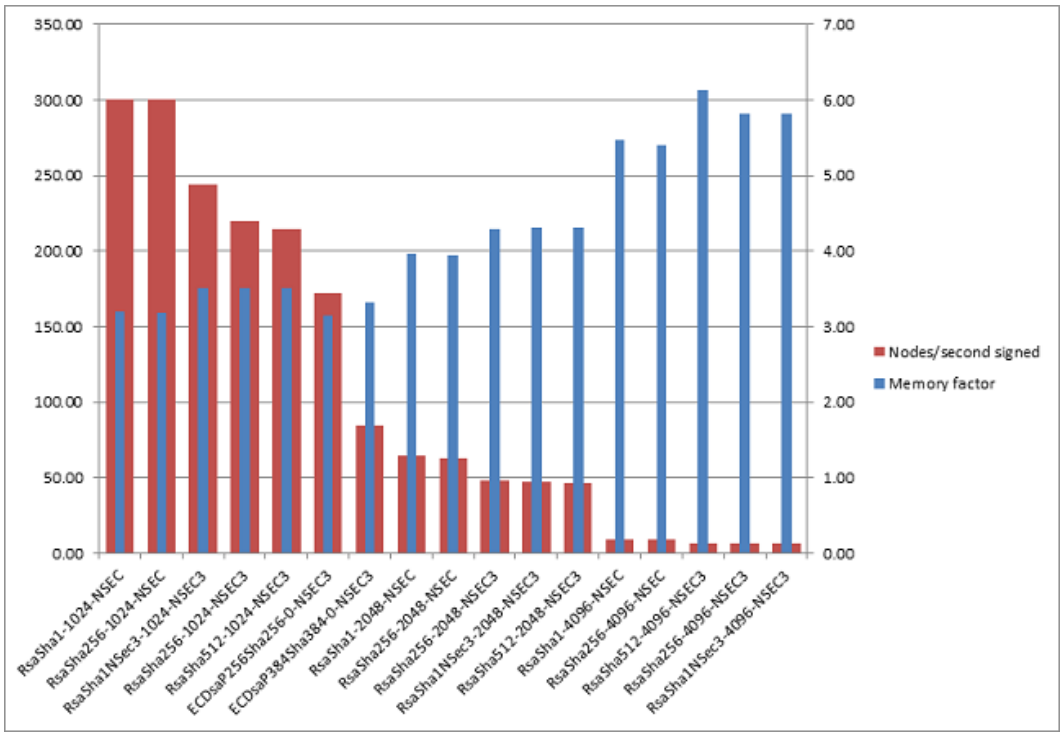

- Red indicates the number of nodes signed per second. In Active Directory, resource records are represented as attribute values on a dnsNode object.
- Blue indicates additional memory required by the DNS service process when the zone is loaded. A value of three means the zone takes up three times as much memory.

As shown in the previous figure, longer key lengths or the use of ECDSAP cryptographic algorithms are associated with significantly greater computational expense and memory consumption.

See also

Why DNSSEC

Stage a DNSSEC Deployment

DNSSEC Requirements

DNSSEC Requirements

Requirements to deploy DNSSEC are the following:

Category	Requirement	Details
Operating system	Windows Server® 2012, or a later operating system is required.	DNS servers must be running Windows Server® 2012 or a later operating system to sign a zone with DNSSEC or validate DNSSEC signatures. To allow for staged migration of server operating systems, DNSSEC deployment does not require that all DNS servers are running Windows Server 2012 or later. A mixed-mode environment is supported, with limitations. A server running Windows Server® 2008 R2 can host a static, offline-signed zone, however DNSSEC settings on a zone signed with Windows Server 2008 R2 cannot be modified by DNS servers running Windows Server 2012 or a later operating system. DNSSEC-related resource records in this type of zone will be loaded and can be validated on DNS servers running a later operating system, but for most operations the zone is treated as if it is unsigned. A zone that is signed on a DNS server running Windows Server 2012 or later will be loaded as an unsigned zone on a secondary or Active Directory-integrated primary DNS server running Windows Server 2008 R2.
Server roles	The DNS Server role is required.	Computers hosting DNSSEC signed zones must have the DNS Server role installed and running. To host an Active Directory-integrated DNS zone, DNS servers must also be running the Active

Category	Requirement	Details
		Directory Domain Services (AD DS) role. Additional roles and role services are optional.
Features	Features are required to perform administrative tasks.	No features are required to host a DNSSEC-signed DNS zone. However, some features are required to perform administrative tasks. For example, Group Policy Management is required to configure the NRPT and DNS Server Tools are required to sign or unsign a zone using DNS Manager, Windows PowerShell, or dnscmd.exe. If a DNS server does not have any features installed, it can still host a zone that was signed on a different DNS server.
Network	The network requirements for DNSSEC are the same as those for the DNS protocol. Note: See EDNS0 requirements.	DNS servers can be on the same network or different networks. DNS servers are recommended to have static IP addresses configured on all network adapters. Network infrastructure must support DNSSEC. See EDNS0 requirements below.
EDNS0	DNS servers must support EDNS0.	DNSSEC requires that EDNS0 is enabled on all DNS servers that will host or validate DNSSEC-signed zones. EDNS0 enables large (greater than 512 byte) UDP packet support in DNS which is required to send DNSSEC-enabled DNS responses. There are known issues with network infrastructure devices such as routers and firewalls dropping,

Category	Requirement	Details
		fragmenting or changing the arrival order of >512 byte UDP packets. This can cause DNS queries to fail. Ensure that your network infrastructure is capable of passing large UDP formatted network packets.
DNS servers	At least one primary, authoritative DNS server is required.	One or more primary, authoritative DNS servers is required to sign or unsign a zone with DNSSEC. At least one primary authoritative DNS server is required to be the Key Master. Additional DNS servers are optional, and can be primary, secondary, or resolving DNS servers.
DNS zones	At least one primary DNS zone is required.	Both forward and reverse lookup zones can be signed with DNSSEC. Zones can be Active Directory-integrated or file-backed. When a zone is signed on a DNS server running Windows Server 2012 or a later operating system, the zone will automatically be signed on all other primary DNS servers running Windows Server 2012 or later. Zone signing on other primary DNS servers is delayed and staggered. A secondary DNS server depends on the primary DNS server to sign the zone, and will transfer a signed version of the zone from the primary that supports DNSSEC validation.
Domain membership	Not required.	DNS servers can be workgroup computers or domain member computers. However, workgroup computers cannot host Active

Category	Requirement	Details
		Directory-integrated DNS zones.

See also

Why DNSSEC

Stage a DNSSEC Deployment

DNSSEC Performance Considerations

Deploy DNSSEC with Windows Server 2012

Use the following concepts and procedures to deploy DNSSEC with Windows Server® 2012 or Windows Server® 2012 R2.

Deploying DNSSEC

To deploy DNSSEC, review DNSSEC conceptual information below, and then use the DNSSEC deployment checklists provided in this guide.

DNSSEC concepts

Overview of DNSSEC: Information about how DNSSEC works is provided.

DNS Servers: DNSSEC support in Windows Server is discussed.

DNS Clients: The behavior of security-aware and non-security-aware DNS clients is discussed.

DNS Zones: Information is provided about zone signing and unsigning with Windows PowerShell or DNS Manager.

Trust Anchors: A trust anchor is a public cryptographic key that must be installed on DNS servers that will validate DNSSEC data.

The NRPT: Introduction and details about the Name Resolution Policy Table (NRPT) is provided.

Why DNSSEC: Risks and benefits of DNSSEC are discussed.

Stage a DNSSEC Deployment: Steps and considerations are provided to help introduce DNSSEC into your environment.

DNSSEC Performance Considerations: The impact of zone signing on a DNS infrastructure is discussed.

DNSSEC Requirements: Requirements for deploying DNSSEC are provided.

DNSSEC deployment checklists

Checklist	Description
Checklist: Deploy DNSSEC	See this parent checklist to get started deploying DNSSEC.
Checklist: Sign a Zone	Sign a DNS zone and verify DNSSEC signing.
Checklist: Distribute Trust Anchors	Export from authoritative DNS servers and import or add trust anchors to validating DNS servers.
Checklist: Deploy DNSSEC Policies to DNS Clients	Configure and verify name resolution policy.
Checklist: Review and Manage a Signed Zone	Administer your signed zone.
Checklist: Revert to an Unsigned Zone	Unsign a zone.
Checklist: Manage Signing Keys	Review and replace zone signing keys.
Checklist: Move the Key Master Role	Change the DNS server that will be the Key Master.
Checklist: Reconfigure Zone Signing Parameters on a Signed Zone	Change zone signing parameters.
Checklist: Perform an Emergency Key Revocation	Unsign a zone and replace signing keys.
Checklist: Perform a Manual Key Rollover	Roll over signing keys manually and update trust anchors.

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Checklist: Deploy DNSSEC

This parent checklist includes cross-reference links to topics that provide important conceptual information about DNSSEC. It also contains links to subordinate checklists that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about DNSSEC, and deciding on a deployment method and DNSSEC parameter values to use.

The DNS server that you use to perform procedures in this checklist is intended to be the Key Master. For more information, see The Key Master. You can also choose to use a different DNS server as the Key Master but you might need to adjust some steps in these procedures accordingly.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy DNSSEC

	Task	Reference
☐	Review DNSSEC concepts, terminology, components, requirements and specifications.	Overview of DNSSEC DNSSEC in Windows Appendix A: DNSSEC Terminology
☐	Decide on a deployment method; identify pilot servers and zones.	DNSSEC Deployment Planning
☐	Sign a zone	Checklist: Sign a Zone
☐	Deploy trust anchors.	Checklist: Distribute Trust Anchors
☐	Configure and deploy DNS client policies.	Checklist: Deploy DNSSEC Policies to DNS Clients
☐	Review zone signing parameters and manage the signed zone.	Checklist: Review and Manage a Signed Zone
☐	Deploy IPsec policy to protect zone transfers	Checklist: Secure Zone Transfers
☐	Review DNS client requirements for DNSSEC validation.	Procedure: Review Name Resolution Policy Settings

See also

Overview of DNSSEC

DNSSEC in Windows
DNSSEC Deployment Planning
Appendix A: DNSSEC Terminology
Appendix B: Windows PowerShell for DNS Server

Procedure: Review Name Resolution Policy Settings

Use the following procedures to verify name resolution policy settings on a DNS server and a DNS client using Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For information about configuring name resolution policy, see The NRPT.

Name resolution policy settings can only be applied to computers that are security-aware. For more information about security-aware computers, see Security-aware client.

See the following procedures:

- Review NRPT configuration: Check settings in the Name Resolution Policy Table (NRPT).
- Review effective name resolution policy: Verify whether or not DNSSEC validation is required for a given namespace.

Review NRPT configuration

Name resolution policy is configured in the NRPT, either in domain Group Policy, local Group Policy, or both. The NRPT can be configured to require or not require that DNSSEC validation be performed for DNS queries within a given namespace.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review NRPT configuration

1. Open an elevated Windows PowerShell prompt on a computer with the Group Policy Management feature installed.
2. To view a list of Group Policy Objects (GPOs) in the domain, use the **Get-GPO** cmdlet. See the following example:

```
PS C:\> Get-GPO -All -DomainName contoso.com | fl -Property
DisplayName,GpoStatus


DisplayName : Default Domain Policy
GpoStatus   : AllSettingsEnabled
```

```
DisplayName : Default Domain Controllers Policy
GpoStatus   : AllSettingsEnabled

DisplayName : NRPT_settings
GpoStatus   : AllSettingsEnabled
```

In the previous example, the Format-List (fl) parameter is used to display only the **DisplayName** and **GpoStatus** properties. The **NRPT_settings** GPO is a custom GPO created in this example to contain all name resolution policies for the domain. This GPO name is only an example. GPO names for your domain will be different. The example GPO is used in the following example.

3. To review name resolution policy settings in a GPO, use the **Get-DnsClientNrptRule** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptRule -GpoName
contoso.com\NRPT_settings

Name                                : Rule1
Version                             : 1
Namespace                           : {.secure.contoso.com}
IPsecCARestriction                  :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessProxyName               :
DirectAccessQueryIPsecEncryption    :
DirectAccessQueryIPsecRequired      :
NameServers                         :
DnsSecEnabled                       : True
DnsSecQueryIPsecEncryption          :
DnsSecQueryIPsecRequired            : False
DnsSecValidationRequired            : True
NameEncoding                        : Disable
DisplayName                         :
Comment                             :
```

In the previous example, the *secure.contoso.com* namespace displays **True** next to **DnsSecValidationRequired**, which means that all DNS clients that are security-aware,

and that have this GPO applied, will require that queries for the *secure.contoso.com* zone are validated.

Review effective name resolution policy

Although a GPO can be configured with NRPT settings, these settings will not affect a DNS client unless they are correctly applied. To verify that settings are applied, review the effective name resolution policy.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review effective name resolution policy

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To review effective name resolution policy, use the **Get-DnsClientNrptPolicy** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptPolicy -Effective


Namespace                       : .secure.contoso.com
QueryPolicy                     : QueryIPv6Only
SecureNameQueryFallback         : FallbackPrivate
DirectAccessIPsecCARestriction  :
DirectAccessProxyName           :
DirectAccessDnsServers          :
DirectAccessEnabled             : False
DirectAccessProxyType           :
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired  :
NameServers                     :
DnsSecIPsecCARestriction        :
DnsSecQueryIPsecEncryption      :
DnsSecQueryIPsecRequired        : False
DnsSecValidationRequired        : True
NameEncoding                    :
```

In the previous example, **True** is displayed next to **DnsSecValidationRequired** for the *secure.contoso.com* namespace, which means that DNSSEC validation will be required for DNS queries in the *secure.contoso.com* zone. Output from the **Get-**

DnsClientNrptPolicy cmdlet also includes any name resolution policy that might be configured in local Group Policy.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Sign a Zone

Checklist: Deploy DNSSEC > **Checklist: Sign a Zone**

This checklist includes procedures to help you sign a zone with DNSSEC.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about DNSSEC and deciding the deployment method and settings to use.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Sign a Zone

	Task	Reference
☐	Sign the zone using one of three available options to specify zone signing parameters: custom, existing, or default.	Procedure: Customize Zone Signing Parameters Procedure: Sign the Zone with Parameters of an Existing Zone Procedure: Use Default Settings to Sign the Zone
☐	Verify that the zone was successfully signed using the specified parameter values.	Procedure: Verify Zone Signing

See also

Overview of DNSSEC

DNSSEC in Windows

Procedure: Customize Zone Signing Parameters

Use the following procedures to customize zone signing parameters in preparation to sign a zone with DNSSEC. Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Also see DNS Zones for additional information about zone signing and DNSSEC parameter values.

Choose DNS Manager or Windows PowerShell to customize zone signing parameters:

- Customize zone signing parameters in DNS Manager
- Customize zone signing parameters in Windows PowerShell

Customize zone signing parameters in DNS Manager

The **Zone Signing Wizard** is used to sign a zone in DNS Manager. For examples of the wizard, see Zone Signing Wizard.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To customize zone signing parameters in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click an unsigned zone, point to **DNSSEC**, and then click **Sign the Zone**.
3. Click **Next**, on the **Signing Options** page choose **Customize zone signing parameters**, and then click **Next**.
4. On the **Key Master** page, choose **The DNS server <local server name> is the Key Master** if you wish to make the local server the Key Master. If you wish to make a different server the Key Master, choose **Select another primary server as the Key Master**, and then select a DNS server from the list of available DNS servers.
5. Click **Yes** on the alert that is displayed to wait for the local server to build a list of DNS servers that are capable of being the Key Master for this zone. Alternatively, you can type the fully qualified domain name of a DNS server and click **Next**. If you enter the name of a Key Master manually, you must click **Yes** when prompted to validate the Key Master.

Important

To choose a different DNS server as the Key Master, the local server must be able to contact other primary, authoritative DNS servers that are capable of DNSSEC online signing. The currently signed in user must also have administrative privileges on these DNS servers. If no other DNS servers can be selected as the Key Master, you must choose the local server as the Key Master. The Key Master can be changed after zone signing, if desired.

6. Click **Next** to view the first **Key Signing Key (KSK)** page. Review information on this page and then click **Next**.
7. On the second **Key Signing Key (KSK)** page, click **Add** to begin configuring KSK parameters.

   Note

   If the zone was previously signed and then unsigned, any KSKs that were used previously are listed and will be used again if they are not removed. You can **Add** new KSKs, **Edit** the existing KSKs, or **Remove** these KSKs.
8. Enter or select values for all KSK parameters available in the **New Key Signing Key (KSK)** dialog box. Default parameter values are provided. The **Guid** parameter is generated automatically. Click **OK** when finished.

   Note

   At least one KSK is required. You can add additional KSKs if desired. If you select **Enable automatic rollover**, then a standby KSK is automatically generated to support the double signature rollover method for KSKs. Do not add additional KSKs to support manual rollover. All signing keys that are added in this dialog box will be used during zone signing.

   Important

   To support automatic key rollover, an additional signing key is generated for each ZSK and KSK that is added to the zone. For automatic KSK rollover, only the Double Signature rollover method is supported. For automatic ZSK rollover, only the Pre-publish rollover method is supported. For more information, see **Manage Signing Keys**.
9. Click **Next** to view the first **Zone Signing Key (ZSK)** page. Review information on this page and then click **Next**.
10. On the second **Zone Signing Key (ZSK)** page, click **Add** to begin configuring KSK parameters.

    Note

    If the zone was previously signed and then unsigned, any ZSKs that were used previously are listed and will be used again if they are not removed. You can **Add** new ZSKs, **Edit** the existing ZSKs, or **Remove** these ZSKs.
11. Enter or select values for all ZSK parameters available in the **New Zone Signing Key (ZSK)** dialog box. Default parameter values are provided. The **Guid** parameter is

generated automatically. Click **OK** when finished.

Note

At least one ZSK is required. You can add additional ZSKs if desired. If you select **Enable automatic rollover**, then the next ZSK is automatically generated to support the pre-publish rollover method for ZSKs. Do not add additional ZSKs to support manual rollover. All signing keys that are added in this dialog box will be used during zone signing.

12. Click **Next** to view the **Next Secure (NSEC)** page. You can choose NSEC or NSEC3 on this page, and configure NSEC3 parameters.
13. Click **Next** to view the **Trust Anchors (TAs)** page. You can enable distribution of trust anchors and automatic update of trust anchors on this page.

    Tip

    Trust anchors are not required on DNS servers that are authoritative for the zone. You should only enable distribution of trust anchors in the Active Directory forest if other DNS servers running on domain controllers will provide non-authoritative responses for the zone that require DNSSEC validation.
14. Click **Next** to view the **Signing and Polling Parameters** page. You can configure values for DNSSEC signing and polling on this page.
15. Click **Next** to view the summary page. You can review current choices for all DNSSEC parameter values on this page. Click **Back** to configure different parameter values, click **Cancel** to close the zone signing wizard without signing the zone, or click **Next** to begin signing the zone.
16. Verify that **The zone has been successfully signed** is displayed, and then click **Finish**.

Customize zone signing parameters in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Important

To use DNSSEC-related Windows PowerShell cmdlets on a DNS zone, the zone must first be configured with a Key Master. This is automatically done the first time a zone is signed, and the Key Master setting is maintained even if a zone is unsigned later. For this reason, you can use cmdlets like **Set-DnsServerDnsSecZoneSetting** on unsigned zones provided they were signed at some time in the past. To use this cmdlet on a zone that has never been signed, you must forcibly assign a Key Master server by seizing the role with the **Reset-DnsServerZoneKeyMasterRole** cmdlet and the **SeizeRole** parameter.

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the zone you wish to sign with DNSSEC. Also replace the example computer names with names of the DNS servers you will use.

To customize zone signing parameters in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To add a Key Master for the zone, use the **Reset-DnsServerZoneKeyMasterRole** cmdlet. See the following example.

   ```
   PS C:\> Reset-DnsServerZoneKeyMasterRole -ZoneName
   secure.contoso.com -KeyMasterServer dc2.contoso.com -
   SeizeRole -Force
   ```

   In this example, the server dc2.contoso.com is assigned the Key Master role for the secure.contoso.com zone. This step is only required if the zone has never been signed and has not already been assigned a Key Master.
3. To configure custom zone signing parameters, use the **Set-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

   ```
   Set-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com -
   SecureDelegationPollingPeriod 8:00:00
   ```

   In this example, the **SecureDelegationPollingPeriod** parameter is set to a custom value. Parameters that are not specified will use default values.
4. To add a custom KSK, use the **Add-DnsServerSigningKey** cmdlet. See the following example:

   ```
   Add-DnsServerSigningKey -ZoneName secure.contoso.com -Type
   KeySigningKey -CryptoAlgorithm RsaSha1NSec3 –KeyLength 3072
   ```

   In this example, the **CryptoAlgorithm** and **KeyLength** parameters are set to a custom value. Parameters that are not specified will use default values.
5. To add a custom ZSK, use the **Add-DnsServerSigningKey** cmdlet. See the following example:

   ```
   Add-DnsServerSigningKey -ZoneName secure.contoso.com -Type
   ZoneSigningKey -CryptoAlgorithm RsaSha1NSec3
   ```

   In this example, the **CryptoAlgorithm** parameter is set to a custom value. Parameters that are not specified will use default values.
6. To sign the zone with the parameters that have been provided, use the **Invoke-DnsServerZoneSign** cmdlet. See the following example:

   ```
   Invoke-DnsServerZoneSign -ZoneName secure.contoso.com -Force
   ```

   In this example, the **Force** parameter is used to skip confirmation and immediately sign the zone.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Sign the Zone with Parameters of an Existing Zone

Use the following procedures to sign a DNS zone using an existing set of DNSSEC parameter values. Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Also see DNS Zones for additional information about zone signing and DNSSEC parameter values.

Choose DNS Manager or Windows PowerShell to sign a zone with existing parameter values:

- Sign a zone with existing zone signing parameters in DNS Manager
- Sign a zone with existing zone signing parameters in Windows PowerShell

Sign a zone with existing zone signing parameters in DNS Manager

The **Zone Signing Wizard** is used to sign a zone in DNS Manager. For examples of the wizard, see Zone Signing Wizard.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To sign a zone with existing zone signing parameters in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click an unsigned zone, point to **DNSSEC**, and then click **Sign the Zone**.
3. Click **Next**, on the **Signing Options** page choose **Customize zone signing parameters**, and then click **Next**.
4. On the **Key Master** page, choose **The DNS server <local server name> is the Key Master** if you wish to make the local server the Key Master. If you wish to make a different server the Key Master, choose **Select another primary server as the Key Master**, and then select a DNS server from the list of available DNS servers.
5. Click **Yes** on the alert that is displayed to wait for the local server to build a list of DNS servers that are capable of being the Key Master for this zone. Alternatively, you can type the fully qualified domain name of a DNS server and click **Next**. If you enter the name of

a Key Master manually, you must click **Yes** when prompted to validate the Key Master.

Important
To choose a different DNS server as the Key Master, the local server must be able to contact other primary, authoritative DNS servers that are capable of DNSSEC online signing. The currently signed in user must also have administrative privileges on these DNS servers. If no other DNS servers can be selected as the Key Master, you must choose the local server as the Key Master. The Key Master can be changed after zone signing, if desired.

6. Click **Next** to view the first **Key Signing Key (KSK)** page. Review information on this page and then click **Next**.
7. On the second **Key Signing Key (KSK)** page, click **Add** to begin configuring KSK parameters.

Note
If the zone was previously signed and then unsigned, any KSKs that were used previously are listed and will be used again if they are not removed. You can **Add** new KSKs, **Edit** the existing KSKs, or **Remove** these KSKs.

8. Enter or select values for all KSK parameters available in the **New Key Signing Key (KSK)** dialog box. Default parameter values are provided. The **Guid** parameter is generated automatically. Click **OK** when finished.

Note
At least one KSK is required. You can add additional KSKs if desired. If you select **Enable automatic rollover**, then a standby KSK is automatically generated to support the double signature rollover method for KSKs. Do not add additional KSKs to support manual rollover. All signing keys that are added in this dialog box will be used during zone signing.

Important
To support automatic key rollover, an additional signing key is generated for each ZSK and KSK that is added to the zone. For automatic KSK rollover, only the Double Signature rollover method is supported. For automatic ZSK rollover, only the Pre-publish rollover method is supported. For more information, see **Manage Signing Keys**.

9. Click **Next** to view the first **Zone Signing Key (ZSK)** page. Review information on this page and then click **Next**.
10. On the second **Zone Signing Key (ZSK)** page, click **Add** to begin configuring KSK parameters.

Note
If the zone was previously signed and then unsigned, any ZSKs that were used previously are listed and will be used again if they are not removed. You can **Add** new ZSKs, **Edit** the existing ZSKs, or **Remove** these ZSKs.

11. Enter or select values for all ZSK parameters available in the **New Zone Signing Key (ZSK)** dialog box. Default parameter values are provided. The **Guid** parameter is generated automatically. Click **OK** when finished.

    Note
    At least one ZSK is required. You can add additional ZSKs if desired. If you select **Enable automatic rollover**, then the next ZSK is automatically generated to support the pre-publish rollover method for ZSKs. Do not add additional ZSKs to support manual rollover. All signing keys that are added in this dialog box will be used during zone signing.
12. Click **Next** to view the **Next Secure (NSEC)** page. You can choose NSEC or NSEC3 on this page, and configure NSEC3 parameters.
13. Click **Next** to view the **Trust Anchors (TAs)** page. You can enable distribution of trust anchors and automatic update of trust anchors on this page.

    Tip
    Trust anchors are not required on DNS servers that are authoritative for the zone. You should only enable distribution of trust anchors in the Active Directory forest if other DNS servers running on domain controllers will provide non-authoritative responses for the zone that require DNSSEC validation.
14. Click **Next** to view the **Signing and Polling Parameters** page. You can configure values for DNSSEC signing and polling on this page.
15. Click **Next** to view the summary page. You can review current choices for all DNSSEC parameter values on this page. Click **Back** to configure different parameter values, click **Cancel** to close the zone signing wizard without signing the zone, or click **Next** to begin signing the zone.
16. Verify that **The zone has been successfully signed** is displayed, and then click **Finish**.

Sign a zone with existing zone signing parameters in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the zone you wish to sign with DNSSEC. Also replace the example computer names with names of the DNS servers you will use.

To sign a zone with existing zone signing parameters in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.

    Note
    If the local DNS server is not the Key Master, you must specify the Key Master

using the **ComputerName** parameter.

2. To sign a zone with an existing parameter set, use the **Get-DnsServerDnsSecZoneSetting** and **Set-DnsServerDnsSecZoneSetting** cmdlets. If both the source and destination zones are file-backed, you can also use the **SigningMetadata** and **IncludeKSKMetadata** parameters to export signing key information and immediately sign the zone. This procedure will not work if one or both of the zones are Active Directory-integrated. See the following example for file-backed zones:

   ```
   Get-DnsServerDnsSecZoneSetting -SigningMetadata -ZoneName
   ext.contoso.com -IncludeKSKMetadata | Set-
   DnsServerDnsSecZoneSetting -ZoneName fin.contoso.com
   ```

   In this example, the DNSSEC settings for the signed zone *ext.contoso.com* are used to sign the *fin.contoso.com* zone.

   Important

   After running this command, the *fin.contoso.com* zone will be signed.

3. If one or both zones are Active Directory-integrated, you can use the **Get-DnsServerDnsSecZoneSetting** cmdlet to export settings. In this case, additional steps are required to sign the zone.

4. First, you must set the Key Master role on an unsigned zone. This is only required if the zone has never been signed. See the following example

   ```
   PS C:\> Reset-DnsServerZoneKeyMasterRole -ZoneName
   secure.contoso.com -KeyMasterServer dc2.contoso.com -
   SeizeRole -Force
   ```

   In this example, the server *dc2.contoso.com* was assigned the Key Master role for the zone *secure.contoso.com*. The **SeizeRole** parameter was used because the zone did not previously have an assigned Key Master.

5. Next, export and import DNSSEC parameter values. See the following example

   ```
   Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com |
   Set-DnsServerDnsSecZoneSetting -ZoneName ad.contoso.com
   ```

   In this example, DNSSEC settings (but not signing keys) are exported from the *secure.contoso.com* zone and imported to the *ad.contoso.com* zone. This step will not sign the zone.

6. To add signing key settings from an existing signed zone to an unsigned zone, use the **Get-DnsServerSigningKey** and **Add-DnsServerSigningKey** cmdlets. See the following example:

   ```
   Get-DnsServerSigningKey -ZoneName secure.contoso.com | Add-
   DnsServerSigningKey -ZoneName ad.contoso.com
   ```

   In this example, signing keys and all associated parameter values are exported from the signed zone: *secure.contoso.com*, and imported to the unsigned zone: *ad.contoso.com*.

7. Finally, sign the ad.contoso.com zone using the **Invoke-DnsServerZoneSign** cmdlet.

See the following example:

```
Invoke-DnsServerZoneSign -ZoneName ad.contoso.com -Force
```

In this example, the **Force** parameter is used to skip confirmation and immediately sign the zone.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Use Default Settings to Sign the Zone

Use the following procedures to sign a zone with default settings. Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Also see DNS Zones for additional information about zone signing and DNSSEC parameter values.

Choose DNS Manager or Windows PowerShell to sign the zone with default parameter values:

- Sign a zone using default parameters in DNS Manager
- Sign a zone using default parameters in Windows PowerShell

Sign a zone using default parameters in DNS Manager

The **Zone Signing Wizard** is used to sign a zone in DNS Manager. For examples of the wizard, see Zone Signing Wizard.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To sign a zone using default parameters in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. In the console tree, right-click an unsigned zone, point to **DNSSEC**, and then click **Sign the Zone**.
3. Click **Next**, on the **Signing Options** page choose **Use default settings to sign the zone**, and then click **Next**.
4. You can review current choices for all DNSSEC parameter values on this page. Click

Back to configure different parameter values, click **Cancel** to close the zone signing wizard without signing the zone, or click **Next** to begin signing the zone.

5. Verify that **The zone has been successfully signed** is displayed, and then click **Finish**.

Sign a zone using default parameters in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the zone you wish to sign with DNSSEC. Also replace the example computer names with names of the DNS servers you will use.

To sign a zone using default parameters in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To sign the zone with default parameters, use the **Invoke-DnsServerZoneSign** cmdlet with the **SignWithDefault** parameter. See the following example:

   ```
   Invoke-DnsServerZoneSign -ZoneName secure.contoso.com –
   SignWithDefault -Force
   ```

   In this example, the **Force** parameter is used to skip confirmation and immediately sign the zone. If you do not use the **Force** parameter, the default zone signing parameter values are displayed and you must confirm that you wish to sign the zone.

By default, the local DNS server will be the Key Master.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Verify Zone Signing

Use the following procedures to verify a zone has been signed with DNSSEC using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

The procedure does not verify that resource record signatures are valid. To verify the validity of zone signing, see Procedure: Verify Name Resolution Policy.

You can use DNS Manager or Windows PowerShell to verify zone signing:

- Verify zone signing in DNS Manager

- Verify zone signing in Windows PowerShell

Verify zone signing in DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To verify zone signing in DNS Manager

1. Open DNS Manager on an authoritative DNS server, or connect to an authoritative DNS server with DNS Manager.
2. In the console tree, under **Forward Lookup Zones** or under **Reverse Lookup Zones**, click the zone you wish to verify.
3. In the right-hand pane, review resource records in the zone and verify that DNSSEC-related resource records are present. If the zone was recently signed, you might need to right-click the zone and click **Refresh**.

   For a zone to be validly signed, it must contain RRSIG, DNSKEY, and NSEC or NSEC3 records. See the following examples of a zone before and after zone signing.

   The following simple, example zone has four resource records prior to zone signing: SOA, NS, A, and MX:

Name	Type	Data
(same as parent folder)	Start of Authority (SOA)	[1], dc2.contoso.com., hostmaster.contoso.com.
(same as parent folder)	Name Server (NS)	dc2.contoso.com.
mail	Host (A)	192.168.0.100
(same as parent folder)	Mail Exchanger (MX)	[10] mail.corp.contoso.com

   The following example shows the same zone after zone signing with default parameter values:

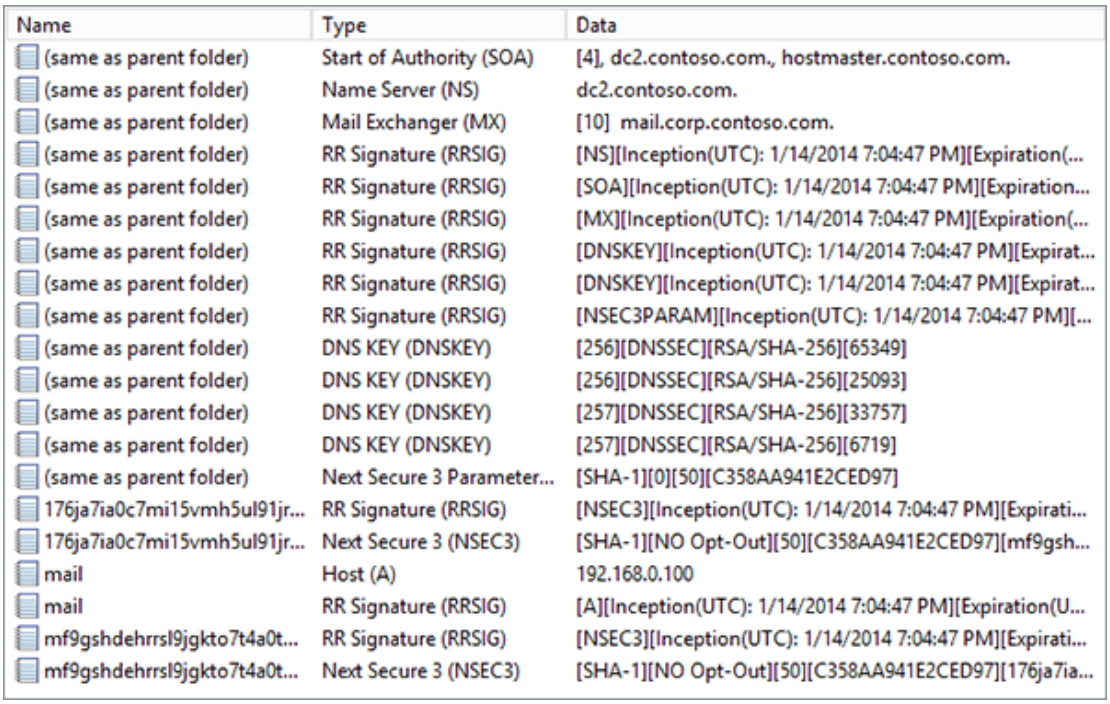

Name	Type	Data
(same as parent folder)	Start of Authority (SOA)	[4], dc2.contoso.com., hostmaster.contoso.com.
(same as parent folder)	Name Server (NS)	dc2.contoso.com.
(same as parent folder)	Mail Exchanger (MX)	[10] mail.corp.contoso.com.
(same as parent folder)	RR Signature (RRSIG)	[NS][Inception(UTC): 1/14/2014 7:04:47 PM][Expiration(...
(same as parent folder)	RR Signature (RRSIG)	[SOA][Inception(UTC): 1/14/2014 7:04:47 PM][Expiration...
(same as parent folder)	RR Signature (RRSIG)	[MX][Inception(UTC): 1/14/2014 7:04:47 PM][Expiration(...
(same as parent folder)	RR Signature (RRSIG)	[DNSKEY][Inception(UTC): 1/14/2014 7:04:47 PM][Expirat...
(same as parent folder)	RR Signature (RRSIG)	[DNSKEY][Inception(UTC): 1/14/2014 7:04:47 PM][Expirat...
(same as parent folder)	RR Signature (RRSIG)	[NSEC3PARAM][Inception(UTC): 1/14/2014 7:04:47 PM][...
(same as parent folder)	DNS KEY (DNSKEY)	[256][DNSSEC][RSA/SHA-256][65349]
(same as parent folder)	DNS KEY (DNSKEY)	[256][DNSSEC][RSA/SHA-256][25093]
(same as parent folder)	DNS KEY (DNSKEY)	[257][DNSSEC][RSA/SHA-256][33757]
(same as parent folder)	DNS KEY (DNSKEY)	[257][DNSSEC][RSA/SHA-256][6719]
(same as parent folder)	Next Secure 3 Parameter...	[SHA-1][0][50][C358AA941E2CED97]
176ja7ia0c7mi15vmh5ul91jr...	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 1/14/2014 7:04:47 PM][Expirati...
176ja7ia0c7mi15vmh5ul91jr...	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][C358AA941E2CED97][mf9gsh...
mail	Host (A)	192.168.0.100
mail	RR Signature (RRSIG)	[A][Inception(UTC): 1/14/2014 7:04:47 PM][Expiration(U...
mf9gshdehrrsl9jgkto7t4a0t...	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 1/14/2014 7:04:47 PM][Expirati...
mf9gshdehrrsl9jgkto7t4a0t...	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][C358AA941E2CED97][176ja7ia...

During zone signing, a matching resource record of **Type** RRSIG will be added to the zone for each of the resource records that were present prior to zone signing. The matching resource record type is displayed under **Data** (ex: [SOA][Inception…]), and is also displayed by viewing properties of the RRSIG resource record.

In addition, DNSKEY and NSEC3 resource records are added. Because NSEC3 was used in this example, an NSEC3PARAM resource record is also present. This record will not be present if the zone was instead signed with NSEC.

Notes:

- If you are viewing the zone on a primary authoritative DNS server, the icon for the zone in the console tree will display a lock symbol when the zone is signed. To see this icon, you might need to refresh the console view.
- If you are viewing the zone on a secondary authoritative DNS server, the icon for the zone in the console tree will not display a lock symbol even if the zone is signed.
- If the zone was signed on a server running Windows Server 2008 R2, DNSSEC-related resource records will be present, but the icon for the zone in the console tree will not display a lock symbol on either primary or secondary authoritative DNS servers. This is true even if you are viewing the zone on a server running Windows Server 2012 or a later operating system.
- If the DNSSEC-signed zone is Active Directory-integrated, and was recently signed, it will not be immediately signed on all other primary authoritative DNS servers. For performance reasons, zone signing on other DNS servers in the domain is staggered. To verify a zone is signed immediately after zone signing, view the zone on the Key Master.

Verify zone signing in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To verify zone signing in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on an authoritative DNS server.
2. To verify a zone is signed, use the **Get-DnsServerZone** cmdlet. See the following example:

```
PS C:\> Get-DnsServerZone -Name corp.contoso.com


ZoneName                              ZoneType
IsAutoCreated   IsDsIntegrated  IsReverseLookupZone  IsSigned
--------                              --------        ---------
----   --------------  -------------------  --------
corp.contoso.com                      Primary         False
False           False                 True
```

Verify that **True** is displayed in the **IsSigned** column.

Note: If the zone was signed on a server running Windows Server 2008 R2, a value of False will be displayed under IsSigned even though the zone is signed. This is because static, offline-signed zones are not automatically recognized as signed zones by computers running Windows Server 2012 or a later. DNSSEC validation is supported on these zones, however. To verify a zone that was signed on a DNS server running Windows Server 2008 R2, use the **Get-DnsServerResourceRecord** cmdlet. See the following example:

```
PS C:\> Get-DnsServerResourceRecord -ZoneName
dc1sec.contoso.com


HostName                  RecordType Timestamp
TimeToLive      RecordData
--------                  ---------- ---------            ---
-------      ----------
@                         NS         0
01:00:00        dc1.contoso.com.
@                         NS         0
01:00:00        dc2.contoso.com.
```

```
@                          SOA        0
01:00:00
[3][dc1.contoso.com.][hostmaster.contoso.com.]
@                          RRSIG      0
01:00:00        [DNSKEY][RsaSha1][35518]
@                          RRSIG      0
01:00:00        [DNSKEY][RsaSha1][43498]
@                          RRSIG      0
01:00:00        [NSEC][RsaSha1][35518]
@                          RRSIG      0
01:00:00        [NS][RsaSha1][35518]
@                          RRSIG      0
01:00:00        [SOA][RsaSha1][35518]
@                          NSEC       0
01:00:00        [dc1.dc1sec.contoso.com.][NS SOA RRSIG NSEC
DNS...
@                          DNSKEY     0
01:00:00        [ZoneKey][RsaSha1][35518]
@                          DNSKEY     0
01:00:00        [SEP][ZoneKey][RsaSha1][43498]
dc1                        A          0
01:00:00        192.168.0.1
dc1                        RRSIG      0
01:00:00        [NSEC][RsaSha1][35518]
dc1                        RRSIG      0
01:00:00        [A][RsaSha1][35518]
dc1                        NSEC       0
01:00:00        [www.dc1sec.contoso.com.][A RRSIG NSEC]
www                        A          0
01:00:00        192.168.0.100
www                        RRSIG      0
01:00:00        [NSEC][RsaSha1][35518]
www                        RRSIG      0
01:00:00        [A][RsaSha1][35518]
www                        NSEC       0
01:00:00        [dc1sec.contoso.com.][A RRSIG NSEC]
```

Verify that RRSIG, DNSKEY, and NSEC resource records are present in the zone.

Do not use Windows PowerShell cmdlets like **Get-DnsServerDnsSecZoneSetting** to verify that a zone is signed. A zone can have DNSSEC settings even if it is not currently signed.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Distribute Trust Anchors

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > **Checklist: Distribute Trust Anchors**

This checklist includes procedures to help you distribute trust anchors for a signed zone.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about DNSSEC and signing a zone with the settings you choose. You cannot distribute trust anchors until after a zone is signed with DNSSEC.

You must re-distribute trust anchors each time a zone is re-signed unless re-signing occurs as part of an automatic key rollover and trust anchors are distributed automatically on key rollover (RFC 5011). Trust anchors can also be distributed automatically in Active Directory to all Active Directory-integrated DNS servers within the replication scope for the zone.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Distribute Trust Anchors

	Task	Reference
☐	Review concepts for managing trust anchors.	Trust Anchors
☐	Enable automatic update of trust anchors on key rollover.	Procedure: Enable Automatic Update of Trust Anchors on Key Rollover **Important** Changes made to this setting do not take effect until the next automatic key rollover.

	Task	Reference
☐	Enable distribution of trust anchors in Active Directory.	Procedure: Distribute Trust Anchors in Active Directory
☐	Export and import trust anchors.	Procedure: Export a Trust Point Procedure: Import a Trust Point
☐	Manually add a trust anchor.	Procedure: Add a Trust Point
☐	Deploy a root trust anchor.	Procedure: Deploy a Root Trust Point

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Enable Automatic Update of Trust Anchors on Key Rollover

Use the following procedures to enable automatic update of trust anchors on key rollover using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

The following procedures enable automatic trust anchor updates for a zone that is already signed with DNSSEC.

If the zone is not signed and you want to enable automatic trust anchor updates when the zone is signed, see Procedure: Customize Zone Signing Parameters or Procedure: Use Default Settings to Sign the Zone. This setting is enabled by default.

Important
If it is enabled, automatic update of trust anchors only occurs during an automatic key rollover. Trust anchors are not updated if you perform a manual key rollover.

Choose to use DNS Manager or Windows PowerShell:

- Enable automatic update of trust anchors on key rollover using DNS Manager
- Enable automatic update of trust anchors on key rollover using Windows PowerShell

Enable automatic update of trust anchors on key rollover using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To enable automatic update of trust anchors on key rollover using DNS Manager

1. Open DNS Manager on a primary authoritative DNS server, or connect to a primary authoritative DNS server with DNS Manager.
2. In the console tree, right-click the zone where you will enable automatic update of trust anchors, point to **DNSSEC**, and then click **Properties**.
3. Click the **KSK** tab.
4. For each KSK that is listed under **Key signing keys (KSKs)**, click the KSK, click **Edit**, and in the **Key Rollover** section, select the **Enable automatic rollover** checkbox.
5. Edit the **Rollover frequency** if desired, and then click **OK**.
6. Click the **ZSK** tab.
7. For each ZSK that is listed under **Zone signing keys (ZSKs)**, click the ZSK, click **Edit**, and in the **Key Rollover** section, select the **Enable automatic rollover** checkbox.
8. Edit the **Rollover frequency** if desired, and then click **OK**.
9. Click **OK** to close the DNSSEC properties dialog.
10. In the popup dialog box that appears, click **Yes** to confirm the changes that were made to the zone.
11. Verify that **Configuration of DNSSEC parameters was successful** is displayed, and then click **OK**.

Enable automatic update of trust anchors on key rollover using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To enable automatic update of trust anchors on key rollover using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. You can use the **Get-DnsServerSigningKey** cmdlet to retrieve the **KeyId** values for signing keys in the zone, and then use the **Enable-DnsServerSigningKeyRollover** cmdlet to enable key rollover for all signing keys. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| Enable-DnsServerSigningKeyRollover -Force
```

In the previous example, the **Force** parameter is used to skip confirmation.

3. If you wish to enable rollover for some signing keys and not others, you must the specify individual KeyId of each signing key. See the following examples:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


KeyId                                   KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                   -------         ------
---------    ---------  -------------   -----------------
5f7d05be-1bb6-4a0c-b923-f9be300bf6ea   KeySigningKey
RsaSha1NSec3       2048       True            False
a2a72d54-f9d7-4b4b-a115-a4097092f5bf   ZoneSigningKey
RsaSha1NSec3       1024       True            False


PS C:\> Enable-DnsServerSigningKeyRollover -ZoneName
secure.contoso.com -KeyId 5f7d05be-1bb6-4a0c-b923-
f9be300bf6ea -Force


PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


KeyId                                   KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                   -------         ------
---------    ---------  -------------   -----------------
5f7d05be-1bb6-4a0c-b923-f9be300bf6ea   KeySigningKey
RsaSha1NSec3       2048       True            True
a2a72d54-f9d7-4b4b-a115-a4097092f5bf   ZoneSigningKey
RsaSha1NSec3       1024       True            False
```

See also

[DNSSEC in Windows]

[DNSSEC Deployment Planning]

Procedure: Distribute Trust Anchors in Active Directory

Use the following procedures to distribute trust anchors in Active Directory using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

The following procedures enable automatic distribution of trust anchors in Active Directory for a zone that is already signed with DNSSEC.

If the zone is not signed and you want to enable automatic trust anchor distribution when the zone is signed, see Procedure: Customize Zone Signing Parameters. This setting is not enabled by default.

Important
If this setting is enabled and the DNS server is also a domain controller, trust anchors will be distributed to all DNS servers running on domain controllers in the forest. If the DNS server is not a domain controller, a trust anchor for the zone will be added only to the local trust anchor store, located by default at **%windir%\system32\dns\TrustAnchors.dns**.

Choose to use DNS Manager or Windows PowerShell:

- Enable trust anchor distribution in Active Directory using DNS Manager
- Enable trust anchor distribution in Active Directory using Windows PowerShell

Enable trust anchor distribution in Active Directory using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To enable trust anchor distribution in Active Directory using DNS Manager

1. Open DNS Manager on a primary authoritative DNS server, or connect to a primary authoritative DNS server with DNS Manager.
2. In the console tree, right-click the zone where you will enable trust anchor distribution, point to **DNSSEC**, and then click **Properties**.
3. Click the **Trust Anchor** tab.
4. Select the **Enable the distribution of trust anchors for this zone** checkbox.
5. Click **OK** to close the DNSSEC properties dialog.
6. In the popup dialog box that appears, click **Yes** to confirm the changes that were made to the zone.
7. Verify that **Configuration of DNSSEC parameters was successful** is displayed, and

then click **OK**.

Enable trust anchor distribution in Active Directory using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To enable trust anchor distribution in Active Directory using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. You can use the **Set-DnsServerDnsSecZoneSetting** cmdlet with the **DistributeTrustAnchor** parameter to enable distribution of trust anchors in Active Directory. See the following example:

```
PS C:\> Set-DnsServerDnsSecZoneSetting -ZoneName
secure.contoso.com -DistributeTrustAnchor DnsKey
```

3. You can use the **Get-DnsServerDnsSecZoneSetting** cmdlet to confirm this setting. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName
secure.contoso.com | fl -Property
ZoneName,DistributeTrustAnchor


ZoneName              : secure.contoso.com
DistributeTrustAnchor : {DnsKey}
```

   In the previous example, the **Format-List** (fl) parameter was used to display only the **ZoneName** and **DistributeTrustAnchor** properties.

Only DNSKEY trust anchors can be distributed in Active Directory. You cannot distribute a DS trust anchor.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Export a Trust Point

Use the following procedures to export a trust point from a DNS server.

When you sign a zone on a DNS server running Windows Server 2008 R2, the DNSKEY RRSET (also called the Keyset) is saved on the local server. When you sign a zone on a DNS server running Windows Server 2012 or later, the DNSKEY RRSET and the DS RRSET (also called the dsset) are automatically saved in a text file on the Key Master. By default these files are available in the **%windir%\system32\dns** directory.

For more information about trust points, see Trust Anchors.

Procedures are provided using the File Explorer and Windows PowerShell. You cannot export a trust point using DNS Manager. When you have completed the procedures in this topic, return to the parent checklist.

Choose File Explorer or Windows PowerShell to perform this procedure:

- Export a trust point using File Explorer
- Export a trust point using Windows PowerShell

Export a trust point using File Explorer

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To export a trust point using File Explorer

1. Open File Explorer on the server where you wish to export the trust point, or connect to the DNS server from another location using File Explorer. The server must be the Key Master for the DNSSEC-signed zone.
2. Navigate to the **%windir%\system32\dns** directory on the server.

   If the trust point that you wish to export is a KEYSET (DNSKEY), then locate the *keyset-<zone.name.com>* file. For example: *keyset-secure.contoso.com*. If the trust point that is available is a DSSET (DS), then locate the *dsset-<zone.name.com>* file. For example: *dsset-secure.contoso.com*.

   Tip
   If File Explorer is currently configured (on the View menu) not to display file name extensions, then the filenames will appear to be their missing ".com" suffix.
3. Copy (do not move) the keyset- or dsset- file to a trusted location where it can be imported to other DNS servers. For example, you might wish to copy files to a network share, USB flash drive, or other storage device.

Export a trust point using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To export a trust point using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the server where you wish to import the trust point.
2. To export a trust point, use the **Export-DnsServerDnsSecPublicKey** cmdlet. See the following example.

   ```
   Export-DnsServerDnsSecPublicKey -ComputerName DC2.contoso.com
   -ZoneName secure.contoso.com -Path \\Myshare\keys
   ```

   In this example, a DNSKEY trust anchor for the zone *secure.contoso.com* is exported from the primary, authoritative DNS server DC2.contoso.com. The \\Myshare\keys location must be accessible from DC2.contoso.com. To export a DS trust anchor, add the **DigestType** parameter. See the following example:

   ```
   Export-DnsServerDnsSecPublicKey -ComputerName DC2.contoso.com
   -ZoneName secure.contoso.com -Path \\Myshare\keys -DigestType
   sha1
   ```

After exporting DNSKEY or DS trust anchors, text files will be saved in the location specified by the **Path** parameter. You can use these file to import a trust anchor on another DNS server. For more information about importing a trust anchor, see Procedure: Import a Trust Point.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Import a Trust Point

Use the following procedures to import a trust point for a zone that has been signed with DNSSEC. To import a trust point, the DSSET or KEYSET records for a zone must be available in a file. You can create this file manually, or export it from the Key Master.

When you sign a zone on a DNS server running Windows Server 2008 R2 or later, the DNSKEY RRSET (also called the Keyset) and the DS RRSET (also called the dsset) are stored in text files on the Key Master for use as trust points. By default these files are available in the

%windir%\system32\dns directory. To import a trust point on a DNS server, these files must be made accessible to the DNS server. For more information, see Procedure: Export a Trust Point.

For more information about trust points, see Trust Anchors.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Import a trust point using DNS Manager
- Import a trust point using Windows PowerShell

Import a trust point using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To import a trust point using DNS Manager

1. Open DNS Manager on the server where you wish to import the trust point, or connect to the DNS server from another location using DNS Manager.
2. In the console tree, right-click **Trust Points**, point to **Import**, and then click **DNSKEY** or click **DS**.

   If the trust point that you have made available is a KEYSET, then point to **DNSKEY**. If the trust point that is available is a DSSET, point to **DS**.
3. Click **Browse** or type the location of the KEYSET or DSSET file, and then click **OK**.

If a trust point for the zone did not previously exist, a new folder will be created in the DNS Manager console tree under Trust Points. If the trust point for this zone had previously been imported or added to the DNS server, the trust point will be updated.

Import a trust point using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To import a trust point using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the server where you wish to import the trust point.
2. To import a trust point, use the **Import-DnsServerTrustAnchor** cmdlet. See the following example.

```
PS C:\> Import-DnsServerTrustAnchor -KeySetFile
"\\File1\DNSKeys\keyset-secure.contoso.com"
```

In this example, a DNSKEY trust anchor for the zone *secure.contoso.com* is imported on the local DNS server. To import a DS trust anchor, use the **DSSetFile** parameter in place of **KeySetFile** and specify a filename of *dsset-secure.contoso.com*.

To verify that the trust point was successfully imported, you can use the **Get-DnsServerTrustAnchor** cmdlet. See the following example:

```
PS C:\> Get-DnsServerTrustAnchor -Name secure.contoso.com


TrustAnchorName                TrustAnchorType      TrustAnchorState
TrustAnchorData
---------------                ---------------      ----------------     -------------
--
secure.contoso.com.            DNSKEY               Valid
[8000][DnsSec][RsaSha256][AwEAAY/kENeHkSmtd...
secure.contoso.com.            DNSKEY               Valid
[62111][DnsSec][RsaSha256][AwEAAbMK6im/lrLB...
```

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Add a Trust Point

Use the following procedures to add a trust point for a zone that has been signed with DNSSEC. You can also import a trust point. To import a trust point, see Procedure: Import a Trust Point.

You can add a DNSKEY trust point, or a DS trust point. To add a DNSKEY trust point, you must be able to provide the zone name, cryptographic algorithm, and public key. To add a DS trust point, you must be able to provide the zone name, key tag, cryptographic algorithm, digest type, and digest. This information is provisioned in a trust point automatically when you import it from a file. For more information about trust points, see Trust Anchors.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Add a trust point in DNS Manager
- Add a trust point in Windows PowerShell

Add a trust point in DNS Manager

Trust points are stored on DNS servers as text files. By default,

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To add a trust point in DNS Manager

1. Open DNS Manager on the DNS server where you wish to add the trust anchor, or connect to this DNS server using DNS Manager.
2. In the console tree, right-click **Trust Points**, point to **Add** and then click **DNSKEY** or click **DS**, depending on the type of trust point you wish to add.
3. Under **Name**, type the name of the zone.
4. Under **Algorithm**, select the cryptographic algorithm used for the signing key.
5. If the new trust anchor is a DS resource record, type the key tag next to **Key Tag**.
6. By default, both the Zone Key and Secure Entry Point checkboxes are selected. If the signing key used is not a secure entry point (SEP), clear the Secure Entry Point checkbox.
7. Type the digest or public key under **Digest** or **Public Key** for a DNSKEY or DS trust anchor, respectively, and then click **OK**.

To view DNSKEY and DS resource record data on a DNS server, you can use the **Get-DnsServerResourceRecord** cmdlet. See the following examples:

```
PS C:\> Get-DnsServerResourceRecord -ZoneName secure.contoso.com -RRType DNSKEY -
ComputerName 192.168.0.2 | fl


DistinguishedName :
DC=@,DC=secure.contoso.com,cn=MicrosoftDNS,DC=DomainDnsZones,DC=contoso,DC=com
HostName          : @
RecordType        : DNSKEY
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
[SEP][ZoneKey][RsaSha256][1175][DnsSec][AwEAAat76D//jxQQHmPNhkfjufvoz74oj3qktZU7Byc+eD
7e3OcVhmaA/p+vu+R86ei8r2GDp7EaYPZ


PQMOUbNhrYPBp8GCliYq1tWy0N4tzL6mRKk+lflS6AG2guvvdCvk0VVQ14RXVggYY9y9/TDAVqOZVS4vpK2m9h
dHK9dK81fMXoEcflDWYjSA7RGsZoDwW2c
```

4Uozuopwi4AVW8g6ThfLwV7Q5rTp67IItqB3+7MgLfxmiHZZD33Xh3GAsbBhpvZF61ouLTMZJV/sOreyZT2ktv
Qm4YpDVzRoIfetChxnBzavcR19DbrJA/N
57D0trxuoa3N9LDCfAhGENvgNSstNE=]

DistinguishedName :
DC=@,DC=secure.contoso.com,cn=MicrosoftDNS,DC=DomainDnsZones,DC=contoso,DC=com
HostName : @
RecordType : DNSKEY
RecordClass : IN
TimeToLive : 01:00:00
Timestamp : 0
RecordData :
[SEP][ZoneKey][RsaSha256][21046][DnsSec][AwEAAbOkjGoP3OJ1OhMqdPJHf5RnDE+QO/oP/iMkJnr3K
I4SOoHgodgJb3a83mkDQp+jThxqxyTEPZ

LxhIezCw5ZLmhTYPmAj66yU3iLSvj/W79ycfv5eToPCORq8q4THo60TItHQya4uR1Odp3gu22GWDFnFpor31s/
DneroLLNzE447WqUq4SAvX8J00dcb2u82

0sHyzzVXgo8CdNVCDNbImn9HI8/AdNlxDRwlogw+lDx4z4360rCXjuVsNsXcejua3Vz47Urf7wLiuqnq42C4JB
GFkaypig21M7dUNCBHkr4ElkkaZvmzMtS
N35rKOLQ//9AT1T77DxDh+1UPVlfcwE=]

DistinguishedName :
DC=@,DC=secure.contoso.com,cn=MicrosoftDNS,DC=DomainDnsZones,DC=contoso,DC=com
HostName : @
RecordType : DNSKEY
RecordClass : IN
TimeToLive : 01:00:00
Timestamp : 0
RecordData :
[ZoneKey][RsaSha256][24879][DnsSec][AwEAAbOHlfH8Zs2rKF7Kv+YckMsSF0p5r35iWn44fulAPLynqA
J+vjgrAomKrNfmYKLU0hchl8QZMKsv8Hi

tUHDlg++esqvUvVdSatz9vXNgQsNDeGqgsl0SV1Uuqalh37j2d1r8xH7HK3TeppsPNU7MjC2bsqys0puoeMjaB
xII+FgX]

```
DistinguishedName :
DC=@,DC=secure.contoso.com,cn=MicrosoftDNS,DC=DomainDnsZones,DC=contoso,DC=com
HostName          : @
RecordType        : DNSKEY
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
[ZoneKey][RsaSha256][56980][DnsSec][AwEAAZ41CYznoxLtCU5wRGS3xK628ksj5BgisqVG2TUJPvxIEp
9rY8NBWFwtdqETk/FK/owkJY9qNbeyMiV

wUakGwsslXzh/hPsvb75lohKMsjSfgsvUV/E61EYlid68wBoy8Oxbf3sDFdIPmcDReQZ78ZnhCPAmRqa/seCsI
jripjLx]



PS C:\> Get-DnsServerResourceRecord -ZoneName adatum.com -RRType DS -ComputerName
192.168.0.2 | fl

DistinguishedName : corp
HostName          : corp
RecordType        : DS
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
[63513][RsaSha1NSec3][Sha256][E3F069CBEAC9C5168C4950FFE3E44B50E0BDB985CF2BC54D42C27CF3
B80CAEF8]

DistinguishedName : corp
HostName          : corp
RecordType        : DS
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
```

```
[63513][RsaSha1NSec3][Sha1][ED0260968FF6ED47942E1993827145D54991E4CF]


DistinguishedName : corp
HostName          : corp
RecordType        : DS
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
[58555][RsaSha1NSec3][Sha256][6EBECD6D28399291D0BE4C20082C3FD1C9CDEE1A1E34E707ABD22977
1AE12D7E]


DistinguishedName : corp
HostName          : corp
RecordType        : DS
RecordClass       : IN
TimeToLive        : 01:00:00
Timestamp         : 0
RecordData        :
[58555][RsaSha1NSec3][Sha1][27468EAEA69B9E04C094478FD2D398B7AF7CD48A]
```

In the previous examples, the DNSKEY records are displayed for the zone *secure.contoso.com* and the DS records are displayed for the zone *corp.adatum.com*.

Add a trust point in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To add a trust point in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To add a trust point, use the **Add-DnsServerTrustAnchor** cmdlet. See the following example.

```
PS C:\> Add-DnsServerTrustAnchor -Name corp.adatum.com -
```

```
KeyTag 63513 -CryptoAlgorithm RsaSha1NSec3 -DigestType Sha1 -
Digest ED0260968FF6ED47942E1993827145D54991E4CF
```

In this example, a DS trust point for the zone *corp.adatum.com* is added on the local DNS server.

3. Verify that the trust anchor was successfully added to the DNS server using the Get-DnsServerTrustAnchor cmdlet. See the following example:

```
PS C:\> Get-DnsServerTrustAnchor -Name corp.adatum.com


TrustAnchorName                TrustAnchorType
TrustAnchorState     TrustAnchorData
---------------                ---------------      ---------
-------     ---------------
corp.adatum.com.               DNSKEY
AddPending
[58555][DnsSec][RsaSha1NSec3][AwEAAe1ZjQR+f...
corp.adatum.com.               DNSKEY               Valid
[63513][DnsSec][RsaSha1NSec3][AwEAAbdN9OIL3...
```

Adding a DS trust point causes priming of the DNS server to acquire the current DNSKEY RRSET.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Deploy a Root Trust Point

Use the following procedures to deploy an Internet root (.) trust point.

Do not use this procedure if your network uses a local DNS server as a root DNS server. Instead, use the procedures for importing or adding a trust point and specify the zone name to be the root zone (.). For more information, see Procedure: Import a Trust Point or Procedure: Add a Trust Point.

The following procedures are used to deploy a root trust anchor from the Internet Assigned Numbers Authority (IANA) using the URL specified in **RootTrustAnchorsURL** property of the DNS server. You can view the **RootTrustAnchorsURL** property by using the **Get-DnsServerSetting** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSetting -All | Select RootTrustAnchorsURL

RootTrustAnchorsURL
```

```
------------------

https://data.iana.org/root-anchors/root-anchors.xml
```

To modify the **RootTrustAnchorsURL** property, use the **Set-DnsServerSetting cmdlet**.

Procedures are provided using Windows PowerShell.

Important

The **Root** parameter is not available for use with the **Add-DnsServerTrustAnchor** cmdlet in Windows Server 2012. This parameter is new in Windows Server 2012 R2. To deploy a root trust anchor on a DNS server running Windows Server 2012, you can use the **Dnscmd.exe** utility (ex: Dnscmd.exe /RetrieveRootTrustAnchors).

When you have completed the procedures in this topic, return to the parent checklist.

Deploy a root trust point using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

Important

To add a root trust point, the computer must be connected to the Internet.

Deploy a root trust point using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the resolving DNS server where you wish to add the root trust anchor.
2. To add a root trust anchor, use the **Add-DnsServerTrustAnchor** cmdlet with the **Root** parameter. See the following example.

```
PS C:\> Add-DnsServerTrustAnchor -Root
```

To verify that the root trust point was successfully added, you can use the **Get-DnsServerTrustPoint** cmdlet. If you do not specify a value for the **Name** parameter, then all current trust points are displayed. See the following example:

```
PS C:\> Get-DnsServerTrustPoint


TrustPointName                  TrustPointState      LastActiveRefreshTime
NextActiveRefreshTime
--------------                  ---------------      ---------------------
---------------------
.                               Active               1/23/2014 1:02:01 PM
```

```
1/23/2014 1:02:17 PM

adatum.com.                              Active              1/23/2014 1:02:01 PM
1/23/2014 1:02:17 PM
```

Warning

If the status of a DS trust point is DS Pending, this means that the DNS server currently does not have access to the zone's DNSKEY resource records. You might also need to verify that the DS trust point is current and uses a supported cryptographic algorithm.

After adding a valid DS trust point, the DNS server will automatically acquire the current DNSKEY RRSET and use it to update the trust anchor.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Deploy DNSSEC Policies to DNS Clients

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > **Checklist: Deploy DNSSEC Policies to DNS Clients**

This checklist includes procedures to help you deploy name resolution policy to DNS clients on your network using the Name Resolution Policy Table (NRPT).

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about DNSSEC, signing a zone, and deploying trust anchors. If you deploy name resolution policy to DNS clients before you have signed a zone and deployed trust anchors, you can break DNS resolution on your network.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy DNSSEC Policies to DNS Clients

	Task	Reference
☐	Review concepts about security-aware DNS clients and	DNS Clients The NRPT

	Task	Reference
	the NRPT.	
☐	Configure name resolution policy.	Procedure: Configure the NRPT
☐	Review name resolution policy settings.	Procedure: Verify Name Resolution Policy

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Configure the NRPT

Use the following procedures to configure the Name Resolution Policy Table (NRPT). The examples that are provided demonstrate how to add a rule in an existing Group Policy Object (GPO) that requires DNSSEC validation for a namespace (zone).

Procedures are provided using the Group Policy Management console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For more information about NRPT rule processing, see The NRPT.

Choose the Group Policy Management console Windows PowerShell to configure the NRPT:

- Configure the NRPT with the Group Policy Management console
- Configure the NRPT with Windows PowerShell

Configure the NRPT with the Group Policy Management console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure the NRPT with the Group Policy Management console

1. On a computer with the Group Policy Management feature installed, click **Tools** on the Server Manager menu and then click **Group Policy Management**.

2. In the Group Policy Management console tree, navigate to **Forest: contoso.com\Domains\contoso.com\Group Policy Objects**.
3. Right-click the Group Policy Object (GPO) that you wish to configure, and then click **Edit**.
4. In the Group Policy Management Editor tree, navigate to **Computer Configuration\Policies\Windows Settings\Name Resolution Policy**.
5. Under **Create Rules**, choose the part of the namespace to which the rule will apply from the drop-down list. For example, choose **Suffix** to create a rule that applies to all names that end in the text that you provide.

   Tip
   To view an explanation of each choice in the drop-down list, choose an item and then click in the text box to the right. A description of the choice that is made will be displayed to the left under **Description**.
6. Next to the drop-down list choice (ex: Suffix), type the namespace to which the rule will apply. For example, to apply a rule to all devices in the *secure.contoso.com* zone, choose **Suffix** and type **secure.contoso.com**.
7. On the **DNSSEC** tab, select the **Enable DNSSEC in this rule** checkbox, and then select the **Require DNS clients to check that name and address data has been validated by the DNS server checkbox**.
8. At the bottom of this section of the NRPT, click **Create**.
9. Under **Name Resolution Policy Table**, verify that the new rule is displayed.
10. Close the Group Policy Management Editor.
11. Close the Group Policy Management console.

Configure the NRPT with Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the zone you wish to sign with DNSSEC. Also replace the example computer names with names of the DNS servers you will use.

To configure the NRPT with Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a computer with the Group Policy Management feature installed.
2. To add an NRPT rule to an existing Group Policy Object (GPO), use the **Add-DnsClientNrptRule** cmdlet. See the following example.

```
PS C:\> Add-DnsClientNrptRule -GpoName NRPT_settings -
NameSpace secure.contoso.com -DnsSecEnable -
DnsSecValidationRequired
```

In this example, the GPO named **NRPT_settings** is configured with a rule requiring DNSSEC validation for the secure.contoso.com namespace.

3. To confirm that the GPO was successfully configured, use the **Get-DnsClientNrptRule** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptRule -GpoName NRPT_settings

Name                                : Rule1
Version                             : 1
Namespace                           : {.adatum.com}
IPsecCARestriction                  :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessProxyName               :
DirectAccessQueryIPsecEncryption    :
DirectAccessQueryIPsecRequired      :
NameServers                         :
DnsSecEnabled                       : True
DnsSecQueryIPsecEncryption          :
DnsSecQueryIPsecRequired            : False
DnsSecValidationRequired            : True
NameEncoding                        : Disable
DisplayName                         :
Comment                             :

Name                                : {77BCF3D4-8442-49F1-9E73-
4715CAD9888A}
Version                             : 1
Namespace                           : {secure.contoso.com}
IPsecCARestriction                  :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessProxyName               :
```

```
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired   :
NameServers                      :
DnsSecEnabled                    : True
DnsSecQueryIPsecEncryption       :
DnsSecQueryIPsecRequired         : False
DnsSecValidationRequired         : True
NameEncoding                     : Disable
DisplayName                      :
Comment                          :
```

In this example, a second rule was added to the **NRPT_settings** GPO. The rule for the *adatum.com* namespace was already present, and the rule for the *secure.contoso.com* namespace was added.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Verify Name Resolution Policy

Use the following procedures to verify name resolution policy on a DNS server and a DNS client using Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For information about configuring name resolution policy, see The NRPT.

Name resolution policy settings can only be applied to computers that are security-aware. For more information about security-aware computers, see Security-aware client.

See the following procedures:

- Review NRPT configuration: Check settings in the Name Resolution Policy Table (NRPT).
- Review effective name resolution policy: Verify whether or not DNSSEC validation is required for a given namespace.
- Verify DNSSEC validation: If DNSSEC validation has been enabled or disabled recently, verify that name resolution is working as expected.

Review NRPT configuration

Name resolution policy is configured in the NRPT, either in domain Group Policy, local Group Policy, or both. The NRPT can be configured to require or not require that DNSSEC validation be performed for DNS queries within a given namespace.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review NRPT configuration

1. Open an elevated Windows PowerShell prompt on a computer with the Group Policy Management feature installed.
2. To view a list of Group Policy Objects (GPOs) in the domain, use the **Get-GPO** cmdlet. See the following example:

```
PS C:\> Get-GPO -All -DomainName contoso.com | fl -Property
DisplayName,GpoStatus


DisplayName : Default Domain Policy
GpoStatus   : AllSettingsEnabled


DisplayName : Default Domain Controllers Policy
GpoStatus   : AllSettingsEnabled


DisplayName : NRPT_settings
GpoStatus   : AllSettingsEnabled
```

   In the previous example, the Format-List (fl) parameter is used to display only the **DisplayName** and **GpoStatus** properties. The **NRPT_settings** GPO is a custom GPO created in this example to contain all name resolution policies for the domain. This GPO name is only an example. GPO names for your domain will be different. The example GPO is used in the following example.
3. To review name resolution policy settings in a GPO, use the **Get-DnsClientNrptRule** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptRule -GpoName
contoso.com\NRPT_settings


Name                            : Rule1
Version                         : 1
Namespace                       : {.secure.contoso.com}
```

```
IPsecCARestriction                :
DirectAccessDnsServers            :
DirectAccessEnabled               : False
DirectAccessProxyType             :
DirectAccessProxyName             :
DirectAccessQueryIPsecEncryption  :
DirectAccessQueryIPsecRequired    :
NameServers                       :
DnsSecEnabled                     : True
DnsSecQueryIPsecEncryption        :
DnsSecQueryIPsecRequired          : False
DnsSecValidationRequired          : True
NameEncoding                      : Disable
DisplayName                       :
Comment                           :
```

In the previous example, the *secure.contoso.com* namespace displays **True** next to **DnsSecValidationRequired**, which means that all DNS clients that are security-aware, and that have this GPO applied, will require that queries for the *secure.contoso.com* zone are validated.

Review effective name resolution policy

Although a GPO can be configured with NRPT settings, these settings will not affect a DNS client unless they are correctly applied. To verify that settings are applied, review the effective name resolution policy.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review effective name resolution policy

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To review effective name resolution policy, use the **Get-DnsClientNrptPolicy** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptPolicy -Effective

Namespace                         : .secure.contoso.com
QueryPolicy                       : QueryIPv6Only
```

```
SecureNameQueryFallback          : FallbackPrivate
DirectAccessIPsecCARestriction   :
DirectAccessProxyName            :
DirectAccessDnsServers           :
DirectAccessEnabled              : False
DirectAccessProxyType            :
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired   :
NameServers                      :
DnsSecIPsecCARestriction         :
DnsSecQueryIPsecEncryption       :
DnsSecQueryIPsecRequired         : False
DnsSecValidationRequired         : True
NameEncoding                     :
```

In the previous example, **True** is displayed next to **DnsSecValidationRequired** for the *secure.contoso.com* namespace, which means that DNSSEC validation will be required for DNS queries in the *secure.contoso.com* zone. Output from the **Get-DnsClientNrptPolicy** cmdlet also includes any name resolution policy that might be configured in local Group Policy.

Verify DNSSEC validation

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To verify DNSSEC validation

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To verify DNSSEC validation, use the **Resolve-DnsName** cmdlet. First, verify that name resolution is successful when querying an authoritative DNS server. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
dc2.contoso.com -dnssecok


Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
```

```
--                -----------------           ------------
secure.contoso.com          SOA  3600  Authority
dc2.contoso.com             hostmaster.contoso.com      6

Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 3600
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {56, 217, 3, 225...}


Name      :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 3600
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
TTL         : 3600
Section     : Authority
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
```

```
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {} PS C:\> Get-DnsServerZone -Name
corp.contoso.com

ZoneName                              ZoneType
IsAutoCreated   IsDsIntegrated  IsReverseLookupZone  IsSigned
--------                              --------        ---------
----   --------------  -------------------  --------
corp.contoso.com                      Primary         False
False           False                 True
```

In this example, the primary authoritative DNS server *dc2.contoso.com* is queried. Since authoritative responses are always valid, the DNS query will succeed. DNSSEC-related resource records are displayed because the zone is signed. If a zone is not signed, such as the contoso.com zone in the following example, RRSIG records are not displayed in the output:

```
PS C:\> Resolve-DnsName contoso.com -Server dc2.contoso.com -
dnssecok

Name                                           Type   TTL
Section    IPAddress
----                                           ----   ---   -
------    ---------
contoso.com                                    A      600
Answer     192.168.0.2
contoso.com                                    A      600
Answer     192.168.0.3
contoso.com                                    A      600
```

```
Answer     192.168.0.1


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

3. Next, use the **Resolve-DnsName** cmdlet to query a non-authoritative, resolving DNS server with a valid trust anchor installed. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok

Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
--                ------------------           ------------
secure.contoso.com             SOA  830   Authority
dc2.contoso.com                hostmaster.contoso.com      6


Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {56, 217, 3, 225...}


Name      :
```

```
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 830
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

In this example, a non-domain-joined non-authoritative recursive DNS server at *192.168.0.4* is queried. Since this DNS server has a valid trust anchor installed for the *secure.contoso.com* zone, the DNS query will succeed. Again, DNSSEC-related resource records are displayed because the zone is signed.

If DNSSEC validation does not succeed, an error will be displayed, see the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok
Resolve-DnsName : secure.contoso.com : DNS server failure
```

```
At line:1 char:1
+ Resolve-DnsName secure.contoso.com -Server 192.168.0.4 -
dnssecok
+
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~~~
    + CategoryInfo          : ResourceUnavailable:
(secure.contoso.com:String) [Resolve-DnsName], Win32Exception
    + FullyQualifiedErrorId :
RCODE_SERVER_FAILURE,Microsoft.DnsClient.Commands.ResolveDnsN
ame
```

In the previous example, the DNS server at 192.168.0.4 does not have a valid trust anchor installed for the secure.contoso.com domain.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Review and Manage a Signed Zone

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Review and Manage a Signed Zone**

This parent checklist includes links to subordinate checklists that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about DNSSEC and signing at least one DNS zone.

Most of the tasks in this checklist are optional. Depending on your environment, you might need to perform some or all of these procedures.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Review and Manage a Signed Zone

	Task	Reference
☐	Review DNSSEC concepts for managing a signed zone.	DNSSEC in Windows
☐	Review the status of your signed zone.	Procedure: Review DNSSEC Parameters and Settings
☐	(Optional) Modify zone signing parameters.	Checklist: Reconfigure Zone Signing Parameters on a Signed Zone
☐	(Optional) Perform a manual key rollover.	Checklist: Perform a Manual Key Rollover
☐	(Optional) Revoke signing keys.	Checklist: Perform an Emergency Key Revocation
☐	(Optional) Move the Key Master role to another DNS server.	Checklist: Move the Key Master Role
☐	(Optional) Unsign a zone.	Checklist: Revert to an Unsigned Zone
☐	(Optional) Perform debug log rollover.	Procedure: Enable DNS Diagnostic Events

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Review DNSSEC Parameters and Settings

Use the following procedures to display and review DNSSEC parameter values and settings using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

You can use DNS Manager or Windows PowerShell to review DNSSEC parameters and settings:

- Review DNSSEC parameters and settings in DNS Manager
- Review DNSSEC parameters and settings in Windows PowerShell

Review DNSSEC parameters and settings in DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review DNSSEC parameters and settings in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. Review settings on the following tabs: **Key Master**, **Next Secure (NSEC)**, **Trust Anchor**, and **Advanced**.
4. To review KSK settings, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then review information under **Summary**. You can also click **Edit** to review KSK settings.
5. To review ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then review information under **Summary**. You can also click **Edit** to review ZSK settings.
6. Click **Cancel** when finished.

Review DNSSEC parameters and settings in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To review DNSSEC parameters and settings in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To review a summary of zone settings, use the **Get-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName
secure.contoso.com
```

```
ZoneName                          : secure.contoso.com
IsKeyMasterServer                 : True
KeyMasterServer                   : DC2.contoso.com
KeyMasterStatus                   : Online
DenialOfExistence                 : NSec3
NSec3HashAlgorithm                : RsaSha1
NSec3Iterations                   : 50
NSec3OptOut                       : False
IsNSec3SaltConfigured             : True
NSec3RandomSaltLength             : 8
NSec3UserSalt                     : -
DnsKeyRecordSetTTL                : 01:00:00
DSRecordSetTTL                    : 01:00:00
DSRecordGenerationAlgorithm       : {Sha1, Sha256}
DistributeTrustAnchor             : {None}
EnableRfc5011KeyRollover          : True
ParentHasSecureDelegation         : False
SecureDelegationPollingPeriod     : 12:00:00
PropagationTime                   : 2.00:00:00
SignatureInceptionOffset          : 01:00:00
```

3. To review a summary of signing keys, use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


KeyId                                    KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                    -------         ------
---------    ---------  -------------   -----------------
1adabf48-5039-4e71-907c-99ff87f43564   KeySigningKey
RsaSha256          2048       True            True
08bf6bd0-b242-49d5-8674-eff9d338ea2f   ZoneSigningKey
RsaSha256          1024       True            True
```

4. You can also review detailed information for signing keys using Format-List (fl). See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| fl *


KeyId                          : 1adabf48-5039-4e71-907c-
99ff87f43564
IsRolloverEnabled              : True
ActiveKey                      : {B4F136C0-4CF8-4CF9-BDBE-
73AC28780767}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
CurrentState                   : Active
DnsKeySignatureValidityPeriod  : 7.00:00:00
DSSignatureValidityPeriod      : 7.00:00:00
InitialRolloverOffset          : 00:00:00
KeyLength                      : 2048
KeyStorageProvider             : Microsoft Software Key
Storage Provider
KeyType                        : KeySigningKey
LastRolloverTime               :
NextKey                        :
NextRolloverAction             : Normal
NextRolloverTime               : 1/17/2016 3:56:40 PM
RolloverPeriod                 : 755.00:00:00
RolloverType                   : DoubleSignature
StandbyKey                     : {4F39AD5A-7FB7-43EF-BCD3-
F33494ABE523}
StoreKeysInAD                  : True
ZoneName                       : secure.contoso.com
ZoneSignatureValidityPeriod    : 10.00:00:00
PSComputerName                 :
CimClass                       :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties          : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties            :
```

```
Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                          : 08bf6bd0-b242-49d5-8674-
eff9d338ea2f
IsRolloverEnabled              : True
ActiveKey                      : {B426354D-9209-4CFD-9049-
6E503ADA14D0}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
CurrentState                   : Active
DnsKeySignatureValidityPeriod  : 7.00:00:00
DSSignatureValidityPeriod      : 7.00:00:00
InitialRolloverOffset          : 00:00:00
KeyLength                      : 1024
KeyStorageProvider             : Microsoft Software Key
Storage Provider
KeyType                        : ZoneSigningKey
LastRolloverTime               :
NextKey                        : {8265DEA2-A0EB-45B4-B1E8-
69F61777EB0E}
NextRolloverAction             : Normal
NextRolloverTime               : 3/23/2014 4:56:40 PM
RolloverPeriod                 : 90.00:00:00
RolloverType                   : PrePublish
StandbyKey                     :
StoreKeysInAD                  : True
ZoneName                       : secure.contoso.com
ZoneSignatureValidityPeriod    : 10.00:00:00
PSComputerName                 :
CimClass                       :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties          : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties            :
Microsoft.Management.Infrastructure.CimSystemProperties
```

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Enable DNS Diagnostic Events

DNS server debug logging is enabled by default with individual diagnostic events disabled. You can use the procedures in this topic to enable diagnostic event logging and change other event log parameters.

Tip

By default, the DNS debug log is located in the **%windir%\system32\dns** directory.

In this section

To view event logging status

To enable diagnostic events in the event log

To enable debug log rollover

View and modify the event logging status

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To view event logging status

1. Open an elevated Windows PowerShell prompt on a DNS server.
2. Use the **Get-DnsServerDiagnostics** cmdlet to view the status of individual diagnostic events. See the following example:

```
PS C:\> Get-DnsServerDiagnostics


SaveLogsToPersistentStorage          : False
Queries                              : False
Answers                              : False
Notifications                        : False
Update                               : False
QuestionTransactions                 : False
UnmatchedResponse                    : False
```

```
SendPackets                           : False
ReceivePackets                        : False
TcpPackets                            : False
UdpPackets                            : False
FullPackets                           : False
FilterIPAddressList                   :
EventLogLevel                         : 4
UseSystemEventLog                     : False
EnableLoggingToFile                   : True
EnableLogFileRollover                 : False
LogFilePath                           :
MaxMBFileSize                         : 500000000
WriteThrough                          : False
EnableLoggingForLocalLookupEvent      : False
EnableLoggingForPluginDllEvent        : False
EnableLoggingForRecursiveLookupEvent  : False
EnableLoggingForRemoteServerEvent     : False
EnableLoggingForServerStartStopEvent  : False
EnableLoggingForTombstoneEvent        : False
EnableLoggingForZoneDataWriteEvent    : False
EnableLoggingForZoneLoadingEvent      : False
```

To enable diagnostic events in the event log

1. Open an elevated Windows PowerShell prompt on the DNS server where you wish to enable event logging.
2. Use the **Set-DnsServerDiagnostics** cmdlet to enable individual diagnostic events, or you can enable all diagnostic events at once. See the following example:

   ```
   PS C:\> Set-DnsServerDiagnostics -All $true
   ```

   The previous command will set the status of all diagnostic events to **True**. You can verify this with the **Get-DnsServerDiagnostics** cmdlet. See the following example:

   ```
   PS C:\> Get-DnsServerDiagnostics


   SaveLogsToPersistentStorage           : True
   Queries                               : True
   Answers                               : True
   ```

```
Notifications                           : True
Update                                  : True
QuestionTransactions                    : True
UnmatchedResponse                       : True
SendPackets                             : True
ReceivePackets                          : True
TcpPackets                              : True
UdpPackets                              : True
FullPackets                             : True
FilterIPAddressList                     :
EventLogLevel                           : 7
UseSystemEventLog                       : False
EnableLoggingToFile                     : True
EnableLogFileRollover                   : True
LogFilePath                             :
MaxMBFileSize                           : 500000000
WriteThrough                            : True
EnableLoggingForLocalLookupEvent        : True
EnableLoggingForPluginDllEvent          : True
EnableLoggingForRecursiveLookupEvent    : True
EnableLoggingForRemoteServerEvent       : True
EnableLoggingForServerStartStopEvent    : True
EnableLoggingForTombstoneEvent          : True
EnableLoggingForZoneDataWriteEvent      : True
EnableLoggingForZoneLoadingEvent        : True
```

3. You can also enable only selected diagnostic events. See the following example:

```
PS C:\> Set-DnsServerDiagnostics -Queries $true -Answers
$true -Notifications $true -SendPackets $true -ReceivePackets
$true -TcpPackets $true -UdpPackets $true
PS C:\> Get-DnsServerDiagnostics

SaveLogsToPersistentStorage          : False
Queries                              : True
```

```
Answers                                 : True
Notifications                           : True
Update                                  : False
QuestionTransactions                    : False
UnmatchedResponse                       : False
SendPackets                             : True
ReceivePackets                          : True
TcpPackets                              : True
UdpPackets                              : True
FullPackets                             : False
FilterIPAddressList                     :
EventLogLevel                           : 0
UseSystemEventLog                       : False
EnableLoggingToFile                     : False
EnableLogFileRollover                   : False
LogFilePath                             :
MaxMBFileSize                           : 500000000
WriteThrough                            : False
EnableLoggingForLocalLookupEvent        : False
EnableLoggingForPluginDllEvent          : False
EnableLoggingForRecursiveLookupEvent    : False
EnableLoggingForRemoteServerEvent       : False
EnableLoggingForServerStartStopEvent    : False
EnableLoggingForTombstoneEvent          : False
EnableLoggingForZoneDataWriteEvent      : False
EnableLoggingForZoneLoadingEvent        : False
```

To enable debug log rollover

1. Open an elevated Windows PowerShell prompt on the DNS server where you wish to enable event logging.
2. Use the **Set-DnsServerDiagnostics** cmdlet to enable debug log rollover. See the following example:

   ```
   PS C:\> Set-DnsServerDiagnostics - EnableLogFileRollover
   $true
   ```

3.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Review Key Rollover Status

Use the following procedures to display and review rollover status of signing keys using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

You can use DNS Manager or Windows PowerShell to review rollover status:

- Review key rollover status in DNS Manager
- Review key rollover status in Windows PowerShell

Review key rollover status in DNS Manager

You can use DNS Manager to review basic key rollover status. However, Windows PowerShell provides additional settings for review.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review key rollover status in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. To review KSK status, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then review information under **Rollover status**.
4. To review ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then review information under **Rollover status**.
5. Click **Cancel** when finished.

Review key rollover status in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To review key rollover status in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To review signing key status, use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| fl *


KeyId                          : 0316fff3-03a8-451d-a4ad-
db2f93e4070a
IsRolloverEnabled              : True
ActiveKey                      : {8F7CB3C7-A9A1-496A-96E6-
D6D845D2F0F1}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
CurrentState                   : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod      : 7.00:00:00
InitialRolloverOffset          : 00:00:00
KeyLength                      : 2048
KeyStorageProvider             : Microsoft Software Key
Storage Provider
KeyType                        : KeySigningKey
LastRolloverTime               :
NextKey                        :
NextRolloverAction             : Normal
NextRolloverTime               : 2/10/2016 2:27:53 PM
RolloverPeriod                 : 755.00:00:00
RolloverType                   : DoubleSignature
StandbyKey                     : {1F4FBEBD-D66F-485C-9076-
187DD40B1993}
StoreKeysInAD                  : True
ZoneName                       : secure.contoso.com
ZoneSignatureValidityPeriod    : 10.00:00:00
PSComputerName                 :
```

```
CimClass                          :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties             : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties               :
Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                             : 60b46b4d-e1cb-433a-ac33-
5d46687e860e
IsRolloverEnabled                 : True
ActiveKey                         : {93BA0FC4-3931-4DBF-A205-
F4EFB24C4E4E}
CryptoAlgorithm                   : RsaSha256
CurrentRolloverStatus             : NotRolling
CurrentState                      : Active
DnsKeySignatureValidityPeriod     : 7.00:00:00
DSSignatureValidityPeriod         : 7.00:00:00
InitialRolloverOffset             : 00:00:00
KeyLength                         : 1024
KeyStorageProvider                : Microsoft Software Key
Storage Provider
KeyType                           : ZoneSigningKey
LastRolloverTime                  :
NextKey                           : {D1FDEAF4-F9AB-4D59-AF5A-
A33856B2B0EE}
NextRolloverAction                : Normal
NextRolloverTime                  : 4/16/2014 3:27:53 PM
RolloverPeriod                    : 90.00:00:00
RolloverType                      : PrePublish
StandbyKey                        :
StoreKeysInAD                     : True
ZoneName                          : secure.contoso.com
ZoneSignatureValidityPeriod       : 10.00:00:00
PSComputerName                    :
CimClass                          :
```

```
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties
```

In the previous example, the Format-List (fl) parameter is used to display extended signing key status information.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Revert to an Unsigned Zone

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Revert to an Unsigned Zone**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Revert to an Unsigned Zone

	Task	Reference
☐	Review concepts about the consequences of unsigning a zone. If trust anchors are deployed, or name resolution policy requires validation, DNS resolution can fail when a zone is unsigned.	Trust Anchors The NRPT
☐	Configure the NRPT so that validation is not required for the	Procedure: Configure the

	Task	Reference
	zone.	NRPT
☐	Verify that name resolution policy does not require DNSSEC validation for the zone.	Procedure: Verify Name Resolution Policy
☐	Remove trust anchors for the zone from validating DNS servers.	Procedure: Remove a Trust Point
☐	Unsign the zone	Procedure: Unsign a Zone
☐	Verify that DNS clients are able to resolve names in the domain.	Procedure: Verify DNS Resolution

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Configure the NRPT

Use the following procedures to configure the Name Resolution Policy Table (NRPT). The examples that are provided demonstrate how to add a rule in an existing Group Policy Object (GPO) that requires DNSSEC validation for a namespace (zone).

Procedures are provided using the Group Policy Management console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For more information about NRPT rule processing, see The NRPT.

Choose the Group Policy Management console Windows PowerShell to configure the NRPT:

- Configure the NRPT with the Group Policy Management console
- Configure the NRPT with Windows PowerShell

Configure the NRPT with the Group Policy Management console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure the NRPT with the Group Policy Management console

1. On a computer with the Group Policy Management feature installed, click **Tools** on the Server Manager menu and then click **Group Policy Management**.
2. In the Group Policy Management console tree, navigate to **Forest: contoso.com\Domains\contoso.com\Group Policy Objects**.
3. Right-click the Group Policy Object (GPO) that you wish to configure, and then click **Edit**.
4. In the Group Policy Management Editor tree, navigate to **Computer Configuration\Policies\Windows Settings\Name Resolution Policy**.
5. Under **Create Rules**, choose the part of the namespace to which the rule will apply from the drop-down list. For example, choose **Suffix** to create a rule that applies to all names that end in the text that you provide.

   Tip
   To view an explanation of each choice in the drop-down list, choose an item and then click in the text box to the right. A description of the choice that is made will be displayed to the left under **Description**.
6. Next to the drop-down list choice (ex: Suffix), type the namespace to which the rule will apply. For example, to apply a rule to all devices in the *secure.contoso.com* zone, choose **Suffix** and type **secure.contoso.com**.
7. On the **DNSSEC** tab, select the **Enable DNSSEC in this rule** checkbox, and then select the **Require DNS clients to check that name and address data has been validated by the DNS server checkbox**.
8. At the bottom of this section of the NRPT, click **Create**.
9. Under **Name Resolution Policy Table**, verify that the new rule is displayed.
10. Close the Group Policy Management Editor.
11. Close the Group Policy Management console.

Configure the NRPT with Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the zone you wish to sign with DNSSEC. Also replace the example computer names with names of the DNS servers you will use.

To configure the NRPT with Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a computer with the Group Policy Management feature installed.
2. To add an NRPT rule to an existing Group Policy Object (GPO), use the **Add-DnsClientNrptRule** cmdlet. See the following example.

```
PS C:\> Add-DnsClientNrptRule -GpoName NRPT_settings -
NameSpace secure.contoso.com -DnsSecEnable -
DnsSecValidationRequired
```

   In this example, the GPO named **NRPT_settings** is configured with a rule requiring DNSSEC validation for the secure.contoso.com namespace.
3. To confirm that the GPO was successfully configured, use the **Get-DnsClientNrptRule** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptRule -GpoName NRPT_settings

Name                                : Rule1
Version                             : 1
Namespace                           : {.adatum.com}
IPsecCARestriction                  :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessProxyName               :
DirectAccessQueryIPsecEncryption    :
DirectAccessQueryIPsecRequired      :
NameServers                         :
DnsSecEnabled                       : True
DnsSecQueryIPsecEncryption          :
DnsSecQueryIPsecRequired            : False
DnsSecValidationRequired            : True
NameEncoding                        : Disable
DisplayName                         :
Comment                             :

Name                                : {77BCF3D4-8442-49F1-9E73-
4715CAD9888A}
```

```
Version                              : 1
Namespace                            : {secure.contoso.com}
IPsecCARestriction                   :
DirectAccessDnsServers               :
DirectAccessEnabled                  : False
DirectAccessProxyType                :
DirectAccessProxyName                :
DirectAccessQueryIPsecEncryption     :
DirectAccessQueryIPsecRequired       :
NameServers                          :
DnsSecEnabled                        : True
DnsSecQueryIPsecEncryption           :
DnsSecQueryIPsecRequired             : False
DnsSecValidationRequired             : True
NameEncoding                         : Disable
DisplayName                          :
Comment                              :
```

In this example, a second rule was added to the **NRPT_settings** GPO. The rule for the *adatum.com* namespace was already present, and the rule for the *secure.contoso.com* namespace was added.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Verify Name Resolution Policy

Use the following procedures to verify name resolution policy on a DNS server and a DNS client using Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For information about configuring name resolution policy, see The NRPT.

Name resolution policy settings can only be applied to computers that are security-aware. For more information about security-aware computers, see Security-aware client.

See the following procedures:

- Review NRPT configuration: Check settings in the Name Resolution Policy Table (NRPT).
- Review effective name resolution policy: Verify whether or not DNSSEC validation is required for a given namespace.
- Verify DNSSEC validation: If DNSSEC validation has been enabled or disabled recently, verify that name resolution is working as expected.

Review NRPT configuration

Name resolution policy is configured in the NRPT, either in domain Group Policy, local Group Policy, or both. The NRPT can be configured to require or not require that DNSSEC validation be performed for DNS queries within a given namespace.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review NRPT configuration

1. Open an elevated Windows PowerShell prompt on a computer with the Group Policy Management feature installed.
2. To view a list of Group Policy Objects (GPOs) in the domain, use the **Get-GPO** cmdlet. See the following example:

```
PS C:\> Get-GPO -All -DomainName contoso.com | fl -Property
DisplayName,GpoStatus


DisplayName : Default Domain Policy
GpoStatus   : AllSettingsEnabled


DisplayName : Default Domain Controllers Policy
GpoStatus   : AllSettingsEnabled


DisplayName : NRPT_settings
GpoStatus   : AllSettingsEnabled
```

   In the previous example, the Format-List (fl) parameter is used to display only the **DisplayName** and **GpoStatus** properties. The **NRPT_settings** GPO is a custom GPO created in this example to contain all name resolution policies for the domain. This GPO name is only an example. GPO names for your domain will be different. The example GPO is used in the following example.
3. To review name resolution policy settings in a GPO, use the **Get-DnsClientNrptRule** cmdlet. See the following example:

```
PS C:\> Get-DnsClientNrptRule -GpoName
```

```
contoso.com\NRPT_settings


Name                                : Rule1
Version                             : 1
Namespace                           : {.secure.contoso.com}
IPsecCARestriction                  :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessProxyName               :
DirectAccessQueryIPsecEncryption    :
DirectAccessQueryIPsecRequired      :
NameServers                         :
DnsSecEnabled                       : True
DnsSecQueryIPsecEncryption          :
DnsSecQueryIPsecRequired            : False
DnsSecValidationRequired            : True
NameEncoding                        : Disable
DisplayName                         :
Comment                             :
```

In the previous example, the *secure.contoso.com* namespace displays **True** next to **DnsSecValidationRequired**, which means that all DNS clients that are security-aware, and that have this GPO applied, will require that queries for the *secure.contoso.com* zone are validated.

Review effective name resolution policy

Although a GPO can be configured with NRPT settings, these settings will not affect a DNS client unless they are correctly applied. To verify that settings are applied, review the effective name resolution policy.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review effective name resolution policy

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To review effective name resolution policy, use the **Get-DnsClientNrptPolicy** cmdlet.

See the following example:

```
PS C:\> Get-DnsClientNrptPolicy -Effective


Namespace                           : .secure.contoso.com
QueryPolicy                         : QueryIPv6Only
SecureNameQueryFallback             : FallbackPrivate
DirectAccessIPsecCARestriction      :
DirectAccessProxyName               :
DirectAccessDnsServers              :
DirectAccessEnabled                 : False
DirectAccessProxyType               :
DirectAccessQueryIPsecEncryption    :
DirectAccessQueryIPsecRequired      :
NameServers                         :
DnsSecIPsecCARestriction            :
DnsSecQueryIPsecEncryption          :
DnsSecQueryIPsecRequired            : False
DnsSecValidationRequired            : True
NameEncoding                        :
```

In the previous example, **True** is displayed next to **DnsSecValidationRequired** for the *secure.contoso.com* namespace, which means that DNSSEC validation will be required for DNS queries in the *secure.contoso.com* zone. Output from the **Get-DnsClientNrptPolicy** cmdlet also includes any name resolution policy that might be configured in local Group Policy.

Verify DNSSEC validation

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To verify DNSSEC validation

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To verify DNSSEC validation, use the **Resolve-DnsName** cmdlet. First, verify that name resolution is successful when querying an authoritative DNS server. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
dc2.contoso.com -dnssecok
```

```
Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
--                ------------------           ------------
secure.contoso.com             SOA  3600  Authority
dc2.contoso.com                hostmaster.contoso.com      6

Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 3600
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {56, 217, 3, 225...}


Name      :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 3600
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
TTL         : 3600
Section     : Authority
```

```
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {} PS C:\> Get-DnsServerZone -Name
corp.contoso.com

ZoneName                              ZoneType
IsAutoCreated   IsDsIntegrated  IsReverseLookupZone  IsSigned
--------                              --------        ---------
----   --------------  -------------------  --------
corp.contoso.com                      Primary         False
False           False                 True
```

In this example, the primary authoritative DNS server *dc2.contoso.com* is queried. Since authoritative responses are always valid, the DNS query will succeed. DNSSEC-related resource records are displayed because the zone is signed. If a zone is not signed, such as the contoso.com zone in the following example, RRSIG records are not displayed in the output:

```
PS C:\> Resolve-DnsName contoso.com -Server dc2.contoso.com -
dnssecok

Name                                           Type   TTL
Section    IPAddress
----                                           ----   ---   -
------    ---------
contoso.com                                    A      600
```

```
Answer     192.168.0.2
contoso.com                                    A      600
Answer     192.168.0.3
contoso.com                                    A      600
Answer     192.168.0.1


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

3. Next, use the **Resolve-DnsName** cmdlet to query a non-authoritative, resolving DNS server with a valid trust anchor installed. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok

Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
--               -----------------           ------------
secure.contoso.com             SOA  830   Authority
dc2.contoso.com                hostmaster.contoso.com      6


Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
```

```
Signature   : {56, 217, 3, 225...}


Name      :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 830
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

In this example, a non-domain-joined non-authoritative recursive DNS server at *192.168.0.4* is queried. Since this DNS server has a valid trust anchor installed for the *secure.contoso.com* zone, the DNS query will succeed. Again, DNSSEC-related resource records are displayed because the zone is signed.

If DNSSEC validation does not succeed, an error will be displayed, see the following

example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok
Resolve-DnsName : secure.contoso.com : DNS server failure
At line:1 char:1
+ Resolve-DnsName secure.contoso.com -Server 192.168.0.4 -
dnssecok
+
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~~~
    + CategoryInfo          : ResourceUnavailable:
(secure.contoso.com:String) [Resolve-DnsName], Win32Exception
    + FullyQualifiedErrorId :
RCODE_SERVER_FAILURE,Microsoft.DnsClient.Commands.ResolveDnsN
ame
```

In the previous example, the DNS server at 192.168.0.4 does not have a valid trust anchor installed for the secure.contoso.com domain.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Remove a Trust Point

Use the following procedures to delete trust anchors from a DNS server. Before deleting a trust anchor, you might wish to export trust points for purposes of backup. For information about exporting a trust point, see Procedure: Export a Trust Point.

For more information about trust points, see Trust Anchors.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Delete a trust point using DNS Manager
- Delete a trust point using Windows PowerShell

Delete a trust point using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To delete a trust point using DNS Manager

1. Open DNS Manager on the server where you wish to delete the trust point, or connect to the DNS server from another location using DNS Manager.
2. In the console tree, navigate to the location under **Trust Points** corresponding to the zone for which you will delete a trust point. For example to delete a trust point for the zone: secure.contoso.com, navigate to Trust Points > com > contoso > secure.
3. In the right-hand pane, right-click the trust anchor you wish to remove, and then click **Delete**.

When the last trust anchor is removed for a zone, right-click **Trust Points** in DNS Manager and then click **Refresh** to remove the container (folder) in DNS Manager for that namespace.

Delete a trust point using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To delete a trust point using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the server where you wish to delete the trust point.
2. To delete a trust point, use the **Remove-DnsServerTrustAnchor** cmdlet. See the following example:

   ```
   Remove-DnsServerTrustAnchor -Name secure.contoso.com -Type DS
   -Force
   ```

   In the previous example, the **Type** parameter was used to specify that only dsset trust anchors should be deleted from the local server for the *secure.contoso.com zone*.
3. To delete DNSKEY trust anchors, specify a Type of DnsKey. See the following example:

   ```
   Remove-DnsServerTrustAnchor -Name fb.contoso.com -Type DnsKey
   -Force
   ```
4. To delete all trust anchors, do not specify a Type.

When the last trust anchor is removed for a zone using Windows PowerShell, the container (folder) in DNS Manager for that namespace is also removed.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Unsign a Zone

Use the following procedures to unsign a zone that has been signed with DNSSEC on Windows Server® 2012 or Windows Server® 2012 R2. Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Unsign a zone in DNS Manager
- Unsign a zone in Windows PowerShell

Unsign a zone in DNS Manager

The **Unsign zone** wizard is used to unsign a zone in DNS Manager. For examples of the wizard, see Unsign Zone Wizard.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To unsign a zone in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click an signed zone, point to **DNSSEC**, and then click **Unsign the Zone**.
3. Click **Next** to immediately unsign the zone.
4. Verify that **The zone has been successfully unsigned** is displayed, and then click **Finish**.
5. Right-click the zone and click **Refresh** to refresh the console view. DNSSEC-related resource records will be removed from the zone.

If the zone is Active Directory-integrated, it will not be immediately unsigned on all primary authoritative DNS servers. For performance reasons, zone unsigning is staggered.

Unsign a zone in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To unsign a zone in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To unsign a zone, use the **Invoke-DnsServerZoneUnsign** cmdlet. See the following example.

   ```
   PS C:\> Invoke-DnsServerZoneUnsign -ZoneName
   secure.contoso.com -Force
   ```

   In this example, the zone *secure.contoso.com* is unsigned on the local DNS server. The **Force** parameter is used to skip confirmation.

If the zone is Active Directory-integrated, it will not be immediately unsigned on all primary authoritative DNS servers. For performance reasons, zone unsigning is staggered.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Verify DNS Resolution

Use the following procedures to verify DNS resolution on a DNS client using Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

For information about configuring name resolution policy, see The NRPT.

Name resolution policy settings can only be applied to computers that are security-aware. For more information about security-aware computers, see Security-aware client.

See the Verify DNSSEC validation procedure in this topic for steps to verify that name resolution is working as expected.

Verify DNSSEC validation

To use the example Windows PowerShell commands that are provided, replace the server and zone names with the names of the servers and zones used in your environment.

To verify DNSSEC validation

1. Open a Windows PowerShell prompt on a DNS client computer.
2. To verify DNSSEC validation, use the **Resolve-DnsName** cmdlet. First, verify that name resolution is successful when querying an authoritative DNS server. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
dc2.contoso.com -dnssecok

Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
--               ------------------           ------------
secure.contoso.com             SOA  3600  Authority
dc2.contoso.com                hostmaster.contoso.com      6

Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 3600
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {56, 217, 3, 225...}


Name      :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 3600
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
```

```
TTL         : 3600
Section     : Authority
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {} PS C:\> Get-DnsServerZone -Name
corp.contoso.com

ZoneName                              ZoneType
IsAutoCreated   IsDsIntegrated  IsReverseLookupZone  IsSigned
--------                              --------        ---------
----   --------------  -------------------  --------

corp.contoso.com                      Primary         False
False           False                 True
```

In this example, the primary authoritative DNS server *dc2.contoso.com* is queried. Since authoritative responses are always valid, the DNS query will succeed. DNSSEC-related resource records are displayed because the zone is signed. If a zone is not signed, such as the contoso.com zone in the following example, RRSIG records are not displayed in the output:

```
PS C:\> Resolve-DnsName contoso.com -Server dc2.contoso.com -
dnssecok

Name                                           Type   TTL
Section    IPAddress
----                                           ----   ---   -
```

```
------    ---------
contoso.com                                A      600
Answer     192.168.0.2
contoso.com                                A      600
Answer     192.168.0.3
contoso.com                                A      600
Answer     192.168.0.1

Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

3. Next, use the **Resolve-DnsName** cmdlet to query a non-authoritative, resolving DNS server with a valid trust anchor installed. See the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok

Name                           Type TTL   Section
PrimaryServer                  NameAdministrator
SerialNumber
----                           ---- ---   -------    -----------
--               -----------------           ------------
secure.contoso.com             SOA  830   Authority
dc2.contoso.com                hostmaster.contoso.com      6

Name        : secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : SOA
Algorithm   : 8
LabelCount  : 3
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
```

```
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {56, 217, 3, 225...}


Name      :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType : NSEC3
TTL       : 830
Section   : Authority


Name        :
vfo5uo7g4u40gj25v1s0p3mh3pp7aai2.secure.contoso.com
QueryType   : RRSIG
TTL         : 830
Section     : Authority
TypeCovered : NSEC3
Algorithm   : 8
LabelCount  : 4
OriginalTtl : 3600
Expiration  : 1/26/2014 10:27:53 PM
Signed      : 1/16/2014 9:27:53 PM
Signer      : secure.contoso.com
Signature   : {97, 195, 245, 11...}


Name      : .
QueryType : OPT
TTL       : 32768
Section   : Additional
Data      : {}
```

In this example, a non-domain-joined non-authoritative recursive DNS server at *192.168.0.4* is queried. Since this DNS server has a valid trust anchor installed for the *secure.contoso.com* zone, the DNS query will succeed. Again, DNSSEC-related resource

records are displayed because the zone is signed.

If DNSSEC validation does not succeed, an error will be displayed, see the following example:

```
PS C:\> Resolve-DnsName secure.contoso.com -Server
192.168.0.4 -dnssecok
Resolve-DnsName : secure.contoso.com : DNS server failure
At line:1 char:1
+ Resolve-DnsName secure.contoso.com -Server 192.168.0.4 -
dnssecok
+
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~~~
    + CategoryInfo          : ResourceUnavailable:
(secure.contoso.com:String) [Resolve-DnsName], Win32Exception
    + FullyQualifiedErrorId :
RCODE_SERVER_FAILURE,Microsoft.DnsClient.Commands.ResolveDnsN
ame
```

In the previous example, the DNS server at 192.168.0.4 does not have a valid trust anchor installed for the secure.contoso.com domain.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Manage Signing Keys

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Manage Signing Keys**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Manage Signing Keys

	Task	Reference
☐	Review important concepts about choosing and managing DNSSEC signing keys.	DNS Zones Cryptographic algorithms Trust Anchors
☐	(Optional) Perform a manual key rollover and revocation.	Checklist: Perform an Emergency Key Revocation
☐	(Optional) Perform a manual key rollover without revocation.	Checklist: Perform a Manual Key Rollover
☐	Review the status of signing keys.	Procedure: Review Signing Keys
☐	Retire a signing key on the next key rollover.	Procedure: Retire a Signing Key

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Review Signing Keys

Use the following procedures to display and review DNSSEC signing keys using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

You can use DNS Manager or Windows PowerShell to review DNSSEC signing keys:

- Review DNSSEC signing keys in DNS Manager
- Review DNSSEC signing keys in Windows PowerShell

Review DNSSEC signing keys in DNS Manager

You can use DNS Manager to review some signing key settings. However, Windows PowerShell provides additional settings for review.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review DNSSEC signing keys in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. To review KSK status, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then review information under **Rollover status** and **Summary**. You can also click **Edit** to review KSK settings.
4. To review ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then review information under **Rollover status** and **Summary**. You can also click **Edit** to review ZSK settings.
5. Click **Cancel** when finished.

Review DNSSEC signing keys in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To review DNSSEC signing keys in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To review signing key status, use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| fl *


KeyId                          : 0316fff3-03a8-451d-a4ad-
db2f93e4070a
IsRolloverEnabled              : True
ActiveKey                      : {8F7CB3C7-A9A1-496A-96E6-
D6D845D2F0F1}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
```

```
CurrentState                    : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                       : 2048
KeyStorageProvider              : Microsoft Software Key
Storage Provider
KeyType                         : KeySigningKey
LastRolloverTime                :
NextKey                         :
NextRolloverAction              : Normal
NextRolloverTime                : 2/10/2016 2:27:53 PM
RolloverPeriod                  : 755.00:00:00
RolloverType                    : DoubleSignature
StandbyKey                      : {1F4FBEBD-D66F-485C-9076-
187DD40B1993}
StoreKeysInAD                   : True
ZoneName                        : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                  :
CimClass                        :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties             :
Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                           : 60b46b4d-e1cb-433a-ac33-
5d46687e860e
IsRolloverEnabled               : True
ActiveKey                       : {93BA0FC4-3931-4DBF-A205-
F4EFB24C4E4E}
CryptoAlgorithm                 : RsaSha256
CurrentRolloverStatus         : NotRolling
CurrentState                    : Active
```

```
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                     : 1024
KeyStorageProvider            : Microsoft Software Key
Storage Provider
KeyType                       : ZoneSigningKey
LastRolloverTime              :
NextKey                       : {D1FDEAF4-F9AB-4D59-AF5A-
A33856B2B0EE}
NextRolloverAction            : Normal
NextRolloverTime              : 4/16/2014 3:27:53 PM
RolloverPeriod                : 90.00:00:00
RolloverType                  : PrePublish
StandbyKey                    :
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties
```

In the previous example, the Format-List (fl) parameter is used to display extended signing key status information.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Retire a Signing Key

Use the following procedures to retire (revoke) a signing key using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Important
To revoke a signing key, it must not be the last ZSK or KSK available in the zone. If only one signing key of each type (ZSK and KSK) is currently configured in the zone, you must add an additional signing key to replace the key that will be retired.

When you retire a signing key, it will remain active in the zone until the next key rollover. To expedite key retirement, you can perform a manual key rollover. For more information, see Procedure: Perform a Manual Signing Key Rollover.

Important
When you rollover a ZSK, the rollover will wait for the MaxZone TTL before starting rollover, even if you attempt to manually roll the ZSK. If you have configured one or more resource records in the zone with a long TTL, rollover will be delayed by this period of time. To expedite key rollover, change the TTL of zone records to no more than 1 hour.

You can use DNS Manager or Windows PowerShell to retire DNSSEC signing keys:

- Revoke a signing key using DNS Manager
- Revoke a signing key using Windows PowerShell

Revoke a signing key using DNS Manager

You can use DNS Manager to review some signing key settings. However, Windows PowerShell provides additional settings for review.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To revoke a signing key using DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. Click the **KSK** tab or click the **ZSK** tab, click the signing key in the list that you will retire, and then click **Retire**.
4. When you are prompted that the key will be marked for retirement, click **OK**, and then click **OK** again.

Revoke a signing key using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To revoke a signing key using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. Use the **Get-DnsServerSigningKey** cmdlet to retrieve signing key IDs. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


KeyId                                   KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                   -------         ------
---------    ---------  -------------   -----------------
96e2236a-4a4e-4409-afe7-c3aa8ffc254e   KeySigningKey
RsaSha256          2048       True            True
bddf9bee-bfdd-41fc-a5c9-e5df8966dbaa   ZoneSigningKey
RsaSha256          1024       True            True
925795c2-4772-48d8-b0f0-0d79c8a7e5c7   ZoneSigningKey
RsaSha256          1024       True            True
```

   In the previous example, there are two ZSKs, therefore one can be retired.
3. To retire a signing key, use the **Set-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Set-DnsServerSigningKey -ZoneName secure.contoso.com
-KeyID bddf9bee-bfdd-41fc-a5c9-e5df8966dbaa –
NextRolloverAction Retire
```

To retire all but two signing keys in a zone, you can combine both cmdlets used in the previous procedure. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com | Set-DnsServerSigningKey
-NextRolloverAction Retire
```

An error will be displayed if you run this command because the last signing key of each type (KSK or ZSK) cannot be retired. All other keys are marked for retirement, however. The result of this command is that the zone will be signed with one KSK and one ZSK after the next key rollover.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Move the Key Master Role

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Move the Key Master Role**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Move the Key Master Role

	Task	Reference
☐	Review important concepts about the Key Master.	The Key Master
☐	Move the Key Master role when the current Key Master is online.	Procedure: Move the Key Master Role
☐	Move the Key Master role when the current Key Master is offline.	Procedure: Seize the Key Master Role

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Move the Key Master Role

Use the following procedures to move the Key Master role for a zone that has been signed with DNSSEC when the current Key Master server is online and the DNS service is running. To move the Key Master role when the current Key Master is offline or the DNS service is unavailable, see Procedure: Seize the Key Master Role.

To be a Key Master for a zone, a DNS server must be a primary, authoritative DNS server for that zone, and it must be capable of online signing of dynamic DNS zones. You cannot move the Key Master role for zones that are not Active Directory-integrated because these zones can have only one primary, authoritative DNS server.

For more information about the Key Master, see The Key Master.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Move the Key Master role using DNS Manager
- Move the Key Master role using Windows PowerShell

Move the Key Master role using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To move the Key Master role using DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server or connect to a primary, authoritative DNS server from another location using DNS Manager.
2. In the console tree, right-click the zone for which you plan to move the Key Master role, point to **DNSSEC**, and then click **Properties**.
3. On the **Key Master** tab, select **Use the following DNS server as the Key Master**.
4. Click the drop-down list, and then click **Yes** in the popup that appears. The local DNS server will create a list of DNS servers that are capable of being the new Key Master for the zone.
5. Choose a DNS server from the list for the new Key Master, and click **OK**.
6. Click **Yes** in the dialog box that appears to move the Key Master role to the DNS server that was chosen. If you click **No**, the operation will be cancelled and the Key Master will remain unchanged.
7. Confirm that **The Key Master for the zone <zone name> has been updated successfully** is displayed, and then click **OK**.

This procedure will immediately change the Key Master. You do not need to wait for Active Directory replication or re-sign the zone.

Move the Key Master role using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To move the Key Master role using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To move the Key Master role, use the **Reset-DnsServerZoneKeyMasterRole** cmdlet. See the following example.

   ```
   PS C:\> Reset-DnsServerZoneKeyMasterRole -ZoneName
   secure.contoso.com -KeyMasterServer dc3.contoso.com -Force
   ```

   In this example, the Force parameter is used to skip confirmation steps. The Key Master role for the zone *secure.contoso.com* is transferred to the server *dc3.contoso.com*.

To verify that the Key Master role was successfully transferred, you can use the **Get-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com


ZoneName                       : secure.contoso.com
IsKeyMasterServer              : False
KeyMasterServer                : DC3.contoso.com
KeyMasterStatus                : Online
DenialOfExistence              : NSec3
NSec3HashAlgorithm             : RsaSha1
NSec3Iterations                : 50
NSec3OptOut                    : False
IsNSec3SaltConfigured          : True
NSec3RandomSaltLength          : 8
NSec3UserSalt                  : -
DnsKeyRecordSetTTL             : 01:00:00
DSRecordSetTTL                 : 01:00:00
DSRecordGenerationAlgorithm    : {Sha1, Sha256}
DistributeTrustAnchor          : {None}
```

```
EnableRfc5011KeyRollover      : True
ParentHasSecureDelegation     : False
SecureDelegationPollingPeriod : 12:00:00
PropagationTime               : 2.00:00:00
SignatureInceptionOffset      : 01:00:00
```

This procedure will immediately change the Key Master. You do not need to wait for Active Directory replication or re-sign the zone.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Seize the Key Master Role

Use the following procedures to seize the Key Master role for a zone that has been signed with DNSSEC when the current Key Master server is offline and the DNS service is not available. To move the Key Master role when the current Key Master is online and available, see Procedure: Move the Key Master Role.

Important
When you seize the Key Master role, the new Key Master must have access to private key material for the zone. If the new Key Master does not have this access, all signing keys in the zone must be replaced with new keys and the zone must be re-signed. If trust anchors have been distributed for the zone, these trust anchors must also be manually updated after zone re-signing. For more information, see Seizing the Key Master role.

To be a Key Master for a zone, a DNS server must be a primary, authoritative DNS server for that zone, and it must be capable of online signing of dynamic DNS zones. You cannot move the Key Master role for zones that are not Active Directory-integrated because these zones can have only one primary, authoritative DNS server.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Seize the Key Master role using DNS Manager
- Seize the Key Master role using Windows PowerShell

Seize the Key Master role using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To seize the Key Master role using DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server or connect to a primary, authoritative DNS server from another location using DNS Manager.
2. In the console tree, right-click the zone for which you plan to move the Key Master role, point to **DNSSEC**, and then click **Properties**. An alert will be displayed indicating that DNSSEC settings for the zone could not be loaded from the Key Master.
3. In the pop-up alert that appears, click **Yes**.

   Note
   If you do not see an alert, verify that the Key Master is offline. If the Key Master is online and you still wish to move the Key Master role, see Procedure: Move the Key Master Role to move the Key Master. Do not use this procedure to seize the Key Master role.
4. On the **Key Master** tab, select **Use the following DNS server as the Key Master**.
5. Click the drop-down list, and then click **Yes** in the popup that appears. The local DNS server will create a list of DNS servers that are capable of being the new Key Master for the zone.
6. Choose a DNS server from the list for the new Key Master, and click **OK**.
7. Click **Yes** in the dialog box that appears to move the Key Master role to the DNS server that was chosen. If you click **No**, the operation will be cancelled and the Key Master will remain unchanged.
8. If you see a dialog box warning that the Key Master cannot be contacted to transfer private key material, this means that the current DNS server does not have access to private key material. In this scenario, transferring the Key Master role also requires that the zone be re-signed and trust anchors must be updated. Click **Yes** to continue.
9. Confirm that **The Key Master for the zone <zone name> has been updated successfully** is displayed, and then click **OK**.

This procedure will immediately change the Key Master.

Seize the Key Master role using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To seize the Key Master role using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To seize the Key Master role, use the **Reset-DnsServerZoneKeyMasterRole** cmdlet with the **SeizeRole** parameter. See the following example.

   ```
   PS C:\> Reset-DnsServerZoneKeyMasterRole -ZoneName
   secure.contoso.com -KeyMasterServer dc3.contoso.com –
   SeizeRole -Force
   ```

   In this example, the **SeizeRole** parameter is used to seize, rather than move, the Key Master role, and the **Force** parameter is used to skip confirmation steps. The Key Master role for the zone *secure.contoso.com* is transferred to the server *dc3.contoso.com*.

To verify that the Key Master role was successfully transferred, you can use the **Get-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName secure.contoso.com


ZoneName                       : secure.contoso.com
IsKeyMasterServer              : False
KeyMasterServer                : DC3.contoso.com
KeyMasterStatus                : Online
DenialOfExistence              : NSec3
NSec3HashAlgorithm             : RsaSha1
NSec3Iterations                : 50
NSec3OptOut                    : False
IsNSec3SaltConfigured          : True
NSec3RandomSaltLength          : 8
NSec3UserSalt                  : -
DnsKeyRecordSetTTL             : 01:00:00
DSRecordSetTTL                 : 01:00:00
DSRecordGenerationAlgorithm    : {Sha1, Sha256}
DistributeTrustAnchor          : {None}
EnableRfc5011KeyRollover       : True
ParentHasSecureDelegation      : False
SecureDelegationPollingPeriod : 12:00:00
PropagationTime                : 2.00:00:00
```

```
SignatureInceptionOffset      : 01:00:00
```

This procedure will immediately change the Key Master.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Reconfigure Zone Signing Parameters on a Signed Zone

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Reconfigure Zone Signing Parameters on a Signed Zone**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Reconfigure Zone Signing Parameters on a Signed Zone

	Task	Reference
☐	Review important concepts about zone signing.	DNS Zones
☐	Change the current zone signing parameter values to different values.	Procedure: Modify Zone Signing Parameters
☐	Verify zone signing parameters.	Procedure: Review DNSSEC Parameters and Settings

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Modify Zone Signing Parameters

Use the following procedures to review and modify DNSSEC parameter values and settings using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Note

You can modify DNSSEC settings using DNS Manager on any primary, authoritative DNS server. However, settings are loaded from the Key Master and any changes are subsequently made on the Key Master. If the Key Master is not available, settings cannot be modified. When modifying settings using Windows PowerShell, you must either specify the Key Master using the ComputerName parameter, or run Windows PowerShell on the Key Master.

You can use DNS Manager or Windows PowerShell to modify DNSSEC parameters and settings:

- Modify DNSSEC parameters and settings in DNS Manager
- Modify DNSSEC parameters and settings in Windows PowerShell

Modify DNSSEC parameters and settings in DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To modify DNSSEC parameters and settings in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. You can modify some, but not all settings on the following tabs: **Key Master**, **Next Secure (NSEC)**, **Trust Anchor**, and **Advanced**.
4. To review and edit KSK settings, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then click **Edit** to review or change KSK settings.

5. To review and edit ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then click **Edit** to review or change ZSK settings.
6. Click **OK** when finished.

Modify DNSSEC parameters and settings in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To modify DNSSEC parameters and settings in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the Key Master or another primary authoritative DNS server. If the local server is not the Key Master, you must add the **ComputerName** parameter to specify the Key Master.
2. To modify signed zone settings, use the **Set-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

   ```
   PS C:\> Set-DnsServerDnsSecZoneSetting -ZoneName
   secure.contoso.com -DnsKeyRecordSetTTL 02:00:00
   ```

   Use the **Get-DnsServerDnsSecZoneSetting** cmdlet to verify settings. See the following example:

   ```
   PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName
   secure.contoso.com | select DnsKeyRecordSetTTL


   DnsKeyRecordSetTTL
   ------------------
   02:00:00
   ```
3. To modify signing keys, use the **Set-DnsServerSigningKey** cmdlet. See the following example:

   ```
   PS C:\> Set-DnsServerSigningKey -ZoneName secure.contoso.com
   -KeyId 925795c2-4772-48d8-b0f0-0d79c8a7e5c7 -
   ZoneSignatureValidityPeriod 8.00:00:00
   ```

   Use the **Get-DnsServerSigningKey** cmdlet to verify settings. See the following example:

   ```
   PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
   | fl
   ```

```
ZoneName                          : secure.contoso.com
KeyId                             : 96e2236a-4a4e-4409-afe7-
c3aa8ffc254e
KeyType                           : KeySigningKey
CryptoAlgorithm                   : RsaSha256
KeyLength                         : 2048
KeyStorageProvider                : Microsoft Software Key
Storage Provider
StoreKeysInAD                     : True
CurrentState                      : Active
IsRolloverEnabled                 : True
RolloverType                      : DoubleSignature
RolloverPeriod                    : 755.00:00:00
InitialRolloverOffset             : 00:00:00
CurrentRolloverStatus             : NotRolling
NextRolloverAction                : Normal
LastRolloverTime                  :
NextRolloverTime                  : 2/10/2016 5:09:37 PM
DnsKeySignatureValidityPeriod     : 7.00:00:00
DSSignatureValidityPeriod         : 7.00:00:00
ZoneSignatureValidityPeriod       : 10.00:00:00
ActiveKey                         : {5A004FCE-55FE-41A2-84BA-
90D0F8BF7CE2}
StandbyKey                        : {85F701CF-08B2-484F-8954-
B737876B0A62}
NextKey                           :

ZoneName                          : secure.contoso.com
KeyId                             : 925795c2-4772-48d8-b0f0-
0d79c8a7e5c7
KeyType                           : ZoneSigningKey
CryptoAlgorithm                   : RsaSha256
KeyLength                         : 1024
KeyStorageProvider                : Microsoft Software Key
Storage Provider
```

```
StoreKeysInAD                   : True
CurrentState                    : Active
IsRolloverEnabled               : True
RolloverType                    : PrePublish
RolloverPeriod                  : 90.00:00:00
InitialRolloverOffset           : 00:00:00
CurrentRolloverStatus           : NotRolling
NextRolloverAction              : Normal
LastRolloverTime                :
NextRolloverTime                : 4/16/2014 6:09:37 PM
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
ZoneSignatureValidityPeriod   : 8.00:00:00
ActiveKey                       : {80FB1558-BA2A-4374-A99B-
534828FACBB6}
StandbyKey                      :
NextKey                         : {08A68540-47EB-4F89-AF57-
9ABC2B9A3A3A}
```

Some settings cannot be modified when the zone is currently signed or after signing keys have been provisioned. For example, you cannot change the key storage provider for an existing signing key. To change this parameter you must add a new signing key. For a list of parameters that can be modified, view the syntax for the **Set-DnsServerDnsSecZoneSetting** and **Set-DnsServerSigningKey** cmdlets (ex: Get-Command Set-DnsServerDnsSecZoneSetting -Syntax).

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Review DNSSEC Parameters and Settings

Use the following procedures to display and review DNSSEC parameter values and settings using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

You can use DNS Manager or Windows PowerShell to review DNSSEC parameters and settings:

- Review DNSSEC parameters and settings in DNS Manager
- Review DNSSEC parameters and settings in Windows PowerShell

Review DNSSEC parameters and settings in DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review DNSSEC parameters and settings in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. Review settings on the following tabs: **Key Master**, **Next Secure (NSEC)**, **Trust Anchor**, and **Advanced**.
4. To review KSK settings, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then review information under **Summary**. You can also click **Edit** to review KSK settings.
5. To review ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then review information under **Summary**. You can also click **Edit** to review ZSK settings.
6. Click **Cancel** when finished.

Review DNSSEC parameters and settings in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To review DNSSEC parameters and settings in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To review a summary of zone settings, use the **Get-DnsServerDnsSecZoneSetting** cmdlet. See the following example:

```
PS C:\> Get-DnsServerDnsSecZoneSetting -ZoneName
secure.contoso.com
```

```
ZoneName                        : secure.contoso.com
IsKeyMasterServer               : True
KeyMasterServer                 : DC2.contoso.com
KeyMasterStatus                 : Online
DenialOfExistence               : NSec3
NSec3HashAlgorithm              : RsaSha1
NSec3Iterations                 : 50
NSec3OptOut                     : False
IsNSec3SaltConfigured           : True
NSec3RandomSaltLength           : 8
NSec3UserSalt                   : -
DnsKeyRecordSetTTL              : 01:00:00
DSRecordSetTTL                  : 01:00:00
DSRecordGenerationAlgorithm     : {Sha1, Sha256}
DistributeTrustAnchor           : {None}
EnableRfc5011KeyRollover        : True
ParentHasSecureDelegation       : False
SecureDelegationPollingPeriod : 12:00:00
PropagationTime                 : 2.00:00:00
SignatureInceptionOffset        : 01:00:00
```

3. To review a summary of signing keys, use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


KeyId                                   KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                   -------         ------
---------    ---------  -------------   -----------------
1adabf48-5039-4e71-907c-99ff87f43564   KeySigningKey
RsaSha256          2048       True            True
08bf6bd0-b242-49d5-8674-eff9d338ea2f   ZoneSigningKey
RsaSha256          1024       True            True
```

4. You can also review detailed information for signing keys using Format-List (fl). See the

following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| fl *


KeyId                          : 1adabf48-5039-4e71-907c-
99ff87f43564
IsRolloverEnabled              : True
ActiveKey                      : {B4F136C0-4CF8-4CF9-BDBE-
73AC28780767}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
CurrentState                   : Active
DnsKeySignatureValidityPeriod  : 7.00:00:00
DSSignatureValidityPeriod      : 7.00:00:00
InitialRolloverOffset          : 00:00:00
KeyLength                      : 2048
KeyStorageProvider             : Microsoft Software Key
Storage Provider
KeyType                        : KeySigningKey
LastRolloverTime               :
NextKey                        :
NextRolloverAction             : Normal
NextRolloverTime               : 1/17/2016 3:56:40 PM
RolloverPeriod                 : 755.00:00:00
RolloverType                   : DoubleSignature
StandbyKey                     : {4F39AD5A-7FB7-43EF-BCD3-
F33494ABE523}
StoreKeysInAD                  : True
ZoneName                       : secure.contoso.com
ZoneSignatureValidityPeriod    : 10.00:00:00
PSComputerName                 :
CimClass                       :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties          : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
```

```
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                         : 08bf6bd0-b242-49d5-8674-
eff9d338ea2f
IsRolloverEnabled             : True
ActiveKey                     : {B426354D-9209-4CFD-9049-
6E503ADA14D0}
CryptoAlgorithm               : RsaSha256
CurrentRolloverStatus         : NotRolling
CurrentState                  : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                     : 1024
KeyStorageProvider            : Microsoft Software Key
Storage Provider
KeyType                       : ZoneSigningKey
LastRolloverTime              :
NextKey                       : {8265DEA2-A0EB-45B4-B1E8-
69F61777EB0E}
NextRolloverAction            : Normal
NextRolloverTime              : 3/23/2014 4:56:40 PM
RolloverPeriod                : 90.00:00:00
RolloverType                  : PrePublish
StandbyKey                    :
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties           :
```

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Perform an Emergency Key Revocation

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Perform an Emergency Key Revocation**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Perform an Emergency Key Revocation

	Task	Reference
☐	Review conceptual information about DNSSEC signing keys and key rollover.	Signing keys Trust Anchors
☐	Replace signing keys and re-sign the zone	Procedure: Replace Signing Keys
☐	Distribute trust anchors	Procedure: Import a Trust Point

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology
Appendix B: Windows PowerShell for DNS Server

Procedure: Import a Trust Point

Use the following procedures to import a trust point for a zone that has been signed with DNSSEC. To import a trust point, the DSSET or KEYSET records for a zone must be available in a file. You can create this file manually, or export it from the Key Master.

When you sign a zone on a DNS server running Windows Server 2008 R2 or later, the DNSKEY RRSET (also called the Keyset) and the DS RRSET (also called the dsset) are stored in text files on the Key Master for use as trust points. By default these files are available in the **%windir%\system32\dns** directory. To import a trust point on a DNS server, these files must be made accessible to the DNS server. For more information, see Procedure: Export a Trust Point.

For more information about trust points, see Trust Anchors.

Procedures are provided using the DNS Manager console and Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Choose DNS Manager or Windows PowerShell to perform this procedure:

- Import a trust point using DNS Manager
- Import a trust point using Windows PowerShell

Import a trust point using DNS Manager

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To import a trust point using DNS Manager

1. Open DNS Manager on the server where you wish to import the trust point, or connect to the DNS server from another location using DNS Manager.
2. In the console tree, right-click **Trust Points**, point to **Import**, and then click **DNSKEY** or click **DS**.

   If the trust point that you have made available is a KEYSET, then point to **DNSKEY**. If the trust point that is available is a DSSET, point to **DS**.
3. Click **Browse** or type the location of the KEYSET or DSSET file, and then click **OK**.

If a trust point for the zone did not previously exist, a new folder will be created in the DNS Manager console tree under Trust Points. If the trust point for this zone had previously been imported or added to the DNS server, the trust point will be updated.

Import a trust point using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To use the Windows PowerShell examples that are provided, replace server and zone names with the names you will use in your environment.

To import a trust point using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on the server where you wish to import the trust point.
2. To import a trust point, use the **Import-DnsServerTrustAnchor** cmdlet. See the following example.

   ```
   PS C:\> Import-DnsServerTrustAnchor -KeySetFile
   "\\File1\DNSKeys\keyset-secure.contoso.com"
   ```

   In this example, a DNSKEY trust anchor for the zone *secure.contoso.com* is imported on the local DNS server. To import a DS trust anchor, use the **DSSetFile** parameter in place of **KeySetFile** and specify a filename of *dsset-secure.contoso.com*.

To verify that the trust point was successfully imported, you can use the **Get-DnsServerTrustAnchor** cmdlet. See the following example:

```
PS C:\> Get-DnsServerTrustAnchor -Name secure.contoso.com


TrustAnchorName                TrustAnchorType      TrustAnchorState
TrustAnchorData
---------------                ---------------      ----------------     -------------
--
secure.contoso.com.            DNSKEY               Valid
[8000][DnsSec][RsaSha256][AwEAAY/kENeHkSmtd...
secure.contoso.com.            DNSKEY               Valid
[62111][DnsSec][RsaSha256][AwEAAbMK6im/lrLB...
```

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Replace Signing Keys

Use the following procedures to replace one or more signing keys using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Important
To remove a signing key immediately, the zone must be unsigned and then re-signed. In an emergency, such as when signing keys have been compromised, do not attempt to retire a signing key and perform a manual key rollover. Signing key rollover includes a waiting period so that validity periods can expire. It you wish to replace a signing key through retirement, see Procedure: Retire a Signing Key.

When you replace a signing key, the zone will immediately be re-signed.

You can use DNS Manager or Windows PowerShell to replace DNSSEC signing keys:

- Replace a signing key using DNS Manager
- Replace a signing key using Windows PowerShell

Replace a signing key using DNS Manager

You can use DNS Manager to review some signing key settings. However, Windows PowerShell provides additional settings for review.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replace a signing key using DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Unsign the Zone**.
3. Click **Next** and then click **Finish**.
4. Right-click the zone, point to **DNSSEC**, and then click **Sign the Zone**.
5. Click **Next**, choose **Customize zone signing parameters**, and then click **Next** until you get to the **Key Signing Key (KSK)** or **Zone Signing Key (ZSK)** page where you wish to replace a signing key.
6. Click the signing key you wish to replace and then click **Remove**.
7. Click **Add**, provide parameter values for a replacement signing key, and then click **OK**.
8. Click Next and provide any additional parameter values as needed, and continue clicking Next until you reach the **DNS Security Extensions (DNSSEC)** page indicating you have successfully configured zone signing parameters.
9. Click **Next**, verify that **The zone has been successfully signed** is displayed, and then click **Finish**.

After replacing signing keys, all trust anchors that have been distributed must also be replaced.

Replace a signing key using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To replace a signing key using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. First, you must unsign the zone. To unsign a zone, use the **Invoke-DnsServerZoneUnsign** cmdlet. See the following example.

   ```
   PS C:\> Invoke-DnsServerZoneUnsign -ZoneName
   secure.contoso.com -Force
   ```

   In this example, the zone *secure.contoso.com* is unsigned on the local DNS server. The **Force** parameter is used to skip confirmation.
3. Next, use the **Get-DnsServerSigningKey** cmdlet to retrieve signing key IDs. See the following example:

   ```
   PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com


   KeyId                                    KeyType
   CryptoAlgorithm    KeyLength  StoreKeysInAD
   IsRolloverEnabled
   -----                                    -------         ------
   ---------    ---------  -------------   -----------------
   488a4688-39bd-4d15-a2bb-56474d665e38   KeySigningKey
   RsaSha256          2048       True            True
   7d1f81e4-230e-4704-a88c-d4bc9c9b5ab9   ZoneSigningKey
   RsaSha256          1024       True            True
   ```

   In the previous example, there is ZSK and one KSK.
4. To replace a signing key, a new signing key must first be added. To add a signing key, use the **Add-DnsServerSigningKey** cmdlet. See the following example:

   ```
   Add-DnsServerSigningKey -ZoneName secure.contoso.com -Type
   KeySigningKey -CryptoAlgorithm RsaSha1NSec3 –KeyLength 3072
   ```
5. In this example, the **CryptoAlgorithm** and **KeyLength** parameters are set to a custom value. Parameters that are not specified will use default values.

6. To add a custom ZSK, use the **Add-DnsServerSigningKey** cmdlet. See the following example:

   ```
   Add-DnsServerSigningKey -ZoneName secure.contoso.com -Type
   ZoneSigningKey -CryptoAlgorithm RsaSha1NSec3
   ```

   In this example, the **CryptoAlgorithm** parameter is set to a custom value. Parameters that are not specified will use default values.

7. When you have completed adding new signing keys, remove the signing key that will be replaced with the **Remove-DnsServerSigningKey** cmdlet. See the following example:

   ```
   PS C:\> Remove-DnsServerSigningKey -ZoneName
   secure.contoso.com -KeyId 488a4688-39bd-4d15-a2bb-
   56474d665e38 -Force
   ```

   In this example, the **Force** parameter is used to skip confirmation.

8. When you have completed adding signing keys, sign the zone. To sign the zone, use the **Invoke-DnsServerZoneSign** cmdlet. See the following example:

   ```
   Invoke-DnsServerZoneSign -ZoneName secure.contoso.com -Force
   ```

   In this example, the **Force** parameter is used to skip confirmation and immediately sign the zone.

After replacing signing keys, all trust anchors that have been distributed must also be replaced.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Checklist: Perform a Manual Key Rollover

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Perform a Manual Key Rollover**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Perform a Manual Key Rollover

	Task	Reference
☐	Review conceptual information about DNSSEC signing keys and key rollover.	Signing keys Trust Anchors
☐	Initiate a manual key rollover.	Procedure: Perform a Manual Signing Key Rollover
☐	Verify that key rollover was successful.	Procedure: Review Key Rollover Status

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Perform a Manual Signing Key Rollover

Use the following procedures to perform a manual signing key rollover using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

Caution

When you perform a manual signing key (ZSK or KSK) rollover, trust anchors that have been distributed for the zone must also be manually updated. Automatic trust anchor updates (RFC 5011) are only supported for automatic key rollover.

Key rollover does not occur immediately:

- For rollover of ZSKs, the DNS server will wait for a minimum of 1 hour to allow any existing RRSIG resource records that are cached on resolving DNS servers to expire. ZSK rollover can require a longer period of time if the DNSKEY TTL or maximum TTL configured for other resource records in the zone is greater than 1 hour.
- For KSK rollover, the DNS server will wait for trust anchor updates. If the current rollover status is **KskWaitingForDSUpdate** and you know that the parent zone has been updated,

you can use the **Step-DnsServerSigningKeyrollover** cmdlet to force key rollover to continue.

You are not allowed to add new signing keys or reconfigure signing key parameters during this waiting period.

You can use DNS Manager or Windows PowerShell to manually rollover signing keys:

- Perform a manual signing key rollover using DNS Manager
- Perform a manual signing key rollover using Windows PowerShell

Perform a manual signing key rollover using DNS Manager

Use the following procedure to perform a manual key rollover using the DNS Manager console.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To perform a manual signing key rollover using DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. Click the **KSK** tab or click the **ZSK** tab, click the signing key in the list that you will rollover, and then click **Rollover**.
4. When you are prompted to confirm rollover on the selected key, click **Yes**.
5. Click **OK** to close DNSSEC properties of the zone.

Perform a manual signing key rollover using Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To perform a manual signing key rollover using Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. Use the **Get-DnsServerSigningKey** cmdlet to retrieve signing key IDs. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
```

```
KeyId                                    KeyType
CryptoAlgorithm    KeyLength  StoreKeysInAD
IsRolloverEnabled
-----                                    -------         ------
---------    ---------  -------------   -----------------

96e2236a-4a4e-4409-afe7-c3aa8ffc254e   KeySigningKey
RsaSha256          2048       True            True
925795c2-4772-48d8-b0f0-0d79c8a7e5c7   ZoneSigningKey
RsaSha256          1024       True            True
```

3. To roll over a signing key, use the **Invoke-DnsServerSigningKeyRollover** cmdlet. See the following example:

```
PS C:\> Invoke-DnsServerSigningKeyRollover -ZoneName
secure.contoso.com -KeyID 925795c2-4772-48d8-b0f0-
0d79c8a7e5c7 -Force
```

   In the previous example, the **Force** parameter is used to skip confirmation.

To review status of the key rollover, you can use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com -KeyId fd708ec9-e309-
45ea-9cb8-f47992350c7b | fl *



KeyId                          : 925795c2-4772-48d8-b0f0-0d79c8a7e5c7
IsRolloverEnabled              : True
ActiveKey                      : {EA067E8E-7A43-4C06-BF5A-385A9270BCB4}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : ZskWaitingForDnsKeyTtl
CurrentState                   : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod      : 7.00:00:00
InitialRolloverOffset          : 00:00:00
KeyLength                      : 1024
KeyStorageProvider             : Microsoft Software Key Storage Provider
KeyType                        : ZoneSigningKey
LastRolloverTime               :
NextKey                        : {8B36FED4-93A3-4E7F-B894-54B9EDAC4493}
```

```
NextRolloverAction            : Normal
NextRolloverTime              : 1/29/2014 3:00:48 PM
RolloverPeriod                : 90.00:00:00
RolloverType                  : PrePublish
StandbyKey                    : {5F7AE18D-BDC3-4BAC-9784-4D6B1AAC45F6}
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      : root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm, CurrentRolloverStatus,
CurrentState...}
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties
```

In the previous example, the value of **CurrentRolloverStatus** is **ZskWaitingForDnsKeyTtl**.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Review Key Rollover Status

Use the following procedures to display and review rollover status of signing keys using the DNS Manager console or Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist.

You can use DNS Manager or Windows PowerShell to review rollover status:

- Review key rollover status in DNS Manager
- Review key rollover status in Windows PowerShell

Review key rollover status in DNS Manager

You can use DNS Manager to review basic key rollover status. However, Windows PowerShell provides additional settings for review.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To review key rollover status in DNS Manager

1. Open DNS Manager on a primary, authoritative DNS server, or connect to a primary, authoritative DNS server with DNS Manager.
2. Right-click a DNSSEC-signed zone, point to **DNSSEC**, and then click **Properties**.
3. To review KSK status, click the **KSK** tab, click a signing key under **Key signing keys (KSKs)**, and then review information under **Rollover status**.
4. To review ZSK settings, click the **ZSK** tab, click a signing key under **Zone signing keys (ZSKs)**, and then review information under **Rollover status**.
5. Click **Cancel** when finished.

Review key rollover status in Windows PowerShell

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

In the Windows PowerShell commands that are used, replace *secure.contoso.com* with the name of the DNSSEC-signed zone you wish to review.

To review key rollover status in Windows PowerShell

1. Open an elevated Windows PowerShell prompt on a primary, authoritative DNS server.
2. To review signing key status, use the **Get-DnsServerSigningKey** cmdlet. See the following example:

```
PS C:\> Get-DnsServerSigningKey -ZoneName secure.contoso.com
| fl *


KeyId                          : 0316fff3-03a8-451d-a4ad-
db2f93e4070a
IsRolloverEnabled              : True
ActiveKey                      : {8F7CB3C7-A9A1-496A-96E6-
D6D845D2F0F1}
CryptoAlgorithm                : RsaSha256
CurrentRolloverStatus          : NotRolling
CurrentState                   : Active
```

```
DnsKeySignatureValidityPeriod : 7.00:00:00
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                     : 2048
KeyStorageProvider            : Microsoft Software Key
Storage Provider
KeyType                       : KeySigningKey
LastRolloverTime              :
NextKey                       :
NextRolloverAction            : Normal
NextRolloverTime              : 2/10/2016 2:27:53 PM
RolloverPeriod                : 755.00:00:00
RolloverType                  : DoubleSignature
StandbyKey                    : {1F4FBEBD-D66F-485C-9076-
187DD40B1993}
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties

KeyId                         : 60b46b4d-e1cb-433a-ac33-
5d46687e860e
IsRolloverEnabled             : True
ActiveKey                     : {93BA0FC4-3931-4DBF-A205-
F4EFB24C4E4E}
CryptoAlgorithm               : RsaSha256
CurrentRolloverStatus         : NotRolling
CurrentState                  : Active
DnsKeySignatureValidityPeriod : 7.00:00:00
```

```
DSSignatureValidityPeriod     : 7.00:00:00
InitialRolloverOffset         : 00:00:00
KeyLength                     : 1024
KeyStorageProvider            : Microsoft Software Key
Storage Provider
KeyType                       : ZoneSigningKey
LastRolloverTime              :
NextKey                       : {D1FDEAF4-F9AB-4D59-AF5A-
A33856B2B0EE}
NextRolloverAction            : Normal
NextRolloverTime              : 4/16/2014 3:27:53 PM
RolloverPeriod                : 90.00:00:00
RolloverType                  : PrePublish
StandbyKey                    :
StoreKeysInAD                 : True
ZoneName                      : secure.contoso.com
ZoneSignatureValidityPeriod   : 10.00:00:00
PSComputerName                :
CimClass                      :
root/Microsoft/Windows/DNS:DnsServerSigningKey
CimInstanceProperties         : {ActiveKey, CryptoAlgorithm,
CurrentRolloverStatus, CurrentState...}
CimSystemProperties           :
Microsoft.Management.Infrastructure.CimSystemProperties
```

In the previous example, the Format-List (fl) parameter is used to display extended signing key status information.

See also

[DNSSEC in Windows]

[DNSSEC Deployment Planning]

Checklist: Secure Zone Transfers

Checklist: Deploy DNSSEC > Checklist: Sign a Zone > Checklist: Distribute Trust Anchors > Checklist: Deploy DNSSEC Policies to DNS Clients > **Checklist: Secure Zone Transfers**

This parent checklist includes links to procedures that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist.

Note

Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Secure Zone Transfers

	Task	Reference
☐	Review concepts about DNSSEC.	Overview of DNSSEC DNS Zones DNS Servers
☐	Deploy certificates to secure zone transfers	Procedure: Deploy Certificates for DNS Server Authentication
☐	Deploy policies to secure zone transfers	Procedure: Deploy IPsec Policy to DNS Servers

See also

Overview of DNSSEC

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

Appendix B: Windows PowerShell for DNS Server

Procedure: Deploy Certificates for DNS Server Authentication

Use the following procedures to configure and publish server certificates for IPsec authentication, and to enable certificate auto-enrollment on your DNS servers.

Important

When choosing a certification authority (CA) to issue certificates for IPsec authentication, it is important that you consider whether or not other certificates will be issued by the same CA. For more information, see Certificate Selection.

You can deploy certificates to DNS servers through one of the following mechanisms:

- **Domain Controllers organizational unit (OU)**: If the DNS servers in your domain are Active Directory-integrated, you can deploy IPsec policy settings using the Domain Controllers OU. This option is recommended to make configuration and deployment easier.
- **DNS Server OU or security group**: If you have DNS servers that are not domain controllers, then consider creating a separate OU or a security group with the machine accounts of your DNS servers.
- **Local firewall configuration**: Use this option if you have DNS servers that are not domain members or if you have a small number of DNS servers that you want to configure locally.

Use the following procedures to deploy DNS Server certificates to the Domain Controllers OU using auto-enrollment. If you wish to deploy DNS Server certificates manually or issue certificates to a different group of computers, modify the permission settings on the security tab in certificate template properties.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Configure certificate templates

A certificate template must be created to provide authentication of DNSSEC protected communications. This certificate template will be configured with the Domain Name System (DNS) Server Trust application policy.

Important

To configure and publish a certificate template for DNSSEC, the computer must be running Windows Server Enterprise Edition with the Active Directory Certificate Services (AD CS) role installed and configured as an Enterprise certification authority (CA).

To configure a certificate template

1. Click **Start**, click **Run**, type **certtmpl.msc**, and then press ENTER.
2. In the details pane, right-click **RAS and IAS Server**, and then click **Duplicate Template**.
3. Select **Windows Server 2003 Enterprise**, and then click **OK**.
4. Under **Template display name**, type **DNS Server**, and then select the **Publish certificate in Active Directory** check box.
5. Click the **Request Handling** tab, and select the **Allow private key to be exported** check box.
6. Click the **Extensions** tab, and then click **Application Policies**.

7. If the certification authority is running Windows Server® 2008 R2:
   a. Click **Edit**, click **Add**, click **Domain Name System (DNS) Server Trust**, and then click **OK**.
8. If the certification authority is running Windows Server® 2008 or an earlier operating system:
   a. Click **Edit**, click **Add**, and then click **New**.
   b. In the **New Application Policy** dialog box, under **Name**, type **Domain Name System (DNS) Server Trust**, and under **Object identifier**, type **1.3.6.1.4.1.311.64.1.1**, and then click **OK** twice.

   Important
   Ensure that you enter the object identifier exactly as specified. This is required for servers to successfully negotiate IPsec with clients.
9. Click **Add**, click **IP security IKE intermediate**, and then click **OK**.
10. In **Edit Application Policies Extension**, verify that the following are listed under **Application policies**.
    - Client Authentication
    - Domain Name System (DNS) Server Trust
    - IP security IKE intermediate
    - Server Authentication
11. Click **OK**, click **Key Usage**, and then click the **Edit**.
12. Under **Signature**, select the **Digital signature** and **Signature is proof or origin (nonrepudiation)** check boxes.
13. Under **Encryption**, choose **Allow key exchange only with key encryption (key encipherment)**, select the **Allow encryption of user data** and **Make this extension critical** check boxes, and then click **OK**.
14. Click the **Security** tab, click **RAS and IAS Servers**, and then click **Remove**.
15. Click Add, type Domain Controllers, and then click OK. ‘Alternatively, type the name of the OU or security group you will use to deploy DNSSEC server policy.
16. Select the Allow permission for **Enroll** and **Autoenroll** for the **Domain Controllers** security group or another security group you will use to deploy DNSSEC server policy.
17. Click **OK** and then close the certificate templates console.

Publish certificate templates

Use the following procedure to allow the CA to issue the new certificate template.

To publish certificate templates

1. Click **Start**, click **Run**, type **certsrv.msc**, and then press ENTER.
2. In the console tree, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.

3. Click **DNS Server**, and then click **OK**.
4. In the console tree, click **Certificate Templates**, and in the details pane under **Name**, verify that **DNS Server** is displayed.
5. Close the Certification Authority console.

Enable certificate auto-enrollment

If you will use auto-enrollment to issue certificates for IPsec authentication, you must enable auto-enrollment in Group Policy. You do not need to perform this procedure to enroll certificates manually on your DNS servers.

To enable certificate auto-enrollment

1. On a domain controller or a computer with the Group Policy Management feature installed, click **Start**, click **Run**, type **gpme.msc**, and then press ENTER.
2. In the Browse for a Group Policy Object dialog box, double-click **Domain Controllers.<domain.com>**, click **Default Domain Controllers Policy**, and then click **OK**. The Group Policy Management Editor will open.

   Important
   If you are using a different GPO to manage DNSSEC settings for your DNS servers, then enable certificate auto-enrollment in this GPO instead.
3. In the console tree, open **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.
4. In the details pane, double-click **Certificate Services Client – Auto-Enrollment**.
5. In the Certificate Services Client – Auto-Enrollment Properties dialog box, next to **Configuration Model** select **Enabled** from the drop-down menu, select the check boxes next to **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**, and then click **OK**.
6. Close the Group Policy Management Editor.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Procedure: Deploy IPsec Policy to DNS Servers

Use the following procedure to configure IP Security (IPsec) rules for DNS servers in your organization that will provide DNS resolution for client computers. IPsec rules are configured to request authentication for all DNS queries.

You can deploy IPsec rules through one of the following mechanisms:

- **Domain Controllers organizational unit (OU):** If the DNS servers in your domain are Active Directory-integrated, you can deploy IPsec policy settings using the Domain Controllers OU. This option is recommended to make configuration and deployment easier.
- **DNS Server OU or security group:** If you have DNS servers that are not domain controllers, then consider creating a separate OU or a security group with the computer accounts of your DNS servers.
- **Local firewall configuration:** Use this option if you have DNS servers that are not domain members or if you have a small number of DNS servers that you want to configure locally.

Use the following procedure to deploy IPsec policy to the Domain Controllers OU. If you wish to deploy IPsec policy to a different group of computers, use a different OU or create a security group and use security group filtering to apply IPsec settings to your DNSSEC Group Policy object (GPO).

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Configure IPsec policy

- To configure IPsec policy using DNS Manager
- To configure IPsec policy using the command line

In the following procedure, IPsec policy is deployed on all domain controllers because it is assumed that domain controllers are also DNS servers. To deploy this policy on DNS servers that are not domain controllers, create and use a custom OU or security group. To deploy this policy on computers that are not domain members, use Local Group Policy to perform the following procedures.

Note

Complete the following procedure twice. First create a rule for UDP connections, then create a rule for TCP connections.

To configure IPsec policy using DNS Manager

1. On a domain controller or a computer with the Group Policy Management feature installed, click **Start**, click **Run**, type **gpme.msc**, and then press ENTER.
2. In the Browse for a Group Policy Object dialog box, double-click **Domain**

Controllers.<domain.com>, click **Default Domain Controllers Policy**, and then click **OK**. The Group Policy Management Editor will open.

3. In the console tree, open **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security - LDAP**.
4. Right-click **Connection Security Rules**, and then click **New Rule**. The New Connection Security Rule Wizard will open.
5. On the **Rule Type** page, choose **Custom**, and then click **Next**.
6. On the **Endpoints** page, choose **Any IP address** for endpoint 1 and **Any IP Address** for endpoint 2, and then click **Next**.
7. On the **Requirements** page, choose **Request authentication for inbound and outbound connections**, and then click **Next**.
8. On the **Authentication Method** page, choose **Advanced**, and then click **Customize**.
9. In the **Customize Advanced Authentication Methods** dialog box, under **First authentication**, click **Add**.
10. In the **Add First Authentication Method** dialog box, choose **Computer certificate from this certification authority (CA)**. Verify that **Signing algorithm** is **RSA** and **Certificate store type** corresponds to the type of CA you are using, either **Root CA** or **Intermediate CA**. Click **Browse**, and select the name of the CA that you used in the previous procedure to create and deploy a certificate for DNS server authentication, and then click **OK**.

    Important
    If multiple certificates from the same CA are present on the DNS server, IPsec authentication might fail due to an incorrect certificate being chosen. For more information, see Certificate Selection.
11. Click **OK** to close the **Add First Authentication Method** dialog box, click **OK** in **Customize Advanced Authentication Methods**, and then click **Next**.
12. On the **Protocol and Ports** page, next to **Protocol type**, choose **UDP**. Next to **Endpoint 1 port**, choose **Specific ports** and then type **53**. Next to **Endpoint 2 port** choose **All ports**, and then click **Next**.
13. On the **Profile** page, verify that the **Domain**, **Private** and **Public** check boxes are selected, and then click **Next**.
14. Type a name and description for the rule. Use a name that will be easy to recognize, for example, **DNSSEC UDP**.
15. Click **Finish** to create the rule.

Next, create an identical rule for DNS TCP connections by repeating this procedure and using TCP as the protocol type. You can also create a new rule using an existing rule as a template. To duplicate a rule, right-click an existing rule, click **Copy**, right-click inside the console details pane, and then click **Paste**. Edit the duplicate rule to provide a unique name and settings.

To configure IPsec policy using the command line

1. Open an elevated command prompt.
2. Enter the following command twice:

```
netsh advfirewall consec add rule name="DNSSEC UDP"
endpoint1=any endpoint2=any action=requestinrequestout
port1=53 port2=any protocol=<protocol>
auth1=computerkerb,computercert auth1ca=<CaName>
```

   The first time you enter this command, replace <protocol> with **UDP**, and replace <CaName> with the name of the CA being used. The second time you enter the command, use a different rule name such as ”DNSSEC TCP”, replace <protocol> with **TCP** and replace <CaName> with the name of the CA being used. See the following examples.

Example UDP rule

```
netsh advfirewall consec add rule name="DNSSEC UDP" endpoint1=any endpoint2=any
action=requestinrequestout port1=53 port2=any protocol=UDP
auth1=computerkerb,computercert auth1ca=”DC=com, DC=woodgrovebank, CN=woodgrovebank-
DC1-CA”
```

Example TCP rule

```
netsh advfirewall consec add rule name="DNSSEC TCP" endpoint1=any endpoint2=any
action=requestinrequestout port1=53 port2=any protocol=TCP
auth1=computerkerb,computercert auth1ca=” DC=com, DC=woodgrovebank, CN=woodgrovebank-
DC1-CA”
```

Use the following command to verify that rules were created successfully:

```
netsh advfirewall consec show rule name=all type=dynamic
```

Certificate Selection

The following options are available to ensure that the correct certificate on a DNS server is selected for IPsec authentication. For information about deploying this certificate, see Procedure: Deploy Certificates for DNS Server Authentication.

- Use a different CA to issue DNS server certificates than the one used to issue other certificates. To accomplish this, install Active Directory Certificate Services (AD CS) on a domain controller or member server and use this CA only for issuing DNS Server authentication certificates.
- If you have deployed Network Access Protection (NAP) on your network, you can add the **Domain Name System (DNS) Server Trust**, **IP security IKE intermediate**, and **Server Authentication** application policies to NAP exemption certificates that are provisioned on DNS servers. To use a NAP exemption certificate with DNS Server authentication, choose the **Computer health certificate from this certification authority (CA)** option instead of the **Computer certificate from this certification authority (CA)**.

- If you have not deployed NAP, you can still add the **System Health Authentication** application policy to the certificate that you use for DNS Server authentication and then configure IPsec policy to require a computer health certificate. You should only use this method if you must use the same CA to issue multiple certificates to your DNS servers.

See also

DNSSEC in Windows

DNSSEC Deployment Planning

Appendix A: DNSSEC Terminology

The following table displays a list of DNSSEC-related terms used in this guide.

Term	Definition	Reference
Authenticated Data (AD) bit	The AD (Authenticated Data) bit indicates in a response that all data included in the answer and authority sections of the response have been authenticated by the DNS server according to the policies of that server.	RFC 4035 [7] Section 3
Authentication chain	A chain of signed and validated DNS records starting from a preconfigured trust anchor to some child zone in the DNS tree.	RFC 4035 [7] Section 5
Delegation Signer (DS) record	A DNSSEC record type used to secure a delegation.	RFC 4034 [6] Section 5
DNS Extension (EDNS0)	A DNS record that carries extended DNS header information, such as the DO bit and maximum UDP packet size.	RFC 2671 [3]
DNS Security Extensions (DNSSEC)	Extensions to the DNS service that provide mechanisms for signing and securely resolving DNS data.	RFCs 4033 [5], 4034 [6], and 4035 [7]
DNSKEY record	A DNSSEC record type used to	RFC 4034 [6] Section 2

	store a public key.	
DNSSEC OK (DO) bit	A bit in the EDNS0 portion of a DNS request that signals that the client is DNSSEC-aware.	RFC 4035 [7] Section 3.2.1
Double signature	Double signature is a key rollover method where a future key and signatures generated using it are published into the DNS zone before the existing key and its signatures are deleted, thus leaving two sets of keys and associated signatures in the zone at any given time.	RFC 4641
Island of Security	A signed zone that does not have an authentication chain from its delegating parent zone.	RFC 4033 [5] Section 2
Key Master	The Key Master is the DNS server that is responsible for key generation and management, and acts as the server that the administrator interacts with to set up DNSSEC.	N/A
Key Signing Key (KSK)	The KSK is an authentication key that corresponds to a private key used to sign one or more other signing keys for a given zone. Typically, the private key corresponding to a KSK will sign a ZSK, which in turn has a corresponding private key that will sign other zone data. Local policy may require that the ZSK be changed frequently, while the KSK may have a longer validity period in order to provide a more stable secure entry point into the zone. Designating an authentication key as a KSK is purely an operational issue: DNSSEC	RFC 4033 section 2

	validation does not distinguish between KSKs and other DNSSEC authentication keys, and it is possible to use a single key as both a KSK and a ZSK.	
Next Secure (NSEC) record	A DNSSEC record type used to prove non-existence of a DNS name.	RFC 4034 [6] Section 4
Next Secure 3 (NSEC3)	The NSEC3 resource record provides hashed authenticated denial of existence for DNS resource record sets.	RFC 5155 Section 3
Non-validating security-aware stub resolver	A security-aware stub resolver that trusts one or more security-aware DNS servers to perform DNSSEC validation on its behalf.	RFC 4033 [5] Section 2
Offline zone signing	The process of signing or re-signing a file that contains the DNS data for a zone on a computer that is not an authoritative DNS server for the zone. Ideally the computer on which zone signing is performed is not network-connected. The signed zone file created by this process can then be transferred to an authoritative DNS server and loaded.	RFC 4035 [7] Section 2
Online zone signing	The process of signing or incrementally re-signing a zone while it is loaded by the DNS server. This requires that private keys be accessible to the DNS server process, and therefore is less secure than offline zone signing.	RFC 4035 [7] Section 2
Pre-publication	Pre-publication is a key rollover method where a key for future use is added and published in	RFC 4641

	the DNS zone without generating any signatures.	
Resource Record Set (RRSet)	Each DNS Resource Record (RR) has a label, class, type, and data. It is meaningless for two records to ever have label, class, type and data all equal - servers should suppress such duplicates if encountered. It is however possible for most record types to exist with the same label, class and type, but with different data. Such a group of records is hereby defined to be a Resource Record Set (RRSet).	RFC 2181
Resource Record Signature (RRSIG) record	A DNSSEC record type used to hold a signature which covers a set of DNS records for a particular name and type.	See RFC 4034 [6] Section 3
Secure entry point (SEP) key	SEP keys are a subset of public keys within the DNSKEY RRset. A SEP key is used either to generate a DS RR or is distributed to resolvers that use the key as a trust anchor.	RFC 3757
Security-aware DNS server	A DNS server that understands the DNS security extensions defined in RFCs 4033 [5], 4034 [6], and 4035 [7]. In particular, a security-aware DNS server is an entity that receives DNS queries, sends DNS responses, supports the EDNS0 [3] message size extension and the DO bit, and supports the DNSSEC record types and message header bits.	RFCs 4033 [5], 4034 [6], and 4035 [7]
Security-aware resolver	A resolver that understands the DNS security extensions defined in RFCs 4033 [5], 4034	RFC 4033 [5] Section 2

	[6], and 4035 [7]. In particular, a security-aware resolver is an entity that sends DNS queries, receives DNS responses, supports the EDNS0 [3] message size extension and the DO bit, and supports the DNSSEC record types and message header bits.	
Signed zone	A zone whose records are signed as defined by RFC 4035 [7] Section 2. A signed zone can contain DNSKEY, NSEC, NSEC3, NSEC3PARAM, RRSIG, and DS records. These records allow DNS data to be validated by resolvers.	RFC 4035 [7] Section 2
Trust anchor	A preconfigured public key associated with a particular zone. A trust anchor allows a DNS resolver to validate signed DNSSEC records for that zone and to build authentication chains to child zones	RFC 4033 [5] Section 2
Unsigned zone	Any DNS zone that has not been signed as defined by RFC 4035 [7] Section 2.	RFC 4035 [7] Section 2
Zone Signing Key (ZSK)	An authentication key that corresponds to a private key used to sign a zone. Typically, a zone signing key will be part of the same DNSKEY RRset as the key signing key whose corresponding private key signs this DNSKEY RRset, but the zone signing key is used for a slightly different purpose and may differ from the key signing key in other ways, such as validity lifetime. Designating an authentication key as a zone	RFC 4033 section 2

	signing key is purely an operational issue; DNSSEC validation does not distinguish between zone signing keys and other DNSSEC authentication keys, and it is possible to use a single key as both a key signing key and a zone signing key.	

Appendix B: Windows PowerShell for DNS Server

Windows PowerShell support for DNSSEC was added in Windows Server 2012 and extended in Windows Server 2012 R2. For a list of all Windows PowerShell cmdlets for DNS server, see Domain Name System (DNS) Server Cmdlets in Windows PowerShell. Several DNS server cmdlets are specifically used with DNSSEC-signed zones, including the following:

Table 1: DNSSEC related Windows PowerShell cmdlets in Windows Server 2012 and Windows Server 2012 R2.

Cmdlet	Description
Add-DnsServerResourceRecordDnsKey	Adds a DNSKEY resource record to a zone.
Add-DnsServerResourceRecordDS	Adds a DS resource record to a zone.
Add-DnsServerSigningKey	Adds a KSK or ZSK to a signed zone.
Add-DnsServerTrustAnchor	Adds a trust anchor to a DNS server.
Disable-DnsServerSigningKeyRollover	Disables key rollover on a specified key.
Enable-DnsServerSigningKeyRollover	Enables rollover on a specified key.
Export-DnsServerDnsSecPublicKey	Exports DS and DNSKEY information for a DNSSEC–signed zone.
Get-DnsServerDnsSecZoneSetting	Gets DNSSEC settings for a zone.
Get-DnsServerSigningKey	Gets zone signing keys.
Get-DnsServerTrustAnchor	Gets trust anchors on a DNS server.
Get-DnsServerTrustPoint	Gets trust points on a DNS server.
Import-DnsServerTrustAnchor	Imports a trust anchor for a DNS server.

Cmdlet	Description
Invoke-DnsServerSigningKeyRollover	Initiates rollover of signing keys for the zone.
Invoke-DnsServerZoneSign	Initiates zone signing.
Invoke-DnsServerZoneUnsign	Initiates zone unsigning.
Remove-DnsServerSigningKey	Removes signing keys from a zone.
Remove-DnsServerTrustAnchor	Removes a trust anchor from a DNS server.
Reset-DnsServerZoneKeyMasterRole	Transfers the role of Key Master for a DNS zone.
Set-DnsServerDnsSecZoneSetting	Changes DNSSEC settings for a zone.
Set-DnsServerSigningKey	Changes settings of a signing key.
Show-DnsServerKeyStorageProvider	Returns a list of key storage providers on a DNS server.
Step-DnsServerSigningKeyRollover	Forces rollover for a KSK that is waiting for a parent DS update.
Test-DnsServerDnsSecZoneSetting	Validates DNSSEC settings for a zone.
Update-DnsServerTrustPoint	Updates all trust points on a DNS server.

The previous table is not a comprehensive list of all DNS server cmdlets that can be used with signed zones. Other generic DNS server cmdlets can also be used to create, view, and modify DNSSEC-signed resource records. For example, Get-DnsServerResourceRecord will display resource records from both signed and unsigned zones. A list of all DNS server cmdlets is provided below.

Table 2: DNS server Windows PowerShell cmdlets in Windows Server 2012 and Windows Server 2012 R2 categorized by noun.

Noun	Verb	Description	Cmdlet / Reference
	Get Set Test	Retrieves a DNS server configuration. Overwrites a DNS server configuration. Tests that a specified computer is a functioning	Get-DnsServer Set-DnsServer Test-DnsServer

Noun	Verb	Description	Cmdlet / Reference
		DNS server.	
Cache	Clear Get Set Show	Clears resource records from a cache on the DNS server. Gets DNS server cache settings. Modifies cache settings for a DNS server. Shows the records in a DNS Server Cache.	Clear-DnsServerCache Get-DnsServerCache Set-DnsServerCache Show-DnsServerCache
ConditionalForwarderZone	Add Set	Adds a conditional forwarder to a DNS server. Changes settings for a DNS conditional forwarder.	Add-DnsServerConditionalForwarderZone Set-DnsServerConditionalForwarderZone
Diagnostics	Get Set	Retrieves DNS event logging details. Sets debugging and logging parameters.	Get-DnsServerDiagnostics Set-DnsServerDiagnostics
DirectoryPartition	Add Get Register Remove	Creates a DNS application directory	Add-DnsServerDirectoryPartition Get-DnsServerDirectoryPartition Register-DnsServerDirectoryPartition Remove-DnsServerDirectoryPartition

Noun	Verb	Description	Cmdlet / Reference
	Unregister	partition. Gets a DNS application directory partition. Registers a DNS server in a DNS application directory partition. Removes a DNS application directory partition. Deregisters a DNS server from a DNS application directory partition.	Unregister-DnsServerDirectoryPartition
DnsSecPublicKey	Export	Exports DS and DNSKEY information for a DNSSEC–signed zone.	Export-DnsServerDnsSecPublicKey
DnsSecZoneSetting	Get Set Test	Gets DNSSEC settings for a zone. Changes settings for DNSSEC for a zone. Validates DNSSEC settings for a zone.	Get-DnsServerDnsSecZoneSetting Set-DnsServerDnsSecZoneSetting Test-DnsServerDnsSecZoneSetting

Noun	Verb	Description	Cmdlet / Reference
DsSetting	Get Set	Retrieves DNS Server Active Directory settings. Modifies DNS Active Directory settings.	Get-DnsServerDsSetting Set-DnsServerDsSetting
EDns	Get Set	Gets EDNS configuration settings on a DNS sever. Changes EDNS settings on a DNS server.	Get-DnsServerEDns Set-DnsServerEDns
Forwarder	Add Get Remove Set	Adds server level forwarders to a DNS server. Gets forwarder configuration settings on a DNS server. Removes server level forwarders from a DNS server. Changes forwarder settings on a DNS server.	Add-DnsServerForwarder Get-DnsServerForwarder Remove-DnsServerForwarder Set-DnsServerForwarder
GlobalNameZone	Get Set	Retrieves DNS server GlobalName zone	Get-DnsServerGlobalNameZone Set-DnsServerGlobalNameZone

Noun	Verb	Description	Cmdlet / Reference
		configuration details. Changes configuration settings for a GlobalNames zone.	
GlobalQueryBlockList	Get Set	Gets a global query block list. Changes settings of a global query block list.	Get-DnsServerGlobalQueryBlockList Set-DnsServerGlobalQueryBlockList
KeyStorageProvider	Show	Returns a list of key storage providers.	Show-DnsServerKeyStorageProvider
PrimaryZone	Add ConvertTo Restore Set	Adds a primary zone to a DNS server. Converts a zone to a DNS primary zone. Restores primary DNS zone contents from Active Directory or from a file. Changes settings for a DNS primary zone.	Add-DnsServerPrimaryZone ConvertTo-DnsServerPrimaryZone Restore-DnsServerPrimaryZone Set-DnsServerPrimaryZone
Recursion	Get Set	Retrieves DNS server recursion settings.	Get-DnsServerRecursion Set-DnsServerRecursion

Noun	Verb	Description	Cmdlet / Reference
		Modifies recursion settings for a DNS server.	
ResourceRecord	Add Get Remove Set	Adds a resource record of a specified type to a specified DNS zone. Gets resource records from a specified DNS zone. Removes specified DNS server resource records from a zone. Changes a resource record in a DNS zone.	Add-DnsServerResourceRecord Get-DnsServerResourceRecord Remove-DnsServerResourceRecord Set-DnsServerResourceRecord
ResourceRecordA	Add	Adds a type A resource record to a DNS zone.	Add-DnsServerResourceRecordA
ResourceRecordAAAA	Add	Adds a type AAAA resource record to a DNS server.	Add-DnsServerResourceRecordAAAA
ResourceRecordAging	Set	Begins aging of resource records in a specified DNS zone.	Set-DnsServerResourceRecordAging

Noun	Verb	Description	Cmdlet / Reference
ResourceRecordCName	Add	Adds a type CNAME resource record to a DNS zone.	Add-DnsServerResourceRecordCName
ResourceRecordDnsKey	Add	Adds a type DNSKEY resource record to a DNS zone.	Add-DnsServerResourceRecordDnsKey
ResourceRecordDS	Add Import	Adds a type DS resource record to a DNS zone. Imports DS resource record information from a file.	Add-DnsServerResourceRecordDS Import-DnsServerResourceRecordDS
ResourceRecordMX	Add	Adds an MX resource record to a DNS zone.	Add-DnsServerResourceRecordMX
ResourceRecordPtr	Add	Adds a type PTR resource record to a DNS zone.	Add-DnsServerResourceRecordPtr
RootHint	Add Get Import Remove Set	Adds root hints on a DNS server. Gets root hints on a DNS server. Copies root hints from a DNS server. Removes root hints from a	Add-DnsServerRootHint Get-DnsServerRootHint Import-DnsServerRootHint Remove-DnsServerRootHint Set-DnsServerRootHint

Noun	Verb	Description	Cmdlet / Reference
		DNS server. Replaces a list of root hints.	
Scavenging	Get Set Start	Gets DNS aging and scavenging settings. Changes DNS server scavenging settings. Notifies a DNS server to attempt a search for stale resource records.	Get-DnsServerScavenging Set-DnsServerScavenging Start-DnsServerScavenging
SecondaryZone	Add ConvertTo Restore Set	Adds a DNS server secondary zone. Converts a primary zone or stub zone to a secondary zone. Restores secondary zone information from its source. Change settings for a DNS secondary zone.	Add-DnsServerSecondaryZone ConvertTo-DnsServerSecondaryZone Restore-DnsServerSecondaryZone Set-DnsServerSecondaryZone
Setting	Get	Retrieves	Get-DnsServerSetting

Noun	Verb	Description	Cmdlet / Reference
	Set	DNS server settings. Modifies DNS server settings.	Set-DnsServerSetting
SigningKey	Add Get Remove Set	Adds a KSK or ZSK to a signed zone. Gets zone signing keys. Removes signing keys. Changes settings of a signing key.	Add-DnsServerSigningKey Get-DnsServerSigningKey Remove-DnsServerSigningKey Set-DnsServerSigningKey
SigningKeyRollover	Disable Enable Invoke Step	Disables key rollover on an input key. Enables rollover on the input key. Initiates rollover of signing keys for the zone. Rolls over a KSK that is waiting for a parent DS update.	Disable-DnsServerSigningKeyRollover Enable-DnsServerSigningKeyRollover Invoke-DnsServerSigningKeyRollover Step-DnsServerSigningKeyRollover
Statistics	Clear Get	Clears all DNS server statistics or statistics for zones. Retrieves DNS server statistics or	Clear-DnsServerStatistics Get-DnsServerStatistics

Noun	Verb	Description	Cmdlet / Reference
		statistics for zones.	
StubZone	Add Set	Adds a DNS stub zone. Changes settings for a DNS server stub zone.	Add-DnsServerStubZone Set-DnsServerStubZone
TrustAnchor	Add Get Import Remove	Adds a trust anchor to a DNS server. Gets trust anchors on a DNS server. Imports a trust anchor for a DNS server. Removes a trust anchor from a DNS server.	Add-DnsServerTrustAnchor Get-DnsServerTrustAnchor Import-DnsServerTrustAnchor Remove-DnsServerTrustAnchor
TrustPoint	Get Update	Gets trust points on a DNS server. Updates all trust points in a DNS trust anchor zone.	Get-DnsServerTrustPoint Update-DnsServerTrustPoint
Zone	Export Get Remove Resume Suspend Sync	Exports contents of a zone to a file. Gets details of DNS zones on a DNS server. Removes a zone from a DNS server. Resumes	Export-DnsServerZone Get-DnsServerZone Remove-DnsServerZone Resume-DnsServerZone Suspend-DnsServerZone Sync-DnsServerZone

Noun	Verb	Description	Cmdlet / Reference
		name resolution on a suspended zone. Suspends a zone on a DNS server. Checks the DNS server memory for changes, and writes them to persistent storage.	
ZoneAging	Get Set	Gets DNS aging settings for a zone. Configures DNS aging settings for a zone.	Get-DnsServerZoneAging Set-DnsServerZoneAging
ZoneDelegation	Add Get Remove Set	Adds a new delegated DNS zone to an existing zone. Gets the zone delegations of a DNS server zone. Removes a name server or delegation from a DNS zone. Changes delegation settings for a child zone.	Add-DnsServerZoneDelegation Get-DnsServerZoneDelegation Remove-DnsServerZoneDelegation Set-DnsServerZoneDelegation

Noun	Verb	Description	Cmdlet / Reference
ZoneKeyMasterRole	Reset	Transfers the Key Master role for a DNS zone.	Reset-DnsServerZoneKeyMasterRole
ZoneSign	Invoke	Signs a DNS zone.	Invoke-DnsServerZoneSign
ZoneTransfer	Start	Starts a zone transfer for a secondary DNS zone from master servers.	Start-DnsServerZoneTransfer
ZoneUnsign	Invoke	Unsigns a DNS zone.	Invoke-DnsServerZoneUnsign

See also

Appendix A: DNSSEC Terminology

Step-by-Step: Demonstrate DNSSEC in a Test Lab

Domain Name System Security Extensions (DNSSEC) is a suite of extensions that add security to the DNS protocol. With DNSSEC, non-authoritative DNS servers are able to validate the responses they receive when they query other DNS servers. In addition, computers running Windows® 7 or later can be configured to require this validation be performed.

In this guide

This guide provides step-by-step instructions for deploying DNSSEC in a test lab using two server computers or, optionally, three server computers and one client computer. Software and hardware requirements are provided, as well as an overview of DNSSEC.

Important

The following instructions are for configuring a test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended

configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

DNSSEC overview

The DNS protocol is vulnerable to attack due to an inherent lack of authentication and integrity checking of data that is exchanged between DNS servers or provided to DNS clients. DNSSEC adds security to DNS responses by providing the ability for DNS servers to validate DNS responses. With DNSSEC, resource records are accompanied by digital signatures. These digital signatures are generated when DNSSEC is applied to a DNS zone using a process called zone signing. When a resolver issues a DNS query for resource record in a signed zone, a digital signature is returned with the response so that validation can be performed. If validation is successful, this proves that the data has not been modified or tampered with in any way.

DNS threats

DNS spoofing is a type of attack that involves impersonation of DNS server responses in order to introduce false information. In a spoofing attack, a malicious user attempts to guess that a DNS client or server has sent a DNS query and is waiting for a DNS response. A successful spoofing attack will insert a fake DNS response into the DNS server's cache, a process known as cache poisoning. A spoofed DNS server has no way of verifying that DNS data is authentic, and will reply from its cache using the fake information. An attacker can also set the time to live (TTL) on fake DNS data to a very long interval, causing the DNS server cache to remain poisoned for many hours or days. It is also possible to send a DNS spoofing attack directly to a DNS client; however these attacks are less persistent than cache poisoning attacks. Both of these types of attacks can be prevented with DNSSEC by requiring that DNS responses are validated as authentic. See the following figure.

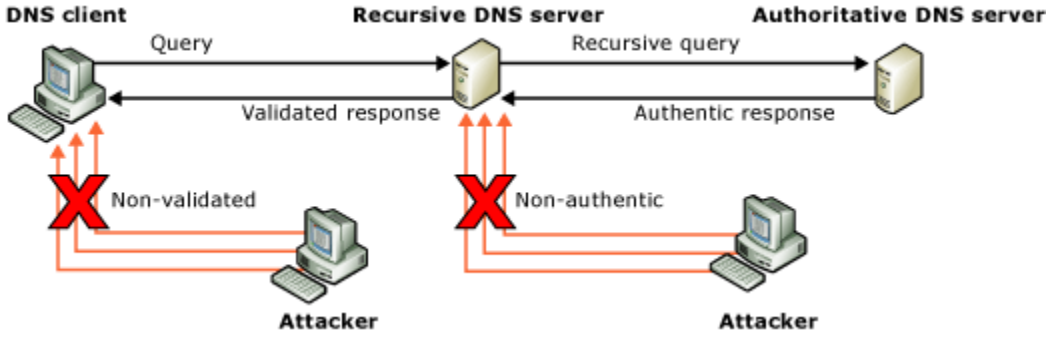

How DNSSEC works

DNSSEC uses digital signatures and cryptographic keys to validate that DNS responses are authentic. The following topics briefly discuss how these signatures are managed and validation is performed.

Digital signatures

Signatures generated with DNSSEC are contained within the DNS zone itself in the new resource records. These new resource records are called RRSIG (resource record signature) records. When a resolver issues a query for a name, the RRSIG record is returned in the response. A public cryptographic key called a DNSKEY is needed to verify the signature. The DNSKEY is retrieved by a DNS server during the validation process.

Zone signing

When you sign a zone with DNSSEC, you are individually signing all the records contained in the zone. This makes it possible to add, modify, or delete records in the zone without re-signing the entire zone. It is only necessary to re-sign the updated records.

Authenticated denial of existence

What if a DNS query is for a record that does not exist? If the DNS server responds that no record was found, this response also needs to be validated as authentic. However, since there is no resource record, then there is no RRSIG record. The answer to this problem is the Next Secure (NSEC) record. NSEC records create a chain of links between signed resource records. To create NSEC records, the zone is sorted and NSEC records are created such that each NSEC record has a pointer to the next NSEC record. The last NSEC record points back to the first record. When a query is submitted for a nonexistent record, the DNS server returns the NSEC record prior to where the nonexistent record would have been in the order. This allows for something called authenticated denial of existence.

NSEC3 is a replacement or alternative to NSEC that has the additional benefit of preventing “zone walking” which is the process of repeating NSEC queries in order to retrieve all the names in a zone. A DNS server running Windows Server® 2012 supports both NSEC and NSEC3. A zone can be signed with either NSEC or NSEC3, but not both.

Trust anchors

A trust anchor is a preconfigured public key associated with a specific zone. A validating DNS server must be configured with one or more trust anchors in order to perform validation. If the DNS server is running on a domain controller, trust anchors are stored in the forest directory partition in Active Directory Domain Services (AD DS) and can be replicated to all domain controllers in the forest. On standalone DNS servers, trust anchors are stored in a file named **TrustAnchors.dns**. A DNS server running Windows Server 2012 also displays configured trust anchors in the DNS Manager console tree in the **Trust Points** container. You can also use Windows PowerShell or Dnscmd.exe to view trust anchors.

DNSSEC key management

DNSSEC key management strategy includes planning for key generation, key storage, key expiration, and key replacement. Together, key expiration and replacement in DNSSEC is called

key rollover. In Windows Server 2012 key management is made easier with simple and flexible key generation, Active Directory storage and replication, an automated key rollover.

DNSSEC-aware clients

In Windows 8 and Windows Server 2012 the DNS Client service continues to be non-validating and security-aware, the same as computers running Windows 7 and Windows Server® 2008 R2. When the DNS client issues a query, it can indicate to the DNS server that it understands DNSSEC. However, the client is non-validating. When issuing queries, the DNS client relies on the local DNS server to indicate that validation was successful. If the server fails to perform validation, or reports that validation was unsuccessful, the DNS Client service can be configured to return no results.

NRPT

The Name Resolution Policy Table (NRPT) is a table that contains rules you can configure to specify DNS settings or special behavior for names or namespaces. The NRPT can be configured using Group Policy or by using the Windows Registry.

When performing DNS name resolution, the DNS Client service checks the NRPT before sending a DNS query. If a DNS query or response matches an entry in the NRPT, it is handled according to settings in the policy. Queries and responses that do not match an NRPT entry are processed normally. You can use the NRPT to require that the DNS Client service perform DNSSEC validation of DNS responses for the namespaces that you specify.

Scenario overview

This test lab demonstrates new DNSSEC functionality in Windows Server 2012. Three server computers and one client computer are used. See the following figure.

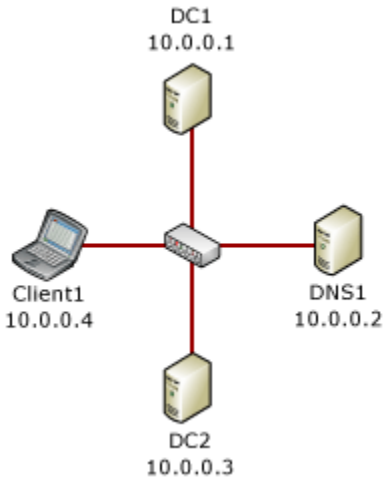

Hardware and software requirements

Two server computers are required to complete the test lab. Optionally, three server computers and one client computer can be used to complete additional steps in the test lab.

The following are required components of the test lab:

1. The product disc or other installation media for Windows Server 2012.
2. Two computers that meet the minimum hardware requirements for Windows Server 2012.

The following components are recommended but not required to complete the test lab:

1. The product disc or other installation media for Windows® 8.
2. Two computers that meet the minimum hardware requirements for Windows Server 2012 and Windows 8.

Steps for configuring the test lab

The following procedures are used to configure computers for the demonstration portion of the test lab:

- Configure DC1: DC1 is a domain controller and Active Directory-integrated authoritative DNS server.
- Configure DNS1: DNS1 is a non-authoritative, caching DNS server.
- Configure DC2: DC2 is a secondary domain controller and Active Directory-integrated DNS server.
- Configure Client1: Group Policy settings for DNS are applied to Client1 and it is used to issue DNS client queries.

Configure DC1

DC1 is a computer running Windows Server 2012, providing the following services:

- A domain controller for the contoso.com Active Directory domain.
- An authoritative DNS server for the contoso.com DNS zone.
- A DNSSEC Key Master for the contoso.com DNS zone.

Initial configuration of DC1 consists of the following steps:

- Install the operating system and configure TCP/IP on DC1
- Install Active Directory and DNS on DC1
- Configure the sec.contoso.com DNS zone
- Enable Remote Desktop on DC1

Additional tasks will be performed on DC1 during the demonstration portion of the test lab, including NRPT configuration, DNSSEC zone signing, distribution of trust anchors, and demonstration of DNSSEC key rollover.

Install the operating system and configure TCP/IP on DC1

To install the operating system and configure TCP/IP on DC1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. Click **Start**, type **ncpa.cpl**, and then press **ENTER**. The **Network Connections** control panel will open.

   Tip
   The previous step demonstrates new functionality in Windows Server 2012 that enables you to search and run applications, settings, and files by clicking **Start** and then typing a search term. You can also open the **Network Connections** control panel by clicking next to **Wired Ethernet Connection** in Server Manager using the **Local Server** view. For more information, see Common Management Tasks and Navigation in Windows Server 2012 (http://go.microsoft.com/fwlink/p/?LinkId=242147).
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.1** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install Active Directory and DNS on DC1

DC1 will serve as the primary domain controller and DNS server for the contoso.com Active Directory domain.

To configure DC1 as a domain controller and DNS server

1. The Server Manager Dashboard is displayed by default. In the navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DC1** under **Computer name**.
4. Click **OK** twice, and then click **Close**.
5. When you are prompted to restart the computer, click **Restart Now**.

6. After restarting the computer, sign-in using the local Administrator account.
7. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
8. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **Active Directory Domain Services** checkbox.
9. When you are prompted to add required features, click **Add Features**.
10. Click **Next** three times, and then click **Install**.
11. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DC1** is displayed, and then click **Close**.
12. Click the Notification flag and then click **Promote this server to a domain controller**.

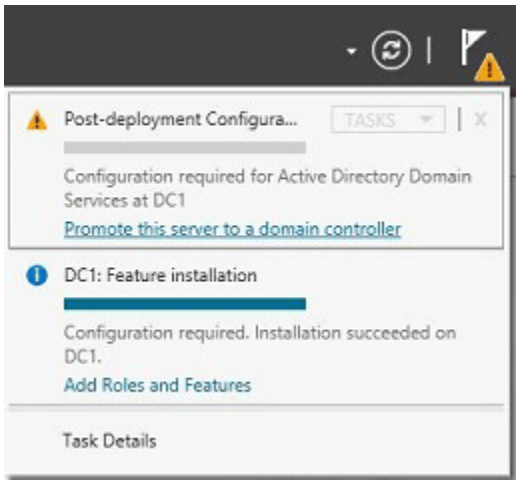

13. In the **Active Directory Domain Services Configuration Wizard**, on the **Deployment Configuration** page, choose **Add a new forest** and then next to **Root domain name**, type **contoso.com**.
14. Click **Next**, and then on the **Domain Controller Options** page, under **Type the Directory Services Restore Mode (DSRM) password**, type a password next to **Password** and **Confirm password**. Confirm that **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected, and then click **Next**.
15. Click **Next** five times and then click **Install**.
16. The computer will restart automatically to complete the installation process.
17. Sign in using the local Administrator account.
18. Next, a domain administrator account must be created to use when performing procedures in the test lab.

> **Tip**
> You can use the CONTOSO\Administrator account in this test lab and skip creation of a domain administrator account if desired. This account has domain administrator privileges, and other privileges. However, it is a best practice to disable or rename this account. For more information, see Active Directory Best Practices(http://go.microsoft.com/fwlink/p/?LinkID=243071).

To create a domain administrator account

1. On the Server Manager menu bar, click **Tools**, and then click **Active Directory Users and Computers**.
2. In the **Active Directory Users and Computers** console tree, double-click **contoso.com**, right-click **Users**, point to **New**, and then click **User**.
3. In the **New Object – User** dialog box, type **user1** under **User logon name** and next to **Full name**, then click **Next**.
4. Next to **Password** and **Confirm password**, type a password for the user1 account.
5. Clear the checkbox next to **User must change password at next logon**, select the **Password never expires** checkbox, click **Next**, and then click **Finish**.
6. Double-click **user1** and then click the **Member Of** tab.
7. Click **Add**, type **domain admins** under **Enter the object names to select**, click **OK** twice, and then close the **Active Directory Users and Computers** console.
8. Click **Start**, click **Administrator**, and then click **Sign out**.
9. Sign in to the computer using the user1 credentials by clicking the left arrow next to **CONTOSO\Administrator** and then clicking **Other user**.

Configure the sec.contoso.com DNS zone

Next, configure a new DNS zone: sec.contoso.com. This zone will be used to demonstrate DNSSEC zone signing.

Warning
The domain zone (contoso.com) can also be signed with DNSSEC. However, the test lab includes validation failure scenarios that add complexity if the domain zone is signed.

To configure the sec.contoso.com DNS zone

1. On the Server Manager menu, click **Tools**, and then click **DNS**.
2. In the DNS Manager console tree, right-click **Forward Lookup Zones** and then click **New Zone**.
3. In the **New Zone Wizard**, click **Next** three times, and then under **Zone name** type **sec.contoso.com**.
4. Click **Next** twice, and then click **Finish**.
5. Verify that the zone "**sec.contoso.com**" is displayed under **Forward Lookup Zones**.
6. Next, add one or more DNS resource records to the sec.contoso.com zone.
7. Leave the DNS Manager console open.

To add DNS resource records to the sec.contoso.com zone

1. Right-click **sec.contoso.com** and then click **New Host (A or AAAA)**.
2. In the **New Host** dialog box, type **dc1** under **Name**, type **10.0.0.1** under **IP address**, and

then click **Add Host**. The IP address of dc1.contoso.com is used here to help demonstrate DNSSEC success and failure scenarios.

3. Confirm that **The host record dc1.sec.contoso.com was successfully added** is displayed, and then click **OK**.
4. Add additional resource records to the zone if desired, and then click **Done**.

Enable Remote Desktop on DC1

DC1 will be used to demonstrate functionality of a network application in an environment with DNSSEC.

To enable Remote Desktop on DC1

1. In the Server Manager navigation pane, click **Local Server**.
2. Click the word Disabled next to Remote Desktop.
3. In the System Properties dialog box, on the Remote tab, click Allow connections from computers running any version of Remote Desktop (less secure), and then click OK.

Configure DNS1

DNS1 is a computer running Windows Server® 2012, providing the following services:

- A non-authoritative, recursive DNS server.
- A DNS client computer (optional: If Client1 is not used).

Initial configuration of DNS1 consists of the following steps:

- Install the operating system and configure TCP/IP on DNS1
- Install and configure DNS on DNS1
- Optional: Install Network Monitor on DNS1

During the demonstration portion of the test lab, DNS1 will be used to perform recursive DNS queries, host a trust anchor for the contoso.com domain, and provide DNSSEC validation for DNS client queries. Optionally (if a separate DNS client computer is not used), DNS1 will be used to issue DNS client queries.

Install the operating system and configure TCP/IP on DNS1

Tip

The procedure below is identical to the steps used to install the operating system and configure TCP/IP on DC1, with the exception that DNS1 is configured with an IP address of 10.0.0.2.

To install the operating system and configure TCP/IP on DNS1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language,

and regional settings, and provide a password for the local Administrator account.

3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. In the Server Manager navigation pane, click **Local Server** and then click the IP address next to **Wired Ethernet Connection**. The **Network Connections** control panel will open.
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.2** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install and configure DNS on DNS1

DNS1 is a domain member server running the DNS Server role service. DNS1 is not a domain controller.

To install and configure DNS on DNS1

1. In the Server Manager Dashboard navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DNS1** under **Computer name**.
4. Under **Member of**, select **Domain**, type **contoso.com**, and then click **OK**.
5. When you are prompted to provide credentials to join the domain, enter the credentials for the user1 account that was created previously.
6. Confirm that computer name and domain changes were successful, click **OK**, and then click **Close**.
7. When you are prompted to restart the computer, click **Restart Now**.
8. After restarting the computer, sign-in using the CONTOSO\user1 account.
9. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
10. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **DNS Server** checkbox.
11. When you are prompted to add required features, click **Add Features**.
12. Click **Next** three times, and then click **Install**.
13. Wait for the installation process to complete, verify on the **Installation progress** page that **Installation succeeded on DNS1.contoso.com** is displayed, and then click **Close**.
14. On the Server Manager menu bar, click **Tools** and then click **DNS**.
15. In the DNS Manager console tree, right-click **DNS1** and then click **Properties**.

16. Click the **Forwarders** tab, click **Edit**, type **10.0.0.1**, and then click **OK** twice.
17. Leave the DNS Manager console open.

Optional: Install Network Monitor on DNS1

Network Monitor will be used to view detailed information about DNS queries. Installation of Network Monitor is optional in the test lab. Results of network traffic analysis using Network Monitor are provided in the Appendix: Network Monitor results section.

To install Network Monitor

1. Download the latest version of Network Monitor from the Microsoft Download Center: Network Monitor 3.4 (http://go.microsoft.com/fwlink/p/?LinkId=103158).
2. Double-click the installation file, click **Yes** when prompted to continue, click **Next**, read and accept the license agreement, and then click **Next**.
3. Choose **Use Microsoft Update when I check for updates (recommended)** and then click **Next**.
4. Choose the **Complete** installation type, and then click **Install**.
5. Click **Yes** in the **User Account Control** dialog box.
6. Click **Finish** and click **Yes** in the **User Account Control** dialog box.

Configure DC2

DC2 is a computer running Windows Server 2012, providing the following services:

- A secondary domain controller for the contoso.com Active Directory domain.
- An authoritative DNS server for the contoso.com DNS zone.

Note
Installing and configuring DC2 is recommended but optional. DC2 is required to perform some, but not all, steps in the test lab. If you must limit the number of computers used in the test lab, skip the steps for installing and configuring DC2. DC2 is required for demonstration of Active Directory replication and transferring the Key Master role from DC1 to DC2.

Initial configuration of DC2 consists of the following steps:

- Install the operating system and configure TCP/IP on DC2
- Install Active Directory and DNS on DC2

During the demonstration portion of the test lab, DC2 will be used to demonstrate Active Directory replication of a DNSSEC signed zone, and how to transfer the Key Master role on DC1 to another authoritative DNS server.

Install the operating system and configure TCP/IP on DC2

Tip

The procedure below is identical to the steps used to install the operating system and configure TCP/IP on DC1, with the exception that DC2 is configured with an IP address of 10.0.0.3.

To install the operating system and configure TCP/IP on DC2

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. In the Server Manager navigation pane, click **Local Server** and then click the IP address next to **Wired Ethernet Connection**. The **Network Connections** control panel will open.
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.3** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install Active Directory and DNS on DC2

The steps to install Active Directory and DNS on DC2 are nearly the same as those for DC1. All steps are provided below for completeness.

To install Active Directory and DNS on DC2

1. The Server Manager Dashboard navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DC2** under **Computer name**.
4. Under **Member of**, select **Domain**, type **contoso.com**, and then click **OK**.
5. When you are prompted to provide credentials to join the domain, enter the credentials for the user1 account.
6. Confirm that computer name and domain changes were successful, click **OK**, and then click **Close**.
7. When you are prompted to restart the computer, click **Restart Now**.
8. After restarting the computer, sign-in using the CONTOSO\user1 account.

9. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
10. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **Active Directory Domain Services** checkbox.
11. When you are prompted to add required features, click **Add Features**.
12. Click **Next** three times, and then click **Install**.
13. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DC2.contoso.com** is displayed, and then click **Close**.
14. In Server Manager, click the Notification flag and then click **Promote this server to a domain controller**.
15. In the **Active Directory Domain Services Configuration Wizard**, on the **Deployment Configuration** page, choose **Add a domain controller to an existing domain**, confirm that the name displayed next to Domain is **contoso.com**, supply credentials for the CONTOSO\user1 account, and then click **Next**.
16. On the **Domain Controller Options** page, select the **Domain Name System (DNS) server** and **Global Catalog (GC)** checkboxes, type the Directory Services Restore Mode (DSRM) password next to **Password** and **Confirm password**, click **Next** five times, and then click Install.
17. Confirm that installation was successful. The computer will restart automatically.
18. After the computer has restarted, sign in with the CONTOSO\user1 credentials.

Configure Client1

Client1 is a computer running Windows® 8 that is acting as a DNS client. Configuration of Client1 consists of the following steps:

- Install the operating system and configure TCP/IP on Client1
- Join Client1 to the contoso.com domain
- Pin Windows PowerShell to the taskbar

During the demonstration portion of the test lab, Client1 will receive NRPT settings from Group Policy and be used to perform DNS queries.

Install the operating system and configure TCP/IP on Client1

To install the operating system and configure TCP/IP on Client1

1. Start your computer using the Windows 8 product disc or other digital media.
2. When prompted, enter a product key and accept license terms.
3. When prompted to enter a computer name, type **Client1** and click **Next**.
4. Click **Use express settings**.
5. On the **Sign in to your PC** page, click **Don’t want to sign in with a Microsoft account** and then click **Local account**.

6. Next to **User name**, type **user1**, enter a password and password hint, and then click **Finish**.
7. On the **Start** page, type **ncpa.cpl**, and then press **ENTER**. The **Network Connections** control panel will open.
8. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
9. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
10. On the **General** tab, choose **Use the following IP address**.
11. Next to **IP address** type **10.0.0.4** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
12. Next to **Preferred DNS server**, type **10.0.0.2**.

    Important
    The DNS client computer should use a non-authoritative DNS server for this test lab. The IP address used for **Preferred DNS server** should correspond to DNS1 (10.0.0.2), not a DNS server running on a domain controller (DC1 or DC2).
13. Click **OK** twice, and then close the **Network Connections** control panel.

Join Client1 to the contoso.com domain

In order for Client1 to receive domain Group Policy settings, it must be joined to the contoso.com domain.

To join Client1 to the contoso.com domain

1. Click **Start**, type **sysdm.cpl** and then press ENTER.
2. In the **System Properties** dialog box, click **Change**.
3. Under **Member of**, choose **Domain**, type **contoso.com** and then click **OK**.
4. When you are prompted to enter an account with permission to join the domain, provide the credentials for the CONTOSO\user1 account, and then click **OK**.
5. Confirm that **Welcome to the contoso.com domain** is displayed, click **OK** twice, and then click **Close**.
6. When you are prompted to restart the computer, click **Restart Now**.
7. After restarting the computer, press **Ctrl+Alt+Delete**, click the left arrow, click **Other user**, and sign in using credentials for the CONTOSO\user1 account.

Pin Windows PowerShell to the taskbar

The DNSSEC demonstration on Client1 makes use of Windows PowerShell to query DNS servers. To make Windows PowerShell more easily accessible, it will be pinned to the taskbar.

To pin Windows PowerShell to the taskbar

1. On the **Start** page, type **powershell**, right-click Windows PowerShell, and then click Pin

to taskbar. Press the ESC key to return to the desktop.

2. Confirm that Windows PowerShell is pinned to the taskbar.

DNSSEC demonstration

For the DNSSEC demonstration portion of the test lab, you can use DNS1 instead of Client1 to perform DNS client queries if the Client1 computer is not available. If the DC2 computer is not available, you must skip some of the procedures below.

A demonstration of DNSSEC functionality on Windows Server 2012 consists of the following procedures:

1. Query an unsigned zone without DNSSEC validation required.
2. Sign a zone on DC1 and distribute trust anchors
3. Query a signed zone without DNSSEC validation required
4. Query a signed zone with DNSSEC validation required.
5. Unsign the zone and then re-sign the zone with custom parameters.
6. Demonstrate failed validation.
7. Demonstrate Active Directory replication of DNSSEC signed resource records.
8. Transfer the Key Master role for sec.contoso.com to DC2

Query an unsigned zone without DNSSEC validation required

First, use the resolve-dnsname cmdlet to query an unsigned zone when validation is not required.

To query an unsigned zone without DNSSEC validation required

1. On Client1, click **Windows PowerShell** on the taskbar, type **cd** and press ENTER.
2. Start a Network Monitor capture if desired. Stop the capture after issuing the following command, and then save the capture using the name: **Capture1**.
3. Type the following command, and then press ENTER:

```
resolve-dnsname dc1.sec.contoso.com –server dns1 –dnssecok
```

```
PS C:\> resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
dc1.sec.contoso.com                            A      3600  Answer     10.0.0.1

Name       : .
QueryType  : OPT
TTL        : 32768
Section    : Additional
Data       : {}
```

Tip

The **dnssecok** option in the command above tells the DNS server the client understands DNSSEC and the server can send these additional records. Since

the zone is not yet signed, no signature (RRSIG) records are displayed in the response.

4. Leave the Windows PowerShell prompt open for the following procedures.

To verify remote connections to dc1.sec.contoso.com

1. Type the following command and press ENTER:

```
mstsc /v:dc1.sec.contoso.com
```

2. Enter the password for the user1 account and click **OK**.
3. When you are prompted that there is a problem with security of the remote computer, click **Yes**.
4. Verify that you are able to successfully connect to dc1.sec.contoso.com, and then close the remote session.

Sign a zone on DC1 and distribute trust anchors

Next, sign the sec.contoso.com zone and distribute a trust anchor for the zone. Trust anchor distribution is manual for DNS servers that are not running on domain controllers, such as DNS1. Automatic trust anchor distribution can be enabled for Active Directory-integrated DNS servers such as DC2.

To sign a zone on DC1

1. In the DNS Manager console tree on DC1, navigate to **Forward Lookup Zones > sec.contoso.com**.
2. Right-click **sec.contoso.com**, point to **DNSSEC**, and then click **Sign the Zone**.

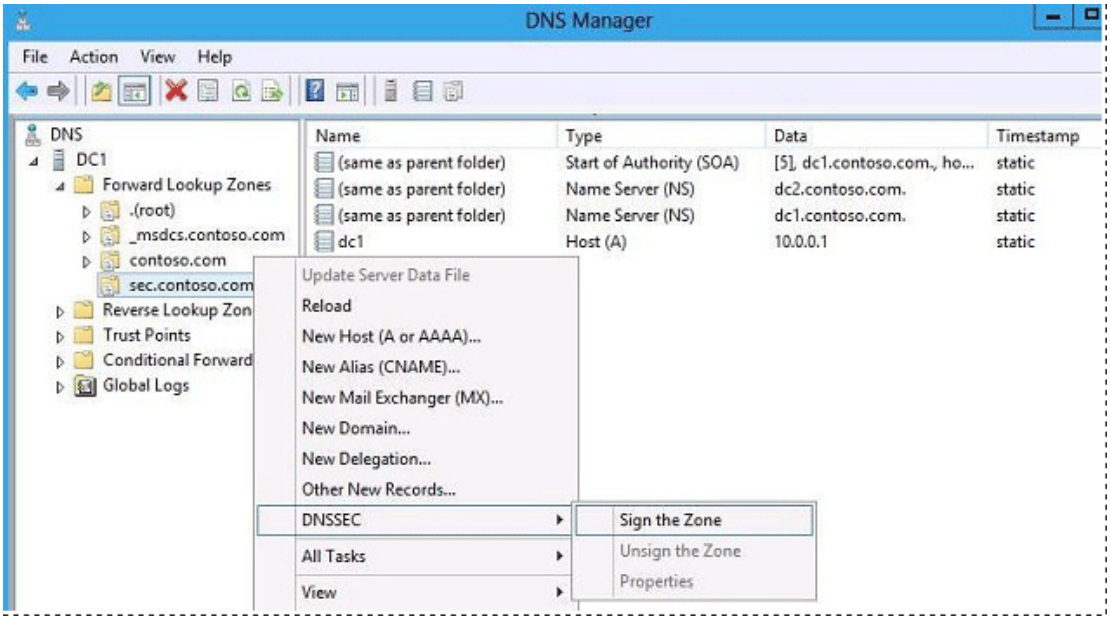

3. In the **Zone Signing Wizard**, click **Next**, and then choose **Use recommended settings to sign the zone**.

4. Click **Next** twice, confirm that **The zone has been successfully signed** is displayed, and then click **Finish**.
5. Refresh the DNS Manager console and verify that a new icon is displayed for the sec.contoso.com zone, indicating that it is currently signed with DNSSEC.
6. Click the sec.contoso.com zone and review the new resource records that are present, including DNSKEY, RRSIG and NSEC3 records.

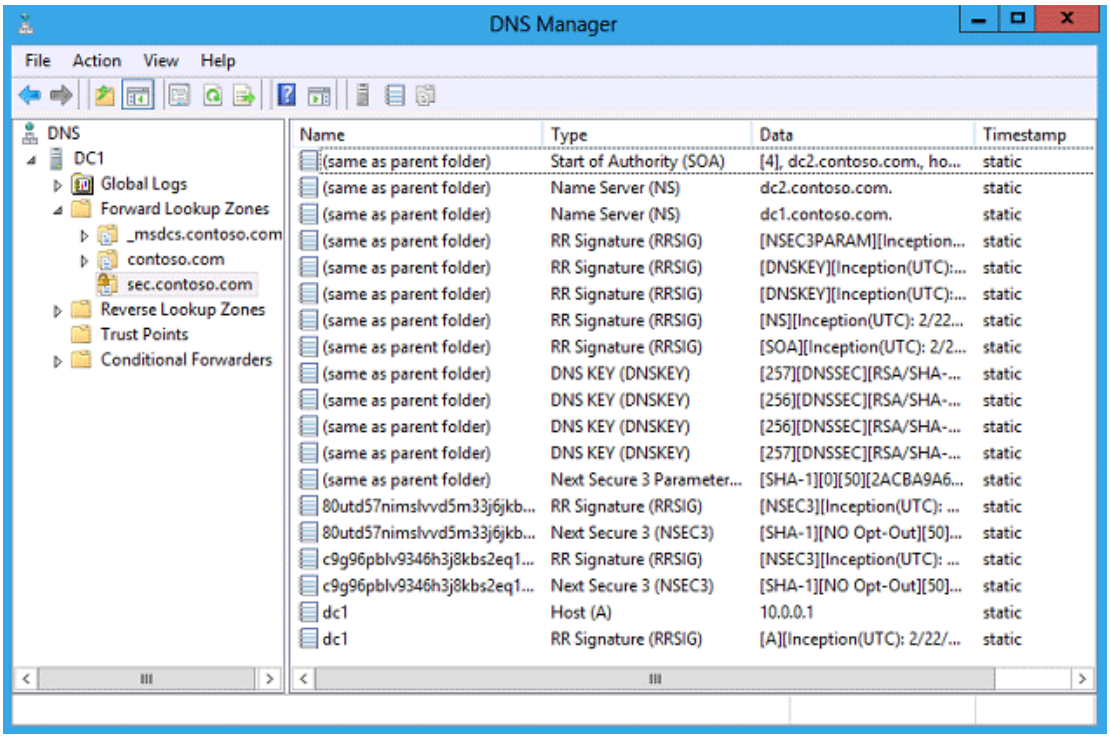

7. Leave the DNS Manager console open.

To distribute a trust anchor to DNS1

1. On DC1, click **Windows Explorer** on the taskbar.
2. Navigate to **C:\Windows\System32**, right-click the **dns** folder, point to **Share with**, and then click **Advanced sharing**.
3. In the **dns Properties** dialog box, click **Advanced Sharing**, select the **Share this folder** checkbox, verify the **Share name** is **dns**, and then click **OK**.

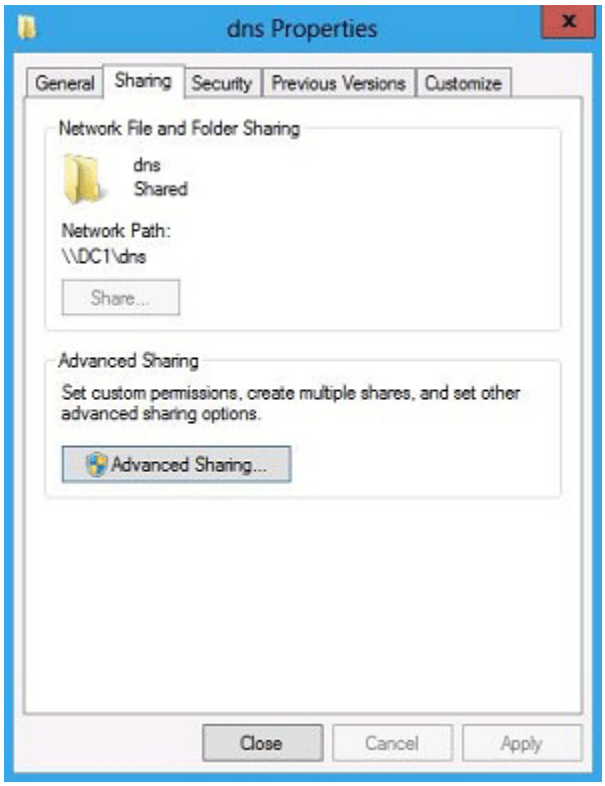

4. Click **Close** and then close Windows Explorer.
5. On DNS1, in the DNS Manager console tree, navigate to the **Trust Points** folder.
6. Right-click **Trust Points**, point to **Import**, and then click **DNSKEY**.
7. In the **Import DNSKEY** dialog box, type **\\dc1\dns\keyset-sec.contoso.com** and then click **OK**.

To verify trust anchors

1. In the console tree, navigate to **Trust Points > com > contsoso > sec** and verify that import was successful.

   Tip
   Two DNSKEY trust points are displayed, one for the active key and one for the standby key.

2. On any computer, click **Windows PowerShell**, type the following command and then press ENTER:

   ```
   resolve-dnsname -name sec.contoso.com.trustanchors -type
   dnskey -server dns1
   ```

   Verify that two trust anchors are displayed.
3. On DNS1, right-click **Windows PowerShell** and then click **Run as Administrator**.
4. Type the following command and then press ENTER:

```
get-dnsservertrustanchor sec.contoso.com
```

Verify that two trust anchors are displayed.

Delete and re-distribute trust anchors using Windows PowerShell

1. On DNS1, in the Administrator Windows PowerShell window, type the following command and press ENTER twice:

   ```
   remove-dnsservertrustanchor –name sec.contoso.com
   ```

2. Type the following command and then press ENTER:

   ```
   get-dnsservertrustanchor sec.contoso.com
   ```

   Verify that “Failed to enumerate the trust anchors” is displayed.

3. Type the following command and then press ENTER twice:

   ```
   remove-dnsserverzone –name trustanchors
   ```

   Important
   The **trustanchors** zone is deleted using the remove-dnsserverzone cmdlet so that the add-dnsserverprimaryzone cmdlet can be demonstrated. It is not typically required to remove and restore the **trustanchors** zone after deleting trust anchors.

4. Type the following command and then press ENTER:

   ```
   add-dnsserverprimaryzone –computername dns1 trustanchors –
   zonefile trustanchors.dns
   ```

5. Type the following command and then press ENTER:

   ```
   get-dnsserverresourcerecord –zonename sec.contoso.com –rrtype
   dnskey –computername dc1 | %{ $_.recorddata | add-
   dnsservertrustanchor -name sec.contoso.com }
   ```

6. Type the following command and then press ENTER:

   ```
   get-dnsservertrustanchor sec.contoso.com
   ```

   Verify that two trust anchors are again displayed.

To distribute a trust anchor to DC2

1. On DC1, in the DNS Manager console tree, navigate to **Forward Lookup Zones > sec.contoso.com**.
2. Right click **sec.contoso.com**, point to **DNSSEC**, and then click **Properties**.
3. Click the **Trust Anchor** tab.
4. Select the **Enable the distribution of trust anchors for this zone** checkbox, and then click **OK**.
5. When you are prompted to confirm changes to the zone, click **Yes**.
6. When you are prompted that configuration was successful, click **OK**.
7. On DC2, refresh the view in DNS Manager and confirm that trust anchors for

sec.contoso.com are present.

Important
You might need to wait a few minutes for replication to occur on DC2.

Query a signed zone without DNSSEC validation required

Additional DNSSEC related information is displayed for signed resource records. Compare query results for dc1.contoso.com to query results for dc1.sec.contoso.com if desired.

To query a signed zone without DNSSEC validation required

1. Start a Network Monitor capture if desired. Stop the capture after issuing the following command, and then save the capture using the name: **Capture2**.
2. On Client1, at the Windows PowerShell prompt, type the following command and then press ENTER:

```
resolve-dnsname dc1.sec.contoso.com –server dns1 –dnssecok
```

```
PS C:\> resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok

Name                                           Type   TTL   Section    IPAddress
----                                           ----   ---   -------    ---------
dc1.sec.contoso.com                            A      3600  Answer     10.0.0.1

Name          : dc1.sec.contoso.com
QueryType     : RRSIG
TTL           : 3600
Section       : Answer
TypeCovered   : A
Algorithm     : 7
LabelCount    : 4
OriginalTtl   : 3600
Expiration    : 2/9/2012 8:22:15 PM
Signed        : 1/30/2012 7:22:15 PM
Signer        : sec.contoso.com
Signature     : {32, 134, 10, 216...}


Name       : .
QueryType  : OPT
TTL        : 32768
Section    : Additional
Data       : {}
```

3. To verify that DNSSEC validation is not currently required, type the following command and press ENTER:

```
get-dnsclientnrptpolicy
```

4. Confirm that no NRPT policy for the sec.contoso.com namespace is currently applied to the client computer.
5. Leave the Windows PowerShell prompt open.

Query a signed zone with DNSSEC validation required

The Name Resolution Policy Table (NRPT) is used to require DNSSEC validation. The NRPT can be configured in local Group Policy for a single computer, or domain Group Policy for some or all computers in the domain. The following procedure uses domain Group Policy.

To require DNSSEC validation be performed

1. On DC1, on the Server Manager menu bar, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management console tree, under **Domains > contoso.com > Group Policy Objects**, right-click **Default Domain Policy**, and then click **Edit**.
3. In the Group Policy Management Editor console tree, navigate to **Computer Configuration > Policies > Windows Settings > Name Resolution Policy**.
4. In the details pane, under **Create Rules** and **To which part of the namespace does this rule apply**, choose **Suffix** from the drop-down list and type **sec.contoso.com** next to **Suffix**.
5. On the **DNSSEC** tab, select the **Enable DNSSEC in this rule** checkbox and then under **Validation** select the **Require DNS clients to check that name and address data has been validated by the DNS server** checkbox.
6. In the bottom right corner, click **Create** and then verify that a rule for **sec.contoso.com** was added under **Name Resolution Policy Table**.

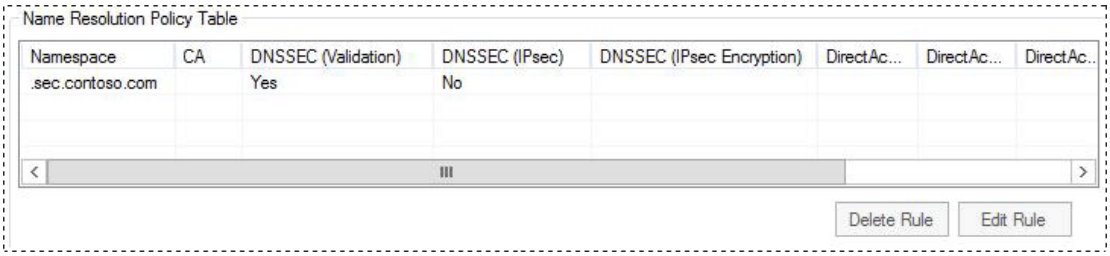

7. Click **Apply**, and then close the Group Policy Management Editor.
8. On DC1, type the following commands at the Windows PowerShell prompt, and then press ENTER:

```
gpupdate /force
get-dnsclientnrptpolicy
```

9. Verify that computer and user policy updates were successful, and that the value of **DnsSecValidationRequired** is **True** for the **.sec.contoso.com** namespace.
10. Repeat the Group Policy update (gpupdate /force) and verify NRPT policy on client1.

```
PS C:\> gpupdate /force
Updating policy...

Computer Policy update has completed successfully.
User Policy update has completed successfully.

PS C:\> get-dnsclientnrptpolicy


Namespace                       : .sec.contoso.com
QueryPolicy                     :
SecureNameQueryFallback         :
DirectAccessIPsecCARestriction  :
DirectAccessProxyName           :
DirectAccessDnsServers          :
DirectAccessEnabled             :
DirectAccessProxyType           : NoProxy
DirectAccessQueryIPsecEncryption :
DirectAccessQueryIPsecRequired  : False
NameServers                     :
DnsSecIPsecCARestriction        :
DnsSecQueryIPsecEncryption      :
DnsSecQueryIPsecRequired        : False
DnsSecValidationRequired        : True
NameEncoding                    : Utf8WithoutMapping
```

To query a signed zone with DNSSEC validation required

1. Start a Network Monitor capture if desired. Stop the capture after issuing the following command, and then save the capture using the name: **Capture3**.
2. On client1, at the Windows PowerShell prompt, type the following command and then press ENTER:

   ```
   resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok
   ```

3. Confirm that the same results are returned as previously before validation was required. Because a valid trust anchor is present on DNS1, the query is successful even if validation is required.

Query DNSSEC records in the sec.contoso.com zone

Before unsigning and re-signing the zone, issue some queries for DNSSEC records. These types of queries can be useful when troubleshooting DNSSEC.

To query DNSSEC records in the sec.contoso.com zone

1. On client1, at the Windows PowerShell prompt, type the following command and then press ENTER:

   ```
   resolve-dnsname -name dc1.sec.contoso.com -type soa -server
   dns1 -dnssecok
   ```

2. On client1, at the Windows PowerShell prompt, type the following command and then press ENTER:

   ```
   resolve-dnsname -name sec.contoso.com -type dnskey -server
   ```

```
dns1 -dnssecok
```

Tip

To see the syntax for resolve-dnsname, including all possible values for the **Type** parameter, type **get-help resolve-dnsname**.

Unsign the zone and then re-sign the zone with custom parameters

DNSSEC signing will be removed from the sec.contoso.com zone and then the zone will be re-signed using custom DNSSEC parameters.

To unsign the zone

1. On DC1, in the DNS Manager console tree, navigate to **Forward Lookup Zones > sec.contoso.com**.
2. Right-click **sec.contoso.com**, point to **DNSSEC**, and then click **Unsign the Zone**.
3. In the Unsign zone wizard, click **Next**.
4. Verify that **The zone has been successfully unsigned** is displayed, and then click **Finish**.
5. Refresh the view in DNS Manager and verify that the **sec.contoso.com** zone no longer contains DNSSEC signed records, and the icon next to the zone indicates it is not currently signed.

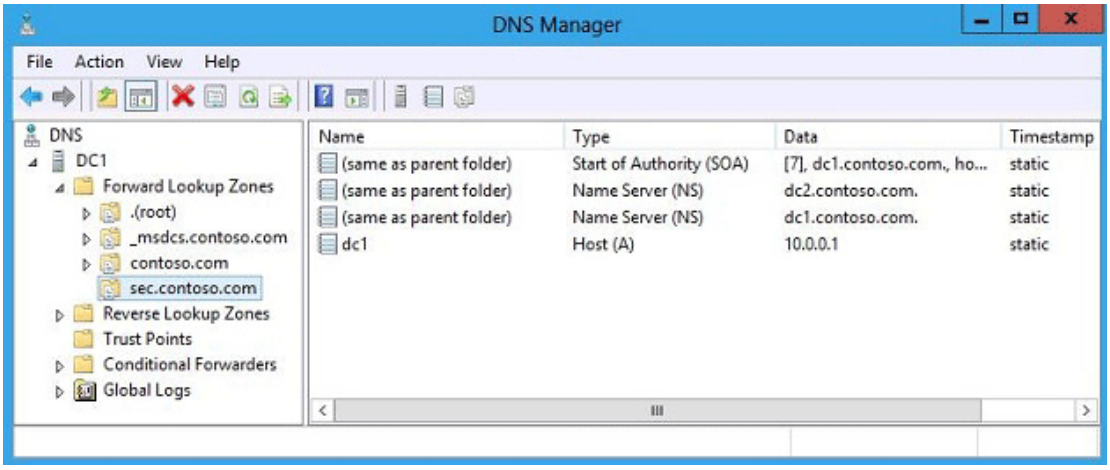

To re-sign the zone with custom parameters

1. On DC1, right-click **sec.contoso.com**, point to **DNSSEC**, and then click **Sign the Zone**.
2. In the **Zone Signing Wizard**, click **Next**.
3. **Customize zone signing parameters** is chosen by default. Click **Next**.
4. On the **Key Master** page, **The DNS server DC1 is the Key Master** is chosen by default, because zone signing is being performed on DC1.

If you have configured DC2 in this test lab, review options available when **Select another primary server as the Key Master** is chosen. Do not choose this option, but verify that dc2.contoso.com is also available as a possible Key Master for this zone. When you are alerted that all authoritative servers capable of DNSSEC online signing will be loaded, click **Yes**.

5. Ensure that DC1 is chosen as the Key Master and then click **Next** twice.
6. On the **Key Signing Key (KSK)**, page, click the existing KSK (with key length of 2048), and then click **Remove**.
7. To add a new KSK, click **Add**.
8. In the **New Key Signing Key (KSK)** dialog box, under **Key Properties**, click the drop-down next to **Cryptographic algorithm** and select **RSA/SHA-512**.
9. Under **Key Properties**, click the drop-down next to **Key length (Bits)** and select **4096** and then click **OK**.

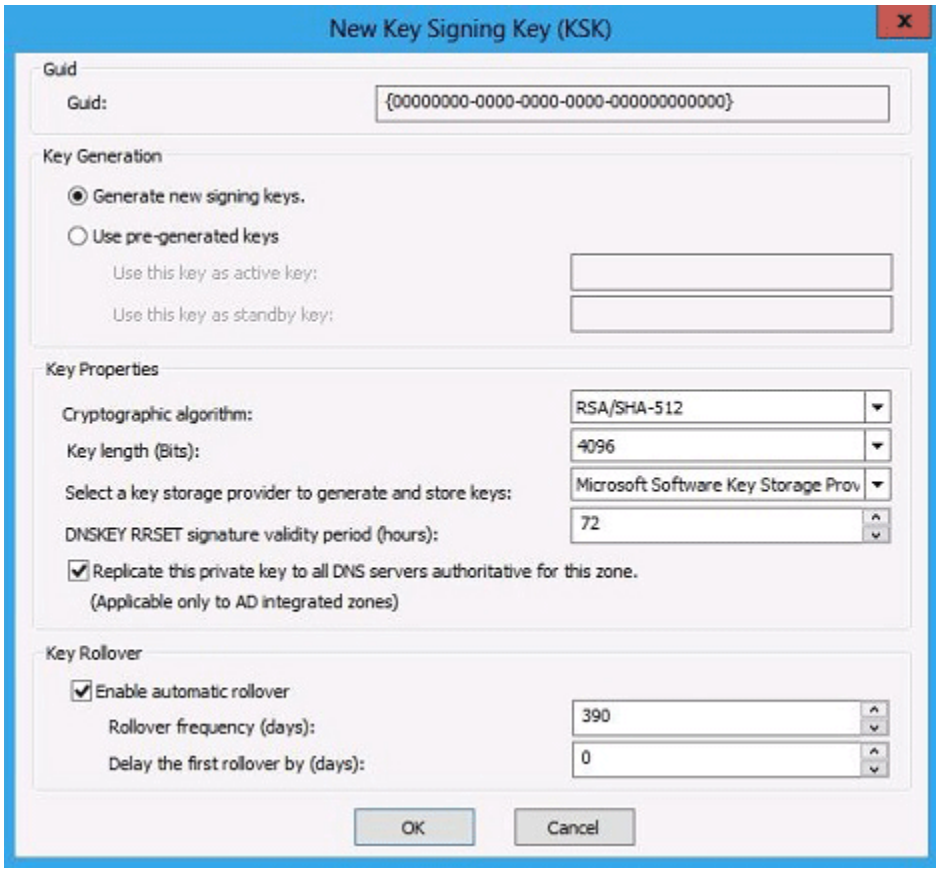

10. Click **Next** until **You have successfully configured the following parameters to sign the zone** is displayed.
11. Review the parameters you have chosen and then click **Next** to start the zone signing process.
12. Confirm that **The zone has been successfully signed** is displayed, click **Finish**, and then refresh the view in DNS Manager to verify the zone is signed again.
13. Refresh the view for the **Trust Points** folder and verify that new DNSKEY trust points are present that use the RSA/SHA-512 algorithm.
14. At an Administrator Windows PowerShell prompt, type the following commands and press ENTER:

```
Get-dnsservertrustanchor -name sec.contoso.com -computername dns1
Get-dnsservertrustanchor -name sec.contoso.com -computername dc1
Get-dnsservertrustanchor -name sec.contoso.com -computername dc2
```

Note that DC1 and DC2 are using the new trust anchors, but DNS1 has the old trust anchors. You might need to wait a few minutes for automatic distribution of the new trust

anchors to DC2.

Demonstrate failed validation

Because the trust anchor that was distributed to DNS1 is no longer valid, DNSSEC validation will fail when resource records are queried in the sec.contoso.com zone.

To demonstrate failed validation

1. On DNS1, view the currently installed **Trust Points** for sec.contoso.com and verify that the old trust anchor that uses the RSA/SHA-1 algorithm is present.
2. To flush the DNS server cache, right-click **DNS1** and then click **Clear Cache**.

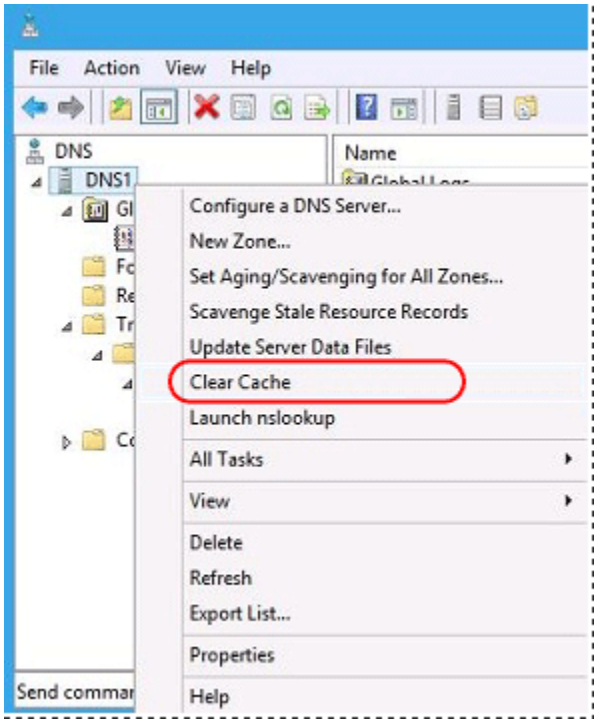

3. Start a Network Monitor capture if desired. Stop the capture after issuing the following command, and then save the capture using the name: **Capture4**.
4. On client1, type the following command at the Windows PowerShell prompt and then press ENTER:

```
resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok
```

```
PS C:\> resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok
Resolve-DnsName : dc1.sec.contoso.com : DNS server failure
At line:1 char:1
+ resolve-dnsname dc1.sec.contoso.com -server dns1 -dnssecok
+ ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    + CategoryInfo          : ResourceUnavailable: (dc1.sec.contoso.com:String) [Resolve-DnsName], Win32Exception
    + FullyQualifiedErrorId : RCODE_SERVER_FAILURE,Microsoft.DnsClient.Commands.ResolveDnsName
```

Important

Automatic updating of trust anchors on a non-authoritative, validating DNS server (per RFC 5011) only occurs during key rollover. If you unsign and re-sign the zone manually with new keys, you must also distribute a new trust anchor manually.

If a validating DNS server has an incorrect trust anchor, DNS queries that require validation will indicate a server failure.

When no trust anchor is present, queries will also appear to fail validation. Since no trust anchor is present, the server does not attempt to validate the response. In this scenario, an unsecure packet error is displayed:

To demonstrate an unsecure response

1. On DNS1, at the Administrator Windows PowerShell prompt, type the following command and then press ENTER twice:

   ```
   remove-dnsservertrustanchor sec.contoso.com
   ```

2. Start a Network Monitor capture if desired. Stop the capture after issuing the following command, and then save the capture using the name: **Capture5**.
3. Type the following command and press ENTER:

   ```
   resolve-dnsname –name dc1.sec.contoso.com –server dns1 -
   dnssecok
   ```

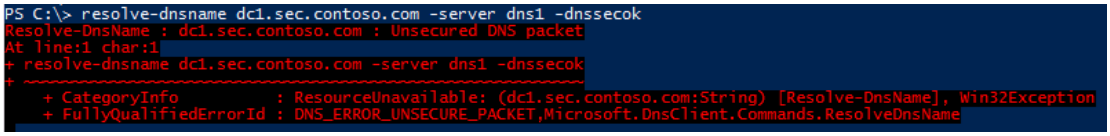

Demonstrate Remote Desktop failure

Because DNSSEC validation fails, you cannot connect to dc1.sec.contoso.com using Remote Desktop.

To demonstrate Remote Desktop failure

1. On client1, type the following commands at the Windows PowerShell prompt, and then press ENTER:

   ```
   ipconfig /flushdns
   ```

   ```
   mstsc /v:dc1.sec.contoso.com
   ```

2. Verify that **Remote Desktop can’t find the computer “dc1.sec.contoso.com”** is displayed.

Demonstrate Active Directory replication of DNSSEC signed resource records

When DNS servers are Active Directory-integrated, trust anchors and signed resource records are updated automatically even if the zone is unsigned and re-signed manually.

To demonstrate Active Directory replication of DNSSEC signed resource records

1. On DC2, in DNS Manager, view the contents of the **Trust Points** folder. Refresh the view if necessary to view the current trust anchors.
2. Verify that the DNSKEY trust anchors for **sec.contoso.com** are automatically updated to use the **RSA/SHA-512** algorithm.
3. In the DNS Manager console tree, click **Global Logs > DNS Events** and review event ID 7653 which states that the DNS server has detected that zone signing parameters for the zone sec.contoso.com have been changed and the zone will be re-signed. No event is displayed after zone signing is completed.
4. Click **Forward Lookup Zones > sec.contoso.com** in the console tree and verify that **Secure Entry Point** DNSKEY records are present that use the **RSA/SHA-512** algorithm.

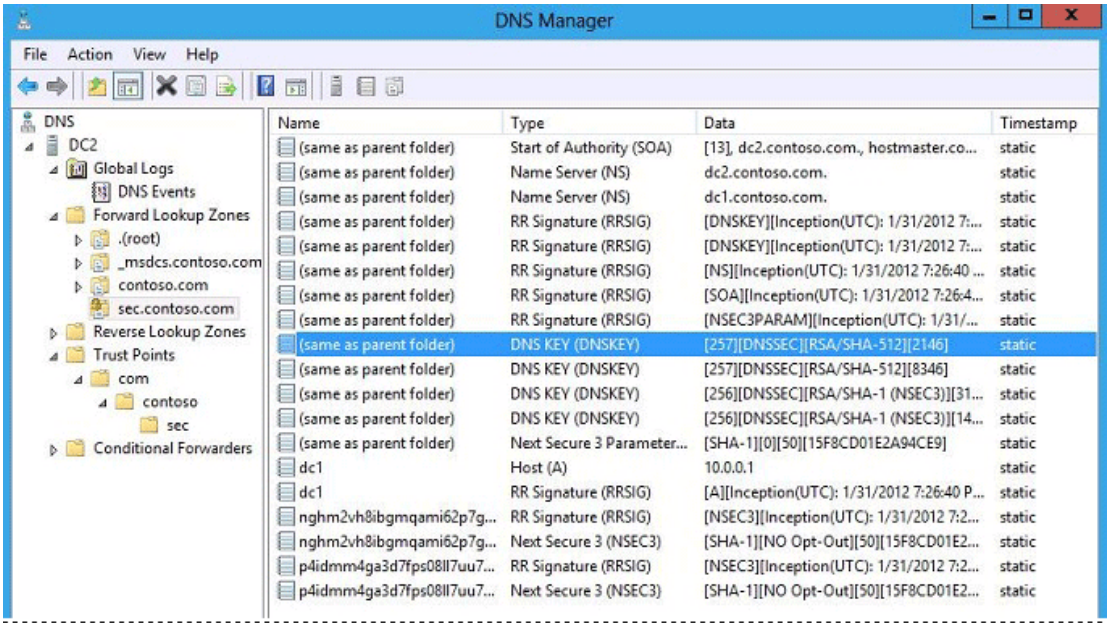

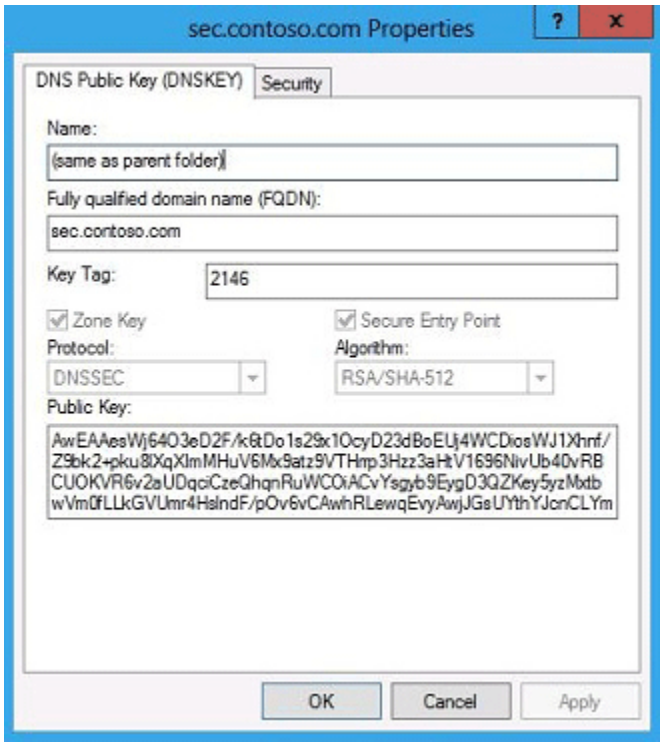

5. On DC1, in DNS Manager, add a new host (A) record for dns1.sec.contoso.com with an IP address of 10.0.0.2.
6. Refresh the view in DNS Manager and verify that an RR Signature (RRSIG) record for dns1 is automatically created.
7. On DC2, refresh the view in DNS Manager and verify that the new signed record has replicated to this server.

> **Tip**
> Adding or editing existing records in a zone does not trigger zone re-signing. Only the new or updated resource records are signed with the updated start of authority (SOA) record for the zone.

Transfer the Key Master role for sec.contoso.com to DC2

If might be necessary to transfer the Key Master role for a zone to another DNS server. The role transfer can be performed from any authoritative DNS server, and the current Key Master can be online or offline. In the following example, the current Key Master is online.

To transfer the Key Master role for sec.contoso.com to DC2

1. On DC1 or DC2, in DNS Manager, right-click the **sec.contoso.com** zone, point to **DNSSEC**, and then click **Properties**.
2. On the **Key Master** tab, select **Use the following DNS server as the Key Master**.

3. Click the drop-down list and when you are alerted that all authoritative DNS servers will be loaded, click **Yes**.
4. Choose **dc2.contoso.com** from the list and then click **OK**.
5. When you are alerted that the Key Master setting will be changed, click **Yes**.
6. Verify that **Key Master for the zone sec.contoso.com is updated successfully** is displayed.
7. Verify that DNS event ID 7649 is displayed on the new Key Master and DNS event ID 7648 is displayed on the previous Key Master.

Appendix: Network Monitor results

The following sections provide information about the results of Network Monitor (netmon) captures during the DNSSEC demonstration portion of the test lab. Analysis of network traffic uses the **Network Conversations**, **Frame Summary**, and **Frame Details** views.

Packet capture results

In all test lab captures, two IPv4 network conversations are displayed. The IPv4 network conversations include queries for **dns1.contoso.com** and **dc1.sec.contoso.com**. Queries for dc1.sec.contoso.com will also include **A** record queries and **AAAA** record queries. For the purposes of the test lab, you can ignore the IPv6 network conversation, and all queries for host records for dns1.contoso.com and AAAA queries for dc1.sec.contoso.com. The two IPv4 network conversations displayed are:

1. **10.0.0.4 – 10.0.0.2**: A packet exchange between Client1 (10.0.0.4) and DNS1 (10.0.0.2). Locate the query for the host (A) record for dc1.sec.contoso.com. There are two phases under **Frame Summary**, one with a source of Client1 and one with a source of DNS1. Locate the packets with the source of DNS1 (with destination = Client1).

   Things to note in this network conversation are:

   - In **Frame Details**, under **Dns\Flags**, the **AuthenticatedData** (AD) flag will be on (“1”) or off (“0”) depending on whether or not information that was returned was validated as authentic.
   - In **Frame Details**, under **Dns\Flags\ARecord**, RRSIG data is returned when the zone is signed as long as the value of **Rcode** under **Dns\Flags** is **Success**.
2. **10.0.0.2 – 10.0.0.1**: A packet exchange between DNS1 (10.0.0.2) and DC1 (10.0.0.1).

   Things to note are:

   - In **Frame Details**, under **Dns\Flags\ARecord**, DNS1 issues a query for sec.contoso.com of type **DNSKEY** when a trust anchor is present on DNS1.

If you started, stopped, and saved a network conversation at each of the suggested points in the guide, you will have the following captures:

Capture file name	Zone signed	Validation required	AD bit	DNSKEY request
Capture1	No	No	0	No
Capture2	Yes (no TA)	No	0	No
Capture3	Yes (valid TA)	Yes	1	Yes
Capture4	Yes (invalid TA)	Yes	0	Yes
Capture5	Yes (no TA)	Yes	0	No

Dynamic Host Configuration Protocol (DHCP) Overview

This topic provides overview information about the Dynamic Host Configuration Protocol (DHCP) Server role in Windows Server® 2012 and Windows Server 2012 R2. Also see the following topics:

- What's New in DHCP in Windows Server 2012 R2: Provides information about changes in functionality for DHCP in Windows Server 2012 R2.
- What's New in DHCP in Windows Server 2012: Provides information about changes in functionality for DHCP in Windows Server 2012.
- Understand and Deploy DHCP Failover: Provides detailed information about a new option for DHCP high availability in Windows Server 2012.
- Step-by-Step: Configure DHCP Using Policy-based Assignment: Provides step by step instructions for configuring DHCP policies in a test lab with Windows Server 2012.
- Step-by-Step: Configure DHCP for Failover: Provides step by step instructions for configuring DHCP failover in a test lab with Windows Server 2012.
- Windows Server 2012 Virtual Labs combines IP address management (IPAM) in a virtual lab with DNSSEC and DHCP failover in Windows Server 2012. See Building a Resilient Network Infrastructure to start the virtual lab. Note: The virtual lab will take some time to start.

Role description

DHCP is an Internet Engineering Task Force (IETF) standard designed to reduce the administration burden and complexity of configuring hosts on a TCP/IP-based network, such as a private intranet. Using the DHCP Server service, the process of configuring TCP/IP on DHCP clients is automatic.

Windows Server 2012 and Windows Server 2012 R2 provide several enhancements to the DHCP Server service.

Practical applications

DHCP in Windows Server 2012 and Windows Server 2012 R2 provide the following:

1. Safe, easy and reliable: Valid configuration parameters for all clients on the network can be dynamically provided. DHCP clients renew their leased address allocation automatically in the background. With DHCP failover, a new feature in Windows Server 2012, DHCP servers can provide a highly resilient DHCP service that allows the DHCP client to extend the lease on its current IP address by contacting another DHCP server on the enterprise network.
2. Flexible assignment: DHCP supports a large and extensible set of client configuration parameters. With client reservations, it is possible to reserve a specific IP address for permanent use by a DHCP client. A new feature in Windows Server 2012 called policy based assignment allows for even greater flexibility.
3. DHCP integration with DNS enables DNS resource records to be dynamically updated for new computers and devices, or when IP addresses of existing devices change on the network.

New and changed functionality

- What's New in DHCP in Windows Server 2012 R2: Provides information about changes in functionality for DHCP in Windows Server 2012 R2.
- What's New in DHCP in Windows Server 2012: Provides information about changes in functionality for DHCP in Windows Server 2012.

Server Manager information

Installation of the DHCP Server role can be performed using Server Manager. The following features and tools are installed automatically when you install DHCP Server:

Feature or Tool	Description
Remote Server Administration Tools	DHCP Server Tools are required to manage the DHCP Server role, but do not have to be installed on the same server. The DHCP Management console is installed automatically when you install DHCP Server unless you choose to cancel installation of Remote Server Administration Tools.

See Also

Step-by-Step: Configure DHCP for Failover

Step-by-Step: Configure DHCP Using Policy-based Assignment

What's New in DHCP in Windows Server 2012 R2

Dynamic Host Configuration Protocol (DHCP) in Windows Server® 2012 R2 provides new features and capabilities over previous versions. This document describes new deployment, manageability, and capabilities added to the DHCP Server role in Windows Server 2012 R2.

Role/Feature description

DHCP is an Internet Engineering Task Force (IETF) standard designed to reduce the administration burden and complexity of configuring hosts on a TCP/IP-based network, such as a private intranet. Using the DHCP Server service, the process of configuring TCP/IP on DHCP clients is automatic. For more information, see Dynamic Host Configuration Protocol (DHCP) Overview.

New and changed functionality

New functionality in DHCP Server for Windows Server 2012 R2 includes the following.

Feature/functionality	New or improved	Description
DNS registration enhancements	New	You can use DHCP policies to configure conditions based on the fully qualified domain name (FQDN) of DHCP clients, and to register workgroup computers using a guest DNS suffix.
DNS PTR registration options	New	You can enable DNS registration of address (A) and pointer (PTR) records, or just enable registration of A records.
Windows PowerShell for DHCP server	Improved	New Windows PowerShell cmdlets are available.

DNS registration enhancements

DHCP policies have been extended to allow users to configure conditions based on the fully qualified domain name (FQDN) of clients. DHCP policies can also be configured to register

DHCP clients using a specific DNS suffix, overriding the DNS suffix that is configured on the client.

Using suitable FQDN-based conditions and a DNS suffix, DHCP policies can be configured to enable full control of DNS registration for computers and devices on the network, including workgroup computers and guest devices, or clients with a specific attribute.

DNS PTR registration options

This feature enables you to configure a DHCP server to register only address (A) resource records of DHCP clients with the DNS server. This can be helpful when a reverse lookup zone has not been configured, causing attempts to register pointer (PTR) resource records to fail. By disabling PTR registration, DHCP servers can be prevented from repeated failed attempts to register PTR records. PTR registration can be disabled for all clients of a DHCP server, or only for clients on a specified subnet or with a specified attribute.

Previously, it was possible to disable DNS registration for a specified group of DHCP clients. However, it was necessary to disable both A and PTR registration.

Windows PowerShell for DHCP server

New Windows PowerShell cmdlets are available in Windows Server 2012 R2. You can use these cmdlets to perform tasks such as creating DHCP security groups, setting DNS credentials, managing superscopes, and managing multicast scopes. Improvements have also been made in existing Windows PowerShell cmdlets for DHCP server to help perform tasks such as changing the mode of a failover relationship and configuring DNS settings for a DHCP policy.

Cmdlet name	New or improved	Description
Add-DhcpServerSecurityGroup	New	Adds security groups to a DHCP server.
Add-DhcpServerv4MulticastExclusionRange	New	Adds a range of addresses to exclude from a multicast scope.
Add-DhcpServerv4MulticastScope	New	Adds a multicast scope on the DHCP server.
Add-DhcpServerv4Policy	Improved	Adds a new policy either at the server level or at the scope level. This cmdlet can now be used to specify lease duration and also add FQDN-based policies.

Cmdlet name	New or improved	Description
Get-DhcpServerDnsCredential	New	Gets an account that the DHCP Server service uses to register or deregister client records on a DNS server.
Get-DhcpServerv4DnsSetting	Improved	This cmdlet can now be used to display DNS settings of DHCP policies.
Get-DhcpServerv4MulticastExclusionRange	New	Retrieves the exclusion range for a specified multicast scope.
Get-DhcpServerv4MulticastLease	New	Retrieves multicast leases for a specified scope name.
Get-DhcpServerv4MulticastScope	New	Gets multicast scope objects.
Get-DhcpServerv4MulticastScopeStatistics	New	Gets multicast scope statistics.
Get-DhcpServerv4SuperscopeStatistics	New	Returns statistics for superscopes.
Remove-DhcpServerDnsCredential	New	Removes the credential that the DHCP Server service uses to register or deregister client records on a DNS server.
Remove-DhcpServerv4MulticastExclusionRange	New	Removes a range of addresses previously excluded from a multicast scope.
Remove-DhcpServerv4MulticastLease	New	Removes one or more multicast scope leases for a specified multicast scope or IP address.
Remove-DhcpServerv4MulticastScope	New	Removes multicast scopes.
Rename-DhcpServerv4Superscope	New	Renames a superscope.

Cmdlet name	New or improved	Description
Repair-DhcpServerv4IPRecord	New	Reconciles inconsistent lease records in the DHCP database.
Set-DhcpServerDnsCredential	New	Sets credentials that the DHCP Server service uses to register or deregister client records on a DNS server.
Set-DhcpServerv4DnsSetting	Improved	Configures how the Dynamic Host Configuration Protocol (DHCP) server service updates the DNS server with the client-related information. Ths cmdlet can now be used to set the DNS settings of policies.
Set-DhcpServerv4Failover	Improved	Modifies the attributes of an existing failover relationship. This cmdlet can now be used to change the mode of a failover relationship.
Set-DhcpServerv4MulticastScope	New	Modifies the properties of a multicast scope.
Set-DhcpServerv4Policy	Improved	Sets the properties of an existing policy either at the server level or at the specified scope level. This cmdlet can now be used to set lease duration of a policy and also modify FQDN-based policies.

For more information, see DHCP Server Cmdlets in Windows PowerShell.

What's New in DHCP in Windows Server 2012

Dynamic Host Configuration Protocol (DHCP) in Windows Server® 2012 provides new features and capabilities over previous versions. This document describes new deployment, manageability, and capabilities added to the DHCP Server role in Windows Server 2012.

Role/Feature description

DHCP is an Internet Engineering Task Force (IETF) standard designed to reduce the administration burden and complexity of configuring hosts on a TCP/IP-based network, such as a private intranet. Using the DHCP Server service, the process of configuring TCP/IP on DHCP clients is automatic. For more information, see Dynamic Host Configuration Protocol (DHCP) Overview.

New and changed functionality

New functionality in DHCP Server for Windows Server 2012 includes the following.

- DHCP failover: This feature provides the ability to have two DHCP servers serve IP addresses and option configuration to the same subnet or scope, providing for continuous availability of DHCP service to clients. The two DHCP servers replicate lease information between them, allowing one server to assume responsibility for servicing of clients for the entire subnet when the other server is unavailable. It is also possible to configure failover in a load-balancing configuration with client requests distributed between the two servers in a failover relationship. For more information about DHCP failover, see Step-by-Step: Configure DHCP for Failover
- Policy based assignment: With policy based assignment, the DHCP the server evaluates DHCP requests against policies that you define. Policies are applicable for a specific scope with a defined processing order. Policies applicable at a scope can be configured at the scope or inherited from server wide policies. A single client request can match multiple policies, and policies can be associated with multiple address ranges. For more information about DHCP policy based assignment, see Step-by-Step: Configure DHCP Using Policy-based Assignment.
- Windows PowerShell for DHCP Server: Windows Server 2012 provides task-oriented Windows PowerShell cmdlets for DHCP server management. The DHCP server configuration operations are implemented in Windows Server 2012 as Common Information Model (CIM) based cmdlets. Windows PowerShell uses object pipelining to eliminate the need for parsing and manipulation of text output. Windows PowerShell for DHCP Server includes 103 cmdlets, covering multiple aspects of DHCP server configuration and management.

See Also

Step-by-Step: Configure DHCP Using Policy-based Assignment

Step-by-Step: Configure DHCP for Failover

Understand and Deploy DHCP Failover

DHCP failover is a new feature available in Windows Server® 2012 ensuring continuous availability of DHCP service to clients.

With DHCP failover, two DHCP servers share DHCP scope and lease information, enabling one server to provide DHCP leases to DHCP clients if the other server is unavailable.

This guide provides an overview of DHCP failover, including deployment information, and steps to migrate from split scope or clustered DHCP to DHCP failover.

DHCP high availability options

In Windows Server® 2008 R2 and previous operating systems, two high availability options are available for DHCP server deployment:

- **DHCP in a Windows failover cluster**. This option places a DHCP server in a cluster with another DHCP server that assumes the load if the primary DHCP server fails. The clustering deployment option uses a single shared storage. The storage unit is a single point of failure, and requires an additional investment to achieve redundancy. Clustering can be relatively complex to configure and maintain.
- **Split scope DHCP**. Split scope DHCP uses two independent DHCP servers that share responsibility for a scope. Typically, 70% of the addresses in the scope are assigned to the primary server and the remaining 30% are assigned to a backup server. If clients cannot reach the primary server then they can acquire an IP address lease from the secondary server. Split scope deployment does not provide IP address continuity and is unusable in scenarios where the DHCP scope is already running with high address space utilization, a common scenario for IPv4 addressing.

In Windows Server® 2012 and later operating systems, a new DHCP high availability option is also available:

- **DHCP failover**. DHCP failover provides redundancy and load balancing for DHCP services, enabling administrators to deploy a highly resilient DHCP service. With DHCP failover, Internet Protocol version 4 (IPv4) scopes with associated leases, reservations, options, and settings, are shared by two Windows DHCP servers. DHCP failover supports large scale DHCP deployments without the challenges of split scope and clustered DHCP.

  Benefits of DHCP failover include:

  Simple: A wizard is provided to create DHCP failover relationships between DHCP servers. The wizard automatically replicates scopes and settings from the primary server to the failover partner.

  Flexible: DHCP failover can be configured to provide redundancy in hot standby mode, or with load balancing mode, client requests can be distributed between two DHCP servers.

Seamless: DHCP servers share lease information, allowing one server to assume responsibility for servicing of clients if the other server is unavailable. DHCP clients can keep the same IP address when a lease is renewed, even if the lease is issued by a different DHCP server.

Muti-site: DHCP failover supports a deployment architecture that includes multiple sites. DHCP failover partner servers do not need to be located at the same physical site.

Note
Microsoft's implementation of DHCP failover is based on the Internet Engineering Task Force (IETF) DHCP Failover Protocol draft.

In this guide

The following topics are available in this guide:

Topic	Description
What is DHCP Failover?	An overview of DHCP failover.
DHCP Failover Relationships	Details about how two DHCP servers share scope information and settings.
DHCP Failover Modes	Information about the two types of deployment options for DHCP failover.
DHCP Failover Communications	Information about messaging between DHCP failover partner servers.
DHCP Failover Settings	Details about individual settings used for DHCP failover.
DHCP Failover Examples	Examples of how DHCP failover works to maintain network continuity for DHCP clients during a service outage.
DHCP Failover Events and Performance	Information about new DHCP failover related events and performance counters.
Deploy DHCP Failover	DHCP failover design information and configuration steps.
DHCP Failover Requirements	Software and hardware requirements to deploy DHCP failover.
DHCP Failover Architecture	Design options for DHCP failover.
Checklist: Deploy DHCP Failover	A list of steps to deploy DHCP failover.
Configure DHCP Failover using the Command	Procedures to configure and deploy DHCP

Topic	Description
Line	failover with Windows PowerShell.
Configure DHCP Failover using the DHCP Console	Procedures to configure and deploy DHCP failover.
Replicate DHCP Failover Using the Command Line	Procedures to replicate DHCP settings between DHCP failover servers using Windows PowerShell.
Replicate DHCP Failover Using the DHCP Console	Procedures to replicate DHCP settings between DHCP failover servers using the DHCP console.
Migrate to DHCP Failover	Steps and recommendations to migrate from spit scope or clustered DHCP to DHCP failover.

See Also

Step-by-Step: Configure DHCP for Failover

What is DHCP Failover?

DHCP failover in Windows Server® 2012 is a new feature that enables two Microsoft DHCP servers to share service availability information with each other, providing DHCP high availability. DHCP failover works by replicating IP address leases and settings in one or more DHCP scopes from a primary DHCP server to a failover partner server.

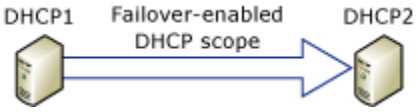

All scope information is shared between the two DHCP servers, including active leases. This enables either DHCP server to assume responsibility for DHCP clients if the other server becomes unavailable.

See the following sections in this topic for an overview of DHCP failover:

- Introduction to DHCP failover
- DHCP failover specifications
- DHCP failover and IPv6
- DHCP failover and Windows Failover Clustering
- DHCP failover and DNS dynamic updates

The following topics in this guide provide more detailed information about DHCP failover concepts, examples, and troubleshooting:

- DHCP Failover Relationships
- DHCP Failover Modes
- DHCP Failover Communications
- DHCP Failover Settings
- DHCP Failover Examples
- DHCP Failover Events and Performance

Introduction to DHCP failover

With DHCP failover, DHCPv4 scopes can be replicated from a primary DHCP server to a partner DHCP server, enabling redundancy and load balancing of DHCP services. DHCP servers that share a failover-enabled DHCP scope are called failover partners. Microsoft's implementation of DHCP failover is based on the Internet Engineering Task Force (IETF) DHCP Failover Protocol draft.

When two DHCP servers are configured for failover, they will share scope information, including all active leases. This enables both DHCP servers to provide leases to the same subnet for load balancing or redundancy purposes. Scope settings are replicated when you first configure DHCP failover, and can be replicated again later if configuration changes are made.

The following figure illustrates how components and settings for a failover-enabled DHCP scope are shared between two DHCP servers.

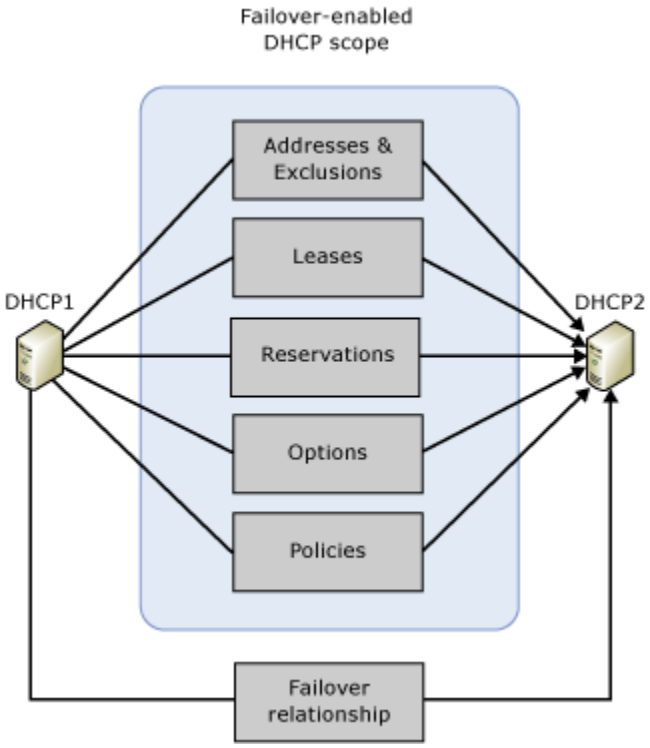

The scopes and settings used with DHCP servers that are configured for DHCP failover are shared using a new object called the DHCP failover relationship. For more information, see DHCP Failover Relationships.

Several configuration options are available with DHCP failover. You can configure failover on all scopes that exist on a DHCP server, or only on some scopes. You can also easily use the same DHCP failover settings for many scopes by adding them to the same failover relationship. A failover relationship is always between only two DHCP servers. However, a server can have many failover relationships, and each failover relationship can be with a different DHCP server.

Important
If changes are made to a failover-enabled scope, you must manually replicate these changes to the partner server in order to synchronize scopes on both DHCP servers. Replication copies scope settings from the DHCP server where replication is initiated to the partner server, overwriting settings on the partner server. Therefore, it is important to always initiate replication from the server that has DHCP scope settings you wish to use.

DHCP failover specifications

The following specifications apply to DHCP failover:

- You cannot configure DHCP failover on a DHCP scope to include more than two DHCP servers.
- DHCP failover supports DHCPv4 scopes only. DHCPv6 scopes cannot be failover-enabled.
- DHCP failover partners must both be running Windows Server 2012 or a later operating system.
- DHCP failover can be configured, and settings can be modified without the need to pause, stop, or restart the DHCP Server service.
- If parameters of a failover-enabled scope are modified, these settings must be manually replicated to the partner DHCP server. Note: Automatic replication of scope settings is available if you use IP address management (IPAM) in Windows Server 2012 R2 to modify failover-enabled scope settings.
- Replication of scope settings can be initiated from either DHCP server to its failover partner server.
- Clustered DHCP is supported in conjunction with DHCP failover. For purposes of failover, a DHCP cluster is considered a single DHCP server. See DHCP failover and Windows Failover Clustering for more information.
- DHCP clients must be able to communicate with both DHCP failover partner servers, either directly or using a DHCP relay.
- DHCP servers configured as failover partners can be located on different subnets, but this is not required.
- When DHCP failover is enabled on a DHCP scope, the DHCP server that renews a DHCP client lease can be different from the DHCP server that initially granted the lease.
- Two DHCP servers configured as failover partners will attempt to maintain a persistent TCP/IP connection.

- Two separate, synchronized client lease databases are maintained independently by each DHCP failover partner server.
- DHCP servers configured as failover partners are both aware of the status of the DHCP service on the other server, and are informed of any change in that status with minimal delay.
- If two DHCP servers configured as failover partners are unable to communicate, precautions are taken to avoid the same IP address lease being issued to two different DHCP clients.
- If a DHCP server becomes unavailable before it is able to successfully synchronize all DHCP client information with its failover partner, precautions are taken to ensure DHCP lease continuity for DHCP clients.

Important
When replicating settings between DHCP failover partner servers that have different operating system versions (ex: Windows Server 2012 and Windows Server 2012 R2), always modify settings and initiate replication from the DHCP server with the more recent operating system version. This ensures that settings are recognized by both failover partners and replicated consistently.

DHCP failover can be configured using Windows PowerShell or using the DHCP console. For information about using Windows PowerShell, see DHCP Server Cmdlets in Windows PowerShell (http://technet.microsoft.com/library/jj590751.aspx). For detailed guidance on configuring DHCP failover, see Deploy DHCP Failover.

DHCP failover and IPv6

DHCP failover is not supported for Internet Protocol version 6 (IPv6) scopes. Network adapters using IPv6 typically determine their own IPv6 address using stateless IP auto-configuration. In this mode, the DHCP server delivers only the DHCP option configuration, and the server does not maintain any lease state information. A high availability deployment for stateless DHCPv6 is possible by simply setting up two servers with identical option configuration. Even in a stateful DHCPv6 deployment, the scopes do not run under high address utilization, which makes split scope a viable solution for high availability.

DHCP failover and Windows Failover Clustering

DHCP failover is supported with clustered DHCP in the following configurations:

- A single DHCP server can have a failover relationship with a DHCP failover cluster.
- A DHCP failover cluster can have a failover relationship with another DHCP failover cluster.

In both cases, you must configure DHCP failover to use the name or IP address of the cluster, not the name or IP address of an cluster node. If an individual cluster node is configured as the failover partner, the primary server will enter a communications interrupted state if the DHCP Server service moves to a different node in the cluster.

Important

If you use a shared secret, you must manually replicate the shared secret to all cluster nodes. See the following procedure for more information.

Replicate a shared secret on DHCP failover cluster nodes

1. On the Server Manager menu, click **Tools** and then click **Failover Cluster Manager**.
2. In Failover Cluster Manager, determine the currently active node. To determine the active node, you can click on each node under **Nodes** until you see the **DHCP Server** role with a status of **Running**. You can also type the following command at a Windows PowerShell prompt:

   ```
   Get-ClusterResource –Name DHCPMCSDhcp | select OwnerNode
   ```

   In this example, the cluster name is **DHCPMCSDhcp**. Output from this command will be the name of the currently active node.
3. On the currently active cluster node, type the following command at an elevated Windows PowerShell prompt, and then press ENTER.

   ```
   Set-DhcpServerv4Failover –Name ”DHCPMCSDhcp”  -SharedSecret
   “123”
   ```

   In this example, the cluster name is **DHCPMCSDhcp** and the shared secret is **123**.
4. Next, move the DHCP Server role to another node. To move a role, click **Action** on the Failover Cluster Manager menu, point to **Move**, and then click **Select Node**.
5. In the **Move Clustered Role** dialog box, click on a node under **Cluster nodes** and then click **OK**.
6. Repeat steps 3-5 until you have set the shared secret on all nodes in the cluster.

The reason you must replicate the shared secret is because it is encrypted in the Windows Registry.

Windows Failover Clustering will automatically replicate registry parameters related to the DHCP Server service when it moves to a new node, but encryption of a shared secret is specific to each node. Therefore, each node can only decrypt a shared secret if it was configured on the same node.

DHCP failover and DNS dynamic updates

If DHCP servers are configured to perform DNS dynamic updates on behalf of the client computer, both DHCP servers in a DHCP failover relationship must use the same DNS credentials to update DNS records. If the failover partner attempts to use different credentials to update DNS resource records, this update will fail.

The following steps describe how DNS dynamic updates might fail when a client computer uses a different DHCP server:

1. A Windows DHCP server performs a dynamic update on behalf of a DHCP client.
2. The DHCP server creates the client’s DNS name and becomes the owner of that name.
3. Now only the DHCP server itself can update the DNS records for the client’s name.

4. The original server fails and a second backup DHCP server comes online; now the second server cannot update the client name because it is not the name's owner.

Also see DNS Record Ownership and the DnsUpdateProxy Group for a discussion of this scenario.

See also

Deploy DHCP Failover

Step-by-Step: Configure DHCP for Failover

DHCP Failover Relationships

See the following sections in this topic for more information about working with failover relationships.

Understand failover relationships

Create failover relationships

View failover relationships

Edit failover relationships

Delete failover relationships

Understand failover relationships

Failover relationships are settings used by two DHCP servers to share one or more failover-enabled DHCP scopes.

A single failover relationship can be associated to multiple DHCP scopes. You do not need to create a new failover relationship every time you enable failover on a DHCP scope, unless a new DHCP server will be used as the failover partner or a new set of failover settings are required. It is recommended to reuse existing failover relationships whenever possible to avoid unnecessary duplication of failover settings.

Characteristics of failover relationships include the following.

- A single DHCP server can have no more than 31 failover relationships.
- A single failover relationship is always shared between exactly two DHCP servers.
- Multiple failover relationships can exist between the same two DHCP servers.
- Multiple DHCP scopes can use the same failover relationship.
- A single DHCP server can have failover relationships with multiple other DHCP servers.
- A single DHCP scope can only be associated to one failover relationship at a time. However, scopes can be removed from one failover relationship and added to a different failover relationship.
- A failover relationship can exist that is not associated to any DHCP scopes.

- Failover relationships can exist between two DHCP servers, a single DHCP server and clustered DHCP, or between two DHCP clusters.
- Deleting a failover relationship will remove failover from all scopes associated with that relationship.
- It is not necessary to pause, stop, or restart the DHCP Server service when adding, deleting, or modifying failover relationships.
- Failover relationships do not support IPv6 scopes.

See the following example.

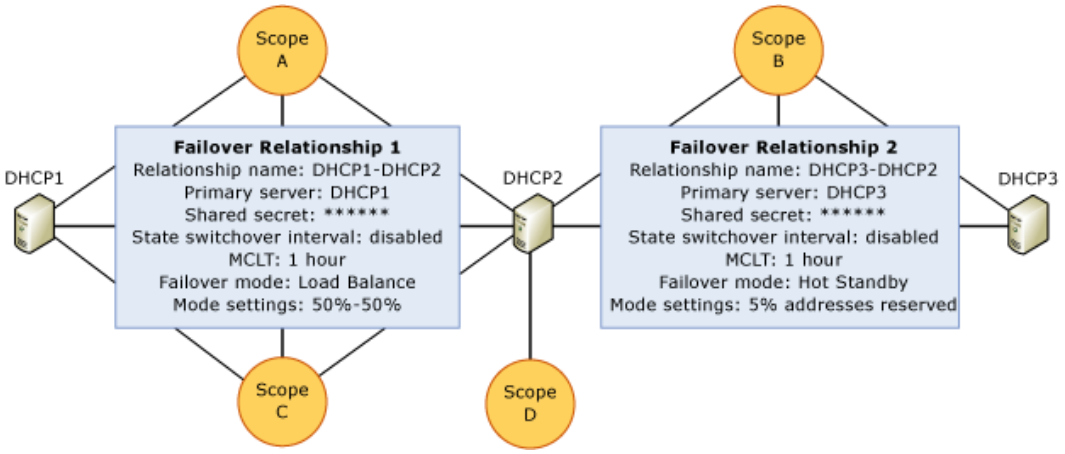

In this example, scopes A and C are configured on DHCP1 and DHCP2 to use a 50-50 load balancing failover relationship named “DHCP1-DHCP2.” Scope B is configured on DHCP2 and DHCP3 and is associated to the hot standby failover relationship “DHCP3-DHCP2” wherein DHCP3 is the active server and DHCP2 is the standby server. Scope D on DHCP2 not currently configured for failover.

Failover relationship 1 can be viewed and edited on DHCP1 or DHCP2. Likewise, failover relationship 2 can be viewed and edited on DHCP2 or DHCP3. The name of a failover relationship is established when it is created, and must be unique for both servers. You can type a custom name, or use the default name that is supplied.

Important

A single DHCP server can support up to 31 failover relationships. To avoid exceeding this limit, reuse existing failover relationships when possible and delete failover relationships that are not in use.

Failover relationships can exist with no association to a DHCP scope. This occurs if you deconfigure failover on the last DHCP scope associated to a failover relationship, and do not also delete the failover relationship. You might do this if you plan to reuse the failover relationship later.

Create failover relationships

To create a failover relationship on a DHCP server, you can type **Add-DhcpServerv4Failover** at an elevated Windows PowerShell prompt, or you can use the DHCP console. If you use the DHCP console, you can right-click **IPv4**, or right-click an IPv4 scope and then click **Configure Failover** to start the Configure Failover wizard.

The following minimum information is required to create a new failover relationship:

1. A unique name for the failover relationship
2. A network-connected partner DHCP server name or IP address
3. At least one scope ID for a DHCP scope that exists on the local server

Additional parameters are optional. The default DHCP failover mode is load balance with 50% of leases being assigned by the local server and 50% being assigned by the partner server. Authentication is also enabled by default, but this setting is not required. If authentication is enabled, you must also supply a shared secret. The scope ID provided must correspond to a scope configured on the local server that is not already failover-enabled, and does not already exist on the partner server. If DHCP failover is being configured remotely, the “local server” refers to the server specified on the command line or added to the management console.

You cannot enable DHCP failover on a DHCP scope if the scope already exists on both DHCP servers in a failover relationship, as might occur in a split scope configuration. If a DHCP scope exists on the partner DHCP server it must be deleted prior to configuring DHCP failover. See the following examples:

Example 1: Configure DHCP failover using Windows PowerShell.

```
PS C:\> Add-DhcpServerv4Failover –Name dc1-dhcp1 –ScopeID 192.168.10.0 –PartnerServer
dhcp1.contoso.com
```

Example 2: Configure DHCP failover using the DHCP console.

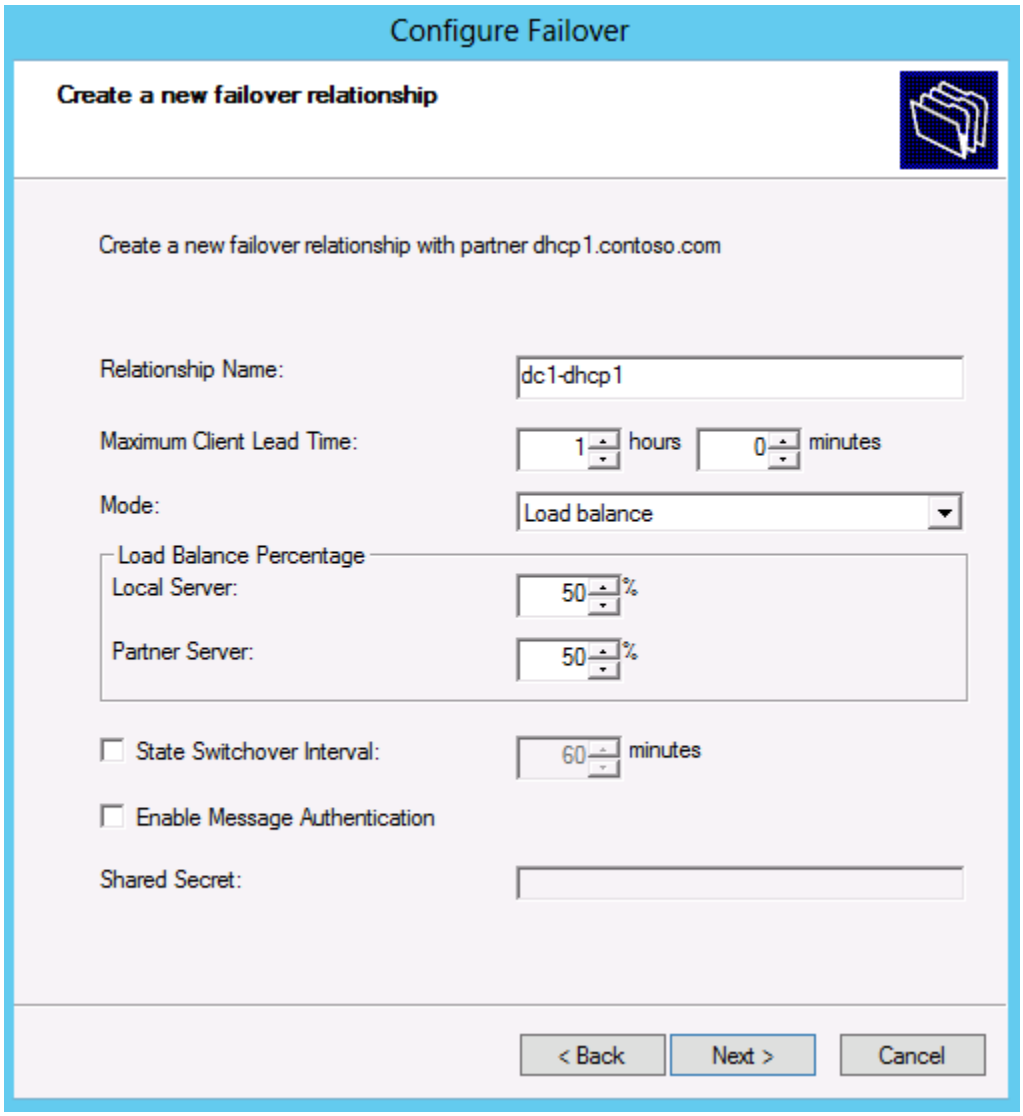

In the previous examples, a new failover relationship named **dc1-dhcp1** is established between the local server and dhcp1.contoso.com to use a 50-50 load balancing percentage for the DHCP scope with ID 192.168.10.0.

For more information about additional settings used to configure failover relationships, see Deploy DHCP Failover.

View failover relationships

To view failover relationships on a DHCP server, you can type **Get-DhcpServerv4Failover** at an elevated Windows PowerShell prompt, or you can open the DHCP console, right-click **IPv4**, click **Properties**, and then click the **Failover** tab. See the following examples:

Example 1: View DHCP failover using Windows PowerShell.

```
PS C:\> Get-DhcpServerv4Failover



Name                : dc1.contoso.com-192.168.0.254
PartnerServer       : 192.168.0.254
Mode                : HotStandby
LoadBalancePercent  :
ServerRole          : Active
ReservePercent      : 5
MaxClientLeadTime   : 01:00:00
StateSwitchInterval :
State               : Normal
ScopeId             : {192.168.50.0, 192.168.0.0}
AutoStateTransition : False
EnableAuth          : False

Name                : dhcp1-dc1 load balance
PartnerServer       : dhcp1.contoso.com
Mode                : LoadBalance
LoadBalancePercent  : 50
ServerRole          :
ReservePercent      :
MaxClientLeadTime   : 01:00:00
StateSwitchInterval :
State               : Normal
ScopeId             :
AutoStateTransition : False
EnableAuth          : False
```

Example 2: View DHCP failover using the DHCP console.

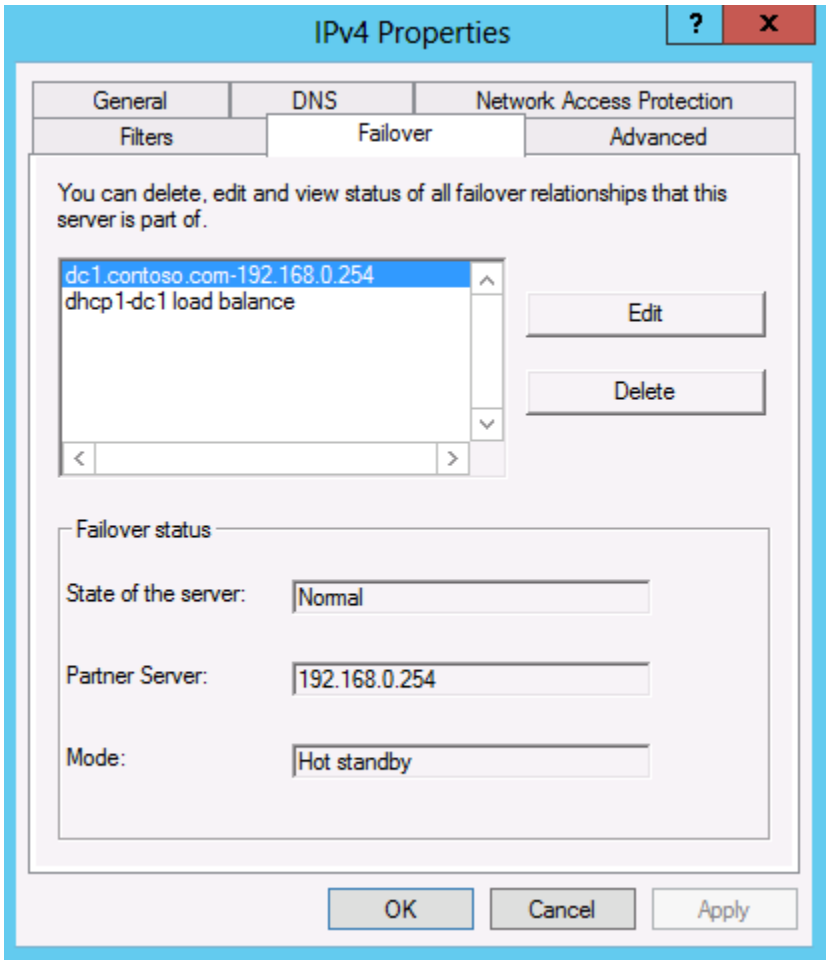

An advantage to using Windows PowerShell to view failover relationships is that you can also view the DHCP scopes that are associated with each failover relationship. This information is provided next to the **ScopeId** parameter.

Edit failover relationships

To edit failover relationships on a DHCP server, you can type **Set-DhcpServerv4Failover** at an elevated Windows PowerShell prompt, or you can open the DHCP console, right-click **IPv4**, click **Properties**, and then click the **Failover** tab. See the following examples:

Example 1: Edit DHCP failover using Windows PowerShell.

```
PS C:\> Set-DhcpServerv4Failover -Name "dhcp1-dc1 load balance" -LoadBalancePercent 70
```

In the previous example, the load balancing percentage on the local server was changed to 70%.

Example 2: To edit properties of the failover relationship using the DHCP console, choose a failover relationship and then click **Edit**.

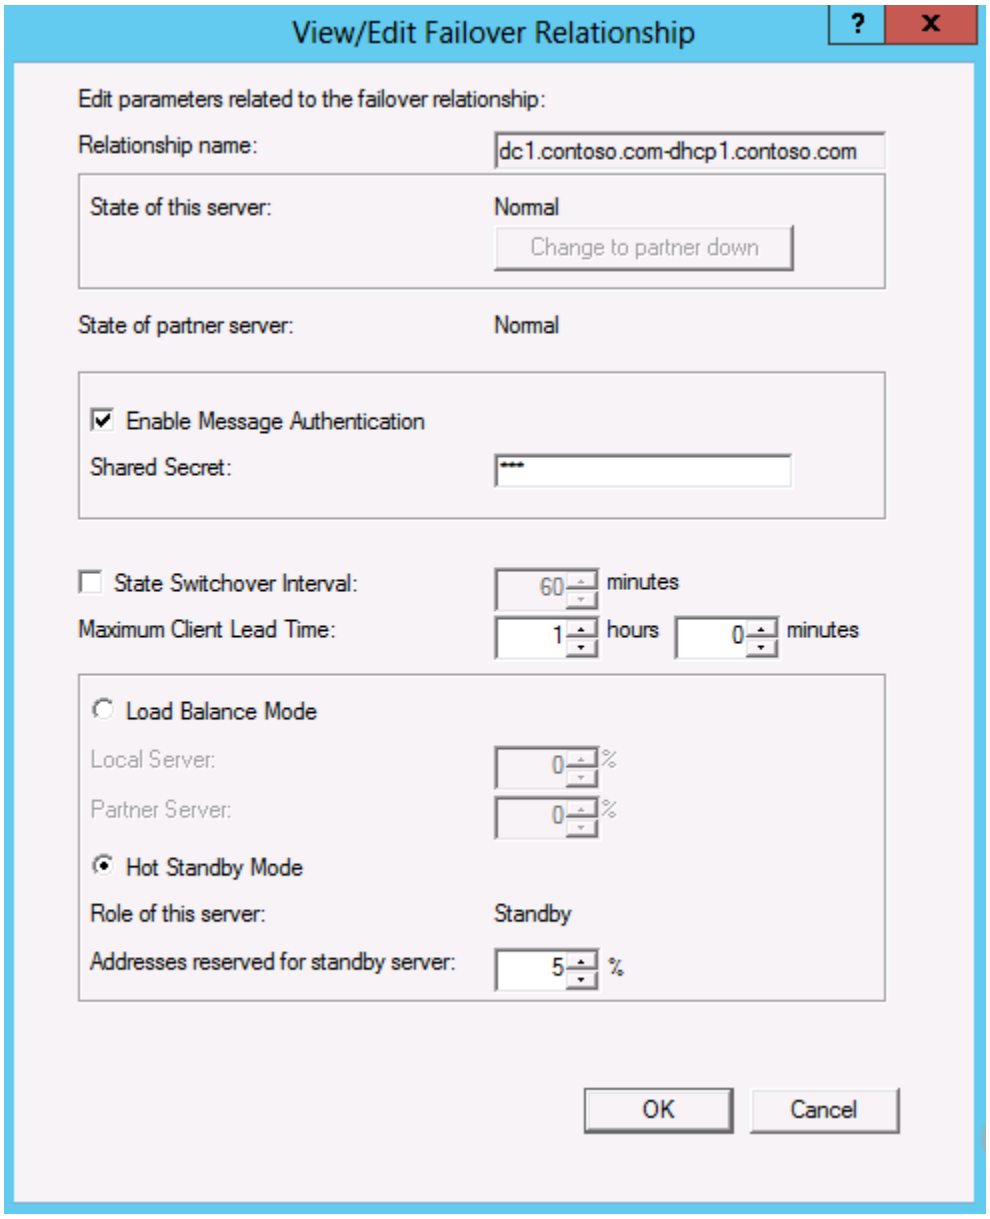

For more information about configuring DHCP failover relationships, see Deploy DHCP Failover.

Delete failover relationships

To delete a failover relationship on a DHCP server, you can type **Remove-DhcpServerv4Failover** at an elevated Windows PowerShell prompt, or you can use the DHCP console. You must supply the failover relationship name. See the following example:

Example 1: Delete a DHCP failover relationship using Windows PowerShell.

```
PS C:\> Remove-DhcpServerv4Failover -Name dhcp1.contoso.com-dhcp2-dns1.contoso.com
```

To delete a failover relationship using the DHCP console, right-click IPv4 and then click **Properties**. Highlight the relationship name on the **Failover** tab and then click **Delete**. You will be prompted to confirm the deletion, and if any scopes are currently associated with the failover relationship, these will be listed. See the following example:

Example 2: Delete a DHCP failover relationship using the DHCP console.

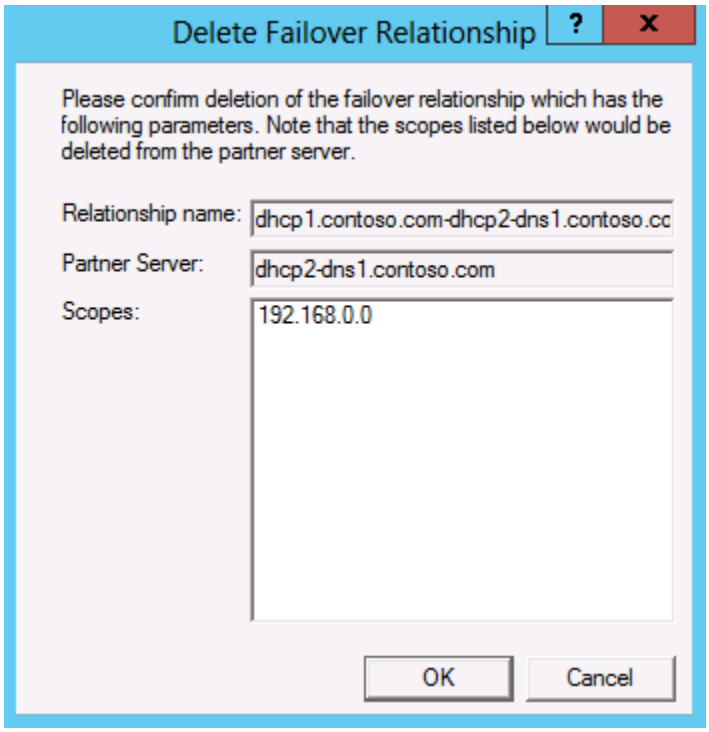

Click OK to confirm deletion of the failover relationship. Deleting a failover relationship that has DHCP scopes associated with it will remove those scopes from the partner DHCP server. The scopes will remain on the DHCP server where the failover relationship was deleted.

Note

If you remove the last DHCP scope from a failover relationship (by “deconfiguring” failover on the scope) using the DHCP console, you will be notified that the failover relationship can be deleted. The failover relationship is not automatically removed. If you do not plan to reuse the relationship later, manually delete this relationship from either the primary DHCP server or the failover partner.

See also

DHCP Failover Modes

DHCP Failover Communications

DHCP Failover Settings

Deploy DHCP Failover

DHCP Failover Modes

Two DHCP failover modes are available to use when you create a DHCP failover relationship:

Hot standby mode: This mode provides redundancy for DHCP services.

Load balance mode: This mode allocates DHCP client leases across two servers.

Important
Redundancy for DHCP is also provided in a load-balancing configuration. DHCP clients can renew their leases with the failover partner server if their assigned DHCP server does not respond.

You can switch between hot standby and load balance mode if desired, but you can only use one mode at a time with a single DHCP scope.

For details about settings used for hot standby and load sharing, see DHCP Failover Settings.

For DHCP failover examples using hot standby and load sharing modes, see DHCP Failover Examples.

For information about deploying DHCP failover in hot standby and load sharing mode, see DHCP Failover Architecture.

Hot standby mode

In hot standby mode, two servers operate in a failover relationship where an active server is responsible for leasing IP addresses and configuration information to all clients in a scope or subnet. The partner server assumes a standby role, with responsibility to issue leases to DHCP clients only if the active server becomes unavailable. Hot standby mode is ideal for scenarios where the failover partner is only intended to be used temporarily when the active server is unavailable.

A server is active or standby in the context of a failover relationship. For instance, a server that has the role of active for a given relationship could be a standby server for another relationship.

By default, the server that is used to create the failover relationship is the active server, but this is not required.

When you choose hot standby, you must also configure the percentage of IP addresses on the active server that are reserved for use on the standby server in the event that the active server does not respond. By default, this reserve percentage is 5%.

The reserve percentage is used for new DHCP leases. If a DHCP client attempts to renew a DHCP lease with the standby server that is unable to contact the active server (COMMUNICATIONS INTERRUPTED or PARTNER DOWN state exists), the same IP address that was previously assigned to the DHCP client will be renewed. In this situation, a temporary lease is granted for the maximum client lead time (MCLT) duration, not the full scope lease time.

In a COMMUNICATIONS INTERRUPTED or PARTNER DOWN state, if the standby server issues all its available reserve percentage leases to new DHCP clients before the MCLT expires, it will refuse to issue new DHCP leases, but it will continue to renew existing leases. After the MCLT has expired, the standby server will be permitted to use the entire available IP address pool for new DHCP leases, provided the server is in PARTNER DOWN state. If the server is still in COMMUNICATIONS INTERRUPTED state, it will not use the entire available IP address pool for new DHCP leases.

For more information about DHCP failover settings and states, see DHCP Failover Settings.

Load balance mode

Load balance mode is the default mode of deployment. In this mode, two DHCP servers simultaneously serve IP addresses and options to clients on a given subnet. DHCP client requests are load balanced and shared between the two DHCP servers. The default load balancing ratio between the two servers is 50:50, but this can be customized to any ratio from 0 to 100%.

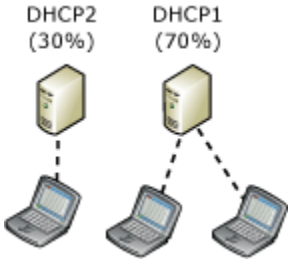

The load-balancing mechanism is defined in RFC 3074, in which a hash is computed from the MAC address contained in each DHCP client request. A range of hash values (also called the hash bucket) is assigned to each DHCP server based on the load balancing percentages that are configured. Servers determine if they are designated to respond to the client based on their assigned hash bucket.

In load balancing mode, when a DHCP server loses contact with its failover partner it will begin granting leases to all DHCP clients. If it receives a lease renewal request from a DHCP client that

is assigned to its failover partner, it will temporarily renew the same IP address lease for the duration of the MCLT. If it receives a request from a client that was not previously assigned a lease, it will grant a new lease from its free IP address pool until this is exhausted, and then it will begin using the free IP address pool of its failover partner. If the DHCP server enters a partner down state, it will wait for the MCLT duration and then assume responsibility for 100% of the IP address pool.

See also

What is DHCP Failover?

DHCP Failover Examples

DHCP Failover Settings

DHCP Failover Communications

Both DHCP servers in a failover relationship must maintain a persistent TCP connection with each other. DHCP failover partners establish and maintain this connection on port 647, and use it to exchange operational state information and lease information.

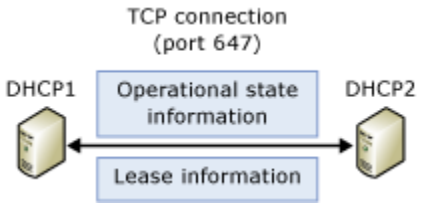

The primary DHCP server will initiate one connection for every DHCP failover relationship that is configured on the server. The connection is used to ensure lease information is consistent across both servers, and to enable each server to monitor the operational state of its partner. DHCP servers are notified if the connection is broken for any reason.

To ensure a persistent TCP connection between DHCP failover partners is maintained, it is important to use a static IP address on all DHCP server network interfaces. If the static IP address of a DHCP server needs to be changed, for example during DHCP migration, you must first delete all DHCP failover relationships that exist on that server, and then recreate the relationships when the new IP address is active.

For examples of how DHCP failover partners communicate before and after service interruption, see DHCP Failover Examples.

DHCP failover messages

The following table provides a summary of DHCP failover message types defined by the Internet Engineering Task Force (IETF) DHCP Failover Protocol draft. A short description of each message is provided. For more information, see the IETF DHCP Failover Protocol draft.

Important
Microsoft's implementation of DHCP failover does not display DHCP servers as primary or secondary in the DHCP console interface. In the following section, **primary** refers to the DHCP server on which a DHCP failover relationship is first configured, or the server specified to be the **active** server in a hot standby configuration.

Type	Function	Description
BNDUPD	Update a failover partner with binding information	The BNDUPD (binding update) message is used to send binding database changes (also known as binding update transactions) to a partner server. The partner server responds with a BNDACK (binding acknowledgment) message when it has successfully committed the requested changes to its own database. BNDUPD and BNDACK messages can contain multiple binding update transactions. The BNDUPD message is also used to distribute addresses for address allocation. Note: Microsoft's implementation of DHCP failover does not use POOLREQ and POOLRESP messages as specified in the IETF DHCP failover protocol draft.
BNDACK	Acknowledge receipt of a binding update	A server sends a BNDACK (binding acknowledgment) message when it has processed a BNDUPD message and successfully

Type	Function	Description
		committed binding database changes (made as a result of processing the BNDUPD message). A BNDACK message might be used to accept or to reject a BNDUPD message. If a BNDACK message is used to reject a BNDUPD message, it will contain a reject reason. BNDUPD and BNDACK messages can contain multiple binding update transactions. A BNDACK message can only be sent in response to a BNDUPD message using the same TCP connection from which the BNDUPD message was received. If a connection to a partner server goes down, a server with unprocessed BNDUPD messages might discard these messages or it might process the messages. However, it will not send any BNDACK messages in response.
CONNECT	Establish a connection with the failover partner	The connect message is used to establish a TCP connection between failover partners. The connect message provides source information for the connection and required configuration information. It is the first message sent using a newly established connection, and it is only sent only by the primary server.
CONNECTACK	Respond to an attempt to establish a connection with the	The CONNECTACK (connect acknowledgment) message is

Type	Function	Description
	failover partner	sent to accept or reject a CONNECT message. It is sent by the secondary server which received a CONNECT message.
UPDREQALL	Request a full transfer of binding information	The UPDREQALL (update request all) message is used by one server to request that its partner send it all of the IP address lease information. This message is used to allow one server to restore its binding database in its entirety from the other server.
UPDDONE	Indicate that all binding update information has been sent	The UPDDONE (update done) message is used by a server receiving an UPDREQ or UPDREQALL message to signify that it has sent all of the requested BNDUPD messages and that it has received a corresponding BNDACK for each of those messages. A BNDACK message must be received for each BNDUPD message prior to the transmission of the UPDDONE message.
UPDREQ	Request transfer of new binding information	The UPDREQ (update request) message is used by one server to request that its partner send it all of the binding database information that it has not already seen. An UPDDONE message is sent after all UPDREQ messages have been received and processed.
STATE	Inform the failover partner of the current state or of a state	The state (STATE) message is used to communicate the

Type	Function	Description
	change	current failover state to the partner server. No response is required from the partner server. For more information about failover states, see DHCP Failover Settings.
CONTACT	Probe communications integrity with the failover partner	The CONTACT message is sent to verify communications integrity with a failover partner. The CONTACT message is sent when no messages have been sent to the failover partner for a default period of time.

See also

What is DHCP Failover?

DHCP Failover Examples

DHCP Failover Settings

DHCP Failover Modes

Deploy DHCP Failover

DHCP Failover Settings

The following table provides a summary of settings used with DHCP failover. Additional details about each setting are provided below.

Setting	Description	More Information
Name	Name of the failover relationship. The name of the failover relationship must be unique on the local server.	Name
Partner server	Name or IP address of the server that shares the failover relationship with the local DHCP server.	Partner server

Setting	Description	More Information
Mode	The type of DHCP failover mode chosen.	Mode
Load balance percentage	The percentage of DHCP client leases that will be evaluated by the local server.	Load balance percentage
Server role	The role of the local server in hot standby mode, either active or standby.	Server role
Reserve percentage	The percentage of IP addresses in a scope that are reserved for the hot standby server.	Reserve percentage
Maximum client lead time (MCLT))	The maximum amount of time that one server can extend a lease for a DHCP client beyond the time known by the partner server.	Maximum client lead time (MCLT)
State switchover interval	The interval after which a DHCP server automatically transitions its failover partner to a partner down state after loss of communication.	State switchover interval
State	The current state of the local server.	State
Scope IDs	A list of DHCP scope IDs that are associated with the failover relationship.	Scope IDs
Automatic state transition	If a state switchover interval is configured, then automatic state transition is enabled. By default, automatic state transition is disabled.	Automatic state transition
Enable authentication	Whether or not a shared secret is required between failover partners.	Enable authentication

Name

You must provide a unique name for each failover relationship on a DHCP server. The name must be unique on both failover partner servers. The relationship name is limited to 126 characters. Special characters (!@#$%^&*) and spaces are allowed, but if special characters are configured in a name using Windows PowerShell, the relationship name must be enclosed in quotes.

Partner server

You can provide the NETBIOS name, DNS name, or IP address of the partner DHCP server. If the partner DHCP server is a member of a DHCP cluster, provide the name or IP address of the cluster. If you configure the partner DHCP server to be a cluster node, contact with the partner server will be lost if the DHCP service is moved to a different node in the cluster.

If you use the DHCP console to configure DHCP failover, a drop-down list of previously configured partner servers is provided. You can also choose to reuse existing failover relationships by selecting the checkbox next to **Reuse existing failover relationships configured with this server (if any exist)**.

Mode

Choose load balance or hot standby mode. For more information see DHCP Failover Modes.

Load balance percentage

If you choose the load balance mode, you must provide the percentage of DHCP client requests that will be serviced by each DHCP server. By default, a 50-50 load balancing percentage is assigned.

Server role

If you choose the hot standby mode, you must designate the role of the partner server to be either active or standby. By default, the local server is active and the partner server is standby.

Reserve percentage

In a failover relationship configured in hot standby mode, administrators can specify a percentage of the address range of the scope as reserved for the hot standby server. A number of addresses, in proportion to the percentage value configured, are assigned to the hot standby server. The hot standby server will use these addresses to service new clients after the primary server goes down, during the time interval before the standby server assumes control over the entire IP address range of a scope. The hot standby server assumes control over the entire IP address range only after it transitions into partner down state and a certain time (defined by MCLT) has

elapsed after moving into the partner down state.If an administrator sets this parameter to zero, no addresses are reserved for the hot standby server, and the failover partner server cannot grant new client leases until the time that the hot standby assumes control over the entire IP address range. The default value for reserve address percentage is 5%.

Maximum client lead time (MCLT)

The maximum amount of time that one server can extend a lease for a DHCP client beyond the time known by the partner server. The MCLT defines the temporary lease period given by a failover partner server, and also determines the amount of time that a server in a failover relationship will wait in partner down state before assuming control over the entire IP address range.

The MCLT cannot be set to zero, and the default setting is 1 hour.

State switchover interval

If automatic state transition is enabled, a DHCP server in communication interrupted state will automatically transition to partner down state after a defined period of time. This period of time is defined by the state switchover interval.

A server that loses communication with a partner server transitions into a communication interrupted state. The loss of communication may be due to a network outage or the partner server may have gone offline. By default, since there is no way for the server to detect the reason for loss of communication with its partner, the server will continue to remain in communication interrupted state until the administrator manually changes the state to partner down. However, if you enable automatic state transition, DHCP failover will automatically transition to partner down state when the auto state switchover interval expires. The default value for auto state switchover interval is 60 minutes.

State

In Windows PowerShell, state refers to the condition of the local server, or the server specified using the **ComputerName** parameter. The DHCP console displays the state of both the local server and the partner server next to **State of this server** and **State of partner server** when you view or edit a DHCP failover relationship.

The following table describes available states. These states are defined in section 5 (Protocol Overview) and section 9 (Failover Endpoint States), in the IETF DHCP Failover Protocol draft.

State	Description
NORMAL	When operating in the NORMAL state, each DHCP server in a failover relationship services DHCP clients as defined by the failover mode that is configured and its associated

State	Description
	parameters.
COMMUNICATION INTERRUPTED	When operating in the COMMUNICATIONS INTERRUPTED state, each server functions independently. Servers do not assume that their partner is not operating, but only that they are unable to communicate. Each server responds to all DHCP client messages that it receives in a manner such that graceful reintegration its partner server is possible when communication is restored.
PARTNER DOWN	When operating in PARTNER DOWN state, a server assumes that its failover partner is not operating. The server responds to all DHCP client requests that it receives.
POTENTIAL CONFLICT	This state indicates that the two servers are attempting to reintegrate with each other, but at least one of them was running in a state that did not guarantee automatic reintegration would be possible. In the POTENTIAL CONFLICT state, the servers may determine that the same IP address has been offered and accepted by two different DHCP clients. It is a goal of DHCP failover to minimize occurrences of the POTENTIAL CONFLICT state. A server in the POTENTIAL CONFLICT state does not process any incoming DHCP requests.
STARTUP	A server operating in STARTUP state assumes that failover is operational, and it spends a short time attempting to contact the failover partner. The server is unresponsive to DHCP client requests during this time. This state exists in order to give the server a chance to determine if its partner has changed state since it was last in communications, and to react to that changed state (if any) prior to responding to DHCP client requests.
RESOLUTION INTERRUPTED	This state indicates that the two servers were attempting to reintegrate with each other in

State	Description
	POTENTIAL-CONFLICT state, but communications failed prior to completion of re-integration. A server in this state will respond to DHCP client requests using its own IP address pool.
CONFLICT DONE	This state indicates that the primary server has received all of the updates from the secondary server during the reintegration process. A server in the CONFLICT DONE state responds to all client requests, as opposed to the NORMAL state where it might be running under the load balancing algorithm.
RECOVER	A server operating in RECOVER state assumes that it is reintegrating with a server that has been operating in PARTNER DOWN state, and that it needs to update its bindings database before it services DHCP client requests. A server may also operate in RECOVER state in order to fully recover its bindings database from its partner server.
RECOVER WAIT	This state indicates that the server has done an UPDREQ or UPDREQALL and has received the UPDDONE message indicating that it has received all outstanding binding update information. In the RECOVER WAIT state the server will wait for the MCLT in order to ensure that any processing that this server might have done prior to losing its stable storage will not cause future difficulties. A server in RECOVER WAIT does not respond to DHCP client requests.
RECOVER DONE	This state exists to allow an interlocked transition for one server from RECOVER state and another server from PARTNER DOWN or COMMUNICATIONS INTERRUPTED state into NORMAL state. A server in RECOVER DONE state MUST respond only to DHCPREQUEST/RENEWAL and DHCPREQUEST/REBINDING DHCP messages. When a server in RECOVER DONE

State	Description
	state determines that its partner server has entered NORMAL or RECOVER DONE state, then it will transition into NORMAL state. If communications fail while in the RECOVER DONE state, a server will stay in the RECOVER DONE state.
NO STATE	If a server's state is unknown, it will be displayed as NO STATE.

Scope IDs

A list of DHCP scope IDs that are associated with the selected failover relationship.

Automatic state transition

If enabled, automatic state transition will occur from the COMMUNICATIONS INTERRUPTED state to PARTNER DOWN state when the state switchover interval expires.

Enable authentication

You can choose to require authentication (a shared secret) in order to secure communications between failover partners. This is enabled in the DHCP console by selecting the checkbox next to **Enable Message Authentication**. If you choose to enable authentication, you must provide a shared secret.

Windows Server 2012 implements failover message authentication using the Secure Hash Algorithm 2 (SHA-2) cryptographic standard. By default, DHCP failover uses the SHA-256 algorithm.

See also

What is DHCP Failover?

DHCP Failover Modes

DHCP Failover Communications

DHCP Failover Examples

Deploy DHCP Failover

DHCP Failover Examples

The following examples illustrate the way in which DHCP failover works to provide load balancing or redundancy of DHCP services.

Hot standby example

Load balancing example

Partner down example

For information about the types of messages exchanged between DHCP failover partner servers, see DHCP Failover Communications.

Hot standby example

The following illustration and its corresponding steps provide a detailed description of the processes involved when a DHCP client acquires a DHCP lease from DHCP servers configured for DHCP failover in hot standby mode.

1. A DHCP client in the INIT state broadcasts a DHCP discover message.
2. Since DHCP2 is in standby status, it does not respond to the client. DHCP1 is in active status and responds by offering a lease to the client.
3. The client requests a lease from DHCP1.
4. DHCP1 provides an IP address lease to the client for a lease duration of MCLT.
5. DHCP1 sends a binding update (BNDUPD) message to DHCP2 letting it know about the new lease.
6. DHCP2 responds to DHCP1 and acknowledges that it received the binding update (BNDACK).
7. At 50% of the lease time (MCLT/2), the client requests a renewal for its lease from DHCP1.
8. DHCP1 renews the client's DHCP lease for the full scope lease duration.
9. DHCP1 sends a binding update (BNDUPD) message to DHCP2 informing it of the lease renewal.
10. DHCP2 acknowledges that it received the binding update message from DHCP1 (BNDACK).
11. DHCP1 becomes unavailable.
12. DHCP2 receives an error when attempting to communicate with

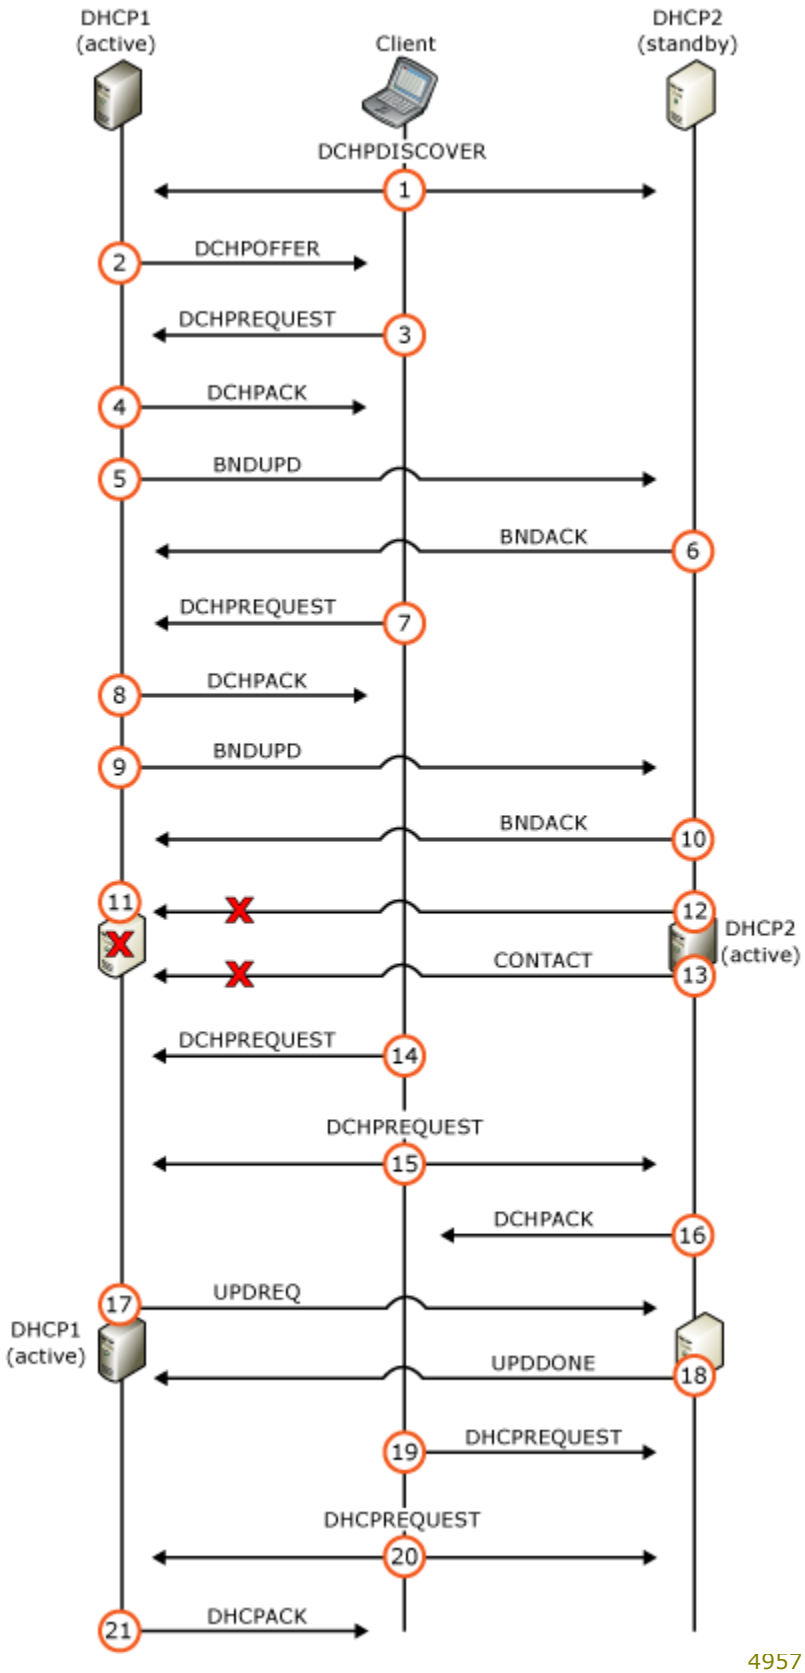

Load balancing example

The following illustration and its corresponding steps provide a detailed description of the processes involved when a DHCP client acquires a DHCP lease from DHCP servers configured for DHCP failover in load balancing mode.

1. A DHCP client in the INIT state broadcasts a DHCP discover message.
2. DHCP1 and DHCP2 each compute a hash of the client ID in the range 0-255. It is determined that the hash value belongs to the hash bucket assignment for DHCP1, so only DHCP1 responds to the client.
3. The client requests a lease from DHCP1.
4. DHCP1 provides an IP address lease to the client for a lease duration of MCLT.
5. DHCP1 sends a binding update (BNDUPD) message to DHCP2 letting it know about the new lease.
6. DHCP2 responds to DHCP1 and acknowledges that it received the binding update (BNDACK).
7. At 50% of the lease time (MCLT/2), the client requests a renewal for its lease from DHCP1.
8. DHCP1 renews the client's DHCP lease for the full scope lease duration.
9. DHCP1 sends a binding update (BNDUPD) message to DHCP2 informing it of the lease renewal.
10. DHCP2 acknowledges that it received the binding update message from DHCP1 (BNDACK).
11. DHCP1 becomes unavailable.
12. DHCP2 receives an error when attempting to communicate with DHCP1 or a timer expires because no messages have been received from DHCP1.

    DHCP2 drops the connection, enters the communication interrupted state, and begins accepting leases requests for all clients.
13. DHCP2 attempts and fails to reestablish a connection with DHCP1.
14. At 50% of the scope lease time, the client requests a renewal for its lease from DHCP1, but DHCP1 is not available so the lease it not renewed

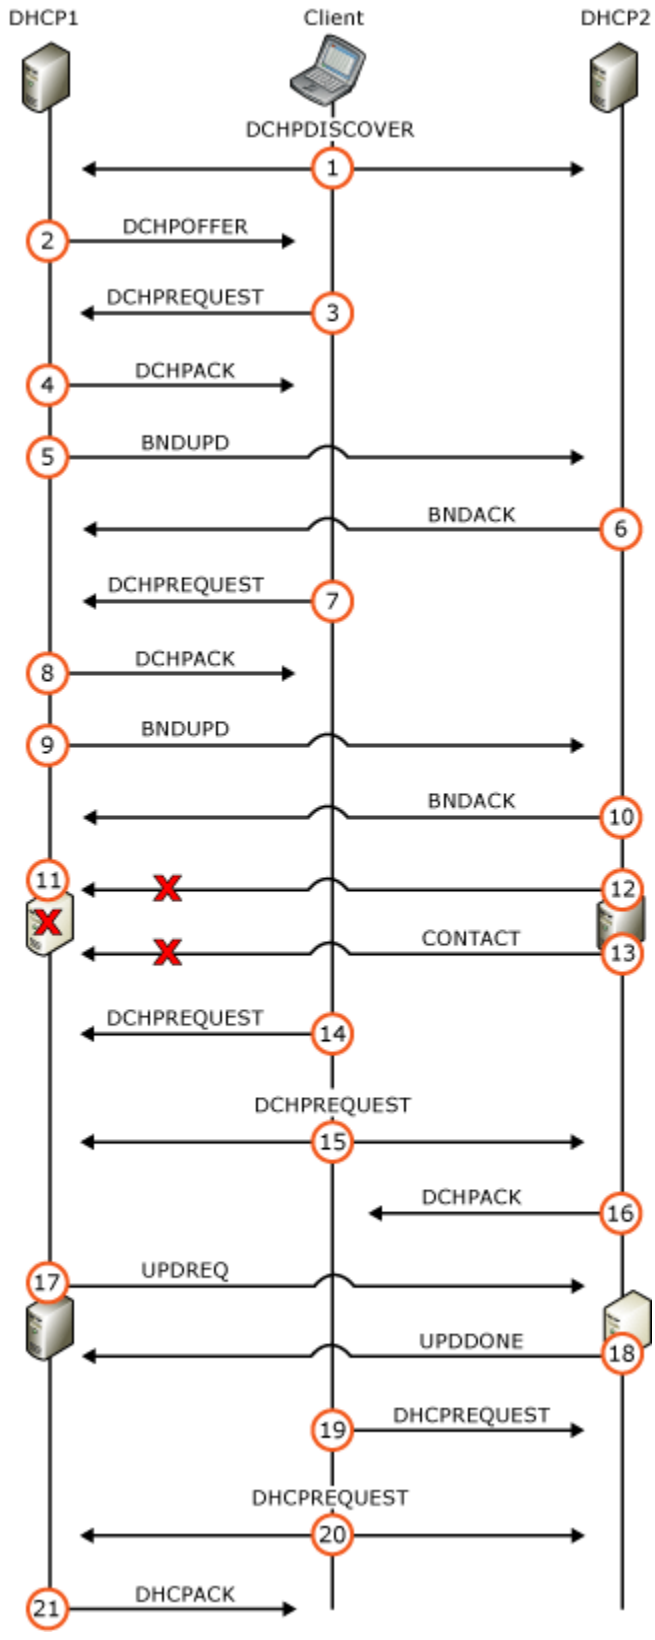

4959

Partner down example

The following illustration and its corresponding steps provide a detailed description of the processes involved when a DHCP failover partner server enters the partner down state in load balance mode.

1. A DHCP client in the INIT state broadcasts a DHCP discover message.
2. DHCP1 and DHCP2 each compute a hash of the client ID in the range 0-255. It is determined that the hash value belongs to the hash bucket assignment for DHCP1, so only DHCP1 responds to the client.
3. The client requests a lease from DHCP1.
4. DHCP1 provides an IP address lease to the client for a lease duration of MCLT.
5. DHCP1 sends a binding update (BNDUPD) message to DHCP2 letting it know about the new lease.
6. DHCP2 responds to DHCP1 and acknowledges that it received the binding update (BNDACK).
7. At 50% of the lease time (MCLT/2), the client requests a renewal for its lease from DHCP1.
8. DHCP1 renews the client's DHCP lease for the full scope lease duration.
9. DHCP1 sends a binding update (BNDUPD) message to DHCP2 informing it of the lease renewal.
10. DHCP2 acknowledges that it received the binding update message from DHCP1 (BNDACK).
11. DHCP1 becomes unavailable.
12. DHCP2 receives an error when attempting to communicate with DHCP1 or a timer expires because no messages have been received from DHCP1.

    DHCP2 drops the connection, enters the communication interrupted state, and begins accepting leases requests for all clients.
13. DHCP2 attempts and fails to reestablish a connection with DHCP1.
14. At 50% of the scope lease time, the client requests a renewal for its lease from DHCP1, but DHCP1 is not available so the lease it not renewed.

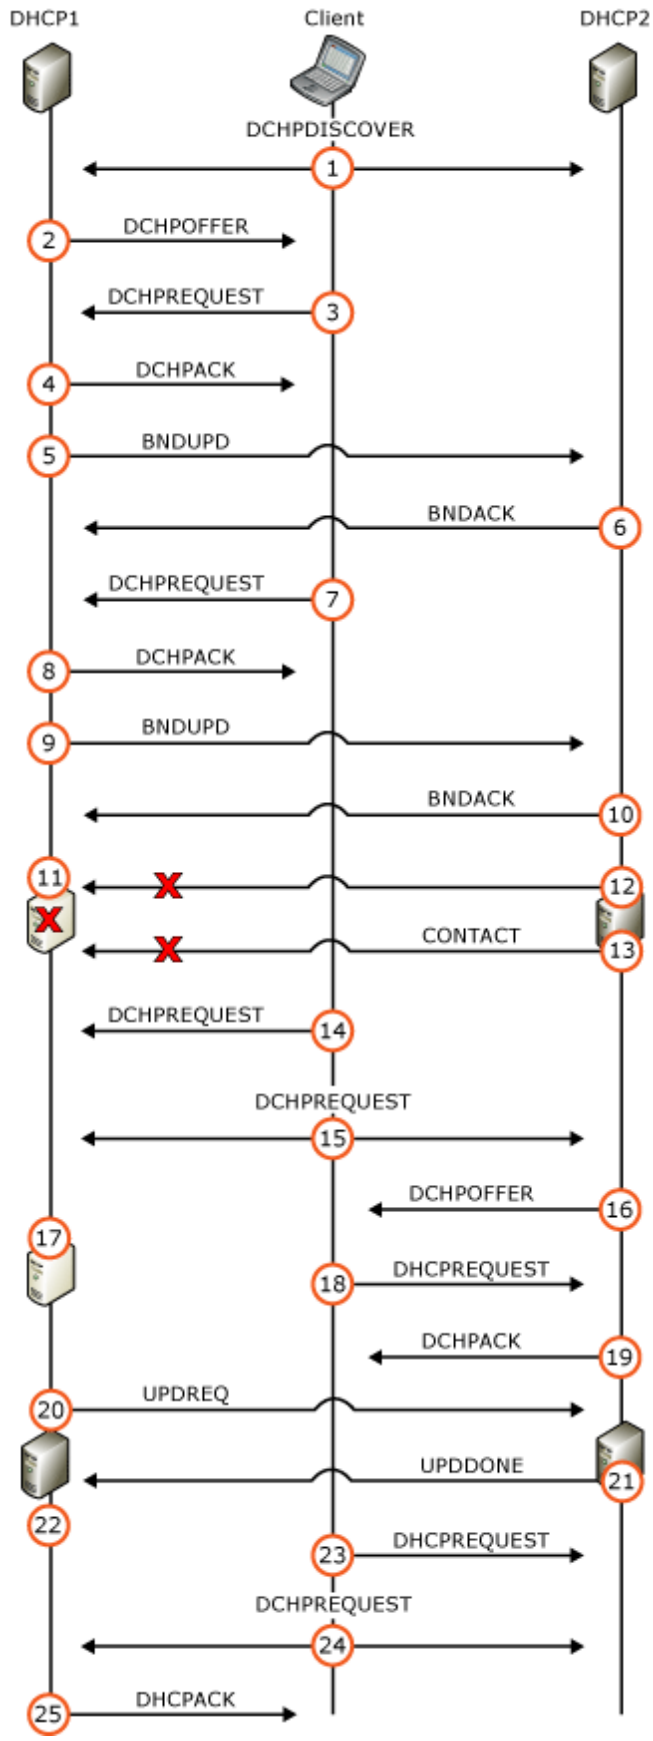

See also

What is DHCP Failover?

DHCP Failover Relationships

DHCP Failover Modes

DHCP Failover Communications

DHCP Failover Settings

DHCP Failover Events and Performance

Deploy DHCP Failover

DHCP Failover Events and Performance

Windows Server® 2012 introduces new event logging and performance counters to support monitoring and troubleshooting of DHCP failover. See the following sections in this topic for more information:

1. DHCP server event channels
2. Performance counters

DHCP server event channels

The following DHCP Server event channels are available using Event Viewer with the path: **Applications and Services Logs\Microsoft\Windows\DHCP-Server**.

1. **Admin** channel (Microsoft-Windows-DHCP Server Events/Admin): This channel provides DHCP server administrative event logging. For example, events are logged to this channel if a DHCP failover state of a server changes. For more information, see Administrative event logging.
2. **FilterNotifications** channel (Microsoft-Windows-DHCP Server Events/FilterNotifications): This channel provides DHCP server link-layer based filtering event logging. This event channel is not used for DHCP failover related events.
3. **Operational** channel (Microsoft-Windows-DHCP Server Events/Operational): This channel provides DHCP configuration auditing. For example, an event is logged if a scope on the DHCP server is added or removed from a DHCP failover relationship. For more information, see Operational event logging.

New DHCP failover-specific events have been added to these channels for administrative and operational (configuration audit) event logging.

Administrative event logging

Event Tracing for Windows (ETW) events are logged for state transition of the local server, or its partner server. These events are logged to the DHCP server **Admin** channel. The following data is provided as part of each state change event.

1. The previous state of the server
2. The current state of the server
3. The failover relationship name
4. The server host name or IP address

In addition to logging its own state change, the server will log changes in the state of its failover partner server. This information is communicated by state messages from the failover partner. The following table lists the values populated for the various fields in the event log.

Event Field	Value
Event ID	20252 (example ID)
Date and Time	Time of the state change
Computer	DHCP server host name or IP address
User	NETWORK SERVICE
Description	The failover state of server: <host name> for failover relationship: <relationship name> changed from: <previous state> to <current state>. The server IP address or host name corresponds to the server that has changed state. This can be the host name or IP address of the local DHCP server or the failover partner. **Note** If the previous state of the server is not known (this happens in the case of new relationship creation and when communication is interrupted), the following description will be displayed for the event: The failover state of server: <host name> for failover relationship changed to <current state>.
OpCode	FailoverStateChange
Task Category	DHCP Failover
Level	• Error (if transitioning out of the NORMAL

Event Field	Value
	state or when entering into the PARTNER DOWN state) • Information (if transitioning into the NORMAL state) • Warning (for transitions between other states)
Keywords	

See the following example event.

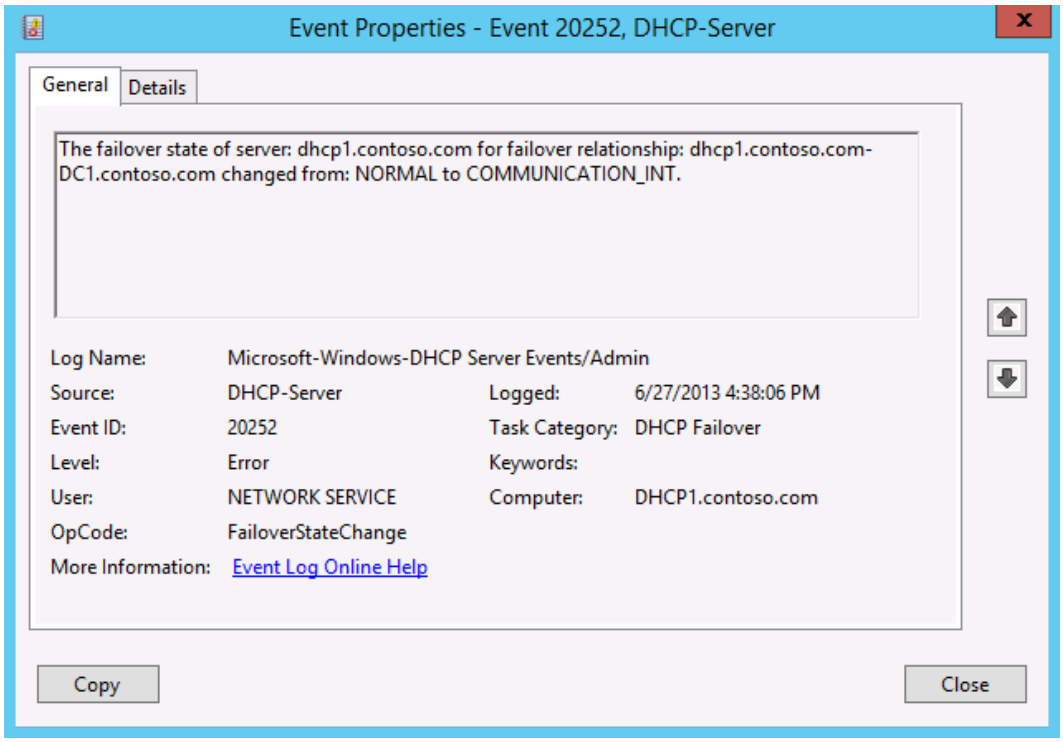

The following type of event data is logged when a server detects that it is out of time synchronization with its partner server.

Event Field	Value
Event ID	20253
Date and Time	Time of detection of time being out of sync with partner server

Event Field	Value
Computer	DHCP server host name
User	NETWORK SERVICE
Description	The server detected that it is out of time synchronization with partner server: <host name> for failover relationship: <relationship name>. The time is out of sync by: <# of seconds> seconds.
OpCode	TimeOutOfSync
Task Category	DHCP Failover
Level	Error
Keywords	

An event will also be logged when a connection is established and when a connection is lost between DHCP failover partners.

Level	Opcode	Task Category	Description
Error	CommDown	DHCP Failover	Server has lost contact with failover partner server: <host name> for relationship: <relationship name>.
Information	CommUp	DHCP Failover	Server has established contact with failover partner server: <host name> for relationship: <relationship name>.

Operational event logging

The following information-level configuration audit events are logged to the DHCP server **Operational** channel.

Event Type	OpCode	Description
Creation of Failover Relationship	CreateFailover	A failover relationship has been created between servers: <host name> and

Event Type	OpCode	Description
		<host name> with the following configuration parameters: Name: <value> Mode: <mode>, Maximum Client Lead Time: <value> seconds, Load Balance percentage on this server/Reserve Address percentage on standby server: <value>, Auto state switchover interval: <value> seconds Standby server: <host name>
Deletion of Failover Relationship	DeleteFailover	Failover relationship <relationship name> between <host name> and <host name> has been deleted
Addition of a scope to an existing failover relationship	AddFailoverScope	Scope <scope ID> has been added to the failover relationship <relationship name> with server <host name>
Removal of a scope from an existing failover relationship	RemoveFailoverScope	Scope <scope ID> has been removed from the failover relationship <relationship name> with server <host name>
Change MCLT	ChangeFailoverConfig	The failover configuration parameter MCLT for failover relationship <relationship name> with server <host name> has been changed from <old value> seconds to <new value> seconds
Change Auto state switchover interval	ChangeFailoverConfig	The failover configuration parameter Auto State switchover interval for failover relationship <relationship name> with server <host

Event Type	OpCode	Description
		name> has been changed from <old value> seconds to <new value> seconds
Change Reserve Address percentage	ChangeFailoverConfig	The failover configuration parameter Reserve Address Percentage on standby server for failover relationship <relationship name> with server <host name> has been changed from <old value> to <new value>
Change Mode of failover relationship	ChangeFailoverConfig	The failover configuration parameter Mode for failover relationship <relationship name> with server <host name> has been changed from <old value> to <new value>
Change Load Balance percentage	ChangeFailoverConfig	The failover configuration parameter Load Balance percentage for failover relationship <relationship name> with server <host name> has been changed from <old value> to <new value> on this server

See the following example event.

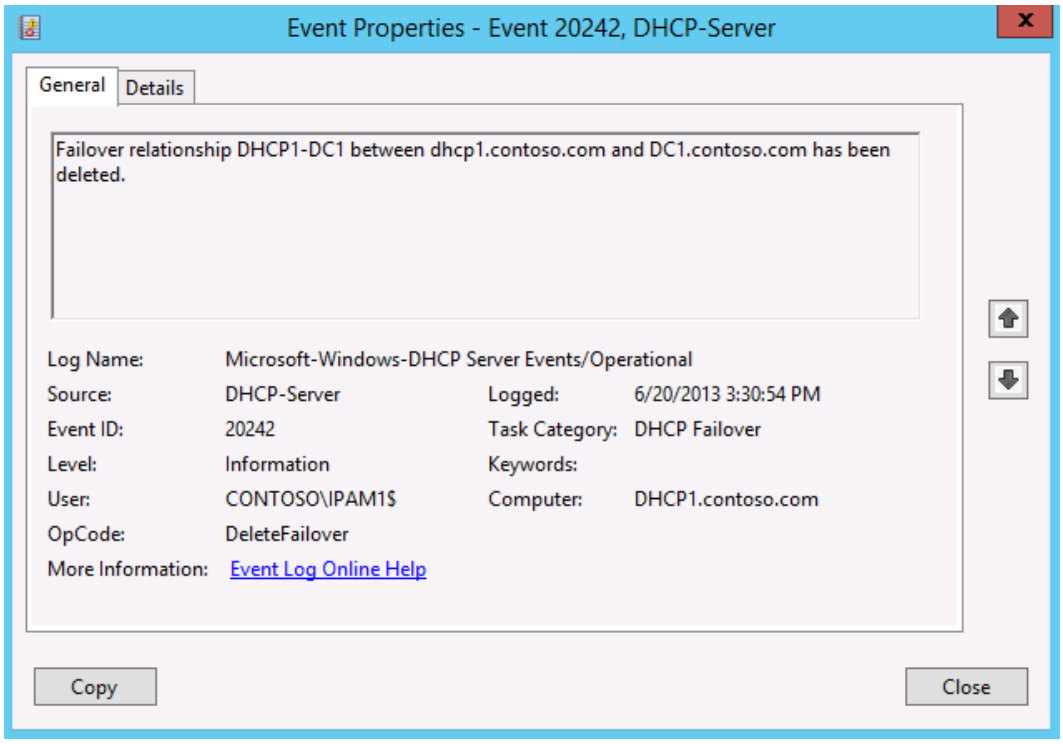

Performance counters

Windows Server 2012 introduces the following new DHCP failover-related performance counters to the DHCP server counters:

- Binding Updates sent per minute
- Binding Acks received per minute
- Binding Updates received per minute
- Binding Acks sent per minute
- Number of pending outbound binding updates on this server
- Number of transitions into COMMUNICATION INTERRUPTED state
- Number of transitions into PARTNER DOWN state
- Number of transitions into RECOVER state
- Number of Binding Update queue running at 90% of the maximum queue size
- Number of Binding Updates dropped

See also

What is DHCP Failover?

DHCP Failover Relationships

DHCP Failover Modes

DHCP Failover Communications

DHCP Failover Settings

DHCP Failover Examples

Deploy DHCP Failover

Deploy DHCP Failover

The following topics provide information you can use to deploy DHCP failover:

DHCP Failover Requirements

DHCP Failover Architecture

Checklist: Deploy DHCP Failover

Configure DHCP Failover using the Command Line

Configure DHCP Failover using the DHCP Console

Replicate DHCP Failover Using the Command Line

Replicate DHCP Failover Using the DHCP Console

Migrate to DHCP Failover

A brief overview of how to configure DHCP failover using the DHCP console or Windows PowerShell is also provided in this topic.

Overview: Configure DHCP failover using the DHCP console

To configure DHCP failover using the DHCP console, right-click a DHCP scope or right-click **IPv4** and then click **Configure Failover**.

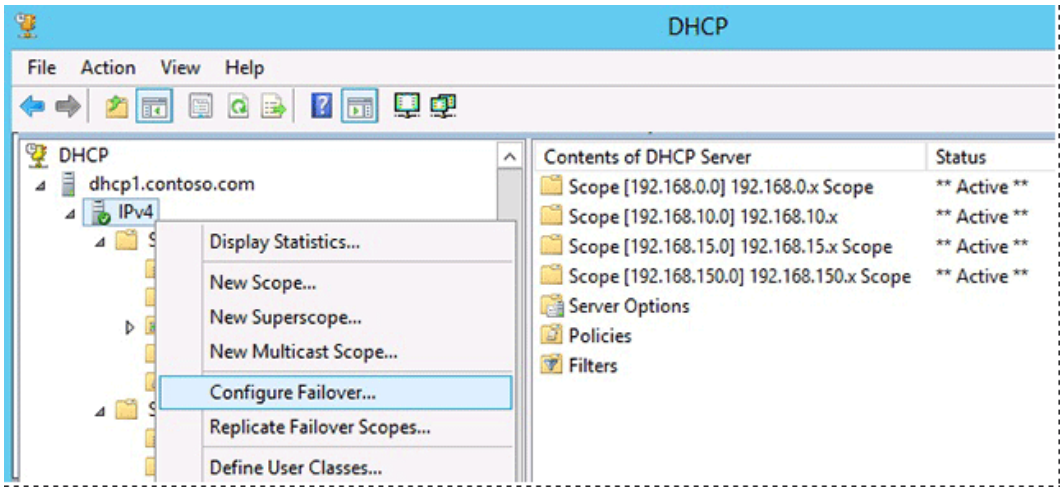

The **Configure Failover** wizard guides you through configuring DHCP failover on the selected scope.

Overview: Configure DHCP failover using Windows PowerShell

To configure DHCP failover using the command line, use the **Add-DhcpServerv4Failover** Windows PowerShell cmdlet to create a new DHCP failover relationship and add a DHCP scope, or use the **Add-DhcpServerv4FailoverScope** cmdlet to add a new scope to an existing failover relationship. An example of the **Add-DhcpServerv4Failover** cmdlet is provided below.

```
Add-DhcpServerv4Failover -ComputerName dhcp1.contoso.com -PartnerServer dhcp2.contoso.com
-Name dhcp1-dhcp2 -ScopeID 10.10.10.0 -LoadBalancePercent 70 -SharedSecret sEcReT -Force
```

For information more information about configuring DHCP failover with Windows PowerShell, see DHCP Server Cmdlets in Windows PowerShell.

See also

What is DHCP Failover?

DHCP Failover Relationships

DHCP Failover Modes

DHCP Failover Communications

DHCP Failover Settings

DHCP Failover Examples

DHCP Failover Events and Performance

DHCP Failover Requirements

Requirements to deploy DHCP failover are the following:

Item	Requirement	Details
Operating system	Windows Server® 2012, or a later operating system is required.	Both DHCP servers in a DHCP failover relationship must be running Windows Server® 2012 or a later operating system. Servers do not need to both be running the same operating system, but this is

Item	Requirement	Details
		recommended to ensure consistent replication of settings.
Role services	The DHCP Server role service is required.	Both computers participating in a DHCP failover relationship must have the DHCP Server role installed and running. The DHCP Server service can be paused, but must not be stopped. Additional roles, role services, and features are optional.
Network	DHCP servers can be on the same network or on different networks. The network connection between DHCP failover partners must be uninterrupted. Both DHCP servers must be able to communicate with each other, and with all DHCP clients that will receive DHCP leases from failover-enabled scopes. Both servers must also be time synchronized to within one minute of each other.	DHCP servers can communicate with each other directly, or through IP routing. DHCP servers can communicate with DHCP clients either directly or using DHCP relay. For more information about configuring DHCP relay agents for a DHCP failover deployment, see DHCP Failover Architecture.
IP addresses	DHCP servers should be configured with static IP addresses.	To ensure a persistent TCP connection between DHCP failover partners is maintained, it is important to use a static IP address on all DHCP server network interfaces. If the static IP address of a DHCP server needs to be changed, for example during DHCP migration, you must first delete all DHCP failover relationships that exist on that

Item	Requirement	Details
		server, and then recreate the relationships when the new IP address is active. For more information about communication between DHCP failover partners, see DHCP Failover Communications.
DHCP scopes	At least one IPv4 DHCP scope must be configured on the primary DHCP server. The same DHCP scope ID, or an overlapping scope, must not be configured on the failover partner.	Other DHCP scopes that are not overlapping can be configured on the failover partner, but are not required. Scopes that are not configured for DHCP failover are not affected. The DHCP scope does not need to be active. An inactive DHCP scope that is replicated to a DHCP failover partner server will also be inactive on the failover partner. You cannot configure DHCP failover using two DHCP servers that are already configured for a split-scope DHCP deployment because the same DHCP scope ID is present on both servers. Initial replication of a DHCP scope to the failover partner will fail if the scope ID already exists on the destination server.
Domain membership	Not required.	DHCP servers can be workgroup computers or domain member computers. However, workgroup computers cannot be authorized in Active Directory. For more information, see Authorizing DHCP servers.

Prerequisite checks

The following prerequisite checks are made before enabling DHCP failover:

Check performed	Error message displayed
DNS name resolution for the failover partner server	Unable to resolve the specified DNS name.
Valid IP address specified for the failover partner server	The specified IP address <x.x.x.x> is invalid.
The server has a network connection to the failover partner server	The specified DHCP server is not reachable. Please provide a DHCP server that is reachable.
The operating system on the failover partner server	The version of specified DHCP server does not support failover.
The user is a member of the DHCP Administrators group, or equivalent, on the failover partner server	You do not have permissions to perform this operation on the remote DHCP server.
The maximum number of failover relationships for either DHCP server	The local/partner server already has 31 (maximum allowed) failover relationships. A server cannot have more than 31 failover relationships.
The DHCP Server service is running on the failover partner server	DHCP server is not running on the specified server. Please ensure that DHCP server is running on the specified server.
Time is synchronized between both servers	The time difference between this server and the specified partner server is greater than the permissible value of x minutes. It is recommended to ensure that both servers are time synchronized before configuring failover. You could setup Network Time Protocol (NTP) service on both servers to ensure time synchronization. **Tip** A maximum permissible time difference can be configured in the Windows Registry. The registry value will be read from both the servers and the minimum of the values is used to perform this check. If registry value has not been

Check performed	Error message displayed
	configured, the default value is used.
Are scopes already present on the failover partner server	Following scopes already exist on the specified partner server. These scope(s) will need to be deleted on the partner server before configuring failover.

See also

What is DHCP Failover?

DHCP Failover Architecture

Checklist: Deploy DHCP Failover

DHCP Failover Architecture

See the following sections in this topic:

1. DHCP failover design
2. Deployment considerations

DHCP failover design

DHCP failover can be deployed using either hot standby mode or load sharing mode. You can also use both modes on the same DHCP server if you configure multiple failover relationships. For details about hot standby and load sharing modes, see DHCP Failover Modes.

You can customize your deployment based on the physical architecture of your network. See the following sections for more information.

Hot standby design

The hot standby mode of operation is best suited to deployments where a central office or data center server acts as a standby backup server to a server at a remote site, which is local to the DHCP clients (ex: hub and spoke deployment). In such deployments, it is undesirable to have a remote standby server service any clients unless the local DHCP server becomes unavailable. The following figure is an example of a hub and spoke deployment.

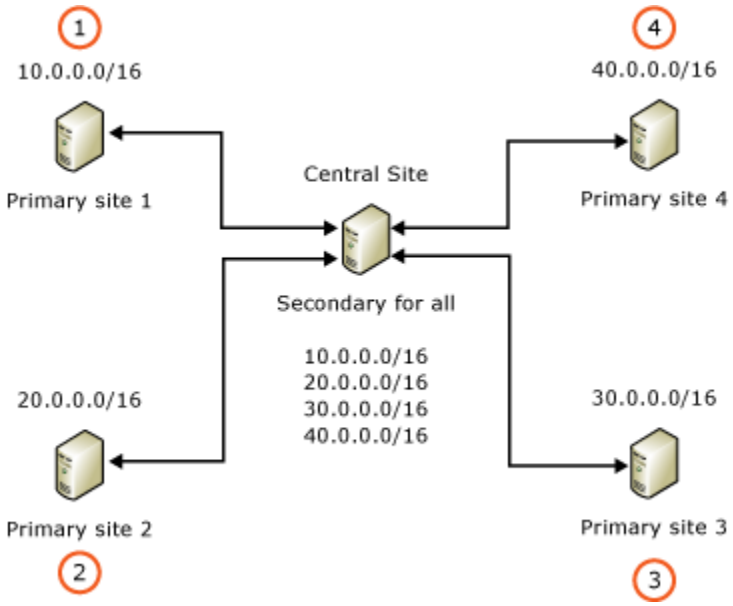

In this example, four remote sites each have a primary (active) DHCP server on-site with the standby server located remotely at a central site. The standby server at the central site might also be an active DHCP server for a subnet at the central site, or its only function can be as a standby server for the sites 1-4. Client leases for all four sites can be viewed on the central DHCP server, even if it is not actively providing these leases.

Load balance design

The load balance mode of operation is best suited to deployments where both servers in a failover relationship are located at the same physical site. Both servers respond to DHCP client requests based on the load distribution ratio configured by the administrator.

Load balancing can be used with a single subnet, or multiple subnets. See the following examples.

Load balance in a single site with a single subnet

In the following example, two DHCP servers simultaneously provide IP addressing to clients on a single subnet.

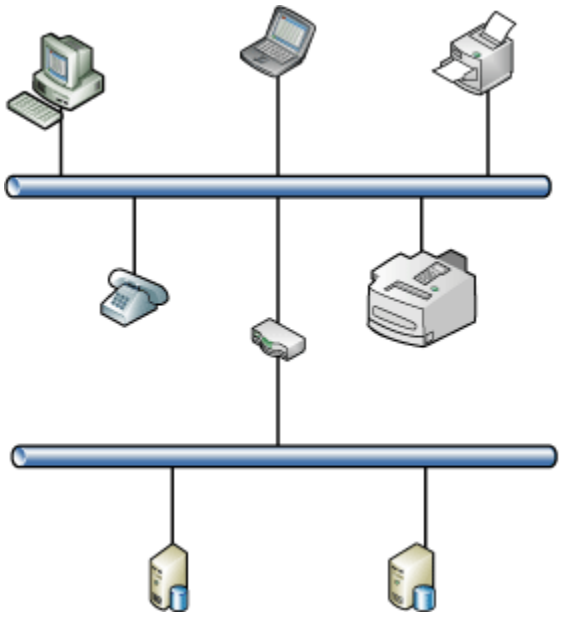

Legend

DHCP Failover — Single Subnet	
Symbol	Description
	Ethernet
	DHCP Server
	DHCP Relay
	Printer
	PC
	Laptop Computer
	IP Phone
	Multi-function device

Load balance in a single site with multiple subnets

In the following example, two DHCP servers simultaneously provide IP addressing to clients on multiple subnets.

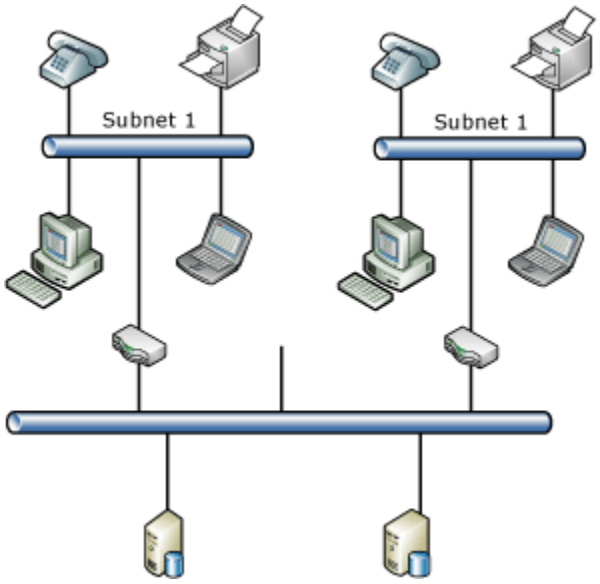

Legend

DHCP Failover — Single site, Multiple Subnets	
Symbol	Description
	Ethernet
	DHCP Server
	DHCP Relay
	Printer
	PC
	Laptop Computer
	Telephone

As previously noted, a load balance design also provides redundancy. If one of the servers in the previous examples becomes unavailable, the other DHCP server will assume responsibility for client leases.

For more information about redundancy in hot standby and load balancing modes, see DHCP Failover Examples.

Deployment considerations

Before you deploy DHCP failover, consider the following:

Time synchronization

For DHCP failover to function correctly, time must be kept synchronized between the two servers in a failover relationship. Time synchronization can be maintained by deployment of the Network Time Protocol (NTP) or any other alternative mechanism. When the failover configuration wizard is run, it will compare the current time on the servers being configured for failover. If the time difference between the servers is greater than one minute, the failover setup process will halt with a critical error and request that time on the servers be synchronized.

Each failover protocol message includes a time field that is populated with the Coordinated Universal Time (UTC) when the source server transmitted the message. For each message, the receiving server will perform a check of the time difference between the time stamp field in the packet and the time at the receiving server. If this time difference is found to be greater than one minute, the receiving server will log a critical event indicating that the two servers are not time synchronized.

A binding is a collection of configuration parameters managed by DHCP servers, including at least an IP address, associated with or "bound to" a DHCP client. A binding update transaction refers to the set of information (contained in options), necessary to perform an update for a single IP address. Any change in state of an IP address lease (e.g. a new lease, renew, expiry, release) leads to a binding update message from the server on which the state change happened to a partner server. Any binding update message received by the server while time is out of synchronization is rejected with an error code. Determination of a time out-of-synch condition will not cause any change in the failover state of the server. Both servers will continue to operate in the same state as before determination of a time out-of-synch condition.

BOOTP support

DHCP failover is not supported for BOOTP in Windows Server 2012. Support for BOOTP clients is provided in Windows Server 2012 R2.

Policy based assignment

Windows Server 2012 includes a new policy based IP address assignment feature, which allows a Windows DHCP administrator to group the DHCP clients by a specific attribute of the client, such as vendor class, user class, client identifier, or MAC address. By grouping the clients based on these attributes, an administrator is able to assign parameters such as IP address, default gateway, DNS server and other DHCP options to a specific grouping of clients. This allows the administrator to exercise greater control on the configuration parameters delivered to end hosts.

This feature introduces the concept of multiple IP address ranges within a single scope. To accommodate this, DHCP failover address distribution in load sharing mode is done on a per IP address range basis.

Windows Firewall

DHCP failover uses TCP port 647 to listen for failover messages between two failover partner servers. For this traffic to be allowed by the Windows firewall, the following inbound and outbound firewall rules are added then you install the DHCP Server role:

- Microsoft-Windows-DHCP-Failover-TCP-In
- Microsoft-Windows-DHCP-Failover-TCP-Out

Relay agents

Initial DHCPDISCOVER messages are broadcast by DHCP clients on the subnet to which they belong. Since routers typically do not forward broadcast traffic, a mechanism is required to enable DHCP clients to communicate with DHCP servers if the DHCP server is not located on the same subnet. Relay agents (typically provided on a router) are designed to perform this function, relaying DHCP and BOOTP messages between clients and servers on different subnets. Relay agents are commonly configured with on a network device, or you can configure DHCP relay on a Windows Server with the Remote Access role installed. See Configure the DHCP Relay Agent for more information.

If your DHCP relay is configured on a network device, consult your vendor's documentation for details. The **helper-address** command is commonly used to configure DHCP relay on a network device, for example:

```
ip helper-address 10.0.1.1
```

When you deploy DHCP failover, a single DHCP relay address might not be sufficient, since DHCP clients must always be able to communicate with both the primary DHCP server and the failover partner server. If both DHCP servers are located on a different subnet than DHCP clients, this requires at least two DHCP relay agents. For example:

```
ip helper-address 10.0.1.1

ip helper-address 10.0.1.2
```

In this example, both DHCP servers are on the same subnet (10.0.1.0/24). The primary DHCP server's IP address is 10.0.1.1 and 10.0.1.2 is the IP address for the failover partner server. If both DHCP servers are located on the same subnet, you can also configure the subnet broadcast address (ex: 10.0.1.255) as a single DHCP relay. The use of a subnet broadcast address as a single DHCP relay is not possible if DHCP servers are located on separate subnets.

Duplicate relay agents

Virtual Router Redundancy Protocol (VRRP) is another failover protocol used to enable redundancy on network devices. An example of VRRP includes Hot Standby Router Protocol (HSRP), which is a Cisco proprietary VRRP. If VRRP/HSRP is configured on a network device

that is also configured with one or more DHCP relays, this can cause duplicate DHCP relay messages to be sent to the same DHCP failover server.

If a single DHCP server configured for DHCP failover receives duplicate lease requests, this can cause inconsistent client lease durations, and clients might lease IP addresses that belong to other clients. Consult your vendor documentation to determine if the router redundancy protocol requires a specific configuration to support DHCP relay. For example, Cisco provides DHCP relay support for the HSRP protocol using virtual router groups.

An update is also available from Microsoft to resolve this issue. See Microsoft Knowledge Base article 2831920 and Microsoft update 2845533 for more information.

See also

What is DHCP Failover?

DHCP Failover Requirements

Checklist: Deploy DHCP Failover

Checklist: Deploy DHCP Failover

This checklist includes cross-reference links to important concepts about deploying DHCP failover. It also contains links to procedures you can use to configure DHCP failover.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the steps in a procedure so that you can continue with the remaining tasks in this checklist.

Checklist: Deploy DHCP Failover

	Task	Reference
☐	Review DHCP failover concepts and components; identify your design goals.	What is DHCP Failover? DHCP Failover Relationships DHCP Failover Modes DHCP Failover Communications DHCP Failover Settings DHCP Failover Examples DHCP Failover Events and Performance

	Task	Reference
☐	Review DHCP failover requirements and specifications; decide on a deployment topology; identify pilot sites; document deployment decisions and processes.	DHCP Failover Requirements DHCP Failover Architecture
☐	Configure DHCP failover	Configure DHCP Failover using the Command Line Configure DHCP Failover using the DHCP Console
☐	Replicate DHCP failover settings (optional)	Replicate DHCP Failover Using the Command Line Replicate DHCP Failover Using the DHCP Console
☐	Migrate to DHCP failover (optional)	Migrate to DHCP Failover

See also

Step-by-Step: Configure DHCP for Failover

Configure DHCP Failover using the Command Line

Use the following procedures to configure DHCP failover using the command line. When you have completed the procedures in this topic, return to the parent checklist: Checklist: Deploy DHCP Failover.

Configuring DHCP failover with the command line uses Windows PowerShell. To configure DHCP failover using the Windows interface instead, see Configure DHCP Failover using the DHCP Console.

Additional considerations:

- When you configure DHCP failover, you can accept default values for several parameters, or you can specify custom values. For more information about DHCP failover settings, see DHCP Failover Settings.
- DHCP failover can be configured on a single DHCP scope, or you can configure multiple scopes simultaneously if they will all use the same failover relationship. The procedures

below describe how to specify a new failover relationship for use with a single DHCP scope. Multiple scopes can be configured using the command line by providing a comma-delimited list of scope IDs after the **ScopeID** parameter.

- In the following procedures, the **Add-DhcpServerv4Failover** cmdlet is used to create a new DHCP failover relationship. You can also add scopes to an existing DHCP failover relationship using the **Add-DhcpServerv4FailoverScope** cmdlet. For more information about Windows PowerShell cmdlets for configuring and managing DHCP failover relationships, scopes, and server settings, see DHCP Server Cmdlets in Windows PowerShell.

To configure DHCP failover, you must choose load balance or hot standby mode:

- Configure DHCP failover in load balance mode using the command line
- Configure DHCP failover in hot standby mode using the command line

Configure DHCP failover in load balance mode using the command line

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure DHCP failover in load balance mode using the command line

1. On a computer with the Windows PowerShell DHCP Server module installed, right-click **Windows PowerShell** and then click **Run as Administrator**.
2. Type the following at the command line, and then press ENTER. Replace the example values with values used in your environment.

```
Add-DhcpServerv4Failover –ComputerName dhcp1.contoso.com –
PartnerServer dhcp2.contoso.com –Name dhcp1-dhcp2 –ScopeID
10.10.10.0 –LoadBalancePercent 70 –SharedSecret sEcReT -Force
```

In this example, the primary DHCP server is named **dhcp1.contoso.com** and the secondary DHCP server is named **dhcp2.contoso.com**. DHCP failover is enabled in load balancing mode with 70% of leases allocated to the primary DHCP server. The failover relationship name is **dhcp1-dhcp2**, the DHCP scope ID is **10.10.10.0**, and the shared secret is **sEcReT**. A shared secret is optional. If specified, message digest authentication is turned on for the newly created failover relationship.

Tip
Windows PowerShell DHCP Server module can be installed using the Add Roles and Features Wizard under Features\Remote Server Administration Tools\Role Administration Tools\DHCP Server Tools. The DHCP Server module is also installed by default when you install the DHCP Server role service.

The following table provides a description of parameters used in this example.

Parameter	Description
ComputerName	The name of the primary DHCP server.
PartnerServer	The name of the secondary DHCP server.
Name	The name of the new DHCP failover relationship.
ScopeID	The ID of the DHCP scope on the primary DHCP server to be configured for failover.
LoadBalancePercent	The load balancing percentage assigned to the primary DHCP server.
SharedSecret	The shared secret for the failover relationship.

Configure DHCP failover in hot standby mode using the command line

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure DHCP failover in hot standby mode using the command line

1. On a computer with the Windows PowerShell DHCP Server module installed, right-click **Windows PowerShell** and then click **Run as Administrator**.
2. Type the following at the command line, and then press ENTER. Replace the example values with values used in your environment.

```
Add-DhcpServerv4Failover -ComputerName dhcp1.contoso.com -
PartnerServer dhcp2.contoso.com -Name dhcp1-dhcp2 -ScopeID
10.10.10.0 -ServerRole Active -SharedSecret sEcReT -Force
```

In this example, the primary DHCP server is named **dhcp1.contoso.com** and the secondary DHCP server is named **dhcp2.contoso.com**. DHCP failover is enabled in hot standby mode with the primary DHCP server assigned as the active server. The failover relationship name is **dhcp1-dhcp2**, the DHCP scope ID is **10.10.10.0**, and the shared secret is **sEcReT**. A shared secret is optional. If specified, message digest authentication is turned on for the newly created failover relationship.

Tip

Windows PowerShell DHCP Server module can be installed using the Add Roles and Features Wizard under Features\Remote Server Administration Tools\Role Administration Tools\DHCP Server Tools. The DHCP Server module is also installed

by default when you install the DHCP Server role service.

The following table provides a description of parameters used in this example.

Parameter	Description
ComputerName	The name of the primary DHCP server.
PartnerServer	The name of the secondary DHCP server.
Name	The name of the new DHCP failover relationship.
ScopeID	The ID of the DHCP scope on the primary DHCP server to be configured for failover.
ServerRole	Specifies the role of the target DHCP server in hot standby mode. Acceptable values are Active or Standby. The target computer is specified by the ComputerName parameter.
SharedSecret	The shared secret for the failover relationship.

See also

What is DHCP Failover?

Deploy DHCP Failover

Step-by-Step: Configure DHCP for Failover

Configure DHCP Failover using the DHCP Console

Use the following procedures to configure DHCP failover using the DHCP console. When you have completed the procedures in this topic, return to the parent checklist: Checklist: Deploy DHCP Failover.

Windows Server 2012 and later operating systems adds the item **Configure Failover** to the list of selections you can make when you right-click the IPv4 node or right-click an individual scope in the console tree, and the **Configure Failover** wizard is provided to guide you through the configuration process.

If DHCP failover is already configured on a scope, choices of **Deconfigure Failover**, **Replicate Scope**, and **Replicate Relationship** are available. In addition, when you right-click the IPv4 node or right-click an individual scope in the console tree and then select **Properties**, the **Failover** is available, where you can view or edit DHCP failover settings. The **Failover** tab is not available for

a scope if the scope is not configured for DHCP failover. For more information viewing and editing DHCP failover settings, see DHCP Failover Relationships.

Use the following procedures to configure DHCP failover using the DHCP console. To configure DHCP failover using the command line instead, see Configure DHCP Failover using the Command Line.

Additional considerations:

- When you configure DHCP failover, you can accept default values for several parameters, or you can specify custom values. For more information about DHCP failover settings, see DHCP Failover Settings.
- DHCP failover can be configured on a single DHCP scope, or you can configure multiple scopes simultaneously if they will all use the same failover relationship. The procedures below describe how to specify a new failover relationship for use with a single DHCP scope. To configure multiple scopes using the DHCP console, choose these scopes in the **Configure Failover** wizard under **Available scopes**. Scopes that are available must be present on the local DHCP server and not already enabled for DHCP failover.

To configure DHCP failover, you must choose load balance or hot standby mode:

- Configure DHCP failover in load balance mode using the DHCP console
- Configure DHCP failover in hot standby mode using the DHCP console

Configure DHCP failover in load balance mode using the DHCP console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure DHCP failover in load balance mode using the DHCP console

1. On a DHCP server, open Server Manager, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click **IPv4**, and then click **Configure Failover**.
3. A list of available DHCP scopes is displayed. You can choose multiple scopes if desired. In this example, a single DHCP scope will be configured using a new failover relationship.
4. To configure failover on a single scope, clear the **Select all** checkbox and then click the scope ID of the scope you wish to configure for failover under **Available scopes**. Do not clear this checkbox if you wish to configure failover on all scopes. Alternatively, you can clear the checkbox and then choose multiple, selected scopes.

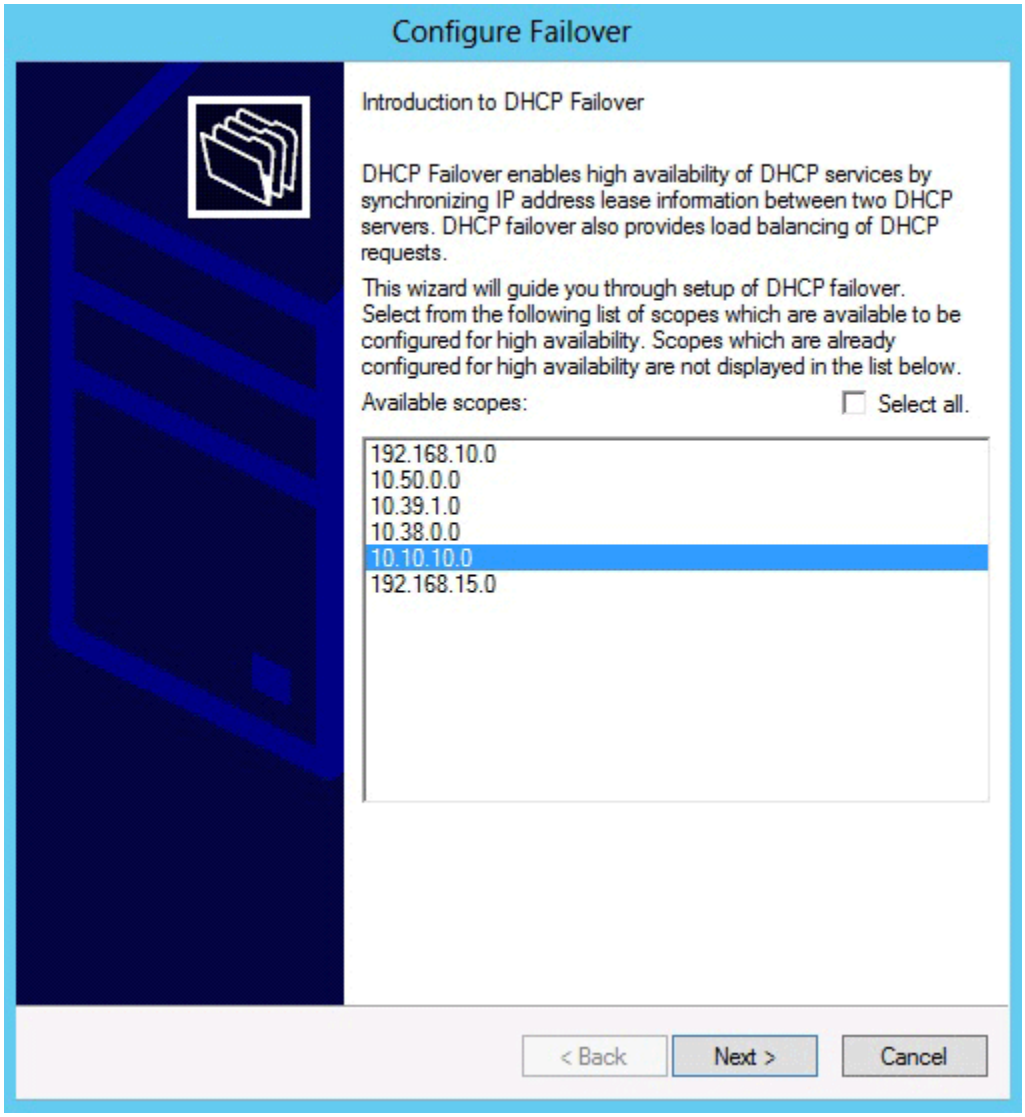

5. Click **Next**, and then on the **Specify the partner server to use for failover page**, next to **Partner Server** type the name or IP address of the failover partner DHCP server. You can also use the drop-down list to choose from a list of authorized DHCP servers in the domain.
6. To reuse an existing DHCP failover relationship, click **Next**. Alternatively, to create a new DHCP failover relationship, clear the **Reuse existing failover relationships…** checkbox, and then click **Next**.

Important

A maximum of 31 DHCP failover relationships can be configured. Reuse existing failover relationships to avoid exceeding this limit. When you reuse existing failover relationships and more than one failover relationship exists between the local DHCP server and the failover partner that was chosen, you are prompted to

select a relationship name from the list of available relationships. Choosing a relationship name will display parameters for that failover relationship. If only one relationship is configured, this relationship is automatically used.

7. On the **Create a new failover relationship page**, type a name for the failover relationship next to **Relationship Name** or accept the default name.
8. Confirm that **Load balance** is displayed by default next to **Mode**.
9. By default, the **Enable Message Authentication** checkbox is enabled. Clear this checkbox if you do not wish to use a shared secret. Otherwise, type a shared secret next to **Shared Secret**. By default, load balancing is set to 50% for both the local server and the partner server. You can adjust these values if desired.

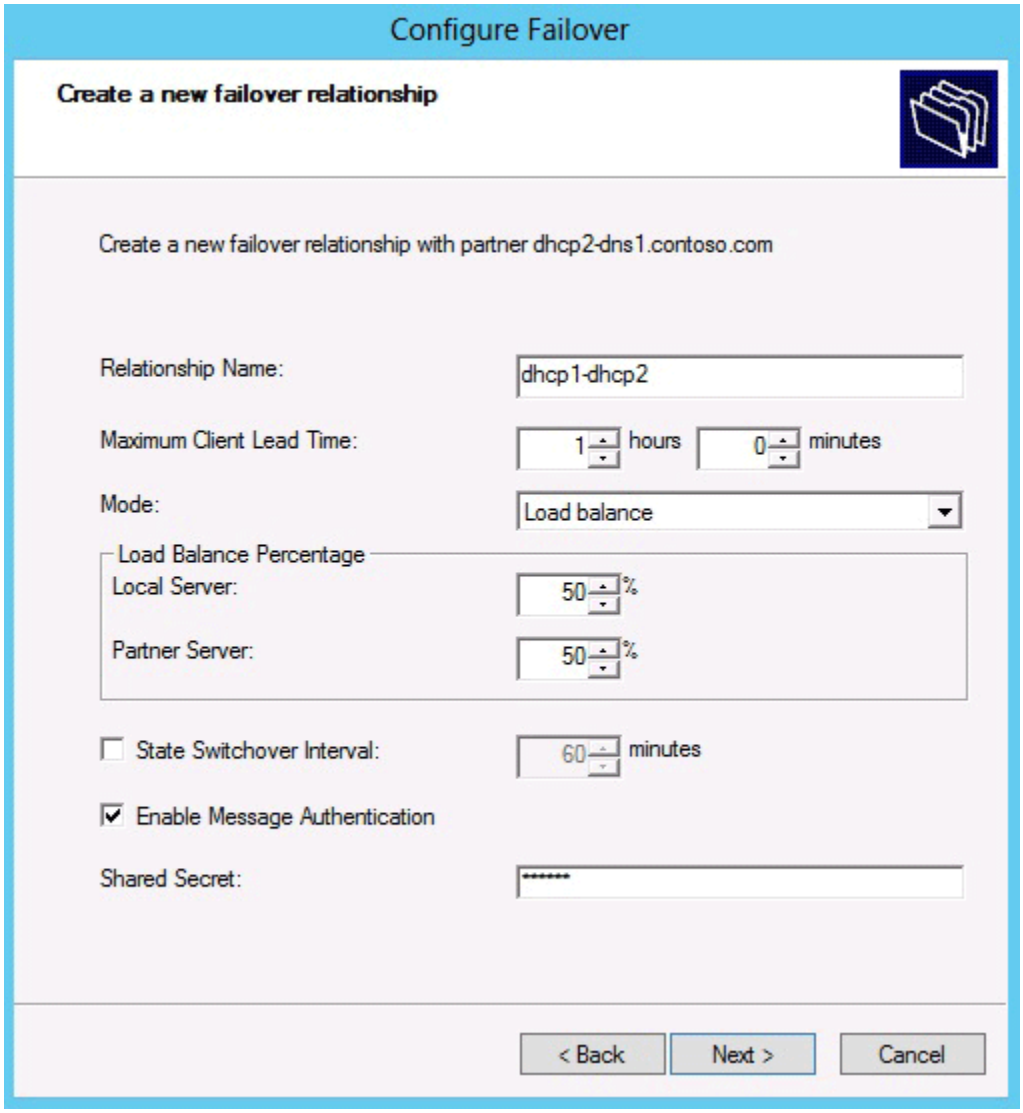

10. Click **Next**, click **Finish**, verify that the status of all the tasks displayed is **Successful**, and then click **Close**.

Configure DHCP failover in hot standby mode using the DHCP console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To configure DHCP failover in hot standby mode using the DHCP console

1. On a DHCP server, open Server Manager, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click **IPv4**, and then click **Configure Failover**.
3. A list of available DHCP scopes is displayed. You can choose multiple scopes if desired. In this example, a single DHCP scope will be configured using a new failover relationship.
4. Clear the **Select all** checkbox and then click the scope ID of the scope you wish to configure for failover under **Available scopes**. Do not clear this checkbox if you wish to configure failover on all scopes. Alternatively, you can clear the checkbox and then choose multiple, selected scopes.
5. Click **Next**, and then on the **Specify the partner server to use for failover page**, next to **Partner Server** type the name or IP address of the failover partner DHCP server. You can also use the drop-down list to choose from a list of authorized DHCP servers in the domain.
6. To reuse an existing DHCP failover relationship, click **Next**. Alternatively, to create a new DHCP failover relationship, clear the **Reuse existing failover relationships…** checkbox, and then click **Next**.

   Important
   A maximum of 31 DHCP failover relationships can be configured. Reuse existing failover relationships to avoid exceeding this limit. When you reuse existing failover relationships and more than one failover relationship exists between the local DHCP server and the failover partner that was chosen, you are prompted to select a relationship name from the list of available relationships. Choosing a relationship name will display parameters for that failover relationship. If only one relationship is configured, this relationship is automatically used.
7. On the **Create a new failover relationship page**, type a name for the failover relationship next to **Relationship Name** or accept the default name.
8. Choose **Hot standby** from the drop-down list next to **Mode**.
9. By default, the **Enable Message Authentication** checkbox is enabled. Clear this checkbox if you do not wish to use a shared secret. Otherwise, type a shared secret next to **Shared Secret**.

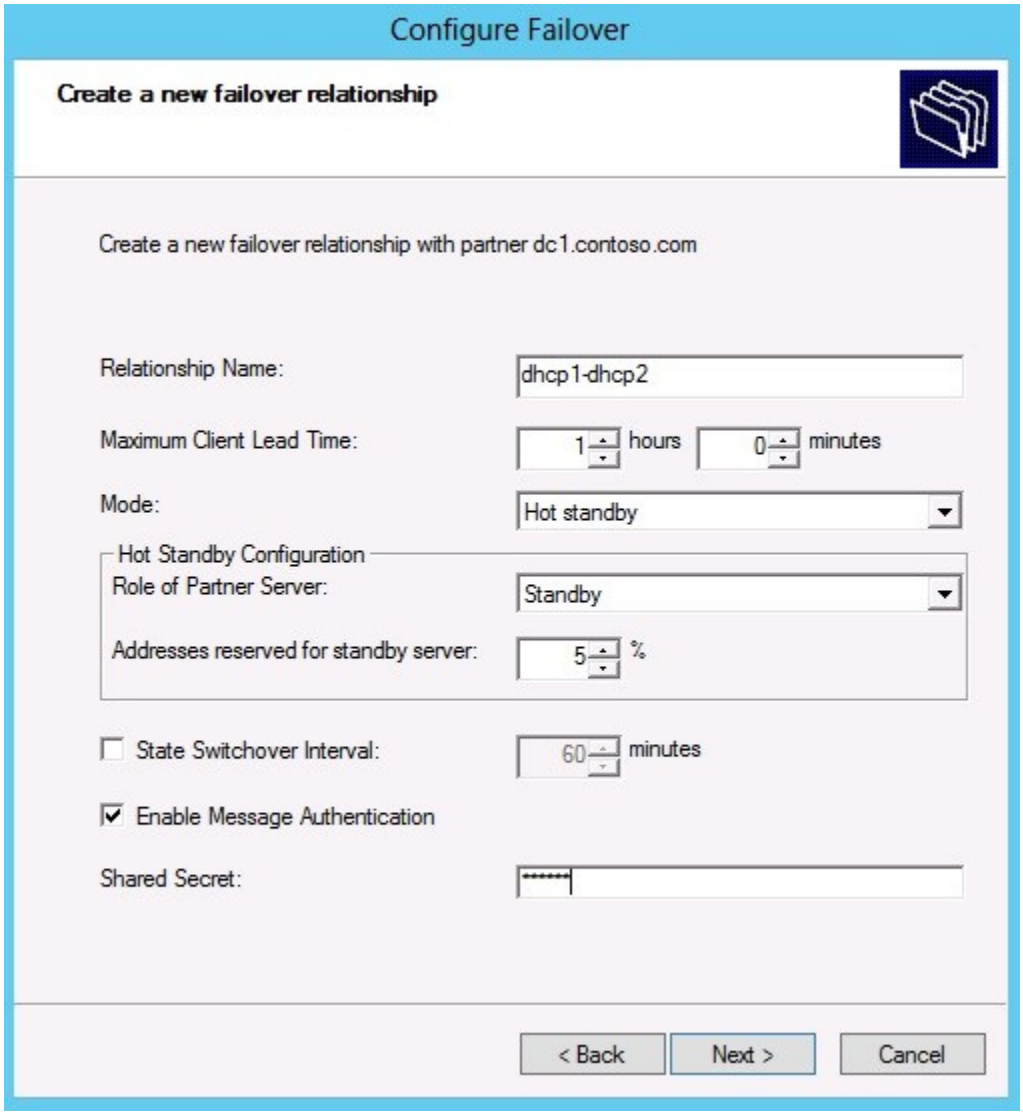

10. Click **Next**, click **Finish**, verify that the status of all the tasks displayed is **Successful**, and then click **Close**.

See also

[What is DHCP Failover?]

[Deploy DHCP Failover]

[Step-by-Step: Configure DHCP for Failover]

Replicate DHCP Failover Using the Command Line

Use the following procedures to replicate DHCP settings using Windows PowerShell. When you have completed the procedures in this topic, return to the parent checklist: Checklist: Deploy DHCP Failover.

Replication is the process of copying settings from one or more failover-enabled DHCP scopes to on a DHCP server to the server's failover partner.

Important
Replication can be initiated from either DHCP server in a failover relationship. Settings on the failover partner server will be overwritten with the settings that are configured on the server where replication is initiated. Therefore, it is important to always ensure replication is initiated on the server that is configured with the settings you wish to use.

This topic provides steps to replicate scope settings using the command line. For steps to replicate scope settings using the DHCP console, see Replicate DHCP Failover Using the DHCP Console.

Three types of replication are available:

- **Server replication**: This type of replication copies the configuration of all failover-enabled scopes on the local server to its failover partner servers.
- **Relationship replication**: This type of replication copies the configuration of all DHCP scopes that are part of a selected failover relationship to the failover partner server.
- **Scope replication**: This type of replication copies the configuration of a selected DHCP scope to the failover partner server.

Choose one of the following procedures based on the type of replication you wish to perform:

- Replicate failover settings at the server level using the command line
- Replicate failover settings at the relationship level using the command line
- Replicate failover settings at the scope level using the command line

Replicate failover settings at the server level using the command line

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the server level using the command line

1. On a computer with the Windows PowerShell DHCP Server module installed, right-click **Windows PowerShell** and then click **Run as Administrator**.
2. Type the following at the command line, and then press ENTER. Replace the example

values with values used in your environment.

```
Invoke-DhcpServerv4FailoverReplication –ComputerName
dhcp1.contoso.com -Force
```

3. Verify that the scope IDs are listed in the command output for all failover enabled scopes on the DHCP server.

In this example, the primary DHCP server is named **dhcp1.contoso.com**. If the **ComputerName** parameter is not specified, the command is run on the local server.

Tip
Windows PowerShell DHCP Server module can be installed using the Add Roles and Features Wizard under Features\Remote Server Administration Tools\Role Administration Tools\DHCP Server Tools. The DHCP Server module is also installed by default when you install the DHCP Server role service.

Replicate failover settings at the relationship level using the command line

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the relationship level using the command line

1. On a computer with the Windows PowerShell DHCP Server module installed, right-click **Windows PowerShell** and then click **Run as Administrator**.
2. Type the following at the command line, and then press ENTER. Replace the example values with values used in your environment.

```
Invoke-DhcpServerv4FailoverReplication –ComputerName
dhcp1.contoso.com –Name dhcp1-dhcp2 -Force
```

3. Verify that the scope IDs are listed in the command output for all failover enabled scopes that are configured to use the specified failover relationship.

In this example, the primary DHCP server is named **dhcp1.contoso.com** and the failover relationship name is **dhcp1-dhcp2**. If the **ComputerName** parameter is not specified, the command is run on the local server.

Tip
Windows PowerShell DHCP Server module can be installed using the Add Roles and Features Wizard under Features\Remote Server Administration Tools\Role Administration Tools\DHCP Server Tools. The DHCP Server module is also installed by default when you install the DHCP Server role service.

Replicate failover settings at the scope level using the command line

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the scope level using the command line

1. On a computer with the Windows PowerShell DHCP Server module installed, right-click **Windows PowerShell** and then click **Run as Administrator**.
2. Type the following at the command line, and then press ENTER. Replace the example values with values used in your environment.

```
Invoke-DhcpServerv4FailoverReplication -ComputerName
dhcp1.contoso.com -ScopeID 10.0.10.0,10.0.20.0 -Force
```

3. Verify that the scope IDs are listed in the command output for all scopes that were specified on the command line.

In this example, the primary DHCP server is named **dhcp1.contoso.com**, and the DHCP scope IDs specified are **10.0.10.0** and **10.0.20.0**. If the **ComputerName** parameter is not specified, the command is run on the local server.

Tip

Windows PowerShell DHCP Server module can be installed using the Add Roles and Features Wizard under Features\Remote Server Administration Tools\Role Administration Tools\DHCP Server Tools. The DHCP Server module is also installed by default when you install the DHCP Server role service.

See also

What is DHCP Failover?

Deploy DHCP Failover

Step-by-Step: Configure DHCP for Failover

Replicate DHCP Failover Using the DHCP Console

Use the following procedures to replicate DHCP settings using the DHCP console. When you have completed the procedures in this topic, return to the parent checklist: Checklist: Deploy DHCP Failover.

Replication is the process of copying settings from one or more failover-enabled DHCP scopes to on a DHCP server to the server's failover partner.

Important

Replication can be initiated from either DHCP server in a failover relationship. Settings on the failover partner server will be overwritten with the settings that are configured on the server where replication is initiated. Therefore, it is important to always ensure replication is initiated on the server that is configured with the settings you wish to use.

This topic provides steps to replicate scope settings using the DHCP console. For steps to replicate scope settings using the command line, see Replicate DHCP Failover Using the Command Line.

Three types of replication are available:

- **Server replication**: This type of replication copies the configuration of all failover-enabled scopes on the local server to its failover partner servers.
- **Relationship replication**: This type of replication copies the configuration of all DHCP scopes that are part of a selected failover relationship to the failover partner server.
- **Scope replication**: This type of replication copies the configuration of a selected DHCP scope to the failover partner server.

Choose one of the following procedures based on the type of replication you wish to perform:

- Replicate failover settings at the server level using the DHCP console
- Replicate failover settings at the relationship level using the DHCP console
- Replicate failover settings at the scope level using the DHCP console

Replicate failover settings at the server level using the DHCP console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the server level using the DHCP console

1. On a DHCP server, open Server Manager, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click **IPv4**, and then click **Replicate Failover Scopes**.
3. Click **OK** in the alert that appears.
4. In the **Failover Scope Configuration Replication** dialog box, verify that replication was successful, and then click **Close**.

Replicate failover settings at the relationship level using the DHCP console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the relationship level using the DHCP console

1. On a DHCP server, open Server Manager, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click any failover-enabled scope, and then click **Replicate Relationship**.
3. Click **OK** in the alert that appears.
4. In the **Failover Scope Configuration Replication** dialog box, verify that replication was successful, and then click **Close**.

Replicate failover settings at the scope level using the DHCP console

Membership in the **Administrators** group, or equivalent, is the minimum required to complete these procedures. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To replicate failover settings at the scope level using the DHCP console

1. On a DHCP server, open Server Manager, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click any failover-enabled scope, and then click **Replicate Scope**.
3. In the **Failover Scope Configuration Replication** dialog box, verify that replication was successful, and then click **Close**.

See also

What is DHCP Failover?

Deploy DHCP Failover

Step-by-Step: Configure DHCP for Failover

Migrate to DHCP Failover

Use the following procedures to migrate from split-scope or clustered DHCP to DHCP failover. When you have completed the procedures in this topic, return to the parent checklist: Checklist: Deploy DHCP Failover.

The following table lists example DHCP server host names used in the procedures described in this topic:

Host name	Information
DHCP2012-1	The DHCP server that will be used to initially configure DHCP failover.
DHCP2012-2	The DHCP failover partner server for DHCP2012-1.
DHCP2008-1	The currently active DHCP server, providing 80% of the leases from a DHCP split-scope, or serving as the active node in a clustered DHCP deployment. DHCP2008-1 might also be the only DHCP server if no DHCP high availability configuration is currently deployed.
DHCP2008-2	The current DHCP server providing 20% of the leases from a DHCP split-scope with DHCP2008-1, or serving as the backup node for DHCP2008-1 in a clustered DHCP deployment.

Migrate to DHCP failover

Migration consists of the following steps:

1. Install the DHCP Server role on DHCP2012-1 and DHCP2012-2
2. Export DHCP settings and leases from DHCP2008-1 and DHCP2008-2
3. Import DHCP settings and leases to DHCP2012-1
4. Import server level settings to DHCP2012-2
5. Configure DHCP failover on DHCP2012-1
6. Update DHCP relay agents and complete migration

Install the DHCP Server role on DHCP2012-1 and DHCP2012-2

You can install the DHCP server role using the command line or with the DHCP console.

1. To install the DHCP server role using the command line
2. To install the DHCP server role using the DHCP console

To install the DHCP server role using the command line

1. At an elevated Windows PowerShell prompt, type the following command, and press ENTER:

   ```
   Add-WindowsFeature –IncludeManagementTools Dhcp
   ```

2. Verify that **True** is displayed under **Success**.

3. Type the following command and then press ENTER:

   ```
   Netsh dhcp add securitygroups
   ```

   Verify that **Command completed successfully** is displayed.
4. Type the following command and then press ENTER:

   ```
   Add-DhcpServerInDC dhcp2012-1 192.168.0.5
   ```
5. At this point, you can launch the DHCP console and see that the DHCP server is authorized in Active Directory. However, Server Manager might still contain an alert that post-deployment configuration is required. To remove this alert from Server Manager, you can click **Complete DHCP configuration** and follow the steps in the DHCP Post-Install configuration wizard, or you can type the following command and press ENTER:

   ```
   Set-ItemProperty –Path
   registry::HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ServerManager
   \Roles\12 –Name ConfigurationState –Value 2
   ```
6. Refresh the view in Server Manager and verify that the DHCP post-deployment configuration alert is no longer displayed.

To install the DHCP server role using the DHCP console

1. In Server Manager, click **Add roles and features**.
2. Click **Next** three times, and then on the **Select server roles** page, select the **DHCP Server** checkbox.
3. Click **Add Features** in the popup that appears.
4. Click **Next** three times, and then on the **Confirm installation selections** page, click **Install**.
5. Confirm that **Installation succeeded…** is displayed and then click **Close**.
6. Click the alert flag in Server Manager and then click **Complete DHCP configuration**.
7. Click **Next**, provide credentials to authorize the server in Active Directory, and then click **Commit**.
8. Verify that authorization was successful and then click **Close**.

Export DHCP settings and leases from DHCP2008-1 and DHCP2008-2

If you are using clustered DHCP, you only need to export DHCP settings from DHCP2008-1. Perform the following procedure for both DHCP2008-1 and DHCP2008-2 only if you are using a split-scope design and need to replicate DHCP leases from both DHCP servers.

Important

You cannot import DHCP leases for the same scope from two different DHCP servers without first consolidating the exported files into a single file. You can also configure DHCP2008-1 to issue 100% of the leases for a split-scope and allow leases on DHCP2008-2 to expire before performing the next procedure.

Also see Migrate split scope DHCP deployment to Windows Server 2012 DHCP Failover in the Microsoft TechNet Script Center. This script automates procedures in this topic, including export and import of leases from your existing DHCP servers, and creation of a DHCP failover relationship on the new DHCP servers. The script does not handle all error conditions or perform automatic rollback if an error is encountered. This script is provided to demonstrate automation with Windows PowerShell for migration to DHCP failover in a test environment. It is not recommended for use in a production environment.

To export DHCP settings and leases

1. On DHCP2012-1, open an elevated Windows PowerShell prompt, type the following commands, and then press ENTER:

   ```
   mkdir C:\export
   mkdir C:\backup
   Export-DhcpServer –ComputerName DHCP2008-1 –Leases –File
   C:\export\DHCP2008-1exp.xml –Verbose
   ```

   Note that DHCP configuration settings and the lease database is being exported from a remote DHCP server to a file on the local server. Alternatively, you can export to a file on DHCP2008-1 and then import settings over the network to DHCP2012-1.
2. If needed, perform the export step again specifying a **Computername** of DHCP2008-2 and exporting to the file **DHCP2008-2exp.xml**. When you are finished, you should have one XML file for every DHCP server that has leases you wish to import to DHCP2012-1.
3. If you have multiple files that contain leases for the same scope ID, these files must be consolidated before importing data to DHCP2012-1.

To consolidate the lease database

a. At an elevated Windows PowerShell prompt, type the following commands and press ENTER:

   ```
   notepad C:\export\DHCP2008-1exp.xml
   notepad C:\export\DHCP2008-2exp.xml
   ```

b. Determine which file will be the primary file to use for import. Copy information from the secondary file and paste this information into the primary file.

c. Click the secondary file's notepad window and type CTRL-F to bring up the **Find** dialog box.

d. Next to **Find what**, type **<Leases>** and then highlight all text that is between **<Leases>** and **</Leases>**. See the following example.

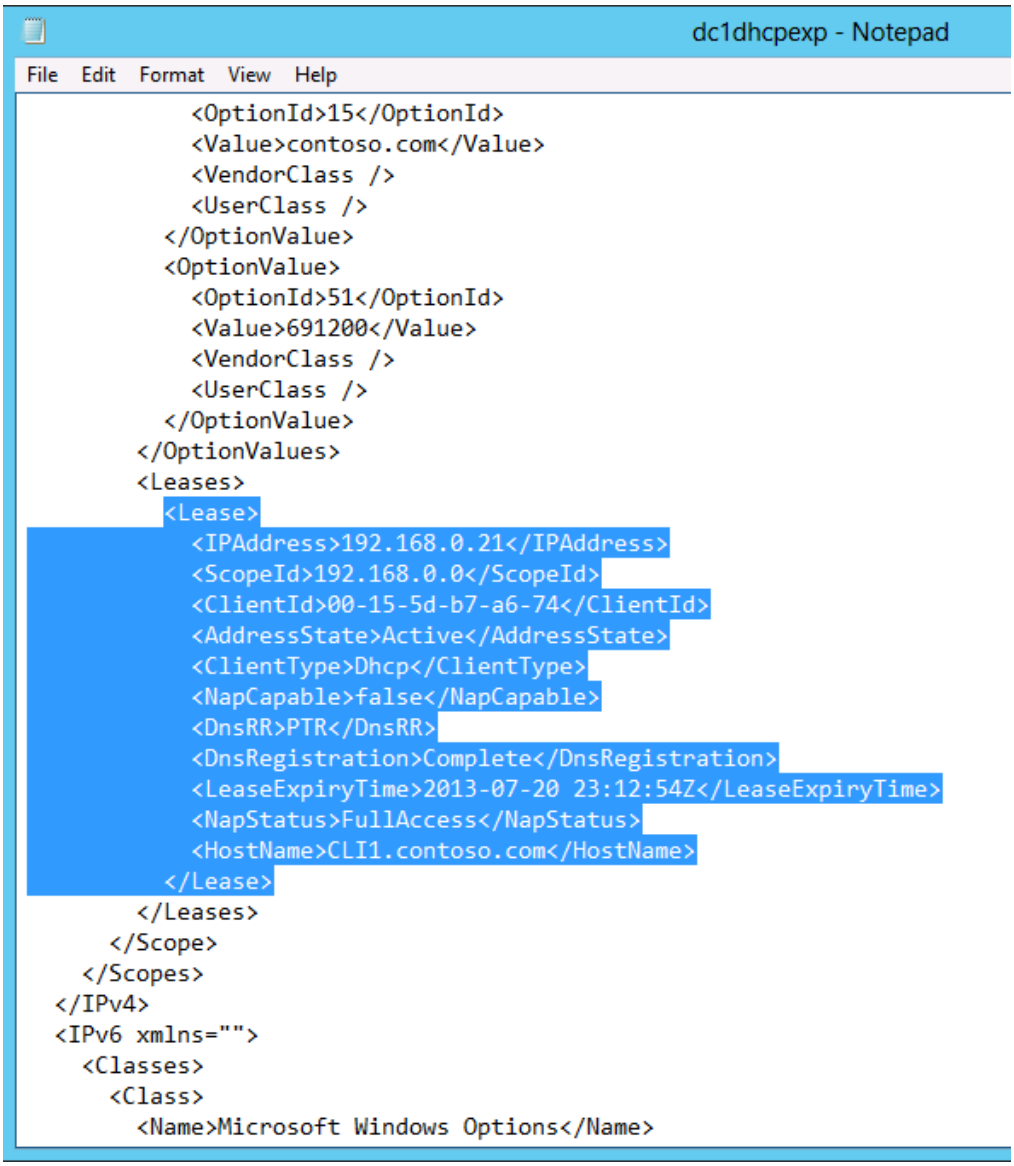

e. Type CTRL-C to copy all the text from this file, or click **Edit** on the menu and then click **Copy**.

f. Click the primary file's notepad window and locate the Leases section by searching for **<Leases>**.

g. Place the mouse cursor after the last </Lease> and before the </Leases> tag in this file and then type CTRL-V or click **Edit** on the menu and then click **Paste**.

h. Save the primary file by typing CTRL-S or clicking **File** on the menu and then clicking **Save**.

When you have completed this procedure, all active leases will be consolidated in the **DHCP2008-1exp.xml** file in preparation for import to DHCP2012-1.

Import DHCP settings and leases to DHCP2012-1

Before you perform this procedure, verify that you have consolidated and saved leases and settings from your current DHCP servers in a file on the new DHCP server. The following procedure assumes DHCP settings and leases are saved in the **C:\export\DHCP2008-1exp.xml** file on a DHCP server with a host name of **DHCP2012-1**.

To import DHCP settings and leases to DHCP2012-1

1. At an elevated Windows PowerShell prompt, type the following commands and press ENTER:

```
Import-DhcpServer -ComputerName DHCP2012-1 -Leases -File
C:\export\DHCP2008-1exp.xml -BackupPath C:\backup\ -Verbose -
Force
```

Note

The **–BackupPath** parameter is used to specify a path where the current DHCP server's database is backed up prior to making any configuration changes as part of the import operation. To roll back the import operation, use the **Restore-DhcpServer** cmdlet and specify **C:\backup** for the value of the **-Path** parameter. For more information, see Restore-DHCPServer.

In addition to leases, this command imports the server-level configuration of DHCP2008-1 to DHCP2012-1. Server-level settings include the following:

- Custom vendor classes and user classes
- Custom option definitions
- Server-level option values
- Server-level policies
- MAC address filters
- Custom conflict detection settings

By default, all scopes will be imported. If a scope already exists on the destination server it will not be imported. To specify only some scopes for import, use the **–ScopeId** parameter. For example:

```
Import-DhcpServer -ComputerName DHCP2012-1 -Leases -File C:\export\DHCP2008-1exp.xml -
BackupPath C:\backup\ -Verbose -ScopeId 10.1.1.0, 10.1.2.0 -Force
```

These server-level settings must also be imported to the new DHCP failover partner server (DHCP2012-2) with the following procedure.

Important

Verify that options and settings on DHCP2012-1 are correct before proceeding. Options such as DNS servers, reservations, or filters also might need to be updated on DHCP2012-1.

Import server level settings to DHCP2012-2

Use the following procedure to import server-level settings from DHCP2008-1 to the failover partner server DHCP2012-2.

To import server level settings to DHCP2012-2

1. At an elevated Windows PowerShell prompt on DHCP2012-2, type the following commands and press ENTER:

```
mkdir C:\backup
Import-DhcpServer –ComputerName DHCP2012-2 –File
C:\export\DHCP2008-1exp.xml -BackupPath C:\backup\ -Verbose -
Force
```

   The DHCP server's database is backed up locally on the same server where the DHCP Server service is running even if you run the Import-DhcpServer cmdlet remotely.

Configure DHCP failover on DHCP2012-1

Next, configure DHCP failover for some or all scopes on DHCP2012-1, specifying DHCP2012-2 as the failover partner server. For procedures to configure DHCP failover, see Configure DHCP Failover using the Command Line or Configure DHCP Failover using the DHCP Console.

Update DHCP relay agents and complete migration

Finally, update DHCP relay agents to use the new DHCP servers and complete DHCP migration. For more information about relay agents and DHCP failover, see DHCP Failover Architecture.

After leases and settings have been imported to DHCP2012-1 and DHCP2012-2, settings are verified, and DHCP relays are updated, you can stop the DHCP Server service on DHCP2008-1 and DHCP2008-2 so that new DHCP leases are only issued from DHCP2012-1 or DHCP2012-2.

See also

What is DHCP Failover?

Deploy DHCP Failover

DHCP Failover Requirements

DHCP Policy Based Assignment (PBA) Scenario Guide

This guide is intended for use by a network or systems administrator. The guide provides information about a new functionality in Windows Server® 2012 called policy based assignment

(PBA) for the DHCP Server role. This guide discusses some key scenarios that are enabled by DHCP PBA.

In this guide

The following topics are available in this guide:

Introduction to DHCP Policies discusses benefits and provides an overview of DHCP PBA.

Scenarios discussed in this guide:

- Scenario: Manage the network configuration of virtual machines
- Scenario: Secure a subnet to a specific set of clients
- Scenario: Customize lease duration based on device type

For each scenario you will find a problem description with an example of conditions and settings that can be used to enable a solution. This guide is also updated periodically to incorporate new scenarios, information, and corrections.

See also

Step-by-Step: Configure DHCP Using Policy-based Assignment

Introduction to DHCP Policies

The DHCP Server role in Windows Server® 2012 introduces a new feature that allows you to create IPv4 policies that specify custom IP address and option assignments for DHCP clients based on a set of conditions.

The policy based assignment (PBA) feature allows you to group DHCP clients by specific attributes based on fields contained in the DHCP client request packet. PBA enables targeted administration and greater control of the configuration parameters delivered to network devices with DHCP.

Why DHCP PBA?

Consider the following scenarios:

1. A subnet has a mix of different types of clients: desktop computers, printers, IP phones, and other devices. You want different types of clients to get IP addresses from different IP address ranges within the subnet. This is possible using DHCP policies if the devices have different vendors. For example:

   Printers can get IP addresses from 10.10.10.1 to 10.10.10.9.

   IP phones can get IP addresses from 10.10.10.10 to 10.10.10.49.

   Desktop computers can be assigned IP addresses from 10.10.10.50 to 10.10.10.239.

   Additional devices can be assigned IP addresses of 10.10.10.240 to 10.10.10.254.

By specifying a different IP address range for different device types, you can more easily identify and manage devices on the network.

2. In a subnet which has a mix of wired and mobile computers, you might want to assign a shorter, 4 hour lease duration to mobile computers and longer, 4 day lease duration to wired computers.
3. You want to control who gets access to the network by providing a DHCP lease to only a known set of clients based on MAC address.
4. Employees bring in their own devices such as smartphones and tablets to work and you want to manage network traffic or control network access based on device type.
5. You want to provide a different set of scope options to different types of devices. For example, IP phones can get a different **Boot Server Host Name** (TFTP server) and **Bootfile Name** option.

DHCP policies provide a very useful tool to achieve these goals. See the following example.

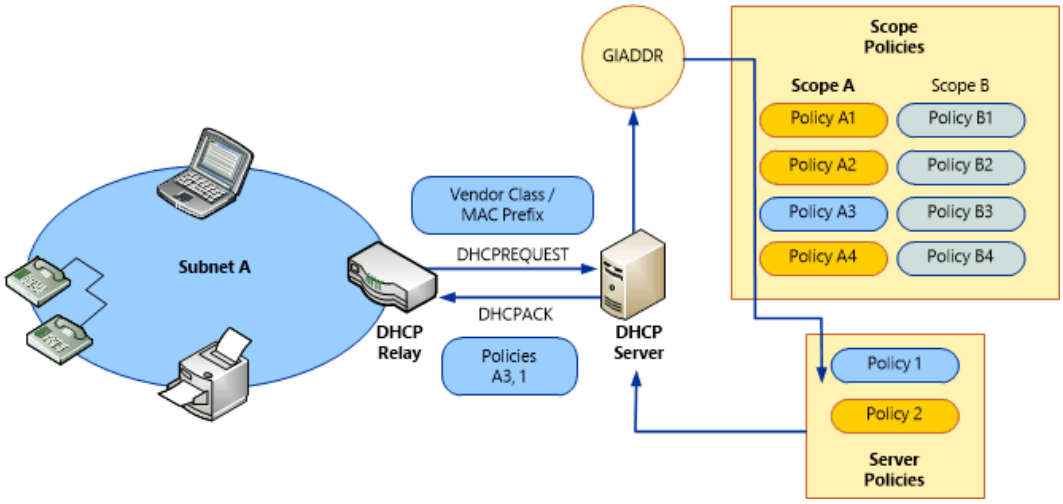

In this example:

- Subnet A contains DHCP client devices of several different types including workstations, printers, and IP phones.
- A DHCP server on another subnet is configured to provide leases to these devices from scope A.
- Polices are configured at the scope level to control IP address range and at the server level to specify lease duration.

DHCP client requests are processed as follows:

1. A client on subnet A submits a DHCPREQUEST that is sent to the DHCP server via DHCP relay.
2. The client's vendor class and MAC prefix are included in the DHCPREQUEST packet along with the gateway IP address (GIADDR).
3. The DHCP server uses the GIADDR to determine that the client requires a lease from scope A, and begins processing policies in that scope.

4. Since scope B does not apply, these policies are ignored.
5. Based on the vendor class and MAC prefix values provided, the client request matches conditions of policy A3.
6. After all scope polices are processed, server level policies are processed and the client also matches conditions of policy 1.
7. After all policies are processed, the DHCP server returns an IP address configuration to the client using the settings specified in policies A3 and 1.

Based on the client's MAC address it is determined that the device is a printer (it matches policy A3). It is assigned the first available IP address in the IP address range 10.10.10.1 to 10.10.10.9, with a lease duration of 14 days.

In Windows Server 2008 R2 and previous operating systems, if you want to specify the IP address range for a specific set of clients or devices, or assign different option values based on device type, the only way to achieve this is to configure a scope with individual reservations. This method can require high effort, and is difficult to manage on an ongoing basis.

DHCP policies in Windows Server 2012 provide much more flexibility to assign unique IP addresses and options to specific DHCP clients in a single subnet, or in multiple subnets.

Note
See Policy processing to understand how settings are applied when they are configured in multiple policies, in reservations, at the scope level, or at the server level.

How DHCP PBA works

DHCP policies are rules that you can define for DHCP clients. You can define a single policy, or several. Characteristics of DHCP policies include:

- **Policy level**: Polices can apply at the server level or the scope level. Server level policies are processed for all DHCP client requests received by the server. Scope level policies are processed only for DHCP client requests that apply to a specific scope.
- **Processing order**: Each policy has an associated processing order that is unique within a server or scope. Policies with a lower numbered processing order are evaluated before higher number policies. If both scope and server level policies apply to a client, the scope level policies are always processed before any server level policies.
- **Conditions**: The conditions specified in a policy enable you to evaluate clients based on fields that are present in the DHCP client request. If a client request matches the conditions in the policy, the settings associated with a policy will be applied to the client by the DHCP server when it responds to the DHCP request.
- **Settings**: Settings are network configuration parameters (ex: IP address, options, lease duration) that are provided to DHCP clients in the DHCP server response. Settings enable you to group clients by applying the same set of network parameters to them.
- **Enabled/Disabled**: Policies at the scope or server level can also be enabled or disabled. A policy that is disabled is skipped when processing incoming DHCP client requests.

To create a policy at the server level using the Windows interface, open the DHCP console, navigate to IPv4, right-click **Policies** and then click **New Policy**.

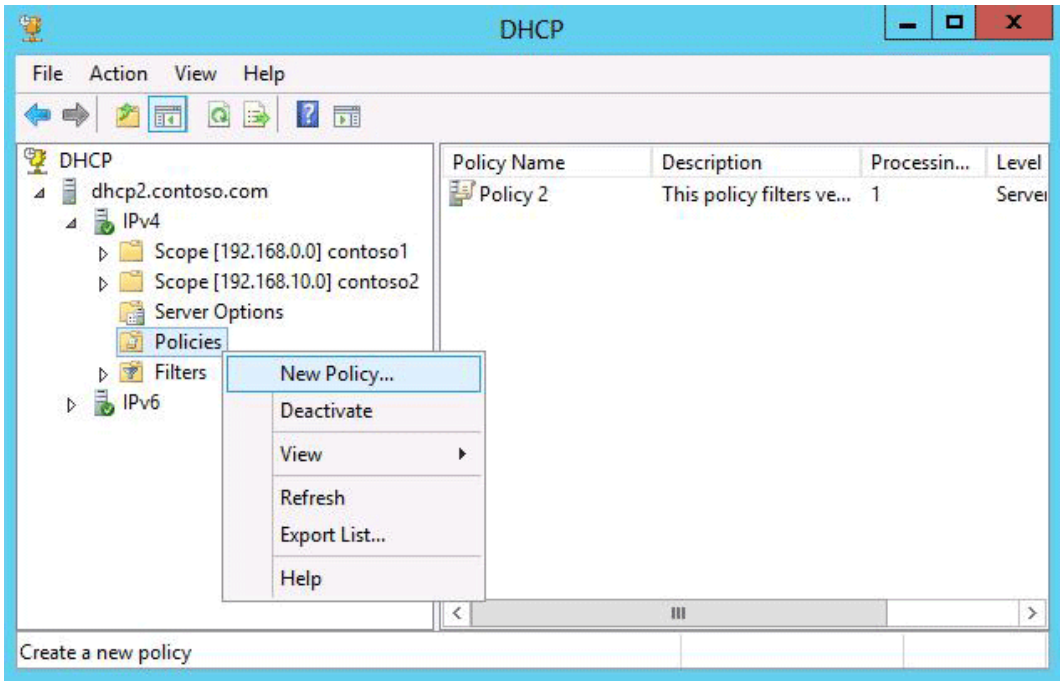

If other server level policies exist, they are displayed in the details pane and can be modified by right-clicking the policy and then clicking **Move Up**, **Move Down**, **Disable**, **Enable**, **Delete**, or **Properties**.

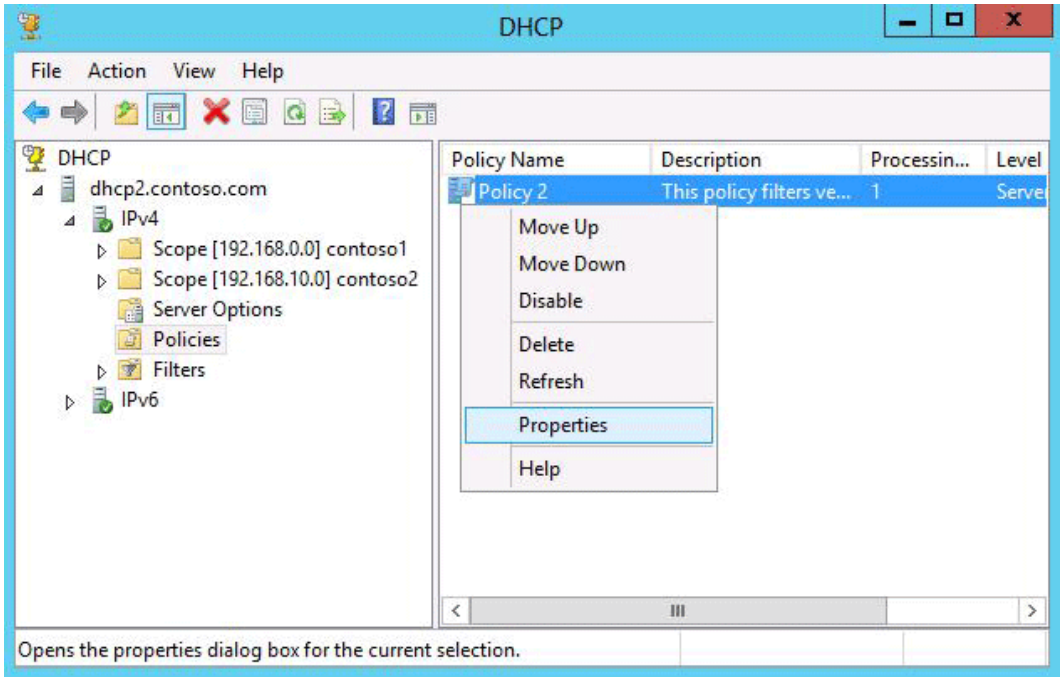

To create a policy at the scope level using the Windows interface, open the DHCP console, navigate to an IPv4 scope, right-click **Policies** and then click **New Policy**. If other scope level policies exist, they are displayed along with any server level policies that exist. You can modify existing scope level policies by right-clicking them. You cannot modify a server level policy at the scope level.

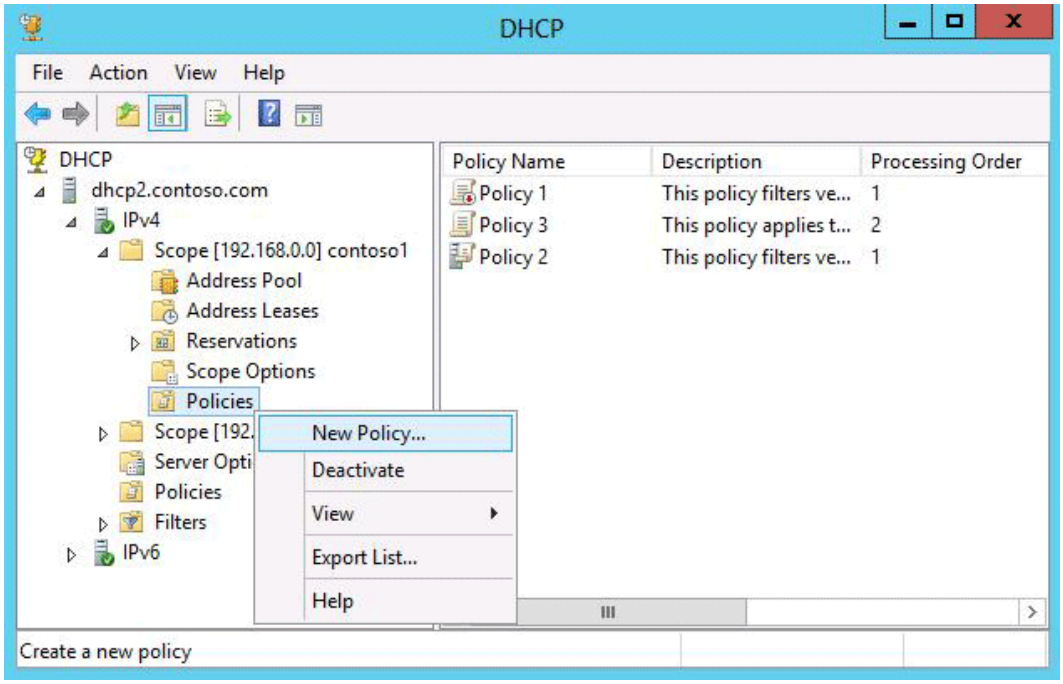

You must provide a unique policy name when creating a new policy. A policy description is optional. A policy must have at least one condition.

Policy settings are optional, but DNS settings are included by default so it is not possible to have a policy with no settings. To view DNS settings for a policy, right-click the policy, click **Properties**, and then click the **DNS** tab.

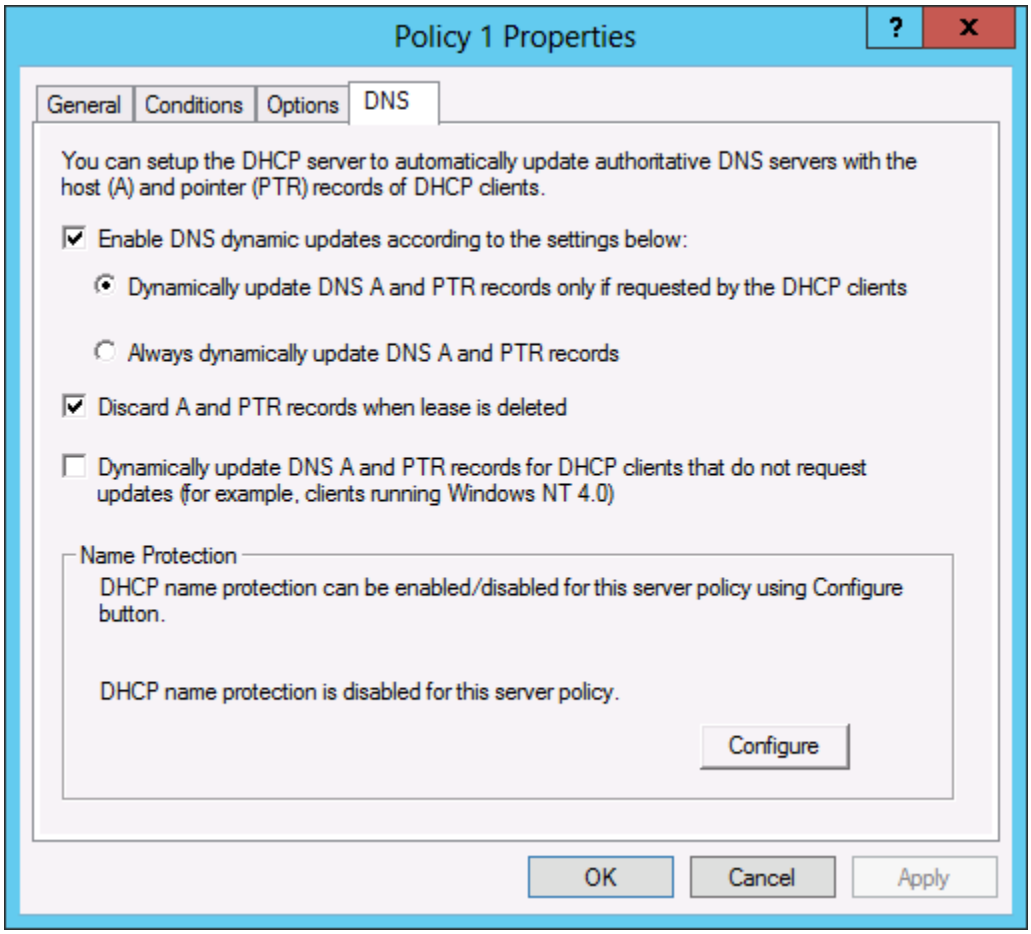

DHCP policy conditions and settings

The following conditions and settings are available when creating a policy:

Conditions: Vendor Class, User Class, MAC Address, Client Identifier, Relay Agent Information.

Settings: IP Address Range, Standard DHCP Options, Vendor Specific DHCP Options.

Conditions

In Windows Server 2012, you can specify five conditional criteria to evaluate and group DHCP clients:

MAC Address: The media access control (MAC) address or link-layer address of the client.

Vendor Class: Vendor managed DHCP option assignments.

User Class: Non-standard DHCP option assignments.

Client Identifier: The client identifier (ClientID) is typically a MAC address. In the case of PXE clients, it can be the GUID of the network interface card (NIC).

Relay Agent Information, including sub-options: **Agent Circuit ID**, **Agent Remote ID**, and **Subscriber ID**: Information inserted into DHCP client requests by a DHCP relay using option 82.

The operators that can be used with these conditions are **equals** and **not equals**. You can also use a trailing wildcard with MAC address, Vendor Class, User Class and Client Identifier conditions to perform a partial match. By combining the equals or not equals with a wildcard in the condition you can effectively achieve a **starts with** or **does not start with** condition.

You can either have a single condition in a policy or a set of conditions which can be OR'ed or AND'ed. For example, "Vendor Class Equals Cisco IP Phone 7940" is a condition (**Cisco IP Phone 7940** is the value of Vendor Class for Cisco IP Phone version 7940). Also, a grouping such as "User Class Equals LabComputers" AND "MAC Address Not Equals 00-11-22*" is a group of two conditions. Each policy is created with either a single condition or a set of such conditions.

Important

Using multiple criterion values: When you list multiple values for a single criterion, such as "User Class Equals (valueA, valueB, valueC)" or "MAC Address Not Equals (value1, value2, value3)" these values are interpreted as being OR'd if the EQ (equals) operator is used, but they are AND'd if the NEQ (not equals) operator is used.

An incoming client request for an IP address and options from the DHCP server matches a policy if the client satisfies the cumulative set of conditions in the policy. A client that does not match conditions of any policy is granted an IP address lease from the rest of the IP address range of the scope, exclusive of all the policy IP address ranges, and is assigned the default option values configured in the scope.

Settings

In Windows Server 2012, three types of policy settings are available that can be applied to DHCP clients:

IP Address Range: A specified sub-range of IP addresses within the scope range. The IP address range setting cannot be specified in a server-level policy.

Standard DHCP Options: Standard DHCP options like default gateway (003 Router) and preferred DNS servers list (006 DNS Servers).

Vendor Specific DHCP Options: Vendor managed DHCP option assignments.

In addition, you can also specify the following settings in policy properties:

DNS settings: DNS registration and Name Protection settings can be specified on the **DNS** tab.

Lease duration: The lease duration can be specified on the **General** tab. See the following example

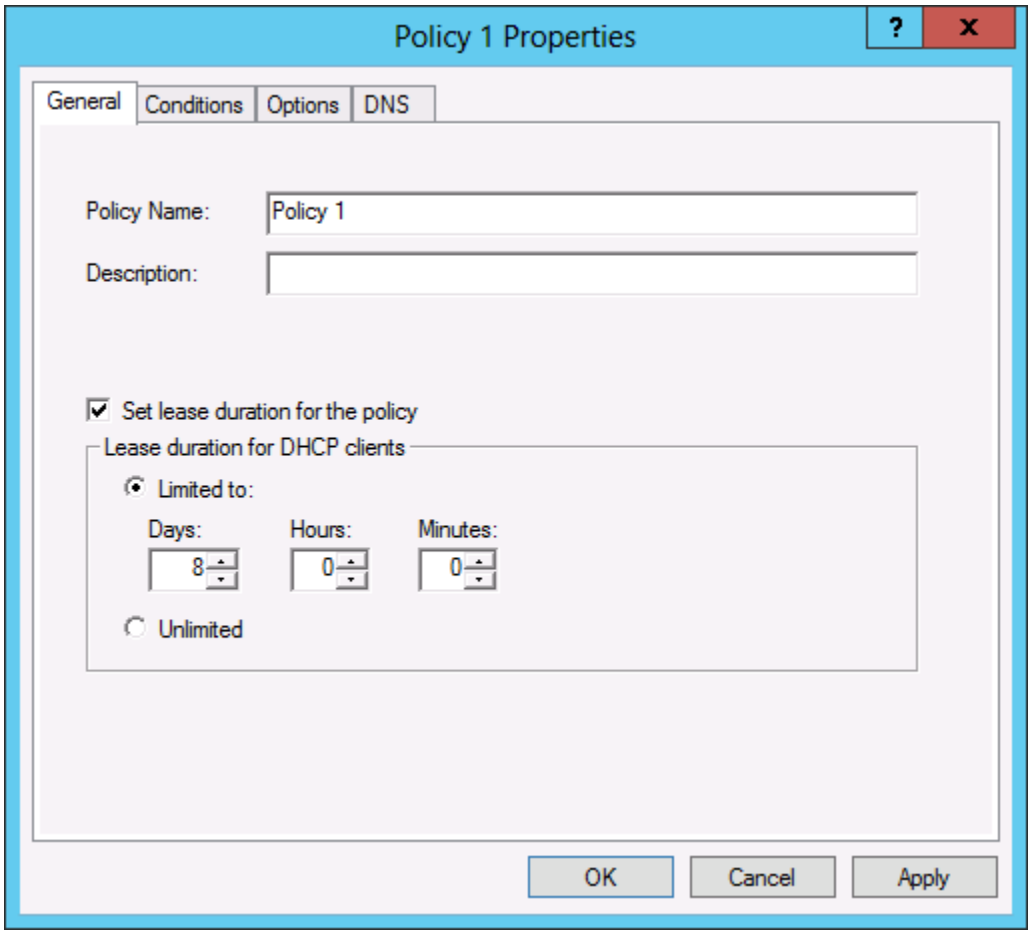

When a client matches the conditions of a policy, the DHCP server responds to the client and includes settings in that policy, provided these settings are not already applied in a higher priority policy or using a reservation. See Policy processing for more information.

A policy can specify an IP address range with no options, or it can specify options with no IP address range, or it can specify both, or it can specify neither. A policy can also specify multiple standard options, vendor-specific options, or both.

Policy processing

Since you can configure multiple policies at both the scope level and server level, each policy is assigned a processing order. The processing order can also be modified, assuming more than a single policy exists. The following conditions exist:

1. When processing DHCP client requests, the DHCP server evaluates each client request against the conditions in all applicable policies, based on their processing order.

2. Scope level policies are processed first by the DHCP server, followed by server wide policies. Theoretically, a client can match the conditions of several scope policies and also several server policies.
3. If a client satisfies the conditions of more than 1 policy, it will get the combined settings from all policies that it matched. If the same option setting is provided in multiple policies, the client will use the setting from the first policy that is processed.

For example, assume that policy-1 has an option value for **003 Router** and policy-2 has an option value for **006 DNS Servers**, and a client request matches both policies. The DHCP server will assign a default gateway value (003 Router) using policy-1 and a DNS server value using policy-2. However, if policy-1 has the higher processing priority (a value of "1") and also has an option value for DNS server, the client will get both the router and DNS server option values from policy-1. The DNS server option value in policy-2 is ignored because policy-2 has a lower processing priority (a value of "2").

A policy does not need to be configured with all option values that you have already configured at the scope or server level. If a policy client has requested an option which is not present in the policy but has been configured in scope level or server level options, these options are applied to the client in the server response. However, if you wish to specify options for certain clients, you can include these option settings in policies and they will have a higher priority than scope or server level options. The only type of option setting that has a higher priority than those configured in polices are options that you configure for a reservation.

The priority for options settings is **reservation** > **scope policy** > **server policy** > **scope-level** > **server-level**. See the following figure.

If a DHCP client obtains option settings because it matched a reservation, it will ignore the same options settings if they are present in any scope or server polices, or configured globally at the scope or server level.

Deploying DHCP policies

A common reason to deploy DHCP policies is to provide unique settings to different types of devices on the network. Two common methods used to identify device type include:

1. **Vendor class**: A text string is sent in option 60 by most DHCP clients that identifies the vendor and therefore the type of the device.
2. **MAC address prefix**: The first three bytes of a MAC address is called the organizationally unique identifier (OUI), and can be used to identify the vendor or manufacturer of a device.

For example, you might decide to group DHCP clients on the network by device type. After assigning IP address ranges to devices, you can configure your router to handle network traffic from each IP address range differently. In effect, you can configure network access control for a

class of devices using DHCP policies. You might also manage network traffic by configuring route options such as default gateway (option 003) and classless static routes (option 121) based on device type.

It is often desirable to configure a short lease duration for wireless devices, and grant a longer lease to wired devices. Since wireless access points are typically capable of behaving as a DHCP relay agent, or are connected to a DHCP relay, they can provide DHCP option 82 (DHCP relay agent). Presence of a specific value in the relay agent option can therefore indicate that the DHCP client is a wireless device.

With DHCP policies, you can configure a policy with a condition based on the relay agent information option value that identifies wireless clients and provides a shorter lease duration. Other DHCP clients in the scope will continue to be provided with the longer lease duration configured at the scope level.

These scenarios and others are discussed in detail in this guide.

See also

Scenario: Manage the network configuration of virtual machines

Scenario: Secure a subnet to a specific set of clients

Scenario: Customize lease duration based on device type

Scenario: Manage the network configuration of virtual machines

See the following sections in this topic:

Problem description

MAC addressing in Hyper-V

Configure a MAC address based DHCP policy

Conclusion

Problem description

More and more organizations are leveraging the advantages of virtualization in today's business environment. However, virtual networks also require physical resources, including:

- **Compute resources**: Physical servers that host virtual machines, such as a Hyper-V cluster are required.
- **Storage resources**: Clustered file servers such as the Scale-Out File Server in Windows Server 2012 are a critical component of any virtualization solution.
- **Network resources**: Physical devices such as routers and switches that connect virtual networks to other networks are necessary.

- **Management resources**: Physical servers and devices that provide automation, access control, monitoring, and scalability for virtual networks, such as System Center Virtual Machine Manager are a critical component.

As a consequence, organizations often need to manage a hybrid environment consisting of both physical and virtual resources. See the following diagram.

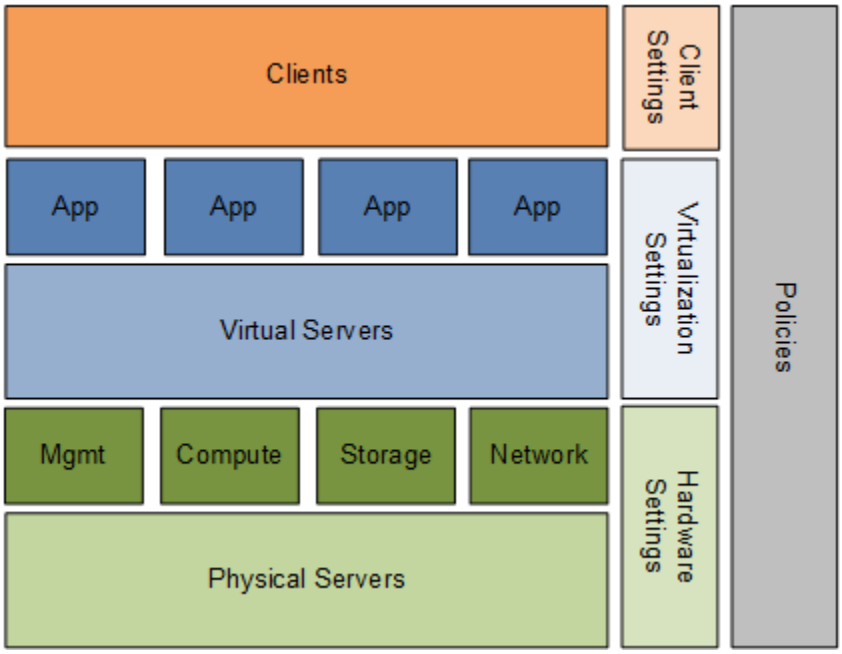

Placing virtual and physical devices on separate subnets enables customization of network settings but adds complexity. Larger, flat networks are typically more efficient and less complex to manage than a design that includes multiple smaller, partitioned networks. DHCP policies can help to simplify management of hybrid environments by enabling greater flexibility and control of the network configuration of DHCP clients.

Virtual machines

An advantage of virtual machines is that they can be added and removed dynamically from the network depending upon load requirements. However, this capability can bring with it some unique requirements. Due to their dynamic nature, virtual machines might require:

A different, often shorter, lease duration

A different default gateway for use in traffic management

A different domain name or DNS server

A different IP address range

These configuration parameters can be managed by creating a DHCP policy based on MAC address prefix. DHCP clients that match the policy are assigned IP addresses from a specific IP

address range, a short lease duration, and different DHCP options such as default gateway and DNS server.

MAC addressing in Hyper-V

With Hyper-V, you can configure a virtual machine to use a dynamic or a static MAC address. The default option is to use a dynamic MAC address, which means that Hyper-V will generate a MAC address for the network adapter. To view MAC address settings in Hyper-V, right-click a virtual machine, click **Settings**, and then under **Network Adapter**, click **Advanced Features**. See the following example.

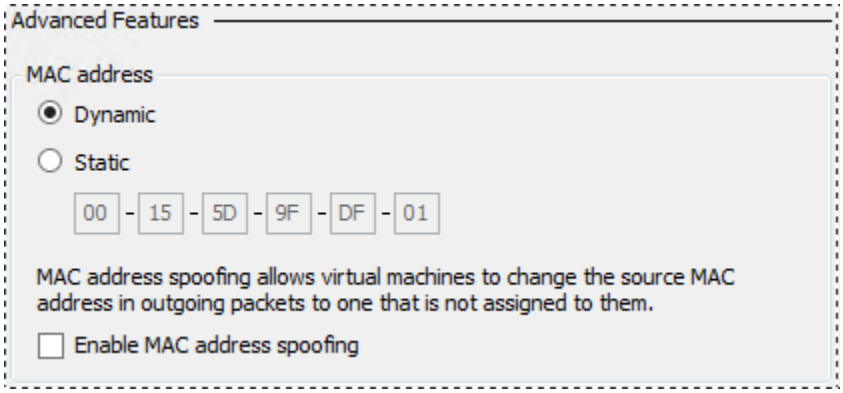

In this example, the Hyper-V server is using a dynamic MAC address range of 00-15-5D-9F-DF-00 to 00-15-5D-9F-DF-FF. To use a static MAC address, select **Static** and then specify a MAC address.

For dynamic MAC addressing in Hyper-V, a range of 256 possible MAC addresses is used by default. The starting and ending addresses are derived from the lowest two octets in an IPv4 address on the Hyper-V server, and will begin with Microsoft's Organizationally Unique Identifier (OUI) of 00-15-5D. The range can be modified, if desired, by editing the **MaximumMacAddress** and **MinimumMacAddress** registry keys located under **HKLM\Software\Microsoft\Windows\NT\CurrentVersion\Virtualization**. This kind of support is provided by all the virtualization platforms.

To group virtual machines in a DHCP policy, create a policy with conditions based on MAC address or MAC address prefix.

For example, if MAC addresses are assigned from the range 00-15-5D-9F-DF-00 to 00-15-5D-9F-DF-FF, you can create a policy based on the MAC address prefix 00-15-5D-9F-DF. If you would like the policy to cater to all virtual machines regardless of the host or host group, use the MAC address prefix 00-15-5D with a wildcard (00155D*). The following procedure demonstrates how to configure this policy.

Configure a MAC address based DHCP policy

Use the following procedure to configure a policy for virtual machines based on MAC addresses. In this procedure, the example MAC address range used is 00-15-5D-9F-DF-00 to 00-15-5D-9F-DF-FF. You must replace the values used in this example with the MAC address range used in your environment. By default, the Hyper-V server will use a MAC address range that includes values of 00 to FF for the sixth octet, with the same first five octets. For more information, see MAC addressing in Hyper-V.

To configure a MAC address based DHCP policy

1. In the DHCP console, choose a scope that will be used to issue leases to both physical and virtual DHCP clients.
2. Under the DHCP scope, right click **Policies** and then click **New Policy**.
3. In the DHCP Policy Configuration Wizard, type **MAC Prefix Policy** next to **Policy Name**, and then click **Next**. You can type a different policy name and enter a description if desired.
4. On the Configure Conditions for the policy page, click **Add**.
5. Next to **Criteria**, choose **MAC Address** from the drop-down list and choose **Equals** next to **Operator**.
6. Next to **Value**, type **00155D9FDF**, select the **Append wildcard** checkbox, click **Add**, and then click **OK**.
7. Click **Next**, and provide an IPv4 address range to be used for clients that match the policy. To provide a range, enter a start and end IP address next to **Start IP address** and **End IP address**, respectively. The percentage of IP address used by this IP address range compared to all available IP addresses in the scope is displayed next to **Percentage of IP address range**. If you do not want to specify a range of IP addresses for the policy, select **No** next to **Do you want to configure an IP address range for the policy**.

   Note
   The IP address range specified must be within the IP address range for the scope. An IP address range setting cannot be specified in server level policies.
8. Click **Next**, and specify DHCP options to be used by clients that match the policy, such as the default route or list of DNS servers.
9. Click **Next**, and on the Summary page review the conditions and settings for this policy. Click **Finish** to create the policy.
10. To configure a custom lease duration for virtual machines, right-click the policy you just created and then click **Properties**. On the **General** tab, select the **Set lease duration for the policy** checkbox and then specify a lease duration.
11. Review the conditions and settings for the policy by clicking each available tab. DNS registration settings can be modified on the **DNS** tab. Click **OK** to apply changes or click **Cancel** to discard changes.

After you have configured this policy, you can also right-click the policy and click **Move Up** or **Move Down** to change its processing order.

Conclusion

Using this policy, new DHCP requests for IPv4 addresses in this scope that are sent from DHCP clients with a MAC address prefix of 00-15-5D-9F-DF will get the IP address and options that are defined in the policy. In this example, only virtual machines that are hosted by the specified Hyper-V server will have this MAC address prefix, and will receive these settings.

You can specify additional MAC address prefix values in the policy by adding them to policy conditions. To enable more than one MAC address prefix to receive the same policy settings, add additional MAC address prefixes on the **Conditions** tab in policy properties and select the **OR** radio button.

See also

Introduction to DHCP Policies

Scenario: Secure a subnet to a specific set of clients

See the following sections in this topic:

Problem description

Creating a scope-based allow-filter list using DHCP policies

Creating a scope-based deny-filter list using DHCP policies

Windows PowerShell

Problem description

Link-layer (MAC address-based) filtering is a feature introduced in Windows Server® 2008 R2 for the DHCP Server role that enabled restriction of network access to DHCP clients based on their MAC address. A limitation of this feature is that it is only possible at the server level. With DHCP policies in Windows Server® 2012 R2, link-layer filtering is now possible at the scope level.

Consider a situation where you want to allow only a specific set of computers or devices on a subnet. To control which clients can be allowed on the subnet, you can configure a policy on the DHCP server to provide IP addresses only to specific clients. You can do this by creating an allow list, or a deny list.

Creating a scope-based allow-filter list using DHCP policies

To accomplish this, create a scope-based policy for the subnet with a condition of MAC Address Equals a1, a2, a3 where a1, a2, and a3 are the MAC addresses of known clients. See the following example.

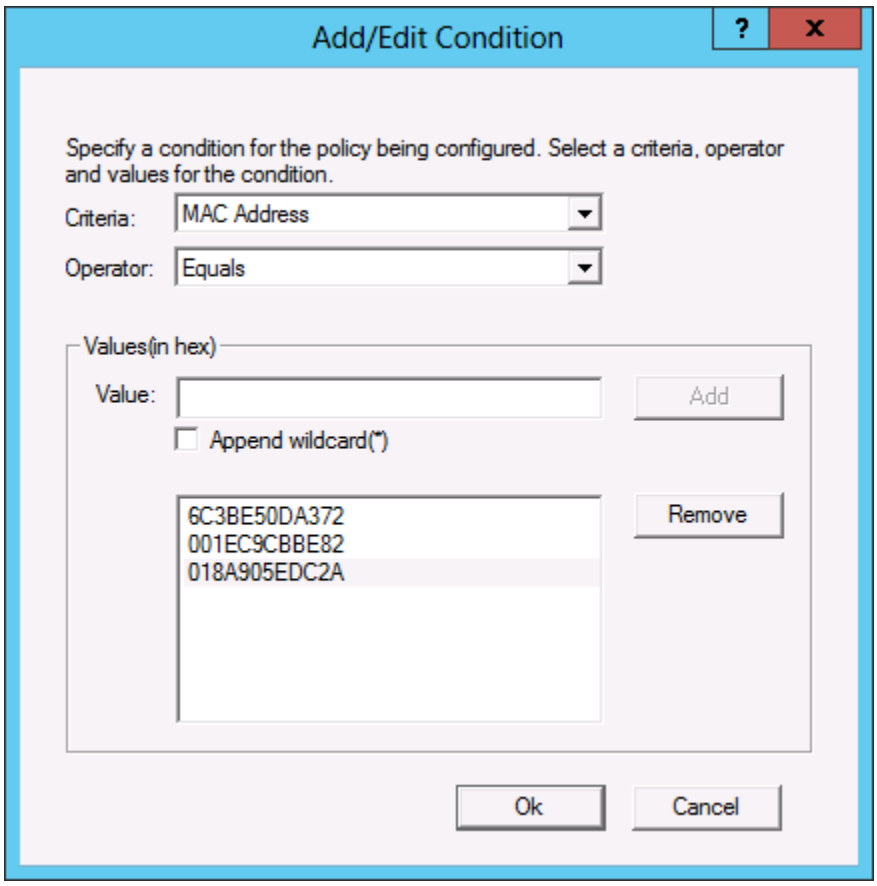

The conditions of the policy will be MAC address = a1, a2, a3 which has a logical meaning of MAC-address = a1 OR MAC-address = a2 OR MAC address = a3. When you are finished entering MAC addresses, click **OK** and you will see a single MAC address condition. You can edit this condition later if desired to add or remove MAC addresses from the list.

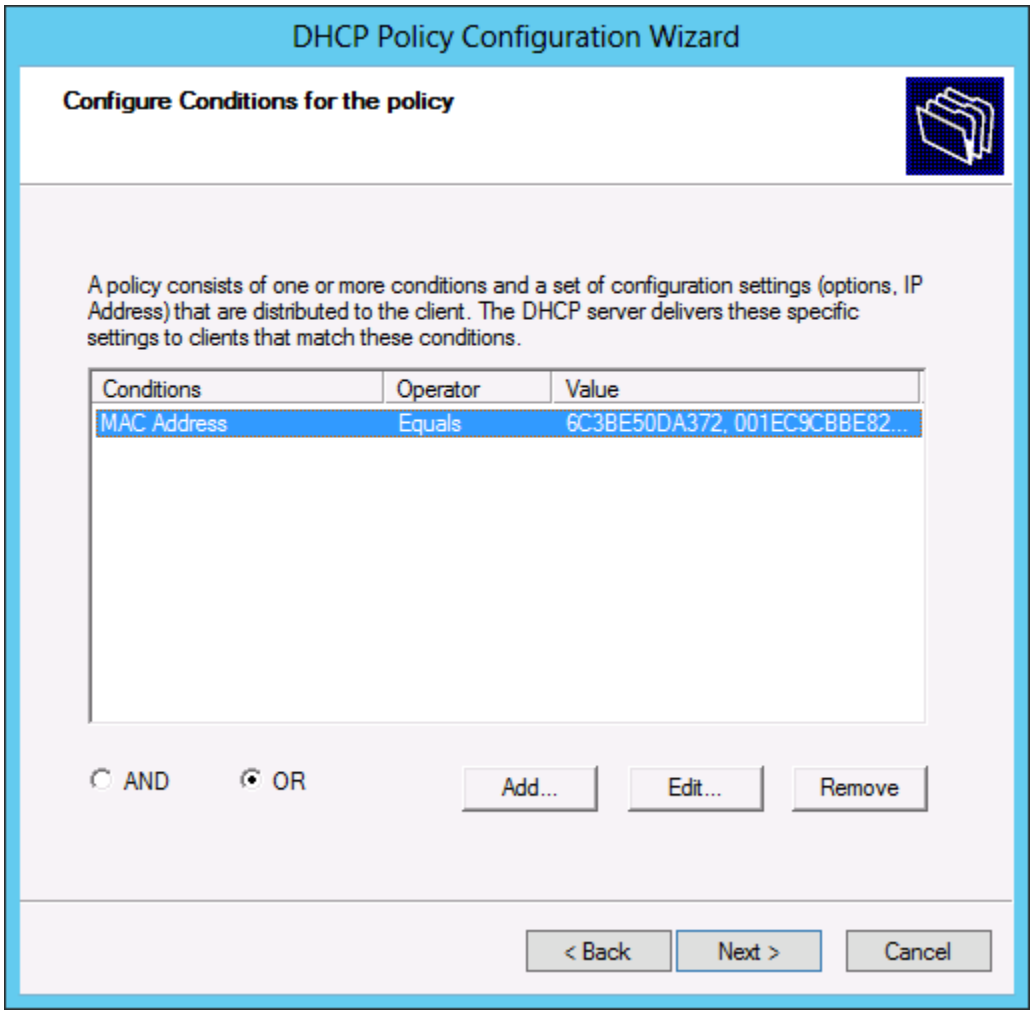

You can also add each MAC address as an individual condition, but you must make sure that the **OR** radio button is selected on the conditions page. Click **Next** to continue. You can also use wildcards if desired.

The IP address range for the policy will be the entire scope IP address range. The **Percentage of IP address range** value should be **100.0**. This also means that you cannot have other policies in this scope that specify an IP address range because overlapping IP address range settings are not allowed.

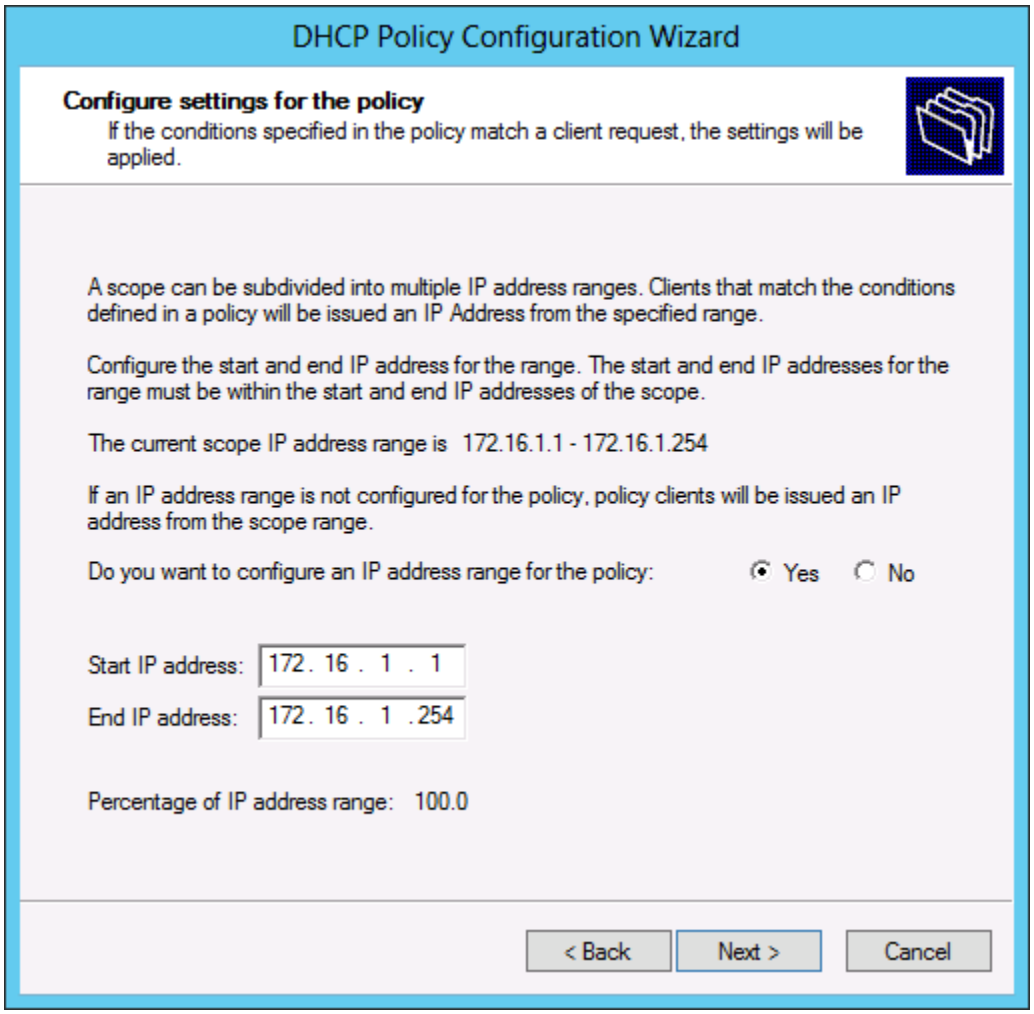

By choosing the entire scope range for this policy, no IP addresses will be available to clients that do not match the policy conditions. A client whose MAC address does not match any of the MAC addresses specified in the policy will be denied an IP address lease by the DHCP server.

Click **Next**, specify any additional DHCP options if desired, then click **Next** and click **Finish**.

Creating a scope-based deny-filter list using DHCP policies

If you wish to block specific MAC addresses from receiving an IP address from the given scope you can configure a deny list instead of an allow list. The only difference in creating a deny list vs. the allow list configured in the previous procedure is in the conditions of the policy. Instead of specifying a MAC address list with an operator of Equals, choose the Not Equals operator instead. See the following example.

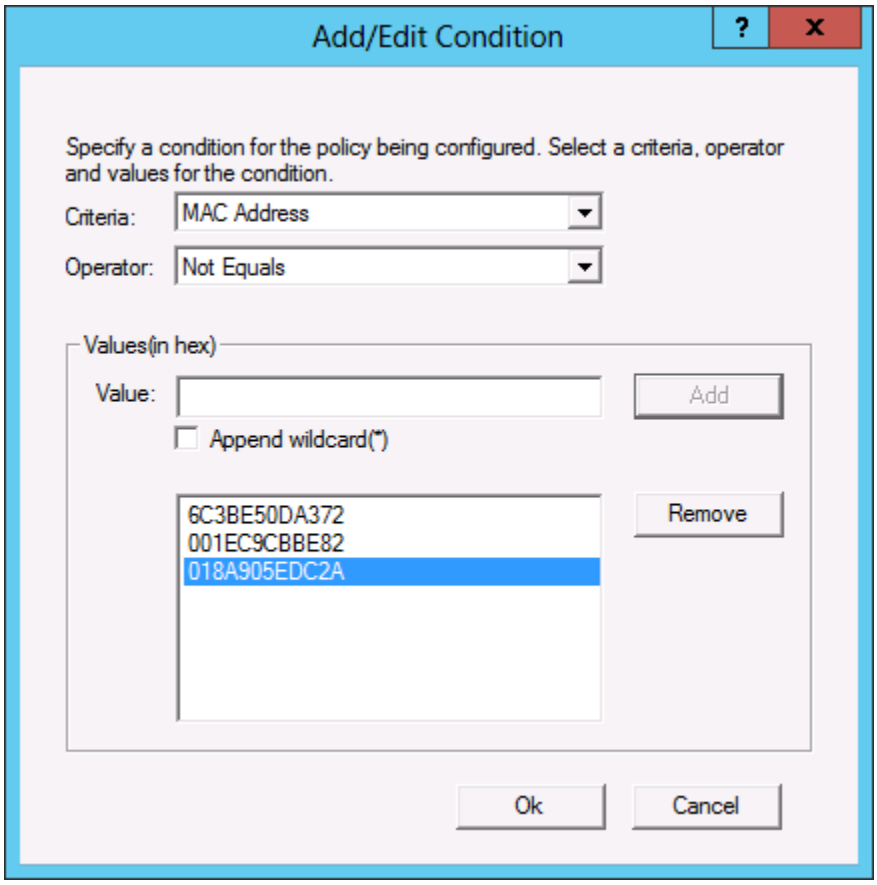

Important

In this case, if you specify individual MAC addresses in separate conditions, you must use the **AND** operator on the conditions page instead of the **OR** operator.

The “not equal to” MAC address conditions and the entire IP address range setting will ensure that incoming client DHCP requests that are sent from computers or devices that match one of the MAC addresses on the list will not receive a lease from the DHCP server.

Caution

Longs lists of conditions and values used in a “Not Equals” policy can negatively impact performance.

Windows PowerShell

If you have a large number of MAC addresses that need to be added to an allow or deny list in a policy, consider using a script to create your DHCP policy. See Scope based Link Layer filtering using DHCP Policies in the Microsoft Script Center for more information. The script takes a file as

input that contains the list of (line separated) MAC addresses, the scope ID, and whether an allow list or deny list should be created. This script will create the appropriate policy for you.

See also

Introduction to DHCP Policies

Scenario: Customize lease duration based on device type

See the following sections in this topic.

Problem description

Grouping devices on the network

Configure DHCP policies for different lease durations

Use Windows PowerShell to configure policies with different lease duration

Problem description

An organization has the following known device types on the network:

- Desktop computers
- Desktop phones
- Servers
- Printers
- Mobile devices

Servers, printers, and desktop phones are expected to be on the network for long periods of time. Mobile devices such as smartphones and tablets are expected to move in and out of the network frequently. Desktop computers are considered standard devices that will be on the network for 1 day to several days. The default lease duration for the DHCP scope is 8 days.

To optimize IP address utilization, the network administrator would like to specify the lease duration based on device type as follows:

Device Type	Lease Duration
Desktop computers	8 days
Desktop phones	30 days
Servers	60 days
Printers	60 days

Device Type	Lease Duration
Mobile devices	8 hours

Grouping devices on the network

Conditions you can use with DHCP policies to group devices on the network include:

1. **Vendor Class**: Most devices send this DHCP option with a DHCP lease request. The vendor class identifier is a string that is unique for the device type and the vendor. For example, Nortel phones might have a Vendor Class of Nortel-i2004-A.
2. **MAC address**: Every DHCP lease request contains the MAC address of the DHCP client, which is a unique 48 bit identifier for that device. The first three bytes of the MAC address contains the Organizationally Unique Identifier (OUI) which can be used to identify the device vendor. For example, devices manufactured by Apple will have a MAC address that begins with 00-1B-63, and HTC uses F8-DB-7F as one of its OUIs. Therefore, you can use the MAC address to identify device types.

By defining conditions of a DHCP policy to be based on Vendor Class identifier or MAC address prefix, you can group clients into different device types. Using the previous example, devices might be grouped using the following conditions in DHCP policies:

- If your organization uses Hewlett-Packard (HP) printers with JetDirect network adapter cards, you can define a new vendor class for **HP printers** using a vendor class definition of **Hewlett-Packard JetDirect**.
- Servers can be grouped based on MAC addresses of the network adapter of each server, and desk phones might be grouped using a Vendor Class such as **Nortel-i2004-A** or **Cisco Systems, Inc. IP Phone CP-7940G**.
- Mobile phones can typically be identified either by the Vendor Class identifier (Android phones tend to use a Vendor Class starting with **dhcpcd**) or by the OUI of the MAC address that specifies a vendor such as **Apple** or **HTC**.

Configure DHCP policies for different lease durations

In the following example, a server-based DHCP policy will be configured that enables different types of device to receive an IP address with a unique lease duration. For some devices, a new Vendor Class will be added.

Important
The value of the new Vendor Class must be the same as the value that is sent in the Vendor Class field by these devices in the DHCP packet when requesting for an address. Consult your vendor documentation to determine the appropriate vendor class.

Create a new Vendor Class

1. In the DHCP console tree, right-click **IPv4** and then click **Define Vendor Classes**.
2. Click **Add**, and then type the display name and description for the new vendor class under **Display name** and **Description**, respectively. For example, type **Nortel Phones** and **Desk phones**.
3. Under **ASCII**, type the vendor class identifier. For example, type **Nortel-i2004-A**. The corresponding **ID** and **Binary** values will be automatically added. When you are finished, click **OK**. See the following example.

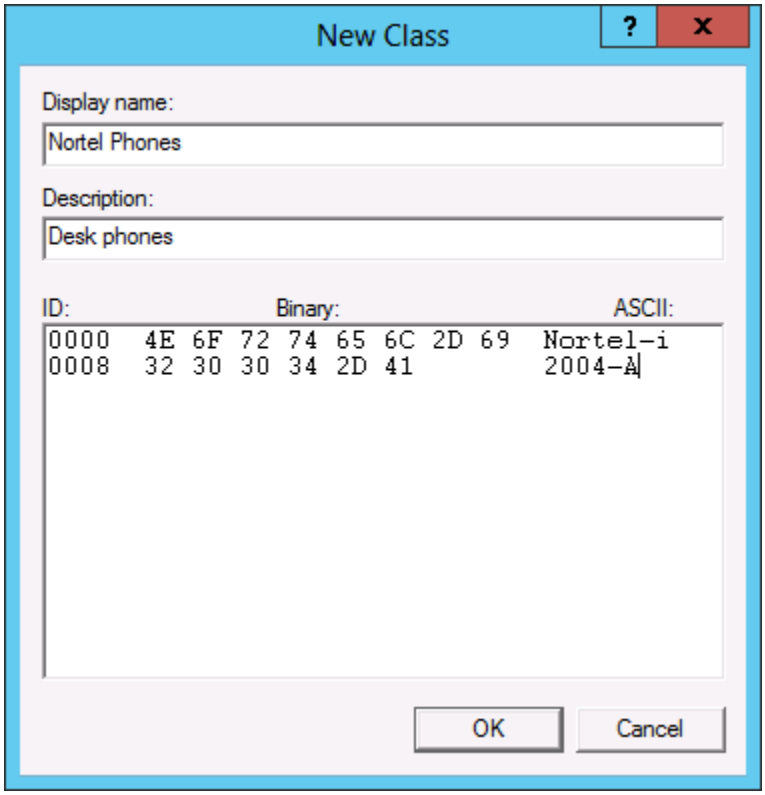

4. Verify that the new Vendor Class was created successfully and then click **Close**. See the following example.

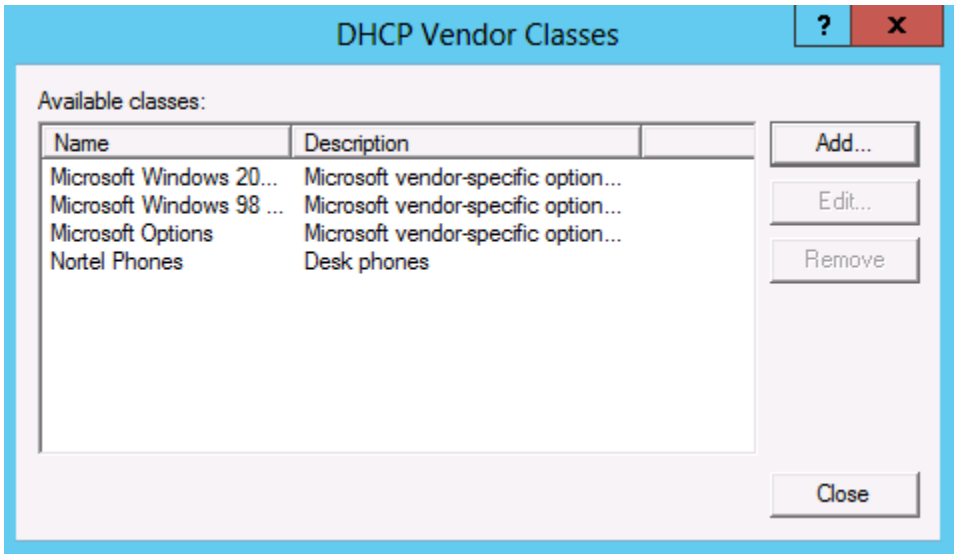

Important

The characters you type under ASCII must be an exact match. This includes using the same character casing and position and number of spaces.

Similarly, you might create a Vendor Class for HP printers with **HP Printer** as the Display name and **Hewlett-Packard JetDirect** in the ASCII data field, and so on for other types of devices.

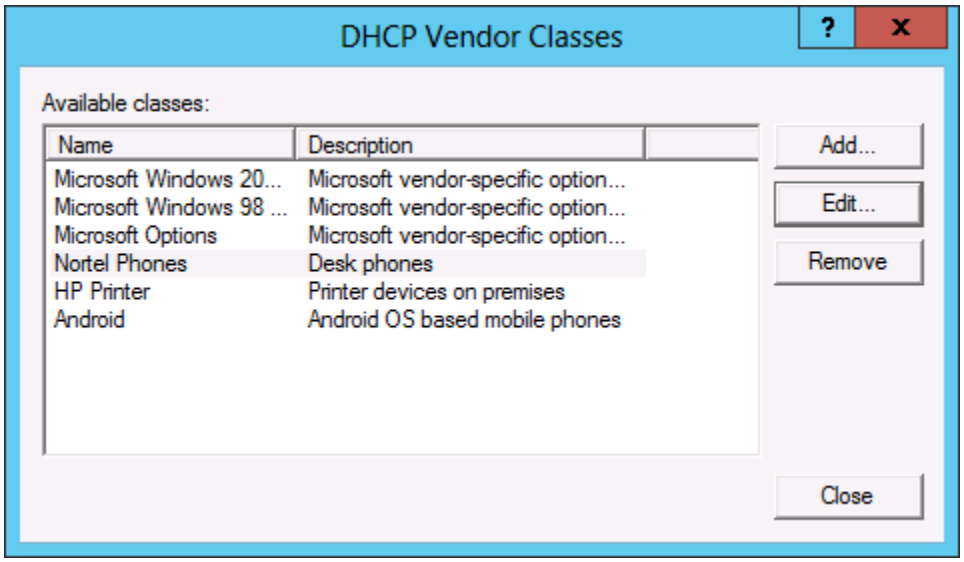

Add a vendor class for each device type that you wish to use in a DHCP policy and specify a custom configuration, such as lease duration. These vendor classes can be used for providing custom network configuration to different devices, such as a different lease duration.

Next, create different policies for each type of device. The following procedure demonstrates how to create one such policy at the server level. You can also specify policies at the scope level which can also be used to group different device types by IP address range. If you use scope level policies instead of server level policies, you must configure similar policies in each scope to group devices in that subnet. If you use server level policies, they will apply to all scopes, but any scope level policies that exist will have precedence. For more information about policy processing, see Policy processing.

Create a server-level Vendor Class based policy

1. In the DHCP console tree, under **IPv4**, right-click **Policies** and then click **New Policy**.
2. Next to **Policy Name**, type a name for the policy. For example, type **Printer Policy** and then click **Next**.
3. Click **Add**, choose **Vendor Class** next to **Criteria**, choose **Equals** next to **Operator**, choose **HP Printer** (or another value you added) next to **Value**, and then click **Add**. See the following example.

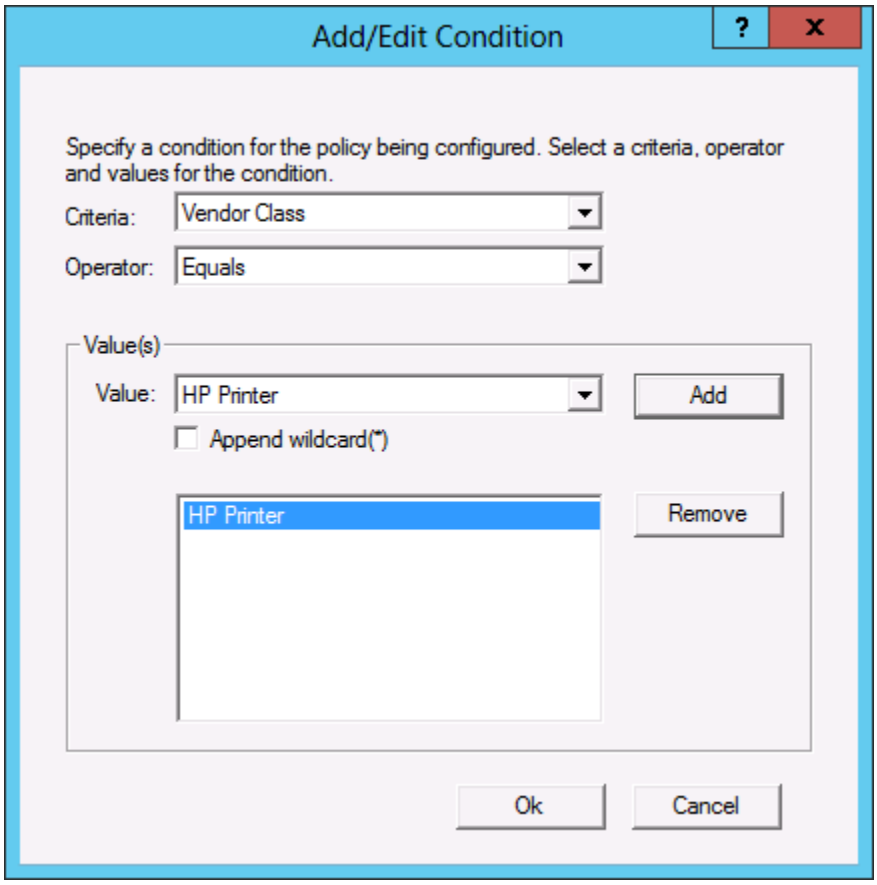

4. Click **OK**, click **Next** twice, and then click **Finish**.
5. Right-click **Printer Policy** and then click **Properties**.
6. On the **General** tab, select the **Set lease duration for the policy** checkbox and specify

a lease of **60** days. See the following example.

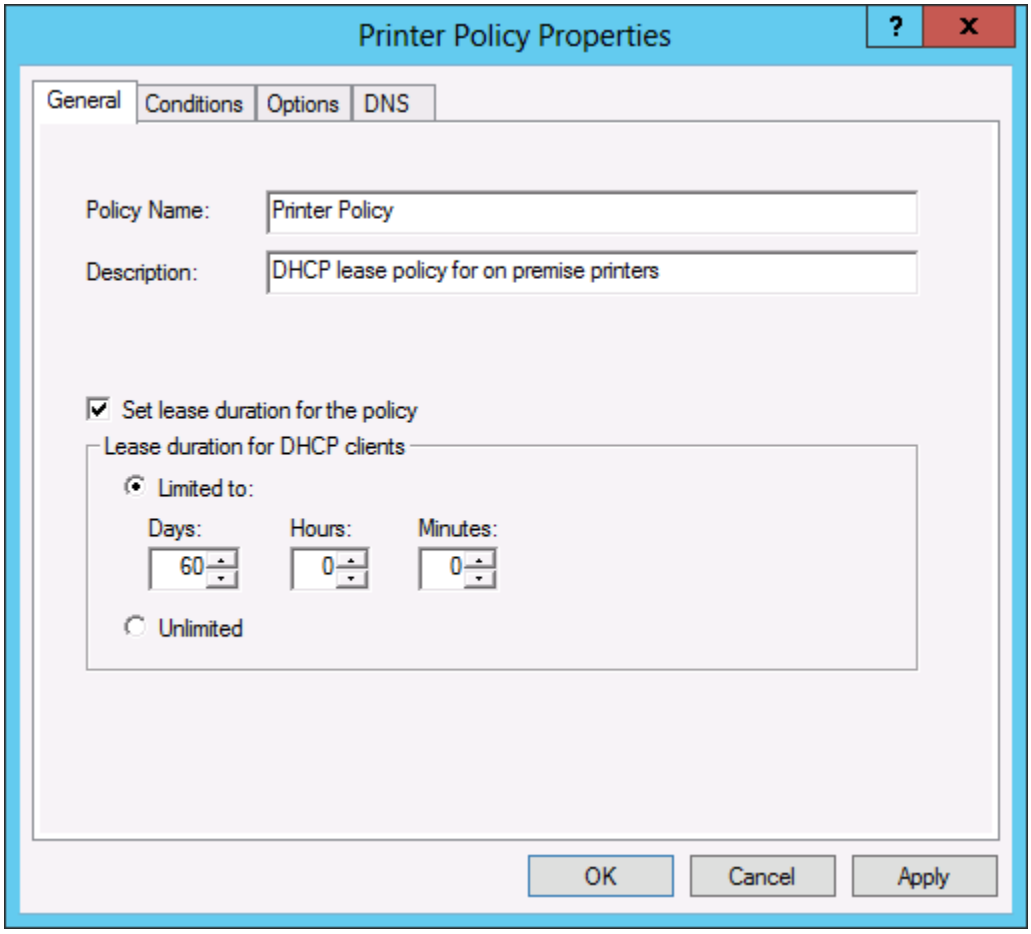

7. Click **OK** to complete configuration of the policy.

You can create policies for the other type of devices and set the policy lease duration as shown above. To create a policy for servers, you can create a MAC address based policy with all of the servers' MAC addresses in the condition of the policy. Configuring a MAC address based condition is discussed in Scenario: Secure a subnet to a specific set of clients.

Similarly, policies for mobile phones can have conditions that match one of the possible phone vendors. You can attempt to identify Android phones by their Vendor Class identifier and other vendors such as HTC and Apple by their MAC address prefix (OUI).

For each of these policies, you can configure the lease duration to the appropriate values.

Use Windows PowerShell to configure policies with different lease duration

You can also use Windows PowerShell to create DHCP policies with a vendor class or MAC address based condition. To customize lease duration in these policies, you must set option value 51 (lease duration) in each policy.

You can also specify new vendor classes with Windows PowerShell. For example, to add a new Vendor Class named **Printers** with ASCII data of **Hewlett-Packard JetDirect**, type the following at an elevated Windows PowerShell prompt and then press ENTER:

```
Add-DhcpServerv4Class -Name "HP Printer" -Type Vendor -Data "Hewlett-Packard JetDirect"
```

To create a policy for printers similar to that in the previous procedure, type the following at an elevated Windows PowerShell prompt and then press ENTER:

```
Add-DhcpServerv4Policy -Name "Printer Policy" -Description "Policy for printers" -
Condition "OR" -VendorClass "EQ", "HP Printer"
```

To create the same policy at the scope level, you must specify a scope ID. The following example specifies a scope with ID of 10.0.0.0:

```
Add-DhcpServerv4Policy -Name "Printers" -Description "Policy for printers" -ScopeId
10.0.0.0 -Condition "OR" -VendorClass "EQ", "HP Printer"
```

The following example shows how to create a MAC address based policy at the scope level with Windows PowerShell:

```
Add-DhcpServerv4Policy -Name "Servers" -Description "Policy for specific servers" -
ScopeId 10.0.0.0 -Condition "OR" -MACAddress "EQ", "0012ABCD34D2", "EQ", "0123AB43CD6E",
"EQ", "0A1234ABCDEF"
```

To add a lease duration setting to a policy, specify an **OptionId** of 51 with a lease duration in seconds. See the following example:

```
Set-DhcpServerv4OptionValue -ScopeId 10.0.0.0 -PolicyName "Printers" -OptionId 51 -Value
5184000
```

The value of 5184000 seconds can be calculated by multiplying 60 days * (24 hrs/day * 60 min/hr * 60 sec/min).

See also

Introduction to DHCP Policies

Step-by-Step: Configure DHCP Using Policy-based Assignment

The Dynamic Host Configuration Protocol (DHCP) server role in Windows Server 2012 introduces a new policy based IP address assignment feature. Policy based assignment (PBA) allows an administrator to group DHCP clients by specific attributes based on fields contained in the DHCP client request packet. This feature allows for targeted administration and greater control of configuration parameters delivered to network devices.

In this guide

This guide provides step-by-step instructions for deploying DHCP policy based assignment in a test lab using one server computer and two client computers. Software and hardware requirements are provided, as well as an overview of DHCP policy based assignment.

Important
The following instructions are for configuring a test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

DHCP policy based assignment overview

With a DHCP server running Windows Server 2012, administrators can define an address assignment policy at the server level or scope level. A policy contains a set of conditions to evaluate when processing client requests. Policy based assignment enables flexibility for some common scenarios, including:

1. **Multiple device types**: A network includes many different DHCP client devices, such as printers, IP phones, and desktops. Administrators need the ability to classify these devices using different IP address ranges. This enables router policies and quality of service (QoS) based on IP address range policies to control network access or traffic. For example, you can add a vendor class of “Hewlett-Packard JetDirect” or Cisco Systems, Inc. IP Phone CP-7940G and configure printer and IP-phone policies to assign a specific IP address range to these devices.
2. **Multiple roles**: A network includes different types of computers, such as laptops, desktops, and servers in the same subnet. Depending on the type of client, the administrator might wish to provide different lease duration settings. All the wireless clients that connect via a specific relay agent can be assigned a four-hour lease duration. DNS dynamic update protocol can be disabled for clients matching this policy. Similarly, a server policy can be created using a list of server MAC addresses. Servers can be assigned a 12-day lease duration

3. **Virtualization**: A data center network employs virtualization for different workloads and applications. Virtual machines are added and removed dynamically depending upon load requirements at a given time. An administrator wishing to route traffic on the network differently for VMs can create a policy based on MAC address prefix to assign a short lease duration, specific IP address range, and different default gateway.

The following fields in the DHCP client request are available when defining policies.

- Vendor Class
- User Class
- MAC address
- Client Identifier
- Relay Agent Information

Policy settings and evaluation

Policy settings can be of three types.

- **IP address range**: The IP address range within a scope from which to assign an IP address to a client. A server level policy cannot have a setting for an IP address range.
- **Standard DHCP options**: One or more standard DHCP options to send to a client in the response based on the options requested by the client in the parameter request list.
- **Vendor specific DHCP options**: One or more vendor specific DHCP options to send to the client based on the vendor class field in the client request.

The DHCP server evaluates policies sequentially according to an assigned processing order. The DHCP administrator assigns the processing order to the policies. If policies exist at the server and scope levels, the server applies both sets of policies and evaluates the scope policies before the server policies. The processing order for a scope level policy defines the order of evaluation within the scope. If there are no policies defined at the scope level, the policies at the server level apply to the scope.

Address assignment

The DHCP server determines the scope to which a DHCP client belongs based on the gateway IP address of the relay agent or the interface of the DHCP server on which it receives the DHCP client packet. Once the server determines the client scope, the server evaluates the DHCP packet against the policies applicable for the scope in the processing order specified. The policies applicable at a scope are those configured at the scope and those inherited from the server. A single client request can match multiple policies.

If a client request matches the conditions of a policy for which a specific IP address range is associated, the server will assign the first free IP address from the range as determined by the rule. If a policy is associated with multiple address ranges, the server will assign IP addresses by first attempting to assign an IP from the lowest address range. If no IP addresses are available to use from the lower address range, the server will then look for a free IP address from the higher address ranges. If no IP addresses are free from any of the address ranges associated with the policy, the server will process the next matched policy as defined by the processing order.

If none of the matched policies has a free IP address, the server will drop the client packet and log an event. If a DHCP client packet does not match any of the policies applicable for the scope, or none of the matched policies for a client packet is associated with an IP address range, the server will lease the client an IP address from the IP address range configured for the scope exclusive of any policy-specific IP address ranges.

Option assignment

A DHCP client uses the parameter request list field in a DHCP packet to request a list of standard options from the server. The option assignment processing for a client is similar to that of IP address assignment. The DHCP server evaluates the fields in the client request against each policy applicable for the scope in the processing order specified. If the client request matches the conditions of any of the policies applicable for the scope, and its settings include specific options, the server returns these options to the client. If multiple policies match the client request, the server returns the sum of the options specified for each of the matched policies. The DHCP server sends vendor class options to the client based on the vendor class contained in the DHCP client request.

Scenario overview

This test lab demonstrates new DHCP functionality in Windows Server 2012. One server computer and two client computers are used. See the following figure.

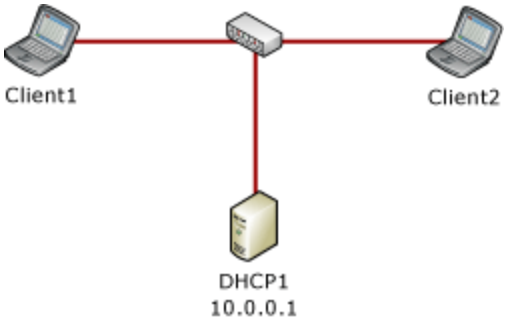

Hardware and software requirements

One server computer and two client computers are required to complete the test lab.

The following are required components of the test lab:

1. The product disc or other installation media for Windows Server 2012.
2. One computer that meets the minimum hardware requirements for Windows Server 2012.
3. At least one DHCP client computer is required.

Note

The lab uses two computers running Windows® 8. If only one client computer is available, or clients are not running Windows 8, you must alter some of the procedures in the test lab accordingly.

Steps for configuring the test lab

The following procedures are used to configure computers for the demonstration portion of the test lab:

1. Configure DHCP1: DHCP1 is a domain controller, DNS server, and DHCP server for the contoso.com Active Directory domain.
2. Configure Client1: Client1 is a DHCP client computer.
3. Configure Client2: Client2 is a DHCP client computer.

Configure DHCP1

DHCP1 is a computer running Windows Server 2012, providing the following services:

- A domain controller for the contoso.com Active Directory domain.
- An authoritative DNS server for the contoso.com DNS zone.
- A DHCP server.

Initial configuration of DHCP1 consists of the following steps:

- Install the operating system and configure TCP/IP on DHCP1
- Install AD DS, DNS Server, and DHCP Server
- Create a domain administrator account
- Create a DHCP scope on DHCP1

Install the operating system and configure TCP/IP on DHCP1

To install the operating system and configure TCP/IP on DHCP1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. Click **Start**, type **ncpa.cpl**, and then press **ENTER**. The **Network Connections** control panel will open.

   Tip
   The previous step demonstrates new functionality in Windows Server 2012 that enables you to search and run applications, settings, and files by clicking **Start** and then typing a search term. You can also open the **Network Connections** control panel by clicking next to **Wired Ethernet Connection** in Server Manager

using the **Local Server** view. For more information, see Common Management Tasks and Navigation in Windows Server 2012 (http://go.microsoft.com/fwlink/p/?LinkId=242147).

6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.1** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install AD DS, DNS Server, and DHCP Server

DHCP1 will serve as a domain controller, DNS server, and DHCP server for the contoso.com Active Directory domain.

To configure DHCP1 as a domain controller, DNS server, and DHCP server

1. The Server Manager Dashboard is displayed by default. In the navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DHCP1** under **Computer name**.
4. Click **OK** twice, and then click **Close**.
5. When you are prompted to restart the computer, click **Restart Now**.
6. After restarting the computer, sign-in using the local Administrator account.
7. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
8. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **Active Directory Domain Services** checkbox.
9. When you are prompted to add required features, click **Add Features**.
10. Select the **DHCP Server** checkbox.
11. When you are prompted to add required features, click **Add Features**.
12. Select the **DNS Server** checkbox.
13. When you are prompted to add required features, click **Add Features**.
14. Click **Next** five times, and then click **Install**.
15. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DHCP1** is displayed, and then click **Close**.
16. Click the Notification flag and then click **Promote this server to a domain controller**.

See the following example.

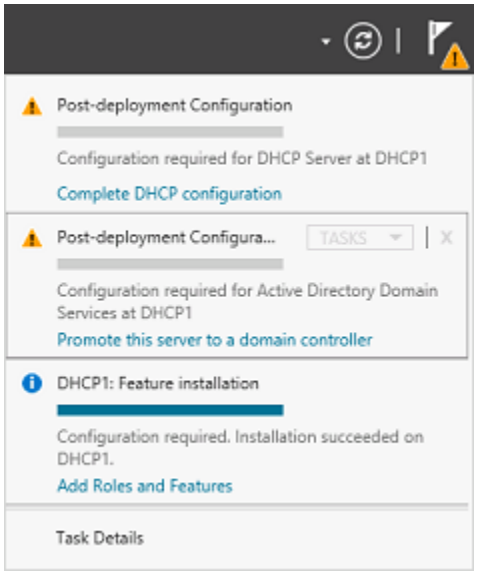

17. In the **Active Directory Domain Services Configuration Wizard**, on the **Deployment Configuration** page, choose **Add a new forest** and then next to **Root domain name**, type **contoso.com**.
18. Click **Next**, and then on the **Domain Controller Options** page, under **Type the Directory Services Restore Mode (DSRM) password**, type a password next to **Password** and **Confirm password**. Confirm that **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected, and then click **Next**.
19. Click **Next** four times, verify that **All prerequisite checks passed successfully** is displayed, and then click **Install**.
20. The computer will restart automatically to complete the installation process.
21. Sign in using the local Administrator account.

Create a domain administrator account

A domain administrator account is required to configure settings in the test lab.

Tip

You can use the CONTOSO\Administrator account in this test lab and skip creation of a domain administrator account if desired. This account has domain administrator privileges, and other privileges. However, it is a best practice to disable or rename this account. For more information, see Active Directory Best Practices(http://go.microsoft.com/fwlink/p/?LinkID=243071).

To create a domain administrator account

1. On the Server Manager menu bar, click **Tools**, and then click **Active Directory Users and Computers**.

2. In the **Active Directory Users and Computers** console tree, double-click **contoso.com**, right-click **Users**, point to **New**, and then click **User**.
3. In the **New Object – User** dialog box, type **user1** under **User logon name** and next to **Full name**, then click **Next**.
4. Next to **Password** and **Confirm password**, type a password for the user1 account.
5. Clear the checkbox next to **User must change password at next logon**, select the **Password never expires** checkbox, click **Next**, and then click **Finish**.
6. Double-click **user1** and then click the **Member Of** tab.
7. Click **Add**, type **domain admins** under **Enter the object names to select**, click **OK** twice, and then close the **Active Directory Users and Computers** console.
8. Click **Start**, click **Administrator**, and then click **Sign out**.
9. Sign in to the computer using the user1 credentials by clicking the left arrow next to **CONTOSO\Administrator** and then clicking **Other user**.

Create a DHCP scope on DHCP1

Next, create a DHCP scope on DHCP1.

To create a DHCP scope on DHCP1

1. On the Server Manager menu bar, click **Tools** and then click **DHCP**. THE DHCP console opens.
2. In the DHCP console tree, navigate to **IPv4**. Right-click **IPv4** and then click **New Scope**. The **New Scope Wizard** opens.
3. Click **Next** and then type a name for the new scope next to **Name** (ex: Contoso-scope1).
4. Click **Next** and then in **IP Address Range**, type **10.0.0.1** next to **Start IP address**, type **10.0.0.254** next to **End IP address**, and type **24** next to **Length**. The value of Subnet mask will change automatically to **255.255.255.0**.
5. Click **Next**, and then in **Add Exclusions and Delay** type **10.0.0.1** under **Start IP address**, type **10.0.0.10** under **End IP address**, and then click **Add**. This allows the first ten IP addresses in the 10.0.0.0/24 subnet to be used for static addressing of servers on the network.
6. Click **Next** and then in **Lease Duration** under **Limited to** enter **0 Days**, **0 Hours**, and **2 Minutes**. This very short lease duration will simplify the DHCP demonstration.
7. Click **Next** three times, and then in **Domain Name and DNS Servers**, verify that the **Parent domain** is **contoso.com** and **10.0.0.1** is listed as the only DNS server.
8. Click **Next** twice, and then in **Activate Scope** select **Yes, I want to activate this scope now**.
9. Click **Next**, and then click **Finish**.
10. In the DHCP console tree, right-click **dhcp1.contoso.com**, and then click **Authorize**.
11. Refresh the view in the DHCP console and verify that DHCP1 is authorized and that the Contoso-scope1 is active.

    Note: To review scopes on the current server using Windows PowerShell, right-click

Windows PowerShell, click **Run as Administrator**, click **Yes** in the User Account Control alert that appears, and then type the following command at the Windows PowerShell prompt, and then press ENTER.

```
get-dhcpserverv4scope
```

```
PS C:\Windows\system32> get-dhcpserverv4scope


ScopeId         SubnetMask      Name            State    StartRange      EndRange
LeaseDuration
-------         ----------      ----            -----    ----------      --------
-------------
10.0.0.0        255.255.255.0   Contoso-scope1 Active   10.0.0.1        10.0.0.254
00:02:00
```

Configure Client1

Client1 is a computer running Windows® 8 that is acting as a DHCP client.

Configuration of Client1 consists of the following steps:

- Install the operating system on Client1
- Pin Windows PowerShell to the taskbar

During the demonstration portion of the test lab, Client1 will be used as a DHCP client.

Install the operating system on Client1

To install the operating system on Client1

1. Start your computer using the Windows 8 product disc or other digital media.
2. When prompted, enter a product key and accept license terms.
3. When prompted to enter a computer name, type **Client1** and click **Next**.
4. Click **Use express settings**.
5. On the **Sign in to your PC** page, click **Don't want to sign in with a Microsoft account** and then click **Local account**.
6. Next to **User name**, type **user1**, enter a password and password hint, and then click **Finish**.

Pin Windows PowerShell to the taskbar

The DHCP failover demonstration on Client1 makes use of Windows PowerShell to verify DHCP lease information. To make Windows PowerShell more easily accessible, it will be pinned to the taskbar.

To pin Windows PowerShell to the taskbar

1. The **Start** menu is displayed by default. If **Start** is not displayed, move the mouse cursor to the lower left corner of the screen until **Start** is displayed, and then click **Start**.
2. Type **powershell** and then under **Results for "powershell"** right-click **Windows PowerShell** and then click **Pin to taskbar**.Confirm that Windows PowerShell is pinned to the taskbar.

Note

Client1 can also be joined to the contoso.com domain; however this is not required to complete the test lab.

Configure Client2

Client2 is a computer running Windows 8 that is acting as a DHCP client.

Configuration of Client2 is nearly identical to Client1. To configure Client2, repeat the identical procedures used to configure Client1, except when you enter a computer name, type Client2 instead of Client1.

DHCP policy based assignment demonstration

For the DHCP policy based assignment demonstration portion of the test lab, a MAC address based policy will be configured to assign unique DHCP options and IP address ranges to Client1 and Client2.

A demonstration of DHCP policy based assignment on Windows Server 2012 consists of the following procedures:

1. Determine MAC addresses
2. Create policies
3. Demonstrate policies

Determine MAC addresses

Unique MAC addresses on Client1 and Client2 will be used to match different DHCP policies.

To determine MAC addresses

1. On DHCP1, open the DHCP console and navigate to **IPv4\Scope [10.0.0.0] Contoso-scope1\Address Leases**.
2. Click **Address Leases** and then in the details pane write down the MAC addresses listed for Client1 and Client2 under **Unique ID**.

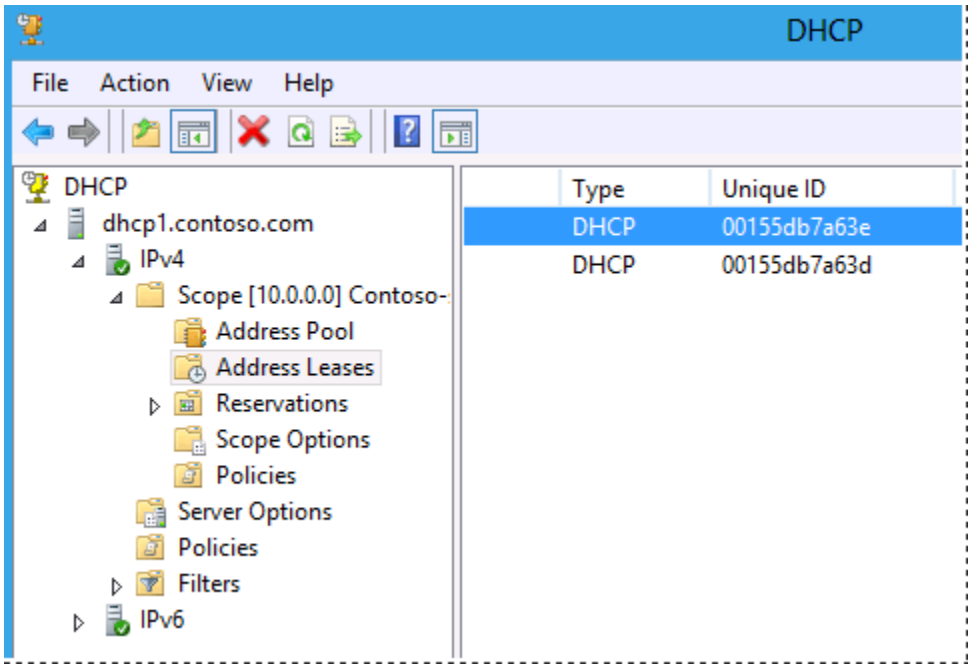

Create policies

The DHCP Policy Configuration Wizard will be used to create a unique policy for Client1 and another policy for Client2. A policy configured for an individual computer is not typical and is only configured for demonstration purposes. On a corporate network, you can us wildcards and other conditions to match multiple DHCP client devices.

To create policies

1. Right-click **Policies** and then click **New Policy**.
2. Next to **Policy Name**, type **Client1 Policy**, and then click **Next**.
3. On the **Configure Conditions for the policy** page, click **Add**.
4. In the **Add/Edit Condition** dialog box, choose **MAC Address** next to **Criteria**, type the MAC address for Client1 next to **Value** (001DB7A63D in this example), and then click **OK**.

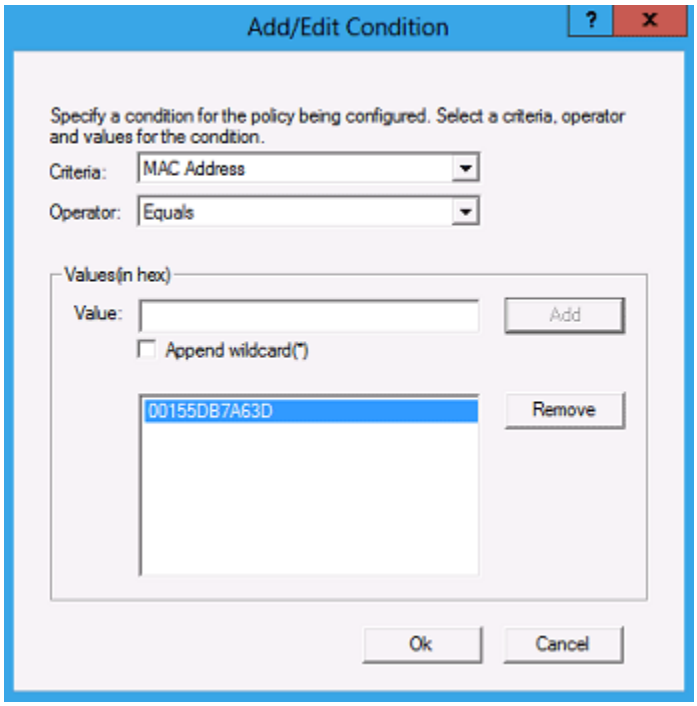

5. Click **Next**, and then in **Configure settings for the policy**, type **10.0.0.100** next to **Start IP address** and type **10.0.0.199** next to **End IP address**.

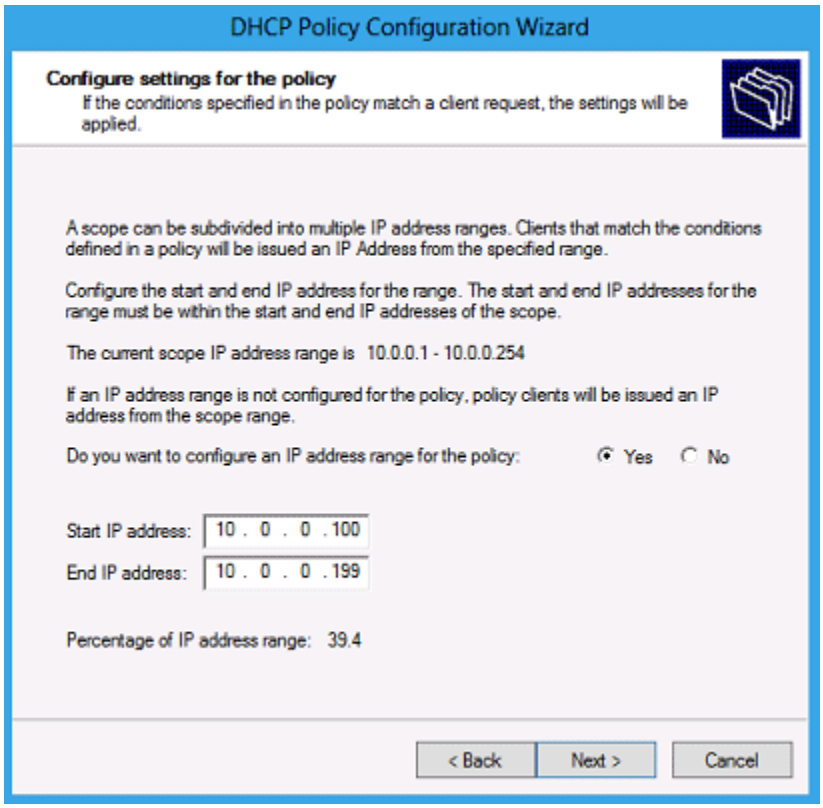

6. Click **Next**, and then under **Available Options**, select **003 Router**, type **10.0.0.7** under **IP address**, and click **Add**.

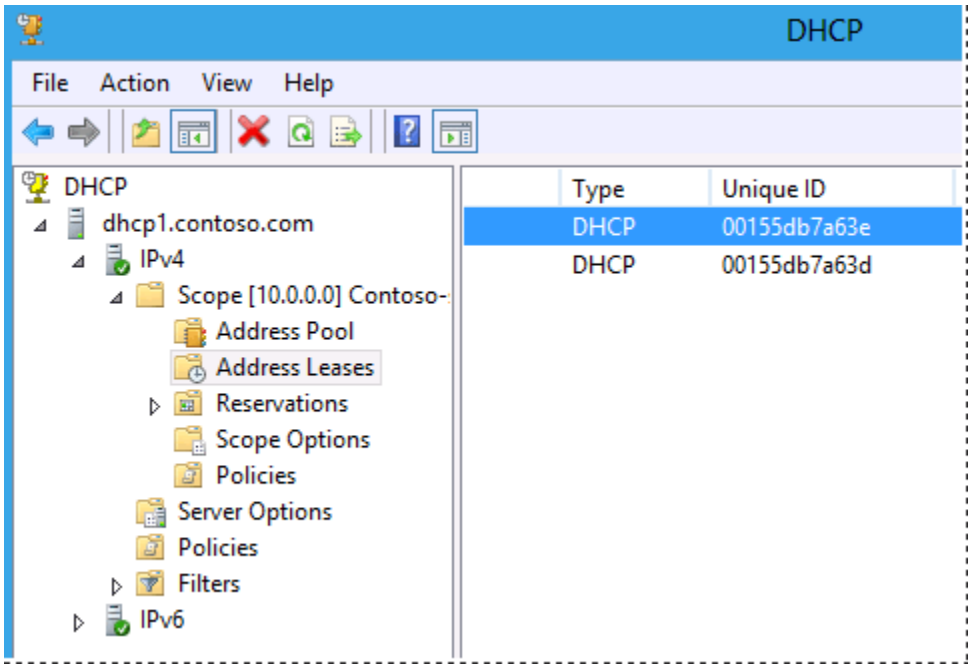

7. Click **Next**, and then click **Finish**.
8. Repeat the previous steps for Client2 using the following conditions, IP address ranges, and options:
   a. **Policy Name**: Client2 Policy
   b. **Condition**: MAC Address equals (in this example) 00155DB7A63E.
   c. **Start IP address**: 10.0.0.200
   d. **End IP address**: 10.0.0.254
   e. **003 Router**: 10.0.0.8

Demonstrate policies

Next, review the effect that these policies have on the IP address configuration of Client1 and Client2.

To demonstrate policies

1. In the details pane, under **Policy Name**, right-click one of the two policies you just created.
2. Note that you can move the policy up or down in the processing order, delete the policy, or disable the policy.
3. Click **Properties**.
4. Review the available parameters that you can edit on the **General**, **Conditions**, **IP**

Address Range, **Options**, and **DNS** tabs.

5. Click **OK**, and then on Client1 type **ipconfig /all** at the Windows PowerShell prompt.

```
PS C:\Users\user1> ipconfig /all

Windows IP Configuration

   Host Name . . . . . . . . . . . . : Client1
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
   DNS Suffix Search List. . . . . . : contoso.com

Ethernet adapter Wired Ethernet Connection:

   Connection-specific DNS Suffix  . : contoso.com
   Description . . . . . . . . . . . : Microsoft Hyper-V Network Adapter
   Physical Address. . . . . . . . . : 00-15-5D-B7-A6-3D
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::1d47:817b:70fb:6f42%12(Preferred)
   IPv4 Address. . . . . . . . . . . : 10.0.0.100(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Monday, February 13, 2012 7:11:35 PM
   Lease Expires . . . . . . . . . . : Monday, February 13, 2012 7:16:34 PM
   Default Gateway . . . . . . . . . : 10.0.0.7
   DHCP Server . . . . . . . . . . . : 10.0.0.1
   DHCPv6 IAID . . . . . . . . . . . : 251663709
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-16-C7-3F-98-00-15-5D-B7-A6-3D
   DNS Servers . . . . . . . . . . . : 10.0.0.1
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

6. Client1 has been assigned the first IP address in the 10.0.0.100 – 10.0.0.199 range, and a default gateway of 10.0.0.7, as determined by the policy based assignment.

   Repeat the previous step on Client2.

```
PS C:\Users\user1> ipconfig /all

Windows IP Configuration

   Host Name . . . . . . . . . . . . : Client2
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
   DNS Suffix Search List. . . . . . : contoso.com

Ethernet adapter Wired Ethernet Connection:

   Connection-specific DNS Suffix  . : contoso.com
   Description . . . . . . . . . . . : Microsoft Hyper-V Network Adapter
   Physical Address. . . . . . . . . : 00-15-5D-B7-A6-3E
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::c51f:bb33:b79c:8834%12(Preferred)
   IPv4 Address. . . . . . . . . . . : 10.0.0.200(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Monday, February 13, 2012 7:12:59 PM
   Lease Expires . . . . . . . . . . : Monday, February 13, 2012 7:14:59 PM
   Default Gateway . . . . . . . . . : 10.0.0.8
   DHCP Server . . . . . . . . . . . : 10.0.0.1
   DHCPv6 IAID . . . . . . . . . . . : 251663709
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-16-CB-55-2D-00-15-5D-B7-A6-3E
   DNS Servers . . . . . . . . . . . : 10.0.0.1
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

7. Client2 has been assigned the first IP address in the 10.0.0.200 – 10.0.0.254 range, and a default gateway of 10.0.0.8.

Policy based assignment in Windows Server 2012 DHCP allows you to create simple yet powerful rules to administer DHCP on your network.

Step-by-Step: Configure DHCP for Failover

Dynamic Host Configuration Protocol (DHCP) failover in Windows Server 2012 is a new method for ensuring continuous availability of DHCP service to clients.

In this guide

This guide provides step-by-step instructions for deploying DHCP failover in a test lab using two server computers and one client computer. Software and hardware requirements are provided, as well as an overview of DHCP failover.

Important

The following instructions are for configuring a test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

DHCP failover overview

In Windows Server® 2008 R2, there are two high availability options available for DHCP Server deployment. Each of these options is associated with some challenges.

1. **DHCP in a Windows failover cluster**. This option places the DHCP server in a cluster with an additional server configured with the DHCP service that assumes the load if the primary DHCP server fails. The clustering deployment option uses a single shared storage. This makes the storage a single point of failure, and requires additional investment in redundancy for storage. In addition, clustering involves relatively complex setup and maintenance.
2. **Split scope DHCP**. Split scope DHCP uses two independent DHCP servers that share responsibility for a scope. Typically 70% of the addresses in the scope are assigned to the primary server and the remaining 30% are assigned to the backup server. If clients cannot reach the primary server then they can get an IP configuration from the secondary server. Split scope deployment does not provide IP address continuity and is unusable in scenarios where the scope is already running at high utilization of address space, which is very common with Internet Protocol version 4 (IPv4).

DHCP failover in Windows Server 2012 enables administrators to deploy a highly resilient DHCP service to support a large enterprise without the challenges of the options discussed earlier. The main goals of the feature are the following:

- Provide DHCP service availability at all times on the enterprise network.
- If a DHCP server is no longer reachable, the DHCP client is able to extend the lease on its current IP address by contacting another DHCP server on the enterprise network.

The DHCP server failover feature provides the ability to have two DHCP servers provide IP addresses and option configuration to the same subnet or scope, providing for continuous availability of DHCP service to clients. The two DHCP servers replicate lease information between them, allowing one server to assume responsibility for servicing of clients for the entire subnet when the other server is unavailable. It is also possible to configure failover in a load-balancing configuration with client requests distributed between the two servers in a failover relationship.

DHCP failover in Windows Server 2012 provides support for a maximum of two DHCP servers, and the failover relationship is limited to IPv4 scopes and subnets. Network nodes using Internet Protocol version 6 (IPv6) typically determine their own IPv6 address using stateless IP auto configuration. In this mode, the DHCP server delivers only the DHCP option configuration, and the server does not maintain any lease state information. A high availability deployment for stateless DHCPv6 is possible by simply setting up two servers with identical option configuration. Even in a stateful DHCPv6 deployment, the scopes do not run under high address utilization, which makes split scope a viable solution for high availability.

DHCP failover architecture

Administrators can deploy DHCP servers running Windows Server 2012 as failover partners in either hot standby mode or load sharing mode.

Hot standby mode

In hot standby mode, two servers operate in a failover relationship where an active server is responsible for leasing IP addresses and configuration information to all clients in a scope or subnet. The secondary server assumes this responsibility if the primary server becomes unavailable. A server is primary or secondary in the context of a subnet. For instance, a server that has the role of a primary for a given subnet could be a secondary server for another subnet.

Hot standby mode of operation is best suited to deployments where a central office or data center server acts as a standby backup server to a server at a remote site, which is local to the DHCP clients (ex: hub and spoke deployment). In such deployments, it is undesirable to have a remote standby server service any clients unless the local DHCP server becomes unavailable. The figure below is an example of a hub and spoke deployment.

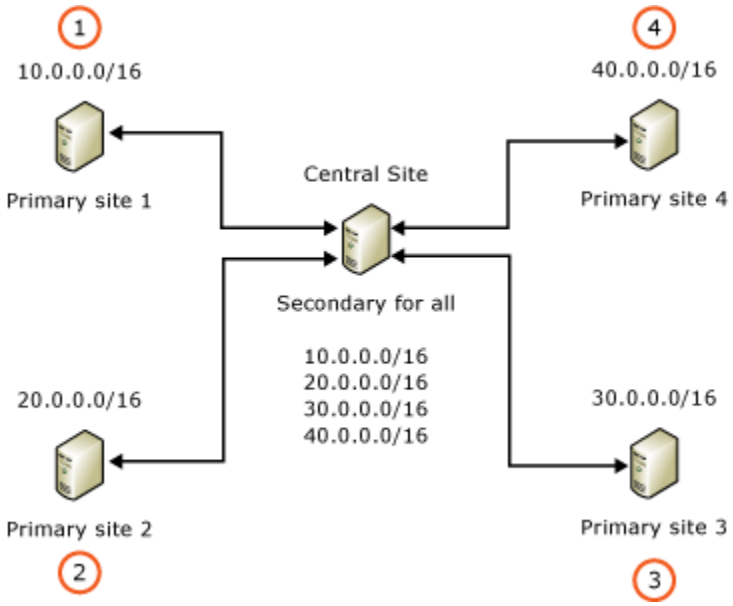

Load sharing mode

In a load sharing mode deployment, which is the default mode of operation, the two servers simultaneously serve IP addresses and options to clients on a given subnet. The client requests are load balanced and shared between the two servers.

The load sharing mode of operation is best suited to deployments where both servers in a failover relationship are located at the same physical site. Both servers respond to DHCP client requests based on the load distribution ratio configured by the administrator. See the following examples.

Load sharing in a single site with a single subnet

In the following example, two DHCP servers simultaneously provide IP addressing to clients on a single subnet.

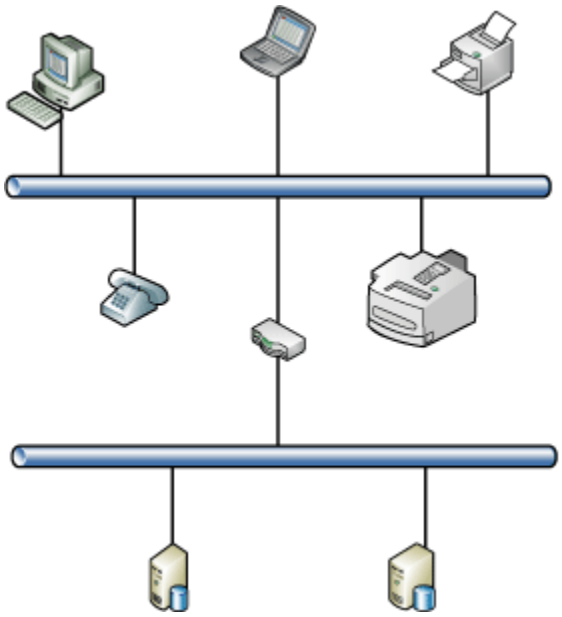

Legend

DHCP Failover — Single Subnet	
Symbol	Description
	Ethernet
	DHCP Server
	DHCP Relay
	Printer
	PC
	Laptop Computer
	IP Phone
	Multi-function device

Load sharing in a single site with multiple subnets

In the following example, two DHCP servers simultaneously provide IP addressing to clients on multiple subnets.

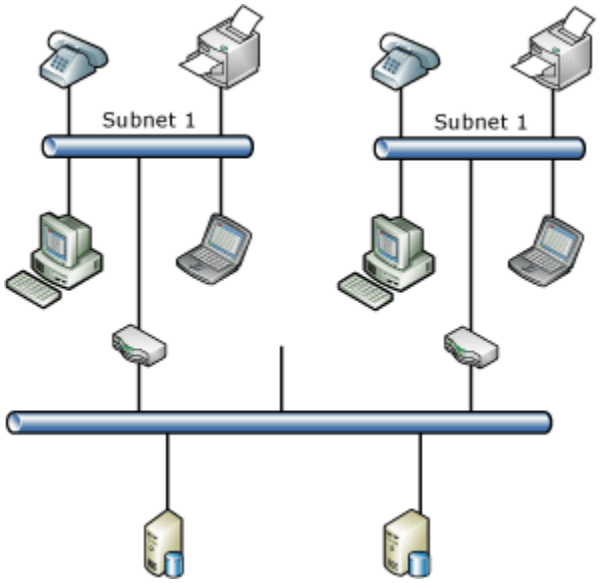

Legend

DHCP Failover — Single site, Multiple Subnets	
Symbol	Description
	Ethernet
	DHCP Server
	DHCP Relay
	Printer
	PC
	Laptop Computer
	Telephone

Scenario overview

This test lab demonstrates new DHCP functionality in Windows Server 2012. Two server computers and one client computer are used. See the following figure.

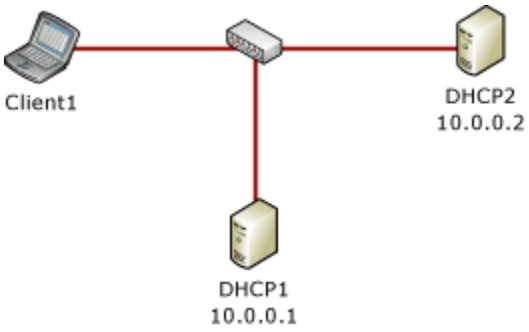

Hardware and software requirements

Two server computers and one client computer are required to complete the test lab.

The following are required components of the test lab:

1. The product disc or other installation media for Windows Server 2012.
2. Two computers that meet the minimum hardware requirements for Windows Server 2012.
3. The product disc or other installation media for Windows® 8.

Number of servers: DHCP failover is not supported for more than two DHCP servers. The failover relationship is always comprised of two DHCP servers.

Domain membership: In this guide, DHCP servers are domain member servers. You can also configure DHCP failover on workgroup computers.

Time synchronization: For DHCP failover to function correctly, time must be kept synchronized between the two servers in a failover relationship. Time synchronization can be maintained by deployment of the Network Time Protocol (NTP) or any alternate mechanism. When the failover configuration wizard is run, it will compare the current time on the servers being configured for failover. If the time difference between the servers is greater than one minute, the failover setup process will halt with a critical error instructing the administrator to synchronize the time on the servers.

Steps for configuring the test lab

The following procedures are used to configure computers for the demonstration portion of the test lab:

1. Configure DHCP1: DHCP1 is a domain controller, DNS server, and DHCP server for the contoso.com Active Directory domain.

   Important
   Domain controller and DNS server roles are not required for DHCP failover. These roles are installed on DHCP1 to mimic a domain environment. DHCP failover can also be configured on a workgroup computer (not demonstrated in this test lab).
2. Configure DHCP2: DHCP2 is a DHCP server and domain member computer.
3. Configure Client1: Client1 is a DHCP client computer.

Configure DHCP1

DHCP1 is a computer running Windows Server 2012, providing the following services:

- A domain controller for the contoso.com Active Directory domain.
- An authoritative DNS server for the contoso.com DNS zone.
- A DHCP server.

Initial configuration of DHCP1 consists of the following steps:

- Install the operating system and configure TCP/IP on DHCP1
- Install AD DS, DNS Server, and DHCP Server
- Create a domain administrator account
- Authorize DHCP1 in Active Directory

Note

It is not necessary to configure a DHCP scope on DHCP1. A DHCP scope will be configured automatically when a failover relationship is created with DHCP2.

Install the operating system and configure TCP/IP on DHCP1

To install the operating system and configure TCP/IP on DHCP1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. Click **Start**, type **ncpa.cpl**, and then press **ENTER**. The **Network Connections** control panel will open.

   Tip

   The previous step demonstrates new functionality in Windows Server 2012 that enables you to search and run applications, settings, and files by clicking **Start** and then typing a search term. You can also open the **Network Connections** control panel by clicking next to **Wired Ethernet Connection** in Server Manager using the **Local Server** view. For more information, see Common Management Tasks and Navigation in Windows Server 2012 (http://go.microsoft.com/fwlink/p/?LinkId=242147).
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.1** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.

10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install AD DS, DNS Server, and DHCP Server

DHCP1 will serve as a domain controller, DNS server, and DHCP server for the contoso.com Active Directory domain.

To configure DHCP1 as a domain controller, DNS server, and DHCP server

1. The Server Manager Dashboard is displayed by default. In the navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DHCP1** under **Computer name**.
4. Click **OK** twice, and then click **Close**.
5. When you are prompted to restart the computer, click **Restart Now**.
6. After restarting the computer, sign-in using the local Administrator account.
7. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
8. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **Active Directory Domain Services** checkbox.
9. When you are prompted to add required features, click **Add Features**.
10. Select the **DHCP Server** checkbox.
11. When you are prompted to add required features, click **Add Features**.
12. Select the **DNS Server** checkbox.
13. When you are prompted to add required features, click **Add Features**.
14. Click **Next** five times, and then click **Install**.
15. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DHCP1** is displayed, and then click **Close**.
16. Click the Notification flag and then click **Promote this server to a domain controller**. See the following example.

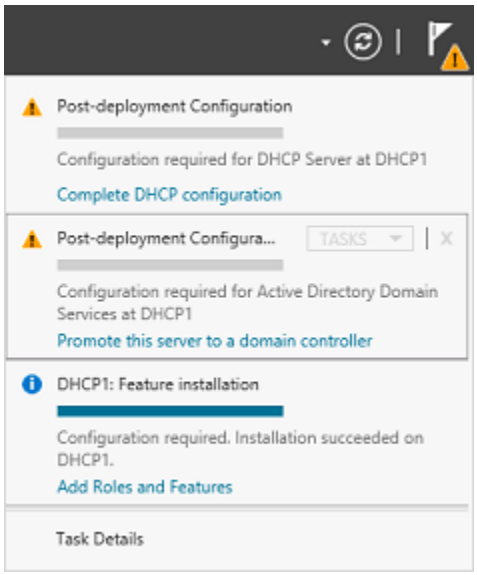

17. In the **Active Directory Domain Services Configuration Wizard**, on the **Deployment Configuration** page, choose **Add a new forest** and then next to **Root domain name**, type **contoso.com**.
18. Click **Next**, and then on the **Domain Controller Options** page, under **Type the Directory Services Restore Mode (DSRM) password**, type a password next to **Password** and **Confirm password**. Confirm that **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected, and then click **Next**.
19. Click **Next** four times, verify that **All prerequisite checks passed successfully** is displayed, and then click **Install**.
20. The computer will restart automatically to complete the installation process.
21. Sign in using the local Administrator account.

Create a domain administrator account

A domain administrator account is required to configure settings in the test lab.

Tip

You can use the CONTOSO\Administrator account in this test lab and skip creation of a domain administrator account if desired. This account has domain administrator privileges, and other privileges. However, it is a best practice to disable or rename this account. For more information, see Active Directory Best Practices(http://go.microsoft.com/fwlink/p/?LinkID=243071).

To create a domain administrator account

1. On the Server Manager menu bar, click **Tools**, and then click **Active Directory Users and Computers**.
2. In the **Active Directory Users and Computers** console tree, double-click **contoso.com**,

right-click **Users**, point to **New**, and then click **User**.

3. In the **New Object – User** dialog box, type **user1** under **User logon name** and next to **Full name**, then click **Next**.
4. Next to **Password** and **Confirm password**, type a password for the user1 account.
5. Clear the checkbox next to **User must change password at next logon**, select the **Password never expires** checkbox, click **Next**, and then click **Finish**.
6. Double-click **user1** and then click the **Member Of** tab.
7. Click **Add**, type **domain admins** under **Enter the object names to select**, click **OK** twice, and then close the **Active Directory Users and Computers** console.
8. Click **Start**, click **Administrator**, and then click **Sign out**.
9. Sign in to the computer using the user1 credentials by clicking the left arrow next to **CONTOSO\Administrator** and then clicking **Other user**.

Authorize DHCP1 in Active Directory

So that DHCP1 can be configured as a failover partner, it must be authorized in Active Directory.

To authorize DHCP1 in Active Directory

1. On the Server Manager menu, click **Tools** and then click **DHCP**.
2. In the DHCP console, right-click **dhcp1.contoso.com** and then click **Authorize**.
3. Refresh the DHCP console and verify that DHCP1 was successfully authorized.
4. On the Server Manager menu, click the Notification flag and then click **Complete DHCP configuration**.
5. In the **DHCP Post-Install configuration wizard**, click **Commit** and then click **Close**.

Configure DHCP2

DHCP2 is a computer running Windows Server® 2012, providing the following services:

- A DHCP server.

Initial configuration of DHCP2 consists of the following steps:

- Install the operating system and configure TCP/IP on DHCP2
- Install and configure DHCP on DHCP2

During the demonstration portion of the test lab, DHCP2 will be used to create a failover relationship with DHCP1.

Install the operating system and configure TCP/IP on DHCP2

Tip

The procedure below is identical to the steps used to install the operating system and configure TCP/IP on DHCP1, with the exception that DHCP2 is configured with an IP address of 10.0.0.2.

To install the operating system and configure TCP/IP on DHCP2

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. In the Server Manager navigation pane, click **Local Server** and then click the IP address next to **Wired Ethernet Connection**. The **Network Connections** control panel will open.
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.2** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install and configure DHCP on DHCP2

DHCP2 is a domain member server running the DHCP Server role service.

To install DHCP on DHCP2

1. In the Server Manager navigation pane, click **Local Server** and then click the name next to **Computer name**. The **System Properties** control panel will open.
2. On the **Computer Name** tab, click **Change** and then type **DHCP2** under **Computer name**.
3. Under **Member of**, select **Domain**, type **contoso.com**, and then click **OK**.
4. When you are prompted to provide credentials to join the domain, enter the credentials for the user1 account that was created previously and then click **OK**.
5. Confirm that computer name and domain changes were successful, click **OK** twice, and then click **Close**.
6. When you are prompted to restart the computer, click **Restart Now**.
7. After restarting the computer, sign-in using the CONTOSO\user1 account.
8. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
9. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **DHCP Server** checkbox.

10. When you are prompted to add required features, click **Add Features**.
11. Click **Next** three times, and then click **Install**.
12. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DHCP2.contoso.com** is displayed, and then click **Close**.

To configure DHCP on DHCP2

1. On the Server Manager menu bar, click **Tools** and then click **DHCP**. THE DHCP console opens.
2. In the DHCP console tree, navigate to **IPv4**. Right-click **IPv4** and then click **New Scope**. The **New Scope Wizard** opens.
3. Click **Next** and then type a name for the new scope next to **Name** (ex: Contoso-scope1).
4. Click **Next** and then in **IP Address Range**, type **10.0.0.1** next to **Start IP address**, type **10.0.0.254** next to **End IP address**, and type **24** next to **Length**. The value of Subnet mask will change automatically to **255.255.255.0**.
5. Click **Next**, and then in **Add Exclusions and Delay** type **10.0.0.1** under **Start IP address**, type **10.0.0.10** under **End IP address**, and then click **Add**. This allows the first ten IP addresses in the 10.0.0.0/24 subnet to be used for static addressing of servers on the network.
6. Click **Next** and then in **Lease Duration** under **Limited to** enter **0 Days**, **0 Hours**, and **2 Minutes**. This very short lease duration will simplify the DHCP failover demonstration.
7. Click **Next** three times, and then in **Domain Name and DNS Servers**, verify that the **Parent domain** is **contoso.com** and **10.0.0.1** is listed as the only DNS server.
8. Click **Next** twice, and then in **Activate Scope** select **Yes, I want to activate this scope now**.
9. Click **Next**, and then click **Finish**.
10. In the DHCP console tree, right-click dhcp2.contoso.com, and then click Authorize.
11. Refresh the view in the DHCP console and verify that DHCP2 is authorized and that the Contoso-scope1 is active.

    Note: To review scopes on the current server using Windows PowerShell, right-click **Windows PowerShell**, click **Run as Administrator**, click **Yes** in the User Account Control alert that appears, and then type the following command at the Windows PowerShell prompt, and then press ENTER.

```
get-dhcpserverv4scope
```

```
PS C:\Windows\system32> get-dhcpserverv4scope


ScopeId         SubnetMask      Name           State    StartRange      EndRange
LeaseDuration
-------         ----------      ----           -----    ----------      --------
-------------
```

```
10.0.0.0        255.255.255.0   Contoso-scope1 Active   10.0.0.1        10.0.0.254
00:02:00
```

Configure Client1

Client1 is a computer running Windows® 8 that is acting as a DHCP client.

Configuration of Client1 consists of the following steps:

- Install the operating system on Client1
- Pin Windows PowerShell to the taskbar

During the demonstration portion of the test lab, Client1 will be used as a DHCP client.

Install the operating system on Client1

To install the operating system on Client1

1. Start your computer using the Windows 8 product disc or other digital media.
2. When prompted, enter a product key and accept license terms.
3. When prompted to enter a computer name, type **Client1** and click **Next**.
4. Click Use express settings.
5. On the **Sign in to your PC** page, click **Don't want to sign in with a Microsoft account** and then click **Local account**.
6. Next to **User name**, type **user1**, enter a password and password hint, and then click **Finish**.

Pin Windows PowerShell to the taskbar

The DHCP failover demonstration on Client1 makes use of Windows PowerShell to verify DHCP lease information. To make Windows PowerShell more easily accessible, it will be pinned to the taskbar.

To pin Windows PowerShell to the taskbar

1. The **Start** menu is displayed by default. If **Start** is not displayed, move the mouse cursor to the lower left corner of the screen until **Start** is displayed, and then click **Start**.
2. Type **power** and then under **Results for "power"** right-click **Windows PowerShell** and then click **Pin to taskbar**.Confirm that Windows PowerShell is pinned to the taskbar.

Note

Client1 can also be joined to the contoso.com domain, however this is not required to complete the test lab.

DHCP failover demonstration

For the DHCP failover demonstration portion of the test lab, a failover relationship will be created using DHCP1 and DHCP2.

A demonstration of DHCP failover on Windows Server 2012 consists of the following procedures:

1. Configure a failover relationship
2. View or edit properties of the failover configuration
3. Edit properties of the failover relationship and demonstrate load balancing
4. Edit properties of the failover relationship and demonstrate hot standby mode

Configure a failover relationship

First, configure a failover relationship using DHCP1 and DHCP2.

To configure a failover relationship

1. On DHCP2, open the DHCP console, right-click the **Contoso-scope1** DHCP scope and then click **Configure failover**.

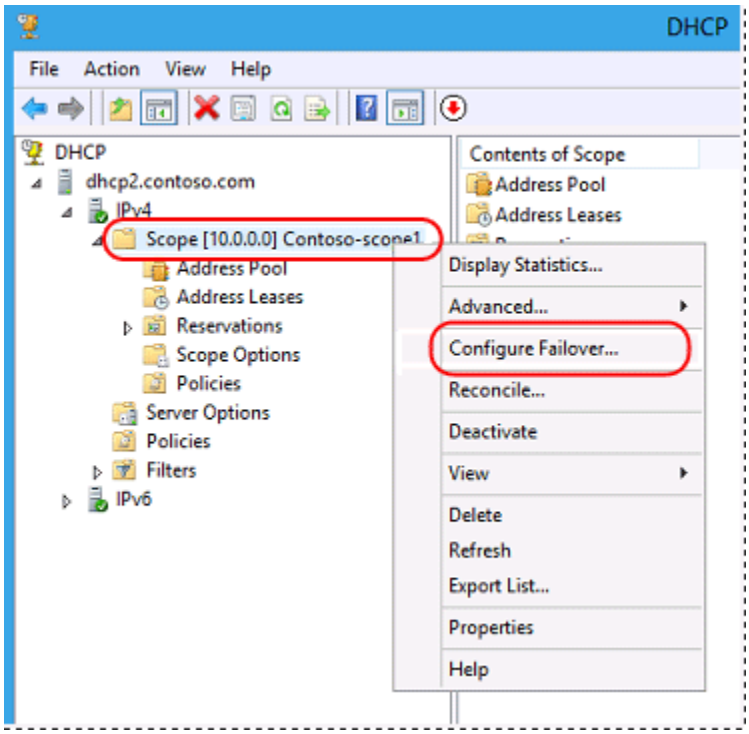

2. In the **Configure Failover** wizard, click **Next**.
3. In **Specify the partner server to use for failover**, next to **Partner Server**, type **dhcp1.contoso.com** and then click **Next**.
4. In **Create a new failover relationship**, type a name next to **Relationship Name**, or accept the default name that is displayed (dhcp2.contoso.com-dhcp1.contoso.com).

5. Type a shared secret for this failover relationship next to **Shared Secret** (ex: secret).
6. Change the value next to **Maximum Client Lead Time** to **0** hours and **1** minute.

   Important

   The Maximum Client Lead Time (MCLT) is additional time provided to a DHCP client after expiration of a DHCP lease. The MCLT is transmitted from the primary to the secondary server in the CONNECT message, and is the maximum amount of time that one server can extend a lease for a client beyond the time known by the partner server.

   The 1 minute MCLT value used here is for test lab purposes only, to prompt lease renewal by the client. In a production environment, you should use a longer MCLT, such as 1 hour.
7. Review the options available in the drop-down menu next to **Mode**. You can choose **Load balance** or **Hot standby**. By default, **Load balance** mode is chosen.

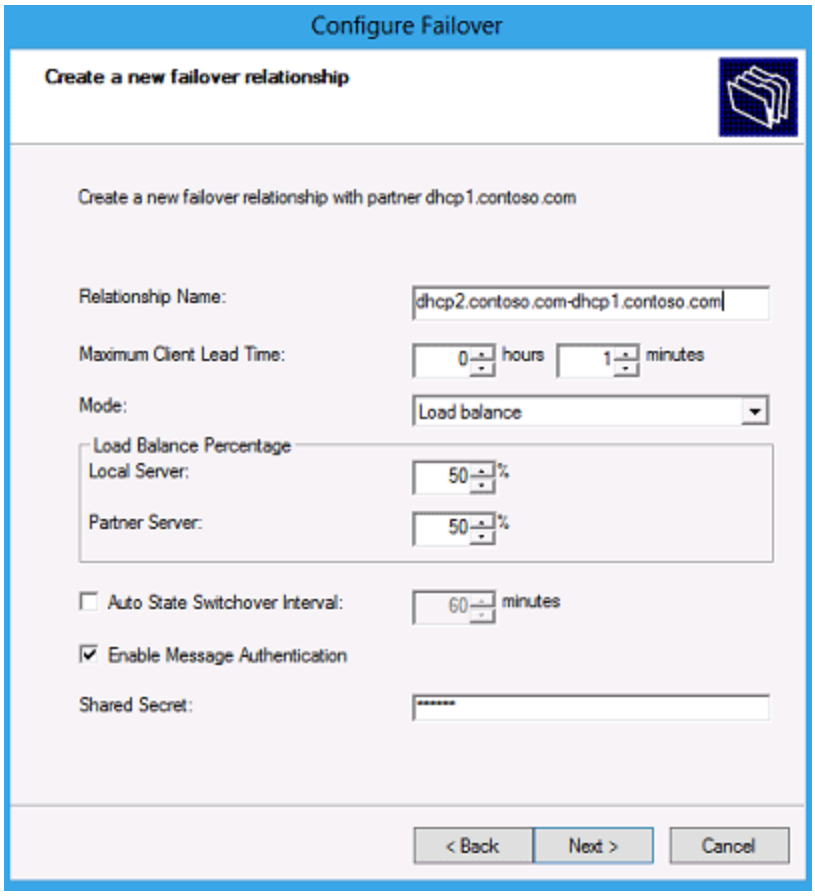

8. Click **Next** and then click **Finish**.
9. Verify that failover configuration was successful, and then click **Close**.

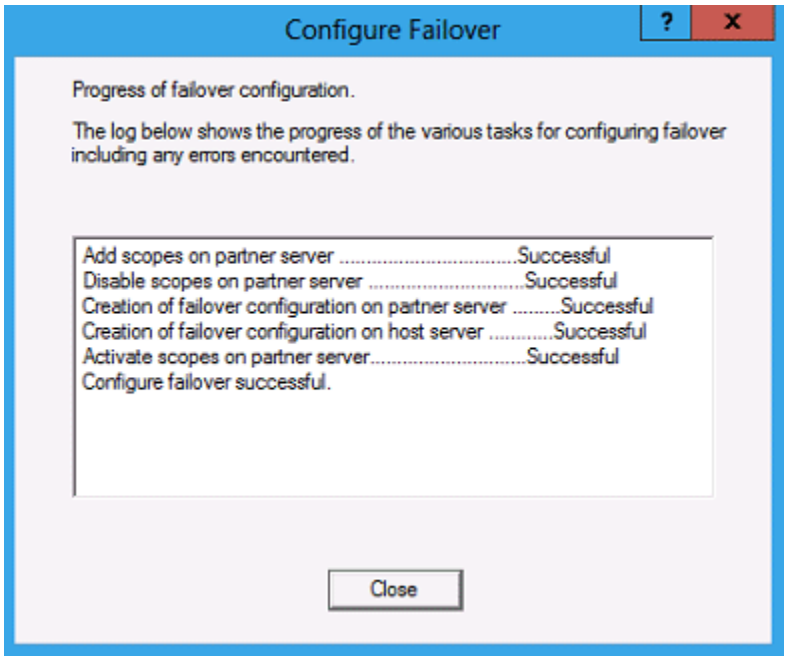

10. On DHCP1, refresh the DHCP console and verify that the same DHCP scope configuration that is present on DHCP2 is now present on DHCP1.

View or edit properties of the failover configuration

After you configure a failover relationship on a DHCP server, details for the failover relationship are displayed in the DHCP console.

To view or edit properties of the failover relationship

1. On DHCP1 or DHCP2, in the DHCP console, right-click the **Contoso-scope1** DHCP scope and then click **Properties**.
2. Click the **Failover** tab and review the information displayed. Verify that **Normal** is displayed next to **State of this Server** and also next to **State of Partner Server**.

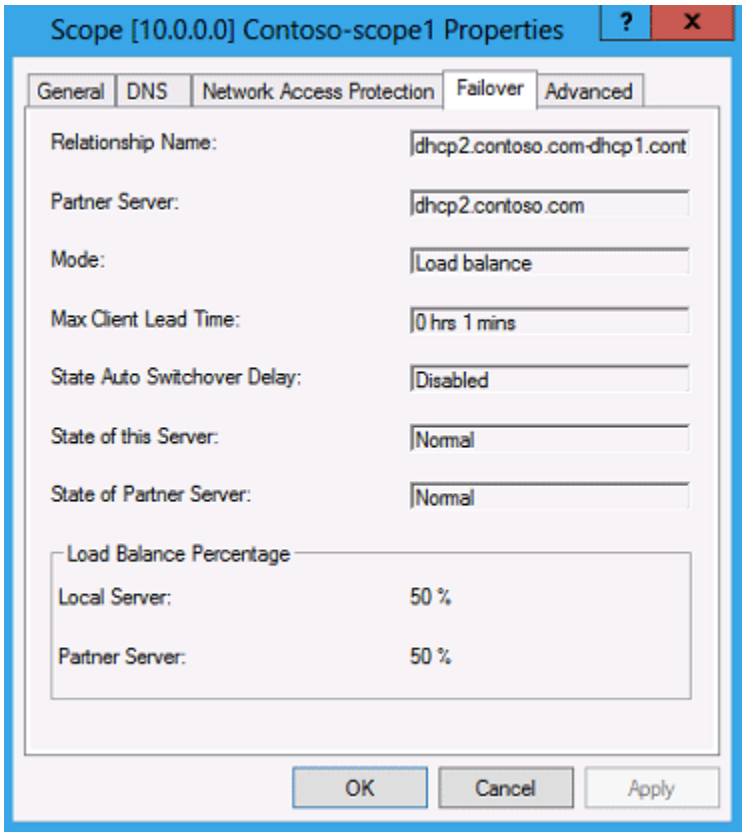

3. Note that you can edit or delete the failover relationship.

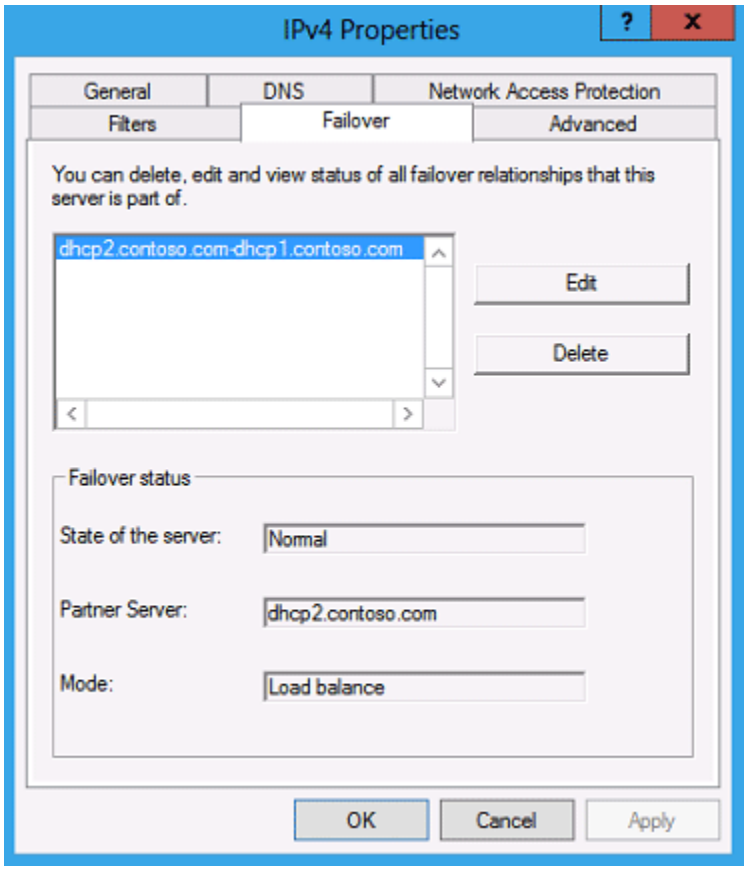

4. Click **Edit** and review properties of the failover relationship that are available to edit.
5. Leave the dialog box open for the following procedure.

Edit properties of the failover relationship and demonstrate load balancing

To demonstrate dynamic load balancing properties of the failover relationship, the load balancing percentage will be changed for an active DHCP scope.

To edit properties of the failover relationship and demonstrate load balancing

1. On Client1, click **Windows PowerShell** and type the following command.

   ```
   ipconfig /all
   ```

2. In the command output, note the DHCP server that is currently supplying an IP address configuration to Client1. The IP address of the DHCP server is displayed next to **DHCP Server**.
3. In the **View/Edit Failover Relationship** dialog box DHCP1 or DHCP2 that was opened in the previous procedure, change the values under **Load Balance Mode** next to **Local Server** and **Partner Server** so that **100%** is assigned to the DHCP server that is currently not supplying an IP address to Client1. The server that is currently supplying an

IP address to Client1 will have a value assigned of **0%**.

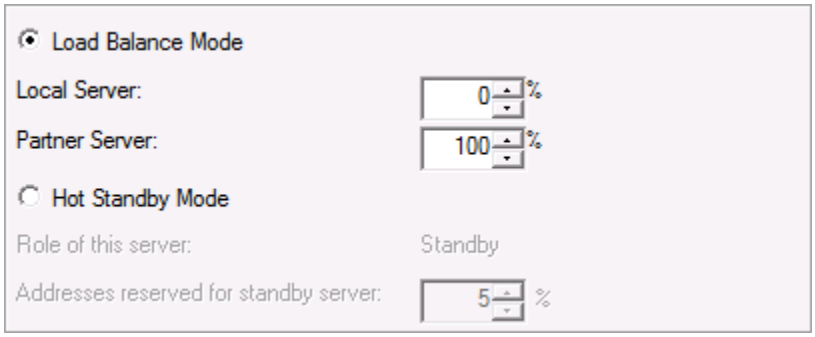

4. Click **OK** twice, wait until the current DHCP lease is expired on Client1, and then type **ipconfig /all** again at the Windows PowerShell prompt.
5. Note that the DHCP server that is supplying an IP address configuration to Client1 has changed to the server that you assigned 100% weight in load balancing mode.

Edit properties of the failover relationship and demonstrate hot standby mode

To demonstrate hot standby mode, the DHCP Server service on one of the failover partners will be stopped.

To edit properties of the failover relationship and demonstrate hot standby mode

1. On DHCP1 or DHCP2, in the DHCP console, right-click the **Contoso-scope1** DHCP scope and then click **Properties**.
2. Click the **Failover** tab.
3. Click **Edit** and then choose **Hot Standby Mode**.
4. Depending on which DHCP server you are configuring, the local server will be assigned either **Active** or **Standby** status. The status is displayed next to **Role of this server**.

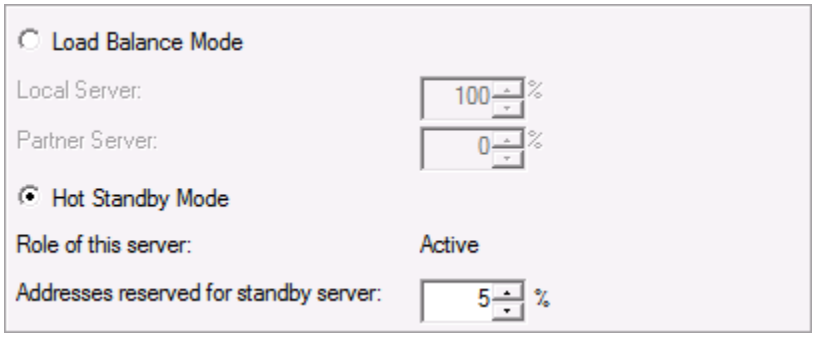

Tip

The server that is designated to be Active in hot standby mode is the server that you used to create the failover relationship.

5. Click **OK** twice and then wait 2 minutes for the DHCP lease on Client1 to renew.
6. On Client1, type **ipconfig /all** at the Windows PowerShell prompt and verify that the server that is assigned as Active is supplying an IP addresses configuration to Client1.
7. In the DHCP console on the DHCP server that is marked as **Active** for the hot standby failover relationship and is currently supplying an IP address to Client1, right-click the server name, point to All Tasks, and then click Stop.
8. Verify that the DHCP service is stopped on the active DHCP server.

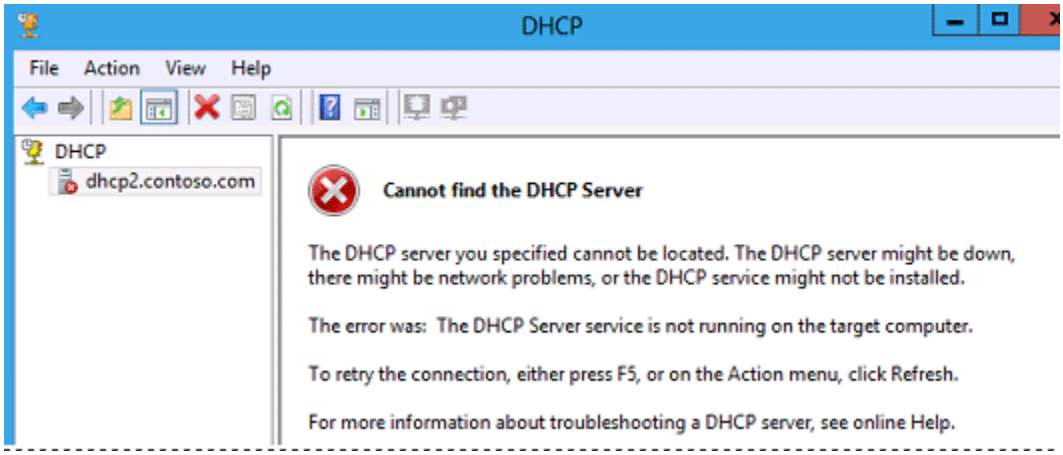

9. Wait for the DHCP lease to renew on Client1, type **ipconfig /all** at the Windows PowerShell prompt, and verify that the standby DHCP server is supplying an IP address to Client1.

Conclusion

DHCP failover provides high availability of DHCP services without the challenges of clustering or split scope DHCP. Benefits of DHCP failover include:

1. **Simple**: A wizard is provided to create DHCP failover relationships between DHCP servers. The wizard automatically replicates scopes and settings from the primary server to the failover partner.
2. **Flexible**: DHCP failover can also be configured for load balancing, with client requests distributed between both DHCP servers in a failover relationship based on the values you choose.
3. **Seamless**: DHCP servers share lease information, allowing one server to assume responsibility for servicing of clients if the other server is unavailable. DHCP clients can keep the same IP address when a lease is renewed, even if the lease is issued by a different DHCP server.

See also

Dynamic Host Configuration Protocol (DHCP) Overview

What's New in DHCP in Windows Server 2012

Step-by-Step: Configure DHCP Using Policy-based Assignment

Extensible Authentication Protocol (EAP) for Network Access Overview

The Extensible Authentication Protocol (EAP) is an architectural framework that provides extensibility for authentication methods for commonly used protected network access technologies, such as IEEE 802.1X-based wireless access, IEEE 802.1X-based wired access, and Point-to-Point Protocol (PPP) connections such as Virtual Private Networking (VPN). EAP is not an authentication method like MS-CHAP v2, but rather a framework on the access client and authentication server that allows networking vendors to develop and easily install new authentication methods known as EAP methods.

For setting details about EAP, see Extensible Authentication Protocol (EAP) Settings for Network Access.

Extensible Authentication Protocol (EAP) Settings for Network Access

This topic presents information about the Extensible Authentication Protocol (EAP) default settings that you can use to configure computers running Windows® 8, Windows® 7, and Windows Vista®.

Authentication methods

This topic contains configuration information specific to the following authentication methods in EAP.

Protected EAP (PEAP)

Protected EAP Properties configuration items

This section contains configuration information for the two default inner EAP methods that are provided with PEAP.

- **EAP-Transport Layer Security (TLS)**

  Appears as **Smart Card or other Certificate Properties** in the operating system. EAP-TLS can be deployed as an inner method for PEAP or as a standalone EAP method. When it is configured as an inner authentication method, the configuration settings for EAP-TLS are identical to the settings that are used to deploy EAP-TLS as an outer method, except that it is configured to operate within PEAP. For configuration details, see Smart Card or other Certificate Properties configuration items.

- **EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2)**

  Secure password (EAP-MSCHAP v2) Properties configuration items

  Note
  EAP authentication methods that are used within tunneled EAP methods are commonly known as inner methods, and they are also referred to as **EAP types** in some documentation.

EAP-TLS

Smart Card or other Certificate Properties configuration items

This section about Smart Card or other Certificate Properties includes information about the following configurations:

- Configure New Certificate Selection configuration items
- Select EKUs
- Add or Edit EKU

EAP-Tunneled Transport Layer Security (TTLS)

TTLS configuration items

EAP-Subscriber Identity Module (SIM), EAP-Authentication and Key Agreement (AKA), and EAP-AKA Prime (AKA')

EAP-SIM, EAP-AKA, and EAP-AKA’

Enables authentication by using SIM cards, and is implemented when a customer purchases a wireless broadband service plan from a mobile network operator. As part of the plan, the customer commonly receives a wireless profile that is preconfigured for SIM authentication.

This section provides information about the following configuration settings:

- EAP-SIM
- EAP-AKA
- EAP-AKA’

EAP-TLS, PEAP, and EAP-TTLS

You can access the EAP properties for 802.1X authenticated wired and wireless access in the following ways:

- By configuring the Wired Network (IEEE 802.3) Policies and Wireless Network (IEEE 802.11) Policies extensions in Group Policy.
- By manually configuring wired or wireless connections on client computers.

You can access the EAP properties for virtual private network (VPN) connections in the following ways:

- By using Connection Manager Administration Kit (CMAK) to configure VPN connections.

- By manually configuring VPN connections on client computers.

By default, you can configure EAP settings for the following network authentication methods for 802.1X authenticated wired access, 802.1X authenticated wireless access, and VPN:

- **Microsoft: Smart Card or other Certificate** (EAP-TLS)
- **Microsoft: Protected EAP (PEAP)**
- **Microsoft: EAP-TTLS**

Additionally, the MS-CHAP-V2 network authentication method is available for VPN by default.

Protected EAP Properties configuration items

Important
Deploying the same type of authentication method for PEAP and EAP creates a security vulnerability. When you deploy both PEAP and EAP (which is not protected), do not use the same authentication type. For example, if you deploy PEAP-TLS, do not also deploy EAP-TLS.

Item	Details
Verify the Server's identity by validating the certificate	Specifies that the client verifies that server certificates presented to the client computer have the correct signatures, have not expired, and were issued by a trusted root certification authority (CA). **Important** If you disable this check box, client computers cannot verify the identity of your servers during the authentication process. If server authentication does not occur, users are exposed to severe security risks, including the possibility that users might unknowingly connect to a rogue network. Default = enabled
Connect to these servers	Allows you to specify the name for Remote Authentication Dial-In User Service (RADIUS) servers that provide network authentication and authorization. **Note** You must type the name exactly as it appears in the **Subject** field of each

Item	Details
	RADIUS server certificate, or use regular expressions to specify the server name. The complete syntax of the regular expression can be used to specify the server name, but to differentiate a regular expression with the literal string, you must use at least one "*" in the string specified. For example, you can specify *nps*.example.com* to specify the RADIUS server nps1.example.com or nps2.example.com. Defaults: Wired and wireless = not enabled VPN = enabled **Note** Even if no RADIUS servers are specified, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
Trusted Root Certification Authorities	Lists the trusted root certification authorities. The list of trusted root certification authorities is built from the trusted root CAs that are installed in the computer and in the user certificate stores. You can specify which trusted root CA certificates supplicants use to determine whether they trust your servers, such as your server running Network Policy Server (NPS) or your provisioning server. If no trusted root CAs are selected, the 802.1X client verifies that the computer certificate of the RADIUS server was issued by an installed trusted root CA. If one or multiple trusted root CAs are selected, the 802.1X client verifies that the computer certificate of the RADIUS server was issued by a selected trusted root CA. If you have a public key infrastructure (PKI) on your network, and you use your CA to issue certificates to your RADIUS servers, your CA certificate is automatically added to the list of

Item	Details
	trusted root CAs. You can also purchase a CA certificate from a non-Microsoft vendor. Some non-Microsoft trusted root CAs provide software with your purchased certificate that automatically installs the purchased certificate into the **Trusted Root Certification Authorities** certificate store. In this case, the trusted root CA automatically appears in the list of trusted root CAs. **Notes** Do not specify a trusted root CA certificate that is not already listed in client computers' **Trusted Root Certification Authorities** certificate stores for **Current User** and **Local Computer**. If you designate a certificate that is not installed on client computers, authentication will fail. Default = not enabled, no trusted root CAs selected. **Note** Even if no trusted root CAs are selected, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
Notifications before connecting	Specifies whether the user is notified if the server name or root certificate is not specified, or whether the server's identity cannot be verified. By default, the following options are provided: Case 1: **Do not ask user to authorize new servers or trusted CAs** specifies that if: a. The server name is not in the **Connect to these servers** list b. or the root certificate is found but is not selected in the list of **Trusted Root Certification Authorities** in **PEAP**

Item	Details
	Properties c. or the root certificate is not found on the computer then the user is not notified, and the connection attempt fails. Case 2: **Tell user if the server name or root certificate is not specified** specifies that if: a. the server name is not in the **Connect to these servers** list b. or the root certificate is found but is not selected in the list of **Trusted Root Certification Authorities** in **PEAP Properties** the user is prompted whether to accept the root certificate. If the user accepts the certificate, authentication proceeds. If the user rejects the certificate, the connection attempt fails. **Note** In this option, if the root certificate is not present on the computer, the user is not notified and the connection attempts fails. Case 3: **Tell user if the server's identity cannot be verified** Specifies that if: a. the server name is not in the **Connect to these servers** list b. or the root certificate is found but is not selected in the list of **Trusted Root Certification Authorities** in **PEAP Properties** c. or the root certificate is not found on the computer the user is prompted whether to accept the root certificate. If the user accepts the certificate, authentication proceeds. If the user rejects the certificate, the connection

Item	Details
	attempt fails.
Select Authentication Method	Allows you to select the EAP type to use with PEAP for network authentication. **Note** By default, two EAP types are available, **Secured password (EAP-MSCHAPv2)** and **Smart card or other certificate** (EAP-TLS). However, EAP is a flexible protocol that allows inclusion of additional EAP methods, and it is not restricted to these two types. For information about **Secured password (EAP-MSCHAPv2)** or **Smart card or other certificate** (EAP-TLS) configuration settings, see: • Secure password (EAP-MSCHAP v2) Properties configuration items • Smart Card or other Certificate Properties configuration items Default = **Secured password (EAP-MSCHAP v2)**
Configure	Provides access to property settings for the specified EAP type.
Enable Fast Reconnect	Enables the ability to create a new or refreshed security association more efficiently or in a smaller number of round- trips, in the case where a security association was previously established. For VPN connections, Fast Reconnect uses IKEv2 technology to provide seamless and consistent VPN connectivity, when users temporarily lose their Internet connections. Users who connect by using wireless mobile broadband will benefit most from this capability. An example of this benefit is a common scenario in which a user is traveling on a train, uses a wireless mobile broadband card to

Item	Details
	connect to the Internet, and then establishes a VPN connection to the corporate network. As the train passes through a tunnel, the Internet connection is lost. When the train is outside the tunnel, the wireless mobile broadband card automatically reconnects to the Internet. In client versions prior to Windows 7, VPN does not automatically reconnect. The user must repeat the multistep process to connect to the VPN each time Internet connectivity is interrupted. This can quickly become time consuming for mobile users with intermittent connectivity disruptions. In Windows 8, Fast Reconnect automatically re-establishes active VPN connections when Internet connectivity is re-established. Although the reconnection might take several seconds to occur, it is performed transparently to users. Default = enabled
Enforce Network Access Protection	Specifies that before connections to a network are permitted, system health checks are performed on EAP supplicants to determine if they meet system health requirements. Default = not enabled
Disconnect if server does not present cryptobinding TLV	Specifies that connecting clients must end the network authentication process if the RADIUS server does not present cryptobinding Type-Length-Value (TLV). **Note** Cryptobinding TLV increases the security of the TLS tunnel in PEAP by combining the inner method and the outer method authentications together so that attackers cannot perform man-in-the-middle attacks by redirecting an MS-CHAP v2 authentication by using the PEAP channel.

Item	Details
	Default = not enabled
Enable Identity Privacy	Specifies that clients are configured so that they cannot send their identity before the client has authenticated the RADIUS server, and optionally, provides a place to type an anonymous identity value. For example, if you select **Enable Identity Privacy**, and then type “guest” as the anonymous identity value, the identity response for a user with identity *alice@example* is *guest@example*. If you select **Enable Identity Privacy** but do not provide an anonymous identity value, the identity response for the user *alice@example* is *@example*. This setting applies only to computers running Windows 7 and Windows 8. Default = not enabled

Secure password (EAP-MSCHAP v2) Properties configuration items

Secure password EAP-MS-CHAP v2 is an EAP type that can be used with PEAP, for password-based network authentication. EAP-MsCHAPv2 can also be used as a standalone method for VPN, but only as a PEAP inner method for wireless.

Item	Details
Automatically use my Windows logon name and password (and domain if any)	Specifies that the current user-based Windows sign in name and password are used as network authentication credentials. Defaults: • Authenticated Wired and Wireless Access = enabled • VPN = not enabled

Smart Card or other Certificate Properties configuration items

Item	Details
Use my smart card	Specifies that clients making authentication

Item	Details
	requests must present a smart card certificate for network authentication. Defaults: Wired and wireless = not enabled VPN = enabled
Use a certificate on this computer	Specifies that authenticating clients must use a certificate located in the **Current User** or **Local Computer** certificate stores. Defaults: Wired and wireless = enabled VPN = not enabled
Use simple certificate selection (Recommended)	Specifies whether Windows filters out certificates that are unlikely to meet authentication requirements. This serves to limit the list of available certificates when prompting the user to select a certificate. Defaults: Wired and wireless = enabled VPN = not enabled
Advanced	Opens the **Configure Certificate Selection** dialog box. For more information about **Configure Certificate Selection**, see Configure New Certificate Selection configuration items.
Verify the server's identity by validating the certificate	Specifies that the client verifies that the server certificates presented to the client computer have the correct signatures, have not expired, and were issued by a trusted root certification authority (CA). **Important** Do not disable this check box or client computers cannot verify the identity of your servers during the authentication process. If server authentication does not occur, users are exposed to severe security risks, including the possibility that users might unknowingly connect

Item	Details
	to a rogue network. Default = enabled
Connect to these servers	Allows you to specify the name for RADIUS servers that provide network authentication and authorization. **Note** You must type the name exactly as it appears in the **Subject** field of each RADIUS server certificate, or use regular expressions to specify the server name. The complete syntax of the regular expression can be used to specify the server name, but to differentiate a regular expression with the literal string, you must use at least one "*" in the string that is specified. For example, you can specify *nps*.example.com* to specify the RADIUS server nps1.example.com or nps2.example.com. Defaults: Wired and wireless = not enabled VPN = enabled **Note** Even if no RADIUS servers are specified, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
Trusted Root Certification Authorities	The list in **Trusted Root Certification Authorities** is built from the trusted root CAs that are installed in the computer and user certificate stores. You can specify which trusted root CA certificates that supplicants use to determine whether they trust your servers, such as your server running NPS or your provisioning server. If no trusted root CAs are selected, the 802.1X client verifies that the computer certificate of the RADIUS server was

Item	Details
	issued by an installed trusted root CA. If one or multiple trusted root CAs are selected, the 802.1X client verifies that the computer certificate of the RADIUS server was issued by a selected trusted root CA. If you have a public key infrastructure (PKI) on your network, and you use your CA to issue certificates to your RADIUS servers, your CA certificate is automatically added to the list of trusted root CAs. You can also purchase a CA certificate from a non-Microsoft vendor. Some non-Microsoft trusted root CAs provide software with your purchased certificate that automatically installs the purchased certificate into the **Trusted Root Certification Authorities** certificate store. In this case, the trusted root CA automatically appears in the list of trusted root CAs. Do not specify a trusted root CA certificate that is not already listed in client computers' **Trusted Root Certification Authorities** certificate stores for **Current User** and **Local Computer**. **Note** If you designate a certificate that is not installed on client computers, authentication will fail. Default = not enabled, no trusted root CAs selected. **Note** Even if no trusted root CAs are selected, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
View Certificate	Enables you to view the properties of the selected certificate.
Do not prompt user to authorize new servers or trusted certification authorities	Prevents the user from being prompted to trust a server certificate if that certificate is

Item	Details
	incorrectly configured, is not already trusted, or both (if enabled). It is recommended that you select this check box to simplify the user experience and to prevent users from inadvertently choosing to trust a server that is deployed by an attacker. Default = not enabled
Use a different user name for the connection	Specifies whether to use a user name for authentication that is different from the user name in the certificate. Default = not enabled

Configure New Certificate Selection configuration items

Use **New Certificate Selection** to configure the criteria that client computers use to automatically select the right certificate on the client computer for the purpose of authentication. When the configuration is provided to network client computers through the Wired Network (IEEE 802.3) Policies, the Wireless Network (IEEE 802.11) Policies, or through Connection Manager Administration Kit (CMAK) for VPN, clients are automatically provisioned with the specified authentication criteria.

Item	Details
Certificate Issuer	Specifies whether Certificate Issuer filtering is enabled. Default = not selected
Certificate Issuer list	Used to specify one or multiple certificate issuers for the certificates. Lists the names of all of the issuers for which corresponding certification authority (CA) certificates are present in the **Trusted Root Certification Authorities** or **Intermediate Certification Authorities** certificate store of local computer account. • Includes all the root certification authorities and intermediate certification authorities. • Contains only those issuers for which

Item	Details
	there are corresponding valid certificates that are present on the computer (for example, certificates that are not expired or not revoked). The final list of certificates that are allowed for authentication contains only those certificates that were issued by any of the issuers selected in this list. Default = none selected
Extended Key Usage (EKU)	You can select **All Purpose**, **Client Authentication**, **Any purpose**, or any combination of these. Specifies that when a combination is selected, all the certificates satisfying at least one of the three conditions are considered valid certificates for the purpose of authenticating the client to the server. If EKU filtering is enabled, one of the choices must be selected; otherwise, the **OK** command control is disabled. Default = not enabled
All Purpose	Specifies that (when selected) certificates having the **All Purpose** EKU are considered valid certificates for the purpose of authenticating the client to the server. Default = selected when **Extended Key Usage (EKU)** is selected.
Client Authentication	Specifies that (when selected) certificates having the **Client Authentication** EKU, and the specified list of EKUs are considered valid certificates for the purpose of authenticating the client to the server. Default = selected when **Extended Key Usage (EKU)** is selected.
Add	Opens the Select EKUs dialog box, which enables you to add standard, custom, or vendor-specific EKUs to the **Client Authentication** list. Default = no EKUs listed

Item	Details
Remove	Removes the selected EKU item from the **Client Authentication** list. Default = N/A
Any Purpose	Specifies that (when selected) all certificates having **Any Purpose** EKU and the specified list of EKUs are considered valid certificates for the purpose of authenticating the client to the server. Default = selected when **Extended Key Usage (EKU)** is selected.
Add	Opens the Select EKUs dialog box, which enables you to add standard, custom, or vendor-specific EKUs into the **Any Purpose** list. Default = no EKUs listed
Remove	Removes the selected EKU item from the **Any Purpose** list. Default = N/A

Note
When both **Certificate Issuer** and **Extended Key Usage (EKU)** are enabled, only those certificates that satisfy both conditions are considered valid for the purpose of authenticating the client to the server.

Select EKUs

You can select an EKU from the list provided, or add a new EKU.

Item	Details
Add	Opens the Add or Edit EKU dialog box, which enables you to define and add custom EKUs. In **Select the EKUs from the list below**, select an EKU in the list, and then click **OK** to add that EKU to the **Client Authentication** or the **Any Purpose** list.
Edit	Opens the Add or Edit EKU dialog box, and enables you to edit custom EKUs that you have added. You cannot edit the default, predefined

Item	Details
	EKUs.
Remove	Removes the selected custom EKU from the list of EKUs in the Select EKUs dialog box. You cannot remove the default, predefined EKUs.

Add or Edit EKU

Item	Details
Enter the name of the EKU	Provides a place to type the name of the custom EKU.
Enter the EKU OID:	Provides a place to type the OID for the EKU. Only numeric digits, separators, and "." are allowed. Wild cards are permitted, in which case all of the child OIDs in the hierarchy are allowed. For example, entering **1.3.6.1.4.1.311.*** allows 1.3.6.1.4.1.311.42 and 1.3.6.1.4.1.311.42.2.1

TTLS configuration items

EAP-TTLS is a standards-based EAP tunneling method that supports mutual authentication and provides a secure tunnel for client inclusion authentication by using EAP methods and other legacy protocols. The addition of EAP-TTLS in Windows Server 2012 provides only client-side support, for the purpose of supporting interoperation with the most commonly-deployed RADIUS servers that support EAP-TTLS.

Item	Description
Enable Identity Privacy	Specifies that clients are configured so that they cannot send their identity before the client has authenticated the RADIUS server, and optionally, provides a place to type an anonymous identity value. For example, if you select **Enable Identity Privacy**, and then type "guest" as the anonymous identity value, the identity response for a user with identity *alice@example* is *guest@example*. If you select **Enable Identity Privacy** but do not

Item	Description
	provide an anonymous identity value, the identity response for the user *alice@example* is *@example*. This setting applies only to computers running Windows 8. Default = not enabled
Connect to these servers	Enables you to specify the name for RADIUS servers that provide network authentication and authorization. **Note** You must type the name exactly as it appears in the **Subject** field of each RADIUS server certificate, or use regular expressions to specify the server name. The complete syntax of the regular expression can be used to specify the server name. But to differentiate a regular expression with the literal string, you must use at least one '*' in the string specified. For example, you can specify *nps*.example.com* to specify the RADIUS server nps1.example.com or nps2.example.com. Defaults: = none **Note** Even if no RADIUS servers are specified, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
Trusted Root Certification Authorities	The list in **Trusted Root Certification Authorities** is built from the trusted root CAs that are installed in the computer and user certificate stores. You can specify which trusted root CA certificates that supplicants use to determine whether they trust your servers, such as your server running NPS or your provisioning server. If no trusted root CAs are selected, the 802.1X client verifies that the

Item	Description
	computer certificate of the RADIUS server was issued by an installed trusted root CA. If one or multiple trusted root CAs are selected, the 802.1X client verifies that the computer certificate of the RADIUS server was issued by a selected trusted root CA. If you have a public key infrastructure (PKI) on your network, and you use your CA to issue certificates to your RADIUS servers, your CA certificate is automatically added to the list of trusted root CAs. If selected, your root CA certificate is installed on a client computer when the computers are joined to the domain. You can also purchase a CA certificate from a non-Microsoft vendor. Some non-Microsoft trusted root CAs provide software with your purchased certificate that automatically installs the purchased certificate into the **Trusted Root Certification Authorities** certificate store. In this case, the trusted root CA automatically appears in the list of trusted root CAs. Do not specify a trusted root CA certificate that is not already listed in client computers' **Trusted Root Certification Authorities** certificate stores for **Current User** and **Local Computer**. **Note** If you designate a certificate that is not installed on client computers, authentication will fail. Default = not enabled, no trusted root CAs selected. **Note** Even if no trusted root CAs are selected, the client will verify that the RADIUS server certificate was issued by a trusted root CA.
Do not prompt user if unable to authorize server	Specifies (when not selected) that if server certificate validation fails due to any of the

Item	Description
	following reasons, the user is prompted to accept or reject the server: • A root certificate for the server certificate is not found or not selected in the **Trusted Root Certification Authorities** list. • One or more of the intermediate root certificates in the certificate chain is not found. • The subject name in the server certificate does not match any of the servers that are specified in the **Connect to these servers** list. Default = not selected
Select a non-EAP method for authentication	Specifies whether a non-EAP or an EAP type is used for authentication. If **Select a non-EAP method for authentication** is selected, **Select an EAP method for authentication** is disabled. If **Select a non-EAP method for authentication** is selected, the following non-EAP authentication types are provided in the drop-down list: • PAP • CHAP • MS-CHAP • MS-CHAP v2 Defaults: • **Select a non-EAP method for authentication** = enabled • Non-EAP type = PAP
Automatically use my Windows logon name and password	Uses Windows sign in credentials when enabled. This check-box is enabled only if **MS-CHAP v2** is selected in the **Select a non-EAP method for authentication** drop-down list. **Automatically use my Windows logon name and password** is disabled for PAP, CHAP and MS-CHAP authentication types.
Select an EAP method for authentication	Specifies whether an EAP type or a non-EAP type is used for authentication. If **Select an EAP method for authentication** is selected,

Item	Description
	Select a non-EAP method for authentication is disabled. If **Select a non-EAP method for authentication** is selected, by default, the following non-EAP authentication types are provided in the drop-down list: • Microsoft: Smart Card or other Certificate • Microsoft: MS-CHAP v2 • MS-CHAP • MS-CHAP v2 **Note** The **Select an EAP method for authentication** drop-down list will enumerate all the EAP methods that are installed on the server, except for PEAP and FAST tunnel methods. The EAP types are listed in the order that they are discovered by the computer.
Configure	Opens the properties dialog box of the specified EAP type. For details about the default EAP-types, see Smart Card or other Certificate Properties configuration items or Secure password (EAP-MSCHAP v2) Properties configuration items.

EAP-SIM, EAP-AKA, and EAP-AKA’

The following tables list the configuration settings for:

- EAP-SIM
- EAP-AKA
- EAP-AKA’

EAP-SIM

EAP Subscriber Identity Module (SIM) is used for authentication and session key distribution for the Global System for Mobile Communications (GSM). EAP-SIM is defined in RFC 4186.

EAP-SIM Properties

Item	Description
Use strong Cipher keys	Specifies that if selected, the profile uses strong encryption.
Do not reveal real identity to server when pseudonym identity is available	When enabled, forces the client to fail the authentication if server requests for permanent identity though the client have a pseudonym identity with it. Pseudonym identities are used for identity privacy so that the actual or permanent identity of a user is not revealed during authentication.
Enable usage of realms	Provides a place to type the realm name. If this field is left blank with **Enabled usage of realms** selected, the realm is derived from the International Mobile Subscriber Identity (IMSI) using the realm 3gpp.org, as described in the 3rd Generation Partnership Project (3GPP) standard 23.003 V6.8.0.
Specify a realm:	Provides a place to type the realm name.

EAP-AKA

EAP Authentication and Key Agreement (AKA) for Universal Mobile Telecommunications System (UMTS) is used for authentication and session key distribution by using the UMTS Universal Subscriber Identity Module (USIM). EAP AKA is defined in RFC 4187.

EAP-AKA Properties

Item	Description
Do not reveal real identity to server when pseudonym identity is available	When enabled, forces the client to fail the authentication if server requests for permanent identity though the client have a pseudonym identity with it. Pseudonym identities are used for identity privacy so that the actual or permanent identity of a user is not revealed during authentication.
Enable usage of realms	Provides a place to type the realm name. If this field is left blank with **Enabled usage of realms** selected, the realm is derived from the International Mobile Subscriber Identity (IMSI) using the realm 3gpp.org, as described in the

Item	Description
	3rd Generation Partnership Project (3GPP) standard 23.003 V6.8.0.
Specify a realm	Provides a place to type the realm name.

EAP-AKA’

EAP- AKA Prime (AKA') is a modified version of EAP-AKA that is used to enable access to the 3rd-Generation Partnership Project (3GPP)-based networks by using non-3GPP standards, such as:

- WiFi - sometimes referred to as wireless fidelity
- Evolution-Data Optimized (EVDO)
- Worldwide Interoperability for Microwave Access (WiMax)

EAP-AKA' is defined in RFC 5448.

EAP-AKA’ Properties

Item	Description
Do not reveal real identity to server when pseudonym identity is available	When enabled, forces the client to fail the authentication if server requests for permanent identity though the client have a pseudonym identity with it. Pseudonym identities are used for identity privacy so that the actual or permanent identity of a user is not revealed during authentication.
Enable usage of realms	Provides a place to type the realm name. If this field is left blank with **Enabled usage of realms** selected, the realm is derived from the International Mobile Subscriber Identity (IMSI) using the realm 3gpp.org, as described in the 3rd Generation Partnership Project (3GPP) standard 23.003 V6.8.0.
Specify a realm	Provides a place to type the realm name.
Ignore network name mismatch	The client compares the name of network known to it, with the name sent by the RADIUS server during authentication. If there is mismatch: • If selected, the mismatch is ignored. • If not selected, authentication fails.

Item	Description
Enable Fast Reauthentication	Specifies that fast reauthentication is enabled. Fast Reauthentication is useful when SIM authentication happens frequently. The encryption keys that are derived from full authentication are reused. As a result, the SIM algorithm is not required to run for every authentication attempt, and the number of network operations that result from frequent authentication attempts is reduced.

Additional resources

For additional information about authenticated wireless settings in Group Policy, see Managing the New Wireless Network (IEEE 802.11) Policies Settings

For additional information about authenticated wired settings in Group Policy, see Managing the New Wired Network (IEEE 802.3) Policies Settings

For information about advanced settings for authenticated wired access and authenticated wireless access, see Advanced Security Settings for Wired and Wireless Network Policies.

Extensible Authentication Protocol (EAP)

High-Speed Networking

High-speed networking features include TCP Chimney Offload, Virtual Machine Chimney, Virtual Machine Queue, receive-side scaling, and Network Direct Memory Access (NetDMA). These features can improve network and operating system performance.

In Windows Server 2012, high-speed networking is unchanged from high-speed networking in Windows Server 2008 R2.

For more information, see High-Speed Networking Features in the Windows Server 2008 and Windows Server 2008 R2 Technical Library.

Hyper-V Network Virtualization Overview

In Windows Server 2012 R2 and System Center 2012 R2 Virtual Machine Manager, Microsoft provides an end-to-end network virtualization solution. There are four major components that comprise Microsoft's network virtualization solution:

- **Windows Azure Pack for Windows Server** provides a tenant facing portal to create virtual networks.

- **System Center 2012 R2 Virtual Machine Manager** (VMM) provides centralized management of the virtual networks.
- **Hyper-V Network Virtualization** provides the infrastructure needed to virtualize network traffic.
- **Hyper-V Network Virtualization gateways** provides connections between virtual and physical networks.

For more information about a network virtualization solution that uses all of these components, see **Hybrid Cloud Multi-Tenant Networking Solution Guide**.

This topic introduces concepts and explains the key benefits and capabilities of Hyper-V Network Virtualization (one part of the overall network virtualization solution) in Windows Server 2012 R2. It explains how network virtualization benefits both private clouds looking for enterprise workload consolidation and public cloud service providers of Infrastructure as a Service (IaaS).

For more technical details about networking virtualization, see Hyper-V Network Virtualization technical details.

Did you mean…

- **Hyper-V FAQ**
- Hyper-V Virtual Switch Overview

Feature description

Hyper-V Network Virtualization provides “virtual networks” (called a VM network) to virtual machines similar to how server virtualization (hypervisor) provides “virtual machines” to the operating system. Network virtualization decouples virtual networks from the physical network infrastructure and removes the constraints of VLAN and hierarchical IP address assignment from virtual machine provisioning. This flexibility makes it easy for customers to move to IaaS clouds and efficient for hosters and datacenter administrators to manage their infrastructure, while maintaining the necessary multi-tenant isolation, security requirements, and supporting overlapping Virtual Machine IP addresses.

Customers want to seamlessly extend their datacenters to the cloud. Today there are technical challenges in making such seamless hybrid cloud architectures. One of the biggest hurdles customers face is reusing their existing network topologies (subnets, IP addresses, network services, and so on.) in the cloud and bridging between their on-premise resources and their cloud resources. Hyper-V Network Virtualization provides the concept of a VM Network that is independent of the underlying physical network. With this concept of a VM Network, composed of one or more Virtual Subnets, the exact location in the physical network of virtual machines attached to a virtual network is decoupled from the virtual network topology. As a result, customers can easily move their virtual subnets to the cloud while preserving their existing IP addresses and topology in the cloud so that existing services continue to work unaware of the physical location of the subnets. That is, Hyper-V Network Virtualization enables a seamless hybrid cloud.

In addition to hybrid cloud, many organizations are consolidating their datacenters and creating private clouds to internally get the efficiency and scalability benefit of cloud architectures. Hyper-V

Network Virtualization allows better flexibility and efficiency for private clouds by decoupling a business unit's network topology (by making it virtual) from the actual physical network topology. In this way, the business units can easily share an internal private cloud while being isolated from each other and continue to keep existing network topologies. The datacenter operations team has flexibility to deploy and dynamically move workloads anywhere in the datacenter without server interruptions providing better operational efficiencies and an overall more effective datacenter.

For workload owners, the key benefit is that they can now move their workload "topologies" to the cloud without changing their IP addresses or re-writing their applications. For example, the typical three-tier LOB application is composed of a front end tier, a business logic tier, and a database tier. Through policy, Hyper-V Network Virtualization allows customer onboarding all or parts of the three tiers to the cloud, while keeping the routing topology and the IP addresses of the services (i.e. virtual machine IP addresses), without requiring the applications to be changed.

For infrastructure owners, the additional flexibility in virtual machine placement makes it possible to move workloads anywhere in the datacenters without changing the virtual machines or reconfiguring the networks. For example Hyper-V Network Virtualization enables cross subnet live migration so that a virtual machine can live migrate anywhere in the datacenter without a service disruption. Previously live migration was limited to the same subnet restricting where virtual machines could be located. Cross subnet live migration allows administrators to consolidate workloads based on dynamic resource requirements, energy efficiency, and can also accommodate infrastructure maintenance without disrupting customer workload up time.

Practical applications

With the success of virtualized datacenters, IT organizations and hosting providers (providers who offer colocation or physical server rentals) have begun offering more flexible virtualized infrastructures that make it easier to offer on-demand server instances to their customers. This new class of service is referred to as Infrastructure as a Service (IaaS). Windows Server 2012 R2 provides all the required platform capabilities to enable enterprise customers to build private clouds and transition to an IT as a service operational model. Windows Server 2012 R2 also enables hosters to build public clouds and offer IaaS solutions to their customers. When combined with Virtual Machine Manager to manage Hyper-V Network Virtualization policy, Microsoft provides a powerful cloud solution.

Windows Server 2012 R2 Hyper-V Network Virtualization provides policy-based, software-controlled network virtualization that reduces the management overhead faced by enterprises when they expand dedicated IaaS clouds, and it provides cloud hosters better flexibility and scalability for managing virtual machines to achieve higher resource utilization.

An IaaS scenario that has virtual machines from different organizational divisions (dedicated cloud) or different customers (hosted cloud) requires secure isolation. Today's solution, virtual local area networks (VLANs), can present significant disadvantages in this scenario.

VLANs Currently, VLANs are the mechanism that most organizations use to support address space reuse and tenant isolation. A VLAN uses explicit tagging (VLAN ID) in the Ethernet frame

headers, and it relies on Ethernet switches to enforce isolation and restrict traffic to network nodes with the same VLAN ID. The main disadvantages with VLANs are as follows:

- Increased risk of an inadvertent outage due to cumbersome reconfiguration of production switches whenever virtual machines or isolation boundaries move in the dynamic datacenter.
- Limited in scalability because there is a maximum of 4094 VLANs and typical switches support no more than 1000 VLAN IDs.
- Constrained within a single IP subnet, which limits the number of nodes within a single VLAN and restricts the placement of virtual machines based on physical locations. Even though VLANs can be expanded across sites, the entire VLAN must be on the same subnet.

IP address assignment In addition to the disadvantages that are presented by VLANs, virtual machine IP address assignment presents issues, which include:

- Physical locations in datacenter network infrastructure determine virtual machine IP addresses. As a result, moving to the cloud typically requires changing IP addresses of the service workloads.
- Policies are tied to IP addresses, such as firewall rules, resource discovery and directory services, and so on. Changing IP addresses requires updating all the associated policies.
- Virtual machine deployment and traffic isolation are dependent on the topology.

When datacenter network administrators plan the physical layout of the datacenter, they must make decisions about where subnets will be physically placed and routed. These decisions are based on IP and Ethernet technology that influence the potential IP addresses that are allowed for virtual machines running on a given server or a blade that is connected to a particular rack in the datacenter. When a virtual machine is provisioned and placed in the datacenter, it must adhere to these choices and restrictions regarding the IP address. Therefore, the typical result is that the datacenter administrators assign new IP addresses to the virtual machines.

The problem with this requirement is that in addition to being an address, there is semantic information associated with an IP address. For instance, one subnet may contain given services or be in a distinct physical location. Firewall rules, access control policies, and IPsec security associations are commonly associated with IP addresses. Changing IP addresses forces the virtual machine owners to adjust all their policies that were based on the original IP address. This renumbering overhead is so high that many enterprises choose to deploy only new services to the cloud, leaving legacy applications alone.

Hyper-V Network Virtualization decouples virtual networks for customer virtual machines from the physical network infrastructure. As a result, it enables customer virtual machines to maintain their original IP addresses, while allowing datacenter administrators to provision customer virtual machines anywhere in the datacenter without reconfiguring physical IP addresses or VLAN IDs. The next section summarizes the key functionality.

Important functionality

The following is a list of the key functionality, benefits, and capabilities of Hyper-V Network Virtualization in Windows Server 2012 R2:

- **Enables flexible workload placement – Network isolation and IP address re-use without VLANs**

  Hyper-V Network Virtualization decouples the customer's virtual networks from the physical network infrastructure of the hosters, providing freedom for workload placements inside the datacenters. Virtual machine workload placement is no longer limited by the IP address assignment or VLAN isolation requirements of the physical network because it is enforced within Hyper-V hosts based on software-defined, multitenant virtualization policies.

  Virtual machines from different customers with overlapping IP addresses can now be deployed on the same host server without requiring cumbersome VLAN configuration or violating the IP address hierarchy. This can streamline the migration of customer workloads into shared IaaS hosting providers, allowing customers to move those workloads without modification, which includes leaving the virtual machine IP addresses unchanged. For the hosting provider, supporting numerous customers who want to extend their existing network address space to the shared IaaS datacenter is a complex exercise of configuring and maintaining isolated VLANs for each customer to ensure the coexistence of potentially overlapping address spaces. With Hyper-V Network Virtualization, supporting overlapping addresses is made easier and requires less network reconfiguration by the hosting provider.

  In addition, physical infrastructure maintenance and upgrades can be done without causing a down time of customer workloads. With Hyper-V Network Virtualization, virtual machines on a specific host, rack, subnet, VLAN, or entire cluster can be migrated without requiring a physical IP address change or major reconfiguration.

- **Enables easier moves for workloads to a shared IaaS cloud**

  With Hyper-V Network Virtualization, IP addresses and virtual machine configurations remain unchanged. This enables IT organizations to more easily move workloads from their datacenters to a shared IaaS hosting provider with minimal reconfiguration of the workload or their infrastructure tools and policies. In cases where there is connectivity between two datacenters, IT administrators can continue to use their tools without reconfiguring them.

- **Enables live migration across subnets**

  Live migration of virtual machine workloads traditionally has been limited to the same IP subnet or VLAN because crossing subnets required the virtual machine's guest operating system to change its IP address. This address change breaks existing communication and disrupts the services running on the virtual machine. With Hyper-V Network Virtualization, workloads can be live migrated from servers running Windows Server 2012 in one subnet to servers running Windows Server 2012 in a different subnet without changing the workload IP addresses. Hyper-V Network Virtualization ensures that virtual machine location changes due to live migration are updated and synchronized among hosts that have ongoing communication with the migrated virtual machine.

- **Enables easier management of decoupled server and network administration**

  Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, and network administrators can focus on overall network infrastructure and traffic management. This enables datacenter server administrators

to deploy and migrate virtual machines without changing the IP addresses of the virtual machines. There is reduced overhead because Hyper-V Network Virtualization allows virtual machine placement to occur independently of network topology, reducing the need for network administrators to be involved with placements that might change the isolation boundaries.

- **Simplifies the network and improves server/network resource utilization**

  The rigidity of VLANs and the dependency of virtual machine placement on a physical network infrastructure results in overprovisioning and underutilization. By breaking the dependency, the increased flexibility of virtual machine workload placement can simplify the network management and improve server and network resource utilization. Note that Hyper-V Network Virtualization supports VLANs in the context of the physical datacenter. For example, a datacenter may want all Hyper-V Network Virtualization traffic to be on a specific VLAN.

- **Is compatible with existing infrastructure and emerging technology**

  Hyper-V Network Virtualization can be deployed in today's datacenter, yet it is compatible with emerging datacenter "flat network" technologies.

- **Provides for interoperability and ecosystem readiness**

  Hyper-V Network Virtualization supports multiple configurations for communication with existing resources, such as cross premise connectivity, storage area network (SAN), non-virtualized resource access, and so on. Microsoft is committed to working with ecosystem partners to support and enhance the experience of Hyper-V Network Virtualization in terms of performance, scalability, and manageability.

- **Uses Windows PowerShell and WMI**

  Hyper-V Network Virtualization supports Windows PowerShell and Windows Management Instrumentation (WMI) for configuring the network virtualization and isolation policies. The Windows PowerShell cmdlets for Hyper-V Network Virtualization enable administrators to build command-line tools or automated scripts to configure, monitor, and troubleshoot network isolation policies.

Software requirements

Hyper-V Network Virtualization requires Windows Server 2012 or Windows Server 2012 R2 and the Hyper-V role.

See also

To learn more about Hyper-V Network Virtualization see the following links:

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2

Content type	References
	• How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2)
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide • Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	• Windows Server Gateway • Hyper-V Virtual Switch Overview

What's New in Hyper-V Network Virtualization in Windows Server 2012 R2

This topic describes the new or changed features and functionality for Hyper-V Network Virtualization (HNV) in Windows Server 2012 R2.

Feature description

HNV provides a virtual network abstraction on top of a physical network. This abstraction provides virtual networks to virtual machines which operate the same as a physical network. This is similar to the abstraction hypervisors provide to the operating system running in virtual machines. HNV provides this abstraction through an overlay network on top of the physical network for each VM network. For a HNV overview, see Hyper-V Network Virtualization Overview.

New and changed functionality

The following HNV features have been added or updated for Windows Server 2012 R2.

Feature/functionality	Windows Server 2012	Windows Server 2012 R2
Inbox HNV Gateway		X
HNV interoperability with Hyper-V Virtual Switch Extensions	X	X
HNV VM Network Diagnostics	X	X
HNV Architecture	X	X
Dynamic IP Address Learning		X
HNV + Windows NIC Teaming	X	X
NVGRE Encapsulated Task Offload	X	X

Inbox HNV Gateway

The inbox HNV gateway is a multi-tenant gateway that performs Site-to-Site (VPN), NAT, and Forwarding functions.

What value does this change add?

It is now easier to set up a gateway which can connect multiple tenant VPN connections in a hybrid cloud scenario, supports multiple tenants connecting to the Internet, and forwards network traffic from a datacenter network to virtual networks in a private cloud scenario.

What works differently?

- You can use System Center 2012 R2 Virtual Machine Manager to fully manage the HNV gateway.
- Supports guest clustering for high availability
- Includes BGP for dynamic routes update

HNV Architecture

The HNV filter moved from being an NDIS lightweight filter (LWF) to being part of the Hyper-V virtual switch.

What value does this change add?

Forwarding switch extensions can co-exist with HNV allowing multiple network virtualization solutions (one provided by HNV and another provided by the forwarding switch extension) to co-exist on the same host running Hyper-V.

What works differently?

- Improved interoperability with switch extensions
- The HNV NDIS LWF does not have to be bound to network adapters anymore. Once you attach a network adapter to the virtual switch you can enable HNV simply by assigning a Virtual Subnet ID to a particular virtual network adapter. For those using Virtual Machine Manager to manage VM networks this is transparent, but for anyone using Windows PowerShell this will save an often-missed step.

HNV interoperability with Hyper-V Virtual Switch Extensions

Switch extensions work in both the Customer Address (CA) and Provider Address (PA) space.

What value does this change add?

Third-party switch extensions can now work with HNV traffic, because the extensions have visibility into both the provider address (PA) space, and the customer address (CA) space.

What works differently?

- The HNV module was moved to inside the virtual switch so that extensions can see both the provider (PA) and virtual (CA) IP address spaces. This allows forwarding and other types of extensions to make decisions with knowledge of both address spaces.
- Hybrid forwarding is implemented. Hybrid forwarding directs packets to different forwarding agents, based upon the packet type. In the Windows Server 2012 R2 implementation, an NVGRE packet is forwarded by the HNV module. A packet that is not NVGRE is forwarded normally by the forwarding extension.

HNV VM Network Diagnostics

Some new diagnostic tools have been included.

What value does this change add?

This enhances your ability to diagnose HNV networks.

What works differently?

- Enhanced ping.exe (ping –p) to allow pinging to and from provider addresses
- Two new Windows PowerShell cmdlets (**Test-VMNetworkAdapter** and **Select-NetVirtualizationNextHop**) that enables diagnostics of HNV policy and the Customer Address space.
- Added the ability for Message Analyzer to decode NVGRE packets

For more details, check out the New Networking Diagnostics with PowerShell in Windows Server R2 blog post.

Dynamic IP Address Learning

HNV learns about the IP address of a virtual machine that has been assigned manually or set via DHCP on the virtual network.

What value does this change add?

- Enables high availability scenarios for both virtual machines on a VM network and the HNV gateway.
- Allows you to run DHCP, DNS, and Active Directory in your VM networks.

What works differently?

- For broadcast or multicast packets in a VM network, a PA multicast IP address is used if configured. If a PA multicast address is not available an intelligent PA unicast replication is used.

  Packets are unicasted only to PA addresses that are configured for the particular virtual subnet the packet is on. In addition, only one unicast packet is sent per host no matter how many relevant virtual machines are on the host.
- Once a host learns a new IP address it will then notify Virtual Machine Manager. At this point, the learned IP address becomes part of the centralized policy that Virtual Machine Manager pushes out. This allows for both rapid dissemination of HNV routing policy and limits the network overhead for disseminating this HNV routing policy.
- Includes support for Duplicate Address Detection (DAD), network unreachability detection (NUD) and Address Resolution Protocol (ARP) packets in the CA address space for both IPv4 and IPv6. The HNV filter also provides a reliable ARP proxy for any known routing policies once again reducing the amount of control traffic that goes out on the physical network.
- Fully supported in Windows PowerShell.

HNV + Windows NIC Teaming

Integrates HNV and Windows NIC Teaming to allow multiple network adapters to be placed into a team for the purposes of bandwidth aggregation, and/or traffic failover to maintain connectivity in the event of a network component failure.

What value does this change add?

Integrating HNV with Windows NIC Teaming increases HNV network throughput and reliability.

What works differently?

Both inbound and outbound spread of virtualized traffic on a NIC Team is enabled. This means that both traffic leaving a host or coming into a host can utilize all the network adapters in the NIC team.

NVGRE Encapsulated Task Offload

In NDIS 6.30 and later (Windows Server 2012 and later), Network Virtualization using Generic Routing Encapsulation (NVGRE) task offload makes it possible to use Generic Routing Encapsulation (GRE)-encapsulated packets with:

- Large Send Offload (LSO)
- Receive Side Scaling (RSS)
- Virtual Machine Queue (VMQ)

What value does this change add?

To increase its performance, NVGRE can offload tasks to a network adapter that has the appropriate task offload capabilities.

What works differently?

Two partners have announced that their next generation network adapters will support NVGRE Encapsulated Task Offload. You can read the press releases from Mellanox and Emulex for more details.

Microsoft is continuing to work with additional network vendors to enable NVGRE Task Offload. More announcements will be made in the future.

See also

The following table provides additional resources for evaluating HNV.

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2 • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2) • Test Lab Guide: Windows Server 2012 R2 Hyper-V Network Virtualization with System Center 2012 R2 VMM
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide • Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	Hyper-V Virtual Switch Overview

Hyper-V Network Virtualization technical details

Server virtualization enables multiple server instances to run concurrently on a single physical host; yet server instances are isolated from each other. Each virtual machine essentially operates as if it is the only server running on the physical computer. Network virtualization provides a similar capability, in which multiple virtual network infrastructures run on the same physical network (potentially with overlapping IP addresses), and each virtual network infrastructure operates as if it is the only virtual network running on the shared network infrastructure. Figure 1 shows this relationship.

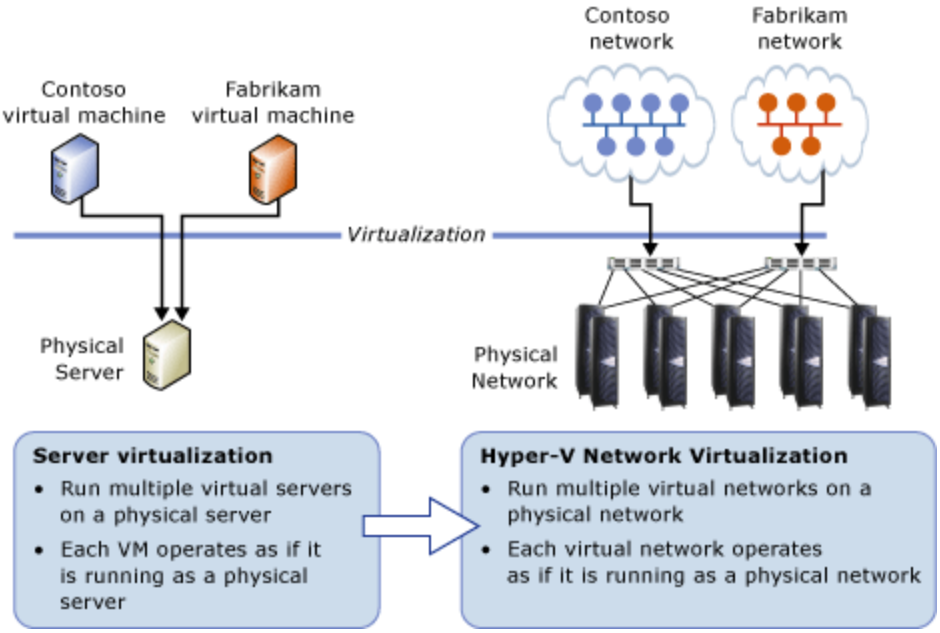

Figure 1: Server virtualization versus network virtualization

Hyper-V Network Virtualization Concepts

In Hyper-V Network Virtualization (HNV), a customer is defined as the “owner” of a group of virtual machines that are deployed in a datacenter. A customer can be a corporation or enterprise in a multitenant public datacenter, or a division or business unit within a private datacenter. Each customer can have one or more VM networks in the datacenter, and each VM network consists of one or more virtual subnets.

VM network

- Each VM network consists of one or more virtual subnets. A VM network forms an isolation boundary where the virtual machines within a VM network can communicate with each other. As a result, virtual subnets in the same VM network must not use overlapping IP address prefixes.

- Each VM network has a Routing Domain which identifies the VM network. The Routing Domain ID (RDID), which identifies the VM network, is assigned by datacenter administrators or datacenter management software, such as System Center 2012 R2 Virtual Machine Manager (VMM). The RDID is a Windows GUID — for example, “{11111111-2222-3333-4444-000000000000}”.

Virtual subnets

- A virtual subnet implements the Layer 3 IP subnet semantics for the virtual machines in the same virtual subnet. The virtual subnet is a broadcast domain (similar to a VLAN). Virtual machines in the same virtual subnet must use the same IP prefix.
- Each virtual subnet belongs to a single VM network (RDID), and it is assigned a unique Virtual Subnet ID (VSID). The VSID must be unique within the datacenter and is in the range 4096 to 2^24-2).

A key advantage of the VM network and routing domain is that it allows customers to bring their network topologies to the cloud. Figure 2 shows an example where the Contoso Corp has two separate networks, the R&D Net and the Sales Net. Because these networks have different routing domain IDs, they cannot interact with each other. That is, Contoso R&D Net is isolated from Contoso Sales Net even though both are owned by Contoso Corp. Contoso R&D Net contains three virtual subnets. Note that both the RDID and VSID are unique within a datacenter.

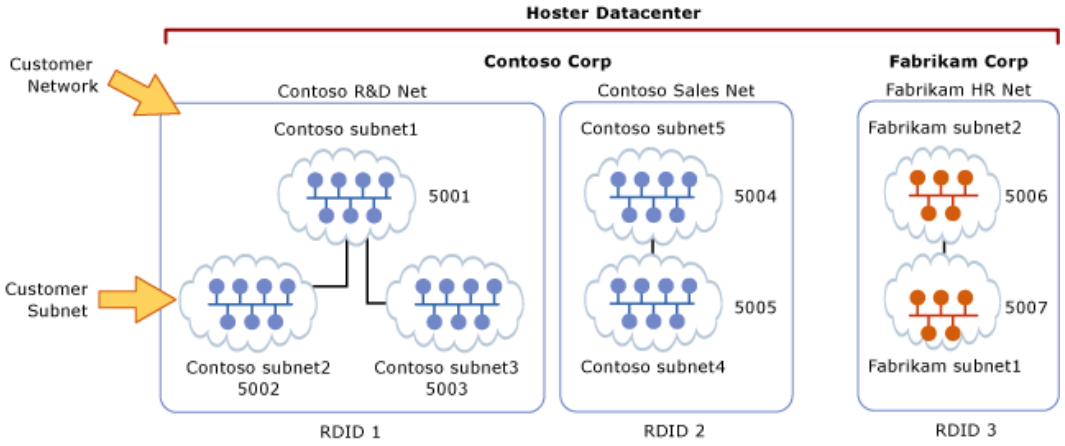

Figure 2: Customer networks and virtual subnets

In Figure 2, the virtual machines with VSID 5001 can have their packets routed or forwarded by HNV to virtual machines with VSID 5002 or VSID 5003. Before delivering the packet to the Hyper-V switch, HNV will update the VSID of the incoming packet to the VSID of the destination virtual machine. This will only happen if both VSIDs are in the same RDID. If the VSID that is associated with the packet does not match the VSID of the destination virtual machine, the packet will be dropped. Therefore, virtual network adapters with RDID1 cannot send packets to virtual network adapters with RDID2.

Note

In the packet flow description above, the term “virtual machine” actually means the “virtual network adapter” on the virtual machine. The common case is that a virtual

machine only has a single virtual network adapter. In this case, the words virtual machine and virtual network adapter can conceptually mean the same thing. Because a virtual machine can have multiple virtual network adapters, and these virtual network adapters can have different VirtualSubnetIDs (VSIDs) or RoutingDomainIDs (RDIDs), HNV specifically focuses on the packets sent and received between virtual network adapters.

Each virtual subnet defines a Layer 3 IP subnet and a Layer 2 (L2) broadcast domain boundary similar to a VLAN. When a virtual machine broadcasts a packet, this broadcast is limited to the virtual machines that are attached to switch ports with the same VSID. Each VSID can be associated with a multicast address in the PA. All broadcast traffic for a VSID is sent on this multicast address.

Note

HNV does NOT depend on broadcast or multicast. For broadcast or multicast packets in a VM network, a PA multicast IP address is used if configured. However, the many datacenter operators do not enable multicast in their environments. As a result, when a PA multicast address is not available an intelligent PA unicast replication is used. This means that packets are unicasted only to PA addresses that are configured for the particular virtual subnet the packet is on. In addition, only one unicast packet per host is sent no matter how many relevant virtual machines are on the host.

In addition to being a broadcast domain, the VSID provides isolation. A virtual network adapter in HNV is connected to a Hyper-V switch port that has a VSID ACL. If a packet arrives on this Hyper-V switch port with a different VSID the packet is dropped. Packets will only be delivered on a Hyper-V switch port if the VSID of the packet matches the VSID of the switch port. This is the reason, in the above example of Figure 2, that packets flowing from VSID 5001 to 5003 must have the VSID in the packet modified before delivery to the destination virtual machine.

If the Hyper-V switch port does not have a VSID ACL, the virtual network adapter that is attached to that switch port is not part of a HNV virtual subnet. Packets sent from a virtual network adapter that does not have a VSID ACL will pass unmodified through the Hyper-V switch.

When a virtual machine sends a packet, the VSID of the Hyper-V switch port is associated with this packet. On the receiving side, HNV delivers to the Hyper-V switch the VSID in the OOB along with the decapsulated packet. On the receiving end, HNV performs a policy lookup and adds the VSID to the OOB data before the packet is passed to the Hyper-V switch.

Note

Hyper-V Switch Extensions can operate in both the Provider Address (PA) space and the Customer Address (CA) space. This means the VSID is available to the switch extensions. This allows the switch extension to become multitenant aware. For example, a firewall switch extension can differentiate CA IP address 10.1.1.5 with OOB containing VSID 5001 from the same CA IP address with VSID 6001.

Routing in Hyper-V Network Virtualization

As in physical networks, routing is an important part of HNV. There are two key aspects to understand: how packets are routed between virtual subnets and how packets are routed outside a virtual network.

Routing Between Virtual Subnets

In a physical network, a subnet is the Layer 2 (L2) domain where computers (virtual and physical) can directly communicate with each other without having to be routed. In Windows, if you statically configure a network adapter you can set a "default gateway" which is the IP address to send all traffic that is going out of the particular subnet so that it can be routed appropriately. This is typically the router for your physical network. HNV uses a built in router that is part of every host to form a distributed router for a virtual network. This means that every host, in particular the Hyper-V Virtual Switch, acts as the default gateway for all traffic that is going between Virtual Subnets that are part of the same VM network. In Windows Server 2012 and Windows Server 2012 R2 the address used as the default gateway is the lowest entry for the subnet (as an example, it is the ".1" address for a /24 subnet prefix). This address is reserved in each virtual subnet for the default gateway and cannot be used by virtual machines in the virtual subnet.

HNV acting as a distributed router allows for a very efficient way for all traffic inside a VM Network to be routed appropriately because each host can directly route the traffic to the appropriate host without needing an intermediary. This is particularly true when two virtual machines in the same VM Network but different Virtual Subnets are on the same physical host. As you will see later in this section, the packet never has to leave the physical host.

Routing Outside a Virtual Network

Most customer deployments will require communication from the HNV environment to resources that are not part of the HNV environment. Network Virtualization gateways are required to allow communication between the two environments. Scenarios requiring a HNV Gateway include Private Cloud and Hybrid Cloud. Basically, HNV gateways are required for VPNs and routing.

Gateways can come in different physical form factors. They can be built on Windows Server 2012 R2, incorporated into a Top of Rack (TOR) switch, a load balancer, put into other existing network appliances, or can be a new stand-alone network appliance.

Note

For more information about gateways, see Hyper-V Network Virtualization Gateway Architectural Guide and Simple demo with gateway script.

Private Cloud (Routing)

Large enterprises may be either hesitant, or for compliance reasons, unable to move some of their services and data to a public cloud hoster. However, enterprises still want to obtain the benefits of the cloud provided by HNV by consolidating their datacenter resources into a private

cloud. In a private cloud deployment, overlapping IP addresses may not be needed because corporations typically have sufficient internal non-routable address (e.g. 10.x.x.x or 192.x.x.x) space. Consider the example shown in Figure 3.

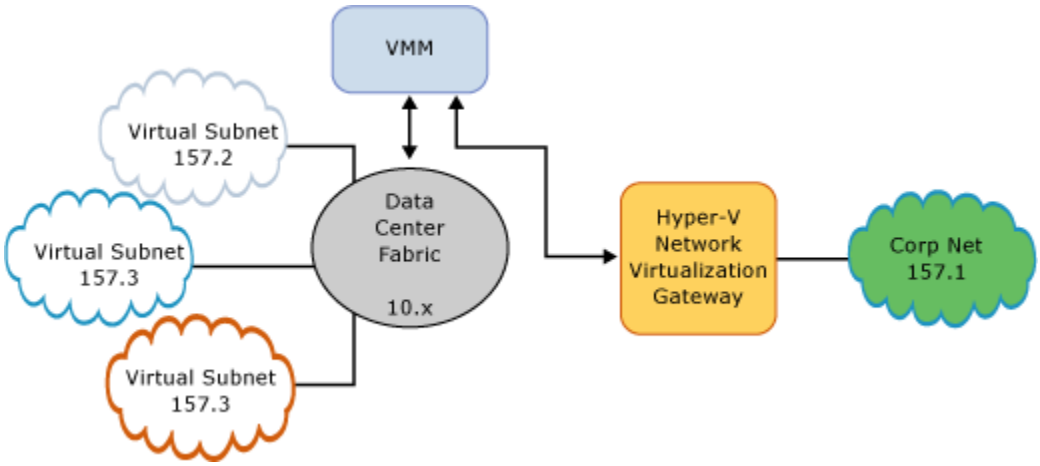

Figure 3: Private Cloud Deployment

Notice in this example that the Customer Addresses in the virtual subnets are 157.x addresses while the IP addresses in the non-network virtualized part of the network (Corp Net) are also 157.x addresses. In this case the PA addresses for the virtual subnets in the datacenter are 10.x IP addresses. This deployment allows the enterprise to take advantage of HNV's ability to offer flexibility in both virtual machine placement and cross-subnet live migration in the datacenter fabric. This increases datacenter efficiency thereby reducing both Operational Expenses (OpEx) and Capital Expenses (CapEx). In this scenario the HNV gateway provides routing between the 10.x and 157.1 IP addresses.

Hybrid Cloud (Site to site VPN)

A key advantage of HNV is that it can seamlessly extend an on-premise datacenter to a Windows Server 2012 based cloud datacenter. This is called a Hybrid Cloud model as shown in Figure 4.

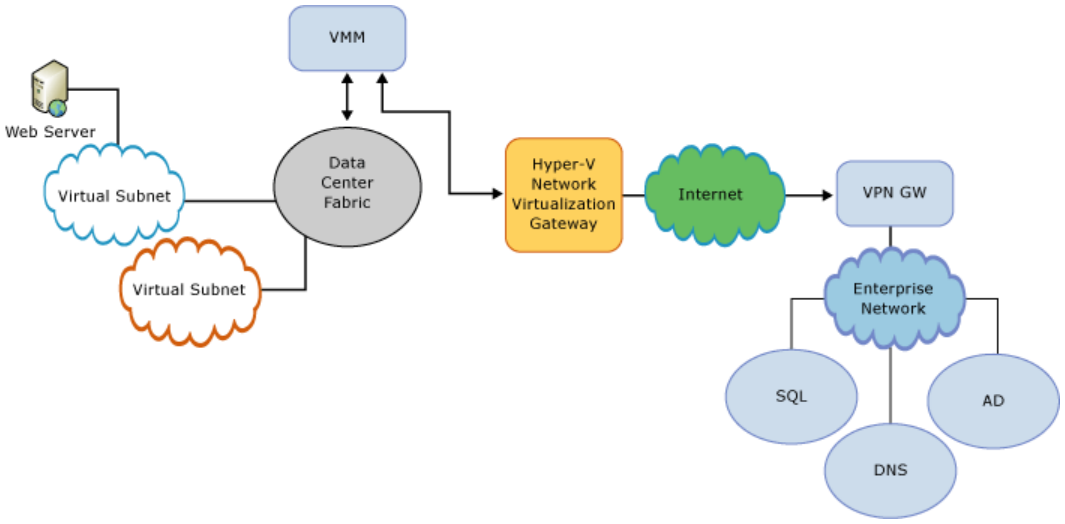

Figure 4: Hybrid Cloud Deployment

In this scenario an internal subnet, such as the subnet containing web servers is moved from the Enterprise Network into a Cloud Hoster's datacenter. Taking advantage of *Bring Your Own IP Address* offered by the Hoster, the Enterprise does not need to change the network configuration of the Web Server virtual machine or any other network endpoint that references that Web Server. The Hoster provides a secure link via a HNV Gateway Appliance. The Enterprise admins need only configure their on-premise VPN with the appropriate IP address. The Web Server virtual machine is unaware that it has been moved to the cloud. It remains domain-joined with Active Directory (AD) and uses the Enterprise's DNS server. The Web Server virtual machine also continues to interact with other servers in the enterprise such as a SQL Server.

The HNV Gateway can support multiple site to site (S2S) VPN tunnels as shown in Figure 5. Note that a VMM is not pictured in the diagram but is required for HNV deployments.

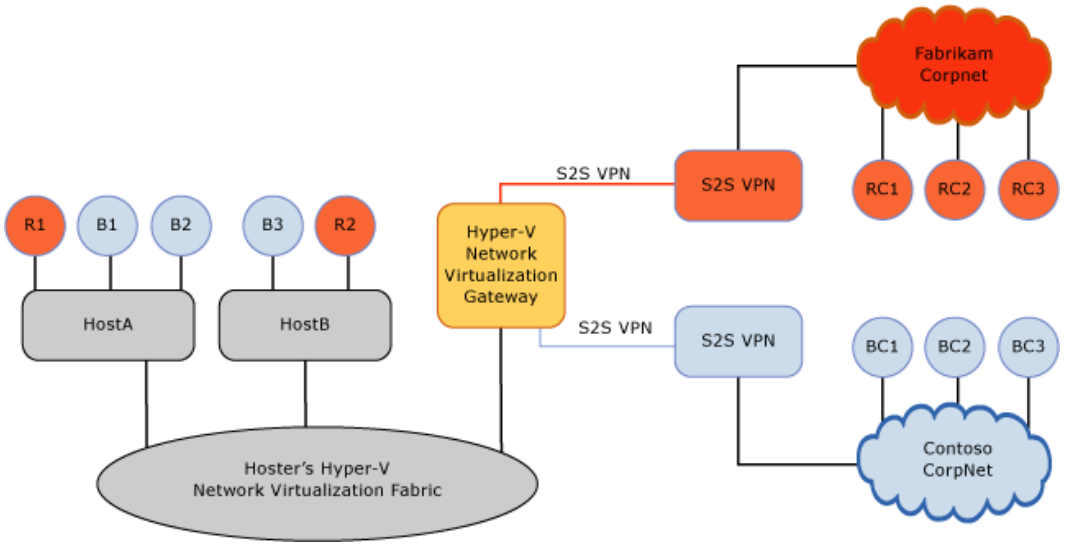

Figure 5: HNV Gateway

Packet Encapsulation

Each virtual network adapter in HNV is associated with two IP addresses:

- **Customer Address** (CA) The IP address that is assigned by the customer, based on their intranet infrastructure. This address enables the customer to exchange network traffic with the virtual machine as if it had not been moved to a public or private cloud. The CA is visible to the virtual machine and reachable by the customer.
- **Provider Address** (PA) The IP address that is assigned by the hoster or the datacenter administrators based on their physical network infrastructure. The PA appears in the packets on the network that are exchanged with the server running Hyper-V that is hosting the virtual machine. The PA is visible on the physical network, but not to the virtual machine.

The CAs maintain the customer's network topology, which is virtualized and decoupled from the actual underlying physical network topology and addresses, as implemented by the PAs. The following diagram shows the conceptual relationship between virtual machine CAs and network infrastructure PAs as a result of network virtualization.

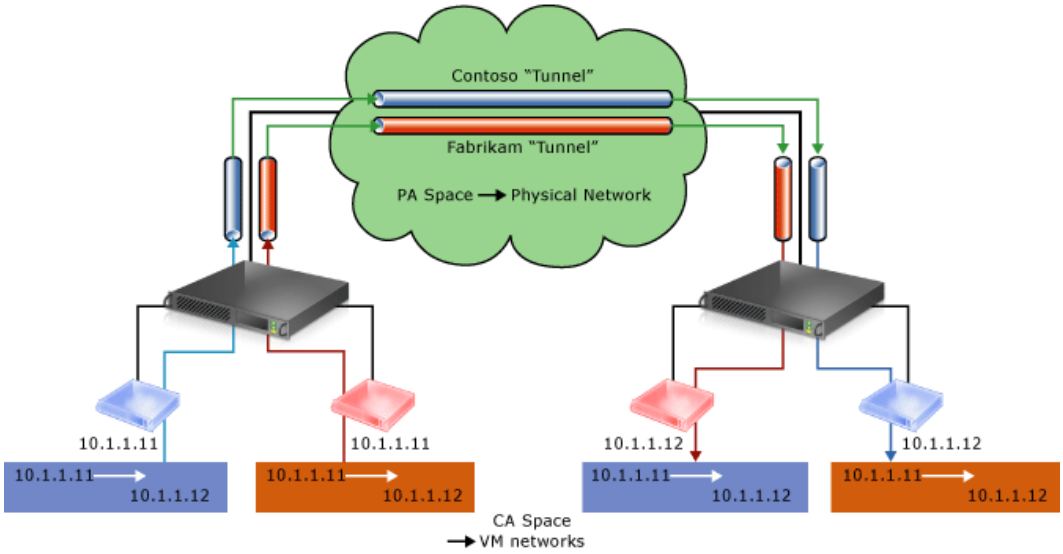

Figure 6: Conceptual diagram of network virtualization over physical infrastructure

In the diagram, customer virtual machines are sending data packets in the CA space, which traverse the physical network infrastructure through their own virtual networks, or “tunnels”. In the example above, the tunnels can be thought of as “envelopes” around the Contoso and Fabrikam data packets with green shipping labels (PA addresses) to be delivered from the source host on the left to the destination host on the right. The key is how the hosts determine the “shipping addresses” (PA’s) corresponding to the Contoso and the Fabrikam CA’s, how the “envelope” is put around the packets, and how the destination hosts can unwrap the packets and deliver to the Contoso and Fabrikam destination virtual machines correctly.

This simple analogy highlighted the key aspects of network virtualization:

- Each virtual machine CA is mapped to a physical host PA. There can be multiple CAs associated with the same PA.

- Virtual machines send data packets in the CA spaces, which are put into an “envelope” with a PA source and destination pair based on the mapping.
- The CA-PA mappings must allow the hosts to differentiate packets for different customer virtual machines.

As a result, the mechanism to virtualize the network is to virtualize the network addresses used by the virtual machines. The next section describes the actual mechanism of address virtualization.

Network virtualization through address virtualization

HNV supports Network Virtualization for Generic Routing Encapsulation (NVGRE) as the mechanism to virtualize the IP Address:

Generic Routing Encapsulation This network virtualization mechanism uses the Generic Routing Encapsulation (NVGRE) as part of the tunnel header. In NVGRE, the virtual machine’s packet is encapsulated inside another packet. The header of this new packet has the appropriate source and destination PA IP addresses in addition to the Virtual Subnet ID, which is stored in the Key field of the GRE header, as shown in Figure 7.

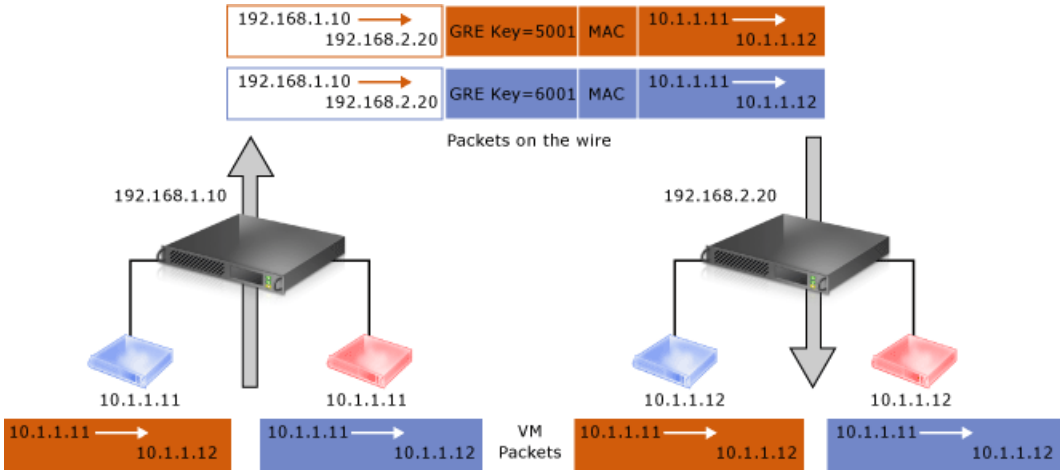

Figure 7: Network virtualization - NVGRE encapsulation

The Virtual Subnet ID allows hosts to identify the customer virtual machine for any given packet, even though the PA’s and the CA’s on the packets may overlap. This allows all virtual machines on the same host to share a single PA, as shown in Figure 7.

Sharing the PA has a big impact on network scalability. The number of IP and MAC addresses that need to be learned by the network infrastructure can be substantially reduced. For instance, if every end host has an average of 30 virtual machines, the number of IP and MAC addresses that need to be learned by the networking infrastructure is reduced by a factor of 30.The embedded Virtual Subnet IDs in the packets also enable easy correlation of packets to the actual customers.

With Windows Server 2012 and later, HNV fully supports NVGRE out of the box; it does NOT require upgrading or purchasing new network hardware such as NICs (Network Adapters), switches, or routers. This is because the NVGRE packet on the wire is a regular IP packet in the PA space, which is compatible with today's network infrastructure.

Windows Server 2012 made working with standards a high priority. Along with key industry partners (Arista, Broadcom, Dell, Emulex, Hewlett Packard, and Intel) Microsoft published a draft RFC that describes the use of Generic Routing Encapsulation (GRE), which is an existing IETF standard, as an encapsulation protocol for network virtualization. For more information, see the following Internet Draft: Network Virtualization using Generic Routing Encapsulation. As NVGRE-aware becomes commercially available the benefits of NVGRE will become even greater.

Multitenant deployment example

The following diagram shows an example deployment of two customers moving in a cloud datacenter with the CA-PA relationship defined by the HNV policies.

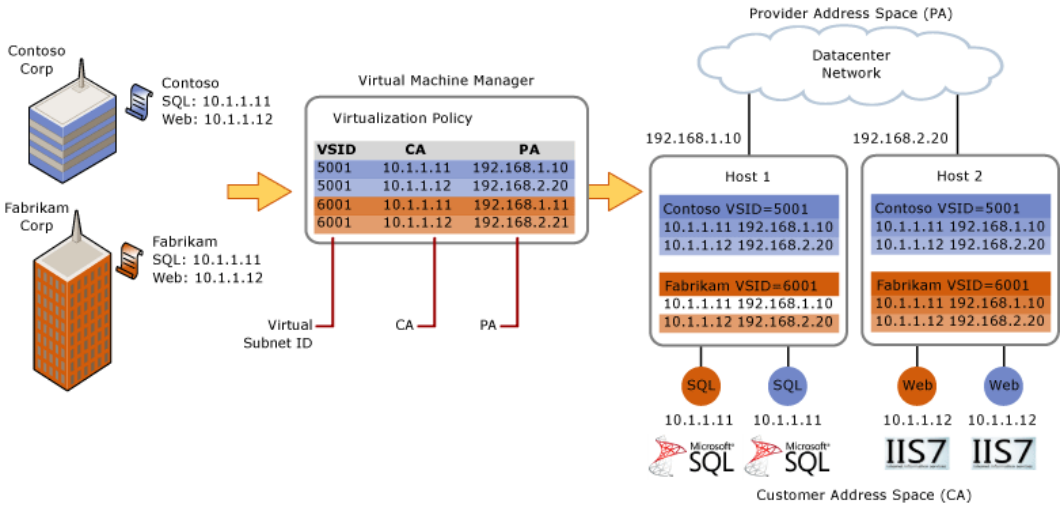

Figure 8: Multi-tenant deployment example

Consider the example in Figure 8. Prior to moving to the hosting provider's shared IaaS service:

- Contoso Corp ran a SQL Server (named **SQL**) at the IP address 10.1.1.11 and a web server (named **Web**) at the IP address 10.1.1.12, which uses its SQL Server for database transactions.
- Fabrikam Corp ran a SQL Server, also named **SQL** and assigned the IP address 10.1.1.11, and a web server, also named **Web** and also at the IP address 10.1.1.12, that uses its SQL Server for database transactions.

Contoso Corp and Fabrikam Corp move their respective SQL Servers and web servers to the same hosting provider's shared IaaS service where, coincidentally, they run the **SQL** virtual machines on Hyper-V Host 1 and the **Web** (IIS7) virtual machines on Hyper-V Host 2. All virtual machines maintain their original intranet IP addresses (their CAs).

Both companies are assigned the following Virtual Subnet ID (VSID) and PAs by their hosting provider when the virtual machines are provisioned:

- PAs of Contoso Corp's virtual machines: **VSID** is 5001, **SQL** is 192.168.1.10, **Web** is 192.168.2.20
- PAs of Fabrikam Corp's virtual machines: **VSID** is 6001, **SQL** is 192.168.1.10, **Web** is 192.168.2.20

The hosting provider creates policy settings, consisting of a customer virtual subnet for Fabrikam Corp that maps the CAs of the Fabrikam Corp virtual machines to their assigned PAs and VSID, and a separate customer virtual subnet for Contoso Corp that maps the CAs of the Contoso Corp virtual machines to their assigned PAs and VSID. The provider applies these policy settings to Hyper-V Host 1 and Hyper-V Host 2.

When the Contoso Corp Web virtual machine on Hyper-V Host 2 queries its SQL Server at 10.1.1.11, the following happens:

Hyper-V Host 2, based on its policy settings, translates the addresses in the packet from:

- Source: 10.1.1.12 (the CA of Contoso Corp **Web**)
- Destination: 10.1.1.11 (the CA of Contoso Corp **SQL**)

The encapsulated packet contains:

- GRE header with VSID: 5001
- Outer source: 192.168.2.20 (the PA for Contoso Corp Web)
- Outer destination: 192.168.1.10 (the PA for Contoso Corp SQL)

When the packet is received at Hyper-V Host 1, based on its policy settings, it will decapsulate the NVGRE packet with:

- Outer source: 192.168.2.20 (the PA for Contoso Corp **Web**)
- Outer destination: 192.168.1.10 (the PA for Contoso Corp **SQL**)
- GRE header with VSID: 5001

The decapsulated packet (the original packet sent from the Contoso Corp **Web** virtual machine) is delivered to the Contoso Corp **SQL** virtual machine:

- Source: 10.1.1.12 (the CA of Contoso Corp **Web**)
- Destination: 10.1.1.11 (the CA of Contoso Corp **SQL**)

When the Contoso Corp **SQL** virtual machine on Hyper-V Host 1 responds to the query, the following happens:

Hyper-V Host 1, based on its policy settings, translates the addresses in the packet from:

- Source: 10.1.1.11 (the CA of Contoso Corp **SQL**)
- Destination 10.1.1.12 (the CA of Contoso Corp **Web**)

The packet is encapsulated with:

- GRE header with VSID: 5001
- Outer source: 192.168.1.10 (the PA for Contoso Corp SQL)
- Outer destination: 192.168.2.20 (the PA for Contoso Corp Web)

When it is received at Hyper-V Host 2, it will, based on its policy settings, decapsulate the packet:

- Source: 192.168.1.10 (the PA for Contoso Corp **SQL**)
- Destination: 192.168.2.20 (the PA for Contoso Corp **Web**)
- GRE header with VSID: 5001

The decapsulated packet is delivered to the Contoso Corp **Web** virtual machine with:

- Source: 10.1.1.11 (the CA of Contoso Corp **SQL**)
- Destination: 10.1.1.12 (the CA of Contoso Corp **Web**)

A similar process for traffic between the Fabrikam Corp **Web** and **SQL** virtual machines uses the HNV policy settings for the Fabrikam Corp. As a result, with HNV, Fabrikam Corp and Contoso Corp virtual machines interact as if they were on their original intranets. They can never interact with each other, even though they are using the same IP addresses.

The separate addresses (CAs and PAs), the policy settings of the Hyper-V hosts, and the address translation between the CA and the PA for inbound and outbound virtual machine traffic isolate these sets of servers. Furthermore, the virtualization mappings and transformation decouples the virtual network architecture from the physical network infrastructure. Although Contoso **SQL** and **Web** and Fabrikam **SQL** and **Web** reside in their own CA IP subnets (10.1.1/24), their physical deployment happens on two hosts in different PA subnets, 192.168.1/24 and 192.168.2/24, respectively. The implication is that cross-subnet virtual machine provisioning and live migration become possible with HNV.

Note
For more information about packet flow, download the Hyper-V Network Virtualization Packet Flow PowerPoint presentation at http://www.microsoft.com/en-us/download/details.aspx?id=34782.

Hyper-V Network Virtualization architecture

In Windows Server 2012, HNV policy enforcement and IP virtualization was performed by the Network Driver Interface Specification (NDIS) Lightweight Filter (LWF) called Windows Network Virtualization (WNV). The WNV filter was located below the Hyper-V switch as shown in Figure 9. A side effect of this architecture was that switch extensions could only see the CA address space traffic, but not for the PA space.

In Windows Server 2012 R2, HNV is now a part of the virtual switch, enabling extensions to gain visibility into both the CA and PA space addresses.

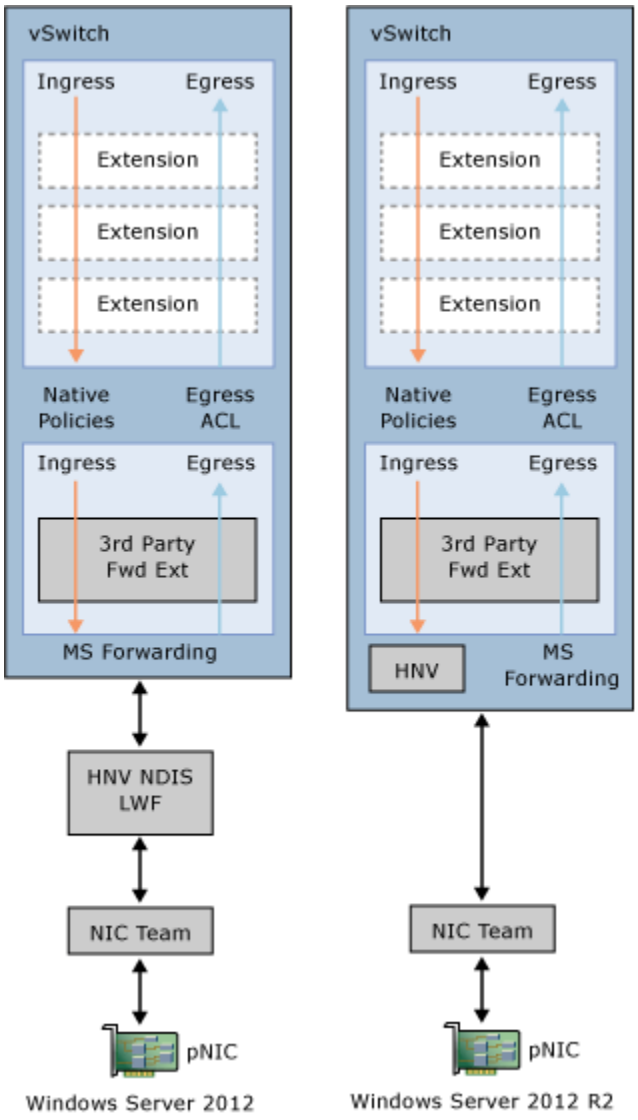

Figure 9: HNV Architecture

Each virtual machine network adapter is configured with an IPv4, and/or an IPv6 address. These are the CAs that will be used by the virtual machines to communicate with each other, and they are carried in the IP packets from the virtual machines. HNV virtualizes the CAs to PAs based on the network virtualization policies.

A virtual machine sends a packet with source address CA1, which is virtualized based on HNV policy in the Hyper-V switch. A special network virtualization access control list based on VSID isolates the virtual machine from other virtual machines that are not part of the same virtual subnet or part of the same routing domain.

Hyper-V Network Virtualization Policy Management

The Windows platform provides public APIs for datacenter management software to manage HNV. Virtual Machine Manager is one such datacenter management product. The management software contains all of the HNV policies. Because the virtual machine manager must be aware of virtual machines and more importantly provisions virtual machines and complete customer virtual networks in the datacenter and must be multi-tenant aware, managing HNV policy is a natural extension for policy-based networking.

Summary

Cloud-based datacenters can provide many benefits such as improved scalability and better resource utilization. To realize these potential benefits requires a technology that fundamentally addresses the issues of multi-tenant scalability in a dynamic environment. HNV was designed to address these issues and also improve the operational efficiency of the datacenter by decoupling the virtual network topology for the physical network topology. Building on an existing standard, HNV runs in today's datacenter and as NVGRE-aware hardware becomes available the benefits will continue to increase. Customers, with HNV can now consolidate their datacenters into a private cloud or seamlessly extend their datacenters to a hoster's environment with a hybrid cloud.

See also

To learn more about HNV see the following links:

Content type	References
Product evaluation	• Deep Dive on Hyper-V Network Virtualization in Windows Server 2012 R2 • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 1 of 2) • How to Design and Configure Networking in Microsoft System Center - Virtual Machine Manager and HyperV (Part 2 of 2)
Architecture	Hyper-V Network Virtualization Gateway Architectural Guide
Solution Guidance	**Hybrid Cloud Multi-Tenant Networking Solution Guide**
Community Resources	• Windows Server 2012 Hyper-V Network Virtualization Survival Guide

Content type	References
	• Private Cloud Architecture Blog • Ask questions: cloudnetfb@microsoft.com
Sample Scripts	• Simple demo script • Simple demo with gateway script
RFC	NVGRE Draft RFC
Related Technologies	Hyper-V Virtual Switch Overview

Hyper-V Network Virtualization Gateway Architectural Guide

Windows Server 2012 introduces a scalable, secure multi-tenant solution called Hyper-V Network Virtualization. This technology can be used to build cloud datacenters and makes it easier for customers to incrementally move their network infrastructure to private, hybrid, and public clouds.

For an overview of Hyper-V Network Virtualization, see Hyper-V Network Virtualization Overview, Hyper-V Network Virtualization Technical Details, and Windows Server® 2012 Hyper-V Survival Guide at the Microsoft TechNet web site.

Hyper-V Network Virtualization is deployable with existing server and switch hardware. However, gateway functionality, in addition to a datacenter management solution, is required to complete key scenarios including private cloud, cross-premise, and hybrid cloud deployments.

This article provides requirements for the different types of Hyper-V Network Virtualization gateways and gateway functionality. Note that gateway functionality is specifically called out because existing network appliances may incorporate Hyper-V Network Virtualization gateway functionality. Additional technical details regarding the packet encapsulation format used by Hyper-V Network Virtualization may be found in the NVGRE draft RFC.

Gateway Scenarios

Most customer deployments require communication from the network virtualized environment to the non-network virtualized environment. Therefore Hyper-V Network Virtualization gateways are required to bridge the two environments.

Gateways can come in different form factors. They can be built upon Windows Server 2012, incorporated into a Top of Rack (TOR) switch, put into an existing network appliance, or can be a stand-alone network appliance. The routing functionality is required for the private cloud scenario described in the Private Cloud section. In addition to routing, VPN functionality is required for the hybrid cloud scenario (described in the Hybrid Cloud section). A third scenario incorporates the gateway functionality into a load balancer (described in the Load Balancer section).

Private Cloud (routing)

Large enterprises may either be hesitant, or for compliance reasons, unable to move some of their services and data to a public cloud hoster. However, enterprises still want to obtain the benefits of the cloud and network virtualization by consolidating their datacenter resources into a private cloud. In a private cloud deployment, overlapping IP addresses may not be needed because corporations have sufficient internal non-routable address space (for example, 10.x.x.x or 192.x.x.x). Consider the example shown in Figure 1.

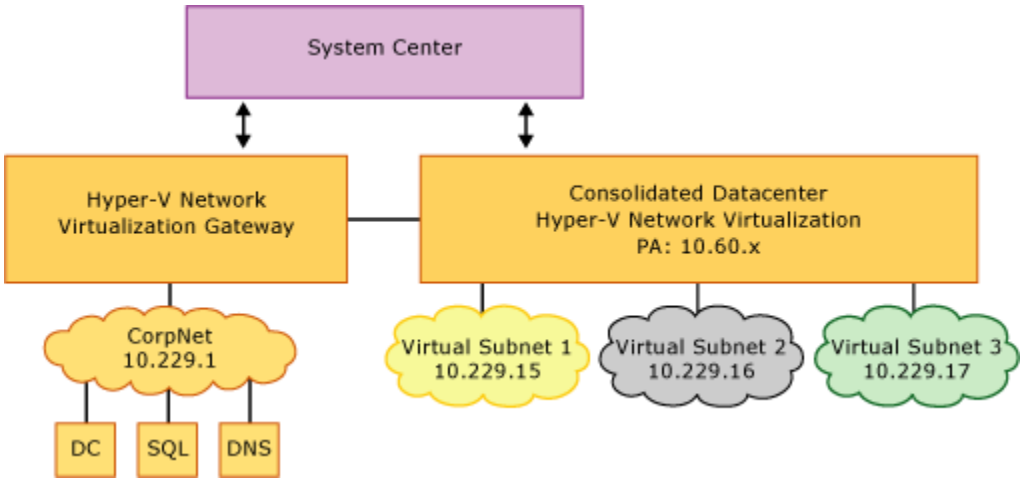

Figure 1: Private cloud deployment

Notice in this example that the customer addresses in the virtual subnets are 10.229.x IP addresses while the IP addresses in the non-network virtualized part of the network (CorpNet) are 10.229.1 addresses. In this case the PA addresses for the virtual subnets in the datacenter are 10.60.x IP addresses. This deployment allows the enterprise to take advantage of Hyper-V Network Virtualization's ability to offer flexibility in both virtual machine placement and cross-subnet live migration in the datacenter fabric. This increases datacenter efficiency thereby reducing both Operational Expenses (OpEx) and Capital Expenses (CapEx).

The Private Cloud scenario shown in Figure 1 requires a private cloud gateway appliance or gateway functionality incorporated into an existing network appliance. The purpose of the private cloud gateway is to provide routing and transition between physical and virtual addresses.

Hybrid Cloud (S2S VPN)

A key advantage of Hyper-V Network Virtualization is that it can seamlessly extend an on-premise datacenter to a Windows Server 2012 based cloud datacenter. This is called a Hybrid Cloud model as shown in Figure 2.

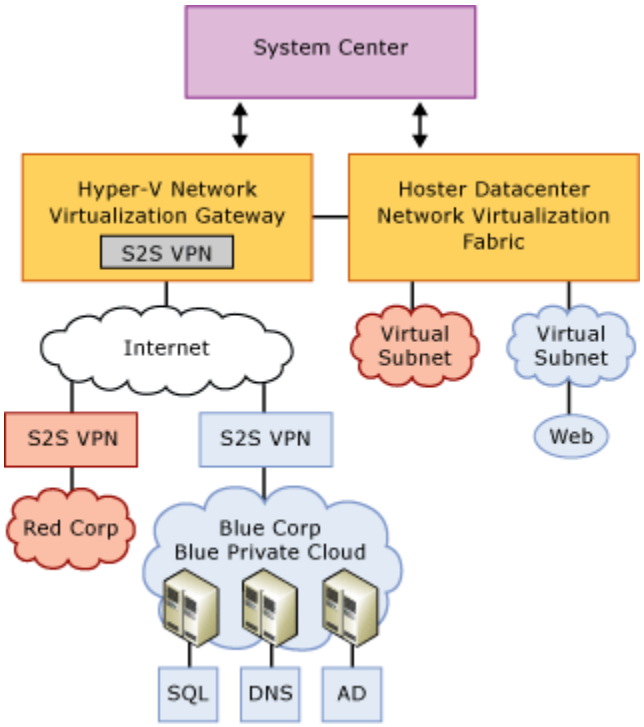

Figure 2: Hybrid cloud deployment

In this scenario an internal server, such as a web server, is moved from the enterprise network into a cloud hoster's datacenter. Taking advantage of *Bring Your Own IP Address* offered by the hoster, the enterprise does not need to change the network configuration of the web server virtual machine or any other network endpoint that references that web server. The hoster provides a secure link via a VPN gateway appliance. The enterprise administrators need to only configure their on-premise VPN with the appropriate IP address. The web server virtual machine is unaware that it has been moved to the cloud. It remains domain-joined with Active Directory (AD) and uses the enterprise's DNS server. The web server virtual machine also continues to interact with other servers in the enterprise such as a SQL Server. In this example, all three of these services (AD, DNS, SQL) remain on-premise. Network diagnostic tools such as tracert, which counts the number of network hops between a source and destination, show that the network traffic between the web server virtual machine and SQL is no longer local but routed across the Internet.

The Hybrid Cloud scenario shown in Figure 2 requires a multi-tenant VPN appliance such as a network virtualization gateway appliance. The VPN appliance gateway must interact with the datacenter orchestrator (for example, System Center Virtual Machine Manager) to obtain the appropriate network virtualization policies.

Load Balancer

A load balancer provides the illusion of a single IP address (Virtual IP address or VIP) to clients requesting a service. The service is implemented on multiple different IP addresses (Direct IP

address or DIP). A traditional load balancer maps network traffic from VIP to DIP via Network Address Translation (NAT). The load balancer re-writes the destination IP address (VIP) of the incoming packet with one of the IP addresses (DIP) behind the load balancer that performs the service. In the case where the load balancer is acting as a bi-directional NAT, the DIP sends a return packet to the load balancer and the load balancer rewrites the source IP (DIP) with the VIP. In this case the client only knows about the VIP and is unaware of the DIP.

In the context of Hyper-V Network Virtualization the load balancer emits NVGRE packets to get the packets to their destination virtual machines.

Consider the following scenario where the Blue Company and the Red Company have their virtual machines in the same multi-tenant datacenter that provides load balancing for the tenants. Blue Company has four virtual machines offering a service behind a load balancer. The Red Company has two virtual machines offering a different service that they want load balanced.

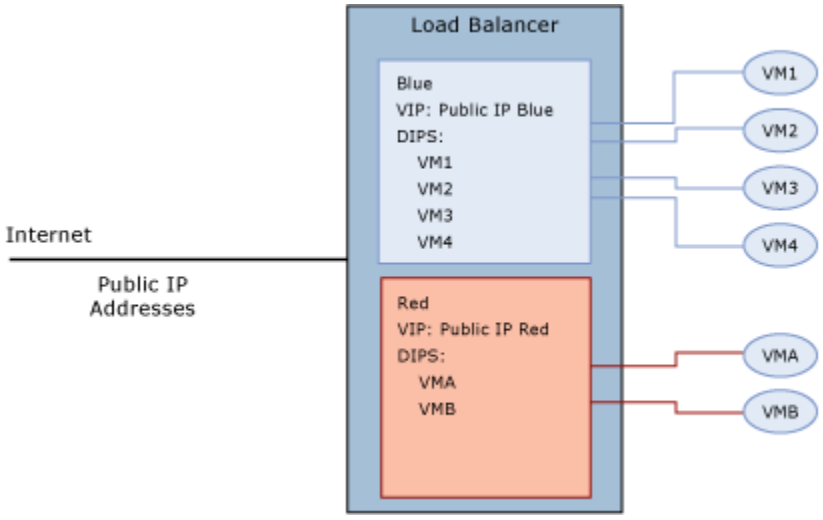

Figure 3: Multi-tenant Load Balancer Example

Figure 3 shows a virtual topology where publically routable IP addresses are load balanced across multiple destination virtual machines. For example, clients from the Internet send traffic to the Blue public IP (VIP). The load balancer, based on Network Virtualization policy, emits an NVGRE packet to get the packet to the appropriate Hyper-V host. They Hyper-V host will decapsulate the packet and then deliver it to the destination virtual machine.

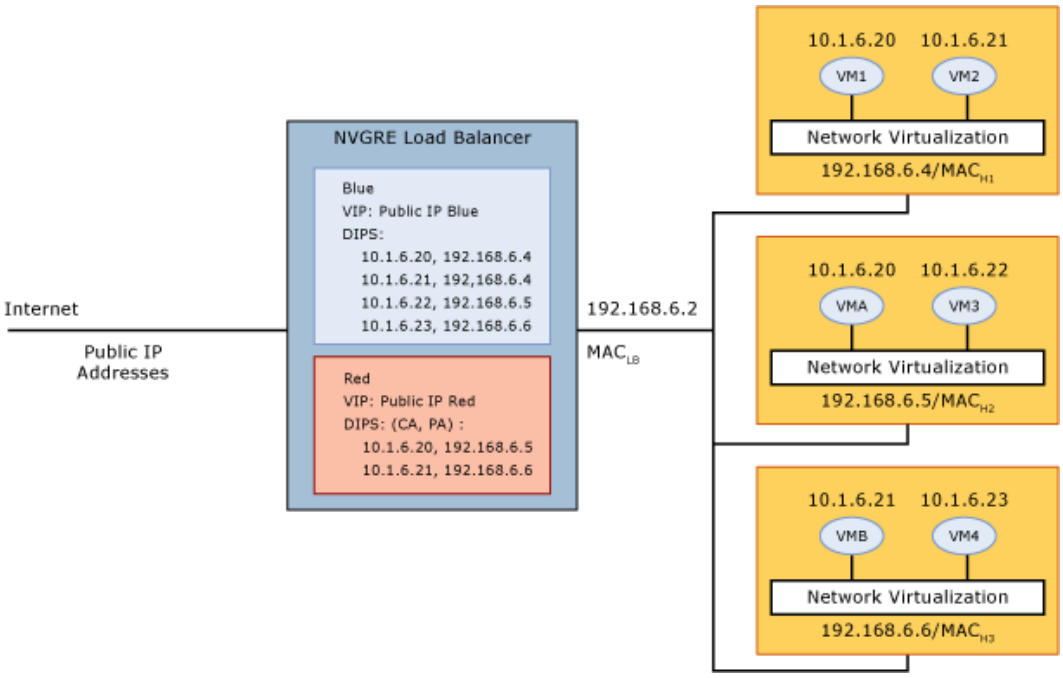

Figure 4: Multi-tenant Load Balancer deployment

Figure 4 show a possible deployment of Figure 3. In this case the PA space of the load balanced virtual machines is 192.168.6.x. Blue Customer has four virtual machines that are running on three physical hosts. Red Customer has two virtual machines running on two different hosts. Both Blue and Red are hosting web services. The outside world accesses these services via the appropriate publically routable VIP (Blue VIP and Red VIP). The load balancer, using network virtualization policy, generates the appropriate NVGRE packet.

Consider the case where a connection initiated from ClientIP and sent to the Blue VIP is load balanced to Blue VM2. The load balancer needs to send a packet to Blue VM2 which has a VSID of BlueVSID1, CA of 10.1.6.21 and a PA of 192.168.6.4. The load balancer's PA is 192.168.6.2. Therefore the load balancer generates the NVGRE packet header:

MAC_{LB} --> MAC_{H1}	192.168.6.2 --> 192.168.6.4	BlueVSID1	MAC_{ext} --> MAC_{VM2}	ClientIP --> 10.1.6.21

Now consider the scenario shown in Figure 5 where Blue and Red customers have multi-tier applications that require load balancing. From the perspective of the virtual machines they communicate to the next tier via VIP1. For example virtual machines in Tier B1 communicate to virtual machines in Tier B2 via Blue VIP1. Blue VIP1 is a Blue CA IP address which will also have a corresponding PA address on the load balancer. A possible deployment topology is shown in Figure 6.

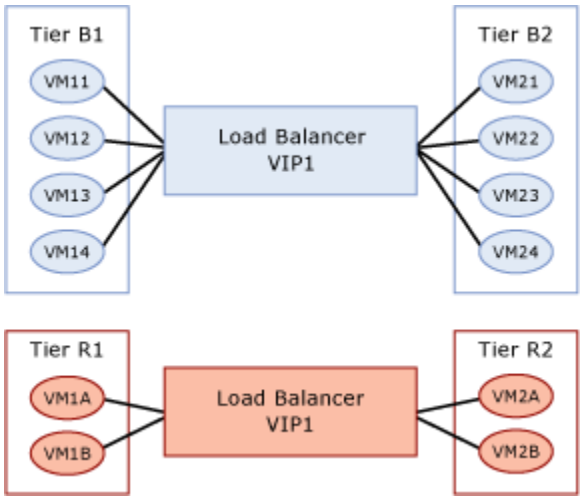

Figure 5: Load balancing between tiers (logical perspective)

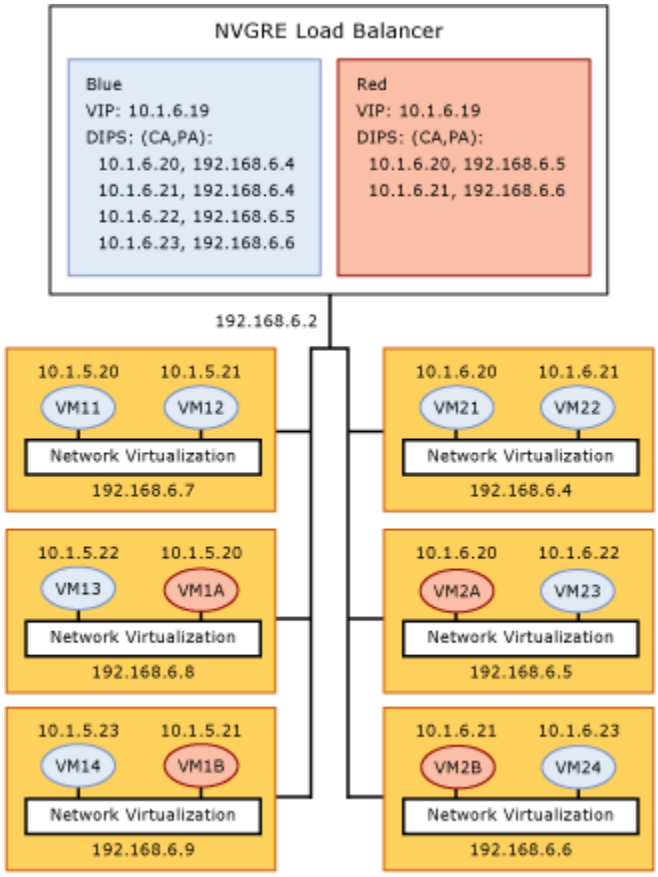

Figure 6: Load Balancer placed internally in the datacenter

The Blue CA VIP and the Red CA VIP on the load balancer both have the PA of 192.168.6.2. The load balancer needs to look at the VSID in the NVGRE packet and then determine the appropriate VIP to DIP mapping. For example Blue virtual machine 11 will send a packet to Blue

CA VIP. The NVGRE packet header will have Source IP: 192.168.6.7, Destination IP: 192.168.6.2, GRE Key Field Blue VSID1, Inner Source IP: 10.1.5.20, Inner Destination IP: Blue CA VIP, and the rest of the original packet. The load balancer's policy must match on both the VSID and the VIP because with overlapping CA IP addresses the VIP by itself may not be unique.

External Router Gateway Configuration

Gateways must contain Hyper-V Network Virtualization policy and perform actions similar to a Windows Server 2012 host. In the typical deployment the Windows Server 2012 host is aware of all the policy information for every routing domain for all virtual machines located on that host. For example, if a host has a virtual machine with routing domain X and another virtual machine from routing domain Y the host has all the policy information for all virtual subnets contained in both routing domain X and routing domain Y. Even in the case of a single routing domain there could be a substantial amount of policy that needs to be on every Hyper-V host.

Consider the case of a private cloud shown in Figure 1. In this scenario a corporation has an internal cloud with a single routing domain. In this case every Hyper-V host must have policy for all network virtualized virtual machines. This can place a burden on the datacenter orchestrator to maintain this policy on all hosts. Note that it is rare for a virtual machine to actually communicate with every other virtual subnet in the corporation so much of this policy on a given host is in practice not needed.

To reduce the amount of policy needed on each host the Windows Server 2012 host can be put into a special mode where it routes traffic within a given Virtual Subnet based on its own host policy. Any cross virtual-subnet traffic is sent to an external router for a network virtualization policy check and potentially forwarding the packet along the path to its eventual destination. In such a deployment the host only needs to be aware of policy for the virtual subnets of virtual machines running on that host and not for every virtual subnet of the routing domains of the virtual machines running on the host. In this scenario the Hyper-V Network Virtualization Gateway makes all cross-subnet routing decisions. The Hyper-V Network Virtualization Gateway must have all policy information for all virtual machines in the virtual subnets to which it routes packets.

Management

For any Hyper-V Network Virtualization deployment a datacenter orchestrator is required. One such datacenter orchestrator is System Center Virtual Machine Manager (VMM). The datacenter orchestrator is responsible for providing a mechanism to get Hyper-V Network Virtualization policies to the appropriate gateway appliance. Note that PA IP routing is not managed by the datacenter orchestrator.

Gateway Console

The Hyper-V Network Virtualization Gateway partner should provide a management console for basic configuration of the gateway appliance or appliance functionality. This management console should handle operations such as PA assignment, high availability, monitoring, and

authentication. The management console will be responsible for configuring the aspects of the gateway not managed by a datacenter orchestrator such as VMM.

Multiple Gateways

High Availability by using multiple physical Hyper-V Network Virtualization gateways is outside the scope of this article. The datacenter orchestrator (for example VMM) will need to configure and manage the gateway. Scenarios in which the datacenter deploys multiple gateways include scalability and resiliency. For scalability a single physical gateway may not be able to handle the load required for a datacenter. Additionally a single gateway introduces a single point of failure. The partner's management console should handle deployments consisting of multiple gateways.

System Center Virtual Machine Manager

In the VMM model the Hyper-V Network Virtualization Gateway is managed via a PowerShell plug-in module. Partners building Hyper-V Network Virtualization gateways need to create a PowerShell plug-in module which physically runs on the VMM server. This plug-in module will communicate policy to the gateway. Figure 7 shows a block diagram of VMM managing a Hyper-V Network Virtualization deployment. Note that a partner plug-in runs inside the VMM server. This plug-in communicates to the gateway appliances. The protocol used for this communication is not specified here. The partner may determine the appropriate protocol. Note that VMM uses the Microsoft implementation of WS-Management Protocol called Windows Remote Management (WinRM) and Windows Management Instrumentation (WMI) to manage the Windows Server 2012 hosts and update network virtualization policies.

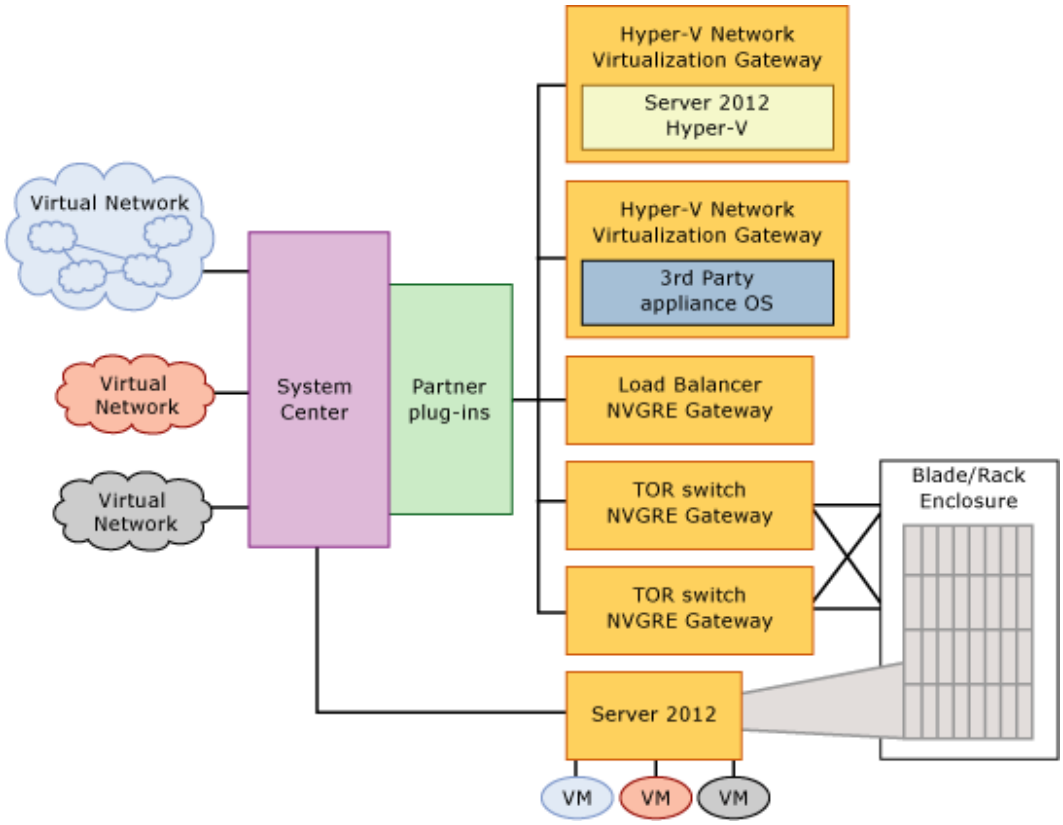

Figure 7: Using System Center to manage Hyper-V Network Virtualization

Windows-Based Gateway Appliances

Windows Server 2012 can be used as the base platform for a Hyper-V Network Virtualization Gateway appliance.

Private Cloud Router Architecture

Figure 8 shows the architecture for the Private Cloud Router using Windows Server 2012 as the base platform. In this scenario all virtual machines are in the same routing domain. The red, blue, and purple virtual machines have a CA IP prefix of 157.4, 157.3, 157.2, respectively. The services on CorpNet have IP addresses with a prefix of 157.1.

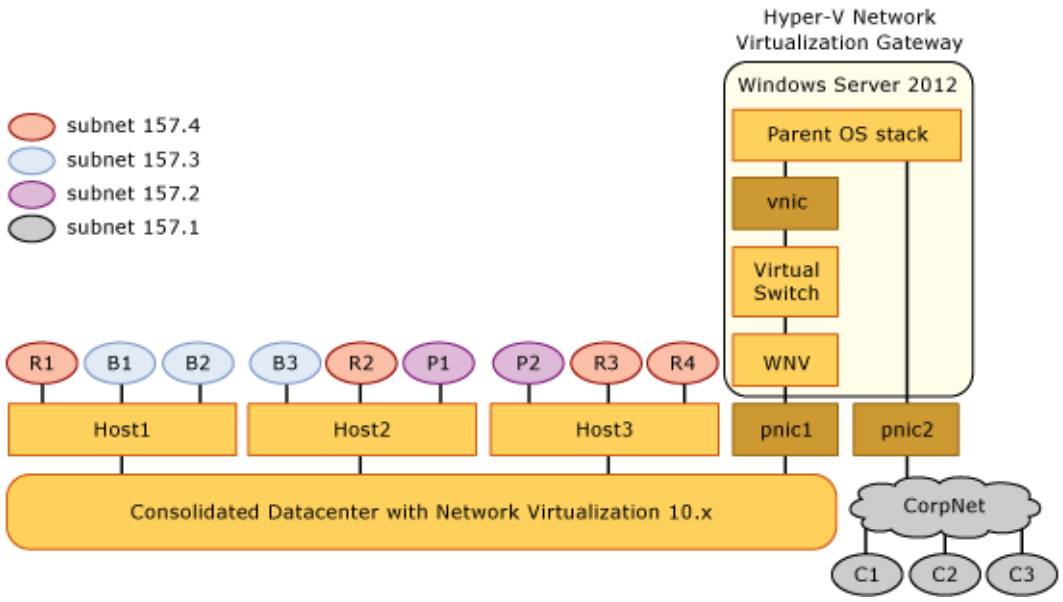

Figure 8: Private Cloud Gateway using Windows Server 2012

The WNV filter in the gateway receives traffic from the red, blue, and purple virtual subnets and sends it to CorpNet via the parent(Management) operating system stack. The virtual switch has a parent (host) VNIC for communication to the host operating system stack from the network virtualized world. The other network interface is bound to the physical network interface (pnic2) connected to CorpNet. Note that there is no requirement that there be two physical network interfaces in the Hyper-V Network Virtualization Gateway shown in Figure 8. A single physical network interface could be used along with two parent virtual network interfaces. However, datacenter administrators may prefer a physical separation of traffic. Two physical network interfaces in the gateway allow the consolidated datacenter and CorpNet traffic to interact with the gateway via different physical switches. Two physical network interfaces, one for the network virtualized world and one for the non-network virtualized world provides more flexibility in customer deployments.

Cross Premise Gateway

The Hybrid Cloud scenario enables an enterprise to seamlessly expand their on-premises datacenter into the cloud. This requires a site to site VPN tunnel. This can be accomplished with Windows Server 2012 as the host platform and a per tenant Windows Server 2012 guest virtual machine running a Site To Site (S2S) VPN tunnel connecting the cloud datacenter with various on-premise datacenters. Windows Server 2012 S2S VPN supports IKEv2 and configuration of remote policy can be accomplished via PowerShell/WMI. In addition Windows Server 2012 guest virtual machines support new network interface offload capabilities that enhance the performance and scalability of the gateway appliance. These offload capabilities are discussed below in the Hardware Considerations section.

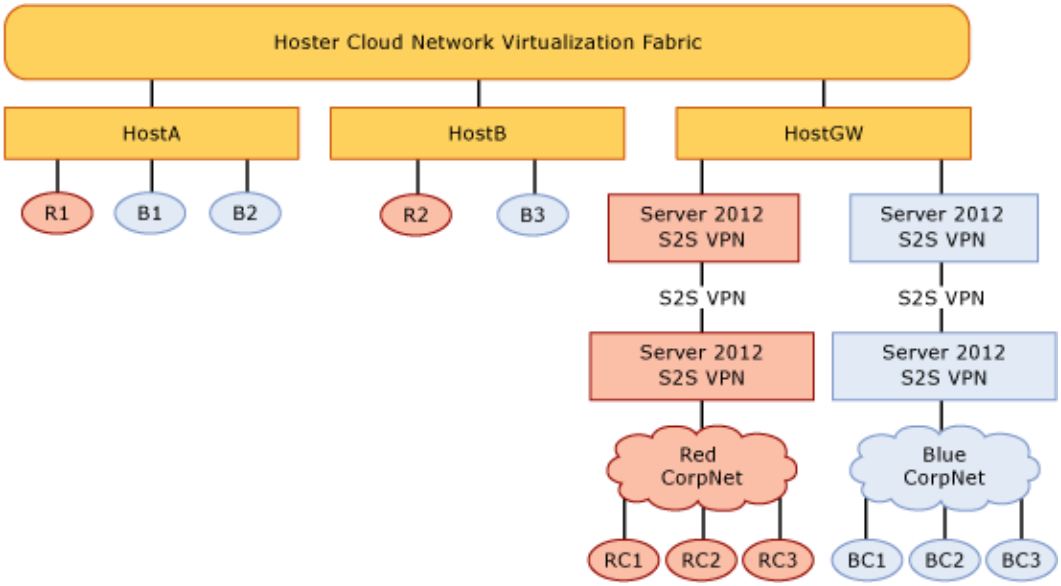

Figure 9: Hybrid Cloud with Windows Server 2012 based Gateway

Figure 9 shows a scenario where Red Corp and Blue Corp are customers of Hoster Cloud. Red Corp and Blue Corp seamlessly extend their datacenter into Hoster Cloud has deployed Windows Server 2012 based per tenant virtual machine gateways allowing Red Corp and Blue Corp to seamlessly extend their on-premise datacenters. In Figure 10 there is no requirement that Red Corp or Blue Corp run Windows Server 2012 S2S VPN, only that the customer's on premise S2S VPN support IKEv2 to interact with corresponding Windows Server 2012 S2S virtual machines running on HostGW.

Figure 9 shows the internal architecture for HostGW. Each Routing Domain requires its own virtual machine. The technical reason for this is that a vmnic can only be associated with a single Virtual Subnet (VSID) and a VSID can only be part of a single routing domain. The VSID switch port ACL does not support trunking of VSIDs. Therefore the simplest way to provide isolation is with a per tenant (Routing Domain) gateway virtual machine.

Each of the virtual machines is dual homed which means they have two virtual network interfaces. One of the virtual network interfaces has the appropriate VSID associated with it. The other virtual network interface has a VSID of 0 which means traffic is not modified by the WNV filter. The Windows Server 2012 virtual machine is running RRAS and using IKEv2 to create a secure tunnel between Hoster Cloud and the customer's on premise gateway.

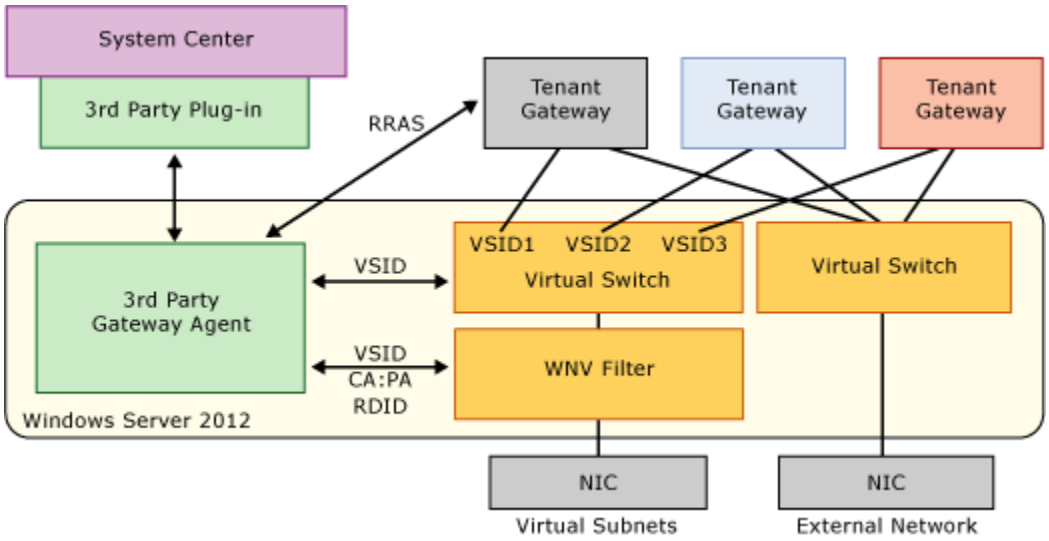

Figure 10: Hybrid Cloud with Windows Server 2012 based Per Tenant VM Gateways

Figure 10 shows the architecture where VMM is managing a Hyper-V Network Virtualization deployment. The partner has a plug-in that runs in the VMM server. When using Windows Server 2012 as a Hyper-V Network Virtualization gateway appliance a local management process running in Windows is required as the end point for this communication from the plug-in running in the VMM server. This is how the plug-in is able to communicate network virtualization policy to the WNV filter running on HostGW.

Hardware Considerations

The scalability of the Windows Server 2012 based gateway appliance will be determined by the server hardware including:

- NUMA capabilities
- number of CPU cores
- amount of RAM
- offload capabilities and performance of the network interfaces

In general, a dual socket server class platform provided by major OEMs should be a sufficient base platform. The number of virtual machines that can effectively run on the server appliance is typically determined by the amount of RAM in the server.

The choice of network interfaces can also have a great impact on the performance and scalability of the gateway appliance. Figure 10 shows a two network interface configuration. Note that ALL traffic sent on the virtual subnet path will be NVGRE traffic. For encapsulated packets traditional offloads such as Large Send Offload (LSO) on the send path and Virtual Machine Queues (VMQ) on the receive path will not provide the expected benefit because they operate on the outer packet header. The outer header for NVGRE traffic makes it appear as though all traffic is generated using the same IP address and destined for a single IP address. Network interface offloads provide substantial benefit to the Hyper-V switch performance. A new feature in Windows

Server 2012 is GRE Task Offload where the network interface will operate on the inner packet header for standard offloads such that LSO and VMQ provide their expected performance and scalability benefits. Therefore the network interface on the path for Virtual Subnets in Figure 10 should support GRE Task Offload. Such network interfaces are expected to be commercially available by the end of 2012.

Another new feature in Windows Server 2012 is IPsec Task Offload for guest virtual machines. IPsec is used for securing site to site VPN tunnels. Both sides of the S2S tunnel must support IKEv2. If the Cross-Premise gateway is NOT behind a NAT then IPsec Task Offload (IPsecTO) can provide better performance by having the External Network network interface in Figure 9 perform the per packet encryption operations on behalf of the per Tenant RRAS virtual machine. Both sides of the S2S tunnel would need to use AES-GCM128 to take advantage of IPsecTO. Therefore, the following conditions must be met for the RRAS virtual machine to take advantage of IPsecTO:

- Both sides of the S2S tunnel must use AES-GCM128
- The gateway appliance must not be behind a NAT

Because IPsec operations can consume a large amount of CPU, IPsecTO can provide a significant CPU reduction and also an increase in network throughput.

SR-IOV is a new Windows Server 2012 feature that directly exposes a physical network interface to the virtual machine. SR-IOV bypasses the virtual switch. Unfortunately, current IPsecTO network interfaces on the market do not support IPsecTO on the SR-IOV path. The Windows Server 2012 Hyper-V Switch introduces new policy-based features such as ACLs and QoS. If any of these Hyper-V switch policies are set on a virtual network interface , the virtual network interface traffic will not take the SR-IOV path and instead will need to go through the Hyper-V Switch path for policy enforcement. When traffic does take the SR-IOV path instead of the Hyper-V Switch path there will be a significant reduction in CPU utilization.

For the Private Cloud gateway similar considerations should be evaluated regarding hardware requirements.

Additional Resources

Additional information regarding Hyper-V Network Virtualization and System Center Virtual Machine Manager can be found in the following articles (articles without links are still under development):

- Windows Server 2012 Hyper-V Network Virtualization Survival Guide
- Hyper-V Network Virtualization Overview
- Network Virtualization Technical Details
- //BUILD talk on Hyper-V Network Virtualization
- //BUILD Keynote Demo (starts 13.5 minutes into video)
- NVGRE Draft RFC

Hyper-V Virtual Switch Overview

This topic describes Hyper-V Virtual Switch in Windows Server 2012. This topic also lists some practical uses for the Hyper-V virtual switch, hardware and software requirements, and provides links to additional information.

Notes

In addition to this topic, the following Hyper-V Virtual Switch documentation is also available.

- What's New in Hyper-V Virtual Switch in Windows Server 2012 R2
- What's New in Hyper-V Virtual Switch in Windows Server 2012
- Create Security Policies with Extended Port Access Control Lists for Windows Server 2012 R2
- Hyper-V: Configure VLANs and VLAN Tagging
- Hyper-V: The WFP virtual switch extension should be enabled if it is required by third party extensions
- Create networks with VMM 2012
- Unified Tracing Overview

Hyper-V Virtual Switch

The Hyper-V Virtual Switch is a software-based layer-2 network switch that is available in Hyper-V Manager when you install the Hyper-V server role. The switch includes programmatically managed and extensible capabilities to connect virtual machines to both virtual networks and the physical network. In addition, Hyper-V Virtual Switch provides policy enforcement for security, isolation, and service levels.

The Hyper-V Virtual Switch in Windows Server® 2012 introduces several new features and enhanced capabilities for tenant isolation, traffic shaping, protection against malicious virtual machines, and simplified troubleshooting.

With built-in support for Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers, the Hyper-V Virtual Switch enables independent software vendors (ISVs) to create extensible plug-ins (known as Virtual Switch Extensions) that can provide enhanced networking and security capabilities. Virtual Switch Extensions that you add to the Hyper-V Virtual Switch are listed in the Virtual Switch Manager feature of Hyper-V Manager.

In the following illustration, a Virtual Machine (VM) has a virtual NIC that is connected to the Hyper-V Virtual Switch through a switch port.

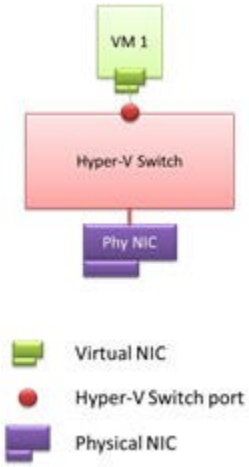

The capabilities provided in the Hyper-V Virtual Switch mean that organizations have more options for enforcing tenant isolation, shaping and controlling network traffic, and employing protective measures against malicious virtual machines.

Practical applications

Displaying statistics: A developer at a hosted cloud vendor implements a management package that displays the current state of the Hyper-V virtual switch. The management package can query switch-wide current capabilities, configuration settings, and individual port network statistics using WMI. The status of the switch is then displayed to give administrators a quick view of the state of the switch.

Resource tracking: A hosting company is selling hosting services priced according to the level of membership. Various membership levels include different network performance levels. The administrator allocates resources to meet the SLAs in a manner that balances network availability. The administrator programmatically tracks information such as the current usage of bandwidth assigned, and the number of virtual machine (VM)—assigned virtual machine queue (VMQ) or IOV channels. The same program also periodically logs the resources in use in addition to the per-VM resources assigned for double entry tracking or resources.

Managing the order of switch extensions: An enterprise has installed extensions on their Hyper-V host to both monitor traffic and report intrusion detection. During maintenance, some extensions may be updated causing the order of extensions to change. A simple script program is run to reorder the extensions after updates.

Forwarding extension manages VLAN ID: A major switch company is building a forwarding extension that applies all policies for networking. One element that is managed is virtual local area network (VLAN) IDs. The virtual switch cedes control of the VLAN to a forwarding extension. The switch company's installation programmatically call a Windows Management Instrumentation (WMI) application programming interface (API) that turns on the transparency, telling the Hyper-V Virtual Switch to pass and take no action on VLAN tags.

Important functionality

Some of the principal features that are included in the Hyper-V Virtual Switch are:

1. **ARP/ND Poisoning (spoofing) protection:** Provides protection against a malicious VM using Address Resolution Protocol (ARP) spoofing to steal IP addresses from other VMs. Provides protection against attacks that can be launched for IPv6 using Neighbor Discovery (ND) spoofing.
2. **DHCP Guard protection:** Protects against a malicious VM representing itself as a Dynamic Host Configuration Protocol (DHCP) server for man-in-the-middle attacks.
3. **Port ACLs:** Provides traffic filtering based on Media Access Control (MAC) or Internet Protocol (IP) addresses/ranges, which enables you to set up virtual network isolation.
4. **Trunk mode to a VM:** Enables administrators to set up a specific VM as a virtual appliance, and then direct traffic from various VLANs to that VM.
5. **Network traffic monitoring:** Enables administrators to review traffic that is traversing the network switch.
6. **Isolated (private) VLAN:** Enables administrators to segregate traffic on multiple vlans, to more easily establish isolated tenant communities.

Following is a list of capabilities that enhance Hyper-V Virtual Switch usability:

1. **Bandwidth limit and burst support:** Bandwidth minimum guarantees amount of bandwidth reserved. Bandwidth maximum caps the amount of bandwidth a VM can consume.
2. **ENC marking support:** Explicit Congestion Notification (ECN) marking—also known as Data CenterTCP (DCTCP)—enables the physical switch and operating system to regulate traffic flow such that the buffer resources of the switch are not flooded, which results in increased traffic throughput.
3. **Diagnostics:** Diagnostics allow easy tracing and monitoring of events and packets through the virtual switch.

Hyper-V virtual switch

The Hyper-V Virtual Switch features described in the previous—Important functionality—section of this topic enable administrators to configure security and isolation options, and monitor traffic in ways not previously provided. The extensible nature of the switch enables ISV to provide an additional layer of customization.

What value does this change add?

The recent increased utilization of virtualization has resulted in many hosting companies placing VMs for multiple clients on the same computer, increasing the need for isolation and protection. While Windows Server 2008 R2 does provide the default protection of MAC spoofing, server releases up through Windows Server 2008 R2 provide only minimal security protection for virtualized network traffic. In Windows Server® 2012, traffic that flows between VMs on the same physical host computer is more secure because of enhancements that protect against malicious virtual machines.

In Windows Server® 2012, the new Hyper-V Virtual Switch provides more security, including functionality that will allow customers to readily monitor and move traffic through the switch.

Additionally, the Hyper-V Virtual Switch supports an interface in which ISVs can extend the switch functionality.

Hardware requirements

Hyper-V Virtual Switch requires a 64-bit processor that includes the following:

- Product disk or files for Windows Server® 2012
- Physical computer for hosting Windows Server® 2012
- Hardware-assisted virtualization. This is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.
- Hardware-enforced Data Execution Prevention (DEP) must be available and enabled. Specifically, you must enable Intel XD bit or AMD NX.

See also

Following is a list of resources related to the Hyper-V virtual switch.

Content type	References
Product evaluation	Understand and Troubleshoot Hyper-V Virtual Network Switch in Windows Server® 2012
Developer Resources	MSDN: Hyper-V Extensible Switch
Community resources	Microsoft Server and Cloud Platform Blog: Windows Server 2012: Introducing Hyper-V Extensible Switch | Virtual PC Guy's Blog: Hyper-V Extensible Switch in Windows Server 2012 | Windows Lifestyle: Hyper-V Extensible Switch in Windows Server 2012
Related technologies	Hyper-V Overview | Microsoft Private Cloud Solutions

Related documentation

- Hyper-V Overview in Windows Server 2012
- Hyper-V in Windows Server 2008 and Windows Server 2008 R2 technical library on TechNet
- Developer resources for Hyper-V Extensible Switch in the Microsoft Developer Network (MSDN) library
- Protected 802.3 Ethernet connections using switch-based 802.1X Authenticated Wired Access

What's New in Hyper-V Virtual Switch in Windows Server 2012 R2

This topic provides information about the new features for the Hyper-V Virtual Switch in Windows Server® 2012 R2.

This topic contains the following sections.

- Hyper-V Virtual Switch Extended Port ACLs
- Dynamic Load Balancing of Network Traffic
- Hyper-V Network Virtualization coexists with third party forwarding extensions for the Hyper-V Virtual Switch
- Traffic bottlenecks to VMs are reduced with vRSS
- Network tracing is streamlined and provides more detail

Hyper-V Virtual Switch Extended Port ACLs

Enterprises and Cloud Service Providers (CSPs) can configure the Hyper-V Virtual Switch Extended Port Access Control Lists (ACLs) to provide firewall protection and enforce security policies for the tenant VMs in their datacenters. Because the port ACLs are configured on the Hyper-V Virtual Switch rather than within the VMs, you can manage security policies for all tenants in a multitenant environment.

Following are the new features for extended port ACLs:

- ACLs now include the socket port number. In Windows Server 2012, you were able to specify both source and destination MAC and IP addresses for IPv4 and IPv6. For Windows Server 2012 R2 you can also specify the port number when you create rules.
- You can now configure stateful rules that are unidirectional and provide a timeout parameter. With a stateful firewall rule, traffic is allowed, and two traffic flows are created dynamically. The two traffic flows are one outbound rule that matches five attributes in outbound packets, and one inbound rule that also match the same five attributes. After a stateful rule is utilized successfully one time, the two traffic flows are allowed without having to be looked up against the rule again for a period of time that you designate using the timeout attribute. When the firewall rule exceeds the timeout attribute, traffic flows are inspected against rules again.

In addition, extended port ACLs provide the following benefits:

- In multitenant environments, you can protect datacenter resources and provide security policy enforcement for your tenants.
- Compatibility with Hyper-V Network Virtualization.
- A management interface that allows you to easily configure firewall rules by using Windows PowerShell.
- Logging and diagnostics capabilities so that you can confirm firewall operation and detect any possible misconfiguration of the port ACLs.

- Configurable as a stateless firewall by filtering packets based on five attributes in the packet; with a stateless firewall configuration you can apply any firewall rule to either inbound or outbound network traffic, and the rule can either allow or deny traffic.

For more information, see the following topics.

- Create Security Policies with Extended Port Access Control Lists for Windows Server 2012 R2
- For more information about port ACLs introduced in Windows Server 2012, see the section "Port Access Control Lists (ACLs)" in the topic What's New in Hyper-V Virtual Switch for Windows Server 2012.

Dynamic Load Balancing of Network Traffic

Windows Server® 2012 provided simultaneous load distribution and failover, but did not ensure load distribution between the NICs in a NIC team in a balanced manner.

In Windows Server® 2012 R2, dynamic load balancing continuously and automatically moves traffic streams from NIC to NIC within the NIC team to share the traffic load as equitably as possible.

For more information on NIC Teaming, see NIC Teaming Overview.

Hyper-V Network Virtualization coexists with third party forwarding extensions for the Hyper-V Virtual Switch

Forwarding Hyper-V Virtual Switch extensions that are installed on the Hyper-V Virtual Switch in a Hyper-V Network Virtualization (HNV) environment can now forward packets for either the VM customer address (CA) space or the physical address (PA) space, because switch extensions now coexist seamlessly with Network Virtualization, which uses Network Virtualization Generic Routing Encapsulation (NVGRE).

When you have a third party forwarding extension installed, Hyper-V Virtual Switch now performs hybrid forwarding. With hybrid forwarding, network traffic that is NVGRE encapsulated is forwarded by the HNV module within the switch, while all non-NVGRE network traffic is forwarded by the third party forwarding extensions that you have installed.

In addition to forwarding, a third party forwarding extension can still apply other policies, such as ACLs and QoS, to both the NVGRE and the non NVGRE-encapsulated traffic. The forwarding extension that you install must be able to process both types of network traffic based on their intended destinations. For example, PA address visibility is necessary for extensions that perform switch team load balancing.

The policies and capabilities of the Hyper-V Virtual Switch and third party extensions do not displace each other – instead, they are mutually available.

For more information about Hyper-V Network Virtualization, see Hyper-V Network Virtualization Overview.

Traffic bottlenecks to VMs are reduced with vRSS

In Windows Server 2012, Receive Side Scaling (RSS) over SR-IOV is supported; now in Windows Server 2012 R2, virtual RSS (vRSS) is supported on the VM network path, allowing VMs to sustain a greater networking traffic load.

In the past, VMs might have trouble achieving network throughput approaching 10Gbps due to the processing load on a single CPU core. vRSS alleviates this problem by spreading the processing across multiple cores on the host and multiple cores on the VM.

To take advantage of vRSS, VMs must be configured to use multiple cores, and they must support RSS. vRSS is enabled automatically when the VM uses RSS on the VM network path.

For more information about RSS, see Receive Side Scaling (RSS).

Network tracing is streamlined and provides more detail

Network traces now contain switch and port configuration information, and tracing packets through the Hyper-V Virtual Switch and any forwarding extensions you have installed are easier to use and read.

For information about Unified Tracing, see Unified Tracing Overview and Netsh Commands for Network Trace in the Windows Server 2012 and Windows Server 2012 R2 Technical Library.

What's New in Hyper-V Virtual Switch in Windows Server 2012

In Windows Server® 2012, Hyper-V Virtual Switch includes new features for manageability, security, isolation, visibility, and scalability to provide more effective Hyper-V deployment in your on-site or private, hybrid, or public cloud environments.

Following are the primary new features for Hyper-V Virtual Switch.

- Windows PowerShell commands for Hyper-V Virtual Switch
- Multiple Virtual NICs
- New security features for Hyper-V Virtual Switch
- Port Virtual Local Area Network (PVLAN) and Trunk Mode
- Port Mirroring
- Receive Side Scaling (RSS) and Dynamic Virtual Machine Queue (dVMQ)

Windows PowerShell commands for Hyper-V Virtual Switch

You can now use Windows PowerShell commands, run manually or in scripts, to configure Hyper-V Virtual Switch and related features.

PowerShell syntax consists primarily of <verb>-<noun> pairs. For example, the following command displays all of the Hyper-V Virtual Switches that are configured on your server.

```
Get-VMSwitch
```

Windows PowerShell nouns that are provided for the Hyper-V Virtual Switch are:

```
VMNetworkAdapter

VMNetworkAdapterAcl

VMNetworkAdapterVlan

VMSwitch
```

A VMNetworkAdapter includes the following elements:

- The virtual network interface card (NIC) in the VM
- Hyper-V Virtual Switch port

VMNetworkAdapter is the primary noun that you can use to manage various security features, Quality of Service (QoS), port mirroring, and other items. For more information on these features, type the following into Windows PowerShell, and then press ENTER:

```
Get-Help Set-VMNetworkAdapter
```

To view all of the Windows PowerShell verbs for VMNetworkAdapter, type the following command into Windows PowerShell, and then press ENTER:

```
Get-Help *-VMNetworkAdapter
```

To view the current state of a VMNetworkAdapter, type the following command into Windows PowerShell, and then press ENTER:

```
Get-VMNetworkAdapter –VMName MyVM | format-list *
```

Following are the Windows PowerShell commands that you can use to configure the Virtual Switch and virtual network adapters. A full command reference for these and other Hyper-V Windows PowerShell cmdlets, including syntax and examples, is available at Hyper-V Cmdlets in Windows PowerShell.

Add-VMNetworkAdapter	Get-VMSystemSwitchExtensionPortFeature
Add-VMNetworkAdapterAcl	Get-VMSystemSwitchExtensionSwitchFeature
Add-VMSwitch	New-VMSwitch
Add-VMSwitchExtensionPortFeature	Remove-VMNetworkAdapter
Add-VMSwitchExtensionSwitchFeature	Remove-VMNetworkAdapterAcl
Connect-VMNetworkAdapter	Remove-VMSwitch

Disable-VMSwitchExtension	Remove-VMSwitchExtensionPortFeature
Enable-VMSwitchExtension	Remove-VMSwitchExtensionSwitchFeature
Get-VMNetworkAdapter	Rename-VMNetworkAdapter
Get-VMNetworkAdapterAcl	Rename-VMSwitch
Get-VMNetworkAdapterFailoverConfiguration	Set-VMNetworkAdapter
Get-VMNetworkAdapterVlan	Set-VMNetworkAdapterFailoverConfiguration
Get-VMSwitch	Set-VMNetworkAdapterVlan
Get-VMSwitchExtension	Set-VMSwitch
Get-VMSwitchExtensionPortData	Set-VMSwitchExtensionPortFeature
Get-VMSwitchExtensionPortFeature	Set-VMSwitchExtensionSwitchFeature
Get-VMSwitchExtensionSwitchData	
Get-VMSwitchExtensionSwitchFeature	
Get-VMSystemSwitchExtension	

Note

When using Windows PowerShell commands on physical NICs, use the NetAdapter family of cmdlets; and when using Windows PowerShell commands on virtual NICs, use the VMNetworkAdapter family of cmdlets. If you only have a single NIC or a team of NICs bound to the Hyper-V Virtual Switch, you can use multiple parent virtual NICs as well.

Multiple Virtual NICs

In previous versions of Hyper-V, only one parent virtual NIC was supported, however in Windows Server 2012 Hyper-V, multiple NICs are supported. In addition, you can share the physical NIC that is bound to the Hyper-V Switch with the management operating system.

Tip

To enable this functionality, open Hyper-V Manager. In the **Actions** pane, select **Virtual Switch Manager**. In Virtual Switch Manager, select a switch. In the details pane for the switch, which is illustrated below, select **Allow management operating system to share this network adapter**.

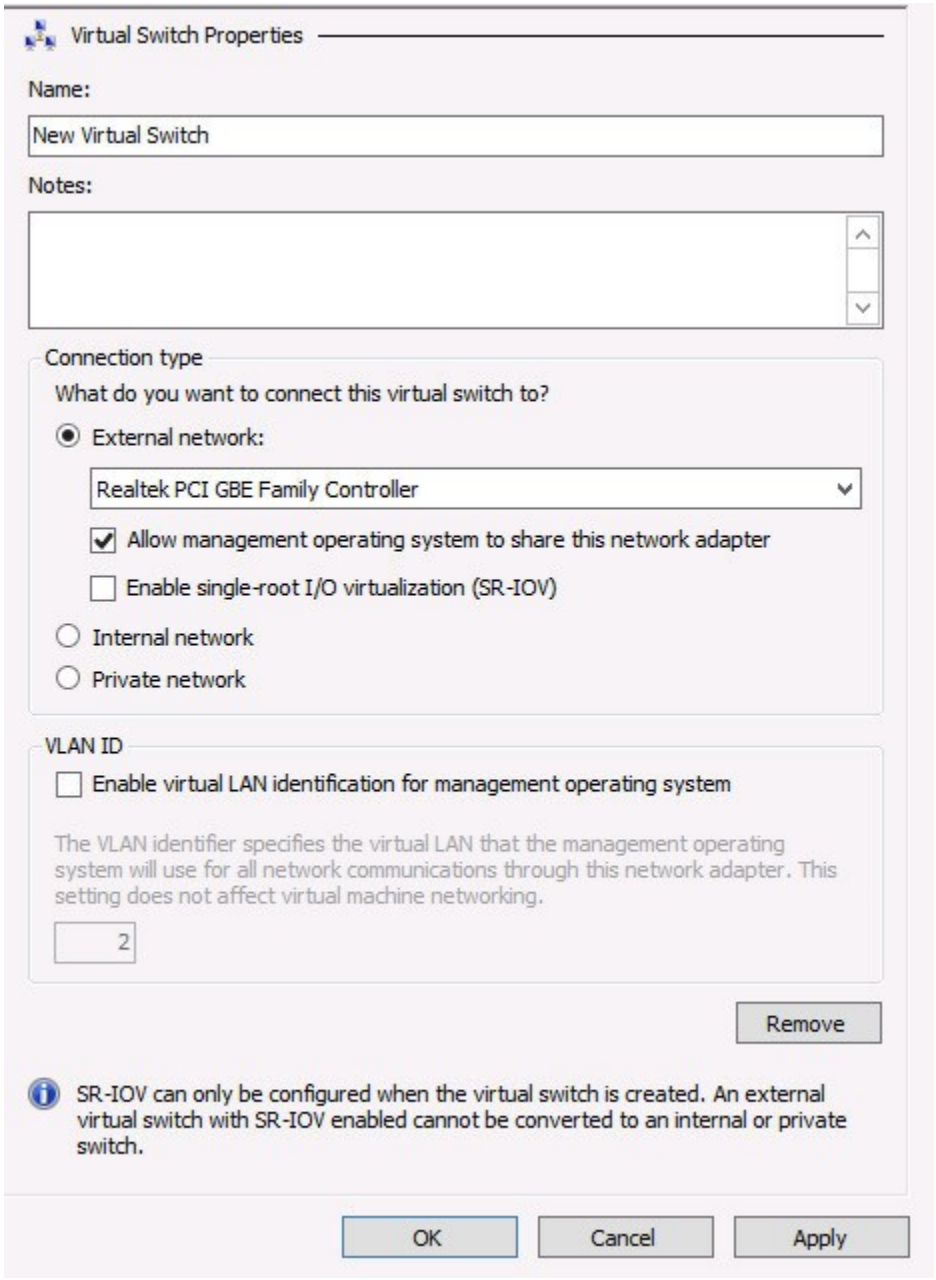

You can create multiple parent virtual NICS that you use for live migration, storage, and management; and you can assign each virtual NIC to a different virtual Local Area Network (VLAN). You can also create different Quality of Service (QoS) polices for each virtual NIC.

A multiple NIC configuration is illustrated below:

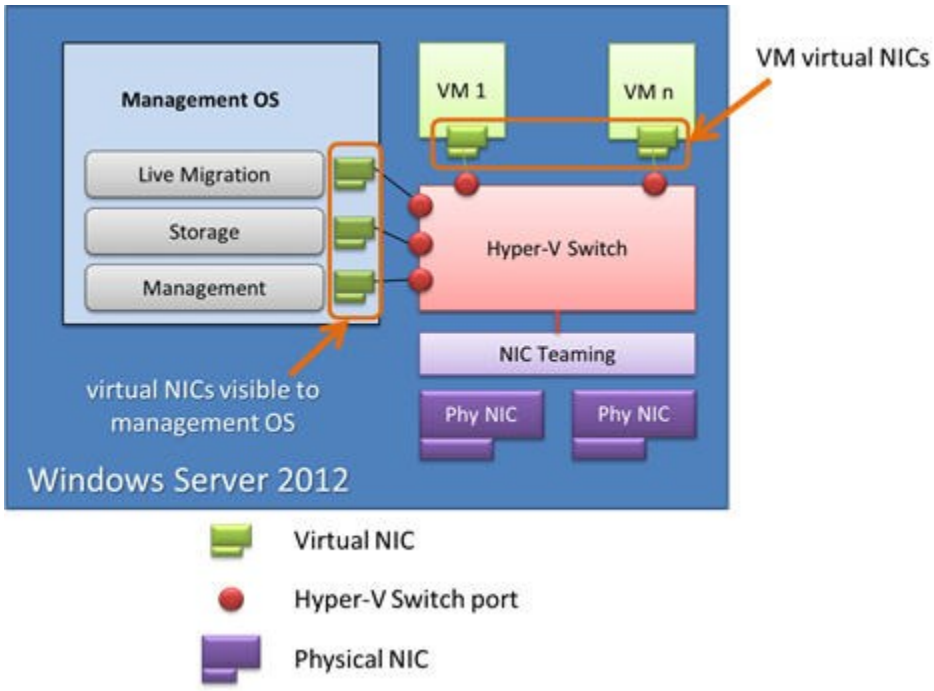

To create the virtual NICS for the management operating system, type the following commands at the Windows PowerShell prompt, and then press ENTER:

```
Add-VMNetworkAdapter –ManagementOS –Name Management

Add-VMNetworkAdapter –ManagementOS –Name Storage

Add-VMNetworkAdapter –ManagementOS –Name "Live Migration"
```

New security features for Hyper-V Virtual Switch

Hyper-V Virtual Switch includes the following new security features in Windows Server 2012.

- Port Access Control Lists (ACLs)
- MacAddressSpoofing
- RouterGuard
- DHCPGuard
- IPsec Task Offload (IpsecTO)

Most of these features provide you with the ability to specify the types of packets that are allowed or blocked when sent to or from the virtual machines that are connected to the virtual switch. The following sections provide additional details.

Port Access Control Lists (ACLs)

A port ACL is a rule that you can apply to a Hyper-V switch port. The rule specifies whether a packet is allowed or denied on the way into or out of the VM. ACLs have three elements with the following structure: **Local or Remote Address | Direction | Action**.

You can specify either a local address or a remote address in **Local or Remote Address**, but you cannot specify both. The value that you supply for **Local or Remote Address** can be an IPv4 address, an IPv6 address, or a media access control (MAC) address. Optionally you can use an IP address range if you provide the range prefix.

You can configure multiple port ACLs for a Hyper-V switch port. During operations, the port ACL whose rules match the incoming or outgoing packet is used to determine whether the packet is allowed or denied.

For example, to configure the virtual NIC attached to a port to always use either a specific MAC address or a broadcast address, you can create rules that are similar to the following:

Local or Remote Address	Direction	Action
12-34-56-78-9A-BC	Inbound	Allow
12-34-56-78-9A-BC	Outbound	Allow
FF-FF-FF-FF-FF-FF	Inbound	Allow
Any	Inbound	Deny
Any	Outbound	Deny

In this example, **Any** is a wildcard. Note that ACLs are evaluated on the longest prefix that is matched. Therefore the MAC address **12-34-56-78-9A-BC** with the **Allow** rule takes precedence over the wildcard rule that has an address of **Any**.

Bi-directional port ACLs

In many cases, you will want to apply the same permissions for both Inbound and Outbound traffic in the same port ACLs. For example, you want to allow both Inbound and Outbound traffic; or you want to deny both Inbound and Outbound traffic.

Because of this, there are bi-directional port ACLs. These ACL entries save some administrative time by allowing you to configure Inbound and Outbound permissions with one port ACL entry. You can accomplish the same effect as the first example by using bi-directional port ACLs in the following way:

Local or Remote Address	Direction	Action
12-34-56-78-9A-BC	Both	Allow
FF-FF-FF-FF-FF-FF	Inbound	Allow
Any	Both	Deny

The following commands demonstrate how you can create this configuration by using Windows PowerShell:

```
Add-VMNetworkAdapterAcl  -VMName MyVM -LocalMacAddress 12-34-56-78-9A--Direction Both -Action Allow

Add-VMNetworkAdapterAcl  -VMName MyVM -LocalMacAddress FF-FF-FF-FF-FF-FF  -Direction InBound -Action Allow

Add-VMNetworkAdapterAcl  -VMName MyVM -LocalMacAddress Any  -Direction Both -Action Deny
```

The above examples focus on ACLs that are defined by using the Local Address. You can also define ACLs by using the Remote Address. This is useful for blocking a particular address by using a remote address with a deny action.

Meter port ACLs

Meter port ACLs allow you to measure how much traffic is sent to or from the virtual NIC to a specified address range.

For example, the following meter port ACL measures the traffic that is sent from the VM named MyVM to the IPv4 subnet 192.168.0.0/16.

```
Add-VMNetworkAdapterAcl -VMName MyVM -RemoteIPAddress 192.168.0.0/16 -Direction Outbound -Action Meter
```

MacAddressSpoofing

MacAddressSpoofing allows you to specify whether a VM is allowed to change its source MAC address for outgoing packets.

If you have a VM for which you want to allow MacAddressSpoofing, you can enable it using the following cmdlet.

```
Set-VMNetworkAdapter -VMName MyVM -MacAddressSpoofing On
```

For all other VMs, you can block the ability to spoof MAC addresses with the following cmdlet.

```
Set-VMNetworkAdapter -VMName CustomerVM -MacAddressSpoofing Off
```

For more information, see Set-VMNetworkAdapter.

RouterGuard

RouterGuard allows you to specify whether the router advertisement and redirection messages from unauthorized VMs should be dropped.

If you have a VM for which you want to allow router advertisement and redirection messages from unauthorized VMs, you can use the following cmdlet.

```
Set-VMNetworkAdapter -VMName MyVM -RouterGuard Off
```

For all other VMs, you can block router advertisement and redirection messages from unauthorized VMs with the following cmdlet.

```
Set-VMNetworkAdapter –VMName CustomerVM –RouterGuard On
```

For more information, see Set-VMNetworkAdapter.

DHCPGuard

DHCPGuard allows you to specify whether DHCP server messages coming from a VM should be dropped. For VMs that are running an authorized instance of the DHCP server role, you can turn DHCPGuard off by using the following cmdlet:

```
Set-VMNetworkAdapter –VMName MyDhcpServer1 –DhcpGuard Off
```

For all other VMs that are not authorized DHCP servers, you can prevent them from becoming a rogue DHCP server by turning DHCPGuard on, using the following cmdlet.

```
Set-VMNetworkAdapter –VMName CustomerVM –DhcpGuard On
```

IPsec Task Offload (IpsecTO)

Many encryption algorithms are processor-intensive, which can slow the performance of VMs. Hyper-V Virtual Switch now provides VMs with the ability to use IPsecTO, which allows the VM to offload encryption processes to the NIC. Offloading the per-packet encryption operations from the VM to the NIC results in substantial CPU savings.

IPsec includes a Security Association (SA) with which it performs encryption, and when you enable IPsecTO, the VM offloads the SA to the NIC for processing. IPsecTO-capable NICs have a limited number of SAs that can be offloaded, so you can use Windows PowerShell to designate the number of SAs that the VM can offload to the NIC. To manage the number of SAs that can be offloaded by a VM you can use the following cmdlet.

```
Set-VMNetworkAdapter –VMName MyVM -IPsecOffloadMaximumSecurityAssociation 200
```

This cmdlet limits the number of SAs that can be offloaded by the VM named MyVM to 200. After the limit is reached, the Windows Networking stack in the VM intelligently manages the SAs so that the more active SAs are offloaded, while less active SAs are uploaded and performed by the VM.

Requirements for using IPsecTO

Following are the requirements for using IPsecTO.

- Only VMs running Windows Server® 2008 R2 and Windows Server 2012 are supported, because the VM's network stack must support IPsecTO.
- The physical NIC must also support IPsecTO.

Port Virtual Local Area Network (PVLAN) and Trunk Mode

Network isolation is related to security, but unlike IPsec - which encrypts the network traffic - isolation logically segments the traffic.

VLANs, however, suffer scalability issues. A VLAN ID is a 12-bit number, and VLANs are in the range 1-4095. In a multi-tenant data center, if you want to isolate each tenant by using a VLAN, configuration is complex and difficult. These scalability issues of VLANs are solved when you deploy Hyper-V Network Virtualization, where tenants each have multiple virtual subnets. However, a simple solution when each tenant only has a single VM is to use PVLAN.

PVLAN addresses some of the scalability issues of VLANs. PVLAN is a switch port property. With PVLAN there are two VLAN IDs, a primary VLAN ID and a secondary VLAN ID. A PVLAN may be in one of three modes.

PVLAN Mode	Description
Isolated	Communicates only with Promiscuous ports in the PVLAN
Promiscuous	Communicates with all ports in the PVLAN
Community	Communicates with ports in the same community and any promiscuous ports in the PVLAN

PVLAN can be used to create an environment where VMs may only interact with the Internet and not have visibility into other VMs' network traffic. To accomplish this put all VMs (actually their Hyper-V switch ports) into the same PVLAN in isolated mode. Therefore, using only two VLAN IDs, primary and secondary, all VMs are isolated from each other. The following PowerShell script puts a VM's switch port into PVLAN isolated mode.

```
Set-VMNetworkAdapterVlan –VMName MyVM –Isolated –PrimaryVlanId 10 –SecondaryVlanId 200
```

Important

When using any Hyper-V VLAN feature, the physical switches must also be properly configured with the appropriate VLAN IDs.

Trunk mode

In addition to PVLAN, Hyper-V Virtual Switch also provides support for VLAN trunk mode. Trunk mode provides network services or network appliances on a VM with the ability to see traffic from multiple VLANs.

In trunk mode, a switch port receives traffic from all VLANs that you configure in an allowed VLAN list. You can also configure a switch port that is connected to a VM - but is not bound to the underlying NIC - for trunk mode.

In the following example cmdlet, MyVM can send or receive traffic on any VLAN in the allowed list. If there is no VLAN specified in the packet, the packet is treated as if it is from VLAN 10.

```
Set-VMNetworkAdapterVlan –VMName MyVM –Trunk –AllowedVlanIdList 1-100 –NativeVlanId 10
```

Port Mirroring

With Port Mirroring, traffic sent to or from a Hyper-V Virtual Switch port is copied and sent to a mirror port. There are a range of applications for port mirroring - an entire ecosystem of network visibility companies exist that have products designed to consume port mirror data for performance management, security analysis, and network diagnostics. With Hyper-V Virtual Switch port mirroring, you can select the switch ports that are monitored as well as the switch port that receives copies of all the traffic.

The following examples configure port mirroring so that all traffic that is sent and received by both MyVM and MyVM2 is also sent to the VM named MonitorVM.

```
Set-VMNetworkAdapter –VMName MyVM –PortMirroring Source

Set-VMNetworkAdapter –VMName MyVM2 –PortMirroring Source

Set-VMNetworkAdapter –VMName MonitorVM –PortMirroring Destination
```

Receive Side Scaling (RSS) and Dynamic Virtual Machine Queue (dVMQ)

In networking it is important to resolve circumstances where network traffic is blocked or slowed down, thereby causing latency.

For native traffic, Receive Side Scaling (RSS) processes incoming network traffic so that it isn't slowed down by a single CPU. RSS processes the IP source and destination fields and TCP source and destination ports to spread the receive traffic across multiple CPU cores.

For receive network traffic coming externally from the server and into the Hyper-V Virtual Switch, Dynamic Virtual Machine Queue (dVMQ) performs a function similar to RSS. With dVMQ, the destination MAC address is hashed to put the traffic destined for a virtual NIC into a specific queue. The interrupts to the CPU cores are also distributed to avoid being slowed by a single CPU core. If your VMs on a Hyper-V Virtual Switch receive a lot of external network traffic, it's a good idea to use dVMQ.

dVMQ also includes dynamic load balancing. Previously, the MAC hashing was done statically, and it was difficult to manage dVMQ. Management of dVMQ is now simple – it is enabled by default, and no other management steps are required.

If for some reason you have disabled dVMQ, you can enable it again by using the following cmdlet. The NIC in this example is a physical NIC that is bound to the virtual switch, and it is named GuestTrafficNic.

```
Enable-NetAdapterVmq GuestTrafficNic
```

For more information, see Enable-NetAdapterVmq.

Create Security Policies with Extended Port Access Control Lists for Windows Server 2012 R2

One of the new features for the Hyper-V Virtual Switch in Windows Server® 2012 R2 is extended port Access Control Lists (ACLs). You can configure extended ACLs on the Hyper-V Virtual Switch to allow and block network traffic to and from the virtual machines (VMs) that are connected to the switch via virtual network adapters.

This topic provides information about new capabilities that are provided by extended ACLs, and contains the following sections.

- Detailed ACL rules
- Stateful ACL rules

Detailed ACL rules

Hyper-V Virtual Switch extended ACLs allow you to create more detailed rules than those possible in Windows Server® 2012, and you can apply these rules to individual VM network adapters that are connected to the Hyper-V Virtual Switch. This ability allows Enterprises and Cloud Service Providers (CSPs) to address network-based security threats in a multitenant shared server environment.

With extended ACLs, rather than having to create broad rules that block or allow all traffic from all protocols to or from a VM, you can now block or allow the network traffic of individual protocols that are running on VMs. You can create extended ACL rules in Windows Server 2012 R2 that include the following 5-tuple set of parameters: source IP address, destination IP address, protocol, source port, and destination port. In addition, each rule can specify network traffic direction (in or out), and the action the rule supports (block or allow traffic).

For example, you can configure port ACLs for a VM to allow all incoming and outgoing HTTP and HTTPS traffic on port 80, while blocking the network traffic of all other protocols on all ports.

This ability to designate the protocol traffic that can or cannot be received by tenant VMs provides flexibility when you configure security policies.

Configuring ACL rules with Windows PowerShell

To configure an extended ACL, you must use the Windows PowerShell command **Add-VMNetworkAdapterExtendedAcl**. This command has four different syntaxes, with a distinct use for each syntax:

1. Add an extended ACL to all of the network adapters of a named VM – which is specified by the first parameter, -VMName. Syntax:

   Note

If you want to add an extended ACL to one network adapter rather than all, you can specify the network adapter with the parameter –VMNetworkAdapterName.

```
    Add-VMNetworkAdapterExtendedAcl [-VMName] <string[]> [-
Action] <VMNetworkAdapterExtendedAclAction> {Allow | Deny}
    [-Direction] <VMNetworkAdapterExtendedAclDirection> {Inbound
| Outbound} [[-LocalIPAddress] <string>]
    [[-RemoteIPAddress] <string>] [[-LocalPort] <string>] [[-
RemotePort] <string>] [[-Protocol] <string>] [-Weight]
    <int> [-Stateful <bool>] [-IdleSessionTimeout <int>] [-
IsolationID <int>] [-Passthru] [-VMNetworkAdapterName
    <string>] [-ComputerName <string[]>] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

2. Add an extended ACL to a specific virtual network adapter on a specific VM. Syntax:

```
    Add-VMNetworkAdapterExtendedAcl [-VMNetworkAdapter]
<VMNetworkAdapterBase[]> [-Action]
    <VMNetworkAdapterExtendedAclAction> {Allow | Deny} [-
Direction] <VMNetworkAdapterExtendedAclDirection> {Inbound |
    Outbound} [[-LocalIPAddress] <string>] [[-RemoteIPAddress]
<string>] [[-LocalPort] <string>] [[-RemotePort]
    <string>] [[-Protocol] <string>] [-Weight] <int> [-Stateful
<bool>] [-IdleSessionTimeout <int>] [-IsolationID
    <int>] [-Passthru] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

3. Add an extended ACL to all virtual network adapters that are reserved for use by the Hyper-V host management operating system.

Note

If you want to add an extended ACL to one network adapter rather than all, you can specify the network adapter with the parameter –VMNetworkAdapterName.

```
    Add-VMNetworkAdapterExtendedAcl [-Action]
<VMNetworkAdapterExtendedAclAction> {Allow | Deny} [-Direction]
    <VMNetworkAdapterExtendedAclDirection> {Inbound | Outbound}
[[-LocalIPAddress] <string>] [[-RemoteIPAddress]
    <string>] [[-LocalPort] <string>] [[-RemotePort] <string>]
[[-Protocol] <string>] [-Weight] <int> -ManagementOS
```

```
    [-Stateful <bool>] [-IdleSessionTimeout <int>] [-IsolationID
<int>] [-Passthru] [-VMNetworkAdapterName <string>]
    [-ComputerName <string[]>] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

4. Add an extended ACL to a VM object that you have created in Windows PowerShell, such as **$vm = get-vm "my_vm"**. In the next line of code you can run this command to create an extended ACL with the following syntax:

```
    Add-VMNetworkAdapterExtendedAcl [-VM] <VirtualMachine[]> [-
Action] <VMNetworkAdapterExtendedAclAction> {Allow |
    Deny} [-Direction] <VMNetworkAdapterExtendedAclDirection>
{Inbound | Outbound} [[-LocalIPAddress] <string>]
    [[-RemoteIPAddress] <string>] [[-LocalPort] <string>] [[-
RemotePort] <string>] [[-Protocol] <string>] [-Weight]
    <int> [-Stateful <bool>] [-IdleSessionTimeout <int>] [-
IsolationID <int>] [-Passthru] [-VMNetworkAdapterName
    <string>] [-WhatIf] [-Confirm]  [<CommonParameters>]
```

Detailed ACL rule examples

Following are several examples of how you can use the **Add-VMNetworkAdapterExtendedAcl** command to configure extended port ACLs and to create security policies for VMs.

- Enforce application-level security
- Enforce both user-level and application-level security
- Provide security support to a non-TCP/UDP application

Note

The values for the rule parameter **Direction** in the tables below are based on traffic flow to or from the VM for which you are creating the rule. If the VM is receiving traffic, the traffic is inbound; if the VM is sending traffic, the traffic is outbound. For example, if you apply a rule to a VM that blocks inbound traffic, the direction of inbound traffic is from external resources to the VM. If you apply a rule that blocks outbound traffic, the direction of outbound traffic is from the local VM to external resources.

Enforce application-level security

Because many application servers use standardized TCP/UDP ports to communicate with client computers, it is easy to create rules that block or allow access to an application server by filtering traffic going to and coming from the port designated to the application.

For example, you might want to allow a user to login to an application server in your datacenter by using Remote Desktop Connection (RDP). Because RDP uses TCP port 3389, you can quickly set up the following rule:

Source IP	Destination IP	Protocol	Source Port	Destination Port	Direction	Action
*	*	TCP	*	3389	In	Allow

Following are two examples of how you can create rules with Windows PowerShell commands. The first example rule blocks all traffic to the VM named “ApplicationServer.” The second example rule, which is applied to the network adapter of the VM named “ApplicationServer,” allows only inbound RDP traffic to the VM.

Note

When you create rules, you can use the **–Weight** parameter to determine the order in which the Hyper-V Virtual Switch processes the rules. Values for **–Weight** are expressed as integers; rules with a higher integer are processed before rules with lower integers. For example, if you have applied two rules to a VM network adapter, one with a weight of 1 and one with a weight of 10, the rule with the weight of 10 is applied first.

```
Add-VMNetworkAdapterExtendedAcl –VMName “ApplicationServer” –Action “Deny” –Direction “Inbound” –Weight 1

Add-VMNetworkAdapterExtendedAcl –VMName “ApplicationServer” –Action “Allow” –Direction “Inbound” –LocalPort 3389 –Protocol “TCP” –Weight 10
```

Enforce both user-level and application-level security

Because a rule can match a 5-tuple IP packet (Source IP, Destination IP, Protocol, Source Port, and Destination Port), the rule can enforce a more detailed security policy than a Port ACL.

For example, if you want to provide DHCP service to a limited number of client computers using a specific set of DHCP servers, you can configure the following rules on the Windows Server 2012 R2 running Hyper-V where the user VMs are hosted:

Source IP	Destination IP	Protocol	Source Port	Destination Port	Direction	Action
*	255.255.255.255	UDP	*	67	Out	Allow
*	10.175.124.0/25	UDP	*	67	Out	Allow
10.175.124.0/25	*	UDP	*	68	In	Allow

Following are examples of how you can create these rules with Windows PowerShell commands.

```
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Deny” –Direction “Outbound”
–Weight 1
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Allow” –Direction
“Outbound” –RemoteIPAddress 255.255.255.255 –RemotePort 67 –Protocol “UDP”–Weight 10
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Allow” –Direction
“Outbound” –RemoteIPAddress 10.175.124.0/25 –RemotePort 67 –Protocol “UDP”–Weight 20
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Allow” –Direction “Inbound”
–RemoteIPAddress 10.175.124.0/25 –RemotePort 68 –Protocol “UDP”–Weight 20
```

Provide security support to a non-TCP/UDP application

While most network traffic in a datacenter is TCP and UDP, there is still some traffic that utilizes other protocols. For example, if you want to permit a group of servers to run an IP-multicast application that relies on Internet Group Management Protocol (IGMP), you can create the following rule.

Note
IGMP has a designated IP protocol number of 0x02.

Source IP	Destination IP	Protocol	Source Port	Destination Port	Direction	Action
*	*	0x02	*	*	In	Allow
*	*	0x02	*	*	Out	Allow

Following is an example of how you can create these rules with Windows PowerShell commands.

```
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Allow” –Direction “Inbound”
–Protocol 2 –Weight 20
Add-VMNetworkAdapterExtendedAcl –VMName “ServerName” –Action “Allow” –Direction
“Outbound” –Protocol 2 –Weight 20
```

Stateful ACL rules

Another new capability of extended ACLs allows you to configure stateful rules. A stateful rule filters packets based on five attributes in a packet – Source IP, Destination IP, Protocol, Source Port, and Destination Port.

Stateful rules have the following capabilities:

- They always allow traffic and are not used to block traffic.

- If you specify that the value for the parameter **Direction** is inbound and traffic matches the rule, Hyper-V Virtual Switch dynamically creates a matching rule that allows the VM to send outbound traffic in response to the external resource.
- If you specify that the value for the parameter **Direction** is outbound and traffic matches the rule, Hyper-V Virtual Switch dynamically creates a matching rule that allows the external resource inbound traffic to be received by the VM.
- They include a timeout attribute that is measured in seconds. When a network packet arrives at the switch and the packet matches a stateful rule, Hyper-V Virtual Switch creates a state so that all subsequent packets in both directions of the same flow are allowed. The state expires if there is no traffic in either direction in the period of time that is specified by the timeout value.

Following is an example of how you can use stateful rules.

Allow inbound remote server traffic only after it is contacted by the local server

In some cases, a stateful rule must be employed because only a stateful rule can keep track of a known, established connection, and distinguish the connection from other connections.

For example, if you want to allow a VM application server to initiate connections on port 80 to web services on the Internet, and you want the remote Web servers to be able to respond to the VM traffic, you can configure a stateful rule that allows initial outbound traffic from the VM to the Web services; because the rule is stateful, return traffic to the VM from the Web servers is also allowed. For security reasons, you can block all other inbound network traffic to the VM.

To achieve this rule configuration, you can use the settings in the table below.

Note

Due to formatting restrictions and the amount of information in the table below, the information is displayed differently than in previous tables in this document.

Parameter	Rule 1	Rule 2	Rule 3
Source IP	*	*	*
Destination IP	*	*	*
Protocol	*	*	TCP
Source Port	*	*	*
Destination Port	*	*	80
Direction	In	Out	Out
Action	Deny	Deny	Allow
Stateful	No	No	Yes
Timeout (in seconds)	N/A	N/A	3600

The stateful rule allows the VM application server to connect to a remote Web server. When the first packet is sent out, Hyper-V Virtual switch dynamically creates two flow states to allow all packets sent to and all returning packets from the remote Web server. When the flow of packets between the servers stops, the flow states time out in the designated timeout value of 3600 seconds, or one hour.

Following is an example of how you can create these rules with Windows PowerShell commands.

```
Add-VMNetworkAdapterExtendedAcl –VMName “ApplicationServer” –Action “Deny” –Direction “Inbound” –Weight 1

Add-VMNetworkAdapterExtendedAcl –VMName “ApplicationServer” –Action “Deny” –Direction “Outbound” –Weight 1

Add-VMNetworkAdapterExtendedAcl –VMName “ApplicationServer” –Action “Allow” –Direction “Outbound” 80 “TCP” –Weight 100 –Stateful –Timeout 3600
```

See Also

Hyper-V Virtual Switch Overview

Hyper-V Cmdlets in Windows PowerShell

Unified Tracing Overview

This topic provides an overview of Unified Tracing in Windows Server® 2012 and Windows Server® 2012 R2.

In Windows Server 2012, Unified Tracing introduced the new parameter **capturetype** for the **netsh trace** command. This parameter extends tracing capabilities and enables network administrators to more efficiently capture network traffic on both physical and virtual networks, making the process of troubleshooting network issues more effective and efficient.

In an environment where there are multiple virtual networks on a single physical computer, Unified Tracing can capture traffic between VMs on the same network, as well as between VMs on different virtual networks.

If you have a server that is running Hyper-V that has multiple virtual machines (VMs) and you want to analyze the network traffic between VMs, you can execute the **netsh trace** command using the **capturetype** parameter on the Hyper-V host. By starting the capture on the host machine, tracing is enabled on each of the VMs simultaneously, and you’re able to analyze the traffic between the VMs.

The parameter **capturetype** allows you to specify the type of traffic to capture:

- When you specify the value **physical** for the **capturetype** parameter, you can capture traffic on the physical computer, and specifically, traffic through the physical network adapter.

- When you specify the value **vmswitch** for the **capturetype** parameter, you can capture traffic that traverses Hyper-V Virtual Switches.
- When you specify the value **both** for the **capturetype** parameter, you can capture all traffic on the physical computer and that traverses the Hyper-V Virtual Switch.

For the purpose of backward compatibility, the default value of the **capturetype** parameter is **physical**.

When the value for the **capturetype** parameter is specified as **vmswitch**, only traffic traversing the switch is captured. If the same traffic traverses the physical network adapter (for example, the traffic is also sent to a different physical computer), then it is not captured again at the physical network adapter level.

For details about Netsh Trace in Windows Server 2012, see Netsh Commands for Network Trace in the Windows Server 2012 Technical Library.

See also

Following is a list of related technologies:

- Hyper-V Virtual Switch in the Windows Server® 2012 Technical Library.
- Using Netsh to Manage Traces in the Microsoft Developer Network (MSDN) library
- Netsh Commands for Network Trace in the Windows Server® 2008 and Windows Server® 2008 R2 Technical Library.
- Netsh technical reference in the Windows Server® 2008 and Windows Server® 2008 R2 Technical Library.

IP Address Management (IPAM) Overview

This topic provides overview information about the IP Address Management (IPAM) Server feature in Windows Server® 2012 and Windows Server 2012 R2. Also see the following:

- What's New in IPAM in Windows Server 2012 R2: Provides information about changes in functionality for IPAM in Windows Server 2012 R2.
- Walkthrough: Demonstrate IPAM in Windows Server 2012 R2: Provides a demonstration of new functionality for IPAM in Windows Server 2012 R2.
- Walkthrough: Demonstrate IPAM in Windows Server 2012: Provides procedures to set up IPAM in a test lab with a computer running Windows Server 2012.
- IPAM Planning and Design Guide: Provides an introduction to IPAM in Windows Server 2012 with design and planning information.
- IPAM Deployment Guide: Provides detailed, click-by-click deployment instructions to deploy IPAM in Windows Server 2012.
- IPAM Operations Guide: Provides operational, troubleshooting, and best practices guidance for IPAM in Windows Server 2012.

- Windows Server 2012 Virtual Labs includes an IPAM virtual lab you can use to test drive the IPAM feature in Windows Server 2012. See Managing Your Network Infrastructure with IP Address Management to start the virtual lab. Note: The virtual lab will take some time to start.
- Windows Server 2012 Virtual Labs also combines IPAM in a virtual lab with DNSSEC and DHCP failover in Windows Server 2012. See Building a Resilient Network Infrastructure to start the virtual lab. Note: The virtual lab will take some time to start.

Feature description

IPAM in Windows Server 2012 is a new built-in framework for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. IPAM provides for administration and monitoring of servers running Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS). IPAM includes components for:

- Automatic IP address infrastructure discovery: IPAM discovers domain controllers, DHCP servers, and DNS servers in the domains you choose. You can enable or disable management of these servers by IPAM.
- Custom IP address space display, reporting, and management: The display of IP addresses is highly customizable and detailed tracking and utilization data is available. IPv4 and IPv6 address space is organized into IP address blocks, IP address ranges, and individual IP addresses. IP addresses are assigned built-in or user-defined fields that can be used to further organize IP address space into hierarchical, logical groups.
- Audit of server configuration changes and tracking of IP address usage: Operational events are displayed for the IPAM server and managed DHCP servers. IPAM also enables IP address tracking using DHCP lease events and user logon events collected from Network Policy Server (NPS), domain controllers, and DHCP servers. Tracking is available by IP address, client ID, host name, or user name.
- Monitoring and management of DHCP and DNS services: IPAM enables automated service availability monitoring for Microsoft DHCP and DNS servers across the forest. DNS zone health is displayed, and detailed DHCP server and scope management is available using the IPAM console.

IPAM architecture

An IPAM server is a domain member computer.

Important
You cannot install the IPAM feature on an Active Directory domain controller.

There are two general methods to deploy IPAM servers:

1. Distributed: An IPAM server deployed at every site in an enterprise.
2. Centralized: One IPAM server in an enterprise.
3. Hybrid: A central IPAM server deployed with dedicated IPAM servers at each site.

The following example displays the distributed IPAM deployment method, with one IPAM server located at the corporate headquarters and also at each branch office. There is no communication or database sharing between different IPAM servers in the enterprise. If multiple IPAM servers

are deployed, you can customize the scope of discovery for each IPAM server, or filter the list of managed servers. A single IPAM server might manage a specific domain or location, perhaps with a second IPAM server configured as a backup.

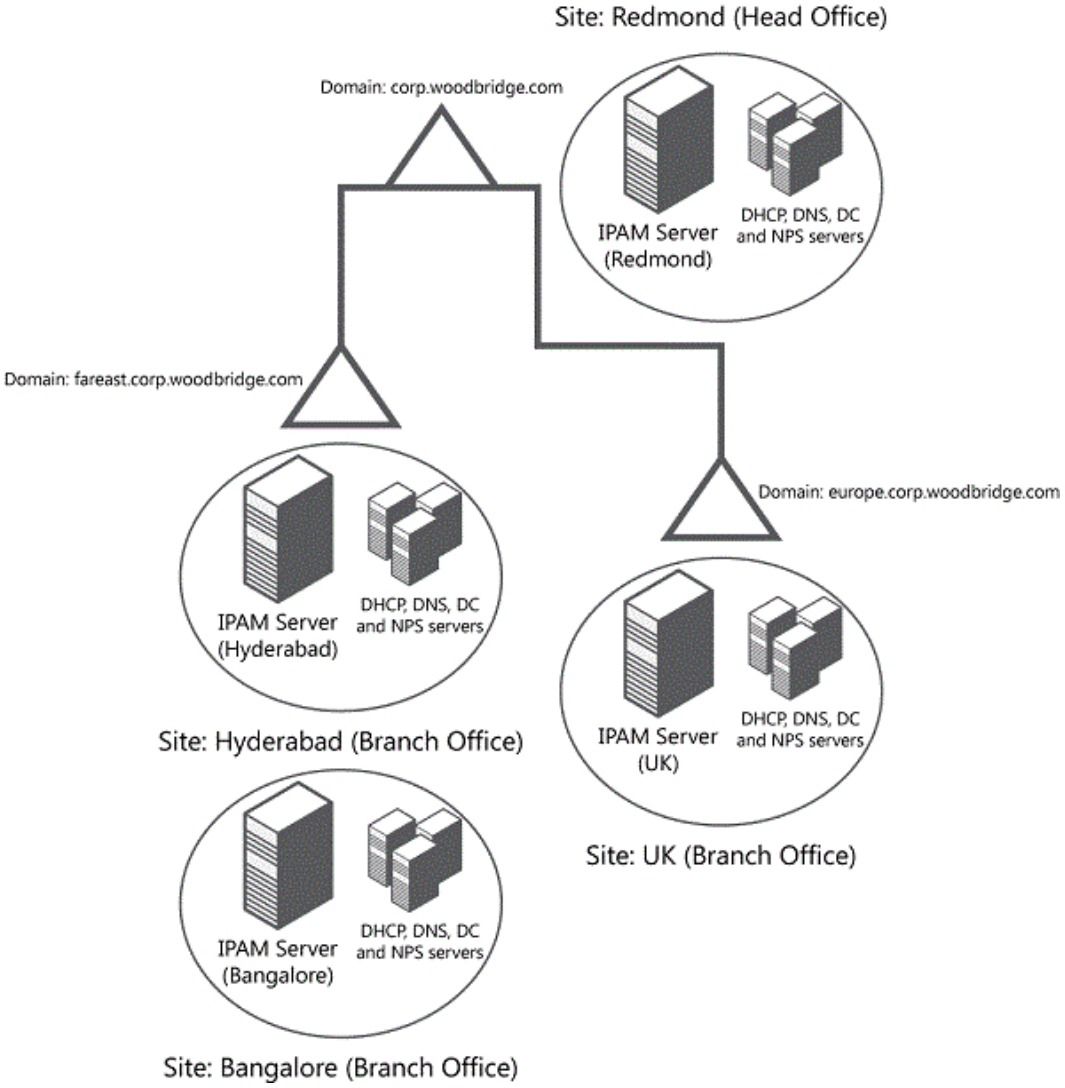

IPAM will periodically attempt to locate domain controllers, DNS, and DHCP servers on the network that are within the scope of discovery that you specify. You must choose whether these servers are managed by IPAM or unmanaged. In this way, you can select different groups of servers that are managed or not managed by IPAM.

To be managed by IPAM, security settings and firewall ports on a server must be configured to allow the IPAM server access so that it can perform required monitoring and configuration functions. You can configure these settings manually, or automatically using Group Policy Objects (GPOs). If you choose the automatic method, settings are applied when a server is marked as managed and settings are removed when it is marked as unmanaged.

The IPAM server will communicate with managed servers using an RPC or WMI interface. IPAM monitors domain controllers and NPS servers for IP address tracking purposes. In addition to monitoring functions, several DHCP server and scope properties can be configured from the IPAM console. Zone status monitoring and a limited set of configuration functions are also available for DNS servers. See the following figure.

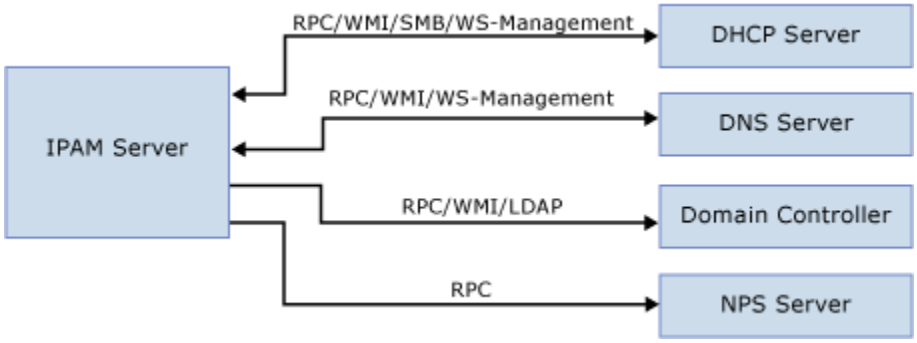

IPAM specifications

The scope of IPAM server discovery is limited to a single Active Directory forest. The forest itself may be comprised of a mix of trusted and untrusted domains. IPAM requires membership in an Active Directory domain, and relies on a functional network infrastructure environment to integrate with other server installations across the AD forest.

IPAM has the following specifications:

1. IPAM supports only Microsoft domain controllers, DHCP, DNS, and NPS servers running Windows Server® 2008 and above.
2. IPAM supports only domain joined DHCP, DNS and NPS servers in a single AD forest.
3. In its recommended configuration, IPAM is installed on a standalone server. You cannot install IPAM on a domain controller. If IPAM is installed on the same server with the DHCP Server role service, automatic discovery of DHCP servers on the network will be disabled.
4. IPAM does not support management and configuration of non-Microsoft network elements. However, you can use Windows PowerShell to import and manage IP address data from non-Microsoft devices.
5. IPAM in Windows Server 2012 does not support external databases. Only a Windows Internal Database is supported.
6. A single IPAM server has been tested to support up to 150 DHCP servers and 500 DNS servers.
7. A single IPAM server has been tested to support up to 40,000 DHCP scopes and 350 DNS zones.
8. IPAM has been tested to store 3 years of forensics data (IP address leases, host MAC addresses, user login/logoff information) for 100,000 users in a Windows Internal Database. Data is not purged automatically. An administrator must purge data manually as needed.
9. IP address utilization trends are provided only for IPv4.
10. IP address reclaiming support is provided for IPv4 and IPv6.

11. IPAM does not check for IP address consistency with routers and switches.
12. IPAM does not support auditing of IPv6 stateless address auto configuration on an unmanaged machine to track the user.
13. IPAM supports integration with System Center Virtual Machine Manager (VMM) using a Windows PowerShell script that is packaged and shipped with System Center VMM. This integration enables IPAM to display detailed utilization and inventory data for IP addresses and IP address ranges used in System Center VMM.

Practical applications

Monitoring and managing the IP address infrastructure on a corporate network is a critical part of network administration, and has become increasingly challenging as networks grow more dynamic and complex. Many IT administrators still track IP address allocation and utilization manually, using spreadsheets or custom database applications. This can be very time consuming and resource intensive, and is inherently prone to user error. IPAM in Windows Server 2012 provides a platform to manage the following IP administration needs.

1. Planning: IPAM replaces manual tools and scripts that can introduce added time, inconsistency and expense into the planning process when business expansions and alterations occur, or new technology and scenario adoptions are required.
2. Managing: IPAM provides a single management platform for IP address administration on the network. IPAM also allows for optimized utilization and capacity planning for DHCP and DNS services in distributed environment.
3. Tracking: IPAM enables tracking and forecasting of IP address utilization. As the demand for public IPv4 address space continues to grow in an environment with limited supply, this can be of critical importance to an organization.
4. Auditing: IPAM assists with compliance requirements such as HIPAA and Sarbanes-Oxley, and provides reporting for forensics and change management.

New and changed functionality

See What's New in IPAM in Windows Server 2012 R2.

Server Manager information

The installation of the IPAM Server feature can be performed through the Server Manager. The following features and tools are installed automatically when you install IPAM Server:

Feature or Tool	Description
Remote Server Administration Tools	DHCP and DNS Server Tools and IP Address Management (IPAM) Client provides for remotely managing DHCP, DNS and IPAM servers.

Feature or Tool	Description
Windows Internal Database	Windows Internal Database is a relational data store that can be used only by Windows roles and features.
Windows Process Activation Service	Windows Process Activation Service generalizes the IIS process model, removing the dependency on HTTP.
Group Policy Management	Group Policy Management is a scriptable Microsoft Management Console (MMC), providing a single administrative tool for managing Group Policy.
.NET Framework 4.5 Features	.NET Framework 4.5 provides a programming model for building and running applications designed for several different platforms.

See also

What's New in IPAM in Windows Server 2012 R2

IPAM Planning and Design Guide

IPAM Deployment Guide

IPAM Operations Guide

Walkthrough: Demonstrate IPAM in Windows Server 2012

Walkthrough: Demonstrate IPAM in Windows Server 2012 R2

What's New in IPAM in Windows Server 2012 R2

IP Address Management (IPAM) is feature which was first introduced in Windows Server® 2012 that provides highly customizable administrative and monitoring capabilities for the IP address infrastructure on a corporate network. IPAM in Windows Server® 2012 R2 includes many enhancements.

Role/Feature description

IPAM in Windows Server 2012 R2 is a built-in framework for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. IPAM provides for administration and monitoring of servers running Dynamic Host Configuration Protocol (DHCP)

and Domain Name System (DNS). For more information, see IP Address Management (IPAM) Overview

New and changed functionality

New functionality in IPAM for Windows Server 2012 R2 includes the following:

Feature/functionality	New or improved	Description
Role based access control	New	Role based access control enables you to customize the types of operations and access permissions for users and groups of users on specific objects.
Virtual address space management	New	IPAM streamlines management of physical and virtual IP address space in System Center Virtual Machine Manager.
Enhanced DHCP server management	Improved	Several new operations are available in IPAM to enhanced the monitoring and management of the DHCP Server service on the network.
External database support	New	In addition to Windows Internal Database (WID), IPAM also optionally supports the use of a Microsoft SQL database.
Upgrade and migration support	New	If you installed IPAM on Windows Server 2012, your data is maintained and migrated when you upgrade to Windows Server 2012 R2.
Enhanced Windows PowerShell support	Improved	Windows PowerShell support for IPAM is greatly enhanced to provide extensibility, integration, and automation support.

Role based access control

Role based access control provides you with the ability to customize roles, access scopes, and access policies. Thus, you have the ability to define and establish fine-grained control for users and groups, enabling them to perform a specific set of administrative operations on specific objects managed by IPAM.

Roles: A role is a collection of IPAM operations. You can associate a role with a user or group in Windows using an access policy. Eight built-in administrator roles are provided for convenience, but you can also create customized roles to meet your business requirements.

Access scopes: An access scope determines the objects that a user has access to. You can use access scopes to define administrative domains in IPAM. For example, you might create access scopes based on geographical location. By default, IPAM includes an access scope of Global. All other access scopes are subsets of the Global access scope. Users or groups that are assigned to the Global access scope have access to all objects in IPAM that are permitted by their assigned role.

Access Policies: An access policy combines a role with an access scope to assign permission to a user or group. For example, you might define an access policy for a user with a role of IP Block Admin and an access scope of Global\Asia. Therefore, this user will have permission to edit and delete IP address blocks that are associated to the Asia access scope. This user will not have permission to edit or delete any other IP address blocks in IPAM.

The following default access scope and roles are provided:

Type	Name	Description
Role	DNS record administrator	Manages DNS resource records
Role	IP address record administrator	Manages IP addresses but not IP address spaces, ranges, blocks, or subnets.
Role	IPAM administrator	Manages all settings and objects in IPAM
Role	IPAM ASM administrator	Completely manages IP addresses
Role	IPAM DHCP administrator	Completely manages DHCP servers
Role	IPAM DHCP reservations administrator	Manages DHCP reservations
Role	IPAM DHCP scope administrator	Manages DHCP scopes
Role	IPAM MSM administrator	Completely manages DHCP

Type	Name	Description
		and DNS servers
Access scope	Global	By default, all objects in IPAM are included in the global access scope. All additional scopes that are configured are subsets of the global access scope.

Virtual address space management

IPAM offers a unified, centralized administrative experience for network administrators to manage IP address space on a corporate network and in Microsoft-powered cloud networks. IPAM enables network administrators to completely streamline the IP address space administration of both physical (fabric) and virtual networks. The integration between IPAM and System Center 2012 R2 Virtual Machine Manager provides end-to-end IP address space automation for Microsoft-powered cloud networks. IPAM integration with Virtual Machine Manager enables a single IPAM server to detect and prevent IP address space conflicts, duplicates, and overlaps across multiple instances of Virtual Machine Manager deployed in the large datacenter.

To view virtual address space in IPAM, click the new **VIRTUALIZED ADDRESS SPACE** node in the upper navigation pane of the IPAM console.

Enhanced DHCP server management

DHCP server management with IPAM is greatly enhanced in Windows Server 2012 R2, including multiple new operations for DHCP scope and DHCP servers, and views for the following objects:

- DHCP failover
- DHCP policies
- DHCP superscopes
- DHCP filters
- DHCP reservations

External database support

During the IPAM provisioning process, you have the option of choosing a WID or Microsoft SQL Server for the IPAM database. With Microsoft SQL Server, the IPAM database can be collocated on the IPAM server, or it can be located on a remote computer. Support for SQL enables additional scalability, disaster recovery, and reporting scenarios.

Upgrade and migration support

The IPAM database can be migrated seamlessly when you upgrade from Windows Server 2012 to Windows Server 2012 R2.

Enhanced Windows PowerShell support

55 new Windows PowerShell cmdlets are available for IPAM in Windows Server 2012 R2. For more information, see IPAM Server Cmdlets in Windows PowerShell.

See also

Walkthrough: Demonstrate IPAM in Windows Server 2012 R2

Walkthrough: Demonstrate IPAM in Windows Server 2012 R2

IPAM in Windows Server® 2012 R2 provides a single console to plan, design and administer network services and IP address spaces, both physical and virtual. In this lab, you will learn more about how you can use IPAM in your organization to manage physical and virtual address space, delegate permissions in a multi-user environment, perform advanced DHCP configuration tasks, and leverage Windows PowerShell cmdlets for IPAM to automate routine operations.

Prerequisites

This lab has the following prerequisites:

Computer	Roles	Operating system	Notes
DC1	Domain controller, Active Directory-integrated DNS server, and secondary DHCP server.	Windows Server 2012 R2	The minimum required operating system to demonstrate DHCP failover is Windows Server 2012. You cannot install IPAM on a domain controller.
IPAM1	IPAM server and IPAM client.	Windows Server 2012 R2	IPAM1 must be running Windows Server 2012 R2 to demonstrate features in this lab.
DHCP1	Primary DHCP server.	Windows Server 2012	The minimum required operating system to

Computer	Roles	Operating system	Notes
		R2	demonstrate DHCP failover is Windows Server 2012. Do not install IPAM on the same server with DHCP. This computer is required to demonstrate DHCP failover but optional for other lab objectives.

Lab setup

The following table summarizes procedures used to configure this lab. Perform the procedures to configure DC1 first, DHCP1 second, and IPAM1 third. If you are not using DHCP1 to demonstrate DHCP failover operations then skip these procedures.

Computer	Procedures
DC1	1. Assign an IP address and DNS server address of **192.168.0.1**, a host name of DC1, then reboot. 2. Install Active Directory Domain Services, DNS Server, and DHCP Server roles. 3. Promote the server to a domain controller for the **contoso.com** domain. The computer will reboot automatically. 4. Sign in as CONTOSO\Administrator. 5. Complete DHCP configuration by authorizing the DHCP Server service in Active Directory. 6. Create a new user account in contoso.com named **user1**.
DHCP1	1. Assign an IP address of **192.168.0.2**, a DNS server address of **192.168.0.1**, a host name of DHCP1, join the contoso.com domain, and then reboot. 2. Sign in as CONTOSO\Administrator. 3. Install the DHCP Server role.

Computer	Procedures
	4. Complete DHCP configuration by authorizing the DHCP Server service in Active Directory.
IPAM1	1. Assign an IP address of **192.168.0.3**, a DNS server address of **192.168.0.1**, a host name of IPAM1, join the contoso.com domain, and then reboot. 2. Sign in as CONTOSO\Administrator. 3. Install the IPAM Server feature. IPAM Client will also be installed by default. 4. Provision the IPAM server to use Windows Internal Database (WID) and Group Policy Based provisioning with a GPO name prefix of IPAM1. 5. Use the **Invoke-IpamGpoProvisioning** cmdlet to create IPAM GPOs. Confirm the GPOs were successfully created. 6. Configure IPAM server discovery by adding the contoso.com domain, and then start server discovery. 7. When the server discovery task has completed, view the server inventory and mark DC1 and DHCP1 as managed servers. 8. Update Group Policy; verify that IPAM GPOs are applied, and that DC1 and DHCP1 are unblocked.

For detailed procedures to install and configure IPAM, see the IPAM Deployment Guide. Detailed steps for installing and configuring DHCP, DNS, and IPAM are also provided in Walkthrough: Demonstrate IPAM in Windows Server 2012.

Note

Procedures used for installation of IPAM and provisioning of managed servers procedures are identical in Windows Server 2012 R2 and Windows Server 2012 when IPAM is used with Windows Internal Database (WID). In Windows Server 2012 R2, you can also use SQL Server 2012 to manage the IPAM database. This option is not demonstrated in this lab.

Objectives

Objective 1: Demonstrate role based access control and delegated administration

Objective 2: Manage DHCP policy based assignment with IPAM

Objective 3: Automate IP address lifecycle management

Objective 4: Administer DHCP failover with IPAM

Objective 5: Manage DHCP MAC address filters with IPAM

Objective 6: Manage DHCP superscopes with IPAM

Objective 1: Demonstrate role based access control and delegated administration

Concepts

Role based access control is new in IPAM in Windows Server 2012 R2. Role based access control is comprised of roles, access scopes, and access policies:

Roles: A role is a collection of IPAM operations. You can associate a role with a user or group in Windows using an access policy. Several built-in roles are provided, but you can also create customized roles to meet your business requirements.

Access scopes: An access scope determines the objects that a user has access to. You can use access scopes to define administrative domains in IPAM. For example, you might create access scopes based on geographical location. By default, IPAM includes an access scope of Global. All other access scopes are subsets of the Global access scope. Users or groups that are assigned to the Global access scope have access to all objects in IPAM that are permitted by their assigned role.

Access policies: An access policy combines a role with an access scope to assign permission to a user or group. For example, you might define an access policy for user1 with a role of IP Block Admin and an access scope of Global\Asia. Therefore, user1 will have permission to edit and delete IP address blocks that are associated to the Asia access scope. This user will not have permission to edit or delete any other IP address blocks in IPAM.

Procedures

1. Create a new DHCP Scope named Lab DHCP Scope
2. Create a second DHCP Scope named Contoso DHCP Scope
3. Create new user role named DHCP Scope Editor with only Edit DHCP Scope permission
4. Create a new access scope named Test Lab under the Global assess scope
5. Set the Test Lab access scope on the Lab DHCP Scope
6. Create an access policy for the user contoso\user1 and assign the DHCP Scope Editor role and Test Lab access scope

7. Sign in as contoso\user1 and verify that this user is only able to perform the operations assigned by the DHCP Scope Editor role and only on the DHCP Scope: Test Lab DHCP Scope

Create a new DHCP Scope named Lab DHCP Scope

1. Click **DNS and DHCP Servers**, right-click DHCP1, and then click **Create DHCP Scope**.
2. Type **Lab DHCP Scope** next to **Scope name**.
3. Type **40.40.1.0** and **40.40.1.100** next to **Start IP address** and **End IP address**, respectively.
4. Leave the subnet mask unchanged, and click **OK**.

Create a second DHCP Scope named Contoso DHCP Scope

1. Click **DNS and DHCP Servers**, right-click DHCP1, and then click **Create DHCP Scope**.
2. Type **Contoso DHCP Scope** next to **Scope name**.
3. Type **50.50.1.0** and **50.50.1.100** next to **Start IP address** and **End IP address**, respectively.
4. Leave the subnet mask unchanged, and click **OK**.

Create new user role named DHCP Scope Editor with only Edit DHCP Scope permission

1. Click **ACCESS CONTROL**, right-click **Roles**, and then click **Add User Role**.
2. Type **DHCP Scope Editor** next to **Name**.
3. Expand **DHCP scope operations** to view the list of all available operations.
4. Select only the **Edit DHCP scope** operation, and then click **OK**.

Create a new access scope named Test Lab under the Global assess scope

1. Right-click **Access Scopes** and then click **Add Access Scope**.
2. Click **New**, type **Test Lab** next to **Name**, click **Add**, and then click **OK**.

Set the Test Lab access scope on the Lab DHCP Scope

1. Click **DHCP Scopes**, right-click the **Lab DHCP Scope**, and then click **Set Access Scope**.
2. Clear the **Inherit access scope** from parent checkbox.
3. Click the **Test Lab** access scope under **Global**, and then click **OK**.
4. Verify that **\Global\Test Lab** is displayed under **Access Scope**.

Create an access policy for the user contoso\user1 and assign the DHCP Scope Editor role and Test Lab access scope

1. Click **ACCESS CONTROL**, right-click **Access Policies**, and then click **Add Access Policy**.

2. Click **Add**, type **contoso\user1**, and then click **OK**.

   Tip
   If you didn't create a user1 account in contoso.com, create it before proceeding.
3. Under **Access Settings**, click **New**.
4. Under **Select role**, choose **DHCP Scope Editor**.
5. Under **Select the access scope for the role**, click **Test Lab**, click **Add Setting**, and then click **OK**.

Sign in as contoso\user1 and verify that this user is only able to perform the operations assigned by the DHCP Scope Editor role and only on the DHCP Scope: Test Lab DHCP Scope

1. Sign out on the IPAM server, and then sign in with the user1 account you created earlier.
2. Launch Server Manager and click **IPAM**.
3. Click **DHCP Scopes**, right-click the **Lab DHCP Scope**, and then click **Deactivate DHCP Scope**.
4. Review the error message displayed.
5. Right-click the **Contoso DHCP Scope** and click **Edit DHCP Scope**.
6. Change the **End IP address** to **50.50.1.101** and then click **OK**.
7. Review the error message displayed, and then click **Cancel**.
8. Edit the Lab DHCP Scope and change the End IP address to 40.40.1.101.
9. Verify that the operation is successful.
10. Sign out and then sign in again as contoso\administrator.

Objective 2: Manage DHCP policy based assignment with IPAM

Concepts

DHCP Policy Based Assignment (PBA) is a powerful feature for IPV4 networks that was first introduced in Windows Server 2012 for the DHCP server role. PBA enables control over the network and the devices accessing it. DHCP policies can be used to identify and group together devices based on attributes like MAC Address, Vendor Class and User Class. You can then control the leases and DHCP options that are assigned to these devices. For example, you can use DHCP policies to match the MAC address of clients and thus ensure that all virtual machines accessing your network are assigned addresses from a specific IP address range, or are assigned specific DHCP options.

DHCP policies can be configured at server level or the scope (subnet) level. Previously, policies were accessible only for individual DHCP servers by using the DHCP management console or with Windows PowerShell. Using IPAM in Windows Server 2012 R2, you can create and manage

policies centrally across multiple DHCP servers. You can create a policy for multiple servers or scopes in a single operation. You can also copy policies from one server or scope to another.

The following procedures demonstrate how IPAM enables you to centrally manage DHCP policies in an efficient manner.

Procedures

1. Configure new DHCP policies and import existing DHCP policies
2. Manage existing DHCP policies in IPAM

Configure new DHCP policies and import existing DHCP policies

1. Click **DNS and DHCP Servers**.
2. Right-click DHCP1 and then click **Configure DHCP Policy**.
3. Next to **Name**, type **Printer Policy**.
4. Under **Policy Conditions**, click **New**.
5. Under **New Condition**, next to **Criteria**, choose **Fully Qualified Domain Name**.
6. Next to **Value**, type **PRN**, select the **Use wildcard** checkbox, choose **Append** from the drop-down list, click **Add**, click **Add Condition**, and then click **OK**.
7. In the **DNS and DHCP Servers** view, choose **DHCP** next to **Server Type** and then choose **Policies** next to View.
8. In the **Details View**, click the **Conditions** tab to view details about the policy.
9. Choose **Server Properties** next to **View**, right-click DC1, and then click **Import DHCP Policy**.
10. Choose **dhcp1.contoso.com** next to **Select server**, choose the **Printer Policy** to import, and then click **OK**.
11. Choose **Policies** next to **View** and verify that **Printer Policy** is present on both DC1 and DHCP1.

Manage existing DHCP policies in IPAM

1. Create another policy on DHCP1 called **Workgroup Policy** using the **Fully Qualified Domain Name** criteria and the **Is Single Label** operator.
2. Using the **Policies** view, right-click **Workgroup Policy** and then click **Edit DHCP Policy**.
3. Under **DNS Dynamic Updates**, next to **Enable DNS dynamic updates**, choose **Yes**.
4. Next to **Enable name protection**, choose **No**.
5. Next to **Dynamically update DNS records**, choose **Always**.
6. Next to **Discard DNS records when lease is deleted**, choose **Yes**.
7. Next to **Dynamically update DNS records for DHCP clients that do not request updates**, choose **Yes**.
8. Next to **Disable dynamic updates for DNS PTR records**, choose **No**.
9. Next to **Register DHCP clients using a different DNS suffix**, choose **Yes**.

10. Type **contoso.com** next to **Use the following DNS suffix**, and then click **OK**.
11. Right-click **Workgroup Policy** and then click **Move Up Processing Order**.

Objective 3: Automate IP address lifecycle management

Concepts

New Windows PowerShell cmdlets for IPAM make it easier to automate IP address lifecycle management. A rich set of cmdlets are available that enable you to perform all management functions for your IP addresses, ranges and blocks. You can leverage these cmdlets to write scripts and integrate IPAM with various physical and virtual devices or servers. This saves time, eliminates manual intervention, and reduces operating cost.

Procedures

1. Create an IP address range in IPAM
2. View IP address ranges in IPAM
3. Track utilization of IP ranges
4. Find a free IP address and assign it to a new printer
5. Deprovision the added printer and reclaim the IP address

Create an IP address range in IPAM

1. Type the following command at a Windows PowerShell prompt and press ENTER:

   ```
   PS C:\>Add-IpamRange –NetworkId 192.168.0.0/24 –
   StartIPAddress 192.168.0.1 –EndIPAddress 192.168.0.254 -
   CreateSubnetIfNotFound
   ```

   This command creates an IP address range for servers in the test lab. Also use the next command to add individual IP addresses for the three computers used in this lab.
2. Type the following command at a Windows PowerShell prompt and press ENTER:

   ```
   PS C:\>Add-IpamAddress –IpAddress 192.168.0.1
   PS C:\>Add-IpamAddress –IpAddress 192.168.0.2
   PS C:\>Add-IpamAddress –IpAddress 192.168.0.3
   ```

View IP address ranges in IPAM

1. Type the following command at a Windows PowerShell prompt and press ENTER:

   ```
   PS C:\>Get-IpamRange –AddressFamily IPv4 –AddressCategory
   Private
   ```
2. Verify that the IP address range that was recently added is displayed. Also try specifying an AddressCategory of Public to view the two public IP address ranges used for DHCP

scopes.

3. To view more detailed information, type the following commands at a Windows PowerShell prompt and press ENTER:

```
PS C:\>$a= Get-IpamRange –AddressFamily IPv4 –AddressCategory
Private
PS C:\>$a[0]|fl *
```

Track utilization of IP ranges

1. Type the following command at a Windows PowerShell prompt and press ENTER:

```
PS C:\>Get-IpamRange –AddressFamily IPv4 –AddressCategory
Private|where-object {$_.PercentageUtilized –gt 80}
```

   This command will display IP address ranges where the utilization is over 80%.

2. Type the following commands at a Windows PowerShell prompt and press ENTER:

```
PS C:\> Get-IpamRange –AddressFamily IPv4 –AddressCategory
Private|where-object {$_.PercentageUtilized –gt 80}|Export-
Csv –Path
"C:\Users\Administrator\Desktop\OverUtilizedRanges.csv" –
NoTypeInformation -Force
PS C:\>notepad
"C:\Users\Administrator\Desktop\OverUtilizedRanges.csv"
```

   This will open a notepad file containing all the over-utilized public IPv4 address ranges (currently blank).

3. Type the following command at a Windows PowerShell prompt and press ENTER:

```
PS C:\> Get-IpamRange –AddressFamily IPv4 –AddressCategory
Private|where-object {$_.Overlapping –eq "True"}
```

   The previous command will display any overlapping IP address ranges, if they exist.

Find a free IP address and assign it to a new printer

1. Type the following command at a Windows PowerShell prompt and press ENTER:

```
PS C:\>$range = Get-IpamRange –StartIPAddress 192.168.0.1 -
EndIPAddress 192.168.0.254
PS C:\>Find-IpamFreeAddress –InputObject $range -
TestReachability
```

2. Confirm that the first available IP address is 192.168.0.4.

3. Type the following command at a Windows PowerShell prompt and press ENTER:

```
PS C:>$range = Get-IpamRange –StartIPAddress 192.168.0.1 -
EndIPAddress 192.168.0.254
PS C:\>$freeip = Find-IpamFreeAddress –InputObject $range –
```

```
TestReachability
PS C:\>$ip = Add-IpamAddress –IpAddress $freeIp.Address –
ManagedByService $range.ManagedByService –ServiceInstance
$range.ServiceInstance –DeviceType Printer –AssignmentType
Dynamic –MacAddress ”AA-BB-CC-DD-EE-FF” –ReservationServer
$range.DhcpServerName –ReservationName “B3_F1_Printer_HP” –
ReservationType Both –ReservationDescription “Reservation for
printer on first floor of building 3” -PassThru
PS C:\>$ip|fl *
```

4. Verify that the 192.168.0.4 IP address is assigned to a printer.

In this case, the IP address range does not correspond to a DHCP scope it is not possible to add this reservation to a DHCP scope. If this scope were present on DC1 or DHCP2, you can use the Add-DhcpServerv4Reservation cmdlet to add this IP address. See the following example:

```
PS C:>Add-DhcpServerv4Reservation –ComputerName $ip.ReservationServer –IPAddress
$ip.IPAddress –ClientId $ip.MacAddress –ScopeId $ip.ReservationScopeDetails –Name
$ip.ReservationName –Description $ip.ReservationDescription -PassThru
```

Use the following command to delete the reservation:

```
PS C:\> Remove-DhcpServerv4Reservation –ComputerName $ip.ReservationServer –IPAddress
$ip.IPAddress -PassThru
```

Deprovision the added printer and reclaim the IP address

1. To delete the IP address from IPAM, type the following command at a Windows PowerShell prompt and press ENTER:

```
PS C:\>Remove-IpamAddress –InputObject $ip -Force
```

Objective 4: Administer DHCP failover with IPAM

Concepts

IPAM enables you to view, monitor, and manage DHCP failover relationships in the managed network. DHCP failover is a new feature introduced with Windows Server 2012. For more information about DHCP failover, see Step-by-Step: Configure DHCP for Failover.

Procedures

1. Create a failover relationship
2. View the details of a failover relationship
3. Edit a failover relationship

4. Replicate Scopes
5. Demonstrate auto-sync
6. Remove scopes from a failover relationship
7. Delete a failover relationship

Create a failover relationship

1. Click **DHCP Scopes** and choose **Scope Properties** next to **Current view**.
2. Select both the **Contoso DHCP Scope** and the **Lab DHCP Scope** (using CTRL or SHIFT), right-click these scopes and then click **Configure DHCP Failover**.
3. Next to **Partner server**, choose **DC1.contoso.com**.
4. Next to **Relationship name**, type **DHCP1-DC1**.
5. Next to **Relationship name**, type **DHCP1-DC1**.
6. Next to **Secret**, type a shared secret, for example: **secret**.
7. The **Load Balance** mode is chosen by default with 50% of leases being issued by each DHCP server. Click **OK** to accept these parameters.
8. Verify that two new scopes have been created with **DC1.contoso.com** displayed in the **Server Name** column.

View the details of a failover relationship

1. Click **DNS and DHCP Servers**, choose **DHCP** next to **Server Type**, and choose **Failover Relationships** next to **View**.
2. In the **Details View**, review information on the **Relationship Properties** tab and the **DHCP Scopes** tab.

Edit a failover relationship

1. Right-click the **DHCP1-DC1** failover relationship and click **Edit DHCP Failover Relationship**.
2. Under **Mode**, choose **Hot Standby**, and then click **Apply**.
3. Verify that **Success** is displayed under **Status**, and then click **OK**.
4. Verify that **Hot Standby** is displayed in the **Mode** column for this failover relationship.

Replicate Scopes

1. Right-click the **DHCP1-DC1** failover relationship and then click **Replicate DHCP Failover Relationship**.
2. In the **Replicate Failover Relationship** dialog box, choose **dhcp1.contoso.com** as the originating server for replication, and then click **OK**. This action replicates all scopes in the failover relationship.
3. Since no changes have been made to DHCP scopes on DC1 or DHCP1, the information displayed under **Replicate progress** will indicate that scopes are identical. Click **Close**.
4. Choose **Server Properties** next to **View**.

5. Right-click DHCP1 and then click **Replicate DHCP Server**. This action replicates all scopes in all DHCP failover relationships on the server. Click **OK** to continue with the replication process.
6. Review the information under **Replicate progress**, and then click **Close**.
7. Choose **Scope Properties** next to **View**.
8. Right click a single scope and then click **Replicate DHCP Scope**. You can also choose multiple scopes to replicate but all scopes must belong to the same DHCP server.
9. Click **OK** to confirm the replication and then click **Close**.

Demonstrate auto-sync

1. Right-click the column header in **Scope Properties** view and select the **Failover Config Sync Status** header.
2. Choose **Server Properties** next to **View**.
3. Choose **Server Properties** next to **View**.
4. Right-click DHCP1 and then click **Launch MMC**.
5. Click **dhcp1.contoso.com** once, then right-click **dhcp1.contoso.com**, point to **All Tasks**, and then click **Stop**.
6. In IPAM, right-click DHCP1, and click **Retrieve Server Data**.
7. When the **IPAM ServerAvailability** task has completed running, refresh the console view and verify that DHCP1 displays **Stopped** in the **Server Availability** column.
8. Choose **Scope Properties** next to **View**, right-click the **Lab DHCP Scope** on DC1 and then click **Edit DHCP Scope**.
9. Change the **Lease duration for DHCP clients** from 8 days to 7 days and then click **OK**. Since the scope is part of a DHCP failover relationship and the partner server is down, an error is displayed.
10. View the **Failover Config Sync Status** column and verify that **Out of sync** is displayed for the Lab DHCP Scope on DC1.
11. Start the DHCP Server service on DHCP1 and **Retrieve Server Data** again.
12. Refresh the IPAM console view and verify that DHCP1 displays **Running** again under **Server Availability**.
13. Choose **Scope Properties** view, right-click the **Lab DHCP Scope** on DC1, and then click **Replicate DHCP Scope**.
14. Under **Replicate progress**, verify that the new option values were replicated successfully, and then click **Close**. You also have the option of editing the scope on the partner server manually and then clearing the config sync error.

Remove scopes from a failover relationship

1. Click **DHCP Scopes**, right-click the column header, point to **Group by**, and then click **Server Name**. This is a useful way of displaying scopes on a per-server basis.
2. Right-click the **Lab DHCP Scope** on DHCP1 and then click **Remove DHCP Failover Configuration**. Click **OK** to confirm the removal. This will keep the scope on DHCP1 and

remove it from the failover partner (DC1).

Tip
Under **Progress of Failover configuration removal**, verify that all processes are successful, and then click **Close**.

3. Verify that no **Failover Relationship Name** is displayed for the **Lab DHCP Scope** on DHCP1 and the scope no longer exists on DC1.

Delete a failover relationship

1. Click **DNS and DHCP Servers** and switch to **Failover Relationships** view.
2. Right-click the **DHCP1-DC1** failover relationship and then click **Delete**.
3. Next to **Select server**, choose the server that will keep the scopes that are still associated to this failover relationship. By default, the primary DHCP server is selected (DHCP1).
4. Click **OK** and verify that the remaining scope was deleted from DC1.
5. Click **Close** and verify that the failover relationship was deleted.

Objective 5: Manage DHCP MAC address filters with IPAM

Concepts

Media access control (MAC) address based filtering or link layer-based filtering for DHCP enables you to control network access based on MAC addresses, providing a low-level security method. You can create MAC address based filters to specify which MAC addresses are allowed on the network and which are denied access.

A DHCP server maintains **Allow** and **Deny** lists of MAC addresses. If you add MAC addresses to the allow list and enable the list, only these MAC addresses will be granted an IP address by the DHCP server. If you add MAC addresses to deny list and enable the list, these MAC addresses will be denied service by the DHCP server. You can enable both allow and deny lists, in which case the deny list takes precedence. This means that the DHCP server provides DHCP services only to clients whose MAC addresses are in the allow list, provided that no corresponding matches are in the deny list.

You can use wildcards to allow or deny network access based on vendor MAC prefixes. Link layer filtering is currently available for IPv4 address only.

This section demonstrates how IPAM allows you to centrally manage MAC address filters.

Procedures

1. Add DHCP MAC address filters
2. Manage DHCP MAC address filters in IPAM

Add DHCP MAC address filters

1. Click **DNS and DHCP Servers**, choose **DHCP** next to **Server Type** and choose **Server Properties** next to **View**.
2. Verify that the **Server Properties** tab in **Details View** shows that both the **Allow** and **Deny MAC Address Filters** are currently **Disabled**.
3. Select one or both DHCP servers, right-click and then click **Edit DHCP Server Properties**.
4. Click **MAC Address Filters** and notice that you can enable or disable MAC address filter lists here. Click **OK**.
5. Right-click DHCP1 and then click **Add DHCP MAC Address Filter**.
6. Next to **Filter Type** choose **Deny**.
7. Next to **MAC Address** type **AA-BB-CC-***.
8. Click **Add MAC Address Filter** and then click **OK**.
9. Choose **Filters** next to View and verify that a **Deny** filter was added to **dhcp1.contoso.com** for the **MAC Address** of **AA-BB-CC-*-*-***.

Manage DHCP MAC address filters in IPAM

1. Right-click the MAC address filter created in the previous procedure and then click **Move to Allow**.
2. Verify that the **Filter Type** changes to **Allow**.
3. Right-click the filter and then click **Delete**.

Objective 6: Manage DHCP superscopes with IPAM

Concepts

Consider the scenario where the available IP addresses in currently active scopes are nearly depleted, but more computers are expected to be added to the network. In this scenario, you can use superscopes to allow a DHCP server to provide leases to clients from more than one scope on a single physical network. DHCP scopes in the same superscope can share IP addresses and gives leases to clients on each other's subnet.

Superscopes can also help you resolve deployment issues such as migrating clients to a new scope (for example, when renumbering the network). Superscopes are available for IPv4 address only.

This section demonstrates how IPAM enables you to create and manage DHCP superscopes.

Procedures

1. Adding scopes to a superscopes

2. Manage DHCP superscopes

Adding scopes to a superscopes

1. Click DNS and DHCP Servers, and choose DHCP next to Server Type and Scope Properties next to View.
2. On the **Scope Properties** tab in **Details View**, the names of superscopes are displayed next to **Superscope Name**. This is also displayed in the **Superscope Name** column. Verify that currently this field is blank for all scopes.
3. Right-click the **Lab DHCP Scope** and then click **Add to DHCP Superscope**.
4. Next to **Superscope name**, type **Lab Superscope** and then click **OK**.
5. Verify that **Lab Superscope** is displayed next to **Superscope Name**.
6. Right-click the **Contoso DHCP Scope** and then click **Add to DHCP Superscope**.
7. Choose **Use existing superscope**, select **Lab Superscope**, and then click **OK**.
8. To remove a scope, Right-click the scope and then click **Remove from DHCP Superscope** If all scopes are removed from a superscope, it is automatically deleted.

Manage DHCP superscopes

1. Next to **View**, choose **Superscope Properties**.
2. Review information in the **Details View** using the **Superscope Properties** tab and the **DHCP Scopes** tab.
3. Right-click a superscope to perform several management actions.

What's New in IPAM in Windows Server 2012

IP Address Management (IPAM) is an entirely new feature in Windows Server® 2012 that provides highly customizable administrative and monitoring capabilities for the IP address infrastructure on a corporate network.

Role/Feature description

IPAM in Windows Server 2012 is a new built-in framework for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. IPAM provides for administration and monitoring of servers running Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS). For more information, see IP Address Management (IPAM) Overview

New and changed functionality

The IPAM/DHCP Integration Module is now available on the Microsoft Script Center. This Windows PowerShell script imports DHCP leases and reservations into the IPAM database.

See also

Walkthrough: Demonstrate IPAM in Windows Server 2012

Walkthrough: Demonstrate IPAM in Windows Server 2012

IP Address Management (IPAM) in Windows Server® 2012 is a framework for discovering, monitoring, managing and auditing IP address space on a corporate network. IPAM provides the following features:

- Automatic IP address infrastructure discovery
- Highly customizable IP address space display, reporting, and management
- Configuration change auditing for DHCP and IPAM services
- Monitoring and management of DHCP and DNS services
- IP address lease tracking

In this guide

This guide provides step-by-step instructions for deploying IPAM in a test lab using three server computers and one client computer. Software and hardware requirements are provided, as well as an overview of IPAM.

Important
The following instructions are for configuring a test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Guide contents:

IPAM overview

- IPAM discovery
- IP address space management
- Multi-server management and monitoring
- Operational auditing and IP address tracking

IPAM architecture
- IPAM security groups
- IPAM tasks
- Privacy
- IPAM requirements

Scenario overview
- Hardware and software requirements

Configuring the test lab
- Configure DC1
- Configure DHCP1
- Configure Client1
- Configure IPAM1

IPAM demonstration
- Address space management
- Infrastructure monitoring and management
- Review audit logs and events

IPAM overview

The IPAM feature consists of four primary modules. The following sections provide a brief description of these modules.

IPAM discovery

IPAM discovery requires access to Active Directory in order to discover network infrastructure servers. This discovery is necessary to enable IPAM services. Discovery allows administrators to enumerate servers running Windows Server® 2008 or later with the DNS Server, DHCP Server and AD DS role services installed. Administrators can also manually add or delete servers to define a custom scope of administrative control. The scope of discovery can be modified in real-time by selecting or removing domains and specific server roles.

IP address space management

The IPAM address space management (ASM) feature provides the ability to efficiently view, monitor, and manage IP address space on the network. ASM supports IPv4 public and private addresses, and IP addresses can be dynamically issued on the network or provided as static IP addresses. Sorting can be based on custom fields, such as region, Regional Internet Registries (RIR), device type, or customer name. A network administrator can track IP address utilization and threshold-crossing status, or display utilization trends. IPAM ASM tools address the IP address space management problem in a growing distributed environment by ensuring better planning, accountability, and control. IPAM also enables an administrator to detect overlapping IP

address ranges defined on different DHCP servers, find free IP addresses within a range, create DHCP reservations, and create DNS records.

Multi-server management and monitoring

IPAM enables administrators to monitor and manage multiple DHCP servers and monitor multiple DNS servers spread across various regions from a centralized console. Administrative tasks are frequently repetitive across multiple servers. The ability to execute these tasks uniformly across servers reduces both the effort involved as well as the probability of error. The multi-server management (MSM) feature enables an administrator to easily edit and configure key properties of multiple DHCP servers and scopes across the organization. IPAM also facilitates monitoring and tracking of DHCP service status and utilization of DHCP scopes. IPAM also enables tagging of servers with built-in and user-defined custom field values and to visualize these servers and group them into logical groups and sub-groups. IPAM helps to monitor the health of a DNS zone on multiple DNS servers by displaying the aggregated status of a zone across all authoritative DNS servers. IPAM also tracks the service status of the DNS and DHCP servers on the network.

Operational auditing and IP address tracking

Auditing tools enable tracking potential configuration problems on IP infrastructure servers. IPAM provides the ability to view consolidated configuration changes on managed DHCP servers and the IPAM server. Details are tracked such as server name, user name, and the date and time a configuration change was made. IP address lease tracking is available to aid forensics investigations by collecting lease logs from DHCP, DC and NPS servers. IPAM enables history tracking for IP address leases and user logins. This enables tracking of IP address activity correlated with MAC addresses, user names, host names and other parameters.

IPAM architecture

An IPAM server is a domain member computer. You cannot install IPAM on an Active Directory domain controller.

There are three general methods to deploy IPAM servers:

1. **Distributed:** An IPAM server is deployed at every site in the enterprise.
2. **Centralized:** One IPAM server is deployed in the enterprise.

   Hybrid: A central IPAM server deployed with dedicated IPAM servers at each site.

There is no communication or database sharing between different IPAM servers in the enterprise. If multiple IPAM servers are deployed, you can customize the scope of discovery for each IPAM server, or filter the list of managed servers. A single IPAM server might manage a specific domain or location, perhaps with a second IPAM server configured as a backup.

IPAM will periodically attempt to locate network policy servers, domain controllers, DNS servers, and DHCP servers on the network that are within the scope of discovery that you specify. You must choose whether these servers are managed by IPAM or unmanaged. In this way, you can select different groups of servers that are managed or not managed by IPAM. To be managed by

IPAM, server security settings and firewall ports must be configured to allow the IPAM server access to perform required monitoring and configuration functions. You can choose to configure these settings manually, or automatically using Group Policy Objects (GPOs). If you choose the automatic method, then settings are applied when a server is marked as managed and settings are removed when it is marked as unmanaged. The IPAM server will communicate with managed servers using an RPC or WMI interface. IPAM monitors domain controllers and NPS servers for IP address tracking purposes. In addition to monitoring functions, several DHCP server and scope properties can be configured from using IPAM. Zone status monitoring and a limited set of configuration functions are also available for DNS servers.

IPAM security groups

The following local IPAM security groups are created when you install IPAM.

- **IPAM Users**: Members of this group can view all information in server discovery, IP address space, and server management. They can view IPAM and DHCP server operational events, but cannot view IP address tracking information.
- **IPAM MSM Administrators**: IPAM multi-server management (MSM) administrators have IPAM Users privileges and can perform IPAM common management tasks and server management tasks.
- **IPAM ASM Administrators**: IPAM address space management (ASM) administrators have IPAM Users privileges and can perform IPAM common management tasks and IP address space tasks.
- **IPAM IP Audit Administrators**: Members of this group have IPAM Users privileges and can perform IPAM common management tasks and can view IP address tracking information.
- **IPAM Administrators**: IPAM Administrators have the privileges to view all IPAM data and perform all IPAM tasks.

IPAM tasks

IPAM launches the following tasks upon installation with the specified periodicity. These tasks can be viewed in Task Scheduler by navigating to **Microsoft > Windows > IPAM**.

Task Name	Description	Default Frequency	Duration
AddressExpiry	Tracks IP address expiry state and logs notifications.	1 day	Indefinite
AddressUtilization	Collects IP address space usage data from DHCP servers to display current and historical utilization.	2 hours	Indefinite
Audit	Collects DHCP and	1 day	Indefinite

Task Name	Description	Default Frequency	Duration
	IPAM server operational events. Also collects events from domain controllers, NPS, and DHCP servers for IP address tracking.		
ServerAvailability	Collects service status information from DHCP and DNS servers.	15 minutes	Indefinite
ServerConfiguration	Collects configuration information from DHCP and DNS servers for display in IP address space and server management functions.	6 hours	Indefinite
ServerDiscovery	Automatically discovers the domain controllers, DHCP servers, and DNS servers in the domains you select.	1 day	Indefinite
ServiceMonitoring	Collects DNS zone status events from DNS servers.	30 minutes	Indefinite

Privacy

The IP address audit functionality in IPAM audit provides tracking of IP address, hostname and Client Identifier (MAC address in IPv4, DUID in IPv6) information of computers and devices on a network in addition to user login information. The IPAM server collects audit logs and events from DHCP servers, domain controllers and network policy servers, and stores the IP address, hostname, client identifier and user name of a network user in the IPAM database on the computer running the IPAM Server feature. An IPAM audit administrator or IPAM administrator can search logs based on IP address, client identifier, hostname, or user name.

Information collected, processed, or transmitted

IP addresses, client identifiers, and host names are collected from audit logs on DHCP servers that are managed by IPAM.

- User names and IP addresses are collected from events on domain controllers.
- User names and client identifiers are collected from events on network policy servers.
- No information is sent to Microsoft.

Audit control

IPAM is not enabled by default and must be installed as a server feature. When the IPAM Server feature is installed, IP address audit functionality is automatically enabled.

To disable IP address audit, start Task Scheduler on the IPAM server, navigate to Microsoft\Windows\IPAM and disable the audit task.

IPAM requirements

The scope of IPAM server discovery is limited to a single Active Directory forest. The forest may be comprised of a mix of trusted and untrusted domains. IPAM requires membership in an Active Directory domain, and is reliant on a prerequisite functional network infrastructure environment in order to integrate with existing DHCP, DNS, domain controller, and network policy server installations across the forest.

IPAM has the following specifications:

- IPAM supports only Microsoft DHCP, DNS, domain controllers, and network policy servers running Windows Server® 2008 and above.
- IPAM supports only domain joined servers in a single Active Directory forest.
- A single IPAM server can support up to 150 DHCP servers and 500 DNS servers.
- A single IPAM server can support up to 6000 DHCP scopes and 150 DNS zones.
- IPAM stores 3 years of forensics data (IP address leases, host MAC addresses, user login and logoff information) for 100,000 users in a Windows Internal Database. There is no database purge policy provided, and the administrator must purge the data manually as needed.
- IPAM does not support management and configuration of non-Microsoft network elements (such as WINS, DHCP relays, or proxies).
- IPAM supports only Windows Internal Database. No external database is supported.
- IP address utilization trends are provided only for IPv4.
- IP address reclaiming support is provided only for IPv4.
- No special processing is done for IPv6 stateless address auto configuration private extensions.
- No special processing for virtualization technology or virtual machine migration.
- IPAM does not check for IP address consistency with routers and switches.
- IPAM does not support auditing of IPv6 address (stateless address auto configuration) on an unmanaged machine to track the user.

Scenario overview

This test lab demonstrates IPAM functionality in Windows Server 2012. Three server computers and one client computer are used. See the following figure.

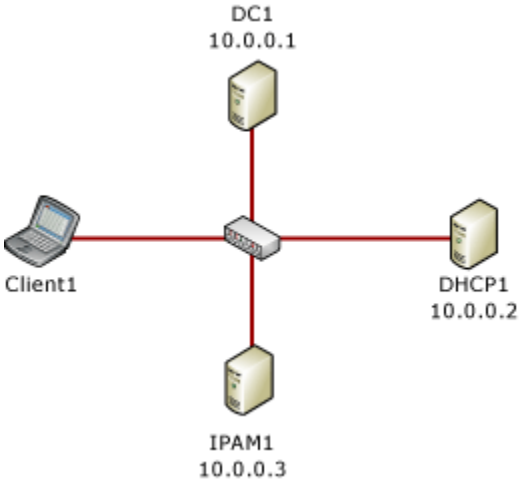

Hardware and software requirements

Three server computers and one client computer are required to complete the test lab.

Note

You can install DHCP on the same server with AD DS and DNS if desired and adjust procedures in the test lab accordingly. DHCP and DNS roles are separated in the test lab to demonstrate discovery and management of multiple servers providing different services on the network. The IPAM feature must be installed on a separate, domain member computer. A client computer is required to demonstrate IP address audit functionality.

The following are required components of the test lab:

1. The product disc or other installation media for Windows Server 2012.
2. Three computers that meet the minimum hardware requirements for Windows Server 2012.
3. The product disc or other installation media for Windows® 8.
4. One computer that meets the minimum hardware requirements for Windows 8.

Configuring the test lab

The following procedures provide instructions to install the operating system, configure TCP/IP, and add required role services and features on computers in the test lab.

1. Configure DC1
2. Configure DHCP1
3. Configure Client1

4. Configure IPAM1

Configure DC1

DC1 is a computer running Windows Server 2012, providing the following services:

- A domain controller for the contoso.com Active Directory domain.
- An authoritative DNS server for the contoso.com DNS zone.

Initial configuration of DC1 consists of the following steps:

- Install the operating system and configure TCP/IP on DC1
- Install Active Directory and DNS on DC1
- Create a domain administrator account

Additional tasks will be performed on DC1 during the demonstration portion of the test lab.

Install the operating system and configure TCP/IP on DC1

To install the operating system and configure TCP/IP on DC1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. Click **Start**, type **ncpa.cpl**, and then press ENTER. The **Network Connections** control panel will open.

   Tip
   The previous step demonstrates new functionality in Windows Server 2012 that enables you to search and run applications, settings, and files by clicking **Start** and then typing a search term. You can also open the **Network Connections** control panel by clicking next to **Wired Ethernet Connection** in Server Manager using the **Local Server** view. For more information, see Common Management Tasks and Navigation in Windows Server 2012 (http://go.microsoft.com/fwlink/p/?LinkId=242147).
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.1** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install Active Directory and DNS on DC1

DC1 will serve as the primary domain controller and DNS server for the contoso.com Active Directory domain.

To configure DC1 as a domain controller and DNS server

1. In the Server Manager navigation pane, click **Configure this local server**.
2. Under **PROPERTIES**, click the name next to **Computer name**. The **System Properties** dialog box will open.
3. On the **Computer Name** tab, click **Change** and then type **DC1** under **Computer name**.
4. Click **OK** twice, and then click **Close**.
5. When you are prompted to restart the computer, click **Restart Now**.
6. After restarting the computer, sign-in using the local Administrator account.
7. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
8. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **Active Directory Domain Services** checkbox.
9. When you are prompted to add required features, click **Add Features**.
10. Select the **DNS Server** checkbox.
11. When you are prompted to add required features, click **Add Features**.
12. Click **Next** four times, and then click **Install**.
13. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DC1** is displayed, and then click **Close**.
14. Click the Notification flag and then click **Promote this server to a domain controller**. See the following example.

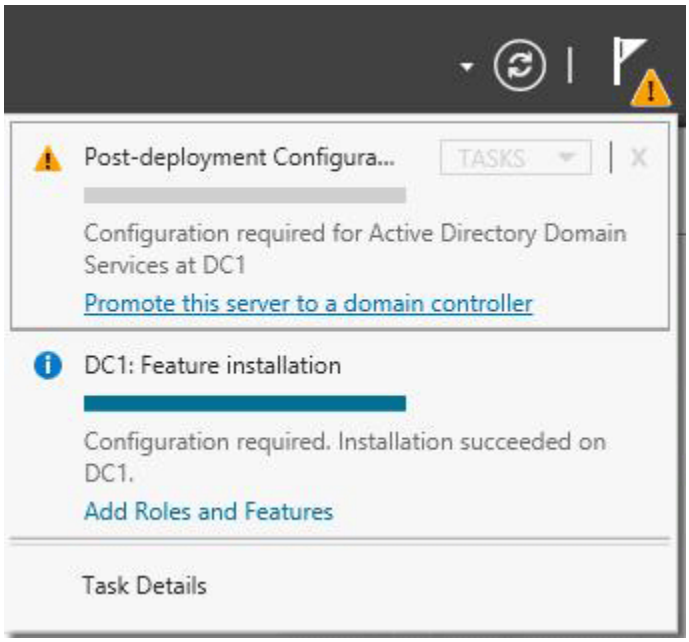

Note

There is a link displayed on the **Installation progress** page of the **Add Roles and Features Wizard** to promote the server to a domain controller after installation of AD DS is complete. However, if you close the **Installation progress** page, additional configuration tasks can always be accessed by clicking the Notification flag.

15. In the **Active Directory Domain Services Configuration Wizard**, on the **Deployment Configuration** page, choose **Add a new forest** and then next to **Root domain name**, type **contoso.com**.
16. Click **Next**, and then on the **Domain Controller Options** page, under **Type the Directory Services Restore Mode (DSRM) password**, type a password next to **Password** and **Confirm password**. Confirm that **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected, and then click **Next**.
17. Click **Next** five times and then click **Install**.

Tip

If **An error was detected in the DNS configuration** is displayed on the DNS Options page, you can ignore this message.

18. The computer will restart automatically to complete the installation process.
19. Sign in using the CONTOSO\Administrator account.

Create a domain administrator account

A user account that is a member of **Domain Admins** is required to complete the test lab.

Tip

You can use the CONTOSO\Administrator account in this test lab and skip creation of a domain administrator account if desired. This account has domain administrator privileges, and other privileges. However, it is a best practice to disable or rename this account. For more information, see Active Directory Best Practices(http://go.microsoft.com/fwlink/p/?LinkID=243071).

To create a domain administrator account

1. On the Server Manager menu bar, click **Tools**, and then click **Active Directory Users and Computers**.
2. In the **Active Directory Users and Computers** console tree, double-click **contoso.com**, right-click **Users**, point to **New**, and then click **User**.
3. In the **New Object – User** dialog box, type **user1** under **User logon name** and next to **Full name**, then click **Next**.
4. Next to **Password** and **Confirm password**, type a password for the user1 account.
5. Clear the checkbox next to **User must change password at next logon**, select the **Password never expires** checkbox, click **Next**, and then click **Finish**.

6. Double-click **user1** and then click the **Member Of** tab.
7. Click **Add**, type **domain admins** under **Enter the object names to select**, click **OK** twice, and then close the **Active Directory Users and Computers** console.
8. Click **Start**, click **Administrator**, and then click **Sign out**.
9. Sign in to the computer using the user1 credentials by clicking the left arrow next to **CONTOSO\Administrator** and then clicking **Other user**.

Configure DHCP1

DHCP2 is a computer running Windows Server 2012, providing the following services:

- A DHCP server.

Initial configuration of DHCP1 consists of the following steps:

- Install the operating system and configure TCP/IP on DHCP1
- Install and configure DHCP on DHCP1

Install the operating system and configure TCP/IP on DHCP1

Tip

The procedure below is identical to the steps used to install the operating system and configure TCP/IP on DC1, with the exception that DHCP1 is configured with an IP address of 10.0.0.2.

To install the operating system and configure TCP/IP on DHCP1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.
5. In the Server Manager navigation pane, click **Local Server** and then click the IP address next to **Wired Ethernet Connection**. The **Network Connections** control panel will open.
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.0.2** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install and configure DHCP on DHCP1

DHCP1 is a domain member server running the DHCP Server role service.

To install DHCP on DHCP1

1. In the Server Manager navigation pane, click **Local Server** and then click the name next to **Computer name**. The **System Properties** control panel will open.
2. On the **Computer Name** tab, click **Change** and then type **DHCP1** under **Computer name**.
3. Under **Member of**, select **Domain**, type **contoso.com**, and then click **OK**.
4. When you are prompted to provide credentials to join the domain, enter the credentials for the user1 account that was created previously and then click **OK**.
5. Confirm that computer name and domain changes were successful, click **OK** twice, and then click **Close**.
6. When you are prompted to restart the computer, click **Restart Now**.
7. After restarting the computer, sign-in using the CONTOSO\user1 account.
8. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
9. In the **Add Roles and Features Wizard**, click **Next** three times, and then on the **Select server roles** page select the **DHCP Server** checkbox.
10. When you are prompted to add required features, click **Add Features**.
11. Click **Next** three times, and then click **Install**.
12. Wait for the installation process to complete, verify on the **Installation progress** page that **Configuration required. Installation succeeded on DHCP1.contoso.com** is displayed, and then click **Close**.

To configure DHCP on DHCP1

1. On the Server Manager menu, click the Notification flag and then click **Complete DHCP configuration**.
2. In the DHCP Post-Install configuration wizard, click **Next** and then click **Commit**.
3. On the Server Manager menu bar, click **Tools** and then click **DHCP**. THE DHCP console opens.
4. In the DHCP console tree, navigate to **IPv4**. Right-click **IPv4** and then click **New Scope**. The **New Scope Wizard** opens.
5. Click **Next** and then type a name for the new scope next to **Name** (ex: Contoso-scope1).
6. Click **Next** and then in **IP Address Range**, type **10.0.0.1** next to **Start IP address**, type **10.0.0.254** next to **End IP address**, and type **24** next to **Length**. The value of Subnet mask will change automatically to **255.255.255.0**.
7. Click **Next**, and then in **Add Exclusions and Delay** type **10.0.0.1** under **Start IP address**, type **10.0.0.10** under **End IP address**, and then click **Add**. This allows the first ten IP addresses in the 10.0.0.0/24 subnet to be used for static addressing of servers on the network.
8. Click **Next** and then in **Lease Duration** under **Limited to** enter **0 Days**, **0 Hours**, and **2**

Minutes. This very short lease duration will simplify generation of more DHCP leases to review for the IP address auditing demonstration.

9. Click **Next** three times, and then in **Domain Name and DNS Servers**, verify that the **Parent domain** is **contoso.com** and **10.0.0.1** is listed as the only DNS server.
10. Click **Next** twice, and then in **Activate Scope** select **Yes, I want to activate this scope now**.
11. Click **Next**, and then click **Finish**.
12. Refresh the view in the DHCP console and verify that DHCP1 is authorized and that the Contoso-scope1 is active.

    Note: To review scopes on the current server using Windows PowerShell, right-click **Windows PowerShell**, click **Run as Administrator**, click **Yes** in the User Account Control alert that appears, and then type the following command at the Windows PowerShell prompt, and then press ENTER.

```
get-dhcpserverv4scope
```

```
PS C:\Windows\system32> get-dhcpserverv4scope


ScopeId         SubnetMask      Name            State    StartRange      EndRange
LeaseDuration
-------         ----------      ----            -----    ----------      --------
-------------
10.0.0.0        255.255.255.0   Contoso-scope1 Active   10.0.0.1        10.0.0.254
00:02:00
```

Configure Client1

Client1 is a computer running Windows® 8 that is acting as a DHCP client. Configuration of Client1 consists of the following steps:

- Install the operating system and configure TCP/IP on Client1
- Join Client1 to the contoso.com domain

During the demonstration portion of the test lab, Client1 will receive DHCP leases from DHCP1.

Install the operating system and configure TCP/IP on Client1

To install the operating system and configure TCP/IP on Client1

1. Start your computer using the Windows 8 product disc or other digital media.
2. When prompted, enter a product key and accept license terms.
3. When prompted to enter a computer name, type **Client1** and click **Next**.
4. Click **Use express settings**.
5. On the **Sign in to your PC** page, click **Don't want to sign in with a Microsoft account**

and then click **Local account**.

6. Next to **User name**, type **user1**, enter a password and password hint, and then click **Finish**.

Join Client1 to the contoso.com domain

In order for Client1 to receive Group Policy settings, it must be joined to the contoso.com domain.

To join Client1 to the contoso.com domain

1. Click **Start**, type **sysdm.cpl** and then press ENTER.
2. In the **System Properties** dialog box, click **Change**.
3. Under **Member of**, choose **Domain**, type **contoso.com** and then click **OK**.
4. When you are prompted to enter an account with permission to join the domain, provide the credentials for the user1 account, and then click **OK**.
5. Confirm that **Welcome to the contoso.com domain** is displayed, click **OK** twice, and then click **Close**.
6. When you are prompted to restart the computer, click **Restart Now**.
7. After restarting the computer, press **Ctrl+Alt+Delete**, click the left arrow, click **Other user**, and sign in using credentials for the CONTOSO\user1 account.

Configure IPAM1

IPAM1 is a computer running Windows Server 2012, providing the following services:

- An IPAM server.

Initial configuration of IPAM1 consists of the following steps:

- Install the operating system and configure TCP/IP on IPAM1
- Install and configure IPAM on IPAM1

Install the operating system and configure TCP/IP on IPAM1

Tip

The procedure below is identical to the steps used to install the operating system and configure TCP/IP on DC1 and DHCP1, with the exception that IPAM1 is configured with an IP address of 10.0.0.3.

To install the operating system and configure TCP/IP on IPAM1

1. Start your computer using the Windows Server 2012 product disc or other digital media.
2. When prompted, enter a product key, accept license terms, configure clock, language, and regional settings, and provide a password for the local Administrator account.
3. Press **Ctrl+Alt+Delete** and sign-in using the local Administrator account.
4. If you are prompted to enable Windows Error Reporting, click **Accept**.

5. In the Server Manager navigation pane, click **Local Server** and then click the IP address next to **Wired Ethernet Connection**. The **Network Connections** control panel will open.
6. In **Network Connections**, right-click **Wired Ethernet Connection** and then click **Properties**.
7. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
8. On the **General** tab, choose **Use the following IP address**.
9. Next to **IP address** type **10.0.03** and next to **Subnet mask** type **255.255.255.0**. It is not necessary to provide an entry next to **Default gateway**.
10. Next to **Preferred DNS server**, type **10.0.0.1**.
11. Click **OK** twice, and then close the **Network Connections** control panel.

Install and configure IPAM on IPAM1

IPAM1 is a domain member server running the IPAM feature. The following procedure provides steps to install the IPAM feature using Server Manager. First, the computer will be renamed and joined to the contoso.com domain.

To install IPAM on IPAM1

1. In the Server Manager navigation pane, click **Local Server** and then click the name next to **Computer name**. The **System Properties** control panel will open.
2. On the **Computer Name** tab, click **Change** and then type **IPAM1** under **Computer name**.
3. Under **Member of**, select **Domain**, type **contoso.com**, and then click **OK**.
4. When you are prompted to provide credentials to join the domain, enter the credentials for the user1 account that was created previously and then click **OK**.
5. Confirm that computer name and domain changes were successful, click **OK** twice, and then click **Close**.
6. When you are prompted to restart the computer, click **Restart Now**.
7. After restarting the computer, sign-in using the CONTOSO\user1 account.
8. In Server Manager, under **Configure this local server**, click **Add Roles and Features**.
9. In the **Add Roles and Features Wizard**, click **Next** four times, and then on the **Select features** page select the **IP Address Management (IPAM) Server** checkbox.
10. When you are prompted to add required features, click **Add Features**.
11. Click **Next**, and then click **Install**.
12. Wait for the installation process to complete, verify on the **Installation progress** page that **Installation succeeded on IPAM1.contoso.com** is displayed, and then click **Close**.

To configure IPAM

1. In the Server Manager navigation pane, click **IPAM**. The IPAM Overview page is displayed. By default, the IPAM client is connected to the local server.
2. Click **provision the IPAM server**. The **Provision IPAM** wizard will launch.

3. Click **Next**. By default, the **Group Policy Based** provisioning method is chosen.
4. Next to **GPO name prefix**, type **IPAM1** and then click **Next**.
5. On the **Summary** page, confirm that the GPO names displayed are **IPAM1_DHCP**, **IPAM1_DNS**, and **IPAM1_DC_NPS**, and then click **Apply**.
6. Wait for provisioning to complete, and then on the **Completion** page verify that **IPAM provisioning completed successfully is displayed**.

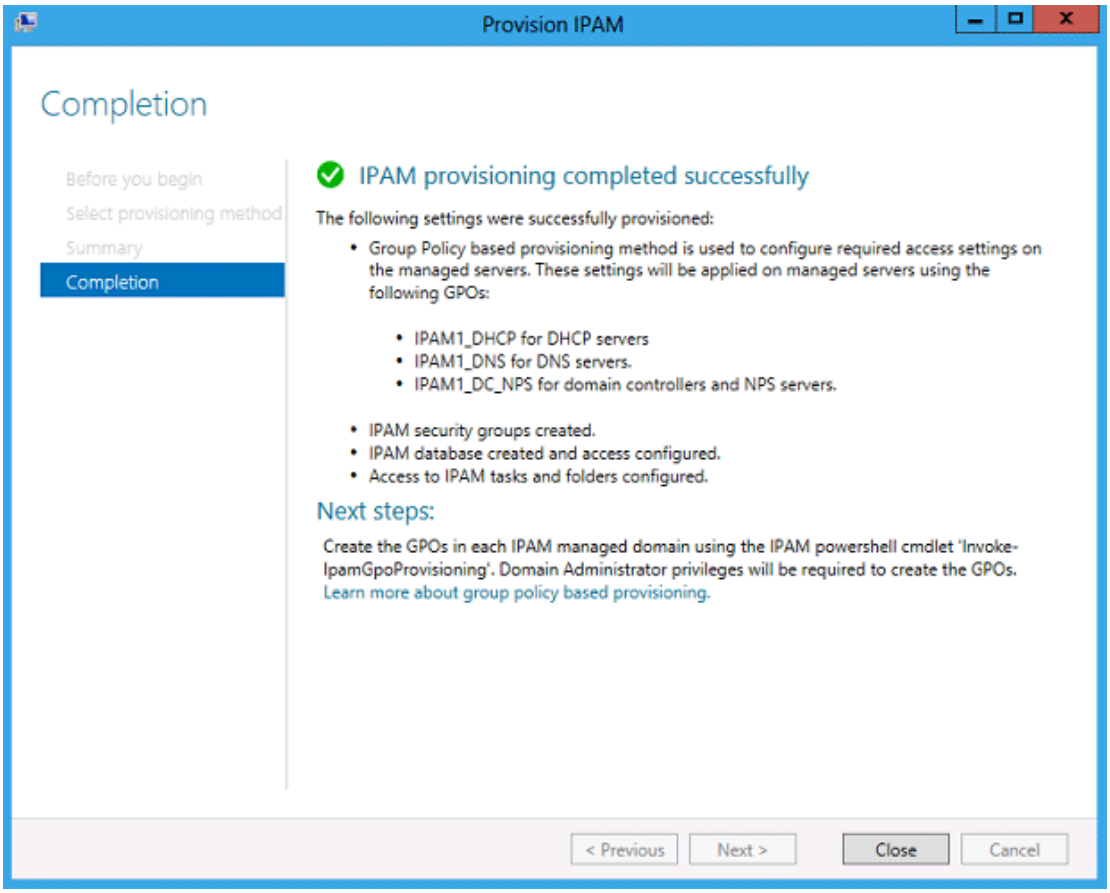

7. Click **Close**, and then in the **Overview** pane, click **configure server discovery**.
8. In the **Configure Discovery Settings** dialog box, under **Select IPAM domains for discovery**, click **Add** next to **(root domain) contoso.com**.
9. Verify that the contoso.com domain was added and the server roles selected include Domain controller, DHCP server, and DNS server.
10. When you are prompted that the discovery scope has been updated, click OK.

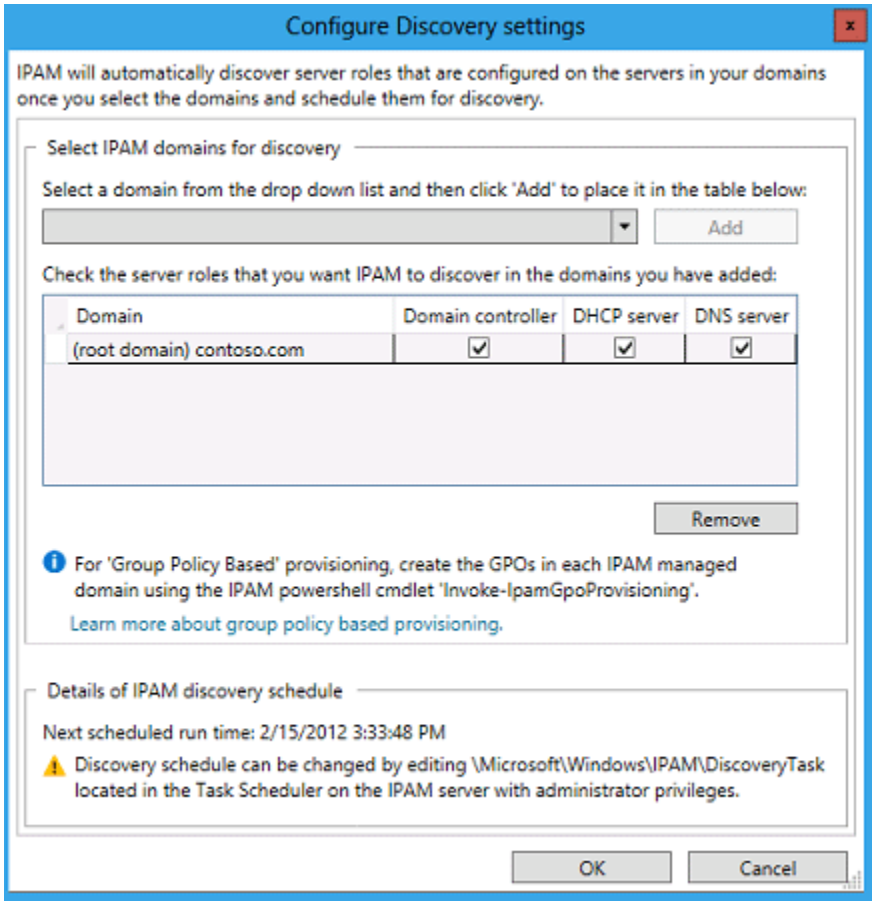

11. Click **start server discovery**, and then click **OK** when you are prompted that discovery has started.
12. Click the Notification flag and then click **Task Details**.
13. Wait for the **IPAM ServerDiscovery task** to display a status of **Complete** under **Stage**, and then close the **Task Details** dialog box.
14. In IPAM OVERVIEW, click **select or add servers to manage and verify IPAM access**. If no servers are displayed, click the **Refresh IPv4** icon located next to the Notification flag. The DHCP1 and DC1 servers will be displayed with a manageability status of **Unspecified** and an IPAM access status of **Blocked**.

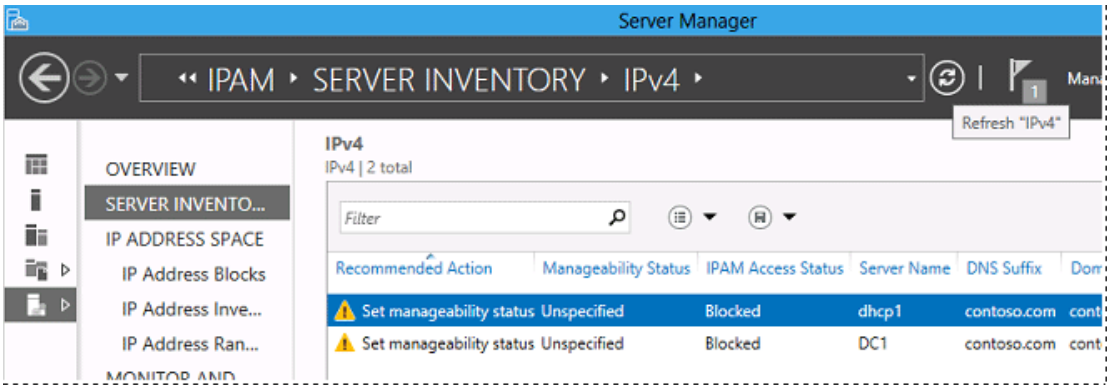

Next, IPAM1 must be granted permission to manage DHCP1 and DC1 using Group Policy Objects (GPOs).

15. On IPAM1, right-click **Windows PowerShell** and then click **Run as Administrator**. Click **Yes** in the User Account Control alert that is displayed.
16. Type the following command at the Windows PowerShell prompt, and then press ENTER.

```
Invoke-IpamGpoProvisioning -Domain contoso.com -GpoPrefixName
IPAM1 -DelegatedGpoUser user1 -IpamServerFqdn
ipam1.contoso.com
```

17. When you are prompted to confirm the action, press ENTER.
18. On the Server Manager menu, click **Tools** and then click **Group Policy Management**.
19. In the Group Policy Management console tree, navigate to **contoso.com\Group Policy Objects** and verify that three GPOs have been created named **IPAM1_DC_NPS**, **IPAM1_DHCP**, and **IPAM1_DNS**.

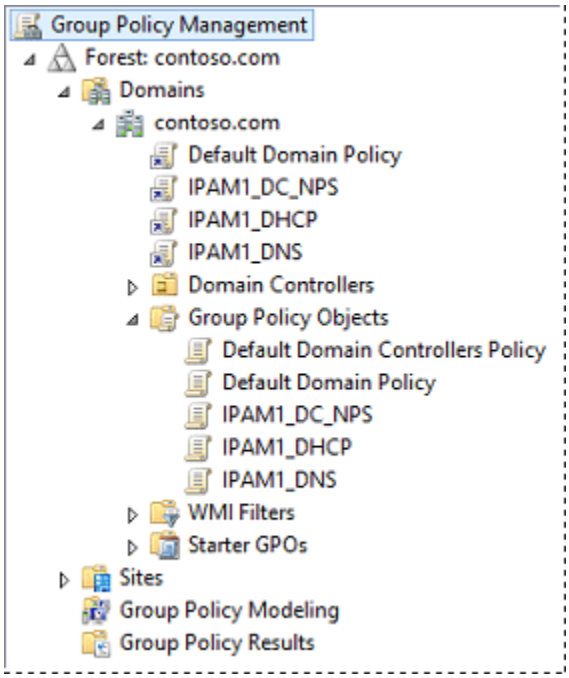

20. Close the Group Policy Management console.
21. In the IPAM > SERVER INVENTORY pane, right-click **dhcp1** and then click **Edit Server**.
22. In the **Add or Edit Server** dialog box, next to **Manageability status**, choose **Managed**, and then click **OK**.
23. Right-click **DC1**, click **Edit Server**, choose **Managed**, and then click **OK**.
24. On DHCP1, click **Windows PowerShell**, type **gpupdate /force**, and then press ENTER.
25. On DC1, click **Windows PowerShell**, type **gpupdate /force**, and then press ENTER.
26. On IPAM1, click the Refresh IPv4 icon and verify that **UnBlocked** is displayed under **IPAM Access Status** for DHCP1 and DC1.

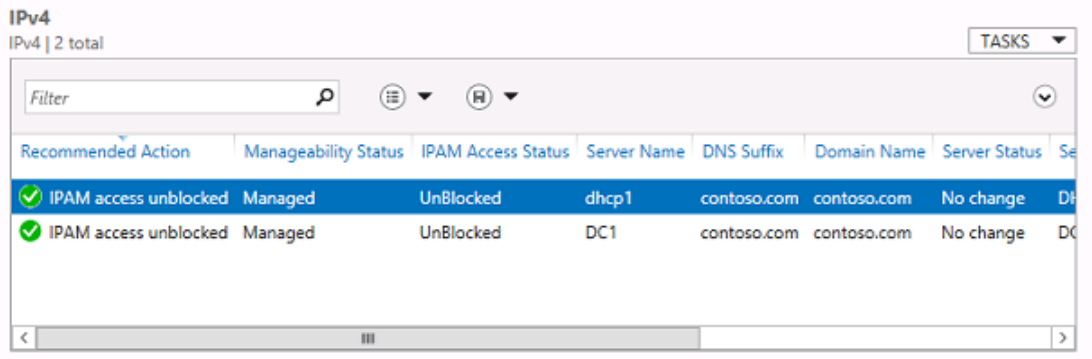

Tip

You might need to wait a few minutes and refresh the IPAM console view for IPAM access status to be updated on managed servers after changing

manageability status.

27. In IPAM > OVERVIEW, click **retrieve data from managed servers**.
28. Click the Notification flag and wait for all tasks to complete.

IPAM demonstration

A demonstration of IPAM on Windows Server 2012 includes the following procedures:

29. 0.1. Address space management
    - Create, delete, import and export IP addresses
    - Find available IP addresses and create reservations
    - Create custom logical groups
2. Infrastructure monitoring and management
3. Review audit logs and events

Address space management

In IPAM, IP address blocks are large chunks of IP addresses that are used for organization of address space. IP address ranges are smaller chunks of IP addresses that typically correspond to a DHCP scope. IP address ranges are mapped to IP address blocks.

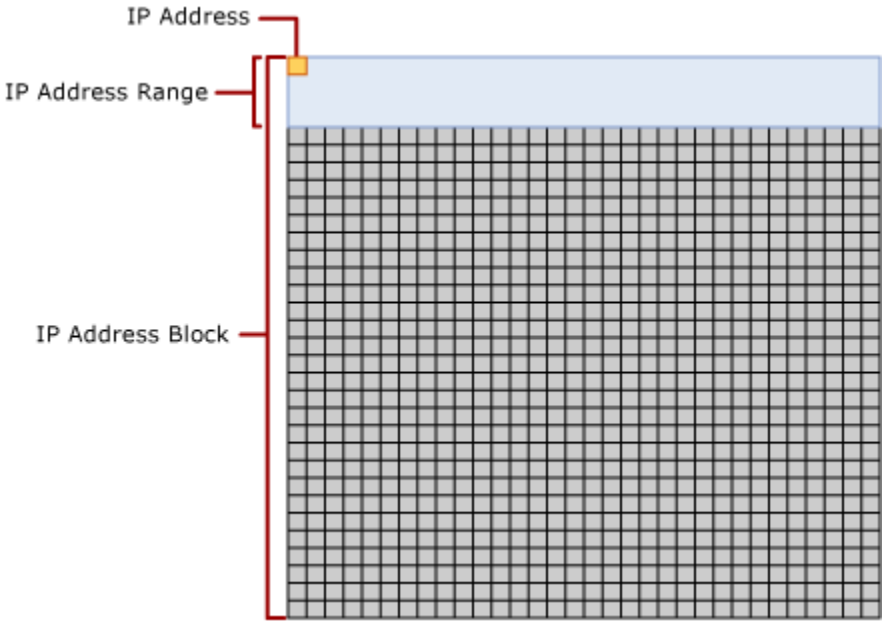

IP addresses can be entered into IPAM manually, or by importing from a comma-delimited file. Addresses can also be exported to a file in comma-delimited format.

Create, delete, import and export IP addresses

The following procedure demonstrates how IP address blocks, ranges, and addresses can be created, deleted, exported, and imported in IPAM.

To create, delete, import, and export IP addresses

1. In the upper IPAM navigation pane, click **IP Address Blocks**.
2. In the lower navigation pane, right-click **IPv4** and then click **Add IP Address Block**.

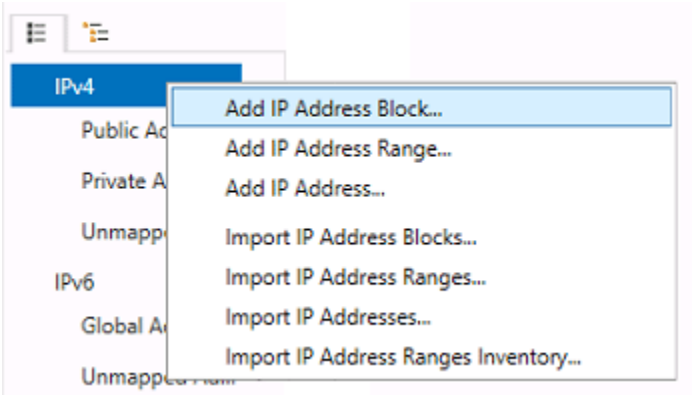

> **Note**
> The IP address block you create is automatically added to public or private address space according to the start and end IP addresses you specify.

3. In the **Add or Edit IPv4 Address Block** dialog box, next to **Network ID**, type **10.0.0.0**.
4. Next to **Prefix Length**, choose **8**. This is the /8 corresponding to the /24 subnet that is being dynamically allocated by DHCP1.
5. Click **OK**, and then next to **Current View** choose **IP Address Blocks**.

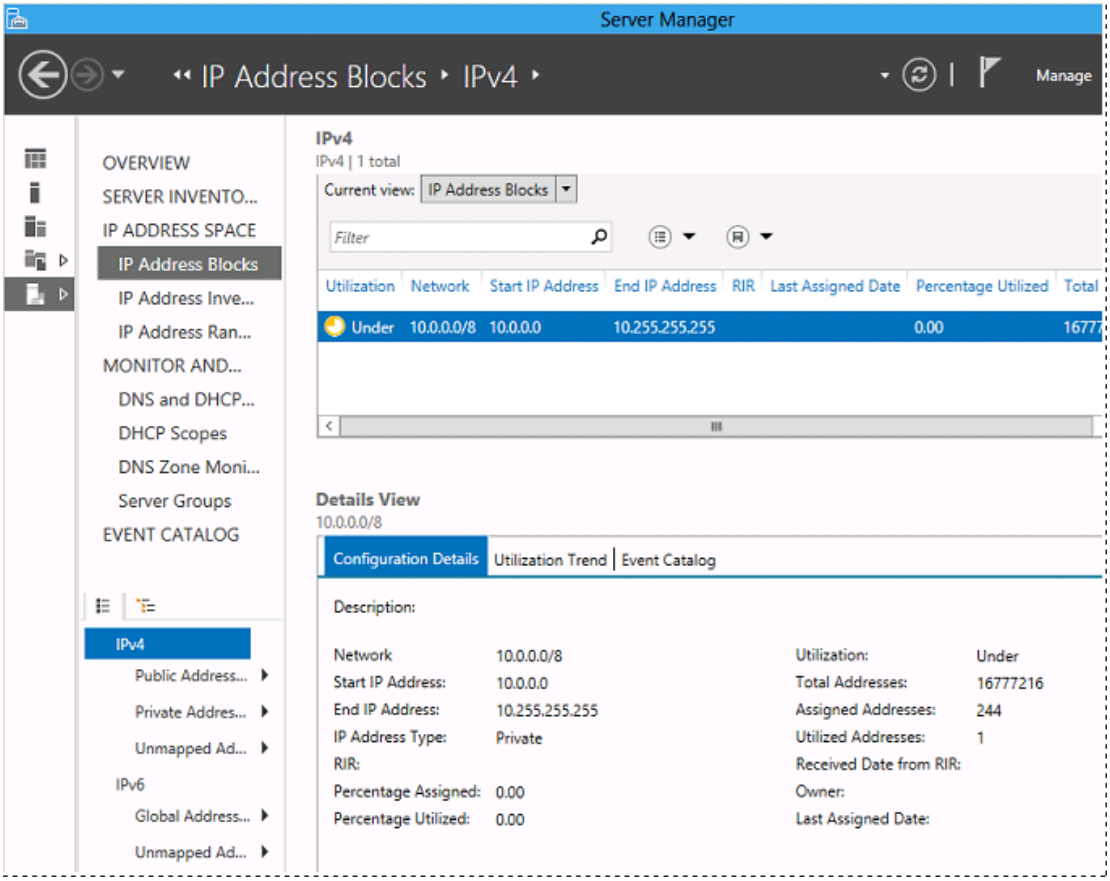

6. On the **Configuration Details** tab, next to **Utilized Addresses**, note that one IP address is currently in use. This corresponds to the lease issued by DHCP1 for Client1.
7. Next to **Current view**, choose **IP Address Ranges**.
8. On the **Configuration Details** tab, review the information displayed. Details are provided for Contoso-scope1 supplied by dhcp1.contoso.com.

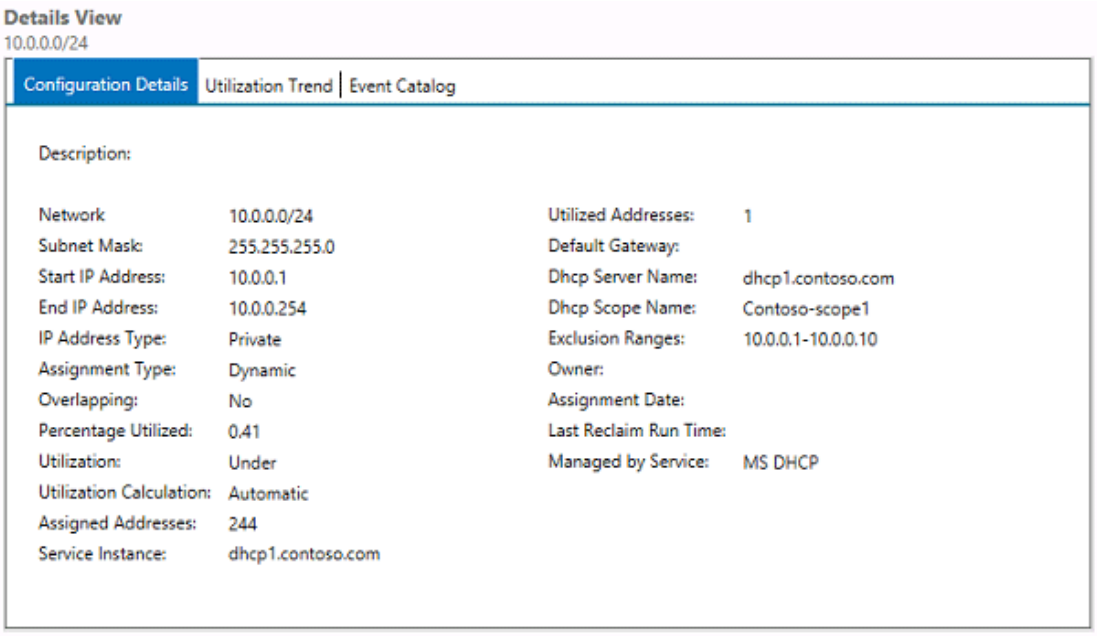

9. In the lower navigation pane, right-click **IPv6** and then click **Add IP Address Block**.
10. Under **Specify the Network ID**, type **21da:d3:0:2f3b::** and then move the slider next to **Specify Prefix length** to that the prefix is **64**, and then click **OK**.

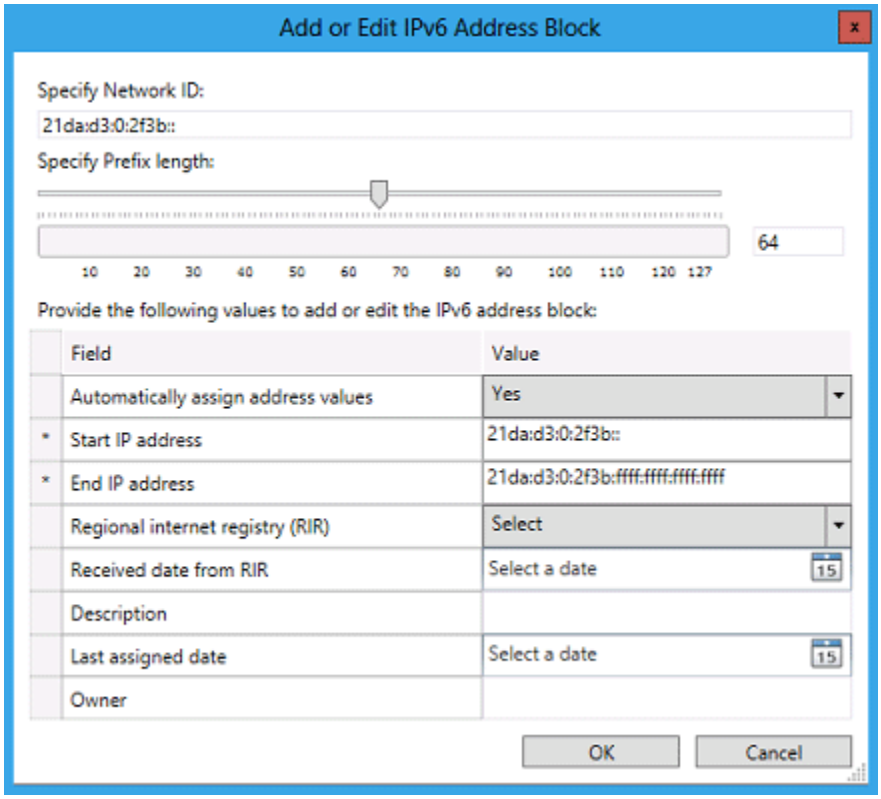

11. Choose **IP Address Blocks** next to **Current view** and confirm that the **21da:d3:0:2f3b::/64** block was successfully added.
12. Right-click **IPv4** and add the following IP address blocks:
    - **192.168.0.0/24**
    - **192.168.1.0/24**
13. Right-click **IPv4** and add the **207.46.0.0/16** address block. Since this is public address space, you must choose a regional Internet registry. Choose **ARIN**, and if desired, supply dates and a description for this block of public IP address space.

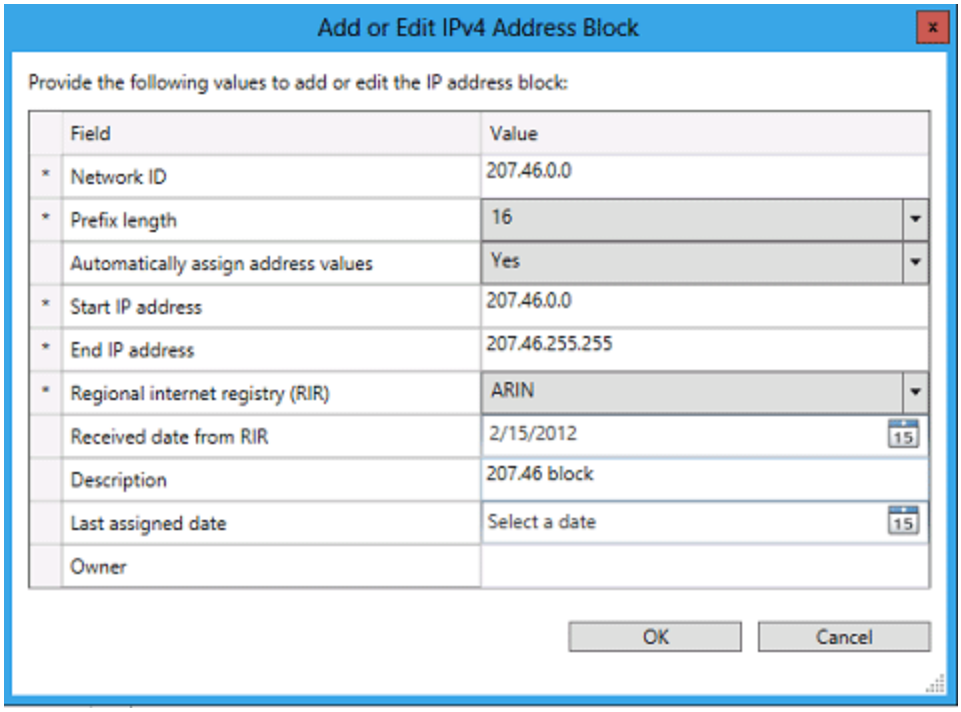

14. Ensure that the **Current view** selected is **IP Address Blocks** and click the **Network** field to sort by highest to lowest network ID. Also try sorting by some other fields.
15. In the lower navigation pane, under **IPv4**, click **Public Address Space** and verify that the **207.46.0.0/16** IP address block is displayed.
16. Right-click **IPv4** and then click **Add IP Address Range**.
17. Next to **Network ID**, type **192.168.0.0**, choose **25** next to **Prefix length**, and then click **OK**.

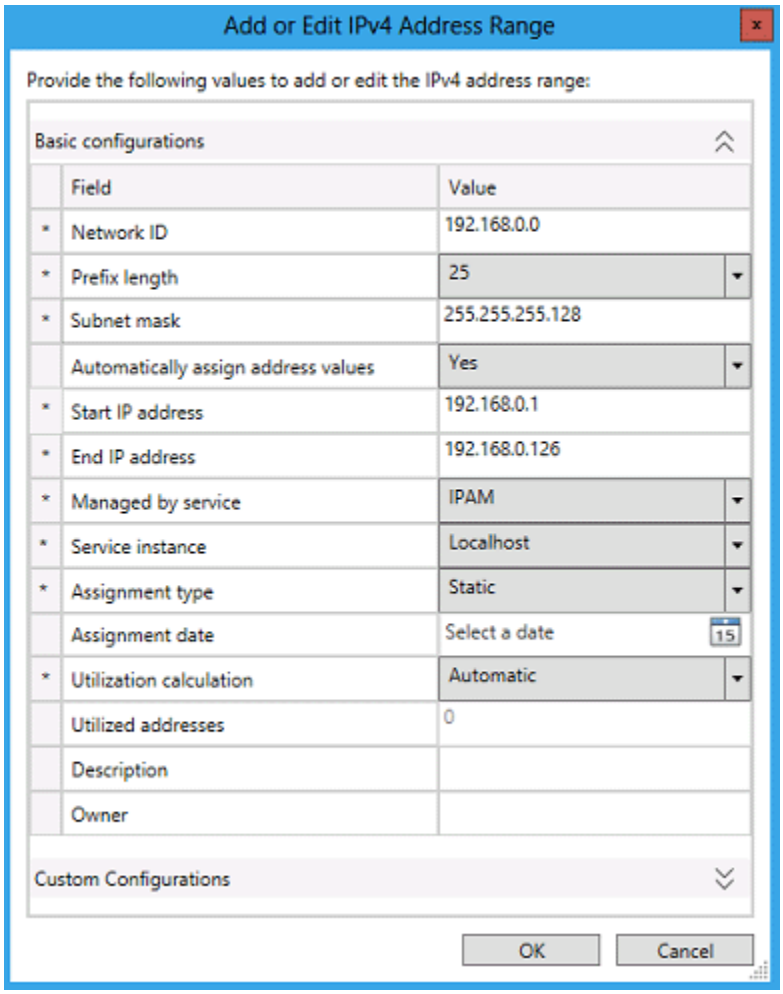

18. Right-click IPv4 and add the following IP address ranges:
    - **192.168.0.128/25**
    - **192.168.1.0/25**
    - **192.168.1.128/25**
19. Right-click **IPv4**, and then click **Add IP Address**.
20. In the **Add IP Address** dialog box, next to **IP address**, type **192.168.0.1**.
21. Next to **MAC address**, type **112233445566** and then click **OK**.
22. Next to **Current view**, choose **IP Addresses** and verify that the static IP address **192.168.0.1** was added, and that it is automatically assigned to the **192.168.0.1-192.168.0.126** range.
23. With the current view set to IP Addresses, click TASKS and then click Export.

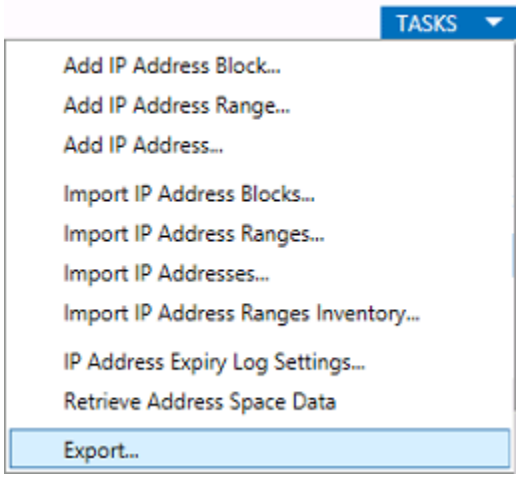

24. Choose a location where you want to save the file.
25. In the **Save As** dialog box, type **ip-addresses** next to **File name** and then click **Save**.
26. Right-click the **ip-addresses.csv** file and then click **Edit**.
27. Highlight the line containing the 192.168.0.1 IP address, right-click the line, and then click **Copy**.

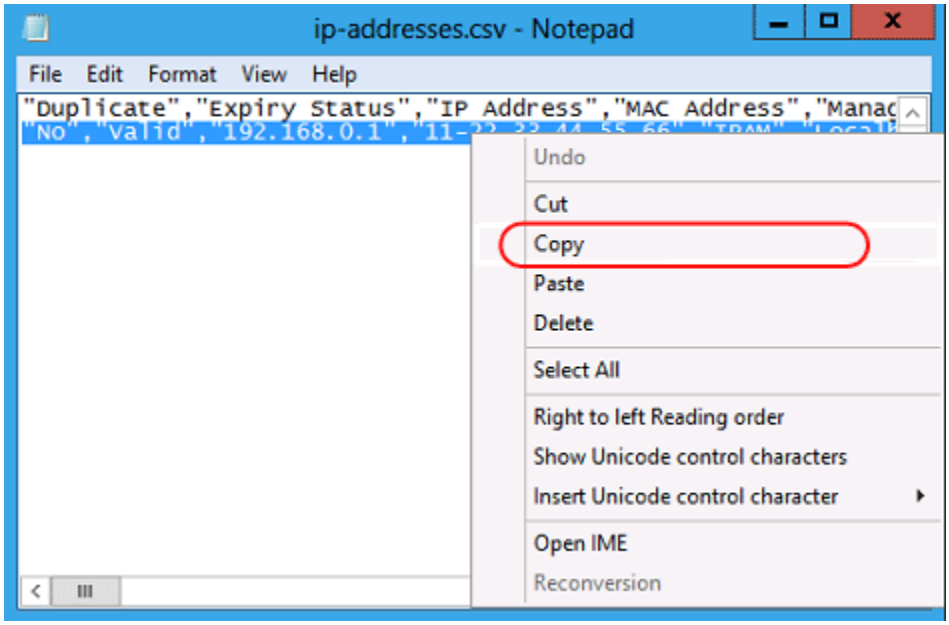

28. Paste the contents of the copied line underneath the text four times, so that you create a total of six rows of text, with the first row containing the column headers.
29. Change the IP address in all five lines from 192.168.0.1 to values ranging from 192.168.0.2 – 192.168.0.6 and then save the file.

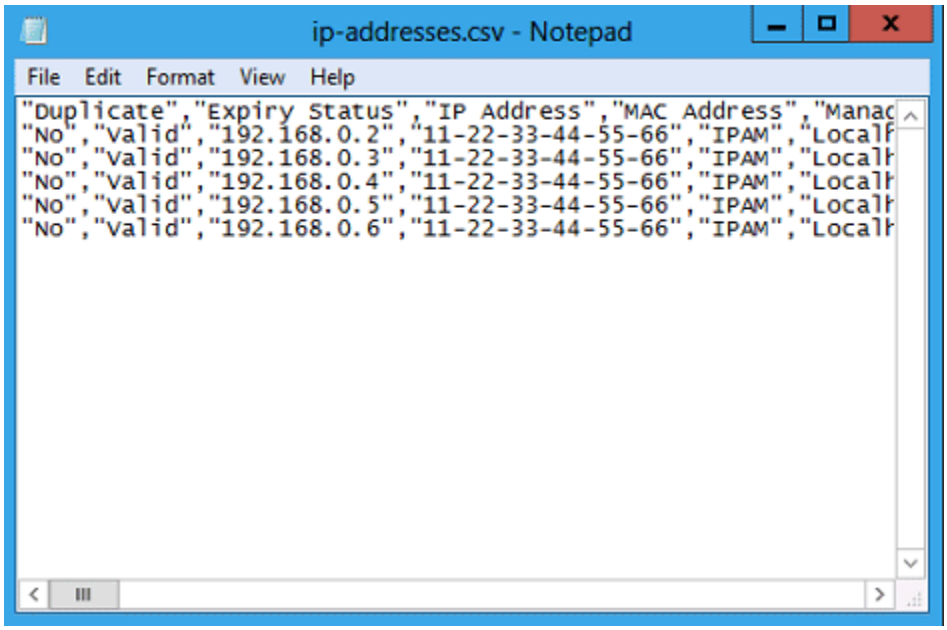

30. Right-click **IPv4** and then click **Import IP Addresses**.
31. Select the **ip-addresses.csv** file and then click **Open**.
32. In the **Import IP Addresses** dialog box, verify that **5 out of 5 records successfully imported** is displayed, and then click OK.
33. Verify that five new IP addresses were added to the **192.168.0.1-192.168.0.126** range.
34. Right-click the **192.168.0.6** IP address and then click **Delete**.
35. Verify that the 192.168.0.6 IP address was removed from the list.

Find available IP addresses and create reservations

A network administrator might wish to locate an available IP address and use it for static assignment to a network device. The following steps demonstrate how to use the **Find and Allocate Available IP Address** function in IPAM for this scenario.

To find, reserve, and reclaim IP addresses

1. With the **Current view** set to **IP Address Ranges**, right-click the 10.0.0.1/24 range that is assigned by DHCP1 and then click **Find and Allocate Available IP Address**.
2. Because the first ten IP addresses are reserved in the Contoso-scope1 DHCP scope, and Client1 has been allocated the first available IP address, the first available IP address will be 10.0.0.12.
3. Wait a few seconds for **Ping Reply Status** and **DNS Record Status** to resolve and display **No Reply** and **Not Found**, respectively.

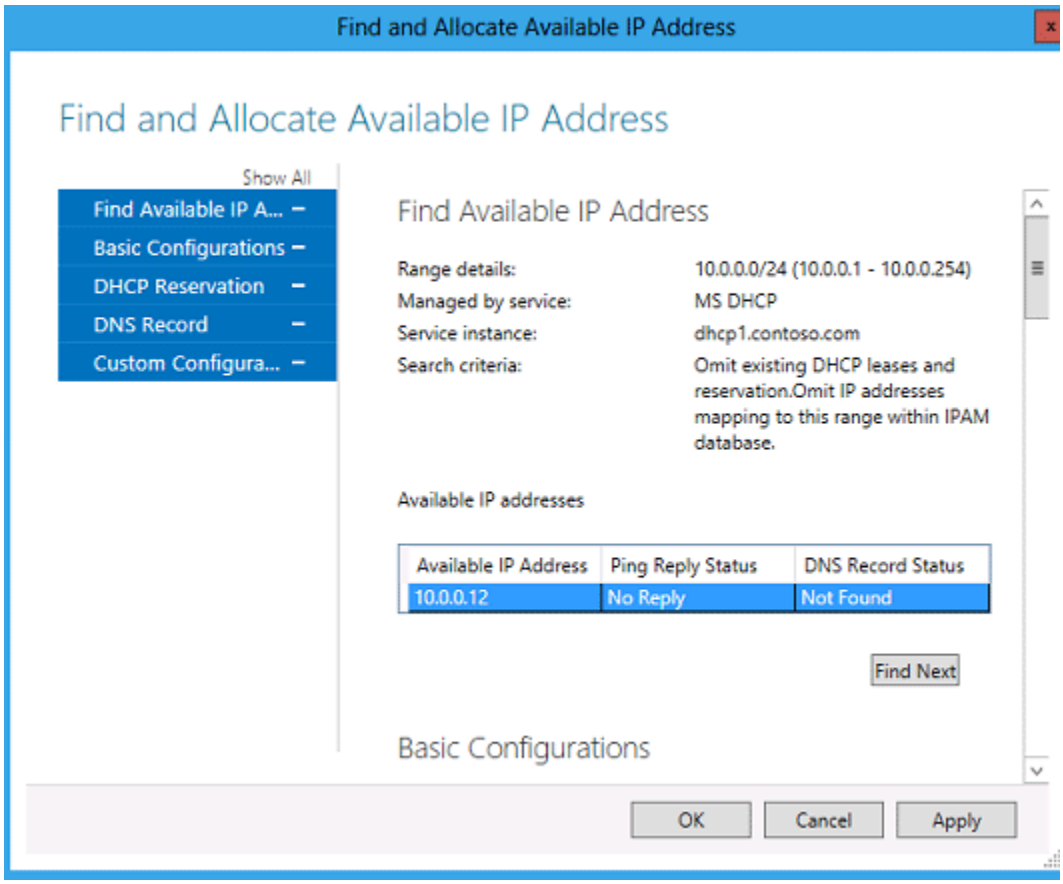

4. Click **Basic Configurations**.
5. Next to **MAC address**, type **112233445566** and next to **Device type** choose **VOIP Gateway**.
6. Click **DHCP Reservation**.
7. Next to **Reservation server name**, choose **dhcp1.contosoc.com**. The **Reservation scope name** will automatically display **Contoso-scope1**.
8. Next to **Reservation name**, type **voip-gw**, and then next to **Reservation type**, choose **Both**.
9. Click **DNS Record**.
10. Next to **Device name**, type **voip-gw**, next to **Forward lookup zone**, choose **contoso.com**, and then next to **Forward lookup primary server** choose **DC1.contoso.com**.
11. If a reverse lookup zone has not been created yet, no in-addr.arpa zone is available to select.
12. Click **OK**, and then in the IPAM navigation pane under **IP ADDRESS SPACE**, click **IP Address Inventory**.
13. In the lower navigation pane, click the arrow next to **IPv4** to expand IPv4 and then click **VOIP Gateway**.

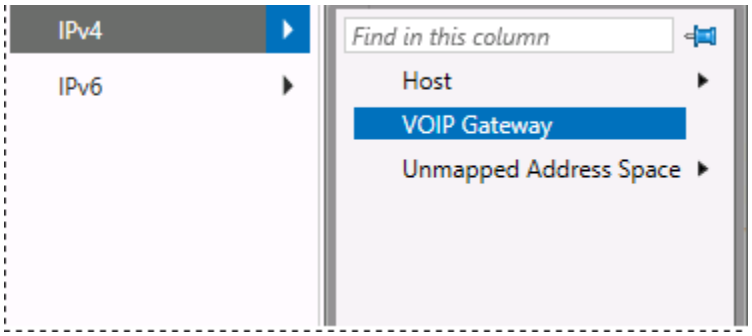

14. Verify that the **10.0.0.12** IP address is displayed.

> **Tip**
> So far, changes have only been made to the IPAM database. The following steps will be used to create a DHCP reservation and DNS host record.

15. Right-click the 10.0.0.12 IP address and then click **Create DNS Host Record**.

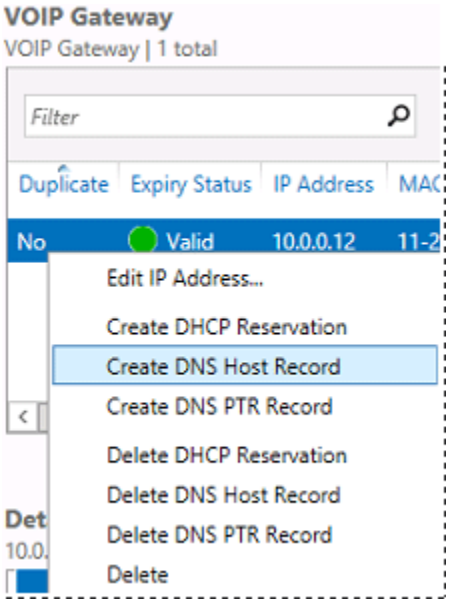

16. Right-click the 10.0.0.12 IP address and then click **Create DHCP Reservation**.
17. On the **Configuration Details** tab, verify that **Create Success** is displayed next to **DHCP reservation sync** and **DNS Host Record sync**.
18. On DHCP1, in the DHCP console, verify that the reservation is present in the Contoso-scope1 DHCP scope.

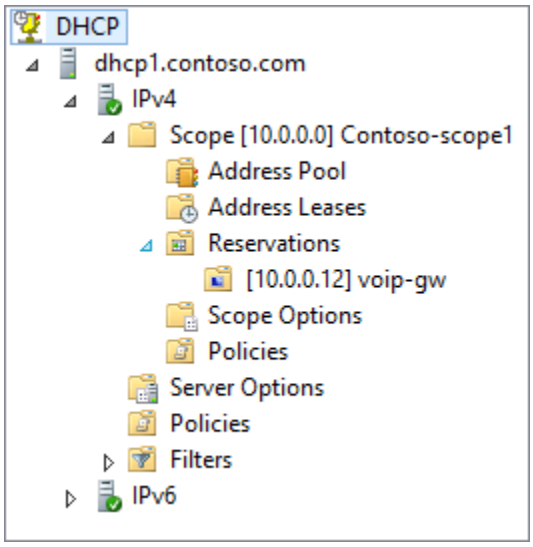

19. On DC1, in DNS Manager, verify that the host record is present.

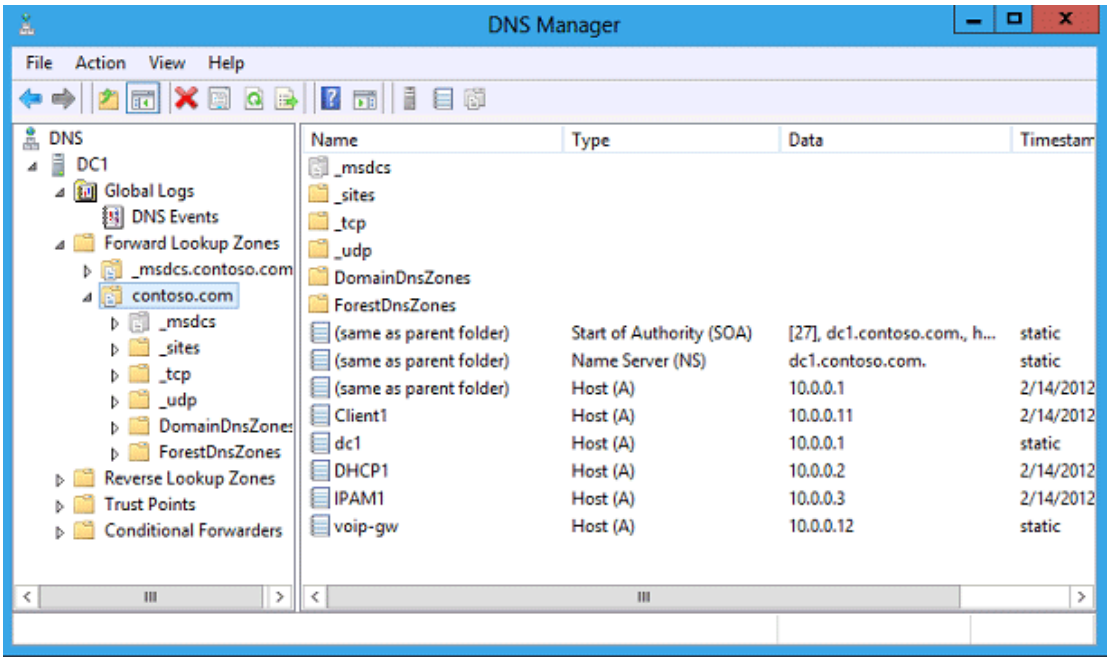

20. On IPAM1, right-click the 10.0.0.12 IP address and then click **Edit IP Address**.
21. Under **Basic Configurations**, click **Select a date** next to **Assignment date** and enter today's date.
22. Click **Select a date** next to **Expiry date**, select a date one month from today, and then click **OK**.

Important

Expiry settings are alerts you can create for objects in the IPAM database. When a reserved IP address passes the expiry date, it is not removed from reservations

on the DHCP server, but IPAM will provide events and alerts when the expiry date is close.

23. Verify that **Valid** is displayed under **Expiry Status**.
24. Click **TASKS** and then click **IP Address Expiry Log Settings**.
25. Under **Expiry Alert Threshold**, type **31**.
26. Under **Logging Frequency**, choose **Log all expiry status messages periodically** and then click OK.

    Tip
    By default, expiry logging begins 10 days before the expiration date. When you choose to log alerts periodically, they will be logged each time the expiry task runs. The expiry task runs once each day by default, but can be configured to run more or less frequently.

27. Refresh the IPAM console view and verify that **Expiry Due** is displayed under **Expiry Status**.
28. Edit the IP address again and change the assignment date and expiry date to one week in the past. Verify that the address is now displayed as **Expired**.
29. Right-click the 10.0.0.12 address and then click **Delete DHCP Reservation**. This removes the DHCP reservation from the DHCP server.
30. Right-click the 10.0.0.12 address and then click **Delete DNS Host Record**. This removes the forward lookup record from the authoritative DNS server.
31. Click **IP Address Blocks** in the IPAM navigation pane and change the current view to **IP Address Ranges**.
32. Highlight all the available ranges by holding down the SHIFT key and clicking the top and bottom ranges.
33. Right-click the highlighted IP address ranges, and then click **Reclaim IP Addresses**.
34. Under **Select IP addresses to be reclaimed**, select the checkbox next to the 10.0.0.12 address, click **Reclaim** and then click **Close**. This removes the IP address from the IPAM database.

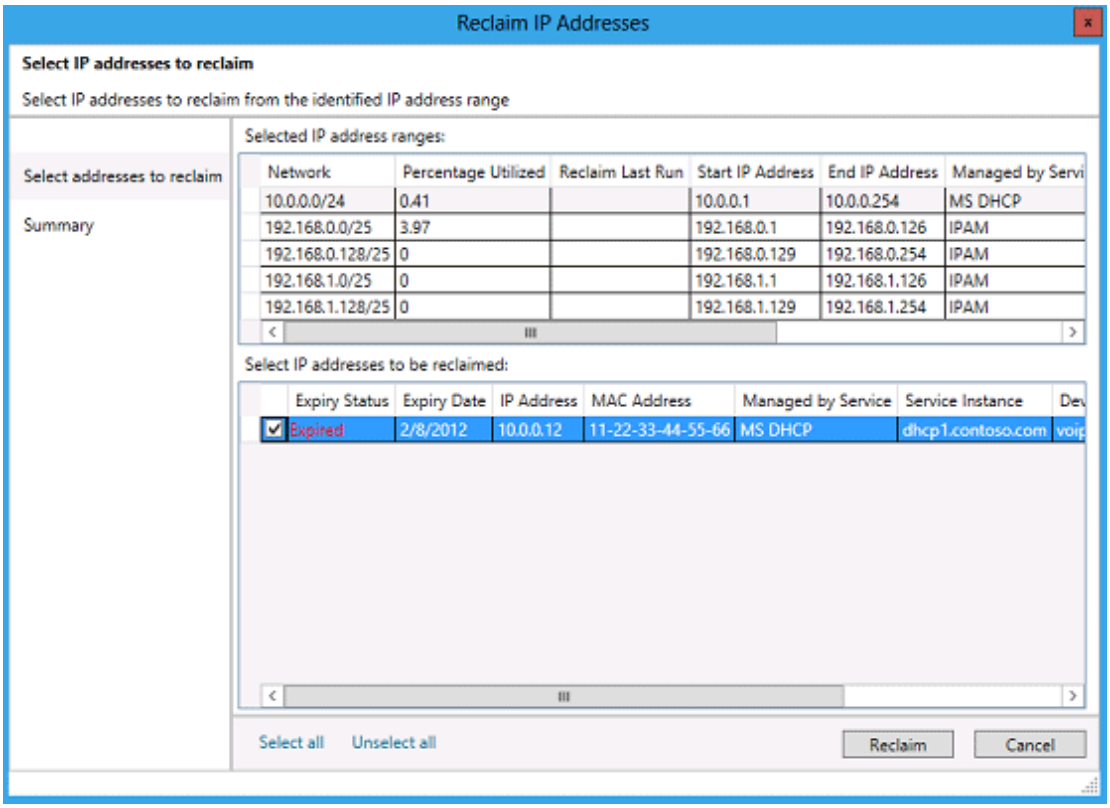

Tip

Reclaiming IP addresses allows you visualize expiry status and delete multiple IP addresses. You can also right-click one or more IP addresses and click **Delete** to remove IP addresses from the IPAM database.

Create custom logical groups

The **IP Address Inventory** group is a built-in group with IP addresses organized by device type. In addition, IPAM allows you to create custom logical groups. To create custom groups:

To create custom logical groups

1. In the IPAM navigation pane, under **IP ADDRESS SPACE**, click **IP Address Range Groups**.
2. On the Server Manager menu, click **Manage** and then click **IPAM settings**.
3. In the **IPAM settings** dialog box, click **Configure custom fields**.
4. In the **Configure Custom Fields** dialog box, under **Add custom fields below**, scroll to the bottom of the list, type **Building** for the **Custom Field Name**, and then select **Yes** under **Multi-Value**.
5. Press ENTER or TAB to commit the new custom field name. A blank line will open that can be used for additional custom fields.

6. Click **Building** and then under **Custom Field Value** type the following values. Press ENTER after you type each one:
   a. **Headquarters**
   b. **Operations**
   c. **Sales**
   d. **Data Center**

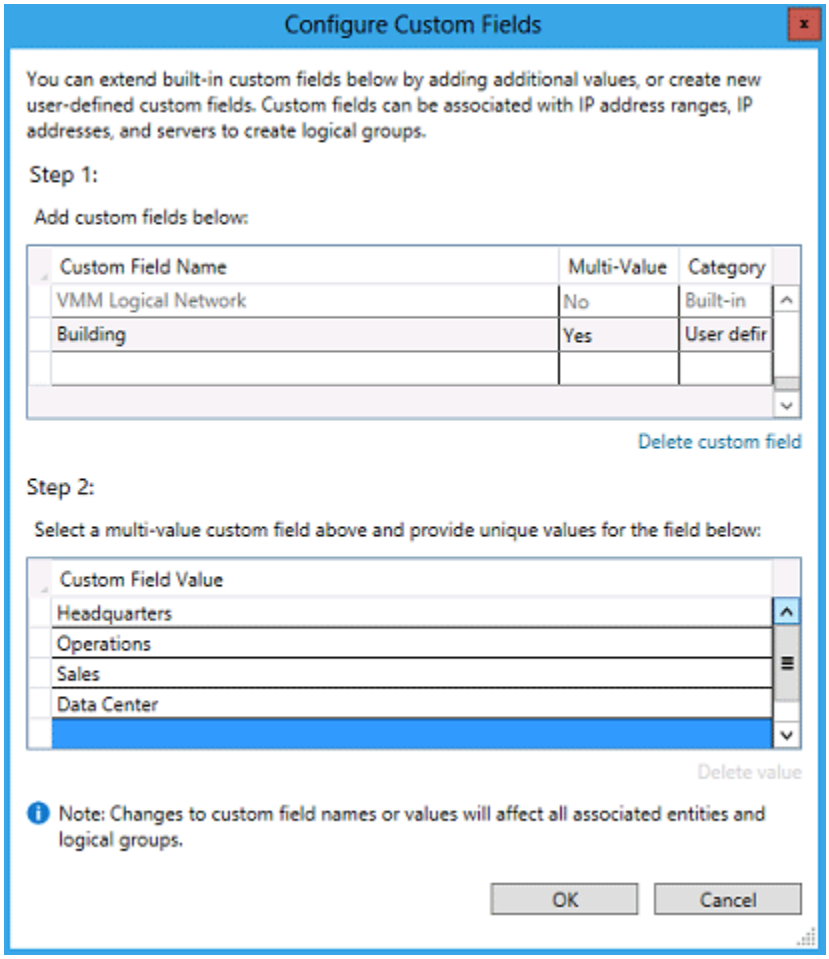

7. Repeat the previous step to add another custom field named **Floor** with the following two custom field values:
   - **First**
   - **Second**
8. Click **OK** twice, and then click **Close**.
9. Click **IP Address Ranges**, right-click the **192.168.0.0/25** range, and then click **Edit IP Address Range**.
10. Click **Custom Configurations**, and then next to **Custom field to configure**, choose **Building**.
11. Next to Specify a value, choose **Headquarters** and then click **Add**.

12. Choose **Floor** next to **Custom field to configure**, choose **First**, and then click **Add**.
13. Edit the other three IP address ranges and add a unique building and a floor to each.

    Tip
    You can also select multiple IP address ranges and add custom fields to all the ranges in one step.
14. Refresh the IP Address Ranges view, right-click the column header and then select **Building** and **Floor** two of the fields to display. The building and floor is now displayed with each IP address range in the list.
15. Right-click **IPv4** and then click **Add IP Address Range Group**.
16. Under **Provide name of the address range group**, type **Building/Floor**.
17. Under **Custom Fields**, select **Building** and then select **Floor** so that items are grouped first by Building and then by Floor.
18. Click OK, and then click the arrow next to IPv4.
19. Verify that you can view IP address ranges by building and floor.

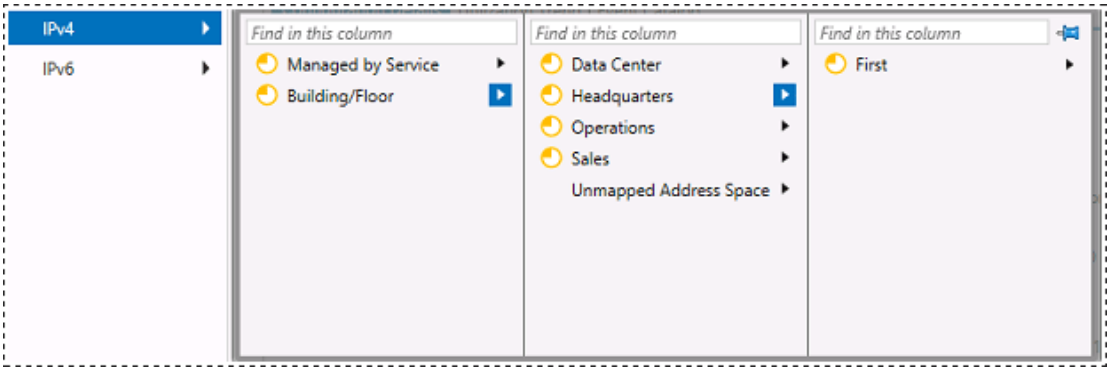

Infrastructure monitoring and management

The following procedure demonstrates how DHCP and DNS servers can be monitored from the IPAM server.

To monitor and manage DHCP and DNS servers

1. In the IPAM navigation pane, under **MONITOR AND MANAGE**, click **DNS and DHCP Servers**.
2. Next to **Server Type**, note that you can choose **DNS**, **DHCP**, or **DNS and DHCP**. The server availability, duration in current state, server name, server role, domain name, and IP address is displayed.
3. Click dhcp1.contoso.com, and then under **Details View** review the information provided on the **Server Properties**, **Options**, and **Event Catalog** tabs.
4. Right-click the DHCP server and note that you can directly configure the DHCP server from the IPAM console.

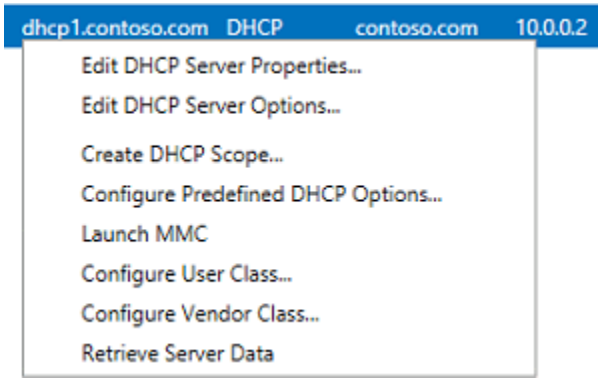

5. Next to **Server Type** choose **DHCP** and then next to **View** choose **Scope properties**.
6. Right-click the **Contoso-scope1** DHCP scope and then click **Duplicate DHCP Scope**.
7. In the **Duplicate DHCP Scope** dialog box, change the **Scope name** to **Contoso-scope2**.
8. Type the following values under **General Properties**:
   - **Start IP address**: **10.0.1.1**
   - **End IP address**: **10.0.1.254**
   - **Subnet mask**: **255.255.255.0**

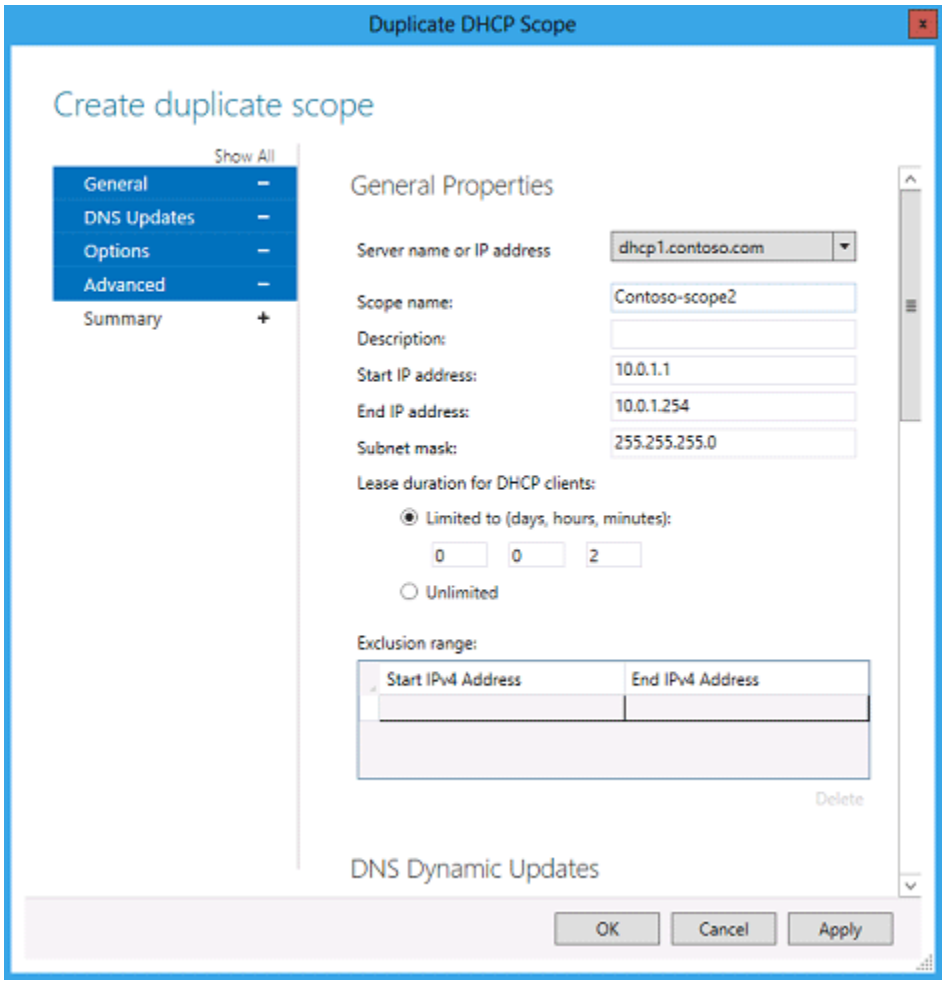

9. In the left pane, click **DNS Updates**, click **Options**, and click **Advanced**. Note that all the scope properties have already been configured identically to the Contoso-scope1 DHCP scope. You can also edit these values if desired.
10. Click **OK** and verify that a new DHCP scope is displayed in the list with the **Scope Name Contoso-scope2**.
11. Refresh the DHCP console on DHCP1 and verify that the Contoso-scope2 DHCP scope is configured and activated.
12. On IPAM1, select both DHCP scopes using SHIFT, right-click the scopes and then click **Edit DHCP Scope**.
13. In the left pane, click Options.
14. Next to **Configuration action**, choose **Add**, and next to **Option** choose **003 Router**.
15. Under IP Address, click **0.0.0.0** and click **Delete**.
16. Under **IP Address**, type **10.0.0.10**, press ENTER, and then click **Add to list**.
17. Click **OK** and verify that a new 003 Router option has been added to both DHCP scopes. The next time that Client1 renews a DHCP lease it will receive this configuration option.

Tip

You can use this method to bulk-edit options on multiple DHCP scopes at once. In the previous example, the **Add** function was used. You can also choose **Overwrite**, **Find and replace**, or **Delete**.

18. Next to **Server Type**, choose **DNS**.
19. Under **Details View**, review the information provided on the **Server Properties**, **DNS Zones**, and **Event Catalog** tabs.
20. Right-click **DC1.contosol.com** and then click **Launch MMC**. Note that you can directly configure zones on DC1.
21. In the IPAM navigation pane, click **DHCP Scopes** and review the information under **Details View** on the **Scope properties** and **Options** tabs.
22. In the IPAM navigation pane, click **DNS Zone Monitoring** and review the information on the **Zone Properties** and **Authoritative Servers** tabs.
23. In the IPAM navigation pane, click **Server Groups**.
24. Right-click **IPv4** and note that the same logical group functionality that is available for IP address ranges is also available for managed servers. Custom fields are available by editing server properties and selecting the Custom Configuration menu item. This provides a highly customizable managed server display.

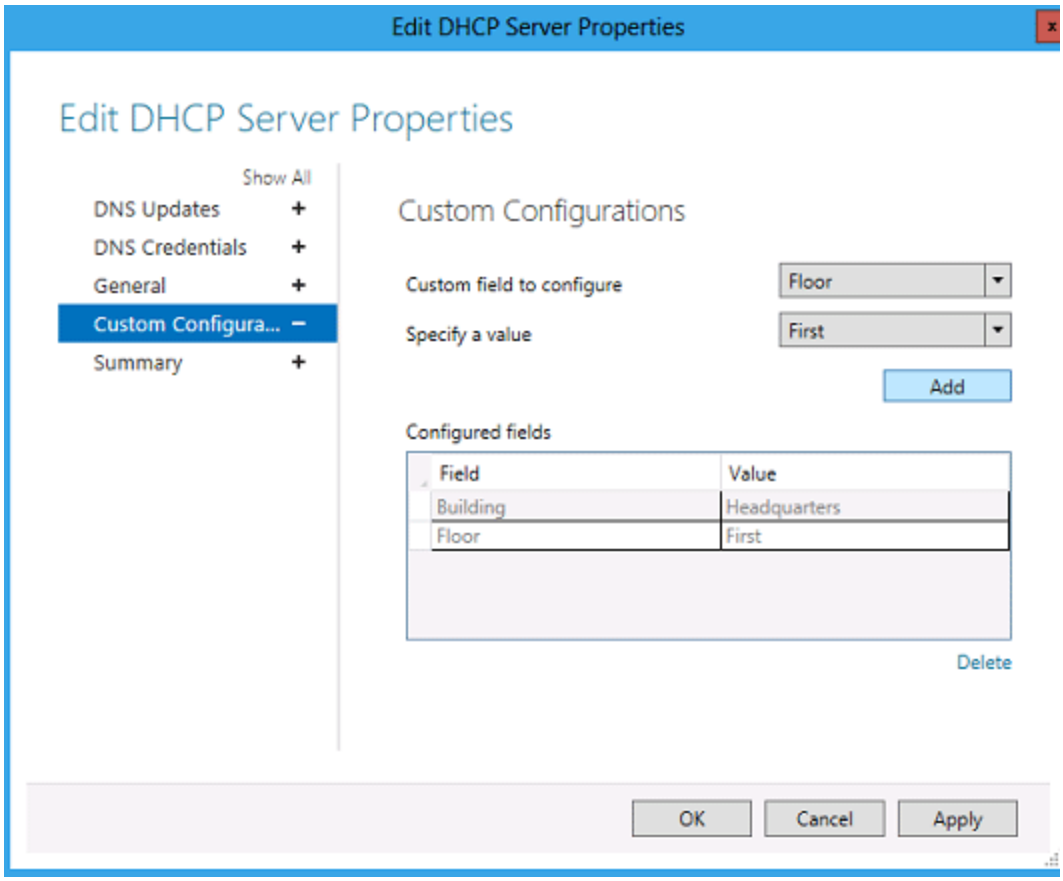

Review audit logs and events

IPAM also allows you to track several types of events on DNS and DHCP servers, including both client and server data.

To review audit logs and events

1. In the IPAM navigation menu, click **EVENT CATALOG**.
2. By default, **IPAM Configuration Events** is selected in the lower navigation pane. Review the events that are displayed.
3. Click **DHCP Configuration Events** in the lower navigation pane and review the DHCP events that are displayed.
4. Under **IP Address Tracking**, click **By Host Name**.
5. Type **Client1** in the search box, and then type dates in the two text boxes next to **and DHCP lease events between these dates** in the format of month/day/year. Enter a range of dates that includes today, and then click **Search**.

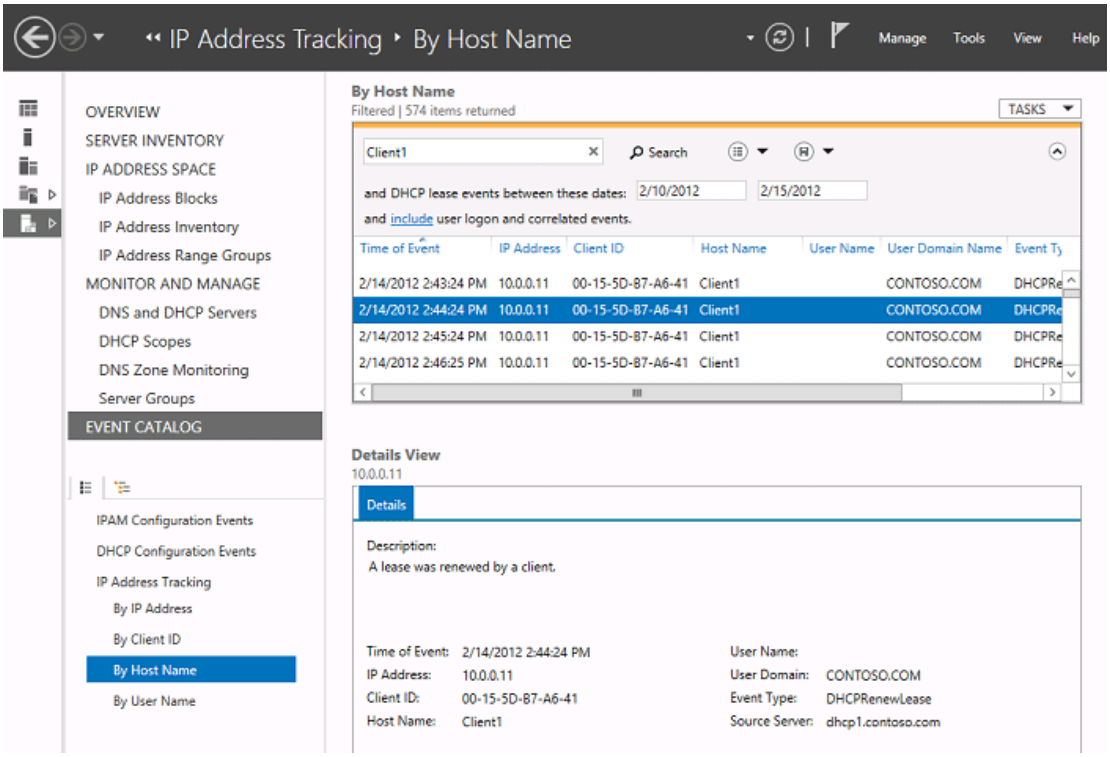

6. Click **TASKS** and then click **Export**.
7. In the **Save As** dialog box, type **client1-events** next to **File name** and then click **Save**.
8. Open the client1-events.csv file in notepad or Excel to view the list of events.

See also

[IP Address Management (IPAM) Overview]

What's New in IPAM in Windows Server 2012 R2

Walkthrough: Demonstrate IPAM in Windows Server 2012 R2

IPAM Planning and Design Guide

Microsoft introduces the IP Address Management (IPAM) feature in Windows Server® 2012. IPAM provides a built-in framework for discovering, monitoring, auditing, and managing IP address space and infrastructure servers on a corporate network.

For detailed, step-by-step deployment guidance, see the IPAM Deployment Guide.

For IPAM administration, troubleshooting, and best practices guidance, see the IPAM Operations Guide.

In this guide

This guide provides an overview of IPAM in Windows Server 2012, including design and planning information. The following topics are available in this guide:

- What is IPAM?: Provides an overview of the key features of IPAM.
- IPAM Terminology: Displays a list of new concepts and terms used in IPAM.
- Getting Started with IPAM: Provides a quick start for configuring and using IPAM.
- IPAM Architecture: Describes the primary components of IPAM and how they interact.
- IPAM Deployment Planning: Provides IPAM requirements, options and specifications.

A brief description of the benefits of IPAM is also provided below.

Benefits of IPAM

Planning, integrating, troubleshooting, and managing the IP infrastructure devices, systems, users, and networks can be a daunting task in business today. The number and variety of network enabled devices is growing exponentially, and increased adoption of technologies such as IPv6, VoIP, and network virtualization makes the job of network administration more complex than ever. Administrators are commonly required to manage multiple data centers at central offices, branch offices, and other remote locations.

IPAM provides a central and integrated experience for IP address management that replaces manual, work intensive tools such as spreadsheets and home-grown scripts that can be cumbersome, unreliable and difficult to scale.

Benefits and uses of IPAM include:

Server administration and management

- DHCP and DNS record creation and deletion
- Service and zone monitoring of DNS services
- Support for concurrent client sessions
- Remote administration support using Windows Server 2012 and Windows 8
- Support for backup, restore, and migration scenarios for IPAM server
- Multi-entity management and monitoring of DHCP services and DHCP scopes
- Server and server group custom views and tagging
- Role based access control
- Agentless management of server roles

For more information about using IPAM to administer DNS and DHCP servers, see Multi-server Management.

Planning and organization of address space

- IPv4 and IPv6 address space allocation
- Static IP address inventory and lifetime management
- IPv4 and IPv6 address space custom views and tagging
- IP address space utilization statistics and trend monitoring
- Organizing and visualizing of data into user-defined hierarchical logical groups
- Built-in relational database support leveraging Windows Internal Database (WID)
- Graphical reporting with Windows PowerShell export support
- Extensive support for user-defined and built-in custom fields

For more information about using IPAM to manage IP address space on your network, see Managing IP Address Space.

Flexibility and automation

- Automatic server role discovery using Active Directory
- Automatic server configuration data collection and dynamic address space discovery
- Group Policy based automated deployment
- Detailed data collection tasks with configurable periodicity
- Flexible support for import and export of IP address space from spreadsheets and management tools

- Support for update of IP address space from systems such as System Center Virtual Machine Manager (SCVMM) and non-Microsoft DHCP servers using Windows PowerShell

For more information about server discovery and automated data collection with IPAM, see IPAM Architecture and Managing Server Inventory.

Forensics

- IP address lease and logon event tracking
- Advanced search and filter support
- Configuration change event auditing for DHCP and IPAM services
- DHCP lease tracking by IP address, client ID, host name, and user name

For more information about using IPAM to monitor configuration changes and track IP address on the network, see Operational Event Tracking and IP Address Tracking.

See also

IPAM Operations Guide

IPAM Deployment Guide

Walkthrough: Demonstrate IPAM in Windows Server 2012

What is IPAM?

IP Address Management (IPAM) in Windows Server® 2012 is an integrated suite of tools to enable end-to-end planning, deploying, managing and monitoring of your IP address infrastructure, with a rich user experience. IPAM automatically discovers IP address infrastructure servers on your network and enables you to manage them from a central interface.

IPAM includes components for:

1. Address Space Management
2. Multi-Server Management and Monitoring
3. Network Audit

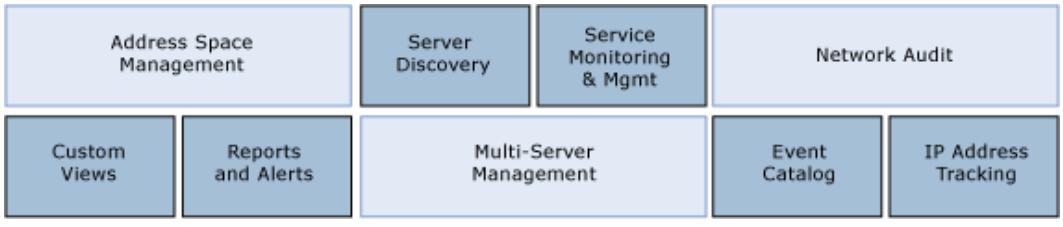

Address Space Management

IPAM's address space management (ASM) feature enables you to gain visibility into all aspects of your IP address infrastructure from a single console. With IPAM, you can create a highly customized, multi-level hierarchy of address space on your network and use it to manage IPv6 addresses and IPv4 public and private addresses. The ASM feature includes a robust reporting capability that enables detailed tracking of IP address utilization trends with customized thresholds and alerts.

Key features of ASM include the following.

- Integrated management of dynamic and static IP address space
- Detection and management of conflicts, overlaps, and duplicates in address space across systems
- Highly customizable inventory view of IP address space
- Centralized monitoring and reporting of address utilization statistics and trends
- Support for IPv4 and stateless IPv6 address utilization monitoring
- Automated discovery of IP address ranges from DHCP scopes
- Export and import of IP addresses and IP address ranges with Windows PowerShell support
- IP address usage alerts and notifications with custom thresholds
- Detection and assignment of available IP addresses

The following example shows how IPAM's ASM feature enables you to monitor IP address utilization. In this example, 7 days of utilization data is displayed for the 10.72.144.0/22 IP address range.

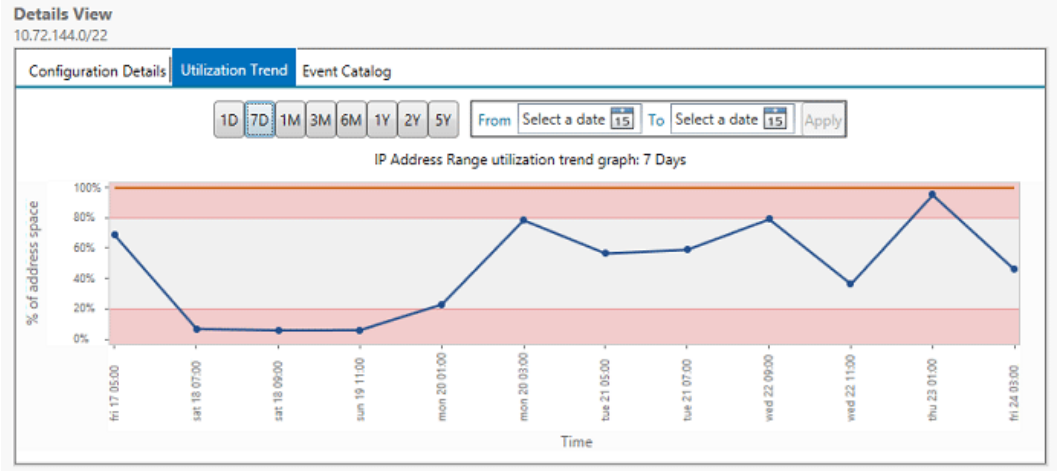

For more information, see Managing IP Address Space.

Multi-Server Management and Monitoring

IPAM's multi-server management (MSM) feature enables you to automatically discover DHCP and DNS servers on the network, monitor service availability, and centrally manage their configuration. Using the Group Policy provisioning mode, IPAM provides quick and painless provisioning of agentless IPAM access settings on managed servers. A manual provisioning mode is also available.

Key features of MSM include the following.

- Discovery of Microsoft DHCP and DNS servers automatically across an Active Directory forest
- Manual addition or removal of managed servers
- End-to end configuration and management of DHCP servers and scopes
- Support for advanced constructs to enable add, delete, overwrite, or find and replace operations on multiple DHCP scopes and servers
- Simultaneous update of common settings across multiple DHCP scopes or DHCP servers
- Availability monitoring for DHCP and DNS services and DNS zones
- Management of Microsoft DHCP and DNS servers running Windows 2008 or later operating systems
- Addition of custom information to servers enabling visualization using logical groups based on business logic
- Monitoring of DHCP scope utilization
- Automatic and on-demand retrieval of server data from managed DHCP and DNS servers
- DNS zone status monitoring based on DNS zone events
- Classify discovered servers and roles as managed or unmanaged

The following example shows how IPAM's MSM feature enables you to monitor IP DHCP scopes on the network. In this example, detailed data is displayed for the scope US_SEA_zzz3.

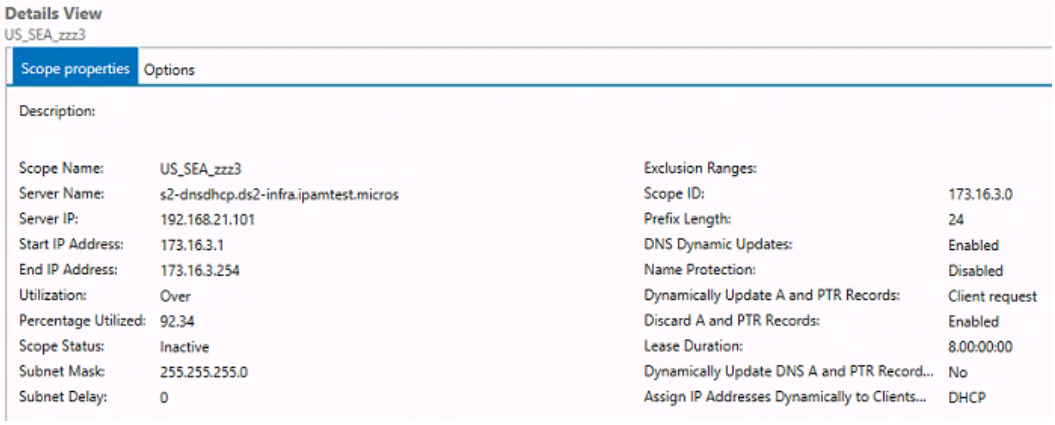

For more information, see Multi-server Management.

Network Audit

IPAM's audit feature provides a centralized repository for all configuration changes performed on DHCP servers and the IPAM server, and for IP addresses issued on the network. IPAM audit tools enable you to view potential configuration problems on DHCP servers by actively tracking and reporting all administrative actions. Detailed IP address tracking data is also provided, including client IP addresses, client ID, host name, and user name. Advanced search capabilities enable you to selectively search for events and obtain results that associate user logons to specific devices and times.

Key features of network audit include the following.

- Query the event catalog for DHCP configuration changes across multiple servers from a single console
- Track users, devices, and IP addresses for specified intervals with advanced queries using DHCP lease logs and logon events from domain controllers and network policy servers
- Track and report changes made to the IPAM server
- Export audit findings and create reports
- Quickly resolve configuration problems and track service level agreements

The following example shows how IPAM's network audit feature enables you to track IP addresses on the network. In this example, details are displayed for a lease event in the contoso.com domain.

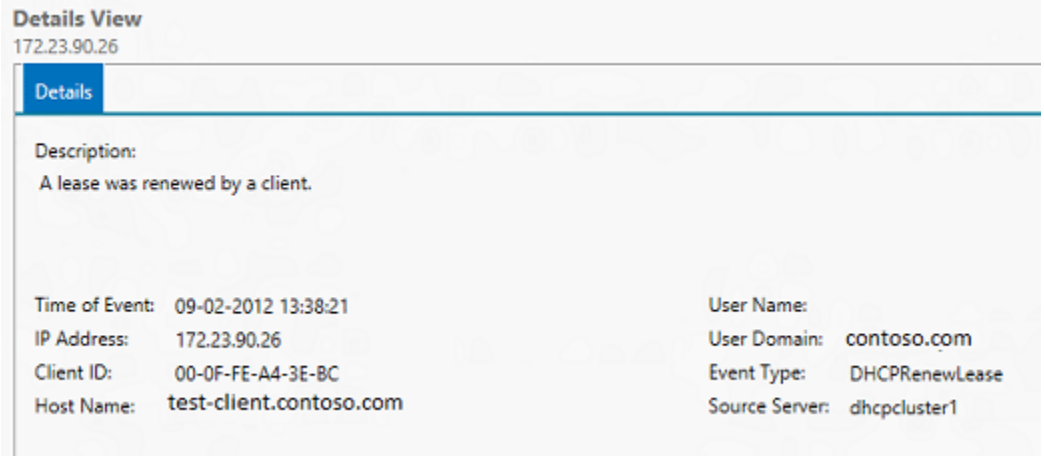

For more information, see IP Address Tracking and Operational Event Tracking.

See also

IPAM Terminology

Getting Started with IPAM

IPAM Deployment Planning

IPAM Terminology

Several new concepts and terms are used to describe IPAM. The following table displays a list of IPAM-related terms used in this guide.

Term	Definition
ASM	Address space management (ASM) in IPAM provides the ability to efficiently view, monitor, and manage IP address space on the network. The ASM feature supports IPv4 public and private address space, and IPv6 address space. You can manage dynamic IP addresses and static IP addresses on the network using IPAM ASM.
IP Address Block	An IP address block is the largest unit used for address space management. IPAM automatically labels IPv4 address blocks that you create as belonging to public or private address space as defined by the IANA. IP address blocks are comprised of smaller units called IP address ranges.
IP Address Inventory	IP Address Inventory is a node in the IPAM console that displays a built-in IP address range group (logical group) that is organized by the device type field of IP addresses.
IP Address Range	An IP address range is typically a DHCP scope or static IP address pool used to assign IP addresses to hosts. IP address ranges in IPAM are assigned to parent IP address blocks. IP address ranges do not contain other IP address ranges, but are comprised of individual IP addresses.
IP Address Range Groups	IP address range groups are one or more IP address ranges that are grouped together based on built-in or user-defined custom fields.
IPAM Administrators	IPAM Administrators is a local security group on an IPAM server that is created when you

Term	Definition
	install the IPAM feature. Members of this group have privileges to view all IPAM data and perform all IPAM tasks.
IPAM ASM Administrators	IPAM ASM Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group, and can perform IP address space tasks in addition to IPAM common management tasks.
IPAM IP Audit Administrators	IPAM IP Audit Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group. They can view IP address tracking data and perform IPAM common management tasks.
IPAM MSM Administrators	IPAM MSM Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group, and can perform server monitoring and management tasks in addition to IPAM common management tasks.
IPAM Users	IPAM Users is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have read-only access and cannot launch IPAM scheduled tasks. They can view all information in server inventory, IP address space, and the monitor and manage IPAM console nodes. IPAM Users can view IPAM and DHCP operational events under in the Event Catalog node, but cannot view IP address tracking data.
Managed By Service	Managed By Service is a built-in mandatory field for IP addresses and ranges that identifies the service that provides IP addresses in this range. An example value is **MS DHCP**, indicating that the IP address or range is

Term	Definition
	supplied by a Microsoft Windows DHCP server.
Managed Server	Managed servers are listed in the server inventory node of the IPAM console with a status of Managed. Domain controllers, DNS servers, DHCP servers, and network policy servers can be managed servers. The IPAM server must have access rights enabled on a managed server to perform management and monitoring actions. These access rights are enabled during the IPAM provisioning process.
MSM	Multi-server management (MSM) in IPAM provides periodic service monitoring and configuration of managed DHCP and DNS servers using the IPAM console.
Provisioning method	Provisioning method refers to the processes used to enable IPAM server access to managed servers. Available provisioning methods are: manual and automatic. The manual method requires that group memberships, file shares, service access settings, and firewall ports be configured manually on each managed server. The automatic provisioning method uses Group Policy Objects to configure these access settings automatically when a server is marked as managed in the server inventory view. IPAM provides a wizard that enables you to choose the provisioning method during the initial setup of the IPAM server. The provisioning method cannot be changed after you complete the wizard.
Service Instance	Service instance is a built-in mandatory field for IP addresses and ranges that identifies the device that provides IP addresses in this range. The service instance for a Microsoft Windows DHCP server is the fully qualified domain name of the server (ex: dhcp.contoso.com).
Server Groups	Server groups are logical groups used to organize managed DHCP and DNS servers.

Term	Definition
	Servers are members of a server group based on values of custom fields that are assigned to the server.
Unmanaged server	Unmanaged servers are listed in the server inventory node of the IPAM console with a status of Unmanaged. The IPAM server typically does not have access rights enabled to perform management and monitoring of unmanaged servers.
Unmapped Address Space	Unmapped address space is an IP address range or individual IP address that does not map to a parent IP address block. To automatically map a child IP address range to a parent IP address block, add the parent block to the IPAM console.

See also

What is IPAM?

Getting Started with IPAM

IPAM Architecture

IPAM Deployment Planning

Getting Started with IPAM

IPAM provides a dynamic view of your IP infrastructure, and the view is continually refreshed by periodic tasks that run on the IPAM server. IPAM also enables administrators to perform several configuration actions directly from the IPAM console.

IPAM is not enabled by default and must be installed as a server feature. You can install IPAM using the **Add Roles and Features Wizard** in Server Manager, or using Windows PowerShell. This topic summarizes how to begin using IPAM, and provides information about IPAM concepts and components. For detailed step by step guidance to install and provision IPAM, see the IPAM Deployment Guide.

The following steps summarize how to begin using IPAM:

#	Step	Summary	More Information
1.	Choose an IPAM server	Connect to an IPAM server using the IPAM client console.	The IPAM Client
2.	Choose a provisioning method	Use the IPAM provisioning wizard to select one of two provisioning methods: GPO-based (automatic) or manual.	Provisioning IPAM
3.	Configure the scope of discovery	Choose one or more domains in the forest that contain servers you wish to manage with the selected IPAM server.	Scope of discovery
4.	Start server discovery	Launch the server discovery task to retrieve a list of domain controllers, DNS servers, and DHCP servers from Active Directory.	Discovering servers
5.	Configure settings on managed servers	Enable IPAM access by creating and applying GPOs to configure managed server settings, or configure settings manually.	Configuring managed server access settings
6.	Select manageability status on managed servers	Choose servers and role services to manage from the list of discovered servers displayed in the server inventory.	Select managed servers
7.	Verify IPAM access	Verify that the IPAM server has access to manage the selected services on managed servers.	Verify managed server access

#	Step	Summary	More Information
8.	Retrieve data from managed servers	Launch data collection tasks on the IPAM server to retrieve data from managed servers.	Retrieve data
9.	Visualize data	Customize the IPAM client display to meet your needs.	Using the IPAM Client Console

The relationship between different components of IPAM is illustrated below.

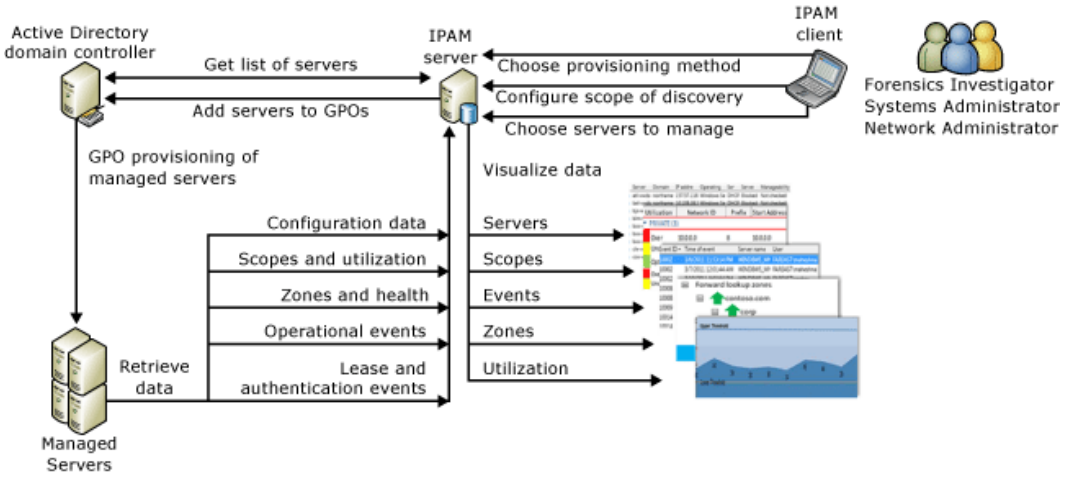

See the following topics for more information about IPAM components and processes:

The IPAM Client

The IPAM client console is integrated with Server Manager. IPAM Client is installed by default when you install IPAM Server on a computer running Windows Server® 2012. The client automatically connects to the local IPAM server after installation, but can also be used to manage a different IPAM server. You can also install IPAM Client on a computer running Windows® 8. For more information about requirements for installing IPAM Client, see IPAM Architecture.

After installing IPAM Client, you must add at least one IPAM server to manage using Server Manager. You can view the IPAM client console by clicking **IPAM** in the Server Manager navigation pane. An IPAM client computer can connect to any IPAM server that is available in the pool of servers managed by the Server Manager console. You cannot connect to multiple IPAM

servers simultaneously. To manage an IPAM server, click **Connect to IPAM server** in the IPAM console overview pane. When a successful connection is made, the fully qualified domain name (FQDN) of the IPAM server with the current domain and user ID is displayed in the IPAM OVERVIEW pane.

The IPAM client console provides administrators and users with multiple functions and custom views of IP address infrastructure services on the network. For more information about IPAM client console functionality, see Using the IPAM Client Console.

Provisioning IPAM

In IPAM, provisioning is the process of enabling required permissions, files shares, and access settings so that the IPAM server can monitor and manage IP infrastructure servers (called managed servers) on the network. Provisioning has two steps:

1. Choose a provisioning method.
2. Configure managed server access settings.

After choosing a provisioning method, you can begin discovering servers on the network. It is not required to configure managed server access settings prior to performing server discovery. However, you will not be able to monitor and manage servers until managed server access settings are applied. If you choose the automatic GPO-based provisioning method, you can create these GPOs immediately. No settings are applied until a server is assigned a status of managed and Group Policy is updated.

Choosing a provisioning method

To choose a provisioning method, click **Provision the IPAM server** in the IPAM console overview pane to launch the **Provision IPAM** wizard. The wizard allows you to choose either the manual or Group Policy based provisioning method. These provisioning methods are described below.

1. **Manual**: The manual provisioning method is typically used when the number of managed servers is small. If you choose the manual provisioning method, access settings must be configured individually on each managed server. Settings must also be removed manually if the server becomes unmanaged. You can use Group Policy to apply settings to managed servers even if the manual provisioning method is chosen, but you must apply and remove GPOs manually.

   The manual provisioning method is not preferred because it is more complex and less consistent than the Group Policy based method.

2. **Group Policy Based**: The Group Policy based method is preferred because it is simpler and less prone to errors. If you choose the Group Policy based provisioning method, GPOs are applied automatically to servers when they are assigned a status of managed in the IPAM console. GPOs are also removed automatically if the status of a server changes from managed to unmanaged.

   If you chose the Group Policy based method, you must also provide a GPO name prefix in the provisioning wizard. After providing a GPO name prefix, the wizard will display the GPO

names that must be created in domains that will be managed by IPAM. The following role-based GPOs are required in each domain that contains managed servers. The wizard does not create these GPOs.

- **<GPO-prefix>_DHCP**: This GPO is used to apply settings that allow IPAM to monitor, manage, and collect information from managed DHCP servers on the network.
- **<GPO-prefix>_DNS**: This GPO is used to apply settings that allow IPAM to monitor and collect information from managed DNS servers on the network.
- **<GPO-prefix>_DC_NPS**: This GPO is used to apply settings that allow IPAM to collect information from managed domain controllers and network policy servers (NPS) on the network for IP address tracking purposes.

Important
You cannot change the provisioning method for the current installation of IPAM Server after it has been selected. To change the provisioning method, you must uninstall and reinstall IPAM Server on the computer.

The following figure illustrates choices for applying settings to managed servers.

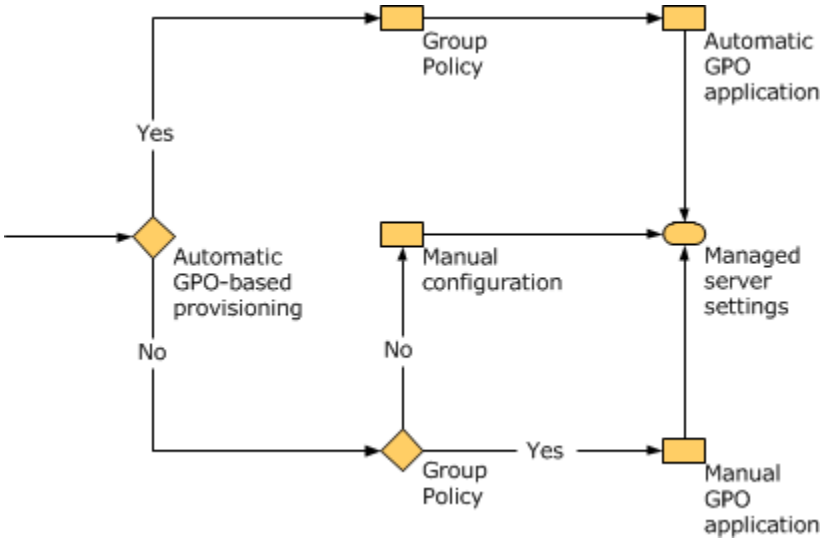

Configuring managed server access settings

The next step in IPAM provisioning is to configure group memberships, file shares, service access settings, and firewall ports on managed servers.

Depending on the provisioning method that was chosen, managed server access settings must either be configured individually on each managed server, or using IPAM provisioning Group Policy Objects (GPOs). The preferred method to create IPAM provisioning GPOs is with Windows PowerShell using the **Invoke-IpamGpoProvisioning** cmdlet. To create IPAM provisioning GPOs, you must have permission to configure Group Policy in the domain.

See the following example of how the **Invoke-IpamGpoProvisioning** cmdlet can be used to create IPAM provisioning GPOs. In this example, three GPOs are created (IPAM1_DHCP, IPAM1_DNS, IPAM1_DC_NPS) and linked to the contoso.com domain. These GPOs enable access for the server **ipam1.contoso.com** using the domain administrator account **user1**. Note: In this example, the hostname of the IPAM server is used as a GPO prefix, however this is not required.

```
Invoke-IpamGpoProvisioning –Domain contoso.com –GpoPrefixName IPAM1 –IpamServerFqdn
ipam1.contoso.com –DelegatedGpoUser user1
```

For more information about IPAM Server Cmdlets in Windows PowerShell, see http://go.microsoft.com/fwlink/p/?linkid=262977.

The following table summarizes access settings that are enabled on managed servers when you apply IPAM provisioning GPOs. These settings must be configured manually on managed servers if IPAM provisioning GPOs are not used. Servers that will not be managed by IPAM do not require any configuration changes.

Managed Server Role	IPAM Access Status	Access Settings
DHCP	DHCP RPC	The computer account for the IPAM server must be a member of the DHCP Users security group. The following firewall rules must be enabled: • DHCP Server (RPC-In) • DHCP Server (RPCSS-In)
DHCP	DHCP audit share	A network file share named **Dhcpaudit** must be created using the DHCP audit file folder, with Read access enabled for the computer account of the IPAM server. The following firewall rules must be enabled: • File and Printer Sharing (NB-Session-In) • File and Printer Sharing (SMB-In)
DNS	DNS RPC	Read access for the computer account of the IPAM server must be added to DNS discretionary access control list (DACL). The following firewall rules must be enabled: • DNS Service RPC • DNS Service RPC Endpoint Mapper
DHCP, DNS, domain controller, NPS	Event log	The computer account of the IPAM server must be a member of the Event Log Readers security group. The computer account for the IPAM server must be granted read access in the ACL that is maintained by the following registry key on the DNS server:

Managed Server Role	IPAM Access Status	Access Settings
		MACHINE\System\CurrentControlSet\Services\Eventlog\DNS Server\CustomSD. This only required on DNS servers. The following firewall rules must be enabled: • Remote Event Log Management (RPC) • Remote Event Log Management (RPC-EPMAP)
DHCP, DNS	DHCP and DNS service monitoring	The following firewall rules must be enabled: • Remote Service Management (RPC) • Remote Service Management (RPC-EPMAP)

As previously described, IPAM GPOs are applied automatically when the Group Policy based provisioning method is chosen. However, if you have chosen the manual provisioning method, you can also create IPAM provisioning GPOs and apply these GPOs manually. This ensures that the settings applied to managed servers in the domain are consistent. To apply GPOs manually, add managed servers to security filtering for each GPO depending on the role services to be managed.

After you have chosen a provisioning method, you can configure the scope of discovery and begin adding servers to manage with IPAM.

Server discovery

Server discovery has two steps:

1. Define the scope of discovery.
2. Discover servers on the network.

Scope of discovery

IPAM leverages Active Directory to define the scope of servers to be managed. To begin discovering servers, click **Configure server discovery** in the IPAM client console and choose at least one domain in the forest to discover. To choose a domain, select it from the drop-down list and then click **Add**.

In each domain chosen, specify the type of servers to discover. By default, domain controllers, DHCP servers, and DNS servers are discovered. In the following example, the root domain has been chosen with all managed server roles enabled for discovery.

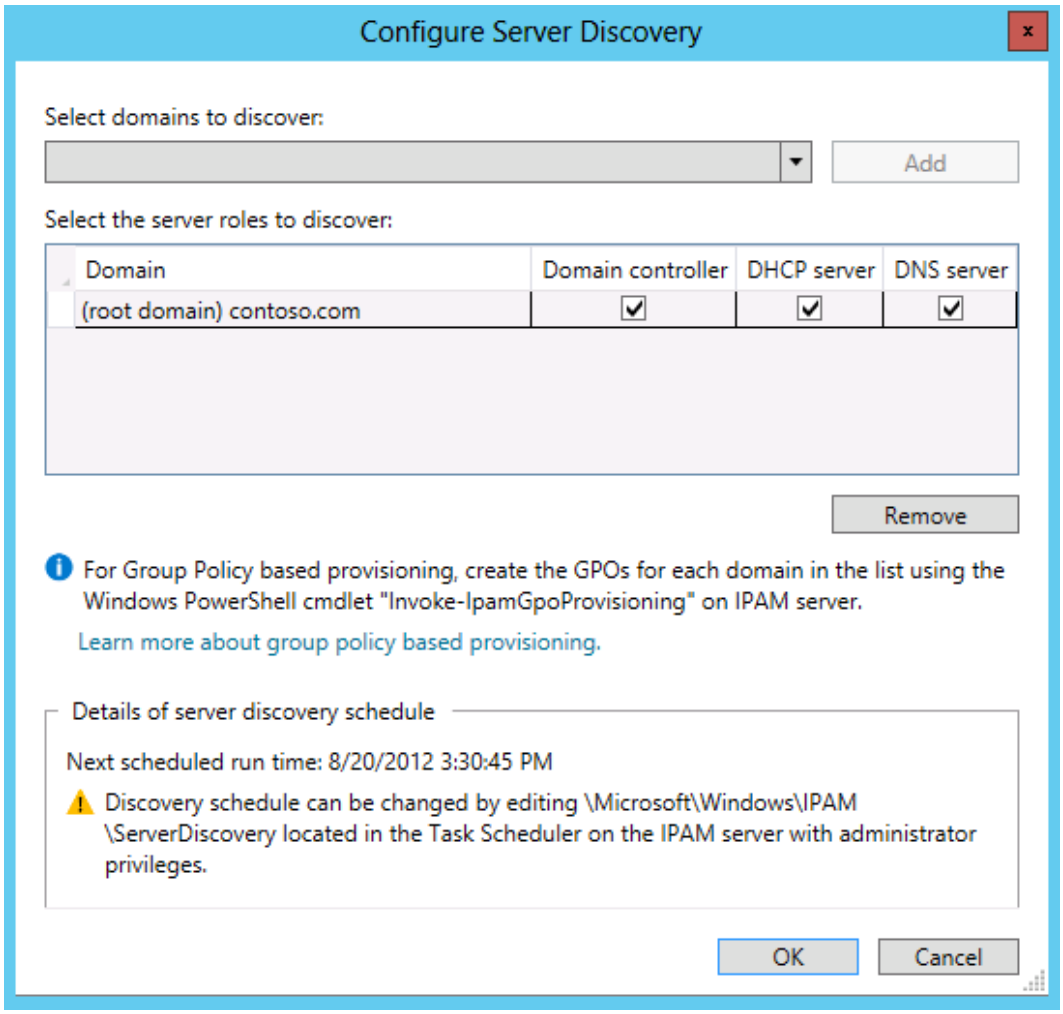

When you have completed adding domains and server roles to discover, click **OK**. The number of domains selected will be displayed in the IPAM console. You can add or remove domains from the scope of discovery at any time.

Discovering servers

To discover server roles, the server must be running Windows Server 2008 or later. To begin discovering servers on the network, click **Start server discovery** to launch the IPAM **ServerDiscovery** task. You can also click **Manage** on the Server Manager menu, and then click **Start Server Discovery**. This task runs automatically at a specified interval, and can also be started on-demand. For more information about IPAM tasks, see **Scheduled tasks**.

When the discovery task completes, click **Select or add servers to manage and verify IPAM access** to view the current list of servers in the IPAM server inventory. The following sections discuss how IPAM discovers different server roles on the network.

Domain controllers

DHCP servers

DNS servers

NPS servers

Domain controllers

Domain controllers are discovered in the domains you specify by querying Active Directory. If a domain controller is not discovered, verify that it is found in Active Directory. The following is an example query that displays a list of domain controllers running Windows Server 2008 or later. Replace contoso.com with the domain that is being discovered.

```
Get-ADDomainController -Filter {OperatingSystemVersion -ge "6.0"} -Server contoso.com
```

DHCP servers

IPAM discovers DHCP servers that are authorized in the Active Directory domains you specify and that respond to a **DHCPInform** message.

Important
If the DHCP server role is installed on the same server with IPAM, DHCP servers will not be discovered on the network.

To discover DHCP servers, the IPAM server reads the DHCP server list stored in the **DHCPServers** group found in the **NetServices** container (CN=NetServices,CN=Services,CN=Configuration,DC=domain,DC=com) in Active Directory. If a DHCP server is not discovered, verify that it is found in Active Directory and that the DHCP service is responding to requests. The following is an example query for DHCP servers that are authorized in Active Directory for the contoso.com domain.

```
Get-DhcpServerInDC | where-object {$_.dnsname -match "contoso.com"}
```

DNS servers

IPAM discovers authoritative DNS servers in the domains you specify by issuing a DNS query to the default DNS server configured on the IPAM server's network interface. If a DNS server is not discovered, verify that it is found in DNS. The following is an example query for authoritative DNS servers in the contoso.com domain.

```
nslookup -q=ns contoso.com
```

DNS servers must be authoritative and have a name server (NS) record present in the domain zone. Caching-only DNS servers are not discovered and cannot be managed with IPAM. Also note that if a DNS server is not running on a domain controller, additional procedures are required

to enable monitoring and management access by the IPAM server. See Verify managed server access for more information.

NPS servers

NPS servers are not automatically discovered by IPAM. However, you can add NPS servers to the server inventory by clicking **TASKS** and then clicking **Add Server**. In the **Add or Edit Server** dialog box, type the fully qualified domain name next to **Server name**, click **Verify**, select **NPS server**, choose the **Manageability status** and then click **OK**.

When you add a server manually to the server inventory, IPAM verifies that the server name exists in DNS, and that the IP address does not conflict with an existing server in the inventory.

Select managed servers

After auto-discovering or manually adding servers to the server inventory, you must choose whether or not they will be managed by the IPAM server.

Clicking **Select or add servers to manage and verify IPAM access** will display the current server inventory. To specify manageability status, right-click one or more servers in the inventory and then click **Edit Server**. The **Add or Edit Server** dialog box is displayed. Press and hold SHIFT or CTRL to select multiple servers.

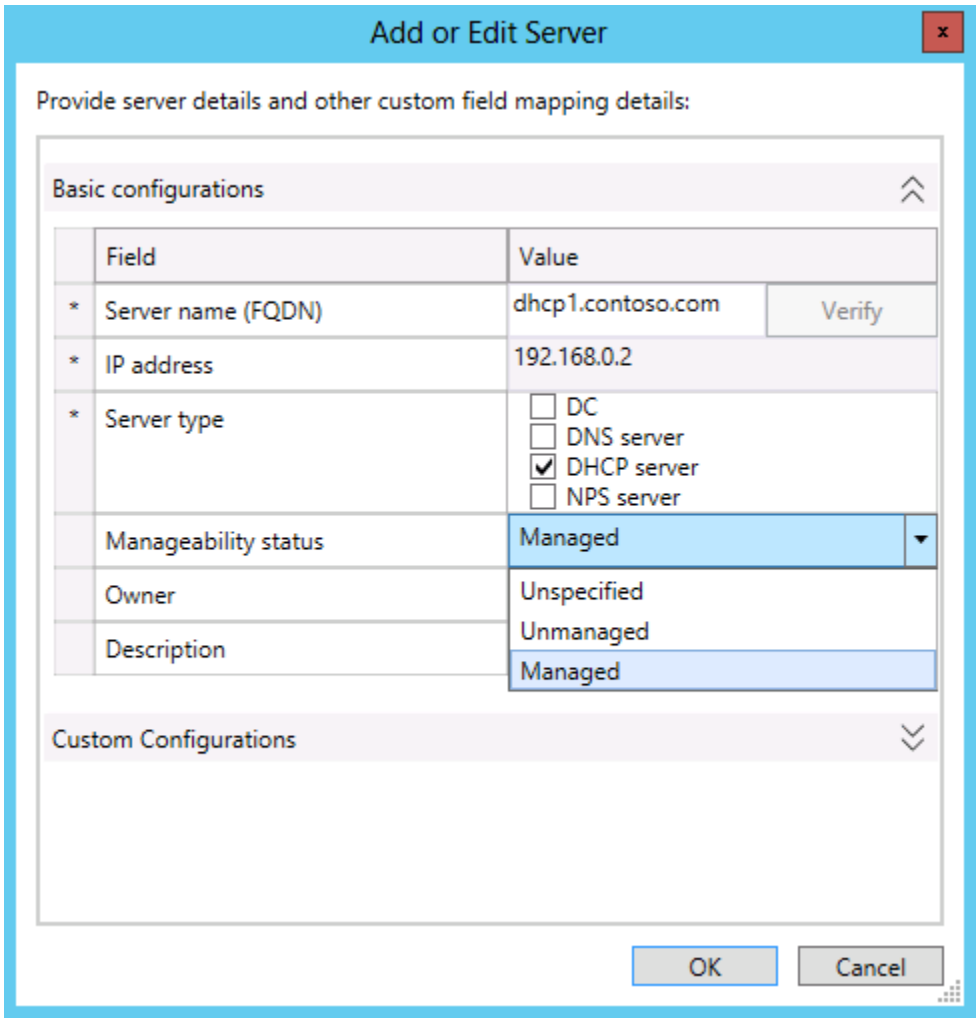

You can choose the following settings next to **Manageability status**:

1. **Unspecified**. This is the default setting and is intended to be temporary. Servers that are Unspecified typically do not have access settings applied. IPAM will not attempt to collect data from these servers, but retains information about the server in its database. You can use this setting when a server is offline temporarily, such as during a maintenance cycle.
2. **Unmanaged**. Choose this setting if a server will not be monitored and managed by the current IPAM server. Servers that are Unmanaged typically do not have access settings applied. IPAM does not attempt to collect data from these servers, and data from these servers is not maintained in the IPAM database.
3. **Managed**. Choose this setting if the server will be monitored and managed by the current IPAM server. Servers that are Managed typically have access settings applied, or will have access settings applied in the future. IPAM attempts to collect data from these servers

periodically based on the server roles that are selected, and maintains historical data for managed servers in the IPAM database.

If you are using the manual provisioning method, permissions and settings that enable access by the IPAM server can be applied to managed servers either before or after you choose their manageability status. If you are using the automatic GPO-based provisioning method, settings will be dynamically applied or removed depending on the manageability status you select.

After you have chosen the servers that will be managed by IPAM, the next step is to verify that these managed server have the correct access settings applied to allow the IPAM server permission to monitor and manage the server roles you have selected.

Verify managed server access

Access status is summarized in the **IPAM Access Status** column in server inventory. Possible values are:

1. **Not checked**: This status is displayed if a server is recently added or edited. This access status will be displayed until the **ServerDiscovery** task is run and the server inventory view is refreshed. To run the **ServerDiscovery** task, select one or more servers in the server inventory, right-click, and then click **Refresh Server Access Status**.
2. **Unblocked**: In order for access status to be unblocked, access to all settings required for the all roles selected next to **Server type** must be unblocked.
3. **Blocked**: This status indicates that one or more access settings are blocked. If some, but not all, access settings are unblocked, access status will be blocked.

Recommended Action	Manageability Status	IPAM Access Status	Server Name	DNS Suffix	Domain Name	Server Status	Server Type
IPAM Access Unblocked	Managed	Unblocked	DC1	contoso.com	contoso.com	No change	DC, DNS
IPAM Access Unblocked	Managed	Unblocked	dhcp1	contoso.com	contoso.com	No change	DHCP
IPAM Access Unblocked	Managed	Unblocked	dhcp2	contoso.com	contoso.com	No change	DNS, DHCP
IPAM Access Unblocked	Managed	Unblocked	nps1	contoso.com	contoso.com	Changed	NPS

To view the current status of all access settings, click a server in the server inventory view and review the information on the **Details** tab.

Tip

If access status is applied using GPOs, you might need to wait for Group Policy to be updated, or you can update Group Policy immediately using **gpupdate**. Verify that the appropriate GPOs are applied by typing **gpresult /r** at an elevated command prompt.

If managed server access settings have been recently applied, run the **ServerDiscovery** task and refresh the server inventory view. You can also wait for the view to refresh automatically.

The **Details View** displays additional information about IPAM access status that specifically applies to the managed server roles that are selected, including:

Status	Possible values
Manageability Status	Managed, Unmanaged, Unspecified
IPAM Access Status	Unblocked, Blocked, Not checked
Recommended Action	IPAM Access Unblocked, Unblock IPAM Access, Set Manageability Status
DHCP RPC Access Status	Unblocked, Blocked, Not applicable, Not checked
DHCP Audit Share Access Status	Unblocked, Blocked, Not applicable, Not checked
DNS RPC Access Status	Unblocked, Blocked, Not applicable, Not checked
Event Log Access Status	Unblocked, Blocked, Not applicable, Not checked

The IPAM access status for a managed server is a summary of all the applicable role-based access statuses displayed. For example, if the server type is [DNS, DHCP], and DHCP access status is unblocked but DNS access status is blocked, then the overall IPAM access status will be blocked. To determine why the IPAM access status is blocked, it is important to review information in the details view to determine the specific type of access that is blocked.

Verifying managed DHCP servers

DHCP service access settings require that the IPAM server is added to the DHCP Users security group. This is done automatically if you are using the Group Policy based provisioning method. However, since permissions for this group are applied when the DHCP Server service starts, you must restart the DHCP service if the server running IPAM does not already have this permission.

Note
IPAM operational event logging is not supported on DHCP servers running Windows Server 2008. In the server inventory view, these servers will display a status of Blocked. However, you can still create and edit scopes and server options on DHCP servers running Windows Server 2008.

Verifying managed DNS servers

To enable IPAM access to managed DNS servers, read access for the computer account of the IPAM server must be added to DNS discretionary access control list (DACL). However, the DACL is only applicable to Active Directory-integrated DNS servers. Therefore, you cannot perform this procedure if the DNS server is not also a domain controller. In order for IPAM to manage the DNS Server service when it is not running on a domain controller, you must add the computer account

of the IPAM server to the local administrators group on the DNS server. If the IPAM server is not a member of local administrators on the DNS server, then DNS RPC access will be blocked.

Retrieve data

Data needs to be collected from managed servers before you can begin visualizing this information. Data collection tasks run automatically on the IPAM server at regular intervals. For more information about these tasks, see IPAM Architecture.

You can also choose to run all IPAM data collection tasks on all managed servers immediately. To run these tasks, click **Retrieve data from managed servers** in the IPAM overview pane. To run all data collection tasks only on selected servers, select the servers from the server inventory view, right-click, and then click **Retrieve All Server Data**.

See also

What is IPAM?

IPAM Terminology

IPAM Architecture

IPAM Deployment Planning

IPAM Architecture

An IPAM infrastructure includes the following primary components:

IPAM Client: A computer running Windows Server® 2012 or Windows® 8.

IPAM Server: A computer running Windows Server 2012.

Managed servers: Domain controllers, NPS, DNS, and DHCP servers running Windows Server® 2008 or later.

IPAM components and interactions are illustrated in the following figure.

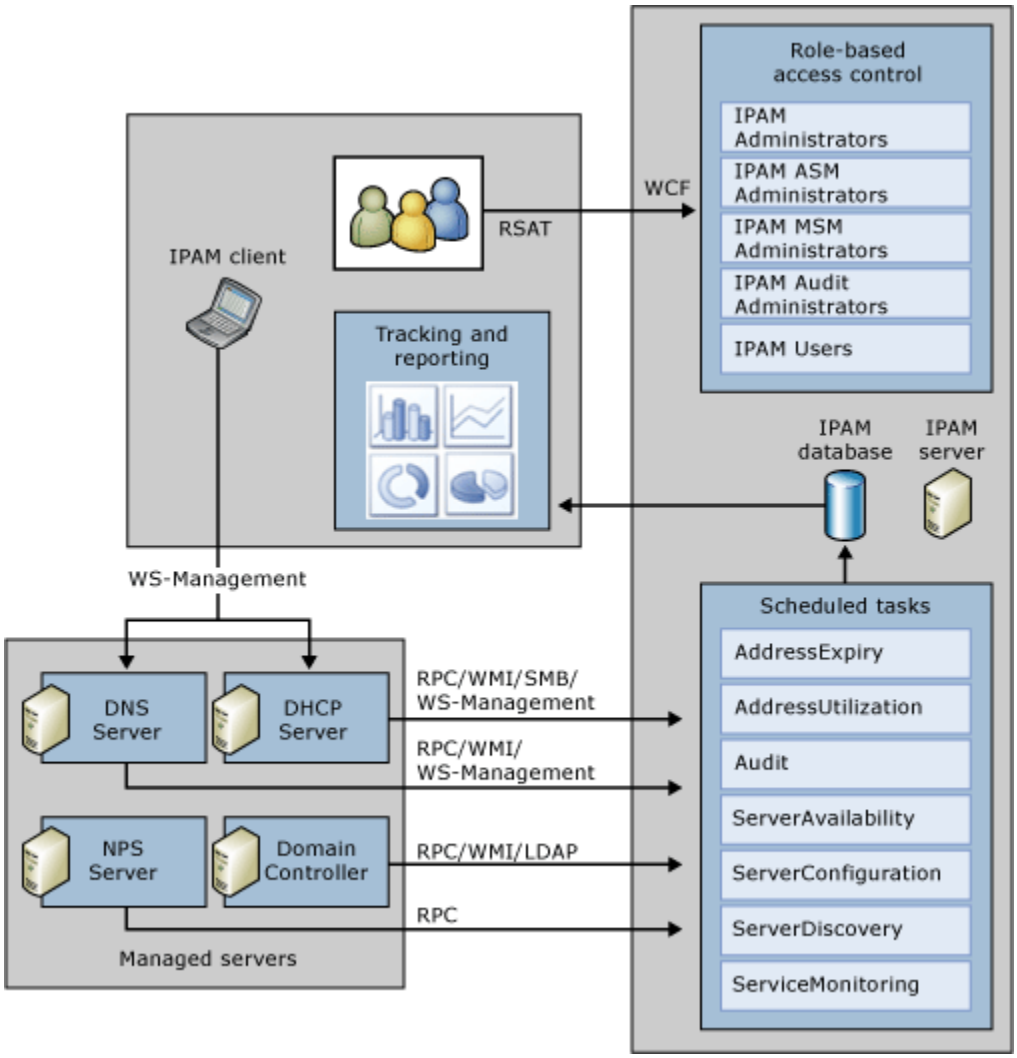

See the following topics for more information.

IPAM Client

An IPAM client computer is a computer running Windows Server 2012 or Windows 8 with the Remote Server Administration Tools (RSAT) for Windows 8 installed. If the computer is running Windows Server 2012, IPAM Client is automatically selected for installation under **Features > Remote Server Administration Tools > Feature Administration Tools > IP Address Management (IPAM) Client** when you install IPAM Server.

In its default configuration, IPAM Client is installed on the same computer with IPAM Server, however this is not required. You can also install IPAM Client on a computer running Windows 8 by installing the Remote Server Administration Tools (RSAT) for Windows 8

(http://go.microsoft.com/fwlink/p/?LinkID=238560). You can also install IPAM Client on more than one computer. IPAM Server supports multiple concurrent connections from different IPAM Clients. A single IPAM Client can connect to one IPAM Server at a time.

Available IPAM installation options are summarized in the following table.

Operating system	Configuration	Available IPAM connections
Windows Server 2012	IPAM Server + IPAM Client	Connect to the local IPAM server or a remote IPAM server from the local client
Windows Server 2012	IPAM Client only	Connect to a remote IPAM server from the local client
Windows Server 2012	IPAM Server only	Connect to the local server from a remote client
Windows 8	RSAT for Windows 8	Connect to a remote IPAM server from the local client

Client communications

The IPAM client communicates with an IPAM server using the Windows Communication Foundation (WCF) protocol with TCP as the transport method. By default, TCP binding is performed on port 48885 on the IPAM server. If there is a port conflict, or there is a need to reconfigure the server port, the port number on the server can be configured using the Windows PowerShell cmdlets **Get-IpamConfiguration** and **Set-IpamConfiguration**. For more information about IPAM and Windows PowerShell, see Using Windows PowerShell with IPAM.

IPAM Server

An IPAM server is a domain member computer running Windows Server® 2012 or a later operating system. You cannot install IPAM Server on a domain controller, and it is not recommended to install IPAM Server on the same server with DHCP Server. For more information about specifications and requirements for IPAM Server, see IPAM Deployment Planning.

Server communications

The IPAM server uses Lightweight Directory Access Protocol (LDAP) to discover domains and DHCP servers in the Active Directory forest. The Session Message Block (SMB) protocol is used to retrieve IP address lease information from the audit log file on DHCP servers.

The IPAM server communicates with managed DHCP servers to get DHCP scope utilization, scope configuration, server configuration information, and DHCP operational events.

Configuration information is acquired by Windows PowerShell for DHCP Server using [MS-DHCPM]. Specific communication ports must be enabled in Windows Firewall for this communication to be successful.

Note
For information about [MS-DHCPM] and other Microsoft protocol specifications, see the Technical Specification Cross-Reference Matrix on MSDN.

IPv4 and IPv6 address lease information is obtained from an audit log file on the DHCP server using SMB, and the IPAM server also regularly retrieves service status for DHCP and DNS servers using [MS-SCMR]. The IPAM server communicates with DNS servers to get server configuration and DNS zone settings with Windows PowerShell commands for DNS Server using [MS-DNSP], respectively.

The IPAM server communicates with domain controllers and NPS using [MS-EVEN6] to collect user and computer logon events that are generated when a network access request is authenticated.

Scheduled tasks

The IPAM server maintains a local database that is dynamically updated by regularly scheduled data collection tasks. These tasks run on the IPAM server at configurable intervals and communicate with managed servers (domain controllers, NPS, DNS, and DHCP servers) using an agentless architecture. IPAM scheduled data collection tasks are summarized below. These tasks can be viewed and modified in Task Scheduler by navigating to **Microsoft > Windows > IPAM**.

Task Name	Description	Default Frequency	Duration
AddressExpiry	Tracks IP address expiry state and logs notifications.	1 day	Indefinite
AddressUtilization	Collects IP address space usage data from DHCP servers to display current and historical utilization.	2 hours	Indefinite
Audit	Collects DHCP and IPAM server operational events. Also collects events from domain controllers, NPS, and DHCP servers for IP address tracking.	1 day	Indefinite

Task Name	Description	Default Frequency	Duration
ServerAvailability	Collects service status information from DHCP and DNS servers.	15 minutes	Indefinite
ServerConfiguration	Collects configuration information from DHCP and DNS servers for display in IP address space and server management functions.	6 hours	Indefinite
ServerDiscovery	Automatically discovers the domain controllers, DHCP servers, and DNS servers in the domains you select.	1 day	Indefinite
ServiceMonitoring	Collects DNS zone status events from DNS servers.	30 minutes	Indefinite

The IPAM database

Information in the IPAM database is regularly updated by data collection tasks, and can be modified by administrators. For example, when an administrator creates a custom logical group or assigns an expiration date to an IP address, this information is stored in the IPAM database and associated to specific managed servers.

The IPAM database leverages the Windows Internal Database (WID) feature, a relational data store that is automatically installed when you install IPAM Server. IPAM does not support the use of an external database. For more information about WID, see Windows Internal Database Overview.

Role-based access control

When you install IPAM, local security groups are created on the IPAM server to provide role-based access control for different sets of IPAM administrators and users. IPAM uses these role-based access controls to determine what information is displayed in the IPAM client console. For example, viewing of IP address lease data can be restricted to a specific set of administrators by adding their user account to the IPAM IP Audit Administrators or IPAM Administrators group.

The following local user groups are created when you install IPAM:

Group Name	Description
IPAM Users	IPAM Users is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group can view all information in server inventory, IP address space, and the monitor and manage IPAM console nodes. IPAM Users can view IPAM and DHCP operational events under in the Event Catalog node, but cannot view IP address tracking data.
IPAM MSM Administrators	IPAM MSM Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group, and can perform server monitoring and management tasks in addition to IPAM common management tasks.
IPAM ASM Administrators	IPAM ASM Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group, and can perform IP address space tasks in addition to IPAM common management tasks.
IPAM IP Audit Administrators	IPAM IP Audit Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have all the privileges of the IPAM Users security group. They can view IP address tracking data and perform IPAM common management tasks.
IPAM Administrators	IPAM Administrators is a local security group on an IPAM server that is created when you install the IPAM feature. Members of this group have privileges to view all IPAM data and perform all IPAM tasks.

Managed servers

Managed servers must be running Windows Server 2008 or above.

The IPAM server communicates on a periodic basis with DHCP servers, DNS servers, domain controllers, and NPS using remote management technologies including: remote procedure call (RPC), Windows Management Instrumentation (WMI), Server Message Block (SMB) and Web Services for Management (WS-Management). This eliminates the need for deploying and maintaining dedicated management agents on each server, but requires that specific firewall ports be made available to the IPAM server.

The following table summarizes communications between the IPAM server and the specified managed server roles.

Server Role	Communication Methods	Actions
DHCP	RPC, WMI, SMB, WS-Management	• Monitor: IP address utilization • Monitor: Service status • Configure: Servers and scopes • Audit: IP address lease events • Audit: DHCP configuration events
DNS	RPC, WMI, WS-Management	• Discover: DNS servers • Monitor: DNS zone health • Configure: DNS zones
Domain Controller	RPC, WMI, LDAP	• Discover: DNS and DHCP servers • Validate: Managed servers • Audit: User and device logon events
NPS	RPC	• Audit: User and device logon events

Protocol specifications used by the IPAM server to communicate with managed servers include Microsoft Dynamic Host Configuration Protocol (DHCP) Server Management Protocol Specification [MS-DHCPM], Domain Name Service (DNS) Server Management Protocol Specification [MS-DNSP], EventLog Remoting Protocol Version 6.0 Specification [MS-EVEN6],

and Service Control Manager Remote Protocol Specification [MS-SCMR]. These specifications use RCP as the underlying communication protocol.

The IPAM client communicates with managed servers to remotely manage DHCP and DNS servers using Windows PowerShell provider with [MS-DHCPM] and [MS-DNSP]. The IPAM client also uses Windows PowerShell with [MS-PSRP] to query the IPAM server port configuration and [MS-GPOL] to configure Group Policy.

See also

What is IPAM?

IPAM Terminology

Getting Started with IPAM

IPAM Deployment Planning

IPAM Backup and Restore

IPAM Deployment Planning

IPAM is an agentless multi-server, multi-service management feature that leverages standard Windows remote management protocols to manage, monitor and collect data from IP address infrastructure servers. IPAM relies on a host of remote management technologies to provide full functionality. Communication with multiple network elements throughout the enterprise is required for data gathering and configuration management. Depending on the scope of managed elements, this communication may need to traverse multiple security boundaries or domains.

Note
IPAM must be installed on a domain member computer. You cannot install IPAM on a domain controller. If IPAM is installed on the same server with DHCP, then DHCP server discovery will be disabled.

An IPAM server provides support for a single Active Directory forest. Multi-forest topologies are not supported. Multiple IPAM servers can support a single domain, or a single IPAM server can support all domains in an Active Directory forest.

In this topic

Deployment topologies

Hardware and software requirements

IPAM specifications

Capacity planning

Deployment topologies

IPAM supports the following topologies for deployment in the enterprise:

- **Distributed**: An IPAM server deployed at every site in the enterprise.
- **Centralized**: One IPAM server in the enterprise.
- **Hybrid**: A central IPAM server deployed with dedicated IPAM servers at each site.

The following diagram illustrates an example of the hybrid deployment model.

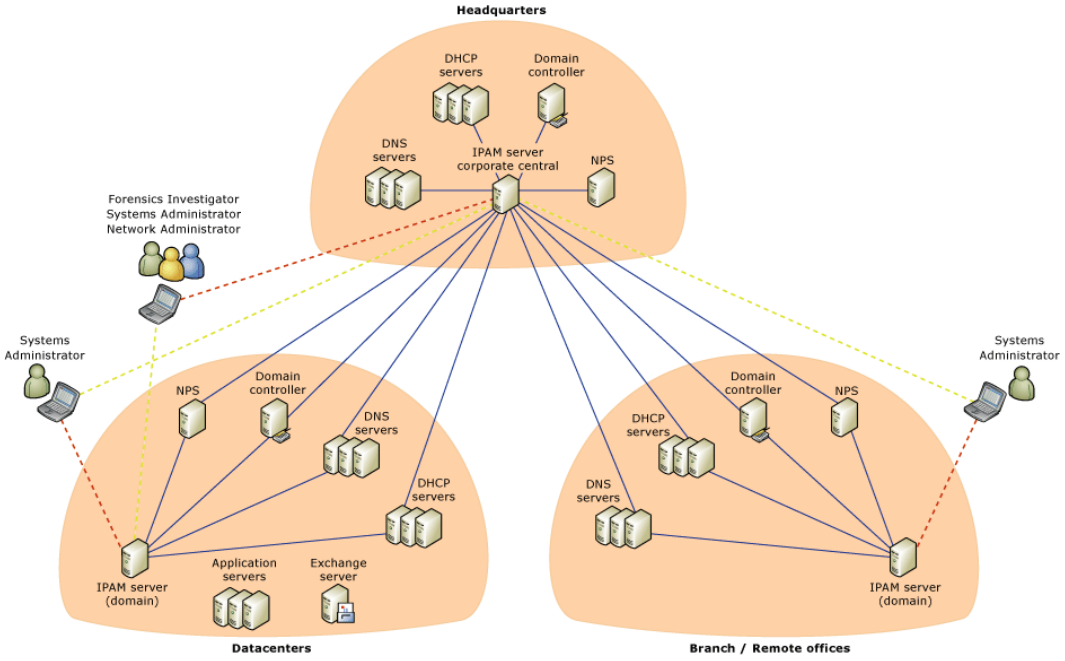

There is no automatic built-in communication or database sharing between different IPAM servers. If multiple IPAM servers are deployed, you can customize the scope of discovery for each IPAM server, or filter the list of managed servers.

If desired, you can leverage export and import functions in Windows PowerShell for IPAM to periodically update IP address range and address information between multiple IPAM servers.

You can also customize the role of different IPAM servers that are deployed. For example, a single IPAM server might be implemented to manage IP addressing for the entire enterprise. A different IPAM server might be used to monitor DNS zone health or configure DHCP scopes. Alternatively, you can limit the discovery and management scope to create a dedicated IPAM server that will perform all functions but only for a specific group of managed servers. The scope of management assigned to an IPAM server is flexible and can be updated if the need arises by adding or removing managed servers and domains.

Hardware and software requirements

IPAM Server must be installed on a computer running Windows Server® 2012. IPAM Client must be installed on a computer running Windows Server 2012 or Windows® 8. To install IPAM Client

on a computer running Windows 8, you must first install the Remote Server Administration Tools for Windows 8 (http://go.microsoft.com/fwlink/p/?LinkID=238560).

Active Directory: An IPAM server must be joined to a domain as a domain member server. Installation in a workgroup environment is not supported, and installation on a domain controller is not supported.

Network: An IPAM server requires a functional networking environment that includes IPv4 and IPv6 network connectivity to integrate with existing network services in the Active Directory forest. Server discovery requires that network settings on the IPAM server be configured to provide access to at least one domain controller and authoritative DNS server. Discovery of IPv6 address space requires that IPv6 is enabled on the IPAM server. The IPAM server must also have network connectivity to all servers that are marked as managed in the server inventory.

Other roles or features: An IPAM server is intended as a single-purpose server. It is not recommended to collocate other network infrastructure roles such as DNS or DHCP on the same server. IPAM installation is not supported on a domain controller, and discovery of DHCP servers will be disabled if you install IPAM on a server that is also running the DHCP Server service. The following features and tools are automatically installed when you install IPAM Server.

Feature or Tool	Description
Remote Server Administration Tools	DHCP and DNS Server Tools and IP Address Management (IPAM) Client provides for remotely managing DHCP, DNS and IPAM servers.
Windows Internal Database	Windows Internal Database is a relational data store that can be used only by Windows roles and features.
Windows Process Activation Service	Windows Process Activation Service generalizes the IIS process model, removing the dependency on HTTP.
Group Policy Management	Group Policy Management is a scriptable Microsoft Management Console (MMC), providing a single administrative tool for managing Group Policy.
.NET Framework 4.5 Features	.NET Framework 4.5 provides a programming model for building and running applications designed for several different platforms.

The following are the minimum and recommended hardware requirements for IPAM Server.

Component	Requirement
Processor	Minimum: 1.4 GHz (x64 processor) Recommended: quad-core, 2.66 GHz or faster
Memory	Minimum: 2 GB RAM Recommended: 4 GB RAM or greater
Disk Space	Minimum: 10 GB Recommended: 80 GB or greater*

*Use a fast storage device to host the IPAM database (on the root drive) to significantly improve IPAM performance.

Actual hardware requirements will vary based on the number of managed servers that are monitored and managed by the IPAM server.

IPAM specifications

IPAM Server has the following specifications:

1. The scope of IPAM server discovery is limited to a single Active Directory forest. The forest itself may be comprised of a mix of trusted and untrusted domains.
2. IPAM supports only Microsoft domain controllers, DHCP, DNS, and NPS servers running Windows Server® 2008 and above.
3. DHCP operational event auditing is supported for DHCP servers running Windows Server® 2008 R2 and above.
4. IPAM installation on a DHCP server is not recommended. The IPAM server discovery feature will not be able to discover DHCP roles if DHCP Server is installed on the same computer.
5. IPAM supports only domain joined DHCP, DNS and NPS servers in a single Active Directory forest.
6. IPAM does not support management and configuration of non-Microsoft network elements.
7. IPAM does not support external databases. Only a Windows Internal Database is supported.
8. A single IPAM server can support up to 150 DHCP servers and 500 DNS servers.
9. A single IPAM server has been tested to support up to 6000 DHCP scopes and 150 DNS zones.
10. IPAM stores 3 years of forensics data (IP address leases, host MAC addresses, user login/logoff information) for 100,000 users in a Windows Internal Database. There is no database purge policy provided, and the administrator must purge data manually as needed.
11. IP address utilization trends are provided only for IPv4.
12. IP address reclaiming support is provided only for IPv4.
13. No special processing is done for IPv6 stateless address auto configuration private extensions.
14. No special processing for virtualization technology or virtual machine migration.

15. IPAM does not check for IP address consistency with routers and switches.
16. IPAM does not support auditing of IPv6 stateless address auto configuration on an unmanaged machine to track the user.
17. IPAM users must be logged in using domain credentials. Do not log sign in to the IPAM server using the local Administrator account or another local user account on the IPAM server.
18. If you are accessing the IPAM server from a remote IPAM client, you must be a member of the **WinRMRemoteWMIUsers** group on the IPAM server, in addition to being a member of the appropriate local IPAM security group.
19. If the Group Policy based provisioning method is used, users must have domain administrator privileges to mark servers as managed or unmanaged in the server inventory.

Capacity planning

When planning disk space requirements and determining the number of IPAM servers to use on your network, consider the following questions:

1. How many IP address ranges will be managed by IPAM?
2. How many DHCP-enabled devices are connected to the network?

If the number of IP address ranges that you plan to manage with IPAM is less than 20,000 each for IPv4 and IPv6, you can deploy a single IPAM server to manage your IP address space. This is assuming a typical IP address range prefix is /24 for IPv4 and /64 for IPv6.

However, you might wish to deploy more than one IPAM server to manage an expanding network. For example, if the organization has 30,000 IP address ranges and is increasing in size, you should deploy at least two IPAM servers. You might also plan to deploy more than one IPAM server if you will assign specific roles to different servers or you will deploy IPAM in a distributed or hybrid topology. For more information, see Planning to Deploy IPAM.

Planning disk capacity

IPAM uses several data collection tasks to gather data from managed servers. This data includes information such as DHCP scopes, DHCP scope utilization, DNS zones, DNS zone events, DHCP lease logs, IPAM and DHCP configuration events, and network authentication events. Data collection tasks run in background and regularly update the local IPAM database, increasing its size. The IPAM database is located on the system root drive (the operating system drive), therefore it is important to ensure enough disk space is available on the system root drive to accommodate this data.

Note

For purposes of planning disk capacity, a system lifetime of 5 years is assumed.

Disk space requirements for the IPAM database can be evaluated based on three categories of stored data:

1. **Base database size**: This is the disk space required to hold IP address blocks, IP address ranges, IP address records, custom fields, DHCP configuration data, DHCP scopes, and

other static managed server information. This type of data doesn't increase over time; it grows only when more records are created in IPAM by administrators or if new managed servers are added. You should allocate 1.0 GB of free space on the root drive to accommodate this data.

2. **Utilization data**: IPAM keeps a track of utilization for IP address blocks, IP address ranges and DHCP scopes by periodically sampling and storing utilization statistics for these items. Each sample is stored in the local IPAM database to enable graphical display of utilization trends. The amount of data collected and stored depends on the number of IP address ranges in the system. Monthly usage is about 1.0 GB of data for every 10,000 IP address ranges. There is no simple method for clearing or purging utilization data from the system. Therefore, you should plan disk space consumption for a period of 5 years or the anticipated system lifetime. For example, if your IPAM server will manage 2000 IP Address ranges for a period of 5 years, you must allocate 1*5*12 = 60 GB of disk space on the system root drive to accommodate utilization data.

# of IP Address Ranges	Disk Space Required
10,000 or less	1 GB / month
10,001 to 20,000	2 GB / month
20,001 to 40,000	3 GB / month

3. **Event catalog data**: IPAM collects DHCP lease logs, DHCP configuration events, IPAM configuration events, and authentication events from all managed NPS servers and domain controllers. Depending on the size of your network, this data can require several GB of disk space. The amount of disk space required for storing event catalog data depends upon frequency at which these events occur on the network. To calculate disk space requirements for event catalog data, you must estimate the number of events per month together with the length of time you wish to maintain these events in IPAM database. Approximately 0.6 GB of free disk space is required for one million events.

   For example, on a network with 2500 users where each user has a laptop, a desktop, and a smart phone, you can allocate 4 IP addresses to each user: one IP address each for the desktop computer, one for the phone and two for the laptop (wired + wireless). This means you will have about 10,000 IP addresses to manage. Also, assuming that a DHCP lease on wireless interfaces will be renewed every 8 hours and every 4 days on wired interfaces, you can expect three lease events per day (24/8) for the wireless devices and one lease event on a wired interface in 4 days. Over a period of four days you can estimate a total of 2500*2*3*4 (wireless) + 2500*2*1*4 (wired) = 80,000 DHCP lease events corresponding to 80000*30/4 =600000 DHCP lease events per month.

   Similarly, if you have 2500 users on the network, and expect each user to login, lock, and unlock their device 20 times per day, you can expect 2500*20*22 = 1100000 user authentication events per month (assuming 22 working days in month). This corresponds to 600000+1100000= 1700000 records per month. To retain data initially for six months, and then purge three month old data on a quarterly basis, the disk must be able to store

1700000*6 = 10200000 records, corresponding to 0.6*10200000/1000000= 6.12 GB of disk space.

Note
These calculations do not take into account DHCP and IPAM configuration change events, machine authentication events, and DHCP events such as granting and expiration of leases. These events do not typically impact disk space requirements. However, you can take these events into consideration if required.

To calculate overall disk space requirements, use the following formula:

Disk space = Base database size + Utilization data + Event catalog data

In the example used previously, a network with 2500 users will require 1 GB + 60 GB + 6.12 GB = 67.12 GP of free disk space for the IPAM database, storing utilization data for 5 years.

Organization Size	Base Size	Utilization Data	Event Catalog Data	Total
25,000 users	1 GB	60 GB	59 GB	120 GB
50,000 users	1 GB	60 GB	117 GB	178 GB
100,000 users	1 GB	60 GB	234 GB	295 GB

Assumptions:

- 4 IP addresses are allocated per user.
- All IP addresses are issued by managed DHCP servers. Half of the IP addresses are wireless with leases refreshing every 8 hours; the other half are wired interfaces with leases refreshing every 4 days.
- Each user generates 20 domain controller authentication events and 1 NPS authentication event per working day.
- Event catalog data is purged every three months. Data is initially collected for 6 months and then three month old data is purged quarterly.
- There is an average of 22 working days per month.

For information about purging event catalog data, see EVENT CATALOG.

See also

What is IPAM?

IPAM Terminology

Getting Started with IPAM

IPAM Architecture

IPAM Backup and Restore

IPAM Deployment Guide

IP Address Management is a built-in framework provided in Windows Server® 2012 for discovering, monitoring, auditing, and managing the IP address space used on a corporate network. With IPAM, you can gain visibility and control over all aspects of your IP address infrastructure from a single console.

About this guide

This guide is intended for use by system administrators and system engineers. It provides detailed guidance for deploying IPAM Server and IPAM Client based on a design that has been chosen. If your organization has not yet selected a design, see IPAM Deployment Planning. You can then use this guide to deploy IPAM in your production environment.

This guide provides steps for deploying IPAM using a distributed, centralized, or hybrid design using a Group Policy based or manual provisioning method.

Use the checklists in Implementing Your IPAM Design Plan to determine how best to use the instructions in this guide to deploy a particular design and provisioning method. For information about hardware and software requirements for deploying IPAM, see the Hardware and software requirements section in the IPAM Deployment Planning topic.

In this guide

Planning to Deploy IPAM

Implementing Your IPAM Design Plan

Checklist: Deploy IPAM

Deploying IPAM Server

Deploying IPAM Client

Assigning IPAM Server and Administrator Roles

Terminology used in this guide

For a list of IPAM-related terms, see IPAM Terminology.

See Also

IPAM Planning and Design Guide

IPAM Operations Guide

Walkthrough: Demonstrate IPAM in Windows Server 2012

Planning to Deploy IPAM

Consider the following when planning your IPAM deployment:

- **Deployment topology**: One or more IPAM servers can be deployed at every site in the organization, or a single IPAM server can be used to manage multiple sites. Choose a design that best fits your needs.
- **IPAM staging strategy**: You can deploy IPAM gradually in your environment by specifying the scope of discovery and selecting a subset of managed servers from the server inventory. Expand your IPAM deployment by including additional domains and servers to the management scope of existing IPAM servers, or by adding additional IPAM servers.
- **Scope of discovery**: Determine the domains that will be managed by each IPAM server. You must deploy more than one IPAM server to manage a multi-forest environment.
- **Role management**: Determine the primary role of each IPAM server. Each IPAM server can perform IP address management, role-specific server monitoring and management, and network audit for the domains you choose. Alternatively, each IPAM server can perform a subset of these functions.
- **Server placement**: Based on the scope of discovery and role management decisions you make, evaluate the best physical topology for placement of IPAM servers to optimally monitor and manage your network.
- **Training and administration**: Considering IPAM's built-in role-based security groups, conduct training and assign specific permissions and tasks to administrators in your organization.

If the IPAM design team is different from the IPAM deployment team, ensure that members of the deployment team understand the reasons the selected IPAM design is being deployed and how IPAM clients, IPAM servers, and managed servers will be affected. Ensure that members of the deployment team also understand the scope of the IPAM deployment, and what roles and responsibilities will be assigned to IPAM servers and administrators.

Reviewing IPAM concepts

For more information about how IPAM works, also see the following resources:

- What is IPAM?
- IPAM Terminology
- Getting Started with IPAM
- IPAM Architecture
- IPAM Deployment Planning
- Walkthrough: Demonstrate IPAM in Windows Server 2012

Implementing Your IPAM Design Plan

After you have decided on an IPAM design, the next step in implementing your design is to determine in what order each of the deployment tasks must be performed. This guide uses checklists to help you walk through the various server and application deployment tasks that are required to implement your design plan. As shown in the following illustration, parent and child checklists are used as necessary to represent the order in which tasks for a specific IPAM design must be performed. For IPAM, many of the tasks are similar across the three different design choices, and only differ in terms of server placement and management scope.

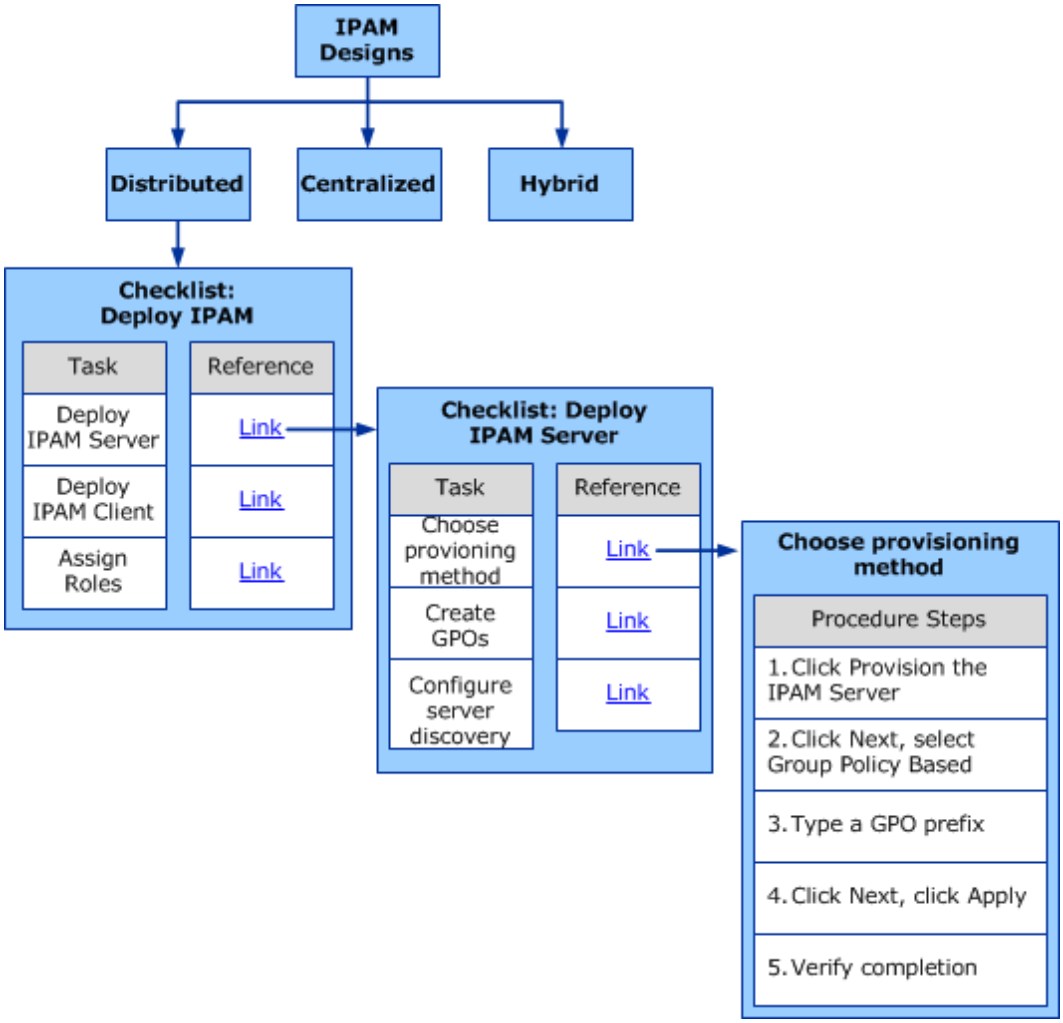

This guide contains the following deployment tasks and checklists you can use to implement your organization's IPAM design:

- Checklist: Deploy IPAM

- Checklist: Deploy IPAM Server
  - Install IPAM Server
  - Choose an IPAM Provisioning Method
  - Configure Server Discovery
  - Discover Servers on the Network
  - Manually Add a Server to Server Inventory
  - Create IPAM Provisioning GPOs
  - Manually Configure DHCP Access Settings
  - Manually Configure DNS Access Settings
  - Manually Configure DC and NPS Access Settings
  - Choose Managed Servers
  - Verify Managed Server Access
  - Retrieve Data from Managed Servers
- Checklist: Deploy IPAM Client
  - Install RSAT
  - Install IPAM Client
  - Connect to an IPAM Server
- Checklist: Assign Roles
  - Assign Administrator Roles
  - Configure IPAM Server Roles

See Also

Planning to Deploy IPAM

Checklist: Deploy IPAM

This parent checklist includes cross-reference links to important concepts about deploying an IP address management (IPAM) server. It also contains links to subordinate checklists that will help you complete the required tasks.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy IPAM

	Task	Reference
☐	Review IPAM concepts, terminology, and components; identify and prioritize goals.	What is IPAM? IPAM Terminology Getting Started with IPAM
☐	Review IPAM requirements and specifications; decide on a deployment topology; identify pilot sites; document deployment decisions and processes.	IPAM Architecture IPAM Deployment Planning
☐	Deploy IPAM Servers	Checklist: Deploy IPAM Server
☐	Deploy IPAM Clients	Checklist: Deploy IPAM Client
☐	Assign IPAM server and administrator roles	Checklist: Assign Roles
☐	Begin using IPAM	Using the IPAM Client Console Managing Server Inventory Managing IP Address Space Multi-server Management IP Address Tracking

Deploying IPAM Server

At least one IPAM server is required to deploy IPAM. In order for a computer to function as an IPAM server, it must be a domain member computer running Windows Server® 2012 or later, and the IPAM Server feature must be installed and configured. You cannot install IPAM Server on a domain controller.

Configuring an IPAM server

Before you install and configure an IPAM server, be sure that you have reviewed the following conceptual topics in the Understanding IP Address Management (IPAM) Guide:

- What is IPAM?: Information about the key features and capabilities of IPAM.
- IPAM Terminology: Terms and concepts used to refer to IPAM are defined.
- Getting Started with IPAM: IPAM deployment is summarized.

- IPAM Architecture: IPAM client, IPAM server, and managed servers are discussed.
- IPAM Deployment Planning: Deployment topologies, hardware and software requirements, specifications, and best practices.

After you have reviewed these topics, complete the tasks in Checklist: Deploy IPAM Server. Depending on your IPAM design choices, you will also be required to configure and deploy managed server access settings on IP infrastructure servers (domain controllers, NPS, DHCP, and DNS servers) to enable a functional IPAM environment.

See Also

Deploying IPAM Client

Checklist: Deploy IPAM Server

Checklist: Deploy IPAM > **Checklist: Deploy IPAM Server**

This parent checklist includes cross-reference links to important concepts about deploying an IP address management (IPAM) server. It also contains links to subordinate checklists that will help you complete the required tasks.

Before completing the tasks in this checklist, ensure you have performed the prerequisite tasks in the parent checklist, such as reviewing conceptual information about IPAM and deciding on a deployment topology. The IPAM server that you deploy using procedures in this checklist can be a central IPAM server that monitors and manages several sites, or it can be a remote server that is responsible for managing a single site or a small number of sites. The role that is assigned to an IPAM server affects procedures in this checklist, such as configuring the scope of discovery and assigning the manageability status of discovered servers.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy IPAM Server

	Task	Reference
☐	Install the IPAM Server feature	Install IPAM Server
☐	Provision the IPAM server	Choose an IPAM Provisioning Method
☐	Configure the scope of discovery	Configure Server Discovery

	Task	Reference
☐	Start server discovery	✓ Discover Servers on the Network ✓ Manually Add a Server to Server Inventory
☐	Configure settings on managed servers	✓ Create IPAM Provisioning GPOs ✓ Manually Configure DHCP Access Settings ✓ Manually Configure DNS Access Settings ✓ Manually Configure DC and NPS Access Settings
☐	Select manageability status on managed servers	✓ Choose Managed Servers
☐	Verify IPAM access to managed servers	✓ Verify Managed Server Access
☐	Retrieve data from managed servers	✓ Retrieve Data from Managed Servers

See Also

Checklist: Deploy IPAM Client

Checklist: Assign Roles

Install IPAM Server

Checklist: Deploy IPAM Server > **Install IPAM Server**

The IPAM Server feature can be installed using the Add Roles and Features wizard in Server Manager, or using Windows PowerShell.

Important
IPAM Server must be installed on a domain member computer running Windows Server® 2012 or a later operating system. The IPAM server is intended as a single purpose server and should not be installed with other network infrastructure roles such as DNS or DHCP. You cannot install IPAM on a domain controller. If IPAM Server is running on a computer

that is also running the DHCP Server role, discovery of DHCP servers on the network will be disabled.

Ensure that the IPAM server has network connectivity to at least one Active Directory domain controller, and all IP infrastructure servers that it will monitor and manage.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Installing IPAM Server

- Using the Windows interface
- Using a command line

To install IPAM Server using the Windows interface

1. On the Server Manager Dashboard page, click **Add roles and features**, click **Next** twice, and select the IPAM server from the server pool. By default, the current server is selected.
2. Click **Next** twice and then on the **Select features** page, select the **IP Address Management (IPAM) Server** checkbox.
3. In the **Add features that are required for IP Address Management (IPAM) Server** dialog box, click **Add Features**.
4. Click **Next**, and then click **Install**.
5. On the **Installation progress** page, verify that installation succeeded and then click **Close**.

To install IPAM Server using the command line

1. Open an elevated Windows PowerShell prompt, type the following command, and then press ENTER.

```
Install-WindowsFeature IPAM -IncludeManagementTools
```

2. In the output, verify that the value of **Exit Code** displayed is **Success**.

See Also

Choose an IPAM Provisioning Method

Choose an IPAM Provisioning Method

Provisioning is the process of enabling required permissions, files shares, and access settings on managed servers so that the IPAM server can communicate with them. You can choose either the manual or Group Policy Based method.

Both the manual provisioning method and the Group Policy based method require that the same settings be applied to managed servers. However, the Group Policy Based method is preferred due to its ease of use and consistency. When you choose the Group Policy based method, no settings are applied to managed servers unless you specifically assign a status of “managed” to the server. Settings are also removed more easily when you use the Group Policy based method. For more information, see Getting Started with IPAM.

Important

The choice of a provisioning method is permanent for the current installation of IPAM Server. To change the provisioning method, you must uninstall and reinstall IPAM Server.

This procedure assumes that you have already connected to the IPAM server that you will provision using the IPAM client console. To connect to a different IPAM server, click **Connect to IPAM server** on the IPAM Overview page.

If you have installed IPAM Server without IPAM Client, you must perform the procedures in Checklist: Deploy IPAM Client before you can connect to the IPAM server and choose a provisioning method. IPAM Client is installed by default when you install IPAM Server, but these features can be installed on separate computers.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Choosing a provisioning method

Use the following procedure to choose an IPAM provisioning method.

To choose an IPAM provisioning method

1. On the IPAM Overview page, click **Provision the IPAM server**. The **Provision IPAM** wizard will open.
2. Click **Next** and then on the **Select provisioning method** page, choose either the **Manual** or **Group Policy Based** radio button.
3. If the Group Policy Based method was chosen, type a GPO prefix next to **GPO name prefix**. The GPO prefix that you provide should allow you to identify GPOs that are created in the domain to be applicable to the current IPAM server. Write down the GPO prefix that you choose to reference later when creating IPAM provisioning GPOs. If the manual provisioning method is chosen, proceed to the next step.

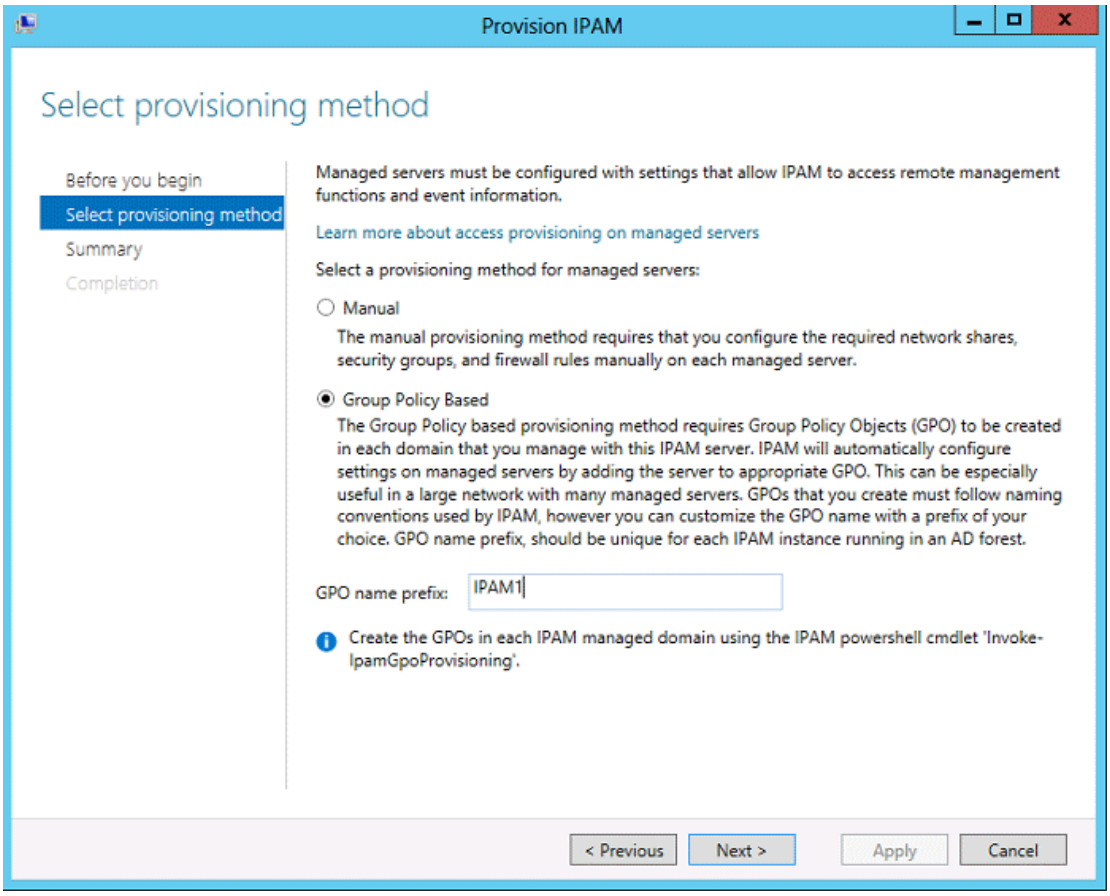

4. Click **Next**, and then on the Summary page review the settings that will be applied below **Confirm the Settings**. If the Group Policy Based method was chosen, the unique GPO names that must be created in domains managed by the current IPAM server are displayed.
5. Click **Apply**, verify that **IPAM provisioning completed successfully** is displayed, and then click **Close**.

See Also

Configure Server Discovery

Configure Server Discovery

After you complete selection of an IPAM provisioning method, you can configure server discovery. When you configure server discovery, you are defining the domains that the IPAM server can monitor and manage (also called the scope of discovery). An IPAM server can monitor and manage multiple domains as long as they are part of the same Active Directory forest as the IPAM server.

You can modify the scope of discovery at any time. For more information, see Getting Started with IPAM.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Configuring server discovery

Use the following procedure to configure server discovery.

To configure server discovery

1. On the IPAM Overview page, click **Configure server discovery**. The **Configure Server Discovery** dialog box will open.
2. Choose each domain that you will manage with the current IPAM server by selecting it from the drop-down list and then clicking **Add**.

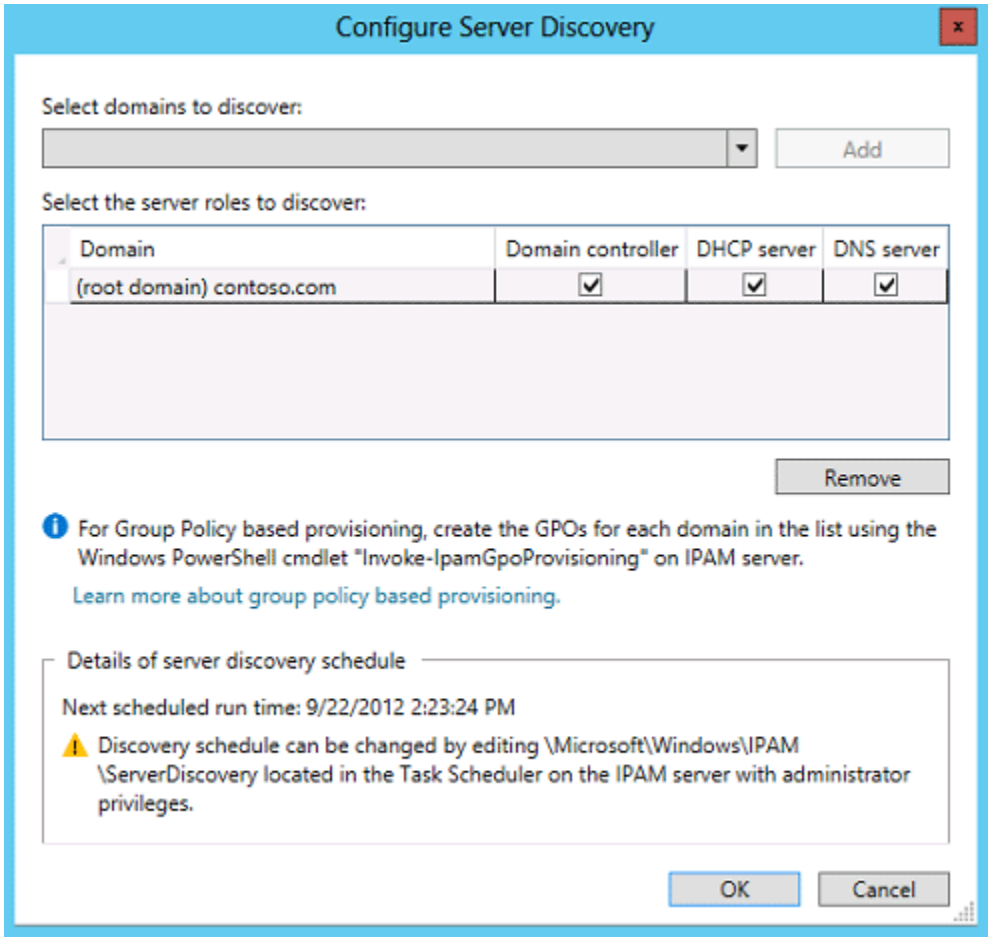

3. To remove a domain from the scope of discovery, click the domain and then click **Remove**.
4. By default all server roles are enabled in the domains you select. To remove a server role from the scope of discovery for a specific domain, de-select the checkbox under the appropriate server role.
5. Click **OK** when you are finished.

See Also

Discover Servers on the Network

Manually Add a Server to Server Inventory

Discover Servers on the Network

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > **Discover Servers on the Network**

After you have configured the scope of discovery, you can begin discovering servers on the network. When a server is discovered on the network, it is added to the IPAM database and displayed in the server inventory view of the IPAM client console.

If this is the first time you are discovering servers, or a new server has recently been added to the network, you can discover servers by running the server discovery task. This task also runs regularly on the IPAM server with a default frequency of once per day, and will discover new servers automatically provided that they meet the conditions for discovery. Conditions for discovery depend on the type of role services that are installed and running on managed servers. For example, a DHCP server will not be discovered if it is not authorized in Active Directory or does not have any DHCP scopes configured. DHCP scopes do not need to be active to be discovered. A DNS server will not be discovered unless it is authoritative for an Active Directory domain configured in the scope of discovery. For more information about conditions for discovering servers on the network, see Getting Started with IPAM. For more information about the server discovery task and other IPAM data collection tasks, see IPAM Architecture.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Discovering servers on the network

Domain controllers, DHCP servers, and DNS servers can be discovered on the network, provided they are running Windows Server® 2008 or a later operating system. Computers running NPS must be added manually to the server inventory. To manually add a server to the server inventory, perform the steps in Manually Add a Server to Server Inventory.

To discover servers on the network

1. On the IPAM Overview page, click **Start server discovery**. This will start the IPAM **ServerDiscovery** task. Alternatively, you can click **Manage** on the IPAM console menu, and then click **Start Server Discovery**.
2. Wait for the task to complete. You can click the notification flag to view status of the **ServerDiscovery** task if desired.
3. When the task has completed running, view the **Server inventory** page to display the list of discovered servers.
4. If the list of discovered servers is incomplete, verify that the correct node is selected in the lower navigation pane. By default, IPv4 is selected. You can click **Refresh** to ensure the view is current.

See Also

Manually Add a Server to Server Inventory

Manually Add a Server to Server Inventory

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | **Manually Add a Server to Server Inventory**

After you have configured the scope of discovery, you can add servers to the server inventory. Domain controllers, DNS servers, and DHCP servers are automatically discovered by IPAM if they are running Windows Server® 2008 or a later operating system. Network Policy Servers (NPS) must be manually added to the server inventory.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Manually adding a server

When you manually add a managed server to the server inventory, IPAM checks that the server name exists in Active Directory, and that the name or IP address does not already exist in the IPAM database.

To manually add a server

1. Click **SERVER INVENTORY** in the upper IPAM navigation tree, click **TASKS**, and then click **Add Server**.
2. In the **Add or Edit Server** dialog box, next to **Server name (FQDN)** type the name or IP address of the managed server and then click **Verify**.

   Note
   The IP address of the server is automatically added when you click **Verify**, and the name will be resolved to a fully qualified domain name when you close the **Add or Edit Server** dialog box. The DNS server that is configured on the IPAM server's network interface must be able to resolve the name or IP address that you provide.
3. Next to **Server type**, select at least one role for the managed server.
4. Optionally, type information next to **Owner** and **Description**, and provide a custom configuration.
5. Click **OK** to finish.

See Also

Choose Managed Servers

Create IPAM Provisioning GPOs

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | Manually Add a Server to Server Inventory > **Create IPAM Provisioning GPOs**

If you have chosen the Group Policy based provisioning method, you must also provide a GPO name prefix in the provisioning wizard. After providing a GPO name prefix, the wizard will display the GPO names that must be created in domains that will be managed by IPAM. The following GPOs are not created by the provisioning wizard and must be created in each domain that will be managed by the IPAM server:

<GPO-prefix>_DHCP: For managed DHCP servers.

<GPO-prefix>_DNS: For managed DNS servers.

<GPO-prefix>_DC_NPS: For managed domain controllers and NPS servers.

You must create GPOs with these names in order for them to be automatically applied by IPAM when a server in the server inventory is marked as managed. GPOs are also removed automatically when a server is marked as unmanaged. Addition and removal of these GPOs is accomplished by modifying security filtering for the GPO. Server names are added when they are marked as managed, or removed if they are marked as unmanaged.

If you have chosen the manual provisioning method, you can also create IPAM provisioning GPOs and use them to apply access settings to managed servers. In this case, specific GPO names are not required. However, you must apply GPOs manually to managed servers by adding or removing the managed server to GPO security filtering as needed. For steps to modify GPO security filtering, see Configuring IPAM GPO security filtering. This procedure is not necessary if you have chosen the automatic GPO-based provisioning method.

For more information about choosing a provisioning method, see Choosing a provisioning method.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Creating IPAM provisioning GPOs

In the following procedure, the **Invoke-IpamGpoProvisioning** Windows PowerShell cmdlet is used to create IPAM provisioning GPOs.

To create IPAM provisioning GPOs

1. Run the Invoke-IpamGpoProvisioning cmdlet at an elevated Windows PowerShell prompt. Invoke-IpamGpoProvisioning has the following parameters:

```
Invoke-IpamGpoProvisioning [-Domain] <String> [-
GpoPrefixName] <String> [-DelegatedGpoGroup <String[]> ] [-
DelegatedGpoUser <String[]> ] [-DomainController <String> ]
[-Force] [-IpamServerFqdn <String> ] [-PassThru] [
<CommonParameters>]
```

2. When prompted in the command output, press ENTER to confirm creation of the GPOs. You can also add the –Force parameter to bypass this confirmation. A full list of available parameters is provided in the following table.
3. Confirm that no errors are displayed.

Value	Description
Invoke-IpamGpoProvisioning	The Windows PowerShell cmdlet for creating and linking group policies in the specified domain for provisioning required access settings on the servers managed by the specified IPAM server.
-Domain	Required. Specifies the FQDN of the domain for which the IPAM GPOs are created and linked. Only one domain name should be specified. This cmdlet needs to be invoked for each domain in the Active Directory (AD) forest that needs to be managed by the IPAM server.
-GpoPrefixName	Required. Specifies the unique GPO prefix name used while creating the Group Policy Objects. The GPO prefix name specified here must be same as the one specified in the IPAM provisioning wizard if the Group Policy Based provisioning method was chosen. Three IPAM GPOs are created with the suffixes _DHCP, _DNS, and _DC_NPS appended to this parameter value. These suffixes signify the three different types of access settings that are applied depending on the type of server roles selected on the managed server.
-DelegatedGpoGroup	Optional. Specifies the comma separated list of security groups which are delegated edit permissions on the three IPAM GPOs created by this cmdlet. The standard GPO delegation

Value	Description
	privilege of **Edit settings**, **Delete**, and **Modify Security** is enabled for these groups. Delegation for users or groups without domain administrator privileges can also be enabled from GPMC or GPO cmdlets for the existing IPAM GPOs. This delegation is leveraged by the IPAM server to permit editing of the manageability status in a security context of the logged in user who is modifying manageability status from the IPAM server inventory view.
-DelegatedGpoUser	Optional. Specifies the comma separated list of domain users who are delegated edit permissions on the three IPAM GPOs created by this cmdlet. The standard GPO delegation privilege of **Edit settings**, **Delete**, and **Modify Security** is enabled for these groups. Delegation for users or groups without domain administrator privileges can also be enabled from GPMC or GPO cmdlets for the existing IPAM GPOs. This delegation is leveraged by the IPAM server to permit editing of the manageability status in a security context of the logged in user who is modifying manageability status from the IPAM server inventory view.
-DomainController	Optional. Specifies the name of the domain controller that is contacted to complete the GPO operation. Either the FQDN or simply the host name of the domain controller can be specified. By default the computer running the IPAM server contacts the PDC emulator of the specified domain unless this parameter is explicitly provided.
-Force	Optional. Forces the command to run without asking for user confirmation.
-IpamServerFqdn	Optional. Specifies the fully qualified domain name (FQDN) of the IPAM server for which the GPOs enable access settings. By default IPAM uses the FQDN of the IPAM server that runs the cmdlet to create GPOs. The specified computer account is added to the universal

Value	Description
	group named **IPAMUG** created by this cmdlet.
-PassThru	Optional. Returns an object representing the item with which you are working. By default, this cmdlet does not generate any output.
CommonParameters	Optional. This cmdlet supports the common parameters: -Verbose, -Debug, -ErrorAction, -ErrorVariable, -OutBuffer, and -OutVariable.

Sample Usage

In the following example, three GPOs are created (IPAM1_DHCP, IPAM1_DNS, IPAM1_DC_NPS) and linked to the contoso.com domain. These GPOs enable access for the server ipam1.contoso.com using the domain administrator account user1. Note: In this example, the hostname of the IPAM server is used as a GPO prefix, however this is not required.

```
Invoke-IpamGpoProvisioning -Domain contoso.com -GpoPrefixName IPAM1 -DelegatedGpoUser user1 -IpamServerFqdn ipam1.contoso.com
```

Configuring IPAM GPO security filtering

It is not necessary to configure IPAM GPO security filtering if you are using the automatic Group Policy based provisioning method. If you are using the manual provisioning method with IPAM GPOs, you must add and remove individual servers from the appropriate GPOs by editing GPO security filters.

To configure IPAM GPO security filtering

1. On a domain controller, IPAM server, or other domain member server with the Group Policy Management feature installed, type **gpmc.msc** at an elevated command prompt and press ENTER.
2. In in the Group Policy Management console tree, navigate to Forest\Domains\<domain>\Group Policy Objects\
3. Click the GPO to configure. For example, if you are adding a managed DHCP server, click the GPO name ending in _DHCP.
4. On the **Scope** tab, under **Security Filtering**, click **Add** to add a new server to this GPO.

   If you are removing a managed server from security filtering, click the name of the managed server, click **Remove**, and then click **OK** when you are prompted to confirm removal of this delegation privilege. You can close Group Policy Management if you are finished configuring security filters.
5. To continue adding a managed server to security filtering, in **Select User, Computer, or Group** click **Object Types**, select the **Computers** checkbox, and then click **OK**.

6. Under **Enter the object name to select**, type the name of the managed server and click **Check Names**.
7. Verify that the server name is underlined, and then click **OK**. The managed server name will be displayed under **Security Filtering**.
8. Close **Group Policy Management**.

See Also

Manually Configure DHCP Access Settings

Manually Configure DNS Access Settings

Manually Configure DC and NPS Access Settings

Manually Configuring Managed Server Access Settings

If you have chosen the manual provisioning method in the Provision IPAM wizard, you must apply managed server access settings manually. You can apply access settings by manually applying IPAM provisioning GPOs, or by configuring access settings individually on each managed server.

Manually configuring managed server access settings

Before you attempt to configure IPAM access settings manually on managed servers, review the following conceptual topics:

- Getting Started with IPAM
- IPAM Architecture
- IPAM Deployment Planning

After you have reviewed these topics, complete the following procedures to configure managed server access settings manually. Choose the procedure that corresponds to the server roles that will be managed by the IPAM server.

Manually Configure DHCP Access Settings

Manually Configure DNS Access Settings

Manually Configure DC and NPS Access Settings

See Also

Create IPAM Provisioning GPOs

Manually Configure DHCP Access Settings

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | Manually Add a Server to Server Inventory > Create IPAM Provisioning GPOs | **Manually Configure DHCP Access Settings**

If you have chosen the manual provisioning method, you must configure access permissions and settings individually on managed servers. Two general methods for configuring these settings are available:

Manual application of IPAM GPOs: With this method, IPAM GPOs are created but are not automatically applied to managed servers. Application of GPOs is done manually by adding managed servers to GPO security filtering. To use this method, first perform the steps in Create IPAM Provisioning GPOs. When you have created these GPOS, configure security filtering by performing the steps in Configuring IPAM GPO security filtering.

Manual configuration of individual settings: GPOs are not used with this method. Instead, access settings on each managed server are configured manually by an administrator. Steps to configure these access settings are provided in this topic.

For more information about choosing an IPAM provisioning method, see Choosing a provisioning method.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Manually configure managed DHCP server access settings

Use the following procedures to configure IPAM access settings on a managed DHCP server:

Configure Windows Firewall on a managed DHCP server

Configure security groups on a managed DHCP server

Configure a DHCP audit share on a managed DHCP server

Restart the DHCP Server service

A summary of required settings is provided in the table below.

Setting Type	Setting Name	Function
Inbound firewall rule	DHCP Server Management	Access configuration data
Inbound firewall rule	Remote Service Management	Access configuration data
Inbound firewall rule	File and Printer Sharing	Access utilization data
Inbound firewall rule	Remote Event Log	Access server logs

Setting Type	Setting Name	Function
	Management	
Security group	DHCP Users	Access configuration data
Security group	Event Log Readers	Access server event logs
Network share	DHCP Audit	Access utilization data

Configure Windows Firewall on a managed DHCP server

Use the following procedure to enable access by the IPAM server to inbound firewall ports on a managed DHCP server. If the managed DHCP server becomes unmanaged, delete these inbound firewall rules. All of the following procedures are performed on the managed DHCP server.

To configure Windows Firewall on a managed DHCP server

1. On the Server Manager menu, click **Tools** and then click **Windows Firewall with Advanced Security**.
2. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
3. In **Rule Type**, select **Predefined**, choose **DHCP Server Management** from the list, and then click **Next**.
4. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

   DHCP Server (RPCSS-In)

   DHCP Server (RPC-In)
5. Click **Next**, choose **Allow the connection**, and then click **Finish**.
6. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
7. In **Rule Type**, select **Predefined**, choose **File and Printer Sharing** from the list, and then click **Next**.
8. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

   File and Printer Sharing (NB-Session-In)

   File and Printer Sharing (SMB-In)
9. Click **Next**, choose **Allow the connection**, and then click **Finish**.
10. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
11. In **Rule Type**, select **Predefined**, choose **Remote Event Log Management** from the list, and then click **Next**.

12. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

    Remote Event Log Management (RPC)

    Remote Event Log Management (RPC-EPMAP)
13. Click **Next**, choose **Allow the connection**, and then click **Finish**.
14. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
15. In **Rule Type**, select **Predefined**, choose **Remote Service Management** from the list, and then click **Next**.
16. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

    Remote Service Management (RPC)

    Remote Service Management (RPC-EPMAP)
17. Click **Next**, choose **Allow the connection**, and then click **Finish**.
18. Close the Windows Firewall with Advanced Security console.

Configure security groups on a managed DHCP server

To access configuration data and server event logs, the IPAM server must be a member of the domain IPAM Users Group (IPAMUG). The IPAM server must also be a member of the local DHCP Users and Event Log Readers security groups.

Important
It is necessary to restart the DHCP service after changing security group memberships in order for these changes to become active.

After performing these procedures, refresh the server access status for the managed DHCP server in the server inventory view on the IPAM server.

To configure the IPAMUG security group

1. On a domain controller, click **Tools** on the Server Manager menu, and then click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console tree, right-click the **Users** container under the managed DHCP server's domain, point to **New**, and then click **Group**.
3. In the **New Object – Group** dialog box, under **Group name**, type **IPAMUG**.
4. Under **Group scope**, select **Universal** and under **Group type** select **Security**, and then click **OK**.

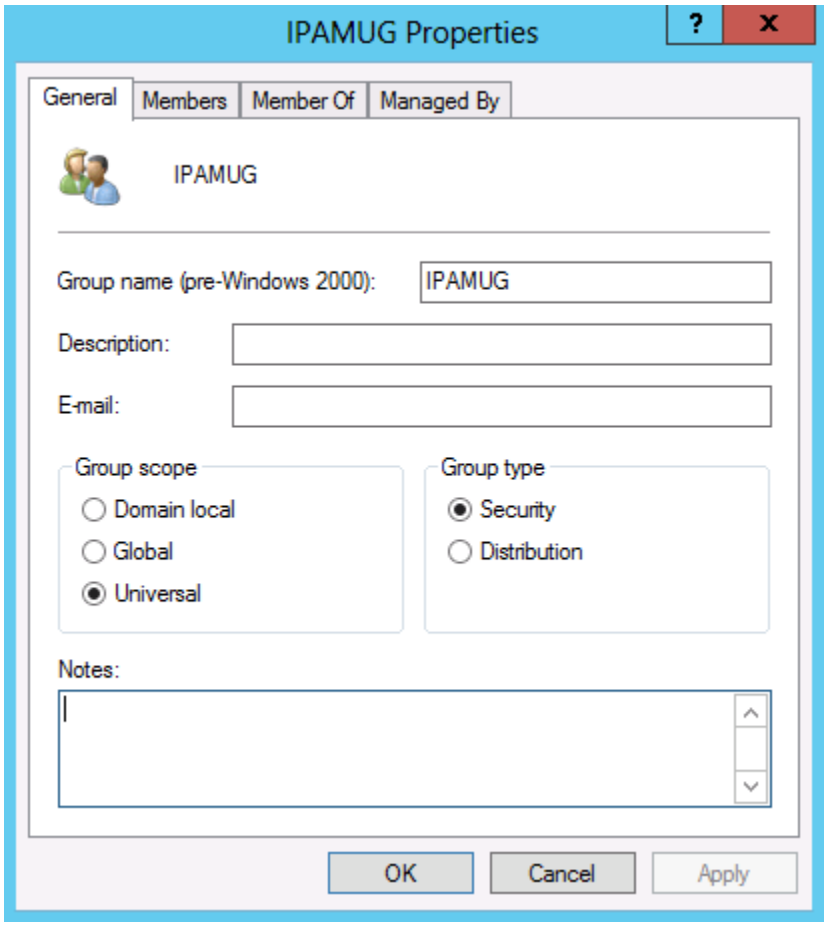

5. Double-click the IPAMUG group and then click the **Members** tab.
6. Click **Object Types**, select the **Computers** checkbox, and then click **OK**.
7. Under **Enter the object names to select**, type the name of the IPAM server and then click **OK**.
8. Verify that the correct IPAM server is displayed under **Name** and then click **OK**.

To configure the DHCP Users and Event Log Readers security groups

1. On the managed DHCP server, click **Tools** on the Server Manager menu and then click **Computer Management**.
2. In the Computer Management navigation tree, click **Local Users and Groups** and then click **Groups**.
3. In the list of groups, double-click **DHCP Users** and then click **Add**.
4. Under **Enter the object names to select**, type **IPAMUG** and then click **OK**.
5. Verify that **<domain>\IPAMUG** is displayed under **Members** and then click **OK**.

6. In the list of groups, double-click **Event Log Readers** and then click **Add**.
7. Under **Enter the object names to select**, type **IPAMUG** and then click **OK**.
8. Verify that **<domain>\IPAMUG** is displayed under **Members** and then click **OK**.
9. Leave the Computer Management console open for the following procedure.

Configure a DHCP audit share on a managed DHCP server

To audit IP address utilization, IPAM requires access to the DHCP database file located in the **windows\system32\dhcp** directory. The following procedure provides access to the DHCP database by creating a network share.

To configure a DHCP audit share on a managed DHCP server

1. In the Computer Management navigation tree, click **Shared Folders** and then click **Shares**.
2. Right-click **Shares** and then click **New Share**.
3. In the **Create A Shared Folder Wizard**, click **Next**.
4. Next to Folder Path, type the absolute path to the windows\system32\dhcp directory, for example **C:\windows\system32\dhcp**, and then click **Next**.
5. Next to **Share name**, type **dhcpaudit** and then click **Next**.
6. Under **Set the kind of permissions you want for the shared folder**, choose **Customize permissions** and then click **Custom**.
7. In the **Customize Permissions** dialog box, click **Everyone**, click **Remove**, and then click **Add**.
8. Under **Enter the object names to select**, type **IPAMUG** and then click **OK**.
9. Verify that **IPAMUG** is displayed under **Group or user names**, and that **Read** is allowed under **Permissions for IPAMUG**, click **OK**, and then click **Finish** twice.

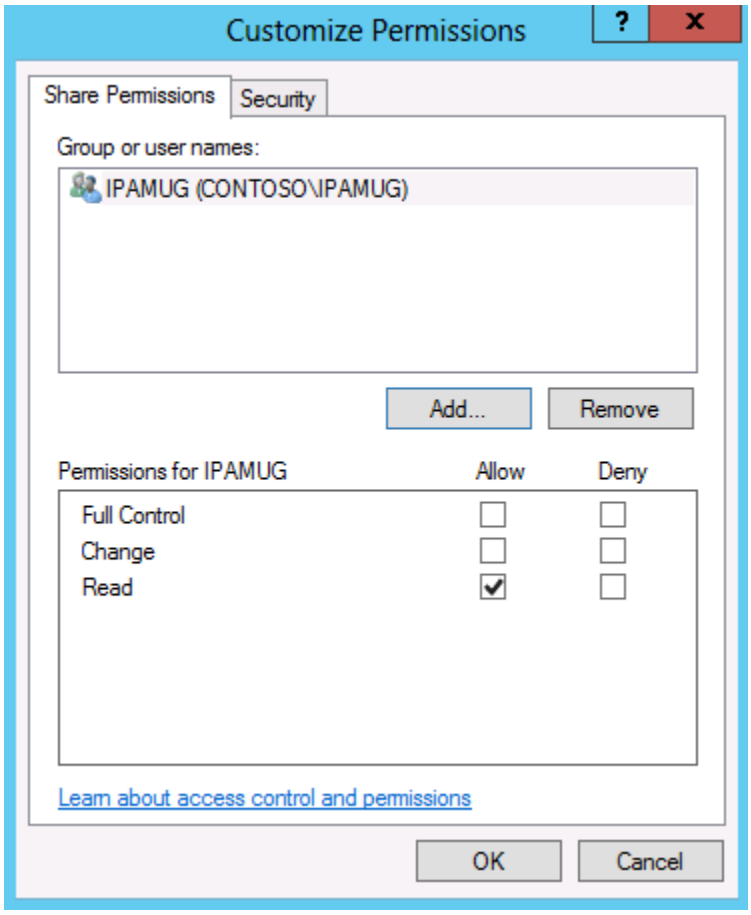

Restart the DHCP Server service

To activate new security group memberships on the DHCP server, you must restart the DHCP Server service.

To restart the DHCP Server service

1. In the Server Manager navigation pane, click **DHCP**.
2. Under **SERVICES**, right-click the name of the managed DHCP server, and then click **Restart Services**.
3. To monitor status of the DHCP service, click the Server Manager notification flag. Verify that the service stops and starts successfully.

See Also

Manually Configure DNS Access Settings

Manually Configure DC and NPS Access Settings

Manually Configure DNS Access Settings

If you have chosen the manual provisioning method, you must configure access permissions and settings individually on managed servers. Two general methods for configuring these settings are available:

Manual application of IPAM GPOs: With this method, IPAM GPOs are created but are not automatically applied to managed servers. Application of GPOs is done manually by adding managed servers to GPO security filtering. To use this method, first perform the steps in Create IPAM Provisioning GPOs. When you have created these GPOS, configure security filtering by performing the steps in Configuring IPAM GPO security filtering.

Manual configuration of individual settings: GPOs are not used with this method. Instead, access settings on each managed server are configured manually by an administrator. Steps to configure these access settings are provided in this topic.

For more information about choosing an IPAM provisioning method, see Choosing a provisioning method.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Manually configure managed DNS server access settings

Use the following procedures to configure IPAM access settings on a managed DHCP server:

Configure Windows Firewall on a managed DNS server

Configure security groups on a managed DNS server

Enable event log monitoring on a managed DNS server

Configure the DNS DACL on a managed DNS server

A summary of required settings is provided in the table below.

Setting Type	Setting Name	Function
Inbound firewall rule	DNS Service	Discover DNS servers
Inbound firewall rule	Remote Service Management	Manage DNS servers
Inbound firewall rule	Remote Event Log Management	Monitor DNS zones and services

Setting Type	Setting Name	Function
Security Group	Event Log Readers	Monitor DNS servers
Event log access	Event log monitoring	Monitor DNS zones and services
DNS administration	DNS DACL	Administer DNS zones and servers

After performing these procedures, refresh the server access status for the managed DNS server in the server inventory view on the IPAM server.

Configure Windows Firewall on a managed DNS server

Use the following procedure to enable access by the IPAM server to inbound firewall ports on a managed DNS server. If the managed DNS server becomes unmanaged, delete these inbound firewall rules. All of the following procedures are performed on the managed DNS server.

To configure Windows Firewall on a managed DNS server

1. On the Server Manager menu, click **Tools** and then click **Windows Firewall with Advanced Security**.
2. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
3. In **Rule Type**, select **Predefined**, choose **DNS Service** from the list, and then click **Next**.
4. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

   RPC (TCP, Incoming)

   DNS (UDP, Incoming)

   DNS (TCP, Incoming)

   RPC Endpoint Mapper (TCP, Incoming)
5. Click **Next**, choose **Allow the connection**, and then click **Finish**.
6. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
7. In **Rule Type**, select **Predefined**, choose **Remote Service Management** from the list, and then click **Next**.
8. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

   Remote Service Management (RPC-EPMAP)

   Remote Service Management (NP-In)

   Remote Service Management (RPC)
9. Click **Next**, choose **Allow the connection**, and then click **Finish**.

10. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
11. In **Rule Type**, select **Predefined**, choose **Remote Event Log Management** from the list, and then click **Next**.
12. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

    Remote Event Log Management (RPC-EPMAP)

    Remote Event Log Management (RPC)
13. Click **Next**, choose **Allow the connection**, and then click **Finish**.
14. Close the Windows Firewall with Advanced Security console.

Configure security groups on a managed DNS server

To access configuration data and server event logs, the IPAM server must be a member of the domain IPAM Users Group (IPAMUG). The IPAM server must also be a member of the Event Log Readers security group.

Warning

If the DNS Server service is running on a domain controller, some procedures below are different than if the DNS server is a domain member computer that is not running on the same server with Active Directory Domain Services (AD DS). If procedures are unique to a type of DNS server, this is specified in the procedure title.

To configure the IPAMUG security group

1. On a domain controller, click **Tools** on the Server Manager menu, and then click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console tree, right-click the **Users** container under the managed DHCP server's domain, point to **New**, and then click **Group**.
3. In the **New Object – Group** dialog box, under **Group name**, type **IPAMUG**.
4. Under **Group scope**, select **Universal** and under **Group type** select **Security**, and then click **OK**.

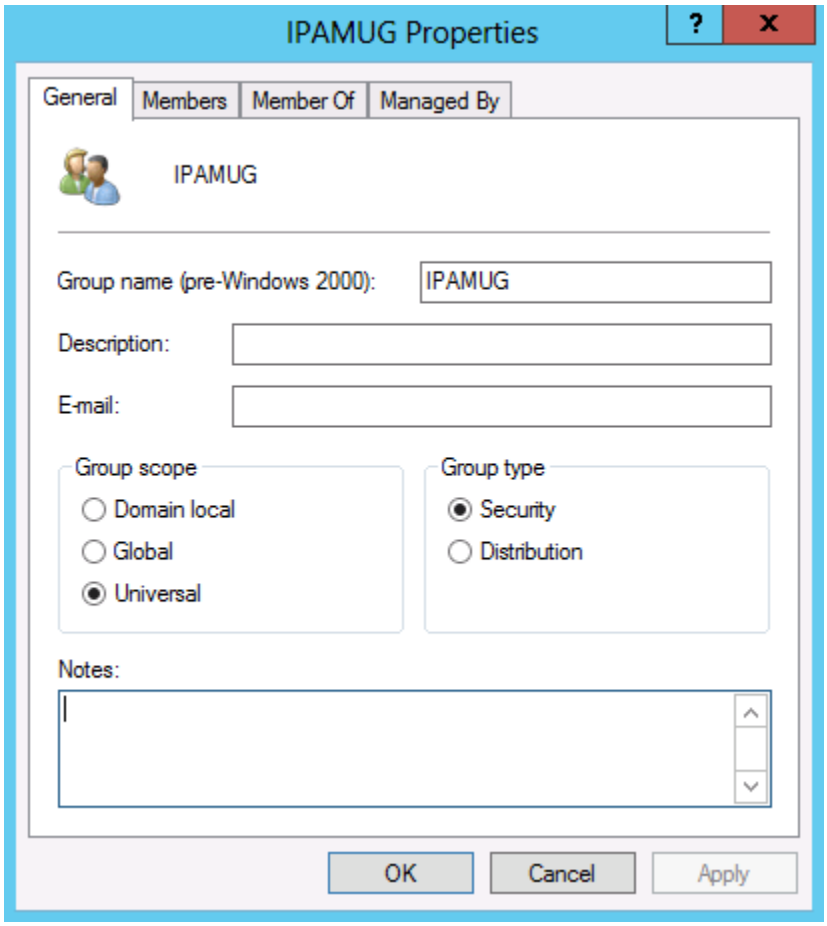

5. Double-click the IPAMUG group and then click the **Members** tab.
6. Click **Object Types**, select the **Computers** checkbox, and then click **OK**.
7. Under **Enter the object names to select**, type the name of the IPAM server and then click **OK**.
8. Verify that the correct IPAM server is displayed under **Name** and then click **OK**.

To configure the Event Log Readers security group on a DC/DNS server

1. On the managed DC/DNS server, click **Tools** on the Server Manager menu and then click **Active Directory Users and Computers**.
2. In the navigation tree, click the **<domain>\Builtin** container and then double-click **Event Log Readers**.
3. On the **Members** tab, click **Add**.
4. Under **Enter the object names to select**, type **IPAMUG**, and then click **OK** twice.

To configure the Event Log Readers security group on a standalone DNS server

1. On the managed DNS server, click **Tools** on the Server Manager menu and then click **Computer Management**.
2. In the Computer Management navigation tree, click **Local Users and Groups** and then click **Groups**.
3. In the list of groups, double-click **Event Log Readers** and then click **Add**.
4. Under **Enter the object names to select**, type **IPAMUG** and then click **OK**.
5. Verify that **<domain>\IPAMUG** is displayed under **Members** and then click **OK**.

Enable event log monitoring on a managed DNS server

Event log monitoring is enabled through the **HKLM\SYSTEM\CurrentControlSet\Services\EventLog\DNS Server** registry key.

To enable event log monitoring on a managed DNS server

1. On the managed DNS server, type **regedit.exe** at an elevated command prompt, and then press ENTER.
2. In Registry Editor, navigate to the **HKLM\SYSTEM\CurrentControlSet\Services\EventLog\DNS Server** registry key.
3. Double-click **CustomSD** in the right-hand pane and then under **Value data**, position the mouse cursor at the end of the string.

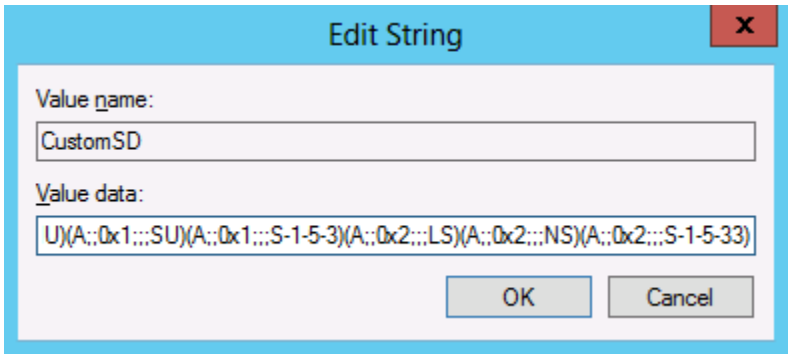

4. Open a Windows PowerShell prompt and type **Get-ADComputer <IPAM server name>** and press ENTER. Replace IPAM server name with the host name of your IPAM server.

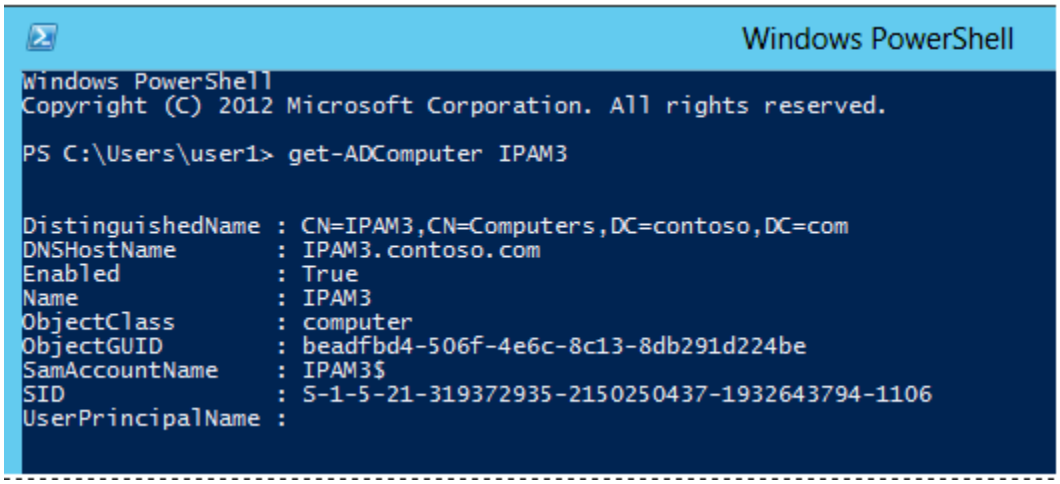

Tip

The get-ADComputer cmdlet is not available unless the Active Directory module for Windows PowerShell is installed. This module is installed by default on a domain controller. If the DNS server is a standalone domain member server, you can install the Active Directory module with the Add Roles and Features Wizard by selecting a Role-based or feature-based installation and then selecting **Features\Remote Server Administration Tools\Role Administration Tools\AD DS and AD LDS Tools\Active Directory module for Windows PowerShell**. Alternatively, you can acquire the SID of the IPAM server on a different computer and copy it to the standalone DNS server.

5. In the command output, highlight the string next to **SID** and press ENTER. This will copy the string to the clipboard.
6. In Registry Editor, at the end of the CustomSD Value data string, type **(A;;0x1;;;** then paste the SID from the clipboard, and then type a closed parens. For example, if the SID that was copied is S-1-5-21-319372935-2150250437-1932643794-1106 (as shown in the example above), then the end of the Value data string is **(A;;0x1;;; S-1-5-21-319372935-2150250437-1932643794-1106)**. Be careful not to replace the entire value data string. Only append the extra characters to the end of the string.
7. Click **OK** and then close Registry Editor.

Configure the DNS DACL on a managed DNS server

The IPAM server must be added to the **DnsAdmins** security group. This enables the IPAM server to perform DNS administration functions. You can use DNS Manager to configure this security setting on a managed DNS server if it is also a domain controller.

Important
To enable DNS administration on a standalone domain member DNS server, it is required to add the computer account of the IPAM server to the local Administrators group on the DNS server. This procedure is provided below, but is not recommended due to security concerns.

After performing these procedures, refresh the server access status for the managed DNS server in the server inventory view on the IPAM server.

To enable DNS administration on a managed DC/DNS server

1. On the managed DNS server, open an elevated command prompt, type **dnsmgmt.msc**, and press ENTER.
2. In DNS Manager, right-click the name of the DNS server and then click **Properties**.
3. Click the **Security** tab, click **Add**, type **IPAMUG**, and then click **OK** twice.

To enable DNS administration on a standalone DNS server

1. On a standalone DNS server, open an elevated command prompt, type **lusrmgr.msc**, and then press ENTER.
2. In the **Local Users and Groups** navigation tree, click **Groups** and then double-click **Administrators** in the right-hand pane.
3. In **Administrators Properties**, click **Add**.
4. Click **Object Types**, select the **Computers** checkbox, and then click **OK**.
5. Under **Enter the object names to select**, type the name of the IPAM server and then click **OK**.
6. Verify that the IPAM server is displayed under **Members**, and then click **OK**.

See Also

Manually Configure DHCP Access Settings

Manually Configure DC and NPS Access Settings

Manually Configure DC and NPS Access Settings

If you have chosen the manual provisioning method, you must configure access permissions and settings individually on managed servers. Two general methods for configuring these settings are available:

Manual application of IPAM GPOs: With this method, IPAM GPOs are created but are not automatically applied to managed servers. Application of GPOs is done manually by adding managed servers to GPO security filtering. To use this method, first perform the steps in Create IPAM Provisioning GPOs. When you have created these GPOS, configure security filtering by performing the steps in Configuring IPAM GPO security filtering.

Manual configuration of individual settings: GPOs are not used with this method. Instead, access settings on each managed server are configured manually by an administrator. Steps to configure these access settings are provided in this topic.

For more information about choosing an IPAM provisioning method, see Choosing a provisioning method.

Membership in the **Domain Admins** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Manually configure managed DC and NPS server access settings

Use the following procedures to configure IPAM access settings on a managed domain controller or NPS server. If the server is also a managed DNS or DHCP server, and you have previously enabled IPAM access for these services, you have already enabled the necessary access. If you have not already enabled IPAM access using Windows Firewall and the Event Log Readers security group, perform the following procedures:

Configure Windows Firewall on a managed DC or NPS server

Configure security groups on a managed DC or NPS server

A summary of required settings is provided in the table below.

Setting Type	Setting Name	Function
Inbound firewall rule	Remote Event Log Management	Audit authentication events
Security group	Event log monitoring	Audit authentication events

Configure Windows Firewall on a managed DC or NPS server

Use the following procedure to enable access by the IPAM server to inbound firewall ports on a managed domain controller or NPS server. If the managed server becomes unmanaged, delete

these inbound firewall rules. All of the following procedures are performed on the managed server.

To configure Windows Firewall on a managed DC or NPS server

1. On the Server Manager menu, click **Tools** and then click **Windows Firewall with Advanced Security**.
2. Right-click **Inbound Rules**, and then click **New Rule**. The **New Inbound Rule Wizard** will launch.
3. In **Rule Type**, select **Predefined**, choose **Remote Event Log Management** from the list, and then click **Next**.
4. In **Predefined Rules**, under **Rules**, select the checkboxes next to the following rules:

   Remote Event Log Management (RPC-EPMAP)

   Remote Event Log Management (RPC)
5. Click **Next**, choose **Allow the connection**, and then click **Finish**.
6. Close the Windows Firewall with Advanced Security console.

Configure security groups on a managed DC or NPS server

To access configuration data and server event logs, the IPAM server must be a member of the domain IPAM Users Group (IPAMUG). The IPAM server must also be a member of the Event Log Readers security group.

Warning

If you have already created the IPAMUG security group in Active Directory and added the IPAM server as a member of the group, it is not necessary to perform this procedure again.

The steps to add the IPAM server to the Event Log Readers security group is different on a managed domain controller than a managed NPS server. Both procedures are provided below. If the server is both a domain controller and NPS server, perform the procedures for a managed domain controller.

After performing these procedures, refresh the server access status for the managed domain controller or NPS server in the server inventory view on the IPAM server.

To configure the IPAMUG security group

1. On a domain controller, click **Tools** on the Server Manager menu, and then click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console tree, right-click the **Users** container under the managed DHCP server's domain, point to **New**, and then click **Group**.
3. In the **New Object – Group** dialog box, under **Group name**, type **IPAMUG**.

4. Under **Group scope**, select **Universal** and under **Group type** select **Security**, and then click **OK**.

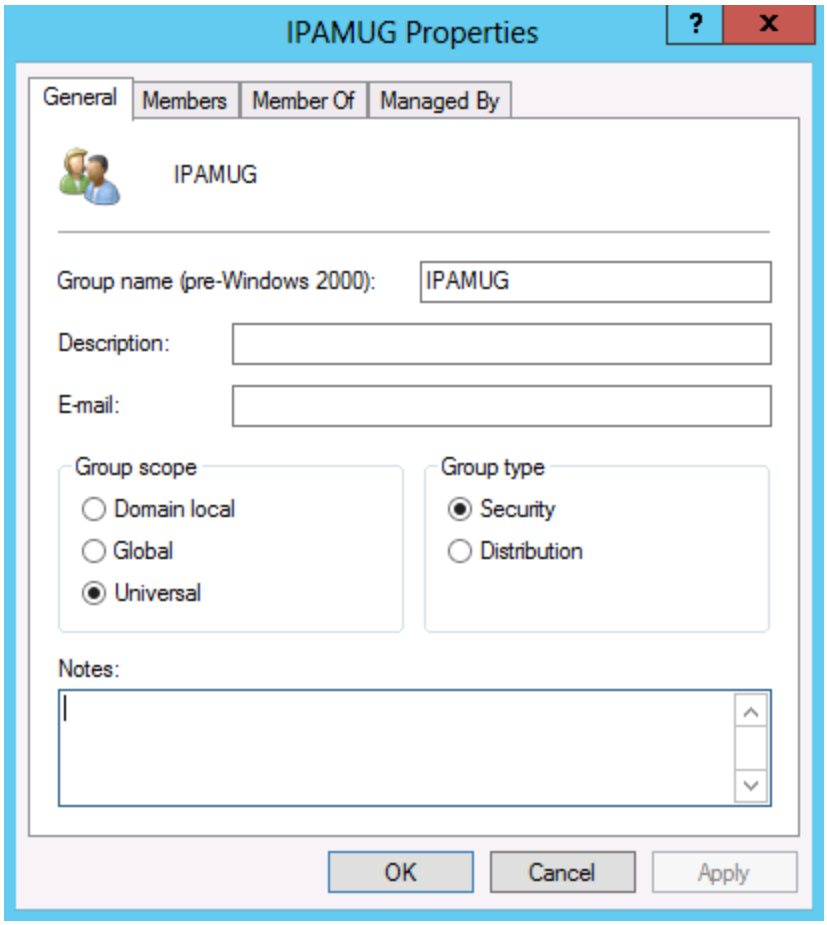

5. Double-click the IPAMUG group and then click the **Members** tab.
6. Click **Object Types**, select the **Computers** checkbox, and then click **OK**.
7. Under **Enter the object names to select**, type the name of the IPAM server and then click **OK**.
8. Verify that the correct IPAM server is displayed under **Name** and then click **OK**.

To configure the Event Log Readers security group on a domain controller

1. On the managed domain controller, click **Tools** on the Server Manager menu and then click **Active Directory Users and Computers**.
2. In the navigation tree, click the **<domain>\Builtin** container and then double-click **Event Log Readers**.
3. On the **Members** tab, click **Add**.

4. Under **Enter the object names to select**, type **IPAMUG**, and then click **OK** twice.

To configure the Event Log Readers security group on an NPS server

1. On the managed NPS server, click **Tools** on the Server Manager menu and then click **Computer Management**.
2. In the Computer Management navigation tree, click **Local Users and Groups** and then click **Groups**.
3. In the list of groups, double-click **Event Log Readers** and then click **Add**.
4. Under **Enter the object names to select**, type **IPAMUG** and then click **OK**.
5. Verify that **<domain>\IPAMUG** is displayed under **Members** and then click **OK**.

See Also

Manually Configure DHCP Access Settings

Manually Configure DNS Access Settings

Choose Managed Servers

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | Manually Add a Server to Server Inventory > Create IPAM Provisioning GPOs | Manually Configure DHCP Access Settings | Manually Configure DNS Access Settings | Manually Configure DC and NPS Access Settings > **Choose Managed Servers**

After you have discovered servers on the network and configured IPAM access settings, you can choose managed servers. Choosing a managed server is done by assign manageability status. If the automatic, GPO-based provisioning method is used, a manageability status of **Managed** means that the server will automatically be added to security filtering in the appropriate IPAM GPO. If the manual provisioning method is chosen, the manageability status of servers must still be selected to enable IPAM data collection and management functions.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Choosing managed servers

The **SERVER INVENTORY** view is used to choose managed servers and to specify managed role services. All discovered servers are displayed in this view, together with the current manageability status and IPAM access status. By default, manageability status of servers is **Unspecified**. For more information about manageability settings and IPAM access status, see Getting Started with IPAM.

To choose managed servers

1. In the upper IPAM navigation tree, click **SERVER INVENTORY**. The list of servers that have been discovered or manually added is displayed.
2. Select one or more servers by clicking the appropriate row in the list. To select multiple servers, hold down SHIFT or CTRL when clicking.
3. Right-click the server or servers that are selected and then click **Edit Server**.
4. In the **Add or Edit Server** dialog box next to **Manageability status**, select **Managed** from the drop-down list if the server will be managed by the current IPAM server. Select **Unmanaged** if the server will not be managed by the current IPAM server, and then click **OK**.

To choose managed services

1. From the **SERVER INVENTORY** view, right-click a single server in the list and then click **Edit Server**.
2. In the **Add or Edit Server** dialog box next to **Server type**, select one or more checkboxes next to the role services that will be managed by the current IPAM server, and then click **OK**.

See Also

Verify Managed Server Access

Retrieve Data from Managed Servers

Verify Managed Server Access

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | Manually Add a Server to Server Inventory > Create IPAM Provisioning GPOs | Manually Configure DHCP Access Settings | Manually Configure DNS Access Settings | Manually Configure DC and NPS Access Settings > Choose Managed Servers > **Verify Managed Server Access**

After you have selected servers that will be managed by IPAM, you can verify IPAM access settings. If the automatic, GPO-based provisioning method is used, selecting a manageability status of **Managed** will automatically add a managed server to security filtering in the appropriate role-based IPAM GPOs. If the manual provisioning method is chosen, you can manually apply IPAM provisioning GPOs, or manually configure managed server access settings. In some cases, application of IPAM access settings requires that role services are restarted. For more information about applying and verifying IPAM access settings, see Getting Started with IPAM.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Verifying managed server access

The **SERVER INVENTORY** view is used to display the IPAM access status for managed and unmanaged servers. Servers will display a status of **Blocked**, **Unblocked**, or **Not checked**. Depending on the manageability status that is selected, a recommended action is also displayed. If the recommended action is **Unblock IPAM Access**, review information in the **Details View** to understand what type of access status is currently blocked, verify that the correct access settings are applied, and refresh server access status.

To verify managed server access

1. In the upper IPAM navigation tree, click **SERVER INVENTORY**. The list of servers that have been discovered or manually added is displayed.
2. Under **Recommended Action**, verify that **IPAM Access Unblocked** is displayed.
3. If **Unblock IPAM Access** is displayed, verify that IPAM access settings have been applied and refresh server access status.
4. Verify IPAM access settings.

To verify IPAM access settings

a. If IPAM access settings are applied using GPOs, update Group Policy and display the Resultant Set of Policy (RSoP) information on the target managed server by typing the following commands at an elevated command prompt on the target server, and then pressing ENTER:

```
gpupdate /force
```

In the command output, verify that **Computer Policy update has completed successfully** is displayed.

```
gpresult /r
```

In the command output, under **COMPUTER SETTINGS**, verify that the GPO name that is applicable to the managed server is displayed under **Applied Group Policy Objects**. For example, if the Group Policy based provisioning method is used with a GPO prefix of IPAM1, and the managed server type is DHCP, then IPAM1_DHCP should be displayed.

5. Refresh server access status.

To refresh server access status

a. In the **SERVER INVENTORY** view, select one or more servers to refresh. Hold down CTRL or SHIFT to select multiple servers.
b. Right-click the selection and then click **Refresh Server Access Status**. The IPAM **ServerDiscovery** task will run.
c. Wait for the IPAM **ServerDiscovery** task to complete, and then refresh the **SERVER INVENTORY** view (click **Refresh** next to the notification flag in

Server Manager).

See Also

Retrieve Data from Managed Servers

Retrieve Data from Managed Servers

Checklist: Deploy IPAM Server > Install IPAM Server > Choose an IPAM Provisioning Method > Configure Server Discovery > Discover Servers on the Network | Manually Add a Server to Server Inventory > Create IPAM Provisioning GPOs | Manually Configure DHCP Access Settings | Manually Configure DNS Access Settings | Manually Configure DC and NPS Access Settings > Choose Managed Servers > Verify Managed Server Access > **Retrieve Data from Managed Servers**

After you have verified the IPAM server has access to managed servers, you can retrieve data from managed servers to begin populating the IPAM database. You might also want to retrieve server data after a configuration change or just to obtain updated information. You do not have to manually retrieve data from managed servers because scheduled tasks on the IPAM server will perform this task automatically. Perform the following procedure if you do not want to wait for these scheduled tasks to run. For more information about IPAM scheduled data collection tasks, see IPAM Architecture.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Retrieving data from managed servers

Use the **SERVER INVENTORY** view to retrieve all server data. Specific types of server data can also be retrieved using views under **IP ADDRESS SPACE** (for address space data), **MONITOR AND MANAGE** (for configuration and monitoring data), and **EVENT CATALOG** (for audit data). You can retrieve data for a single managed server or for multiple managed servers.

To retrieve data from managed servers

1. In the upper IPAM navigation tree, click **SERVER INVENTORY**.
2. Select the managed servers from which you want to collect data. Press and hold CTRL or SHIFT to select multiple servers.
3. Right-click the servers that are selected, and then click **Retrieve All Server Data**.
4. The following data collection tasks will run immediately on the selected servers: **AddressExpiry**, **AddressUtilization**, **Audit**, **ServerAvailability**, **ServiceMonitoring**, **ServerConfiguration**.

5. Wait for the data collection tasks to complete.

Deploying IPAM Client

At least one IPAM client is required to deploy and configure IPAM. In order for a computer to function as an IPAM client, it must be running Windows Server® 2012, Windows® 8, or a later operating system.

Note
IPAM Client is installed by default on the same computer when you install IPAM Server. It is not required to install an additional instance of IPAM Client unless you choose not to install IPAM Client on the same computer with IPAM Server.

To connect to an IPAM server, a computer running IPAM Server must be present in the list of servers managed by Server Manager. It is not required for IPAM Client to be installed on the same computer with IPAM Server.

Installing IPAM Client

For a list of tasks and procedures required to install IPAM Client, see Checklist: Deploy IPAM Client.

See Also

IPAM Client

Deploying IPAM Server

Checklist: Deploy IPAM Client

This parent checklist includes cross-reference links to important concepts about deploying an IP address management (IPAM) client. It also contains links to subordinate checklists that will help you complete the required tasks.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy IPAM Client

	Task	Reference
☐	Review operating system and software requirements for IPAM Client, and choose an installation option.	IPAM Client
☐	Install Remote Server Administration Tools (RSAT) for Windows 8 (optional). If the IPAM client computer is running Windows® 8, you must download and install RSAT from Microsoft. If the computer is running Windows Server® 2012, RSAT is installed automatically. Continue to the next task.	Install RSAT
☐	Install IPAM Client. Perform this task if the IPAM client computer is running Windows Server 2012 and IPAM Client has not already been installed in parallel with the IPAM Server feature. If the IPAM client computer is running Windows 8 and you have installed RSAT for Windows 8, IPAM Client is already available. Continue to the next task.	Install IPAM Client
☐	Ensure the IPAM Server is available in the Server Manager server pool, and connect to the IPAM Server.	Connect to an IPAM Server

See Also

Checklist: Deploy IPAM Server

Install RSAT

Checklist: Deploy IPAM Client > **Install RSAT**

The Remote Server Administration Tools (RSAT) for Windows 8 must be installed on a computer running Windows® 8. You cannot install RSAT for Windows 8 on a computer running Windows® 7 or an earlier operating system.

RSAT includes tools for remotely managing multiple Windows Server roles and features. On a computer running Windows Server® 2012, RSAT is automatically available.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Installing RSAT

To install RSAT on a computer running **win8_client_2**, download the Microsoft Update Standalone Package (.msu) from the Microsoft Download Center, and then run the installation package.

To install RSAT on a computer running Windows 8

1. Depending on the architecture of your computer, select Windows6.2-KB2693643-x64.msu (for x64 systems) or Windows6.2-KB2693643-x86.msu (for x86 systems) from the Microsoft Download Center at Remote Server Administration Tools for Windows 8 (http://go.microsoft.com/fwlink/p/?LinkID=238560) by clicking the download button next to the appropriate file.
2. Double-click the file and then click **Open** if a security warning is displayed.
3. In the **Windows Update Standalone Installer** dialog box, click **Yes**.
4. Read the license terms and then click **I Accept** to continue the installation.
5. Wait for the installation to complete, and then click **Close**.
6. Verify that **Server Manager** is displayed on **Start**.

Important
If the signed-in user account on the computer changes, Server Manager might not be displayed on **Start**. You can find Server Manager in the list of available **Apps**.

See Also

Install IPAM Server

Install IPAM Client

Checklist: Deploy IPAM Client > Install RSAT > **Install IPAM Client**

IPAM Client is installed by default on the local computer when you install IPAM Server. IPAM Client is also available by default on a computer running Windows® 8 when you install the Remote Server Administration Tools (RSAT) for Windows 8. If you are using a remote computer running Windows Server® 2012 to connect to an IPAM server, use the following procedure to install IPAM Client.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

The server pool and IPAM Client

If the local computer is running Windows Server 2012, the local server is available in the server pool by default. If the local computer is running Windows 8, you must add at least one server to the pool that supports remote management.

If the IPAM Server computer is already available in the Server Manager server pool, you can skip the procedure below to add a server to the pool, and proceed directly to install IPAM Client. If the IPAM server is not currently available in the server pool, you can use the following basic procedure to add a server to the pool. For detailed information about adding servers to Server Manager, see Add Servers to Server Manager (http://technet.microsoft.com/library/hh831453.aspx).

To add a server to the server pool

1. On the Server Manager menu, click **Manage** and then click **Add Servers**.
2. In **Add Servers**, on the **Active Directory** tab, next to **Location** choose the domain of the destination server, next to **Operating System** choose **All**, type the name of the server next to **Name (CN)**, and then click **Find Now**.
3. Click the name of the server under **Name**, click the add icon, verify the server name is displayed under **Selected**, and then click **OK**.

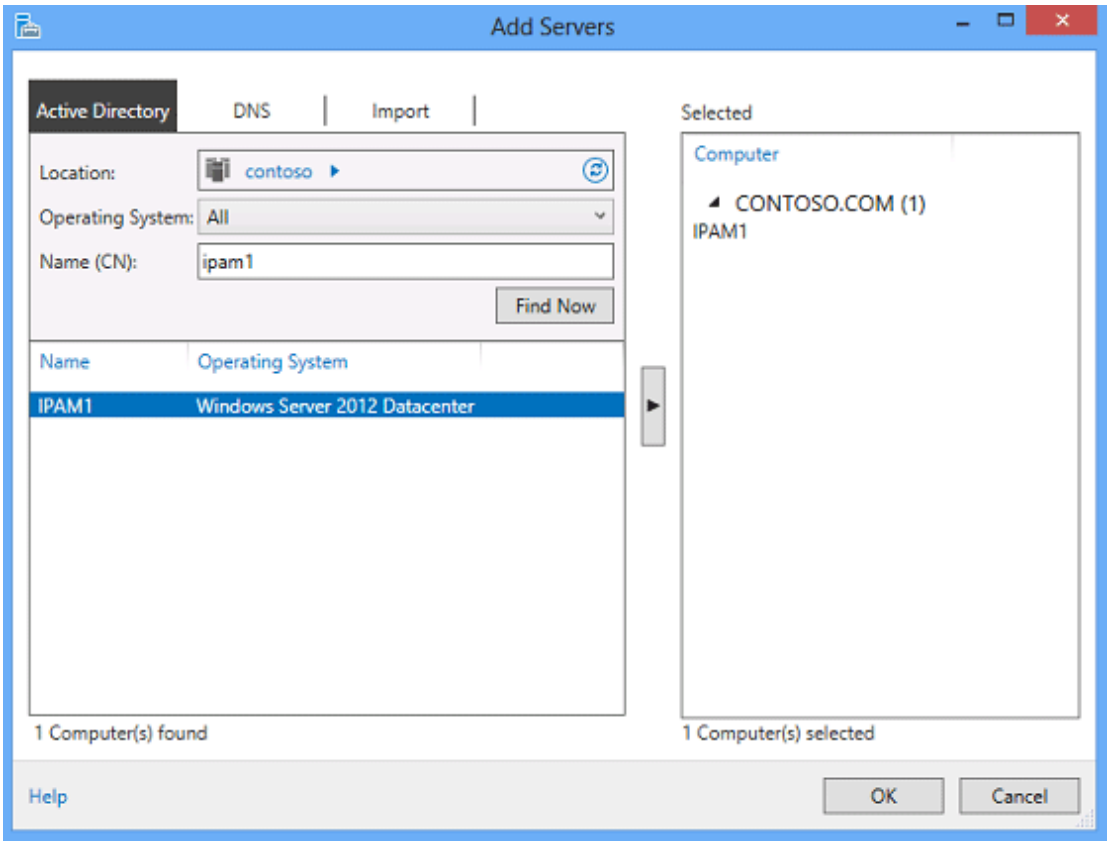

To install IPAM Client

1. In the Server Manager welcome tile, click **Add roles and features**.
2. In the Add Roles and Features Wizard, click **Next**, choose **Role-based or feature-based installation**, and then click **Next**.
3. On the **Select destination server** page, choose the server you will use to run IPAM Client and then click **Next** twice (skip the select roles page).
4. On the **Select features** page, select the checkbox next to **Remote Server Administration Tools\Feature Administration Tools\IP Address Management (IPAM) Client**.
5. If a popup dialog box is displayed requesting additional features to be installed, click **Add Features**.
6. Click **Next**, and then click **Install**.

Tip

To view the IPAM Client console, there must be at least one computer running IPAM Server available in the Server Manager server pool.

See Also

Install IPAM Server

Connect to an IPAM Server

Checklist: Deploy IPAM Client > Install RSAT > Install IPAM Client > **Connect to an IPAM Server**

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Connecting to an IPAM server

An IPAM client computer can connect to any IPAM server that is available in the Server Manager server pool. To view the list of servers in the pool, click **All Servers** on the Server Manager navigation pane.

To connect to an IPAM server

1. If the IPAM server is not already available in the Server Manager server pool, click **Manage** on the Server Manager menu, and then click **Add Servers**.
2. On the **Active Directory** tab, next to **Name (CN)**, type the hostname or FQDN of the IPAM server and then click **Find Now**.
3. Click the desired server in the list and add it to the list of selected computers.

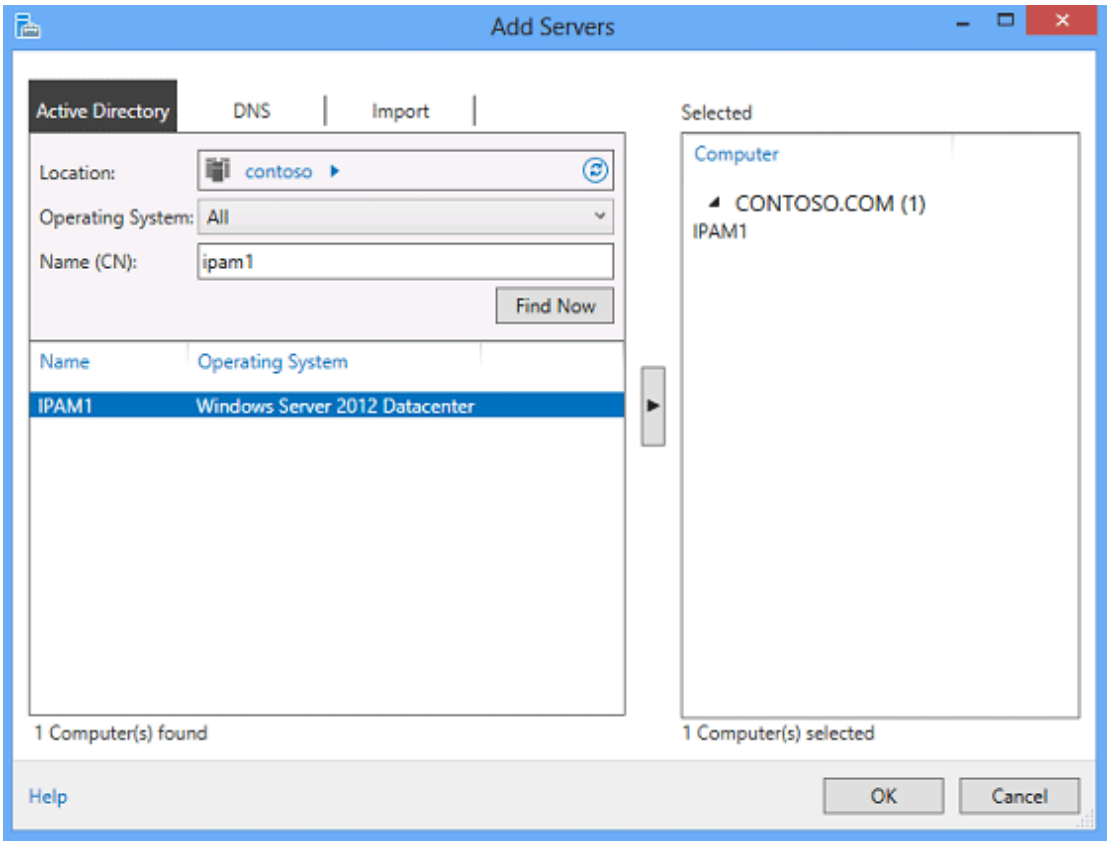

4. Click **OK** and then wait for IPAM to be displayed in the Server Manager navigation pane.
5. When the desired IPAM server is available in the Server Manager server pool, click **IPAM**, and then on the **IPAM OVERVEW** page, click **Connect to IPAM server**.
6. In the Connect to IPAM server dialog box, click the name of the IPAM server and then click **OK**.
7. Verify that Connected to <IPAM server> is displayed under **Connect to IPAM server**.

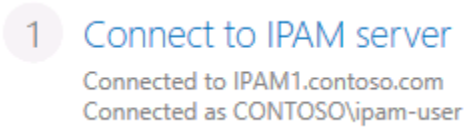

See Also

[Install IPAM Server]

Assigning IPAM Server and Administrator Roles

A single IPAM server can perform multiple management functions for multiple domains. You can also restrict an IPAM server to specific roles or limit the scope of servers to be managed. IPAM security groups can also be used to assign unique responsibilities to Administrators. See the following topics for information about configuring these roles and responsibilities.

Configuring IPAM server roles

Some reasons you might want to configure unique IPAM server roles include:

1. IPAM servers are deployed in the organization using a distributed or hybrid deployment topology.
2. Multiple IPAM servers are deployed at a single site.
3. A site managed by an IPAM server has unique requirements.
4. You are expanding your IPAM deployment.
5. You are adding new sites or domains to your organization.
6. You are configuring the scope of management for an IPAM server as part of a staged deployment.

IPAM server roles can be assigned by configuring the scope of server discovery and customizing server manageability. IPAM servers can also be assigned to specific functions such as address space management, multi-server management, or network audit by using IPAM security groups. To perform these tasks, see Configure IPAM Server Roles.

Configuring IPAM security groups

When you install IPAM Server, the following local role-based IPAM security groups are created:

- IPAM Users
- IPAM MSM Administrators
- IPAM ASM Administrators
- IPAM IP Audit Administrators
- IPAM Administrators

For more information about IPAM security groups, see IPAM Server. To configure these security groups, see Assign Administrator Roles.

See Also

What is IPAM?

Checklist: Assign Roles

This parent checklist includes cross-reference links to important concepts about assigning IPAM server and administrator roles. It also contains links to subordinate checklists that will help you complete the required tasks.

Note
Complete the tasks in this checklist in order. When a reference link takes you to a conceptual topic or to a subordinate checklist, return to this topic after you review the conceptual topic or you complete the tasks in the subordinate checklist so that you can proceed with the remaining tasks in this checklist.

Checklist: Deploy Assign Roles

	Task	Reference
☐	Review IPAM Server features and architecture.	What is IPAM? IPAM Server
☐	Configure domains, servers, and server roles to be managed by the IPAM server.	Configure IPAM Server Roles
☐	Assign IPAM administrators roles.	Assign Administrator Roles

See Also

Checklist: Deploy IPAM Server

Configure IPAM Server Roles

Checklist: Assign Roles > **Configure IPAM Server Roles**

Before you can begin using IPAM, you must configure server discovery (also called the scope of discovery) and select the manageability status of servers in the server inventory. For more information, see Getting Started with IPAM.

You can also update these settings when (for example) you add new servers or domains to your organization. The procedures to update the scope of discovery and managed server roles are the same as those used to initially configure these settings. The procedures with some additional details are also provided in this topic for convenience, and to emphasize that the management scope of an IPAM server is dynamic and can be modified at any time to adapt to your environment.

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Configuring IPAM server roles

Use the following procedures to update server discovery and managed server role settings:

1. To configure the scope of discovery
2. To configure managed servers and roles

To configure the scope of discovery

1. On the IPAM Overview page, click **Configure server discovery**. Alternatively, you can click **Manage** on the Server Manager menu, click **IPAM Settings**, and then click **Configure Server Discovery**. The **Configure Server Discovery** dialog box will open.
2. To add a domain to manage with the current IPAM server, selecting it from the drop-down list and then click **Add**.

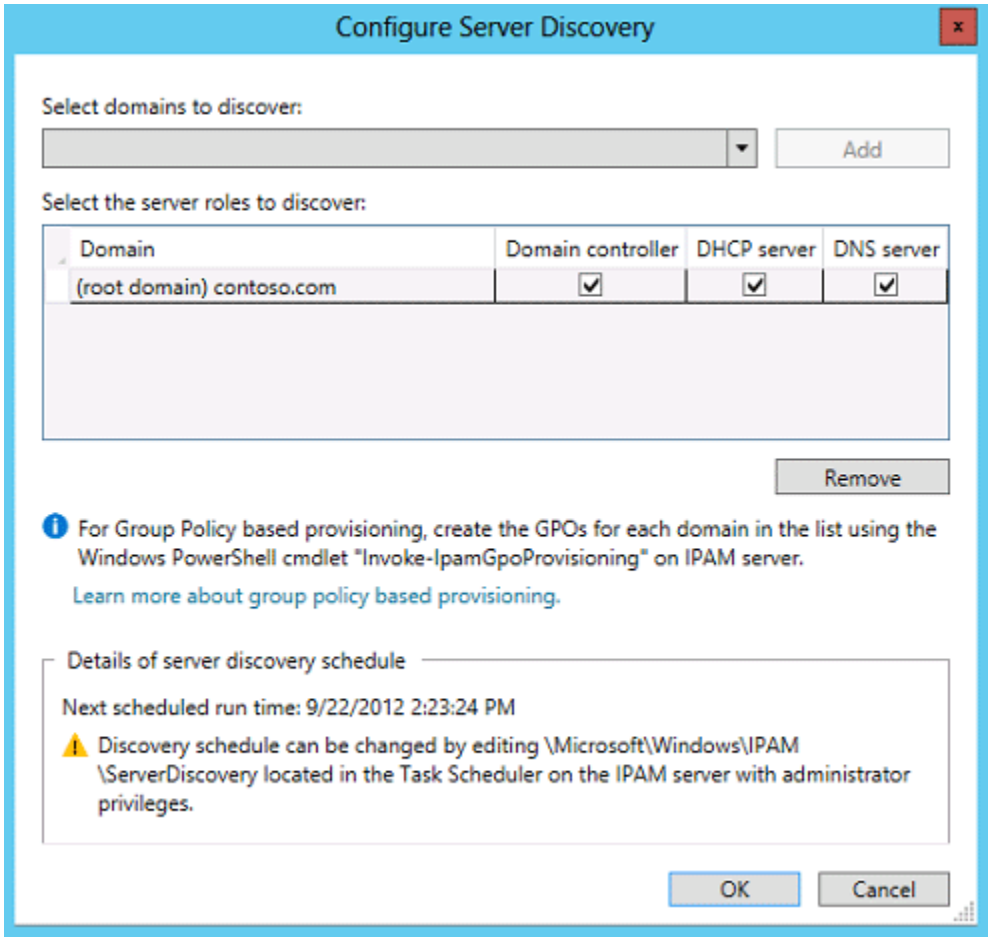

Tip

To sort items in the list by name, click a column header.

3. To remove a domain from the scope of discovery, click the domain and then click **Remove**. Removing domains from the scope of discovery will not remove managed servers from the server inventory list if they have already been discovered or added. To remove these servers, you must delete them manually.
4. By default all server roles are enabled in the domains you select. To remove a server role from the scope of discovery for a specific domain, de-select the checkbox under the appropriate server role. Selecting a server role to be discovered does not automatically enable IPAM management for the role, it only allows discovery of the role on the network.

Note

NPS servers can also be managed servers, but this server role is not automatically discovered on the network.

5. Click **OK** when you are finished choosing domains and server roles to discover.

To configure managed servers and roles

1. In the upper IPAM navigation tree, click **SERVER INVENTORY**. The list of servers that have been discovered or manually added is displayed.
2. Select one or more servers by clicking the appropriate row in the list. To select multiple servers, hold down SHIFT or CTRL when clicking.
3. Right-click the server or servers that are selected and then click **Edit Server**.

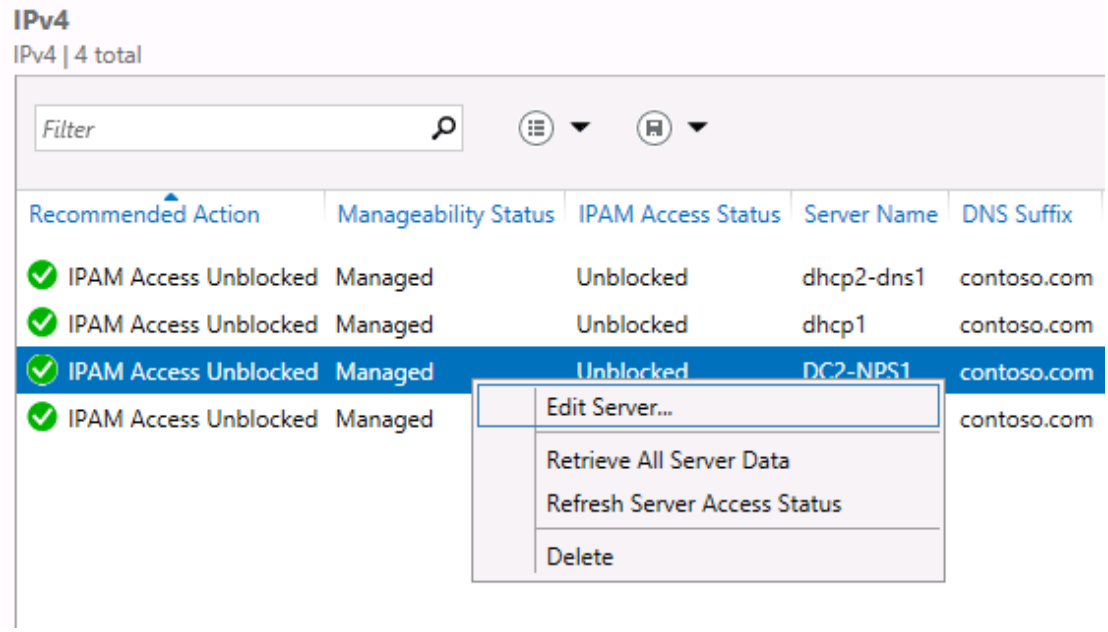

4. In the **Add or Edit Server** dialog box next to **Manageability status**, select **Managed**,

Unmanaged, or **Unspecified** from the drop-down list.

5. Select the server roles that will be managed by the current IPAM server by selecting the appropriate checkboxes next to **Server type**, and then click **OK**.

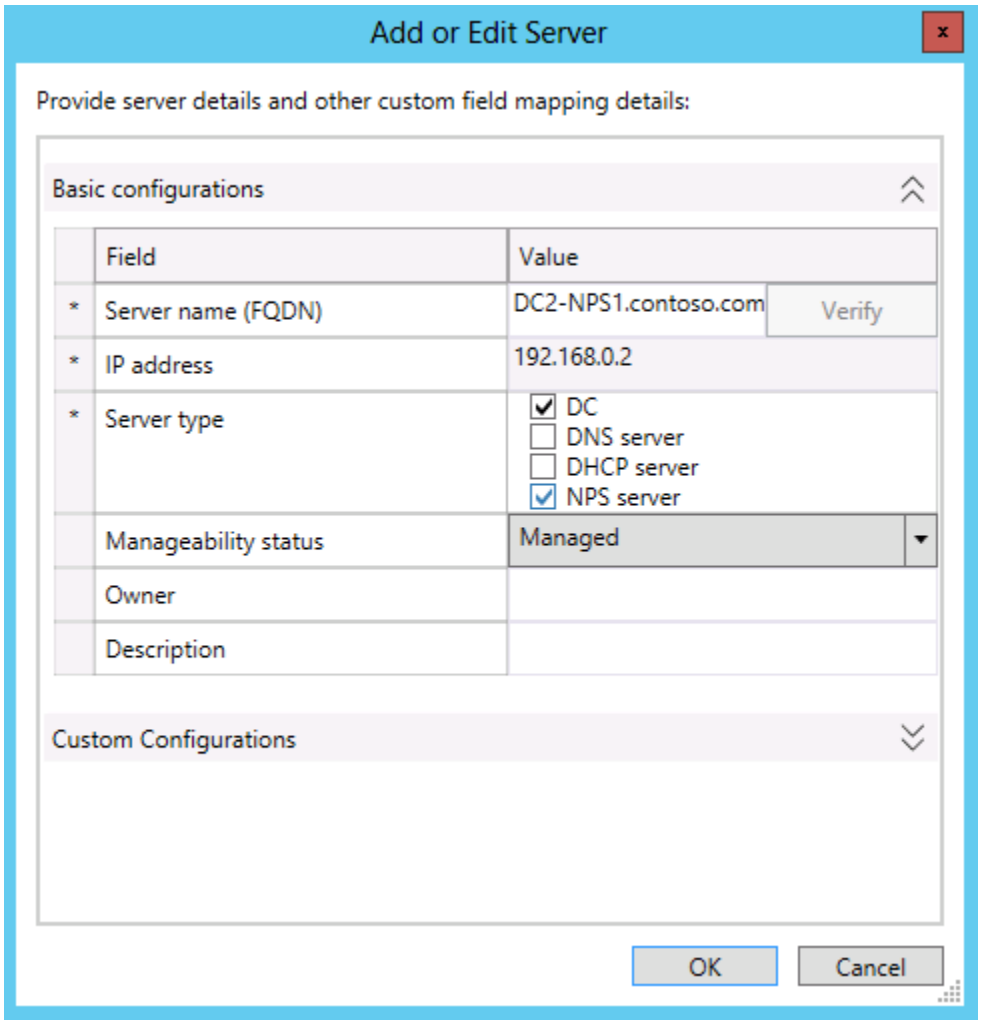

Note

If you have selected the manual provisioning method, you must apply or remove IPAM access settings to managed servers manually if you add or remove servers and server roles from the managed server inventory.

6. If a new managed server role was added, a recommended action of **Retrieve IPAM Access Status** will be displayed. To refresh access status, right-click the server in the list and then click **Refresh Server Access Status**.

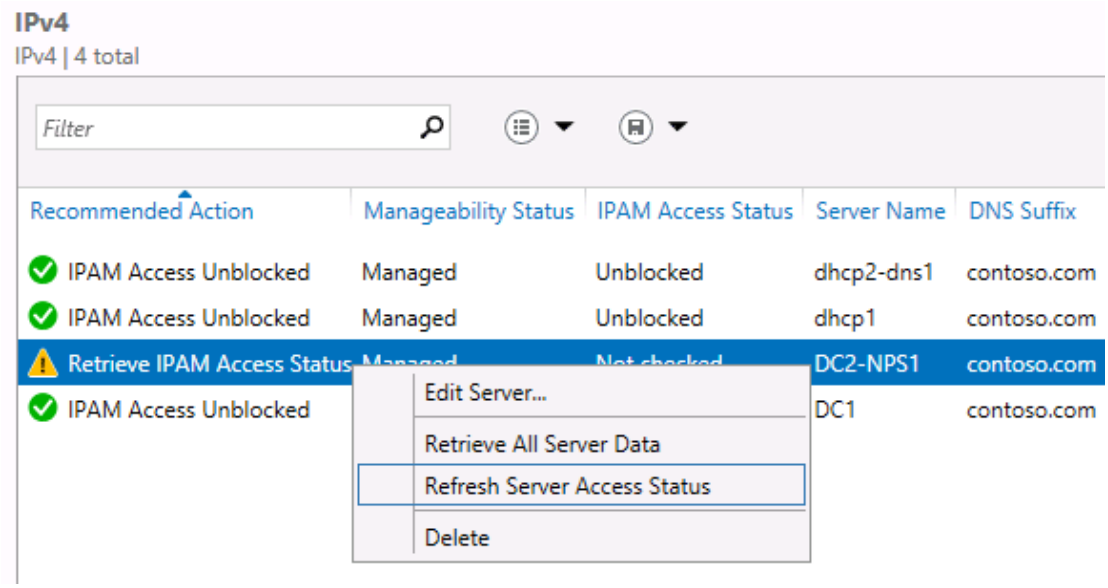

7. Wait for the IPAM **ServerDiscovery** task to complete, and then refresh the console view.

To quickly view managed server roles for all servers in the inventory, view the **Server Type** column.

See Also

Configure Server Discovery

Choose Managed Servers

Assign Administrator Roles

Checklist: Assign Roles > Configure IPAM Server Roles > **Assign Administrator Roles**

Membership in the **Administrators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Assigning IPAM administrator roles

Use the following procedures to update server discovery and managed server role settings:

To assign IPAM administrator roles

1. On the IPAM server, click **Tools** on the Server Manager menu and then click **Computer Management**.

2. In the Computer Management console tree, open **Local Users and Groups** and then click **Groups**. The following local IPAM security groups are displayed:
   - **IPAM Administrators**: IPAM administrators can view all IPAM data and manage all IPAM features.
   - **IPAM ASM Administrators**: IPAM address space management (ASM) administrators can manage IP address blocks, ranges, and addresses.
   - **IPAM IP Audit Administrators**: IPAM IP audit administrators can view IP address tracking data.
   - **IPAM MSM Administrators**: IPAM multi-server management (MSM) administrators can manage DNS and DHCP servers.
   - **IPAM Users**: IPAM users can view information in IPAM, but cannot manage IPAM features or view IP address tracking data.
3. To add a user or group to the group, double-click the group, click **Add**, type the username under **Enter the object names to select**, and then click **OK**.
4. Verify that the user you added is displayed under **Members**.

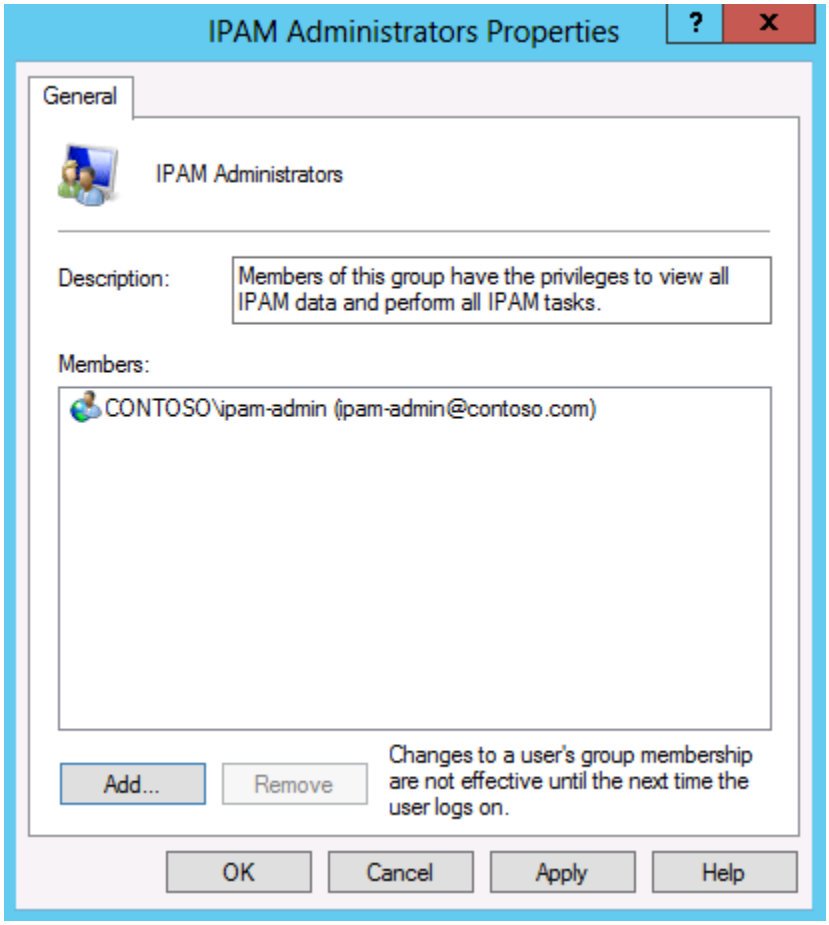

5. To remove a user or group from the group, click the user under **Members** and then click

Remove.

6. To add or remove another user or group, repeat this procedure.
7. Click **OK** when you are finished, and then close the Computer Management console.

The following table summarizes IPAM actions and functions that are permitted or denied with membership in the specified security group. Access rights are additive if a user or group is a member of multiple security groups.

Security Group	Server Inventory	IP Address Space	Monitor and Manage	Event Catalog	IP Address Tracking	Common Management Tasks
IPAM Administrators	Manage	Manage	Manage	View	View	Manage
IPAM ASM Administrators	Manage*	Manage	View	View	Denied	Manage
IPAM IP Audit Administrators	Manage*	View	View	View	View	Manage
IPAM MSM Administrators	Manage*	View	Manage	View	Denied	Manage
IPAM Users	View	View	View	View	Denied	View

*If the Group Policy-based automatic provisioning method is used, GPO access permission must be delegated to add or remove servers from GPO security filtering.

Common management tasks include:

- Connect to an IPAM server
- IPAM settings:
  - Configure server discovery
  - Configure custom fields
  - Configure utilization threshold
- Starts server discovery
- Retrieve all server data

Members of the local **Administrators** group on the IPAM server also have permission to modify the server inventory and perform common management tasks.

If a user does not have sufficient privileges to perform an action, they will receive an alert that is unique to the action they are attempting to perform.

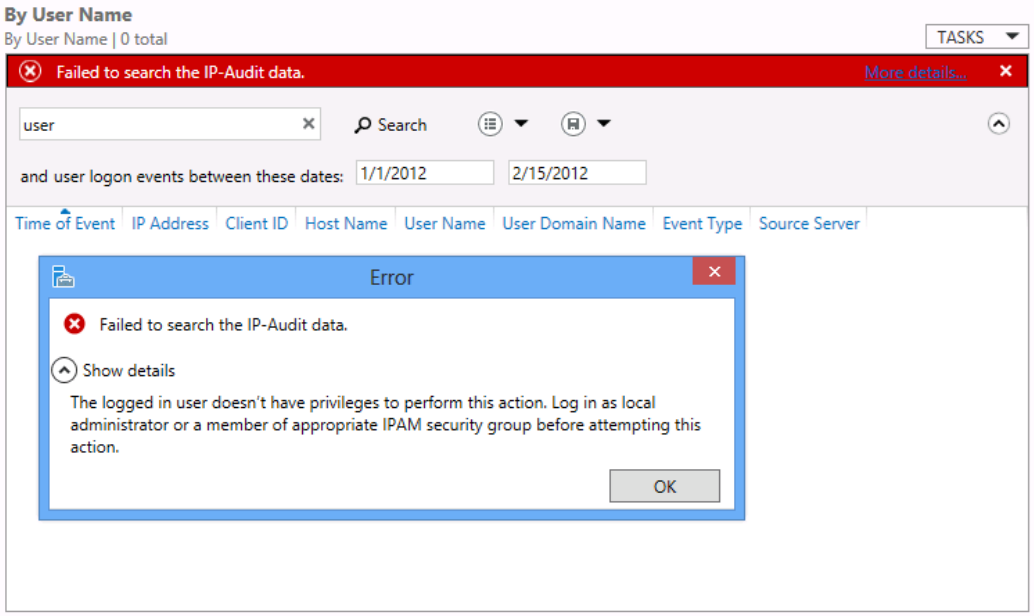

See Also

IPAM Server

IPAM Operations Guide

This guide provides administration and troubleshooting information for the IP Address Management (IPAM) feature in Windows Server® 2012.

In this guide

This guide includes the following sections:

1. Using the IPAM Client Console describes the controls, menus, and types of displays that are available in the IPAM client console. Specific items and sections in the console are identified, and their functionality is described. Information is provided about how to customize items displayed in the console.
2. Managing Server Inventory: Provides information about how to customize and organize the server inventory display.
3. Managing IP Address Space: Provides information about how to use IPAM's IP address management (ASM) feature.
4. Multi-server Management: Provides information about how to use IPAM's multi-server management (MSM) feature.

5. IP Address Tracking: Provides information about how to use IPAM's IP address tracking feature.
6. Operational Event Tracking: Provides information about how to use IPAM's operational event tracking feature.
7. Using Windows PowerShell with IPAM: Provides information about Windows PowerShell support for IPAM.
8. Best Practices: Provides recommendations for deploying and using IPAM.
9. Troubleshooting IPAM: Provides information about IPAM common problems and solutions.
10. IPAM Backup and Restore: Provides guidance for backup and restore of an IPAM server using Windows Server Backup.

See also

IPAM Planning and Design Guide

IPAM Deployment Guide

Walkthrough: Demonstrate IPAM in Windows Server 2012

Using the IPAM Client Console

The IPAM client console contains the following controls, lists and menus. Follow the links below for more information about how to use each element of the console.

1. Upper navigation pane
2. Lower navigation pane
3. Lower navigation list view
4. Lower navigation tree view
5. Display pane search view
6. Display pane details view
7. Display pane refresh
8. Notifications area
9. Server Manager menu bar
10. TASKS menu

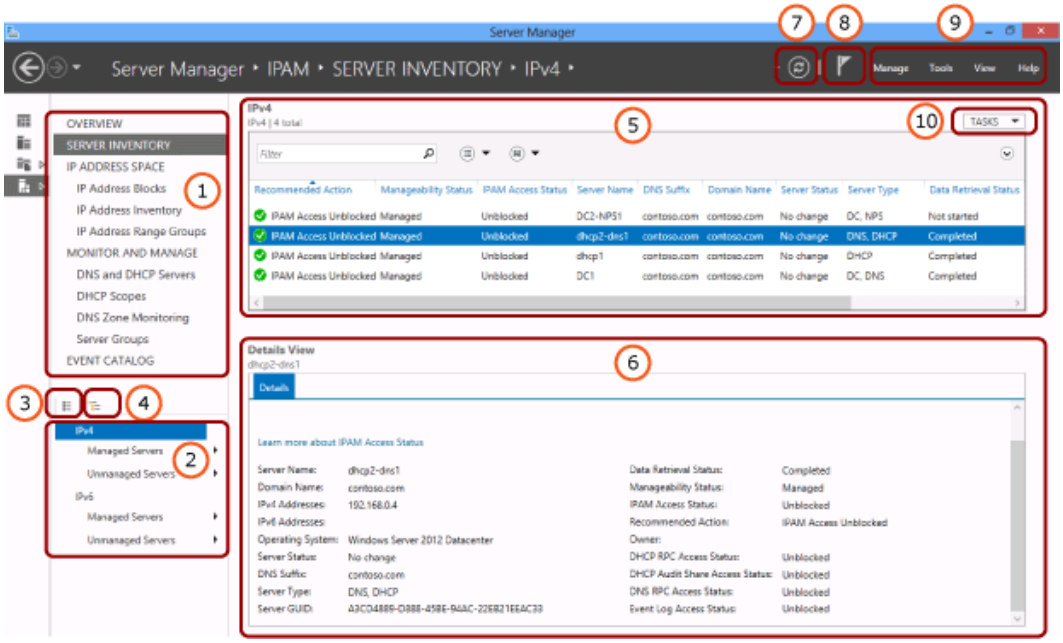

Upper navigation pane

The upper navigation pane has five parent nodes:

1. OVERVIEW
2. SERVER INVENTORY
3. IP ADDRESS SPACE
4. MONITOR AND MANAGE
5. EVENT CATALOG

Some nodes also contain child nodes that are discussed in this topic.

OVERVIEW

Use the overview node to quickly view a configuration summary for the IPAM server, perform essential management tasks, and view IPAM knowledge articles.

When you click the **OVERVIEW** node in the upper navigation pane, the following four tiles are available in the display pane:

1. The IPAM SERVER TASKS tile has three buttons:

   The **QUICK START** button provides links to perform administrative tasks required to begin using IPAM. You can also use these links after IPAM is deployed to update IPAM settings, connect to other IPAM servers, and run IPAM data collection tasks.

   The **ACTIONS** button provides quick links to multiple management actions (shown below).

   The **LEARN MORE** button provides links to several IPAM help topics.

2. The MANAGED NETWORK tile displays the current IPAM server name and a list of the domains managed by this IPAM server.
3. The CONFIGURATION SUMMARY tile provides details about the managed server access provisioning method that was chosen, IPAM scheduled data collection tasks, IPAM local security groups, and IPAM communication settings.
4. The SCHEDULED TASKS tile displays the name, status, next run time, and last run time for each of the IPAM scheduled data collection tasks.

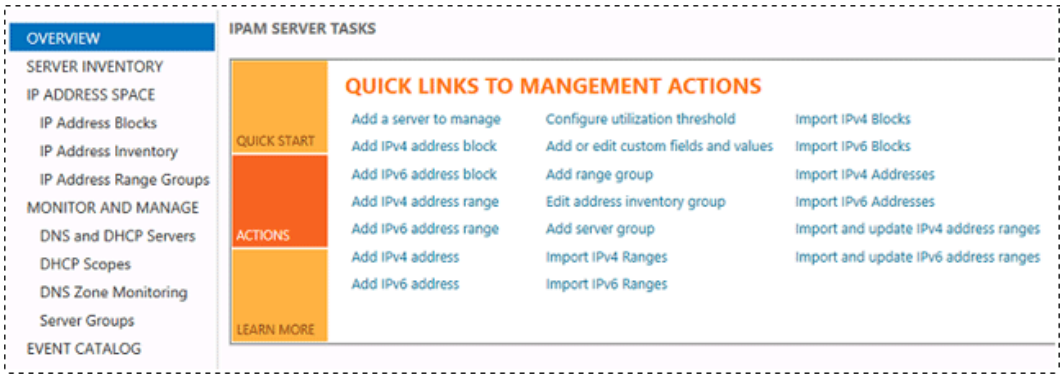

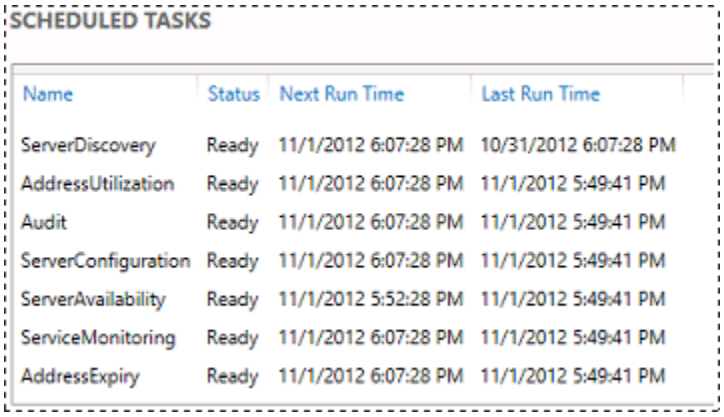

SERVER INVENTORY

Use the server inventory node to add or remove managed servers, and to update their manageability status.

When you click the **SERVER INVENTORY** node in the upper navigation pane, two tiles are available in the display pane:

1. The upper display pane search view tile provides the server inventory for servers with IPv4 connections to the IPAM server. Using the lower navigation pane, you can also choose to list

only managed or unmanaged servers, or servers with an IPv4 or IPv6 connection. Using the lower navigation pane, you can also further specify only servers with addresses in a specified range.

Note

Servers must have a valid static or dynamic IPv6 address to be displayed in the server inventory. Computers with only a link-local IPv6 address are not displayed.

2. The lower display pane details view provides configuration information for servers that are selected in the display pane search view tile. The lower details view is blank if you choose more than one server in the upper search view.

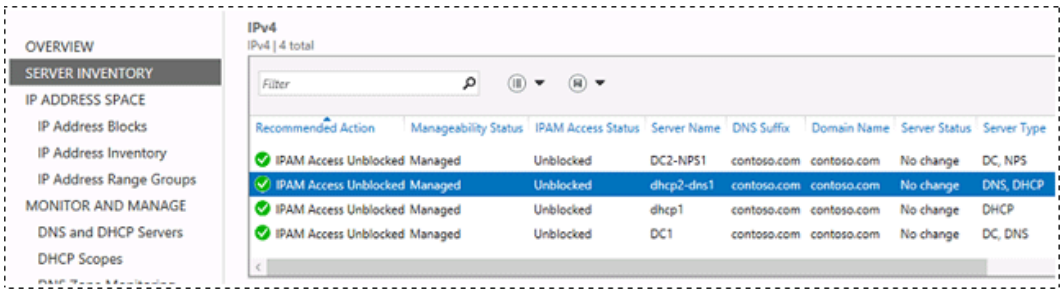

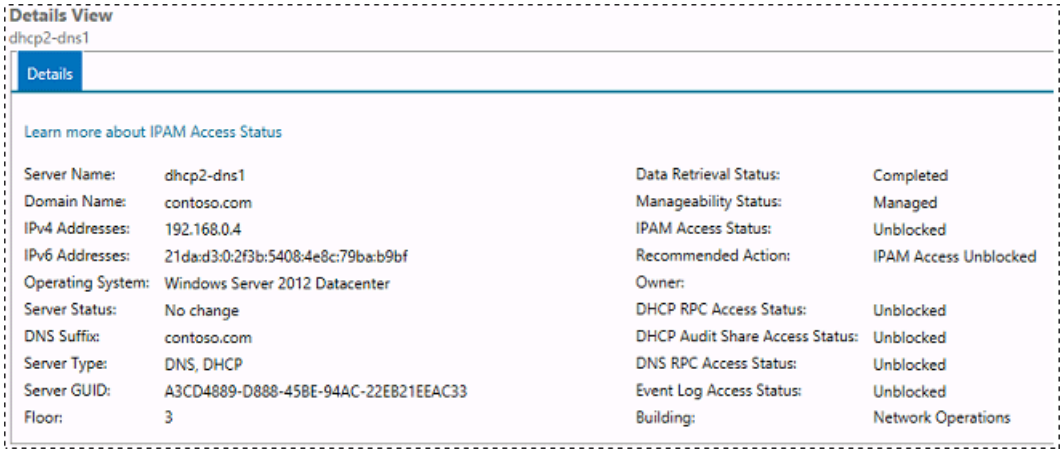

IP ADDRESS SPACE

Use the IP address space node and child nodes to configure, view, monitor and manage IP address space on your network.

IP ADDRESS SPACE contains the following child nodes:

IP Address Blocks

IP Address Inventory

IP Address Range Groups

When you click the **IP ADDRESS SPACE** node in the upper navigation pane, a summary of IPAM's address space management (ASM) feature is provided with information about each of the three child nodes. The functionality of these child nodes is also summarized below.

Context-sensitive menus

Several management actions are available in each of the **IP ADDRESS SPACE** child nodes when you right-click one or more items in the list. The actions that are available depend on the type of view that is selected (IP Address Blocks, IP Address Ranges, IP Addresses), whether multiple items are selected, and whether or not the IP address space is managed by a Microsoft DHCP server.

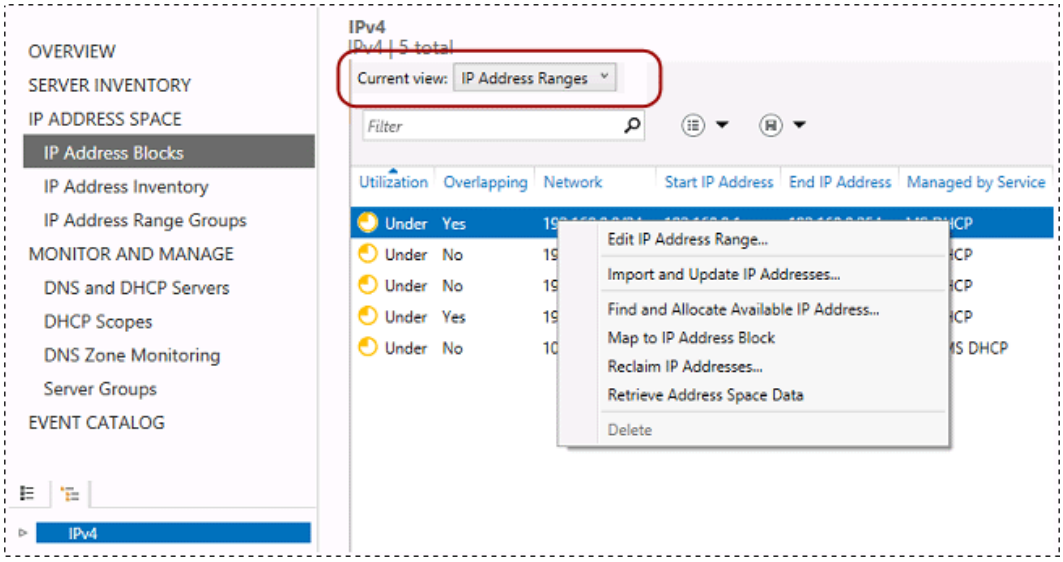

The following context-sensitive actions are available in the **IP ADDRESS SPACE** child nodes:

Current view	Action	Multiple items	MS DHCP	Non-MS DHCP
IP Address Blocks	Edit IP Address Block	✓	N/A	N/A
	Add IP Address Block		N/A	N/A
	Delete	✓	N/A	N/A
IP Address Ranges	Edit IP Address Range	✓	✓	✓
	Import and Update IP Addresses		✓	✓
	Find and Allocate		✓	✓

Current view	Action	Multiple items	MS DHCP	Non-MS DHCP
	Available IP Address			
	Map to IP Address Block		✓	✓
	Reclaim IP Addresses	✓	✓	✓
	Retrieve Address Space Data	✓	✓	✓
	Delete	✓		✓
IP Addresses	Edit IP Address	✓	✓	✓
	Create DHCP Reservation	✓	✓	✓
	Create DNS Host Record	✓	✓	✓
	Create DNS PTR Record	✓	✓	✓
	Delete DHCP Reservation	✓	✓	✓
	Delete DNS Host Record	✓	✓	✓
	Delete DNS PTR Record	✓	✓	✓
	Delete	✓	✓	✓

IP Address Blocks

Use the IP address blocks node to search and display IP address blocks, IP address ranges, and individual IP addresses on the network.

Important

When you click the **IP Address Blocks** node in the upper navigation pane, the address space that is displayed is based on queries to the IPAM database. IP addresses, ranges, and blocks in the IPAM database can be different from IP address space on the network if the IPAM database has not been recently synchronized with the network. IP address

space data is synchronized automatically by scheduled tasks if IP addresses are provided by Microsoft DHCP servers, and these tasks can be run on-demand. IP addresses on the network that are not provided by Microsoft DHCP servers must be updated manually, or automated through the use of custom scripts. Windows PowerShell for IPAM provides cmdlets that can be used to update IP address space data in the IPAM database. For more information about IPAM automation with Windows PowerShell, see Using Windows PowerShell with IPAM.

The following actions will add an IP address range to the IPAM database:

1. An active DHCP scope is discovered on the network.
2. An administrator adds an IP address range manually using the **Add or Edit IPv4 Address Range** or the **Add or Edit IPv6 Address Range** dialog box.
3. An administrator imports an IP address range manually from a comma delimited file using the IPAM console.
4. An administrator imports an IP address range manually from a comma delimited file using Windows PowerShell.
5. A regularly scheduled task runs a Windows PowerShell script that imports and updates IP address ranges automatically from DHCP and DNS servers on the network.

The information that is provided in the display pane search view depends on the node that is selected in the lower navigation tree and also the type of view that is selected next to **Current view**. You can choose the following views from the drop-down list: **IP Address Ranges**, **IP Addresses**, and **IP Address Blocks**.

These views can be further customized by choosing the type of IPv4 or IPv6 address space to display using the lower navigation pane. IPv4 address space is automatically assigned as public or private based on RFC 1918, “Address Allocation for Private Internets.”

Note

An IP address range or individual IP address is considered **Unmapped** if there is no parent IP address block or IP address range found in the IPAM database, respectively. IP address ranges are automatically mapped to parent IP address blocks unless there is a conflict with an existing IP address range. Individual IP addresses are automatically mapped to parent IP address ranges if they have the same values for **Managed by service** and **Service instance**. For more information about organizing IP address space in IPAM, see Managing IP Address Space.

IP Address Inventory

Use the IP address inventory view to display individual IP addresses on the network organized by device type.

Note

Active DHCP address leases and static IP address assignments on the network are not automatically displayed in the IP address inventory due to network and server load constraints. This information can be added to the IPAM database manually or automatically. For information about how to enter this information into the IPAM database, see Managing IP Address Space.

When you click the **IP Address Inventory** node in the upper navigation pane, the IP addresses displayed are the same as those displayed using the **IP Address Blocks** node when you choose **IP Addresses** next to **Current view**. However, different choices are available in the lower navigation pane. Instead of the public, private, and unmapped address space choices available in the IP address blocks view, the IP address inventory view organizes IP addresses by device type. To choose a specific device type, click the name of the device under IPv4 or IPv6 in the lower navigation pane.

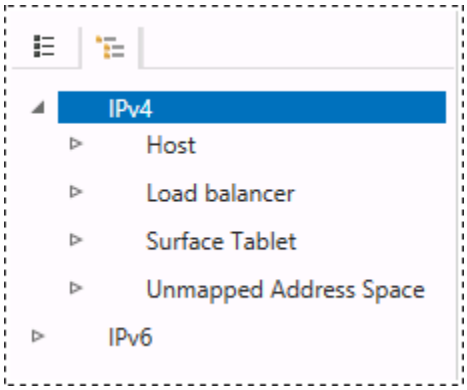

You can also add custom devices by clicking **Manage**, clicking **IPAM Settings**, and then clicking **Configure Custom Fields**. Several built-in device types are available (these cannot be deleted). You can add new device types by clicking **Device Type** in the list of custom fields and then typing a name for the new device type under **Custom Field Value**. To delete a custom device type, click the device name and then click **Delete value**.

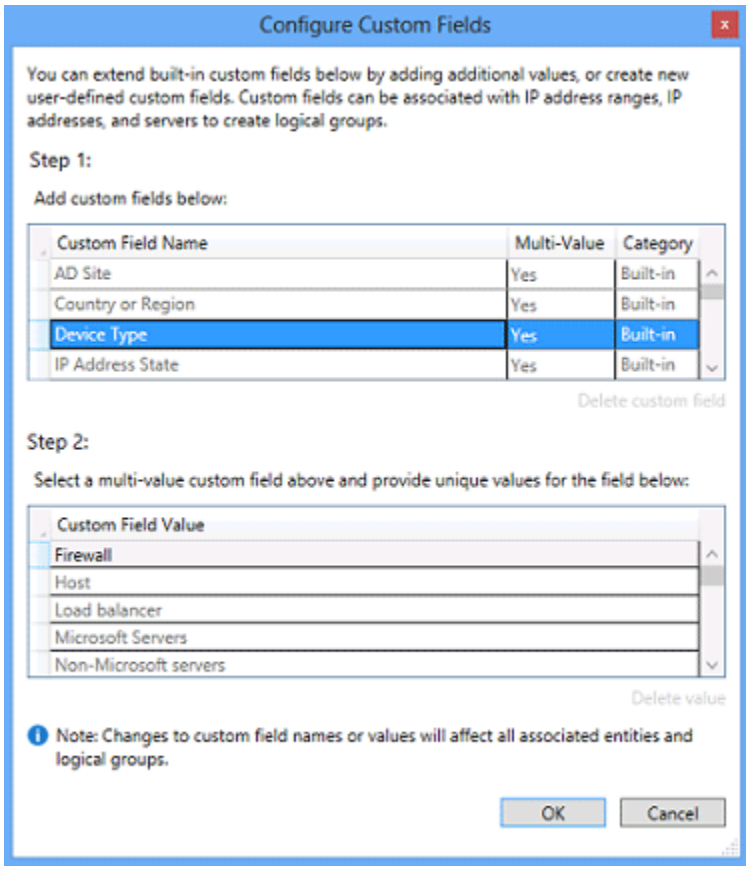

IP Address Range Groups

Use the IP address range groups view to display IP addresses ranges, IP addresses, and IP address range groups on the network organized into groups. By default, address space is organized by the **Managed by Service** field. However, you can choose to organize address space in other ways. For example, to organize IP address ranges by region, click **TASKS**, click **Add IP Address Range Group**, Type **Region** for the name of the address range group, and then choose to group items by the **Region** field. The Region field must also be defined under **Custom Configurations** for the IP address ranges you wish to display. If the **Region** field is not yet defined, you must configure it by editing the IP address range in the IP address blocks view.

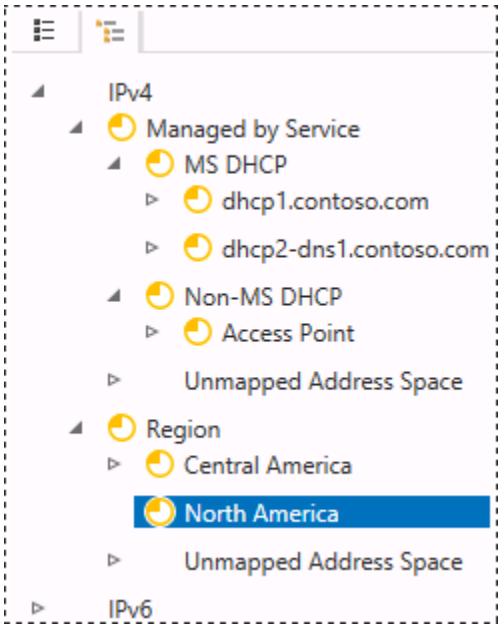

When you select **IP Address Range Groups** from the drop-down list next to **Current view**, information is displayed for each defined **Group Name** according to the selection you have made in the lower navigation pane. For example, if you choose **MS DHCP** under **Managed by Service** in the lower navigation pane, then all the service instances (host names) will be displayed for Microsoft DHCP servers on the network.

When you select **IP Address Ranges** from the drop-down list next to **Current view**, information is displayed for each IP address range according to the selection you have made in the lower navigation pane. Information displayed includes the network address of the range, starting and ending IP addresses, managed by service, service instance, assignment type, and utilization data.

The display pane details view also provides information about each IP address range that is selected in the display pane search view. Information that is displayed under **Details View** includes **Configuration Details**, **Utilization Trend**, and **Event Catalog**.

When you select **IP Addresses** from the drop-down list next to **Current view**, information is displayed for individual IP addresses on the network according to the selection you have made in the lower navigation pane. Information displayed includes the expiry status, IP address, MAC address, managed by service, service instance, parent IP address range, device name and type, address state (inactive, in use, or reserved), assignment type, and whether or not DHCP and DNS synchronization has been attempted for the IP address. Synchronization information can be particularly important when attempting to reconcile the IPAM database with configuration data on the network.

The display pane details view also provides information about each IP address range that is selected in the display pane search view. Information that is displayed under **Details View** includes **Configuration Details** and **Event Catalog**

MONITOR AND MANAGE

Use the monitor and manage node and child nodes to configure, view, monitor and manage IP infrastructure servers on your network.

MONITOR AND MANAGE contains the following child nodes:

DNS and DHCP Servers

DHCP Scopes

DNS Zone Monitoring

Server Groups

When you click the **MONITOR AND MANAGE** node in the upper navigation pane, a summary of IPAM's multi-server management (MSM) feature is provided with information about the four child nodes (DNS and DHCP Servers, DHCP Scopes, DNS Zone Monitoring, Server Groups). The functionality of these child nodes is also summarized below.

DNS and DHCP Servers

Use the DNS and DHCP servers node to search and display information about DNS and DHCP services on the network.

When you click the **DNS and DHCP Servers** node in the upper navigation pane, the information that is provided in the display pane search view depends on the node that is selected in the lower navigation tree, and also the type of view that is selected next to **Server Type** and **View**. You can choose the following views from these drop-down lists:

Choosing a server type of **DNS and DHCP** provides a view of **Server properties**.

Choosing a server type of **DHCP** provides a view of **Server properties** or **Scope properties**.

Choosing a server type of **DNS** provides a view of **Server properties**.

The server properties view provides information about the state of service availability, duration in the current state, and configuration details such as server name, role, domain, and IP address.

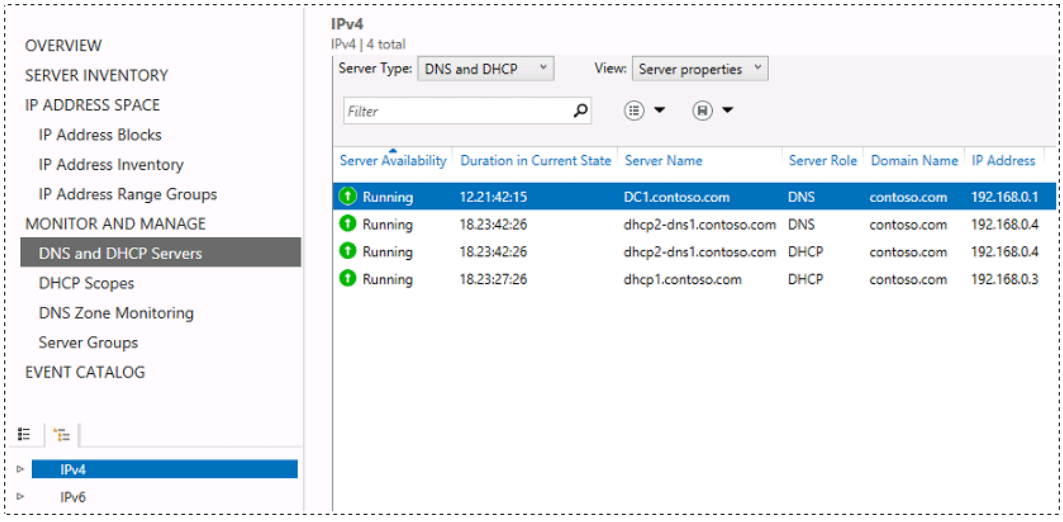

Tip

The information displayed in the **DNS and DHCP Servers** view when you choose **DHCP** next to **Server type** and **Scope properties** next to **View** is the same as the information displayed by default in the DHCP Scopes view.

Different management actions are available depending on the view that is selected. In a **Server Properties** view, the following options are available for the available choices of **Server Type**. Some options are also available when multiple servers are selected, allowing you to retrieve data and configure settings identically on all selected servers:

Server type	Configuration options
DNS and DHCP	Configuration options depend on the type of server or servers selected. If a server is running both DNS and DHCP, it is displayed twice in the list: once for the DNS server role and once for the DHCP server role. If multiple server roles (DNS and DHCP servers) are selected, then **Retrieve Server Data** is the only available option. If all selected servers are the same type (DHCP or DNS), then the available options are the same as those available if you choose **DHCP** or **DNS** for the value of **Server type**.
DHCP	**Edit DHCP Server Properties**: Configure a single server or multiple servers. **Edit DHCP Server Options**: Configure a single server or multiple servers. **Create DHCP Scope**: Configure a single server. **Configure Predefined DHCP Options**: Configure a single server or multiple servers. **Launch MMC**: Configure a single server. **Configure User Class**: Configure a single server or multiple servers. **Configure Vendor Class**: Configure a single server or multiple servers. **Retrieve Server Data**: Retrieve data from a single server or multiple servers.
DNS	**Launch MMC**: Configure a single server. **Retrieve Server Data**: Retrieve data from a

Server type	Configuration options
	single server or multiple servers.

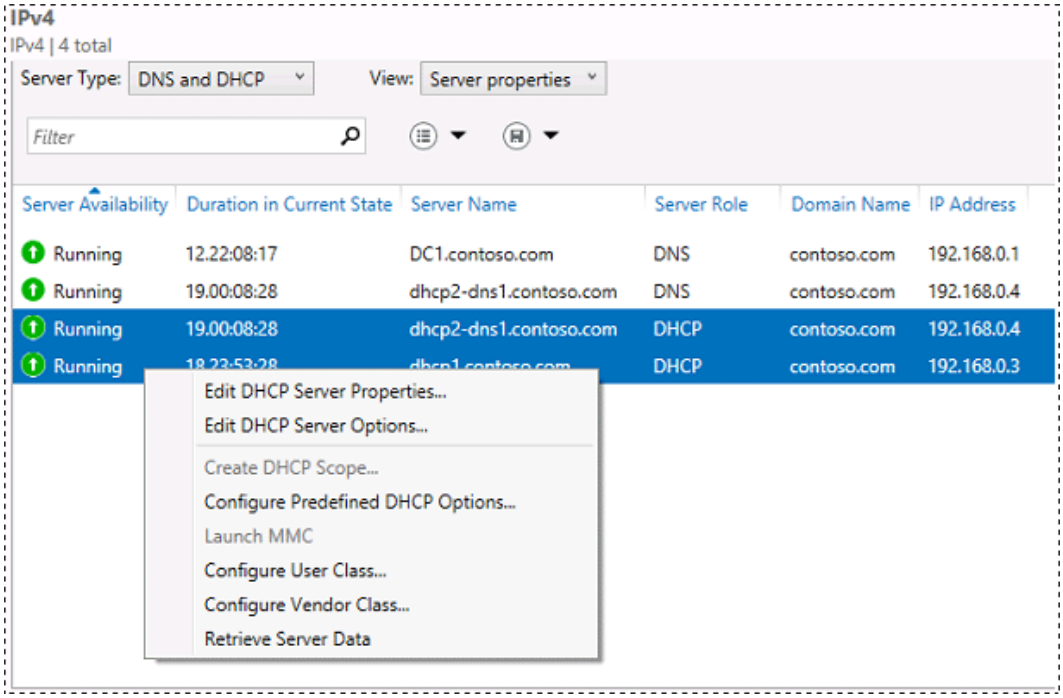

Tip

You cannot use the **Create DHCP Scope** configuration option on multiple DHCP servers. To duplicate a DHCP scope, create the scope on at least one DHCP server and then use the **Duplicate DHCP Scope** configuration option available in the **DHCP Scopes** view.

For more information about configuration options for DNS and DHCP servers using the IPAM console, see Multi-server Management.

DHCP Scopes

Use the DHCP scopes node to search and display information about IPv4 and IPv6 DHCP scopes on the network. DHCP scopes can be active or inactive.

When you click the **DHCP Scopes** node in the upper navigation pane, the DHCP scopes displayed in the display pane search view depend on the node that is selected in the lower navigation tree. By default, IPv4 is selected.

The **DHCP Scopes** view provides information about the utilization and status of DHCP scopes and summarizes configuration information. The **Details View** provides a summary of scope properties, scope options, and scope configuration events (shown below).

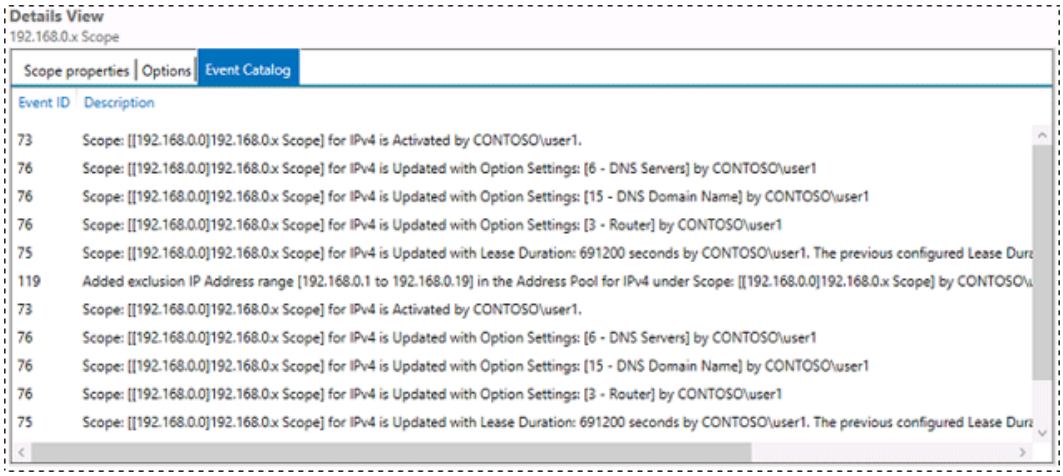

Configuration options that are available for DHCP scopes include:

- **Edit DHCP Scope**: Configure a single scope or multiple scopes.
- **Duplicate DHCP Scope**: Duplicate a single scope.
- **Activate DHCP Scope**: Activate a single scope or multiple scopes.
- **Deactivate DHCP Scope**: Deactivate a single scope or multiple scopes.
- **Delete**: Delete a single scope or multiple scopes from the DHCP server.

For more information about configuring DHCP scopes using the IPAM console, see Multi-server Management.

DNS Zone Monitoring

Use the DNS zone monitoring node to search and display information about DNS zones hosted by managed DNS servers on the network.

When you click the **DNS Zone Monitoring** node in the upper navigation pane, DNS zones hosted on managed, authoritative DNS servers are displayed. The zones that are displayed depend on the node that is selected in the lower navigation tree. By default, **Forward Lookup** zones are selected.

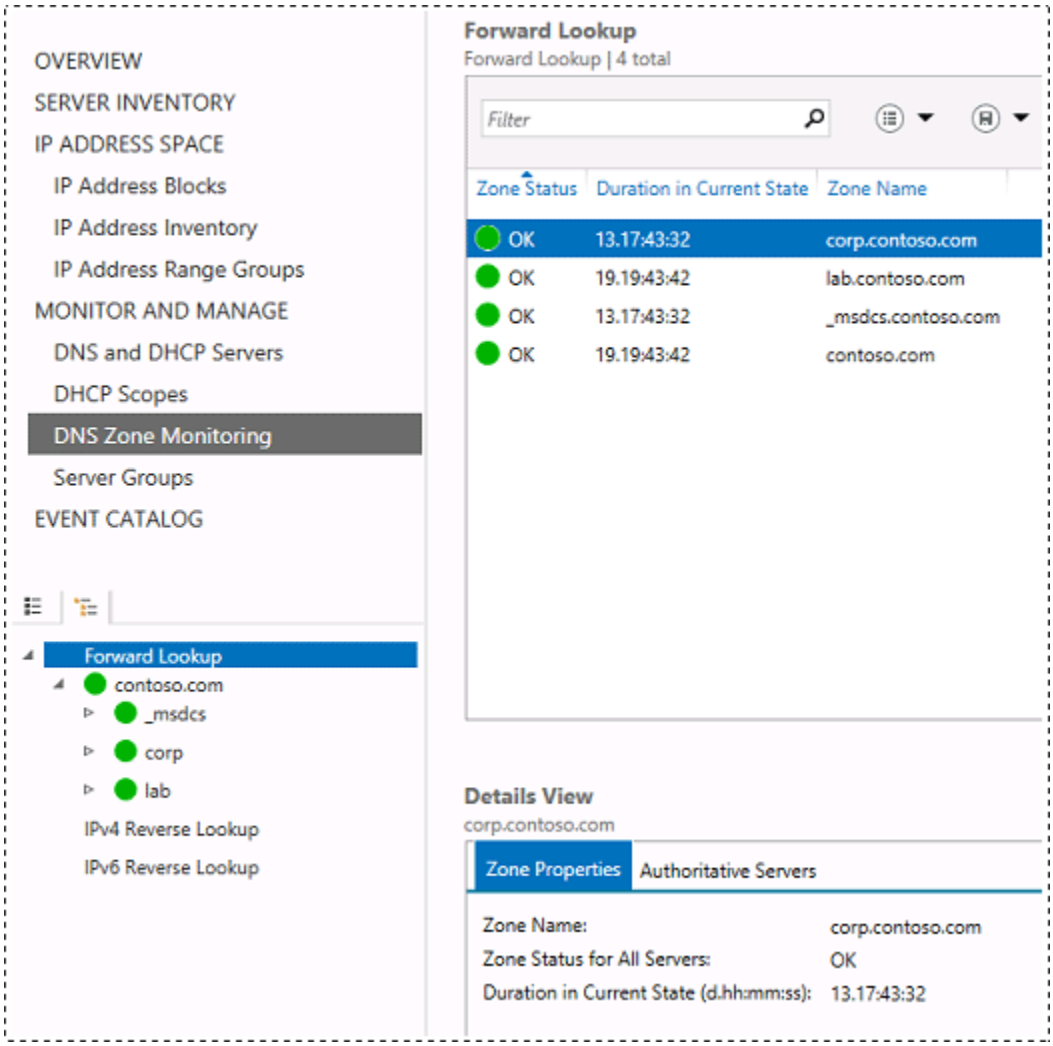

Information is provided about the health of forward DNS zones. IPv4 and IPv6 reverse DNS zones are displayed, but health information is not provided. The **Zone Status** (health) of a zone is considered OK if all DNS servers that host the zone respond to a DNS name server query for the zone (ex: **nslookup –q=ns contoso.com dc1.contoso.com**). The **Details View** provides a summary of zone properties and authoritative servers for the zone that is selected.

Server Groups

Use the server groups node to search and display information about DNS and DHCP servers that you have assigned to unique, logical groups. You can create multiple logical groups, and use this functionality to organize DNS and DHCP servers by characteristics such as location, hardware specifications, operating system version, and administrative control. If a managed server is not a DHCP server or a DNS server, it is not displayed.

If you have not yet created custom server groups, the view of DNS and DHCP servers that is displayed when you click the **Server Groups** node in the upper navigation pane is the same as

that shown when you click **DNS and DHCP Servers**. To customize this view, you must add server groups. To add a server group, right-click **IPv4** or **IPv6** in the lower navigation pane, and then click **Add Server Group**.

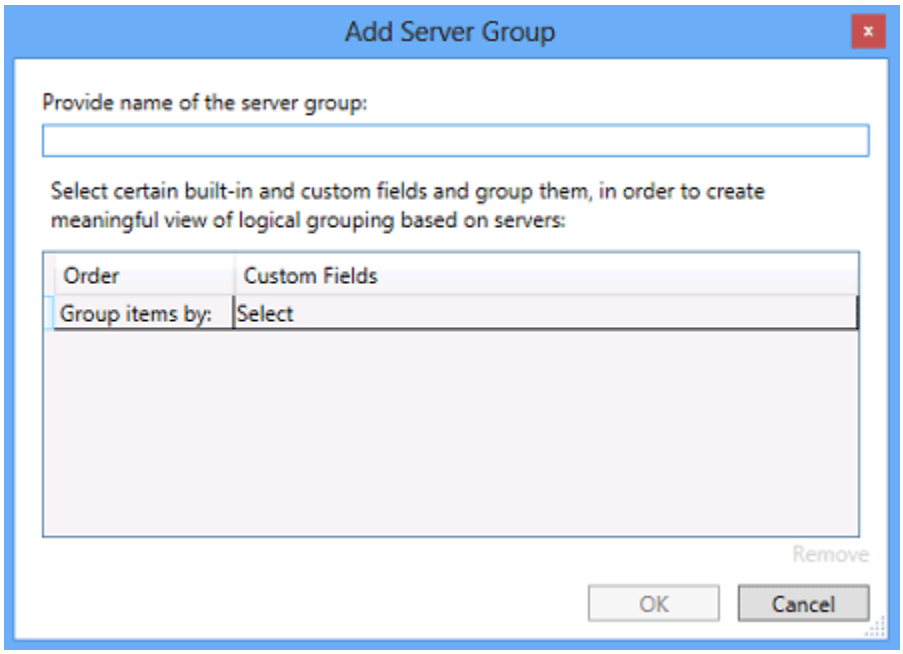

Important

This procedure is not the same as clicking **Manage** on the Server Manager menu and then clicking **Create Server Group**.

Server groups can contain multiple levels of nested grouping. To specify custom grouping values, use the Configure Custom Fields dialog box. After configuring custom fields and adding server groups, use the **SERVER INVENTORY** view to edit managed servers and specify custom configurations for each server. You can add custom configurations to multiple servers simultaneously by selecting multiple servers in the server inventory. When you have added server groups and custom configurations, you can use the lower navigation pane to view servers that belong to specific groups.

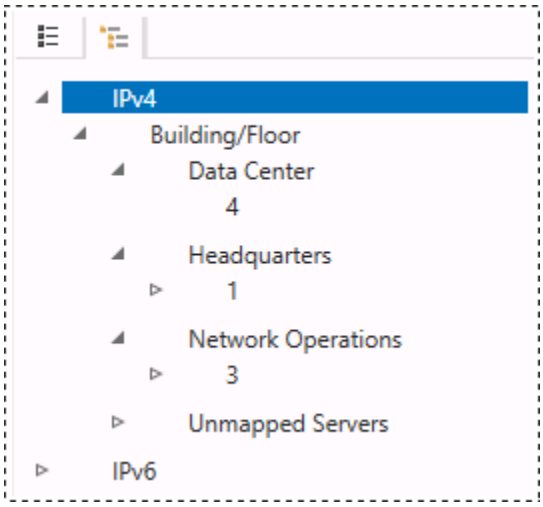

To edit or delete existing server groups, right-click the name of the server group and then click **Edit Server Group** or **Delete**.

EVENT CATALOG

EVENT CATALOG contains the following child nodes in the lower navigation pane:

IPAM Configuration Events

DHCP Configuration Events

IP Address Tracking

Use the event catalog node to view IPAM and DHCP configuration events, and to search and display DHCP lease events and user logon events on managed servers.

When you click the **EVENT CATALOG** node in the upper navigation pane, IPAM configuration events are displayed by default. Use the lower navigation pane to change the display to DHCP configuration events or IP address tracking.

To update the list of events, click **TASKS** and then click **Retrieve Event Catalog Data**, which will run the IPAM audit task. To delete events from IPAM and Event Viewer, click **TASKS** and then click **Purge Event Catalog Data**. You will be prompted to choose the date and type of events to purge from the event catalog.

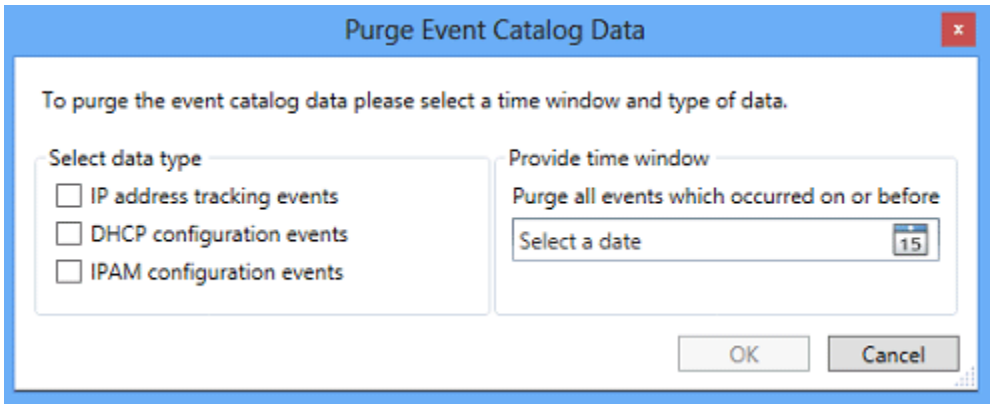

IPAM Configuration Events

When you click **IPAM Configuration Events** in the lower navigation pane, a list of configuration events is displayed in the (upper) display pane search view. These events are the same events as displayed in Event Viewer on the IPAM server under **Applications and Services Logs>Microsoft>Windows>IPAM>ConfigurationChange**. You can also filter these events by the user name of the user that made the configuration change on the IPAM server.

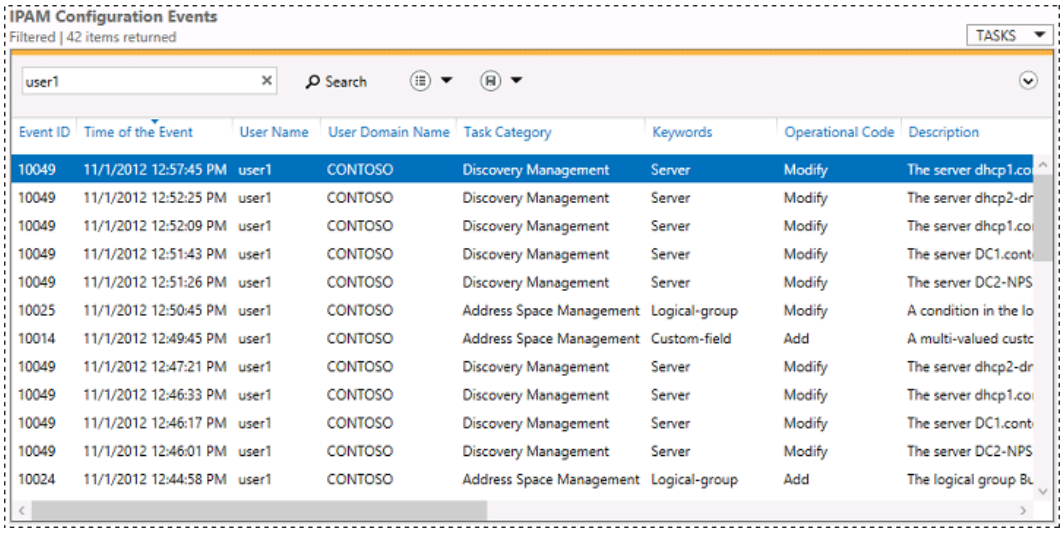

Clicking a specific event in this list will display detailed information in the (lower) display pane **Details View** (not shown here).

DHCP Configuration Events

When you click **DHCP Configuration Events** in the lower navigation pane, a list of configuration events is displayed in the (upper) display pane search view. These events are the same events as displayed in Event Viewer on the IPAM server under **Applications and Services Logs>Microsoft>Windows>DHCP-Server>Microsoft-Windows-DHCP Server**

Events/Operational. You can also filter these events by the user name of the user that made the configuration change on the IPAM server.

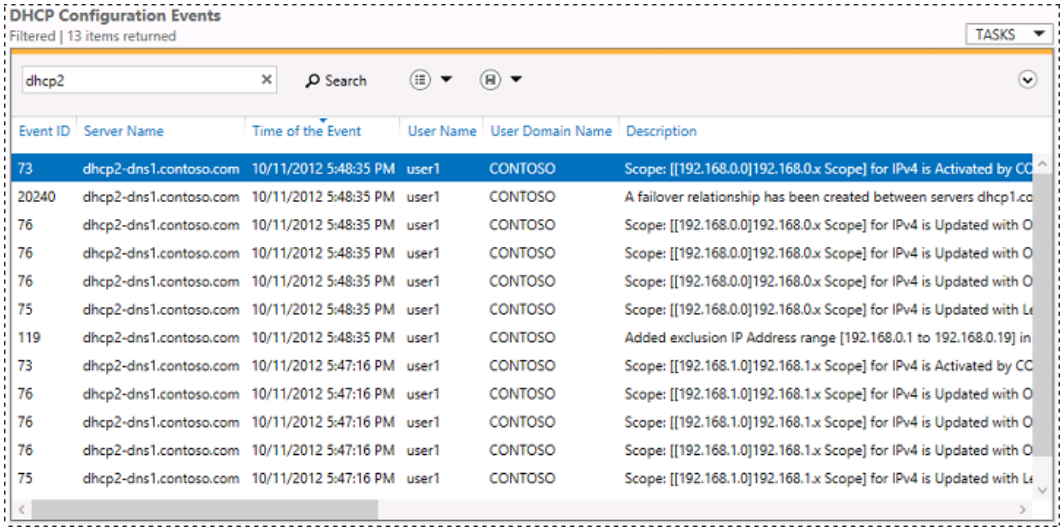

Clicking a specific event in this list will display detailed information in the (lower) display pane **Details View** (not shown here).

IP Address Tracking

When you click **IP Address Tracking** in the lower navigation pane, a summary of IPAM's IP address tracking feature is provided with information about the four types of address tracking that are available (**By IP Address**, **By Client ID**, **By Host Name**, **By User Name**). You also have the option of saving queries for later reuse. Wildcards and pattern-matching syntax is not supported. Some example queries are shown below.

By IP Address

To display DHCP lease events by IP address, type an IP address in the search box and enter a date range to search. You can also choose to include or exclude correlated logon events found on domain controllers and NPS servers.

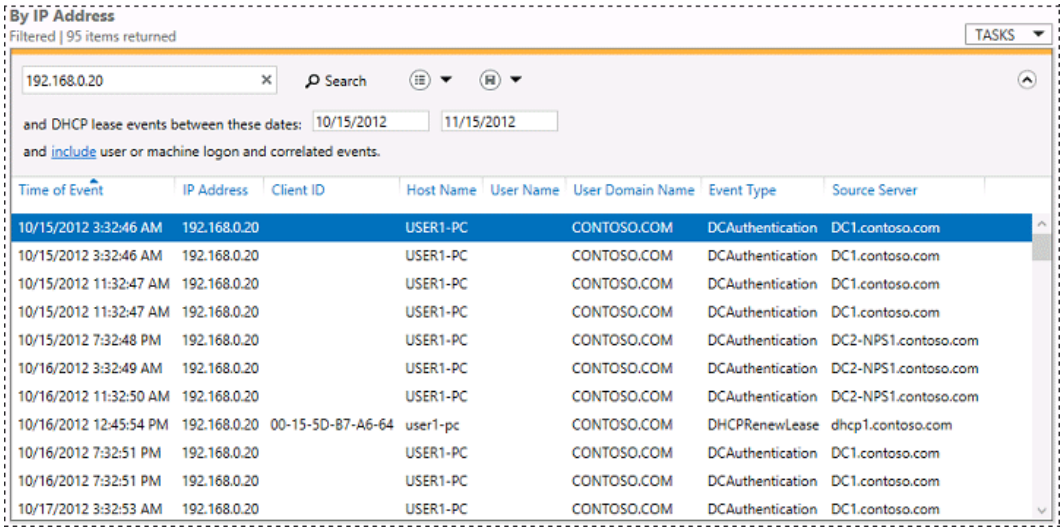

By Client ID

To display DHCP lease events by client ID, type a MAC address in the search box and enter a date range to search. You can also choose to include or exclude correlated logon events found on domain controllers and NPS servers.

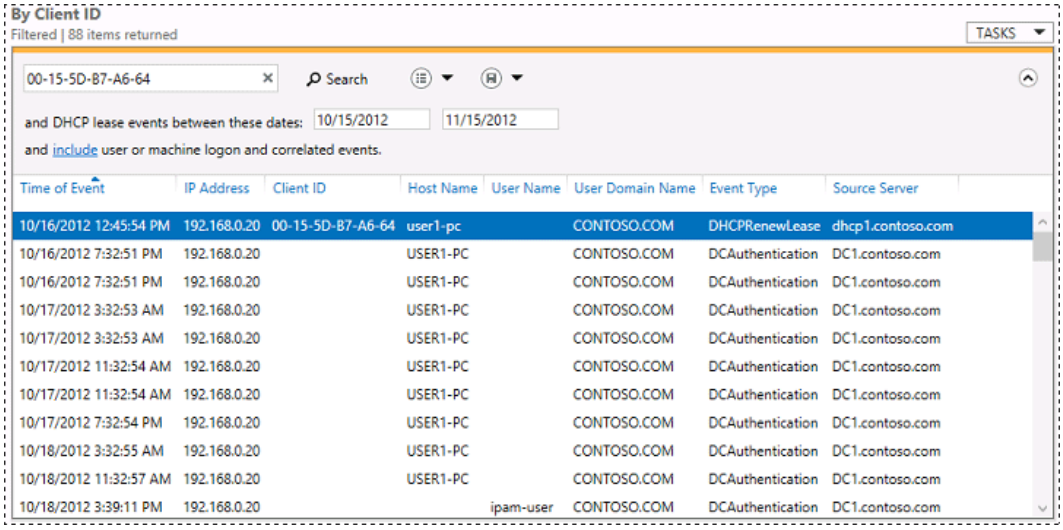

By Host Name

To display DHCP lease events by host name, type a host name in the search box and enter a date range to search. You can also choose to include or exclude correlated logon events found on domain controllers and NPS servers.

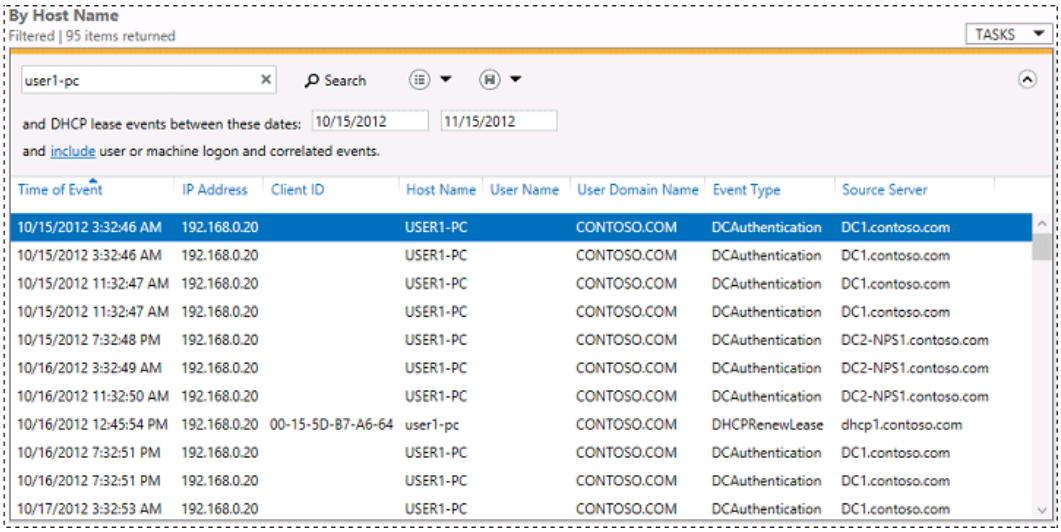

By User Name

To display DHCP lease events by user name, type user name in the search box and enter a date range to search. When searching by user name, you must include correlated logon events found on domain controllers and NPS servers.

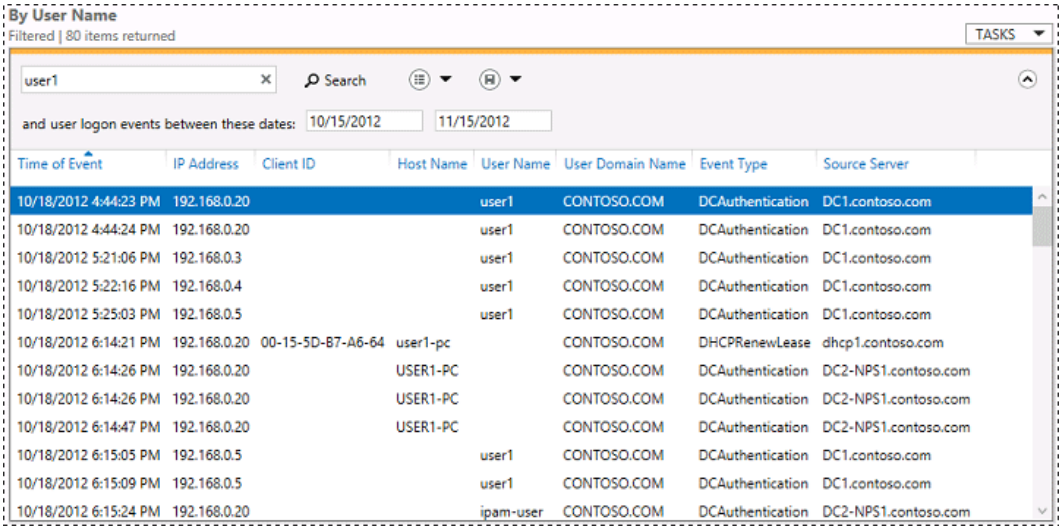

Lower navigation pane

Several examples for uses of the lower navigation pane have been provided in this topic when describing uses of the upper navigation pane. The lower navigation pane is used to group and narrow the selection of items that are displayed in the (upper) display pane search view when you click a node in the upper navigation pane.

Depending on the node that is selected in the upper navigation pane, different grouping choices are available in the lower navigation pane. For example, if you choose **IP Address Inventory** in

the upper navigation pane, the lower navigation pane will display **IPv4** and **IPv6** parent nodes with child nodes classified by device type. If you choose **IP Address Range Groups** in the upper navigation pane, the lower navigation pane will display **IPv4** and **IPv6** parent nodes, with child nodes corresponding to range groups that you have created. IPv4 and IPv6 are the typical parent classifications for items in the lower navigation pane, with exceptions of **DNS Zone Monitoring** and **EVENT CATALOG**, which have the unique parent nodes described previously.

The lower navigation pane has two types of available views, the list view and the tree view.

Lower navigation list view

List view in the lower navigation pane allows you to separate child nodes into unique columns. You can search these columns and display only specified items. You can search any part of the name.

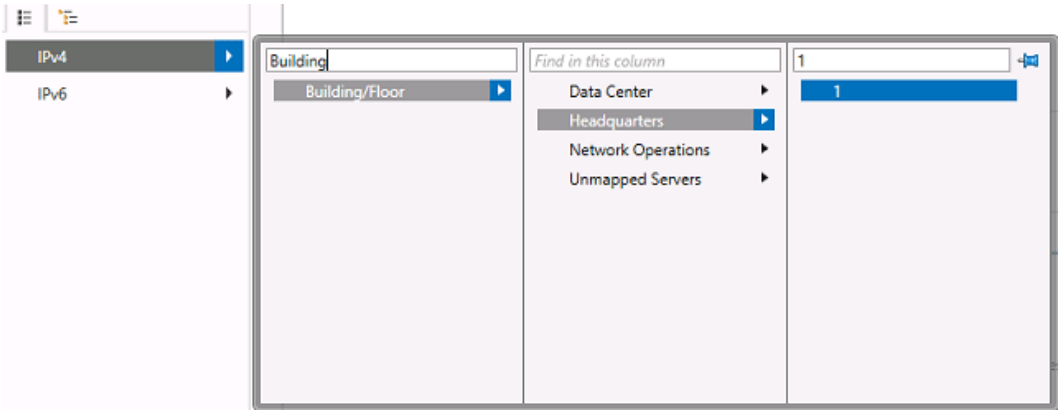

Lower navigation tree view

Tree view in the lower navigation pane allows you to expand or collapse all child nodes as needed. You cannot perform searches when using the tree view.

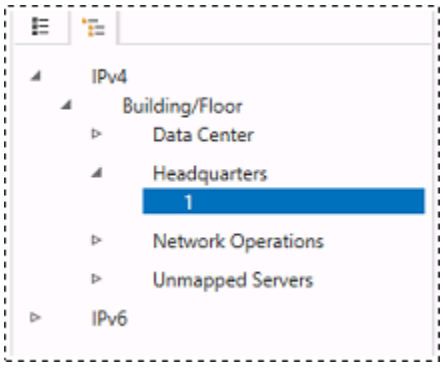

Display pane search view

The display pane search view provides a list of items found in the IPAM database that match the specific nodes that are selected in the upper and lower navigation panes, and (if applicable) the current selection next to **Current view**. To sort items in the list by a specific column header, left-click the column header. Click the column header again to reverse the sorting order.

Tip

You can customize the columns that are displayed by adding or removing column headers, and by changing the order of columns. You can also choose to group items in the list by specific fields.

To add or remove a column header, right-click and existing column header and select the headers you wish to display. To change the order of column headers, click and hold an existing header and then move it to the desired location, either before or after another header. To group items in the list, right-click an existing column header, point to Group by, and then select a grouping category. You can only choose one grouping category at a time. To remove item grouping, de-select the existing grouping category by clicking it again.

The display pane search view provides the capability to filter results based on a query. You can also save queries for later use.

Display pane details view

The display pane details view provides detailed information applicable to the current selection in the display pane search view. Depending on the item that is selected, information might be presented using multiple tabs. For example, when viewing IP address blocks and ranges, the **Details View** displays **Configuration Details**, **Utilization Trend**, and **Event Catalog** tabs. When viewing DHCP server properties, the **Details View** displays **Server Properties**, **Options**, and **Event Catalog** tabs.

Display pane refresh

Server Manager provides a control to refresh information that is displayed in specific console view. Data in the IPAM client console might need to be refreshed after a configuration change on the IPAM server or on a managed server, or when data collection tasks have recently completed running. Clicking the display pane refresh icon ensures that the information that is displayed is current.

Notifications area

Server Manager provides a notification control that displays the status of important tasks. For example, if you click Manage and then click Retrieve All Server Data, the status of multiple IPAM data collection tasks will be displayed in the notifications area.

Click the notification area and then click **Task Details** to see a recent history of tasks and errors.

Server Manager menu bar

When viewing the IPAM console, **Manage** on the Server Manager menu bar contains selections that are used for general Server Manager tasks as well as selections that are used with IPAM. Items that are specific to IPAM you can choose by clicking **Manage** include:

- Connect to an IPAM Server
- IPAM Settings
- Start Server Discovery
- Retrieve All Server Data

For more information about these management tasks, see IPAM Architecture, Managing Server Inventory, and Managing IP Address Space.

TASKS menu

The selections that are displayed on the TASKS menu are context-specific and depend on the node that is selected in the upper navigation pane. The following actions are available for specified nodes in the navigation pane:

1. **SERVER INVENTORY**: For more information, see Managing Server Inventory.
   - Add Server
   - Retrieve All Server Data
   - Export
2. **IP ADDRESS SPACE** child nodes: For more information, see Managing IP Address Space.
   - Add IP Address Block
   - Add IP Address Range
   - Add IP Address
   - Import IP Address Blocks
   - Import IP Address Ranges
   - Import IP Addresses
   - Import and Update IP Address Ranges
   - IP Address Expiry Log Settings
   - Retrieve Address Space Data
   - Export
3. **MONITOR AND MANAGE** child nodes: For more information, see Multi-server Management.
   - Retrieve Server Data
   - Export
4. **EVENT CATALOG**: For more information, see IP Address Tracking and Operational Event Tracking.
   - Purge Event Catalog Data
   - Retrieve Event Catalog Data
   - Export

See also

IP Address Tracking

Using Windows PowerShell with IPAM

Managing Server Inventory

This topic contains information and procedures you can use to manage the server inventory in IPAM. To review best practices for server discovery and configuring managed server settings, see Discovery and Provisioning.

Also see the following topics in the IPAM Planning and Design Guide for information about discovering, adding and removing servers from the server inventory, choosing a manageability status, and enabling IPAM access:

Scope of discovery

Discovering servers

Configuring managed server access settings

Select managed servers

Verify managed server access

See the following topics in the IPAM Deployment Guide to review step by step procedures for managing the server inventory:

Configure Server Discovery

Discover Servers on the Network

Manually Add a Server to Server Inventory

Create IPAM Provisioning GPOs

Manually Configuring Managed Server Access Settings

Choose Managed Servers

Verify Managed Server Access

Retrieve Data from Managed Servers

Server inventory

The IPAM server inventory is not the same as the pool of computers that are available for remote management in Server Manager. The IPAM server inventory is a list of computers located on IPAM's managed network that have been discovered or manually added to the IPAM database. The server inventory can include computers that are managed, and computers that are not managed by the IPAM server.

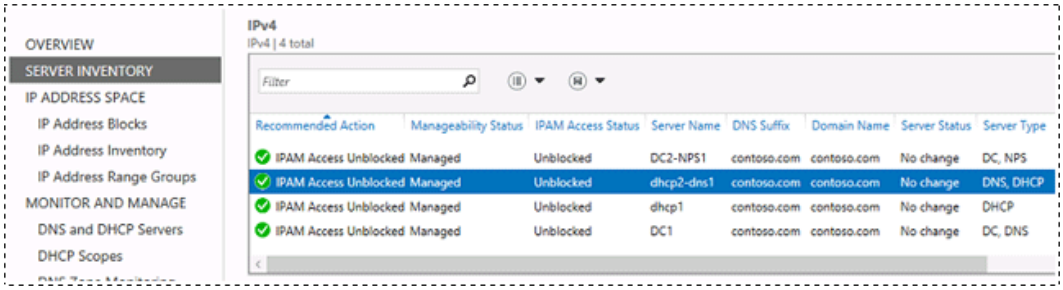

Use the lower navigation pane to view only managed servers or unmanaged servers.

Adding servers to the inventory

IPAM leverages Active Directory to discover and validate the computer accounts of servers in the server inventory. There are two methods to add servers:

1. **Discovery**: The IPAM **ServerDiscovery** task will automatically add domain controllers, DNS servers, and DHCP servers to the server inventory that are within the scope of discovery and found in Active Directory. NPS servers are not automatically discovered and must be added manually. By default, the task runs once every 60 minutes.

   You can run the **ServerDiscovery** task on-demand by clicking **Manage** and then clicking **Start Server Discovery**. To configure a different time interval for the **ServerDiscovery** task, open Task Scheduler on the IPAM server and navigate to **Microsoft>Windows>IPAM**.
2. **Manual**: If a computer account is present in Active Directory, it can be added to the server inventory. The computer's fully qualified domain name must also resolve to an IPv4 or IPv6 address using DNS. If you add a server manually, at least one server role (domain controller, DNS server, DHCP server, or NPS server) must also be selected manually. To add a server to the inventory manually, click **TASKS** and then click **Add Server**.

Important

If IPAM Server is installed on the same computer with DHCP Server, then IPAM will not discover DHCP servers on the network. Other computers running DHCP Server must be added manually. Also, in order for IPAM to manage DHCP on the local server, **Network Service** (instead of IPAMUG or the IPAM server's computer account) must be added to the local **DHCP Users** and **Event Log Readers** security groups.

Manageability status

Servers that are added to the server inventory will have a default management status of **Unspecified**. Other choices for manageability status are **Managed** and **Unmanaged**. Choosing a manageability status for computers in the server inventory has the following effects:

1. **Access settings**: If you are using the automatic GPO-based provisioning method, choosing a manageability status of **Managed** will automatically add the computer name to security filtering for the appropriate GPOs based on server roles that are selected next to **Server type**. Selecting a status of **Unmanaged** will remove the computer name from security filtering for all GPOs. Selecting a status of **Unspecified** will make no changes to the current security filtering configuration. If you are using the manual provisioning method, you are required to make these changes manually. Changes to Group Policy will not take effect until Group Policy is refreshed on managed or unmanaged servers.
2. **IPAM database**: When a computer is added to the server inventory, information about the computer is added to the IPAM database by data collection tasks. This information includes general posture data such as the computer's IPAM access status and can also include role-based details such as DHCP scope properties, IP address utilization statistics, DNS zone health, and event catalog data.

If a computer is assigned a manageability status of **Managed**, then IPAM gathers and stores historical role based data. If a computer is assigned a manageability status of **Unmanaged**, then IPAM does not store role based data and will purge any data that current exists for this computer. If a computer is marked as **Unspecified**, IPAM does not attempt to gather new role-based data from the computer, but it will maintain all data that currently exists.

IPAM access status

See Verify managed server access and Troubleshooting IPAM for information about applying and verifying IPAM managed server access settings.

Use the details view to review specific information about computers in the server inventory.

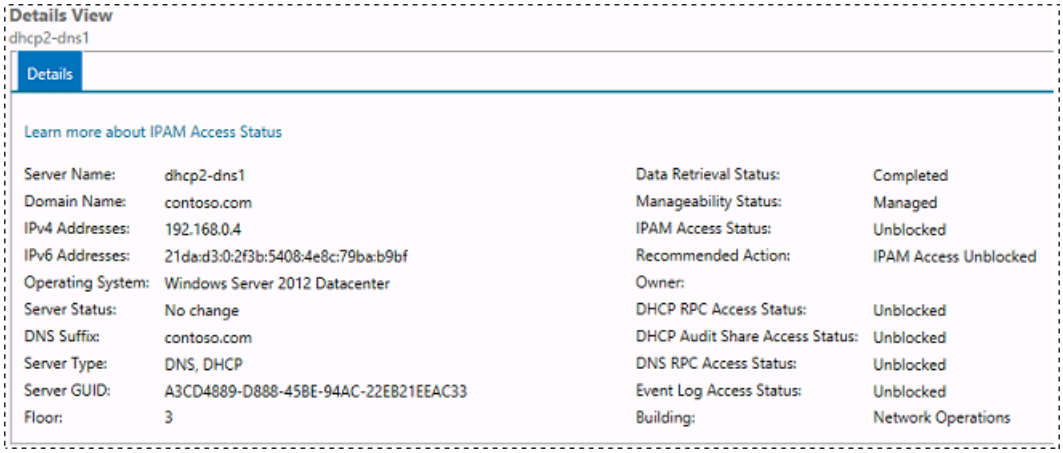

Some quick fixes to unblock IPAM access on managed servers include:

1. Use Group Policy Management to verify that the computer's domain and host name is listed in the appropriate GPOs under **Security Filtering**, or verify the configuration of manually configured settings.
2. Use Group Policy Results (gpresult.exe) to verify that Group Policy has been applied to managed servers.
3. Refresh server access status information by running the IPAM **ServerDiscovery** task. To run this task, right-click the computer in server inventory and then click **Refresh Server Access Status**.
4. Refresh the IPAM client console view.
5. If DHCP or DNS security groups were modified, restart the DHCP Server or DNS Server services to apply new settings and permissions.

See also

Walkthrough: Demonstrate IPAM in Windows Server 2012

Managing IP Address Space

This topic contains detailed information and procedures you can use to manage IP address space with IPAM. To review best practices for IP address space management, see IP address management.

See the following sections in this topic.

IP address space mapping
- IP address blocks
- IP address ranges
- IP addresses

IP address operations
- Add or edit IP address space
  - IPv4 blocks
  - IPv6 blocks
  - IPv4 ranges
  - IPv6 ranges
  - IPv4 addresses
  - IPv6 addresses
- Import
- Import and update
- Export
- Find and allocate available IP address
- Reclaim IP addresses
- IP address expiry log settings

IP address inventory
- Adding IP addresses to the inventory
  - Basic configurations
  - Managing DHCP reservations
  - Managing DNS records
  - Custom configurations
  - Importing DHCP lease data from the network

IP address range groups

IP address space mapping

In IPAM, IP address space is divided into blocks, ranges, and individual addresses. See the following figure.

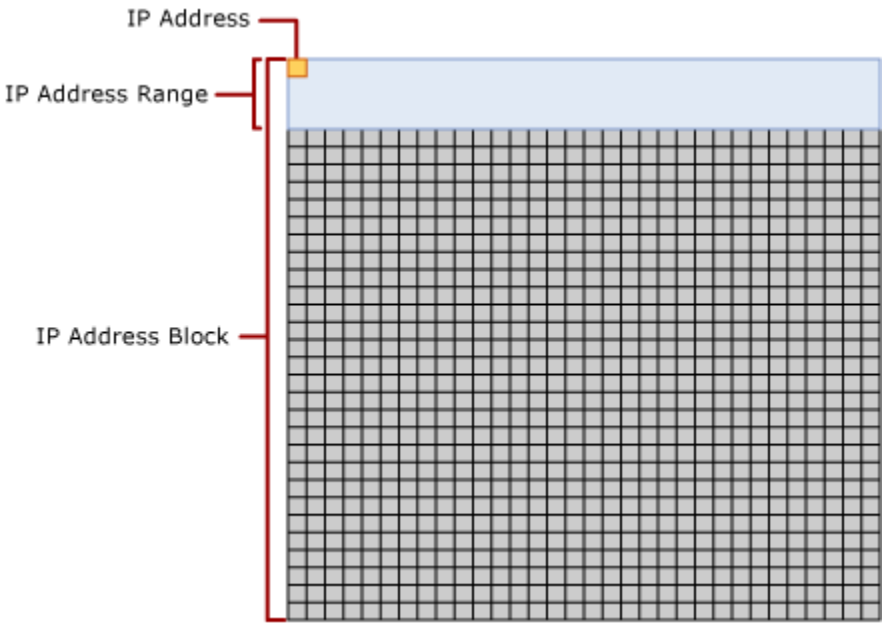

IP address blocks

IP address blocks are large chunks of IP addresses that are used to organize address space at a high level.

For example, you might use one IP address block for all private IP addresses in your organization, and another block for public IP addresses. Alternatively, you might assign smaller IP address blocks based on criteria such as location, purpose, or owner. You can think of IP address blocks as containers that hold IP address ranges. IP address blocks are not deployed and managed on the network like IP address ranges or individual IP addresses.

IP address ranges

IP address ranges are smaller chunks of IP addresses that typically correspond to a DHCP scope. IP address ranges are contained within, or “mapped to” IP address blocks. IP address ranges cannot map to multiple IP address blocks, and ranges that map to the same block cannot overlap.

Mapping IP addresses to ranges and IP address ranges to blocks enables accurate, broad, and detailed visualization of IP address utilization statistics.

IP addresses

Just as IP address ranges map to an IP address block, individual IP addresses map to a single IP address range. When you map an IP address to a range, it enables actions to be taken on a range that affect all IP addresses in the range, such as adding, updating, or deleting IP address fields.

IP address operations

IP address ranges are automatically entered into the IPAM database when DHCP scopes are discovered on managed DHCP servers. IP address blocks and individual IP addresses are not automatically added to the IPAM database, and IP address ranges that are not DHCP scopes on managed Microsoft DHCP servers are not automatically discovered. IPAM provides several dialogs to allow you to enter and edit IP address data manually, and you can also import IP address data directly from a file. Data can also be exported from IPAM to a file.

Add or edit IP address space

To add a new IP address block, range, or address, click **IPv4** or **IPv6** in the lower navigation pane, click **TASKS**, and then click **Add IP Address [Range/Block]**.

To edit existing address space, right-click the unit of address space (address, range, or block) and then click **Edit IP Address [Range/Block]**.

If you choose **No** next to **Automatically assign address values** when creating a new address block or range, then you must manually provide the **Start IP address** and **End IP address**. This allows you to work with IP address space that does not always start and end on network boundaries.

You cannot specify a **Managed by service** value of **MS DHCP** for IP address ranges that are manually added with the user interface. The **MS DHCP** value is reserved for DHCP scopes that are discovered on the network. Discovered DHCP scopes are added to IP address ranges automatically. Most values for these ranges are automatically populated based on discovery data, and cannot be modified.

IPv4 blocks

Type the network ID next to **Network ID** and choose a prefix length using the drop-down list next to **Prefix length**. If the IP address space is private, you will not be able to provide RIR information. If the IP address space is public, RIR information is optional.

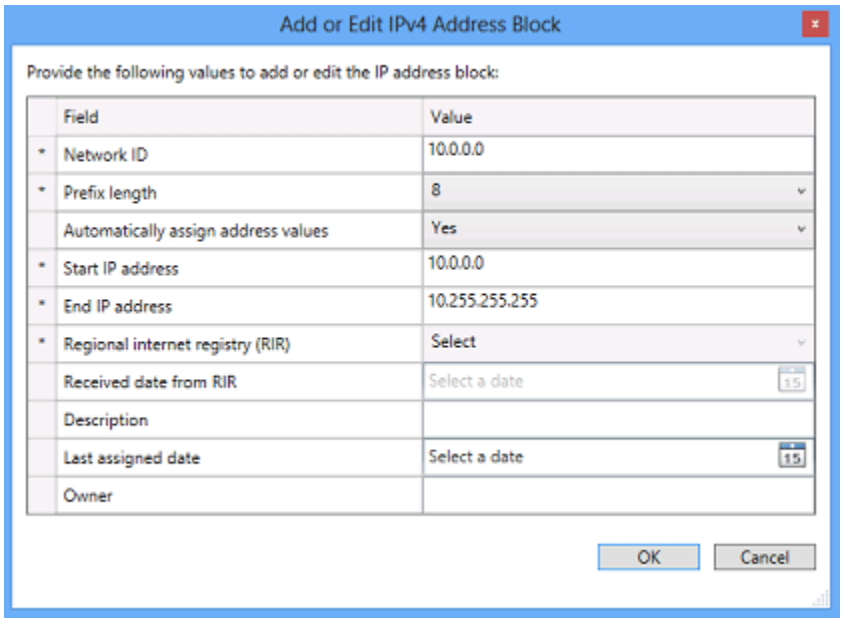

IPv6 blocks

Type the network ID under **Specify Network ID** and then slide the pointer to the desired prefix length, or type the prefix length next to the slider under **Specify Prefix length**. Optionally provide additional details in the fields provided.

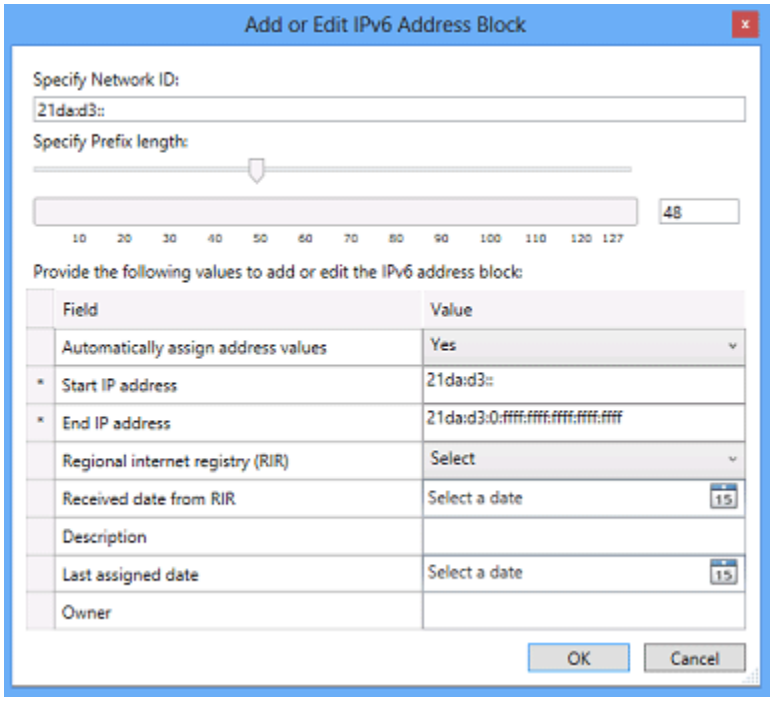

IPv4 ranges

Type the network ID next to **Network ID** and choose a prefix length using the drop-down list next to **Prefix length**. A value of **IPAM** is selected for **Managed by service** and a value of **Localhost** is selected for **Service instance** by default. You can choose other built-in values for these fields, or you can add your own custom values. For information about how to add custom values, see IP Address Inventory.

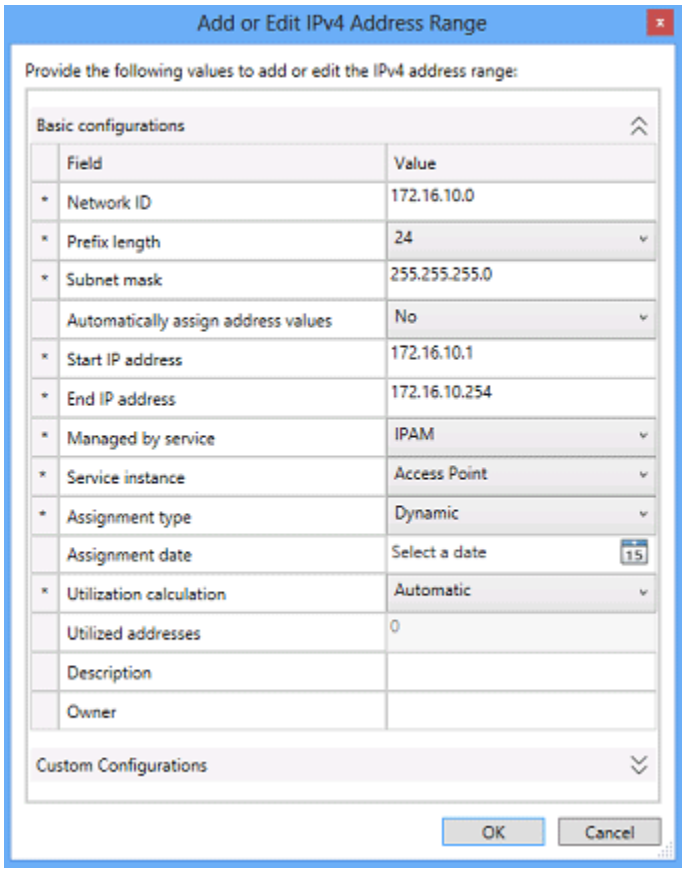

IPv6 ranges

Type the network ID under **Specify Network ID** and then slide the pointer to the desired prefix length, or type the prefix length next to the slider under **Specify Prefix length**. Optionally provide additional details in the fields provided. A value of **IPAM** is selected for **Managed by service** and a value of **Localhost** is selected for **Service instance** by default. You can choose other built-in values for these fields, or you can add your own custom values. For information about how to add custom values, see IP Address Inventory

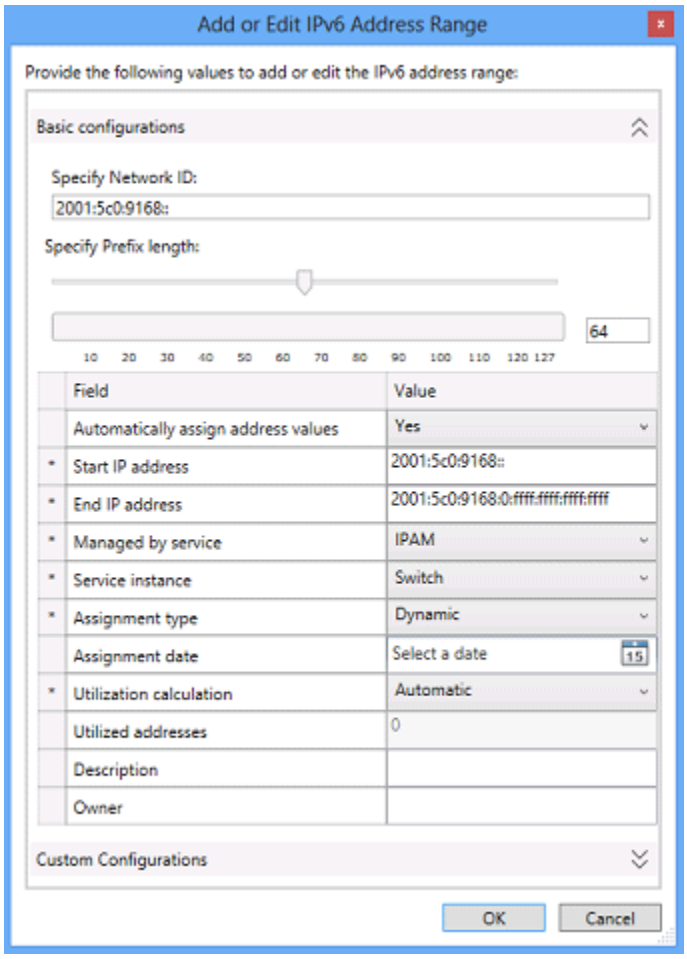

IPv4 addresses

See Adding IP addresses to the inventory.

IPv6 addresses

See Adding IP addresses to the inventory.

Import

When importing address space data from a file, required fields are the same as those that are required when adding data to IPAM with the user interface. You must include a list of valid field names on line 1 of the file, followed by one line per entry of corresponding data. Custom field names can be used, but these must be defined prior to import. Fields can be in any order, as long as you ensure that the data values are also in the same order. Data must be valid for the corresponding field type. Data and field names can be enclosed in quotes, are not case-sensitive, and can include spaces. For example, the following data can be used to import two IP addresses

into the IPAM database, assuming that dhcp1.contoso.com is a valid **Service instance** on the network:

```
IP address,managed by service,service instance,device type,ip address state,assignment type
192.168.0.25,ms dhcp,dhcp1.contoso.com,host,in-use,static
192.168.0.26,ms dhcp,dhcp1.contoso.com,host,in-use,static
```

For IP address ranges and blocks, the network ID and network prefix length are combined in a single field named **Network**. For example, the following data can be used to import an IP address block of 65.52.0.0/14. This example includes optional spaces between the field names and data values:

```
Network, start IP address, end IP address, RIR
65.52.0.0/14, 65.52.0.0, 65.52.255.255, ARIN
```

Because 65.52.0.0/14 is public IP address space, the RIR field is required. Also note that blocks begin on a network ID and end on a broadcast address, unlike IP address ranges which start and end on usable IP addresses.

If a required field is missing or contains incompatible data, an error report is automatically created in the current user's **Documents** folder. For example, the following data will generate an error if a **Managed by service** value of **MS DHCP** is specified. This value is reserved for DHCP scopes on managed DHCP servers. To avoid this error, use a value of **IPAM** for **Managed by service**. This example includes optional quotes around the field names and data values:

```
"Network", "Start IP address", "End IP address", "Managed by service", "Service instance", "Assignment Type"
"192.168.100.0/24", "192.168.100.1", "192.168.100.254", "IPAM", "router", "dynamic"
```

IP address import mandatory fields: IP address, managed by service, service instance, device type, IP address state, assignment type.

IP address range import mandatory fields: Network, start IP address, end IP address, managed by service, service instance, assignment type

IP address block import mandatory fields: Network, start IP address, end IP address, RIR

You can also import data into IPAM using Windows PowerShell. See Using Windows PowerShell with IPAM for more information.

Import and update

The import and update process is only available for IP address ranges. When importing IP address range data, you have the option to update all range data, including deletion of ranges that are no longer present. If you use this process, you do not need to manually delete or edit ranges that are no longer used or have changed. If you perform an import only (not import and update), then no ranges are deleted if they are not present in the data file. The import and update process is specific for a defined **Managed by service** and **Service instance** pair.

For example, if you have a data file that contains 5 IP address ranges with a **Managed by service** value of **Non-MS DHCP** and a **Service instance** value of **Switch**, and you currently have 7 IP address ranges in the IPAM database with values of **Non-MS DHCP** and **Switch** for these two fields, then performing an import and update with the data file will result in the deletion of at least 2 IP address ranges. Depending on the data, ranges might also be replaced with different ranges. The old dataset is overwritten by the data contained in the file.

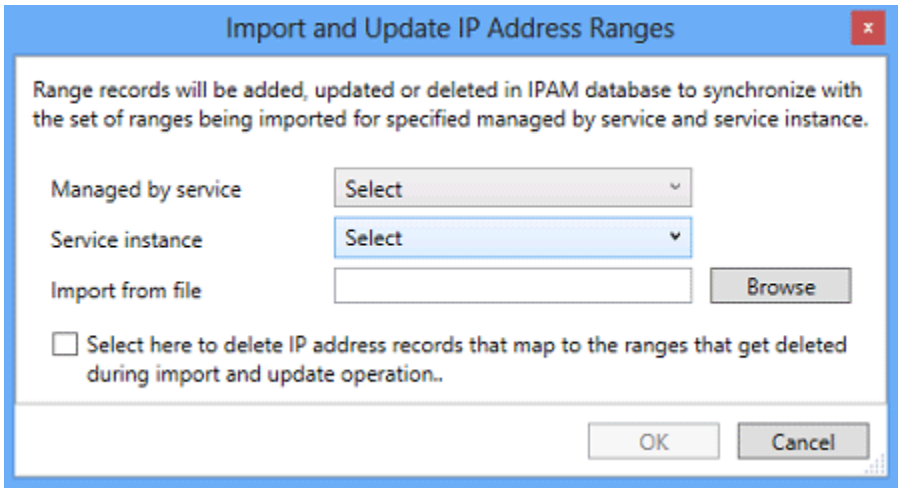

Export

IP address data that can be exported from IPAM includes IP addresses, IP address ranges, IP address blocks, IP address range groups, and IP address inventory. This data can include the results of custom data views and searches, or the entire unfiltered dataset. Data is exported to a text file in comma-separated-value (.csv) format.

Tip

Other types of data can also be exported from IPAM that might be used in IP address space operations, such as IP address tracking data. See IP Address Tracking for more information.

To export IP address data from IPAM, view the data that you wish to export, click **TASKS**, and then click **Export**. Type a file name, and then click **Save**.

Find and allocate available IP address

IPAM enables you to assign IP addresses based on availability. IPAM considers an IP address to be available if:

1. The IP address does not currently exist in the IPAM database.
2. The IP address is not reserved on the managed DHCP server providing the IP address range.
3. The IP address is not excluded on the managed DHCP server providing the IP address range.

4. The IP address does not respond to a ping request from the IPAM server.
5. A DNS PTR record is not found for the IP address.

To find an available IP address, right-click an IP address range and then click **Find and Allocate Available IP Address**.

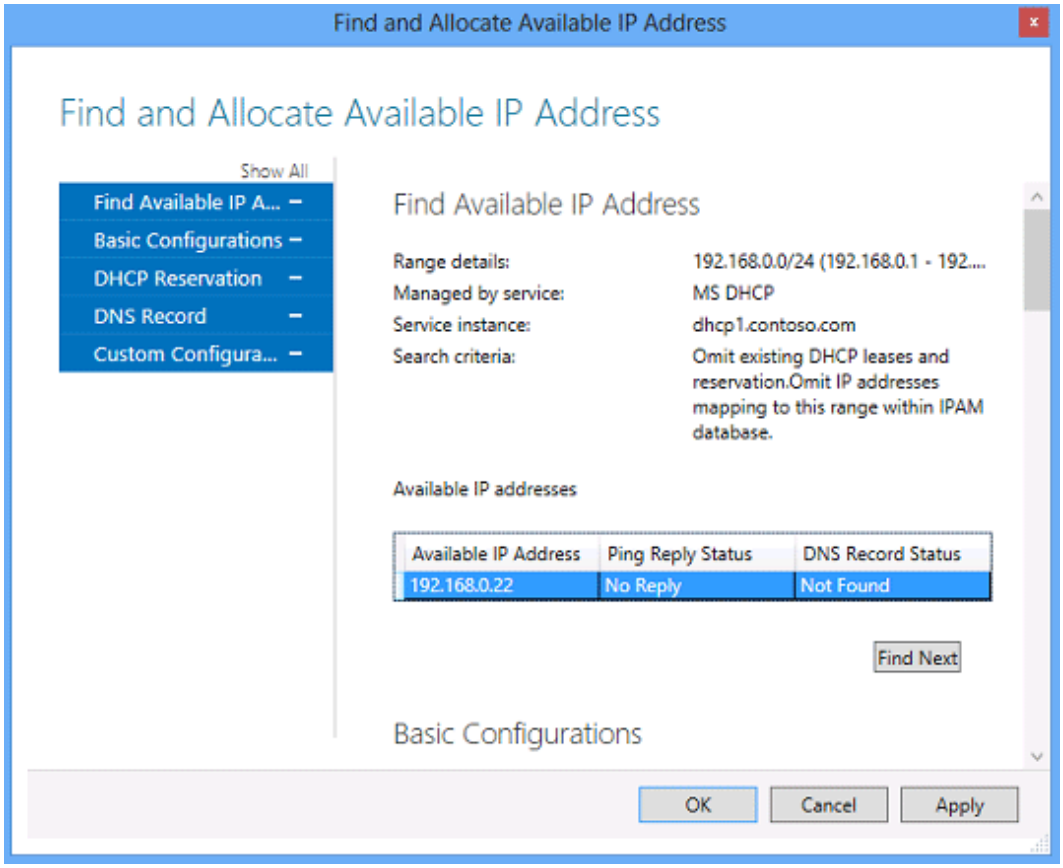

IPAM will search for an available address by starting with the first IP address that is not assigned in IPAM. If the address range being searched is a DHCP scope on a managed DHCP server, IPAM ignores all reserved or excluded IP addresses. When the ping request and DNS PTR query is complete, you can click **Find Next** to look for another available IP address. By default, the last IP address in the list is added to **Basic Configurations** with default settings. You can choose a different IP address from the list by clicking on it, or you can type an IP address manually next to **IP address**. If you type an IP address manually, it must be within a range provided by an available managed DHCP server in order to also create a DHCP reservation.

Optionally, you can complete The **DHCP Reservation** and **DNS Record** sections. Adding information to these sections only affects the IPAM database, it does not configure DHCP or DNS servers. IPAM database information is displayed in the IP address view, along with the DNS and DHCP synchronization status.

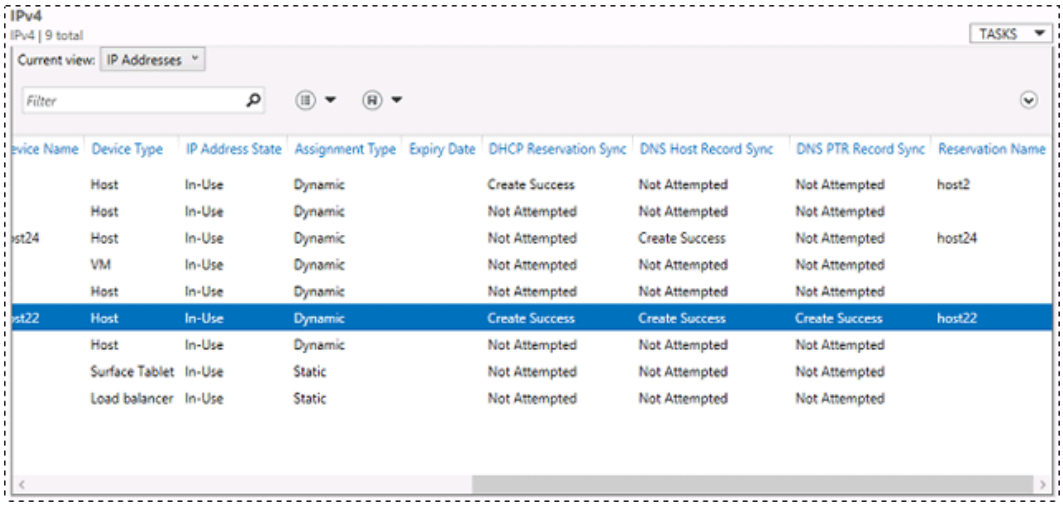

To add DHCP reservations to a DHCP server, you must complete the **Reservation server name**, **Reservation scope name**, and **Reservation name** fields in the **Edit IP Address** dialog box, or in the **Find and Allocate Available IP Address** dialog box. When this required information has been populated in the IPAM database, right-click one or more IP addresses in IP address view and then click **Create DHCP Reservation**. For more information, see Managing DHCP reservations.

To add DNS host records or DNS PTR records to a DNS server, you must complete the **Device name** field. **Forward lookup zone** and **Forward lookup primary server** fields are required for DNS host records. **Reverse lookup zone** and **Reverse lookup primary server** fields are required for DNS PTR records. When this required information has been populated in the IPAM database, right-click one or more IP addresses in IP address view and then click **Create DNS Host record** or **Create DNS PTR record**. For more information, see Managing DNS records.

Reclaim IP addresses

When you reclaim IP addresses in IPAM, they are deleted from the IPAM database. DHCP reservations and DNS records are not affected. If you also want to delete DHCP reservations and DNS records with the IPAM console, you must select one or more IP addresses using the IP addresses view, right-click, and then click **Delete DHCP Reservation**, **Delete DNS Host Record**, or **Delete DNS PTR Record**.

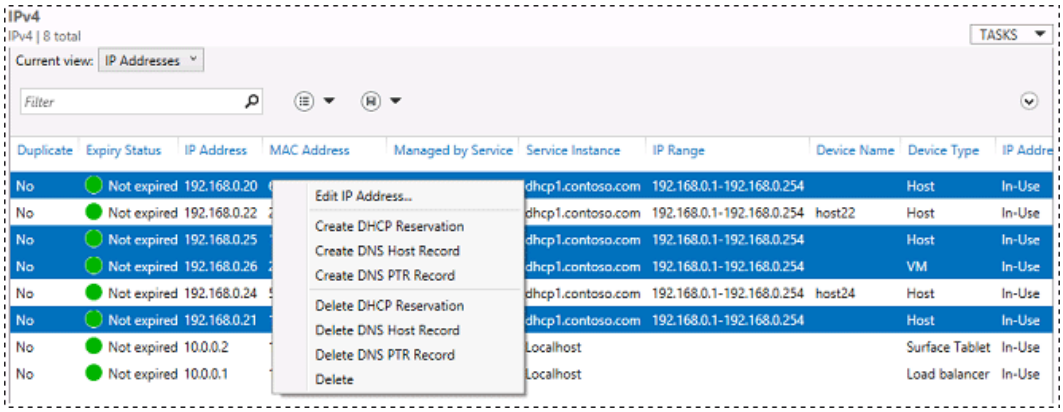

If you reclaim IP addresses or right-click and delete them before removing any DHCP reservations or DNS records that exist, you must delete these reservations and DNS records manually using the DNS or DHCP consoles, or the command line.

To reclaim IP addresses, right-click one or more IP address ranges and then click **Reclaim IP Addresses**.

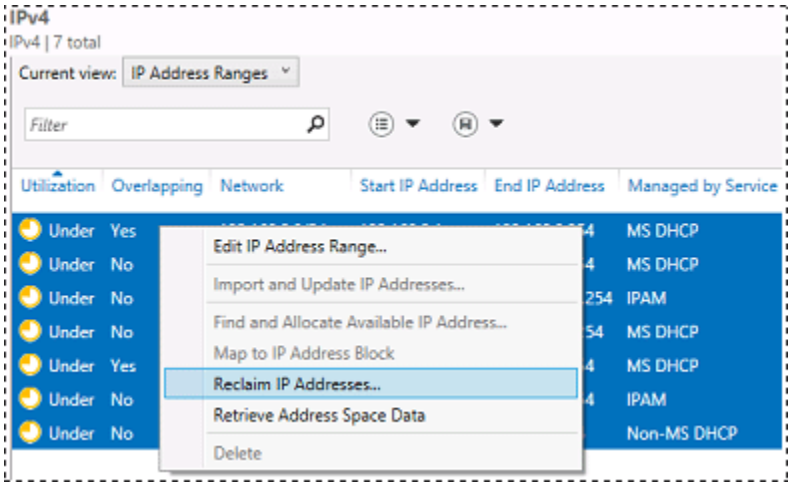

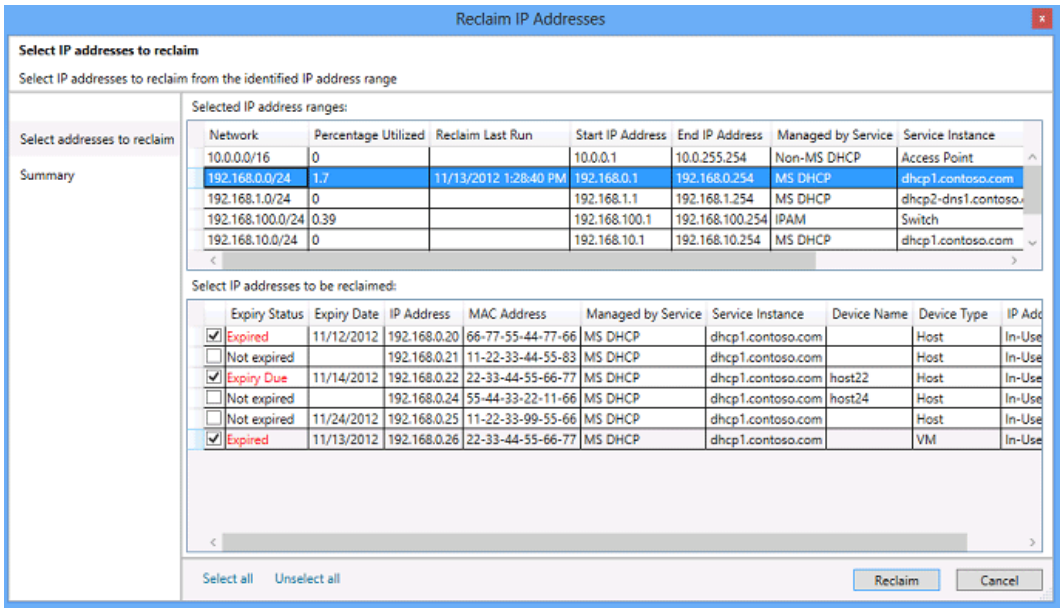

Click a range at the top and then select IP addresses from the bottom to reclaim. You can reclaim IP addresses from multiple IP address ranges simultaneously.

When you are finished selecting IP addresses, click **Reclaim**. A summary of IP addresses that were reclaimed is displayed.

Important

You cannot reclaim IP addresses that are not mapped to an IP address range. In order to map to a range, the IP address and the IP address range must have the same values for **Managed by Service** and **Service Instance**. You can review the status of IP address mapping to ranges by viewing the **IP Range** column in IP addresses view.

IP address expiry log settings

IP address expiry settings enable lifetime management for IP addresses on the network. To use expiry settings, an IP address must be configured with an expiry date. Expiry dates are optional.

To configure expiry log settings, verify that you have selected a child node under **IP ADDRESS SPACE**, click **TASKS**, and then click **IP Address Expiry Log Settings**.

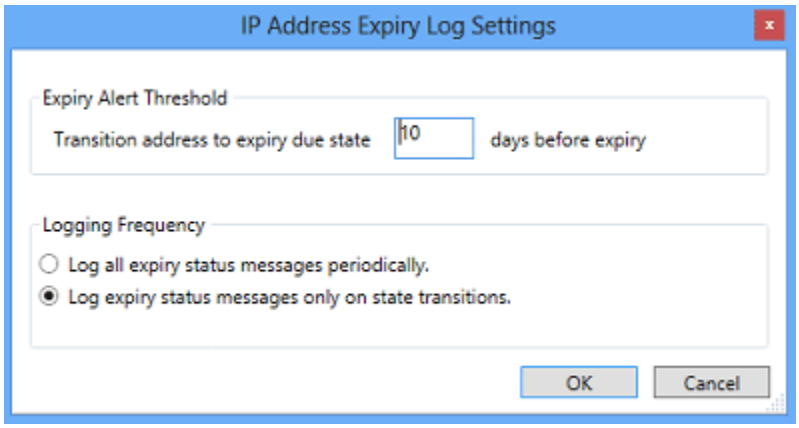

The three available IP address expiry states are: **Not expired**, **Expiry Due**, and **Expired**. By default, the expiry state of an IP address will transition from **Not expired** to **Expiry Due** 10 days before the expiry date. To configure a different transition time, type the number of days under **Expiry Alert Threshold**. You can also configure the **Logging Frequency**, which determines when the IPAM operational event log will receive expiry alerts. By default, events are generated only when address state transitions occur. You can also choose to log expiry status messages periodically. If you choose this setting, events are generated once every 24 hours for IP addresses with an expiry status of **Expiry Due** or **Expired**.

IP address inventory

You can view IP addresses in the IPAM database by clicking **IP Address Blocks** or by clicking **IP Address Inventory**. Using **IP Address Blocks**, you must choose **IP Addresses** next to **Current view**. **IPv4** addresses are displayed by default. Choose **IPv6** in the lower navigation pane to switch to the IPv6 addresses view. If you use **IP Address Inventory**, only IP addresses are displayed (ranges and blocks are not shown). You also must choose either IPv4 or IPv6 addresses in the lower navigation pane. The IP address blocks view organizes IP addresses based on whether they are mapped to a public or private IP address block. The IP address inventory view organizes IP addresses by device type, including custom IP address inventory groups. For more information about creating and using custom IP address inventory groups, see IP Address Inventory.

Adding IP addresses to the inventory

To add an IP address to the IPAM database using the IPAM client console, choose the **IP Address Blocks** or **IP Address Inventory** view, right-click **IPv4** or **IPv6** in the lower navigation pane, and then click **Add IP Address**. Alternatively, you can click **IPv4** or **IPv6**, click **TASKS**, and then click **Add IP Address**.

Tip
To specify the type of address to add (IPv4 or IPv6) you must use the lower navigation pane. By default, IPv4 is selected.

For information about adding IP addresses with Windows PowerShell, see Using Windows PowerShell with IPAM.

Basic configurations

Default values are assigned for all mandatory fields, with the exception of **IP address**. If you also choose to create a DHCP reservation for this IP address in IPAM, then the **MAC address** field is also required.

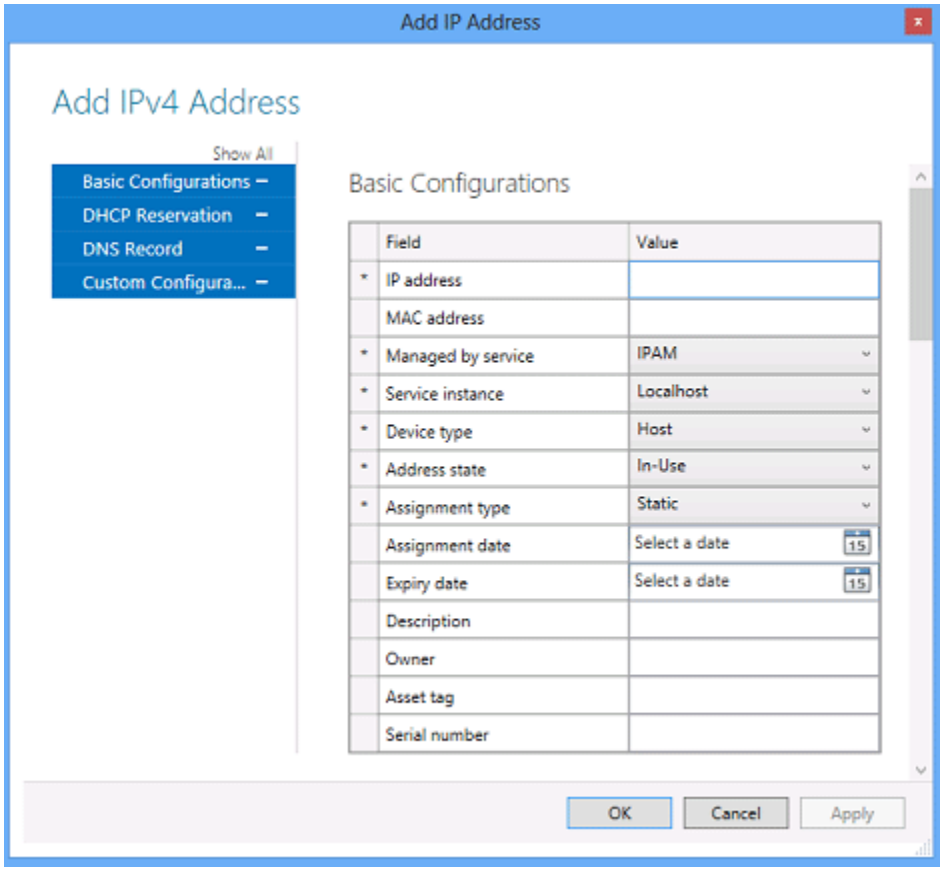

Basic configuration fields are identical for IPv4 and IPv6 addresses.

Managing DHCP reservations

To create an IPv4 address reservation in IPAM, you must first ensure that a valid IP address and MAC address are entered under **Basic Configurations**. Choose a **Managed by service** value of **MS DHCP** and a **Service instance** value corresponding to a valid managed DHCP server (ex: dhcp1.contoso.com). When these fields have been configured, click **DHCP Reservation** and choose a value for **Reservation server name** from the drop-down list. **Reservation scope name** and **Reservation scope detail** fields will be automatically populated. Enter required values for **Reservation name** and **Reservation type**. Reservation description is optional. Click **OK** to enter

the DHCP reservation into the IPAM database, or click **Apply** to commit the current values and continue entering more information.

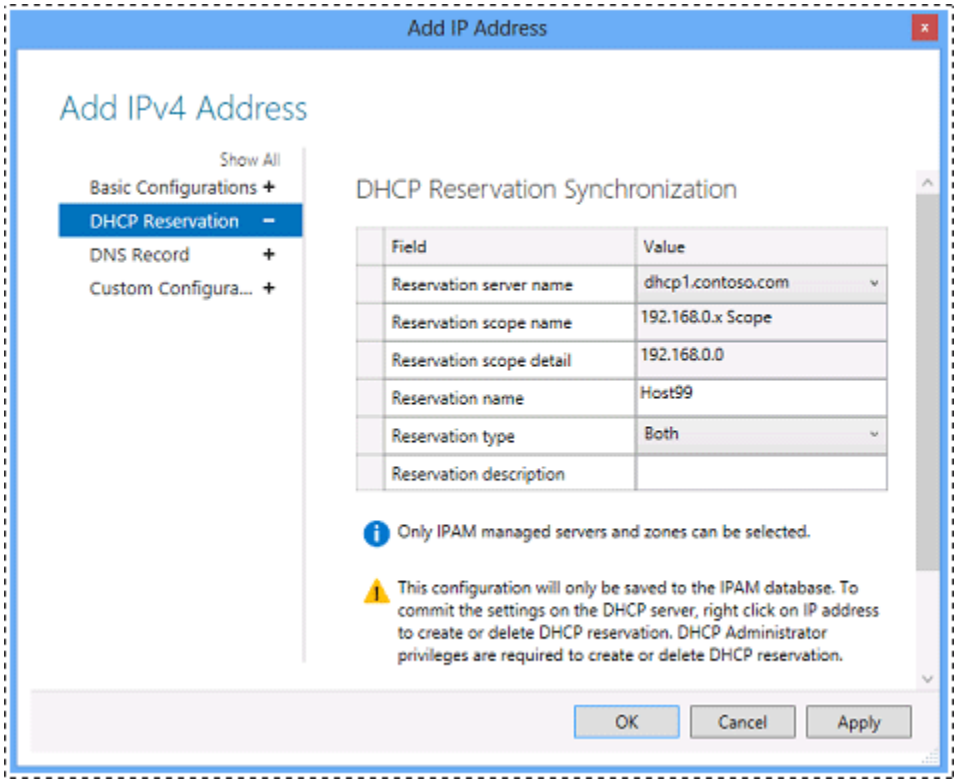

Important

Creating a DHCP reservation in the IPAM database does not automatically add this DHCP reservation to a managed DHCP server.

To create DHCP reservations on a managed DHCP server

1. Select one or more IP addresses and verify that each has the correct information in all required DHCP reservation fields.
2. Right-click the selected IP addresses, and then click **Create DHCP Reservation**.
3. Verify that **Create Success** is displayed under **DHCP Reservation Sync** in the display pane. Status is also displayed on the **Configuration Details** tab.

Warning

DHCP reservations created this way will replace an existing DHCP reservation, if it exists, for the same IP address.

To create an IPv6 address reservation in IPAM, you must first ensure that a valid IP address and MAC address are entered under **Basic Configurations**. Choose a **Managed by service** value of **MS DHCP** and a **Service instance** value corresponding to a valid managed DHCP server (ex:

dhcp1.contoso.com). When these fields have been configured, click **DHCP Reservation** and choose a value for **Reservation server name** from the drop-down list. **Reservation scope name** and **Reservation scope detail** fields will be automatically populated. Enter required values for **DUID**, **IAID**, and **Reservation name**. Reservation description is optional. Click **OK** to enter the DHCP reservation into the IPAM database, or click **Apply** to commit the current values and continue entering more information.

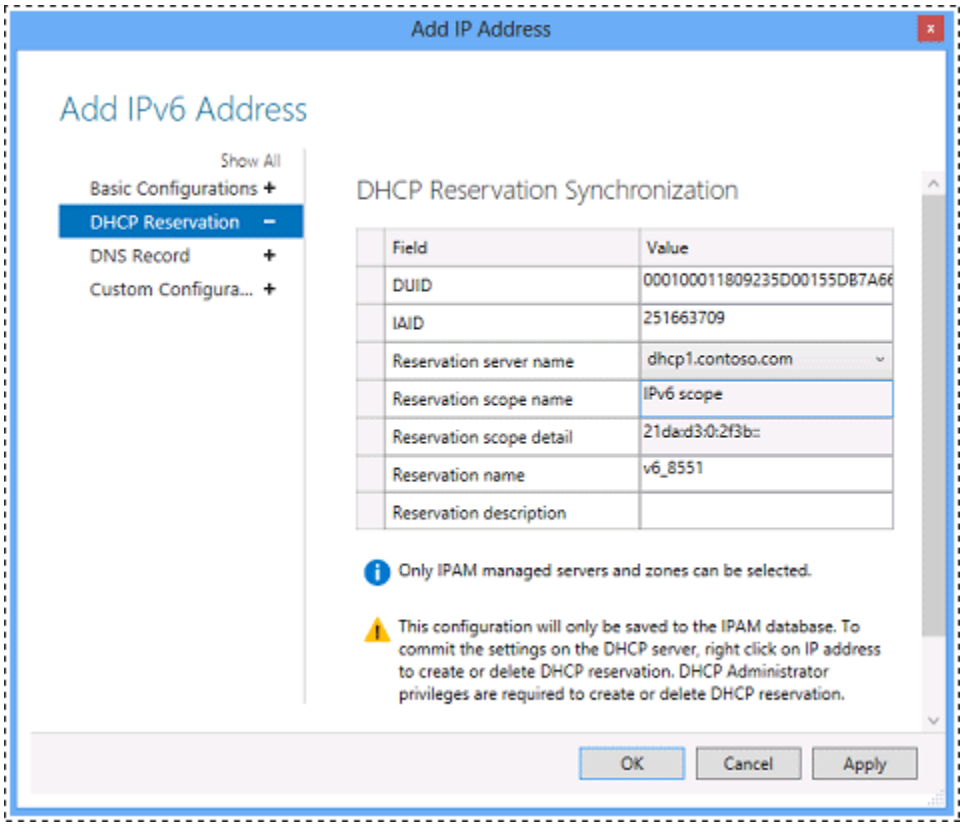

To complete creating the IPv6 DHCP reservation, add this reservation to the managed DHCP server using the same procedure as for IPv4 addresses.

Managing DNS records

To create DNS records in the IPAM database, you must ensure that an IP address is entered into the **Basic Configurations** dialog. This is the only basic configuration field that is required to add DNS records. Note: you cannot create a DNS PTR record if the IP address that is entered does not correspond to a reverse lookup zone hosted on a managed DNS server. For example, you can create a host record for the IP address 192.168.30.2 in any of the forward lookup zones hosted on a managed DNS server. However, the reverse lookup zone 30.168.192.in-addr.arpa must exist as a primary zone on a managed DNS server in order to create a PTR record for this IP address.

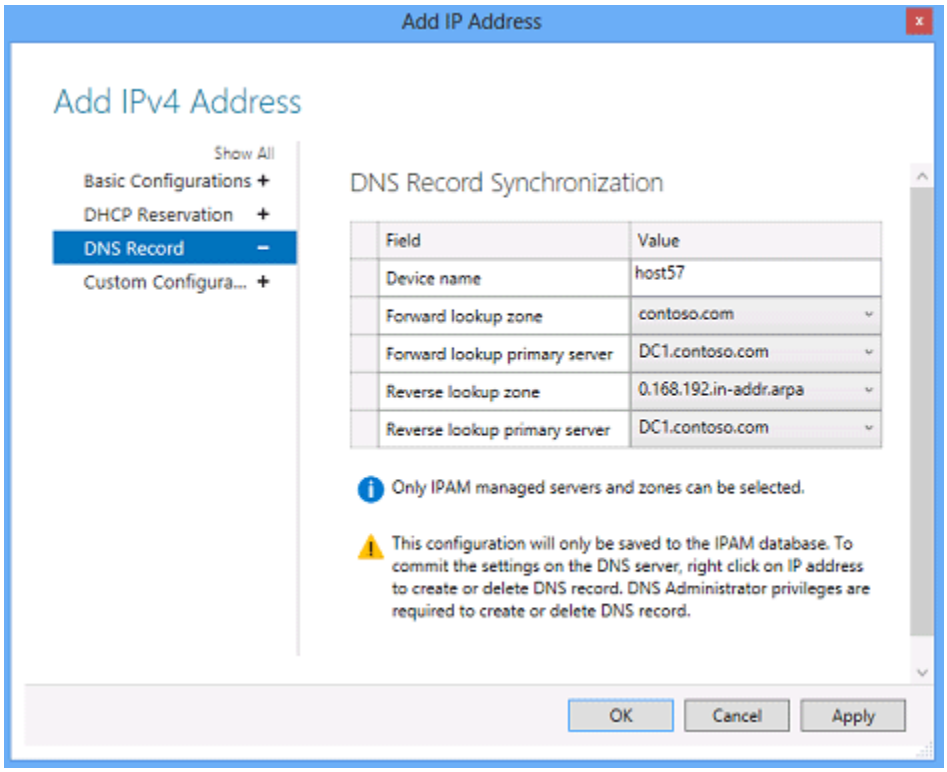

DNS configuration fields and requirements are identical for IPv4 and IPv6 addresses.

Important
Creating a DNS record in the IPAM database does not automatically add this DNS record to a managed DNS server.

To create DNS records on a managed DNS server

1. Select one or more IP addresses and verify that each has the correct information in all required DHCP record fields.
2. Right-click the selected IP addresses, and then click **Create DNS Host Record** to create a forward lookup record.
3. Right-click the selected IP addresses, and then click **Create DNS PTR Record** to create a reverse lookup record.
4. Verify that **Create Success** is displayed under **DNS Host Record Sync** or **DNS PTR Record Sync** in the display pane. Status is also displayed on the **Configuration Details** tab.

Custom configurations

Before you add custom configurations to IP addresses, ensure that you have configured custom fields. To configure custom fields, click **Manage** on the Server Manager menu bar, and then click

IPAM Settings. In the IPAM Settings dialog box, click **Configure Custom Fields**. In the **Configure Custom Fields** dialog box, choose a field to configure below **Custom Field Name** and then enter values under **Custom Field Value**. For example, you can click **AD Site** and then provide a list of Active Directory sites in your organization. Several built-in fields are available, and you can add custom field names by entering these manually under **Add custom fields below**.

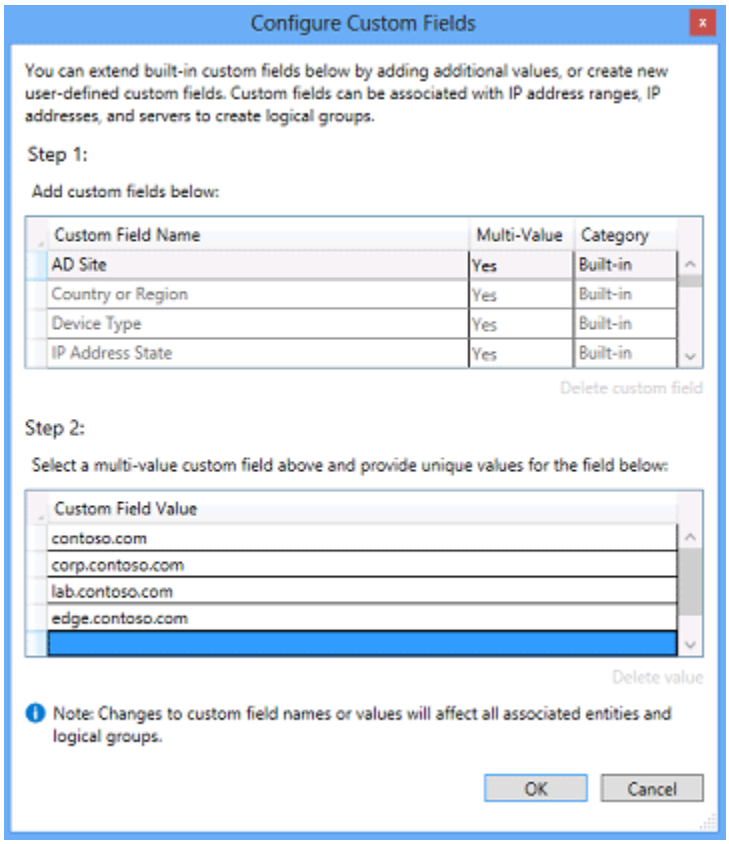

When you have completed entering custom fields and custom field values, you are ready to assign these field/value pairs to IP addresses using the custom configurations dialog. To add custom fields and values, choose a field next to **Custom field to configure**, choose a value next to **Specify a value**, and then click **Add**. To remove a field/value pair, click the row displaying the pair under **Configured fields** and then click **Delete**. Click **OK** when you are done.

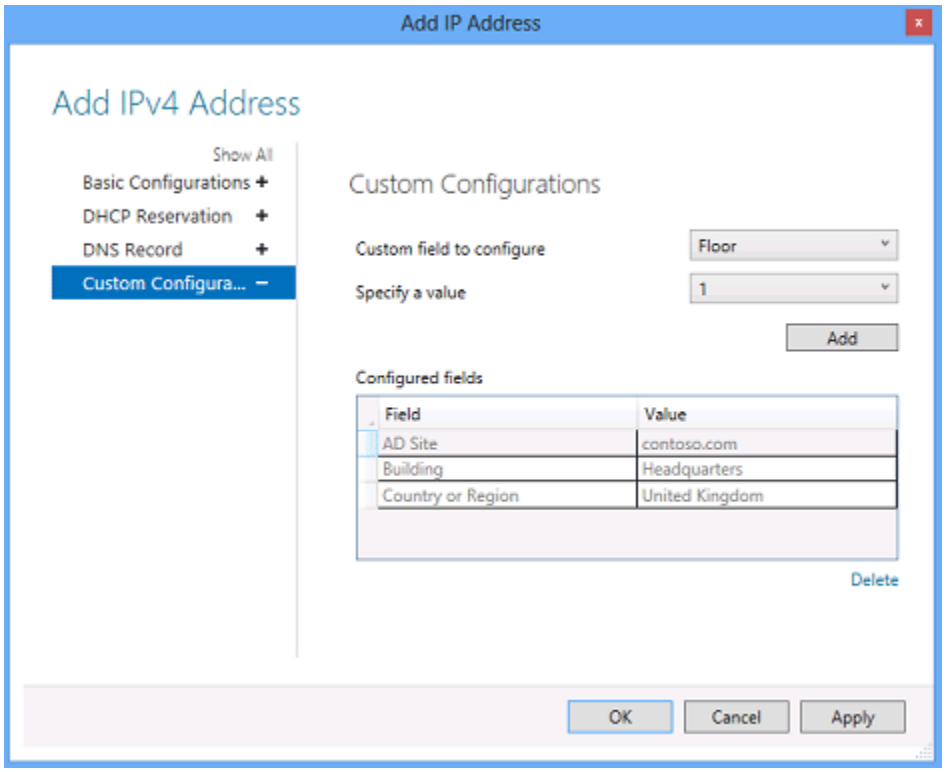

Importing DHCP lease data from the network

IPAM does not automatically add DHCP leases to the IPAM database as IP addresses. You can add these addresses using the IPAM client console, and you can automate importing of DHCP leases using Windows PowerShell. Use the following procedure to import DHCP leases.

To import DHCP leases

1. Download IpamIntegration_dhcp.ps1 from the Microsoft Script Center at Windows Server 2012 IPAM/DHCP integration module (http://gallery.technet.microsoft.com/scriptcenter/Windows-Server-2012-f44cefce).
2. Ensure that you have sufficient permissions to add IP addresses to the IPAM database. To enable this permission, you must be a member of the IPAM ASM Administrators group, or the IPAM Administrators group.
3. Launch an elevated Windows PowerShell console on a domain member computer.
4. Use the Set-Executionpolicy cmdlet to temporarily allow running unsigned scripts:

   ```
   Set-ExecutionPolicy Bypass -Force
   ```

5. Run the IPAM integration script to make Invoke-IpamDhcpLease available in your Windows PowerShell session.

   ```
   .\IpamIntegration_dhcp.ps1
   ```

6. Run Invoke-IpamDhcpLease and verify no errors are displayed. Replace the IPAM and

DHCP server names in the example below.

```
Invoke-IpamDhcpLease –IpamServerName ipam1.contoso.com –
DhcpServerFqdn dhcp1.contoso.com -Force
```

7. Verify that DHCP leases and reservations were successfully imported as IP addresses in the IPAM database.
8. If desired, create a scheduled task to import DHCP data periodically.

   Important
   The computer used to create the periodic task must be available on the network for the scheduled task to run. For this reason, it is recommended to create this task on the IPAM server or a DHCP server.
9. To create a scheduled task, run the Invoke-IpamDhcpLease with the –Periodic parameter.

   ```
   Invoke-IpamDhcpLease –IpamServerName ipam1.contoso.com –
   DhcpServerFqdn dhcp1.contoso.com -Periodic -Force
   ```
10. You will be prompted to enter Windows PowerShell credentials. Type **network service** next to **User name**, and then click **OK**.
11. By default, the task runs at the current time and again every six hours for one day. The default task name is **IpamDhcpIntegration**.
12. Monitor the task for one day, to ensure that records are imported successfully at an appropriate time and interval.
13. To modify the task to run indefinitely, or to change the time and interval, open Task Scheduler, click **Task Scheduler Library**, right-click the **IpamDhcpIntegration** task, and then click **Properties**.
14. Configure the task name on the **General** tab.
15. Configure task intervals on the **Triggers** tab. For example, to configure the task to run indefinitely, double-click the **Daily** trigger and then choose **Indefinitely** from the drop-down list next to **for a duration of**.

IP address range groups

Use the **IP Address Range Groups** node to organize IP address ranges into logical groups. You can use built-in logical groups and extend them by adding additional user-defined fields, or you can create new logical groups and fields. A pre-defined group called **Managed by Service** is available that organizes IP address ranges by this field and cannot be deleted.

To add an IP address range group, you must provide a name and a grouping criterion. The name can be anything, as long as it is unique. The grouping criterion is defined by custom fields. To add an IP address range group, click the **IP Address Range Groups** node in the upper navigation tree and click **TASKS** or right-click a node in the lower navigation tree, then click **Add IP Address Range Group**. Type a name for the group and then select grouping criterion from the list of fields at the bottom. An IP address range group must have at least one grouping criterion.

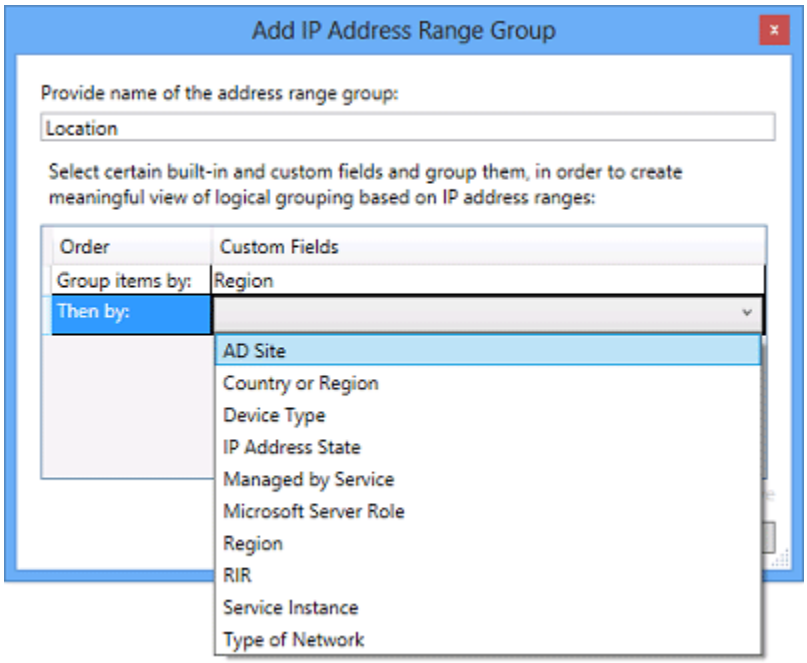

To use new custom fields or values with an IP address range group, you must first configure the custom fields. To define or edit custom fields, click **Manage**, click **IPAM Settings**, and then click **Configure Custom Fields**. Add new custom fields at the bottom of the list under **Step 1**. If a field has multiple values, select **Yes** under **Multi-Value** and then enter values under **Step 2**. Click **OK** when you are finished, and the new user-defined fields will now be available for use when creating IP address range groups.

To edit or delete a user-defined IP address range group, right-click the IP address range group in the lower navigation pane then then click **Edit IP Address Range Group**, or click **Delete**.

Tip

To edit user-defined fields that are assigned to IP addresses, IP address ranges, or servers in **Custom Configurations**, you must first remove these fields from the custom configurations to delete or modify them. To determine which servers or addresses are assigned custom fields, add the field as a display column in the display pane search view.

See also

[Walkthrough: Demonstrate IPAM in Windows Server 2012]

Multi-server Management

This topic contains detailed information and procedures you can use to manage Microsoft DNS and DHCP servers with IPAM. To review best practices for multi-server management, see Monitor and manage.

IPAM provides the following options for managing Microsoft DNS and DHCP servers on the network:

DNS and DHCP server monitoring

DNS zone monitoring

DNS and DHCP record synchronization

DHCP server management

DHCP scope management

DNS and DHCP server monitoring

Click the **DNS and DHCP Servers** node in the upper navigation pane to quickly view the status of managed DNS and DHCP services. Servers with IPv4 addresses are displayed by default. Use the lower navigation pane to display servers with IPv6 addresses, or to choose a specific IPv4 or IPv6 network. The **Server Availability** and **Duration in Current State** fields provide information about the health of DNS and DHCP services.

For example, if the DHCP Server service on a managed DHCP server is **Paused** or **Stopped**, this is displayed under **Server Availability**.

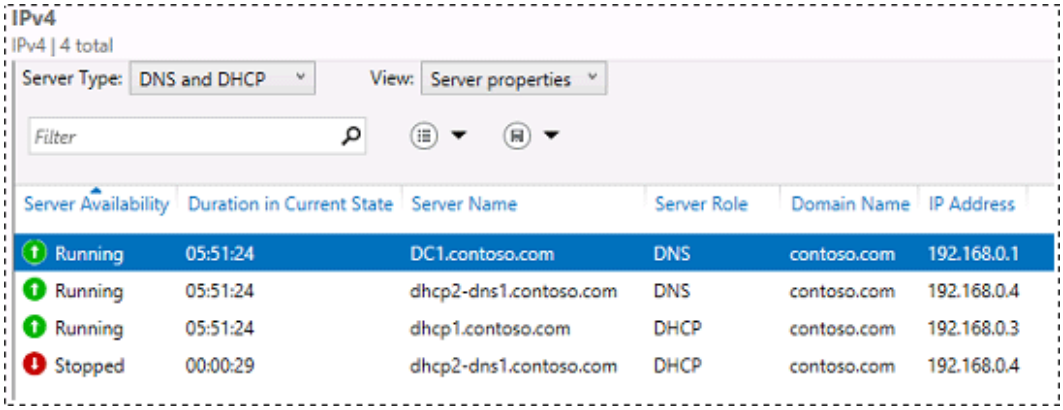

To immediately refresh the status of a server, right-click the server name, click **Retrieve Server Data**, wait for the **ServerAvailability** task to complete, and then refresh the console view.

To view when service status was most recently refreshed, view the **Last Refreshed** field in the display pane details view. You can also add this column to the upper display pane search view (see below). This column is not displayed by default.

If a server is not available due to network problems, it will display a status of **Not Reachable** for all services that the server provides.

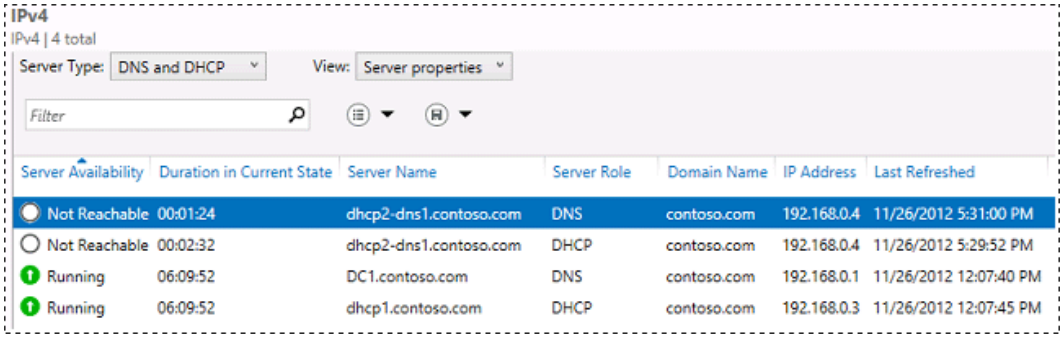

If only some services are not reachable, for example due to permission or firewall settings, then only those services will be **Not Reachable**.

DNS zone monitoring

Click the **DNS Zone Monitoring** node in the upper navigation pane to view a list of all zones hosted by managed DNS servers. Forward lookup zones are displayed by default. To view reverse DNS zones, click **IPv4 Reverse Lookup** or **IPv6 Reverse Lookup** in the lower navigation pane. The health of reverse lookup zones is not displayed. The lower navigation pane also displays a nested list of forward DNS zones with their associated status icon.

The health of a forward DNS zone is displayed in the **Zone Status** column next to the **Duration in Current State**. The zone status is a summary of status for all DNS servers that host the zone. Click the **Authoritative Servers** tab in the **Details View** to view a list of authoritative servers.

When a DNS zone is paused, the status of the zone will remain as **OK** because the DNS server will continue to respond to DNS queries for a paused zone. However, if the DNS Server service on one or more DNS servers is stopped, paused or not responding for another reason, zones hosted by the DNS server will display a status of **Error** or **Warning**.

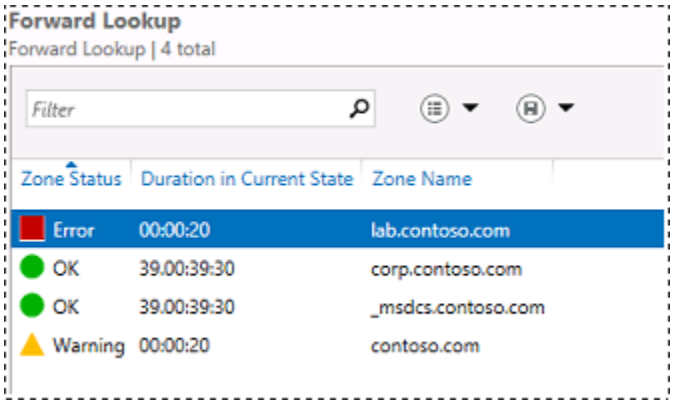

An **Error** status displays as a red square, and indicates that none of the authoritative DNS servers for the zone are responding to DNS queries.

A **Warning** status displays as a yellow triangle, and indicates that some of the authoritative DNS servers for the zone are not responding to DNS queries.

Zone status will update periodically with information from scheduled data collection tasks. To update the status of a zone immediately, right-click the zone, click **Reset Zone Status**, click **TASKS**, click **Retrieve Server Data**, and then refresh the console view.

DNS and DHCP record synchronization

You can use the **IP ADDRESS SPACE** node in the upper navigation pane to synchronize DNS and DHCP data in IPAM with data on managed DNS and DHCP servers. The following methods are available:

1. To import records from managed servers into the IPAM database, see Importing DHCP lease data from the network.
2. To manually create records in the IPAM database that corresponds to existing data on managed servers, unmanaged servers, and other IP address infrastructure devices, see Add or edit IP address space.
3. To use IP address data in IPAM to create records on managed servers, see Managing DHCP reservations and Managing DNS records.

DHCP server management

To manage DHCP server properties, click **DNS and DHCP Servers** in the upper navigation pane. Choose **DHCP** or **DNS and DHCP** from the drop-down list next to **Server Type**, and choose **Server properties** from the drop-down list next to **View**. Select one or more DHCP servers in the list, choose IPv4 or IPv6 in the lower navigation pane, and then right-click the selected servers to access the following DHCP server management options:

- **Edit DHCP Server Properties** configures the following options on one DHCP server or on multiple DHCP servers simultaneously:

  DNS dynamic update settings and name protection options. These options can also be configured on the **DNS** tab in the DHCP console.

  DNS dynamic update credentials used by the DHCP server to register names on DNS servers. This option can also be configured on the **Advanced** tab in the DHCP console.

  Enable or disable DHCP audit logging. This option can also be configured on the **General** tab in the DHCP console.

- **Edit DHCP Server Options** configures the following options on one DHCP server or on multiple DHCP servers simultaneously:

  Vendor class

  User class

  Option

Value

These options can also be configured by right-clicking **Server Options** under **IPv4** or **IPv6** in the DHCP console, and clicking **Configure Options**, and then clicking the **Advanced** tab.

- **Create DHCP Scope** adds a new DHCP scope to a single DHCP server. This option is not available when multiple DHCP servers are selected. To duplicate a DHCP scope, see DHCP scope management.
- **Configure Predefined DHCP Options** adds, edits, or deletes predefined options and values on one DHCP server or on multiple DHCP servers simultaneously. These options can also be configured by right-clicking **IPv4** or **IPv6** and then clicking **Set Predefined Options** in the DHCP console.
- **Launch MMC** opens the DHCP console to configure a single managed DHCP server. This option is not available when multiple DHCP servers are selected.
- **Configure User Class** creates, overwrites, or deletes DHCP user classes on one DHCP server or on multiple DHCP servers simultaneously. User classes can also be configured by right-clicking **IPv4** or **IPv6** and then clicking **Define User Classes** in the DHCP console.
- **Configure Vendor Class** creates, overwrites, or deletes DHCP vendor classes on one DHCP server or on multiple DHCP servers simultaneously. Vendor classes can also be configured by right-clicking **IPv4** or **IPv6** and then clicking **Define User Classes** in the DHCP console
- **Retrieve Server Data** runs the following IPAM data collection tasks immediately on the selected managed DHCP servers:

  ServerConfiguration

  ServiceMonitoring

  ServerAvailability

  Audit

DHCP scope management

To manage DHCP scope properties, click **DHCP Scopes** in the upper navigation pane. You can also choose **DNS and DHCP Servers** in the upper navigation pane and then choose **DHCP** from the drop-down list next to **Server Type**, and then choose **Scope properties** from the drop-down list next to **View**.

Tip

IPv4 and **IPv6** nodes in the lower navigation pane refer to IP addresses assigned to scopes if you are using the DHCP Scopes view. These nodes refer to interface addresses for managed DHCP servers if you are using the DNS and DHCP Servers view.

Select one or more scopes from the list, and then right-click the selected scopes to access the following DHCP scope management options:

- **Edit DHCP Scope** configures the following options on one DHCP server or on multiple DHCP servers simultaneously.

- **General Properties**:

  Scope name

  Description

  Start IP address

  End IP address

  Subnet mask

  Lease duration

  Exclusion range

  If multiple scopes are selected, you can only configure the description and lease duration.
- **DNS Dynamic Updates**:

  Enable DNS dynamic updates

  Enable name protection

  Dynamically update DNS A and PTR records

  Discard DNS A and PTR records when lease is deleted

  Dynamically update DNS A and PTR records for DHCP clients that do not request updates

  If multiple scopes are selected, you can only enable or disable DNS dynamic updates.
- **DHCP Scope Options**: Add, overwrite, find and replace, or delete vendor class and user class options and values.

  All DHCP scope options can be configured simultaneously for multiple scopes.
- **Advanced Properties**:

  Assign IP addresses dynamically to DHCP clients, BOOTP clients, or both

  Lease duration for BOOTP clients

  Subnet delay

  All advanced properties can be configured simultaneously for multiple scopes.

- **Duplicate DHCP Scope** creates a new scope on a single DHCP server that is already configured with the same DNS dynamic update options, scope options, and advanced properties of the scope that is selected for duplication. The new scope is created on the same DHCP server. You are required to provide a start IP address, end IP address, and subnet mask. DHCP failover relationships are not duplicated.

  Important

  The new DHCP scope will be active by default, even if the original scope was inactive.
- **Activate DHCP Scope** activates the selected scopes on managed DHCP servers. All scopes that you select must be inactive to use this option.
- **Deactivate DHCP Scope** deactivates the selected scopes on managed DHCP servers. All scopes that you select must be active to use this option.
- **Delete** deletes the selected scopes on managed DHCP servers.

See also

Using the IPAM Client Console

IP Address Tracking

IPAM enables IP address tracking through correlation of DHCP lease events on managed DHCP servers with user and machine authentication events on managed domain controllers and NPS servers.

IPAM supports searching IP address records using four parameters:

Tracking by IP address

Tracking by client ID

Tracking by host name

Tracking by user name

Searches using wildcards or regular expressions are not supported. Future dates are allowed when entering a date range to search. You can also choose to include or exclude user and machine logon events, and additional correlated events. If you exclude logon and correlated events, only DHCP lease data is searched. Search results can be exported to a text file in comma-separated-value (.csv) format. You can also save and reuse queries.

To search the most recent lease and logon events, run IPAM data collection tasks on managed servers before performing a search of event data.

Tracking by IP address

To search by IP address, click **By IP Address** and enter a complete IPv4 address and valid date range. IPv6 address tracking is not supported.

To track by IP address

1. Click **EVENT CATALOG** in the upper navigation pane.
2. Under **IP Address tracking** in the lower navigation pane, click **By IP Address**.
3. Type an IPv4 address next to **Search**.
4. Type a start and end date next to **and DHCP lease events between these dates**.
5. Click **include** or **exclude** to view or not to view events, respectively, from domain controllers and NPS servers.
6. Click **Search**.

For example, a search by IP address might reveal that a DHCP client received a DHCP lease from two different DHCP servers during a defined time frame, perhaps due to a DHCP failover relationship between the two DHCP servers. This could be verified by reviewing the **Server**

Availability and **Duration in Current State** fields in the **DNS and DHCP Servers** node under **MONITOR AND MANAGE**.

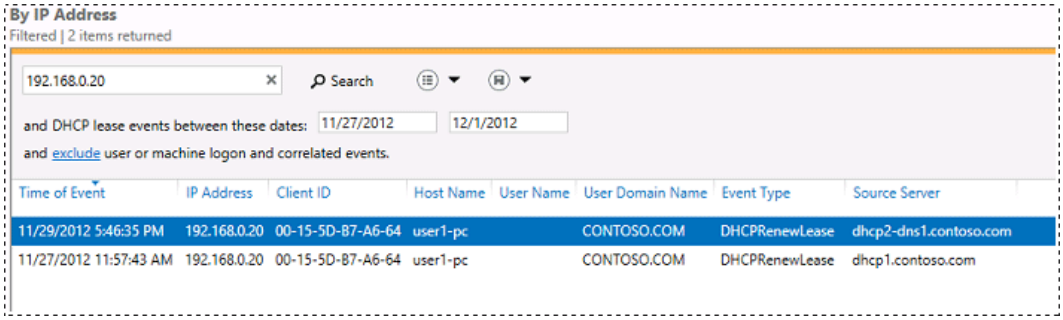

Tracking by client ID

To search by client ID (MAC address), click **By Client ID** and enter a client ID (MAC address) and valid date range. Dashes are allowed in the client ID, but are not required.

To track by client ID

1. Click **EVENT CATALOG** in the upper navigation pane.
2. Under **IP Address tracking** in the lower navigation pane, click **By Client ID**.
3. Type a client ID next to **Search**.
4. Type a start and end date next to **and DHCP lease events between these dates**.
5. Click **include** or **exclude** to view or not to view events, respectively, from domain controllers and NPS servers.
6. Click **Search**.

For example, a search by client ID can reveal different IP addresses that have been assigned to a specific device on the network during a defined time frame.

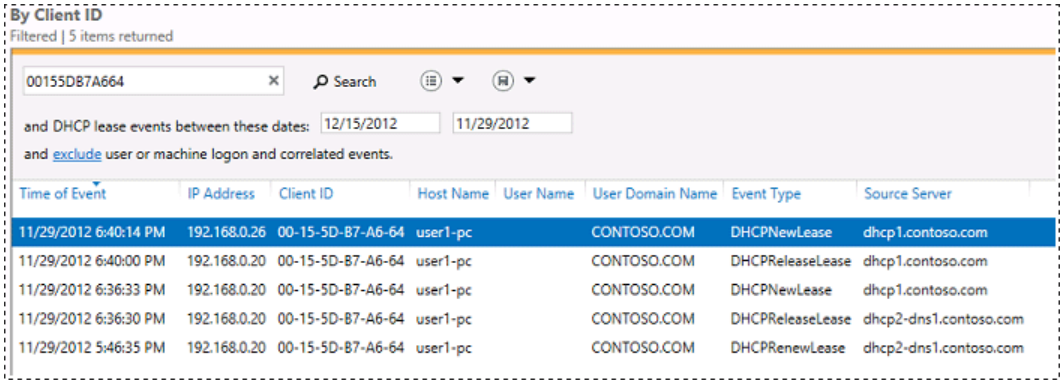

Tracking by host name

To search by host name, click **By Host Name** and enter a host name and valid date range. Case is ignored for host names.

To track by host name

1. Click **EVENT CATALOG** in the upper navigation pane.
2. Under **IP Address tracking** in the lower navigation pane, click **By Host Name**.
3. Type a host name next to **Search**.
4. Type a start and end date next to **and DHCP lease events between these dates**.
5. Click **include** or **exclude** to view or not to view events, respectively, from domain controllers and NPS servers.
6. Click **Search**.

For example, a search by host name can reveal the users that have signed into a specific device on the network during a defined time frame. To more easily view user data, sort by the **User Name** field.

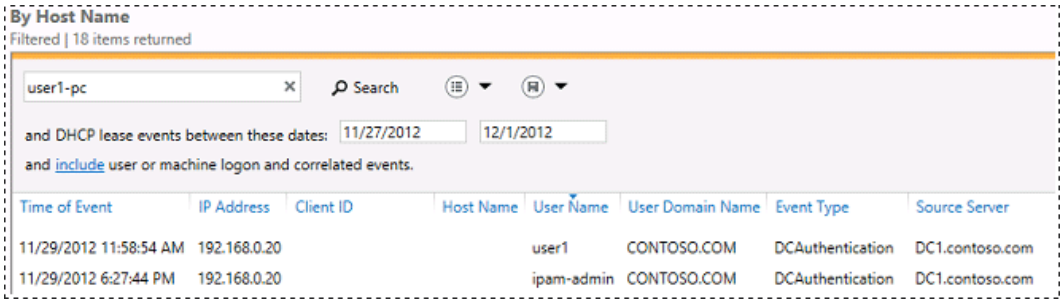

Tracking by user name

To search by user name, click **By User Name** and enter a host name and valid date range. Case is ignored for user names. Since user names are not included in DHCP lease events, you must include logon events in the search.

To track by user name

1. Click **EVENT CATALOG** in the upper navigation pane.
2. Under **IP Address tracking** in the lower navigation pane, click **By User Name**.
3. Type a host name next to **Search**.
4. Type a start and end date next to **and DHCP lease events between these dates**.
5. Click **Search**.

For example, a search by user name can reveal when the domain Administrator account was used on the network.

See also

Operational Event Tracking

IPAM Audit Logic

Operational Event Tracking

This topic contains detailed information and procedures you can use to review DHCP server and IPAM server operational events in IPAM. To review best practices for managing the event catalog in IPAM, see Event catalog.

You can use the **EVENT CATALOG** node in the upper navigation pane to monitor configuration events on managed DHCP servers and the IPAM server. Configuration event monitoring is not available for DNS servers.

DHCP configuration events

Click **DHCP Configuration Events** in the lower navigation pane to display events from all managed DHCP servers. Events are displayed from the **Operational** event log. These events can also be viewed in Event Viewer on individual DHCP servers by navigating to **Applications and Services Logs>Microsoft>Windows>DHCP-Server>Microsoft-Windows-DHCP Server Events/Operational**.

Events are retrieved from managed DHCP servers by the IPAM **Audit** task, which runs every 24 hours by default. To retrieve events immediately, click **TASKS** and then click **Retrieve Event Catalog Data**. To search event log data, type a search term next to **Search** and press ENTER. Searching with wildcards or regular expressions is not supported, but you can search for part of a name in the server FQDN.

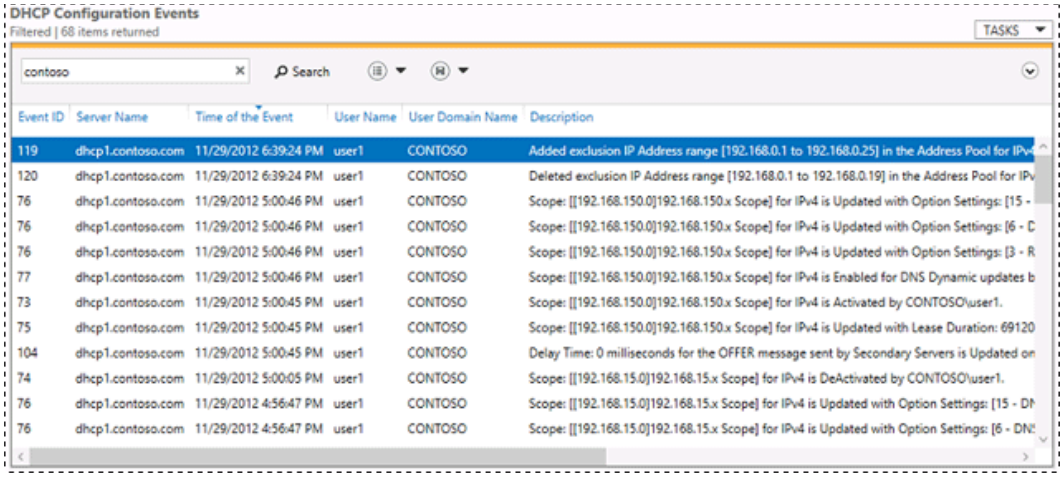

You can also use several different criteria to limit the scope of search results. To add criteria, click the expand icon in the right-hand corner of the DHCP Configuration Events window and choose from the following criteria:

- Event ID
- Server Name
- Time of the Event
- User Name
- User Domain Name
- Scope Name (in description)
- Scope ID (in description)
- Option Number (in description)
- Option Name (in description)
- Reservation Address (in description)

To purge event log data, click **TASKS** and then click **Purge Event Catalog Data**. The Purge Event Catalog Data dialog box allows you to choose the type of events to purge and the age of events to purge.

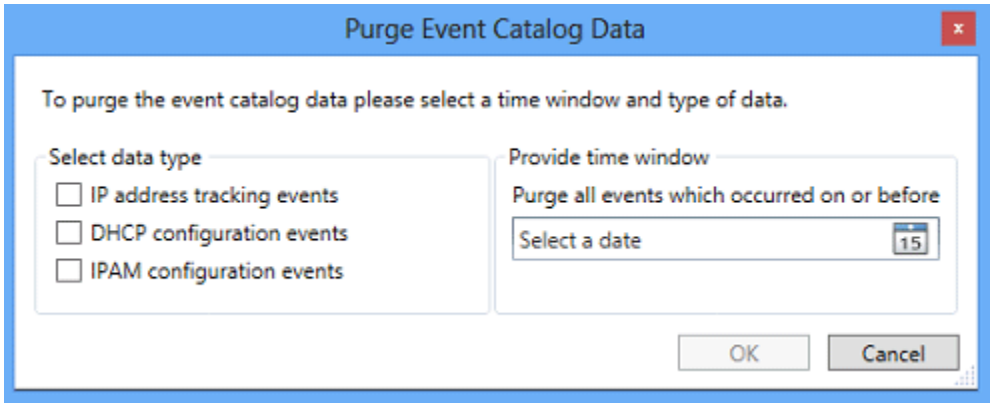

After choosing the data type and time window to use for purging data, click **OK** and then click **Yes** to confirm data deletion. After the IPAM Purge Event Catalog task completes, refresh the console view to only events that were not purged.

Tip

Purging events from the IPAM database does not clear events from Event Viewer on managed DHCP servers. However, subsequent event catalog data retrieval only adds newly generated events in IPAM. Data that was purged is not re-introduced.

IPAM configuration events

Click **IPAM Configuration Events** in the lower navigation pane to display events from the current IPAM server. Events are displayed from the **ConfigurationChange** event log. These events can also be viewed in Event Viewer on the IPAM server by navigating to **Applications and Services Logs>Microsoft>Windows>IPAM>ConfigurationChange**.

You cannot add criteria when searching IPAM configuration events.

Retrieval and purging of event catalog data is the same for IPAM configuration events as for DHCP configuration events.

See also

IP Address Tracking

Using Windows PowerShell with IPAM

This topic provides a summary of Windows PowerShell IPAM Server cmdlets available in Windows Server® 2012. For more information about these cmdlets, including examples and parameter descriptions, see IPAM Server Cmdlets in Windows PowerShell (http://technet.microsoft.com/library/jj553807.aspx).

IpamGpoProvisioning

Use this cmdlet to create and link IPAM provisioning Group Policy Objects (GPOs) in the specified domain. These GPOs can be used to configure required IPAM access settings on managed domain controllers, NPS servers, DNS servers, and DHCP servers. For more information about using this cmdlet, see Provisioning IPAM. The following IpamGpoProvisioning cmdlets are available: **Invoke-IpamGpoProvisioning**.

```
    Invoke-IpamGpoProvisioning [-Domain] <string> [-GpoPrefixName] <string> [-
IpamServerFqdn <string>] [-DelegatedGpoUser <string[]>]
    [-DelegatedGpoGroup <string[]>] [-DomainController <string>] [-PassThru] [-Force] [-
WhatIf] [-Confirm]  [<CommonParameters>]
```

IpamCustomValue

Use this cmdlet to add a new custom value to a multi-valued, custom field in the IPAM database. For more information about custom values in IPAM, see Custom configurations. The following IpamCustomValue cmdlets are available: **Add-IpamCustomValue**.

```
    Add-IpamCustomValue [-Name] <string> [-Value] <string> [-PassThru] [-WhatIf] [-
Confirm]  [<CommonParameters>]
```

IpamAddress

Use this cmdlet set to import and export IP address data to and from the IPAM database. Imported data must be formatted as comma separated values (.csv). Exported IP address data can also include an array of Windows PowerShell® objects. For more information about importing and exporting IP address data, see Import and Export. The following IpamAddress cmdlets are available: **Import-IpamAddress**, **Export-IpamAddress**.

```
    Import-IpamAddress [-Path] <string> [-AddressFamily] <string> {IPv4 | IPv6} [-
ErrorPath <string>] [-Force] [-WhatIf] [-Confirm]
    [<CommonParameters>]
```

```
    Import-IpamAddress [-Path] <string> [-AddressFamily] <string> {IPv4 | IPv6} [-
ManagedByService] <string> [-ServiceInstance] <string>
    [-NetworkId] <string> [-StartIPAddress <ipaddress>] [-EndIPAddress <ipaddress>] [-
ErrorPath <string>] [-Force] [-WhatIf] [-Confirm]
    [<CommonParameters>]
```

```
    Export-IpamAddress [-AddressFamily] <string> {IPv4 | IPv6} [-Path <string>] [-
PassThru] [-Force]  [<CommonParameters>]
```

IpamCustomField

Use this cmdlet set to display or add custom field information to the IPAM database. For more information about custom fields in IPAM, see Custom configurations. The following IpamCustomField cmdlets are available: **Get-IpamCustomField**, **Add-IpamCustomField**.

```
    Get-IpamCustomField [[-Name] <string>]  [<CommonParameters>]
```

```
    Add-IpamCustomField [-Name] <string> [-Multivalue] [-PassThru] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

IpamRange

Use this cmdlet set to import and export IP address range data to and from the IPAM database. Imported data must be formatted as comma separated values (.csv). Exported IP address data can also include an array of Windows PowerShell® objects. For more information about importing and exporting IP address data, see Import and Export. The following IpamRange cmdlets are available: **Import-IpamRange**, **Export-IpamRange**.

```
    Import-IpamRange [-Path] <string> [-AddressFamily] <string> {IPv4 | IPv6} [-ErrorPath
<string>] [-Force] [-WhatIf] [-Confirm]
    [<CommonParameters>]

    Import-IpamRange [-Path] <string> [-AddressFamily] <string> {IPv4 | IPv6} [-
ManagedByService] <string> [-ServiceInstance] <string>
    [-AddManagedByService] [-AddServiceInstance] [-DeleteMappedAddresses] [-ErrorPath
<string>] [-Force] [-WhatIf] [-Confirm]
    [<CommonParameters>]

    Export-IpamRange [-AddressFamily] <string> {IPv4 | IPv6} [-Path <string>] [-PassThru]
[-Force]  [<CommonParameters>]
```

IpamConfiguration

Use the cmdlet set to display or configure the TCP port number used by the IPAM server to communicate with an IPAM client running Remote Server Administration Tools (RSAT). The IPAM version is also displayed. For more information about IPAM client communication, see IPAM Client. The following IpamConfiguration cmdlets are available: **Get-IpamConfiguration**, **Set-IpamConfiguration**.

```
    Get-IpamConfiguration  [<CommonParameters>]

    Set-IpamConfiguration [-Port] <uint16> [-Force] [-PassThru] [-WhatIf] [-Confirm]
[<CommonParameters>]
```

See also

Importing DHCP lease data from the network

Best Practices

IPAM best practices

The following are some best practices for deploying and operating IPAM.

General

Provisioning

Discovery

IP address management

Monitor and manage

Event catalog

General

- IPAM users and administrators should be added to the appropriate IPAM security groups based on their roles and administrative privileges. Do not make all users a member of the IPAM Administrators group unless this is required.
- User-defined custom fields that are required to be used for a logical grouping must be added as multivalue fields.
- Refresh the IPAM user interface periodically to ensure that the view is up-to-date with any changes in the database. Changes can occur due to periodic data collection tasks and concurrent user sessions.
- Leverage advanced features like logical grouping, saving and loading filter queries, group by, and viewing customizations to visualize and track entities and fields of interest.
- The IPAM database should be backed up periodically using Windows Server Backup.
- Tasks that require a long time to complete and operations such as database purging should be scheduled during non-business hours when the load on the network is low.

Provisioning

- The Group Policy based provisioning method is recommended over the manual provisioning method due to the number of security settings that need to be enabled for each server role that is managed by IPAM. Using Group Policy also ensures that settings are regularly updated.
- A domain-wide group named **IpamGpoAdmins** should be created, and IPAM GPO administration should be delegated to this group at the time of GPO creation. Any user who needs to mark servers as managed or unmanaged can be added to this group.
- Ensure that there are no conflicting Group Policies in your environment for the settings that are required by IPAM. IPAM access can be disabled if a firewall port that is required by IPAM is blocked by another Group Policy, or if membership in a security group that is required by IPAM is overwritten.

- When multiple servers are being selected for management, such as during initial IPAM provisioning, select all servers simultaneously to set their manageability status. This ensures that incremental edits to IPAM GPOs are minimal.
- For large deployments, linking the GPO to the root of the domain may not be ideal. Administrators should use organizational units to deploy Group Policy for infrastructure servers.
- Rollback of IPAM access settings is recommended from the servers that are being marked as unmanaged. However, updating the domain-wide DNS ACL will enable DNS RPC access for the IPAM server on all domain controllers that are co-located with DNS servers in the domain. This setting is not updated on a per-server basis.
- For standalone DNS servers that are not co-located on domain controller, the computer for the IPAM server must be added to the local Administrators group because the domain-wide DNS ACL will not be applicable on these servers. If this is not done, the standalone DNS server will display an IPAM access status of blocked.

Discovery

- Do not install the IPAM feature on a server that is running the on a DHCP Server role service because it affects IPAM's discovery of DHCP servers in the network.
- Do not add servers individually (for roles other than NPS), to avoid having to manually maintain the server inventory. Leverage the automated discovery mechanism as much as possible.
- Avoid keeping the server manageability status as **Unspecified** except when the server is temporarily out of service. Ideally the server should be marked as **Managed** or **Unmanaged**.
- If only a subset of the roles that are discovered on a physical server need to be managed, make sure that only the roles that are actively managed are selected on the server. Other roles should be cleared. IPAM periodic data collection tasks will collect data from all roles that are selected.

IP address management

- Configure utilization thresholds appropriately to select the most effective checkpoints to identify under and over utilized address space.
- Over utilization threshold crossing events that are logged by IPAM should be monitored for timely detection of potential address space exhaustion.
- Overutilized and underutilized address space must be identified and effectively redistributed to achieve optimal utilization across the entire address space. Utilization trends can be used as effective tools to predict the short and long term address space requirements.
- If exclusion ranges are being used within DHCP scopes to make static allocations or reservations, then IPAM recommends explicitly creating a static IP address range that corresponds to the exclusion range to track the utilization statistics for the entire subnet.
- For a homogenous and consistent rollout of address space allocations, arrange IPv4 address ranges in size units of between /22 and /24 and IPv6 ranges of /64. Use IPv4 address blocks of size /16, and IPv6 address blocks of /48.

- Create static IP addresses with a finite expiry date to ensure that stale addresses can be reclaimed from address ranges at the correct intervals. When an IP address transitions to the expiry due state, administrators must try to ensure that the address is alive in the network and extend the address life or reclaim the address.
- Leverage IPAM's overlap and duplicate detection feature to identify and track intentional and unintentional IP address ranges on your network.
- If you have multiple instances of overlapping address space configured in your environment, then using logical groups will be more effective to track overall utilization. Address blocks only map one overlapping range for utilization roll-up.
- For tracking IP address leases in the network, the IP address tracking feature in the Event Catalog view can be more efficient than pulling all dynamic DHCP lease records into IPAM.
- Leverage IPAM's Find an Available IP Address feature with in-line DHCP reservation and DNS record updates to simplify static IP address allocations.

Monitor and manage

- Create and save queries to quickly identify services and zones that are not in a healthy state.
- Leverage advanced constructs in multiedit scenarios to add, overwrite, delete, or find-and-replace to cater to the exact scenario requirement when you edit options across multiple DHCP scopes and servers.
- Leverage advanced constructs in multiedit scenarios to add, overwrite, or delete to cater to the exact scenario requirement when you configure vendor and user classes across multiple DHCP servers.
- Use the duplicate scope functionality to create new scopes with similar settings. The typical scenarios where this can be leveraged are migrating scopes from one server to another and configuring split scopes.
- Use the overall forward lookup zone view to identify potential issues and to identify servers that might have a problem. Isolate whether the issue is due to a zone event or server availability state.
- Monitor DHCP scope utilization percentage and utilization status to identify overutilized and underutilized scopes. Take necessary actions to align scope utilization to an optimal state, keeping utilization trend history in mind.
- For settings that are not supported by IPAM, launch the DHCP or DNS MMC from within the IPAM console to complete the configuration scenario.

Event catalog

- For best performance and disk space management, an IPAM event catalog data purge should be performed periodically, to maintain a six-month old database.
- Auditing of account logon events should be enabled on domain controllers and NPS servers, and the security event log size should be large enough to allow the periodic audit task to complete data collection before it is rolled over.
- The DHCP audit log file size must be large enough to hold the audit events for one day, to ensure that no lease events are lost because of size overruns.

- The audit log file location for both DHCP IPv4 and IPv6 leases must be configured in a common folder. The IPAM audit task selects the log files from one network share per server.
- The time period that is selected for a query should be optimal. Typically, a query interval between 3 days to 15 days is optimal.

See also

Troubleshooting IPAM

Troubleshooting IPAM

See the following sections in this topic.

Event logs

Common problems and solutions

- Connecting to the IPAM server
- Server access
- Server discovery
- IP address space management
- Monitoring and management

Event logs

IPAM event logs are available in Event Viewer using the following path: Application and Services Logs > Microsoft > Windows > IPAM. The following log channels are available:

- **Admin**: This channel captures events that are related to IPAM user actions and IPAM periodic tasks.
- **ConfigurationChange**: This channel captures events that are related to configuration changes made to the IPAM server.
- **Operational**: This channel captures informational events that are related to IPAM periodic tasks. Event logging on this channel is disabled by default.
- **Analytic and Debug**: These channels are hidden and disabled by default. Use these channels for debugging purposes. To show them, click **View** and then click **Show Analytic and Debug Logs**.

Events in the Admin and Operational channels can also be viewed from the IPAM tile in the Server Manager Dashboard.

Common problems and solutions

The following section lists common problems and solutions for IPAM.

Connecting to the IPAM server

- Problem: You are not able to connect to the IPAM server.
- Solution: View the **Description** field on Task Manager's **Services** tab, and verify that the following services are running on the IPAM server:
  - Windows Process Activation Service
  - Windows Internal Database

  If they are not running, start these services. Also verify that you are a signed in as a domain user and that you are a member of either the local Administrators security group on the IPAM server, or a member of the appropriate IPAM security group on the IPAM server.

Server access

- Problem: The access status of a managed server does not change to Unblocked after completing IPAM provisioning.
- Solution: Review the details pane in Server Inventory view to identify the types of access status that are blocked. If you are using the manual IPAM provisioning method, verify that the access settings that correspond to this status are enabled. For Group Policy-based provisioning, verify that the required GPOs have been created and linked, and that the managed server is a member of the security filter list of appropriate GPOs. This issue can also occur if Group Policy has not been updated on the target servers. Open an elevated command prompt on the managed server and type **gpupdate /force** to force the application of Group Policy.
- Problem: You are unable to add a managed server to an IPAM GPO.
- Solution: Verify that you are signed in as a member of the Domain Admins security group, or higher, or that you have been delegated privileges to edit the GPO. GPO delegation for users or groups can be enabled from the IPAM **Invoke-IpamGpoProvisioning** cmdlet at the time the GPO is created, or from the Group Policy Management Console if the GPO is already created.

Server discovery

- Problem: You are unable to configure server discovery.
- Solution: Verify that you are signed in as a domain user and not as a local user on the IPAM server.
- Problem: A DNS server is not discovered.
- Solution: Verify that the DNS server configured on the IPAM server's network interface can query name server (NS) records for the domain zone, and that an NS record exists for the DNS server. All NS records for the domain zone and NS records for allowed DNS suffixes are queried during the DNS server discovery operation.
- Problem: A DHCP server is not discovered.
- Solution: Verify that the DHCP server role is not installed on the IPAM server. Verify that at least one IPv4 scope is configured on a DHCP server, and that the IPAM server has a

TCP/IP connection to the DHCP server. Also verify that DHCP INFORM request messages sent by IPAM server are not filtered on the network.

- Problem: You are unable to manually add a server to IPAM.
- Solution: Verify DNS resolution of the IP address of the server. Also verify that the server name is present in the Active Directory Global Catalog.

IP address space management

- Problem: You are unable to find an available IP address from a Microsoft DHCP-managed IP address range.
- Solution: Verify that DHCP RPC firewall ports are enabled on the target DHCP server, and that your user account is a member of the DHCP Users group on the DHCP server.
- Problem: You are unable to import a comma-separated value file with IP address records into IPAM.
- Solution: Verify that all the header field names in the file match the existing built-in and custom field names in IPAM. Create new custom field names in IPAM if required.
- Problem: Importing records from a comma-separated value file is successful, but some of the records are not displayed in IPAM.
- Solution: Review the error log for the import operation. By default, this log is created in the Documents folder, and it contains details of any records that failed to import. Use the error information provided to fix each failed record instance, and then repeat the import operation.
- Problem: You are unable to select DHCP servers or DNS zones in the IP address add or edit dialog box.
- Solution: Verify that the DHCP and DNS servers that are hosting the scope or zone are managed by IPAM. Also verify that you have administrative privileges on the target DHCP or DNS server.

Monitoring and management

- Problem: A managed server availability state is not reachable.
- Solution: Verify TCP/IP connectivity between the IPAM server and the target server. Verify that the DNS Server service or DHCP Server service is running on the target server, and that the IPAM server has been successfully provisioned with Read access on the service ACLs. Also verify that remote service management ports are open on the target server.
- Problem: You are unable to make configuration changes on a DHCP server or scope.
- Solution: Verify that DHCP RPC firewall ports are enabled on the target DHCP server, and that you are signed in with an account that has DHCP Administrators privileges on the target DHCP server.

See also

Best Practices

IPAM Backup and Restore

This topic provides instructions for a performing backup and restore of an IPAM server, using Windows Server Backup tools in Windows Server® 2012. You can use the steps in this topic to:

1. Configure a regularly scheduled backup of IPAM data.
2. Recover the IPAM database on an IPAM server.

Requirements for using Windows Server Backup

Windows Server Backup is available on all editions of Windows Server 2012. However, the Windows Server Backup snap-in is not available when you use the Server Core installation option. To run backups for computers with a Server Core installation, you must use the command line or manage backups remotely from another computer that is running a full installation of Windows Server 2012. For more information, see Windows Server Backup Cmdlets in Windows PowerShell (http://technet.microsoft.com/library/ee706683.aspx).

If Windows Server Backup feature is not installed on your system, you can install the feature as described in Install Windows Server Backup Tools (http://technet.microsoft.com/library/cc732081.aspx).

Membership in the **Administrators** group or **Backup Operators** group, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

Backup and restore IPAM

Important
If the version of IPAM changes or patches are installed between backup and restore procedures, you may need to convert the database schema to higher version so that it is compatible.

The following procedures assume that a network share is available on another server (for example: \\MyBackupSrv\ipambackup) to store IPAM backup data.

To create a scheduled backup

1. On the IPAM server, open **Administrative Tools** and then click **Windows Server Backup**.
2. In the navigation pane, right-click **Local Backup** and then click **Backup Schedule**.

   Tip
   Choosing **Backup Schedule** does not create a backup immediately. After using **Backup Schedule** to create a regularly scheduled backup, you can also choose to create a backup immediately with the **Backup Once** option.

3. In the Backup Schedule Wizard, click **Next**.
4. On the **Select Backup Configuration** page, choose **Custom** and then click **Next**.
5. On the **Select Items for Backup** page, click **Add Items**.
6. In **Select Items**, browse to **Windows\System32\ipam\Database** and select **ipam.mdf** and **ipam_log.ldf**.

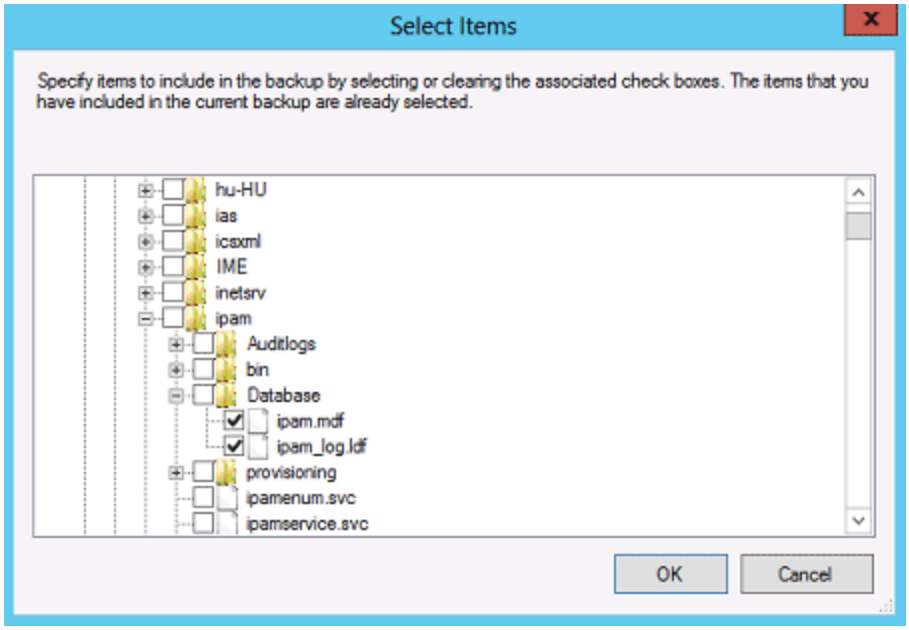

7. Click **OK** and then click **Next**.
8. On the **Specify Backup Time** page, choose when and how often to perform a backup of the IPAM database, and then click **Next**.
9. On the **Specify Destination Type** page, choose a location to store backup files and then click **Next**. In this example, a shared network folder is used. If you choose another option, you must have a hard disk available to store files that does not contain operating system or application files.
10. If you chose to backup files to a network share, click **OK** in the notification dialog box.
11. Type a location for the network share under Location (for example: **\\MyBackupSrv\ipambackup**) and then click **Next**.
12. In the Windows Security dialog box, provide credentials that have permission to write to the network share and then click **OK**.
13. On the **Confirmation** page, click **Finish**.
14. Verify that **You have successfully created the backup schedule** is displayed, and then click **Close**.

To restore from a backup

1. Temporarily stop the Windows Internal Database service on the IPAM server by typing **services.msc** in an elevated command prompt, pressing ENTER, right-clicking

Windows Internal Database, and then clicking **Stop**.

Leave the Services console open so that you can start the Windows Internal Database service again after the IPAM database has been restored from backup.

2. On the IPAM server, open **Administrative Tools** and then click **Windows Server Backup**.
3. Right-click **Local Backup** and then click **Recover**.
4. On the **Getting Started** page, choose the location where a backup is stored. In this example, a backup is stored on a network share so **A backup stored on another location** is chosen. Click **Next** to continue.
5. On the **Specify Location Type** page, choose **Remote shared folder**. Choose **Local drives** only if your backup was not stored remotely as shown in this example. Click **Next** to continue.
6. On the Specify Remote Folder page, type the location of your network share (for example: **\\MyBackupSrv\ipambackup**) and then click **Next**.
7. On the **Select Backup Date** page, choose a date to use for the restore point of the IPAM database and then click **Next**.
8. On the **Select Recovery Type** page, choose **Files and folders** and then click **Next**.
9. Browse to **Windows\System32\ipam\Database**, click the **Database** folder, and then select both of the available files: **ipam.mdf** and **ipam_log.ldf**. Click **Next** to continue.

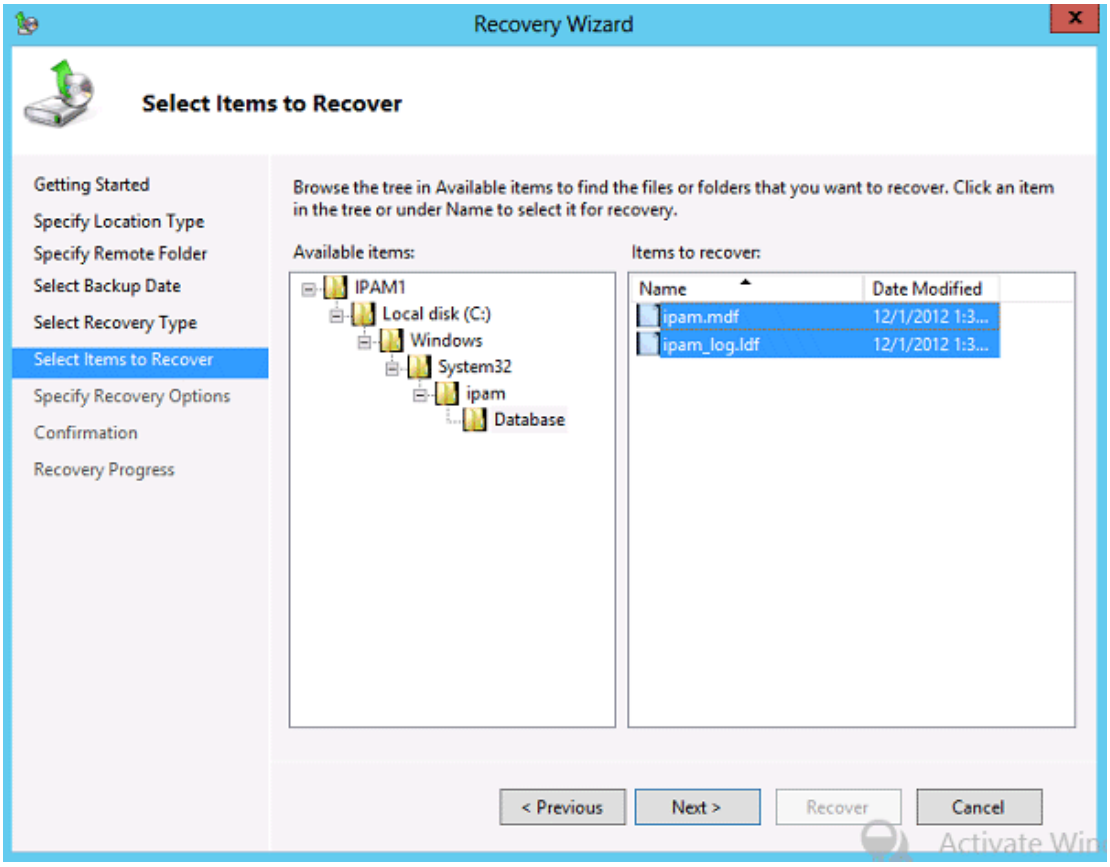

10. On the **Specify Recovery Options** page, click **Browse**, browse to the **Windows\System32\ipam\Database** folder, and click **OK**
11. Under **When the wizard find items…** choose **Overwrite the existing versions with the recovered versions**, ensure that the **Restore access control…** checkbox is selected, and then click **Next**.

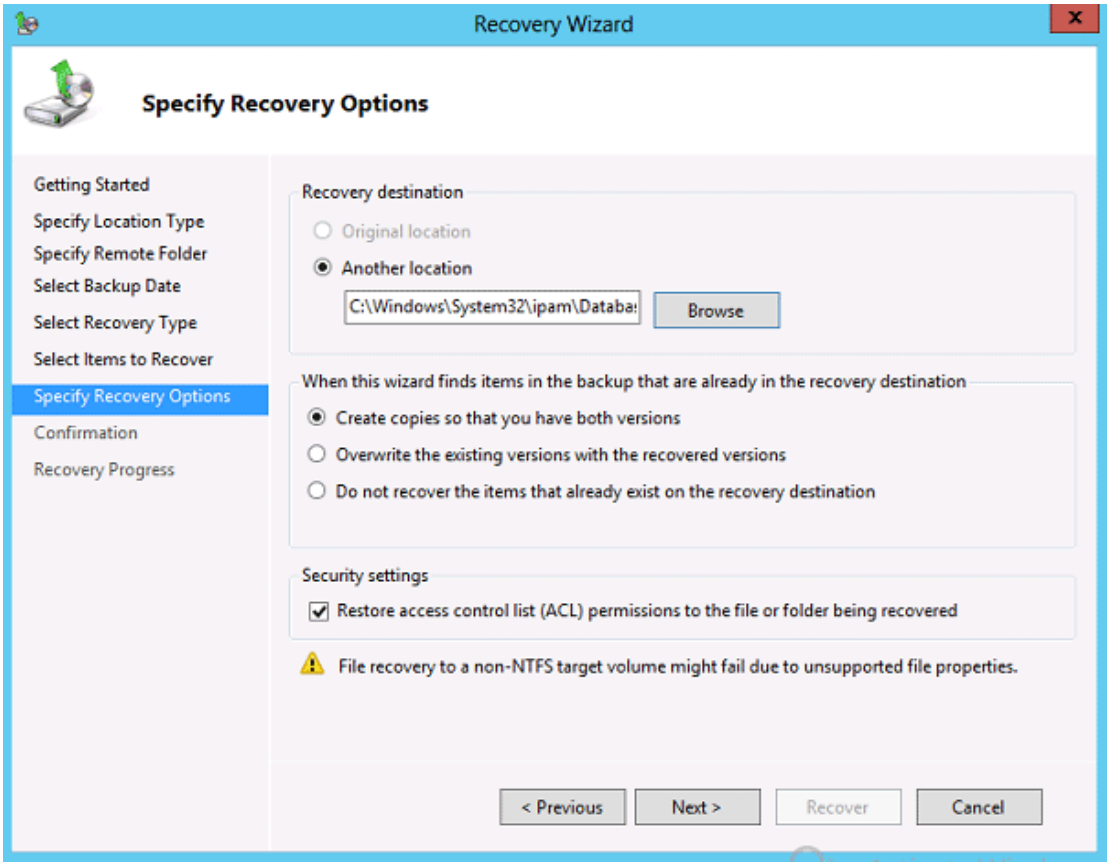

12. On the **Confirmation** page, click **Recover**.
13. Verify that **Completed** is displayed next to **Status**, and then click **Close**.
14. In the Services console, right-click **Windows Internal Database**, and then click **Start**.

Tip

If the Windows Internal Database service fails to start, or IPAM is not able to connect to the database after a restore operation, you can review the error log file at **%SystemRoot%\WID\Log**.

See also

Capacity planning

Low Latency Workloads Technologies

This section provides overviews for the following group of technologies that are designed for or which were improved in Windows Server® 2012 to address low latency computing scenarios.

1. Data Center Bridging
2. Data Center Transmission Control Protocol (DCTCP)
3. Kernel Mode Remote Direct Memory Access (kRDMA)
4. Network Interface Card (NIC) Teaming
5. NetworkDirect
6. Receive Segment Coalescing (RSC)
7. Receive Side Scaling (RSS)
8. Registered Input/Output (RIO) API Extensions
9. Transmission Control Protocol (TCP) Loopback Optimization
10. Low Latency Workloads Management and Operations

Latency means delay, and it refers to the length of time that elapses between two specific events, such as the amount of time between the transmission and the reception of a network message between two computers over a network path. Latency has a variety of possible causes, including electrical propagation delays, processing delays, and queuing effects.

A variety of processing workloads require that the time spent on inter-machine communications is reduced to the lowest amount possible. These workloads include distributed computing algorithms whose convergence time is bound by the network latency. Examples of such systems include distributed consensus and agreement protocols, Message Passing Interface (MPI) workloads, and distributed caching. Stock trading and other financial markets workloads also require that the latency incurred by network communications is reduced to the greatest degree possible.

Low latency computing environments typically contain applications that require very fast inter-process communication (IPC) and inter-computer communications, a high degree of predictability regarding latency and transaction response times, and the ability to handle very high message rates. The following section contains information about technologies that you can use to improve performance in low latency computing scenarios.

Data Center Bridging

Data Center Bridging

Data Center Bridging (DCB) is a suite of Institute of Electrical and Electronics Engineers (IEEE) standards that enable Converged Fabrics in the data center, where storage, data networking, cluster IPC and management traffic all share the same Ethernet network infrastructure. DCB provides hardware-based bandwidth allocation to a specific type of traffic and enhances Ethernet transport reliability with the use of priority-based flow control.

Hardware-based bandwidth allocation is essential if traffic bypasses the operating system and is offloaded to a converged network adapter, which might support Internet Small Computer System Interface (iSCSI), Remote Direct Memory Access (RDMA) over Converged Ethernet, or Fiber Channel over Ethernet (FCoE). Priority-based flow control is essential if the upper layer protocol, such as Fiber Channel, assumes a lossless underlying transport.

For more information, see Data Center Bridging (DCB) Overview and Overview of Data Center Bridging (Windows Drivers)

Data Center Transmission Control Protocol (DCTCP)

Data Center Transmission Control Protocol (DCTCP)

Data centers host diverse applications, mixing on the same network a variety of workflows that require small, predictable latency, while other applications require large, sustained throughput. In this environment, today's state-of-the-art Transmission Control Protocol (TCP) congestion control mechanisms do not provide sufficiently detailed congestion control settings. This results in queue formations in network switches leading to delays, fluctuations in latency, and timeouts.

To reduce this problem, Windows Server 2012 introduces DCTCP, which uses Explicit Congestion Notification (ECN) to estimate the extent of the congestion at the source, and reduce the sending rate only to the extent of the congestion. This provides a more detailed control over network traffic, allowing DCTCP to operate with very low buffer occupancies while still achieving high throughput.

The following illustration demonstrates the effectiveness of DCTCP in achieving full throughput while taking up a very small footprint in an Ethernet switch packet buffer, as compared to traditional TCP. The graph depicts the queue length in the network switch when using DCTCP and TCP. Two separate 1 gigabits per second (Gbps) TCP/IP streams are directed into two separate switch ports, and are combined into a single outgoing 1 Gbps port.

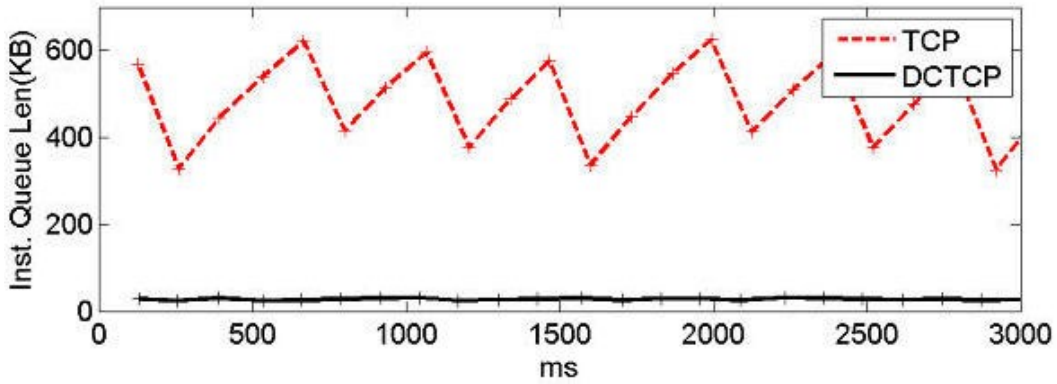

With traditional TCP, long-lived, large volume TCP flows cause the length of the bottleneck queue to grow until packets are dropped, resulting in the sawtooth pattern of the TCP traffic. When long and short flows traverse the same queue, two impairments occur. First, packet loss on the short flows may occur as described above. Second, there is a queue buildup even when no packets are lost. The short flows experience increased latency as they are queued behind packets from the large flows.

DCTCP, however, provides earlier, more detailed responses to congestion, which effectively fine-tune the sending rates at each source to operate with much smaller message queue build-up in the switch, while maintaining the same aggregate throughput. The much lower queue lengths exhibited with DCTCP avoid the latency and variation in latency that occur with TCP.

When used with commodity, shallow-buffered switches, DCTCP delivers the same or better throughput than TCP, while using 90% less buffer space in the network infrastructure. Unlike TCP, it also provides high burst tolerance and low latency for short flows. While the limitations of TCP cause the mistransmission of traffic sent today, DCTCP enables the applications to handle ten times the current background traffic without impacting foreground traffic. In addition, a tenfold increase in foreground traffic does not cause any timeouts, thus largely eliminating problems that can occur when computers retransmit messages as a result of previous timeouts.

Kernel Mode Remote Direct Memory Access (kRDMA)

Kernel Mode Remote Direct Memory Access (kRDMA)

kRDMA is a new feature for Windows Server 2012. Kernel mode Remote Direct Memory Access (RDMA) is an accelerated input-output (I/O) delivery model that allows application software to bypass most software layers to communicate directly with the computer hardware, which improves application performance and reduces delay.

For more information, see the Windows Dev Center topic kRDMA Reference.

Network Interface Card (NIC) Teaming

NIC teaming

NIC teaming is a built-in feature of Windows Server 2012, and allows you to provide fault tolerance for your network adapters. NIC teaming allows multiple network interfaces to work together as a team, preventing connectivity loss if one network interface card fails. It allows a server to tolerate network interface card and port failure up to the first switch segment. NIC teaming also allows you to aggregate bandwidth from multiple network adapters, for example, so four 1-gigabit (Gb) network adapters can provide an aggregate of 4 Gb/second of throughput.

The advantages of a built-in teaming solution are that it works with all network interface card vendors, spares you from any potential problems that proprietary solutions cause, and provides a common set of management tools for all adapter types.

You can configure outbound traffic to be distributed among the available network adapters by using modes for network traffic distribution. These modes keep all packets that are associated with a single flow (TCP stream) on a single interface, which minimizes the possibility of out-of-order packet arrival scenarios. The NIC teaming feature uses the following modes for network traffic distribution.

- Originating media access control (MAC) address mode. In cases where virtual machines (VMs) have independent MAC addresses, you can configure the MAC address of each VM as the basis for dividing network traffic. There is a major advantage in using this mode with virtualization: because the adjacent switch sees specific source MAC addresses on one connected interface, the switch balances the egress load (the network traffic from the switch to the host) on multiple links based on the destination MAC address per VM. This is particularly useful when Virtual Machine Queue (VMQ) is used as a queue and can be placed on a specific network interface card where the traffic is expected to arrive. However, this mode might not be detailed enough to get a well-balanced distribution and will always limit a single VM to the bandwidth available on a single interface.
- Hashing mode. This mode creates a hash that is based on components of the packet. The mode then assigns packets that have that hash value to one of the available interfaces, which keeps all packets from the same TCP stream on the same interface. Frequently, this mechanism alone is sufficient to create a reasonable balance across the available interfaces. The components that can be used as inputs to the hashing mode include:
  - Source and destination MAC addresses
  - Source and destination IP addresses with or without considering the MAC addresses (2-tuple hash)
  - Source and destination TCP ports, usually used along with the IP addresses (4-tuple hash)

The 4-tuple hash creates a finer distribution of traffic streams, resulting in smaller streams that can be independently moved between interfaces. This algorithm is not compatible with IPsec; if IPsec is deployed, this algorithm reverts to a 2-tuple hash.

- Broadcast mode. This mode broadcasts all packets on all interfaces. It does not balance the network traffic load and it causes a need for significant extra processing to drop duplicate packets at the receive side.

In addition, the NIC teaming feature uses the following two sets of configuration modes for use with switches:

- Switch-independent modes. These modes make it possible for network interface card team members to connect to different switches because the switch does not know that the interface is part of a team at the host. These modes do not require the switch to participate in the teaming.
- Switch-dependent modes. These modes require the switch to participate in the NIC teaming. Here, all interfaces of the team are connected to the same switch.

NIC teaming in virtual machines

Windows Server 2012 allows a virtual machine to have virtual network adapters connected to more than one Hyper-V switch and still have connectivity even if the network interface card under that switch becomes disconnected.

Because failover between network adapters in a virtual machine might result in traffic being sent with the MAC address of another interface, each Hyper-V switch port associated with a virtual machine that is using NIC teaming must be set to allow MAC spoofing. Because some solutions that utilize Single-Root Input/Output Virtualization (SR-IOV) do not support MAC spoofing in the hardware, NIC teaming — and particularly the failover aspects of NIC teaming — might not work when SR-IOV network adapters are present.

Incompatibilities. The NIC teaming feature is compatible with networking capabilities in Windows Server 2012 with three exceptions:

- SR-IOV
- Remote Direct Memory Access (RDMA)
- TCP Chimney

For SR-IOV and RDMA, data is delivered directly to the network interface card without passing it through the networking stack. Therefore, the network interface card team cannot see or redirect the data to another path in the team. In this release, TCP Chimney is not supported with NIC teaming.

For more information about NIC teaming, see NIC Teaming Overview.

NetworkDirect

NetworkDirect

NetworkDirect is a user-mode programming interface specification for Remote Direct Memory Access (RDMA). RDMA is provided by RDMA-enabled network adapters. Because NetworkDirect is fabric agnostic, it can be used on Infiniband, iWARP, and RoCE. NetworkDirect allows RDMA-enabled network interface card manufacturers to expose the RDMA functionality of their network adapters in Windows Server 2012.

RDMA is a kernel bypass technique which makes it possible to transfer large amounts of data quite rapidly. Because the transfer is performed by the DMA engine on the network adapter, the CPU is not used for the memory movement, which frees the CPU to perform other work.

NetworkDirect is widely used for high-performance computing (HPC) applications in which computational workloads are distributed to large numbers of servers for parallel processing. In addition, various financial markets trading workloads also require extremely low latency and extremely high message rates, which RDMA can provide.

Windows Server 2012 extends support for NetworkDirect to non-HPC server versions, adds numerous management features (performance monitor counters and event tracing) and a new Kernel Mode interface which is used as the underlying transport mechanism for Server Message Block 2 and RDMA.

For more information, see the following Windows Dev Center topics.

- OID_NDK_CONNECTIONS (Windows Drivers)
- OID_NDK_LOCAL_ENDPOINTS (Windows Drivers)
- NDK_OBJECT_TYPE enumeration (Windows Drivers)
- NDK_FN_CONNECT routine (Windows Drivers)
- NDIS_QOS_CLASSIFICATION_ELEMENT structure (Windows Drivers)
- NDISTest 6.5 -
- NDK_FN_LISTEN routine (Windows Drivers)
- NDK_FN_CREATE_SHARED_ENDPOINT routine (Windows Drivers)
- NDK_FN_CONNECT_WITH_SHARED_ENDPOINT routine (Windows Drivers)

Receive Segment Coalescing (RSC)

Receive Segment Coalescing (RSC)

RSC is a stateless offload technology that helps reduce CPU utilization for network processing on the receive side by offloading tasks from the CPU to an RSC-capable network adapter. CPU saturation due to networking-related processing can limit server scalability. This problem in turn

reduces the transaction rate, raw throughput, and efficiency. RSC enables an RSC-capable network interface card to do the following:

- Parse multiple TCP/IP packets and strip the headers from the packets while preserving the payload of each packet.
- Join the combined payloads of the multiple packets into one packet.
- Send the single packet, which contains the payload of multiple packets, to the network stack for subsequent delivery to applications.

The network interface card performs these tasks based on rules that are defined by the network stack subject to the hardware capabilities of the specific network adapter. This ability to receive multiple TCP segments as one large segment significantly reduces the per-packet processing overhead of the network stack. Because of this, RSC significantly improves the receive-side performance of the operating system (by reducing the CPU overhead) under network I/O intensive workloads.

Note

Examples of other stateless offloads are Internet Protocol (IP) and Transmission Control Protocol (TCP) checksum, and Large Send Offload, which is also called Segmentation Offload.

When your server has an RSC-capable network interface card installed, RSC provides benefits under many circumstances, including the following:

- Hosted cloud deployments. RSC minimizes the CPU cycles that are spent for network storage and live migration.
- I/O intensive database applications and database replication.
- File servers that are deployed with the Windows Server 2012 File Services server role. In addition, if your file server is also configured as a BranchCache-enabled content server, BranchCache performance is improved by RSC.
- I/O intensive server workloads.

Important

Because RSC is designed to improve performance and reduce latency for receive-intensive workloads, it does not significantly improve performance for send-intensive workloads, such as Web servers that send HTML files to Web browsers running on client computers and other devices. For receive-intensive workloads, RSC allows virtual and physical servers to effectively service a larger number of clients and connections.

Here are some key facts about RSC:

- RSC is enabled by default for clean installations of all editions of Windows Server 2012 on computers that have RSC-capable network adapters. You do not need to take any action to enable RSC when you have an RSC-capable network interface card installed in your physical computer or used by a virtual machine that is running Windows Server 2012. In addition, you can track RSC status and usage by using relevant Performance Monitor counters.

  Important

When you perform an operating system upgrade from a previous version of Windows to Windows Server 2012, RSC functionality is disabled by default. This ensures that server functionality remains consistent despite the upgrade of the operating system.

- You can enable or disable RSC globally or on a per-interface basis.
- RSC does not cause any difference in performance between IPv4 and IPv6 networks.
- RSC works seamlessly with third-party drivers that support Network Driver Interface Specification (NDIS) 6.30 and above (NDIS Lightweight Filters, NDIS Intermediates and Windows Filtering Platform [WFP] callout drivers).
- RSC is automatically disabled for older incompatible drivers.
- RSC works well in virtualized environments. The availability of RSC is limited to the parent partition for storage and live migration, and to virtual machines (VMs) running Windows® 8 with SR-IOV capable network adapters. RSC functionality is not available for VMs that are not running Windows 8 and VMs that do not run SR-IOV.
- RSC is defined only for TCP connections. RSC does not work with User Datagram Protocol (UDP) and other transport layer protocols.
- RSC is a stateless offload technology and does not hold the state of TCP connections. It is not a replacement for TCP Chimney, because RSC-capable network adapters are not expected to handle the TCP state computer.
- RSC does not function with IPsec encrypted traffic, because network adapters currently cannot coalesce IPsec packets.
- RSC provides performance counters that capture the following for each interface: the number of TCP connections on which RSC is active; the number of packets that are coalesced per Deferred Procedure Call (DPC); the number of RSC aborts that are issued by the network adapter; the maximum number of packets coalesced in one DPC.

For more information, see the following Windows Dev Center topics.

- NDIS_RSC_NBL_INFO union (Windows Drivers)
- NDIS_TCP_RECV_SEG_COALESCE_OFFLOAD structure (Windows Drivers)
- OID_TCP_RSC_STATISTICS (Windows Drivers)
- NDIS_OFFLOAD structure (Windows Drivers)
- NDIS_OFFLOAD_PARAMETERS structure (Windows Drivers)
- Standardized INF Keywords for RSC (Windows Drivers)
- FWPS_CALLOUT2 structure (Windows Drivers)

Receive Side Scaling (RSS)

Receive Side Scaling (RSS)

Note

For information about virtual Receive-side Scaling (vRSS), see Virtual Receive-side Scaling in Windows Server 2012 R2.

RSS enables network adapters to distribute the kernel-mode network processing load across multiple processor cores in multi-core computers. The distribution of this processing makes it possible to support higher network traffic loads than would be possible if only a single core were to be used. In Windows Server 2012, RSS has been enhanced, including computers with more than sixty-four processors. RSS achieves this by spreading the network processing load across many processors and actively load balancing TCP terminated traffic.

In addition, RSS in Windows Server 2012 provides automatic load balancing capabilities for non-TCP traffic (including UDP unicast, multicast, and IP-forwarded traffic), improved RSS diagnostics and management capabilities, improved scalability across Non-Uniform Memory Access (NUMA) nodes, and improved resource utilization and resource partitioning through facilitation of kernel-scheduler and networking stack alignment.

RSS diagnostics and management capabilities include:

- Event tracing. RSS logs events for the following cases:
  - RSS configuration information (base RSS CPU, number of queues, and indirection table)
  - Indirection table change notification
  - Change in processor set
- NICStats functionality through Windows Management Instrumentation (WMI) and Windows PowerShell.
- Improved per-interface default settings.
- RSS profiles. RSS now supports the following predefined profiles:
  a. **Default:** With this profile selected, new RSS features are disabled and it performs very similarly to the way that RSS performs in Windows Server 2008 R2.
  b. **No dynamic load balancing:** With this profile selected, RSS distributes but does not load balance at runtime.
  c. **NUMA scalability:** With this profile selected, RSS CPUs are assigned per NUMA node on a round robin basis. This is done to enable applications running on multi-NUMA servers to scale well. For this profile, RSS assigns at least one RSS CPU per NUMA node before assigning the second RSS CPU for each NUMA node. While in NUMA scalability mode, RSS uses the limits that are set for scaling across multiple NUMA nodes.
  d. **NUMA scalability with no dynamic load balancing:** With this profile selected, the profile NUMA Scalability is used, however RSS does not perform load balancing.
  e. **Conservative Scaling:** With this profile selected, RSS uses as few processors as possible to sustain the load. This option assists in reducing the number of interrupts.

RSS provides the following additional benefits for low latency scenarios:

- **Parallel Receive Processing:** Receive packets from a single network interface card can be indicated via generating interrupts and DPCs concurrently on multiple CPUs.
- **Preserving in-order packet delivery:** Received packets for a specific stream from a single network interface card are delivered in order to the TCP/IP protocol driver.

- **Dynamic Load Balancing:** As the load on the host varies, RSS rebalances the network processing load between the processors.
- **Cache Locality:** Because packets from a single connection are mapped to a specific processor state for a particular connection stays resident in the cache of the processor, thereby eliminating cache thrashing and improving performance.
- **Send Side Scaling:** The RSS hash value is passed by the TCP/IP protocol to the network interface card in each packet on the egress path thus allowing for the send completions to be indicated on the same CPU. This enables better scaling on the send side.
- **Toeplitz Hash:** The default generated RSS signature is statistically secure, making it more difficult for malicious remote hosts to force the system into an unbalanced state.

Following are some application-specific advantages of RSS:

- RSS enables Windows Server 2012 to scale optimally on large servers that have more than 64 processors.
- RSS allows multi-NUMA node servers, such as web servers, file servers, and servers running other workloads, to scale efficiently.
- For DirectAccess, RSS allows improved scalability by actively load balancing forwarded traffic.
- RSS enables financial applications, such as ticker plants and algorithmic trading systems using UDP as their transport, to scale with increasing load.
- RSS provides applications with an input/output control to align themselves with the stack.
- RSS is compatible with the Windows load balancing and failover solution.
- RSS provides administrators with the ability to deploy, manage, and diagnose RSS functioning through standard management interfaces and diagnostics tools.
- RSS optimally scales all workloads by aligning the networking stack with the kernel scheduler. RSS uses the processor list provided by the kernel scheduler while making processor selection decisions.

For more information, see Introduction to Receive Side Scaling (Windows Drivers).

Registered Input/Output (RIO) API Extensions

Registered Input/Output (RIO) API Extensions

The RIO API is a new extension to Windows Sockets (Winsock) and provides an opportunity for you to reduce network latency, increase message rates and improve the predictability of response times for applications that require very high performance, very high message rates, and predictability. RIO API extensions allow applications that process large numbers of small messages to achieve higher I/O operations per second (IOPS) with reduced jitter and latency. Server loads with high message rates and low latency requirements benefit most from RIO API extensions, including applications for financial services trading and high speed market data

reception and dissemination. In addition, RIO API extensions provide high IOPS when you deploy many Hyper-V virtual machines (VMs) on a single physical computer.

RIO enables send and receive operations to be performed with pre-registered buffers using queues for requests and completions. Send and receive operations are queued to a request queue that is associated with a Winsock socket. Completed I/O operations are inserted into a completion queue, and many different sockets can be associated with the same completion queue. Completion queues can also be split between send and receive completions. Completion operations, such as polling, can be performed entirely in user-mode and without making system calls.

The use of registered buffers streamlines the network related processing, reduces jitter, and additionally makes it possible for application developers to specify the NUMA node affinity of networking buffers used by the protocol stack — further enhancing overall performance, and reducing latency and jitter characteristics.

RIO API extensions support Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and multicast UDP, as well as both IPv4 and IPv6.

You can use RIO API extensions if you want to achieve any of the following:

- Scale up your server to minimize CPU utilization per message
- Reduce the latency contribution and jitter of the networking stack to a minimum
- Handle very high rates of multicast or UDP traffic

Use of the RIO API extensions have the following additional benefits:

- RIO works on all editions of Windows Server 2012.
- RIO is compatible with normal network adapters and does not require special network adapters or RDMA.
- RIO is fully compatible with existing Windows networking features, including RSS, RSC, network interface card teaming, and static offloads.
- RIO works with virtualization when you deploy Hyper-V in Windows Server 2012.
- RIO sockets use the standard Windows networking stack and standard TCP/IP and UDP protocols.

For more information, see What's New for Windows Sockets (Windows) and RIO_EXTENSION_FUNCTION_TABLE structure.

Transmission Control Protocol (TCP) Loopback Optimization

Transmission Control Protocol (TCP) loopback optimization

Many applications use the TCP/IP loopback path to provide interprocess communications. While convenient from a development and deployment perspective, the TCP/IP loopback path can become a bottleneck for applications that require very high performance or very low latency.

To address the needs of these kinds of applications, Windows Server 2012 introduces a TCP Loopback performance enhancement that provides an optimized fast path. The optimized fast path increases the achievable TCP send and receive rates, and reduces latency and jitter for low latency scenarios.

Optimization of TCP loopback provides the following benefits:

- TCP loopback optimization provides an optimized low-latency loopback path for performance-critical applications that rely on loopback for inter-process communication.
- TCP loopback optimization does not cause any disparity in performance between IPv4 and IPv6 networks.
- The optimized loopback path works in virtualized environments. Loopback is defined as an inter-process communication mechanism for both processes on the same computer. Because of this, VM-to-VM optimization — when both VMs are installed in the same physical host — is not provided. The boundary for optimization between VMs in this case is the individual VM boundary.
- With TCP loopback optimization, you are provided the choice to use either registered (Winsock RIO) sockets or Winsock to make use of the optimized fast path. You can observe when the connections take the slow path rather than the optimized fast path.
- The optimized loopback path is available for developers to programmatically control on a per-connection basis.

In addition, TCP loopback fast path is an optimization for newly written applications, and does not provide performance improvements for existing applications.

TCP loopback optimization is available on all versions of Windows Server 2012 and Windows 8. This feature is turned on whenever the Socket option for fast path is set for a connection, and utilizes the optimized fast path whenever the application has opted-in and the Windows Filtering Platform (WFP) filter driver is not present for that connection.

The TCP loopback fast path is enabled by the socket IOCTL SIO_LOOPBACK_FASTPATH applied to the sending and receiving sockets.

Low Latency Workloads Management and Operations

New Performance Counters

Following are some of the new performance counters for Windows Server 2012.

UDP Datagram Loss

Windows Server 2012 includes new performance counters that make it easier to measure UDP datagram and multicast receive loss due to insufficient socket buffering. If UDP datagrams or multicast datagrams arrive more rapidly than the receive application can process them, or if the application has not specified sufficient receive socket buffering, these counters record the number of datagrams that are dropped and the number of datagrams that are dropped per second. The counters are part of the Microsoft Winsock BSP counter set in performance monitor.

Rejected Connections

This new performance counter provides the total number of rejected TCP connections, and the rate of rejected TCP connections per second. The counters are part of the Microsoft Winsock BSP counter set in performance monitor.

NetworkDirect Activity

These new performance counters allow you to track NetworkDirect (RDMA) activity, including traffic rates, errors, and control plane activity.

Per Processor Network interface card Activity

These new performance monitor counters allow you to monitor how RSS is distributing the workload across multiple cores, automatic rebalancing of workloads, and RSC packets coalescing activity.

NUMA systems synchronization

NUMA systems synchronization performance monitor counters track lock contention and related information, such as the generation of inter-processor interrupts (IPIs), to simplify performance monitoring of NUMA systems.

Windows PowerShell management of networking features

Windows PowerShell now supports a wide variety of network configuration tasks and management tasks on both local and remote computers, including:

- Networking stack configuration
- Network interface card configuration
- RSS configuration
- SR-IOV configuration
- Windows Firewall configuration
- RDMA configuration
- Network interface card teaming configuration

Because it is now possible to examine and manage these parameters locally and remotely, Windows PowerShell makes it much easier to ensure that large server installations are correctly and consistently configured using scripts and automation.

Network Load Balancing Overview

This topic provides an overview of the Network Load Balancing (NLB) feature. By managing two or more servers as a single virtual cluster, NLB enhances the availability and scalability of Internet server applications such as those used on web, FTP, firewall, proxy, virtual private network (VPN), and other mission-critical servers. This topic describes the NLB feature and provides links to additional guidance about creating, configuring, and managing NLB clusters.

Did you mean…

- Failover Clustering Overview
- Scale-Out File Server for Application Data Overview
- **Hyper-V Replica Overview**

Feature description

The Network Load Balancing (NLB) feature distributes traffic across several servers by using the TCP/IP networking protocol. By combining two or more computers that are running applications into a single virtual cluster, NLB provides reliability and performance for web servers and other mission-critical servers.

The servers in an NLB cluster are called *hosts*, and each host runs a separate copy of the server applications. NLB distributes incoming client requests across the hosts in the cluster. You can configure the load that is to be handled by each host. You can also add hosts dynamically to the cluster to handle increased load. NLB can also direct all traffic to a designated single host, which is called the *default host*.

NLB allows all of the computers in the cluster to be addressed by the same set of IP addresses, and it maintains a set of unique, dedicated IP addresses for each host. For load-balanced

applications, when a host fails or goes offline, the load is automatically redistributed among the computers that are still operating. When it is ready, the offline computer can transparently rejoin the cluster and regain its share of the workload, which allows the other computers in the cluster to handle less traffic.

Practical applications

NLB is useful for ensuring that stateless applications, such as web servers running Internet Information Services (IIS), are available with minimal downtime, and that they are scalable (by adding additional servers as the load increases). The following sections describe how NLB supports high availability, scalability, and manageability of the clustered servers that run these applications.

High availability

A high availability system reliably provides an acceptable level of service with minimal downtime. To provide high availability, NLB includes built-in features that can automatically:

- Detect a cluster host that fails or goes offline, and then recover.
- Balance the network load when hosts are added or removed.
- Recover and redistribute the workload within ten seconds.

Scalability

Scalability is the measure of how well a computer, service, or application can grow to meet increasing performance demands. For NLB clusters, scalability is the ability to incrementally add one or more systems to an existing cluster when the overall load of the cluster exceeds its capabilities. To support scalability, you can do the following with NLB:

- Balance load requests across the NLB cluster for individual TCP/IP services.
- Support up to 32 computers in a single cluster.
- Balance multiple server load requests (from the same client or from several clients) across multiple hosts in the cluster.
- Add hosts to the NLB cluster as the load increases, without causing the cluster to fail.
- Remove hosts from the cluster when the load decreases.
- Enable high performance and low overhead through a fully pipelined implementation. Pipelining allows requests to be sent to the NLB cluster without waiting for a response to a previous request.

Manageability

To support manageability, you can do the following with NLB:

- Manage and configure multiple NLB clusters and the cluster hosts from a single computer by using NLB Manager or the NLB Windows PowerShell cmdlets.

- Specify the load balancing behavior for a single IP port or group of ports by using port management rules.
- Define different port rules for each website. If you use the same set of load-balanced servers for multiple applications or websites, port rules are based on the destination virtual IP address (using virtual clusters).
- Direct all client requests to a single host by using optional, single-host rules. NLB routes client requests to a particular host that is running specific applications.
- Block undesired network access to certain IP ports.
- Enable Internet Group Management Protocol (IGMP) support on the cluster hosts to control switch port flooding (where incoming network packets are sent to all ports on the switch) when operating in multicast mode.
- Start, stop, and control NLB actions remotely by using Windows PowerShell commands or scripts.
- View the Windows Event Log to check NLB events. NLB logs all actions and cluster changes in the event log.

Important functionality

Note

The NLB functionality in Windows Server 2012 R2 and Windows Server 2012 is generally the same as in Windows Server 2008 R2. However, some task details are changed. For information about new ways to do tasks, see Common Management Tasks and Navigation in Windows Server 2012 R2 and Windows Server 2012.

NLB is installed as a standard Windows networking driver component. Its operations are transparent to the TCP/IP networking stack. The following figure shows the relationship between NLB and other software components in a typical configuration:

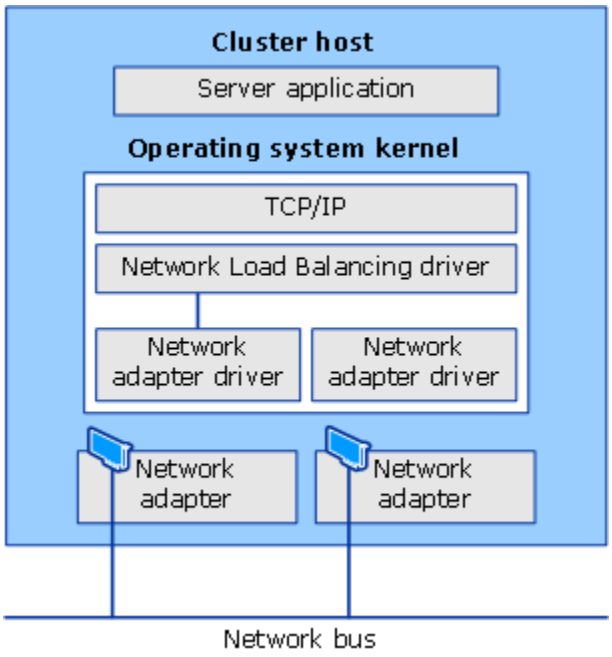

The NLB features include the following:

- Requires no hardware changes to run.
- Provides Network Load Balancing Tools to configure and manage multiple clusters and all of the hosts from a single remote or local computer.
- Enables clients to access the cluster by using a single, logical Internet name and virtual IP address, which is known as the cluster IP address (it retains individual names for each computer). NLB allows multiple virtual IP addresses for multihomed servers.

  Note
  NLB does not require servers to be multihomed to have multiple virtual IP addresses in the case of virtual clusters.
- Enables NLB to be bound to multiple network adapters, which enables you to configure multiple independent clusters on each host. Support for multiple network adapters differs from virtual clusters in that virtual clusters allow you to configure multiple clusters on a single network adapter.
- Requires no modifications to server applications so that they can run in an NLB cluster.
- Can be configured to automatically add a host to the cluster if that cluster host fails and is subsequently brought back online. The added host can start handling new server requests from clients.
- Enables you to take computers offline for preventive maintenance without disturbing the cluster operations on the other hosts.

Hardware requirements

To run an NLB cluster, the following are hardware requirements:

- All hosts in the cluster must reside on the same subnet.
- There is no restriction on the number of network adapters on each host, and different hosts can have a different number of adapters.
- Within each cluster, all network adapters must be either multicast or unicast. NLB does not support a mixed environment of multicast and unicast within a single cluster.
- If you use the unicast mode, the network adapter that is used to handle client-to-cluster traffic must support changing its media access control (MAC) address.

Software requirements

To run an NLB cluster, the following are software requirements:

- Only TCP/IP can be used on the adapter for which NLB is enabled on each host. Do not add any other protocols (for example, IPX) to this adapter.
- The IP addresses of the servers in the cluster must be static.

Note

NLB does not support Dynamic Host Configuration Protocol (DHCP). NLB disables DHCP on each interface that it configures.

Server Manager information

In Server Manager, use the Add Roles and Features Wizard to add the Network Load Balancing feature. Optionally you can install the Network Load Balancing Tools to manage a local or remote NLB cluster. The tools include Network Load Balancing Manager and the NLB Windows PowerShell cmdlets. For more information about installing features, see Install or Uninstall Roles, Role Services, or Features.

To open Network Load Balancing Manager in Server Manager, click **Tools**, and then click **Network Load Balancing Manager**.

See also

The following table provides links to additional information about the NLB feature that is available on the web.

Content type	References
Deployment	Network Load Balancing Deployment Guide | Configuring Network Load Balancing with Terminal Services
Operations	Managing Network Load Balancing Clusters |

Content type	References
	Setting Network Load Balancing Parameters | Controlling Hosts on Network Load Balancing Clusters
Troubleshooting	Troubleshooting Network Load Balancing Clusters | NLB Cluster Events and Errors
Security	Network Load Balancing Best Practices
Tools and settings	Network Load Balancing Windows PowerShell cmdlets
Community resources	High Availability (Clustering) Forum | Failover Clustering and Network Load Balancing Team Blog
Related technologies	Web Server (IIS) | Application Server | NIC Teaming Overview

Network Policy and Access Services Overview

This topic provides an overview of Network Policy and Access Services in Windows Server® 2012, including the specific role services of Network Policy Server (NPS), Health Registration Authority (HRA), and Host Credential Authorization Protocol (HCAP). Use the Network Policy and Access Services server role to deploy and configure Network Access Protection (NAP), secure wired and wireless access points, and RADIUS servers and proxies.

Did you mean…

- NAP Client Configuration Overview
- Network Policy and Access Services Overview for Windows Server 2008 R2

Role description

Network Policy and Access Services provides the following network connectivity solutions:

Solution	Description
Network Access Protection (NAP)	NAP is a client health policy creation, enforcement, and remediation technology. With NAP, system administrators can establish and automatically enforce health policies, which can

Solution	Description
	include software requirements, security update requirements, and other settings. Client computers that are not in compliance with health policy can be provided restricted network access until their configuration is updated and brought into compliance with policy.
802.1X authenticated wired and wireless access	When you deploy 802.1X-capable wireless access points and Ethernet switches, you can use Network Policy Server (NPS) to deploy certificate-based authentication methods that are more secure than password-based authentication. Deploying 802.1X-capable hardware with NPS allows you to ensure that intranet users are authenticated before they can connect to the network or obtain an IP address from a DHCP server.
Central network policy management with RADIUS server and proxy	Rather than configuring network access policy at each network access server, you can create policies in a single location that specify all aspects of network connection requests, including who is allowed to connect, when they can connect, and the level of security they must use to connect to your network.

Note
Network Policy and Access Services can be run in Windows Azure VMs, but there are no recommended scenarios for running it on Windows Azure VMs.

New and changed functionality

The following table lists the primary differences in the Network Policy and Access Services server role by operating system:

Feature/functionality	Windows Server® 2008 R2 and Windows Server® 2008	Windows Server 2012
Support for Windows PowerShell®	None	X

Support for Windows PowerShell

You can now use Windows PowerShell to automate the installation of the Network Policy and Access Services server role. You can also deploy and configure some aspects of Network Policy Server by using Windows PowerShell. For more information, see Windows PowerShell for Network Policy and Access Services.

Removed functionality

In Windows Server® 2008 R2 and Windows Server® 2008, Network Policy and Access Services included the Routing and Remote Access Service (RRAS) role service. In Windows Server 2012, RRAS is now a role service in the Remote Access server role.

Deprecated functionality

With the release of Windows Server 2012 R2, NAP is deprecated. NAP is fully supported in Windows Server 2012 R2 and Windows 8.1. For more information about support lifecycles, see Microsoft Support Lifecycle.

Note
Central network policy management with RADIUS server and proxy, and 802.1X authenticated wired and wireless access are not deprecated

For the health policy creation, enforcement, and remediation features provided by NAP, as well as for monitoring, consider using System Center Configuration Manager to replace and enhance NAP's monitoring functionality:

- Windows Update status—System Center Configuration Manager integrates with Windows Server Update Services (WSUS) and has its own software updates feature. See Software Updates in Configuration Manager, in particular Monitor software updates.
- Windows Firewall settings—System Center Endpoint Protection can be used to manage Windows Firewall, see Endpoint Protection in Configuration Manager, in particular How to Create and Deploy Windows Firewall Policies for Endpoint Protection in Configuration Manager.
- Computer settings: registry settings, files, custom scripts—System Center Configuration Manager can help assess, track, and remediate the configuration compliance of client computers in the enterprise, see Compliance Settings in Configuration Manager.
- Reporting—System Center Configuration Manager also provides a set of tools and resources that help you use the advanced reporting capabilities of Microsoft SQL Server Reporting Services in the Configuration Manager console. See Reporting in Configuration Manager.
- Bitlocker—Microsoft BitLocker Administration and Monitoring (MBAM) lets you manage Bitlocker drive encryption throughout your enterprise, see Microsoft BitLocker Administration and Monitoring 2 Administrator's Guide.
- Managing mobile devices. See How to Manage Mobile Devices by Using Configuration Manager and Windows Intune and How to Manage Mobile Devices by Using Configuration Manager and Exchange.

To provide an always managed and always compliant experience for remote devices, you can use Remote Access, see Manage DirectAccess Clients Remotely. This way you can ensure the clients are always healthy, not only when they try to access resources in the corporate network.

NAP enables you to provide full internal network access to your users. However, if you only require the ability to provide access to specific applications and services in your internal network, you can use Web Application Proxy. Web Application Proxy enables you to provide this type of specific access to end users with domain-joined laptops or using their own devices; home computers, tablets, or personal smartphones. See Web Application Proxy Overview.

Server Manager information

The following role services can be installed with this role.

Role service	Description
Network Policy Server (NPS)	You can use NPS to centrally manage network access through a variety of network access servers, including RADIUS-compliant 802.1X-capable wireless access points, VPN servers, dial-up servers, and 802.1X-capable Ethernet switches. In addition, you can use NPS to deploy secure password authentication with Protected Extensible Authentication Protocol (PEAP)-MS-CHAP v2 for wireless connections. NPS also contains key components for deploying NAP on your network.
Health Registration Authority (HRA)	HRA is a NAP component that issues health certificates to clients that pass the health policy verification that is performed by NPS using the client SoH. HRA is used only with the NAP IPsec enforcement method.
Host Credential Authorization Protocol (HCAP)	HCAP allows you to integrate your Microsoft NAP solution with Cisco Network Access Control Server. When you deploy HCAP with NPS and NAP, NPS can perform client health evaluation and the authorization of Cisco 802.1X access clients.

Running Network Policy and Access Services

How do I deploy and configure Network Policy and Access Services using Windows PowerShell?

You can use to deploy and configure some aspects of Network Policy and Access Services. For more information about Windows PowerShell® cmdlets and scripts that you can use to deploy and manage Network Policy and Access Services, see Windows PowerShell for Network Policy and Access Services.

How do I deploy and configure this role in a multi-server environment?

You can deploy NPS servers for different functions. For example, you can deploy one NPS server as a RADIUS server for authentication, another as a RADIUS proxy, in order to distribute policy evaluation between servers with different roles, and another as a NAP policy server. For more information about multi-server management of Network Policy and Access Services, see Network Policy Server Overview.

Can I run this role on virtual machines?

Yes, you can run Network Policy and Access Services on Hyper-V virtual machines.

Can I run this role in a clustered environment?

No, Network Policy and Access Services cannot be run in a server cluster.

Special considerations for managing this role remotely

You can manage Network Policy and Access Services remotely. For more information about running Network Policy and Access Services from a remote computer, see Administer NPS by Using Tools.

Special considerations for managing the role on the Server Core installation option

You cannot install or run Network Policy and Access Services on the Server Core installation option of Windows Server 2012.

See also

The following table provides links to more content about Network Policy and Access Services.

Content type	References
Product evaluation	Network Access Protection Benefits
Planning	Network Access Protection Design Guide

Content type	References
Deployment	Deploying NPS | Checklist for deploying an HRA server | NAP Deployment Guide
Operations	NPS Operations Guide
Troubleshooting	Best Practices Analyzer for Network Policy and Access Services | Network Policy Server Infrastructure (Errors and Events) | NAP Infrastructure (Errors and Events) | Network Access Protection Troubleshooting Guide | Tools for Troubleshooting NAP | Troubleshooting HRA Guide
Tools and settings	Windows PowerShell for Network Policy and Access Services | Netsh Commands for Network Policy Server | Netsh Commands for Health Registration Authority | Netsh Commands for Network Access Protection (NAP) Client
Community resources	NAP Team Blog | NAP TechNet Forum | TechNet Wiki

Network Shell (Netsh) Overview

Network shell (netsh) is a command-line utility that allows you to configure and display the status of various network communications server roles and components after they are installed on computers running Windows Server® 2012.

Some client technologies, such as Network Access Protection (NAP) client and Dynamic Host Configuration Protocol (DHCP) client, also provide netsh commands that allow you to configure client computers running Windows® 8.

In most cases, netsh commands provide the same functionality that is available when using the Microsoft Management Console (MMC) snap-in for each server role or component. For example, you can configure Network Policy Server (NPS) by using either the NPS MMC snap-in or the netsh commands in the **netsh nps** context.

In addition, there are netsh commands for network technologies, such as for IPv6, network bridge, and remote procedure call (RPC), that are not available in Windows as an MMC snap-in.

Network Shell (Netsh) Technical Reference

The Netsh Technical Reference provides a comprehensive netsh command reference, including syntax, parameters, and examples for netsh commands. You can use the Netsh Technical Reference to build scripts and batch files by using netsh commands for local or remote management of network technologies on computers running Windows Server 2012.

Content availability

- Available for download in Windows Help format from the Microsoft Download Center: http://go.microsoft.com/fwlink/?LinkId=193003
- Available in HTML format in the Windows Server 2008 Technical Library: http://go.microsoft.com/fwlink/?LinkId=110582.

Formatting Legend

You can use the following formatting legend to understand the meaning of the text formatting used in the netsh command references.

Formatting legend

The text formatting described below details how to interpret the text formatting of all netsh commands in this reference so that you type the command syntax correctly when you run the command at the netsh prompt or in a batch file or script.

Format	Meaning
Italic	Information that the user must supply.
Bold	Elements that the user must type exactly as shown.
Ellipsis (...)	Parameter that can be repeated several times in a command line.
Between brackets ([])	Optional items.
Between braces ({}); choices separated by pipe (|). Example: {even|odd}	Set of choices from which the user must choose only one.
`Courier font`	Code or program output.

Netsh Commands for Trace

For information about Netsh Trace commands for computers running Windows® 7 and Windows Server® 2008 R2, see **Netsh Commands for Trace** in the Windows Server 2008 and Windows Server 2008 R2 technical library on TechNet.

On computers running Windows Server® 2012 and Windows® 8, network trace is a command context within Netsh that extends the Network Diagnostic Framework (NDF) to provide an efficient method to troubleshoot network connectivity issues. Network trace enables you to collect targeted diagnostic and event information by utilizing Event Tracing for Windows (ETW) to log network events and packets in a single file.

To run these commands from the command prompt, you must either enter the **netsh trace** context or prepend the context to the command. You can enter the **netsh trace** context by typing **netsh** and pressing ENTER, and then typing **trace** and pressing ENTER. Alternately, if you are at the command prompt but have not entered the trace context, you can type:

netsh trace *command*

Where *command* is the command that you want to run, including all of the required parameters for the command, and any optional parameters that you want to include.

Learn More

- **Netsh Commands for Network Trace in Windows Server 2008**

Netsh Commands for Network Trace

For information about Netsh Trace commands for computers running Windows® 7 and Windows Server® 2008 R2, see **Netsh Commands for Network Trace in Windows Server 2008 R2 and Windows 7** in the Windows Server 2008 and Windows Server 2008 R2 technical library on TechNet.

On computers running Windows Server® 2012 and Windows® 8, you can use the Netsh Trace context from a command prompt to enable and configure network tracing to assist you when troubleshooting network connectivity problems.

In this topic

- Using Netsh trace commands
- Commands in the Netsh trace context
  - convert
  - correlate
  - diagnose
  - show CaptureFilterHelp

- show globalKeywordsAndLevels
- show helperclass
- show interfaces
- show provider
- show providers
- show scenario
- show scenarios
- show status
- start
- stop

Using Netsh trace commands

This section provides an overview of Netsh trace usage.

Note

The examples in this section, Using Netsh trace commands, are documented as if you have entered the Netsh trace context. If you have not already entered the Netsh trace context, you can run these commands from the command prompt by prepending **Netsh trace** to each example.

Identifying scenarios

The Netsh trace context contains predefined sets of trace providers, known as scenarios, which you can enable for troubleshooting. To view a complete list of scenarios and a brief description of each scenario's purpose, type **show scenarios**.

Following is an example of the results that are rendered by running the Netsh trace **show scenarios** command:

```
netsh trace>show scenarios


Available scenarios (13):

----------------------------------------------------------------------

AddressAcquisition       : Troubleshoot address acquisition-related issues

DirectAccess             : Troubleshoot DirectAccess related issues

FileSharing              : Troubleshoot common file and printer sharing problems

InternetClient           : Diagnose web connectivity issues

InternetServer           : Troubleshoot server-side web connectivity issues

L2SEC                    : Troubleshoot layer 2 authentication related issues

LAN                      : Troubleshoot wired LAN related issues

Layer2                   : Troubleshoot layer 2 connectivity related issues
```

```
NDIS                    : Troubleshoot network adapter related issues
NetConnection           : Troubleshoot issues with network connections
RPC                     : Troubleshoot issues related to RPC framework
WFP-IPsec               : Troubleshoot Windows Filtering Platform and IPsec related
issues
WLAN                    : Troubleshoot wireless LAN related issues

netsh trace>
```

When you have identified the scenario that appears to be most relevant to your issues, you can see a list of all of the providers included in that scenario. For example, to see all of the providers enabled under the **LAN** scenario, type **show scenario lan**.

Following is an example of the results that are rendered by running the Netsh trace **show scenario lan** command:

```
netsh trace>show scenario lan

Description: Troubleshoot wired LAN related issues

Providers
---------

Provider:
Name:               Microsoft-Windows-L2NACP
Provider Guid:      {85FE7609-FF4A-48E9-9D50-12918E43E1DA}
Default Level:      4 (win:Informational)
Default Keywords:   0x0

Provider:
Name:               Microsoft-Windows-Wired-AutoConfig
Provider Guid:      {B92CF7FD-DC10-4C6B-A72D-1613BF25E597}
Default Level:      4 (win:Informational)
Default Keywords:   0x0

Provider:
Name:               Microsoft-Windows-NDIS
Provider Guid:      {CDEAD503-17F5-4A3E-B7AE-DF8CC2902EB9}
```

```
Default Level:      4 (win:Informational)
Default Keywords:   0x0

Provider:
Name:               Microsoft-Windows-EapHost
Provider Guid:      {6EB8DB94-FE96-443F-A366-5FE0CEE7FB1C}
Default Level:      4 (win:Informational)
Default Keywords:   0x0

Provider:
Name:               Microsoft-Windows-OneX
Provider Guid:      {AB0D8EF9-866D-4D39-B83F-453F3B8F6325}
Default Level:      4 (win:Informational)
Default Keywords:   0x0

netsh trace>
```

Obtaining trace provider details

To obtain a complete list of providers, you can type **show providers** from within the Netsh trace context. The **show providers** command lists – by name and GUID - every available provider, including providers that are not relevant to networking.

Following is a short excerpt of the list of providers that are rendered as a result of running the **show providers** command:

```
{0F67E49F-FE51-4E9F-B490-6F2948CC6027}  Microsoft-Windows-Kernel-Processor-Power
{122EE297-BB47-41AE-B265-1CA8D1886D40}  Microsoft-Windows-LoadPerf
{127E0DC5-E13B-4935-985E-78FD508B1D80}  Microsoft-Windows-TerminalServices-RdpSoundDriver
{134EA407-755D-4A93-B8A6-F290CD155023}  Microsoft-Windows-HomeGroup-ControlPanel
{13B197BD-7CEE-4B4E-8DD0-59314CE374CE}  Microsoft-Windows-Perflib
{1418EF04-B0B4-4623-BF7E-D74AB47BBDAA}  Microsoft-Windows-WMI-Activity
{155CB334-3D7F-4FF1-B107-DF8AFC3C0363}  Microsoft-Windows-MSDTC Client 2
{15A7A4F8-0072-4EAB-ABAD-F98A4D666AED}  Microsoft-Windows-Dhcp-Client
{15CA44FF-4D7A-4BAA-BBA5-0998955E531E}  Microsoft-Windows-Kernel-Boot
{17D6E590-F5FE-11DC-95FF-0800200C9A66}  Microsoft-Windows-ApplicationExperience-
SwitchBack
{17E92E2A-3D08-413E-BAEB-A79A262BF486}  Microsoft-Windows-MsiServer
```

```
{17F14A23-551D-40CC-A086-E4194D64ED4C}  Microsoft-Windows-WABSyncProvider

{18F4A5FD-FD3B-40A5-8FC2-E5D261C5D02E}  Microsoft-Windows-ApplicationExperience-
LookupServiceTrigger

{192EDE41-9175-4C86-AC02-9D003C9D43AB}  Microsoft-Windows-DisplaySwitch

{199FE037-2B82-40A9-82AC-E1D46C792B99}  LsaSrv
```

You can obtain details about a specific provider by typing **show provider** followed by the provider name. For example, to see details about the **Microsoft-Windows-Dhcp-Client** provider, type **show provider name=Microsoft-windows-dhcp-client**.

Following is an example of the results that are rendered by running the **show provider name=Microsoft-windows-dhcp-client** command:

```
netsh trace>show provider name=Microsoft-windows-dhcp-client


Name:          Microsoft-Windows-Dhcp-Client

Provider Guid: {15A7A4F8-0072-4EAB-ABAD-F98A4D666AED}


Levels (value) - 3 element(s)

----------------------------------------------------------------------

win:Error (2)

win:Warning (3)

win:Informational (4)


Keywords (mask) - 4 element(s)

----------------------------------------------------------------------

win:ResponseTime (0x1000000000000)

Microsoft-Windows-Dhcp-Client/Operational (0x8000000000000000)

Microsoft-Windows-Dhcp-Client/Admin (0x4000000000000000)

System (0x2000000000000000)


netsh trace>
```

Using filters to limit ETL trace file details

When troubleshooting, it is frequently beneficial to target tracing results by limiting irrelevant tracing details. For example, if you are running traces over an extended period of time, or if you are running tracing scenarios in which multiple providers are enabled, the resulting ETL trace file can become very large. If the ETL trace file becomes too large, the ETW buffer constraints may

cause some traces to be dropped. Additionally, reducing the amount of data in the ETL trace file can simplify troubleshooting by reducing the amount of data to review.

You can use Netsh trace filters to reduce the ETL trace file size. The trace filters that you can apply to individual providers are:

- **ETW levels**
- **ETW keywords**

The Netsh trace context also supports packet filtering capability that is similar to Network Monitor. See the Remarks section within the Netsh trace **start** command section in this topic for information about trace packet filter parameters and usage.

To see a list of filters which can be applied, type **show CaptureFilterHelp**.

The following example is an excerpt of the filter information that is displayed by running the **show CaptureFilterHelp** command.

```
netsh trace>show capturefilterhelp


  Capture Filters:

        Capture filters are only supported when capture is explicitly

        enabled with capture=yes. Supported capture filters are:


        CaptureInterface=<interface name or GUID>

         Enables packet capture for the specified interface name or GUID. Use

         'netsh trace show interfaces' to list available interfaces.

        e.g. CaptureInterface={716A7812-4AEE-4545-9D00-C10EFD223551}

        e.g. CaptureInterface=!{716A7812-4AEE-4545-9D00-C10EFD223551}

        e.g. CaptureInterface="Local Area Connection"


        Ethernet.Address=<MAC address>

         Matches the specified filter against both source and destination

         MAC addresses.

        e.g. Ethernet.Address=00-0D-56-1F-73-64


         Wifi.Type=<Management|Data>

         Matches the specified filter against the Wifi type. Allowed values

         are 'Management' and 'Data'. If not specified, the Wifi.Type filter

         is not applied.

         Note: This capture filter does not support ranges, lists or negation.
```

```
        e.g. Wifi.Type=Management


        Protocol=<protocol>
         Matches the specified filter against the IP protocol.
        e.g. Protocol=6
        e.g. Protocol=!(TCP,UDP)
        e.g. Protocol=(4-10)
netsh trace>
```

Example Netsh trace filter parameters and usage

Following is an example start command for Netsh trace that includes filter parameters.

start InternetClient provider=Microsoft-Windows-TCPIP level=5 keywords=ut:ReceivePath,ut:SendPath

In this example, the level is set to 5, which means that 5 events is the maximum number of events that will be shown. The keywords **ut:ReceivePath** and **ut:SentPath** filter the events to show only the events that are traced on the receive or send path.

Note

255 is the maximum that is supported by ETW. ETW maximums can vary by provider, within the range of 0 – 255. In the example, the value of **5** is the highest level supported by Microsoft-Windows-TCPIP.

See the Remarks section within the Netsh trace start command in this topic for additional details concerning trace filter parameters and usage.

Starting and stopping trace

You can start a trace for all of the providers in a given scenario. For example, to start a trace for all of the providers enabled under the InternetClient scenario, type **start scenario=internetclient**. You can also start a trace for additional providers not included in that particular scenario. For example, you might want to start traces for all of the providers enabled under the WLAN scenario and also for the DHCP provider. To do this, you can type **start scenario=wlan provider=Microsoft-Windows-Dhcp-Client**.

To stop tracing, type **stop** from within the Netsh trace context.

Using the files rendered by trace

When tracing is stopped, two files are generated by default:

- An **Event Trace Log (ETL)** file. The ETL file collects trace events, which can be viewed using tools such as Network Monitor. By default, the ETL is saved as **nettrace.etl**.

  Note

You can specify a different name by including the **tracefile=***FileName***.etl** parameter when starting the trace, where *FileName* is the name you specify for the saved .etl file.

- A **.cab** file. The .cab file contains detailed information about the software and hardware on the system, such as the adapter information, build, operating system, and wireless settings. By default, the .cab file will be named nettrace.cab by default, unless another name was specified as indicated above. The .cab file will contain two files, which will always be named as follows:
  - **Report.etl**. Report.etl is a duplicate copy of the information included in nettrace.etl.
  - **Report.html**. The report.html file includes additional information about the trace events, as well as other information that was collected by trace.

    Tip
    To receive the most detailed information, include the command parameter **report=yes** when starting a trace.

Commands in the Netsh trace context

This section provides usage details for commands in the Netsh trace context.

convert

Converts a trace file in ETL format into the specified format.

Syntax

convert input=*Tracefilename.etl* [[**output=**]*FileName*] [[**dump=**]{**csv**|**xml**|**evtx**|**txt**|**no**}] [[**report=**]{**yes**|**no**}] [[**overwrite=**]{**yes**|**no**}] [[**tmfpath=**]*PathAndFileName*]

Parameters

input

Required. Specifies the input ETL trace file.

output

Optional. Specifies the name of the converted file. If unspecified, the output file name defaults to the input name.

dump

Optional. Specifies the file format of the converted trace file: .CSV .XML .EVTX .TXT No. If unspecified, the converted file output defaults to .txt.

Report

Optional. Specifies whether to generate an HTML report. If unspecified, the parameter defaults to “no.”

overwrite

Optional. Specifies whether this instance of the trace conversion command overwrites files that were rendered from previous trace conversions. If unspecified, the parameter defaults to “no.”

tmfpath

Optional. Specifies the path (folder location) of the tmf files that are used for decoding WPP traces.

Example command

- **convert input=ExampleTestTrace.etl output=TestTrace dump=XML report=yes**

correlate

Normalizes or filters a trace file to a new output file.

Syntax

correlate input=*tracefilename.etl* **output=***newtracefilename.etl* [[**filter=**]*Activity_ID*] [[**overwrite=**]{**yes**|**no**}] [[**retaincorrelationevents=**]{**yes**|**no**}] [[**retainpii=**]{**yes**|**no**}] [[**retainglobalevents=**]{**yes**|**no**}]

Parameters

input

Required. Specifies the input trace file name.

output

Required. Specifies the file name of the output file.

filter

Optional. Specifies that trace output is rendered only for events that are related to the specific Activity GUID. If unspecified, the default entry for filter is null; the filter parameter is ignored.

overwrite

Optional. Overwrites existing files. If unspecified, the default entry for overwrite is no.

retaincorrelationevents

Optional. Retains correlation events. If unspecified, the default entry for retaincorrevents is no.

retainpii

Optional. Specifies that events containing personally identifiable information (pii) are kept. If unspecified, the default entry for retainpii is yes.

retainglobalevents

Optional. Specifies that global events are kept. If unspecified, the default entry for retainglobalevents is yes.

Remarks

Filter activity ID is a GUID expressed in the form {xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx}.

Personally identifiable information includes packet capture events.

Example command

- **correlate input=exampleTraceFileName.etl output=NewTraceFileName.etl overwrite=yes retaincorrelationevents=no retainpii=yes retainglobalevents=yes**

diagnose

Starts a diagnostic session.

Syntax

Diagnose scenario=*ScenarioName* **namedAttribute=***attributeValue* **[[saveSessionTrace=]{yes|no}] [[report=]yes|no] [[capture=]yes|no]**

Parameters

scenario

Required. Specifies the name of the scenario.

namedAttribute

Required. Specifies the attribute value.

Tip

namedAttributes change depending on the scenario, and can be discovered by running the Netsh Trace **show scenario** command.

saveSessionTrace

Optional. If parameter saveSessionTrace is unspecified and parameter report is set to yes, then the default entry for saveSessionTrace is yes. If parameter saveSessionTrace

is unspecified and parameter report is set to no or is unspecified, then the default entry for saveSessionTrace is no.

report

Optional. If unspecified, the default entry for report is no.

capture

Optional. If unspecified, the default entry for capture is no.

Remarks

Diagnose command behavior for parameters capture and report:

The following table summarizes the results that are rendered by the use of the various combinations of the capture and report parameters:

capture=	report=	Results:
not specified	not specified	The diagnostic results will display, but neither the ETL nor CAB files are created.
yes	no	The diagnostic results will display, and only ETL file is created. The ETL file is created with both traces and packets.
yes	yes	The diagnostic results will display, and both the ETL and CAB files are saved. The ETL file is saved with both traces and packets.
no	yes	The diagnostic results will display, and both ETL and CAB files are created. Only events are saved in the ETL file; no packets are saved.
no	no	The diagnostic results will display, ETL file is created with trace events only. The CAB file is not created.

1. If parameters capture and report are not specified the diagnostic results will display, but neither the ETL nor CAB files are saved.
2. If parameter capture=yes and parameter report=no the diagnostic results will display, and only ETL file is saved. The ETL file is saved with both traces and packets.
3. If parameter capture=yes and parameter report=yes, then the diagnostic results will display, and both the ETL and CAB files are saved. The ETL file is saved with both traces and packets.
4. If parameter capture=no and parameter report=yes, then the diagnostic results will display, and both ETL and CAB files are saved. Only events are saved in the ETL file; no packets are saved.
5. If parameter capture=no and parameter report=no, the diagnostic results will display, ETL file is saved with trace events only. The CAB file is not saved.

Tip

In all cases, if the ETL file is saved and capture=yes, ETL file will contain both packet captures and trace events

Example command

- **Diagnose scenario=internetclient namedAttribute=http://www.Microsoft.com saveSessionTrace=yes report=no**

show CaptureFilterHelp

Displays the supported capture filters and filter usage.

Syntax

show CaptureFilterHelp

Parameters

There are no parameters associated with this command

Example command

- **show CaptureFilterHelp**

show globalKeywordsAndLevels

Displays a list of global keywords and the levels that may be used with the start command.

Syntax

show globalKeywordsAndLevels

Parameters

There are no parameters associated with this command.

Remarks

Following is an example of the global levels and keywords that are displayed by running the command **show globalKeywordsAndLevels**.

```
netsh trace>show globalkeywordsandlevels
```

```
Levels (value) - 5 element(s)
------------------------------------------------------------------------
win:Critical (1)
win:Error (2)
win:Warning (3)
win:Informational (4)
win:Verbose (5)

Keywords (mask) - 15 element(s)
------------------------------------------------------------------------
ut:SendPath (0x100000000)
ut:ReceivePath (0x200000000)
ut:L3ConnectPath (0x400000000)
ut:L2ConnectPath (0x800000000)
ut:ClosePath (0x1000000000)
ut:Authentication (0x2000000000)
ut:Configuration (0x4000000000)
ut:Global (0x8000000000)
ut:Dropped (0x10000000000)
ut:PiiPresent (0x20000000000)
ut:Packet (0x40000000000)
ut:Address (0x80000000000)
ut:StdTemplateHint (0x100000000000)
ut:StateTransition (0x200000000000)
ut:Diagnostic (0x400000000000)

netsh trace>
```

You can specify the global levels and keywords in the start command as shown in the following example of the start command: **start InternetClient provider=Microsoft-Windows-TCPIP level=5 keywords=ut:ReceivePath,ut:SendPath**. This is explained in greater detail in the section: **Example Netsh trace filter parameters and usage**

Example command

- **show globalKeywordsAndLevels**

show helperclass

Displays the helper class name, description and lists the dependent helper classes. Additionally, this command also lists possible root causes and repairs that can be returned by the helper class.

Syntax

show helperclass name=*helperClassName*

Parameters

name

Required. Specifies the name of the helperclass for which to display information.

Remarks

Helper class names are case sensitive. For example, you cannot substitute wininethelperclass, if WinInetHelperClass is the correct, case-sensitive name.

Example commands

- **Show helperclass name=WinInetHelperClass**

show interfaces

Displays a list of network interfaces.

Syntax

show interfaces

Parameters

There are no parameters associated with this command.

Remarks

Example command

- **show interfaces**

show provider

For the provider specified, displays the keywords and levels that can be used with the start command.

Syntax

show provider name=*providerIdOrName*

Parameters

name

Required. Specifies the provider name or provider identification (Id) GUID.

Remarks

This displays the keywords and levels that the specified provider supports, which can be used when starting a tracing session.

Example commands

- **show provider name=Microsoft-Windows-TCPIP**
- **show provider name=B92CF7FD-DC10-4C6B-A72D-1613BF25E597**

show providers

Displays a list of the providers that can be used with the start command.

Syntax

show providers

Parameters

There are no parameters associated with this command.

Example command

- **Show providers**

show scenario

Displays information for the specified scenario.

Syntax

show scenario name=*scenarioName*

Parameters

name

Required. Specifies the name of the scenario for which to display information.

Remarks

This command displays information about a scenario, including:

- Attributes that might be required by the diagnose command.
- The list of providers that would be enabled for the specified scenario.
- The default keywords and levels that would be applied for the specified scenario.

Example command

- **show scenario name=DirectAccess**

show scenarios

Displays a list of the scenarios that can be used with the start command.

Syntax

show scenarios

Parameters

There are no parameters associated with this command.

Example command

- **show scenarios**

show status

Displays the current tracing configuration.

Syntax

show status

Parameters

There are no parameters associated with this command.

Example command

- **show status**

start

Starts a network tracing session.

Syntax

start [[**scenario=**]*Scenario1,Scenario2*] [[**globalKeywords=**]*keywords*] [[**globalLevel**=]*level*] [[**capture=**]{**yes**|**no**}][[**capturetype=***physical*|*vmswitch*|*both*] [[**report=**]{**yes**|**no**}] [[**persistent=**]{**yes**|**no**}] [[**traceFile=**]*Path\Filename*] [[**maxSize=**]*MaxFileSizeInMB*] [[**fileMode=**]{**single**|**circular**|**append**}] [[**overwrite=**]{**yes**|**no**}] [[**correlation=**]{**yes**|**no**|**disabled**}] [[**provider=**]*ProviderIdOrName*] [[**keywords=**]*KeywordMaskOrSet*] [[**level=**]*level*] [[**provider=**]*Provider2IdOrName*] [[**keywords=**]*Keyword2MaskOrSet*] [[**level=**]*level2*]

Parameters

scenario

Optional. Specifies the scenarios that are applied to the trace session.

globalKeywords

Optional. Specifies the global keywords that are applied to the trace session.

globalLevel

Optional. Species at what level the traces should be collected across all the providers. For example, the command **netsh trace start scenario=internetclient globallevel=win:error** will capture only error events; the command **netsh trace show globalkeywordsandlevel** displays list of options.

capture

Optional. Specifies whether packet capture is enabled in addition to trace events. If unspecified, the default entry for capture is no.

capturetype

Specifies whether packet capture needs to be enabled for physical network adapters

only, virtual switch only, or both network adapters and virtual switch.

report

Optional. Specifies whether a complementing report will be generated in addition to the trace file report. If unspecified, the default entry for report is no.

persistent

Optional. Specifies whether the tracing session resumes upon restarting the computer, and continues to function until the “Netsh trace stop” command is issued. If unspecified, the default entry for persistent is no.

traceFile

Optional. Specifies the location to save the output file. If unspecified, the output file is saved using the default path and filename: “traceFile=%LOCALAPPDATA%\Temp\NetTraces\NetTrace.etl”

maxSize

Optional. Specifies the maximum size in MB for saved trace files. If the maxSize parameter is unspecified, the default value of 250 MB is used. MB 0=no maximum. To specify the maxSize=0 for the no tracing file max size option, you must also specify filemode=single. By default, filemode=circular.

fileMode

Optional. Specifies which file mode is applied when tracing output is generated: single|circular|append. If unspecified, the default entry “fileMode=circular” is used.

overwrite

Optional. Specifies whether an existing trace output file – in either the specified or default location - will be overwritten. If parameter traceFile is not specified, then the default location and filename is applied, and any pre-existing version of the trace file is automatically overwritten. See parameter traceFile for default path and filename details.

correlation

Optional. Specifies whether related events will be correlated and grouped together. If unspecified, the default entry “correlation=yes” is applied.

provider

Optional. Specifies additional providers to enable for tracing.

keywords

Optional. Specifies the keywords to enable for the corresponding provider.

level

Optional. Specifies the level to enable for the corresponding provider.

Remarks

Capture filters are only supported when capture is explicitly enabled with capture=yes. You can use the **show CaptureFilterHelp** command to display the list of supported capture filters and their usage.

Follow is a list and brief explanation of the ETW levels that you can specify by using the level parameter.

1. **Critical**. Captures only critical events.
2. **Errors**. Captures only critical events and error events.
3. **Warnings**. Captures only critical events, error events, and warning events.
4. **Informational**. Captures only critical events, error events, warning events, and informational events.
5. **Verbose** Captures all events.

You can view a complete list of keywords for a specific provider by typing **show provider**, followed by the provider name. For example, you can type **show provider Microsoft-Windows-TCPIP** to display information about the Microsoft-Windows-TCPIP provider, including the list of supported keywords.

Example commands

- **start scenario=InternetClient capture=yes**
  - Starts tracing for the InternetClient scenario and dependent providers with packet capture enabled for physical network adapters only.
  - Tracing will stop when the "Netsh trace stop" command is issued or when the system reboots.
  - Default location and name will be used for the output file. If an old file exists, it will be overwritten.
- **start provider=Microsoft-windows-wlan-autoconfig keywords=state,ut:authentication**
  - Starts tracing for the Microsoft-windows-wlan-autoconfig provider.
  - Tracing will stop when the "Netsh trace stop" command is issued or when the system reboots.
  - Default location and name will be used for the output file.

- If an old file exists, it will be overwritten.
- Trace will only log events with the keywords 'state' or 'ut:authentication.'

  You can use the Netsh trace show provider command to display the supported keywords and levels.

The Netsh trace context also supports packet filtering capability that is similar to Network Monitor. You can turn on packet capturing by specifying **capture=yes** in the Netsh trace start command. You can use packet filtering to capture specific packets in a trace file. For example, you can type **start capture=yes ipv4.address=***x.x.x.x*, (where x.x.x.x is the IP address), to only capture only those packets containing ipv4 traffic using that specific source or destination address.

To obtain additional information about how to use packet filtering, you can type **show capturefilterHelp**.

stop

Stops a network tracing session that is currently in progress.

Syntax

stop

Parameters

There are no parameters associated with this command.

Example command

- **stop**

Network Subsystem Performance Tuning

Performance tuning the network subsystem, particularly for network intensive workloads, can involve each layer of the network architecture, which is also called the network stack. These layers are broadly divided into the following sections.

1. **Network interface**. This is the lowest layer in the network stack, and contains the network driver that communicates directly with the network adapter.
2. **Network Driver Interface Specification (NDIS)**. NDIS exposes interfaces for the driver below it and for the layers above it, such as the Protocol Stack.
3. **Protocol Stack**. The protocol stack implements protocols such as TCP/IP and UDP/IP. These layers expose the transport layer interface for layers above them.
4. **System Drivers**. These are typically clients that use a transport data extension (TDX) or Winsock Kernel (WSK) interface to expose interfaces to user-mode applications. The WSK interface was introduced in Windows Server® 2008 and Windows Vista, and it is exposed by AFD.sys. The interface improves performance by eliminating the switching between user mode and kernel mode.
5. **User-Mode Applications**. These are typically Microsoft solutions or custom applications.

The table below provides a vertical illustration of the layers of the network stack, including examples of items that run in each layer.

5	**User-Mode Applications**	WMS	DNS	IIS
4	**System Drivers**	AFD.sys	HTTP.sys	
3	**Protocol Stack**	TCP/IP	UDP/IP	VPN
2	**NDIS**	Network Driver Interface Specification (NDIS)		
1	**Network interface**	Network driver		

The following sections of this guide provide performance tuning recommendations for network devices and the network stack.

- Choosing a Network Adapter
- Performance Tuning Network Adapters
- Network-Related Performance Counters
- Performance Tools for Network Workloads

Choosing a Network Adapter

This topic contains the following sections.

- Offload Capabilities
- Receive Side Scaling (RSS)
- Receive Segment Coalescing (RSC)
- Network Adapter Resources
- Message-Signaled Interrupts (MSI/MSI-X)
- Interrupt Moderation
- Suggested Network Adapter Features for Server Workloads

Network-intensive applications require high-performance network adapters. This section explores some considerations for choosing network adapters, as well as how to configure different network adapter settings to achieve the best network performance.

Tip

You can configure network adapter settings by using Windows PowerShell. For more information, see Network Adapter Cmdlets in Windows PowerShell.

Offload Capabilities

Offloading tasks from the central processing unit (CPU) to the network adapter can reduce CPU usage on the server, which improves the overall system performance. The network stack in Microsoft products can offload one or more tasks to a network adapter if you select a network adapter that has the appropriate offload capabilities. The following table provides details about different offload capabilities that are available in Windows Server 2012.

Offload type	Description
Checksum calculation	The network stack can offload the calculation and validation of Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) checksums on send and receive code paths. It can also offload the calculation and validation of IPv4 and IPv6 checksums on send and receive code paths.
IP security authentication and encryption	The TCP/IP transport layer can offload the calculation and validation of encrypted checksums for authentication headers and Encapsulating Security Payloads (ESPs). The TCP/IP transport layer can also offload the encryption and decryption of ESPs.
Segmentation of large TCP packets	The TCP/IP transport layer supports Large Send Offload v2 (LSOv2). With LSOv2, the TCP/IP transport layer can offload the segmentation of large TCP packets to the network adapter.
Receive Side Scaling (RSS)	RSS is a network driver technology that enables the efficient distribution of network receive processing across multiple CPUs in multiprocessor systems.
Receive Segment Coalescing (RSC)	RSC is the ability to group packets together to minimize the header processing that is necessary for the host to perform. A maximum of 64 KB of received payload can be coalesced into a single larger packet for processing.

Receive Side Scaling (RSS)

Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 support Receive Side Scaling (RSS). Some servers are configured with multiple logical processors that share hardware resources (such as a physical core) and which are treated as Simultaneous Multi-Threading (SMT) peers. Intel Hyper-Threading Technology is an example. RSS directs network processing to up to one logical processor per core. For example, on a server with Intel Hyper-Threading, 4 cores, and 8 logical processors, RSS uses no more than 4 logical processors for network processing.

RSS distributes incoming network I/O packets among logical processors so that packets which belong to the same TCP connection are processed on the same logical processor, which preserves ordering. RSS also load balances UDP unicast and multicast traffic, and it routes related flows (which are determined by hashing the source and destination addresses) to the same logical processor, preserving the order of related arrivals. This helps improve scalability and performance for receive-intensive scenarios for servers that have fewer network adapters than they do eligible logical processors.

Configuring RSS

In Windows Server 2012, you can configure RSS by using Windows PowerShell cmdlets and RSS profiles. You can define RSS profiles by using the **–Profile** parameter of the Set-NetAdapterRss Windows PowerShell cmdlet.

Windows PowerShell commands for RSS configuration

The following cmdlets allow you to see and modify RSS parameters per network adapter.

Note

For a detailed command reference for each cmdlet, including syntax and parameters, you can click the following links. In addition, you can pass the cmdlet name to Get-Help at the Windows PowerShell prompt for details on each command.

- Disable-NetAdapterRss. This command disables RSS on the network adapter that you specify.
- Enable-NetAdapterRss. This command enables RSS on the network adapter that you specify.
- Get-NetAdapterRss. This command retrieves RSS properties of the network adapter that you specify.
- Set-NetAdapterRss. This command sets the RSS properties on the network adapter that you specify.

RSS profiles

You can use the **–Profile** parameter of the Set-NetAdapterRss cmdlet to specify which logical processors are assigned to which network adapter. Available values for this parameter are:

- **Closest**. Logical processor numbers that are near the network adapter's base RSS processor are preferred. With this profile, the operating system might rebalance logical processors dynamically based on load.

- **ClosestStatic**. Logical processor numbers near the network adapter's base RSS processor are preferred. With this profile, the operating system does not rebalance logical processors dynamically based on load.
- **NUMA**. Logical processor numbers are generally selected on different NUMA nodes to distribute the load. With this profile, the operating system might rebalance logical processors dynamically based on load.
- **NUMAStatic**. This is the **default profile**. Logical processor numbers are generally selected on different NUMA nodes to distribute the load. With this profile, the operating system will not rebalance logical processors dynamically based on load.
- **Conservative**. RSS uses as few processors as possible to sustain the load. This option helps reduce the number of interrupts.

Depending on the scenario and the workload characteristics, you can also use other parameters of the Set-NetAdapterRss Windows PowerShell cmdlet to specify the following:

- On a per-network adapter basis, how many logical processors can be used for RSS.
- The starting offset for the range of logical processors.
- The node from which the network adapter allocates memory.

Following are the additional Set-NetAdapterRss parameters that you can use to configure RSS:

Note

In the example syntax for each parameter below, the network adapter name Ethernet is used as an example value for the **–Name** parameter of the Set-NetAdapterRss command. When you run the cmdlet, ensure that the network adapter name that you use is appropriate for your environment.

- *** MaxProcessors**: Sets the maximum number of RSS processors to be used. This ensures that application traffic is bound to a maximum number of processors on a given interface. Example syntax:

```
Set-NetAdapterRss –Name "Ethernet" –MaxProcessors <value>
```

- *** BaseProcessorGroup**: Sets the base processor group of a NUMA node. This impacts the processor array that is used by RSS. Example syntax:

```
Set-NetAdapterRss –Name "Ethernet" –BaseProcessorGroup <value>
```

- *** MaxProcessorGroup**: Sets the Max processor group of a NUMA node. This impacts the processor array that is used by RSS. Setting this would restrict a maximum processor group so that load balancing is aligned within a k-group. Example syntax:

```
Set-NetAdapterRss –Name "Ethernet" –MaxProcessorGroup <value>
```

- *** BaseProcessorNumber**: Sets the base processor number of a NUMA node. This impacts the processor array that is used by RSS. This allows partitioning processors across network adapters. This is the first logical processor in the range of RSS processors that is assigned to each adapter. Example syntax:

```
Set-NetAdapterRss –Name "Ethernet" –BaseProcessorNumber <Byte Value>
```

- *** NumaNode**: The NUMA node that each network adapter can allocate memory from. This can be within a k-group or from different k-groups. Example syntax:

```
Set-NetAdapterRss –Name "Ethernet" –NumaNodeID <value>
```

- *** NumberofReceiveQueues**: If your logical processors seem to be underutilized for receive traffic (for example, as viewed in Task Manager), you can try increasing the number of RSS queues from the default of 2 to the maximum that is supported by your network adapter. Your network adapter may have options to change the number of RSS queues as part of the driver. Example syntax:

  ```
  Set-NetAdapterRss –Name “Ethernet” –NumberOfReceiveQueues <value>
  ```

For more information, click the following link to download Scalable Networking: Eliminating the Receive Processing Bottleneck—Introducing RSS in Word format.

Understanding RSS Performance

Tuning RSS requires understanding the configuration and the load-balancing logic. To verify that the RSS settings have taken effect, you can review the output when you run the **Get-NetAdapterRss** Windows PowerShell cmdlet. Following is example output of this cmdlet.

```
PS C:\Users\Administrator> get-netadapterrss

Name                                            : testnic 2

InterfaceDescription                  : Broadcom BCM5708C NetXtreme II GigE (NDIS VBD
Client) #66

Enabled                                         : True

NumberOfReceiveQueues         : 2

Profile                                             : NUMAStatic

BaseProcessor: [Group:Number]                   : 0:0

MaxProcessor: [Group:Number]                    : 0:15

MaxProcessors                                 : 8


IndirectionTable: [Group:Number]                :

     0:0    0:4    0:0    0:4    0:0    0:4    0:0    0:4

…

(# indirection table entries are a power of 2 and based on # of processors)

…

                                                  0:0    0:4    0:0    0:4    0:0    0:4
0:0    0:4
```

In addition to echoing parameters that were set, the key aspect of the output is the indirection table output. The indirection table displays the hash table buckets that are used to distribute incoming traffic. In this example, the n:c notation designates the Numa K-Group:CPU index pair

that is used to direct incoming traffic. We see exactly 2 unique entries (0:0 and 0:4), which represent k-group 0/cpu0 and k-group 0/cpu 4, respectively.

There is only one k-group for this system (k-group 0) and a n (where n <= 128) indirection table entry. Because the number of receive queues is set to 2, only 2 processors (0:0, 0:4) are chosen - even though maximum processors is set to 8. In effect, the indirection table is hashing incoming traffic to only use 2 CPUs out of the 8 that are available.

To fully utilize the CPUs, the number of RSS Receive Queues must be equal to or greater than Max Processors. In the previous example, the Receive Queue should be set to 8 or greater.

RSS and Virtualization

RSS provides hashing and scalability to host interface only. RSS does not provide any interaction with virtual machines, instead users can configure VMQ in those scenarios.

RSS can be enabled for guest virtual machines in the case of SR-IOV because the virtual function driver supports RSS capability. In this case, the guest and the host will have the benefit of RSS. Note that the host does not get RSS capability because the virtual switch is enabled with SR-IOV.

NIC Teaming and RSS

RSS can be enabled on a network adapter that is teamed with another network interface card using NIC Teaming. In this scenario, only the underlying physical network adapter can be configured to use RSS. A user cannot set RSS cmdlets on the teamed network adapter.

Receive Segment Coalescing (RSC)

Receive Segment Coalescing (RSC) helps performance by reducing the number of IP headers that are processed for a given amount of received data. It should be used to help scale the performance of received data by grouping (or coalescing) the smaller packets into larger units.

This approach can affect latency with benefits mostly seen in throughput gains. RSC is recommended to increase throughput for received heavy workloads. Consider deploying network adapters that support RSC. On these network adapters, ensure that RSC is on (this is the default setting), unless you have specific workloads (for example, low latency, low throughput networking) that show benefit from RSC being off.

In Windows Server 2012, the following Windows PowerShell cmdlets allow you to configure RSC capable network adapters:

Understanding RSC Diagnostics

You can diagnose RSC by using the Windows PowerShell cmdlets Get-NetAdapterRsc and Get-NetAdapterStatistics.

Following is example output when you run the Get-NetAdapterRsc cmdlet.

```
PS C:\Users\Administrator> Get-NetAdapterRsc
```

```
Name                      IPv4Enabled  IPv6Enabled  IPv4Operational IPv6Operational
IPv4FailureReason             IPv6Failure
                                          Reason
----                      -----------  -----------  --------------- ---------------
----------------- ------------
Ethernet                  True         False        True            False
NoFailure       NicProperties
```

The **Get** cmdlet shows whether RSC is enabled in the interface and whether TCP enables RSC to be in an operational state. The failure reason provides details about the failure to enable RSC on that interface.

In the previous scenario, IPv4 RSC is supported and operational in the interface. To understand diagnostic failures, one can see the coalesced bytes or exceptions caused. This provides an indication of the coalescing issues.

Following is example output when you run the Get-NetAdapterStatistics cmdlet.

```
PS C:\Users\Administrator> $x = Get-NetAdapterStatistics “myAdapter”
PS C:\Users\Administrator> $x.rscstatistics


CoalescedBytes       : 0
CoalescedPackets     : 0
CoalescingEvents     : 0
CoalescingExceptions : 0
```

RSC and Virtualization

RSC is only supported in the physical host when the host network adapter is not bound to the Hyper-V Virtual Switch. RSC is disabled by the operating system when the host is bound to the Hyper-V Virtual Switch. In addition, virtual machines do not get the benefit of RSC because virtual network adapters do not support RSC.

RSC can be enabled for a virtual machine when Single Root Input/Output Virtualization (SR-IOV) is enabled. In this case, virtual functions support RSC capability; hence, virtual machines also receive the benefit of RSC.

Network Adapter Resources

A few network adapters actively manage their resources to achieve optimum performance. Several network adapters let the administrator manually configure resources by using the

Advanced Networking tab for the adapter. For such adapters, you can set the values of a number of parameters, including the number of receive buffers and send buffers.

In Windows Server 2012, configuring network adapter resources is simplified by the use of the following Windows PowerShell cmdlets.

- Get-NetAdapterAdvancedProperty
- Set-NetAdapterAdvancedProperty
- Enable-NetAdapter
- Enable-NetAdapterBinding
- Enable-NetAdapterChecksumOffload
- Enable-NetAdapterLso
- Enable-NetAdapterIPSecOffload
- Enable-NetAdapterPowerManagement
- Enable-NetAdapterQos
- Enable-NetAdapterRDMA
- Enable-NetAdapterSriov

Message-Signaled Interrupts (MSI/MSI-X)

Network adapters that support MSI/MSI-X can target their interrupts to specific logical processors. If the adapters also support RSS, then a logical processor can be dedicated to servicing interrupts and deferred procedure calls (DPCs) for a given TCP connection. This preserves the cache locality of TCP structures and greatly improves performance.

Interrupt Moderation

To control interrupt moderation, some network adapters expose either different interrupt moderation levels, or buffer coalescing parameters (sometimes separately for send and receive buffers), or both of these. You should consider buffer coalescing or batching when the network adapter does not perform interrupt moderation. Interrupt moderation helps reduce overall CPU utilization by minimizing the per-buffer processing cost, but the moderation of interrupts and buffer batching can have a negative impact on latency-sensitive scenarios.

Suggested Network Adapter Features for Server Workloads

The following table lists high-performance network adapter features – including Checksum Offload (ChO), Large Sum Offload (LSO), RSS, and RSC - that can improve performance in terms of throughput, latency, or scalability for some server workloads.

Note

When there is an **X** in a feature column, it means that for the server role, installing a network adapter with this feature enabled is recommended.

Server workload	ChO	LSO	RSS	RSC
File server	**X**	**X**	**X**	**X**
Web server	**X**	**X**	**X**	
Mail server (short-lived connections)	**X**		**X**	
Database server	**X**	**X**	**X**	
FTP server	**X**	**X**		**X**
Media server	**X**		**X**	**X**

Performance Tuning Network Adapters

This topic contains the following sections.

- Enabling Offload Features
- Enabling Receive Side Scaling (RSS) for Web Servers
- Increasing Network Adapter Resources
- Performance Tuning for Low Latency Packet Processing
- System Management Interrupts
- Performance Tuning TCP

Determining the correct tuning settings for your network adapter depend on the following variables:

- The network adapter and its feature set
- The type of workload performed by the server
- The server hardware and software resources
- Your performance goals for the server

If your network adapter provides tuning options, you can optimize network throughput and resource usage to achieve optimum throughput based on the parameters described above.

The following sections describe some of your performance tuning options.

Enabling Offload Features

Turning on network adapter offload features is usually beneficial. Sometimes, however, the network adapter is not powerful enough to handle the offload capabilities with high throughput. For example, enabling segmentation offload can reduce the maximum sustainable throughput on

some network adapters because of limited hardware resources. However, if the reduced throughput is not expected to be a limitation, you should enable offload capabilities, even for this type of network adapter.

Note
Some network adapters require offload features to be independently enabled for send and receive paths.

Enabling Receive Side Scaling (RSS) for Web Servers

RSS can improve web scalability and performance when there are fewer network adapters than logical processors on the server. When all the web traffic is going through the RSS-capable network adapters, incoming web requests from different connections can be simultaneously processed across different CPUs.

It is important to note that due to the logic in RSS and Hypertext Transfer Protocol (HTTP) for load distribution, performance might be severely degraded if a non-RSS-capable network adapter accepts web traffic on a server that has one or more RSS-capable network adapters. In this circumstance, you should use RSS-capable network adapters or disable RSS on the network adapter properties **Advanced Properties** tab. To determine whether a network adapter is RSS-capable, you can view the RSS information on the network adapter properties **Advanced Properties** tab.

RSS Profiles and RSS Queues

RSS predefined profiles are new in Windows Server 2012.

The default profile is NUMA Static, which changes the default behavior from previous versions of the operating system. To get started with RSS Profiles, you can review the available profiles to understand when they are beneficial and how they apply to your network environment and hardware.

For example, if you open Task Manager and review the logical processors on your server, and they seem to be underutilized for receive traffic, you can try increasing the number of RSS queues from the default of 2 to the maximum that is supported by your network adapter. Your network adapter might have options to change the number of RSS queues as part of the driver.

Increasing Network Adapter Resources

For network adapters that allow manual configuration of resources, such as receive and send buffers, you should increase the allocated resources. Some network adapters set their receive buffers low to conserve allocated memory from the host. The low value results in dropped packets and decreased performance. Therefore, for receive-intensive scenarios, we recommend that you increase the receive buffer value to the maximum.

Note

If a network adapter does not expose manual resource configuration, it either dynamically configures the resources, or the resources are set to a fixed value that cannot be changed.

Enabling Interrupt Moderation

To control interrupt moderation, some network adapters expose different interrupt moderation levels, buffer coalescing parameters (sometimes separately for send and receive buffers), or both.

You should consider interrupt moderation for CPU-bound workloads, and consider the trade-off between the host CPU savings and latency versus the increased host CPU savings because of more interrupts and less latency. If the network adapter does not perform interrupt moderation, but it does expose buffer coalescing, increasing the number of coalesced buffers allows more buffers per send or receive, which improves performance.

Performance Tuning for Low Latency Packet Processing

Many network adapters provide options to optimize operating system-induced latency. Latency is the elapsed time between the network driver processing an incoming packet and the network driver sending the packet back. This time is usually measured in microseconds. For comparison, the transmission time for packet transmissions over long distances is usually measured in milliseconds (an order of magnitude larger). This tuning will not reduce the time a packet spends in transit.

Following are some performance tuning suggestions for microsecond-sensitive networks.

- Set the computer BIOS to **High Performance**, with C-states disabled. However, note that this is system and BIOS dependent, and some systems will provide higher performance if the operating system controls power management. You can check and adjust your power management settings from Control Panel or by using the **powercfg** command. For more information, see Powercfg Command-Line Options
- Set the operating system power management profile to **High Performance System**. Note that this will not work properly if the system BIOS has been set to disable operating system control of power management.
- Enable Static Offloads, for example, UDP Checksums, TCP Checksums, and Send Large Offload (LSO).
- Enable RSS if the traffic is multi-streamed, such as high-volume multicast receive.
- Disable the **Interrupt Moderation** setting for network card drivers that require the lowest possible latency. Remember, this can use more CPU time and it represents a tradeoff.
- Handle network adapter interrupts and DPCs on a core processor that shares CPU cache with the core that is being used by the program (user thread) that is handling the packet. CPU affinity tuning can be used to direct a process to certain logical processors in conjunction with

RSS configuration to accomplish this. Using the same core for the interrupt, DPC, and user mode thread exhibits worse performance as load increases because the ISR, DPC, and thread contend for the use of the core.

System Management Interrupts

Many hardware systems use System Management Interrupts (SMI) for a variety of maintenance functions, including reporting of error correction code (ECC) memory errors, legacy USB compatibility, fan control, and BIOS controlled power management. The SMI is the highest priority interrupt on the system and places the CPU in a management mode, which preempts all other activity while it runs an interrupt service routine, typically contained in BIOS.

Unfortunately, this can result in latency spikes of 100 microseconds or more. If you need to achieve the lowest latency, you should request a BIOS version from your hardware provider that reduces SMIs to the lowest degree possible. These are frequently referred to as “low latency BIOS” or “SMI free BIOS.” In some cases, it is not possible for a hardware platform to eliminate SMI activity altogether because it is used to control essential functions (for example, cooling fans).

Note
The operating system can exert no control over SMIs because the logical processor is running in a special maintenance mode, which prevents operating system intervention.

Performance Tuning TCP

You can performance tune TCP using the following items.

- TCP Receive Window Auto-Tuning
- Windows Filtering Platform
- TCP Parameters

Details are provided in the following sections.

TCP Receive Window Auto-Tuning

Prior to Windows Server 2008, the network stack used a fixed-size receive-side window that limited the overall potential throughput for connections. One of the most significant changes to the TCP stack is TCP receive window auto-tuning. You can calculate the total throughput of a single connection when you use this fixed size default as:

Total achievable throughput in bytes = TCP window * (1 / connection latency)

For example, the total achievable throughput is only 51 Mbps on a 1 GB connection with 10 ms latency – which is a reasonable value for a large corporate network infrastructure.

With auto-tuning, however, the receive-side window is adjustable, and it can grow to meet the demands of the sender. It is entirely possible for a connection to achieve the full line rate of a 1 GB connection. Network usage scenarios that might have been limited in the past by the total achievable throughput of TCP connections can now fully use the network.

Windows Filtering Platform

The Windows Filtering Platform (WFP) that was introduced in Windows Vista and Windows Server 2008 provides APIs to non-Microsoft independent software vendors (ISVs) to create packet processing filters. Examples include firewall and antivirus software.

Note
A poorly written WFP filter can significantly decrease a server's networking performance.

For more information, see Windows Filtering Platform in the Windows Dev Center.

TCP Parameters

The following registry keywords from Windows Server 2003 are no longer supported, and they are ignored in Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008.

- TcpWindowSize

  ```
  HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters
  ```

- NumTcbTablePartitions

  ```
  HKEY_LOCAL_MACHINE\system\CurrentControlSet\Services\Tcpip\Parameters
  ```

- MaxHashTableSize

  ```
  HKEY_LOCAL_MACHINE\system\CurrentControlSet\Services\Tcpip\Parameters
  ```

Network-Related Performance Counters

This topic lists the counters that are relevant to managing network performance, and contains the following sections.

- Resource Utilization
- Potential Network Problems
- Receive Side Coalescing (RSC) performance

Resource Utilization

The following performance counters are relevant to network resource utilization.

- IPv4, IPv6
  - Datagrams Received/sec
  - Datagrams Sent/sec
- TCPv4, TCPv6
  - Segments Received/sec

    - Segments Sent/sec
    - Segments Retransmitted/sec
- Network Interface(*), Network Adapter(*)
    - Bytes Received/sec
    - Bytes Sent/sec
    - Packets Received/sec
    - Packets Sent/sec
    - Output Queue Length

    This counter is the length of the output packet queue (in packets). If this is longer than 2, delays occur. You should find the bottleneck and eliminate it if you can. Because NDIS queues the requests, this length should always be 0.
- Processor Information
    - % Processor Time
    - Interrupts/sec
    - DPCs Queued/sec

    This counter is an average rate at which DPCs were added to the logical processor's DPC queue. Each logical processor has its own DPC queue. This counter measures the rate at which DPCs are added to the queue, not the number of DPCs in the queue. It displays the difference between the values that were observed in the last two samples, divided by the duration of the sample interval.

Potential Network Problems

The following performance counters are relevant to potential network problems.

- Network Interface(*), Network Adapter(*)
    - Packets Received Discarded
    - Packets Received Errors
    - Packets Outbound Discarded
    - Packets Outbound Errors
- WFPv4, WFPv6
    - Packets Discarded/sec
- UDPv4, UDPv6
    - Datagrams Received Errors
- TCPv4, TCPv6
    - Connection Failures
    - Connections Reset
- Network QoS Policy
    - Packets dropped
    - Packets dropped/sec

- Per Processor Network Interface Card Activity
  - Low Resource Receive Indications/sec
  - Low Resource Received Packets/sec
- Microsoft Winsock BSP
  - Dropped Datagrams
  - Dropped Datagrams/sec
  - Rejected Connections
  - Rejected Connections/sec

Receive Side Coalescing (RSC) performance

The following performance counters are relevant to RSC performance.

- Network Adapter(*)
  - TCP Active RSC Connections
  - TCP RSC Average Packet Size
  - TCP RSC Coalesced Packets/sec
  - TCP RSC Exceptions/sec

Performance Tools for Network Workloads

This topic contains the following sections about performance tools for networking.

- Tuning for NTttcp
- TCP/IP Window Size
- Microsoft Server Performance Advisor 3.0

Tuning for NTttcp

NTttcp is a Winsock-based port of ttcp to Windows. It helps measure network driver performance and throughput on different network topologies and hardware setups. It provides the customer with a multithreaded, asynchronous performance workload for measuring an achievable data transfer rate on an existing network setup.

When setting up NTttcp, consider the following:

- A single thread should be sufficient for optimal throughput.
- Multiple threads are required only for single-to-many clients.
- Posting enough user receive buffers (by increasing the value passed to the -a option) reduces TCP copying.
- You should not excessively post user receive buffers because the first buffers that are posted would return before you need to use other buffers.

- It is best to bind each set of threads to a logical processor (the second delimited parameter in the **-m** option).
- Each thread creates a logical processor that connects to (listens) a different port.

The following table provides example syntax for NTttcp Sender and Receiver.

Syntax	Details
Example Syntax for a Sender: NTttcps –m 1,0,10.1.2.3 –a 2	Single thread; Bound to CPU 0; Connects to a computer that uses IP 10.1.2.3; Posts two send-overlapped buffers; Default buffer size: 64 K; Default number of buffers to send: 20 K.
Example Syntax for a Receiver: NTttcpr –m 1,0,10.1.2.3 –a 6 –fr	Single thread; Bound to CPU 0; Binds on local computer to IP 10.1.2.3; Posts six receive-overlapped buffers; Default buffer size: 64 KB; Default number of buffers to receive: 20 K; Posts full-length (64 K) receive buffers.

Note
Ensure that you enable all offloading features on the network adapter.

For more information, see How to Use NTttcp to Test Network Performance in the Windows Dev Center.

TCP/IP Window Size

For 1 GB adapters, the settings shown in the previous table should provide good throughput because NTttcp sets the default TCP window size to 64 K through a specific logical processor option (SO_RCVBUF) for the connection. This provides good performance on a low-latency network.

In contrast, for high-latency networks or for 10 GB adapters, the default TCP window size value for NTttcp yields less than optimal performance. In both cases, you must adjust the TCP window size to allow for the larger bandwidth delay product.

You can statically set the TCP window size to a large value by using the **-rb** option. This option disables TCP Window Auto-Tuning, and we recommend using it only if the user fully understands the resultant change in TCP/IP behavior. By default, the TCP window size is set at a sufficient value and adjusts only under heavy load or over high-latency links.

Microsoft Server Performance Advisor 3.0

Microsoft Server Performance Advisor (SPA) 3.0 helps IT administrators collect metrics to identify, compare, and diagnose potential performance issues in a Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 deployment. SPA generates comprehensive

diagnostic reports and charts, and it provides recommendations to help you quickly analyze issues and develop corrective actions.

For more information, see Microsoft Server Performance Advisor 3.0 in the Windows Dev Center.

NIC Teaming Overview

NIC Teaming, also known as load balancing and failover (LBFO), allows multiple network adapters on a computer to be placed into a team for the following purposes:

- Bandwidth aggregation
- Traffic failover to prevent connectivity loss in the event of a network component failure

This feature has been a requirement for independent hardware vendors (IHVs) to enter the server network adapter market, but until now NIC Teaming has not been included in Windows Server operating systems.

Requirements

NIC Teaming requires the presence of a single Ethernet network adapter, which can be used for separating traffic that is using VLANs. All modes that provide fault protection through failover require at least two Ethernet network adapters. Windows Server® 2012 supports up to 32 network adapters in a team.

Technical overview

This guide includes the following sections that describe the NIC Teaming technology in Windows Server 2012:

- NIC Teaming architecture
- Bandwidth aggregation (also known as load balancing) mechanisms
- Traffic distribution algorithms
- NIC Teaming in virtual machines
- Incompatibilities
- Requirements

NIC Teaming architecture

Today, all NIC Teaming solutions on the market have a similar architecture, as shown in Figure 1.

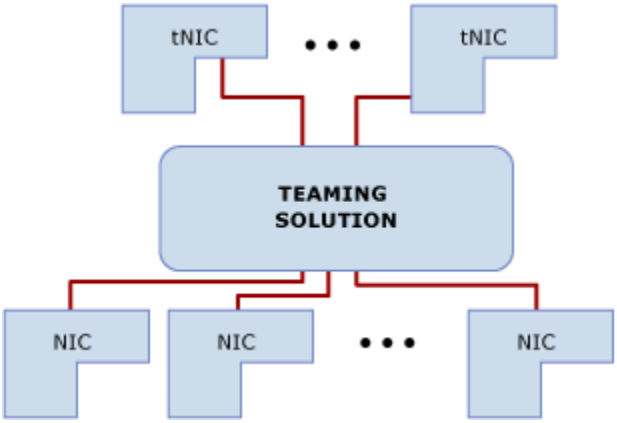

Figure 1 Teaming solution architecture

Two or more physical network adapters are connected to the NIC Teaming solution multiplexing unit, which then presents one or more virtual adapters (also known as *team network adapters*) to the operating system. There are several different algorithms that distribute inbound and outbound traffic between the physical network adapters. In current non-Microsoft NIC Teaming solutions, the team network adapters divide traffic by virtual LAN (VLAN) so that applications can connect to different VLANs at the same time. Technically, this separation of traffic is not part of NIC Teaming. However, because other commercial implementations of NIC Teaming have this capability, the Windows Server 2012 implementation also includes it.

NIC Teaming configurations

There are two basic sets of algorithms that are used for NIC Teaming:

- Algorithms that require the switch to participate in the teaming, also known as *switch-dependent modes*. These algorithms usually require all the network adapters of the team to be connected to the same switch.
- Algorithms that do not require the switch to participate in the teaming, also referred to as *switch-independent modes*. Because the switch does not know that the network adapter is part of a team, the team network adapters can be connected to different switches. Switch-independent modes do not require that the team members connect to different switches, they merely make it possible.

There are two common choices for switch-dependent modes of NIC Teaming:

- **Generic or static teaming (IEEE 802.3ad draft v1).** This mode requires configuration on the switch and the computer to identify which links form the team. Because this is a statically configured solution, no additional protocol assists the switch and the computer to identify incorrectly plugged cables or other errors that could cause the team to fail. This mode is typically supported by server-class switches.
- **Dynamic teaming (IEEE 802.1ax, LACP).** IEEE 802.1ax uses the Link Aggregation Control Protocol (LACP) to dynamically identify links between the computer and a specific switch. This enables the automatic creation of a team and, in theory, the expansion and reduction of a team simply by the transmission or receipt of LACP from the peer network adapter. Typical

server-class switches support IEEE 802.1ax, but most switches require manual administration to enable LACP on the port.

> **Note**
> IEEE 802.1ax is also commonly known as IEEE 802.3ad because it was developed by the IEEE 802.3ad committee before being published as IEEE 802.1ax.

Both modes should result in inbound and outbound traffic approaching the practical limits of the aggregated bandwidth because the pool of links in the team functions as a single pipeline.

Traffic distribution algorithms

Outbound traffic can be distributed among the available links in many ways. Try to keep all packets that are associated with a single flow (TCP stream) on a single network adapter. This is needed to minimize out-of-order packet arrival scenarios.

NIC Teaming in Windows Server 2012 supports the following traffic distribution methods:

- **Hyper-V switch port.** In cases where virtual machines have independent media access control (MAC) addresses, the MAC address of the virtual machine can provide the basis for dividing traffic. There is an advantage in using this scheme in virtualization. Because the adjacent switch can determine that specific source MAC addresses are on only one connected network adapter, the switch will balance the egress load (the traffic from the switch to the computer) on multiple links, based on the destination MAC address for the virtual machine. This is particularly helpful when used with virtual machine queue. However, this mode might not be specific enough to get a well-balanced distribution, and it limits a single virtual machine to the bandwidth that is available on a single network adapter.

  > **Note**
  > Windows Server 2012 uses the Hyper-V switch port as the identifier rather than the source MAC address, because in some instances, a virtual machine might be using more than one MAC address on a switch port.

- **Hashing.** This algorithm creates a hash based on components of the packet, and then it assigns packets that have that hash value to one of the available network adapters. This keeps all packets from the same TCP stream on the same network adapter. Hashing alone usually creates balance across the available network adapters. Some NIC Teaming solutions that are available on the market monitor the distribution of the traffic and reassign specific hash values to different network adapters in an attempt to better balance the traffic. The dynamic redistribution is known as *smart load balancing* or *adaptive load balancing*.

The components that can be used as inputs to the hashing function include:

- Source and destination MAC addresses
- Source and destination IP addresses, with or without considering the MAC addresses (2-tuple hash)
- Source and destination TCP ports, usually used along with the IP addresses (4-tuple hash)

The quadruple hash creates a finer distribution of traffic streams, which results in smaller streams that can be independently moved between network adapters. However, it cannot be used for traffic that is not TCP or UDP traffic or where the TCP and UDP ports are hidden from the stack,

such as traffic that is protected by Internet Protocol security (IPsec). In these cases, the hash falls back to a 2-tuple hash. If the traffic is not IP traffic, the hash generator will use the source and destination MAC addresses.

NIC Teaming in virtual machines

NIC Teaming in Windows Server 2012 also works within a virtual machine. This allows a virtual machine to have virtual network adapters that are connected to more than one Hyper-V switch and still have connectivity even if the network adapter under that switch gets disconnected. This is particularly important when working with features such as Single Root I/O Virtualization (SR-IOV) because SR-IOV traffic does not go through the Hyper-V switch. Thus, it cannot be protected by a team that is under a Hyper-V switch. With the virtual machine teaming option, an administrator can set up two Hyper-V switches, each connected to its own SR-IOV-capable network adapter. At that point:

- Each virtual machine can then install a virtual function from one or both SR-IOV network adapters. Then, in the event of a network adapter disconnect, the virtual machine can fail over from the primary virtual function to the backup virtual function.
- Alternately, the virtual machine might have a virtual function from one network adapter and a non-virtual function network adapter to the other switch. If the network adapter that is associated with the virtual function gets disconnected, the traffic can fail over to the other switch without loss of connectivity.

Note

Because failover between network adapters in a virtual machine might result in traffic being sent with the MAC address of the other network adapter, each Hyper-V switch port that is associated with a virtual machine that is using NIC Teaming must be set to allow MAC spoofing or must have the “AllowTeaming=On” parameter set using the **Set-VmNetworkAdapter** PowerShell cmdlet.

Incompatibilities

NIC Teaming is compatible with all networking capabilities in Windows Server 2012 with three exceptions: SR-IOV, remote direct memory access (RDMA), and TCP Chimney. For SR-IOV and remote direct memory access (RDMA), data is delivered directly to the network adapter without passing through the networking stack. Therefore, it is not possible for the network adapter team to look at or redirect the data to another path in the team. TCP Chimney is not supported with NIC Teaming in Windows Server 2012.

Requirements

NIC Teaming requires the presence of a single Ethernet network adapter, which can be used for separation of traffic using VLANs. All modes that provide fault protection through failover require at least two Ethernet network adapters. The Windows Server 2012 implementation supports up to 32 NICs in a team.

Configuring NIC Teaming

You can configure NIC Teaming in Windows Server 2012 through the NIC Teaming management interface or through the rich PowerShell interface. The NIC Teaming management UI is a simple interface that uses the PowerShell cmdlets to manage teams. See below for the complete guide to NIC Teaming PowerShell cmdlets. There are multiple ways to invoke the NIC Teaming management interface.

From Server Manager:

1. Open Server Manager.
2. In the console tree, click **Local Server**.
3. In the details pane, in the **Properties** section, click **NIC Teaming Administration** under **Remote Desktop**.

From the desktop:

1. Click **Start**, type **lbfoadmin**, and then press **Enter**. Alternately, click **Start**, click **Run**, type **lbfoadmin**, and then press **Enter**.

To manage multiple servers at a time you can add servers through the “Add Server” task item, or you can use the following command:

lbfoadmin /servers servername1 servername2 …

This will run the NIC Teaming management interface with the list of specified server names.

To configure NIC Teaming on a server

1. Click the server name in the list of servers (even if there is only a single server).
2. From the **Tasks** drop-down menu in the **Teams** section, click **New Team**.
3. In the **Add Team** dialog box, type a team name and select the network adapters in the team.
4. If you are using VLANs, clear the **Default** check box and specify the VLAN ID. By default, the team will deliver up all traffic received regardless of VLAN ID (though the VLAN ID is passed with the packet so that other components in the stack can sort them appropriately).
5. If you want to use a mode other than the default, click **Advanced** and specify the teaming mode (Switch Independent, Static Teaming, or LACP) and load distribution mode (Address Hash) as needed.
6. Click **OK** to create the team.

   Note
   The Address Hash option for the load distribution mode is the 4-tuple hash previously described. PowerShell cmdlets for NIC Teaming allow you to configure other address hashing modes.

   To modify the settings of a network adapter team, click the team name in the list of

teams, and then click **Modify team** in the list of tasks.

NIC Teaming PowerShell commands

Windows Server 2012 NIC Teaming has PowerShell cmdlets that operate on the following object types:

Objects to be managed	Supported operations on the object
1. NetLbfoTeam	Get, New, Remove, Rename, Set
2. NetLbfoTeamMember	Add, Get, Remove, Set
3. NetLbfoTeamNic	Get, New, Remove, Set

For further information and complete descriptions of the NIC Teaming PowerShell cmdlets, see the PowerShell cmdlet reference for NIC Teaming.

Quality of Service (QoS) Overview

QoS is a set of technologies that provide you with the ability to cost-effectively manage network traffic and to enhance user experiences in enterprise environments, small offices, and even home networking environments.

This document provides introductory information about Quality of Service (QoS) in Windows Server® 2012.

Notes
In addition to this topic, the following Windows Server 2012 QoS documentation is available.

- QoS Common Configurations
- QoS Minimum Bandwidth Best Practices
- Policy-based Quality of Service (QoS)
- Configuring Policy-based Quality of Service (QoS)
- Quality Windows Audio Video Experience (qWave)
- QoS Frequently Asked Questions

QoS features

QoS technologies allow you to meet the service requirements of a workload or an application by measuring network bandwidth, detecting changing network conditions (such as congestion or availability of bandwidth), and prioritizing - or throttling - network traffic. For example, you can use

QoS to prioritize traffic for latency-sensitive applications (such as voice or video streaming), and to control the impact of latency-insensitive traffic (such as bulk data transfers).

QoS provides the following features.

- Bandwidth management
- Classification and tagging
- Priority based flow control
- Policy-based QoS and Hyper-V QoS

Information on these features is provided in the following sections.

Note

QoS is included in all versions of Windows Server 2012 and Windows® 8, and supports all QoS APIs and features that are provided by QoS in Windows Server® 2008 R2 and Windows® 7. QoS is not available in Windows® RT.

Bandwidth management

In previous versions of QoS, as with QoS in Windows Server 2012, Maximum Bandwidth is a feature that sets an absolute ceiling - a throttling rate - for a workload. Windows Server 2012 introduces Minimum Bandwidth.

Minimum Bandwidth provides a specified level of service to a workload when network congestion occurs while still permitting higher bandwidth utilization by this workload in circumstances where there is no network congestion. Network congestion occurs on a computer system when multiple workloads compete for access to an external network through one network adapter.

Because the capacity of any network adapter is limited, when the sum of the workloads exceeds that capacity, the network throughput of each workload becomes nondeterministic. With Minimum Bandwidth, you can define how bandwidth is shared by competing applications.

For example, you might have Storage, Live Migration, and Cluster sharing a 10 GB NIC on a computer that is running the Hyper-V server role. If you assign a weight of 50, 30 and 20 to Storage, Live Migration, and Cluster, respectively, they receive corresponding shares of bandwidth when the NIC is congested.

Both Minimum Bandwidth and Maximum Bandwidth enable you to enforce predictable network throughput for each workload. The main difference between them, from an administrator's point of view, is how they utilize network resources.

Network resource utilization

When Maximum Bandwidth is set for a workload, the workload can never exceed the ceiling even if no other workloads in the system are using the network bandwidth. When Minimum Bandwidth is set for a workload, the workload can use as much bandwidth as it can - until network congestion occurs. Using the previous example, when Live Migration is not occurring, both the Storage and Cluster workloads can exceed their predefined shares of network bandwidth.

You can choose to enable Maximum Bandwidth, Minimum Bandwidth, or both based on the requirements of a specific scenario. For example, to prevent a hosted VM from unexpectedly

using all available Internet bandwidth, a public hoster can impose a QoS Maximum Bandwidth setting on the VM. Conversely, to maximize the bandwidth utilization on any computer running the Hyper-V server role within the datacenter, you can impose a minimum bandwidth setting on the server.

Server Message Block Direct (SMB Direct) and Data Center Bridging (DCB).

Minimum Bandwidth and Maximum Bandwidth are enforced by QoS Packet Scheduler. They work well with traffic that flows through the Windows Server networking stack.

But there is network traffic that bypasses the networking stack. The most notable example is Server Message Block Direct (SMB Direct), a Windows Server 2012 feature that builds on Remote Direct Memory Access (RDMA). SMB Direct offloads the SMB traffic directly to an RDMA-capable NIC to reduce latency and the number of CPU cycles that are spent on networking. Because SMB Direct still competes for bandwidth with other network traffic at the NIC hardware, one can impact or even starve the other if they are left unchecked.

To resolve this issue, you can use Data Center Bridging (DCB). DCB is implemented by some NIC vendors in network adapters. DCB works in a similar way as Minimum Bandwidth: each class of traffic, regardless of whether it's offloaded or not, has bandwidth allocation; in the event of network congestion, each class gets its share - otherwise, each class gets as much bandwidth as is available.

Classification and tagging

Before the bandwidth for a workload is managed, the workload must be classified or filtered out so that QoS Packet Scheduler or a DCB capable NIC can act upon it. Windows has sophisticated traffic classification capability. Classification can be based on 5-tuples, user type, or URI. Windows Server 2012 simplifies the management task so that you can invoke built-in filters in Windows PowerShell to classify some of the most common workloads. The table below lists these built-in filters.

Workload	Built-in Filter (Parameter) in Windows PowerShell	Filter Implementation
iSCSI	-iSCSI	Match TCP or UDP port 3260
NFS	-NFS	Match TCP or UDP port 2049
SMB	-SMB	Match TCP or UDP port 445
Live Migration	-LiveMigration	Match TCP port 660
SMB Direct	-NetDirect<port>	Match the specified port
Wild Card	-Default	Any traffic that is not otherwise classified

Notes

- SMB Direct offloads SMB traffic to a NIC. When “-NetDirect <port>” is specified in a QoS policy, Windows sends the instruction to any NIC that supports SMB Direct; the compatible NICs then filter out the network traffic flow based on the specified port.
- Wild Card allows you to capture any traffic that is not filtered out or classified. For example, you can use this wild card filter to create a class of traffic for Management on a computer that is running the Hyper-V server role.

Bandwidth Management is not the only action that you can specify in a QoS policy. For example, you can prioritize and tag traffic so that QoS is enforced from end to end across a data center.

Previous versions of Windows support priority tagging in the IP header of a network packet. The priority bits are known as Differentiation Service Code Point (DSCP). Windows Server 2012 adds the support of 802.1p tagging. 802.1p is a 3-bit priority value in the layer 2 Ethernet frame. Therefore it applies to non-IP packets such as RDMA over Ethernet.

Priority based flow control

Workloads, such as RDMA, require lossless transport. If RDMA is built on top of Ethernet directly, which is known as RDMA over Converged Ethernet (ROCE), the Ethernet transport must be lossless. Traditionally, link level flow control, which relies on the 802.3 Pause frame, is a solution. But link level flow control causes problems such as head of line blocking.

This issue is resolved by Priority based Flow Control (PFC), which is one standard defined by the Institute of Electrical and Electronics Engineers (IEEE) DCB workgroup. Windows Server 2012 allows you to enable PFC as long as it is supported by a NIC. When PFC is enabled for ROCE on both ends of an Ethernet link, only the virtual link designated for ROCE, which is denoted by a priority value, becomes lossless, and other workloads on the same physical link do not suffer from head of line blocking.

Policy-based QoS and Hyper-V QoS

For network administrators, QoS in Windows Server® 2012 is designed to help manage network traffic on the physical network and on the virtual network. Policy-based QoS is designed to manage traffic on the physical network. And a new functionality in QoS, called Hyper-V QoS, allows you to manage traffic on the virtual network.

Policy-based QoS

You can use Policy-based QoS to manage network traffic on the physical network. Policy-based QoS enables you to specify network bandwidth control based on application type, users, and computers. You can use Policy-based QoS to manage traffic to help control bandwidth costs, negotiate service levels with bandwidth providers or business departments, and to offer better end-user experiences. Because Policy-based QoS is built into Group Policy, it is part of your existing management infrastructure, and is therefore a cost-effective solution to implement.

Policy-based QoS provides the ability to:

- Enforce minimum bandwidth for a traffic flow, which is identified by a 5-tuple traffic filter.

Note

5-tuple is a term used in computer networking that describes a set of five different values that constitute a Transmission Control Protocol/Internet Protocol (TCP/IP) connection.

- Configure and query or view QoS policies, which enforce minimum and maximum bandwidth and 802.1p or Differentiated Services Code Point (DSCP) marking on filtered packets, through WMI and PowerShell.

Notes

- DSCP is a field in an IPv4 or IPv6 header in computer networking.
- 802.1p is an Institute of Electrical and Electronics Engineers (IEEE) specification that enables Layer-2 switches to perform traffic prioritization and dynamic multicast filtering.

- Use Group Policy editors to configure an Extended Quality of Service (eQoS) policy to tag packets with an 802.1p value.
- Configure local QoS policies on computers that are not joined to a domain.

For more information, see Policy-based Quality of Service (QoS).

Hyper-V QoS

You can use Hyper-V QoS to manage network traffic on the virtual network. In Windows Server® 2012, QoS includes new bandwidth management features that enable cloud hosting providers and enterprises to provide services that deliver predictable network performance to virtual machines on a server that is running the Hyper-V server role. Hyper-V QoS supports the management of upper-allowed and lower-allowed bandwidth limits, commonly referred to as maximum bandwidth and minimum bandwidth.

In hosted environments, Hyper-V QoS enables you to guarantee specific performance levels based on the service level agreements (SLAs) to which you have agreed with your customers. Hyper-V QoS helps ensure that your customers are not impacted or compromised by other customers on their shared infrastructure, which can include computing, storage, and network resources.

In addition, enterprise networks might require similar functionality. By using Hyper-V QoS in your enterprise, you can run multiple virtual machine-based application servers on a host server that is running Hyper-V, and have confidence that each application server delivers predictable performance.

Hyper-V QoS provides the ability to:

- Enforce minimum bandwidth and maximum bandwidth for a traffic flow, which is identified by a Hyper-V Virtual Switch port number.
- Configure minimum bandwidth and maximum bandwidth per Hyper-V virtual switch port by using either PowerShell cmdlets or Windows Management Instrumentation (WMI).
- Configure multiple virtual network adapters in Hyper-V and specify QoS on each virtual network adapter individually.

Windows Server 2012 Hyper-V QoS can also use hardware that is compatible with data center bridging (DCB) to converge multiple types of network traffic on a single network adapter with a guaranteed level of service provided to each type of traffic. With Windows PowerShell, you can configure these new features manually or enable automation in a script to manage a group of servers, regardless of whether they are joined to a domain.

QoS Common Configurations

When you configure Quality of Service (QoS), you must consider the following variables.

- The number of network interface cards (NICs) that are installed on the server computer.
- Whether or not each NIC is dedicated to a specific workload or is shared by multiple workloads.
- Whether NICs are teamed with the NIC teaming feature.
- The number of NIC teams that you configure on the server.
- Whether you plan on using offload technologies like Single Root Input-Output Virtualization (SR-IOV) or Remote Direct Memory Access (RDMA).

The following sections examine a few common server configurations and illustrate where QoS can be enforced.

- 2 NICs without NIC Teaming
- 2 NICs with NIC Teaming
- 4 NICs in two NIC teams
- 4 NICs with a standard NIC team and two RDMA NICs
- Alternate configuration of 4 NICs with a standard NIC team and two RDMA NICs

2 NICs without NIC Teaming

This configuration prioritizes the separation of the host traffic from the guest traffic.

Note

In this document, the term *management operating system* refers to the operating system of the computer upon which you are running Hyper-V.

The following illustration depicts this configuration.

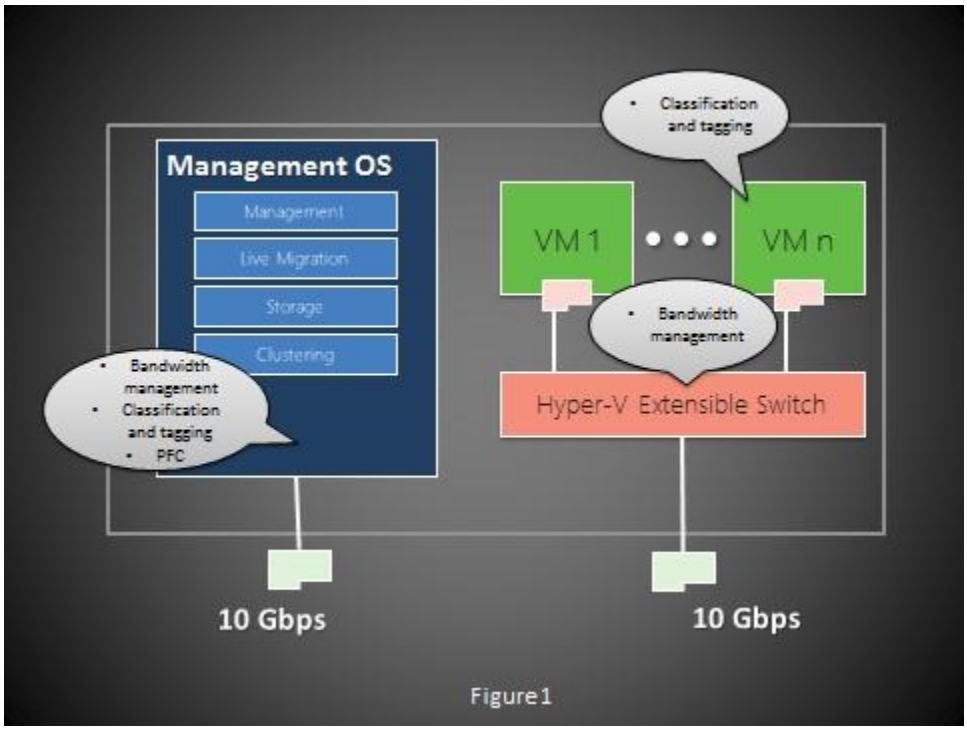

Figure 1

With this configuration, you can apply all QoS features, including Bandwidth Management, Classification and Tagging, and Priority-based Flow Control, as long as the feature is supported by the NIC.

The following code, which is run in the management operating system Windows PowerShell environment, provides an example of creating this configuration by using Windows PowerShell. In this code sample, the lines that begin with a number or pound sign (#) are remarks that explain the code in the subsequent line.

```
# Use the inbox filter for Live Migration

New-NetQosPolicy "Live Migration" -LiveMigration -MinBandwidthWeight 30 -Priority 5

# Use the inbox filter for SMB

New-NetQosPolicy "SMB" -SMB -MinBandwidthWeight 50 -Priority 3

# Create a policy for the cluster heartbeat traffic sent on port 3343

New-NetQosPolicy "Cluster"-IPDstPort 3343 -MinBandwidthWeight 10 -Priority 6

# Use the inbox filter to capture all the rest of the traffic.

New-NetQosPolicy "Management" -Default -MinBandwidthWeight 10


# Note that the management traffic is deliberately configured not to be

# tagged.
```

You can also enable Bandwidth Management on the Hyper-V Virtual Switch for each individual virtual network adapter. The following code presents an example of how to create this configuration.

```
# Create a Hyper-V Virtual Switch that supports Minimum Bandwidth by weight
New-VMSwitch "10GbE switch" –NetAdapterName "XYZ's 10GbE" –MinimumBandwidthMode Weight –AllowManagementOS $false
# Assign an even weight to all the VMs here as well as a hard cap at 1Gbps
Set-VMNetworkAdapter –VMName * –MinimumBandwidthWeight 1 –MaximumBandwidth 1000000000
```

Optionally, if a VM is trusted – such as in an enterprise environment where you can trust the Administrator of the VM - you can enable QoS classification and tagging from within the VM. The following code, which is run in the Windows PowerShell environment on a VM, presents an example of how to create this configuration.

```
# Create a QoS policy to tag (in the IP header) a critical application
# running on this VM
New-NetQosPolicy "Line of Business app" –ApplicationName "%ProgramFiles%\crm.exe" –DSCP 46
```

2 NICs with NIC Teaming

This configuration prioritizes high availability for all of the workloads on a computer that is running the Hyper-V server role.

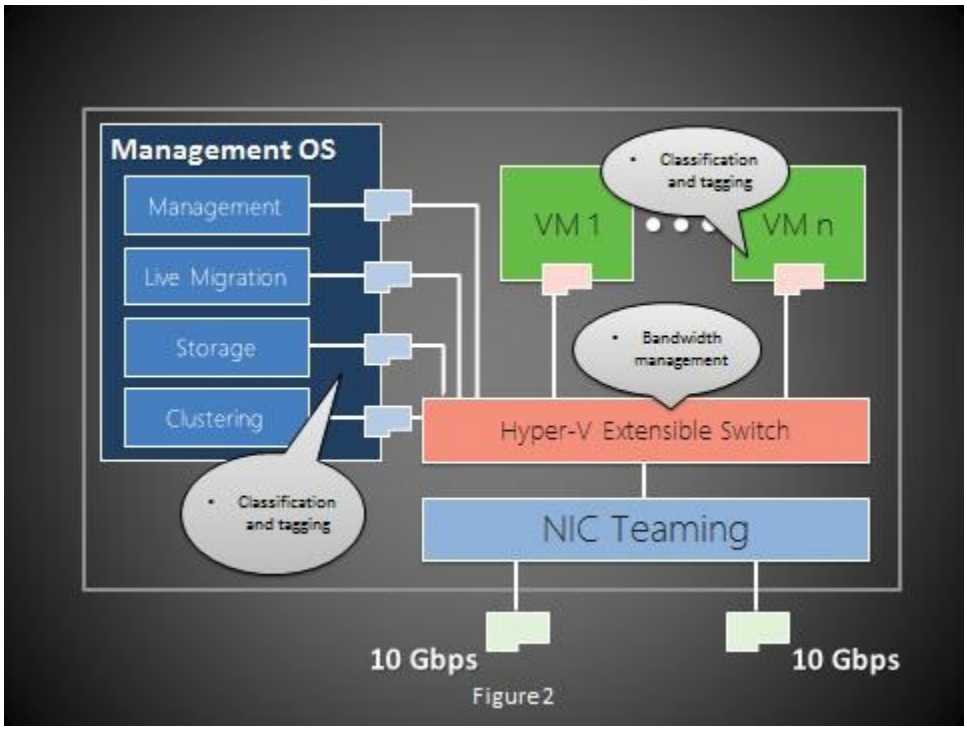

Figure2

You can enable Bandwidth Management on the Hyper-V Virtual Switch for both the VMs and the workloads in the management operating system. The following example commands are run in the Windows PowerShell environment in the management operating system.

```
# Create a NIC team that consists of two 10GbE NICs

New-NetLbfoTeam “2x10GbE Team” –TeamMembers “10GbE NIC1”,”10GbE NIC2” –TeamNicName “2x10GbE”

# Create a Hyper-V Virtual Switch that binds to the NIC team

New-VMSwitch “20GbE switch” –NetAdapterName “2x10GbE” –MinimumBandwidthMode Weight –AllowManagementOS $false

# Create a virtual NIC in the management operating system for SMB (storage)

Add-VMNetworkAdapter –ManagementOS –Name “SMB” –SwitchName “20GbE switch”

# Create a virtual NIC in the management operating system for Live Migration

Add-VMNetworkAdapter –ManagementOS –Name “LM” –SwitchName “20GbE switch”

# Create a virtual NIC in the management operating system for Cluster

Add-VMNetworkAdapter –ManagementOS –Name “Cluster” –SwitchName “20GbE switch”

# Create a virtual NIC in the management operating system for Management

Add-VMNetworkAdapter –ManagementOS –Name “Management” –SwitchName “20GbE switch”


# Assign a weight to SMB
```

```
Set-VMNetworkAdapter –ManagementOS –Name “SMB” –MinimumBandwidthWeight 40
# Assign a weight to Live Migration
Set-VMNetworkAdapter –ManagementOS –Name “LM” –MinimumBandwidthWeight 20
# Assign a weight to Cluster
Set-VMNetworkAdapter –ManagementOS –Name “Cluster” –MinimumBandwidthWeight 5
# Assign a weight to Management
Set-VMNetworkAdapter –ManagementOS –Name “Management” –MinimumBandwidthWeight 5

# Assign a weight to each VM
Set-VMNetworkAdapter –VMName * –MinimumBandwidthWeight 1

# Usually the workloads in the management operating system are put in different IP
subnets
# and VLANs, to assign a VLAN to a virtual network adapter in the Management
# OS
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName SMB –Access –VlanId 100
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName LM –Access –VlanId 200
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName Cluster –Access –VlanId 300
Set-VMNetworkAdapterVlan –ManagementOS –VMNetworkAdapterName Management –Access –VlanId
400
```

Optionally, you can classify and tag traffic from within the management operating system and, in a trusted environment, from within the VMs.

```
# Use the inbox filter for Live Migration
New-NetQosPolicy “Live Migration” –LiveMigration –Priority 5
# Use the inbox filter for SMB
New-NetQosPolicy “SMB” –SMB –Priority 3
# Create a policy for the cluster heartbeat traffic sent on port 3343
New-NetQosPolicy “Cluster”-IPDstPort 3343 –Priority 6

# Alternatively, if these workloads are in different IP subnets they can
# be classified based on their IP subnet address
# Assume Live Migration is on 10.1.0.0/16
New-NetQosPolicy “Live Migration” –IPDstPrefix 10.1.0.0/16 –Priority 5
```

```
# Assume SMB is on 10.2.0.0/16
New-NetQosPolicy “SMB” –IPDstPrefix 10.2.0.0/16 –Priority 3
# Assume Cluster is on 10.3.0.0/16
New-NetQosPolicy “Cluster”-IPDssPrefix 10.3.0.0/16 –Priority 6

# Note that no explicit policy is created for the Management traffic, if it
# doesn’t require to be tagged.


# Enable priority tagged traffic to go through the Hyper-V Virtual Switch
Set-VMNetworkAdapter –ManagementOS –IeeePriorityTag On

# Note that the name of the virtual network adapter in the management operating system
# is deliberately omitted in the above command so that the configuration applies
# to all virtual network adapters in the management operating system

# Also note that if workloads are tagged in the IP header like the example
# shown in “2 NICs without NIC Teaming,” no additional configuration is required
# to let the Hyper-V Virtual Switch pass such DSCP tagged traffic.
```

4 NICs in two NIC teams

This configuration provides separation of the host traffic and the guest traffic, as well as high availability for all the workloads. This configuration doubles the required number of NICs to four, but they do not need to all be 10 GbE NICs. For example, if a server has a dual port 10GbE NIC and two 1GbE LOM ports, you can team the two 10GbE NICs for the management operating system and the two 1GbE LOM ports for the VMs.

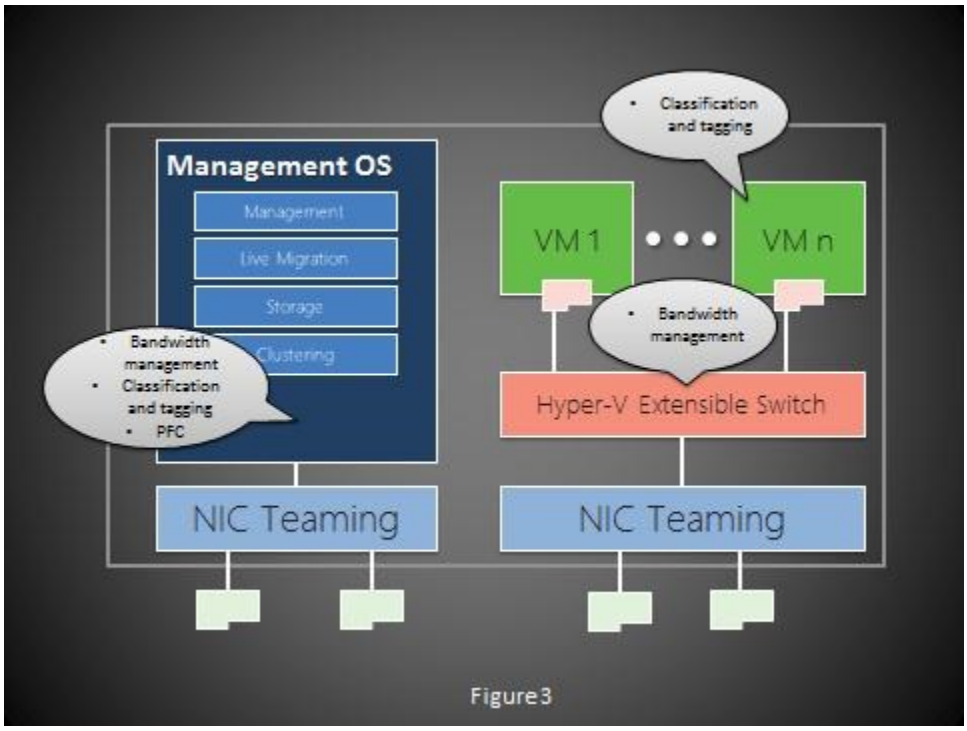

Figure3

For this configuration, you can use the same Windows PowerShell commands as those provided in the section 2 NICs without NIC Teaming.

4 NICs with a standard NIC team and two RDMA NICs

This configuration emphasizes the use of RDMA. To converge other workloads such as Live Migration, Cluster and Management on the same RDMA NICs, the NICs must also support Data Center Bridging (DCB). To provide high availability to Storage, you can enable Microsoft Multipath I/O (MPIO).

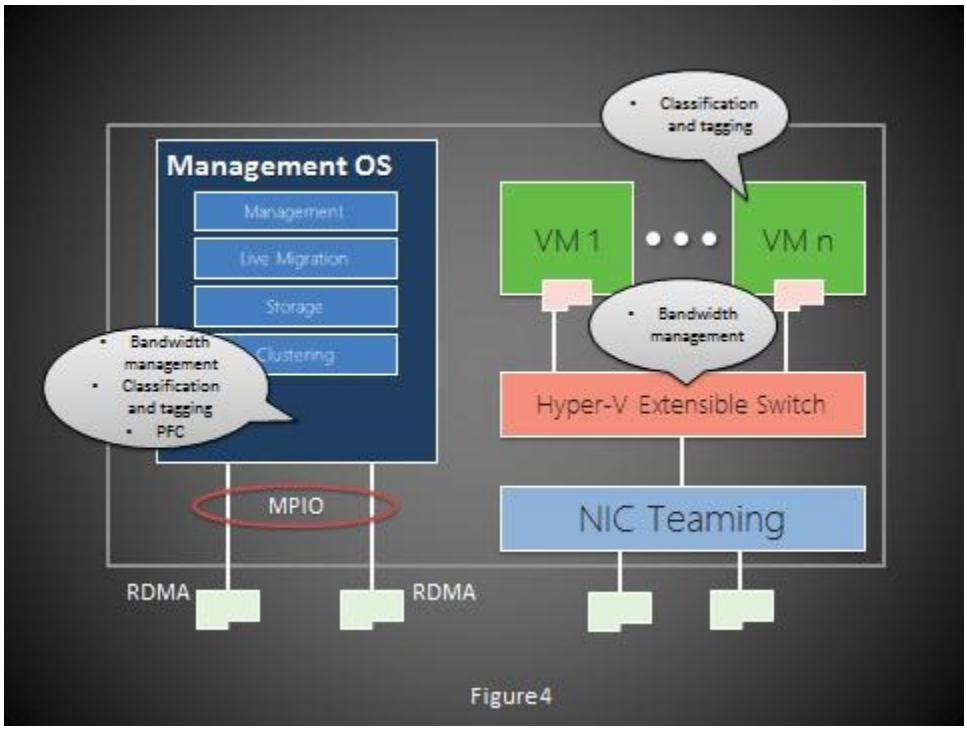

Figure4

With this configuration, you can apply all QoS features, including Bandwidth Management, Classification and Tagging, and PFC in the management operating system.

```
# Use the inbox filter for Live Migration

New-NetQosPolicy "Live Migration" -LiveMigration -Priority 5

# Create a policy for SMB Direct traffic sent on port 445

New-NetQosPolicy "SMB Direct" -NetDirectPort 445 -Priority 3

# Create a policy for the cluster heartbeat traffic sent on port 3343

New-NetQosPolicy "Cluster"-IPDstPort 3343 -Priority 6


# Install DCB on the server

Install-WindowsFeature Data-Center-Bridging

# Create a traffic class on Data Center Bridging (DCB) NICs for Live Migration

New-NetQosTrafficClass "Live Migration" -Priority 5 -Algorithm ETS -Bandwidth 30

# Create a traffic class on DCB NICs for SMB Direct (ROCE)

New-NetQosTrafficClass "SMB Direct" -Priority 3 -Algorithm ETS -Bandwidth 50

# Create a traffic class on DCB NICs for Cluster

New-NetQosTrafficClass "Cluster" -Priority 6 -Algorithm ETS -Bandwidth 10
```

```
# Note that there’s a default traffic class, which owns 100% of the total
# bandwidth. As bandwidth is allocated to user-created traffic classes, the
# default traffic class owns whatever is left. In this example, it has 10%
# left for all the rest of traffic in the management operating system.

# Enable PFC for ROCE
Enable-NetQosFlowControl –Priority 3

# Note that PFC should be enabled for priority 3 from end (e.g. a computer running
# Hyper-V) to end (e.g. a File Server) at both ends (i.e. switch ports and
# NICs) of all the links.

# Enable DCB settings on the NICs
Set-NetQosDcbxSetting –Willing $false

Enable-NetAdapterQos “DCB-capable 10Gbe NIC1”
Enable-NetAdapterQos “DCB-capable 10Gbe NIC2”
```

You can enable Bandwidth Management on the Hyper-V Virtual Switch for the VMs and enable Classification and Tagging on the VMs, if they are trusted.

Alternate configuration of 4 NICs with a standard NIC team and two RDMA NICs

MPIO provides redundancy to Storage only. If you want to provide redundancy for Management, Live Migration, and Cluster you can modify the previous configuration slightly as follows.

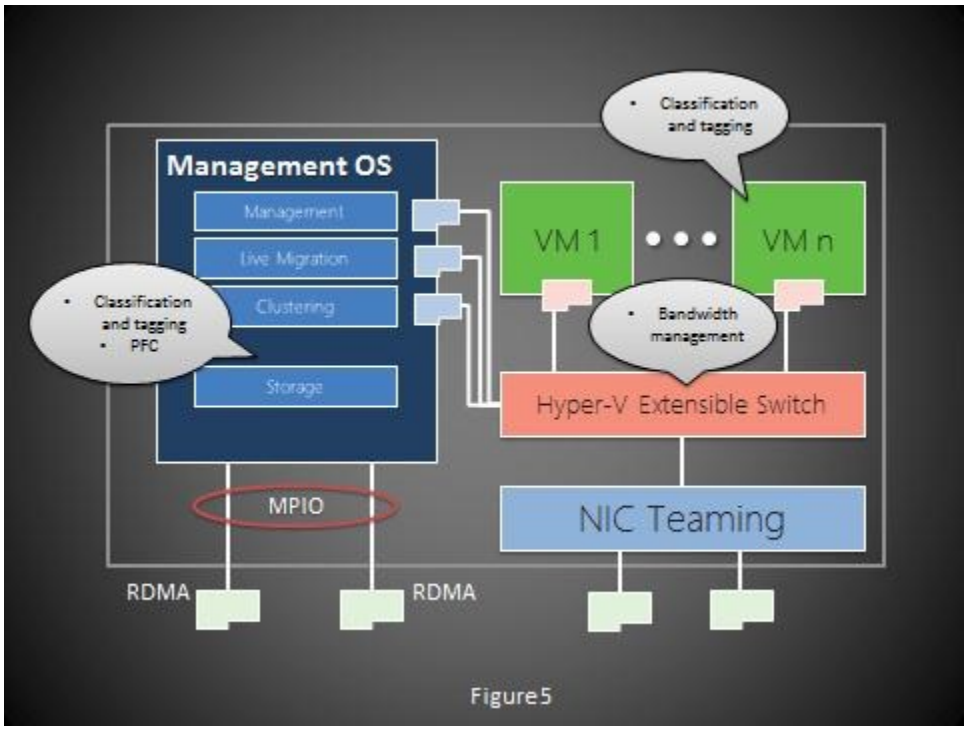

Figure5

Because the RDMA NICs are dedicated to Storage, no Bandwidth Management is required in the management operating system. Minimum Bandwidth and Maximum Bandwidth are instead configured on the Hyper-V Virtual Switch because Management, Live Migration, and Cluster each have their own virtual network adapter that is connected to the switch.

You can enable Classification and Tagging and PFC in the management operating system. You can also enable Classification and Tagging from within the VMs, if they are trusted.

QoS Minimum Bandwidth Best Practices

This topic provides information about configuring QoS Minimum Bandwidth, and includes the following sections.

- Guidelines for using Minimum Bandwidth
- Minimum Bandwidth Modes
- Minimum Bandwidth and Data Center Bridging (DCB)
- Configuring QoS and NIC Teaming
- Related topics

Guidelines for using Minimum Bandwidth

Following are guidelines for configuring Minimum Bandwidth.

1. **Keep the sum of the weights near or under 100**. The larger the value for weight, the higher rounding error or quantization error there is when QoS Packet Scheduler partitions the bandwidth among workloads. The best example to illustrate this guidance is if there are 20 VMs that you want to share the network bandwidth, rather than assigning a weight of 10 to each of them (in which case the sum would be 200), assign each VM a weight of 1.
2. **Assign a relatively large weight to critical workloads even if they don't require that percentage of bandwidth**. The best example to illustrate this guidance is to think of Management and Cluster heartbeat traffic on a computer that is running the Hyper-V server role. They rarely require more than 1 or 2 percent of bandwidth over a 10GbE NIC. But they're critical. So instead of assigning to each of them a weight of 1 or 2, you should assign each a weight of 5 or more.
3. **Gap the weight assignment to differentiate the level of service to be provided**. In other words, when assigning weight values to VMs, do not use consecutive numbers, such as 1, 2, and 3. For example, if you have VMs named Gold, Silver, and Bronze on the same Hyper-V Virtual Switch, assign them a respective weight of 5 for Gold, 3 for Silver, and 1 for Bronze - rather than a weight of 3, 2 and 1, respectively.
4. **Ensure that traffic that is not specifically filtered out is also accounted for with a weight assignment**. For example, review the configuration 2 NICs without NIC Teaming, where Storage, Live Migration, and Cluster are filtered out with a workload-specific filter. In addition, the rest of the network traffic is grouped with a wild card filter as one single flow that has a weight assignment.

Minimum Bandwidth Modes

When you create a Hyper-V Virtual Switch by using Windows PowerShell, you can choose between one of the following Minimum Bandwidth configuration modes, which determine how you can quantify Minimum Bandwidth for a workload.

- By weight, which is an integer in the range between 1 and 100
- By bits per second (BPS)

It is recommended that you configure Minimum Bandwidth by using weight rather than BPS. Minimum Bandwidth specified by weight is more flexible and it is compatible with other features, such as Live Migration and NIC Teaming.

For example, consider a VM that is configured to have a Minimum Bandwidth of 500 MBPS. In the Live Migration case, if a destination computer running Hyper-V doesn't have 500 MBPS spare bandwidth to satisfy this VM, Live Migration fails because the service is disrupted. If Minimum Bandwidth is specified by weight, this problem does not occur because weight specifies a percentage of bandwidth rather than a specific number of BPS that are required for the traffic. Live-migrating a VM with a Minimum Bandwidth weight of 3, for example, only increases the sum of the weights by 3 on the destination computer that is running Hyper-V. In this configuration, the Live Migration service does not fail.

If you must specify minimum bandwidth in an absolute number in BPS, use the following guidelines.

- The minimum unit is 1% of the link capacity. For example, if a Hyper-V Virtual Switch is connected to a 10 GB NIC, the minimum amount of bandwidth that you can allocate to a virtual network adapter connected to the switch is 100 MBPS.
- The number that you specify for BPS is rounded to nearest percent point. For example, 234Mbps is rounded down to 200Mbps, which is 2% of a 10GbE NIC throughput.

When you create a Hyper-V Virtual Switch by using Windows PowerShell, you can use the –MinimumBandwidthMode parameter of the New-VMSwitch cmdlet to specify whether Minimum Bandwidth is configured by weight or by BPS. The possible values for the –MinimumBandwidthMode parameter are “Weight,” to specify weight, and “Absolute,” to specify BPS.

The following example Windows PowerShell command illustrates how to create a Hyper-V Virtual Switch that allows Minimum Bandwidth to be specified by weight.

```
New-VMSwitch "QoS Switch by Weight" -NetAdapter "External NIC" -MinimumBandwidthMode Weight
```

The following example Windows PowerShell command illustrates how to create a Hyper-V Virtual Switch that allows Minimum Bandwidth to be specified by bits per second.

```
New-VMSwitch "QoS Switch by BPS" -NetAdapter "External NIC" -MinimumBandwidthMode Absolute
```

For more information, see the Windows PowerShell Hyper-V reference topic New-VMSwitch.

Minimum Bandwidth and Data Center Bridging (DCB)

Minimum Bandwidth, which is enforced by QoS Packet Scheduler in Windows, and DCB, which is enforced by a capable NIC, have the same functionality – to ensure that a workload receives its fair share of bandwidth when network congestion occurs. QoS Minimum Bandwidth and DCB are not designed to work together, however.

Important
Do not enable both Minimum Bandwidth and DCB for workloads that share the same networking stack or NIC.

The figure below illustrates that you can enable either Minimum Bandwidth or DCB on a computer running Hyper-V, but you cannot enable both technologies at the same time.

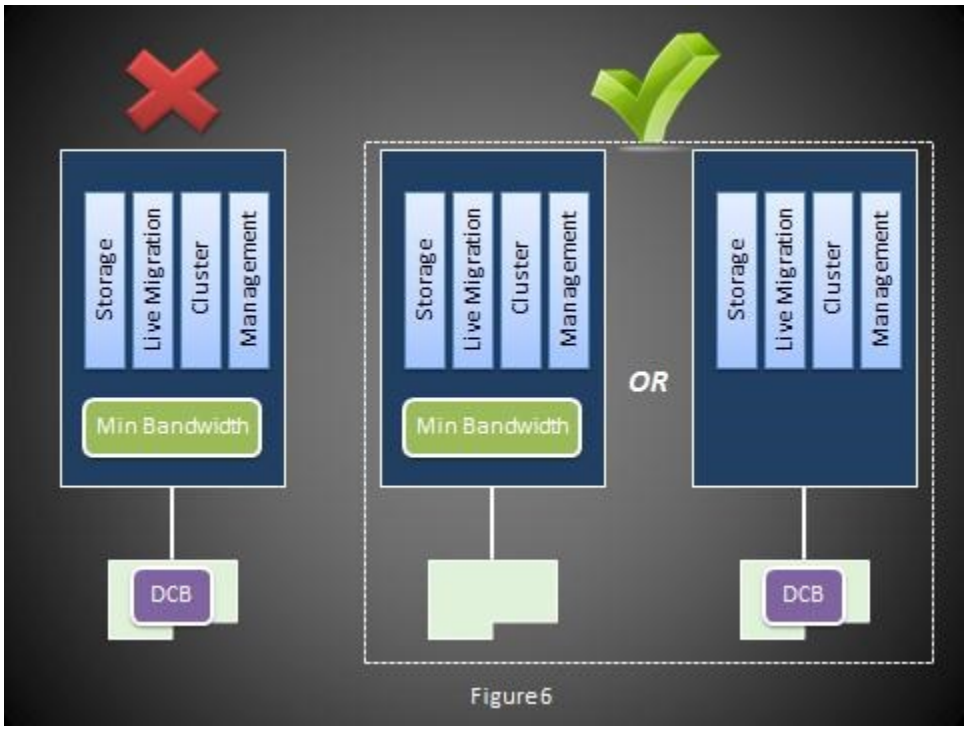

Figure6

You can enable both Minimum Bandwidth and DCB on one server running Windows Server 2012 for two separate networking stacks or NICs.

Configuring QoS and NIC Teaming

NIC Teaming is a new feature in Windows Server 2012. When you configure multiple NICs into a NIC team, the following benefits are provided.

- Bandwidth aggregation (i.e. bandwidth distribution across team members)
- Traffic failover to maintain connectivity in the event of a network component failure

NIC teaming supports two distribution algorithms.

- Hyper-V Virtual Switch port, where the switch port to which a VM's or the management operating system's virtual network adapter is connected is used for hashing traffic that is received on the switch port.
- Address hashing, where the MAC address, the IP address, and the Transport port in the network packet header are used for hashing the traffic.

QoS functionality and NIC Teaming work together in the following ways.

- **Classification and Tagging**. This QoS feature works seamlessly with NIC Teaming. Network traffic is classified and priority-tagged by Windows network components before it is hashed to members of a NIC team.
- **Priority based Flow Control**. This QoS feature works seamlessly with NIC Teaming because PFC occurs at the NIC hardware. It is highly recommended though that it is enabled or disabled consistently across all member NICs of a team.

- **Bandwidth Management**. The interaction between Bandwidth Management and NIC Teaming is as follows.
  - **Maximum Bandwidth**. Maximum Bandwidth works seamlessly with NIC Teaming. QoS Packet Scheduler throttles a traffic flow, which can be an application, a service or a workload, before the traffic, which already conforms to the throttling rate, is handed over to NIC Teaming, which hashes it to a team member for transmission.
  - **Minimum Bandwidth**. As in the case of throttling, QoS Packet Scheduler partitions bandwidth among traffic flows according to the weight assigned to each flow before the traffic is hashed to a team member for transmission. There is a problem, however, when the traffic distribution algorithm of the NIC team that is bound to a Hyper-V Virtual Switch is Hyper-V switch port. In such a case, a traffic flow contains all traffic from a VM (assuming that this VM has only one virtual network adapter).

    It is possible for NIC Teaming to hash VMs with large weights to one team member and VMs with small weights to another. When that occurs, the amount of bandwidth each VM receives when network congestion occurs does not always reflect the bandwidth weight it is assigned. For example, consider a server that hosts four VMs, two with a minimum bandwidth weight of 3 each and two with a weight of 1 each. If the first two VMs are hashed to one team member of a two-NIC team and the last two VMs to another team member, essentially each VM gets equal share (25%) of the total bandwidth of this NIC team when both team members are congested.

    The problem described above is a less concern when there are many VMs on a Hyper-V Virtual Switch and/or the traffic distribution algorithm of a NIC team is Address hashing. In those cases, there are many flows for NIC Teaming to distribute (e.g. each TCP connection is a flow in Address hashing mode), and statistically speaking, VM traffic is more likely evenly distributed across team members of the NIC team. When half of the VMs with weight of 3 and half of the VMs with weight of 1 are hashed to one team member, the rest to the other team member, Minimum Bandwidth is enforced more effectively when congestion occurs.

    To overcome the problem caused by hashing when the traffic distribution algorithm is Hyper-V Virtual Switch port, you can carefully script a specific order to bring up VMs on a computer running Hyper-V. Because NIC Teaming hashes VM's virtual network adapters to team members in a round robin fashion, as long as groups of VMs with the same weight are brought up at the same time, the weight distribution across a NIC team is equal.
  - **Hardware Enforced Minimum Bandwidth**. The same problem described above can occur when traffic distribution is skewed across multiple DCB capable NICs, even though bandwidth is managed by the NIC hardware.

The table below summarizes how QoS works with NIC Teaming.

QoS Feature	Functionality with NIC Teaming
Maximum Bandwidth	Works seamlessly

QoS Feature	Functionality with NIC Teaming
Minimum Bandwidth	Does not work seamlessly; uneven traffic distribution can affect weight enforcement
Hardware Enforced Minimum Bandwidth (DCB)	Does not work seamlessly; uneven traffic distribution can affect weight enforcement
Classification and Tagging	Works seamlessly
Priority-based Flow Control	Works seamlessly

When you configure Minimum Bandwidth in absolute numbers, there is another level of complexity when the deployment involves NIC Teaming. For example, if a Hyper-V switch is bound to a two 1Gb NIC team, and the NIC team appears to provide 2 Gbps available bandwidth. However you cannot allocate more than 1 Gbps bandwidth to VMs. If VMs were configured with more bandwidth, they wouldn't be able to satisfy the requirements of those VMs in the event of a NIC failure.

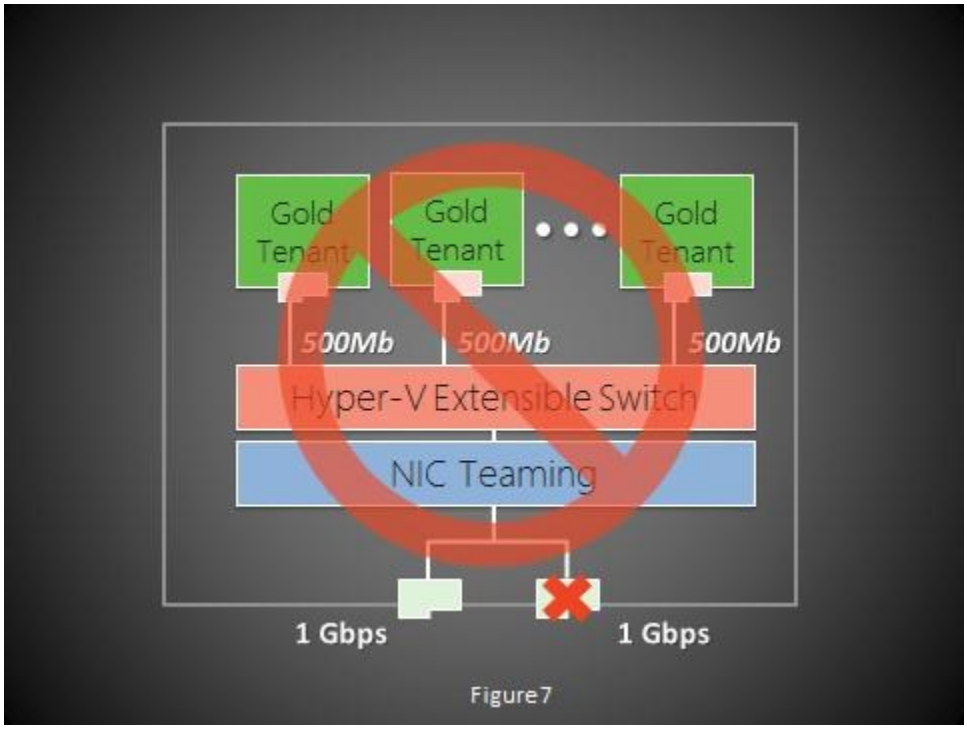

Figure7

In contrast, if Minimum Bandwidth were specified by weight in the above configuration, there would be no issue because all traffic would be hashed to the remaining NIC, and bandwidth for each VM would be partitioned proportionally according to its weight. This is another reason that for most deployments you should use the weight mode when configuring Minimum Bandwidth.

Related topics

Following are links to content that will assist you in configuring QoS and related technologies.

- Network Quality of Service (QoS) Cmdlets in Windows PowerShell
- Data Center Bridging (DCB) Quality of Service (QoS) Cmdlets in Windows PowerShell
- Network Adapter Cmdlets in Windows PowerShell
- New-VMSwitch
- Set-VMNetworkAdapter
- Hyper-V Cmdlets in Windows PowerShell

Policy-based Quality of Service (QoS)

Notes

In addition to this topic, the following Policy-based QoS documentation is also available.

- Configuring Policy-based Quality of Service (QoS)

As traffic increases on a network, it becomes increasingly important for IT departments to balance network performance with the cost of service. However, network traffic is not easily prioritized and managed. Mission-critical and latency-sensitive applications must compete for bandwidth against lower priority traffic. At the same time, some users and computers with specific network performance requirements might require differentiated service levels.

Such challenges of providing cost-effective, predictable network performance levels often first appear over wide area network (WAN) connections or with latency-sensitive applications, like voice over IP (VoIP) and video. However, the end-goal of providing predictable network service levels applies to any network environment (for example, an enterprises' local area network), and to more than VoIP applications, such as your company's custom line-of-business applications.

For computers running Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows 7, Windows Server 2008, and Windows Vista, you can use Policy-based Quality of Service (QoS) to manage traffic to offer better user experiences, control bandwidth costs, or more finely negotiate service levels with bandwidth providers or business departments. Policy-based QoS provides network control based on applications, users, and computers. Applications do not need to be written for specific application programming interfaces (APIs), providing you with the ability to use QoS with existing applications. Additionally, Policy-based QoS takes advantage of your existing management infrastructure, because Policy-based QoS is built into Group Policy.

You can specify QoS policies that define priority through a Differentiated Services Code Point (DSCP) value. The DSCP applies a value (0–63) within the Type of Service (TOS) field in an IPv4 packet's header and within the Traffic Class field in IPv6. This DSCP value provides classification at the Internet Protocol (IP) level, which routers can use to decide queuing behavior. You can also limit an application's outbound network traffic by specifying a throttle rate.

For example, you can configure routers to place packets with specific DSCP values into one of three queues: high-priority, best-effort, or lower-than-best effort. Therefore, mission-critical

network traffic, which is in the high-priority queue, gets preference before other traffic. The QoS policy that defines throttling limits the rate of outbound network traffic. For example, an IT department might implement a service level agreement that specifies that a file server can never provide downloads beyond a specific rate to manage WAN costs.

You can also use Policy-based QoS to apply DSCP values and throttle rates for outbound network traffic to the following:

- Sending application and directory path
- Source and destination IP addresses, including support for address prefixes
- Source and destination ports and port ranges
- Protocol (Transmission Control Protocol [TCP] and User Datagram Protocol [UDP])
- Specific groups of users or computers (through deployment in Group Policy)

By using these controls, you can specify a QoS policy with a DSCP value of 46 for a VoIP application, enabling routers to place those packets in a low-latency queue, or you can use a QoS policy to throttle a set of servers' outbound traffic to 512 kilobytes per second (KBps) when sending from TCP port 443. Or, as in the expanded example in the next section, QoS policy can be applied to a particular application that has special bandwidth requirements.

Advantages of Policy-based QoS

With Policy-based QoS, IT administrators can configure and enforce QoS policies that cannot be configured on routers and switches.

1. **Level of detail:** It is difficult to create user-level QoS policies on routers or switches, especially if the user's computer is either configured by using dynamic IP address assignment or if the computer is not connected to fixed switch or router ports, as is frequently the case with portable computers. In contrast, Policy-based QoS makes it easier to configure a user-level QoS policy on a domain controller and propagate it to the user's computer, regardless of where or how the computer connects to the network. For example, you can configure a user-level QoS policy that will apply to a computer, no matter where the user logs on: for example, in either a main office or in a branch office. Similarly, the same user level QoS policy will apply no matter how the user connects to the network: for example, by either the wired Ethernet network or by Wi-Fi.
2. **Security:** If your IT department encrypts users' traffic from end to end by using Internet Protocol security (IPsec), you cannot classify the traffic on routers based on any information above the IP layer in the packet (for example, a TCP port). However, by using Policy-based QoS, you can classify packets at the end device to indicate the priority of the packets in the IP header before the IP payloads are encrypted and the packets sent out.
3. **Performance:** Some QoS functions, such as throttling, are better performed when they are closer to the source. Policy-based QoS moves such QoS functions to where they can be closest to the source.
4. **Manageability:** Policy-based QoS enhances the network manageability in two ways:

First, because it is based on Group Policy, you can use Policy-based QoS to configure and manage a set of user/computer QoS policies whenever necessary, and on one central domain-controller computer.

Second, Policy-based QoS facilitates user/computer configuration by providing a mechanism to specify policies by URL, as opposed to specifying policies based on the IP addresses of each of the servers where QoS policies need to be applied.

For example, assume your network has a cluster of servers that share a common URL. By using Policy-based QoS, you need only to create one policy based on that common URL, as opposed to creating one policy for each server in the cluster, with each policy based on the IP address of each server.

Scenario 1: Prioritizing a line-of-business application

In this scenario, an IT department adds QoS to provide better network performance for a key set of users and the mission-critical applications users depend on, but also needs to minimize the WAN link costs. The IT department decides to prioritize specific applications by using DSCP values to classify network traffic, and to configure its routers to provide preferential treatment for higher priority traffic. Later, they might consider using this classification to negotiate service levels for the leased WAN links.

In addition to DSCP values, the QoS policies can specify a throttle rate. Throttling limits all outbound traffic that matched the QoS policy to a specific send rate.

The first mission-critical application to use Policy-based QoS is a company-wide enterprise resource planning (ERP) application. The ERP application is hosted on several computers running Windows Server 2012 in the data center, which are part of an organization unit (OU) for computers. While many groups within the company access the ERP application, the finance group requires differentiated performance, because the finance group depends on this application when dealing with customers. The client-side component for the ERP application is installed on computers running Windows 8, Windows 7, and Windows Vista.

The following example illustrates the clients and servers in a prioritization scenario. The subsequent section titled "Configuring Policy-based QoS" walks through the wizard steps for creating QoS policies. In this example, the IT administrator selects the Group Policy Object (GPO) on which the QoS policy will be deployed. Through the QoS policy wizard, the IT administrator creates a QoS policy for the group of servers called "Server LOB policy" that specifies a high-priority DSCP value of 44 for all applications, any IP address, TCP and UDP, and port number. The QoS policy is applied only to the line-of-business (LOB) servers by linking the GPO to the OU that contains only these servers, via the Group Policy Management Console (GPMC) tool. This initial server LOB policy applies the high-priority DSCP value whenever the computer sends network traffic. This QoS policy can later be edited (in the Group Policy Object Editor tool) to include the ERP application's port numbers, which limits the policy to apply only when the specified port number is used.

To ensure that the finance group can support their customers, a QoS policy needs to classify these users' traffic as higher priority. However, the policy should not apply when members of the finance group use applications other than the ERP application. Thus, the IT department defines a second QoS policy called "Client LOB policy" in the Group Policy Object Editor tool that applies a DSCP value of 60 when the finance user group runs the ERP application.

A separate backup application is running on all computers. To ensure the backup application's traffic does not use all available network resources, a backup data policy is created. This backup policy specifies a DSCP value of 1 based on the executable name for this backup application, backup.exe. A third GPO is created and deployed for all client computers in the domain. Whenever the backup application sends data, the low-priority DSCP value is applied, even if it originates from computers in the finance department.

Note that traffic without a QoS policy sends with a DSCP value of 0.

The following table summarizes the QoS policies for this scenario.

Policy name	DSCP value	Throttle rate	Applied to organization units	Description
[No policy]	0	None	[No deployment]	Best effort (default) treatment for unclassified traffic.
Backup data	1	None	All clients	Applies a low-priority DSCP value for this bulk data.
Server LOB	44	None	Computer OU for ERP servers	Applies high-priority DSCP for ERP server traffic
Client LOB	60	None	Finance user group	Applies high-priority DSCP for ERP client traffic

Note

DSCP values are represented in decimal form.

With QoS policies defined and applied by using Group Policy, outbound network traffic receives the policy-specified DSCP value. Routers then provide differential treatment based on these DSCP values by using queuing. For this IT department, the routers are configured with four queues: high-priority, middle-priority, best-effort, and low-priority.

When traffic arrives at the router with DSCP values from "Server LOB policy" and "Client LOB policy," the data is placed into high-priority queues. Traffic with a DSCP value of 0 receives a best-effort level of service. Packets with a DSCP value of 1 (from the backup application) receive low-priority treatment.

Prerequisites for prioritizing a line-of-business application

To complete this task, ensure that you meet the following requirements:

- The computers involved are running Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, or Windows Vista.
- The computers involved are members of an Active Directory Domain Services (AD DS) domain so that they can be controlled by using Group Policy.
- TCP/IP networks are set up with routers configured for DSCP (RFC 2474). For more information, see http://www.ietf.org/rfc/rfc2474.txt.
- Administrative credentials requirements are met.

Administrative credentials

To complete this task, you must be able to create and deploy Group Policy Objects.

Setting up the test environment for prioritizing a line-of-business application

To set up the test environment, complete the following tasks:

- Create an AD DS domain controller with clients and users grouped into organization units.
- Configure the routers to differentially queue based on DSCP values. For example, DSCP value 44 enters a "Platinum" queue and all others are weighted-fair-queued.

> **Note**
> DSCP values can be viewed by using network captures, such as tools like NetMon, and by observing the TOS field.

Steps for prioritizing a line-of-business application

To prioritize a line-of-business application, complete the following tasks:

1. Create and link a Group Policy Object (GPO) with a QoS policy.
2. Configure the routers to differentially treat a line-of-business application (by using queuing) based on the selected DSCP values. The procedures of this task will vary depending upon the type of routers you have.

Scenario 2: Prioritizing an HTTP server application

Many enterprise applications are developed for and hosted on Internet Information Services (IIS) web servers. Typically, the application is accessed from browsers on client computers. For these deployments, IT departments can benefit from prioritizing the network traffic that is associated with web-based applications. In Windows Server 2012, Windows 8, Windows 7, and Windows Server 2008 R2, Policy-based QoS provides a new feature, known as URL-based Policies, that enables administrators to place HTTP responses—to applications that are built on top of HTTP—subject to QoS control.

In this scenario, assume that you manage a set of IIS servers that host training videos for all your organization's employees. Your objective is to ensure that the traffic from these video servers won't overwhelm your network, and ensure that video traffic is differentiated from voice and data traffic on the network. The task is similar to the task in Scenario 1. You will design and configure the traffic management settings, such as the DSCP value for the video traffic, and the throttling rate the same as you would for the line-of-business applications. Only when specifying the traffic, instead of providing the application name, you need only to enter the URL to which your HTTP server application will respond: for example, https://hrweb/training.

Note

URL–based QoS policies apply only to computers running Windows Server 2012, Windows 8, Windows 7, and Windows Server 2008 R2. You cannot use URL-based QoS policies to prioritize network traffic for computers running Windows operating systems that were released prior to Windows 7 and Windows Server 2008 R2.

Precedence rules for URL-based policies

All the following URLs are valid and can be specified in Policy-based QoS and applied simultaneously to a computer or a user:

- http://video
- https://training.hr.mycompany.com
- http://10.1.10.249:8080/tech
- https://*/ebooks

But which one will receive precedence? The rules are simple. URL-based policies are prioritized in a left-to-right reading order. So, from the highest priority to the lowest priority, the URL fields are:

1. 1. URL scheme
2. 2. URL host
3. 3. URL port
4. 4. URL path

Details are as follows:

1. URL scheme

https:// has a higher priority than http://.

2. URL host

From the highest priority to the lowest, they are:

1. Hostname
2. IPv6 address
3. IPv4 address
4. Wildcard

In the case of hostname, a hostname with more dotted elements (more depth) has a higher priority than a hostname with fewer dotted elements. For example, among the following hostnames:

- video.internal.training.hr.mycompany.com (depth = 6)
- selfguide.training.mycompany.com (depth = 4)
- training (depth = 1)
- library (depth = 1)

video.internal.training.hr.mycompany.com has the highest priority, and **selfguide.training.mycompany.com** has the next highest priority. **Training** and **library** share the same lowest priority.

3. URL port

A specific or an implicit port number has a higher priority than a wildcard port.

4. URL path

Like a hostname, a URL path may consist of multiple elements. The one with more elements always has a higher priority than the one with less. For example, the following paths are listed by priority:

1. /ebooks/tech/windows/networking/qos
2. /ebooks/tech/windows/
3. /ebooks
4. /

If a user chooses to include all subdirectories and files following a URL path, this URL path will have a lower priority than it would have if the choice were not made.

A user may also choose to specify a destination IP address in a URL-based policy. The destination IP address has a lower priority than any of the four URL fields described previously.

URL-based policy vs. Quintuple policy

A Quintuple policy is specified by protocol ID, source IP address, source port, destination IP address, and destination port. A Quintuple policy always has a higher precedence than any URL-based policy. If a Quintuple policy is already applied for a user, a new URL-based policy will not cause conflicts on any of that user's client computers.

Configuring Policy-based QoS

Based on the previous line-of-business application example, you can see that QoS for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for mission-critical applications. By using QoS mechanisms, network administrators can use existing resources efficiently and ensure the required level of service without reactively expanding or over-provisioning their networks.

In Windows operating systems, Policy-based QoS combines the functionality of standards-based QoS with the manageability of Group Policy. Configuration of this combination makes for easy application of QoS policies to Group Policy Objects. Windows includes a Policy-based QoS Wizard to help you:

- Create a QoS policy.
- View, edit, or delete a QoS policy.

Create a QoS policy

Before you create a QoS policy, it is important that you understand the two key QoS controls that are used to manage network traffic:

- DSCP value
- Throttle rate

Prioritizing traffic with DSCP

As noted in the previous line-of-business application example, you can define the priority of outbound network traffic by using **Specify DSCP Value** to configure a QoS policy with a specific DSCP value. As described in RFC 2474, DSCP allows values from 0 to 63 to be specified within the TOS field of an IPv4 packet and within the Traffic Class field in IPv6. Network routers use the DSCP value to classify network packets and to queue them appropriately.

Note

By default, Windows traffic has a DSCP value of 0.

The number of queues and their prioritization behavior needs to be designed as part of your organization's QoS strategy. For example, your organization may choose to have five queues: latency-sensitive traffic, control traffic, business-critical traffic, best-effort traffic, and bulk-data-transfer traffic.

Throttling traffic

Along with DSCP values, throttling is another key control for managing network bandwidth. As mentioned earlier, you can use the **Specify Throttle Rate** setting to configure a QoS policy with a specific throttle rate for outbound traffic. By using throttling, a QoS policy limits the outgoing network traffic to a specified throttle rate. Both DSCP marking and throttling can be used together to manage traffic effectively.

Note

By default, the **Specify Throttle Rate** check box is not selected.

To create a QoS policy, edit the settings of a Group Policy Object (GPO) from within the Group Policy Management Console (GPMC) tool. GPMC then opens the Group Policy Object Editor.

QoS policy names must be unique. How policies are applied to servers and end users depends on where the QoS policy is stored in the Group Policy Object Editor:

- A QoS policy in Computer Configuration\Windows Settings\Policy-based QoS applies to computers, regardless of the user that is currently logged on. You typically use computer-based QoS policies for server computers.
- A QoS policy in User Configuration\Windows Settings\Policy-based QoS applies to users after they have logged on, regardless of which computer they have logged on to.

To create a new QoS policy with the Policy-based QoS wizard

- In Group Policy Object Editor, right-click either of the **Policy-based QoS** nodes, and then click **Create a new policy**.

Wizard page 1—Policy Profile

On the first page of the Policy-based QoS wizard, you can specify a policy name and configure how QoS controls outgoing network traffic.

To configure the Policy Profile page of the QoS-based Policy wizard

1. In **Policy name**, type a name for the QoS policy. The name must uniquely identify the policy.
2. Optionally, use **Specify DSCP Value** to enable DSCP marking, and then configure a DSCP value between 0 and 63.
3. Optionally, use **Specify Throttle Rate** to enable traffic throttling and configure the throttle rate. The throttle rate value must be greater than 1 and you can specify units of kilobytes per second (KBps) or megabytes per second (MBps).
4. Click **Next**.

Wizard page 2—Application Name

In the second page of the Policy-based QoS wizard you can apply the policy to all applications, to a specific application as identified by its executable name, to a path and application name, or to the HTTP server applications that handle requests for a specific URL.

- **All applications** specifies that the traffic management settings on the first page of the Policy-based QoS wizard apply to all applications.
- **Only applications with this executable name** specifies that the traffic management settings on the first page of the Policy-based QoS wizard are for a specific application. The executable file name and must end with the .exe file name extension.
- **Only HTTP server applications responding to requests for this URL** specifies that the traffic management settings on the first page of the Policy-based QoS wizard apply to certain HTTP server applications only.

Optionally, you can enter the application path. To specify an application path, include the path with the application name. The path can include environment variables. For example, %ProgramFiles%\My Application Path\MyApp.exe, or c:\program files\my application path\myapp.exe.

Note

The application path cannot include a path that resolves to a symbolic link.

The URL must conform to RFC 1738 (http://tools.ietf.org/html/rfc1738), that is, in the form of "http[s]://<hostname>:<port>/<url-path>". You can use a wildcard, '*', for <hostname> and/or <port>, e.g. http://training.*/, https://*.*, but the wildcard cannot denote a substring of <hostname> or <port>.

In other words, neither http://my*site/ nor https://*training*/ is valid. Optionally, you can check **Include subdirectories and files** to perform matching on all subdirectories and files following a URL. For example, if this option is checked and the URL is "http://training", Policy-based QoS will consider requests for "http://training/video" a good match.

To configure the Application Name page of the Policy-based QoS wizard

1. In **This QoS policy applies to**, select either **All applications** or **Only applications with this executable name**.
2. If you select **Only applications with this executable name**, specify an executable name ending with the .exe file name extension.
3. Click **Next**.

Wizard page 3—IP Addresses

In the third page of the Policy-based QoS wizard you can specify IP address conditions for the QoS policy, including the following:

- All source IPv4 or IPv6 addresses or specific source IPv4 or IPv6 addresses
- All destination IPv4 or IPv6 addresses or specific destination IPv4 or IPv6 addresses

If you select **Only for the following source IP address** or **Only for the following destination IP address**, you must type one of the following:

- An IPv4 address, such as 192.168.1.1
- An IPv4 address prefix using network prefix length notation, such as 192.168.1.0/24
- An IPv6 address, such as 3ffe:ffff::1
- An IPv6 address prefix, such as 3ffe:ffff::/48

If you select both **Only for the following source IP address** and **Only for the following destination IP address**, both addresses or address prefixes must be either IPv4- or IPv6-based.

If you specified the URL for HTTP server applications in the previous wizard page, you'll notice that the source IP address for the QoS policy on this wizard page is grayed out. That's because the source IP address is the HTTP server address and it is not configurable here. On the other hand, you can still customize the policy by specifying the destination IP address. This makes it possible for you to create different policies for different clients by using the same HTTP server applications.

To configure the IP Addresses page of the Policy-based QoS wizard

1. In **This QoS policy applies to** (source), select **Any source IP address** or **Only for the following IP source address**.
2. If you selected **Only the following IP source address**, specify an IPv4 or IPv6 address or prefix.
3. In **This QoS policy applies to** (destination), select **Any destination address** or **Only for the following IP destination address.**
4. If you selected **Only for the following IP destination address**, specify an IPv4 or IPv6 address or prefix that corresponds to the type of address or prefix specified for the source address.
5. Click **Next**.

Wizard page 4—Protocols and Ports

On the fourth page of the Policy-based QoS wizard, you can specify the types of traffic and the ports that are controlled by the settings on the first page of the wizard. You can specify:

- TCP traffic, UDP traffic, or both
- All source ports, a range of source ports, or a specific source port
- All destination ports, a range of destination ports, or a specific destination port

To configure the Protocols and Ports page of the Policy-based QoS wizard

1. In **Select the protocol this QoS policy applies to**, select **TCP**, **UDP**, or **TCP and UDP**.
2. In **Specify the source port number**, select **From any source port** or **From this source port number**.
3. If you selected **From this source port number**, type a port number between 1 and 65535.

   Optionally, you can specify a port range, in the format of "*Low*:*High*," where *Low* and *High* represent the lower bounds and upper bounds of the port range, inclusively. *Low* and *High* each must be a number between 1 and 65535. No space is allowed between the colon (:) character and the numbers.
4. In **Specify the destination port number**, select **To any destination port** or **To this destination port number**.
5. If you selected **To this destination port number** in the previous step, type a port number between 1 and 65535.

To complete the creation of the new QoS policy, click **Finish** on the **Protocols and Ports** page of the Policy-based QoS wizard. When completed, the new QoS policy is listed in the details pane of the Group Policy Object Editor.

To apply the QoS policy settings to users or computers, link the GPO in which the QoS policies are located to an Active Directory Domain Services container, such as a domain, a site, or an organizational unit (OU).

View, edit, or delete a QoS policy

The pages of the Policy-based QoS wizard described previously correspond to the properties pages that are displayed when you view or edit the properties of a policy.

To view the properties of a QoS policy

- Right-click the policy name in the details pane of the Group Policy Object Editor, and then click **Properties**.

  The Group Policy Object Editor displays the properties page with the following tabs:

  - Policy Profile
  - Application Name
  - IP Addresses
  - Protocols and Ports

To edit a QoS policy

- Right-click the policy name in the details pane of the Group Policy Object Editor, and then click **Edit existing policy**.

  The Group Policy Object Editor displays the **Edit an existing QoS policy** dialog box.

To delete a QoS policy

- Right-click the policy name in the details pane of the Group Policy Object Editor, and then click **Delete policy**.

QoS Policy GPMC Reporting

After you have applied a number of QoS policies across your organization, it may be useful or necessary to periodically review how the policies are applied. A summary of the QoS policies for a specific user or computer can be viewed by using GPMC reporting.

To run the Group Policy Results Wizard for a report of QoS policies

- In GPMC, right-click the **Group Policy Results** node, and then select the menu option for **Group Policy Results Wizard.**

After Group Policy results are generated, click the **Settings** tab. On the **Settings** tab, the QoS policies can be found under the "Computer Configuration\Windows Settings\Policy-based QoS" and "User Configuration\Windows Settings\Policy-based QoS" nodes.

On the **Settings** tab, the QoS policies are listed by their QoS policy names with their DSCP value, throttle rate, policy conditions, and winning GPO listed in the same row..

The Group Policy results view uniquely identifies the winning GPO. When multiple GPOs have QoS policies with the same QoS policy name, the GPO with the highest GPO precedence is applied. This is the winning GPO. Conflicting QoS policies (identified by policy name) that are attached to a lower priority GPO are not applied. Note the GPO priorities define which QoS

policies are deployed in the site, domain, or OU, as appropriate. After deployment, at a user or computer level, the QoS Policy Precedence Rules determine which traffic is allowed and blocked.

The QoS policy's DSCP value, throttle rate, and policy conditions are also visible in Group Policy Object Editor (GPOE)

Advanced settings for roaming and remote users

With Policy-based QoS, the goal is to manage traffic on an enterprise's network. In mobile scenarios, users might be sending traffic on or off the enterprise network. Because QoS policies are not relevant while away from the enterprise's network, QoS policies are enabled only on network interfaces that are connected to the enterprise for Windows 8, Windows 7, or Windows Vista.

For example, a user might connect her portable computer to her enterprise's network via virtual private network (VPN) from a coffee shop. For VPN, the physical network interface (such as wireless) will not have QoS policies applied. However, the VPN interface will have QoS policies applied because it connects to the enterprise. If the user later enters another enterprise's network that does not have an AD DS trust relationship, QoS policies will not be enabled.

Note that these mobile scenarios do not apply to server workloads. For example, a server with multiple network adapters might sit on the edge of an enterprise's network. The IT department might choose to have QoS policies throttle traffic that egresses the enterprise; however, this network adapter that sends this egress traffic does not necessarily connect back to the enterprise network. For this reason, QoS policies are always enabled on all network interfaces of a computer running Windows Server 2012.

Note

Selective enablement only applies to QoS policies and not to the Advanced QoS settings discussed next in this document.

Advanced QoS settings

Advanced QoS settings provide additional controls for IT administrators to manage computer network consumption and DSCP markings. Advanced QoS settings apply only at the computer level, whereas QoS policies can be applied at both the computer and user levels.

To configure advanced QoS settings

1. Click **Computer Configuration**, and then click **Windows Settings in Group Policy**.
2. Right-click **Policy-based QoS**, and then click **Advanced QoS Settings**.

   The following figure shows the two advanced QoS settings tabs: **Inbound TCP Traffic** and **DSCP Marking Override**.

Note

Advanced QoS Settings are computer-level Group Policy settings.

Advanced QoS settings: inbound TCP traffic

Inbound TCP Traffic controls the TCP bandwidth consumption on the receiver's side, whereas QoS policies affect the outbound TCP and UDP traffic. By setting a lower throughput level on the **Inbound TCP Traffic** tab, TCP will limit the size of its advertised TCP receive window. The effect of this setting will be increased throughput rates and link utilization for TCP connections with higher bandwidths or latencies (bandwidth delay product). By default, computers running Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, and Windows Vista are set to the maximum throughput level.

The TCP receive window has changed in Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, and Windows Vista from previous versions of Windows. Previous versions of Windows limited the TCP receive-side window to a maximum of 64 kilobytes (KB), whereas Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, and Windows Vista dynamically size the receive-side window up to 16 megabytes (MB). In the Inbound TCP Traffic control, you can control the inbound throughput level by setting the maximum value to which the TCP receive-window can grow. The levels correspond to the following maximum values.

Inbound throughput level	Maximum
0	64 KB
1	256 KB
2	1 MB
3	16 MB

The actual window size may be a value equal to or smaller than the maximum, depending on network conditions.

To set the TCP receive-side window

1. In Group Policy Object Editor, click **Local Computer Policy**, click **Windows Settings**, right click **Policy-based QoS**, and then click **Advanced QoS Settings**.
2. In **TCP Receiving Throughput**, select **Configure TCP Receiving Throughput**, and then select the level of throughput that you want.
3. Link the GPO to the OU.

Advanced QoS settings: DSCP Marking Override

DSCP Marking Override restricts the ability of applications to specify—or "mark"—DSCP values other than those specified in QoS policies. By specifying that applications are allowed to set DSCP values, applications can set non-zero DSCP values. By specifying **Ignore**, applications that use QoS APIs will have their DSCP values set to zero, and only QoS policies can set DSCP values. By default, computers running Windows Server 2012, Windows 8, Windows

Server 2008 R2, Windows Server 2008, and Windows Vista and allow applications to specify DSCP values; applications and devices that do not use the QoS APIs are not overridden.

Wireless Multimedia and DSCP values

The Wi-Fi Alliance has established a certification for Wireless Multimedia (WMM) that defines four access categories (WMM_AC) for prioritizing network traffic transmitted on a Wi-Fi wireless network. The access categories include (in order of highest-to-lowest priority): voice, video, best effort, and background; respectively abbreviated as VO, VI, BE, and BK. The WMM specification defines which DSCP values correspond with each of the four access categories:

DSCP Value	WMM Access Category
48-63	Voice (VO)
32-47	Video (VI)
24-31, 0-7	Best effort (BE)
8-23	Background (BK)

In Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, and Windows Vista QoS policies can be created that use these DSCP values to ensure that portable computers with Wi-Fi Certified™ for WMM wireless adapters receive prioritized handling when associated with Wi-Fi Certified for WMM access points.

QoS Policy Precedence Rules

Similar to GPO's priorities, QoS policies have precedence rules to resolve conflicts when multiple QoS policies apply to a specific set of traffic. For outbound TCP or UDP traffic, only one QoS policy can be applied at a time, which means that QoS policies do not have a cumulative effect, such as where throttle rates would be summed.

In general, the QoS policy with the most matching conditions wins. When multiple QoS policies apply, the rules fall into three categories: user-level versus computer-level; application versus the network quintuple; and among the network quintuple.

By *network quintuple*, we mean the source IP address, destination IP address, source port, destination port, and protocol (TCP/UDP).

User-level QoS policy takes precedence over computer-level QoS policy

This rule greatly facilitates network administrators' management of QoS GPOs, particularly for user group–based policies. For example, if the network admin wants to define a QoS policy for a user group, they can just create and distribute a GPO to that group. They don't have to worry about which computers those users are logged on to and whether those computers will have conflicting QoS policies defined, because, if a conflict exists, the user-level policy always takes precedence.

Note

A user-level QoS policy is only applicable to traffic that is generated by that user. Other users of a specific computer, and the computer itself, will not be subject to any QoS policies that are defined for that user.

Application specificity and taking precedence over network quintuple

When multiple QoS policies match the specific traffic, the more specific policy is applied. Among policies that identify applications, a policy that includes the sending application's file path is considered more specific than another policy that only identifies the application name (no path). If multiple policies with applications still apply, the precedence rules use the network quintuple to find the best match.

Alternatively, multiple QoS policies might apply to the same traffic by specifying non-overlapping conditions. Between the conditions of applications and the network quintuple, the policy that specifies the application is considered more specific and is applied. For example, policy_A only specifies an application name (app.exe), and policy_B specifies the destination IP address 192.168.1.0/24. When these QoS policies conflict (app.exe sends traffic to an IP address within the range of 192.168.4.0/24), policy_A gets applied.

More specificity takes precedence within the network quintuple

For policy conflicts within the network quintuple, the policy with the most matching conditions takes precedence. For example, assume policy_C specifies source IP address "any", destination IP address 10.0.0.1, source port "any", destination port "any", and protocol "TCP". Next, assume policy_D specifies source IP address "any", destination IP address 10.0.0.1, source port "any", destination port 80, and protocol "TCP". Then policy_C and policy_D both match connections to destination 10.0.0.1:80. Because Policy-based QoS applies the policy with the most specific matching conditions, policy_D takes precedence in this example.

However, QoS policies might have an equal number of conditions. For example, several policies may each specify only one (but not the same) piece of the network quintuple. Among the network quintuple, the following order is from higher to lower precedence:

- Source IP address
- Destination IP address
- Source port
- Destination port
- Protocol (TCP or UDP)

Within a specific condition, such as IP address, a more specific IP address is treated with higher precedence; for example, an IP address 192.168.4.1 is more specific than 192.168.4.0/24.

Note

Generally your QoS policies should be designed as specifically as possible to simplify your organization's understanding of which policies are in effect.

Reference

Error and Event Messages

Following are the error and event messages that are associated with Policy-based QoS.

Informational Messages

MessageId	16500
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_MACHINE_POLICY_REFRESH_NO_CHANGE
Language	English
Message	Computer QoS policies successfully refreshed. No changes detected.

MessageId	16501
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_MACHINE_POLICY_REFRESH_WITH_CHANGE
Language	English
Message	Computer QoS policies successfully refreshed. Policy changes detected.

MessageId	16502
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_USER_POLICY_REFRESH_NO_CHANGE
Language	English
Message	User QoS policies successfully refreshed. No changes detected.

MessageId	16503
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_USER_POLICY_REFRESH_WITH_CHANGE
Language	English

Message	User QoS policies successfully refreshed. Policy changes detected.

MessageId	16504
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_TCP_AUTOTUNING_NOT_CONFIGURED
Language	English
Message	The Advanced QoS Setting for inbound TCP throughput level successfully refreshed. Setting value is not specified by any QoS policy. Local computer default will be applied.

MessageId	16505
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_TCP_AUTOTUNING_OFF
Language	English
Message	The Advanced QoS Setting for inbound TCP throughput level successfully refreshed. Setting value is Level 0 (minimum throughput).

MessageId	16506
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_TCP_AUTOTUNING_HIGHLY_RESTRICTED
Language	English
Message	The Advanced QoS Setting for inbound TCP throughput level successfully refreshed. Setting value is Level 1.

MessageId	16507
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_TCP_AUTOTUNING_RESTRICTED
Language	English

Message	The Advanced QoS Setting for inbound TCP throughput level successfully refreshed. Setting value is Level 2.

MessageId	16508
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_TCP_AUTOTUNING_NORMAL
Language	English
Message	The Advanced QoS Setting for inbound TCP throughput level successfully refreshed. Setting value is Level 3 (maximum throughput).

MessageId	16509
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_APP_MARKING_NOT_CONFIGURED
Language	English
Message	The Advanced QoS Setting for DSCP marking overrides successfully refreshed. Setting value is not specified. Applications can set DSCP values independently of QoS policies.

MessageId	16510
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_APP_MARKING_IGNORED
Language	English
Message	The Advanced QoS Setting for DSCP marking overrides successfully refreshed. Application DSCP marking requests will be ignored. Only QoS policies can set DSCP values.

MessageId	16511

Severity	Informational
SymbolicName	EVENT_EQOS_INFO_APP_MARKING_ALLOWED
Language	English
Message	The Advanced QoS Setting for DSCP marking overrides successfully refreshed. Applications can set DSCP values independently of QoS policies.

MessageId	16512
Severity	Informational
SymbolicName	EVENT_EQOS_INFO_LOCAL_SETTING_DONT_USE_NLA
Language	English
Message	Selective application of QoS policies based on domain network category has been disabled. QoS policies will be applied to all network interfaces.

Warning Messages

MessageId	16600
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_TEST_1
Language	English
Message	EQOS: ***Testing***[, with one string] "%2".

MessageId	16601
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_TEST_2
Language	English
Message	EQOS: ***Testing***[, with two strings, string1 is] "%2"[, string2 is] "%3".

MessageId	16602

Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_MACHINE_POLICY_VERSION
Language	English
Message	The computer QoS policy "%2" has an invalid version number. This policy will not be applied.

MessageId	16603
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_USER_POLICY_VERSION
Language	English
Message	The user QoS policy "%2" has an invalid version number. This policy will not be applied.

MessageId	16604
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_MACHINE_POLICY_PROFILE_NOT_SPECIFIED
Language	English
Message	The computer QoS policy "%2" does not specify a DSCP value or throttle rate. This policy will not be applied.

MessageId	16605
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_USER_POLICY_PROFILE_NOT_SPECIFIED
Language	English
Message	The user QoS policy "%2" does not specify a DSCP value or throttle rate. This policy will not be applied.

MessageId	16606

Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_MACHINE_POLICY_QUOTA_EXCEEDED
Language	English
Message	Exceeded the maximum number of computer QoS policies. The QoS policy "%2" and subsequent computer QoS policies will not be applied.

MessageId	16607
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_USER_POLICY_QUOTA_EXCEEDED
Language	English
Message	Exceeded the maximum number of user QoS policies. The QoS policy "%2" and subsequent user QoS policies will not be applied.

MessageId	16608
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_MACHINE_POLICY_CONFLICT
Language	English
Message	The computer QoS policy "%2" potentially conflicts with other QoS policies. See documentation for rules about which policy will be applied.

MessageId	16609
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_USER_POLICY_CONFLICT
Language	English
Message	The user QoS policy "%2" potentially conflicts with other QoS policies. See documentation for rules about which policy will be applied.

MessageId	16610

Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_MACHINE_POLICY_NO_FULLPATH_APPNAME
Language	English
Message	The computer QoS policy "%2" was ignored because the application path cannot be processed. The application path may be invalid, contain an invalid drive letter, or contain a network mapped drive.

MessageId	16611
Severity	Warning
SymbolicName	EVENT_EQOS_WARNING_USER_POLICY_NO_FULLPATH_APPNAME
Language	English
Message	The user QoS policy "%2" was ignored because the application path cannot be processed. The application path may be invalid, contain an invalid drive letter, or contain a network mapped drive.

Error Messages

MessageId	16700
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_MACHINE_POLICY_REFERESH
Language	English
Message	Computer QoS policies failed to refresh. Error code: "%2".

MessageId	16701
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_USER_POLICY_REFERESH
Language	English
Message	User QoS policies failed to refresh. Error code: "%2".

MessageId	16702
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_OPENING_MACHINE_POLICY_ROOT_KEY
Language	English
Message	QoS failed to open the machine-level root key for QoS policies. Error code: "%2".

MessageId	16703
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_OPENING_USER_POLICY_ROOT_KEY
Language	English
Message	QoS failed to open the user-level root key for QoS policies. Error code: "%2".

MessageId	16704
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_MACHINE_POLICY_KEYNAME_TOO_LONG
Language	English
Message	A computer QoS policy exceeds the maximum allowed name length. The offending policy is listed under the machine-level QoS policy root key, with index "%2".

MessageId	16705
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_USER_POLICY_KEYNAME_TOO_LONG
Language	English
Message	A user QoS policy exceeds the maximum allowed name length. The offending policy is listed under the user-level QoS policy root key, with index "%2".

MessageId	16706
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_MACHINE_POLICY_KEYNAME_SIZE_ZERO
Language	English
Message	A computer QoS policy has a zero length name. The offending policy is listed under the machine-level QoS policy root key, with index "%2".

MessageId	16707
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_USER_POLICY_KEYNAME_SIZE_ZERO
Language	English
Message	A user QoS policy has a zero length name. The offending policy is listed under the user-level QoS policy root key, with index "%2".

MessageId	16708
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_OPENING_MACHINE_POLICY_SUBKEY
Language	English
Message	QoS failed to open the registry subkey for a computer QoS policy. The policy is listed under the machine-level QoS policy root key, with index "%2".

MessageId	16709
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_OPENING_USER_POLICY_SUBKEY
Language	English
Message	QoS failed to open the registry subkey for a user QoS policy. The policy is listed under the user-level QoS policy root key, with index "%2".

MessageId	16710
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_PROCESSING_MACHINE_POLICY_FIELD
Language	English
Message	QoS failed to read or validate the "%2" field for the computer QoS policy "%3".

MessageId	16711
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_PROCESSING_USER_POLICY_FIELD
Language	English
Message	QoS failed to read or validate the "%2" field for the user QoS policy "%3".

MessageId	16712
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_SETTING_TCP_AUTOTUNING
Language	English
Message	QoS failed to read or set inbound TCP throughput level, error code: "%2".

MessageId	16713
Severity	Error
SymbolicName	EVENT_EQOS_ERROR_SETTING_APP_MARKING
Language	English
Message	QoS failed to read or set the DSCP marking override setting, error code: "%2".

Additional resources

- For more an overview of QoS, see Quality of Service (QoS) Overview.

- For more information about QoS, see the Windows Server TechCenter web site at http://go.microsoft.com/fwlink/?LinkId=65245.
- For more information on configuring Policy-based QoS, see Configuring Policy-based Quality of Service (QoS).
- For more about Policy-based QoS, see the Cable Guy article on the Microsoft web site at http://go.microsoft.com/fwlink/?LinkId=65246.
- For more information about Group Policy, see the Windows Server TechCenter on the Microsoft web site at http://go.microsoft.com/fwlink/?LinkID=17530.
- For background on the Group Policy Management Console, see Enterprise Management with the Group Policy Management Console on the Microsoft web site at http://go.microsoft.com/fwlink/?LinkID=8630.

FAQ

What operating system does my domain controller need to be running?

Your domain controller can be running Windows 2000 Server, Windows Server 2003, Windows Server 2008, or Windows 8.

Note
GPMC must be the version available in Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, or Windows Vista.

Do I need Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to deploy QoS policies?

No, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 is not required. However, QoS policies only apply to users or computers running Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, and Windows Vista.

Do QoS policies apply to the sender or receiver of traffic?

QoS policies must be applied on the sending computer to affect its outbound traffic. In order to affect the bidirectional traffic of two computers, QoS policies need to be deployed to both computers.

What happens if conflicting policies are deployed to the same computer?

If multiple policies apply, the more specific QoS policy takes precedence. For example, a policy that states a host address (192.168.4.12) gets applied instead of a network address (192.168.0.0/16). If a computer-level and user-level policy have the same specificity, the user-level QoS policy is applied instead of the computer-level QoS policy.

Additional resources

- For more information about QoS, see the Quality of Service (QoS) Overview.
- For more information about configuring Policy-based QoS, see Configuring Policy-based Quality of Service (QoS).

Configuring Policy-based Quality of Service (QoS)

Quality of Service (QoS) for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for critical applications. By using QoS mechanisms, network administrators can use existing resources efficiently and ensure the required level of service without reactively expanding or over-provisioning their networks.

In Windows, policy-based QoS combines the functionality of standards-based QoS with the manageability of Group Policy. This combination allows for easy application of QoS policies to Group Policy Objects. Windows includes a Policy-based QoS Wizard to help you configure QoS in Group Policy.

This topic contains the following sections:

- QoS policies overview
- Create a QoS policy
  - Policy Profile - Wizard page 1
  - Application Name - Wizard page 2
  - IP Addresses - Wizard page 3
  - Protocols and Ports - Wizard page 4
- View, edit, or delete a QoS policy
- Additional resources

QoS policies overview

To create a QoS policy, edit the settings of a Group Policy Object (GPO) from within the Group Policy Management Console (GPMC). QoS policy names must be unique. Behavior depends on where the QoS policy is stored in the GPMC:

- A QoS policy in Computer Configuration\Windows Settings\Policy-based QoS applies to computers, regardless of the user that is currently logged on. You typically use computer-based QoS policies for server computers.
- A QoS policy in User Configuration\Windows Settings\Policy-based QoS applies to users after they have logged on, regardless of which computer they have logged on to.

Prioritizing traffic by using DSCP

To define the priority of outbound network traffic, use **Specify DSCP Value** to configure a QoS policy with a specific Differentiated Services Code Point (DSCP) value. As described in RFC 2474, DSCP allows values from 0 to 63 to be specified within the Type of Service (TOS) field of an IPv4 packet. Network routers use the DSCP value to classify network packets and determine the appropriate queue. A higher value indicates a higher priority for the packet. By default, the **Specify DSCP Value** check box is selected and its value is 0.

The number of queues and their prioritization behavior needs to be designed as part of your organization's QoS strategy. For example, you can create four QoS policies so that packets with specific DSCP values are kept in five queues: latency-sensitive traffic, control traffic, business critical traffic, best effort traffic, and bulk data transfer traffic.

Throttling traffic

To manage the use of network bandwidth, you can use **Specify Throttle Rate** to configure a QoS policy with a throttle rate for outbound traffic. By default, the **Specify Throttle Rate** check box is not selected. With throttling, a QoS policy will limit the outgoing network traffic to a specified rate. Both DSCP marking and throttling can be used together to manage traffic effectively.

Create a QoS policy

To create a new QoS policy with the Policy-based QoS wizard

- In the Group Policy Management Console, right-click either of the **Policy-based QoS** nodes, and then click **Create a new policy**.

Policy Profile - Wizard page 1

On the first page of the Policy-based QoS wizard, you can specify a policy name and configure how QoS will control outgoing network traffic.

To configure the Policy Profile page of the QoS-based Policy wizard

1. In **Policy name**, type a name for the QoS policy. The name must uniquely identify the policy within the GPO node of its creation.
2. Optionally, use **Specify DSCP Value** to enable DSCP marking and configure a DSCP value between 0 and 63.
3. Optionally, use **Specify Throttle Rate** to enable traffic throttling and configure the throttle rate. The throttle rate value must be greater than 1 and you can specify units of kilobytes per second (KBps) or megabytes per second (MBps).
4. Click **Next**.

QoS policies overview

Application Name - Wizard page 2

The second page of the Policy-based QoS wizard allows you to apply the policy to all applications, to a specific application as identified by its executable name, or to a path and application name. **All applications** specifies that the traffic management settings on the first page of the Policy-based QoS wizard apply to all applications. **Only applications with this executable name** specifies that the traffic management settings on the first page of the Policy-based QoS wizard are for a specific application. The executable file name and must end with the .exe file name extension.

Optionally, you can enter the application path. To specify an application path, include the path with the application name. The path may include environment variables. For example, %ProgramFiles%\My Application Path\MyApp.exe, or c:\program files\my application path\myapp.exe. Note the application path cannot include a symbolic link.

To configure the Application Name page of the Policy-based QoS wizard

1. Under **This QoS policy applies to** select either **All applications** or **Only applications with this executable name**.
2. If you select **Only applications with this executable name**, specify an executable name ending with the .exe file name extension.
3. Click **Next**.

QoS policies overview

IP Addresses - Wizard page 3

The third page of the Policy-based QoS wizard allows you to specify that the traffic management settings specified on the first page of the Policy-based QoS wizard apply to the following:

- All source IPv4 or IPv6 addresses or specific source IPv4 or IPv6 addresses
- All destination IPv4 or IPv6 addresses or specific destination IPv4 or IPv6 addresses.

If you select **Only for the following source IP address** or **Only for the following destination IP address**, you must type one of the following:

- An IPv4 address, such as 192.168.1.1
- An IPv4 address prefix using network prefix length notation, such as 192.168.1.0/24
- An IPv6 address, such as 3ffe:ffff::1
- An IPv6 address prefix, such as 3ffe:ffff::/48

If you select both **Only for the following source IP address** and **Only for the following destination IP address**, both addresses or address prefixes must be either IPv4 or IPv6-based.

To configure the IP Addresses page of the Policy-based QoS wizard

1. In **This QoS policy applies to** (source), select **Any source IP address** or **Only for the following IP source address**.
2. If you selected **Only the following IP source address:**, specify an IPv4 or IPv6 address

or prefix.

3. In **This QoS policy applies to** (destination), select **Any destination address**, or **Only for the following IP destination address.**
4. If you selected **Only for the following IP destination address:**, specify an IPv4 or IPv6 address or prefix that corresponds to the type of address or prefix specified for the source address.
5. Click **Next**.

QoS policies overview

Protocols and Ports - Wizard page 4

On the fourth page of the Policy-based QoS wizard you can specify the types of traffic and the ports that are controlled by the settings on the first page of the wizard. You can specify:

- TCP traffic, UDP traffic, or both
- All source ports, a range of source ports, or a specific source port
- All destination ports or a specific destination port

To configure the Protocols and Ports page of the Policy-based QoS wizard

1. In **Select the protocol this QoS policy applies to**, select **TCP**, **UDP**, or **TCP and UDP**.
2. In **Specify the source port number**, select **From any source port** or **From this source port number**.
3. If you selected **From this source port number** in the previous step, type a port number between 1 and 65535. Optionally, you can specify a port range, in the format of "Low:High," where Low and High represent the lower bounds and upper bounds of the port range, inclusively. Low and High each must be a number between 1 and 65535. No space is allowed between the colon (:) character and the numbers.
4. In **Specify the destination port number**, select **To any destination port** or **To this destination port number**.
5. If you selected **To this destination port number** in the previous step, type a port number between 1 and 65535.

To complete the creation of the new QoS policy, click **Finish** on the Protocols and Ports page of the Policy-based QoS wizard. When completed, the new QoS policy is listed in the details pane of the GPMC.

To apply the QoS policy settings to users or computer, link the GPO in which the QoS policies are located to an Active Directory container, such as a domain, a site, or an organizational unit (OU).

QoS policies overview

View, edit, or delete a QoS policy

The pages of the Policy-based QoS wizard described above correspond to the properties pages that are displayed when you view the properties of a policy or edit the policy.

To view the properties of a QoS policy

- Right-click the policy name in the details pane of the Group Policy Management Console, and then click **Properties**.

  The GPMC displays the properties page with the following tabs:

  - Policy Profile
  - Application Name
  - IP Addresses
  - Protocols and Ports

To edit a QoS policy

- Right-click the policy name in the details pane of the Group Policy Management Console, and then click **Edit existing policy**.

  The GPMC displays the **Edit an existing QoS policy** dialog box.

To delete a QoS policy

- Right-click the policy name in the details pane of the Group Policy Management Console, and then click **Delete policy**.

QoS policies overview

Additional resources

- For more information about QoS, see the Windows Server TechCenter Web site at http://go.microsoft.com/fwlink/?LinkId=65245.
- For more about Policy-based QoS, see the Cable Guy article on the Microsoft Web site at http://go.microsoft.com/fwlink/?LinkId=65246.

Quality Windows Audio Video Experience (qWave)

qWave is a collection of QoS-related software modules that run on qWave-enabled devices in non-enterprise environments, such as a home network or small-office network. qWave supports multiple audio/video streams (real-time flows requiring QoS) and data streams (best-effort flows, such as e-mail) simultaneously over the network, while providing a high-quality audio/video (A/V) user experience. The devices cooperate to ensure that data streams (for which the application has requested QoS) receive the level of quality that they were granted upon admission to the network. Because qWave is targeted for home and small-office A/V scenarios, it is typically disabled in other environments, such as enterprise networks.

qWave Overview

QoS refers to the mechanisms that are implemented to provide a desired level of network service to an application on IP-based networks. The network service parameters might include bandwidth, jitter, latency, and packet loss. The desired level of network service can be achieved by explicitly reserving bandwidth on the network for a service, or by providing preferential service to selected network traffic while providing best-effort service to all other traffic.

qWave is a collection of QoS-related software modules—first included on computers running Windows Vista®—in non-enterprise environments, such as a home network or small-office network. qWave supports multiple A/V streams (real-time flows requiring QoS), and data streams (best-effort flows, such as e-mail) simultaneously over the network, while providing a high-quality A/V user experience. The devices cooperate to ensure that data streams for which the application has requested QoS receive the level of quality that they were granted upon admission to the network. Because qWave is targeted for home and small-office A/V scenarios, it is typically disabled in other environments, such as enterprise networks.

Reservation techniques--such as RSVP--are not implemented for home A/V streaming. Bandwidth reservation guarantees are meaningless if the Wi-Fi station moves out of the broadcast range of the wireless access point (wireless AP) or if some obstruction or interference causes a drop in link bandwidth. Additionally, for price-sensitive home networking markets, reservation-based QoS features might make devices cost-prohibitive.

Instead, qWave provides preferential service for traffic prioritization and other throughput, range, and stability-enhancement mechanisms. These mechanisms are available in many wireless APs, wireless network adapters and other devices. These mechanisms include 802.1p support for switches, hubs, and network adapters, as well as Wireless Multimedia Extensions (WME) support in wireless APs.

qWave provides the best experience when both the source and the target devices are qWave-enabled, and when the network devices support traffic prioritization. However, qWave also provides an improved experience without priority support in either the network or legacy devices.

On Windows Server platforms, qWave provides only rate-of-flow and prioritization services.

See Also

Quality of Service (QoS) Overview

QoS Frequently Asked Questions

Frequently Asked Questions

Following are frequently asked questions – and answers to those questions – for QoS configuration.

1. **Is QoS enabled by default?**

No, QoS is not enabled by default. You must configure Bandwidth Management, Classification and Tagging, or PFC explicitly. For more information, see **Related topics**.

2. **Is there a recommended default QoS configuration?**

   No. Due to variations in server configurations and workloads, there is no recommended default QoS configuration. However, the examples in this document can serve as a basis for QoS configurations for many deployments in datacenters.

3. **Can DCB be leveraged on a NIC that is bound to a Hyper-V Virtual Switch?**

   Yes. By default, when a DCB-capable NIC is bound to a Hyper-V Virtual Switch, DCB is disabled. To take advantage of the hardware bandwidth management functionality, administrators can enable this component by using the following Windows PowerShell command in the management operating system. In addition, you can refer to the example 2 NICs with NIC Teaming for information on how to classify and tag traffic, and how to allow tagged traffic to go through the Hyper-V Virtual Switch.

   ```
   # Enable the DCB component for the DCB NIC that is bound to a
   Hyper-V switch
   Enable-NetAdapterBinding "DCB NIC bound to Hyper-V switch" -
   ComponentID ms_lldp
   ```

4. **Can multiple virtual network adapters be created for one workload and still have QoS (particularly Bandwidth Management) applied to them?**

   Yes. In fact, to leverage MPIO, you must create multiple virtual network adapters in the management operating system. In addition, you must configure Minimum Bandwidth and Maximum Bandwidth on a per-virtual network adapter basis. Also, these virtual network adapters must be placed in different IP subnets.

Remote Access (DirectAccess, Routing and Remote Access) Overview

An overview of the Remote Access technology in Windows Server 2012, its new and changed functionality, and links to additional resources.

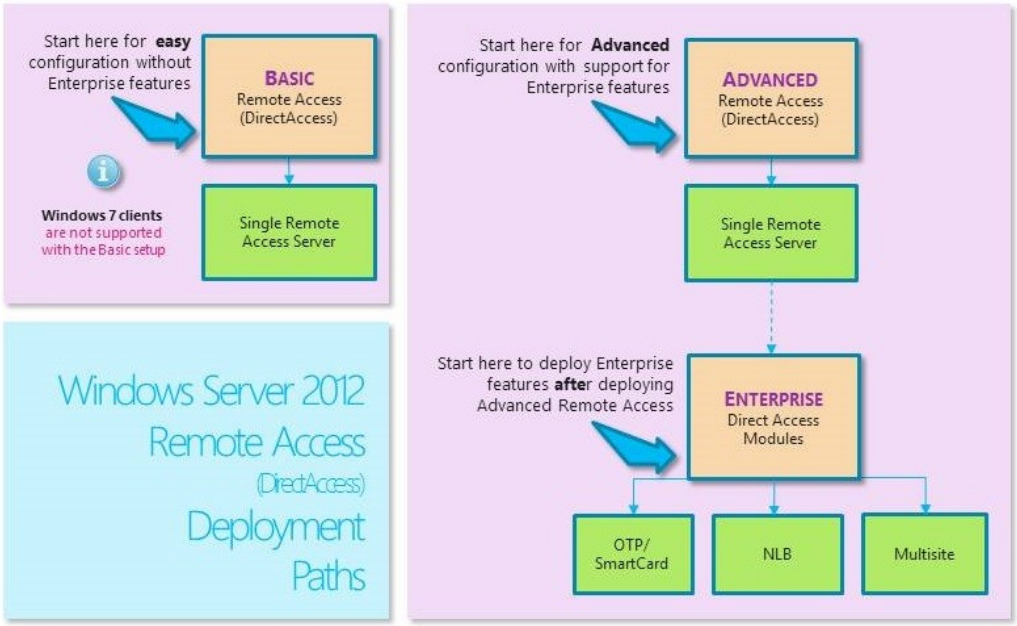

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning

- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Overview of the Remote Access Technologies

Remote Access combines two networking services into one unified server role in Windows Server 2012.

DirectAccess and RRAS Unified Server Role

Windows Server 2008 R2 introduced DirectAccess, a new remote access feature that allows connectivity to corporate network resources without the need for traditional Virtual Private Network (VPN) connections. DirectAccess provides support only for domain-joined Windows 7 Enterprise and Ultimate edition clients. The Windows Routing and Remote Access Server (RRAS) provides traditional VPN connectivity for legacy clients, non-domain joined clients, and third party VPN clients. RRAS also provides site-to-site connections between servers. RRAS in Windows Server 2008 R2 cannot coexist on the same edge server with DirectAccess, and must be deployed and managed separately from DirectAccess.

In Windows Server 2012 the DirectAccess feature and the RRAS role service were combined into a new unified server role. This new Remote Access server role allows for centralized administration, configuration, and monitoring of both DirectAccess and VPN-based remote access services. Additionally, Windows Server 2012 DirectAccess provided multiple updates and improvements to address deployment blockers and provide simplified management.

DirectAccess in Windows Server 2012 Essentials

DirectAccess is also available in Windows Server 2012 Essentials, and likewise enables seamless connectivity to your organization's network from any Internet-equipped remote location without having to establish a virtual private network (VPN) connection.

To learn more about DirectAccess in Windows Server 2012 Essentials, see **Configure DirectAccess in Windows Server 2012 Essentials**.

What works differently?

The unified server role for DirectAccess and RRAS provides a single point of configuration and management for remote access server deployment. Windows Server 2012 and Windows Server 2012 R2 includes the following improvements over Windows 7 DirectAccess and RRAS.

- DirectAccess and RRAS coexistence
- Simplified DirectAccess Deployment
- Removal of public key infrastructure (PKI) deployment as a DirectAccess prerequisite
- Built-in NAT64 and DNS64 support for accessing IPv4-only resources
- Support for DirectAccess server behind a NAT device
- Simplified network security policy
- Load balancing support
- Support for multiple domains
- Support for OTP (token based authentication)
- Automated support for force tunneling
- IP-HTTPS interoperability and performance improvements
- DirectAccess manage-out to clients support
- Multisite support
- Support for Server Core
- Windows PowerShell support
- User monitoring
- Server operations status
- Diagnostics

- Accounting and reporting
- Site-to-site IKEv2 IPsec tunnel mode VPN

DirectAccess and RRAS coexistence

Both DirectAccess and RRAS implement security features to protect the server from hostile inbound traffic. Previously these security feature settings conflicted with each other if both services attempt to run on the same server, preventing either DirectAccess or RRAS from functioning as expected.

DirectAccess relies on Internet Protocol version six (IPv6) transition technologies to establish client connections. RRAS implements Internet Key Exchange version 2 (IKEv2) Internet Protocol security (IPsec), and configures incoming and outgoing packet filters to drop all packets using transition technologies. This results in DirectAccess traffic being blocked if RRAS is installed and VPN access is deployed with IKEv2.

DirectAccess implements IPsec Denial of Service Protection (DoSP) to protect resources on the corporate network. DoSP drops all IPv4 traffic, and all IPv6 traffic that is not protected by IPsec, except ICMPv6 packets. This results in all IPv4 packets and non-IPsec-protected IPv6 packets forwarded by RRAS being blocked if DirectAccess is installed.

Windows Server 2012DirectAccess and RRAS unified server role solves these problems by modifying IKEv2 policies to allow IPv6 transition technology traffic, and by modifying IPsec DoSP to allow VPN traffic. These changes allow both DirectAccess and RRAS to coexist on the same server.

Simplified DirectAccess Deployment

Windows Server 2012 includes features to facilitate deployment, particularly for small and medium size organizations. These new features include a simplified prerequisite list, removal of the need for a full PKI deployment, integrated certificate provisioning, and removal of the requirement for two consecutive public IPv4 addresses. Each of these features is discussed in more detail in the following sections.

Administrators can now deploy DirectAccess using a new Getting Started Wizard, which presents a greatly simplified configuration experience. The Getting Started Wizard masks the complexity of DirectAccess, and allows for an automated setup in a few simple steps. The administrator no longer requires an understanding of the technical details of IPv6 transition technologies and Network Location Server (NLS) deployment.

Removal of PKI Deployment as a DirectAccess Prerequisite

One major deployment blocker for Windows 7 DirectAccess is the requirement of a PKI for server and client certificate-based authentication. DirectAccess relies on IPsec AuthIP policies for

authenticating and securing traffic from Internet-connected clients. In order to authenticate to domain resources using Kerberos, the client must first establish connectivity to DNS servers and Domain Controllers (DCs).

Windows 7 DirectAccess enables this connectivity by implementing two authentication methods in the AuthIP policies. The infrastructure IPsec tunnel is established using computer certificate as the first authentication method and user NTLM as the second method. Once this tunnel is established and a DC is available, the client can obtain a Kerberos token and establish the intranet IPsec tunnel using computer certificate and user Kerberos as the first and second authentication methods.

This implementation requires that the DirectAccess server and all clients be provisioned with computer certificates for mutual authentication. Managing an internal PKI is considered difficult by many small and medium organizations. Windows Server 2012 DirectAccess makes PKI deployment optional to simplify configuration and management.

This functionality is achieved by implementing an HTTPS based Kerberos proxy. Client authentication requests are sent to a Kerberos proxy service running on the DirectAccess server. The Kerberos proxy then sends Kerberos requests to Domain Controllers on behalf of the client.

The new Getting Started wizard provides a seamless experience for the administrator by configuring this solution automatically. In this simplified DirectAccess deployment, user level configuration options such as force tunneling, Network Access Protection (NAP) integration, and two-factor authentication are not available. This deployment requires only one IPsec tunnel to be established, and has the following requirements:

- The DirectAccess server must have TCP port 443 open on its firewall
- The DirectAccess server should have a server authentication certificate for TLS issued by a Certification Authority (CA) that is trusted by the DirectAccess clients. This can be a public CA and does not require an internal PKI deployment. If no certificate is available, the DirectAccess server setup process will configure the necessary IP-HTTPS and KDC proxy certificate automatically as a self-signed certificate.

NAT64 and DNS64 Support for Accessing IPv4-only Resources

To allow access to internal IPv4-only resources, Windows Server 2012 and Windows Server 2012 R2 DirectAccess includes native support for a protocol translation (NAT64) and name resolution (DNS64) gateway to convert the IPv6 communication from a DirectAccess client to IPv4 for the internal servers. IPv4-only intranet computers cannot initiate connections to DirectAccess clients for remote management because the translation done with NAT64 is unidirectional (for traffic initiated by the DirectAccess client).

There are three primary instances where IPv6-only DirectAccess does not allow full access to corporate intranet resources.

- Intranet servers that are not fully IPv6 capable and support only IPv4, such as Windows Server 2003 file servers
- Environments where IPv6 has been administratively disabled on the network

- Applications running on IPv6 capable servers (such as Windows Server 2008) which are not IPv6 capable themselves (such as applications that are not able to listen and respond to traffic on the IPv6 interface)

To access these resources over DirectAccess, protocol translation must be done between the DirectAccess server and the internal IPv4-only resources, with subsequent translation back to IPv6 for responses sent to DirectAccess clients. NAT64 receives IPv6 traffic from the client and converts it into IPv4 traffic to the intranet. NAT64 is used in combination with DNS64. DNS64 intercepts DNS queries from clients, and sends responses after converting IPv4 answers into associated IPv6 mappings on the NAT64.

Note

Prior to Windows Server 2012 DirectAccess, the only method available to provide protocol translation for DirectAccess is through deployment of Microsoft Forefront Unified Access Gateway DirectAccess.

The DirectAccess setup wizard will seamlessly configure protocol translation components as a background operation, without any need for administrative interaction. There are no configuration options exposed to the administrator. The setup wizard will automatically enable NAT64 and DNS64 if the internal interface of the DirectAccess server has an IPv4 address assigned. To support this functionality, the setup wizard will configure an IPv6 network prefix for NAT64. The wizard assigns the NAT64 prefix automatically, and applies it to all IPv4 ranges in the enterprise.

Support for DirectAccess Server behind a NAT Device

A Windows Server 2008 R2 DirectAccess server requires two network interfaces with two consecutive public IPv4 addresses assigned to the external interface. This is required so that it can act as a Teredo server. In order for clients behind a NAT to determine the Teredo server and the type of NAT device, the Teredo server requires two consecutive IPv4 addresses.

This requirement presents difficulty for small and medium organizations that do not have access to consecutive, public IPv4 addresses. In the future this has the potential to become a deployment blocker as the available IPv4 address space is exhausted. Windows Server 2012 and Windows Server 2012 R2 DirectAccess provides the ability to deploy the DirectAccess server behind a NAT device, with support for a single network interface or multiple interfaces, and removes the public IPv4 address prerequisite.

When the Remote Access Services setup Getting Started Wizard or Remote Access Setup Wizard is run, it will check the status of network interfaces on the server to determine if the DirectAccess server is located behind a NAT device. In this configuration, only IP over HTTPS (IP-HTTPS) will be deployed. The IP-HTTPS protocol is an IPv6 transition technology that allows for a secure IP tunnel to be established using a secure HTTP connection.

Simplified Network Security Policy

Windows Server 2008 R2 DirectAccess uses two IPsec tunnels to establish connectivity to the corporate network. The DirectAccess client requires the infrastructure tunnel to access infrastructure resources such as DNS, DC, and Management servers. These infrastructure servers are all listed as endpoints in the infrastructure tunnel IPsec policy. Then the intranet tunnel provides access to all other corporate intranet resources.

The endpoints listed in the infrastructure tunnel policy require periodic updates as the infrastructure changes, such as when DNS servers or Domain Controllers are added to or removed from the production network. Clients can lose connectivity to the domain when their IPsec policies are not updated to reflect the current infrastructure server endpoints, and this loss of connectivity will prevent them from receiving group policy updates to correct the failure.

In Windows Server 2012 and Windows Server 2012 R2, the Simplified DirectAccess model provides a way to deploy DirectAccess over a single IPsec tunnel, which eliminates this problem. However, Simplified DirectAccess does not support certain capabilities, which rely on certificate-based authentication. Examples are two-factor authentication with smart cards, multi-site, and NAP integration. Enterprises requiring a full featured DirectAccess experience will still need to deploy the two-tunnel model.

If the two-tunnel model is required for full functionality, there is additional functionality available to enable administrators to refresh the list of servers that are made accessible via the infrastructure tunnel. New domain controllers and System Center Configuration Manager servers are discovered and added to the list. Servers that no longer exist are removed from the list, and entries for servers whose IP addresses have changed are updated.

Load Balancing Support

Windows Server 2008 R2 DirectAccess does not provide a full high availability solution, and is limited to single-server deployments. To provide limited hardware redundancy, DirectAccess can be configured inside a Hyper-V Failover cluster configured for Hyper-V Live Migration. However, only one server node may be online at any time.

DirectAccess deployments have quickly grown beyond the point where a single server can provide adequate processing power. Enterprises need the flexibility to deploy additional servers quickly and transparently to meet changing load requirements. Additionally, the Network Location Server used for inside/outside detection must be highly available to prevent major outages for DirectAccess clients connected to the intranet.

Windows Server 2012 and Windows Server 2012 R2 DirectAccess address these issues through built-in support for Windows Network Load Balancing (NLB) to achieve high availability and scalability for both DirectAccess and RRAS VPN. The NLB configuration is simple to setup and automate through the new deployment wizard interface. The setup process also provides integrated support for third party external hardware-based load balancer solutions.

Important

Windows Server 2012 and Windows Server 2012 R2 DirectAccess provides a basic failover solution using Network Load Balancing for up to eight nodes. Although server load will be shared across all NLB nodes, existing connections will not automatically be transferred to other servers when one server becomes unavailable.

Support for Multiple Domains

The DirectAccess setup wizard in Windows Server 2008 R2 can be used to configure DirectAccess for a single domain only. This means that remote clients from a different domain from the DirectAccess server will not be able to use DirectAccess. In addition, if application servers are in a different domain, remote clients will not be able to access them remotely via DirectAccess.

Although administrators can manually configure multiple domain support in Windows Server 2008 R2, the deployment requires manual edit of the DirectAccess policies after setup is completed. Windows Server 2012 and Windows Server 2012 R2 DirectAccess provides integrated multiple domain support to allow remote client access to enterprise resources located in different domains.

Support for OTP (Token Based Authentication)

To increase login security, many organizations have deployed One-Time Password (OTP) two-factor authentication, and mandate its use for remote access connections. Windows Server 2008 R2 DirectAccess provided support for two-factor authentication with Smart Cards, but was not capable of integrating with OTP vendor solutions, such as RSA SecurID. This prevented DirectAccess deployment in organizations that require this level of security.

Windows Server 2012 and Windows Server 2012 R2 DirectAccess supports two-factor authentication with Smart Cards or OTP token based solutions. This feature requires a PKI deployment, so if the option is selected in the DirectAccess Setup Wizard, the **Use computer certificates** option is automatically selected and enforced.

In addition to support for standard smart card authentication, DirectAccess can use the Trusted Platform Module (TPM)-based virtual smart card capabilities available in Windows Server 2012 and Windows Server 2012 R2. The TPM of client computers can act as a virtual smart card for two-factor authentication, thus removing the overhead and costs incurred in smart card deployment.

Automated Support for Force Tunneling

By default, DirectAccess clients are able to access the Internet, the corporate intranet, and local LAN resources simultaneously. Since only connections made to the corporate intranet are sent over the DirectAccess IPsec tunnels, this is known as a split-tunnel configuration. Split tunneling provides an optimal user experience when accessing resources on the Internet, while still providing strong security for traffic intended for the intranet.

Although split tunneling is not a security risk, some organizations have a requirement to force all traffic through a corporate proxy so that it can be inspected by their IDS. With legacy VPN connections, the potential exists for users to bridge traffic between networks, such as a home network and the corporate network, effectively making the client operate as a router. For this reason, it is common practice for administrators to disable split tunneling for VPN connections, forcing all network traffic to be routed through the VPN connection. This results in decreased performance when accessing Internet resources, since all traffic must traverse the VPN tunnel and then be proxied out to the Internet. It also consumes significant additional bandwidth on the corporate network.

The perceived security risk of split tunneling is not valid in a DirectAccess scenario, since the IPsec rules that enable DirectAccess require authentication by the client endpoint. If another endpoint attempts to route through the DirectAccess client, it will not be an authenticated source, and IPsec will prevent the connection. However, since some organizations have a requirement to force all traffic through the corporate proxy server so that it can be inspected, the DirectAccess Force Tunneling option provides this ability.

The Force Tunneling option was provided in Windows Server 2008 R2 DirectAccess, but required manual steps to enable it via group policy setting. Windows Server 2012and Windows Server 2012 R2 DirectAccess integrate the Force Tunneling option with the Setup Wizard and management UI to automate the required settings. Enabling the Force Tunneling option limits the DirectAccess client to using only the IP-HTTPS protocol for connectivity, and by default uses the DirectAccess server as the NAT64/DNS64 server to translate IPv6 resources to send to the IPv4 proxy server.

IP-HTTPS Interoperability and Performance Improvements

On certain networks, Internet firewall settings may prevent successful client connections using the 6to4 or Teredo IPv6 transition technologies. IP-HTTPS is an IPv6 transition technology introduced in Windows 7 to ensure that DirectAccess clients can connect to the corporate network even when all other IPv6 transition technologies fail. IP-HTTPS assigns a unique, globally routable IPv6 address to an IPv4 host, encapsulates the IPv6 packets within IPv4 for transmission over an HTTPS tunnel, and routes IPv6 traffic between the host and other globally routable IPv6 nodes.

Windows Server 2012 and Windows Server 2012 R2 provide several improvements to the implementation of IP-HTTPS. Changes to the technology allow IP-HTTPS clients to obtain proxy configuration information, and authenticate to an HTTP proxy if authentication is required. The Windows 7 requirement to deploy client certificates to each IP-HTTPS client has been removed.

IP-HTTPS works by creating an SSL/TLS connection between the client and server, then passing IP traffic across the connection. This data is encrypted by IPsec, which means that the data is encrypted twice, first by IPsec, and again by SSL. The result is poor performance and a negative user experience compared to the other IPv6 transition technologies 6to4 and Teredo, and limits the scalability of the DirectAccess server.

Windows Server 2012 and Windows Server 2012 R2 DirectAccess include several performance improvements for IP-HTTPS to increase scalability and reduce the overhead associated with this connectivity method. These optimizations include changes to batched send behavior and receive buffers, reduced lock contention, and the option to implement SSL with NULL encryption.

IP-HTTPS runs in a system context rather than a user context. This context can cause connection issues. For example, if a DirectAccess client computer is located in the network of a partner company that uses a proxy for Internet access, and WPAD auto detection is not used, the user must manually configure proxy settings in order to access the Internet. These settings are configured in Internet Explorer on a per user basis, and cannot be retrieved in an intuitive way on behalf of IP-HTTPS. In addition, if the proxy requires authentication, the client provides credentials for Internet access, but IP-HTTPS will not provide the credentials required to authenticate to DirectAccess. In Windows Server 2012, a new feature solves these issues. Specifically, the user can configure IP-HTTPS to work when behind a proxy that is not configured using WPAD and IP-HTTPS will request and provide the proxy credentials needed to IP-HTTPS request authenticated, and relay it to the DirectAccess server.

When configuring IP-HTTPS in DirectAccess on the server, you can use a certificate issued by a certification authority (CA), or you can specify that DirectAccess should automatically generate a self-signed certificate.

DirectAccess Manage-out to Clients Support

DirectAccess clients establish connectivity to the corporate intranet whenever an Internet connection is available, even if there is no user logged in. This allows IT administrators to manage remote machines for patching and compliance enforcement even when they are not in the office. Some customers see this as the primary benefit to DirectAccess, and choose to keep their existing remote access solution in place for user connectivity, while using DirectAccess just for remote management.

Windows Server 2008 R2 DirectAccess does not provide an automated method to limit the deployment to manage-out only, and administrators must manually edit the policies created by the setup wizard. Windows Server 2012 and Windows Server 2012 R2 DirectAccess provide support for a manage-out only configuration through a deployment wizard option that limits the creation of policies to only those needed for remote management of client computers. In this deployment, user level configuration options such as force tunneling, NAP integration, and two-factor authentication are not available.

Note

Manage-out support requires ISATAP deployment or management servers with v6 addresses.

Multisite Support

DirectAccess servers can be deployed in multiple sites to increase capacity and provide more efficient access to the nearest entry point for intranet resources. This works well if clients remain

in their respective sites and do not need to travel to different sites within the enterprise. However, setting up multisite DirectAccess requires careful planning and design if clients will roam between sites, to ensure that they connect through DirectAccess servers via the most efficient route.

There are many challenges to consider in a multisite environment, such as making sure the client locates the closest IP-HTTPS server, Teredo server, DNS server, and Domain Controller. Windows Server 2012 and Windows Server 2012 R2 DirectAccess provide a solution that allows for deployment of multiple DirectAccess entry points across geographic locations, and allows clients regardless of their physical location to access resources within corpnet in an efficient manner.

Windows Server 2012 and Windows Server 2012 R2 DirectAccess servers can be configured in a multisite deployment that allows remote users in dispersed geographical locations to connect to the multisite entry point closest to them. For client computers running Windows 8, entry points can be assigned automatically, or selected manually by the client. For Windows 7 client computers, entry points can be allocated statically. Windows 7 clients can also get the benefit of automatic entry point selection if the organization deploys a Global Server Load Balancing (GSLB) solution. Traffic across the multisite deployment can be distributed and balanced with an external Traffic across the multisite deployment can be distributed and balanced with an external global load balancer.

Support for Server Core

Server Core is a minimal server installation option designed to reduce disk space, servicing, and management requirements, and decrease the operating system attack surface. The Server Core system does not support a Graphical User Interface, and administrators must use command line or remote management tools to accomplish all necessary configuration tasks.

A Server Core installation supports only a limited subset of the features available on a full installation of Windows Server, and currently does not include support for the DirectAccess feature or the Remote Access server role. A Windows Server 2012 or Windows Server 2012 R2 Server Core installation includes support for the unified server role for both DirectAccess and RRAS.

The new Remote Access server role has complete Windows PowerShell support in Windows Server 2012 and Windows Server 2012 R2 that may be used for installation, configuration and monitoring. The Remote Access server role can also be configured through remote server management.

PowerShell Support

DirectAccess in Windows Server 2008 R2 lacks a complete scripting and command line interface for configuration options. Windows Server 2012 provides full Windows PowerShell support for the setup, configuration, management, monitoring and troubleshooting of the Remote Access Server Role.

User Monitoring

Monitoring and diagnostics capabilities for both RRAS server and DirectAccess are limited in Windows Server 2008 R2. For DirectAccess, the monitoring capabilities only include basic health monitoring of DirectAccess and its components. The monitoring data and statistics available are of little significance or relevance to administrators.

User and server health monitoring introduced in Windows Server 2012 allows the administrator to understand the behavior of the DirectAccess clients and connections. The monitoring console is used to keep track of the load on the DirectAccess server, user activity, and current resource usage. The administrator uses this information to identify potentially undesirable or inappropriate usage activities. Diagnostic tracing can be enabled from the monitoring console as well.

User Monitoring

Administrators of remote access solutions require the ability to monitor not only which users are connected, but also which resources they are accessing. If users complain that a particular server or file share is inaccessible while remote, the administrator currently has no way to determine if other users are successfully accessing the resource from the remote access console. Multiple tools and applications are typically needed to troubleshoot issues such as particular users consuming excessive bandwidth.

The Dashboard is accessed from the new Remote Access server management console by selecting the Dashboard tab in the navigation pane. The dashboard displays overall operational status and remote client activity and status. The administrator can also view quick reports directly from the dashboard.

The Monitoring Dashboard shows a summary of remote client connectivity status for the following items. The information is generated from the relevant Performance Monitor counters and accounting data.

- Total number of active remote clients connected – includes both DirectAccess and VPN remote clients
- Total number of active DirectAccess clients connected – only the total number of clients connected using DirectAccess
- Total number of active VPN clients connected – only the total number of clients connected using VPN
- Total unique users connected – includes both DirectAccess and VPN users, based on the active connections
- Total number of cumulative connections – the total number of connections serviced by the Remote Access Server since the last server restart
- Maximum number of remote clients connected – the maximum concurrent remote users connected to the Remote Access Server since the last server restart
- Total data transferred – the total inbound and outbound traffic from the Remote Access Server for both DirectAccess and VPN since the last server restart
    a. Inbound traffic – Total bytes/traffic into the remote access server/gateway
    b. Outbound traffic – Total bytes/traffic out of the remote access server/gateway

In a cluster deployment, the Remote Client Activity and Status summary on the Remote Access Dashboard displays total values for all of the cluster nodes.

Administrators can see a list of all remote users currently connected, and can display a listing of all resources being accessed by clicking on a particular remote user. Administrators can display remote user statistics by selecting the Remote Client Status link in the Remote Access Management Console. The user statistics can be filtered based on criteria selections using the following fields:

Field Name	Value
Username	The user name or alias of the remote user. Wildcards may be used to select a group of users, such as contoso* or *\administrator. If Comment [A5]: Move this table to the monitoring scenario doc no domain is specified, then *\username is assumed.
Hostname	The computer account name of the remote user. An IPv4 or IPv6 address can be specified as well.
Type	Either DirectAccess or VPN. If DirectAccess is selected, then all remote users connecting using DirectAccess are listed. If VPN is selected, then all remote users connecting using VPN are listed.
ISP address	The IPv4 or IPv6 address of the remote user
IPv4 address	The inner IPv4 address of the tunnel connecting the remote user to the corporate network
IPv6 address	The inner IPv6 address of the tunnel connecting the remote user to the corporate network
Protocol/Tunnel	The transitioning technology used by the remote client – Teredo, 6to4 or IP-HTTPS in case of DirectAccess users, and PPTP, L2TP, SSTP or IKEv2 in case of VPN users
Resource Accessed	All users accessing a particular corporate resource or an endpoint. The value corresponding to this field is the hostname/IP address of the server/endpoint

Server	The Remote Access server to which clients are connected. This is relevant only for cluster and multisite deployments.

Server Operational Status

This feature allows administrators to manage and monitor the status of remote access deployments from a centralized monitoring console. The feature shows alerts to administrators whenever an issue requiring attention is detected. The interface displays detailed diagnostic information with steps to provide resolution.

The **Dashboard** node of the console tree shows the status of the Remote Access Server, including the status of remote access infrastructure and related components, as well as whether the configuration is properly distributed across entry points.

The Server **Operations Status** node of the console tree shows the status of the Remote Access Server, including the status of remote access infrastructure and related components. By clicking on a particular component, administrators can see the state, change history, and monitoring details for that component.

If remote access servers are deployed in a cluster or multisite deployment, all servers in the cluster or multisite deployment are evaluated asynchronously, and are listed with their overall status. Administrators can scroll through the list of servers and expand or collapse the view to display DirectAccess and VPN server components.

The Remote Access components with status monitors displayed in the Server Operations Status pane are listed below.

- 6to4
- DNS
- DNS64
- Domain controller
- IP-HTTPS
- IPsec
- ISATAP
- Kerberos
- Management Servers
- NAT64
- Network Adapters
- Network Location Server
- Network Security (IPsec DoSP)
- Services
- Teredo
- Load Balancing

- VPN addressing
- VPN connectivity

Diagnostics

Troubleshooting remote access connectivity failures for both RRAS and DirectAccess can be very complex due to the limited logging capabilities currently provided. Administrators typically require Network Monitor captures and RRAS tracing for troubleshooting, since Event Viewer logs are not very useful or prescriptive.

Windows Server 2012 and Windows Server 2012 R2 provide the following diagnostic feature improvements for remote access troubleshooting.

- Detailed event logging for DirectAccess

  Administrators can use improved event logging to identify problems and perform capacity and performance analysis. The event logs are standardized to ensure a consistent experience with other networking components.

- Tracing and Packet Capture

  Integrated tracing makes it easy for administrators to gather trace logs and network packet captures with a single click. Both tracing with packet capture and log correlation are done as part of a single process when the administrator clicks the **Start tracing** task in the Tasks pane.

- Log Correlation

  This feature provides automated collection and correlation of logs for different DirectAccess components with a single click, leveraging the Unified Tracing feature of Windows. The events gathered from different components are consolidated into a single file through correlation of Activity IDs. Activity IDs are GUIDs that identify a particular task or action. When a component logs an event, it associates an Activity ID with the event. The component then passes either this ID or a transfer event mapped to the ID to the component that performs the next task in the scenario, which associates its activity ID to log events. When analyzing the resulting trace file, the relationship between the various components relevant to a scenario can be reconstructed.

- Enabling/Viewing Logs

  Tracing can be enabled from the Tasks pane of the Monitoring Dashboard, or from the command line, which also controls logging levels, keywords and filters. The resultant Unified Tracing ETL files generated can be read and viewed using Network Monitor.

Accounting and Reporting

A Windows Server 2012and Windows Server 2012 R2 Remote Access servers can leverage an existing RADIUS server deployment or Windows Internal Database (WID) for accounting purposes. Information and historical data for load and server status are available through system Performance Monitor counters, and are stored in the WID accounting store. Whenever any connection is received or disconnected on the remote access server, all the remote user statistics

(including the endpoints accessed) is saved in the accounting store as one session. This allows session details to later be accessed for reporting and auditing purposes.

The accounting and reporting functionality provided in the Remote Access Server Role includes the ability to measure specific metrics. Available metrics include the number of users connected to a particular DirectAccess server, and total bytes transferred. Administrators can create custom reports to identify traffic and usage patterns, including a history of these patterns.

DirectAccess and RRAS reporting capabilities enable administrators to generate rich usage reports on various statistics such as remote user statistics, server availability and load. The inbox accounting store is leveraged to generate the usage report. Inbox accounting to a local WID database must be enabled in order to generate usage reports. NPS/RADIUS accounting cannot be used for generating reports.

The usage report provides a display of usage history including which users established remote connections, what resources they accessed, the total number of unique users, and maximum server load generated. The administrator can select a specific timeframe from which to generate the data.

Site-to-site IKEv2 IPsec tunnel mode VPN

Cross Premise Connectivity is a Windows Server 2012 and Windows Server 2012 R2 feature that provides the network connectivity to enable service hosting providers to migrate their applications and infrastructure to the cloud. This feature includes a site-to site Internet Key Exchange version 2 (IKEv2) tunnel mode VPN connectivity solution and management interface. Windows Server 2008 R2 introduced IKEv2 support in RRAS for VPN connections. An IKEv2 VPN provides resilience to the VPN client when the client moves from one network to another or when it switches from a wireless to a wired connection. The use of IKEv2 and IPsec allows support for strong authentication and encryption methods. Windows Server 2012 and Windows Server 2012 R2 RRAS provides added feature enhancements to enable IKEv2 for site-to-site VPN connections.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID

Content type	References
Deployment	Remote Access (DirectAccess, Routing and Remote Access) Overview
Tools and settings	Remote Access PowerShell cmdlets
Community resources	RRAS Product Team blog | Remote Access TechNet Forum DirectAccess Wiki entries
Related technologies	How IPv6 works

What's New in Remote Access in Windows Server 2012 R2

New Features in Windows Server 2012 R2

There are a number of new Remote Access server and client features in Windows Server® 2012 R2 and Windows® 8.1. The new server features include:

- Multi-tenant Site-to-site VPN Gateway
- Multi-tenant Remote Access VPN Gateway
- Border Gateway Protocol (BGP)
- Web Application Proxy

The new Windows 8.1 Remote Access client features include:

- Auto-triggered VPN
- Enhanced VPN Client PowerShell configuration
- Enhanced VPN IPsec
- Create and Edit VPN profiles in PC settings

Multi-tenant Site-to-site VPN Gateway

With Windows Server 2012 R2, hosters can deploy multi-tenant Site-to-site (S2S) Gateways (GW) to provide cross-premises connectivity from networks at the tenant sites to virtual networks dedicated per tenant in the hoster's network. The virtual network of the tenant could be built on top of Hyper-V Network Virtualization (HNV) or VLAN at the hoster. A single GW instance is capable of serving multiple tenants with overlapping IP address spaces, maximizing efficiency for the hoster as compared to deploying separate GW instance per tenant. The Routing and Remote Access (RRAS) GW is a software only solution that can be deployed in multiple instances of multi-tenant RRAS servers to balance the load.

Multi-tenant Remote Access VPN Gateway

With Windows Server 2012, hosters will able to allow transparent VPN access to virtual machines replicated in the cloud even after a failure when the entire site of the tenant goes down. Windows Server 2012 reduces the CAPEX and OPEX for hosters with a single RRAS Gateway that can service multiple tenants with overlapping IP address spaces. The RRAS GW is a software only solution that can be deployed in multiple instances of multi-tenant RRAS servers to balance the load.

Border Gateway Protocol (BGP)

Windows Server 2012 Border Gateway Protocol (BGP) enables dynamic distribution and learning of routes by site-to-site (S2S) interfaces of RRAS. This feature enables hosters (primarily Infrastructure as a service (IaaS) providers) to deploy BGP on multi-tenant RRAS S2S gateway, so that the gateway can learn what packets need to be routed to the Internet, tenant premises, and tenant virtual network at the hoster, and route them accordingly. RRAS gateway with BGP enabled can also be deployed by enterprises at their premises edge to distribute internal routes to other edge gateways (of the same enterprise in physical or virtual networks) over secure tunnels.

Web Application Proxy

Web Application Proxy is a new Remote Access role service in Windows Server 2012 R2. Web Application Proxy provides reverse proxy functionality for web applications inside your corporate network to allow users on any device to access them from outside the corporate network. Organizations can provide selective or conditional access to these web applications, according to organizational requirements. Web Application Proxy also provides built-in AD FS proxy capabilities. For more information, see Web Application Proxy Overview.

Auto-triggered VPN

Auto-triggered VPN, or Application-triggered VPN, in Windows 8.1 allows pre-defined applications to automatically connect to corporate networks by opening a VPN connection when the application is started. You can define the applications you want to make available for auto-triggering and restrict remote access based on the user identity and the computer identity from which the user is accessing the resource. This provides a secure and reliable way of accessing corporate resources from various devices. For more information, see Windows Server 2012 R2 Test Lab Guide: Demonstrate VPN Auto trigger.

SSL VPN plugin from 3rd Party VPN vendors

Windows 8.1 for x86, amd64, and ARM (Windows 8.1 RT) supports the inbox SSL VPN plugin from the following third party VPN vendors: **Checkpoint Software Technologies Ltd**, **DELL SonicWall Inc**., **F5 Networks, Inc**., and **Juniper Networks, Inc.** The 3rd party VPN plugin also supports a similar Windows PowerShell cmdlet to configure auto-triggered VPN. The VPN

connection profiles can be configured for auto-triggered based on a DNS name or for an application by Windows PowerShell or Microsoft Mobile device Management.

Support for 3rd party VPN plugins is provided by the respective 3rd party VPN partners.

The following table provides links to documents from the respective 3rd party VPN partners

VPN Plugin Provider	Support Document
SonicWall	SonicWall Mobile Connect for Windows 8.1 User Guide
Juniper	Windows In-Box Junos Pulse Client- Pdf - Link
F5	BIG-IP APM Client Compatibility Matrix|Configuration Notes: Inbox F5 VPN Client for Microsoft Windows 8.1
Checkpoint	Check Point Mobile VPN for Windows 8.1 Release Notes Check Point Mobile VPN for Windows 8.1 Administration Guide

Advanced VPN Client PowerShell configuration

Advanced VPN configuration support in Windows Windows 8.1 enables use of a single set of VPN configuration PowerShell cmdlets to configure the VPN connections instead of using multiple scripts. This feature also helps with better Microsoft Mobile Device Management allowing all the required settings from the server to be pushed out to users without executing any additional scripts.

Enhanced VPN IPsec

Enhanced IPsec in in Windows 8.1 allows Suite-B cryptography standards, or other custom cryptography configurations for IPSec based tunnels. VPN clients can be configured to use these standards and configurations with easy-to-use PowerShell cmdlets. And it also allows you to decide which client certificate is selected for any particular VPN connection.

Create and Edit VPN profiles in PC settings

In Windows 8, creating or editing VPN profiles had to be done in the **Desktop** view. Now, in Windows 8.1, you can create and edit VPN connection profiles in **PC Settings** rather than going to the **Desktop** and navigating through a number of menus. This new VPN profile editing

experience is optimized for touch screens and simplified. Advanced properties can still be accessed through the desktop when they are needed.

DirectAccess Offline Domain Join

This guide explains the steps to perform an offline domain join with DirectAccess. During an offline domain join, a computer is configured to join a domain without physical or VPN connection.

This guide includes the following sections:

- **Offline domain join overview**
- **Requirements for offline domain join**
- **Offline domain join process**
- **Steps for performing an offline domain join**

Offline domain join overview

Introduced in Windows Server 2008 R2, domain controllers include a feature called Offline Domain Join. A command line utility named Djoin.exe lets you join a computer to a domain without physically contacting a domain controller while completing the domain join operation. The general steps for using Djoin.exe are:

1. Run **djoin /provision** to create the computer account metadata. The output of this command is a .txt file that includes a base-64 encoded blob.
2. Run **djoin /requestODJ** to insert the computer account metadata from the .txt file into the Windows directory of the destination computer.
3. Reboot the destination computer, and the computer will be joined to the domain.

Offline domain join with DirectAccess policies scenario overview

DirectAccess offline domain join is a process that computers running Windows Server 2012 and Windows 8can use to join a domain without being physically joined to the corporate network, or connected through VPN. This makes it possible to join computers to a domain from locations where there is no connectivity to a corporate network. Offline domain join for DirectAccess provides DirectAccess policies to clients to allow remote provisioning.

A domain join creates a computer account and establishes a trust relationship between a computer running a Windows operating system and an Active Directory® domain.

Prepare for offline domain join

1. Create the machine account.
2. Inventory the membership of all security groups to which the machine account belongs.

3. Gather the required computer certificates, group policies, and group policy objects to be applied to the new client(s).

. The following sections explain operating system requirements and credential requirements for performing a DirectAccess offline domain join using Djoin.exe.

Operating system requirements

You can run Djoin.exe for DIrectAccess only on computers that run Windows Server 2012 or Windows 8. The computer on which you run Djoin.exe to provision computer account data into AD DS must be running Windows Server 2012 or Windows 8. The computer that you want to join to the domain must also be running Windows Server 2012 or Windows 8.

Credential requirements

To perform an offline domain join, you must have the rights that are necessary to join workstations to the domain. Members of the Domain Admins group have these rights by default. If you are not a member of the Domain Admins group, a member of the Domain Admins group must complete one of the following actions to enable you to join workstations to the domain:

- Use Group Policy to grant the required user rights. This method allows you to create computers in the default Computers container and in any organizational unit (OU) that is created later (if no Deny access control entries (ACEs) are added).
- Edit the access control list (ACL) of the default Computers container for the domain to delegate the correct permissions to you.
- Create an OU and edit the ACL on that OU to grant you the **Create child – Allow** permission. Pass the **/machineOU** parameter to the **djoin /provision** command.

The following procedures show how to grant the user rights with Group Policy and how to delegate the correct permissions.

Granting user rights to join workstations to the domain

You can use the Group Policy Management Console (GPMC) to modify the domain policy or create a new policy that has settings that grant the user rights to add workstations to a domain.

Membership in **Domain Admins**, or equivalent, is the minimum required to grant user rights. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups (http://go.microsoft.com/fwlink/?LinkId=83477).

To grant rights to join workstations to a domain

1. Click **Start**, click **Administrative Tools**, and then click **Group Policy Management**.
2. Double-click the name of the forest, double-click **Domains**, double-click the name of the domain in which you want to join a computer, right-click **Default Domain Policy**, and then click **Edit**.
3. In the console tree, double-click **Computer Configuration**, double-click **Policies**, double-click **Windows Settings**, double-click **Security Settings**, double-click **Local**

Policies, and then double-click **User Rights Assignment**.

4. In the details pane, double-click **Add workstations to domain**.
5. Select the **Define these policy settings** check box, and then click **Add User or Group**.
6. Type the name of the account that you want to grant the user rights to, and then click **OK** twice.

Offline domain join process

Run Djoin.exe at an elevated command prompt to provision the computer account metadata. When you run the provisioning command, the computer account metadata is created in a binary file that you specify as part of the command.

For more information about the NetProvisionComputerAccount function that is used to provision the computer account during an offline domain join, see NetProvisionComputerAccount Function (http://go.microsoft.com/fwlink/?LinkId=162426). For more information about the NetRequestOfflineDomainJoin function that runs locally on the destination computer, see NetRequestOfflineDomainJoin Function (http://go.microsoft.com/fwlink/?LinkId=162427).

Steps for performing a DirectAccess offline domain join

The offline domain join process includes the following steps:

1. Create a new computer account for each of the remote clients and generate a provisioning package using the Djoin.exe command from an already domain joined computer in the corporate network.
2. Add the client computer to the DirectAccessClients security group
3. Transfer the provisioning package securely to the remote computers(s) that will be joining the domain.
4. Apply the provisioning package and join the client to the domain.
5. Reboot the client to complete the domain join and establish connectivity.

There are two options to consider when creating the provisioning packet for the client. If you used the Getting Started Wizard to install DirectAccess without PKI, then you should use option 1 below. If you used the Advanced Setup Wizard to install DirectAccess with PKI, then you should use option 2 below.

Complete the following steps to perform the offline domain join:

Option1: Create a provisioning package for the client without PKI

1. At a command prompt of your Remote Access server, type the following command to provision the computer account:

```
Djoin /provision /domain <your domain name> /machine <remote
machine name> /policynames DA Client GPO name /rootcacerts
```

```
/savefile c:\files\provision.txt /reuse
```

Option2: Create a provisioning package for the client with PKI

1. At a command prompt of your Remote Access server, type the following command to provision the computer account:

```
Djoin /provision /machine <remote machine name> /domain <Your
Domain name> /policynames <DA Client GPO name> /certtemplate
<Name of client computer cert template> /savefile
c:\files\provision.txt /reuse
```

Add the client computer to the DirectAccessClients security group

1. On your Domain Controller, from **Start** screen, type **Active** and select **Active Directory Users and Computers** from **Apps** screen.
2. Expand the tree under your domain, and select the **Users** container.
3. In the details pane, right-click **DirectAccessClients**, and click **Properties**.
4. On the **Members** tab, click **Add**.
5. Click **Object Types…**, select **Computers**, and then click **OK**.
6. Type the client name to add, and then click **OK**.
7. Click **OK** to close the **DirectAccessClients** Properties dialog, and then close **Active Directory Users and Computers**.

Copy and then apply the provisioning package to the client computer

1. Copy the provisioning package from c:\files\provision.txt on the Remote Access Server, where it was saved, to c:\provision\provision.txt on the client computer.
2. On the client computer, open an elevated command prompt, and then type the following command to request the domain join:

```
Djoin /requestodj /loadfile C:\provision\provision.txt
/windowspath %windir% /localos
```

3. Reboot the client computer. The computer will be joined to the domain. Following the reboot, the client will be joined to the domain and have connectivity to the corporate network with DirectAccess.

See Also

NetProvisionComputerAccount Function

NetRequestOfflineDomainJoin Function

Remote Access (DirectAccess) Known Issues

Recommended hotfixes and updates for Windows Server 2012 DirectAccess

The following link lists Microsoft Technical Support documents for DirectAccess that you should review and apply before you start your deployment to avoid an unusable configuration.

Recommended hotfixes and updates for Windows Server 2012 DirectAccess

Remote Access (DirectAccess) Prerequisites

The following table lists the prerequisites necessary for using the configuration wizards to deploy DirectAccess.

Scenario	Prerequisites
Deploy a Single Remote Access Server using the Getting Started Wizard	• Windows Firewall must be enabled on all profiles • Only supported for clients running Windows 8.1 Enterprise and Windows 8 Enterprise. • ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6. • A public key infrastructure is not required. • Not supported for deploying two-factor authentication. Domain credentials are required for authentication. • Automatically deploys DirectAccess to all mobile computers in the current domain. • Traffic to the Internet does not go through DirectAccess. Force tunnel configuration is not supported. • DirectAccess server is the network location server. • Network Access Protection (NAP) is not supported. • Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is

	not supported. • For a multisite configuration, now or in the future, first follow the guidance in Deploy a Single Remote Access Server with Advanced Settings.

Deploy a Single Remote Access Server with Advanced Settings	• A public key infrastructure must be deployed.
	For more information, see Test Lab Guide Mini-Module: Basic PKI for Windows Server 2012.
	• Windows Firewall must be enabled on all profiles.
	• ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
	• Computers that are running the following operating systems are supported as DirectAccess clients.
	  • Windows Server® 2012 R2
	  • Windows 8.1 Enterprise
	  • Windows Server® 2012
	  • Windows 8 Enterprise
	  • Windows Server® 2008 R2
	  • Windows 7 Ultimate
	  • Windows 7 Enterprise
	• Force tunnel configuration is not supported with KerbProxy authentication.
	• Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is not supported.
	• Separating NAT64/DNS64 and IPHTTPS server roles on another server is not supported.

Deploy Remote Access in a Cluster	• Default load balancing is through the Network Load Balancing (NLB) feature in Windows Server. • External load balancers are supported. • Unicast mode is the default and recommended mode for NLB.

	• Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is not supported. • When NLB or an external load balancer is used, the IPHTTPS prefix must remain /59. • Load balanced nodes must be in the same IPv4 subnet. • In external load balancer deployments, if remote management is needed, DirectAccess clients cannot use Teredo. Only IPHTTPS can be used for end-to-end communication. • All known hotfixes for Network Load Balancing external load balancing must be installed. • ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
Deploy Multiple Remote Access Servers in a Multisite Deployment	• For a multisite configuration, now or in the future, first follow the guidance in Deploy a Single Remote Access Server with Advanced Settings. • Clients running Windows 7 always connect to a specific site. They cannot connect to the closest site based on the location of the client (unlike clients running Windows 8.1 and Windows 8). • Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is not supported. • A public key infrastructure must be deployed. For more information, see Test Lab Guide Mini-Module: Basic PKI for Windows Server 2012. • The corporate network must be using IPv6. If you are using ISATAP, you should remove it and use native IPv6.
Deploy Remote Access with OTP Authentication	• Before you deploy a one-time password authentication, follow the guidance in Deploy a single Remote Access server with

	advanced settings . • Clients running Windows 7 Enterprise and Windows 7 Ultimate need to use DCA 2.0 to support one-time password authentication. • One-time password authentication does not support a PIN change. • A public key infrastructure must be deployed. For more information, see Test Lab Guide Mini-Module: Basic PKI for Windows Server 2012. • Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is not supported.
Deploy Remote Access in a Multi-Forest Environment	• Two-way trust is required.

Manage DirectAccess Clients Remotely	• Windows Firewall must be enabled on all profiles.
	• ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
	• Computers that are running the following operating systems are supported as DirectAccess clients.
	  • Windows Server® 2012 R2
	  • Windows 8.1 Enterprise
	  • Windows Server® 2012
	  • Windows 8 Enterprise
	  • Windows Server® 2008 R2
	  • Windows 7 Ultimate
	  • Windows 7 Enterprise
	• Changing policies by using a feature other than the DirectAccess management console or Windows PowerShell cmdlets is not supported.

Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012	• ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6. • If NAP is used in Forefront Unified Access

	Gateway, NAP requires a separate Network Policy Server. NAP was deprecated in Windows Server 2012 R2. This means that NAP may not be supported in future versions of Windows. New deployments with NAP are not recommended.

Remote Access (DirectAccess) Unsupported Configurations

Review the following list of unsupported DirectAccess configurations before you start your deployment to avoid having to start your deployment again.

Network Access Protection for DirectAccess clients

Network Access Protection (NAP) is used to determine whether remote client computers meet IT policies before they are granted access to the corporate network. NAP was deprecated in Windows Server 2012 R2. This means that NAP may not be supported in future versions of Windows. For this reason, starting a new deployment of DirectAccess with NAP is not recommended. A different method of end point control for the security of DirectAccess clients is recommended.

Multisite support for Windows 7 clients

When DirectAccess is configured in a multisite deployment, Windows 8 and Windows 8.1 clients have the capability of connecting to the nearest site. Windows 7 client computers do not have the same capability. Site selection for Windows 7 clients is set to a particular site at the time of policy configuration, and these clients will always connect to that designated site, regardless of their location.

User-based access control

DirectAccess policies are computer based, not user based. Specifying DirectAccess user policies to control access to the corporate network is not supported.

Customizing DirectAccess policy

DirectAccess can be configured by using the DirectAccess Setup Wizard, the Remote Access Management console, or the Remote Access Windows PowerShell cmdlets. Using any means other than the DirectAccess Setup Wizard to configure DirectAccess, such as modifying DirectAccess Group Policy Objects directly or manually modifying the default policy settings on the server or client, is not supported. These modifications may result in an unusable configuration.

KerbProxy authentication

When a DirectAccess server is configured with the Getting Started Wizard, the server uses KerbProxy authentication for computer and user authentication. As such, the Getting Started Wizard should be used only for single-site deployments that run only Windows 8 or Windows 8.1 clients. Consider the following issues when you are planning to configure DirectAccess:

The following features should not be used with KerbProxy authentication:

- Load balancing (by using an external load balancer or Windows Load Balancer)
- Two-factor authentication where smart cards or a one-time password are needed in the deployment

The following deployment plans are not supported if you configure KerbProxy authentication:

- Multisite
- Force tunneling
- DirectAccess support for Windows 7 clients

Note

For the previous deployments, you should use the Advanced Configuration Wizard, which uses a two-tunnel configuration with a certificate-based computer and user authentication. For more information, see Deploy a Single Remote Access Server with Advanced Settings.

Using ISATAP

ISATAP is a transition technology that provides IPv6 connectivity in IPv4-only corporate networks. It is limited to small- and medium-sized organizations with a single DirectAccess server deployment, and it allows remote management of DirectAccess clients. If ISATAP is deployedin a multisite, load balancing, or multidomain environment, you must remove it or move it to a native IPv6 deployment before you configure DirectAccess.

Remote Access (DirectAccess) Troubleshooting

Follow these steps to troubleshoot Remote Access (DirectAccess) issues.

Issue	Resolution

Remote Access management console is unable to show the DirectAccess configuration	**To restore missing configuration information**
	• If you are troubleshooting a multisite deployment, ensure that the domain controller closest to the entry point is available.
	• Use the **Get-DAEntrypointDC** cmdlet to retrieve the name of the domain controller closest to the entry point. If the domain controller is not running, use the **Set-DAEntryPointDC** cmdlet to point to another domain controller.
	• Run **gpresult** from an elevated command prompt on the server to ensure the server is getting the DirectAccess Group Policy Objects.
	• Enable user interface (UI) logging.
	• Use the following command to start Windows PowerShell logging:
	`logman create trace ETWTrace -ow -o c:\ETWTrace.etl -p {AAD4C46D-56DE-4F98-BDA2-B5EAEBDD2B04} 0xffffffffffffffff 0xff -nb 16 16 -bs 1024 -mode 0x2 -max 2048 -ets`
	`logman update trace ETWTrace -p {62DFF3DA-7513-4FCA-BC73-25B111FBB1DB} 0xffffffffffffffff 0xff -ets`
	`<repro>`
	• Close and reopen the user interface.
	• Disable Windows Powershell logging. Collect the Event Trace Log files. Also, collect all the logs from the **%windir%/tracing** folder.

Applying the DirectAccess	**To refresh the DirectAccess configuration**

configuration fails	1.
	• If you are troubleshooting a multisite deployment, ensure that the domain controller closest to the entry point is available.
	• Use the **Get-DAEntrypointDC** cmdlet to retrieve the name of the domain controller closest to the entry point. If the domain controller is not running, use the **Set-DAEntryPointDC** cmdlet to point to another domain controller.
	• Use the following command to start Windows Powershell logging:
	`logman create trace ETWTrace -ow -o c:\ETWTrace.etl -p {AAD4C46D-56DE-4F98-BDA2-B5EAEBDD2B04} 0xffffffffffffffff 0xff -nb 16 16 -bs 1024 -mode 0x2 -max 2048 -ets`
	`logman update trace ETWTrace -p {62DFF3DA-7513-4FCA-BC73-25B111FBB1DB} 0xffffffffffffffff 0xff -ets`
	`<repro>`
	• Click **Apply**.
	• After the failure occurs, disable Windows Powershell logging, and collect the Event Trace Log.

DirectAccess is configured, but clients are not able to connect to internal resources	**To troubleshoot client connection issues** • Click the **Operations Status** tab in the Remote Access Management console, and ensure that all the components show a green icon. If not, check the error details and follow the resolution steps. • Run the Remote Access Server Best Practices Analyzer (BPA). If there are any warnings or errors, follow the resolution steps to resolve the issue.
Encountering issues related to a multisite configuration (for example, enabling a multisite, adding entry points, or setting the domain controller for an entry point)	Follow the steps in Troubleshoot a Multisite Deployment.
Configuration status	Follow the steps in Monitor the configuration distribution status of the

tile on the dashboard shows a warning or error	Remote Access server.
Encountering issues related to configuring load balancing (for example, the configuration fails when you enable load balancing, or there are issues when you add or remove servers from a cluster)	If you were enabling load balancing or adding a node, and the configuration refreshed when you clicked **Apply**, but the cluster didn't form correctly on the server, run the following command: **cmd.exe /c "reg add HKLM\SYSTEM\CurrentControlSet\Services\RaMgmtSvc\Parameters /f /v DebugFlag /t REG_DWORD /d ""0xffffffff"" "** to collect the user interface logs on the new server.
Operations status shows an error or warning after following steps to correct the situation	If the operations status is showing incorrect information (such as errors—even after you fix them): • Enable the registry key **cmd.exe /c "reg add HKLM\SYSTEM\CurrentControlSet\Services\RaMgmtSvc\Parameters /f /v EnableTracing /t REG_DWORD /d ""5"" "**. • Refresh the operations status and collect the logs from **%windir%/tracing**.

DirectAccess Capacity Planning

This document is a report on Windows Server 2012 DirectAccess server performance. Testing was performed to determine throughput capacity using high-end computer hardware and low-end computer hardware. High and low-end CPU performance was dependent on the network traffic throughput and the types of clients used. A typical DirectAccess deployment (and the basis for these tests) consists of 1/3 (30%) IPHTTPS clients, and 2/3 (70%) Teredo clients. Teredo clients outperform IPHTTPS clients in part because Windows Server 2012 utilizes Receive Side Scaling (RSS) which allows use of all CPU cores. In these tests, since RSS is enabled, Hyper threading is disabled. In addition, TCP/IP in Windows Server 2012 supports UDP traffic allowing Teredo clients to load balance across CPUs.

Data was collected from both a low-end (4 core, 4 Gig) server, and from hardware which is expected to be a more typical in a high-end (8 core, 8 Gig) server. Below is a screen shot of the new Windows 8 task manager on low-end hardware with 750 clients (562 Teredo, 188 IPHTTPS) running ~77 Mbits/sec. This is to simulate users who do not present smart card credentials.

These test results indicate that Teredo performs better than IPHTTPS in Windows 8, but that both Teredo and IPHTTPS bandwidth usage has improved when compared to Windows 7.

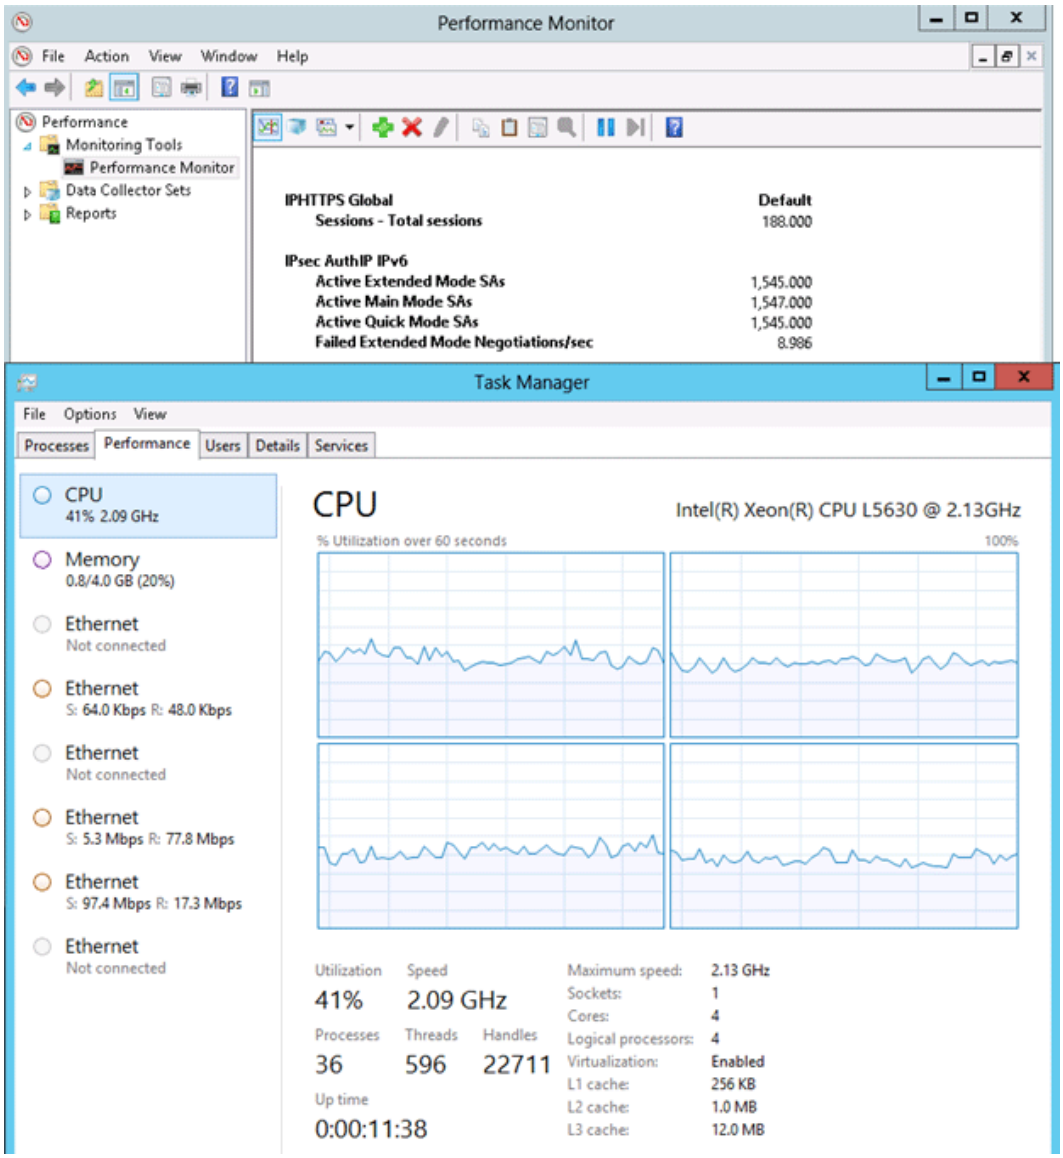

High-end hardware test environment

The following chart shows the results of the high-end hardware performance test environment. All test results and analysis are detailed in this document.

Configuration - Hardware	Low-end Hardware (4GB ram, 4 core)	High-end Hardware (8 GB, 8 core)
Double Tunnel • PKI • Including DNS64/NAT64	750 concurrent connections at 50% CPU, 50 % Memory with Corpnet NIC throughput 75	1500 concurrent connections at 50% CPU, 50 % Memory with Corpnet NIC throughput

	Mbps. Stretch target is 1000 users @ 50% CPU.	150 Mbps.

Test Environment

Perf Bench Topology

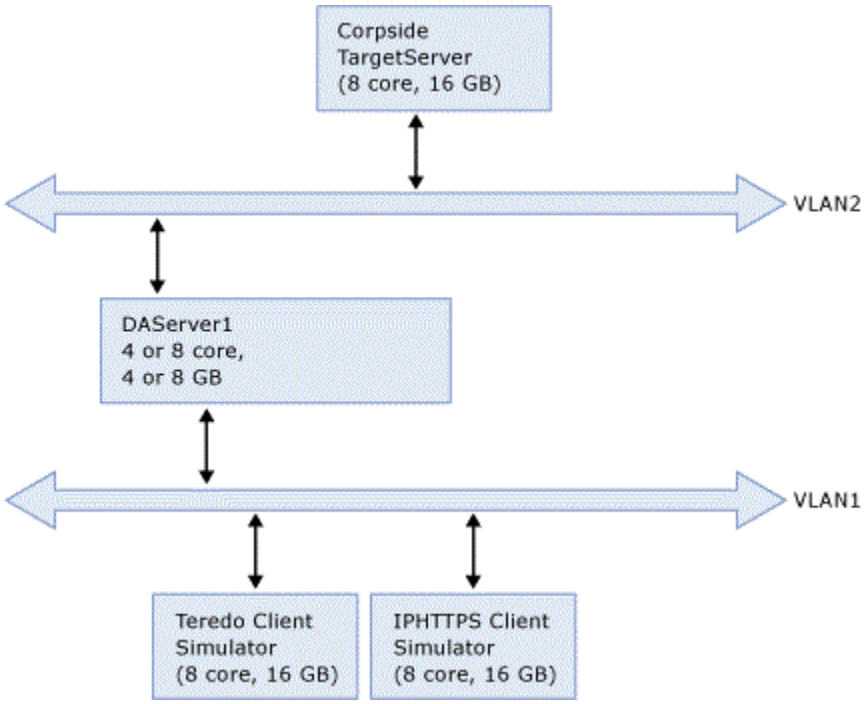

The performance test environment is a 5 machine bench. For the low-end test, one 4-core 4 Gig DirectAccess server was used and for the high-end hardware test, one 8-core, 16 Gig DirectAccess server was used. For low-end and high-end test environments the following was used: one Back end Server (the sender), and two client computers (the receivers). Receivers are split among the two client computers. Otherwise, the receivers would be CPU bound and limit the number of clients and bandwidth. On the receiving side a simulator to simulate hundreds of clients (either HTTPS or Teredo clients are simulated). IPsec, DOSp are both configured. RSS is enabled on the DirectAccess server. RSS queue size is set to 8. Without configuring RSS, a single processor will get pegged at a high utilization while the other cores are underutilized. Also of note is that the DirectAccess server is a 4 core machine with hyper threading turned off. Hyper threading is off because RSS only works on physical cores and use of hyper threading produces skewed results. (This means that not all the cores will be uniformly loaded).

Testing results for low-end hardware:

Testing was performed both with 1000 & with 750 clients. In all cases traffic split was 70% Teredo and 30% IPHTTPS. All tests involved TCP traffic over Nat64 using 2 IPsec tunnels per client. In all tests, memory utilization was light and CPU utilization was acceptable.

Individual Test Results:

The following sections describe individual tests. Each section title highlights the key elements of each test followed by a summary description of the results and then a chart showing the detailed results data.

Low-end Perf: 750 clients, 70/30 split, 84.17 Mbits/sec throughput:

The following three tests represent low-end hardware. In the below test runs, there were 750 clients with a throughput of 84.17 Mbits/sec and a traffic split of 562 Teredo and 188 IPHTTPS. Teredo MTU was set to 1472, and Teredo Shunt was enabled. CPU utilization averaged 46.42% across the three tests, and average memory utilization, expressed as a percentage of committed bytes of the total available memory of 4Gb, was 25.95%.

Scenario	CPUAvg (from counter)	Mbit/s (Corp Side)	Mbit/s (internet Side)	Active QMSA	Active MMSA	Mem Utilization (4 Gig system)
Low-end HW. 562 Teredo clients. 188 IPHTTPS clients.	47.7472542	84.3	119.13	1502.05	1502.1	26.27%
Low-end HW. 562 Teredo clients. 188 IPHTTPS clients.	46.3889778	84.146	118.73	1501.25	1501.2	25.90%
Low-end HW. 562 Teredo clients. 188 IPHTTPS clients.	45.113082	84.0494	118.43	1546.14	1546.1	25.68%

1000 clients, 70/30 split, 78 Mbits/sec throughput:

The following three tests represent performance on low-end hardware. In the test runs below, there were 1000 clients with an average throughput of ~78.64 Mbits/sec and a traffic split of 700 Teredo and 300 IPHTTPS. Teredo MTU was set to 1472 and Teredo Shunt was enabled. CPU utilization averaged ~50.7%, and average memory utilization, expressed as a percentage of committed bytes of the total available memory of 4Gb, was ~28.7%.

Scenario	CPUAvg (from counter)	Mbit/s (Corp Side)	Mbit/s (internet Side)	Active QMSA	Active MMSA	Mem Utilization (4 Gig system)
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	51.28406247	78.6432	113.19	2002.42	1502.1	25.59%
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	51.06993128	78.6402	113.22	2001.4	1501.2	30.87%
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	49.75235617	78.6387	113.2	2002.6	1546.1	30.66%

1000 clients, 70/30 split, 109 Mbits/sec throughput:

In the following three test runs there were 1000 clients with an average throughput of ~109.2 Mbits/sec and a traffic split of 700 Teredo and 300 IPHTTPS. Teredo MTU was set to 1472 and Teredo Shunt was enabled. CPU utilization averaged ~59.06% across the three tests, and average memory utilization, expressed as a percentage of committed bytes of the total available memory of 4Gb, was ~27.34%.

Scenario	CPUAvg (from	Mbit/s (Corp	Mbit/s (internet	Active	Active	Mem Utilization (4

	counter)	Side)	Side)	QMSA	MMSA	Gig system)
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	59.81640675	108.305	153.14	2001.64	2001.6	24.38%
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	59.46473798	110.969	157.53	2005.22	2005.2	28.72%
Low-end HW. 700 Teredo clients. 300 IPHTTPS clients.	57.89089768	108.305	153.14	1999.53	2018.3	24.38%

Testing results for high-end hardware:

Testing was performed with 1500 clients. Traffic split was 70% Teredo and 30% IPHTTPS. All tests involved TCP traffic over Nat64 using 2 IPsec tunnels per client. In all tests, memory utilization was light and CPU utilization was acceptable.

Individual Test Results:

The following sections describe individual tests. Each section title highlights the key elements of each test followed by a summary description of the results and then a chart containing the detailed results data.

1500 clients, 70/30 split, 153.2 Mbits/sec throughput

The following five tests represent high-end hardware. In the below test runs there were 1500 clients with an average throughput of 153.2 Mbits/sec and a traffic split of 1050 Teredo and 450 IPHTTPS. CPU utilization averaged 50.68% across the five tests, and average memory utilization, expressed as a percentage of committed bytes of the total available memory of 8Gb, was 22.25%.

Scenario	CPUAvg (from counter)	Mbit/s (Corp Side)	Mbit/s (internet Side)	Active QMSA	Active MMSA	Mem Utilization (4 Gig system)
High-end HW. 1050 Teredo clients. 450 IPHTTPS clients.	51.712437	157.029	216.29	3000.31	3046	21.58%
High-end HW. 1050 Teredo clients. 450 IPHTTPS clients.	48.86020205	151.012	206.53	3002.86	3045.3	21.15%
High-end HW. 1050 Teredo clients. 450 IPHTTPS clients.	52.23979519	155.511	213.45	3001.15	3002.9	22.90%
High-end HW. 1050 Teredo clients. 450 IPHTTPS clients.	51.26269767	155.09	212.92	3000.74	3002.4	22.91%
High-end HW. 1050 Teredo clients. 450 IPHTTPS clients.	50.15751307	154.772	211.92	3000.9	3002.1	22.93%
High-end HW. 1050	49.83665607	145.994	201.92	3000.51	3006	22.03%

Teredo clients. 450 IPHTTPS clients.						

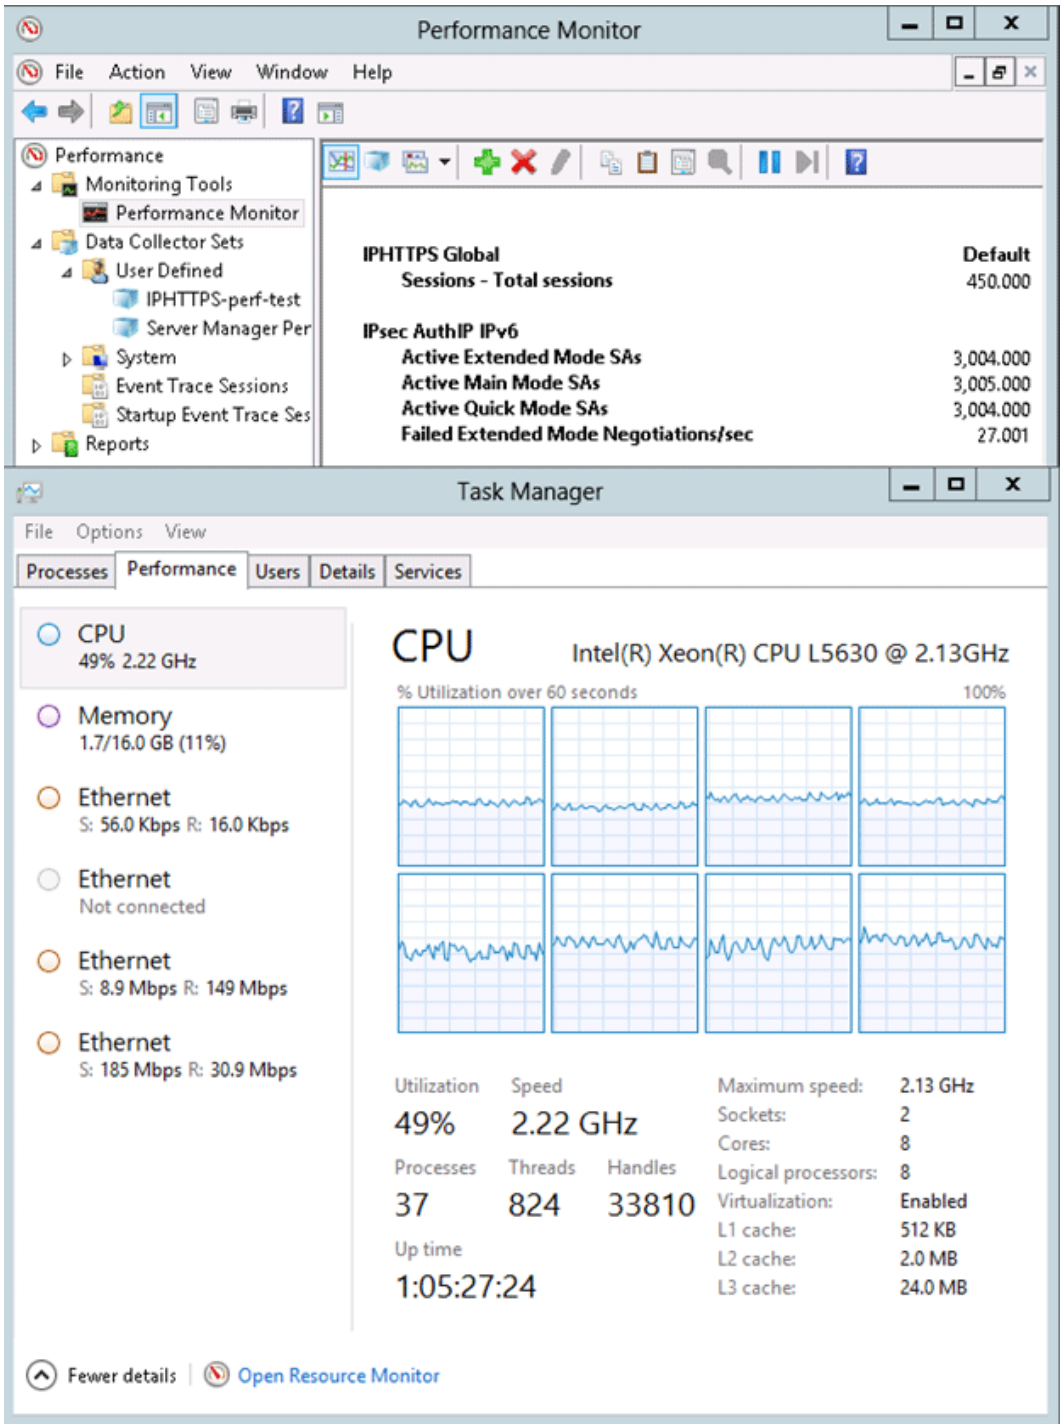
Performance Monitor
File Action View Window Help
Performance
Monitoring Tools
Performance Monitor
Data Collector Sets
User Defined
IPHTTPS-perf-test
Server Manager Per
System
Event Trace Sessions
Startup Event Trace Ses
Reports
IPHTTPS Global
Default
Sessions - Total sessions
450.000
IPsec AuthIP IPv6
Active Extended Mode SAs
3,004.000
Active Main Mode SAs
3,005.000
Active Quick Mode SAs
3,004.000
Failed Extended Mode Negotiations/sec
27.001
Task Manager
File Options View
Processes Performance Users Details Services
CPU
49% 2.22 GHz
Memory
1.7/16.0 GB (11%)
Ethernet
S: 56.0 Kbps R: 16.0 Kbps
Ethernet
Not connected
Ethernet
S: 8.9 Mbps R: 149 Mbps
Ethernet
S: 185 Mbps R: 30.9 Mbps
CPU
Intel(R) Xeon(R) CPU L5630 @ 2.13GHz
% Utilization over 60 seconds
100%
Utilization
Speed
49%
2.22 GHz
Processes
Threads
Handles
37
824
33810
Up time
1:05:27:24
Maximum speed: 2.13 GHz
Sockets: 2
Cores: 8
Logical processors: 8
Virtualization: Enabled
L1 cache: 512 KB
L2 cache: 2.0 MB
L3 cache: 24.0 MB
Fewer details
Open Resource Monitor

Deploy a Single Remote Access Server using the Getting Started Wizard

Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This topic provides an introduction to the Remote Access advanced scenario, which is used to set up a single Remote Access server in a few easy steps.

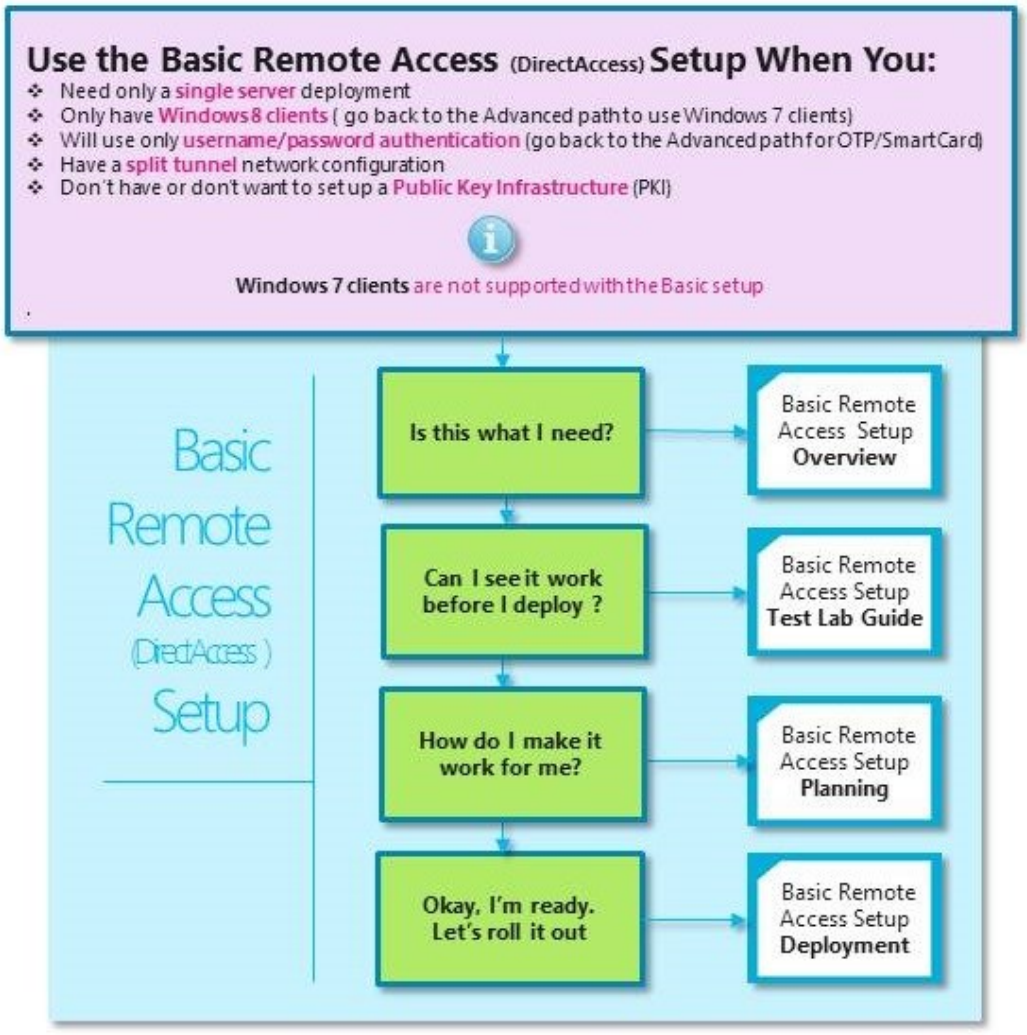

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Scenario description

In this scenario a single computer running Windows Server 2012 is configured as a Remote Access server in a few easy steps.

1. In the simple scenario, Remote Access is set up with default settings in a few easy wizard steps, without any need to configure infrastructure settings such as a certification authority (CA) or Active Directory security groups.
2. If you want to configure an advanced deployment with custom settings, see Deploy a Single Remote Access Server with Advanced Settings
3. If you want to configure Remote Access with enterprise features such as a load-balanced cluster, multisite deployment, or two-factor client authentication, complete the scenario described in this topic to set up a single server, and then set up the required enterprise scenario using Deploy Remote Access in an Enterprise.

In this scenario

To set up a simple Remote Access server, several planning and deployment steps are required.

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- Windows firewall must be enabled on all profiles
- Only supported for Windows 8, Windows 8.1 clients.
- ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
- A Public Key Infrastructure is not required.
- Not supported for deploying two factor authentication. Domain credentials are required for authentication.
- Automatically deploys DirectAccess to all mobile computers in the current domain.
- Traffic to Internet does not go over the DirectAccess tunnel. Force tunnel configuration is not supported.
- DirectAccess server is the Network Location Server.
- Network Access Protection (NAP) is not supported.

- Changing policies outside of the DirectAccess management console or PowerShell cmdlets is not supported.
- To deploy multisite, now or in the future, first Deploy a Single Remote Access Server with Advanced Settings.

Planning steps

Planning is divided into two phases:

1. Planning for the Remote Access infrastructure—This phase describes the planning required to set up the network infrastructure before beginning the Remote Access deployment. It includes planning the network and server topology, and the DirectAccess network location server.
2. Planning for the Remote Access deployment—This phase describes the planning steps required to prepare for the Remote Access deployment. It includes planning for Remote Access client computers, server and client authentication requirements, VPN settings, infrastructure servers, and management and application servers.

For detailed planning steps, see Plan an Advanced Remote Access Deployment.

Deployment steps

Deployment is divided into three phases:

1. Configuring the Remote Access infrastructure—This phase includes configuring network and routing, configuring firewall settings if required, configuring certificates, DNS servers, Active Directory and GPO settings, and the DirectAccess network location server.
2. Configuring Remote Access server settings—This phase includes steps for configuring Remote Access client computers, the Remote Access server, infrastructure servers, management and application servers.
3. Verifying the deployment—This phase includes steps to verify that the deployment is working as required.

For detailed deployment steps, see Install and Configure Basic Remote Access.

Practical applications

Deploying a single Remote Access server provides the following:

- Ease-of-access—Managed client computers running Windows 8 and Windows 7 can be configured as DirectAccess client computers. These clients can access internal network resources via DirectAccess any time they are located on the Internet without needing to log in to a VPN connection. Client computers not running one of these operating systems can connect to the internal network via VPN.
- Ease-of-management—DirectAccess client computers located on the Internet can be remotely managed by Remote Access administrators over DirectAccess, even when the client computers are not located in the internal corporate network. Client computers that do not meet corporate requirements can be remediated automatically by management servers. Both DirectAccess and VPN are managed in the same console and with the same set of

wizards. Additionally, one or more Remote Access servers can be managed from a single Remote Access Management console

Roles and features included in this scenario

The following table lists the roles and features required for the scenario:

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the Server Manager console or Windows PowerShell. This role encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. The Remote Access Server Role is dependent on the following server roles/features: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server on the Remote Access server, and the default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface and Windows PowerShell cmdlets. • It can be optionally installed on a server not

Role/feature	How it supports this scenario
	running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI • Remote Access module for Windows PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

- Server requirements:
  - A computer that meets the hardware requirements for Windows Server 2012.
  - The server must have at least one network adapter installed, enabled, and connected to the internal network. When two adapters are used, there should be one adapter connected to the internal corporate network, and one connected to the external network (Internet, or private network).
  - At least one domain controller. The Remote Access server and DirectAccess clients must be domain members.
- Client requirements:
  - A client computer must be running Windows 8.

Important
If you are using Windows 7 client computers, or a mix of Windows 7 and Windows 8 clients, you must use the Advanced Setup Wizard. The Getting Started Setup Wizard described in this document does not support Windows 7 clients. See Deploy a Single Remote Access Server with Advanced Settings for instructions on how to use Windows 7 clients with DirectAccess.

Note

Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

- Infrastructure and management server requirements:
  - If VPN is enabled, a DHCP server is required to allocate IP addresses automatically to VPN clients, if a static address pool is not used.
- A DNS server running Windows Server 2008 SP2; Windows Server 2008 R2; or Windows Server 2012 is required.

Software requirements

There are a number of requirements for this scenario:

- Server requirements:
  - The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
  - If the Remote Access server is located behind an edge firewall or NAT device, the device must be configured to allow traffic to and from the Remote Access server.
  - The person deploying remote access on the server requires local administrator permissions on the server, and domain user permissions. In addition, the administrator requires permissions for the GPOs used in DirectAccess deployment. To take advantage of the features that restricts DirectAccess deployment to mobile computers only, permissions to create a WMI filter on the domain controller are required.
- Remote Access client requirements:
  - DirectAccess clients must be domain members. Domains containing clients can belong to the same forest as the Remote Access server, or have a two-way trust with the Remote Access server forest.
  - An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients. If a security group is not specified when configuring DirectAccess client settings, by default the client GPO is applied on all laptop computers in the Domain Computers security group. Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB

Content type	References
	Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
Deployment	Remote Access (DirectAccess, Routing and Remote Access) Overview
Tools and settings	Remote Access PowerShell cmdlets
Community resources	DirectAccess Wiki entries
Related technologies	How IPv6 works

Plan a Basic Remote Access Deployment

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This topic describes the planning steps required in order to deploy a single Windows Server® 2012 Remote Access server with basic features:

1. Step 1: Plan the Remote Access Infrastructure—Plan network and server topology, firewall settings, certificate requirements, DNS and Active Directory.
2. Step 2: Plan the Remote Access Deployment—Plan client and server deployment.

Next steps

After you have completed these planning steps, you can begin the server deployment. For instructions, see Install and Configure Basic Remote Access.

Step 1: Plan the Remote Access Infrastructure

The first step of planning for a basic Remote Access deployment on a single server is to perform planning for the infrastructure required for the deployment. This topic describes the infrastructure planning steps:

Task	Description
Plan network topology and settings	Decide where to place the Remote Access server (at the edge, or behind a Network

Task	Description
	Address Translation (NAT) device or firewall), and plan IP addressing and routing.
Plan firewall requirements	Plan for allowing Remote Access through edge firewalls.
Plan certificate requirements	Remote Access can use Kerberos or certificates for client authentication. In this basic Remote Access deployment a Kerberos Proxy is automatically configured and authentication is accomplished using Active Directory credentials.
Plan DNS requirements	Plan DNS settings for the Remote Access server, infrastructure servers, and client connectivity.
Plan Active Directory	Plan your domain controllers and Active Directory requirements.
Plan Group Policy Objects	Decide what GPOs are required in your organization and how to create or edit the GPOs.

The planning tasks do not need to be done in a specific order.

Plan network topology and settings

Plan network adapters and IP addressing

1. Identify the network adapter topology you want to use. Remote Access can be set up with either of the following:
   - With two network adapters—Either at the edge with one network adapter connected to the Internet and the other to the internal network, or behind a NAT, firewall, or router device, with one network adapter connected to a perimeter network and the other to the internal network.
   - Behind a NAT device with one network adapter—The Remote Access server is installed behind a NAT device, and the single network adapter is connected to the internal network.
2. Identity your IP addressing requirements:

   DirectAccess uses IPv6 with IPsec to create a secure connection between DirectAccess client computers and the internal corporate network. However, DirectAccess does not necessarily require connectivity to the IPv6 Internet or native IPv6 support on internal networks. Instead, it automatically configures and uses IPv6 transition technologies to tunnel

IPv6 traffic across the IPv4 Internet (6to4, Teredo, IP-HTTPS) and across your IPv4-only intranet (NAT64 or ISATAP). For an overview of these transition technologies, see the following resources:

- IPv6 Transition Technologies
- IP-HTTPS Tunneling Protocol Specification

3. Configure required adapters and addressing according to the following table. For deployments behind a NAT device using a single network adapter, configure your IP addresses using only the 'Internal network adapter' column.

	External network adapter	Internal network adapter[1]	Routing requirements
IPv4 intranet and IPv4 Internet	Configure the following: • One static public IPv4 address with the appropriate subnet mask. • A default gateway IPv4 address of your Internet firewall or local Internet service provider (ISP) router.	Configure the following: • An IPv4 intranet address with the appropriate subnet mask. • A connection-specific DNS suffix of your intranet namespace. A DNS server must also be configured on the internal interface. • Do not configure a default gateway on any intranet interfaces.	To configure the Remote Access server to reach all subnets on the internal IPv4 network do the following: 1. List the IPv4 address spaces for all the locations on your intranet. 2. Use the **route add -p** or **netsh interface ipv4 add route** commands to add the IPv4 address spaces as static routes in the IPv4 routing table of the Remote Access server.
IPv6 Internet and IPv6 intranet	Configure the following: • Use the autoconfigured address configuration provided by your ISP. • Use the **route print** command to ensure that a	Configure the following: • If you are not using default preference levels, configure your intranet interfaces with the **netsh interface ipv6 set InterfaceIndex ignoredefaultroutes=enabled** command. This command ensures that additional default routes pointing to intranet routers will not be added to the IPv6 routing	If you have an IPv6 intranet, to configure the Remote Access server to reach all of the IPv6 locations, do the following: 3. 0.1. List the IPv6 address spaces for all

	External network adapter	Internal network adapter[1]	Routing requirements
	default IPv6 route pointing to the ISP router exists in the IPv6 routing table. • Determine whether the ISP and intranet routers are using default router preferences described in RFC 4191, and using a higher default preference than your local intranet routers. If both of these are true, no other configuration for the default route is required. The higher preference for the ISP router ensures that the active default IPv6 route of the Remote Access server points to the IPv6 Internet. Because the Remote Access server is an IPv6 router, if you have a native IPv6 infrastructure, the Internet interface can also reach the domain controllers on the intranet. In this case, add packet filters to the domain controller in the perimeter network that prevent	table. You can obtain the InterfaceIndex of your intranet interfaces from the display of the netsh interface show interface command.	the locations on your intranet. 2. Use the **netsh interface ipv6 add route** command to add the IPv6 address spaces as static routes in the IPv6 routing table of the Remote Access server.

	External network adapter	Internal network adapter[1]	Routing requirements
	connectivity to the IPv6 address of the Internet-facing interface of the Remote Access server.		
IPv4 Internet and IPv6 intranet	The Remote Access server forwards default IPv6 route traffic using the Microsoft 6to4 Adapter interface to a 6to4 relay on the IPv4 Internet. You can configure a Remote Access server for the IPv4 address of the Microsoft 6to4 relay on the IPv4 Internet (used when native IPv6 is not deployed in the corporate network) with the following command : netsh interface ipv6 6to4 set relay name=192.88.99.1 state=enabled command.		

Notes

Note the following:

a. If the DirectAccess client has been assigned a public IPv4 address, it will use the 6to4 transition technology to connect to the intranet. If the DirectAccess client cannot connect to the DirectAccess server with 6to4, it will use IP-HTTPS.

b. Native IPv6 client computers can connect to the Remote Access server over native IPv6, and no transition technology is required.

Plan firewall requirements

If the Remote Access server is behind an edge firewall, the following exceptions will be required for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound.
- If you are deploying Remote Access with a single network adapter, and installing the network location server on the Remote Access server, TCP port 62000 should also be exempted.

Note

This exemption is on the Remote Access server. All the other exceptions are on the edge firewall.

The following exceptions will be required for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic

Plan certificate requirements

Certificate requirements for IPsec include a computer certificate used by DirectAccess client computers when establishing the IPsec connection between the client and the Remote Access server, and a computer certificate used by Remote Access servers to establish IPsec connections with DirectAccess clients. For DirectAccess in Windows Server® 2012 the use of these IPsec certificates is not mandatory. The Getting Started Wizard configures the Remote Access server to act as a Kerberos proxy to perform IPsec authentication without requiring certificates.

1. **IP-HTTPS server**—When you configure Remote Access, the Remote Access server is automatically configured to act as the IP-HTTPS web listener. The IP-HTTPS site requires a website certificate, and client computers must be able to contact the certificate revocation list (CRL) site for the certificate. The Enable DirectAccess wizard tries to use the SSTP VPN certificate. If SSTP is not configured, it checks if a certificate for IP-HTTPS is present in the machine personal store. If none is available, it automatically creates a self-signed certificate.
2. **Network location server**—The network location server is a website used to detect whether client computers are located in the corporate network. The network location server requires a website certificate. DirectAccess clients must be able to contact the CRL site for the certificate. The Enable DirectAccess wizard checks if a certificate for Network Location Server is present in the machine personal store. If not present, it automatically creates a self-signed certificate.

The certification requirements for each of these are summarized in the following table:

IPsec authentication	IP-HTTPS server	Network location server
An internal CA is required to issue computer certificates to the Remote Access server and clients for IPsec authentication when you don't use the Kerberos proxy for authentication	Public CA—It is recommended to use a public CA to issue the IP-HTTPS certificate, this ensures that the CRL distribution point is available externally.	Internal CA—You can use an internal CA to issue the network location server website certificate. Make sure that the CRL distribution point is highly available from the internal network.
	Internal CA—You can use an internal CA to issue the IP-HTTPS certificate; however, you must make sure that the CRL distribution point is available externally.	Self-signed certificate—You can use a self-signed certificate for the network location server website; however, you cannot use a self-signed certificate in multisite deployments.
	Self-signed certificate—You can use a self-signed certificate for the IP-HTTPS server; however, you must make sure that the CRL distribution point is available externally. A self-signed certificate cannot be used in a multisite deployment.	

Plan certificates for IP-HTTPS and network location server

If you want to provision a certificate for these purposes, refer to Deploy a Single Remote Access Server with Advanced Settings. If no certificates are available, the Getting Started wizard automatically creates self-signed certificates for these purposes.

Note

If you provision certificates for IP-HTTPS and the network location server manually, ensure that the certificates have a subject name. If the certificate does not have a subject name, but does have an alternative name, it will not be accepted by the Remote Access wizard.

Plan DNS requirements

In a Remote Access deployment, DNS is required for the following:

- **DirectAccess client requests**—DNS is used to resolve requests from DirectAccess client computers that are not located on the internal network. DirectAccess clients attempt to

connect to the DirectAccess network location server in order to determine whether they are located on the Internet, or on the corporate network: If the connection is successful, then clients are determined to be on the intranet and DirectAccess is not used, and client requests are resolved using the DNS server configured on the network adapter of the client computer. If the connection does not succeed, clients are assumed to be on the Internet. DirectAccess clients will use the name resolution policy table (NRPT) to determine which DNS server to use when resolving name requests. You can specify that clients should use DirectAccess DNS64 to resolve names, or an alternative internal DNS server. When performing name resolution, the NRPT is used by DirectAccess clients to identify how to handle a request. Clients request an FQDN or single-label name such as http://internal. If a single-label name is requests, a DNS suffix is appended to make an FQDN. If the DNS query matches an entry in the NRPT, and DNS4 or an intranet DNS server is specified for the entry, then the query is sent for name resolution using the specified server. If a match exists but no DNS server is specified, then this indicates an exemption rule and normal name resolution is applied.

Note that when a new suffix is added to the NRPT in the Remote Access Management console, the default DNS servers for the suffix can be automatically discovered by clicking the **Detect** button. Auto detection works as follows:

a. If the corporate network is IPv4-based, or IPv4 and IPv6, the default address is the DNS64 address of the internal adapter on the Remote Access server.

b. If the corporate network is IPv6-based, the default address is the IPv6 address of DNS servers in the corporate network.

- **Infrastructure servers**—

  a. **Network location server**—DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, but must be prevented from resolving the name when they are located on the Internet. To ensure this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT. In addition, when you configure Remote Access, the following rules are created automatically:

     i. A DNS suffix rule for root domain or the domain name of the Remote Access server, and the IPv6 addresses corresponding to the intranet DNS servers configured on the Remote Access server. For example, if the Remote Access server is a member of the corp.contoso.com domain, a rule is created for the corp.contoso.com DNS suffix.

     ii. An exemption rule for the FQDN of the network location server. For example, if the network location server URL is https://nls.corp.contoso.com, an exemption rule is created for the FQDN nls.corp.contoso.com.

     IP-HTTPS server—The Remote Access server acts as an IP-HTTPS listener and uses its server certificate to authenticate to IP-HTTPS clients. The IP-HTTPS name must be resolvable by DirectAccess clients using public DNS servers.

     Connectivity verifiers—Remote Access creates a default web probe that is used by DirectAccess client computers use to verify connectivity to the internal network. To ensure the probe works as expected the following names must be registered manually in DNS:

i. directaccess-webprobehost—should resolve to the internal IPv4 address of the Remote Access server, or to the IPv6 address in an IPv6-only environment.

ii. directaccess-corpconnectivityhost—should resolve to localhost (loopback) address. A and AAAA record should be created, A record with value 127.0.0.1 and AAAA record with value constructed out of NAT64 prefix with the last 32 bits as 127.0.0.1. The NAT64 prefix can be retrieved by running the cmdlet get-netnattransitionconfiguration.

You can create additional connectivity verifiers using other web addresses over HTTP or PING. For each connectivity verifier, a DNS entry must exist.

DNS server requirements

- For DirectAccess clients, you must use either a DNS server running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or any DNS server that supports IPv6.

Plan the network location server

The network location server is a website used to detect whether DirectAccess clients are located in the corporate network. Clients in the corporate network do not use DirectAccess to reach internal resources, but instead connect directly.

The Getting Started Wizard automatically sets up network location server on the Remote Access server and the website is created automatically when you deploy Remote Access. This allows for a simple installation without the use of a certificate infrastructure.

If you want to deploy a Network Location Server and not use self-signed certificates, refer to Deploy a Single Remote Access Server with Advanced Settings.

Plan Active Directory

Remote Access uses Active Directory and Active Directory Group Policy Objects as follows:

- **Authentication**—Active Directory is used for authentication. The DirectAccess tunnel uses Kerberos authentication for the user to access internal resources.
- **Group policy objects**—Remote Access gathers configuration settings into group policy objects that are applied to Remote Access servers, and clients.
- **Security groups**—Remote Access uses security groups to gather together and identify DirectAccess client computers, and Remote Access servers. The group policies are applied to the required security group.

Active Directory Requirements

When planning Active Directory for a Remote Access deployment, the following is required:

- At least one domain controller installed on the Windows Server 2012, Windows Server 2008 R2 Windows Server 2008, or Windows Server 2003 operating systems.

  If the domain controller is on a perimeter network (and therefore reachable from the Internet-facing network adapter of Remote Access server) prevent the Remote Access server from

reaching it by adding packet filters on the domain controller, to prevent connectivity to the IP address of the Internet adapter.

- The Remote Access server must be a domain member.
- DirectAccess clients must be domain members. Clients can belong to:
  - Any domain in the same forest as the Remote Access server.
  - Any domain that has a two-way trust with the Remote Access server domain.
  - Any domain in a forest that has a two-way trust with the forest to which the Remote Access domain belongs.

Notes

- The Remote Access server cannot be a domain controller.
- The Active Directory domain controller used for Remote Access must not be reachable from the external Internet adapter of the Remote Access server (the adapter must not be in the domain profile of Windows Firewall).

Plan Group Policy Objects

DirectAccess settings configured when you configure Remote Access are collected into Group Policy Objects (GPO). Two different GPOs are populated with DirectAccess settings, and distributed as follows:

- **DirectAccess client GPO**—This GPO contains client settings, including IPv6 transition technology settings, NRPT entries, and Windows Firewall with Advanced Security connection security rules. The GPO is applied to the security groups specified for the client computers.
- **DirectAccess server GPO**—This GPO contains the DirectAccess configuration settings that are applied to any server configured as a Remote Access server in your deployment. It also contains Windows Firewall with Advanced Security connection security rules.

GPOs can be configured in two ways:

1. **Automatically**—You can specify that they are created automatically. A default name is specified for each GPO. GPOs are created automatically by the Getting Started Wizard.
2. **Manually**—You can use GPOs that have been predefined by the Active Directory administrator.

Note that once DirectAccess is configured to use specific GPOs, it cannot be configured to use different GPOs.

Important

Whether you are using automatically or manually configured GPOs, you will need to add a policy for slow link detection if your clients will use 3G. The Group Policy path for **Policy: Configure Group Policy slow link detection** is: **Computer configuration / Polices / Administrative Templates / System / Group Policy**.

Caution

Use the following procedure to backup all Remote Access Group Policy Objects before executing DirectAccess cmdlets: Back up and Restore Remote Access Configuration

Automatically-created GPOs

Note the following when using automatically-created GPOs:

Automatically created GPOS are applied according to the location and link target parameter, as follows:

- For the DirectAccess server GPO, both the location and link parameters point to the domain containing the Remote Access server.
- When client GPOs are created, the location is set to a single domain in which the GPO will be created. The GPO name is looked up in each domain, and filled with DirectAccess settings if it exists. The link target is set to the root of the domain in which the GPO was created. A GPO is created for each domain that contains client computers, and the GPO is linked to the root of its respective domain.

When using automatically created GPOs, to apply DirectAccess settings, the Remote Access server administrator requires the following permissions:

- GPO create permissions for each domain.
- Link permissions for all the selected client domain roots.
- Link permissions for the server GPO domain roots.
- Create, edit, delete, and modify security permissions are required for the GPOs.
- It is recommended that the Remote Access administrator has GPO read permissions for each required domain. This enables Remote Access to verify that GPOs with duplicate names do not exist when creating GPOs.

Note that if the correct permissions for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but linking will not occur. If this warning is issued links will not be created automatically, even after the permissions are added later. Instead the administrator will need to create the links manually.

Manually-created GPOs

Note the following when using manually-created GPOs:

- The GPOs should exist before running the Remote Access Getting Started wizard.
- When using manually-created GPOs, to apply DirectAccess settings the Remote Access administrator requires full GPO permissions (Edit, Delete, Modify security) on the manually-created GPOs.
- When using manually created GPOs a search is made for a link to the GPO in the entire domain. If the GPO is not linked in the domain then a link is automatically created in the domain root. If the required permissions to create the link are not available a warning is issued.

Note that if the correct permissions for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but linking will not occur. If this warning is issued links will not be created automatically, even when the permissions are added later. Instead the administrator will need to create the links manually.

Recovering from a deleted GPO

If a Remote Access server, client, or application server GPO has been deleted by accident and there is no backup available, you must remove the configuration settings and re-configure again. If a backup is available, you can restore the GPO from the backup.

Remote Access Management will display the following error message: **GPO <GPO name> cannot be found**. To remove the configuration settings, take the following steps:

1. Run the PowerShell cmdlet **Uninstall-remoteaccess**.
2. Re-open **Remote Access Management**.
3. You will see an error message that the GPO is not found. Click **Remove configuration settings**. After completion, the server will be restored to an un-configured state.

See also

- Step 2: Plan the Remote Access Deployment

Step 2: Plan the Remote Access Deployment

After planning the Remote Access infrastructure, the next step in deploying Remote Access on a single server with basic settings is to plan the settings for the Getting Started Wizard.

Task	Description
Planning for client deployment	By default, the Getting Started Wizard deploys DirectAccess to all laptops and notebook computers in the domain by applying a WMI filter to the client settings GPO
Planning for Remote Access server deployment	Plan how to deploy the Remote Access server.

Planning for client deployment

There are two decisions to make when planning your client deployment:

1. Will DirectAccess be available to mobile computers only, or to any computer?

   When you configure DirectAccess clients in the Getting Started Wizard, you can choose to allow only mobile computers in the specified security groups to connect using DirectAccess. If you restrict access to mobile computers, Remote Access automatically configures a WMI filter to ensure that the DirectAccess client GPO is applied only to mobile computers in the specified security groups. The Remote Access administrator requires permissions to create or modify group policy WMI filters to enable this setting.
2. What security groups will contain the DirectAccess client computers?

DirectAccess settings are contained in the DirectAccess client GPO. The GPO is applied to computers that are part of the security groups that you specify in the Getting Started wizard. You can specify security groups contained in any supported domain. Before you configure Remote Access, the security groups should be created. You can add computers to the security group after completing the Remote Access deployment, but note that if you add client computers that reside in a different domain to the security group, the client GPO will not be applied to these clients. For example, if you created SG1 in domain A for DirectAccess clients, and later add clients from domain B to this group, the client GPO will not be applied to clients in domain B. To avoid this issue, create a new client security group for each domain that contains client computers. Alternatively, if you do not want to create a new security group, run the Add-DAClient cmdlet with the name of the new GPO for the new domain.

Planning for Remote Access server deployment

There are a number of decisions to make when planning to deploy your Remote Access server:

- **Network topology**—There are two topologies available when deploying a Remote Access server:
  - **Two adapters**—With two network adapters, Remote Access can be configured with one network adapter connected directly to the Internet, and the other is connected to the internal network. Or alternatively the server is installed behind an edge device such as a firewall or a router. In this configuration one network adapter is connected to the perimeter network, the other is connected to the internal network.
  - **Single network adapter**—In this configuration the Remote Access server is installed behind an edge device such as a firewall or a router. The network adapter is connected to the internal network.
- **Network adapters**—The Remote Access Getting Started wizard automatically detects the network adapters configured on the Remote Access server. You can make sure that the correct adapters are selected from the **Review** page.
- **IP-HTTPS certificate**—Since there is no PKI deployment in this deployment, the wizard will automatically provision self-signed certificates for IP-HTTPS and the Network Location Server (if no certificates are present), and will automatically enable Kerberos proxy. The wizard will also enable NAT64 and DNS64 for protocol translation in the IPv4-only environment. After the wizard successfully completes applying the configuration, click **Close**.
- **Windows 7 clients**—You cannot enable support for Windows 7 clients from the Getting Started wizard. This can be enabled from the Advanced Setup Wizard. For more details, see Deploy a Single Remote Access Server with Advanced Settings.
- **VPN configuration**—Before you configure Remote Access, decide if you are going to provide VPN access to remote clients. You should provide VPN access if you have client computers in your organization that do not support DirectAccess connectivity (either because they are unmanaged or run an operating system for which DirectAccess is not supported). The Getting Started Wizard configures VPN IP address assignment using DHCP and configures VPN clients to be authenticated using Active Directory.
- **Force Tunneling** — If you plan to use Force Tunneling, or might add it in the future, you should use Deploy a Single Remote Access Server with Advanced Settings to deploy a two

tunnel configuration. Because of security considerations, Force Tunneling in a single tunnel configuration is not supported.

See also

- Step 1: Plan the Remote Access Infrastructure

Install and Configure Basic Remote Access

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the configuration steps required in order to deploy a single Windows Server 2012 Remote Access server with basic settings.

- Step 1: Configure the Remote Access Infrastructure. This step includes configuring network and server settings, DNS settings and Active Directory settings.
- Step 2: Configure the Remote Access Server. This step includes configuring DirectAccess client computers, server settings.
- Step 3: Verify the Deployment. This step includes steps for verifying the deployment.

Before starting the deployment, verify the planning steps described in Plan a Basic Remote Access Deployment.

Step 1: Configure the Remote Access Infrastructure

This topic describes how to configure the infrastructure required for a basic Remote Access deployment using a single Remote Access server in a mixed IPv4 and IPv6 environment. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 1: Plan the Remote Access Infrastructure.

Task	Description
Configure server network settings	Configure the server network settings on the Remote Access server.
Configure routing in the corporate network	Configure routing in the corporate network to make sure traffic is appropriately routed.
Configure firewalls	Configure additional firewalls, if required.
Configure the DNS server	Configure DNS settings for the Remote Access server.

Task	Description
Configure Active Directory	Join client computers and the Remote Access server to the Active Directory domain.
Configure GPOs	Configure GPOs for the deployment, if required.
Configure security groups	Configure security groups that will contain DirectAccess client computers, and any other security groups required in the deployment.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Configure server network settings

The following network interface settings are required for a single server deployment in an environment with IPv4 and IPv6. All IP addresses are configured by using **Change adapter settings** in the **Windows Networking and Sharing Center**.

- Edge topology
  - One Internet-facing public static IPv4 or IPv6 addresses.

    Note
    Two consecutive public IPv4 addresses are required for Teredo. If you are not using Teredo, you can configure a single public static IPv4 address.
  - A single internal static IPv4 or IPv6 address.
- Behind NAT device (two network adapters)
  - A single internal network-facing static IPv4 or IPv6 address.
- Behind NAT device (one network adapter)
  - A single static IPv4 or IPv6 address.

Notes
In the event that the Remote Access server has two or more network adapters (one classified in the domain profile and the other in a public/private profile), but a single NIC topology will be used, then the recommendation is as follows:

1. Ensure that the 2nd, and any additional, NIC is also classified in the domain profile - Recommended.
2. If the 2nd NIC cannot be configured for the domain profile for any reason, then the DirectAccess IPsec policy must be manually scoped to all profiles using the following Windows PowerShell commands:

```
$gposession = Open-NetGPO –PolicyStore <Name of the server GPO>
```

```
Set-NetIPsecRule –DisplayName <Name of the IPsec policy> –
GPOSession $gposession –Profile Any
Save-NetGPO –GPOSession $gposession
```

The names of the IPsec policies are DirectAccess-DaServerToInfra and DirectAccess-DaServerToCorp.

Configure routing in the corporate network

Configure routing in the corporate network as follows:

- When native IPv6 is deployed in the organization, add a route so that the routers on the internal network route IPv6 traffic back through the Remote Access server.
- Manually configure organization IPv4 and IPv6 routes on the Remote Access servers. Add a published route so that all traffic with an organization (/48) IPv6 prefix is forwarded to the internal network. In addition, for IPv4 traffic, add explicit routes so that IPv4 traffic is forwarded to the internal network.

Configure firewalls

When using additional firewalls in your deployment, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required.

  Note
  This exemption has to be configured on the remote access server. All the other exemptions have to be configured on the edge firewall.

Note
For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server. For IP-HTTPS the exceptions need only be applied for the address to which the external name of the server resolves.

When using additional firewalls, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound

- TCP/UDP for all IPv4/IPv6 traffic

Configure the DNS server

You must manually configure a DNS entry for the network location server website for the internal network in your deployment.

To create the network location server and NCSI probe DNS records

1. On the internal network DNS server: On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the left pane of the **DNS Manager** console, expand the forward lookup zone for your domain. Right click the domain and click **New Host (A or AAAA)**.
3. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the network location server website (this is the name the DirectAccess clients use to connect to the network location server). In the **IP address** box, enter the IPv4 address of the network location server, and then click **Add Host**. On the **DNS** dialog box, click **OK**.
4. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the web probe (the name for the default web probe is directaccess-webprobehost). In the **IP address** box, enter the IPv4 address of the web probe, and then click **Add Host**. Repeat this process for directaccess-corpconnectivityhost and any manually created connectivity verifiers. On the **DNS** dialog box, click **OK**.
5. Click **Done**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-DnsServerResourceRecordA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv4Address <network_location_server_IPv4_address>

Add-DnsServerResourceRecordAAAA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv6Address <network_location_server_IPv6_address>
```

You must also configure DNS entries for the following:

- **The IP-HTTPS server**—DirectAccess clients must be able to resolve the DNS name of the Remote Access server from the Internet.
- **CRL revocation checking**—DirectAccess uses certificate revocation checking for the IP-HTTPS connection between DirectAccess clients and the Remote Access server, and for the HTTPS-based connection between the DirectAccess client and the network location server. In both cases, DirectAccess clients must be able to resolve and access the CRL distribution point location.

Configure Active Directory

The Remote Access server and all DirectAccess client computers must be joined to an Active Directory domain. DirectAccess client computers must be a member of one of the following domain types:

- Domains that belong to the same forest as the Remote Access server.
- Domains that belong to forests with a two-way trust with the Remote Access server forest.
- Domains that have a two-way domain trust to the Remote Access server domain.

To join the Remote Access server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. On the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In **Computer Name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
5. When you see a dialog box welcoming you to the domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

To join client computers to the domain

1. On the **Start** screen, type **explorer.exe**, and then press ENTER.
2. Right-click the Computer icon, and then click **Properties**.
3. On the **System** page, click **Advanced system settings**.
4. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
5. In **Computer name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
6. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
7. When you see a dialog box welcoming you to the domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click Close. Click **Restart Now** when prompted.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note that you must supply domain credentials after entering the Add-Computer command below.

```
Add-Computer -DomainName <domain_name>

Restart-Computer
```

Configure GPOs

To deploy Remote Access, you require a minimum of two Group Policy Objects: one Group Policy Object contains settings for the Remote Access server and one contains settings for DirectAccess client computers. When you configure Remote Access, the wizard automatically creates the required Group Policy Objects. However, if your organization enforces a naming convention, or you do not have the required permissions to create or edit Group Policy Objects, they must be created prior to configuring Remote Access.

To create Group Policy Objects, see Create and Edit a Group Policy Object.

Important

The administrator can manually link the DirectAccess Group Policy Objects to an Organizational Unit using these steps:

1. Before configuring DirectAccess, link the created GPOs to the respective Organizational Units.
2. Configure DirectAccess, specifying a security group for the client computers.
3. The administrator may or may not have permissions to link the Group Policy Objects to the domain. In either case, the Group Policy Objects will be configured automatically. If the GPOs are already linked to an OU, the links will not be removed. Nor will the GPOs be linked to the domain. For server GPO, the OU must contain the server computer object, else the GPO will be linked to the root of the domain.
4. If the linking to OU has not been done before running the DirectAccess wizard, after the configuration is complete, the administrator can link the DirectAccess Group Policy Objects to the required Organizational Units. The link to the domain can be removed. Steps for linking a Group Policy Object to an Organization Unit can be found here

Note

If a Group Policy Object was created manually, it is possible during the DirectAccess configuration that the Group Policy Object will not be available. The Group Policy Object may not have been replicated to the closest Domain Controller to the management computer. In this event, the administrator can wait for replication to complete, or force the replication.

Configure security groups

The DirectAccess settings contained in the client computer Group Policy Object are applied only to computers that are members of the security groups that you specify when configuring Remote Access.

To create a security group for DirectAccess clients

1. On the **Start** screen, type **dsa.msc**, and then press ENTER. In the **Active Directory Users and Computers** console, in the left pane, expand the domain that will contain the security group, right-click **Users**, point to **New**, and then click **Group**.
2. On the **New Object - Group** dialog box, under **Group name**, enter the name for the security group.
3. Under **Group scope**, click **Global**, under **Group type**, click **Security**, and then click **OK**.
4. Double-click the DirectAccess client computers security group, and on the properties dialog box, click the **Members** tab.
5. On the **Members** tab, click **Add**.
6. On the **Select Users, Contacts, Computers, or Service Accounts** dialog box, select the client computers that you want to enable for DirectAccess, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADGroup -GroupScope global -Name <DirectAccess_clients_group_name>

Add-ADGroupMember -Identity DirectAccess_clients_group_name -Members <computer_name>
```

Step 2: Configure the Remote Access Server

This topic describes how to configure the client and server settings required for a basic Remote Access deployment. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 2: Plan the Remote Access Deployment.

Task	Description
Install the Remote Access role	Install the Remote Access role.
Configure DirectAccess Using the Getting Started Wizard	The new Getting Started Wizard presents a greatly simplified configuration experience. The wizard masks the complexity of DirectAccess, and allows for an automated setup in a few simple steps. The wizard provides a seamless experience for the administrator by configuring

Task	Description
	Kerberos proxy automatically to eliminate the need for an internal PKI deployment.
Update clients with the DirectAccess configuration	To receive the DirectAccess settings, clients must update group policy while connected to the intranet.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Install the Remote Access role

To deploy Remote Access, you must install the Remote Access role on a server in your organization that will act as the Remote Access server.

To install the Remote Access role

1. On the Remote Access server, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. On the **Select server roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature RemoteAccess -IncludeManagementTools
```

Configure DirectAccess with the Getting Started Wizard

To configure DirectAccess using the Getting Started Wizard

1. In Server Manager click **Tools**, and then click **Remote Access Management**.
2. In the Remote Access Management console, click **Run the Getting Started Wizard**.

3. Click **Deploy DirectAccess only**.
4. Select the topology of your network configuration and type the public name to which remote access clients will connect. Click **Next**.

   Note
   By default, the Getting Started Wizard deploys DirectAccess to all laptops and notebook computers in the domain by applying a WMI filter to the client settings GPO.
5. Click **Finish**.
6. Since there is no PKI used in this deployment, if certificates are not found, the wizard will automatically provision self-signed certificates for IP-HTTPS and the Network Location Server, and will automatically enable Kerberos proxy. The wizard will also enable NAT64 and DNS64 for protocol translation in the IPv4-only environment. After the wizard successfully completes applying the configuration, click **Close**.
7. In the console tree of the Remote Access Management console, select **Operations Status**. Wait until the status of all monitors display as "Working". In the Tasks pane under Monitoring, click **Refresh** periodically to update the display.

Update clients with the DirectAccess configuration

To update DirectAccess clients

1. Click the **Search** charm, and then click **Apps**. Right-click **Windows PowerShell**, click **Advanced**, and then click **Run as administrator**.
2. In the PowerShell window, type **gpupdate** and then press **ENTER**.
3. Wait for the computer policy update to complete successfully.
4. Type **Get-DnsClientNrptPolicy** and then press **ENTER**

   The Name Resolution Policy Table (NRPT) entries for DirectAccess are displayed. Note that the NLS server exemption is displayed. The Getting Started wizard automatically created this DNS entry for the DirectAccess server, and provisioned an associated self-signed certificate so that the DirectAccess server can function as the Network Location Server.
5. Type **Get-NCSIPolicyConfiguration** and then press **ENTER**. The network connectivity status indicator settings deployed by the wizard are displayed. Note the value of DomainLocationDeterminationURL. Whenever this network location server URL is accessible, the client will determine that it is inside the corporate network, and NRPT settings will not be applied.
6. Type **Get-DAConnectionStatus** and then press **ENTER**. Since the client can reach the network location server URL, the status will display as **ConnectedLocally**.

See also

- Step 3: Verify the Deployment

Step 3: Verify the Deployment

This topic describes how to verify that you have correctly configured your basic DirectAccess deployment.

To verify access to internal resources through DirectAccess

1. Connect a DirectAccess client computer to the corporate network and obtain the group policy.
2. Click the **Network connections** icon in the notification area to access the DA Media Manager.
3. Click the **DirectAccess Connection**, and you will see that the status is **Connected Locally**.
4. Connect the client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

Deploy a Single Remote Access Server with Advanced Settings

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This topic provides an introduction for setting up a single Remote Access server with basic and advanced features.

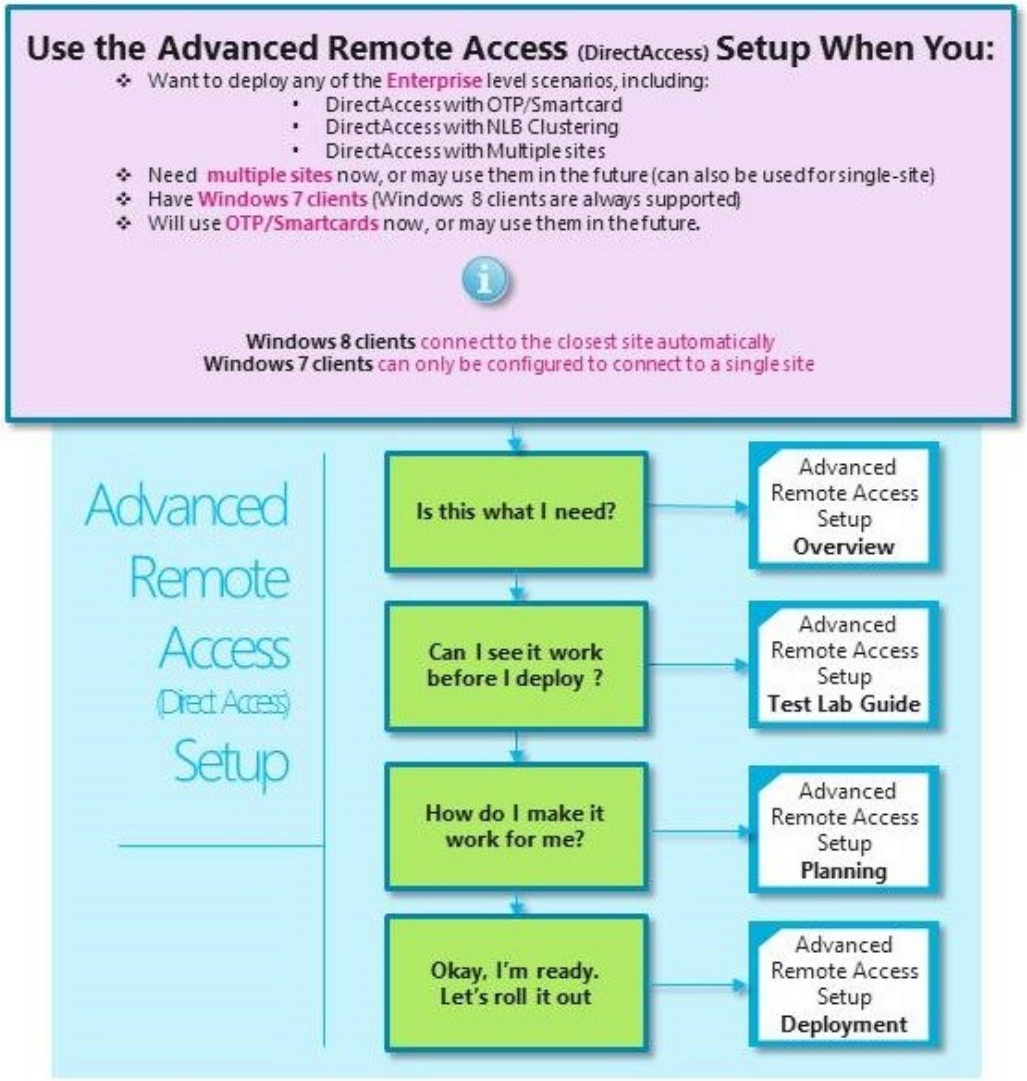

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**

- Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Scenario description

In this scenario, a single computer running Windows Server 2012 is configured as a Remote Access server with advanced settings.

- If you want to configure a basic deployment with simple settings only, see Deploy a Single Remote Access Server using the Getting Started Wizard. In the simple scenario, Remote Access is set up with default settings by using a wizard, without any need to configure infrastructure settings such as a certification authority (CA) or Active Directory security groups.
- If you want to configure Remote Access with enterprise features such as a load-balanced cluster, multisite deployment, or two-factor client authentication, complete the scenario described in this topic to set up a single server, and then set up the required enterprise scenario as described in Deploy Remote Access in an Enterprise.

In this scenario

To set up a single Remote Access server with advanced settings, a number of planning and deployment steps are required.

Prerequisites

Before you begin this deployment scenario, review the following list for important requirements:

1. Windows firewall must be enabled on all profiles.
2. This scenario is only supported for clients running Windows 8.1 or Windows 8.
3. ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
4. A public key infrastructure is not required.
5. This scenario is not supported for deploying two-factor authentication. Domain credentials are required for authentication.
6. Automatically deploys DirectAccess to all mobile computers in the current domain.
7. Traffic to the Internet does not go through the DirectAccess tunnel. Force tunnel configuration is not supported.
8. The DirectAccess server is the network location server.
9. Network Access Protection (NAP) is not supported.
10. Changing policies outside of the DirectAccess management console or by using Windows PowerShell cmdlets is not supported.

Planning steps

Planning is divided into two phases:

1. Planning for the Remote Access infrastructure—This phase describes the planning required to set up the network infrastructure before beginning the Remote Access deployment. It

includes planning the network and server topology, certificate planning, DNS, Active Directory and Group Policy object (GPO) configuration, and the DirectAccess network location server.

2. Planning for the Remote Access deployment—This phase describes the planning steps required to prepare for the Remote Access deployment. It includes planning for Remote Access client computers, server and client authentication requirements, VPN settings, infrastructure servers, and management and application servers.

For detailed planning steps, see Plan an Advanced Remote Access Deployment.

Deployment steps

Deployment is divided into three phases:

1. Configuring the Remote Access infrastructure—This phase includes configuring network and routing, configuring firewall settings if required, configuring certificates, DNS servers, Active Directory and GPO settings, and the DirectAccess network location server.
2. Configuring Remote Access server settings—This phase includes steps for configuring Remote Access client computers, the Remote Access server, infrastructure servers, management and application servers.
3. Verifying the deployment—This phase includes steps to verify that the deployment is working as required.

For detailed deployment steps, see Install and Configure Advanced Remote Access.

Practical applications

Deploying a single Remote Access server provides the following:

- Ease-of-access—Managed client computers running Windows 8 and Windows 7 can be configured as DirectAccess client computers. These clients can access internal network resources via DirectAccess any time they are located on the Internet without needing to log in to a VPN connection. Client computers not running one of these operating systems can connect to the internal network via VPN.
- Ease-of-management—DirectAccess client computers located on the Internet can be remotely managed by Remote Access administrators over DirectAccess, even when the client computers are not located in the internal corporate network. Client computers that do not meet corporate requirements can be remediated automatically by management servers. Both DirectAccess and VPN are managed in the same console and with the same set of wizards. Additionally, one or more Remote Access servers can be managed from a single Remote Access Management console

Roles and features required for this scenario

The following table lists the roles and features that are required for this scenario:

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the

Role/feature	How it supports this scenario
	Server Manager console or Windows PowerShell. This role encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. The Remote Access Server Role is dependent on the following server roles/features: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server on the Remote Access server, and the default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface and Windows PowerShell cmdlets. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI • Remote Access module for Windows

Role/feature	How it supports this scenario
	PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

- Server requirements:
  - A computer that meets the hardware requirements for Windows Server 2012.
  - The server must have at least one network adapter installed, enabled, and connected to the internal network. When two adapters are used, there should be one adapter connected to the internal corporate network, and one connected to the external network (Internet, or private network).
  - If Teredo is required as an IPv4 to IPv6 transition protocol, the external adapter of the server requires two consecutive public IPv4 addresses. If a single IP address is available, then only IP-HTTPS can be used as the transition protocol.
  - At least one domain controller. The Remote Access server and DirectAccess clients must be domain members.
  - A certification authority (CA) is required if you do not want to use self-signed certificates for IP-HTTPS or the network location server, or if you want to use client certificates for client IPsec authentication. Alternatively, you can request certificates from a public CA.
  - If the network location server is not located on the Remote Access server, a separate web server is required to run it.
- Client requirements:
  - A client computer must be running Windows 8 or Windows 7.

    Note
    Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.
- Infrastructure and management server requirements:
  - During remote management of DirectAccess client computers, clients initiate communications with management servers such as domain controllers, System Center Configuration Servers, and Health Registration Authority (HRA) servers for services that include Windows and antivirus updates and Network Access Protection (NAP) client

compliance. The required servers should be deployed before beginning the Remote Access deployment.

- If Remote Access requires client NAP compliance, NPS and HRS servers should be deployed before beginning remote access deployment
- If VPN is enabled, a DHCP server is required to allocate IP addresses automatically to VPN clients, if a static address pool is not used.

Software requirements

There are a number of requirements for this scenario:

- Server requirements:
  - The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
  - If the Remote Access server is located behind an edge firewall or NAT device, the device must be configured to allow traffic to and from the Remote Access server.
  - The person deploying remote access on the server requires local administrator permissions on the server, and domain user permissions. In addition, the administrator requires permissions for the GPOs used in DirectAccess deployment. To take advantage of the features that restricts DirectAccess deployment to mobile computers only, permissions to create a WMI filter on the domain controller are required.
- Remote Access client requirements:
  - DirectAccess clients must be domain members. Domains containing clients can belong to the same forest as the Remote Access server, or have a two-way trust with the Remote Access server forest or domain.
  - An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients. If a security group is not specified when configuring DirectAccess client settings, by default the client GPO is applied on all laptop computers in the Domain Computers security group. Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

    Note
    It is recommended that you create a security group for each domain containing computers that will be configured as DirectAccess clients.

    Important
    If you have enabled Teredo in your DirectAccess deployment, and you want to provide access to Windows 7 clients, make sure that the clients are upgraded to Windows 7 with SP1. Clients using Windows 7 RTM will not be able to connect over Teredo. However, these clients will still be able to connect to the corporate network over IP-HTTPS.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
Deployment	Remote Access (DirectAccess, Routing and Remote Access) Overview
Tools and settings	Remote Access PowerShell cmdlets
Community resources	DirectAccess Wiki entries
Related technologies	How IPv6 works

Plan an Advanced Remote Access Deployment

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This topic lists the planning steps that are required to deploy a single Remote Access server running Windows Server 2012 with a full range of basic and advanced features. The planning phase includes the following steps:

- Step 1: Plan the Remote Access Infrastructure

  In this phase, you plan your network and server topology, firewall settings, certificate requirements, the Domain Name System (DNS), Active Directory, DirectAccess management servers, and the DirectAccess network location server.
- Step 2: Plan the Remote Access Deployment

  In this phase, you plan for the client and server deployment, including the DirectAccess infrastructure and the application servers.
- Step 3: Record the plan

  Record your planning decisions for your advanced Remote Access deployment. This record can be used as a job aid for everyone involved in completing the deployment steps.

See also

- Install and Configure Advanced Remote Access
- Deploy a Single Remote Access Server with Advanced Settings

Step 1: Plan the Remote Access Infrastructure

The first step of planning for an advanced Remote Access deployment on a single server is to plan the infrastructure that is required for the deployment. This topic describes the infrastructure planning steps. These planning tasks do not need to be completed in a specific order.

Task	Description
1.1 Plan network topology and settings	Decide where to place the Remote Access server (at the edge or behind a Network Address Translation (NAT) device or firewall), and plan IP addressing, routing, and force tunneling.
1.2 Plan firewall requirements	Plan for allowing Remote Access through edge firewalls.
1.3 Plan certificate requirements	Decide if you will use Kerberos protocol or certificates for client authentication, and plan your website certificates. IP-HTTPS is a transition protocol that is used by DirectAccess clients to tunnel IPv6 traffic over IPv4 networks. Decide whether to authenticate to the IP-HTTPS server by using a certificate that is issued by a certification authority (CA), or by using a self-signed certificate that is issued automatically by the Remote Access server.
1.4 Plan DNS requirements	Plan Domain Name System (DNS) settings for the Remote Access server, infrastructure servers, local name resolution options, and client connectivity.
1.5 Plan the network location server	Decide where to place the network location server website in your organization (on the Remote Access server or an alternative server), and plan the certificate requirements if the network location server will be located on

Task	Description
	the Remote Access server.
1.6 Plan management servers	Plan for management servers (such as update servers) that administrators can use to remotely manage DirectAccess client computers that are located outside the corporate network..
1.7 Plan Active Directory	Plan for your domain controllers, your Active Directory requirements, client authentication, and multiple domains.
1.8 Plan Group Policy Objects	Decide what GPOs are required in your organization and how to create or edit the GPOs.

1.1 Plan network topology and settings

This section explains how to plan for your network, including:

- 1.1.1 Plan network adapters and IP addressing
- Plan IPv6 intranet connectivity
- Plan force tunneling

1.1.1 Plan network adapters and IP addressing

1. Identify the network adapter topology you want to use. Remote Access can be set up by using either of the following topologies:
   - With two network adapters—The Remote Access server can be installed at the edge with one network adapter that is connected to the Internet and the other to the internal network, or it can be installed behind a NAT, firewall, or router device, with one network adapter connected to a perimeter network and the other to the internal network.
   - With one network adapter—The Remote Access server is installed behind a NAT device, and the single network adapter is connected to the internal network.
2. Identify your IP addressing requirements:

   DirectAccess uses IPv6 with IPsec to create a secure connection between DirectAccess client computers and the internal corporate network. However, DirectAccess does not necessarily require connectivity to the IPv6 Internet or native IPv6 support on internal networks. Instead, it automatically configures and uses IPv6 transition technologies to tunnel IPv6 traffic across the IPv4 Internet (by using 6to4, Teredo, or IP-HTTPS) and across your IPv4-only intranet (by using NAT64 or ISATAP). For an overview of these transition technologies, see the following resources:

   - IPv6 Transition Technologies
   - IP-HTTPS Tunneling Protocol Specification

3. Configure required adapters and addresses according to the following table. For deployments that use a single network adapter and are set up behind a NAT device, configure your IP addresses by using only the **Internal network adapter** column.

	External network adapter	Internal network adapter	Routing requirements
IPv4 Internet and IPv4 intranet	Configure the following: • Two static consecutive public IPv4 addresses with the appropriate subnet masks (required for Teredo only). • The default gateway IPv4 address of your Internet firewall or local Internet service provider (ISP) router. **Note** The Remote Access server requires two consecutive public IPv4 addresses so that it can act as a Teredo server and Windows-based clients can use the Remote Access server to detect the type of NAT device that they are behind.	Configure the following: • An IPv4 intranet address with the appropriate subnet mask. • The connection-specific DNS suffix of your intranet namespace. A DNS server should also be configured on the internal interface. **Caution** Do not configure a default gateway on any intranet interfaces.	To configure the Remote Access server to reach all subnets on the internal IPv4 network, do the following: • List the IPv4 address spaces for all the locations on your intranet. • Use the **route add -p** or the**netsh interface ipv4 add route** command to add the IPv4 address spaces as static routes in the IPv4 routing table of the Remote Access server.
IPv6 Internet and IPv6 intranet	Configure the following: • Use the address configuration that is provided by your ISP. • Use the **Route Print** command to ensure	Configure the following: If you are not using the default preference levels, you can configure your intranet interfaces by using the following command**netsh interface ipv6**	If you have an IPv6 intranet, to configure the Remote Access server to reach all of the IPv6

	External network adapter	Internal network adapter	Routing requirements
	that a default IPv6 route exists, and it is pointing to the ISP router in the IPv6 routing table. • Determine whether the ISP and intranet routers are using the default router preferences described in RFC 4191, and using a higher default preference than your local intranet routers. If both of these are true, no other configuration for the default route is required. The higher preference for the ISP router ensures that the active default IPv6 route of the Remote Access server points to the IPv6 Internet. Because the Remote Access server is an IPv6 router, if you have a native IPv6 infrastructure, the Internet interface can also reach the domain controllers on the intranet. In this case, add packet filters to the domain controller in the perimeter network that prevent connectivity to the IPv6 address of the Internet-facing interface of the	**set InterfaceIndex ignoredefaultroutes=enabled**. This command ensures that additional default routes that point to intranet routers will not be added to the IPv6 routing table. You can obtain the interface index of your intranet interfaces by using the following command: **netsh interface ipv6 show**.	locations, do the following: • List the IPv6 address spaces for all the locations on your intranet. • Use the **netsh interface ipv6 add route** command to add the IPv6 address spaces as static routes in the IPv6 routing table of the Remote Access server.

	External network adapter	Internal network adapter	Routing requirements
	Remote Access server.		
IPv4 Internet and IPv6 intranet	The Remote Access server forwards default IPv6 route traffic through the Microsoft 6to4 adapter to a 6to4 relay on the IPv4 Internet. You can configure a Remote Access server for the IPv4 address of the Microsoft 6to4 adapter by using the following command: **netsh interface ipv6 6to4 set relay name=<ipaddress> state=enabled**.		

Notes

- If the DirectAccess client has been assigned a public IPv4 address, it will use the 6to4 transition technology to connect to the intranet. If it is assigned a private IPv4 address, it will use Teredo. If the DirectAccess client cannot connect to the DirectAccess server with 6to4 or Teredo, it will use IP-HTTPS.
- To use Teredo, you must configure two consecutive IP addresses on the external facing network adapter.
- You cannot use Teredo if the Remote Access server has only one network adapter.
- Native IPv6 client computers can connect to the Remote Access server over native IPv6, and no transition technology is required.

1.1.2 Plan IPv6 intranet connectivity

ISATAP is required to use “manage out” DirectAccess capabilities. This allows DirectAccess management servers to connect to DirectAccess clients that are located on the Internet for the purpose of remote management. ISATAP is not required to support connections tha are initiated by DirectAccess client computers to IPv4 resources on the corporate network. NAT64/DNS64 is used for this purpose. If your deployment requires ISATAP, use the following table to identify your requirements.

Notes

- If you are not managing out to DirectAccess clients, ISATAP is not required.
- ISATAP is not supported in external load-balanced clusters.

- In a multisite deployment, the ISATAP DNS name should be registered with the internal IPv4 address of all entry points.
- In a cluster deployment, the ISATAP DNS name should be registered with the internal virtual IPv4 address and the dedicated IPv4 addresses of all cluster nodes.
- When you use ISATAP, enable IPv6 host (AAAA) resource record queries for DNS64 with the following Windows PowerShell command: **Set-NetDnsTransitionConfiguration -OnlySendAQuery $false**.

ISATAP deployment scenario	Requirements
Existing native IPv6 intranet (no ISATAP is required)	With an existing native IPv6 infrastructure, you specify the prefix of the organization during Remote Access deployment, and the Remote Access server does not configure itself as an ISATAP router. Do the following: • To ensure that DirectAccess clients are reachable from the intranet, you must modify your IPv6 routing so that default route traffic is forwarded to the Remote Access server. If your intranet IPv6 address space uses something other than a single IPv6 address prefix, you must specify the relevant organization IPv6 prefix during deployment. • If you are currently connected to the IPv6 Internet, you must configure your default route traffic so that it is forwarded to the Remote Access server, and then configure the appropriate connections and routes on the Remote Access server, so that the default route traffic is forwarded to the device that is connected to the IPv6 Internet.
Existing ISATAP deployment	If you have an existing ISATAP infrastructure with native IPv6 connectivity, during deployment you are prompted for the 48-bit prefix of the organization, and the Remote Access server does not configure itself as an ISATAP router. To ensure that DirectAccess clients are reachable from the intranet, you must modify your IPv6 routing infrastructure so that default route traffic is forwarded to the Remote Access server. This change needs to be done on the

ISATAP deployment scenario	Requirements
	existing ISATAP router to which the intranet clients already forward the default traffic. If there is no native IPv6 connectivity between the Remote Access server and the ISATAP router, the Remote Access server will act as an ISATAP host. In this case, the Remote Access server's internal IPv4 address needs to be added to the potential ISATAP router list by adding it to DNS along with the existing ISATAP router DNS entry.
No existing IPv6 connectivity	When the Remote Access Setup Wizard detects that the server has no native or ISATAP-based IPv6 connectivity, it automatically derives a 6to4-based 48-bit prefix for the intranet (a 6to4-based prefix is generated only if the server has a public IPv4 address, otherwise a ULA prefix is automatically generated). The wizard also configures the Remote Access server as an ISATAP router to provide IPv6 connectivity to ISATAP hosts across your intranet. To use ISATAP do the following: 1. Register the name ISATAP on a domain DNS server for each domain on which you want to enable ISATAP based connectivity, so that the ISATAP name is resolvable by the internal DNS server to the internal IPv4 address of the Remote Access server. 2. By default, DNS servers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 block the resolution of the ISATAP name with the global query block list. To enable ISATAP, you must remove the ISATAP name from the block list. For more information, see Remove ISATAP from the DNS Global Query Block List. Windows-based ISATAP hosts that can resolve the ISATAP name automatically configure addresses with the Remote Access server, which results in automatic configuration of the

ISATAP deployment scenario	Requirements
	following: • An ISATAP-based IPv6 address on an ISATAP tunneling interface. • A 64-bit route that provides connectivity to the other ISATAP hosts on the intranet. • A default IPv6 route that points to the Remote Access server. The default route ensures that intranet ISATAP hosts can reach DirectAccess clients. When your Windows-based ISATAP hosts obtain an ISATAP-based IPv6 address, they begin to use ISATAP-encapsulated traffic to communicate, if the destination is also an ISATAP host. Because ISATAP uses a single 64-bit subnet for the entire intranet, your communication goes from a segmented, multisubnet IPv4 model of communication, to a flat, single subnet communication model with IPv6. This can affect the way that some Active Directory Domain Services (AD DS), and other applications that rely on your Active Directory Sites and Services configuration, behave. For example, if you used the Active Directory Sites and Services snap-in to configure sites, IPv4-based subnets, and intersite transports for forwarding requests to servers within sites, this configuration is not used by ISATAP hosts. • To configure Active Directory sites and services for forwarding within sites for ISATAP hosts, for each IPv4 subnet object, you must configure an equivalent IPv6 subnet object, in which the IPv6 address prefix for the subnet expresses the same range of ISATAP host addresses as the IPv4 subnet. For example, for the IPv4 subnet 192.168.99.0/24 and the 64-bit ISATAP address prefix 2002:836b:1:8000::/64, the equivalent IPv6 address prefix for the IPv6 subnet object is 2002:836b:1:8000:0:5efe:192.168.99.0/120. For an arbitrary IPv4 prefix length (set to 24 in the example), you can determine the

ISATAP deployment scenario	Requirements

	corresponding IPv6 prefix length from the following formula: 96 + IPv4PrefixLength.
	• For the IPv6 addresses of DirectAccess clients, add the following:
	  • For Teredo-based DirectAccess clients: An IPv6 subnet for the range 2001:0:WWXX:YYZZ::/64, in which WWXX:YYZZ is the colon-hexadecimal version of the selected first Internet-facing IPv4 address of the Remote Access server.
	  • For IP-HTTPS-based DirectAccess clients: An IPv6 subnet for the range 2002:WWXX:YYZZ:1000::/56, in which WWXX:YYZZ is the colon-hexadecimal version of the selected first Internet-facing IPv4 address (w.x.y.z) of the Remote Access server.
	  • For 6to4-based DirectAccess clients: A series of 6to4-based IPv6 prefixes that begin with 2002: and represent the regional public IPv4 address prefixes that are administered by Internet Assigned Numbers Authority (IANA) and regional registries. The 6to4-based prefix for a public IPv4 address prefix w.x.y.z/n is 2002:WWXX:YYZZ::/[16+n], in which WWXX:YYZZ is the colon-hexadecimal version of w.x.y.z.
	  For example, the 7.0.0.0/8 range is administered by the American Registry for Internet Numbers (ARIN) for North America. The corresponding 6to4-based prefix for this public IPv6 address range is 2002:700::/24. For information about the IPv4 public address space, see IANA IPv4 Address Space Registry.

Important

If you are deploying ISATAP in your internal network, ensure that you do not have public IP addresses on the internal interface of the DirectAccess server. If you have public IP addresses on the internal interface, connectivity through ISATAP may fail.

1.1.3 Plan for force tunneling

With IPv6 and the Name Resolution Policy Table (NRPT), by default, DirectAccess clients separate their intranet and Internet traffic as follows:

- DNS name queries for intranet fully qualified domain names (FQDNs) and all intranet traffic is exchanged over the tunnels that are created with the DirectAccess server or directly with intranet servers. Intranet traffic from DirectAccess clients is IPv6 traffic.
- DNS name queries for FQDNs that correspond to exemption rules or do not match the intranet namespace, and all traffic to Internet servers, is exchanged over the physical interface that is connected to the Internet. Internet traffic from DirectAccess clients is typically IPv4 traffic.

In contrast, by default, some remote access virtual private network (VPN) implementations, including the VPN client, send all intranet and Internet traffic over the remote access VPN connection. Internet-bound traffic is routed by the VPN server to intranet IPv4 web proxy servers for access to IPv4 Internet resources. It is possible to separate the intranet and Internet traffic for remote access VPN clients by using split tunneling. This involves configuring the Internet Protocol (IP) routing table on VPN clients so that traffic to intranet locations is sent over the VPN connection, and traffic to all other locations is sent by using the physical interface that is connected to the Internet.

You can configure DirectAccess clients to send all of their traffic through the tunnels to the DirectAccess server with force tunneling. When force tunneling is configured, DirectAccess clients detect that they are on the Internet, and they remove their IPv4 default route. With the exception of local subnet traffic, all traffic sent by the DirectAccess client is IPv6 traffic that goes through tunnels to the DirectAccess server.

Important

If you plan to use force tunneling, or you might add it in the future, you should deploy a two tunnel configuration. Because of security considerations, force tunneling in a single tunnel configuration is not supported.

Enabling force tunneling has the following consequences:

- DirectAccess clients use only IP-HTTPS to obtain IPv6 connectivity to the DirectAccess server over the IPv4 Internet.
- The only locations that a DirectAccess client can reach by default with IPv4 traffic are those on its local subnet. All other traffic that is sent by applications and services running on the DirectAccess client is IPv6 traffic, which is sent over the DirectAccess connection. Therefore, IPv4-only applications on the DirectAccess client cannot be used to reach Internet resources, except those on the local subnet.

Important

For force tunneling through DNS64 and NAT64, IPv6 Internet connectivity must be implemented. One way to achieve this is by making the IP-HTTPS prefix globally routable, so that ipv6.msftncsi.com will be reachable over IPv6, and the response from the Internet server to the IP-HTTPS clients is able to return through the Remote Access server.

Because this is not feasible in most cases, the best option is to create virtual NCSI servers inside the corporate network as follows:

1. Add an NRPT entry for ipv6.msftncsi.com and resolve it against DNS64 to an internal website (this can be IPv4 website).
2. Add an NRPT entry for dns.msftncsi.com and resolve it against a corporate DNS server to return the IPv6 host (AAAA) resource record fd3e:4f5a:5b81::1. (Using DNS64 to only send host (A) resource record queries for this FQDN may not work because it is configured in IPv4 only deployments, so you should configure it to resolve against corporate DNS directly.)

1.2 Plan firewall requirements

If the Remote Access server is behind an edge firewall, the following exceptions are required for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- Teredo traffic—User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound.
- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound.
- If you are deploying Remote Access with a single network adapter, and installing the network location server on the Remote Access server, TCP port 62000 should also be exempted.

  Note
  This exemption is on the Remote Access server, and all other exemptions are on the edge firewall.

For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server. For IP-HTTPS, the exceptions need to be applied on the address that is registered on the public DNS server.

The following exceptions are required for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol ID 50
- UDP destination port 500 inbound, and UDP source port 500 outbound
- ICMPv6 traffic inbound and outbound (when using Teredo only)

When you use additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4 and IPv6 traffic
- ICMP for all IPv4 and IPv6 traffic (when using Teredo only)

1.3 Plan certificate requirements

There are three scenarios that require certificates when you deploy a single Remote Access server:

- Plan computer certificates for IPsec authentication

  Certificate requirements for IPsec include a computer certificate that is used by DirectAccess client computers when they establish the IPsec connection between the client and the Remote Access server, and a computer certificate that is used by Remote Access servers to establish IPsec connections with DirectAccess clients.

  For DirectAccess in Windows Server 2012 the use of these IPsec certificates is not mandatory. As an alternative the Remote Access server can act as a Kerberos proxy to perform IPsec authentication without requiring certificates. If Kerberos protocol is used, it works over SSL, and the Kerberos proxy uses the certificate that is configured for IP-HTTPS for this purpose. Some enterprise scenarios (including multisite deployment and one-time password (OTP) client authentication) require the use of certificate authentication, and not the Kerberos protocol.

- Plan computer certificates for IP-HTTPS

  When you configure Remote Access, the Remote Access server is automatically configured to act as the IP-HTTPS listener. The IP-HTTPS site requires a website certificate, and client computers must be able to contact the certificate revocation list (CRL) site for the certificate.

- Plan website certificates for the network location server

  The network location server is a website that is used to detect whether client computers are located in the corporate network. The network location server requires a website certificate. DirectAccess clients must be able to contact the CRL site for the certificate.

The certification authority (CA) requirements for each scenario are summarized in the following table.

IPsec authentication	IP-HTTPS server	Network location server
An internal CA is required to issue computer certificates to the Remote Access server and clients for IPsec authentication when you don't use the Kerberos proxy for authentication	Internal CA: You can use an internal CA to issue the IP-HTTPS certificate; however, you must make sure that the CRL distribution point is available externally.	Internal CA: You can use an internal CA to issue the network location server website certificate. Make sure that the CRL distribution point has high availability from the internal network.
	Self-signed certificate: You can use a self-signed certificate for the IP-HTTPS server; however, you must	Self-signed certificate: You can use a self-signed certificate for the network location server website.

IPsec authentication	IP-HTTPS server	Network location server
	make sure that the CRL distribution point is available externally. A self-signed certificate cannot be used in multisite deployments.	A self-signed certificate cannot be used in multisite deployments.
	Recommended Public CA: It is recommended to use a public CA to issue the IP-HTTPS certificate. This ensures that the CRL distribution point is available externally.	

1.3.1 Plan computer certificates for IPsec authentication

If you are using certificate-based IPsec authentication, the Remote Access server and clients are required to obtain a computer certificate. The simplest way to install the certificates it to configure Group Policy-based automatic enrollment for computer certificates. This ensures that all domain members obtain a certificate from an enterprise CA. If you do not have an enterprise CA set up in your organization, see Active Directory Certificate Services.

This certificate has the following requirements:

- The certificate should have a client authentication extended key usage (EKU).
- The client certificate and the server certificate should chain to the same root certificate. This root certificate must be selected in the DirectAccess configuration settings.

1.3.2 Plan certificates for IP-HTTPS

The Remote Access server acts as an IP-HTTPS listener, and you must manually install an HTTPS website certificate on the server. Consider the following when you are planning:

- Using a public CA is recommended, so that CRLs are readily available.
- In the **Subject** field, specify the IPv4 address of the Internet adapter of Remote Access server or the FQDN of the IP-HTTPS URL (the ConnectTo address). If the Remote Access server is located behind a NAT device, the public name or address of the NAT device should be specified.
- The common name of the certificate should match the name of the IP-HTTPS site.
- For the **Enhanced Key Usage** field, use the server authentication object identifier (OID).
- For the **CRL Distribution Points** field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.

- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcard characters in the name.

If you plan to use IP-HTTPS on a non-standard port, perform the following steps on the Remote Access server:

1. Remove the existing certificate binding for 0.0.0.0:443, and replace it with a certificate binding for your chosen port. For purposes of this example, port 44500 is used. Prior to deleting the certificate binding, show and copy the **appid**.
   a. To delete the certificate binding, enter:

      ```
      netsh http delete ssl ipport=0.0.0.0:443
      ```

   b. To add the new certificate binding, enter:

      ```
      netsh http add ssl ipport=0.0.0.0:44500 certhash=<use the
      thumbprint from the Remote Access server SSL cert> appid=<use the
      appid from the binding that was deleted>
      ```

2. To modify the IP-HTTPS URL on the server, enter:

   ```
   Netsh int http set int url=https://<Remote Access server name
   (for example server.contoso.com)>:44500/IPHTTPS
   Net stop iphlpsvc & net start iphlpsvc
   ```

3. Change the URL reservation for kdcproxy.
   a. To delete the existing URL reservation, enter:

      ```
      netsh http del urlacl url=https://+:443/KdcProxy/
      ```

   b. To add a new URL reservation, enter:

      ```
      netsh http add urlacl
      url=https://+:44500/KdcProxy/ sddl=D:(A;;GX;;;NS)
      ```

4. Add the setting to make kppsvc listen on the non-standard port. To add the registry entry, enter:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\KPSSVC\Sett
   ings /v HttpsUrlGroup /t REG_MULTI_SZ /d +:44500 /f
   ```

5. To restart the kdcproxy service on the domain controller, enter:

   ```
   net stop kpssvc & net start kpssvc
   ```

To use IP-HTTPS on a non-standard port, perform the following steps on the domain controller:

1. Modify the IP-HTTPS setting in the client GPO.
   a. Open the Group Policy editor.
   b. Navigate to Computer configuration=>Policies=>Administrative Templates=> Network=>TCPIP settings =>IPv6 transition technologies.
   c. Open the IP-HTTPS state setting and change the URL to **https://<Remote Access server name (for example server.contoso.com)>:44500/IPHTTPS**.
   d. Click **Apply**.

2. Modify the Kerberos proxy client settings in the client GPO.
    a. In the Group Policy editor, navigate to Computer configuration=>Policies=>Administrative Templates=> System=>Kerberos => Specify the KDC proxy servers for Kerberos clients.
    b. Open the IPHTTPS state setting and change the URL to **https://<Remote Access server name (for example server.contoso.com)>:44500/IPHTTPS**.
    c. Click **Apply**.
3. Modify the client IPsec policy settings to use ComputerKerb and UserKerb.
    a. In the Group Policy editor, navigate to Computer configuration=>Policies=> Windows Settings=> Security Settings=> Windows Firewall with Advanced security.
    b. Click **Connection security rules**, and then double-click **IPsec rule**.
    c. On the **Authentication** tab, click **Advanced**.
    d. For Auth1: rRmove the existing authentication method, and replace it with ComputerKerb. For Auth2: Remove the existing authentication method, and replace it with UserKerb.
    e. Click **Apply**, and then **OK**.

To complete the manual process for using an IP-HTTPS non-standard port, run **gpupdate /force** on the client computers and the Remote Access server.

1.3.3 Plan website certificates for the network location server

When you plan for the network location server website, consider the following:

- In the **Subject** field, specify the IP address of the intranet interface of the network location server or the FQDN of the network location URL.
- In the **Enhanced Key Usage** field, use the Server Authentication OID.
- In the **CRL Distribution Points** field, use a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet. This CRL distribution point should not be accessible from outside the internal network.
- If you are later planning to configure a multisite or cluster deployment, the name of the certificate should not match the internal name of any Remote Access server that will be added to the deployment.

> **Note**
> Ensure that the certificates for IP-HTTPS and the network location server have a **Subject Name**. If the certificate does not have a **Subject Name**, but it has an **Alternative Name**, it will not be accepted by the Remote Access Wizard.

1.4 Plan DNS requirements

This section explains the DNS requirements for DirectAccess client requests and infrastructure servers in a Remote Access deployment. It includes the following subsections:

- Plan for DNS server requirements
- Plan for DNS local name resolution

DirectAccess client requests

DNS is used to resolve requests from DirectAccess client computers that are not located on the internal (or corporate) network. DirectAccess clients attempt to connect to the DirectAccess network location server to determine whether they are located on the Internet, or on the internal network.

- If the connection is successful, clients are identified as being located on the internal network, DirectAccess is not used, and client requests are resolved by using the DNS server that is configured on the network adapter of the client computer.
- If the connection does not succeed, clients are assumed to be located on the Internet, and DirectAccess clients will use the name resolution policy table (NRPT) to determine which DNS server to use when resolving name requests.

You can specify that clients use DirectAccess DNS64 to resolve names, or an alternative internal DNS server. When performing name resolution, the NRPT is used by DirectAccess clients to identify how to handle a request. Clients request an FQDN or single-label name such as http://internal. If a single-label name is requested, a DNS suffix is appended to make an FQDN. If the DNS query matches an entry in the NRPT, and DNS64 or a DNS server on the internal network is specified for the entry, the query is sent for name resolution using the specified server. If a match exists, but no DNS server is specified, this indicates an exemption rule, and normal name resolution is applied.

Note

Note that when a new suffix is added to the NRPT in the Remote Access Management Console, the default DNS servers for the suffix can be automatically discovered by clicking **Detect**.

Auto detection works as follows:

- If the corporate network is IPv4-based, or it uses IPv4 and IPv6, the default address is the DNS64 address of the internal adapter on the Remote Access server.
- If the corporate network is IPv6-based, the default address is the IPv6 address of DNS servers in the corporate network.

Infrastructure servers

- **Network location server**

  DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, but they must be prevented from resolving the name when they are located on the Internet. To ensure that this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT. In addition, when you configure Remote Access, the following rules are created automatically:

  - A DNS suffix rule for the root domain or the domain name of the Remote Access server, and the IPv6 addresses that correspond to the DNS64 address. In IPv6-only corporate networks, the intranet DNS servers are configured on the Remote Access server. For example, if the Remote Access server is a member of the corp.contoso.com domain, a rule is created for the corp.contoso.com DNS suffix.

- An exemption rule for the FQDN of the network location server. For example, if the network location server URL is https://nls.corp.contoso.com, an exemption rule is created for the FQDN nls.corp.contoso.com.

- **IP-HTTPS server**

  The Remote Access server acts as an IP-HTTPS listener, and it uses its server certificate to authenticate to IP-HTTPS clients. The IP-HTTPS name must be resolvable by DirectAccess clients that are using public DNS servers.

- **CRL revocation checking**

  DirectAccess uses certificate revocation checking for the IP-HTTPS connection between DirectAccess clients and the DirectAccess server, and for the HTTPS-based connection between the DirectAccess client and the network location server. In both cases, DirectAccess clients must be able to resolve and access the CRL distribution point location.

- **ISATAP**

  ISATAP enables corporate computers to acquire an IPv6 address, and it encapsulates IPv6 packets within an IPv4 header. It is used by the Remote Access server to provide IPv6 connectivity to ISATAP hosts across an intranet. In a non-native IPv6 network environment, the Remote Access server configures itself automatically as an ISATAP router.

  DNS servers must be configured to answer ISATAP queries. By default, the DNS server service blocks name resolution for the ISATAP name through the DNS Global Query Block List. If you are using ISATAP on your intranet, you must remove the ISATAP name from the query block list, and then manually create a host (A) resource record for the ISATAP name in the appropriate DNS server. For example, for the corp.contoso.com domain, create a host (A) resource record for isatap.corp.contoso.com, with the internal IPv4 address of the Remote Access server.

- **Connectivity verifiers**

  Remote Access creates a default web probe that is used by DirectAccess client computers to verify connectivity to the internal network. To ensure that the probe works as expected, the following names must be registered manually in DNS:

  - **directaccess-webprobehost**—Should resolve to the internal IPv4 address of the Remote Access server, or to the IPv6 address in an IPv6-only environment.
  - **directaccess-corpconnectivityhost**—Should resolve to the local host (loopback) address. The following host (A) and (AAAA) resource records should be created: a host (A) resource record with value 127.0.0.1, and a host (AAAA) resource record with value constructed out of NAT64 prefix with the last 32 bits as 127.0.0.1. The NAT64 prefix can be retrieved by running the Windows PowerShell command **get-netnattransitionconfiguration**.

    Note

    This is valid only in an IPv4-only environment. In an IPv4 plus IPv6, or an IPv6-only environment, only a host (AAAA) resource record should be created with the loopback IP address ::1.

You can create additional connectivity verifiers by using other web addresses over HTTP or by using **ping**. For each connectivity verifier, a DNS entry must exist.

1.4.1 Plan for DNS server requirements

- For DirectAccess clients, you must use a DNS server running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, or any DNS server that supports IPv6.
- We recommend that you use a DNS server that supports dynamic updates. You can use DNS servers that do not support dynamic updates, but you must manually update entries.
- The FQDN for your Internet-accessible CRL distribution points must be resolvable by using Internet DNS servers. For example, if URL http://crl.contoso.com/crld/corp-DC1-CA.crl is in the **CRL Distribution Points** field of the IP-HTTPS certificate of the Remote Access server, you must ensure that the FQDN crld.contoso.com is resolvable by using Internet DNS servers.

1.4.2 Plan for local name resolution

When you plan for local name resolution, consider the following issues:

- **NRPT**

  You may need to create additional NRPT rules in the following cases:

  - If you need to add more DNS suffixes for your intranet namespace.
  - If the FQDN of your CRL distribution points are based on your intranet namespace, you must add exemption rules for the FQDNs of the CRL distribution points.
  - If you have a split-brain DNS environment, you must add exemption rules for the names of resources for which you want DirectAccess clients located on the Internet to access the Internet version, rather than the intranet version.
  - If you are redirecting traffic to an external website through your intranet web proxy servers, the external website is available only from the intranet, and it uses the addresses of your web proxy servers to permit the inbound requests. In this case, add an exemption rule for the FQDN of the external website, and specify that the rule use your intranet web proxy server rather than the IPv6 addresses of intranet DNS servers.

    For example, if you are testing an external website named test.contoso.com, this name is not resolvable through Internet DNS servers, but the Contoso web proxy server knows how to resolve the name and to direct requests for the website to the external web server. To prevent users who are not on the Contoso intranet from accessing the site, the external website allows requests only from the IPv4 Internet address of the Contoso web proxy. Thus, intranet users can access the website because they are using the Contoso web proxy, but DirectAccess users cannot access it because they are not using the Contoso web proxy. By configuring an NRPT exemption rule for test.contoso.com that uses the Contoso web proxy, webpage requests for test.contoso.com are routed to the intranet web proxy server over the IPv4 Internet.
- **Single label names**

Single label names, such as http://paycheck, are sometimes used for intranet servers. If a single label name is requested, and a DNS suffix search list is configured, the DNS suffixes in the list will be appended to the single label name. For example, when a user on a computer that is a member of the corp.contoso.com domain types http://paycheck in the web browser, the FQDN that is constructed as the name is paycheck.corp.contoso.com. By default the appended suffix is based on the primary DNS suffix of the client computer.

Note

In a disjoint name space scenario (where one or more domain computers has a DNS suffix that does not match the Active Directory domain to which the computers belong), you should ensure that the search list is customized to include all the required suffixes. By default, the Remote Access Wizard will configure the Active Directory DNS name as the primary DNS suffix on the client. Make sure to add the DNS suffix that is used by clients for name resolution.

If multiple domains and Windows Internet Name Service (WINS) are deployed in your organization and you are connecting remotely, single-names can be resolved as follows:

- Deploy a WINS forward lookup zone in the DNS. When trying to resolve computername.dns.zone1.corp.contoso.com, the request is directed to the WINS server that is using only computername. The client thinks it is issuing a regular DNS host (A) resource record, but it is actually a NetBIOS request. For more information, see Managing a Forward Lookup Zone.
- Add a DNS suffix, for example dns.zone1.corp.contoso.com, to the default domain policy GPO.

- **Split-brain DNS**

  Split-brain DNS is the use of the same DNS domain for Internet and intranet name resolution.

  For split-brain DNS deployments, you must list the FQDNs that are duplicated on the Internet and intranet, and decide which resources the DirectAccess client should reach—the intranet or the Internet version. For each name that corresponds to a resource for which you want DirectAccess clients to reach the Internet version, you must add the corresponding FQDN as an exemption rule to the NRPT for your DirectAccess clients.

  In a split-brain DNS environment, if you want both versions of the resource to be available, configure your intranet resources with alternate names, that are not duplicates of the names used on the Internet, and instruct your users to use the alternate name when on the Intranet. For example, configure the alternate name www.internal.contoso.com for the internal name www.contoso.com.

  In a non-split-brain DNS environment, the Internet namespace is different from the intranet namespace. For example, the Contoso Corporation uses contoso.com on the Internet and corp.contoso.com on the intranet. Because all intranet resources use the corp.contoso.com DNS suffix, the NRPT rule for corp.contoso.com routes all DNS name queries for intranet resources to intranet DNS servers. DNS name queries for names with the contoso.com suffix do not match the corp.contoso.com intranet namespace rule in the NRPT, and are sent to Internet DNS servers. With a non-split-brain DNS deployment, because there is no

duplication of FQDNs for intranet and Internet resources, there is no additional configuration needed for the NRPT. DirectAccess clients can access both Internet and intranet resources for the organization.

- **Local name resolution behavior for DirectAccess clients**

  If a name cannot be resolved with DNS, to resolve the name on the local subnet, the DNS Client service in Windows Server 2012 , Windows Server 2008 R2, Windows 8, and Windows 7 can use local name resolution, with the Link-Local Multicast Name Resolution (LLMNR) and NetBIOS over TCP/IP protocols.

  Local name resolution is typically needed for peer-to-peer connectivity when the computer is located on private networks, such as single subnet home networks. When the DNS Client service performs local name resolution for intranet server names and the computer is connected to a shared subnet on the Internet, malicious users can capture LLMNR and NetBIOS over TCP/IP messages to determine intranet server names. On the DNS page of the Infrastructure Server Setup Wizard, you configure the local name resolution behavior based on the types of responses received from intranet DNS servers. The following options are available:

  - **Use local name resolution if the name does not exist in DNS**—This option is the most secure because the DirectAccess client performs local name resolution only for server names that cannot be resolved by intranet DNS servers. If the intranet DNS servers can be reached, the names of intranet servers are resolved. If the intranet DNS servers cannot be reached, or if there are other types of DNS errors, the intranet server names are not leaked to the subnet through local name resolution.
  - **Use local name resolution if the name does not exist in DNS or DNS servers are unreachable when the client computer is on a private network (recommended)**—This option is recommended because it allows the use of local name resolution on a private network only when the intranet DNS servers are unreachable.
  - **Use local name resolution for any kind of DNS resolution error (least secure)**—This is the least secure option because the names of intranet network servers can be leaked to the local subnet through local name resolution.

1.5 Plan the network location server

The network location server is a website that is used to detect whether DirectAccess clients are located in the corporate network. Clients in the corporate network do not use DirectAccess to reach internal resources, but instead, they connect directly.

The network location server website can be hosted on the Remote Access server or on another server in your organization. If you host the network location server on the Remote Access server, the website is created automatically when you deploy Remote Access. If you host the network location server on another server in your organization running a Windows operating system, you must make sure that Internet Information Services (IIS) is installed on that server, and that the website is created. Remote Access does not configure settings on the remote network location server.

Make sure that the network location server website meets the following requirements:

- It is a website with an HTTPS server certificate.
- DirectAccess client computers must trust the CA that issued the server certificate to the network location server website.
- DirectAccess client computers on the internal network must be able to resolve the name of the network location server site.
- The network location server site must have high availability to computers on the internal network.
- The network location server must not be accessible to DirectAccess client computers on the Internet.
- The server certificate must be checked against a Certificate Revocation List (CRL).

1.5.1 Plan certificates for the network location server

When you are obtaining the website certificate to use for the network location server, consider the following:

1. In the **Subject** field, specify an IP address of the intranet interface of the network location server or the FQDN of the network location URL.
2. In the **Enhanced Key Usage** field, use the Server Authentication OID.
3. In the **CRL Distribution Points** field, use a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet. This CRL distribution point should not be accessible from outside the internal network.

1.5.2 Plan DNS for the network location server

DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, but they must be prevented from resolving the name when they are located on the Internet. To ensure that this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT.

1.6 Plan management servers

DirectAccess clients initiate communications with management servers that provide services such as Windows Update and antivirus updates. DirectAccess clients also use the Kerberos protocol to contact domain controllers to authenticate before they access the internal network. During remote management of DirectAccess clients, management servers communicate with client computers to perform management functions such as software or hardware inventory assessments. Remote Access can automatically discover some management servers, including:

- Domain controllers—Remote Access automatically discovers domain controllers for all domains in the same forest as the Remote Access server and client computers.
- System Center Configuration Manager servers—Remote Access automatically discovers System Center Configuration Manager servers for all domains in the same forest as the Remote Access server and client computers.

Domain controllers and System Center Configuration Manager servers are automatically detected the first time that DirectAccess is configured. Detected domain controllers are not displayed in the console, but settings can be retrieved by using the Windows PowerShell cmdlet **Get-DAMgmtServer –Type All**. If domain controller or System Center Configuration Manager servers are modified, clicking **Refresh Management Servers** in the Remote Access Management console refreshes the management server list.

Management server requirements

- Management servers must be accessible over the first (infrastructure) tunnel. When you configure Remote Access, adding servers to the management servers list automatically makes them accessible over this tunnel.
- Management servers that initiate connections to DirectAccess clients must fully support IPv6, by means of a native IPv6 address or by using one that is assigned by ISATAP.

1.7 Plan Active Directory

This section explains how Remote Access uses Active Directory, and it includes the following subsections:

- Plan client authentication
- Plan multiple domains

Remote Access uses Active Directory and Active Directory Group Policy Objects (GPOs) as follows:

- **Authentication**

  Active Directory is used for authentication. The infrastructure tunnel uses NTLMv2 authentication for the computer account that is connecting to the Remote Access server, and the account must be listed in an Active Directory domain. The intranet tunnel uses Kerberos authentication for the user to create the second tunnel.

- **Group Policy Objects**

  Remote Access gathers configuration settings into GPOs that are applied to Remote Access servers, clients, and internal application servers.

- **Security groups**

  Remote Access uses security groups to gather and identify DirectAccess client computers. The GPOs are applied to the required security group.

- **Extended IPsec policies**

  Remote Access can use IPsec authentication and encryption between clients and the Remote Access server. You can extend IPsec authentication and encryption from the client to the specified internal application servers. To do this, add the required application servers into a security group.

Active Directory requirements

When you plan Active Directory for a Remote Access deployment, consider the following requirements:

- At least one domain controller must be installed with the Windows Server 2012, Windows Server 2008 R2 Windows Server 2008, or Windows Server 2003 operating system.

  If the domain controller is on a perimeter network (and therefore reachable from the Internet-facing network adapter of Remote Access server), you must prevent the Remote Access server from reaching it by adding packet filters on the domain controller to prevent connectivity to the IP address of the Internet adapter.
- The Remote Access server must be a domain member.
- DirectAccess clients must be domain members. Clients can belong to:
  - Any domain in the same forest as the Remote Access server.
  - Any domain that has a two-way trust with the Remote Access server domain.
  - Any domain in a forest that has a two-way trust with the forest to which the Remote Access domain belongs.

Notes

- The Remote Access server cannot be a domain controller.
- The Active Directory domain controller that is used for Remote Access must not be reachable from the external Internet adapter of the Remote Access server (that is, the adapter must not be in the domain profile of Windows Firewall).

1.7.1 Plan client authentication

Remote Access in Windows Server 2012 allows you to choose between using certificates for IPsec computer authentication or using a built-in Kerberos proxy that authenticates by using user names and passwords.

When choosing to use Active Directory credentials for authentication, DirectAccess uses one security tunnel that uses Computer Kerberos for the first authentication and User Kerberos for the second authentication. When using this mode for authentication, DirectAccess uses a single security tunnel that provides access to the DNS server, the domain controller, and to other servers on the internal network.

When choosing to use two-factor authentication or Network Access Protection, DirectAccess uses two security tunnels. The Remote Access Setup Wizard configures Windows Firewall with Advanced Security connection security rules that specify the use of the following types of credentials when negotiating the IPsec security associations for the tunnels to the Remote Access server:

- The infrastructure tunnel uses Computer Kerberos credentials for the first authentication and User Kerberos for the second authentication.
- The intranet tunnel uses computer certificate credentials for the first authentication and User Kerberos for the second authentication.

When DirectAccess is choosing to allow access to clients running Windows 7 or in a multisite deployment, it uses two security tunnels. The Remote Access Setup Wizard configures Windows Firewall with Advanced Security connection security rules that specify the use of the following

types of credentials when negotiating the IPsec security associations for the tunnels to the Remote Access server:

- The infrastructure tunnel uses computer certificate credentials for the first authentication and NTLMv2 for the second authentication. NTLMv2 credentials force the use of Authenticated Internet Protocol (AuthIP), and they provide access to a DNS server and domain controller before the DirectAccess client can use Kerberos credentials for the intranet tunnel.
- The intranet tunnel uses computer certificate credentials for the first authentication and User Kerberos for the second authentication.

1.7.2 Plan multiple domains

The management servers list should include domain controllers from all domains that contain security groups that include DirectAccess client computers. It should contain all domains that contain user accounts that might use computers that are configured as DirectAccess clients. This ensures that users who are not located in the same domain as the client computer they are using are authenticated with a domain controller in the user domain. This is done automatically if domains are in the same forest.

Note

If there are computers in the security groups that are used for client computers or application servers in different forests, the domain controllers of those forests are not detected automatically. You can run the task **Refresh Management Servers** in the Remote Access Management console to detect these domain controllers.

Where possible, common domain name suffixes should be added to the Name Resolution Policy Table (NRPT) during the Remote Access deployment. For example, if you have two domains, domain1.corp.contoso.com and domain2.corp.contoso.com, instead of adding two entries into the NRPT, you can add a common DNS suffix entry, where the domain name suffix is corp.contoso.com. This happens automatically for domains in the same root, but domains that are not in the same root must be added manually.

If Windows Internet Name Service (WINS) is deployed in a multiple domain environment, you must deploy a WINS forward lookup zone in DNS. For more information, see **Single label names** in the 1.4.2 Plan for local name resolution section earlier in this document.

1.8 Plan Group Policy Objects

This section explains the role of Group Policy Objects (GPOs) in your Remote Access infrastructure, and it includes the following subsections:

- Configure automatically created GPOs
- Configure manually created GPOs
- Manage GPOs in a multi-domain controller environment
- Manage Remote Access GPOs with limited permissions
- Recover from a deleted GPO

DirectAccess settings that are configured when you configure Remote Access are collected into GPOs. The following types of GPOs are populated with DirectAccess settings, and they are distributed as follows:

- **DirectAccess client GPO**

  This GPO contains client settings, including IPv6 transition technology settings, NRPT entries, and Windows Firewall with Advanced Security connection security rules. The GPO is applied to the security groups that are specified for the client computers.

- **DirectAccess server GPO**

  This GPO contains the DirectAccess configuration settings that are applied to any server that is configured as a Remote Access server in your deployment. It also contains Windows Firewall with Advanced Security connection security rules.

- **Application servers GPO**

  This GPO contains settings for selected application servers to which you optionally extend authentication and encryption from DirectAccess clients. If authentication and encryption are not extended, this GPO is not used.

GPOs can be configured in two ways:

- **Automatically**—You can specify that they are created automatically. A default name is specified for each GPO.
- **Manually**—You can use GPOs that have been predefined by the Active Directory administrator.

Note

After DirectAccess is configured to use specific GPOs, it cannot be configured to use different GPOs.

Whether you are using automatically or manually configured GPOs, you need to add a policy for slow link detection if your clients will use 3G networks. The path for **Policy: Configure Group Policy slow link detection** is: **Computer configuration/Polices/Administrative Templates/System/Group Policy**.

Caution

Use the following procedure to back up all Remote Access GPOs before you run DirectAccess cmdlets: Back up and Restore Remote Access Configuration.

If the correct permissions (which are listed in the following sections) for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but linking will not occur. If this warning is issued, links will not be created automatically, even when the permissions are added later. Instead the administrator needs to create the links manually.

1.8.1 Configure automatically created GPOs

Consider the following when you use automatically-created GPOs.

Automatically created GPOs are applied according to the location and link target parameter, as follows:

- For the DirectAccess server GPO, the location parameter and the link parameter point to the domain that contains the Remote Access server.
- When client and application server GPOs are created, the location is set to a single domain in which the GPO will be created. The GPO name is looked up in each domain, and it is filled with DirectAccess settings if it exists. The link target is set to the root of the domain in which the GPO was created. A GPO is created for each domain that contains client computers or application servers, and the GPO is linked to the root of its respective domain.

When you use automatically created GPOs to apply DirectAccess settings, the Remote Access administrator requires the following permissions:

- GPO create permissions for each domain
- Link permissions for all the selected client domain roots
- Link permissions for the server GPO domain roots

In addition, the following permissions are needed:

- Create, Edit, Delete, and Modify security permissions are required for the GPOs.
- We recommend that the Remote Access administrator has GPO Read permissions for each required domain. This enables Remote Access to verify that GPOs with duplicate names do not exist when creating GPOs.

1.8.2 Configure manually created GPOs

Consider the following when using manually created GPOs:

- The GPOs should exist before running the Remote Access Setup Wizard.
- To apply DirectAccess settings, the Remote Access administrator requires full GPO permissions (Edit, Delete, Modify security permissions) on the manually created GPOs.
- A search is made in the entire domain for a link to the GPO. If the GPO is not linked in the domain, a link is automatically created in the domain root. If the required permissions to create the link are not available, a warning is issued.

1.8.3 Manage GPOs in a multi-domain controller environment

Each GPO is managed by a specific domain controller, as follows:

- The server GPO is managed by one of the domain controllers in the Active Directory site that is associated with the server. If domain controllers in that site are Read-only, the server GPO is managed by the Write-enabled domain controller that is closest to the Remote Access server.
- Client and application server GPOs are managed by the domain controller that is running as the primary domain controller (PDC).

If you want to manually modify GPO settings in a multi-domain controller environment, consider the following:

- For the server GPO, to identify which domain controller is associated with the Remote Access server, from an elevated command prompt on the Remote Access server, run **nltest /dsgetdc: /writable**.

- By default, when you make changes with networking Windows PowerShell cmdlets or you make changes from the Group Policy Management console, the domain controller that is acting as the PDC is used.

In addition, if you modify settings on a domain controller that is not the domain controller associated with the Remote Access server (for the server GPO) or the PDC (for client and application server GPOs), consider the following:

- Before you modify the settings, ensure that the domain controller is replicated with an up-to-date GPO, and back up your GPO settings. For more information, see Back up and Restore Remote Access Configuration. If the GPO is not updated, merge conflicts during replication might occur, which can result in a corrupt Remote Access configuration.
- After you modify the settings, you must wait for changes to replicate to the domain controllers that are associated with the GPOs. Do not make additional changes by using the Remote Access Management console or Remote Access PowerShell cmdlets until replication is complete. If a GPO is edited on two domain controllers before replication is complete, merge conflicts might occur, which can result in a corrupt Remote Access configuration.

Alternatively, you can change the default setting by using the **Change Domain Controller** dialog box in the Group Policy Management console, or by using the Windows PowerShell cmdlet **Open-NetGPO**, so that the changes use the domain controller you specify.

- To do this in the Group Policy Management console, right-click the domain or sites container, and click **Change Domain Controller**.
- To do this in Windows PowerShell, specify the **DomainController** parameter for the **Open-NetGPO** cmdlet. For example, to enable the private and public profiles in Windows Firewall on a GPO named domain1\DA_Server_GPO _Europe by using a domain controller named europe-dc.corp.contoso.com, enter the following:

```
$gpoSession = Open-NetGPO -PolicyStore "domain1\DA_Server_GPO
_Europe" -DomainController "europe-dc.corp.contoso.com"
Set-NetFirewallProfile -GpoSession $gpoSession -Name
@("Private","Public") -Enabled True
Save-NetGPO -GpoSession $gpoSession
```

1.8.4 Manage Remote Access GPOs with limited permissions

To manage a Remote Access deployment, the Remote Access administrator requires full GPO permissions (Read, Edit, Delete, and Modify security permissions) on the GPOs that are used in the deployment. This is because the Remote Access Management console and the Remote Access PowerShell modules read the configuration from and write it to the Remote Access GPOs (that is, client, server, and application server GPOs).

In many organizations, the domain administrator who is in charge of GPO operations is not the same person as the Remote Access administrator who is in charge of the Remote Access configuration. These organizations may have policies that restrict the Remote Access administrator from having full permissions on GPOs in the domain. The domain administrator may

be also required to review the policy configuration before applying it to any computer in the domain.

To accommodate these requirements, the domain administrator should create two copies of each GPO: staging and production. The Remote Access administrator is given full permissions on the staging GPOs. The Remote Access administrator specifies the staging GPOs in the Remote Access Management console and in Windows PowerShell cmdlets as the GPOs used for the Remote Access deployment. This enables the Remote Access administrator to read and modify the Remote Access configuration as and when required.

The domain administrator must make sure that the staging GPOs are not linked to any scope-of-management in the domain and that the Remote Access administrator does not have GPO link permissions in the domain. This ensures that changes made by the Remote Access administrator to the staging GPOs do not have any effect on computers in the domain.

The domain administrator links the production GPOs to the required scope-of-management and applies the appropriate security filtering. This ensures that changes to these GPOs are applied to the computers in the domain (client computers, Remote Access servers, and application servers). The Remote Access administrator has no permissions on the production GPOs.

When changes are made to the staging GPOs, the domain administrator can review the policy configuration in these GPOs to make sure that it satisfies the security requirements in the organization. The domain administrator then exports the settings from the staging GPOs by using the backup feature, and imports the settings to the corresponding production GPOs, which will be applied to the computers in the domain.

The following diagram shows this configuration.

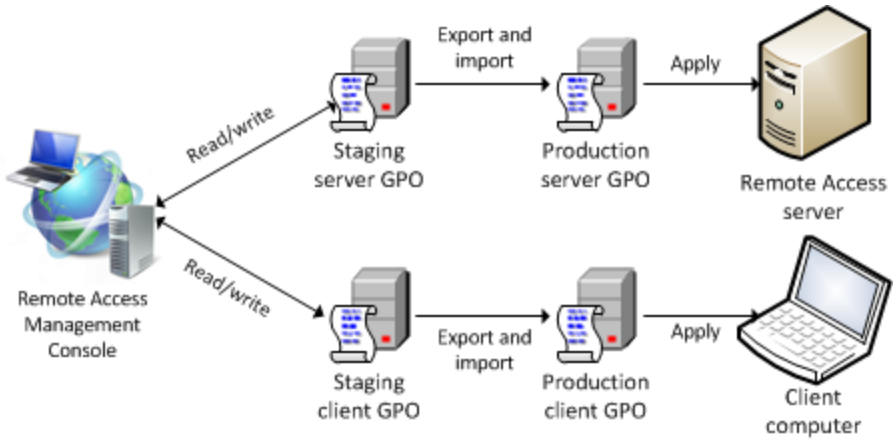

1.8.5 Recover from a deleted GPO

If a client, Remote Access server, or application server GPO has been deleted accidentally and there is no backup available, you must remove the configuration settings and reconfigure them. If a backup is available, you can restore the GPO from the backup.

The Remote Access Management console will display the following error message: **GPO <GPO name> cannot be found**. To remove the configuration settings, take the following steps:

1. Run the Windows PowerShell cmdlet **Uninstall-remoteaccess**.
2. Open the Remote Access Management console.
3. You will see an error message that the GPO is not found. Click **Remove configuration settings**. After completion, the server will be restored to an unconfigured state.

See also

- Step 2: Plan the Remote Access Deployment
- Plan an Advanced Remote Access Deployment

Step 2: Plan the Remote Access Deployment

After you plan the Remote Access infrastructure, the next step in deploying advanced Remote Access on a single server with IPv4 and IPv6 is to plan the settings for the Remote Access Setup Wizard.

Task	Description
2.1 Plan for client deployment	Plan how to allow client computers to connect by using DirectAccess. Decide which managed computers will be configured as DirectAccess clients, and plan to deploy the Network Connectivity Assistant or DirectAccess Connectivity Assistant on client computers.
2.2 Plan for Remote Access server deployment	Plan how to deploy the Remote Access server.
2.3 Plan infrastructure servers	Plan the infrastructure servers for your Remote Access deployment, including the DirectAccess network location server, Domain Name System (DNS) servers, and DirectAccess management servers.
2.4 Plan application servers	Plan for IPv4 and IPv6 application servers, and optionally consider whether end-to-end authentication between DirectAccess client computers and internal application servers is required.
2.5 Plan DirectAccess and third-party VPN clients	When deploying DirectAccess with third-party VPN clients, it may be necessary to set a registry value to enable a seamless coexistence of the two remote access solutions.

2.1 Plan for client deployment

There are three decisions to make when you are planning your client deployment:

1. Will DirectAccess be available to mobile computers only, or to any computer?

   When you configure DirectAccess clients in the DirectAccess Client Setup Wizard, you can choose to allow only mobile computers in the specified security groups to connect by using DirectAccess. If you restrict access to mobile computers, Remote Access automatically configures a WMI filter to ensure that the DirectAccess client GPO is applied only to mobile computers in the specified security groups. The Remote Access administrator requires Create or Modify security permissions to create or modify Group Policy Object (GPO) WMI filters to enable this setting.

2. What security groups will contain the DirectAccess client computers?

   DirectAccess client settings are contained in the DirectAccess client GPO. The GPO is applied to computers that are part of the security groups that you specify in the DirectAccess Client Setup Wizard. You can specify that security groups be contained in any supported domain. For more information, see section 1.7 Plan Active Directory.

   Before you configure Remote Access, you should create the security groups. You can add computers to the security group after you complete the Remote Access deployment, but if you add client computers that reside in a different domain than the security group, the client GPO will not be applied to those clients. For example, if you created SG1 in domain A for DirectAccess clients, and later add clients from domain B to this group, the client GPO will not be applied to clients from domain B. To avoid this issue, create a new client security group for each domain that contains client computers. Alternatively, if you do not want to create a new security group, run the Windows PowerShell cmdlet **Add-DAClient** with the name of the new GPO for the new domain.

3. What settings will you configure for the network connectivity assistant?

   The network connectivity assistant runs on client computers, and it provides additional information about the DirectAccess connection to end users. In the DirectAccess Client Setup Wizard, you can configure the following:

   - **Connectivity verifiers**

     A default web probe is created that clients use to validate connectivity to the internal network. The default name is:

     http://directaccess-WebProbeHost.<domain_name>

     The name should be registered manually in DNS. You can create other connectivity verifiers by using other web addresses over HTTP or **ping**. For each connectivity verifier, a DNS entry must exist.

   - **A Help Desk email address**

     If end users experience DirectAccess connectivity issues, they can send an email that contains diagnostic information to the Remote Access administrator to troubleshoot the issue.

- **A DirectAccess connection name**

  Specify a DirectAccess connection name to help end users identify the DirectAccess connection on their computer.

- **Allow DirectAccess clients to use local name resolution**

  Clients require a way to resolve names locally. If you allow DirectAccess clients to use local name resolution, end users can use local DNS servers to resolve names. When end users select to use local DNS servers for name resolutions, DirectAccess does not send resolution requests for single label names to the internal corporate DNS server. It uses local name resolution instead (by using the Link-Local Multicast Name Resolution (LLMNR) and NetBIOS over TCP/IP protocols).

2.2 Plan for Remote Access server deployment

Consider the following decisions when you are planning to deploy your Remote Access server:

- **Network topology**

  There are a couple of topologies available when deploying a Remote Access server:

  - **Two adapters**—With two network adapters, Remote Access can be configured with one network adapter connected directly to the Internet, and the other connected to the internal network. Or alternatively, the server is installed behind an edge device such as a firewall or a router. In this configuration, one network adapter is connected to the perimeter network, and the other is connected to the internal network.
  - **Single network adapter**—In this configuration, the Remote Access server is installed behind an edge device such as a firewall or a router. The network adapter is connected to the internal network.

  When selecting the topology for your deployment, see also 1.1 Plan network and server topology and settings.

- **ConnectTo address**

  Client computers use the ConnectTo address to connect to the Remote Access server. The address that you choose must match the subject name of the IP-HTTPS certificate that you deploy for the IP-HTTPS connection, and it must be available in the public DNS. For more information, see Plan certificates for IP-HTTPS.

- **Network adapters**

  The Remote Access Server Setup Wizard automatically detects the network adapters that are configured on the Remote Access server. You must make sure that the correct adapters are selected.

- **IP-HTTPS certificate**

  The Remote Access Server Setup Wizard automatically detects a certificate that is suitable for the IP-HTTPS connection. The subject name of the certificate that you select must match the ConnectTo address. If you are using self-signed certificates, you can select to use a certificate that is created automatically by the Remote Access server.

- **IPv6 prefixes**

If the Remote Access Server Setup Wizard detects that IPv6 has been deployed on the network adapters, it automatically populates IPv6 prefixes for the internal network, an IPv6 prefix to assign to DirectAccess client computers, and an IPv6 prefix to assign to VPN client computers. If the automatically generated prefixes are not correct for your native IPv6 infrastructure, you must manually change them. For more information, see 1.1 Planning network and server topology and settings.

- **Authentication**

  Decide how DirectAccess clients will authenticate to the Remote Access server:

  - **User authentication**—You can enable users to authenticate with Active Directory credentials or with two-factor authentication. For more information about authenticating with two-factor authentication, see Deploy Remote Access with OTP Authentication.
  - **Computer authentication**—You can configure computer authentication to use certificates or to use the Remote Access server as a Kerberos proxy on behalf of the client. For more information, see 1.2 Planning certification authorities and certificates.
  - **Windows 7 clients**—By default, Windows 7 client computers cannot connect to a Windows Server 2012 Remote Access deployment. If you have clients in your organization running Windows 7, and they require remote access to internal resources, you can allow them to connect. Any client computers that you want to allow to access internal resources must be a member of a security group that you specify in the DirectAccess Client Setup Wizard.

    Note
    Allowing clients running Windows 7 to connect by using DirectAccess requires that you use computer certificate authentication.

- **VPN configuration**

  Before you configure Remote Access, decide if you are going to provide VPN access to remote clients. You should provide VPN access if you have client computers in your organization that do not support DirectAccess connectivity (because they are unmanaged or because run an operating system for which DirectAccess is not supported). The Remote Access Server Setup Wizard allows you to configure how IP addresses are assigned (by using DHCP or from a static address pool) and how VPN clients are authenticated (by using Active Directory or a RADIUS server).

2.3 Plan infrastructure servers

Remote Access requires three types of infrastructure servers:

- **DNS servers**— For more information, see 1.4 Plan DNS requirements
- **Network location server**—For more information, see 1.5 Plan the network location server
- **Management servers**— For more information, see 1.6 Plan management servers

2.4 Plan application servers

Application servers are the servers on the corporate network that are accessible by client computers over a DirectAccess connection. Application servers are identified by adding them into a security group. The application server GPO is then applied to servers in that group.

Note

Adding the application servers to a security group is required only if you require end-to-end authentication and encryption.

You can optionally require end-to-end authentication and encryption between DirectAccess client and selected internal application servers. If you configure end-to-end authentication, DirectAccess clients use an IPsec transport policy. This policy requires that the authentication and traffic protection of IPsec sessions is terminated at the specified application servers. In this case, the Remote Access server forwards the authenticated and protected IPsec sessions to the application servers.

By default, when you extend authentication to application servers, the data payload is encrypted between the DirectAccess client and the application server. You can choose to not encrypt traffic, and use authentication only. However, this is less secure than using authentication and encryption, and it is supported only for application servers running the Windows Server 2008 R2, or Windows Server 2012 operating systems.

2.5 Plan DirectAccess and third-party VPN clients

Some third-party VPN clients do not create connections in the Network Connections folder. This can cause DirectAccess to determine that it has no intranet connectivity when the VPN connection is established and connectivity to the intranet exists. This occurs when third-party VPN clients register their interfaces by defining them as Network Device Interface Specification (NDIS) endpoint types. You can enable coexistence with these types of VPN clients by setting the following registry value to 1 on DirectAccess clients:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\NlaSvc\Parameters\ShowDomainEndpointInterfaces (REG_DWORD)

Some third-party VPN clients use a split-tunneling configuration, which allows the VPN client computer to access the Internet directly, without the need to send the traffic through the VPN connection to the intranet.

Split-tunnel configurations typically leave the default gateway setting on the VPN client as not configured or as all zeros (0.0.0.0). You can confirm this by establishing a successful VPN connection to your intranet and using the Ipconfig.exe command line tool to view the resulting configuration.

If the VPN connection lists its default gateway as empty or all zeros (0.0.0.0), your VPN client is configured in this way. By default, the DirectAccess client does not identify split-tunnelf configurations. To configure DirectAccess clients to detect these types of VPN client configurations and coexist with them, set the following registry value to 1:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\NlaSvc\Parameters\Internet\EnableNoGatewayLocationDetection (REG_DWORD)

See also

- Step 1: Plan the Remote Access Infrastructure
- Plan an Advanced Remote Access Deployment

Install and Configure Advanced Remote Access

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This overview lists the configuration steps that are required to deploy a single Remote Access server running Windows Server 2012 with IPv4 and IPv6.

- Step 1: Configure the Remote Access Infrastructure.

  In this phase, you configure network and server settings, certificate requirements, Domain Name System (DNS) settings, the network location server deployment, DirectAccess management servers, Active Directory settings, and Group Policy Objects (GPOs).
- Step 2: Configure the Remote Access Server.

  In this phase, you configure the DirectAccess client computers, server settings, infrastructure servers, and application servers.
- Step 3: Verify the Deployment.

  This step describes steps for verifying the deployment.

Before you start your deployment, verify the planning steps that are described in Plan an Advanced Remote Access Deployment.

See also

1. Deploy a Single Remote Access Server with Advanced Settings

Step 1: Configure the Remote Access Infrastructure

This topic describes how to configure the infrastructure that is required for an advanced Remote Access deployment that uses a single Remote Access server in a mixed IPv4 and IPv6 environment. Before you begin the deployment steps, ensure that you have completed the planning steps that are described in Plan an Advanced Remote Access Deployment.

Task	Description
1.1 Configure server network settings	Configure the server network settings on the Remote Access server.
1.2 Configure force tunneling	Configure force tunneling.
1.3 Configure routing in the corporate network	Configure routing in the corporate network.
1.4 Configure firewalls	Configure additional firewalls, if required.
1.5 Configure CAs and certificates	Configure a certification authority (CA), if required, and any other certificate templates that are required in the deployment.
1.6 Configure the DNS server	Configure the Domain Name System (DNS) settings for the Remote Access server.
1.7 Configure Active Directory	Join client computers and the Remote Access server to the Active Directory domain.
1.8 Configure GPOs	Configure GPOs for the deployment, if required.
1.9 Configure security groups	Configure security groups that will contain DirectAccess client computers, and any other security groups that are required in the deployment.
1.10 Configure the network location server	Configure the network location server, including installing the network location server website certificate.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1 Configure server network settings

The following network interface settings are required for a single server deployment in an environment that is using IPv4 and IPv6. All IP addresses are configured by using **Change adapter settings** in the **Windows Networking and Sharing Center**.

Edge topology

- Two Internet-facing consecutive public static IPv4 or IPv6 addresses

Note

Two public addresses are required for Teredo. If you are not using Teredo, you can configure a single public static IPv4 address.

- A single internal static IPv4 or IPv6 address

Behind NAT device (with two network adapters)

- A single Internet-facing static IPv4 or IPv6 address
- A single internal network-facing static IPv4 or IPv6 address

Behind NAT device (with one network adapter)

- A single internal network-facing static IPv4 or IPv6 address

Notes

If a Remote Access server with two or more network adapters (one classified in the domain profile and the other in a public or private profile) is configured with a single network adapter topology, we recommend the following:

- Ensure that the second network adapter and any additional network adapters are classified in the domain profile.
- If the second network adapter cannot be configured for the domain profile, the DirectAccess IPsec policy must be manually scoped to all profiles by using the following Windows PowerShell command after DirectAccess is configured:

```
$gposession = Open-NetGPO –PolicyStore <Name of the server GPO>
Set-NetIPsecRule –DisplayName <Name of the IPsec policy> –
GPOSession $gposession –Profile Any
Save-NetGPO –GPOSession $gposession
```

1.2 Configure force tunneling

Force tunneling can be configured through the Remote Access Setup Wizard. It is presented as a check box in the Configure Remote Clients Wizard. This setting only affects DirectAccess clients. If VPN is enabled, VPN clients will by default use force tunneling. Administrators can change the setting for VPN clients from the client profile.

Selecting the check box for force tunneling does the following:

- Enables force tunneling on DirectAccess clients
- Adds an **Any** entry in the Name Resolution Policy Table (NRPT) for DirectAccess clients, which means that all DNS traffic will go to the internal network DNS servers
- Configures DirectAccess clients to always use the IP-HTTPS transition technology

To make Internet resources available to DirectAccess clients that use force tunneling, you can use a proxy server, which can receive IPv6-based requests for Internet resources and translate them to requests for IPv4-based Internet resources. To configure a proxy server for Internet resources, you need to modify the default entry in NRPT to add the proxy server. You can accomplish this by using the Remote Access PowerShell cmdlets or the DNS PowerShell cmdlets. For example, use the Remote Access PowerShell cmdlet as follows:

```
Set-DAClientDNSConfiguration –DNSSuffix '.' –ProxyServer <Name of the proxy server:port>
```

Note

If DirectAccess and VPN are enabled on the same server, and VPN is in force-tunnel mode, and the server is deployed in an edge topology or a behind NAT topology (with two network adapters, one connected to the domain and one to a private network), VPN Internet traffic cannot be forwarded through the external interface of the Remote Access server. To enable this scenario, organizations must deploy Remote Access on the server behind a firewall in single network adapter topology. Alternatively, organizations can use a separate proxy server in the internal network to forward the Internet traffic from VPN clients.

Note

If an organization is using a web proxy for DirectAccess clients to access Internet resources, and the corporate proxy is not capable of handling internal network resources, DirectAccess clients will not be able to access internal resources if they are outside the intranet. In such a scenario, to enable DirectAccess clients to access internal resources, manually create NRPT entries for the internal network suffixes by using the DNS page of the infrastructure wizard. Do not apply proxy settings on these NRPT suffixes. The suffixes should be populated with default DNS server entries.

1.3 Configure routing in the corporate network

Configure routing in the corporate network as follows:

- When native IPv6 is deployed in the organization, add a route so that the routers on the internal network route IPv6 traffic back through the Remote Access server.
- Manually configure the organization's IPv4 and IPv6 routes on the Remote Access servers. Add a published route so that all traffic with an organization (/48) IPv6 prefix is forwarded to the internal network. For IPv4 traffic, add explicit routes so that IPv4 traffic is forwarded to the internal network.

1.4 Configure firewalls

When using additional firewalls in your deployment, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- Teredo traffic—User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound.
- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required.

Note

This exemption must be configured on the Remote Access server, while all the other exemptions have to be configured on the edge firewall.

Note

For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server. For IP-HTTPS the exceptions need only be applied to the address where the public name of the server resolves.

When using additional firewalls, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.
- Internet Control Message Protocol for IPv6 (ICMPv6) traffic inbound and outbound – for Teredo implementations only.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic
- ICMP for all IPv4/IPv6 traffic

1.5 Configure CAs and certificates

Remote Access in Windows Server 2012 allows you to choose between using certificates for computer authentication or using a built in Kerberos proxy that authenticates using user names and passwords. You must also configure an IP-HTTPS certificate on the Remote Access server.

For more information, see Active Directory Certificate Services.

1.5.1 Configure IPsec authentication

A computer certificate is required on the Remote Access server and on all DirectAccess clients to use IPsec authentication. The certificate must be issued by an internal certification authority (CA), and Remote Access servers and DirectAccess clients must trust the CA chain that issues root and intermediate certificates.

To configure IPsec authentication

1. In the internal CA, decide if you will use the **Computer** certificate template, or if you will create a new certificate template as described in Creating Certificate Templates.

   Note

   If you create a new template, it must be configured for Client Authentication.
2. Deploy the certificate template, if required. For more information, see Deploying Certificate Templates.

3. Configure the certificate template for autoenrollment, if required. For more information, see Configure Certificate Autoenrollment.

1.5.2 Configure certificate templates

When you use an internal CA to issue certificates, you must configure a certificate template for the IP-HTTPS certificate and the network location server website certificate.

To configure a certificate template

1. In the internal CA, create a certificate template as described in Creating Certificate Templates.
2. Deploy the certificate template as described in Deploying Certificate Templates.

1.5.3 Configure the IP-HTTPS certificate

Remote Access requires an IP-HTTPS certificate to authenticate IP-HTTPS connections to the Remote Access server. There are three certificate options that are available for IP-HTTPS authentication:

Public certificate

A public certificate is supplied by a third party. If the certificate subject name does not contain wildcard characters, it must be the externally resolvable fully qualified domain name (FQDN) URL that is used only for the Remote Access server IP-HTTPS connections.

Private certificate

If you use a private certificate, the following are required, if they do not already exist:

- A website certificate that is used for IP-HTTPS authentication. The certificate subject should be an externally resolvable FQDN that is reachable from the Internet. The certificate is based on the certificate template that you created by following the instructions in 1.5.2 Configure certificate templates.
- A certificate revocation list (CRL) distribution point that is reachable from a publicly resolvable FQDN.

Self-signed certificate

If you use a self-signed certificate, the following are required, if they do not already exist:

- A website certificate that is used for IP-HTTPS authentication. The certificate subject should be an externally resolvable FQDN that is reachable from the Internet.
- A CRL distribution point that is reachable from a publicly resolvable FQDN.

Note
Self-signed certificates cannot be used in multisite deployments.

Make sure that the website certificate that is used for IP-HTTPS authentication meets the following requirements:

- The common name of the certificate should match the name of the IP-HTTPS site.

- In the **Subject** field, specify the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL.
- For the **Enhanced Key Usage** field, use the server authentication object identifier (OID).
- For the **CRL Distribution Points** field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.
- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcard characters in the name.

To install the IP-HTTPS certificate from an internal CA

1. On the Remote Access server: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, select the check box for the certificate template that you previously created (for more information, see 1.5.2 Configure certificate templates). If required, click **More information is required to enroll for this certificate**.
8. In the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common Name**.
9. In **Value**, specify the IPv4 address of the external facing adapter of the Remote Access server or the FQDN of the IP-HTTPS URL, and then click **Add**.
10. In the **Alternative name** area, in **Type**, select **DNS**.
11. In **Value**, specify the IPv4 address of the external facing adapter of the Remote Access server or the FQDN of the IP-HTTPS URL, and then click **Add**.
12. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
13. On the **Extensions** tab, click the arrow next to **Extended Key Usage**, and make sure that **Server Authentication** appears in the **Selected options** list.
14. Click **OK**, click **Enroll**, and then click **Finish**.
15. In the details pane of the Certificates snap-in, verify that the new certificate was enrolled with Intended Purposes of Server Authentication.

1.6 Configure the DNS server

You must manually configure a DNS entry for the network location server website for the internal network in your deployment.

Do this step using Windows PowerShell

To create the network location server

1. On the internal network DNS server: On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the left pane of the **DNS Manager** console, expand the forward lookup zone for your domain. Right click the domain, and click **New Host (A or AAAA)**.
3. In the **New Host** dialog box, In the **IP address** box:
   - In the **Name (uses parent domain name if blank)** box, enter the DNS name for the network location server website (this is the name that the DirectAccess clients use to connect to the network location server).
   - Enter the IPv4 or IPv6 address of the network location server, and then click **Add Host**, and then click **OK**.
4. In the **New Host** dialog box:
   - In the **Name (uses parent domain name if blank)** box, enter the DNS name for the web probe (the name for the default web probe is **directaccess-webprobehost**).
   - In the **IP address** box, enter the IPv4 or IPv6 address of the web probe, and then click **Add Host**.
   - Repeat this process for **directaccess-corpconnectivityhost** and any manually created connectivity verifiers.
5. In the **DNS** dialog box, click **OK**, and then click **Done**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-DnsServerResourceRecordA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv4Address <network_location_server_IPv4_address>

Add-DnsServerResourceRecordAAAA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv6Address <network_location_server_IPv6_address>
```

You must also configure DNS entries for the following:

- **The IP-HTTPS server**

  DirectAccess clients must be able to resolve the DNS name of the Remote Access server from the Internet.
- **CRL revocation checking**

  DirectAccess uses certificate revocation checking for the IP-HTTPS connection between DirectAccess clients and the Remote Access server, and for the HTTPS-based connection between the DirectAccess client and the network location server. In both cases, DirectAccess clients must be able to resolve and access the CRL distribution point location.
- **ISATAP**

Intrasite Automatic Tunnel Addressing Protocol (ISATAP) uses tunneling to enable DirectAccess clients to connect to the Remote Access server over the IPv4 Internet, encapsulating IPv6 packets within an IPv4 header. It is used by Remote Access to provide IPv6 connectivity to ISATAP hosts across an intranet. In a non-native IPv6 network environment, the Remote Access server configures itself automatically as an ISATAP router. Resolution support for the ISATAP name is required.

1.7 Configure Active Directory

The Remote Access server and all DirectAccess client computers must be joined to an Active Directory domain. DirectAccess client computers must be a member of one of the following domain types:

- Domains that belong in the same forest as the Remote Access server.
- Domains that belong to forests with a two-way trust with the Remote Access server forest.
- Domains that have a two-way domain trust to the Remote Access server domain.

Do this step using Windows PowerShell

To join the Remote Access server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. In the **System Properties** dialog box, click the **Computer Name** tab, and then click **Change**.
3. In **Computer Name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server (for example, corp.contoso.com), and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
5. When you see a dialog box that welcomes you to the domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. In the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

To join client computers to the domain

1. On the **Start** screen, type **explorer.exe**, and then press ENTER.
2. Right-click the Computer icon, and then click **Properties**.
3. On the **System** page, click **Advanced system settings**.
4. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
5. In **Computer name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server (for example,

corp.contoso.com), and then click **OK**.

6. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
7. When you see a dialog box that welcomes you to the domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. In the **System Properties** dialog box, click **Close**.
10. When you are prompted to restart the computer, click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note

You must supply domain credentials when you enter the following **Add-Computer** command.

```
Add-Computer -DomainName <domain_name>

Restart-Computer
```

1.8 Configure GPOs

A minimum of two Group Policy Objects are required to deploy Remote Access:

- One contains settings for the Remote Access server
- One contains settings for DirectAccess client computers

When you configure Remote Access, the wizard automatically creates the required Group Policy Objects. However, if your organization enforces a naming convention, you can type a name in the GPO dialog box in the Remote Access Management console. For more information, see 2.7. Configuration summary and alternate GPOs. If you have created permissions, the GPO will be created. If you do not have the required permissions to create GPOs, they must be created prior to configuring Remote Access.

To create Group Policy Objects, see Create and Edit a Group Policy Object.

Important

Administrators can manually link the DirectAccess Group Policy Objects to an organizational unit (OU) by following these steps:

1. Before you configure DirectAccess, link the created GPOs to the respective OUs.
2. When you configure DirectAccess, specify a security group for the client computers.
3. The Remote Access administrator may or may not have permissions to link the Group Policy Objects to the domain. In either case, the Group Policy Objects will be configured automatically. If the GPOs are already linked to an OU, the links will not be removed, and the GPOs will not be linked to the domain. For a server GPO, the OU must contain the server computer object, or the GPO will be linked to the root of the domain.

4. If you did not link to the OU before running the DirectAccess Wizard, after the configuration is complete, the domain administrator can link the DirectAccess Group Policy Objects to the required OUs. The link to the domain can be removed. For more information, see Link a Group Policy Object.

Note
If a Group Policy Object was created manually, it is possible that the Group Policy Object will not be available during the DirectAccess configuration. The Group Policy Object may not have been replicated to the closest domain controller to the management computer. In this event, the administrator can wait for replication to complete, or force the replication.

1.8.1 Configure Remote Access GPOs with limited permissions

In a deployment that uses staging and production GPOs (see 1.8.4 Managing Remote Access GPOs with limited permissions), the domain administrator should do the following:

1. Obtain the list of required GPOs for the Remote Access deployment from the Remote Access administrator. For more information, see 1.8 Plan Group Policy Objects.
2. For each GPO that is requested by the Remote Access administrator, create a pair of GPOs with different names. The first will be used as the staging GPO, and the second as the production GPO.

   To create Group Policy Objects, see Create and Edit a Group Policy Object.
3. To link the production GPOs, see Link a Group Policy Object.
4. Grant the Remote Access administrator **Edit settings, delete and modify security** permissions on all of the staging GPOs. For more informantion, see Delegate Permissions for a Group or User on a Group Policy Object.
5. Deny the Remote Access administrator permissions to link GPOs in all domains (or verify that the Remote Access administrator doesn't have such permissions). For more information, see Delegate Permissions to Link Group Policy Objects.

When Remote Access administrators configure Remote Access, they should always specify only the staging GPOs (not the production GPOs). This is true in the initial configuration of Remote Access and when performing additional configuration operations where additional GPOs are required; for example, when adding entry points in a multisite deployment or enabling client computers in additional domains.

After the Remote Access administrator completes any changes to the Remote Access configuration, the domain administrator should review the settings in the staging GPOs, and use the following procedure to copy the settings to the production GPOs.

Tip
Perform the following procedure after each change of the Remote Access configuration.

Do this step using Windows PowerShell

To copy settings to the production GPOs

1. Verify that all of the staging GPOs in the Remote Access deployment have been replicated to all of the domain controllers in the domain. This is required to ensure the most up-to-date configuration is imported to the production GPOs. For more information, see Check Group Policy Infrastructure Status.
2. Export the settings by backing up all of the staging GPOs in the Remote Access deployment. For more information, see Back Up a Group Policy Object.
3. For each production GPO, change the security filters to match the security filters of the corresponding staging GPO. For more information, see Filter Using Security Groups.

   Note
   This is required because **Import Settings** does not copy the security filter of the source GPO.
4. For each production GPO, import the settings from the backup of the corresponding staging GPO as follows:
   a. In the Group Policy Management Console (GPMC), expand the Group Policy Objects node in the forest and domain that contains the production Group Policy Object into which the settings will be imported.
   b. Right-click the GPO, and click **Import Settings**.
   c. In the **Import Settings Wizard**, on the **Welcome** page, click **Next**.
   d. On the **Backup GPO** page, click **Backup**.
   e. In the **Back up Group Policy Object** dialog box, in the **Location** box, enter the path for the location where you want to store the GPO backups, or click **Browse** to locate the folder.
   f. In the **Description** box, type a description for the production GPO, and then click **Back Up**.
   g. When the backup completes, click **OK**, and then on the **Backup GPO** page, click **Next**.
   h. On the **Backup location** page, in the **Backup folder** box, enter the path for the location in which the backup of the corresponding staging GPO was stored in Step 2, or click **Browse** to locate the folder, and then click **Next**.
   i. On the **Source GPO** page, select the **Show only the latest version of each GPO** check box to hide older backups, and select the corresponding staging GPO. Click **View Settings** to review the Remote Access settings before applying them to the production GPO, and then click **Next**.
   j. On the **Scanning Backup** page, click **Next**, and then click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

- To back up the staging client GPO “DirectAccess Client Settings - Staging” in domain “corp.contoso.com” to the backup folder “C:\Backups\”:

```
$backup = Backup-GPO -Name 'DirectAccess Client Settings -
Staging' -Domain 'corp.contoso.com' -Path 'C:\Backups\'
```

- To see the security filtering of the staging client GPO “DirectAccess Client Settings - Staging” in domain “corp.contoso.com”:

```
Get-GPPermission -Name 'DirectAccess Client Settings - Staging'
-Domain 'corp.contoso.com' -All | ?{ $_.Permission -eq
'GpoApply'}
```

- To add the security group “corp.contoso.com\DirectAccess clients” to the security filter of the production client GPO “DirectAccess Client Settings – Production” in domain “corp.contoso.com”:

```
Set-GPPermission -Name 'DirectAccess Client Settings -
Production' -Domain 'corp.contoso.com' -PermissionLevel GpoApply
-TargetName 'corp.contoso.com\DirectAccess clients' -TargetType
Group
```

- To import settings from the backup to the production client GPO “DirectAccess Client Settings – Production” in domain “corp.contoso.com”:

```
Import-GPO -BackupId $backup.Id -Path $backup.BackupDirectory -
TargetName 'DirectAccess Client Settings - Production' -Domain
'corp.contoso.com'
```

1.9 Configure security groups

The DirectAccess settings that are contained in the client computer Group Policy Object are applied only to computers that are members of the security groups that you specify when you configure Remote Access. In addition, if you are using security groups to manage your application servers, create a security group for these servers.

Do this step using Windows PowerShell

To create a security group for DirectAccess clients

1. On the **Start** screen, type **dsa.msc**, and then press ENTER. In the **Active Directory Users and Computers** console, in the left pane, expand the domain that will contain the security group, right-click **Users**, point to **New**, and then click **Group**.
2. In the **New Object - Group** dialog box, under **Group name**, enter the name for the security group.
3. Under **Group scope**, click **Global**, and under **Group type**, click **Security**, and then click **OK**.
4. Double-click the DirectAccess client computers security group, and in the properties dialog box, click the **Members** tab.
5. On the **Members** tab, click **Add**.
6. In the **Select Users, Contacts, Computers, or Service Accounts** dialog box, select the client computers that you want to enable for DirectAccess, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADGroup -GroupScope global -Name <DirectAccess_clients_group_name>

Add-ADGroupMember -Identity DirectAccess_clients_group_name -Members <computer_name>
```

1.10 Configure the network location server

The network location server should be a server with high availability, and it should have a valid SSL certificate that is trusted by the DirectAccess clients. There are two certificate options for the network location server certificate:

- **Private certificate**

  This certificate is based on the certificate template that you created by following the instructions in 1.5.2 Configure certificate templates.
- **Self-signed certificate**

  Note
  Self-signed certificates cannot be used in multisite deployments.

The following are required for either type of certificate, if they do not already exist:

- A website certificate that is used for the network location server. The certificate subject should be the URL of the network location server.
- A CRL distribution point that has high availability from the internal network.

Note
If the network location server website is located on the Remote Access server, a website is created automatically when you configure Remote Access. This site is bound to the server certificate that you provide.

To install the network location server certificate from an internal CA

1. On the server that will host the network location server website: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, select the check box for the certificate template that you created by following the instructions in 1.5.2 Configure certificate templates. If

required, click **More information is required to enroll for this certificate**.

8. In the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common Name**.
9. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
10. In the **Alternative name** area, in **Type**, select **DNS**.
11. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
12. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
13. Click **OK**, click **Enroll**, and then click **Finish**.
14. In the details pane of the Certificates snap-in, verify that new certificate was enrolled with Intended Purposes of Server Authentication.

To configure the network location server

1. Set up a website on a high availability server. The website does not require any content, but when you test it, you might define a default page that provides a message when clients connect.

   Note
   This step is not required if the network location server website is hosted on the Remote Access server.

2. Bind an HTTPS server certificate to the website. The common name of the certificate should match the name of the network location server site. Ensure that DirectAccess clients trust the issuing CA.

   Note
   This step is not required if the network location server website is hosted on the Remote Access server.

3. Set up a CRL site that has high availability from the internal network.

   CRL distribution points can be accessed through:

   - Web servers by using an HTTP-based URL, such as: http://crl.corp.contoso.com/crld/corp-APP1-CA.crl
   - File servers that are accessed through a universal naming convention (UNC) path, such as \\crl.corp.contoso.com\crld\corp-APP1-CA.crl

   If the internal CRL distribution point is reachable only over IPv6, you must configure a Windows Firewall with Advanced Security connection security rule to exempt IPsec protection from the IPv6 address of your intranet to the IPv6 addresses of your CRL distribution points.

4. Ensure that DirectAccess clients on the internal network can resolve the name of the network location server. Ensure that the name is not resolvable by DirectAccess clients on the Internet.

See also

- Step 2: Configure the Remote Access Server
- Install and Configure Advanced Remote Access
- Deploy a Single Remote Access Server with Advanced Settings

Step 2: Configure the Remote Access Server

This topic describes how to configure the client and server settings that are required for an advanced Remote Access deployment that uses a single Remote Access server in a mixed IPv4 and IPv6 environment. Before you begin the deployment steps, ensure that you have completed the planning steps that are described in Step 2: Plan the Remote Access Deployment.

Task	Description
2.1. Install the Remote Access role	Install the Remote Access role.
2.2. Configure the deployment type	Configure the deployment type as DirectAccess and VPN, DirectAccess only, or VPN only.
2.3. Configure DirectAccess clients	Configure the Remote Access server with the security groups that contain DirectAccess clients.
2.4. Configure the Remote Access server	Configure Remote Access server settings.
2.5. Configure the infrastructure servers	Configure the infrastructure servers that are used in the organization.
2.6. Configure application servers	Configure application servers so that they require authentication and encryption.
2.7. Configuration summary and alternate GPOs	View the Remote Access configuration summary, and modify the GPOs if desired.
2.8. How to configure the Remote Access server by using Windows PowerShell	Configure Remote Access by using Windows PowerShell.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

2.1. Install the Remote Access role

To deploy Remote Access, you must install the Remote Access role on a server in your organization that will act as the Remote Access server.

Do this step using Windows PowerShell

To install the Remote Access role

1. On the Remote Access server, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the **Select server roles** screen.
3. On the **Select server roles** page, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Confirm installation selections** page, click **Install**.
6. On the **Installation progress** page, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature RemoteAccess -IncludeManagementTools
```

2.2. Configure the deployment type

Remote Access can be deployed by using the Remote Access Management console in three ways:

- DirectAccess and VPN
- DirectAccess only
- VPN only

This guide uses a DirectAccess-only deployment in the example procedures.

To configure the deployment type

1. On the Remote Access server, open the Remote Access Management console: On the **Start** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, in the middle pane, click **Run the Remote Access Setup Wizard**.
3. In the **Configure Remote Access** dialog box, click the whether to deploy DirectAccess and VPN, DirectAccess only, or VPN only.

2.3. Configure DirectAccess clients

For a client computer to be provisioned to use DirectAccess, it must belong to the selected security group. After DirectAccess is configured, client computers in the security group are provisioned to receive the DirectAccess Group Policy Object (GPO). You can also configure the deployment scenario, which allows you to configure DirectAccess for client access and remote management, or for remote management only.

To configure DirectAccess clients

1. In the middle pane of the Remote Access Management console, in the **Step 1 Remote Clients** area, click **Configure**.
2. In the DirectAccess Client Setup Wizard, on the **Deployment Scenario** page, click the deployment scenario that you want to use in your organization (**Full DirectAccess** or **Remote management only**) , and then click **Next**.
3. On the **Select Groups** page, click **Add**.
4. In the **Select Groups** dialog box, select the security groups that contain your DirectAccess client computers.

   Note
   If the security group is located in a different forest than the Remote Access server, after you complete the Remote Access Setup Wizard, click **Refresh Management Servers** in the **Tasks** pane to discover the domain controllers and System Center Configuration Manager servers in the new forest.
5. Select the **Enable DirectAccess for mobile computers only** check box to allow only mobile computers to access the internal network, if required.
6. Select the **Use force tunneling** check box to route all client traffic (to the internal network and to the Internet) through the Remote Access server, if required.
7. Click **Next**.
8. On the **Network Connectivity Assistant** page:
   - In the table, add resources that will be used to determine connectivity to the internal network. A default web probe is created automatically if no other resources are configured.

     Caution
     When you configure the web probe locations for determining connectivity to the Enterprise network, ensure that you have at least one HTTP-based probe configured. Configuring a **ping** probe only is not sufficient, and this could lead to inaccurate determination of connectivity status. This is because **ping** is exempt from IPsec, and as a result, it does not ensure that the IPsec tunnels are properly established.
   - Add a Help Desk email address to allow users to send information if they experience connectivity issues.
   - Provide a friendly name for the DirectAccess connection. This name appears in the

network list when users click the network icon in the notification area.

- Select the **Allow DirectAccess clients to use local name resolution** check box, if required.

  Note
  When local name resolution is enabled, users who run the Network Connectivity Assistant can select to resolve names by using DNS servers that are configured on the DirectAccess client computer.

9. Click **Finish**.

2.4. Configure the Remote Access server

To deploy Remote Access, you need to configure the Remote Access server with the correct network adapters, a public URL for the Remote Access server to which client computers can connect (the ConnectTo address), an IP-HTTPS certificate with a subject that matches the ConnectTo address, IPv6 settings, and client computer authentication.

To configure the Remote Access server

1. In the middle pane of the Remote Access Management console, in the **Step 2 Remote Access Server** area, click **Configure**.
2. In the Remote Access Server Setup Wizard, on the **Network Topology** page, click the deployment topology that will be used in your organization. In **Type the public name or IPv4 address used by clients to connect to the Remote Access server**, enter the public name for the deployment (this name matches the subject name of the IP-HTTPS certificate, for example, edge1.contoso.com), and then click **Next**.
3. On the **Network Adapters** page, the wizard automatically detects the network adapters for the networks in your deployment. If the wizard does not detect the correct network adapters, manually select the correct adapters. The wizard also automatically detects the IP-HTTPS certificate, based on the public name for the deployment set in the previous step of the wizard. If the wizard does not detect the correct IP-HTTPS certificate, click **Browse** to manually select the correct certificate, and then click **Next**.
4. On the **Prefix Configuration** page (this appears only if IPv6 is deployed in the internal network), the wizard automatically detects the IPv6 settings that are used in the internal network. If your deployment requires additional prefixes, configure the IPv6 prefixes for the internal network, an IPv6 prefix to assign to DirectAccess client computers, and an IPv6 prefix to assign to VPN client computers.

   Note
   You can specify multiple internal IPv6 prefixes by using a semicolon delimited list, for example, 2001:db8:1::/48;2001:db8:2::/48.

5. On the **Authentication** page:
   - In **User Authentication**, click **Active Directory credentials**. To configure a deployment by using two-factor authentication, click **Two-factor authentication**. For

more information, see Deploy Remote Access with OTP Authentication.

- For multisite and two-factor authentication deployments, you must use computer certificate authentication. Select the **Use computer certificates** check box to use computer certificate authentication, and select the IPsec root certificate.
- To enable Windows 7 client computers to connect through DirectAccess, select the **Enable Windows 7 client computers to connect via DirectAccess** check box.

  Note
  You must also use computer certificate authentication in this type of deployment.

6. Click **Finish**.

2.5. Configure the infrastructure servers

To configure the infrastructure servers in a Remote Access deployment, you must configure the network location server, DNS settings (including the DNS suffix search list), and management servers that are not automatically detected by Remote Access.

To configure the infrastructure servers

1. In the middle pane of the Remote Access Management console, in the **Step 3 Infrastructure Servers** area, click **Configure**.
2. In the Infrastructure Server Setup Wizard, on the **Network Location Server** page, click the option that corresponds to the location of the network location server in your deployment. If the network location server is on a remote web server, enter the URL and click **Validate** before you continue. If the network location server is on the Remote Access server, click **Browse** to locate the relevant certificate, and then click **Next**.
3. On the **DNS** page, in the table, enter any additional name suffixes that will be applied as Name Resolution Policy Table (NRPT) exemptions. Select a local name resolution option, and then click **Next**.
4. On the **DNS Suffix Search List** page, the Remote Access server automatically detects any domain suffixes in the deployment. Use the **Add** and **Remove** buttons to add and remove domain suffixes from the list of domain suffixes to use. To add a new domain suffix, in **New Suffix**, enter the suffix, and then click **Add**. Click **Next**.
5. On the **Management** page, add any management servers that are not detected automatically, and then click **Next**. Remote Access automatically adds domain controllers and System Center Configuration Manager servers.

   Note
   Although the servers are added automatically, they don't appear in the list. After you apply the configuration the first time, the System Center Configuration Manager servers appear in the list.

6. Click **Finish**.

2.6. Configure application servers

In a Remote Access deployment, configuring application servers is an optional task. Remote Access enables you to require authentication for selected application servers, which is determined by their inclusion in an application servers security group. By default, traffic to application servers that require authentication is also encrypted; however, you can choose to not encrypt traffic to application servers and use authentication only.

Note
Authentication without encryption is supported only on application servers running Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2.

To configure application servers

1. In the middle pane of the Remote Access Management console, in the **Step 4 Application Servers** area, click **Configure**.
2. In the DirectAccess Application Server Setup Wizard, to require authentication to selected application servers, click **Extend authentication to selected application servers**. Click **Add** to select the application server security group.
3. To limit access to only the servers in the application server security group, select the **Allow access only to servers included in the security groups** check box.
4. To use authentication without encryption, select the **Do not encrypt traffic. Use authentication only** check box.
5. Click **Finish**.

2.7. Configuration summary and alternate GPOs

When the Remote Access configuration is complete, the **Remote Access Review** is displayed. You can review all of the settings that you previously selected, including:

1. **GPO Settings**: The DirectAccess server GPO name and client GPO name are listed. Additionally, you can click the **Change** link next to the **GPO Settings** heading to modify the GPO settings.
2. **Remote Clients**: The DirectAccess client configuration is displayed, including the security group, force tunneling status, connectivity verifiers, and DirectAccess connection name.
3. **Remote Access Server**: The DirectAccess configuration is displayed including the public name/address, network adapter configuration, certificate information, and OTP information if configured.
4. **Infrastructure Servers**: This list includes the network location server URL, DNS suffixes that are used by DirectAccess clients, and management server information.
5. **Application Servers**: The DirectAccess remote management status is displayed, in addition to the status of the end-to-end authentication to specific application servers.

2.8. How to configure the Remote Access server by using Windows PowerShell

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To perform a full installation in an edge topology of Remote Access for DirectAccess only in a domain with the root **corp.contoso.com** and using the following parameters: server GPO: **DirectAccess Server Settings**, client GPO: DirectAccess Client Settings, internal network adapter: **Corpnet**, external network adapter: **Internet**, ConnectTto address: **edge1.contoso.com**, and network location server: **nls.corp.contoso.com**:

```
Install-RemoteAccess -Force -PassThru -ServerGpoName 'corp.contoso.com\DirectAccess
Server Settings' -ClientGpoName 'corp.contoso.com\DirectAccess Client Settings' -
DAInstallType 'FullInstall' -InternetInterface 'Internet' -InternalInterface 'Corpnet' -
ConnectToAddress 'edge1.contoso.com' -NlsUrl 'https://nls.corp.contoso.com/'
```

To configure the Remote Access server to use computer certificate authentication, with an IPsec root certificate that is issued by the certification authority named CORP-APP1-CA:

```
$certs = Get-ChildItem Cert:\LocalMachine\Root
$IPsecRootCert = $certs | Where-Object {$_.Subject -Match "corp-APP1-CA"}
Set-DAServer -IPsecRootCertificate $IPsecRootCert
```

To add the security group that contains DirectAccess clients named **DirectAccessClients**, and to remove the default Domain Computers security group:

```
Add-DAClient -SecurityGroupNameList @('corp.contoso.com\DirectAccessClients')
Remove-DAClient -SecurityGroupNameList @('corp.contoso.com\Domain Computers')
```

To enable Remote Access for all computers (not only notebooks and laptops), and to enable Remote Access for Windows 7 clients:

```
Set-DAClient -OnlyRemoteComputers 'Disabled' -Downlevel 'Enabled'
```

To configure the DirectAccess client experience, including the friendly connection name and the web probe URL:

```
Set-DAClientExperienceConfiguration -FriendlyName 'Contoso DirectAccess Connection' -
PreferLocalNamesAllowed $False -PolicyStore 'corp.contoso.com\DirectAccess Client
Settings' -CorporateResources @('HTTP:http://directaccess-WebProbeHost.corp.contoso.com')
```

See also

- Step 3: Verify the Deployment
- Step 1: Configure the Remote Access Infrastructure
- Install and Configure Advanced Remote Access

- Deploy a Single Remote Access Server with Advanced Settings

Step 3: Verify the Deployment

This topic describes how to verify that you have correctly configured your DirectAccess deployment.

To verify access to internal resources through DirectAccess

1. Connect a DirectAccess client computer to the corporate network and obtain the Group Policy Object.
2. Click the **Network connections** icon in the notification area to access the DirectAcess Media Manager.
3. Click **DirectAccess Connection**, and you will see that the status is **Connected Locally**.
4. Connect the client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

Deploy Remote Access in an Enterprise

Remote Access in Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single role. This overview provides an introduction to the main enterprise deployment scenarios for Remote Access.

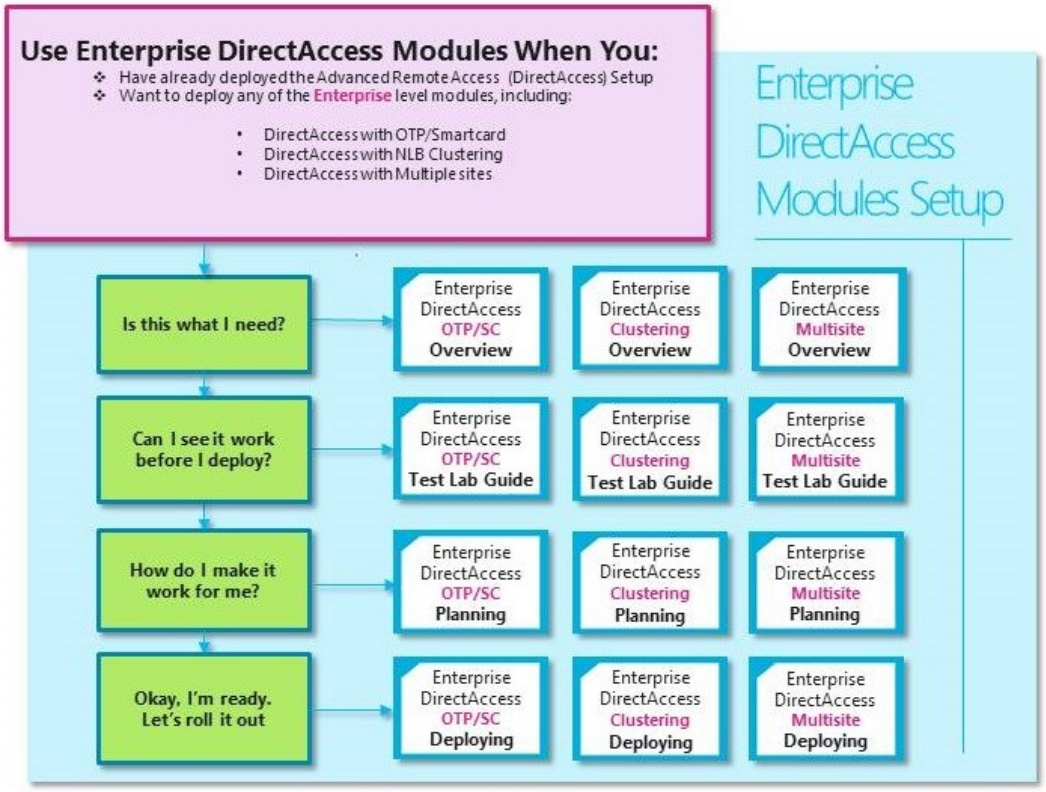

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- **Deploy Remote Access with Network Access Protection**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Scenario description

Remote access includes a number of enterprise features, including deploying multiple Remote Access servers in a cluster load balanced with Windows Network Load Balancing (NLB) or an external load balancer, setting up a multisite deployment with Remote Access servers situated in dispersed geographical locations, and deploying DirectAccess with two-factor client authentication using a one-time password (OTP).

In this scenario

Each enterprise scenario is described in a document that includes planning and deployment instructions. For more information, see:

- Deploy Remote Access in a Cluster
- Deploy Multiple Remote Access Servers in a Multisite Deployment
- Deploy Remote Access with OTP Authentication
- Deploy Remote Access in a Multi-Forest Environment
- **Deploy Remote Access with Network Access Protection**

Practical applications

Remote access enterprise scenarios provide the following:

- **Increased availability**—Deploying multiple Remote Access servers in a cluster provides scalability and increases the capacity for throughput and number of users. Load balancing the cluster provides high availability. If a server in the cluster fails, remote users can continue to access the internal corporate network via a different server in the cluster. Failover is transparent as clients connect to the cluster using a virtual IP (VIP) address.
- **Ease-of-management**—A cluster or multisite deployment can be configured and managed as a single entity using the Remote Access Management console running on one of the cluster servers. In addition, a multisite deployment allows administrators to align Remote Access deployment to Active Directory sites, providing a simplified architecture. Shared settings can easily be set across cluster servers or on all multisite entry point servers. Remote Access settings can be managed from any of the servers in the cluster or deployment, or remotely using Remote Server Administration Tools (RSAT). In addition, the entire cluster or multisite deployment can be monitored from a single Remote Access Management console.
- **Cost efficiency**—A Remote Access multisite deployment allows enterprises to deploy Remote Access servers in multiple sites corresponding to client locations. This provides a predictable access experience for remote clients regardless of location, and reduces costs and intranet bandwidth by routing client traffic over the Internet to the closest Remote Access server.
- **Security**—Deploying strong client authentication with a one-time password (OTP) instead of standard Active Directory password increases security.

Roles and features included in this scenario

The following table lists the roles and features used in the enterprise scenario.

Role/feature	How it supports this scenario
Remote Access server role	The role is installed and uninstalled using the Server Manager console. This role encompasses both DirectAccess, which was

Role/feature	How it supports this scenario
	previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. The Remote Access Server Role is dependent on the following server features: • Internet Information Services (IIS) – This feature is required to configure the network location server and default web probe. • Group Policy Management Console feature – feature is required by DirectAccess to create and manage the Group Policy Objects (GPOs) in Active Directory and must be installed as a required feature for the server role.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: 1. Remote Access GUI and Command Line Tools 2. Remote Access module for Windows PowerShell

Role/feature	How it supports this scenario
	Dependencies include: 1. Group Policy Management Console 2. RAS Connection Manager Administration Kit (CMAK) 3. Windows PowerShell 3.0 4. Graphical Management Tools and Infrastructure
Windows NLB	This feature allows the load balancing of multiple Remote Access servers.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Demonstrate DirectAccess in a cluster with NLB Demonstrate a DirectAccess multisite deployment Demonstrate a DirectAccess multisite deployment
Deployment	Remote Access
Tools and settings	Remote Access PowerShell cmdlets
Community resources	RRAS Product Team blog | Remote Access TechNet Forum DirectAccess Wiki entries
Related technologies	How IPv6 works

Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB

Remote Access is a server role in the Windows Server® 2012 operating system that enables remote users to securely access internal network resources using DirectAccess or RRAS VPN. This guide contains step-by-step instructions for extending the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 to demonstrate DirectAccess Network Load Balancing and cluster configuration.

About this guide

This guide contains instructions for configuring and demonstrating Remote Access using six servers and two client computers. The completed Remote Access test lab with NLB simulates an intranet, the Internet, and a home network and demonstrates Remote Access functionality in different Internet connection scenarios.

Important
This lab is a proof of concept using the minimum number of computers. The configuration detailed in this guide is for test lab purposes only, and is not to be used in a production environment.

Known issues

The following are known issues when configuring a cluster scenario:

- After configuring DirectAccess in an IPv4-only deployment with a single network adapter, and after the default DNS64 (the IPv6 address which contains ":3333::") is automatically configured on the network adapter, attempting to enable load-balancing via the Remote Access Management console causes a prompt for the user to supply an IPv6 DIP. If an IPv6 DIP is supplied, the configuration fails after clicking **Commit** with the error: The parameter is incorrect.

  To resolve this issue:

  a. Download the backup and restore scripts from Back up and Restore Remote Access Configuration.

  b. Back up your Remote Access GPOs using the downloaded script Backup-RemoteAccess.ps1

  c. Attempt to enable load balancing until the step at which it fails. On the Enable Load Balancing dialog box, expand the details area, right-click in the details area, and then click **Copy Script**.

  d. Open Notepad, and paste the contents of the clipboard. For example:

```
Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress
@('10.244.4.19/255.255.255.0','fdc4:29bd:abde:3333::2/128') -
InternetVirtualIPAddress @('fdc4:29bd:abde:3333::1/128',
```

```
'10.244.4.21/255.255.255.0') -ComputerName
'DA1.domain1.corp.contoso.com' -Verbose
```

e. Close any open Remote Access dialog boxes and close the Remote Access Management console.

f. Edit the pasted text and remove the IPv6 addresses. For example:

```
Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress
@('10.244.4.19/255.255.255.0') -InternetVirtualIPAddress
@('10.244.4.21/255.255.255.0') -ComputerName
'DA1.domain1.corp.contoso.com' -Verbose
```

g. In an elevated PowerShell window, run the command from the previous step.

h. If the cmdlet fails while it is running (not due to incorrect input values), run the command Restore-RemoteAccess.ps1 and follow instructions to make sure that the integrity of your original configuration is maintained.

i. You can now open the Remote Access Management console again.

Overview of the Test Lab Scenario

In this test lab scenario, DirectAccess is deployed with:

- **DC1**—A server that is configured as a domain controller, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server.
- **EDGE1**—A server on the internal network that is configured as the first Remote Access server in a Remote Access server cluster. This server has two network adapters; one connected to the internal network, and the other connected to the external network.
- **EDGE2**—A server on the internal network that is configured as the second Remote Access server in a Remote Access server cluster. This server has two network adapters; one connected to the internal network, and the other connected to the external network.
- **APP1**—A server on the internal network that is configured as a web and file server, and as an enterprise root certification authority (CA)
- **APP2**—A computer on the internal network that is configured as an IPv4 only web and file server. This computer is used to highlight the NAT64/DNS64 capabilities. This guide uses Windows Server 2003 SP2; however, you can use any computer that is configured as an IPv4 only computer.
- **INET1**—A server that is configured as an Internet DNS and DHCP server.
- **NAT1**—A client computer that is configured as a network address translator (NAT) device using Internet Connection Sharing.
- **CLIENT1**—A client computer that is configured as a DirectAccess client that will be used to test DirectAccess connectivity when moving between the internal network, the simulated Internet, and a home network.

The test lab consists of three subnets that simulate the following:

- A home network named Homenet (192.168.137.0/24) connected to the Internet by a NAT.
- The external network represented by the Internet subnet (131.107.0.0/24).
- An internal network named Corpnet (10.0.0.0/24; 2001:db8:1::/64) separated from the Internet by the Remote Access server.

Computers on each subnet connect using either a physical or virtual hub or switch, as shown in the following figure.

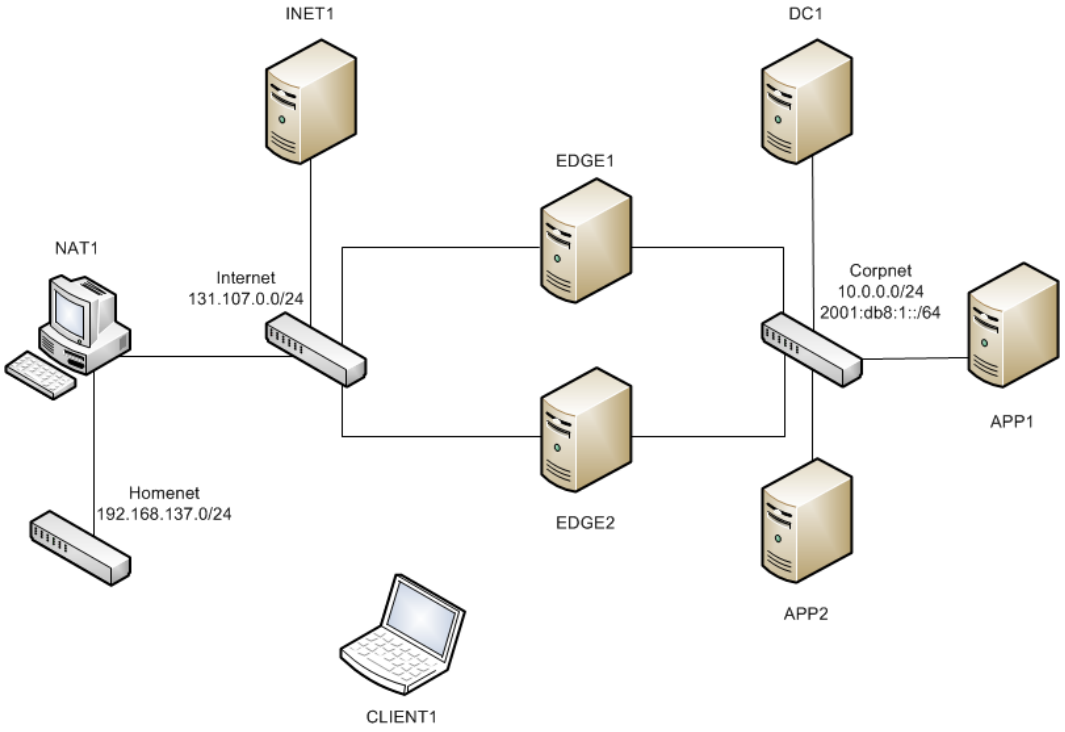

Configuration Requirements

The following components are required for configuring DirectAccess in the test lab:

- The product disc or files for Windows Server® 2012.
- Five computers or virtual machines that meet the minimum hardware requirements for Windows Server 2012; two of these computers have two network adapters installed.
- The product disc or files for Windows® 8.
- Two computers or virtual machines that meet the minimum hardware requirements for Windows 8; one of these computers has two network adapters installed.
- The product disc or files for Windows Server 2003 Enterprise SP2.
- One computer or virtual machine that meets the minimum hardware requirements for Windows Server 2003 SP2.

Steps for Configuring the Test Lab

The following steps describe how to configure the Remote Access infrastructure, configure the Remote Access servers and clients, and test DirectAccess connectivity from the Internet and Homenet subnets.

In this test lab guide you will build a Network Load Balancing (NLB) enabled Remote Access cluster by performing the following steps:

- STEP 1: Complete the DirectAccess Configuration—Complete all the steps in the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6.
- STEP 2: Configure EDGE1—Configure the Remote Access role on EDGE1 for load balancing.
- STEP 3: Install and Configure EDGE2—EDGE2 acts as the second Remote Access server in a Remote Access cluster.
- STEP 4: Create the Network Load Balanced Remote Access Cluster—EDGE1 is configured as the first server in a Remote Access cluster. EDGE2 is joined to the cluster and NLB is configured for the cluster.
- STEP 5: Test DirectAccess Connectivity from the Internet and Through the Cluster—After NLB and cluster configuration is complete you can test DirectAccess client connectivity through the load balanced cluster.
- STEP 6: Test DirectAccess Client Connectivity from Behind a NAT Device—Move the client computer behind a NAT device to simulate testing DirectAccess client connectivity from behind a home router.
- STEP 7: Test Connectivity When Returning to the Corpnet—Make sure that the client computer can still access corporate resources when returning to Corpnet.
- STEP 8: Snapshot the Configuration—After completing the test lab, take a snapshot of the working Remote Access NLB cluster so that you can return to it later to test additional scenarios.

STEP 1: Complete the DirectAccess Configuration

The first step is to complete all the steps in the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6. If you have already completed the steps in this test lab guides and saved a snapshot or disk image of the test lab, you can restore the snapshot or image and begin with the next step.

STEP 2: Configure EDGE1

The following procedure is performed on the DirectAccess server:

- Configure DirectAccess on EDGE1—Configure DirectAccess for your organization.

Configure DirectAccess on EDGE1

1. On the **Start** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the left pane, click **Configuration**.
3. In the middle pane of the console, in the **Step 2 Remote Access Server** area, click **Edit**.
4. In the **Remote Access Server Setup** wizard, click **Prefix Configuration**. On the **Prefix Configuration** page, in **IPv6 prefix assigned to DirectAccess client computers**, enter **2001:db8:1:1000::/59**, and then click **Next**.
5. Click **Finish**.
6. In the middle pane of the console, click **Finish**.
7. On the **Remote Access Review** dialog box, review the configuration settings, and then click **Apply**. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.

STEP 3: Install and Configure EDGE2

EDGE2 is the second member of a Remote Access cluster. EDGE2 is installed and configured before enabling the cluster configuration.

You will perform the following steps to configure EDGE2:

1. Install the operating system on EDGE2—Install Windows Server® 2012 on EDGE2.
2. Configure TCP/IP properties—After installing the operating system on EDGE2, configure static IP addressing information on its internal and external network adapters.
3. Rename EDGE2 and join it to the domain—Rename the computer to EDGE2 and join it to the CORP domain.
4. Install the IP-HTTPS certificate—To accept incoming IP-HTTPS requests, the EDGE2 Remote Access cluster member must have an IP-HTTPS website certificate installed in its machine certificate store. The certificate must have the same name as the IP-HTTPS certificate on EDGE1.
5. Install the Remote Access role on EDGE2—Before you can join EDGE2 to the Remote Access cluster, you must install the Remote Access server role on EDGE2.

Install the operating system on EDGE2

1. On EDGE2, start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Log on using the local Administrator account.

3. Connect EDGE2 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012, and then disconnect from the Internet.
4. Connect one network adapter to the Corpnet subnet or the virtual switch representing the corpnet subnet and the other to the Internet subnet or virtual switch representing the Internet subnet.

Configure TCP/IP properties

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In the **Network Connections** window, right-click the network connection that is connected to the Corpnet subnet or virtual switch, and then click **Rename**.
3. Type **Corpnet**, and then press ENTER.
4. Right-click **Corpnet**, and then click **Properties**.
5. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
6. Click **Use the following IP address**. In **IP address**, type **10.0.0.8**. In **Subnet mask**, type **255.255.255.0**.
7. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
8. Click **Advanced**, and then click the **DNS** tab.
9. In **DNS suffix for this connection**, type **corp.contoso.com**, click **OK** twice.
10. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
11. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:1::8**. In **Subnet prefix length**, type **64**.
12. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **2001:db8:1::1**.
13. Click **Advanced**, and then click the **DNS** tab.
14. In **DNS suffix for this connection**, type **corp.contoso.com**, click **OK** twice, and then click **Close**.
15. In the **Network Connections** window, right-click the network connection that is connected to the Internet subnet, and then click **Rename**.
16. Type **Internet**, and then press ENTER.
17. Right-click **Internet**, and then click **Properties**.
18. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
19. Click **Use the following IP address**. In **IP address**, enter **131.107.0.8**. In **Subnet mask**, enter **255.255.255.0**.
20. Click the **DNS** tab
21. In **DNS suffix for this connection**, type **isp.example.com**, and then click **OK** twice, and then click **Close**.
22. Close the **Network Connections** window.
23. To check network communication between EDGE2 and DC1, click **Start**, type **cmd**, and

then press ENTER.

24. In the Command Prompt window, type **ping dc1.corp.contoso.com** and press ENTER. Verify that there are four responses from `10.0.0.1` or the IPv6 address `2001:db8:1::1`
25. Close the Command Prompt window.

Rename EDGE2 and join it to the domain

1. In the Server Manager console, in **Local Server**, in the **Properties** area, next to **Computer name**, click the link.
2. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
3. On the **Computer Name/Domain Changes** dialog box, in the **Computer name** box, type **EDGE2**. In the **Member of** area, click **Domain**, and in the text box, enter **corp.contoso.com**, and then click **OK**.
4. When you are prompted for a user name and password, type **User1** and its password, and then click **OK**.
5. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.
9. After restarting, login as CORP\User1.

Install the IP-HTTPS certificate

1. On the **Start** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Finish**, and then click **OK**.
4. In the left pane of the console, navigate to **Certificates (Local Computer)\Personal\Certificates**. Right click the **Certificates** node, point to **All Tasks**, and then click **Request New Certificate**.
5. On the Certificate Enrollment wizard, click **Next** twice.
6. On the **Request Certificates** page, select the **Web Server** check box, and then click **More information is required to enroll for this certificate**.
7. On the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in the **Type** list, click **Common Name**.
8. In **Value**, type **edge1.contoso.com**, and then click **Add**.
9. In the **Alternative name** area, in the **Type** list, click **DNS**.
10. In **Value**, type **edge1.contoso.com**, and then click **Add**.
11. On the **General** tab, in **Friendly name**, type **IP-HTTPS Certificate**.
12. Click **OK**, click **Enroll**, and then click **Finish**.

13. In the details pane of the Certificates snap-in, verify that a new certificate with the name edge1.contoso.com was enrolled with Intended Purposes of Server Authentication.
14. Close the console window. If you are prompted to save settings, click **No**.

Install the Remote Access role on EDGE2

1. In the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. On the **Select server roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

STEP 4: Create the Network Load Balanced Remote Access Cluster

Windows Server® 2012 enables you to create clusters of Remote Access servers. A cluster acts as a single logical server and provides centralized configuration and management for the servers in the cluster. When using Network Load Balancing (NLB) there is support for up to 8 Remote Access members in a single cluster. Remote Access clusters provide high availability and load balancing of connections from DirectAccess clients to the internal network.

The following procedures enable you to create and test a Remote Access cluster:

1. Install the Network Load Balancing feature on EDGE1 and EDGE2—Before enabling load balancing, you must install the Network Load Balancing feature on both EDGE1 and EDGE2.
2. Enable load balancing on EDGE1—EDGE1 was originally installed in single server mode. To enable load balancing, you configure new external and internal dedicated IP addresses (DIPs) for EDGE1. The previous DIPs on EDGE1 are automatically configured as virtual IP addresses (VIPs) for the cluster. The new external DIP is 131.107.0.10, the new internal IPv4 DIP is 10.0.0.10, the new internal IPv6 DIP is 2001:db8:1::10. The cluster VIPs are 131.107.0.2 and 131.107.0.3 (external), and 10.0.0.2 and 2001:db8:1::2 (internal).
3. Add EDGE2 to the load balanced cluster—After enabling load balancing, you can now add EDGE2 to the cluster to provide load balancing and high availability for DirectAccess client connections.

Prerequisites

If you are creating this test lab on virtual machines, you must enable MAC address spoofing on EDGE1 and EDGE2.

Enable MAC address spoofing on EDGE1 and EDGE2

1. Perform a graceful shutdown on EDGE1 and EDGE2.
2. On the machine hosting your virtual machines, in **Hyper-V Manager**, right-click EDGE1, and then click **Settings**.
3. On the **Settings** dialog box, in the **Hardware** list, click the network adapter connected to the corpnet network, and then in the details pane, select the **Enable spoofing of MAC addresses** check box.
4. In the **Hardware** list, click the network adapter connected to the Internet network, and then in the details pane, select the **Enable spoofing of MAC addresses** check box.
5. On the **Settings** dialog box, click **OK**.
6. Repeat this procedure from step 2 on EDGE2.

Install the Network Load Balancing feature on EDGE1 and EDGE2

To configure EDGE1 and EDGE2 in a cluster, you must install the Network Load Balancing feature on both EDGE1 and EDGE2.

To install Network Load Balancing

1. On EDGE1, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** four times to get to the server feature selection screen.
3. On the **Select features** dialog, select **Network Load Balancing**, click **Add Features**, click **Next**, and then click **Install**.
4. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.
5. Repeat this procedure on EDGE2.

Enable load balancing on EDGE1

Use this procedure to enable load balancing and configure the new DIPs on EDGE1.

Enable load balancing

1. On EDGE1, click **Start**, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the left pane, click **Configuration**, and then in the **Tasks** pane, click **Enable Load Balancing**.
3. In the Enable Load Balancing Wizard, click **Next**.

4. On the **Load Balancing Method** page, click **Use Windows Network Load Balancing (NLB)**, and then click **Next**.
5. On the **External Dedicated IP Addresses** page, in the **IPv4 address** box, type **131.107.0.10**, in the **Subnet mask** box, verify the subnet prefix is **255.255.255.0**, and then click **Next**.
6. On the **Internal Dedicated IP Addresses** page, do the following, and then click **Next**:
    a. In the **IPv4 address** box, type **10.0.0.10** and in the **Subnet mask** box, verify the subnet prefix is **255.255.255.0**.
    b. In the **IPv6 address** box, type **2001:db8:1::10** and in the Subnet prefix length, verify the value is **64**.
7. On the **Summary** page, click **Commit**.
8. On the **Enable Load Balancing** dialog box, click **Close**.
9. In the Enable Load Balancing Wizard, click **Close**.

Add EDGE2 to the load balanced cluster

Use this procedure to add EDGE2 to the NLB cluster.

Note
You should wait two minutes after completing the previous steps before proceeding. After enabling NLB, the RAConfigTask runs and configures the machine with NLB settings. This might take a few minutes to complete, and if the administrator runs another NLB related configuration before the task ends, that configuration will fail.

Add EDGE2 to the cluster

1. On the EDGE1 computer or virtual machine, in the Remote Access Management Console, in the **Tasks** pane, under **Load Balanced Cluster**, click **Add or Remove Servers**.
2. On the **Add or Remove Servers** dialog box, click **Add Server**.
3. On the **Add a Server** wizard, on the **Select Server** page, type **EDGE2**, and then click **Next**.
4. On the **Network Adapters** page, in **External adapter**, make sure that **Internet** is selected, and in **Internal adapter**, make sure that **Corpnet** is selected. Click **Browse**, on the **Windows Security** dialog box, make sure that **IP-HTTPS Certificate** is selected, click **OK**, and then click **Next**.
5. On the **Summary** page, click **Add**.
6. On the **Completion** page, click **Close**.
7. On the **Add or Remove Servers** dialog box, click **Commit**.
8. On the **Adding and Removing Servers** dialog box, click **Close**.
9. On the **Start** screen, type **nlbmgr.exe**, and the press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

10. In the **Network Load Balancing Manager**, click **Internal DA cluster**. In the details pane, make sure that both **EDGE1(Corpnet)** and **EDGE2(Corpnet)** have the status **Converged**.
11. If a server is not **Converged**, in the console tree, right-click the server, point to **Control Host**, and then click **Start**.
12. In the **Network Load Balancing Manager**, click **Internet DA cluster**. Make sure that in the details pane, both **EDGE1(Internet)** and **EDGE2(Internet)** have the status **Converged**.
13. If a server is not **Converged**, in the console tree, right-click the server, point to **Control Host**, and then click **Start**.

STEP 5: Test DirectAccess Connectivity from the Internet and Through the Cluster

CLIENT1 is now ready for DirectAccess testing.

- Test DirectAccess connectivity from the Internet—Connect CLIENT1 to the simulated Internet. When connected to the simulated Internet, the client is assigned public IPv4 addresses. When a DirectAccess client is assigned a public IPv4 address, it tries to establish a connection to the Remote Access server using an IPv6 transition technology.
- Test DirectAccess client connectivity through the cluster—Test cluster functionality. Before you begin testing, we recommend that you shut down both EDGE1 and EDGE2 for at least five minutes. There are a number of reasons for this, which include ARP cache timeouts and changes related to NLB. When validating NLB configuration in a test lab, you will need to be patient as changes in configuration will not be immediately reflected in connectivity until after a period of time has elapsed. This is important to keep in mind when you carry out the following tasks.

  Tip
  We recommend that you clear the Internet Explorer cache before performing this procedure and each time you test the connection through a different Remote Access server to make sure that you are testing the connection and not retrieving the webpages from the cache.

Test DirectAccess connectivity from the Internet

1. Unplug CLIENT1 from the corpnet switch and connect it to the Internet switch. Wait for 30 seconds.
2. In an elevated Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This flushes name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the corpnet.
3. In the Windows PowerShell window, type **Get-DnsClientNrptPolicy** and press ENTER.

The output shows the current settings for the Name Resolution Policy Table (NRPT). These settings indicate that all connections to .corp.contoso.com should be resolved by the Remote Access DNS server, with the IPv6 address 2001:db8:1::2. Also, note the NRPT entry indicating that there is an exemption for the name nls.corp.contoso.com; names on the exemption list are not answered by the Remote Access DNS server. You can ping the Remote Access DNS server IP address to confirm connectivity to the Remote Access server; for example, you can ping 2001:db8:1::2.

4. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address for APP1, which in this case is `2001:db8:1::3`.
5. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is `fdc9:9f4e:eb1b:7777::a00:4`.

   The ability to ping APP2 is important, because success indicates that you were able to establish a connection using NAT64/DNS64, as APP2 is an IPv4 only resource.
6. Leave the Windows PowerShell window open for the next procedure.
7. Open Internet Explorer, in the Internet Explorer address bar, enter **http://app1/** and press ENTER. You will see the default IIS website on APP1.
8. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP2.
9. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file.

   This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource in the resource domain.
10. On the **Start** screen, type **wf.msc**, and then press ENTER.
11. In the **Windows Firewall with Advanced Security** console, notice that only the **Private** or **Public Profile** is active. The Windows Firewall must be enabled for DirectAccess to work correctly. If the Windows Firewall is disabled, DirectAccess connectivity does not work.
12. In the left pane of the console, expand the **Monitoring** node, and click the **Connection Security Rules** node. You should see the active connection security rules: **DirectAccess Policy-ClientToCorp**, **DirectAccess Policy-ClientToDNS64NAT64PrefixExemption**, **DirectAccess Policy-ClientToInfra**, and **DirectAccess Policy-ClientToNlaExempt**. Scroll the middle pane to the right to show the **1st Authentication Methods** and **2nd Authentication Methods** columns. Notice that the first rule (ClientToCorp) uses Kerberos V5 to establish the intranet tunnel and the third rule (ClientToInfra) uses NTLMv2 to establish the infrastructure tunnel.
13. In the left pane of the console, expand the **Security Associations** node, and click the **Main Mode** node. Notice the infrastructure tunnel security associations using NTLMv2 and the intranet tunnel security association using Kerberos V5. Right-click the entry that shows **User (Kerberos V5)** as the **2nd Authentication Method** and click **Properties**. On the **General** tab, notice the **Second authentication Local ID** is **CORP\User1**, indicating that User1 was able to successfully authenticate to the CORP domain using Kerberos.

Test DirectAccess client connectivity through the cluster

1. Perform a graceful shutdown on EDGE2.

   You can use the Network Load Balancing Manager to view the status of the servers when running these tests.
2. On CLIENT1, in the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This flushes name resolution entries that may still exist in the client DNS cache.
3. In the Windows PowerShell window, ping APP1, and APP2. You should receive replies from both of these resources.
4. On the **Start** screen, type **\\app2\files**. You should see the shared folder on the APP2 computer. The ability to open the file share on APP2 indicates that the second tunnel, which requires Kerberos authentication for the user, is working correctly.
5. Open Internet Explorer, and then open the websites http://app1/ and http://app2/. The ability to open both websites confirms that both the first and second tunnels are up and functioning. Close Internet Explorer.
6. Start the EDGE2 computer.
7. On EDGE1 perform a graceful shutdown.
8. Wait for 5 minutes, and then return to CLIENT1. Perform steps 2-5. This confirms that CLIENT1 was able to transparently fail over to EDGE2 after EDGE1 became unavailable.

STEP 6: Test DirectAccess Client Connectivity from Behind a NAT Device

When a DirectAccess client is connected to the Internet from behind a NAT device or a web proxy server, the DirectAccess client uses either Teredo or IP-HTTPS to connect to the Remote Access server. If the NAT device enables outbound UDP port 3544 to the Remote Access server's public IP address, then Teredo is used. If Teredo access is not available, the DirectAccess client falls back to IP-HTTPS over outbound TCP port 443, which enables access through firewalls or web proxy servers over the traditional SSL port. If the web proxy requires authentication, the IP-HTTPS connection will fail. IP-HTTPS connections also fail if the web proxy performs outbound SSL inspection, due to the fact that the HTTPS session is terminated at the web proxy instead of the Remote Access server. In this section you will perform the same tests performed when connecting using a 6to4 connection in the previous section.

The following procedures are performed on both client computers:

1. Test Teredo connectivity—The first set of tests are performed when the DirectAccess client is configured to use Teredo. This is the automatic setting when the NAT device allows outbound access to UDP port 3544.

2. Test IP-HTTPS connectivity—The second set of tests are performed when the DirectAccess client is configured to use IP-HTTPS. In order to demonstrate IP-HTTPS connectivity, Teredo is disabled on the client computers.

Tip

We recommend that you clear the Internet Explorer cache before performing these procedures to make sure that you are testing the connection and not retrieving the website pages from the cache.

Prerequisites

Before performing these tests, unplug CLIENT1 from the Internet switch and connect it to the Homenet switch. If asked what type of network you want to define the current network, select **Home Network**.

Start EDGE1 and EDGE2 if they are not already running.

Test Teredo connectivity

1. On CLIENT1, open an elevated Windows PowerShell window, type **ipconfig /all** and press ENTER.
2. Examine the output of the ipconfig command.

   CLIENT1 is now connected to the Internet from behind a NAT device and is assigned a private IPv4 address. When the DirectAccess client is behind a NAT device and assigned a private IPv4 address, the preferred IPv6 transition technology is Teredo. If you look at the output of the ipconfig command, you should see a section for `Tunnel adapter Teredo Tunneling Pseudo-Interface` and then a description `Microsoft Teredo Tunneling Adapter`, with an IP address that starts with 2001: consistent with being a Teredo address. If you do not see the Teredo section, enable Teredo with the following command: **netsh interface Teredo set state enterpriseclient** and then rerun the ipconfig command. You will not see a default gateway listed for the Teredo tunnel adapter.
3. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER.

   This will flush name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the Internet.
4. In the Windows PowerShell window, type **Get-DnsClientNrptPolicy** and press ENTER.

   The output shows the current settings for the Name Resolution Policy Table (NRPT). These settings indicate that all connections to .corp.contoso.com should be resolved by the Remote Access DNS server, with the IPv6 address of 2001:db8:1::2. Also, note the NRPT entry indicating that there is an exemption for the name nls.corp.contoso.com; names on the exemption list are not answered by the Remote Access DNS server. You can ping the Remote Access DNS server IP address to confirm connectivity to the Remote Access server; for example, you can ping 2001:db8:1::2 in this example.
5. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address of APP1, 2001:db8:1::3.

6. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fdc9:9f4e:eb1b:7777::a00:4.
7. Leave the Windows PowerShell window open for the next procedure.
8. Open Internet Explorer, in the Internet Explorer address bar, enter **http://app1/** and press ENTER. You will see the default IIS website on APP1.
9. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP2.
10. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file. This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource on an IPv4 only host.

Test IP-HTTPS connectivity

1. Open an elevated Windows PowerShell window, type **netsh interface teredo set state disabled** and press ENTER. This disables Teredo on the client computer and enables the client computer to configure itself to use IP-HTTPS. An **Ok** response appears when the command completes.
2. In the Windows PowerShell window, type **ipconfig /all** and press ENTER.
3. Examine the output of the ipconfig command. This computer is now connected to the Internet from behind a NAT device and is assigned a private IPv4 address. Teredo is disabled and the DirectAccess client falls back to IP-HTTPS. When you look at the output of the ipconfig command, you see a section for `Tunnel adapter iphttpsinterface` with an IP address that starts with 2001:db8:1:100 consistent with this being an IP-HTTPS address based on the prefix that was configured when setting up DirectAccess. You will not see a default gateway listed for the IP-HTTPS tunnel adapter.
4. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This will flush name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the corpnet.
5. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address of APP1, 2001:db8:1::3.
6. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fdc9:9f4e:eb1b:7777::a00:4.
7. Open Internet Explorer, in the Internet Explorer address bar, enter **http://app1/** and press ENTER. You will see the default IIS site on APP1.
8. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP2.
9. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file. This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource on an IPv4 only host.

STEP 7: Test Connectivity When Returning to the Corpnet

Many of your users will move between remote locations and the corpnet, so it's important that when they return to the corpnet that they are able to access resources without having to make any configuration changes. Remote Access makes this possible because when the DirectAccess client returns to the corpnet, it is able to make a connection to the network location server. Once the HTTPS connection is successfully established to the network location server, the DirectAccess client disables the DirectAccess client configuration and uses a direct connection to corpnet.

Test connectivity on CLIENT1

1. Shut down CLIENT1 and then unplug CLIENT1 from the Homenet subnet or virtual switch and connect it to the Corpnet subnet or virtual switch. Turn on CLIENT1, and log on as CORP\User1.
2. Open an elevated Windows PowerShell window, type **ipconfig /all**, and press ENTER. The output will indicate that CLIENT1 has a local IP address, and that there is no active 6to4, Teredo, or IP-HTTPS tunnel.
3. Test connectivity to the network share on APP2. On the **Start** screen, type **\\APP2\Files**, and then press ENTER. You will be able to open the file in that folder.

STEP 8: Snapshot the Configuration

This completes the DirectAccess test lab. To save this configuration so that you can quickly return to a working DirectAccess with NLB cluster configuration from which you can test other DirectAccess modular test lab guides, test lab guide extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots DirectAccess cluster and NLB. If your lab uses physical computers, create disk images to save the DirectAccess test lab configuration.

Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment

Remote Access is a server role in the Windows Server® 2012 operating system that enables remote users to securely access internal network resources using DirectAccess or RRAS VPN. This guide contains step-by-step instructions for extending the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 to demonstrate Remote Access in a multisite scenario.

Deploying Remote Access in a multisite scenario enables you to configure Remote Access servers in geographically diverse locations. Previously, remote users were required to always connect to the corporate network through a particular DirectAccess server. With Windows Server 2012 and Windows® 8, you can configure entry points for each geographic location in your deployment. Each entry point can be a single Remote Access server or a cluster of Remote Access servers. Remote users have the option to connect to any of the organization's Remote Access entry points. For example, if a remote user usually connects to the Remote Access entry point located in Asia, but then goes on a business trip to Europe, the client computer automatically connects to the closest Remote Access entry point.

About this guide

This guide contains instructions for configuring and demonstrating Remote Access using nine servers and three client computers. The completed Remote Access multisite test lab simulates an intranet, the Internet, and a home network and demonstrates Remote Access functionality in different Internet connection scenarios.

Important
This lab is a proof of concept using the minimum number of computers. The configuration detailed in this guide is for test lab purposes only, and is not to be used in a production environment.

Overview of the Test Lab Scenario

In this test lab scenario, DirectAccess is deployed with:

- **DC1**—A server that is configured as a domain controller, DNS server, and DHCP server for the corp.contoso.com domain.
- **2-DC1**—A server that is configured as a domain controller and DNS server for the corp2.corp.contoso.com domain.
- **EDGE1 and 2-EDGE1**—Two servers on the internal network that are configured as Remote Access servers. Each server has two network adapters; one connected to the internal network, and the other connected to the external network.

- **APP1 and 2-APP1**—Two servers on the internal network that are configured as web and file servers.
- **APP2**—A computer on the internal network that is configured as an IPv4 only web and file server. This computer is used to highlight the NAT64/DNS64 capabilities. This guide uses Windows Server 2003 SP2; however, you can use any computer that is configured as an IPv4 only computer.
- **ROUTER1**—A server that is configured to provide routing between the two corporate internal networks.
- **INET1**—A server that is configured as an Internet DNS and DHCP server.
- **NAT1**—A client computer that is configured as a network address translator (NAT) device using Internet Connection Sharing.
- **CLIENT1 and CLIENT2**—Two client computers that are configured as DirectAccess clients that will be used to test DirectAccess connectivity when moving between the internal network, the simulated Internet, and a home network. **CLIENT2** is a Windows® 7 client.

The test lab consists of four subnets that simulate the following:

- A home network named Homenet (192.168.137.0/24) connected to the Internet by a NAT.
- The external network represented by the Internet subnet (131.107.0.0/24).
- An internal network named Corpnet (10.0.0.0/24; 2001:db8:1::/64) separated from the Internet by the EDGE1 Remote Access server.
- An internal network named 2-Corpnet1 (10.2.0.0/24; 2001:db8:2::/64) separated from the Internet by the 2-EDGE1 Remote Access server.

Computers on each subnet connect using either a physical or virtual hub or switch, as shown in the following figure.

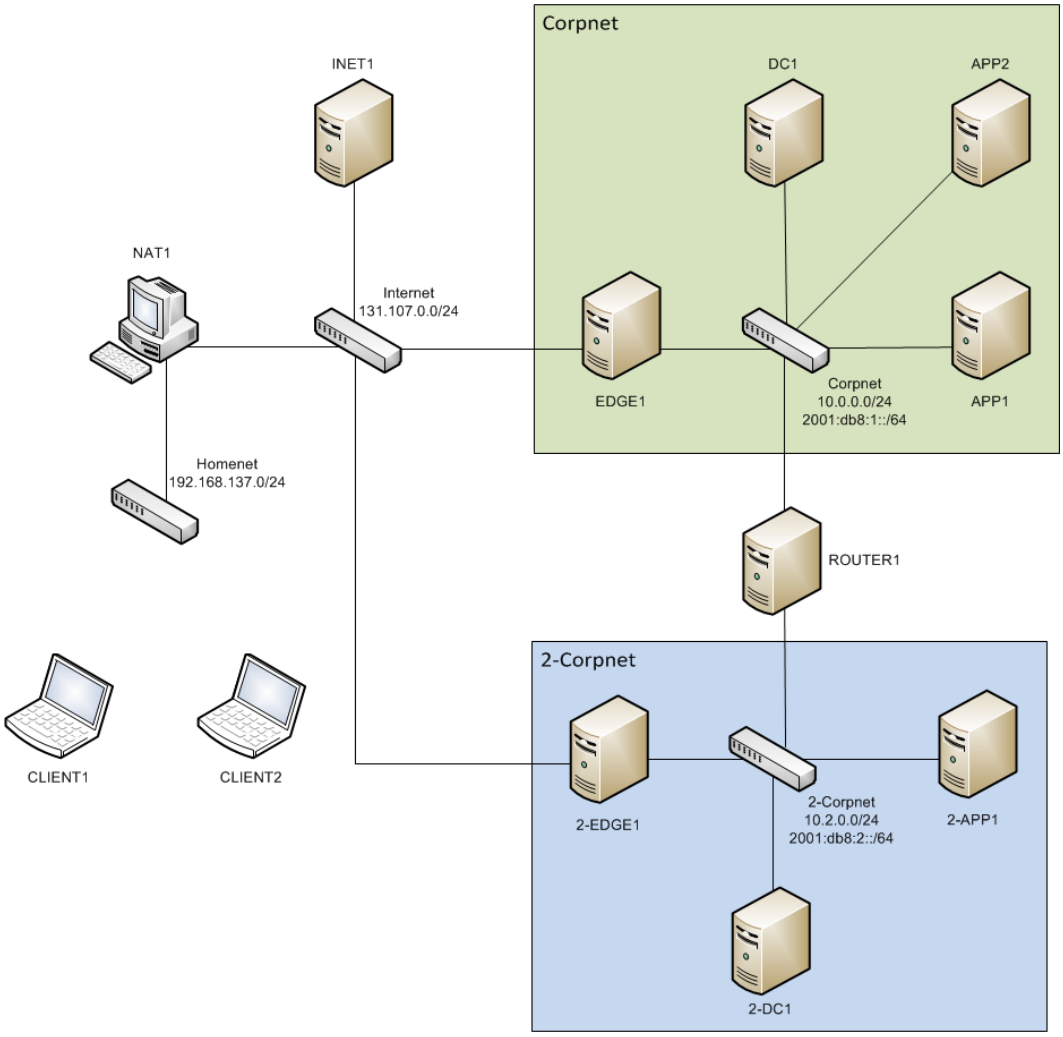

Configuration Requirements

The following components are required to configure Remote Access in the test lab:

- The product disc or files for Windows Server® 2012.
- Eight computers or virtual machines that meet the minimum hardware requirements for Windows Server 2012; three of these computers have two network adapters installed.
- The product disc or files for Windows® 8.
- The product disc or files for Windows 7 Ultimate.
- Three computers or virtual machines that meet the minimum hardware requirements for Windows 7 or Windows 8; one of these computers has two network adapters installed.
- The product disc or files for Windows Server 2003 Enterprise SP2.

- One computer or virtual machine that meets the minimum hardware requirements for Windows Server 2003 SP2.

Steps for Configuring the Test Lab

The following steps describe how to configure the Remote Access infrastructure, configure the Remote Access servers and clients and test DirectAccess connectivity from the Internet and Homenet subnets.

In this test lab guide you will build a multisite Remote Access deployment by performing the following steps:

- STEP 1: Complete the DirectAccess Configuration—Complete all the steps in the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6.
- STEP 2: Install and Configure ROUTER1—ROUTER1 provides routing and forwarding functionality between the Corpnet and 2-Corpnet subnets.
- STEP 3: Install and Configure CLIENT2—CLIENT2 is a Windows 7 client computer that is used to demonstrate the backwards compatibility of a Windows Server® 2012 Remote Access deployment.
- STEP 4: Configure APP1—Configure APP1 with ROUTER1 as the default gateway and 2-DC1 as the alternate DNS server.
- STEP 5: Configure DC1—Configure DC1 with an additional Active Directory site and additional security groups for Windows 7 client computers.
- STEP 6: Install and Configure 2-DC1—In a multisite deployment, you have two or more domains and sites. 2-DC1 provides domain controller and DNS services for the corp2.corp.contoso.com domain.
- STEP 7: Install and Configure 2-APP1—2-APP1 is a web and file server in the 2-Corpnet network.
- STEP 8: Configure INET1—INET1 simulates the Internet in this test lab guide. You must configure a DNS entry that resolves to the public IP address of 2-EDGE1.
- STEP 9: Configure EDGE1—Configure the 2-Corpnet DNS server and routing on EDGE1.
- STEP 10: Install and Configure 2-EDGE1—Two Remote Access servers are required in a multisite deployment. 2-EDGE1 provides Remote Access services for the second domain.
- STEP 11: Configure the Multisite Deployment—After configuring both Remote Access servers, you can configure your multisite deployment.
- STEP 12: Test DirectAccess Connectivity—Test DirectAccess connectivity from both client computers from the Internet subnet through EDGE1 and 2-EDGE1.
- STEP 13: Test DirectAccess Connectivity from Behind a NAT Device—Test DirectAccess connectivity from behind a NAT device.
- STEP 14: Snapshot the Configuration—After completing the test lab, take a snapshot of the working Remote Access multisite deployment so that you can return to it later to test additional scenarios.

STEP 1: Complete the DirectAccess Configuration

The first step is to complete all the steps in the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6. If you have already completed the steps in this test lab guides and saved a snapshot or disk image of the test lab, you can restore the snapshot or image and begin with the next step.

STEP 2: Install and Configure ROUTER1

In this multisite test lab guide, the router computer provides an IPv4 and IPv6 bridge between the Corpnet and 2-Corpnet subnets, and acts as a router for IP-HTTPS and Teredo traffic.

- Install the operating system on ROUTER1
- Configure TCP/IP properties and rename the computer
- Turn off the firewall
- Configure routing and forwarding

Install the operating system on ROUTER1

First, install Windows Server® 2012.

To install the operating system on ROUTER1

1. Start the installation of Windows Server 2012 (full installation).
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect ROUTER1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012, and then disconnect from the Internet.
4. Connect ROUTER1 to the Corpnet and 2-Corpnet subnets.

Configure TCP/IP properties and rename the computer

Configure TCP/IP settings on the router and rename the computer to ROUTER1.

To configure TCP/IP properties and rename the computer

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.

2. In the **Network Connections** window, right-click the network adapter that is connected to Corpnet, click **Rename**, type **Corpnet**, and press ENTER.
3. Right-click **Corpnet**, and then click **Properties**.
4. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. Click **Use the following IP address**. In **IP address**, type **10.0.0.254**. In **Subnet mask**, type **255.255.255.0**, and then click **OK**.
6. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
7. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:1::fe**. In **Subnet prefix length**, type **64**, and then click **OK**.
8. On the **Corpnet Properties** dialog box click **Close**.
9. In the **Network Connections** window, right-click the network adapter that is connected to 2-Corpnet, click **Rename**, type **2-Corpnet**, and press ENTER.
10. Right-click **2-Corpnet**, and then click **Properties**.
11. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
12. Click **Use the following IP address**. In **IP address**, type **10.2.0.254**. In **Subnet mask**, type **255.255.255.0**, and then click **OK**.
13. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
14. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:2::fe**. In **Subnet prefix length**, type **64**, and then click **OK**.
15. On the **2-Corpnet Properties** dialog box click **Close**.
16. Close the **Network Connections** window.
17. In the Server Manager console, in **Local Server**, in the **Properties** area, next to **Computer name**, click the link.
18. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
19. On the **Computer Name/Domain Changes** dialog box, in **Computer name**, type **ROUTER1**, and then click **OK**.
20. When you are prompted that you must restart the computer, click **OK**.
21. On the **System Properties** dialog box, click **Close**.
22. When you are prompted to restart the computer, click **Restart Now**.
23. After the computer has restarted, log on with the local Administrator account.

Turn off the firewall

This computer is configured only to provide routing between the Corpnet and 2-Corpnet subnets; therefore, the firewall must be turned off.

To turn off the firewall

1. On the **Start** screen, type **wf.msc**, and then press ENTER.
2. In Windows Firewall with Advanced Security, in the **Actions** pane, click **Properties**.
3. On the **Windows Firewall with Advanced Security** dialog box, on the **Domain Profile**

tab, in **Firewall state**, click **Off**.

4. On the **Windows Firewall with Advanced Security** dialog box, on the **Private Profile** tab, in **Firewall state**, click **Off**.
5. On the **Windows Firewall with Advanced Security** dialog box, on the **Public Profile** tab, in **Firewall state**, click **Off**, and then click **OK**.
6. Close Windows Firewall with Advanced Security.

Configure routing and forwarding

To provide routing and forwarding services between the Corpnet and 2-Corpnet subnets, you must enable forwarding on the network interfaces and configure static routes between the subnets.

To configure static routes

1. On the **Start** screen, type **cmd.exe**, and then press ENTER.
2. Enable forwarding on both the IPv4 and IPv6 interfaces of both network adapters using the following commands. After entering each command, press ENTER.

```
netsh interface IPv4 set interface Corpnet forwarding=enabled
netsh interface IPv4 set interface 2-Corpnet
forwarding=enabled
netsh interface IPv6 set interface Corpnet forwarding=enabled
netsh interface IPv6 set interface 2-Corpnet
forwarding=enabled
```

3. Enable IP-HTTPS routing between the Corpnet and 2-Corpnet subnets.

```
netsh interface IPv6 add route 2001:db8:1:1000::/59 Corpnet
2001:db8:1::2
netsh interface IPv6 add route 2001:db8:2:2000::/59 2-Corpnet
2001:db8:2::20
```

4. Enable Teredo routing between the Corpnet and 2-Corpnet subnets.

```
netsh interface IPv6 add route 2001:0:836b:2::/64 Corpnet
2001:db8:1::2
netsh interface IPv6 add route 2001:0:836b:14::/64 2-Corpnet
2001:db8:2::20
```

5. Close the Command Prompt window.

STEP 3: Install and Configure CLIENT2

CLIENT2 is a Windows® 7 computer that is used to demonstrate the backwards compatibility of Remote Access running on Windows Server® 2012 servers.

1. To install the operating system on CLIENT2—Install Windows® 7 Enterprise or Windows® 7 Ultimate on CLIENT2.
2. To join CLIENT2 to the CORP domain—Join CLIENT2 to the corp.contoso.com domain.

To install the operating system on CLIENT2

1. Start the installation of Windows 7.
2. When you are prompted for a user name, type **User1**. When you are prompted for a computer name, type **CLIENT2**.
3. When you are prompted for a password, type a strong password twice.
4. When you are prompted for protection settings, click **Use recommended settings**.
5. When you are prompted for your computer's current location, click **Work network**.
6. Connect CLIENT2 to a network that has Internet access and run Windows Update to install the latest updates for Windows 7, and then disconnect from the Internet.
7. Connect CLIENT2 to the Corpnet subnet.

To join CLIENT2 to the CORP domain

1. Click **Start**, right-click **Computer**, and then click **Properties**.
2. On the **System** page, in the **Computer name, domain, and workgroup settings** area, click **Change settings**.
3. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
4. On the **Computer Name/Domain Changes** dialog box, click **Domain**, type **corp.contoso.com**, and then click **OK**.
5. When you are prompted for a user name and password, type the user name and password for the User1 domain account, and then click **OK**.
6. When you see a dialog box that welcomes you to the corp.contoso.com domain, click **OK**.
7. When you see a dialog box that prompts you to restart the computer, click **OK**.
8. On the **System Properties** dialog box, click **Close**, and when you see a dialog that prompts you to restart the computer, click **Restart Now**.
9. After the computer restarts, log on as CORP\User1.

User account control

When you configure the Windows 7 operating system, you are required to click **Continue** on the **User Account Control** (UAC) dialog box for some tasks. Several of the configuration tasks

require UAC approval. When you are prompted, always click **Continue** to authorize these changes.

STEP 4: Configure APP1

Configure static IPv6 addressing and gateway settings to enable APP1 access to the 2-Corpnet subnet.

- To configure the default gateway and DNS server—The multisite configuration uses the ROUTER1 computer as a default gateway. Configure the default gateway on APP1.

To configure the default gateway and DNS server

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In the **Network Connections** window, right-click **Wired Ethernet Connection**, and then click **Properties**.
3. On the **Wired Ethernet Connection Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. In **Default gateway**, type **10.0.0.254**, and in **Alternate DNS server**, type **10.2.0.1**, and then click **OK**.
5. On the **Wired Ethernet Connection Properties** dialog box, click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
6. In **Default gateway**, type **2001:db8:1::fe**. In **Alternate DNS server**, type **2001:db8:2::1**, and then click **OK**.
7. On the **Wired Ethernet Connection Properties** dialog box, click **Close**, and then close the **Network Connections** window.

STEP 5: Configure DC1

DC1 acts as a domain controller, DNS server, and DHCP server for the corp.contoso.com domain.

To configure Remote Access to use a multisite topology, it is necessary to add an additional Active Directory Domain Services (AD DS) site for the second domain controller 2-DC1, and to configure routing between the subnets.

1. To configure the default gateway on the domain controller—Configure the default gateway on DC1.
2. Create security groups for Windows 7 DirectAccess clients on DC1—When DirectAccess is configured, it automatically creates Group Policy Objects (GPOs) and GPO settings that are

applied to DirectAccess clients and servers. The DirectAccess client GPO is applied to specific Active Directory security groups.

3. To add a new AD DS site—Create a second AD DS site.

To configure the default gateway on the domain controller

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In the Network Connections window, right-click **Wired Ethernet Connection**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. In **Default gateway**, type **10.0.0.254**, and in **Alternate DNS server**, type **10.2.0.1**, and then click **OK**.
5. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
6. In **Default gateway**, type **2001:db8:1::fe**, and in **Alternate DNS server**, type **2001:db8:2::1**, and then click **OK**.
7. On the **Wired Ethernet Connection Properties** dialog box, click **Close**.
8. Close the **Network Connections** window.

To add a new AD DS site

1. On the **Start** screen, type **dssite.msc**, and then press ENTER.
2. In the Active Directory Sites and Services console, in the console tree, right-click **Sites**, and then click **New Site**.
3. On the **New Object - Site** dialog box, in the **Name** box, type **Second-Site**.
4. In the list box, click **DEFAULTIPSITELINK**, and then click **OK** twice.
5. In the console tree, expand **Sites**, right-click **Subnets**, and then click **New Subnet**.
6. On the **New Object - Subnet** dialog box, under **Prefix**, type **10.0.0.0/24**, in the **Select a site object for this prefix** list, click **Default-First-Site-Name**, and then click **OK**.
7. In the console tree, right-click **Subnets**, and then click **New Subnet**.
8. On the **New Object – Subnet** dialog box, under **Prefix**, type **2001:db8:1::/64**, in the **Select a site object for this prefix** list, click **Default-First-Site-Name**, and then click **OK**.
9. In the console tree, right-click **Subnets**, and then click **New Subnet**.
10. On the **New Object - Subnet** dialog box, under **Prefix**, type **10.2.0.0/24**, in the **Select a site object for this prefix** list, click **Second-Site**, and then click **OK**.
11. In the console tree, right-click **Subnets**, and then click **New Subnet**.
12. On the **New Object – Subnet** dialog box, under **Prefix**, type **2001:db8:2::/64**, in the **Select a site object for this prefix** list, click **Second-Site**, and then click **OK**.
13. Close Active Directory Sites and Services.

Create security groups for Windows 7 DirectAccess clients on DC1

Create the DirectAccess security groups for Windows® 7 with the following procedure.

Windows 7 client computers must be members of separate security groups because they are able to connect to internal resources through a single entry point only. When enabling Multisite support or adding entry points, if Windows 7 support is requested, then a separate GPO will be automatically created by DirectAccess for Windows 7 clients for each entry point.

Create security groups

1. On the **Start** screen, type **dsa.msc**, and then press ENTER.
2. In the left pane, expand **corp.contoso.com**, click **Users**, then right-click **Users**, point to **New**, and then click **Group**.
3. On the **New Object - Group** dialog box, under **Group name**, enter **Win7_Clients_Site1**.
4. Under **Group scope**, click **Global**, under **Group type**, click **Security**, and then click **OK**.
5. Double-click the **Win7_Clients_Site1** security group, and on the **Win7_Clients_Site1 Properties** dialog box, click the **Members** tab.
6. On the **Members** tab, click **Add**.
7. On the **Select Users, Contacts, Computers, or Service Accounts** dialog box, click **Object Types**. On the **Object Types** dialog box, select **Computers**, and then click **OK**.
8. In **Enter the object names to select**, type **client2**, and then click **OK**, and then on the **Win7_Clients_Site1 Properties** dialog box click **OK**.
9. In the **Active Directory Users and Computers** console, in the left pane, right-click **Users**, point to **New**, and then click **Group**.
10. On the **New Object - Group** dialog box, under **Group name**, enter **Win7_Clients_Site2**.
11. Under **Group scope**, click **Global**, under **Group type**, click **Security**, and then click **OK**.
12. Close the **Active Directory Users and Computers** console.

STEP 6: Install and Configure 2-DC1

2-DC1 provides the following services:

- A domain controller for the corp2.corp.contoso.com Active Directory Domain Services (AD DS) domain.
- A DNS server for the corp2.corp.contoso.com DNS domain.

2-DC1 configuration consists of the following:

- Install the operating system on 2-DC1
- Configure TCP/IP properties
- Configure 2-DC1 as a domain controller and DNS server

- Provide Group Policy permissions to CORP\User1
- Allow CORP2 computers to obtain computer certificates
- Force replication between DC1 and 2-DC1

Install the operating system on 2-DC1

First, install Windows Server® 2012.

To install the operating system on 2-DC1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect 2-DC1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012, and then disconnect from the Internet.
4. Connect 2-DC1 to the 2-Corpnet subnet.

Configure TCP/IP properties

Configure the TCP/IP protocol with static IP addresses.

To configure TCP/IP on 2-DC1

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Click **Use the following IP address**. In **IP address**, type **10.2.0.1**. In **Subnet mask**, type **255.255.255.0**. In **Default gateway**, type **10.2.0.254**. Click **Use the following DNS server addresses**, in **Preferred DNS server**, type **10.2.0.1**, and in **Alternate DNS server**, type **10.0.0.1**.
5. Click **Advanced**, and then click the **DNS** tab.
6. In **DNS suffix for this connection**, type **corp2.corp.contoso.com**, and then click **OK** twice.
7. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
8. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:2::1**. In **Subnet prefix length**, type **64**. In **Default gateway**, type **2001:db8:2::fe**. Click **Use the following DNS server addresses**, in **Preferred DNS server**, type **2001:db8:2::1**, and in **Alternate DNS server**, type **2001:db8:1::1**.
9. Click **Advanced**, and then click the **DNS** tab.
10. In **DNS suffix for this connection**, type **corp2.corp.contoso.com**, and then click **OK**

twice.

11. On the **Wired Ethernet Connection Properties** dialog box, click **Close**.
12. Close the **Network Connections** window.
13. In the Server Manager console, in **Local Server**, in the **Properties** area, next to **Computer name**, click the link.
14. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
15. On the **Computer Name/Domain Changes** dialog box, in **Computer name**, type **2-DC1**, and then click **OK**.
16. When you are prompted that you must restart the computer, click **OK**.
17. On the **System Properties** dialog box, click **Close**.
18. When you are prompted to restart the computer, click **Restart Now**.
19. After restarting, login using the local administrator account.

Configure 2-DC1 as a domain controller and DNS server

Configure 2-DC1 as a domain controller for the corp2.corp.contoso.com domain and as a DNS server for the corp2.corp.contoso.com DNS domain.

To configure 2-DC1 as a domain controller and DNS server

1. In the Server Manager console, on the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen
3. On the **Select server roles** page, select **Active Directory Domain Services**. Click **Add Features** when prompted, and then click **Next** three times.
4. On the **Confirm installation selections** page, click **Install**.
5. When the installation completes successfully, click **Promote this server to a domain controller**.
6. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, click **Add a new domain to an existing forest**.
7. In **Parent domain name**, type **corp.contoso.com**, in **New domain name**, type **corp2**.
8. Under **Supply the credentials to perform this operation**, click **Change**. On the **Windows Security** dialog box, in **User name**, type **corp.contoso.com\Administrator**, and in **Password**, enter the corp\Administrator password, click **OK**, and then click **Next**.
9. On the **Domain Controller Options** page, make sure that the **Site name** is **Second-Site**. Under **Type the Directory Services Restore Mode (DSRM) password**, in **Password** and **Confirm password**, type a strong password twice, and then click **Next** five times.
10. On the **Prerequisites Check** page, after the prerequisites are validated, click **Install**.
11. Wait until the wizard completes the configuration of Active Directory and DNS services, and then click **Close**.

12. After the computer restarts, log in to the CORP2 domain using the Administrator account.

Provide Group Policy permissions to CORP\User1

Use this procedure to provide the CORP\User1 user with full permissions to create and change corp2 Group Policy Objects.

To provide Group Policy permissions

1. On the **Start** screen, type **gpmc.msc**, and then press ENTER.
2. In the Group Policy Management console, open **Forest: corp.contoso.com/Domains/corp2.corp.contoso.com**.
3. In the details pane, click the **Delegation** tab. In the **Permission** drop-down list, click **Link GPOs**.
4. Click **Add**, and on the new **Select User, Computer, or Group** dialog box, click **Locations**.
5. On the **Locations** dialog box, in the **Location** tree, click **corp.contoso.com**, and click **OK**.
6. In **Enter the object name to select** type **User1**, click **OK**, and on the **Add Group or User** dialog box, click **OK**.
7. In the Group Policy Management console, in the tree, click **Group Policy Objects**, and in the details pane click the **Delegation** tab.
8. Click **Add**, and on the new **Select User, Computer, or Group** dialog box, click **Locations**.
9. On the **Locations** dialog box, in the **Location** tree, click **corp.contoso.com**, and click **OK**.
10. In **Enter the object name to select** type **User1**, click **OK**.
11. In the Group Policy Management console, in the tree, click **WMI Filters**, and in the details pane click the **Delegation** tab.
12. Click **Add**, and on the new **Select User, Computer, or Group** dialog box, click **Locations**.
13. On the **Locations** dialog box, in the **Location** tree, click **corp.contoso.com**, and click **OK**.
14. In **Enter the object name to select** type **User1**, click **OK**. On the **Add Group or User** dialog box, make sure that **Permissions** are set to **Full control**, and then click **OK**.
15. Close the Group Policy Management console.

Allow CORP2 computers to obtain computer certificates

Computers in the CORP2 domain must obtain computer certificates from the certification authority on APP1. Perform this procedure on APP1.

To allow CORP2 computers to automatically obtain computer certificates

1. On APP1, click **Start**, type **certtmpl.msc**, and then press ENTER.
2. In the **Certificates Template Console**, in the middle pane, double-click **Client-Server Authentication**.
3. On the **Client-Server Authentication Properties** dialog box, click the **Security** tab.
4. Click **Add**, and on the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Locations**.
5. On the **Locations** dialog box, in **Location**, expand **corp.contoso.com**, click **corp2.corp.contoso.com**, and then click **OK**.
6. In **Enter the object names to select**, type **Domain Admins; Domain Computers** and then click **OK**.
7. On the **Client-Server Authentication Properties** dialog box, in **Group or user names**, click **Domain Admins (CORP2\Domain Admins)**, and in **Permissions for Domain Admins**, in the **Allow** column, select **Write** and **Enroll**.
8. In **Group or user names**, click **Domain Computers (CORP2\Domain Computers)**, and in **Permissions for Domain Computers**, in the **Allow** column, select **Enroll** and **Autoenroll**, and then click **OK**.
9. Close the Certificate Templates Console.

Force replication between DC1 and 2-DC1

Before you can enroll for certificates on 2-EDGE1, you must force the replication of settings from DC1 to 2-DC1. This operation should be done on DC1.

To force replication

1. On DC1, click **Start**, and then click **Active Directory Sites and Services**.
2. In the Active Directory Sites and Services console, in the tree, expand **Inter-Site Transports**, and then click **IP**.
3. In the details pane, double-click **DEFAULTIPSITELINK**.
4. On the **DEFAULTIPSITELINK Properties** dialog box, in **Cost**, type **1**, in **Replicate every**, type **15**, and then click **OK**. Wait for 15 minutes for replication to complete.
5. To force replication now in the console tree, expand **Sites\Default-First-Site-name\Servers\DC1\NTDS Settings**, in the details pane, right-click **<automatically generated>**, click **Replicate Now**, and then on the **Replicate Now** dialog box, click **OK**.
6. To ensure replication has completed successfully do the following:
   a. On the **Start** screen, type **cmd.exe**, and then press ENTER.
   b. Type the following command, and then press ENTER.

      ```
      repadmin /syncall /e /A /P /d /q
      ```

   c. Make sure that all partitions are synchronized with no errors. If not, then rerun the command until no errors are reported before proceeding.

7. Close the Command Prompt window.

STEP 7: Install and Configure 2-APP1

2-APP1 provides web and file sharing services. 2-APP1 configuration consists of the following:

- Install the operating system on 2-APP1
- Configure TCP/IP properties
- Join 2-APP1 to the CORP2 domain
- Install the Web Server (IIS) role on 2-APP1
- Create a shared folder on 2-APP1

Install the operating system on 2-APP1

First, install Windows Server® 2012.

To install the operating system on 2-APP1

1. Start the installation of Windows Server 2012 (full installation).
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect 2-APP1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012, and then disconnect from the Internet.
4. Connect 2-APP1 to the 2-Corpnet subnet.

Configure TCP/IP properties

Configure TCP/IP properties on 2-APP1.

To configure TCP/IP properties

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In the **Network Connections** window, right-click **Wired Ethernet Connection**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Click **Use the following IP address**. In **IP address**, type **10.2.0.3**. In **Subnet mask**, type **255.255.255.0**. In **Default gateway**, type **10.2.0.254**.
5. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.2.0.1**.
6. Click **Advanced**, and then click the **DNS** tab. In **DNS suffix for this connection**, type

corp2.corp.contoso.com, and click **OK** twice.

7. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
8. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:2::3**. In **Subnet prefix length**, type **64**. In **Default gateway**, type **2001:db8:2::fe**. Click **Use the following DNS server addresses**, and in **Preferred DNS server**, type **2001:db8:2::1**.
9. Click **Advanced**, and then click the **DNS** tab.
10. In **DNS suffix for this connection**, type **corp2.corp.contoso.com**, and then click **OK** twice.
11. On the **Wired Ethernet Connection Properties** dialog box click **Close**.
12. Close the **Network Connections** window.

Join 2-APP1 to the CORP2 domain

Join 2-APP1 to the corp2.corp.contoso.com domain.

To join 2-APP1 to the CORP2 domain

1. In the Server Manager console, in **Local Server**, in the **Properties** area, next to **Computer name**, click the link.
2. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
3. In **Computer name**, type **2-APP1**. In **Member of**, click **Domain**, type **corp2.corp.contoso.com**, and then click **OK**.
4. When you are prompted for a user name and password, type **Administrator** and its password, and then click **OK**.
5. When you see a dialog box welcoming you to the corp2.corp.contoso.com domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.
9. After the computer restarts, click **Switch User**, and then click **Other User** and log on to the CORP2 domain with the Administrator account.

Install the Web Server (IIS) role on 2-APP1

Install the Web Server (IIS) role to make 2-APP1 a web server.

To install the Web Server (IIS) role

1. In the Server Manager console, on the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen
3. On the **Select server roles** page, select **Web Server (IIS)**, and then click **Next** four times.

4. On the **Confirm installation selections** page, click **Install**.
5. Verify that the installation was successful, and then click **Close**.

Create a shared folder on 2-APP1

Create a shared folder and a text file within the folder on 2-APP1.

To create a shared folder

1. On the **Start** screen, type **explorer.exe**, and then press ENTER.
2. Click **Computer**, then double-click **Local Disk (C:)**.
3. Click **New Folder**, type **Files**, and then press ENTER. Leave the **Local Disk** window open.
4. On the **Start** screen, type **notepad.exe**, right-click **Notepad**, click **Advanced**, and then click **Run as administrator**.
5. In the **Untitled – Notepad** window, type **This is a shared file on 2-APP1**.
6. Click **File**, click **Save**, click **Computer**, double-click **Local Disk (C:)**, and then double-click the **Files** folder.
7. In **File name**, type **example.txt**, and then click **Save**. Close Notepad.
8. In the **Local Disk** window, right-click the **Files** folder, point to **Share with**, and then click **Specific people**.
9. On the **File Sharing** dialog box, in the drop-down list, click **Everyone**, and then click **Add**. In **Permission Level** for **Everyone**, click **Read/Write**.
10. Click **Share**, and then click **Done**.
11. Close the **Local Disk** window.

STEP 8: Configure INET1

To enable client computers to connect to Remote Access servers over the Internet, you must configure a DNS entry for 2-EDGE1 on INET1.

To create the 2-EDGE1 DNS entry

1. On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the console tree, open **Forward Lookup Zones**, click **contoso.com**, then right-click **contoso.com**, and then click **New Host (A or AAAA)**.
3. In **Name**, type **2-EDGE1**. In **IP address**, type **131.107.0.20**. Click **Add Host**, click **OK**, and then click **Done**.

STEP 9: Configure EDGE1

The following procedures are performed on the EDGE1 server:

1. Configure the DNS servers on EDGE1—It is necessary to configure the DNS server from the corp2.corp.contoso.com domain on EDGE1.
2. Configure routing between subnets—Configure routing on EDGE1 to enable communication between the Corpnet and 2-Corpnet subnets.

Configure the DNS servers on EDGE1

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Corpnet**, click the link.
2. In the Network Connections window, right-click **Corpnet**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. In **Alternate DNS server**, type **10.2.0.1**. and then click **OK**.
5. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
6. In **Alternate DNS server**, type **2001:db8:2::1** and then click **OK**.
7. On the **Corpnet Properties** dialog box, click **Close**.
8. Close the **Network Connections** window.

Configure routing between subnets

1. On the **Start** screen, type **cmd.exe**, right-click **cmd**, click **Advanced**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Command Prompt window, enter the following commands. After entering each command, press ENTER.

```
netsh interface IPv4 add route 10.2.0.0/24 Corpnet 10.0.0.254
netsh interface IPv6 add route 2001:db8:2::/64 Corpnet
2001:db8:1::fe
```

3. Close the Command Prompt window.

STEP 10: Install and Configure 2-EDGE1

2-EDGE1 configuration consists of the following:

- Install the operating system on 2-EDGE1—Install Windows Server® 2012 on 2-EDGE1.
- Configure TCP/IP properties—Configure 2-EDGE1 with static addresses on both network interfaces.

- Configure routing between subnets—To enable communication between the Corpnet and 2-Corpnet subnets, you must configure routing.
- Join 2-EDGE1 to the CORP2 domain—Join 2-EDGE1 to the corp2.corp.contoso.com domain.
- Obtain certificates on 2-EDGE1—Certificates are required for the IPsec connection between DirectAccess clients and the Remote Access server, and to authenticate the IP-HTTPS listener when clients connect over HTTPS.
- Provide access to CORP\User1—The user CORP\User1 is the Remote Access administrator. To enable this user to make changes to 2-EDGE1 from EDGE1, you must grant access to the user.
- Install the Remote Access role on 2-EDGE1—To enable a multisite deployment, you must install the Remote Access role on 2-EDGE1.

2-EDGE1 must have two network adapters installed.

Install the operating system on 2-EDGE1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect 2-EDGE1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012, and then disconnect from the Internet.
4. Connect one network adapter to the 2-Corpnet subnet and the other to the simulated Internet.

Configure TCP/IP properties

1. In the Server Manager console, click **Local Server**, and then in the **Properties** area, next to **Wired Ethernet Connection**, click the link.
2. In **Network Connections**, right-click the network connection that is connected to the 2-Corpnet subnet, click **Rename**, type **2-Corpnet**, and then press ENTER.
3. Right-click **2-Corpnet**, and then click **Properties**.
4. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. Click **Use the following IP address**. In **IP address**, type **10.2.0.20**, in **Subnet mask**, type **255.255.255.0**.
6. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.2.0.1**, and in **Alternate DNS server**, type **10.0.0.1**.
7. Click **Advanced**, and then click the **DNS** tab.
8. In **DNS suffix for this connection**, type **corp2.corp.contoso.com**, and then click **OK** twice.
9. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
10. Click **Use the following IPv6 address**. In **IPv6 address**, type **2001:db8:2::20**, in

Subnet prefix length, type **64**. Click **Use the following DNS server addresses**, and in **Preferred DNS server**, type **2001:db8:2::1**, in **Alternate DNS server**, type **2001:db8:1::1**.

11. Click **Advanced**, and then click the **DNS** tab.
12. In **DNS suffix for this connection**, type **corp2.corp.contoso.com**, and then click **OK** twice.
13. On the **2-Corpnet Properties** dialog box, click **Close**.
14. In the **Network Connections** window, right-click the network connection that is connected to the Internet subnet, click **Rename**, type **Internet**, and then press ENTER.
15. Right-click **Internet**, and then click **Properties**.
16. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
17. Click **Use the following IP address**. In **IP address**, type **131.107.0.20**. In **Subnet mask**, type **255.255.255.0**.
18. Click **Advanced**. On the **IP Settings** tab, in the **IP addresses** area, click **Add**. On the **TCP/IP Address** dialog box, in **IP address** type **131.107.0.21**, in **Subnet mask** type **255.255.255.0**, and then click **Add**.
19. Click the **DNS** tab.
20. In **DNS suffix for this connection**, type **isp.example.com**, click **OK** twice, and then click **Close**.
21. Close the **Network Connections** window.

Configure routing between subnets

1. On the **Start** screen, type **cmd.exe**, and then press ENTER.
2. In the Command Prompt window, enter the following commands. After entering each command, press ENTER.

```
netsh interface IPv4 add route 10.0.0.0/24 2-Corpnet
10.2.0.254
netsh interface IPv6 add route 2001:db8:1::/64 2-Corpnet
2001:db8:2::fe
```

3. To check network communication between 2-EDGE1 and DC1, type **ping dc1.corp.contoso.com**.
4. Verify that there are four responses from either the IPv4 address, 10.0.0.1, or from the IPv6 address, 2001:db8:1::1.
5. Close the Command Prompt window.

Join 2-EDGE1 to the CORP2 domain

1. In the Server Manager console, in **Local Server**, in the **Properties** area, next to **Computer name**, click the link.
2. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.

3. On the **Computer Name/Domain Changes** dialog box, in **Computer name**, type **2-EDGE1**. In **Member of**, click **Domain**, type **corp2.corp.contoso.com**, and then click **OK**.
4. If you are prompted for a user name and password, type **Administrator** and its password, and then click **OK**.
5. When you see a dialog box welcoming you to the corp2.corp.contoso.com domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.
9. After the computer has restarted, click **Switch User**, and then click **Other User** and log on to the CORP2 domain with the Administrator account.

Obtain certificates on 2-EDGE1

1. On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal**.
5. Right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, select the **Client-Server Authentication** and the **Web Server** check boxes, and then click **More information is required to enroll for this certificate**.
8. On the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common name**.
9. In **Value**, type **2-edge1.contoso.com**, and then click **Add**.
10. In the **Alternative name** area, in **Type**, select **DNS**.
11. In **Value**, enter **2-edge1.contoso.com**, and then click **Add**.
12. On the **General** tab, in **Friendly name**, type **IP-HTTPS Certificate**.
13. Click **OK**, click **Enroll**, and then click **Finish**.
14. In the details pane of the Certificates snap-in, verify that a new certificate with the name 2-edge1.contoso.com was enrolled with Intended Purposes of Server Authentication, and a new certificate with the name 2-edge1.corp2.corp.contoso.com was enrolled with Intended Purposes of Client Authentication and Server Authentication.
15. Close the console window. If you are prompted to save settings, click **No**.

Provide access to CORP\User1

1. On the **Start** screen, type **compmgmt.msc**, and then press ENTER.

2. In the left pane, click **Local Users and Groups**.
3. Double-click **Groups**, and then double-click **Administrators**.
4. On the **Administrators Properties** dialog box, click **Add**, and on the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Locations**.
5. On the **Locations** dialog box, in the **Location** tree, click **corp.contoso.com**, and then click **OK**.
6. In the **Enter the object names to select** type **User1**, and then click **OK**.
7. On the **Administrators Properties** dialog box, click **OK**.
8. Close the Computer Management window.

Install the Remote Access role on 2-EDGE1

1. In the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. On the **Select server roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

STEP 11: Configure the Multisite Deployment

To configure a multisite deployment, make changes to the current Remote Access configuration wizard on EDGE1, enable the multisite feature, and then add 2-EDGE1 as a second entry point.

- Configure Remote Access on EDGE1
- Enable multisite configuration on EDGE1
- Add 2-EDGE1 as a second entry-point

Configure Remote Access on EDGE1

1. On the **Start** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**.
3. In the middle pane of the console, in the **Step 2 Remote Access Server** area, click **Edit**.
4. Click **Prefix Configuration**. On the **Prefix Configuration** page, in **Internal network IPv6 prefixes**, enter **2001:db8:1::/64; 2001:db8:2::/64**. In **IPv6 prefix assigned to DirectAccess client computers**, enter **2001:db8:1:1000::/64**, click **Next**, and then click **Finish**.

5. In the middle pane of the console, in the **Step 3 Infrastructure Servers** area, click **Edit**.
6. Click **DNS Suffix Search List**. On the **DNS Suffix Search List** page, make sure that the **Configure DirectAccess clients with DNS client suffix search list** check box is selected and that the **corp.contoso.com** and **corp2.corp.contoso.com** domain suffixes appear in the **Domain suffixes to use** list, click **Next**, and then click Finish.
7. In the middle pane of the console, click **Finish**.
8. On the **Remote Access Review** dialog box, review the configuration settings, and then click **Apply**. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.
9. In the **Tasks** pane, click **Refresh Management Servers**, and click **Close** when finished.

Enable multisite configuration on EDGE1

1. In the Remote Access Management Console, in the **Tasks** pane, click **Enable Multisite**.
2. In the Enable Multisite Deployment wizard, on the **Before You Begin** page, click **Next**.
3. On the **Deployment Name** page, in **Multisite deployment name**, type **Contoso**, in **First entry point name**, type **Edge1-Site**, and then click **Next**.
4. On the **Entry Point Selection** page, click **Assign entry points automatically, and allow clients to select manually**, and then click **Next**.
5. On the **Global Load Balancing** page, click **No, do not use global load balancing**, and then click **Next**.
6. On the **Client Support** page, click **Allow client computers running Windows 7 to access this entry point**, and click **Add**.
7. On the **Select Groups** dialog box, in **Enter the object names to select**, type **Win7_Clients_Site1**, click **OK**, and then click **Next**.
8. On the **Client GPO Settings** page, click **Next**.
9. On the **Summary** page, click **Commit**.
10. On the **Enabling Multisite Deployment** dialog box, click **Close** and then on the Enable Multisite Deployment wizard, click **Close**.

Add 2-EDGE1 as a second entry-point

1. In the Remote Access Management Console, in the **Tasks** pane, click **Add an Entry Point**.
2. In the Add an Entry Point Wizard, on the **Entry Point Details** page, in **Remote Access server**, type **2-edge1.corp2.corp.contoso.com**, in **Entry point name**, type **2-Edge1-Site**, and then click **Next**.
3. On the **Network Topology** page, click **Edge**, and then click **Next**.
4. On the **Network Name or IP Address** page, in **Type in the public name or IP address used by clients to connect to the Remote Access server**, type **2-edge1.contoso.com**, and then click **Next**.
5. On the **Network Adapters** page, make sure that the **External adapter** is **Internet**, the **Internal adapter** is **2-Corpnet**, the certificate is **CN=2-edge1.contoso.com**, and then

click **Next**.

6. On the **Prefix Configuration** page, in **IPv6 prefix assigned to DirectAccess client computers**, type **2001:db8:2:2000::/64**, and then click **Next**.
7. On the **Client Support** page, click **Allow client computers running Windows 7 to access this entry point**, and click **Add**.
8. On the **Select Groups** dialog box, in **Enter the object names to select**, type **Win7_Clients_Site2**, click **OK**, and then click **Next**.
9. On the **Client GPO Settings** page, click **Next**.
10. On the **Server GPO Settings** page, click **Next**.
11. On the **Summary** page click **Commit**.
12. On the **Adding Entry Point** dialog box, click **Close** and then on the Add an Entry Point wizard, click **Close**.

STEP 12: Test DirectAccess Connectivity

Before you can test connectivity from the client computers when they are located on the Internet or Homenet networks, you must make sure they have the correct group policy settings.

- Verify clients have the correct group policy
- Test DirectAccess connectivity from the Internet through EDGE1
- Move CLIENT2 to the Win7_Clients_Site2 security group
- Test DirectAccess connectivity from the Internet through 2-EDGE1

Prerequisites

Connect both client computers to the Corpnet network, and then restart both client computers.

Verify clients have the correct group policy

1. On CLIENT1, click **Start**, type **powershell.exe**, right-click **powershell**, click **Advanced**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Windows PowerShell window, type **ipconfig** and press ENTER.

   Make sure that the Corpnet adapter IPv4 address starts with 10.0.0.
3. In the Windows PowerShell window type **Get-DnsClientNrptPolicy** and press ENTER. The Name Resolution Policy Table (NRPT) entries for DirectAccess are displayed.
   - .corp.contoso.com—These settings indicate that all connections to corp.contoso.com should be resolved by one of the DirectAccess DNS servers, with the IPv6 address 2001:db8:1::2 or 2001:db8:2::20.
   - nls.corp.contoso.com—These settings indicate that there is an exemption for the

name nls.corp.contoso.com.

4. Leave the Windows PowerShell window open for the next procedure.
5. On CLIENT2, click **Start**, click **All Programs**, click **Accessories**, click **Windows PowerShell**, right-click **Windows PowerShell**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
6. In the Windows PowerShell window, type **ipconfig** and press ENTER.

   Make sure that the Corpnet adapter IPv4 address starts with 10.0.0.
7. In the Windows PowerShell window, type **netsh namespace show policy** and press ENTER.

   In the output, there should be two sections:

   - .corp.contoso.com—These settings indicate that all connections to corp.contoso.com should be resolved by the DirectAccess DNS server, with the IPv6 address 2001:db8:1::2.
   - nls.corp.contoso.com—These settings indicate that there is an exemption for the name nls.corp.contoso.com.
8. Leave the Windows PowerShell window open for the next procedure.

Test DirectAccess connectivity from the Internet through EDGE1

1. Unplug 2-EDGE1 from the Internet network.
2. Unplug CLIENT1 and CLIENT2 from the corpnet switch and connect them to the Internet switch. Wait for 30 seconds.
3. On CLIENT1, in the Windows PowerShell window, type **ipconfig /all** and press ENTER.
4. Examine the output from the ipconfig command.

   The client computer is now connected to the Internet and has a public IPv4 address. When the DirectAccess client has a public IPv4 address, it uses the Teredo or IP-HTTPS IPv6 transition technologies to tunnel the IPv6 messages over an IPv4 Internet between the DirectAccess client and Remote Access server. Note that Teredo is the preferred transition technology.
5. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This flushes name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the corpnet.
6. Disable the Teredo interface to ensure that the client computer uses IP-HTTPS to connect to corpnet with the following command:

   ```
   netsh interface teredo set state disable
   ```
7. Ensure you are connected through EDGE1. Type **netsh interface httpstunnel show interfaces** and press ENTER.

   The output should contain `URL : https://edge1.contoso.com:443/IPHTTPS`.

   Tip

   On CLIENT1, you can also run the following Windows PowerShell command:

Get-NetIPHTTPSConfiguration. The output shows the available server URL connections and the currently active profile.

8. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address assigned to APP1, which in this case is 2001:db8:1::3.
9. In the Windows PowerShell window, type **ping 2-app1** and press ENTER. You should see replies from the IPv6 address assigned to 2-APP1, which in this case is 2001:db8:2::3.
10. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fd**c9:9f4e:eb1b**:7777::a00:4. Note that the bold values will vary due to how the address is generated.

    The ability to ping APP2 is important, because success indicates that you were able to establish a connection using NAT64/DNS64, as APP2 is an IPv4 only resource.
11. Open Internet Explorer, in the Internet Explorer address bar, enter **http://app1/** and press ENTER. You will see the default IIS website on APP1.
12. In the Internet Explorer address bar, enter **http://2-app1/** and press ENTER. You will see the default website on 2-APP1.
13. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP2.
14. On the **Start** screen, type **\\2-App1\Files**, and then press ENTER. Double-click the example text file.

    This demonstrates that you were able to connect to the file server in the corp2.corp.contoso.com domain when connected through EDGE1.
15. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file.

    This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource in the resource domain.
16. On the **Start** screen, type **wf.msc**, and then press ENTER.
17. In the **Windows Firewall with Advanced Security** console, notice that only the **Public Profile** is active. The Windows Firewall must be enabled for DirectAccess to work correctly. If the Windows Firewall is disabled, DirectAccess connectivity does not work.
18. In the left pane of the console, expand the **Monitoring** node, and click the **Connection Security Rules** node. You should see the active connection security rules: **DirectAccess Policy-ClientToCorp**, **DirectAccess Policy-ClientToDNS64NAT64PrefixExemption**, **DirectAccess Policy-ClientToInfra**, and **DirectAccess Policy-ClientToNlaExempt**. Scroll the middle pane to the right to show the **1st Authentication Methods** and **2nd Authentication Methods** columns. Notice that the first rule (ClientToCorp) uses Kerberos V5 to establish the intranet tunnel and the third rule (ClientToInfra) uses NTLMv2 to establish the infrastructure tunnel.
19. In the left pane of the console, expand the **Security Associations** node, and click the **Main Mode** node. Notice the infrastructure tunnel security associations using NTLMv2 and the intranet tunnel security association using Kerberos V5. Right-click the entry that shows **User (Kerberos V5)** as the **2nd Authentication Method** and click **Properties**.

On the **General** tab, notice the **Second authentication Local ID** is **CORP\User1**, indicating that User1 was able to successfully authenticate to the CORP domain using Kerberos.

20. Repeat this procedure from step 3 on CLIENT2.

Move CLIENT2 to the Win7_Clients_Site2 security group

1. On DC1, click **Start**, type **dsa.msc**, and then press ENTER.
2. In the Active Directory Users and Computers console, open **corp.contoso.com/Users** and double-click **Win7_Clients_Site1**.
3. On the **Win7_Clients_Site1 Properties** dialog box, click the **Members** tab, click **CLIENT2**, click **Remove**, click **Yes**, and then click **OK**.
4. Double-click **Win7_Clients_Site2**, and then on the **Win7_Clients_Site2 Properties** dialog box, click the **Members** tab.
5. Click **Add**, and on the **Select Users, Contacts, Computers, or Service Accounts** dialog box, click **Object Types**, select **Computers**, and then click **OK**.
6. In **Enter the object names to select**, type **CLIENT2**, and then click **OK**.
7. Restart CLIENT2 and log on using the corp/User1 account.
8. On CLIENT2, open an elevated Windows PowerShell window, type **netsh namespace show policy** and press ENTER.

   In the output, there should be two sections:

   - .corp.contoso.com—These settings indicate that all connections to corp.contoso.com should be resolved by the DirectAccess DNS server, with the IPv6 address 2001:db8:2::20.
   - nls.corp.contoso.com—These settings indicate that there is an exemption for the name nls.corp.contoso.com.

Test DirectAccess connectivity from the Internet through 2-EDGE1

1. Connect 2-EDGE1 to the Internet network.
2. Unplug EDGE1 from the Internet network.
3. On CLIENT1, open an elevated Windows PowerShell window.
4. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This flushes name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the corpnet.
5. Ensure you are connected through 2-EDGE1. Type **netsh interface httpstunnel show interfaces** and press ENTER.

   The output should contain `URL : https://2-edge1.contoso.com:443/IPHTTPS`.

   Tip
   On CLIENT1, you can also run the following command: **Get-NetIPHTTPSConfiguration**. The output shows the available server URL connections and the currently active profile.

Note

CLIENT1 automatically changes the server through which it connects to corporate resources. If the output of the command shows a connection to EDGE1, wait for approximately five minutes and then try again.

6. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address assigned to APP1, which in this case is 2001:db8:1::3.
7. In the Windows PowerShell window, type **ping 2-app1** and press ENTER. You should see replies from the IPv6 address assigned to 2-APP1, which in this case is 2001:db8:2::3.
8. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fd**c9:9f4e:eb1b**:7777::a00:4. Note that the bold values will vary due to how the address is generated.

   The ability to ping APP2 is important, because success indicates that you were able to establish a connection using NAT64/DNS64, as APP2 is an IPv4 only resource.
9. Open Internet Explorer, in the Internet Explorer address bar, enter **http://app1/** and press ENTER. You will see the default IIS website on APP1.
10. In the Internet Explorer address bar, enter **http://2-app1/** and press ENTER. You will see the default website on APP2.
11. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP3.
12. On the **Start** screen, type **\\App1\Files**, and then press ENTER. Double-click the example text file.

    This demonstrates that you were able to connect to the file server in the corp.contoso.com domain when connected through 2-EDGE1.
13. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file.

    This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource in the resource domain.
14. Repeat this procedure on CLIENT2 from step 3.

STEP 13: Test DirectAccess Connectivity from Behind a NAT Device

When a DirectAccess client is connected to the Internet from behind a NAT device or a web proxy server, the DirectAccess client uses either Teredo or IP-HTTPS to connect to the Remote Access server. If the NAT device enables outbound UDP port 3544 to the Remote Access server's public IP address, then Teredo is used. If Teredo access is not available, the DirectAccess client falls

back to IP-HTTPS over outbound TCP port 443, which enables access through firewalls or web proxy servers over the traditional SSL port. If the web proxy requires authentication, the IP-HTTPS connection will fail. IP-HTTPS connections also fail if the web proxy performs outbound SSL inspection, due to the fact that the HTTPS session is terminated at the web proxy instead of the Remote Access server.

The following procedures are performed on both client computers:

1. Test Teredo connectivity—The first set of tests are performed when the DirectAccess client is configured to use Teredo. This is the automatic setting when the NAT device allows outbound access to UDP port 3544. First run the tests on CLIENT1 and then run the tests on CLIENT2.
2. Test IP-HTTPS connectivity—The second set of tests are performed when the DirectAccess client is configured to use IP-HTTPS. In order to demonstrate IP-HTTPS connectivity, Teredo is disabled on the client computers. First run the tests on CLIENT1 and then run the tests on CLIENT2.

Prerequisites

Start EDGE1 and 2-EDGE1 if they are not already running, and make sure they are connected to the Internet subnet.

Before performing these tests, unplug CLIENT1 and CLIENT2 from the Internet switch and connect them to the Homenet switch. If asked what type of network you want to define the current network, select **Home network**.

Test Teredo connectivity

1. On CLIENT1, open an elevated Windows PowerShell window.
2. Enable the Teredo adapter, type **netsh interface teredo set state enterpriseclient**, and then press ENTER.
3. In the Windows PowerShell window, type **ipconfig /all** and press ENTER.
4. Examine the output of the ipconfig command.

   This computer is now connected to the Internet from behind a NAT device and is assigned a private IPv4 address. When the DirectAccess client is behind a NAT device and assigned a private IPv4 address, the preferred IPv6 transition technology is Teredo. If you look at the output of the ipconfig command, you should see a section for `Tunnel adapter Teredo Tunneling Pseudo-Interface` and then a description `Microsoft Teredo Tunneling Adapter`, with an IP address that starts with 2001:0 consistent with being a Teredo address. You should see the default gateway listed for the Teredo tunnel adapter as '::'.
5. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER.

   This will flush name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the Internet.
6. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address of APP1, 2001:db8:1::3.

7. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fd**c9:9f4e:eb1b**:7777::a00:4. Note that the bold values will vary due to how the address is generated.
8. In the Windows PowerShell window, type **ping 2-app1** and press ENTER. You should see replies from the IPv6 address of 2-APP1, 2001:db8:2::3.
9. Open Internet Explorer, in the Internet Explorer address bar, enter **http://2-app1/** and press ENTER. You will see the default IIS website on 2-APP1.
10. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see the default website on APP2.
11. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file. This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource on an IPv4 only host.
12. Repeat this procedure on CLIENT2.

Test IP-HTTPS connectivity

1. On CLIENT1, open an elevated Windows PowerShell window, and type **netsh interface teredo set state disabled** and press ENTER. This disables Teredo on the client computer and enables the client computer to configure itself to use IP-HTTPS. An **Ok** response appears when the command completes.
2. In the Windows PowerShell window, type **ipconfig /all** and press ENTER.
3. Examine the output of the ipconfig command. This computer is now connected to the Internet from behind a NAT device and is assigned a private IPv4 address. Teredo is disabled and the DirectAccess client falls back to IP-HTTPS. When you look at the output of the ipconfig command, you see a section for `Tunnel adapter iphttpsinterface` with an IP address that starts with 2001:db8:1:1000 or 2001:db8:2:2000 consistent with this being an IP-HTTPS address based on the prefixes that were configured when setting up DirectAccess. You will not see a default gateway listed for the IPHTTPSInterface tunnel adapter.
4. In the Windows PowerShell window, type **ipconfig /flushdns** and press ENTER. This will flush name resolution entries that may still exist in the client DNS cache from when the client computer was connected to the corpnet.
5. In the Windows PowerShell window, type **ping app1** and press ENTER. You should see replies from the IPv6 address of APP1, 2001:db8:1::3.
6. In the Windows PowerShell window, type **ping app2** and press ENTER. You should see replies from the NAT64 address assigned by EDGE1 to APP2, which in this case is fd**c9:9f4e:eb1b**:7777::a00:4. Note that the bold values will vary due to how the address is generated.
7. In the Windows PowerShell window, type **ping 2-app1** and press ENTER. You should see replies from the IPv6 address of 2-APP1, 2001:db8:2::3.
8. Open Internet Explorer, in the Internet Explorer address bar, enter **http://2-app1/** and press ENTER. You will see the default IIS website on 2-APP1.
9. In the Internet Explorer address bar, enter **http://app2/** and press ENTER. You will see

the default website on APP2.

10. On the **Start** screen, type **\\App2\Files**, and then press ENTER. Double-click the New Text Document file. This demonstrates that you were able to connect to an IPv4 only server using SMB to obtain a resource on an IPv4 only host.
11. Repeat this procedure on CLIENT2.

STEP 14: Snapshot the Configuration

This completes the DirectAccess multisite test lab. To save this configuration so that you can quickly return to a working DirectAccess multisite configuration from which you can test other DirectAccess modular test lab guides, test lab guide extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots TLG DirectAccess Multisite. If your lab uses physical computers, create disk images to save the DirectAccess test lab configuration.

Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID

Remote Access is a server role in the Windows Server® 2012 operating system that enables remote users to securely access internal network resources using DirectAccess or RRAS VPN. This guide contains step-by-step instructions for extending the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 to demonstrate a Remote Access one-time password (OTP) configuration.

About this guide

Remote Access in Windows Server 2012 adds support for client authentication with OTP. For the purposes of this test lab only RSA SecurID is used to demonstrate OTP functionality with Remote Access. Other RADIUS based OTP solutions are supported as well, but are outside the scope of this test lab. This guide contains instructions for configuring and demonstrating Remote Access using six servers and two client computers. The completed Remote Access with OTP test lab simulates an intranet, the Internet, and a home network, and demonstrates Remote Access functionality in different Internet connection scenarios.

Important

This lab is a proof of concept using the minimum number of computers. The configuration detailed in this guide is for test lab purposes only, and is not to be used in a production environment.

Overview of the Test Lab Scenario

In this test lab scenario, DirectAccess is deployed with:

- **DC1**—A server computer that is configured as a domain controller, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server.
- **EDGE1**—A server on the internal network that is configured as the Remote Access server. This server has two network adapters; one connected to the internal network, and the other connected to the external network.
- **RSA**—A server on the internal network that is configured as the RSA Authentication Manager server.
- **APP1**—A server on the internal network that is configured as a web and file server, and an enterprise root certification authority (CA).
- **APP2**—A server on the internal network that is configured as an IPv4 only web and file server. This server is used to highlight the NAT64/DNS64 capabilities. This guide uses Windows Server 2003 SP2; however, you can use any computer that is configured as an IPv4 only computer.
- **INET1**—A server that is configured as an Internet DNS and DHCP server.
- **NAT1**—A client computer that is configured as a network address translator (NAT) device using Internet Connection Sharing.
- **CLIENT1**—A client computer that is configured as a DirectAccess client, and will be used to test DirectAccess connectivity when moving between the internal network, the simulated Internet, and a home network.

The test lab consists of three subnets that simulate the following:

- A home network named Homenet (192.168.137.0/24) connected to the Internet by a NAT.
- The external network represented by the Internet subnet (131.107.0.0/24).
- An internal network named Corpnet (10.0.0.0/24; 2001:db8:1::/64) separated from the Internet by the Remote Access server.

Computers on each subnet connect using either a physical or virtual hub or switch, as shown in the following figure.

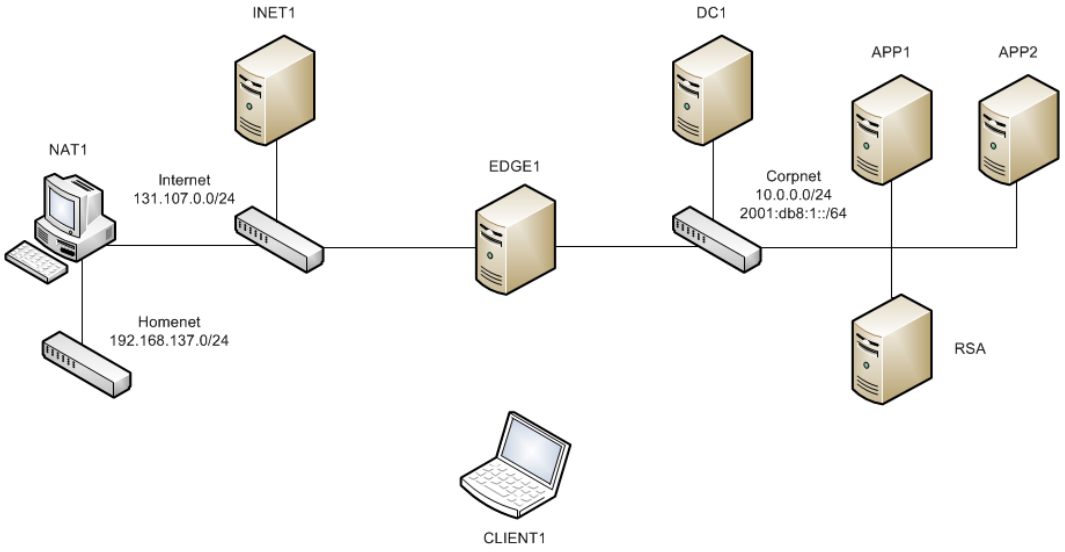

Configuration Requirements

The following components are required for configuring Remote Access in the test lab:

- The product disc or files for Windows Server® 2012.
- Four computers or virtual machines that meet the minimum hardware requirements for Windows Server 2012; one of these computers has two network adapters installed.
- The product disc or files for of Windows® 8.
- Two computers or virtual machines that meet the minimum hardware requirements for Windows 8; one of these computers has two network adapters installed.
- The product disc or files for Windows Server 2008 R2.
- One computer or virtual machine that meets the minimum hardware requirements for Windows Server 2008 R2; for the purposes of RSA this computer additionally meets the minimum memory requirements of 3GB RAM and 25GB hard disk with at least 10GB free disk space on drive C.
- The product disc or files for Windows Server 2003 SP2.
- One computer or virtual machine that meets the minimum hardware requirements for Windows Server 2003 SP2.
- Installation media or files for RSA Authentication Manager 7.1 SP4.
- License file for RSA Authentication Manager 7.1.
- Xml file with SecurID tokens initial seeds.
- Installation media or files for RSA SecurID software token.

Steps for Configuring the Test Lab

The following steps describe how to configure the Remote Access infrastructure, configure the Remote Access server and client, and test DirectAccess connectivity from the Homenet and Internet subnets.

In this test lab guide you will build a Remote Access with OTP environment by performing the following steps:

- STEP 1: Complete the DirectAccess Configuration—Complete all the steps in the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6.
- STEP 2: Configure APP1—Configure APP1 with OTP certificate templates for use by EDGE1.
- STEP 3: Configure DC1—Verify User Principal Name defined on DC1.
- STEP 4: Install and Configure RSA and EDGE1—Install and Configure RSA, a RADIUS and OTP server, and configure EDGE1 for OTP.
- STEP 5: Verify OTP Health on EDGE1—Make sure that the status of OTP is healthy on the Remote Access server.
- STEP 6: Test DirectAccess Connectivity from the Homenet Subnet—Test DirectAccess OTP functionality from behind a NAT device.
- STEP 7: Test DirectAccess Connectivity from the Internet—Test DirectAccess client connectivity from the Internet.
- STEP 8: Snapshot the Configuration—After completing the test lab, take a snapshot of the working DirectAccess with OTP configuration so that you can return to it later to test additional scenarios.

STEP 1: Complete the DirectAccess Configuration

The first step is to complete the Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6. If you have already completed the steps in that test lab guide and saved a snapshot or disk image of the test lab, you can restore the snapshot or image and begin with the next step.

STEP 2: Configure APP1

Use the following steps to prepare APP1 for OTP support:

1. To create and deploy a certificate template used to sign OTP certificate requests—Configure a certificate template used to sign OTP certificate requests.
2. To create and deploy a certificate template for OTP certificates issued by the corporate CA—Configure a certificate template for OTP certificates issued by the corporate CA.

To create and deploy a certificate template used to sign OTP certificate requests

1. On the **Start** screen, type **certtmpl.msc**, and then press ENTER.
2. In the Certificate Templates Console, in the details pane, right-click the **Computer** template, and click **Duplicate Template**.
3. On the **Properties of New Template** dialog box, on the **Compatibility** tab, in the **Certification Authority** list, click **Windows Server 2012**, and in the **Resulting changes** dialog box click **OK**. In the **Certificate recipient** list click **Windows 8/ Windows Server 2012**, and in the **Resulting changes** dialog box click **OK**.
4. On the **Properties of New Template** dialog box, click the **General** tab.
5. On the **General** tab, in **Template display name**, type **DAOTPRA**. Set the **Validity period** to 2 days, and set the **Renewal Period** to 1 day. If the **Certificate Templates** warning is displayed, click **OK**.
6. Click the **Security** tab, and then click **Add**.
7. On the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**. On the **Object Types** dialog box, select **Computers**, and then click **OK**. In the **Enter the object names to select** box, type **EDGE1**, click **OK**, and in the **Allow** column, select the **Read**, **Enroll**, and **Autoenroll** check boxes. Click **Authenticated Users**, select the **Read** check box under the **Allow** column, and clear all other check boxes. Click **Domain Computers**, and uncheck **Enroll** under the **Allow** column. Click **Domain Admins** and **Enterprise Admins** and click **Full Control** under the **Allow** column for both. Click **Apply**.
8. Click the **Subject Name** tab, and then click **Build from this Active Directory information**. In the **Subject name format:** list select **DNS name**, make sure that the **DNS name** box is checked, and click **Apply**.
9. Click the **Extensions** tab, select **Application Policies** and then click **Edit**. Remove all existing application policies. Click **Add**, and on the **Add Application Policy** dialog box, click **New**, enter **DA OTP RA** in the **Name:** field and **1.3.6.1.4.1.311.81.1.1** in the **Object identifier:** field, and click **OK**. On the **Add Application Policy** dialog box, click **OK**. On the **Edit Application Policies Extension**, click **OK**. On the **Properties of New Template** dialog box, click **OK**.

To create and deploy a certificate template for OTP certificates issued by the corporate CA

1. In the Certificate Templates Console, in the details pane, right-click the **Smartcard Logon** template, and click **Duplicate Template**.
2. On the **Properties of New Template** dialog box, on the **Compatibility** tab in the **Certification Authority** list, click **Windows Server 2012**, and in the **Resulting changes** dialog box, click **OK**. In the **Certificate recipient** list click **Windows 8/ Windows Server 2012**, and in the **Resulting changes** dialog box click **OK**.
3. On the **Properties of New Template** dialog box, click the **General** tab.
4. On the **General** tab, in **Template display name**, type **DAOTPLogon**. In **Validity period**, in the drop-down list, click **hours**, on the **Certificate Templates** dialog box, click **OK**, and

make sure that the number of hours is set to 1. In **Renewal period**, type **0**.

Note

In situations where the CA server is a Windows Server 2003 computer, then the template must be configured on a different computer. This is due to the fact that setting the **Validity period** in hours is not possible when running Windows versions prior to 2008/Vista. If the computer that you use to configure the template does not have the Certification Service role installed, or it is a client computer, then you may need to install the Certificate Templates snap-in. For more information on this subject click here.

5. Click the **Security** tab, select **Authenticated Users**, in the **Allow** column, and select the **Read** and **Enroll** check boxes. Click **OK**. Click **Full Control** under the **Allow** column. Click **Domain Admins** and **Enterprise Admins**, and click **Full Control** under the **Allow** column for both. Click **Apply**.
6. Click the **Subject Name** tab, and then click **Build from this Active Directory information**. In the **Subject name format:** list select **Fully distinguished name**, make sure that the **User principal name (UPN)** box is checked, and click **Apply**.
7. Click the **Server** tab, select the **Do not store certificates and requests in the CA database** check box, clear the **Do not include revocation information in issued certificates** check box, and then on the **Properties of New Template** dialog box, click **Apply**.
8. Click the **Issuance Requirements** tab, select the **This number of authorized signatures:** check box, set the value to 1. In the **Policy type required in signature:** list select **Application policy**, and in the **Application policy** list select **DA OTP RA**. On the **Properties of New Template** dialog box, click **OK**.
9. Click the **Extensions** tab, and on **Application Policies** click **Edit**. Delete **Client Authentication**, keep **SmartCardLogon**, and click **OK** twice.
10. Close the Certificate Templates Console.
11. On the **Start** screen, type **certsrv.msc**, and then press ENTER.
12. In the Certification Authority console tree, expand **corp-APP1-CA-1**, click **Certificate Templates**, right-click **Certificate Templates**, point to **New**, and click **Certificate Template to Issue**.
13. In the list of certificate templates, click **DAOTPRA** and **DAOTPLogon**, and click **OK**.
14. In the details pane of the console you should see the **DAOTPRA** certificate template with an **Intended purpose** of **DA OTP RA** and the **DAOTPLogon** certificate template with an **Intended Purpose** of **Smart Card Logon, Client Authentication**.
15. Restart the services.
16. Close the Certification Authority console.
17. Open an elevated command prompt. Type **CertUtil.exe –SetReg DBFlags +DBFLAGS_ENABLEVOLATILEREQUESTS**, and press ENTER.
18. Leave the Command Prompt window open for the next step.

STEP 3: Configure DC1

DC1 acts as a domain controller, DNS server, and DHCP server for the corp.contoso.com domain. Configure DC1 as follows:

Verify User1 has a User Principal Name defined on DC1

1. On DC1, open Server Manager, and click **AD DS** in the left pane. Right-click **DC1** and select **Active Directory Users and Computers**. In the left pane expand **corp.contoso.com\Users**, and double-click User1.
2. On the **Account** tab verify that **User logon name** is set to User1. If not, then enter **User1** in the **User logon name** field.
3. Click **OK**. Close the **Active Directory Users and Computers** console.

STEP 4: Install and Configure RSA and EDGE1

RSA is the RADIUS and OTP server, and is installed prior to configuring RADIUS and OTP.

You will perform the following steps to configure the RSA deployment:

1. Install the operating system on the RSA server—Install Windows Server 2008 R2 on the RSA server.
2. Configure TCP/IP on RSA—Configure TCP/IP settings on the RSA server.
3. Copy Authentication Manager installation files to the RSA server—After installing the operating system on RSA, copy the Authentication Manager files to the RSA computer.
4. Join the RSA server to the CORP domain—Join RSA to the CORP domain.
5. Disable Windows Firewall on RSA—Disable the Windows Firewall on the RSA server.
6. Install RSA Authentication Manager on the RSA server—Install RSA Authentication Manager.
7. Configure RSA Authentication Manager—Configure Authentication Manager.
8. Create DAProbeUser—Create a user account for probing purposes.
9. Install RSA SecurID software token on CLIENT1—Install RSA SecurID software token on CLIENT1.
10. Configure EDGE1 as an RSA Authentication Agent—Configure RSA Authentication Agent on EDGE1.
11. Configure EDGE1 to support OTP authentication—Configure OTP for DirectAccess, and verify the configuration.

Install the operating system on the RSA server

1. On RSA, start the installation of Windows Server 2008 R2.

2. Follow the instructions to complete the installation, specifying Windows Server 2008 R2 Standard (Full Installation) and a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RSA to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2008 R2, and then disconnect from the Internet.
4. Connect RSA to the Corpnet subnet.

Configure TCP/IP on RSA

1. In Initial Configuration Tasks, click **Configure networking**.
2. In **Network Connections**, right-click **Local Area Connection**, and then click **Properties**.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Click **Use the following IP address**. In **IP address**, type **10.0.0.5**. In **Subnet mask**, type **255.255.255.0**. In **Default Gateway**, type **10.0.0.2**. Click **Use the following DNS server addresses**, in **Preferred DNS server**, type **10.0.0.1**.
5. Click **Advanced**, and then click the **DNS** tab.
6. In **DNS suffix for this connection**, type **corp.contoso.com**, and then click **OK** twice.
7. On the **Local Area Connection Properties** dialog box, click **Close**.
8. Close the **Network Connections** window.

Copy Authentication Manager installation files to the RSA server

1. On the RSA server create the folder C:\RSA Installation.
2. Copy the contents of the RSA Authentication Manager 7.1 SP4 media to the C:\RSA Installation folder.
3. Create the subfolder C:\RSA Installation\License and Token.
4. Copy the RSA license files to C:\RSA Installation\License and Token.

Join the RSA server to the CORP domain

1. Right-click **My Computer**, and click **Properties**.
2. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
3. In **Computer Name**, type **RSA**. In **Member of**, click **Domain**, type **corp.contoso.com**, and click **OK**.
4. When you are prompted for a user name and password, type **User1** and its password, and click **OK**.
5. On the domain welcoming dialog box click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.
9. After the computer has restarted, type **User1** and the password, select CORP in the **Log on to:** drop down list, and click **OK**.

Disable Windows Firewall on RSA

1. Click **Start**, click **Control Panel**, click **System and Security**, and click **Windows Firewall**.
2. Click **Turn Windows Firewall on or off**.
3. **Turn off Windows Firewall** for all settings.
4. Click **OK** and close Windows Firewall.

Install RSA Authentication Manager on the RSA server

1. If the Security Warning message appears at any time during this process, click **Run** to continue.
2. Open the C:\RSA Installation folder and double-click **autorun.exe**.
3. Click **Install Now**, click **Next**, select the top option for the Americas, and click **Next**.
4. Select **I accept the terms of the license agreement**, and click **Next**.
5. Select **Primary Instance**, and click **Next**.
6. In the **Directory Name:** field type **C:\RSA**, and click **Next**.
7. Verify that the server name (RSA.corp.contoso.com) and IP address are correct, and click **Next**.
8. Browse to C:\RSA Installation\License and Token, and click **Next**.
9. On the **Verify license file** page, click **Next**.
10. In the **User ID** field type **Administrator**, and in the **Password** and **Confirm Password** fields type a strong password. Click **Next**.
11. On the log selection screen, accept the defaults and click **Next**.
12. On the summary screen, click **Install**.
13. After installation is complete, click **Finish**.

Configure RSA Authentication Manager

1. If the RSA Security Console does not open automatically, then on the RSA computer desktop double-click “RSA Security Console”.
2. If the security certificate warning / security alert appears, click **Continue to this website** or click **Yes** to proceed, and add this site to trusted sites, if requested.
3. In the **User ID** field type **Administrator** and click **OK**.
4. In the **Password** field type the password for the Administrator account and click **Log On**.
5. Insert Token information.
    a. In the **RSA Security Console** click **Authentication** and click **SecurID Tokens**.
    b. Click **Import Tokens Job**, and then click **Add New**.
    c. In the **Import Options** section click **Browse**. Browse to and select the tokens XML file in the C:\ RSA Installation\License and Token folder and click **Open**.
    d. Click **Submit Job** on the bottom of the page.

6. Create OTP new user.
    a. In the **RSA Security Console** click the **Identity** tab, click **Users**, and click **Add New**.
    b. In the **Last Name:** section type **User**, and in the **User ID:** section type **User1** (UserID must be the same as the AD username used for this lab). In the **Password:** and **Confirm Password:** sections type a strong password. Clear the **'Require user to change password at next logon'** check box and click **Save**.
7. Assign User1 to one of the imported tokens.
    a. On the **Users** page click **User1** and click **SecurID Tokens**.
    b. Click **SecurID Tokens** and click **Assign Token**.
    c. Under the **Serial Number** heading click the first number listed, and click **Assign**.
    d. Click the assigned token, and click **Edit**. In the **SecurID PIN Management** section for **User Authentication Requirement**, select **Do not require PIN (only tokencode)**.
    e. Click **Save and Distribute Token**.
    f. On the **Distribute Software Token** page in the **Basics** section, click **Issue Token File (SDTID)**.
    g. On the **Distribute Software Token** page in the **Token File Options** section, clear the **Enable copy protection** check box. Click **No Password** and **Next**.
    h. On the **Distribute Software Token** page in the **Download File** section, click **Download Now**. Click **Save**. Browse to C:\RSA Installation and click **Save** and **Close**.
    i. Minimize the **RSA Security Console** for use later.
8. Configure Authentication Manager as RADIUS server.
    a. On the RSA computer desktop double-click **"RSA Security Operations Console"**.
    b. If the security certificate warning / security alert appears, click **Continue to this website** or click **Yes** to proceed, and add this site to trusted sites if requested.
    c. Enter the User ID and Password and click **Log On**.
    d. Click **Deployment Configuration – RADIUS – Configure Server**.
    e. On the **Additional Credentials Required** page enter the administrator User ID and Password and click **OK**.
    f. On the **Configure RADIUS Server** page enter the same password used for the administrator user for the **Secrets** and **Master Password**. Enter the Administrator User ID and Password, and click **Configure**.
    g. Verify that the message **'Successfully configured RADIUS server'** is displayed. Click **Done**. Close the **RSA Operations Console**.
    h. Switch back to the **"RSA Security Console"**.
    i. On the **RADIUS** tab click **RADIUS Servers**. Verify that rsa.corp.contoso.com is listed.
9. Configure RSA server as RSA Authentication Client.
    a. On the **RADIUS** tab, click **RADIUS Clients** and **Add New**.
    b. Click the **ANY RADIUS Client** check box.
    c. Type a strong password of your choice in the **Shared Secret** field. You will use this

same password later when configuring EDGE1 for OTP.

d. Leave the **IP Address** field blank, and the **Make / Model** entry as **Standard RADIUS**.

e. Click **Save without RSA Agent**.

10. Create files required for configuring EDGE1 as a RSA Authentication Agent.
    a. On the **Access** tab, highlight **Authentication Agents**, and click **Add New**.
    b. Type **EDGE1** in the **Hostname** field, and click **Resolve IP**.
    c. Notice that the IP address for EDGE1 is now displayed in the **IP Address** field. Click **Save**.
11. Generate a configuration file for the EDGE1 server (AM_Config.zip).
    a. On the **Access** tab, highlight **Authentication Agents**, and click **Generate Configuration File**.
    b. On the **Generate Configuration File** page click **Generate Config File**, and then click **Download Now**.
    c. Click **Save**, browse to C:\ RSA Installation, and click **Save**.
    d. Click **Close** on the **Download Complete** dialog.
12. Generate a node secret file for the EDGE1 server (EDGE1_NodeSecret.zip).
    a. On the **Access** tab, highlight **Authentication Agents**, and click **Manage Existing**.
    b. Click the current configured node EDGE1, and click **Manage Node Secret**.
    c. Check the **Create a new random node secret, and export the node secret to a file** check box.
    d. Enter the same password used for the administrator user in the **Encryption Password** and **Confirm Encryption Password** fields, and click **Save**.
    e. On the **Node Secret File Generated** page click **Download Now**.
    f. On the **File Download** dialog click **Save**, browse to C:\RSA Installation, and click **Save**. Click **Close** on the **Download Complete** dialog.
    g. From the RSA Authentication Manager media copy \auth_mgr\windows-x86_64\am\rsa-ace_nsload\win32-5.0-x86\agent_nsload.exe to C:\RSA Installation.

Create DAProbeUser

1. In the **RSA Security Console** click the **Identity** tab, click **Users**, and click **Add New**.
2. In the **Last Name:** section type **Probe**, and in the **User ID:** section type **DAProbeUser**. In the **Password:** and **Confirm Password:** sections type a strong password. Clear the **'Require user to change password at next logon'** check box and click **Save**.

Install RSA SecurID software token on CLIENT1

Use this procedure to install SecurID software token on CLIENT1.

Install SecurID software token

1. On the CLIENT1 computer, create the folder C:\RSA Files. Copy the file Software_Tokens.zip from C:\RSA Installation on the RSA computer to C:\RSA Files. Extract the file User1_000031701832.SDTID to C:\RSA Files on CLIENT1.
2. Access the RSA SecurID software token media source, and double-click RSASECURIDTOKEN410 in the **SecurID SoftwareToken client app** folder to start the RSA SecurID installation. If the **Open File – Security Warning** message appears, then click **Run**.
3. On the **RSA SecurID Software Token – InstallShield Wizard** dialog click **Next** twice.
4. Accept the license agreement, and click **Next**.
5. On the **Setup Type** dialog select **Typical**, click **Next**, and click **Install**.
6. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
7. Select the **Launch RSA SecurID Software Token** check box, and click **Finish**.
8. Click **Import from File**.
9. Click **Browse**, select C:\RSA Files\User1_000031701832.SDTID, and click **Open**.
10. Click **OK** twice.

Configure EDGE1 as an RSA Authentication Agent

Use this procedure to configure EDGE1 to perform RSA authentication.

Configure the RSA Authentication Agent

1. On EDGE1 open Windows Explorer and create the folder C:\RSA Files. Browse to the RSA ACE Installation media.
2. Copy the files agent_nsload.exe, AM_Config.zip and EDGE1_NodeSecret.zip from the RSA media to C:\RSA Files.
3. Extract the contents of both zip files to the following locations:
   a. C:\Windows\system32\
   b. C:\Windows\SysWOW64\
4. Copy agent_nsload.exe to C:\Windows\SysWOW64\.
5. Open an elevated command prompt and navigate to C:\Windows\SysWOW64.
6. Type **agent_nsload.exe –f nodesecret.rec –p <password>** where <password> is the strong password that you created during the initial RSA configuration. Press Enter.
7. Copy C:\Windows\SysWOW64\securid to C:\Windows\System32.

Configure EDGE1 to support OTP authentication

Use this procedure to configure OTP for DirectAccess, and verify the configuration.

Configure OTP for DirectAccess

1. On EDGE1, open Server Manager, and click **REMOTE ACCESS** in the left pane.
2. Right-click **EDGE1** in the SERVERS pane, and select **Remote Access Management**.
3. Click **Configuration**.
4. In the **DirectAccess Setup** window, under **Step 2 – Remote Access Server**, click **Edit**.
5. Click **Next** three times, and in the **Authentication** section select **Two factor authentication** and **Use OTP**, and ensure that **Use computer certificates** is checked. Verify that the root CA is set to **CN=corp-APP1-CA**. Click **Next**.
6. In the **OTP RADIUS Server** section, double-click the blank **Server Name** field.
7. In the **Add a RADIUS Server** dialog, type **RSA** in the **Server name** field. Click **Change** next to the **Shared secret** field, and type the same password that you used when configuring the RADIUS clients on the RSA server in the **New secret** and **Confirm new secret** fields. Click **OK** twice, and click **Next**.

   Note
   If the RADIUS server is in a domain that is different than the Remote Access server, then the **Server Name** field must specify the FQDN of the RADIUS server.
8. In the **OTP CA Servers** section select APP1.corp.contoso.com, and click **Add**. Click **Next**.
9. On the **OTP Certificate Templates** page click **Browse** to select a certificate template used for the enrollment of certificates that are issued for OTP authentication, and on the **Certificate Templates** dialog box select **DAOTPLogon**. Click **OK**. Click **Browse** to select a certificate template used to enroll the certificate used by the Remote Access server to sign OTP certificate enrollment requests, and on the **Certificate Templates** dialog box select **DAOTPRA**. Click **Ok**. Click **Next**.
10. On the **Remote Access Server Setup** page click **Finish**, and click **Finish** on the **DirectAccess Expert Wizard**.
11. On the **Remote Access Review** dialog box click **Apply**, wait for the DirectAccess policy to be updated, and click **Close**.
12. On the **Start** screen, type **powershell.exe**, right-click **powershell**, click **Advanced**, and click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
13. In the Windows PowerShell window, type **gpupdate /force** and press ENTER.
14. Close and reopen the Remote Access Management Console and verify that all OTP settings are correct.

STEP 5: Verify OTP Health on EDGE1

The following procedures verify that OTP is configured and functioning correctly using DirectAccess Server Health Monitoring on EDGE1.

Verify OTP Health on EDGE1 using DirectAccess Server Health Monitoring

1. On EDGE1, open the **Remote Access Management** console.
2. Click **Operations Status**.
3. Verify that the status of OTP is **Working**.

STEP 6: Test DirectAccess Connectivity from the Homenet Subnet

The DirectAccess one-time password (OTP) deployment is now complete and you can start to test connectivity from the Homenet subnet.

To test OTP functionality from the Homenet subnet on CLIENT1

1. On CLIENT1, make sure that you are logged on as **User1**.
2. On the **Start** screen, type **powershell.exe**, right-click **powershell**, click **Advanced**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. In the Windows PowerShell window, type **gpupdate /force**, and press ENTER.
4. Unplug CLIENT1 from the Corpnet subnet and connect it to the Homenet subnet.
5. On CLIENT1, open Internet Explorer, and in the address bar, type **http://app1.corp.contoso.com/** and press ENTER. Press F5.

   The site should not open.
6. On the **Start** screen, type **RSA**, and click **RSA SecurID Token**.
7. Wait until the RSA SecurID token changes the One-time password, and then click **Copy**.
8. Click the **Network connections** icon in the notification area to access the DA Media Manager.
9. Click **Contoso DirectAccess Connection**, and click **Continue**.
10. Press Control+Alt+Delete, and click the **One-time password (OTP)** tile.
11. Paste the previously copied eight digit tokencode, and click **OK**. Wait for authentication to complete. The DirectAccess Workplace Connection status will now be **Connected**.
12. In Internet Explorer, in the address bar, type **http://app1.corp.contoso.com/** and press ENTER. Press F5. You will see the default IIS website on APP1.
13. In the Internet Explorer address bar, type **http://app2.corp.contoso.com/** and press ENTER. Press F5. You will see the default IIS website on APP2.
14. On the **Start** screen, type **\\app1\files**, and press ENTER.
15. In the **Files** shared folder window, double-click the **Example.txt** file. You will see the contents of the Example.txt file.

16. On the **Start** screen, type **\\app2\files**, and press ENTER.
17. In the **Files** shared folder window, double-click the **New Text Document.txt** file. You will see the contents of the New Text Document.txt file.

STEP 7: Test DirectAccess Connectivity from the Internet

The DirectAccess one-time password (OTP) deployment has been tested from the Homenet subnet, and can now be tested from the Internet.

To test OTP functionality from the Internet on CLIENT1

1. On CLIENT1 make sure that you are logged on as **User1**. Connect CLIENT1 to the Corpnet subnet.
2. On the **Start** screen, type **powershell.exe**, right-click **powershell**, click **Advanced**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. In the Windows PowerShell window, type **gpupdate /force** and press ENTER.
4. Unplug CLIENT1 from the Homenet subnet, connect it to the Internet, and restart the computer.
5. On CLIENT1, open Internet Explorer, and in the address bar, type **http://app1.corp.contoso.com/** and press ENTER. Press F5.

   The site should not open.
6. On the **Start** screen, type **RSA**, and click **RSA SecurID Token**.
7. Wait until the RSA SecurID token changes the One-time password, and then click **Copy**.
8. Click the **Network connections** icon in the notification area to access the DA Media Manager.
9. Click **Workplace Connection**, and click **Continue**.
10. Press Control+Alt+Delete, and click the **One-time password (OTP)** tile.
11. Paste the previously copied eight digit tokencode, and click **OK**. Wait for authentication to complete. The DirectAccess Workplace Connection status will now be **Connected**.
12. In Internet Explorer, in the address bar, type **http://app1.corp.contoso.com/** and press ENTER. Press F5. You will see the default IIS website on APP1.
13. In the Internet Explorer address bar, type **http://app2.corp.contoso.com/** and press ENTER. Press F5. You will see the default IIS website on APP2.
14. On the **Start** screen, type **\\app1\files**, and press ENTER.
15. In the **Files** shared folder window, double-click the **Example.txt** file. You will see the contents of the Example.txt file.
16. On the **Start** screen, type **\\app2\files**, and press ENTER.

17. In the **Files** shared folder window, double-click the **New Text Document.txt** file. You will see the contents of the New Text Document.txt file.

STEP 8: Snapshot the Configuration

This completes the DirectAccess with OTP Configuration test lab. To save this configuration so that you can quickly return to a working DirectAccess with OTP configuration from which you can test other DirectAccess modular test lab guides, test lab guide extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots **OTP DA Configuration**. If your lab uses physical computers, create disk images to save the DirectAccess with OTP Configuration.

Deploy Remote Access in a Cluster

Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. Remote Access can be deployed in a number of enterprise scenarios. This overview provides an introduction to the enterprise scenario for deploying multiple Remote Access servers in a cluster load balanced with Windows Network Load Balancing (NLB) or with an external load balancer (ELB), such as F5 Big-IP.

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**

- Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Scenario description

A cluster deployment gathers multiple Remote Access servers into a single unit, which then acts as a single point of contact for remote client computers connecting over DirectAccess or VPN to the internal corporate network using the external virtual IP (VIP) address of the Remote Access cluster. Traffic to the cluster is load balanced using Windows NLB or with an external load balancer (such as F5 Big-IP).

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- Default load balancing through Windows NLB.
- External load balancers are supported.
- Unicast mode is the default and recommended mode for NLB.
- Changing policies outside of the DirectAccess management console or PowerShell cmdlets is not supported.
- When NLB or an external load balancer is used, the IPHTTPS prefix cannot be changed to anything other than /59.

- The load balanced nodes must be in the same IPv4 subnet.
- In ELB deployments, if manage out is needed, then DirectAccess clients cannot use Teredo. Only IPHTTPS can be used for end-to-end communication.
- Ensure all known NLB/ELB hotfixes are installed.
- ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.

In this scenario

The cluster deployment scenario includes a number of steps:

1. **Deploy a Single Remote Access Server with Advanced Settings**—A single Remote Access server with advanced settings must be deployed before setting up a cluster deployment.
2. **Plan a Remote Access Cluster Deployment**—To build a cluster from a single server deployment a number of additional steps are required, including preparing certificates for the cluster deployment.
3. **Configure a Remote Access Cluster**—This consists of a number of configuration steps, including preparing the single server for Windows NLB or the external load balancer, preparing additional servers to join the cluster, and enabling load balancing.

Practical applications

Gathering multiple servers into a server cluster provides the following:

- Scalability—A single Remote Access server provides a limited level of server reliability and scalable performance. By grouping the resources of two or more servers into a single cluster, you increase the capacity for number of users and throughput.
- High availability—A cluster provides high availability for always-on access. If one server in the cluster fails, remote users can continue to access the corporate network via a different server in the cluster. All servers in the cluster have the same set of cluster virtual IP (VIP) addresses, while still maintaining a unique, dedicated IP addresses for each server.
- Ease-of-management—A cluster allows management of multiple servers as a single entity. Shared settings can easily be set across cluster server. Remote Access settings can be managed from any of the servers in the cluster, or remotely using Remote Server Administration Tools (RSAT). In addition, the entire cluster can be monitored from a single Remote Access Management console.

Roles and features included in this scenario

The following table lists the roles and features required for the scenario:

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the Server Manager console. It encompasses both

Role/feature	How it supports this scenario
	DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services (RRAS), which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: • DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. • RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. Dependencies are as follows: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server and default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI and Command Line Tools • Remote Access module for Windows PowerShell Dependencies include: • Group Policy Management Console

Role/feature	How it supports this scenario
	• RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure
Network Load Balancing	This feature provides load balancing in a cluster using Windows NLB.

Hardware requirements

Hardware requirements for this scenario include the following:

- At least two computers that meet the hardware requirements for Windows Server 2012.
- For the External Load Balancer scenario, dedicated hardware is required (i.e. F5 BigIP).
- In order to test the scenario, at least one computer running Windows 8 or Windows 7, and configured as a DirectAccess client is required.

Software requirements

There are a number of requirements for this scenario:

- Software requirements for single server deployment. For more information see Deploy a Single Remote Access Server with Advanced Settings.
- In addition to software requirements for a single server there are a number of cluster-specific requirements:
  - On each cluster server the IP-HTTPS certificate subject name must match the ConnectTo address. A cluster deployment supports a mix of wildcard and non-wildcard certificates on cluster servers.
  - If the network location server is installed on the Remote Access server, on each cluster server the network location server certificate must have the same subject name. In addition, the name of the network location server certificate must not have the same name as any server in the DirectAccess deployment.
  - IP-HTTPS and network location server certificates must be issued using the same method with which the certificate to the single server was issued. For example, if the single server uses a public certification authority (CA) then all servers in the cluster must have a certificate issued by a public CA. Or if the single server uses a self-signed certificate for IP-HTTPS then all servers in the cluster must do likewise.
  - The IPv6 prefix assigned to DirectAccess client computers on server clusters must be 59 bits. If VPN is enabled, the VPN prefix must also be 59 bits.

Known issues

The following are known issues when configuring a cluster scenario:

- After configuring DirectAccess in an IPv4-only deployment with a single network adapter, and after the default DNS64 (the IPv6 address which contains ":3333::") is automatically configured on the network adapter, attempting to enable load-balancing via the Remote Access Management console causes a prompt for the user to supply an IPv6 DIP. If an IPv6 DIP is supplied, the configuration fails after clicking **Commit** with the error: The parameter is incorrect.

  To resolve this issue:

  a. Download the backup and restore scripts from Back up and Restore Remote Access Configuration.

  b. Back up your Remote Access GPOs using the downloaded script Backup-RemoteAccess.ps1

  c. Attempt to enable load balancing until the step at which it fails. On the Enable Load Balancing dialog box, expand the details area, right-click in the details area, and then click **Copy Script**.

  d. Open Notepad, and paste the contents of the clipboard. For example:

  ```
  Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress
  @('10.244.4.19 /255.255.255.0','fdc4:29bd:abde:3333::2/128') -
  InternetVirtualIPAddress @('fdc4:29bd:abde:3333::1/128',
  '10.244.4.21 /255.255.255.0') -ComputerName
  'DA1.domain1.corp.contoso.com' -Verbose
  ```

  e. Close any open Remote Access dialog boxes and close the Remote Access Management console.

  f. Edit the pasted text and remove the IPv6 addresses. For example:

  ```
  Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress
  @('10.244.4.19 /255.255.255.0') -InternetVirtualIPAddress
  @('10.244.4.21 /255.255.255.0') -ComputerName
  'DA1.domain1.corp.contoso.com' -Verbose
  ```

  g. In an elevated PowerShell window, run the command from the previous step.

  h. If the cmdlet fails while it is running (not due to incorrect input values), run the command Restore-RemoteAccess.ps1 and follow instructions to make sure that the integrity of your original configuration is maintained.

  i. You can now open the Remote Access Management console again.

Plan a Remote Access Cluster Deployment

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the planning steps required in order to deploy a cluster of Windows Server 2012 Remote Access servers.

- Step 1: Plan an Advanced Single Server Deployment. This step includes planning for the infrastructure required to deploy a single server. It includes planning for network and server settings, certificate requirements, DNS settings, network location server deployment, DirectAccess management servers, Active Directory settings, and Group Policy Objects (GPOs).
- Step 2: Plan Cluster Servers .
- Step 3: Plan a Load-Balanced Cluster Deployment.
- Step 4: Record your planning decisions for Remote Access advanced deployment. This record can be used as a job aid for everyone involved in completing the deployment steps.

After you have completed these planning steps, see Configure a Remote Access Cluster. For instructions on configuring a cluster deployment as a proof of concept in a lab environment, see Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB.

Step 1: Plan an Advanced Single-Server Deployment

The first step in planning a Windows Server® 2012 Remote Access load-balanced cluster is to plan and configure an advanced single server deployment.

Plan a single server deployment

Before you deploy a Remote Access load-balanced cluster, make sure that you have completed all the planning steps to deploy a single Remote Access server. See Plan an Advanced Remote Access Deployment. Then continue to Step 2: Plan Cluster Servers.

Step 2: Plan Cluster Servers

After deploying a single Remote Access server, plan to add additional servers to the cluster.

Task	Description
2.1 Installing roles and features.	For each server that will be added to the cluster, plan for installation of the Remote

Task	Description
	Access role and the Windows NLB feature (if needed), plan the topology, IP addressing, routing and forwarding.
2.2 Configure server settings	Configure settings for each server that will be added to the cluster. Note that you can configure a load balanced cluster of servers using virtual machines. In order for routing and connectivity to work correctly, you must configure the virtual machines to use MAC address spoofing.

2.1 Installing roles and features

For each server you want to join to the cluster, plan to install the Remote Access role. In addition plan to install the Network Load Balancing (NLB) feature if you want to load balance traffic to the cluster using Windows NLB. For more information see Network Load Balancing overview.

2.2 Configure server settings

For each server that will be added to the cluster, plan IP address and domain settings. Note the following:

1. Servers in the cluster must all belong to the same domain.
2. Servers in the cluster must be located on the same subnet.
3. Each server in the cluster should have the same number of network adapters in use for the DirectAccess deployment.

For more information about configuring server settings, see the single server planning guide — Step 1: Plan the Remote Access Infrastructure.

When you load balance the cluster using Windows NLB the following Windows NLB settings are applied:

1. Operation mode—Unicast. This can be changed to multicast using NLB Manager. This setting cannot be modified in the Remote Access Management console.
2. Load weight factor—Defined as Equal, where all cluster servers have equal load.
3. Filtering mode—Traffic will be load balanced across multiple hosts.
4. Affinity—Single affinity is defined.
5. Protocols—Both

See also

- Step 3: Plan a Load-Balanced Cluster Deployment

Step 3: Plan a Load-Balanced Cluster Deployment

The next step is to plan the load-balancing configuration and cluster deployment.

Task	Description
3.1 Plan load balancing	Decide whether to use Windows Network Load Balancing (NLB), or an external load balancer (ELB).
3.2 Plan IP-HTTPS	If a self-signed certificate is not used, the Remote Access server needs an SSL certificate on each server in the cluster, in order to authenticate IP-HTTPS connections.
3.3 Plan for VPN client connections	Note the requirements for VPN client connections.
3.4 Plan the network location server	If the network location server website is hosted on the Remote Access server, and a self-signed certificate is not used, ensure that each server in the cluster has a server certificate to authenticate the connection to the website.

3.1 Plan load balancing

Remote Access can be deployed on a single server, or on a cluster of Remote Access servers. Traffic to the cluster can be load balanced to provide high availability and scalability for DirectAccess clients. There are two load balancing options:

- **Windows NLB**—Windows NLB is a Windows Server® 2012 server feature. To use it, you do not require additional hardware because all the servers in the cluster are responsible for managing the traffic load. Windows NLB supports a maximum of eight servers in a Remote Access cluster.
- **External load balancer**—Using an external load balancer requires external hardware to manage the traffic load between the Remote Access cluster servers. In addition, using an external load balancer supports a maximum of 32 Remote Access servers in a cluster. Some points to keep in mind when configuring external load balancing are:
  - The administrator must ensure that the Virtual IPs configured through the Remote Access load balancing wizard are used on the external load balancers (like F5 Big-Ip Local Traffic Manager system). When external load balancing is enabled, the IP addresses on the external and internal interfaces will be promoted to Virtual IP addresses, and have to be plumbed on the load balancers. This is done so that the administrator does not have

to change the DNS entry for the public name of the cluster deployment. Also, the IPsec tunnel endpoints are derived from the server IPs. If the administrator provides separate Virtual IPs, then the client will not be able to connect to the server. See example for configuring DirectAccess with External Load Balancing in 3.1.1 External Load Balancer configuration example.

- Many external load balancers (including F5) do not support load balancing of 6to4 and ISATAP. If the Remote Access server is an ISATAP router, the ISATAP function should be moved to a different computer. Also, when ISATAP function is on a different computer, the DirectAccess servers must have native IPv6 connectivity with the ISATAP router. Note that this connectivity should be present before configuring DirectAccess.
- For external load balancing, if Teredo has to be used, all the Remote Access servers must have two consecutive public IPv4 addresses as dedicated IP addresses. The Virtual IPs of the cluster also must have two consecutive public IPv4 addresses. This is not true for Windows NLB where only the Virtual IPs of the cluster must have two consecutive public IPv4 addresses. If Teredo is not used, two consecutive IP addresses are not required.
- The administrator can switch from Windows NLB to external load balancer and vice versa. Note that the administrator cannot switch from external load balancer to Windows NLB if he has more than 8 servers in the external load balancer deployment.

3.1.1 External Load Balancer configuration example

This section describes the configuration steps for enabling an external load balancer on a new Remote Access deployment. When using an external load balancer, the Remote Access cluster may look like the figure below, where the Remote Access servers are connected to the corporate network through a load balancer on the internal network, and to the Internet through a load balancer connected to the external network:

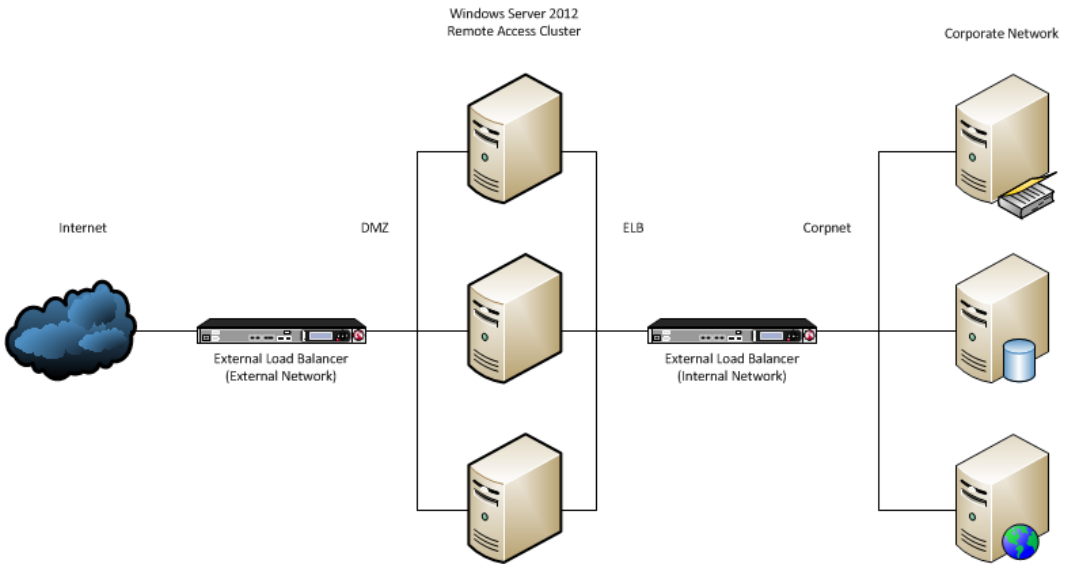

Planning information

1. External VIPs (IPs that the client will use to connect to Remote Access) were decided to be 131.107.0.102, 131.107.0.103
2. Load balancer on external network self-IPs – 131.107.0.245 (Internet), 131.107.1.245

   The perimeter network (also known as demilitarized zone and DMZ) is between the load balancer on the external network and the Remote Access server.
3. IP addresses for the Remote Access server on the perimeter network – 131.107.1.102, 131.107.1.103
4. IP addresses for the Remote Access server on the ELB network (i.e. between the Remote Access server and the load balancer on the internal network) – 30.11.1.101, 2006:2005:11:1::101
5. Load balancer on internal network self-IPs – 30.11.1.245 2006:2005:11:1::245 (ELB), 30.1.1.245 2006:2005:1:1::245 (Corpnet)
6. Internal VIPs (IP addresses used for client web-probe and for network location server, if installed on the Remote Access servers) were decided to be 30.1.1.10, 2006:2005:1:1::10

Steps

1. Configure the Remote Access server's external network adapter (that is connected to the perimeter network) with addresses 131.107.0.102, 131.107.0.103. This step is required for the DirectAccess configuration to detect the correct IPsec tunnel endpoints.
2. Configure the Remote Access server's internal network adapter (that is connected to the ELB network) with the web-probe/network location server IP addresses (30.1.1.10, 2006:2005:1:1::10). This step is required for allowing clients to access the web-probe IP, so the network connectivity assistant correctly indicates the connection status to DirectAccess. This step also allows access to the network location server, if it is configured on the DirectAccess server.

   Note

   Make sure that the domain controller is reachable from the Remote Access server with this configuration.
3. Configure DirectAccess single server on the Remote Access server.
4. Enable external load balancing in the DirectAccess configuration. Use 131.107.1.102 as the external dedicated IP address (DIP) (131.107.1.103 will be selected automatically), use 30.11.1.101, 2006:2005:11:1::101 as the internal DIPs.
5. Configure the external virtual IPs (VIP) on the external load balancer with addresses 131.107.0.102 and 131.107.0.103. Also, configure the internal VIPs on the external load balancer with addresses 30.1.1.10 and 2006:2005:1:1::10.
6. The Remote Access server will now be configured with the planned IP addresses, and the external and internal IP addresses for the cluster will be configured according to the planned IP addresses.

3.2 Plan IP-HTTPS

1. **Certificate requirements**—During deployment of the single Remote Access server you selected to use either an IP-HTTPS certificate issued by a public or internal certification authority (CA), or a self-signed certificate. For the cluster deployment, you must use identical type of certificate on each member of the Remote Access cluster. That is, if you used a certificate issued by a public CA (recommended), you must install a certificate issued by a public CA on each member of the cluster. The subject name of the new certificate should be identical to the subject name of the IP-HTTPS certificate currently used in your deployment. Note that if you are using self-signed certificates these will be configured automatically on each server during cluster deployment.
2. **Prefix requirements**—Remote Access enables load balancing of both SSL-based traffic and DirectAccess traffic. To load balance all IPv6 based DirectAccess traffic, Remote Access must examine the IPv4 tunneling for all transition technologies. Because IP-HTTPS traffic is encrypted, examining the content of the IPv4 tunnel is not possible. To enable IP-HTTPS traffic to be load balanced, you must allocate a wide enough IPv6 prefix so that a different IPv6 /64 prefix can be assigned to every cluster member. You can configure a maximum of 32 servers in a load balanced cluster; therefore, you must specify a /59 prefix. This prefix must be routable to the internal IPv6 address of the Remote Access cluster, and is configured in the Remote Access Server Setup wizard.

   Note
   The prefix requirements are relevant only in an IPv6 enabled internal network (IPv6-only, or IPV4+IPv6). In an IPv4 only corporate network, the client prefix is automatically configured and the administrator cannot change it.

3.3 Plan for VPN client connections

There are a number of considerations for VPN client connections:

- VPN client traffic cannot be load balanced if VPN client addresses are allocated using DHCP. A static address pool is required.
- RRAS can be enabled on a load-balanced cluster that has been deployed for DirectAccess only, using **Enable VPN** on the Tasks pane of the Remote Access Management console.
- Any VPN changes completed in the Routing and Remote Access Management console (rrasmgmt.msc) will have to be replicated manually on all Remote Access servers in the cluster.
- To enable VPN IPv6 client traffic to be load balanced, you must specify a 59-bit IPv6 prefix.

3.4 Plan the network location server

If you are running the network location server website on the single Remote Access server, during deployment you selected to use either a certificate issued by an internal certification authority (CA), or a self-signed certificate. Note the following:

1. Each member of the Remote Access cluster must have a certificate for the network location server that corresponds to the DNS entry for the network location server website.

2. The certificate for each cluster server must be issued in the same way as the certificate on the single Remote Access server current network location server certificate. For example, if you used a certificate issued by an internal CA, you must install a certificate issued by the internal CA on each member of the cluster.
3. If you used a self-signed certificate, a self-signed certificate will be configured automatically for each server during cluster deployment.
4. The subject name of the certificate must not be identical to the name of any of the servers in the Remote Access deployment.

See also

- Configure a Remote Access Cluster

Configure a Remote Access Cluster

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the configuration steps required to deploy a single Windows Server® 2012 Remote Access server in a load-balanced cluster.

- Step 1: Implement a Single-Server Remote Access Deployment.
- Step 2: Prepare Cluster Servers.
- Step 3: Configure a Load-Balanced Cluster.
- Step 4: Verify the Cluster.

Step 1: Implement a Single-Server Remote Access Deployment

The first configuration step to deploy Remote Access in a load-balanced cluster is to implement an advanced single server deployment and then plan to add servers to the cluster.

Implement a single server deployment

Before you can configure a load-balanced cluster, you must configure an advanced single server Remote Access deployment as described in Deploy a Single Remote Access Server with Advanced Settings.

See also

- Step 2: Prepare Cluster Servers

Step 2: Prepare Cluster Servers

Before you can configure a cluster deployment, you prepare additional servers to add to the cluster.

Task	Description
2.1 Configure the Remote Access infrastructure	On each server you want to add to the cluster, configure the server topology, IP addressing, routing, and forwarding. If you configure a load balanced cluster of virtual machines, you must configure the virtual machines to use MAC address spoofing. In addition, join each server to the same domain, and connect all the servers to the same subnet.
2.2 Install the Remote Access role	On each additional server you want to add to the cluster, install the Remote Access role
2.3 Install NLB	On the deployed Remote Access server and on each additional server you want to add to the cluster, install the NLB feature. Note that this step is not required when using an External Load Balancer.

2.1 Configure the Remote Access infrastructure

To configure a Remote Access cluster, you must configure the server topology, IP addressing, routing, and forwarding on every server that will form part of the cluster.

To configure the Remote Access infrastructure

1. Configure each of the servers that will be in the cluster with the same topology as the first Remote Access server.
2. Configure each of the servers that will be in the cluster with appropriate IP addressing, routing, and forwarding based on the configuration of the first Remote Access server. Note that all the servers in the cluster must be connected to the same subnet.
3. Join each of the servers that will be in the cluster to the same domain as the first Remote Access server.

To install the Remote Access role on DirectAccess servers

1. On the DirectAccess server, in the Server Manager console, in the **Dashboard**, click **Add**

roles and features.

2. Click **Next** three times to get to the server role selection screen.
3. On the **Select Server Roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.
7. Repeat this procedure on all servers that you want to be cluster members.

2.2 Install the Remote Access role

To configure a Remote Access cluster, you must install the Remote Access role on every server that will form a part of the cluster.

2.3 Install NLB

To configure a Remote Access cluster, you must install the Network Load Balancing feature on every server that will form a part of the cluster.

Note
This step is not required if an external load balancer is used.

To install the NLB role

1. On the DirectAccess server, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** four times to get to the server feature selection screen.
3. On the **Select features** dialog, select **Network Load Balancing**, click **Add Features**, click **Next**, and then click **Install**.
4. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.
5. Repeat this procedure on all servers that you want to be cluster members.

See also

- Step 3: Configure a Load-Balanced Cluster

Step 3: Configure a Load-Balanced Cluster

After preparing servers for the cluster, configure load-balancing on the single server, configure the required certificates, and deploy the cluster.

Task	Description
3.1 Configure the IPv6 prefix	If the corporate environment is IPv4+IPv6, or IPv6-only, then on the single Remote Access server, ensure that the IPv6 prefix assigned to DirectAccess client computers is large enough to cover all of the servers in your cluster.
3.2 Enable load balancing	Enable load balancing on the single Remote Access server.
3.3 Install the IP-HTTPS certificate	Each server in the cluster requires a server certificate to authenticate IP-HTTPS connection. Export the IP-HTTPS certificate from the single Remote Access server and deploy it on each server you will add to the cluster. This is required only if using non-self-signed certificates.
3.4 Install the network location server certificate	If the single server has the network location server deployed locally, then you will need to deploy the network location server certificate on each server in the cluster. If the network location server is hosted on an external server, a certificate on each server is not required. This is required only if using non-self-signed certificates.
3.5 Add servers to the cluster	Add all servers to the cluster. Remote Access must not be configured on the servers to be added.
3.6 Remove a server from the cluster	Instructions for removing a server from the cluster.
3.7 Disable load balancing	Instructions for disabling load balancing.
3.8 Change from unicast to multicast	Instructions for changing from unicast to multicast.

Note

The IP address that is selected for the DIP must not be in use on the network adapters of the first Remote Access server in the cluster. Beginning the DirectAccess deployment with both VIP and DIP added to the network adapter will result in failure.

Note

Make sure not to use a DIP that is already present on another machine on the network.

3.1 Configure the IPv6 prefix

To configure the prefix

1. On the Remote Access server, click **Start**, and then click **Remote Access Management**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**.
3. In the middle pane of the console, in the **Step 2 DirectAccess Server** area, click **Edit**.
4. Click **Prefix Configuration**. On the **Prefix Configuration** page, in **IPv6 prefix assigned to DirectAccess client computers**, enter the IPv6 prefix used for DirectAccess client computers with a subnet length of 59, for example, **2001:db8:1:1000::/59**. If VPN were also enabled with IPv6, then an IPv6 prefix would be displayed, and the subnet length would need to be changed to 59. Click **Next**.
5. In the middle pane of the console, click **Finish**.
6. On the **Remote Access Review** dialog box, review the configuration settings, and then click **Apply**. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.

3.2 Enable load balancing

Do this step using Windows PowerShell

To enable load balancing

1. On the configured DirectAccess server, click **Start**, and then click **Remote Access Management**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the left pane, click **Configuration**, and then in the **Tasks** pane, click **Enable Load Balancing**.
3. In the Enable Load Balancing Wizard, click **Next**.
4. Depending on what you chose in planning steps:
   a. Windows NLB: On the **Load Balancing Method** page, click **Use Windows Network Load Balancing (NLB)**, and then click **Next**.
   b. External load balancer: On the **Load Balancing Method** page, click **Use an external load balancer**, and then click **Next**.
5. In a single network adapter deployment, on the **Dedicated IP Addresses** page, do the following, and then click **Next**:
   a. In the **IPv4 address** box, enter the new IPv4 address for this Remote Access server; the current IPv4 address will be the virtual IP address (VIP) of the load-balanced cluster. In the **Subnet mask** box, enter the subnet mask.

b. If the corporate environment is native IPv6, then in the **IPv6 address** box, enter the new IPv6 address for this Remote Access server; the current IPv6 address will be the VIP of the load-balanced cluster. In the **Subnet prefix length** box, enter the subnet prefix length.

6. In a two network adapter deployment, on the **External Dedicated IP Addresses** page, do the following, and then click **Next**:
   a. In the **IPv4 address** box, enter the new external IPv4 address for this Remote Access server; the current IPv4 address will be the virtual IP address (VIP) of the load balancing cluster. In the **Subnet mask** box, enter the subnet mask.
   b. If there are currently native IPv6 addresses configured on the internet-facing network adapter of the Remote Access server, in the **IPv6 address** box, enter the new external IPv6 address for this Remote Access server; the current IPv6 address will be the VIP of the load balancing cluster. In the **Subnet prefix length** box, enter the subnet prefix length.
7. In a two network adapter deployment, on the **Internal Dedicated IP Addresses** page, do the following, and then click **Next**:
   a. In the **IPv4 address** box, enter the new internal IPv4 address for this Remote Access server; the current IPv4 address will be the VIP of the load balancing cluster. In the **Subnet mask** box, enter the subnet mask.
   b. If the corporate environment is native IPv6, then in the **IPv6 address** box, enter the new internal IPv6 address for this Remote Access server; the current IPv6 address will be the VIP of the load balancing cluster. In the **Subnet prefix length** box, enter the subnet prefix length.
8. On the **Summary** page, click **Commit**.
9. On the **Enable Load Balancing** dialog box, click **Close**.
10. In the Enable Load Balancing Wizard, click **Close**.

Note

If external load balancing is being used, note the Virtual IPs and provide them as on the external load balancers.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

If you chose to use Windows NLB in the planning steps, then execute the following:

```
Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress "2.1.1.20/255.255.255.0" -
InternalDedicatedIPAddress @("10.1.1.30/255.255.255.0","3ffe::20/64") -
InternetVirtualIPAddress @("2.1.1.1/255.255.255.0","2.1.1.2/255.255.255.0") -
InternalVirtualIPAddress @("10.1.1.2/255.255.255.0","3ffe::2/64")
```

If you chose to use an external load balancer in the planning steps: then execute the following:

```
Set-RemoteAccessLoadBalancer -InternetDedicatedIPAddress "2.1.1.20/255.255.255.0" -
InternalDedicatedIPAddress @("10.1.1.30/255.255.255.0","3ffe::20/64") -
UseThirdPrtyLoadBalancer
```

Note

It is recommended to not include changes to load-balancer settings with changes to any other settings, if you are using staging GPOs, see 1.8.4 Managing Remote Access GPOs with limited permissions. Any changes to load-balancer settings must be applied first and then other configuration changes should be made. Also, after configuring load-balancer on a new DirectAccess server, please allow some time for the IP changes to be applied and replicated across the DNS servers in the enterprise, before you change other DirectAccess settings related to the new cluster.

3.3 Install the IP-HTTPS certificate

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To install the IP-HTTPS certificate

1. On the configured Remote Access server, click **Start**, type **mmc** and press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Finish**, and then click **OK**.
4. In the left pane of the console, navigate to **Certificates (Local Computer)\Personal\Certificates**. Right-click the IP-HTTPS certificate, point to **All Tasks** and click **Export**.
5. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
6. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
7. On the **Export File Format** page, click **Personal Information Exchange – PKCS #12 (.PFX)**, and then click **Next**.
8. On the **Security** page, select the **Password** check box, enter a password in the **Password** box and confirm the password, and then click **Next**.
9. On the **File to Export** page, enter a name for the certificate file and save it to the desktop, and then click **Next**.
10. On the **Completing the Certificate Export Wizard** page, click **Finish**.
11. On the **Certificate Export Wizard** dialog box, click **OK**.
12. Copy the certificate to all servers that you want to be cluster members.
13. On the new DirectAccess server, click **Start**, type **mmc** and press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want,

and then click **Yes**.

14. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
15. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Finish**, and then click **OK**.
16. In the left pane of the console, navigate to **Certificates (Local Computer)\Personal\Certificates**. Right click the **Certificates** node, point to **All Tasks**, and then click **Import**.
17. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
18. On the **File to Import** page, click **Browse** to locate the certificate. Select the certificate and then click **Next**.
19. On the **Private key protection** page, in the **Password** box, type the password, and then click **Next**.
20. On the **Certificate Store** page, click **Next**.
21. On the **Completing the Certificate Import Wizard** page, click **Finish**.
22. On the **Certificate Import Wizard** dialog box, click **OK**.
23. Repeat steps 13-22 on all servers that you want to be cluster members.

3.4 Install the network location server certificate

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To install a certificate for network location

1. On the Remote Access server, click **Start**, type **mmc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, click the Web Server certificate template, and then click **More information is required to enroll for this certificate**.

   If the Web Server certificate template does not appear, ensure that the Remote Access server computer account has enroll permissions for the Web Server certificate template. For more information, see **Configure Permissions on the Web Server Certificate Template**.
8. On the **Subject** tab of the **Certificate Properties** dialog box, in **Subject name**, for **Type**, select **Common name**.

9. In **Value**, type the fully qualified domain name (FQDN) for the intranet name of the network location server website (for example, nls.corp.contoso.com), and then click **Add**.
10. Click **OK**, click **Enroll**, and then click **Finish**.
11. In the details pane of the Certificates snap-in, verify that a new certificate with the FQDN was enrolled with **Intended Purposes** of **Server Authentication**.
12. Right-click the certificate, and then click **Properties**.
13. In **Friendly Name**, type **Network Location Certificate**, and then click **OK**.

    Tip
    Steps 12 and 13 are optional, but make it easier for you to select the certificate for network location when configuring Remote Access.
14. Repeat this procedure on all servers that you want to be cluster members.

3.5 Add servers to the cluster

Do this step using Windows PowerShell

To add servers to the cluster

1. On the configured DirectAccess server, click **Start**, and then click **Remote Access Management**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**. In the **Tasks** pane, under **Load Balanced Cluster**, click **Add or Remove Servers**.
3. On the **Add or Remove Servers** dialog box, click **Add Server**.
4. On the **Add a Server** dialog box, on the **Select Server** page, enter the name of the additional Remote Access server, and then click **Next**.
5. On the **Network Adapters** page, do one the following:
   - If you are deploying a topology with two network adapters, in **External adapter**, select the adapter that is connected to the external network. In **Internal adapter**, select the adapter that is connected to the internal network.
   - If you are deploying a topology with one network adapter, in **Network adapter**, select the adapter that is connected to the internal network.
6. On the **Network Adapters** page, in **Select the certificate used to authenticate IP-HTTPS connections**, click **Browse** to locate and select the IP-HTTPS certificate, and then click **Next**.
7. On the **Network Location Server** page, click **Browse** to select the certificate for the network location server website running on the Remote Access server, and then click **Next**.

   Note
   The **Network Location Server** page appears only when the network location server website is running on the Remote Access server.

Note

If VPN were also configured on the Remote Access server, then you would be asked to add the VPN IP address pool information at this point.

8. On the **Summary** page, click **Add**.
9. On the **Completion** page, click **Close**.
10. Repeat this procedure for all Remote Access servers to be added to the cluster.
11. On the **Add or Remove Servers** dialog box, click **Commit**.
12. On the **Adding and Removing Servers** dialog box, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-RemoteAccessLoadBalancerNode -RemoteAccessServer <server name>
```

Note

If VPN has not been enabled in a load balanced cluster, you should not provide any VPN address ranges when adding a new server to the cluster using Windows PowerShell cmdlets. If you have done so by mistake, remove the server from the cluster, and add it again to the cluster without specifying the VPN address ranges.

3.6 Remove a server from the cluster

Do this step using Windows PowerShell

To remove a server from the cluster

1. On the configured Remote Access server, click **Start**, and then click **Remote Access Management**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**. In the **Tasks** pane, under **Load Balanced Cluster**, click **Add or Remove Servers**.
3. On the **Add or Remove Servers** dialog box, select the Remote Access server you want to remove, and then click **Remove Server**.
4. On the **Remove Server Warning** dialog box, make sure you chose the right server, and then click **OK**.
5. Repeat this procedure for all Remote Access servers to be removed from the cluster.
6. On the **Add or Remove Servers** dialog box, click **Commit**.
7. On the **Adding and Removing Servers** dialog box, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Remove-RemoteAccessLoadBalancerNode -RemoteAccessServer <server name>
```

3.7 Disable load balancing

Do this step using Windows PowerShell

To disable load balancing

1. On the configured DirectAccess server, click **Start**, and then click **Remote Access Management**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**. In the **Tasks** pane, under **Load Balanced Cluster**, click **Disable Load Balancing**.
3. On the **Disable Load Balancing** dialog box, click **Ok**.
4. On the **Disable Load Balancing** dialog box, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
set-RemoteAccessLoadBalancer -disable
```

Disabling load balancing will remove Remote Access settings and NLB settings (if configured) from all servers except the server from which it is being executed. On this Remote Access server, NLB settings will be removed (if it was configured) but Remote Access settings will remain.

Clicking **Remove configuration settings** will remove Remote Access and NLB (if configured) from all servers in the deployment.

Notes

- If Remote Access is uninstalled when load balancing is deployed, all the servers are left with DIPs. The VIPs are removed. This causes all routes in the corporate network that are targeted to the VIPs addresses to fail. This also affects DNS entries which were resolved to the VIPs, such as the network location server certificate subject name. To avoid this issue, disable load balancing, which leaves the VIPs on the last Remote Access server, and then uninstall Remote Access.
- After using the **Set-RemoteAccessLoadBalancer** cmdlet to disable load balancing, wait for 2 minutes before running any other cmdlet. This should also be done in any scripts that run another cmdlet after the **Set-RemoteAccessLoadBalancer -disable** cmdlet.

- Disabling load balancing changes the virtual IP address of the cluster into a dedicated IP address. As a result, any operation that queries for the name of the server will fail until the cached DNS entry on the server expires. Make sure that you do not run any Remote Access PowerShell cmdlets after disabling load balancing until the cache on the server has expired. This problem is more common if you try to disable load balancing on a machine from another machine that is in another domain. This also occurs if you disable load balancing from the Remote Access Management console and may prevent the configuration from loading. The configuration will load after the cache has expired or has been flushed.

3.8 Change from unicast to multicast

Note
This procedure applies only to NLB.

To change from unicast to multicast

1. Open the NLB Manager.
2. Stop the Internal and the Internet Clusters.
3. Change both clusters to use Multicast. Make changes to change the Internal Cluster first, followed by the Internet Cluster.
4. Once changed, start the Internal Cluster and wait for both nodes to enter the converged state.
5. Start the Internet Cluster.

Changes to the Internal Cluster need to precede changes to the Internet Cluster because of the BDA master-subordinate (also known as slave) concept.

See also

- Step 4: Verify the Cluster

Step 4: Verify the Cluster

This topic describes how to verify that you have correctly configured your DirectAccess cluster deployment.

To verify access to internal resources through the cluster

1. Connect a DirectAccess client computer to the corporate network and obtain the group policy.
2. Connect the client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

3. Test connectivity through each server in the cluster by turning off, or disconnecting from the external network, all but one of the cluster servers. On the client computer, attempt to access corporate resources. Repeat the test on a different cluster server.

   You should be able to access all corporate resources through each cluster server.

Deploy Multiple Remote Access Servers in a Multisite Deployment

Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. Remote Access can be deployed in a number of enterprise scenarios. This overview provides an introduction to the enterprise scenario for deploying Remote Access servers in a multisite configuration.

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**

    - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Scenario description

In a multisite deployment two or more Remote Access servers or server clusters are deployed and configured as different entry points in a single location, or in dispersed geographical locations. Deploying multiple entry points in a single location allows for server redundancy, or for the alignment of Remote Access servers with existing network architecture. Deployment by geographical location ensures efficient use of resources, as remote client computers can connect to internal network resources using an entry point closest to them. Traffic across a multisite deployment can be distributed and balanced with an external global load balancer.

A multisite deployment supports client computers running Windows 8 or Windows 7. Client computers running Windows 8 automatically identify an entry point, or the user can manually select an entry point. Automatic assignment occurs in the following priority order:

1. Use an entry point selected manually by the user.
2. Use an entry point identified by an external global load balancer if one is deployed.
3. Use the closest entry point identified by the client computer using an automatic probe mechanism.

Support for clients running Windows 7 must be manually enabled on each entry point, and selection of an entry point by these clients is not supported.

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- Deploy a Single Remote Access Server with Advanced Settings must be deployed before a multisite deployment.
- Windows 7 clients will always connect to a specific site. They will not be able to connect to the closest site based on location of the client (unlike Windows 8, windows 8.1 clients).
- Changing policies outside of the DirectAccess management console or PowerShell cmdlets is not supported.
- A Public Key Infrastructure must be deployed.

For more information see: Test Lab Guide Mini-Module: Basic PKI for Windows Server 2012.

- The corporate network must be IPv6 enabled. If you are using ISATAP, you should remove it and use native IPv6.

In this scenario

The multisite deployment scenario includes a number of steps:

1. **Deploy a Single Remote Access Server with Advanced Settings**—A single Remote Access server with advanced settings must be deployed before setting up a multisite deployment.
2. **Plan a Multisite Deployment**—To build a multisite deployment from a single server a number of additional planning steps are required, including compliance with multisite prerequisites, and planning for Active Directory security groups, Group Policy Objects (GPOs), DNS, and client settings.
3. **Configure a Multisite Deployment**—This consists of a number of configuration steps, including preparation of the Active Directory infrastructure, configuration of the existing Remote Access server, and the addition of multiple Remote Access servers as entry points to the multisite deployment.
4. **Troubleshoot a Multisite Deployment**—This troubleshooting section describes a number of the most common errors that can occur when deploying Remote Access in a multisite deployment.

Practical applications

A multisite deployment provides the following:

- Improved performance—A multisite deployment allows client computers accessing internal resources using Remote Access to connect using the closest and most suitable entry point. Client access internal resources efficiently, and the speed of client Internet requests routed via DirectAccess is improved. Traffic across entry points can be balanced using an external global load balancer.
- Ease-of-management—Multisite allows administrators to align the Remote Access deployment to an Active Directory sites deployment, providing a simplified architecture. Shared settings can easily be set across entry point servers or clusters. Remote Access settings can be managed from any of the servers in the deployment, or remotely using Remote Server Administration Tools (RSAT). In addition, the entire multisite deployment can be monitored from a single Remote Access Management console.

Roles and features included in this scenario

The following table lists the roles and features used in this scenario.

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the

Role/feature	How it supports this scenario
	Server Manager console. It encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services (RRAS), which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: • DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. • RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. Dependencies are as follows: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server and default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI and Command Line Tools • Remote Access module for Windows PowerShell Dependencies include:

Role/feature	How it supports this scenario
	• Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

- At least two Remote Access computers to be gathered into a multisite deployment. Hardware requirements for these computers are described in Deploy a Single Remote Access Server with Advanced Settings.
- In order to test the scenario, at least one computer running Windows 8 and configured as a DirectAccess client is required. To test the scenario for clients running Windows 7 at least one computer running Windows 7 is required.
- To load balance traffic across entry point servers, a third-party external global load balancer is required.

Software requirements

Software requirements for this scenario include the following:

- Software requirements for single server deployment. For more information see Deploy a Single Remote Access Server with Advanced Settings.
- In addition to software requirements for a single server there are a number of multisite-specific requirements:
    - IPsec authentication requirements—In a multisite deployment DirectAccess must be deployed using IPsec machine certificate authentication. The option to perform IPsec authentication using the Remote Access server as a Kerberos proxy is not supported. An internal CA is required to deploy the IPsec certificates.
    - IP-HTTPS and network location server requirements—Certificates required for IP-HTTPS and the network location server must be issued by a CA. The option to use certificates that are automatically issued and self-signed by the Remote Access server is not supported. Certificates can be issued by an internal CA or by a third-party external CA.
    - Active Directory requirements—At least one Active Directory site is required. The Remote Access server should be located in the site. For faster update times, it is recommended that each site has a writeable domain controller, though this is not mandatory.
    - Security group requirements—Requirements are as follows:

- A single security group is required for all Windows 8 client computers from all domains. It is recommended to create a unique security group of these clients for each domain.
- A unique security group containing Windows 7 computers is required for each entry point configured to support Windows 7 clients. It is recommended to have a unique security group for each entry point in each domain.
- Computers should not be included in more than one security group that includes DirectAccess clients. If clients are included in multiple groups, name resolution for client requests will not work as expected.

- GPO requirements—GPOs can be manually created before configuring Remote Access, or automatically created during Remote Access deployment. Requirements are as follows:
  - A unique client GPO is required for each domain.
  - A server GPO is required for each entry point, in the domain in which the entry point is located. So if multiple entry points are located in the same domain, there will be multiple server GPOs (one for each entry point) in the domain.
  - A unique Windows 7 client GPO is required for each entry point enabled for Windows 7 client support, for each domain.

Known issues

The following are known issues when configuring a multisite scenario:

- **Multiple entry points in the same IPv4 subnet**—Adding multiple entry points in the same IPv4 subnet will result in an IP address conflict message, and the DNS64 address for the entry point will not be configured as expected. This issue occurs when IPv6 has not been deployed on the internal interfaces of the servers on the corporate network. To prevent this problem run the following Windows PowerShell command on all current and future Remote Access servers:

```
Set-NetIPInterface –InterfaceAlias <InternalInterfaceName> –
AddressFamily IPv6 –DadTransmits 0
```

- If the public address specified for DirectAccess clients to connect to the Remote Access server has a suffix included in NRPT, DirectAccess might not work as expected. Ensure that the NRPT has an exemption for the public name. In a multisite deployment, exemptions should be added for the public names of all entry points. Note that if force tunneling is enabled these exemptions are added automatically. They are removed if force tunneling is disabled.
- When using the Windows PowerShell cmdlet **Disable-DAMultiSite**, the WhatIf and Confirm parameters have no effect, and Multisite will be disabled and Windows 7 GPOs will be removed.
- When Windows 7 clients using DCA in a Multisite deployment are upgraded to Windows 8, the Network Connectivity Assistant will not function. This issue can be resolved in advance of the client upgrade by modifying the Windows 7 GPOs using the following Windows PowerShell cmdlets:

```
Set-GPRegistryValue -Name <Windows7GpoName> -Domain <DomainName>
-Key
"HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\NetworkC
onnectivityAssistant" -ValueName "TemporaryValue" -Type Dword -
Value 1
Remove-GPRegistryValue -Name <Windows7GpoName> -Domain
<DomainName> -Key
"HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\NetworkC
onnectivityAssistant"
```

In the event that the client has already been upgraded, then move the client computer to the Windows 8 security group.

- When modifying domain controller settings using the Windows PowerShell cmdlet **Set-DAEntryPointDC**, if the ComputerName parameter specified is a Remote Access server in an entry point other than the last one added to the Multisite deployment, a warning will be displayed indicating that the server specified will not be updated until the next policy refresh. The actual server(s) that was not updated can be seen using the **Configuration Status** in the **DASHBOARD** of the **Remote Access Management Console**. This will not cause any functional problems, however, you can run **gpupdate /force** on the server(s) that was not updated to get the configuration status updated immediately.
- When multisite is deployed in an IPv4-only corporate network, changing the internal network IPv6 prefix also changes the DNS64 address, but does not update the address on firewall rules that allow DNS queries to the DNS64 service. To resolve this issue, run the following Windows PowerShell commands after changing the internal network IPv6 prefix:

```
$dns64Address = (Get-DAClientDnsConfiguration).NrptEntry | ?{
$_.DirectAccessDnsServers -match ':3333::1' } | Select-Object -
First 1 -ExpandProperty DirectAccessDnsServers

$serverGpoName = (Get-RemoteAccess).ServerGpoName

$serverGpoDc = (Get-DAEntryPointDC).DomainControllerName

$gpoSession = Open-NetGPO -PolicyStore $serverGpoName -
DomainController $serverGpoDc

Get-NetFirewallRule -GPOSession $gpoSession | ? {$_.Name -in
@('0FDEEC95-1EA6-4042-8BA6-6EF5336DE91A', '24FD98AA-178E-4B01-
9220-D0DADA9C8503')} |  Set-NetFirewallRule -LocalAddress
$dns64Address
```

```
Save-NetGPO -GPOSession $gpoSession
```

- If DirectAccess was deployed when an existing ISATAP infrastructure was present, when removing an entry point that was an ISATAP host, the IPv6 address of the DNS64 service will be removed from the DNS server addresses of all DNS suffixes in the NRPT.

  To resolve this issue, in the **Infrastructure Server Setup** wizard, on the **DNS** page, remove the DNS suffixes that were modified and add them again with the correct DNS server addresses, by clicking **Detect** on the **DNS Server Addresses** dialog box.

Plan a Multisite Deployment

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the planning steps required in order to deploy Windows Server® 2012 Remote Access in a multisite configuration.

1. Step 1: Plan an Advanced Single Server Deployment—This step includes planning for the infrastructure required to deploy a single server. It includes planning for network and server settings, certificate requirements, DNS settings, network location server deployment, DirectAccess management servers, Active Directory settings, and Group Policy objects (GPOs).
2. Step 2: Plan the Multisite Infrastructure—This step includes Active Directory and GPO planning, and DNS configuration.
3. Step 3: Plan the Multisite Deployment—This step includes planning certificate settings, network location server configuration, client entry point settings, IPv6 prefix settings, and optionally global load balancing settings.

Note
Record your planning decisions for Remote Access advanced deployment. This record can be used as a job aid for everyone involved in completing the deployment steps.

After you have completed these planning steps, see Configure a Multisite Deployment. For instructions on configuring a multisite deployment as a proof of concept in a lab environment, see Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment.

Step 1: Plan an Advanced Single Server Deployment

The first step in deploying Windows Server® 2012 Remote Access in a multisite topology is to plan and configure an advanced single server deployment, and then plan to add servers to each multisite entry point.

Plan a single server deployment

Before you deploy Remote Access in a multisite topology, make sure that you have completed all the planning steps to deploy a single Remote Access server. See Plan an Advanced Remote Access Deployment. Note the following:

- The single Remote Access server must be configured to use IPsec client certificate authentication. The option to use the Remote Access server as a Kerberos proxy for client authentication is not supported.
- The single server deployment should use a certification authority (CA) to issue the certificate for the network location server and IP-HTTPS websites.

If you have already deployed a single Remote Access without these settings, configure them before beginning the multisite deployment.

See also

- Step 2: Plan the Multisite Infrastructure

Step 2: Plan the Multisite Infrastructure

The next step in deploying Remote Access in a multisite topology is to complete the multisite infrastructure planning; including, Active Directory, security groups, and Group Policy Objects.

Task	Description
2.1 Plan Active Directory	Plan your Active Directory topology including Active Directory sites and Remote Access entry points.
2.2 Plan security groups	Plan to create additional security groups if required.
2.3 Plan Group Policy Objects	Plan additional Group Policy Objects (GPOs) for your deployment.
2.4 Plan DNS	Plan additional DNS entries for each entry point in your deployment.

2.1 Plan Active Directory

A Remote Access multisite deployment can be configured in a number of topologies:

- **Single Active Directory site, multiple entry points**—In this topology, you have a single Active Directory site for your entire organization with fast intranet links throughout the site, but you have multiple Remote Access servers deployed throughout your organization, each

acting as an entry point. A geographical example of this topology is to have a single Active Directory site for the United States with entry points on the East coast and the West coast.

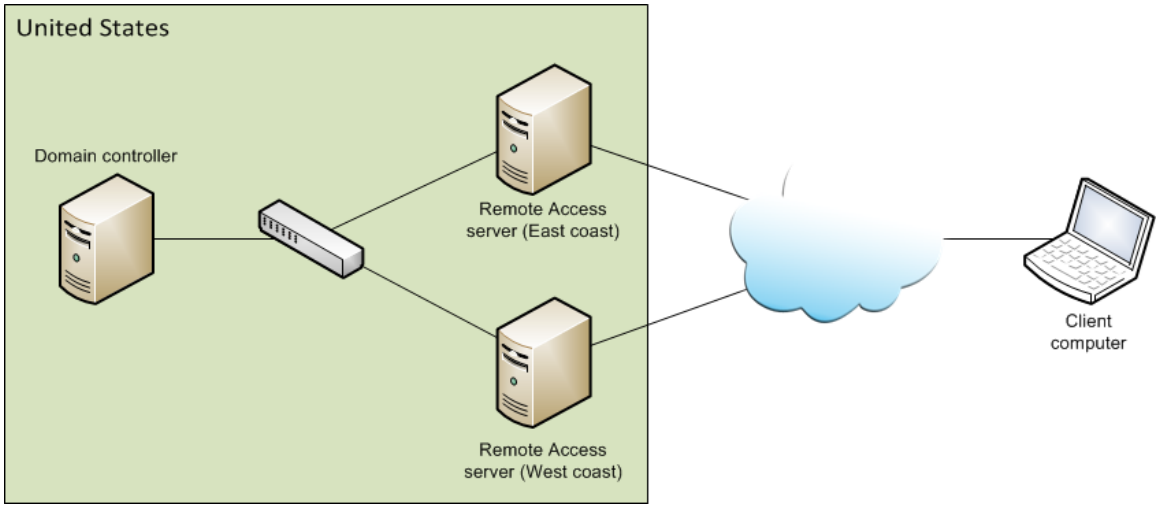

- **Multiple Active Directory sites, multiple entry points**—In this topology, you have two or more Active Directory sites with a Remote Access server deployed as an entry point for each site. Each Remote Access server is associated with the Active Directory domain controller for the site. A geographical example of this topology is to have an Active Directory site for the United States and one for Europe with a single entry point for each site. Note that if you have multiple Active Directory sites you do not need to have an entry point associated with each site. In addition, some Active Directory sites can have more than one entry point associated with it.

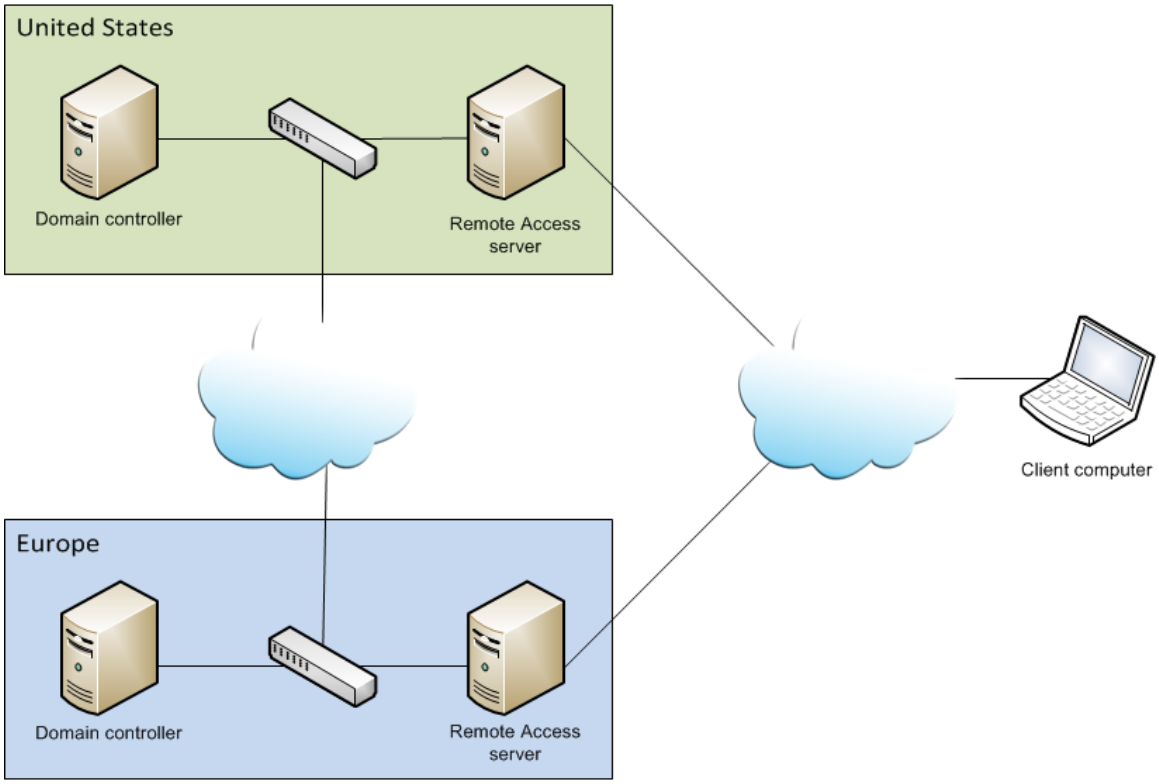

In a multisite entry point, you can configure a single Remote Access server, multiple Remote Access servers, or a Remote Access server clusters. For information on planning and deploying a cluster, see **Deploying Remote Access in a cluster with Windows NLB**.

Active Directory best practices and recommendations

Note the following recommendations and constraints for Active Directory deployment in a multisite scenario:

1. Each Active Directory site can contain one or more Remote Access servers, or a server cluster, functioning as multisite entry points for client computers. However an Active Directory site is not required to have an entry point.
2. A multisite entry point can only be associated with a single Active Directory site. When client computers running Windows® 8 connect to a specific entry point, they are considered as belonging to the Active Directory site associated with that entry point.
3. It is recommended that each Active Directory site has a domain controller. The domain controller can be read-only.
4. If each Active Directory site contains a domain controller, The GPO for a server in the entry point is managed by one of the domain controllers in the Active Directory site associated with the endpoint. If there are no write-enabled domain controllers in that site, then the GPO for a server is managed on a write-enabled domain controller that is found closest to the first Remote Access server that is configured in the entry point. Closest is determined by a link cost calculation. Note that in this scenario, after making configuration changes there might be

a delay when replicating between the domain controller managing the GPO and the read-only domain controller in the server's Active Directory site.

5. Client GPOs and optional application server GPOs are managed on the domain controller running as the Primary Domain Controller (PDC) emulator. This means that client GPOs are not necessarily managed in the Active Directory site containing the entry point to which clients connect.
6. If the domain controller for an Active Directory site is unreachable, the Remote Access server will connect to an alternate domain controller in the site (if available). If not, it connects to the domain controller for another site to retrieve an updated GPO and to authenticate clients. In both cases, the Remote Access Management console and PowerShell cmdlets cannot be used to retrieve or modify configuration settings until the domain controller is available. Note the following:
    a. If the server running as the PDC emulator is unavailable, you must designate an available domain controller that has updated GPOs as the PDC emulator.
    b. If the domain controller that manages a server GPO is unavailable, use the Set-DAEntryPointDC PowerShell cmdlet to associate a new domain controller with the entry point. The new domain controller should have up-to-date GPOs before running the cmdlet. For more information about unavailable domain controllers, see Domain controller maintenance and downtime.

2.2 Plan security groups

During deployment of a single server with advanced settings, all client computers accessing the internal network via DirectAccess were gathered into a security group. In a multisite deployment, this security group is used for Windows 8 client computers only. For a multisite deployment, Windows 7 client computers will be gathered into separate security groups for each entry point in the multisite deployment. For example, if you previously grouped all client computers in the group DA_Clients, you must now remove any Windows 7 computers from that group and put them in a different security group. For example, in the multiple Active Directory sites, multiple entry points topology, you create a security group for the United States entry point (DA_Clients_US) and one for the Europe entry point (DA_Clients_Europe). Place any Windows 7 client computers located in the United States in the DA_Clients_US group and any located in Europe in the DA_Clients_Europe group. If you do not have any Windows 7 client computers, you do not need to plan security groups for Windows 7 computers.

Required security groups are as follows:

- One security group for all Windows 8 client computers. It is recommended to create a unique security group for these clients for each domain.
- A unique security group containing Windows 7 client computers for each entry point. It is recommended to create a unique group for each domain. For example: Domain1\DA_Clients_Europe; Domain2\DA_Clients_Europe; Domain1\DA_Clients_US; Domain2\DA_Clients_US.

2.3 Plan Group Policy Objects

DirectAccess settings configured during Remote Access deployment are collected into GPOs. Your single server deployment already uses GPOs for DirectAccess clients, the Remote Access server, and optionally for application servers. A multisite deployment requires the following GPOs:

- A server GPO for each entry point.
- A GPO for each domain containing client computers running Windows 8.
- A GPO for each entry point and each domain containing Windows 7 client computers.

GPOs can be configured as follows:

- **Automatically**—You can specify that GPOs are created automatically by Remote Access. A default name is specified for each GPO, and can be modified.
- **Manually**—You can use GPOs that have been created manually by the Active Directory administrator.

Note

Once DirectAccess is configured to use specific GPOs, it cannot be configured to use different GPOs.

2.3.1 Automatically-created GPOs

Note the following when using automatically-created GPOs:

- Automatically created GPOs are applied according to the location and link target parameter, as follows:
  - For the server GPO, both the location and link parameters point to the domain containing the Remote Access server.
  - For client GPOs, the link target is set to the root of the domain in which the GPO was created. A GPO is created for each domain that contains client computers, and the GPO will be linked to the root of each domain. .
- For automatically-created GPOs, to apply DirectAccess settings the Remote Access server administrator requires the following permissions:
  - GPO create permissions for required domains.
  - Link permissions for all the selected client domain roots.
  - Permission to create the WMI filter for GPOs is required if DirectAccess was configured to work on mobile computers only.
  - Link permissions for the roots of domains associated with the entry points (the server GPO domains)
  - It is recommended that the Remote Access administrator has GPO read permissions for each required domain. This enables Remote Access to verify that GPOs with duplicate names do not exist when creating GPOs for the multisite deployment.
  - Active Directory replication permissions on domains associated with entry points. This is because when initially adding entry points, the server GPO for the entry point is created on the domain controller closest to that entry point. However, since link creation is supported on the PDC emulator only, the GPO must be replicated from the domain

controller on which it was created to the domain controller running as the PDC emulator before creating the link.

Note that if the correct permissions for replication and linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but replication and linking will not occur. If a link warning is issued links will not be created automatically, even after the permissions are in place. Instead the administrator will need to create the links manually.

2.3.2 Manually created GPOs

Note the following when using manually-created GPOs:

- The following GPOs should be created manually for the multisite deployment:
  - **Server GPO**—A server GPO for each entry point (in the domain in which the entry point is located). This GPO will be applied on each Remote Access server in the entry point.
  - **Client GPO (Windows 7)**—A GPO for each entry point and each domain containing Windows 7 client computers that will connect to entry points in the multisite deployment. For example Domain1\DA_W7_Clients_GPO_Europe; Domain2\DA_W7_Clients_GPO_Europe; Domain1\DA_W7_Clients_GPO_US; Domain2\DA_W7_Clients_GPO_US. If no Windows 7 client computers will connect to entry points, GPOs are not required.
- There is no requirement to create additional GPOs for Windows 8 client computers. A GPO for each domain containing client computers was already created when the single Remote Access server was deployed. In a multisite deployment these client GPOs will function as the GPOs for Windows 8 clients.
- The GPOs should exist before clicking the Commit button in the multisite deployment wizards.
- When using manually created GPOs a search is made for a link to the GPO in the entire domain. If the GPO is not linked in the domain then a link is automatically created in the domain root. If the required permissions to create the link are not available a warning is issued.
- When using manually-created GPOs, to apply DirectAccess settings the Remote Access server administrator requires full GPO permissions (Edit, Delete, Modify security) on the manually-created GPOs.

Note that if the correct permissions for replication and linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but replication and linking will not occur. If a link warning is issued links will not be created automatically, even after the permissions are in place. Instead the administrator will need to create the links manually.

2.3.3 Managing GPOs in a multi-domain controller environment

Each GPO will be managed by a specific domain controller, as follows:

- The server GPO is managed by one of the domain controllers in the Active Directory site associated with the server, or if domain controllers in that site are read-only, by a write-enabled domain controller closest to the first server in the entry point.
- Client GPOs will be managed by the domain controller running as the PDC emulator.

If you want to manually modify GPO settings note the following:

- For server GPOs, in order to identify which domain controller is associated with a specific entry point, run the PowerShell cmdlet `Get-DAEntryPointDC -EntryPointName <name of entry point>`.
- By default, when you make changes with networking PowerShell cmdlets or the Group Policy Management console, the domain controller acting as the PDC emulator is used.
- If you modify settings on a domain controller that is not the domain controller associated with the entry point (for server GPOs) or the PDC emulator (for client GPOs), note the following:
    a. Before modifying the settings, ensure that the domain controller is replicated with an up-to-date GPO, and back up GPO settings, before making changes. If the GPO is not updated, merge conflicts during replication might occur, resulting in a corrupt Remote Access configuration.
    b. After modifying the settings, you must wait for changes to replicate to the domain controller that is associated with the GPOs. Do not make additional changes using the Remote Access Management console or Remote Access PowerShell cmdlets until replication is complete. If a GPO is edited on two different domain controllers before replication is complete, merge conflicts might occur, resulting in a corrupt configuration
- Alternatively you can change the default setting using the **Change Domain Controller** dialog box in the Group Policy Management console, or using the **Open-NetGPO** PowerShell cmdlet, so that changes made using the console or the networking cmdlets use the domain controller you specify.
    a. To do this in the Group Policy Management console, right-click the domain or sites container and click **Change Domain Controller**.
    b. To do this in PowerShell, specify the DomainController parameter for the Open-NetGPO cmdlet. For example, to enable the private and public profiles in Windows Firewall on a GPO named domain1\DA_Server_GPO _Europe using a domain controller named europe-dc.corp.contoso.com, do the following:

```
$gpoSession = Open-NetGPO –PolicyStore “domain1\DA_Server_GPO
_Europe” –DomainController “europe-dc.corp.contoso.com”
Set-NetFirewallProfile –GpoSession $gpoSession –Name
@(“Private”,”Public”) –Enabled True
Save-NetGPO –GpoSession $gpoSession
```

Modifying domain controller association

To maintain the configuration consistency in a multisite deployment, it is important to make sure that each GPO is managed by a single domain controller. In some scenarios, it may be required to assign a different domain controller for a GPO:

- **Domain controller maintenance and downtime**—When the domain controller that manages a GPO is not available, it may be required to manage the GPO on a different domain controller. See Domain controller maintenance and downtime.

- **Optimization of configuration distribution**—After network infrastructure changes, it may be required to manage the server GPO of an entry point on a domain controller in the same Active Directory site as the entry point. See Optimization of configuration distribution.

2.4 Plan DNS

Note the following when planning DNS for a multisite deployment:

1. Client computers use the ConnectTo address in order to connect to the Remote Access server. Each entry point in your deployment requires a different ConnectTo address. Each entry point ConnectTo address must be available in the public DNS and the address that you choose must match the subject name of the IP-HTTPS certificate that you deploy for the IP-HTTPS connection. See Planning website certificates for IP-HTTPS.
2. In addition, Remote Access enforces symmetric routing; therefore, only Teredo and IP-HTTPS clients can connect when a multisite deployment is enabled. To allow native IPv6 clients to connect, the ConnectTo address (the IP-HTTPS URL) must resolve to both the external (Internet-facing) IPv6 and IPv4 addresses on the Remote Access server.
3. Remote Access creates a default web probe that is used by DirectAccess client computers to verify connectivity to the internal network. During configuration of the single server the following names should have been registered in DNS:
   a. directaccess-WebProbeHost—should resolve to the internal IPv4 address of the Remote Access server, or to the IPv6 address in an IPv6-only environment.
   b. directaccess-CorpConnectivityHost—should resolve to a localhost (loopback) address (either ::1 or 127.0.0.1, depending on whether IPv6 is deployed in the corporate network).

   In a multisite deployment an additional DNS entry for directaccess-WebProbeHost must be created for each entry point. When adding an entry point, Remote Access tries to automatically create this additional directaccess-WebProbeHost entry. However, if it fails a warning will be shown and you must create the entry manually.

   Note
   When DNS scavenging is enabled in your DNS infrastructure, it is recommended to disable scavenging on the DNS entries created automatically by Remote Access.

Note
DNS requirements for the network location server are discussed in Step 3: Plan the Multisite Deployment.

See also

- Step 3: Plan the Multisite Deployment
- Step 1: Plan an Advanced Single Server Deployment
- Deploy Multiple Remote Access Servers in a Multisite Deployment
- **Deploying Remote Access in an enterprise**

Step 3: Plan the Multisite Deployment

After planning the multisite infrastructure, plan any additional certificate requirements, how client computers select entry points, and IPv6 addresses assigned in your deployment.

Task	Description
3.1 Plan IP-HTTPS certificates	Plan the certificate to use for the IP-HTTPS website for each entry point.
3.2 Plan the network location server	Decide where to host the network location server website and the certificate to use.
3.3 Plan the IPsec root certificate for all Remote Access servers	Plan to configure all Remote Access servers with the same IPsec root certificate.
3.4 Plan global server load balancing	Decide whether to use global server load balancing.
3.5 Plan DirectAccess client entry point selection	Plan how DirectAccess client computers will select the entry point to connect to.
3.6 Plan prefixes and routing	Plan the IPv6 prefixes for your deployment.
3.7 Plan the transition to IPv6 when multisite Remote Access is deployed	Plan for a transition to an IPv6-only multisite deployment.

3.1 Plan IP-HTTPS certificates

When you configure the entry points, you configure each entry point with a specific ConnectTo address. The IP-HTTPS certificate for each entry point must match the ConnectTo address. Note the following when obtaining the certificate:

- You cannot use self-signed certificates in a multisite deployment.
- Using a public CA is recommended, so that CRLs are readily available.
- In the subject field, specify either the IPv4 address of the external adapter of Remote Access server (if the ConnectTo address has been specified as an IP address and not a DNS name), or the FQDN of the IP-HTTPS URL.
- The common name of the certificate should match the name of the IP-HTTPS website. Use of a wildcard URL that matches the ConnectTo DNS name is also supported.
- IP-HTTPS certificates can use wildcards in the subject name. The same wildcard certificate can be used for all entry points.
- For the Enhanced Key Usage field, use the Server Authentication object identifier (OID).
- If you are supporting client computers running Windows 7 in the multisite deployment, in the CRL Distribution Points field, specify a CRL distribution point that is accessible by

DirectAccess clients that are connected to the Internet. This is not required for clients running Windows 8 (by default the CRL revocation check is disabled for IP-HTTPS on these clients).

- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store of the computer, and not of the user.

3.2 Plan the network location server

The network location server website can be hosted on the Remote Access server or on another server in your organization. If you host the network location server on the Remote Access server, the website is created automatically when you deploy Remote Access. If you host the network location server on another server running a Windows operating system in your organization, you must make sure that Internet Information Services (IIS) is installed to create the website.

3.2.1 Certificate requirements for the network location server

Make sure that the network location server website meets the following requirements for certificate deployment:

- It requires an HTTPS server certificate.
- If the network location server is located on the Remote Access server, and you selected to use a self-signed certificate when you deployed the single Remote Access server, you must reconfigure the single server deployment to use a certificate issued by an internal CA.
- DirectAccess client computers must trust the CA that issued the server certificate to the network location server website.
- DirectAccess client computers on the internal network must be able to resolve the name of the network location server website.
- The network location server website must be highly available to computers on the internal network.
- The network location server must not be accessible to DirectAccess client computers on the Internet.
- The server certificate must be checked against a Certificate Revocation List (CRL).
- Wildcard certificates are not supported when the network location server is hosted on the Remote Access server.

When obtaining the website certificate to use for the network location server, note the following:

1. In the Subject field, specify either an IP address of the intranet interface of the network location server or the FQDN of the network location URL. Note that you should not specify an IP address if the network location server is hosted on the Remote Access server. This is because the network location server must use the same subject name for all entry points, and not all entry points have the same IP address.
2. For the Enhanced Key Usage field, use the Server Authentication OID.
3. For the CRL Distribution Points field, use a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet.

3.2.2DNS for the network location server

If you host the network location server on the Remote Access server, you must add a DNS entry for the network location server website for every entry point in your deployment. Note the following:

- The subject name of the first network location server certificate in the multisite deployment is used as the network location server URL for all entry points, therefore the subject name and the network location server URL cannot be the same as the computer name of the first Remote Access server in the deployment. It must be an FQDN dedicated for the network location server.
- The service provided by the network location server traffic is balanced across entry points using DNS, and therefore there should be a DNS entry with the same URL for each entry point, configured with the internal IP address of the entry point.
- All entry points must be configured with a network location server certificate with the same subject name (that matches the network location server URL).
- The network location server infrastructure (DNS and certificate settings) for an entry point must be created before adding the entry point.

3.3 Plan the IPsec root certificate for all Remote Access servers

Note the following when planning for IPsec client authentication in a multisite deployment:

1. If you opted to use the built-in Kerberos proxy for computer authentication when you set up the single Remote Access server, you must change the setting to use computer certificates issued by an internal CA, since Kerberos proxy is not supported for a multisite deployment.
2. If you used a self-signed certificate, you must reconfigure the single server deployment to use a certificate issued by an internal CA.
3. For IPsec authentication to succeed during client authentication, all Remote Access servers must have a certificate issued by the IPsec root or intermediate CA, and with the Client Authentication OID for Enhanced Key Usage.
4. The same IPsec root or intermediate certificate must be installed on all of the Remote Access servers in the multisite deployment.

3.4 Plan global server load balancing

In a multisite deployment, you can additionally configure a global server load balancer. A global server load balancer can be useful to your organization if your deployment covers a large geographical distribution because it can distribute traffic load between the entry points. The global server load balancer can be configured to provide DirectAccess clients with the entry point information of the closest entry point. The process works as follows:

1. Client computers running Windows 8 have a list of global server load balancer IP addresses, each associated with an entry point.

2. The Windows 8 client computer attempts to resolve the FQDN of the global server load balancer in the public DNS to an IP address. If the resolved IP address is listed as the global server load balancer IP address of an entry point, the client computer automatically selects that entry point and connects to its IP-HTTPS URL (ConnectTo address), or to its Teredo server IP address. Note that the IP address of the global server load balancer does not need to be identical to the ConnectTo address or to the Teredo server address of the entry point, since the client computers never try to connect to the global server load balancer IP address.
3. If the client computer is behind a web proxy (and cannot use DNS resolution), or if the global server load balancer FQDN does not resolve to any configured global server load balancer IP address, then an entry point will be selected automatically using an HTTPS probe to the IP-HTTPS URLs of all entry points. The client will connect to server that responds first.

For a list of global server load balancing devices that support Remote Access, go to the Find a Partner page at Microsoft Server and Cloud Platform.

3.5 Plan DirectAccess client entry point selection

When you configure a multisite deployment, by default, Windows 8 client computers are configured with the information required to connect to all entry points in the deployment and to automatically connect to a single entry point based on a selection algorithm. You can also configure your deployment to allow Windows 8 client computers to manually select the entry point to which they will connect. If a Windows 8 client computer is currently connected to the United States entry point and automatic entry point selection is enabled, if the United States entry point becomes unreachable, after a few minutes the client computer will attempt to connect through the Europe entry point. Using automatic entry point selection is recommended; however, allowing manual entry point selection enables end users to connect to a different entry point based on current network conditions. For example, if a computer is connected to the United States entry point and the connection to the internal network becomes much slower than expected. In this situation, the end user can manually select to connect to the Europe entry point to improve the connection to the internal network.

Note

After an end user selects an entry point manually, the client computer will not revert to automatic entry point selection. That is, if the manually selected entry point becomes unreachable, the end user must either revert to automatic entry point selection, or manually select another entry point.

Windows 7 client computers are configured with the information required to connect to a single entry point in the multisite deployment. They cannot store the information for multiple entry points simultaneously. For example, a Windows 7 client computer can be configured to connect to the United States entry point, but not the Europe entry point. If the United States entry point is unreachable, the Windows 7 client computer will lose connectivity to the internal network until the entry point is reachable. The end user cannot make any changes to attempt to connect to the Europe entry point.

3.6 Plan prefixes and routing

Internal IPv6 prefix

During deployment of the single Remote Access server you planned the internal network IPv6 prefixes Note the following in a multisite deployment:

1. If you included all of your Active Directory sites when you configured your single server Remote Access deployment, the internal network IPv6 prefixes will already be defined in the Remote Access Management console.
2. If you create additional Active Directory sites for multisite deployment, you must plan new IPv6 prefixes for the additional sites and define them in Remote Access. Note that IPv6 prefixes can only be configured using the Remote Access Management console or PowerShell cmdlets if IPv6 is deployed in the internal corporate network.

IPv6 prefix for DirectAccess client computers (IP-HTTPS prefix)

1. If IPv6 is deployed in the internal corporate network, you must plan an IPv6 prefix to assign to DirectAccess client computers in any additional entry points in your deployment.
2. Ensure that the IPv6 prefixes to assign to DirectAccess client computers in each entry point are distinct and that there is no overlap in the IPv6 prefixes.
3. If IPv6 is not deployed in the corporate network, an IP-HTTPS prefix for each entry point will be automatically selected when adding the entry point.

IPv6 prefix for VPN clients

If you deployed VPN on the single Remote Access server, note the following:

1. Adding an IPv6 VPN prefix to an entry point is only required if you want to allow VPN client IPv6 connectivity to the corporate network.
2. The VPN prefix can only be configured on an entry point using the Remote Access Management console or PowerShell cmdlet if IPv6 is deployed in the internal corporate network, and if VPN is enabled on the entry point.
3. The VPN prefix should be unique in each entry point and should not overlap with other VPN or IP-HTTPS prefixes.
4. If IPv6 is not deployed in the corporate network, VPN clients connecting to the entry point will not be assigned an IPv6 address.

Routing

In a multisite deployment symmetric routing is enforced using Teredo and IP-HTTPS. When IPv6 is deployed in the corporate network note the following:

1. The Teredo and IP-HTTPS prefixes of each entry point must be routable across the corporate network to their associated Remote Access server.
2. The routes must be configured in the corporate network routing infrastructure.
3. For each entry point there should be one to three routes in the internal network:

a. IP-HTTPS prefix—This prefix is chosen by the administrator in the Add an Entry Point wizard.
b. VPN IPv6 prefix (optional)—This prefix can be chosen after enabling VPN for an entry point
c. Teredo prefix (optional)—This prefix is relevant only if the Remote Access server is configured with two consecutive public IPv4 addresses on the external adapter. The prefix is based on the first public IPv4 address of the address pair. For example if the external addresses are:
    i. www.xxx.yyy.zzz
    ii. www.xxx.yyy.zzz+1

    Then the Teredo prefix to configure is 2001:0:WWXX:YYZZ::/64, where WWXX:YYZZ is the hexadecimal representation of the IPv4 address www.xxx.yyy.zzz.

    Note that you can use the following script to calculate the Teredo prefix:

```
$TeredoIPv4 = (Get-NetTeredoConfiguration).ServerName # Use for a
Remote Access server that is already configured
$TeredoIPv4 = "20.0.0.1" # Use for an IPv4 address

    [Byte[]] $TeredoServerAddressBytes = `

[System.Net.IPAddress]::Parse("2001::").GetAddressBytes()[0..3] +
`
    [System.Net.IPAddress]::Parse($TeredoIPv4).GetAddressBytes()
+ `
    [System.Net.IPAddress]::Parse("::").GetAddressBytes()[0..7]

Write-Host "The server's Teredo prefix is
$([System.Net.IPAddress]$TeredoServerAddressBytes)/64"
```

d. All of the above routes must be routed to the IPv6 address on the internal adapter of the Remote Access server (or to the internal virtual IP (VIP) address for a load balanced entry point).

Note

When IPv6 is deployed in the corporate network and Remote Access server administration is performed remotely over DirectAccess, routes for the Teredo and IP-HTTPS prefixes of all other entry points must be added to each Remote Access server so that the traffic will be forwarded to the internal network.

Active Directory site-specific IPv6 prefixes

When a client computer running Windows 8 is connected to an entry point, the client computer is immediately associated with the Active Directory site of the entry point, and is configured with IPv6 prefixes associated with the entry point. The preference is for client computers to connect to resources using these IPv6 prefixes, since they are configured dynamically in the IPv6 prefix policy table with higher precedence when connecting to an entry point.

If your organization uses an Active Directory topology with site-specific IPv6 prefixes (for example an internal resource FQDN app.corp.com is hosted in both North America and Europe with a site-specific IP address in each location) this is not configured by default using the Remote Access console, and site-specific IPv6 prefixes are not configured for each entry point. If you do want to enable this optional scenario, you need to configure each entry point with the specific IPv6 prefixes that should be preferred by client computers connecting to a specific entry point. Do this as follows:

1. For each GPO used for Windows 8 client computers, run the Set-DAEntryPointTableItem PowerShell cmdlet
2. Set the EntryPointRange parameter for the cmdlet with the site-specific IPv6 prefixes. For example, to add the site-specific prefixes 2001:db8:1:1::/64 and 2001:db:1:2::/64 to an entry point called Europe, run the following

```
$entryPointName = "Europe"
$prefixesToAdd = @("2001:db8:1:1::/64", "2001:db8:1:2::/64")
$clientGpos = (Get-DAClient).GpoName
$clientGpos | % { Get-DAEntryPointTableItem –EntryPointName
$entryPointName -PolicyStore $_ | %{ Set-DAEntryPointTableItem -
PolicyStore $_.PolicyStore -EntryPointName $_.EntryPointName -
EntryPointRange ($_.EntryPointRange) + $prefixesToAdd}}
```

3. When modifying the EntryPointRange parameter, ensure that you do not remove the existing 128-bit prefixes which belong to the IPsec tunnel endpoints and the DNS64 address.

3.7 Plan the transition to IPv6 when multisite Remote Access is deployed

Many organizations use the IPv4 protocol on the corporate network. With the exhaustion of available IPv4 prefixes, many organizations are making the transition from IPv4-only to IPv6-only networks.

This transition is most likely to take place in two stages:

1. From an IPv4-only to an IPv6+IPv4 corporate network.
2. From an IPv6+IPv4 to an IPv6-only corporate network.

In each part, the transition may be performed in stages. In each stage only one subnet of the network may be changed to the new network configuration. Therefore, a DirectAccess multisite deployment is required to support a hybrid deployment where, for example, some of the entry

points belong to an IPv4-only subnet and others belong to an IPv6+IPv4 subnet. In addition, configuration changes during the transition processes must not break client connectivity via DirectAccess.

Transition from an IPv4-only to an IPv6+IPv4 corporate network

When adding IPv6 addresses to an IPv4-only corporate network, you may want to add an IPv6 address to an already deployed DirectAccess server. In addition, you may want to add an entry point or a node to a load balanced cluster with both IPv4 and IPv6 addresses to the DirectAccess deployment.

Remote Access allows you to add servers with both IPv4 and IPv6 addresses to a deployment that was originally configured with only IPv4 addresses. These servers are added as IPv4-only servers and their IPv6 addresses are ignored by DirectAccess; consequently, your organization cannot take advantage of the benefits of native IPv6 connectivity on these new servers.

To transform the deployment to an IPv6+IPv4 deployment and take advantage of the native IPv6 capabilities, you must reinstall DirectAccess. To maintain client connectivity throughout the reinstallation see Transition from an IPv4-only to an IPv6-only deployment using dual DirectAccess deployments.

Note

As with an IPv4-only network, in a mixed IPv4+IPv6 network, the address of the DNS server that is used to resolve client DNS requests must be configured with the DNS64 that is deployed on Remote Access servers themselves, and not with a corporate DNS.

Transition from an IPv6+IPv4 to an IPv6-only corporate network

DirectAccess enables you to add IPv6-only entry points only if the first Remote Access server in the deployment originally had either both IPv4 and IPv6 addresses or only an IPv6 address. That is, you cannot transition from an IPv4-only network to an IPv6-only network in a single step without reinstalling DirectAccess. To transition directly from an IPv4-only network to an IPv6-only network, see Transition from an IPv4-only to an IPv6-only deployment using dual DirectAccess deployments.

After you have completed the transition from an IPv4-only deployment to an IPv6+IPv4 deployment, you can transition to an IPv6-only network. During, and after the transition, note the following:

- If any IPv4-only backend servers remain in the corporate network, they will not be reachable to clients that connect through the IPv6-only entry points.
- When adding IPv6-only entry points to an IPv4+IPv6 deployment, DNS64 and NAT64 will not be enabled on the new servers. Clients that connect to these entry points will be automatically configured to use the corporate DNS servers.
- If you need to delete IPv4 addresses from a deployed server, you must remove the server from the DirectAccess deployment, remove its IPv4 corporate network address, and add it again to the deployment.

To support client connectivity to the corporate network, you must ensure that the network location server can be resolved by the corporate DNS to its IPv6 address. An additional IPv4 address can be set as well, but isn't required.

Transition from an IPv4-only to an IPv6-only deployment using dual DirectAccess deployments

The transition from an IPv4-only to an IPv6-only corporate network cannot be done without reinstalling the DirectAccess deployment. To maintain client connectivity during the transition, you can use another DirectAccess deployment. Dual deployment is required when the first transition stage finishes (IPv4-only network upgraded to IPv4+IPv6) and you intend to prepare for a future transition to an IPv6-only corporate network or to take advantage of the native IPv6 connectivity benefits. The dual deployment is described in the following general steps:

1. Install a second DirectAccess deployment. You can install DirectAccess on new servers, or remove servers from the first deployment and use them for the second deployment.

   Notes

   When installing an additional DirectAccess deployment alongside a current one, make sure that no two entry points share the same client prefix.

   If you install DirectAccess using the Getting Started Wizard or with the cmdlet `Install-RemoteAccess`, Remote Access automatically sets the client prefix of the first entry point in the deployment to a default value of <IPv6 subnet_prefix>:1000::/64. If needed, you must change the prefix.
2. Remove the chosen client security groups from the first deployment.
3. Add the client security groups to the second deployment.

   Important

   To maintain client connectivity throughout the process, you must add the security groups to the second deployment immediately after removing them from the first. This ensures that clients will not be updated with either two or zero DirectAccess GPOs. Clients will start using the second deployment once they retrieve and update their client GPO.
4. Optional: Remove the DirectAccess entry points from the first deployment and add those servers as new entry points in the second.

When you have completed the transition, you can uninstall the first DirectAccess deployment. When uninstalling, the following issues may occur:

- If the deployment was configured to support only clients on mobile computers, the WMI filter will be deleted. If the client security groups of the second deployment include desktop computers, the DirectAccess client GPO will not filter desktop computers and may cause issues on them. If a mobile computers filter is needed, recreate it by following the instructions on Create WMI Filters for the GPO.
- If both deployments were originally created on the same Active Directory domain, the DNS probe entry which points to localhost will be deleted and may cause client connectivity issues.

For example, clients may connect using IP-HTTPS rather than Teredo, or switch between DirectAccess multisite entry points. In this case, you must add the following DNS entry to the corporate DNS:

- Zone: domain name
- Name: directaccess-corpConnectivityHost
- IP address: ::1
- Type: AAAA

See also

- Configure a Multisite Deployment

Configure a Multisite Deployment

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the configuration steps required in order to deploy a single Windows Server 2012 Remote Access multisite deployment.

- Step 1: Implement a Single Server Remote Access Deployment—Install and configure a single Remote Access server. The multisite deployment requires you to install a single server before configuring a multisite deployment. For instructions, see Deploy a Single Remote Access Server with Advanced Settings.
- Step 2: Configure the Multisite Infrastructure—For a multisite deployment you must configure additional Active Directory sites and domain controllers. Additional security groups and Group Policy Objects (GPOs) are also required if you are not using automatically configured GPOs.
- Step 3: Configure the Multisite Deployment—Install the Remote Access role on additional Remote Access servers, enable the multisite deployment, and configure the additional servers as entry points for the deployment.
- Step 4: Verify the Multisite Deployment.

Step 1: Implement a Single Server Remote Access Deployment

The first configuration step to deploy Remote Access in a multisite topology is to implement an advanced single server deployment and then plan to add servers to each multisite entry point.

Implement a single server deployment

Before you can configure a multisite deployment, you must configure an advanced single server Remote Access deployment as described in Deploy a Single Remote Access Server with Advanced Settings.

See also

- Step 2: Configure the Multisite Infrastructure
- Configure a Multisite Deployment
- Deploy Multiple Remote Access Servers in a Multisite Deployment

Step 2: Configure the Multisite Infrastructure

To configure a multisite deployment, there are a number of steps required to modify network infrastructure settings including: configuring additional Active Directory sites and domain controllers, configuring additional security groups, and configuring Group Policy Objects (GPOs) if you are not using automatically configured GPOs.

Task	Description
2.1. Configure additional Active Directory sites	Configure additional Active Directory sites for the deployment.
2.2. Configure additional domain controllers	Configure additional Active Directory domain controllers as required.
2.3. Configure security groups	Configure security groups for any Windows 7 client computers.
2.4. Configure GPOs	Configure additional Group Policy Objects as required.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

2.1. Configure additional Active Directory sites

As discussed in 2.1 Plan Active Directory, all entry points can reside in a single Active Directory site. Therefore, at least one Active Directory site is required for the implementation of Remote Access servers in a multisite configuration. Use this procedure if you need to create the first Active Directory site, or if you desire to use additional Active Directory sites for the multisite

deployment. Use the Active Directory Sites and Services snap-in to create new sites in your organization's network.

Membership in the **Enterprise Admins** group in the forest or the **Domain Admins** group in the forest root domain, or equivalent, at a minimum is required to complete this procedure. Review details about using the appropriate accounts and group memberships at Local and Domain Default Groups.

For more information, see Adding a Site to the Forest.

Do this step using Windows PowerShell

To configure additional Active Directory sites

1. On the primary domain controller, click **Start**, and then click **Active Directory Sites and Services**.
2. In the Active Directory Sites and Services console, in the console tree, right-click **Sites**, and then click **New Site**.
3. On the **New Object - Site** dialog box, in the **Name** box, enter a name for the new site.
4. In **Link Name**, click a site link object, and then click **OK** twice.
5. In the console tree, expand **Sites**, right-click **Subnets**, and then click **New Subnet**.
6. On the **New Object - Subnet** dialog box, under **Prefix**, type the IPv4 or IPv6 subnet prefix, in the **Select a site object for this prefix** list, click the site to associate with this subnet, and then click **OK**.
7. Repeat steps 5 and 6 until you have created all the subnets required in your deployment.
8. Close Active Directory Sites and Services.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To install the Windows Feature 'Active Directory module for Windows PowerShell':

```
Install-WindowsFeature –Name RSAT-AD-PowerShell
```

or add the 'Active Directory PowerShell Snap-In' via OptionalFeatures.

If running the following cmdlets on Windows® 7 or Windows Server 2008 R2, then the Active Directory PowerShell module must be imported:

```
Import-Module ActiveDirectory
```

To configure an Active Directory site named 'Second-Site' using the built-in DEFAULTIPSITELINK:

```
New-ADReplicationSite -Name "Second-Site"
```

```
Set-ADReplicationSiteLink -Identity "DEFAULTIPSITELINK" -sitesIncluded @{Add="Second-Site"}
```

To configure IPv4 and IPv6 subnets for the Second-Site:

```
New-ADReplicationSubnet -Name "10.2.0.0/24" -Site "Second-Site"
New-ADReplicationSubnet -Name "2001:db8:2::/64" -Site "Second-Site"
```

2.2. Configure additional domain controllers

To configure a multisite deployment in a single domain, it is recommended that you have at least one writeable domain controller for each site in your deployment.

To perform this procedure, at a minimum you must be a member of the Domain Admins group in the domain in which the domain controller is being installed.

For more information, see Installing an Additional Domain Controller.

To configure additional domain controllers

1. On the server that will act as a domain controller, in **Server Manager**, on the **Dashboard**, click **add roles and features**.
2. Click **Next** three times to get to the server role selection screen
3. On the **Select Server Roles** page, select **Active Directory Domain Services**. Click **Add Features** when prompted, and then click **Next** three times.
4. On the **Confirmation** page, click **Install**.
5. When the installation completes successfully, click **Promote this server to a domain controller**.
6. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, click **Add a domain controller to an existing domain**.
7. In **Domain**, enter the domain name; for example, corp.contoso.com.
8. Under **Supply the credentials to perform this operation**, click **Change**. On the **Windows Security** dialog box, provide the user name and password for an account that can install the additional domain controller. To install an additional domain controller, you must be a member of the Enterprise Admins group or the Domain Admins group. When you are finished providing credentials, click **Next**.
9. On the **Domain Controller Options** page, do the following:
   a. Make the following selections:
      - **Domain Name System (DNS) server**—This option is selected by default so that your domain controller can function as a Domain Name System (DNS) server. If you do not want the domain controller to be a DNS server, clear this option.

   If the DNS server role is not installed on the Primary Domain Controller (PDC) emulator in the forest root domain, then the option to install DNS server on an additional domain controller is not available. As a workaround in this situation, you

can install the DNS server role before or after the AD DS installation.

> **Note**
> If you select the option to install DNS server, you might receive a message that indicates that a DNS delegation for the DNS server could not be created and that you should manually create a DNS delegation to the DNS server to ensure reliable name resolution. If you are installing an additional domain controller in either the forest root domain or a tree root domain, you do not have to create the DNS delegation. In this case, click **Yes** and disregard the message.

- **Global Catalog (GC)**—This option is selected by default. It adds the global catalog, read-only directory partitions to the domain controller, and it enables global catalog search functionality.
- **Read-only domain controller (RODC)**—This option is not selected by default. It makes the additional domain controller read only; that is, it makes the domain controller an RODC.

b. In **Site name**, select a site from the list.

c. Under **Type the Directory Services Restore Mode (DSRM) password**, in **Password** and **Confirm password**, type a strong password twice, and then click **Next**. This password must be used to start AD DS in DSRM for tasks that must be performed offline.

10. On the **DNS Options** page, select the **Update DNS delegation** check box if you want to update DNS delegation during role installation, and then click **Next**.
11. On the **Additional Options** page, type or browse to the volume and folder locations for the database file, the directory service log files, and the system volume (SYSVOL) files. Specify replication options as required, and then click **Next**.
12. On the **Review Options** page, review the installation options, and then click **Next**.
13. On the **Prerequisites Check** page, after the prerequisites are validated, click **Install**.
14. Wait until the wizard completes the configuration, and then click **Close**.
15. Restart the computer if it did not restart automatically.

2.3. Configure security groups

A multisite deployment requires an additional security group for Windows 7 client computers for every entry point in the deployment that allows access to Windows 7 client computers. If there are multiple domains containing Windows 7 client computers, then it is recommended to create a security group in each domain for the same entry point. Alternatively, one universal security group containing the client computers from both domains can be used. For example, in an environment with two domains, if you want to allow access to Windows 7 client computers in entry points 1 and 3, but not in entry point 2, then create two new security groups to contain the Windows 7 client computers for each entry point in each of the domains.

Do this step using Windows PowerShell

To configure additional security groups

1. On the primary domain controller, click **Start**, and then click **Active Directory Users and Computers**.
2. In the console tree, right-click the folder in which you want to add a new group, for example, corp.contoso.com/Users. Point to **New**, and then click **Group**.
3. On the **New Object - Group** dialog box, under **Group name**, type the name of the new group, for example, Win7_Clients_Entrypoint1.
4. Under **Group scope**, click **Universal**, under **Group type**, click **Security**, and then click **OK**.
5. To add computers to the new security group, double-click the security group, and on the **<Group_Name> Properties** dialog box, click the **Members** tab.
6. On the **Members** tab, click **Add**.
7. Select the Windows 7 computers to add to this security group, and then click **OK**.
8. Repeat this procedure to create a security group for every entry point as required.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To install the Windows Feature 'Active Directory module for Windows PowerShell':

```
Install-WindowsFeature –Name RSAT-AD-PowerShell
```

or add the 'Active Directory PowerShell Snap-In' via OptionalFeatures.

If running the following cmdlets on Windows® 7 or Windows Server 2008 R2, then the Active Directory PowerShell module must be imported:

```
Import-Module ActiveDirectory
```

To configure a security group named Win7_Clients_Entrypoint1 and to add a client computer named CLIENT2:

```
New-ADGroup -GroupScope universal -Name Win7_Clients_Entrypoint1

Add-ADGroupMember -Identity Win7_Clients_Entrypoint1 -Members CLIENT2$
```

2.4. Configure GPOs

A multisite Remote Access deployment requires the following Group Policy Objects:

- A GPO for every entry point for the Remote Access server.
- A GPO for any Windows 8 client computers for each domain.
- A GPO in each domain that contains Windows 7 client computers for each entry point configured to support Windows 7 clients.

Note

If you do not have any Windows 7 client computers, you do not need to create GPOs for Windows 7 computers.

When you configure Remote Access, the wizard automatically creates the required Group Policy Objects if they don't already exist. If you do not have the required permissions to create Group Policy Objects, they must be created prior to configuring Remote Access. The DirectAccess administrator must have full permissions on the GPOs (Edit + modify security + delete).

Important

After manually creating the GPOs for Remote Access you must allow sufficient time for Active Directory and DFS replication to the domain controller in the Active Directory site that is associated with the Remote Access server. If Remote Access automatically created the Group Policy Objects, then no wait time is required.

To create Group Policy Objects, see Create and Edit a Group Policy Object.

Domain controller maintenance and downtime

When a domain controller running as the PDC emulator, or domain controllers managing server GPOs experience downtime, it is not possible to load or modify the Remote Access configuration. This does not affect client connectivity if other domain controllers are available.

To load or modify the Remote Access configuration, you can transfer the PDC emulator role to a different domain controller for the client or application server GPOs; for server GPOs, change the domain controllers that manage the server GPOs.

Important

This operation can be performed only by a domain administrator. The impact of changing the primary domain controller is not confined to Remote Access; therefore, use caution when transferring the PDC emulator role.

Note

Before modifying domain controller association, make sure that all of the GPOs in the Remote Access deployment have been replicated to all of the domain controllers in the domain. If the GPO is not synchronized, recent configuration changes may be lost after modifying domain controller association, which may lead to a corrupt configuration. To verify GPO synchronization, see Check Group Policy Infrastructure Status.

To transfer the PDC emulator role

1. On the **Start** screen, type **dsa.msc**, and then press ENTER.
2. In the left pane of the Active Directory Users and Computers console, right-click **Active Directory Users and Computers**, and then click **Change Domain Controller**. On the Change Directory Server dialog box, click **This Domain Controller or AD LDS instance**, in the list click the domain controller that will be the new role holder, and then click **OK**.

Note

You must perform this step if you are not on the domain controller to which you want to transfer the role. Do not perform this step if you are already connected to the domain controller to which you want to transfer the role.

3. In the console tree, right-click **Active Directory Users and Computers**, point to **All Tasks**, and then click **Operations Masters**.
4. On the Operations Masters dialog box, click the **PDC** tab, and then click **Change**.
5. Click **Yes** to confirm that you want to transfer the role, and then click **Close**.

To change the domain controller that manages server GPOs

- Run the Windows PowerShell cmdlet `Set-DAEntryPointDC` on the Remote Access server and specify the unreachable domain controller name for the *ExistingDC* parameter. This command modifies the domain controller association for the server GPOs of the entry points that are currently managed by that domain controller.
  - To replace the unreachable domain controller 'dc1.corp.contoso.com' with the domain controller 'dc2.corp.contoso.com', do the following:

    ```
    Set-DAEntryPointDC –ExistingDC 'dc1.corp.contoso.com' –NewDC
    'dc2.corp.contoso.com' –ErrorAction Inquire
    ```

  - To replace the unreachable domain controller 'dc1.corp.contoso.com' with a domain controller in the closest Active Directory site to the Remote Access server 'DA1.corp.contoso.com', do the following:

    ```
    Set-DAEntryPointDC –ExistingDC 'dc1.corp.contoso.com' –
    ComputerName 'DA1.corp.contoso.com' –ErrorAction Inquire
    ```

Change two or more domain controllers that manage server GPOs

In a minimal number of cases, two or more domain controllers that manage server GPOs are unavailable. If this occurs, more steps are required to change the domain controller association for the server GPOs.

Domain controller association information is stored both in the registry of the Remote Access servers and in all server GPOs. In the following example, there are two entry points with two Remote Access servers, 'DA1' in 'Entry point 1' and 'DA2' in 'Entry point 2'. The server GPO of 'Entry point 1' is managed in the domain controller 'DC1', while the server GPO of 'Entry point 2' is managed in the domain controller 'DC2'. Both 'DC1' and 'DC2' are unavailable. A third domain controller is still available in the domain, 'DC3', and the data from 'DC1' and 'DC2' was already replicated to 'DC3'.

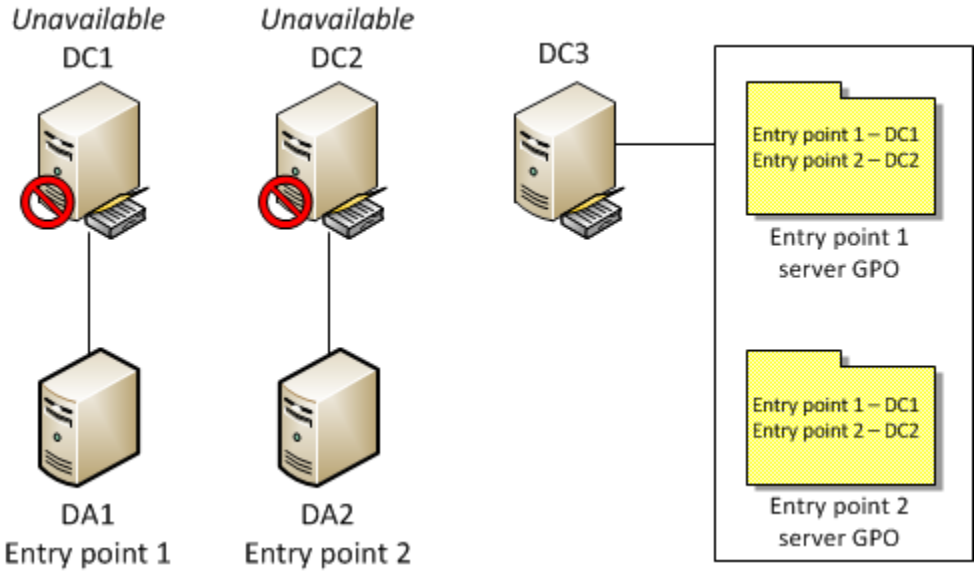

To change two or more domain controllers that manage server GPOs

1. To replace the unavailable domain controller ‘DC2’ with the domain controller ‘DC3’, run the following command:

```
Set-DAEntryPointDC –ExistingDC 'DC2' –NewDC 'DC3' –
ComputerName 'DA2' –ErrorAction Continue
```

   This command updates the domain controller association for the ‘Entry point 2’ server GPO in the registry of DA2 and in the ‘Entry point 2’ server GPO itself; however, it does not update the ‘Entry point 1’ server GPO because the domain controller that manages it is unavailable.

   Tip
   This command uses the Continue value for the *ErrorAction* parameter, which updates the ‘Entry point 2’ server GPO despite the failure to update ‘Entry point 1’ server GPO.

   The resulting configuration is shown in the following diagram.

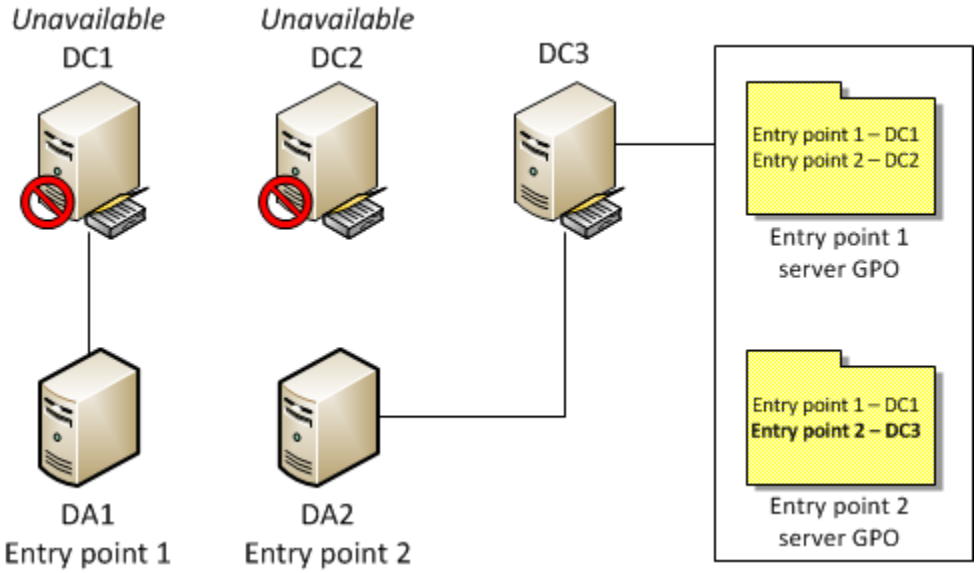

2. To replace the unavailable domain controller ‘DC1’ with the domain controller ‘DC3’, run the following command:

```
Set-DAEntryPointDC –ExistingDC 'DC1' –NewDC 'DC3' –
ComputerName 'DA2' –ErrorAction Continue
```

This command updates the domain controller association for the ‘Entry point 1’ server GPO in the registry of DA1 and in the ‘Entry point 1’ and ‘Entry point 2’ server GPOs. The resulting configuration is shown in the following diagram.

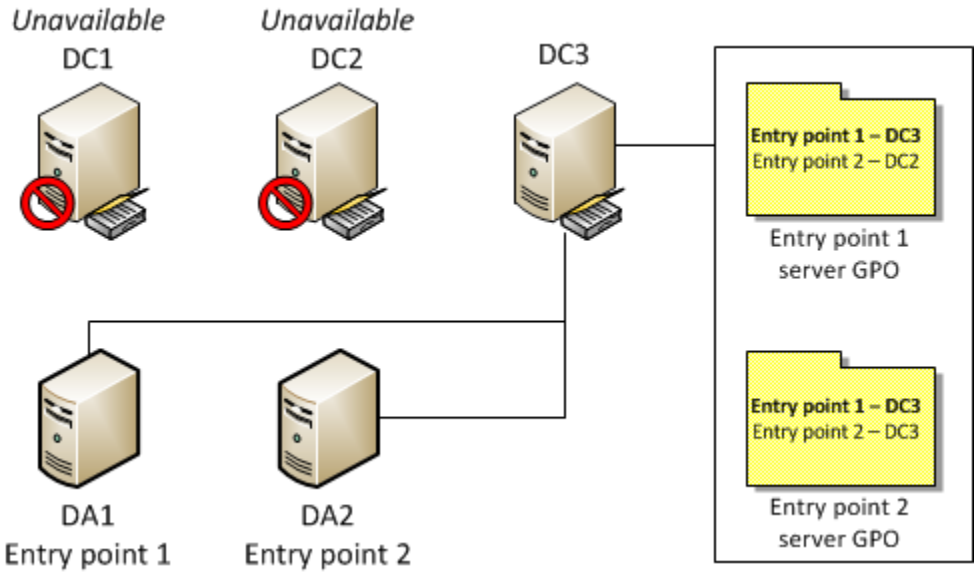

3. To synchronize the domain controller association for the ‘Entry point 2’ server GPO in the ‘Entry point 1’ server GPO, run the command to replace ‘DC2’ with ‘DC3’, and specify the

Remote Access server whose server GPO is not synchronized, in this case ‘DA1’, for the *ComputerName* parameter.

```
Set-DAEntryPointDC –ExistingDC 'DC2' –NewDC 'DC3' –
ComputerName 'DA1' –ErrorAction Continue
```

The final configuration is shown in the following diagram.

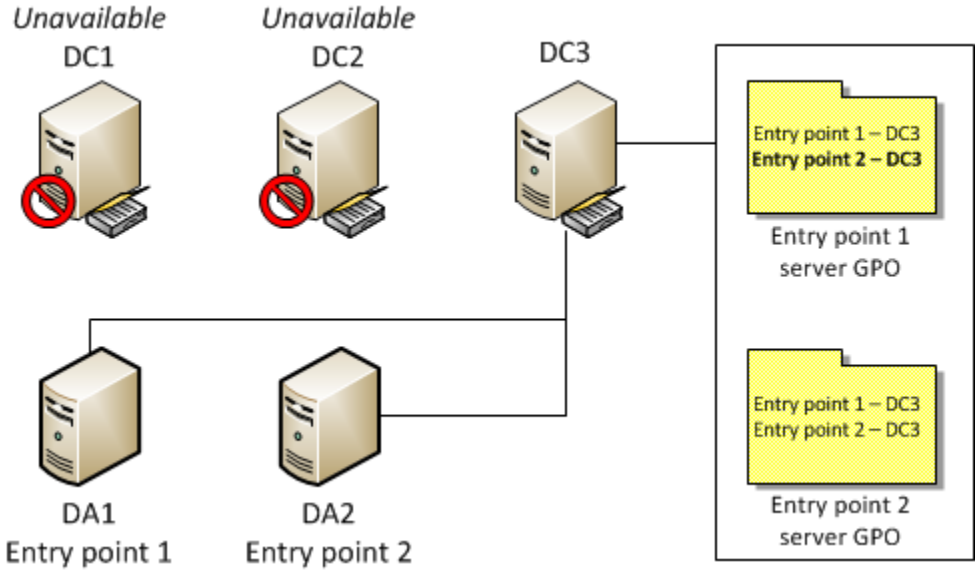

Optimization of configuration distribution

When making configuration changes, the changes are applied only after the server GPOs propagate to the Remote Access servers. To reduce the configuration distribution time, Remote Access automatically selects a writable domain controller which is closest to the Remote Access server when creating its server GPO.

In some scenarios, it may be required to manually modify the domain controller that manages a server GPO in order to optimize configuration distribution time:

- There were no writable domain controllers in the Active Directory site of a Remote Access server at the time of adding it as an entry point. A writable domain controller is now being added to the Active Directory site of the Remote Access server.
- An IP address change, or an Active Directory Sites and Subnets change may have moved the Remote Access server to a different Active Directory site.
- The domain controller association for an entry point was manually modified due to maintenance work on a domain controller, and now the domain controller is back online.

In these scenarios, run the PowerShell cmdlet `Set-DAEntryPointDC` on the Remote Access server and specify the name of the entry point you want to optimize using the parameter *EntryPointName*. You should do this only after the GPO data from the domain controller currently storing the server GPO was already fully replicated to the desired new domain controller.

Note

Before modifying domain controller association, make sure that all of the GPOs in the Remote Access deployment have been replicated to all of the domain controllers in the domain. If the GPO is not synchronized, recent configuration changes may be lost after modifying domain controller association, which may lead to a corrupt configuration. To verify GPO synchronization, see Check Group Policy Infrastructure Status.

To optimize the configuration distribution time, do one of the following:

- To manage the server GPO of entry point 'Entry point 1' on a domain controller in the closest Active Directory site to the Remote Access server 'DA1.corp.contoso.com', run the following command:

```
Set-DAEntryPointDC –EntryPointName 'Entry point 1' –ComputerName
'DA1.corp.contoso.com' –ErrorAction Inquire
```

- To manage the server GPO of entry point 'Entry point 1' on the domain controller 'dc2.corp.contoso.com', run the following command:

```
Set-DAEntryPointDC –EntryPointName 'Entry point 1' –NewDC
'dc2.corp.contoso.com' –ComputerName 'DA1.corp.contoso.com' –
ErrorAction Inquire
```

Note

When modifying the domain controller associated with a specific entry point, you must specify a Remote Access server that is a member of that entry point for the *ComputerName* parameter.

See also

- Step 3: Configure the Multisite Deployment
- Step 1: Implement a Single Server Remote Access Deployment
- Configure a Multisite Deployment
- Deploy Multiple Remote Access Servers in a Multisite Deployment

Step 3: Configure the Multisite Deployment

After configuring the multisite infrastructure follow these steps to set up the Windows Server 2012 Remote Access multisite deployment.

Task	Description
3.1. Configure Remote Access servers	Configure additional Remote Access servers by

Task	Description
	setting up IP addresses, joining them to the domain, and installing the Remote Access role.
3.2. Grant administrator access	Grant privileges on the additional Remote Access servers to the DirectAccess administrator.
3.3. Configure IP-HTTPS for a multisite deployment	Configure the IP-HTTPS certificate used in a multisite deployment.
3.4. Configure the network location server for a multisite deployment	Configure the network location server certificate used in a multisite deployment.
3.5. Configure DirectAccess clients for a multisite deployment	Remove Windows 7 client computers from Windows 8 security groups.
3.6. Enable the multisite deployment	Enable the multisite deployment on the first Remote Access server.
3.7. Add entry points to the multisite deployment	Add additional entry points to the multisite deployment.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

3.1. Configure Remote Access servers

Do this step using Windows PowerShell

To install the Remote Access role

1. Make sure that each Remote Access server is configured with the proper deployment topology (Edge, behind a NAT, single network interface), and corresponding routes as described in Step 1: Implement a Single Server Remote Access Deployment.
2. Configure the IP addresses on each Remote Access server according to the site topology and your organization's IP addressing scheme.
3. Join each Remote Access server to an Active Directory domain.
4. In the Server Manager console, in the **Dashboard**, click **Add roles and features**.
5. Click **Next** three times to get to the server role selection screen.
6. On the **Select Server Roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
7. Click **Next** five times.
8. On the **Confirm installation selections** dialog, click **Install**.

9. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

Steps 1 – 3 must be performed manually, and are not accomplished using this Windows PowerShell cmdlet.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature RemoteAccess -IncludeManagementTools
```

3.2. Grant administrator access

To grant administrator permissions

1. On the Remote Access server in the additional entry point: On the **Start** screen, type **Computer Management**, and then press ENTER.
2. In the left pane, click **Local Users and Groups**.
3. Double-click **Groups**, and then double-click **Administrators**.
4. On the **Administrators Properties** dialog box, click **Add**, and on the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Locations**.
5. On the **Locations** dialog box, in the **Location** tree, click the location that contains the user account of the DirectAccess administrator, and then click **OK**.
6. In the **Enter the object names to select**, enter the user name of the DirectAccess administrator, and then click **OK** twice.
7. On the **Administrators Properties** dialog box, click **OK**.
8. Close the Computer Management window.
9. Repeat this procedure on all Remote Access servers that will be part of the multisite deployment.

3.3. Configure IP-HTTPS for a multisite deployment

On each Remote Access server that will be added to the multisite deployment, an SSL certificate is required to verify the HTTPS connection to the IP-HTTPS web server. Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To obtain an IP-HTTPS certificate

1. On each Remote Access server: On the **Start** screen, type **mmc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, click the Web Server certificate template, and then click **More information is required to enroll for this certificate**.

   If the Web Server certificate template does not appear, ensure that the Remote Access server computer account has enroll permissions for the Web Server certificate template. For more information, see **Configure Permissions on the Web Server Certificate Template**.
8. On the **Subject** tab of the **Certificate Properties** dialog box, in **Subject name**, for **Type**, select **Common name**.
9. In **Value**, type the fully qualified domain name (FQDN) of the Internet name of the Remote Access server (for example, Europe.contoso.com), and then click **Add**.
10. Click **OK**, click **Enroll**, and then click **Finish**.
11. In the details pane of the Certificates snap-in, verify that a new certificate with the FQDN was enrolled with **Intended Purposes** of **Server Authentication**.
12. Right-click the certificate, and then click **Properties**.
13. In **Friendly Name**, type **IP-HTTPS Certificate**, and then click **OK**.

    Tip
    Steps 12 and 13 are optional, but make it easier for you to select the certificate for IP-HTTPS when configuring Remote Access.
14. Repeat this procedure on all Remote Access servers in your deployment.

3.4. Configure the network location server for a multisite deployment

If you selected to set up the network location server website on the Remote Access server when you set up your first server using Deploy a Single Remote Access Server with Advanced Settings, each new Remote Access server that you add needs to be configured with a Web Server certificate that has the same subject name that was selected for the network location server for the first server. Each server requires a certificate to authenticate the connection to the network location server, and client computers located in the internal network must be able to resolve the name of the website in DNS.

To install a certificate for network location

1. On the Remote Access server: On the **Start** screen, type **mmc**, and then press ENTER.

If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.

   Note
   You may also import the same certificate that was used for the network location server for the first Remote Access server.
6. Click **Next** twice.
7. On the **Request Certificates** page, click the Web Server certificate template, and then click **More information is required to enroll for this certificate**.

   If the Web Server certificate template does not appear, ensure that the Remote Access server computer account has enroll permissions for the Web Server certificate template. For more information, see **Configure Permissions on the Web Server Certificate Template**.
8. On the **Subject** tab of the **Certificate Properties** dialog box, in **Subject name**, for **Type**, select **Common name**.
9. In **Value**, type the fully qualified domain name (FQDN) that was configured for the network location server certificate of the first Remote Access server (for example, nls.corp.contoso.com), and then click **Add**.
10. Click **OK**, click **Enroll**, and then click **Finish**.
11. In the details pane of the Certificates snap-in, verify that a new certificate with the FQDN was enrolled with **Intended Purposes** of **Server Authentication**.
12. Right-click the certificate, and then click **Properties**.
13. In **Friendly Name**, type **Network Location Certificate**, and then click **OK**.

    Tip
    Steps 12 and 13 are optional, but make it easier for you to select the certificate for network location when configuring Remote Access.
14. Repeat this procedure on all Remote Access servers in your deployment.

To create the network location server DNS records

1. On the DNS server: On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the left pane of the **DNS Manager** console, open the forward lookup zone for the internal network. Right click the relevant zone and click **New Host (A or AAAA)**.
3. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the name that was used for the network location server for the first Remote Access server. In the **IP address** box, enter the intranet-facing IPv4 address of the Remote

Access server, and then click **Add Host**. On the **DNS** dialog box, click **OK**.

4. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the name that was used for the network location server for the first Remote Access server. In the **IP address** box, enter the intranet-facing IPv6 address of the Remote Access server, and then click **Add Host**. On the **DNS** dialog box, click **OK**.
5. Repeat steps 3 and 4 for every Remote Access server in your deployment.
6. Click **Done**.
7. Repeat this procedure before adding servers as additional entry points in the deployment.

3.5. Configure DirectAccess clients for a multisite deployment

DirectAccess Windows client computers must be members of security group(s) which define their DirectAccess association. Before multisite is enabled, these security group(s) can contain both Windows 8 clients and Windows 7 clients (if the appropriate “downlevel” mode was selected). Once multisite is enabled, existing client security group(s), in single server mode, are converted to security group(s) for Windows 8 only. After multisite is enabled, DirectAccess Windows 7 client computers must be moved to corresponding dedicated Windows 7 client security groups (which are associated with specific entry-points), or they will not be able to connect over DirectAccess. The Windows 7 clients must first be removed from the existing security groups which are now Windows 8 security groups. Caution: Windows 7 client computers that are members of both Windows 7 and Windows 8 client security groups will lose remote connectivity, and Windows 7 clients without SP1 installed will lose corporate connectivity as well. Therefore, all Windows 7 client computers must be removed from Windows 8 security groups.

Remove Windows 7 clients from Windows 8 security groups

1. On the primary domain controller, click **Start**, and then click **Active Directory Users and Computers**.
2. To remove computers from the security group, double-click the security group, and on the **<Group_Name> Properties** dialog box, click the **Members** tab.
3. Select the Windows 7 client computer, and click **Remove**.
4. Repeat this procedure to remove the Windows 7 client computers from the Windows 8 security groups.

Important

When enabling a Remote Access multisite configuration all client computers (Windows 7 and Windows 8) will lose remote connectivity until they are able to connect to the corporate network directly or by VPN to update their group policies. This is true when enabling multisite functionality for the first time, and also when disabling multisite.

3.6. Enable the multisite deployment

To configure a multisite deployment, enable the multisite feature on your existing Remote Access server. Before enabling multisite in your deployment, make sure you have the following information:

1. Global load balancer settings and IP addresses if you want to load balance DirectAccess client connections across all entry-points in your deployment.
2. The security group(s) containing Windows® 7 client computers for the first entry point in your deployment, if you want to enable Remote Access for Windows 7 client computers.
3. Group Policy Object names, if you are required to use non-default Group Policy Objects, which are applied on Windows 7 client computers for the first entry point in your deployment, if you require support for Windows 7 client computers.

Do this step using Windows PowerShell

To enable a multisite configuration

1. On your existing Remote Access server: On the **Start** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**, and then in the **Tasks** pane, click **Enable Multisite**.
3. In the **Enable Multisite Deployment** wizard, on the **Before You Begin** page, click **Next**.
4. On the **Deployment Name** page, in **Multisite deployment name**, enter a name for your deployment. In **First entry point name**, enter a name to identify the first entry point which is the current Remote Access server, and then click **Next**.
5. On the **Entry Point Selection** page, do one of the following:
   - Click **Assign entry points automatically, and allow clients to select manually** to automatically route client computers to the most suitable entry point, while also allowing client computers to select an entry point manually. Manual entry point selection is available only for Windows® 8 computers. Click **Next**.
   - Click **Assign entry points automatically** to automatically route client computers to the most suitable entry point, and then click **Next**.
6. On the **Global Load Balancing** page, do one of the following:
   - Click **No, do not use global load balancing** if you do not want to use a global load balancing, and then click **Next**.

     Note
     When selecting this option client computers connect to their closest entry point automatically.
   - Click **Yes, use global load balancing** if you want to load balance the traffic globally between all entry points. In **Type the global load balancing FQDN to be used by all entry points**, enter the global load balancing FQDN, and in **Type the global load balancing IP address for this entry point** that contains the first Remote Access server, enter the global load balancing IP address for this entry point, and then click

Next.

7. On the **Client Support** page, do one of the following:
   - To limit access to client computers running Windows 8 or later operating systems, click **Limit access to client computers running Windows 8 or a later operating system**, and then click **Next**.
   - To allow client computers running Windows 7 to access this entry point, click **Allow client computers running Windows 7 to access this entry point**, and click **Add**. On the **Select Groups** dialog box, select the security group(s) that contains the Windows 7 client computers, click **OK**, and then click **Next**.
8. On the **Client GPO Settings** page, accept the default GPO for Windows 7 client computers for this entry point, type the name of the GPO that want Remote Access to create automatically, or click **Browse** to locate the GPO for Windows 7 client computers, and then click **Next**.

   Notes
   - The **Client GPO Settings** page appears only when you configure the entry point to allow Windows 7 client computers to access the entry point.
   - You can optionally click **Validate GPOs** to ensure that you have the proper permissions for the selected GPO or GPOs for this entry point. If the GPO does not exist and will be automatically created, then create and link permissions are required. In the case where the GPOs were created manually, then edit, modify security, and delete permissions are required.
9. On the **Summary** page, click **Commit**.
10. On the **Enabling Multisite Deployment** dialog box, click **Close** and then on the Enable Multisite Deployment wizard, click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To enable a multisite deployment named ‘Contoso’ on the first entry point named ‘Edge1-US’. The deployment allows clients to manually select the entry point, and does not use a global load balancer.

```
Enable-DAMultiSite -Name 'Contoso' -EntryPointName 'Edge1-US' -
ManualEntryPointSelectionAllowed 'Enabled'
```

To allow Windows 7 client computers access through the first entry point through the security group DA_Clients_US and using the GPO DA_W7_Clients_GPO_US.

```
Add-DAClient -EntrypointName 'Edge1-US' -DownlevelSecurityGroupNameList
@('corp.contoso.com\DA_Clients_US') -DownlevelGpoName
@('corp.contoso.com\DA_W7_Clients_GPO_US)
```

3.7. Add entry points to the multisite deployment

After enabling multisite in your deployment, you can add additional entry points using the Add an Entry Point wizard. Before adding entry points, make sure you have the following information:

- Global load balancer IP addresses for each new entry point if you are using global load balancing.
- The security group(s) containing Windows® 7 client computers for each entry point that will be added if you want to enable Remote Access for Windows 7 client computers.
- Group Policy Object names, if you are required to use non-default Group Policy Objects, which are applied on Windows 7 client computers for each entry point that will be added, if you require support for Windows 7 client computers.
- In the case where IPv6 is deployed in the organization's network you will need to prepare the IP-HTTPS prefix for the new entry point.

Do this step using Windows PowerShell

To add entry-points to your multisite deployment

1. On your existing Remote Access server: On the **Start** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, click **Configuration**, and then in the **Tasks** pane, click **Add an Entry Point**.
3. In the Add an Entry Point Wizard, on the **Entry Point Details** page, in **Remote Access server**, enter the fully qualified domain name (FQDN) of the server to add. In **Entry point name**, enter the name of the entry point, and then click **Next**.
4. On the **Global Load Balancing Settings** page, enter the global load balancing IP address of this entry point, and then click **Next**.

   Note
   The **Global Load Balancing Settings** page appears only when the multisite configuration uses a global load balancer.
5. On the **Network Topology** page, click the topology that corresponds with the network topology of the Remote Access server that you are adding, and then click **Next**.
6. On the **Network Name or IP Address** page, in **Type in the public name or IP address used by clients to connect to the remote access server**, enter the public name or IP address used by clients to connect to the Remote Access server. The public name corresponds with the subject name of the IP-HTTPS certificate. In the case where Edge network topology was implemented, the IP address is that of the external adapter of the Remote Access server. Click **Next**.
7. On the **Network Adapters** page, do one the following:
   - If you are deploying a topology with two network adapters, in **External adapter**, select the adapter that is connected to the external network. In **Internal adapter**, select the adapter that is connected to the internal network.
   - If you are deploying a topology with one network adapter, in **Network adapter**, select

the adapter that is connected to the internal network.

8. On the **Network Adapters** page, in **Select the certificate used to authenticate IP-HTTPS connections**, click **Browse** to locate and select the IP-HTTPS certificate. Click **Next**.
9. If IPv6 is configured on the corporate network, on the **Prefix Configuration** page, in **IPv6 prefix assigned to client computers**, enter an IP-HTTPS prefix to assign IPv6 addresses to DirectAccess client computers, and click **Next**.
10. On the **Client Support** page, do one of the following:
    - To limit access to client computers running Windows 8 or later operating systems, click **Limit access to client computers running Windows 8 or a later operating system**, and then click **Next**.
    - To allow client computers running Windows 7 to access this entry point, click **Allow client computers running Windows 7 to access this entry point**, and click **Add**. On the **Select Groups** dialog box, select the security group(s) that contain the Windows 7 client computers that will connect to this entry point, click **OK**, and click **Next**.
11. On the **Client GPO Settings** page, accept the default GPO for Windows 7 client computers for this entry point, type the name of the GPO that you want Remote Access to create automatically, or click **Browse** to locate the GPO for Windows 7 client computers, and click **Next**.

    Notes
    - The **Client GPO Settings** page appears only when you configure the entry point to allow Windows 7 client computers to access the entry point.
    - You can optionally click **Validate GPOs** to ensure that you have the proper permissions for the selected GPO or GPOs for this entry point. If the GPO does not exist and will be automatically created, then create and link permissions are required. In the case where the GPOs were created manually, then edit, modify security, and delete permissions are required.
12. On the **Server GPO Settings** page, accept the default GPO for this Remote Access server, type the name of the GPO that you want Remote Access to create automatically, or click **Browse** to locate the GPO for this server, and then click **Next**.
13. On the **Network Location Server** page, click **Browse** to select the certificate for the network location server website running on the Remote Access server, and then click **Next**.

    Note

    The **Network Location Server** page appears only when the network location server website is running on the Remote Access server.
14. On the **Summary** page, review the entry point settings, and then click **Commit**.
15. On the **Adding Entry Point** dialog box, click **Close** and then on the Add an Entry Point Wizard, click **Close**.

    Note

If the entry point that was added is in a different forest than the existing entry points or client computers, then it is required to click **Refresh Management Servers** in the **Tasks** pane to discover the domain controllers and System Center Configuration Manager in the new forest.

16. Repeat this procedure from step 2 for every entry point that you want to add to your multisite deployment.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To add the computer edge2 from the corp2 domain as a second entry point named Edge2-Europe. The entry point configuration is: a client IPv6 prefix '2001:db8:2:2000::/64', a connect to address (the IP-HTTPS certificate on the edge2 computer) 'edge2.contoso.com', a server GPO named "DirectAccess Server Settings – Edge2-Europe", and the internal and external interfaces named Internet and Corpnet2 respectively:

```
Add-DAEntryPoint -RemoteAccessServer 'edge2.corp2.corp.contoso.com' -Name 'Edge2-Europe'
-ClientIPv6Prefix '2001:db8:2:2000::/64' -ConnectToAddress 'Europe.contoso.com' -
ServerGpoName 'corp2.corp.contoso.com\DirectAccess Server Settings – Edge2-Europe' -
InternetInterface 'Internet' -InternalInterface 'Corpnet2'
```

To allow Windows 7 client computers access through the second entry point through the security group DA_Clients_Europe and using the GPO DA_W7_Clients_GPO_Europe.

```
Add-DAClient -EntrypointName 'Edge2-Europe' -DownlevelGpoName @('corp.contoso.com\
DA_W7_Clients_GPO_Europe') -DownlevelSecurityGroupNameList
@('corp.contoso.com\DA_Clients_Europe')
```

See also

- Step 2: Configure the Multisite Infrastructure
- Configure a Multisite Deployment
- Deploy Multiple Remote Access Servers in a Multisite Deployment

Step 4: Verify the Multisite Deployment

This topic describes how to verify that you have correctly configured your Remote Access multisite deployment.

To verify access to internal resources through the multisite deployment

1. Connect a DirectAccess client computer to the corporate network and obtain the group

policy.

2. Connect the client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

3. Test connectivity through each server in the multisite deployment by turning off, or disconnecting from the external network, all but one of the Remote Access servers. On the client computer, attempt to access corporate resources. Repeat the test on a different multisite server. It may take up to 10 minutes for the client computer to connect to the new entry point. This is because probing is turned off for 10 minutes for an entry point after it is deemed unreachable, in order to optimize bandwidth and battery life. Alternatively, you can switch between the various entry-points manually by choosing the desired entry-point from the combo-box shown when executing **daprop.exe**.

   You should be able to access all corporate resources through each multisite server.

4. Connect a Windows® 7 client computer to the corporate network and obtain the group policy.

5. Connect the Windows 7 client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

6. Test connectivity for Windows 7 clients through each server in the multisite deployment by accessing the Active Directory Users and Computers console, and moving the client computer to the security group that corresponds to each server. After the changes have replicated throughout the domain, restart the client computer while connected to the corporate network to obtain the new group policy. Attempt to access corporate resources. Repeat the test on a different multisite server.

   You should be able to access all corporate resources through each multisite server.

   In a production environment this method may not be feasible due to the amount of time required for changes to replicate throughout the domain. You may want to force replication where possible. Testing can also be done from multiple different Windows 7 client computers that are already members of the different Windows 7 security groups in the multisite deployment.

Troubleshoot a Multisite Deployment

This topic describes how to troubleshoot the most common errors that may occur when configuring a multisite Remote Access deployment. For a list of known issues when configuring multisite Remote Access, see Known issues.

- Troubleshooting Enabling Multisite
- Troubleshooting Adding Entry Points
- Troubleshooting Setting the Entry Point Domain Controller

- Troubleshooting Web Probe URLs
- Troubleshooting General Issues

Troubleshooting Enabling Multisite

This topic contains troubleshooting information for issues related to the `Enable-DAMultisite` command. To confirm that the error you received is related to enabling multisite, check in the Windows Event log for the event ID 10051.

User connectivity issues

Users may experience connectivity issues when you enable multisite if the configuration is not correct.

Cause

In a multisite deployment, Windows 8 client computers are able to roam between different entry points. Windows 7 client computers must be associated with a specific entry point in the multisite deployment. If client computers are not in the correct security group, they may receive the wrong group policy settings.

Solution

DirectAccess requires at least one security group for all Windows 8 client computers; we recommend using one security group for all Windows 8 computers per domain. DirectAccess also requires a security group for Windows 7 client computers for each entry point. Each client computer should be included in only one security group. Therefore, you should make sure that the security groups for Windows 8 clients contain only computers running Windows 8, and that each Windows 7 client computer belongs to a single dedicated security group for the relevant entry point and that no Windows 8 clients belong to the Windows 7 security groups.

Configure the Windows 8 security groups on the **Select Groups** page of the **DirectAccess Client Setup** wizard. Configure Windows 7 security groups on the **Client Support** page of the **Enable Multisite Deployment** wizard, or on the **Client Support** page of the **Add an Entry Point** wizard.

Kerberos proxy authentication

Error received—Kerberos proxy authentication is not supported in a multisite deployment. You must enable the use of computer certificates for IPsec user authentication.

Cause

Computer certificate authentication must be enabled before enabling multisite.

Solution

To enable computer certificate authentication:

1. In the Remote Access Management console, in the details pane, under **Step 2 Remote Access Server**, click **Edit**.
2. In the **Remote Access Server Setup** wizard, on the **Authentication** page, select the **Use computer certificates** check box, and select the root or intermediate certification authority that issues certificates in your deployment.

To enable computer certificate authentication using Windows PowerShell, use the `Set-DAServer` cmdlet and specify the *IPsecRootCertificate* parameter.

IP-HTTPS certificates

Error received—The DirectAccess server uses a self-signed IP-HTTPS certificate. Configure IP-HTTPS to use a signed certificate from a known CA.

Cause

The IP-HTTPS certificate is self-signed. You cannot use self-signed certificates in a multisite deployment.

Solution

To select an IP-HTTPS certificate:

1. In the Remote Access Management console, in the details pane, under **Step 2 Remote Access Server**, click **Edit**.
2. In the **Remote Access Server Setup** wizard, on the **Network Adapters** page, under **Select the certificate used to authenticate IP-HTTPS connections**, make sure the **Use a self-signed certificate created automatically by DirectAccess** check box is cleared, and click **Browse** to select a certificate issued by a trusted CA.

Network location server

- **Issue 1**

  Error received—DirectAccess is configured to use a self-signed certificate for the network location server. Configure the network location server to use a signed certificate from a CA.

  Cause

  The network location server is deployed on the Remote Access server and uses a self-signed certificate. You cannot use self-signed certificates in a multisite deployment.

  Solution

  To select a network location server certificate:

  a. In the Remote Access Management console, in the details pane, under **Step 3 Infrastructure Servers**, click **Edit**.

  b. In the **Infrastructure Server Setup** wizard, on the **Network Location Server** page, under **The network location server is deployed on the Remote Access server**, make sure the **Use a self-signed certificate** check box is cleared, and click **Browse** to select a certificate issued by an Enterprise CA.

- **Issue 2**

Error received—To deploy a network load balanced cluster or multisite deployment, obtain a certificate for the network location server with a subject name that is different from the internal name of the Remote Access server.

Cause

The subject name of the certificate used for the network location server website is the same as the internal name of the Remote Access server. This will cause name resolution issues.

Solution

Obtain a certificate with a subject name that is not the same as the internal name of the Remote Access server.

To configure the network location server:

a. In the Remote Access Management console, in the details pane, under **Step 3 Infrastructure Servers**, click **Edit**.

b. In the **Infrastructure Server Setup** wizard, on the **Network Location Server** page, under **The network location server is deployed on the Remote Access server**, click **Browse** to select the certificate you previously obtained. The certificate must have a subject name that is different from the internal name of the Remote Access server.

Windows 7 client computers

Warning received—When enabling multisite, the security groups configured for DirectAccess clients must not contain Windows 7 computers. To support Windows 7 client computers in a multisite deployment, select a security group containing the clients for each entry point.

Cause

In the existing DirectAccess deployment, Windows 7 client support was enabled.

Solution

DirectAccess requires at least one security group for all Windows 8 client computers and a security group for Windows 7 client computers for each entry point. Each client computer should be included in only one security group. Therefore, you should make sure that the security group for Windows 8 clients contains only computers running Windows 8, and that each Windows 7 client computer belongs to a single dedicated security group for the relevant entry point and that no Windows 8 clients belong to the Windows 7 security groups.

Active Directory site

Error received—The server <server_name> is not associated with an Active Directory Site.

Cause

DirectAccess could not determine the Active Directory site. In the Active Directory Sites and Services console you can configure the different subnets for your network and associate each subnet with the relevant Active Directory site. This error may occur if the IP address of the Remote Access server does not belong to any of the subnets, or that the subnet to which the IP address belongs is not defined with an Active Directory site.

Solution

Confirm that this is the issue by running the command `nltest /dsgetsite` on your Remote Access server. If this is the issue, the command will return ERROR_NO_SITENAME. To solve this issue, on your domain controller, make sure that a subnet which contains the internal server IP address exists and that it is defined with an Active Directory site.

Saving server GPO settings

Error received—An error occurred while saving Remote Access settings to GPO <GPO_name>.

Cause

Changes to the server GPO couldn't be saved either due to connectivity issues or if there is a sharing violation on the registry.pol file, for example, another user has locked the file.

Solution

Make sure that there is connectivity between the Remote Access server and the domain controller. If there is connectivity, check on the domain controller if the registry.pol file is locked to another user, and if necessary, end that user session to unlock the file.

Internal error occurred

Error received—An internal error occurred.

Cause

This may be caused by an unexpected entry point table configuration in the client GPO. This can occur if the administrator uses DirectAccess client cmdlets to edit the entry point table in the client GPO.

Solution

Review the entry point table configuration in all client GPOs and fix any inconsistencies in the multisite configuration between the different instances of the client GPOs and the DirectAccess configuration. Use the `Get-DaEntryPointTableItem` cmdlet with the name of the client GPO to get the entry point table on the client. Use the `Get-NetIPHttpsConfiguration` cmdlet to get all the IP-HTTPS profiles for all the entry points.

For further information, see DirectAccess Client Cmdlets in Windows PowerShell.

Troubleshooting Adding Entry Points

This topic contains troubleshooting information for issues related to the `Add-DAEntryPoint` command. To confirm that the error you received is related to adding an entry point, check in the Windows Event log for the event ID 10067.

Internal error occurred

Error received—An internal error occurred.

To troubleshoot this error, see Internal error occurred.

Missing RemoteAccessServer parameter

Error received—You must supply a value for the parameter RemoteAccessServer.

Cause

When adding a new entry point to a multisite deployment, you must specify the parameter *RemoteAccessServer*, which is the name of the server which you want to add as the new entry point.

Solution

Run the command and make sure to specify the *RemoteAccessServer* parameter with the name of the server to be added as an entry point.

Remote Access is not configured

Error received—Remote Access is not configured on <server_name>. Specify the name of a server that belongs to a multisite deployment.

Cause

Remote Access is not configured on the computer specified by the *ComputerName* parameter, or on the computer on which you run the command.

When adding a new entry point to a multisite deployment, you must specify two parameters: *ComputerName* and *RemoteAccessServer*. The *ComputerName* is the name of a server which is already part of the multisite deployment, the *RemoteAccessServer* is the name of the server which you want to add as the new entry point. If you run from a computer that is part of the multisite deployment, the ComputerName parameter is not required.

Solution

Run the command and make sure to specify the *ComputerName* parameter with the name of the server that is already configured as part of the multisite deployment, or run the command from a computer that is part of the multisite deployment.

Multisite not enabled

Error received—You must enable a multisite deployment before performing this operation. Use the `Enable-DAMultiSite` cmdlet to do this.

Cause

Multisite is not enabled on the server specified by the *ComputerName* parameter. To add a new entry point to a Remote Access deployment, you must first enable multisite.

Solution

Enable multisite using the `Enable-DaMultiSite` cmdlet. For more information, see Deploy multisite Remote Access.

IPv6 prefix issues

- **Issue 1**

  Error received—IPv6 is deployed in the internal network, but you have not specified a client IPv6 prefix.

  Cause

  IPv6 is deployed on the corporate network, and an IP-HTTPS prefix is required. However, a prefix was not specified in the *ClientIPv6Prefix* parameter for the new entry point.

  Solution

  a. Assign a unique IP-HTTPS prefix to the new entry point and ensure that packets targeted to an IP address under this prefix will be routed to the server you are adding.

  b. Run the `Add-DAEntryPoint` cmdlet and specify the IP-HTTPS prefix in the *ClientIPv6Prefix* parameter.

- **Issue 2**

  Error received—The client IPv6 prefix is already in use by another entry point. Specify an alternate value.

  Cause

  The IP-HTTPS prefix specified in the *ClientIPv6Prefix* parameter is already used by a different entry point

  Solution

  a. Assign a unique IP-HTTPS prefix to the new entry point and ensure that packets targeted to an IP address under this prefix will be routed to the server you are adding.

  b. Run the `Add-DAEntryPoint` cmdlet and specify the IP-HTTPS prefix in the *ClientIPv6Prefix* parameter.

ConnectTo address

Error received—The address (<connect_to_address>) to which DirectAccess clients connect on the RemoteAccess server is the same as the network location server address. Specify an alternate value.

Cause

The ConnectTo address and the network location server address are the same.

Solution

The ConnectTo address should be resolvable over the Internet to allow client machines to connect over IP-HTTPS. The network location server address should be resolvable over the corporate network but should not be resolvable over the Internet. Make sure that the network

location server and the ConnectTo addresses are not the same. Select different addresses and try again.

DirectAccess or VPN already installed

Error received—A VPN installation was detected on the server <server_name>. Specify an alternate server that does not have Remote Access installed, or remove the VPN configuration from the server.

Or

Remote Access is already installed on server <server_name>. Specify an alternate server not running DirectAccess, or remove the existing DirectAccess configuration from the server.

Cause

DirectAccess or VPN is already configured on the new entry point. You cannot add a configured entry point to a multisite deployment.

Solution

To add a server to a multisite deployment, you must install the Remote Access role on the server but DirectAccess and VPN should not be configured.

Run the command and make sure that the server you specify in the *RemoteAccessServer* parameter does not have DirectAccess or VPN configured.

IPsec root certificate

Error received—The configured IPsec root certificate cannot be located on server <server_name>.

Cause

The certificate of the root or intermediate certification authority (CA) that issues computer certificates could not be found on the server you are trying to add to the deployment.

Solution

In the Remote Access Management console, in **Step 2 Remote Access Server**, click **Edit**, and on the **Authentication** page, under **Use computer certificates**, make sure that the certificate selected is valid. If the certificate is valid, make sure that it is located under the trusted root CA on the server you want to add and try again.

Note

The certificate must be the same certificate with the same Thumbprint.

If the certificate is not valid, select a valid certificate which is configured as the trusted root CA on all the Remote Access servers.

Mixing IPv6 and IPv4 entry points

When DirectAccess is installed for the first time, the internal network adapter is inspected to determine if the network contains IPv4 addresses only (IPv4-only network), IPv6 and IPv4 addresses, or IPv6 addresses only (IPv6-only network). The information is used to determine the deployment type (IPv4 only, IPv6+IPv4, or IPv6 only).

- **Issue 1**

  Warning received—The Remote Access server being added is configured with both IPv4 and IPv6 addresses. This is an IPv4 only deployment, and Remote Access will ignore the IPv6 addresses.

  Cause

  When this deployment was first installed, the internal network was detected as an IPv4 only network. In a multisite deployment, different entry points are assumed to be located in different subnets with different characteristics. Therefore, although the deployment is configured as an IPv4 only deployment, it can contain an entry point located in an IPv6+IPv4 subnet. However, although the entry point will be added to the deployment, DirectAccess will ignore the IPv6 addresses configured on the new entry point's internal interface.

  Solution

  If the entire internal network is configured with IPv6 and IPv4 addresses, consider moving to an IPv6+IPv4 deployment to benefit from IPv6 technologies. See "Transitioning from a pure IPv4 to an IPv6+IPv4 corporate network" in Step 3: Plan the Multisite Deployment.

- **Issue 2**

  Error received—The internal network adapters of the Remote Accces servers in this multisite deployment are configured with IPv4 addresses. The entry point you are adding must also be configured with an IPv4 address on the internal network adapter.

  Cause

  When this deployment was first installed, the internal network was detected as an IPv4 only network. Remote Access detected that the entry point that you are trying to add is configured with only IPv6 addresses on its internal network. This is not allowed in an IPv4 only deployment.

  Solution

  If the entire network is already configured with IPv6 addresses, you should move to an IPv6+IPv4 or IPv6-only deployment. See "Plan the transition to IPv6 when multisite Remote Access is deployed".

- **Issue 3**

  Error received—This entry point is located in an IPv4 network, but previous entry points are located in an IPv6 network. Connect this entry point to the IPv6 network before adding it to the same multisite deployment.

  Cause

When this deployment was first installed, it was detected that the internal network is IPv6+IPv4 or IPv6-only. It was detected that only IPv4 addresses are configured on the internal network on the new entry point you are trying to add. This is not allowed in IPv6+IPv4, or IPv6-only deployments.

Solution

Configure the new entry point with IPv6 addresses and then add the entry point to the multisite deployment.

- **Issue 4**

  Warning received—The internal network adapter on the Remote Access server is not configured with an IPv4 address. DNS64 and NAT64 will not be configured on this server. DirectAccess clients can access IPv6 internal servers only.

  Cause

  When this deployment was first installed, it was detected that the internal network is IPv6+IPv4. In this deployment mode, DNS64 and NAT64 are enabled to allow client computers to access machines on the internal network that are configured with only IPv4 addresses.

  When adding the new entry point, Remote Access detected that the internal interface on the new computer has only IPv6 addresses. To configure DNS64 and NAT64, an IPv4 address is required in order to route packets from the Remote Access server to the IPv4 only computer. Since no such IP exists on the new computer, NAT64 and DNS64 won't be configured on the Remote Access server. Therefore, client machines accessing the corporate network over DirectAccess using this entry point won't be able to access IPv4 only servers on the internal network. For information on how to transition to an IPv6+IPv4 network, or an IPv6-only network, see “Plan the transition to IPv6 when multisite Remote Access is deployed”.

  Solution

  Add an IPv4 address to the new Remote Access server to ensure that DNS64 and NAT64 work correctly.

Domain issues with the ServerGpoName

- **Issue 1**

  Error received—The domain specified in the ServerGpoName parameter <server_GPO> does not exist. Specify the domain <domain_name> instead.

  Cause

  The domain name part of the server GPO name which was sent by the administrator was not found.

  Solution

  Make sure that you typed the domain name correctly. If the domain name is spelled correctly, try again using the fully qualified domain name (FQDN).

- **Issue 2**

Error received—The server GPO must be located in the Remote Access server domain. Specify the domain <domain_name> in the ServerGpoName parameter.

Cause

The domain of the server GPO is not the same as the one to which the Remote Access server belongs.

Solution

The server GPO should be located in the same domain as the Remote Access server. Use the server's domain name for the server GPO and try again.

Split-brain DNS

Warning received—The NRPT entry for the DNS suffix <DNS_suffix> contains the public name used by client computers to connect to the Remote Access server. Add the name <connect_to_address> as an exemption in the NRPT.

Cause

You are using split brain DNS. To allow clients to connect using IP-HTTPS, you should make sure that the ConnectTo address selected is exempt in the NRPT rules.

Solution

If you have a multisite deployment, make sure that all connect to addresses of the different entry points are exempt from the NRPT rules.

To exempt an address in the NRPT rules:

1. In the Remote Access Management console, under **Step 3 Infrastructure Servers**, click **Edit**.
2. In the **Infrastructure Server Setup** wizard, on the **DNS** page, double-click the table to enter a new name suffix.
3. On the **DNS Server Addresses** dialog box, in DNS suffix, enter the ConnectTo address of the entry point, and then click **Apply**.

When you add name suffixes without specifying a server address, the suffix is treated as an NRPT exemption.

Saving server GPO settings

Error received—An error occurred while saving Remote Access settings to GPO <GPO_name>.

To troubleshoot this error, see Saving server GPO settings.

GPO updates cannot be applied

Warning received—GPO updates cannot be applied on <server_name>. Changes will not take effect until the next policy refresh.

Cause

An error occurred while trying to refresh policies on the specified computer. Therefore, changes made will not take effect until the next policy refresh.

Solution

To force a policy refresh, run `gpupdate /force` on the specified computer.

Troubleshooting Setting the Entry Point Domain Controller

This topic contains troubleshooting information for issues related to the `Set-DAEntryPointDC` command. To confirm that the error you received is related to setting the entry point domain controller, check in the Windows Event log for the event ID 10065.

Saving server GPO settings

Error received—An error occurred while saving Remote Access settings to GPO <GPO_name>.

To troubleshoot this error, see Saving server GPO settings.

Remote Access is not configured

Error received—Remote Access is not configured on <server_name>. Specify the name of a server that belongs to a multisite deployment.

Or

Remote Access is not configured on the server <server_name>. Specify a computer with DirectAccess enabled.

Cause

Remote Access is not configured on the computer specified by the *ComputerName* parameter.

The `Set-DaEntryPointDC` cmdlet is available only on servers that are part of a configured multisite deployment.

Solution

Run the command and make sure to specify the *ComputerName* parameter with the name of the server that is already configured as part of the multisite deployment.

Multisite is not enabled

Error received—You must enable a multisite deployment before performing this operation. Use the `Enable-DAMultiSite` cmdlet to do this.

Cause

Multisite is not enabled on the server specified by the *ComputerName* parameter.

The `Set-DaEntryPointDC` cmdlet is available only on servers that are part of a configured multisite deployment.

Solution

Run the command and make sure to specify the *ComputerName* parameter with the name of the server that is already configured as part of the multisite deployment.

Entry point and domain controller not provided in cmdlet

The `Set-DaEntryPointDC` cmdlet enables you to change the domain controller that is associated with different entry points, for example, if a particular domain controller is no longer available. You can update a specific entry point to use a different domain controller, or you can update all entry points which use a specific domain controller to use a new domain controller. In the first case, you should use the *EntryPointName* parameter to specify which entry point should be updated. In the second case, you should use the *ExistingDC* parameter to specify which domain controller should be replaced. You can specify only one of these parameters.

Error received—No required parameters were specified. Provide the name of an entry point or an existing domain controller.

Or

Cmdlet `Set-DaEntryPointDC` is missing all required parameters.

Cause

The *EntryPointName* or *ExistingDC* parameters were not specified, or both parameters were specified, for the `Set-DaEntryPointDC` cmdlet.

Solution

Run the command and make sure to specify either the *EntryPointName* parameter or the *ExistingDC* parameter.

Could not locate domain controller

Error received—Unable to locate a new domain controller automatically. Retry later or verify domain controller settings.

Cause

The computer specified with the *ComputerName* parameter is not reachable over RPC or the domain does not contain any available writable domain controllers.

Solution

Make sure that the remote computer is accessible over RPC and that there is a writable domain controller available for the domain. If a writable domain controller is available for the domain, you can also specify its name explicitly using the *NewDC* parameter.

Could not connect to domain controller

- **Issue 1**

  Error received—The domain controller <domain_controller> cannot be reached. Check network connectivity and server availability.

  Cause

  The domain controller cannot be reached. This occurs only when the administrator specifies a domain controller in the *NewDC* or *ExistingDC* parameters.

  Solution

  Make sure that the domain controller's name is spelled correctly. If you used a short name to specify the name, use the FQDN and try again.

- **Issue 2**

  Error received—The domain controller <domain_controller> cannot be contacted.

  Cause

  There may be a network issue that means the domain controller specified in the *NewDC* parameter, or any other existing domain controller in the configuration cannot be reached.

  Solution

  Make sure that the domain controller's name is spelled correctly, make sure it exists, is running, is writable, and that there is a trust relationship between the domain controller and the domain.

- **Issue 3**

  Error received—Domain controller <domain_controller> cannot be reached for %2!s!.

  Cause

  To maintain the configuration consistency in a multisite deployment, it is important to make sure that each GPO is managed by a single domain controller. When the domain controller that manages an entry point's server GPO is not available, Remote Access configuration settings cannot be read or modified.

  Solution

  Follow the procedure "To change the domain controller that manages server GPOs" described in 2.4. Configure GPOs.

- **Issue 4**

  Error received—The primary domain controller in domain <domain_name> cannot be reached.

  Cause

  To maintain the configuration consistency in a multisite deployment, it is important to make sure that each GPO is managed by a single domain controller. Client GPOs are managed on the primary domain controller. If the primary domain controller is not available, Remote Access configuration settings cannot be read or modified.

  Solution

Follow the procedure “To transfer the PDC emulator role” described in 2.4. Configure GPOs.

Read-only domain controller

Error received—The domain controller <domain_controller> is read-only. Specify a domain controller that is not read-only.

Cause

The domain controller specified with the *NewDC* parameter is read-only.

Solution

When using the `Set-DAEntryPointDC`, the *NewDC* parameter is used to update the domain controller associated with a particular entry point, or to update all entry points associated with a domain controller. Therefore, the new domain controller must be writable. Specify a writable domain controller in the *NewDC* parameter and try again.

Cannot retrieve GPO

- **Issue 1**

  Error received—GPO <GPO_name> on domain controller <previous_domain_controller> cannot be retrieved from domain controller <replacement_domain_controller> because they are not in the same domain.

  Cause

  The Remote Access server and the domain controller are not in the same domain; therefore, the GPO cannot be retrieved.

  Solution

  If you tried to update a specific entry point, make sure that the new domain controller is in the same domain as the entry point server. If you tried to update a specific domain controller, make sure that the new domain controller is in the same domain as the one you are trying to replace.

- **Issue 2**

  Error received—GPO <GPO_name> on domain controller <previous_domain_controller> cannot be retrieved from domain controller <replacement_domain_controller>. Wait until domain replication completes and then try again.

  Cause

  When trying to update an entry point domain controller, the cmdlet tries to read the server GPO from the new domain controller; however, the GPO cannot be found on the new domain controller because it has not yet replicated.

  Solution

  The server GPO does not exist on the new domain controller. Make sure that the GPOs have replicated successfully to the new domain controller and try again.

- **Issue 3**

Error received—You do not have permissions to access GPO <GPO_name>.

Cause

When trying to update an entry point domain controller, the cmdlet tries to read the server GPO from the new domain controller; however, the GPO cannot be read on the new domain controller because you do not have the correct permissions.

Solution

The GPO exists on the domain controller, but it cannot be read. Make sure that you have the required permissions and try again.

Entry point not part of multisite deployment

Error received—Entry point <entry_point_name> is not part of the multisite deployment. Specify an alternate value.

Cause

The entry point name you specified was not found.

Solution

Make sure that the entry point name is spelled correctly and that GPOs are replicated to the required domain controllers, and then try again. To view the assigned domain controller for each entry point, use `Get-DAEntryPointDC`.

Remote Access server settings

- **Issue 1**

  Error received—Server <server_name> in entry point <entry_point_name> cannot be accessed.

  Cause

  When trying to update an entry point domain controller, the cmdlet tries to read and write the entry point domain controller from all relevant Remote Access servers. The cmdlet was not able to read the data from one or more Remote Access servers.

  Solution

  Make sure that all relevant Remote Access servers are running and that you have local administrator permissions on all of them and then try again.

- **Issue 2**

  Error received—Settings cannot be saved to the registry on server <server_name> in entry point <entry_point_name>.

  Cause

  When trying to update an entry point domain controller, the cmdlet tries to read and write the entry point domain controller from all relevant Remote Access servers. The cmdlet was not able to write the data to one or more Remote Access servers.

  Solution

Make sure that all relevant Remote Access servers are running and that you have local administrator permissions on all of them and then try again.

- **Issue 3**

  Error received—GPO updates cannot be applied on <server_name>. Changes will not take effect until the next policy refresh.

  Cause

  When using the cmdlet `Set-DAEntryPointDC`, the *ComputerName* parameter specified is a Remote Access server in an entry point other than the last one added to the Multisite deployment.

  Solution

  Any servers that were not updated can be seen using the **Configuration Status** in the **DASHBOARD** of the Remote Access Management Console. This does not cause any functional problems; however, you can run `gpupdate /force` on any servers that were not updated to get the configuration status updated immediately.

Problem resolving FQDN

Error received—Server <server_name> in entry point <entry_point_name> cannot be accessed.

Cause

While getting the list of DirectAccess servers to modify, the cmdlet was not able to resolve the fully qualified domain name (FQDN) of one of the servers from its computer SID.

Solution

The entry point specified in the error message is associated with a domain controller. Make sure that the domain controller is available for the entry point. If the computer to which the specified SID belongs was removed from the domain, ignore this message and then remove the server from the multisite deployment.

No entry points to update

Warning received—Domain controller settings were not modified. If you think changes are required, ensure that cmdlet parameters are configured correctly, and that GPOs are replicated to the required domain controllers.

Cause

When calling the `Set-DaEntryPointDC` cmdlet with the *ExistingDC* parameter, DirectAccess checks all the entry points and updates the entry points that are associated with the specified domain controller. However, no entry point uses the specified *ExistingDC*.

Solution

To see the list of entry points and their associated domain controllers, use the `Get-DAEntryPointDC` cmdlet. If changes should have been made, make sure that the cmdlet

parameters are spelled correctly, and that the GPOs are replicated to the required domain controllers, and then try again.

Troubleshooting Web Probe URLs

When remote client computers connect to the corporate network over DirectAccess, a web probe URL is used to validate that the connection was successful and that the remote client is able to access corporate network resources. To load balance web probe requests in a multisite deployment, a web probe URL is added for each entry point.

This topic contains troubleshooting information for issues related to web probe URL.

Adding and removing web probe URLs

- **Issue 1**

  Warning received—A DNS entry for web probe <web_probe_hostname> and IP address <IP_address> cannot be added. Add the entry manually.

  Cause

  Adding a DNS entry for the web probe URL for the new entry point failed.

  Solution

  Manually add a DNS record with the name and IP addresses specified in the warning message.

- **Issue 2**

  Warning received—A DNS entry for web probe <web_probe_hostname> and IP address <IP_address> cannot be removed. Remove the entry manually.

  Cause

  Removing the DNS entry for the web probe URL for the entry point you are removing failed.

  Solution

  Manually remove the DNS record with the name and IP addresses specified in the warning message.

IP address of management server associated with web probe URL

Warning received—One or more IP addresses of management server <management_server> cannot be added because they are associated with the web probe URL used to verify DirectAccess client connectivity. The server will be added without the conflicting IP addresses. Ensure that IP addresses associated with the probe URL are not configured on any management servers.

Cause

The IP address of the management server you are trying to add conflicts with one of the web probe URLs. Accessing the web probe server should not be available on the infrastructure tunnel.

Solution

Make sure that any IP addresses associated with the web probe URL are not configured on any management server.

Cannot verify IP addresses

Warning received—The check to verify that management server IP addresses do not conflict with the IP addresses associated with the web probe URL cannot be completed. The action will continue without the check. Ensure that the URL is configured correctly in DNS, and that IP addresses associated with the web probe are not configured on any management servers.

Cause

Validating that the web probe does not conflict with IP addresses of the management servers failed. Accessing the web probe server should not be available on the infrastructure tunnel.

Solution

Make sure that any IP addresses associated with the web probe URL are not configured on any management server.

DNS entry cannot be updated

Warning received—A DNS entry for web probe <web_probe_hostname> and IP address <previous_web_probe_IP_address> cannot be updated to IP address <new_web_probe_address>. Update the entry manually.

Cause

Updating the web probe DNS record failed.

Solution

Update the DNS record manually according to the warning message instructions.

Troubleshooting General Issues

This topic contains troubleshooting information for general issues related to Remote Access.

GPO retrieval error

Error received—DirectAccess server GPO settings cannot be retrieved. Ensure you have edit permissions for the GPO.

The Remote Access management console is unresponsive after receiving this error.

Cause

DirectAccess cannot access the GPO of one of the entry points in the deployment and as a result the configuration fails to load.

Solution

Make sure that each entry point in the deployment has a corresponding GPO on its domain controller and verify that the logged on user has read and write permissions for all GPOs configured in the Remote Access deployment.

As a workaround, use the configuration cmdlets instead of using the Remote Access management console; for example, using `Get-RemoteAccess` and `Get-DAEntryPoint`.

Note

This scenario does not occur when the server GPO of the current entry point isn't available.

You can use the `Get-DAEntryPointDC` cmdlet to list all domain controllers that store server GPOs and `Get-DAMultiSite` in conjunction with `Get-RemoteAccess` to retrieve a complete list of the server GPOs in the deployment. For example:

```
$ServerGpos = Get-DAEntryPointDC | ForEach-Object {

   @{

        GpoName = (Get-RemoteAccess -EntryPoint $_.EntryPointName).ServerGpoName;

         DC = $_.DomainControllerName }

}

$ServerGpos | ForEach-Object { $GpoName = $_['GpoName'] ; $DC = $_['DC'] ; Write-Host

"Server GPO '$GpoName' on DC '$DC'" }
```

Windows 7 to Windows 8 client upgrade

Symptom—After a Windows 7 client upgrades to Windows 8 in a multisite deployment, the DirectAccess connection is not visible in the Networks list.

Cause

The Windows 7 GPOs in a multisite deployment do not contain the Windows 8 Network Connectivity Assistant configuration.

Windows 7 clients should use DirectAccess Connectivity Assistant to monitor their DirectAccess connectivity status which requires a separate manual configuration in the Windows 7 client GPOs. When Windows 7 clients are upgraded to Windows 8, the Network Connectivity Assistant will not function if the Windows 7 client GPO is still applied.

Solution

If DirectAccess Connectivity Assistant settings are configured in the Windows 7 GPOs, you can resolve this issue before upgrading client computers by modifying the Windows 7 GPOs using the following PowerShell cmdlets:

```
Set-GPRegistryValue -Name <Windows7GpoName> -Domain <DomainName> -Key
"HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\NetworkConnectivityAssistant" -
ValueName "TemporaryValue" -Type Dword -Value 1

Remove-GPRegistryValue -Name <Windows7GpoName> -Domain <DomainName> -Key
"HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\NetworkConnectivityAssistant"
```

If a client has already been upgraded or the DCA is not configured, move the client computer to the Windows 8 security group.

General cmdlet errors

- **Issue 1**

  Error received—Domain controller <domain_controller> cannot be reached for <server_name or entry_point_name>.

  Cause

  To maintain the configuration consistency in a multisite deployment, it is important to make sure that each GPO is managed by a single domain controller. When the domain controller that manages an entry point's server GPO is not available, Remote Access configuration settings cannot be read or modified.

  Solution

  Follow the procedure "To change the domain controller that manages server GPOs" described in 2.4. Configure GPOs.

- **Issue 2**

  Error received—The primary domain controller in domain <domain_name> cannot be reached.

  Cause

  To maintain the configuration consistency in a multisite deployment, it is important to make sure that each GPO is managed by a single domain controller. Client GPOs are managed on the primary domain controller. If the primary domain controller is not available, Remote Access configuration settings cannot be read or modified.

  Solution

  Follow the procedure "To transfer the PDC emulator role" described in 2.4. Configure GPOs.

Deploy Remote Access with OTP Authentication

Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. Remote Access can be deployed in a number of enterprise scenarios. This overview provides an introduction to the enterprise scenario for

deploying Windows Server 2012 DirectAccess with one-time password (OTP) user authentication.

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Scenario description

In this scenario a Remote Access server with DirectAccess enabled is configured to authenticate DirectAccess client users with two-factor OTP authentication, in addition to standard Active Directory credentials.

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- Deploy a Single Remote Access Server with Advanced Settings must be deployed before deploying OTP.
- Windows 7 Clients need to use DCA 2.0 to support OTP.
- OTP does not support PIN change.
- A Public Key Infrastructure must be deployed.

  For more information see: Test Lab Guide Mini-Module: Basic PKI for Windows Server 2012.
- Changing policies outside of the DirectAccess management console or PowerShell cmdlets is not supported.

In this scenario

The OTP authentication scenario includes a number of steps:

1. **Deploy a Single Remote Access Server with Advanced Settings**—A single Remote Access server must be deployed before configuring OTP. Planning and deploying a single server includes designing and configuring a network topology, planning and deploying certificates, setting up DNS and Active Directory, configuring Remote Access server settings, deploying DirectAccess clients, and preparing intranet servers.
2. **Plan Remote Access with OTP Authentication**—In addition to the planning required for a single server, OTP requires planning for a Microsoft certification authority (CA) and certificate templates for OTP; and a RADIUS-enabled OTP server. Planning might also include a requirement for security groups to exempt specific users from strong (OTP or smart card) authentication. For information regarding the configuration of OTP in a multi-forest environment, see Configure a Multi-Forest Deployment.
3. **Configure DirectAccess with OTP Authentication**—OTP deployment consists of a number of configuration steps, including preparing the infrastructure for OTP authentication, configuring the OTP server, configuring OTP settings on the Remote Access server, and updating DirectAccess client settings.
4. **Troubleshoot an OTP Deployment**—This troubleshooting section describes a number of the most common errors that can occur when deploying Remote Access with OTP authentication.

Practical applications

Increase security—Using OTP increases the security of your DirectAccess deployment. A user requires OTP credentials in order to gain access to the internal network. A user supplies OTP credentials via the Workplace Connections available in the network connections on the Windows 8 client computer, or using DirectAccess Connectivity Assistant (DCA) on client computers running Windows 7. The OTP authentication process works as follows:

1. The DirectAccess client enters domain credentials to access DirectAccess infrastructure servers (over the infrastructure tunnel). If no connection to the internal network is available, due to a specific IKE failure, Workplace Connection on the client computer notifies the user that credentials are required. On client computers running Windows 7, a pop-up requesting smart card credentials appears.
2. After the OTP credentials have been entered, they are sent over SSL to the Remote Access server, together with a request for a short-term smart card logon certificate.
3. The Remote Access server initiates validation of the OTP credentials with the RADIUS-based OTP server.
4. If successful, the Remote Access server signs the certificate request using its registration authority certificate, and sends it back to the DirectAccess client computer
5. The DirectAccess client computer forwards the signed certificate request to the CA and stores the enrolled certificate for use by the Kerberos SSP/AP.
6. Using this certificate the client computer transparently performs standard smart card Kerberos authentication.

Roles and features included in this scenario

The following table lists the roles and features required for the scenario:

Role/feature	How it supports this scenario
Remote Access Management role	The role is installed and uninstalled using the Server Manager console. This role encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are

Role/feature	How it supports this scenario
	managed in the legacy Routing and Remote Access console. The Remote Access role is dependent on the following server features: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server, utilize OTP authentication, and configure the default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI and Command Line Tools • Remote Access module for Windows PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

- A computer that meets the hardware requirements for Windows Server 2012.

- In order to test the scenario, at least one computer running Windows 8 or Windows 7 configured as a DirectAccess client is required.
- An OTP server that supports PAP over RADIUS.
- An OTP hardware or software token.

Software requirements

There are a number of requirements for this scenario:

1. Software requirements for single server deployment. For more information, see Deploy a Single Remote Access Server with Advanced Settings.
2. In addition to software requirements for a single server there are a number of OTP-specific requirements:
    a. CA for IPsec authentication—In an OTP deployment DirectAccess must be deployed using IPsec machines certificates issued by a CA. IPsec authentication using the Remote Access server as a Kerberos proxy is not supported in an OTP deployment. An internal CA is required.
    b. CA for OTP authentication—A Microsoft Enterprise CA (running on Windows 2003 Server or later) is required to issue the OTP client certificate. The same CA used for issuing certificates for IPsec authentication can be used. The CA server must be available over the first infrastructure tunnel.
    c. Security group—To exempt users from strong authentication, an Active Directory security group containing these users is required.
    d. Client-side requirements—For Windows 8 client computers, the Network Connectivity Assistant (NCA) service is used to detect whether OTP credentials are required. If they are the DirectAccess Media Manager prompts for credentials. NCA is included in the Windows 8 operating system, and no installation or deployment is required. For Windows 7 client computers DirectAccess Connectivity Assistant (DCA) 2.0 is required. This is available as a download on the Microsoft Download Center.
    e. Note the following:
        i. OTP authentication can be used in parallel with smart card and Trusted Platform Module (TPM)-based authentication. Enabling OTP authentication in the Remote Access Management console also enables use of smartcard authentication.
        ii. During Remote Access configuration users in a specified security group can be exempt from two-factor authentication, and thus authenticate with username/password only.
        iii. OTP new PIN and next tokencode modes are not supported
        iv. In a Remote Access multisite deployment, OTP settings are global and identify for all entry points. If multiple RADIUS or CA servers are configured for OTP, they are sorted by each Remote Access server according to availability and proximity.
        v. When configuring OTP in a Remote Access multi-forest environment, OTP CAs should be from the resource forest only, and certificate enrollment should be configured across forest trusts. For more information, see AD CS: Cross-forest Certificate Enrollment with Windows Server 2008 R2.

vi. Users who are using a KEY FOB OTP token should insert the PIN followed by the tokencode (without any separators) in the DirectAccess OTP dialog. Users who are using PIN PAD OTP token should insert only the tokencode in the dialog.

vii. When the WEBDAV is enabled then OTP should not be enabled.

Known Issues

The following are known issues when configuring an OTP scenario:

- Remote Access uses a probe mechanism to verify connectivity to RADIUS-based OTP servers. In some cases this might cause an error to be issued on the OTP server. To avoid this issue, do the following on the OTP server:
  - Create a user account that matches the username and password configured on the Remote Access server for the probe mechanism. The username should not define an Active Directory user.

    By default the username on the Remote Access server is DAProbeUser and the password is DAProbePass. These default settings can be modified using the following values in the registry on the Remote Access server:

    - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\DirectAccess\OTP\RadiusProbeUser
    - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\DirectAccess\OTP\RadiusProbePass
- If you change the IPsec root certificate in a configured and running DirectAccess deployment, OTP stops working. To resolve this issue, on each DirectAccess server, in a Windows PowerShell window, run the command: `iisreset`

Plan Remote Access with OTP Authentication

Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the planning steps required in order to deploy Windows Server 2012 Remote Access with one-time password (OTP) client authentication.

- Step 1: Plan an Advanced Single Server Deployment. This step includes planning for the infrastructure required to deploy a single server. It includes planning for network and server settings, certificate requirements, DNS settings, network location server deployment, DirectAccess management servers, Active Directory settings, and Group Policy objects (GPOs).
- Step 2: Plan the RADIUS Server Deployment.
- Step 3: Plan OTP Certificate Deployment.
- Step 4: Plan for OTP on the Remote Access Server

After you have completed these planning steps, see Configure DirectAccess with OTP Authentication. For information on configuring a multisite deployment as a proof of concept in a lab environment, see Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID.

Step 1: Plan an Advanced Single Server Deployment

The first step in planning a Windows Server® 2012 Remote Access with one-time password (OTP) client authentication is to plan and configure an advanced single server deployment.

Plan a single server deployment

Before you deploy Remote Access with OTP, make sure that you have completed all the steps to deploy a single Remote Access server. See Deploy a Single Remote Access Server with Advanced Settings. Then continue to Step 2: Plan the RADIUS Server Deployment.

Step 2: Plan the RADIUS Server Deployment

After deploying a single Remote Access server, plan for the one-time password (OTP) authentication server.

Task	Description
2.1 Plan the RADIUS server	For the OTP authentication server, Remote Access in Windows Server 2012 supports any RADIUS-enabled OTP server that supports the password authentication protocol (PAP).

2.1 Plan the RADIUS server

Note the following when planning a RADIUS server for OTP authentication:

- For most types of OTP deployments, you must configure the Remote Access server as a RADIUS agent. For more information, see the OTP vendor documentation.
- For all OTP deployments, you must synchronize your Active Directory users with your RADIUS server.
- The RADIUS server does not need to be a domain member.

- When you deploy the RADIUS server, you configure a shared secret and the port number for RADIUS traffic. Make a note of these details; they are required when you configure the Remote Access server.

You can view an example test lab guide that sets up OTP authentication with an RSA SecurID server in Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID.

See also

- Step 3: Plan OTP Certificate Deployment

Step 3: Plan OTP Certificate Deployment

After planning the RADIUS server, you must plan for certification authority (CA) requirements, including the CA that will issue one-time password (OTP) certificates, the OTP certificate template, and the registration authority certificate used by the Remote Access server to sign all DirectAccess client OTP certificate requests. These certificates are used as follows:

1. The DirectAccess client requests an OTP certificate, and the Remote Access server receives the request.
2. The Remote access server verifies the OTP credentials and if they are valid, the server acts as a registration authority, and signs the OTP certificate enrollment request using a short-lived signing certificate.
3. The Remote Access server sends the signed certificate enrollment request back to the DirectAccess client
4. The client then enrolls the OTP certificate from the CA using the certificate enrollment requests signed by the server.
5. The CA verifies the credentials and the request.

Task	Description
3.1 Plan the OTP CA	Plan the certification authority (CA) to use to issue certificates to DirectAccess clients for OTP authentication.
3.2 Plan the OTP certificate template	Plan the OTP certificate template.
3.3 Plan the registration authority certificate	Plan the registration authority certificate to sign all OTP authentication certificate requests.

3.1 Plan the OTP CA

To deploy DirectAccess using one-time password authentication (OTP), you require an internal CA to issue the OTP authentication certificates to DirectAccess client computers. For this

purpose, you can use the same internal CA that you use to issue the certificates that are used for regular IPsec computer authentication.

3.2 Plan the OTP certificate template

Each DirectAccess client requires an OTP authentication certificate in order to gain access to the internal network. You must configure a template on your internal CA for the OTP certificate. Note the following when configuring the OTP certificate template:

- All users who need to perform OTP authentication must have read and enroll permissions for this template.
- The subject name should be built from Active Directory information, to ensure that the subject name matches the OTP user name, and not the name of the Remote Access server that performs the certificate request. The subject name must be in the fully distinguished name format, and the subject alternative name must be in UPN format. This ensures that the enrolled OTP certificate is valid for Smartcard Kerberos authentication.
- The intended purpose of the certificate must be Smart Card Logon
- Issuance must require one authorized signature. The signature must be configured with the predefined DirectAccess OTP Application Policy set in the registration authority signing certificate template.
- The validity period should be set to one hour.

  Note
  In situations where the CA server is a Windows Server 2003 computer, then the template must be configured on a different computer. This is due to the fact that setting the **Validity period** in hours is not possible when running Windows versions prior to 2008/Vista. If the computer that you use to configure the template does not have the Certification Service role installed, or it is a client computer, then you may need to install the Certificate Templates snap-in. For more information on this subject click here.
- The renewal period should be set to 0.
- (Optional) Certificates and requests should not be stored in the CA database.
- The certificate Enhanced Key Usage parameter must be set correctly, as follows:
  - For the DirectAccess registration signing certificate template use the key 1.3.6.1.4.1.311.81.1.1.
  - For the OTP authentication certificate template use the key 1.3.6.1.4.1.311.20.2.2 key.

3.3 Plan the registration authority certificate

When DirectAccess clients request an OTP certificate, the Remote Access server receives the request from the client. The Remote Access server signs all OTP certificate requests from clients using the registration authority certificate. The CA issues certificates only if the request is signed by the registration authority certificate on the Remote Access server. The certificate must be issued by an internal CA, the certificate cannot be self-signed. It does not have to be issued by

the CA that issued the OTP certificates but the CA that issues the OTP certificates must trust the CA that issues the registration authority signing certificate.

See also

- Step 4: Plan for OTP on the Remote Access Server

Step 4: Plan for OTP on the Remote Access Server

After planning for the one-time password (OTP) RADIUS server and certificate settings, the final step in planning a Remote Access OTP deployment is to plan for client OTP settings on the Remote Access server.

Task	Description
4.1 Plan for OTP client exemptions	Plan for exemptions for users that you do not require to authentication using OTP.
4.2 Plan for Windows 7 clients	Plan to deploy the DirectAccess Connectivity Assistant (DCA) 2.0 to Windows 7 client computers.
4.3 Plan for smart cards	Plan for the use of smart cards for additional authorization.

4.1 Plan for OTP client exemptions

When OTP authentication is enabled, by default all users are required to authenticate using a combination of user name and password, and OTP credentials. However, you can allow selected users to authenticate using a user name and password only, without OTP. To do this, create a security group and add any users that you want to exempt from OTP authentication.

Note
Only client computers from a single forest may be exempted due to the fact that only one security group can be selected for client exemptions.

4.2 Plan for Windows 7 clients

By default, Windows 7 client computers cannot authenticate using OTP. Windows 7 client computers require DCA 2.0 to authenticate using OTP in a Windows Server 2012 Remote Access

deployment. For more information about DCA 2.0, see DirectAccess Connectivity Assistant 2.0 on the Microsoft Download Center.

4.3 Plan for smart cards

When OTP authentication is enabled, the option to enable the use of smart cards for additional authorization is available. Create a security group to allow temporary access in the event that a user's smart card is not functional.

See also

- Configure DirectAccess with OTP Authentication

Configure DirectAccess with OTP Authentication

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the configuration steps required to deploy a single Windows Server® 2012 Remote Access server with one-time password (OTP) authentication.

- Step 1: Implement a Single Server Remote Access Deployment. Install and configure a single Remote Access server. For instructions, see Deploy a Single Remote Access Server with Advanced Settings.
- Step 2: Configure the RADIUS Server.
- Step 3: Configure the Remote Access Server for OTP.
- Step 4: Verify DirectAccess with OTP.

Step 1: Implement a Single Server Remote Access Deployment

The first configuration step to deploy Remote Access with OTP is to implement an advanced single server deployment and then plan to add the RADIUS server.

Implement a single server deployment

Before you can configure Remote Access with OTP, you must configure an advanced single server Remote Access deployment as described in Deploy a Single Remote Access Server with Advanced Settings.

See also

- Step 2: Configure the RADIUS Server

Step 2: Configure the RADIUS Server

Before you configure the Remote Access server to support DirectAccess with OTP support, you configure the RADIUS server.

Task	Description
2.1 Configure the RADIUS software distribution tokens	On the RADIUS server configure software distribution tokens.
2.2 Configure the RADIUS security information	On the RADIUS server configure the ports and shared secret to be used.
2.3 Adding user account for OTP probing	On the RADIUS server create a new user account for OTP probing.
2.4 Synchronize with Active Directory	On the RADIUS server create user accounts synchronized with Active Directory accounts.
2.5 Configure the RADIUS authentication agent	Configure the Remote Access server as a RADIUS authentication agent.

2.1 Configure the RADIUS software distribution tokens

The RADIUS server must be configured with the necessary license and software and/or hardware distribution tokens to be used by DirectAccess with OTP. This process will be specific to each RADIUS vendor implementation.

2.2 Configure the RADIUS security information

The RADIUS server uses UDP ports for communication purposes, and each RADIUS vendor has its own default UDP ports for incoming and outgoing communication. For the RADIUS server to work with the Remote Access server, make sure that all firewalls in the environment are configured to allow UDP traffic between the DirectAccess and OTP servers over the required ports as needed.

The RADIUS server uses a shared secret for authentication purposes. Configure the RADIUS server with a strong password for the shared secret, and note that this will be used when

configuring the DirectAccess server's client computer configuration for use with DirectAccess with OTP.

2.3 Adding user account for OTP probing

On the RADIUS server create a new user account called **DAProbeUser** and give it the password **DAProbePass**.

2.4 Synchronize with Active Directory

The RADIUS server must have user accounts that correspond to the users in Active Directory that will be using DirectAccess with OTP.

To synchronize the RADIUS and Active Directory users

1. Record the user information from Active Directory for all DirectAccess with OTP users.
2. Use the vendor specific procedure to create identical user **domain\username** accounts in the RADIUS server that were recorded.

2.5 Configure the RADIUS authentication agent

The Remote Access server must be configured as a RADIUS authentication agent for the DirectAccess with OTP implementation. Follow the RADIUS vendor instructions to configure the Remote Access server as a RADIUS authentication agent.

Step 3: Configure the Remote Access Server for OTP

Once the RADIUS server has been configured with software distribution tokens, communication ports are open, a shared secret has been created, user accounts corresponding to Active Directory have been created on the RADIUS server, and the Remote Access server has been configured as a RADIUS authentication agent, then the Remote Access server needs to be configured to support OTP.

Task	Description
3.1 Exempt users from OTP authentication (optional)	If specific users will be exempted from DirectAccess with OTP authentication, then follow these preliminary steps.
3.2 Configure the Remote Access server to support OTP	On the Remote Access server update the Remote Access configuration to support OTP

Task	Description
	two-factor authentication.
3.3 Smart cards for additional authorization	Additional information regarding the use of smart cards.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

3.1 Exempt users from OTP authentication (optional)

If specific users are to be exempted from OTP authentication, then these steps must be taken prior to the Remote Access configuration:

Note
You must wait for replication between domains to complete when configuring the OTP exemption group.

Create user exemption security group

1. Create a security group in Active Directory for the purpose OTP exemption.
2. Add all users to be exempted from OTP authentication to the security group.

   Note
   Make sure to include only user accounts, and not computer accounts, in the OTP exemption security group.

3.2 Configure the Remote Access server to support OTP

To configure Remote Access to use two-factor authentication and OTP with the RADIUS server and certificate deployment from the previous sections, use the following steps:

Configure Remote Access for OTP

1. Open **Remote Access Management** and click **Configuration**.
2. In the **DirectAccess Setup** window, under **Step 2 – Remote Access Server**, click **Edit**.
3. Click **Next** three times, and in the **Authentication** section select both **Two factor authentication** and **Use OTP**, and ensure that **Use computer certificates** is checked.

   Note

After OTP has been enabled on the Remote Access server, if you disable OTP by deselecting **Use OTP**, the ISAPI and CGI extensions will be uninstalled on the server.

4. If Windows 7 support is required, select the **Enable Windows 7 client computers to connect via DirectAccess** check box. Note: As discussed in the planning section, Windows 7 clients must have DCA 2.0 installed to support DirectAccess with OTP.
5. Click **Next**.
6. In the **OTP RADIUS Server** section, double-click the blank **Server Name** field.
7. In the **Add a RADIUS Server** dialog, type the name of the RADIUS server in the **Server name** field. Click **Change** next to the **Shared secret** field, and type the same password that you used when configuring the RADIUS server in the **New secret** and **Confirm new secret** fields. Click **OK** twice, and click **Next**.

   Note
   If the RADIUS server is in a domain that is different than the Remote Access server, then the **Server Name** field must specify the FQDN of the RADIUS server.
8. In the **OTP CA Servers** section select the CA servers to be used for the enrollment of OTP client authentication certificates, and click **Add**. Click **Next**.
9. In the **OTP Certificate Templates** section click **Browse** to select the certificate template used for the enrollment of certificates that are issued for OTP authentication.

   Note
   The certificate template for OTP certificates issued by the corporate CA must be configured without the "Do not include revocation information in issued certificates" option. If this option is selected during the certificate template creation, then OTP client computers will fail to logon properly.

   Click **Browse** to select a certificate template used to enroll the certificate used by the Remote Access server to sign OTP certificate enrollment requests. Click **OK**. Click **Next**.
10. If exempting specific users from DirectAccess with OTP is required, then in the **OTP Exemptions** section select **Do not require users in the specified security group to authenticate using two-factor authentication**. Click **Security Group** and select the security group that was created for OTP exemptions.
11. On the **Remote Access Server Setup** page click **Finish**.
12. In the **DirectAccess Setup** window, under **Step 3 – Infrastructure Servers**, click **Edit**.
13. Click **Next** two times, and in the **Management** section double-click the **Management Servers** field.
14. Enter the **Computer name** or **Address** of the CA server that is configured to issue OTP certificates, and click **OK**.
15. In the **Remote Access Setup** windows click **Finish**.
16. Click **Finish** on the **DirectAccess Expert Wizard**.
17. On the **Remote Access Review** dialog box click **Apply**, wait for the DirectAccess policy

to be updated, and click **Close**.

18. On the **Start** screen, type **powershell.exe**, right-click **powershell**, click **Advanced**, and click **Run as administrator**. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
19. In the Windows PowerShell window, type **gpupdate /force** and press ENTER.

To configure Remote Access for OTP using PowerShell commands:

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To configure Remote Access to use two-factor authentication on a deployment that currently uses computer certificate authentication:

```
Set-DAServer -UserAuthentication TwoFactor
```

To configure Remote Access to use OTP authentication using the following settings:

- An OTP server named OTP.corp.contoso.com.
- A CA server named APP1.corp.contoso.com\corp-APP1-CA1.
- A certificate template named DAOTPLogon used for the enrollment of certificates that are issued for OTP authentication.
- A certificate template named DAOTPRA used to enroll the Registration Authority certificate used by the Remote Access server to sign OTP certificate enrollment requests.

```
Enable-DAOtpAuth -CertificateTemplateName 'DAOTPLogon' -OTPSignCertificateTemplateName
'DAOTPRA' -CAServer @('APP1.corp.contoso.com\corp-APP1-CA1') -RadiusServer
OTP.corp.contoso.com -SharedSecret Abcd123$
```

After executing the PowerShell commands complete steps 12-19 from the previous procedure to configure the Remote Access server to support OTP.

Note

Make sure to verify that you have applied the OTP settings on the Remote Access server before you add an entry point.

3.3 Smart cards for additional authorization

On the Authentication page of Step 2 in the Remote Access Setup Wizard, you can require the use of smart cards for access to the internal network. When this option is selected, the Remote Access Setup Wizard configures the IPsec connection security rule for the intranet tunnel on the DirectAccess server to require tunnel mode authorization with smart cards. Tunnel mode authorization allows you to specify that only authorized computers or users can establish an inbound tunnel.

To use smart cards with IPsec tunnel mode authorization for the intranet tunnel, you must deploy a public key infrastructure (PKI) with smart cards infrastructure.

Because your DirectAccess clients are using smart cards for access to the intranet, you can also use authentication mechanism assurance, a feature of Windows Server 2008 R2, to control access to resources, such as files, folders, and printers, based on whether the user logged on with a smart card-based certificate. Authentication mechanism assurance requires a domain functional level of Windows Server 2008 R2.

Allowing access for users with unusable smart cards

To allow temporary access for users with smart cards that are unusable, do the following:

1. Create an Active Directory security group to contain the user accounts of users who temporarily cannot use their smart cards.
2. For the DirectAccess server Group Policy Object, configure global IPsec settings for IPsec tunnel authorization and add the Active Directory security group to the list of authorized users.

To grant access to a user that cannot use their smart card, temporarily add their user account to the Active Directory security group. Remove the user account from the group when the smart card is usable.

Under the covers: Smart card authorization

Smart card authorization works by enabling tunnel mode authorization on the intranet tunnel connection security rule of the DirectAccess server for a specific Kerberos-based security identifier (SID). For smart card authorization, this is the well-known SID (S-1-5-65-1), which maps to smart card-based logons. This SID is present in a DirectAccess client's Kerberos token and is referred to as "This Organization Certificate" when configured in the global IPsec tunnel mode authorization settings.

When you enable smart card authorization in Step 2 of the DirectAccess Setup Wizard, the DirectAccess Setup Wizard configures the global IPsec tunnel mode authorization setting with this SID for the DirectAccess server Group Policy Object. To view this configuration in the Windows Firewall with Advanced Security snap-in for the DirectAccess server Group Policy Object, do the following:

1. Right click Windows Firewall with Advanced Security, and then click Properties.
2. On the IPsec Settings tab, in IPsec tunnel authorization, click Customize.
3. Click the Users tab. You should see the "NT AUTHORITY\This Organization Certificate" as an authorized user.

Step 4: Verify DirectAccess with OTP

This topic describes how to verify that you have correctly configured your DirectAccess with OTP deployment.

To verify OTP health on the Remote Access server

1. On the Remote Access server open the **Remote Access Management** console.
2. Under **Remote Access Servers** click the Remote Access server that has been configured for OTP support.
3. Click **Operations Status**.
4. Verify that the status of OTP displays the green icon and is Working.

> **Note**
> The health status update interval will be a maximum of the sum of the values from the registry key HKLM\SYSTEM\CCS\Services\Ramgmtsvc\parameters\HealthRefreshTimeout and the **Time interval for server activity** that was set in the Remote Access configuration.

To verify access to internal resources using OTP authentication

1. Connect a DirectAccess client computer to the corporate network and run **gpupdate /force** from the command prompt to obtain the group policy.
2. Disconnect the client computer from the corporate network, connect to the external network, and attempt to access internal resources. You should not have access to the internal resources.
3. In the case of a software token, access the OTP client token using the vendor instructions, and note the current token code. When a hardware token is used, follow the vendor instructions for authentication.
4. Click the **Network connections** icon in the notification area to access the DA Media Manager.
5. Click the **DirectAccess Connection**, and click **Continue**.
6. Enter the token code noted previously, and click **OK**. Wait for authentication to complete. The DirectAccess Workplace Connection status will now be **Connected**.
7. Attempt to access internal resources. You should be able to access all corporate resources.

Troubleshoot an OTP Deployment

This topic describes how to troubleshoot the most common errors that may occur when configuring a Remote Access deployment using OTP authentication.

For a list of known issues when deploying Remote Access with OTP authentication, see Known Issues.

- Troubleshooting Authentication Issues
- Troubleshooting Enabling OTP

Troubleshooting Authentication Issues

This topic contains troubleshooting information for issues related to problems users may have when attempting to connect to DirectAccess using OTP authentication. DirectAccerss OTP related events are logged on the client computer in Event Viewer under **Applications and Services Logs/Microsoft/Windows/OtpCredentialProvider**. Make sure that this log is enabled when troubleshooting issues with DirectAccess OTP.

Failed to access the CA that issues OTP certificates

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—OTP certificate enrollment for user <username> failed on CA server <CA_name>, request failed, possible reasons for failure: CA server name cannot be resolved, CA server cannot be accessed over the first DirectAccess tunnel or the connection to the CA server cannot be established.

Cause

The user provided a valid one-time password and the DirectAccess server signed the certificate request; however, the client computer cannot contact the CA that issues OTP certificates to finish the enrollment process.

Solution

On the DirectAccess server, run the following Windows PowerShell commands:

1. Get the list of configured OTP issuing CAs and check the value of ‘CAServer’: `Get-DAOtpAuthentication`
2. Make sure that the CAs are configured as a management servers: `Get-DAMgmtServer -Type All`
3. Make sure that the client computer has established the infrastructure tunnel: In the Windows Firewall with Advanced Security console, expand **Monitoring/Security Associations**, click **Main Mode**, and make sure that the IPsec security associations appear with the correct remote addresses for your DirectAccess configuration.

DirectAccess server connectivity issues

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)

One of the following errors:

- A connection cannot be established to Remote Access server <DirectAccess_server_hostname> using base path <OTP_authentication_path> and port <OTP_authentication_port>. Error code: <internal_error_code>.

- User credentials cannot be sent to Remote Access server <DirectAccess_server_hostname> using base path <OTP_authentication_path> and port <OTP_authentication_port>. Error code: <internal_error_code>.
- A response was not received from Remote Access server <DirectAccess_server_hostname> using base path <OTP_authentication_path> and port <OTP_authentication_port>. Error code: <internal_error_code>.

Cause

The client computer cannot access the DirectAccess server over the Internet, due to either network issues or to a misconfigured IIS server on the DirectAccess server.

Solution

Make sure that the Internet connection on the client computer is working, and make sure that the DirectAccess service is running and accessible over the Internet.

Failed to enroll for the DirectAccess OTP logon certificate

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—Certificate enrollment from CA <CA_name> failed. The request was not signed as expected by the OTP signing certificate, or the user does not have permission to enroll.

Cause

The one-time password provided by the user was correct, but the issuing certification authority (CA) refused to issue the OTP logon certificate. The certificate request may not be properly signed with the correct EKU (OTP registration authority application policy), or the user does not have the “Enroll” permission on the DA OTP template.

Solution

Make sure that DirectAccess OTP users have permission to enroll for the DirectAccess OTP logon certificate and that the proper “Application Policy” is included in the DA OTP registration authority signing template. Also make sure that the DirectAccess registration authority certificate on the Remote Access server is valid. See 3.2 Plan the OTP certificate template and 3.3 Plan the registration authority certificate.

Missing or invalid computer account certificate

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)— OTP authentication cannot be completed because the computer certificate required for OTP cannot be found in local machine certificate store.

Cause

DirectAccess OTP authentication requires a client computer certificate to establish an SSL connection with the DirectAccess server; however, the client computer certificate was not found or is not valid, for example, if the certificate expired.

Solution

Make sure that the computer certificate exists and is valid:

1. On the client computer, in the MMC certificates console, for the Local Computer account, open **Personal/Certificates**.
2. Make sure that there is a certificate issued that matches the computer name and double-click the certificate.
3. On the **Certificate** dialog box, on the **Certificate Path** tab, under **Certificate status**, make sure that it says “This certificate is OK.”

If a valid certificate is not found, delete the invalid certificate (if it exists) and re-enroll for the computer certificate by either running `gpupdate /Force` from an elevated command prompt or restarting the client computer.

Missing CA that issues OTP certificates

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—OTP authentication cannot be completed because the DA server did not return an address of an issuing CA.

Cause

Either there are no CAs that issue OTP certificates configured, or all of the configured CAs that issue OTP certificates are unresponsive.

Solution

1. Use the following command to get the list of CAs that issue OTP certificates (the CA name is shown in CAServer): `Get-DAOtpAuthentication`.
2. If no CAs are configured:
    a. Use either the command `Set-DAOtpAuthentication` or the Remote Access Management console to configure the CAs that issue the DirectAccess OTP logon certificate.
    b. Apply the new configuration and force the clients to refresh the DirectAccess GPO settings by running `gpupdate /Force` from an elevated command prompt or restarting the client machine.
3. If there are CAs configured, make sure they’re online and responding to enrollment requests.

Misconfigured DirectAccess server address

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—OTP authentication cannot complete as expected. The name or address of the Remote Access server cannot be determined. Error code: <error_code>. DirectAccess settings should be validated by the server administrator.

Cause

The address of the DirectAccess server is not configured properly.

Solution

Check the configured DirectAccess server address using `Get-DirectAccess` and correct the address if it is misconfigured.

Make sure the latest settings are deployed on the client computer by running `gpupdate /force` from an elevated command prompt or restart the client machine.

Failed to generate the OTP logon certificate request

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—The certificate request for OTP authentication cannot be initialized. Either a private key cannot be generated, or user <username> cannot access certificate template <OTP_template_name> on the domain controller.

Cause

There are two possible causes for this error:

- The user doesn’t have permission to read the OTP logon template.
- The user’s computer can’t access the domain controller because of network issues.

Solution

- Review the permissions setting on the OTP logon template and make sure that all users provisioned for DirectAccess OTP have ‘Read’ permission.
- Make sure that the domain controller is configured as a management server and that the client machine can reach the domain controller over the infrastructure tunnel. See 3.2 Plan the OTP certificate template.

No connection to the domain controller

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—A connection with the domain controller for the purpose of OTP authentication cannot be established. Error code: <error_code>.

Cause

There are two possible causes for this error:

- The user’s computer has no network connectivity.
- The domain controller isn’t accessible over the infrastructure tunnel.

Solution

- Make sure that the domain controller is configured as a management server by running the following command from a PowerShell prompt: `Get-DAMgmtServer -Type All`.
- Make sure that the client computer can reach the domain controller over the infrastructure tunnel.

OTP provider requires challenge/response

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—OTP authentication with Remote Access server (<DirectAccess_server_name>) for user (<username>) required a challenge from the user.

Cause

The OTP provider used requires the user to provide additional credentials in the form of a RADIUS challenge/response exchange, which is not supported by Windows Server 2012 DirectAccess OTP.

Solution

Configure the OTP provider to not require challenge/response in any scenario.

Incorrect OTP logon template used

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—The CA template from which user <username> requested a certificate is not configured to issue OTP certificates.

Cause

The DirectAccess OTP logon template was replaced and the client computer is attempting to authenticate using an older template.

Solution

Make sure the client computer is using the latest OTP configuration by performing one of the following:

- Force a Group Policy update by running the following command from an elevated command prompt: `gpupdate /Force`.
- Restart the client machine.

Missing OTP signing certificate

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (client event log)—An OTP signing certificate cannot be found. The OTP certificate enrollment request cannot be signed.

Cause

The DirectAccess OTP signing certificate cannot be found on the Remote Access server; therefore, the user certificate request can't be signed by the Remote Access server. Either there is no signing certificate, or the signing certificate has expired and was not renewed.

Solution

Perform these steps on the Remote Access server.

1. Check the configured OTP signing certificate template name by running the PowerShell cmdlet `Get-DAOtpAuthentication` and inspect the value of `SigningCertificateTemplateName`.
2. Use the Certificates MMC snap-in to make sure that a valid certificate enrolled from this template exists on the computer.
3. If no such certificate exists, delete the expired certificate (if one exists) and enroll for a new certificate based on this template.

To create the OTP signing certificate template see 3.3 Plan the registration authority certificate.

Missing or incorrect UPN/DN for the user

Scenario—User fails to authenticate using OTP with the error: "Authentication failed due to an internal error"

Error received (client event log)

One of the following errors:

- User <username> cannot be authenticated with OTP. Ensure that a UPN is defined for the user name in Active Directory. Error code: <error_code>.
- User <username> cannot be authenticated with OTP. Ensure that a DN is defined for the user name in Active Directory. Error code: <error_code>.

Error received (server event log)

The user name <username> specified for OTP authentication does not exist.

Cause

The user does not have the User Principal Name (UPN) or Distinguished Name (DN) attributes properly set in the user account, these properties are required for proper functioning of DirectAccess OTP.

Solution

Use the Active Directory Users and Computers console on the domain controller to verify that both of these attributes are properly set for the authenticating user.

OTP certificate is not trusted for login

Scenario—User fails to authenticate using OTP with the error: "Authentication failed due to an internal error"

Cause

The CA that issues OTP certificates is not in the enterprise NTAuth store; therefore, enrolled certificates can't be used for logon. This can occur in multi domain and multiforest environments where cross domain CA trust is not established.

Solution

Make sure that the certificate of the root of the CA hierarchy that issues OTP certificates is installed in the enterprise NTAuth Certificate store of the domain to which the user is attempting to authenticate.

Windows could not verify user credentials

Scenario—User fails to authenticate using OTP with the error: “Authentication failed due to an internal error”

Error received (Client computer)—Something went wrong while Windows was verifying your credentials. Try again, or ask your administrator for help.

Cause

The Kerberos authentication protocol does not work when the DirectAccess OTP logon certificate does not include a CRL. The DirectAccess OTP logon certificate does not include a CRL because either:

- The DirectAccess OTP logon template was configured with the option **Do not include revocation information in issued certificates**.
- The CA is configured not to publish CRLs.

Solution

1. To confirm the cause for this error, in the Remote Access Management console, in **Step 2 Remote Access Server**, click **Edit**, and then in the **Remote Access Server Setup** wizard, click **OTP Certificate Templates**. Make a note of the certificate template used for the enrollment of certificates that are issued for OTP authentication. Open the Certification Authority console, in the left pane, click **Certificate Templates**, double-click the OTP logon certificate to view the certificate template properties.

   To solve this issue, configure a certificate for the OTP logon certificate and do not select the **Do not include revocation information in issued certificates** check box on the **Server** tab of the template properties dialog box.
2. On the CA server, open the Certification Authority MMC, right click the issuing CA and click **Properties**. On the **Extensions** tab make sure that CRL publishing is correctly configured.

Troubleshooting Enabling OTP

This topic contains troubleshooting information for issues related to enabling DirectAccess OTP authentication using either the `Enable-DAOtpAuthentication` PowerShell cmdlet or the Remote Access Management console.

Failed to enroll the OTP signing certificate

Error received (server event log)—An OTP signing certificate cannot be enrolled using certificate template <OTP_signing_template_name>

Cause

There are three possible causes for this error:

- The template doesn't exist.
- The permissions set on the template do not allow the DirectAccess server to enroll.
- There is no network connectivity to the issuing certification authority (CA).

Solution

1. Make sure that the OTP signing certificate template with the given name:
   a. Exists and has the proper permissions.
   b. Is set to be issued by at least one CA that can issue certificates to the DirectAccess server.
2. If the template doesn't exist, create it as described in 3.3 Plan the registration authority certificate, or if another matching template exists reconfigure DirectAccess OTP with the new template name.

Failed to enable DirectAccess OTP when WebDAV is installed

Scenario—While attempting to apply the DirectAccess OTP configuration in the Remote Access Management console or by using the `Enable-DAOtpAuthentication` PowerShell cmdlet, the operation fails.

Error received (server event log)—DirectAccess OTP settings cannot be applied because the WebDAV IIS extension is running on the server. Remove WebDAV and apply the settings again.

Cause

The DirectAccess OTP service is incompatible with the WebDAV Publishing feature and cannot be enabled while WebDAV is installed.

Solution

Uninstall the WebDAV role:

1. In the Server Manager console, in the left-pane, click **IIS**.
2. In the main pane, scroll to **ROLES AND FEATURES**.
3. Right-click **WebDAV Publishing**, and then click **Remove Role or Feature**.
4. Complete the Remove Roles and Features Wizard.
5. Re-apply the DirectAccess OTP configuration.

No templates available in the Remote Access Management console

Scenario—While configuring OTP or registration authority certificate templates using the Remote Access Management console, some, or all of the templates are missing from the selection windows.

Cause

There are two possible causes for this error:

- The template is not configured according to the DirectAccess OTP requirements and so it cannot be selected.
- The selected CAs under **OTP CA Servers** are not configured to issue the required templates.

Solution

1. Make sure that the OTP logon template and the OTP signing certificate template are configured properly as described in 3.2 Plan the OTP certificate template and 3.3 Plan the registration authority certificate.
2. Make sure that the configured CAs in the **OTP CA Servers** list are configured to issues the relevant templates:
   a. On the CA server, open the Certification Authority console.
   b. In the left pane, expand the chosen CA server.
   c. Click **Certificate Templates** and make sure the required templates are enabled. If not, right-click **Certificate Templates**, click **New**, click **Certificate Template to issue**, and then select the templates you want to enable.

Cannot set renewal period of OTP template to 1 hour

Scenario—When configuring the DirectAccess OTP logon template using Windows 2003 CA, it is not possible to set the renewal period of the template to 1 hour.

Cause

The Certificate Templates MMC snap-in in Windows Server 2003 doesn't allow you to set the renewal period of a template to 1 hour.

Solution

Install Certificate Templates snap-in on a post-Windows Server 2003 server and use it to configure the OTP logon template, see Install the Certificate Templates Snap-In.

Deploy Remote Access in a Multi-Forest Environment

The Remote Access configuration tools (the Remote Access Management console and the Windows PowerShell cmdlets) are designed to work well in a single forest environment containing one or more domains. However, when Remote Access is deployed in a multi-forest environment, the Remote Access administrator must perform some manual configuration for a successful deployment. This guide describes the planning and configuration steps for a multi-forest deployment; including when one-time password (OTP) authentication is used.

- Plan a Multi-Forest Deployment
- Configure a Multi-Forest Deployment

Plan a Multi-Forest Deployment

This topic describes the planning steps required when configuring Remote Access in a multi-forest deployment.

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- Two-way trust is required.

Plan trust between forests

When you decide to enable access to resources from a new forest, allow clients from the new forest to use DirectAccess, or add Remote Access servers from the new forest as entry points to the Remote Access deployment, you must make sure that a full trust; that is, a two-way transitive trust, is configured between the two forests, see Trust types. Full trust between forests is a prerequisite for a multi-forest deployment to allow administrators to perform operations such as editing GPOs in the new forest, using security groups from new forest as the client security group, performing remote calls (WinRM, RPC) to computers in the new forest, and authenticating remote clients from the new forest.

Plan Remote Access administrator permissions

When you configure Remote Access it updates and sometimes creates GPOs in each of the domains that contain Remote Access servers or clients. In a multi-forest environment, as in single-forest environments, the Remote Access administrator must have permissions to write and modify DirectAccess GPOs and their security filters and optionally have permissions to create

links for the DirectAccess GPOs in all involved forests. These permissions are required regardless of the forest to which the Remote Access administrator belongs.

In addition, the Remote Access administrator must be a local administrator on all Remote Access servers including on the Remote Access servers from the new forest that are added as entry points to the original Remote Access deployment.

Plan client security groups

You must configure at least one security group in the new forest for DirectAccess client machines in the new forest. This is because a single security group cannot contain accounts from several forests.

Notes

- DirectAccess requires at least one Windows® 8 client security group for each forest. However, it is recommended to have one Windows 8 client security group for each domain that contains Windows 8 clients.
- When multisite is enabled, DirectAccess requires at least one Windows® 7 client security group per forest for each DirectAccess entry point in which Windows 7 client computers are supported. However, it is recommended to have a separate Windows 7 client security group for each entry point for each domain that contains Windows 7 clients.

  For DirectAccess to be applied on client computers in additional domains, client GPOs must be created in those domains. Adding security groups triggers writing new client GPOs for the new domains; therefore, if you add a new security group from a new domain to the list of DirectAccess client security groups, a client GPO will be automatically created on the new domain and client computers from the new domain will get the DirectAccess settings via the client GPO.

  Note that if you add a client from a new domain to an existing security group that is already configured as a DirectAccess client security group, the client GPO will not be created automatically by DirectAccess on the new domain. The client in the new domain will not receive the DirectAccess settings and will not be able to connect using DirectAcecss.

Plan certification authorities

If the DirectAccess deployment is configured to use One-Time Password (OTP) authentication, each forest contains the same signing certificate templates but with different Oid values. This results in the forests not being able to be configured as a single configuration unit. To resolve this issue and configure OTP in a multi-forest environment, see Configure OTP in a multi-forest deployment.

When using IPsec machine certificate authentication, all client and server computers must have a computer certificate issued by the same root or intermediate certification authority, regardless of the forest to which they belong.

Plan OTP exemptions

If you are using DirectAccess OTP authentication, note that the OTP exemption security group is limited to users of a single forest. This is because each security group can contain only users from a single forest and only one such security group can be configured.

Configure a Multi-Forest Deployment

This topic describes how to configure a Remote Access multi-forest deployment in several possible scenarios. All of the scenarios assume that DirectAccess is currently deployed on a single forest called Forest1, and that you are configuring DirectAccess to work with a new forest called Forest2.

- Access resources from Forest2
- Enable clients from Forest2 to connect via DirectAccess
- Add entry points from Forest2
- Configure OTP in a multi-forest deployment

Access resources from Forest2

In this scenario, DirectAccess is already deployed on Forest1, and is configured to allow clients from Forest1 to access the corporate network. By default, clients connected via DirectAccess can access only resources in Forest1 and cannot access any servers in Forest2.

To enable DirectAccess clients to access resources from Forest2

1. If the DNS suffix of Forest2 is not part of the DNS suffix of Forest1, add NRPT rules with the suffixes of the domains in Forest2, and optionally, add the suffixes of the domains in Forest2 to the DNS suffix search list. See Add NRPT rules and DNS suffixes.
2. Add the relevant internal IPv6 prefixes in Forest2 if IPv6 is deployed on the internal network. See Add internal IPv6 prefix.

Enable clients from Forest2 to connect via DirectAccess

In this scenario, you configure the Remote Access deployment to allow clients from Forest2 to access the corporate network. It is assumed that you have created the required security groups for client computers in Forest2. For information about the required security groups, see Plan client security groups.

To allow clients from Forest2 to access the corporate network

1. Add the security group of the clients from Forest2. See Add client security groups.

2. If the DNS suffix of Forest2 is not part of the DNS suffix of Forest1, add NRPT rules with the suffixes of the clients' domain in Forest2 to enable access to the domain controllers for authentication, and optionally, add the suffixes of the domains in Forest2 to the DNS suffix search list. See Add NRPT rules and DNS suffixes.
3. Add the internal IPv6 prefixes in Forest2 to enable DirectAccess to create the IPsec tunnel to the domain controllers for authentication. See Add internal IPv6 prefix.
4. Refresh the management servers list. See Refresh the management servers list.

Add entry points from Forest2

In this scenario, DirectAccess is deployed in a multisite configuration on Forest1, and you want to add a Remote Access server, named DA2, from Forest2 as an entry point to the existing DirectAccess multisite deployment.

To add a Remote Access server from Forest2 as an entry point

1. Make sure that the Remote Access administrator has sufficient permissions to write GPOs on the domain of DA2, and that the Remote Access administrator is a local administrator on DA2.
2. Add DA2 as an entry point. See 3.7. Add entry points to the multisite deployment.
3. Add NRPT rules with the suffixes of the domains in Forest2 to enable access to the domain controllers for authentication, and optionally, add the suffixes of the domains in Forest2 to the DNS suffix search list. See Add NRPT rules and DNS suffixes.
4. Add the relevant internal IPv6 prefixes in Forest2, if required, to enable Remote Access to establish the IPsec tunnel to the corporate resources and to make sure that NCSI probes work correctly. See Add internal IPv6 prefix.
5. Refresh the management servers list. See Refresh the management servers list.

Configure OTP in a multi-forest deployment

Note the following terms when configuring OTP in a multi-forest deployment:

- Root CA—The forest(s) main PKI tree CA.
- Enterprise CA—All other CAs.
- Resource Forest—The forest that contains the Root CA, and is considered to be the 'Managing forest\domain'.
- Account Forest—All other forests in the topology.

The PowerShell script, PKISync.ps1, is required for this procedure. See AD CS: PKISync.ps1 Script for Cross-forest Certificate Enrollment.

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Configure CAs as certificate publishers

1. Enable LDAP referral support on all of the enterprise CAs in all of the forests by running the following command from an elevated command prompt:

   ```
   certutil -setreg Policy\EditFlags +EDITF_ENABLELDAPREFERRALS
   ```

2. Add all enterprise CA computer accounts to the Active Directory Cert Publishers security groups in each of the account forests.
3. Restart all of the certsvc services on all of the CA computers in all of the forests by running the following command from an elevated command prompt:

   ```
   net stop certsvc && net start certsvc
   ```

4. Extract the root CA certificate by running the following command from an elevated command prompt:

   ```
   certutil -config <Computer-Name>\<Root-CA-Name> -ca.cert
   <root-ca-cert-filename.cer>
   ```

   (If you run the command on the root CA you can omit the connection information, -config <Computer-Name>\<Root-CA-Name>)

   a. Import the Root CA certificate from the previous step on the Account Forest CA by running the following command from an elevated command prompt:

      ```
      certutil -dspublish -f <root-ca-cert-filename.cer> RootCA
      ```

   b. Grant Resource Forest certificate templates Read/Write permissions to the <Account Forest>\<Administrator account>.

   c. Extract all resource forest enterprise CA certificates by running the following command from an elevated command prompt:

      ```
      certutil -config <Computer-Name>\<Enterprise-CA-Name> -ca.cert
      <enterprise-ca-cert-filename.cer>
      ```

      (If you run the command on the root CA you can omit the connection information, -config <Computer-Name>\<Root-CA-Name>)

   d. Import the Enterprise CA certificates from the previous step on the Account Forest CA by running the following commands from an elevated command prompt:

      ```
      certutil -dspublish -f <enterprise-ca-cert-filename.cer>
      NTAuthCA
      ```

      ```
      certutil -dspublish -f <enterprise-ca-cert-filename.cer> SubCA
      ```

   e. Remove account forest OTP certificate templates from the list of issued Certificate Templates.

Delete and import OTP certificate templates

1. Delete OTP certificate templates from the Account Forest; that is, Forest2.
2. Copy certificate templates and Oid objects from the Resource Forest to each of the Account Forests using the following PowerShell commands:

```
.\PKISync.ps1 -sourceforest <resource forest DNS> -
targetforest <account forest DNS> -type Template -cn <DA OTP
registration authority template common name>.
.\PKISync.ps1 -sourceforest <resource forest DNS> -
targetforest <account forest DNS> -type Template -cn <Secure
DA OTP logon certificate template common name>.
.\PKISync.ps1 -sourceforest <resource forest DNS> -
targetforest <account forest DNS> -type Oid –f
```

Publish OTP certificate templates

- Issue the newly imported certificate templates on all Account Forests CAs.

Extract and Synchronize the CA

1. Extract all Enterprise CA certificates from the Account Forests by running the following commands from an elevated command prompt:

   ```
   certutil -config <Computer-Name>\<Enterprise-CA-Name> -
   ca.cert <enterprise-ca-cert-filename.cer>
   ```

2. Synchronize CAs across forests from the Account Forests to the Resource Forest using the following PowerShell command:

   ```
   .\PKISync.ps1 -sourceforest <account forest DNS> -
   targetforest <resource forest DNS> -type CA -cn <enterprise
   CA sanitized name> –f
   ```

3. Synchronize CAs across forests from the Resource Forest to the Account Forests using the following PowerShell command:

   ```
   .\PKISync.ps1 -sourceforest <resource forest DNS> -
   targetforest <account forest DNS> -type CA -cn <enterprise CA
   sanitized name> –f
   ```

Configuration procedures

The following sections contain the configuration procedures for the above scenario deployments. After completing a procedure, return to the scenario to continue.

Add NRPT rules and DNS suffixes

Clients that connect via DirectAccess to the corporate network use the Name Resolution Policy Table (NRPT) to determine which DNS server should be used to resolve the address of different resources. This allows the client to resolve corporate resource addresses and helps the client maintain a proper inside-corporate/outside-corporate classification, which is required to keep DirectAccess working. The DirectAccess configuration tools automatically detect the root DNS

suffix of Forest1 and add it to the NRPT table. However, the FQDN suffixes of Forest2 are not added automatically to the NRPT table, and the Remote Access administrator must add them manually.

The DNS suffix search list allows the clients to use short label names instead of FQDNs. The Remote Access configuration tools automatically add all the domains in Forest1 to the DNS suffix search list. If you want to enable clients to use short label names for resources in Forest2, you need to add them manually.

To add a DNS suffix to the NRPT table and domain suffixes to the DNS suffix search list

1. In the middle pane of the Remote Access Management console, in the **Step 3 Infrastructure Servers** area, click **Edit**.
2. On the **Network Location Server** page, click **Next**.
3. On the **DNS** page, in the table, enter any additional name suffixes that are part of the corporate network in Forest 2. In **DNS Server Address**, enter the DNS server address manually, or by clicking **Detect**. If you don't enter the address, the new entries are applied as NRPT exemptions. Then click **Next**.
4. Optional: On the **DNS Suffix Search List** page, add any DNS suffixes by entering the suffix in the **New Suffix** box, and then clicking **Add**. Then click **Next**.
5. On the **Management** page, click **Finish**.
6. In middle pane of the Remote Access Management console, click **Finish**.
7. On the **Remote Access Review** dialog box, click **Apply**.
8. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.

Add internal IPv6 prefix

Note
Adding an internal IPv6 prefix is relevant only when IPv6 is deployed on the internal network.

Remote Access manages a list of IPv6 prefixes for corporate resources. Only resources with these IPv6 prefixes may be accessed by clients that are connected via DirectAccess. Because the Remote Access Management console and Windows PowerShell commands automatically add the IPv6 prefixes of Forest1, and might not add the prefixes of other forests, you must add any missing prefixes of Forest2 manually.

To add an IPv6 prefix

1. In the middle pane of the Remote Access Management console, in the **Step 2 Remote Access Server** area, click **Edit**.
2. In the Remote Access Server Setup wizard, click **Prefix Configuration**.
3. On the **Prefix Configuration** page, in **Internal network IPv6 prefixes**, add any additional IPv6 prefixes separated by semicolons, for example, 2001:db8:1::/64;2001:db8:2::/64. Then click **Next**.

4. On the **Authentication** page, click **Finish**.
5. In middle pane of the Remote Access Management console, click **Finish**.
6. On the **Remote Access Review** dialog box, click **Apply**.
7. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.

Add client security groups

To enable Windows 8 client computers from Forest2 to access resources through DirectAccess, you must add the security group from Forest2 to the Remote Access deployment.

To add Windows 8 client security groups

1. In the middle pane of the Remote Access Management console, in the **Step 1 Remote Clients** area, click **Edit**.
2. In the DirectAccess Client Setup wizard, click **Select Groups**, and then on the **Select Groups** page, click **Add**.
3. On the **Select Groups** dialog box, select the security groups containing DirectAccess client computers. Then click **Next**.
4. On the **Network Connectivity Assistant** page, click **Finish**.
5. In middle pane of the Remote Access Management console, click **Finish**.
6. On the **Remote Access Review** dialog box, click **Apply**.
7. On the **Applying Remote Access Setup Wizard Settings** dialog box, click **Close**.

To enable Windows 7 client computers from Forest2 to access resources through DirectAccess when multisite is enabled, you must add the security group from Forest2 to the Remote Access deployment for each entry point. For information about adding Windows 7 security groups, see the description of the **Client Support** page in 3.6. Enable the multisite deployment.

Refresh the management servers list

Remote Access automatically discovers the infrastructure servers in all the forests that contain DirectAccess configuration GPOs. If DirectAccess was deployed on a server from Forest1, the server GPO will be written to its domain in Forest1. If you enabled access to DirectAccess for clients from Forest2, the client GPO will be written to a domain in Forest2.

The automatic discovery process of infrastructure servers is required to allow access through DirectAccess to the domain controllers and System Center Configuration Manager. You must manually start the discovery process.

To refresh the management servers list

1. In the Remote Access Management console, click **Configuration**, and then in the **Tasks** pane, click **Refresh Management Servers**.
2. On the **Refreshing Management Servers** dialog box, click **Close**.

Deploy Remote Access in the Cloud

Scenario description

The Remote Access Hosted Cloud Scenario provides a secure way for enterprise to access resources in the private cloud over the Internet. This scenario is designed for enterprise organizations that want to take advantage of the services offered by placing some of their infrastructure in private, hosted clouds, while ensuring high availability.

Cross-premises connectivity to a hosted cloud enables enterprises to connect to private subnets in a hosted cloud network. It also enables connectivity between geographically separate enterprise locations. With cross-premises connectivity, enterprises can use their existing networking infrastructure to connect to hosting providers using the industry standard IKEv2-IPsec protocol.

In this scenario

The Hosted Cloud scenario includes several possible cross-premise configurations from which to choose.

Network and server topology—Decide which cross-premise architecture fits the needs of your organization:

- **Site-to-Site**—With this configuration, you can provide both network and site redundancy without the use of a network load balancing device.The Remote Access server is configured with two network adapters that are connected to two separate Internet service provider sites. Each of the ISP sites have two network adapters connected to the on-site Remote Access server. Subnets are hosted outside the customer premises. By creating site-to-site tunnels with the appropriate routes and you can fully optimize all the links and sites (assuming the load across all the subnets is same).
- **Network Load Balancing and Failover**—Using this configuration, you can provide both network and site redundancy without the use of a network load balancing device.The Remote Access server is configured with two network adapters that are connected to two separate Internet service provider sites. Each of the ISP sites has two network adapters connected to the on-site Remote Access server adapters. Multiple subnets are hosted outside the customer premises, and mirrored in the ISP's clouds. By creating site-to-site tunnels with the appropriate routes, you can fully optimize all the links and sites (assuming the load across all the subnets is same).

  The Remote Access server site-to-site gateway can be deployed using Windows Network Load Balancing (NLB), or third party load balancing devices. Network Adapter Teaming (also referred to as NIC teaming) is supported if multiple adapters are available and teaming is supported on the server.

Practical applications

The Remote Access Hosted Cloud Scenarios utilize a private cloud infrastructure using IKEv2 and IPSec to make enterprise resources exclusively available to employees or partners of your company. Authorized users are able to access the services on the internet, and if they are outside the company, via site-to-site Virtual Private Networks (VPN). When your resources are hosted by a third-party service provider (for example, hosters, outsourcers, and increasingly, telecommunications companies), the cloud is referred to as a hosted private cloud. Cloud Bursting and Data Recovery are two practical applications for site-to-site VPNs in the cloud

Cloud Bursting

Cloud bursting, is the function of using virtual private clouds from service providers to meet peak computing demand situations. When resources are stretched during peak times or seasons in the corporate data center, Cloud Bursting allows the excess data to be moved across the cloud to a provider network that can absorb the extra data capacity. The diagram below shows how sit-to-site VPNs are used to facilitate Cloud Bursting.

Disaster Recovery

Disaster Recovery uses virtual private clouds from service providers as backup infrastructure. As the Figure 1 below illustrates, a corporation can use the hoster's infrastructure with Hyper-V Replica and VPN, to replicate a mission critical application to its hoster. In the event of failure, everything, including the IP address infrastructure fails over to the hosted cloud. In this way clients traffic can seamlessly be routed to the service that's currently up and running (the hoster's failover site). Cross-Premises connectivity ensures secure connectivity to resources in the cloud.

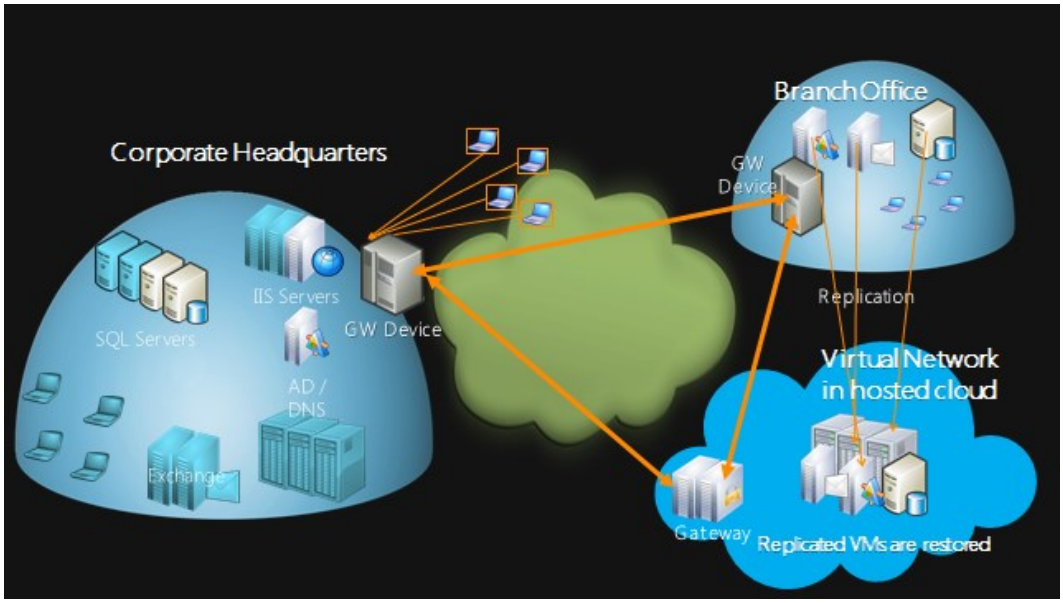

Figure 1: Disaster Recovery Topology

Roles and features included in this scenario

The following table lists the roles and features required for the scenario:

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the Server Manager console. This role encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. The role depends on the following: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server and default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools role	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: 1. Remote Access GUI and Command Line

Role/feature	How it supports this scenario
	Tools 2. Remote Access module for Windows PowerShell The role depends on the following: 1. Group Policy Management Console 2. RAS Connection Manager Administration Kit (CMAK) 3. Windows PowerShell 3.0 4. Graphical Management Tools and Infrastructure

Hardware requirements

The hardware requirements for this scenario will depend on which configuration you choose. Requirements will be specified in detail in the scenario planning guide.

Software requirements

The software requirements for this scenario will depend on which configuration you choose. Requirements will be specified in detail in the scenario planning guide.

See also

The following table provides links to additional information about Remote Access, including DirectAccess and VPN.

Content type	References
Product evaluation	Remote Access TechCenter | Remote Access test lab guides, when published
Planning	Links to the other Remote Access deployment scenarios when published.
Deployment	Links to the Remote Access deployment scenarios when published.
Tools and settings	Windows PowerShell cmdlets for Remote Access, when published.
Community resources	RRAS Product Team blog

Content type	References
Related technologies	IKEv2-IPsec

Test Lab Guide : Corp to Cloud: Configure an IKEv2-based Site-to-site VPN Connection

Test Lab Guides (TLGs) allow you to get hands-on experience with new products and technologies using a pre-defined and tested methodology that results in a working configuration. When you use a TLG to create a test lab, instructions tell you what servers to create, how to configure the operating systems and platform services, and how to install and configure any additional products or technologies. A TLG experience enables you to see all of the components and the configuration steps on both the front-end and back-end that go into a single- or multi-product or technology solution.

Depending on how you deploy your test lab environment, you can image the drives for the Windows Server 2012 Corp-to-cloud test lab if you are using physical computers or you can create snapshots of the Corp-to-cloud test lab virtual machines. This enables you to easily return to baseline configuration where most of the routine client, server, and networking services have already been configured so that you can focus on building out a test lab for the products or technologies of interest. For this reason, make sure that you perform a disk image on each computer if you're using physical computers, or perform virtual machine snapshots if you are using virtual machines after completing all the steps in this TLG.

About this guide

This document contains instructions for setting up the Windows Server 2012Corp-to-cloud test lab by deploying six server computers running Windows Server 2012 and one client computer running Windows® 8. The resulting configuration simulates a private intranet (contoso) connected to another private intranet (fabrikam) over the Internet.

Important

The following instructions are for configuring the Windows Server 2012 Corp-to-cloud test lab. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Warning

If you are able to work from a computer-based copy of this document during the lab exercises, and you are running virtual machines in Hyper-V, leverage the Hyper-V clipboard integration feature to paste commands. This will minimize potential errors with mistyped command strings.

1. Highlight and right-click a command from this document listed in **bold** text
2. Click **Copy**.
3. From the virtual machine menu bar, click **Clipboard**, and then click **Type clipboard text**.

Overview of the test lab scenario

The Windows Server® 2012 The Site-to-site VPN Corp-to-cloud test lab consists of the following:

Base Configuration test lab guide

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server.
- One intranet member server running Windows Server 2012 named APP1 that is configured as a general application and web server.
- One member client computer running Windows® 8 named CLIENT1 that will switch between Internet and intranet subnets.
- One intranet member server running Windows Server 2012 named EDGE1 that is configured as an Internet edge server.
- One standalone server running Windows Server 2012 named INET1 that is configured as an Internet DNS server, web server, and DHCP server.

The Windows Server 2012 Base Configuration test lab consists of two subnets that simulate the following:

- The Internet, referred to as the Internet subnet (131.107.0.0/24).
- An intranet, referred to as the Corpnet subnet (10.0.0.0/24), separated from the Internet subnet by EDGE1.

Computers on each subnet connect using a physical hub, switch, or virtual switch. See the following figure for the configuration of the Windows Server 2012 Base Configuration test lab.

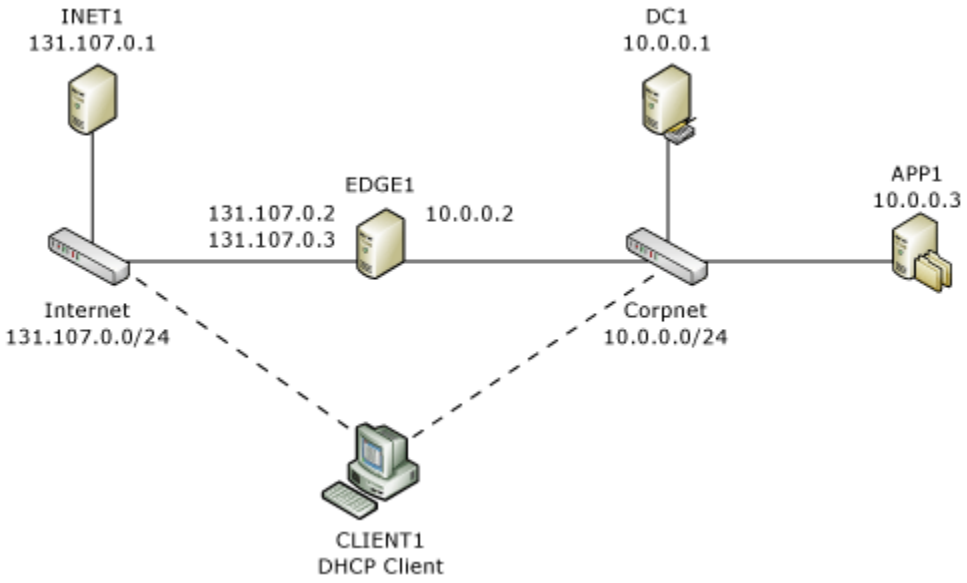

Corp-to-cloud test lab guide

Once the Base Configuration is in place, you will add the following three computers to create the cloud provider infrastructure:

1. One intranet member server running Windows Server 2012 named 3-EDGE1 that is configured as an Internet edge server.
2. One intranet member server running Windows Server 2012 named 3DC-1 that is configured as a domain controller.

The Windows Server 2012 The Corp-to-cloud test lab guide adds a subnet that simulates an intranet, referred to as Cloudnet subnet (10.6.0.0/24) connected by site to site VPN to Corpnet by EDGE1 and 3-EDGE1.

Computers on each subnet connect using a physical hub, switch, or virtual switch. See the following figure for the configuration of the Windows Server 2012 Corp-to-cloud test lab:

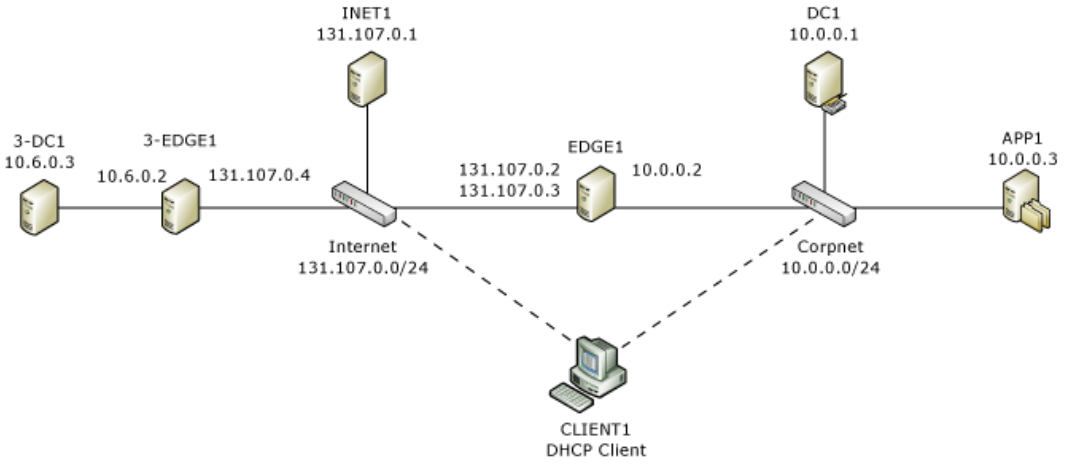

In addition, you will configure EDGE1 and 3-EDGE1 as site-to-site VPN servers to complete the Corp-to-Cloud test lab.

Configuration component requirements

The following components are required to configure the site-to-site VPN Cloud test lab:

- The product disc or files for Windows Server® 2012.
- Six computers or virtual machines that meet the minimum hardware requirements for Windows Server 2012; two of these computers have two network adapters installed.
- The product disc or files for of Windows® 8.
- One computer or virtual machine that meet the minimum hardware requirements for Windows 8.

Steps for configuring the test lab

There are 6 steps to setting up site to site VPN in the Cloud with Windows Server® 2012.

- STEP 1: Complete the base configuration
- STEP 2: Install and Configure 3-EDGE1
- STEP 3: Configure routing and forwarding between EDGE-1 and 3-EDGE1 using pre-shared key authentication
- STEP 4: Install a basic PKI certificate infrastructure
- STEP 5: Configure EDGE1 and 3-EDGE1 to use EAP for site-to-site authentication
- STEP 6: Configure EDGE1 and 3-EDGE1 to use certificates for site-to-site authentication

Note
You must be logged on as a member of the Domain Admins group or a member of the local Administrators group on each computer to complete the tasks described in this guide.

STEP 1: Complete the base configuration

STEP 1: Complete the Base Configuration

The first step is to complete all the steps in the **Test Lab Guide: Deploy the base configuration**. After completing the steps in that test lab guide you will have the core infrastructure required to complete the steps that follow.

STEP 2: Install and Configure 3-EDGE1

STEP 2: Install and Configure 3-EDGE1

- Install the operating system on 3-EDGE1
- Configure TCP/IP properties
- Install the Remote Access server role

3-EDGE1 must have two network adapters installed.

First, install Windows Server® 2012 as a standalone server.

To install the operating system on 3-EDGE1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying Windows Server 2012 Enterprise Edition (full installation) and a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect 3-EDGE1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.
4. Connect one network adapter to the cloudnet subnet and the other to the Internet subnet.

Next, configure the TCP/IP protocol with static IPv4 addresses on both network interfaces.

To configure TCP/IP properties

1. In **Server Manager**, click **Local Server**.
2. In the **Properties** pane, click **IPv4 address assigned by DHCP, IPv6 enabled**.
3. In **Network Connections**, right-click the network connection that is connected to the Corpnet subnet, click **Rename**, type **Corpnet**, and then press ENTER.
4. Right-click **Corpnet**, and then click **Properties**.
5. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
6. Click **Use the following IP address**. In **IP address**, type **10.6.0.2**. In **Subnet mask**, type **255.255.255.0**.
7. In the **Network Connections** window, right-click the network connection that is connected to the Internet subnet, click **Rename**, type **Internet**, and then press ENTER.
8. Right-click **Internet**, and then click **Properties**.
9. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
10. Click **Use the following IP address**. In **IP address**, type **131.107.0.4**. In **Subnet mask**, type **255.255.255.0**.
11. Click the **DNS** tab.
12. In **DNS suffix for this connection**, type **isp.example.com**, click **OK** twice, and then click **Close**.

13. Close the **Network Connections** window.
14. From the Start screen, type command, and then click Command Prompt.
15. To check network communication between 3-EDGE1 and EDGE1, type ping EDGE1.contoso.com in the command prompt window and hit ENTER.
16. Verify that there are four responses from 131.107.0.2.
17. Close the Command Prompt window.

The Remote Access server role in Windows Server 2012 combines the DirectAccess feature and the RRAS role service into a new unified server role. This new Remote Access server role allows for centralized administration, configuration, and monitoring of both DirectAccess and VPN-based remote access services. Use the following procedure to install the Remote Access role on 3-EDGE1.

To install the Remote Access role on 3-EDGE1

1. In the Server Manager console, in the **Dashboard**, click **add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. On the **Select server roles** dialog, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. On the **Role Services** screen select **Routing**, and then click **Next**
6. On the **Confirm installation selections** dialog, click **Install**.
7. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
ipmo servermanager

Install-WindowsFeature RemoteAccess -IncludeManagementTools

add-windowsFeature -name  routing -IncludeManagementTools
```

STEP 3: Configure routing and forwarding between EDGE-1 and 3-EDGE1 using pre-shared key authentication

---

STEP 3: Configure routing and forwarding between EDGE1 and 3-EDGE1 using EAP and (optionally) preshared key authentication

To create an IKE-V2-based site-to-site connection, perform the following tasks:

1. Configure site-to-site on the fabrikam (hoster) RRAS Server, 3-EDGE1.
2. Configure the demand-dial interface on 3-EDGE1.
3. Configure site-to-site on the contoso (corporate) RRAS Server, EDGE1.
4. Configure the demand-dial interface on EDGE1.
5. Confirm remote access policy configuration on EDGE1 and 3-EDGE1.
6. Initiate the site-to-site connection.
7. Test the site-to-site connection.

Configure site-to-site on the Corpnet router (3-EDGE1)

1. On **3-EDGE1**, click **Start**, and then click **Routing and Remote Access**.
2. In Routing and Remote Access, right-click **3-EDGE1 (local)** in the console tree, and then click **Configure and Enable Routing and Remote Access**.
3. The **Routing and Remote Access Server Setup Wizard** appears. Click **Next**.
4. On the **Configuration** page, select **Secure connection between two private networks. Connect this network to a remote network such as a branch office.**, and then click **Next**.
5. On the **Demand-Dial Connections** page, verify that **Yes** is checked and then click **Next**.
6. Click **Next**. On the **IP Address Assignment** page, select **Automatically**.
7. Click **Finish**

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
ipmo remoteaccess

install-remoteaccess -vpntype vpns2s
```

Configure the demand-dial interface on the fabrikam (hoster) router (3-EDGE1)

1. The Demand-Dial Interface Wizard will start. To complete the Demand-Dial Interface Wizard, click Next.
2. On the **Interface Name** page, type **3-EDGE1**. The interface name must match the user account name of the calling router.
3. Click **Next**. On the **Connection Type** page, select **Connect using virtual private networking (VPN)**.
4. Click **Next**. On the **VPN Type** page, select **IKEv2**.
5. Click **Next**. On the **Destination Address** page, type **131.107.0.2** in the **Host name or IP address** box.
6. Click **Next**. On the **Protocols and Security** page, do the following:
   a. Select **Route IP packets on this interface**.
   b. Select **Add a user account so a remote router can dial in**.
7. Click **Next**. On the **Static Routes for Remote Networks** page, click **Add**.
8. In the **Static Route** dialog box, do the following:
   a. In the **Destination** box, type **10.0.0.0**.
   b. In the **Network Mask** box, type **255.255.255.0**.
   c. In the **Metric** box type **1**.
9. Click **OK**. On the **Static Routes for Remote Networks** page, click **Next**.
10. Click **Next**. On the **Dial In Credentials** page, type **Pass@word1** for the **EDGE1** user account, and then retype **Pass@word1** in the **Confirm password** box. The **User name** box is pre-populated with the value **3-EDGE1**.

    Note
    The username for dial-In credentials should be the same as the name of the site to site interface if the peer is Windows server site-to-site VPN.
11. Click **Next**. On the **Dial Out Credentials** page, do the following:
    a. In the **User name** box, type **3-EDGE1**.
    b. In the **Domain** box, type **corp.contoso.com**.
    c. In the **Password** box, type **Pass@word1** for the fabrikam user account that will be created on 3-EDGE1.
    d. In the **Confirm password** box, retype the password for the **fabrikam** user account.

       Note
       The username for dial-out credentials should be the same as the name of the site to site interface if the peer is Windows server site-to-site VPN.
12. Click **Next**. On the last **Demand-Dial Interface Wizard** page, click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-VpnS2SInterface  EDGE1  131.107.0.2 -Protocol IKEv2 -AuthenticationMethod EAP -
UserName Contoso\3-EDGE1 -password Pass@word1 -IPv4Subnet 10.0.0.0/24:1
```

Optionally you can change the authentication method to preshared key. However, if you change the authentication method on 3-EDGE1, you will also have to change it correspondingly on EDGE1.

(Optional) Change the authentication method to preshared key

1. In the Routing and Remote Access snap-in, expand **3-EDGE1**, and then click **Network Interfaces**.
2. Right-click **Remote Router** and then click **Properties**.
3. Click the **Security** tab and then select **Use preshared key for authentication**.
4. Type **abc** in **Key:**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-VpnS2SInterface 3-EDGE1  -Protocol IKEv2 –AuthenticationMethod –SharedSecret “abc”
```

Configure site-to-site on the VPN-Corpnet router (EDGE1)

1. On **EDGE1**, click **Start**, and then click **Routing and Remote Access**.
2. In Routing and Remote Access, right-click **EDGE1 (local)** in the console tree, and then click **Configure and Enable Routing and Remote Access**.
3. The Routing and Remote Access Server Setup Wizard appears. Click **Next**.
4. On the **Configuration** page, select **Secure connection between two private networks. Connect this network to a remote network such as a branch office.**, and then click **Next**.
5. On the **Demand-Dial Connections** page, verify that **Yes** is checked and then click **Next**.
6. Click **Next**. On the **IP Address Assignment** page, select **Automatically**.
7. Click **Finish**

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
ipmo servermanager

add-windowsFeature -name  routing -IncludeManagementTools

ipmo remoteaccess

install-remoteaccess -vpntype vpns2s –IPv6Prefix 2001:db8:1:200::/64 -
```

Configure the demand-dial interface on the VPN-Corpnet router (EDGE1)

1. The Demand-Dial Interface Wizard will start. To complete the Demand-Dial Interface Wizard, click Next.
2. On the **Interface Name** page, type **EDGE1**. The interface name must match the user account name of the calling router.
3. Click **Next**. On the **Connection Type** page, select **Connect using virtual private networking (VPN)**.
4. Click **Next**. On the **VPN Type** page, select **IKEv2**.
5. Click **Next**. On the **Destination Address** page, type **131.107.0.4** in the **Host name or IP address** box.
6. Click **Next**. On the **Protocols and Security** page, do the following:
   a. Select **Route IP packets on this interface**.
   b. Select **Add a user account so a remote router can dial in**.
7. Click **Next**. On the **Static Routes for Remote Networks** page, click **Add**.
8. In the **Static Route** dialog box, do the following:
   a. In the **Destination** box, type **192.168.1.0**.
   b. In the **Network Mask** box, type **255.255.255.0**.
   c. In the **Metric** box type **1**.
9. Click **OK**. On the **Static Routes for Remote Networks** page, click **Next**.
10. Click **Next**. On the **Dial In Credentials** page, type **Pass@word1** for the **EDGE1** user account, and then retype **Pass@word1** in the **Confirm password** box. The **User name** box is pre-populated with the value **EDGE1**.
11. Click **Next**. On the **Dial Out Credentials** page, do the following:
    a. In the **User name** box, type **3-EDGE1**.
    b. In the **Domain** box, type **corp.fabrikam.com**.
    c. In the **Password** box, type **Pass@word1** for the VPN_Corpnet user account that will be created on 3-EDGE1.
    d. In the **Confirm password** box, retype the password for the VPN_Corpnet user account.
12. Click **Next**. On the last **Demand-Dial Interface Wizard** page, click **Finish**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-VpnS2SInterface  3-EDGE1 131.107.0.4 -Protocol IKEv2 -AuthenticationMethod EAP -
UsecName Fabrikam\EDGE1 -password Pass@word1 -IPv4Subnet 192.168.1.0 /24:1
```

Optionally you can change the authentication method to preshared key. However, if you change the authentication method on EDGE1, you should have also changed it correspondingly on 3-EDGE1.

(Optional) Change the authentication method to preshared key

1. In the Routing and Remote Access snap-in, expand **3-EDGE1**, and then click **Network Interfaces**.
2. Right-click **Remote Router** and then click **Properties**.
3. Click the **Security** tab and then select **Use preshared key for authentication**.
4. Type **abc** in **Key:**, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-VpnS2SInterface  EDGE1  -Protocol IKEv2 –AuthenticationMethod –SharedSecret “abc”
```

Initiate the VPN connection by performing the following steps on 3-EDGE1.

Initiate the site-to-site connection

1. On **3-EDGE1**, in the console tree of the Routing and Remote Access snap-in, click **Network Interfaces**.
2. In the details pane, right-click **contoso**, and then click **Connect**.
3. Confirm that the connection state of **VPN_Corpnet** is connected.

Test the site-to-site connection

1. On **DC1** at the **Start** menu, type **cmd** and then hit **Enter**. At the command prompt, type **ping 10.6.0.2**. Verify that there are four replies from 10.6.0.2

STEP 4: Install a basic PKI certificate infrastructure

Step 4: Install a basic PKI certificate infrastructure for site-to-site certificate authentication method.

Of the three authentication types available for site-to-site connections, EAP, preshared key and certificates, use of certificates is recommended. In the next step you will change the site-to-site authentication method to certificates.

There are three steps to adding a certificate infrastructure.

1. Install and configure 3-DC1 as a domain controller.
2. Install an enterprise root certification authority (CA) on 3-DC1.
3. Enable computer certificate auto-enrollment for the corp.fabrikam.com domain.
4. Configure the client-server authentication template.
5. Join 3-APP1 and 3-EDGE1 to the corp.fabrikam.com domain

DC1 provides the following services:

- A domain controller for the corp.fabrikam.com Active Directory Domain Services (AD DS) domain
- A DNS server for the corp.fabrikam.com DNS domain

DC1 configuration consists of the following:

- Install the operating system
- Configure TCP/IP
- Install active Directory and DNS
- Create a user account in Active Directory

First, install Windows Server 2012 as a standalone server.

To install the operating system on 3-DC1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying Windows Server 2012 (full installation) and a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect 3-DC1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.
4. Connect DC1 to the Corpnet subnet.

Next, configure the TCP/IP protocol with a static IP address of 10.6.0.3 and the subnet mask of 255.255.255.0.

To configure TCP/IP on 3-DC1

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection**.

   Note
   The link may not immediately appear. Wait for the network interfaces to be enumerated.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In IP address, type **10..6.0.3**. In **Subnet mask**, type **255.255.255.0**. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **127.0.0.1**.
5. Click **OK** and then close the **Wired Ethernet Properties** dialog.
6. Close the **Network Connections** window.
7. In Server Manager, click **Local Server** in the console tree. Click the link next to **Computer name** in the Properties tile.
8. On the **Computer Name** tab of the **System Properties** dialog, click **Change**.
9. In **Computer name**, type **3-DC1**, click **OK** twice, and then click **Close**. When you are prompted to restart the computer, click **Restart Now**.
10. After restarting, login using the local Administrator account.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.6.0.3 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
127.0.0.1

Rename-Computer 3-DC1

Restart-Computer
```

Next, configure 3-DC1 as a domain controller and DNS server for the corp.fabrikam.com domain.

To configure 3-DC1 as a domain controller and DNS server

1. Launch **Server Manager**.
2. On the **Dashboard** screen, under **Configure this local server**, click **Add roles and features**.

3. Click **Next** three times to get to the server role selection screen.
4. In the **Select Server Roles** dialog, select **Active Directory Domain Services**. Click **Add Features** when prompted, and then click **Next**.
5. In the **Select features** dialog, click **Next**.
6. In the **Active Directory Domain Services** dialog, click **Next**.
7. In the **Confirm installation selections** dialog, click **Install**. Wait for the installation to complete.
8. In the **Installation Progress** dialog, click the **Promote this server to a Domain Controller** link.

   Note
   If you close the "Installation Progress" dialog before it presents the promotion link, click the gray Tasks flag in the upper right section of **Server Manager**. When the installation is complete you will see the **Promote this server to a Domain Controller** link.
9. In the **Deployment Configuration** dialog, select **Add a new forest**. In the **Root domain name** field, type **corp.fabrikam.com**. Click **Next**.
10. In the **Domain Controller Options** dialog, leave the default values, specify a strong DSRM password twice, and then click **Next** four times to accept default settings for DNS, NetBIOS, and directory paths.
11. In the **Review Options** dialog, review your selections and then click **Next**.

    Note
    You can also click the **View script** button to review and save the PowerShell commands that Server Manager will run during DC Promotion.
12. In the **Prerequisites Check** dialog, allow the validation to complete and verify that no errors are reported. Since this is the first DNS server deployment in the forest, you can safely ignore all warnings regarding DNS delegation. Click **Install** to start the domain controller promotion. Allow the installation to complete.
13. Allow the domain controller to restart. After the server restarts, logon using the CORP\Administrator credentials.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints..

```
Install-WindowsFeature AD-Domain-Services -IncludeManagementTools

Install-ADDSForest -DomainName corp.fabrikam.com
```

Note
Windows PowerShell in Windows Server 2012 implements dynamic module loading. Using the **Import-Module** cmdlet is no longer required; instead, simply invoking the

cmdlet, alias, or function automatically loads the module. To see loaded modules, use the **Get-Module** cmdlet.

Next, create a user account in Active Directory that will be used when logging in to CORP domain member computers.

To create a user account in Active Directory

1. From the Start screen, click **Active Directory Administrative Center**.
2. In the console tree, click the arrow to expand **corp (local)**, and then double-click **Users**. This adds Users as a recent navigation link in the console tree.
3. In the **Tasks** pane, click **New**, and then click **User**.On the **New Object - User** dialog box, in **Full name**, type **EDGE1**, and in **User logon name**, type **EDGE1**, and then click **Next**.
4. In the **Create User** dialog, type **EDGE1** next to **Full name** and type **EDGE1** next to **User SamAccountName logon: corp**.
5. In **Password**, type the password that you want to use for this account, and in **Confirm password**, type the password again.
6. Under **Password options**, select **Other password options**, and select **Password never expires**.
7. Scroll down to access the **Member of** section of the **Create User** dialog, and click **Add**. Type **Domain Admins; Enterprise Admins**, and then click **OK**.
8. Click **OK** to close the **Create User** dialog.
9. Exit the Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-ADUser -SamAccountName EDGE1 -AccountPassword (read-host "Set user password" -
assecurestring) -name "EDGE1" -enabled $true -PasswordNeverExpires $true -
ChangePasswordAtLogon $false


Add-ADPrincipalGroupMembership -Identity "CN=EDGE1,CN=Users,DC=corp,DC=fabrikam,DC=com" -
MemberOf "CN=Enterprise Admins,CN=Users,DC=corp,DC=contoso,DC=com","CN=Domain
Admins,CN=Users,DC=corp,DC=fabrikam,DC=com"
```

To install the Certification Services server role on 3-DC1

1. On the **Server Manager Dashboard** screen, under **Configure this local server**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.

3. On the **Select Server Roles** page, select **Active Directory Certificate Services**, click **Add Features** when prompted, and then click **Next**.
4. Click **Next** three times to accept the default settings, and then click **Install**.
5. Wait for the installation to complete.
6. In the **Installation Progress** dialog, click the **Configure Active Directory Certificate Services** on the destination server link.
7. On the **Credentials** screen, click **Next**.
8. On the **Role Services** page, select **Certification Authority**, and click **Next**.
9. Click **Next** seven times to accept the default configuration options for Enterprise Root CA.
10. On the confirmation screen, click **Configure**.
11. Verify that configuration succeeded, and click **Close**.
12. Click **Close** in the **Add Roles and Features Wizard**.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature AD-Certificate -IncludeManagementTools

Install-AdcsCertificationAuthority -Force
```

Next, configure Group Policy on 3-DC1 so that domain members automatically request computer certificates.

To enable computer certificate auto-enrollment

1. On 3-DC1, from the **Start** screen, click **Group Policy Management**.
2. In the console tree, open **Forest: corp.contoso.com\Domains\corp.contoso.com**.
3. In the console tree, right-click **Default Domain Policy**, and then click **Edit**.
4. In the console tree of the **Group Policy Management Editor**, open **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.
5. In the details pane, double-click **Certificate Services Client – Auto-Enrollment**. In **Configuration Model**, select **Enabled**.
6. Select **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**. Click **OK**.

Next, configure a custom client-server authentication template.

To configure the client-server authentication template

1. On **3-DC1**, from the **Start** screen, click **Certification Authority**.
2. In the details pane, expand **corp-3-DC1-CA**.
3. Right-click **Certificate Templates** and select **Manage**.
4. In the Certificate Templates console, right-click **Workstation Authentication** and click

Duplicate Template.

5. On the **General** tab, change the **Template display name** to **Client-Server Authentication** and select **Publish certificate in Active Directory**.
6. On the **Extensions** tab, click **Application Policies** and then click **Edit**. Click **Add**, and then select **Server Authentication**. Click **OK**. Click **Add** and select **IP security IKE intermediate**. Click **OK** twice to return to the **Properties of New Template** dialog..
7. Click the **Security** tab. For **Domain Computers**, select the checkbox to Allow **Autoenroll**. Click **OK**. Close the **Certificate Templates Console**.
8. In the **Certification Authority** snap-in console tree, right-click **Certificate Templates** and select **New**, and then **Certificate Template to Issue**.On the **Extensions** tab, **click Application Policie**s and then click **Edit**. Click **Add,** and then select **Server Authentication**. Click **OK** twice to return to the **Properties of New Template** dialog.
9. Select **Client-Server Authentication** and then click **OK**.
10. Close the **Certification Authority console**.

Next, join 3-EDGE1 to the corp.fabrikam.com domain.

To join 3-EDGE to the corp.fabrikam.com domain

1. In **Server Manager**, click **Local Server** in the console tree. Click the link next to **Computer name** in the **Properties** tile.
2. In the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In Computer Name, type **APP1**. Under **Member of**, click **Domain**, and then type **corp.contoso.com**.
4. Click **OK**.
5. When you are prompted for a user name and password, type **User1** and its password, and then click **OK**.
6. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
7. When you are prompted that you must restart the computer, click **OK**.
8. On the **System Properties** dialog box, click **Close**.
9. When you are prompted to restart the computer, click **Restart Now**.
10. After the computer restarts, click the **Switch User arrow icon**, then click **Other User** and log on to the CORP domain with the **User1** account.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName 3-EDGE1

Restart-Computer

Add-Computer -DomainName corp.fabrikam.com
```

```
Restart-Computer
```

To obtain an additional certificate for 3-EDGE1

1. On 3-EDGE1, click **Start**, type **mmc** and then click **mmc**. Click **Yes** at the **User Account Control** prompt.
2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the **Certificates** snap-in, open **Certificates (Local Computer)\Personal\Certificates**.Click **Next** three times to accept the default settings, and then click **Install**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice
7. On the **Request Certificates** page, click **Client Server Authentication** and then click **More information is required to enroll for this certificate**.
8. On the **Subject** tab of the **Certificate Properties** dialog box, in **Subject name**, for **Type**, select **Common Name**.
9. In **Value**, type **3-EDGE1.corp.fabrikam.com**, and then click **Add**.
10. Click **OK**, click **Enroll**, and then click **Finish**.
11. In the **Details** pane of the **Certificates** snap-in, verify that a new certificate with the name **3-EDGE1.corp.fabrikam.com** was enrolled with **Intended Purposes of Client Authentication, Server Authentication, IP security IKE intermediate**.
12. Close the console window. If you are prompted to save settings, click **No**.

To install the Certification Services server role on DC1

1. On the **Server Manager Dashboard** screen, under **Configure this local server**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. On the **Select Server Roles** page, select **Active Directory Certificate Services**, click **Add Features** when prompted, and then click **Next**.
4. Click **Next** three times to accept the default settings, and then click **Install**.
5. Wait for the installation to complete.
6. In the **Installation Progress** dialog, click the **Configure Active Directory Certificate Services** on the destination server link.
7. On the **Credentials** screen, click **Next**.
8. On the **Role Services** page, select **Certification Authority**, and click **Next**.
9. Click **Next** seven times to accept the default configuration options for Enterprise Root CA.
10. On the confirmation screen, click **Configure**.

11. Verify that configuration succeeded, and click **Close**.
12. Click **Close** in the **Add Roles and Features Wizard**.

Next, configure Group Policy on DC1 so that domain members automatically request computer certificates.

To enable computer certificate auto-enrollment

1. On DC1, from the **Start** screen, click **Group Policy Management**.
2. In the console tree, open **Forest: corp.contoso.com\Domains\corp.contoso.com**.
3. In the console tree, right-click **Default Domain Policy**, and then click **Edit**.
4. In the console tree of the **Group Policy Management Editor**, open **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.
5. In the details pane, double-click **Certificate Services Client – Auto-Enrollment**. In **Configuration Model**, select **Enabled**.
6. Select **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**. Click **OK**.

Next, configure a custom client-server authentication template.

To configure the client-server authentication template

1. On **DC1**, from the **Start** screen, click **Certification Authority**.
2. In the details pane, expand **corp-DC1-CA**.
3. Right-click **Certificate Templates** and select **Manage**.
4. In the Certificate Templates console, right-click **Workstation Authentication** and click **Duplicate Template**.
5. On the **General** tab, change the **Template display name** to **Client-Server Authentication** and select **Publish certificate in Active Directory**.
6. On the **Extensions** tab, click **Application Policies** and then click **Edit**. Click **Add**, and then select **Server Authentication**. Click **OK**. Click **Add** and select **IP security IKE intermediate**. Click **OK** twice to return to the **Properties of New Template** dialog..
7. Click the **Security** tab. For **Domain Computers**, select the checkbox to Allow **Autoenroll**. Click **OK**. Close the **Certificate Templates Console**.
8. In the **Certification Authority** snap-in console tree, right-click **Certificate Templates** and select **New**, and then **Certificate Template to Issue**.On the **Extensions** tab, **click Application Policie**s and then click **Edit**. Click **Add,** and then select **Server Authentication**. Click **OK** twice to return to the **Properties of New Template** dialog.
9. Select **Client-Server Authentication** and then click **OK**.
10. Close the **Certification Authority console**.

Next, join 3-EDGE1 to the corp.fabrikam.com domain.

To join 3-EDGE to the corp.fabrikam.com domain

1. In **Server Manager**, click **Local Server** in the console tree. Click the link next to **Computer name** in the **Properties** tile.
2. In the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In Computer Name, type **APP1**. Under **Member of**, click **Domain**, and then type **corp.contoso.com**.
4. Click **OK**.
5. When you are prompted for a user name and password, type **User1** and its password, and then click **OK**.
6. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
7. When you are prompted that you must restart the computer, click **OK**.
8. On the **System Properties** dialog box, click **Close**.
9. When you are prompted to restart the computer, click **Restart Now**.
10. After the computer restarts, click the **Switch User arrow icon**, then click **Other User** and log on to the CORP domain with the **User1** account.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName 3-EDGE1

Restart-Computer

Add-Computer -DomainName corp.fabrikam.com

Restart-Computer
```

Next, create a user account in Active Directory that will be used when logging in to CORP domain member computers.

To create a user account in Active Directory of DC1

1. From the Start screen, click **Active Directory Administrative Center**.
2. In the console tree, click the arrow to expand **corp (local)**, and then double-click **Users**. This adds Users as a recent navigation link in the console tree.
3. In the **Tasks** pane, click **New**, and then click **User**.On the **New Object - User** dialog box, in **Full name**, type **3-EDGE1**, and in **User logon name**, type **EDGE1**, and then click **Next**.
4. In the **Create User** dialog, type **3-EDGE1** next to **Full name** and type **3-EDGE1** next to **User SamAccountName logon: corp**.
5. In **Password**, type the password that you want to use for this account, and in **Confirm password**, type the password again.

6. Under **Password options**, select **Other password options**, and select **Password never expires**.
7. Scroll down to access the **Member of** section of the **Create User** dialog, and click **Add**. Type **Domain Admins; Enterprise Admins**, and then click **OK**.
8. Click **OK** to close the **Create User** dialog.
9. Exit the Active Directory Administrative Center.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-ADUser -SamAccountName EDGE1 -AccountPassword (read-host "Set user password" -
assecurestring) -name "3-EDGE1" -enabled $true -PasswordNeverExpires $true -
ChangePasswordAtLogon $false


Add-ADPrincipalGroupMembership -Identity "CN=3-EDGE1,CN=Users,DC=corp,DC=contoso,DC=com"
-MemberOf "CN=Enterprise Admins,CN=Users,DC=corp,DC=contoso,DC=com","CN=Domain
Admins,CN=Users,DC=corp,DC=contoso,DC=com"
```

To obtain an additional certificate for EDGE1

1. On EDGE1, click **Start**, type **mmc** and then click **mmc**. Click **Yes** at the **User Account Control** prompt.
2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the **Certificates** snap-in, open **Certificates (Local Computer)\Personal\Certificates**.Click **Next** three times to accept the default settings, and then click **Install**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice
7. On the **Request Certificates** page, click **Client Server Authentication** and then click **More information is required to enroll for this certificate**.
8. On the **Subject** tab of the **Certificate Properties** dialog box, in **Subject name**, for **Type**, select **Common Name**.
9. In **Value**, type **EDGE1.corp.contoso.com**, and then click **Add**.
10. Click **OK**, click **Enroll**, and then click **Finish**.
11. In the **Details** pane of the **Certificates** snap-in, verify that a new certificate with the name **EDGE1.corp.contoso.com** was enrolled with **Intended Purposes of Client**

Authentication, Server Authentication, IP security IKE intermediate.

12. Close the console window. If you are prompted to save settings, click **No**.

STEP 5: Configure EDGE1 and 3-EDGE1 to use EAP for site-to-site authentication

Change the authentication method to EAP in EDGE1

1. On EDGE1, In the **Routing and Remote Access** snap-in, expand **EDGE1**, and then click **Network Interfaces**.
2. Right-click **EDGE1@corp.contoso.com** and then click **Properties**.
3. Click the **Security** tab and then select **Use Extensible Authentication Protocol (EAP)**.
4. Right-click the **IPsec** template in the list, and then click **Duplicate Template**.
5. Select **EAP-MSCHAP V2** in the dropdown menu
6. Click Ok.
7. When prompted for username\password, enter **3-EDGE1**, **Pass@word1**

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
Set-VpnS2SInterface  EDGE1  -Protocol IKEv2 –AuthenticationMethod –EAP –UserName "3edge1"
–Password Pass@word1
```

Change the authentication method to EAP in 3-EDGE1

1. On EDGE1, In the **Routing and Remote Access** snap-in, expand **EDGE1**, and then click **Network Interfaces**.
2. Right-click **EDGE1@corp.fabrikam.com** and then click **Properties**
3. Click the **Security** tab and then select **Use Extensible Authentication Protocol (EAP)**.
4. Select **EAP-MSCHAP V2** in the dropdown menu.
5. Click **Ok**.
6. When prompted for username\password, enter **edge1**, **Pass@word1**

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-

wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
Set-VpnS2SInterface  EDGE1  -Protocol IKEv2 -AuthenticationMethod -EAP -UserName "edge1"
-Password Pass@word1
```

Initiate the site-to-site connection

1. **On 3-EDGE1**, in the console tree of the **Routing and Remote Access** snap-in, click **Network Interfaces**.
2. In the **Details** pane, right-click **edge1@corp.fabrikam.com**, and then click **Connect**.
3. Confirm that the connection state of **VPN_Corpnet** is connected.

Test the site-to-site connection

1. On **DC1** at the **Start** menu, type **cmd** and then hit **Enter**. At the command prompt, type **ping 10.6.0.2**. Verify that there are four replies from 10.6.0.2. On 3-EDGE1, in the **Internet Explorer** address bar, type **http://localhost/certsrv**, and then press **ENTER**.

STEP 6: Configure EDGE1 and 3-EDGE1 to use certificates for site-to-site authentication

Configure computer certificate authentication on EDGE1

To configure computer certificate authentication on EDGE1

1. In **Routing and Remote Access**, right click **EDGE1(local)**. Click **Properties**.
2. In the **Security** tab, click **Authentication Methods**. Check **Allow machine certificate authentication for IKEv2**.
3. Click **OK** twice.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
PS C:\> $cert1 = ( Get-ChildItem -Path cert:LocalMachine\root | Where-Object -
FilterScript { $_.Subject -Like "*CN=Contoso Root Certification Authority,*" } )
```

```
Set-VpnAuthProtocol -UserAuthProtocolAccepted Certificate -RootCertificateNameToAccept
$cert1 -PassThru
```

Note

This change is applicable to both L2TP and IKEv2 VPN (site-to-site and Remote Accesss) connections as well. If you change the authentication method on the server from certificates to pre-shared key, it impacts VPN client connections as well. If a server is configured to advertise only pre-shared key, VPN clients validating the server using certificates will fail to authenticate the server.

Configure name resolution

To configure name resolution

1. In the **Start** screen, type **notepad**, right-click **Notepad** and then click **Run as administrator**.
2. Click **Yes** in the **User Account Control** dialog.
3. In **Notepad**, click **File** and then click **Open**.
4. In the **Open** dialog, change **file type** to **All Files**. Navigate to **C:/Windows/System32/drivers/etc**.
5. Open the **hosts** file.
6. Go to end of the hosts file and type, **131.107.0.4 3-EDGE1.corp.fabrikam.com**.
7. Close the file and click **Save** when prompted.

Export the root certificate (corp-DC1-CA) from EDGE1

To Export the root certificate (corp-DC1-CA) from EDGE1

1. On EDGE1, click **Start**, type **mmc**, and then click **mmc**. Click **Yes** at the **User Account Control** prompt.
2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the **Certificates** snap-in, open **Certificates (Local Computer)\Trusted Root Certification Authorities\Certificates**.
5. Right click **corp-DC1-CA**. Point to **All Tasks** and click **Export**.
6. On **Certificate Export Wizard**, click **Next** twice. Click **Browse** and select **Desktop**. In **File Name** box type **DC1cert**. Click **OK**.
7. Click **Next** and then click **Finish**.

8. Click **OK** on the dialog indicating the export was successful.
9. Go to **Desktop**. Right-click **DC1cert**, point to **Share with** and click **Specific people**.
10. In **File Sharing** dialog, select **Everyone** from dropdown list. Click **Share**. Click **Done**.

Warning
It is not advisable to share this file with everyone in a non-test production environment.

Configure computer certificate authentication on 3-EDGE1

To configure computer certificate authentication on 3-EDGE1

1. In **Routing and Remote Access**, right click **3-EDGE1(local)**. Click **Properties**.
2. In the **Security** tab, click **Authentication Methods**. Check **Allow machine certificate authentication for IKEv2**
3. Click **OK** twice.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
PS C:\> $cert1 = ( Get-ChildItem -Path cert:LocalMachine\root | Where-Object -
FilterScript { $_.Subject -Like "*CN=Contoso Root Certification Authority,*" } )


Set-VpnAuthProtocol -UserAuthProtocolAccepted Certificate -RootCertificateNameToAccept
$cert1 -PassThru
```

Note
This change is applicable to both L2TP and IKEv2 VPN (site-to-site and Remote Accesss) connections as well. If you change the authentication method on the server from certificates to pre-shared key, it impacts VPN client connections as well. If a server is configured to advertise only pre-shared key, VPN clients validating the server using certificates will fail to authenticate the server.

Configure name resolution

To configure name resolution

1. In the **Start** screen, type **notepad**, right-click **Notepad** and then click **Run as administrator**.

2. Click **Yes** in the **User Account Control** dialog.
3. In **Notepad**, click **File** and then click **Open**.
4. In the **Open** dialog, change **file type** to **All Files**. Navigate to **C:/Windows/System32/drivers/etc**.
5. Open the **hosts** file.
6. Go to end of the hosts file and type, **131.107.0.2 3-EDGE1.corp.contoso.com**.
7. Close the file and click **Save** when prompted.

Export the root certificate (corp-3-DC1-CA) from 3-EDGE1

To Export the root certificate (corp-3-DC1-CA) from 3-EDGE1

1. On 3-EDGE1, click **Start**, type **mmc**, and then click **mmc**. Click **Yes** at the **User Account Control** prompt.
2. Click **File**, and then click **Add/Remove Snap-ins**.
3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the **Certificates** snap-in, open **Certificates (Local Computer)\Trusted Root Certification Authorities\Certificates**.
5. Right click **corp-3-DC1-CA**. Point to **All Tasks** and click **Export**.
6. On **Certificate Export Wizard**, click **Next** twice. Click **Browse** and select **Desktop**. In **File Name** box type **3-DC1cert**. Click **OK**.
7. Click **Next** and then click **Finish**.
8. Click **OK** on the dialog indicating the export was successful.
9. Go to **Desktop**. Right-click **3-DC1cert**, point to **Share with** and click **Specific people**.
10. In **File Sharing** dialog, select **Everyone** from dropdown list. Click **Share**. Click **Done**.

    Warning
    It is not advisable to share this file with everyone in a non-test production environment.

Import the root certificate (corp-DC1-CA) to 3-EDGE1

To Import the root certificate (corp-DC1-CA) to 3-EDGE1

1. On 3-EDGE1, click **Start**, type **mmc**, and then click **mmc**. Click **Yes** at the **User Account Control** prompt.
2. Click **File**, and then click **Add/Remove Snap-ins**.

3. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the **Certificates** snap-in, open **Certificates (Local Computer)\Trusted Root Certification Authorities**.
5. Right click **Trusted Root Certification Authorities**, point to **All Tasks** and click **Import**.
6. On **Certificate Import Wizard**, click **Next**. Click **Browse**. In the address bar in **Open** dialog type **\\131.107.0.2\Users\Desktop** and press **Enter**. Select **DC1cert** and click **Open**.
7. Click **Next** twice and then click **Finish**.

Change the authentication method to certificates on EDGE1

To Change the authentication method to certificates on EDGE1

1. In the **Routing and Remote Access** snap-in, expand **EDGE1**, and then click **Network Interfaces**.
2. Right-click **3-EDGE1.corp.fabrikam.com** and then click **Properties**.
3. Click the **Security** tab, select **Use machine certificates**, and then click **Verify the Name and Usage attributes of the server's certificate**.
4. Click **OK**.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$cert1 = Get-ChildItem -Path cert:LocalMachine\My | Where-Object -FilterScript {
$_.Subject -Like "*CN edge1.contoso.com*" }

Set-VpnS2SInterface –Name edge1 -AuthenticationMethod MachineCertificates -Certificate
$cert1 -ResponderAuthenticationMethod MachineCertificates -Protocol IKEv2
```

Change the authentication method to certificates on 3-EDGE1

To Change the authentication method to certificates on 3-EDGE1

1. In the **Routing and Remote Access** snap-in, expand **3-EDGE1**, and then click **Network Interfaces**.
2. Right-click **EDGE1.corp.contoso.com** and then click **Properties**.
3. Click the **Security** tab, select **Use machine certificates**, and then click **Verify the Name and Usage attributes of the server's certificate**.
4. Click **OK**.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
$cert1 = Get-ChildItem -Path cert:LocalMachine\My | Where-Object –FilterScript {
$_.Subject -Like "*CN=3-edge1.contoso.com*" }
Set-VpnS2SInterface –Name 3-edge1 -AuthenticationMethod MachineCertificates -Certificate
$cert1 -ResponderAuthenticationMethod MachineCertificates -Protocol IKEv2
```

Reinitiate the site-to-site connection

To Reinitiate the site-to-site connection

1. On **3-EDGE1**, in the console tree of the **Routing and Remote Access** snap-in, click **Network Interfaces**.
2. In the **Details** pane, right-click **EDGE1.corp.contoso.com**, and then click **Connect**.
3. Confirm that the connection state of **VPN_Corpnet** is connected.

Test the site-to-site connection

To Test the site-to-site connection

1. On **DC1** at the **Start** menu, type **cmd** and then hit **Enter**. At the command prompt, type **ping 10.6.0.2**. Verify that there are four replies from 10.6.0.2

Snapshot the configuration

This completes the Base Configuration test lab. To save this configuration for additional test labs, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots Windows Server® 2012 VPN Cloud Configuration. If your lab uses physical computers, create disk images to save the Base Configuration.

Manage Remote Access

The DirectAccess Remote Client Management deployment scenario uses DirectAccess to maintain clients over the Internet. This document explains the scenario, including its phases, roles, features, and links to additional resources.

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

Scenario description

DirectAccess client computers are connected to the intranet whenever they are connected to the Internet, regardless of whether the user has signed in to the computer. They can be managed as intranet resources and kept current with Group Policy changes, operating system updates, antimalware updates, and other organizational changes.

In some cases, intranet servers or computers must initiate connections to DirectAccess clients. For example, Help Dtechnicians can use remote desktop connections to connect to and troubleshoot remote DirectAccess clients. This scenario lets you keep your existing remote access solution in place for user connectivity, while using DirectAccess for remote management.

DirectAccess in Windows Server 2012 provides a configuration that supports remote management of DirectAccess clients. You can use a deployment wizard option that limits the creation of policies to only those needed for remote management of client computers.

Note
In this deployment, user-level configuration options such as force tunneling, Network Access Protection (NAP) integration, and two-factor authentication are not available.

In this scenario

The DirectAccess Remote Client Management deployment scenario includes the following steps for planning and configuring.

Plan the deployment

There are only a few computer and network requirements for planning this scenario. They include:

- **Network and server topology**: With DirectAccess, you can place your Remote Access server at the edge of your intranet or behind a network address translation (NAT) device or a firewall.
- **DirectAccess network location server**: The network location server is used by DirectAccess clients to determine whether they are located on the internal network. The network location server can be installed on the DirectAccess server or on another server.
- **DirectAccess clients**: Decide which managed computers will be configured as DirectAccess clients.

Configure the deployment

Configuring your deployment consists of a number of steps. These include:

1. **Configure the infrastructure**: Configure DNS settings, join the server and client computers to a domain if required, and configure Active Directory security groups.

   In this deployment scenario, Group Policy Objects (GPOs) are created automatically by Remote Access. For advanced certificate GPO options, see Deploy a Single Remote Access Server with Advanced Settings.
2. **Configure Remote Access server and network settings**: Configure network adapters, IP addresses, and routing.
3. **Configure certificate settings**: In this deployment scenario, the Getting Started Wizard creates self-signed certificates, so there is no need to configure the more advanced certificate infrastructure that is explained in the topic Deploy a Single Remote Access Server with Advanced Settings.
4. **Configure the network location server**: In this scenario, the network location server is installed on the Remote Access server.
5. **Plan DirectAccess management servers**: Administrators can remotely manage DirectAccess client computers that are located outside the corporate network by using the Internet. Management servers include computers that are used during remote client management (such as update servers).
6. **Configure the Remote Access server**: Install the Remote Access role and run the DirectAccess Getting Started Wizard to configure DirectAccess.
7. **Verify the deployment**: Test a client to make sure it is able to connect to the internal network and the Internet by using DirectAccess.

Practical applications

Deploying a single Remote Access server for managing DirectAccess clients provides the following:

- **Ease-of-access**: Managed client computers running Windows 8 or Windows 7 can be configured as DirectAccess client computers. These clients can access internal network resources through DirectAccess any time they are connected to the Internet without needing to sign in to a VPN connection. Client computers not running one of these operating systems can connect to the internal network through VPN. DirectAccess and VPN are managed in the same console and with the same set of wizards.
- **Ease-of-management**: DirectAccess client computers that are connected to the Internet can be remotely managed by remote access administrators by using DirectAccess, even when the client computers are not located in the internal corporate network. Client computers that do not meet corporate requirements can be remediated automatically by management servers. One or more Remote Access servers can be managed from a single Remote Access Management console.

Roles and features included in this scenario

The following table lists the roles and features required for the scenario:

Role or feature	How it supports this scenario
Remote Access role	This role is installed and uninstalled by using the Server Manager console or Windows PowerShell. This role encompasses DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services, which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN: DirectAccess and VPN are managed in the Remote Access Management console. 2. RRAS: Features are managed in the Routing and Remote Access console. The Remote Access server role is dependent on the following features: • Internet Information Services (IIS) Web Server:Required to configure the network location server and default web probe. • Windows internal database: Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • By default on a Remote Access server when the Remote Access role is installed and supports the Remote Management console user interface. • As an option on a server that is not running the Remote Access server role. In this case, it is used for remote management of a Remote Access server running DirectAccess and VPN. This feature consists of the following: • Remote Access GUI and command-line tools • Remote Access module for Windows

Role or feature	How it supports this scenario
	PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

Server requirements

- A computer that meets the hardware requirements for Windows Server 2012. For more information, see Installing Windows Server 2012.
- The server must have at least one network adapter installed and enabled. There should be only one adapter connected to the corporate internal network, and only one connected to the external network (Internet).
- If Teredo is required as an IPv6 to IPv4 transition protocol, the external adapter of the server requires two consecutive public IPv4 addresses. If a single network adapter is available, only IP-HTTPS can be used as the transition protocol.
- At least one domain controller. The Remote Access servers and DirectAccess clients must be domain members.
- A certification authority is required on the server if you do not want to use self-signed certificates for IP-HTTPS or the network location server, or if you want to use client certificates for client IPsec authentication.

Client requirements

- A client computer must be running Windows 8 or Windows 7.

Infrastructure and management server requirements

- During remote management of DirectAccess client computers, clients initiate communications with management servers, such as domain controllers, System Center Configuration Servers, and Health Registration Authority (HRA) servers. These servers provide services that include Windows and antivirus updates and Network Access Protection (NAP) client compliance. You should deploy the required servers before you begin the Remote Access deployment.
- A DNS server running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 with SP2 is required.

Software requirements

Software requirements for this scenario include the following:

Server requirements

- The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
- If the Remote Access server is located behind an edge firewall or NAT device, the device must be configured to allow traffic to and from the Remote Access server.
- Admins who deploy a Remote Access server require local administrator permissions on the server and domain user permissions. In addition, the administrator requires permissions for the GPOs that are used for DirectAccess deployment. To take advantage of the features that restrict DirectAccess deployment to only mobile computers, Domain Admin permissions are required on the domain controller to create a WMI filter.
- If the network location server is not located on the Remote Access server, a separate server to run it is required.

Remote access client requirements

- DirectAccess clients must be domain members. Domains that contain clients can belong to the same forest as the Remote Access server, or they can have a two-way trust with the Remote Access server forest or domain.
- An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients. Computers should not be included in more than one security group that includes DirectAccess clients. If clients are included in multiple groups, name resolution for client requests will not work as expected.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
Deployment	Remote Access (DirectAccess, Routing and Remote Access) Overview

Content type	References
Tools and settings	Remote Access PowerShell cmdlets
Community resources	RRAS Product Team blog | Remote Access TechNet Forum DirectAccess Wiki entries
Related technologies	How IPv6 works

Use Remote Access Monitoring and Accounting

Remote Access monitoring reports remote user activity and status for DirectAccess and VPN connections. It tracks the number and duration of client connections (among other statistics), and monitors the operations status of the server. An easy-to-use monitoring console provides a view of your entire Remote Access infrastructure. Monitoring views are available for single server, cluster, and multisite configurations.

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

In this guide

This document contains instructions for leveraging the monitoring capabilities of Remote Access by using the DirectAccess management console and the corresponding Windows PowerShell cmdlets, which are provided as part of the Remote Access server role.

The following monitoring and accounting scenarios are explained:

1. Monitor the existing load on the Remote Access server
2. Monitor the configuration distribution status of the Remote Access server
3. Monitor the operations status of the Remote Access server and its components
4. Identify and resolve Remote Access server operations issues
5. Monitor connected remote clients for activity and status
6. Generate a usage report for remote clients by using historical data

Understand monitoringand accounting

Before you begin monitoring and accounting tasks for remote clients, you need to understand the difference between the two.

- **Monitoring** shows actively connected users at a given point in time.

- **Accounting** keeps a history of users who have connected to the corporate network, and their usage details (for compliance and auditing purposes).

Remote client monitoring is based on connections. There are two types of tunnel connections that are established by DirectAccess clients:

- **Machine tunnel traffic connections**: This tunnel is established by the computer, in system context, to access servers that are required for name resolution, authentication, remediation updating, and so on.
- **User tunnel traffic connections**: This tunnel is established by the user account on the computer, in a user context, when the user tries to access a resource on the corporate network. Depending on the deployment requirements, a user might have to provide strong credentials (for example, by using a smart card or providing a one-time password) to access the corporate network resources.

For DirectAccess, a connection is uniquely identified by the IP address of the remote client. For example, if a machine tunnel is open for a client computer, and a user is connected from that computer, these would be using the same connection. In a situation where the user disconnects and connects again while the machine tunnel is still active, it is a single connection.

Monitor the existing load on the Remote Access server

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

The term **Load** refers to the statistics that relate to the number of connections on the Remote Access server. Following are the steps required to track the load on the Remote Access server.

You can use the monitoring dashboard that is available in the management console on the Remote Access server to view the load statistics for the server, or you can use Performance Monitor counters to track the statistics.

Note

You must be signed in as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this topic. If you cannot complete a task while you are signed in with an account that is a member of the Administrators group, try performing the task while you are signed in with an account that is a member of the Domain Admins group.

To use the monitoring dashboard to monitor the Remote Access server load

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **DASHBOARD** to navigate to **Remote Access Dashboard** in the **Remote Access Management Console**.
3. On the monitoring dashboard, notice the **Remote Client Status** tile within the **Server Status** tile. This tile lists statistics such as the total number of remote clients that are

connected, the total number of DirectAccess clients that are connected, and the maximum number of users who connected in last 24 hours.

4. You can click **Refresh** under **Tasks** in the right pane to reload the health status. To change the default refresh interval, click **Configure Refresh Interval** under **Tasks**.

To use the Performance Monitor tool to monitor performance counters on the Remote Access server

1. Click **Start**, click **Administrative Tools**, and then double-click **Performance Monitor**.
2. Under **Performance**, click **Performance Monitor**.
3. Click the **Add** button (denoted by a green cross icon) in the **Performance Monitor** toolbar.
4. From the list of **Available Counters**, select all the counters in the **RAS** and **RAmgmtsvc** categories, and then click **Add>>**.
5. Again, from the list of **Available Counters**, select all the counters in the **IPsec Connections** category, and then click **Add>>.**
6. Click **OK** to add the selected counters in the **Performance Monitor** console for tracking.

Performance Monitor will now graphically show the selected server load statistics.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
PS> Get-RemoteAccessConnectionStatisticsSummary
```

Monitor the configuration distribution status of the Remote Access server

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

The Remote Access Management Console compares the configuration versions from all the monitored servers to verify that they match and are using the latest configuration version. This shows whether the latest configuration version (which is specified in the Group Policy Objects or GPOs) was distributed to all of the servers and whether it was successfully applied on the servers.

To use the monitoring dashboard to monitor the configuration distribution

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **DASHBOARD** to navigate to **Remote Access Dashboard** in the **Remote Access**

Management Console.

3. On the monitoring dashboard, notice the **Configuration Status** tile at the top center. This tile shows the current status of the configuration distribution.

The following table shows the messages that are generated by the **Configuration Status** tile, their meanings, and the necessary administrative action (if any).

Severity	Message	Meaning	What to do?
Success	The configuration was distributed successfully.	The configuration in the GPO was successfully applied on the server.	No action needed.
Warning	Configuration for server [*server name*] not retrieved from the domain controller. The GPO is not linked.	The configuration in the GPO did not yet reach the server. This could be because the GPO is not linked to the server.	Link the GPO to a scope of management that is applied to the server, or in a staging GPO scenario, manually export the settings from the staging GPO and import them to the production GPO. For more information about staging GPOs, see **Managing Remote Access GPOs with limited permissions** in Step 1: Plan the Remote Access Infrastructure. For GPO staging steps, see **Configuring Remote Access GPOs with limited permissions** in Step 1: Configure the Remote Access Infrastructure.
Warning	Configuration for server [*server name*] not yet retrieved from the domain controller.	The configuration in the GPO did not yet reach the server. It can take up to 10 minutes to propagate a new configuration.	Allow more time for the policies to update on the server.
Error	Configuration for server [*server name*] cannot be retrieved from the domain controller.	The configuration in the GPO did not reach the server, and more than 10 minutes have	This could happen in one of the following scenarios: • The server has no connectivity to the domain to update the policies. You can run “gpupdate /force” on the server to force a policy update.

		passed since the configuration was changed.	• GPO replication might be required to retrieve the updated configuration. • There is no writable domain controller in the Active Directory site of the Remote Access server. Wait for GPOs to replicate to all domain controllers, and then use the Windows PowerShell cmdlet **Set-DAEntryPointDC** to associate the entry point with a writable domain controller in Active Directory in the Remote Access server.
Warning	Configuration for server [*server name*] retrieved from the domain controller, but not yet applied.	The configuration in the GPO reached the server but is not yet applied. It can take up to 15 minutes before the configuration is applied.	Allow more time for the configuration to be fully applied to the server.
Error	Configuration for server [*server name*] retrieved from the domain controller cannot be applied.	The configuration in the GPO reached the server but is not successfully applied, and more than 15 minutes have passed since the configuration was changed.	This could happen in one of the following scenarios: 1. The configuration is currently in the process of being applied. This is shown as an error because it may have taken a long time to retrieve the configuration from the GPO. To verify whether this is the reason, use **Task Scheduler** and navigate to Microsoft\Windows\RemoteAccess to verify that **RAConfigTask** is currently running. 2. If **RAConfigTask** is not currently running, it may have failed to apply the configuration on the server. Check for errors in **Event Viewer** under the Remote Access server operations

			channel, which is located at \Applications and Services Logs\Microsoft\Windows\RemoteAccess-RemoteAccessServer. Check for errors in **OPERATIONS STATUS** in the Remote Access Management Console. For more information, see Monitor the operations status of the Remote Access server and its components.
Error	Configuration for multisite servers retrieved from the domain controller. The configuration does not match on all servers.	There is an inconsistency between the configuration versions of the server GPOs in the multisite deployment. Ideally, all the server GPOs for all entry points will have the same global configuration, but for some reason, they are out of sync.	This can happen when a configuration change failed and was not rolled back successfully. You should restore the GPOs from a backup state where all server GPOs were synchronized. For information about a script that you can use, see Back up and Restore Remote Access Configuration.

Monitor the operations status of the Remote Access server and its components

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

The management console in the Remote Access server can be used to monitor its operations status.

Note

You must be signed in as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the task described in this topic. If you cannot complete a task while you are signed in with an account that is a member of the Administrators group, try performing the task while you are signed in with an account that is a member of the Domain Admins group.

To monitor the Remote Access server operations status

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **DASHBOARD** to navigate to **Remote Access Reporting** in the **Remote Access Management Console**.
3. On the monitoring dashboard, notice the **Operations Status** tile within the **Server Status** tile. This tile lists the server operations status and the status of all the server's components.
4. Click **Refresh** under **Tasks** in the right pane to reload the operations status. The operations status is automatically refreshed every five minutes, which is the default refresh interval. To change the default refresh interval, click **Configure Refresh Interval**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note

The command for operations status of a cluster is included for reference.

```
PS> Get-RemoteAccessHealth

PS> Get-RemoteAccessHealth -Cluster
```

Identify and resolve Remote Access server operations problems

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

You can using the following procedures to identify Remote Access server operations issues, their root causes, and the resolution required to fix the issues.

Note

You must be signed in as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this topic. If you cannot complete a task while you are signed in with an account that is a member of

the Administrators group, try performing the task while you are signed in with an account that is a member of the Domain Admins group.

This topic includes information about performing the following tasks:

Simulate an operations issue

Identify the operations issue and take corrective action

Restore the IP Helper service

Simulate an operations issue

Caution

Because your Remote Access server is probably configured properly and not experiencing any issues, you can use the following procedure to simulate an operations issue. If your server is currently servicing clients in a production environment, you may not want to take these actions at this time. Rather, you can read through the steps to understand how to address issues that might arise on your Remote Access server in the future.

The IP Helper service (IPHlpSvc) hosts IPv6 transitioning technologies (such as IP-HTTPS, 6to4, or Teredo), and it is required for the DirectAccess server to function properly. To demonstrate a simulated operations issue on the Remote Access server, you must stop the (IPHlpSvc) network service.

To stop the IP Helper service

1. On the **Start** screen of the Remote Access server, click **Administrative Tools**, and then double-click **Services**.
2. In the list of **Services**, scroll down and right-click **IP Helper**, and then click **Stop**.

Identify the operations issue and take corrective action

Turning off the IP Helper service will cause a serious error on the Remote Access server. The monitoring dashboard will show the operations status of the server and the details of the issue.

To identify the details and take corrective action

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **DASHBOARD** to navigate to **Remote Access Dashboard** in the **Remote Access Management Console**.
3. Make sure your Remote Access server is selected in the left pane, and then in the middle pane, click **Operations Status**.
4. You will see the list of components with green or red icons, which indicate their operational status. Click the **IP-HTTPS** row in the list. When you selected a row, the details for the operation are shown in the **Details** pane as follows:

   Error

The IP Helper service (IPHlpSvc) has stopped. DirectAccess might not function as expected. The IP Helper service provides tunnel connectivity by using the connectivity platform, IPv6 transition technologies, and IP-HTTPS.

Causes

a. The IP Helper service has stopped.

b. The IP Helper service is not responding.

Resolution

a. To ensure that the service is running, type **Get-Service iphlpsc** at a Windows PowerShell prompt.

b. To enable the service, type **Start-Service iphlpsvc** from an elevated Windows PowerShell prompt.

c. To restart the service, type **Restart-Service iphlpsvc** from an elevated Windows PowerShell prompt.

Restore the IP Helper service

To restore the IP Helper service on your Remote Access server, you can follow the Resolution steps above to start or restart the service, or you can use the following procedure to reverse the procedure that you used to simulate the IP Helper service failure.

To restart the IP Helper service on the Remote Access server

1. On the **Start** screen, click **Administrative Tools**, and then double-click **Services**.
2. In the list of **Services**, scroll down and right-click **IP Helper**, and then click **Start**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
PS> Get-RemoteAccessHealth | Where-Object {$_.Component –eq “IP-HTTPS”} | Format-List –Property *
```

Monitor connected remote clients for activity and status

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

You can use the management console on the Remote Access server to monitor remote client activity and status.

Note

You must be signed in as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this topic. If you cannot complete a task while you are signed in with an account that is a member of the Administrators group, try performing the task while you are signed in with an account that is a member of the Domain Admins group.

To monitor remote client activity and status

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **REPORTING** to navigate to **Remote Access Reporting** in the **Remote Access Management Console**.
3. Click **Remote Client Status** to navigate to the remote client activity and status user interface in the **Remote Access Management Console**.
4. You will see the list of users who are connected to the Remote Access server and detailed statistics about them. Click the first row in the list that corresponds to a client. When you select a row, the remote user activity is shown in the preview pane.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
PS> Get-RemoteAccessConnectionStatistics
```

The user statistics can be filtered, based on criteria selections, by using the fields in the following table.

Field Name	Value
Username	The user name or alias of the remote user. Wildcard characters can be used to select a group of users, such as contoso* or *\administrator.
Hostname	The computer account name of the remote user. An IPv4 or IPv6 address also can be specified.
Type	DirectAccess or VPN. If DirectAccess is selected, all remote users who are connected by using DirectAccess are listed. If VPN is selected, all remote users who are connected by using VPN are listed.

Field Name	Value
ISP address	The IPv4 or IPv6 address of the remote user.
IPv4 address	The inner IPv4 address of the tunnel that connect the remote user to the corporate network.
IPv6 address	The inner IPv6 address of the tunnel that connects the remote user to the corporate network.
Protocol/Tunnel	The transitioning technology that is used by the remote client. This is Teredo, 6to4, or IP-HTTPS for DirectAccess users, and it is PPTP, L2TP, SSTP, or IKEv2 for VPN users.
Resource Accessed	All users who are accessing a particular corporate resource or an endpoint. The value that corresponds to this field is the hostname/IP address of the server.
Server	The Remote Access server to which clients are connected. This is relevant only for cluster and multisite deployments.

Generate a usage report for remote clients using historical data

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

The management console on the Remote Access server can be used to generate a usage report for the remote clients that are accessing the server. To generate a usage report for remote clients, you first enable accounting on the Remote Access server. After you generate the report, you can use the monitoring dashboard that is available in the management console on the Remote Access server to view the load statistics on the server.

Note

You must be signed in as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this topic. If you cannot complete a task while you are signed in with an account that is a member of the Administrators group, try performing the task while you are signed in with an account that is a member of the Domain Admins group.

To enable accounting on the Remote Access Server

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **REPORTING** to navigate to **Remote Access Reporting** in the **Remote Access Management Console**.
3. Click **Configure Accounting** in the **Remote Access Reporting** task pane.
4. Select the **Use inbox accounting** check box to enable accounting on the Remote Access server.
5. Click **Apply** to enable the accounting configuration on the server, and then click **Close** after the server has applied the configuration successfully.

To generate the usage report

1. In **Server Manager**, click **Tools**, and then click **Remote Access Management**.
2. Click **REPORTING** to navigate to **Remote Access Reporting** in the **Remote Access Management Console**.
3. In the middle pane, click dates in the calendar to select the report duration **Start date:** and **End date:**, and then click **Generate Report**.
4. You will see the list of users that have connected to the Remote Access server within the selected time and detailed statistics about them. Click the first row in the list. When you select a row, the remote user activity is shown in the preview pane.

   Click the **Server Load Statistics** tab in the preview pane to see the historical load on the server.

Notes

Understanding sessions

Remote Access accounting is based on the concept of **sessions**. In contrast to a **connection**, a **session** is uniquely identified by a combination of remote client IP address and user name. For example, if a machine tunnel is formed from the remote client, named Client1, a session will be created and stored in the accounting database. When a user named User1 connects from that client after some time passes (but the machine tunnel is still active), the session is recorded as a separate session. The distinction of sessions is to retain the distinction between machine tunnel and user tunnel.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

In the following script, change the date range for which you want a report in the **–StartDateTime** and **–EndDateTime** parameters.

```
PS> Get-RemoteAccessConnectionStatisticsSummary –StartDateTime “1 October 2010 00:00:00” –EndDateTime “14 October 2010 00:00:00”
```

```
Shows server load statistics.

PS> Get-RemoteAccessUserActivity -HostIPAddress 10.0.0.1 -StartDateTime "1 October 2010
00:00:00" -EndDateTime "14 October 2010 00:00:00"
```

Manage DirectAccess Clients Remotely

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This topic provides an introduction to an advanced scenario that you can use to set up a single Remote Access server for remote management of DirectAccess clients.

Scenario description

In this scenario, a single computer running Windows Server 2012 is configured as a Remote Access server with the only purpose of remotely managing DirectAccess clients. Using this scenario enables the remote management of clients, but it disables other components that you can use if you select a full DirectAccess deployment, including client access to internal networks, force tunneling, strong authentication, and NAP compliance.

Note
If you want to configure a basic deployment with simple settings only, see Deploy a Single Remote Access Server using the Getting Started Wizard. In the simple scenario, you can set up Remote Access with default settings by using a wizard. You do not configure infrastructure settings, such as a certification authority (CA) or Active Directory security groups.

In this scenario

To set up a single Remote Access server to manage clients, a number of planning and deployment steps are required.

Planning steps

Planning for this scenario is divided into two phases:

1. Plan the Remote Access infrastructure: In this phase, you plan the network infrastructure before you begin the Remote Access deployment. It includes planning for the network and server topology, certificates, Domain Name System (DNS), Active Directory, Group Policy Objects (GPOs), and the DirectAccess network location server.
2. **Plan the Remote Access deployment**: In this phase, you prepare for the Remote Access deployment. It includes planning for Remote Access client computers, server and client authentication requirements, VPN settings, infrastructure, and management servers.

For detailed planning steps, see Plan Deployment for Remote Management of DirectAccess Clients.

Prerequisites

Before you begin this scenario, review this list for important requirements:

- Windows firewall must be enabled on all profiles.
- ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
- DirectAccess supports only clients running Windows 8.1, Windows 8, and Windows 7.
- Changing policies outside of the DirectAccess management console or by using Windows PowerShell cmdlets is not supported.

Deployment steps

Deployment for this scenario is divided into three phases:

1. **Configure the Remote Access infrastructure**: In this phase, you configure network and routing, firewall settings if required, certificates, DNS servers, Active Directory and GPO settings, and the DirectAccess network location server.
2. **Configure Remote Access server settings**: In this phase, you configure Remote Access client computers, the Remote Access server, infrastructure servers, and management and application servers.
3. **Verify the deployment**: In this phase, you verify that the deployment is working as required.

For detailed deployment steps, see Install and Configure Deployment for Remote Management of DirectAccess Clients.

Practical applications

Deploying a single Remote Access server for managing DirectAccess clients provides the following benefits:

- **Ease-of-access**: Managed client computers running Windows 8.1, Windows 8, or Windows 7 can be configured as DirectAccess client computers. These clients can access internal network resources through DirectAccess any time they are connected to the Internet, without the need to sign in to a VPN connection. Client computers that are not running one of these operating systems can connect to the internal network through a VPN. DirectAccess and VPN are managed in the same console and with the same set of wizards.
- **Ease-of-management**: DirectAccess client computers that are connected to the Internet can be remotely managed by remote access administrators through DirectAccess, even when the client computers are not located on the internal corporate network. Client computers that do not meet corporate requirements can be remediated automatically by using management servers.

Roles and features included in this scenario

The following table lists the roles and features that are required to plan and deploy this scenario.

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled using the Server Manager console or Windows PowerShell. This role encompasses both DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN—DirectAccess and VPN are managed together in the Remote Access Management console. 2. RRAS Routing—RRAS routing features are managed in the legacy Routing and Remote Access console. The Remote Access Server Role is dependent on the following server roles\features: • Internet Information Services (IIS) Web Server – This feature is required to configure the network location server on the Remote Access server, and the default web probe. • Windows Internal Database—Used for local accounting on the Remote Access server.
Remote Access Management Tools feature	This feature is installed as follows: • It is installed by default on a Remote Access server when the Remote Access role is installed, and supports the Remote Management console user interface and the Windows PowerShell cmdlets. • It can be optionally installed on a server not running the Remote Access server role. In this case it is used for remote management of a Remote Access computer running

Role/feature	How it supports this scenario
	DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI • Remote Access module for Windows PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

- Server requirements:
  - A computer that meets the hardware requirements for Windows Server 2012.
  - The server must have at least one network adapter installed and enabled. When two adapters are used, there should be one adapter connected to the internal corporate network, and one connected to the external network (Internet).
  - If Teredo is required as an IPv4 to IPv6 transition protocol, the external adapter of the server requires two consecutive public IPv4 addresses. If a single IP address is available, then only IP-HTTPS can be used as the transition protocol.
  - At least one domain controller. The Remote Access server and DirectAccess clients must be domain members.
  - A CA server is required if you do not want to use self-signed certificates for IP-HTTPS or the network location server, or if you want to use client certificates for client IPsec authentication. Alternatively, you can request certificates from a public CA.
  - If the network location server is not located on the Remote Access server, a separate web server to run it is required.
- Client requirements:
  - A client computer must be running Windows 8 or Windows 7.

    Note
    Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

- Infrastructure and management server requirements:
  - During remote management of DirectAccess client computers, clients initiate communications with management servers such as domain controllers and System Center Configuration Servers, for services that include Windows and antivirus updates and Network Access Protection (NAP) client compliance. The required servers should be deployed before beginning the Remote Access deployment.
  - If remote access requires client NAP compliance, NPS and HRS servers should be deployed before beginning remote access deployment
  - If VPN is enabled, a DHCP server is required to allocate IP addresses automatically to VPN clients, if a static address pool is not used.
  - A DNS server running Windows Server 2008 SP2; Windows Server 2008 R2; or Windows Server 2012 is required.

Software requirements

There are a number of requirements for this scenario:

- Server requirements:
  - The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
  - If the Remote Access server is located behind an edge firewall or NAT device, the device must be configured to allow traffic to and from the Remote Access server.
  - The person deploying remote access on the server requires local administrator permissions on the server, and domain user permissions. In addition, the administrator requires permissions for the GPOs used in DirectAccess deployment. To take advantage of the features that restricts DirectAccess deployment to mobile computers only, permissions to create a WMI filter on the domain controller are required.
- Remote access client requirements:
  - DirectAccess clients must be domain members. Domains containing clients can belong to the same forest as the Remote Access server, or have a two-way trust with the Remote Access server forest or domain.
  - An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients. If a security group is not specified when configuring DirectAccess client settings, by default the client GPO is applied on all laptop computers in the Domain Computers security group. Note the following:

    We recommend that you create a security group for each domain that contains computers that will be configured as DirectAccess clients.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
Deployment	Remote Access (DirectAccess, Routing and Remote Access) Overview
Troubleshooting	Troubleshooting Remote Access documentation, when available.
Tools and settings	Remote Access PowerShell cmdlets
Community resources	RRAS Product Team blog | Remote Access TechNet Forum DirectAccess Wiki entries
Related technologies	How IPv6 works

Plan Deployment for Remote Management of DirectAccess Clients

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

The following topics provide planning steps for deploying a single Remote Access server running Windows Server 2012 that can be used for remote management of DirectAccess clients.

Step 1: Plan the Remote Access Infrastructure: This topic helps you plan the network topology and server settings, firewall requirements, certificate requirements, Domain Name System requirements, DirectAccess network location server and management servers' configuration, Active Directory requirements, and Group Policy Object creation.

Step 2: Plan the Remote Access Deployment: Plan client and server deployment strategies and infrastructure servers' configurations.

Next steps

After you complete these planning steps, you can begin the server deployment. For instructions, see Install and Configure Deployment for Remote Management of DirectAccess Clients.

Step 1: Plan the Remote Access Infrastructure

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This topic describes the steps for planning an infrastructure that you can use to set up a single Remote Access server for remote management of DirectAccess clients. The following table lists the steps, but these planning tasks do not need to be done in a specific order.

Task	Description
Plan network topology and settings	Decide where to place the Remote Access server (at the edge or behind a Network Address Translation (NAT) device or firewall), and plan IP addressing and routing.
Plan firewall requirements	Plan for allowing Remote Access through edge firewalls.
Plan certificate requirements	Decide if you will use Kerberos protocol or certificates for client authentication, and plan your website certificates. IP-HTTPS is a transition protocol that is used by DirectAccess clients to tunnel IPv6 traffic over IPv4 networks. Decide whether to authenticate IP-HTTPS for the server by using a certificate that is issued by a certification authority (CA), or by using a self-signed certificate that is issued automatically by the Remote Access server.
Plan DNS requirements	Plan the Domain Name System (DNS) settings for the Remote Access server, infrastructure servers, local name resolution options, and client connectivity.
Plan the network location server configuration	Decide where to place the network location server website in your organization (on the Remote Access server or an alternative server),

Task	Description
	and plan the certificate requirements if the network location server will be located on the Remote Access server. **Note** The network location server is used by DirectAccess clients to determine whether they are located on the internal network.
Plan management servers' configuration	Plan for management servers (such as update servers) that are used during remote client management. **Note** Administrators can remotely manage DirectAccess client computers that are located outside the corporate network by using the Internet.
Plan Active Directory requirements	Plan your domain controllers, your Active Directory requirements, client authentication, and multiple domain structure.
Plan Group Policy Object creation	Decide what GPOs are required in your organization and how to create and edit the GPOs.

Plan network topology and settings

When you plan your network, you need to consider the network adapter topology, settings for IP addressing, and requirements for ISATAP.

Plan network adapters and IP addressing

1. Identify the network adapter topology that you want to use. Remote Access can be set up with any of the following topologies:
   - With two network adapters: The Remote Access server is installed at the edge with one network adapter connected to the Internet and the other to the internal network.
   - With two network adapters: The Remote Access server is installed behind a NAT device, firewall, or router, with one network adapter connected to a perimeter network and the other to the internal network.

- With one network adapter: The Remote Access server is installed behind a NAT device, and the single network adapter is connected to the internal network.

2. Identify your IP addressing requirements:

   DirectAccess uses IPv6 with IPsec to create a secure connection between DirectAccess client computers and the internal corporate network. However, DirectAccess does not necessarily require connectivity to the IPv6 Internet or native IPv6 support on internal networks. Instead, it automatically configures and uses IPv6 transition technologies to tunnel IPv6 traffic across the IPv4 Internet (6to4, Teredo, or IP-HTTPS) and across your IPv4-only intranet (NAT64 or ISATAP). For an overview of these transition technologies, see the following resources:

   - IPv6 Transition Technologies
   - IP-HTTPS Tunneling Protocol Specification

3. Configure required adapters and addressing according to the following table. For deployments that are behind a NAT device using a single network adapter, configure your IP addresses by using only the **Internal network adapter** column.

	External network adapter	Internal network adapter[1, above]	Routing requirements
IPv4 Internet and IPv4 intranet	Configure the following: • Two static consecutive public IPv4 addresses with the appropriate subnet masks (required for Teredo only). • A default gateway IPv4 address for your Internet firewall or local Internet service provider (ISP) router. **Note** The Remote Access server requires two consecutive public IPv4 addresses so that it can act as a Teredo server and Windows-based Teredo clients can use the Remote Access server to detect the type of NAT device.	Configure the following: • An IPv4 intranet address with the appropriate subnet mask. • A connection-specific DNS suffix for your intranet namespace. A DNS server should also be configured on the internal interface. **Caution** Do not configure a default gateway on any intranet interfaces.	To configure the Remote Access server to reach all subnets on the internal IPv4 network, do the following: • List the IPv4 address spaces for all the locations on your intranet. • Use the `route add -p` or `netsh interface ipv4 add route` commands to add the IPv4 address spaces as static routes in the IPv4 routing table of the Remote Access server.
IPv6 Internet and IPv6 intranet	Configure the following: • Use the autoconfigured address configuration provided by your ISP. • Use the `route print` command to ensure that a default IPv6 route that points to the	Configure the following: If you are not using default preference levels, configure your intranet interfaces by using the `netsh interface ipv6 set InterfaceIndex ignoredefaultroutes=enabled` command. This command	If you have an IPv6 intranet, to configure the Remote Access server to reach all of the IPv6 locations, do the following:

	External network adapter	Internal network adapter[1, above]	Routing requirements
	ISP router exists in the IPv6 routing table. • Determine if the ISP and intranet routers are using default router preferences as described in RFC 4191, and if they are use a higher default preference than your local intranet routers. If both of these are true, no other configuration for the default route is required. The higher preference for the ISP router ensures that the active default IPv6 route of the Remote Access server points to the IPv6 Internet. Because the Remote Access server is an IPv6 router, if you have a native IPv6 infrastructure, the Internet interface can also reach the domain controllers on the intranet. In this case, add packet filters to the domain controller in the perimeter network that prevent connectivity to the IPv6 address of the Internet interface of the Remote Access server.	ensures that additional default routes that point to intranet routers will not be added to the IPv6 routing table. You can obtain the InterfaceIndex of your intranet interfaces from the display of the `netsh interface show interface` command.	• List the IPv6 address spaces for all the locations on your intranet. • Use the `netsh interface ipv6 add route` command to add the IPv6 address spaces as static routes in the IPv6 routing table of the Remote Access server.
IPv4 Internet and IPv6 intranet	The Remote Access server forwards default IPv6 route traffic by using the Microsoft 6to4 adapter		

	External network adapter	Internal network adapter[1, above]	Routing requirements
	interface to a 6to4 relay on the IPv4 Internet. When native IPv6 is not deployed in the corporate network, you can use the following command to configure a Remote Access server for the IPv4 address of the Microsoft 6to4 relay on the IPv4 Internet: `netsh interface ipv6 6to4 set relay name=<ipaddress> state=enabled`.		

Notes

- If the DirectAccess client has been assigned a public IPv4 address, it will use the 6to4 relay technology to connect to the intranet. If the client is assigned a private IPv4 address, it will use Teredo. If the DirectAccess client cannot connect to the DirectAccess server with 6to4 or Teredo, it will use IP-HTTPS.
- To use Teredo, you must configure two consecutive IP addresses on the external facing network adapter.
- You cannot use Teredo if the Remote Access server has only one network adapter.
- Native IPv6 client computers can connect to the Remote Access server over native IPv6, and no transition technology is required.

Plan ISATAP requirements

ISATAP is required for remote management of DirectAccessclients, so that DirectAccess management servers can connect to DirectAccess clients located on the Internet. ISATAP is not required to support connections that are initiated by DirectAccess client computers to IPv4 resources on the corporate network. NAT64/DNS64 is used for this purpose. If your deployment requires ISATAP, use the following table to identify your requirements.

ISATAP deployment scenario	Requirements
Existing native IPv6 intranet (no ISATAP is required)	With an existing native IPv6 infrastructure, you specify the prefix of the organization during Remote Access deployment, and the Remote Access server does not configure itself as an

ISATAP deployment scenario	Requirements
	ISATAP router. Do the following: 1. To ensure that DirectAccess clients are reachable from the intranet, you must modify your IPv6 routing so that default route traffic is forwarded to the Remote Access server. If your intranet IPv6 address space uses an address other than a single 48-bit IPv6 address prefix, you must specify the relevant organization IPv6 prefix during deployment. 2. If you are currently connected to the IPv6 Internet, you must configure your default route traffic so that it is forwarded to the Remote Access server, and then configure the appropriate connections and routes on the Remote Access server, so that the default route traffic is forwarded to the device that is connected to the IPv6 Internet.
Existing ISATAP deployment	If you have an existing ISATAP infrastructure, during deployment you are prompted for the 48-bit prefix of the organization, and the Remote Access server does not configure itself as an ISATAP router. To ensure that DirectAccess clients are reachable from the intranet, you must modify your IPv6 routing infrastructure so that default route traffic is forwarded to the Remote Access server. This change needs to be done on the existing ISATAP router to which the intranet clients must already be forwarding the default traffic.
No existing IPv6 connectivity	When the Remote Access setup wizard detects that the server has no native or ISATAP-based IPv6 connectivity, it automatically derives a 6to4-based 48-bit prefix for the intranet, and configures the Remote Access server as an ISATAP router to provide IPv6 connectivity to ISATAP hosts across your intranet. (A 6to4-based prefix is used only if the server has public addresses, otherwise the prefix is automatically generated from a unique local

ISATAP deployment scenario	Requirements
	address range.) To use ISATAP do the following: 1. Register the ISATAP name on a DNS server for each domain on which you want to enable ISATAP-based connectivity, so that the ISATAP name is resolvable by the internal DNS server to the internal IPv4 address of the Remote Access server. 2. By default, DNS servers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 block resolution of the ISATAP name by using the global query block list. To enable ISATAP, you must remove the ISATAP name from the block list. For more information, see Remove ISATAP from the DNS Global Query Block List. Windows-based ISATAP hosts that can resolve the ISATAP name automatically configure an address with the Remote Access server as following: 1. An ISATAP-based IPv6 address on an ISATAP tunneling interface 2. A 64-bit route that provides connectivity to the other ISATAP hosts on the intranet 3. A default IPv6 route that points to the Remote Access server. The default route ensures that intranet ISATAP hosts can reach DirectAccess clients When your Windows-based ISATAP hosts obtain an ISATAP-based IPv6 address, they begin to use ISATAP-encapsulated traffic to communicate if the destination is also an ISATAP host. Because ISATAP uses a single 64-bit subnet for the entire intranet, your communication goes from a segmented IPv4 model of communication, to a single subnet communication model with IPv6. This can affect the behavior of some Active Directory Domain Services (AD DS) and applications that rely on your Active Directory Sites and Services configuration. For example, if you used the

ISATAP deployment scenario	Requirements
	Active Directory Sites and Services snap-in to configure sites, IPv4-based subnets, and intersite transports for forwarding requests to servers within sites, this configuration is not used by ISATAP hosts. 1. To configure Active Directory Sites and Services for forwarding within sites for ISATAP hosts, for each IPv4 subnet object, you must configure an equivalent IPv6 subnet object, in which the IPv6 address prefix for the subnet expresses the same range of ISATAP host addresses as the IPv4 subnet. For example, for the IPv4 subnet 192.168.99.0/24 and the 64-bit ISATAP address prefix 2002:836b:1:8000::/64, the equivalent IPv6 address prefix for the IPv6 subnet object is 2002:836b:1:8000:0:5efe:192.168.99.0/120 . For an arbitrary IPv4 prefix length (set to 24 in the example), you can determine the corresponding IPv6 prefix length from the formula 96 + IPv4PrefixLength. 2. For the IPv6 addresses of DirectAccess clients, add the following: • For Teredo-based DirectAccess clients: An IPv6 subnet for the range 2001:0:WWXX:YYZZ::/64, in which WWXX:YYZZ is the colon-hexadecimal version of the first Internet-facing IPv4 address of the Remote Access server. . • For IP-HTTPS-based DirectAccess clients: An IPv6 subnet for the range 2002:WWXX:YYZZ:8100::/56, in which WWXX:YYZZ is the colon-hexadecimal version of the first Internet-facing IPv4 address (w.x.y.z) of the Remote Access server. . • For 6to4-based DirectAccess clients: A series of 6to4-based IPv6 prefixes that begin with 2002: and represent the regional, public IPv4 address prefixes that are administered by Internet Assigned Numbers Authority (IANA)

ISATAP deployment scenario	Requirements
	and regional registries. The 6to4-based prefix for a public IPv4 address prefix w.x.y.z/n is 2002:WWXX:YYZZ::/[16+n], in which WWXX:YYZZ is the colon-hexadecimal version of w.x.y.z. For example, the 7.0.0.0/8 range is administered by American Registry for Internet Numbers (ARIN) for North America. The corresponding 6to4-based prefix for this public IPv6 address range is 2002:700::/24. For information about the IPv4 public address space, see IANA IPv4 Address Space Registry. .

Important
Ensure that you do not have public IP addresses on the internal interface of the DirectAccess server. If you have public IP address on the internal interface, connectivity through ISATAP may fail.

Plan firewall requirements

If the Remote Access server is behind an edge firewall, the following exceptions will be required for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- For IP-HTTPS: Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound.
- For Teredo traffic: User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound.
- For 6to4 traffic: IP Protocol 41 inbound and outbound.

  Notes
  For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server.

  For IP-HTTPS the exceptions need to be applied on the address that is registered on the public DNS server.
- If you are deploying Remote Access with a single network adapter and installing the network location server on the Remote Access server, TCP port 62000.

  Note
  This exemption is on the Remote Access server, and the previous exemptions are on the edge firewall.

The following exceptions are required for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.
- ICMPv6 traffic inbound and outbound (only when using Teredo).

When you are using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- For ISATAP: Protocol 41 inbound and outbound
- For all IPv4/IPv6 traffic: TCP/UD
- For Teredo: ICMP for all IPv4/IPv6 traffic

Plan certificate requirements

There are three scenarios that require certificates when you deploy a single Remote Access server.

- **IPsec authentication**: Certificate requirements for IPsec include a computer certificate that is used by DirectAccess client computers when they establish the IPsec connection with the Remote Access server, and a computer certificate that is used by Remote Access servers to establish IPsec connections with DirectAccess clients.

  For DirectAccess in Windows Server 2012, the use of these IPsec certificates is not mandatory. As an alternative, the Remote Access server can act as a proxy for Kerberos authentication without requiring certificates. If Kerberos authentication is used, it works over SSL, and the Kerberos protocol uses the certificate that was configured for IP-HTTPS. Some enterprise scenarios (including multisite deployment and one-time password client authentication) require the use of certificate authentication, and not Kerberos authentication.
- **IP-HTTPS server**: When you configure Remote Access, the Remote Access server is automatically configured to act as the IP-HTTPS web listener. The IP-HTTPS site requires a website certificate, and client computers must be able to contact the certificate revocation list (CRL) site for the certificate.
- **Network location server**: The network location server is a website that is used to detect whether client computers are located in the corporate network. The network location server requires a website certificate. DirectAccess clients must be able to contact the CRL site for the certificate.

The certification authority (CA) requirements for each of these scenarios is summarized in the following table.

IPsec authentication	IP-HTTPS server	Network location server
An internal CA is required to issue computer certificates to the Remote Access server and clients for IPsec authentication when you don't use the	Internal CA: You can use an internal CA to issue the IP-HTTPS certificate; however, you must make sure that the CRL distribution point is	Internal CA: You can use an internal CA to issue the network location server website certificate. Make sure that the CRL distribution point

IPsec authentication	IP-HTTPS server	Network location server
Kerberos protocol for authentication.	available externally.	is highly available from the internal network.
	Self-signed certificate: You can use a self-signed certificate for the IP-HTTPS server. A self-signed certificate cannot be used in a multisite deployment.	Self-signed certificate: You can use a self-signed certificate for the network location server website; however, you cannot use a self-signed certificate in multisite deployments.
	Public CA: We recommend that you use a public CA to issue the IP-HTTPS certificate, this ensures that the CRL distribution point is available externally.	

Plan computer certificates for IPsec authentication

If you are using certificate-based IPsec authentication, the Remote Access server and clients are required to obtain a computer certificate. The simplest way to install the certificates is to use Group Policy to configure automatic enrollment for computer certificates. This ensures that all domain members obtain a certificate from an enterprise CA. If you do not have an enterprise CA set up in your organization, see Active Directory Certificate Services.

This certificate has the following requirements:

- The certificate should have client authentication extended key usage (EKU).
- The client and the server certificates should relate to the same root certificate. This root certificate must be selected in the DirectAccess configuration settings.

Plan certificates for IP-HTTPS

The Remote Access server acts as an IP-HTTPS listener, and you must manually install an HTTPS website certificate on the server. Consider the following when you are planning:

- Using a public CA is recommended, so that CRLs are readily available.
- In the subject field, specify the IPv4 address of the Internet adapter of Remote Access server or the FQDN of the IP-HTTPS URL (the ConnectTo address). If the Remote Access server is located behind a NAT device, the public name or address of the NAT device should be specified.
- The common name of the certificate should match the name of the IP-HTTPS site.
- For the **Enhanced Key Usage** field, use the Server Authentication object identifier (OID).
- For the **CRL Distribution Points** field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.

Note

This is only required for clients running Windows 7.

- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcard characters in the name.

Plan website certificates for the network location server

Consider the following when you are planning the network location server website:

- In the **Subject** field, specify an IP address of the intranet interface of the network location server or the FQDN of the network location URL.
- For the **Enhanced Key Usage** field, use the Server Authentication OID.
- For the **CRL Distribution Points** field, use a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet. This CRL distribution point should not be accessible from outside the internal network.

Note

Ensure that the certificates for IP-HTTPS and network location server have a subject name. If the certificate uses an alternative name, it will not be accepted by the Remote Access Wizard.

Plan DNS requirements

This section explains the DNS requirements for clients and servers in a Remote Access deployment.

DirectAccess client requests

DNS is used to resolve requests from DirectAccess client computers that are not located on the internal network. DirectAccess clients attempt to connect to the DirectAccess network location server to determine whether they are located on the Internet or on the corporate network.

- If the connection is successful, clients are determined to be on the intranet, DirectAccess is not used, and client requests are resolved by using the DNS server that is configured on the network adapter of the client computer.
- If the connection does not succeed, clients are assumed to be on the Internet. DirectAccess clients will use the name resolution policy table (NRPT) to determine which DNS server to use when resolving name requests. You can specify that clients should use DirectAccess DNS64 to resolve names, or an alternative internal DNS server.

When performing name resolution, the NRPT is used by DirectAccess clients to identify how to handle a request. Clients request an FQDN or single-label name such as http://internal. If a single-label name is requested, a DNS suffix is appended to make an FQDN. If the DNS query matches an entry in the NRPT and DNS4 or an intranet DNS server is specified for the entry, the query is sent for name resolution by using the specified server. If a match exists but no DNS server is specified, an exemption rule and normal name resolution is applied.

When a new suffix is added to the NRPT in the Remote Access Management console, the default DNS servers for the suffix can be automatically discovered by clicking the **Detect** button. Automatic detection works as follows:

- If the corporate network is IPv4-based, or it uses IPv4 and IPv6, the default address is the DNS64 address of the internal adapter on the Remote Access server.
- If the corporate network is IPv6-based, the default address is the IPv6 address of DNS servers in the corporate network.

Infrastructure servers

- **Network location server**

  DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, and they must be prevented from resolving the name when they are located on the Internet. To ensure that this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT. In addition, when you configure Remote Access, the following rules are created automatically:

  - A DNS suffix rule for root domain or the domain name of the Remote Access server, and the IPv6 addresses that correspond to the intranet DNS servers that are configured on the Remote Access server. For example, if the Remote Access server is a member of the corp.contoso.com domain, a rule is created for the corp.contoso.com DNS suffix.
  - An exemption rule for the FQDN of the network location server. For example, if the network location server URL is https://nls.corp.contoso.com, an exemption rule is created for the FQDN nls.corp.contoso.com.
- **IP-HTTPS server**

  The Remote Access server acts as an IP-HTTPS listener and uses its server certificate to authenticate to IP-HTTPS clients. The IP-HTTPS name must be resolvable by DirectAccess clients that use public DNS servers.

Connectivity verifiers

Remote Access creates a default web probe that is used by DirectAccess client computers to verify connectivity to the internal network. To ensure that the probe works as expected, the following names must be registered manually in DNS:

- **directaccess-webprobehost** should resolve to the internal IPv4 address of the Remote Access server, or to the IPv6 address in an IPv6-only environment.
- **directaccess-corpconnectivityhost** should resolve to the local host (loopback) address. You should create A and AAAA records. The value of the A record is 127.0.0.1, and the value of the AAAA record is constructed from the NAT64 prefix with the last 32 bits as 127.0.0.1. The NAT64 prefix can be retrieved by running the **Get-netnatTransitionConfiguration** Windows PowerShell cmdlet.

  Note

  This is valid only in IPv4-only environments. In an IPv4 plus IPv6 or an IPv6-only environment, create only a AAAA record with the loopback IP address ::1.

You can create additional connectivity verifiers by using other web addresses over HTTP or PING. For each connectivity verifier, a DNS entry must exist.

DNS server requirements

- For DirectAccess clients, you must use a DNS server running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, or any DNS server that supports IPv6.
- You should use a DNS server that supports dynamic updates. You can use DNS servers that do not support dynamic updates, but then entries must be manually updated.
- The FQDN for your CRL distribution points must be resolvable by using Internet DNS servers. For example, if URL http://crl.contoso.com/crld/corp-DC1-CA.crl is in the **CRL Distribution Points** field of the IP-HTTPS certificate of the Remote Access server, you must ensure that the FQDN crld.contoso.com is resolvable by using Internet DNS servers.

Plan for local name resolution

Consider the following when you are planning for local name resolution:

NRPT

You may need to create additional name resolution policy table (NRPT) rules in the following situations:

- You need to add more DNS suffixes for your intranet namespace.
- If the FQDNs of your CRL distribution points are based on your intranet namespace, you must add exemption rules for the FQDNs of the CRL distribution points.
- If you have a split-brain DNS environment, you must add exemption rules for the names of resources for which you want DirectAccess clients that are located on the Internet to access the Internet version, rather than the intranet version.
- If you are redirecting traffic to an external website through your intranet web proxy servers, the external website is available only from the intranet. It uses the addresses of your web proxy servers to permit the inbound requests. In this situation, add an exemption rule for the FQDN of the external website, and specify that the rule uses your intranet web proxy server rather than the IPv6 addresses of intranet DNS servers.

  For example, let's say that you are testing an external website named test.contoso.com. This name is not resolvable through Internet DNS servers, but the Contoso web proxy server knows how to resolve the name and how to direct requests for the website to the external web server. To prevent users who are not on the Contoso intranet from accessing the site, the external website allows requests only from the IPv4 Internet address of the Contoso web proxy. Thus, intranet users can access the website because they are using the Contoso web proxy, but DirectAccess users cannot because they are not using the Contoso web proxy. By configuring an NRPT exemption rule for test.contoso.com that uses the Contoso web proxy, webpage requests for test.contoso.com are routed to the intranet web proxy server over the IPv4 Internet.

Single label names

Single label names, such as http://paycheck, are sometimes used for intranet servers. If a single label name is requested and a DNS suffix search list is configured, the DNS suffixes in the list will be appended to the single label name. For example, when a user on a computer that is a member of the corp.contoso.com domain types http://paycheck in the web browser, the FQDN that is constructed as the name is paycheck.corp.contoso.com. By default, the appended suffix is based on the primary DNS suffix of the client computer.

Note

In a disjointed name space scenario (where one or more domain computers has a DNS suffix that does not match the Active Directory domain to which the computers are members), you should ensure that the search list is customized to include all the required suffixes. By default, the Remote Access Wizard, configures the Active Directory DNS name as the primary DNS suffix on the client. Make sure to add the DNS suffix that is used by clients for name resolution.

If multiple domains and Windows Internet Name Service (WINS) are deployed in your organization, and you are connecting remotely, single-names can be resolved as follows:

- By deploying a WINS forward lookup zone in the DNS. When trying to resolve computername.dns.zone1.corp.contoso.com, the request is directed to the WINS server that is only using the computer name. The client thinks it is issuing a regular DNS A records request, but it is actually a NetBIOS request.

  For more information, see Managing a Forward Lookup Zone.
- By adding a DNS suffix (for example, dns.zone1.corp.contoso.com) to the default domain GPO.

Split-brain DNS

Split-brain DNS refers to the use of the same DNS domain for Internet and intranet name resolution.

For split-brain DNS deployments, you must list the FQDNs that are duplicated on the Internet and intranet, and decide which resources the DirectAccess client should reach—the intranet or the Internet version. When you want DirectAccess clients to reach the Internet version, you must add the corresponding FQDN as an exemption rule to the NRPT for each resource.

In a split-brain DNS environment, if you want both versions of the resource to be available, configure your intranet resources with names that do not duplicate the names that are used on the Internet. Then instruct your users to use the alternate name when they access the resource on the intranet. For example, configure www.internal.contoso.com for the internal name of www.contoso.com.

In a non-split-brain DNS environment, the Internet namespace is different from the intranet namespace. For example, the Contoso Corporation uses contoso.com on the Internet and corp.contoso.com on the intranet. Because all intranet resources use the corp.contoso.com DNS suffix, the NRPT rule for corp.contoso.com routes all DNS name queries for intranet resources to

intranet DNS servers. DNS queries for names with the contoso.com suffix do not match the corp.contoso.com intranet namespace rule in the NRPT, and they are sent to Internet DNS servers. With a non-split-brain DNS deployment, because there is no duplication of FQDNs for intranet and Internet resources, there is no additional configuration needed for the NRPT. DirectAccess clients can access both Internet and intranet resources for their organization.

Plan local name resolution behavior for DirectAccess clients

If a name cannot be resolved with DNS, the DNS Client service in Windows Server 2012, Windows 8, Windows Server 2008 R2, and Windows 7 can use local name resolution, with the Link-Local Multicast Name Resolution (LLMNR) and NetBIOS over TCP/IP protocols, to resolve the name on the local subnet. Local name resolution is typically needed for peer-to-peer connectivity when the computer is located on private networks, such as single subnet home networks.

When the DNS Client service performs local name resolution for intranet server names, and the computer is connected to a shared subnet on the Internet, malicious users can capture LLMNR and NetBIOS over TCP/IP messages to determine intranet server names. On the DNS page of the Infrastructure Server Setup Wizard, you can configure the local name resolution behavior based on the types of responses received from intranet DNS servers. The following options are available:

- **Use local name resolution if the name does not exist in DNS**: This option is the most secure because the DirectAccess client performs local name resolution only for server names that cannot be resolved by intranet DNS servers. If the intranet DNS servers can be reached, the names of intranet servers are resolved. If the intranet DNS servers cannot be reached, or if there are other types of DNS errors, the intranet server names are not leaked to the subnet through local name resolution.
- **Use local name resolution if the name does not exist in DNS or DNS servers are unreachable when the client computer is on a private network (recommended)**: This option is recommended because it allows the use of local name resolution on a private network only when the intranet DNS servers are unreachable.
- **Use local name resolution for any kind of DNS resolution error (least secure)**: This is the least secure option because the names of intranet network servers can be leaked to the local subnet through local name resolution.

Plan the network location server configuration

The network location server is a website that is used to detect whether DirectAccess clients are located in the corporate network. Clients in the corporate network do not use DirectAccess to reach internal resources; but instead, they connect directly.

The network location server website can be hosted on the Remote Access server or on another server in your organization. If you host the network location server on the Remote Access server, the website is created automatically when you deploy Remote Access. If you host the network location server on another server running a Windows operating system, you must make sure that Internet Information Services (IIS) is installed on that server, and that the website is created. Remote Access does not configure settings on the network location server.

Make sure that the network location server website meets the following requirements:

- Has an HTTPS server certificate.
- Has high availability to computers on the internal network.
- Is not accessible to DirectAccess client computers on the Internet.

In addition, consider the following requirements for clients when you are setting up your network location server website:

- DirectAccess client computers must trust the CA that issued the server certificate to the network location server website.
- DirectAccess client computers on the internal network must be able to resolve the name of the network location server site.

Plan certificates for the network location server

When you obtain the website certificate to use for the network location server, consider the following:

- In the **Subject** field, specify the IP address of the intranet interface of the network location server or the FQDN of the network location URL.
- For the **Enhanced Key Usage** field, use the Server Authentication OID.
- The network location server certificate must be checked against a certificate revocation list (CRL). For the **CRL Distribution Points** field, use a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet. This CRL distribution point should not be accessible from outside the internal network.

Plan DNS for the network location server

DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, but must be prevented from resolving the name when they are located on the Internet. To ensure this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT.

Plan management servers' configuration

DirectAccess clients initiate communication with management servers that provide services such as Windows Update and antivirus updates. DirectAccess clients also use the Kerberos protocol to authenticate to domain controllers before they access the internal network. During remote management of DirectAccess clients, management servers communicate with client computers to perform management functions such as software or hardware inventory assessments. Remote Access can automatically discover some management servers, including:

- Domain controllers: Automatic discovery of domain controllers is performed for the domains that contain client computers and for all domains in the same forest as the Remote Access server.
- System Center Configuration Manager servers

Domain controllers and System Center Configuration Manager servers are automatically detected the first time DirectAccess is configured. The detected domain controllers are not displayed in the

console, but settings can be retrieved using Windows PowerShell cmdlets. If domain controller or System Center Configuration Manager servers are modified, clicking **Update Management Servers** in the console refreshes the management server list.

Management server requirements

- Management servers must be accessible over the infrastructure tunnel. When you configure Remote Access, adding servers to the management servers list automatically makes them accessible over this tunnel.
- Management servers that initiate connections to DirectAccess clients must fully support IPv6, by means of a native IPv6 address or by using an address that is assigned by ISATAP.

Plan Active Directory requirements

Remote Access uses Active Directory as follows:

- **Authentication**: The infrastructure tunnel uses NTLMv2 authentication for the computer account that is connecting to the Remote Access server, and the account must be in an Active Directory domain. The intranet tunnel uses Kerberos authentication for the user to create the intranet tunnel.
- **Group Policy Objects**: Remote Access gathers configuration settings into Group Policy Objects (GPOs), which are applied to Remote Access servers, clients, and internal application servers.
- **Security groups**: Remote Access uses security groups to gather and identify DirectAccess client computers. GPOs are applied to the required security groups.

When you plan an Active Directory environment for a Remote Access deployment, consider the following requirements:

- At least one domain controller is installed on the Windows Server 2012, Windows Server 2008 R2 Windows Server 2008, or Windows Server 2003 operating system.

  If the domain controller is on a perimeter network (and therefore reachable from the Internet-facing network adapter of Remote Access server), prevent the Remote Access server from reaching it. You need to add packet filters on the domain controller to prevent connectivity to the IP address of the Internet adapter.
- The Remote Access server must be a domain member.
- DirectAccess clients must be domain members. Clients can belong to:
  - Any domain in the same forest as the Remote Access server.
  - Any domain that has a two-way trust with the Remote Access server domain.
  - Any domain in a forest that has a two-way trust with the forest of the Remote Access server domain.

Notes

- The Remote Access server cannot be a domain controller.
- The Active Directory domain controller that is used for Remote Access must not be reachable from the external Internet adapter of the Remote Access server (the adapter must not be in the domain profile of Windows Firewall).

Plan client authentication

In Remote Access in Windows Server 2012, you can choose between using built-in Kerberos authentication, which uses user names and passwords, or using certificates for IPsec computer authentication.

Kerberos authentication: When you choose to use Active Directory credentials for authentication, DirectAccess first uses Kerberos authentication for the computer, and then it uses Kerberos authentication for the user. When using this mode of authentication, DirectAccess uses a single security tunnel that provides access to the DNS server, the domain controller, and any other server on the internal network

IPsec authentication: When you choose to use two-factor authentication or Network Access Protection, DirectAccess uses two security tunnels. The Remote Access Setup Wizard configures connection security rules in Windows Firewall with Advanced Security. These rules specify the following credentials when negotiating IPsec security to the Remote Access server:

- The infrastructure tunnel uses computer certificate credentials for the first authentication and user (NTLMv2) credentials for the second authentication. User credentials force the use of Authenticated Internet Protocol (AuthIP), and they provide access to a DNS server and domain controller before the DirectAccess client can use Kerberos credentials for the intranet tunnel.
- The intranet tunnel uses computer certificate credentials for the first authentication and user (Kerberos V5) credentials for the second authentication.

Plan multiple domains

The management servers list should include domain controllers from all domains that contain security groups that include DirectAccess client computers. It should contain all domains that contain user accounts that might use computers configured as DirectAccess clients. This ensures that users who are not located in the same domain as the client computer they are using are authenticated with a domain controller in the user domain.

This authentication is automatic if the domains are in the same forest. If there is a security group with client computers or application servers that are in different forests, the domain controllers of those forests are not detected automatically. Forests are also not detected automatically. You can run the task **Update Management Servers** in the **Remote Access Management** to detect these domain controllers.

Where possible, common domain name suffixes should be added to the NRPT during Remote Access deployment. For example, if you have two domains, domain1.corp.contoso.com and domain2.corp.contoso.com, instead of adding two entries into the NRPT, you can add a common DNS suffix entry, where the domain name suffix is corp.contoso.com. This happens automatically for domains in the same root. Domains that are not in the same root must be added manually.

Plan Group Policy Object creation

When you configure Remote Access, DirectAccess settings are collected into Group Policy Objects (GPOs). Two GPOs are populated with DirectAccess settings, and they are distributed as follows:

- **DirectAccess client GPO**: This GPO contains client settings, including IPv6 transition technology settings, NRPT entries, and connection security rules for Windows Firewall with Advanced Security. The GPO is applied to the security groups that are specified for the client computers.
- **DirectAccess server GPO**: This GPO contains the DirectAccess configuration settings that are applied to any server that you configured as a Remote Access server in your deployment. It also contains connection security rules for Windows Firewall with Advanced Security.

Note

Configuration of application servers is not supported in remote management of DirectAccess clients because clients cannot access the internal network of the DirectAccess server where the application servers reside. Step 4 in the Remote Access Setup configuration screen is unavailable for this type of configuration.

You can configure GPOs automatically or manually.

Automatically: When you specify that GPOs are created automatically, a default name is specified for each GPO.

Manually: You can use GPOs that have been predefined by the Active Directory administrator.

When you configure your GPOs, consider the following warnings:

- After DirectAccess is configured to use specific GPOs, it cannot be configured to use different GPOs.
- Use the following procedure to back up all Remote Access Group Policy Objects before you run DirectAccess cmdlets:

  Back up and Restore Remote Access Configuration.
- Whether you are using automatically or manually configured GPOs, you need to add a policy for slow link detection if your clients will use 3G. The path for **Policy: Configure Group Policy slow link detection** is:

  Computer configuration/Polices/Administrative Templates/System/Group Policy.
- If the correct permissions for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue, but linking will not occur. If this warning is issued, links will not be created automatically, even if the permissions are added later. Instead the administrator needs to create the links manually.

Automatically created GPOs

Consider the following when using automatically created GPOs:

Automatically created GPOS are applied according to the location and link target, as follows:

- For the DirectAccess server GPO, the location and link target point to the domain that contains the Remote Access server.

- When client and application server GPOs are created, the location is set to a single domain. The GPO name is looked up in each domain, and the domain is filled with DirectAccess settings if it exists.
- The link target is set to the root of the domain in which the GPO was created. A GPO is created for each domain that contains client computers or application servers, and the GPO is linked to the root of its respective domain.

When using automatically created GPOs to apply DirectAccess settings, the Remote Access server administrator requires the following permissions:

- Permissions to create GPOs for each domain.
- Permissions to link to all the selected client domain roots.
- Permissions to link to the server GPO domain roots.
- Security permissions to create, edit, delete, and modify the GPOs.
- GPO read permissions for each required domain. This permission is not required, but it is recommended because it enables Remote Access to verify that GPOs with duplicate names do not exist when GPOs are being created.

Manually created GPOs

Consider the following when using manually created GPOs:

- The GPOs should exist before running the Remote Access Setup Wizard.
- To apply DirectAccess settings, the Remote Access server administrator requires full security permissions to create, edit, delete, and modify the manually created GPOs.
- A search is made for a link to the GPO in the entire domain. If the GPO is not linked in the domain, a link is automatically created in the domain root. If the required permissions to create the link are not available, a warning is issued.

Recovering from a deleted GPO

If a GPO on a Remote Access server, client, or application server has been deleted by accident, the following error message will appear: **GPO <GPO name> cannot be found**.

If a backup is available, you can restore the GPO from the backup. If there is no backup available, you must remove the configuration settings and configure them again.

To remove configuration settings

1. Run the Windows PowerShell cmdlet **Uninstall-RemoteAccess**.
2. Open **Remote Access Management**.
3. You will see an error message that the GPO is not found. Click **Remove configuration settings**. After completion, the server will be restored to an unconfigured state, and you can reconfigure the settings.

See also

- Manage DirectAccess Clients Remotely

- Step 2: Plan the Remote Access Deployment

Step 2: Plan the Remote Access Deployment

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

After you plan for the infrastructure that you intend to use to set up your single Remote Access server for remote management of DirectAccess clients, you are ready to plan the settings that the Remote Access Setup Wizard will use.

Note
Before you continue with these tasks, see Step 1: Plan the Remote Access Infrastructure.

Task	Description
Plan a client deployment strategy	Decide which managed computers will be configured as DirectAccess clients.
Plan a Remote Access server deployment strategy	Plan how to deploy the Remote Access server.
Plan the infrastructure servers' configurations	Plan the infrastructure servers in your Remote Access deployment, including the DirectAccess network location server, DNS servers, and DirectAccess management servers.

Plan a client deployment strategy

There are three decisions to make when you are planning your client deployment:

1. Will DirectAccess be available to mobile computers only, or to every computer in a specified security group?

   When you configure DirectAccess clients in the DirectAccess Client Setup Wizard, you can choose to allow only mobile computers in specified security groups to connect to the server by using DirectAccess. If you restrict access to mobile computers, Remote Access automatically configures a WMI filter to ensure that the DirectAccess client GPO is applied only to mobile computers in the specified security groups. The Remote Access administrator requires permissions to create or modify group policy WMI filters to enable this setting.

2. What security groups will contain the DirectAccess client computers?

   DirectAccess settings are contained in the DirectAccess client Group Policy Object (GPO). The GPO is applied to computers that are part of the security groups that you specify in the DirectAccess Client Setup Wizard. You can specify security groups that are contained in any supported domain. For more information, see **Plan Active Directory requirements**.

Before you configure Remote Access, you need to create the security groups. You can add computers to the security group after you complete the Remote Access deployment. However, if you add client computers that reside in a different domain than the security group, the client GPO will not be applied to those clients. For example, if you created SG1 in domain A for DirectAccess clients, and you later add clients from domain B to this group, the client GPO will not be applied to clients in domain B.

To avoid this issue, create a new client security group for each domain that contains client computers. Alternatively, if you do not want to create a new security group, run the **Add-DAClient** Windows PowerShell cmdlet with the name of the new GPO for the new domain.

3. What settings will you configure for the DirectAccess Network Connectivity Assistant?

   The DirectAccess Network Connectivity Assistant runs on client computers and provides additional information about the DirectAccess connection to end users. In the DirectAccess Client Setup Wizard, you can configure the following:

   - **Connectivity verifiers**

     A default web probe is created that clients use to validate connectivity to the internal network. The default name is http://directaccess-WebProbeHost.<domain_name>. The name should be registered manually in DNS. You can create other connectivity verifiers that use other web addresses over HTTP or PING. For each connectivity verifier, a DNS entry must exist.
   - **Help Desk email address**

     If end users experience DirectAccess connectivity issues, they can send an email that contains diagnostic information to the Remote Access administrator, who can troubleshoot the issue.
   - **DirectAccess connection name**

     You can specify a DirectAccess connection name to help end users identify the DirectAccess connection on their computer.
   - **Allow DirectAccess clients to use local name resolution**

     Clients require a means of resolving names locally. If you allow DirectAccess clients to use local name resolution, end users can use local DNS servers to resolve names. When end users choose to use local DNS servers for name resolution, DirectAccess does not send resolution requests for single label names to the internal corporate DNS server. It uses local name resolution instead (by using the Link-Local Multicast Name Resolution (LLMNR) and NetBios over TCP/IP protocols).

Plan a Remote Access server deployment strategy

Decisions that you need to make when you are planning to deploy your Remote Access server include:

- **Network topology**

  There are two topologies available when deploying a Remote Access server:

- **Two adapters**: With two network adapters, Remote Access can be configured with one network adapter connected directly to the Internet and the other connected to the internal network. Or alternatively, the server is installed behind an edge device, such as a firewall or a router. In this configuration one network adapter is connected to the perimeter network, and the other is connected to the internal network.
- **Single network adapter**: In this configuration the Remote Access server is installed behind an edge device, such as a firewall or a router. The network adapter is connected to the internal network.

For more information, see Plan network topology and server settings.

- **Network adapters**

  The Remote Access Server Setup Wizard automatically detects the network adapters that are configured on the Remote Access server. You must make sure that the correct adapters are selected.

- **IP-HTTPS certificate**

  The Remote Access Server Setup Wizard automatically detects a certificate that is suitable for the IP-HTTPS connection. The subject name of the certificate that you select must match the ConnectTo address. If you are using self-signed certificates, you can choose to use a certificate that is created automatically by the Remote Access server.

- **IPv6 prefixes**

  If the Remote Access Server Setup Wizard detects that IPv6 has been deployed on the network adapters, it automatically populates IPv6 prefixes for the internal network, an IPv6 prefix to assign to DirectAccess client computers, and an IPv6 prefix to assign to VPN client computers. If the automatically generated prefixes are not correct for your native IPv6 or ISATAP infrastructure, you must manually change them. For more information, see **Plan network topology and server settings**.

- **Authentication**

  You can choose one of the following methods for authenticating DirectAccess clients to the Remote Access server:

  - **User authentication**: You can enable users to authenticate with Active Directory credentials or with two-factor authentication.
  - **Computer authentication**: You can configure computer authentication to use certificates. Or the Remote Access server can act as a proxy for Kerberos authentication without requiring certificates.. For more information, see Planning certification authorities and certificates.
  - **Windows 7 clients** By default, client computers running Windows 7 cannot connect to a Remote Access deployment running Windows Server 2012. If you have clients running Windows 7 in your organization that require remote access to internal resources, you can allow them to connect. Any client computers that you want to allow to access internal resources must be a member of a security group that you specify in the DirectAccess Client Setup Wizard.

    Note

Allowing clients running Windows 7 to connect by using DirectAccess requires you to use computer certificate authentication.

- **VPN configuration**

  Before you configure Remote Access, decide if you are going to provide VPN access to remote clients. You should provide VPN access if you have client computers in your organization that do not support DirectAccess connectivity (for example, they are unmanaged or they run an operating system for which DirectAccess is not supported). The Remote Access Server Setup Wizard allows you to configure how IP addresses are assigned (by using DHCP or from a static address pool) and how VPN clients are authenticated (by using Active Directory or a RADIUS server).

Plan the infrastructure servers' configurations

Remote Access requires three types of infrastructure servers:

- **Network location server**: For more information, see **Plan the network location server configuration**.
- **DNS servers**: For more information, see **Plan DNS requirements**.
- **Management servers**: For more information, see **Plan management servers' configurations**.

See also

- Step 1: Plan the Remote Access Infrastructure

Install and Configure Deployment for Remote Management of DirectAccess Clients

Remote_access_combo

This topic introduces the configuration steps that are required to deploy a single Remote Access server running Windows Server 2012 that can be used for remote management of DirectAccess clients.

- Step 1: Configure the Remote Access Infrastructure: This topic describes how to configure network and server settings, certificate requirements, DNS settings, network location server deployment, DirectAccess management servers, Active Directory settings, and Group Policy Objects.
- Step 2: Configure the Remote Access Server: This topic describes how to configure DirectAccess client computers, server settings, and infrastructure and application servers.
- Step 3: Verify the Deployment: This topic describes how to verify the deployment.

Important

Before you start the deployment, complete the planning steps that are described in Plan Deployment for Remote Management of DirectAccess Clients.

See also

Manage DirectAccess Clients Remotely

Step 1: Configure the Remote Access Infrastructure

Remote_access_combo

This topic describes how to configure the infrastructure that is required for an advanced Remote Access deployment using a single Remote Access server in a mixed IPv4 and IPv6 environment. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 1: Plan the Remote Access Infrastructure.

Task	Description
Configure server network settings	Configure the server network settings on the Remote Access server.
Configure routing in the corporate network	Configure routing in the corporate network to make sure traffic is appropriately routed.
Configure firewalls	Configure additional firewalls, if required.
Configure CAs and certificates	Configure a certification authority (CA), if required, and any other certificate templates required in the deployment.
Configure the DNS server	Configure DNS settings for the Remote Access server.
Configure Active Directory	Join client computers and the Remote Access server to the Active Directory domain.
Configure GPOs	Configure Group Policy Objects (GPOs) for the deployment, if required.
Configure security groups	Configure security groups that will contain DirectAccess client computers, and any other security groups that are required in the deployment.
Configure the network location server	Configure the network location server, including installing the network location server website

Task	Description
	certificate.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Configure server network settings

Depending on if you decide to place the Remote Access server at the edge or behind a Network Address Translation (NAT) device, the following network interface address settings are required for a single server deployment in an environment with IPv4 and IPv6. All IP addresses are configured by using **Change adapter settings** in the **Windows Networking and Sharing Center**.

Edge topology:

Requires the following:

- Two Internet-facing consecutive public static IPv4 or IPv6 addresses.

  Note
  Two consecutive public IPv4 addresses are required for Teredo. If you are not using Teredo, you can configure a single public static IPv4 address.

- A single internal static IPv4 or IPv6 address.

Behind NAT device (two network adapters):

Requires a single internal network-facing static IPv4 or IPv6 address.

Behind NAT device (one network adapter):

Requires a single static IPv4 or IPv6 address.

If the Remote Access server has two network adapters (one for the domain profile and the other for a public or private profile), but you are using a single network adapter topology, the recommendation is as follows:

1. Ensure that the second network adapter is also classified in the domain profile.
2. If the second network adapter cannot be configured for the domain profile for any reason, the DirectAccess IPsec policy must be manually scoped to all profiles by using the following Windows PowerShell command:

   ```
   $gposession = Open-NetGPO -PolicyStore <Name of the server GPO>
   Set-NetIPsecRule -DisplayName <Name of the IPsec policy> -
   GPOSession $gposession -Profile Any
   Save-NetGPO -GPOSession $gposession
   ```

   The names of the IPsec policies to use in this command are **DirectAccess-DaServerToInfra** and **DirectAccess-DaServerToCorp**.

Configure routing in the corporate network

Configure routing in the corporate network as follows:

- When native IPv6 is deployed in the organization, add a route so that the routers on the internal network route IPv6 traffic back through the Remote Access server.
- Manually configure organization IPv4 and IPv6 routes on the Remote Access servers. Add a published route so that all traffic with an (/48) IPv6 prefix is forwarded to the internal network. In addition, for IPv4 traffic, add explicit routes so that IPv4 traffic is forwarded to the internal network.

Configure firewalls

Depending on the network settings you chose, when you use additional firewalls in your deployment, apply the following firewall exceptions for Remote Access traffic:

Remote Access server on IPv4 Internet

Apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- **Teredo traffic**

  User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound. Apply this exemption for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server.
- **6to4 traffic**

  IP Protocol 41 inbound and outbound. Apply this exemption for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server.
- **IP-HTTPS traffic**

  Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required. Apply these exemptions only for the address to which the external name of the server resolves.

  Note

  This exemption is configured on the Remote Access server. All the other exemptions are configured on the edge firewall.

Remote Access server on IPv6 Internet

Apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.

- Internet Control Message Protocol for IPv6 (ICMPv6) traffic inbound and outbound – for Teredo implementations only.

Remote Access traffic

Apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP: Protocol 41 inbound and outbound
- TCP/UDP for all IPv4 or IPv6 traffic
- ICMP for all IPv4 or IPv6 traffic

Configure CAs and certificates

With Remote Access in Windows Server 2012, you to choose between using certificates for computer authentication or using a built-in Kerberos authentication that uses user names and passwords. You must also configure an IP-HTTPS certificate on the Remote Access server. This section explains how to configure these certificates.

For information about setting up a public key infrastructure (PKI), see Active Directory Certificate Services.

Configure IPsec authentication

A certificate is required on the Remote Access server and all DirectAccess clients so that they can use IPsec authentication. The certificate must be issued by an internal certification authority (CA). Remote Access servers and DirectAccess clients must trust the CA that issues the root and intermediate certificates.

To configure IPsec authentication

1. On the internal CA, decide if you will use the default computer certificate template, or if you will create a new certificate template as described in Creating Certificate Templates.

   Note
   If you create a new template, it must be configured for client authentication.
2. Deploy the certificate template if required. For more information, see Deploying Certificate Templates.
3. Configure the template for autoenrollment if required.
4. Configure certificate autoenrollment if required. For more information, see Configure Certificate Autoenrollment.

Configure certificate templates

When you use an internal CA to issue certificates, you must configure certificate templates for the IP-HTTPS certificate and the network location server website certificate.

To configure a certificate template

1. On the internal CA, create a certificate template as described in Creating Certificate Templates.
2. Deploy the certificate template as described in Deploying Certificate Templates.

After you prepare your templates, you can use them to configure the certificates. See the following procedures for details:

- Configure the IP-HTTPS certificate
- Configure the network location server

Configure the IP-HTTPS certificate

Remote Access requires an IP-HTTPS certificate to authenticate IP-HTTPS connections to the Remote Access server. There are three certificate options for the IP-HTTPS certificate:

- **Public**

  Supplied by a third party.
- **Private**

  The certificate is based on the certificate template that you created in **Configuring certificate templates**. It requires, a certificate revocation list (CRL) distribution point that is reachable from a publicly resolvable FQDN.
- **Self-signed**

  This certificate requires a CRL distribution point that is reachable from a publicly resolvable FQDN.

  Note
  Self-signed certificates cannot be used in multisite deployments.

Make sure that the website certificate used for IP-HTTPS authentication meets the following requirements:

- The certificate subject name should be the externally resolvable fully qualified domain name (FQDN) of the IP-HTTPS URL (the ConnectTo address) that is used only for the Remote Access server IP-HTTPS connections.
- The common name of the certificate should match the name of the IP-HTTPS site.
- In the subject field, specify the IPv4 address of the external-facing adapter of the Remote Access server or the FQDN of the IP-HTTPS URL.
- For the **Enhanced Key Usage** field, use the Server Authentication object identifier (OID).
- For the **CRL Distribution Points** field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.
- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcard characters in the name.

To install the IP-HTTPS certificate from an internal CA

1. On the Remote Access server: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, click **Request New Certificate**, and then click **Next** twice..
6. On the **Request Certificates** page, select the check box for the certificate template that you created in **Configuring certificate templates**, and if required, click **More information is required to enroll for this certificate**.
7. In the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common Name**.
8. In **Value**, specify the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL, and then click **Add**.
9. In the **Alternative name** area, in **Type**, select **DNS**.
10. In **Value**, specify the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL, and then click **Add**.
11. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
12. On the **Extensions** tab, next to **Extended Key Usage**, click the arrow, and make sure that Server Authentication is in the **Selected options** list.
13. Click **OK**, click **Enroll**, and then click **Finish**.
14. In the details pane of the Certificates snap-in, verify that the new certificate was enrolled with the intended purpose of server authentication.

Configure the DNS server

You must manually configure a DNS entry for the network location server website for the internal network in your deployment.

To add the network location server and web probe

1. On the internal network DNS server: On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the left pane of the **DNS Manager** console, expand the forward lookup zone for your domain. Right-click the domain, and click **New Host (A or AAAA)**.
3. In the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the network location server website (this is the name the

DirectAccess clients use to connect to the network location server). In the **IP address** box, enter the IPv4 address of the network location server, and click **Add Host**, and then click **OK**.

4. In the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the web probe (the name for the default web probe is directaccess-webprobehost). In the **IP address** box, enter the IPv4 address of the web probe, and then click **Add Host**.
5. Repeat this process for directaccess-corpconnectivityhost and any manually created connectivity verifiers. In the **DNS** dialog box, click **OK**.
6. Click **Done**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-DnsServerResourceRecordA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv4Address <network_location_server_IPv4_address>
Add-DnsServerResourceRecordAAAA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv6Address <network_location_server_IPv6_address>
```

You must also configure DNS entries for the following:

- **The IP-HTTPS server**

  DirectAccess clients must be able to resolve the DNS name of the Remote Access server from the Internet.

- **CRL revocation checking**

  DirectAccess uses certificate revocation checking for the IP-HTTPS connection between DirectAccess clients and the Remote Access server, and for the HTTPS-based connection between the DirectAccess client and the network location server. In both cases, DirectAccess clients must be able to resolve and access the CRL distribution point location.

- **ISATAP**

  Intrasite Automatic Tunnel Addressing Protocol (ISATAP) uses tunnels to enable DirectAccess clients to connect to the Remote Access server over the IPv4 Internet, encapsulating IPv6 packets within an IPv4 header. It is used by Remote Access to provide IPv6 connectivity to ISATAP hosts across an intranet. In a non-native IPv6 network environment, the Remote Access server configures itself automatically as an ISATAP router. Resolution support for the ISATAP name is required.

Configure Active Directory

The Remote Access server and all DirectAccess client computers must be joined to an Active Directory domain. DirectAccess client computers must be a member of one of the following domain types:

- Domains that belong in the same forest as the Remote Access server.

- Domains that belong to forests with a two-way trust with the Remote Access server forest.
- Domains that have a two-way domain trust to the Remote Access server domain.

To join the Remote Access server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. In the **System Properties** dialog box, click the **Computer Name** tab, and then click **Change**.
3. In the **Computer Name** box, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server, (for example, corp.contoso.com), and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with permissions to join computers to the domain, and then click **OK**.
5. When you see a dialog box welcoming you to the domain, click **OK**.
6. When you are prompted that you must restart the computer, click **OK**.
7. In the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

To join client computers to the domain

1. On the **Start** screen, type **explorer.exe**, and then press ENTER.
2. Right-click the Computer icon, and then click **Properties**.
3. On the **System** page, click **Advanced system settings**.
4. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
5. In the **Computer name** box, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server (for example, corp.contoso.com), and then click **OK**.
6. When you are prompted for a user name and password, enter the user name and password of a user with permissions to join computers to the domain, and then click **OK**.
7. When you see a dialog box welcoming you to the domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. In the **System Properties** dialog box, click Close.
10. Click **Restart Now** when prompted.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note

You must supply domain credentials after you enter the following command.

```
Add-Computer -DomainName <domain_name>

Restart-Computer
```

Configure GPOs

To deploy Remote Access, you require a minimum of two Group Policy Objects. One Group Policy Object contains settings for the Remote Access server, and one contains settings for DirectAccess client computers. When you configure Remote Access, the wizard automatically creates the required Group Policy Objects. However, if your organization enforces a naming convention, or you do not have the required permissions to create or edit Group Policy Objects, they must be created prior to configuring Remote Access.

To create Group Policy Objects, see Create and Edit a Group Policy Object.

An administrator can manually link the DirectAccess Group Policy Objects to an organizational unit (OU). Consider the following:

1. Link the created GPOs to the respective OUs before you configure DirectAccess.
2. When you configure DirectAccess, specify a security group for the client computers.
3. The GPOs are configured automatically, regardless of if the administrator has permissions to link the GPOs to the domain.
4. If the GPOs are already linked to an OU, the links will not be removed, but they are not linked to the domain.
5. For server GPOs, the OU must contain the server computer object—otherwise, the GPO will be linked to the root of the domain.
6. If the OU has not been linked previously by running the DirectAccess Setup Wizard, after the configuration is complete, the administrator can link the DirectAccess GPOs to the required OUs, and remove the link to the domain.

   For more information, see Link a Group Policy Object.

Note

If a Group Policy Object was created manually, it is possible that the Group Policy Object will not be available during the DirectAccess configuration. The Group Policy Object may not have been replicated to the domain controller closest to the management computer. The administrator can wait for replication to complete or force the replication.

Configure security groups

The DirectAccess settings that are contained in the client computer Group Policy Object are applied only to computers that are members of the security groups that you specify when configuring Remote Access.

To create a security group for DirectAccess clients

1. On the **Start** screen, type **dsa.msc**, and then press ENTER.
2. In the **Active Directory Users and Computers** console, in the left pane, expand the domain that will contain the security group, right-click **Users**, point to **New**, and then click **Group**.
3. In the **New Object - Group** dialog box, under **Group name**, enter the name for the security group.
4. Under **Group scope**, click **Global**, and under **Group type**, click **Security**, and then click **OK**.
5. Double-click the DirectAccess client computers security group, and in the **Properties** dialog box, click the **Members** tab.
6. On the **Members** tab, click **Add**.
7. In the **Select Users, Contacts, Computers, or Service Accounts** dialog box, select the client computers that you want to enable for DirectAccess, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADGroup -GroupScope global -Name <DirectAccess_clients_group_name>

Add-ADGroupMember -Identity DirectAccess_clients_group_name -Members <computer_name>
```

Configure the network location server

The network location server should be on a server with high availability, and it needs a valid Secure Sockets Layer (SSL) certificate that is trusted by the DirectAccess clients.

Note
If the network location server website is located on the Remote Access server, a website will be created automatically when you configure Remote Access and it is bound to the server certificate that you provide.

There are two certificate options for the network location server certificate:

- **Private**

  Note
  The certificate is based on the certificate template that you created in **Configuring certificate templates**.

- **Self-signed**

  Note
  Self-signed certificates cannot be used in multisite deployments.

Whether you use a private certificate or a self-signed certificate, they require the following:

- A website certificate that is used for the network location server. The certificate subject should be the URL of the network location server.

- A CRL distribution point that has high availability on the internal network.

To install the network location server certificate from an internal CA

1. On the server that will host the network location server website: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, click **Request New Certificate**, and then click **Next** twice.
6. On the **Request Certificates** page, select the check box for the certificate template that you created in **Configuring certificate templates**, and if required, click **More information is required to enroll for this certificate**.
7. In the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common Name**.
8. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
9. In the **Alternative name** area, in **Type**, select **DNS**.
10. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
11. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
12. Click **OK**, click **Enroll**, and then click **Finish**.
13. In the details pane of the Certificates snap-in, verify that new certificate was enrolled with the intended purpose of server authentication.

To configure the network location server

1. Set up a website on a high availability server. The website does not require any content, but when you test it, you might define a default page that provides a message when clients connect.

   This step is not required if the network location server website is hosted on the Remote Access server.
2. Bind an HTTPS server certificate to the website. The common name of the certificate should match the name of the network location server site. Ensure that DirectAccess clients trust the issuing CA.

   This step is not required if the network location server website is hosted on the Remote Access server.
3. Set up a CRL site that hass high availability on the internal network.

   CRL distribution points can be accessed through:

   - Web servers that use an HTTP-based URL, such as:

http://crl.corp.contoso.com/crld/corp-APP1-CA.crl

- File servers that are accessed through a universal naming convention (UNC) path, such as \\crl.corp.contoso.com\crld\corp-APP1-CA.crl

If the internal CRL distribution point is reachable only over IPv6, you must configure a Windows Firewall with Advanced Security connection security rule. This exempts IPsec protection from the IPv6 address space of your intranet to the IPv6 addresses of your CRL distribution points.

4. Ensure that DirectAccess clients on the internal network can resolve the name of the network location server, and that DirectAccess clients on the Internet cannot resolve the name.

See also

- Step 2: Configure the Remote Access Server
- Manage DirectAccess Clients Remotely

Step 2: Configure the Remote Access Server

Remote_access_combo

This topic describes how to configure the client and server settings that are required for remote management of DirectAccess clients. Before you begin the deployment steps, ensure that you have completed the planning steps that are described in Step 2: Plan the Remote Access Deployment.

Task	Description
Install the Remote Access role	Install the Remote Access role.
Configure the deployment type	Configure the deployment type as DirectAccess and VPN, DirectAccess only, or VPN only.
Configure DirectAccess clients	Configure the Remote Access server with the security groups that contain DirectAccess clients.
Configure the Remote Access server	Configure the Remote Access server settings.
Configure the infrastructure servers	Configure the infrastructure servers that are used in the organization.
Configure application servers	Configure the application servers to require authentication and encryption.
Configuration summary and alternate GPOs	View the Remote Access configuration

Task	Description
	summary, and modify the GPOs if desired.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Install the Remote Access role

You must install the Remote Access role on a server in your organization that will act as the Remote Access server.

To install the Remote Access role

1. On the server that will act as the Remote Access server, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. In the **Select server roles** dialog box, select **Remote Access**, click **Add Features**, and then click **Next**.
4. Click **Next** five times.
5. In the **Confirm installation selections** dialog box, click **Install**.
6. In the **Installation progress** dialog box, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature RemoteAccess -IncludeManagementTools
```

Configure the deployment type

There are three options that you can use to deploy Remote Access from the Remote Access Management console:

- DirectAccess and VPN
- DirectAccess only
- VPN only

Note
This guide uses the DirectAccess only method of deployment in the example procedures.

To configure the deployment type

1. On the Remote Access server, open the Remote Access Management console: On the **Start** screen, type, type **Remote Access Management Console**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management Console, in the middle pane, click **Run the Remote Access Setup Wizard**.
3. In the **Configure Remote Access** dialog box, select DirectAccess and VPN, DirectAccess only, or VPN only.

Configure DirectAccess clients

For a client computer to be provisioned to use DirectAccess, it must belong to the selected security group. After DirectAccess is configured, client computers in the security group are provisioned to receive the DirectAccess Group Policy Objects (GPOs) for remote management.

To configure DirectAccess clients

1. In the middle pane of the Remote Access Management console, in the **Step 1 Remote Clients** area, click **Configure**.
2. In the DirectAccess Client Setup Wizard, on the **Deployment Scenario** page, click **Deploy DirectAccess for remote management only**, and then click **Next**.
3. On the **Select Groups** page, click **Add**.
4. In the **Select Groups** dialog box, select the security groups that contain the DirectAccess client computers, and then click **Next**.
5. On the **Network Connectivity Assistant** page:
   - In the table, add the resources that will be used to determine connectivity to the internal network. A default web probe is created automatically if no other resources are configured. When configuring the web probe locations for determining connectivity to the enterprise network, ensure that you have at least one HTTP based probe configured. Configuring only a ping probe is not sufficient, and it could lead to an inaccurate determination of connectivity status. This is because ping is exempted from IPsec. As a result, ping does not ensure that the IPsec tunnels are properly established.
   - Add a Help Desk email address to allow users to send information if they experience connectivity issues.
   - Provide a friendly name for the DirectAccess connection.
   - Select the **Allow DirectAccess clients to use local name resolution** check box, if required.

     Note
     When local name resolution is enabled, users who are running the NCA can resolve names by using DNS servers that are configured on the DirectAccess client computer.

6. Click **Finish**.

Configure the Remote Access server

To deploy Remote Access, you need to configure the server that will act as the Remote Access server with the following:

1. Correct network adapters
2. A public URL for the Remote Access server to which client computers can connect (the ConnectTo address)
3. An IP-HTTPS certificate with a subject that matches the ConnectTo address
4. IPv6 settings
5. Client computer authentication

To configure the Remote Access server

1. In the middle pane of the Remote Access Management console, in the **Step 2 Remote Access Server** area, click **Configure**.
2. In the Remote Access Server Setup Wizard, on the **Network Topology** page, click the deployment topology that will be used in your organization. In **Type the public name or IPv4 address used by clients to connect to the Remote Access server**, enter the public name for the deployment (this name matches the subject name of the IP-HTTPS certificate, for example, edge1.contoso.com), and then click **Next**.
3. On the **Network Adapters** page, the wizard automatically detects:
   - Network adapters for the networks in your deployment. If the wizard does not detect the correct network adapters, manually select the correct adapters.
   - IP-HTTPS certificate. This is based on the public name for the deployment that you set during the previous step of the wizard. If the wizard does not detect the correct IP-HTTPS certificate, click **Browse** to manually select the correct certificate.
4. Click **Next**.
5. On the **Prefix Configuration** page (this page is only visible if IPv6 is detected in the internal network), the wizard automatically detects the IPv6 settings that are used on the internal network. If your deployment requires additional prefixes, configure the IPv6 prefixes for the internal network, an IPv6 prefix to assign to DirectAccess client computers, and an IPv6 prefix to assign to VPN client computers.
6. On the **Authentication** page:
   - For multisite and two-factor authentication deployments, you must use computer certificate authentication. Select the **Use computer certificates** check box to use computer certificate authentication and select the IPsec root certificate.
   - To enable client computers running Windows 7 to connect via DirectAccess, select the **Enable Windows 7 client computers to connect via DirectAccess** check box. You must also use computer certificate authentication in this type of deployment.
7. Click **Finish**.

Configure the infrastructure servers

To configure the infrastructure servers in a Remote Access deployment, you must configure the following:

- Network location server
- DNS settings, including the DNS suffix search list
- Any management servers that are not automatically detected by Remote Access

To configure the infrastructure servers

1. In the middle pane of the Remote Access Management console, in the **Step 3 Infrastructure Servers** area, click **Configure**.
2. In the Infrastructure Server Setup Wizard, on the **Network Location Server** page, click the option that corresponds to the location of the network location server in your deployment.
   - If the network location server is on a remote web server, enter the URL, and then click **Validate** before you continue.
   - If the network location server is on the Remote Access server, click **Browse** to locate the relevant certificate, and then click **Next**.
3. On the **DNS** page, in the table, enter additional name suffixes that will be applied as Name Resolution Policy Table (NRPT) exemptions. Select a local name resolution option, and then click **Next**.
4. On the **DNS Suffix Search List** page, the Remote Access server automatically detects domain suffixes in the deployment. Use the **Add** and **Remove** buttons to create the list of domain suffixes that you want to use. To add a new domain suffix, in **New Suffix**, enter the suffix, and then click **Add**. Click **Next**.
5. On the **Management** page, add management servers that are not detected automatically, and then click **Next**. Remote Access automatically adds domain controllers and System Center Configuration Manager servers.
6. Click **Finish**.

Configure application servers

In a full Remote Access deployment, configuring application servers is an optional task. In this scenario for remote management of DirectAccess clients, application servers are not utilized and this step is greyed out to indicate that it is not active. Click **Finish** to apply the configuration.

Configuration summary and alternate GPOs

When the Remote Access configuration is complete, the **Remote Access Review** is displayed. You can review all of the settings that you previously selected, including:

- **GPO Settings**

The DirectAccess server GPO name and Client GPO name are listed. You can click the **Change** link next to the **GPO Settings** heading to modify the GPO settings.

- **Remote Clients**

  The DirectAccess client configuration is displayed, including the security group, connectivity verifiers, and DirectAccess connection name.

- **Remote Access Server**

  The DirectAccess configuration is displayed, including the public name and address, network adapter configuration, and certificate information.

- **Infrastructure Servers**

  This list includes the network location server URL, DNS suffixes that are used by DirectAccess clients, and management server information.

See also

- Manage DirectAccess Clients Remotely
- Manage DirectAccess Clients Remotely

Step 3: Verify the Deployment

Remote_access_combo

This topic describes how to verify that you have correctly configured your deployment for remote management of DirectAccess clients .

To verify proper deployment

1. Connect a DirectAccess client computer to the corporate network and obtain the Group Policy Object.
2. On the client computer, click the **Network connections** icon in the notification area to access the DirectAccess Media Manager.
3. Click **DirectAccess Connection**, and you will see that the status is **Locally Connected**.
4. Remove the computer from the corporate network and connect it to a public network.
5. At a command prompt, type **nltest /dsgetdc: [fully qualified domain name]**. This command will verify that the corporate network is accessible to the client. If the domain controller is not accessible, the following error message will display reporting that the domain does not exist: ERROR_NO_SUCH_DOMAIN.

Deploy VPN When Connecting Remotely with Windows 8

It is becoming very common for computer users to connect to networks by using personal devices such as laptops, slates, and tablets. For organizations that support it, remote users can connect to the corporate network from their personal devices. The users can connect directly through the corporate WiFi by using their domain credentials, or by using virtual private network (VPN) connectivity. This topic describes the detailed steps and recommendations for deploying a VPN so that users running Windows 8 can connect remotely. It is meant to help organizations enable users to use their personal devices for common Windows 8 deployment scenarios.

Deploy Windows 8 VPN client

This section describes several methods for deploying a VPN , and it explains the IT pro and remote user experiences in detail. Some of the common profile management tasks such as editing, deleting, and profile renaming are also discussed briefly.

A VPN profile deployment includes setting up the profile on the client device and preparing the relevant settings that are required to access the corporate network.

Note

Windows RT does not support the installation of third-party applications. However, the Windows VPN client is available on Windows RT, and it is capable of connecting to third-party solutions. The third-party VPN solutions may need to be configured to allow the connection, and third-party VPN vendors own this support. The Windows VPN client can be provisioned by the methods that are described in this section (excluding the Connection Manager Administration Kit).

The following VPN connection profile deployment methods are available for unmanaged client devices (including devices running Windows RT):

- Use self-service portal
- Use Windows PowerShell scripts
- Use Connection Manager Administration Kit

Use self-service portal

On devices running Windows RT, an integrated management client is available that can communicate with a management infrastructure in the cloud to deliver line-of-business applications to the users. There are two parts to this management client:

- **Agent**: The built-in system component.
- **Self-service portal (SSP)**: The application that is used to browse for and install the line-of-business applications made available to them. IT pros can also use this SSP to deploy VPN connections on Windows client computers.

The following subsections explain the IT pro and remote user experiences for a profile that is set up by using the SSP.

IT pro experience

- To set up the profile, simply open the server management interface and provide settings for the VPN profile, then click **Finish**.
- The IT pro can also update the profile (if needed) through the server management interface, and no action is required from the users.

Remote user experience

Authorized users self-deploy the VPN setup on a device running Windows RT by navigating to the Control Panel application called **Company apps**, and then providing their company email address and password.

The agent then automatically configures the VPN and other features on the client device, based on the instructions provided by the organization's management infrastructure.

Use Windows PowerShell scripts

Simple Windows PowerShell scripts can be used by IT pros to help remote users install and configure VPN connections on unmanaged client computers. These scripts serve a dual purpose:

1. To set up VPN profiles on devices
2. To configure network-specific settings (such as prepared routes and network access point (NAP) settings) on the client device

The following subsections explain the IT pro and remote user experiences for installing the VPN profile by using these scripts.

IT pro experience

1. The IT pro creates an external web portal that is accessible to remote users when they connect with domain credentials. This web portal is published to all the corporate network users who intend to connect to the corporate network by using the VPN.
2. The IT pro creates a Windows PowerShell script and adds it to the web portal with setup instructions and guidelines.

The following Windows PowerShell script to create the VPN profiles is very simple. Other Windows PowerShell sample scripts and documentation are available on Microsoft TechNet to help IT pros create corresponding scripts for their network deployment.

Syntax

```
## VPN profile creation cmdlet syntax

Add-VpnConnection -Name "<name of the VPN connection>" -ServerAddress "<FQDN or IP
Address of VPN server>" -TunnelType <tunnel type enum value> -EncryptionLevel <encryption
level enum value> -AuthenticationMethod <authentication method enum value> -
SplitTunneling -EapConfigXmlStream <EAP configuration XML>
```

Syntax for password-based deployment

```
## Generate the required EAP method configuration (if applicable)


# Generate configuration XML for (default) EAP-MSCHAPv2 authentication method
$a = New-EapConfiguration
# Generate configuration XML for PEAP authentication method with EAP-MSCHAPv2 as its inner method
$b = New-EapConfiguration -Peap -FastReconnect $true -TunnledEapAuthMethod $a.EapConfigXmlStream


## VPN cmdlet usage sample with VPN Connection name 'MyVPN' and VPN Server – contoso.com
Add-VpnConnection -Name "My VPN" -ServerAddress "contoso.com" -TunnelType Automatic -EncryptionLevel Required -AuthenticationMethod Eap -SplitTunneling -EapConfigXmlStream $b.EapConfigXmlStream
```

Remote user experience

1. Remote users receive the external web portal information from the IT pro. To configure the VPN on the unmanaged device, users navigate to the portal and sign in with their domain credentials.
2. On the VPN configuration page, the user simply clicks (or touches) the link, and the setup script is downloaded to the client device.
3. The user clicks **Run**, and the profile is installed.
4. The IT pro informs the users to run the script again each time the script is updated.

Advanced scenarios

IT pros can include many advanced configuration settings in this single-click VPN deployment script to provide a complete client deployment experience for remote users. These advanced settings might include the following for the VPN connection:

- Policies for IPsec, publishing routes, or NAP settings
- Web-proxy configuration
- Intranet settings
- Connection-specific DNS suffix configuration

Connection-specific DNS suffix

For a given interface, the connection specific DNS suffix can be configured by using the following Windows PowerShell script:

```
## Configure connection specific DNS suffix
Set-DnsClient -InterfaceAlias "My VPN" -ConnectionSpecificSuffix "contoso.com"
```

However, for VPN connections, there might be situations in which Kerberos authentication occurs before the VPN connection configuration is complete. When this happens, the Kerberos protocol

will not allow authentication for 15 minutes. To work around this potential situation, IT pros can use the following command to clear the Kerberos negative cache and avoid delay:

```
## Clear negative caching for Kerberos

klist purge_bind
```

Get Connected Wizard

The Get Connected Wizard is the VPN connection setup wizard that is available on Windows Server and Windows client operating systems. It is a simple method for creating VPN connections if the VPN deployment is password-based, and it has no complexities like deployment NAP settings or publishing routes. The remote access server name is the only piece of information that is needed to create the connection.

If the profile is more complex, the IT pro should provide detailed instructions and configuration scripts.

For more information, see Set Up and Edit VPN Connections in Windows 8.

Use Connection Manager Administration Kit

The Connection Manager Administration Kit (CMAK) Wizard creates customized VPN profiles according to organizational requirements, and distributes the generated package to clients for a single-click installation. This tool is commonly used for VPN client package generation.

Note
CMAK can be used to deploy VPN profiles on all devices running Windows, except those running Windows RT.

For more information, see Connection Manager Administration Kit Overview.

Manage profiles

For the VPN profiles that are created by using Windows PowerShell scripts, the IT pro can share another single-click script through the web portal to remove the profile and its corresponding settings from a client device.

Windows PowerShell script example to remove a VPN profile

```
## Remove a VPN profile named “My VPN” created for local user only

Remove-VpnConnection -Name "My VPN" -Force
```

For more information, see Set Up and Edit VPN Connections in Windows 8.

Advanced profile editing configurations

1. If IT pros want to make any changes to a VPN profile that was deployed on client systems by using System Center Configuration Manager, they simply open the server user interface and make the required changes. A new VPN profile will be generated and used for all future deployments
2. For script-based profiles, IT pros can write a new script and distribute it through a shared resource or a web portal. As an example, the following scripts show some methods for changing the tunnel type, encryption level, and authentication method for the VPN profile “My VPN”, which was created in the Use Windows PowerShell scripts section earlier.

   The following script shows how an IT pro can delete an existing VPN connection, and then deploy a new one:

```
## 1. Delete older VPN connection, and provision a new one (with
the same name) with updated ##    settings
$vpnConnection = Get-VpnConnection -AllUserConnection


if($vpnConnection.Name -eq "My VPN")
{
    # Remove the connection
    Remove-VpnConnection –Name "My VPN" –AllUserConnection -
Force
}
## Provision a new VPN connection “My VPN” with updated settings
Add-VpnConnection -Name "My VPN" -ServerAddress "contoso.com" -
TunnelType Pptp -EncryptionLevel Optional -AuthenticationMethod
MsChapv2 -SplitTunneling
```

   The following script provides an example for how you can edit an existing connection, and add a new connection if it is not already provisioned:

```
## 2. (a) If already provisioned, update this older VPN
connection;
##    (b) If not provisioned, provision a new VPN connection
with updated settings


$vpnConnection = Get-VpnConnection -AllUserConnection


if($vpnConnection.Name -eq "My VPN")
{
    # Update the connection
```

```
    Set-VpnConnection –Name "My VPN" –AllUserConnection -
TunnelType Pptp -EncryptionLevel Optional -AuthenticationMethod
MSChapv2
}
else
{
    # Provision a new VPN connection “My VPN” with updated
settings
    Add-VpnConnection -Name "My VPN" -ServerAddress
"contoso.com" -TunnelType Pptp -EncryptionLevel Optional -
AuthenticationMethod MsChapv2 -SplitTunneling
}
```

The following script shows how you can distribute separate Install and Uninstall scripts. In this case, the IT pro simply updates the installation script and asks users run this script to uninstall the older VPN connection.

```
## 3. If already provisioned, remove the older VPN connection
$vpnConnection = Get-VpnConnection -AllUserConnection

if($vpnConnection.Name -eq "My VPN")
{
    # Remove the connection
    Remove-VpnConnection –Name "My VPN" –AllUserConnection -
Force
```

For more information, see Set Up and Edit VPN Connections in Windows 8.

Connect VPN by using Windows 8 user interface

The Windows 8 user interface includes the following features:

- Multiple connection profiles are supported
- Credential collection for all authentication methods in Windows 8
- Connection progress status and failure messages
- Seamless connection by using authentication protocol negotiation during connection for default profiles
- Users are not prompted for stored credentials, and stored credentials for a profile can be cleared from the right-click menu

- Silent disconnect when the server initiates a disconnection

Launch Windows 8 user interface for VPN connections

Touch-enabled devices: Slide the user interface from the right corner of the screen.

All other devices: Use the following procedure to launch the Windows 8 user interface:

1. Open Charms (**Windows Key + C**).
2. Click **Settings**.
3. Click the network icon in the right corner of the screen.
4. The Windows 8 user interface appears.

Desktop: Click the **Network** icon in the **Notification** area to launch the Windows 8 user interface.

The following screenshot shows the Windows 8 user interface. Multiple VPN and wireless connections are shown, and one of the VPN connections is selected.

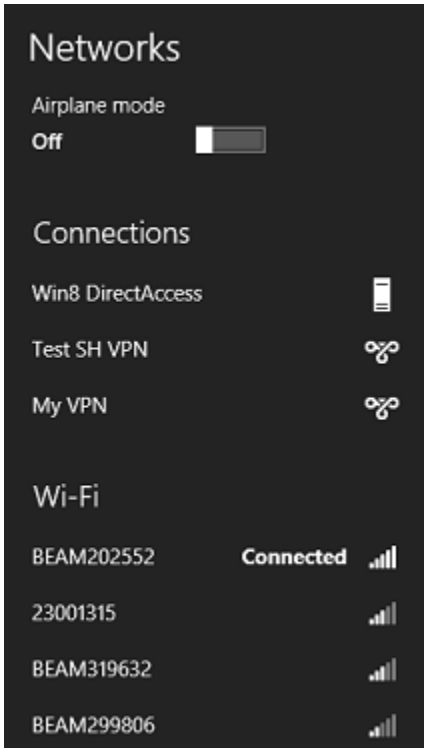

Connect to VPN

Use the following procedure to connect to a password-based VPN:

To connect to a VPN

1. On the **Networks** menu, click the name of a VPN connection, and then click **Connect**.

2.
3. In the **Network Authentication** text boxes, the remote user types a user name and password.
4. If the IT pro enabled the **Verify server identity** option, the user is prompted to confirm, and then clicks **Connect**.

The following images represent the password-based VPN connection process.

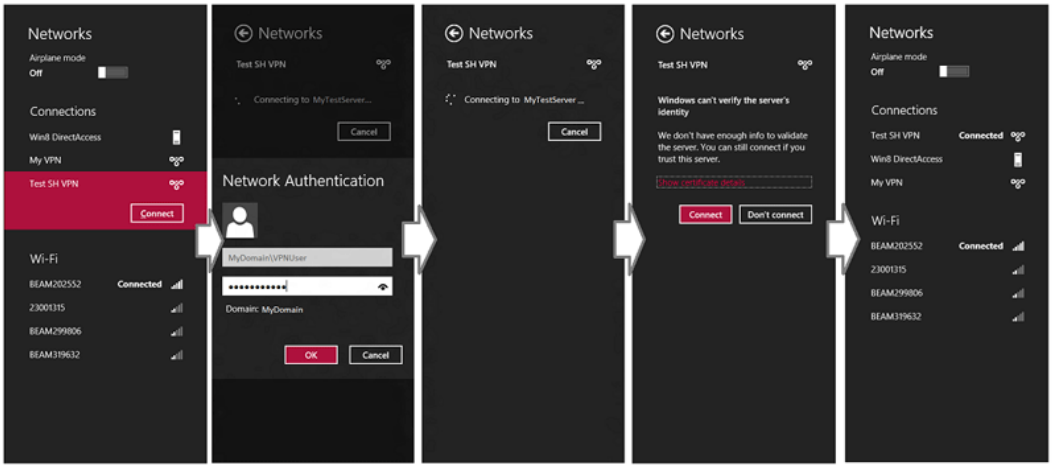

Deploy the VPN

This section describes the options that are available for the VPN deployment.

Password-based deployment

Windows 8 can establish a remote access connection through a variety of password-based point-to-point authentication protocols, which include:

1. Non-EAP methods (PAP, CHAP, or MS-CHAPv2)
2. EAP methods (PEAP or EAP-TTLS)

Non-EAP password-based methods are not considered secure, and they should be deployed with secure TLS methods such as PEAP or EAP-TTLS.

We recommend using PEAP- or EAP-TTLS-based authentication methods for corporate deployments. Both options are available in Windows 8.

Certificate-based deployments

When an organization needs to support enrolling and importing user certificates, there are many ways to deploy the public key infrastructure (PKI). The best deployment method depends on the organization's policy and deployment infrastructure.

For more information, see Certificate Enrollment Web Services in Windows Server 2008 R2.

Enable split tunneling

Windows PowerShell cmdletsand System Center Configuration Manager support enabling and disabling split-tunneling for VPN connections. But disabling split-tunneling (or enabling force tunneling) requires some additional configuration, such as deploying a web proxy for Internet access over the VPN connection.

IT pros can enable split tunneling by using the Windows PowerShell switch parameter **–SplitTunneling** in the **Add-VpnConnection** cmdlet, or by setting the **–SplitTunneling** Boolean parameter to the value **$true** in the **Set-VpnConnection** cmdlet. This is shown in the following commands:

```
## 1. Enabling* Split-Tunneling in a VPN connection while creating VPN connection

Add-VpnConnection -Name "My VPN" -ServerAddress "contoso.com" -SplitTunneling


## 2. Enabling Split-Tunneling in an existing VPN connection

Set-VpnConnection -Name "My VPN" -SplitTunneling $true


## 2. Disabling Split-Tunneling in an existing VPN connection

Set-VpnConnection -Name "My VPN" -SplitTunneling $false
```

Note

By default, **Add-VpnConnection** keeps split tunneling disabled for a VPN connection.

For more information, see Set Up and Edit VPN Connections in Windows 8.

Understand route implications

For split-tunneled VPN connections, the corporate network routes have to be configured on the client device after the VPN connection is established. The following two methods can accomplish this:

- **DHCP-based route publishing**: When a client device is connected to the corporate network by using a VPN connection, it gets its IP address from the DHCP server that is designated for that purpose. This DHCP server can also supply the required routes for intranet access with the configuration packet to the client.
- **One-click script route publishing**: The IT pro can provide a one-click Windows PowerShell script to remote users for scheduled tasks. IT pros use a web portal or external shared resource to provide script for configuration of the corporate network routes on unmanaged client devices. The script adds a scheduled task on the device, which is triggered every time the device establishes a connection with the remote VPN server. The script includes the required routes, which enable the client to access corporate resources. If the IT pro already uses a single-click script to deploy other settings (such as IPsec policies and the DNS suffix), this route information can be included with that script.

Deploy web proxy and intranet settings

In the case of force-tunneling, the web proxy needs to be configured on the unmanaged devices for them to access Internet sites. Also, the intranet settings have to be enabled for a VPN connection to suppress repeated credential prompts while accessing intranet sites. The following sections discuss these configurations for unmanaged devices.

Web proxy

Force-tunneling requires the VPN client to use corporate network web proxy for all the Internet traffic. The IT pro can configure a web-proxy that is specific for the VPN connection by using Windows PowerShell cmdlets for web proxy deployment. These cmdlets can be included in the one-click script used to deploy the VPN connection on the client computer.

Following is the syntax and a sample usage command for a VPN-specific web proxy:

```
## Syntax of PowerShell cmdlet for VPN specific web-proxy configuration

Set-VpnConnectionProxy –Name “<name of the VPN connection>” [-AutoDetect] [-
AutoConfigurationScript “<URL of the proxy configuration WPAD script>”] –[ProxyServer
“<URL or IP Address/port number of the proxy server>”] [-BypassProxyForLocal] [-
ExceptionPrefix “<comma separated list of the proxy exception Protocols and/or URLs
and/or IP Addresses>”]
```

Intranet settings

Enabling intranet settings is required for the user to access corporate sites without repeated credential prompts while connected through the VPN. Intranet settings determine whether websites are a part of the local intranet zone and if they require domain credentials. If intranet settings are enabled for a VPN connection, the VPN credentials are automatically used to authenticate users when they access local intranet sites.

For domain-joined computers, the VPN interface is automatically recognized, and these settings are turned on every time the client is connected to the corporate network through the VPN. But for unmanaged devices, VPN connections are not automatically recognized, which results in a credentials prompt every time the user navigates to a local intranet site. The IT pro should configure the following registry settings on clients running Windows 8 to enable the intranet settings when the client is connected to the corporate network through a VPN connection:

```
# registry settings to enable the Intranet settings while connected to a VPN Server

reg add "HKLM\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap" /v
AutoDetect /t REG_DWORD /d 0 /f

reg add "HKLM\Software\Microsoft\Windows\CurrentVersion\Internet Settings\ZoneMap" /v
UNCAsIntranet /t REG_DWORD /d 1 /f
```

Deploy multisite VPN

Organizations with more than one business location can choose to have multisite VPN server deployments. In this case, Global Server Load Balancing (GSLB) is used to resolve the address and to redirect the client to the nearest VPN server site. If this is not an option, the IT pro can provide VPN connection profiles per site, which provides a choice for remote users to connect to the desired server.

Interoperate with third-party VPN servers

The native VPN client running Windows 8 can interoperate with commonly deployed VPN server devices by using the PPTP, L2TP/IPsec, SSTP, and IKEv2 protocols.

For more information about tunneling, authentication, and encryption protocols that are supported, see the **Third-party VPN interoperability** section in VPN Compatibility and Interoperability in Windows 8 and Windows Server 2012.

Tip
We recommend using L2TP/IPsec and IKEv2 protocols for connection.

Use Connected Standby state

The Connected Standby state is a feature in Windows 8, which enables the computer to enter a standby state and consume very low battery power, but remain connected and fresh.

When a device is connected to the remote VPN server, and it enters the Connected Standby state, the connection remains live, even if the interface is closed due to a low power state. When the device resumes from the Connected Standby state, VPN clients search for the same underlying interface. If the interface is available after resuming, and the VPN credentials are cached, the client seamlessly connects with the remote VPN server. However, if the credentials were not cached, and remote user input (such as a physical smart card or one-time password) is required, the VPN connection is not resumed automatically.

Appendix A: Glossary

The following table explains terms that are used in this topic.

Name	Details
Windows RT management by Configuration Manager	New server role that deploys and manages certain settings on Windows RT devices. It does support managing the complete VPN profile.
Network access point (NAP) infrastructure settings	Settings to enable the NAP agent on the client side. Settings require information such as NAP

	server and SHAs for deploying the NAP service.
VPN profile	Contains all the VPN settings, such as VPN server information, VPN authentication settings (such as 2-factor, EAP settings, and password), tunnel settings (such as IKEv2, SSTP, L2TP, and PPTP), health check settings, and so on.
OTP	RSA is the most popular OTP deployment. It provides a client EAP method in Windows to integrate with RSA OTP server.
Split-tunneling	Accessing the Internet without going through the corporate network firewall.
NAP	NAP configuration for unmanaged devices is not considered a primary requirement, but if required in a corporate network, it can be deployed by a single-click script provided by the IT pro (similar to the methods discussed above for single-click scripts).
FQDN	Fully qualified domain name
GSLB	Global Server Load Balancing
WinRT	Modern Windows 8 environment
WindowsRT	New Windows 8 version for ARM architecture

Experience terminology

Name	Details
VPN: **One click script**	VPN profile is deployed using the PowerShell script **Enabled only in W8 clients:** New VPN client PowerShell
VPN: **System Center Configuration Manager**	VPN profile is deployed using the System Center Configuration Manager Windows RT management feature **Enabled only in Windows RT clients:** New instance based VPN client PowerShell

	(consumes profile setting from XML)
VPN: **Get Connected Wizard**	Remote user creates the profile using built-in client UI **Available in W8 & down-level Windows client:** However, Windows 8 has a connect experience and method discovery option where the remote user is not asked any information except the server name for majority profiles.
NAP: **Single click script**	IT pro provides the script to plumb the NAP policies on unmanaged client devices – details available on TechNet **Available in W8 & down-level Windows client:** No changes made in Windows 8
Routes deployment: **Single click script**	The IT pro provides the script to plumb the intranet routes on unmanaged devices. **Available in W8 & down-level Windows client:** No changes made in Windows 8
Certificate enrollment web-servicve	Organization needs to support enrolling/importing user certificates. Web-enrollment is an intuitive way to do so. More details are available on TechNet.

Manage VPN Connections in Windows 8 by Using Windows PowerShell

To make virtual private network (VPN) profile management easier for administrators, Windows 8 includes new Windows PowerShell cmdlets 8. These cmdlets help administrators create, configure, and remove VPN connections on client computers by using Windows PowerShell scripts.

Windows PowerShell cmdlets enable you to achieve complex tasks with a few lines of script. Third-party vendors can build custom solutions based on the overall Windows PowerShell support that is provided by Windows, specifically in Windows Networking.

You can write simple Windows PowerShell scripts by using these cmdlets with the required configuration to deploy a VPN, and distribute these scripts to client computers through a common

shared resource or a web portal. Knowledge workers then run these single-click scripts on their computers, and the VPN profiles are configured automatically.

Use VPN deployment and management scripts: For IT pros

You can use Windows PowerShell cmdlets to write scripts to deploy VPN connections on computers running Windows 8. The following sections include:

1. Create VPN connections
2. Configure and edit VPN connections
3. Write configuration scripts for VPN connection management
4. Use Windows PowerShell cmdlets for third-party EAP methods

Create VPN connections

The Windows PowerShell cmdlet for creating a VPN connection on the client system is **Add-VpnConnection**.

Syntax

```
Add-VpnConnection [-Name] <string> [-ServerAddress] <string> [-TunnelType <string> {Pptp | L2tp | Sstp | Ikev2 | Automatic}] [-EncryptionLevel <string> {NoEncryption | Optional | Required | Maximum}] [-AuthenticationMethod <string[]> {Pap | Chap | MSChapv2 | Eap}] [-SplitTunneling] [-AllUserConnection] [-L2tpPsk <string>] [-RememberCredential] [-UseWinlogonCredential] [-EapConfigXmlStream <xml>] [-Force] [-PassThru] [-WhatIf] [-Confirm]
```

Properties

1. **Name** is the connection name, which is a mandatory parameter, and it cannot be changed.
2. **ServerAddress** is the IP address (FQDN, or URL) of the remote server to which the connection is to be established. This is a mandatory parameter.
3. The remainder of the properties are optional. Their names and value sets indicate their purpose. For example, **TunnelType** specifies the VPN tunnel selection (such as PPTP or L2TP), and **AuthenticationMethod** specifies the VPN authentication protocol selection.
4. Switch parameters (**SplitTunneling**, **AllUserConnection**, **RememberCredentials**, UseWinlogonCredentials), when specified, enable the property.
5. **EapConfigXmlStream** parameter accepts Extensible Authentication Protocol (EAP) configuration XML. This can be any EAP XML (Microsoft or non-Microsoft). For EAP methods in Windows 8, this XML can be generated by using the New-EapConfiguration, which is described later in this document.

Configure and edit VPN connections

The Windows PowerShell cmdlet for editing a VPN connection on the client system is Set-VpnConnection.

Syntax

```
Set-VpnConnection [-Name] <string> [[-ServerAddress] <string>] [-TunnelType <string>
{Pptp | L2tp | Sstp | Ikev2 | Automatic}] [-EncryptionLevel <string> {NoEncryption |
Optional | Required | Maximum}] [-AuthenticationMethod <string[]> {Pap | Chap | MSChapv2
| Eap}] [-SplitTunneling <bool>] [-AllUserConnection] [-L2tpPsk <string>] [-
RememberCredential <bool>] [-UseWinlogonCredential <bool>] [-EapConfigXmlStream <xml>] [-
PassThru] [-Force] [-WhatIf] [-Confirm]
```

Properties

1. **Name** is the connection name. It is a mandatory parameter, and it cannot be modified.
2. **AllUserConnection** is used to identify the scope of the connection. If this parameter is supplied, the global address list is searched for the connection; otherwise, the local connections are searched.
3. All the other parameters carry the same meaning as explained previously in the information about the **Add-VpnConnection** cmdlet .

New-EapConfiguration

New-EapConfiguration is a VPN connection cmdlet that is used to generate XML for EAP configuration. It can be used with the **Add-VpnConnection** or **Set-VpnConnection** cmdlets.

This cmdlet can generate XML for EAP, and it includes the following parameters: **EAP-MSCHAPv2** (default configuration), **EAP-TLS** (for smart card or other certificates), and **EAP-TTLS** (for internal authentication). It accepts configuration XML from Microsoft or from third-parties for EAP, non-EAP, or PEAP.

Syntax

```
New-EapConfiguration [[-UseWinlogonCredential]] [-WhatIf] [-Confirm]

New-EapConfiguration [-Tls] [-VerifyServerIdentity] [-UserCertificate] [-WhatIf] [-
Confirm]

New-EapConfiguration [-Ttls] [-UseWinlogonCredential] [-TunnledNonEapAuthMethod <string>
{Pap | Chap | MSChap | MSChapv2}] [-TunnledEapAuthMethod <xml>] [-WhatIf] [-Confirm]

New-EapConfiguration [-Peap] [-VerifyServerIdentity] [-TunnledEapAuthMethod <xml>] [-
EnableNap] [-FastReconnect <bool>] [-WhatIf] [-Confirm]
```

Properties

1. If no method switch (**Tls**, **Ttls**, or **Peap**) is supplied, the default **Eap-MsChapv2** method is used.

2. Smart card is the default choice for **EAP-TLS**. To use user certificates, specify **UserCertificate**.
3. In addition to the methods that are provided by Microsoft, administrators can supply configuration XML for any EAP method for the **Ttls** and **Peap** parameter sets in the **TunneledEapAuthMethod** property.

VPN connection lookup

The Windows PowerShell cmdlet for VPN connection lookup on the client system is **Get-VpnConnection**.

Syntax

```
Get-VpnConnection [[-Name] <string[]>] [-AllUserConnection]
```

Properties

1. **Name** is the connection name. It is an optional parameter.
   a. If the connection name is not specified, all VPN connections in the given address list are listed for the user.
   b. More than one connection name can also be looked up in the global address list by using a comma-separated list.
2. **AllUserConnection** parameter is used to identify the scope of the connection. If this switch is supplied, the global address list is searched for the connections; otherwise, the local connections are used.

VPN connection removal

The Windows PowerShell cmdlet for VPN connection deletion on the client system is **Remove-VpnConnection**.

Syntax

1. **Name** is the connection name. It is a mandatory parameter. More than one connection name can be removed at one time by using a comma-separated list.
2. **AllUserConnection** parameter is used to identify the scope of the connection. If this switch is supplied, the connections are removed from the global address list; otherwise, the local connections are removed.
3. **Force** switch is specified to suppress the connection removal confirmation dialogue.

Write configuration scripts for VPN connection management

You can use the VPN management cmdlets described previously to deploy VPN connections on computers running Windows 8, according to the network deployment and corporate requirements.When distributed, this script provides a single-click installation for knowledge workers.

The following examples show how to write a single-click script that can distribute the VPN profile and run specific post-connection actions every time the VPN interface initiates.

```
# VPN Connection 'My VPN' provisioning


#-------------------------------------------------------------
# VPN Connection look-up to check any previous installations
#-------------------------------------------------------------
$isTestVpn = $false
$vpnConnections = Get-VpnConnection -AllUserConnection


if($vpnConnections.Name -eq "My VPN")
{
    Write-Host "'My VPN' connection is already configured on your system." -
ForegroundColor Yellow -BackgroundColor DarkGreen
    Write-Host "If you wish to reinstall, please uninstall the connection and then attemp
again." -ForegroundColor Yellow -BackgroundColor DarkGreen
    Write-Host ""
    Write-Host "Please press any key to exit..."
    $x = read-host
    exit
}


Write-Host "****************************************" -ForegroundColor Black -
BackgroundColor White
Write-Host "      Installing 'My VPN' connection      " -ForegroundColor Black -
BackgroundColor White
Write-Host "________________________________________" -ForegroundColor Black -
BackgroundColor White


try
{
    # Generate Eap-Mschapv2 configuration XML
    $a = New-EapConfiguration
    # Generate Eap-TTLS configuration XML with EAP-Mschapv2 as inner method
    $b = New-EapConfiguration -Ttls -TunnledEapAuthMethod $a.EapConfigXmlStream
```

```
    # Edit properties within the generated configuration XML
    $c = $b.EapConfigXmlStream
    $c.EapHostConfig.Config.EapTtls.Phase1Identity.IdentityPrivacy = "false"

    # Create the VPN connection 'My VPN' with the EAP configuration XML generated above
    Add-VpnConnection -Name "My VPN" -ServerAddress "contoso.com" -EncryptionLevel
Required -AuthenticationMethod Eap -EapConfigXmlStream $c -SplitTunneling -
AllUserConnection
}
catch
{
    Write-Host "Error in connection setup!" -ForegroundColor White -BackgroundColor Red
    Write-Host $_.Exception.Message
    throw
}

Write-Host ""
Write-Host "'My VPN' VPN connection is ready for use." -ForegroundColor Black -
BackgroundColor White
```

The following sample script adds a new task to the system Task Scheduler. It performs the tasks when the VPN connects and disconnects.

```
## This piece of code performs the following actions -
## 1. Dynamically generate support files for VPN connection configuration,
## 2. Creates task in task-scheduler, which triggers post-connect actions on the VPN
connection using
##    the support files generated above

write-host "Generating support files..."
# Create the common directory for support files
$testProfileDir = Test-Path $env:ALLUSERSPROFILE\Microsoft\Network\Connections\Cm\MyVPN\
if(!$testProfileDir)
{
    mkdir $env:ALLUSERSPROFILE\Microsoft\Network\Connections\Cm\MyVPN
}
```

```
$Directory = "$env:ALLUSERSPROFILE\Microsoft\Network\Connections\Cm\MyPN"

# Specify support file content variables, and generate these files in the common location
($Directory)
$RoutesFile = $Directory + "\routes.netsh"
$DnsSuffixFile = $Directory + "\dnssuffix.ps1"
$ConnectTaskFile = $Directory + "\Connect.xml"

# route provisioning file content
$routes = @"
interface ipv4
add route prefix=192.168.0.0/16 interface="My VPN" store=active
exit
"@

# Connection specific DNS suffix provisioning file content
$dnssuffix = @"
Set-DnsClient -InterfaceAlias "My VPN" -ConnectionSpecificSuffix "contoso.com"
"@

# Post-connect task scheduler XML file content
$connecttask = @"
<?xml version="1.0" encoding="UTF-16"?>
<Task version="1.2" xmlns="http://schemas.microsoft.com/windows/2004/02/mit/task">

# XML file content
</Settings>
  <Actions Context="Author">
    <Exec>
      <Command>netsh.exe</Command>
      <Arguments>exec routes.netsh</Arguments>
```

```
<WorkingDirectory>%ALLUSERSPROFILE%\Microsoft\Network\Connections\Cm\MyVPN\</WorkingDirec
tory>
    </Exec>
    <Exec>
      <Command>powershell.exe</Command>
      <Arguments>-Command
"%ALLUSERSPROFILE%\Microsoft\Network\Connections\Cm\MyVPN\dnssuffix.ps1"
      </Arguments>
    </Exec>
  </Actions>
</Task>
"@

$routes|out-file -FilePath $RoutesFile
$dnssuffix|out-file -FilePath $DnsSuffixFile
$connecttask|out-file -FilePath $ConnectTaskFile

write-host "Files generated..."

write-host "Initiating configuration for 'My VPN' ..." -ForegroundColor Black -
BackgroundColor White
SCHTASKS /create /tn "\MyVPN\Connect" /xml
$env:ALLUSERSPROFILE\Microsoft\Network\Connections\Cm\MyVPN\Connect.xml /F

write-host ""
write-host "Configuration completed..." -ForegroundColor Black -BackgroundColor White
```

Use VPN deployment and management scripts: For knowledge workers

After you write the scripts for connection deployment and distribute these scripts to a public shared resource or to an external web portal, knowledge workers can run them by clicking the link you provide. When a user runs the scripts, the VPN profile is configured on the client computer, and it is ready to use on the corporate network.

The Windows PowerShell cmdlets are available in all Windows 8 operating systems, and all types of clients (managed, unmanaged, computers, and tablets) can use these scripts for client deployment.

Use Windows PowerShell cmdlets for third-party EAP methods

Extensible Authentication Protocol (EAP) methods in Windows 8 are supported with the **Add-VpnConnection** , **Set-VpnConnection** , and **Get-VpnConnection** Windows PowerShell cmdlets, but configuration and retrieval of third-party EAP authentication methods for the given VPN connection can be a bit difficult.

If a system is already configured with a VPN connection that uses a third-party EAP method, the **Get-VpnConnection** cmdlet for that VPN connection might fail if the configuration details are not stored in the default global address list. The following procedure describes how to retrieve the EAP configuration settings for third-party EAP methods and use those settings to set up similar new VPN connections.

To use third-party EAP settings

1. On a computer where the third-party EAP method is configured for VPN, click **Start**, type Control Panel, click **Network and Internet**, and then click **Network Connections**.
2. Right-click the VPN connection that is configured with third-party EAP, and select **Properties**.
3. In the VPN **Connection Properties**, click the **Security** tab.
4. In **Authentication - Use Extensible Authentication Protocol (EAP)**, verify that the third-party EAP method is selected.
5. Click **Properties** under the EAP method drop-down list.
6. Click **OK** twice to exit.
7. The third-party EAP configuration details will now be stored in the default global address list, and the subsequent calls to **Get-VpnConnection** will succeed.

   To configure new VPN connections with the identical third-party EAP settings, retrieve the EAP configuration XML by using **Get-VpnConnection** as described previously, and then use **Add/Set-VpnConnection** with the EAP XML.

Set Up and Edit VPN Connections in Windows 8

The virtual private network (VPN) in Windows 8 has been optimized to use with touch devices. Creating a VPN connection and connecting to corporate networks is faster and easier. The VPN

connection icon appears in **View Available Networks** with other network connections. Creating a VPN profile requires only the server information. The Windows operating system automatically discovers authentication methods and tunnels when it connects, which provides a seamless experience for creating a VPN.

The following sections describe how to use Windows 8 to set up and edit VPN connections:

- Create a new VPN connection
- Connect to a corporate network by use an existing VPN connection
- Edit connection properties
- Delete a connection

Create a new VPN connection

In Windows 8, you can use the Get Connected Wizard to create a VPN profile. To launch the wizard, type **vpn** in **Settings**, and then click **Set up a virtual private network (VPN) connection**.

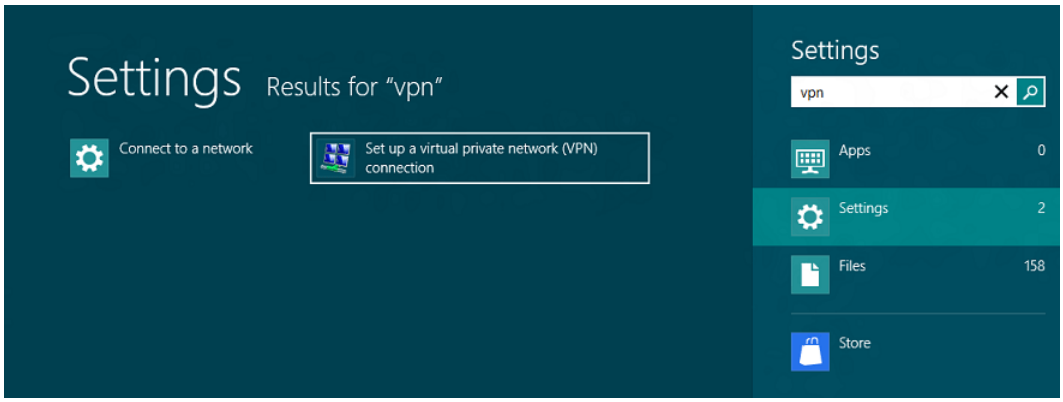

Figure 1: Launch the Get Connected Wizard

In the wizard, you will provide the Internet address of the Remote Access VPN server to which you will connect, and a name for the new connection. You can use an IP address or the fully qualified domain name of the Remote Access server.

To find your server settings, on a computer running Windows 8 with a VPN already configured, view the connection properties as shown in Figure 8. The destination name is the text that will be used for the VPN connection name (profile name), as shown in Figure 2.

You can click **Remember Credentials** to save your credentials on the first successful connection attempt. This creates a new Discovery Profile, which the client can use to connect to the VPN.

Discovery Profile

Figure 2 shows the only information that is needed to create a VPN profile. During the first successful connection attempt, authentication and tunneling protocols are negotiated and configured for user name and password-based VPN deployments. After these protocols are negotiated, you can edit the settings in **Connection Properties**.

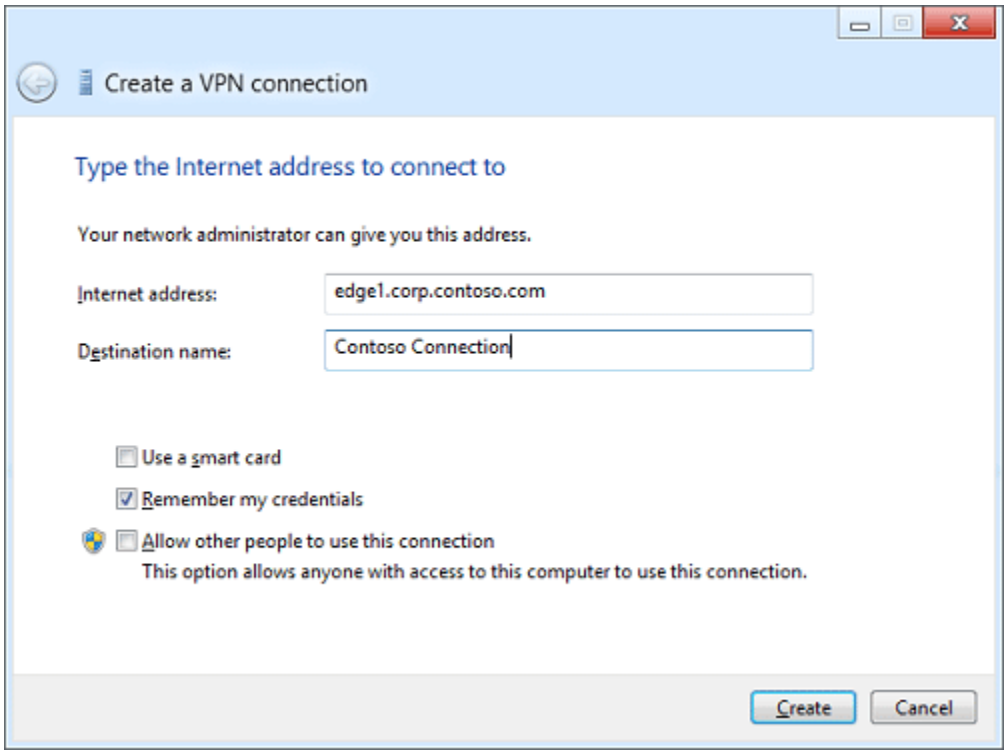

Figure 2: Get Connected Wizard Screen 1

Connect to a corporate network by use an existing VPN connection

To connect with an existing VPN connection

1. Open **Charms** by moving the cursor to the top or bottom right corner of your computer screen. It can also be accessed by using the **Win+C** keyboard shortcut.

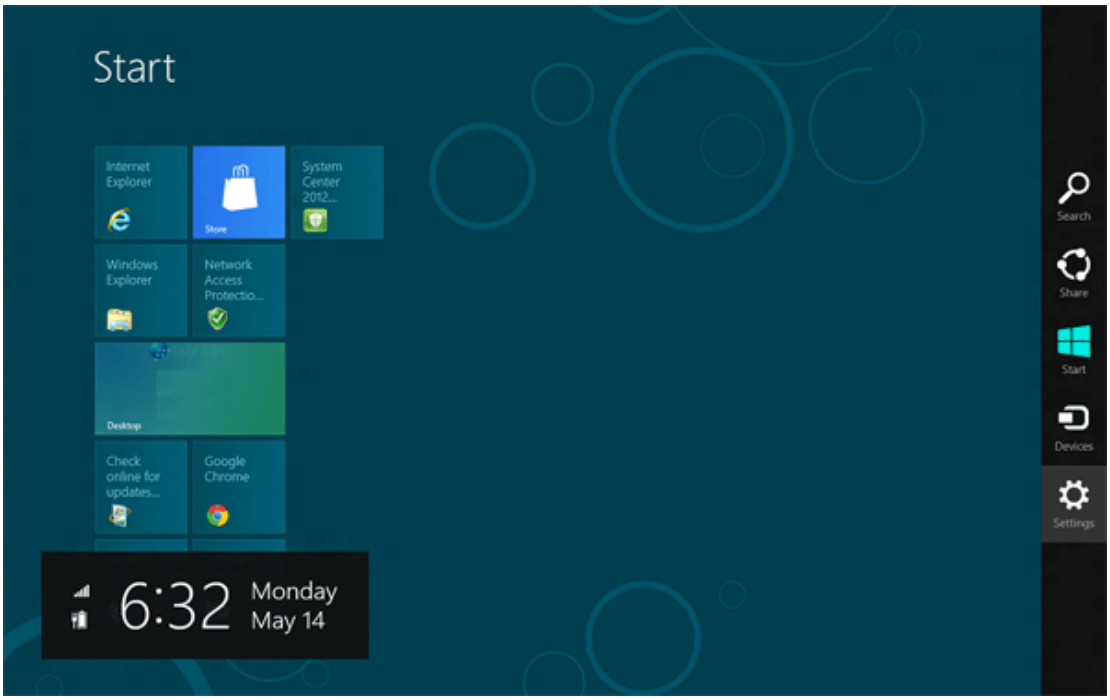

Figure 3: Charms Bar

2. Click **Settings**, and then click the network icon that is displayed in the right corner of the screen, as shown in Figure 4.

Figure 4: Network icon location

3. In **Networks**, click the VPN connection that you want to use, and then click **Connect**.

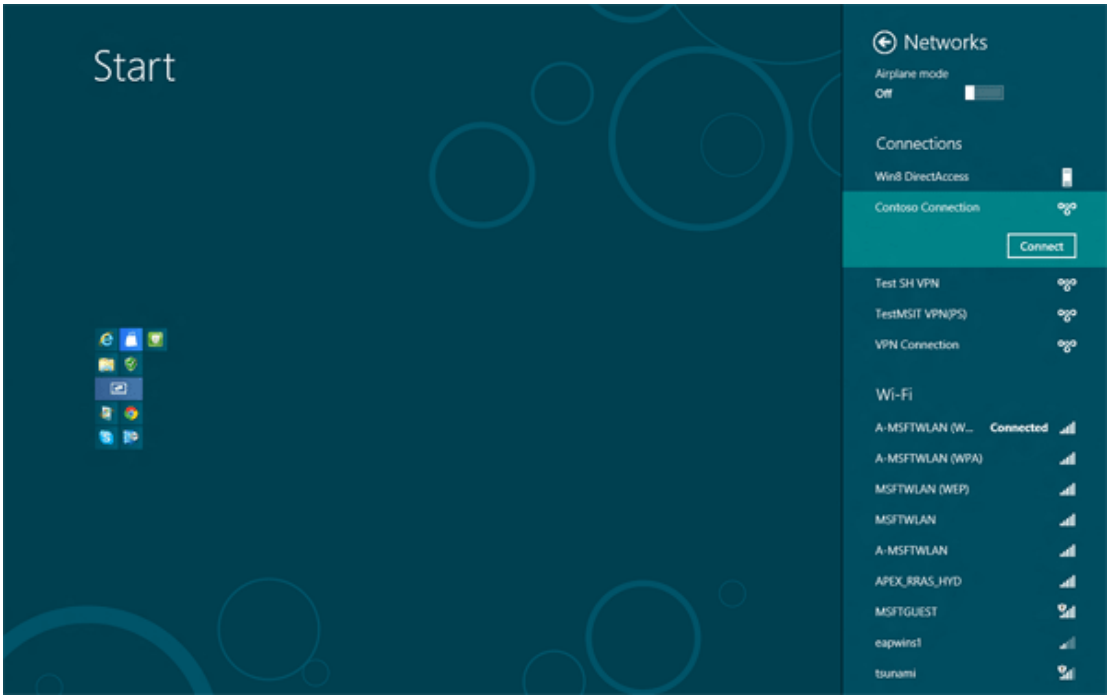

Figure 5: Connect to VPN

When you click **Connect**, a new connection is negotiated with the server by using the saved credentials. If you did not previously select **Remember Credentials** while creating the connection, you will be asked to provide credentials.

Tip

When you select **Remember my credentials** when you create your connection, you won't be asked for the credentials in further connection attempts. To clear the cached credentials, right-click the connection, and then select **Clear Cached Credentials**, as shown in the following figure.

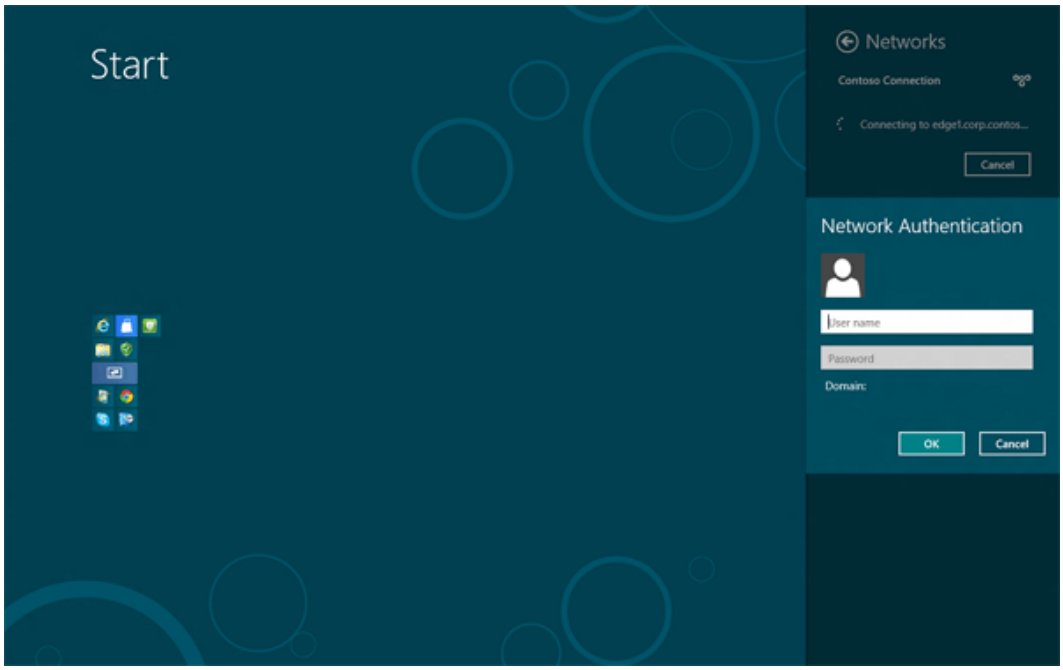

Figure 6: Specify VPN connection credentials

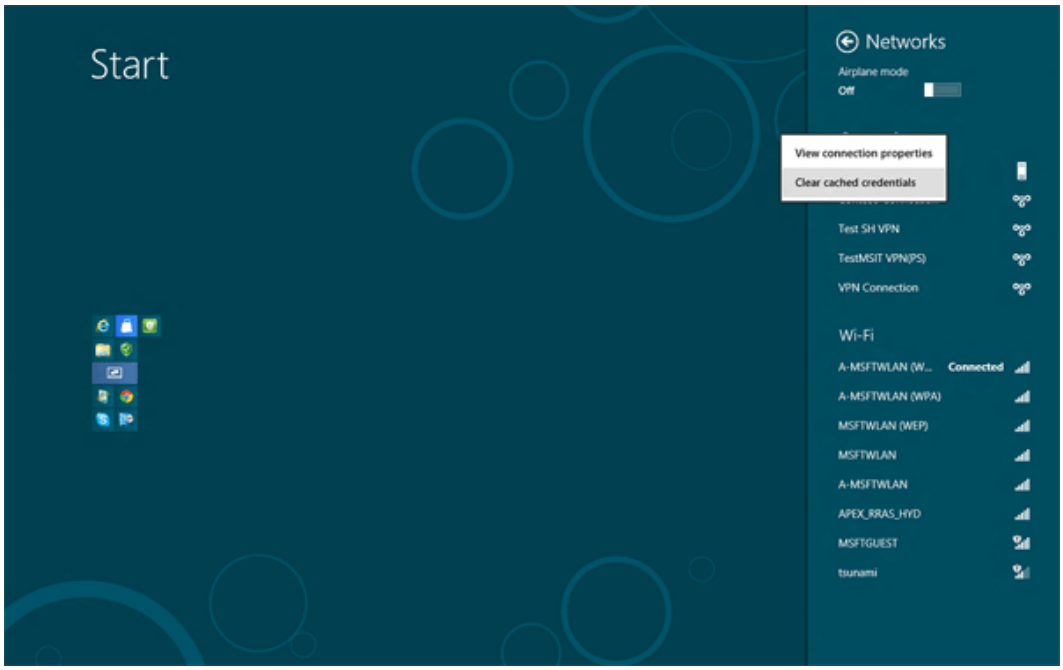

Figure 7: Clear cached credentials

Edit connection properties

To edit a connection that you have already configured, right-click the connection, and then select **Connection Properties**.

To access **Connection Properties**, in the **Networks**, click **View Connection Properties**.

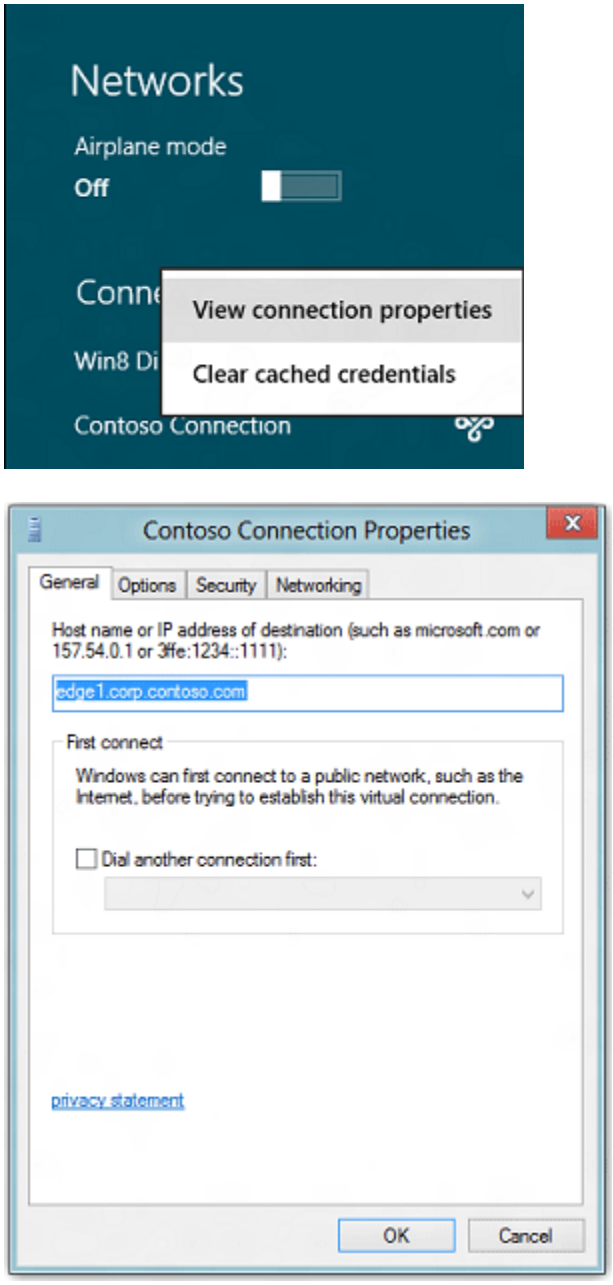

Figure 8: Connection Properties

Connection Properties has four tabs as shown in Figure 8. Click each tab to perform the following tasks:

- **General**

  Change the VPN server host name or IP address. Additionally, configure an interface to public network.
- **Options**

  Configure point-to-point protocol (PPP) settings and enable or disable caching of credentials.
- **Security**

  Configure tunnels, authentication, and encryption settings.
- **Networking**

  Configure transports (or protocols). The most common protocols are TCP/IPv4 and TCP/IPv6.

Delete a connection

On the **Start** screen, click **Search** in **Charms**, type **vpn**, and then click **View Network Connections** under **Settings**.

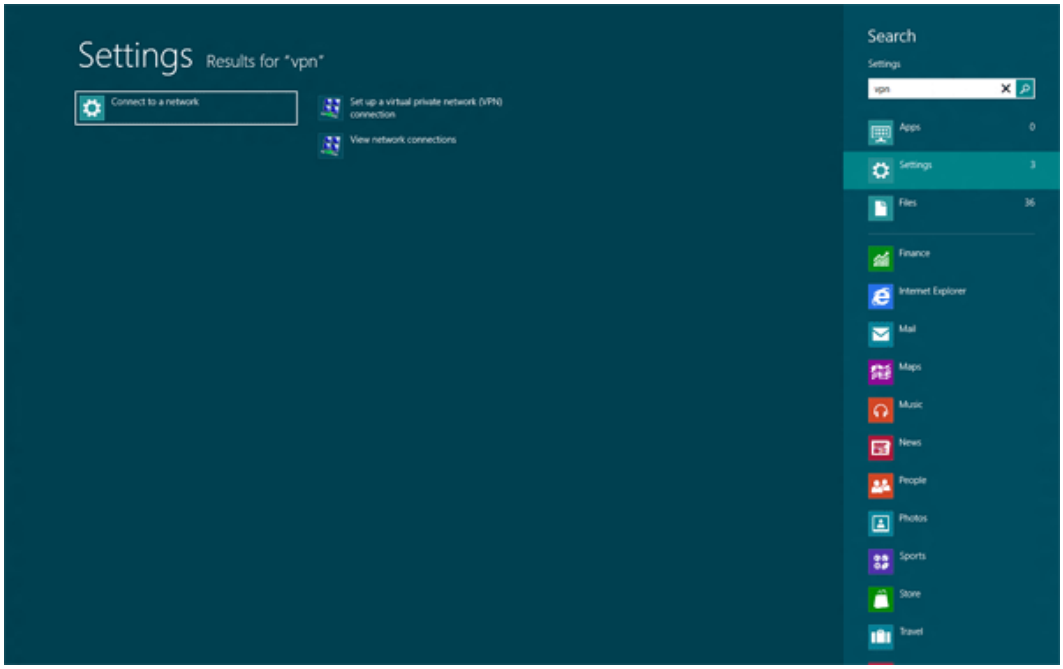

Figure 9: View Network Connections

Right-click a VPN connection, and then click **Delete**, as shown in Figure 10.

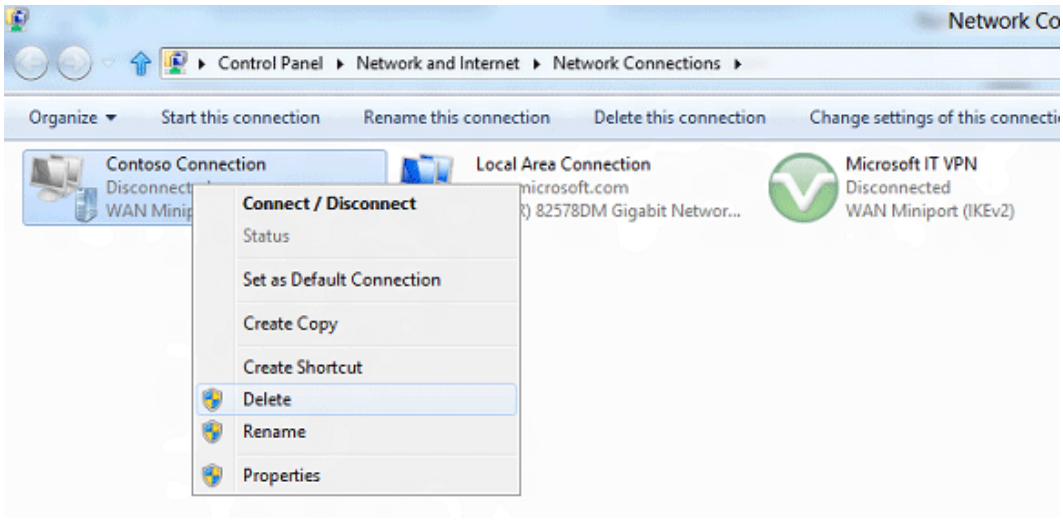

Figure 10: Deleting a connection

VPN Compatibility and Interoperability in Windows 8 and Windows Server 2012

Windows Server 2012 supports industry standards technology to provide customers with an open interoperable virtual private network (VPN). The following downloadable document contains information about third-party VPN software, including:

- Application compatibility of third-party VPN client software with Windows Server 2012.
- Interoperability of clients running Windows 8 for remote access with third-party VPN servers.
- Interoperability of clients running third-party software for remote access with computers running Windows Server 2012.

Virtual Private Networking with Windows Server 2012 and Windows 8: Compatibility and Interoperability

Windows Server 2012 R2 Test Lab Guide: Demonstrate VPN Auto trigger

Windows® 8.1 allows a Microsoft or third party VPN client to automatically connect to the corporate network. From an end user perspective, access to corporate intranet resources works the same as when inside the corporate network.

If auto triggering is enabled for a VPN profile, a remote access connection is triggered whenever a specific corporate resource or application is accessed. If user credentials are cached within the VPN profile, the VPN connection will be automatically established without any user interaction.

In this Guide

This guide provides step-by-step instructions for configuring new Windows VPN capability in a test lab to demonstrate functionality and operation. You will set up and deploy Remote Access based on the Windows Server® 2012 R2 Base Configuration using four server computers and one client computer. The resulting test lab simulates an intranet and the Internet.

Important

The following instructions are for configuring a Remote Access test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Attempting to adapt this Remote Access test lab configuration to a pilot or production deployment can result in configuration or functionality issues.

Test Lab Overview

In this test lab, Remote Access is deployed with:

- One computer running Windows Windows Server 2012 R2 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server
- One intranet member server running Windows Server 2012 R2 named EDGE1 that is configured as a Remote Access server
- One intranet member server running Windows Server 2012 R2 named APP1 that is configured as a general application server and web server
- One standalone server running Windows Server 2012 R2 named INET1 that is configured as an Internet DHCP server, DNS server, and web server
- One non-domain-joined, Internet-connected client computer running Windows 8.1 named CLIENT2 that is configured as a Remote Access VPN client

The Remote Access test lab consists of two subnets that simulate the following:

1. The Internet (131.107.0.0/24)
2. An intranet named Corpnet (10.0.0.0/24) separated from the Internet by EDGE1

Automatic triggering of VPN connections is not supported on domain joined machines. Configuration of auto triggering properties on domain joined machines is possible through Windows PowerShell so that the profile can be exported to another computer. However, auto triggering will not work in a VPN profile if the client is domain joined. This modular test lab demonstrates Remote Access VPN auto trigger functionality using a non-domain-joined Windows 8.1 client. This test client computer is used to show connectivity to corporate resources via a simulated Internet (the Internet subnet).

Hardware and software requirements

The following are required components of the test lab:

- The product disc or files for Windows Server 2012 R2
- The product disc or files for Windows 8.1
- Five computers or virtual machines that meet the minimum hardware requirements for Windows Server 2012 R2

Steps for Configuring the Remote Access VPN Auto-trigger Test Lab

There are four steps to follow when setting up a Remote Access VPN Auto-trigger test lab based on the Test Lab Guide: Base Configuration.

1. Set up the Base Configuration test lab.

   The Remote Access simplified setup test lab requires the Test Lab Guide: Base Configuration with both Corpnet and Internet subnets, as its starting point.

2. Configure EDGE1.

   EDGE1 is already a member server computer. For the Remote Access VPN auto-trigger test lab, EDGE1 must be configured as a Remote Access server with VPN access deployed

3. Configure INET1.

   INET1 is already a Windows Windows Server 2012 R2 that is configured as an Internet DNS server, web server, and DHCP server. For the Remote Access VPN auto-trigger test lab, INET1 must be configured with a DNS record for the internetsite.corp.contoso.com Internet resource

4. Configure CLIENT2.

   Add CLIENT2 as a standalone client computer running Windows Server 2012 R2. For this test lab, CLIENT2 will be used to test and demonstrate remote access VPN auto-trigger functionality and operation.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

This guide provides steps for configuring the computers of the Windows Server 2012 R2: Base Configuration test lab, configuring Remote Access in Windows Server 2012 R2, and demonstrating remote client connectivity. The following sections provide details about how to perform these tasks.

Step 1: Set up the Base Configuration Test Lab

Set up the Base Configuration test lab for both the Corpnet and Internet subnets using the procedures in the “Steps for Configuring the Corpnet Subnet” and “Steps for Configuring the Internet Subnet” sections of the Test Lab Guide: Base Configuration.

Step 2: Configure EDGE1

EDGE1 configuration for the Remote Access VPN Auto-trigger test lab consists of the following procedures:

1. Install the Remote Access server role on EDGE1
2. Deploy VPN access using the Remote Access Configuration wizard

The following sections explain these procedures in detail.

To install the Remote Access server role on EDGE1

1. In the **Dashboard** console of **Server Manager**, under **Configure this local server**, click **Add roles and features**.
2. Click **Next** three times to get to the server role selection screen.
3. In the **Select Server Roles** dialog, select **Remote Access**, and then click **Next**.
4. Click **Next** on the **Select Features** screen.
5. Click **Next** on the **Remote Access** screen.
6. On the **Role Services** screen, select **DirectAccess and VPN (RAS)**. Click **Add Features** when prompted, and then click **Next**.
7. Click **Next** on the **Web Server Role (IIS)** screen, and then click **Next** to accept default role services for IIS.
8. On the **Confirmation** screen, click **Install**.
9. Wait for the feature installations to complete, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
Install-WindowsFeature RemoteAccess -IncludeManagementTools

Add-WindowsFeature -Name DirectAccess-VPN -IncludeManagementTools
```

Deploy VPN access using the Remote Access Configuration wizard

Use the following procedure to deploy VPN Remote Access on EDGE1.

To configure VPN Remote Access on EDGE1

1. From the **Tools** menu in **Server Manager**, click **Routing and Remote Access**.
2. In the **Routing and Remote Access** console, right-click **EDGE1 (local)**, and then click **Configure and Enable Routing and Remote Access**.
3. When the **Routing and Remote Access Server Setup Wizard** launches. Click **Next**.
4. Select **Remote access (dial-up or VPN)**, and then click **Next**.
5. Select the checkbox for **VPN**, and then click **Next**.
6. Under Network interfaces, select **Internet**, and then click **Next**.
7. On the IP Address Assignment screen, select **Automatically**, and then click **Next**.
8. On the Managing Multiple Remote Access Servers screen, select **No, use Routing and Remote Access to authenticate connection requests**, and then click **Next**.
9. On the **Summary** screen, click **Finish**. When you are prompted to enable DHCP Relay Agent, click **OK**.
10. When the Routing and Remote Access Server Setup Wizard completes, expand **EDGE1 (local)** in the Routing and Remote Access console tree, and select **Remote Access Logging & Policies**.
11. When the Routing and Remote Access Server Setup Wizard completes, expand **EDGE1 (local)** in the Routing and Remote Access console tree, and select **Remote Access Logging & Policies**.
12. Right-click **Remote Access Logging and Policies**, and select **Launch NPS**.
13. Double-click the first default policy named **Connections to Microsoft Routing and Remote Access server**.
14. On the **Overview** tab, under **Access Permission**, select **Grant Access**. **Grant Access if the connection request matches this policy radio button**.
15. Click to select **Ignore user account dial-in properties**.
16. Click **OK** to close the Connections to Microsoft Routing and Remote Access server properties window.
17. Close the **Network Policy Server** window.

Step 3: Configure INET1

INET1 is already a Windows Server 2012 R2 server that is configured as an Internet DNS server, web server, and DHCP server. For the Remote Access VPN auto-trigger test lab, INET1 must be configured with a DNS record for the internetsite.corp.contoso.com Internet resource.

INET1 configuration for the Remote Access VPN Auto-trigger test lab consists of the following procedure:

- Create a new DNS host record for internetsite.corp.contoso.com on INET1

The following section explains this procedure in detail.

Create a new DNS host record on INET1

Use the following procedure to create the internetsite record on INET1.

To create a new DNS entry on INET1

1. From the Tools menu in Server Manager, click **DNS**.
2. In the console tree of DNS Manager, expand **INET1**, and click **Forward Lookup Zones**.
3. Right-click **Forward Lookup Zones**, click **New Zone**, and then click **Next**.
4. On the **Zone Type** page, click **Next**.
5. On the **Zone Name** page, type **corp.contoso.com**, and then click **Next**.
6. Click **Next** twice to accept defaults for zone file and dynamic update, and then click **Finish**.
7. In the console tree, expand **Forward Lookup Zones**, right click **corp.contoso.com**, and then click **New Host (A or AAAA)**.
8. In **Name**, type **internetsite**. In **IP address**, type **131.107.0.1**. Click **Add Host**.
9. Click **OK**, and then click **Done**.
10. Close the DNS Manager console.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
Add-DnsServerPrimaryZone -Name corp.contoso.com -ZoneFile corp.contoso.com.dns

Add-DnsServerResourceRecordA -ZoneName corp.contoso.com -Name internetsite -IPv4Address
131.107.0.1
```

Step 4: Add and Configure CLIENT2

Use the following procedures to demonstrate remote access connectivity with CLIENT2:

1. Install the operating system on CLIENT2
2. Connect CLIENT2 to the Internet subnet and configure remote access with VPN
3. Configure and test VPN auto-trigger functionality on CLIENT2
4. Connect CLIENT2 to the Corpnet subnet and verify that auto-trigger functionality is disabled

The following sections explain these procedures in detail.

Install the operating system on CLIENT2

To install the operating system on CLIENT2

1. Start the installation of Windows 8.1.
2. When you are prompted for a PC name, type **CLIENT2**.
3. When you are prompted by the **Settings** dialog, click **Use express settings**.
4. At the Log on prompt, click **Sign in without a Microsoft account**. Click **Local account**. If CLIENT2 does not have Internet access during setup, you will be prompted to **Create a local account**.
5. When you are prompted for a user name, type **User1**. Type a strong password twice, and type a password hint. Click **Finish**.
6. Connect CLIENT2 to the Internet subnet. When prompted to automatically connect to devices on this network, click **Yes**.

Connect CLIENT2 to the Internet subnet and configure remote access with VPN

To configure VPN access to Corpnet from the Internet

1. Connect CLIENT2 to the Internet subnet. Once the network determination process completes, the network icon should indicate Internet access.
2. Open **PC Settings** (click the **Settings** charm, and then click **Change PC Settings**)
3. In the PC Settings console tree, click **Network**.
4. Click **Add a VPN Connection**.
5. In the **Add a VPN connection** dialog, select **Microsoft** as the VPN provider.
6. Under **Connection name**, type **Contoso VPN Connection**.
7. Under **Server name or address**, type **edge1.contoso.com**.
8. Under **Type of sign-in info**, select **User name and password**.
9. Type **CORP\User1** and the password for User1 in the User name and Password fields. Verify that the checkbox is selected for **Remember my sign-in info**.
10. Click **Save** to close the **Add a VPN connection** dialog.

Configure and test VPN auto-trigger functionality on CLIENT2

To configure and test VPN auto-trigger functionality

1. From the desktop on CLIENT2, right-click the **Start** icon. Click **Windows PowerShell (Admin)**.
2. Enable Split-tunneling for the VPN connection by entering the following command, followed by **ENTER**.

   Set-VpnConnection –Name “Contoso VPN connection” –SplitTunneling $true
3. Configure CLIENT2 to automatically connect the Contoso VPN connection whenever resources on the Contoso intranet are resolved. In the Windows PowerShell window, type the following command to configure name-based VPN triggering for resources in the

corp.contoso.com domain, and then press **ENTER**.

Add-VpnConnectionTriggerDnsConfiguration -Name "Contoso VPN Connection" –DnsSuffix corp.contoso.com -DnsIPAddress 10.0.0.1

4. Configure an exemption for a web site in the corp.contoso.com DNS namespace that should be accessed over the Internet without causing the VPN connection to be initiated. In the Windows PowerShell window, type the following command to configure an exemption for the internetsite.corp.contoso.com website, and then press **ENTER**.

   Add-VpnConnectionTriggerDnsConfiguration –Name "Contoso VPN Connection" -DnsSuffix internetsite.corp.contoso.com

5. Configure CLIENT2 to always connect to the Contoso intranet whenever the legacy application Notepad is used. In the Windows PowerShell window, type the following command to configure application-based VPN triggering for the Notepad application, and then press **ENTER**.

   Add-VpnConnectionTriggerApplication -Name "Contoso VPN Connection" –ApplicationID "C:\Windows\System32\notepad.exe"

6. Configure CLIENT2 to always connect to the Contoso intranet whenever the modern application Calculator is used. First, in the Windows PowerShell window, type the following command to determine the application package family name associated with the Windows Calculator app, and then press **ENTER**.

   Get-AppXPackage -Name Microsoft.WindowsCalculator

   Note the string listed for PackageFamilyName.

   Next, type the following command to configure application-based VPN triggering for the Windows Calculator app, and then press **ENTER**.

   Add-VpnConnectionTriggerApplication -Name "Contoso VPN Connection" –ApplicationID Microsoft.WindowsCalculator_8wekyb3d8bbwe

7. View the newly added VPN triggers associated with the connection by typing the following command followed by **ENTER**.

   Get-VpnConnectionTrigger "Contoso VPN Connection"

   Examine the output of Get-VpnApplicationTrigger. The Contoso VPN Connection will be triggered for automatic connection without user interaction whenever resources in corp.contoso.com are accessed, with the exception of the internetsite resource, unless the client is already connected to the Trusted Network corp.contoso.com. Additionally, the Contoso VPN Connection will be triggered for automatic connection whenever the applications notepad.exe or Windows Calculator are launched.

8. From the Windows PowerShell prompt, type **Get-DnsClientNrptPolicy** and then press ENTER. The contents on the Name Resolution Policy Table are displayed.

   Examine the output of Get-DnsClientNrptPolicy. Any name resolution queries directed to the namespace corp.contoso.com will require resolution by the Contoso intranet DNS server DC1 at IP address 10.0.0.1. The only exception to this rule is the internetsite.corp.contoso.com name, which will use the DNS server assigned to the CLIENT2 Internet connection.

9. First, access the NRPT exemption to verify that the VPN connection is not triggered. Click the Internet Explorer icon to launch IE. In the address bar, type **internetsite.corp.contoso.com**, and then press **ENTER**. You should see the default Internet Information Services welcome screen.
10. Click the network icon in the System Notification Area. Note that **Contoso VPN Connection** is not connected. This demonstrates that internetsite.corp.contoso.com is accessible directly over the Internet without a VPN connection. Click **Contoso VPN Connection** and note the new checkbox for **Let apps automatically use this VPN connection**.
11. Next, launch the Notepad application to automatically trigger the VPN connection. From the Windows PowerShell prompt, type **notepad.exe** and then press **ENTER**.
12. Click the network icon in the System Notification Area. Note that **Contoso VPN Connection** is now listed as **Connected**.
13. In the Notepad window, click **File**, and then click **Open**. Next to **File name**, type **\\app1.corp.contoso.com\Files\example.txt**, and then click **Open**.
14. You should see the contents of the Example.txt file shared on the APP1 server.
15. Click the network icon in the System Notification Area. Click **Contoso VPN Connection**, and then click **Disconnect**.
16. Click **Contoso VPN Connection** again. Note that manually disconnecting the VPN disables the auto-connect setting. Select the checkbox for **Let apps automatically use this VPN connection** to re-enable automatic triggering.
17. Click the Internet Explorer icon to launch IE. Verify that you can access the website on **http://app1.corp.contoso.com**.
18. Click the network icon in the System Notification Area. Note that **Contoso VPN Connection** is now listed as **Connected**.

Connect CLIENT2 to the Corpnet subnet and verify that auto-trigger functionality is disabled

To verify that VPN is not triggered when connected to the Trusted Network

1. Connect CLIENT2 to the Corpnet subnet. CLIENT2 will receive a DHCP address with the corp.contoso.com connection-specific DNS suffix.
2. From the Windows PowerShell prompt, type **ipconfig** and then press ENTER. Note that the value of **Connection-specific DNS suffix** is listed as **corp.contoso.com**.
3. Review the VPN triggers configuration associated with the connection by typing the following command followed by E**NTER**.

   Get-VpnConnectionTrigger "Contoso VPN Connection"

   Examine the output of Get-VpnApplicationTrigger. While CLIENT2 is connected to the Corpnet subnet, access to the Trusted Network corp.contoso.com will prevent automatic triggering of the VPN connection.
4. From the Windows PowerShell prompt, type **notepad.exe** and then press **ENTER**.

5. Click the network icon in the System Notification Area. Note that **Contoso VPN Connection** is no longer automatically connected when the Notepad application is launched.
6. Click the Internet Explorer icon to launch IE. Verify that you can access the website on **http://app1.corp.contoso.com**. This website is now accessed directly over the LAN, and there is no need to establish a VPN connection.

Snapshot the Configuration

This completes the Remote Access Auto-trigger VPN test lab. To save this configuration so that you can quickly return to a working Remote Access configuration from which you can test other Remote Access modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots **Remote Access Auto-trigger VPN**. If your lab uses physical computers, create disk images to save the DirectAccess simplified IPv4-only test lab configuration.

Deploy and Configure a VPN for Devices Running Windows RT

Employees who want to use personal devices (such as laptops, slates, and tablets) to connect to corporate networks when they are not joined to the domain is a topic of high interest for businesses. Virtual private network (VPN) connectivity is available in all Windows 8 devices to enable remote access to the corporate network. This capability is documented in detail in a set of recently published Microsoft TechNet articles. These articles explain using a VPN with x86-based devices and devices running Windows RT.

This topic specifically focuses on how to deploy and configure a VPN when employees are using devices running Windows RT. Additionally, this topic identifies limitations in VPN support in Windows RT, describes the VPN configurations that are possible with major third-party VPN servers, and outlines the client-side experiences.

Work with third-party VPN servers

Devices running Windows RT come with a VPN client implementation that is fully compatible with Microsoft Routing and Remote Access Server (RRAS). Third-party VPN solutions from major networking vendors are in widespread use by enterprises. Windows RT (unlike x86-based Windows 8) does not support the installation of native VPN client applications from these vendors. Thus, the only way to connect to third-party VPN servers is by using the VPN client implementation that is provided in Windows RT.

The VPN client in Windows RT can connect by using standard PPTP, L2TP/IPsec, and IKEv2 protocols, which are widely supported. However, SSL-based VPN connectivity to third-party servers is not possible because SSL VPN implementations are proprietary and distinct to each vendor. Network administrators must configure their third-party VPN servers to allow connections by using one of the supported standard protocols.

Verify third-party VPN interoperability

Although connections to any third-party server by using one of the previously mentioned protocols should work, we have specifically verified interoperability of the VPN client in Windows RT with popular versions of Cisco and Juniper server software. Following table depicts the interoperability details.

Third-party VPN server solution	Tested operating system version	Tunnels supported	Authentication methods supported	Crypto Suits supported
CISCO (2951 VPN Server)	IOS 15.1.4	• PPTP • L2TP/IPSec with certificate • IPSec (IKEv2) with certificate	• CHAP • PSK (over v4 and v6) • Computer certificate	**IPsec**: AH auth: HMAC_SHA_1_96, HMAC_MD5_96 **ESP encryption**: AES_128, CBC_3DES, CBC_DES, None IKEv2: Encryption: 3DES, AES_128, AES_192, AES_256 Integrity:SHA1, SHA_256, SHA_384 DH Group: DH2
Juniper (SSG series)	6.2.0r5.0	• L2TP/IPSec with certificate	• CHAP • PSK (over v4 and v6) • Computer certificate	

We recommend using the L2TP/IPSec or IKEv2 protocols to connect to third-party servers.

Note

The Extensible Authentication Protocol (EAP) method is supported by the native VPN client in Windows RT, but it is not tested for interoperability with third-party solutions.

For more information, seeVirtual Private Networking with Windows Server 2012 and Windows 8: Compatibility and Interoperability.

VPN deployment options

The following apply to both Microsoft RRAS as well as 3rd party VPN servers.**Password-based deployment**

Windows RT can establish a remote access connection through a variety of password-based PPP authentication protocols, which includes

- Non-EAP methods like PAP, CHAP, MS-CHAPv2, etc., or
- EAP methods (PEAP, EAP-TTLS etc.)

The password-based methods which are not considered secure (Non-EAP password based methods) should be deployed with secure TLS methods such as PEAP, EAP-TTLS.

Important
Microsoft recommends using PEAP or TTLS based authentication methods (both available in-the-box in Windows RT) for corporate deployments.

Certificate-based deployments

To support this, an organization needs to support enrolling/importing user certificates. There are many options to deploy a Public Key Infrastructure (PKI) for this purpose and it depends upon the organization's policy and deployment capabilities. More details on the Certificate enrollment web service can be found on TechNet.

Multi-site VPN deployment

Organizations with more than one business location may choose to have multi-site VPN server deployments. In this case, organizations should use Global Server Load Balancing (GSLB) for resolution of the address and redirecting the client to the nearest VPN server site. If for some reason, that is not an option, then the IT administrator can provide VPN connection profiles per site, thus providing the user with a choice to connect the desired remote server.

Windows RT VPN client provisioning

VPN profile provisioning includes setting up the profile on the client device and plumbing the relevant settings required for corpnet access. The following describes the VPN connection profile provisioning methods that are available for Windows RT devices.

Get Connected Wizard (GCW)

GCW is the legacy VPN connection setup wizard, and is available on Windows RT. It provides a simple method for creating VPN connections from the client device itself. If the VPN deployment

is simple (password based) and has no complexities like provisioning NAP, Routes, etc., then users can easily create the profile themselves. In this case, a user just configures the remote access server name or IP address as instructed by his network administrator.

If the profile is a more complex then the administrator needs to provide the detailed instructions and/or the configuration scripts.

The TechNet article, Set Up and Edit VPN Connections in Windows 8, provides more details.

Single-Click (PowerShell) scripts

Network administrators can create simple PowerShell scripts to help users install and configure the VPN connections on unmanaged devices. These scripts serve dual purposes:

- Create a VPN profile, and
- Configure network specific settings (e.g. plumbing routes, NAP settings, etc.)

The admin and user experiences for installing the VPN profile using PowerShell scripts are discussed below.

Admin experience

- Admin creates a PowerShell script and setup instructions. The PowerShell cmdlets to create the VPN profiles are very simple and comprise a few lines of code.
    a. PS examples, sample scripts & documentation are available on TechNet to help admins create scripts as per their network deployment.
- Admin delivers the script and instructions to users via email or a web portal. The admin needs to inform users to run the script again each time the script gets updated

User Experience

- The user navigates to the web portal as instructed and logs-in with his domain credentials.
- He simply clicks on the VPN configuration link and the setup script gets downloaded to the client device
- User clicks on the “Run” option from browser download UI and profile gets installed

Advanced Scenarios

Administrators can club many advanced configurations with this single-click VPN provisioning script to provide a complete client provisioning story for users. These advanced settings could involve

- Deploying additional policies like IPsec, routes publishing or NAP for VPN connections
- Web-proxy configuration for VPN connection
- Intranet settings for VPN connection
- Connection specific DNS Suffix plumbing for VPN connectionv

Provisioning via System Center

On Windows RT devices, an integrated management client called **Company Apps** is available that can communicate with the management infrastructure and receive LOB apps and VPN connection configuration delivered by the System center.

The IT Administrators can generate the VPN connection configuration XML and configure it with System Center manager to push down to the client devices enrolling for the service.

The clients can browse to **Company Apps** (in Control Panel) and enroll for the device management infrastructure by using their domain credentials. On successful enrollment, System center management infrastructure pushes the VPN connection configuration information to the client device and the **Self-Service Portal** agent (on WindowsRT devices) configures the VPN connection on the client device without any user intervention.

The admin and user experiences for profile setup via SSP are discussed below.

Admin experience

1. Open System Center management UI and configure VPN profile configuration XML for clients.
2. Admin can update the profile (if needed) again through System Center server management UI and no further action is required from user.

User experience

User needs self-provisioning for WoA machine to get enrolled for the VPN setup.

- Authorized users navigates to the new Control Panel applet called Company apps on their WOA devices and provide their company email address and password.
- The SSP agent then automatically configures VPN and other features on the client device based on the instructions received from organization's management infrastructure.

Note

The above device management solution is not yet released. The release along with documentation is slated for Jan 2013.

Profile management

After a VPN profile is created, it may need to be modified occasionally. This section describes common profile settings and how these settings can be modified by the administrator or the user.

- Edit and delete experiences for GCW profiles are described in the Windows 8 VPN – Get Connected TechNet article.
- For VPN profiles created using powershell scripts (managed by the administrator),
    - The administrator can create and share another single-click script to remove the profile and corresponding settings from client device.
    - The new script may also be used to
    - Delete an existing VPN connection, and then provision a new one, or
    - Edit an existing connection, or add a new connection
- If the IT administrator wants to make any changes to the VPN profile deployed on client systems, he can simply create a new VPN profile configuration XML and update that in management infrastructure. The new VPN profile will be used for provisioning from then on.

The Windows 8 VPN – Get Connected TechNet article provides more details on the client experience for VPN profile management.

Enabling split-tunneling

Split-tunneling, which allows Internet-bound traffic to bypass the VPN connection, can be set up as follows on Windows RT devices:

To enable split-tunneling

1. Navigate to **View Available Networks** by clicking on the network icon in taskbar and select the VPN connection
2. Right-click the VPN connection and select **Properties** from the contex menu
3. Navigate to **Networking** tab; select **Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.
4. In **Properties** click **Advanced**.
5. Under **General** in **Advanced TCP/IP settings**, unselect the option **Use default gateway on remote network to enable split-tunneling**.

Rather than the user executing the above steps, the administrator can also enable split-tunneling using the configuration PowerShell cmdlet as described in this TechNet article.

Implications on routes

For split-tunneled VPN connections, the corporate network routes have to be configured on the client device after the VPN connection is established. The administrator can do this using one of the following two methods:

DHCP based routes publishing – When a client is connected to the corporate network via a VPN connection, it gets its IP address from the DHCP server designated for the same purpose. This DHCP server can also supply the required routes along with other configuration.

One-click script (scheduled tasks) – The administrator can enable the configuration of routes through a script. The script adds a scheduled task in the client device, which is triggered every time the device establishes a connection with the remote VPN server, and plumbs the required routes on the client, enabling it to access corporate resources. If the administrator is already using a single-click script to provision other settings (IPSec policies, DNS suffix etc.), then route-provisioning can also be clubbed with the same script.

Provisioning web proxy and Intranet settings for VPN connection

In the case of force-tunneling, a web proxy needs to be configured on the Windows RT devices for them to be able to access Internet sites. Also, the intranet settings have to be enabled for a VPN connection to suppress repeated credential prompts while accessing intranet sites. The following sections discuss these configurations for Windows RT devices.

Web-proxy

Force-tunneling requires the VPN client to use corpnet web proxy for all the internet traffic. The administrator can either provide the required web proxy information to the users via e-mail and

have them configure the web-proxy themselves, or he can configure a VPN-specific web proxy by using a PowerShell cmdlet. This cmdlet can be included in the one-click script used for VPN connection provisioning or can be distributed separately.

Intranet Settings

Enabling intranet settings is another requirement for accessing corporate sites without repeated credential prompts for the user while connected via a VPN connection. Intranet settings determine whether accessed web sites are a part of the local intranet zone and require domain credentails. If intranet settings are enabled for a VPN connection, the VPN credentials are automatically used for authenticating the user while accessing the local intranet sites.

For unmanaged Windows RT devices, VPN connection is not automatically recognized and intranet settings are not turned ON by default. This results in a credentials prompt every time the user navigates to a local intranet site. To resolve this, users need to turn on the intranet settings by following the procedure below:

To turn on intranet settings

1. Open **Internet Explorer**, and click **Tools**.
2. On the **Tools** menu, click **Internet options**, and then click the **Security** tab.
3. Click **Local Intranet**, and then click **Sites**.
4. In **Local Intranet**, clear the **Automatically detect intranet network** check box, and select the other three options.
5. Click **OK**, and then **Apply**.

Alternatively, the administrator can configure certain registry settings on devices running Windows RT by using Windows PowerShell scripts when the client is connected to corporate network through the VPN.

Connect to the VPN

The Get Connected Wizard is located in the **View Available Networks** UI (with a few exceptions). It provides a seamless connection experience for the user. Users do not need to specify the authentication or encryption protocols because they are negotiated with the server during the connection.

For more information, see Set Up and Edit VPN Connections in Windows 8.

Use one-time password for authentication

An organization can use one-time password (OTP) authentication by completing the following configuration steps on clients and the server.

On clients

Configure the VPN client profile in Windows RT to use TTLS-PAP authentication protocol. TTLS-PAP is a password-based, secure EAP method in Windows 8 that is used to transport the user's user name and password, and the OTP token.

On servers

- Install RSA Authentication Manager, and configure it to accept and process TTL-PAP authentication requests.
- Configure the VPN server (Microsoft RRAS or a third-party server) to redirect all authentication requests to RSA Authentication Manager.

To authenticate with an OTP, the user enters the OTP user name and passcode for the VPN profile. The passcode is the user's password followed by the OTP token code. TTLS-PAP protocol accesses the RSA Authentication Manager, which authenticates the user name, password, and token code.

Note

- This method for OTP authentication is currently being validated by Microsoft on third-party VPN servers.

This authentication method does not provide the following capabilities, which are offered by the regular RSA OTP authentication method.

Feature	Description	Workaround or Solution
PIN changes	During first time authentication or PIN expiry, the user is required to set a new PIN. After the PIN is changed, the authentication is completed and the user is connected.	Use the RSA Security Center, which is a tool provided by RSA with the EAP client. The tool enables users to communicate with the RSA Authentication Manager directly and reset the PIN. The procedure is required to be completed inside the corporate network.
Next token code	If user authentication succeeds after several failed attempts, the authentication manager challenges the user to enter the next token code that flashes on the device and validates it.	This feature can be disabled in the RSA Authentication Manager, and then users do not have to enter the next token code.

Migrate Remote Access to Windows Server 2012

Routing and Remote Access Service (RRAS) was a role service in Windows Server operating systems prior to Windows Server 2012 that enabled you to use a computer as an IPv4 or IPv6 router, as an IPv4 network address translation (NAT) router, or as a remote access server that

hosted dial-up or virtual private network (VPN) connections from remote clients. Now, that feature has been combined with DirectAccess to make up the Remote Access server role in Windows Server 2012 This guide describes how to migrate a server that is hosting the Routing and Remote Access service (in Windows Server 2008 R2 and other down-level versions) to a computer that is running Windows Server 2012.

Note

Your detailed feedback is very important, and helps us to make Windows Server Migration Guides as reliable, complete, and easy to use as possible. Please take a moment to rate this topic, and then add comments that support your rating. Describe what you liked, did not like, or want to see in future versions of the topic. To submit additional suggestions about how to improve Migration guides or utilities, post on the Windows Server Migration forum.

About this guide

Migration documentation and tools ease the migration of server role settings and data from an existing server to a destination server that is running Windows Server 2012. By using the tools that are described in this guide, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help to eliminate possible conflicts that might otherwise occur during the migration process. For more information about installing and using the migration tools on the source and destination servers, see the Windows Server Migration Tools Installation, Access, and Removal Guide (http://go.microsoft.com/fwlink/?LinkId=247607).

Target audience

This document is intended for information technology (IT) administrators, IT professionals, and other knowledge workers who are responsible for the operation and deployment of the Remote Access servers in a managed environment. Some scripting knowledge may be required to perform some of the migration steps that are contained in this guide.

What this guide does not provide

This guide does not describe the architecture or detailed functionality of the Remote Access role. The following scenarios are not supported in this migration guide:

- Any process for an in-place upgrade, in which the new operating system is installed on the existing server hardware by using the **Upgrade** option during setup
- Clustering and multisite scenarios
- Migrating more than one server role

  If your server is running multiple roles, it is recommended that you design a custom migration procedure that is specific to your server environment and based on the information that is provided in this and other role migration guides.

Supported migration scenarios

This guide provides instructions for migrating an existing server to a server that is running Windows Server 2012.

Caution

This guide does not contain instructions for migration when the source server is running multiple roles. If your source server is running multiple roles, some migration steps in this guide, such as those for migrating user accounts and network interface names, can cause other roles that are running on the source server to fail.

Supported operating systems

This guide provides instructions for migrating data and settings from an existing server that is being replaced by a new physical or virtual 64-bit server with a clean-installed operating system, as described in the following table.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	win8_Server_2	x64-based
x86- or x64-based	Windows Server 2003 R2	win8_Server_2	x64-based
x86- or x64-based	Windows Server 2008, full installation option only	win8_Server_2	x64-based
x64-based	Windows Server 2008 R2, full installation option only	win8_Server_2	x64-based
x64-based	Win8_Server_2	win8_Server_2	x64-based

- The versions of operating systems shown in the preceding table are the earliest combinations of operating systems and service packs that are supported. Newer service packs are supported.
- Migration with the destination server already having DirectAccess configured is not supported.
- The Foundation, Standard, Enterprise, and Datacenter editions of the Windows Server operating system are supported as either source or destination servers. This includes migrating across editions. For example, you can migrate from a server running Windows Server 2003 Standard to a server running Windows Server 2012.
- Migrations between physical operating systems and virtual operating systems are supported.
- Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shared resources from a computer that is running

Windows Server 2008 R2 in the French system UI language to a computer that is running Windows Server 2012 in the German system UI language.

Note
The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

- Both x86- and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2008 R2 and win8_server_2 are x64-based.

Supported role configurations

The following is a broad list of the migration scenarios that are supported for Remote Access. All settings under these scenarios are migrated.

- **DirectAccess (supported in Windows Server 2012 to Windows Server 2012 migration only)**
- **VPN server**
- **Dial-up server**
- **Network address translation (NAT)**
- **Routing**, with the following optional components:
  - DHCP Relay Agent
  - Routing Information Protocol (RIP)
  - Internet Group Management Protocol (IGMP)

In addition to the above scenarios, migration also automatically adjusts configuration of the destination server to account for features that are no longer supported and to support features that are new in Windows Server 2012 and not supported on earlier versions of Windows.

Migration dependencies

If a local or remote NPS server that is used for authentication, accounting, or policy management must also be migrated, then migrate the NPS service before migrating Remote Access. For more information, see Migrate Network Policy Server to Windows Server 2012.

If you are upgrading from Windows 2008 R2 DirectAccess to Windows Server 2012, ensure that all the DirectAccess configuration settings have been applied on the Windows 2008 R2 server. It is possible to save settings through the console but not apply them. Before upgrading, ensure that the saved settings have also been applied.

Migration components that are not supported in all operating system versions

The following Remote Access components are not supported by all operating systems:

Component	UI Dialog/Settings	Action

New Components, not available on Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2		
Specifying adapter to obtain DNS/WINS addresses	RAS Properties – IPv4 Tab: Adapter	This component is not supported on Windows Server 2003. The default value should be used for this setting on the target computer.
SSTP	SSTP ports	SSTP is not supported on Windows Server 2003. SSTP ports should be enabled on the target computer. The number depends on the default value for the SKU on the target computer
IPv6	1. RAS Properties – General Tab: IPv6 router check-box and corresponding radio-buttons, IPv6 Remote access server 1. RAS Properties – IPv6 Tab: All settings 1. Demand-Dial (VPN/PPPoE) properties – Networking tab – TCP/IPv6 1. Demand-dial (VPN/PPPoE) – IPv6 filters for connection initiation 1. IPv6 - Router	• IPv6 is not supported on Windows Server 2003. It should be disabled on the target computer if DirectAccess is not deployed (legacy VPN). Under the General tab of RRAS properties, IPv6 Router and IPv6 Remote access server should not be selected. • The adapter setting under the IPv6 tab of RAS properties was introduced in Windows Server 2008 R2 for IKEv2. It was not present in Windows Server 2008. During migration this setting should be set to the default value on the target computer for Windows Server 2008 and be “as is” for Windows Server 2008 R2 and Windows Server 2012
IP Filters under RA Logging	Remote Access Logging &	Windows Server 2003 and

and Policies	Policies – IP Filters	Windows Server 2008 do not have an option to create IP filters under Remote Access Logging and Policies. Hence there would be no filters to migrate.
Automatically obtaining IPv6 address	Demand-Dial (VPN/PPPoE) properties - Networking tab - IPv6 Properties – 'Obtain IP address automatically' radio-button	This setting is not present in Windows Server 2008. The value of this setting should be set to default on the target computer.
IKEv2	1. IKEv2 ports 1. RAS Properties – Security Tab: computer certificate authentication for IKEv2 1. RAS Properties – IKEv2 Tab: All settings	• IKEv2 is not supported on Windows Server 2003 and Windows Server 2008. IKEv2 ports should be enabled on the target computer. The number would depend on the default value for the SKU on the target computer. • Default values should be used for all IKEv2 settings on the target computer.
SSTP Cert. Selection	RAS Properties – Security Tab: 'Use HTTP' check-box, drop-down to select certificate, crypto binding settings	This component is not supported on Windows Server 2003 and Windows Server 2008. Default values should be used for all these settings on the target computer.
VPN Accounting	Remote Access Logging & Policies – Accounting: Logging failure action settings under 'SQL Server Logging Properties' and 'Log File Properties'	These are not present in Windows Server 2008. The default value should be used on the target computer.
Deprecated Features: Not available on Windows Server 2008, Windows Server 2008 R2, or Win8_Server_2		

SPAP, MS-CHAP, EAP-MD5 protocols and related settings	VPN/PPPoE Demand-dial interface properties - Security tab	SPAP, EAP-MD5 and MS-CHAP settings are not supported on Windows Server 2008 R2, or win8_Server_2, and will not be migrated.
Local Area Connection interface configuration under Routing	Routing – General – Local Area Connection properties – Configuration tab	This tab provides settings to configure how an IP address should be obtained for this interface. It is only present on Windows Server 2003 and will not be migrated.
RAS Firewall (integrated with NAT)	1. NAT – Interface – NAT/Basic Firewall Tab 1. NAT – Interface – ICMP Tab	Windows Server 2003 supported RAS firewall functionality which was removed in Windows Server 2008. These settings will not be migrated.
Weak Encryption settings		Weak encryption is supported on Windows Server 2003 but on Windows Server 2008 and Windows Server 2008 R2, it can only be enabled through the registry. During migration from Windows Server 2003 the registry settings will not be created automatically. For Windows Server 2008 and higher versions, if these registry settings happen to be present already, they will be migrated.

Migration components that are not automatically migrated

The following Remote Access elements and settings are not migrated by the Windows PowerShell cmdlets that are supplied with the Windows Server Migration Tools. Instead, you must manually configure the element or setting on the new RRAS server as described in **Completing the required manual migration steps** in this guide.

Important

Perform the manual configuration of these elements only when directed later in this guide.

- **SSL certificate bindings**. SSL Certificate binding and crypto-binding settings for SSTP are migrated as follows:
  a. The migration Wizard looks for a source certificate on the destination computer. If one is found, SSTP uses that certificate.
  b. If a source certificate is not found, the migration wizard will look for a valid certificate with the same trusted root as the source certificate.
  c. If still no certificate is found, then the SSTP configuration on the destination computer is 'Default'.
  d. If self-signed certificates are being used (valid for win8_Server_2), they will be automatically created on the destination computer.
- **User accounts on the local RRAS server**. If you use domain-based user and group accounts, and both the old and new RRAS servers are part of the same domain, no migration of the accounts is required. Local user accounts can be used if **Windows Authentication** is configured on the RRAS source server.
- **Only routing/VPN/DirectAccess when all are installed**. If your Remote Access server configuration includes all of the available services, then all services must be migrated together. Migrating only one of the services to the destination server is not supported.
- **A local or remote server that is running Network Policy Server (NPS)** that provides authentication, accounting, and policy management. This guide does not include the steps that are required to migrate a server that is running NPS. To migrate a server that is running NPS, use Migrate Network Policy Server to Windows Server 2012. NPS migration should be performed when directed later in this guide.

  Note

  If you are not using a server that is running NPS, the default Remote Access policies and accounting settings that are automatically created while configuring RRAS are not migrated.
- **Dial-up based demand-dial connections**. The destination server might have different modems, and there are many demand-dial settings that are specific to the modem or ISDN device that is selected.
- **Certificates** used for authenticating IKEv2, SSTP, and L2TP/IPsec connections.
- **SSL Certificate Binding** for SSTP when the **Use HTTP** check box is not selected.
- **IKEv2 VPN connections that use IPv6 network adapters**. IKEv2 is supported on RRAS servers that are running Windows Server 2008 R2 only. In the RRAS Microsoft Management Console (MMC) on Windows Server 2008 R2, you can specify the network interface used to acquire IPV6 DHCP and DNS addresses that are used for IKEv2 VPN clients. If you migrate RRAS from Windows Server 2008 R2 to another server running Windows Server 2008 R2, the setting is migrated. However, if you migrate from a previous version of Windows, there is no setting to migrate, and the default value of **Allow RAS to select the adapter** is used.

- **Weak encryption**. In Windows Server 2003, weak encryption is enabled, but on later versions of Windows it is disabled by default. You can enable weak encryption only by modifying the registry. During migration from Windows Server 2003 the required registry settings are not created on the new server by the migration process, and they must manually be configured. For later versions of Windows, if these registry settings are present, they are migrated.
- **Admin DLLs and Security DLLs** and their corresponding registry keys. These DLLs are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration.
- **Custom DLLs used for dialing a demand-dial connection**. These DLLs are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration. Any corresponding registry settings also are not migrated.
- **Connection Manager profiles**. The Connection Manager Administration Kit is used to create VPN and dial-up remote access profiles. Profiles created are stored under specific folders on the RRAS server. Profiles that are created on a 32-bit version of Windows do not work on computers that are running a 64-bit version of Windows, and vice versa. For more information about connection profiles, see Connection Manager Administration Kit (http://go.microsoft.com/fwlink/?linkid=55986).
- **The Group Forwarded Fragments setting on NAT**. This setting is enabled if the RRAS server is deployed behind a NAT router running on the Windows operating system. This is required for L2TP/IPsec connections that are using computer certificate authentication to succeed. We recommend that you enable this value to work around a known issue in RRAS.
- The **Log additional Routing and Remote Access information (used for debugging)** setting in the **Routing and Remote Access Properties** dialog box on the **Logging** tab.

Overview of the Routing and Remote Access service migration process

The pre-migration process involves the manual collection of data, followed by running procedures on the destination and source servers. The migration process includes source and destination server procedures that use the **Export** and **Import** cmdlets to automatically collect, store, and migrate server role settings. Post-migration procedures include verifying that the destination server successfully replaced the source server and then retiring or repurposing the source server. If the verification procedure indicates that the migration failed, troubleshooting begins. If troubleshooting fails, rollback instructions are provided to return the network to the use of the original source server.

10afc962-3596-46ec-b83c-7886d9af025b

Impact of migration

During migration, the Remote Access server is not available to accept incoming connections or to route traffic.

- New remote clients cannot connect to the server by using dial-up, VPN or DirectAccess connections. Existing connections on the server are disconnected. If you have multiple remote access servers, then the loss of availability of this server results in a reduction in capacity until the new server is operational again. For demand-dial connections, you must provide alternate connectivity between offices or reconfigure the connections to point to an alternate server.
- Routing and NAT functionalities are not available. If the functionality is required during the migration, an alternate router can be deployed until the new destination server is available.

Post-migration impacts include the following:

- If you plan to reuse the name of the source server as the name of the destination server, the name can be configured on the destination server after the source server is disconnected from the network. Otherwise, there is a name conflict that can affect availability. If you plan to run both servers, then the destination server must be given a unique name.
- A VPN server can be directly connected to the Internet, or it can be placed on a perimeter network that is behind a firewall or NAT router. If the IP address or DNS name of the destination server changes as part of the migration, or after the migration is completed, then the mappings in the firewall or NAT device must be reconfigured to point to the correct address or name. You must also update any intranet or Internet DNS servers with the new name and IP address. Also, remember to provide information about any server name or IP address changes to your users so that they can connect to the correct server. If you use connection profiles that are created by using the Connection Manager Administration Kit, then deploy a new profile with the updated server address information.

Note

For DirectAccess, the source computer and the destination computer must have the same IP addresses and interface names.

We recommend that you advertise the expected date and time of the migration so that users can plan accordingly, and make other arrangements as needed.

Permissions required to complete migration

The following permissions are required on the source server and the destination Remote Access servers:

- Domain user rights are required join the new server to the domain.
- Local administrative rights are required to install and manage the Remote Access role.
- Equivalent administrator permissions for DirectAccess GPOs as were configured on the source computer.
- Write permissions are required to the migration store location. For more information, see Remote Access: Prepare to Migrate in this guide.

Estimated duration

The migration can take two to three hours, including testing.

See Also

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Prepare to Migrate

Perform the following steps before you begin migrating Remote Access from your x86-based or x64-based source server to an x64-based destination server that is running Windows Server® 2012.

- Prepare your destination server
- Prepare your source server
- Install the migration tools

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, you might have to run the following steps by using the **Run as administrator** option.

Prepare your destination server

Complete the following procedures to prepare the destination server for the migration of Remote Access.

Hardware requirements for the destination server

Your destination server should have the same number or more network adapters as your source server. You can have more network adapters on the destination server than the source server, but the migration fails if you have fewer.

Important

The names of the network adapters on the destination server must be the same as those on the source server, and they must have the same intention (for example, Internet facing versus intranet facing). Most Remote Access server components have interface-specific settings and configuration. Having the same number of interfaces, with the same names and intent, helps ensure that the settings are migrated to the right interface. This is critical to a successful migration. If there are more adapters on the destination server than on the source server, you must still have a one-to-one match between the names and intention of the network adapters on the source server and those on the destination server.

Note

DirectAccess configuration can only be migrated from a computer running Windows Server 2012 to another computer running Windows 2012. Migration from Windows Server 2008 R2 DirectAccess to Windows Server 2012 DirectAccess is not supported.

Prepare the destination server for migration

To prepare the destination server

1. Install Windows Server 2012 on the destination server.
2. Whether or not you intend to migrate the source server name to the destination server, give the destination server a temporary computer name at this time.
3. If you store the user accounts for remote access users locally on the Remote Access server instead of in Active Directory, and if you use the Challenge Handshake Authentication Protocol (CHAP) for authentication, then you must perform the following additional steps before migrating Remote Access:
   a. To enable the use of CHAP authentication, you must manually configure a local security policy setting that enables passwords to be stored by using a reversible encryption algorithm.

      Security

      We recommend that you do not use CHAP for authentication, and that you do not enable the setting to store passwords with reversible encryption. These options are not considered secure, and they are provided only for backwards compatibility. Use them only if your environment requires the use of CHAP.

      i. On the destination server, in the **Start** screen, click **Administrative tools**, and then click **Local Security Policy**.
      ii. In the navigation tree, expand **Account Policies**, and then select **Password Policy**.
      iii. In the details pane, double-click **Store passwords using reversible encryption**, click **Enabled**, and then click **OK**.

   b. Migrate the local users and groups from the source server to the destination server. Do this separately and before you begin migrating Remote Access.
4. If the source server that is being replaced is joined to a domain, join the destination server to the same domain.
5. In the dashboard of the Server Manager console click **Add roles and Features**.
6. Click **Next** until you reach the **Select Server Roles** screen.
7. On the **Select Server Roles** screen, select **Remote Access**. Click **Add Required Features**, and then click **Next**.
8. On the **Select features** screen, click **Windows Server Migration Tools**. Click **Next** until you reach the **Select Role Services** screen.
9. In the **Select Role Services** screen, select **Routing** and then click **Next**.

10. On the **Confirm Installation selections** screen, click **Install**.
11. On the **Installation progress** screen, verify that the installation was successful, and then click **Close**.
12. Server roles were introduced in Windows Server 2008, and they are also used in Windows Server 2012. Remote Access is a role service that consists of the Routing service and the DirectAccess and VPN service. In Windows Server 2003, the Routing and VPN services were not separate. If the source server is running Windows Server 2003, ensure that both the Routing service and the DirectAccess and VPN service are installed on the destination server. If the source server is running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012, ensure that the destination server has the same Remote Access services installed as has the source server. If the source server has the routing service and the DirectAccess and VPN service, then you must install all these components on the destination server.
13. If DirectAccess is being migrated, the IP-HTTPS and Network Location certificate must be imported to the destination server. Note that if self-signed certificates were being used on the source server, this step is not required.
14. If DirectAccess is being migrated, run the Windows PowerShell cmdlet **install-remoteaccess –prerequisite** to ensure that the destination server meets all the requirements for DirectAccess.

The destination server is now prepared for migration.

Prepare your source server

Important
Before you begin migration, as a best practice, we recommend that you perform a backup of the source server. If the migration fails, and the recovery steps to restore the source server also fail, this backup can be critical for the quick restoration of service.

Back up your source server

- For information about backing up Windows Server 2003, see Backing up and restoring data in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=163718).
- For information about backing up Windows Server 2008 or Windows Server 2008 R2, see Backup and Recovery in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=163719).

Install the migration tools

Windows Server Migration Tools in Windows Server 2012 allows an administrator to migrate some server roles, features, operating system settings, shares, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2 to computers that are running Windows Server 2012.

Install Windows Server Migration Tools on the source and the destination servers. Complete installation, configuration, and removal instructions for Windows Server Migration Tools are available in Install, Use, and Remove Windows Server Migration Tools.

Important

Before you run the **Import-SmigServerSetting**, **Export-SmigServerSetting**, or **Get-SmigServerFeature** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with domain users or groups who are members of local groups on the source server.

Before you run the **Send-SmigServerData** or **Receive-SmigServerData** cmdlets, verify that during migration, both source and destination servers can contact the domain controller that is associated with those domain users who have rights to files or shares that are being migrated.

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Migrate Remote Access

Complete the following procedures to migrate the Routing and Remote Access service from a source server to a destination server.

- Migrating Remote Access from the source server
- Migrating Remote Access to the destination server
- Completing the required manual migration steps

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, you might have to run the following steps by using the **Run as administrator** option. For more information, see Run a program with administrative credentials in the Windows Server TechCenter (http://go.microsoft.com/fwlink/?LinkId=131210).

Migrating Remote Access from the source server

Follow these steps to capture the configuration of Remote Access on the source server.

To capture Remote Access configuration: Windows Server 2003, Windows Server 2008, Windows Server 2008 R2

1. On the source server, open a Windows PowerShell session with elevated user rights.
2. Load Windows Server Migration Tools into your Windows PowerShell session.

   Note
   If you opened the current Windows PowerShell session by using the Windows Server Migration Tools shortcut on the **Start** menu, skip this step and go to the next step. You should only load the Windows Server Migration Tools snap-in in a Windows PowerShell session that was opened by using another method, and into which the snap-in has not already been loaded.

   To load Windows Server Migration Tools, run the following cmdlet:

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```

3. Remote Access can be running on the source server while you are capturing its configuration. However, if you made configuration changes to Remote Access that require a service restart, then you must stop Remote Access before starting the migration. Use the following PowerShell command to stop the service:

   ```
   stop-service remoteaccess -force
   ```

   Note
   You must use the `-force` parameter because Remote Access has dependent services.

   To verify that the service is stopped, run the following command:

   ```
   get-service remoteaccess
   ```

4. On the source server, from Windows PowerShell, collect the settings from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The following is the syntax for the cmdlet:

   ```
   Export-SmigServerSetting -featureID NPAS-RRAS -User All -
   Group -path StorePath -verbose
   ```

   The **Export-SmigServerSetting** cmdlet parameters can collect all Routing and Remote Access service settings on the source server in a single file (Svrmig.mig). Before you run this command, review the following:

   - When you run the **Export-SmigServerSetting** command, you are prompted to provide a password to encrypt the migration store data. You must provide this same password when you later import from the migration store. Make sure you provide a strong password to encrypt the migration data and that the location of the migration data file is secure.
   - The *StorePath* variable that is provided as the value of the **path** parameter can be an empty or nonempty folder. The actual data file that is placed in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Do not specify a file name. If a migration data file already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must first move the Svrmig.mig file from that location and store it elsewhere, rename it, or delete it.
   - If the path is not a shared location that the destination server can access, you must

manually copy the migration store to the destination server or to a location that the destination server can access.

- Migrating users and groups can be combined with the cmdlets that are used to migrate Remote Access. The **-Users** and **-Group** parameters can be used in the **Export-SmigServerSetting** command to migrate the user and group accounts that are present locally on the Remote Access source server. If you are using an Active Directory domain or RADIUS for authentication, then these parameters are not needed.

  The **-Users** command supports the following parameters:

  - **All**. All user accounts on the source server are included in the migration output file.
  - **Enabled**. Only enabled user accounts on the source server are included in the migration output file.
  - **Disabled**. Only disabled user accounts on the source server are included in the migration output file.

  To prevent migrating any user accounts, do not include the **-Users** parameter in the command.

  The **-Group** command takes no additional parameters. If it is present, then all groups defined locally on the source server are included in the migration output file.

Note
The process described in the Local User and Group Migration Guide does not migrate some settings, including those under the **Dial-in** tab. We recommend that you thoroughly review the Local User and Group Migration Guide to understand which settings are migrated and which are not.

To capture the Routing and Remote Access service configuration: Windows Server 2012

1. From the Start screen, click **Windows Server Migration Tools**.
2.

Note
This step is only required if Routing/VPN is configured on the source computer.

Remote Access can be running on the source server while you are capturing its configuration. However, if you made configuration changes to Remote Access that require a service restart, then you must stop Remote Access before starting the migration. Use the following PowerShell command to stop the service:

```
stop-service remoteaccess -force
```

Note
You must use the `-force` parameter because Remote Access has dependent services.

To verify that the service is stopped, run the following command:

```
get-service remoteaccess
```

3. Before you start to capture the configuration, you must stop the Remote Access Management service. Use the following PowerShell command to stop the service:

   ```
   stop-service Ramgmtsvc
   ```

   To verify that the service is stopped, run the following command:

   ```
   get-service Ramgmtsvc
   ```

   Once the export and migration are complete, you can restart the Remote Access Management service:

   ```
   start-service Ramgmtsvc
   ```

4. On the source server, from Windows PowerShell, collect the settings from the source server by running the **Export-SmigServerSetting** cmdlet as an administrator. The following is the syntax for the cmdlet:

   ```
   Export-SmigServerSetting -featureID DirectAccess-VPN [-User
   All] -Group -path StorePath
   ```

   The **Export-SmigServerSetting** cmdlet parameters can collect all Remote Access settings on the source server in a single file (Svrmig.mig). Before you run this command, review the following:

   - When you run the **Export-SmigServerSetting** command, you are prompted to provide a password to encrypt the migration store data. You must provide this same password when you later import from the migration store. Make sure you provide a strong password to encrypt the migration data and that the location of the migration data file is secure.
   - The *StorePath* variable that is provided as the value of the **path** parameter can be an empty or nonempty folder. The actual data file that is placed in the folder (Svrmig.mig) is created by the **Export-SmigServerSetting** cmdlet. Do not specify a file name. If a migration data file already exists and you want to rerun the **Export-SmigServerSetting** cmdlet, you must first move the Svrmig.mig file from that location and store it elsewhere, rename it, or delete it.
   - If the path is not a shared location that the destination server can access, you must manually copy the migration store to the destination server or to a location that the destination server can access.
   - Migrating users and groups can be combined with the cmdlets that are used to migrate Remote Access. The **-Users** and **-Group** parameters can be used in the **Export-SmigServerSetting** command to migrate the user and group accounts that are present locally on the Remote Access source server. If you are using an Active Directory domain or RADIUS for authentication, then these parameters are not needed.

     The **-Users** command supports the following parameters:

     - **All**. All user accounts on the source server are included in the migration output

file.

- **Enabled**. Only enabled user accounts on the source server are included in the migration output file.
- **Disabled**. Only disabled user accounts on the source server are included in the migration output file.

To prevent migrating any user accounts, do not include the **-Users** parameter in the command.

The **-Group** command takes no additional parameters. If it is present, then all groups defined locally on the source server are included in the migration output file.

Note
The process described in the Local User and Group Migration Guide does not migrate some settings, including those under the **Dial-in** tab. We recommend that you thoroughly review the Local User and Group Migration Guide to understand which settings are migrated and which are not.

Migrating Remote Access to the destination server

Return to the destination server, and use the following procedure to complete the migration:

To import the Routing and Remote Access service configuration to the destination server

1. Before you use the **Import-SmigServerSetting** cmdlet to import the Routing and Remote Access service settings, be aware of the following condition:
   - If you chose to migrate the users and groups on the source computer, you need to specify the **-User** and **-Group** parameters in the **Import-SmigServerSetting** cmdlet on the destination server.
2. On the destination Server that is running Windows Server 2012, from the Start screen, click **Windows Server Migration Tools**.
3. On the destination server, from Windows PowerShell, run the following command, where *StorePath* is the folder that contains the Svrmig.mig file that you exported from the source server. Do not include the name of the file in the path.

   ```
   Import-SmigServerSetting -featureID DirectAccess-VPN [-User
   All] -Group -path StorePath –Force
   ```

   For more information about running the **Import-SmigServerSetting** cmdlet, see the “Using Windows Server Migration Tools” section in the Windows Server Migration Tools Install, Use, and Remove Windows Server Migration Tools guide.
4.

Note

This step is only required if Routing/VPN is configured on the source computer.

Before starting the Remote Access service, you must manually stop the RASMAN service. Run the following command in the Windows PowerShell Command Prompt window:

```
Stop-service -force rasman
```

5.

Note

This step is only required if Routing/VPN is configured on the source computer.

Then run the following command in the Windows PowerShell Command Prompt window to start the Routing and Remote Access service:

```
Start-Service RemoteAccess
```

If a failure occurs while running the **Import-SmigServerSetting** cmdlet, review the Setupact.log, Setuperr.log, and ServerMigration.log files under %localappdata%\SvrMig\Log. Information about how the Remote Access components migrated is included in the Servermigration.log file.

After the script completes, review the following section and adjust any remaining settings that require manual configuration.

Completing the required manual migration steps

Certain settings cannot be migrated by the Windows PowerShell scripts, and they must be configured manually on the destination server. Review the following configuration options, and apply those that are relevant to your environment.

DirectAccess

To ensure that the destination server meets all DirectAccess requirements, run the following Windows PowerShell cmdlet: **Install-remoteaccess –prerequisite**.

Dial-up demand-dial connections

Because of the differences in modem hardware that might exist between the source and destination servers, dial-up connections are not migrated. Use the Demand-Dial Interface Wizard in the Remote Access MMC snap-in.

To create a dial-up demand-dial connection

1. If you are using Server Manager, in Tools click Routing and Remote Access.
2. Right-click the server in the tree, and then click **Configure and Enable Routing and Remote Access**.
3. Follow the steps in the wizard to configure the connection.

Certificates for IKEv2, SSTP, and L2TP/IPsec connections

Certificates can be exported from the source server and imported to the destination server by using the Certificates MMC snap-in.

Routing and Remote Access service policies and accounting settings

If you are not using a local or remote server to run NPS, then default remote access policies and accounting settings are automatically created on the destination server when Remote Access is configured.

To migrate NPS settings, refer to Migrate Network Policy Server to Windows Server 2012.

PEAP, smart card, and other certificate settings on Network Policy Server

If you also migrated a local or remote server running NPS to support the Remote Access server that you are migrating, we recommend that you verify that the server that is running NPS has the correct certificate configuration. Specifically, confirm that any certificates that are associated with Protected Extensible Authentication Protocol (PEAP) and the **Smart card or other certificate** authentication settings are set properly. You can find these settings on the server that is running NPS, in the NPS MMC snap-in under **Connection Request Policies** or **Network Policies** (depending on where the authentication protocols are configured). If no certificates are present, or if the certificates are not configured correctly, perform the following steps:

To reconfigure PEAP or smart card certificates

1. Remove the **PEAP** or **SmartCard or other certificate** methods from the list of authentication methods.
2. Add the method back to the list.
3. Reconfigure the certificate for the specified method.

Weak encryption settings

In Windows Server 2003 weak encryption is enabled, but on later versions of Windows it is disabled by default. You can enable weak encryption only by modifying the registry. During migration from Windows Server 2003, the required registry settings are not created on the new server by the migration process, and they must manually be configured. For later versions of Windows, if these registry settings are present, they are migrated. For more information about the registry entries that Remote Access adds, see “Registry entries that Routing and Remote Access adds in Windows Server 2008”, article 947054 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?linkid=159112). The description of the settings for the weak encryption settings are at the end of the article, and they are named **AllowPPTPWeakCrypto** and **AllowL2TPWeakCrypto**.

Security

Weak encryption includes the use of 40-bit or 56-bit encryption in PPTP, and the use of MD5 or DES for L2TP/IPsec. By default, these weak algorithms are disabled, and we recommend that you do not use them unless they are required.

Connection Manager profiles

Profiles that are created by the Connection Manager Administration Kit (CMAK) can only be created on a computer with the same 32-bit or 64-bit architecture as the client computer on which they are to be run. If your source server is 64-bit, and you have created 64-bit profiles on that source server, you can copy them from the **%PROGRAMFILES%\CMAK\Profiles** folder to the appropriate folder on the destination server.

If the source server is 32-bit, you must use a computer running a 32-bit version of Windows to create and manage the profiles. You can set up a computer running a 32-bit version of Windows 7 or Windows 8, and then install CMAK on it to manage the profiles for your 32-bit client computers. For more information, see Connection Manager Administration Kit in the Windows Server Technical Library (http://go.microsoft.com/fwlink/?linkid=136440).

Group forwarded fragments

The **Group Forwarded Fragments** setting on NAT is enabled if the Remote Access server is deployed behind a NAT device that runs Windows. This is required for L2TP/IPsec connections that are using computer certificate authentication to succeed. We recommend that you enable this setting. Group Forwarded Fragments can be enabled for IPv4 on the Windows NAT computer by running the following command at the command prompt:

```
netsh int ipv4 set global groupforwardfragments=enabled
```

RAS administration and security DLLs

Administration DLLs and security DLLs and their corresponding registry keys are not migrated. This is because they are available in both 32-bit and 64-bit versions, and they do not work in a 32-bit to 64-bit migration. If the source and destination computers are 64-bit, the administration and security DLLs can be reused. For more information, refer to the following topics:

RAS Administration DLL (http://go.microsoft.com/fwlink/?linkid=163778)

RAS Security DLL (http://go.microsoft.com/fwlink/?linkid=163779)

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Verify the Migration

Remote Access: Post-migration Tasks

Remote Access: Verify the Migration

After all the migration steps are completed, you can use the following procedure to verify that the migration of Remote Access was successful. If the migration failed, you can return to the previous valid configuration by following the roll-back steps in Remote Access: Post-migration Tasks.

Verifying the destination server configuration

Membership in the local **Administrators** group or equivalent is the minimum required to complete these procedures. If User Account Control (UAC) is enabled, then you might have to run the following steps by using the **Run as administrator** option.

We recommend that you check the configuration of the destination Remote Access server, from the service start-up to the detailed configuration of individual components. The following sections provide a list of items to check. Depending on which Remote Access components are enabled on your server, only some of these checks might be necessary.

Installation state of Remote Access

The first verification step is to confirm that the Remote Access feature installed successfully.

To verify that Remote Access installed on the destination server

1. Click **Windows Server Migration Tools** on the **Start** screen.
2. View the installation status of the Routing and Remote Access service by running the following command:

   ```
   Get-WindowsFeature RemoteAccess
   ```

   The check box on the left of the **Remote Access** feature name is selected if the service is installed on the destination server. If it is not installed, the check box is clear.

Status of Remote Access Service

Verify that the Remote Access service is running.

To verify that the Routing and Remote Access service is running on the destination server

1. Click **Windows Server Migration Tools** on the **Start** screen.
2. View the service status of the Routing and Remote Access service by running the following command:

   ```
   Get-service RemoteAccess
   ```

3. Examine the **Status** column. It should read **Running**.

Remote access Operations Status

Verify the operations status of the deployment.

To verify the Remote Access operations status

1. In **Server Manager** click **Tools** and then click **Remote Access Management**.
2. Click **OPERATIONS STATUS** to navigate to **Operations Status** in the **Remote Access Management Console**. **Operations Status** lists the server operational status and that of all its components.

DirectAccess configuration

Verify the operations status of the deployment.

To verify the DirectAccess configuration settings

1. In **Server Manager** click **Tools** and then click **Remote Access Management**.
2. Click **CONFIGURATION** to navigate to Configuration tab in **Remote Access Management console**. Step through each of the wizards to ensure that the configuration has been migrated successfully.

VPN configuration

Confirm the configuration settings for the Remote Access server and ports.

To verify the Remote Access configuration settings

1. Start Server Manager.
2. In **Tools**, click **Routing and Remote Access**.
3. Right-click the Remote Access server node, and then click **Properties**.

   On each tab, confirm that the destination server is configured the same as the source server, and then click **OK**.
4. In the navigation pane, select **Ports**.

   Confirm that any modem or ISDN devices that are attached to the computer are included in the list.
5. In the navigation pane, right-click **Remote Access Logging and Policies**, and then click **Launch NPS**. In the Network Policy Server navigation pane, select **Network Policies**.

   Confirm that the NPS policies that are currently configured are those required for your environment. If you migrated them from an NPS source server to an NPS destination server, confirm that you are connected to the destination server and that the policies migrated successfully.

Dial-up configuration

You must confirm that the correct phone lines are attached to the modems or ISDN ports on the destination server.

Demand-dial VPN configuration

Examine all of your demand-dial VPN connections to ensure that they migrated with the correct settings.

To verify the settings for a demand-dial VPN connection

1. Start Server Manager.
2. Click **Routing and Remote Access**, and then select **Network Interfaces**.
3. In the details pane, right-click a demand-dial interface, and then click **Properties**.

   On each tab, confirm that the connection is configured the same as the source server, and then click **OK**.

Router settings

Confirm that the router components installed, and verify that each is configured correctly. The available routing components include:

- **IPv4:** Static Routes, DHCP Relay Agent, IGMP, NAT, and RIPv2
- **IPv6:** Static Routes and DHCPv6 Relay Agent

To verify the routing components

1. Start Server Manager.
2. Click **Routing and Remote Access**.
3. Expand **IPv4**. Examine the list of installed routing components, and ensure that the components required for your deployment are installed.
4. Expand **IPv6** and follow the same process as the previous step.
5. In the navigation pane, under **IPv4**, click **General**.

   The details pane identifies the interfaces that are configured to route packets for each version of IP. Confirm that the list contains the expected interfaces, including configured demand-dial interfaces.
6. In the navigation pane, under **IPv6**, click **General** and follow the same process as the previous step.
7. In the details pane for **General** under **IPv4**, right-click each interface and select **Properties**.

   On each tab confirm that the interface is configured as required for its routing role on the server.
8. Follow the same process as described in the previous step for the interfaces listed on the

under **IPv6** / **General**.

9. Under **IPv4** select **Static Routes** and confirm that the routes to destination networks are correctly configured with the associated interface and destination gateway address.
10. Follow the same process as described in the previous step for the **Static Routes** under **IPv6**.
11. Under **IPv4**, select **NAT**. The details pane shows the interfaces that NAT is configured to use. Right-click each interface and click **Properties**.
    - Confirm that each interface is configured correctly for NAT. There should be at least two interfaces enabled for NAT, one configured as the **Private** interface, and one configured as the **Public** interface.
    - If NAT is responsible for providing IPv4 addresses to clients on the private network, then on **NAT Properties** page, on the **Address Assignment** tab, select the **Automatically assign IP addresses by using the DHCP allocator** check box and enter the address information to be used.
    - If your ISP has provided a pool of addresses to be used by the NAT public interface, ensure that they are configured correctly. The addresses are under **NAT**, on the **Properties** page for the interface, on the **Address Pool** tab. If the addresses that were migrated are not applicable to the target computer, modify the list to use the correct addresses.
    - For each interface under **NAT**, on the interface's **Properties** page on the **Services and Ports** tab, examine the port mappings for services that must be routed to a specific server IP address. Confirm that each service that is to be mapped has the correct address pool entry, private IP address, and port settings configured.
12. Under **IPv4**, select each enabled routing protocol. The details pane shows the interfaces on which the selected routing protocol is enabled. Right-click each interface, and then click **Properties**.

    Confirm that each interface is configured correctly for the selected routing protocol. For example, under **IPv4**/**NAT**, there should be at least two interfaces, one configured as the **Private** interface, and one configured as the **Public** interface.
13. Under **IPv6**, select each enabled routing protocol and follow the same process described in the previous step.
14. Under **IPv4**, right-click each routing protocol, and then select **Properties** to examine the global configuration for that routing protocol.

    Confirm that each protocol is configured correctly for your environment. For example, ensure that the **DHCP Relay Agent** has a list of DHCP server addresses to which it can forward DHCP requests from clients.
15. Under **IPv6**, select each enabled routing protocol and follow the same process described in the previous step.

User and Group accounts

If you migrated the user and group accounts by using the Local User and Group Migration Guide (http://go.microsoft.com/fwlink/?linkid=163774), follow the procedures in its verification section to confirm that the required users and group were migrated successfully.

If you instead used the **-user** and **-group** parameters on the **Import-SmigServerSetting** command, you can manually verify the accounts by using the **Local Users and Groups** MMC snap-in to examine the user and group accounts and confirm that the properties for the accounts are set properly.

Final checks

- If your computer is configured to host VPN/DirectAccess connections, test each type of supported connection to confirm that users can connect.
- If your server is configured to host dial-up connections, verify that client computers can successfully dial-in and connect to the server by using the modems that are installed.
- If your server is configured as an IPv4 or IPv6 router, verify that clients on each attached network can connect through the router to computers on all of the other attached networks. If you use the **ping** command for this test, ensure that Windows Firewall on the router and the client computers is configured to allow ICMP Echo Request and ICMP Echo Reply messages.

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Post-migration Tasks

Remote Access: Post-migration Tasks

Perform the following post-migration tasks to complete your migration:

- Completing the migration
- Configuring firewall rules for VPN
- Configuring firewall rules for DirectAccess
- Restoring Remote Access in the event of migration failure
- Retiring Remote Access on your source server
- Troubleshooting cmdlet-based migration

Note

The post-migration tasks for the source server are optional, depending on your migration scenario.

Completing the migration

Migration is complete as soon as verification efforts demonstrate that the destination server has replaced the source server in serving the network.

If your verification efforts indicate that the migration failed, follow the steps in Restoring Remote Access in the event of migration failure to return to the previous valid configuration. The following list identifies some settings that must be manually configured after the migration is complete.

- **Server name**: If the source server is to be decommissioned, the same name is available to be configured on the destination server. If the IP addresses of the intranet and internet interfaces have changed, it is important to ensure that the intranet and internet DNS servers are updated with the new name and IP addresses. These configuration details also need to be deployed to your users by updating their connection profiles so that they can connect to the correct server.
- **Perimeter network firewall or NAT**: If the destination server is in a perimeter network, the firewall or NAT through which the server is accessed from the Internet must be configured with the new IP address of the RRAS server. Refer to the documentation for your Firewall or NAT router for the relevant configuration instructions.

Configuring firewall rules for VPN

Firewall rules that permit VPN network traffic are included with the Windows Firewall with Advanced Security, but they are disabled by default. The rules that enable the required inbound network traffic must be enabled on the destination server. If there were any other rules explicitly configured on the source server to support RRAS or its roles, they should be configured on the destination server also. If you use non-Microsoft firewall software, refer to the documentation that is provided by the vendor for instructions about how to configure the appropriate rules.

When you migrate RRAS from the source server to the destination server, the firewall rules are not automatically enabled. You must use Windows Firewall with Advanced Security or a Group Policy setting to enable the rules that correspond to the types of RRAS network traffic that must enter your server. The following is a list of rules that you should enable, depending on which RRAS protocols you use:

- Routing and Remote Access (GRE-In)
- Routing and Remote Access (L2TP-In)
- Routing and Remote Access (PPTP-In)
- Secure Socket Tunneling Protocol (SSTP-In)

If you change the default firewall behavior to blocking all traffic that does not match an allow rule then you must enable the outbound allow rules in addition to the inbound rules.

Configuring firewall rules for DirectAccess

When using additional firewalls in your deployment, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- Teredo traffic—User Datagram Protocol (UDP) destination port 3544 inbound, and UDP source port 3544 outbound.
- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required.

  Note
  This exemption must be configured on the Remote Access server, while all the other exemptions have to be configured on the edge firewall.

Note
For Teredo and 6to4 traffic, these exceptions should be applied for both of the Internet-facing consecutive public IPv4 addresses on the Remote Access server. For IP-HTTPS the exceptions need only be applied to the address where the public name of the server resolves.

When using additional firewalls, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.
- Internet Control Message Protocol for IPv6 (ICMPv6) traffic inbound and outbound – for Teredo implementations only.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic
- ICMP for all IPv4/IPv6 traffic

Restoring Remote Access in the event of migration failure

If migration of Remote Access to the destination server fails, you can put the source server back into operation by following these steps:

- If the source server has not been repurposed and the computer name and IP address have not been migrated from the source to the destination server, simply connect the source server to the network and start the Routing and Remote Access service to allow users to connect again.
- If the computer name and IP address have been migrated from the source server to the destination server, rename the destination server to a temporary name and change its IP address to a different IP address. Set the source server computer name and IP address to

the values that were used before the migration, and restart the Routing and Remote Access service on the source server.

- If the previous steps are not valid options, such as when the source server has been repurposed or is otherwise unavailable, use the backup files that were created from the source server, as described in Remote Access: Prepare to Migrate. You can use the backup files any time after migration to restore the original Routing and Remote Access source server if a catastrophic failure occurs.

Estimated time to complete a rollback

You should be able to complete a rollback in one to two hours.

Retiring Remote Access on your source server

After you verify that the migration is complete, the source server can be disconnected from the network and removed from service. Stop RRAS before you remove the computer from the network and turn it off. You can keep this computer as a backup server in the event that you want to revert to your previous Routing and Remote Access configuration.

Troubleshooting cmdlet-based migration

The Windows Server Migration Tools deployment log file is located at %*windir*%\Logs\SmigDeploy.log. Additional Windows Server Migration Tools log files are created at the following locations.

- %*windir*%\Logs\ServerMigration.log
- On Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012: %*localappdata*%\SvrMig\Log
- On Windows Server 2003: %*userprofile*%\Local Settings\Application Data\SvrMig\Log

If migration log files cannot be created in the preceding locations, **ServerMigration.log** and **SmigDeploy.log** are created in %*temp*%, and other logs are created in %*windir*%\System32.

If you are migrating Remote Access and a failure occurs while running the **Import-SmigServerSetting** cmdlet, review the Setupact.log, Setuperr.log, and ServerMigration.log files under %localappdata%\SvrMig\Log. Information about how the Remote Access components migrated is included in the Servermigration.log file.

If a migration cmdlet fails, and the Windows PowerShell session closes unexpectedly with an access violation error message, look for a message similar to the following example in the %*localappdata*%\SvrMig\Logs\setuperr.log file.

```
FatalError [0x090001] PANTHR Exception (code 0xC0000005: ACCESS_VIOLATION) occurred at
0x000007FEEDE9E050 in C:\Windows\system32\migwiz\unbcl.dll (+000000000008E050).  Minidump
attached (317793 bytes).
```

This failure occurs when the server cannot contact domain controllers that are associated with domain users or groups who are members of local groups, or who have rights to files or shares

that are being migrated. When this happens, each domain user or group is displayed in the GUI as an unresolved security identifier (SID). An example of a SID is **S-1-5-21-1579938362-1064596589-3161144252-1006**.

To prevent this problem, verify that required domain controllers or global catalog servers are running, and that network connectivity allows communication between both source and destination servers and required domain controllers or global catalog servers. Then, run the cmdlets again.

If connections between either the source or destination servers and the domain controllers or global catalog servers cannot be restored, do the following.

1. Before you run **Export-SmigServerSetting**, **Import-SmigServerSetting** or **Get-SmigServerFeature** again, remove all unresolved domain users or groups who are members of local groups from the server on which you are running the cmdlet.
2. Before you run **Send-SmigServerData** or **Receive-SmigServerData** again, remove all unresolved domain users or groups who have user rights to files, folders, or shares on the migration source server.

Viewing the content of Windows Server Migration Tools result objects

All Windows Server Migration Tools cmdlets return results as objects. You can save result objects, and query them for more information about settings and data that were migrated. You can also use result objects as input for other Windows PowerShell commands and scripts.

Result object descriptions

The Windows Server Migration Tools **Import-SmigServerSetting** and **Export-SmigServerSetting** cmdlets return results in a list of **MigrationResult** objects. Each **MigrationResult** object contains information about the data or setting that the cmdlet processes, the result of the operation, and any related error or warning messages. The following table describes the properties of a **MigrationResult** object.

Property name	Type	Definition
ItemType	Enum	The type of item being migrated. Values include **General**, **WindowsFeatureInstallation**, **WindowsFeature**, and **OSSetting**.
ID	String	The ID of the migrated item. Examples of values include **Local User**, **Local Group**, and **DHCP**.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise,

Property name	Type	Definition
		False is displayed.
DetailsList	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.

Send-SmigServerData and **Receive-SmigServerData** cmdlets return results in a list of **MigrationDataResult** objects. Each **MigrationDataResult** object contains information about the data or share that the cmdlet processes, the result of the operation, any error or warning messages, and other related information. The following table describes the properties of a **MigrationDataResult** object.

Property name	Type	Definition
ItemType	Enum	The type of migrated item. Values include **File**, **Folder**, **Share**, and **Encrypted File**.
SourceLocation	String	The source location of the item, shown as a path name.
DestinationLocation	String	The destination location of the item, shown as a path name.
Success	Boolean	The value **True** is displayed if migration was successful; otherwise, **False** is displayed.
Size	Integer	The item size, in bytes.
ErrorDetails	List <MigrationResultDetails>	A list of **MigrationResultDetails** objects.
Error	Enum	Errors enumeration for errors that occurred.
WarningMessageList	List <String>	A list of warning messages.

The following table describes the properties of objects within the **MigrationResultDetails** object that are common to both **MigrationResult** and **MigrationDataResult** objects.

Property name	Type	Definition
FeatureId	String	The name of the migration setting that is related to the item. Examples of values

Property name	Type	Definition
		include **IPConfig** and **DNS**. This property is empty for data migration.
Messages	List <String>	A list of detailed event messages.
DetailCode	Integer	The error or warning code associated with each event message.
Severity	Enum	The severity of an event, if events occurred. Examples of values include **Information**, **Error**, and **Warning**.
Title	String	Title of the result object. Examples of values include NIC physical address for IP configuration, or user name for local user migration.

Examples

The following examples show how to store the list of the result objects in a variable, and then use the variable in a query to return the content of result objects after migration is complete.

To store a list of result objects as a variable for queries

1. To run a cmdlet and save the result in variable, type a command in the following format, and then press **Enter**.

   $*VariableName* **=** $(*Cmdlet*)

   The following is an example.

   $ImportResult = $(Import-SmigServerSetting -FeatureId DHCP -User all -Group -Path D:\rmt\DemoStore -force -Verbose)

   This command runs the **Import-SmigServerSetting** cmdlet with several parameters specified, and then saves result objects in the variable **ImportResult**.
2. After the **Import-SmigServerSetting** cmdlet has completed its operations, return the information contained in the result object by typing a command in the following format, and then pressing **Enter**.

   $*VariableName*

   In the following example, the variable is named **ImportResult**.

$*ImportResult*

This command returns information contained in the result objects that were returned by **Import-SmigServerSetting** in the example shown in step 1. The following is an example of the output that is displayed by calling the **ImportResult** variable.

```
            ItemType  ID                                Success
DetailsList
            --------  --                                -------
-----------
          OSSetting  Local User                          True
{Local User, Loc...
          OSSetting  Local Group                         True
{Local Group, Lo...
     WindowsFeature  DHCP                                True
{}
```

Each line of the preceding sample is a migration result for an item that was migrated by using the **Import-SmigServerSetting** cmdlet. The column heading names are properties of **MigrationResult** objects. You can incorporate these properties into another command to return greater detail about result objects, as shown by examples in step 3 and forward.

3. To display a specific property for all result objects in the list, type a command in the following format, and then press **Enter**.

   $<VariableName>| Select-Object -ExpandProperty <PropertyName>

   The following is an example.

   $importResult | Select-Object -ExpandProperty DetailsList

4. You can run more advanced queries to analyze result objects by using Windows PowerShell cmdlets. The following are examples.

   - The following command returns only those details of result objects that have the ID **Local User**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList

   - The following command returns only those details of result objects with an ID of **Local User** that have a message severity equal to **Warning**.

     $ImportResult | Where-Object { $_.ID -eq "Local User" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Severity -eq "Warning") {$_} }

   - The following command returns only the details of result objects with an ID of **Local User** that also have the title **Remote Desktop Users**.

     $ImportResult | Where-Object { $_.ID -eq "Local Group" } | Select-Object -ExpandProperty DetailsList | ForEach-Object { if ($_.Title -eq "Remote

```
DesktopUsers") {$_} }
```

More information about querying results

For more information about the cmdlets that are used in the preceding examples, see the following additional resources.

- Where-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134853).
- Select-Object on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134858).
- ForEach-Object on the Microsoft Script Center Web site (http://www.microsoft.com/technet/scriptcenter/topics/msh/cmdlets/foreach-object.mspx)

For more information about Windows PowerShell scripting techniques, see What Can I Do With Windows PowerShell? - Scripting Techniques on the Microsoft Script Center Web site (http://go.microsoft.com/fwlink/?LinkId=134862).

See Also

Migrate Remote Access to Windows Server 2012

Remote Access: Prepare to Migrate

Remote Access: Migrate Remote Access

Remote Access: Verify the Migration

Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012

This document describes the migration of an existing Forefront UAG SP1 DirectAccess deployment to DirectAccess in Windows Server® 2012. It illustrates migration of a simple scenario that includes either a single Forefront UAG server, or an array of Forefront UAG servers configured in a single domain and a single site using NAT64, and not set up as an ISATAP router. Note that this upgrade is only supported for computers running Forefront UAG SP1.

Windows Server 2012 Remote Access (DirectAccess) Deployment Documentation Set

Following is a listing of the documentation for the three main Remote Access deployment paths: **Basic**, **Advanced** and **Enterprise**. Also listed are the **Managing** and **Migrating** documents available for this release.

Deploy Basic Remote Access

- **Deploy a Single Remote Access Server using the Getting Started Wizard**
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Deploy Advanced Remote Access

- **Deploy a Single Remote Access Server with Advanced Settings**
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Deploy Remote Access in an Enterprise

- DirectAccess Capacity Planning
- **Deploy Remote Access in a Cluster**
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- **Deploy Multiple Remote Access Servers in a Multisite Deployment**
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- **Deploy Remote Access with OTP Authentication**
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- **Deploy Remote Access in a Multi-Forest Environment**
- DirectAccess Offline Domain Join

Manage Remote Access

- **Use Remote Access Monitoring and Accounting**
- **Manage DirectAccess Clients Remotely**

Migrate Remote Access

- **Migrate Remote Access to Windows Server 2012**
- **Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012**

Before you begin deploying, see the list of unsupported configurations, known issues, and prerequisites

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Scenario description

Supported migration scenarios for Forefront UAG SP1 are as follows:

Supported client operating systems	Supported domain controllers	Supported application servers
Windows® 7Windows 7 Windows® 8 Windows Server 2012	Windows 2003 Server Windows Server® 2008 Windows Server® 2008 R2 Windows Server 2012	Windows 2003 Server Windows Server 2008 Windows Server 2008 R2 Windows Server 2012

In this scenario

Two migration scenarios are described:

- **Side-by-side migration**— Use this type of migration to keep the Forefront UAG DirectAccess server running while you deploy DirectAccess in Windows Server 2012. After the deployment is complete, DirectAccess clients use DirectAccess configured on the Windows Server 2012 computer, and the Forefront UAG server is removed from service. This type of migration requires duplication of some settings because FQDN, IP addresses, and certificate settings must be unique on each server.
- **Offline migration**— Use this type of migration to copy the DirectAccess configuration with identical settings from the Forefront UAG DirectAccess server to the Windows Server 2012 computer running as a Remote Access server. Then shut down the Forefront UAG server. Service for DirectAccess clients will not be available until the Windows Server 2012 Remote Access server is up and running.

Prerequisites

Before you begin deploying this scenario, review this list for important requirements:

- ISATAP in the corporate network is not supported. If you are using ISATAP, you should remove it and use native IPv6.
- If NAP is used inbox in UAG, NAP is now going to require a separate NPS server.

  NAP was deprecated in Windows Server 2012 R2. This means that NAP may not be supported in future versions of Windows. New deployments with NAP are not recommended.

Practical applications

This scenario describes how to continue running an existing DirectAccess deployment using Windows Server 2012 instead of Forefront UAG.

Hardware requirements

Hardware requirements include the following:

- One or more Forefront UAG servers successfully running a DirectAccess deployment.
- Windows Server 2012 Remote Access server requirements:
  - A computer that meets the hardware requirements for Windows Server 2012.
- Client requirements for DirectAccess in Windows Server 2012:
  - A client computer must be running Windows® 8 or Windows 7.
- Infrastructure and management server requirements:
  - During remote management of DirectAccess client computers, clients initiate communications with management servers such as domain controllers, System Center Configuration Servers, and Health Registration Authority (HRA) servers for services that include Windows and antivirus updates and Network Access Protection (NAP) client compliance. The required servers should be deployed before beginning the Remote Access deployment.
  - If remote access requires client NAP compliance, NPS and HRS servers should be deployed before beginning remote access deployment
  - A certification authority (CA) server is required if certificates will be issued for authentication of IP-HTTPS and the network location server. Note that DirectAccess in Windows Server 2012 supports the use of self-signed certificates created automatically during DirectAccess deployment.
  - A DNS server running Windows Server 2003, Windows Server 2008 SP2; Windows Server 2008 R2; or Windows Server 2012 is required.

Software requirements

There are a number of requirements for this scenario:

- DirectAccess server requirements in Windows Server 2012:
  - The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
  - The person deploying remote access on the server requires local administrator permissions on the server, and have a domain user account. To prepare the GPOs domain administrator permissions are required.
- Remote access client requirements:
  - DirectAccess clients must be domain members. Domains containing clients can belong to the same forest as the Remote Access server, or have a two-way trust with the Remote Access server forest or domain.
  - An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients.

Using ISATAP

ISATAP is not recommended for use as the IPv6 to IPv4 transition technology in DirectAccess in Windows Server 2012. If Forefront UAG is configured to use ISATAP, it is recommended to disable it, and use NAT64 instead.

With ISATAP disabled DirectAccess clients can initiate connections to computers on the internal network, and the computers on the internal network are able to respond. However, computers on the internal network will not be able to initiate connections to DirectAccess for purposes of remote client management. If you want to be able to remote client management, consider deploying native IPv6 for management servers that will connect to DirectAccess client computers.

Using NAP

Forefront UAG provides complex Network Access Policy (NAP) configuration settings, and the Network Policy Server (NPS) and Health Registration Authority (HRA) roles can be installed on the Forefront UAG server. These settings are not supported for DirectAccess in Windows Server 2012. In Windows Server 2012, you can only specify whether or not client compliance is enforced using NAP during IPsec authentication. NPS and HRA roles are installed on remote servers in the internal network. The HRA server must be accessible over the first DirectAccess tunnel, or over the Internet.

Side-by-Side Migration of Forefront UAG DirectAccess

This topic describes how to perform a side-by-side migration of Forefront UAG DirectAccess to DirectAccess in Windows Server® 2012. A side-by-side migration keeps the Forefront UAG DirectAccess server running while you deploy Windows Server 2012 DirectAccess. After the deployment is complete, DirectAccess clients access internal resources via the Windows Server 2012 Remote Access server, and the Forefront UAG server is removed from service.

The side-by-side migration consists of the following steps:

- Step 1: Forefront UAG DirectAccess configuration settings are exported. If Forefront UAG is deployed in an array, this should be performed on the array manager server.
- Step 2: Note the names of Group Policy Objects (GPOs) used for Forefront UAG DirectAccess, including the server GPO, client GPOs, and application server GPOs if relevant.
- Step 3: Set up the Windows Server 2012 computer as a Remote Access server, by installing the Remote Access role.<*Step name*>
- Step 4: Configure infrastructure settings for the Remote Access server, including IP addresses, DNS settings, certificates, client security groups, and GPOs.

For a complete walkthrough, see Side-by-Side Migration Steps.

The following graphic illustrates the side-by-side migration process.

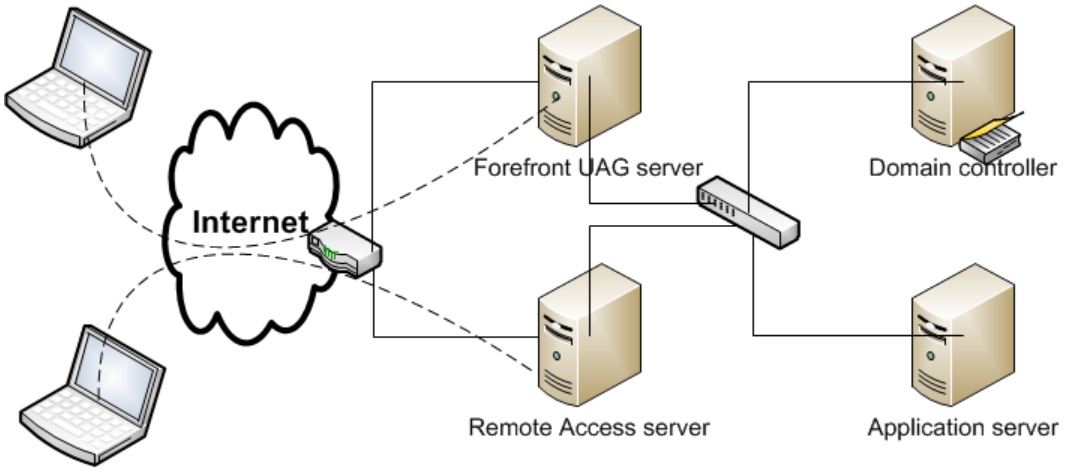

Side-by-Side Migration Steps

This topic describes the steps required to perform a side-by-side migration from Forefront UAG SP1 DirectAccess to DirectAccess in Windows Server® 2012.

Task	Description
Step 1: Configure ISATAP	ISATAP is not recommended for use as the IPv6 to IPv4 transition technology in DirectAccess in Windows Server 2012. If ISATAP is enabled in the Forefront UAG DirectAccess deployment, disable it before exporting the configuration settings. If manage out capability is also required for Windows Server 2012, you can configure the Forefront UAG and the Windows Server 2012 servers as ISATAP routers.
Step 2: Export the Forefront UAG DirectAccess settings	Export DirectAccess settings using the Forefront UAG export feature.
Step 3: Collect DirectAccess GPO settings	Note the names of GPOs used by Forefront UAG DirectAccess.
Step 4: Install the Remote Access role	Set up the computer running Windows Server 2012 as a Remote Access server.
Step 5: Configure server and infrastructure settings	Configure infrastructure and server settings before configuring DirectAccess on the Remote Access server.

Task	Description
Step 6: Obtain an IP-HTTPS certificate	DirectAccess in Windows Server 2012 provides two options for the IP-HTTPS certificate. You can obtain a certificate from a CA, in a similar way to Forefront UAG DirectAccess, or you can configure Windows Server 2012 DirectAccess to automatically issue a self-signed certificate for IP-HTTPS authentication.
Step 7: Create a security group for DirectAccess clients	Create an Active Directory security group that contains managed computers that will be configured as DirectAccess clients.
Step 8: Prepare GPOs	Prepare GPOs in all relevant domains.
Step 9: Configure DirectAccess	Configure DirectAccess settings on the Remote Access server.

Step 1: Configure ISATAP

It is recommended that you do not use ISATAP in a Windows Server 2012 DirectAccess deployment. You should disable ISATAP and use NAT64 instead.

To disable ISATAP

1. Remove the ISATAP entry from DNS.
2. Wait for replication to complete on all domain controllers.
3. Wait an additional 24 hours for all computers to lose their ISATAP addresses.

Note that before configuring Remote Access on the Windows Server 2012 computer, ensure that the computer does not have any ISATAP addresses.

If manage out capability is required only for Forefront UAG run the following command before you configure the Remote Access role on the Windows Server 2012 DirectAccess server:

```
Set-NetIsatapConfiguration -State Disabled
```

If manage out capability is also required for Windows Server 2012, we recommend that you enable native IPv6 in your network, at least for those hosts that need to manage DirectAccess clients. If deploying native IPv6 is not an option and managing DirectAccess clients from within the internal network is required, perform the following procedure:

This procedure configures both the Forefront UAG and Windows Server 2012 servers as ISATAP routers, each handling only traffic to its own IP-HTTPS and Teredo clients.

Note
To allow DNS to propagate, this procedure includes a waiting period of 24 hours.

Note

6to4 clients will no longer be able to connect to Forefront UAG.

To configure Forefront UAG and Windows Server 2012 as ISATAP routers

1. On the Windows Server 2012 server, obtain, and then write down the Forefront UAG ISATAP prefix by running the following command:

```
Get-NetRoute | ?{ $_.DestinationPrefix -match ":8000::/64" }
| Select-Object -ExpandProperty DestinationPrefix
```

2. Prior to configuring the Remote Access role on the Windows Server 2012 DirectAccess server, run the following command to prevent the Windows Server 2012 DirectAccess server from becoming an ISATAP host (and later enable it to become an ISATAP router)

```
Set-NetIsatapConfiguration -State Disabled
```

3. Install Windows Server 2012 and add the Remote Access role. See Step 4: Install the Remote Access role.
4. Configure DirectAccess according to the instructions in this guide.
5. Enable multisite on the Windows Server 2012 DirectAccess server to enable symmetric routing and the removal, at a later stage, of the Forefront UAG server in an ISATAP scenario by following the instructions at 3.6. Enable the multisite deployment.

   Note

   It is necessary to enable multisite for the deployment to enable the changing of the ISATAP prefix, which is not possible in a single site deployment.

6. Change the ISATAP prefix of the Windows Server 2012 DirectAccess server to be the same as the Forefront UAG ISATAP prefix by typing the following command

```
Set-DAServer -InternalIPv6Prefix <UAG ISATAP Prefix>
```

   Where *UAG ISATAP Prefix* is the prefix obtained in step 1.

   Although this cmdlet changes the DNS64 address on the Windows Server 2012 DirectAccess server to use the Forefront UAG prefix, it doesn't update the address on the firewall rule to allow DNS queries to reach DNS64.

   To complete this step run the following commands on the Windows Server 2012 DirectAccess server:

```
$dns64Address = (Get-DAClientDnsConfiguration).NrptEntry | ?{
$_.DirectAccessDnsServers -match ':3333::1' } | Select-Object
-First 1 -ExpandProperty DirectAccessDnsServers


$serverGpoName = (Get-RemoteAccess).ServerGpoName


$serverGpoDc = (Get-DAEntryPointDC).DomainControllerName
```

```
$gpoSession = Open-NetGPO -PolicyStore $serverGpoName -
DomainController $serverGpoDc

Get-NetFirewallRule -GPOSession $gpoSession | ? {$_.Name -in
@('0FDEEC95-1EA6-4042-8BA6-6EF5336DE91A', '24FD98AA-178E-
4B01-9220-D0DADA9C8503')} | Set-NetFirewallRule -LocalAddress
$dns64Address

Save-NetGPO -GPOSession $gpoSession

gpupdate /force
```

7. On the Windows Server 2012 DirectAccess server run the following commands to obtain, and then write down, the IPv6 prefix of the Forefront UAG server's Teredo clients:

```
$UAGIPv4Address="<UAG's first external IPv4 address>"

[Byte[]] $TeredoServerAddressBytes = `
    [System.Net.IPAddress]::Parse("2001::").GetAddressBytes()
[0..3] + `
    [System.Net.IPAddress]::Parse($UAGIPv4Address).GetAddress
Bytes() + `
    [System.Net.IPAddress]::Parse("::").GetAddressBytes()[0..
7]

Write-Host "The UAG's Teredo prefix is
$([System.Net.IPAddress]$TeredoServerAddressBytes)/64"
```

8. On the Forefront UAG server, run the following command to obtain the Teredo interface's index:

```
Netsh int ipv6 show route
```

   Find the interface with a prefix of 2001::/32 and write down its interface index from the Idx column.

9. On the Forefront UAG server run the following command to enable Forefront UAG to publish its Teredo client prefix to the ISATAP subnet:

```
Netsh int ipv6 add route prefix=<teredo
prefix>  interface=<teredo interface index> publish=yes
```

   Where *teredo prefix* was obtained in step 7 and *teredo interface index* was obtained in step 8.

10. Since 6to4 clients will no longer be able to successfully connect to via Forefront UAG to

the backend servers, you should disable 6to4 on all clients using a GPO:

Create a GPO to set the *Computer Configuration/Policies/Administrative Templates/Network/TCP/IP Settings/IPv6 Transition Technologies/6to4 State* setting to Disabled State and apply it to all Forefront UAG client machines. Alternatively you can accomplish this by blocking IP Protocol 41 inbound through your edge firewalls.

In addition, if the IPv4 address of the Windows Server 2012 DirectAccess server's internal network adapter is public, you must disable forwarding on the 6to4 interface of the Windows Server 2012 DirectAccess server so that ISATAP forwarding will work properly.

Note

Running the following command causes an error to be constantly displayed in the Operations Status (health monitor) for 6to4. You can ignore this error.

```
Set-NetIPInterface -InterfaceAlias '6TO4 Adapter' -Forwarding
Disabled
```

11. Create an ISATAP DNS record for the Windows Server 2012 DirectAccess server to publish it as an ISATAP router alongside Forefront UAG.

    On the DNS server, run dnscmd /recordadd command. For example:

    ```
    dnscmd dnssvr1.contoso.com /recordadd contoso.com isatap A
    10.0.0.5
    ```

    Where *dnssvr1.contoso.com* is the name of your DNS server, *contoso.com* is the DNS zone (domain name) and *10.0.0.5* is the internal IPv4 address of the Windows Server 2012 DirectAccess server.

    Note

    You must perform this command for each domain in which ISATAP records already exist.

12. At this stage client connectivity through the Windows Server 2012 DirectAccess server is working over NAT64, but manage out still does not work.
13. Wait 24 hours for the DNS to propagate before adding any client machines to the Windows Server 2012 DirectAccess server.
14. On the Windows Server 2012 DirectAccess server enable AAAA name resolution by running the following command:

    ```
    Set-NetDnsTransitionConfiguration -OnlySendAQuery $false
    ```

15. At this stage manage out functionality also works.
16. When all clients are migrated from Forefront UAG to the Windows Server 2012 DirectAccess server, you can remove the Forefront UAG ISATAP DNS records from the DNS server.

Step 2: Export the Forefront UAG DirectAccess settings

Export the settings using the following procedure.

To export settings

1. Click **Start**, click **All Programs**, click **Microsoft Forefront UAG**, and then click **Forefront UAG Management** to open the console.
2. Click **DirectAccess** in the left-hand node of the console.
3. Click **Export Policy** in the details pane. If there have been changes in corporate infrastructure servers since the last time Forefront UAG was configured, you might be prompted to refresh the management server list. To ensure settings are up to date, do this before exporting your settings.
4. Click **Print Preview**.
5. In the **Select Printer** list, select **Microsoft XPS Document Printer**, and click **Print**.
6. Specify a path and file name to save the configuration, and click **Save**.
7. Click **Cancel** to close the review page.
8. Copy the exported file to a location from which it can easily be retrieved, preferably on the Windows Server 2012 server that will be configured for DirectAccess.

Step 3: Collect Forefront UAG DirectAccess GPO settings

Collect GPO settings using the following procedure.

To collect GPO settings

1. In the Forefront UAG Management console, click **Edit** in **Step 1 Clients and GPOs**.
2. In the Policy Management page, note the names of the GPOs. Then click **Cancel**.

   Note that during a side-by-side migration, the GPOs used on the Windows Server 2012 server must have different names from those used on the Forefront UAG DirectAccess server.

Step 4: Install the Remote Access role

Install the role using the following procedure.

To install the role

1. In the dashboard of the Server Manager console click **Add roles**.
2. Click **Next** until you reach the **Select Server Roles** dialog.
3. On the **Select Server Roles** dialog, select **Remote Access**. Click **Add Required**

Features, and then click **Next**.

4. On the **Select features** dialog, expand **Remote Server Administration Tools**. Expand **Role Administration Tools**, and then select **Remote Access Management Tools**. Click **Next** until you reach the **Confirm installation selections** dialog.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.
7. In addition, turn on the XPS Viewer feature.

Step 5: Configure server and infrastructure settings

Configure server and infrastructure using the following procedure.

To configure server and infrastructure settings

1. Determine the public host name of the Windows Server 2012 Remote Access server. This should be different from the host name of the Forefront UAG server. To determine the public name of the Forefront UAG server, view the value The certificate that the DirectAccess server uses for HTTPS is in the **DirectAccess Server Settings** section of the exported Forefront UAG configuration settings file.
2. Create a public DNS record for this public host FQDN. The record should resolve to the first of the two IPv4 public consecutive addresses on the external adapter of the Remote Access server.

Step 6: Obtain an IP-HTTPS certificate

If you want to use a certificate issued by a CA for IP-HTTPS authentication, obtain a web server certificate with a subject name that matches the FQDN of the Windows 8 server. Optionally, using DirectAccess in Windows Server 2012 you can select to create a self-signed certificate for IP-HTTPS during DirectAccess configuration.

Step 7: Create a security group for DirectAccess clients

Create a new Active Directory security group that contains client computers running Windows® 8 or Windows 7 that will be enabled as DirectAccess clients. Use separate security groups for Windows 8 and Windows 7 computers.

Step 8: Prepare GPOs

Prepare GPOs in all relevant domains for the Remote Access server, DirectAccess clients, and application servers (if needed), and link them to the root of the respective domain. DirectAccess

administrators should have the (Edit settings, delete, modify security) permissions to modify the GPOs.

Step 9: Configure DirectAccess

On the Windows Server 2012 configure DirectAccess using the following procedure.

To configure DirectAccess

1. Before beginning to configure Remote Access, ensure that ISATAP is not enabled on the Windows 8 server. If it is, disable it as follows:
   a. Remove the ISATAP entry from DNS
   b. Wait for replication to complete on all domain controllers
   c. Wait and additional 24 hours for all computers to lose their ISATAP addresses
2. To begin DirectAccess deployment, in the Remote Access Management console, select **Run the Remote Access Setup Wizard**.
3. In the first page of the wizard, select **Deploy DirectAccess only**.
4. Click Step 1: Remote Clients, click **Edit**. Note that if you are logged on as a user that is not a member of the domain administrators security group you will get a warning message stating that you do not have permissions to create the WMI filter. This can be ignored. Click **Next**.
5. On the Deployment Scenario page, select **Deploy full DirectAccess for client access and remote management**. Then click **Next**. This setting allows DirectAccess client computers to access the internal network, and allows DirectAccess administrators to remotely manage DirectAccess client computers located on the Internet.
6. In the Select Groups page, click **Add**. Specify the security groups containing DirectAccess clients. The groups must not currently be used in the Forefront UAG deployment. Click **OK**. Then click **Next**.
7. If Forefront UAG is configured to use force tunneling for DirectAccess clients, select **Use force tunneling**.

   Note
   If you are using these instructions to perform an offline migration, specify the security groups that are defined for DirectAccess clients in Forefront UAG.
8. In the Network Connectivity Assistant page, configure settings for the Network Connectivity Assistant (NCA). NCA replaces the DirectAccess Connectivity Assistant (DCA) used in Forefront UAG. Note that NCA is supported only for Windows 8 computers and cannot be used by client computers running Windows 7. The DCA settings used in Forefront UAG can be verified in the DirectAccess Client Settings section of the exported Forefront UAG configuration settings. Note that NCA supports the use of HTTP and PING probes only. File and HTTPS probes used in DCA are not supported. In addition, administrator scripts and help portal settings are not supported.
9. To add a probe, double-click in the resources list. Specify the first HTTP probe that appears in the Forefront UAG configuration. Then click **OK**.

10. Repeat the last step to add all of the HTTP probes.
11. If **Allow local name resolution** is enabled in the exported Forefront UAG configuration settings, select **Allow DirectAccess clients to use local name resolution**. If the value is disabled, the checkbox should not be enabled.
12. Copy the value specified in **Send log file to email** in the exported Forefront UAG configuration settings file to the **Helpdesk email address** field.
13. Click **Finish** to complete the wizard.
14. In Step 2 Remote Access server, click **Edit**.
15. In the Network Topology page, select **Edge**.
16. In **Type the public name or IPv4 address**, specify the FQDN that clients will use to access the Windows 8 Remote Access server. Ensure that the FQDN is not the same as the FQDN specified for the Forefront UAG IP-HTTPS certificate. Note that the FQDN should be resolvable by a public DNS server, to the public IPv4 address of the Remote Access server. A server certificate is required for the FQDN in order to authenticate the IP-HTTPS connection.

    Note
    If you are using these instructions to configure an offline migration, use the public host name specified for the IP-HTTPS certificate configured in the Forefront UAG deployment.
17. In the Network Adapters page ensure that IP addresses are correct. Verify that the addresses are different than those that appear in the **DirectAccess Server Settings** section of the exported Forefront UAG configuration settings file.

    Note
    If you are using these instructions to configure an offline migration, verify that the addresses are the same as those specified in the **DirectAccess Server Settings** section of the exported Forefront UAG configuration settings file.
18. Click **Browse** to select an IP-HTTPS certificate. You can either purchase a certificate from a commercial CA, or configure Windows Server 2012 DirectAccess to issue a self-signed certificate.

    Note
    If you are using these instructions to configure an offline migration, you can purchase a certificate from a commercial CA, configure Windows Server 2012 DirectAccess to issue a self-signed certificate, or use the exported certificate from the Forefront UAG server.
19. On the Prefix Configuration page, verify settings are correct.
20. On the Authentication page, if standard client authentication is used, select **Active Directory credentials**
21. If Forefront UAG is configured to use smart card authentication, select **Two-factor authentication**.
22. If Forefront UAG is configured to use one-time password (OTP) authentication, note the following:

a. Little OTP information can be shared between the Forefront UAG and the Windows Server 2012 deployments. Windows Server 2012 DirectAccess can only perform OTP authentication over RADIUS. If the Forefront UAG deployment uses RSA SecurID without RADIUS, either the RSA SecurID server must be configured to support RADIUS, or a different server must be used for OTP authentication.
b. To identify OTP servers used by the Forefront UAG, click Admin. Then select Authentication and Authorization Servers.

To configure OTP in the Windows Server 2012 deployment using the same settings, select **Two-factor authentication**. Then select **Use OTP**.

23. Select **Use computer certificates**.
24. Click **Browse** button to select a certificate. Ensure that you use the root certificate that is used in Forefront UAG. To verify Forefront UAG settings verify the value in **root certificate to which remote clients chain** in the **DirectAccess Server Settings** section of the exported Forefront UAG configuration settings file.
25. Select **Enable Windows 7 Clients to connect via DirectAccess**.
26. To use NAP enforcement for the IPsec connection select **Enforce corporate compliance for DirectAccess clients with Network Access Protection (NAP)**. Note the following:
    a. When creating the health policy in NPS, set **Type of network access server** in the **Network Connection Method** section of the **Overview** tab to **Health Registration Authority**.
    b. NAP GPOs must be configured for DirectAccess clients running Windows 8. Ensure that these GPOs are applied only to the DirectAccess clients using the Windows Server 2012 server.
27. If you are not using OTP authentication, click **Finish**. If you have enabled OTP, do the following:

    a. In OTP RADIUS Servers page, specify the name of the OTP RADIUS server, the shared secret, and the port number.
    b. On the OTP CA Servers page, select the OTP CA servers used to issue certificates for OTP authentication. Note that CA servers are detected automatically.
    c. On the OTP Certificate Template page, select a certificate template.
    d. On the OTP Exemptions page, optionally select a security group containing client computers exempt from two-factor authentication.

28. In Step 3 Infrastructure Servers, click **Edit**.
29. On the Network Location Server page, select **The Network Location Server is deployed on a remote web server**. To verify settings in Forefront UAG, check the value **The HTTPS URL that provides clients with location information** in the **Infrastructure Servers (DNS, DC, Management)** section of the exported Forefront UAG configuration

settings file. Click **Validate** to ensure it works, and then click **Next**.

30. On the DNS page, specify the settings that are in the section **Infrastructure Servers (DNS, DC, Management)** of the exported Forefront UAG configuration settings file. Copy any entries that are specified for Forefront UAG that do not appear in the table. In the IP address of DNS server column, fill in one of the following, depending on the type of entry: for exclusion entries the column should remain blank. For DNS64 entries, copy the address that appears in the first row (<Any Suffix>). For internal DNS servers, enter the IPv4 address of the DNS server. In addition, select the local name resolution option to match that specified in the value Local name resolution option in Forefront UAG.
31. On the Management page, specify the management server settings to match the settings in the **Infrastructure Servers (DNS, DC, Management)** section in the exported Forefront UAG configuration settings file. Copy any entry not in the Domain Controllers section, that does not appear in the table. If NAP is enabled, ensure that the HRA server is added to the management servers list. HRA servers are not discovered automatically, so they must be added manually each time a change is made in HRA server deployment settings.
32. Then click **Finish**.
33. In Step 4: Application Servers, click **Edit**.
34. Click **Finish** to keep the default settings.
35. In the details pane of the management console, click **Finish**.
36. On the Review page, in the **GPO Information** section, click **Change**.
37. In **DirectAccess client GPO** click **Browse**.
38. Select the GPO prepared for Windows 8 DirectAccess clients and click Select. Then click OK.
39. Repeat the previous two steps for **DirectAccess server GPO**.
40. In the details pane, click **Apply** to apply the configuration.

See also

- Side-by-Side Migration of Forefront UAG DirectAccess

Offline Migration of Forefront UAG DirectAccess

This topic describes how to perform an offline migration of Forefront UAG DirectAccess to DirectAccess in Windows Server® 2012. In an offline migration the Forefront UAG server is shut down before the Windows Server 2012 Remote Access server is activated. This enables use of existing IP addresses, certificates and FQDNs. However, it creates server downtime, requires all clients to migrate to the new server at once, and is difficult to roll back.

The offline migration consists of the following steps:

- Step 1: Install the Remote Access role on the Windows Server 2012 computer.
- Step 2: Configure server IP addresses.
- Step 3: Obtain a server certificate for IP-HTTPS connections.
- Step 4: Prepare GPOs in relevant domains for the Remote Access server, DirectAccess clients, and the application server if required. DirectAccess administrators should have the correct permissions (edit settings, delete, modify security) to modify the GPOs.
- Step 5: Configure DirectAccess.

For a complete walkthrough, see Offline Migration Steps.

The following graphic illustrates the offline side-by-side migration process.

1

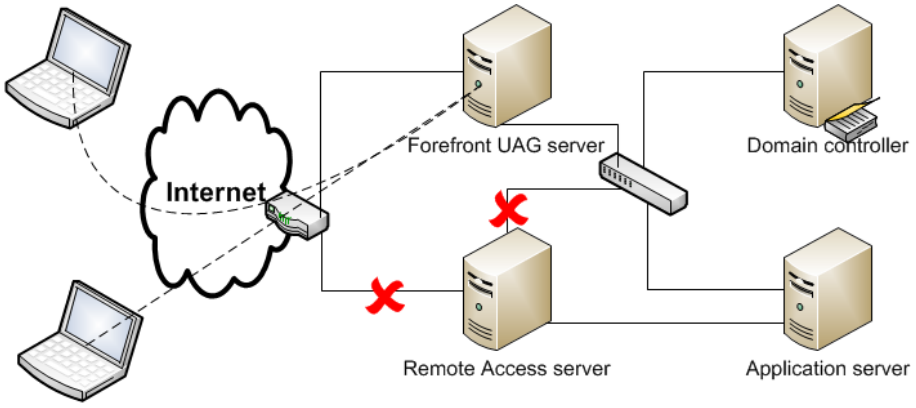

2

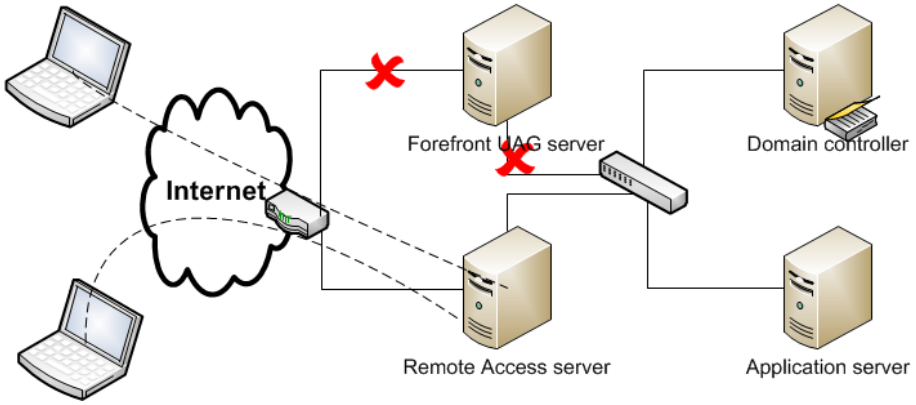

Offline Migration Steps

This topic describes the steps required to perform an offline migration from Forefront UAG SP1 DirectAccess to DirectAccess in Windows Server® 2012.

Task	Description
Step 1: Install the Remote Access role	Configure the Windows Server 2012 computer as a Remote Access server.
Step 2: Configure IP addresses	Configure IP addresses on the Remote Access server.
Step 3: Obtain a server certificate for IP-HTTPS connections	DirectAccess in Windows Server 2012 provides two options for the IP-HTTPS certificate. You can obtain a certificate from a CA, in a similar way to Forefront UAG DirectAccess, or you can configure Windows Server 2012 DirectAccess to automatically issue a self-signed certificate for IP-HTTPS authentication.
Step 4: Prepare GPOs	Prepare the required GPOs.
Step 5: Additional steps	

Step 1: Install the Remote Access role

Export the settings using the following procedure.

To install the role

1. In the dashboard of the Server Manager console click **Add roles**.
2. Click **Next** until you reach the **Select Server Roles** dialog.
3. On the **Select Server Roles** dialog, select **Remote Access**. Click **Add Required Features**, and then click **Next**.
4. On the **Select features** dialog, expand **Remote Server Administration Tools**. Expand **Role Administration Tools**, and then select **Remote Access Management Tools**. Click **Next** until you reach the **Confirm installation selections** dialog.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Step 2: Configure IP addresses

Configure the IP addresses using the following procedure.

To configure addresses

1. On the external network adapter, use the value specified in **DirectAccess server Internet-facing address**, in the **DirectAccess Server Settings** section of the exported

Forefront UAG configuration setting file as the first IP address. For the second IP address, use this address increased by one. For example, 1.2.3.4 and 1.2.3.5.

2. To ensure that ISATAP is not configured, configure an arbitrary IPv6 unique local address (prefix fc00::/7) on the internal network adapter.
3. For the internal network adapter, use the address specified in the **DirectAccess server internal address**, in the **DirectAccess Server Settings** section of the exported Forefront UAG configuration settings file.
4. On the **Select features** dialog, expand **Remote Server Administration Tools**. Expand **Role Administration Tools**, and then select **Remote Access Management Tools**. Click **Next** until you reach the **Confirm installation selections** dialog.
5. On the **Confirm installation selections** dialog, click **Install**.
6. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Step 3: Obtain an IP-HTTPS certificate

Obtain a web server certificate with a subject name that matches the FQDN of the Forefront UAG server. If you want to export the certificate from Forefront UAG and import it to the Remote Access server, see Export a certificate with the private key for instructions. Note that exporting the private key is only possible if the **Make private key exportable** option was checked when the original Forefront UAG certificate was created. Otherwise, the private key cannot be exported, and a new certificate with the same FQDN for the Remote Access server must be created.

Step 4: Prepare GPOs

Prepare GPOs for the Remote Access server, DirectAccess clients, and application servers. DirectAccess administrators should have the correct permissions (edit settings, delete, modify security) to modify the GPOs.

Step 5: Configure DirectAccess

Configure DirectAccess using the instructions described in Side-by-side migration steps.

See also

- Offline Migration of Forefront UAG DirectAccess

Connection Manager Overview

You can use Connection Manager to connect to remote networks by using profiles configured and provided by your network administrator. Most of the required configuration that you need to connect is contained in the profile; you supply your network credentials. Connection Manager

also supports virtual private network (VPN) connections over the Internet by using your broadband network connection or a dial-up modem. For dial-up users, the administrator can provide a phone book with a list of available numbers, allowing you to simply select the best phone number for your current location.

For network administrators, you can use the Connection Manager Administration Kit to create profiles you can distribute to your users.

The documentation for Connection Manager has not been updated for this version of Windows. You can find the documentation for Connection Manager at Connection Manager in the Windows Server TechCenter.

Connection Manager Administration Kit Overview

Connection Manager is a client network connection tool that allows a user to connect to a remote network, such as an Internet service provider (ISP), or a corporate network protected by a virtual private network (VPN) server.

The Connection Manager Administration Kit (CMAK) is a tool that you can use to customize the remote connection experience for users on your network by creating predefined connections to remote servers and networks. To create and customize a connection for your users, you use the CMAK wizard.

The documentation for the Connection Manager Administration Kit has not been updated for this version of Windows. You can find the documentation at Connection Manager Administration Kit in the Windows Server TechCenter.

Add DirectAccess to an Existing Remote Access (VPN) Deployment

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role.

This topic provides an introduction to the Remote Access Enable DirectAccess Wizard, which is used to set up a single Remote Access server with recommended settings after you have already set up a virtual private network (VPN).

Deployment documentation set for Windows Server 2012 Remote Access (DirectAccess)

Following is a list of topics that you can use to deploy Remote Access through three main paths:

- **Basic**

- **Advanced**
- **Enterprise**

Topics related to managing and migrating Remote Access that are available for this release are also listed.

Before you begin your deployment, see the following list of unsupported configurations, known issues, and prerequisites:

- Remote Access (DirectAccess) Unsupported Configurations
- Remote Access (DirectAccess) Known Issues
- Remote Access (DirectAccess) Prerequisites

Basic Remote Access deployment

- Deploy a Single Remote Access Server using the Getting Started Wizard
  - Test Lab Guide: Demonstrate DirectAccess Simplified Setup in an IPv4-only Test Environment in Windows Server 2012

Advanced Remote Access deployment

- Deploy a Single Remote Access Server with Advanced Settings
  - Test Lab Guide: Demonstrate DirectAccess Single Server Setup with Mixed IPv4 and IPv6 in Windows Server 2012

Remote Access deployment in an enterprise

- DirectAccess Capacity Planning
- Deploy Remote Access in a Cluster
  - Test Lab Guide: Demonstrate DirectAccess in a Cluster with Windows NLB
- Deploy Multiple Remote Access Servers in a Multisite Deployment
  - Test Lab Guide: Demonstrate a DirectAccess Multisite Deployment
- Deploy Remote Access with OTP Authentication
  - Test Lab Guide: Demonstrate DirectAccess with OTP Authentication and RSA SecurID
- Deploy Remote Access in a Multi-Forest Environment
- DirectAccess Offline Domain Join

Manage Remote Access

- Use Remote Access Monitoring and Accounting
- Manage DirectAccess Clients Remotely

Migrate Remote Access

- Migrate Remote Access to Windows Server 2012
- Migrate from Forefront UAG SP1 DirectAccess to Windows Server 2012

Scenario description

In this scenario, a single computer running Windows Server 2012 is configured as a Remote Access server with recommended settings after you have already installed and configured VPN. If you want to configure Remote Access with enterprise features such as a load-balanced cluster, multisite deployment, or two-factor client authentication, complete the scenario described in this topic to set up a single server, and then set up the enterprise scenario as described in Deploy Remote Access in an Enterprise.

In this scenario

To set up a single Remote Access server, a number of planning and deployment steps are required.

Planning steps

Planning is divided into two phases:

1. **Plan the Remote Access infrastructure**

   In this phase, you describe the planning that is required to set up the network infrastructure before you begin the Remote Access deployment. It includes planning for the network and server topology, certificates, Domain Name System (DNS), Active Directory and Group Policy Object (GPO) configuration, and the DirectAccess network location server.

2. **Plan the Remote Access deployment**

   In this phase, you describe the planning steps that are required to prepare for the Remote Access deployment. It includes planning for Remote Access client computers, server and client authentication requirements, and infrastructure servers.

Deployment steps

Deployment is divided into three phases:

1. **Configure the Remote Access infrastructure**

   In this phase, you configure the network and routing, firewall settings (if required), certificates, DNS servers, Active Directory and GPO settings, and the DirectAccess network location server.

2. **Configure Remote Access server settings**

   In this phase, you configure the Remote Access client computers, the Remote Access server, and the infrastructure servers.

3. **Verify the deployment**

   In this phase, you verify that the deployment is working as required.

Practical applications

Deploying a single Remote Access server provides the following:

- **Ease of access**

  Managed client computers running Windows 8 and Windows 7 can be configured as DirectAccess client computers. These clients can access internal network resources through DirectAccess any time they are located on the Internet, without the need to sign in to a VPN connection. Client computers that are not running one of these operating systems can connect to the internal network through a VPN. DirectAccess and VPN are managed in the same console and with the same set of wizards.
- **Ease of management**

  DirectAccess client computers that have access to the Internet can be remotely managed by remote access administrators by using DirectAccess, even when the client computers are not located on the internal corporate network. Client computers that do not meet corporate requirements can be remediated automatically by management servers.

Roles and features required for this scenario

The following table lists the roles and features that are required for this scenario:

Role/feature	How it supports this scenario
Remote Access role	The role is installed and uninstalled by using the Server Manager console or Windows PowerShell. This role encompasses DirectAccess, which was previously a feature in Windows Server 2008 R2, and Routing and Remote Access Services, which was previously a role service under the Network Policy and Access Services (NPAS) server role. The Remote Access role consists of two components: 1. DirectAccess and Routing and Remote Access Services (RRAS) VPN: Managed in the Remote Access Management console. 2. RRAS Routing: Managed in the Routing and Remote Access console. The Remote Access Server role is dependent on the following server features: • Internet Information Services (IIS) Web Server:Required to configure the network location server on the Remote Access server, and the default web probe. • Windows Internal Database: Used for local accounting on the Remote Access server.

Role/feature	How it supports this scenario
Remote Access Management Tools feature	This feature is installed as follows: • By default on a Remote Access server when the Remote Access role is installed. Supports the Remote Management console user interface and the Windows PowerShell cmdlets. • Optionally installed on a server not running the Remote Access server role. In this case, it is used for remote management of a Remote Access computer running DirectAccess and VPN. The Remote Access Management Tools feature consists of the following: • Remote Access GUI • Remote Access module for Windows PowerShell Dependencies include: • Group Policy Management Console • RAS Connection Manager Administration Kit (CMAK) • Windows PowerShell 3.0 • Graphical Management Tools and Infrastructure

Hardware requirements

Hardware requirements for this scenario include the following:

Server requirements

- A computer that meets the hardware requirements for Windows Server 2012.
- The server must have at least one network adapter installed, enabled, and joined to the internal network. When two adapters are used, there should be one adapter connected to the internal corporate network, and one connected to the external network (Internet).
- If Teredo is required as an IPv4 to IPv6 transition protocol, the external adapter of the server requires two consecutive public IPv4 addresses. The Enable DirectAccess Wizard does not enable Teredo, even if two consecutive IP addresses are present. To enable Teredo, see Deploy a Single Remote Access Server with Advanced Settings. If a single IP address is available, only IP-HTTPS can be used as the transition protocol.
- At least one domain controller. The Remote Access server and DirectAccess clients must be domain members.

- The Enable DirectAccess Wizard requires certificates for IP-HTTPS and the network location server. If the SSTP VPN is already using a certificate, it is reused for IP-HTTPS. If the SSTP VPN is not configured, you can configure a certificate for IP-HTTPS or use an automatically created self-signed certificate. For the network location server, you can configure a certificate or use an automatically created self-signed certificate.

Client requirements

- A client computer must be running Windows 8 or Windows 7.

  Note
  Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

Infrastructure and management server requirements

- During the remote management of DirectAccess client computers, clients initiate communication with management servers, such as domain controllers, System Center configuration servers, and Health Registration Authority (HRA) servers for services that include Windows and antivirus updates and Network Access Protection (NAP) client compliance. The required servers should be deployed before you begin the Remote Access deployment.
- If remote access requires client NAP compliance, the Network Policy Server (NPS) and the HRA should be deployed before you begin the Remote Access deployment.
- A DNS server running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 with SP2 is required.

Software requirements

Software requirements for this scenario include the following:

Server requirements

- The Remote Access server must be a domain member. The server can be deployed at the edge of the internal network, or behind an edge firewall or other device.
- If the Remote Access server is located behind an edge firewall or network address translation (NAT) device, the device must be configured to allow traffic to and from the Remote Access server.
- The person who deploys remote access on the server requires local administrator permissions on the server, and domain user permissions. In addition, the administrator requires permissions for the GPOs that are used in DirectAccess deployment. To take advantage of the features that restrict a DirectAccess deployment to mobile computers only, permissions to create a WMI filter on the domain controller are required.

Remote access client requirements

- DirectAccess clients must be domain members. Domains that contain clients can belong to the same forest as the Remote Access server, or they can have a two-way trust with the Remote Access server forest or domain.

- An Active Directory security group is required to contain the computers that will be configured as DirectAccess clients. If a security group is not specified when you configure DirectAccess client settings, by default, the client GPO is applied on all laptop computers (that are DirectAccess capable) in the Domain Computers security group. Only the following operating systems can be used as DirectAccess clients: Windows Server 2012, Windows Server 2008 R2, Windows 8 Enterprise, Windows 7 Enterprise, and Windows 7 Ultimate.

  Note
  We recommend that you create a security group for each domain that contains computers that will be configured as DirectAccess clients.

See also

The following table provides links to additional resources.

Content type	References
Remote Access on TechNet	Remote Access TechCenter
Product evaluation	Demonstrate DirectAccess in a cluster with NLB Demonstrate a DirectAccess multisite deployment Demonstrate a DirectAccess multisite deployment
Deployment	Remote Access
Tools and settings	Remote Access PowerShell cmdlets
Community resources	• RRAS Product Team blog • DirectAccess Wiki entries
Related technologies	How IPv6 works

Plan to Enable DirectAccess

Note: Windows Server 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) into a single Remote Access role. This section describes the planning steps that are required to deploy a single Remote Access server running Windows Server 2012 with basic features. The planning phase includes the following steps:

- Step 1: Plan the DirectAccess Infrastructure

  In this phase, you describe the planning that is required to set up the network infrastructure before you begin the Remote Access deployment. It includes planning for the network and

server topology, certificates, Domain Name System (DNS), Active Directory and Group Policy Object (GPO) configuration, and the DirectAccess network location server.

- Step 2: Plan the DirectAccess Deployment

  In this phase, you describe the planning steps that are required to prepare for the Remote Access deployment. It includes planning for Remote Access client computers, server and client authentication requirements, and infrastructure servers.

See also

- Add DirectAccess to an Existing Remote Access (VPN) Deployment
- Enable DirectAccess

Step 1: Plan the DirectAccess Infrastructure

The first step of planning for a basic Remote Access deployment on a single server is to perform planning for the infrastructure required for the deployment. This topic describes the infrastructure planning steps:

Task	Description
Plan network topology and settings	Decide where to place the Remote Access server (at the edge, or behind a Network Address Translation (NAT) device or firewall), and plan IP addressing and routing.
Plan firewall requirements	Plan for allowing Remote Access through edge firewalls.
Plan certificate requirements	Remote Access can use Kerberos or certificates for client authentication. In this basic Remote Access deployment, Kerberos is automatically configured and authentication is accomplished using a self-signed certificate issued automatically by the Remote Access server.
Plan DNS requirements	Plan DNS settings for the Remote Access server, infrastructure servers, local name resolution options, and client connectivity.
Plan Active Directory	Plan your domain controllers and Active Directory requirements.
Plan Group Policy Objects	Decide what GPOs are required in your

Task	Description
	organization and how to create or edit the GPOs.

The planning tasks do not need to be done in a specific order.

Plan network topology and settings

Plan network adapters and IP addressing

1. Identify the network adapter topology you want to use. Remote Access can be set up with either of the following:
   - With two network adapters—Either at the edge with one network adapter connected to the Internet and the other to the internal network, or behind a NAT, firewall, or router device, with one network adapter connected to a perimeter network and the other to the internal network.
   - Behind a NAT device with one network adapter—The Remote Access server is installed behind a NAT device, and the single network adapter is connected to the internal network.
2. Identity your IP addressing requirements:

   DirectAccess uses IPv6 with IPsec to create a secure connection between DirectAccess client computers and the internal corporate network. However, DirectAccess does not necessarily require connectivity to the IPv6 Internet or native IPv6 support on internal networks. Instead, it automatically configures and uses IPv6 transition technologies to tunnel IPv6 traffic across the IPv4 Internet (6to4, Teredo, IP-HTTPS) and across your IPv4-only intranet (NAT64 or ISATAP). For an overview of these transition technologies, see the following resources:
   - IPv6 Transition Technologies
   - IP-HTTPS Tunneling Protocol Specification
3. Configure required adapters and addressing according to the following table. For deployments behind a NAT device using a single network adapter, configure your IP addresses using only the ‘Internal network adapter’ column.

	External network adapter	Internal network adapter[1]	Routing requirements
IPv4 intranet and IPv4 Internet	Configure the following: • One static public IPv4 address with the appropriate subnet masks. • A default gateway IPv4 address of your Internet firewall or local Internet service provider (ISP) router.	Configure the following: • An IPv4 intranet address with the appropriate subnet mask. • A connection-specific DNS suffix of your intranet namespace. The DNS Server must also be configured on the Internal interface. • Do not configure a default gateway on any intranet interfaces.	To configure the Remote Access server to reach all subnets on the internal IPv4 network do the following: 1. List the IPv4 address spaces for all the locations on your intranet. 2. Use the **route add -p** or **netsh interface ipv4 add route** commands to add the IPv4 address spaces as static routes in the IPv4 routing table of the Remote Access server.
IPv6 Internet and IPv6 intranet	Configure the following: • Use the autoconfigured address configuration provided by your ISP. • Use the **route print** command to ensure that a	Configure the following: • If you are not using default preference levels, configure your intranet interfaces with the **netsh interface ipv6 set InterfaceIndex ignoredefaultroutes=enabled** command. This command ensures that additional default routes pointing to intranet routers will not be added to the IPv6 routing	If you have an IPv6 intranet, to configure the Remote Access server to reach all of the IPv6 locations, do the following: 3. 0.1. List the IPv6 address spaces for all

	External network adapter	Internal network adapter[1]	Routing requirements
	default IPv6 route pointing to the ISP router exists in the IPv6 routing table. • Determine whether the ISP and intranet routers are using default router preferences described in RFC 4191, and using a higher default preference than your local intranet routers. If both of these are true, no other configuration for the default route is required. The higher preference for the ISP router ensures that the active default IPv6 route of the Remote Access server points to the IPv6 Internet. Because the Remote Access server is an IPv6 router, if you have a native IPv6 infrastructure, the Internet interface can also reach the domain controllers on the intranet. In this case, add packet filters to the domain controller in the perimeter network that prevent	table. You can obtain the InterfaceIndex of your intranet interfaces from the display of the netsh interface show interface command.	the locations on your intranet. 2. Use the **netsh interface ipv6 add route** command to add the IPv6 address spaces as static routes in the IPv6 routing table of the Remote Access server.

	External network adapter	Internal network adapter[1]	Routing requirements
	connectivity to the IPv6 address of the Internet-facing interface of the Remote Access server.		
IPv6 Internet and IPv4 intranet	The Remote Access server forwards default IPv6 route traffic using the Microsoft 6to4 Adapter interface to a 6to4 relay on the IPv4 Internet. You can configure a Remote Access server for the IPv4 address of the Microsoft 6to4 relay on the IPv4 Internet (used when native IPv6 is not deployed in the corporate network) with the following command : netsh interface ipv6 6to4 set relay name=192.88.99.1 state=enabled command.		

Notes

Note the following:

a. If the DirectAccess client has been assigned a public IPv4 address, it will use the 6to4 transition technology to connect to the intranet. If the DirectAccess client cannot connect to the DirectAccess server with 6to4, it will use IP-HTTPS.

b. Native IPv6 client computers can connect to the Remote Access server over native IPv6, and no transition technology is required.

Plan firewall requirements

If the Remote Access server is behind an edge firewall, the following exceptions will be required for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound.
- If you are deploying Remote Access with a single network adapter, and installing the network location server on the Remote Access server, TCP port 62000 should also be exempted.

The following exceptions will be required for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic

Plan certificate requirements

Certificate requirements for IPsec include a computer certificate used by DirectAccess client computers when establishing the IPsec connection between the client and the Remote Access server, and a computer certificate used by Remote Access servers to establish IPsec connections with DirectAccess clients. For DirectAccess in Windows Server® 2012 the use of these IPsec certificates is not mandatory. The Enable DirectAccess Wizard configures the Remote Access server to act as a Kerberos proxy to perform IPsec authentication without requiring certificates.

1. **IP-HTTPS server**—When you configure Remote Access, the Remote Access server is automatically configured to act as the IP-HTTPS web listener. The IP-HTTPS site requires a website certificate, and client computers must be able to contact the certificate revocation list (CRL) site for the certificate. The Enable DirectAccess wizard tries to use the SSTP VPN certificate. If SSTP is not configured, it checks if a certificate for IP-HTTPS is present in the machine personal store. If none is available, it automatically creates a self-signed certificate.
2. **Network location server**—The network location server is a website used to detect whether client computers are located in the corporate network. The network location server requires a website certificate. DirectAccess clients must be able to contact the CRL site for the certificate. The Enable DirectAccess wizard checks if a certificate for Network Location Server is present in the machine personal store. If not present, it automatically creates a self-signed certificate.

The certification requirements for each of these are summarized in the following table:

IPsec authentication	IP-HTTPS server	Network location server
An internal CA is required to	Public CA—It is recommended	Internal CA—You can use an

IPsec authentication	IP-HTTPS server	Network location server
issue computer certificates to the Remote Access server and clients for IPsec authentication when you don't use the Kerberos proxy for authentication	to use a public CA to issue the IP-HTTPS certificate, this ensures that the CRL distribution point is available externally.	internal CA to issue the network location server website certificate. Make sure that the CRL distribution point is highly available from the internal network.
	Internal CA—You can use an internal CA to issue the IP-HTTPS certificate; however, you must make sure that the CRL distribution point is available externally.	Self-signed certificate—You can use a self-signed certificate for the network location server website; however, you cannot use a self-signed certificate in multisite deployments.
	Self-signed certificate—You can use a self-signed certificate for the IP-HTTPS server; however, you must make sure that the CRL distribution point is available externally. A self-signed certificate cannot be used in a multisite deployment.	

Plan certificates for IP-HTTPS

The Remote Access server acts as an IP-HTTPS listener, and you must manually install an HTTPS website certificate on the server. Note the following when planning:

- Using a public CA is recommended, so that CRLs are readily available.
- In the subject field, specify either the IPv4 address of the Internet adapter of Remote Access server, or the FQDN of the IP-HTTPS URL (the ConnectTo address). If the Remote Access server is located behind a NAT device, the public name or address of the NAT device should be specified.
- The common name of the certificate should match the name of the IP-HTTPS site.
- For the Enhanced Key Usage field, use the Server Authentication object identifier (OID).
- For the CRL Distribution Points field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.
- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcards in the name.

Plan website certificates for the network location server

When planning for the network location server website, note the following:

- In the Subject field, specify either an IP address of the intranet interface of the network location server or the FQDN of the network location URL.
- For the Enhanced Key Usage field, use the Server Authentication OID.
- For the CRL Distribution Points field, a CRL distribution point that is accessible by DirectAccess clients that are connected to the intranet. This CRL distribution point should not be accessible from outside the internal network.
- If you are later planning to configure a multisite or cluster deployment, the name of the certificate should not match the internal name of the Remote Access server.

> **Note**
> Ensure that the certificates for IP-HTTPS and network location server have a **Subject Name**. If the certificate does not have a **Subject Name** but has an **Alternative Name**, then it will not be accepted by the Remote Access wizard.

Plan DNS requirements

In a Remote Access deployment, DNS is required for the following:

- **DirectAccess client requests**—DNS is used to resolve requests from DirectAccess client computers that are not located on the internal network. DirectAccess clients attempt to connect to the DirectAccess network location server in order to determine whether they are located on the Internet, or on the corporate network: If the connection is successful, then clients are determined to be on the intranet and DirectAccess is not used, and client requests are resolved using the DNS server configured on the network adapter of the client computer. If the connection does not succeed, clients are assumed to be on the Internet. DirectAccess clients will use the name resolution policy table (NRPT) to determine which DNS server to use when resolving name requests. You can specify that clients should use DirectAccess DNS64 to resolve names, or an alternative internal DNS server. When performing name resolution, the NRPT is used by DirectAccess clients to identify how to handle a request. Clients request an FQDN or single-label name such as http://internal. If a single-label name is requests, a DNS suffix is appended to make an FQDN. If the DNS query matches an entry in the NRPT, and DNS4 or an intranet DNS server is specified for the entry, then the query is sent for name resolution using the specified server. If a match exists but no DNS server is specified, then this indicates an exemption rule and normal name resolution is applied.

  Note that when a new suffix is added to the NRPT in the Remote Access Management console, the default DNS servers for the suffix can be automatically discovered by clicking the **Detect** button. Auto detection works as follows:

  a. If the corporate network is IPv4-based, or IPv4 and IPv6, the default address is the DNS64 address of the internal adapter on the Remote Access server.
  b. If the corporate network is IPv6-based, the default address is the IPv6 address of DNS servers in the corporate network.
- **Infrastructure servers**—

a. **Network location server**—DirectAccess clients attempt to reach the network location server to determine if they are on the internal network. Clients on the internal network must be able to resolve the name of the network location server, but must be prevented from resolving the name when they are located on the Internet. To ensure this occurs, by default, the FQDN of the network location server is added as an exemption rule to the NRPT. In addition, when you configure Remote Access, the following rules are created automatically:

   i. A DNS suffix rule for root domain or the domain name of the Remote Access server, and the IPv6 addresses corresponding to the intranet DNS servers configured on the Remote Access server. For example, if the Remote Access server is a member of the corp.contoso.com domain, a rule is created for the corp.contoso.com DNS suffix.

   ii. An exemption rule for the FQDN of the network location server. For example, if the network location server URL is https://nls.corp.contoso.com, an exemption rule is created for the FQDN nls.corp.contoso.com.

   IP-HTTPS server—The Remote Access server acts as an IP-HTTPS listener and uses its server certificate to authenticate to IP-HTTPS clients. The IP-HTTPS name must be resolvable by DirectAccess clients using public DNS servers.

   Connectivity verifiers—Remote Access creates a default web probe that is used by DirectAccess client computers use to verify connectivity to the internal network. To ensure the probe works as expected the following names must be registered manually in DNS:

   i. directaccess-webprobehost—should resolve to the internal IPv4 address of the Remote Access server, or to the IPv6 address in an IPv6-only environment.

   ii. directaccess-corpconnectivityhost—should resolve to localhost (loopback) address. A and AAAA record should be created, A record with value 127.0.0.1 and AAAA record with value constructed out of NAT64 prefix with the last 32 bits as 127.0.0.1. The NAT64 prefix can be retrieved by running the cmdlet get-netnattransitionconfiguration.

   Note
   This is valid only in an IPv4-only environment. In an IPv4+IPv6, or IPv6-only environment, only a AAAA record should be created with the loopback IP address ::1.

   You can create additional connectivity verifiers using other web addresses over HTTP or PING. For each connectivity verifier, a DNS entry must exist.

DNS server requirements

- For DirectAccess clients, you must use either a DNS server running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or any DNS server that supports IPv6.

Plan Active Directory

Remote Access uses Active Directory and Active Directory Group Policy Objects as follows:

- **Authentication**—Active Directory is used for authentication. The intranet tunnel uses Kerberos authentication for the user to access internal resources.
- **Group policy objects**—Remote Access gathers configuration settings into group policy objects that are applied to Remote Access servers, clients, and internal application servers.
- **Security groups**—Remote Access uses security groups to gather together and identify DirectAccess client computers, and Remote Access servers. The group policies are applied to the required security group.
- **Extended IPsec policies**—Remote Access can use IPsec authentication and encryption between clients and the Remote Access server. You can extend IPsec authentication and encryption through to specified internal application servers. To do this, see Deploy a Single Remote Access Server with Advanced Settings.

Active Directory Requirements

When planning Active Directory for a Remote Access deployment, the following is required:

- At least one domain controller installed on the Windows Server 2012, Windows Server 2008 R2 Windows Server 2008, or Windows Server 2003 operating systems.

  If the domain controller is on a perimeter network (and therefore reachable from the Internet-facing network adapter of Remote Access server) prevent the Remote Access server from reaching it by adding packet filters on the domain controller, to prevent connectivity to the IP address of the Internet adapter.
- The Remote Access server must be a domain member.
- DirectAccess clients must be domain members. Clients can belong to:
  - Any domain in the same forest as the Remote Access server.
  - Any domain that has a two-way trust with the Remote Access server domain.
  - Any domain in a forest that has a two-way trust with the forest to which the Remote Access domain belongs.

Notes

- The Remote Access server cannot be a domain controller.
- The Active Directory domain controller used for Remote Access must not be reachable from the external Internet adapter of the Remote Access server (the adapter must not be in the domain profile of Windows Firewall).

Plan Group Policy Objects

DirectAccess settings configured when you configure Remote Access are collected into Group Policy Objects (GPO). Three different GPOs are populated with DirectAccess settings, and distributed as follows:

- **DirectAccess client GPO**—This GPO contains client settings, including IPv6 transition technology settings, NRPT entries, and Windows Firewall with Advanced Security connection security rules. The GPO is applied to the security groups specified for the client computers.
- **DirectAccess server GPO**—This GPO contains the DirectAccess configuration settings that are applied to any server configured as a Remote Access server in your deployment. It also contains Windows Firewall with Advanced Security connection security rules.

- **Application servers GPO**—This GPO contains settings for selected application servers to which you optionally extend authentication and encryption from DirectAccess clients. If authentication and encryption are not extended then this GPO is not used.

GPOs are created automatically by the Enable DirectAccess Wizard and a default name is specified for each GPO.

Caution
Use the following procedure to backup all Remote Access Group Policy Objects before executing DirectAccess cmdlets: Back up and Restore Remote Access Configuration

GPOs can be configured in two ways:

1. **Automatically**—You can specify that they are created automatically. A default name is specified for each GPO.
2. **Manually**—You can use GPOs that have been predefined by the Active Directory administrator.

Note that once DirectAccess is configured to use specific GPOs, it cannot be configured to use different GPOs.

Automatically-created GPOs

Note the following when using automatically-created GPOs:

Automatically created GPOS are applied according to the location and link target parameter, as follows:

- For the DirectAccess server GPO, both the location and link parameters point to the domain containing the Remote Access server.
- When client GPOs are created, the location is set to a single domain in which the GPO will be created. The GPO name is looked up in each domain, and filled with DirectAccess settings if it exists. The link target is set to the root of the domain in which the GPO was created. A GPO is created for each domain that contains client computers or application servers, and the GPO is linked to the root of its respective domain.

When using automatically created GPOs, to apply DirectAccess settings, the Remote Access server administrator requires the following permissions:

- GPO create permissions for each domain.
- Link permissions for all the selected client domain roots.
- Link permissions for the server GPO domain roots.
- Create, edit, delete, and modify security permissions are required for the GPOs.
- It is recommended that the Remote Access administrator has GPO read permissions for each required domain. This enables Remote Access to verify that GPOs with duplicate names do not exist when creating GPOs.

Note that if the correct permissions for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but linking will not occur. If this warning is issued links will not be created automatically, even after the permissions are in place. Instead the administrator will need to create the links manually.

Manually-created GPOs

Note the following when using manually-created GPOs:

- The GPOs should exist before running the Remote Access Setup wizard.
- When using manually-created GPOs, to apply DirectAccess settings the Remote Access administrator requires full GPO permissions (Edit, Delete, Modify security) on the manually-created GPOs.
- When using manually created GPOs a search is made for a link to the GPO in the entire domain. If the GPO is not linked in the domain then a link is automatically created in the domain root. If the required permissions to create the link are not available a warning is issued.

Note that if the correct permissions for linking GPOs do not exist, a warning is issued. The Remote Access operation will continue but linking will not occur. If this warning is issued links will not be created automatically, even when the permissions are added later. Instead the administrator will need to create the links manually.

Recovering from a deleted GPO

If a Remote Access server, client, or application server GPO has been deleted by accident and there is no backup available, you must remove the configuration settings and re-configure again. If a backup is available, you can restore the GPO from the backup.

Remote Access Management will display the following error message: **GPO <GPO name> cannot be found**. To remove the configuration settings, take the following steps:

1. Run the PowerShell cmdlet **Uninstall-remoteaccess**.
2. Re-open **Remote Access Management**.
3. You will see an error message that the GPO is not found. Click **Remove configuration settings**. After completion, the server will be restored to an un-configured state.

See also

- Step 2: Plan the Remote Access Deployment [Enable]

Step 2: Plan the DirectAccess Deployment

After planning the Remote Access infrastructure, the next step in enabling DirectAccess is to plan the settings for the Enable DirectAccesss Wizard.

Task	Description
Planning for client deployment	Plan how to allow client computers to connect using DirectAccess. Decide which managed computers will be configured as DirectAccess clients.

Task	Description
Planning for Remote Access server deployment	Plan how to deploy the Remote Access server.

Planning for client deployment

There are two decisions to make when planning your client deployment:

- Will DirectAccess be available to mobile computers only, or to any computer?

  When you configure DirectAccess clients in the Enable DirectAccess wizard, you can choose to allow only mobile computers in the specified security groups to connect using DirectAccess. If you restrict access to mobile computers, Remote Access automatically configures a WMI filter to ensure that the DirectAccess client GPO is applied only to mobile computers in the specified security groups. The Remote Access administrator requires permissions to create or modify group policy WMI filters to enable this setting.

- What security groups will contain the DirectAccess client computers?

  DirectAccess settings are contained in the DirectAccess client GPO. The GPO is applied to computers that are part of the security groups that you specify in the Enable DirectAccess wizard. You can specify security groups contained in any supported domain. Before you configure Remote Access, the security groups should be created. You can add computers to the security group after completing the Remote Access deployment, but note that if you add client computers that reside in a different domain to the security group, the client GPO will not be applied to these clients. For example, if you created SG1 in domain A for DirectAccess clients, and later add clients from domain B to this group, the client GPO will not be applied to clients in domain B. To avoid this issue, create a new client security group for each domain that contains client computers. Alternatively, if you do not want to create a new security group, run the Add-DAClient cmdlet with the name of the new GPO for the new domain.

Planning for Remote Access server deployment

There are a number of decisions to make when planning to deploy your Remote Access server:

- **Network topology**—There are two topologies available when deploying a Remote Access server:
  - **Two adapters**—With two network adapters, Remote Access can be configured with one network adapter connected directly to the Internet, and the other is connected to the internal network. Or alternatively the server is installed behind an edge device such as a firewall or a router. In this configuration one network adapter is connected to the perimeter network, the other is connected to the internal network.
  - **Single network adapter**—In this configuration the Remote Access server is installed behind an edge device such as a firewall or a router. The network adapter is connected to the internal network.

- **Network adapters**—The Enable DirectAccess wizard automatically detects the network adapters configured on the Remote Access server, based on the interfaces used by VPN. You must make sure that the correct adapters are selected.
- **ConnectTo address**—Client computers use the ConnectTo address in order to connect to the Remote Access server. The address that you choose must match the subject name of the IP-HTTPS certificate that you deploy for the IP-HTTPS connection and must be available in the public DNS. See Planning website certificates for IP-HTTPS.
- **IP-HTTPS certificate**—If SSTP VPN is configured, the Enable DirectAccess wizard picks up the certificate used by SSTP for IP-HTTPS. If SSTP VPN is not configured, the wizard will try to see if a certificate for IP-HTTPS has been configured. If not, it will automatically provision self-signed certificates for IP-HTTPS.The wizard automatically enables Kerberos authentication. The wizard will also enable NAT64 and DNS64 for protocol translation in the IPv4-only environment.
- **IPv6 prefixes**—If the wizard detects that IPv6 has been deployed on the network adapters, it automatically creates IPv6 prefixes for the internal network, an IPv6 prefix to assign to DirectAccess client computers, and an IPv6 prefix to assign to VPN client computers. If the automatically generated prefixes are not correct for your native IPv6 or ISATAP infrastructure, you must manually change them. See 1.1 Planning network and server topology and settings.
- **Windows 7 clients**—By default, Windows 7 client computers cannot connect to a Windows Server 2012 Remote Access deployment. If you have Windows 7 client computers in your organization that require remote access to internal resources, you can allow them to connect. Any client computers that you want to allow to access internal resources must be a member of a security group that you specify in the Enable DirectAccess wizard.

  Note
  Allowing Windows 7 client computers to connect using DirectAccess requires you to use computer certificate authentication.
- **Authentication**—The Enable DirectAccess wizard uses Active Directory to authenticate the user credentials. To deploy two-factor authentication, see Deploy Remote Access with OTP Authentication.

See also

- Step 1: Plan the Remote Access infrastructure

Enable DirectAccess

Windows Server® 2012 combines DirectAccess and Routing and Remote Access Service (RRAS) VPN into a single Remote Access role. This overview provides an introduction to the configuration steps required in order to deploy a single Windows Server 2012 Remote Access server with basic settings.

- Step 1: Configure the DirectAccess Infrastructure. This step includes configuring network and server settings, DNS settings and Active Directory settings.
- Step 2: Configure the DirectAccess Server. This step includes configuring DirectAccess client computers, server settings.
- Step 3: Verify the Deployment. This step includes steps for verifying the deployment.

Before starting the deployment, verify the planning steps described in Plan to Enable DirectAccess.

Step 1: Configure the DirectAccess Infrastructure

This topic describes how to configure the infrastructure required for enabling DirectAccess in an existing VPN deployment. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 1: Plan the DirectAccess Infrastructure.

Task	Description
Configure server network settings	Configure the server network settings on the Remote Access server.
Configure routing in the corporate network	Configure routing in the corporate network to make sure traffic is appropriately routed.
Configure firewalls	Configure additional firewalls, if required.
Configure CAs and certificates	The Enable DirectAccess Wizard configures a built in Kerberos proxy that authenticates using user names and passwords. It also configures an IP-HTTPS certificate on the Remote Access server.
Configure the DNS server	Configure DNS settings for the Remote Access server.
Configure Active Directory	Join client computers to the Active Directory domain.
Configure GPOs	Configure GPOs for the deployment, if required.
Configure security groups	Configure security groups that will contain DirectAccess client computers, and any other security groups required in the deployment.
Configure the network location server	The Enable DirectAccess Wizard configures

Task	Description
	the network location server on the DirectAccess server.

Configure server network settings

The following network interface settings are required for a single server deployment in an environment with IPv4 and IPv6. All IP addresses are configured by using **Change adapter settings** in the **Windows Networking and Sharing Center**.

- Edge topology
  - One Internet-facing public static IPv4 or IPv6 address.
  - A single internal static IPv4 or IPv6 address.
- Behind NAT device (two network adapters)
  - A single internal network-facing static IPv4 or IPv6 address.
- Behind NAT device (one network adapter)
  - A single static IPv4 or IPv6 address.

Notes

In the event that the Remote Access server has two network adapters (one classified in the domain profile and the other in a public/private profile), but a single NIC topology will be used, then the recommendation is as follows:

1. Ensure that the 2nd NIC is also classified in the domain profile - Recommended.
2. If the 2nd NIC cannot be configured for the domain profile for any reason, then the DirectAccess IPsec policy must be manually scoped to all profiles using the following Windows PowerShell commands:

```
$gposession = Open-NetGPO –PolicyStore <Name of the server GPO>
Set-NetIPsecRule –DisplayName <Name of the IPsec policy> –
GPOSession $gposession –Profile Any
Save-NetGPO –GPOSession $gposession
```

Configure routing in the corporate network

Configure routing in the corporate network as follows:

- When native IPv6 is deployed in the organization, add a route so that the routers on the internal network route IPv6 traffic back through the Remote Access server.
- Manually configure organization IPv4 and IPv6 routes on the Remote Access servers. Add a published route so that all traffic with an organization (/48) IPv6 prefix is forwarded to the internal network. In addition, for IPv4 traffic, add explicit routes so that IPv4 traffic is forwarded to the internal network.

Configure firewalls

When using additional firewalls in your deployment, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv4 Internet:

- 6to4 traffic—IP Protocol 41 inbound and outbound.
- IP-HTTPS—Transmission Control Protocol (TCP) destination port 443, and TCP source port 443 outbound. When the Remote Access server has a single network adapter, and the network location server is on the Remote Access server, then TCP port 62000 is also required.

When using additional firewalls, apply the following Internet-facing firewall exceptions for Remote Access traffic when the Remote Access server is on the IPv6 Internet:

- IP Protocol 50
- UDP destination port 500 inbound, and UDP source port 500 outbound.

When using additional firewalls, apply the following internal network firewall exceptions for Remote Access traffic:

- ISATAP—Protocol 41 inbound and outbound
- TCP/UDP for all IPv4/IPv6 traffic

Configure CAs and certificates

The Enable DirectAccess Wizard configures a built in Kerberos proxy that authenticates using user names and passwords. It also configures an IP-HTTPS certificate on the Remote Access server.

Configure certificate templates

When you use an internal CA to issue certificates, you must configure a certificate template for the IP-HTTPS certificate and the network location server website certificate.

To configure a certificate template

1. On the internal CA, create a certificate template as described in Creating Certificate Templates.
2. Deploy the certificate template as described in Deploying Certificate Templates.

Configure the IP-HTTPS certificate

Remote Access requires an IP-HTTPS certificate to authenticate IP-HTTPS connections to the Remote Access server. There are three certificate options for the IP-HTTPS certificate:

- **Public**—Supplied by a 3rd party.

  A certificate used for IP-HTTPS authentication. In the case that the certificate subject name is not a wild card, then it must be the externally resolvable FQDN URL used only for the Remote Access server IP-HTTPS connections.

- **Private**—The following are required, if they do not already exist:
  - A website certificate used for IP-HTTPS authentication. The certificate subject should be an externally resolvable fully qualified domain name (FQDN) reachable from the Internet. The certificate is based on the certificate template created in **Configuring certificate templates**.
  - A certificate revocation list (CRL) distribution point that is reachable from a publicly resolvable FQDN.
- **Self-signed**—The following are required, if they do not already exist:

  Note
  Self-signed certificates cannot be used in multisite deployments.
  - A website certificate used for IP-HTTPS authentication. The certificate subject should be an externally resolvable FQDN reachable from the Internet.
  - A CRL distribution point that is reachable from a publicly resolvable fully qualified domain name (FQDN).

Make sure that the website certificate used for IP-HTTPS authentication meets the following requirements:

- The common name of the certificate should match the name of the IP-HTTPS site.
- In the subject field, specify either the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL.
- For the Enhanced Key Usage field, use the Server Authentication object identifier (OID).
- For the CRL Distribution Points field, specify a CRL distribution point that is accessible by DirectAccess clients that are connected to the Internet.
- The IP-HTTPS certificate must have a private key.
- The IP-HTTPS certificate must be imported directly into the personal store.
- IP-HTTPS certificates can have wildcards in the name.

To install the IP-HTTPS certificate from an internal CA

1. On the Remote Access server: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, select the check box for the certificate template created in **Configuring certificate templates**, and if required, click **More information is required to enroll for this certificate**.
8. On the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area,

in **Type**, select **Common Name**.

9. In **Value**, specify either the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL, and then click **Add**.
10. In the **Alternative name** area, in **Type**, select **DNS**.
11. In **Value**, specify either the IPv4 address of the external-facing adapter of the Remote Access server, or the FQDN of the IP-HTTPS URL, and then click **Add**.
12. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
13. On the **Extensions** tab, next to **Extended Key Usage**, click the arrow, and make sure that Server Authentication is in the **Selected options** list.
14. Click **OK**, click **Enroll**, and then click **Finish**.
15. In the details pane of the Certificates snap-in, verify that new certificate was enrolled with Intended Purposes of Server Authentication.

Configure the DNS server

You must manually configure a DNS entry for the network location server website for the internal network in your deployment.

To create the network location server and web probe DNS records

1. On the internal network DNS server: On the **Start** screen, type **dnsmgmt.msc**, and then press ENTER.
2. In the left pane of the **DNS Manager** console, expand the forward lookup zone for your domain. Right click the domain and click **New Host (A or AAAA)**.
3. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the network location server website (this is the name the DirectAccess clients use to connect to the network location server). In the **IP address** box, enter the IPv4 address of the network location server, and then click **Add Host**. On the **DNS** dialog box, click **OK**.
4. On the **New Host** dialog box, in the **Name (uses parent domain name if blank)** box, enter the DNS name for the web probe (the name for the default web probe is directaccess-webprobehost). In the **IP address** box, enter the IPv4 address of the web probe, and then click **Add Host**. Repeat this process for directaccess-corpconnectivityhost and any manually created connectivity verifiers. On the **DNS** dialog box, click **OK**.
5. Click **Done**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-DnsServerResourceRecordA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv4Address <network_location_server_IPv4_address>
Add-DnsServerResourceRecordAAAA -Name <network_location_server_name> -ZoneName
<DNS_zone_name> -IPv6Address <network_location_server_IPv6_address>
```

You must also configure DNS entries for the following:

- **The IP-HTTPS server**—DirectAccess clients must be able to resolve the DNS name of the Remote Access server from the Internet.
- **CRL revocation checking**—DirectAccess uses certificate revocation checking for the IP-HTTPS connection between DirectAccess clients and the Remote Access server, and for the HTTPS-based connection between the DirectAccess client and the network location server. In both cases, DirectAccess clients must be able to resolve and access the CRL distribution point location.

Configure Active Directory

The Remote Access server and all DirectAccess client computers must be joined to an Active Directory domain. DirectAccess client computers must be a member of one of the following domain types:

- Domains that belong in the same forest as the Remote Access server.
- Domains that belong to forests with a two-way trust with the Remote Access server forest.
- Domains that have a two-way domain trust to the Remote Access server domain.

To join client computers to the domain

1. On the **Start** screen, type **explorer.exe**, and then press ENTER.
2. Right-click the Computer icon, and then click **Properties**.
3. On the **System** page, click **Advanced system settings**.
4. On the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
5. In **Computer name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
6. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
7. When you see a dialog box welcoming you to the domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click Close. Click **Restart Now** when prompted.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note that you must supply domain credentials after entering the Add-Computer command below.

```
Add-Computer -DomainName <domain_name>

Restart-Computer
```

Configure GPOs

To deploy Remote Access, you require a minimum of two Group Policy Objects: one Group Policy Object contains settings for the Remote Access server and one contains settings for DirectAccess client computers. When you configure Remote Access, the wizard automatically creates the required Group Policy Objects. However, if your organization enforces a naming convention, or you do not have the required permissions to create or edit Group Policy Objects, they must be created prior to configuring Remote Access.

To create Group Policy Objects, see Create and Edit a Group Policy Object.

Important

The administrator can manually link the DirectAccess Group Policy Objects to an Organizational Unit using these steps:

1. Before configuring DirectAccess, link the created GPOs to the respective Organizational Units.
2. Configure DirectAccess, specifying a security group for the client computers.
3. The Remote Access administrator may or may not have permissions to link the Group Policy Objects to the domain. In either case, the Group Policy Objects will be configured automatically. If the GPOs are already linked to an OU, the links will not be removed, and the GPOs will not be linked to the domain. For a server GPO, the OU must contain the server computer object, or the GPO will be linked to the root of the domain.
4. If the linking to the OU has not been done before running the DirectAccess wizard, then after the configuration is complete, the domain administrator can link the DirectAccess Group Policy Objects to the required Organizational Units. The link to the domain can be removed. Steps for linking a Group Policy Object to an Organization Unit can be found here.

Note

If a Group Policy Object was created manually, it is possible during the DirectAccess configuration that the Group Policy Object will not be available. The Group Policy Object may not have been replicated to the closest Domain Controller to the management computer. In this event, the administrator can wait for replication to complete, or force the replication.

Configure security groups

The DirectAccess settings contained in the client computer Group Policy Object are applied only to computers that are members of the security groups that you specify when configuring Remote Access. In addition, if you are using security groups to manage your application servers, create a security group for these servers.

To create a security group for DirectAccess clients

1. On the **Start** screen, type **dsa.msc**, and then press ENTER. In the **Active Directory Users and Computers** console, in the left pane, expand the domain that will contain the security group, right-click **Users**, point to **New**, and then click **Group**.
2. On the **New Object - Group** dialog box, under **Group name**, enter the name for the security group.
3. Under **Group scope**, click **Global**, under **Group type**, click **Security**, and then click **OK**.
4. Double-click the DirectAccess client computers security group, and on the properties dialog box, click the **Members** tab.
5. On the **Members** tab, click **Add**.
6. On the **Select Users, Contacts, Computers, or Service Accounts** dialog box, select the client computers that you want to enable for DirectAccess, and then click **OK**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
New-ADGroup -GroupScope global -Name <DirectAccess_clients_group_name>

Add-ADGroupMember -Identity DirectAccess_clients_group_name -Members <computer_name>
```

Configure the network location server

The network location server should be on a server with high availability, and a valid SSL certificate trusted by the DirectAccess clients. There are two certificate options for the network location server certificate:

- **Private**—The following are required, if they do not already exist:
  - A website certificate used for the network location server. The certificate subject should be the URL of the network location server. The certificate is based on the certificate template created in **Configuring certificate templates**.
  - A CRL distribution point that is highly available from the internal network.
- **Self-signed**—The following are required, if they do not already exist:

  Note
  Self-signed certificates cannot be used in multisite deployments.

  - A website certificate used for the network location server. The certificate subject should be the URL of the network location server.

Note
If the network location server website is located on the Remote Access server, a website will be created automatically when configuring Remote Access that is bound to the server certificate that you provide.

To install the network location server certificate from an internal CA

1. On the server that will host the network location server website: On the **Start** screen, type **mmc.exe**, and then press ENTER.
2. In the MMC console, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, open **Certificates (Local Computer)\Personal\Certificates**.
5. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
6. Click **Next** twice.
7. On the **Request Certificates** page, select the check box for the certificate template created in **Configuring certificate templates**, and if required, click **More information is required to enroll for this certificate**.
8. On the **Certificate Properties** dialog box, on the **Subject** tab, in the **Subject name** area, in **Type**, select **Common Name**.
9. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
10. In the **Alternative name** area, in **Type**, select **DNS**.
11. In **Value**, enter the FQDN of the network location server website, and then click **Add**.
12. On the **General** tab, in **Friendly name**, you can enter a name that will help you identify the certificate.
13. Click **OK**, click **Enroll**, and then click **Finish**.
14. In the details pane of the Certificates snap-in, verify that new certificate was enrolled with Intended Purposes of Server Authentication.

See also

- Step 2: Configure the Remote Access server
- Install and configure advanced Remote Access
- Deploy advanced Remote Access

Step 2: Configure the DirectAccess Server

This topic describes how to configure the client and server settings required for a basic Remote Access deployment using the Enable DirectAccess Wizard. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 2: Plan the Remote Access deployment.

Task	Description
Configure DirectAccess clients	Configure the Remote Access server with the security groups containing DirectAccess clients.
Configure the Network Topology	Configure Remote Access server settings.
Configure the DNS Suffix Search List	Modify the Suffix search list if desired.
GPO Configuration	Modify the GPOs if desired.

To Start the Enable DirectAcces Wizard

1. In the **Start** screen click **Routing and Remote Access** and the Enable DirectAcces Wizard will start automatically unless you have selected **Do not show this screen again**.
2. If the wizard does not start automatically, right-click the server node in the **Routing and Remote Access** tree, and then click **Enable DirectAccess**.
3. Click **Next**.

Configure DirectAccess clients

For a client computer to be provisioned to use DirectAccess it must belong to the selected security group. After DirectAccess is configured, client computers in the security group are provisioned to receive the DirectAccess group policy.

To configure DirectAccess clients

1. On the **Select Groups** page, click **Add**.
2. On the **Select Groups** dialog box, select the security groups containing DirectAccess client computers.
3. Select the **Enable DirectAccess for mobile computers only** check box to allow only mobile computers to access the internal network.
4. Select the **Use force tunneling** check box to route all client traffic (to the internal network and to the Internet) through the Remote Access server.
5. Click **Next**.

Configure the Network Topology

To deploy Remote Access, you need to configure the Remote Access server with the correct network adapters, a public URL for the Remote Access server to which client computers can connect (the connect to address), and an IP-HTTPS certificate whose subject matches the connect to address.

To configure the network topology

1. On the **Network Topology** page, click the deployment topology that will be used in your organization. In **Type the public name or IPv4 address used by clients to connect to the Remote Access server**, enter the public name for the deployment (this name matches the subject name of the IP-HTTPS certificate, for example, edge1.contoso.com), and then click **Next**.

Configure the DNS Suffix Search List

To configure the DNS suffix search list

1. Select **Configure DirectAccess Clients with DNS client suffix search list** to specify additional suffixes for client name searches.
2. Type a new suffix name in **New Suffix:** and then click **Add**. Additionally, you can change the search order and remove suffixes from **Domain Suffixes to use:**

For DNS clients, you can configure a DNS domain suffix search list that extends or revises their DNS search capabilities. By adding additional suffixes to the list, you can search for short, unqualified computer names in more than one specified DNS domain. Then, if a DNS query fails, the DNS Client service can use this list to append other name suffix endings to your original name and repeat DNS queries to the DNS server for these alternate FQDNs.For computers and servers, the following default DNS search behavior is predetermined and used when completing and resolving short, unqualified names.When the suffix search list is empty or unspecified, the primary DNS suffix of the computer is appended to short unqualified names, and a DNS query is used to resolve the resultant FQDN. If this query fails, the computer can try additional queries for alternate FQDNs by appending any connection-specific DNS suffix configured for network connections.If no connection-specific suffixes are configured or queries for these resultant connection-specific FQDNs fail, then the client can then begin to retry queries based on systematic reduction of the primary suffix (also known as devolution).For example, if the primary suffix were "example.microsoft.com", the devolution process would be able to retry queries for the short name by searching for it in the "microsoft.com" and "com" domains.When the suffix search list is not empty and has at least one DNS suffix specified, attempts to qualify and resolve short DNS names is limited to searching only those FQDNs made possible by the specified suffix list. If queries for all FQDNs formed as a result of appending and trying each suffix in the list are not resolved, the query process fails, producing a "name not found" result.

Warning

If the domain suffix list is used, clients continue to send additional alternate queries based on different DNS domain names when a query is not answered or resolved. Once a name is resolved using an entry in the suffix list, unused list entries are not tried. For this reason, it is most efficient to order the list with the most used domain suffixes first.

Domain name suffix searches are used only when a DNS name entry is not fully qualified. To fully qualify a DNS name, a trailing period (.) is entered at the end of the name.

Note
In a disjoint name space scenario (where one or more domain computers has an DNS suffix that does not match the Active Directory domain to which the computers belong), you should ensure that the search list is customized to include all the required suffixes. The Remote Access wizard will by default configure the Active Directory DNS name as the primary DNS suffix on the client. Admin should ensure that he adds the DNS suffix used by clients for name resolution.

GPO Configuration

GPO Settings – The DirectAccess server GPO name and Client GPO name are listed. Additionally, you can modify the GPO selection settings.

DirectAccess settings configured when you configure Remote Access are collected into Group Policy Objects (GPO). Two GPOs are populated automatically with DirectAccess settings, and distributed as follows:

1. DirectAccess client GPO—This GPO contains client settings, including IPv6 transition technology settings, NRPT entries, and Windows Firewall with Advanced Security connection security rules. The GPO is applied to the security groups specified for the client computers.
2. DirectAccess server GPO—This GPO contains the DirectAccess configuration settings that are applied to any server configured as a Remote Access server in your deployment. It also contains Windows Firewall with Advanced Security connection security rules.

Summary

Once the Remote Access configuration is complete the **Summary** will be displayed. You can change configured setting or click **Finish** to apply the configuration

Step 3: Verify the Deployment

This topic describes how to verify that you have correctly configured your DirectAccess deployment.

To verify access to internal resources through DirectAccess

1. Connect a DirectAccess client computer to the corporate network and obtain the group policy.
2. Click the **Network connections** icon in the notification area to access the DA Media Manager.
3. Click the **DirectAccess Connection**, and you will see that the status is **Connected**

Locally.

4. Connect the client computer to the external network and attempt to access internal resources.

   You should be able to access all corporate resources.

Working with Web Application Proxy

Web Application Proxy is a new Remote Access role service in Windows Server® 2012 R2. Web Application Proxy provides reverse proxy functionality for web applications inside your corporate network to allow users on any device to access them from outside the corporate network. Web Application Proxy preauthenticates access to web applications using Active Directory Federation Services (AD FS), and also functions as an AD FS proxy.

The following resources will help you install and configure the Web Application Proxy:

1. Web Application Proxy Overview
2. Installing and Configuring Web Application Proxy for Publishing Internal Applications
3. Best Practices Analyzer for Web Application Proxy
4. **Web Application Proxy Cmdlets in Windows PowerShell**
5. Video: Installing the Web Application Proxy

Web Application Proxy Overview

This section helps you get started using Web Application Proxy:

1. Getting Started with Web Application Proxy
2. Overview: Connect to Applications and Services from Anywhere with Web Application Proxy
3. Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy

Getting Started with Web Application Proxy

Web Application Proxy is a new Remote Access role service in Windows Server® 2012 R2. Web Application Proxy provides reverse proxy functionality for web applications inside your corporate network to allow users on any device to access them from outside the corporate network. Web Application Proxy preauthenticates access to web applications using Active Directory Federation Services (AD FS), and also functions as an AD FS proxy.

This topic includes the following information:

- **Providing Access to Applications**

- **Protecting Applications from External Threats**
- **Web Application Proxy Technical Overview**

Providing Access to Applications

Web Application Proxy provides organizations with the ability to provide selective access to applications running on servers inside the organization to end users located outside of the organization. The process to make the application available externally is known as publishing. Unlike traditional VPN solutions, when you publish applications through Web Application Proxy end users can gain access only to applications that you publish. However, Web Application Proxy can also be deployed with VPN as part of a Remote Access deployment in your organization. See **Interoperability with Other Remote Access Products**.

- **Publishing Applications**
- **Accessing Applications**

Note

Web Application Proxy can be run on Windows Azure VMs, but there are no recommended scenarios for running it on Windows Azure VMs.

Publishing Applications

Web Application Proxy publishing enables end users to access their organization's applications from their own devices, so that users are not limited to corporate laptops to do their work, they can use their home computer, their tablet, or their smartphone. In addition, end users are not required to install any additional software on their device to access published applications. Web Application Proxy can be used on clients with a standard browser, an Office client or a rich client using OAuth (for example Windows Store apps). Web Application Proxy serves as a reverse proxy for any application that is published through it and as such, the end user experience is the same as if the end user's device connects directly to the application.

Accessing Applications

Web Application Proxy must always be deployed with AD FS. This enables you to leverage the features of AD FS, such as, single sign-on (SSO). This enables users to enter their credentials one time and on subsequent occasions, they will not be required to enter their credentials. SSO is supported by Web Application Proxy for backend servers that use claims-based authentication; for example SharePoint claims-based applications, and Integrated Windows authentication using Kerberos constrained delegation. Integrated Windows authentication-based applications can be defined in AD FS as relying party trusts which can define rich authentication and authorization policies that are enforced in requests to the application.

Protecting Applications from External Threats

Web Application Proxy serves as a barrier between the Internet and your corporate applications. In many organizations, when you deploy Web Application Proxy and publish applications through it, those applications will be available to external users on devices that are not joined to your domain; for example, personal laptops, tablets, or smartphones. These devices are not domain-joined and as such, they are described as unmanaged devices, and are untrusted within the corporate network. Since you want your users to be able to access important information whenever and wherever they are located, you must mitigate the security risk of allowing users access to corporate resources from these unmanaged and untrusted devices. Web Application Proxy provides a number of security features to protect your corporate network from external threats. Web Application Proxy uses AD FS for authentication and authorization to ensure that only users on devices who authenticate and are authorized can access your corporate applications.

- **Defense in Depth**
- **Authentication and Authorization**

Defense in Depth

In the recommended deployment, Web Application Proxy is deployed in a perimeter network between an Internet-facing firewall and a corporate network firewall. However, in addition to the protection provided by the firewalls themselves, Web Application Proxy provides additional protection for your applications from external threats.

- When HTTPS traffic arrives that is directed to an address published by Web Application Proxy, it terminates the traffic and initiates new requests to the published applications. It therefore acts as a session-level buffer between external devices and published applications. That is, when users access published applications, they do not directly access the application, instead, they access the application through Web Application Proxy.
- Any other traffic that arrives at Web Application Proxy is dropped and not forwarded to the published applications. This includes any illegal HTTP or HTTPS requests that might be used as part of denial of service attacks, zero day attacks, SSL attacks, and so on.
- Any authenticated request that arrives at Web Application Proxy containing an authentication token from AD FS will be inspected to make sure that the token received was intended for the client sending the token. This is done by checking that the device (through the Workplace Join certificate) corresponds to the claim within the token that identified the device when authenticated to AD FS.

Authentication and Authorization

To protect access to applications in your organization, it is recommended to allow access only to authenticated and authorized users. When you publish applications through Web Application Proxy, this is achieved through the use of AD FS, which provides authentication and enforces authorization for the published applications.

Note

Web Application Proxy also allows pass-through preauthentication, which enables you to publish applications that do not require preauthentication or whose clients do not support the available authentication capabilities.

- **Authenticating Users and Devices**
- **Authentication Capabilities**

Authenticating Users and Devices

When you publish applications through Web Application Proxy, the process by which users and devices are authenticated before they gain access to applications is known as preauthentication. Web Application Proxy supports two forms of preauthentication:

- AD FS preauthentication—When using AD FS for preauthentication, the user is required to authenticate to the AD FS server before Web Application Proxy redirects the user to the published web application. This ensures that all traffic to your published web applications is authenticated.
- Pass-through preauthentication—Users are not required to enter credentials before they connect to published web applications.

> **Note**
> Pass-through preauthentication has no impact on whether an application requires users to provide credentials to the application. That is, an application configured with pass-through preauthentication does not require users to enter credentials to get into the corporate network, but may require users to enter credentials to view the application content.

To easily access applications published by Web Application Proxy, and to use AD FS preauthentication end users should use one of the following clients:

- Any client that supports HTTP redirects; for example, a web browser. Web Application Proxy performs the appropriate action on the incoming request to redirect the user to an authentication address and back to the original web address, this time with the authentication proof.
- Any client that uses MSOFBA; for example, Word, Excel, or PowerPoint. In this case, a user attempts to access a document from their Recent Documents list that is stored on a server within the corporate network.
- Windows Store apps and RESTful applications with clients that use the Web Authentication Broker for authentication. A user can open an app on their device which obtains a token from AD FS via the Web Authentication Broker, and includes that token in the HTTP Authorization header in subsequent requests to the app.

> **Note**
> Depending on the client used to access the published application, Web Application Proxy decides how to process the request.

Authentication Capabilities

When you use AD FS for authentication, you also benefit from all of the features that AD FS provides:

- Workplace Join—This is a new feature in AD FS in Windows Server 2012 R2. It allows users to join devices to the workplace that would not normally be domain-joined; for example, personal laptops, tablets, and smartphones. When this feature is enabled, the AD FS administrator can configure all applications, or individual applications, to require devices to be registered before they can gain access to published applications. For more information, see Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
- SSO—This allows users to enter credentials once and be authenticated to all supported published applications. See Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications.
- Multifactor authentication (MFA)—AD FS can be configured to require users to authenticate with more than one authentication scheme; for example, a one-time password or a smart card. See Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.
- Multifactor access control—Access control in AD FS is implemented with authorization claim rules that are used to issue a permit or deny claims that will determine whether a user or a group of users will be allowed to access AD FS-secured resources or not. Authorization rules can only be set on relying party trusts. All of the above features can be combined, as required, to provide stricter security for confidential applications, or ignored for less confidential applications. See Overview: Manage Risk with Multi-Factor Access Control.

When you publish applications through Web Application Proxy you are not required to configure the AD FS authentication features mentioned above. This allows you to provide access to devices that are not able to join the workplace, or provide additional factors of authentication, such as kiosks.

Web Application Proxy Technical Overview

When you decide to use Web Application Proxy in your organization, we recommend that you deploy your Web Application Proxy servers behind a frontend firewall to separate it from the Internet, or between two firewalls; a frontend firewall to separate it from the Internet, and a backend firewall to separate it from the corporate network. In this topology, Web Application Proxy provides a protection layer against malicious users that may be coming from the Internet. No other servers are required to be located in this perimeter network; that is, your AD FS servers are located in the corporate network and can only be reached via Web Application Proxy using its built-in AD FS proxy functionality.

The following diagram shows a typical topology for deploying Web Application Proxy in a perimeter network between two firewalls.

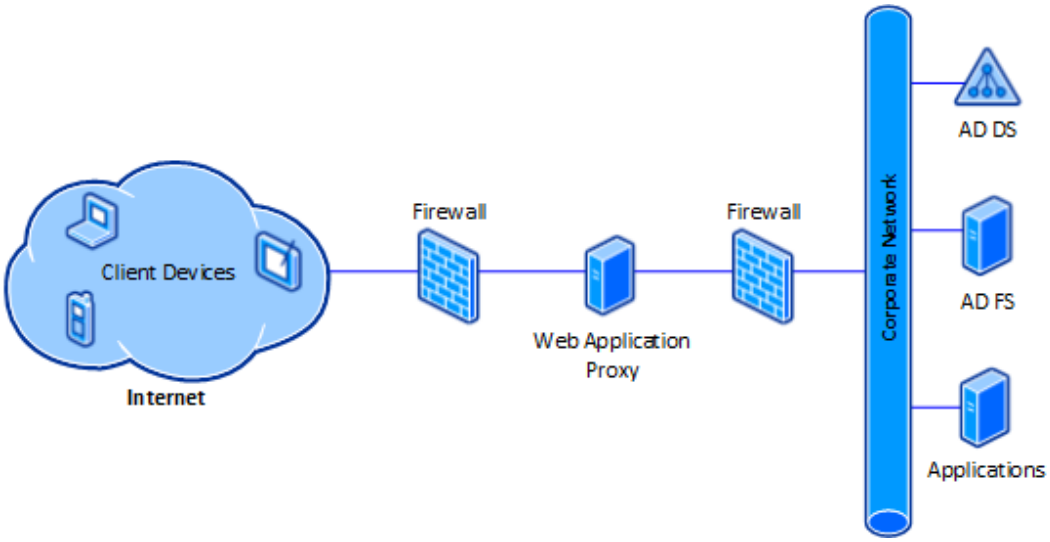

- **Web Application Proxy Configuration Storage**
- **AD FS Proxy Functionality**
- **Managing Web Application Proxy**
- **Interoperability with Other Remote Access Products**

Web Application Proxy Configuration Storage

The Web Application Proxy configuration is stored on the AD FS servers in your organization; therefore, Web Application Proxy servers require connectivity to the AD FS servers. In addition, after configuring the first Web Application Proxy server, you can install additional Web Application Proxy servers to create a cluster deployment. When you install the role service on the new server in the cluster, the configuration is automatically transferred to the new server after completing the Web Application Proxy Configuration Wizard.

Since Web Application Proxy stores its configuration on the AD FS servers it has no locally stored configuration information.

AD FS Proxy Functionality

The Web Application Proxy role service is also an AD FS proxy. That is, Web Application Proxy listens to all of the end-points that AD FS listens to. Web Application Proxy also forwards any requests from the Internet to AD FS and responses from AD FS to the Internet. Note that the Web Application Proxy role service is a replacement for the AD FS proxy role.

Creating a proxy in your organization for your Federation Service adds additional security layers to your AD FS deployment. Consider deploying Web Application Proxy in your organization's perimeter network when you want to:

- Prevent external client computers from directly accessing your AD FS servers. By deploying a Web Application Proxy server in your perimeter network, you effectively isolate your AD FS

servers. Web Application Proxy servers do not have access to the private keys that are used to produce tokens.

- Provide a convenient way to differentiate the sign-in experience for users who are coming from the Internet as opposed to users who are coming from your corporate network using Integrated Windows authentication.

Managing Web Application Proxy

Web Application Proxy uses a number of tools and features provided by Windows Server 2012 R2 to enable you to easily install, deploy, and manage it in your corporate deployments.

- Web Application Proxy is a role service in Windows Server 2012 R2. This allows you to easily install Web Application Proxy in your deployment using Server Manager or Windows PowerShell.
- Web Application Proxy is integrated into the Remote Access Management console, allowing you to manage your Web Application Proxy servers and other Remote Access technologies, such as DirectAccess and VPN from the same Remote Access Management console.
- Web Application Proxy provides full functionality through a set of Windows PowerShell commands and a Windows Management Instrumentation (WMI) API.
- To aid troubleshooting, Web Application Proxy:
  - Writes events to the Windows Event log.
  - Exposes a number of performance counters.
  - Has a dedicated Best Practices Analyzer (BPA).

Interoperability with Other Remote Access Products

Web Application Proxy is a role service of the Remote Access role in Windows Server 2012 R2. You can install Web Application Proxy side-by-side with Remote Access in the following scenarios:

DirectAccess	VPN	Web Application Proxy
Single server deployment	Single server deployment	Single server deployment
Multisite deployment	Multiple server deployment	Not supported on the same server
Not supported on the same server	Multiple server deployment	Multiple server deployment
Cluster deployment[1]	Multiple server deployment	Multiple server deployment[2]

Notes

1—In a pre-existing DirectAccess cluster deployment, you can install Web Application Proxy only using Windows PowerShell.

2—In a pre-existing multiple server Web Application Proxy deployment, you can install DirectAccess only using Windows PowerShell.

Web Application Proxy provides application publishing capabilities, similar to Forefront Unified Access Gateway (UAG). However, Web Application Proxy interacts with other servers and services to provide a more streamlined deployment. This helps you to concentrate on configuring only the necessary parts of your deployment. It is recommended that for any new deployments where you require application publishing capabilities for the scenarios described above, you should use Web Application Proxy.

Overview: Connect to Applications and Services from Anywhere with Web Application Proxy

This scenario describes how you can use Web Application Proxy – a new Remote Access role service in Windows Server® 2012 R2 – to provide access to a sample web application using claims-based authentication and a sample website using Integrated Windows authentication, both websites use Active Directory Federation Services (AD FS) preauthentication. This scenario also uses the following AD FS features:

- **Workplace Join**—Joining devices to the workplace connects these devices with Active Directory in your workplace. When you join personal devices to your workplace they become known devices and will provide seamless second factor authentication and single-sign-on to workplace resources and applications.

  This scenario configures the device registration service (DRS) to enable you to join the client device to the workplace.

- **Multifactor authentication**—This enables you to require users to provide more than one form of authentication when connecting to published applications and services. For example, using one-time passwords or smart cards. You can configure Web Application Proxy and AD FS to use multifactor authentication for all authentication requests, or per-application. In addition, configuring AD FS to allow access to only registered devices creates two-factor seamless authentication because the user must provide credentials, and the device must be registered.

  This scenario uses certificate authentication to provide the additional factor when authenticating.

- **Multifactor access control**—Access control in AD FS is implemented with authorization claim rules that are used to issue a permit or deny claims that will determine whether a user or a group of users will be allowed to access AD FS-secured resources or not. In AD FS in Windows Server 2012 R2, access control is enhanced with multiple factors, including user, device, location, and authentication data. This is made possible by a greater variety of claim types available for the authorization claim rules.

Web Application Proxy also provides built-in AD FS proxy capabilities. The following diagram shows the topology used in this scenario for Web Application Proxy to publish Microsoft applications and other line-of-business (LOB) applications.

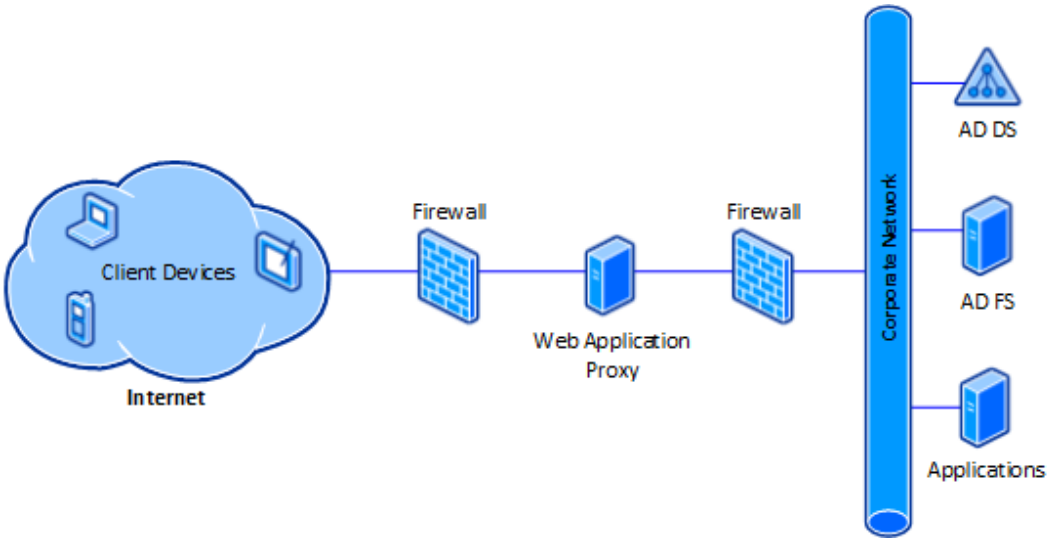

- Providing Access to Applications and Services
- Using Active Directory Federation Services
- AD FS Proxy
- Roles and Features Included in this Scenario
- Scenario Steps

Providing Access to Applications and Services

This scenario demonstrates how to plan and deploy Web Application Proxy in your organization to provide end users located outside of an organization access to applications and services running on servers inside the organization. Web Application Proxy publishing enables end users to access their organization's applications from their own devices, so that users are not limited to corporate laptops to do their work, they can use their home computer, their tablet, or their smartphone. Web Application Proxy can be used on clients with a standard browser, an Office client or a rich client using OAuth (for example Windows Store apps). Web Application Proxy serves as a reverse proxy for any application that is published through it and as such, the end user experience is the same as if the end user's device connects directly to the application.

Using Active Directory Federation Services

This scenario describes the additions and changes that you must make to your AD FS servers to provide the following functionality:

- Application publishing—For all applications and services that you want to publish through Web Application Proxy, you must configure a relying party on the AD FS server.

- Authentication—No specific configuration is required to provide authentication for published applications. However, to use workplace join, MFA, or multifactor access control in your deployment, you must perform additional configuration as described in the following guides:
  - Workplace Join—Overview: Join to Workplace from Any Device for SSO and Seamless Second Factor Authentication Across Company Applications
  - MFA—Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications
  - Multifactor access control—Overview: Manage Risk with Multi-Factor Access Control.

AD FS Proxy

This scenario does not describe using Web Application Proxy as a proxy for AD FS. However, this functionality is enabled by default when you install the Web Application Proxy role service. Any AD FS endpoint that is enabled for proxy publishing is automatically published by Web Application Proxy after completing the Web Application Proxy Configuration Wizard.

Roles and Features Included in this Scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Federation Services Overview	AD FS is required to provide authentication and authorization services to Web Application Proxy and to store the Web Application Proxy configuration.
Remote Access (DirectAccess, Routing and Remote Access) Overview	Remote Access is the role containing the Web Application Proxy role service.
Active Directory Domain Services Overview	Active Directory® Domain Services is required as a prerequisite before you can deploy AD FS.
Web Server (IIS) Overview	The Web Server (IIS) role is used in this scenario to host a sample application that can be published by Web Application Proxy.

Scenario Steps

This scenario contains the following steps:

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy—This walkthrough contains the following steps:

- Step 1: Attempt to access the web application from an Internet client
- Step 2: Configure the Web Application Proxy server and publish the application
- Step 3: Configure and test accessing a site using Integrated Windows authentication
- Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control

See Also

Web Application Proxy Overview

Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy

This walkthrough provides instructions for connecting to applications and services from anywhere with Web Application Proxy. For more information about connecting to applications and services from anywhere with Web Application Proxy and Windows Server® 2012 R2, see Overview: Connect to Applications and Services from Anywhere with Web Application Proxy and the video for Installing the Web Application Proxy.

Before you perform this walkthrough, make sure you have completed the lab setup guide, see Set up the lab environment for AD FS in Windows Server 2012 R2. Also make sure you have completed the following planning steps that are required in order to deploy Web Application Proxy and publish applications through it using Active Directory Federation Services (AD FS) for authentication and authorization, including workplace join, multifactor authentication (MFA), and multifactor access control.

- Plan the Web Application Proxy infrastructure—Before you deploy Web Application Proxy, you must plan the infrastructure. For this walkthrough, this includes:
  - The Web Application Proxy server should be located between the corporate network and the Internet; in this walkthrough, this is achieved by configuring two network adapters on the Web Application Proxy server.
  - For clients to reach a published web application they must be able to resolve the external address of the application that is published by Web Application Proxy. In this walkthrough, this is achieved through the hosts file on the client.
- Plan the Web Application Proxy server
  - It is recommended to join the Web Application Proxy server to the AD DS domain. Joining the Web Application Proxy server to the domain allows you to expand this walkthrough to demonstrate publishing applications that use Integrated Windows authentication.
  - Make sure that the server time of the Web Application Proxy and AD FS servers are synchronized.

- To install the Web Application Proxy role service, you must have local administrator permissions on the server.
- To configure the Web Application Proxy server to connect to AD FS, you are required to enter the credentials of the local administrator on the AD FS server.
- Make sure that there is a certificate in the Personal certificates store on the Web Application Proxy server that covers the federation service name. This certificate is required for AD FS proxy functionality.

  To locate the federation service name value, open the AD FS Management console, click **Service**, in the **Actions** pane, click **Edit Federation Service Properties**, and then find the value in the **Federation Service name** text box.
- Make sure that there is a certificate in the Personal certificates store on the Web Application Proxy server that covers the external address of the sample web application; that is, the address that you will configure when you publish the application through Web Application Proxy. In this walkthrough, you can use a self-signed certificate, or a certificate from an Enterprise certification authority (CA).

Note

In a production environment, it is recommended to use certificates from a public CA on your Web Application Proxy server.

Task	Description
Step 1: Attempt to access the web application from an Internet client	Configure an edge server and demonstrate that before configuring Web Application Proxy and publishing the application, the web application server is protected from the Internet, and clients are unable to access the application.
Step 2: Configure the Web Application Proxy server and publish the application	Configure the edge server as a Web Application Proxy server, publish the web application, and demonstrate that with the default AD FS authentication scheme, a client device can access the published application from the Internet.
Step 3: Configure and test accessing a website using Integrated Windows authentication	Configure an additional website on WebServ1 to use Integrated Windows authentication, configure Kerberos constrained delegation in the lab, publish the application through Web Application Proxy, and demonstrate accessing this application from the Internet.
Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control	Configure Workplace Join, multifactor authentication (MFA), and multifactor access control, and demonstrate how these additional

Task	Description
	authentication factors change the user experience when client devices attempt to access the published web application.

Step 1: Attempt to access the web application from an Internet client

In this step, you will configure an edge server and attempt to access the sample application. This step does not install the Web Application Proxy role service on the server. After configuring the edge server you will modify the hosts file on the client computer and attempt to connect to the sample web application through the edge server.

To install the edge server

1. Create a Windows Server 2012 R2 server with two network adapters, one connected to the corporate network, one connected to the simulated Internet.
2. Configure the static IP addresses on the server. For example, the network adapter on the corporate network should have an IP address of 10.0.0.2, the network adapter on the simulated Internet should have an IP address of 131.107.0.2.
3. Join the server to the AD DS domain.

To access the application

1. In your lab environment, connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
2. On the client computer, open Notepad as the administrator.
3. In the Notepad window, open the hosts file located at %windir%\System32\Drivers\etc.
4. Add two new lines to the file; for example:

```
131.107.0.2       webserv1.contoso.com
131.107.0.2       adfs1.contoso.com
131.107.0.2       enterpriseregistration.contoso.com
```

   - The first line contains the IP address of the Internet adapter of the edge server and the external address of the sample application that you will configure when you publish the application through Web Application Proxy.
   - The second line contains the same IP address and the address of the federation server.
   - The third line contains the same IP address and the address required for Workplace Join.
5. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.

The edge server is not providing proxy or publishing functionality; therefore, you will not be able to access the application.

Step 2: Configure the Web Application Proxy server and publish the application

In this step you will:

- Configure the Web Application Proxy server.
- Connect the Web Application Proxy server to the AD FS server using the Web Application Proxy Configuration Wizard.
- Publish the sample application.
- Attempt to connect to the sample application using the default AD FS authentication scheme.

Before you can publish the application, make sure you have done the following:

- Create a relying party trust for the sample application in the AD FS Management console. If you did not do this when setting up your environment, see Create a relying party trust on your federation server.
- Verify that a certificate on the Web Application Proxy server is suitable for the sample application you want to publish.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To install the Web Application Proxy role service

1. On the edge server, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. In the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. In the **Add Roles and Features Wizard**, click **Next** three times to get to the server role selection screen.
4. On the **Select server roles** dialog, select **Remote Access**, and then click **Next**.
5. Click **Next** twice.
6. On the **Select role services** dialog, select **Web Application Proxy**, click **Add Features**, and then click **Next**.
7. On the **Confirm installation selections** dialog, click **Install**.
8. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature Web-Application-Proxy -IncludeManagementTools
```

To configure Web Application Proxy

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
4. In the **Web Application Proxy Configuration Wizard**, on the **Welcome** dialog, click **Next**.
5. On the **Federation Server** dialog, do the following, and then click **Next**:
   - In the **Federation service name** box, enter the fully qualified domain name (FQDN) of the federation service; for example, adfs1.contoso.com.

     To locate this value, open the AD FS Management console, click **Service**, in the **Actions** pane, click **Edit Federation Service Properties**, and then find the value in the **Federation Service name** text box.
   - In the **User name** and **Password** boxes, enter the credentials of a local administrator account on the AD FS servers.
6. On the **AD FS Proxy Certificate** dialog, in the list of certificates currently installed on the Web Application Proxy server, select a certificate to be used by Web Application Proxy for AD FS proxy functionality, and then click **Next**.
7. On the **Confirmation** dialog, review the settings. If required, you can copy the PowerShell cmdlet to automate additional installations. Click **Configure**.
8. On the **Results** dialog, verify that the configuration was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following command will prompt you to enter credentials of a local administrator account on the AD FS servers.

```
Install-WebApplicationProxy -CertificateThumprint
'1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b' -FederationServiceName adfs1.contoso.com
```

To publish the sample claims-based application

1. On the Web Application Proxy server, in the Remote Access Management console, in the **Navigation** pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Active Directory Federation Services (AD FS)**,

and then click **Next**.

4. On the **Relying Party** page, in the list of relying parties select the relying party for the sample application, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://webserv1.contoso.com/claimapp/.
   - In the **External certificate** list, select a certificate whose subject name covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that application published successfully, and then click **Close**.

   Publishing the application may take a couple of minutes.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WebApplicationProxyApplication

    -BackendServerURL 'https://webserv1.contoso.com/claimapp/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://webserv1.contoso.com/claimapp/'

    -Name 'Claims App'

    -ExternalPreAuthentication ADFS

    -ADFSRelyingPartyName 'App_Relying_Party'
```

To access the application

1. On the client computer, open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
2. On the sign in page, enter the credentials of the test user that you created when you set up the environment.
3. The application appears in the web browser.

Step 3: Configure and test accessing a website using Integrated Windows authentication

In this step you will configure a website on WebServ1 to use Integrated Windows authentication. You can publish this website as an additional application through Web Application Proxy to demonstrate how Web Application Proxy uses Kerberos constrained delegation to authenticate users to non-claims-based applications.

Important
You must join the Web Application Proxy server to the Active Directory® Domain Services domain before you can publish applications that use Integrated Windows authentication.

Install Windows Authentication on WebServ1

When you install the Web Server (IIS) role, by default, it does not install the option to use Integrated Windows authentication (referred to by IIS as Windows Authentication). This procedure installs Windows Authentication.

To install Windows Authentication on WebServ1

1. On the edge server, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. In the **Quick Start** tab of the **Welcome** tile on the **Dashboard** page, click **Add roles and features**. Alternatively, you can click **Add Roles and Features** on the **Manage** menu.
3. In the **Add Roles and Features Wizard**, click **Next** three times to get to the server role selection screen.
4. On the **Select server roles** dialog, expand **Web Server (IIS)**, expand **Web Server**, expand **Security**, select **Windows Authentication**, and then click **Next**.
5. On the **Select features** dialog, click **Next**.
6. On the **Confirm installation selections** dialog, click **Install**.
7. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Create a new website using IIS

The WebServ1 server already hosts a website that uses claims-based authentication. This procedure creates a new website that uses Integrated Windows authentication.

To create a new website

1. Open an Explorer window and go to C:\inetpub\.
2. Create a new folder called NonClaims.
3. Open C:\inetpub\wwwroot\ and copy the contents to the new folder

C:\inetpub\NonClaims\.

Important
Do not copy the web.config file.

4. On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **inetmgr.exe**, and then press ENTER.
5. In the IIS console, in the **Connections** pane, expand **WebServ1**, expand **Sites**, right-click **Default Web Site**, and then click **Add Virtual Directory**.
6. On the **Add Virtual Directory** dialog box, in the **Alias** box, type **NonClaims**.
7. Next to the **Physical Path** box, type **C:\inetpub\NonClaims**, and then click **OK**.
8. In the **Connections** pane, click the **NonClaims** virtual directory.
9. In the middle pane, in the **IIS** section, double-click **Authentication**.
10. In the **Authentication** pane, right-click **Anonymous Authentication**, and then click **Disable**. Right-click **Windows Authentication**, and then click **Enable**.

Create a non-claims-aware relying party trust

To publish the application through Web Application Proxy, you must first create a non-claims-aware relying party trust on the AD FS server.

To create a non-claims-aware relying party trust

1. On the AD FS server, in the AD FS Management console, in the left pane, click **Trust Relationships**.
2. In the **Actions** pane, click **Add Non-Claims-Aware Relying Party Trust**.
3. In the **Add Non-Claims-Aware Relying Party Trust Wizard**, on the **Welcome** page, click **Start**.
4. On the **Specify Display Name** page, in **Display name**, enter a name for this trust; for example, **Non-Claims Application**, and then click **Next**.
5. On the **Configure Identifiers** page, in **Non-claims-aware relying party trust identifier**, enter an arbitrary URL as an identifier for this trust, click **Add**, and then click **Next**.
6. Click **Next** twice.
7. On the **Finish** page, click **Close**.

   The **Edit Claim Rules for <relying party>** dialog box opens.
8. On the **Edit Claim Rules for <relying party>**, click **Add Rule**.
9. In the **Add Issuance Authorization Claim Rule Wizard**, on the **Select Rule Template** page, in the **Claim rule template** list, select **Permit All Users**, and then click **Next**.
10. On the **Configure Rule** page, click **Finish**, and then on the **Edit Claim Rules for <relying party>** dialog box, click **OK**.

Configure Kerberos constrained delegation

To allow users to access applications that use Integrated Windows authentication, the Web Application Proxy server must be able to provide impersonation for users to the published application.

Tip

This procedure assumes that the Web Application Proxy server is named EDGE1.

To configure Kerberos constrained delegation

1. On the domain controller, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. Click **Tools**, and then click **ADSI Edit**.
3. On the **Action** menu, click **Connect To**, and then on the **Connection Settings** dialog box, accept the default settings to connect to the default naming context, and then click **OK**.
4. In the left pane, expand **Default naming context**, expand **DC=contoso,DC=com**, expand **CN=Computers**, right-click **CN=EDGE1**, and then click **Properties**.
5. On the **CN=EDGE1 Properties** dialog box, on the **Attribute Editor** tab, in the **Attributes** list, select **servicePrincipalName**, and then click **Edit**.
6. On the **Multi-valued String Editor** dialog box, in **Value to add**, enter **HTTP/EDGE1.contoso.com** and click **Add**. Then enter **HTTP/EDGE1** and click **Add**.

   The **Values** list now contains two new entries; for example, **HTTP/EDGE1.contoso.com** and **HTTP/EDGE1**.
7. On the **Multi-valued String Editor** dialog box, click **OK**.
8. On the **CN=EDGE1 Properties** dialog box, click **OK**.
9. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
10. In the navigation pane, under **contoso.com**, click **Computers**. In the details pane, right-click the edge server, and then click **Properties**.
11. On the **EDGE1 Properties** dialog box, on the **Delegation** tab, click **Trust this computer for delegation to specified services only**, and then click **Use any authentication protocol**.
12. Click **Add**, and on the **Add Services** dialog box, click **Users or Computers**.
13. On the **Select Users or Computers** dialog box, in **Enter the object names to select**, enter the name of the web server; for example, WebServ1, and then click **OK**.
14. On the **Add Services** dialog box, in the **Available services** list, select the **http** service type, and then click **OK**.
15. On the **EDGE1 Properties** dialog box, click **OK**.

Test accessing the application internally

To make sure that the application is working correctly, you should test accessing the application from the Web Application Proxy server before continuing.

To access the application internally

1. On the Web Application Proxy server, make sure that IE Enhanced Security Configuration is turned off.
2. Open an Internet explorer window and go to the non-claims-based web application; for example, https://webserv1.contoso.com/nonclaims/.
3. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   The default IIS website appears.

Publish the application

You must publish the application through Web Application Proxy to make it available to external users.

To publish the non-claims-based application

1. On the Edge server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER.
2. In the navigation pane, under **Configuration**, click **Web Application Proxy**.
3. In the **Tasks** pane, click **Publish**.
4. In the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
5. On the **Preauthentication** page, make sure that Active Directory Federation Services (AD FS) is selected, and then click **Next**.
6. On the **Relying Party** page, select the non-claims application relying party, and then click **Next**.
7. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application to identify it in the list of published applications; for example, **Non-Claims Application**.
   - In the **External URL** box, enter the external URL for this application; for example, **https://WebServ1.contoso.com/NonClaims/**.
   - In the **External certificate** list, select a certificate whose subject covers the external address.
   - In the **Backend server URL** box, the value is entered automatically and is the same as the external URL.
   - In the **Backend server SPN** box, enter the service principal name for this application; for example, **HTTP/WEBSERV1.contoso.com**.
8. On the **Confirmation** page, click **Publish**.

9. On the **Results** page, click **Close**.

Test accessing the application

After publishing the application through Web Application Proxy, you can test accessing the application from the simulated Internet.

To test accessing the application

1. Connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
2. Open an Internet explorer window and go to the non-claims-based web application; for example, https://webserv1.contoso.com/nonclaims/.
3. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   The default IIS website appears.

Step 4: Demonstrate accessing an application using Workplace Join, MFA, and multifactor access control

In this step you will join the client device to the workplace, and configure AD FS to enforce MFA, and multifactor access control when publishing the sample application through Web Application Proxy, and demonstrate that MFA and multifactor access control are working.

Important

For Workplace Join to succeed, the client computer (Client1) must trust the SSL certificate that was used to configure AD FS in Step 2: Configure the Federation Server with Device Registration Service (ADFS1). It must also be able to validate revocation information for the certificate. If you have any issues with Workplace Join, you can view the event log on Client1.

To see the event log, open Event Viewer, expand **Applications and Services Logs**, expand **Microsoft** , expand **Windows** and click **Workplace Join**.

To Workplace Join your device and access the application

1. On the client computer, on the **Start** screen, open the **Charms** bar and then select the **Settings** charm. Click **Change PC Settings**.
2. On the **PC Settings** page, click **Network**, and then click **Workplace**.
3. In the **Enter your UserID to get workplace access or turn on device management** box, type **RobertH@contoso.com** and then click **Join**.
4. When prompted for credentials, type **roberth@contoso.com**, and enter the password. Click **OK**.

5. You should now see the message: This device has joined your workplace network.
6. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
7. On the sign in page, enter the credentials of the test user (RobertH) that you created when you set up the environment.
8. The application appears in the web browser showing the claims in your security token. Notice that the token contains both user and device claims.

To configure MFA and attempt to access the application

1. Configure MFA as described in Step 3: Configure MFA on your federation server.
2. In your lab environment, connect the client computer to the simulated Internet network and assign a static IP address; for example 131.107.0.10.
3. Open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
4. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   At this point, because of the MFA policy that you configured, the user will be prompted to undergo additional authentication.

   This demonstrates that the additional authentication factor that you configured is now required to access the application.

Tip

In this walkthrough, it is necessary to disable MFA before configuring multifactor access control.

To configure multifactor access control and attempt to access the application

1. Configure multifactor access control as described in Step 3: Configure multi-factor access control policy based on user data.
2. Disable MFA as follows:
   a. On the AD FS server, in the AD FS Management console, navigate to **Authentication Policies/Per Relying Party Trust**, and select the relying party trust that represents your sample application.
   b. In the **Actions** pane, click **Edit Custom Multi-factor Authentication**.
   c. In the **Edit Relying Party Trust for <relying_party_trust_name>** dialog box, in the **Users/Groups** list, select the test group that you previously added, click **Remove**, and then click **OK**.
3. On the client computer, open an Internet explorer window and go to the sample web application; for example, https://webserv1.contoso.com/claimapp/.
4. On the sign in page, enter the credentials of the test user that you created when you set up the environment.

   At this point, because of the access control policy that you set up in the previous steps,

an ‘access denied’ message is displayed for this user, demonstrating that the access control rule is in effect.

See also

- Set up the lab environment for AD FS in Windows Server 2012 R2
- Overview: Connect to Applications and Services from Anywhere with Web Application Proxy

Installing and Configuring Web Application Proxy for Publishing Internal Applications

This scenario describes how you can use Web Application Proxy – a new Remote Access role service in Windows Server® 2012 R2 - to provide reverse proxy functionality for corporate web applications and services. When you use Web Application Proxy with Active Directory Federation Services (AD FS), you can manage the risk of exposing your applications to the Internet by configuring features provided by AD FS, including: Workplace Join, multifactor authentication (MFA), and multifactor access control.

Web Application Proxy also functions as an AD FS proxy.

The following diagram shows the topology used in this scenario for Web Application Proxy to publish Microsoft applications and other line-of-business (LOB) applications.

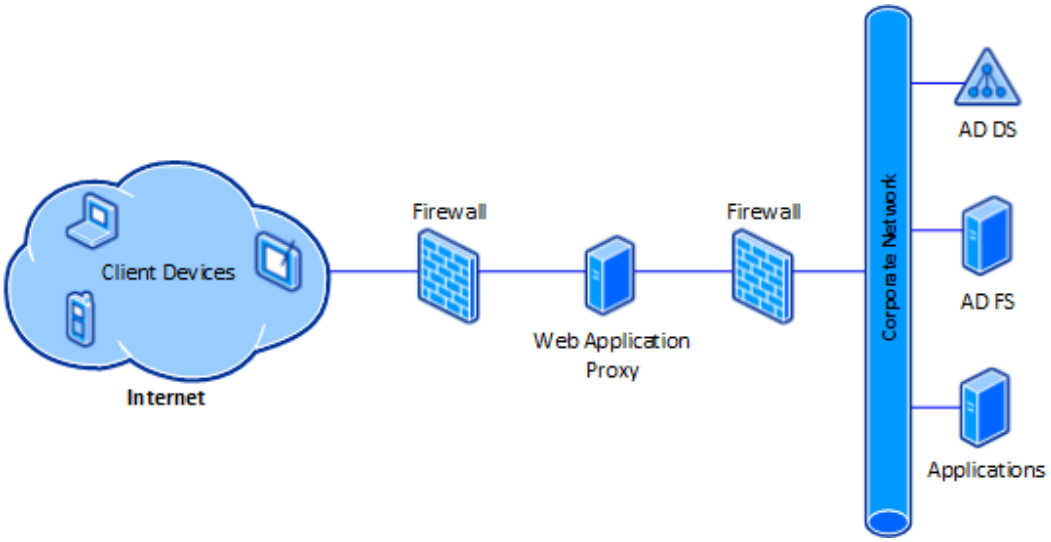

- Providing Access to Applications and Services
- Roles and Features Included in this Scenario
- Scenario Steps
- Hardware requirements
- Software requirements

- Known issues

Providing Access to Applications and Services

This scenario demonstrates how to plan and deploy Web Application Proxy in your organization to provide selective access to applications running on servers inside the organization to end users located outside of the organization. The process to make the application available externally is known as publishing. Web Application Proxy publishing enables end users to access their organization's applications from their own devices, so that users are not limited to corporate laptops to do their work, they can use their home computer, their tablet, or their smartphone. Web Application Proxy serves as a reverse proxy for any application that is published through it and as such, the end user experience is the same as if the end user's device connects directly to the application.

Roles and Features Included in this Scenario

The following table lists the roles and features that are part of this scenario and describes how they support it.

Role/feature	How it supports this scenario
Active Directory Domain Services Overview	Active Directory® Domain Services is required as a prerequisite before you can deploy AD FS. It is also required for Web Application Proxy deployments that use Kerberos constrained delegation.
Active Directory Federation Services Overview	AD FS is required to provide authentication and authorization services to Web Application Proxy and to store the Web Application Proxy configuration.
Remote Access (DirectAccess, Routing and Remote Access) Overview	Remote Access is the role containing the Web Application Proxy role service.

Scenario Steps

This scenario contains the following steps:

- Step 1: Plan the Web Application Proxy Infrastructure
- Step 2: Plan the Web Application Proxy Server
- Step 3: Plan to Publish Applications using AD FS Preauthentication
- Step 5: Plan to Publish Applications using Pass-through Preauthentication
- Step 1: Configure the Web Application Proxy Infrastructure

- Step 2: Install and Configure the Web Application Proxy Server
- Step 3: Publish Applications using AD FS Preauthentication
- Step 5: Publish Applications using Pass-through Preauthentication

Hardware requirements

Hardware requirements for this scenario include the following:

- A computer that meets the hardware requirements for Windows Server 2012 R2 running one of the following server editions: Essentials, Standard, or Datacenter.
- The server must have at least one network adapter installed, enabled, and connected to the internal network either directly, or through a firewall or NAT device. When two adapters are used, there should be one adapter connected to the internal corporate network, and one connected to the external network (Internet, or private network).

Software requirements

Software requirements for this scenario include the following:

- If the Web Application Proxy server is located behind an edge firewall or NAT device, the device must be configured to allow traffic to and from the Web Application Proxy server.
- The person deploying Web Application Proxy on the server requires local administrator permissions on the server. In addition, when you connect the Web Application Proxy server to the AD FS server, you require the credentials of the local administrator on the AD FS servers.
- You must deploy AD FS on a server running Windows Server 2012 R2 in your organization before you can deploy Web Application Proxy.
- If you want to remotely manage Web Application Proxy servers, you must enable remote PowerShell management on the Web Application Proxy servers. See Running Remote Commands.

Known issues

The following are known issues when configuring Web Application Proxy:

- When the token-signing certificate is updated in the Federation Service, the public key of the certificate is not automatically updated on the Web Application Proxy servers. To resolve this issue, you must manually import the new token-signing certificate to the Web Application Proxy servers.
- When you configure the Device Registration Service on your Federation Service, you can also enable device authentication for all relying parties with the following PowerShell command:

  ```
  Set-AdfsGlobalAuthenticationPolicy -DeviceAuthenticationEnabled
  $true
  ```

  When you publish applications with Web Application Proxy, you can publish more than one application using the same host name. You can also configure each application to

authenticate the certificate presented by client devices. To validate client certificates for an application, use the following PowerShell command:

```
Set-WebApplicationProxyApplication -id <Published Application
ID> -ClientCertificateAuthenticationBindingMode
ValidateCertificate
```

Because Web Application Proxy uses the HTTP protocol stack (http.sys), this configuration setting must be identical for all applications that use the same host name.

Plan to Publish Applications through Web Application Proxy

Web Application Proxy is a new role service in Windows Server® 2012 R2. This overview provides an introduction to the planning steps required in order to deploy Web Application Proxy and to publish applications through it. This scenario also describes the available preauthentication methods, including using Active Directory Federation Services (AD FS) for authentication and authorization, which allows you to benefit from the following AD FS features, including: Workplace Join, multifactor authentication (MFA), and multifactor access control.

Important
You must deploy AD FS on a server running Windows Server 2012 R2 in your organization before you can deploy Web Application Proxy. For AD FS planning and deployment information, see Active Directory Federation Services Overview.

- Step 1: Plan the Web Application Proxy Infrastructure—This step includes planning for the infrastructure required before you can deploy Web Application Proxy in your organization. It includes planning for network and server settings, DNS settings, Active Directory, and AD FS.
- Step 2: Plan the Web Application Proxy Server—This step includes planning for the Web Application Proxy server, including planning for deploying multiple servers and planning certificate requirements.
- Step 3: Plan to Publish Applications using AD FS Preauthentication—This step includes planning for publishing applications through Web Application Proxy using AD FS preauthentication.
- Step 5: Plan to Publish Applications using Pass-through Preauthentication—This step includes planning for publishing applications through Web Application Proxy using pass-through preauthentication.

After you have completed these planning steps, see Configure Publishing Applications through Web Application Proxy.

Step 1: Plan the Web Application Proxy Infrastructure

The first step of planning for a Web Application Proxy deployment is to perform planning for the infrastructure required for the deployment. This topic describes the infrastructure planning steps.

Task	Description
1.1. Plan Network Location	Decide where to place the Web Application Proxy server and prepare your firewalls for an Web Application Proxy deployment.
1.2. Plan DNS	Plan DNS settings for the Web Application Proxy server, infrastructure servers, and application servers.
1.3. Plan Load Balancing	Decide what type of load-balancing you require for the Web Application Proxy servers and the web application servers.
1.4. Plan Active Directory	Plan your Active Directory requirements, in particular the requirements if you plan to publish applications that use Integrated Windows authentication.
1.5. Plan Active Directory Federation Services	Plan your Active Directory Federation Services (AD FS) requirements.
1.6. Plan Preauthentication	Decide what type of preauthentication you will use when publishing applications through Web Application Proxy.

1.1. Plan Network Location

Web Application Proxy can be deployed in several topologies. For example, you can deploy Web Application Proxy behind a frontend firewall to separate it from the Internet, or between two firewalls; a frontend firewall to separate it from the Internet, a backend firewall to separate it from the corporate network. In both topologies, Web Application Proxy provides a protection layer against malicious HTTP requests that originate from the Internet through the following features:

- Preauthentication—Make sure that only authenticated traffic can get into the corporate network.
- Network Isolation—Incoming web traffic cannot directly access backend servers.
- Selective Publishing—Only specific applications and paths within these applications are accessible.

- DDoS Protection—Incoming traffic arrives at Web Application Proxy before entering the corporate network. Because Web Application Proxy acts as a buffer, many DDoS attacks can be prevented from reaching the backend servers.

Plan Firewalls

Deploying Web Application Proxy behind a firewall adds network level protection and reduces the attack surface of the Web Application Proxy servers. If the Web Application Proxy server is located in front of a firewall that separates it from the corporate network, you must make sure that the firewall does not block traffic to URLs configured for the backend servers. This could be over HTTP or HTTPS and on any specified port. Note that if you deploy Web Application Proxy on a domain-joined server, it must have connectivity to the domain controller, see 1.4. Plan Active Directory.

All firewall servers must be configured to allow HTTPS traffic because the firewall servers must publish Web Application Proxy using port 443 so that Web Application Proxy in the perimeter network can access the federation server in the corporate network. Port 443 is also required for device registration using Workplace Join.

Note

All federation server communications to and from client devices also occur over HTTPS.

In the federated world of AD FS, client requests are typically made to a specific URL, for example, a federation server identifier URL such as https://fs.fabrikam.com. Because these client requests originate from the Internet, the Internet-facing firewall server must be configured to publish the federation server identifier URL for each Web Application Proxy server that is deployed in the perimeter network.

If you want to use client certificate authentication, you must also configure the firewall to allow traffic on port 49443.

1.2. Plan DNS

DNS planning requirements include the following:

- Web Application Proxy requires internal name resolution to resolve the names of backend servers, and of infrastructure servers such as the AD FS server.
- When publishing web applications via Web Application Proxy, every web application you publish requires an external URL. For clients to reach these web applications, a public DNS server must be able to resolve each external URL that you configure. Note that the external URL must resolve to the external IP address of the Web Application Proxy server, or the external IP address of a firewall or load-balancer placed in front of the Web Application Proxy server.

1.3. Plan Load Balancing

Web Application Proxy does not include integrated load-balancing functionality. If you plan to deploy multiple Web Application Proxy servers, you should consider deploying a load-balancer to

ensure that the external traffic is distributed evenly between Web Application Proxy servers. You can use any hardware or software load-balancer that supports HTTP and HTTPS, including Windows Network Load Balancing.

You can also configure a load-balancer for published web applications. That is, you can deploy a load-balancer between the Web Application Proxy servers and the published web application. You can use any hardware or software load-balancer that supports HTTP and HTTPS, including Windows Network Load Balancing.

Visual Basic

A cluster must be either completely domain joined or not domain joined.

1.4. Plan Active Directory

Web Application Proxy can be deployed without joining the server to an AD DS domain or by joining the Web Application Proxy server to a standalone domain in a perimeter network.

You can deploy Web Application Proxy with a read-only domain controller. However, if you want to deploy Web Application Proxy and DirectAccess on the same server, you cannot use a read-only domain controller.

Plan Integrated Windows authentication and Kerberos constrained delegation

When publishing applications that use Integrated Windows authentication, the Web Application Proxy server uses Kerberos constrained delegation to authenticate users to the published application.

The following diagram shows the connections from Web Application Proxy to the domain controller to get the Kerberos ticket, and the connection from Web Application Proxy to the backend server to perform Integrated Windows authentication.

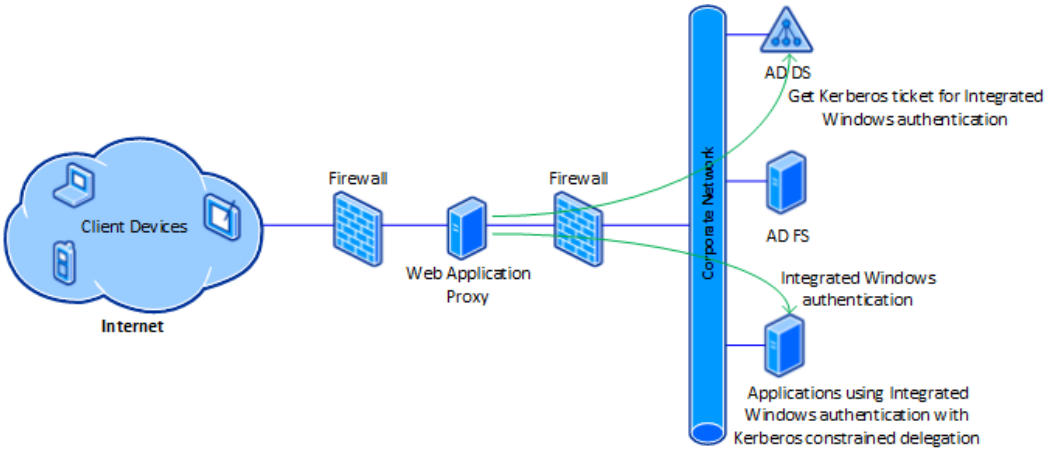

To use Integrated Windows authentication, the Web Application Proxy server must be joined to an AD DS domain. The following lists the domain and forest requirements for a deployment using Integrated Windows authentication with Kerberos constrained delegation.

1. Deployments where users, resources, and Web Application Proxy servers are all in the same forest are supported.
2. In deployments with multiple forests where there is a user forest, a resource forest, and an Web Application Proxy forest, the following deployments are supported:
    a. Users, resources, and Web Application Proxy servers are all in different forests.
    b. Users and Web Application Proxy servers are in the same forest, but resources are in a different forest.
    c. Resources and Web Application Proxy servers are in the same forest, but users are in a different forest.
    d. Users and resources are in the same forest, but Web Application Proxy servers are in a different forest.

In multi-forest deployments:

1. The user forest must trust the Web Application Proxy forest, and the Web Application Proxy forest must trust the resource forest.
2. All of the Active Directory domains in a multi-forest deployment must have at least one Windows Server® 2012 or higher domain controller. For more information, see Kerberos Constrained Delegation across Domains

1.5. Plan Active Directory Federation Services

Web Application Proxy can use AD FS for preauthentication; that is, when you configure web applications to use AD FS preauthentication, unauthenticated client requests are redirected to the AD FS server for authentication and authorization before forwarding the request to the published web application.

AD FS Requirements

When planning AD FS for Web Application Proxy deployments, the following is required:

- At least one AD FS server installed on the Windows Server 2012 R2 operating system.

1.6. Plan Preauthentication

Web Application Proxy enables you to use two forms of preauthentication:

- AD FS—When using AD FS preauthentication, you can also use Integrated Windows authentication. See Step 3: Plan to Publish Applications using AD FS Preauthentication.
- Pass-through preauthentication—When using pass-through preauthentication, Web Application Proxy does not authenticate users before they are allowed to connect to published web applications. See Step 5: Plan to Publish Applications using Pass-through Preauthentication.

See also

- Step 2: Plan the Web Application Proxy Server
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 2: Plan the Web Application Proxy Server

The next step of planning for an Web Application Proxy deployment is to perform planning for the Web Application Proxy server.

Task	Description
2.1. Plan the Web Application Proxy Role Service Installation	Plan the additional server roles that will be installed on the Web Application Proxy server.
2.2. Plan Multiple Servers	Plan adding additional servers to your deployment to create a multiple server Web Application Proxy deployment.
2.3. Plan Certificates	Plan the location, names, and issuers of certificates required by Web Application Proxy.
2.4. Plan NTP	Plan the time synchronization for claims in your deployment.

2.1. Plan the Web Application Proxy Role Service Installation

The following table describes the supported Remote Access role service deployments.

DirectAccess	VPN	Web Application Proxy
Single server deployment	Single server deployment	Single server deployment
Multisite deployment	Multiple server deployment	Not supported on the same server
Not supported on the same server	Multiple server deployment	Multiple server deployment

DirectAccess	VPN	Web Application Proxy
Cluster deployment[1]	Multiple server deployment	Multiple server deployment[2]

Notes

1—In a pre-existing DirectAccess cluster deployment, you can install Web Application Proxy only using Windows PowerShell.

2—In a pre-existing multiple server Web Application Proxy deployment, you can install DirectAccess only using Windows PowerShell.

Important

If you deploy DirectAccess and Web Application Proxy on the same server, you cannot use a read-only domain controller.

You can deploy the Web Application Proxy role service on a server that is also running the Internet Information Services (IIS) role. However, in this type of deployment, you must make sure that IIS is configured to only listen, or be bound, to URLs that are not configured as external URLs on Web Application Proxy.

2.2. Plan Multiple Servers

The Web Application Proxy configuration is stored on the AD FS servers in your organization. After configuring the first Web Application Proxy server, you can install additional Web Application Proxy servers to create a multiple server deployment. When you install the role service on the new server in the multiple server deployment, the configuration is automatically transferred to the new server after completing the Web Application Proxy Configuration Wizard.

2.3. Plan Certificates

The following table describes the certificates that are required when deploying Web Application Proxy, and any other requirements when using those certificates.

Certificate purpose and location[1]	Certificate issuer[2]	Notes
Server authentication for the Web Application Proxy server. Import the certificate to the Personal Certificates store on all Web Application Proxy servers.	Public CA Enterprise CA (internal)	• For clients to be able to connect to published web applications using HTTPS, Web Application Proxy must present a certificate that is trusted by clients. Because clients are not required to be included in your PKI, this usually requires you to acquire a certificate from an external certification authority (CA). • It is not recommended to use an

Certificate purpose and location[1]	Certificate issuer[2]	Notes
		Enterprise CA if users are allowed to use their own devices to access published applications. • You can use single name certificates, subject alternative name (SAN) certificates, or wildcard certificates. • You may require more than one certificate to successfully publish all the required applications.

Server authentication for the federation server.
Import the certificate to the Personal Certificates store on all Web Application Proxy servers.
	Public CA
	Enterprise CA (internal)
		This certificate is required for AD FS proxy functionality.
		• The subject of this certificate must match the Federation Service name value that is specified in the AD FS Management console. To locate this value, open the AD FS Management console, click **Service**, in the **Actions** pane, click **Edit Federation Service Properties**, and then find the value in the **Federation Service name** text box.
		• For Workplace Join, the certificate must have the following subject alternative names (SANs):
		  • <federation service name>.<domain>
		  For example, adfs1.contoso.com.
		  • enterpriseregistration.<domain>
		  For example, enterpriseregistration.contoso.com.
		Note
		Only one certificate is required for server authentication for the federation server.

Notes

1—If any certificate that you use has certificate revocation lists (CRLs), the server with the configured certificate must be able to contact the server that distributes the CRLs; the CRL distribution point (CDP). Clients must be able to reach the CDP. The type of CRL determines what ports are used.

2—In all cases, the client must trust the issuing CA and any intermediate CAs.

2.4. Plan NTP

When using AD FS preauthentication, the time of all Web Application Proxy servers must be identical to the time of the AD FS servers so that the timestamps on claims match. The time of all Web Application Proxy servers must be identical to the time of the applications servers when using Kerberos constrained delegation. It is recommended to enable Network Time Protocol (NTP) on all Web Application Proxy and AD FS servers.

See also

- Step 3: Plan to Publish Applications using AD FS Preauthentication
- Step 1: Plan the Web Application Proxy Infrastructure
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 3: Plan to Publish Applications using AD FS Preauthentication

This topic describes the general preauthentication flow when using Active Directory Federation Services (AD FS) preauthentication and the planning tasks for publishing applications through Web Application Proxy using AD FS preauthentication, including the preauthentication flow for different clients and applications.

For all types of application that you can publish using AD FS preauthentication, you must add a AD FS relying party trust to the Federation Service.

The general AD FS preauthentication flow is as follows:

Note
This authentication flow is not applicable for clients that use Windows Store apps.

1. The client device attempts to access a published web application on a particular resource URL; for example https://app1.contoso.com/.

   The resource URL is a public address on which Web Application Proxy listens for incoming HTTPS requests.

2. Web Application Proxy redirects the HTTPS request to the AD FS server with URL encoded parameters, including the resource URL and the appRealm (a relying party identifier).

   The user authenticates using the authentication method required by the AD FS server; for example, user name and password, two-factor authentication with a one-time password, and so on.

3. After the user is authenticated, the AD FS server issues a security token, the 'edge token', containing the following information and redirects the HTTPS request back to the Web Application Proxy server:
   - The resource identifier that the user attempted to access.
   - The user's identity as a user principal name (UPN).
   - The expiry of the access grant approval; that is, the user is granted access for a limited period of time, after which they are required to authenticate again.
   - Signature of the information in the edge token.
4. Web Application Proxy receives the redirected HTTPS request from the AD FS server with the edge token and validates and uses the token as follows:
   - Validates that the edge token signature is from the federation service that is configured in the Web Application Proxy configuration.
   - Validates that the token was issued for the correct application.
   - Validates that the token has not expired.
   - Uses the user identity when required; for example to obtain a Kerberos ticket if the backend server is configured to use Integrated Windows authentication.
5. If the edge token is valid, Web Application Proxy forwards the HTTPS request to the published web application using either HTTP or HTTPS.
6. The client now has access to the published web application; however, the published application may be configured to require the user to perform additional authentication. If, for example, the published web application is a SharePoint site and does not require additional authentication, the user will see the SharePoint site in the browser.
7. Web Application Proxy saves a cookie on the client device. The cookie is used by Web Application Proxy to identify that this session has already been preauthenticated and that no further preauthentication is required.

Note

Web Application Proxy does not support wildcard domain publishing. That is, you cannot configure an external URL using a wildcard; for example, https://*.contoso.com.

Task	Description
3.1. Plan Claims-Based Applications	Plan any prerequisites that are required before you can publish an application that uses claims-based authentication.
3.2. Plan Integrated Windows Authentication-Based Applications	Plan any prerequisites that are required before you can publish an application that uses Integrated Windows authentication.
3.3. Plan Applications for MSOFBA Clients	Plan any prerequisites that are required before you can publish an application for clients that use Microsoft Office Forms Based Authentication (MSOFBA).

Task	Description
3.4. Plan Applications for Clients that use Windows Store Apps	Plan any prerequisites that are required before you can publish an application for clients that use Windows Store apps.

3.1. Plan Claims-Based Applications

To publish an application that uses claims for authentication, you must add a relying party trust for the application to the Federation Service.

When publishing claims-based applications and accessing the application from a browser, the general authentication flow is as follows:

1. The client attempts to access a claims-based application using a web browser; for example, https://appserver.contoso.com/claimapp/.
2. The web browser sends an HTTPS request to the Web Application Proxy server which redirects the request to the AD FS server.
3. The AD FS server authenticates the user and the device and redirects the request back to Web Application Proxy. The request now contains the edge token. The AD FS server adds a single sign-on (SSO) cookie to the request because the user has already performed authentication against the AD FS server.
4. Web Application Proxy validates the token, adds its own cookie, and forwards the request to the backend server.
5. The backend server redirects the request to the AD FS server to get the application security token.
6. The request is redirected to the backend server by the ADFS server. The request now contains the application token and the SSO cookie. The user is granted access to the application and is not required to enter a user name or password.

3.2. Plan Integrated Windows Authentication-Based Applications

Web Application Proxy can be used to publish applications that uses Integrated Windows authentication; that is, Web Application Proxy performs preauthentication as required, and can then perform SSO to the published application that uses Integrated Windows authentication. To publish an application that uses Integrated Windows authentication you must add a non-claims-aware relying party trust for the application to the Federation Service.

To allow Web Application Proxy to perform single sign-on (SSO) and to perform credentials delegation using Kerberos constrained delegation, the Web Application Proxy server must be joined to a domain. See 1.4. Plan Active Directory.

To allow users to access applications that use Integrated Windows authentication, the Web Application Proxy server must be able to provide delegation for users to the published application.

You can do this on the domain controller for any application. You can also do this on the backend server if it is running on Windows Server 2012 R2 or Windows Server 2012. For more information, see What's New in Kerberos Authentication.

For a walkthrough of how to configure Web Application Proxy to publish an application using Integrated Windows authentication, see Step 3: Configure and test accessing a website using Integrated Windows authentication.

When using Integrated Windows authentication to backend servers, the authentication between Web Application Proxy and the published application is not claims-based, instead it uses Kerberos constrained delegation to authenticate end users to the application. The general flow is described below:

1. The client attempts to access a non-claims-based application using a web browser; for example, https://appserver.contoso.com/nonclaimapp/.
2. The web browser sends an HTTPS request to the Web Application Proxy server which redirects the request to the AD FS server.
3. The AD FS server authenticates the user and redirects the request back to Web Application Proxy. The request now contains the edge token.
4. Web Application Proxy validates the token.
5. If the token is valid, Web Application Proxy obtains a Kerberos ticket from the domain controller on behalf of the user.
6. Web Application Proxy adds the Kerberos ticket to the request as part of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) token, and forwards the request to the backend server. The request contains the Kerberos ticket; therefore, the user is granted access to the application without the need for further authentication.

3.3. Plan Applications for MSOFBA Clients

Web Application Proxy supports access from Microsoft Office clients such as Microsoft Word that access documents and data on backend servers. The only difference between these applications and a standard browser is that the redirection to the STS is done not via regular HTTP redirection but with special MSOFBA headers as specified in here: http://msdn.microsoft.com/en-us/library/dd773463(v=office.12).aspx. Backend application may be claims or IWA.
To publish an application for clients that use MSOFBA, you must add a relying party trust for the application to the Federation Service. Depending on the application, you can use claims-based authentication or Integrated Windows authentication. Therefore, you must add the relevant relying party trust depending on the application.

To allow Web Application Proxy to perform single sign-on (SSO) and to perform credentials delegation using Kerberos constrained delegation, the Web Application Proxy server must be joined to a domain. See 1.4. Plan Active Directory.

There are no additional planning steps if the application uses claims-based authentication. If the application used Integrated Windows authentication, see 3.2. Plan Integrated Windows Authentication-Based Applications.

The authentication flow for clients that use the MSOFBA protocol using claims-based authentication is described below. The authentication for this scenario can either use the application token in the URL, or in the body.

1. The user is working in an Office program, and from the **Recent Documents** list, opens a file on a SharePoint site.
2. The Office program shows a window with a browser control that requires the user to enter credentials.

   Note
   In some cases, the window may not appear because the client is already authenticated.
3. Web Application Proxy redirects the request to the AD FS server, which performs the authentication.
4. The AD FS server redirects the request back to Web Application Proxy. The request now contains the edge token.
5. The AD FS server adds a single sign-on (SSO) cookie to the request because the user has already performed authentication against the AD FS server.
6. Web Application Proxy validates the token and forwards the request to the backend server.
7. The backend server redirects the request to the AD FS server to get the application security token.
8. The request is redirected to the backend server. The request now contains the application token and the SSO cookie. The user is granted access to the SharePoint site and is not required to enter a user name or password to view the file.

3.4. Plan Applications for Clients that use Windows Store Apps

To publish an application for Windows Store apps, you must add a relying party trust for the application to the Federation Service.

To allow Web Application Proxy to perform single sign-on (SSO) and to perform credentials delegation using Kerberos constrained delegation, the Web Application Proxy server must be joined to a domain. See 1.4. Plan Active Directory.

Note
Web Application Proxy supports publishing only for Windows Store apps that use the OAuth 2.0 protocol.

In the AD FS Management console, you must make sure that the OAuth endpoint is proxy enabled. To check if the OAuth endpoint is proxy enabled, open the AD FS Management console, expand **Service**, click **Endpoints**, in the **Endpoints** list, locate the OAuth endpoint and make sure that the value in the **Proxy Enabled** column is **Yes**.

The authentication flow for clients that use Windows Store apps is described below:

Notes

The Web Application Proxy redirects to the AD FS server for authentication. Because Windows Store apps do not support redirects, if you use Windows Store apps, it is necessary to set the URL of the AD FS server using the Set-WebApplicationProxyConfiguration cmdlet and the OAuthAuthenticationURL parameter.

Windows Store apps can be published only using Windows PowerShell.

1. The client attempts to access a published web application using a Windows Store app.
2. The app sends an HTTPS request to the URL published by Web Application Proxy.
3. Web Application Proxy returns an HTTP 401 response to the app containing the URL of the authenticating AD FS server. This process is known as ‘discovery’.

   Note

   If the app knows the URL of the authenticating AD FS server and already has a combo token containing the OAuth token and the edge token, steps 2 and 3 are skipped in this authentication flow.
4. The app sends an HTTPS request to the AD FS server.
5. The app uses the web authentication broker to generate a dialog box in which the user enters credentials to authenticate to the AD FS server. For information about web authentication broker, see Web authentication broker.
6. After successful authentication, the AD FS server creates a combo token that contains the OAuth token and the edge token and sends the token to the app.
7. The app sends an HTTPS request containing the combo token to the URL published by Web Application Proxy.
8. Web Application Proxy splits the combo token into its two parts and validates the edge token.
9. If the edge token is valid, Web Application Proxy forwards the request to the backend server with only the OAuth token. The user is granted access to the published web application.

See also

- Step 2: Plan the Web Application Proxy Server
- Step 5: Plan to Publish Applications using Pass-through Preauthentication
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 4: Plan to Publish Applications using Client Certificate Preauthentication

This topic describes the preauthentication flow when using client certificate preauthentication and the planning tasks for publishing applications through Web Application Proxy using client certificate preauthentication.

Client certificate preauthentication is used in scenarios where a small group of external servers connect to an on-premises server through Web Application Proxy. Client certificate preauthentication should not be used for client devices connecting to published applications.

The general client certificate preauthentication flow is as follows:

1. An external server makes a request to the published resource URL.

   The resource URL is a public address on which Web Application Proxy listens for new HTTPS requests.
2. Web Application Proxy authenticates the external server using the certificate thumbprint of a certificate configured on the external server.
3. After the external server is authenticated, Web Application Proxy forwards the HTTPS request to the corporate network to the published web application using either HTTP or HTTPS.
4. The published web application responds to the request from the external server.

Note
Web Application Proxy does not support wildcard domain publishing. That is, you cannot configure an external URL using a wildcard; for example, https://*.contoso.com.

Task	Description
4.1. Plan the External Servers	Decide which certificate will be used by the external servers for client certificate authentication.
4.2. Plan Applications for Client Certificate Preauthentication	Plan any prerequisites that are required before you can publish an application that uses client certificate preauthentication.

4.1. Plan the External Servers

To use client certificate preauthentication, the external servers must have a certificate issued by a public certification authority (CA). For information about requesting and configuring this certificate, see the documentation for your server. All of the external servers that you use to connect to the published web application must have the same certificate.

The certificate can be issued by any trusted certification authority, and you must be able to obtain the certificate thumbprint.

4.2. Plan Applications for Client Certificate Preauthentication

To publish an application using client certificate preauthentication, you do not require a relying party trust for the application.

Note
Applications that use client certificate preauthentication cannot leverage the additional features that AD FS provides; such as, Workplace Join, multifactor authentication (MFA), and multifactor access control.

See also

- Step 3: Plan to Publish Applications using AD FS Preauthentication
- Step 5: Plan to Publish Applications using Pass-through Preauthentication
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 5: Plan to Publish Applications using Pass-through Preauthentication

This topic describes the preauthentication flow when using pass-through preauthentication and the planning tasks for publishing applications through Web Application Proxy using pass-through preauthentication.

The general pass-through preauthentication flow is as follows:

1. The client device attempts to access a published web application on a particular resource URL.

   The resource URL is a public address on which Web Application Proxy listens for incoming HTTPS requests.
2. Web Application Proxy forwards the HTTPS request directly to the backend server using either HTTP or HTTPS.
3. If required by the backend server, the user authenticates directly to the backend server.
4. After successfully authenticating, the client now has access to the published web application.

Note

Web Application Proxy does not support wildcard domain publishing. That is, you cannot configure an external URL using a wildcard; for example, https://*.contoso.com.

5.1. Plan Applications for Pass-Through Preauthentication

No additional planning is required for applications that use pass-through preauthentication.

Note
Applications that use pass-through preauthentication cannot leverage the additional features that AD FS provides; such as, Workplace Join, multifactor authentication (MFA), and multifactor access control.

See also

- Step 3: Plan to Publish Applications using AD FS Preauthentication
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 6: Plan to Publish SharePoint Server and Exchange Server

This topic describes the planning tasks for publishing SharePoint Server or Exchange Server through Web Application Proxy.

Note
Web Application Proxy does not support wildcard domain publishing. That is, you cannot configure an external URL using a wildcard; for example, https://*.contoso.com.

Task	Description
6.1. Plan to publish SharePoint Server	Plan to publish SharePoint Server, including supported deployments.
6.2. Plan to publish Exchange Server	Plan to publish Exchange Server, including supported Exchange services and preauthentication.

6.1. Plan to publish SharePoint Server

You can publish a SharePoint site through Web Application Proxy when the SharePoint site is configured for claims-based authentication or Integrated Windows authentication. If you want to use Active Directory Federation Services (AD FS) for preauthentication, you must configure a relying party using one of the wizards.

- If the SharePoint site uses claims-based authentication, you must use the Add Relying Party Trust Wizard to configure the relying party trust for the application.
- If the SharePoint site uses Integrated Windows authentication, you must use the Add Non-Claims-Based Relying Party Trust Wizard to configure the relying party trust for the application.

  To allow users to authenticate using Integrated Windows authentication, the Web Application Proxy server must be joined to a domain. See 1.3. Plan Active Directory.

  You must configure the application to support Kerberos constrained delegation. You can do this on the domain controller for any application. You can also configure the application directly on the backend server if it is running on Windows Server 2012 R2 or Windows Server 2012. For more information, see What's New in Kerberos Authentication. You must also make sure that the Web Application Proxy servers are configured for delegation to the service principal names of the backend servers. For a walkthrough of how to configure Web Application Proxy to publish an application using Integrated Windows authentication, see Configure a site to use Integrated Windows authentication.

If your SharePoint site is configured using either alternate access mappings (AAM) or host-named site collections, you can use different external and backend server URLs to publish your application. However, if you do not configure your SharePoint site using AAM or host-named site collections, you must use the same external and backend server URLs.

6.2. Plan to publish Exchange Server

The following table describes the Exchange services that you can publish through Web Application Proxy and the supported preauthentication for these services:

Exchange service	Preauthentication	Notes
Outlook Web App	• AD FS using non-claims-based authentication • Pass-through	—
Exchange Control Panel	Pass-through	—
Outlook Anywhere	Pass-through	You must publish three URLs for Outlook Anywhere to work correctly: • The autodiscover URL. • The external host name of

Exchange service	Preauthentication	Notes
		the Exchange Server; that is, the URL that is configured for clients to connect to. • The internal FQDN of the Exchange Server.
Exchange ActiveSync	Pass-through	—

To publish Outlook Web App using Integrated Windows authentication, you must use the Add Non-Claims-Based Relying Party Trust Wizard to configure the relying party trust for the application.

To allow users to authenticate using Integrated Windows authentication, the Web Application Proxy server must be joined to a domain. See 1.3. Plan Active Directory.

You must configure the application to support Kerberos constrained delegation. You can do this on the domain controller for any application. You can also configure the application directly on the backend server if it is running on Windows Server 2012 R2 or Windows Server 2012. For more information, see What's New in Kerberos Authentication. You must also make sure that the Web Application Proxy servers are configured for delegation to the service principal names of the backend servers. For a walkthrough of how to configure Web Application Proxy to publish an application using Integrated Windows authentication, see Configure a site to use Integrated Windows authentication.

See also

- Step 5: Plan to Publish Applications using Pass-through Preauthentication
- Plan to Publish Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Configure Publishing Applications through Web Application Proxy

Web Application Proxy is a new role service in Windows Server® 2012 R2. This overview provides an introduction to the configuration steps required in order to deploy Web Application Proxy and publish applications through it.

- Step 1: Configure the Web Application Proxy Infrastructure—Describes how to configure the infrastructure required to deploy Web Application Proxy.

- Step 2: Install and Configure the Web Application Proxy Server—Describes how to configure Web Application Proxy servers, including configuring any required certificates, installing the Web Application Proxy role service, and joining Web Application Proxy servers to a domain.
- Step 3: Publish Applications using AD FS Preauthentication—Describes how to publish applications through Web Application Proxy using AD FS preauthentication.
- Step 5: Publish Applications using Pass-through Preauthentication—Describes how to publish applications using pass-through preauthentication.

For planning information you should review before deploying Web Application Proxy, see Plan to Publish Applications through Web Application Proxy.

Step 1: Configure the Web Application Proxy Infrastructure

This topic describes how to configure the infrastructure required for an Web Application Proxy deployment. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 1: Plan the Web Application Proxy Infrastructure.

Task	Description
1.1. Configure server network settings	Configure the server network settings on the Web Application Proxy server.
1.2. Configure firewalls	Configure additional firewalls, if required.
1.3. Configure DNS	Configure DNS settings for the Web Application Proxy deployment.
1.4. Configure Active Directory	Join the Web Application Proxy server to the Active Directory domain, if required.
1.5. Configure Active Directory Federation Services	Configure Active Directory Federation Services (AD FS) servers for authentication and authorization.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

1.1. Configure server network settings

Network interface settings on the Web Application Proxy server are configured by using **Change adapter settings** in the **Windows Networking and Sharing Center**. Specific IP address requirements depend on the topology and any other organizational requirements.

1.2. Configure firewalls

When the Web Application Proxy is separated from client devices or from the internal network by a firewall, you must configure the firewalls to allow traffic for the protocol and port number used by the published web applications. You must also configure the firewall to allow HTTPS traffic on port 443 for clients to communicate with the AD FS server.

Configure the backend firewall to enable connectivity to backend servers, including:

1. A rule to allow connectivity to the AD FS on port 443
2. Open the port on which you are running your apps
3. Connectivity to Active Directory

1.3. Configure DNS

Configure DNS servers internally and externally as required by your published web applications.

- Web Application Proxy requires internal name resolution to resolve the names of backend servers, and infrastructure servers such as the AD FS servers.
- When publishing web applications via Web Application Proxy, every web application you publish requires an external URL. For clients to reach these web applications, a public DNS server must be able to resolve each external URL that you configure. Note that the external URL must resolve to the same IP address as the Web Application Proxy server, or the external IP address of a firewall or load-balancer placed in front of the Web Application Proxy server.

1.4. Configure Active Directory

If your deployment uses Integrated Windows authentication, the Web Application Proxy server must be joined to an AD DS domain. For information about domain and forest trusts when deploying Web Application Proxy, see 1.4. Plan Active Directory.

Do this step using Windows PowerShell

To join the Web Application Proxy server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. On the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In **Computer Name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
5. When you see a dialog box welcoming you to the domain, click **OK**.

6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

Note that you must supply domain credentials after entering the Add-Computer command below.

```
Add-Computer -DomainName <domain_name>

Restart-Computer
```

Configure Kerberos constrained delegation

To allow users to access applications that use Integrated Windows authentication, the Web Application Proxy server must be able to provide delegation for users to the published application.

Tip

This procedure assumes that the Web Application Proxy server is named EDGE1 and the domain is contoso.com.

To configure Kerberos constrained delegation

1. On the domain controller, open Server Manager. To do this, click **Server Manager** on the **Start** screen, or **Server Manager** in the taskbar on the desktop.
2. Click **Tools**, and then click **ADSI Edit**.
3. On the **Action** menu, click **Connect To**, and then on the **Connection Settings** dialog box, accept the default settings to connect to the default naming context, and then click **OK**.
4. In the left pane, expand **Default naming context**, expand **DC=contoso,DC=com**, expand **CN=Computers**, right-click **CN=EDGE1**, and then click **Properties**.
5. On the **CN=EDGE1 Properties** dialog box, on the **Attribute Editor** tab, in the **Attributes** list, select **servicePrincipalName**, and then click **Edit**.
6. On the **Multi-valued String Editor** dialog box, in **Value to add**, enter **HTTP/EDGE1.contoso.com** and click **Add**. Then enter **HTTP/EDGE1** and click **Add**.

   The **Values** list now contains two new entries; for example, **HTTP/EDGE1.contoso.com** and **HTTP/EDGE1**.
7. On the **Multi-valued String Editor** dialog box, click **OK**.
8. On the **CN=EDGE1 Properties** dialog box, click **OK**.
9. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
10. In the navigation pane, under **contoso.com**, click **Computers**. In the details pane, right-click the Web Application Proxy server, and then click **Properties**.
11. On the **EDGE1 Properties** dialog box, on the **Delegation** tab, click **Trust this computer**

for delegation to specified services only, and then click **Use any authentication protocol**.

12. Click **Add**, and on the **Add Services** dialog box, click **Users or Computers**.
13. On the **Select Users or Computers** dialog box, in **Enter the object names to select**, enter the name of the web servers that use Integrated Windows authentication; for example, WebServ1, and then click **OK**.
14. On the **Add Services** dialog box, in the **Available services** list, select the **http** service type, and then click **OK**.
15. On the **EDGE1 Properties** dialog box, click **OK**.

1.5. Configure Active Directory Federation Services

Configuring AD FS is described in detail in the Windows Server 2012 AD FS Deployment Guide.

For all web applications that you intend to publish using AD FS preauthentication, you must configure a relying party trust for each application on the AD FS server. See Create a Relying Party Trust Using Federation Metadata.

For all AD FS endpoints that you require to be published by Web Application Proxy, in the AD FS Management console, you must set the endpoint to be **Proxy Enabled**.

See also

- Step 2: Install and Configure the Web Application Proxy Server
- Configure Publishing Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview

Step 2: Install and Configure the Web Application Proxy Server

This topic describes how to install the Remote Access role with the Web Application Proxy role service and how to configure the Web Application Proxy server to connect to an Active Directory Federation Services (AD FS) server. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 2: Plan the Web Application Proxy Server.

Task	Description
2.1. Configure CAs and certificates	Configure any certificate templates and certificates that might be required in your Web Application Proxy deployment.

Task	Description
2.2. Install the Remote Access role	Install the Remote Access role with the Web Application Proxy role service.
2.3. Configure Web Application Proxy	Configure the Web Application Proxy server to connect to an AD FS server.

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

2.1. Configure CAs and certificates

Web Application Proxy servers require the following certificates in the certificate store on each Web Application Proxy server:

- A certificate whose subject covers the federation service name. If you want to use Workplace Join, the certificate must also contain the following subject alternative names (SANs): <federation service name>.<domain> and enterpriseregistration.<domain>.
- A wildcard certificate, a subject alternative name (SAN) certificate, several SAN certificates, or several certificates whose subjects cover each web application.
- A copy of the certificate issued to external servers when using client certificate preauthentication.

Configure certificate templates

Depending on your deployment and authentication requirements, you might require additional certificate templates on your internal certification authority (CA).

To configure a certificate template

1. On the internal CA, create a certificate template as described in Creating Certificate Templates.
2. Deploy the certificate template as described in Deploying Certificate Templates.

Configure web application certificates

In an Web Application Proxy deployment you require certificates for the published web applications, and for the AD FS proxy if your deployment provides AD FS proxy functionality. For these required certificates, there are two options for the issuing CA:

- **Public**—Supplied by a 3rd party.

  A website certificate used for server authentication. If the certificate subject is not a wildcard, it must be the externally resolvable fully qualified domain name (FQDN) URL that you configure on the Web Application Proxy server for the application.

- **Private**—The following are required, if they do not already exist:
  - A website certificate used for server authentication. The certificate subject should be an externally resolvable FQDN that is reachable from the Internet. The certificate can be based on the certificate template created in Configure certificate templates.
  - A certificate revocation list (CRL) distribution point that is reachable from a publicly resolvable FQDN.

Make sure that the website certificate used for server authentication meets the following requirements:

- The common name of the certificate should match the name that you configure for the external URL of the published web application, or the federation service name.
- For the Enhanced Key Usage field, use the Server Authentication object identifier (OID).
- For the CRL Distribution Points field, specify a CRL distribution point that is accessible by client devices that are connected to the Internet.
- The certificate must have a private key.
- The certificate must be imported directly into the personal store.
- Certificates can have wildcards in the name. A wildcard certificate with the subject name *.contoso.com can be used for web applications in the domain contoso.com, for example, sharepoint.contoso.com and owa.contoso.com. This wildcard certificate cannot be used for the website sharepoint.internal.contoso.com.
- Certificates can be subject alternative name (SAN) certificates. For example, a SAN certificate with the names owa.contoso.com and crm.contoso.com can be used for only those two websites. It cannot be used for sharepoint.contoso.com.

Notes

For Workplace Join, a SAN certificate is required with the following SANs:

- <federation service name>.<domain>

  For example, adfs1.contoso.com.
- enterpriseregistration.<domain>

  For example, enterpriseregistration.contoso.com.

2.2. Install the Remote Access role

To deploy Web Application Proxy, you must install the Remote Access role with the Web Application Proxy role service on a server that will act as the Web Application Proxy server.

Repeat this procedure for all of the servers that you want to deploy as Web Application Proxy servers.

Do this step using Windows PowerShell

To install the Web Application Proxy role service

1. On the Web Application Proxy server, in the Server Manager console, in the **Dashboard**, click **Add roles and features**.

2. In the **Add Roles and Features Wizard**, click **Next** three times to get to the server role selection screen.
3. On the **Select server roles** dialog, select **Remote Access**, and then click **Next**.
4. Click **Next** twice.
5. On the **Select role services** dialog, select **Web Application Proxy**, click **Add Features**, and then click **Next**.
6. On the **Confirm installation selections** dialog, click **Install**.
7. On the **Installation progress** dialog, verify that the installation was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Install-WindowsFeature Web-Application-Proxy -IncludeManagementTools
```

See the instructional video for help Installing the Web Application Proxy

2.3. Configure Web Application Proxy

You must configure Web Application Proxy to connect to an AD FS server.

Repeat this procedure for all of the servers that you want to deploy as Web Application Proxy servers.

Do this step using Windows PowerShell

To configure Web Application Proxy

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the navigation pane, click **Web Application Proxy**.
3. In the Remote Access Management console, in the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
4. On the **Web Application Proxy Configuration Wizard**, on the **Welcome** dialog, click **Next**.
5. On the **Federation Server** dialog, do the following, and then click **Next**:
   - In the **Federation service name** box, enter the fully qualified domain name (FQDN) of the AD FS server; for example, fs.contoso.com.
   - In the **User name** and **Password** boxes, enter the credentials of a local administrator account on the AD FS servers.
6. On the **AD FS Proxy Certificate** dialog, in the list of certificates currently installed on the Web Application Proxy server, select a certificate to be used by Web Application Proxy

for AD FS proxy functionality, and then click **Next**.

The certificate you choose here should be the one that whose subject is the Federation Service name, for example, fs.contoso.com. If you plan on using Workplace Join, this must be a SAN certificate with the SANs described in 2.1. Configure CAs and certificates.

7. On the **Confirmation** dialog, review the settings. If required, you can copy the PowerShell cmdlet to automate additional installations. Click **Configure**.
8. On the **Results** dialog, verify that the configuration was successful, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

The following command will prompt you to enter credentials of a local administrator account on the AD FS servers.

```
Install-WebApplicationProxy -CertificateThumprint
'1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b' -FederationServiceName fs.contoso.com
```

See also

- Step 3: Publish Applications using AD FS Preauthentication
- Step 1: Configure the Web Application Proxy Infrastructure
- Configure Publishing Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview
- Web Application Proxy Cmdlets in Windows PowerShell
- Install-WebApplicationProxy

Step 3: Publish Applications using AD FS Preauthentication

This topic describes how to publish applications through Web Application Proxy using Active Directory Federation Services (AD FS) preauthentication. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 3: Plan to Publish Applications using AD FS Preauthentication.

Important

When configuring the external URL and the backend server URL, make sure you include the fully qualified domain name (FQDN), and not an IP address.

Task	Description
3.1. Publish a Claims-based Application for Web Browser Clients	Publish an application that uses claims-based authentication that will be available to web browser clients.
3.2. Publish an Integrated Windows authenticated-based Application for Web Browser Clients	Publish an application that uses Integrated Windows authentication that will be available to web browser clients.
3.3. Publish an Application that uses MSOFBA	Publish an application that will be available to clients using Microsoft Office Forms Based Authentication (MSOFBA).
3.4. Publish an Application for a Windows Store App	Publish an application that will be available to Windows Store apps.

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

3.1. Publish a Claims-based Application for Web Browser Clients

This procedure describes how to publish a claims-based application, such as a SharePoint site, that will be accessed by web browser clients. Before you begin, make sure that you have done the following:

- Created a relying party trust for the application in the AD FS Management console.
- Verified that a certificate on the Web Application Proxy server is suitable for the application you want to publish.

Do this step using Windows PowerShell

To publish a claims-based application

1. On the Web Application Proxy server, in the Remote Access Management console, in the **Navigation** pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Active Directory Federation Services (AD FS)**, and then click **Next**.
4. On the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:

- In the **Name** box, enter a friendly name for the application.

  This name is used only in the list of published applications in the Remote Access Management console.
- In the **External URL** box, enter the external URL for this application; for example, https://sp.contoso.com/app1/.
- In the **External certificate** list, select a certificate whose subject covers the external URL.
- In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://sp/app1/.

  Note
  Web Application Proxy can translate host names in URLs, but cannot translate path names. Therefore, you can enter different host names, but you must enter the same path name. For example, you can enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of http://app-server/app1/. However, you cannot enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://apps.contoso.com/internal-app1/.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command to set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WebApplicationProxyApplication

    -BackendServerURL 'https://sp.contoso.com/app1/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://sp.contoso.com/app1/'

    -Name 'SP'

    -ExternalPreAuthentication ADFS

    -ADFSRelyingPartyName 'SP_Relying_Party'
```

3.2. Publish an Integrated Windows authenticated-based Application for Web Browser Clients

This procedure describes how to publish an application that uses Integrated Windows authentication, such as Outlook Web App, that will be accessed by web browser clients. Before you begin, make sure that you have done the following:

- Created a non-claims-aware relying party trust for the application in the AD FS Management console.
- Configured the backend server to support Kerberos constrained delegation on the domain controller or by using the Set-ADUser cmdlet with the -PrincipalsAllowedToDelegateToAccount parameter. Note that if the backend server is running on Windows Server 2012 R2 or Windows Server 2012, you can also run this PowerShell command on the backend server.
- Made sure that the Web Application Proxy servers are configured for delegation to the service principal names of the backend servers.
- Verified that a certificate on the Web Application Proxy server is suitable for the application you want to publish.

Do this step using Windows PowerShell

To publish a non-claims-based application

1. On the Web Application Proxy server, in the Remote Access Management console, in the **Navigation** pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Active Directory Federation Services (AD FS)**, and then click **Next**.
4. On the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.

     This name is used only in the list of published applications in the Remote Access Management console.
   - In the **External URL** box, enter the external URL for this application; for example, https://owa.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://owa/.

     Note
     Web Application Proxy can translate host names in URLs, but cannot

translate path names. Therefore, you can enter different host names, but you must enter the same path name. For example, you can enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of http://app-server/app1/. However, you cannot enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://apps.contoso.com/internal-app1/.

- In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/owa.contoso.com.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command to set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WebApplicationProxyApplication

    -BackendServerAuthenticationSpn 'HTTP/owa.contoso.com'

    -BackendServerURL 'https://owa.contoso.com/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://owa.contoso.com/'

    -Name 'OWA'

    -ExternalPreAuthentication ADFS

    -ADFSRelyingPartyName 'Non-Claims_Relying_Party'
```

3.3. Publish an Application that uses MSOFBA

The steps to publish an application that uses MSOFBA are identical to the steps for a claims-based application or a non-claims-based application. For claims-based applications, see 3.1. Publish a Claims-based Application for Web Browser Clients, for non-claims-based applications, see 3.2. Publish an Integrated Windows authenticated-based Application for Web Browser Clients. Web Application Proxy automatically detects the client and will authenticate the user as required.

3.4. Publish an Application for a Windows Store App

This procedure describes how to publish an application for a Windows Store app. This type of application can be published only using Windows PowerShell. Before you begin, make sure that you have done the following:

- Created a relying party trust for the application in the AD FS Management console.
- Made sure that the OAuth endpoint is proxy enabled in the AD FS Management console and make a note of the URL Path.
- Verified that a certificate on the Web Application Proxy server is suitable for the application you want to publish.

Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To set the OAuth authentication URL for a federation server address of fs.contoso.com and a URL path of /adfs/oauth2/:

```
Set-WebApplicationProxyConfiguration –OAuthAuthenticationURL
'https://fs.contoso.com/adfs/oauth2/'
```

To publish the application:

```
Add-WebApplicationProxyApplication
    -BackendServerURL 'https://storeapp.contoso.com/'
    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'
    -ExternalURL 'https://storeapp.contoso.com/'
    -Name 'Windows Store app Server'
    -ExternalPreAuthentication ADFS
    -ADFSRelyingPartyName 'Store_app_Relying_Party'
    -UseOAuthAuthentication
```

See also

- Step 5: Publish Applications using Pass-through Preauthentication
- Step 2: Install and Configure the Web Application Proxy Server
- Configure Publishing Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview
- Web Application Proxy Cmdlets in Windows PowerShell
- Add-WebApplicationProxyApplication
- Set-WebApplicationProxyConfiguration

Step 4: Publish Applications using Client Certificate Preauthentication

This topic describes how to publish applications through Web Application Proxy using client certificate preauthentication. Before beginning the deployment steps, ensure that you have

completed the planning steps described in Step 4: Plan to Publish Applications using Client Certificate Preauthentication.

Important
When configuring the external URL and the backend server URL, you must enter a URL. You cannot enter an IP address.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

4.1. Publish an Application using Client Certificate Preauthentication

This procedure describes how to publish an application using client certificate authentication. This type of application can be published only using Windows PowerShell. Before you begin, make sure that you have done the following:

- Import the website certificate from the backend server, or verify that a certificate on the Web Application Proxy server is suitable for the application you want to publish.
- Import the certificate from the external server.

Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To publish the application:

```
Add-WebApplicationProxyApplication

    -BackendServerURL 'https://clientcertapp.contoso.com/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://clientcertapp.contoso.com/'

    -Name 'Client certificate preauthentication application'

    -ExternalPreAuthentication ClientCertificate

    -ClientCertificatePreauthenticationThumbprint
'123456abcdef123456abcdef123456abcdef12ab'
```

See also

- Step 5: Publish Applications using Pass-through Preauthentication
- Step 3: Publish Applications using AD FS Preauthentication
- Configure Publishing Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview
- Web Application Proxy Cmdlets in Windows PowerShell

- Add-WebApplicationProxyApplication

Step 5: Publish Applications using Pass-through Preauthentication

This topic describes how to publish applications through Web Application Proxy using pass-through preauthentication. Before beginning the deployment steps, ensure that you have completed the planning steps described in Step 5: Plan to Publish Applications using Pass-through Preauthentication.

Important
When configuring the external URL and the backend server URL, you must enter a URL including the fully qualified domain name (FQDN), and not an IP address.

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

5.1. Publish an Application using Pass-Through Preauthentication

This procedure describes how to publish an application that uses pass-through preauthentication. Before you begin, make sure that you have done the following:

- Import the website certificate from the backend server, or verify that a certificate on the Web Application Proxy server is suitable for the application you want to publish.

Do this step using Windows PowerShell

To publish an application using pass-through preauthentication

1. On the Web Application Proxy server, in the Remote Access Management console, in the **Navigation** pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, click **Pass-through**, and then click **Next**.
4. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.

     This name is used only in the list of published applications in the Remote Access Management console.
   - In the **External URL** box, enter the external URL for this application; for example, https://maps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external

URL.

- In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://maps/.

  Note
  Web Application Proxy can translate host names in URLs, but cannot translate path names. Therefore, you can enter different host names, but you must enter the same path name. For example, you can enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of http://app-server/app1/. However, you cannot enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://apps.contoso.com/internal-app1/.

5. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command to set up additional published applications.
6. On the **Results** page, make sure that the application published successfully, and then click **Close**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-WebApplicationProxyApplication

    -BackendServerURL 'https://maps.contoso.com/'

    -ExternalCertificateThumbprint '1a2b3c4d5e6f1a2b3c4d5e6f1a2b3c4d5e6f1a2b'

    -ExternalURL 'https://maps.contoso.com/'

    -Name 'Maps (no preauthentication)'

    -ExternalPreAuthentication PassThrough
```

See also

- Step 3: Publish Applications using AD FS Preauthentication
- Configure Publishing Applications through Web Application Proxy
- Installing and Configuring Web Application Proxy for Publishing Internal Applications
- Web Application Proxy Overview
- Web Application Proxy Cmdlets in Windows PowerShell
- Add-WebApplicationProxyApplication

Best Practices Analyzer for Web Application Proxy

In Windows management, *best practices* are guidelines that are considered the ideal way, under normal circumstances, to configure a server as defined by experts. For example, it is considered a best practice for most server technologies to keep open only those ports required for the technologies to communicate with other networked computers, and block unused ports. While best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Topics in this section can help you bring Web Application Proxy running on Windows Server 2012 R2 into compliance with best practices. Content in this section is most valuable to administrators who have completed a Best Practices Analyzer scan of Web Application Proxy, and who want information about how to interpret and resolve scan results that identify areas of Web Application Proxy that are noncompliant with best practices.

For more information about Best Practices Analyzer and scans, see Run Best Practices Analyzer Scans and Manage Scan Results.

More information about Web Application Proxy

Web Application Proxy enables you to publish selected HTTP- and HTTPS-based applications from your corporate network to client devices outside of the corporate network. It can use Active Directory Federation Services (AD FS) to ensure that users are authenticated before they gain access to published applications. Web Application Proxy also provides proxy functionality for AD FS servers.

Topics in this section

- Web Application Proxy must be configured before it is used
- Web Application Proxy: The external and backend server URLs are different and URL translation is disabled
- Web Application Proxy: The service is not configured to run automatically
- Web Application Proxy: The AD FS Proxy service is not configured to run automatically
- Web Application Proxy: This server is not included in the ConnectedServersName list
- Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high
- Web Application Proxy: Application is using an external certificate that is not yet valid
- Web Application Proxy: Application is using an external certificate that is about to expire
- Web Application Proxy: Application is using an external certificate that has no private key
- Web Application Proxy: Application is using an external certificate that has expired
- Web Application Proxy: Application is configured to use an external certificate that is not present on this server

- Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows authentication but the server is not joined to a domain
- Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed

Web Application Proxy must be configured before it is used

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy was installed but not configured for initial use.

Impact

If Web Application Proxy is not configured it cannot be used to publish applications.

Resolution

Use 'Remote Access Management' in Server Manager to start the configuration wizard or use the Install-WebApplicationProxy PowerShell command.

After installing the Web Application Proxy role service on a server, you must also configure Web Application Proxy on the server. If you do not configure Web Application Proxy, you will be unable to publish applications and the Web Application Proxy server will not provide AD FS proxy functionality.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To configure Web Application Proxy

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the navigation pane, click **Web Application Proxy**.
3. In the Remote Access Management console, in the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
4. On the **Web Application Proxy Configuration Wizard**, on the **Welcome** dialog, click **Next**.
5. On the **Federation Server** dialog, do the following, and then click **Next**:
   - In the **Federation service name** box, enter the fully qualified domain name (FQDN) of the AD FS server; for example, fs.contoso.com.
   - In the **User name** and **Password** boxes, enter the credentials of a local administrator account on the AD FS servers.
6. On the **AD FS Proxy Certificate** dialog, in the list of certificates currently installed on the Web Application Proxy server, select a certificate to be used by Web Application Proxy for AD FS proxy functionality, and then click **Next**.
7. On the **Confirmation** dialog, review the settings. If required, you can copy the PowerShell cmdlet to automate additional installations. Click **Configure**.
8. On the **Results** dialog, verify that the configuration was successful, and then click **Close**.

Web Application Proxy: The external and backend server URLs are different and URL translation is disabled

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

The external and backend server URLs are different and URL translation is disabled.

Impact

The published application might reject requests from the client.

Resolution

Consider changing the application publishing settings.

By default, Web Application Proxy translates the host portion of requests to a backend server. For example, Web Application Proxy will translate the URLs successfully if the external URL is https://apps.contoso.com/ and the backend server URL is https://appsinternal.contoso.com/. However, URL translation is currently disabled, which might cause client requests to be rejected. You can manually enable the translation of host headers by using the DisableTranslateUrlInRequestHeaders parameter.

You may need to remove the published application and republish it to resolve this issue.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page,

in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.

5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.

     Note
     Web Application Proxy can translate host names in URLs, but cannot translate path names. Therefore, you can enter different host names, but you must enter the same path name. For example, you can enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://app-server/app1/. However, you cannot enter an external URL of https://apps.contoso.com/app1/ and a backend server URL of https://apps.contoso.com/internal-app1/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

You can view the settings of a published application and enable URL translation, if required.

To enable URL translation using PowerShell

1. To obtain the application ID of an application, use the following PowerShell command to show the ID, name, and external URL of all applications published by Web Application Proxy, and locate the application ID of the required application in the command output:

```
Get-WebApplicationProxyApplication | Format-Table ID, Name,
ExternalURL
```

2. Use the following PowerShell command to enable URL translation for a specific published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
DisableTranslateUrlInRequestHeaders:$false
```

Web Application Proxy: The service is not configured to run automatically

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Web Application Proxy service (appproxysvc) settings were changed and it is not configured to run automatically.

Impact

After the server restarts, the service will not start and users will not be able to access published applications.

Resolution

Consider changing the Web Application Proxy service startup type to Automatic.

In order for Web Application Proxy to work correctly, the appproxysvc service must be running. It is recommended that this service is set to automatic to ensure that whenever the Web Application Proxy server restarts, the service will restart without any user action.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To set the appproxysvc service startup type to automatic

1. On the Web Application Proxy server, open the Services console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **services.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the details pane, right-click **Web Application Proxy Service**, and then click

Properties.

3. On the **Web Application Proxy Service** dialog, on the **General** tab, in **Startup type**, click **Automatic**, and then click **OK**.

Web Application Proxy: The AD FS Proxy service is not configured to run automatically

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Active Directory Federation Services service (adfssrv) settings were changed and it is not configured to run automatically.

Impact

After the server restarts, the service will not start and users will not be able to authenticate and will not be able to access published applications.

Resolution

Consider changing the Active Directory Federation Services service startup type to Automatic.

In order for Web Application Proxy to work correctly, the adfssrv service must be running. It is recommended that this service is set to automatic to ensure that whenever the Web Application Proxy server restarts, the service will restart without any user action.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To set the adfssrv service startup type to automatic

1. On the Web Application Proxy server, open the Services console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **services.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the details pane, right-click **Active Directory Federation Services**, and then click **Properties**.
3. On the **Active Directory Federation Services** dialog, on the **General** tab, in **Startup type**, click **Automatic**, and then click **OK**.

Web Application Proxy: This server is not included in the ConnectedServersName list

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

This server is not included in the ConnectedServersName list.

Impact

The Remote Access Management console may not work as expected.

Resolution

Use the Set-WebApplicationProxyConfiguration cmdlet to add this server to the ConnectedServersName list.

All Web Application Proxy servers must be included in the ConnectedServersName list, otherwise they will not appear in the Remote Access Management console and attempts to change or configure published applications may result in unexpected behavior. You can manually configure the ConnectedServersName list only using Windows PowerShell.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To add a server to the ConnectedServersName list

1. On the Web Application Proxy server, open a Windows PowerShell window.
2. Use the following command to view the current ConnectedServersName list:

   ```
   Get-WebApplicationProxyConfiguration
   ```
3. Copy the current ConnectedServersName list to the clipboard.
4. Run the following command to add a server to the ConnectedServersName list:

   ```
   Set-WebApplicationProxyConfiguration -ConnectedServersName
   <paste-existing-list-here>,<new_server_name>.<domain_name>
   ```

   Note
   You must provide the existing list of connected servers, and then add each new server as a comma-separated list for this command to be effective.

Web Application Proxy: The ConfigurationChangesPollingIntervalSec value is high

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy

Severity	Warning
Category	Configuration

Issue

When the ConfigurationChangesPollingIntervalSec value is high, any changes that you make are not propagated in a timely manner, or changes are only partially propagated until the next time the server checks for updates.

Impact

Applications might not be published as expected until the next time the server checks for updates.

Resolution

Consider changing ConfigurationChangesPollingIntervalSec to a lower value using the Set-WebApplicationProxyConfiguration cmdlet.

Web Application Proxy servers must periodically check if there have been any configuration changes made to the Web Application Proxy configuration that is stored on the Active Directory Federation Services (AD FS) servers. The time between these checks is set by the ConfigurationChangesPollingIntervalSec parameter. If the value is low, Web Application Proxy checks for changes more frequently, but does increase internal network traffic. However, if the value is high, Web Application Proxy may operate for a significant period of time before starting to use the new configuration. The default value for this parameter is 30 seconds.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To change how frequently Web Application Proxy checks for configuration changes

1. On the Web Application Proxy server, open a Windows PowerShell window.
2. Use the following command to view the current ConfigurationChangesPollingIntervalSec value:

   ```
   Get-WebApplicationProxyConfiguration
   ```

3. Run the following command to change the ConfigurationChangesPollingIntervalSec value to 15 seconds:

   ```
   Set-WebApplicationProxyConfiguration –
   ConfigurationChangesPollingIntervalSec 15
   ```

Web Application Proxy: Application is using an external certificate that is not yet valid

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that is not yet valid.

Impact

If an application is using a certificate that is not yet valid, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Publish this application again with a valid certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication is not yet valid, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.

4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that is about to expire

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue

addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that is about to expire.

You will receive this message for applications whose certificates will expire in under 60 days.

Impact

If an application is using a certificate that has expired, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Issue a new certificate for this address and publish this application again with the new certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication expires, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you

want, and then click **Yes**.

2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.

- In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
- In the **External certificate** list, select a certificate whose subject covers the external URL.
- In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
- In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that has no private key

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy

Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that has no private key.

Impact

If an application is using a certificate that has no private key, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Publish this application again with a valid certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application has no private key, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.

8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend

server; for example, HTTP/apps.contoso.com.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is using an external certificate that has expired

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using a certificate that has expired.

Impact

If an application is using a certificate that has expired, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites.

Resolution

Issue a new certificate for this address and publish this application again with the new certificate.

When you publish an application through Web Application Proxy, a valid certificate with the private key is required to be stored in the Personal certificates store on each Web Application Proxy server. If the certificate used by an application for authentication expires, users will not be able to secure their access to the application and sensitive data will not be encrypted. Some browsers might block access to such sites. To resolve this issue, you must obtain a new certificate for this application. After obtaining the new certificate, you can either change the certificate using the Set-WebApplicationProxyApplication cmdlet with the ExternalCertificateThumbprint parameter, or remove the application and republish it using the new certificate.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.

10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

To remove a published application

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
3. In the details pane, select the application identified in the BPA, and then in the **Tasks** pane, click **Remove**.
4. On the **Remove Applications** dialog box, click **Yes**.

After removing the application, you can republish it.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:
   - In the **Name** box, enter a friendly name for the application.
   - In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
   - In the **External certificate** list, select a certificate whose subject covers the external URL.
   - In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
   - In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.
6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then

click **Close**.

To change the certificate using PowerShell

1. Obtain a new certificate for the application as described above.
2. To obtain the application ID of the required application, use the following PowerShell command to show the ID and externalURL of all applications whose name matches the application name. Locate the application ID in the command output of the relevant application:

```
Get-WebApplicationProxyApplication –Name <application_name> |
Format-Table ID, ExternalURL
```

3. Use the following PowerShell command to change the certificate used for the published application:

```
Set-WebApplicationProxyApplication –ID <application_ID> -
ExternalCertificateThumbprint <New_certificate_thumbprint>
```

Web Application Proxy: Application is configured to use an external certificate that is not present on this server

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Application <application_name> that is published by Web Application Proxy is using an external certificate that is not present on this server.

Impact

If an application uses a certificate that is not present on this server, Web Application Proxy will not publish the application and users will not have access.

Resolution

Obtain the configured external certificate from another server and copy it to this server with its private key.

For every application that you publish through Web Application Proxy, each Web Application Proxy server in a multi-server deployment must have a copy of the certificate used for authentication for the website with the private key.

To export the application certificate

1. On an Web Application Proxy server that has the required certificate, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. In the details pane, right-click the required certificate, click **All Tasks**, and then click **Export**.
8. On the **Certificate Export Wizard**, click **Next**.
9. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
10. On the **Export File Format** page, make sure **Personal Information Exchange** is selected, select the **Include all certificates in the certification path if possible** and **Export all extended properties** check boxes, and then click **Next**.
11. On the **Security** page, decide how to protect this certificate, and then click **Next**.
12. On the **File to Export** page, save the certificate in a location that can be accessed by the Web Application Proxy server that requires this certificate, and then click **Next**, and then click **Finish**.
13. If the export is successful, click **OK**.

After exporting the application certificate, you must import it on the Web Application Proxy server that requires the certificate.

To import the application certificate

1. On the Web Application Proxy server that requires the certificate, open an MMC console as described above.

2. In the Console window, open **Certificates/Personal/Certificates**.
3. Right-click in the details pane, click **All Tasks**, and then click **Import**.
4. On the **Certificate Import Wizard**, click **Next**.
5. On the **File to Import** page, locate the certificate that you exported previously, and then click **Next**.
6. On the **Private key protection** page, enter the password if required, and then click **Next**.
7. On the **Certificate Store** page, click **Next**, and then click **Finish**.
8. If the import is successful, click **OK**.

Web Application Proxy: Some applications are configured to perform backend authentication using Integrated Windows authentication but the server is not joined to a domain

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy can perform backend authentication using Integrated Windows authentication only when it is running on a server that is joined to a domain.

Impact

Users will not be able to access this application from the current server.

Resolution

Join this server to a domain.

When publishing applications that use Integrated Windows authentication, the Web Application Proxy server uses Kerberos constrained delegation to authenticate users to the published application.

To use Integrated Windows authentication, the Web Application Proxy server must be joined to an AD DS domain. The following lists the domain and forest requirements for a deployment using Integrated Windows authentication with Kerberos constrained delegation.

- Deployments where users, resources, and Web Application Proxy servers are all in the same forest are supported.
- In deployments with multiple forests where there is a user forest, a resource forest, and a Web Application Proxy forest, the following deployments are supported:
  - Users, resources, and Web Application Proxy servers are all in different forests.
  - Users and Web Application Proxy servers are in the same forest, but resources are in a different forest.
  - Resources and Web Application Proxy servers are in the same forest, but users are in a different forest.
  - Users and resources are in the same forest, but Web Application Proxy servers are in a different forest.

In multi-forest deployments:

1. The user forest must trust the Web Application Proxy forest, and the Web Application Proxy forest must trust the resource forest.
2. All of the Active Directory domains in a multi-forest deployment must have at least one Windows Server® 2012 or higher domain controller. For more information, see Kerberos Constrained Delegation across Domains

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To join the Web Application Proxy server to a domain

1. In Server Manager, click **Local Server**. In the details pane, click the link next to **Computer name**.
2. On the **System Properties** dialog box, click the **Computer Name** tab. On the **Computer Name** tab, click **Change**.
3. In **Computer Name**, type the name of the computer if you are also changing the computer name when joining the server to the domain. Under **Member of**, click **Domain**, and then type the name of the domain to which you want to join the server; for example, corp.contoso.com, and then click **OK**.
4. When you are prompted for a user name and password, enter the user name and password of a user with rights to join computers to the domain, and then click **OK**.
5. When you see a dialog box welcoming you to the domain, click **OK**.

6. When you are prompted that you must restart the computer, click **OK**.
7. On the **System Properties** dialog box, click **Close**.
8. When you are prompted to restart the computer, click **Restart Now**.

Web Application Proxy: A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

A cluster of Web Application Proxy servers is deployed and DirectAccess is also installed on this server.

Impact

Web Application Proxy does not support this configuration and might not work as expected.

Resolution

Consider uninstalling DirectAccess from this server.

When Web Application Proxy is deployed in a multiple server deployment, it cannot be deployed on a server that also has DirectAccess installed. You must uninstall either DirectAccess, or Web Application Proxy. The following table shows the supported deployments:

DirectAccess	VPN	Web Application Proxy
Single server deployment	Single server deployment	Single server deployment
Multisite deployment	Multiple server deployment	Not supported on the same server
Not supported on the same server	Multiple server deployment	Multiple server deployment
Cluster deployment[1]	Multiple server deployment	Multiple server deployment[2]

Notes

1—In a pre-existing DirectAccess cluster deployment, you can install Web Application Proxy only using Windows PowerShell.

2—In a pre-existing multiple server Web Application Proxy deployment, you can install DirectAccess only using Windows PowerShell.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To uninstall DirectAccess

1. On the Web Application Proxy server, open the Remote Access Management console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **RAMgmtUI.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Remote Access Management console, in the navigation pane, click **DirectAccess and VPN**.
3. In the **Tasks** pane, click **Remove Configuration Settings**.
4. On the **Confirm Remove Configuration** dialog, click **OK**.
5. On the **Removing Configuration Settings** dialog, make sure the configuration was removed, and then click **Close**.

Web Application Proxy: Could not perform Integrated Windows authentication to the backend servers

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue

addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

The backend server rejected the Kerberos ticket that was presented by Web Application Proxy.

Impact

The backend server was not able to authenticate users.

Resolution

Check the Windows event log for event 12008 to identify the application. Check the Kerberos configuration on the domain controller and backend server for each application. This might also occur if the Web Application Proxy server time is not synchronized with the domain controller or the backend server.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To <perform a configuration task>

1. Perform step 1.
2. Perform step 2.

Use explanatory text between procedures as necessary.

To <perform another configuration task>

1. Perform step 1.
2. Perform step 2.

Additional references

Brief supporting information. If there are no additional resources, delete this heading.

Web Application Proxy could not publish an application

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy could not publish an application because it could not configure the HTTP protocol stack (http.sys) for that application.

Impact

Users were not able to access the applications.

Resolution

In the Windows event log, look for events 12019 and 12020 for more details about these applications. This may occur due to misconfiguration or a conflict with other services on the server.

Web Application Proxy could not publish an application due to certificate problems

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy could not publish an application due to certificate problems.

Impact

Users were not able to access the applications.

Resolution

In the Windows event log, look for event 12021 for more details about these applications. Consider issuing a new certificate for these applications.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To view details about this application in the event log

1. On the Web Application Proxy server, open the Windows event viewer: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **eventvwr.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

2. In the navigation pane, open **Applications and Services Logs/Microsoft/Windows/Web Application Proxy/Admin**.
3. In the details pane, locate the newest event with ID 12021.
4. **Do something….**

If there is a problem with the certificate used to publish this application, you may need to request a new certificate.

To request a certificate for the published application from an internal certification authority

1. On the Web Application Proxy server, open an MMC console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **mmc.exe**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the Console window, on the **File** menu, click **Add/Remove Snap-in**.
3. On the **Add or Remove Snap-ins** dialog, double-click **Certificates**.
4. On the **Certificates snap-in** dialog, click **Computer account**, and then click **Next**.
5. On the **Select Computer** dialog, click **Local computer**, click **Finish**, and then click **OK**.
6. In the Console window, open **Certificates/Personal/Certificates**.
7. Right-click in the details pane, click **All Tasks**, and then click **Request New Certificate**.
8. On the **Certificate Enrollment** dialog, click **Next** twice.
9. On the **Request Certificates** page, select the certificate template that has been configured for website authentication, and click **More information is required to enroll for this certificate**.
10. On the **Certificate Properties** dialog, on the **Subject** tab, in **Subject name**, in the **Type** list, click **Common name**, and in the **Value** box, enter a value for this certificate that covers the application that you are attempting to publish, click **Add**, click **OK**, and then click **Enroll**.
11. After successfully enrolling for this certificate, click **Finish**.

After requesting a new certificate for this application, you can attempt to republish the application.

To publish an application

1. In the Remote Access Management console, in the Navigation pane, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.
2. On the **Publish New Application Wizard**, on the **Welcome** page, click **Next**.
3. On the **Preauthentication** page, select the required preauthentication, and then click **Next**.
4. If you selected Active Directory Federation Services (AD FS), on the **Relying Party** page, in the list of relying parties select the relying party for the application that you want to publish, and then click **Next**.
5. On the **Publishing Settings** page, do the following, and then click **Next**:

- In the **Name** box, enter a friendly name for the application.
- In the **External URL** box, enter the external URL for this application; for example, https://apps.contoso.com/.
- In the **External certificate** list, select the certificate whose subject covers the external URL.

  Note
  This is the certificate that you previously requested.
- In the **Backend server URL** box, enter the URL of the backend server. Note that this value is automatically entered when you enter the external URL and you should change it only if the backend server URL is different; for example, http://apps/.
- In the **Backend server SPN** box, enter the service principal name for the backend server; for example, HTTP/apps.contoso.com.

6. On the **Confirmation** page, review the settings, and then click **Publish**. You can copy the PowerShell command set up additional published applications.
7. On the **Results** page, make sure that the application published successfully, and then click **Close**.

Web Application Proxy was not able to check for configuration updates

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Warning
Category	Configuration

Issue

Web Application Proxy was not able to check for configuration updates from AD FS storage.

Impact

Web Application Proxy was temporarily unable to retrieve its configuration from AD FS storage and used the latest copy of the configuration; any updates to the configuration were not applied on this server. When the service starts, there is no copy of the configuration and any previously published applications are temporarily unavailable to users.

Resolution

Check connectivity between Web Application Proxy and the AD FS servers.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To check connectivity between Web Application Proxy and the AD FS servers

1. Perform step 1.
2. Perform step 2.

Web Application Proxy stopped working because it could not retrieve its configuration

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy stopped working because it could not retrieve its configuration from AD FS storage.

Impact

Web Application Proxy stopped working and users were not able to access published applications.

Resolution

Check connectivity between Web Application Proxy and the AD FS servers. If the problem persists, consider running the Install-WebApplicationProxy cmdlet.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To check connectivity between Web Application Proxy and the AD FS servers

1. Perform step 1.
2. Perform step 2.

If the problem persists, it might be required to run the Install-WebApplicationProxy cmdlet.

To install Web Application Proxy

1. On the Web Application Proxy server, open a Windows PowerShell window.
2. Obtain credentials for the local administrator account on the AD FS servers:

   ```
   $fscred = get-credential
   ```

3. Obtain the certificate thumbprint to be used by Web Application Proxy as the AD FS proxy:

   ```
   dir Cert:\LocalMachine\My
   ```

   Locate the certificate in the list and copy the thumbprint to the clipboard.

4. Run the cmdlet to install Web Application Proxy:

   ```
   Install-WebApplicationProxy –FederationServiceCredential
   $fscred -CertificateThumbprint <certificate_thumbprint> -
   FederationServiceName <e.g. fs.contoso.com>
   ```

Web Application Proxy: The AD FS Proxy service is stopped

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Active Directory Federation Services service (adfssrv) is stopped.

Impact

Users will not be able to authenticate and will not be able to access published applications.

Resolution

Look for problems in the Windows event log. If no problems are found, restart the server. If this doesn't solve the problem, reinstall Web Application Proxy.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To view details about this application in the event log

1. On the Web Application Proxy server, open the Windows event viewer console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **eventvwr.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.

2. In the navigation pane, open **Applications and Services Logs/Microsoft/Windows/Web Application Proxy/Admin**.
3. In the details pane, **look for problems**.
4. **What now?**

If the event log indicates that there are no problems, restart the Web Application Proxy server.

To restart the Web Application Proxy server

1. On the Web Application Proxy server, open the **Charms** bar.
2. On the **Charms** bar, click **Settings**, click **Power**, and then click **Restart**.
3. When the server restarts, check if the problem persists.

If restarting the Web Application Proxy server does not solve the problem, reinstall Web Application Proxy. Use the procedures in Step 2: Install and Configure the Web Application Proxy Server to install the Web Application Proxy role service and to configure Web Application Proxy.

Web Application Proxy service is stopped

This topic is intended to address a specific issue identified by a Best Practices Analyzer scan. You should apply the information in this topic only to computers that have had the Web Application Proxy Best Practices Analyzer run against them and are experiencing the issue addressed by this topic. For more information about best practices and scans, see Best Practices Analyzer.

Operating System	Windows Server 2012 R2
Product/Feature	Web Application Proxy
Severity	Error
Category	Configuration

Issue

Web Application Proxy service (appproxysvc) is stopped.

Impact

Users will not be able to access published applications.

Resolution

Look for problems in the Windows event log. If no problems are found, restart the server. If this doesn't solve the problem, reinstall Web Application Proxy.

Additional resolution content is required; provide (1) an introductory paragraph and (2) step-by-step procedures for each task in the best practice configuration. Each resolution should be complete and self-contained. The user should be able to implement the BP without having to go to another page, unless the implementation steps are long and complicated (e.g., Exchange disaster recovery steps). If you do have to take the customer to another topic, this topic should contain a high-level overview of those particular resolution steps.

Credential statement, if necessary.

To view details about this application in the event log

1. On the Web Application Proxy server, open the Windows event viewer console: On the **Start** screen, click the **Apps** arrow. On the **Apps** screen, type **eventvwr.msc**, and then press ENTER. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. In the navigation pane, open **Applications and Services Logs/Microsoft/Windows/Web Application Proxy/Admin**.
3. In the details pane, **look for problems**.
4. **What now?**

If the event log indicates that there are no problems, restart the Web Application Proxy server.

To restart the Web Application Proxy server

1. On the Web Application Proxy server, open the **Charms** bar.
2. On the **Charms** bar, click **Settings**, click **Power**, and then click **Restart**.
3. When the server restarts, check if the problem persists.

If restarting the Web Application Proxy server does not solve the problem, reinstall Web Application Proxy. Use the procedures in Step 2: Install and Configure the Web Application Proxy Server to install the Web Application Proxy role service and to configure Web Application Proxy.

TCP/IP

Transmission Control Protocol/Internet Protocol (TCP/IP) is an industry standard suite of protocols that is designed for large networks consisting of network segments that are connected by routers. TCP/IP is the core protocol suite that is used on the Internet.

In Windows Server 2012, TCP/IP – including both Internet Protocol version 4 (IPv4) and IPv6 – is unchanged from TCP/IP in Windows Server 2008 R2. For more information, see TCP/IP in the Windows Server® 2008 and Windows Server 2008 R2 Technical Library.

Telephony Server and the Telephony API Overview

The Telephony service in Windows Server 2012 provides distributed access to shared telephony devices. For example, a LAN-based server may have multiple connections to a local telephone switch or private branch exchange (PBX). Telephony uses the Telephony Application Programming Interface (TAPI) to enable computer programs to communicate with the shared telephony services. TAPI operations that are invoked by a client are forwarded over the network to the server. The server can then use a non-Microsoft call control between itself and the PBX to fulfill the client's call control requests. TAPI services in Windows Server 2012 enable programs that effectively work the same way over intranets, the Internet, and traditional public phone networks.

TAPI services in Windows Server 2012 have not significantly changed since Windows Server 2003. The documentation for TAPI has not been updated for this version of Windows. You can find all of the documentation for TAPI at see Telephony in the Windows Server TechCenter.

Virtual Receive-side Scaling in Windows Server 2012 R2

Virtual Receive-side scaling is feature in Windows Server® 2012 R2 that allows the load from a virtual network adapter to be distributed across multiple virtual processors in a virtual machine.

Feature description

Receive-side network processing in multi-core virtual machines is conventionally bottlenecked by the fact that a single virtual processor services all the interrupts from a virtual network adapter. Virtual Receive-side scaling solves this problem by enabling a network adapter to distribute its network processing load across multiple virtual processors in multi-core virtual machines.

Virtual Receive-side scaling coexists with the following networking technologies:

- IPv4 and IPv6
- TCP and UDP
- LBFO (NIC Teaming)
- Live Migration
- NVGRE

Virtual Receive-side scaling is not enabled by default. It must be enabled in the virtual machine using PowerShell cmdlets.

Virtual Receive-side scaling in Windows Server 2012 R2 is easily enabled and managed inside the virtual machine, much like a physical host, using the same PowerShell cmdlets that are used for Receive-side scaling in a physical host.

Practical applications

Processor load balancing

Anthony, a network administrator, is setting up a new host for his workplace with 2 NICs that are SR-IOV capable. He is going to use Windows Server 2012 R2 to host a file server system that his workplace is purchasing and implementing. After the hardware and software are installed, Anthony configures his VM to use 8 virtual processors and 4096 MB of memory. Unfortunately, Anthony does not have the option of turning on SR-IOV because his VMs rely on policy enforcement through the vmSwitch. Initially, Anthony assigns 4 virtual processors through PowerShell to be available for use with vRSS. After a week, the service has become extremely popular and Anthony checks the performance. He discovers that all four virtual processors are fully utilized. Anthony changes the RSS processor and assigns two more virtual processors to be available to RSS in the virtual machine to help handle the large network load.

Software load balancing

Sandra, a network administrator, is setting up a single large virtual machine on one of her systems for the sole purpose of being a software load balancer. Sandra has just installed Windows Server 2012 R2 so that she can expand VMQ to use multiple processors per vmNIC. Since Sandra only has one vmNIC in this system she decides she will turn on RSS in the VM to get spreading for processing through the vmSwitch. Since vRSS is disabled by default, Sandra enables RSS using PowerShell cmdlets.

Software requirements

Virtual Receive-side scaling can be used on any computer running Windows Server 2012 R2. The following Windows version virtual machines support Virtual Receive-side scaling:

Virtual Receive-side scaling can be used on any computer running Windows Server 2012 R2. The following Windows version virtual machines support Virtual Receive-side scaling:

1. Windows Server 2012 R2
2. Windows® 8.1
3. Windows Server 2012 R2 with the Windows Server 2012 R2 integration components installed.
4. Windows 8.1 with the Windows Server 2012 R2 integration components installed.

Hardware requirements

Virtual Receive-side scaling requires Virtual Machine Queue (VMQ) support from the physical adapter. If VMQ is disabled or is not supported, then Virtual Receive-side scaling is disabled for the host.

Enable and Configure Virtual Receive-side scaling

Virtual Receive-side scaling requires Virtual Machine Queue (VMQ) support from the physical adapter. If VMQ is disabled, or is not supported then Virtual Receive-side scaling is disabled.

Before enabling Virtual Receive-side scaling

There are 4 steps you should take before implementing Virtual Receive-side scaling:

1. Verify that the network adapter is VMQ capable and has a link speed of >= 10G. If the link speed is less than 10G, VMQ is disabled by default by the vmSwitch even though it will still show as enabled in the PowerShell cmdlet, **Get-NetAdapterVmq**. One way to verify that VMQ is disabled is to use the cmdlet, **Get-NetAdapterVmqQueue**. This will show that there is not a QueueID assigned to the VM or host vNIC
2. Verify that VMQ is enabled on the host machine. Virtual Receive-side scaling will not work if the host does not support VMQ. You can check this by running **Get-VMSwitch** and finding the adapter that the vmSwitch is using. Next, run **Get-NetAdapterVmq** and ensure that the adapter is shown in the results and has VMQ enabled.
3. Verify that an SRIOV Virtual Function (VF) driver is not attached to the VM network interface. This can be done using the **Get-NetAdapterSriov** cmdlet. If a VF driver is loaded, Receive-side scaling will use the scaling settings from this driver instead of those configured by Virtual Receive-side scaling. If the VF driver does not support Receive-side scaling then Virtual Receive-side scaling is disabled.
4. If you are using LBFO (NIC Teaming), it is essential that VMQ be properly configured to work with the LBFO settings. For detailed information about LBFO deployment and management, see NIC Teaming Configuration and Management.

Enabling Virtual Receive-side scaling

Virtual Receive-side scaling can be enabled using PowerShell or with Device Manager:

To enable Virtual Receive-side scaling using Device Manager

1. On the virtual machine, open **Device Manager** (In **Settings** click **Control Panel**, and then click **Device Manager**).
2. Expand **Network adapters**, right-click the network adapter you want to work with, and then click **Properties**.
3. On the **Advanced** tab in the network adapter properties, locate the setting for **Receive-side scaling** and make sure it is enabled.

Note

Some network adapters advertise the number of receive-side scaling queues they support on the **Advanced** tab.

To enable Virtual Receive-side scaling using PowerShell

1. On the virtual machine, open **PowerShell**.
2. At the command line, type

   ```
   Enable-NetAdapterRSS -Name "AdapterName"
   ```

   replacing “AdapterName” with the name of the vmNIC adapter that Virtual Receive-side scaling should be enabled on.

   Alternately, the following command will also enable Virtual Receive-side scaling:

   ```
   Set-NetAdapterRSS -Name "AdapterName" -Enabled $True
   ```

Configuring Virtual Receive-side scaling

Virtual Receive-side scaling settings in the virtual machine are configured using **Set-NetAdapterRss**, which is the same command used for native RSS. You can view and configure the settings of the virtual machine by using the cmdlet, **Get-NetAdapterRSS**. A portion of the PowerShell help for the cmdlet is below, and a full description of the cmdlets that lists all the configurable settings and expected types can be found on TechNet at **Set-NetAdapterRss**.

Note

Setting the profile inside the virtual machine will not impact the scheduling of the work. The hypervisor makes all the scheduling decisions and ignores the profile inside the virtual machine.

Inside the VM:

```
get-help Set-NetAdapterRss

NAME

    Set-NetAdapterRss


SYNOPSIS

    Sets the receive side scaling (RSS) properties on the network adapter.


SYNTAX

    Set-NetAdapterRss [-Name] <String[]> [-AsJob [<SwitchParameter>]] [-
BaseProcessorGroup <UInt16>] [-BaseProcessorNumber <Byte>] [-CimSession <CimSession[]>]
[-Enabled <Boolean>] [-IncludeHidden [<SwitchParameter>]] [-MaxProcessorGroup <UInt16>]
[-MaxProcessorNumber <Byte>] [-MaxProcessors <UInt32>][-NoRestart [<SwitchParameter>]] [-
NumaNode <UInt16>] [-NumberOfReceiveQueues <UInt32>] [-PassThru [<SwitchParameter>]] [-
```

```
Profile <Profile>] [-ThrottleLimit <Int32>] [-Confirm [<SwitchParameter>]] [-WhatIf
[<SwitchParameter>]] [<CommonParameters>]
```

Inside the host, you will need to use the keywords that control the VMQ processors. You can view the current settings inside the VM by using the cmdlet, **Get-NetAdapterVmq**. To configure the settings you will use the **Set-NetAdapterVmq** cmdlet. The output for this cmdlet is slightly different than that of native Receive-side scaling, so a snipped portion of **get-help** for the cmdlet is below. A detailed description of the cmdlet can be found on TechNet at **Set-NetAdapterVmq**.

Inside the Host:

```
get-help set-netadaptervmq

NAMEv

    Set-NetAdapterVmq


SYNOPSIS

    Sets the virtual machine queue (VMQ) properties of the network adapter.


SYNTAX

    Set-NetAdapterVmq [-Name] <String[]> [-AsJob [<SwitchParameter>]] [-
BaseProcessorGroup <UInt16>] [-BaseProcessorNumber <Byte>] [-CimSession <CimSession[]>]
[-Enabled <Boolean>] [-IncludeHidden [<SwitchParameter>]] [-MaxProcessorNumber <Byte>] [-
MaxProcessors <UInt32>] [-NoRestart [<SwitchParameter>]] [-NumaNode <UInt16>] [-PassThru
[<SwitchParameter>]] [-ThrottleLimit <Int32>] [-Confirm [<SwitchParameter>]] [-WhatIf
[<SwitchParameter>]] [<CommonParameters>]
```

Disabling Virtual Receive-side scaling

Virtual Receive-side scaling can be disabled using PowerShell or with Device Manager:

To disable Virtual Receive-side scaling using Device Manager

1. On the virtual machine, open **Device Manager** (In **Settings** click **Control Panel**, and then click **Device Manager**).
2. Expand **Network adapters**, right-click the network adapter you want to work with, and then click **Properties**.
3. On the **Advanced** tab in the network adapter properties, locate the setting for **Receive-side scaling** and make sure it is disabled.

To disable Virtual Receive-side scaling using PowerShell

1. On the virtual machine, open **PowerShell**.
2. At the command line, type

```
Disable-NetAdapterRSS -Name "AdapterName"
```

   replacing “AdapterName” with the name of the vmNIC adapter that Virtual Receive-side scaling should be enabled on.

   Alternately, the following command will also enable Virtual Receive-side scaling:

```
Set-NetAdapterRSS -Name "AdapterName" -Enabled $False
```

Managing Virtual Receive-side scaling FAQ

1. **Virtual Receive-side scaling is on, but how do I know if it is working?**

   You’ll be able to tell Virtual Receive-side scaling is working by opening the task manager in your virtual machine and viewing the virtual processor utilization. If there are multiple connections established to the virtual machine, you will see more than one core above 0% utilization.
2. **Are there any perfmon counters I can check?**

   The same holds true for the host. Virtual Receive-side scaling made changes to the VMQ algorithm to expand virtual machine processing to more than 1 core for vmSwitch processing. If a single core reaches 80% utilization, you should see the traffic start to expand to multiple cores.
3. **Are there any perfmon counters I can check?**

   In **Perfmon**, there are three counters that can help you evaluate Virtual Receive-side scaling. Once in **Perfmon**, right click on the graph and click **Add Counters**…. Click the **Hyper-V Virtual Switch Processor** category and the counters are under:

   a. **Number of VMQs** – The number of VMQ processors on affinitized to that processor
   b. **Packets from External** – Packets indicated to a processor from any external NIC
   c. **Packets from Internal** – Packets indicated to a processor from any internal interfaces, ie vNIC (Host NIC), vmNIC.
4. **I’m looking at the host and not all of the processors are being used. It looks like every other one is being skipped.**

   Check to see if hyper threading is enabled. VMQ and Virtual Receive-side scaling are both designed to skip hyper threaded cores.

Windows Firewall with Advanced Security Overview

This is an overview of the Windows Firewall with Advanced Security (WFAS) and Internet Protocol security (IPsec) features in Windows Server 2012.

Did you mean…

- Windows Firewall with Advanced Security in Windows Server 2008 R2

Feature description

Windows Firewall with Advanced Security is an important part of a layered security model. By providing host-based, two-way network traffic filtering for a computer, Windows Firewall with Advanced Security blocks unauthorized network traffic flowing into or out of the local computer. Windows Firewall with Advanced Security also works with Network Awareness so that it can apply security settings appropriate to the types of networks to which the computer is connected. Windows Firewall and Internet Protocol Security (IPsec) configuration settings are integrated into a single Microsoft Management Console (MMC) named Windows Firewall with Advanced Security, so Windows Firewall is also an important part of your network's isolation strategy.

Practical applications

To help address your organizational network security challenges, Windows Firewall with Advanced Security offers the following benefits:

- **Reduces the risk of network security threats.** Windows Firewall with Advanced Security reduces the attack surface of a computer, providing an additional layer to the defense-in-depth model. Reducing the attack surface of a computer increases manageability and decreases the likelihood of a successful attack. Network Access Protection (NAP), a feature of Windows Server 2012, also helps ensure client computers comply with policies that define the required software and system configurations for computers that connect to your network. The integration of NAP helps prevent communications between compliant and noncompliant computers.
- **Safeguards sensitive data and intellectual property.** With its integration with IPsec, Windows Firewall with Advanced Security provides a simple way to enforce authenticated, end-to-end network communications. It provides scalable, tiered access to trusted network resources, helping to enforce integrity of the data, and optionally helping to protect the confidentiality of the data.
- **Extends the value of existing investments.** Because Windows Firewall with Advanced Security is a host-based firewall that is included with Windows Server 2012, and prior Windows operating systems and because it is tightly integrated with Active Directory® Domain Services (AD DS) and Group Policy, there is no additional hardware or software required. Windows Firewall with Advanced Security is also designed to complement existing non-Microsoft network security solutions through a documented application programming interface (API).

New and changed functionality

The following table lists some of the new features for Windows Firewall with Advanced Security in Windows Server 2012.

Feature/functionality	Windows Server 2008 R2	Windows Server 2012
Internet Key Exchange version 2 (IKEv2) for IPsec transport mode		X
Windows Store app network isolation		X
Windows PowerShell cmdlets for Windows Firewall		X

IKEv2 for IPsec transport mode

In Windows Server 2012, IKEv2 supports additional scenarios including IPsec end-to-end transport mode connections.

What value does this change add?

Windows Server 2012 IKEv2 support provides interoperability for Windows with other operating systems using IKEv2 for end-to-end security, and Supports Suite B (RFC 4869) requirements.

What works differently?

In Windows Server 2008 R2, IKEv2 is available as a virtual private network (VPN) tunneling protocol that supports automatic VPN reconnection. IKEv2 allows the security association to remain unchanged despite changes in the underlying connection.

In Windows Server 2012, IKEv2 support has been expanded.

Windows Store app network isolation

Administrators can custom configure Windows Firewall to fine tune network access if they desire more control of their Windows Store apps.

What value does this change add?

The feature adds the ability to set and enforce network boundaries ensure that apps that get compromised can only access networks where they have been explicitly granted access. This significantly reduces the scope of their impact to other apps, the system, and the network. In addition, apps can be isolated and protected from malicious access from the network.

What works differently?

In addition to firewall rules that you can create for program and services, you can also create firewall rules for Windows Store apps and their various capabilities.

Windows PowerShell cmdlets for Windows Firewall

Windows PowerShell has extensive cmdlets to allow Windows Firewall configuration and management.

What value does this change add?

You can now fully configure and manage Windows Firewall, IPsec, and related features using the very powerful and scriptable Windows PowerShell.

What works differently?

In previous Windows versions, you could use Netsh to perform many configuration and management functions. This capability has been greatly expanded using the more powerful Windows PowerShell scripting language.

See also

See the following topics for more information about Windows Firewall with Advanced Security in Windows Server 2012.

Content type	References
Deployment	Securing End-to-End IPsec Connections by Using IKEv2 in Windows Server 2012 | Isolating Windows Store Apps on Your Network | Windows Firewall with Advanced Security Administration with Windows PowerShell
Troubleshooting	Troubleshooting Windows Firewall with Advanced Security in Windows Server 2012

Isolating Windows Store Apps on Your Network

When you add new computers and devices that are running Windows 8 to your network, you may want to customize your Windows Firewall configuration to isolate the network access of the new Windows Store apps that run on them. Developers who build Windows Store apps can declare certain app capabilities that enable different classes of network access. A developer can decide what kind of network access the app requires and configure this capability for the app. When the app is installed on a computer running Windows 8, appropriate firewall rules are automatically

created to enable access. Administrators can then customize the firewall configuration to further fine-tune this access if they desire more control over the network access for the app.

For example, a developer can decide that their app should only connect to trusted local networks (such as at home or work), and not to the Internet. In this way, developers can define the scope of network access for their app. This network isolation prevents an app from accessing a network and a connection type (inbound or outbound) if the connection has not been configured for the app. Then the network administrator can customize the firewall to further restrict the resources that the app can access.

The ability to set and enforce these network boundaries ensures that apps that get compromised can only access networks where they have been explicitly granted access. This significantly reduces the scope of their impact on other apps, the computer, and the network. In addition, apps can be isolated and protected from malicious access from the network.

When creating new Windows Store apps, a developer can define the following network capabilities for their app:

- **Home\Work Networking**

  Provides inbound and outbound access to intranet networks that the user has designated as a home or a work network, or if the network has an authenticated domain controller.

- **Internet (Client)**

  Provides outbound access to the Internet and untrusted networks, such as airports and coffee shops (for example, intranet networks where the user has designated the network as Public). Most apps that require Internet access should use this capability.

- **Internet (Client and Server)**

  Provides inbound and outbound access to the Internet and untrusted networks, such as airports and coffee shops. This capability is a superset of the **Internet (Client)** capability, and **Internet (Client)** does not need to be enabled if this capability is enabled.

- **Proximity**

  Provides near-field communication (NFC) with devices that are in close proximity to the computer. Proximity may be used to send files or connect with an application on a proximate device.

In this document

To isolate Windows Store apps on your network, you need to use Group Policy to define your network isolation settings and create custom Windows Store app firewall rules.

- Prerequisites
- Step 1: Define your network
- Step 2: Create custom firewall rules

Prerequisites

- A domain controller is installed on your network, and your computers are joined to the Windows domain.

- Your Windows Store app is installed on your client computer.
- The Remote Server Administration Tools (RSAT) are installed on your client computer. When you perform the following steps from your client computer, you can select your Windows Store app when you create Windows Firewall rules.

  Note
  You can install the RSAT on your computer running Windows 8 from the Microsoft Download Center.

Step 1: Define your network

The **Home\Work Networking** capability enables access to intranet resources. Administrators can use Group Policy settings to define the scope of the intranet. This ensures that Windows Store apps can access intranet resources appropriately.

The Windows Store Internet Explorer app that is included with Windows 8 uses the network capabilities to detect which zone it should use. The browser uses the network capabilities to ensure that it operates in the correct security zone.

A network endpoint is considered part of the **Home\Work Network** if:

- It is part of the local subnet of a trusted network.

  For example, home users generally flag their network as Trusted. Local computers will be designated as such.
- A computer is on a network, and it is authenticated to a domain controller.
  - Endpoints within the intranet address space are considered private.
  - Endpoints within the local subnet are considered private.
- The computer is configured for DirectAccess, and the endpoint is part of the intranet address space.

The intranet address space is composed of configured Active Directory sites and subnets, and it is configured for Windows network isolation specifically by using Group Policy. You can disable the usage of Active Directory sites and subnets by using Group Policy by declaring that your subnet definitions are authoritative.

Any proxies that you configure or that are automatically configured with proxy autoconfiguration (by using Web Proxy Auto-Discovery (WPAD) protocol) are exempt from the intranet zone. You can add proxy addresses by using Group Policy.

All other endpoints that do not meet the previously stated criteria are considered endpoints on the Internet.

To configure a GPO that defines your intranet address space

1. Open the Group Policy Management snap-in (gpmc.msc) and edit the Default Domain Policy.
2. From the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Network**, and click **Network**

Isolation.

3. In the right pane, double-click **Private network ranges for apps**.
4. In the **Private network ranges for apps** dialog box, click **Enabled**. In the **Private subnets** text box, type the private subnets for your intranet, separated by commas if necessary.

   For example, if the Contoso intranet is defined as 10.0.0.0 with a subnet mask of 255.255.255.0, you would type 10.0.0.0/24 in the **Private subnets** text box.
5. Double-click **Subnet definitions are authoritative**.

   If you want the subnet definitions that you previously created to be the single source for your subnet definition, click **Enabled**. Otherwise, leave the **Not Configured** default so that you can add additional subnets by using local settings or network isolation heuristics.

To configure the proxy addresses for the intranet and Internet

1. Double-click **Internet proxy servers for apps**. Click **Enabled**, and then in the **Domain Proxies** text box, type the IP addresses of your Internet proxy servers, separated by semicolons.
2. Double-click **Intranet proxy servers for apps**. Click **Enabled**, and then in the IP address text box, type the IP addresses of your intranet proxy servers, separated by semicolons.
3. Double-click **Proxy definitions are authoritative**.

   If you want the proxy definitions that you previously created to be the single source for your proxy definition, click **Enabled**. Otherwise, leave the **Not Configured** default so that you can add additional proxies by using local settings or network isolation heuristics.

Step 2: Create custom firewall rules

Windows Store apps can declare many capabilities in addition to the network capabilities discussed previously. For example, apps can declare capabilities to access user identity, the local file system, and certain hardware devices.

The following table provides a complete list of the possible app capabilities.

Capability	Name	Description
Internet (Client)	internetClient	Your outgoing Internet connection.
Internet (Client & Server)	internetClientServer	Your Internet connection, including incoming unsolicited connections from the Internet The app can send information to or from your computer through a

Capability	Name	Description
		firewall. You do not need to declare **internetClient** if this capability is declared.
Home\Work Networking	privateNetworkClientServer	A home or work network. The app can send information to or from your computer and other computers on the same network.
Document Library Access	documentsLibrary	Your Documents library, including the capability to add, change, or delete files. The package can only access file types that are declared in the manifest. The app cannot access document libraries on HomeGroup computers.
Picture Library Access	picturesLibrary	Your Pictures library, including the capability to add, change, or delete files. This capability also includes Picture libraries on HomeGroup computers and picture file types on locally connected media servers.
Video Library Access	videosLibrary	Your Videos library, including the capability to add, change, or delete files. This capability also includes Video libraries on HomeGroup computers and video file types on locally connected media servers.
Music Library Access	musicLibrary	Your Music library, including the capability to add, change, or delete files. This capability also includes Music libraries on

Capability	Name	Description
		HomeGroup computers and music file types on locally connected media servers.
Default Windows Credentials	defaultWindowsCredentials	Your Windows credentials for access to a corporate intranet. This application can impersonate you on the network.
Removable Storage	removableStorage	A removable storage device, such as an external hard disk, USB flash drive, or MTP portable device, including the capability to add, change, or delete specific files. This package can only access file types that are declared in the manifest.
Shared User Certificates	sharedUserCertificates	Software and hardware certificates or a smart card, which the app uses to identify you. This capability can be used by an employer, a bank, or government services to identify you.
Location	location	Provides access to the user's current location.
Microphone	microphone	Provides access to the microphone's audio feed.
Near-field Proximity	proximity	Required for near-field communication (NFC) between devices in close proximity. NFC can be used to send files or connect with an app on a proximate device.
Text Messaging	sms	Provides access to computer text messaging functionality.

Capability	Name	Description
Webcam	webcam	Provides access to the webcam's video feed.
Other devices (represented by GUIDs)	<GUID>	Includes specialized devices and Windows Portable Devices.

In Windows Server 2012, it is possible to create a Windows Firewall policy that is scoped to a set of apps that use a specified capability or scoped to a specific Windows Store app.

For example, you could create a Windows Firewall policy to block Internet access for any apps on your network that have the Documents Library capability.

To block Internet access for any apps on your network that have the Documents Library capability

1. Open the Group Policy Management snap-in (gpmc.msc).
2. In the left pane, right-click your domain name and click **Create a GPO in this domain, and link it here**.
3. Type a name for the GPO in the **Name** text box, and then click **OK**.
4. Right-click the new GPO, and then click **Edit**.
5. In the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Windows Firewall with Advanced Security**, and click **Windows Firewall with Advanced Security – LDAP://…**
6. Right-click **Outbound Rules**, and then click **New Rule**.
7. Click **Custom**, and then click **Next**.
8. Click **Next** on the **Program** page, the **Protocols and Ports** page, and the **Scope** page.
9. On the **Action** page, ensure that **Block the Connection** is selected, and then click **Next**.
10. On the **Profile** page, click **Next**.
11. On the **Name** page, type a name for your rule, and then click **Finish**.
12. In the right pane, right-click your new rule and click **Properties**.
13. Click the **Local Principals** tab, select the **Only allow connections from these users** check box, and then click **Add**.
14. Click **Application Package Properties**, and then click **OK**.
15. In the **Choose Capabilities** dialog box, click **APPLICATION PACKAGE AUTHORITY\Your documents library**, and then click **OK**.
16. Click the **Scope** tab under **Remote IP addresses**, and then click **Add**.
17. Click **Predefined set of computers**, select **Internet**, and click **OK**.

    This scopes the rule to block traffic to Internet computers.

18. Click the **Programs and Services** tab, and in the **Application Packages** area, click **Settings**.
19. Click **Apply to application packages only**, and then click **OK**.

    Important
    You must do this to ensure that the rule applies only to Windows Store apps and not to other applications and programs. Non-Windows Store applications and programs declare all capabilities by default, and this rule would apply to them if you do not configure it this way.
20. Click **OK** to close the **Properties** dialog box.
21. Close the Group Policy Management Editor.
22. In the Group Policy Management snap-in, ensure that your new GPO is selected, and in the right pane under **Security Filtering**, select **Authenticated Users**. Click **Remove**, and then click **OK**.
23. Under **Security Filtering**, click **Add**.
24. Type **domain computers** in the text box, and then click **OK**.
25. Close the Group Policy Management snap-in.

Use the following procedure if you want to block intranet access for a specific media sharing app on your network.

To block intranet access for a specific media sharing app on your network

1. Open the Group Policy Management snap-in (gpmc.msc).
2. In the left pane, right-click your domain name, and then click **Create a GPO in this domain, and link it here**.
3. Type a name for your GPO in the **Name** text box, and then click **OK**.
4. Right-click your new GPO, and then click **Edit**.
5. From the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Windows Firewall with Advanced Security**, and then click **Windows Firewall with Advanced Security – LDAP://**…
6. Right-click **Outbound Rules**, and then click **New Rule**.
7. Click **Custom**, and then click **Next**.
8. Click **Next** on the **Program** page, the **Protocols and Ports** page, and the **Scope** page.
9. On the **Action** page, ensure **Block the Connection** is selected, and then click **Next**.
10. On the **Profile** page, click **Next**.
11. On the **Name** page, type a name for your rule, and then click **Finish**.
12. In the right pane, right-click your new rule, and then click **Properties**.
13. Click the **Local Principals** tab, select the **Only allow connections from these users** check box, and then click **Add**.
14. Click **Application Package Properties**, and then click **OK**.

15. In the **Choose Capabilities** dialog box, click **APPLICATION PACKAGE AUTHORITY\A home or work network**, and then click **OK**.
16. Click the **Programs and Services** tab under **Application Packages**, and then click **Settings**.
17. Click **Apply to this application package**, select the app in the text box, and then click **OK**.
18. Click **OK** to close the **Properties** dialog box.
19. Close the Group Policy Management Editor.
20. In Group Policy Management, ensure that your new GPO is selected, and in the right pane under **Security Filtering**, select **Authenticated Users**, click **Remove**, and then click **OK**.
21. Under **Security Filtering**, click **Add**.
22. Type **domain computers** in the text box and click **OK**.
23. Close Group Policy Management.

See also

- Windows Firewall with Advanced Security Overview

Securing End-to-End IPsec Connections by Using IKEv2 in Windows Server 2012

In Windows Server 2012, Internet Key Exchange version 2 (IKEv2) support is broadened from previous Windows versions.

For example, in Windows Server 2012, IKEv2 does the following:

- Supports additional scenarios, including IPsec end-to-end transport mode connections
- Provides interoperability for Windows with other operating systems that use IKEv2 for end-to-end security
- Supports Suite B (RFC 4869) requirements
- Coexists with existing policies that deploy AuthIP/IKEv1
- Uses the Windows PowerShell interface exclusively for configuration. You cannot configure IKEv2 through the user interface.
- Uses certificates for the authentication mechanism

In Windows Server 2008 R2, IKEv2 is available as a virtual private network (VPN) tunneling protocol that supports automatic VPN reconnection. IKEv2 allows the security association to remain unchanged despite changes in the underlying connection.

In this document

- Prerequisites
- Computers joined to a domain

- Computers not joined to a domain
- Troubleshooting

Note
This topic includes sample Windows PowerShell cmdlets. For more information, see How to Run a Windows PowerShell Cmdlet.

Prerequisites

These procedures assume that you already have a public key infrastructure (PKI) in place for computer authentication.

Computers joined to a domain

The following Windows PowerShell script establishes a connection security rule that uses IKEv2 for communication between two computers (CLIENT1 and SERVER1) that are joined to the corp.contoso.com domain as shown in Figure 1.

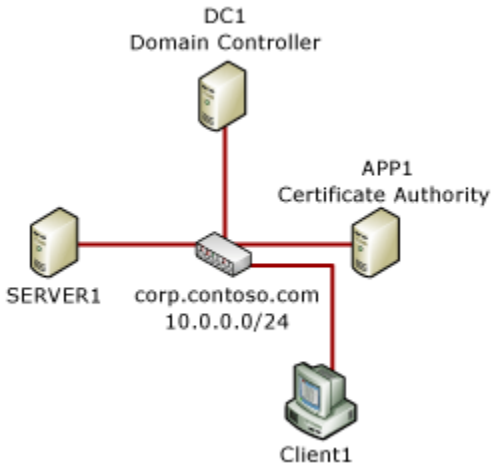

Figure 1 The Contoso corporate network

This script does the following:

- Creates a security group called **IPsec client and servers** and adds CLIENT1 and SERVER1 as members.
- Creates a Group Policy Object (GPO) called **IPsecRequireInRequestOut** and links it to the corp.contoso.com domain.
- Sets the permissions to the GPO so that they apply only to the computers in **IPsec client and servers** and not to **Authenticated Users**.
- Indicates the certificate to use for authentication.

Important
The certificate parameters that you specify for the certificate are case sensitive, so make sure that you type them exactly as specified in the certificate, and place the

parameters in the exact order that you see in the following example. Failure to do so will result in connection errors.

- Creates the IKEv2 connection security rule called **My IKEv2 Rule**.

Windows PowerShell commands

Type each cmdlet on a single line, even though they may appear to wrap across several lines because of formatting constraints.

```
# Create a Security Group for the computers that will get the policy

$pathname = (Get-ADDomain).distinguishedname

New-ADGroup -name "IPsec client and servers" -SamAccountName "IPsec client and servers" `

-GroupCategory security -GroupScope Global -path $pathname


# Add test computers to the Security Group

$computer = Get-ADComputer -LDAPFilter "(name=client1)"

Add-ADGroupMember -Identity "IPsec client and servers" -Members $computer

$computer = Get-ADComputer -LDAPFilter "(name=server1)"

Add-ADGroupMember -Identity "IPsec client and servers" -Members $computer


# Create and link the GPO to the domain

$gpo = New-gpo IPsecRequireInRequestOut

$gpo | new-gplink -target "dc=corp,dc=contoso,dc=com" -LinkEnabled Yes


# Set permissions to security group for the GPO

$gpo | Set-GPPermissions -TargetName "IPsec client and servers" -TargetType Group -
PermissionLevel GpoApply -Replace

$gpo | Set-GPPermissions -TargetName "Authenticated Users" -TargetType Group -
PermissionLevel None -Replace


#Set up the certificate for authentication

$gponame = "corp.contoso.com\IPsecRequireInRequestOut"

$certprop = New-NetIPsecAuthProposal -machine -cert -Authority "DC=com, DC=contoso,
DC=corp, CN=corp-APP1-CA"

$myauth = New-NetIPsecPhase1AuthSet -DisplayName "IKEv2TestPhase1AuthSet" -proposal
$certprop –PolicyStore GPO:$gponame
```

```
#Create the IKEv2 Connection Security rule

New-NetIPsecRule  -DisplayName "My IKEv2 Rule" -RemoteAddress any -Phase1AuthSet
$myauth.InstanceID `

-InboundSecurity Require -OutboundSecurity Request -KeyModule IKEv2 -PolicyStore
GPO:$gponame
```

Computers not joined to a domain

Use a Windows PowerShell script similar to the following to create a local IPsec policy on the computers that you want to include in the secure connection.

Important

The certificate parameters that you specify for the certificate are case sensitive, so make sure that you type them exactly as specified in the certificate, and place the parameters in the exact order that you see in the following example. Failure to do so will result in connection errors.

Windows PowerShell commands

Type each cmdlet on a single line, even though they may appear to wrap across several lines because of formatting constraints.

```
#Set up the certificate

$certprop = New-NetIPsecAuthProposal -machine -cert -Authority "DC=com, DC=contoso,
DC=corp, CN=corp-APP1-CA"

$myauth = New-NetIPsecPhase1AuthSet -DisplayName "IKEv2TestPhase1AuthSet" -proposal
$certprop


#Create the IKEv2 Connection Security rule

New-NetIPsecRule  -DisplayName "My IKEv2 Rule" -RemoteAddress any -Phase1AuthSet
$myauth.InstanceID `

-InboundSecurity Require -OutboundSecurity Request -KeyModule IKEv2
```

Make sure that you install the required certificates on the participating computers.

Notes

- For local computers, you can import the certificates manually if you have administrator access to the computer. For more information, see Import or export certificates and private keys.
- You need a root certificate and a computer certificate on all computers that participate in the secure connection. Save the computer certificate in the **Personal/Certificates** folder.
- For remote computers, you can create a secure website to facilitate access to the script and certificates.

Troubleshooting

Follow these procedures to verify and troubleshoot your IKEv2 IPsec connections:

Use the Windows Firewall with Advanced Security snap-in to verify that a connection security rule is enabled.

1. On the **Start** screen, type **wf.msc**, and then press ENTER.
2. In the left pane of the Windows Firewall with Advanced Security snap-in, click **Connection Security Rules**, and then verify that there is an enabled connection security rule.
3. Expand **Monitoring**, and then click **Connection Security Rules** to verify that your IKEv2 rule is active for your currently active profile.

Use Windows PowerShell cmdlets to display the security associations.

1. Open a Windows PowerShell command prompt.
2. Type **get-NetIPsecQuickModeSA** to display the Quick Mode security associations.
3. Type **get-NetIPsecMainModeSA** to display the Main Mode security associations.

Use netsh to capture IPsec events.

1. Open an elevated command prompt.
2. At the command prompt, type **netsh wfp capture start**.
3. Reproduce the error event so that it can be captured.
4. At the command prompt, type **netsh wfp capture stop**.

   A wfpdiag.cab file is created in the current folder.
5. Open the cab file, and then extract the wfpdiag.xml file.
6. Open the wfpdiag.xml file with your an XML viewer program or Notepad, and then examine the contents. There will be a lot of data in this file. One way to narrow down where to start looking is to search the last “errorFrequencyTable” at the end of the file. There might be many instances of this table, so make sure that you look at the last table in the file. For example, if you have a certificate problem, you might see the following entry in the last table at the end of the file:

   ```
   <item><error>ERROR_IPSEC_IKE_NO_CERT</error>
   <frequency>32</frequency>
   </item>
   ```

   In this example, there are 32 instances of the **ERROR_IPSEC_IKE_NO_CERT** error. So now you can search for **ERROR_IPSEC_IKE_NO_CERT** to get more details regarding this error.

You might not find the exact answer for the issue, but you can find good hints. For example, you might find that there seems to be an issue with the certificates, so you can look at your certificates and the related cmdlets for possible issues.

See also

- Windows Firewall with Advanced Security Overview

Windows Firewall with Advanced Security Administration with Windows PowerShell

The Windows Firewall with Advanced Security Administration with Windows PowerShell Guide provides essential scriptlets for automating Windows Firewall with Advanced Security management in Windows Server 2012. It is designed for IT pros, system administrators, IT managers, and others who use and need to automate Windows Firewall with Advanced Security management in Windows.

In Windows Server 2012 and Windows 8, administrators can use Windows PowerShell to manage their firewall and IPsec deployments. This object-oriented scripting environment will make it easier for administrators to manage policies and monitor network conditions than was possible in Netsh. Windows PowerShell allows network settings to be self-discoverable through the syntax and parameters in each of the cmdlets. This guide demonstrates how common tasks were performed in Netsh and how you can use Windows PowerShell to accomplish them.

Important
The netsh commands for Windows Firewall with Advanced Security have not changed since the previous operating system version. The netsh commands for Windows Firewall with Advanced Security in Windows Server 2012 are identical to the commands that are provided in Windows Server® 2008 R2.

In future versions of Windows, Microsoft might remove the netsh functionality for Windows Firewall with Advanced Security. Microsoft recommends that you transition to Windows PowerShell if you currently use netsh to configure and manage Windows Firewall with Advanced Security.

Windows PowerShell and netsh command references are at the following locations.

- Netsh Commands for Windows Firewall with Advanced Security

Scope

This guide does not teach you the fundamentals of Windows Firewall with Advanced Security, which can be found in Windows Firewall with Advanced Security Overview. It does not teach the fundamentals of Windows PowerShell, and it assumes that you are familiar with the Windows PowerShell language and the basic concepts of Windows PowerShell. For more information

about Windows PowerShell concepts and usage, see the reference topics in the Additional resources section of this guide.

Audience and user requirements

This guide is intended for IT pros, system administrators, and IT managers, and it assumes that you are familiar with Windows Firewall with Advanced Security, the Windows PowerShell language, and the basic concepts of Windows PowerShell.

System requirements

To run the scripts and scriptlets in this guide, install and configure your system as follows:

- Windows Server 2012
- Windows PowerShell 3.0 (included in Windows Server 2012)
- Windows NetSecurity Module for Windows PowerShell (included in Windows Server 2012)
- Windows PowerShell ISE (optional feature in Windows PowerShell 3.0, which is installed by using Server Manager)

Notes

In Windows PowerShell 3.0, modules are imported automatically when you get or use any cmdlet in the module. You can still use the **Import-Module** cmdlet to import a module.

Use **Import-Module** if you are using Windows PowerShell 2.0, or if you need to use a feature of the module before you use any of its cmdlets. For more information, see Import-Module.

Use **Import-PSSnapIn** to use cmdlets in a Windows PowerShell snap-in, regardless of the version of Windows PowerShell that you are running.

In this guide

Topic	Description
Set profile global defaults	Enable and control firewall behavior
Deploy basic firewall rules	How to create, modify, and delete firewall rules
Manage remotely	Remote management by using `-CimSession`
Deploy basic IPsec rule settings	IPsec rules and associated parameters
Deploy secure firewall rules with IPsec	Domain and server isolation
Additional resources	More information about Windows PowerShell

Set profile global defaults

Global defaults set the system behavior in a per profile basis. Windows Firewall with Advanced Security supports Domain, Private, and Public profiles.

Enable Windows Firewall

Windows Firewall drops traffic that does not correspond to allowed unsolicited traffic, or traffic that is sent in response to a request by the computer. If you find that the rules you create are not being enforced, you may need to enable Windows Firewall. Here is how to do this on a local domain computer:

Netsh

```
netsh advfirewall set allprofiles state on
```

Windows PowerShell

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-NetFirewallProfile -Profile Domain,Public,Private -Enabled True
```

Control firewall behavior

The global default settings can be defined through the command-line interface. These modifications are also available through the Windows Firewall with Advanced Security MMC snap-in.

The following scriptlets set the default inbound and outbound actions, specifies protected network connections, and allows notifications to be displayed to the user when a program is blocked from receiving inbound connections. It allows unicast response to multicast or broadcast network traffic, and it specifies logging settings for troubleshooting.

Netsh

```
netsh advfirewall set allprofiles firewallpolicy blockinbound,allowoutbound

netsh advfirewall set allprofiles settings inboundusernotification enable

netsh advfirewall set allprofiles settings unicastresponsetomulticast enable

netsh advfirewall set allprofiles logging filename
%SystemRoot%\System32\LogFiles\Firewall\pfirewall.log
```

Windows PowerShell

```
Set-NetFirewallProfile -DefaultInboundAction Block -DefaultOutboundAction Allow –
NotifyOnListen True –AllowUnicastResponseToMulticast True –LogFileName
%SystemRoot%\System32\LogFiles\Firewall\pfirewall.log
```

Deploy basic firewall rules

This section provides scriptlet examples for creating, modifying, and deleting firewall rules.

Create firewall rules

Adding a firewall rule in Windows PowerShell looks a lot like it did in Netsh, but the parameters and values are specified differently.

Here is an example of how to allow the Telnet application to listen on the network. This firewall rule is scoped to the local subnet by using a keyword instead of an IP address. Just like in Netsh, the rule is created on the local computer, and it becomes effective immediately.

Netsh

```
netsh advfirewall firewall add rule name="Allow Inbound Telnet" dir=in program=
%SystemRoot%\System32\tlntsvr.exe remoteip=localsubnet action=allow
```

Windows PowerShell

```
New-NetFirewallRule -DisplayName "Allow Inbound Telnet" -Direction Inbound -Program
%SystemRoot%\System32\tlntsvr.exe -RemoteAddress LocalSubnet -Action Allow
```

The following scriptlet shows how to add a basic firewall rule that blocks outbound traffic from a specific application and local port to a Group Policy Object (GPO) in Active Directory. In Windows PowerShell, the policy store is specified as a parameter within the **New-NetFirewall** cmdlet. In Netsh, you must first specify the GPO that the commands in a Netsh session should modify. The commands you enter are run against the contents of the GPO, and this remains in effect until the Netsh session is ended or until another set store command is executed.

Here, **domain.contoso.com** is the name of your Active Directory Domain Services (AD DS), and **gpo_name** is the name of the GPO that you want to modify. Quotation marks are required if there are any spaces in the GPO name.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\gpo_name
netsh advfirewall firewall add rule name="Block Outbound Telnet" dir=out
program=%SystemRoot%\System32\telnet.exe protocol=tcp localport=23 action=block
```

Windows PowerShell

```
New-NetFirewallRule -DisplayName "Block Outbound Telnet" -Direction Outbound -Program
%SystemRoot%\System32\tlntsvr.exe -Protocol TCP -LocalPort 23 -Action Block -PolicyStore
domain.contoso.com\gpo_name
```

GPO Caching

To reduce the burden on busy domain controllers, Windows PowerShell allows you to load a GPO to your local session, make all your changes in that session, and then save it back at all once.

The following performs the same actions as the previous example (by adding a Telnet rule to a GPO), but we do so leveraging GPO caching in PowerShell. Changing the GPO by loading it onto your local session and using the *-GPOSession* parameter are not supported in Netsh

Windows PowerShell

```
$gpo = Open-NetGPO -PolicyStore domain.contoso.com\gpo_name
New-NetFirewallRule -DisplayName "Block Outbound Telnet" -Direction Outbound -Program
%SystemRoot%\System32\telnet.exe -Protocol TCP -LocalPort 23 -Action Block -GPOSession
$gpo
Save-NetGPO -GPOSession $gpo
```

Note that this does not batch your individual changes, it loads and saves the entire GPO at once. So if any other changes are made by other administrators, or in a different Windows PowerShell window, saving the GPO overwrites those changes.

Modify an existing firewall rule

When a rule is created, Netsh and Windows PowerShell allow the administrator to change rule properties and influence, but the rule maintains its unique identifier (in Windows PowerShell this is specified with the *-Name* parameter).

For example, you could have a rule **Allow Web 80** that enables TCP port 80 for inbound unsolicited traffic. You can change the rule to match a different remote IP address of a Web server whose traffic will be allowed by specifying the human-readable, localized name of the rule.

Netsh

```
netsh advfirewall firewall set rule name="Allow Web 80" new remoteip=192.168.0.2
```

Windows PowerShell

```
Set-NetFirewallRule -DisplayName "Allow Web 80" -RemoteAddress 192.168.0.2
```

Netsh requires you to provide the name of the rule for it to be changed and we do not have an alternate way of getting the firewall rule. In Windows PowerShell, you can query for the rule using its known properties.

When you run `Get-NetFirewallRule`, you may notice that common conditions like addresses and ports do not appear. These conditions are represented in separate objects called Filters. As shown before, you can set all the conditions in New-NetFirewallRule and Set-NetFirewallRule. If you want to query for firewall rules based on these fields (ports, addresses, security, interfaces, services), you will need to get the filter objects themselves.

You can change the remote endpoint of the **Allow Web 80** rule (as done previously) using filter objects. Using Windows PowerShell you query by port using the port filter, then assuming additional rules exist affecting the local port, you build with further queries until your desired rule is retrieved.

In the following example, we assume the query returns a single firewall rule, which is then piped to the `Set-NetFirewallRule` cmdlet utilizing Windows PowerShell's ability to pipeline inputs.

Windows PowerShell

```
Get-NetFirewallPortFilter | ?{$_.LocalPort -eq 80} | Get-NetFirewallRule | ?{
$_.Direction -eq "Inbound" -and $_.Action -eq "Allow"} | Set-NetFirewallRule -
RemoteAddress 192.168.0.2
```

You can also query for rules using the wildcard character. The following example returns an array of firewall rules associated with a particular program. The elements of the array can be modified in subsequent `Set-NetFirewallRule` cmdlets.

Windows PowerShell

```
Get-NetFirewallApplicationFilter -Program "*svchost*" | Get-NetFirewallRule
```

Multiple rules in a group can be simultaneously modified when the associated group name is specified in a Set command. You can add firewall rules to specified management groups in order to manage multiple rules that share the same influences.

In the following example, we add both inbound and outbound Telnet firewall rules to the group **Telnet Management**. In Windows PowerShell, group membership is specified when the rules are first created so we re-create the previous example rules. Adding rules to a custom rule group is not possible in Netsh.

Windows PowerShell

```
New-NetFirewallRule -DisplayName "Allow Inbound Telnet" -Direction Inbound -Program
%SystemRoot%\System32\tlntsvr.exe -RemoteAddress LocalSubnet -Action Allow -Group "Telnet
Management"
New-NetFirewallRule -DisplayName "Block Outbound Telnet" -Direction Inbound -Program
%SystemRoot%\System32\tlntsvr.exe -RemoteAddress LocalSubnet -Action Allow -Group "Telnet
Management"
```

If the group is not specified at rule creation time, the rule can be added to the rule group using dot notation in Windows PowerShell. You cannot specify the group using `Set-NetFirewallRule` since the command allows querying by rule group.

Windows PowerShell

```
$rule = Get-NetFirewallRule -DisplayName "Allow Inbound Telnet"

$rule.Group = "Telnet Management"

$rule | Set-NetFirewallRule
```

Using the `Set` command, if the rule group name is specified, the group membership is not modified but rather all rules of the group receive the same modifications indicated by the given parameters.

The following scriptlet enables all rules in a predefined group containing remote management influencing firewall rules.

Netsh

```
netsh advfirewall firewall set rule group="windows firewall remote management" new
enable=yes
```

Windows PowerShell

```
Set-NetFirewallRule -DisplayGroup "Windows Firewall Remote Management" –Enabled True
```

There is also a separate `Enable-NetFirewallRule` cmdlet for enabling rules by group or by other properties of the rule.

Windows PowerShell

```
Enable-NetFirewallRule -DisplayGroup "Windows Firewall Remote Management" -Verbose
```

Delete a firewall rule

Rule objects can be disabled so that they are no longer active. In Windows PowerShell, the **Disable-NetFirewallRule** cmdlet will leave the rule on the system, but put it in a disabled state so the rule no longer is applied and impacts traffic. A disabled firewall rule can be re-enabled by **Enable-NetFirewallRule**. This is different from the **Remove-NetFirewallRule**, which permanently removes the rule definition from the system.

The following cmdlet deletes the specified existing firewall rule from the local policy store.

Netsh

```
netsh advfirewall firewall delete rule name="Allow Web 80"
```

Windows PowerShell

```
Remove-NetFirewallRule –DisplayName "Allow Web 80"
```

Like with other cmdlets, you can also query for rules to be removed. Here, all blocking firewall rules are deleted from the system.

Windows PowerShell

```
Remove-NetFirewallRule –Action Block
```

Note that it may be safer to query the rules with the **Get** command and save it in a variable, observe the rules to be affected, then pipe them to the **Remove** command, just as we did for the

Set commands. The following example shows how the administrator can view all the blocking firewall rules, and then delete the first four rules.

Windows PowerShell

```
$x = Get-NetFirewallRule -Action Block

$x

$x[0-3] | Remove-NetFirewallRule
```

Manage remotely

Remote management using WinRM is enabled by default on Windows Server 2012. The cmdlets that support the *CimSession* parameter use WinRM and can be managed remotely by default. This is important because the default and recommended installation mode for Windows Server 2012 is Server Core which does not include a graphical user interface.

The following example returns all firewall rules of the persistent store on a computer named **RemoteComputer**.

Windows PowerShell

```
Get-NetFirewallRule -CimSession RemoteComputer
```

We can perform any modifications or view rules on remote computers by simply using the –*CimSession* parameter. Here we remove a specific firewall rule from a remote computer.

Windows PowerShell

```
$RemoteSession = New-CimSession -ComputerName RemoteComputer

Remove-NetFirewallRule -DisplayName "AllowWeb80" -CimSession $RemoteSession -Confirm
```

Deploy basic IPsec rule settings

An Internet Protocol security (IPsec) policy consists of rules that determine IPsec behavior. IPsec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection. For more information about IPsec, see Windows Firewall with Advanced Security Learning Roadmap.

Windows PowerShell can create powerful, complex IPsec policies like in Netsh and the Windows Firewall with Advanced Security MMC snap-in. However, because Windows PowerShell is object-based rather than string token-based, configuration in Windows PowerShell offers greater control and flexibility.

In Netsh, the authentication and cryptographic sets were specified as a list of comma-separated tokens in a specific format. In Windows PowerShell, rather than using default settings, you first create your desired authentication or cryptographic proposal objects and bundle them into lists in your preferred order. Then, you create one or more IPsec rules that reference these sets. The benefit of this model is that programmatic access to the information in the rules is much easier. See the following sections for clarifying examples.

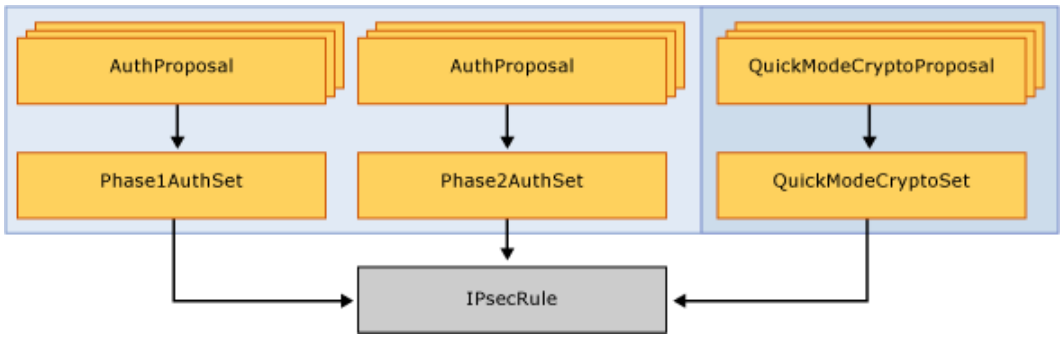

Create IPsec rules

The following cmdlet creates basic IPsec transport mode rule in a Group Policy Object. An IPsec rule is simple to create; all that is required is the display name, and the remaining properties use default values. Inbound traffic is authenticated and integrity checked using the default quick mode and main mode settings. These default settings can be found in the MMC snap-in under Customize IPsec Defaults.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\gpo_name
```

```
netsh advfirewall consec add rule name="Require Inbound Authentication" endpoint1=any endpoint2=any action=requireinrequestout
```

Windows PowerShell

```
New-NetIPsecRule -DisplayName "Require Inbound Authentication" -PolicyStore domain.contoso.com\gpo_name
```

Add custom authentication methods to an IPsec rule

If you want to create a custom set of quick-mode proposals that includes both AH and ESP in an IPsec rule object, you create the associated objects separately and link their associations. For more information about authentication methods, see Choosing the IPsec Protocol .

You can then use the newly created custom quick-mode policies when you create IPsec rules. The cryptography set object is linked to an IPsec rule object.

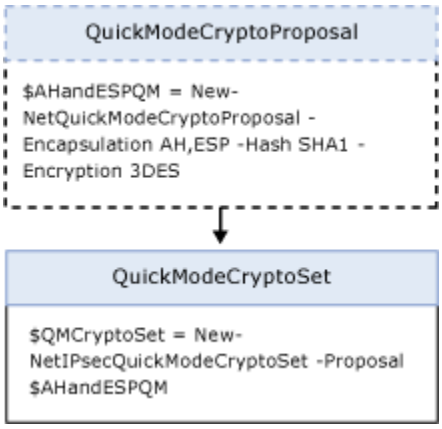

In this example, we build on the previously created IPsec rule by specifying a custom quick-mode crypto set. The final IPsec rule requires outbound traffic to be authenticated by the specified cryptography method.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\gpo_name

netsh advfirewall consec add rule name="Require Outbound Authentication" endpoint1=any endpoint2=any action=requireinrequestout qmsecmethods=ah:sha1+esp:sha1-3des
```

Windows PowerShell

```
$AHandESPQM = New-NetIPsecQuickModeCryptoProposal -Encapsulation AH,ESP –AHHash SHA1 -ESPHash SHA1 -Encryption DES3

$QMCryptoSet = New-NetIPsecQuickModeCryptoSet –DisplayName “ah:sha1+esp:sha1-des3” -Proposal $AHandESPQM –PolicyStore domain.contoso.com\gpo_name

New-NetIPsecRule -DisplayName “Require Inbound Authentication” -InboundSecurity Require -OutboundSecurity Request -QuickModeCryptoSet $QMCryptoSet.Name –PolicyStore domain.contoso.com\gpo_name
```

IKEv2 IPsec transport rules

A corporate network may need to secure communications with another agency. But, you discover the agency runs non-Windows operating systems and requires the use of the Internet Key Exchange Version 2 (IKEv2) standard.

You can leverage IKEv2 capabilities in Windows Server 2012 by simply specifying IKEv2 as the key module in an IPsec rule. This can only be done using computer certificate authentication and cannot be used with phase 2 authentication.

Windows PowerShell

```
New-NetIPsecRule -DisplayName "Require Inbound Authentication" -InboundSecurity Require -
OutboundSecurity Request -Phase1AuthSet MyCertAuthSet -KeyModule IKEv2 -RemoteAddress
$nonWindowsGateway
```

For more information about IKEv2, including scenarios, see Securing End-to-End IPsec Connections by Using IKEv2 in Windows Server 2012.

Copy an IPsec rule from one policy to another

Firewall and IPsec rules with the same rule properties can be duplicated to simplify the task of re-creating them within different policy stores.

To copy the previously created rule from one policy store to another, the associated objects must be also be copied separately. Note that there is no need to copy associated firewall filters. You can query rules to be copied in the same way as other cmdlets.

Copying individual rules is a task that is not possible through the Netsh interface. Here is how you can accomplish it with Windows PowerShell.

Windows PowerShell

```
$Rule = Get-NetIPsecRule -DisplayName "Require Inbound Authentication"
$Rule | Copy-NetIPsecRule -NewPolicyStore domain.costoso.com\new_gpo_name
$Rule | Copy-NetPhase1AuthSet -NewPolicyStore domain.costoso.com\new_gpo_name
```

Handling Windows PowerShell errors

To handle errors in your Windows PowerShell scripts, you can use the *–ErrorAction* parameter. This is especially useful with the **Remove** cmdlets. If you want to remove a particular rule, you will notice that it fails if the rule is not found. When removing rules, if the rule isn't already there, it is generally acceptable to ignore that error. In this case, you can do the following to suppress any "rule not found" errors during the remove operation.

Windows PowerShell

```
Remove-NetFirewallRule -DisplayName "Contoso Messenger 98" -ErrorAction SilentlyContinue
```

Note that the use of wildcards can also suppress errors, but they could potentially match rules that you did not intend to remove. This can be a useful shortcut, but should only be used if you know there aren't any extra rules that will be accidentally deleted. So the following cmdlet will also remove the rule, suppressing any "not found" errors.

Windows PowerShell

```
Remove-NetFirewallRule -DisplayName "Contoso Messenger 98*"
```

When using wildcards, if you want to double-check the set of rules that is matched, you can use the *–WhatIf* parameter.

Windows PowerShell

```
Remove-NetFirewallRule –DisplayName "Contoso Messenger 98*" –WhatIf
```

If you only want to delete some of the matched rules, you can use the *–Confirm* parameter to get a rule-by-rule confirmation prompt.

Windows PowerShell

```
Remove-NetFirewallRule –DisplayName "Contoso Messenger 98*" –Confirm
```

You can also just perform the whole operation, displaying the name of each rule as the operation is performed.

Windows PowerShell

```
Remove-NetFirewallRule –DisplayName "Contoso Messenger 98*" –Verbose
```

Monitor

The following Windows PowerShell commands are useful in the update cycle of a deployment phase.

To allow you to view all the IPsec rules in a particular store, you can use the following commands. In Netsh, this command does not show rules where profile=domain,public or profile=domain,private. It only shows rules that have the single entry domain that is included in the rule. The following command examples will show the IPsec rules in all profiles.

Netsh

```
netsh advfirewall consec show rule name=all
```

Windows PowerShell

```
Show-NetIPsecRule –PolicyStore ActiveStore
```

You can monitor main mode security associations for information such as which peers are currently connected to the computer and which protection suite is used to form the security associations.

Use the following cmdlet to view existing main mode rules and their security associations:

Netsh

```
netsh advfirewall monitor show mmsa all
```

Windows PowerShell

```
Get-NetIPsecMainModeSA
```

Find the source GPO of a rule

To view the properties of a particular rule or group of rules, you query for the rule. When a query returns fields that are specified as **NotConfigured**, you can to determine which policy store a rule originates from.

For objects that come from a GPO (the *–PolicyStoreSourceType* parameter is specified as **GroupPolicy** in the **Show** command), if *–TracePolicyStore* is passed, the name of the GPO is found and returned in the **PolicyStoreSource** field.

Windows PowerShell

```
Get-NetIPsecRule –DisplayName “Require Inbound Authentication” –TracePolicyStore
```

It is important to note that the revealed sources do not contain a domain name.

Deploy a basic domain isolation policy

IPsec can be used to isolate domain members from non-domain members. Domain isolation uses IPsec authentication to require that the domain computer members positively establish the identities of the communicating computers to improve security of an organization. One or more features of IPsec can be used to secure traffic with an IPsec rule object.

To implement domain isolation on your network, the computers in the domain receive IPsec rules that block unsolicited inbound network traffic that is not protected by IPsec. Here we create an IPsec rule that requires authentication by domain members. Through this, you can isolate domain member computers from computers that are non-domain members. In the following examples, Kerberos authentication is required for inbound traffic and requested for outbound traffic.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\domain_isolation

netsh advfirewall consec add rule name=“Basic Domain Isolation Policy” profile=domain
endpoint1=”any” endpoint2=”any” action=requireinrequestout auth1=”computerkerb”
```

Windows PowerShell

```
$kerbprop = New-NetIPsecAuthProposal –Machine –Kerberos


$Phase1AuthSet = New-NetIPsecPhase1AuthSet -DisplayName "Kerberos Auth Phase1" -Proposal
$kerbprop –PolicyStore domain.contoso.com\domain_isolation


New-NetIPsecRule –DisplayName “Basic Domain Isolation Policy” –Profile Domain –
Phase1AuthSet $Phase1AuthSet.Name –InboundSecurity Require –OutboundSecurity Request –
PolicyStore domain.contoso.com\domain_isolation
```

Configure IPsec tunnel mode

The following command creates an IPsec tunnel that routes traffic from a private network (192.168.0.0/16) through an interface on the local computer (1.1.1.1) attached to a public network

to a second computer through its public interface (2.2.2.2) to another private network (192.157.0.0/16). All traffic through the tunnel is checked for integrity by using ESP/SHA1, and it is encrypted by using ESP/DES3.

Netsh

```
netsh advfirewall consec add rule name="Tunnel from 192.168.0.0/16 to 192.157.0.0/16"
mode=tunnel endpoint1=192.168.0.0/16 endpoint2=192.157.0.0/16 localtunnelendpoint=1.1.1.1
remotetunnelendpoint=2.2.2.2 action=requireinrequireout qmsecmethods=esp:sha1-3des
```

Windows PowerShell

```
$QMProposal = New-NetIPsecQuickModeCryptoProposal -Encapsulation ESP -ESPHash SHA1 -
Encryption DES3
$QMCryptoSet = New-NetIPsecQuickModeCryptoSet –DisplayName “esp:sha1-des3” -Proposal
$QMProposal
New-NetIPSecRule -DisplayName “Tunnel from HQ to Dallas Branch” -Mode Tunnel -
LocalAddress 192.168.0.0/16 -RemoteAddress 192.157.0.0/16 -LocalTunnelEndpoint 1.1.1.1 -
RemoteTunnelEndpoint 2.2.2.2 -InboundSecurity Require -OutboundSecurity Require -
QuickModeCryptoSet $QMCryptoSet.Name
```

Deploy secure firewall rules with IPsec

In situations where only secure traffic can be allowed through the Windows Firewall, a combination of manually configured firewall and IPsec rules are necessary. The firewall rules determine the level of security for allowed packets, and the underlying IPsec rules secure the traffic. The scenarios can be accomplished in Windows PowerShell and in Netsh, with many similarities in deployment.

Create a secure firewall rule (allow if secure)

Configuring firewalls rule to allow connections if they are secure requires the corresponding traffic to be authenticated and integrity protected, and then optionally encrypted by IPsec.

The following example creates a firewall rule that requires traffic to be authenticated. The command permits inbound Telnet network traffic only if the connection from the remote computer is authenticated by using a separate IPsec rule.

Netsh

```
netsh advfirewall firewall add rule name="Allow Authenticated Telnet" dir=in
program=%SystemRoot%\System32\tlntsvr.exe security=authenticate action=allow
```

Windows PowerShell

```
New-NetFirewallRule -DisplayName “Allow Authenticated Telnet” -Direction Inbound -Program
%SystemRoot%\System32\tlntsvr.exe -Authentication Required -Action Allow
```

The following command creates an IPsec rule that requires a first (computer) authentication and then attempts an optional second (user) authentication. Creating this rule secures and allows the traffic through the firewall rule requirements for the messenger program.

Netsh

```
netsh advfirewall consec add rule name="Authenticate Both Computer and User"
endpoint1=any endpoint2=any action=requireinrequireout auth1=computerkerb,computerntlm
auth2=userkerb,userntlm,anonymous
```

Windows PowerShell

```
$mkerbauthprop = New-NetIPsecAuthProposal -Machine –Kerberos
$mntlmauthprop = New-NetIPsecAuthProposal -Machine -NTLM
$P1Auth = New-NetIPsecPhase1AuthSet -DisplayName “Machine Auth” –Proposal
$mkerbauthprop,$mntlmauthprop
$ukerbauthprop = New-NetIPsecAuthProposal -User -Kerberos
$unentlmauthprop = New-NetIPsecAuthProposal -User -NTLM
$anonyauthprop = New-NetIPsecAuthProposal -Anonymous
$P2Auth = New-NetIPsecPhase2AuthSet -DisplayName “User Auth” -Proposal
$ukerbauthprop,$unentlmauthprop,$anonyauthprop
New-NetIPSecRule -DisplayName “Authenticate Both Computer and User” -InboundSecurity
Require -OutboundSecurity Require -Phase1AuthSet $P1Auth.Name –Phase2AuthSet $P2Auth.Name
```

Isolate a server by requiring encryption and group membership

To improve the security of the computers in an organization, an administrator can deploy domain isolation in which domain-members are restricted. They require authentication when communicating among each other and reject non-authenticated inbound connections. To improve the security of servers with sensitive data, this data must be protected by allowing access only to a subset of computers within the enterprise domain.

IPsec can provide this additional layer of protection by isolating the server. In server isolation, sensitive data access is restricted to users and computers with legitimate business need, and the data is additionally encrypted to prevent eavesdropping.

Create a firewall rule that requires group membership and encryption

To deploy server isolation, we layer a firewall rule that restricts traffic to authorized users or computers on the IPsec rule that enforces authentication.

The following firewall rule allows Telnet traffic from user accounts that are members of a custom group created by an administrator called “Authorized to Access Server.” This access can additionally be restricted based on the computer, user, or both by specifying the restriction parameters.

A Security Descriptor Definition Language (SDDL) string is created by extending a user or group’s security identifier (SID). For more information about finding a group’s SID, see: Finding the SID for a group account.

Restricting access to a group allows administrations to extend strong authentication support through Windows Firewall/and or IPsec policies.

The following example shows you how to create an SDDL string that represents security groups.

Windows PowerShell

```
$user = new-object System.Security.Principal.NTAccount
("corp.contoso.com\Administrators")
$SIDofSecureUserGroup =
$user.Translate([System.Security.Principal.SecurityIdentifier]).Value
$secureUserGroup = "D:(A;;CC;;;$SIDofSecureUserGroup)"
```

By using the previous scriptlet, you can also get the SDDL string for a secure computer group as shown here:

Windows PowerShell

```
$secureMachineGroup = "D:(A;;CC;;;$SIDofSecureMachineGroup)"
```

For more information about how to create security groups or how to determine the SDDL string, see Working with SIDs.

Telnet is an application that does not provide encryption. This application can send data, such as names and passwords, over the network. This data can be intercepted by malicious users. If an administrator would like to allow the use of Telnet, but protect the traffic, a firewall rule that requires IPsec encryption can be created. This is necessary so that the administrator can be certain that when this application is used, all of the traffic sent or received by this port is encrypted. If IPsec fails to authorize the connection, no traffic is allowed from this application.

In this example, we allow only authenticated and encrypted inbound Telnet traffic from a specified secure user group through the creation of the following firewall rule.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\Server_Isolation
```

```
netsh advfirewall firewall add rule name="Allow Encrypted Inbound Telnet to Group Members Only" program=%SystemRoot%\System32\tlntsvr.exe protocol=TCP dir=in action=allow localport=23 security=authenc rmtusrgrp ="D:(A;;CC;;; S-1-5-21-2329867823-2610410949-1491576313-1735)"
```

Windows PowerShell

```
New-NetFirewallRule -DisplayName "Allow Encrypted Inbound Telnet to Group Members Only" -Program %SystemRoot%\System32\tlntsvr.exe -Protocol TCP -Direction Inbound -Action Allow -LocalPort 23 -Authentication Required -Encryption Required –RemoteUser $secureUserGroup –PolicyStore domain.contoso.com\Server_Isolation
```

Endpoint security enforcement

The previous example showed end to end security for a particular application. In situations where endpoint security is required for many applications, having a firewall rule per application can be cumbersome and difficult to manage. Authorization can override the per-rule basis and be done at the IPsec layer.

In this example, we set the global IPsec setting to only allow transport mode traffic to come from an authorized user group with the following cmdlet. Consult the previous examples for working with security groups.

Windows PowerShell

```
Set-NetFirewallSetting -RemoteMachineTransportAuthorizationList $secureMachineGroup
```

Create firewall rules that allow IPsec-protected network traffic (authenticated bypass)

Authenticated bypass allows traffic from a specified trusted computer or user to override firewall block rules. This is helpful when an administrator wants to use scanning servers to monitor and update computers without the need to use port-level exceptions. For more information, see How to enable authenticated firewall bypass.

In this example, we assume that a blocking firewall rule exists. This example permits any network traffic on any port from any IP address to override the block rule, if the traffic is authenticated as originating from a computer or user account that is a member of the specified computer or user security group.

Netsh

```
netsh advfirewall set store gpo=domain.contoso.com\domain_isolation

netsh advfirewall firewall add rule name="Inbound Secure Bypass Rule" dir=in security=authenticate action="bypass" rmtcomputergrp="D:(A;;CC;;;S-1-5-21-2329867823-2610410949-1491576313-1114)" rmtusrgrp="D:(A;;CC;;; S-1-5-21-2329867823-2610410949-1491576313-1735)"
```

Windows PowerShell

```
New-NetFirewallRule –DisplayName "Inbound Secure Bypass Rule" –Direction Inbound –
Authentication Required –OverrideBlockRules $true -RemoteMachine $secureMachineGroup –
RemoteUser $secureUserGroup –PolicyStore domain.contoso.com\domain_isolation
```

Additional resources

For more information about Windows PowerShell concepts, see the following topics.

- Windows PowerShell Getting Started Guide
- Windows PowerShell User Guide
- Windows PowerShell About Help Topics
- about_Functions
- about_Functions_Advanced
- about_Execution_Policies
- about_Foreach
- about_Objects
- about_Properties
- about_While
- about_Scripts
- about_Signing
- about_Throw
- about_PSSessions
- about_Modules
- about_Command_Precedence

Windows Firewall with Advanced Security Design Guide

Windows Firewall with Advanced Security in Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista is a host firewall that helps secure the computer in two ways. First, it can filter the network traffic permitted to enter the computer from the network, and also control what network traffic the computer is allowed to send to the network. Second, Windows Firewall with Advanced Security supports IPsec, which enables you to require authentication from any computer that is attempting to communicate with your computer. When authentication is required, computers that cannot authenticate cannot communicate with your computer. By using IPsec, you can also require that specific network

traffic be encrypted to prevent it from being read or intercepted while in transit between computers.

The interface for Windows Firewall with Advanced Security is much more capable and flexible than the consumer-friendly interface found in the Windows Firewall Control Panel. They both interact with the same underlying services, but provide different levels of control over those services. While the Windows Firewall Control Panel meets the needs for protecting a single computer in a home environment, it does not provide enough centralized management or security features to help secure more complex network traffic found in a typical business enterprise environment.

For more overview information about Windows Firewall with Advanced Security and see **Windows Firewall with Advanced Security Overview**.

About this guide

This guide provides recommendations to help you to choose or create a design for deploying Windows Firewall with Advanced Security in your enterprise environment. The guide describes some of the common goals for using Windows Firewall with Advanced Security, and then helps you map the goals that apply to your scenario to the designs that are presented in this guide.

This guide is intended for the IT professional who has been assigned the task of deploying firewall and IPsec technologies on an organization's network to help meet the organization's security goals.

Windows Firewall with Advanced Security should be part of a comprehensive security solution that implements a variety of security technologies, such as perimeter firewalls, intrusion detection systems, virtual private networking (VPN), IEEE 802.1X authentication for wireless and wired connections, and IPsec connection security rules.

To successfully use this guide, you need a good understanding of both the capabilities provided by Windows Firewall with Advanced Security, and how to deliver configuration settings to your managed computers by using Group Policy in Active Directory.

You can use the deployment goals to form one of these Windows Firewall with Advanced Security designs, or a custom design that combines elements from those presented here:

- **Basic firewall policy design**. Restricts network traffic in and out of your computers to only that which is needed and authorized.
- **Domain isolation policy design**. Prevents computers that are domain members from receiving unsolicited network traffic from computers that are not domain members. Additional "zones" can be established to support the special requirements of some computers, such as:
  - A "boundary zone" for computers that must be able to receive requests from non-isolated computers.
  - An "encryption zone" for computers that store sensitive data that must be protected during network transmission.
- **Server isolation policy design**. Restricts access to a server to only a limited group of authorized users and computers. Commonly configured as a zone in a domain isolation

design, but can also be configured as a stand-alone design, providing many of the benefits of domain isolation to a small set of computers.

- **Certificate-based isolation policy design**. This design is a complement to either of the previous two designs, and supports any of their capabilities. It uses cryptographic certificates that are deployed to clients and servers for authentication, instead of the Kerberos V5 authentication used by default in Active Directory. This enables computers that are not part of an Active Directory domain, such as computers running operating systems other than Windows, to participate in your isolation solution.

In addition to descriptions and example for each design, you will find guidelines for gathering required data about your environment. You can then use these guidelines to plan and design your Windows Firewall with Advanced Security deployment. After you read this guide, and finish gathering, documenting, and mapping your organization's requirements, you have the information that you need to begin deploying Windows Firewall with Advanced Security using the guidance in the Windows Firewall with Advanced Security Deployment Guide.

You can find the Windows Firewall with Advanced Security Deployment Guide at these locations:

- (Web page)
- (Downloadable Word document)

Terminology used in this guide

The following table identifies and defines terms used throughout this guide.

Term	Definition
Active Directory domain	A group of computers and users managed by an administrator by using Active Directory Domain Services (AD DS). Computers in a domain share a common directory database and security policies. Multiple domains can co-exist in a "forest," with trust relationships that establish the forest as the security boundary.
Authentication	A process that enables the sender of a message to prove its identity to the receiver. For connection security in Windows, authentication is implemented by the IPsec protocol suite.
Boundary zone	A subset of the computers in an isolated domain that must be able to receive unsolicited and non-authenticated network traffic from computers that are not members of the isolated domain. Computers in the boundary zone request but do not require authentication. They

Term	Definition
	use IPsec to communicate with other computers in the isolated domain.
Connection security rule	A rule in Windows Firewall with Advanced Security that contains a set of conditions and an action to be applied to network packets that match the conditions. The action can allow the packet, block the packet, or require the packet to be protected by IPsec. In previous versions of Windows, this was called an *IPsec rule*.
Certificate-based isolation	A way to add computers that cannot use Kerberos V5 authentication to an isolated domain, by using an alternate authentication technique. Every computer in the isolated domain and the computers that cannot use Kerberos V5 are provided with a computer certificate that can be used to authenticate with each other. Certificate-based isolation requires a way to create and distribute an appropriate certificate (if you choose not to purchase one from a commercial certificate provider).
Domain isolation	A technique for helping protect the computers in an organization by requiring that the computers authenticate each other's identity before exchanging information, and refusing connection requests from computers that cannot authenticate. Domain isolation takes advantage of Active Directory domain membership and the Kerberos V5 authentication protocol available to all members of the domain. Also see "Isolated domain" in this table.
Encryption zone	A subset of the computers in an isolated domain that process sensitive data. Computers that are part of the encryption zone have all network traffic encrypted to prevent viewing by non-authorized users. Computers that are part of the encryption zone also typically are subject to the access control restrictions of server isolation.

Term	Definition
Firewall rule	A rule in Windows Firewall with Advanced Security that contains a set of conditions used to determine whether a network packet is allowed to pass through the firewall. By default, the firewall rules in Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista block unsolicited inbound network traffic. Likewise, by default, all outbound network traffic is allowed. The firewall included in previous versions of Windows only filtered inbound network traffic.
Internet Protocol security (IPsec)	A set of industry-standard, cryptography-based protection services and protocols. IPsec protects all protocols in the TCP/IP protocol suite except Address Resolution Protocol (ARP).
IPsec policy	A collection of connection security rules that provide the required protection to network traffic entering and leaving the computer. The protection includes authentication of both the sending and receiving computer, integrity protection of the network traffic exchanged between them, and can include encryption.
Isolated domain	An Active Directory domain (or an Active Directory forest, or set of domains with two-way trust relationships) that has Group Policy settings applied to help protect its member computers by using IPsec connection security rules. Members of the isolated domain require authentication on all unsolicited inbound connections (with exceptions handled by the other zones). In this guide, the term *isolated domain* refers to the IPsec concept of a group of computers that can share authentication. The term *Active Directory domain* refers to the group of computers that share a security database by using Active Directory.

Term	Definition
Server isolation	A technique for using group membership to restrict access to a server that is typically already a member of an isolated domain. The additional protection comes from using the authentication credentials of the requesting computer to determine its group membership, and then only allowing access if the computer account (and optionally the user account) is a member of an authorized group.
Solicited network traffic	Network traffic that is sent in response to a request. By default, Windows Firewall with Advanced Security allows all solicited network traffic through.
Unsolicited network traffic	Network traffic that is not a response to an earlier request, and that the receiving computer cannot necessarily anticipate. By default, Windows Firewall with Advanced Security blocks all unsolicited network traffic.
Zone	A zone is a logical grouping of computers that share common IPsec policies because of their communications requirements. For example, the boundary zone permits inbound connections from non-trusted computers. The encryption zone requires that all connections be encrypted. This is not related to the term zone as used by Domain Name System (DNS).

Next: Understanding the Windows Firewall with Advanced Security Design Process

Understanding the Windows Firewall with Advanced Security Design Process

Designing any deployment starts by performing several important tasks:

- **Identifying Your Windows Firewall with Advanced Security Design Goals**
- **Mapping Your Deployment Goals to a Windows Firewall with Advanced Security Design**

- **Evaluating Windows Firewall with Advanced Security Design Examples**

After you identify your deployment goals and map them to a Windows Firewall with Advanced Security design, you can begin documenting the design based on the processes that are described in the following topics:

- **Designing A Windows Firewall with Advanced Security Strategy**
- **Planning Your Windows Firewall with Advanced Security Design**

Next: Identifying Your Windows Firewall with Advanced Security Design Goals

Identifying Your Windows Firewall with Advanced Security Deployment Goals

Correctly identifying your Windows Firewall with Advanced Security deployment goals is essential for the success of your Windows Firewall with Advanced Security design project. Form a project team that can clearly articulate deployment issues in a vision statement. When you write your vision statement, identify, clarify, and refine your deployment goals. Prioritize and, if possible, combine your deployment goals so that you can design and deploy Windows Firewall with Advanced Security by using an iterative approach. You can take advantage of the predefined Windows Firewall with Advanced Security deployment goals presented in this guide that are relevant to your scenarios.

The following table lists the three main tasks for articulating, refining, and subsequently documenting your Windows Firewall with Advanced Security deployment goals.

Deployment goal tasks	Reference links
Evaluate predefined Windows Firewall with Advanced Security deployment goals that are provided in this section of the guide, and combine one or more goals to reach your organizational objectives.	Predefined deployment goals: • **Protect Computers from Unwanted Network Traffic** • **Restrict Access to Only Trusted Computers** • **Require Encryption When Accessing Sensitive Network Resources** • **Restrict Access to Sensitive Resources to Only Specified Users or Computers**
Map one goal or a combination of the predefined deployment goals to an existing Windows Firewall with Advanced Security design.	• **Mapping Your Deployment Goals to a Windows Firewall with Advanced Security Design**
Based on the status of your current infrastructure, document your deployment goals	• **Designing A Windows Firewall with Advanced Security Strategy**

Deployment goal tasks	Reference links
for your Windows Firewall with Advanced Security design into a deployment plan.	• **Planning Your Windows Firewall with Advanced Security Design**

Next: Protect Computers from Unwanted Network Traffic

Protect Computers from Unwanted Network Traffic

Although network perimeter firewalls provide important protection to network resources from external threats, there are network threats that a perimeter firewall cannot protect against. Some attacks might successfully penetrate the perimeter firewall, and at that point what can stop it? Other attacks might originate from inside the network, such as a computer virus that is brought in on portable media and run on a trusted computer. Portable computers are often taken outside the network and connected directly to the Internet, without adequate protection between the computer and security threats.

Reports of targeted attacks against organizations, governments, and individuals have become more widespread in recent years. For a general overview of these threats, also known as advanced persistent threats (APT), see the Microsoft Security Intelligence Report at http://download.microsoft.com/download/C/9/A/C9A544AD-4150-43D3-80F7-4F1641EF910A/Microsoft_Security_Intelligence_Report_Volume_12_Key_Findings_Summary_English.pdf.

Running a host-based firewall on every computer that your organization manages is an important layer in a "defense-in-depth" security strategy. A host-based firewall can help protect against attacks that originate from inside the network and also provide additional protection against attacks from outside the network that manage to penetrate the perimeter firewall. It also travels with a portable computer to provide protection when it is away from the organization's network.

A host-based firewall helps secure a computer by dropping all network traffic that does not match the administrator-designed rule set for permitted network traffic. This design, which corresponds to **Basic Firewall Policy Design**, provides the following benefits:

- Network traffic that is a reply to a request from the local computer is permitted into the computer from the network.
- Network traffic that is unsolicited, but that matches a rule for allowed network traffic, is permitted into the computer from the network.

  For example, Woodgrove Bank wants a computer that is running SQL Server to be able to receive the SQL queries sent to it by client computers. The firewall policy deployed to the computer that is running SQL Server includes firewall rules that specifically allow inbound network traffic for the SQL Server program.
- Outbound network traffic that is not specifically blocked is allowed on the network.

For example, Woodgrove Bank has a corporate policy that prohibits the use of certain peer-to-peer file sharing programs. The firewall policy deployed to the computers on the network includes firewall rules that block both inbound and outbound network traffic for the prohibited programs. All other outbound traffic is permitted.

The following component is recommended for this deployment goal:

- **Active Directory**: Active Directory supports centralized management of connection security rules by configuring the rules in one or more Group Policy objects (GPOs) that can be automatically applied to all relevant computers in the domain. For more information about Active Directory, see **Additional Resources [lhs]**.

Other means of deploying a firewall policy are available, such as creating scripts that use the **netsh** command-line tool, and then running those scripts on each computer in the organization. This guide uses Active Directory as a recommended means of deployment because of its ability to scale to very large organizations.

Next: Restrict Access to Only Trusted Computers

Restrict Access to Only Trusted Computers

Your organizational network likely has a connection to the Internet. You also likely have partners, vendors, or contractors who attach computers that are not owned by your organization to your network. Because you do not manage those computers, you cannot trust them to be free of malicious software, maintained with the latest security updates, or in any way in compliance with your organization's security policies. These untrustworthy computers both on and outside of your physical network must not be permitted to access your organization's computers except where it is truly required.

To mitigate this risk, you must be able to isolate the computers you trust, and restrict their ability to receive unsolicited network traffic from untrusted computers. By using connection security and firewall rules available in Windows Firewall with Advanced Security, you can logically isolate the computers that you trust by requiring that all unsolicited inbound network traffic be authenticated. Authentication ensures that each computer or user can positively identify itself by using credentials that are trusted by the other computer. Connection security rules can be configured to use IPsec with the Kerberos V5 protocol available in Active Directory, or certificates issued by a trusted certification authority as the authentication method.

Note

Because the primary authentication method recommended for computers that are running Windows is to use the Kerberos V5 protocol with membership in an Active Directory domain, this guide refers to this logical separation of computers as *domain isolation*, even when certificates are used to extend the protection to computers that are not part of an Active Directory domain.

The protection provided by domain isolation can help you comply with regulatory and legislative requirements, such as those found in the Federal Information Security Management Act of 2002

(FISMA), the Sarbanes-Oxley Act of 2002, the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and other government and industry regulations.

The following illustration shows an isolated domain, with one of the zones that are optionally part of the design. The rules that implement both the isolated domain and the different zones are deployed by using Group Policy and Active Directory.

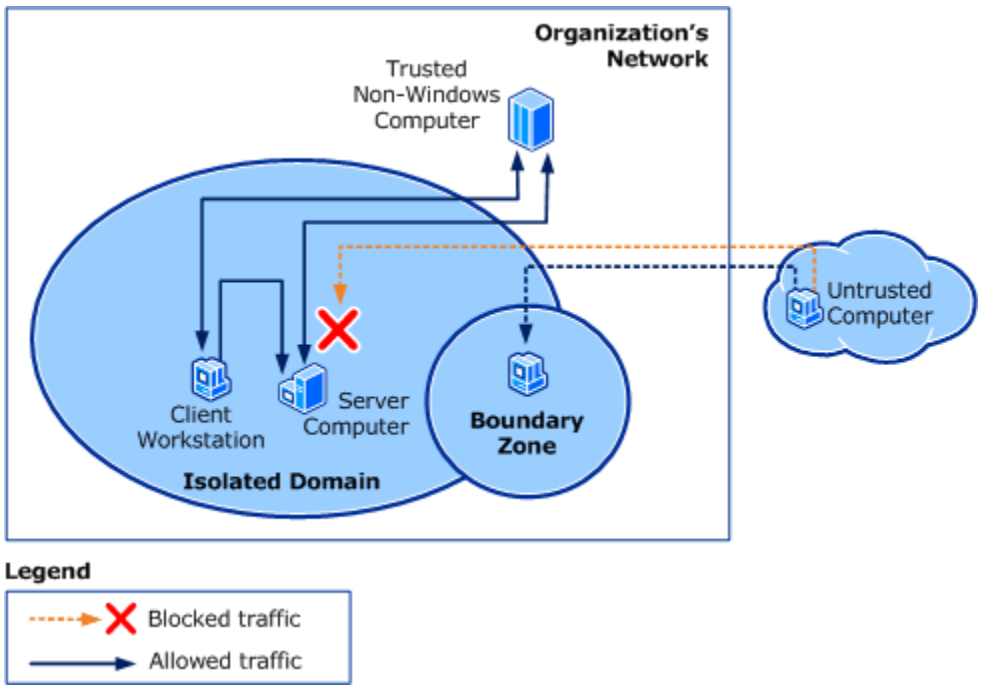

These goals, which correspond to **Domain Isolation Policy Design** and **Certificate-based Isolation Policy Design**, provide the following benefits:

- Computers in the isolated domain accept unsolicited inbound network traffic only when it can be authenticated as coming from another computer in the isolated domain. Exemption rules can be defined to allow inbound traffic from trusted computers that for some reason cannot perform IPsec authentication.

  For example, Woodgrove Bank wants all of its computers to block all unsolicited inbound network traffic from any computer that it does not manage. The connection security rules deployed to domain member computers require authentication as a domain member or by using a certificate before an unsolicited inbound network packet is accepted.

- Computers in the isolated domain can still send outbound network traffic to untrusted computers and receive the responses to the outbound requests.

  For example, Woodgrove Bank wants its users at client computers to be able to access Web sites on the Internet. The default Windows Firewall with Advanced Security settings for outbound network traffic allow this. No additional rules are required.

These goals also support optional zones that can be created to add customized protection to meet the needs of subsets of an organization's computers:

- Computers in the "boundary zone" are configured to use connection security rules that request but do not require authentication. This enables them to receive unsolicited inbound network traffic from untrusted computers, and also to receive traffic from the other members of the isolated domain.

  For example, Woodgrove Bank has a server that must be accessed by its partners' computers through the Internet. The rules applied to computers in the boundary zone use authentication when the client computer can support it, but do not block the connection if the client computer cannot authenticate.
- Computers in the "encryption zone" require that all network traffic in and out must be encrypted to secure potentially sensitive material when it is sent over the network.

  For example, Woodgrove Bank wants the computers running SQL Server to only transmit data that is encrypted to help protect the sensitive data stored on those computers.

The following components are required for this deployment goal:

- **Active Directory**: Active Directory supports centralized management of connection security rules by configuring the rules in one or more GPOs that can be automatically applied to all relevant computers in the domain. For more information about Active Directory, see **Additional Resources [lhs]**.

Next: Require Encryption When Accessing Sensitive Network Resources

Require Encryption When Accessing Sensitive Network Resources

The use of authentication in the previously described goal (Restrict Access to Only Trusted Computers) enables a computer in the isolated domain to block traffic from untrusted computers. However, it does not prevent an untrusted computer from eavesdropping on the network traffic shared between two trusted computers, because by default network packets are not encrypted.

For computers that share sensitive information over the network, Windows Firewall with Advanced Security allows you to require that all such network traffic be encrypted. Using encryption can help you comply with regulatory and legislative requirements such as those found in the Federal Information Security Management Act of 2002 (FISMA), the Sarbanes-Oxley Act of 2002, the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and other government and industry regulations. By creating connection security rules that apply to computers that host and exchange sensitive data, you can help protect the confidentiality of that data by encrypting it.

The following illustration shows an encryption zone in an isolated domain. The rules that implement both the isolated domain and the different zones are deployed by using Group Policy and Active Directory.

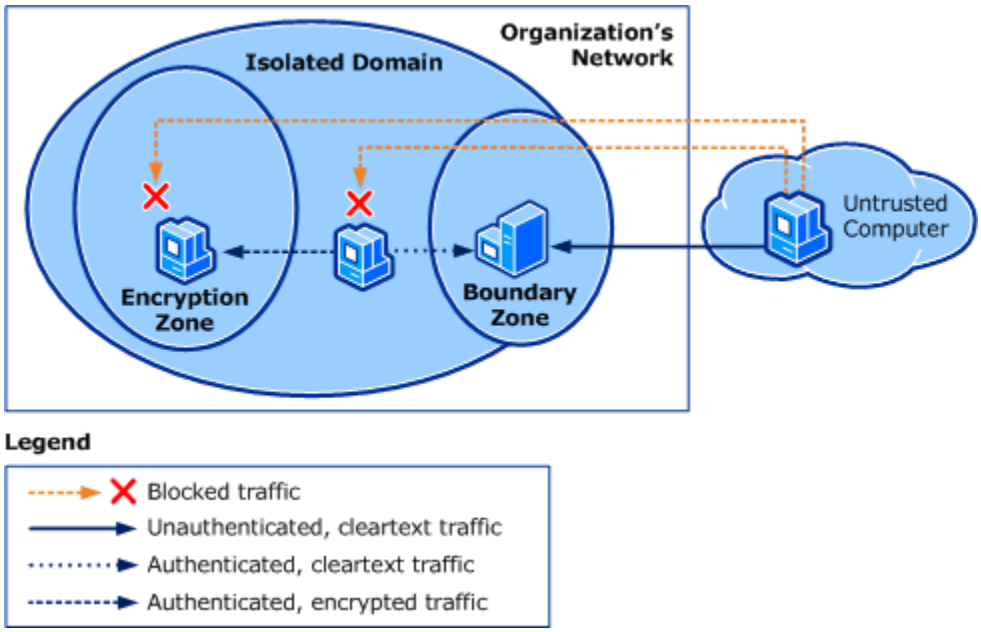

This goal provides the following benefits:

- Computers in the encryption zone require authentication to communicate with other computers. This works no differently from the domain isolation goal and design. For more information, see Restrict Access to Only Trusted Computers.
- Computers in the encryption zone require that all inbound and outbound network traffic be encrypted.

  For example, Woodgrove Bank processes sensitive customer data on a computer that must be protected from eavesdropping by computers on the network. Connection security rules specify that all traffic must be encrypted by a sufficiently complex encryption algorithm to help protect the data.
- Computers in the encryption zone are often good candidates for server isolation, where access is limited to only computer accounts and user accounts that are members of an authorized access group. In many organizations, the encryption zone and the server isolation zone are one and the same. For more information, see Restrict Access to Only Specified Users or Computers.

The following components are required for this deployment goal:

- **Active Directory**: Active Directory supports centralized management of connection security rules by configuring the rules in one or more GPOs that can be automatically applied to all relevant computers in the domain. For more information about Active Directory, see Additional Resources.

Next: Restrict Access to Only Specified Users or Computers

Restrict Access to Only Specified Users or Computers

Domain isolation (as described in the previous goal **Restrict Access to Only Trusted Computers**) prevents computers that are members of the isolated domain from accepting network traffic from untrusted computers. However, some computers on the network might host sensitive data that must be additionally restricted to only those users and computers that have a business requirement to access the data.

Windows Firewall with Advanced Security enables you to restrict access to computers and users that are members of domain groups authorized to access that computer. These groups are called *network access groups (NAGs)*. When a computer authenticates to a server, the server checks the group membership of the computer account and the user account, and grants access only if membership in the NAG is confirmed. Adding this check creates a virtual "secure zone" within the domain isolation zone. You can have multiple computers in a single secure zone, and it is likely that you will create a separate zone for each set of servers that have specific security access needs. Computers that are part of this server isolation zone are often also part of the encryption zone (see **Require Encryption When Accessing Sensitive Network Resources**).

Restricting access to only users and computers that have a business requirement can help you comply with regulatory and legislative requirements, such as those found in the Federal Information Security Management Act of 2002 (FISMA), the Sarbanes-Oxley Act of 2002, the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and other government and industry regulations.

Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista enable you to restrict access by specifying either computer or user credentials.

The following illustration shows an isolated server, and examples of computers that can and cannot communicate with it. Computers that are outside the Woodgrove corporate network, or computers that are in the isolated domain but are not members of the required NAG, cannot communicate with the isolated server.

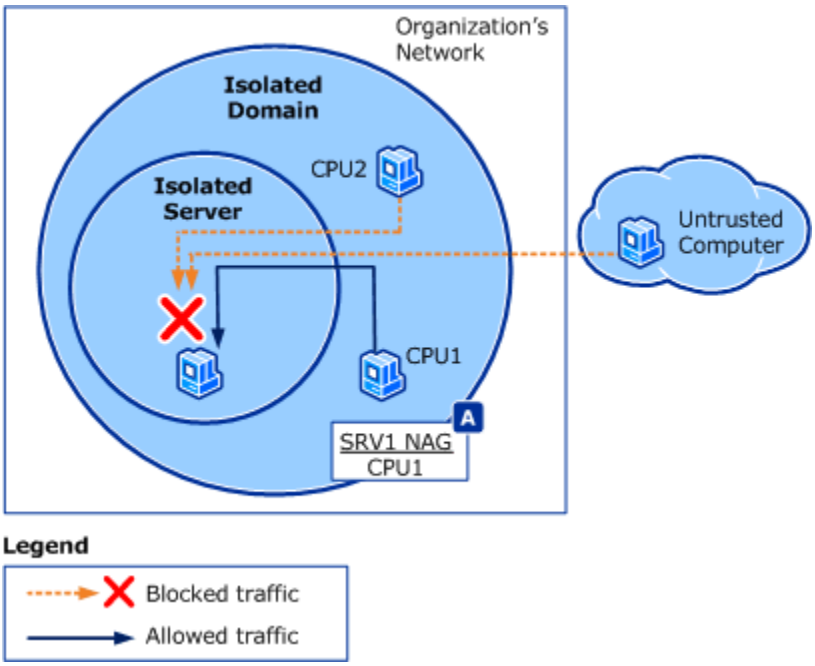

This goal, which corresponds to Server Isolation Policy Design, provides the following features:

- Isolated servers accept unsolicited inbound network traffic only from computers or users that are members of the NAG.
- Isolated servers can be implemented as part of an isolated domain, and treated as another zone. Members of the zone group receive a GPO with rules that require authentication, and that specify that only network traffic authenticated as coming from a member of the NAG is allowed.
- Server isolation can also be configured independently of an isolated domain. To do so, configure only the computers that must communicate with the isolated server with connection security rules to implement authentication and check NAG membership.
- A server isolation zone can be simultaneously configured as an encryption zone. To do this, configure the GPO with rules that force encryption in addition to requiring authentication and restricting access to NAG members. For more information, see Require Encryption When Accessing Sensitive Network Resources.

The following components are required for this deployment goal:

- **Active Directory**: Active Directory supports centralized management of connection security rules by configuring the rules in one or more GPOs that can be automatically applied to all relevant computers in the domain. For more information about Active Directory, see Additional Resources.

Next: Mapping Your Deployment Goals to a Windows Firewall with Advanced Security Design

Mapping Your Deployment Goals to a Windows Firewall with Advanced Security Design

After you finish reviewing the existing Windows Firewall with Advanced Security deployment goals and you determine which goals are important to your specific deployment, you can map those goals to a specific Windows Firewall with Advanced Security design.

Important
The first three designs presented in this guide build on each other to progress from simpler to more complex. Therefore during deployment, consider implementing them in the order presented. Each deployed design also provides a stable position from which to evaluate your progress, and to make sure that your goals are being met before you continue to the next design.

Use the following table to determine which Windows Firewall with Advanced Security design maps to the appropriate combination of Windows Firewall with Advanced Security deployment goals for your organization. This table refers only to the Windows Firewall with Advanced Security designs as described in this guide. However, you can create a hybrid or custom Windows Firewall with Advanced Security design by using any combination of the Windows Firewall with Advanced Security deployment goals to meet the needs of your organization.

Deployment Goals	Basic Firewall Policy Design	Domain Isolation Policy Design	Server Isolation Policy Design	Certificate-based Isolation Policy Design
Protect Computers from Unwanted Network Traffic	Yes	Yes	Yes	Yes
Restrict Access to Only Trusted Computers	-	Yes	Yes	Yes
Restrict Access to Only Specified Users or Computers	-	-	Yes	Yes
Require Encryption When Accessing Sensitive Network Resources	-	Optional	Optional	Optional

To examine details for a specific design, click the design title at the top of the column in the preceding table.

Next: Basic Firewall Policy Design

Basic Firewall Policy Design

Many organizations have a network perimeter firewall that is designed to prevent the entry of malicious traffic in to the organization's network, but do not have a host-based firewall enabled on each computer in the organization.

The Basic Firewall Policy Design helps you to protect the computers in your organization from unwanted network traffic that gets through the perimeter defenses, or that originates from inside your network. In this design, you deploy firewall rules to each computer in your organization to allow traffic that is required by the programs that are used. Traffic that does not match the rules is dropped.

Traffic can be blocked or permitted based on the characteristics of each network packet: its source or destination IP address, its source or destination port numbers, the program on the computer that receives the inbound packet, and so on. This design can also be deployed together with one or more of the other designs that add IPsec protection to the network traffic permitted.

Many network administrators do not want to tackle the difficult task of determining all the appropriate rules for every program that is used by the organization, and then maintaining that list over time. In fact, most programs do not require specific firewall rules. The default behavior of Windows and most contemporary applications makes this task easy:

- On client computers, the default firewall behavior already supports typical client programs. Programs designed for Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista create any required rules for you as part of the installation process. You only have to create a rule if the client program must be able to receive unsolicited inbound network traffic from another computer.
- When you install a server program that must accept unsolicited inbound network traffic, the installation program likely creates or enables the appropriate rules on the server for you.

  For example, when you install a server role in Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008, the appropriate firewall rules are created and enabled automatically.
- For other standard network behavior, the predefined rules that are built into Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista can easily be configured in a GPO and deployed to the computers in your organization.

  For example, by using the predefined groups for Core Networking and File and Printer Sharing you can easily configure GPOs with rules for those frequently used networking protocols.

With few exceptions, the firewall can be enabled on all configurations of Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista. Therefore, we recommended that you enable the firewall on every computer in your organization. This includes servers in your perimeter network, on mobile and remote clients that connect to the network, and on all servers and clients in your internal network.

Caution

Stopping the service associated with Windows Firewall with Advanced Security is not supported by Microsoft.

By default, in new installations, Windows Firewall is turned on in Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista.

If you turn off the Windows Firewall with Advanced Security service you lose other benefits provided by the service, such as the ability to use IPsec connection security rules, Windows Service Hardening, and network protection from forms of attacks that use network fingerprinting. For more information about Windows Service Hardening, see http://go.microsoft.com/fwlink/?linkid=104976.

Third-party firewall software that is compatible with Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista can programmatically disable only the parts of Windows Firewall with Advanced Security that might need to be disabled for compatibility. This is the recommended approach for third-party firewalls to coexist with the Windows Firewall; third-party party firewalls that comply with this recommendation have the certified logo from Microsoft.

An organization typically uses this design as a first step toward a more comprehensive Windows Firewall with Advanced Security design that adds server isolation and domain isolation.

After implementing this design, your administrative team will have centralized management of the firewall rules applied to all computers that are running Windows in your organization.

Important

If you also intend to deploy the Domain Isolation Policy Design, or the Server Isolation Policy Design, we recommend that you do the design work for all three designs together, and then deploy in layers that correspond with each design.

The basic firewall design can be applied to computers that are part of an Active Directory forest. Active Directory is required to provide the centralized management and deployment of Group Policy objects that contain the firewall settings and rules.

For more information about this design:

- This design coincides with the deployment goal to Protect Computers from Unwanted Network Traffic.
- To learn more about this design, see Firewall Policy Design Example.
- Before completing the design, gather the information described in Designing a Windows Firewall with Advanced Security Strategy.

- To help you make the decisions required in this design, see Planning Settings for a Basic Firewall Policy.
- For a list of detailed tasks that you can use to deploy your basic firewall policy design, see "Checklist: Implementing a Basic Firewall Policy Design" in the Windows Firewall with Advanced Security Deployment Guide at http://go.microsoft.com/fwlink/?linkid=98308.

Next: Domain Isolation Policy Design

Domain Isolation Policy Design

In the domain isolation policy design, you configure the computers on your network to accept only connections coming from computers that are authenticated as members of the same isolated domain.

This design typically begins with a network configured as described in the Basic Firewall Policy Design section. For this design, you then add connection security and IPsec rules to configure computers in the isolated domain to accept only network traffic from other computers that can authenticate as a member of the isolated domain. After implementing the new rules, your computers reject unsolicited network traffic from computers that are not members of the isolated domain.

The isolated domain might not be a single Active Directory domain. It can consist of all the domains in a forest, or domains in separate forests that have two-way trust relationships configured between them.

By using connection security rules based on IPsec, you provide a logical barrier between computers even if they are connected to the same physical network segment.

The design is shown in the following illustration, with the arrows that show the permitted communication paths.

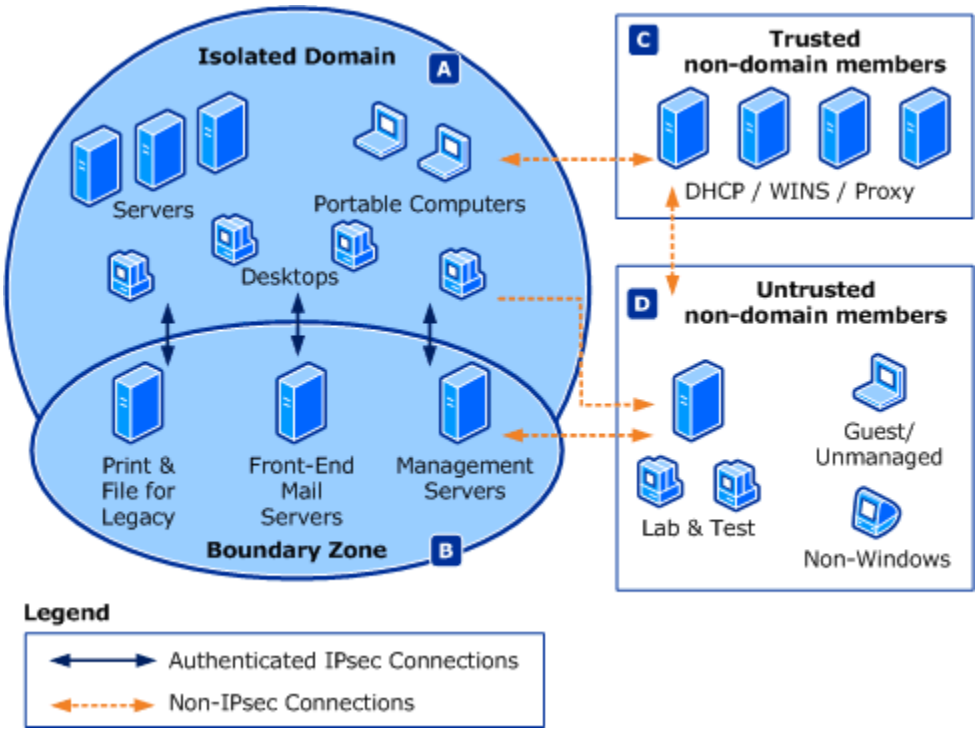

Characteristics of this design, as shown in the diagram, include the following:

- Isolated domain (area A) - Computers in the isolated domain receive unsolicited inbound traffic only from other members of the isolated domain or from computers referenced in authentication exemption rules. Computers in the isolated domain can send traffic to any computer. This includes unauthenticated traffic to computers that are not in the isolated domain. Computers that cannot join an Active Directory domain, but that can use certificates for authentication, can be part of the isolated domain. For more information, see the Certificate-based Isolation Policy Design.
- Boundary zone (area B) - Computers in the boundary zone are part of the isolated domain but are allowed to accept inbound connections from untrusted computers, such as clients on the Internet.

  Computers in the boundary zone request but do not require authentication to communicate. When a member of the isolated domain communicates with a boundary zone member the traffic is authenticated. When a computer that is not part of the isolated domain communicates with a boundary zone member the traffic is not authenticated.

  Because boundary zone computers are exposed to network traffic from untrusted and potentially hostile computers, they must be carefully managed and secured. Put only the computers that must be accessed by external computers in this zone. Use firewall rules to ensure that network traffic is accepted only for services that you want exposed to non-domain member computers.
- Trusted non-domain members (area C) - Computers on the network that are not domain members or that cannot use IPsec authentication are allowed to communicate by configuring

authentication exemption rules. These rules enable computers in the isolated domain to accept inbound connections from these trusted non-domain member computers.

- Untrusted non-domain members (area D) - Computers that are not managed by your organization and have an unknown security configuration must have access only to those computers required for your organization to correctly conduct its business. Domain isolation exists to put a logical barrier between these untrusted computers and your organization's computers.

After implementing this design, your administrative team will have centralized management of the firewall and connection security rules applied to the computers that are running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, and Windows Vista in your organization.

Important
This design builds on the Basic Firewall Policy Design, and in turn serves as the foundation for the Server Isolation Policy Design. If you plan to deploy all three, we recommend that you do the design work for all three together, and then deploy in the sequence presented.

This design can be applied to computers that are part of an Active Directory forest. Active Directory is required to provide the centralized management and deployment of Group Policy objects that contain the connection security rules.

In order to expand the isolated domain to include computers that cannot be part of an Active Directory domain, see the Certificate-based Isolation Policy Design.

For more information about this design:

- This design coincides with the deployment goals to Protect Computers from Unwanted Network Traffic, Restrict Access to Only Trusted Computers, and optionally Require Encryption When Accessing Sensitive Network Resources.
- To learn more about this design, see the Domain Isolation Policy Design Example.
- Before completing the design, gather the information described in Designing a Windows Firewall with Advanced Security Strategy.
- To help you make the decisions required in this design, see Planning Domain Isolation Zones and Planning Group Policy Deployment for Your Isolation Zones.
- For a list of tasks that you can use to deploy your domain isolation policy design, see "Checklist: Implementing a Domain Isolation Policy Design" in the Windows Firewall with Advanced Security Deployment Guide at http://go.microsoft.com/fwlink/?linkid=xxxxx.

Next: Server Isolation Policy Design

Server Isolation Policy Design

In the server isolation policy design, you assign servers to a zone that allows access only to users and computers that authenticate as members of an approved network access group (NAG).

This design typically begins with a network configured as described in the Domain Isolation Policy Design section. For this design, you then create zones for servers that have additional security requirements. The zones can limit access to the server to only members of authorized groups, and can optionally require the encryption of all traffic in or out of these servers. This can be done on a per server basis, or for a group of servers that share common security requirements.

You can implement a server isolation design without using domain isolation. To do this, you use the same principles as domain isolation, but instead of applying them to an Active Directory domain, you apply them only to the computers that must be able to access the isolated servers. The GPO contains connection security and firewall rules that require authentication when communicating with the isolated servers. In this case, the NAGs that determine which users and computers can access the isolated server are also used to determine which computers receive the GPO.

The design is shown in the following illustration, with arrows that show the permitted communication paths.

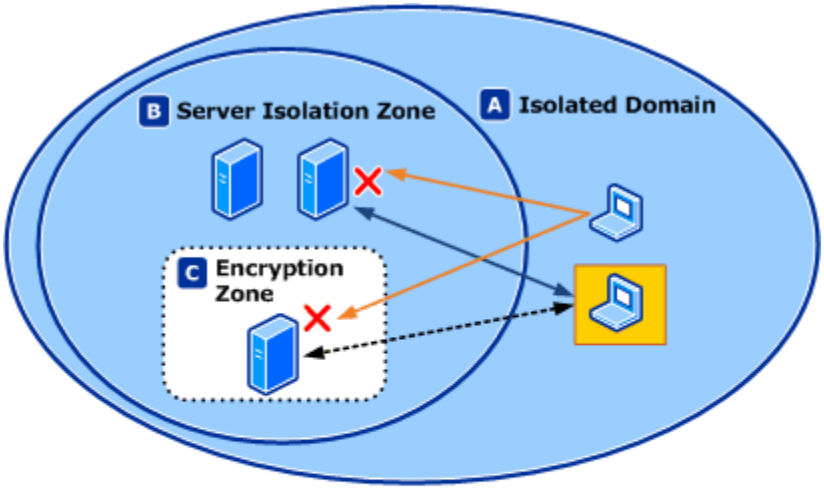

Characteristics of this design include the following:

- Isolated domain (area A) - The same isolated domain described in the Domain Isolation Policy Design section. If the isolated domain includes a boundary zone, then computers in the boundary zone behave just like other members of the isolated domain in the way that they interact with computers in server isolation zones.
- Isolated servers (area B) - Computers in the server isolation zones restrict access to computers, and optionally users, that authenticate as a member of a network access group (NAG) authorized to gain access.

- Encryption zone (area C) - If the data being exchanged is sufficiently sensitive, the connection security rules for the zone can also require that the network traffic be encrypted. Encryption zones are most often implemented as rules that are part of a server isolation zone, instead of as a separate zone. The diagram illustrates the concept as a subset for conceptual purposes only.

To add support for server isolation, you must ensure that the authentication methods are compatible with the requirements of the isolated server. For example, if you want to authorize user accounts that are members of a NAG in addition to authorizing computer accounts, you must enable both user and computer authentication in your connection security rules.

Important
This design builds on the Domain Isolation Policy Design, which in turn builds on the Basic Firewall Policy Design. If you plan to deploy all three designs, do the design work for all three together, and then deploy in the sequence presented.

This design can be applied to computers that are part of an Active Directory forest. Active Directory is required to provide the centralized management and deployment of Group Policy objects that contain the connection security rules.

For more information about this design:

- This design coincides with the deployment goals to Protect Computers from Unwanted Network Traffic, Restrict Access to Only Trusted Computers, Restrict Access to Only Specified Users or Computers, and Require Encryption When Accessing Sensitive Network Resources.
- To learn more about this design, see Server Isolation Policy Design Example.
- Before completing the design, gather the information described in Designing a Windows Firewall with Advanced Security Strategy.
- To help you make the decisions required in this design, see Planning Server Isolation Zones and Planning Group Policy Deployment for Your Isolation Zones.
- For a list of tasks that you can use to deploy your server isolation policy design, see "Checklist: Implementing a Standalone Server Isolation Policy Design" in the Windows Firewall with Advanced Security Deployment Guide at http://go.microsoft.com/fwlink/?linkid=xxxx.

Next: Certificate-based Isolation Policy Design

Certificate-based Isolation Policy Design

In the certificate-based isolation policy design, you provide the same types of protections to your network traffic as described in the Domain Isolation Policy Design and Server Isolation Policy Design sections. The only difference is the method used to share identification credentials during the authentication of your network traffic.

Domain isolation and server isolation help provide security for the computers on the network that run Windows and that can be joined to an Active Directory domain. However, in most corporate

environments there are typically some computers that must run another operating system, such as Linux or UNIX. These computers cannot join an Active Directory domain, without a third-party package being installed. Also, some computers that do run Windows cannot join a domain for a variety of reasons. To rely on Kerberos V5 as the authentication protocol, the computer needs to be joined to the Active Directory and (for non-windows computers) support Kerberos as an authentication protocol.

To authenticate with non-domain member computers, IPsec supports using standards-based cryptographic certificates. Because this authentication method is also supported by many third-party operating systems, it can be used as a way to extend your isolated domain to computers that do not run the Windows operating system.

The same principles of the domain and server isolation designs apply to this design. Only computers that can authenticate (in this case, by providing a specified certificate) can communicate with the computers in your isolated domain.

For computers that run Windows and that are part of an Active Directory domain, you can use Group Policy to deploy the certificates required to communicate with the computers that are trusted but are not part of the Active Directory domain. For other computers, you will have to either manually configure them with the required certificates, or use a third-party program to distribute the certificates in a secure manner.

For more information about this design:

- This design coincides with the deployment goals to Protect Computers from Unwanted Network Traffic, Restrict Access to Only Trusted Computers, and optionally Require Encryption When Accessing Sensitive Network Resources.
- To learn more about this design, see Certificate-based Isolation Policy Design Example.
- Before completing the design, gather the information described in Designing a Windows Firewall with Advanced Security Strategy.
- To help you make the decisions required in this design, see Planning Certificate-based Authentication.
- For a list of tasks that you can use to deploy your certificate-based policy design, see "Checklist: Implementing a Certificate-based Isolation Policy Design" in the Windows Firewall with Advanced Security Deployment Guide at http://go.microsoft.com/fwlink/?linkid=98308.

Next: Evaluating Windows Firewall with Advanced Security Design Examples

Evaluating Windows Firewall with Advanced Security Design Examples

The following Windows Firewall with Advanced Security design examples illustrate how you can use Windows Firewall with Advanced Security to improve the security of the computers connected to the network. You can use these topics to evaluate how the firewall and connection security rules work across all Windows Firewall with Advanced Security designs and to determine which design or combination of designs best suits the goals of your organization.

- **Firewall Policy Design Example**
- **Domain Isolation Policy Design Example**
- **Server Isolation Policy Design Example**
- **Certificate-based Isolation Policy Design Example**

Firewall Policy Design Example

In this example, the fictitious company Woodgrove Bank is a financial services institution.

Woodgrove Bank has an Active Directory domain that provides Group Policy-based management for all their Windows-based computers. The Active Directory domain controllers also host Domain Name System (DNS) for host name resolution. Separate computers host Windows Internet Name Service (WINS) for network basic input/output system (NetBIOS) name resolution. A set of computers that are running UNIX provide the Dynamic Host Configuration Protocol (DHCP) services for automatic IP addressing.

Woodgrove Bank is in the process of migrating their computers from Windows Vista and Windows Server 2008 to Windows 8 and Windows Server 2012. A significant number of the computers at Woodgrove Bank continue to run Windows Vista and Windows Server 2008. Interoperability between the previous and newer operating systems must be maintained. Wherever possible, security features applied to the newer operating systems must also be applied to the previous operating systems.

A key line-of-business program called WGBank consists of a client program running on most of the desktop computers in the organization. This program accesses several front-end server computers that run the server-side part of WGBank. These front-end servers only do the processing — they do not store the data. The data is stored in several back-end database computers that are running Microsoft SQL Server.

Design requirements

The network administrators want to implement Windows Firewall with Advanced Security throughout their organization to provide an additional security layer to their overall security strategy. They want to create firewall rules that allow their business programs to operate, while blocking network traffic that is not wanted.

The following illustration shows the traffic protection needs for this design example.

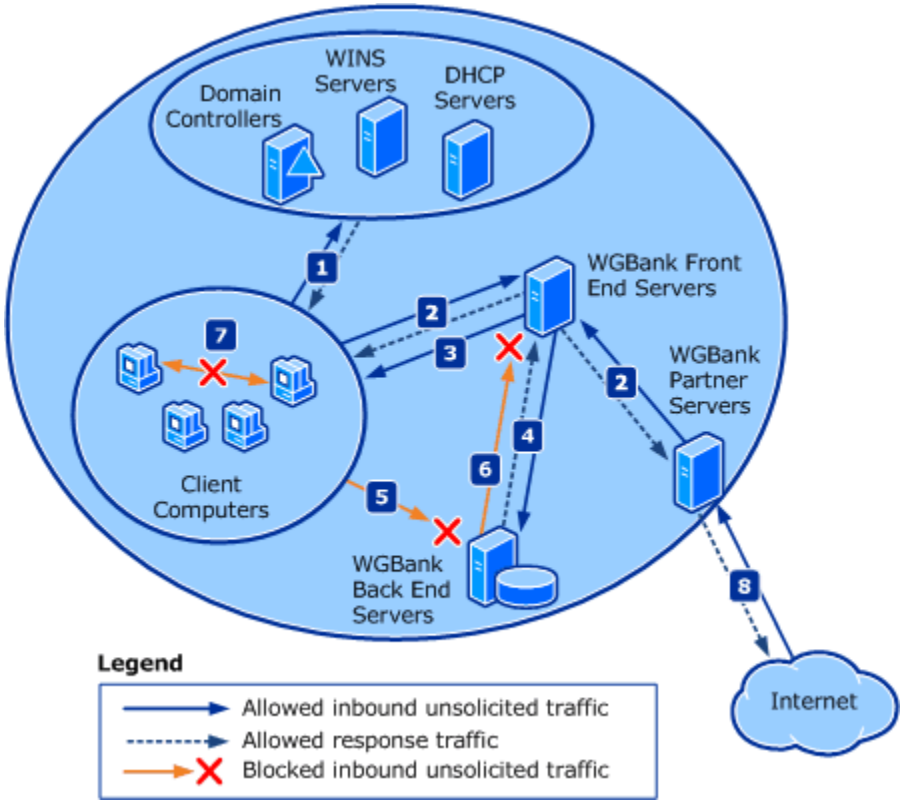

1. The network infrastructure servers that are running services, such as Active Directory, DNS, DHCP, or WINS, can receive unsolicited inbound requests from network clients. The network clients can receive the responses from the infrastructure servers.
2. The WGBank front-end servers can receive unsolicited inbound traffic from the client computers and the WGBank partner servers. The WGBank client computers and partner servers can receive the response.
3. The WGBank front-end servers can send updated information to the client computers to support real-time display. The clients do not poll for this unsolicited traffic, but must be able to receive it.
4. The WGBank back-end servers can receive SQL query requests from the WGBank front-end servers. The WGBank front-end servers can receive the corresponding responses.
5. There is no direct communications between the client computers and the WGBank back-end computers.
6. There is no unsolicited traffic from the WGBank back-end computers to the WGBank front-end servers.
7. Company policy prohibits the use of peer-to-peer file transfer software. A recent review by the IT staff found that although the perimeter firewall does prevent most of the programs in this category from working, two programs are being used by staff members that do not require an outside server. Firewall rules must block the network traffic created by these programs.
8. The WGBank partner servers can receive inbound requests from partner computers through the Internet.

Other traffic notes:

- Computers are not to receive any unsolicited traffic from any computer other than specifically allowed above.
- Other outbound network traffic from the client computers not specifically identified in this example is permitted.

Design details

Woodgrove Bank uses Active Directory groups and Group Policy Objects to deploy the firewall settings and rules to the computers on their network. They know that they must deploy policies to the following collections of computers:

- Client computers that run Windows 8, Windows 7, or Windows Vista
- WGBank front-end servers that run Windows Server 2012 or Windows Server 2008 R2 (there are none in place yet, but their solution must support adding them)
- WGBank partner servers that run Windows Server 2008
- WGBank back-end SQL Server computers that run Windows Server 2008 (there are none in place yet, but their solution must support adding them)
- Infrastructure servers that run Windows Server 2008
- Active Directory domain controllers that run Windows Server 2008 R2 or Windows Server 2012
- DHCP servers that run the UNIX operating system

After evaluating these sets of computers, and comparing them to the Active Directory organizational unit (OU) structure, Woodgrove Bank network administrators determined that there was not a good one-to-one match between the OUs and the sets. Therefore the firewall GPOs will not be linked directly to OUs that hold the relevant computers. Instead, the GPOs are linked to the domain container in Active Directory, and then WMI and group filters are attached to the GPO to ensure that it is applied to the correct computers.

Setting up groups as described here ensures that you do not have to know what operating system a computer is running before assigning it to a group. A combination of WMI filters and security group filters are used to ensure that members of the group receive the GPO appropriate for the version of Windows running on that computer. For some groups, you might have four or even five GPOs.

The following groups were created by using the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in, and all computers that run Windows were added to the correct groups:

- **CG_FIREWALL_ALLCOMPUTERS**. Add the predefined and system managed **Domain computers** group as a member of this group. All members of the FIREWALL_ALLCOMPUTERS group receive an operating system-specific GPO with the common firewall rules applied to all computers.

  The two computer types (client and server) are distinguished by using a WMI filters to ensure that only the policy intended for computers that are running a client version of Windows can be applied to that computer. A similar WMI filter on the server GPO ensures that only

computers that are running server versions of Windows can apply that GPO. Each of the GPOs also have security group filters to prevent members of the group FIREWALL_NO_DEFAULT from receiving either of these two GPOs.

- Client computers receive a GPO that configures Windows Firewall with Advanced Security to enforce the default Windows Firewall behavior (allow outbound, block unsolicited inbound). The client default GPO also includes the built-in firewall rule groups Core Networking and File and Printer Sharing. The Core Networking group is enabled for all profiles, whereas the File and Printer Sharing group is enabled for only the Domain and Private profiles. The GPO also includes inbound firewall rules to allow the WGBank front-end server dashboard update traffic, and rules to prevent company-prohibited programs from sending or receiving network traffic, both inbound and outbound.
- Server computers receive a GPO that includes similar firewall configuration to the client computer GPO. The primary difference is that the rules are enabled for all profiles (not just domain and private). Also, the rules for WGBank dashboard update are not included, because it is not needed on server computers.

All rules are scoped to allow network traffic only from computers on Woodgrove Bank's corporate network.

- **CG_FIREWALL_NO_DEFAULT**. Members of this group do not receive the default firewall GPO. Computers are added to this group if there is a business requirement for it to be exempted from the default firewall behavior. The use of a group to represent the exceptions instead of the group members directly makes it easier to support the dynamic nature of the client computer population. A new computer joined to the domain is automatically given the appropriate default firewall GPO, unless it is a member of this group.
- **CG_FIREWALL_WGB_FE**. This group contains the computer accounts for all the WGBank front-end server computers. Members of this group receive a GPO that configures Windows Firewall with Advanced Security with inbound firewall rules to allow unsolicited WGBank client traffic. Computers in this group also receive the default firewall GPO.
- **CG_FIREWALL_WGB_SQL**. This group contains the computer accounts for all the WGBank back-end computers that run SQL Server. Members of this group receive a GPO that configures Windows Firewall with Advanced Security with inbound firewall rules to allow the SQL Server program to receive unsolicited queries only from the WGBank front-end servers. Computers in this group also receive the default firewall GPO.
- **CG_FIREWALL_BOUNDARY_WGBANKFE**. This group contains the computer accounts for the servers that host Web services that can be accessed from the Internet. Members of this group receive a GPO that adds an inbound firewall rule to allow inbound HTTP and HTTPS network traffic from any address, including the Internet. Computers in this group also receive the default firewall GPO.
- **CG_FIREWALL_WINS**. This group contains the computer accounts for all the WINS server computers. Members of this group receive a GPO that configures Windows Firewall with Advanced Security with an inbound firewall rule to allow unsolicited inbound requests from WINS clients. Computers in this group also receive the default firewall GPO.
- **CG_FIREWALL_ADDC**. This group contains all the computer accounts for the Active Directory domain controller server computers. Members of this group receive a GPO that configures Windows Firewall with Advanced Security with inbound firewall rules to allow

unsolicited Active Directory client and server-to-server traffic. Computers in this group also receive the default firewall GPO.

In your own design, create a group for each computer role in your organization that requires different or additional firewall rules. For example, file servers and print servers require additional rules to allow the incoming network traffic for those functions. If a function is ordinarily performed on most computers on the network, you might consider adding computers performing those roles to the common default firewall GPO set, unless there is a security reason not to include it there.

Next: Domain Isolation Policy Design Example

Domain Isolation Policy Design Example

This design example continues to use the fictitious company Woodgrove Bank, and builds on the example described in the Firewall Policy Design Example section. See that example for an explanation of the basic corporate network infrastructure at Woodgrove Bank with diagrams.

Design Requirements

In addition to the basic protection provided by the firewall rules in the previous design example, the administrators of the network want to implement domain isolation to provide another layer of security to their networked computers. They want to create firewall and connection security rules that use authentication to reduce the risk of communicating with untrusted and potentially hostile computers.

The following illustration shows the traffic protection needed for this design example.

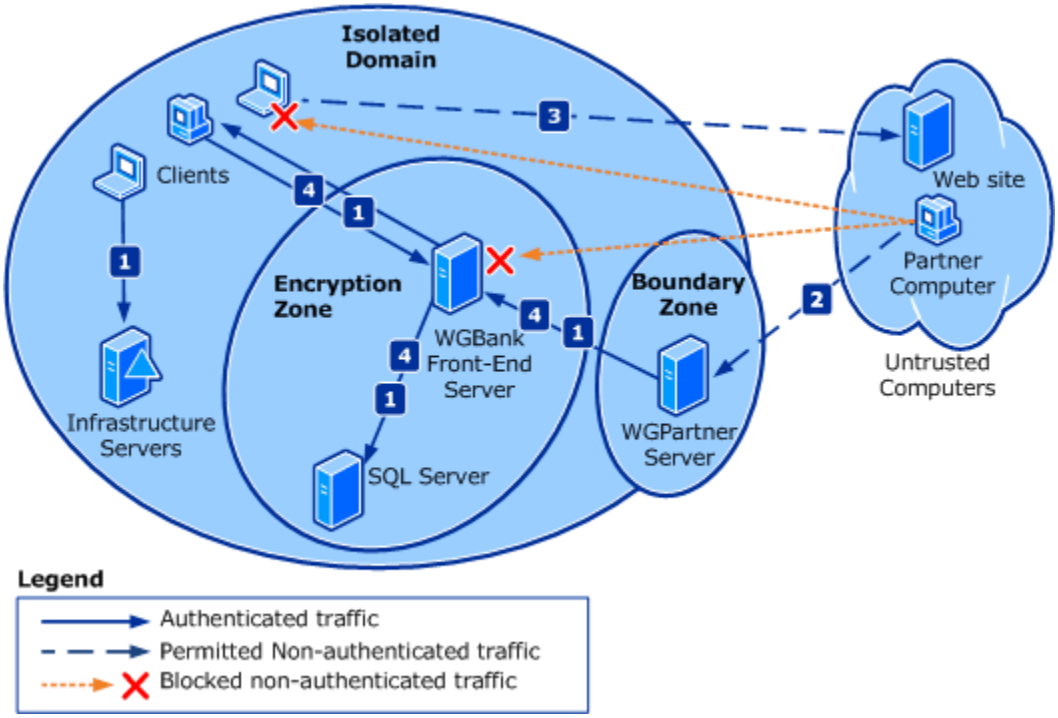

1. All computers on the Woodgrove Bank corporate network that are Active Directory domain members must authenticate inbound network traffic as coming from another computer that is a member of the domain. Unless otherwise specified in this section, Woodgrove Bank's computers reject all unsolicited inbound network traffic that is not authenticated. If the basic firewall design is also implemented, even authenticated inbound network traffic is dropped unless it matches an inbound firewall rule.
2. The servers hosting the WGPartner programs must be able to receive unsolicited inbound traffic from computers owned by its partners, which are not members of Woodgrove Bank's domain.
3. Client computers can initiate non-authenticated outbound communications with computers that are not members of the domain, such as browsing external Web sites. Unsolicited inbound traffic from non-domain members is blocked.
4. Computers in the encryption zone require that all network traffic inbound and outbound must be encrypted, in addition to the authentication already required by the isolated domain.

Other traffic notes:

- All of the design requirements described in the Firewall Policy Design Example section are still enforced.

Design Details

Woodgrove Bank uses Active Directory groups and GPOs to deploy the domain isolation settings and rules to the computers on its network.

Setting up groups as described here ensures that you do not have to know what operating system a computer is running before assigning it to a group. As in the firewall policy design, a combination of WMI filters and security group filters are used to ensure that members of the group receive the GPO appropriate for the version of Windows running on that computer. For some groups, you might have four or even five GPOs.

The following groups were created by using the Active Directory Users and Computers MMC snap-in, all computers that run Windows were added to the correct groups, and then the appropriate GPO are applied to the group. To include a computer in the isolated domain or any one of its subordinate zones, simply add the computer's account in the appropriate group.

- **CG_DOMISO_ISOLATEDDOMAIN**. The members of this group participate in the isolated domain. After an initial pilot period, followed by a slowly increasing group membership, the membership of this group was eventually replaced with the entry **Domain Computers** to ensure that all computers in the domain participate by default. The WMI filters ensure that the GPO does not apply to domain controllers. GPOs with connection security rules to enforce domain isolation behavior are linked to the domain container and applied to the computers in this group. Filters ensure that each computer receives the correct GPO for its operating system type. The rules in the domain isolation GPO require Kerberos v5 authentication for inbound network connections, and request (but not require) it for all outbound connections.
- **CG_DOMISO_NO_IPSEC**. This group is denied read or apply permissions on any of the domain isolation GPOs. Any computer that cannot participate in domain isolation, such as a DHCP server running UNIX, is added to this group.
- **CG_DOMISO_BOUNDARY**. This group contains the computer accounts for all the computers that are part of the boundary group able to receive unsolicited inbound traffic from untrusted computers. Members of the group receive a GPO that configures connection security rules to request (but not require) both inbound and outbound authentication.
- **CG_DOMISO_ENCRYPTION**. This group contains the computer accounts for all the computers that require all inbound and outbound traffic to be both authenticated and encrypted. Members of the group receive a GPO that configures connection security and firewall rules to require both authentication and encryption on all inbound and outbound traffic.

Note

If you are designing GPOs for only Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2, you can design your GPOs in nested groups. For example, you can make the boundary group a member of the isolated domain group, so that it receives the firewall and basic isolated domain settings through that nested membership, with only the changes supplied by the boundary zone GPO. However, computers that are running older versions of Windows can only support a single IPsec policy being active at a time. The policies for each GPO must be complete (and to a great extent redundant with each other), because you cannot layer them as you can in the newer versions of Windows. For simplicity, this guide describes the techniques used to create the independent, non-layered policies. We recommend that you create and periodically run a script that compares the memberships

of the groups that must be mutually exclusive and reports any computers that are incorrectly assigned to more than one group.

Next: Server Isolation Policy Design Example

Server Isolation Policy Design Example

This design example continues to use the fictitious company Woodgrove Bank, as described in the Firewall Policy Design Example section and the Domain Isolation Policy Design Example section.

In addition to the protections provided by the firewall and domain isolation, Woodgrove Bank wants to provide additional protection to the computers that are running Microsoft SQL Server for the WGBank program. They contain personal data, including each customer's financial history. Government and industry rules and regulations specify that access to this information must be restricted to only those users who have a legitimate business need. This includes a requirement to prevent interception of and access to the information when it is in transit over the network.

The information presented by the WGBank front-end servers to the client computers, and the information presented by the WGPartner servers to the remote partner computers, are not considered sensitive for the purposes of the government regulations, because they are processed to remove sensitive elements before transmitting the data to the client computers.

In this guide, the examples show server isolation layered on top of a domain isolation design. If you have an isolated domain, the client computers are already equipped with GPOs that require authentication. You only have to add settings to the isolated server(s) to require authentication on inbound connections, and to check for membership in the NAG. The connection attempt succeeds only if NAG membership is confirmed.

Server isolation without domain isolation

Server isolation can also be deployed by itself, to only the computers that must participate. The GPO on the server is no different from the one discussed in the previous paragraph for a server in an existing isolated domain. The difference is that you must also deploy a GPO with supporting connection security rules to the clients that must be able to communicate with the isolated server. Because those computers must be members of the NAG, that group can also be used in a security group filter on the client GPO. That GPO must contain rules that support the authentication requirements of the isolated server.

In short, instead of applying the client GPO to all clients in the domain, you apply the GPO to only the members of the NAG.

If you do not have an Active Directory domain then you can manually apply the connection security rules to the client computers, or you can use a netsh command-line script (or Windows PowerShell in Windows 8 and Windows Server 2012) to help automate the configuration of the rules on larger numbers of computers. If you do not have an Active Directory domain, you cannot

use the Kerberos V5 protocol, but instead must provide the clients and the isolated servers with certificates that are referenced in the connection security rules.

Design requirements

In addition to the protection provided by the firewall rules and domain isolation described in the previous design examples, the network administrators want to implement server isolation to help protect the sensitive data stored on the computers that run SQL Server.

The following illustration shows the traffic protection needs for this design example.

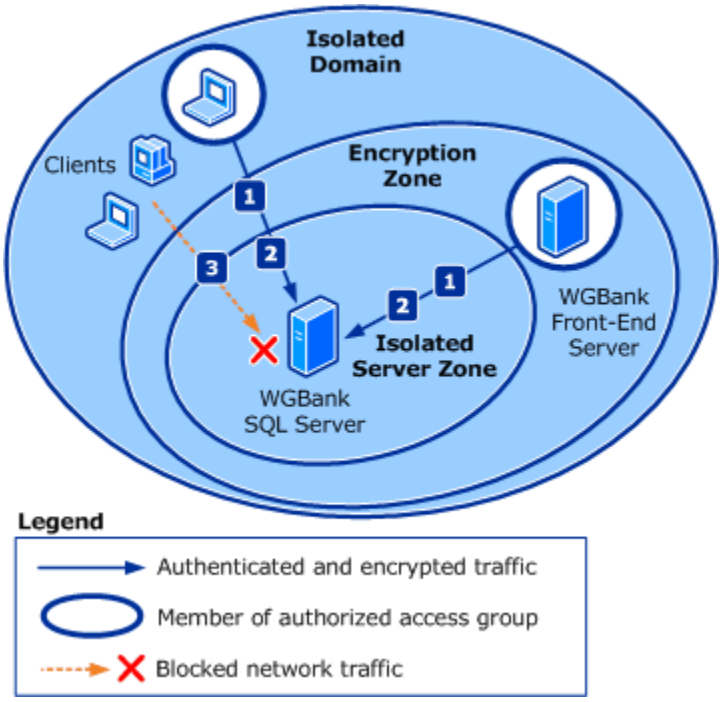

1. Access to the SQL Server computers must be restricted to only those computer or user accounts that have a business requirement to access the data. This includes the service accounts that are used by the WGBank front-end servers, and administrators of the SQL Server computers. In addition, access is only granted when it is sent from an authorized computer. Authorization is determined by membership in a network access group (NAG).
2. All network traffic to and from the SQL Server computers must be encrypted.
3. Client computers or users whose accounts are not members of the NAG cannot access the isolated servers.

Other traffic notes:

- All of the design requirements shown in the Firewall Policy Design Example section are still enforced.
- All of the design requirements shown in the Domain Isolation Policy Design Example section are still enforced.

Design details

Woodgrove Bank uses Active Directory groups and GPOs to deploy the server isolation settings and rules to the computers on its network.

As in the previously described policy design examples, GPOs to implement the domain isolation environment are linked to the domain container in Active Directory, and then WMI filters and security group filters are attached to GPOs to ensure that the correct GPO is applied to each computer. The following groups were created by using the Active Directory Users and Computers snap-in, and all computers that run Windows were added to the correct groups.

- **CG_SRVISO_WGBANK_SQL**. This group contains the computer accounts for the computers that run SQL Server. Members of this group receive a GPO with firewall and connections security rules that require that only users who are members of the group CG_NAG_SQL_USERS can access the server, and only when they are using a computer that is a member of the group CG_NAG_SQL_COMPUTERS.

Note

If you are designing GPOs for only Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2, you can design your GPOs in nested groups. For example, you can make the boundary group a member of the isolated domain group, so that it receives the firewall and basic isolated domain settings through that nested membership, with only the changes supplied by the boundary zone GPO. However, computers that are running older versions of Windows can only support a single IPsec policy being active at a time. The policies for each GPO must be complete (and to a great extent redundant with each other), because you cannot layer them as you can in the newer versions of Windows. For simplicity, this guide describes the techniques used to create the independent, non-layered policies. We recommend that you create and periodically run a script that compares the memberships of the groups that must be mutually exclusive and reports any computers that are incorrectly assigned to more than one group.

Network access groups (NAGs) are not used to determine which GPOs are applied to a computer. Instead, these groups determine which users and computers can access the services on the isolated server.

- **CG_NAG_SQL_COMPUTERS**. This network access group contains the computer accounts that are able to access the computers running SQL Server hosting the WGBank data. Members of this group include the WGBank front-end servers, and some client computers from which SQL Server administrators are permitted to work on the servers.
- **CG_NAG_SQL_USERS**. This network access group contains the user accounts of users who are permitted to access the SQL Server computers that host the WGBank data. Members of this group include the service account that the WGBank front-end program uses to run on its computers, and the user accounts for the SQL Server administration team members.

Note

You can use a single group for both user and computer accounts. Woodgrove Bank chose to keep them separate for clarity.

If Woodgrove Bank wants to implement server isolation without domain isolation, the CG_NAG_SQL_COMPUTERS group can also be attached as a security group filter on the GPOs that apply connection security rules to the client computers. By doing this, all the computers that are authorized to access the isolated server also have the required connection security rules.

You do not have to include the encryption-capable rules on all computers. Instead, you can create GPOs that are applied only to members of the NAG, in addition to the standard domain isolation GPO, that contain connection security rules to support encryption.

Next: Certificate-based Isolation Policy Design Example

Certificate-based Isolation Policy Design Example

This design example continues to use the fictitious company Woodgrove Bank, as described in the sections Firewall Policy Design Example, Domain Isolation Policy Design Example, and Server Isolation Policy Design Example.

One of the servers that must be included in the domain isolation environment is a computer running UNIX that supplies other information to the WGBank dashboard program running on the client computers. This computer sends updated information to the WGBank front-end servers as it becomes available, so it is considered unsolicited inbound traffic to the computers that receive this information.

Design requirements

One possible solution to this is to include an authentication exemption rule in the GPO applied to the WGBank front-end servers. This rule would instruct the front-end servers to accept traffic from the non-Windows computer even though it cannot authenticate.

A more secure solution, and the one selected by Woodgrove Bank, is to include the non-Windows computer in the domain isolation design. Because it cannot join an Active Directory domain, Woodgrove Bank chose to use certificate-based authentication. Certificates are cryptographically-protected documents, encrypted in such a way that their origin can be positively confirmed.

In this case, Woodgrove Bank used Microsoft Certificate Services, included with Windows Server 2008, to create the appropriate certificate. They might also have acquired and installed a certificate from a third-party commercial certification authority. They then used Group Policy to deploy the certificate to the front-end servers. The GPOs applied to the front-end servers also include updated connection security rules that permit certificate-based authentication in addition to Kerberos V5 authentication. They then manually installed the certificate on the UNIX server.

The UNIX server is configured with firewall and IPsec connection security rules using the tools that are provided by the operating system vendor. Those rules specify that authentication is performed by using the certificate.

The creation of the IPsec connection security rules for a non-Windows computer is beyond the scope of this document, but support for a certificate that can be used to authenticate such a non-Windows computer by using the standard IPsec protocols is the subject of this design.

The non-Windows computer can be effectively made a member of the boundary zone or the encryption zone based on the IPsec rules applied to the computer. The only constraint is that the main mode and quick mode encryption algorithms supported by the UNIX computer must also be supported by the Windows-based computers with which it communicates.

Other traffic notes:

- None of the capabilities of the other designs discussed in this guide are compromised by the use of certificate authentication by a non-Windows computer.

Design details

Woodgrove Bank uses Active Directory groups and GPOs to deploy the domain isolation settings and rules to the computers in their organization.

The inclusion of one or more non-Windows computers to the network requires only a simple addition to the GPOs for computers that must communicate with the non-Windows computer. The addition is allowing certificate-based authentication in addition to the Active Directory–supported Kerberos V5 authentication. This does not require including new rules, just adding certificate-based authentication as an option to the existing rules.

When multiple authentication methods are available, two negotiating computers agree on the first one in their lists that match. Because the majority of the computers in Woodgrove Bank's network run Windows, Kerberos V5 is listed as the first authentication method in the rules. Certificate-based authentication is added as an alternate authentication type.

By using the Active Directory Users and Computers snap-in, Woodgrove Bank created a group named NAG_COMPUTER_WGBUNIX. They then added the computer accounts to this group for Windows computers that need to communicate with the non-Windows computers. If all the computers in the isolated domain need to be able to access the non-Windows computers, then the **Domain Computers** group can be added to the group as a member.

Woodgrove Bank then created a GPO that contains the certificate, and then attached security group filters to the GPO that allow read and apply permissions to only members of the NAG_COMPUTER_WGBUNIX group. The GPO places the certificate in the **Local Computer / Personal / Certificates** certificate store. The certificate used must chain back to a certificate that is in the **Trusted Root Certification Authorities** store on the local computer.

Next: Designing a Windows Firewall with Advanced Security Strategy

Designing a Windows Firewall with Advanced Security Strategy

To select the most effective design for helping to protect the network, you must spend time collecting key information about your current computer environment. You must have a good understanding of what tasks the computers on the network perform, and how they use the network to accomplish those tasks. You must understand the network traffic generated by the programs running on the computers.

- Gathering the Information You Need
- Determining the Trusted State of Your Computers

The information that you gather will help you answer the following questions. The answers will help you understand your security requirements and select the design that best matches those requirements. The information will also help you when it comes time to deploy your design, by helping you to build a deployment strategy that is cost effective and resource efficient. It will help you project and justify the expected costs associated with implementing the design.

- What traffic must always be allowed? What are characteristics of the network traffic generated and consumed by the business programs?
- What traffic must always be blocked? Does your organization have policies that prohibit the use of specific programs? If so, what are the characteristics of the network traffic generated and consumed by the prohibited programs?
- What traffic on the network cannot be protected by IPsec because the computers or devices sending or receiving the traffic do not support IPsec?
- For each type of network traffic, does the default configuration of the firewall (block all unsolicited inbound network traffic, allow all outbound traffic) allow or block the traffic as required?
- Do you have an Active Directory domain (or forest of trusted domains) to which all your computers are joined? If you do not, then you cannot use Group Policy for easy mass deployment of your firewall and connection security rules. You also cannot easily take advantage of Kerberos V5 authentication that all domain clients can use.
- Which computers must be able to accept unsolicited inbound connections from computers that are not part of the domain?
- Which computers contain data that must be encrypted when exchanged with another computer?
- Which computers contain sensitive data to which access must be restricted to specifically authorized users and computers?
- Does your organization have specific network troubleshooting devices or computers (such as protocol analyzers) that must be granted unlimited access to the computers on the network, essentially bypassing the firewall?

If you already have firewall or IPsec rules deployed

Windows Firewall with Advanced Security in Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 has many new capabilities that are not available in earlier versions of Windows.

If you already have a domain and/or server isolation deployment in your organization then you can continue to use your existing GPOs and apply them to computers running Windows 8 and Windows Server 2012.

Note
Computers running Windows XP and Windows Server 2003 will not be able to participate in this domain and/or server isolation deployment plan.

This guide describes how to plan your groups and GPOs for an environment with a mix of operating systems, starting with Windows Vista and Windows Server 2008. Windows XP and Windows Server 2003 are not discussed in this guide. Details can be found in the section Planning Group Policy Deployment for Your Isolation Zones later in this guide.

Next: Gathering the Information You Need

Gathering the Information You Need

Before starting the planning process for a Windows Firewall with Advanced Security deployment, you must collect and analyze up-to-date information about the network, the directory services, and the computers that are already deployed in the organization. This information enables you to create a design that accounts for all possible elements of the existing infrastructure. If the gathered information is not accurate, problems can occur when devices and computers that were not considered during the planning phase are encountered during implementation.

Review each of the following topics for guidance about the kinds of information that you must gather:

- Gathering Information about Your Current Network Infrastructure
- Gathering Information about Your Active Directory Deployment
- Gathering Information about Your Computers
- Gathering Other Relevant Information

Gathering Information about Your Current Network Infrastructure

Perhaps the most important aspect of planning for Windows Firewall with Advanced Security deployment is the network architecture, because IPsec is layered on the Internet Protocol itself.

An incomplete or inaccurate understanding of the network can prevent any Windows Firewall with Advanced Security solution from being successful. Understanding subnet layout, IP addressing schemes, and traffic patterns are part of this effort, but accurately documenting the following components are important to completing the planning phase of this project:

- **Network segmentation**. This includes IP addressing maps, showing how your routers separate each network segment. It includes information about how the routers are configured, and what security filters they impose on network traffic flowing through them.
- Network address translation (NAT). NAT is a means of separating network segments by using a device that maps all of the IP addresses on one side of the device to a single IP address accessible on the other side.
- Network infrastructure devices. This includes the routers, switches, hubs, and other network equipment that makes communications between the computers on the network possible.
- **Current network traffic model.** This includes the quantity and the characteristics of the network traffic flowing through your network.
- Intrusion Detection System (IDS) devices. You will need to identify if you have any IDS devices on your network that might be negatively impacted by any encryption introduced in an Encryption Zone.

The goal is to have enough information to be able to identify an asset by its network location, in addition to its physical location.

Do not use a complex and poorly documented network as a starting point for the design, because it can leave too many unidentified areas that are likely to cause problems during implementation.

This guidance helps obtain the most relevant information for planning Windows Firewall with Advanced Security implementation, but it does not try to address other issues, such as TCP/IP addressing or virtual local area network (VLAN) segmentation.

Network segmentation

If your organization does not have its current network architecture documented and available for reference, such documentation should be obtained as soon as possible before you continue with the design and deployment. If the documented information is not current or has not been validated recently, you have two options:

- Accept that the lack of accurate information can cause risk to the project.
- Undertake a discovery project, either through manual processes or with network analysis tools that can provide the information you need to document the current network topology.

Although the required information can be presented in many different ways, a series of schematic diagrams is often the most effective method of illustrating and understanding the current network configuration. When creating network diagrams, do not include too much information. If necessary, use multiple diagrams that show different layers of detail. Use a top-level diagram that illustrates the major sites that make up your organization's network, and then break out each site into a more detailed diagram that captures a deeper level of detail. Continue until you reach the individual IP subnet level, and so have the means to identify the network location of every computer in your organization.

During this process, you might discover some network applications and services that are not compatible with IPsec. For example, IPsec breaks network-based prioritization and port/protocol-based traffic management. If traffic management or prioritization must be based on ports or protocol, the host itself must be able to perform any traffic management or prioritization.

Other examples of incompatibility include:

- Cisco NetFlow on routers cannot analyze packets between IPsec members based on protocol or port.
- Router-based Quality of Service (QoS) cannot use ports or protocols to prioritize traffic. However, using firewall rules that specify IP addresses to prioritize traffic are not affected by this limitation of QoS. For example, a rule that says "From anyone to anyone using port 80 prioritize" does not work, but a rule that says "From anyone to 10.0.1.10 prioritize" works.
- Weighted Fair Queuing and other flow-based router traffic priority methods might fail.
- Devices that do not support or allow IP protocol 50, the port that is used by Encapsulating Security Payload (ESP).
- Router access control lists (ACLs) cannot examine protocol and port fields in ESP-encrypted packets, and therefore the packets are dropped. ACLs based only on IP address are forwarded as usual. If the device cannot parse ESP, any ACLs that specify port or protocol rules will not be processed on the ESP packets. If the device has an ESP parser and uses encryption, ACLs that specify port or protocol rules will not be processed on the ESP packets.
- Network monitoring tools might be unable to parse ESP packets that are not encrypted (ESP-Null).

  Note
  Network Monitor added an ESP parser starting in version 2.1 to aid troubleshooting of unencrypted IPsec packets. The latest version of Network Monitor is available as a free download from Microsoft (http://go.microsoft.com/fwlink/?linkid=94770).

Network address translation (NAT)

IPsec NAT traversal (NAT-T) enables IPsec peers that are behind NATs to detect the presence of NATs, negotiate IPsec security associations (SAs), and send ESP-protected data even though the addresses in the IPsec-protected IPv4 packets change. IPsec NAT-T does not support the use of AH across NAT devices.

IPsec NAT-T is supported by Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, Windows Server 2008 R2,

For detailed information about how IPsec NAT-T works, see "IPsec NAT Traversal Overview" in the August 2002 Cable Guy article at http://go.microsoft.com/fwlink/?LinkId=45080.

Network infrastructure devices

The devices that make up the network infrastructure (routers, switches, load balancers, and firewalls) must be able communicate using IPsec after the solution is implemented. For this

reason, you have to examine the following characteristics of these network devices to ensure that they can handle the technical and physical requirements of the design:

- **Make/model**. You can use this information to determine the features that the device supports. In addition, check the BIOS version or software running on the device to ensure that IPsec is supported.
- **Amount of RAM**. This information is useful when you are analyzing capacity or the impact of IPsec on the device.
- **Traffic analysis**. Information, such as peak usage and daily orweekly trends, is helpful to have. The information helps provide a baseline snapshot of the device and how it is used over time. If problems occur after IPsec is implemented, the information can help determine whether the root cause is related to greater usage of the device.
- **Router ACLs that affect IPsec directly**. ACLs directly affect the ability of specific protocols to function. For example, blocking the Kerberos V5 protocol (UDP and TCP port 88) or IP protocol 50 or 51 prevents IPsec from working. Devices must also be configured to allow IKE traffic (UDP port 500) if using NAT-T (UDP port 4500).
- **Networks/subnets connected to device interfaces**. This information provides the best picture of what the internal network looks like. Defining the boundary of subnets based on an address range is straightforward and helps identify whether other addresses are either unmanaged or foreign to the internal network (such as IP addresses on the Internet).
- **VLAN segmentation**. Determining how VLANs are implemented on the network can help you understand traffic patterns and security requirements, and then help to determine how IPsec might augment or interfere with these requirements.
- **The maximum transmission unit (MTU) size on device interface(s)**. The MTU defines the largest datagram that can be transmitted on a particular interface without being divided into smaller pieces for transmission (a process also known as *fragmentation*). In IPsec communications, the MTU is necessary to anticipate when fragmentation occurs. Packet fragmentation must be tracked for Internet Security Association and Key Management Protocol (ISAKMP) by the router. IPsec configures the MTU size on the session to the minimum-discovered MTU size along the communication path being used, and then set the Don't Fragment bit (DF bit) to 1.

  Note
  If Path MTU (PMTU) discovery is enabled and functioning correctly, you do not have to gather the MTU size on device interfaces. Although sources, such as the Windows Server 2003 Hardening Guide, recommend disabling PMTU discovery, it must be enabled for IPsec to function correctly.
- **Intrusion detection system (IDS) in use**. Your IDS must have an IPsec-compatible parser to detect ESP packets. If the IDS does not have such a parser, it cannot determine if data in those packets is encrypted.

After you obtain this information, you can quickly determine whether you must upgrade the devices to support the requirements of the project, change the ACLs, or take other measures to ensure that the devices can handle the loads needed.

Current network traffic model

After gathering the addressing and network infrastructure information, the next step is to examine the communications flow. For example, if a department such as Human Resources (HR) spans several buildings, and you want to use server isolation with encryption to help protect information in that department, you must know how those buildings are connected to determine the level of "trust" to place in the connection. A highly secured building that is connected by an unprotected cable to another building that is not secured can be compromised by an eavesdropping or information replay attack. If such an attack is considered a threat, IPsec can help by providing strong mutual authentication and traffic encryption for trusted hosts. IPsec allows you to more securely communicate across untrusted links such as the Internet.

When you examine traffic flow, look closely at how all managed and unmanaged devices interact. This includes non-Windows-based computers running Linux, UNIX, and Macintosh. Ask yourself such questions as:

- Do specific communications occur at the port and protocol level, or are there many sessions between the same hosts across many protocols?
- How do servers and clients communicate with each other?
- Are there security devices or projects currently implemented or planned that could affect an isolation deployment? For example, if you use Windows Firewall on your computers to "lock down" specific ports, such as UDP 500, IKE negotiations fail.

Some of the more common applications and protocols are as follows:

- **NetBIOS over TCP/IP (NetBT) and server message block (SMB)**. On a LAN, it is common to have ports 137, 138, and 139 enabled for NetBT and port 445 enabled for SMB. These ports provide NetBIOS name resolution services and other features. Unfortunately, they also allow the creation of *null sessions*. A null session is a session that is established on a host that does not use the security context of a known user or entity. Frequently, these sessions are anonymous.
- **Remote procedure call (RPC)**. RPC operates by listening on a port known as the *endpoint mapper*, TCP port 135. The response to a query on this port is an instruction to begin communication on another port in the ephemeral range (ports numbered over 1024). In a network that is segmented by firewalls, RPC communication presents a configuration challenge because it means opening the RPC listener port and all ports greater than 1024. Opening so many ports increases the attack surface of the whole network and reduces the effectiveness of the firewalls. Computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 reduce this risk by introducing stateful inspection of RPC traffic. Because many applications depend on RPC for basic functionality, any firewall and connection security policy must take RPC requirements into account.
- **Other traffic**. Windows Firewall with Advanced Security can help secure transmissions between computers by providing authentication of the packets in addition to encrypting the data that they contain. The important thing to do is to identify what must be protected, and the threats that must be mitigated. Examine and model other traffic or traffic types that must be secured.

Next: Gathering Information about Your Active Directory Deployment

Gathering Information about Your Active Directory Deployment

Active Directory is another important item about which you must gather information. You must understand the forest structure. This includes domain layout, organizational unit (OU) architecture, and site topology. This information makes it possible to know where computers are currently placed, their configuration, and the impact of changes to Active Directory that result from implementing Windows Firewall with Advanced Security. Review the following list for information needed:

- **Names and number of forests**. The forest (not the domain) is the security boundary in an Active Directory implementation. You must understand the current Active Directory architecture to determine the most effective strategy for deploying your firewall and connection security rules using Group Policy. It also enables you to understand which computers can be isolated and how best to accomplish the required degree of isolation.
- **Names and number of domains**. Authentication in server and domain isolation uses the IKE negotiation process with the Kerberos V5 protocol. This protocol assumes that computers are domain members.
- **Number and types of trusts**. Trusts affect the logical boundaries of domain isolation and define whether IKE negotiation can occur between computers in different Active Directory domains.
- **Names and number of sites**. Site architecture is usually aligned with the network topology. Understanding how sites are defined in Active Directory will help provide insight into replication and other details. Site architecture can provide a better understanding of the current Active Directory deployment.
- **OU structure**. OUs are logical constructs and can therefore be molded to fit many different requirements and goals. The OU structure is an ideal place to examine how Group Policy is currently used and how the OUs are laid out. You do not have to redesign an already implemented OU structure in order to effectively deploy firewall and connection security policy, but an understanding of the structure helps you know what WMI or group filtering is required to apply each GPO to the correct computers.
- **Existing IPsec policy**. Because this project culminates in the implementation of IPsec policy, you must understand how the network currently uses IPsec (if at all). Windows Firewall with Advanced Security connection security rules for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 are not compatible with earlier versions of Windows. If you already have IPsec policies deployed to computers running Windows XP and Windows Server 2003 in your organization, you must ensure that the new IPsec policies you deploy enable computers using either the old or new IPsec policies to communicate with each other.

Next: Gathering Information about Your Computers

Gathering Information about Your Computers

One of the most valuable benefits of conducting an asset discovery project is the large amount of data that is obtained about the client and server computers on the network. When you start designing and planning your isolation zones, you must make decisions that require accurate information about the state of all hosts to ensure that they can use IPsec as planned.

Capture the following information from each computer:

- **Computer name**. This name is the computer's NetBIOS or DNS name that identifies the computer on the network. Because a computer can have more than one media access control (MAC) or IP address, the computer's name is one of the criteria that can be used to determine uniqueness on the network. Because computer names can be duplicated under some circumstances, the uniqueness should not be considered absolute.
- **IP address for each network adapter**. The IP address is the address that is used with the subnet mask to identify a host on the network. An IP address is not an effective way to identify an asset because it is often subject to change.
- **Operating system, service pack, and hotfix versions**. The operating system version is a key factor in determining the ability of a host to communicate by using IPsec. It is also important to track the current state of service packs and updates that might be installed, because these are often used to determine that minimum security standards have been met.
- **Domain membership**. This information is used to determine whether a computer can obtain IPsec policy from Active Directory or whether it must use a local IPsec policy.
- **Physical location**. This information is just the location of the device in your organization. It can be used to determine whether a device can participate in a specific isolation group based on its location or the location of the devices that it communicates with regularly.
- **Hardware type or role**. Some tools that perform host discovery can provide this information by querying the hardware information and running applications to determine its type, such as server, workstation, or portable computer. You can use this information to determine the appropriate IPsec policy to assign, whether a specific computer can participate in isolation, and in which isolation group to include the computer.

After collecting all this information and consolidating it into a database, perform regular discovery efforts periodically to keep the information current. You need the most complete and up-to-date picture of the managed hosts on their networks to create a design that matches your organization's requirements.

You can use various methods to gather data from the hosts on the network. These methods range from high-end, fully automated systems to completely manual data collection. Generally, the use of automated methods to gather data is preferred over manual methods for reasons of speed and accuracy.

Automated Discovery

Using an automated auditing network management system such as Microsoft System Center Configuration Manager (formerly known as Systems Management Server) provides valuable information about the current state of the IT infrastructure.

For more information about how System Center Configuration Manager 2007 can help perform automated information gathering, see http://go.microsoft.com/fwlink/?linkid=110412.

Manual Discovery

The biggest difference between manual discovery methods and automated methods is time.

You can use the Windows Script Host (WSH), VBScript, and Windows Management Instrumentation (WMI) to create a script file that can collect the system configuration information. VBScript and WMI are built-in to Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. Starting with Windows Server 2008, Windows PowerShell is included with the operating system. For more information, see “Scripting with Windows PowerShell” (http://go.microsoft.com/fwlink/?linkid=110413).

Whether you use an automatic, manual, or hybrid option to gather the information, one of the biggest issues that can cause problems to the design is capturing the changes between the original inventory scan and the point at which the implementation is ready to start. After the first scan has been completed, make support staff aware that all additional changes must be recorded and the updates noted in the inventory.

This inventory will be critical for planning and implementing your Windows Firewall with Advanced Security design.

Next: Gathering Other Relevant Information

Gathering Other Relevant Information

This topic discusses several other things that you should examine to see whether they will cause any complications in your ability to deploy Windows Firewall with Advanced Security policies in your organization.

Capacity considerations

Because IPsec uses mathematically intensive cryptographic techniques, it can consume significant overhead on a computer. Areas to watch:

- **Encryption.** You might use 256-bit Advanced Encryption Standard (AES-256) and 384-bit Secure Hash Algorithm (SHA-384) to check integrity in situations that require the strongest available encryption and key exchange protection. If you have NICs that support IPsec Task Offload, you can reduce the effect that encryption has on network throughput. For more information, see IPsec Task Offload at http://technet.microsoft.com/network/dd277647.aspx

- **Security association (SA) negotiation.** You can use a shorter lifetime for the main mode SA, such as three hours, but then you might need to make tradeoffs. Because each main mode SA occupies approximately 5 KB of RAM, situations in which a server brokers tens of thousands of concurrent connections can lead to overutilization.
- **NAT devices.** As discussed earlier, NAT does not allow Authentication Header (AH) conversations between hosts. If NAT devices exist on the internal network, ESP must be selected instead of AH.
- **Switches and routers.** Proper capacity planning for the implementation of IPsec is more about thorough testing and expected traffic loads than exact calculations. You might have to upgrade or reconfigure switches or routers that currently exceed 75 percent usage to allow for increased traffic on the device and still provide some extra usage for bursts of traffic.
- **Other factors.** These include CPU usage on network infrastructure servers, increased overhead on servers and workstations running IPsec (especially servers, because they usually contain more main mode SAs than clients), and increased network latency because of IPsec negotiation.

  Note
  When Microsoft deployed its own domain isolation solution, it found a one to three percent increase in usage on the network as a direct result of IPsec.

Group Policy deployment groups and WMI filters

You do not have to rearrange the organization unit (OU) hierarchy of your Active Directory domains to effectively deploy Windows Firewall with Advanced Security GPOs. Instead, you can link your GPOs at the domain level (or another high level container), and then use security group filtering or WMI filtering to ensure that only the appropriate computers or users can apply the GPO settings. Because the firewall and connection security rules have evolved significantly from Windows 2000 Server to Windows XP and Windows Server 2003, and now with Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2, we recommend that you use WMI filtering to dynamically ensure that GPOs apply only to computers that are running the correct operating system. It is not necessary to use this technique if your network consists of computers running Windows Vista or later.

Different Active Directory trust environments

When you design a domain isolation policy, consider any logical boundaries that might affect IPsec-secured communications. For example, the trust relationships between your domains and forests are critical in determining an appropriate IKE authentication method.

Kerberos V5 authentication is recommended for use in a two-way (mutual) domain and forest trust environment. You can use Kerberos V5 for IKE authentication across domains that have two-way trusts established, if the domains are in the same forest or different forests. If the two domains are in different forests, you must configure two external trusts, one for each direction, between the domains. The external trusts must use the fully qualified domain name (FQDN) of the domains, and IPsec policy must allow an IKE initiator in one domain to communicate with any

domain controller in the forest domain hierarchy, so that the initiator can obtain a Kerberos V5 ticket from a domain controller in the responder's domain. If firewalls separate the domains then you must configure the firewall to allow Kerberos V5 traffic over UDP destination port 88, TCP destination port 88, and UDP destination port 389.

For more information, see "Active Directory in Networks Segmented by Firewalls" at http://go.microsoft.com/fwlink/?LinkId=45087.

If the use of Kerberos V5 authentication is not possible because two-way trusts across forests cannot be established as in some large enterprise environments, you can use a public key infrastructure (PKI) and digital certificates to establish IPsec-trusted communication. For an example of how Microsoft deployed their PKI, see "Deploying PKI Inside Microsoft" at http://go.microsoft.com/fwlink/?LinkId=45088.

Creating firewall rules to permit IKE, AH, and ESP traffic

In some cases, IPsec-secured traffic might have to pass through a router, perimeter firewall, or other filtering device. In the case of a router, unless the router filters TCP and UDP traffic or other upper-level protocol headers, no special configuration is required to allow the IPsec traffic to be forwarded.

In the case of a filtering router or a firewall, you must configure these devices to allow IPsec traffic to be forwarded. Configure the firewall to allow IPsec traffic on UDP source and destination port 500 (IKE), UDP source and destination port 4500 (IPsec NAT-T), and IP Protocol 50 (ESP). You might also have to configure the firewall to allow IPsec traffic on IP protocol 51 (AH) to allow troubleshooting by IPsec administrators and to allow the IPsec traffic to be inspected.

For more information, see "How to Enable IPsec Traffic Through a Firewall" at http://go.microsoft.com/fwlink/?LinkId=45085.

Network load balancing and server clusters

There are challenges implementing connection security for network traffic going to and from network load balancing (NLB) clusters and server clusters. NLB enables multiple servers to be clustered together to provide high availability for a service by providing automatic failover to other nodes in the cluster. Because IPsec matches a security association to a specific computer, it prevents different computers from handling the same client connection. If a different node in the cluster responds to an IPsec connection that was originally established by another node, the traffic will be dropped by the client computer as untrusted.

This means that NLB in "no affinity" mode is not supported by IPsec at all. If you must use "no affinity" mode in the cluster then consider including the servers that make up the cluster in your IPsec exemption group, and allowing clients to communicate with the servers without IPsec.

IPsec improvements for clusters running Windows Server 2008

Starting with Windows Server 2008 and Windows Vista, IPsec is much more tightly integrated into TCP/IP than in earlier versions of Windows. When a TCP connection is dropped because of a cluster node failover, IPsec on a computer that is running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 detects the TCP connection failure and removes the IPsec SAs for that connection. When the new TCP connection is established to another node, IPsec can negotiate new SAs immediately without having to wait for the obsolete SAs to time out.

Network inspection technologies

Within a TCP/IP packet, IPsec without encryption changes the offsets for the destination ports and protocols. These changes can adversely affect applications that are running on network devices such as routers that monitor and manage traffic on the network. While some network applications have been updated to support IPsec, some are not yet compatible. Check with the vendor of your device to see whether the changes in the protocol and port fields caused by IPsec are compatible with the device.

Any device designed to view network traffic, such as hardware protocol analyzers or Microsoft Network Monitor, cannot parse ESP-encrypted traffic. Only the destination computer, with which the originating computer negotiated the connection, can decrypt the traffic.

In general, IPsec defeats network-based prioritization and port- or protocol-based traffic management. For encrypted packets, there is no workaround; the host itself must handle any traffic management functions. For unencrypted, authenticated-only packets, the devices and applications must be aware of how IPsec changes packets to be able to do anything with them other than route them to the correct host. If you cannot upgrade monitoring or management devices to support IPsec, it is important that you record this information and figure it into your domain or server isolation design.

Network Monitor includes parsers for the ISAKMP (IKE), AH, and ESP protocols. Network Monitor parsers for ESP can parse inside the ESP packet only if ESP null-encryption is being used. Network Monitor cannot parse the encrypted parts of IPsec ESP traffic when encryption is performed in software. However, if encryption is performed by an IPsec hardware offload network adapter, the ESP packets can be decrypted when Network Monitor captures them on either the source or the destination and, therefore, they can be parsed. To diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPsec policy or connection security rule on both computers.

Network Monitor is available as a free download from Microsoft at http://go.microsoft.com/fwlink/?linkid=94770.

Next: Determining the Trusted State of Your Computers

Determining the Trusted State of Your Computers

After obtaining information about the computers that are currently part of the IT infrastructure, you must determine at what point a computer is considered trusted. The term *trusted* can mean different things to different people. Therefore, you must communicate a firm definition for it to all stakeholders in the project. Failure to do this can lead to problems with the security of the trusted environment, because the overall security cannot exceed the level of security set by the least secure client that achieves trusted status.

Note
In this context, the term *trust* has nothing to do with an Active Directory trust relationship between domains. The trusted state of your computers just indicates the level of risk that you believe the computer brings to the network. Trusted computers bring little risk whereas untrusted computers can potentially bring great risk.

Trust states

To understand this concept, consider the four basic states that apply to computers in a typical IT infrastructure. These states are (in order of risk, lowest risk first):

- Trusted
- Trustworthy
- Known, untrusted
- Unknown, untrusted

The remainder of this section defines these states and how to determine which computers in your organization belong in each state.

Trusted state

Classifying a computer as trusted means that the computer's security risks are managed, but it does not imply that it is perfectly secure or invulnerable. The responsibility for this managed state falls to the IT and security administrators, in addition to the users who are responsible for the configuration of the computer. A trusted computer that is poorly managed will likely become a point of weakness for the network.

When a computer is considered trusted, other trusted computers can reasonably assume that the computer will not initiate a malicious act. For example, trusted computers can expect that other trusted computers will not run a virus that attacks them, because all trusted computers are required to use mechanisms (such as antivirus software) to mitigate the threat of viruses.

Spend some time defining the goals and technology requirements that your organization considers appropriate as the minimum configuration for a computer to obtain trusted status.

A possible list of technology requirements might include the following:

- **Operating system.** A trusted client computer should run Windows 8, Windows 7, or Windows Vista. A trusted server should run Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008.
- **Domain membership.** A trusted computer will belong to a managed Active Directory domain, which means that the IT department has security management rights and can configure member computers by using Group Policy.
- **Management client.** All trusted computers must run a specific network management client to allow for centralized management and control of security policies, configurations, and software. Microsoft System Center Configuration Manager is one such management system with an appropriate client. For more information, see System Center Configuration Manager at http://technet.microsoft.com/systemcenter/bb507744.aspx.
- **Antivirus software.** All trusted computers will run antivirus software that is configured to check for and automatically update the latest virus signature files daily. Microsoft ForeFront Endpoint Protection is one such antivirus software program. For more information, see ForeFront Endpoint Protection at http://technet.microsoft.com/forefront/ee822838.aspx.
- **File system.** All trusted computers will be configured to use the NTFS file system.
- **BIOS settings.** All trusted portable computers will be configured to use a BIOS-level password that is under the management of the IT support team.
- **Password requirements.** Trusted clients must use strong passwords.

It is important to understand that the trusted state is not constant; it is a transient state that is subject to changing security standards and compliance with those standards. New threats and new defenses emerge constantly. For this reason, the organization's management systems must continually check the trusted computers to ensure ongoing compliance. Additionally, the management systems must be able to issue updates or configuration changes if they are required to help maintain the trusted status.

A computer that continues to meet all these security requirements can be considered trusted. However it is possible that most computers that were identified in the discovery process discussed earlier do not meet these requirements. Therefore, you must identify which computers can be trusted and which ones cannot. To help with this process, you use the intermediate *trustworthy* state. The remainder of this section discusses the different states and their implications.

Trustworthy state

It is useful to identify as soon as possible those computers in your current infrastructure that can achieve a trusted state. A *trustworthy state* can be assigned to indicate that the current computer can physically achieve the trusted state with required software and configuration changes.

For each computer that is assigned a trustworthy status, make an accompanying configuration note that states what is required to enable the computer to achieve trusted status. This information is especially important to both the project design team (to estimate the costs of adding the computer to the solution) and the support staff (to enable them to apply the required configuration).

Generally, trustworthy computers fall into one of the following two groups:

- **Configuration required.** The current hardware, operating system, and software enable the computer to achieve a trustworthy state. However, additional configuration changes are required. For example, if the organization requires a secure file system before a computer can be considered trusted, a computer that uses a FAT32-formatted hard disk does not meet this requirement.
- **Upgrade required.** These computers require upgrades before they can be considered trusted. The following list provides some examples of the type of upgrade these computers might require:
    - **Operating system upgrade required.** If the computer's current operating system cannot support the security needs of the organization, an upgrade would be required before the computer could achieve a trusted state.
    - **Software required.** A computer that is missing a required security application, such as an antivirus scanner or a management client, cannot be considered trusted until these applications are installed and active.
    - **Hardware upgrade required.** In some cases, a computer might require a specific hardware upgrade before it can achieve trusted status. This type of computer usually needs an operating system upgrade or additional software that forces the required hardware upgrade. For example, security software might require additional hard disk space on the computer.
    - **Computer replacement required.** This category is reserved for computers that cannot support the security requirements of the solution because their hardware cannot support the minimum acceptable configuration. For example, a computer that cannot run a secure operating system because it has an old processor (such as a 100-megahertz [MHz] x86-based computer).

Use these groups to assign costs for implementing the solution on the computers that require upgrades.

Known, untrusted state

During the process of categorizing an organization's computers, you will identify some computers that cannot achieve trusted status for specific well-understood and well-defined reasons. These reasons might include the following types:

- **Financial.** The funding is not available to upgrade the hardware or software for this computer.
- **Political.** The computer must remain in an untrusted state because of a political or business situation that does not enable it to comply with the stated minimum security requirements of the organization. It is highly recommended that you contact the business owner or independent software vendor (ISV) for the computer to discuss the added value of server and domain isolation.
- **Functional.** The computer must run a nonsecure operating system or must operate in a nonsecure manner to perform its role. For example, the computer might be required to run an older operating system because a specific line of business application will only work on that operating system.

There can be multiple functional reasons for a computer to remain in the known untrusted state. The following list includes several examples of functional reasons that can lead to a classification of this state:

- **Computers that run unsupported versions of Windows.** This includes Windows XP, Windows Millennium Edition, Windows 98, Windows 95, or Windows NT. Computers that run these versions of the Windows operating system cannot be classified as trustworthy because these operating systems do not support the required security infrastructure. For example, although Windows NT does support a basic security infrastructure, it does not support “deny” ACLs on local resources, any way to ensure the confidentiality and integrity of network communications, smart cards for strong authentication, or centralized management of computer configurations (although limited central management of user configurations is supported).
- **Stand-alone computers.** Computers running any version of Windows that are configured as stand-alone computers or as members of a workgroup usually cannot achieve a trustworthy state. Although these computers fully support the minimum required basic security infrastructure, the required security management capabilities are unlikely to be available when the computer is not a part of a trusted domain.
- **Computers in an untrusted domain.** A computer that is a member of a domain that is not trusted by an organization's IT department cannot be classified as trusted. An untrusted domain is a domain that cannot provide the required security capabilities to its members. Although the operating systems of computers that are members of this untrusted domain might fully support the minimum required basic security infrastructure, the required security management capabilities cannot be fully guaranteed when computers are not in a trusted domain.

Unknown, untrusted state

The unknown, untrusted state should be considered the default state for all computers. Because computers in this state have a configuration that is unknown, you can assign no trust to them. All planning for computers in this state must assume that the computer is an unacceptable risk to the organization. Designers of the solution should strive to minimize the impact that the computers in this state can have on their organizations.

Capturing upgrade costs for current computers

The final step in this part of the process is to record the approximate cost of upgrading the computers to a point that they can participate in the server and domain isolation design. You must make several key decisions during the design phase of the project that require answers to the following questions:

- Does the computer meet the minimum hardware requirements necessary for isolation?
- Does the computer meet the minimum software requirements necessary for isolation?
- What configuration changes must be made to integrate this computer into the isolation solution?

- What is the projected cost or impact of making the proposed changes to enable the computer to achieve a trusted state?

By answering these questions, you can quickly determine the level of effort and approximate cost of bringing a particular computer or group of computers into the scope of the project. It is important to remember that the state of a computer is transitive, and that by performing the listed remedial actions you can change the state of a computer from untrusted to trusted. After you decide whether to place a computer in a trusted state, you are ready to begin planning and designing the isolation groups, which the next section Planning Domain Isolation Zones discusses.

The following table is an example of a data sheet that you could use to help capture the current state of a computer and what would be required for the computer to achieve a trusted state.

Computer name	Hardware reqs met	Software reqs met	Configuration required	Details	Projected cost
CLIENT001	No	No	Upgrade hardware and software.	Current operating system is Windows XP. Old hardware is not compatible with Windows 8.	$??
SERVER001	Yes	No	Join trusted domain and upgrade from Windows Server 2003 to Windows Server 2012.	No antivirus software present.	$??

In the previous table, the computer CLIENT001 is currently "known, untrusted" because its hardware must be upgraded. However, it could be considered trustworthy if the required upgrades are possible. However, if many computers require the same upgrades, the overall cost of the solution would be much higher.

The computer SERVER001 is "trustworthy" because it meets the hardware requirements but its operating system must be upgraded. It also requires antivirus software. The projected cost is the amount of effort that is required to upgrade the operating system and install antivirus software, along with their purchase costs.

With the other information that you have gathered in this section, this information will be the foundation of the efforts performed later in the Planning Domain Isolation Zones section.

The costs identified in this section only capture the projected cost of the computer upgrades. Many additional design, support, test, and training costs should be accounted for in the overall project plan.

For more information about how to configure firewalls to support IPsec, see "Configuring Firewalls" at http://go.microsoft.com/fwlink/?linkid=69783.

For more information about WMI, see "Windows Management Instrumentation" at http://go.microsoft.com/fwlink/?linkid=110483.

Next: Planning Your Windows Firewall with Advanced Security Design

Planning Your Windows Firewall with Advanced Security Design

After you have gathered the relevant information in the previous sections, and understand the basics of the designs as described earlier in this guide, you can select the design (or combination of designs) that meet your needs.

Basic firewall design

We recommend that you deploy at least the basic firewall design. As discussed in the Protect Computers from Unwanted Network Traffic section, host-based firewalls are an important element in a defense-in-depth strategy and complement most other security measures you put in place in your organization.

When you are ready to examine the options for firewall policy settings, see the Planning Settings for a Basic Firewall Policy section.

Algorithm and method support and selection

To create a domain isolation or server isolation design, you must understand the algorithms available in each version of Windows, as well as their relative strengths. To review the algorithms and methods supported in versions of the Windows operating system, see IPsec Algorithms and Methods Supported in Windows (http://go.microsoft.com/fwlink/?linkid=129230).

IPsec performance considerations

Although IPsec is critically important in securing network traffic going to and from your computers, there are costs associated with its use. The mathematically intensive cryptographic algorithms require a significant amount of computing power, which can prevent your computer from making use of all of the available bandwidth. For example, an IPsec-enabled computer using the AES encryption protocols on a 10 gigabits per second (Gbps) network link might see a throughput of 4.5 Gbps. This is due to the demands placed on the CPU to perform the cryptographic functions required by the IPsec integrity and encryption algorithms.

IPsec task offload is a Windows technology that supports network adapters equipped with dedicated cryptographic processors to perform the computationally intensive work required by IPsec. This frees up a computer's CPU and can dramatically increase network throughput. For

the same network link as above, the throughput with IPsec task offload enabled improves to about 9.2 Gbps. For more information, see Improving Network Performance by Using IPsec Task Offload (http://go.microsoft.com/fwlink/?linkid=129229).

Domain isolation design

Include this design in your plans:

- If you have an Active Directory domain of which most of the computers are members.
- If you want to prevent the computers in your organization from accepting any unsolicited network traffic from computers that are not part of the domain.

If you plan on including the basic firewall design as part of your deployment, we recommend that you deploy the firewall policies first to confirm that they work properly. Also plan to enable your connection security rules in request mode at first, instead of the more restrictive require mode, until you are sure that the computers are all correctly protecting network traffic with IPsec. If something is wrong, request mode still allows communications to continue while you are troubleshooting.

When you are ready to examine the options for creating an isolated domain, see the Planning Domain Isolation Zones section.

Server isolation design

Include this design in your plans:

- If you have an isolated domain and you want to additionally restrict access to specific servers to only authorized users and computers.
- You are not deploying an isolated domain, but want to take advantage of similar benefits for a few specific servers. You can restrict access to the isolated servers to only authorized users and computers.

If you plan to include domain isolation in your deployment, we recommend that you complete that layer and confirm its correct operation before you implement the additional server isolation elements.

When you are ready to examine the options for isolating servers, see the Planning Server Isolation Zones section.

Certificate-based authentication design

Include this design in your plans:

- If you want to implement some of the elements of domain or server isolation on computers that are not joined to an Active Directory domain, or do not want to use domain membership as an authentication mechanism.
- You have an isolated domain and want to include a server that is not a member of the Active Directory domain because the computer is not running Windows, or for any other reason.

- You must enable external computers that are not managed by your organization to access information on one of your servers, and want to do this in a secure way.

If you plan to include domain or server isolation in your deployment, we recommend that you complete those elements and confirm their correct operation before you add certificate-based authentication to the computers that require it.

When you are ready to examine the options for using certificate-based authentication, see the Planning Certificate-based Authentication section.

Documenting your design

After you finish selecting the designs that you will use, you must assign each of your computers to the appropriate isolation zone and document the assignment for use by the deployment team.

- Documenting the Zones

Designing groups and GPOs

After you have selected a design and assigned your computers to zones, you can begin laying out the isolation groups for each zone, the network access groups for isolated server access, and the GPOs that you will use to apply the settings and rules to your computers.

When you are ready to examine the options for the groups, filters, and GPOs, see the Planning Group Policy Deployment for Your Isolation Zones section.

Next: Planning Settings for a Basic Firewall Policy

Planning Settings for a Basic Firewall Policy

After you have identified your requirements, and have the information about the network layout and computers available, you can begin to design the GPO settings and rules that will enable you to enforce your requirements on the computers.

The following is a list of the firewall settings that you might consider for inclusion in a basic firewall design, together with recommendations to serve as a starting point for your analysis:

- **Profile selection**. The firewall rules can be configured for any of the network location profiles that you see in the Network and Sharing Center: **Domain**, **Public**, and **Private** (on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2). Most settings are enforced in the Domain profile, without an option for the user to change them. However, you might want to leave the profile settings configurable by the user on computers that can be taken from the organization's physical network and joined to a public or home network. If you lock down the public and private profiles, you might prevent a user from accessing a required network program or service. Because they are not on the organization's network, you cannot fix a connectivity problem by deploying rule changes in a GPO. For each section that follows, consider each profile and apply the rules to those profiles that make sense for your organization.

Important
We recommend that on server computers that you set all rules for all profiles to prevent any unexpected profile switch from disrupting network connectivity. You might consider a similar practice for your desktop computers, and only support different profiles on portable computers.

- **Firewall state: On**. We recommend that you prevent the user from turning it off.
- **Default behavior for Inbound connections: Block**. We recommend that you enforce the default behavior of blocking unsolicited inbound connections. To allow network traffic for a specific program, create an inbound rule that serves as an exception to this default behavior.
- **Default behavior for Outbound connections: Allow**. We recommend that you enforce the default behavior of allowing outbound connections.
- **Allow unicast response: Yes**. We recommend that you use the default setting of **Yes** unless you have specific requirements to do otherwise.
- **Apply local firewall rules: Yes**. We recommend that you allow users to create and use local firewall rules. If you set this to **No**, then when a user clicks **Allow** on the notification message to allow traffic for a new program, Windows does not create a new firewall rule and the traffic remains blocked.

  If you and the IT staff can create and maintain the list of firewall rules for all permitted applications and deploy them by using GPOs then you can set this value to **No**.
- **Apply local connection security rules: No**. We recommend that you prevent users from creating and using their own connection security rules. Connection failures caused by conflicting rules can be difficult to troubleshoot.
- **Logging**. We recommend that you enable logging to a file on the local hard disk. Be sure to limit the size, such as 4096 KB, to avoid causing performance problems by filling the user's hard disk. Be sure to specify a folder to which the Windows Firewall service account has write permissions.
- **Inbound rules**. Create inbound rules for programs that must be able to receive unsolicited inbound network packets from another computer on the network. Make the rules as specific as possible to reduce the risk of malicious programs exploiting the rules. For example, specify both program and port numbers. Specifying a program ensures that the rule is only active when the program is actually running, and specifying the port number ensures that the program cannot receive unexpected traffic on a different port.

  Inbound rules are common on servers, because they host services to which client computers connect. When you install programs and services on a server, the installation program typically creates and enables the rules for you. Examine the rules to ensure that they do not open up more ports than are required.

  Important
  If you create inbound rules that permit RPC network traffic by using the **RPC Endpoint Mapper** and **Dynamic RPC** rule options, then all inbound RPC network traffic is permitted because the firewall cannot filter network traffic based on the UUID of the destination application.

- **Outbound rules**. Only create outbound rules to block network traffic that must be prevented in all cases. If your organization prohibits the use of certain network programs, you can support that policy by blocking the known network traffic used by the program. Be sure to test the restrictions before you deploy them to avoid interfering with traffic for needed and authorized programs.

Next: Planning Domain Isolation Zones

Planning Domain Isolation Zones

After you have the required information about your network, Active Directory, and client and server computers, you can use that information to make decisions about the isolation zones you want to use in your environment.

The bulk of the work in planning server and domain isolation is determining which computers to assign to each isolation zone. Correctly choosing the zone for each computer is important to providing the correct level of security without compromising performance or the ability a computer to send or receive required network traffic.

The zones described in this guide include the following:

- Exemption List
- Isolated Domain
- Boundary Zone
- Encryption Zone

Exemption List

When you implement a server and domain isolation security model in your organization, you are likely to find some additional challenges. Key infrastructure servers such as DNS servers and DHCP servers typically must be available to all computers on the internal network, yet secured from network attacks. However, if they must remain available to all computers on the network, not just to isolated domain members, then these servers cannot require IPsec for inbound access, nor can they use IPsec transport mode for outbound traffic.

In addition to the infrastructure servers mentioned earlier, there might also be other servers on the network that trusted computers cannot use IPsec to access, which would be added to the exemption list.

Generally, the following conditions are reasons to consider adding a computer to the exemption list:

- If the computer must be accessed by trusted computers but it does not have a compatible IPsec implementation.
- If the computer must provide services to both trusted and untrusted computers, but does not meet the criteria for membership in the boundary zone.

- If the computer must be accessed by trusted computers from different isolated domains that do not have an Active Directory trust relationship established with each other.
- If the computer is a domain controller running version of Windows earlier than Windows Server 2008, or if any of its clients are running a version of Windows earlier than Windows Vista.
- If the computer must support trusted and untrusted computers, but cannot use IPsec to help secure communications to trusted computers.

For large organizations, the list of exemptions might grow very large if all the exemptions are implemented by one connection security rule for the whole domain or for all trusted forests. If you can require all computers in your isolated domain to run at least Windows Vista or Windows Server 2008, you can greatly reduce the size of this list. A large exemption list has several unwanted effects on every computer that receives the GPO, including the following:

- Reduces the overall effectiveness of isolation.
- Creates a larger management burden (because of frequent updates).
- Increases the size of the IPsec policy, which means that it consumes more memory and CPU resources, slows down network throughput, and increases the time required to download and apply the GPO containing the IPsec policy.

To keep the number of exemptions as small as possible, you have several options:

- Carefully consider the communications requirements of each isolation zone, especially server-only zones. They might not be required to communicate with every exemption in the domain-level policy for clients.
- Consolidate server functions. If several exempt services can be hosted at one IP address, the number of exemptions is reduced.
- Consolidate exempted hosts on the same subnet. Where network traffic volume allows, you might be able to locate the servers on a subnet that is exempted, instead of using exemptions for each IP address.

As with defining the boundary zone, create a formal process to approve hosts being added to the exemption list. For a model of processing requests for exemptions, see the decision flowchart in the Boundary Zone section.

Next: Isolated Domain

Isolated Domain

The isolated domain is the primary zone for trusted computers. The computers in this zone use connection security and firewall rules to control the communications that can be sent between computers in the zone.

The term *domain* in this context means a boundary of communications trust instead of an Active Directory domain. In this solution the two constructs are very similar because Active Directory domain authentication (Kerberos V5) is required for accepting inbound connections from trusted computers. However, many Active Directory domains (or forests) can be linked with trust relationships to provide a single, logical, isolated domain. In addition, computers

that authenticate by using certificates can also be included in an isolated domain without joining the Active Directory domain.

For most implementations, an isolated domain will contain the largest number of computers. Other isolation zones can be created for the solution if their communication requirements differ from those of the isolated domain. Examples of these differences are what result in the boundary and encryption zones described in this guide. Conceptually, the isolated domain is just the largest isolation zone, and a superset to the other zones.

You must create a group in Active Directory to contain members of the isolated domain. You then apply one of several GPOs that contain connection security and firewall rules to the group so that authentication on all inbound network connections is enforced. Creation of the group and how to link the GPOs that apply the rules to its members are discussed in the **Planning Group Policy Deployment for Your Isolation Zones** section.

The GPOs for the isolated domain should contain the following connection security rules and settings.

GPO settings for isolated domain members running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008

GPOs for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 should include the following:

- IPsec default settings that specify the following options:
  a. Exempt all ICMP traffic from IPsec.
  b. Key exchange (main mode) security methods and algorithm. We recommend that you use at least DH4, AES and SHA2 in your settings. Use the strongest algorithm combinations that are common to all your supported operating systems.
  c. Data protection (quick mode) algorithm combinations. We recommend that you do not include DES, or MD5 in any setting. They are included only for compatibility with previous versions of Windows. Use the strongest algorithm combinations that are common to all your supported operating systems.

     If any NAT devices are present on your networks, use ESP encapsulation. If isolated domain members must communicate with hosts in the encryption zone, ensure that you include algorithms that are compatible with the requirements of the encryption mode policies.
  d. Authentication methods. Include at least computer-based Kerberos V5 authentication. If you want to use user-based access to isolated servers, then also include user-based Kerberos V5 as an optional authentication method. Likewise, if any of your isolated domain members cannot use Kerberos V5 authentication, then include certificate-based authentication as an optional authentication method.
- The following connection security rules:

- A connection security rule that exempts all computers on the exemption list from authentication. Be sure to include all your Active Directory domain controllers on this list. Enter subnet addresses, where possible, instead of discrete addresses, if applicable in your environment.
- A connection security rule, from any IP address to any, that requires inbound and requests outbound authentication by using Kerberos V5 authentication.

  Important
  Be sure to begin operations by using request in and request out behavior until you are sure that all the computers in your IPsec environment are communicating successfully by using IPsec. After confirming that IPsec is operating as expected, you can change the policy to require in, request out.

- A registry policy that includes the following values:
  a. Enable PMTU discovery. Enabling this setting allows TCP/IP to dynamically determine the largest packet size supported across a connection. The value is found at HKLM\System\CurrentControlSet\Services\TCPIP\Parameters\EnablePMTUDiscovery (dword). The sample GPO preferences XML file in Appendix A: Sample GPO Template Files for Settings Used in this Guide sets the value to **1**.

  Note
  For a sample template for these registry settings, see Appendix A: Sample GPO Template Files for Settings Used in this Guide.

Next: Boundary Zone

Boundary Zone

In most organizations, some computers must be able to receive network traffic from computers that are not part of the isolated domain, and therefore cannot authenticate. To accept communications from untrusted computers, create a boundary zone within your isolated domain.

Computers in the boundary zone are trusted computers that can accept communication requests both from other isolated domain member computers and from untrusted computers. Boundary zone computers try to authenticate any incoming request by using IPsec, initiating an IKE negotiation with the originating computer.

The GPOs you build for the boundary zone include IPsec or connection security rules that request authentication for both inbound and outbound network connections, but do not require it.

Because these boundary zone computers can receive unsolicited inbound communications from untrusted computers that use plaintext, they must be carefully managed and secured in other ways. Mitigating this additional risk is an important part of deciding whether to add a computer to the boundary zone. For example, completing a formal business justification process before adding each computer to the boundary zone can help ensure that the additional risk is minimized. The following illustration shows a sample process that can help make such a decision.

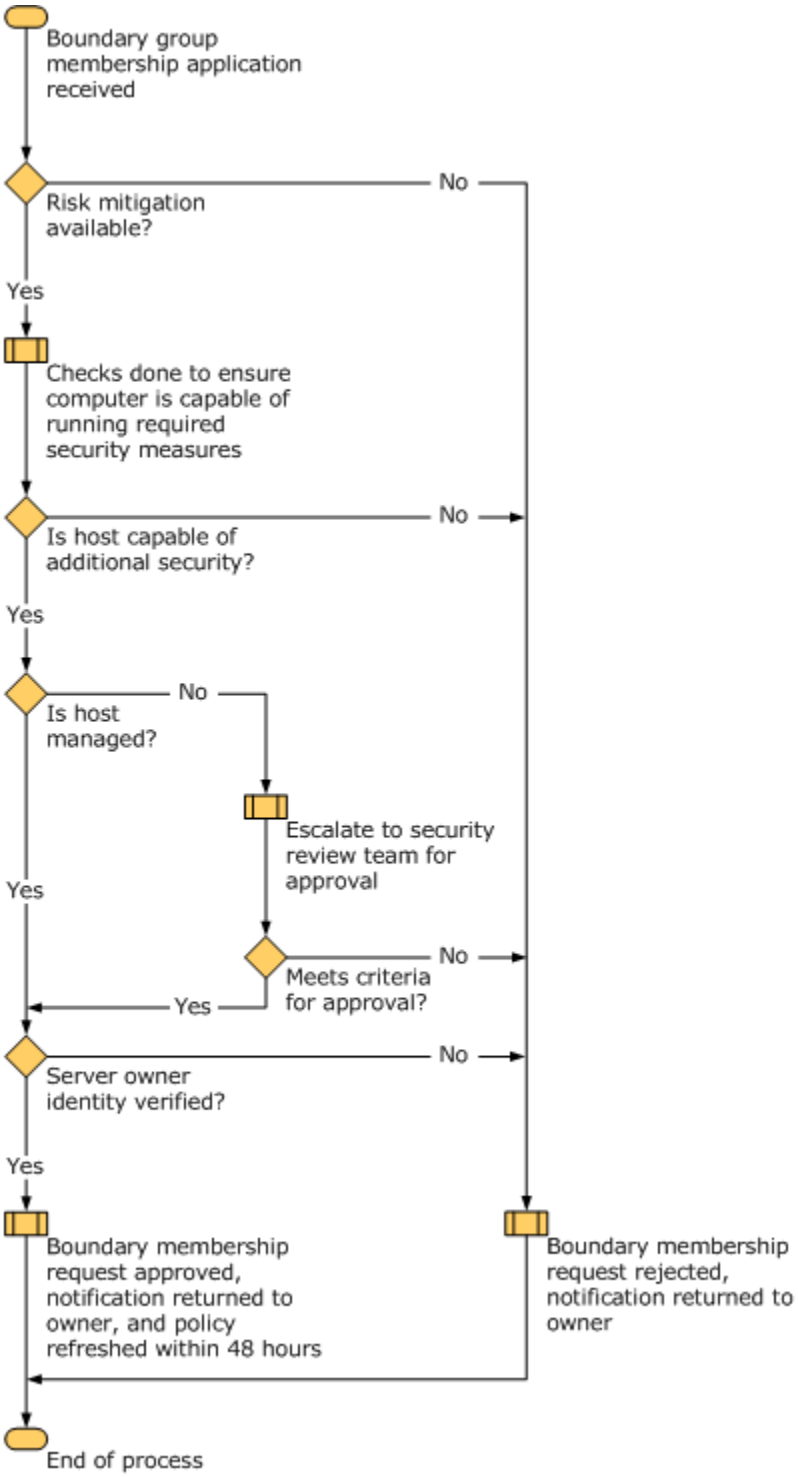

The goal of this process is to determine whether the risk of adding a computer to a boundary zone can be mitigated to a level that makes it acceptable to the organization. Ultimately, if the risk cannot be mitigated, membership must be denied.

You must create a group in Active Directory to contain the members of the boundary zones. The settings and rules for the boundary zone are typically very similar to those for the isolated domain, and you can save time and effort by copying those GPOs to serve as a starting point. The primary difference is that the authentication connection security rule must be set to request authentication for both inbound and outbound traffic, instead of requiring inbound authentication and requesting outbound authentication as used by the isolated domain.

Creation of the group and how to link it to the GPOs that apply the rules to members of the group are discussed in the Planning Group Policy Deployment for Your Isolation Zones section.

GPO settings for boundary zone servers running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

The boundary zone GPO for computers running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2 should include the following:

- IPsec default settings that specify the following options:
  a. Exempt all ICMP traffic from IPsec.
  b. Key exchange (main mode) security methods and algorithm. We recommend that you use at least DH4, AES and SHA2 in your settings. Use the strongest algorithm combinations that are common to all your supported operating systems.
  c. Data protection (quick mode) algorithm combinations. We recommend that you do not include DES or MD5 in any setting. They are included only for compatibility with previous versions of Windows. Use the strongest algorithm combinations that are common to all your supported operating systems..

     If any NAT devices are present on your networks, use ESP encapsulation. If isolated domain members must communicate with hosts in the encryption zone, ensure that you include algorithms that are compatible with the requirements of the encryption mode policies.
  d. Authentication methods. Include at least computer-based Kerberos V5 authentication. If you want to use user-based access to isolated servers then you must also include user-based Kerberos V5 authentication as an optional authentication method. Likewise, if any of your domain isolation members cannot use Kerberos V5, you must include certificate-based authentication as an optional authentication method.
- The following connection security rules:
  - A connection security rule that exempts all computers on the exemption list from authentication. Be sure to include all your Active Directory domain controllers on this list. Enter subnet addresses, if applicable in your environment.
  - A connection security rule, from **Any IP address** to **Any IP address**, that requests inbound and outbound authentication.
- A registry policy that includes the following values:

a. Enable PMTU discovery. Enabling this setting allows TCP/IP to dynamically determine the largest packet size supported across a connection. The value is found at HKLM\System\CurrentControlSet\Services\TCPIP\Parameters\EnablePMTUDiscovery (dword). The sample GPO preferences XML file in Appendix A: Sample GPO Template Files for Settings Used in this Guide sets the value to **1**.

Note

For a sample template for these registry settings, see Appendix A: Sample GPO Template Files for Settings Used in this Guide

Next: Encryption Zone

Encryption Zone

Some servers in the organization host data that is very sensitive, including medical, financial, or other personally identifying data. Government or industry regulations might require that this sensitive information must be encrypted when it is transferred between computers.

To support the additional security requirements of these servers, we recommend that you create an encryption zone to contain the computers and that requires that the sensitive inbound and outbound network traffic be encrypted.

You must create a group in Active Directory to contain members of the encryption zone. The settings and rules for the encryption zone are typically similar to those for the isolated domain, and you can save time and effort by copying those GPOs to serve as a starting point. You then modify the security methods list to include only algorithm combinations that include encryption protocols.

Creation of the group and how to link it to the GPOs that apply the rules to members of the group are discussed in the Planning Group Policy Deployment for Your Isolation Zones section.

GPO settings for encryption zone servers running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

The GPO for computers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 should include the following:

- IPsec default settings that specify the following options:
  a. Exempt all ICMP traffic from IPsec.
  b. Key exchange (main mode) security methods and algorithm. We recommend that you use at least DH4, AES and SHA2 in your settings. Use the strongest algorithm combinations that are common to all your supported operating systems.
  c. Data protection (quick mode) algorithm combinations. Check **Require encryption for all connection security rules that use these settings**, and then specify one or more integrity and encryption combinations. We recommend that you do not include DES or

MD5 in any setting. They are included only for compatibility with previous versions of Windows. Use the strongest algorithm combinations that are common to all your supported operating systems.

If any NAT devices are present on your networks, use ESP encapsulation..

d. Authentication methods. Include at least computer-based Kerberos V5 authentication. If you want to use user-based access to isolated servers then you must also include user-based Kerberos V5 authentication as an optional authentication method. Likewise, if any of your domain isolation members cannot use Kerberos V5 authentication, then you must include certificate-based authentication as an optional authentication method.

- The following connection security rules:
  - A connection security rule that exempts all computers on the exemption list from authentication. Be sure to include all your Active Directory domain controllers on this list. Enter subnet addresses, if applicable in your environment.
  - A connection security rule, from any IP address to any, that requires inbound and requests outbound authentication using the default authentication specified earlier in this policy.

    Important
    Be sure to begin operations by using request in and request out behavior until you are sure that all the computers in your IPsec environment are communicating successfully by using IPsec. After confirming that IPsec is operating as expected, you can change the GPO to require in, request out.

- A registry policy that includes the following values:

  a. Enable PMTU discovery. Enabling this setting allows TCP/IP to dynamically determine the largest packet size supported across a connection. The value is found at HKLM\System\CurrentControlSet\Services\TCPIP\Parameters\EnablePMTUDiscovery (dword). The sample GPO preferences XML file in Appendix A: Sample GPO Template Files for Settings Used in this Guide sets the value to **1**.

  Note
  For a sample template for these registry settings, see Appendix A: Sample GPO Template Files for Settings Used in this Guide.

- If domain member computers must communicate with computers in the encryption zone, ensure that you include in the isolated domain GPOs quick mode combinations that are compatible with the requirements of the encryption zone GPOs.

Next: Planning Server Isolation Zones

Planning Server Isolation Zones

Sometimes a server hosts data that is sensitive. If your servers host data that must not be compromised, you have several options to help protect that data. One was already addressed: adding the server to the encryption zone. Membership in that zone prevents the server from being

accessed by any computers that are outside the isolated domain, and encrypts all network connections to server.

The second option is to additionally restrict access to the server, not just to members of the isolated domain, but to only those users or computers who have business reasons to access the resources on the server. You can specify only approved users, or you can additionally specify that the approved users can only access the server from approved computers.

To grant access, you add the approved user and computer accounts to network access groups (NAGs) that are referenced in a firewall rule on this server. When the user sends a request to the server, the standard domain isolation rules are invoked. This causes IKE to use Kerberos V5 to exchange credentials with the server. The additional firewall rule on the server causes Windows to check the provided computer and user accounts for group membership in the NAGs. If either the user or computer is not a member of a required NAG then the network connection is refused.

Isolated domains and isolated servers

If you are using an isolated domain, the client computers already have the IPsec rules to enable them to authenticate traffic when the server requires it. If you add an isolated server, it must have a GPO applied to its group with the appropriate connection security and firewall rules. The rules enforce authentication and restrict access to only connections that are authenticated as coming from an authorized computer or user.

If you are not using an isolated domain, but still want to isolate a server that uses IPsec, you must configure the client computers that you want to access the server to use the appropriate IPsec rules. If the client computers are members of an Active Directory domain, you can still use Group Policy to configure the clients. Instead of applying the GPO to the whole domain, you apply the GPO to only members of the NAG.

Creating multiple isolated server zones

Each set of servers that must be accessed by different sets of users should be set up in its own isolated server zone. After one set of GPOs for one isolated server zone has been successfully created and verified, you can copy the GPOs to a new set. You must change the GPO names to reflect the new zone, the name and membership of the isolated server zone group to which the GPOs are applied, and the names and membership of the NAG groups that determine which clients can access the servers in the isolated server zone.

Creating the GPOs

Creation of the groups and how to link them to the GPOs that apply the rules to members of the groups are discussed in the Planning Group Policy Deployment for Your Isolation Zones section.

An isolated server is often a member of the encryption zone. Therefore, copying that GPO set serves as a good starting point. You then modify the rules to additionally restrict access to only NAG members.

GPO settings for isolated servers running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008

GPOs for computers running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 should include the following:

Note

The connection security rules described here are identical to the ones for the encryption zone. If you do not want to encrypt access and also restrict access to NAG members, you can use connection security rules identical to the main isolated domain. You must still add the firewall rule described at the end of this list to change it into an isolated server zone.

- IPsec default settings that specify the following options:
  a. Exempt all ICMP traffic from IPsec.
  b. Key exchange (main mode) security methods and algorithm. We recommend that you do not include Diffie-Hellman Group 1, DES, or MD5 in any setting. They are included only for compatibility with previous versions of Windows. Use the strongest algorithm combinations that are common to all your supported operating systems.
  c. Data protection (quick mode) algorithm combinations. Check **Require encryption for all connection security rules that use these settings**, and then specify one or more integrity and encryption combinations. We recommend that you do not include DES or MD5 in any setting. They are included only for compatibility with previous versions of Windows. Use the strongest algorithm combinations that are common to all your supported operating systems.

     If any NAT devices are present on your networks, do not use AH because it cannot traverse NAT devices. If isolated servers must communicate with hosts in the encryption zone, include an algorithm that is compatible with the requirements of the encryption zone GPOs.
  d. Authentication methods. Include at least computer-based Kerberos V5 authentication for compatibility with the rest of the isolated domain. If you want to restrict access to specific user accounts, also include user-based Kerberos V5 authentication as an optional authentication method. Do not make the user-based authentication method mandatory, or else computers that cannot use AuthIP instead of IKE, including Windows XP and Windows Server 2003, cannot communicate. Likewise, if any of your domain isolation members cannot use Kerberos V5, include certificate-based authentication as an optional authentication method.
- The following connection security and firewall rules:
  - A connection security rule that exempts all computers on the exemption list from authentication. Be sure to include all your Active Directory domain controllers on this list. Enter subnet addresses, if applicable in your environment.
  - A connection security rule, from **Any IP address** to **Any IP address**, that requires inbound and requests outbound authentication by using Kerberos V5 authentication.

    Important

Be sure to begin operations by using request in and request out behavior until you are sure that all the computers in your IPsec environment are communicating successfully by using IPsec. After confirming that IPsec is operating as expected, you can change the GPO to require in, request out.

- A firewall rule that specifies **Allow only secure connections**, **Require encryption**, and on the **Users and Computers** tab includes references to both computer and user network access groups.

- A registry policy that includes the following values:

  a. Enable PMTU discovery. Enabling this setting allows TCP/IP to dynamically determine the largest packet size supported across a connection. The value is found at HKLM\System\CurrentControlSet\Services\TCPIP\Parameters\EnablePMTUDiscovery (dword). The sample GPO preferences XML file in Appendix A: Sample GPO Template Files for Settings Used in this Guide sets the value to **1**.

Note

For a sample template for these registry settings, see Appendix A: Sample GPO Template Files for Settings Used in this Guide.

Next: Planning Certificate-based Authentication

Planning Certificate-based Authentication

Sometimes a computer cannot join an Active Directory domain, and therefore cannot use Kerberos V5 authentication with domain credentials. However, the computer can still participate in the isolated domain by using certificate-based authentication.

The non-domain member server, and the clients that must be able to communicate with it, must be configured to use cryptographic certificates based on the X.509 standard. These certificates can be used as an alternate set of credentials. During IKE negotiation, each computer sends a copy of its certificate to the other computer. Each computer examines the received certificate, and then validates its authenticity. To be considered authentic, the received certificate must be validated by a certification authority certificate in the recipient's Trusted Root Certification Authorities store on the local computer.

Certificates can be acquired from commercial firms, or by an internal certificate server set up as part of the organization's public key infrastructure (PKI). Microsoft provides a complete PKI and certification authority solution with Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 Active Directory Certificate Services (AD CS). For more information about creating and maintaining a PKI in your organization, see Active Directory Certificate Services Overview at http://technet.microsoft.com/library/hh831740.aspx.

Deploying certificates

No matter how you acquire your certificates, you must deploy them to clients and servers that require them in order to communicate.

Using Active Directory Certificate Services

If you use AD CS to create your own user and computer certificates in-house, then the servers designated as certification authorities (CAs) create the certificates based on administrator-designed templates. AD CS then uses Group Policy to deploy the certificates to domain member computers. Computer certificates are deployed when a domain member computer starts. User certificates are deployed when a user logs on.

If you want non-domain member computers to be part of a server isolation zone that requires access by only authorized users, make sure to include certificate mapping to associate the certificates with specific user accounts. When certificate mapping is enabled, the certificate issued to each computer or user includes enough identification information to enable IPsec to match the certificate to both user and computer accounts.

AD CS automatically ensures that certificates issued by the CAs are trusted by the client computers by putting the CA certificates in the correct store on each domain member computer.

Using a commercially purchased certificate for computers running Windows

You can import the certificates manually onto each computer if the number of computers is relatively small. For a deployment to more than a handful of computers, use Group Policy.

You must first download the vendor's root CA certificate, and then import it to a GPO that deploys it to the Local Computer\Trusted Root Certification Authorities store on each computer that applies the GPO.

You must also import the purchased certificate into a GPO that deploys it to the Local Computer\Personal store on each computer that applies the GPO.

Using a commercially purchased certificate for computers running a non-Windows operating system

If you are installing the certificates on an operating system other than Windows, see the documentation for that operating system.

Configuring IPsec to use the certificates

When the clients and servers have the certificates available, you can configure the IPsec and connection security rules to include those certificates as a valid authentication method. The authentication method requires the subject name of the certificate, for example: **DC=com,DC=woodgrovebank,CN=CorporateCertServer**. Optionally, select **Enable certificate to account mapping** to support using these credentials for restricting access to users or computers that are members of authorized groups in a server isolation solution.

Starting in Windows Server 2012, the Administrator can configure certificate selection criteria so the desired certificate is selected and/or validated. Enhanced Key Usage (EKU) criteria can be configured, as well as name restrictions and certificate thumbprints. This is configured using the

Advanced button when choosing certificates for the authentication method in the user interface, or through Windows PowerShell.

Next: Documenting the Zones

Documenting the Zones

Generally, the task of determining zone membership is not complex, but it can be time-consuming. Use the information generated during the Designing a Windows Firewall with Advanced Security Strategy section of this guide to determine the zone in which to put each host. You can document this zone placement by adding a Group column to the inventory table shown in the Designing a Windows Firewall with Advanced Security Strategy section. A sample is shown here:

Host name	Hardware reqs met	Software reqs met	Configuration required	Details	Projected cost	Group
CLIENT001	No	No	Upgrade hardware and software.	Current operating system is Windows XP. Old hardware not compatible with Windows 8.	$??	Isolated domain
SERVER002	Yes	No	Join trusted domain, upgrade from Windows Server 2008 to Windows Server 2012	No antivirus software present.	$??	Encryption
SENSITIVE001	Yes	Yes	Not required.	Running Windows Server 2012. Ready for inclusion.	$0	Isolated server (in zone by itself)
PRINTSVR1	Yes	Yes	Not required.	Running Windows Server 2008 R2. Ready for inclusion.	$0	Boundary

Next: Planning Group Policy Deployment for Your Isolation Zones

Planning Group Policy Deployment for Your Isolation Zones

After you have decided on the best logical design of your isolation environment for the network and computer security requirements, you can start the implementation plan.

You have a list of isolation zones with the security requirements of each. For implementation, you must plan the groups that will hold the computer accounts in each zone, the network access groups that will be used to determine who can access an isolated server, and the GPOs with the connection security and firewall rules to apply to corresponding groups. Finally you must determine how you will ensure that the policies will only apply to the correct computers within each group.

- Planning Isolation Groups for the Zones
- Planning Network Access Groups
- Planning the GPOs
- Planning GPO Deployment

Planning Isolation Groups for the Zones

Isolation groups in Active Directory are how you implement the various domain and server isolation zones. A computer is assigned to a zone by adding its computer account to the group which represents that zone.

Caution

Do not add computers to your groups yet. If a computer is in a group when the GPO is activated then that GPO is applied to the computer. If the GPO is one that requires authentication, and the other computers have not yet received their GPOs, the computer that uses the new GPO might not be able to communicate with the others.

Universal groups are the best option to use for GPO assignment because they apply to the whole forest and reduce the number of groups that must be managed. However, if universal groups are unavailable, you can use domain global groups instead.

The following table lists typical groups that can be used to manage the domain isolation zones discussed in the Woodgrove Bank example in this guide:

Group name	Description
CG_DOMISO_No_IPsec	A universal group of computer accounts that do not participate in the IPsec environment.

Group name	Description
	Typically consists of infrastructure computer accounts that will also be included in exemption lists. This group is used in security group filters to ensure that GPOs with IPsec rules are not applied to group members.
CG_DOMISO_IsolatedDomain	A universal group of computer accounts that contains the members of the isolated domain. During the early days of testing, this group might contain only a very small number of computers. During production, it might contain the built-in **Domain Computers** group to ensure that every computer in the domain participates. Members of this group receive the domain isolation GPO that requires authentication for inbound connections.
CG_DOMISO_Boundary	A universal group of computer accounts that contains the members of the boundary zone. Members of this group receive a GPO that specifies that authentication is requested, but not required.
CG_DOMISO_Encryption	A universal group of computer accounts that contains the members of the encryption zone. Members of this group receive a GPO that specifies that both authentication and encryption are required for all inbound connections.
CG_SRVISO_*ServerRole*	A universal group of computer accounts that contains the members of the server isolation group. Members of this group receive the server isolation GPO that requires membership in a network access group in order to connect. There will be one group for each set of servers that have different user and computer restriction requirements.

Multiple GPOs might be delivered to each group. Which one actually becomes applied depends on the security group filters assigned to the GPOs in addition to the results of any WMI filtering assigned to the GPOs. Details of the GPO layout are discussed in the section Planning the GPOs.

If multiple GPOs are assigned to a group, and similar rules are applied, the rule that most specifically matches the network traffic is the one that is used by the computer. For example, if one IPsec rule says to request authentication for all IP traffic, and a second rule from a different GPO says to require authentication for IP traffic to and from a specific IP address, then the second rule takes precedence because it is more specific.

Next: Planning Network Access Groups

Planning Network Access Groups

A network access group (NAG) is used to identify users and computers that have permission to access an isolated server. The server is configured with firewall rules that allow only network connections that are authenticated as originating from a computer, and optionally a user, whose accounts are members of its NAG. A member of the isolated domain can belong to as many NAGs as required.

Minimize the number of NAGs to limit the complexity of the solution. You need one NAG for each server isolation group to restrict the computers or users that are granted access. You can optionally split the NAG into two different groups: one for authorized computers and one for authorized users.

The NAGs that you create and populate become active by referencing them in the **Users and Computers** tab of the firewall rules in the GPO assigned to the isolated servers. The GPO must also contain connection security rules that require authentication to supply the credentials checked for NAG membership.

For the Woodgrove Bank scenario, access to the computers running SQL Server that support the WGBank application are restricted to the WGBank front-end servers and to approved administrative users logged on to specific authorized administrative computers. They are also only accessed by the approved admin users and the service account that is used to the run the WGBank front end service.

NAG Name	NAG Member Users, Computers, or Groups	Description
CG_NAG_*ServerRole*_Users	Svr1AdminA Svr1AdminB Group_AppUsers AppSvcAccount	This group is for all users who are authorized to make inbound IPsec connections to the isolated servers in this zone.

NAG Name	NAG Member Users, Computers, or Groups	Description
CG_NAG_*ServerRole*_Computers	Desktop1 Desktop2 AdminDT1 AppAdminDT1	This group contains all computers that are authorized to make inbound IPsec connections to the isolated servers in this zone.

Note
Membership in a NAG does not control the level of IPsec traffic protection. The IKE negotiation is only aware of whether the computer or user passed or failed the Kerberos V5 authentication process. The connection security rules in the applied GPO control the security methods that are used for protecting traffic and are independent of the identity being authenticated by Kerberos V5.

Next: Planning the GPOs

Planning the GPOs

When you plan the GPOs for your different isolation zones, you must complete the layout of the required zones and their mappings to the groups that link the computers to the zones.

General considerations

A few things to consider as you plan the GPOs:

- Do not allow a computer to be a member of more than one isolation zone. A computer in more than one zone receives multiple and possibly contradictory GPOs. This can result in unexpected, and difficult to troubleshoot behavior.

  The examples in this guide show GPOs that are designed to prevent the requirement to belong to multiple zones.
- Ensure that the IPsec algorithms you specify in your GPOs are compatible across all the versions of Windows. The same principle applies to the data integrity and encryption algorithms. We recommend that you include the more advanced algorithms when you have the option of selecting several in an ordered list. The computers will negotiate down from the top of their lists, selecting one that is configured on both computers. So a computer that is running Windows Vista that is connected to a server that is running Windows Server 2012 can communicate by using a much more secure algorithm.
- The primary difference in your domain isolation GPOs is whether the rules request or require authentication.

Caution

It is **critical** that you begin with all your GPOs set to request authentication instead of requiring it. Since the GPOs are delivered to the computers over time, applying a require policy to one computer breaks its ability to communicate with another computer that has not yet received its policy. Using request mode at the beginning enables computers to continue communicating by using plaintext connections if required. After you confirm that your computers are using IPsec where expected, you can schedule a conversion of the rules in the GPOs from requesting to requiring authentication, as required by each zone.

- Windows Firewall with Advanced Security in Windows Vista and Windows Server 2008 only support one network location profile at a time. If you add a second network adapter that is connected to a different network, or not connected at all, you could unintentionally change the profile that is currently active on the computer. If your GPO specifies different firewall and connection security rules based on the current network location profile, the behavior of how the computer handles network traffic will change accordingly. We recommend for stationary computers, such as desktops and servers, that you assign any rule for the computer to all profiles. Apply GPOs that change rules per network location to computers that must move between networks, such as your portable computers. Consider creating a separate domain isolation GPO for your servers that uses the same settings as the GPO for the clients, except that the server GPO specifies the same rules for all network location profiles. For more information, see Network Location Types at http://go.microsoft.com/fwlink/?linkid=110826.

  Note

  Computers running Windows 8, Windows 7, Windows Server 2012, and Windows Server 2008 R2 support different network location types, and therefore profiles, for each network adapter at the same time. Each network adapter is assigned the network location appropriate for the network to which it is connected. Windows Firewall then enforces only those rules that apply to that network type's profile. So certain types of traffic are blocked when coming from a network adapter connected to a public network, but those same types might be permitted when coming from a private or domain network.

After considering these issues, document each GPO that you require, and the details about the connection security and firewall rules that it needs.

Woodgrove Bank example GPOs

The Woodgrove Bank example uses the following set of GPOs to support its domain isolation requirements. This section only discusses the rules and settings for server and domain isolation. GPO settings that affect which computers receive the GPO, such as security group filtering and WMI filtering, are discussed in the Planning GPO Deployment section.

In this section you can find information about the following:

- Firewall GPOs
- Isolated Domain GPOs

- Boundary Zone GPOs
- Encryption Zone GPOs
- Server Isolation GPOs

Firewall GPOs

All the computers on Woodgrove Bank's network that run Windows are part of the isolated domain, except domain controllers. To configure firewall rules, the GPO described in this section is linked to the domain container in the Active Directory OU hierarchy, and then filtered by using security group filters and WMI filters.

The GPO created for the example Woodgrove Bank scenario include the following:

- GPO_DOMISO_Firewall

GPO_DOMISO_Firewall

This GPO is authored by using the Windows Firewall with Advanced Security interface in the Group Policy editing tools. The User Configuration section of the GPO is disabled. It is intended to only apply to computers that are running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2.

Firewall settings

This GPO provides the following settings:

- Unless otherwise stated, the firewall rules and settings described here are applied to all profiles.
- The firewall is enabled, with inbound, unsolicited connections blocked and outbound connections allowed.
- Under the domain profile, the settings **Display notifications to the user**, **Apply local firewall rules**, and **Apply local connection security rules** are all set to **No**. These settings are applied only to the domain profile because the computers can only receive an exception rule for a required program from a GPO if they are connected to the domain. Under the public and private profiles, those settings are all set to **Yes**.

  Note

  Enforcing these settings requires that you define any firewall exceptions for programs, because the user cannot manually permit a new program. You must deploy the exception rules by adding them to this GPO. We recommend that you do not enable these settings until you have tested all your applications and have tested the resulting rules in a test lab and then on pilot computers.

Firewall rules

This GPO provides the following rules:

- Built-in firewall rule groups are configured to support typically required network operation. The following rule groups are set to **Allow the connection**:
  - Core Networking
  - File and Printer Sharing
  - Network Discovery
  - Remote Administration
  - Remote Desktop
  - Remote Event Log Management
  - Remote Scheduled Tasks Management
  - Remote Service Management
  - Remote Volume Management
  - Windows Firewall Remote Management
  - Windows Management Instrumentation (WMI)
  - Windows Remote Management
- A firewall exception rule to allow required network traffic for the WGBank dashboard program. This inbound rule allows network traffic for the program Dashboard.exe in the %ProgramFiles%\WGBank folder. The rule is also filtered to only allow traffic on port 1551. This rule is applied only to the domain profile.

Next: Isolated Domain GPOs

Isolated Domain GPOs

All of the computers in the isolated domain are added to the group CG_DOMISO_IsolatedDomain. You must create multiple GPOs to align with this group, one for each Windows operating system that must have different rules or settings to implement the basic isolated domain functionality that you have in your isolated domain. This group is granted Read and Apply Group Policy permissions on all the GPOs described in this section.

Each GPO has a security group filter that prevents the GPO from applying to members of the group GP_DOMISO_No_IPsec. A WMI filter is attached to each GPO to ensure that the GPO is applied to only the specified version of Windows. For more information, see the Planning GPO Deployment section.

The GPOs created for the Woodgrove Bank isolated domain include the following:

- GPO_DOMISO_IsolatedDomain_Clients
- GPO_DOMISO_IsolatedDomain_Servers

GPO_DOMISO_IsolatedDomain_Clients

This GPO is authored by using the Windows Firewall with Advanced Security interface in the Group Policy editing tools. The User Configuration section of the GPO is disabled. It is intended to only apply to client computers that are running Windows 8, Windows 7, or Windows Vista.

Because client computers can sometimes be portable, the settings and rules for this GPO are applied to only the domain profile.

General settings

This GPO provides the following settings:

- No firewall settings are included in this GPO. Woodgrove Bank created separate GPOs for firewall settings (see the Firewall GPOs section) in order to share them with all clients in all isolation zones with minimum redundancy.
- The ICMP protocol is exempted from authentication requirements to support easier network troubleshooting.
- Diffie-Hellman Group 2 is specified as the key exchange algorithm. This is the strongest algorithm available that is supported by all the operating systems that are being used at Woodgrove Bank. After Woodgrove Bank has completed the upgrade to versions of Windows that support stronger algorithms, such as Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2, they can remove the weaker key exchange algorithms, and use only the stronger ones.
- The registry settings shown in the following table. For more information, see the description of the registry settings in Isolated Domain.

Setting	Value
Enable PMTU Discovery	1
IPsec Exemptions	3

- The main mode security method combinations in the order shown in the following table.

Integrity	Encryption
Secure Hash Algorithm (SHA-1)	Advanced Encryption Standard (AES-128)
SHA-1	3DES

- The following quick mode security data integrity algorithms combinations in the order shown in the following table.

Protocol	Integrity	Key Lifetime (minutes/KB)
ESP	SHA-1	60/100,000

- The quick mode security data integrity and encryption algorithm combinations in the order shown in the following table.

Protocol	Integrity	Encryption	Key Lifetime (minutes/KB)
ESP	SHA-1	AES-128	60/100,000
ESP	SHA-1	3DES	60/100,000

Note
Do not use the MD5 and DES algorithms in your GPOs. They are included only for compatibility with previous versions of Windows.

Connection Security Rules

This GPO provides the following rules:

- A connection security rule named **Isolated Domain Rule** with the following settings:
  - From **Any IP address** to **Any IP address**.
  - **Require inbound and request outbound** authentication requirements.

    Important
    On this, and all other GPOs that require authentication, Woodgrove Bank first chose to only request authentication. After confirming that the computers were successfully communicating by using IPsec, they switched the GPOs to require authentication.
  - For **First authentication methods**, select **Computer Kerberos v5** as the primary method. Add certificate-based authentication from **DC=com,DC=woodgrovebank,CN=CorporateCertServer** for computers that cannot run Windows or cannot join the domain, but must still participate in the isolated domain.
  - For **Second authentication**, select **User Kerberos v5**, and then select the **Second authentication is optional** check box.
- A connection security rule to exempt computers that are in the exemption list from the requirement to authenticate:
  - The IP addresses of all computers on the exemption list must be added individually under **Endpoint 2**.
  - Authentication mode is set to **Do not authenticate**.

Next: GPO_DOMISO_IsolatedDomain_Servers

GPO_DOMISO_IsolatedDomain_Servers

This GPO is authored by using the Windows Firewall with Advanced Security interface in the Group Policy editing tools. The User Configuration section of the GPO is disabled. It is intended to only apply to server computers that are running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2.

Because so many of the settings and rules for this GPO are common to those in the GPO for Windows 8, Windows 7 and Windows Vista, you can save time by exporting the Windows Firewall with Advanced Security piece of the GPO for Windows 8, Windows 7 and Windows Vista, and importing it to the GPO for Windows Server 2012, Windows Server 2008 and Windows Server 2008 R2. After the import, change only the items specified here:

- This GPO applies all its settings to all profiles: Domain, Private, and Public. Because a server is not expected to be mobile and changing networks, configuring the GPO in this way prevents a network failure or the addition of a new network adapter from unintentionally switching the computer to the Public profile with a different set of rules (in the case of a server running Windows Server 2008).

  Important
  Windows Vista and Windows Server 2008 support only one network location profile at a time. The profile for the least secure network type is applied to the computer. If you attach a network adapter to a computer that is not physically connected to a network, the public network location type is associated with the network adapter and applied to the computer.

Next: Boundary Zone GPOs

Boundary Zone GPOs

All the computers in the boundary zone are added to the group CG_DOMISO_Boundary. You must create multiple GPOs to align with this group, one for each operating system that you have in your boundary zone. This group is granted Read and Apply permissions in Group Policy on the GPOs described in this section.

Note
If you are designing GPOs for only Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2, you can design your GPOs in nested groups. For example, you can make the boundary group a member of the isolated domain group, so that it receives the firewall and basic isolated domain settings through that nested membership, with only the changes supplied by the boundary zone GPO. For simplicity, this guide describes the techniques used to create the independent, non-layered policies. We recommend that you create and periodically run a script that compares the memberships of the groups that must be mutually

exclusive and reports any computers that are incorrectly assigned to more than one group.

This means that you create a GPO for a boundary group for a specific operating system by copying and pasting the corresponding GPO for the isolated domain, and then modifying the new copy to provide the behavior required in the boundary zone.

The boundary zone GPOs discussed in this guide are only for server versions of Windows because client computers are not expected to participate in the boundary zone. If the need for one occurs, either create a new GPO for that version of Windows, or expand the WMI filter attached to one of the existing boundary zone GPOs to make it apply to the client version of Windows.

In the Woodgrove Bank example, only the GPO settings for a Web service on Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 are discussed.

- GPO_DOMISO_Boundary_WS2008

GPO_DOMISO_Boundary_WS2008

This GPO is authored by using the Windows Firewall with Advanced Security interface in the Group Policy editing tools. Woodgrove Bank began by copying and pasting the GPO for the Windows Server 2008 version of the isolated domain GPO, and then renamed the copy to reflect its new purpose.

This GPO supports the ability for computers that are not part of the isolated domain to access specific servers that must be available to those untrusted computers. It is intended to only apply to server computers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008.

IPsec settings

The copied GPO includes and continues to use the IPsec settings that configure key exchange, main mode, and quick mode algorithms for the isolated domain when authentication can be used.

Connection security rules

Rename the **Isolated Domain Rule** to **Boundary Zone Rule**. Change the authentication mode to **Request inbound and request outbound**. In this mode, the computer uses authentication when it can, such as during communication with a member of the isolated domain. It also supports the "fall back to clear" ability of request mode when an untrusted computer that is not part of the isolated domain connects.

Registry settings

The boundary zone uses the same registry settings as the isolated domain to optimize IPsec operation. For more information, see the description of the registry settings in Isolated Domain.

Firewall rules

Copy the firewall rules for the boundary zone from the GPO that contains the firewall rules for the isolated domain. Customize this copy, removing rules for services not needed on servers in this zone, and adding inbound rules to allow the network traffic for the services that are to be accessed by other computers. For example, Woodgrove Bank added a firewall rule to allow inbound network traffic to TCP port 80 for Web client requests.

Make sure that the GPO that contains firewall rules for the isolated domain does not also apply to the boundary zone to prevent overlapping, and possibly conflicting rules.

Next: Encryption Zone GPOs

Encryption Zone GPOs

Handle encryption zones in a similar manner to the boundary zones. A computer is added to an encryption zone by adding the computer account to the encryption zone group. Woodgrove Bank has a single service that must be protected, and the computers that are running that service are added to the group CG_DOMISO_Encryption. This group is granted Read and Apply Group Policy permissions in on the GPO described in this section.

The GPO is only for server versions of Windows. Client computers are not expected to participate in the encryption zone. If the need for one occurs, either create a new GPO for that version of Windows, or expand the WMI filter attached to one of the existing encryption zone GPOs to make it apply to the client version of Windows.

- GPO_DOMISO_Encryption_WS2008

GPO_DOMISO_Encryption_WS2008

This GPO is authored by using the Windows Firewall with Advanced Security interface in the Group Policy editing tools. Woodgrove Bank began by copying and pasting the GPO for the Windows Server 2008 version of the isolated domain GPO, and then renamed the copy to reflect its new purpose.

This GPO supports the ability for servers that contain sensitive data to require encryption for all connection requests. It is intended to only apply to server computers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008.

IPsec settings

The copied GPO includes and continues to use the IPsec settings that configure key exchange, main mode, and quick mode algorithms for the isolated domain The following changes are made to encryption zone copy of the GPO:

The encryption zone servers require all connections to be encrypted. To do this, change the IPsec default settings for the GPO to enable the setting **Require encryption for all connection security rules that use these settings**. This disables all integrity-only algorithm combinations.

Connection security rules

Rename the **Isolated Domain Rule** to **Encryption Zone Rule**. Leave the authentication mode setting on **Require inbound and request outbound**. In this mode, the computer forces authentication for all inbound network traffic, and uses it when it can on outbound traffic.

Registry settings

The encryption zone uses the same registry settings as the isolated domain to optimize IPsec operation. For more information, see the description of the registry settings in Isolated Domain.

Firewall rules

Copy the firewall rules for the encryption zone from the GPO that contains the firewall rules for the isolated domain. Customize this copy, removing rules for services not needed on servers in this zone, and adding inbound rules to allow the network traffic for the services that are to be accessed by other computers. For example, Woodgrove Bank added a firewall rule to allow inbound network traffic to TCP port 1433 for SQL Server client requests.

Change the action for every inbound firewall rule from **Allow the connection** to **Allow only secure connections**, and then select **Require the connections to be encrypted**.

Make sure that the GPO that contains firewall rules for the isolated domain does not also apply to the boundary zone to prevent overlapping, and possibly conflicting rules.

Next: Server Isolation GPOs

Server Isolation GPOs

Each set of computers that have different users or computers accessing them require a separate server isolation zone. Each zone requires one GPO for each version of Windows running on computers in the zone. The Woodgrove Bank example has an isolation zone for their computers that run SQL Server. The server isolation zone is logically considered part of the encryption zone. Therefore, server isolation zone GPOs must also include rules for encrypting all isolated server traffic. Woodgrove Bank copied the encryption zone GPOs to serve as a starting point, and renamed them to reflect their new purpose.

All of the computer accounts for computers in the SQL Server server isolation zone are added to the group CG_SRVISO_WGBANK_SQL. This group is granted Read and Apply Group Policy permissions in on the GPOs described in this section. The GPOs are only for server versions of Windows. Client computers are not expected to be members of the server isolation zone, although they can access the servers in the zone by being a member of a network access group (NAG) for the zone.

GPO_SRVISO_WS2008

This GPO is identical to the GPO_DOMISO_Encryption_WS2008 GPO with the following changes:

- The firewall rule that enforces encryption is modified to include the NAGs on the **Users and Computers** tab of the rule. The NAGs granted permission include CG_NAG_SQL_Users and CG_NAG_SQL_Computers.

  Important

  Earlier versions of Windows support only computer-based authentication. If you specify that user authentication is mandatory, only users on computers that are running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 can connect.

Next: Planning GPO Deployment

Planning GPO Deployment

You can control which GPOs are applied to computers in Active Directory in a combination of three ways:

- **Active Directory organizational unit hierarchy**. This involves linking the GPO to a specific OU in the Active Directory OU hierarchy. All computers in the OU and its subordinate containers receive and apply the GPO.

  Controlling GPO application through linking to OUs is typically used when you can organize the OU hierarchy according to your domain isolation zone requirements. GPOs can apply settings to computers based on their location within Active Directory. If a computer is moved from one OU to another, the policy linked to the second OU will eventually take effect when Group Policy detects the change during polling.
- **Security group filtering**. This involves linking the GPOs to the domain level (or other parent OU) in the OU hierarchy, and then selecting which computers receive the GPO by using permissions that only allow correct group members to apply the GPO.

  The security group filters are attached to the GPOs themselves. A group is added to the security group filter of the GPO in Active Directory, and then assigned Read and Apply Group Policy permissions. Other groups can be explicitly denied Read and Apply Group Policy permissions. Only those computers whose group membership are granted Read and Apply Group Policy permissions without any explicit deny permissions can apply the GPO.

- **WMI filtering**. A WMI filter is a query that is run dynamically when the GPO is evaluated. If a computer is a member of the result set when the WMI filter query runs, the GPO is applied to the computer.

  A WMI filter consists of one or more conditions that are evaluated against the local computer. You can check almost any characteristic of the computer, its operating system, and its installed programs. If all of the specified conditions are true for the computer, the GPO is applied; otherwise the GPO is ignored.

This guide uses a combination of security group filtering and WMI filtering to provide the most flexible options. If you follow this guidance, even though there might be five different GPOs linked to a specific group because of operating system version differences, only the correct GPO is applied.

General considerations

- Deploy your GPOs before you add any computer accounts to the groups that receive the GPOs. That way you can add your computers to the groups in a controlled manner. Be sure to add only a few test computers at first. Before adding many group members, examine the results on the test computers and verify that the configured firewall and connection security rules have the effect that you want. See the following sections for some suggestions on what to test before you continue.

Test your deployed groups and GPOs

After you have deployed your GPOs and added some test computers to the groups, confirm the following before you continue with more group members:

- Examine the GPOs that are both assigned to and filtered from the computer. Run the **gpresult** tool at a command prompt.
- Examine the rules deployed to the computer. Open the Windows Firewall with Advanced Security MMC snap-in, expand the **Monitoring** node, and then expand the **Firewall** and **Connection Security** nodes.
- Verify that communications are authenticated. Open the Windows Firewall with Advanced Security MMC snap-in, expand the **Monitoring** node, expand the **Security Associations** node, and then click **Main Mode**.
- Verify that communications are encrypted when the computers require it. Open the Windows Firewall with Advanced Security MMC snap-in, expand the **Monitoring** node, expand the **Security Associations** node, and then select **Quick Mode**. Encrypted connections display a value other than **None** in the **ESP Confidentiality** column.
- Verify that your programs are unaffected. Run them and confirm that they still work as expected.

After you have confirmed that the GPOs have been correctly applied, and that the computers are now communicating by using IPsec network traffic in request mode, you can begin to add more computers to the group accounts, in manageable numbers at a time. Continue to monitor and confirm the correct application of the GPOs to the computers.

Do not enable require mode until deployment is complete

If you deploy a GPO that requires authentication to a computer before the other computers have a GPO deployed, communication between them might not be possible. Wait until you have all the zones and their GPOs deployed in request mode and confirm (as described in the previous section) that the computers are successfully communicating by using IPsec.

If there are problems with GPO deployment, or errors in configuration of one or more of the IPsec GPOs, computers can continue to operate, because request mode enables any computer to fall back to clear communications.

Only after you have added all of the computers to their zones, and you have confirmed that communications are working as expected, you can start changing the request mode rules to require mode rules where it is required in the zones. We recommend that you enable require mode in the zones one zone at a time, pausing to confirm that they are functioning properly before you continue. Turn the required mode setting on for the server isolation zones first, then the encryption zone, and then the isolated domain.

Do not change the boundary zone GPO, because it must stay in request mode for both inbound and outbound connections.

If you create other zones that require either inbound or outbound require mode, make the setting change in a manner that applies the setting in stages from the smaller groups of computers to the larger groups.

Example Woodgrove Bank deployment plans

Woodgrove Bank links all its GPOs to the domain level container in the Active Directory OU hierarchy. It then uses the following WMI filters and security group filters to control the application of the GPOs to the correct subset of computers. All of the GPOs have the User Configuration section disabled to improve performance.

GPO_DOMISO_Firewall_2008_Win7-Vista

- **WMI filter**. The WMI filter allows this GPO to apply only to computers that match the following WMI query:

  ```
  select * from Win32_OperatingSystem where Version like "6.%" and ProductType <> "2"
  ```

  Note
  This excludes domain controllers (which report a ProductType value of 2). Do not include domain controllers in the isolated domain if there are computers running versions of Windows earlier than Windows Vista and Windows Server 2008.

- **Security filter**. This GPO grants Read and Apply Group Policy permissions only to computers that are members of the group CG_DOMISO_IsolatedDomain. The GPO also explicitly denies Read and Apply Group Policy permissions to members of the CG_DOMISO_NO_IPSEC.

GPO_DOMISO_IsolatedDomain_Clients_Win7Vista

- **WMI filter**. The WMI filter allows this GPO to apply only to computers that match the following WMI query:

```
select * from Win32_OperatingSystem where Version like "6.%" and ProductType = "1"
```

- **Security filter**. This GPO grants Read and Apply Group Policy permissions only to computers that are members of the group CG_DOMISO_IsolatedDomain. The GPO also explicitly denies Read and Apply Group Policy permissions to members of the group CG_DOMISO_NO_IPSEC.

GPO_DOMISO_IsolatedDomain_Servers_WS2008

- **WMI filter**. The WMI filter allows this GPO to apply only to computers that match the following WMI query:

```
select * from Win32_OperatingSystem where Version like "6.%" and ProductType = "3"
```

  Note

  This excludes domain controllers (which report a ProductType value of 2). Do not include domain controllers in the isolated domain if there are computers that are running versions of Windows earlier than Windows Vista and Windows Server 2008.

- **Security filter**. This GPO grants Read and Apply Group Policy permissions only to computers that are members of the group CG_DOMISO_IsolatedDomain. The GPO also explicitly denies Read and Apply Group Policy permissions to members of the group CG_DOMISO_NO_IPSEC.

GPO_DOMISO_Boundary_WS2008

- **WMI filter**. The WMI filter allows this GPO to apply only to computers that match the following WMI query:

```
select * from Win32_OperatingSystem where Version like "6.%" and ProductType = "3"
```

  Note

  This excludes domain controllers (which report a ProductType value of 2). Do not include domain controllers in the isolated domain if there are computers that are running versions of Windows earlier than Windows Vista and Windows Server 2008.

- **Security filter**. This GPO grants Read and Apply Group Policy permissions only to computers that are members of the group CG_DOMISO_Boundary. The GPO also explicitly denies Read and Apply Group Policy permissions to members of the group CG_DOMISO_NO_IPSEC.

GPO_DOMISO_Encryption_WS2008

- **WMI filter**. The WMI filter allows this GPO to apply only to computers that match the following WMI query:

```
select * from Win32_OperatingSystem where Version like "6.%" and ProductType = "3"
```

Note
This excludes domain controllers (which report a ProductType value of 2). Do not include domain controllers in the isolated domain if there are computers that are running versions of Windows earlier than Windows Vista and Windows Server 2008.

- **Security filter**. This GPO grants Read and Apply permissions in Group Policy only to computers that are members of the group CG_DOMISO_Encryption. The GPO also explicitly denies Read and Apply permissions in Group Policy to members of the group CG_DOMISO_NO_IPSEC.

Appendix A: Sample GPO Template Files for Settings Used in this Guide

You can import an XML file containing customized registry preferences into a Group Policy Object (GPO) by using the Preferences feature of the Group Policy Management Console (GPMC). Creating registry setting preferences as described here was first implemented in Windows Server 2008 and Windows Vista with Service Pack 1 (SP1).

To manually create the file, build the settings under **Computer Configuration**, **Preferences**, **Windows Settings**, **Registry**. After you have created the settings, drag the container to the desktop. An .xml file is created there.

To import an .xml file to GPMC, drag it and drop it on the **Registry** node under **Computer Configuration**, **Preferences**, **Windows Settings**. If you copy the following sample XML code to a file, and then drag and drop it on the **Registry** node, it creates a **Server and Domain Isolation** collection with the six registry keys discussed in this guide.

The following sample file uses item-level targeting to ensure that the registry keys are applied only on the versions of Windows to which they apply.

Note
The file shown here is for sample use only. It should be customized to meet the requirements of your organization's deployment. To customize this file, import it into a test GPO, modify the settings, and then drag the Server and Domain Isolation Settings node to your desktop. The new file will contain all of your customization.

```
<?xml version="1.0" encoding="utf-8"?>


<Collection clsid="{53B533F5-224C-47e3-B01B-CA3B3F3FF4BF}" name="Server and Domain
Isolation Settings">


<Registry
         clsid="{9CD4B2F4-923D-47f5-A062-E897DD1DAD50}"
         name="Enable PMTU Discovery"
```

```
         status="EnablePMTUDiscovery"
         image="12"
         changed="2008-05-30 20:37:37"
         uid="{52C38FD7-A081-404C-A8EA-B24A9614D0B5}"
         desc="&lt;b&gt;Enable PMTU Discovery&lt;/b&gt;&lt;p&gt;
            This setting configures whether computers can use PMTU
            discovery on the network.&lt;p&gt;
            &lt;b&gt;1&lt;/b&gt; --  Enable&lt;br&gt;
            &lt;b&gt;0&lt;/b&gt; --  Disable"
         bypassErrors="1">
   <Properties
         action="U"
         displayDecimal="1"
         default="0"
         hive="HKEY_LOCAL_MACHINE"
         key="System\CurrentControlSet\Services\TCPIP\Parameters"
         name="EnablePMTUDiscovery" type="REG_DWORD" value="00000001"/>
</Registry>

<Registry
         clsid="{9CD4B2F4-923D-47f5-A062-E897DD1DAD50}"
         name="IPsec Default Exemptions (Vista and W2K8)"
         status="NoDefaultExempt"
         image="12"
         changed="2008-05-30 20:33:32"
         uid="{AE5C505D-283E-4060-9A55-70659DFD56B6}"
         desc="&lt;b&gt;IPsec Default Exemptions for Windows Server 2008
            and later&lt;/b&gt;&lt;p&gt;
            This setting determines which network traffic type is exempt
            from any IPsec authentication requirements.&lt;p&gt;
            &lt;b&gt;0&lt;/b&gt;: Exempts multicast, broadcast, RSVP, Kerberos,
ISAKMP&lt;br&gt;
            &lt;b&gt;1&lt;/b&gt;: Exempts multicast, broadcast, ISAKMP&lt;br&gt;
            &lt;b&gt;2&lt;/b&gt;: Exempts RSVP, Kerberos, ISAKMP&lt;br&gt;
```

```
            &lt;b&gt;3&lt;/b&gt;: Exempts ISAKMP only"
         bypassErrors="1">
   <Properties
         action="U"
         displayDecimal="1"
         default="0"
         hive="HKEY_LOCAL_MACHINE"
         key="SYSTEM\CurrentControlSet\Services\PolicyAgent"
         name="NoDefaultExempt"
         type="REG_DWORD"
         value="00000003"/>
   <Filters>
      <FilterOs
         bool="AND" not="0"
         class="NT" version="VISTA"
         type="NE" edition="NE" sp="NE"/>
      <FilterOs
         bool="OR" not="0"
         class="NT" version="2K8"
         type="NE" edition="NE" sp="NE"/>
   </Filters>
</Registry>


</Collection>
```

Additional Resources

For more information about the technologies discussed in this guide, see topics referenced in the following sections.

Windows Firewall with Advanced Security

- Windows Firewall with Advanced Security Overview (http://technet.microsoft.com/library/hh831365)

This TechNet page contains links to a variety of documents available for Windows Firewall with Advanced Security.

IPsec

- IPsec (http://technet.microsoft.com/network/bb531150.aspx)

  This TechNet page contains links to a variety of documents currently available for Internet Protocol security (IPsec) for Windows available as connection security rules.

Server and Domain Isolation

- Server and Domain Isolation (http://technet.microsoft.com/network/bb545651.aspx)

  This TechNet page contains links to documentation about the most common uses for IPsec: server isolation and domain isolation.

Group Policy

Group Policy is a key method for implementing firewall and server and domain isolation designs.

For more information about Group Policy and related technologies, see:

- **Group Policy** Group Policy Overview (http://technet.microsoft.com/library/hh831791)

  This page contains links to the documents currently available for Group Policy.
- WMI Filtering Using GPMC (http://technet.microsoft.com/library/6237b9b2-4a21-425e-8976-2065d28b3147)
- HOWTO: Leverage Group Policies with WMI Filters (http://support.microsoft.com/kb/555253)

  This article describes how to create a WMI filter to set the scope of a GPO based on computer attributes, such as operating system.

Active Directory Domain Services

Organizations can use AD DS to manage users and resources, such as computers, printers, or applications, on a network. Server isolation and domain isolation also require AD DS to use the Kerberos V5 protocol for IPsec authentication.

For more information about AD DS and related technologies, see:

- Active Directory Domain Services Overview (http://technet.microsoft.com/library/hh831484)

Windows Firewall with Advanced Security Deployment Guide

You can use the Windows Firewall with Advanced Security MMC snap-in in Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 to help protect the computers and the data that they share across a network.

You can use Windows Firewall to control access to the computer from the network. You can create rules that allow or block network traffic in either direction based on your business requirements. You can also create IPsec connection security rules to help protect your data as it travels across the network from computer to computer.

About this guide

This guide is intended for use by system administrators and system engineers. It provides detailed guidance for deploying a Windows Firewall with Advanced Security design that you or an infrastructure specialist or system architect in your organization has selected.

Begin by reviewing the information in Planning to Deploy Windows Firewall with Advanced Security.

If you have not yet selected a design, we recommend that you wait to follow the instructions in this guide until after you have reviewed the design options in the Windows Firewall with Advanced Security Design Guide and selected the one most appropriate for your organization.

After you select your design and gather the required information about the zones (isolation, boundary, and encryption), operating systems to support, and other details, you can then use this guide to deploy your Windows Firewall with Advanced Security design in your production environment. This guide provides steps for deploying any of the following primary designs that are described in the Design Guide:

- Basic Firewall Policy Design
- Domain Isolation Policy Design
- Server Isolation Policy Design
- Certificate-based Isolation Policy Design

Use the checklists in Implementing Your Windows Firewall with Advanced Security Design Plan to determine how best to use the instructions in this guide to deploy your particular design.

Caution

We recommend that you use the techniques documented in this guide only for GPOs that must be deployed to the majority of the computers in your organization, and only when the OU hierarchy in your Active Directory domain does not match the deployment needs of these GPOs. These characteristics are typical of GPOs for server and domain isolation scenarios, but are not typical of most other GPOs. When the OU hierarchy supports it, deploy a GPO by linking it to the lowest level OU that contains all of the accounts to which the GPO applies.

In a large enterprise environment with hundreds or thousands of GPOs, using this technique with too many GPOs can result in user or computer accounts that are members of an excessive number of groups; this can result in network connectivity problems if network protocol limits are exceeded. For more information about the problems associated with excessive group membership, see the following articles in the Microsoft Knowledge Base:

- Article 327825, “New resolution for problems with Kerberos authentication when users belong to many groups” (http://go.microsoft.com/fwlink/?linkid=128367)
- Article 263693 “Group Policy may not be applied to users belonging to many groups” (http://go.microsoft.com/fwlink/?linkid=128368)
- Article 328889 “Users who are members of more than 1,015 groups may fail logon authentication” (http://go.microsoft.com/fwlink/?linkid=128369)

What this guide does not provide

This guide does not provide:

- Guidance for creating firewall rules for specific network applications. For this information, see Planning Settings for a Basic Firewall Policy in the Windows Firewall with Advanced Security Design Guide.
- Guidance for setting up Active Directory Domain Services (AD DS) to support Group Policy. For more information, see Active Directory Domain Services (http://go.microsoft.com/fwlink/?linkid=102573) and Group Policy (http://go.microsoft.com/fwlink/?linkid=93542).
- Guidance for setting up certification authorities (CAs) to create certificates for certificate-based authentication. For this information, see Active Directory Certificate Services (http://go.microsoft.com/fwlink/?linkid=110820).

Overview of Windows Firewall with Advanced Security

Windows Firewall with Advanced Security in Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 is a stateful host firewall that helps secure the computer by allowing you to create rules that determine which network traffic is permitted to enter the computer from the network and which network traffic the computer is allowed to send to the network. Windows Firewall with Advanced Security also supports Internet Protocol security (IPsec), which you can use to require authentication from any computer that is attempting to communicate with your computer. When authentication is required, computers that cannot be authenticated as a trusted computer cannot communicate with your computer. You can also use IPsec to require that certain network traffic is encrypted to prevent it from being read by network packet analyzers that could be attached to the network by a malicious user.

The Windows Firewall with Advanced Security MMC snap-in is more flexible and provides much more functionality than the consumer-friendly Windows Firewall interface found in the Control

Panel. Both interfaces interact with the same underlying services, but provide different levels of control over those services. While the Windows Firewall Control Panel program can protect a single computer in a home environment, it does not provide enough centralized management or security features to help secure more complex network traffic found in a typical business enterprise environment.

For more information about Windows Firewall with Advanced Security, see Windows Firewall with Advanced Security Overview at http://technet.microsoft.com/library/hh831365.aspx.

Planning to Deploy Windows Firewall with Advanced Security

After you collect information about your environment and decide on a design by following the guidance in the Windows Firewall with Advanced Security Design Guide, you can begin to plan the deployment of your design. With the completed design and the information in this topic, you can determine which tasks to perform to deploy Windows Firewall with Advanced Security in your organization.

Reviewing your Windows Firewall with Advanced Security Design

If the design team that created the Windows Firewall with Advanced Security design for your organization is different from the deployment team that will implement it, make sure that the deployment team reviews the final design with the design team. Review the following points:

- The design team's strategy for determining how WMI and security group filters attached to the GPOs will determine which computers apply to which GPO. The deployment team can refer to the following topics in the Windows Firewall with Advanced Security Design Guide:
  - Planning Isolation Groups for the Zones
  - Planning the GPOs
  - Planning GPO Deployment
- The communication to be allowed between members of each of the zones in the isolated domain and computers that are not part of the isolated domain or members of the isolated domain's exemption list.
- The recommendation that domain controllers are exempted from IPsec authentication requirements. If they are not exempt and authentication fails, then domain clients might not be able to receive Group Policy updates to the IPsec connection security rules from the domain controllers.
- The rationale for configuring all IPsec authentication rules to request, not require, authentication until the successful negotiation of IPsec has been confirmed. If the rules are set to require authentication before confirming that authentication is working correctly, then communications between computers might fail. If the rules are set to request authentication

only, then an IPsec authentication failure results in fall-back-to-clear behavior, so communications can continue while the authentication failures are investigated.

- The requirement that all computers that must communicate with each other share a common set of:
  - Authentication methods
  - Main mode key exchange algorithms
  - Quick mode data integrity algorithms

  If at least one set of each does not match between two computers, then the computers cannot successfully communicate.

After the design and deployment teams agree on these issues, they can proceed with the deployment of the Windows Firewall with Advanced Security design. For more information, see Implementing Your Windows Firewall with Advanced Security Design Plan.

Implementing Your Windows Firewall with Advanced Security Design Plan

The following are important factors in the implementation of your Windows Firewall with Advanced Security design plan:

- **Group Policy**. The Windows Firewall with Advanced Security designs make extensive use of Group Policy deployed by Active Directory Domain Services (AD DS). A sound Group Policy infrastructure is required to successfully deploy the firewall and IPsec settings and rules to the computers on your network. Group Policy Analysis and Troubleshooting Overview (http://technet.microsoft.com/library/jj134223.aspx) can help you review and change, if necessary, your Group Policy infrastructure.
- **Perimeter firewall**. Most organizations use a perimeter firewall to help protect the computers on the network from potentially malicious network traffic from outside of the organization's network boundaries. If you plan a deployment that includes a boundary zone to enable external computers to connect to computers in that zone, then you must allow that traffic through the perimeter firewall to the computers in the boundary zone.
- **Computers running operating systems other than Windows**. If your network includes computers that are not running the Windows operating system, then you must make sure that required communication with those computers is not blocked by the restrictions put in place by your design. You must do one of the following:
  - Include those computers in the isolated domain or zone by adding certificate-based authentication to your design. Many other operating systems can participate in an isolated domain or isolated server scenario, as long as certificate-based authentication is used.
  - Include the computer in the authentication exemption list included in your design. You can choose this option if for any reason the computer cannot participate in the isolated domain design.

How to implement your Windows Firewall with Advanced Security design using this guide

The next step in implementing your design is to determine in what order each of the deployment steps must be performed. This guide uses checklists to help you accomplish the various deployment tasks that are required to implement your design plan. As the following diagram shows, checklists and subchecklists are used as necessary to provide the end-to-end procedure for deploying a design.

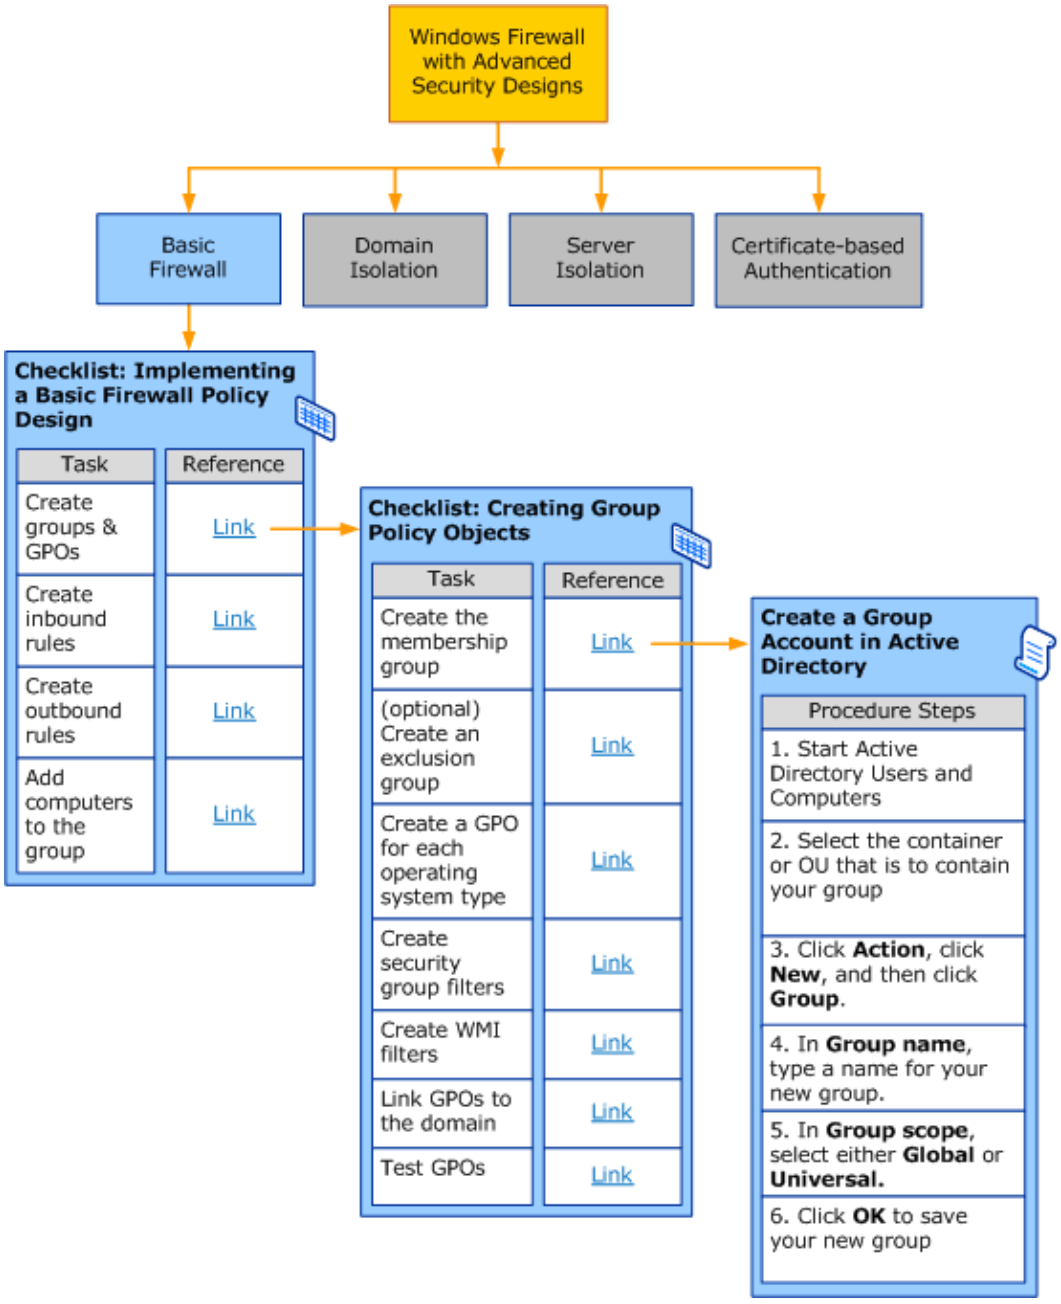

Use the following parent checklists in this section of the guide to become familiar with the deployment tasks for implementing your organization's Windows Firewall with Advanced Security design.

- Checklist: Implementing a Basic Firewall Policy Design
- Checklist: Implementing a Domain Isolation Policy Design
- Checklist: Implementing a Domain Isolation Policy Design
- Checklist: Implementing a Certificate-based Isolation Policy Design

The procedures in these checklists use the Group Policy MMC snap-in interfaces to configure firewall and connection security rules in GPOs, but you can also use Windows PowerShell. For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell. This guide recommends using GPOs in a specific way to deploy the rules and settings for your design. For information about deploying your GPOs, see Planning Group Policy Deployment for Your Isolation Zones and the checklist Checklist: Creating Group Policy Objects.

Checklist: Creating Group Policy Objects

To deploy firewall or IPsec settings or firewall or connection security rules, we recommend that you use Group Policy in AD DS. This section describes a tested, efficient method that requires some up-front work, but serves an administrator well in the long run by making GPO assignments as easy as dropping a computer into a membership group.

The checklists for firewall, domain isolation, and server isolation include a link to this checklist.

About membership groups

For most GPO deployment tasks, you must determine which computers must receive and apply which GPOs. Because different versions of Windows can support different settings and rules to achieve similar behavior, you might need multiple GPOs: one for each operating system that has settings different from the others to achieve the same result. For example, Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 use rules and settings that are incompatible with Windows 2000, Windows XP, and Windows Server 2003. Therefore, if your network included those older operating systems you would need to create a GPO for each set of operating systems that can share common settings. To deploy typical domain isolation settings and rules, you might have five different GPOs for the versions of Windows discussed in this guide. By following the procedures in this guide, you only need one membership group to manage all five GPOs. The membership group is identified in the security group filter for all five GPOs. To apply the settings to a computer, you make that computer's account a member of the membership group. WMI filters are used to ensure that the correct GPO is applied.

About exclusion groups

A Windows Firewall with Advanced Security design must often take into account domain-joined computers on the network that cannot or must not apply the rules and settings in the GPOs. Because these computers are typically fewer in number than the computers that must apply the GPO, it is easier to use the Domain Members group in the GPO membership group, and then place these exception computers into an exclusion group that is denied Apply Group Policy permissions on the GPO. Because deny permissions take precedence over allow permissions, a computer that is a member of both the membership group and the exception group is prevented from applying the GPO. Computers typically found in a GPO exclusion group for domain isolation include the domain controllers, DHCP servers, and DNS servers.

You can also use a membership group for one zone as an exclusion group for another zone. For example, computers in the boundary and encryption zones are technically in the main domain isolation zone, but must apply only the GPO for their assigned role. To do this, the GPOs for the main isolation zone deny Apply Group Policy permissions to members of the boundary and encryption zones.

Checklist: Creating Group Policy objects

	Task	Reference
☐	Review important concepts and examples for deploying GPOs in a way that best meets the needs of your organization.	Identifying Your Windows Firewall with Advanced Security Deployment Goals Planning Group Policy Deployment for Your Isolation Zones
☐	Create the membership group in AD DS that will be used to contain computer accounts that must receive the GPO. If some computers in the membership group are running an operating system that does not support WMI filters, such as Windows 2000, create an exclusion group to contain the computer accounts for the computers that cannot be blocked by using a WMI filter.	Create a Group Account in Active Directory
☐	Create a GPO for each version of Windows that has different implementation requirements.	Create a Group Policy Object

	Task	Reference
☐	Create security group filters to limit the GPO to only computers that are members of the membership group and to exclude computers that are members of the exclusion group.	Assign Security Group Filters to the GPO
☐	Create WMI filters to limit each GPO to only the computers that match the criteria in the filter.	Create WMI Filters for the GPO
☐	If you are working on a GPO that was copied from another, modify the group memberships and WMI filters so that they are correct for the new zone or version of Windows for which this GPO is intended.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Before adding any rules or configuring the GPO, add a few test computers to the membership group, and make sure that the correct GPO is received and applied to each member of the group.	Add Test Computers to the Membership Group for a Zone

Checklist: Implementing a Basic Firewall Policy Design

This parent checklist includes cross-reference links to important concepts about the basic firewall policy design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

Notes

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

The procedures in this section use the Group Policy MMC snap-in interfaces to configure the GPOs, but you can also use Windows PowerShell. For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell at http://technet.microsoft.com/library/hh831755.aspx.

Checklist: Implementing a basic firewall policy design

	Task	Reference
☐	Review important concepts and examples for the basic firewall policy design to determine if this design meets the needs of your organization.	Identifying Your Windows Firewall with Advanced Security Deployment Goals Basic Firewall Policy Design Firewall Policy Design Example Planning Settings for a Basic Firewall Policy
☐	Create the membership group and a GPO for each set of computers that require different firewall rules. Where GPOs will be similar, such as for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2, create one GPO, configure it by using the tasks in this checklist, and then make a copy of the GPO for the other version of Windows. For example, create and configure the GPO for Windows 8, make a copy of it for Windows Server 2012, and then follow the steps in this checklist to make the few required changes to the copy.	Checklist: Creating Group Policy Objects Copy a GPO to Create a New GPO
☐	If you are working on a GPO that was copied from another, modify the group membership and WMI	Modify GPO Filters to Apply to a Different Zone or Version of Windows

	Task	Reference
	filters so that they are correct for the computers for which this GPO is intended.	
☐	Configure the GPO with firewall default settings appropriate for your design.	Checklist: Configuring Basic Firewall Settings
☐	Create one or more inbound firewall rules to allow unsolicited inbound network traffic.	Checklist: Creating Inbound Firewall Rules
☐	Create one or more outbound firewall rules to block unwanted outbound network traffic.	Checklist: Creating Outbound Firewall Rules
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add test computers to the membership group, and then confirm that the computers receive the firewall rules from the GPOs as expected.	Add Test Computers to the Membership Group for a Zone
☐	According to the testing and roll-out schedule in your design plan, add computer accounts to the membership group to deploy the completed firewall policy settings to your computers.	Add Production Computers to the Membership Group for a Zone

Checklist: Configuring Basic Firewall Settings

This checklist includes tasks for configuring a GPO with firewall defaults and settings that are separate from the rules.

✓ **Checklist: Configuring firewall defaults and settings**

	Task	Reference
☐	Turn the firewall on and set the default inbound and outbound behavior.	Turn on Windows Firewall and Configure Default Behavior
☐	Configure the firewall to not display notifications to the user when a program is blocked, and to ignore locally defined firewall and connection security rules.	Configure Windows Firewall to Suppress Notifications When a Program Is Blocked
☐	Configure the firewall to record a log file.	Configure the Windows Firewall Log

Checklist: Creating Inbound Firewall Rules

This checklist includes tasks for creating firewall rules in your GPOs.

Checklist: Creating inbound firewall rules for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

	Task	Reference
☐	Create a rule that allows a program to listen for and accept inbound network traffic on any ports it requires.	Create an Inbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
☐	Create a rule that allows inbound network traffic on a specified port number.	Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
☐	Create a rule that allows inbound ICMP network traffic.	Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
☐	Create rules that allow inbound RPC network traffic.	Create Inbound Rules to Support RPC on Windows 8,

	Task	Reference
		Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
☐	Enable a predefined rule or a group of predefined rules. Some predefined rules for basic network services are included as part of the installation of Windows; others can be created when you install a new application or network service.	Enable Predefined Inbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Checklist: Creating Outbound Firewall Rules

This checklist includes tasks for creating outbound firewall rules in your GPOs. Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 support the use of outbound rules.

Important
By default, in Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2, outbound filtering is disabled. Because all outbound network traffic is permitted, outbound rules are typically used to block traffic that is not wanted on the network. However, it is a best practice for an administrator to create outbound allow rules for those applications that are approved for use on the organization's network. If you do this, then you have the option to set the default outbound behavior to block, preventing any network traffic that is not specifically authorized by the rules you create.

Checklist: Creating outbound firewall rules for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

	Task	Reference
☐	Create a rule that allows a program to send any outbound network traffic on any port it requires.	Create an Outbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

	Task	Reference
☐	Create a rule that allows outbound network traffic on a specified port number.	Create an Outbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2
☐	Enable a predefined rule or a group of predefined rules. Some predefined rules for basic network services are included as part of the installation of Windows; others can be created when you install a new application or network service.	Enable Predefined Outbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Checklist: Implementing a Domain Isolation Policy Design

This parent checklist includes cross-reference links to important concepts about the domain isolation policy design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

Notes

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

The procedures in this section use the Group Policy MMC snap-ins to configure the GPOs, but you can also use Windows PowerShell to configure GPOs. For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell at http://technet.microsoft.com/library/hh831755.aspx.

For more information about the security algorithms and authentication methods available in each version of Windows, see IPsec Algorithms and Methods Supported in Windows at http://technet.microsoft.com/library/dd125380.aspx.

Checklist: Implementing a domain isolation policy design

	Task	Reference
☐	Review important concepts and examples for the domain	Identifying Your Windows Firewall with Advanced Security

	Task	Reference
	isolation policy design, determine your Windows Firewall with Advanced Security deployment goals, and customize this design to meet the needs of your organization.	Deployment Goals Domain Isolation Policy Design Domain Isolation Policy Design Example Planning Domain Isolation Zones
☐	Create the GPOs and connection security rules for the isolated domain.	Checklist: Configuring Rules for the Isolated Domain
☐	Create the GPOs and connection security rules for the boundary zone.	Checklist: Configuring Rules for the Boundary Zone
☐	Create the GPOs and connection security rules for the encryption zone.	Checklist: Configuring Rules for the Encryption Zone
☐	Create the GPOs and connection security rules for the isolated server zone.	Checklist: Configuring Rules for an Isolated Server Zone
☐	According to the testing and roll-out schedule in your design plan, add computer accounts to the membership group to deploy rules and settings to your computers.	Add Production Computers to the Membership Group for a Zone
☐	After you confirm that network traffic is authenticated by IPsec, you can change authentication rules for the isolated domain and encryption zone from request to require mode.	Change Rules from Request to Require Mode

Checklist: Configuring Rules for the Isolated Domain

The following checklists include tasks for configuring connection security rules and IPsec settings in your GPOs to implement the main zone in the isolated domain.

Checklist: Configuring isolated domain rules for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Note

The GPOs for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 are usually similar. If this is true for your design, create one GPO, configure it by using the tasks in this checklist, and then make a copy of the GPO for the other operating system. For example, create and configure the GPO for Windows 8, make a copy of it for Windows Server 2012, and then follow the steps in this checklist to make the few required changes to the copy.

	Task	Reference
☐	Create a GPO for the computers in the isolated domain running one of the operating systems. After you have finished the tasks in this checklist and configured the GPO for that version of Windows, you can create a copy of it.	Checklist: Creating Group Policy Objects Copy a GPO to Create a New GPO
☐	If you are working on a GPO that was copied from another GPO, modify the group memberships and WMI filters so that they are correct for the isolated domain zone and the version of Windows for which this GPO is intended.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Configure IPsec to exempt all ICMP network traffic from IPsec protection.	Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows

	Task	Reference
		Server 2008 R2
☐	Create a rule that exempts all network traffic to and from computers on the exemption list from IPsec.	Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the key exchange (main mode) security methods and algorithms to be used.	Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the data protection (quick mode) algorithm combinations to be used.	Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the authentication methods to be used.	Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create the rule that requests authentication for all inbound network traffic.	Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Link the GPO to the domain level of the AD DS organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test computers to the membership group for the isolated domain. Be sure to add at least one for each operating	Add Test Computers to the Membership Group for a Zone

	Task	Reference
	system supported by a different GPO in the group.	
☐	Verify that the connection security rules are protecting network traffic to and from the test computers.	Verify That Network Traffic Is Authenticated

Do not change the rules for any of your zones to require authentication until all of the zones have been set up and are operating correctly.

Checklist: Configuring Rules for the Boundary Zone

The following checklists include tasks for configuring connection security rules and IPsec settings in your GPOs to implement the boundary zone in an isolated domain.

Rules for the boundary zone are typically the same as those for the isolated domain, with the exception that the final rule is left to only request, not require, authentication.

Checklist: Configuring boundary zone rules for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

A GPO for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 can simply be copied and then customized. This checklist assumes that you have already created the GPO for the isolated domain as described in Checklist: Implementing a Domain Isolation Policy Design. After you create a copy for the boundary zone, make sure that you do not change the rule from request authentication to require authentication when you create the other GPOs.

	Task	Reference
☐	Make a copy of the domain isolation GPO for this version of Windows to serve as a starting point for the GPO for the boundary zone. Unlike the GPO for the main isolated domain zone, this copy is not changed after deployment to require authentication.	Copy a GPO to Create a New GPO

	Task	Reference
☐	If you are working on a copy of a GPO, modify the group memberships and WMI filters so that they are correct for the boundary zone and version of Windows for which this GPO is intended.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test computers to the membership group for the boundary zone. Be sure to add at least one for each operating system supported by a different GPO in the group.	Add Test Computers to the Membership Group for a Zone
☐	Verify that the connection security configuration is protecting network traffic with authentication when it can, and that unauthenticated traffic is accepted.	Verify That Network Traffic Is Authenticated

Checklist: Configuring Rules for the Encryption Zone

This checklist includes tasks for configuring connection security rules and IPsec settings in your GPOs to implement the encryption zone in an isolated domain.

Rules for the encryption zone are typically the same as those for the isolated domain, with the exception that the main rule requires encryption in addition to authentication.

Checklist: Configuring encryption zone rules for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

A GPO for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 can simply be copied and then customized. This checklist assumes that you have already created the GPO for the isolated domain as described in Checklist: Implementing a Domain Isolation Policy Design. You can then copy those GPOs for

use with the encryption zone. After you create the copies, modify the main rule to require encryption in addition to the authentication required by the rest of the isolated domain.

	Task	Reference
☐	Make a copy of the domain isolation GPOs to serve as a starting point for the GPOs for the encryption zone.	Copy a GPO to Create a New GPO
☐	Modify the group memberships and WMI filters so that they are correct for the encryption zone and the version of Windows for which this GPO is intended.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Add the encryption requirements for the zone.	Configure the Rules to Require Encryption on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test computers to the membership group for the encryption zone. Be sure to add at least one for each operating system supported by a different GPO in the group.	Add Test Computers to the Membership Group for a Zone
☐	Verify that the connection security rules are protecting network traffic.	Verify That Network Traffic Is Authenticated

Checklist: Configuring Rules for an Isolated Server Zone

The following checklists include tasks for configuring connection security rules and IPsec settings in your GPOs for servers in an isolated server zone that are part of an isolated domain. For

information about creating a standalone isolated server zone that is not part of an isolated domain, see Checklist: Implementing a Standalone Server Isolation Policy Design.

In addition to requiring authentication and optionally encryption, servers in an isolated server zone can be accessed only by users or computers who are authenticated members of a network access group (NAG). Computers that are running Windows 2000, Windows XP, or Windows Server 2003 can restrict access in IPsec only to computers that are members of the NAG, because IPsec and IKE in those versions of Windows do not support user-based authentication. If you include user accounts in the NAG, then the restrictions can still apply; they are just enforced at the application layer, rather than the IP layer.

Computers that are running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 can identify both computers and users in the NAG because IPsec in these versions of Windows supports AuthIP in addition to IKE. AuthIP adds support for user-based authentication. For more information, see “AuthIP in Windows Vista” (http://go.microsoft.com/fwlink/?linkid=69843).

The GPOs for an isolated server or group of servers are similar to those for the isolated domain itself or the encryption zone, if you require encryption to your isolated servers. This checklist refers you to procedures for creating rules as well as restrictions that allow only members of the NAG to connect to the server.

Checklist: Configuring rules for isolated servers for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Note

The GPOs for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 are usually similar. If this is true for your design, create one GPO, configure it by using the tasks in this checklist, and then make a copy of the GPO for the other operating system. For example, create and configure the GPO for Windows 8, make a copy of it for Windows Server 2012, and then follow the steps in this checklist to make the few required changes to the copy.

	Task	Reference
☐	Create a GPO for the computers that need to have access restricted to the same set of client computers. If there are multiple servers and they run different versions of the Windows operating system, then start by creating the GPO for one version of Windows. After you have finished the tasks in this checklist and configured the GPO for that	Copy a GPO to Create a New GPO

	Task	Reference
	version of Windows, you can create a copy of it. Copy the GPO from the isolated domain or from the encryption zone to serve as a starting point. Where your copy already contains elements listed in the following checklist, review the relevant procedures and compare them to your copied GPO's element to make sure it is constructed in a way that meets the needs of the server isolation zone.	
☐	Configure the security group filters and WMI filters on the GPO so that only members of the isolated server zone's membership group that are running the specified version of Windows can read and apply it.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Configure IPsec to exempt all ICMP network traffic from IPsec protection.	Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the key exchange (main mode) security methods and algorithms to be used.	Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the data protection (quick mode) algorithm combinations to be used. If you require encryption for the isolated server zone, then make sure that you choose only algorithm combinations that include encryption.	Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

	Task	Reference
☐	Configure the authentication methods to be used.	Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create a rule that exempts all network traffic to and from computers on the exemption list from IPsec.	Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create a rule that requests authentication for all network traffic. **Important** Just as in an isolated domain, do not set the rules to require authentication for inbound traffic until you have completed testing. That way, if the rules do not work as expected, communications are not affected by a failure to authenticate.	Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create the NAG to contain the computer or user accounts that are allowed to access the servers in the isolated server zone.	Create a Group Account in Active Directory
☐	Create a firewall rule that permits inbound network traffic only if authenticated as a member of the NAG.	Restrict Server Access to Members of a Group Only
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test server to the	Add Test Computers to the

	Task	Reference
	membership group for the isolated server zone. Be sure to add at least one server for each operating system supported by a GPO in the group.	Membership Group for a Zone

Do not change the rules for any of your zones to require authentication until all of the zones have been set up and are operating correctly.

Checklist: Implementing a Standalone Server Isolation Policy Design

This checklist contains procedures for creating a server isolation policy design that is not part of an isolated domain. For the steps required to create an isolated server zone within an isolated domain, see Checklist: Configuring Rules for an Isolated Server Zone.

This parent checklist includes cross-reference links to important concepts about the domain isolation policy design. It also contains links to subordinate checklists that will help you complete the tasks that are required to implement this design.

Notes

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist.

The procedures in this section use the Group Policy MMC snap-in interfaces to configure the GPOs, but you can also use Windows PowerShell. For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell at http://technet.microsoft.com/library/hh831755.aspx.

Checklist: Implementing a standalone server isolation policy design

	Task	Reference
☐	Review important concepts and examples for the server isolation policy design to determine if this design meets your deployment goals and the needs of your organization.	Identifying Your Windows Firewall with Advanced Security Deployment Goals Server Isolation Policy Design Server Isolation Policy Design Example Planning Server Isolation

	Task	Reference
		Zones
☐	Create the GPOs and connection security rules for isolated servers.	Checklist: Configuring Rules for Servers in a Standalone Isolated Server Zone
☐	Create the GPOs and connection security rules for the client computers that must connect to the isolated servers.	Checklist: Creating Rules for Clients of a Standalone Isolated Server Zone
☐	Verify that the connection security rules are protecting network traffic on your test computers.	Verify That Network Traffic Is Authenticated
☐	After you confirm that network traffic is authenticated by IPsec as expected, you can change authentication rules for the isolated server zone to require authentication instead of requesting it.	Change Rules from Request to Require Mode
☐	According to the testing and roll-out schedule in your design plan, add computer accounts for the client computers to the membership group so that you can deploy the settings.	Add Production Computers to the Membership Group for a Zone

Checklist: Configuring Rules for Servers in a Standalone Isolated Server Zone

This checklist includes tasks for configuring connection security rules and IPsec settings in your GPOs for servers in a standalone isolated server zone that is not part of an isolated domain. In addition to requiring authentication and optionally encryption, servers in a server isolation zone are accessible only by users or computers that are authenticated as members of a network access group (NAG). The GPOs described here apply only to the isolated servers, not to the client computers that connect to them. For the GPOs for the client computers, see Checklist: Creating Rules for Clients of a Standalone Isolated Server Zone.

The GPOs for isolated servers are similar to those for an isolated domain. This checklist refers you to those procedures for the creation of some of the rules. The other procedures in this checklist are for creating the restrictions that allow only members of the server access group to connect to the server.

Checklist: Configuring rules for isolated servers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Note

The GPOs for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 are usually similar. If this is true for your design, create one GPO, configure it by using the tasks in this checklist, and then create a copy of the GPO for the other operating system. For example, create and configure the GPO for Windows 8, make a copy of it for Windows Server 2012, and then follow the steps in this checklist to make the few required changes to the copy.

	Task	Reference
☐	Create a GPO for the computers that need to have access restricted to the same set of client computers. If there are multiple servers running different versions of the Windows operating system, start by creating the GPO for one version of Windows. After you have finished the tasks in this checklist and configured the GPO for that version of Windows, you can create a copy of it.	Checklist: Creating Group Policy Objects Copy a GPO to Create a New GPO
☐	If you are working on a copy of a GPO, modify the group memberships and WMI filters so that they are correct for the computers for which this GPO is intended.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Configure IPsec to exempt all ICMP network traffic from IPsec protection.	Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

	Task	Reference
☐	Create a rule that exempts all network traffic to and from computers on the exemption list from IPsec.	Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the key exchange (main mode) security methods and algorithms to be used.	Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the data protection (quick mode) algorithm combinations to be used.	Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the authentication methods to be used. This procedure sets the default settings for the computer. If you want to set authentication on a per-rule basis, this procedure is optional.	Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create a rule that requests authentication for all inbound network traffic. **Important** Just as in an isolated domain, do not set the rules to require authentication until your testing is complete. That way, if the rules do not work as expected, communications are not affected by a failure	Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

	Task	Reference
	to authenticate.	
☐	If your design requires encryption in addition to authentication for access to the isolated servers, then modify the rule to require it.	Configure the Rules to Require Encryption on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create the NAG to contain the computer or user accounts that are allowed to access the isolated servers. If you have multiple groups of isolated servers that are accessed by different client computers, then create a NAG for each set of servers.	Create a Group Account in Active Directory
☐	Create a firewall rule that allows inbound network traffic only if it is authenticated from a user or computer that is a member of the zone's NAG.	Restrict Server Access to Members of a Group Only
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test server to the membership group for the isolated server zone. Be sure to add at least one for each operating system supported by a different GPO in the group.	Add Test Computers to the Membership Group for a Zone

Do not change the rules for any of your zones to require authentication until all zones have been set up and thoroughly tested.

Checklist: Creating Rules for Clients of a Standalone Isolated Server Zone

This checklist includes tasks for configuring connection security rules and IPsec settings in the GPOs for client computers that must connect to servers in an isolated server zone.

Checklist: Configuring isolated server zone client rules for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Note

The GPOs for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 are usually similar. If this is true for your design, create one GPO, configure it by using the tasks in this checklist, and then create a copy of the GPO. For example, create and configure the GPO for Windows 8, create a copy of it for Windows Server 2012, and then follow the steps in this checklist to make the required changes (if any) to the copy.

	Task	Reference
☐	Create a GPO for the client computers that must connect to servers in the isolated server zone, and that are running one of the versions of Windows. After you have finished the tasks in this checklist, you can make a copy of it.	Checklist: Creating Group Policy Objects Copy a GPO to Create a New GPO
☐	To determine which computers receive the GPO, assign the NAG for the isolated servers to the security group filter for the GPO. Make sure that each GPO has the WMI filter for the correct version of Windows.	Modify GPO Filters to Apply to a Different Zone or Version of Windows
☐	Configure IPsec to exempt all ICMP network traffic from IPsec protection.	Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create a rule that exempts all	Create an Authentication

	Task	Reference
	network traffic to and from computers on the exemption list from IPsec.	Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the key exchange (main mode) security methods and algorithms to be used.	Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the data protection (quick mode) algorithm combinations to be used.	Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Configure the authentication methods to be used.	Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Create a rule that requests authentication for network traffic. Because fallback-to-clear behavior in Windows Vista and Windows Server 2008 has no delay when communicating with computers that cannot use IPsec, you can use the same any-to-any rule used in an isolated domain.	Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2
☐	Link the GPO to the domain level of the Active Directory organizational unit hierarchy.	Link the GPO to the Domain
☐	Add your test computers to the NAG for the isolated server zone. Be sure to add at least one	Add Test Computers to the Membership Group for a Zone

	Task	Reference
	for each operating system supported by a different GPO in the group.	

Checklist: Implementing a Certificate-based Isolation Policy Design

This parent checklist includes cross-reference links to important concepts about using certificates as an authentication option in either a domain isolation or server isolation design.

Notes

Complete the tasks in this checklist in order. When a reference link takes you to a procedure, return to this topic after you complete the steps in that procedure so that you can proceed with the remaining tasks in this checklist

The procedures in this section use the Group Policy MMC snap-in interfaces to configure the GPOs, but you can also use Windows PowerShell. For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell at http://technet.microsoft.com/library/hh831755.aspx.

Checklist: Implementing certificate-based authentication

	Task	Reference
☐	Review important concepts and examples for certificate-based authentication to determine if this design meets your deployment goals and the needs of your organization.	Identifying Your Windows Firewall with Advanced Security Deployment Goals Certificate-based Isolation Policy Design Certificate-based Isolation Policy Design Example Planning Certificate-based Authentication
☐	Install the Active Directory Certificate Services (AD CS) role as an enterprise root issuing certification authority (CA). This step is required only if you have not already deployed a CA on your network.	Install Active Directory Certificate Services

	Task	Reference
☐	Configure the certificate template for workstation authentication certificates.	Configure the Workstation Authentication Certificate Template
☐	Configure Group Policy to automatically deploy certificates based on your template to workstation computers.	Configure Group Policy to Autoenroll and Deploy Certificates
☐	On a test computer, refresh Group Policy and confirm that the certificate is installed.	Confirm That Certificates Are Deployed Correctly

Procedures Used in This Guide

The procedures in this section appear in the checklists found earlier in this document. They should be used only in the context of the checklists in which they appear. They are presented here in alphabetical order.

Add Production Computers to the Membership Group for a Zone

Add Test Computers to the Membership Group for a Zone

Assign Security Group Filters to the GPO

Change Rules from Request to Require Mode

Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Configure Group Policy to Autoenroll and Deploy Certificates

Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Configure the Rules to Require Encryption on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Configure the Windows Firewall Log

Configure the Workstation Authentication Certificate Template

Configure Windows Firewall to Suppress Notifications When a Program Is Blocked

Confirm That Certificates Are Deployed Correctly

Copy a GPO to Create a New GPO

Create a Group Account in Active Directory

Create a Group Policy Object

Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Create an Inbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Create an Outbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

Create an Outbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

Create Inbound Rules to Support RPC on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Create WMI Filters for the GPO

Enable Predefined Inbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Enable Predefined Outbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

Install Active Directory Certificate Services

Link the GPO to the Domain

Modify GPO Filters to Apply to a Different Zone or Version of Windows

Open the Group Policy Management Console to IP Security Policies

Open the Group Policy Management Console to Windows Firewall

Open the Group Policy Management Console to Windows Firewall with Advanced Security

Open Windows Firewall with Advanced Security

Restrict Server Access to Members of a Group Only

Start a Command Prompt as an Administrator

Turn on Windows Firewall and Configure Default Behavior

Verify That Network Traffic Is Authenticated

Add Production Computers to the Membership Group for a Zone

After you test the GPOs for your design on a small set of computers, you can deploy them to the production computers.

Caution
For GPOs that contain connection security rules that prevent unauthenticated connections, be sure to set the rules to request, not require, authentication during testing. After you deploy the GPO and confirm that all of your computers are successfully communicating by using authenticated IPsec, then you can modify the GPO to require authentication. Do not change the boundary zone GPO to require mode.

The method discussed in this guide uses the **Domain Computers** built-in group. The advantage of this method is that all new computers that are joined to the domain automatically receive the isolated domain GPO. To do this successfully, you must make sure that the WMI filters and security group filters exclude computers that must not receive the GPOs. Use computer groups that deny both read and apply Group Policy permissions to the GPOs, such as a group used in the CG_DOMISO_NOIPSEC example design. Computers that are members of some zones must also be excluded from applying the GPOs for the main isolated domain. For more information, see the "Prevent members of a group from applying a GPO" section in Assign Security Group Filters to the GPO.

Without such a group (or groups), you must either add computers individually or use the groups containing computer accounts that are available to you.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the membership of the group for the GPO.

In this topic:

- Add the group Domain Computers to the GPO membership group
- Refresh Group Policy on the computers in the membership group
- Check which GPOs apply to a computer

To add domain computers to the GPO membership group

1. On a computer that has the Active Directory management tools installed, click the **Start** charm, then click the **Active Directory Users and Computers** tile.
2. In the navigation pane, expand **Active Directory Users and Computers**, expand *YourDomainName*, and then the container in which you created the membership group.
3. In the details pane, double-click the GPO membership group to which you want to add computers.
4. Select the **Members** tab, and then click **Add**.
5. Type **Domain Computers** in the text box, and then click **OK**.

6. Click **OK** to close the group properties dialog box.

After a computer is a member of the group, you can force a Group Policy refresh on the computer.

To refresh Group Policy on a computer

- For a computer that is running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, Start a Command Prompt as an Administrator, and then type the following command:

```
gpupdate /target:computer /force
```

After Group Policy is refreshed, you can see which GPOs are currently applied to the computer.

To see which GPOs are applied to a computer

- For a computer that is running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, Start a Command Prompt as an Administrator, and then type the following command:

```
gpresult /r /scope:computer
```

Add Test Computers to the Membership Group for a Zone

Before you deploy your rules to large numbers of computers, you must thoroughly test the rules to make sure that communications are working as expected. A misplaced WMI filter or an incorrectly typed IP address in a filter list can easily block communications between computers. Although we recommend that you set your rules to request mode until testing and deployment is complete, we also recommend that you initially deploy the rules to a small number of computers only to be sure that the correct GPOs are being processed by each computer.

Add at least one computer of each supported operating system type to each membership group. Make sure every GPO for a specific version of Windows and membership group has a computer among the test group. After Group Policy has been refreshed on each test computer, check the output of the **gpresult** command to confirm that each computer is receiving only the GPOs it is supposed to receive.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the membership of the group for the GPO.

In this topic:

- Add the test computers to the GPO membership groups
- Refresh Group Policy on the computers in each membership group

- Check which GPOs apply to a computer

To add test computers to the GPO membership groups

1. On a computer that has the Active Directory management tools installed, click the **Start** charm, then click the **Active Directory Users and Computers** tile.
2. In the navigation pane, expand **Active Directory Users and Computers**, expand *YourDomainName*, and then expand the container that holds your membership group account.
3. In the details pane, double-click the GPO membership group to which you want to add computers.
4. Select the **Members** tab, and then click **Add**.
5. Type the name of the computer in the text box, and then click **OK**.
6. Repeat steps 5 and 6 for each additional computer account or group that you want to add.
7. Click **OK** to close the group properties dialog box.

After a computer is a member of the group, you can force a Group Policy refresh on the computer.

To refresh Group Policy on a computer

- For a computer that is running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, Start a Command Prompt as an Administrator, and then type the following command:

```
gpupdate /target:computer /force
```

After Group Policy is refreshed, you can see which GPOs are currently applied to the computer.

To see which GPOs are applied to a computer

- For a computer that is running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, Start a Command Prompt as an Administrator, and then type the following command:

```
gpresult /r /scope:computer
```

Assign Security Group Filters to the GPO

To make sure that your GPO is applied to the correct computers, use the Group Policy Management MMC snap-in to assign security group filters to the GPO.

Important

This deployment guide uses the method of adding the Domain Computers group to the membership group for the main isolated domain after testing is complete and you are ready to go live in production. To make this method work, you must prevent any computer that is a member of either the boundary or encryption zone from applying the GPO for the main isolated domain. For example, on the GPOs for the main isolated domain, deny Read and Apply Group Policy permissions to the membership groups for the boundary and encryption zones.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the relevant GPOs.

In this topic:

- Allow members of a group to apply a GPO
- Prevent members of a group from applying a GPO

Use the following procedure to add a group to the security filter on the GPO that allows group members to apply the GPO.

To allow members of a group to apply a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, find and then click the GPO that you want to modify.
3. In the details pane, under **Security Filtering**, click **Authenticated Users**, and then click **Remove**.

   Note
   You must remove the default permission granted to all authenticated users and computers to restrict the GPO to only the groups you specify.
4. Click **Add**.
5. In the **Select User, Computer, or Group** dialog box, type the name of the group whose members are to apply the GPO, and then click **OK**. If you do not know the name, you can click **Advanced** to browse the list of groups available in the domain.

Use the following procedure to add a group to the security filter on the GPO that prevents group members from applying the GPO. This is typically used to prevent members of the boundary and encryption zones from applying the GPOs for the isolated domain.

To prevent members of group from applying a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, find and then click the GPO that you want to modify.
3. In the details pane, click the **Delegation** tab.
4. Click **Advanced**.

5. Under the **Group or user names** list, click **Add**.
6. In the **Select User, Computer, or Group** dialog box, type the name of the group whose members are to be prevented from applying the GPO, and then click **OK**. If you do not know the name, you can click **Advanced** to browse the list of groups available in the domain.
7. Select the group in the **Group or user names** list, and then select the box in the **Deny** column for both **Read** and **Apply group policy**.
8. Click **OK**, and then in the **Windows Security** dialog box, click **Yes**.
9. The group appears in the list with **Custom** permissions.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Change Rules from Request to Require Mode

After you confirm that network traffic is being correctly protected by using IPsec, you can change the rules for the domain isolation and encryption zones to require, instead of request, authentication. Do not change the rules for the boundary zone; they must stay in request mode so that computers in the boundary zone can continue to accept connections from computers that are not part of the isolated domain.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

In this topic:

- Convert a rule in a GPO for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
- **Convert a rule for an earlier version of Windows**
- Refresh policy on the client computers to receive the modified GPOs

To convert a rule from request to require mode for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Connection Security Rules**.
3. In the details pane, double-click the connection security rule that you want to modify.
4. Click the **Authentication** tab.
5. In the **Requirements** section, change **Authenticated mode** to **Require inbound and request outbound**, and then click **OK**.

To apply the modified GPOs to the client computers

1. The next time each computer refreshes its Group Policy, it will receive the updated GPO and apply the modified rule. To force an immediate refresh, Start a Command Prompt as an Administrator and run the following command:

   ```
   gpupdate /force
   ```

2. To verify that the modified GPO is correctly applied to the client computers, you can run one of the following commands:

   On computers that are running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, run the following command:

   ```
   gpresult /r /scope computer
   ```

3. Examine the command output for the list of GPOs that are applied to the computer, and make sure that the list contains the GPOs you expect to see on that computer.

Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

This procedure shows you how to configure the authentication methods that can be used by computers in an isolated domain or standalone isolated server zone.

Note

If you follow the steps in the procedure in this topic, you alter the system-wide default settings. Any connection security rule can use these settings by specifying **Default** on the **Authentication** tab.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To configure authentication methods

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane on the main Windows Firewall with Advanced Security page, click **Windows Firewall Properties**.
3. On the **IPsec Settings** tab, click **Customize**.
4. In the **Authentication Method** section, select the type of authentication that you want to use from among the following:
   a. **Default**. Selecting this option tells the computer to use the authentication method

currently defined by the local administrator in Windows Firewall with Advanced Security or by Group Policy as the default.

b. **Computer and User (using Kerberos V5)**. Selecting this option tells the computer to use and require authentication of both the computer and the currently logged-on user by using their domain credentials. This authentication method works only with other computers that can use Authenticated IP (AuthIP), including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.

c. **Computer (using Kerberos V5)**. Selecting this option tells the computer to use and require authentication of the computer by using its domain credentials. This option works with other computers that can use IKE v1, including earlier versions of Windows.

d. **User (using Kerberos V5)**. Selecting this option tells the computer to use and require authentication of the currently logged-on user by using his or her domain credentials. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.

e. **Computer certificate from this certification authority**. Selecting this option and entering the identification of a certification authority (CA) tells the computer to use and require authentication by using a certificate that is issued by the selected CA. If you also select **Accept only health certificates**, then only certificates that include the system health authentication enhanced key usage (EKU) typically provided in a Network Access Protection (NAP) infrastructure can be used for this rule.

f. **Advanced**. Click **Customize** to specify a custom combination of authentication methods required for your scenario. You can specify both a **First authentication method** and a **Second authentication method**.

   The first authentication method can be one of the following:

   - **Computer (Kerberos V5)**. Selecting this option tells the computer to use and require authentication of the computer by using its domain credentials. This option works with other computers that can use IKE v1, including earlier versions of Windows.
   - **Computer (NTLMv2)**. Selecting this option tells the computer to use and require authentication of the computer by using its domain credentials. This option works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
   - **Computer certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to use and require authentication by using a certificate that is issued by that CA. If you also select **Accept only health certificates**, then only certificates issued by a NAP server can be used.

- **Preshared key (not recommended)**. Selecting this method and entering a preshared key tells the computer to authenticate by exchanging the preshared keys. If they match, then the authentication succeeds. This method is not recommended, and is included only for backward compatibility and testing purposes.

If you select **First authentication is optional**, then the connection can succeed even if the authentication attempt specified in this column fails.

The second authentication method can be one of the following:

- **User (Kerberos V5)**. Selecting this option tells the computer to use and require authentication of the currently logged-on user by using his or her domain credentials. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
- **User (NTLMv2)**. Selecting this option tells the computer to use and require authentication of the currently logged-on user by using his or her domain credentials, and uses the NTLMv2 protocol instead of Kerberos V5. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
- **User health certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to use and require user-based authentication by using a certificate that is issued by the specified CA. If you also select **Enable certificate to account mapping**, then the certificate can be associated with a user in Active Directory for purposes of granting or denying access to specified users or user groups.
- **Computer health certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to use and require authentication by using a certificate that is issued by the specified CA. If you also select **Accept only health certificates**, then only certificates that include the system health authentication EKU typically provided in a NAP infrastructure can be used for this rule.

If you select **Second authentication is optional**, then the connection can succeed even if the authentication attempt specified in this column fails.

Important
Make sure that you do not select the check boxes to make both first and second authentication optional. Doing so allows plaintext connections whenever authentication fails.

5. Click **OK** on each dialog box to save your changes and return to the Group Policy Management Editor.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

This procedure shows you how to configure the data protection (quick mode) settings for connection security rules in an isolated domain or a standalone isolated server zone.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To configure quick mode settings

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane on the main Windows Firewall with Advanced Security page, click **Windows Firewall Properties**.
3. On the **IPsec Settings** tab, click **Customize**.
4. In the **Data protection (Quick Mode)** section, click **Advanced**, and then click **Customize**.
5. If you require encryption for all network traffic in the specified zone, then check **Require encryption for all connection security rules that use these settings**. Selecting this option disables the **Data integrity** section, and forces you to select only integrity algorithms that are combined with an encryption algorithm. If you do not select this option, then you can use only data integrity algorithms. Before selecting this option, consider the performance impact and the increase in network traffic that will result. We recommend that you use this setting only on network traffic that truly requires it, such as to and from computers in the encryption zone.
6. If you did not select **Require encryption**, then select the data integrity algorithms that you want to use to help protect the data sessions between the two computers. If the data integrity algorithms displayed in the list are not what you want, then do the following:
   a. From the left column, remove any of the data integrity algorithms that you do not want by selecting the algorithm and then clicking **Remove**.
   b. Add any required data integrity algorithms by clicking **Add**, selecting the appropriate protocol (ESP or AH) and algorithm (SHA1 or MD5), selecting the key lifetime in minutes or sessions, and then clicking **OK**. We recommend that you do not include MD5 in any combination. It is included for backward compatibility only. We also recommend that you use ESP instead of AH if you have any devices on your network that use network address translation (NAT).

c. In **Key lifetime (in sessions)**, type the number of times that the quick mode session can be rekeyed. After this number is reached, the quick mode SA must be renegotiated. Be careful to balance performance with security requirements. Although a shorter key lifetime results in better security, it also reduces performance because of the more frequent renegotiating of the quick mode SA. We recommend that you use the default value unless your risk analysis indicates the need for a different value.

d. Click **OK** to save your algorithm combination settings.

e. After the list contains only the combinations you want, use the up and down arrows to the right of the list to rearrange them in the correct order for your design. The algorithm combination that is first in the list is tried first, and so on.

7. Select the data integrity and encryption algorithms that you want to use to help protect the data sessions between the two computers. If the algorithm combinations displayed in the list are not what you want, then do the following:

   a. From the second column, remove any of the data integrity and encryption algorithms that you do not want by selecting the algorithm combination and then clicking **Remove**.

   b. Add any required integrity and encryption algorithm combinations by clicking **Add**, and then doing the following:

   c. Select the appropriate protocol (ESP or AH). We recommend that you use ESP instead of AH if you have any devices on your network that use NAT.

   d. Select the appropriate encryption algorithm. The choices include, in order of decreasing security: AES-256, AES-192, AES-128, 3DES, and DES. We recommend that you do not include DES in any combination. It is included for backward compatibility only.

   e. Select the appropriate integrity algorithm (SHA1 or MD5). We recommend that you do not include MD5 in any combination. It is included for backward compatibility only.

   f. In **Key lifetime (in minutes)**, type the number of minutes. When the specified number of minutes has elapsed, any IPsec operations between the two computers that negotiated this key will require a new key. Be careful to balance performance with security requirements. Although a shorter key lifetime results in better security, it also reduces performance because of the more frequent rekeying. We recommend that you use the default value unless your risk analysis indicates the need for a different value.

8. Click **OK** three times to save your settings.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Configure Group Policy to Autoenroll and Deploy Certificates

You can use this procedure to configure Group Policy to automatically enroll client computer certificates and deploy them to the workstations on your network. Follow this procedure for each GPO that contains IPsec connection security rules that require this certificate.

Administrative credentials

To complete these procedures, you must be a member of both the Domain Admins group in the root domain of your forest and a member of the Enterprise Admins group.

To configure Group Policy to autoenroll certificates

1. On a computer that has the Group Policy Management feature installed, click **Start**, click **Administrative Tools**, and then click **Group Policy Management**.
2. In the navigation pane, expand **Forest:** *YourForestName*, expand **Domains**, expand *YourDomainName*, expand **Group Policy Objects**, right-click the GPO you want to modify, and then click **Edit**.
3. In the navigation pane, expand the following path: **Computer Configuration**, **Policies**, **Windows Settings**, **Security Settings**, **Public Key Policies**.
4. Double-click **Certificate Services Client - Auto-Enrollment**.
5. In the **Properties** dialog box, change **Configuration Model** to **Enabled**.
6. Select both **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**.
7. Click **OK** to save your changes. Computers apply the GPO and download the certificate the next time Group Policy is refreshed.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Configure Key Exchange (Main Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

This procedure shows you how to configure the main mode key exchange settings used to secure the IPsec authentication traffic.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To configure key exchange settings

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane on the main Windows Firewall with Advanced Security page, click **Windows Firewall Properties**.
3. On the **IPsec Settings** tab, click **Customize**.
4. In the **Key exchange (Main Mode)** section, click **Advanced**, and then click **Customize**.
5. Select the security methods to be used to help protect the main mode negotiations between the two computers. If the security methods displayed in the list are not what you want, then do the following:

   Important
   In Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, you can specify only one key exchange algorithm. This means that if you want to communicate by using IPsec with another computer running Windows 8 or Windows Server 2012, then you must select the same key exchange algorithm on both computers.

   Also, if you create a connection security rule that specifies an option that requires AuthIP instead of IKE, then only the one combination of the top integrity and encryption security method are used in the negotiation. Make sure that all of your computers that run Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 have the same methods at the top of the list and the same key exchange algorithm selected.

   Note
   When AuthIP is used, no Diffie-Hellman key exchange protocol is used. Instead, when Kerberos V5 authentication is requested, the Kerberos V5 service ticket secret is used in place of a Diffie-Hellman value. When either certificate authentication or NTLM authentication is requested, a transport level security (TLS) session is established, and its secret is used in place of the Diffie-Hellman value. This happens no matter which Diffie-Hellman key exchange protocol you select.

   a. Remove any of the security methods that you do not want by selecting the method and then clicking **Remove**.
   b. Add any required security method combinations by clicking **Add**, selecting the appropriate encryption algorithm and integrity algorithm from the lists, and then clicking **OK**.

      Caution
      We recommend that you do not include MD5 or DES in any combination.

They are included for backward compatibility only.

c. After the list contains only the combinations you want, use the up and down arrows to the right of the list to arrange them in the order of preference. The combination that appears first in the list is tried first, and so on.

6. From the list on the right, select the key exchange algorithm that you want to use.

   Caution
   We recommend that you do not use Diffie-Hellman Group 1. It is included for backward compatibility only.

7. In **Key lifetime (in minutes)**, type the number of minutes. When the specified number of minutes has elapsed, any IPsec operation between the two computers requires a new key.

   Note
   You need to balance performance with security requirements. Although a shorter key lifetime results in better security, it also reduces performance.

8. In **Key lifetime (in sessions)**, type the number of sessions. After the specified number of quick mode sessions have been created within the security association protected by this key, IPsec requires a new key.
9. Click **OK** three times to save your settings.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Configure the Rules to Require Encryption on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

If you are creating a zone that requires encryption, you must configure the rules to add the encryption algorithms and delete the algorithm combinations that do not use encryption.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To modify an authentication request rule to also require encryption

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Connection Security Rules**.
3. In the details pane, double-click the connection security rule you want to modify.

4. On the **Name** page, rename the connection security rule, edit the description to reflect the new use for the rule, and then click **OK**.
5. In the navigation pane, right-click **Windows Firewall with Advanced Security – LDAP://CN={*guid*}**, and then click **Properties**.
6. Click the **IPsec Settings** tab.
7. Under **IPsec defaults**, click **Customize**.
8. Under **Data protection (Quick Mode)**, click **Advanced**, and then click **Customize**.
9. Click **Require encryption for all connection security rules that use these settings**.

   This disables the data integrity rules section. Make sure the **Data integrity and encryption** list contains all of the combinations that your client computers will use to connect to members of the encryption zone. The client computers receive their rules through the GPO for the zone to which they reside. You must make sure that those rules contain at least one of the data integrity and encryption algorithms that are configured in this rule, or the client computers in that zone will not be able to connect to computers in this zone.
10. If you need to add an algorithm combination, click **Add**, and then select the combination of encryption and integrity algorithms. The options are described in Configure Data Protection (Quick Mode) Settings on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2.

    Notes

    Not all of the algorithms available in Windows 8 or Windows Server 2012 can be selected in the Windows Firewall with Advanced Security user interface. To select them, you can use Windows PowerShell.

    Quick mode settings can also be configured on a per-rule basis, but not by using the Windows Firewall with Advanced Security user interface. Instead, you can create or modify the rules by using Windows PowerShell.

    For more information, see Windows Firewall with Advanced Security Administration with Windows PowerShell
11. During negotiation, algorithm combinations are proposed in the order shown in the list. Make sure that the more secure combinations are at the top of the list so that the negotiating computers select the most secure combination that they can jointly support.
12. Click **OK** three times to save your changes.

Configure the Windows Firewall Log

To configure Windows Firewall to log dropped packets or successful connections, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

In this topic:

To configure Windows Firewall logging for Windows Vista or Windows Server 2008

To configure Windows Firewall logging for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane, in the **Overview** section, click **Windows Firewall Properties**.
3. For each network location type (Domain, Private, Public), perform the following steps.
   a. Click the tab that corresponds to the network location type.
   b. Under **Logging**, click **Customize**.
   c. The default path for the log is **%windir%\system32\logfiles\firewall\pfirewall.log**. If you want to change this, clear the **Not configured** check box and type the path to the new location, or click **Browse** to select a file location.

      Important
      The location you specify must have permissions assigned that permit the Windows Firewall service to write to the log file.

   d. The default maximum file size for the log is 4,096 kilobytes (KB). If you want to change this, clear the **Not configured** check box, and type in the new size in KB, or use the up and down arrows to select a size. The file will not grow beyond this size; when the limit is reached, old log entries are deleted to make room for the newly created ones.
   e. No logging occurs until you set one of following two options:
      - To create a log entry when Windows Firewall drops an incoming network packet, change **Log dropped packets** to **Yes**.
      - To create a log entry when Windows Firewall allows an inbound connection, change **Log successful connections** to **Yes**.
   f. Click **OK** twice.

Configure the Workstation Authentication Certificate Template

This procedure describes how to configure a certificate template that Active Directory Certification Services (AD CS) uses as the starting point for computer certificates that are automatically enrolled and deployed to workstations in the domain. It shows how to create a copy of a template, and then configure the template according to your design requirements.

Administrative credentials

To complete these procedures, you must be a member of both the Domain Admins group in the root domain of your forest, and a member of the Enterprise Admins group.

To configure the workstation authentication certificate template and autoenrollment

1. On the computer where AD CS is installed, click the **Start** charm, and then click **Certification Authority**.
2. In the navigation pane, right-click **Certificate Templates**, and then click **Manage**.
3. In the details pane, click the **Workstation Authentication** template.
4. On the **Action** menu, click **Duplicate Template**. In the **Duplicate Template** dialog box, select the template version that is appropriate for your deployment, and then click **OK**. For the resulting certificates to have maximum compatibility with the available versions of Windows, we recommended that you select **Windows Server 2003**.
5. On the **General** tab, in **Template display name**, type a new name for the certificate template, such as **Domain Isolation Workstation Authentication Template**.
6. Click the **Subject Name** tab. Make sure that **Build from this Active Directory information** is selected. In **Subject name format**, select **Fully distinguished name**.
7. Click the **Request Handling** tab. You must determine the best minimum key size for your environment. Large key sizes provide better security, but they can affect server performance. We recommended that you use the default setting of 2048.
8. Click the **Security** tab. In **Group or user names**, click **Domain Computers**, under **Allow**, select **Enroll** and **Autoenroll**, and then click **OK**.

   Note
   If you want do not want to deploy the certificate to every computer in the domain, then specify a different group or groups that contain the computer accounts that you want to receive the certificate.
9. Close the Certificate Templates Console.
10. In the Certification Authority MMC snap-in, in the left pane, right-click **Certificate Templates**, click **New**, and then click **Certificate Template to Issue**.
11. In the **Enable Certificate Templates** dialog box, click the name of the certificate template you just configured, and then click **OK**.

Configure Windows Firewall to Suppress Notifications When a Program Is Blocked

To configure Windows Firewall to suppress the display of a notification when it blocks a program that tries to listen for network traffic and to prohibit locally defined rules, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in.

Caution

If you choose to disable alerts and prohibit locally defined rules, then you must create firewall rules that allow your users' programs to send and receive the required network traffic. If a firewall rule is missing, then the user does not receive any kind of warning, the network traffic is silently blocked, and the program might fail.

We recommend that you do not enable these settings until you have created and tested the required rules.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

In this topic:

[To configure Windows Firewall to suppress the display of a notification for a blocked program and to ignore locally defined rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2]

To configure Windows Firewall to suppress the display of a notification for a blocked program and to ignore locally defined rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane, in the **Overview** section, click **Windows Firewall Properties**.
3. For each network location type (Domain, Private, Public), perform the following steps.
   a. Click the tab that corresponds to the network location type.
   b. Under **Settings**, click **Customize**.
   c. Under **Firewall settings**, change **Display a notification** to **No**.
   d. Under **Rule merging**, change **Apply local firewall rules** to **No**.
   e. Although a connection security rule is not a firewall setting, you can also use this tab to prohibit locally defined connection security rules if you are planning to deploy IPsec rules as part of a server or domain isolation environment. Under **Rule merging**, change **Apply local connection security rules** to **No**.
   f. Click **OK** twice.

Confirm That Certificates Are Deployed Correctly

After configuring your certificates and autoenrollment in Group Policy, you can confirm that the policy is being applied as expected, and that the certificates are being properly installed on the workstation computers.

In these procedures, you refresh Group Policy on a client computer, and then confirm that the certificate is deployed correctly.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

In this topic:

- Refresh Group Policy on a computer
- Verify that a certificate is installed

To refresh Group Policy on a computer

1. On a computer running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2, Start a Command Prompt as an Administrator, and then type the following command:

   ```
   gpupdate /target:computer /force
   ```

After Group Policy is refreshed, you can see which GPOs are currently applied to the computer.

To verify that a certificate is installed

1. Click the **Start** charm, type **certmgr.msc**, and then press ENTER.
2. In the navigation pane, expand **Trusted Root Certification Authorities**, and then click **Certificates**.

   The CA that you created appears in the list.

Copy a GPO to Create a New GPO

To create the GPO for the boundary zone computers, make a copy of the main domain isolation GPO, and then change the settings to request, instead of require, authentication. To make a copy of a GPO, use the Active Directory Users and Computers MMC snap-in.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to create new GPOs.

To make a copy of a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click **Group Policy Management** tile.
2. In the navigation pane, expand **Forest:***YourForestName*, expand **Domains**, expand *YourDomainName*, and then click **Group Policy Objects**.
3. In the details pane, right-click the GPO you want to copy, and then click **Copy**.
4. In the navigation pane, right-click **Group Policy Objects** again, and then click **Paste**.
5. In the **Copy GPO** dialog box, click **Preserve the existing permissions**, and then click **OK**. Selecting this option preserves any exception groups to which you denied Read and Apply GPO permissions, making the change simpler.
6. After the copy is complete, click **OK**. The new GPO is named **Copy of** *original GPO name*.
7. To rename it, right-click the GPO, and then click **Rename**.
8. Type the new name, and then press ENTER.
9. You must change the security filters to apply the policy to the correct group of computers. To do this, click the **Scope** tab, and in the **Security Filtering** section, select the group that grants permissions to all members of the isolated domain, for example **CG_DOMISO_IsolatedDomain**, and then click **Remove**.
10. In the confirmation dialog box, click **OK**.
11. Click **Add**.
12. Type the name of the group that contains members of the boundary zone, for example **CG_DOMISO_Boundary**, and then click **OK**.
13. If required, change the WMI filter to one appropriate for the new GPO. For example, if the original GPO is for client computers running Windows 8, and the new boundary zone GPO is for computers running Windows Server 2012, then select a WMI filter that allows only those computers to read and apply the GPO.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create a Group Account in Active Directory

To create a security group to contain the computer accounts for the computers that are to receive a set of Group Policy settings, use the Active Directory Users and Computers MMC snap-in.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to create new group accounts.

To add a new membership group in Active Directory

1. On a computer that has Active Directory management tools installed, click the **Start**

charm, and then click the **Active Directory Users and Computers** tile.

2. In the navigation pane, select the container in which you want to store your group. This is typically the **Users** container under the domain.
3. Click **Action**, click **New**, and then click **Group**.
4. In the **Group name** text box, type the name for your new group.

   Note
   Be sure to use a name that clearly indicates its purpose. Check to see if your organization has a naming convention for groups.
5. In the **Description** text box, enter a description of the purpose of this group.
6. In the **Group scope** section, select either **Global** or **Universal**, depending on your Active Directory forest structure. If your group must include computers from multiple domains, then select **Universal**. If all of the members are from the same domain, then select **Global**.
7. In the **Group type** section, click **Security**.
8. Click **OK** to save your group.

Create a Group Policy Object

To create a new GPO, use the Active Directory Users and Computers MMC snap-in.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to create new GPOs.

To create a new GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, expand **Forest:***YourForestName*, expand **Domains**, expand *YourDomainName*, and then click **Group Policy Objects**.
3. Click **Action**, and then click **New**.
4. In the **Name** text box, type the name for your new GPO.

   Note
   Be sure to use a name that clearly indicates the purpose of the GPO. Check to see if your organization has a naming convention for GPOs.
5. Leave **Source Starter GPO** set to **(none)**, and then click **OK**.
6. If your GPO will not contain any user settings, then you can improve performance by disabling the **User Configuration** section of the GPO. To do this, perform these steps:
   a. In the navigation pane, click the new GPO.

b. In the details pane, click the **Details** tab.

c. Change the **GPO Status** to **User configuration settings disabled**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Authentication Exemption List Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

In almost any isolated server or isolated domain scenario, there are some computers or devices that cannot communicate by using IPsec. This procedure shows you how to create rules that exempt those computers from the authentication requirements of your isolation policies.

Important
Adding computers to the exemption list for a zone reduces security because it permits computers in the zone to send network traffic that is unprotected by IPsec to the computers on the list. As discussed in the Windows Firewall with Advanced Security Design Guide, you must add only managed and trusted computers to the exemption list.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create a rule that exempts specified hosts from authentication

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Connection Security Rules**.
3. Click **Action**, and then click **New Rule**.
4. On the **Rule Type** page of the New Connection Security Rule Wizard, click **Authentication exemption**, and then click **Next**.
5. On the **Exempt Computers** page, to create a new exemption, click **Add**. To modify an existing exemption, click it, and then click **Edit**.
6. In the **IP Address** dialog box, do one of the following:
   - To add a single IP address, click **This IP address or subnet**, type the IP address of the host in the text box, and then click **OK**.
   - To add an entire subnet by address, click **This IP address or subnet**, and then type the IP address of the subnet, followed by a forward slash (/) and the number of bits in the corresponding subnet mask. For example, **10.50.0.0/16** represents the class B subnet that begins with address 10.50.0.1, and ends with address **10.50.255.254**.

Click **OK** when you are finished.

- To add the local computer's subnet, click **Predefined set of computers**, select **Local subnet** from the list, and then click **OK**.

  Note
  If you select the local subnet from the list rather than typing the subnet address in manually, the computer automatically adjusts the active local subnet to match the computer's current IP address.

- To add a discrete range of addresses that do not correspond to a subnet, click **This IP address range**, type the beginning and ending IP addresses in the **From** and **To** text boxes, and then click **OK**.
- To exempt all of the remote hosts that the local computer uses for a specified network service, click **Predefined set of computers**, select the network service from the list, and then click **OK**.

7. Repeat steps 5 and 6 for each exemption that you need to create.
8. Click **Next** when you have created all of the exemptions.
9. On the **Profile** page, check the profile for each network location type to which this set of exemptions applies, and then click **Next**.

   Caution
   If all of the exemptions are on the organization's network and that network is managed by an Active Directory domain, then consider restricting the rule to the Domain profile only. Selecting the wrong profile can reduce the protection for your computer because any computer with an IP address that matches an exemption rule will not be required to authenticate.

10. On the **Name** page, type the name of the exemption rule, type a description, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Authentication Request Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

After you have configured IPsec algorithms and authentication methods, you can create the rule that requires the computers on the network to use those protocols and methods before they can communicate.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create the authentication request rule

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, right-click **Connection Security Rules**, and then click **New Rule**.
3. On the **Rule Type** page, select **Isolation**, and then click **Next**.
4. On the **Requirements** page, select **Request authentication for inbound and outbound connections**.

   Caution
   Do not configure the rule to require inbound authentication until you have confirmed that all of your computers are receiving the correct GPOs, and are successfully negotiating IPsec and authenticating with each other. Allowing the computers to communicate even when authentication fails prevents any errors in the GPOs or their distribution from breaking communications on your network.
5. On the **Authentication Method** page, select the authentication option you want to use on your network. To select multiple methods that are tried in order until one succeeds, click **Advanced**, click **Customize**, and then click **Add** to add methods to the list. Second authentication methods require Authenticated IP (AuthIP), which is supported only on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2.
   a. **Default**. Selecting this option tells the computer to request authentication by using the method currently defined as the default on the computer. This default might have been configured when the operating system was installed or it might have been configured by Group Policy. Selecting this option is appropriate when you have configured system-wide settings by using the Configure Authentication Methods on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 procedure.
   b. **Computer and User (Kerberos V5)**. Selecting this option tells the computer to request authentication of both the computer and the currently logged-on user by using their domain credentials. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
   c. **Computer (Kerberos V5)**. Selecting this option tells the computer to request authentication of the computer by using its domain credentials. This option works with other computers than can use IKE v1, including earlier versions of Windows.
   d. **Advanced**. Click **Customize** to specify a custom combination of authentication methods required for your scenario. You can specify both a **First authentication method** and a **Second authentication method**.

The **First authentication method** can be one of the following:

- **Computer (Kerberos V5)**. Selecting this option tells the computer to request authentication of the computer by using its domain credentials. This option works with other computers than can use IKE v1, including earlier versions of Windows.
- **Computer (NTLMv2)**. Selecting this option tells the computer to use and require authentication of the computer by using its domain credentials. This option works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
- **Computer certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to request authentication by using a certificate that is issued by the specified CA. If you also select **Accept only health certificates**, then only certificates issued by a NAP server can be used for this rule.
- **Preshared key (not recommended)**. Selecting this method and entering a pre-shared key tells the computer to authenticate by exchanging the pre-shared keys. If the keys match, then the authentication succeeds. This method is not recommended, and is included for backward compatibility and testing purposes only.

If you select **First authentication is optional**, then the connection can succeed even if the authentication attempt specified in this column fails.

The **Second authentication method** can be one of the following:

- **User (Kerberos V5)**. Selecting this option tells the computer to use and require authentication of the currently logged-on user by using his or her domain credentials. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using Kerberos V5 is not supported by IKE v1.
- **User (NTLMv2)**. Selecting this option tells the computer to use and require authentication of the currently logged-on user by using his or her domain credentials, and uses the NTLMv2 protocol instead of Kerberos V5. This authentication method works only with other computers that can use AuthIP, including Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. User-based authentication using NTLMv2 is not supported by IKE v1.
- **User health certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to request user-based authentication by using a certificate that is issued by the specified CA. If you also select **Enable certificate to account mapping**, then the certificate can

be associated with a user in Active Directory for purposes of granting or denying access to certain users or user groups.

- **Computer health certificate from this certification authority (CA)**. Selecting this option and entering the identification of a CA tells the computer to use and require authentication by using a certificate that is issued by the specified CA. If you also select **Accept only health certificates**, then only certificates issued by a NAP server can be used for this rule.

If you check **Second authentication is optional**, the connection can succeed even if the authentication attempt specified in this column fails.

Important
Make sure that you do not select the boxes to make both first and second authentication optional. Doing so allows plaintext connections whenever authentication fails.

6. After you have configured the authentication methods, click **OK** on each dialog box to save your changes and close it, until you return to the **Authentication Method** page in the wizard. Click **Next**.
7. On the **Profile** page, select the check boxes for the network location type profiles to which this rule applies.
   - On portable computers, consider clearing the **Private** and **Public** boxes to enable the computer to communicate without authentication when it is away from the domain network.
   - On computers that do not move from network to network, consider selecting all of the profiles. Doing so prevents an unexpected switch in the network location type from disabling the rule.

   Click **Next**.
8. On the **Name** page, type a name for the connection security rule and a description, and then click **Finish**.

   The new rule appears in the list of connection security rules.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To allow inbound Internet Control Message Protocol (ICMP) network traffic, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create

firewall rules. This type of rule allows ICMP requests and responses to be sent and received by computers on the network.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

This topic describes how to create a port rule that allows inbound ICMP network traffic. For other inbound port rule types, see:

- Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
- Create Inbound Rules to Support RPC on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To create an inbound ICMP rule

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Inbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Custom**, and then click **Next**.
5. On the **Program** page, click **All programs**, and then click **Next**.
6. On the **Protocol and Ports** page, select **ICMPv4** or **ICMPv6** from the **Protocol type** list. If you use both IPv4 and IPv6 on your network, you must create a separate ICMP rule for each.
7. Click **Customize**.
8. In the **Customize ICMP Settings** dialog box, do one of the following:
   - To allow all ICMP network traffic, click **All ICMP types**, and then click **OK**.
   - To select one of the predefined ICMP types, click **Specific ICMP types**, and then select each type in the list that you want to allow. Click **OK**.
   - To select an ICMP type that does not appear in the list, click **Specific ICMP types**, select the **Type** number from the list, select the **Code** number from the list, click **Add**, and then select the newly created entry from the list. Click **OK**
9. Click **Next**.
10. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
11. On the **Action** page, select **Allow the connection**, and then click **Next**.
12. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

    Note
    If this GPO is targeted at server computers running Windows Server 2008 that

never move, consider modifying the rules to apply to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

13. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To allow inbound network traffic on only a specified TCP or UDP port number, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create firewall rules. This type of rule allows any program that listens on a specified TCP or UDP port to receive network traffic sent to that port.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

This topic describes how to create a standard port rule for a specified protocol or TCP or UDP port number. For other inbound port rule types, see:

- Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
- Create Inbound Rules to Support RPC on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To create an inbound port rule

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Inbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Custom**, and then click **Next**.

   Note
   Although you can create rules by selecting **Program** or **Port**, those choices limit

the number of pages presented by the wizard. If you select **Custom**, you see all of the pages, and have the most flexibility in creating your rules.

5. On the **Program** page, click **All programs**, and then click **Next**.

   Note
   This type of rule is often combined with a program or service rule. If you combine the rule types, you get a firewall rule that limits traffic to a specified port and allows the traffic only when the specified program is running. The specified program cannot receive network traffic on other ports, and other programs cannot receive network traffic on the specified port. If you choose to do this, follow the steps in the Create an Inbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 procedure in addition to the steps in this procedure to create a single rule that filters network traffic using both program and port criteria.

6. On the **Protocol and Ports** page, select the protocol type that you want to allow. To restrict the rule to a specified port number, you must select either **TCP** or **UDP**. Because this is an incoming rule, you typically configure only the local port number.

   If you select another protocol, then only packets whose protocol field in the IP header match this rule are permitted through the firewall.

   To select a protocol by its number, select **Custom** from the list, and then type the number in the **Protocol number** box.

   When you have configured the protocols and ports, click **Next**.

7. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
8. On the **Action** page, select **Allow the connection**, and then click **Next**.
9. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

   Note
   If this GPO is targeted at server computers running Windows Server 2008 that never move, consider modifying the rules to apply to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

10. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Inbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To allow inbound network traffic to a specified program or service, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create firewall rules. This type of rule allows the program to listen and receive inbound network traffic on any port.

Note

This type of rule is often combined with a program or service rule. If you combine the rule types, you get a firewall rule that limits traffic to a specified port and allows the traffic only when the specified program is running. The program cannot receive network traffic on other ports, and other programs cannot receive network traffic on the specified port. To combine the program and port rule types into a single rule, follow the steps in the Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2 procedure in addition to the steps in this procedure.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create an inbound firewall rule for a program or service

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Inbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Custom**, and then click **Next**.

   Note

   Although you can create rules by selecting **Program** or **Port**, those choices limit the number of pages presented by the wizard. If you select **Custom**, you see all of the pages, and have the most flexibility in creating your rules.
5. On the **Program** page, click **This program path**.
6. Type the path to the program in the text box. Use environment variables, where applicable, to ensure that programs installed in different locations on different computers work correctly.
7. Do one of the following:
   - If the executable file contains a single program, click **Next**.

- If the executable file is a container for multiple services that must all be allowed to receive inbound network traffic, click **Customize**, select **Apply to services only**, click **OK**, and then click **Next**.
- If the executable file is a container for a single service or contains multiple services but the rule only applies to one of them, click **Customize**, select **Apply to this service**, and then select the service from the list. If the service does not appear in the list, click **Apply to service with this service short name**, and then type the short name for the service in the text box. Click **OK**, and then click **Next**.

Important

To use the **Apply to this service** or **Apply to service with this service short name** options, the service must be configured with a security identifier (SID) with a type of **RESTRICTED** or **UNRESTRICTED**. To check the SID type of a service, run the following command:

sc qsidtype <*ServiceName*>

If the result is **NONE**, then a firewall rule cannot be applied to that service.

To set a SID type on a service, run the following command:

sc sidtype <*Type*> <*ServiceName*>

In the preceding command, the value of <*Type*> can be **UNRESTRICTED** or **RESTRICTED**. Although the command also permits the value of **NONE**, that setting means the service cannot be used in a firewall rule as described here. By default, most services in Windows are configured as **UNRESTRICTED**. If you change the SID type to **RESTRICTED**, the service might fail to start. We recommend that you change the SID type only on services that you want to use in firewall rules, and that you change the SID type to **UNRESTRICTED**. For more information, see Vista Services (http://go.microsoft.com/fwlink/?linkid=141454) and the "Service Security Improvements" section of Inside the Windows Vista Kernel (http://go.microsoft.com/fwlink/?linkid=141455).

8. It is a best practice to restrict the firewall rule for the program to only the ports it needs to operate. On the **Protocols and Ports** page, you can specify the port numbers for the allowed traffic. If the program tries to listen on a port different from the one specified here, it is blocked. For more information about protocol and port options, see Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2. After you have configured the protocol and port options, click **Next**.
9. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
10. On the **Action** page, select **Allow the connection**, and then click **Next**.
11. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

Note

If this GPO is targeted at server computers running Windows Server 2008 that never move, consider applying the rule to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

12. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Outbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

By default, Windows Firewall with Advanced Security allows all outbound network traffic unless it matches a rule that prohibits the traffic. To block outbound network traffic on a specified TCP or UDP port number, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create firewall rules. This type of rule blocks any outbound network traffic that matches the specified TCP or UDP port numbers.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create an outbound port rule

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Outbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Outbound Rule wizard, click **Custom**, and then click **Next**.

   Note
   Although you can create rules by selecting **Program** or **Port**, those choices limit the number of pages presented by the wizard. If you select **Custom**, you see all of the pages, and have the most flexibility in creating your rules.
5. On the **Program** page, click **All programs**, and then click **Next**.
6. On the **Protocol and Ports** page, select the protocol type that you want to block. To

restrict the rule to a specified port number, you must select either **TCP** or **UDP**. Because this is an outbound rule, you typically configure only the remote port number.

If you select another protocol, then only packets whose protocol field in the IP header match this rule are blocked by Windows Firewall. Network traffic for protocols is allowed as long as other rules that match do not block it.

To select a protocol by its number, select **Custom** from the list, and then type the number in the **Protocol number** box.

When you have configured the protocols and ports, click **Next**.

7. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
8. On the **Action** page, select **Block the connection**, and then click **Next**.
9. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

   Note

   If this GPO is targeted at server computers running Windows Server 2008 that never move, consider applying the rules to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.
10. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create an Outbound Program or Service Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2

By default, Windows Firewall with Advanced Security allows all outbound network traffic unless it matches a rule that prohibits the traffic. To block outbound network traffic for a specified program or service, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create firewall rules. This type of rule prevents the program from sending any outbound network traffic on any port.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create an outbound firewall rule for a program or service

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Outbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Outbound Rule Wizard, click **Custom**, and then click **Next**.

   Note
   Although you can create many rules by selecting **Program** or **Port**, those choices limit the number of pages presented by the wizard. If you select **Custom**, you see all of the pages, and have the most flexibility in creating your rules.
5. On the **Program** page, click **This program path**.
6. Type the path to the program in the text box. Use environment variables as appropriate to ensure that programs installed in different locations on different computers work correctly.
7. Do one of the following:
   - If the executable file contains a single program, click **Next**.
   - If the executable file is a container for multiple services that must all be blocked from sending outbound network traffic, click **Customize**, select **Apply to services only**, click **OK**, and then click **Next**.
   - If the executable file is a container for a single service or contains multiple services but the rule only applies to one of them, click **Customize**, select **Apply to this service**, and then select the service from the list. If the service does not appear in the list, then click **Apply to service with this service short name**, and type the short name for the service in the text box. Click **OK**, and then click **Next**.
8. If you want the program to be allowed to send on some ports, but blocked from sending on others, then you can restrict the firewall rule to block only the specified ports or protocols. On the **Protocols and Ports** page, you can specify the port numbers or protocol numbers for the blocked traffic. If the program tries to send to or from a port number different from the one specified here, or by using a protocol number different from the one specified here, then the default outbound firewall behavior allows the traffic. For more information about the protocol and port options, see Create an Outbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2. When you have configured the protocol and port options, click **Next**.
9. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
10. On the **Action** page, select **Block the connection**, and then click **Next**.
11. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

Note

If this GPO is targeted at server computers running Windows Server 2008 that never move, consider modifying the rules to apply to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

12. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create Inbound Rules to Support RPC on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To allow inbound remote procedure call (RPC) network traffic, use the Windows Firewall with Advanced Security node in the Group Policy Management MMC snap-in to create two firewall rules. The first rule allows incoming network packets on TCP port 135 to the RPC Endpoint Mapper service. The incoming traffic consists of requests to communicate with a specified network service. The RPC Endpoint Mapper replies with a dynamically-assigned port number that the client must use to communicate with the service. The second rule allows the network traffic that is sent to the dynamically-assigned port number. Using the two rules configured as described in this topic helps to protect your computer by allowing network traffic only from computers that have received RPC dynamic port redirection and to only those TCP port numbers assigned by the RPC Endpoint Mapper.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

This topic describes how to create rules that allow inbound RPC network traffic. For other inbound port rule types, see:

- Create an Inbound Port Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2
- Create an Inbound ICMP Rule on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

In this topic:

- To create a rule to allow inbound network traffic to the RPC Endpoint Mapper service
- To create a rule to allow inbound network traffic to RPC-enabled network services

To create a rule to allow inbound network traffic to the RPC Endpoint Mapper service

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Inbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Custom**, and then click **Next**.
5. On the **Program** page, click **This Program Path**, and then type **%systemroot%\system32\svchost.exe**.
6. Click **Customize**.
7. In the **Customize Service Settings** dialog box, click **Apply to this service**, select **Remote Procedure Call (RPC)** with a short name of **RpcSs**, click **OK**, and then click **Next**.
8. On the warning about Windows service-hardening rules, click **Yes**.
9. On the **Protocol and Ports** dialog box, for **Protocol type**, select **TCP**.
10. For **Local port**, select **RPC Endpoint Mapper**, and then click **Next**.
11. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
12. On the **Action** page, select **Allow the connection**, and then click **Next**.
13. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

    Note
    If this GPO is targeted at server computers running Windows Server 2008 that never move, consider applying the rules to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.
14. On the **Name** page, type a name and description for your rule, and then click **Finish**.

To create a rule to allow inbound network traffic to RPC-enabled network services

1. On the same GPO you edited in the preceding procedure, click **Action**, and then click **New rule**.
2. On the **Rule Type** page of the New Inbound Rule Wizard, click **Custom**, and then click **Next**.
3. On the **Program** page, click **This Program Path**, and then type the path to the executable file that hosts the network service. Click **Customize**.
4. In the **Customize Service Settings** dialog box, click **Apply to this service**, and then

select the service that you want to allow. If the service does not appear in the list, then click **Apply to service with this service short name**, and then type the short name of the service in the text box.

5. Click **OK**, and then click **Next**.
6. On the **Protocol and Ports** dialog box, for **Protocol type**, select **TCP**.
7. For **Local port**, select **RPC Dynamic Ports**, and then click **Next**.
8. On the **Scope** page, you can specify that the rule applies only to network traffic to or from the IP addresses entered on this page. Configure as appropriate for your design, and then click **Next**.
9. On the **Action** page, select **Allow the connection**, and then click **Next**.
10. On the **Profile** page, select the network location types to which this rule applies, and then click **Next**.

    Note
    If this GPO is targeted at server computers running Windows Server 2008 that never move, consider applying the rules to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.
11. On the **Name** page, type a name and description for your rule, and then click **Finish**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Create WMI Filters for the GPO

To make sure that each GPO associated with a group can only be applied to computers running the correct version of Windows, use the Group Policy Management MMC snap-in to create and assign WMI filters to the GPO. Although you can create a separate membership group for each GPO, you would then have to manage the memberships of the different groups. Instead, use only a single membership group, and let WMI filters automatically ensure the correct GPO is applied to each computer.

- To create a WMI filter that queries for a specified version of Windows
- To link a WMI filter to a GPO

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

First, create the WMI filter and configure it to look for a specified version (or versions) of the Windows operating system.

To create a WMI filter that queries for a specified version of Windows

1. On a computer that has the Group Policy Management feature installed, click **Start**, click **Administrative Tools**, and then click **Group Policy Management**.
2. In the navigation pane, expand **Forest:** *YourForestName*, expand **Domains**, expand *YourDomainName*, and then click **WMI Filters**.
3. Click **Action**, and then click **New**.
4. In the **Name** text box, type the name of the WMI filter.

   Note
   Be sure to use a name that clearly indicates the purpose of the filter. Check to see if your organization has a naming convention.
5. In the **Description** text box, type a description for the WMI filter. For example, if the filter excludes domain controllers, you might consider stating that in the description.
6. Click **Add**.
7. Leave the **Namespace** value set to **root\CIMv2**.
8. In the **Query** text box, type:

   ```
   select * from Win32_OperatingSystem where Version like "6.%"
   ```

   This query will return **true** for computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. To set a filter for just Windows 8 and Windows Server 2012, use `"6.2%"`. To specify multiple versions, combine them with `or`, as shown in the following:

   ```
   ... where Version like "6.1%" or Version like "6.2%"
   ```

   To restrict the query to only clients or only servers, add a clause that includes the `ProductType` parameter. To filter for client operating systems only, such as Windows 8 or Windows 7, use only `ProductType="1"`. For server operating systems that are not domain controllers, use `ProductType="3"`. For domain controllers only, use `ProductType="2"`. This is a useful distinction, because you often want to prevent your GPOs from being applied to the domain controllers on your network.

   The following clause returns **true** for all computers that are not domain controllers:

   ```
   ... where ProductType="1" or ProductType="3"
   ```

   The following complete query returns **true** for all computers running Windows 8, and returns **false** for any server operating system or any other client operating system.

   ```
   select * from Win32_OperatingSystem where Version like "6.2%"
   and ProductType="1"
   ```

   The following query returns **true** for any computer running Windows Server 2012, except domain controllers:

   ```
   select * from Win32_OperatingSystem where Version like "6.2%"
   and ProductType="3"
   ```
9. Click **OK** to save the query to the filter.

10. Click **Save** to save your completed filter.

After you have created a filter with the correct query, link the filter to the GPO. Filters can be reused with many GPOs simultaneously; you do not have to create a new one for each GPO if an existing one meets your needs.

To link a WMI filter to a GPO

1. On a computer that has the Group Policy Management feature installed, click **Start**, click **Administrative Tools**, and then click **Group Policy Management**.
2. In the navigation pane, find and then click the GPO that you want to modify.
3. Under **WMI Filtering**, select the correct WMI filter from the list.
4. Click **Yes** to accept the filter.

Enable Predefined Inbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

Windows Firewall with Advanced Security includes many predefined rules for common networking roles and functions. When you install a new server role on a computer or enable a network feature on a client computer, the installer typically enables the rules required for that role instead of creating new ones. When deploying firewall rules to the computers on the network, you can take advantage of these predefined rules instead of creating new ones. Doing this helps to ensure consistency and accuracy, because the rules have been thoroughly tested and are ready for use.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To deploy predefined firewall rules that allow inbound network traffic for common network functions

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Inbound Rules**.
3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Predefined**, select the rule category from the list, and then click **Next**.

5. On the **Predefined Rules** page, the list of rules defined in the group is displayed. By default, they are all selected. For rules that you do not want to deploy, clear the check boxes next to the rules, and then click **Next**.
6. On the **Action** page, select **Allow the connection**, and then click **Finish**.

   The selected rules are added to the GPO and applied to the computers to which the GPO is assigned the next time Group Policy is refreshed.

   Note

   If this GPO is targeted at server computers running Windows Server 2008 that never move, consider modifying the rules to apply to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Enable Predefined Outbound Rules on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

By default, Windows Firewall with Advanced Security allows all outbound network traffic unless it matches a rule that prohibits the traffic. Windows Firewall with Advanced Security includes many predefined outbound rules that can be used to block network traffic for common networking roles and functions. When you install a new server role on a computer or enable a network feature on a client computer, the installer can install, but typically does not enable, outbound block rules for that role. When deploying firewall rules to the computers on the network, you can take advantage of these predefined rules instead of creating new ones. Doing this helps to ensure consistency and accuracy, because the rules have been thoroughly tested and are ready for use.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To deploy predefined firewall rules that block outbound network traffic for common network functions

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the navigation pane, click **Outbound Rules**.

3. Click **Action**, and then click **New rule**.
4. On the **Rule Type** page of the New Inbound Rule Wizard, click **Predefined**, select the rule category from the list, and then click **Next**.
5. On the **Predefined Rules** page, the list of rules defined in the group is displayed. They are all selected by default. For rules that you do not want to deploy, clear the check boxes next to the rules, and then click **Next**.
6. On the **Action** page, select **Block the connection**, and then click **Finish**.

   The selected rules are added to the GPO.

   Note

   If this GPO is targeted at server computers running Windows Server 2008 that never move, consider modifying the rules to apply to all network location type profiles. This prevents an unexpected change in the applied rules if the network location type changes due to the installation of a new network card or the disconnection of an existing network card's cable. A disconnected network card is automatically assigned to the Public network location type.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Exempt ICMP from Authentication on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2

This procedure shows you how to add exemptions for any network traffic that uses the ICMP protocol.

Important

Because of its usefulness in troubleshooting network connectivity problems, we recommend that you exempt all ICMP network traffic from authentication requirements unless your network risk analysis indicates a need to protect this traffic.

Administrative credentials

To complete this procedure, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To exempt ICMP network traffic from authentication

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. On the main Windows Firewall with Advanced Security page, click **Windows Firewall**

Properties.

3. On the **IPsec settings** tab, change **Exempt ICMP from IPsec** to **Yes**, and then click **OK**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Install Active Directory Certificate Services

To use certificates in a server isolation or domain isolation design, you must first set up the infrastructure to deploy the certificates. This is called a public key infrastructure (PKI). The services required for a PKI are available in Windows Server 2012 in the form of the Active Directory Certificate Services (AD CS) role.

Caution

Creation of a full PKI for an enterprise environment with all of the appropriate security considerations included in the design is beyond the scope of this guide. The following procedure shows you only the basics of installing an issuing certificate server; it is appropriate for a test lab environment only. For more information about deploying AD CS in a production environment, see **Active Directory Certificate Services Overview** in the Windows Server 2012 Technical Library (http://technet.microsoft.com/library/hh831740.aspx).

To perform this procedure, the computer on which you are installing AD CS must be joined to an Active Directory domain.

Administrative credentials

To complete this procedure, you must be a member of both the Domain Admins group in the root domain of your forest, and a member of the Enterprise Admins group.

To install AD CS

1. Log on as a member of both the Enterprise Admins group and the root domain's Domain Admins group.
2. Click **Server Manager** in the taskbar. The Server Manager console opens. Click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, ensure **Role-based or feature-based installation** is selected and click **Next**.
5. On the **Select destination server** page, ensure your server is selected and click **Next**.
6. On the **Select Server Roles** page, select **Active Directory Certificate Services**, and then click **Add Features** and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Certificate Services** page, click **Next**.
9. On the **Select role services** page, ensure **Certification Authority** is selected and click

Next.

10. On the **Confirm installation selections** page, click **Install**.

    After installation completes, click close.

11. On the Server Manager Dashboard, click the Notifications flag icon and then click **Configure Active Directory Certificate Services on the destination server**.
12. On the **Credentials** page, ensure the default user account is a member of both the local Administrators group and the Enterprise Admins group and then click **Next**.
13. On the **Role Services** page, click **Certification Authority**, and click **Next**.
14. On the **Setup Type** page, ensure **Enterprise CA** is selected, and click **Next**.
15. On the **CA Type** page, ensure **Root CA** is selected, and then click **Next**.
16. On the **Private Key** page, ensure **Create a new private key** is selected, and then click **Next**.
17. On the **Cryptography for CA** page, keep the default settings for CSP (**RSA#Microsoft Software Key Storage Provider**) and hash algorithm (**sha1**), and determine the best key character length for your deployment. Large key character lengths provide optimal security, but they can affect server performance. It is recommended that you keep the default setting of 2048 or, if appropriate for your deployment, reduce key character length to 1024. Click **Next**.
18. On the **CA Name** page, keep the suggested common name for the CA or change the name according to your requirements, and then click **Next**.
19. On the **Validity Period** page, in **Specify the validity period**, type the number and select a time value (Years, Months, Weeks, or Days). The default setting of five years is recommended. Click **Next**.
20. On the **CA Database** page, in **Certificate database location** and **Certificate database log location**, specify the folder location for these items. If you specify locations other than the default locations, make sure that the folders are secured with access control lists (ACLs) that prevent unauthorized users or computers from accessing the CA database and log files.
21. Click **Next**, click **Configure**, and then click **Close**.

Link the GPO to the Domain

After you create the GPO and configure it with security group filters and WMI filters, you must link the GPO to the container in Active Directory that contains all of the target computers.

If the filters comprehensively control the application of the GPO to only the correct computers, then you can link the GPO to the domain container. Alternatively, you can link the GPO to a site container or organizational unit if you want to limit application of the GPO to that subset of computers.

Administrative credentials

To complete this procedure, you must be a member of the Domain Admins group, or otherwise be delegated permissions to modify the GPOs.

To link the GPO to the domain container in Active Directory

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, expand **Forest:** *YourForestName*, expand **Domains**, and then expand *YourDomainName*.
3. Right-click *YourDomainName*, and then click **Link an Existing GPO**.
4. In the **Select GPO** dialog box, select the GPO that you want to deploy, and then click **OK**.
5. The GPO appears in the **Linked Group Policy Objects** tab in the details pane and as a linked item under the domain container in the navigation pane.
6. You can adjust the order of the linked GPOs to ensure that the higher priority GPOs are processed last. Select a GPO and click the up or down arrows to move it. The GPOs are processed by the client computer from the highest link order number to the lowest.

Modify GPO Filters to Apply to a Different Zone or Version of Windows

You must reconfigure your copied GPO so that it contains the correct security group and WMI filters for its new role. If you are creating the GPO for the isolated domain, use the Block members of a group from applying a GPO procedure to prevent members of the boundary and encryption zones from incorrectly applying the GPOs for the main isolated domain.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

In this topic:

- Change the security group filter for a GPO
- Block members of a group from applying a GPO
- Remove a block for members of a group from applying a GPO

Use the following procedure to change a group to the security filter on the GPO that allows group members to apply the GPO. You must remove the reference to the original group, and add the group appropriate for this GPO.

To change the security group filter for a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start**

charm, and then click the **Group Policy Management** tile.

2. In the navigation pane, find and then click the GPO that you want to modify.
3. In the details pane, under **Security Filtering**, click the currently assigned security group, and then click **Remove**.
4. Now you can add the appropriate security group to this GPO. Under **Security Filtering**, click **Add**.
5. In the **Select User, Computer, or Group** dialog box, type the name of the group whose members are to apply the GPO, and then click **OK**. If you do not know the name, you can click **Advanced** to browse the list of groups available in the domain.

Use the following procedure if you need to add a group to the security filter on the GPO that blocks group members from applying the GPO. This can be used on the GPOs for the main isolated domain to prevent members of the boundary and encryption zones from incorrectly applying the GPOs for the main isolated domain.

To block members of group from applying a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, find and then click the GPO that you want to modify.
3. In the details pane, click the **Delegation** tab.
4. Click **Advanced**.
5. Under the **Group or user names** list, click **Add**.
6. In the **Select User, Computer, or Group** dialog box, type the name of the group whose members are to be prevented from applying the GPO, and then click **OK**. If you do not know the name, you can click **Advanced** to browse the list of groups available in the domain.
7. Select the group in the **Group or user names** list, and then select the boxes in the **Deny** column for both **Read** and **Apply group policy**.
8. Click **OK**, and then in the **Windows Security** dialog box, click **Yes**.
9. The group appears in the list with custom permissions.

To remove a block for members of group from applying a GPO

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, find and then click the GPO that you want to modify.
3. In the details pane, click the **Delegation** tab.
4. In the **Groups and users** list, select the group that should no longer be blocked, and then click **Remove**.
5. In the message box, click **OK**.

If you arrived at this page by clicking a link in a checklist, use your browser's **Back** button to return to the checklist.

Open the Group Policy Management Console to IP Security Policies

Procedures in this guide that refer to GPOs for earlier versions of the Windows operating system instruct you to work with the IP Security Policy section in the Group Policy Management Console (GPMC).

To open a GPO to the IP Security Policies section

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, expand **Forest:** *YourForestName*, expand **Domains**, expand *YourDomainName*, expand **Group Policy Objects**, right-click the GPO you want to modify, and then click **Edit**.
3. In the navigation pane of the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click **IP Security Policies on Active Directory (***YourDomainName***)**.

Open the Group Policy Management Console to Windows Firewall

To open a GPO to Windows Firewall

1. Open **Active Directory Users and Computers**.
2. In the navigation pane, expand *YourDomainName*, right-click the container that your GPO is linked to, and then click **Properties**.
3. Click the **Group Policy** tab, select your GPO, and then click **Edit**.
4. In the navigation pane of the Group Policy Object Editor, expand **Computer Configuration**, expand **Administrative Templates**, expand **Network**, expand **Network Connections**, and then expand **Windows Firewall**.

Open the Group Policy Management Console to Windows Firewall with Advanced Security

Most of the procedures in this guide instruct you to use Group Policy settings for Windows Firewall with Advanced Security.

To open a GPO to Windows Firewall with Advanced Security

1. On a computer that has the Group Policy Management feature installed, click the **Start** charm, and then click the **Group Policy Management** tile.
2. In the navigation pane, expand **Forest:** *YourForestName*, expand **Domains**, expand *YourDomainName*, expand **Group Policy Objects**, right-click the GPO you want to modify, and then click **Edit**.
3. In the navigation pane of the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Windows Firewall with Advanced Security**, and then expand **Windows Firewall with Advanced Security - LDAP://cn={***GUID***},cn=…**.

Open Windows Firewall with Advanced Security

This procedure shows you how to open the Windows Firewall with Advanced Security MMC snap-in.

Administrative credentials

To complete this procedure, you must be a member of the Administrators group. For more information, see Additional considerations.

Opening Windows Firewall with Advanced Security

- Using the Windows interface
- Using a command line

To open Windows Firewall with Advanced Security by using the Windows interface

- Click the **Start** charm, right-click the Start page, click **All Apps**, and then click the **Windows Firewall with Advanced Security** tile.

To open Windows Firewall with Advanced Security from a command prompt

1. Open a command prompt window.
2. At the command prompt, type:

```
wf.msc
```

Additional considerations

Although standard users can start the Windows Firewall with Advanced Security MMC snap-in, to change most settings the user must be a member of a group with the permissions to modify those settings, such as Administrators.

Restrict Server Access to Members of a Group Only

After you have configured the IPsec connection security rules that force client computers to authenticate their connections to the isolated server, you must configure the rules that restrict access to only those computers or users who have been identified through the authentication process as members of the isolated server's access group.

The way in which you restrict access to the isolated server depends on which version of the Windows operating system the server is running.

- If the server is running Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012, then you create a firewall rule that specifies the user and computer accounts that are allowed. The authentication method used in the connection must support the account type specified. Remember that only Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2 support user-based authentication.

In this topic:

- Create a firewall rule to access isolated servers running Windows Server 2008 or later

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To create a firewall rule that grants access to an isolated server running Windows Server 2008 or later

1. **Open the Group Policy Management Console to Windows Firewall with Advanced Security**. You must edit the GPO that applies settings to servers in the isolated server zone.
2. In the navigation pane, right-click **Inbound Rules**, and then click **New Rule**.
3. On the **Rule Type** page, click **Custom**, and then click **Next**.
4. If you must restrict access to a single network program, then you can select **This program path**, and specify the program or service to which to grant access. Otherwise, click **All programs**, and then click **Next**.
5. If you must restrict access to only some TCP or UDP port numbers, then enter the port numbers on the **Protocol and Ports** page. Otherwise, set **Protocol type** to **Any**, and then click **Next**.
6. On the **Scope** page, select **Any IP address** for both local and remote addresses, and then click **Next**.

7. On the **Action** page, click **Allow the connection if it is secure**. If required by your design, you can also click **Customize** and select **Require the connections to be encrypted**. Click **Next**.
8. On the **Users and Computers** page, select the check box for the type of accounts (computer or user) you want to allow, click **Add**, and then enter the group account that contains the computer and user accounts permitted to access the server.

   Caution
   Remember that if you specify a user group on the Users page, your authentication scheme must include a method that uses user-based credentials. User-based credentials are only supported on versions of Windows that support AuthIP, such as Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2. Earlier versions of Windows and other operating systems that support IKE v1 only do not support user-based authentication; computers running those versions or other operating systems will not be able to connect to the isolated server through this firewall rule.

Start a Command Prompt as an Administrator

This topic describes how to open a command prompt with full administrator permissions. If your user account is a member of the Administrators group, but is not the Administrator account itself, then, by default, the programs that you run only have standard user permissions. You must explicitly specify that you require the use of your administrative permissions by using one of the procedures in this topic.

Administrative credentials

To complete these procedures, you must be a member of the Administrators group.

To start a command prompt as an administrator

- Right-click the **Start** charm, and then click **Command Prompt (Admin)**.

To start a command prompt as an administrator (alternative method)

1. Click the **Start** charm.
2. Type **cmd**, right-click the **Command Prompt** tile, and then click **Run as administrator**.

Turn on Windows Firewall and Configure Default Behavior

To enable Windows Firewall and configure its default behavior, use the Windows Firewall with Advanced Security node (for Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, and Windows Server 2008 R2) in the Group Policy Management MMC snap-in.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

To enable Windows Firewall and configure the default behavior on Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

1. Open the Group Policy Management Console to Windows Firewall with Advanced Security.
2. In the details pane, in the **Overview** section, click **Windows Firewall Properties**.
3. For each network location type (Domain, Private, Public), perform the following steps.

   Note
   The steps shown here indicate the recommended values for a typical deployment. Use the settings that are appropriate for your firewall design.

   a. Click the tab that corresponds to the network location type.
   b. Change **Firewall state** to **On (recommended)**.
   c. Change **Inbound connections** to **Block (default)**.
   d. Change **Outbound connections** to **Allow (default)**.

Verify That Network Traffic Is Authenticated

After you have configured your domain isolation rule to request, rather than require, authentication, you must confirm that the network traffic sent by the computers on the network is being protected by IPsec authentication as expected. If you switch your rules to require authentication before all of the computers have received and applied the correct GPOs, or if there are any errors in your rules, then communications on the network can fail. By first setting the rules to request authentication, any network connections that fail authentication can continue in clear text while you diagnose and troubleshoot.

In these procedures, you confirm that the rules you deployed are working correctly. Your next steps depend on which zone you are working on:

- **Main domain isolation zone.** Before you convert your main domain isolation IPsec rule from request mode to require mode, you must make sure that the network traffic is protected according to your design. By configuring your rules to request and not require authentication at the beginning of operations, computers on the network can continue to communicate even when the main mode authentication or quick mode integrity and encryption rules are not working correctly. For example, if your encryption zone contains rules that require a certain encryption algorithm, but that algorithm is not included in a security method combination on the clients, then those clients cannot successfully negotiate a quick mode security association, and the server refuses to accept network traffic from the client. By first using request mode only, you have the opportunity to deploy your rules and then examine the network traffic to see if they are working as expected without risking a loss of communications.
- **Boundary zone.** Confirming correct operation of IPsec is the last step if you are working on the boundary zone GPO. You do not convert the GPO to require mode at any time.
- **Encryption zone.** Similar to the main isolation zone, after you confirm that the network traffic to zone members is properly authenticated and encrypted, you must convert your zone rules from request mode to require mode.

Note

In addition to the steps shown in this procedure, you can also use network traffic capture tools such as Microsoft Network Monitor, which can be downloaded from http://go.microsoft.com/fwlink/?linkid=94770. Network Monitor and similar tools allow you to capture, parse, and display the network packets received by the network adapter on your computer. Current versions of these tools include full support for IPsec. They can identify encrypted network packets, but they cannot decrypt them.

Administrative credentials

To complete these procedures, you must be a member of the Domain Administrators group, or otherwise be delegated permissions to modify the GPOs.

For computers running Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008, or Windows Server 2008 R2

To verify that network connections are authenticated by using the Windows Firewall with Advanced Security MMC snap-in

1. Click the **Start** charm, type **wf.msc**, and then press ENTER.

   Windows Firewall with Advanced Security opens.
2. In the navigation pane, expand **Monitoring**, and then click **Connection Security Rules**.

   The details pane displays the rules currently in effect on the computer.
3.

To display the Rule Source column

a. In the **Actions** pane, click **View**, and then click **Add/Remove Columns**.
b. In the **Available columns** list, select **Rule Source**, and then click **Add**.
c. Use the **Move up** and **Move down** buttons to rearrange the order. Click **OK** when you are finished.

   It can take a few moments for the list to be refreshed with the newly added column.

4. Examine the list for the rules from GPOs that you expect to be applied to this computer.

   Note
   If the rules do not appear in the list, then troubleshoot the GPO security group and the WMI filters that are applied to the GPO. Make sure that the local computer is a member of the appropriate groups and meets the requirements of the WMI filters.

5. In the navigation pane, expand **Security Associations**, and then click **Main Mode**.

   The current list of main mode associations that have been negotiated with other computers appears in the details column.

6. Examine the list of main mode security associations for sessions between the local computer and the remote computer. Make sure that the **1st Authentication Method** and **2nd Authentication Method** columns contain expected values. If your rules specify only a first authentication method, then the **2nd Authentication Method** column displays **No authentication**. If you double-click the row, then the **Properties** dialog box appears with additional details about the security association.
7. In the navigation pane, click **Quick mode**.
8. Examine the list of quick mode security associations for sessions between the local computer and the remote computer. Make sure that the **AH Integrity**, **ESP integrity**, and **ESP Confidentiality** columns contain expected values.

Additional Resources

For more information about the technologies discussed in this guide, see topics referenced in the following sections.

Windows Firewall with Advanced Security

- Windows Firewall with Advanced Security Overview (http://technet.microsoft.com/library/hh831365.aspx)

This TechNet page contains links to a variety of documents available for Windows Firewall with Advanced Security in Windows Server 2012.

- Troubleshooting Windows Firewall with Advanced Security in Windows Server 2012

  This wiki article describes how Windows Firewall with Advanced Security works, what the common troubleshooting situations are, and which tools you can use for troubleshooting. The community is encouraged to add their troubleshooting and experiences to this article.

IPsec

- IPsec (http://www.microsoft.com/ipsec)

  This TechNet page contains links to a variety of documents currently available for Internet Protocol security (IPsec) in Windows.

Group Policy

Group Policy is a key method for implementing firewall and server and domain isolation designs.

For more information about Group Policy and related technologies, see:

- Group Policy Overview (http://technet.microsoft.com/library/hh831791.aspx)

  This page contains links to the documents currently available for Group Policy.
- WMI Filtering Using GPMC (http://go.microsoft.com/fwlink/?linkid=93188)
- HOWTO: Leverage Group Policies with WMI Filters (http://go.microsoft.com/fwlink/?linkid=93760)

  This article describes how to create a WMI filter to set the scope of a GPO based on computer attributes, such as operating system.

Active Directory Domain Services

In Windows 8 and Windows Server 2012, organizations can use AD DS to manage users and resources, such as computers, printers, or applications, on a network. Server isolation and domain isolation also require AD DS to use the Kerberos V5 protocol for IPsec authentication.

For more information about AD DS and related technologies, see:

- Active Directory Domain Services Overview (http://technet.microsoft.com/library/hh831484.aspx)

Windows Internet Name Service (WINS) Overview

Windows Internet Name Service (WINS) provides a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses that are used on your network. The WINS service on Windows Server 2012 enables a server to act as a NetBIOS name server that

registers and resolves names for WINS-enabled client computers on your network as described in the NetBIOS over TCP/IP standards.

WINS primarily supports clients that run early versions of the Windows operating system and applications that require NetBIOS. Later versions of Windows including Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012, Windows 2000, Windows XP, Windows Vista, Windows 7, Windows 8 use Domain Name System (DNS) names in addition to NetBIOS names. Environments that include older computers that continue to use NetBIOS names on the same network as computers that use DNS names must use WINS servers in addition to the DNS servers.

If all of your network computers are running one of these later versions of Windows, and you are not supporting any applications that require NetBIOS names, we recommend that you establish DNS as your only method of name resolution. However, before you consider decommissioning WINS servers, identify any computers or applications that rely on NetBIOS, and determine the impact of removing NetBIOS name support from the network. You might find that a critical application relies on NetBIOS (with no alternative currently available) in which case, you must continue to use WINS. For example, older versions of certain applications, such as Microsoft Systems Management Server (SMS) or Microsoft BackOffice Server client/server mail configurations might require NetBIOS naming.

The documentation for WINS has not been updated for this version of Windows. Aside from some software engineering changes that make WINS less vulnerable to malicious attack, WINS functionality has not changed from that available in Windows Server 2003.

You can find all of the documentation for WINS at Windows Internet Name Service in the Windows Server TechCenter.

Windows PowerShell Cmdlets for Networking

You can use the following Windows PowerShell cmdlets to configure networking technologies in Windows Server® 2012.

- BranchCache Cmdlets in Windows PowerShell
- Data Center Bridging (DCB) Quality of Service (QoS) Cmdlets in Windows PowerShell
- DHCP Server Cmdlets in Windows PowerShell
- DirectAccess Client Cmdlets in Windows PowerShell
- DNS Client Cmdlets in Windows PowerShell
- Hyper-V Cmdlets in Windows PowerShell

> **Note**
> Hyper-V includes cmdlets to configure virtual networking features, such as virtual network adapters and the Hyper-V Virtual Switch, in addition to cmdlets that are not related to networking.

- IPAM Server Cmdlets in Windows PowerShell
- Net TCP/IP Cmdlets in Windows PowerShell
- Network Adapter Cmdlets in Windows PowerShell
- Network Connectivity Status Cmdlets in Windows PowerShell
- Network Load Balancing Cmdlets in Windows PowerShell
- Network Quality of Service (QoS) Cmdlets in Windows PowerShell
- Network Security Cmdlets in Windows PowerShell
- Network Switch Team Cmdlets in Windows PowerShell
- NIC Teaming (NetLBFO) Cmdlets in Windows PowerShell
- Remote Access Cmdlets in Windows PowerShell

Windows Server Gateway

This topic, which is intended for Information Technology (IT) professionals, provides overview information about Windows Server Gateway, including Windows Server Gateway capabilities and features.

Notes

In addition to this topic, the following Windows Server Gateway documentation is available.

- Windows Server Gateway Hardware and Configuration Requirements
- Enable the forwarding function on Windows Server Gateway: a use case study
- Test Lab Guide: Windows Server 2012 R2 Hyper-V Network Virtualization with System Center 2012 R2 VMM
- How to Use a Server Running Windows Server 2012 R2 as a Gateway with VMM
- Configuring VM Networks and Gateways in VMM

Who will be interested in Windows Server Gateway?

If you are a system administrator, network architect, or other IT professional, Windows Server Gateway might interest you under one or more of the following circumstances:

- You design or support IT infrastructure for an organization that is using or planning to use Hyper-V to deploy virtual machines (VMs) on virtual networks.
- You design or support IT infrastructure for an organization that has deployed or is planning to deploy cloud technologies.
- You want to provide full network connectivity between physical networks and virtual networks.
- You want to provide your organization’s customers with access to their virtual networks over the Internet.

This topic includes the following sections.

- What is Windows Server Gateway?
- Windows Server Gateway integration with Hyper-V Network Virtualization

- Clustering Windows Server Gateway for high availability
- Windows Server Gateway as a forwarding gateway for private cloud environments
- Windows Server Gateway as a site-to-site VPN gateway for hybrid cloud environments
- Multitenant Network Address Translation (NAT) for VM Internet access
- Multitenant remote access VPN connections

What is Windows Server Gateway?

Windows Server Gateway (Windows Server Gateway) is a virtual machine (VM)-based software router that allows Cloud Service Providers (CSPs) and Enterprises to enable datacenter and cloud network traffic routing between virtual and physical networks, including the Internet.

Virtual networks are created by using Hyper-V Network Virtualization, which is a technology that was introduced in Windows Server® 2012.

Hyper-V Network Virtualization provides the concept of a virtual machine (VM) network that is independent of the underlying physical network. With this concept of VM networks, which are composed of one or more virtual subnets, the exact physical location of an IP subnet is decoupled from the virtual network topology. As a result, organizations can easily move their subnets to the cloud while preserving their existing IP addresses and topology in the cloud. This ability to preserve infrastructure allows existing services to continue to work, unaware of the physical location of the subnets. That is, Hyper-V Network Virtualization enables a seamless hybrid cloud.

In both private and hybrid cloud environments using Windows Server 2012, however, it was difficult to provide connectivity between VMs on the virtual network and resources on physical networks at local and remote sites, creating a circumstance where virtual subnets were islands separated from the rest of the network.

In Windows Server 2012 R2, Windows Server Gateway routes network traffic between the physical network and VM network resources, regardless of where the resources are located. You can use Windows Server Gateway to route network traffic between physical and virtual networks at the same physical location or at many different physical locations. For example, if you have both a physical network and a virtual network at the same physical location, you can deploy a computer running Hyper-V that is configured with a Windows Server Gateway VM to act as a forwarding gateway and route traffic between the virtual and physical networks. In another example, if your virtual networks exist in the cloud, your CSP can deploy a Windows Server Gateway so that you can create a virtual private network (VPN) site-to-site connection between your VPN server and the CSP's Windows Server Gateway; when this link is established you can connect to your virtual resources in the cloud over the VPN connection.

Windows Server Gateway integration with Hyper-V Network Virtualization

Windows Server Gateway is integrated with Hyper-V Network Virtualization, and is able to route network traffic effectively in circumstances where there are many different customers – or tenants – who have isolated virtual networks in the same datacenter.

Multi-tenancy is the ability of a cloud infrastructure to support the virtual machine workloads of multiple tenants, but isolate them from each other, while all of the workloads run on the same infrastructure. The multiple workloads of an individual tenant can interconnect and be managed remotely, but these systems do not interconnect with the workloads of other tenants, nor can other tenants remotely manage them.

For example, an Enterprise might have many different virtual subnets, each of which is dedicated to servicing a specific department, such as Research and Development or Accounting. In another example, a CSP has many tenants with isolated virtual subnets existing in the same physical datacenter. In both cases, Windows Server Gateway can route traffic to and from each tenant while maintaining the designed isolation of each tenant. This capability makes the Windows Server Gateway multitenant-aware.

Hyper-V Network Virtualization is a network overlay technology using Network Virtualization Generic Routing Encapsulation (NVGRE), which allows tenants to bring their own address space and allows CSPs better scalability than is possible by using VLANs for isolation.

Note

For more information about Hyper-V Network Virtualization and Hyper-V Virtual Switch in Windows Server 2012, see Hyper-V Network Virtualization Overview and Hyper-V Virtual Switch Overview in the Windows Server 2012 Technical Library.

Clustering Windows Server Gateway for high availability

Windows Server Gateway is deployed on a dedicated computer that is running Hyper-V and that is configured with one VM. The VM is then configured as a Windows Server Gateway.

For high availability of network resources, you can deploy Windows Server Gateway with failover by using two physical host servers running Hyper-V that are each also running a virtual machine (VM) that is configured as a gateway. The gateway VMs are then configured as a cluster to provide failover protection against network outages and hardware failure.

When you deploy Windows Server Gateway, the host servers running Hyper-V and the VMs that you configure as gateways must be running Windows Server 2012 R2.

Unless otherwise noted in the illustrations that are provided in the sections below, the following icon represents two Hyper-V hosts, each of which is running a VM configured as a Windows Server Gateway. In addition, both the servers running Hyper-V and the VMs on each server are running Windows Server 2012 R2, and the gateway VMs are clustered.

Windows Server Gateway as a forwarding gateway for private cloud environments

Private cloud is a computing model that uses infrastructure dedicated to your organization. A private cloud shares many of the characteristics of public cloud computing including resource pooling, self-service, elasticity, and metered services delivered in a standardized manner with the additional control and customization available from dedicated resources.

The only fundamental difference between a private cloud and a public cloud is that a public cloud provides cloud resources to multiple organizations, while the private cloud hosts resources for a single organization. However, a single organization may have multiple business units and divisions which can lend itself to being multi-tenant in nature. In these circumstances, private cloud shares many of the security and isolation requirements of public cloud.

For Enterprises that deploy an on-premises private cloud, Windows Server Gateway can act as a forwarding gateway and route traffic between virtual networks and the physical network. For example, if you have created virtual networks for one or more of your departments, such as Research and Development or Accounting, but many of your key resources (such as Active Directory Domain Services, SharePoint, or DNS) are on your physical network, Windows Server Gateway can route traffic between the virtual network and the physical network to provide employees working on the virtual network with all of the services that they need.

In the illustration below, the physical and virtual networks are at the same physical location. Windows Server Gateway is used to route traffic between the physical network and virtual networks.

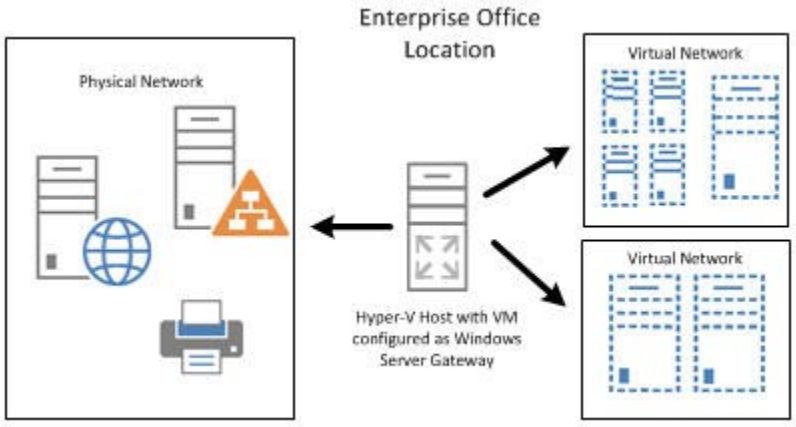

Windows Server Gateway as a site-to-site VPN gateway for hybrid cloud environments

For CSPs that host many tenants in their datacenter, Windows Server Gateway provides a multitenant gateway solution that allows your tenants to access and manage their resources over site-to-site VPN connections from remote sites, and that allows network traffic flow between virtual resources in your datacenter and their physical network.

In the illustration below, a CSP provides datacenter network access to multiple tenants, some of whom have multiple sites across the Internet. In this example, tenants use third party VPN servers at their corporate sites, while the CSP uses Windows Server Gateway for the site-to-site VPN connections.

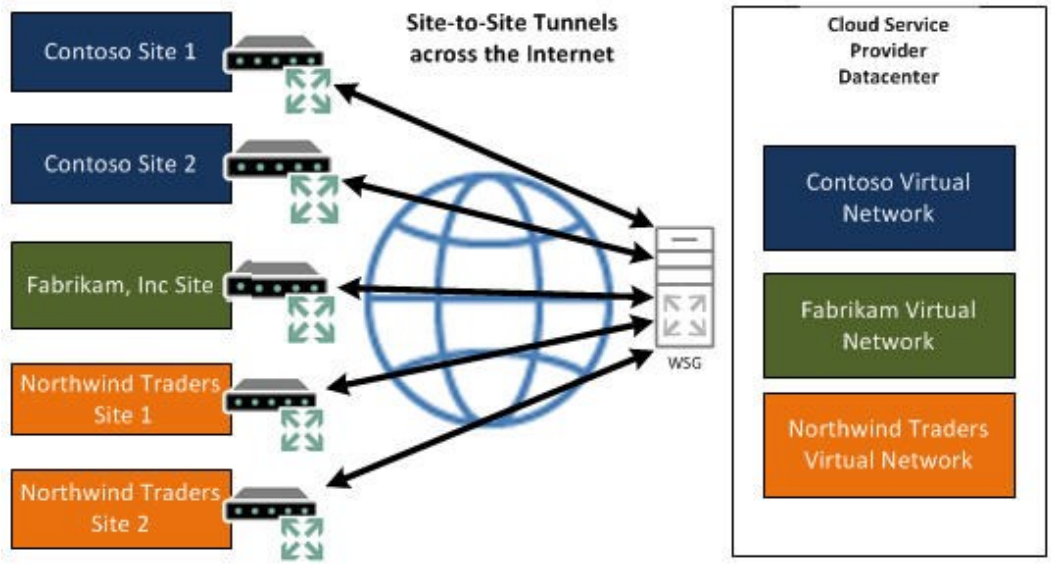

Multitenant Network Address Translation (NAT) for VM Internet access

In the illustration below, a home user running a Web browser on their computer makes a purchase on the Internet from a Contoso Web server that is a VM on the Contoso Virtual Network. During the purchasing process, the Web app verifies the credit card information provided by the home user by connecting to a Financial Services company on the Internet. This ability to connect from the virtual network to Internet resources is provided when NAT is enabled on the CSP Windows Server Gateway.

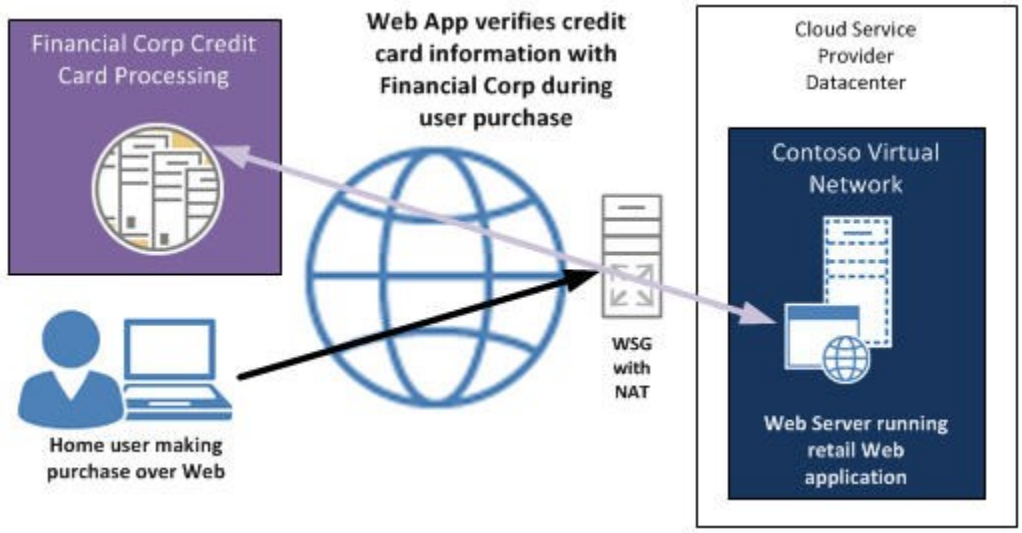

Multitenant remote access VPN connections

In the illustration below, Administrators use VPN dial-in connections to administer VMs on their corporate virtual networks. The Administrator from Contoso initiates the VPN connection from an Internet-enabled branch office, and connects through the CSP Windows Server Gateway to the Contoso Virtual Network.

Similarly, the Northwind Traders Administrator establishes a VPN connection from a residence office to manage VMs on the Northwind Traders Virtual Network.

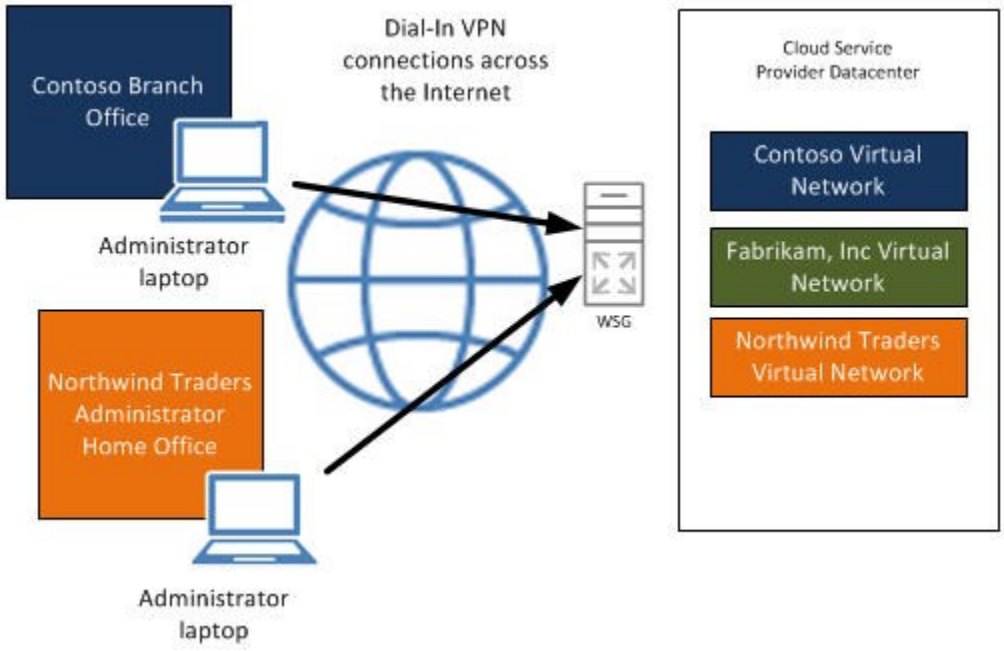

Windows Server Gateway Hardware and Configuration Requirements

You can use the information in this topic to choose the correct server hardware configuration when you are deploying Windows Server Gateway in a hybrid cloud or private cloud deployment. This topic provides hardware specifications and configuration recommendations for servers that are running Hyper-V, in addition to configuration parameters for Windows Server Gateway virtual machines (VMs).

On each server running Hyper-V, you can configure either two or four Windows Server Gateway VMs, depending on the requirements for your deployment. For example, if your organization is an Enterprise with a private cloud deployment, you might need only two Windows Server Gateway VMs.

In another example, if your organization is a Cloud Service Provider (CSP) with two hundred tenants in your datacenter, you can use four Windows Server Gateway VMs, with each Windows Server Gateway VM providing routing services for fifty tenants.

The following sections contain hardware and configuration requirements when you deploy Windows Server Gateway.

- Hyper-V hardware recommendations
- Failover Clustering and Load Balancing
- Hyper-V host configuration
- Windows Server Gateway VM configuration

Hyper-V hardware recommendations

Following is the recommended minimum hardware configuration for each server that is running Windows Server® 2012 R2 and Hyper-V.

Server component	Specification
Central Processing Unit (CPU)	Non-Uniform Memory Architecture (NUMA) nodes: 2
Cores per NUMA node	8
Hyper-threading	Disabled. Hyper-Threading does not improve the performance of Windows Server Gateway.
Random Access Memory (RAM)	48 GB
Network Interface Cards (NICs)	Two 10 GB NICs

Ensure that the number of virtual processors that are assigned to a Windows Server Gateway VM does not exceed the number of processors on the NUMA node. For example, if a NUMA node

has 8 cores, the number of virtual processors should be less than or equal to 8. To find out the number of NUMA nodes and the number of cores per NUMA node, run the following Windows PowerShell script on each Hyper-V host:

```
$nodes = [object[]] $(gwmi –Namespace root\virtualization\v2 –Class MSVM_NumaNode)

$cores = ($nodes | Measure-Object NumberOfProcessorCores –sum).Sum

$lps = ($nodes | Measure-Object NumberOfLogicalProcessors –sum).Sum



Write-Host "Number of NUMA Nodes: ", $nodes.count

Write-Host ("Total Number of Cores: ", $cores)

Write-Host ("Total Number of Logical Processors: ", $lps)
```

Important

Allocating virtual processors across NUMA nodes might have a negative performance impact on Windows Server Gateway. Running multiple VMs, each of which has virtual processors from one NUMA node, likely provides better aggregate performance than a single VM to which all virtual processors are assigned.

Following are the recommendations when you select the number of VMs to install on each Hyper-V host when each NUMA node has eight cores:

- Two gateway VMs with eight virtual processors each and at least 8GB RAM
- Four gateway VMs with four virtual processors each and at least 8GB RAM

Failover Clustering and Load Balancing

For load balancing and failover, it is recommended that you deploy two servers running Hyper-V and that you configure the computers in a failover cluster, which is depicted in the following illustration.

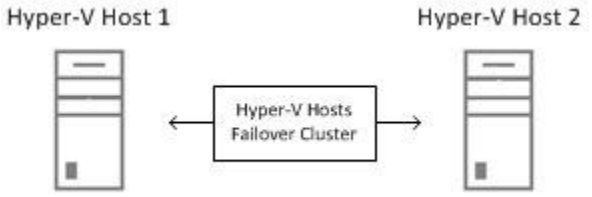

In addition, you must configure the Windows Server Gateway VMs on each Hyper-V host in individual guest clusters, so that there are either two Windows Server Gateway VM guest clusters or four Windows Server Gateway VM guest clusters. The illustration below depicts eight VMs, four on each Hyper-V host, that are configured in four guest clusters.

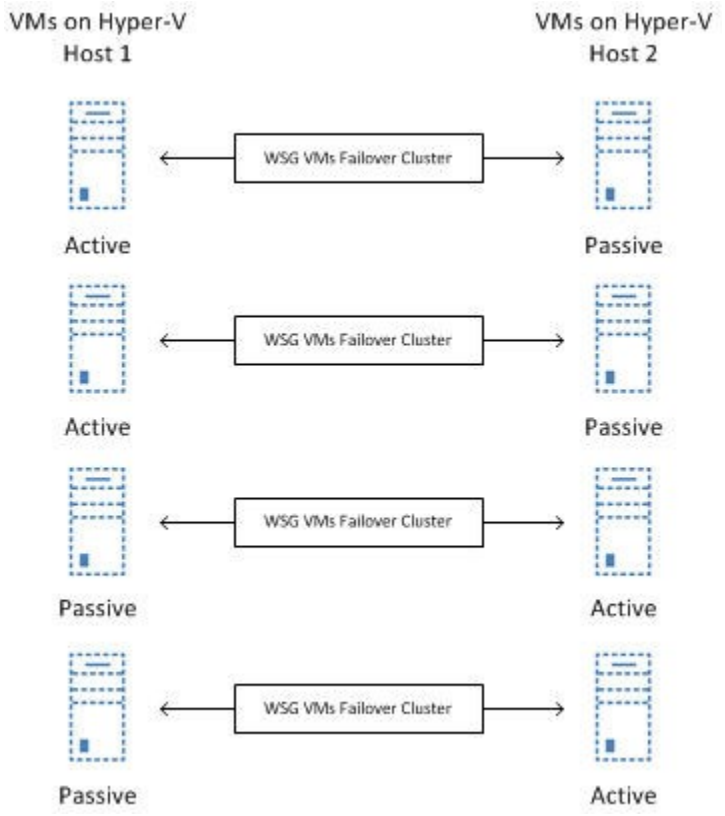

When more than one gateway VM cluster is hosted on a two node Hyper-V host cluster, you can distribute active gateways across the two Hyper-V hosts to load balance the workloads.

Hyper-V host configuration

Following is the recommended configuration for each server that is running Windows Server® 2012 R2 and Hyper-V and whose workload is to run Windows Server Gateway VMs. These configuration instructions include the use of Windows PowerShell command examples. These examples contain placeholders for actual values that you need to provide when you run the commands in your environment. For example, network adapter name placeholders are “NIC1” and “NIC2.” When you run commands that use these placeholders, utilize the actual names of the network adapters on your servers rather than using the placeholders, or the commands will fail.

Note
To run the following Windows PowerShell commands, you must be a member of the Administrators group.

Configuration item	Windows PowerShell Configuration
NIC Teaming	In our test environment, both NICs are

Configuration item	Windows PowerShell Configuration
	connected to the same physical switch. You can use the following command to create a NIC team; ensure that you run the command with parameter values that are appropriate for your deployment. `New-NetLbfoTeam "Team" -TeamMember "NIC1","NIC2" -TeamingMode SwitchIndependent -LoadBalancingAlgorithm Dynamic`
Interrupt Moderation on physical NICs	Use default settings. To check the configuration you can use the following Windows PowerShell command: `Get-NetAdapterAdvancedProperty`
Receive Buffers size on physical NICs	You can verify whether the physical NICs support the configuration of this parameter by running the command **Get-NetAdapterAdvancedProperty**. If they do not support this parameter, the output from the command does not include the property "Receive Buffers." If NICs do support this parameter, you can use the following Windows PowerShell command to set the Receive Buffers size: `Set-NetAdapterAdvancedProperty "NIC1" -DisplayName "Receive Buffers" -DisplayValue 3000`
Send Buffers size on physical NICs	You can verify whether the physical NICs support the configuration of this parameter by running the command **Get-NetAdapterAdvancedProperty**. If the NICs do not support this parameter, the output from the command does not include the property "Send Buffers." If NICs do support this parameter, you can use the following Windows PowerShell command to set the Send Buffers size: `Set-NetAdapterAdvancedProperty "NIC1" -DisplayName "Transmit Buffers" -DisplayValue`

Configuration item	Windows PowerShell Configuration
	`3000`

Receive Side Scaling (RSS) on physical NICs	You can verify whether your physical NICs have RSS enabled by running the Windows PowerShell command **Get-NetAdapterRss**. You can use the following Windows PowerShell commands to enable and configure RSS on your network adapters.
	`Enable-NetAdapterRss "NIC1","NIC2"`
	`Set-NetAdapterRss "NIC1","NIC2" -NumberOfReceiveQueues 16 -MaxProcessors`

Virtual Machine Queue (VMQ) on the NIC Team	You can enable VMQ on your NIC team by using the following Windows PowerShell command. `Enable-NetAdapterVmq "Team"`

Windows Server Gateway VM configuration

On both Hyper-V hosts, you can configure two or four VMs that are configured as gateways with Windows Server Gateway. You can use Virtual Switch Manager to create a Hyper-V Virtual Switch that is bound to the NIC team on the Hyper-V host.

After you have created Windows Server Gateway VMs, configure each Windows Server Gateway VM with four virtual network adapters that are connected to the Hyper-V Virtual Switch that is bound to the NIC Team.

Following is the recommended configuration for each Windows Server Gateway VM.

Configuration item	Windows PowerShell configuration
Memory	8 GB
Number of virtual network adapters	4 NICs with the following specific uses: 1 for Management that is used by the management operating system 1 for Clustering 1 External that provides access to external networks 1 that is Internal that provides access to internal networks only.
Receive Side Scaling (RSS)	You can keep the default RSS settings for the

Configuration item	Windows PowerShell configuration
	Management NIC and the Cluster NIC. The following example configuration is for a VM that has 8 virtual processors. For the External and Internal NICs, you can enable RSS with BaseProcNumber set to 0 and MaxRssProcessors set to 8 using the following Windows PowerShell command. `Set-NetAdapterRss "Internal","External" -BaseProcNumber 0 -MaxProcessorNumber 8`
Send Side Buffer	You can keep the default Send Side Buffer settings for the Management NIC and the Cluster NIC. For both the Internal and External NICs you can configure the Send Side Buffer with 32 MB of RAM by using the following Windows PowerShell command. `Set-NetAdapterAdvancedProperty "Internal","External" -DisplayName "Send Buffer Size" -DisplayValue "32MB"`
Receive Side Buffer	You can keep the default Receive Side Buffer settings for the Management NIC and the Cluster NIC. For both the Internal and External NICs you can configure the Receive Side Buffer with 16 MB of RAM by using the following Windows PowerShell command. `Set-NetAdapterAdvancedProperty "Internal","External" -DisplayName "Receive Buffer Size" -DisplayValue "16MB"`
Forward Optimization	You can keep the default Forward Optimization settings for the Management NIC and the Cluster NIC. For both the Internal and External NICs you can enable Forward Optimization by using the following Windows PowerShell command. `Set-NetAdapterAdvancedProperty "Internal","External" -DisplayName "Forward`

Configuration item	Windows PowerShell configuration
	`Optimization" -DisplayValue "1"`

See Also

Windows Server Gateway

Print and Document Services Overview

This is an overview of Print and Document Services, including Print Server, Distributed Scan Server and Fax Server in Windows Server 2012.

Notes

In addition to this topic, the following Print and Document Services documentation is available.

- What's New in Print and Document Services in Windows Server 2012 R2
- Install Print and Document Services
- Install Distributed Scan Server
- Fax Server Step-by-Step Guide
- Configure Print and Document Services
- Branch Office Direct Printing Overview
- Branch Office Direct Print (BODP) Logging in Windows Server 2012 R2
- Migrate Print and Document Services to Windows Server 2012
- Print Server Scalability and Capacity Planning
- High Availability Printing Overview
- Client Name Instrumentation in Windows Server 2012 R2

Did you mean…

- Print and Document Services for Windows Server 2008 R2

Role description

Print and Document Services enables you to centralize print server and network printer tasks. With this role, you can also receive scanned documents from network scanners and route the documents to a shared network resource, Windows SharePoint Services site, or email addresses.

Fax Server sends and receives faxes and allows you to manage fax resources such as jobs, settings, reports, and fax devices on your fax server.

Practical applications

Print Server helps you monitor print queues and receive notifications when print queues stop processing print jobs. It also enables you to migrate print servers and deploy printer connections using Group Policy.

Distributed Scan Server provides the service which receives scanned documents from network scanners and routes them to the correct destinations. It also includes the Scan Management snap-in, which you can use to manage network scanners and configure scan processes.

The Fax Server server role can save the fax administrator a significant amount of time. Instead of installing fax machines on client computers and managing and monitoring fax machines, you can complete those tasks in 2 or 3 steps on multiple computers simultaneously and remotely.

By using Fax Service Manager, you can automatically make fax connections available to users and computers in your organization.

New and changed functionality

The following are new or changed features in Windows Server 2012 Print and Document Services.

Feature/functionality	Windows Server 2008 R2	Windows Server 2012
Type 4 drivers		X
Branch Office Direct Printing		X
Print Management Windows PowerShell module		X
WSD Secure printing		X
High Availability Printing	X	X
OpenXPS format support		X

Server Manager information

The following role services can be installed with this role.

Role service	Description
Print Server	Includes the Print Management snap-in, which is used for managing multiple printers or print servers and migrating printers to and from other Windows print servers.

Role service	Description
Distributed Scan Server	Provides the service which receives scanned documents from network scanners and routes them to the correct destinations. It also includes the Scan Management snap-in, which you can use to manage network scanners and configure scan processes.
Internet Printing	Creates a Web site where users can manage print jobs on the server. It also enables users who have Internet Printing Client installed to use a Web browser to connect and print to shared printers on the server by using the Internet Printing Protocol (IPP).
LPD Service	Line Printer Daemon (LPD) Service enables UNIX-based computers or other computers using the Line Printer Remote (LPR) service to print to shared printers on this server.

Important
The LPD and LPR Services are deprecated starting with Windows Server 2012. Eventually, they will be completely removed from the product, but they are still available in this release. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features. For more information about features or functionalities that have either been removed from the product in the current release or are planned for potential removal in subsequent releases, see Features Removed or Deprecated in Windows Server 2012.

See also

Refer to the following content for more information about Print and Document Services in Windows Server 2012.

Content type	References
Deployment	Install Print and Document Services | Configure Print and Document Services | High Availability Printing Overview | Install Distributed Scan Server | Fax Server Step-by-Step Guide
Architecture	Print and Document Services Architecture
Migration	Migrate Print and Document Services to

Content type	References
	Windows Server 2012
Troubleshooting	Troubleshooting Windows Server 2012 Printing
Test labs	Test Lab Guide: Demonstrate Windows Server 2012 Printing | Test Lab Guide: Demonstrate Windows Server 2012 High Availability Printing
Related technologies	Branch Office Direct Printing Overview | Changing the Default Format for Microsoft XPS Document Writer
Tools and Settings	Print Management Cmdlets in Windows PowerShell

What's New in Print and Document Services in Windows Server 2012 R2

This topic provides information about the new features for Print and Document Services in Windows Server 2012 R2.

Role and technology description

Print and Document Services enables you to centralize print server and network printer management tasks. With this role, you can also receive scanned documents form network scanners and route the documents to a shared network resource, Windows SharePoint Services site, or email addresses.

New and changed functionality

The following table provides a list of the new features for Print and Document Services.

Feature/functionality	New or updated?	Description
Event Logging for Branch Office Direct Printing	Updated	Printing events are now written to the print server when BODP is enabled for the queue
Printer Migration for Web Services for Devices (WSD) print	Updated	You can now use the Printer Migration Wizard or command line tools to export and import

Feature/functionality	New or updated?	Description
devices		printer settings for WSD print devices
Roaming Settings include Printer Connections	New	Printer settings and some connections are now available in roaming user profiles when you log on using different computers and devices
Easier Printing in Windows RT	New	Windows RT users can now easily find and use printers
Near Field Communication (NFC) Connections to Printers	New	You can configure shared printers to advertise NFC wireless connections to Printers so that users with compatible devices can easily locate, install, and use printers
Common framework for PIN-protected printing support by IHVs	New	A common framework is provided for PIN-protected printing, with an integrated user interface and easier implementation for IHVs
Print and Fax services now include user access logging	Updated	Print and fax service logging now includes user name and computer name logging in addition to IP address logging.

Event Logging for Branch Office Direct Printing

You can now deploy Branch Office Direct Printing (BODP) with event logging.

With BODP, the print job is rendered on the client computer and then sent directly to the local printer rather than using the remote print server as an intermediary. In previous operating system versions, print jobs using BODP were not logged and did not produce events written to Event Viewer on the print server. In Windows Server 2012 R2, BODP print jobs are logged and events are written to Event Viewer on the remote print server.

For example, if you have enabled BODP for a print queue on the print server and a user in a remote office prints to the queue, the print job goes directly to the user's local printer in the branch office. While the print job goes to the remote printer, the print event is written to the print server.

If there is a problem with the print job and it fails, the failure is logged in Event Viewer on the print server, including information such as job owner, time, which queue included the print job, and reasons for the print failure. This additional logging information provides you with the ability to identify and fix printing issues at remote office locations.

If connectivity is down between the main office that contains the print server and a branch office where a user initiates a print job on a BODP-enabled queue, the job prints at the local printer, while logging and event information is automatically cached until office connectivity is restored.When you use BODP:

- Logging is enabled by default for any BODP enabled queue.
- You can enable and disable logging on any queue by using Windows PowerShell commands.
- Sending of cached print job events do not negatively impact performance of the print server.

Printer Migration for Web Services for Devices (WSD) print devices

With Windows Server 2012 R2, you can now backup, restore, and migrate WSD print devices by using the Print Management console or by using the command line tool PrintBRM.exe.

Note

To view command line help for PrintBRM.exe, open command prompt or Windows PowerShell and change directories to %WINDIR%\System32\Spool\Tools\. Type .\PrintBrm.exe /?, and then press ENTER.

You can use Printer Migration to backup and restore the local computer, or you can target a remote WSD print server for backup and restore. When you use Printer Migration:

- Backing up a WSD device with Printer Migration results in a backup of all needed information to restore the device later, regardless of whether or not you can currently connect to the device.
- Restoring a WSD device to a computer running Windows Server 2012 R2 when you can connect to the server results in one WSD queue on the target server that is mapped to the correct port and contains the correct printer information and configuration.
- Restoring a WSD device to a computer running Windows Server 2012 R2 when you are unable to connect to the server results in one WSD queue on the target server that is mapped to the correct port and contains the correct printer information and configuration.

If you want to backup or restore a print server by using Printer Migration, open the Print Management console, click **Action**, and then click **Migrate Printers**. In the **Printer Migration** dialog box, select one of the following options:

- Export printer queues and printer drivers to a file. This option exports printer queues, printer ports, and printer drivers for both TCP/IP and WSD devices to the file that you specify.
- Import printer queues and printer drivers from a file. This option imports printer queues, printer ports, and printer drivers for both TCP/IP and WSD devices to the file that you specify.

After your selection, click **Next** and follow the instructions to conclude the process.

Roaming Settings include Printer Connections

Printer connections are now included in the roaming settings of users who log in with Microsoft accounts, making it easier for them to successfully find and connect to the shared printers that they have connected to in the past. This allows users to more easily find the devices on other computers during the device installation process.

Users can control roaming of printer connections using the PC Settings application and toggling the **Other Windows Settings** group. In addition, IT administrators can disable this feature by configuring the **Other Windows Settings** Group Policy setting.

If roaming is enabled, any printer connections that the user creates or deletes are reflected on all of their devices. For example, when a user deletes a roamed printer connection on one device, the printer connection is removed from all other devices that they use.

Note

Roamed printer settings and connections are also supported on Windows® RT.

Easier Printing in Windows RT

As a result of the included print class drivers provided with the operating system, Windows RT supports a wide variety of printers, but in order to connect to shared printers, the print server must be configured with a version 4 (v4) print driver. With Windows 8.1, it's easier to connect Windows RT clients to shared printers, even if the printer is configured using a version 3 (v3) print driver. Windows RT utilizes the local print class driver to perform client side rendering for compatible printers. This feature works best if the print server is running Windows 8.1 or Windows Server 2012 R2.

Near Field Communication (NFC) Connections to Printers

Windows 8.1 now includes connections to printers by using Near Field Communication (NFC) Tap to Pair wireless technology that provides short range connections between devices, such as smartphones, tablets, laptops, and printers.

IT administrators can create NFC tags for existing hardware that advertise a printer and allow easy discovery and installation of the printer by users. Printer installation by using programmable NFC tags is available for both WSD printers and shared printers within an Enterprise. NFC can also be used to connect to Wi-Fi Direct printers, but this requires that the NFC radio is built-into the Wi-Fi Direct device.

You can configure your printers by using the **Write-PrinterNfcTag** and **Read-PrinterNfcTag** Windows PowerShell commands on an NFC-enabled device that is running Windows 8.1 using the following steps:

1. Buy NFC forum approved tag types that are at least 700B in capacity
2. On a device that supports NFC and is running Windows 8.1, open Windows PowerShell. Run one of the following commands to program either a print share or a WSD device that does not have an intermediary print server.

- **Write-PrinterNfcTag –SharePath** *\\server\printer*
- **Write-PrinterNfcTag –WsdAddress** *printerhostname*

3. When prompted, tap the tag against the NFC radio on your device. This programs the tag. If you have hardware and tags that support it, you can also optionally specify the **–Lock** flag on any of the previous commands to prevent subsequent modification of the tag.

This completes programming of the tag. Any client computer or device can now tap the tag to initiate installation. After the tap, the user is prompted for consent; when the user provides consent, the printer is installed in the background.

Common framework for PIN-protected printing support by IHVs

Independent Hardware Vendors (IHVs) who provide PIN-protected printing can now benefit from the new common framework for PIN-protected printing support included in Windows 8.1. This new feature makes the delivery of PIN-protected printing easier for IHVs, and the usage of PIN-protected printing easier for users.

To prevent unauthorized use of printing resources, you can modify printing queue defaults by using the Print Management console or Windows PowerShell. To enable PIN-protected printing, set **PIN-protected printing** to **On**, and then assign a default PIN number.

You can also configure PIN-protected printing to use a unique PIN for each print job. To do so, do not assign a default PIN number. This will cause users to be prompted to enter a PIN when printing from Windows Store applications.

Print and Fax services now include user access logging

In previous operating system versions, you were able to perform accounting by measuring Client Access Licenses (CALs) for Print and Document Services only by using IP addresses, because only IP addresses were logged when users sent print jobs or sent and received faxes.

In Windows Server 2012 R2, you can now log the number of unique users that utilize print and fax services by both IP address and user name, allowing you to collect accurate volume license data for the Print and Document Services server role.

This logging of the user name allows you to verify that the CALs that are used on a print/fax server are aligned with your terms of agreement. You can validate the correct license usage while also collecting information on usage trends for future licensing decisions.

You can use the following Windows PowerShell User Access Logging (UAL) Cmdlets to view logging information.

- Get-UalUserAccess
- Get-UalDailyUserAccess

See Also

Print and Document Services Overview

Install Print and Document Services

Before installing Print and Document Services in Windows Server 2012, consider the Printing and Scanning scenarios that your organization will benefit from and what kind of advanced features are required. In addition to local installation of role services, Windows Server 2012 Server Manager supports remote installation and configuration of role services and can be used to install Print and Document Services on a remote computer in your organization. You can then remotely manage your print server from your local computer. Start the Print Management console on your local computer, and then add the remote print server to the console.

The following sections outline the deployment scenarios for Print and Document Services and how to use Server Manager to install and perform the initial configuration of the features that make up a specific scenario.

Deployment Scenarios

Windows Server 2012 Print and Document Services consists of several different role services that can be installed individually to support basic functionality, or can be combined with other role services to provide more advanced features. When installing Print and Document Services features to support advanced functionality it is helpful to think of the scenarios that are enabled by those advanced features. Windows Server 2012 Server Manager will prompt the administrator to select the role services to be installed at the time of installation and to choose which role services are installed on each server in the deployment.

Print Server

The Print Server role is used to create a Windows Print Server and includes the Print Management Console snap-in, which is useful for managing multiple printers or print servers and migrating printers to and from other Windows print servers.

Internet Printing

The Internet Printing service creates a web site where users can manage print jobs on a Print Server. It also enables users who have the Internet Printing Client installed to use a web browser to connect and print to shared printers by using the Internet Printing Protocol (IPP).

LPD Service

The Line Printer Daemon service enables UNIX-based computers or other computers using the Line Printer Remote (LPR) service to print to shared printers on a Windows Print Server.

Important

The LPD and LPR Services are deprecated starting with Windows Server 2012. Eventually, they will be completely removed from the product, but they are still available

in this release. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features. For more information about features or functionalities that have either been removed from the product in the current release or are planned for potential removal in subsequent releases, see Features Removed or Deprecated in Windows Server 2012.

Installation Process

The following section provides details of the new Server Manager user interface and how to use it to install the necessary role services needed for the chosen Print and Document role services.

Installation UI/Wizard

Installation of all Print and Document Services roles are performed using the new Server Manager user interface for Windows Server "8" Beta. This new interface provides both local and remote server management and uses a series of plugins to manage specific server functions, role services, and deployment types.

The following is an example of the Server Manager Dashboard page, which is the first page that will be displayed when Server Manager is opened:

Figure 1: Server Manager Dashboard

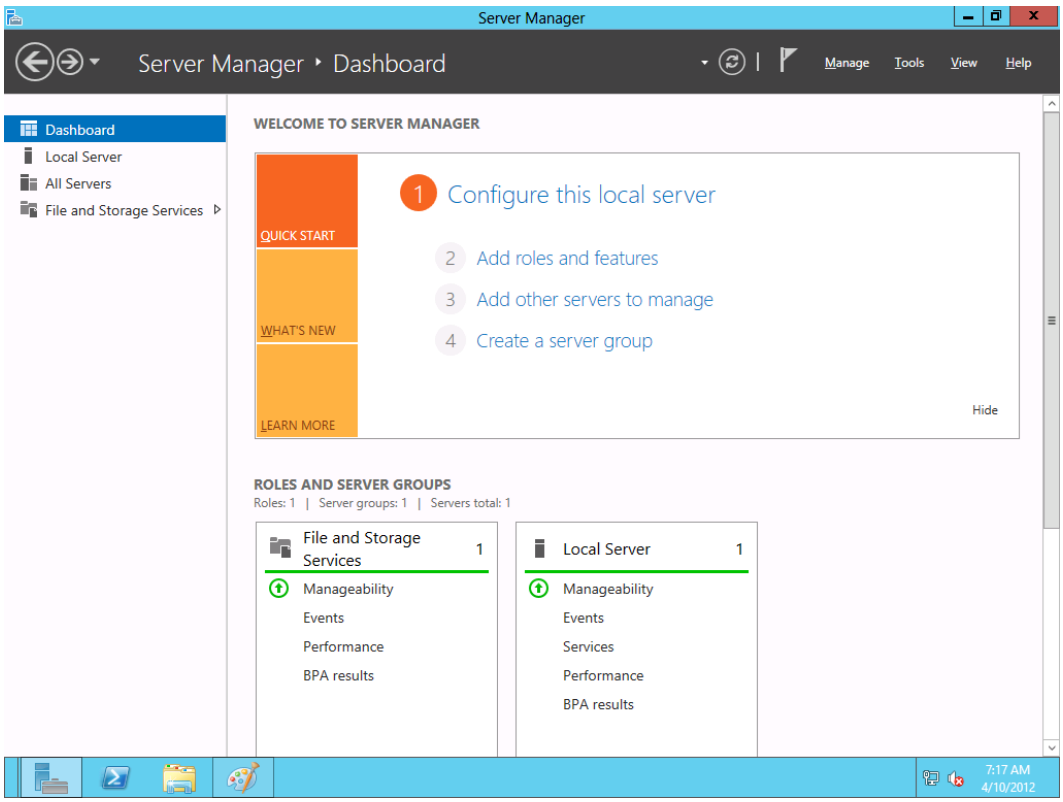

Installing Print and Document Services

Once the server that will host Print and Document services has been configured with initial settings, the role services can be installed using the Add Roles and Features method of Server Manager.

To install Print and Document Services

1. Open Server Manager and click **All Servers** in the navigation pane.
2. Click **Manage** in the Menu Bar and then click **Add Roles and Features**.

   Important
   Server Manager prompts you to verify that the following prerequisites have been met:
   - The Administrator account has a strong password
   - Network settings, such as static IP addresses are configured
   - The latest security updates from Windows Update are installed

   Ensure these prerequisites have been met before continuing the installation.
3. Click **Next**, select **Role or feature-based Installation**, and then click **Next**.
4. On the **Select destination server** page, select the server to install the Print and Document Services on. The default server is the local server. Then click **Next**.
5. On the **Select server roles** page, select **Print and Document Services**. Click **Add Features** to add the features that are required for Print and Document Services, then click **Next**.
6. Click **Next** on the **Select features** page.
7. On the **Print and Document Services** page, review the notes for the administrator and then click **Next**.
8. On the **Select role services** page, choose the role services that you wish to install. By default, **Print Server** is selected automatically. If you select Internet Printing you must also add the required features by clicking **Add Features**. Click **Next**.
9. If you added Internet Printing, verify the Web Server Role settings and click **Next** until the **Confirm installation selections** page is displayed. Click **Install** to install the required role services.
10. The Installation progress page displays the status of the installation.

Figure 2: Installation progress

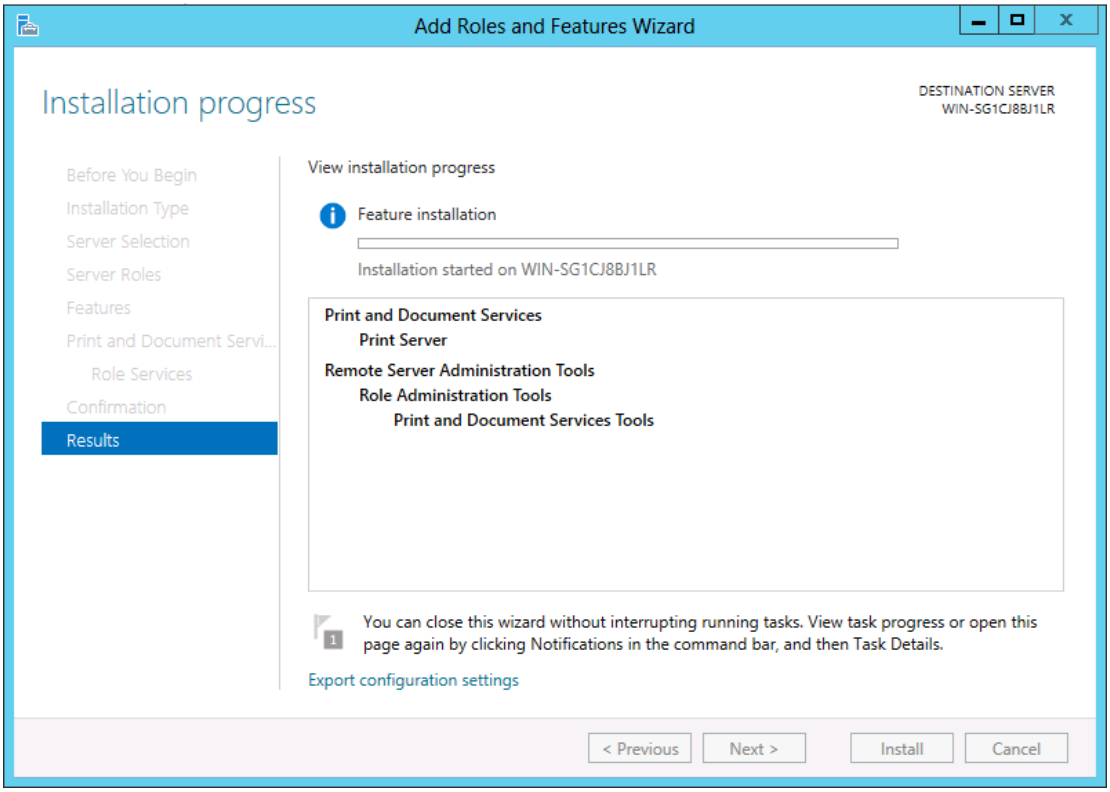

11. When installation has finished, restart the server to complete the installation and configuration. To see the status of each component installation, open the Task Details by clicking on the Notifications flag.

Verifying Installation

To verify installation of the Print and Document role services, open Server Manager and click **Print Services** in the navigation pane. If the role services were successfully installed, Server Manager should display the name of the server in the list of Servers, and the **Print Spooler** service displays as Running in the Services tile:

Figure 6: Print and Document Services Installation Verification

Figure 3: Print and Document Services Installation Verification

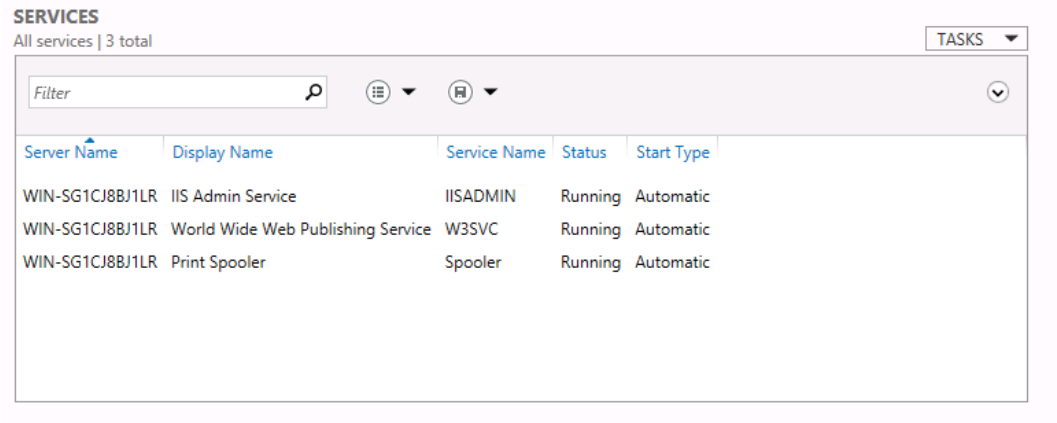

The IIS Admin Service and the World Wide Web Publishing Service are required services for Internet Printing, so these services would also be present and a running on a server that has the Internet Printing role service installed.

Installing on Server Core

The Print Server role can be installed on Server Core installations, and Windows Server 2012 now provides XPS and v4 driver support on both Full Server and Server Core installations.

To install the Print Server role on a Server Core server

1. Server Core installations will start a command prompt by default. Type the following commands to start Windows PowerShell, load the Server Manager provider, and install the Print-Server role.

   Windows PowerShell commands

   ```
   PowerShell

   ipmo ServerManager

   add-WindowsFeature Print-Server
   ```

2. Restart the server. Before you can remotely configure the server you must follow the steps in the Remote Management of Server Core Print Servers section of this article.
3. Printers and drivers can also be installed using Windows PowerShell if the server is to be managed locally. For more information, see Print Management cmdlets in the Microsoft TechNet Library web site.

Installing with Minimal Server Interface

You can install a Print Server with the Minimal Server Interface, with the following limitations:

- Print Management Console email notifications will not work.
- Extended View functions will not work.

- Extended view includes printer web pages as well as viewing documents printing from the queue in Print Management console. However, you can still open the print queue to see jobs.

  If email notifications and Extended View functions are required, you should install the **Server with a GUI** option.

Note
For more information about server installations options, see Windows Server Installation Options.

Uninstalling Role Services

Uninstalling role services with Server Manager must be done manually and on one server at a time. There is no Scenario Based Uninstall that will uninstall role services from multiple remote computers using a single wizard, but role services can be uninstalled from a remote computer by running Server Manager on any server in the deployment.

To uninstall Print and Document Services role services

1. Open Server Manager and click All Servers in the navigation pane.
2. Click **Manage** in the Menu Bar and then click **Remove Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select destination server** page, select the server that you want to remove the role services from and then click **Next**.
5. Select the role services that you would like to remove on the Remove server roles screen and then click **Next**.
6. Select the Features that you would like to remove (if any) on the **Remove features** screen and then click **Next**.
7. Click **Remove** on the **Confirm removal selections** page.
8. Repeat steps 1-7 for each server that you wish to remove role services from.

See Also

Configure Print and Document Services

Install Distributed Scan Server

Fax Server Step-by-Step Guide

Install Distributed Scan Server

This guide describes how to install and configure Distributed Scan Server on a single computer running Windows Server 2012. Distributed Scan Server is a server role included in the Print and Document Services server role. You can use Distributed Scan Server to help you automate

document processes in your organization and provide a central administration point for sharing and managing network scanners.

Important
To install Disturbed Scan Server on a computer, it must be a member of a domain.

For information about Print Servers, see Print and Document Services Overview.

New for Windows Server 2012, the Scan Server now supports Kerberos user authentication for connections to the scan server. In addition, when adding a scan process you have the option to allow the user to enter a document name at the scanner, and the option to send a copy of the scan to the user's email.

Important
The Distributed Scan Server role requires scanners that support Distributed Scan Management (DSM). Windows Server 2012 implements DSM and uses Web Services on Devices (WSD) technologies to integrate scanning devices into the system. For a list of scanners available that support DSM, see Distributed Scan Management (DSM) Scanners. For more information about DSM see Distributed Scan Management (DSM) in the MSDN Technical Library at http://go.microsoft.com/fwlink/?LinkId=204841.

You can install role services using the Add Roles and Features Wizard. However, after you complete the wizard, you must add, share, and configure network scanners separately. As part of installation, associated Microsoft Management Console (MMC) snap-ins, services, and other tools are installed. You can use these tools to perform the additional sharing and configuration tasks.

Important
Distributed Scan Server requires a **Server with a GUI installation**. You cannot install Distributed Scan Server on a Server Core Installation or a Minimal Server Interface installation. For more information about server installation options, see Windows Server Installation Options.

To perform all tasks described in this topic, you must be a member of the Administrators group, or you must have been delegated the appropriate permissions. For more information about configuring user and group permissions, see Change Permissions for Scan Users and Groups for scan servers (http://go.microsoft.com/fwlink/?LinkId=171019).

This guide contains the following sections:

- Step 1: Install software
- Step 2: Configure the server
- Step 3: Add and share scan devices

Step 1: Install software

You can use the Add Roles and Features Wizard in Server Manager to install the Print and Document Services role (needed for the Distributed Scan Server role service) and associated role services, MMC snap-ins, and tools.

The following MMC snap-ins are installed:

- **Print Management**. This snap-in enables you to manage printers, print queues, printer drivers, and printer connections.
- **Scan Management**. This snap-in enables you to manage scanners and scan processes. Scan processes allow you to define how to process scanned documents, and then route them to network folders, SharePoint sites, and to e-mail recipients. Remote management is not supported. To administer a Distributed Scan Server, you must do it from the local console.

  Important
  The Distributed Scan Server role requires scanners that support Distributed Scan Management (DSM). Windows Server 2012 implements DSM and uses Web Services on Devices (WSD) technologies to integrate scanning devices into the system. For more information about DSM see Distributed Scan Management (DSM) in the MSDN Technical Library at http://go.microsoft.com/fwlink/?LinkId=204841.

You can also use these snap-ins to define user and group security permissions for accessing and using network printers and scan processes.

To install the Print and Document Services role

1. Click **Server Manager** on the taskbar.
2. Click **Add roles and features** on the Dashboard.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the Select installation type page, verify that **Role-based or feature-based installation** is selected and click **Next**.
5. On the Select destination server page, verify that your target server is selected and click **Next**.
6. On the **Select server roles** page, select **Print and Document Services** and click **Add Features** and then click **Next**.
7. On the Print and Document Services overview page, click **Next**.
8. On the Select features page, click **Next**.
9. On the Select role services page, click the **Distributed Scan Server** check box and verify that the **Print Server** box is checked, and then click **Next**.
10. On the Confirm installation selections page, click **Install**.
11. Click **Close** to close the wizard and then restart the server.

Step 2: Configure the server

Part of the configuration of the server takes place during installation. However, there are a few issues that might apply to your environment that involve post-configuration, including the following:

- To detect Web Services on Devices (WSD) printers and scanners to view and monitor them on your network, network discovery must be enabled. To detect WSD printers, the Device Association Service must also be running.
- Write and List permissions are needed to run scan processes. The Read permission is needed to read scan processes stored in Active Directory Domain Services (AD DS). You should consider removing any unneeded permissions from the Distributed Scan Server service account.

This section contains the following instructions:

- To detect WSD printers and scanners on the network
- To configure settings for the scan server

To detect WSD printers and scanners on the network

1. To enable network discovery of printers and scanners, right-click the **Start** charm, click **Control Panel**, and then click **Network and Internet**.
2. On the **Network and Internet** page, click **Network and Sharing Center**.
3. On the **Network and Sharing Center** page, click **Change advanced sharing settings**.
4. On the **Advanced sharing settings** page, click the **Domain** drop-down arrow, click **Turn on network discovery**, and then click **Save changes**.
5. Then, to start the Device Association Service, right-click the **Start** charm, click **Computer Management**, under **Services and Applications** click **Services**.
6. In the center pane, right-click **Device Association Service**, and then click **Start**.

To configure settings for the scan server

1. Open Server Manager.

   Important
   The Distributed Scan Server role requires scanners that support Distributed Scan Management (DSM). Windows Server 2008 R2 implements DSM and uses Web Services on Devices (WSD) technologies to integrate scanning devices into the system. For more information about DSM see Distributed Scan Management (DSM) in the MSDN Technical Library at http://go.microsoft.com/fwlink/?LinkId=204841.
2. In the menu bar, click the notification flag. You will see a post-deployment notification for Distributed Scan Server. Click Perform additional configuration. The Scan Management console opens.
3. Ensure that **Scan Servers** is selected in the left pane. In the right pane under Actions, click **Configure local scan server**. The **Scan Server Configuration Wizard** opens.
4. Follow the instructions in the Scan Server Configuration Wizard to change the scan server service account, e-mail server information, temporary scan folder and size, SSL encryption certificate, user authentication (Kerberos or client certificates) for the scan server security option.

Important

If you choose **Require User Authentication** for your security setting option, you must ensure that the Scan Operators group has write access to the scan server's computer object in Active Directory. To do this, use the following procedure:

a. Open Active Directory User and Computers.

b. Click **View** and click **Advanced Features**.

c. Expand your domain, and then expand either the **Computers** or **Domain Controllers** folder depending on where your scan server computer is located.

d. Locate your scan server computer account, then right-click it and click **Properties**.

e. Click the **Security** tab, and then click **Add**.

f. In the **Enter the object names to select** text box, type **scan** and then click **Check Names**. The Scan Operators group should appear in the text box. Click **OK**.

g. In the Permissions for Scan Operators area, click the **Write** check box and then click **OK**.

h. Close Active Directory Users and Computers.

5. After the scan server is configured, download the Active Directory Schema Extensions LDF file to your domain controller if you have a Windows Server 2003 or Windows Server 2008 domain environment. This file is available from the Microsoft Download Center at Windows Server 2008 R2 Distributed Scan Management Active Directory Schema update (http://go.microsoft.com/fwlink/?LinkId=150121). This LDF file extends the AD DS schema to include scan process container objects. Scan process information is stored in AD DS. After you apply this schema, the scan servers you set up in the domain will work with the new schema.

Note

As part of this configuration process, you will need a certificate. You should always select a certificate issued by a certification authority (CA) that is trusted by clients and devices connecting to the scan server. The subject name of the certificate must match the fully qualified domain name (FQDN) of the scan server.

Note

If you define a scan process that delivers scanned documents to a SharePoint site, you must manually enter the SharePoint URL in the wizard, rather than copying and pasting the SharePoint URL from your browser address bar. Also, we recommend that you limit the number of scan processes to a number that is manageable by end users. Defining too many scan processes can be confusing to end users.

Step 3: Add and share scan devices

Now you are ready to add and share devices.

You can use the MMC snap-ins to manage printers and scanners that are located on the same subnet as your server.

You can use the Scan Management snap-in to add scanners that you want to manage on your network. Remote management is not supported. To administer a Scan Server, you must do it from the local console.

Note
To add WSD printers, you must already have network discovery enabled. For instructions, see "To detect WSD printers and scanners on the network" in this document.

To add a network scanner to the scan server

1. Click the **Start** charm, and then click **Scan Management**.

   Important
   The Distributed Scan Server role requires scanners that support Distributed Scan Management (DSM). Windows Server 2012 implements DSM and uses Web Services on Devices (WSD) technologies to integrate scanning devices into the system. For a list of scanners available that support DSM, see Distributed Scan Management (DSM) Scanners. For more information about DSM see Distributed Scan Management (DSM) in the MSDN Technical Library at http://go.microsoft.com/fwlink/?LinkId=204841.

2. In the left pane, click **Scan Management**, right-click **Managed Scanners**, and then click **Manage**.
3. To add a scanner, in the **Add or Remove Scanners** dialog box, type the host name, IP address, or URI of the scanner, and then click **Add**.

   Note
   Scanners must support WSD. TCP/IP and USB scanners are not supported.

Additional resources

For more information, see the following resources on the Microsoft Web site:

- For information about Print and Document Services, see Print and Document Services Overview
- For information about Fax Server, see Fax Server Step-by-Step Guide
- To read and participate in discussions with Microsoft and industry professionals, see the Print/Fax Forum, which is included in the Windows Server Forums (http://go.microsoft.com/fwlink/?LinkId=171028).
- For information about WSD for scanners and printers, see Implementing Web Services on Devices for Printing and Scanning (http://go.microsoft.com/fwlink/?LinkId=144068).

Distributed Scan Management (DSM) Scanners

The following scanners support Microsoft Distributed Scan Management (DSM):

- Ricoh Aficio MP C2800/3300

  For more information, see the Ricoh Americas Corporation website at http://go.microsoft.com/fwlink/?LinkId=209825

Fax Server Step-by-Step Guide

There are two tools you can use to administer a fax server in Windows Server 2012: Server Manager and Fax Service Manager. You can use Server Manager to install the Fax Server server role, the fax printer, and the Fax Service Manager. Server Manager also displays fax-related events from the Event Viewer and includes an instance of the Fax Management snap-in, which can administer the local fax devices only.

Fax Service Manager provides a single interface that administrators can use to efficiently administer fax resources and is the primary focus of this article.

What is a fax server?

A fax server is comprised of four different parts: the Fax Server role, Fax service, Fax Service Manager, and Windows Fax and Scan.

- **Fax Server server role**. To create a fax server, you must install the Fax Server server role from Server Manager. This installs Fax Service Manager, Windows Fax and Scan, the fax printer, and the Fax service. After you install the Fax Server server role, you can access a role page in Server Manager that provides a single point from which you can view fax events, view status information for the Fax service, and access resources and additional information for fax servers. To access this page in Server Manager, click **Fax Server** in the left pane. The role page appears in the right pane.
- **Fax service**. After you install the Fax Server server role, the Fax service appears in the Services snap-in and you can start or stop the service directly from there or from the Fax Server server role page in Server Manager.
- **Fax Service Manager**. Fax Service Manager, a Microsoft Management Console (MMC) snap-in, provides a central administration point for configuring and managing your fax resources. After you install the Fax Server server role, you can access Fax Service Manager from Server Manager by clicking **Tools**, clicking **Fax Service Manager**. You can also access this snap-in by adding it to an MMC console.
- **Windows Fax and Scan**. After you install the Fax Server server role, you can access Windows Fax and Scan by clicking the **Start** charm, and then clicking **Windows Fax and Scan**. On a fax server, you can use Windows Fax and Scan to send faxes, add accounts, and to monitor the incoming fax queue, the inbox, and the outbox. Users who are using

computers running windows client operating systems can use this feature to send faxes and configure fax receipts.

Note

Users who are using computers running versions of Windows XP and Windows Server 2003 can send a fax by using Fax Console. For more information about how to use Fax Console and manage fax servers in earlier versions of Windows, see How to enable and configure the fax service in Windows XP on the Microsoft Support Web site at http://go.microsoft.com/fwlink/?LinkId=90750.

Who should use a fax server?

This guide is targeted at the following audiences:

- Fax administrators and help desk professionals.
- Information technology (IT) planners and analysts who are evaluating the product.
- Enterprise IT planners and designers.

Benefits of a fax server

The Fax Server server role can save the fax administrator a significant amount of time. Instead of installing fax machines on client computers and managing and monitoring fax machines, you can complete those tasks in 2 or 3 steps on multiple computers simultaneously and remotely.

By using Fax Service Manager, you can automatically make fax connections available to users and computers in your organization.

Requirements for using the Fax service

All of the fax-related parts discussed in the previous section are available in all versions of Windows Server 2012 and are compatible with its basic system requirements. To determine your fax capacity needs, use performance counters, testing in the lab, data from existing hardware in a production environment, and pilot rollouts.

Important

Fax Server requires a **Server with a GUI installation**. You cannot install Fax Server on a Server Core Installation or a Minimal Server Interface installation. For more information about server installation options, see Windows Server Installation Options.

Deploying a fax server

The following sections provide information about how to deploy a fax server and attach fax devices:

Step 1: Install and open Fax Service Manager

Step 2: Install and share a fax printer

Step 3: Add and remove fax devices
Step 4: Configure fax devices to send and receive
Step 5: Configure fax settings
Step 6: Set up user accounts

Step 1: Install and open Fax Service Manager

To use Fax Service Manager, you must install the Fax Server server role.

Before you begin installing the Fax Server server role, make sure that any modem devices have been installed on the server. If you plan to install a new modem device, you can save time by installing it before you set up the Fax Server role. We recommend that you install the Fax Server role locally and not by using a Remote Desktop connection. You can install the Fax Server role remotely, but you need to make sure that local resource sharing is turned off. Remote management is not supported. To administer a Fax Server, you must do it from the local console.

To install the Fax Server server role

1. In the task bar, click **Server Manager**.

   In the menu bar, click **Manage**, and then click **Add Roles and Features**.
2. On the Before you begin page, read the text and click **Next**.
3. On the Select installation type page, ensure **Role-based or feature-based installation** is selected and click **Next**.
4. On the Select destination server page, ensure the target server is selected and click **Next**.
5. On the **Select Server Roles** page, select the **Fax Server** check box. When you do this, a dialog box displays the tools and services that also need to be installed with the Fax Server server role. Click **Add Features**, then click **Next** on the Select server roles page.
6. On the Select features page, click **Next**.
7. On the Fax Server page, read the text and then click **Next**.
8. On the Confirm installation selections page, click **Install**.
9. When the installation completes, click **Close** and restart the server.
10. From Server Manager, click the Notifications flag, and click **Perform additional configuration** in the Fax Server notification. The Microsoft Fax Service Manager opens.
11. Right-click **Accounts**, select **New**, and click **Account**.
12. Add any users that you want to allow to send and receive faxes using this server.

    Note
    The members of the local Administrators group are added automatically to the Fax Users group and cannot be removed. (The local Administrators group is also separately granted full permissions for the Fax service.) Membership in this group can be modified later by using Fax Service Manager.
13. Right-click the **Start** charm, and click **Control Panel**.

14. In the Control Panel search box, type **phone** and then click **Phone and Modem**.
15. In the **Location Information** dialog, type the information for your country/region, area/city code, carrier code, information to dial an outside line, and whether you use tone or pulse dialing.

Step 2: Install and share a fax printer

When you install the Fax Server server role, if you have not attached a fax device to the computer, a local fax printer connection, **Fax**, is automatically created in the **Printers** folder in Control Panel. Follow the steps in this procedure to share the printer so that users can connect.

This procedure describes how to create a new fax printer, and then share the printer so that your users can find it.

Note

To install fax devices, such as a fax printer, you must be a member of the Administrators group or you must have been delegated the appropriate authority.

To create and share a new fax printer

1. Click the **Start** charm, right-click the **Windows Fax and Scan** tile, and then click **Run as administrator**.
2. Click **Tools**, and then click **Fax Accounts**.
3. In Fax Accounts, click **Add** to open Fax Setup.
4. On the **Choose a fax modem or server** page, click **Connect to a fax modem**.
5. You might be asked to install a printer. To do this, follow the instructions in the Add Hardware Wizard.
6. On the **Choose a modem name page**, type a name for the fax printer, and then click **Next**. The default name is **Fax Modem**.
7. On the **Choose how to receive faxes** page, click the option that you prefer.
8. The new fax printer should appear in **Fax Accounts**, under **Account Name**.
9. Then, to share the printer that has been created, right-click the **Start** charm, and then click **Control Panel**.
10. Under **Hardware**, click **View devices and printers**.
11. In the list of printers, right-click **Fax**, click **Printer properties**, click the **Sharing** tab, select **Share this printer**, and then type a name for the printer that you want your network users to see.
12. If you want to enable users with computers running different versions of Windows to use this printer, click **Additional Drivers** to install the needed drivers.
13. In **Additional Drivers**, select the check box for the versions of Windows that you want to support. You are prompted to provide a path to the driver. Provide the path to the %windir%\System32\DriverStore\FileRepository\prnms002.inf* file on a computer that is running the version of Windows that you want to support.
14. To confirm that the files were copied properly, in Windows Explorer navigate to

%windir%\System32\spool\drivers\ and look for the folder that contains the files for the selected architectures.

Important

Because the Fax service uses remote procedure calls (RPCs), you must ensure that the firewall on the server will allow RPCs. Typically, when you add an account to the Fax service, Fax Setup in Windows Fax and Scan will ask the user to unblock RPC so that the firewall is unblocked for fax communication. You can also configure the firewall directly at any other time in Windows Firewall by verifying that **Windows Fax Service** is in the exceptions list. If you are sharing the Fax service, make sure that **File and Printer Sharing** are also added to the exceptions list. For more information see Windows Firewall with Advanced Security Overview.

Step 3: Add and remove fax devices

Windows recognizes any newly installed Plug and Play fax device during service startup. After the device is installed, restart the Fax service and refresh the display. The new device will automatically appear in Fax Service Manager. To remove a device, remove it as you would any device using Device Manager, restart the Fax service, and refresh the display.

To refresh the Device list and view the updated display

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Stop**.
3. In the left pane, right-click **Fax**, and then click **Start**.
4. Press **F5** to refresh the display.
5. Double-click **Devices and Providers**, and then click **Devices**.
6. In the right pane, view the list of devices.

Step 4: Configure fax devices to send and receive

By default, a fax device is enabled to send when it is installed. However, if needed, you can also enable a device to send later. As part of this process, you can specify an optional Transmitting Subscriber ID (TSID) string, a text string that identifies the device sending the fax. The identification information typically appears in the banner in a fax message to help the recipient determine where the fax originated.

You must specifically enable a fax device to receive—it is not enabled by default. A fax device can be configured to answer incoming calls automatically or manually. (The manual option is useful if your fax and telephone share the same incoming line.) And you can specify the number of rings before answering. (A higher number of rings before answering is useful if your fax and telephone share the same incoming line, to give you time to answer the telephone.)

As part of this process, you can also specify an optional Called Subscriber ID (CSID) string, a text string used to identify a fax recipient to a fax sender.

Important

To configure fax devices, such as a fax printer, you must be a member of the Administrators group or you must have been delegated the appropriate authority.

To enable a fax device to send

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Devices and Providers**, and then click **Devices**.
3. In the right pane, right-click the device you want to configure, and then click **Send**.
4. To configure a Transmitting Subscriber ID (TSID), right-click the device you want to configure, and then click **Properties**.
5. On the **General** tab, in **Transmitting Subscriber ID (TSID)**, type a TSID string. This string is limited to 20 characters.

Note

Later, if you want to disable the fax device from sending faxes, right-click the device, and then click **Send** to clear the check mark.

To enable a fax device to receive

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Devices and Providers**, and then click **Devices**.
3. In the right pane, right-click the device you want to configure, click **Properties**, and then click the **General** tab.
4. To configure whether calls are answered automatically or manually, select the Receive faxes check box, and then do one of the following:
   - Click **Automatic Answer** to enable the fax device to answer calls automatically, and then type a number for the value in the **Rings before answering** field. The default number of rings is 5. If your fax and telephone do not share a line, you should decrease this number. If your fax and telephone do share a line, and you decide to increase this number, do not increase it so much that the fax sender stops transmission.
   - Click **Manual Answer** to enable the fax device to answer calls manually.
5. To configure a Called Subscriber ID (CSID), in **Called Subscriber ID (CSID)**, type a CSID string. This string is limited to 20 characters.

Important

Later, if you want to disable the fax device from receiving faxes, right-click the device, and then clear either the **Auto Receive** or **Manual Receive** check box.

Step 5: Configure fax settings

You can use Fax Service Manager to configure fax settings, including routing rules and Simple Mail Transfer Protocol (SMTP) settings. For outgoing faxes, you can also configure receipts,

whether you will enable your users to create and use personal cover pages, and what your organization will include in the banners of sent faxes.

Configure incoming fax settings

Fax administrators have full control over the incoming fax queue. This queue is the collective queue for all receiving fax devices that are managed by the fax server—it corresponds to the Incoming folder. Users who are not members of the Administrators group need to be assigned permissions to view or manage the incoming fax queue.

A fax remains in the incoming fax queue until it is received successfully and routed in accordance with the policy specified by the fax administrator. Only then is it moved to the inbox. If a fax is not routed successfully for some reason, it remains in the incoming fax queue until it is deleted manually—or it is handled in accordance with the automatic deletion policy that you have configured for the fax queue.

You can control incoming fax jobs handled by the Fax service by disabling the reception of faxes into the incoming fax queue. If the reception of a fax is stopped while being received, but a portion of the fax is successfully received, the fax is assigned the **Partially received** status and is moved to the Inbox folder.

Note
You can view the Incoming and Inbox folders by using Windows Fax and Scan.

Start or stop fax reception

You can start or stop the reception of faxes for all devices on a fax server.

To start or stop fax reception

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. On the **General** tab, select **Disable reception of new faxes** to block new faxes from being added to the incoming fax queue.

Configure incoming fax routing

Administrators can configure incoming fax routing methods to route incoming faxes to recipients on the network. In Fax Service Manager, there are incoming fax routing extensions that are global (applied to all devices) and others that are associated with just individual fax devices. For global methods, you can set the priority order in which they are applied to an incoming fax. Non-Microsoft vendors can extend the global methods for incoming faxes. Extensions could include, for example, converting incoming files to specific file formats.

Individual incoming fax routing methods are configured per device. After a method is configured, it can then be enabled or disabled. More than one incoming fax routing method can be applied to incoming faxes, and if a method is disabled, the device skips that method and processes the incoming fax in order of global-level priority. A disabled incoming fax routing method appears with an icon (an arrow pointing down) over it in the right pane.

The following default incoming fax routing methods can be configured and prioritized:

- **Route through e-mail**. Specify the e-mail address for receiving incoming faxes. To change SMTP server or authentication settings, use Fax Service Manager for configuring delivery receipts for sent faxes.
- **Store in a folder**. Specify the local or network path of the folder to which a copy of incoming faxes will be saved. If you choose this method, make sure that the folder you use has access to the Network Service account.
- **Print**. Specify the Universal Naming Convention (UNC) path of the printer location (\\<*network_location*>\<*printername*>) to which incoming faxes will be printed.

You can use incoming fax routing methods to control what happens to faxes that are received by a fax device managed by the fax server. You can use any combination of routing methods to forward received faxes to an e-mail address, store faxes in a folder, or print faxes.

To configure incoming fax routing methods

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Devices and Providers**, double-click **Devices**, double-click the device you want to configure, and then click **Incoming Methods**.
3. In the right pane, right-click the method you want to configure, click **Properties**, and then do the following:
    a. To configure the **Store in a folder** method, click the **Store in Folder** tab, and then either type the UNC path of the folder (\\<*network_location*>\<*folder*>), or click **Browse** to select the folder.
    b. To configure the **Route through e-mail** method, click the **E-mail** tab, and then type the e-mail address to which you want incoming faxes to be delivered.
    c. For the **Print** method, click the **Print** tab, and then type the UNC path of the printer for incoming faxes.
4. After the incoming methods have been configured, in the details pane, right-click each method, and then click **Enable** to enable the method, or **Disable** to disable it. (You might want to disable a method if, for example, you want to temporarily disable printing because the printer is not working and faxes are getting stuck in the inbox.)

Configure SMTP settings for e-mail routing

You should use this procedure if you have chosen to configure the **Route through e-mail** method for a fax device.

To configure SMTP settings for e-mail routing

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. To enable incoming faxes to be routed to recipients by e-mail, click the **Receipts** tab and then do the following:
    a. Select the **Use these SMTP settings for the Route Through E-mail incoming**

routing method check box.

b. In **From e-mail address**, type the e-mail address to which all incoming faxes will be received.

c. In **Server address**, type the address of the SMTP server to which incoming faxes will be received.

d. In **Port**, type the port number.

4. Also on the **Receipts** tab, click **Authentication**, and then, based on how your SMTP server is set up, do one of the following:

- Click **Anonymous access** if this is sufficient to access the SMTP server.
- Click **Basic authentication**, click **Credentials**, and then type a user name and password.
- Click **Integrated Windows Authentication** if NTLM authentication is needed to access the SMTP server, click **Credentials**, and then type a user name and password.

Configuring outgoing fax settings

You can use Fax Service Manager to configure and manage how faxes are sent and archived, what is included as part of the fax transmission, whether receipts are issued, and rules for routing.

Note

You can monitor outgoing faxes in the Outbox folder in Windows Fax and Scan.

Start or stop outgoing faxes

You can start or stop sending faxes for all devices on a fax server.

To start or stop outgoing faxes

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. On the **General** tab, do the following:

a. Select **Disable submission of new outgoing faxes** to prevent users from submitting new faxes to the Outgoing folder. Clear this check box to restart fax submissions.

b. Select **Disable transmission of outgoing faxes** to prevent faxes from being sent from the Outgoing folder. Clear this check box to restart fax transmissions.

Configure the outbox and outgoing transmissions

The outbox folder is the queue of all faxes that are waiting to be sent or are in the process of being sent. If you send faxes using multiple fax devices, the outbox is the collective queue for all these devices. After a fax is sent successfully, it is moved to the archive.

You can view the Outbox folder by using Windows Fax and Scan.

Administrators can use Fax Service Manager to configure the following properties for the outbox and outgoing faxes:

- **Include banner**. You can add a banner to outgoing faxes. A banner is text that is added to each page of a transmitted fax. The banner format cannot be modified. However, if all the banner information does not fit, a shorter version appears. Information in a banner includes:
  - Date and time of transmission
  - TSID of sender
  - Recipient phone number (full version of banner only)
  - Page of total pages
- **Allow use of personal cover pages**. You can enable users to attach custom cover pages to outgoing faxes. If you do not enable this feature, then users connecting to a fax server with a remote fax printer connection can only choose cover page templates located on the fax server computer.
- **Use device Transmitting Subscriber ID (TSID)**. You can specify that the TSID specified for the device should be used. If you do not enable this option, the TSID used in the outgoing fax is the sender's fax number, as specified on the **Sender Information** properties page. However, if you send a fax programmatically, the TSID parameter specified in the application programming interface (API) will be used instead of the sender's fax number.
- **Number of retries**. You can set the number of retries for attempting to resend faxes following an unsuccessful connection attempt or transmission failure. If the maximum number of retries is reached, the outgoing fax remains in the outbox until it is deleted or until you restart the failed transmission.
- **Retry after**. You can set the number of minutes between attempts to retry fax transmission. When the maximum number of retries has been exhausted, the user can restart the transmission, which resets the transmission attempts.
- **Discount rate start and stop**. You can specify times when discount rates for the fax line apply and send faxes during these times. A discount rate fax job can never start transmission outside the specified discount period. However, for a long fax, a discount rate fax job might have some part of its transmission outside the specified discount period.
- **Automatically delete faxes older than**. You can specify that failed faxes in the fax queue should be automatically deleted periodically. The fax queue consists of the Outbox folder and Incoming folder.

To configure the outbox and outgoing fax transmissions

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. On the **Outbox** tab, do the following:
   a. Select **Include banner** to print transmission information on the edge of each page of outgoing faxes.
   b. Select **Allow use of personal cover pages** to allow users to use cover pages from their personal cover pages folder when sending faxes.
   c. Select **Use device Transmitting Subscriber ID (TSID)** to include identification

information about a fax device to a fax recipient.

d. In **Number of retries**, enter the number of times to attempt to resend a failed outgoing fax.

e. In **Retry after**, enter the number of minutes to wait before trying to resend a failed outgoing fax.

f. In **Discount rate start**, specify the time at which discount line rates will be taken into account when sending faxes.

g. In **Discount rate stop**, specify the time at which discount line rates will no longer be taken into account when sending faxes.

h. Select **Automatically delete faxes older than** <*number*> **days** to set a duration to automatically delete failed faxes that are not sent and remain in the outbox.

Configure fax delivery receipts

Delivery notification receipts provide users with information about the success or failure of their sent faxes. The following options can be configured for delivery receipts:

- **Enable SMTP e-mail receipts delivery**. An e-mail receipt is sent to the user who sent the fax, to confirm whether the fax has been sent correctly. To configure this option, you must specify an e-mail address from which receipts will be sent to fax senders, the IP address of the SMTP server, and the port number of the SMTP service. You must also select the appropriate authentication method for accessing the SMTP server. The following authentication methods can be chosen:
  - **Anonymous access**. All users can access the SMTP server, which in effect disables authentication.
  - **Basic authentication**. You can specify a user name and password for access to the SMTP server. The password is sent over the network in standard text.
  - **Integrated Windows Authentication**. You can specify a user name and password for access to the SMTP server using Integrated Windows Authentication.

After this receipt option is enabled and configured, a user who sends a fax by using Windows Fax and Scan (or using Fax Console in Windows XP) can request an e-mail to confirm that the fax was sent without error.

To configure delivery receipts for sent faxes

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. On the **Receipts** tab, select the **Enable SMTP e-mail receipts delivery** check box and then do the following:
   a. In **From e-mail address**, type the e-mail address from which the receipt will be sent.
   b. In **Server address**, type the IP address of the SMTP server from which the receipts will be sent.
   c. In **Port**, type the port number.
4. Also on the **Receipts** tab, click **Authentication**, and then, based on how your SMTP

server is set up, do one of the following:

- Click **Anonymous access** if this is sufficient to access the SMTP server.
- Click **Basic authentication**, click **Credentials**, and then type a user name and password.
- Click **Integrated Windows Authentication** if NTLM authentication is needed to access the SMTP server, click **Credentials**, and then type a user name and password.

Configure outgoing routing

You can use Outgoing Routing in Fax Service Manager to configure routing rules for outgoing faxes to optimize use of available fax devices. To do this, you can create rules to associate a device or group of devices with faxes sent to a specific domestic area code or to a specific country/region.

Create a fax device group

You can use Fax Service Manager to create groups of one or more fax devices, and then associate outgoing rules with these groups. For example, you can designate specific fax devices to send faxes to only specific destinations.

By default, there is a group, **All devices**, that cannot be deleted. You can add or remove groups, or devices from groups, at any time. A device can also belong to one or more groups simultaneously.

When you create a new group, it is not active until you add a fax device to it. Groups that do not contain devices display a mark over the group folder icon.

Outgoing faxes are routed to groups according to the outgoing routing rules. Fax Service Manager notes the destination of the fax, routes it to the appropriate group, and accesses the first available device in that group for transmission.

To create or delete a fax device group

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Fax**, double-click **Outgoing Routing**, right-click **Groups**, point to **New**, and then click **Group**.
3. In the **Add New Group** dialog box, type a group name then click **OK**.
4. To later delete the group, right-click the group name, and then click **Delete**.

Add or remove a fax device in a group

To perform this procedure, you must first have created a group. You can add or remove devices from groups at any time. A device can also belong to one or more groups simultaneously. Fax devices in a group can be sorted into priority order, so that the highest priority device in a group will attempt to send a fax first. If it is busy, the next highest priority device will attempt to send.

To add or remove a fax device in a group

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Fax**, double-click **Outgoing Routing**, and then double-click **Groups**.
3. Right-click the group to which you want to add a device, point to **New**, and then click **Device(s)**.
4. In **Add New Device(s)**, click the device that you want to add, and then click **OK**.
5. To change the priority order of devices within a group, in the right pane, right-click a device, click **Move Up** or **Move Down**, and then click **Yes** to confirm.
6. To later remove the device from the group, in the right pane, right-click the device, and then click **Delete**.

Add or delete a fax routing rule

You can use Fax Service Manager to create rules for routing outgoing faxes. For example, you can create a group for local destinations and another for domestic long distance destinations, or a group for a specific foreign destination.

Using these rules helps you ensure that your fax devices are used efficiently, and outgoing faxes do not spend idle time in queues. If all of the devices in a group are constantly busy, or if faxes accumulate within queues, check the group and rule configurations.

To add or delete a routing rule for outgoing faxes

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Fax**, and then double-click **Outgoing Routing**.
3. Right-click **Rules**, point to **New**, and then click **Rule**.
4. In **Add New Rule**, if you know the country/region code, type it in **Country/region code**. If you are not sure of the code, click **Select**. In the list, click a country/region, and then click **OK**.
5. In **Area code**, select **Specific** to limit the rule within the selected country/region, and then type the area code, or select **All areas** to apply the rule to the entire country or region.
6. In **Target device**, if you want the rule to apply to a specific device or routing group, do one of the following:
   - If you want the rule to apply to a device, select **Device**, and then select a device from the list.
   - If you want the rule to apply to a routing group, select **Routing group**, and then select a group from the list.

Step 6: Set up user accounts

User accounts

Windows Server 2012 provides fax-related user accounts for more privacy and better organization and management of faxes. Users can use accounts for accessing different types of fax services such as the local fax modem and a fax server.

Your users must have an account to be able to access a fax server running Windows Server 2012. You can configure a fax server so that when a user uses Windows Fax and Scan for the first time to send a fax, an account is automatically created. If you use this option, as soon as a user connects to the fax server, the account is automatically created.

If you want to limit users connecting to the fax server, you can disable this setting, in which case you must manually create the user accounts for all the users who need to access the fax server. In such a setting, if a user tries to connect to the server without having an account already, the connection attempt fails.

To automatically create a user account for a fax server

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, right-click **Fax**, and then click **Properties**.
3. On the **Accounts** tab, select the **Auto-create accounts on connection** check box so that an Administrator does not need to explicitly create each user account.
4. In **Reassign Setting**, do one of the following:
   - Click **On** if you want messages that are received to be reassigned to individual accounts.
   - Click **Off** if you want to allow all accounts to access the fax server inbox and view incoming messages.

You can also add or delete user accounts directly.

To manually add a user account for a fax server

1. Open Server Manager, click **Tools**, and then click **Fax Service Manager**.
2. In the left pane, double-click **Fax**, and then click **Accounts**.
3. In the right pane, view the existing accounts, and then determine which new accounts you want to add.
4. To add a new account, right-click **Accounts**, click **New**, and then click **Account**.
5. In **Create New Account**, provide a user name and domain, and then click **OK**.
6. To later delete an account, right-click the account, and then click **Delete**.

Additional Resources

For more information about Print and Document Services, see the following resources on the Microsoft Web site:

- Print and Document Services Overview

Configure Print and Document Services

This article describes how to configure and manage Print and Document Services using both the Print Management Console and the associated Windows PowerShell cmdlets. Management of a Windows Server 2012 Print Server is done using the Print Management Console (PMC) MMC snap-in, which has been updated to support the new Windows Server 2012 features. You can remotely manage your print server, but you need to first configure it from the Print Server's local console.

Internet Printing is also managed using the PMC, and LPD Printing uses Windows PowerShell tools for configuration and remains unchanged functionally from Windows Server 2008 R2.

Important

The LPD and LPR Services are deprecated starting with Windows Server 2012. Eventually, they will be completely removed from the product, but they are still available in this release. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features. For more information about features or functionalities that have either been removed from the product in the current release or are planned for potential removal in subsequent releases, see Features Removed or Deprecated in Windows Server 2012.

Windows Server 2012 introduces a new set of Windows PowerShell cmdlets that can be used to manage print servers and print queues from the command line. Wherever possible, examples on how to use these cmdlets will be shown for common print management scenarios. For more information, see Print Management Windows PowerShell cmdlets.

In this document

- Step 1: Install v4 drivers
- Step 2: Install v3 drivers (if necessary)
- Step 3: Create a shared print queue
- Step 4: Connect to the print queue
- Step 5: Print from Windows
- Step 6: Print from a Windows app

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Install v4 drivers

Using v4 drivers with Windows Server 2012 Print Servers is recommended wherever possible. If Windows 8 client computers are being used, users and IT administrators will receive the new

printing experience and receive all of the benefits of the v4 driver model. For some devices, using a v3 driver will be necessary and differences between Windows Server 2012 and previous versions of Windows will also be discussed.

Do this step using Windows PowerShell

To install v4 drivers using the Print Management Console

1. Open the Print Management Console by opening Server Manager, click **Tools**, and then click **Print Management**.
2. Expand **Print Servers**, and then expand the Print Server name. Right click **Drivers** and select **Add Drivers**.
3. To add a v4 driver for a device, select the driver that has v4 or Class Driver in the name.
4. After the driver is installed, it should display in the list of installed drivers.

Figure 1: Installed Drivers

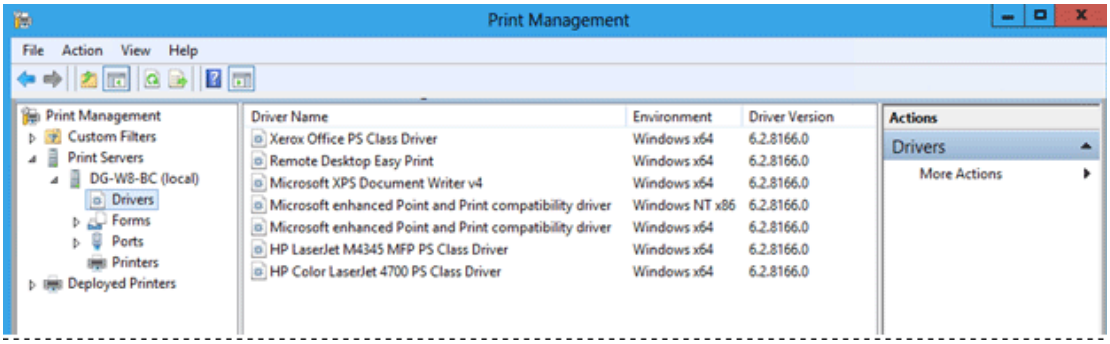

Once installed, V4 drivers are identified by the Version field displayed in the Driver Properties:

Figure 2: Driver Properties

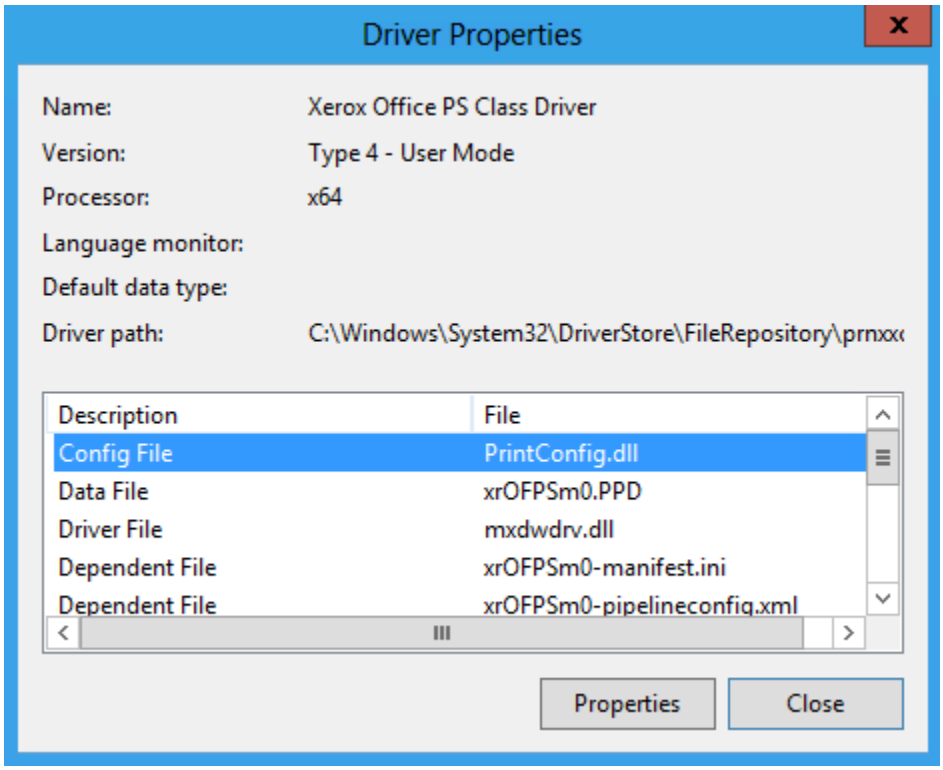

The driver name will state **Class Driver**, the **Config File** should show **PrintConfig.dll**, and the driver path should be `%systemroot%\system32\DriverStore`.

Class Drivers

V4 drivers that ship with Windows Server 2012 are known as Class Drivers. Drivers of this type should always display **Class Driver** in the name.

Model Specific Drivers

V4 drivers that are downloaded directly from a printer manufacturer website or downloaded from Windows Update are known as model specific drivers.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Add-PrinterDriver -Name "HP Color LaserJet 5550 PS Class Driver”
```

Step 2: Install v3 drivers (if necessary)

In many cases it will be necessary to use v3 drivers on a Windows Server 2012 Print Server. In these cases, these drivers will install and share exactly as they did before in Windows Server 2008 R2. To make installation and maintenance of these drivers easier than they were in

previous operating systems, the 32-bit core driver files and the ntprint.inf driver definition file from 32-bit Windows are included with Windows Server 2012. If a 32-bit driver has a dependency on a core driver, an administrator can install a 32-bit v3 driver without needing to use an installed 32-bit operating system as a source for the ntprint.inf file.

Step 3: Create a shared print queue

Creating shared print queues in Windows Server 2012 follows the same process as Windows Server 2008 R2. A local printer must be installed on the server and connected to a port and a driver, and the print queue must be shared in order for client computers to discover and connect to it. When using a v4 driver, an administrator does not need to install a driver for any additional processor architectures.

Windows PowerShell equivalent commands

```
Add-printer -Name "HP Color LaserJet 4700" -DriverName "HP Color LaserJet 5550 PS Class Driver" -shared -ShareName "HP Color LaserJet 4700" -PortName "192.168.100.100"
```

Important

The Port and Driver must already exist or the Add-Printer command will fail. Use the `Add-PrinterDriver` and `Add-PrinterPort` cmdlets to install the driver and the port before running the `Add-Printer` command.

Step 4: Connect to the print queue

When connecting to a print queue using a v4 driver from a client computer, the client computer will install the print queue using the enhanced Point and Print process described previously. If the client computer is Windows 8, the driver will be downloaded from either from the local driver store, Windows Update, or Windows Software Update Services (WSUS).

Windows 8 and v4 drivers

The user will not be prompted to provide credentials to install the v4 print driver. The v4 driver installs from the local driver store and the user will see a printer installation indicator with a status bar. After the installation completes, the printer is ready to use with no further action needed.

Microsoft Enhanced Point and Print Compatibility Driver

A client running a previous version of Windows can install and print to a Windows Server 2012 print queue that is using a v4 driver by using the Microsoft enhanced Point and Print Compatibility Driver. Windows 7 clients do not support the v4 driver model so they need to download and install the Microsoft enhanced Point and Print Compatibility Driver from the Windows Server 2012 Print Server. Windows Server 2012 comes with both a 64-bit and 32-bit compatibility driver pre-installed.

After the printer is installed, its properties will show the driver as the Microsoft enhanced Point and Print compatibility driver instead of the v4 Class Driver or Model Specific driver that is installed on the server.

Figure 3: Enhanced Point and Print Compatibility Driver

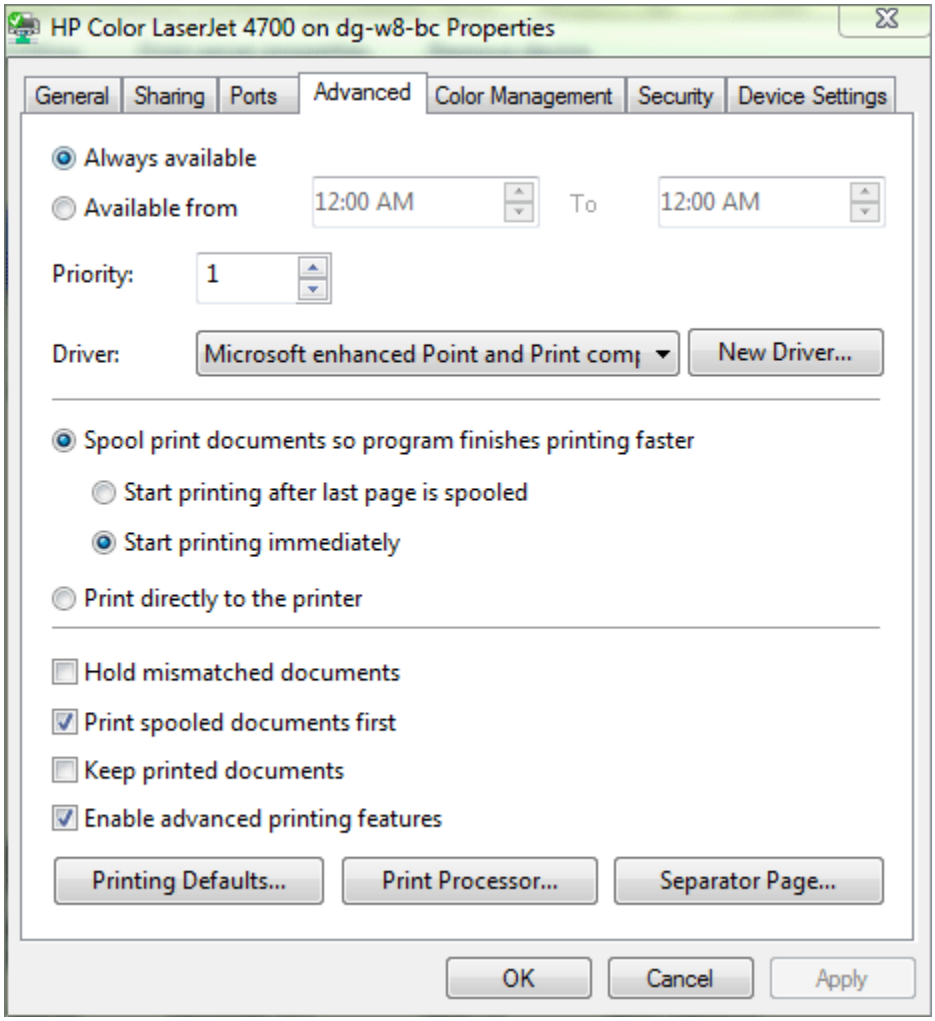

The Microsoft enhanced Point and Print Compatibility Driver works cross-platform as well, as a 32-bit Windows 7 client downloads and installs the 32-bit version of the driver from the server running Windows Server 2012 and uses this driver to print to all v4 print queues.

Step 5: Print from Windows

The Windows 8 Shell includes a new user interface that is designed to help users easily discover and install devices. Print, Fax, and Scan devices are installed from the Settings or Devices charms:

Figure 4: Charms Bar with Devices and Settings

To install a Printer using the Windows 8 Shell

1. Activate the Charms bar by swiping in from the right edge of a touch screen or moving the mouse to the upper or lower right corner of the display.
2. Click or touch **Settings**, and then **More PC Settings**.
3. Click or touch **Devices**, and then click or touch **Add a device**.
4. Windows scans the network for devices and displays a list of printers that were found. If the administrator published printers to the Active Directory then they should show up at the top of the list.

Figure 5: Searching for Devices

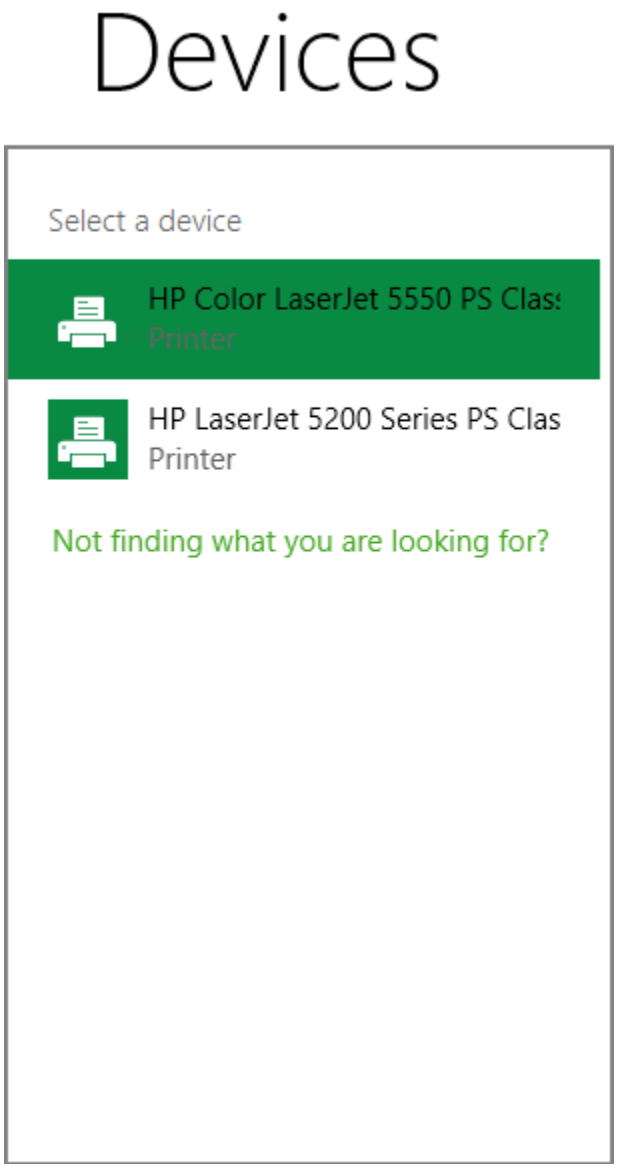

5. Click or touch the device that you want to install and it will be added to the list of devices. If the printer is using a v4 driver then no other user interaction is necessary to install the device.

Note

In some cases you may need to add a printer through the Control Panel, Devices and Printers, Add Printer user interface. This works the same way as in previous Windows versions.

Important

When printing to a Windows Server 2012 print queue, only users with administrator permissions on the print queue (full system administrator or delegated print administrator, see Assign Delegated Print Administrator and Printer Permission Settings in Windows Server 2012) will see readable document names when viewing the print queue. Users can still view the queue from their client computer and will see their own document name in full, but for other user's print jobs they will see **Print Document** for the document name.

Additionally, in most cases in the event log on the server, the document name will also appear as **Print Document**. However, in some job failure cases, the document name will be listed in the failure event.

Step 6: Print from a Windows app

Windows 8 has a new user interface for discovering and printing to print devices on a network. A Windows app that runs in the Windows 8 modern shell may have a custom printing interface that is accessible via the application bar or it may provide a printing link similar to the Internet Explorer app. Additionally, the `CTRL + P` hotkey will activate the new print dialog from any Windows app.

Note
Printing from desktop applications remains unchanged and works the same as previous Windows versions.

To print a document from a Windows app

1. Click or tap the Devices charm or use the `CTRL + P` hotkey to activate the Printing interface.
2. Choose a device to print to by clicking or touching the Print device icon:

Figure 6: Choosing a print device

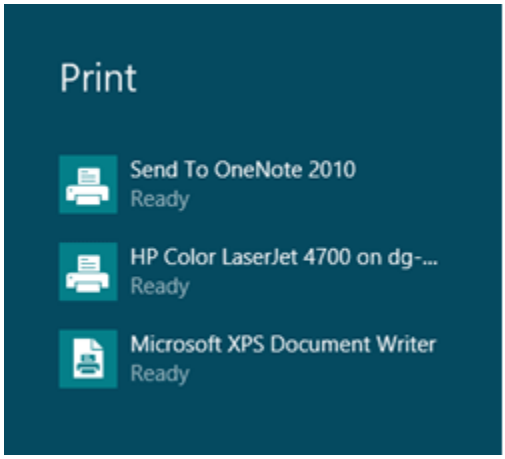

3. A print preview and basic print settings are displayed. The document can be printed using the **Print** button:

Figure 7: Basic print settings

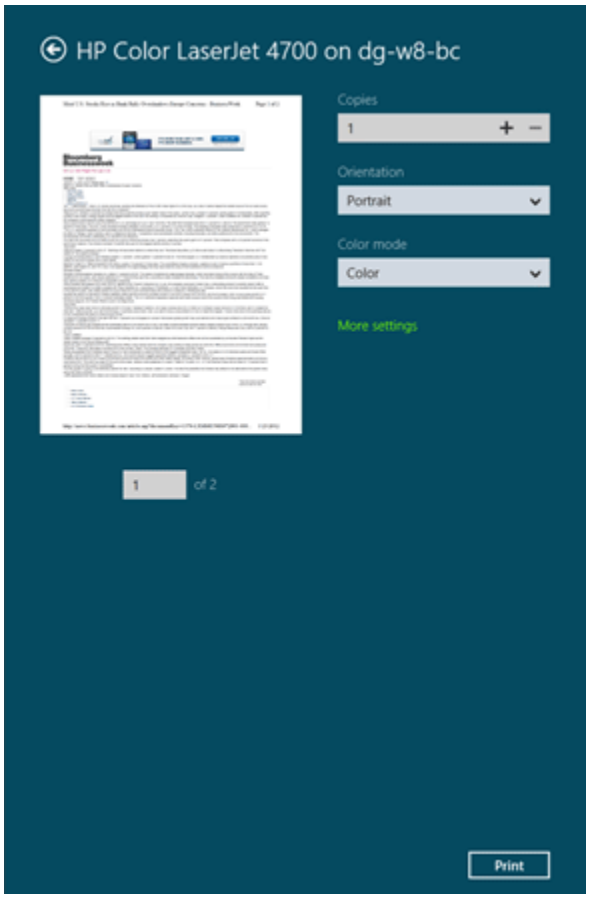

4. Click **More Settings** to activate the advanced print settings dialog where page layout, paper and quality, and output options can be specified:

Figure 8: More Settings

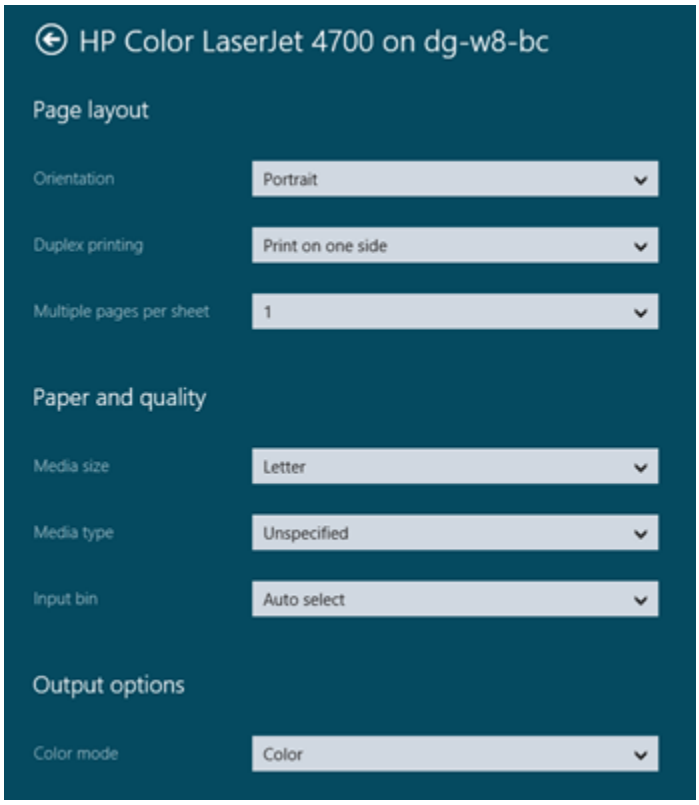

See also

- Create Custom Separator Pages in Windows Server 2012
- Assign Delegated Print Administrator and Printer Permission Settings in Windows Server 2012
- Print and Document Services Overview
- Print and Document Services Architecture
- Install Print and Document Services
- Print Management Windows PowerShell cmdlets

Create Custom Separator Pages in Windows Server 2012

Windows allows you to specify a separator page file to print at the start of each print job. Several separator page files are included with Windows in the **%windir%\system32** folder. In addition, you can create custom separator page files to suit your individual needs.

In a Branch Office Direct Printing environment, the specified separator page is copied automatically from the print server to the print client if the file size is less than 64 KB. The print client then uses the separator page when printing directly to the local printer. For more information about Branch Office Direct Printing, see Branch Office Direct Printing Overview.

Create a custom separator page file

Use a text editor such as Notepad to create a custom separator page file. On the first line, type a single character, and then press ENTER. The character on this line defines the character used as an escape character. For example, the examples in this article assume that the escape character is the forward slash character (/).

Use the following procedure to create a custom separator page file.

To create a custom separator page file

1. Open your separator page file.
2. Type your escape character on the first line.
3. Type the escape codes for the functions you want, and then save the file with an .sep extension in the Windows **SYSTEM32** subfolder.
4. In Print Management, right-click the printer that you want to use the separator page file with, and then click **Properties**.
5. Click the **Advanced** tab, and then click **Separator Page**.
6. Type the path for your separator page file, or click **Browse** to browse for the file and then click **OK**.
7. Click **OK** to close the Properties page.

The following table describes the escape codes and their functions that can be used in a separator page file:

/J	Prints the domain name of the person that submitted the job. Available only to print queues located on Windows Server 2012 or Windows 8. Previous operating systems will ignore this escape code. In a Branch Office Direct Printing environment, /J only works if the print job originates from a computer running Windows Server 2012 or Windows 8. For more information about Branch Office Direct Printing, see Branch Office Direct Printing Overview.
/N	Prints the user name of the person that

	submitted the job
/I	Prints the job number.
/D	Prints the date the job was printed. The representation of the date is the same as the Date Format in the International section in Control Panel.
/T	Prints the time the job was printed. The representation of the time is the same as the Time Format in the International section in Control Panel.
/L**xxxx**	Prints all the characters (**xxxx**) following it until another escape code is encountered.
/F**pathname**	Prints the contents of the file specified by **pathname**, starting on an empty line. The contents of this file are copied directly to the printer without any processing. The file will print only if the path specifies a file in one of the following locations: • In the **%WINDIR%\System32** folder. • In a **SepFiles** folder that is under the **%WINDIR%\system32** folder, or in any subfolder of that **SepFiles** folder. • In the same **SepFiles** folder as the custom separator page file, or in any subfolder of that **SepFiles** folder. **Note** In a Branch Office Direct Printing environment this escape code is ignored.
/H**nn**	Sets a printer-specific control sequence, where **nn** is a hexadecimal ASCII code sent directly to the printer. To determine the specific numbers, see your printer manual.
/W**nn**	Sets the width of the separator page. The default width is 80; the maximum width is 256. Any printable characters beyond this width are truncated.
/U	Turns off block character printing.

/B/S	Prints text in single-width block characters until @U is encountered.
/B/M	Prints text in double-width block characters until /U is encountered.
/E	Ejects a page from the printer. Use this code to start a new separator page or to end the separator page file. If you get an extra blank separator page when you print, remove this code from your separator page file.
/**n**	Skips **n** number of lines (from 0 through 9). Skipping 0 lines moves printing to the next line.

See also

- Print and Document Services Overview
- Branch Office Direct Printing Overview

Assign Delegated Print Administrator and Printer Permission Settings in Windows Server 2012

Planning security for your print servers and determining how to restrict access to them is an important part of print server administration. In Windows Vista® and Windows Server® 2008, only full system administrators were able to perform print administrative tasks. In Windows Server 2008 R2 and Windows Server 2012, you can delegate print management tasks directly to users who are not system administrators. You can also define default printer security settings that are inherited when you add new printers to your print server.

These changes enable the following improvements for printer and print server administration:

- You can control access to resources and balance workloads by delegating specific print administrative tasks to users without adding them to the Administrators security group.
- You can manage permission settings through the improved user interface of the Security tab in the Print Server Properties dialog box.
- You can manage your printer infrastructure by configuring default printer security settings which new printers inherit automatically when you add them. You can configure the settings per server so that you do not have to configure the printers individually.

Configuring security settings

This section covers the following:

- The print server security user interface
- Setting permissions in Print Server Properties
- Creating a delegated print administrator
- Print-related permissions and the tasks they enable
- Designing and creating print security groups

Note
Print server security can be configured only by members of the Administrators group.

The print server security user interface

In Windows Server 2012, users in the Administrators group can configure the print security settings directly by editing the print server access control list (ACL) permissions in the Print Management Microsoft Management Console (MMC) snap-in. (To view the ACL permissions for your printer server, open Server Manager, click **Tools**, and then click **Print Management**. In the left pane, click **Print Servers**, right-click the applicable print server and then click **Properties**. In the **Print Server Properties** dialog box, click the **Security** tab.)

Figure 1 shows the user interface of the **Security** tab that is opened by a user who is a member of the Administrators group.

Figure 1: Print Server Properties Security tab

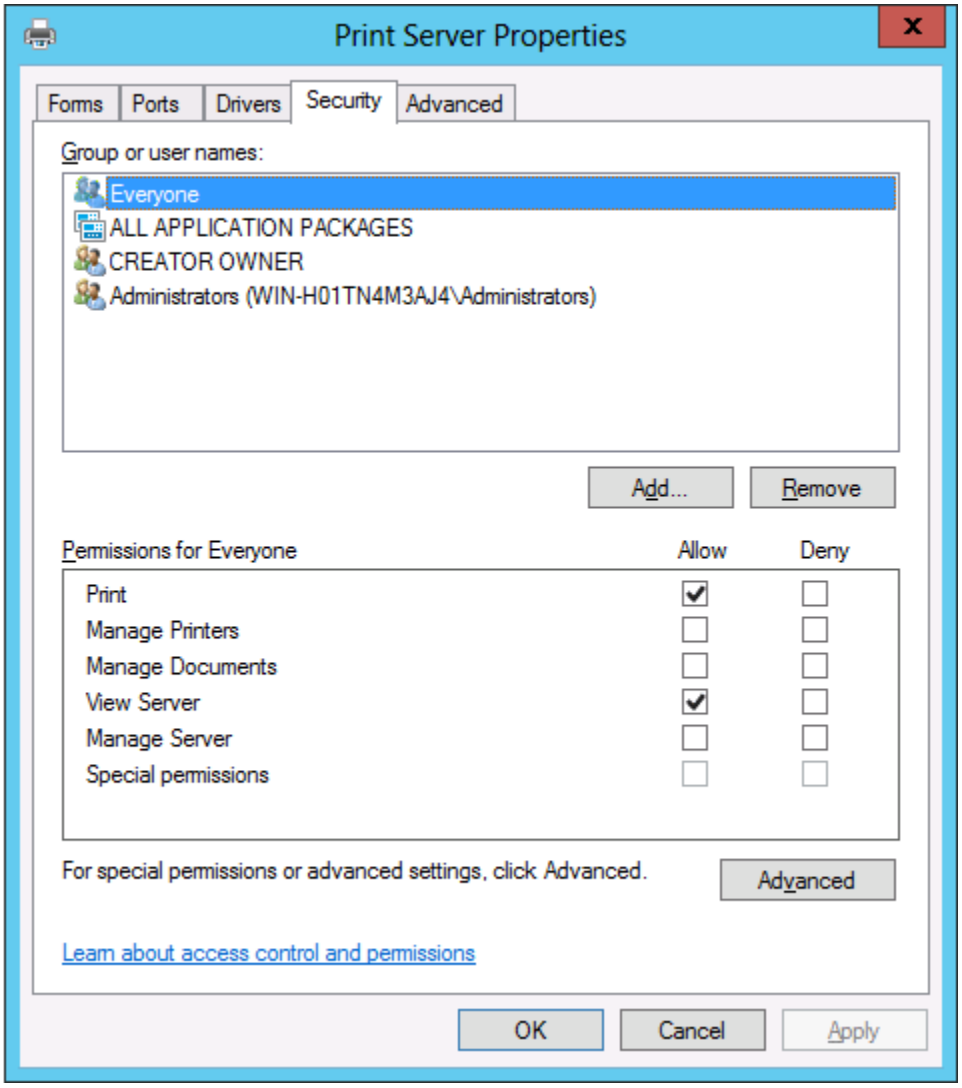

In a domain, members of the Administrators group can remotely configure the print server security settings. You can do this by using the Print Management snap-in. The remote functionality for users to view the print server security user interface is supported for certain earlier operating systems, including Windows Server 2008, Windows Vista with SP1, and Windows Vista with SP2. However, the delegated print administrator functionality is currently only available on Windows Server 2008 R2 and Windows Server 2012.

Setting permissions in Print Server Properties

Print server permissions control the levels of access for users on a particular print server. Printer permissions control which printing tasks users can perform on newly added printers that are managed by the print server. Administrators should assign these permissions as needed to users who are not system administrators.

After an administrator customizes the security settings for the print server, all newly added printers to this print server automatically inherit these security settings. (The security settings for the existing printers on the server are not altered.)

The two levels of print server permissions are:

- **View Server**

  The View Server permission assigns the ability to view the print server. Without the View Server permission, users cannot see the printers that are managed by the server. By default, this permission is given to members of the Everyone group.

- **Manage Server**

  The Manage Server permission assigns the ability to create and delete print queues (with already installed drivers), add or delete ports, and add or delete forms. A standard user with this permission is called a “delegated print administrator.”

  Note

  Only users who have Manage Server access and are members of the Administrators group can add printer drivers.

The three levels of printer permissions are:

- **Print**

  The Print permission assigns the ability for users to connect to printers and to print, pause, resume, start, and cancel their own documents. By default, this permission is given to members of the Everyone group when a print queue is created.

- **Manage Documents**

  The Manage Documents permission assigns the ability to control job settings for all documents and to pause, restart, and delete all documents.

- **Manage Printers**

  The Manage Printer permission assigns the ability to pause and restart the printer, change spooler settings, share a printer, adjust printer permissions, and change printer properties.

The ability to assign access to a printer on a per-user or a per-group basis makes it possible to manage printers from a central location. For example, an administrator could limit access to a printer in a public area while managing the printer from a more secure, central location.

In Windows Server 2012, the default print server and printer security settings are as follows:

	Everyone	Creator Owner	Administrators
Print	Allow		Allow
Manage Documents		Allow	Allow
Manage Printers			Allow
View Server	Allow		Allow

	Everyone	Creator Owner	Administrators
Manage Server			Allow

Creating a delegated print administrator

Members of the Administrators group can create a full delegated print administrator by assigning the Manage Server permission to a user. When the Manage Server permission is assigned, the View Server permission is also automatically assigned. You can also delegate a subset of these permissions to create a partial delegated print administrator.

To create a full delegated print administrator

1. Open Server Manager, click **Tools**, and then click **Print Management**.

   In the left pane, click **Print Servers**, right-click the applicable print server, and then click **Properties**.

   In **Print Server Properties**, click the **Security** tab.
2. To configure permissions for a new group or user, click **Add**. Type the name of the group or user that you want to set permissions for by using the following format: *domain name\username*. Click **OK** to close the dialog box.

   Tip
   Before adding any printers to the server, you should create a group of users who can perform delegated print tasks, and then configure the proper permissions. If you do this, all newly added printers automatically inherit these settings, and you do not have to individually configure existing printers for the print server.
3. Highlight the user or group name that you just added, and in **Permissions for** *<user or group name>*, click **Allow** for the **Manage Server** permission. (The **View Server** permission is assigned too.)
4. Select the **Allow** check boxes for the **Print**, **Manage Documents**, and **Manage Printers** permissions.

To create a partial delegated print administrator

- To enable an administrator to add printers:

  Follow the previous instructions, but select the **Allow** check boxes for the **Manage Server** and **Print** permissions. (**View Server** permission is assigned automatically too.)
- To enable an administrator to manage existing print queues:

  Follow the previous instructions, but select the **Allow** check boxes for the **View Server**, **Print**, **Manage Documents**, and **Manage Printer** permissions.

Print-related permissions and the tasks they enable

The following table lists the print tasks that a user can perform when assigned the corresponding permissions from the **Print Server Properties Security** tab.

	Print	Manage Printers	Manage Documents	View Server	Manage Server
View the print queue (on the local server)				Yes	
Print owned documents to the queue	Yes				
View, pause, restart, and cancel all print jobs in a queue			Yes		
Update installed or included drivers, and drivers available from Windows Update, to an existing queue **Note** This does not apply to clustered print environments.		Yes			
Add or delete a form in a queue		Yes			
View the printer properties				Yes	
View the print server proprieties				Yes	
Configure printer security permissions in a print queue		Yes			
Manage the print server security descriptor *setServerSecurityDescriptor flag*					

	Print	Manage Printers	Manage Documents	View Server	Manage Server
Add a print queue to a print server					Yes, when the drivers are already installed.
Delete a print queue from a print server		Yes, but only the queue they have permissions for.			
Add a print driver to a print server					Yes, but locally only. The user must be a member of the Administrators group to add drivers (including remotely) to the print server.
Delete a print driver from a print server					Yes, but only for drivers (not driver packages).
Add, delete, and configure ports on a print server					Yes
Add and delete a form on a print server		A user who is assigned Manage Printers, but not Manage Server, permissions can add a form when AllowUserManageForms is set in the Windows registry to a non-zero value. A user can add forms up to the specified value for AllowUserManageForms. A			Yes

	Print	Manage Printers	Manage Documents	View Server	Manage Server
		user can only add user forms and delete user forms. However, a user with SERVER_ACCESS_ADMINISTER permission can add and delete printer and user forms with no limitations.			
Share the printer		Yes, if you have Manage Printer permissions on the print server and the **File and Printer Sharing*** exceptions have been enabled in Windows Firewall with Advanced Security.			Yes, if you have Manage Printer permissions on the print server and the **File and Printer Sharing*** exceptions have been enabled in Windows Firewall with Advanced Security.

Designing and creating print security groups

Following is a list of suggested print security groups and their associated permissions:

- **System Administrators Group**: Consists of members of the Administrators security group.
- **Print Administrators Group**: Consists of members of the System Administrators group and users who have been assigned some set of delegated print administrator rights. Depending on what rights you assign, members of this group may be considered full delegated administrators or partial delegated administrators.

Note

If you want to mitigate the ability of members of the Administrators group to perform print management tasks, instead of adding whole groups to these print security groups, you can add members individually, and then assign the proper permissions.

The following table demonstrates which actions can be performed depending on the permissions assigned:

	Standard Users: Can connect to printers and print their documents (Permissions: Print, View Server)	**Partial Delegated Administrators: Can add printers (Permissions: Print, View Server, Manage Server)**	**Partial Delegated Administrators: Can manage existing queues (Permissions: Print, View Server, Manage Printers, Manage Documents)**	**Full Delegated Administrators: Can perform all administrative print tasks (Permissions: Print, Manage Documents, Manage Printers, View Server, Manage Server)**	**System Administrators: Can fully administer the system (Permissons: Print, Manage Documents, Manage Printers, View Server, Manage Server)**
View the print queue on the local server	Yes	Yes	Yes	Yes	Yes
Print to the queue	Yes	Yes	Yes	Yes	Yes
View, pause, restart, or cancel print jobs owned by the user in a queue	Yes	Yes	Yes	Yes	Yes
Modify all print jobs in a queue			Yes	Yes	Yes
Update an installed or included driver to an existing queue			Yes	Yes	Yes
Add or delete a form in the queue			Yes	Yes	Yes
View the printer properties	Yes	Yes	Yes	Yes	Yes
View the print server proprieties	Yes	Yes	Yes	Yes	Yes
Manage security permission on the print			Yes	Yes	Yes

	Standard Users: Can connect to printers and print their documents (Permissions: Print, View Server)	Partial Delegated Administrators: Can add printers (Permissions: Print, View Server, Manage Server)	Partial Delegated Administrators: Can manage existing queues (Permissions: Print, View Server, Manage Printers, Manage Documents)	Full Delegated Administrators: Can perform all administrative print tasks (Permissions: Print, Manage Documents, Manage Printers, View Server, Manage Server)	System Administrators: Can fully administer the system (Permissons: Print, Manage Documents, Manage Printers, View Server, Manage Server)
queue					
Manage the print server security descriptor *setServerSecurityDescriptor flag*					Yes
Add and delete the print queue on a server		Yes, but you can add a printer using only a preinstalled driver.	Yes, but you can only delete the print queue with the Manage Printer permission.	Yes, but you can add a printer using only a preinstalled driver.	Yes
Add and delete a print driver on a server		Yes, but locally only. The user must be a member of the Administrators group to add non-included drivers or to add drivers		Yes, but locally only. The user must be a member of the Administrators group to add non-included drivers or to add drivers	Yes

	Standard Users: Can connect to printers and print their documents (Permissions: Print, View Server)	**Partial Delegated Administrators: Can add printers (Permissions: Print, View Server, Manage Server)**	**Partial Delegated Administrators: Can manage existing queues (Permissions: Print, View Server, Manage Printers, Manage Documents)**	**Full Delegated Administrators: Can perform all administrative print tasks (Permissions: Print, Manage Documents, Manage Printers, View Server, Manage Server)**	**System Administrators: Can fully administer the system (Permissons: Print, Manage Documents, Manage Printers, View Server, Manage Server)**
		remotely to the print server.		remotely to the print server.	
Add, delete, and configure ports on a print server		Yes		Yes	Yes
Add and delete a form on a print server		Yes		Yes	Yes
Share the printer			Yes, if you have Manage Printer permissions on the print server and the **File and Printer Sharing*** exceptions have been enabled in Windows Firewall with Advanced Security.	Yes, if the **File and Printer Sharing*** exceptions have been enabled in Windows Firewall with Advanced Security.	Yes

Notes

We recommend that only a member of the System Administrators group install drivers. If a delegated print administrator plans to remotely add or manage queues, the System Administrator should install the driver to the following directory by using Windows PowerShell® or manually:

*systemdrive***Windows****System32****spool****drivers***processor_architecture***3**

For more information about Windows PowerShell Print Management cmdlets, see Print Management Cmdlets in Windows PowerShell.

Branch Office Direct Printing Overview

This topic introduces the Branch Office Direct Printing feature in Windows Server 2012 and describes how it can reduce your network costs if you have branch office printers combined with a centralized Print and Document Services environment.

Did you mean…

- Print and Document Services Overview

Feature description

Branch Office Direct Printing can reduce Wide Area Network (WAN) usage by printing directly to a print device instead of a server print queue. This feature can be enabled or disabled on a per printer basis and is transparent to the user.

This feature requires a print server running Windows Server 2012 and clients running Windows 8. It is enabled by an administrator using the Print Management Console or Windows PowerShell on the server.

Practical applications

Branch Office Direct Printing may be able to save your organization money by reducing network traffic on expensive WAN links and by reducing the workload on busy print servers centrally located in your datacenter.

Important functionality

Branch Office Direct Printing may benefit your organization in the following ways:

- Client computers running Windows Server 2012 obtain printer information from the print server, but send the print jobs directly to the printer. The print data no longer travels to the central server and then back to the branch office printer.

- The printer information is cached in the branch office, so that if the print server is unavailable for some reason (for example if the WAN link to the data center is down), then it is still possible for the user to print.
- The client computer renders the print job before sending it to the printer. The major difference between Branch Office Direct Printing and Client-Side Rendering (CSR) is that Branch Office Direct Printing does not send the print job to the server. Instead the job is sent directly to the printer, resulting in a significant reduction in network bandwidth usage when the printers are centrally managed.

  Note
  CSR, first supported with Windows Server 2008, allows print jobs to be fully rendered on the client when targeting a shared printer hosted on a Windows Vista or Windows Server 2008 computer acting as a print server. After being rendered on the client, the print job is sent to the print server for queuing and printing by the print spooler without additional server-side rendering. This feature is enabled by default on Windows Vista and Windows Server 2008, and is not available on earlier versions of Windows.
- Scalability may be improved as server load will be reduced, but exact amount depends on your organization's printing environment and setup.

Software requirements

Branch Office Direct Printing requires the following operating systems:

- Windows Server 2012
- Windows 8

See also

For more information about Branch Office Direct Printing, see the following topics:

Content type	References
Product evaluation	Branch Office Direct Printing Technical Details
Deployment	Configure Branch Office Direct Printing
Related technologies	Print and Document Services Overview

Configure Branch Office Direct Printing

This topic shows you how to configure Branch Office Direct Printing on a print server running Windows Server 2012.

In this document

- **Prerequisites**
- **Step 1: Configure Branch Office Direct Printing**

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

This procedure assumes you have a print server installed with printers installed and shared on Windows Server 2012.

Step 1: Configure Branch Office Direct Printing

You can configure Branch Office Direct Printing using either the Print Management Console or Windows PowerShell.

Do this step using Windows PowerShell

To Configure Branch Office Direct Printing

1. Open the Print Management Console and expand **Print Servers**.
2. Expand the print server where the print queues are installed and then expand **Printers**.
3. Right click the printer that you wish to manage and select **Enable Branch Office Direct Printing**. Multiple printers can be configured at the same time by highlighting each printer prior to this step.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

```
Set-Printer -name <String> -ComputerName <String> -RenderingMode BranchOffice
```

See also

- Branch Office Direct Printing Overview
- Branch Office Direct Printing Technical Details
- Print and Document Services Overview

Branch Office Direct Printing Technical Details

Branch Office Direct Printing is designed to reduce network traffic on expensive WAN links and potentially reduce the workload on busy print servers centrally located in your datacenter. The topic covers the architecture and design of Branch Office Direct Printing in Windows Server 2012.

Overview

Branch Office Direct Printing has the following benefits and limitations:

Benefits of Branch Office Direct Printing

- Works for print queues using printer drivers that support Client Side Rendering (CSR), whether they are v3 OR v4 drivers.
- Functions with both TCP and Web Services on Devices (WSD) print devices (directed discovery only, not multicast)
- Can be enabled or disabled via Windows PowerShell.
- PrintBRM will backup and restore the Branch Office Direct Printing setting on a print queue.
- Saves bandwidth and jobs may print faster since print jobs no longer go through server.
- Scalability *might* improve, as server load will be reduced, but exact amount depends on your organization's printing environment and setup.

Limitations of Branch Office Direct Printing

- Requires a print server running Windows Server 2012, and Windows 8 clients.
- The printer must be a network attached device.
- The printer can't be locked to accept print jobs from a single IP address. It must be configured to accept print jobs from any client.
- Quota, auditing and detailed job logging capabilities are lost if Branch Office Direct Printing is used. For queues that must retain these records, leave Branch Office Direct Printing disabled
- Even if Branch Office Direct Printing is enabled on a print queue, clients running an operating system prior to Windows 8 will not use this setting and will still print to the server.
- Settings that conflict with Branch Office Direct Printing (such as **Keep Printed Jobs**) cannot be used.
- If server side rendering (SSR) is required for some specific reason, Branch Office Direct Printing cannot be used because it depends on CSR.
    - For example, if a client computer running Windows 8 receives and installs a v4 print class driver (a v4 driver that supports a broad set of printers sharing common features) because a model specific v4 driver is not available, then CSR is not possible and thus Branch Office Direct Printing is not supported.

- If printer pooling is enabled on a printer queues, Branch Office Direct Printing is not compatible.
- All third party server-based tools will no longer work properly (quota management, job auditing, and so on.) because print jobs will no longer be routed through the server.
- If Branch Office Direct Printing is enabled, it will override the default SSR mode for systems like laptops and tablets. Typically, these portable computers running Windows 8 will force SSR to conserve battery power while printing. When Branch Office Direct Printing is enabled on a queue, it is assumed that bandwidth is more valuable than extended battery life, and the portable clients will switch to CSR and take advantage of the Branch Office Direct Printing feature.

Architecture

The following diagrams show a visual representation of the data flow for CSR versus Branch Office Direct Printing:

Figure 1: Printing Data Flow for CSR Printing

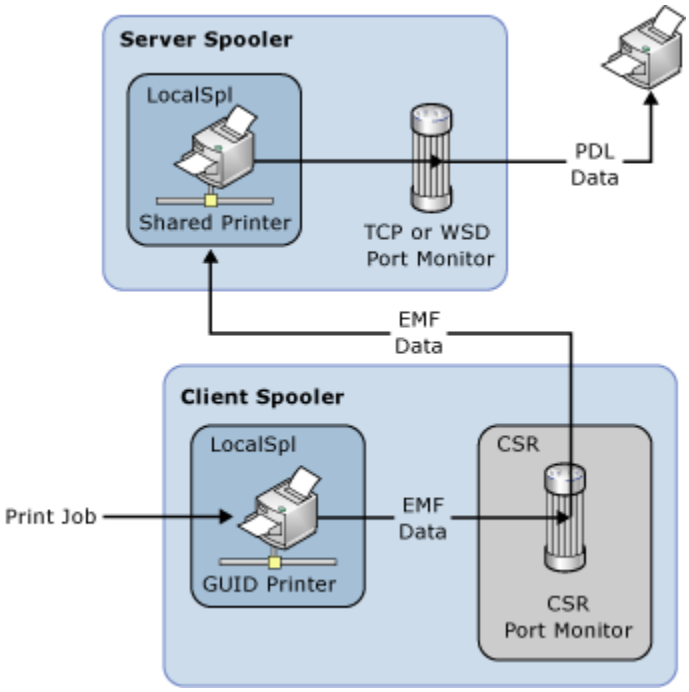

Figure 2: Print Data Flow for Branch Office Direct Printing

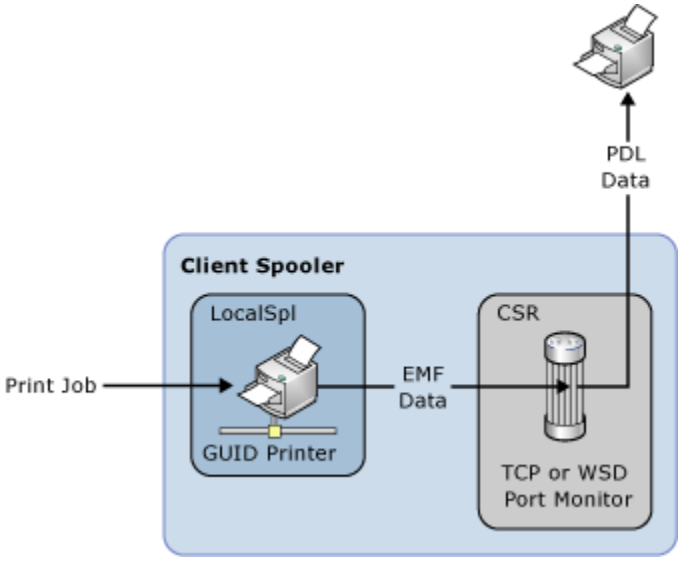

Branch Office Direct Printing depends on the existing CSR print technology. CSR implements a remote printer connection in such a way that the bulk of the work involved in a print job is performed locally on the client. With CSR, the print job is processed on the client and then sent to the print server. The print server then sends the spooled data to the shared print device. In Branch Office Direct Printing mode, the port on the client is replaced so when CSR is done, the job is sent directly to the physical print device instead of the print server.

Implementation Details

The implementation of this feature and its integration into the CSR feature can be broken down into several parts: Configuration, Setup, Synchronization, and Fallback.

Configuration

Branch Office Printing is configured on a per-connection basis on the server – there is no way to enable this feature from the client side. Configuration is accomplished by using the Print Management Console or Windows Windows PowerShell to set a registry value on the print queue:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Print\Printers\<printername>\PrinterDriverData
```

```
valueName: EnableBranchOfficePrinting
```

Values set on the remote printer are synchronized periodically by the client (either after an elapsed time period, or when changes on the server are detected), and when “BypassServer” is set to a non-zero value, Branch Office printing for the connection is enabled. The client may be

forced to ignore this setting if it is incompatible or if conflicting settings are in effect, explained in more detail in the fallback section.

Setup

Creating a connection involves setting up a local GUID printer, and as part of setting up the printer, it must be assigned a port. For both normal and Branch Office Direct Printing, this port defaults to the CSR Port Monitor until a CSR cache synchronization occurs.

As the last step of setup, a synchronous CSR cache sync is forced for the new GUID printer – as a result, this copies over all the remote printer's settings to the GUID printer and at this point the port settings are changed from the CSR Port Monitor to the port settings used on the server.

Synchronization

As described above, during Setup the CSR Cache Synchronization is the point at which CSR chooses how to configure the port for the connection. The print spooler then uses an internal work item that is responsible for keeping a single local GUID printer in sync with the remote print queue the connection it is associated with. This ensures that if the administrator changes port settings on the remote printer that they will be propagated to the client.

The synchronization process determines how to configure the port based on the following:

- Whether branch office printing is enabled or disabled for the connection.
- If Branch Office printing is enabled and if the remote port settings differ from the local ones, the local settings must be synchronized.
- If Branch Office printing is disabled, then the local port must be set to the CSR Port Monitor in order to allow normal operation of the connection.

After these settings are checked the local port is updated accordingly and synchronization is complete.

The following diagram displays the synchronization logic:

Figure 3: Synchronization Logic

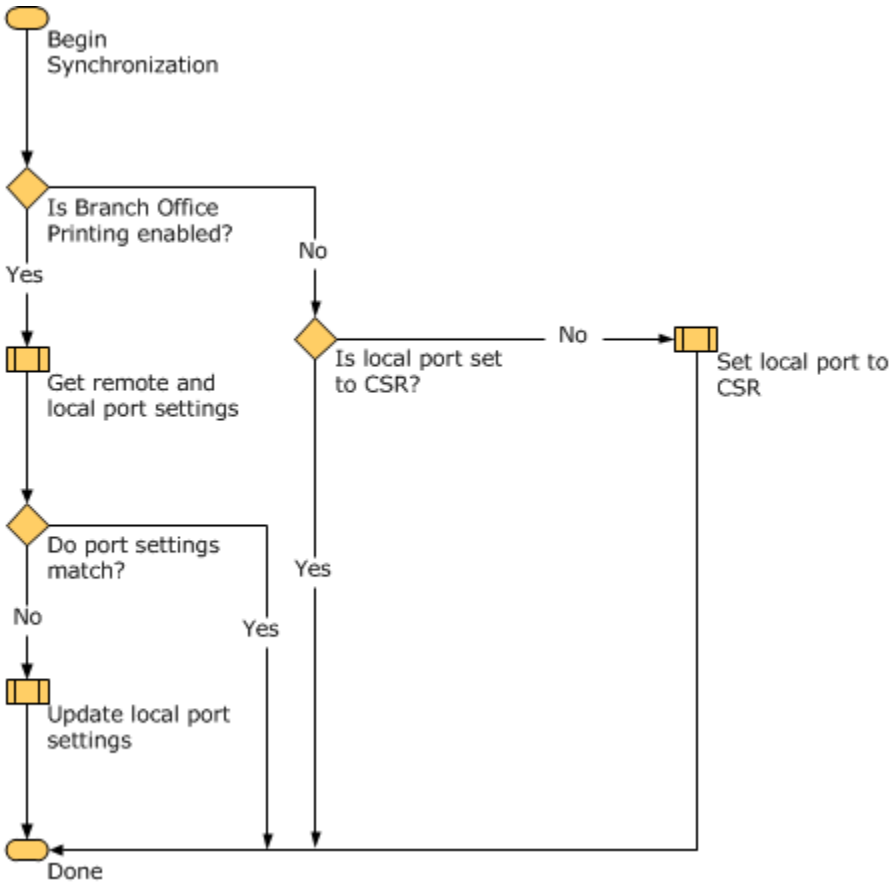

Branch Office Status

The registry value that determines whether Branch Office Printing is enabled for the connection is contained in the GetPrinterData value EnableBranchOfficePrinting. If this value is zero or unset then Branch Office Printing is disabled, if set to a nonzero value then the feature is enabled as long as no other settings conflict with it and the connection hasn't been determined to be incompatible.

A table of conflicting settings and their resolution is defined in the **Fallback** section.

Synchronization Before Printing

The CSR Sync manager in general runs in a polling mode such that it pulls down settings from the server periodically. Since the sync manager is used for synchronizing port settings, there is the potential for a window where the client's port settings are out of sync with the server's. If a client were to print a job to a Branch Office printer during this time the job would fail or go to the wrong place.

To eliminate this issue, a synchronous cache sync is forced on the client at the start of every job. This insures the settings are up to date at the time when they are needed most. To reduce the performance impact, only port information is synced, not a full printer sync.

Fallback

There are mutually exclusive settings, some of which cause automatic fallback (either to CSR or SSR, both through the server). If either condition occurs there are events logged to the event log.

- Printer falls back to Server side rendering if the enhanced point and print driver is used.
- SSR Fallback occurs if the driver in use does not support CSR.
- CSR Fallback occurs if port settings are not WSD or TCP.
- SSR Fallback occurs if admin has manually enabled SSR on the given queue.
- CSR Fallback occurs if the **Keep printed jobs** setting is enabled on the given queue.
- CSR Fallback occurs if printer pooling is enabled on the given queue.

In each of the above cases, the client will fall back to normal printing and an event will be logged explaining why for administrative purposes. To make the fallback apparent to the administrator, warning events are logged on the server as well. The server checks whether incoming jobs are from clients that should have Branch Office Printing enabled but go through the server instead.

See the following tables for more information about these events.

Branch Office Direct Printing Event IDs

The following Branch Office Direct Printing events are logged in the **Applications and Services Logs/Microsoft/PrintService/Admin** event log:

Event ID	Description
858	Connection '%1' has been reconfigured for normal operation because the queue has been configured for Server Side Rendering. To re-enable Branch Office Printing, enable the 'Render Jobs On Client' setting on the server queue.
862	Connection '%1' has been reconfigured for normal operation because the 'Keep Printed Jobs' setting is enabled on the queue. To re-enable Branch Office Printing, disable the 'Keep Printed Jobs' setting on the server queue.

The following Branch Office Direct Printing events are logged in the **Applications and Services Logs/Microsoft/PrintService/Operational** event log:

Event ID	Description
856	Connection '%1' has been reconfigured for normal operation because branch office printing has been disabled.
857	Connection '%1' has been reconfigured for normal operation because the queue is incompatible with branch office printing.
859	Connection '%1' has been reconfigured for normal operation because the client is incompatible with branch office printing.
860	Connection '%1' has been reconfigured for normal operation because the server is incompatible with branch office printing.
861	Connection '%1' has been reconfigured for normal operation because the remote port is incompatible with branch office printing.
863	Connection '%1' has been reconfigured for normal operation due to an internal error, %2.

Branch Office Direct Print (BODP) Logging in Windows Server 2012 R2

Branch Office Direct Print (BODP) Logging Overview

Windows Server 2012 R2 allows admins to take advantage of the Branch Office Direct Printing feature without having to sacrifice print job logging information.

This enables admins to confidently deploy BODP to save their organization bandwidth and improve connectivity without losing the ability to track print jobs and get logs to help with troubleshooting.

Branch Office Direct Print (BODP) Logging Configuration

This feature is enabled by default for all new print queues, but the configuration may be viewed and configured with Windows PowerShell.

The "set-printer" and "get-printer" cmdlets have been extended with two new parameters.

Name	Type	Description
branchOfficeLoggingEnabled	Boolean	Enables or disables logging for branch office queues. $true by default. Cannot be specified for unshared queues or queues that do not have Branch Office enabled.
branchOfficeOfflineLogSize	Integer	Specifies the maximum size of the offline log size on the client, in megabytes. Default is 10mb and a value of zero disables logging. This file persists logging data if the server is offline or unreachable.

Branch Office Direct Print (BODP) Logging Limitations

This feature is limited in the following manner:

- This feature requires a Windows 8.1 client and Windows Server 2012 R2.
- Programmatic configuration of the offline log file size can only be accomplished through Windows PowerShell. A Group Policy is not available.
- The following non-critical events are not logged.
  - Paused (Event ID 308)
  - Resumed (Event ID 309)
  - Deleted (Event ID 310)
  - Timed Out (Event ID 314)

Branch Office Direct Print (BODP) Logging Diagram

The following chart illustrates BODP Logging.

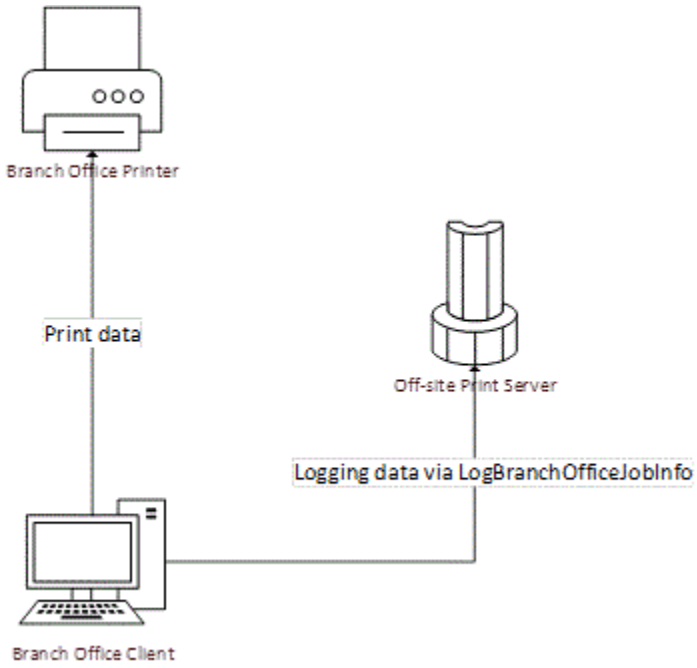

With this feature, a new communication channel for job event data is created by means of a new RPC entry point named LogBranchOfficeJobInfo. An authenticated client can send job events through this API and the server will add those events to its Event Log as if the job had been processed on the server itself.

Print and Document Services Architecture

This article provides an overview of the version 4 (v4) print driver model and the specific types of v4 drivers known as **Print Class Drivers** and **Model Specific Drivers**. Windows Server 2012 fully supports both version 3 (v3) and v4 print driver models. In addition, the different printer sharing types supported by Windows Server 2012 is covered.

Printer driver overview

At the core of the Windows Server 2012 printing experience is a new driver model known as the v4 print driver model. The v4 driver model includes changes to printer sharing known as enhanced Point and Print, eliminating the need to install cross platform drivers and eliminating the scenario where a Print Server is a driver distribution point.

V4 printer drivers

The v4 print driver model provides a simple but flexible management experience. V4 print drivers can be distributed via Windows Update or Windows Software Update Services (WSUS), but are not distributed to print clients from the print server.

Some of the benefits of v4 drivers are:

- Printer sharing does not require the installation of device specific drivers that match the client architecture
- Driver files are isolated from each other, preventing driver file name conflicts
- A single driver can support a single device or multiple devices
- Drivers are typically smaller than version 3 drivers, and installation times are faster than comparable version 3 drivers
- Administrators can deploy printer user interface separately from the printer driver

Print Class Drivers

Using the v4 driver model, IHVs can provide **Print Class Drivers** that support features that are common to a broad set of devices that use the same printer description language, such as PCL, PS, or XPS. The print drivers included with Windows Server 2012 are Print Class Drivers and will always have the text “class driver” in their display name.

Print Class Drivers also provide the following benefits:

- Driver packages are smaller and printer driver files are less likely to cause stability problems
- Administrators have fewer drivers to manage
- Class drivers can support future devices that haven’t been released yet

Model Specific Drivers

You can get model specific v4 drivers directly from the printer manufacturer’s web site or from Windows Update.

V3 printer drivers

Prior to Windows Server 2012, the Windows printer driver model had remained relatively unchanged since the introduction of v3 drivers in Windows 2000. The v3 model relies heavily on OEMs to produce customized drivers for each specific device to ensure that specific features of each print device can be accessed by Windows applications. Managing a printing infrastructure using the v3 driver model requires the administrator to manage a large number of drivers, and to manage drivers on client computers as well as servers, resulting in the requirement to manage both 32-bit and 64-bit print drivers to support both of these client and server architectures.

Changes to the printer driver model

Microsoft works closely with independent hardware vendors (IHVs) to obtain printer drivers for every release of Windows. Windows 7 includes 2100 print drivers with the operating system, and

several thousand more printer drivers are available on Windows Update. Each driver is produced by the IHV, and is branded and customized to provide the best experience possible.

Print drivers have historically comprised a large portion of the files included with Windows, even if the user never connects to or uses a printer. In Windows Vista, print drivers accounted for about 1 GB of the total Windows installed files; in Windows 7, this was reduced to 538MB or about 60% of the overall driver files installed for Windows.

Print drivers take up a lot of disk space because the printer drivers in Windows have historically been model specific. Every device is associated with an individual driver and these drivers have no ability to support devices that will be released in the future. Additionally, the drivers that are available for each device lose relevance over time as new devices enter the market. While there may be code sharing between some devices, small variances between devices generally require that every device has its own data files or DLLs in order to precisely describe every attribute of the device.

DLLs are the largest proportion of the driver files. DLLs include not only code for tasks like PDL rendering or customized UI, they also contain the localized strings and image resources that support that UI. Windows configuration files, including GPD (Generic Printer Description) and PPD (PostScript Printer Description) files are the next largest group. These files describe the features and options a printer supports, and in some cases, how to render PDL for the device. The full breakdown is described in the table below.

Type of file	Size (KB)
DLL	342,711.38
Windows Configuration Files	88,148.78
Data Files	52,389.66
Color Profiles	30,228.26
XML	18,228.42
Setup files	4,581.20
Security catalogs	1,828.67
Fonts	102.23
Help files	14.05
Grand Total	**538,232.67**

Reducing driver size

To reduce the size of the operating system image and increase the long term relevance of the drivers that ship with Windows 8, it was necessary to change the paradigm. Many printers consume common formats known as PDLs (Page Description Languages) like PCL, PostScript or

XPS. While every device needs to get a PDL it understands, the commonality of PDLs introduces an opportunity to reduce the number of drivers overall. Instead of focusing on providing model specific drivers that provide every feature available on the device, printer drivers in Windows 8 focus on providing print class drivers that support only features that are common to a broad set of devices that use the same PDL. By doing this, much of the code needed to support existing devices can be eliminated, while at the same time providing a baseline for new printers to target. New printers can specify support for a print class driver using an identifier known as a **Compatible ID**. The net result is fewer, smaller drivers targeting more devices, and fewer problems for Windows users.

The result is shown in the following diagrams. Whereas the v3 print drivers included in Windows 7 could provide customization (blue boxes) in most layers of the driver model, print class drivers are intentionally limited to providing only critical code—PDL rendering. Windows provides the rest of the functionality using common modules and data files provided by the hardware vendor.

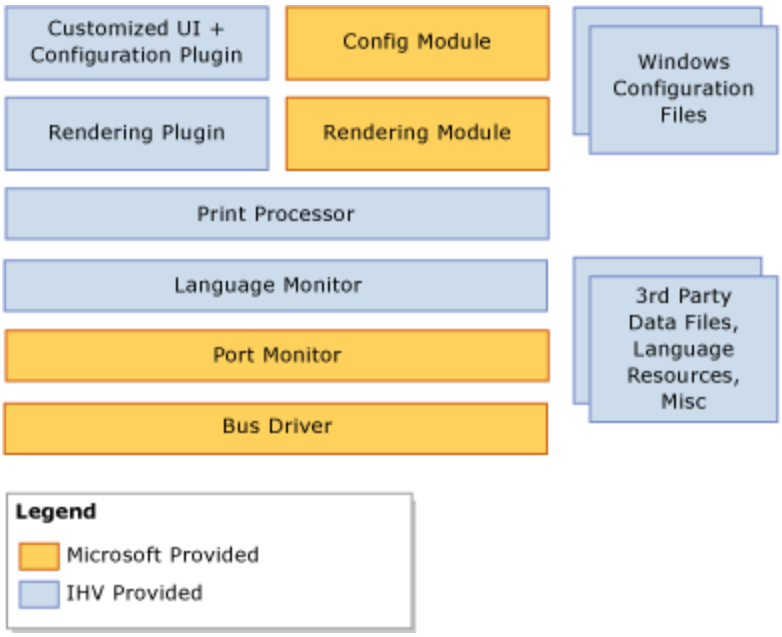

Figure 1: v3 Print Driver Model with GDI Rendering

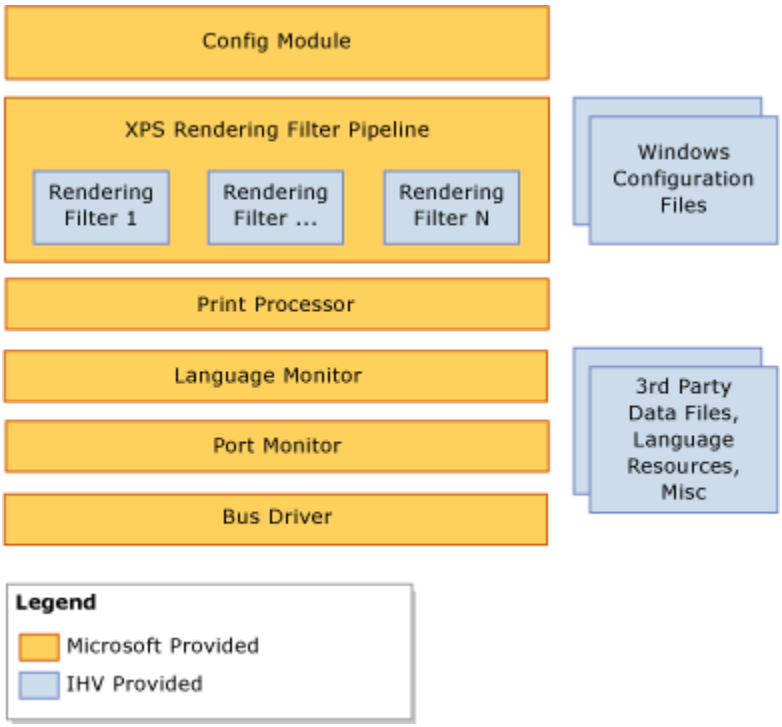

Figure 2: Print Class Driver (v4) Model

Driver isolation

First introduced in Windows 7 and Windows Server 2008 R2, the Windows Printing Driver Isolation feature removed an isolated driver from the print spooler process and loaded it into what is known as a shared sandbox with other isolated drivers, or a completely isolated sandbox to run completely isolated from the print spooler and any other drivers. If an isolated driver would experience a fault or otherwise unexpectedly terminate or crash, this condition would not affect the print spooler process and the other printers and drivers that were loaded. All drivers written for Windows 7 were required to support driver isolation, however drivers released prior to Windows 7 did not have a driver isolation attribute attached to them.

As was the case with Windows 7 and Windows Server 2012 R2, drivers are isolated in Windows Server 2012 in the following scenarios:

- The driver INF has the keyword **DriverIsolation** that indicates it supports driver isolation.

  DriverIsolation=2 indicates that the driver supports driver isolation. Setting **DriverIsolation=0** indicates that the driver does not support driver isolation.
- The administrator specifically enables driver isolation for a particular driver

  This can be done using the Print Management console. Expand **Print Servers**, expand your print server name, select **Drivers**, right-click the driver name and select **Set Driver Isolation**.

  The administrator can configure the driver to use one of the following settings:

  - **Shared**

Run the driver in a process that is shared with other printer drivers but is separate from the spooler process.

- **Isolated**

  Run the driver in a process that is separate from the spooler process and is not shared with other printer drivers.

- **None**

  Run the driver in the spooler process.

- The administrator specifically enables a group policy that controls the driver isolation behavior.

  The group policy name is: Computer Configuration\Policies\Administrative Templates\Printers**Override print driver execution compatibility setting reported by print driver** and it works as follows:

  - If the policy setting is enabled, the print spooler isolates all print drivers that do not explicitly opt out of Driver Isolation.
  - If the policy setting is disabled or not configured, the print spooler uses the Driver Isolation setting indicated by the **DriverIsolation** keyword in the INF file.

By default, if the INF file that installs a printer driver does not indicate that the driver supports driver isolation, the printer class installer configures the driver to run in the spooler process. However, if the INF file indicates that the driver supports driver isolation, the installer configures the driver to run in an isolated process. An administrator can override these configuration settings and specify, for each driver, whether to run the driver in the spooler process or in an isolated process.

Ideally, a printer driver is able to run in shared mode. That is, it runs in an isolated process shared with other printer drivers but separate from the spooler process. A driver might need to run in isolated mode if it can run in a process separate from the spooler process, but has difficulty sharing the process with other drivers. For example, a poorly designed driver might have file names that conflict with those of related drivers or of different versions of the same driver, or the driver might fault frequently or have a memory leak that interferes with the operation of other drivers that run in the same process.

To support troubleshooting, the domain administrator can disable the driver isolation feature on a computer in the domain, or the administrator can force all of the printer drivers on the computer to run in isolated mode. In isolated mode, each driver must run in a process separate from the spooler and from the other printer drivers.

Note

If the driver INF file explicitly states **DriverIsolation=0** then the administrator cannot force driver isolation.

If driver isolation is disabled by group policy, the isolation is off for all printer drivers. The group policy name is Computer Configuration\Policies\Administrative Templates\Printers**Execute print drivers in isolation**.

- If you enable or do not configure this policy setting, the print spooler runs print drivers in an isolated process by default.
- If you disable this policy setting, the print spooler runs print drivers in the print spooler process.

The following chart shows a decision map for choosing the driver isolation mode:

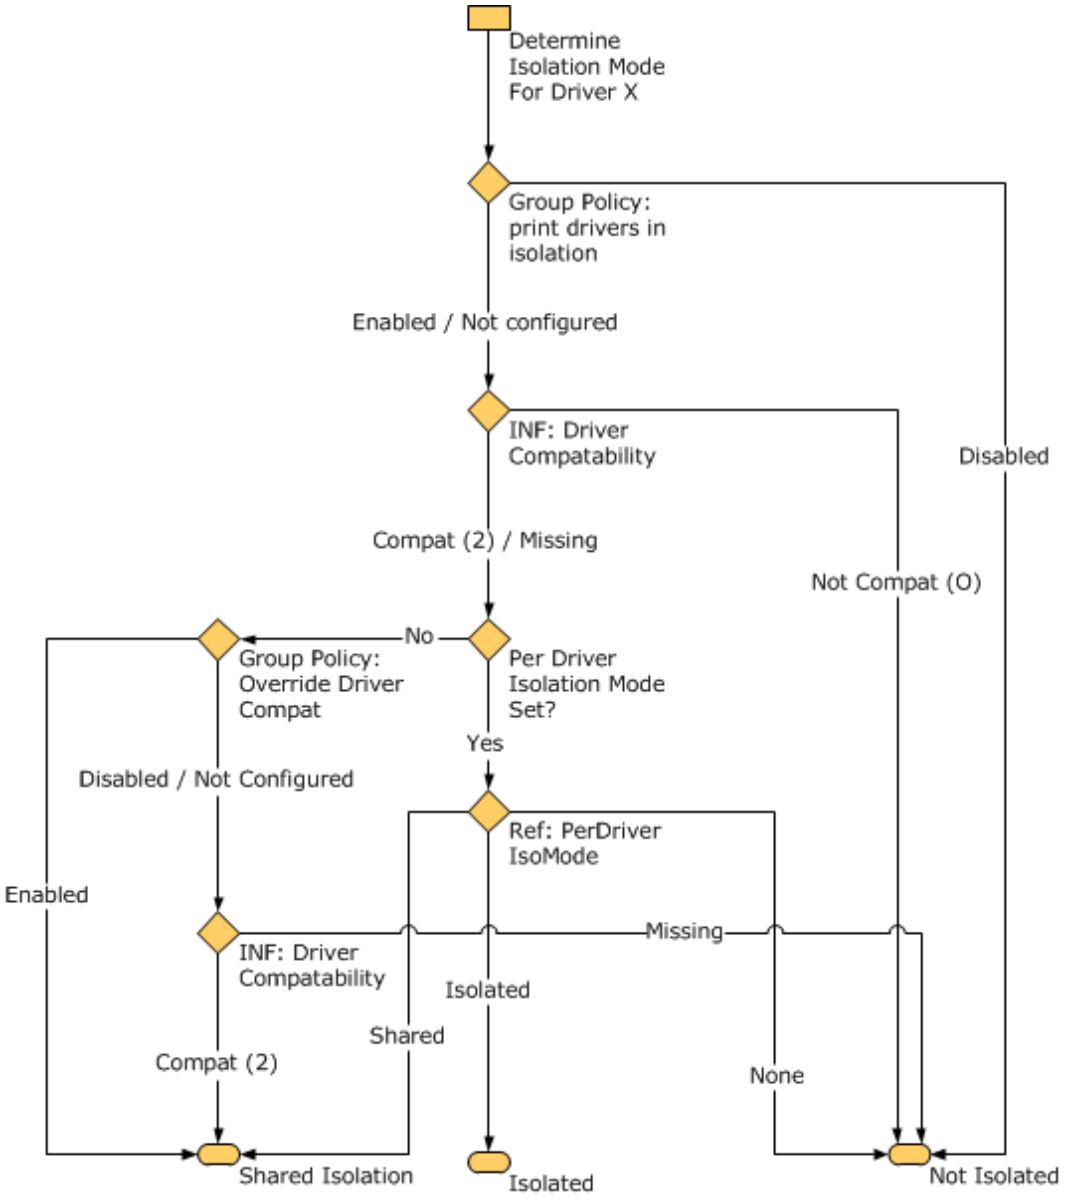

Application isolation

Application isolation isolates an application from a print driver so that if a print driver crashes, the application does not crash. This feature extends an existing Windows 7 feature (SplWoW64) to cover all printing scenarios in Windows Server 2012 and Windows 8. Applications that support

isolation are more stable because the application is separated from the print driver by using a “process isolation” layer. All of the driver DLLs are loaded and hosted in a separate process, isolating the application from print driver crashes

Splwow64 is a 64-bit process that hosts print drivers on behalf of the 32-bit applications. These applications are able to call standard printing functions and act without any knowledge of Splwow64. Behind the scenes, the calls are thunked back and forth between the client application and Splwow64. A side effect of this solution is that these applications are isolated from print driver crashes.

In Windows Server 2012 and Windows 8, with the group policy enabled, Splwow64 hosts the driver DLLs for all applications that support this feature.

Application isolation is controlled by the following Group Policy setting:

Computer Configuration/Administrative Templates/Printers/**Isolate print drivers from applications**

The **Isolate print drivers from applications** setting has the following description:

Determines if print driver components are isolated from applications instead of normally loading them into applications. Isolating print drivers greatly reduces the risk of a print driver failure causing an application crash.

Not all applications support driver isolation. By default, Microsoft Excel 2007, Excel 2010, Word 2007, Word 2010 and certain other applications are configured to support it. Other applications may also be capable of isolating print drivers, depending on whether they are configured for it.

If you enable or do not configure this policy setting, then applications that are configured to support driver isolation will be isolated.

If you disable this policy setting, then print drivers will be loaded within all associated application processes.

V4 driver model design and architecture

The past 10 years have seen the evolution of the v3 driver model from generic printer description (GPD)/PostScript printer description (PPD)-only configuration and rendering to XPSDrv, auto-configuration, and full UI replacement. These changes supported the demands of customers and partners, but also added complexity and management overhead. The following are some, but not all of the difficulties associated with the v3 driver model:

- Printer drivers are a top cause of crashes and hangs in Microsoft Office.
- Administrators spend a lot of time managing print driver updates and consolidating print queues, and end users struggle with cross-architecture drivers.

To help mitigate these issues, many partners have tried universal or class drivers, but the v3 driver model falls short in a number of key scenarios that would enable partner success.

The v4 printer driver model improves the quality of the device experience for all users by mitigating known issues in the v3 driver model and developing easier-to-implement extensibility points. While v4 print drivers continue to use GPDs, PPDs, Bidi, and Generic Descriptor

Language (GDL)-based autoconfiguration, some of the layers available in the past have been removed or replaced.

One new aspect of the v4 print driver model is a focus on providing print class drivers. These special v4 print drivers, which implement only functionality that is common across a broad class of devices, are shipped with Windows and are marked as generic. As a result, Windows can automatically replace the driver with a better option from Windows Update, providing the end user with a better experience as soon as it is available.

V4 driver design

The following were design considerations for the new v4 model:

- **Windows Store Apps**

  The new Windows experience and the new Windows Store app models introduce new design considerations regarding UI behavior and security context. V4 print drivers are very important in the Windows user interface with deep integration, seamless behavior, and rich extensibility.

- **Windows on ARM**

  Running Windows on an ARM processor changes the paradigm with regard to power utilization and driver behaviors. The v4 print driver model supports printers on ARM with print class drivers while still allowing rich end user experiences thanks to a decoupled UI development model. There is no v3 driver support for Windows on ARM.

- **Ease of Driver Development**

  The v4 driver model supports existing investments in v3 and the XPSDrv architectures, while making drivers easier to develop and test. The v4 driver model supports driver development using Visual Studio, and supports a number of templates that make drivers easier to build. V4 print drivers also include support for building customized UI as a simple Windows application, enabling developers to use the most up-to-date frameworks and toolsets to keep their experiences fresh.

- **Printer Sharing**

  Printer sharing is a major value proposition of Windows print servers, and v4 print drivers are designed to make printer sharing even better. In particular, they reduce management costs, eliminate cross-architecture driver management, and support all client operating systems from Windows Vista to Windows 8.

  Note

  Operating systems prior to Windows 8 do not support the v4 driver model but can print to a v4 queue shared from a Windows Server 2012 print server by using the enhanced Point and Print Compatibility Driver which is hosted by any print server running Windows Server 2012.

V4 driver architecture

The following is a high-level representation of a v4 print driver. With the exception of the rendering filters and user interface applications, all code is Microsoft provided. V4 print drivers

rely heavily on data files and JavaScript for extensibility. The blue boxes represent existing files that were used in the v3 driver model, and the green boxes represent new places to plug in.

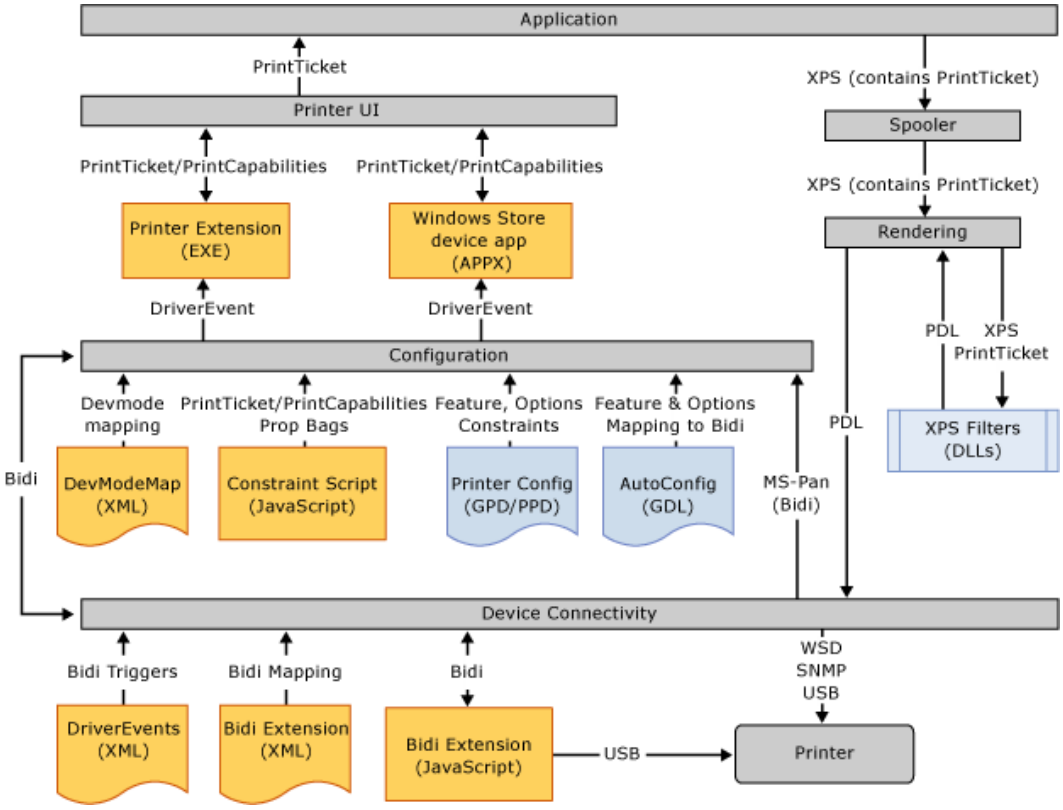

Figure 3: v4 Driver Architecture

Rendering

The v4 driver model uses only the XPSDrv architecture in order to support rendering into device PDL. XPS direct devices do not need any filters, but all others must either include filters that render into the device PDL or take a dependency on an existing print class driver using the RequiredClass directive in the v4 manifest file.

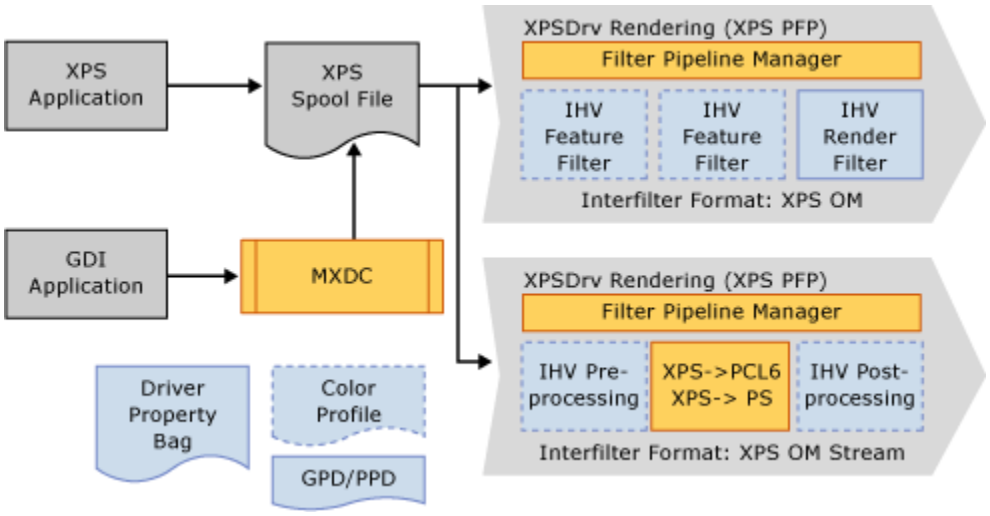

Figure 4: Rendering Architecture

Configuration Layer

The v4 driver model supports a greatly simplified configuration layer. Unlike v3 print drivers where the user interface was strongly coupled to the configuration, v4 print drivers focus on providing PrintTicket, PrintCapabilities, and constraint functionality. A common configuration module, PrintConfig.dll, encapsulates the functionality that was previously available in the UnidrvUI and PS5UI core drivers.

The v4 driver model does not employ configuration plug-ins, so most of the device configuration are expressed in GPD or PPD files. In addition, v4 print drivers may provide a JavaScript file that supports advanced constraint handling as well as PrintTicket and PrintCapabilities support.

The Generic Printer Description (GPD) and PostScript Printer Description (PPD) file formats are unchanged with v4 print drivers. Existing GPD and PPD files are compatible.

Customized User Interfaces

V4 print drivers support customized user interfaces in both the Windows desktop and in the new Windows user interface. Due to the very different nature of these experiences, these UIs must be implemented as two different applications.

Printer extensions support v4 print drivers in the desktop and work with all existing applications. They also work in printer sharing scenarios with the enhanced Point and Print driver. Support is planned for all operating systems from Windows Vista through Windows 8.

Windows Store device apps support v4 print drivers in the new Windows user interface.

Printer sharing overview

The printer sharing implementation in Windows 8 was updated to meet the demands of modern users and to support the new v4 driver model being introduced in Windows 8.

The prior model relied on both the server and the client computers using identical drivers; given the different supported architectures for Windows (x86, x64, and SoC systems) getting the right matching driver can be a difficult proposition. Incompatibilities between driver versions can cause client connection errors, and establishing connections can take a significant amount of time.

The overall goal of both the new v4 driver model and the sharing implementation that supports it is to make the end-user and administrative experience as easy as possible.

There are several features that define the changes that were made to printer sharing in Windows 8. These can be summarized as the following:

- Clients can use enhanced Point and Print to generate print jobs that the server can use without using a device specific driver.
- Servers can encapsulate the configuration and capabilities of the printer and communicate that data to a client computer in a way that the client computers can use without needing a device specific driver.
- The print server is no longer a software distribution mechanism. Previous versions of Windows provided a mechanism by which print clients could obtain a driver from the print server. For reasons of security, compatibility, serviceability, and reliability, this functionality has been removed from the v4 driver model and enhanced Point and Print. Downlevel client computers will still be able to receive the enhanced Point and Print Compatibility driver from Windows Server 2012 servers to enable compatibility with v4 print shares. Client computers running Windows 8 have enhanced Point and Print support built into the operating system, and they can use standard point and print mechanisms to obtain matching v3 drivers if the administrator wants to continue using older drivers. Device specific v4 drivers can also be deployed to Windows 8clients, or downloaded from Windows Update/WSUS in order to provide additional features or capabilities, such as client side rendering when connecting to enhanced Point and Print shares.

Architecture

There are three major changes to the sharing implementation in Windows Server 2012:

1. **Enhanced Point and Print Compatibility Driver**

   Windows Server 2012 servers provide the Microsoft enhanced Point and Print compatibility driver to clients running previous versions of Windows, and these clients can use the driver without the need for any updates or modifications to the client printing subsystem.

2. **Disabling Driver Distribution from the Server**

   Windows Server 2012 servers will not provide v4 drivers to clients (with the exception of the previously mentioned enhanced Point and Print compatibility driver). Clients running a previous version of Windows will still be able to connect to v4 print shares without issue. Servers running Windows Server 2012 using v3 drivers will behave the same way that Windows Server 2008 R2 behaves. No changes have been made to v3 printer sharing in Windows Server 2012.

3. **Enhanced Point and Print**

   Windows 8 clients must be able to connect to Windows Server 2012 servers sharing v4 queues. This involves the ability to procure and install a compatible driver, the ability to

synchronize configuration settings with those on the server, and the ability to print to a printer shared with a v4 driver.

The changes to the spooler necessary to implement these new sharing features were made on both the client side and the server side, and the following sections separate the changes accordingly.

Server Changes

The following are the changes made to Windows Server 2012 for Printer Sharing and Setup.

- **Compatibility Driver**

  To support driver sharing to operating systems prior to Windows 8, the Microsoft enhanced Point and Print compatibility driver is provided with the operating system. The Microsoft enhanced Point and Print compatibility driver is a v3 driver that interfaces with Windows Server 2012 print queues that are shared using v4 drivers, and emits XPS as its PDL. The Microsoft enhanced Point and Print compatibility driver is compatible with Windows 7 and Windows Vista; Windows versions prior to Windows Vista are unsupported and blocked (the server does not attempt to block clients based on version, but the Microsoft enhanced Point and Print compatibility driver determines whether it is compatible with the client during installation.)

  By default, when a client using a previous version of Windows queries the driver (via GetPrinterDriver API) used by a printer shared with a v4 driver, it will appear as the Microsoft enhanced Point and Print compatibility driver for the processor architecture of the client, regardless of what the actual driver in use by the server queue is. As a result, when a connection is made the Microsoft enhanced Point and Print compatibility driver is downloaded from the server and installed on the client.

  The v3 model remains unchanged in Windows Server 2012 and sharing and setup of a print queue using a v3 driver remains unchanged regardless if the client is Windows 8 or a previous version of Windows.

- **Disabling Driver Distribution**

  As described previously, the Microsoft enhanced Point and Print compatibility driver is provided to clients running a previous version of Windows to enable printing to a queue shared with a v4 driver.

  There is no actual API used by existing clients to retrieve driver binaries from the server. Existing servers make driver files available under the \\server\print$ share, and the information about what driver is in use for a queue and what files are needed is retrieved by clients through the GetPrinterDriverEx API. Existing Windows clients call GetPrinterDriverEx for the remote printer to get details about the remote driver, and then simply use filesystem APIs to copy the files from the server share to the client over SMB.

  Instead of disabling the print$ share since it is still necessary for distributing the Microsoft enhanced Point and Print compatibility driver to downlevel clients, GetPrinterDriverEx does not provide the driver information so that clients will be unable to use the returned information to find drivers under print$, and are forced to try other methods of procuring drivers.

- **Enhanced Point and Print**

Complementing the v4 driver model are changes to the printer sharing model. The updated printer sharing mechanism is referred to as enhanced Point and Print, and it allows print clients to print to v4 shares without downloading the manufacturer-provided device driver from the print server.

When a computer running Windows 8 connects to a shared print queue on the server, it checks locally for a v4 print driver that has a HardwareID that is the same as the PrinterDriverID defined in the server's print driver. It may also check Windows Update at connection time or later for a matching driver on Windows Update. If it finds a match, the client downloads that driver and connects using client-side rendering (CSR). If the driver includes a customized UI, then that UI will be shown to the user.

Otherwise, the client connects using the enhanced Point and Print driver. If the client does not have this driver already, it downloads it from the server. At this point, the client downloads configuration data files from the server and associates them with the client print queue. When the user prints, they are provided with a localized, Microsoft standard UI, unless the user acquires a printer extension or a Windows Store device app is automatically downloaded. The print job will be rendered on the client into XPS and include the user's settings as a PrintTicket. This is then sent to the server and rendered into PDL using the server's print driver.

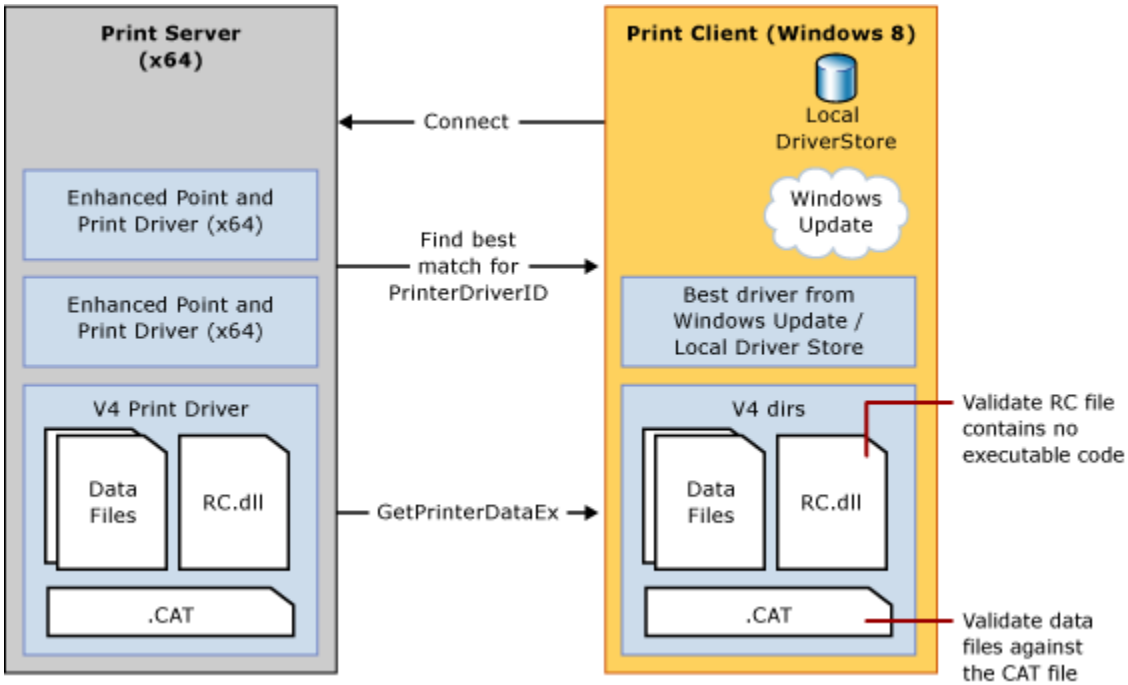

Figure 5: enhanced Point and Print Diagram

Client Changes

The following are the changes made to Windows 8 clients for Printer Sharing and Setup.

- **Driver Acquisition and Update Changes**

To achieve the goal of disabling driver distribution from Windows Server 2012 print servers, and in order to provide an optimal v4 driver experience, Windows 8 clients cannot use legacy v3 Point and Print Mechanisms to obtain a driver. Instead a new Point and Print mechanism has been implemented, leveraging Plug and Play for both the acquisition and installation of compatible print drivers, as well as the servicing/updating of these drivers.

Note
For more information on how to configure Windows Server Update Services (WSUS) to offer print drivers to clients, see Windows Server Update Services Overview.

For more technical details about printer sharing, see Printer Sharing Technical Details.

Appendix A: Terms and Definitions

The following table defines commonly used terms for Print and Document Services.

Term	Definition
Print Queue	A print device (either physical or virtual) managed by a print server. Print Queues may or may not be shared
Printer Connection	A print queue that acts as a proxy to a shared Print Queue on a print server. These are created on client computers as a fundamental part of the printer sharing process.
Compatibility Driver	The Microsoft enhanced Point and Print compatibility Driver. A v4 driver that provides generic baseline print functionality and experience for a client connecting to a print queue with a v4 driver on a server running Windows Server 2012. This driver emits XPS as its PDL which can be used natively by a server running Windows Server 2012, and this driver uses v3 or v4 configuration data to provide configuration options to the client.
v3	The printer driver model used for Windows versions from Windows 2000 through Windows 8.
v4	The new printer driver model in Windows Server 2012 and Windows 8.
CSR	Client-Side Rendering, which indicates that the rendering of a print job into PDL occurs on the

Term	Definition
	client (offloading work from the server).
SSR	Server-Side Rendering, which indicates that the rendering of a print job into PDL occurs on the server.
Localspl	The local print provider which maintains local print queues, as well as GUID printers (see below).
GUID Printer	A local printer object managed by localspl that provides the backing store for a print connection used by CSR for caching, client-side-rendering, offline printing, and more. This GUID printer is created as a fundamental part of connection establishment, and makes use of the driver procured during Point and Print.
Point and Print (P&P)	Not to be confused with Plug and Play, this describes functionality allowing a client computer to create a Printer Connection to a remote Print Queue; both acquiring the driver and setting up the rest of the connection state.
PDL	A page description language (PDL) is a language that describes the appearance of a printed page in a higher level than an actual output bitmap.

Printer Sharing Technical Details

This topic focuses on the various printer sharing technologies supported in Windows 8. It provides a quick overview of the various technologies, and why they might be used in certain scenarios.

There are a number of other features in Windows 8 that can be used to help augment or improve the printer sharing features in Windows 8. These features often provide more flexibility to IT administrators, or allow the user experience for printer sharing to be more tailored to a given environment.

Finally, this topic covers deployment scenarios for the various printer sharing features and technologies supported in Windows 8. This section is designed to allow IT administrators to start working with print servers quickly without needing to cross reference a number of other sources for basic print server functions.

Types of printer sharing in Windows Server 2012

Windows Server 2012 supports a number of different sharing mechanisms that enable deployment of print shares to meet various user and business needs. These include Enhanced Point and Print, Package Aware Point and Print, Legacy Point and Print, and Internet Print Protocol (IPP). In most cases, the administrator is not in direct control over the type of sharing that can be used; they are only able to "share" the printer. The drivers installed on the server, the client operating system version, and the connectivity mechanism dictates what type of sharing is available.

The following table shows which printer sharing technology is available based on operating system and driver versions being used from the print server side.

	V4 printer driver	**V3 Package Aware printer driver**	**V3 Legacy printer driver**
Windows 8/Windows Server 2012	Enhanced Point and Print	Package Aware Point and Print	Legacy Point and Print
Windows 7/Windows Server 2008 R2	N/A	Package Aware Point and Print	Legacy Point and Print
Windows Vista/Windows Server 2008	N/A	Package Aware Point and Print	Legacy Point and Print
Windows XP/Windows Server 2003	N/A	Legacy Point and Print	Legacy Point and Print

The client operating system is also a factor, and can impact the type of connection used beyond what the server makes available. This table shows which mechanism a client attempts to use when connecting to a specific type of print share.

	Enhanced Point and Print on server	**Package Aware Point and Print on server**	**Legacy Point and Print on server**
Windows 8/Windows Server 2012	Enhanced Point and Print	Package Aware Point and Print	Legacy Point and Print
Windows 7/Windows Server 2008 R2	Enhanced Point and Print	Package Aware Point and Print	Legacy Point and Print
Windows Vista/Windows Server 2008	Enhanced Point and Print	Package Aware Point and Print	Legacy Point and Print

Windows XP/Windows Server 2003	n/a	Legacy Point and Print	Legacy Point and Print

When the server has IPP sharing enabled, all print shares will automatically be available through IPP. When a Point and Print connection is possible, all Windows clients will create a Point and Print connection even if the user is attempting to create an IPP connection. Since IPP connections can be manually configured so that clients and servers can communicate even while using different drivers, any client can theoretically connect to any server using any driver capable of being installed on the client. It is important that the drivers are similar enough, so the server driver can accept and route the output of the client driver to the device. There is no programmatic way to test for driver compatibility in IPP scenarios.

Enhanced Point and Print

New to Windows 8/Windows Server 2012, enhanced Point and Print allows a server running Windows 8/Windows Server 2012 to share a printer with any other supported Windows client without the use of additional drivers. This feature is only available when the print server is using a v4 print driver for the print share.

Enhanced Point and Print enables v4 drivers to be used in Point and Print scenarios where client computers may not have any access or ability to use the same driver that is available on the print server. This is accomplished by using an enhanced Point and Print driver on Windows 8 and previous operating systems that can access the configuration of the printer from the print server and mimic the printer's capabilities without using any device specific code. Only Windows 8 and Windows Server 2012 print servers are capable of creating enhanced Point and Print shares, but any Windows client is capable of connecting to them.

A v4 driver can expose features that require user interface to configure. For standard features supported in the operating system, all of these features should be accessible to all client computers through standard user interface provided as part of that operating system's basic printer extension. For non-standard features, custom user interface must be provided by the Independent Hardware Vendor (IHV), and installed so the user can access those options. This user interface can be packaged in a Windows Store device app for use in the new Windows user interface print workflows, or as a Printer Extension when the user is printing through a desktop workflow. Printer Extensions can be bundled in a v4 print driver, or installed as a standalone application. All Windows Store device apps must be installed from the Windows App Store, or distributed through enterprise tools designed to deliver Windows Store apps.

The following is a high level diagram of enhanced Point and Print scenarios. The black arrows indicate the required Point and Print data flow to enable connections. The tan arrows show data flow for optional extensions that can extend or improve the printing experience. Since printer extensions are simply Windows applications, they can be installed from media or deployed just like any other application. The print server can also provide an URL that will be shown to users as part of the printer properties. The intension of this link is to provide a simple mechanism for print administrators to direct their users to Printer Extensions for their print share.

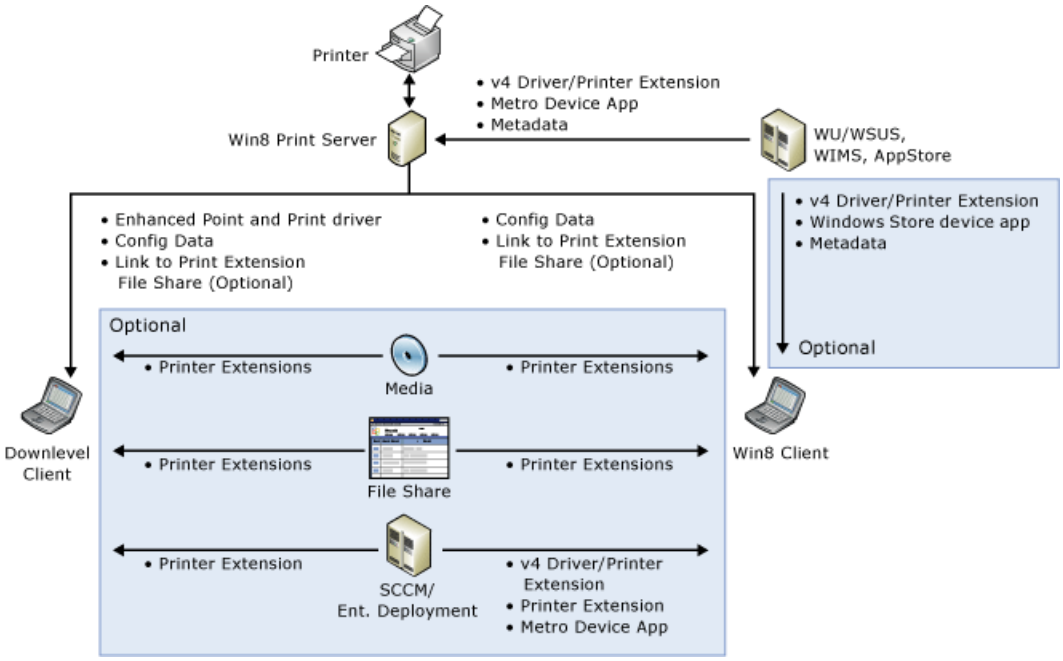

Figure 1: Overview of enhanced Point and Print data flows in Windows 8 printer sharing

Supported Servers

Windows 8/Windows Server 2012 (all versions including x86 and x64 architectures)

Supported Clients

- Windows 8/Windows Server 2012(All)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)

Advantages

- Only mechanism to share printers using v4 print drivers on the server
- Only mechanism to directly share a printer with an ARM architecture Windows client
- Simple driver management
- Enables basic printer sharing without any additional steps
- Integrates with Windows Store apps printing workflows on Windows 8 clients
- The print server no longer distributes device specific drivers to clients
- Windows 8 clients have an inbox Enhanced Point and Print driver to support connections

Limitations

- Not available while using v3 print drivers
- Requires distribution of Printer Extensions to access advanced features of the printer
- Server Side Rendering is forced on all downlevel clients

- Windows 8/Windows Server 2012 still requires access to a matching v4 print driver to enable Client Side Rendering
- Driver distribution for Windows 8/Windows Server 2012 clients handled through Windows Update/Windows Server Update Services (WSUS)

Package Aware Point and Print

Package Aware Point and Print was introduced in Windows Vista to enable print servers to distribute signed driver packages to client machines. This improves security, and eliminates the need for security prompts that are necessary in legacy Point and Print support. This feature is only enabled on Vista and later print servers when a Package Aware print driver is used for the print share. This feature will automatically downgrade to legacy point and print when downlevel clients connect to the share (Windows 2000, Windows XP, Windows Server 2003) or when non-Package Aware drivers are in use.

Supported Servers

- Windows 8/Windows Server 2012 (All versions, x86 and x64 architectures only)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)

Supported Clients

- Windows 8/Windows Server 2012 (All versions, x86 and x64 architectures only)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)
- Windows Server 2003 (All in legacy Point and Print mode)
- Windows XP (All in legacy Point and Print mode)

Advantages

- Allows print servers to distribute full driver packages to Package Aware clients
- Client Side Rendering is an option on all connected clients
- All user interface and rendering capabilities are part of the v3 driver payload

Limitations

- Print server administrator must stage print drivers on the server for all supported architectures
- When used in a Windows Store app print workflow, users are forced into a desktop experience when custom user interface is required
- Print server acts as a distribution point for v3 print drivers
- Does not support v4 print drivers
- Does not support ARM clients

Legacy Point and Print

Point and Print as it works today was introduced in Windows 2000, and remains largely unchanged. It is now referred to as Legacy Point and Print to differentiate the basic functionality from later extensions like Package Aware Point and Print and enhanced Point and Print. In Legacy Point and Print, the server's spooler provides the client with just enough information to copy down the required v3 printer driver files from the server, and what settings to use. Since the driver is not kept intact, there is no way for a client to validate that the driver has not been altered. Architecture specific drivers must be installed on the print server to support clients with different architectures.

Supported Servers

- Windows 8/Windows Server 2012 (All versions, x86 and x64 architectures)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)
- Windows Server 2003 (All)
- Windows XP (All)

Supported Clients

- Windows 8/Windows Server 2012 (All versions, x86 and x64 architectures)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)
- Windows Server 2003 (All)
- Windows XP (All)

Advantages

Does not require any special print driver or operating system support

Limitations

- Print server administrator must stage print drivers on the server for all supported architectures
- When used in a Windows Store app print workflow, users are forced into a desktop experience when custom user interface is required
- Print server acts as a distribution point for v3 print drivers
- Does not support v4 print drivers
- Does not support ARM clients directly (ARM clients may be able to connect to the printer directly if it is a network attached device)

Internet Print Protocol (IPP)

This mechanism for printing over HTTP was introduced in Windows 2000, and has largely been unchanged. Print Servers that have IPP enabled will automatically make a shared printer available through the relevant Point and Print sharing mechanism, as well as IPP. Windows also exposes additional functionality to enhance the usefulness of Windows IPP print servers. Print servers will publish a webpage that lists available IPP shares, and Windows clients connecting to those shares can download available v3 drivers directly from the print server. This is additional functionality that is not a part of the IPP 1.0 standard supported by Windows. Enhanced Point and Print shares will not be published to the print server's webpage in a way that allows the client to download the driver and automatically connect to the IPP share.

Windows clients can install IPP print shares from that webpage if the proper drivers have been staged onto the print server. In the case where proper drivers are not available on the server, or when a v4 print driver is being used, clients can still create an IPP connection by building the connection manually in the Add Printer Wizard (APW). When possible, clients will connect over Point and Print instead of IPP even when an IPP connection is attempted.

Point and Print in all of its variations is preferred by Windows over IPP. IPP is primarily used to solve one of two problems:

1. Allows non-Windows clients to connect to the print share
2. Allows communication over HTTP (Point and Print relies on RPC which may not be available in some network configurations)

Supported Servers

- Windows 8/Windows Server 2012 (Server versions with Print Role and IPP enabled)
- Windows Server 2008 R2 (All versions with Print Role and IPP enabled)
- Windows Server 2008 (All versions with Print Role and IPP enabled)
- Windows Server 2003 (All versions with Print Role and IPP enabled)

Supported Clients

- Windows 8/Windows Server 2012 (All versions, ARM not supported through server page)
- Windows Server 2008 R2 (All)
- Windows 7 (All)
- Windows Server 2008 (All)
- Windows Vista (All)
- Windows Server 2003 (All)
- Windows XP (All)

Advantages

- Allows access to printers over HTTP
- Allows non-Windows clients to connect to a print share

Limitations

- Driver used on client must be compatible with driver used on server

- May be subject to v3 or v4 limitations depending on the client driver used to connect to the print share
- Print server acts as a distribution point for v3 print drivers
- Does not support v4 print drivers through the server page

Technologies to Augment Printer Sharing

Branch Office Direct Printing

New to Windows Server 2012, Branch Office Direct Printing is an extension to existing printer sharing technologies. Standard mechanisms are still used to create and maintain a Point and Print connection (enhanced, Package Aware and legacy Point and Print are supported). Once the connection is established, client computers submit print jobs directly to the print device instead of submitting print jobs to the print server.

The goal of Branch Office Direct Printing is to help support network topologies where the print device and client computer are on the same local network, but the print server is separated by a slow or intermittent network connection. Once a client computer has established a Point and Print connection to the server, it is capable of printing directly to the device even if contact to the print server is lost. The print job is not uploaded to the server through a slow/intermittent connection, only to have it resent across the same network connection back to the device. This significantly reduces network traffic, and improves the resiliency of centralized print servers managing geographically separate branch offices.

If Branch Office Direct Printing is enabled, it overrides the default Server Side Rendering (SSR) mode for systems like laptops and tablets. Typically, these portable systems running Windows 8 force SSR to conserve battery power while printing. When Branch Office Direct Printing is enabled on a queue, it is assumed that bandwidth is more valuable than extended battery life, and the portable client computers will switch to Client Side Rendering (CSR) to take advantage of the Branch Office Direct Printing feature. Connections using enhanced Point and Print must be using the device specific v4 driver, which enables CSR, to work with Branch Office Direct Printing environments. Connections using the enhanced Point and Print driver (SSR only) will always fall back to printing through the print server.

For more information about Branch Office Direct Printing, see Branch Office Direct Printing Overview, and Branch Office Direct Printing Technical Details.

WSD Secure Printing

New in Windows Server 2012, print servers have the option to connect to Web Services for Devices (WSD) devices securely without additional network technologies like IPsec. The print device must support WSD Secure Printing, and a domain is required to manage permissions. The Active Directory for the Domain acts as a trust arbitrator between the print server and the print device. While WSD Secure Print does not directly impact sharing, it does allow for print

servers to create a private secure channel to the device on the network. This prevents *Man in the Middle* attacks from intercepting the print job, and prevents snooping of the printed content on the network between the print server and the print device. This feature has no impact on client connections to the print server.

Supported Servers

- Windows Server 2012 (All versions)

Supported Clients

- N/A

Advantages

- Enables secure network communication between print servers and print devices
- Provides a mechanism to ensure the identity of the network resources involved (print server and print device)

Limitations

- Only available on servers running Windows Server 2012
- Only available with devices that support WSD Secure Printing
- Requires a Domain (Active Directory) to work

AD Printer Publishing

Introduced in Windows 2000, AD printer publishing allows servers joined to a domain to publish available print shares into an AD print directory. This directory is accessible to other client computers in the domain, and they can easily discover published printers through various operating system entry points.

Supported Servers

- Windows Server 2012 and Windows 8 (All versions that support domains)
- Windows Server 2008 R2 (All)
- Windows 7 (All versions that support domains)
- Windows Server 2008 (All)
- Windows Vista (All versions that support domains)
- Windows Server 2003 (All)
- Windows XP (All versions that support domains)
-

Supported Clients

- Windows Server 2012 and Windows 8 (All versions that support domains)
- Windows Server 2008 R2 (All)
- Windows 7 (All versions that support domains)
- Windows Server 2008 (All)
- Windows Vista (All versions that support domains)
- Windows Server 2003 (All)

- Windows XP (All versions that support domains)
-

Advantages

- Allows easy discovery of printers in a domain environment, and is not subject to subnet limitations like network attached printers
- Provides IT administrators a central directory to manage all approved printers

Limitations

- Client computers must be part of the domain to read the directory
- Client computers must still be able to create Point and Print connections to the print server to leverage the published printer

Push Printer Connections

Introduced in Windows 2000, Push Printer Connections allow an IT administrator to create printer connections on client computers in a domain through Group Policy. Pushed printer connections are added to systems in the domain when the user logs in or when the system restarts depending on the specific deployment script used. These connections are not removable by the user, and provide IT administrators an easy way to provide printer access without requiring users to create their own connections manually.

Supported Servers

- See Point and Print sections for viable server list; any Point and Print connection can be pushed

Supported Clients

- Windows Server 2012 and Windows 8 (All versions that support domains)
- Windows Server 2008 R2 (All)
- Windows 7 (All versions that support domains)
- Windows Server 2008 (All)
- Windows Vista (All versions that support domains)
- Windows Server 2003 (All)
- Windows XP (All versions that support domains)
-

Advantages

- Allows IT administrators to ensure all client computers/users have standard printers available automatically
- Users in the domain do not need to manually create printer connections to use domain print shares
- Users cannot delete pushed printer connections without Domain Admin credentials

Limitations

- Pushed printers are created as a per computer connection, and are not managed per user

Deployment Scenarios

Legacy and Package Aware Point and Print

These technologies have not changed in Windows Server 2012, and will continue to operate as they did in Windows 7 and Windows Server 2008 R2. Extensive information is already available on how to setup, maintain, and manage these types of print servers and connections.

IPP

IPP is largely unchanged, but some combinations of print shares and client operating systems can create new deployment issues with the introduction of ARM support, and v4 print drivers.

The IPP protocol does not support a standard for the client and server to share or exchange driver information. Windows has custom extensions it uses to exchange driver information during the setup of an IPP connection. These additional features allow Windows clients and servers to exchange v3 drivers in some situations. This has enabled the automatic deployment of v3 print drivers to support IPP in a fashion very similar to legacy point and print. IPP enabled print servers host a web page that lists available IPP print shares, and allows client computers that support the ActiveX control to install those connections from the web page without any additional configuration. Connections can also be made through the user interface or programmatically without relying on the driver provided by the print server.

The extensions made to IPP to enable v3 driver distribution are not compatible with enhanced Point and Print, and do not extend to v4 drivers. Print shares using v4 drivers will not be installable through the server's IPP print share page. Connections to v4 powered IPP print shares must be done manually through the user interface, or programmatically through an app or script.

Using v3 print drivers for IPP

Publishing IPP printers for queues using v3 print drivers is unchanged from Windows Server 2008 R2. All previous documentation is valid for Windows Server 2012.

Similarly, x86 and x64 client computers running Windows 8 will continue to support IPP connections to v3 enabled print servers without changes. The ActiveX control required to connect to IPP print shares from the server's IPP print share page will not run in Internet Explorer, and must be used from Internet Explorer for the desktop on Windows 8.

Using v4 print drivers on the Server

For print queues with v4 print drivers, sharing the queue via IPP remains unchanged. Once the IPP services are installed and running on the print server, all shared queues with v4 print drivers (enhanced Point and Print) will be made available for IPP connections as well.

Unlike shares with v3 print drivers, client computers will not be able to connect through the server's IPP print share page. All connections to v4 IPP print shares must be done manually through the user interface, or programmatically created through an app or script.

Whether you are manually connecting to a v4 IPP share, or scripting it for broad deployment, the main requirement is that client computers have access to a print driver that can generate device-ready output for the target device. It is possible for the driver on the server to interact with the print information sent by the client computer, so testing is essential before deploying any manually configured IPP connections. As long as the driver on the server and device generate the same PDL (PostScript, PLC6, XPS, and so on.) there should be very little chance for problems, but print drivers can be quite large and complicated. Without testing, there is no guarantee that two different drivers will be able to work with each other through an IPP connection. Ideally, the same driver would be used on both the server and client computer, but this may not be possible depending on the architectures and operating system versions being used. The following is a list of scenarios where manual configuration is required to create connections because the client and server cannot use the same driver.

1. Windows 7 and earlier systems must use a v3 driver to connect to a v4 IPP print share on Windows Server 2012
2. Client computers running Windows 8 must be configured to connect to a v4 IPP print share on Windows Server 2012, but the same driver can be used on the client and server
3. Windows ARM clients must use a v4 driver to connect to a v3 IPP print share on any print server

Even when it is possible for a client computer and server to use the same driver to connect to an IPP print share, it is not required. The same process can be used to specify a driver when creating a manual IPP connection, even when the client could automatically get the driver from the server. This is true for Windows 8, as well as for computers running previous Windows operating systems.

Manually Creating an IPP Connection in the User Interface

IPP Connections can be created manually using the Add Printer Wizard. To start the Add Printer Wizard, type **Advanced Printer Setup** at the start screen, click **Settings**, and the click the **Advanced Printer Setup** tile.

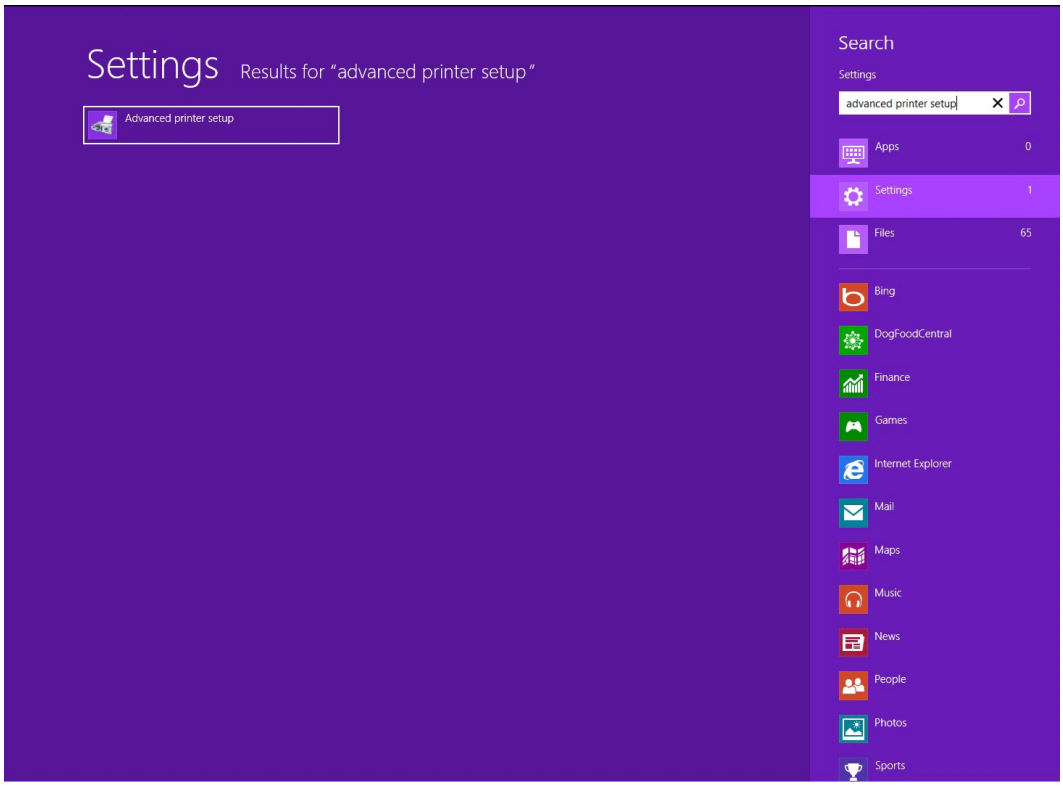

Figure 2: Advanced Printer Setup

There are several different start screens depending on your system's configuration. Click **The printer that I want isn't listed**. Under **Find a printer by other options**, click **Select a shared printer by name** and type the URL for the IPP print share.

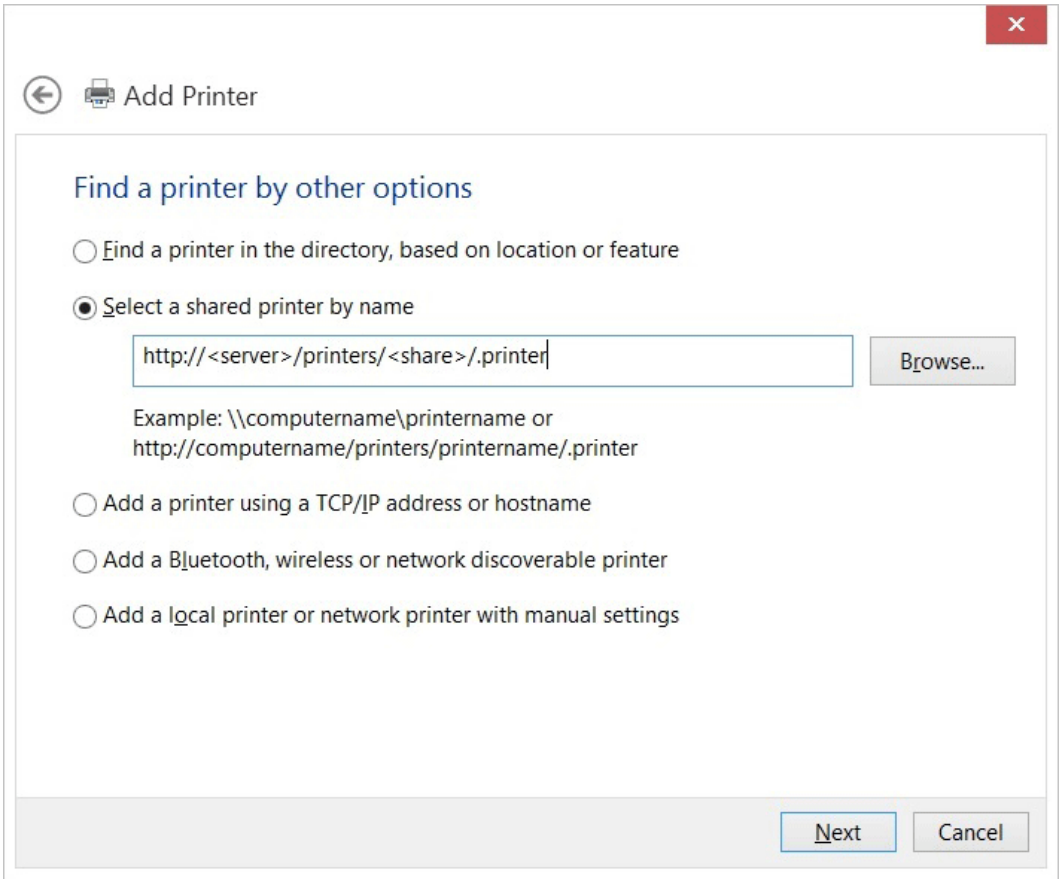

Figure 3: Add Printer

The URL for the IPP share uses the following format:

```
http://<server>/printers/<share>/.printer
```

Where **<server>** is replaced with the name of the print server, and **<share>** is replaced with the name of the print share.

If the server does not have a viable driver for the client (for example, different architecture, the server is using a v4 driver, or if it is not a Windows print server), the user is prompted to choose a driver on the local system to use. After selecting a driver, the connection is complete. Users can choose to print a test page, or make the new IPP printer connection their default printer.

Enhanced Point and Print

Enhanced Point and Print is new in Windows Server 2012, and introduces a lot of new concepts into printer sharing. One of the ultimate goals is to simplify the administration of Print Servers moving forward.

The functionality of enhanced Point and Print can be broken down into three major areas:

Figure 2: Basic v4 sharing

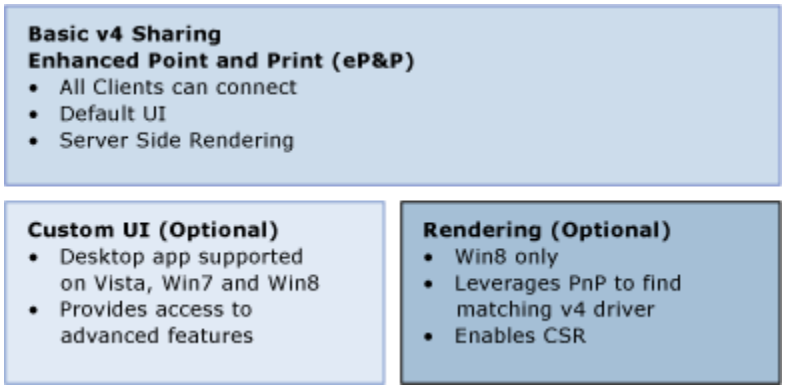

1. **Basic** – This is the foundation of enhanced Point and Print. It provides universal access to common functionality using default user interface and rendering is performed on the print server.
2. **Custom user interface** (Optional) – Advanced features not supported in the default user interface can still be accessed if the client has an application installed that can present the proper user interface. There are separate desktop and Windows Store app versions depending on the workflow being used to print.
3. **Rendering** (Optional) – Windows 8 print clients can identify and install v4 print drivers that match those being used on the print server, enabling client side rendering (CSR).

Basic Deployment of enhanced Point and Print shares

Enhanced Point and Print relies on two drivers:

1. The enhanced Point and Print driver for use on client computers running Windows 8.
2. The enhanced Point and Print compatibility driver for use on client computers running previous Windows operating systems.

Both of these drivers ship with all versions of Windows 8. 64-bit versions of Windows 8 also include the 32-bit version of the enhanced Point and Print compatibility driver, enabling any client computer running a previous Windows operating system to connect to a v4 share. Client computers running a previous Windows operating system will download the enhanced Point and Print compatibility driver directly from the print server, just like they would download a driver from the server for any Point and Print share.

Client computers using either of these drivers to connect to a print server will be given a basic user experience by default. Rendering will always happen on the server (SSR mode). Any print share hosted on a Windows Server 2012 print server using a v4 driver will be shared through this mechanism. If the basic user experience is sufficient, no other work is required to make this print share available to all supported client computers running Windows 8 and previous Windows operating systems, regardless of operating system architecture (IA64 is not supported through enhanced Point and Print).

Printers can be shared through the Print Management Console, print spooler API's, or directly through Printer Properties on the device just as an administrator would share any device using a

v3 driver. There is no choice to enable enhanced Point and Print specifically, since that choice is dictated by the driver model (v3 versus v4) used to connect to the print device locally on the print server.

Custom User Interface for enhanced Point and Print Shares

For situations where additional user interface is required to expose advanced functionality in the device that is not available in the basic user interface, IT administrators will need to deploy additional applications to complete the user experience. There are two types of applications that can add additional user interface to an enhanced Point and Print connection:

1. Printer extension applications for desktop printing workflows
2. Windows Store device apps for the new Windows user interface printing workflows

The additional user interface is provided as a Windows application that can register with the spooler to handle certain aspects of the print driver's user interface. This can extend the functionality of the basic user interface by adding a customized Printer Preferences page, as well as providing custom user interface for dealing with driver defined events. Each app is only used in its respective workflow (desktop or new Windows user interface) for displaying Print Preferences. Custom notifications are always handled by the Windows Store device app first if it is installed since there is a unified notification system across the new Windows user interface and desktop experiences.

Note

Connections to enhanced Point and Print (v4) shares are not supported on Windows XP and Windows Server 2003.

Deploying Desktop Printer Extension Apps as an Application

Printer extension apps are desktop applications that can be deployed to Windows 8, Windows 7, Windows Server 2008 R2, Windows Server 2008, and Windows Vista client computers to extend the basic desktop user interface for an enhanced Point and Print share. These apps can be deployed as part of a client image, or using standard application deployment tools. Printer Extension apps can be installed before or after a connection to the print share is made. Printer extension apps can also be installed manually from a share or media.

Deploying Desktop Printer Extension Apps in a v4 Driver

Client computers running Windows 8 can also receive a Printer extension app as part of a v4 driver installation. This can introduce complications if other deployment mechanisms are also being used. Refer to Enabling CSR for Client Computers Running <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">win8_client_2</token> for more details.

Deploying Desktop Printer Extension Apps in Print Preferences

IT administrators can update the metadata on the print share to provide users an URL to obtain a printer extension app. The clickable link for the URL will be on the Printer Preferences page when users have a connection to an enhanced Point and Print share with this value set. If a

printer extension app is already installed on the client computer, it will receive that user interface instead of being given a link to the share. The registry key for this URL is:

`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Print\Printers\<print queue name>\PrinterDriverData\PrinterExtensionUrl`

All v4 print drivers can specify a default URL provided by the manufacturer that links users to a manufacturer website with instructions for installing a Printer Extension app. By default, any manufacturer specified URL will be provided to users connecting to a v4 print share if they do not have any other Printer Extension installed.

The URL specified can link to a file share with the printer extension on it, a web page with help content, or any other valid URL target.

Blocking Desktop Printer Extensions

Some enterprises may also choose to block printer extensions altogether. This is supported by a group policy called `Computer Configuration\Administrative Templates\Printers\Do not allow v4 printer drivers to show printer extension applications.`

Deploying Windows Store Device Apps

Windows Store device apps are used to extend the basic user interface for printing from a Windows Store app. Since it is specific to Windows Store apps, it can only be deployed to client computers running Windows 8, and it is distributed automatically through the Windows Store.

Enabling CSR for Client Computers Running Windows 8

In addition to using the enhanced Point and Print driver for basic SSR functionality, client computers running Windows 8 can use a v4 driver to render content on the client, as well as gain access to any printer extension app (desktop user interface extensions) that may be packaged in the driver.

All v4 drivers must report a Printer Driver ID to the operating system. When a print queue is shared through enhanced Point and Print, the Printer Driver ID provided by the print driver is shared with the Windows 8 print client. The Windows 8 print client will create a virtual PnP device using the Printer Driver ID as the virtual device's Hardware ID (HWID). This triggers PnP to locate and install a driver for the device.

A matching driver is any v4 driver that reports to Windows 8 in its INF that it supports the Printer Driver ID of the driver used on the server. If a matching v4 driver is found, it is installed using the same mechanisms as a PnP device installation. CSR printing is enabled by default when a v4 driver is running on a client computer running Windows 8, unless CSR is disabled by some other policy (Group Policy, print server settings, power management profile, and so on.).

Since printer extension apps operate on a Last-Installed-Is-Active mode, there can be contention between printer extension apps installed by a driver, and those deployed through other mechanisms. To minimize conflicts in enterprises, it is often preferable to only use one mechanism to distribute printer extension apps to clients.

The following are recommendations that can help eliminate the use of unintended printer extension apps:

- Do not distribute v4 print drivers that include printer extension apps, and always rely on standard software deployment tools to support the distribution of printer extension apps.
- Use standard deployment tools to distribute printer extension apps to client computers running operating systems prior to Windows 8, and leverage the driver package's printer extension app for all client computers running Windows 8.
- Create deployment rules for printer extension apps that can detect when a different printer extension app has been installed, and replace it.
- Only distribute printer extension apps that match (the driver uses the same printer extension app that is being deployed).

Typically, v4 print drivers are published to Windows Update and will be installed automatically when a connection is made if the client has Windows Update access, and if a matching driver is available. In an enterprise, there are other mechanisms that can be used; Windows Server Update Services (WSUS), DevicePath, and standard deployment tools.

Deploying v4 drivers with WSUS

If using the live Windows Update site is not an option for v4 driver distribution, then deploying drivers through Windows Server Update Services (WSUS) is the recommended option. Client computers running Windows 8 can be configured to search WSUS for driver updates, including v4 drivers used in Windows 8 for enhanced Point and Print connections

Client computers running Windows 8 can be configured to check WSUS for available drivers by configuring group policy. The systems must be configured to search Windows Update for drivers automatically, and to use the WSUS server. The new policy for Windows 8 is **Specify the search server for device driver updates**, available under `Computer Configuration>Administrative Templates>System>Device Installation`. By default, WSUS is never searched. The policy can be enabled to allow Windows Update only, WSUS only, and to search WSUS followed by Windows Update if no driver is available on WSUS.

The option to search WSUS only will prevent PnP from searching the Windows Update live site, but still allow the administrator to provide additional drivers on their internal WSUS server. The option to search WSUS first, and then Windows Update allows IT administrators to still leverage the breadth of support on Windows Update while allowing specific drivers to be deployed through their internal WSUS server first.

WSUS server is a standalone server role, and it is used as a feature of System Center Configuration Manager. General information about the installation and use of WSUS is available online at TechNet:

Windows Server Update Services

WSUS allows administrators to easily search and import driver packages from the Microsoft Update Catalog. If the v4 driver that needs to be deployed is not available in the Microsoft Update Catalog, it can still be distributed through WSUS using Local Publishing. Using Local Publishing in WSUS requires a high level of technical understanding and the ability to author update packages. This option is not recommended if a suitable driver is available on the Microsoft Update Catalog or in situations where Basic v4 Sharing with Customized user interface would be sufficient. Information on the use of Local Publishing is available on MSDN:

Local Publishing

For more information about deploying WSUS, see Windows Server Update Services Overview.

Deploying v4 drivers with DevicePath

DevicePath is a feature in PnP behavior that allows additional drivers to be discovered on a file share or local folder. DevicePath is searched after PnP searches the local driver store, but before a Windows Update check is started. DevicePath is always searched for all PnP driver searches, and it is not indexed in any way. Because of this, performance of all PnP driver searches will be degraded if the connection to the folder is slow, or if a large number of drivers are published on the share. DevicePath use should be restricted to cases where a very small number of drivers need to be distributed, and Windows Update/WSUS are not viable options.

The first step in using a DevicePath deployment strategy is to create a file share on the network that is accessible to all of the necessary client computers. Once that is done, each client computer needs to be configured to search the DevicePath. Information on setting the DevicePath is available on MSDN at:

Modify the DevicePath Registry Key

Drivers published on the DevicePath share must still be signed to avoid prompts and elevations during installation, or additional configuration of the clients may be necessary to suppress elevations and prompts.

Deploying v4 drivers with standard deployment tools

Most deployment tools are designed to distribute applications. To deploy a v4 driver as an application, it must be wrapped in an install package that will add the driver to the local driver store when it is run on the client computer. No other actions are required. Once the driver is in the driver store, it will be available for PnP to discover it when the connection is made.

If your deployment tool supports driver deployments directly to client computers, there is no need to create an install package for the v4 driver.

The v4 driver must be in the local driver store when the connection is made to ensure it is used by the local spooler when creating an enhanced Point and Print connection to a print share. Drivers deployed to the driver store after the connection is made will not be used unless the connection is recreated, or a driver update is invoked on the queue.

Windows on ARM in Enterprises

The two main mechanisms for connecting to printers in an enterprise environment for Windows on Arm clients will be to connect to the printer directly, and to connect through a Windows Server 2012 print server using enhanced Point and Print. Other options are available, but extra care must be taken for those mechanisms to function properly.

Connecting Directly to the Printer

Windows 8 on ARM clients can connect directly to printers in enterprise environments using v4 class drivers. To support direct connections, the following conditions must be met:

- Ensure the printer is supported by Windows Server 2012 drivers
- Ensure the device is configured to accept print jobs from any IP address/user on the network
- Ensure the clients can communicate with the printer on the network (they must be on the same network, and not be blocked by any other technology like firewalls, or IPsec).
- Publish the IP address of the printer so that clients can connect to it directly (typically requires the user to go through the Advanced Printer Wizard since printers are rarely on the same subnet as the user). Links to the printer's IP address can be configured to create a connection automatically.

The disadvantages of direct connection is that IT Administrators do not have the option of using a print server to audit print jobs, service drivers, or establish custom default settings profiles. Each client computer independently connects and submits print jobs directly to the print device. Any auditing or archival solutions will need to be enabled directly on the client computer or the print device.

Connecting to v4 Shares

Windows 8 running on ARM systems can connect to any v4 print share using enhanced Point and Print. All of the deployment options available to other Windows 8 clients apply to Windows 8 running on ARM clients, so refer to the section on enhanced Point and Print for more details. The one exception is that ARM systems can only use the print drivers that ship with the operating system as well as basic enhanced Point and Print support. This may restrict the choices available to allow CSR, and deployment of custom user interface will not be possible through the Windows Update/WSUS driver channel.

Since an ARM client cannot join the domain, it will not be able to retrieve the list of published print shares from the Active Directory (AD). Connections to shares can still be established by opening the URL for the share, clicking on a link to the share, or by navigating to the print share in Explorer. Connections can also be created through the APW, but users will still need to know the URL for the print share.

Connecting to v3 Shares

Windows 8 on ARM systems cannot connect to v3 print shares (legacy Point and Print, Package Aware Point and Print). If a connection to a v3 print share is attempted, the Windows 8 on ARM system may prompt the user to provide a driver. Since there are no v3 drivers for ARM, this dialog should be dismissed.

Connecting to IPP Shares

ARM systems can also be configured to print to IPP print shares, if a compatible device driver is available in the local driver store. This follows the same manual process for creating an IPP connections described in the IPP Deployment Scenarios section of this topic. Creating connections to IPP print shares from an ARM client are not supported through the server's print share web page.

Branch Office Printing

For more information about Branch Office Direct Printing, see Branch Office Direct Printing Overview.

WSD Secure Printing

WSD Secure Printing is intended to enable secure channel communication between a Windows Server 2012 print server and a WSD Secure Print enabled device. This does not rely on IPsec, or any other additional network services to provide a secure channel. Instead, it uses Secure WSD to establish a communication channel between the device and the print server. Unlike standard HTTPS communication, the WSD Secure Print feature can leverage self-signed certificates to secure the communication channel by relying on Active Directory (AD) to arbitrate trust between the two end points.

Requirements for deploying a WSD Secure Printer

- A print device that supports the WSD Secure Print feature
- A Windows Server 2012 print server
- An Active Directory Domain that both the server and device can join

Steps to deploy a WSD Secure Printer

To deploy a WSD Secure Printer, an IT administrator must:

1. Ensure the Active Directory has the WSD Secure Print updates. Windows Server 2012 Active Directory already contains the updates. The specific script to add the appropriate fields to an older Active Directory deployment is available in the following MSDN topic: Building a WSD Secure Print Device for Windows
2. Create a machine account in the domain for the WSD Secure Print device.
3. Ensure the WSD Secure Print device is attached to the network, and is functional.
4. Provide the WSD Secure Print device with the computer credentials per the manufacturer's instructions. The device will automatically authenticate with the Active Directory and update its computer account with the necessary information.
5. Add the WSD Secure Print Device to the print server using directed discovery. The print server must be in the same domain as the WSD Secure Print device, and have access to the Active Directory at the time making the connection. Adding a device through directed discovery can be accomplished a number of ways, but the primary user interface path is described below:
    a. In the Print Management Console, select the print server you want to use.
    b. Under **Action** in the tool bar, choose **Add Printer**.
    c. Choose the radio button for **Add a TCP/IP or Web Services Printer by IP address or hostname**, and click **Next**.

d. In the drop down box, select **Web Service Secure Printer**, and specify the IP address or host name of the device in the text box below. When done, click **Next**.

6. Share the WSD Secure Print queue normally, and manages it like any other print queue. Once the connection is made, there are no special steps in managing a Secure WSD Print device.

Notes on WSD Secure Print scenarios

- Multicast discovery is not supported over HTTPS. All connections to WSD Secure Print printers must be done through directed discovery.
- If the WSD Secure Print device is using a self-signed certificate, it will not be possible for the device to provide events back to the print server. The print spooler will poll the WSD Secure Print device for events and status directly. This can introduce minor delays in status updates on the print server and clients. This is expected behavior.
- WSD Secure Print devices do not require the deployment of trusted certificates to secure the connection between the print server and the device. Supporting deployment of trusted certificates can enable two-way communication between the device and print server, providing better status and event handling, but this is an optional feature. Check with your manufacturer if your policies require the use of trusted certificates.

Migrate Print and Document Services to Windows Server 2012

Overview

This article provides guidance to migrate a print server running Windows Server® 2003, Windows Server 2003 R2, Windows Server 2008, or Windows Server 2008 R2 operating systems to a server running the Windows Server 2012 operating system with the Print and Document Services role installed. This includes cross-architecture and stand-alone migrations, as well as configurations for a server in a cluster. This document provides step-by-step instructions for migrating from old hardware to new hardware and consolidating print servers.

Print and Document Services enables print server tools and configures the server to act as a print server. Print and Document Services is not dependent on any other features or roles. However, some specific network configurations, clients, and hardware may require you to install additional features or enable certain services.

This guide provides you with instructions for migrating an existing print server to a server that is running Windows Server 2012. This guide does not contain instructions for migration when the source server is running multiple roles. If your server is running multiple roles, we recommend that you design a custom migration procedure specific to your server environment, based on the information provided in other role migration guides. Migration guides for additional roles are available in the Windows Server TechCenter.

Caution

If your source server is running multiple roles, some migration steps in this guide, such as those for computer name and IP configuration, can cause other roles that are running on the source server to fail.

To manage the migration process, use one of the following:

- The Printer Migration Wizard, which you access through Print Management, a snap-in in Microsoft Management Console (MMC).
- The Printbrm.exe command-line tool.

You can perform the migration locally or remotely, and from either a client computer or server.

Important

As a best practice, run the Printer Migration Wizard or Printbrm.exe from a computer running Windows Server 2012 or Windows 8 to ensure that you are using the newest version of the migration tools that have the latest updates and features. You can run these tools either locally on the server or remotely from any other computer running Windows Server 2012 or Windows 8.

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

For more information about installing and using these tools, see Access the migration tools.

Notes

The Print Management snap-in is not available in Windows Server 2003. However, it is available in Windows Vista Enterprise and Windows Vista Ultimate, which enables you to migrate from Windows Server 2003. It is also available in Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012. For more information about migrating from Windows Server 2003, see Preparing to Migrate.

The Print Management snap-in and the Printbrm.exe command-line tool are not available for the Server Core installation option. To migrate from a print server running on a Server Core installation, use a server running the Printer Migration Wizard, Windows Vista Enterprise, or Windows Vista Ultimate. For more information about migrating from a server running a Server Core installation, see Preparing to Migrate.

There is no equivalent to Print and Document Services for Windows® client operating systems.

You can migrate Print and Document Services from the destination server or from any client with the following:

- The Printer Migration Wizard (provided that the client is running one of the supported operating systems listed in the Supported operating systems matrix).
- Remote access to the destination server.
- Access to the printer settings file created when you back up the source server.

Note

All commands in this guide are case-insensitive unless specifically noted.

About this guide

This guide is designed as a step-by-step tutorial for migrating print servers.

Target audience

This document is intended for IT administrators and other knowledge workers who are responsible for the operation and deployment of print servers in a managed environment.

What this guide does not provide

This document does not provide guidance for the following:

- Upgrading roles on the same computer
- Migrating printer configurations during client installations of Windows
- Migrating settings for a server that is not being used as a print server
- Recovering server information that was not properly saved prior to migration for in-place upgrades
- Instructions for migrating more than Print and Document Services

Supported migration scenarios

You must have access to the Printer Migration Wizard to migrate print servers. For more information about supported scenarios and limitations, see Preparing to Migrate.

Supported operating systems

The following table outlines the supported operating systems for migration covered in this guide.

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2003 with Service Pack 2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x86- or x64-based	Windows Server 2003 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based

Source server processor	Source server operating system	Destination server operating system	Destination server processor
x86- or x64-based	Windows Server 2008, both full and Server Core installation options	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x64-based	Windows Server 2008 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
x64-based	Server Core installation option of Windows Server 2008 R2	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
X64-based	Windows Server 2012	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based
X64-based	Server Core and MinShell installation options of Windows Server 2012	Windows Server 2012, including full, MinShell, and Server Core installation options	x64-based

The versions of operating systems shown in the preceding table are the oldest combinations of operating systems and service packs that are supported. Newer service packs, if available, are supported.

The Foundation, Standard, Enterprise, and Datacenter editions of Windows Server are supported on full, Server Core, and MinShell installation options, as either source or destination servers.

All versions of Windows Server 2012 are x64-based. Migrating to an x86-based server is not supported.

Migrations between physical operating systems and virtual operating systems are supported.

Note
Both x86-based and x64-based migrations are supported for Windows Server 2003 and Windows Server 2008. All editions of Windows Server 2012 are x64-based.

You might prefer the migration process, rather than an upgrade, even when the hardware is native x64-based. An example would be a case where there is increased use of the server and there is a server role split (in which the source server has more than one server role)—and you decide to separate the roles onto several additional x64-based servers. In this case, migration of individual server roles to other servers might be the best solution.

The server administrator can choose which parts of an existing installation to migrate, but along with the server role, this usually includes configuration, data, system identity, and operating system settings.

Migration from a source server to a destination server that is running an operating system in a different system UI language (that is, the installed language) than the source server is not supported. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French language to a computer that is running Windows Server 2012 in the German language.

Note

The system UI language is the language of the localized installation package that was used to set up the Windows operating system.

Cross-architecture migration (such as migrating from an x86-based server to an x64-based server) is supported. The source server must have print queue drivers installed for both the source and destination server architectures. If a print queue does not have a driver for the destination server architecture, then it will not be migrated. Similarly, verify that the destination server contains drivers for each supported architecture.

Supported role configurations

The Supported operating systems matrix provides a complete listing of the supported migration scenarios.

Some migration scenarios require additional preparation. For more information about these scenarios, see Appendix B - Additional Destination Server Scenarios.

Supported role services and features

The Printer Migration Wizard migrates:

- Print queues.
- Shared printer settings.
- Printer drivers in use by the print spooler.
- Any security settings specific to the installed printer.

Migrating from x86-based to x64-based v3 printer drivers

There are several things you must consider when managing migrations using v3 print drivers. The first is that a print queue cannot function without the native printer driver for the server

architecture (x86 or x64) on which it exists. Since Windows Server 2012 is a 64-bit only operating system, it is important that you have 64-bit drivers installed for all of your printers if you are migrating from a 32-bit system. The most difficult transition is from 32-bit to 64-bit servers in an organization with 32-bit clients since it is common to have third-party 32-bit printer drivers that do not have 64-bit equivalents. During the print configuration restoration for cross-platform scenarios, if the backup file does not contain driver binaries for the platform of the target server, the drivers will be installed from the target server's driver store, if available.

As a best practice, when migrating from x86-based to x64-based v3 drivers:

1. Verify that x64-based versions of the drivers are available.
2. If you are unable to verify their availability, back up the source server before the migration.
3. Install the x64-based drivers on the source server so that you can determine if there any problems or conflicts before the migration process.
4. If there are conflicts or problems on the destination server after the migration, roll back the migration. For more information, see "Roll back migration on the source server" in Post-Migration Tasks.

Unsupported scenarios

The Printer Migration Wizard does not migrate the following:

- Other services or settings that specific printers may rely on, such as Line Printer Remote (LPR), Internet Printer Protocol (IPP), or Web Services on Devices (WSD). You must enable or install these features before restoring the source print server configuration. For additional information, see "Roll back migration on the source server" in Post-Migration Tasks.
- Local bus printers (LPT and USB), although they are shown during backup. For additional information, see Appendix B - Additional Destination Server Scenarios.
- Plug and play printers. However, plug and play printer drivers will be migrated. For additional information, see Appendix B - Additional Destination Server Scenarios.
- Any print jobs currently in the printer queue.
- Any system or print administrators, or permissions. If you want to retain the same system or print administrators on the destination server as on the source server, you will need to manually add these administrators to the destination server.

Print and Document Services migration overview

While the original server is still running, use the Printer Migration Wizard or the Printbrm.exe command-line tool to export or back up the print information (such as settings, queues, and drivers) in a printer settings file. Then, import or restore this backup image to a destination server running Windows Server 2012 that has been configured to run as a print server.

Some migration scenarios require additional preparation. For more information about these scenarios, see Appendix B - Additional Destination Server Scenarios.

To migrate printers from a server running Windows Server 2003 or a Server Core installation to a server running Windows Server 2008 R2, you must use a computer running the Printer Migration

Wizard to remotely manage the server running Windows Server 2003 or a Server Core installation. Using this computer, you can store the printer settings file (containing information about the printers you want to migrate, such as settings, queues, and drivers) from the server running Windows Server 2003 or a Server Core installation to a file share. You can then use the Printer Migration Wizard to migrate the printers from the file share to the server running Windows Server 2008 R2.

For more information about accessing the Printer Migration Wizard, see Preparing to Migrate.

Note

The **Printing-Server Core** role must be installed on a server running a Server Core installation from which you want to migrate.

Migrate print servers (overview)

The following list provides an overview of the steps to migrating the print servers.

- Access the migration tools
- Prepare the destination server
- Prepare the source server
- Back up the source server
- Restoration
- Verify the migration
- Post-migration

Impact of migration

The objective of the migration process is that the destination server is able to perform the same functions as the source server did, without client computers on the network being aware that the migration has taken place. The following sections describe the impact of migration.

Impact of migration on the source server

The source server is not impacted by print server migration until the destination server takes over as the active server (typically when the name or IP address of the source server is assigned to the destination server). At that point, the source server no longer services print requests that target the print server.

Impact of migration on other computers in the enterprise

If the destination server replaces the source server in the network (same name or IP address), then there should be no impact to other computers in the enterprise.

If the destination server has a different name or IP address than the source server, then all clients with existing print connections must delete and recreate those print connections so that they target the destination server.

Permissions required to complete migration

Administrative permissions are required on both the source print server and the destination print server.

Permissions required to complete migration on other computers in the enterprise

If the destination server replaces the source server in the network, then no permissions are required on other computers in the enterprise. If the destination server has a different name or IP address, then the permissions required on other computers may vary depending on Group Policy settings, Windows Update access, and driver availability.

Estimated duration

The time required to migrate a print server will vary from server to server, depending on the following:

- The number of queues being migrated.
- The number of individual drivers needed for the queue.
- The size of a given driver, in terms of its file size and the number of files.
- The configuration of the server.

Migrating a single printer queue with a typical x86-based and x64-based driver can range from five seconds to several minutes, depending on the factors listed above. Because of this range, a typical migration can take anywhere from less than an hour to several hours.

See Also

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Preparing to Migrate

Access the migration tools

The Printer Migration Wizard and the Printbrm.exe command-line tool support all migrations to Windows Server 2012.

Note

Although the Printer Migration Wizard supports migrations from servers running Windows Server 2003 or a Server Core installation, it cannot run on these servers directly.

To access the Printer Migration Wizard

Open the Print Management snap-in on computers running operating systems previous to Windows 8

1. If necessary, enable the **Administrative Tools** menu, which is hidden by default on Windows-based client operating systems.
   a. Right-click **Start**, and then click **Properties**. The **Start Menu and Taskbar Properties** dialog box opens.
   b. On the **Start Menu** tab, click **Customize**. The **Customize Start Menu** dialog box opens.
   c. Under **System Administrative Tools**, select **Display on the All Programs menu** or **Display on the All Programs menu and the Start menu**.
2. In the **Administrative Tools** menu, click **Print Management**.

Tip

Alternatively, you can click **Start**, and type **printmanagement.msc** in the search text box.

Open the Print Management snap-in on computers running Windows Server 2012

- Open Server Manager, click **Tools**, and then click **Print Management**.

Open the Print Management snap-in on computers running Windows 8

- From the Start screen, type **printmanagement.msc** and click **printmanagement** when it appears in the search results.

Tip

If you run the Print Management Console remotely (for example, from a client computer), simply add the print server to your console (right-click **Print Servers**, click **Add/Remove Servers**) and continue the migration from there.

Start the Printer Migration wizard

- In Print Management, right-click **Print Management** and click **Migrate Printers**.

Important
The Print Management snap-in filter settings will not be migrated and need to be saved independently of the printer migration.

To access the Printbrm.exe command-line tool

1. To open a Command Prompt window, click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.

   To open a command Prompt window on a computer running Windows 8 or Windows Server 2012, right-click the start charm and click Command Prompt (Admin)
2. Type:

   ```
   %WINDIR%\System32\Spool\Tools\Printbrm.exe
   ```

To view the complete syntax for this command, at a command prompt, type:

```
Printbrm.exe /?
```

For a listing of the available syntax for the **Printbrm.exe** command, see Appendix A - Printbrm.exe Command-Line Tool Details.

Prepare the destination server

The second step in the migration process is to prepare the destination server.

Hardware requirements for the destination server

There are no specific hardware requirements for being a print server beyond those for the version of the server operating system you are using.

The amount of disk space needed for a migration is dependent on the number of print drivers to be migrated and the size of the drivers. Because print drivers can vary widely in size, the amount of disk space can range from one megabyte to several gigabytes.

Software requirements for the destination server

Verify that hard drive space is sufficient on the destination server for the backup.

No additional software is needed other than the necessary drivers required for the printers to be hosted. Migrate these drivers from the source server.

For cross-architecture migrations, verify that the destination server contains drivers for each supported architecture.

Installing the Print and Document Services role on the destination server

You must install the Print and Document Services role on the destination server before you begin the migration process. For more information on installing this and other server roles, see the Print and Document Services overview.

Preparing for cross-architecture migrations

If you are migrating from the x86-based architecture of Windows Server 2003 or Windows Server 2008 to the x64-based architecture of Windows Server 2012, you should install x64-based drivers on the source server before creating the backup file. The migration process copies all installed drivers from the source server to the destination server. It recreates the printer queues on the destination server if the printer settings file contains the x64-based drivers.

Verify that each print queue on the source server has a driver installed for the operating system on the destination server before creating the printer settings file. For example, if you are migrating an x86-based source print server to an x64-based destination print server, verify that each print queue has an x64-based driver installed before you create the printer settings file. Any print queue that does not have a cross-architecture driver installed will not be migrated to the destination server.

To install cross-architecture drivers for a printer, you can use:

- The **Add Printer Driver Wizard**, which is available in the Print Management snap-in.
- The **Printer Properties** dialog box, which is available through the Printers folder in the Control Panel.

As a best practice, you need to install a driver with the same name as the native architecture. To add the x86-based driver to the x64-based destination server, use the x86-based client to remotely open the x64-based server using Windows Explorer and navigate to the remote printer folder and add the driver. To install an x64-based driver on the x86-based source server, use the x64-based client to remotely open the x86-based server using Windows Explorer and navigate to the remote printer folder and add the driver.

Tip

In many cases, it can take you a long time to update all the print drivers for all your print queues. To save time, you may want to update just the most used print queues first, and gradually update the others when you have time. To save time, you can set the existing print queues to the **Generic/Text Only** driver for migration and later switch them to the OEM driver at your convenience. Most printers allow basic printing using the Generic Text driver. For more information, see Cross-Architecture print server migrations: Speeding up the migration process at the Microsoft TechNet Blogs web site.

Preparing for additional scenarios

In the following instances, installing a feature on your destination server may require additional preparation before you migrate to it:

- The server hosts Line Printer Remote (LPR) printers.
- The server offers Internet Printing Protocol (IPP) printer connections.
- The server hosts Web Services on Devices (WSD) printers.
- The server is in a server cluster.
- The server hosts plug and play printers.

For more information on these scenarios, see Appendix B - Additional Destination Server Scenarios.

Prepare the source server

Simple system-to-system migrations require no preparation for the source server. However, additional preparation is required for cross-architecture migrations. If performing the migration as quickly as possible is a priority, remove unused drivers, ports, and queues before starting the migration to improve its speed after verifying with users that the items to remove are no longer in use. In general, however, minimize changes to the source server environment to ensure you can roll back to this environment if necessary.

Caution

If your source server is running multiple roles, renaming the source server or changing its IP address can cause other roles that are running on the source server to fail.

Notes

You should delete native print drivers that are not currently associated with a print queue because these drivers increase the size of the printer settings file unnecessarily. The print spooler will not allow a native print driver that is currently associated with a print queue to be deleted.

The Print Spooler service will use non-native drivers. It routes these drivers to the Print Server service when a non-native client connects to a print queue and has to download a driver. You should remove any unused drivers and print queues.

Do not delete a non-native driver with a corresponding native print driver that is associated with a print queue. In this instance, the Print Spooler service will not prevent the non-native driver from being deleted. If the non-native driver's architecture matches the destination server's architecture, then you must block the driver's deletion. Cross-architecture drivers will never appear to be loaded by the Print Spooler service. Administrators should only delete them after confirming the drivers are no longer needed.

To install cross-architecture drivers using the Print Management snap-in on computers running Windows Vista and Windows Server 2008

1. Open the Print Management snap-in. Click **Start**, click **Administrative Tools**, and then click **Print Management**.
2. In the Print Management tree, under **Print Servers**, click the print server you want.
3. Under the print server, right-click **Drivers** and then select **Add Driver** to open the **Add**

Printer Driver Wizard.

4. Follow the steps as indicated by the wizard.

To install cross-architecture drivers by using only the Printer Properties dialog box on computers running Windows XP and Windows Server 2003

1. Click **Start**, click **Control Panel**, and double-click **Printers**.
2. Select **Printer**. Right-click **Sharing**.
3. Click **Additional Drivers** and select **Processor** from the list.
4. Follow the instructions in the dialog boxes to install the correct driver. Only install the driver associated with the printer you are administering.

Note

You can only add a cross-architecture driver if you have already installed a native architecture version of the same driver.

See Also

Migrate Print and Document Services to Windows Server 2012

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Migrating the Print and Document Services Role

Back up the source server

The fourth step in the migration process is to back up your source server data to a printer settings file using either the Printer Migration Wizard or the **Printbrm.exe** command-line tool in preparation for exporting printer queues, print drivers, and printer settings.

Important

As a best practice, run the Printer Migration Wizard or Printbrm.exe from a computer running Windows Server 2012 or Windows 8 to ensure that you are using the newest version of the migration tools that have the latest updates and features. You can run

these tools either locally on the server or remotely from any other computer running Windows Server 2012 or Windows 8.

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

To back up the source server using the Printer Migration Wizard

1. Open the Print Management snap-in.
2. Do one of the following:
   - In the **Print Management** window, right-click **Print Management**, and then click **Migrate Printers** to open the Printer Migration Wizard. Make sure that **Export printer queues and printer drivers to a file** is selected, and then click **Next**. In the **Select a print server** window, select the print server to be migrated, and then click **Next**.
   - In the **Print Management** tree, under **Print Servers**, right-click the print server that contains the printer queues to migrate, and then click **Export printers to a file** to open the Printer Migration Wizard.
3. Review the list of items to be exported, and then click **Next**.
4. In the **Export printer data to** box, enter the path to the printer settings file to use, or browse to the location where you want to store the file. Click **Next** to export the printer-specific information for the server to this file.
5. Verify that the printer settings file is stored on a resource that will be available to the destination server. Optimally, store it on a network share. Click **Finish**.

The backup file that you create by using either the Printer Migration Wizard or the Printbrm.exe tool inherits the permissions allowed by your user credentials. Only you can access the file if you saved the file directly to a share during the backup file creation process. You must either change the file permissions on the **Security** tab of the file's **Properties** dialog box, or you must perform any restorations or migrations by using that file yourself. If you create the backup file on the computer from which you are running the migration and later copy the file to a share, then file access permissions are inherited from the destination folder.

To back up the source server using the Printbrm.exe command-line tool

1. Open an administrator Command Prompt window.
2. Perform a remote print backup. To do this, type the following command in the **%WINDIR%\System32\Spool\Tools** folder at the command prompt, in which *Source Computer1* is the Universal Naming Convention (UNC) name of the source computer, and *Printer1 Settings* is the name of the printer settings file to back up.

```
Printbrm -s \\<Source Computer1> -b -f <Printer1
Settings>.printerExport
```

Notes

The Printer Migration Wizard and the Printbrm.exe command-line tool only

support a printer settings file that is created by the migration tool you are using. For example, .cab file backups that were created by using the Printer Migration Wizard are not supported. To view the complete syntax for the **Printbrm.exe** command, type **Printbrm.exe /?** in a Command Prompt session.

Only TCP/IP, WSD, and LPR ports will be migrated. The Printer Migration Wizard will not migrate printers attached through USB, LPT, or other local ports. For more information about these scenarios and migrating Plug and Play printers, see Appendix B - Additional Destination Server Scenarios.

Cross-architecture migrations

For cross-architecture migrations, verify that each print queue has a driver installed on the source server that is compatible with the operating system on the destination server before creating the printer settings file on the source server. For example, if you are migrating an x86-based source print server to an x64-based destination print server, verify that each print queue has an x64-based driver installed before you create the printer settings file. Any print queue that does not have a cross-architecture driver installed will not be migrated to the destination server.

Restoration

The fifth step in the migration process is to restore the printers to the destination server, using the printer settings file you created.

Before beginning the migration process, verify that you installed the Print and Document Services role on the destination server as part of your preparation.

To restore printers to the destination server using the Printer Migration Wizard

1. On the source server, stop the Print Spooler service for all printers so you can preserve all print jobs prior to the migration.
   a. Open Computer Management. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then click **Computer Management**.
   b. In the console tree, expand **Services and Applications**.
   c. In the console tree, under **Services and Applications**, click **Services**.
   d. In the details pane, do one of the following to stop the service:
      i. Right-click **Print Spooler** and select **Stop**.
      ii. Double-click **Print Spooler**. On the **General** tab, under **Service Status**, click **Stop**.
2. From the computer that is running the Printer Migration Wizard, on the **Administrative Tools** menu, click **Print Management**.
3. Do one of the following:
   - Right-click **Print Management**, and then click **Migrate Printers** to open the Printer Migration Wizard. Select **Import printer queues and printer drivers from a file,** and then click **Next**.

- In the **Print Management** tree, under **Print Servers**, right-click the destination print server, and then click **Import printers from a file** to open the Printer Migration Wizard.

4. Specify the printer settings file created in the **Back up the source server** section, and then click **Next**.
5. Review the list of items to be imported, and then click **Next**.
6. In the **Import Mode** list, indicate whether you want to keep or overwrite existing printers.

   If the printer settings file contains a printer already on the destination server, the printer is not restored, and the existing printer on the destination server is not changed.
7. In the **List in the directory** list, indicate which printers to list on the destination server.
8. Optionally, indicate whether you want to convert LPR ports to standard port monitors when you migrate.
9. Click **Next** to import the printers.
10. Click **Finish**.

    Note

    It is recommended that you review the Application events that have a PrintBRM source to determine whether any additional actions are needed. The restored printers are shared in the same manner in which they were shared previously.
11. To view details of the migration, click **Open Event Viewer**. For more information, see "Verify the Migration" in Verifying the Migration. If you identify Error 30 in the Event Viewer, see "Troubleshooting" and "Migrating cross-platform driver language monitors" in Post-Migration Tasks for instructions on resolving the error.

To restore printers to the destination server using the Printbrm.exe command-line tool

1. Open an administrator Command Prompt window.
2. Type the following command in the **%WINDIR%\System32\Spool\Tools** folder at the command prompt, in which *Source Computer1* is the UNC name of the source computer, and *Printer1 Settings* is the name of the printer settings file to restore.

```
Printbrm -s \\<Source Computer1> -r -f <Printer1
Settings>.printerExport
```

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Verifying the Migration

Verify the migration

The sixth step in the migration process is to verify that the migration was successful by testing and validating the new print server. The Printer Migration Wizard provides detailed logging of migration events in the Event Viewer.

To verify destination server configuration

1. View event log messages about the migration.
   - If you are managing Print Services migration from a remote client computer, you can view event messages in **Custom Views/Administrative Events** in Event Viewer on the Windows-based client computer.
   - If you are managing migration from the destination server and Print and Document Services is not yet installed, then migration-related events are logged in **Custom Views/Administrative Events** in Event Viewer on the destination server.
   - If you are managing migration from the destination server and Print and Document Services is installed, events are logged to a different location. See the events located at **Applications and Services Logs/Microsoft/PrintBRM/Admin**

     To view details after closing the Printer Migration Wizard, right-click the **Start** charm, and then click **Event Viewer**. In the Event Viewer pane, under **Custom Views**, click **Server Roles**, and then click **Print and Document Services**. In the center pane, click the printer migration event to view details.
2. To verify that each printer queue was migrated to the new server:
   - Manually check the destination server for each printer migrated from the source server.
   - Verify that the printer associated with each printer queue is online.

   To check each printer queue's online status in the Print Management snap-in, under **Print Servers**, click **Printers**. A list of all migrated printers appears in the center pane, listing the printer queue status for each printer. Clients will be unable to print to printers that were not restored successfully, and any connections to a missing printer queue will be invalidated.
3. Check the printer queue settings.
   - Confirm that a printer queue's special settings, permissions, or drivers were preserved during the migration.
   - Check the properties for each queue on the destination server and verify that any special settings are still applicable.
   - If the driver installs any non-standard settings that have been altered as a result of the migration, verify those as well.

   Note

The migration process only preserves printer queue permissions. Other permissions on the source server, such as system permissions (for example, user accounts) and custom permissions, are not migrated using this process.

4. Make any necessary changes, such as adding a port monitor or a new driver.

Rename the destination server to the name of the source server

Temporarily rename the destination server. For example, you can name the destination server the same as the source server with **_NEW** appended to the source server name. After verifying that printers are restored to the destination server, rename the source server (for example, by appending **_OLD** to the source server name), and then use the source server's pre-migration name as the new name of the destination server.

Important

Validating existing printer connections from client computers can only be completed after the destination server name is the same as the pre-migration source server name. Print connections to the server in the period between renaming the source server and renaming the destination server will fail. All migration steps should be complete on the destination server before renaming to ensure that the downtime occurs only between the renaming of the source server and the final renaming of the destination server.

If you are using Active Directory Domain Services (AD DS) to publish printers, do the following to ensure that AD DS does not contain multiple instances of the same printer.

When you restore printers to the destination server, do not publish printers to AD DS. This prevents duplicate printers from being displayed by AD DS before the destination server configuration is verified.

On the source server, you must unpublish printers before renaming the source server. To do this, select all printers in the Print Management snap-in, right-click the selected printers, and then click **Remove from Directory**. This prevents printers from being published twice to AD DS when the source server is renamed.

After renaming the destination server to the source server's original name, you can publish all printers on the destination server to AD DS. To do this, select all printers in the Print Management snap-in, right-click the selected printers, and then click **List in directory**.

To verify configuration of other computers in the enterprise

In most cases, a new print server will not affect other computers in the enterprise. Existing client connections may be corrupted if you make a change to any of the following print server properties:

- The print server name
- The printer name
- The print share name
- The share permissions
- The printer's availability to the server

Print a test page to each printer queue from a client (or set of clients) that had an existing connection to the source server to verify that other computers are not affected by the new print server. In a cross-architecture environment, test each supported architecture.

Print a test job from a client with an existing connection

From a client computer that is configured to print to the source print server, use the existing print queue to print a page to the new server. If you cannot print a test page:

- Determine if one or more of the print server properties listed above have been changed.
- Check whether the destination server is available to the client on the network. Create a new connection to a printer on the destination server to verify that the client and server are communicating normally.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Post-Migration Tasks

Post-migration

The final step in the migration process is determined by whether the migration was successful or unsuccessful.

No post-migration tasks are necessary beyond the standard migration process. If the server state (new settings, drivers, queues, and so on) changed, create and archive a new backup of the server state for recovery purposes.

Success

The next section provides tasks that should be completed after you have successfully migrated the source printer server.

Retire the source server

After taking the source server offline while backing it up, users are unable to print until the migration to the destination server is complete. To minimize the impact, leave the source server

in service while you complete the migration and testing of the destination server. By leaving it in service, you can also add both the source and destination servers to the Print Management snap-in to simplify verifying the restoration.

Once you have validated the installation, rename the source and destination servers and take the source server offline.

1. On the source server, restart the print spooler for all printers so it can finish spooling delayed print jobs. When it finishes, verify that there are no new print jobs.
2. Rename source and destination servers as directed in Verifying the Migration.
3. Follow your enterprise's normal policy for server decommissioning and retirement until retirement of the source server is complete.

Note

If the destination server has already been published in Active Directory Domain Services (AD DS), then the source server must be unpublished.

Failure

The next section provides tasks that should be completed if your migration of the source printer server did not succeede.

Restoring the role in the event of migration failure

Restoring the source server lets you deploy the settings to a new system or use the source server while determining the cause of the failure.

Rollback requirements

Warning

Rollback can only be completed if retirement of the source server has not been started. After you start retiring a source server—that is, you delete any print queues, close any print connections, reformat any drivers, or remove any hardware from the source server—you cannot roll back migration. After you start retiring the source server, the only method of rolling back migration is to restart the Print Services migration process from the beginning.

Estimated time to complete rollback

Rolling back migration involves renaming the source server to its pre-migration name, and renaming the destination server to either its original name, or another name that is not the same as the pre-migration name of the source server. Renaming the source and destination servers can be completed in a few minutes.

Roll back migration on the source server

Rename the source server to its original, pre-migration name. You might have to rename the destination server to a temporary name first.

Roll back migration on the destination server

Rename the destination server, either to its original name, or to another name that is not the same as the original name of the source server.

Troubleshooting

The following sections can help you troubleshoot any migration issues.

Log file locations

Printer migration events are included in the Application log, which is located at %*SystemRoot*%\System32\Winevt\Logs\Application.evtx and can be viewed using Event Viewer. A custom view for Printer Migration Events is available in Event Viewer.

Note

If the Printer Migration Wizard fails, you are directed to Event Viewer to view error messages. If you cannot find an error that explains the failure in Event Viewer, restore the backup by using the Printbrm.exe command-line tool. Error reporting from Printbrm.exe can often provide more detail than what is available in the event log.

Migrating cross-platform driver language monitors

When a cross-architecture migration includes the migration of printer language monitors, an error will occur during the process of restoring the printers to the destination server using the Backup Restore Migration (Printbrm.exe) command-line tool. The reason for the error is that language monitor driver architecture must always be the same as the source server architecture. Therefore, when migrating from x86-based architecture to x64-based platforms, language monitor migration cannot be successful. An error posted to the event log will state that the source architecture is not the same as that of the destination server.

You can recover from the printer restore error on the destination server by manually installing (or reinstalling) the appropriate standard driver for the migrated printer(s) running on that architecture.

Mitigating a failure in the Print Spooler service

If you encounter a failure in the Print Spooler service during print server migration, you can work around the failure. Using policy settings, you can isolate print drivers in separate processes so that print driver failures will not cause the Print Spooler service to fail—which allows the restoration to continue.

To turn on print driver isolation using Group Policy

1. Open the Group Policy Management Console. Right-click a Group Policy Object with the necessary scope, and then click **Edit**.
2. In the console tree under **Computer Configuration**, expand the **Administrative Templates** folder, and then expand the **Printers** folder.
3. Double-click **Execute print drivers in isolated processes**.
4. Click **Enabled**, and then click **OK**.
5. Double-click **Override print driver execution compatibility setting reported by print driver**.
6. Click **Enabled**, and then click **OK**.
7. At a command prompt, type **gpupdate /force** to reapply Group Policy settings.

Additional references

- Install, Deploy, and Migrate to Windows Server 2012
- Windows PowerShell Blog (http://go.microsoft.com/fwlink/?LinkId=128557)

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Appendix A - Printbrm.exe Command-Line Tool Details

Printbrm.exe command-line tool syntax

The following table lists the available printbrm parameters:

Syntax	Description
-s <server name>	Specifies the destination server.

Syntax	Description
-b	Backs up the server to the specified file.
-r	Restores the configuration in the file to the server.
-q	Queries the server or the backup file.
-f <file name>	Specifies the backup file.
-d <directory name>	Unpacks the backup file to the directory (with -r), or repacks a backup file from the directory (with -b).
-o force	Forces overwriting of existing objects.
-p all:org	Publishes all printers in the directory, or publishes printers that were published originally.
-nobin	Omits binary files from the backup.
-lpr2tcp	Converts LPR ports to standard TCP/IP ports on restore.
-c <file name>	Uses the specified configuration file.
-noacl	Removes ACLs from print queues on restore.
-?	Displays Help.

Printbrm enhancements

Printbrm.exe in Windows Server 2012 has several enhancements and improvements, including the following:

- **Supports both v3 and v4 print drivers**

  Windows Server 2012 supports both driver types, so with Printbrm you have full flexibility to back up, restore and configure the drivers you need to support your environment.

- **Supports backup CAB files greater than 2 GB**

  You should use the latest PrintBRM version when performing a migration or restoration. You will avoid CAB file size issues if you use the latest version of Printbrm to migrate and restore your Windows Server 2008 R2 (or previous) servers.

- **General improvements for reporting and error handling conditions during the backup and restore processes**

  These conditions are primarily logged in the Event Log under **Custom Views\Printer Migration Events**.

Printbrm usage scenarios

There are many ways Printprm can be used to make migrating your printers easier and more flexible.

Using the configuration file

You can use a configuration file to customize your printbrm migration for the following purposes:

- Replace printer drivers during a restore operation.

  For example, you might want to import your printers to a new print server, but use new v4 printer drivers.
- Backup / Restore dependent files from third-party Language Monitors
- Backup / Restore dependent files from third-party party Print Processors

For example, to replace your printer drivers, you can backup your printers using the `-nobin` parameter, and then restore the printers using the `-c <file name>` parameter to specify a configuration file with a DriverMap section.

To use a configuration file to specify updated print drivers

1. Backup you printers using the `-nobin` parameter. For example:

   ```
   Printbrm.exe -b -nobin -s \\myoldprintserver -f printers.printerExport
   ```
2. On the new print server computer, manually install the updated printer drivers.
3. Create a BrmConfig.xml configuration file to map the old drivers to the new drivers. For example:

   ```
   <BrmConfig>
   <PLUGINS>
   </PLUGINS>

   <LanguageMonitors>
   </LanguageMonitors>

   <DriverMap>
   <DRV old="OldDriverName1" new="NewDriverName1"/>
   <DRV old="OldDrverName2" new="NewDriverName2"/>
   </DriverMap>

   </BrmConfig>
   ```
4. Restore the printers specifying your configuration file using your configuration file. For

example: `PrintBrm.exe -r -c BrmConfig.xml -f printers.printerExport -o force`

5. Check your installed printers to verify they are installed with the updated printer drivers.

Important

Remember that if you are running printbrm over the network to remote servers, the **Print$** share must exist on both the source and target servers and the Remote Registry Service must be running.

Selectively restoring your printers

After you export the printers from your source server, you can selectively restore the printers and their related objects using the `-d` parameter. You can follow a general procedure to accomplish this:

To selectively restore your printers

1. Export the printer objects from the source server.
2. Restore to a temporary folder using the `-d` parameter.
3. Manually edit the files in the temporary folder.

   Note

   More information about the files created by a printbrm backup is described later.
4. Backup the temporary folder using the `-d` parameter.
5. Import the modified backup file to the target server.

A backup or export operation using PrintBRM produces a compressed file that is used for the restore or import operation. The following XML files are part of the export file in addition to the individual printer driver and configuration files:

- **BrmDrivers** – contains a list of every driver installed on the computer and the driver files for each driver
- **BrmForms** – contains a list of all the installed forms
- **BrmLMons** – will usually contain either Windows NT x86 or Windows x64 as the architecture and a list of port monitors and the port monitor files installed
- **BrmPorts** – contains a list of all the installed printer ports
- **BrmPrinters** – contains a list of all printers that have been installed
- **BrmSpoolerAttrib** – contains information about the spooler directory path and also indicates whether or not the source server was a cluster server

Moving printers to a different domain

If you move printers to a different domain, you will want to prevent the restoration of the print queue's ACLs. Use the `-NoACL` parameter to do this. If you use this parameter, the restored print queues will inherit the permissions of the target print server.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix B - Additional Destination Server Scenarios

Appendix C - Printbrm Event IDs

Appendix B - Additional Destination Server Scenarios

In some instances, your destination server may require additional preparation before you migrate to it.

If your server hosts Line Printer Remote (LPR) printers

To enable the hosting of LPR printers, install the LPR Port Monitor feature on the server:

1. Open Server Manager.
2. In the **Server Manager** dashboard, click **Add roles and features**. The **Add Roles and Features Wizard** opens.
3. Click **Next** on the **Before you begin** page.
4. Ensure **Role-based or feature-based installation** is select on the **Select installation type** page, and click **Next**.
5. Ensure your destination server is selected on the **Select destination server** page and click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select Features** page, click **LPR Port Monitor**, click **Next**, and then follow the instructions to complete the installation.

Important

The LPD and LPR Services are deprecated starting with Windows Server 2012. Eventually, they will be completely removed from the product, but they are still available in this release. You should begin planning now to employ alternate methods for any applications, code, or usage that depend on these features. For more information about features or functionalities that have either been removed from the product in the current

release or are planned for potential removal in subsequent releases, see Features Removed or Deprecated in Windows Server 2012.

If your server offers Internet Printing Protocol (IPP) printer connections

To enable the Internet Printing Protocol (IPP):

- When installing the Print and Document Services role, select **Internet Printing**.

This automatically configures IIS and any other necessary features to support IPP printer hosting.

If your server hosts Web Services on Devices (WSD) printers

To enable WSD printing support:

1. Start the **Network and Sharing Center** from Control Panel and, click **Change advanced sharing settings** and click **Turn on network discovery**.
2. In Computer Management, Services, start the **Function Discovery Provider Host** service.
3. In Computer Management, Services, ensure the **Device Association Service** is started.

The server will be then able to identify and communicate with WSD-enabled printers.

If your print server is a highly available virtual machine

- To create a highly available print environment using Hyper-V and Failover Clustering, see High Availability Printing Overview. For more information about Windows Server 2012 Failover Clustering, see What's New in Failover Clustering in Windows Server 2012 and Clustering and High-Availability.
- Continue with the restoration process on the Printer Server virtual machine on the primary node.

If your server hosts local bus printers (LPT and USB)

The migration of local bus printers (LPT and USB) is not supported, although these printers are shown during backup. After the migration is complete:

1. Share the local bus printers again on the destination server.
2. Verify that each printer's name has not changed.
3. Test the printers to ensure that the shared connections still work.

If your server hosts plug and play printers

The migration of plug-and-play printers is not supported. To migrate plug and play printers:

1. Plug the printer into the destination server. The plug-and-play printer drivers will be installed automatically.
2. Enable printer sharing for the print queues.

See Also

Migrate Print and Document Services to Windows Server 2012

Preparing to Migrate

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix C - Printbrm Event IDs

Appendix C - Printbrm Event IDs

Printbrm Event IDs

The following Printbrm events are logged in the **Applications and Services Logs/Microsoft/PrintBRM/Admin** event log:

Event ID	Description
1	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is beginning a backup of print queues. No user action is required.
2	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is beginning a restore of print queues. No user action is required.
3	Printbrm.exe (the Printer Migration Wizard or the command-line tool) replaced driver map settings %1 with %2 for queue %3. No user action is required.
4	Printer queue %1 will be restored without a separator page. Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create a separator file for this queue. Error: %2.
5	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find a configuration file and is using the default settings. No user action is required.

Event ID	Description
6	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored print queue %1. No user action is required.
7	Printbrm.exe (the Printer Migration Wizard or the command-line tool) restored queue %1, but failed to restore printer driver settings. Error %2. This can occur if the driver on the destination server is newer than the driver in the migration file. Open the printer Properties dialog box on the destination computer and manually specify the appropriate printer settings.
8	While attempting to publish the printer to the Active Directory directory service, Printbrm.exe (the Printer Migration Wizard or the command-line tool) may have failed to adjust publishing settings for %1. This can occur if Printbrm.exe cannot access Active Directory. Manually publish the printer using the printer Properties dialog box.
9	Printer queue %1 already exists on the destination computer. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will update the printer settings to match the settings on the source computer. No user action is required.
10	Printer queue %1 already exists on the destination computer and will not be changed because Printbrm.exe (the Printer Migration Wizard or the command-line tool) was run in 'Keep Existing' mode. No user action is required.
11	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print queue %1 because port %2 is unknown. Printbrm.exe will attempt to restore the print queue on port FILE: instead. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are

Event ID	Description
	incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer and then change the print queue to use the new port.
12	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed printer queue %1 on port FILE:. No user action is required.
13	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore printer queue %1 on port FILE:. Error %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
14	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored %1. No user action is required.
15	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print processor %1 while restoring print queues from a file. Error: %2. Examine the Windows error returned by this event to determine the cause.
16	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not back up a dependent file for a language monitor. Dependent file: %1. Error: %2. This can occur if the file was deleted or moved.
17	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored language monitor %1. No user action is required.
18	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore language monitor %1 while restoring print queues from a file. Check the print queue

Event ID	Description
	backup file and examine the Windows error returned by this event to determine the cause.
19	The language monitors in the backup file are for a different processor architecture than the destination computer. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will not migrate any language monitors. Source architecture: %1. Destination architecture: %2.
20	Printbrm.exe (the Printer Migration Wizard or the command-line tool) restored a driver for a different processor architecture than that of the destination computer. Printbrm.exe will attempt to locate and install a native version of driver %1 on destination %2. Try to print to the print queue, and if necessary, manually install a native version of the driver.
21	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 for the processor architecture of the destination computer. No user action is required.
22	The driver in the backup file is for a different processor architecture than the destination computer, and Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not locate and install a native version of the driver. Driver: %1. Destination architecture: %2. Error: %3. Install a native version of the driver on the destination computer and then retry importing the print queues.
23	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from a cabinet (.cab) file. No user action is required.
24	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) because it was backed up on the source computer without its binary files. Printbrm.exe could not install the driver from the

Event ID	Description
	local driver cabinet (.cab) file on the destination computer. Error: %3. This can occur if the user did not save the driver binary files while backing up the print queue, or when restoring a print queue on a destination computer that uses a different processor architecture than the source computer. Install the driver manually on the destination computer or back up the print queue with its binary files. Then, retry importing the print queues.
25	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from files. No user action is required.
26	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) from files. Error reported: %3. This can occur if the driver requires a file that Printbrm.exe did not back up or if the user does not have permission to install drivers on the destination computer.
27	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to remove driver temporary folder %1. Error %2. Manually delete the temporary files and folder.
28	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to copy %1 to %2. Error %3 occurred while it was restoring print queues from a file. This can occur if the user does not have proper permissions on the destination computer, if the backup file is corrupted, or if the system is unstable. Retry exporting the printer queues on the source computer.
29	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not convert Line Printer Remote (LPR) port %1 to a standard TCP/IP printer port. Cannot get device settings. Error %2. This can occur if the LPR port was

Event ID	Description
	incorrectly configured. Manually create a standard TCP/IP printer port for the printer.
30	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) printer port %1 because the length of the port name is too long. len >= MAX_PORTNAME_LEN: Error %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
31	Port %1 already exists. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will skip restoring this port. No user action is required.
32	Printbrm.exe (the Printer Migration Wizard or the command-line tool) port install status: %1. No user action is required.
33	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is backing up printer forms. No user action is required.
34	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is restoring printer forms. No user action is required.
35	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully saved %1 user forms. No user action is required.
36	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find any user forms to restore. No user action is required.
37	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore one or more printer forms from the backup file. This can occur if the forms already exist on the destination computer or if the user does not have permissions to create forms on the destination computer.

Event ID	Description
38	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not load .xml file %1 while restoring print queues from a file. Error: %2. This can occur if the user does not have proper permissions on the destination computer, if the backup file is corrupted, or if the system is unstable. Retry exporting the printer queues on the source computer.
39	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not find the printer driver migration plug-in DLL %1 for print processor %2. Printbrm.exe will skip this file and attempt to migrate the driver, but the printer might not work. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
40	Printbrm.exe (the Printer Migration Wizard or the command-line tool) changed the spooler folder to: %1. No user action is required.
41	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change spooler folder to %1. This path does not exist, so Printbrm.exe used the default spooler folder location (%WINDIR%\System32\Spool\Printers). This can occur if the configuration of the source computer is different from the destination computer. No user action is required unless you want to use a custom spooler folder location.
42	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change the location of the spooler folder. Error %1. Printbrm.exe used the default spooler folder location (%WINDIR%\System32\Spool\Printers). This

Event ID	Description
	can occur if the configuration of the source computer is different from the destination computer. No user action is required unless you want to use a custom spooler folder location.
43	Printbrm.exe (the Printer Migration Wizard or the command-line tool) changed the spooler log level to: %1. No user action is required.
44	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully restored the print queue, but could not change the spooler log level. Error reported: %1. No user action is required unless you want to change the default spooler log level.
45	Printbrm.exe (the Printer Migration Wizard or the command-line tool) is backing up printer objects on server: %1. No user action is required.
46	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to place a file in cabinet (CAB) file %1 while backing up print queues. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
47	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create destination file %1 while backing up print queues. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
48	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not load printer driver migration plug-in DLL %1. Error: %2. Printbrm.exe will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer

Event ID	Description
	and install a newer version of the driver, if necessary.
49	The printer driver migration plug-in %1 is incompatible with this version of Windows. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
50	The printer driver migration plug-in %1 ran with error: %2. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will attempt to migrate the driver, but the printer might not work because of missing files. This can occur due to a problem with the migration plug-in provided by the printer driver. Try to print to the affected printer and install a newer version of the driver, if necessary.
51	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully ran printer driver migration plug-in %1. No user action is required.
52	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore print processor %1 because it is already installed on the destination computer. No user action is required.
53	While Printbrm.exe (the Printer Migration Wizard or the command-line tool) was restoring a print queue from backup, the following service control function %1 on computer %2 returned error %3. This can occur if Printbrm.exe cannot locate all files for a print processor, cannot stop the print spooler or the Cluster service, or if either of these services become unresponsive during installation of the print processor. If

Event ID	Description
	printing fails, manually install the print processor on the destination computer and examine the Print Spooler service and/or the Cluster service.
54	The service %1 is now running on server %2. No user action is required.
55	The service %1 is now stopped on server %2. No user action is required.
56	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to stop service %1 on computer %2. Error %3. Try printing to the affected printer to determine if there is a problem.
57	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to start service %1 on computer %2. Error %3. Try printing to the affected printer to determine if there is a problem.
58	Printbrm.exe (the Printer Migration Wizard or the command-line tool) did not overwrite print processor file %1 because the file on the destination system has an identical (or newer) time stamp. No user action is required.
59	Printbrm.exe (the Printer Migration Wizard or the command-line tool) copied print processor file %1 to %2. No user action is required.
60	While restoring a print queue from backup, Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to copy print processor file %1 to %2. Error %3. This can occur if Printbrm.exe cannot locate all files for a print queue, or if the Print Spooler service cannot be stopped during installation of the print processor. If printing fails, manually install the print processor on the destination computer.
61	The device %1 %2 supports more than one

Event ID	Description
	port. Printbrm.exe (the Printer Migration Wizard or the command-line tool) will convert the first port on the device from a Line Printer Remote (LPR) port to a standard TCP/IP port. The port on the device might need to be reconfigured.
62	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) port %1 because the LPR port monitor is not installed on the destination computer. Install the LPR port monitor on the destination computer.
63	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the Line Printer Remote (LPR) port %1 while restoring print queues from a previously created backup file. Error code: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
64	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the standard TCP/IP printer port %1. Error: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.
65	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore printer driver settings (%1, %2). Error %3. This can occur if the driver on the destination server is newer than the driver in the migration file. Open the printer Properties dialog box on the destination computer and manually specify the appropriate printer settings.
66	Printbrm.exe (the Printer Migration Wizard or

Event ID	Description
	the command-line tool) cannot apply new settings for printer queue %1 because port %2 that is saved in the cabinet (.cab) file does not exist on the destination server. This can occur if Printbrm.exe failed to migrate the port. Manually install the appropriate port, and then specify the appropriate printer settings or retry importing the print queue.
67	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not open registry key %1. Error code %2. This can occur if an important Windows resource (such as the registry) is unavailable, if the Component Object Model (COM) cannot be initialized, or if Printbrm.exe cannot allocate memory. Examine the Windows error returned by this event to determine the cause.
68	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not create a cabinet (CAB) file at this location: %1. Error reported: %2. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources.
69	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the standard TCP/IP printer port %2 because the length of the host IP address %1 is too long to be restored on operating systems older than Windows Vista.
70	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to save separator file %1 because the file or path does not exist. This can occur when a separator page cannot be found during the backup process on the source computer.
71	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not access the print$ share on %1. Error reported: %2. This

Event ID	Description
	can occur if the destination computer does not have any shared printers.
72	Printbrm.exe (the Printer Migration Wizard or the command-line tool) was unable to access the Remote Registry service on %1. Error reported: %2. This can occur if the Remote Registry service is not started or if the computer is behind a firewall.
73	Printbrm.exe (the Printer Migration Wizard or the command-line tool) reset the availability information for the print queue because the StartTime or UntilTime values were invalid. Error %1. The printer will be always available until you use the Advanced tab of the printer's Properties dialog box to specify the correct availability.
74	Printbrm.exe (the Printer Migration Wizard or the command-line tool) successfully installed driver %1 (%2) from a driver package. No user action is required.
75	Printbrm.exe (the Printer Migration Wizard or the command-line tool) could not restore driver %1 (%2) from a driver package. Error: %3.
76	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to backup driver %1 (%2). The backup process will continue, skipping this driver. Error: %3.
77	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to backup port %1. The backup process will continue, skipping this port. Error: %2.
78	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print processor %1 (%2). Error: %3.
79	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print processor files for architecture %1. Error: %2.

Event ID	Description
80	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to back up print queue %1. The backup process will continue, skipping this queue. Error: %2.
81	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore print queue %1. The restore process will continue, skipping this queue. Error: %2.
82	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to create a new separator file folder while restoring print queue %1. This can occur if the user does not have permission to create a file in the destination location, or if there is insufficient disk space or system resources. The print queue was not restored successfully. Error: %2.
83	Printbrm.exe (the Printer Migration Wizard or the command-line tool) failed to restore the WSD printer port described by the remote URL or global identifier %1. Error: %2. This can occur if the backup file contains incomplete data about the port, or if the port or port settings are incompatible with the version of Windows installed on the destination computer. Recreate the affected port on the destination computer.

See Also

Migrate Print and Document Services to Windows Server 2012

Migrating the Print and Document Services Role

Verifying the Migration

Post-Migration Tasks

Appendix A - Printbrm.exe Command-Line Tool Details

Appendix B - Additional Destination Server Scenarios

Print Server Scalability and Capacity Planning

This guide describes the Microsoft best practices approach for estimating scalability and capacity for Print Server in Windows Server 2012 R2 or Windows Server 2012. It provides an in-depth description of the contributing factors that affect the scalability and performance of Print Server, such as the hardware specifications, network latency, and the content of the print job, and it includes best practices and recommendations for deploying this solution correctly.

This guide contains the following information:

- Analysis of the key factors affecting performance, scalability, and capacity planning
- Case study of a Print Server reference system
- Print Server deployment best practices and recommendations

This guide is intended for readers including administrators and architects who are responsible for creating the architecture and design of Print Server within the enterprise. It assumes that the reader has knowledge of Print Server, Windows Server 2012 R2 and Windows Server 2012, and related Microsoft products and technologies.

What this guide includes

- Factors that influence Print Server performance
- Diagnostic and performance analysis
- Reference systems case study
- Recommendations for deployment
- Print Server performance data

Factors that influence Print Server performance

The following contributing factors are presented in order from those that have the greatest, to those that have the least, effect on the scalability and performance of Print Server.

Number of clients

The number of clients that are connected to a print server can causes a large drain on system resources. System resources are allocated and ready for use immediately when a connection is established between the client and the server. The connection to a Print Server is established each time a client invokes spooler activity, such as when querying for status or submitting a print job.

Depending on the application that the client uses to establish the connection to the print server, and then to print, the client may inadvertently hold connections open, for example, when the user does not terminate the application on their workstation. In this case, the connection might remain open for hours, or even days in extreme cases. In fact, a print server can become constrained

either due to the number of clients that are connected actively when printing, or passively when the client holds the connection open.

Administrators generally find that the total number of users connected to a print server shows light activity at beginning of the workday, an increasing server load and peak during midday, and a gradual decrease at the end of the day.

Client-side rendering or server-side rendering

Print Server in Windows Server 2012 R2 or Windows Server 2012 can be configured to render documents to the device-ready page description language (PDL) format either on the client or on the server. In client-side rendering, the print system on the client translates the document to a printer-specific PDL before it sends the data stream to the print server.

When the queue is configured to use server-side rendering, or the client is not capable of using client-side rendering, the client print system sends the document directly to Print Server. Print Server then renders the document to a printer-specific PDL locally.

To maximize server capacity, we recommend using client-side rendering when possible.

Operating system of the print clients

Depending on the enterprise ecosystem and the mixture of supported clients, the performance of Print Server might vary due to its feature capabilities. For example, in an environment that uses Windows 7 clients and v4 printer drivers on the printing shared resource, server-side rendering is used. Effectively, every print job that is generated by these clients requires processing on the server. Therefore, when there are more simultaneous jobs, the workload on the server increases.

As another example, if Branch Office Direct Printing is turned on in Windows Server 2012, and there is a heavy mixture of Windows 8 clients, performance is improved. Windows 8 performs client-side rendering that offloads the processing to the client instead of to the server. Then it routes the job directly to printer, instead of passing it through a centralized print server, which reduces network traffic.

For more information, see Branch Office Direct Printing Technical Details.

The default rendering options that are used by Windows Server 2012 R2 and Windows Server 2012 are described in the following table.

	V3 printer driver	V4 printer driver	V4 printer driver (Branch Office)
Windows 8.1	Client-side	Client-side	Client-side
Windows 8	Client-side	Client-side	Client-side
Windows 7	Client-side	Server-side	Client-side

Because of the various advantages of the v4 printer driver, administrators are encouraged to promote a client mixture that uses Windows 8.1 or Windows 8 with Windows Server 2012 R2 or Windows Server 2012. In situations where the enterprise has a large number of Windows 7

users, the effect of server-side rendering could require additional benchmarking to ensure acceptable performance.

For client devices that run on ARM processor architecture, we encourage the use of v4 printer drivers to take advantage of the various benefits of the Print and Document Services architecture.

For more information, see Print and Document Services Architecture.

Mix of drivers hosted on print server (v4 or v3 printer drivers)

In Windows 8 and Windows Server 2012, Microsoft introduced a new iteration of the printer driver model, called the v4 printer driver. This driver model addresses gaps in the previous version and it offers a wide range of advantages. Because the v4 printer driver model has fewer components that cause heavy server processing, there is a lower server load when a v4 printer driver is used with Windows 8.1 and Windows 8 clients. Because Windows 7 clients always use server-side rendering with the v4 printer drivers, more processing occurs on the server.

For more information, see V4 Printer Driver.

Types of print jobs

When server-side rendering is used, the server handles the conversion of the job to device-ready PDL. The amount of processing for each print job is hard to predict because it depends on the content of each print job. For complex jobs, such as documents with extensive graphics, documents in PDF format, or documents with multiple types of fonts, more CPU processing is required, as compared to a simple text document.

If client-side rendering is used, the processing occurs on the client-side, and it does not present additional workload to the server. Given the computing capabilities of devices in the market, we recommend using the client-side rendering configuration when possible.

Printer hardware

Printers that support the Open XML Paper Specification (also called XPS or OpenXPS) require only minimal server support, even when the server is configured to use server-side rendering. Printers that support PCL6 or PostScript are less efficient than XPS because they are vector formats. They require fewer resources than a printer that uses a low-end raster-based image editor.

High-end printers generally have built-in hard disk drives and memory. This allows them to process a print job without a protracted spooling process and the associated disk utilization. In contrast, less expensive printers generally have less processing power and internal memory, and they require more processing and disk space on the print server. It is the administrator's responsibility to gather organizational requirements and analyze the tradeoff between device capabilities and cost.

Printing related factors

The network bandwidth and storage for Print Server can be affected by printing related factors including the number, size, and frequency of print jobs. For example, using Print Server with high-frequency smaller jobs has a different storage requirement than using Print Server with large print jobs. If Print Server is configured to save print jobs after printing, additional storage is required and this requirement needs to be considered for the deployment.

System hardware configuration

As with any computing environment, more cores on a processor and more memory enables more computing output. Depending on enterprise requirement and the budget that is available, this is an important factor that administrators should consider when making decisions about deployment.

With the trend toward server consolidation using virtualization, it is common to see print servers running in a virtualization environment. Although there is no functional difference between a print server on a virtual machine and a physical print server, there can be a significant difference in the scalability factor based on the hardware requirements and specifications.

Performance needs to be evaluated based on the allocated number of cores and memory. On physical print servers, the print job offloads processing to the GPU to expedite completion when it uses a v4 printer driver and supports the XPS rasterization service. On virtual machines, this behavior does not occur due to the absence of a GPU.

To verify if a particular v4 printer driver supports the XPS rasterization service, administrators can check the manifest file for the driver package. Specifically, you need to check the filter pipeline configuration XML file.

For more information, see Filter Pipeline Configuration File.

In addition, administrators need to verify that the FilterServiceProvider element contains XpsRasterService.dll. You might also review the details on the XPS rasterization service usage tree.

For more information, see:

- Using the XPS Rasterization Service
- XPSRas GPU Usage Decision Tree

Administrators can also add failover clustering to their print service by using the standard clustering strategy. For more information, see:

- Clustering and High Availability, Creating a Windows Server 2012 Failover Cluster
- Failover Clustering Hardware Requirements and Storage Options

Storage and spool directory

A print server typically performs many I/O operations, while accepting and processing print jobs. Depending on the state of Print Server, excessive page fault can occur due to memory paging during high print activities. The combination of concurrent I/O operations and page fault might cause Print Server performance to degrade.

To minimize this issue, we recommend that administrators store the print spool folder on a drive that is separate from the drive that hosts the operating system and page files. Having this separate I/O channel reduces server activity and workload.

For more information, see HOW TO: Move the Print Spool Folder in Windows 2000.

Non-Microsoft software

Some non-Microsoft software can affect server scalability. For example, if Print Server uses a non-Microsoft accounting package that runs as a system service and tracks print usage information, it would likely consume CPU. Non-Microsoft port monitors may provide similar

effects. Depending on their implementation and interaction with the system, this type of software could create additional workload that reduces performance.

Diagnostic and performance analysis

Because there are many factors that affect the sizing of a print server, the best way to effectively estimate server size is to monitor performance for an existing print server over a period of time in order to identify bottlenecks, and to understand the investment. Because the server load varies over the course of a day or week, we recommend monitoring the print server for at least a day to identify peak resource utilization. This section describes some of the performance counters that are useful for print server analysis.

To review these performance counters, Performance Monitor is provided in Windows Server 2012 R2 and Windows Server 2012. For a complete list of all performance counters, see Using Performance Monitor.

Print-specific counters

The print-specific counters for monitoring performance are described in the following table.

Counter name	Description
Add Network Printer Calls	Total number of calls from other print servers to add shared network printers to this server since the last restart.
Bytes Printed/Sec	Number of bytes printed per second in the print queue. This provides a rough indication of the extent of time that the printer is busy. This counter can be used for bottleneck detection.
Enumerate Network Printer Calls	Total number of calls received from clients to this print server to request network browse lists since the last restart.
Job Errors	Total number of job errors in a print queue since the last restart. Job errors can occur if the connection to the printer has errors due to network issues.
Jobs	Current number of jobs in a print queue. Use this counter to identify excessive use.
Jobs Spooling	Current number of spooling jobs (incoming or incomplete) in a print queue.
Max Jobs Spooling	Maximum number of spooling jobs (incoming or incomplete) in a print queue since the last

Counter name	Description
	restart.
Max References	Peak number of references (open handles) to this printer.
Not Ready Errors	Total number of printer-not-ready errors in a print queue since the last restart.
Out of Paper Errors	Total number of out-of-paper errors in a print queue since the last restart.
References	Current number of references to a print queue. A reference can be a user or a program that is connecting to a printer and opening a print queue.
Total Jobs Printed	Total number of jobs printed from a print queue since the last restart. Note that the counter reports the total number of jobs spooled. In the event a job has not successfully despooled, it is still accounted for by using this counter.
Total Pages Printed	Total number of pages that have printed through the Graphics Device Interface (GDI) in a print queue since the last restart. **Note** Print jobs using the RAW data type are not included in this counter because they do not provide page-count data. In addition, the print volume of a print job using the RAW data type is not captured by this counter. The RAW data type was used by earlier Windows operating systems clients to submit print jobs.

Process-specific counters

The process-specific counters for monitoring performance are described in the following table. In addition, to monitor the performance of a Print Server operation, check the following processes, since they provide the most useful information regarding print activities:

1. Spoolsv.exe
2. Printfilterpipelinesvc.exe

3. Printisolationhost.exe

Counter name	Description
% Processor Time	The percentage of elapsed time that all process threads use to run instructions. An instruction is the basic unit of run time for a computer, a thread is the object that runs instructions, and a process is the object that is created when a program is run. This count includes code that is run to handle some hardware interrupts and trap conditions.
Handle Count	The total number of handles that are currently open by this process. This number is equal to the sum of the handles that are currently open by each thread in this process.
Virtual Bytes	The current size of the virtual address space for the process in bytes. Use of virtual address space does not necessarily imply a corresponding use of disk space or main memory pages. However, virtual address space is finite, and the process can limit the ability of the virtual space to load libraries when it uses too much of it.
Virtual Bytes Peak	The maximum size of virtual address space that the process has used at any one time in bytes. Use of virtual address space does not necessarily imply corresponding use of disk space or main memory pages. However, virtual address space is finite, and the process can limit the ability of the virtual space to load libraries when it uses too much of it.
Pool Paged Bytes	The size of the paged pool in bytes. This is an area of system memory (physical memory that is used by the operating system) for objects that can be written to disk when they are not in use. Memory\\Pool Paged Bytes is calculated differently than Process\\Pool Paged Bytes, and it might not equal Process\\Pool Paged Bytes_Total. This counter displays the last observed value only; it is not an average.

Counter name	Description
Pool Nonpaged Bytes	The size of the nonpaged pool in bytes. This is an area of system memory (physical memory that is used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated. Memory\\Pool Nonpaged Bytes is calculated differently than Process\\Pool Nonpaged Bytes, and it might not equal Process\\Pool Nonpaged Bytes_Total. This counter displays the last observed value only; it is not an average.
Working Set	The current size of the working set of the process in bytes. The working set is the set of memory pages that were recently touched by threads in the process. If free memory in the computer is above a certain threshold, pages are left in the working set of the process, even when not in use. When free memory falls below a certain threshold, pages are trimmed from the working sets. If needed, the pages are then soft-faulted back into the working set before leaving main memory.
Working Set Peak	The maximum size in the working set of the process in bytes at any point in time. If free memory in the computer is above a certain threshold, pages are left in the working set of the process, even when not in use. When free memory falls below a certain threshold, pages are trimmed from the working sets. If needed, the pages are then soft-faulted back into the working set before leaving main memory.

Disk-specific counters

The disk-specific counters for monitoring performance are described in the following table.

Counter name	Description
% Disk Time	Percentage of elapsed time during which the selected drive was busy servicing Read and

Counter name	Description
	Write requests.
Avg Disk Queue Length	Average number of Read and Write requests that were queued for the selected disk during the sample interval.

Memory-specific counters

The memory-specific counters for monitoring performance are described in the following table.

Counter name	Description
Available Mbytes	Amount of physical memory in megabytes. This value is immediately available for allocation to a process or for system use. It is equal to the sum of memory assigned to the standby (cached), free, and zero page lists.
% Committed Bytes in Use	In Use Ratio of Memory\\Committed Bytes to the Memory\\Commit Limit. Committed memory is the physical memory in use for which space has been reserved in the paging file, in the event that it needs to be written to disk. The commit limit is determined by the size of the paging file. If the paging file is enlarged, the commit limit increases and the ratio is reduced. This counter displays the current percentage value only; it is not an average.

CPU-specific counters

The CPU-specific counters for monitoring performance are described in the following table.

Counter name	Description
% Idle Time	Percentage of time that the processor is idle during the sample interval.
% Processor Time	Percentage of elapsed time that the processor spends executing a non-idle thread. This value is equal to the percentage of time the processor spends executing the idle thread, subtracted

Counter name	Description
	from 100 percent. (Each processor has an idle thread that consumes cycles when no other threads are ready to run). This counter is the primary indicator of processor activity. It displays the average percentage of busy time that is observed during the sample interval. The calculation used to determine whether the processor is idle is performed by using an internal sampling interval with the system clock (10 milliseconds). On today's faster processors, % Processor Time can underestimate the processor utilization because the processor may be spending time servicing threads between the sampling intervals. Workload-based timer applications are one example of applications that are more likely to be measured inaccurately because timers are signaled directly after the sample is taken.

Bandwidth-specific Counters

On the **Network Adapter** menu, select the network adapter from which the administrator want to monitor bandwidth.

Counter name	Description
Bytes Received/Sec	Rate at which bytes are received over each network adapter including the framing characters. Network Interface\Bytes Received/Sec is a subset of Network Interface\Bytes Total/Sec.
Bytes Sent/Sec	Rate at which bytes are sent over each network adapter including the framing characters. The Network Interface\Bytes Sent/Sec is a subset of Network Interface\Bytes Total/Sec.

Reference systems case study

Estimating the capacity of Print Server is difficult and it does not adhere to a specific formula. To give administrator's a starting point, this case study describes two reference systems that have been successfully deployed and tested by Microsoft.

Note

These case studies are for reference only. Using them will not guarantee performance. We recommend that you perform benchmarking to understand the system's capabilities in the intended environment as the best practice for estimating the capacity of Print Server.

Reference system 1

Reference system 1 is a standalone physical server running Windows Server 2012 R2. It services approximately 1,300 users across 1,000 queues. These queues host high-end, enterprise-level printers.

Note

The Print Server performance data section later in this guide provides performance data for Reference system 1 during a five-day interval, to provide a performance system overview for memory and processor time.

Name	Description
System manufacturer	Hewlett Packard
System model	ProLiant DL585 G2
System type	x64-based computer
Processor	Dual-core AMD Opteron(tm) processor 8212
BIOS version/date	HP A07, 6/27/2008
Total physical memory	4.0 GB
Total virtual memory	4.69 GB
Available virtual memory	2.23 GB
Network adapters	Two network adapters installed
Broadcom	BCM5706C NetXtreme II GigE (NDIS VBD Client)

Reference system 2

Reference system 2 is a large-scale virtual machine that is running on Windows Server 2012 R2. It covers approximately 8,000 users across 130 queues. These queues host high-end enterprise level printers.

Name	Description
System manufacturer	Microsoft
System model	Virtual machine
System type	x64-based PC
Processor	Intel Xeon processor L5640 2.27 GHz 4 cores 4 logical
BIOS version and date	American Megatrends, Inc. 090006, 5/23/2012
Total physical memory	8.0 GB
Total virtual memory	20 GB
Available virtual memory	15 GB
Network adapters	One Microsoft Hyper-V network adapter installed [01]

Recommendations for deployment

The best practices recommendations that follow are based on the factors described in the previous sections, to help you achieve Print Server scalability, maximize your hardware investment, and provide the best performance.

1. Locate Print Server on a Hyper-V-hosted virtual machine to allow the hardware capacity to be adjusted on demand, and to provide failover clustering as needed. For more information, see System hardware configuration earlier in this guide.
2. Use v4 printer drivers with Print Server, and make the same drivers available in Windows Server Update Services (WSUS) or as part of the default operating system image.

   Note
   If the organization has a large number of Windows 7 clients, use v3 printer drivers with client-side rendering enabled for high-volume queues.
3. Enable Branch Office Direct Printing if feasible within the organization.

   For more information, including the limitations of this feature, see Branch Office Direct Printing Technical Details.
4. Standardize printer models across deployments for easier set up and manageability.

5. Deploy printers that have been signed by Microsoft and certified by the Windows Logo Program.
6. For high-volume servers, consider moving the spool directory to another physical drive.
7. Use Performance Monitor to conduct a performance analysis on an existing Print Server to identify usage trends and bottlenecks.

Administrators can use the information in the following table as a starting point when planning deployments.

Condition	Options	Impact	Suggested action
Volume of print jobs	Low/Medium/High	High	Review the effect of the print job, characteristics of the print job, and spool activity. For more information, see: Printing related factors
Maximum concurrent users	1-50, 51-200, or more than 200	High	Review the effect of the client load. For more information, see: Number of clients
Mixed client operating systems	Windows 8.1, Windows 8, and Windows 7	Medium	Review the effect of mixed client operating systems. For more information, see: Operating system of the print clients
Hosted drivers	Mostly v4 or mostly v3	Medium	Review the effect of mixed drivers. For more information, see: Mix of drivers hosted on print server (v4 or v3 printer drivers)
Number of printers available	1 to 50, 51 to 200, or more than 200	Low	Review the effect of the printer queue. For more information, see: Non-Microsoft software

Print Server performance data

The following images show Print Server performance data for Reference system 1 during a five-day interval. These figures provide a performance system overview for memory and processor time.

Figure 1: Print server performance over a five-day interval

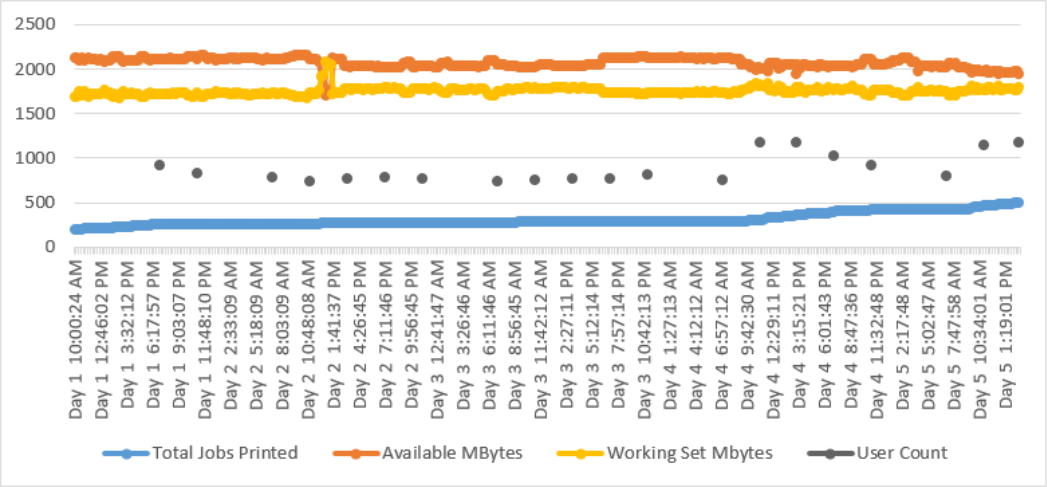

Figure 2: Print server performance showing processor time

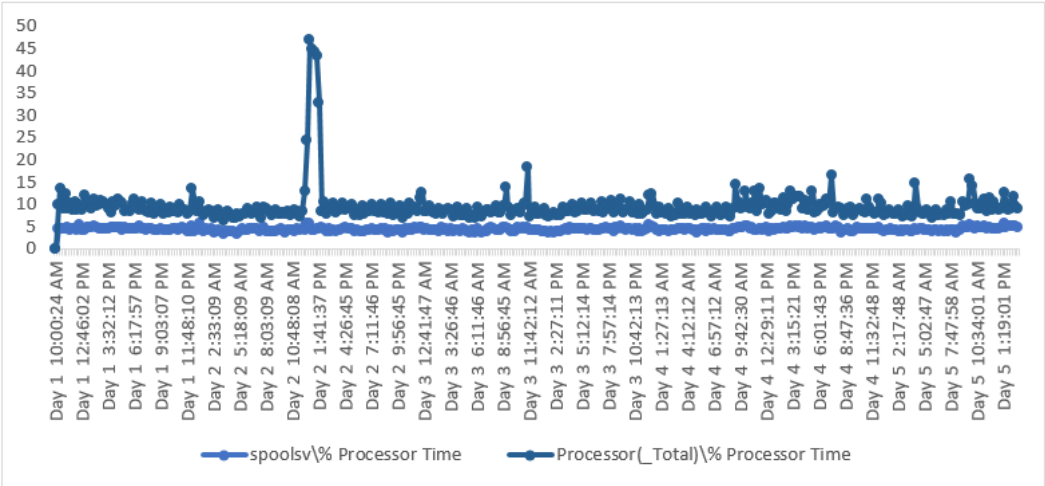

See also

The following resources provide additional information about deploying Print Server.

- Infrastructure Planning and Design Guide Series (Print Server for Windows)

Note
Although this guide describes Print Server for Windows Server 2008 R2, the general deployment methodology remains applicable for Windows Server 2012 R2 and Windows Server 2012.

- Print and Document Services Overview (for Windows Server 2012)
- Print and Document Services Architecture
- Using Performance Monitor
- V4 Print Driver
- Microsoft Support article, HOW TO: Move the Print Spool Folder in Windows 2000
- Filter Pipeline Configuration File
- Using the XPS Rasterization Service
- XPSRas GPU Usage Decision Tree
- Clustering and High Availability, Creating a Windows Server 2012 Failover Cluster
- Failover Clustering Hardware Requirements and Storage Options
- Branch Office Direct Printing Technical Details

High Availability Printing Overview

This topic introduces the High Availability Printing feature in Windows Server 2012 and describes how it can make your print servers highly available using Hyper-V virtual machines running as nodes in a cluster.

Did you mean…

- Print and Document Services Overview
- Print and Document Services in Windows Server 2008 R2

Feature description

High Availability Printing, or HA Printing, is the combination of Microsoft Failover Clustering technologies and Hyper-V technologies deployed to make a Windows Server 2012 Print Server highly available.

A server cluster is a group of independent servers working collectively and running the Microsoft Failover Clustering feature. Server clusters provide high availability, failback, scalability, and manageability for resources and applications.

High Availability Printing Architecture

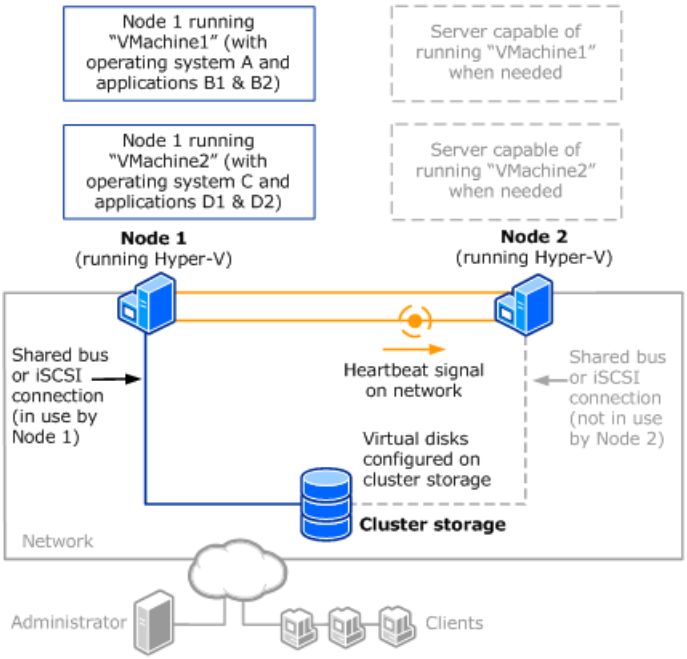

Practical applications

Using server clusters allows client access to the print services in the event of failures and planned outages. If one of the servers in the cluster is unavailable because of a failure or maintenance requirements, print services move to other available cluster nodes.

Important functionality

In contrast to previous versions of Windows Server, Windows Server 2012 defines a highly available print server as a Hyper-V virtual machine running on a node in a cluster. A single virtual machine with the Print Server role installed can then be migrated from one node in the Hyper-V cluster to the other using either manual or automatic methods.

In Windows Server 2012, the print spooler service is no longer a clustered resource and instead the entire virtual machine is migrated from one Hyper-V node to the other. This new model provides the same seamless user experience as previous versions of Windows but with the following added benefits:

- Windows Server 2012 Print Servers can utilize the Live Migration and Quick Migration features of Hyper-V.
- Using Windows Server 2012 as the Hyper-V and failover clustering host allows access to the new Virtual Machine Monitoring feature. For more information, see Virtual machine application monitoring and management, and How to configure VM Monitoring in Windows Server 2012 on MSDN.

- Windows Server 2012 Highly Available Print Servers are easier to deploy and have reduced complexity.
- An HA Print Server can now be deployed on Server Core installations of Windows.
- Problems with print devices and drivers that were not designed to work in a server cluster environment are now greatly reduced, if not completely eliminated.
- Backup, restore, and migration of Highly Available Print Servers is greatly simplified.

  As previously described, Windows Server 2012 defines a highly available print server as a virtual machine running on a clustered server running Hyper-V. A single virtual machine with the Print Server role installed can then be migrated from one cluster node to the other using either manual or automatic methods.
- Drivers do not need to be replicated from one physical node to the other. Only one copy of a driver needs to be installed on the virtual machine.
- Driver compatibility problems on failover clusters are eliminated. No additional driver code is needed to support a clustered print spooler.
- Backup and restore of a print server configuration is greatly simplified as the backup/restore solution no longer needs to be aware of a clustered print server.
- Scripting solutions using WMI no longer need to be modified in order to function on a clustered print server.
- Servicing of printer drivers is simplified as it is no longer necessary to store driver files in a non-default path.
- Cluster specific code in the print spooler process is no longer necessary.
- Running as a virtual machine allows an administrator to take advantage of the Virtual Machine Monitoring feature in Windows Server 2012 and have greater flexibility and control over recovery actions.

Notes

When you deploy the new High Availability Printing feature, you could experience downtimes in the following situations:

- A virtual machine is restarting from Virtual Machine Monitoring. Most likely this will only be only a few minutes.
- A virtual machine needs to be updated which requires a restart.

Hardware requirements

To deploy High Availability printing, you must have a computer capable of running Hyper-V on Windows Server 2012. For more information, see Hyper-V Overview.

Software requirements

A functioning Hyper-V cluster with at least two nodes is required in order to deploy High Availability Printing with Windows Server 2012.

Windows Server 2012 Failover Clustering requires that each cluster node is joined to an Active Directory domain. Although not recommended, a Windows Server 2012 server with the Failover Clustering feature enabled can also act as a domain controller.

To install, configure, and manage Microsoft Failover Clusters and servers running Hyper-V, a user must be an administrator or a member of the local administrators group on each node.

See also

For more information related to High Availability Printing, see the following topics:

Content type	References
Deployment	Install and Configure High Availability Printing
Related technologies	Print and Document Services Overview | Failover Clustering Overview | Hyper-V Overview | VM Monitoring in Windows Server 2012 – Frequently Asked Questions

Install and Configure High Availability Printing

Before attempting to install and configure a highly available Print Server, install and configure Hyper-V and Failover Clustering on at least two physical servers. Once Hyper-V and Failover Clustering have been installed on both servers, create a cluster and make sure the virtual machines are stored on a Cluster Shared Volume. Once the initial setup of Hyper-V and Failover Clustering is complete, install a Windows Server 2012 virtual machine and configure the virtual machine for failover.

For more information about Hyper-V Failover Clustering, see Failover Clustering Overview.

Once the virtual machine that will host Print and Document services has been configured with initial settings, the Print and Document role services can be installed using Add Roles and Features in Server Manager.

For more information about installing Print and Document Services, see Install Print and Document Services.

Configure and manage high availability printing

Management of a Windows Server 2012 Print Server is done using the Print Management Console (PMC) MMC snap-in, which has been updated to support the new Windows Server 2012

features. Internet Printing is also managed using the PMC, and LPD Printing uses command line tools for configuration and remains unchanged from Windows Server 2008 R2.

Windows Server 2012 introduces a new set of Windows PowerShell cmdlets that can be used to manage print servers and print queues from the command line. Wherever possible, examples on how to use these cmdlets will be shown for common print management scenarios. For more information, see Print Management Cmdlets in Windows PowerShell.

Configuration and management of the Print Server role as a highly available print server is identical to that of a standalone print server. For details on how to use the Print Management Console to configure and manage a print server, see Configure Print and Document Services.

The following section will guide you through the steps to configure a highly available Print Server with Virtual Machine Monitoring:

- Step 1: Configure the firewall
- Step 2: Configure Virtual Machine Monitoring
- Step 3: Verify the Virtual Machine Monitoring configuration
- Step 4: Testing Virtual Machine Monitoring

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Step 1: Configure the firewall

Prior to enabling the Virtual Machine Monitoring feature, the firewall must be correctly configured on the highly available print server virtual machine.

To configure the firewall on the HA Print Server virtual machine

1. Open **Server Manager** and then click **Local Server** in the navigation pane.
2. Click the link to the right of **Windows Firewall** to open the Windows Firewall Control Panel. This link will display differently depending on the state of the firewall.
3. Click **Allow an app or feature through Windows Firewall** in the left pane to open the Allowed Apps control panel applet.
4. Click **Change Settings** and scroll down to **Virtual Machine Monitoring**. Check the boxes for **Domain**, **Private**, and **Public** and then click **OK**.

Step 2: Configure Virtual Machine Monitoring

The Virtual Machine Monitoring feature of Windows Server 2012 allows for greater control of a clustered virtual machine by monitoring the Print Spooler service in the virtual machine for failures, or by monitoring specific events on the virtual machine. If the Print Spooler fails or a specific event is logged in the event log, the clustering virtual machine monitoring feature will invoke the following recovery actions:

- By default, the Print Spooler service recovery actions are configured so that the Service Control Manager will handle the first two failures of the Print Spooler service. On the first and second service failures, Service Control Manager will restart the Print Spooler service.
- The Cluster service will handle the third failure of the Print Spooler service and will restart the virtual machine using the Virtual Machine Monitoring feature.
- On the first occurrence of an event, Virtual Machine Monitoring will restart the virtual machine.
- On the second occurrence of an event, Virtual Machine Monitoring will failover the virtual machine to another node in the cluster in the event of a second failure within the RetryPeriod and then restart the virtual machine. By default, the RetryPeriod is 15 minutes.
- If the Cluster service has moved the virtual machine to all nodes of the cluster within the FailoverPeriod then the virtual machine is left in a failed state on the last node that owned the virtual machine. By default, the FailoverPeriod is 6 hours.

The Virtual Machine Monitoring feature is installed by default when the Failover Clustering role is installed on the Hyper-V host but is not configured. The features needed for Virtual Machine Monitoring to take corrective action of a virtual machine are built-in to Windows Server 2012, however the firewall must be correctly configured. Configuration of the Virtual Machine Monitoring feature is done using either the Failover Cluster Manager or Windows PowerShell.

For example, in Windows Server 2012 there are a couple Print Spooler failure scenarios that you may want to monitor with Virtual Machine Monitoring:

- A failure of the Print Spooler service.
- A particular Event ID, such as 373 logged in the Microsoft-Windows-PrintService/Admin event log:

```
Log Name:      Microsoft-Windows-PrintService/Admin
Source:        Microsoft-Windows-PrintService
Date:          2/14/2011 4:05:25 PM
Event ID:      373
Task Category: Printing a document
Level:         Error
Keywords:      Classic Spooler Event,Document Print Job
User:          SYSTEM
Computer:
Description:
The spooler has detected that a component has an unusually large
number of open Graphical Device Interface (GDI) objects. As a
result, some enhanced metafile (EMF) print jobs might not print
until the spooler is restarted.
```

To configure a service to be monitored for failures on the virtual machine, you can use Failover Cluster Manager. Any node in the cluster can be used to configure the monitoring, but this feature cannot be configured using Hyper-V Manager.

To configure the service that will be monitored by the Virtual Machine Monitoring feature

1. Open **Failover Cluster Manager** on one of the nodes of the cluster, and connect to the cluster that you wish to manage.
2. Click the **Roles** node under the cluster name that you wish to configure. Each virtual machine that is clustered will display in the middle pane.
3. Right click the virtual machine that will be monitored and select **More Actions**, and then **Configure Monitoring**.

Configure monitoring

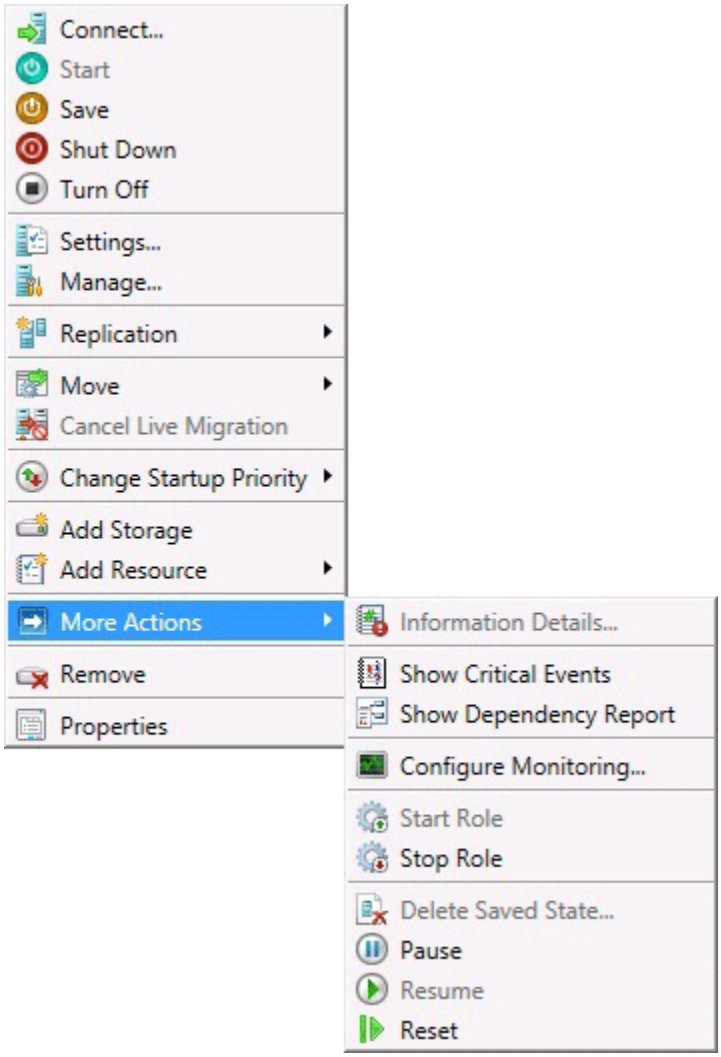

4. Select the **Print Spooler** from the list of services by checking the box next to it.

Select services

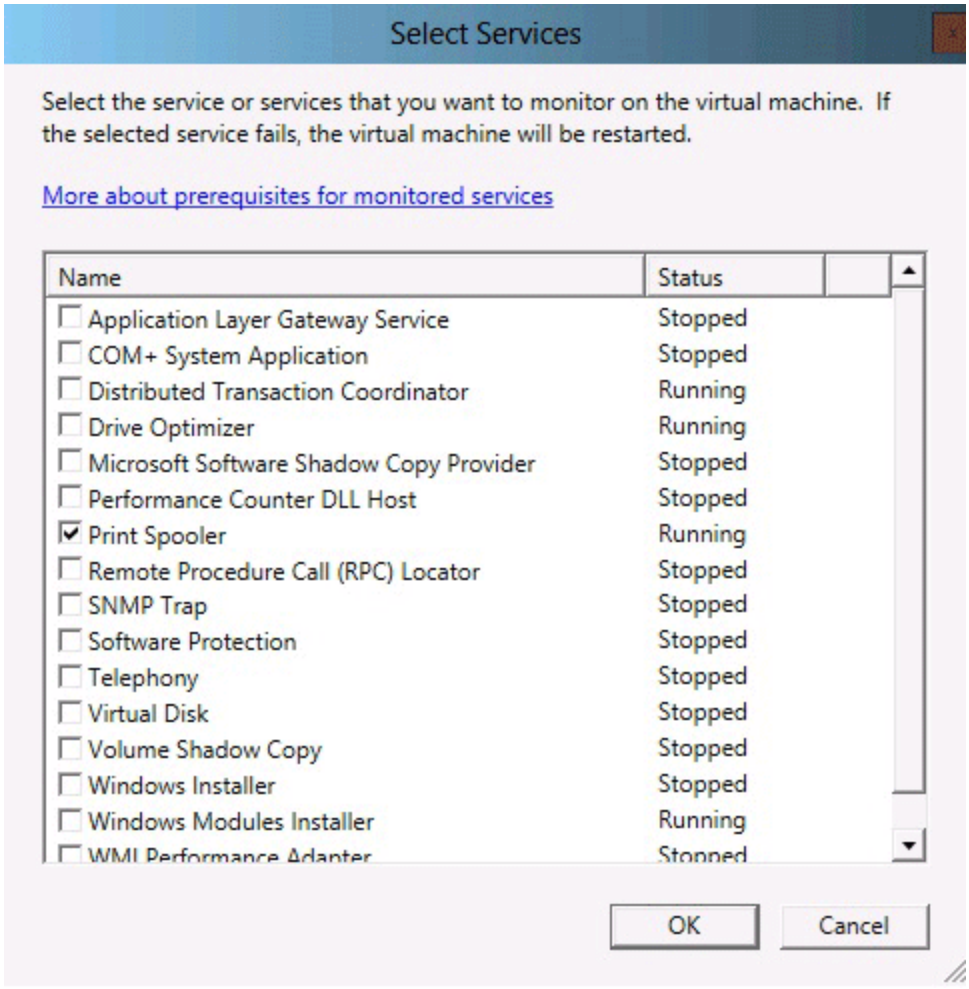

5. When successfully configured, the lower pane of Failover Cluster Manager should show Print Spooler listed next to **Monitored Services** when the HA Print Server is selected.

Monitored services

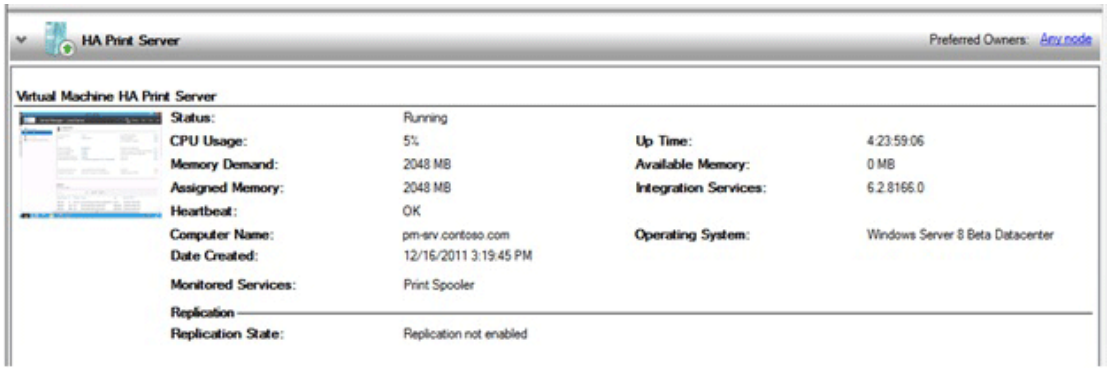

Windows PowerShell equivalent commands

To configure the Virtual Machine Monitoring feature to monitor the Print Server virtual machine Print Spooler service, open an administrative Windows PowerShell command prompt on the Hyper-V host that currently owns the highly available print server virtual machine and type the following command:

Windows PowerShell command

```
Add-ClusterVMMonitoredItem -vm "<name of print server>" -service spooler
```

Where **<name of print server>** is the name of the highly available print server virtual machine. The name should be the Hyper-V virtual machine name, not the NetBIOS or FQDN name of the print server on the network.

The Add-ClusterVMMonitoredItem cmdlet can be used to configure a VM for monitoring of a service and an event. Although service monitoring can also be configured using the Failover Cluster Manager interface, event monitoring can only be configured using Windows PowerShell.

To configure the Virtual Machine Monitoring feature to monitor the Print Server virtual machine event log for a particular event ID (for example 373), open an administrative Windows PowerShell command prompt on the Hyper-V host that currently owns the highly available print server virtual machine and type the following command:

Windows PowerShell command

```
Add-ClusterVMMonitoredItem -vm "<name of print server>" -eventlog "Microsoft-Windows-
PrintService/Admin" -eventsource Microsoft-Windows-PrintService -eventid 373
```

Where **<name of print server>** is the name of the highly available print server virtual machine. The name should be the Hyper-V virtual machine name, not the NetBIOS or FQDN name of the print server on the network.

Step 3: Verify the Virtual Machine Monitoring configuration

The Get-ClusterVMMonitoredItem cmdlet can be used to verify the configuration of Virtual Machine Monitoring for a specific virtual machine. Type the following command to display the Virtual Machine Monitoring configuration for a specific virtual machine:

Windows PowerShell command

```
Get-ClusterVMMonitoredIem "HA Print Server"
```

The output of this command for a correctly configured highly available print server should show that both spooler monitoring and event ID 373 monitoring are enabled:

```
Name

----

Microsoft-Windows-PrintService+Admin,Microsoft-Windows-PrintService,373

Spooler
```

Step 4: Testing Virtual Machine Monitoring

To verify the configuration of Virtual Machine Monitoring it is necessary to simulate a failure of the Print Spooler service or manually generate an event ID 373 using an event generating tool.

To test the functionality of Virtual Machine Monitoring of the highly available Print Server

1. Open Task Manager on the Print Server and terminate the Print Spooler process (spoolsv.exe) three times to simulate a spooler failure. The Service Control Manager will restart the Print Spooler process the first two times and Virtual Machine Monitoring will not take action until the third simulated spooler failure.

   Note
   Stopping the Print Spooler service is not considered a failure and will not trigger a recovery action by the Virtual Machine Monitoring feature.
2. Failover Cluster Manager will indicate **Running (Application in VM Critical)** on the virtual machine status.
3. An Event ID 1250 indicating an embedded failure will be logged in the system event log.
4. The highly available print server virtual machine will restart automatically.
5. After the virtual machine reboots, terminate the Print Spooler process a second time within 15 minutes using Task Manager.
6. Failover Cluster Manager will indicate that the application in the virtual machine is critical and another event 1250 is logged in the system event log.
7. The highly available print server virtual machine is shutdown and then should failover based on the cluster configuration.

See Also

High Availability Printing Overview

Client Name Instrumentation in Windows Server 2012 R2

Client Name Instrumentation Overview

Windows Server 2012 R2 implements unique username tracking to measure the number of unique users connecting to and consuming Print/Fax services.

Client Name Instrumentation Limitations

- User access logging (UAL) uniqueness is not guaranteed for LPD/LPR print jobs.

- IPP and SMB scenarios will always show a user name with the localhost IP.
- Named pipes are not logged.
- Windows 8.1 RT clients connecting to v3 print queues will not be logged.

Client Name Instrumentation UAL Cmdlets

- UAL Cmdlets in Windows PowerShell
- Get-UalUserAccess
- Get-UalDailyUserAccess

Remote Desktop Services Overview

Remote Desktop Services accelerates and extends desktop and application deployments to any device, improving remote worker efficiency, while helping to keep critical intellectual property secure and simplify regulatory compliance. Remote Desktop Services enables virtual desktop infrastructure (VDI), session-based desktops, and applications, allowing users to work anywhere.

Did you mean…

- Terminal Services in Windows Server 2008
- Remote Desktop Services in Windows Server 2008 R2
- Remote Desktop Services resources (all versions)
- Microsoft Server and Cloud Platform

Role description

The Remote Desktop Services role provides technologies that enable users to connect to virtual desktops, RemoteApp programs, and session-based desktops. With Remote Desktop Services, users can access remote connections from within a corporate network or from the Internet.

Practical applications

Remote Desktop Services allows workers to work anywhere. Some of the key benefits of Remote Desktop Services include:

- **Unified administration experience** – Administer your session and virtual desktop collections, configure your RemoteApp programs, manage your virtual desktops, and add servers to the deployment from one centralized console.
- **User personalization** – User profile disks allow you to preserve user personalization settings across session collections and pooled virtual desktop collections.
- **Less expensive storage** – Pooled virtual desktops can use local storage live migration between host computers. Personal virtual desktops can use storage located on network shares.

- **Automated pooled virtual desktop management** – Deploy and manage pooled virtual desktops centrally by using a virtual desktop template. Any changes, such as application installation or security updates, are installed on the virtual desktop template, and the pooled virtual desktops are then recreated from the virtual desktop template.

New and changed functionality for Windows Server 2012

Remote Desktop Services enables the mobile work force to connect to desktop and applications from anywhere. In Windows Server 2012, Remote Desktop Services includes enhancements in the following areas:

- Simplified Virtual Desktop Infrastructure (VDI) deployment and management
- Simplified Session Virtualization deployment and management
- Centralized resource publishing
- Rich user experience with Remote Desktop Protocol (RDP)

The user experience has been enhanced for Remote Desktop Services in Windows Server 2012 in the following ways:

- Rich Windows desktop remoting experience
- Smooth audio and video playback experience
- Rich graphics and video user experience over a WAN
- Enhanced device remoting support with USB Redirection for Session Virtualization and VDI
- True Multi-Touch and gesture remoting
- Email name discovery and subscription to administrator supplied remote resources
- RemoteFX virtualized GPU provides DX 11 support

In addition to these areas of enhancement, Remote Desktop Services in Windows Server 2012 introduces a new management console for managing the majority of Remote Desktop Services-related tasks. For more information about new features and functionality, see What's New in Remote Desktop Services in Windows Server 2012.

New and changed functionality for Windows Server 2012 R2

In Windows Server 2012 R2, Remote Desktop Services includes enhancements in the following areas:

- Monitor and control by using session shadowing
- Reduced storage requirements and improved performance accessing common data
- RemoteApp programs perform more like locally-based applications
- Improved reconnection performance for remote clients
- Improved compression allowing improved usage of network bandwidth
- Display resolution changes are automatically reflected on the remote client

- RemoteFX virtualized GPU supports DX11.1

For more information about new features and functionality, see What's New in Remote Desktop Services in Windows Server 2012 R2.

New Microsoft Remote Desktop Clients

You can use the Microsoft Remote Desktop client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a remote desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps using the RD client.

- Remote Desktop Client for Android
  - Getting Started with Remote Desktop Client on Android
  - Remote Desktop Client on Android: FAQ
- Remote Desktop Client for iOS
  - Getting Started with Remote Desktop Client on iOS
  - Remote Desktop Client on iOS: FAQ
- Remote Desktop Client for Mac
  - Getting Started with Remote Desktop Client on Mac
  - Remote Desktop Client on Mac: FAQ
- Remote Desktop Client for Windows
  - Getting Started with Remote Desktop Client on Windows
  - Remote Desktop Connection: FAQ

For more information about new features and functionality, see Microsoft Remote Desktop Clients.

Removed or deprecated functionality

For a list of deprecated features, see Features Removed or Deprecated in the Windows Server Developer Preview.

Hardware requirements

Remote Desktop Services requires that the Windows Server 2012 operating system be installed. There are no additional hardware or software requirements for running Remote Desktop Services.

There are several hardware requirements that must be met when you deploy RemoteFX virtualized GPU to hardware accelerate your Windows client virtual desktops:

- **SLAT-enabled processor**. The processor on the RemoteFX server must support Second-Level Address Translation (SLAT).

- **GPU**. At least one graphics processing unit (GPU) that is capable of supporting RemoteFX is required on the RemoteFX server. The GPU driver must support DirectX 11.

  Note
  Without a RemoteFX virtualized GPU, applications that require DirectX will still work using a built in Hyper-V specific GPU.

Server Manager information

Remote Desktop Services is a server role that consists of several role services. In Windows Server 2012, the following Remote Desktop Services role services can be installed with this role:

Role service name	Role service description
RD Virtualization Host	Remote Desktop Virtualization Host (RD Virtualization Host) integrates with Hyper-V to deploy pooled or personal virtual desktop collections within your organization.
RD Session Host	Remote Desktop Session Host (RD Session Host) enables a server to host RemoteApp programs or session-based desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers.
RD Connection Broker	Remote Desktop Connection Broker (RD Connection Broker): • Allows users to reconnect to their existing virtual desktops, RemoteApp programs, and session-based desktops. • Enables you to evenly distribute the load among RD Session Host servers in a session collection or pooled virtual desktops in a pooled virtual desktop collection. • Provides access to virtual desktops in a virtual desktop collection.
RD Web Access	Remote Desktop Web Access (RD Web Access) enables users to access RemoteApp and Desktop Connection through the **Start** menu on a computer that is running Windows 8, Windows 7, or through a web browser. RemoteApp and Desktop Connection provides a customized view of RemoteApp programs

Role service name	Role service description
	and session-based desktops in a session collection, and RemoteApp programs and virtual desktops in a virtual desktop collection.
RD Licensing	Remote Desktop Licensing (RD Licensing) manages the licenses required to connect to a Remote Desktop Session Host server or a virtual desktop. You can use RD Licensing to install, issue, and track the availability of licenses.
RD Gateway	Remote Desktop Gateway (RD Gateway) enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device.

See also

The following table provides additional resources for evaluating Remote Desktop Services.

Content type	References
Product evaluation	• Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start • Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment • Test Lab Guide: Virtual Desktop Infrastructure Quick Start • Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment • Test Lab Guide: Managed Pooled Virtual Desktop Collections • Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections • Test Lab Guide: Remote Desktop Services Publishing • Test Lab Guide: Remote Desktop Services Licensing • Remote desktop experience in Windows 8

Content type	References
	• Test Lab Guide: Demonstrate Remote Desktop Services in Windows Server "8" Beta • Test Lab Guide: Demonstrate Remote Desktop Services Desktop Virtualization in Windows Server "8" Beta • Test Lab Guide: Deploying Remote Desktop Licensing in Windows Server "8" Beta • Understand and Troubleshoot Remote Desktop Services Desktop Virtualization in Windows Server “8” Beta • Understand and Troubleshoot Remote Desktop Services in Windows Server “8” Beta

Community resources	• Remote Desktop Services (Terminal Services) Team Blog
	• Remote Desktop Services Web Forum
	• Remote Desktop Services TechNet Wiki
	• Windows Server 2008 R2 Virtualization Web Forum
	• Windows Virtualization Team Blog
	• Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (01) Introduction to Desktop Virtualization
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (02) Optimizing the User Experience
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (03) Deploying Virtual Desktops with Windows Server 2012 and RDS
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (04) Hyper-V for VDI
	  • Microsoft Virtual Academy: Using

Content type	References
	Microsoft VDI to Enable New Workstyles: (05) Capacity Planning and Architecture • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (06) Leveraging Citrix • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (07) Microsoft VDI Licensing
Related technologies	• Terminal Services in Windows Server 2008 • Remote Desktop Services in Windows Server 2008 R2 • Remote Desktop Services resources (all versions)

What's New in Remote Desktop Services in Windows Server 2012 R2

This topic describes the Remote Desktop Services functionality that is new or changed in Windows Server 2012 R2.

The Remote Desktop Services server role in Windows Server® 2012 provided technologies that enabled users to connect to virtual desktops, RemoteApp programs, and session-based desktops. With Remote Desktop Services, users can access remote connections from within a corporate network or from the Internet.

In Windows Server 2012, Remote Desktop Services offers support for the following enhanced features:

- Virtual Desktop Infrastructure (VDI) deployments
- Session Virtualization deployments
- Centralized resource publishing
- Rich user experience with Remote Desktop Protocol (RDP)

For more information about new features and functionality for Windows Server 2012, see What's New in Remote Desktop Services in Windows Server 2012.

In Windows Server 2012 R2, Remote Desktop Services offers enhanced support in the following areas:

- Session Shadowing
- Online Storage Deduplication

- Improved RemoteApp behavior
- Quick reconnect for remote desktop clients
- Improved compression and bandwidth usage
- Dynamic display handling
- RemoteFX virtualized GPU supports DX11.1

Each of these is described in the sections that follow.

Session Shadowing

In Windows Server 2012 R2, Session Shadowing enables you to remotely monitor or control an active session of another user on a Remote Desktop Session Host (RD Session Host) server. The current version includes integration with Server Manager and Remote Desktop Connection (mstsc.exe).

Online Data Deduplication

In Windows Server 2012 R2 the Data Deduplication feature, released in Windows Server 2012, can be used with actively running personal desktop collections when the desktop virtual hard disks (VHDs) are stored on a file server running Windows Server 2012 R2 and accessed by using Server Message Block (SMB). Storage capacity requirements can be dramatically reduced by using the Data Deduplication feature. Data Deduplication on the SMB server caches frequently accessed data, allowing performance improvements on many read-intensive operations, including parallel remote client boot.

Improved RemoteApp behavior

In Windows Server 2012 R2 RemoteApp programs are one step closer to the look and feel of local applications by including support for transparency, live thumbnails, and seamless application move that allows the application content to remain visible while the application is moved on screen.

Quick reconnect for remote desktop clients

In Windows Server 2012 R2 Quick Reconnect improves connection performance enabling users to reconnect to their existing virtual desktops, RemoteApp programs, and session-based desktops more quickly. The connection process for RemoteApp programs has been redesigned for Windows® 8.1 and Windows Server 2012 R2 clients, to be more informative and user friendly.

Improved compression and bandwidth usage

Windows Server 2012 R2 improves performance on the Windows Server 2012 remote experience, by using codecs that enable better compression, delivering bandwidth savings (for

example video content delivery over a WAN utilizes up to 50% less bandwidth compared to Windows Server 2012).

Dynamic display handling

In Windows 8.1 and Windows Server 2012 R2 support has been added for display changes on the client to be automatically reflected on the remote client. This translates to seamless device rotation, and monitor addition and removal (for example connecting to a projector or docking a laptop) for both remote sessions and RemoteApp programs.

RemoteFX virtualized GPU supports DX11.1

The RemoteFX vGPU in Windows Server 2012 R2 includes support for DX11.1 on systems that have DX11.1 capable GPU cards. Graphics intensive applications that rely on DX11.1 features can be virtualized and run in a virtualized environment on Windows 8.1 and Windows Server 2012 R2. Windows Server 2012 R2 introduces the following functionality:

- Non-Uniform Memory Access (NUMA) support: RemoteFX running on NUMA-based platforms will experience improved scaling capabilities.
- Video RAM (VRAM) changes: Adding system memory, to the server running Hyper-V, will allow for a dynamic increase in virtual machine VRAM which can improve performance for applications.

RestrictedAdmin Mode Remote Desktop

Windows 8.1 and Windows Server 2012 R2 support a new remote desktop option to connect in RestrictedAdmin mode. When connecting using RestrictedAdmin mode, the user's credentials are not sent to the host by the remote desktop client. Using this mode with administrator credentials, the remote desktop client attempts to interactively logon to a host that also supports this mode without sending credentials. When the host verifies that the user account connecting to it has administrator rights and supports Restricted Admin mode, the connection is successful. Otherwise, the connection attempt fails. Restricted Admin mode does not at any point send plain text or other re-usable forms of credentials to remote computers.

Note

Once connected to a host in RestrictedAdmin mode, the user will not be able to seamlessly access other network resources from that host using the credentials they provided to the remote desktop client.

To connect using RestrictedAdmin mode requires

- Windows 8.1 or Windows Server 2012 R2 for the remote desktop host
- Windows 8.1 or Windows Server 2012 R2 for the remote desktop client

Remote Desktop Services role services description

Remote Desktop Services is a server role that consists of several role services. In Windows Server 2012 R2, Remote Desktop Services consists of the following role services:

Role service name	Role service description
RD Virtualization Host	Remote Desktop Virtualization Host (RD Virtualization Host) integrates with Hyper-V to deploy pooled or personal virtual desktop collections within your organization.
RD Session Host	Remote Desktop Session Host (RD Session Host) enables a server to host RemoteApp programs or session-based desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers.
RD Connection Broker	Remote Desktop Connection Broker (RD Connection Broker): • Allows users to reconnect to their existing virtual desktops, RemoteApp programs, and session-based desktops. • Enables you to evenly distribute the load among RD Session Host servers in a session collection or pooled virtual desktops in a pooled virtual desktop collection. • Provides access to virtual desktops in a virtual desktop collection.
RD Web Access	Remote Desktop Web Access (RD Web Access) enables users to access RemoteApp and Desktop Connection through the **Start** menu on a computer that is running Windows 8, Windows 7, or through a web browser. RemoteApp and Desktop Connection provides a customized view of RemoteApp programs and session-based desktops in a session collection, and RemoteApp programs and virtual desktops in a virtual desktop collection.
RD Licensing	Remote Desktop Licensing (RD Licensing)

Role service name	Role service description
	manages the licenses required to connect to a Remote Desktop Session Host server or a virtual desktop. You can use RD Licensing to install, issue, and track the availability of licenses.
RD Gateway	Remote Desktop Gateway (RD Gateway) enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device.

Removed or deprecated functionality

For a list of deprecated features, see Features Removed or Deprecated in the Windows Server Developer Preview.

See also

The following table provides additional resources for evaluating Remote Desktop Services.

Content type	References
Product evaluation	• Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start • Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment • Test Lab Guide: Virtual Desktop Infrastructure Quick Start • Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment • Test Lab Guide: Managed Pooled Virtual Desktop Collections • Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections • Test Lab Guide: Remote Desktop Services Publishing • Test Lab Guide: Remote Desktop Services Licensing
Community resources	• Remote Desktop Services (Terminal

Content type	References

	Services) Team Blog
	• Remote Desktop Services Web Forum
	• Remote Desktop Services TechNet Wiki
	• Windows Server 2008 R2 Virtualization Web Forum
	• Windows Virtualization Team Blog
	• Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (01) Introduction to Desktop Virtualization
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (02) Optimizing the User Experience
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (03) Deploying Virtual Desktops with Windows Server 2012 and RDS
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (04) Hyper-V for VDI
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (05) Capacity Planning and Architecture
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (06) Leveraging Citrix
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (07) Microsoft VDI Licensing

Related technologies	• Terminal Services in Windows Server 2008 • Remote Desktop Services in Windows Server 2008 R2 • Remote Desktop Services in Windows Server 2012 • Remote Desktop Services resources (all

Content type	References
	versions)

What's New in Remote Desktop Services in Windows Server 2012

The Remote Desktop Services server role in Windows Server® 2012 provides technologies that enable users to connect to virtual desktops, RemoteApp programs, and session-based desktops. With Remote Desktop Services, users can access remote connections from within a corporate network or from the Internet.

In Windows Server 2012, Remote Desktop Services offers enhanced support for the following scenarios:

- Virtual Desktop Infrastructure (VDI) deployments
- Session Virtualization deployments
- Centralized resource publishing
- Rich user experience with Remote Desktop Protocol (RDP)
- Rich graphics experience with RemoteFX vGPU

Each of these is described in the sections that follow.

Virtual Desktop Infrastructure (VDI) deployment

Remote Desktop Services introduced a VDI deployment in Windows Server 2008 R2. In Windows Server 2012, Remote Desktop Services includes new ways to efficiently configure and manage your virtual desktops. Some of the enhancements include:

- **Unified central experience** – Deploy VDI quickly, and then manage your pooled and personal virtual desktop deployments through a new unified central experience.
- **Automated and simple single-image management** – Take advantage of automated ways to deploy and manage pooled virtual desktops with a virtual desktop template.
- **User personalization** – Preserve user personalization settings for pooled virtual desktop deployments by using user profile disks.
- **Less expensive storage** – Use inexpensive local storage with live migration functionality between host computers for pooled virtual desktops. Personal virtual desktops can use the less expensive SMB central storage.

Session Virtualization deployment

In Windows Server 2012, Session Virtualization deployment in Remote Desktop Services includes new ways to efficiently configure and manage your session-based desktops. In earlier versions of Remote Desktop Services, ongoing management of the RD Session Host servers is

performed at a per server level. By using a Session Virtualization deployment scenario, centralized management and installation is enabled. Session Virtualization in Windows Server 2012 offers the following benefits:

- **Unified central experience** – In Windows Server 2012, you can deploy Session Virtualization quickly and manage your deployments through a new unified central experience.
- **Simplified and centralized deployment** – Simple scenario-based installations allow you to create an entire session collection at one time.
- **User personalization** – User profile disks allow you to preserve user personalization settings for your session collections.
- **Centralized and unified management** – Manage all of the RD Session Host servers in your session collection from a single location.
- **Fairshare experience** – For a predictable user experience in Windows Server 2012 and to ensure that one user does not negatively impact the performance of another user's session, the following features are enabled by default on RD Session Host servers:
    - **Network Fairshare** – Dynamically distributes available bandwidth across sessions based on the number of active sessions to enable equal bandwidth utilization.
    - **Disk Fairshare** – Prevents sessions from over utilizing disk usage by equal distribution of disk I/O among sessions.
    - **CPU Fairshare** – Dynamically distributes processor time across sessions based on the number of active sessions and load on these sessions. This was introduced in Windows Server 2008 R2 and has been improved for heavier loads in Windows Server 2012.

A Session Virtualization deployment consists of RD Session Host servers along with infrastructure servers, such as RD Licensing, RD Connection Broker, RD Gateway, and RD Web Access servers.

A session collection (referred to as a farm in earlier versions of Windows Server) is a grouping of RD Session Host servers for a given session. A session collection is used to publish one of the following resources:

- Session-based desktops
- RemoteApp programs

Session Virtualization is a scenario-based installation within Server Manager that allows you to install, configure, and manage RD Session Host servers from a central location. With the Session Virtualization deployment scenario, you're presented with two deployment types:

- **Quick Start:** Installs all the necessary Remote Desktop Services role services on one computer to enable you to install and configure Remote Desktop Services role services in a test environment.
- **Standard deployment:** Allows you to flexibly deploy the various Remote Desktop Services role services on different servers.

Centralized resource publishing

In Windows Server 2008 R2, publishing and managing applications on both pooled and personal virtual desktops is a very time-consuming and costly process. Because RemoteApp programs only partially integrate with the native Windows® experience, they add to the management cost since there is no way to organize published RemoteApp programs to users.

Remote Desktop Services in Windows Server 2012 enables you to publish and manage resources, such as RemoteApp programs, session-based desktops, and virtual desktops, from a centralized console. Using this new publishing feature, you can get an historic view of resources assigned to end users, change published resources for any given collection, and edit properties of published resources.

In addition to the centralized console, you can now configure a RemoteApp and Desktop Connection URL by using Group Policy, and then give users access to the URL automatically through an email address.

Centralized resource publishing provides end users with an experience that can replace locally installed applications.

Rich user experience with Remote Desktop Protocol (RDP)

Windows Server 2008 R2 with Service Pack 1 and Windows 7 with Service Pack 1 introduced Microsoft RemoteFX, which enables the delivery of a full Windows user experience to a range of client devices, including rich clients, thin clients, and ultrathin clients. Windows Server 2012 builds on this platform to enable a richer and more seamless experience on all types of networks and devices. Specifically, Remote Desktop Protocol (RDP) enables a consistent user experience when connecting to centralized workspaces even on networks where bandwidth is limited and latency is high.

Following are the key benefits for end users and IT professionals in using RDP in Windows 8 for their centralized workspace deployments.

RDP key benefits to end users	RDP key benefits to IT professionals

• Rich end user experience
• Access to corporate desktop and applications from anywhere by using any device
• Secure, managed, and fully supported desktops and applications
	• Flexible deployment options
	  • Virtual desktops
	  • Session-based desktops
	  • RemoteApp programs
	  • Full desktop clients
	  • Thin clients
	  • Corporate networks
	  • Branch offices
	• Diverse workforce support

	• Roaming users • Telecommuters • Bring your own computer • Keep costs down

Rich graphics experience with RemoteFX vGPU

Windows Server 2008 R2 with Service Pack 1 and Windows 7 with Service Pack 1 introduced Microsoft RemoteFX vGPU support, a set of technologies for a rich PC-like experience for virtual desktops. Windows Server 2008 R2 with Service Pack 1 introduced the following functionality:

- Host-side remoting
- Render-capture-encode pipeline
- Highly efficient GPU-based encode
- Throttling based on client activity
- DirectX-enabled virtual graphics processing unit (vGPU)

Windows Server 2012 builds upon this platform and implements support for the vGPU from DX9 to DX11. The user experience is also improved including support for additional monitors at higher resolutions. As hardware acceleration proliferates to more applications like the web browser, the vGPU enables applications to run higher levels of DirectX within Remote Desktop Virtualization Host (RD Virtualization Host).

Maximum monitor resolution for virtual machines per Windows Version:

Maximum resolution	Monitors per virtual machine for Windows 7 with SP 1	Monitors per virtual machine for Windows 8
1024 x 768	4	8
1280 x 1024	4	8
1600 x 1200	3	4
1920 x 1200	2	4
2560 x 1600	-	2

Remote Desktop Services role services description

Remote Desktop Services is a server role that consists of several role services. In Windows Server 2012, Remote Desktop Services consists of the following role services:

Role service name	Role service description
RD Virtualization Host	Remote Desktop Virtualization Host (RD Virtualization Host) integrates with Hyper-V to deploy pooled or personal virtual desktop collections within your organization by using RemoteApp and Desktop Connection.
RD Session Host	Remote Desktop Session Host (RD Session Host) enables a server to host RemoteApp programs or session-based desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers.
RD Connection Broker	Remote Desktop Connection Broker (RD Connection Broker): • Allows users to reconnect to their existing virtual desktops, RemoteApp programs, and session-based desktops. • Enables you to evenly distribute the load among RD Session Host servers in a session collection or pooled virtual desktops in a pooled virtual desktop collection. • Provides access to virtual desktops in a virtual desktop collection.
RD Web Access	Remote Desktop Web Access (RD Web Access) enables users to access RemoteApp and Desktop Connection through the **Start** menu on a computer that is running Windows 8, Windows 7, or through a web browser. RemoteApp and Desktop Connection provides a customized view of RemoteApp programs and session-based desktops in a session collection, and RemoteApp programs and virtual desktops in a virtual desktop collection.
RD Licensing	Remote Desktop Licensing (RD Licensing) manages the licenses required to connect to a Remote Desktop Session Host server or a virtual desktop. You can use RD Licensing to install, issue, and track the availability of licenses.

Role service name	Role service description
RD Gateway	Remote Desktop Gateway (RD Gateway) enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device.

Removed or deprecated functionality

For a list of deprecated features, see Features Removed or Deprecated in the Windows Server Developer Preview.

See also

The following table provides additional resources for evaluating Remote Desktop Services.

Content type	References
Product evaluation	• Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start • Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment • Test Lab Guide: Virtual Desktop Infrastructure Quick Start • Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment • Test Lab Guide: Managed Pooled Virtual Desktop Collections • Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections • Test Lab Guide: Remote Desktop Services Publishing • Test Lab Guide: Remote Desktop Services Licensing
Community resources	• Remote Desktop Services (Terminal Services) Team Blog • Remote Desktop Services Web Forum • Remote Desktop Services TechNet Wiki • Windows Server 2008 R2 Virtualization Web Forum

Content type	References
	• Windows Virtualization Team Blog
Related technologies	• Terminal Services in Windows Server 2008 • Remote Desktop Services in Windows Server 2008 R2 • Remote Desktop Services resources (all versions)

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

In Windows Server 2012, the Microsoft Virtual Desktop Infrastructure (VDI) deployment is a scenario-based installation that allows you to install, configure, and manage your virtual desktops from a central location. With the VDI deployment scenario, you are presented with two deployment types:

Microsoft Virtual Desktop Infrastructure (VDI) Quick Start Installs all the necessary Remote Desktop Services role services on one computer to let you install and configure them in a test environment.

Microsoft Virtual Desktop Infrastructure (VDI) standard deployment Allows you to flexibly deploy the Remote Desktop Services role services to servers. To use the standard deployment instead of the Quick Start, see: Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment.

This test lab walks you through a **VDI Quick Start deployment**. The deployment process completes the following tasks:

- Installs the Remote Desktop Connection Broker (RD Connection Broker), Remote Desktop Virtualization Host (RD Virtualization Host), and Remote Desktop Web Access (RD Web Access) role services on a single computer.
- Creates a pooled virtual desktop collection that is based on a virtual hard disk of the virtual desktop template.
- Creates a Hyper-V network switch named RDS Virtual and assign the pooled virtual desktop to that switch.

In this guide

This topic contains instructions for setting up a test lab based on the Test Lab Guide: Base Configuration test lab and deploying a VDI Quick Start deployment that uses two servers and one

client computer. The resulting test lab demonstrates how to successfully deploy a VDI Quick Start deployment.

Important
The following instructions are for configuring a VDI Quick Start deployment test lab that uses the minimum number of computers. Individual computers are needed to separate the services that are provided on the network and to clearly show the desired functionality. This configuration is not designed to reflect best practices or to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed to work only on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named APP1 that is configured as the VDI Quick Start deployment server.
- One client computer running Windows 8 named CLIENT1 that is used to connect to the VDI Quick Start deployment.

In addition, the VDI Quick Start test lab includes:

- One subnet (10.0.0.0/24) that simulates an intranet named Corpnet.
- Computers connected to each other by using a Hyper-V virtual switch as shown in the following figure.

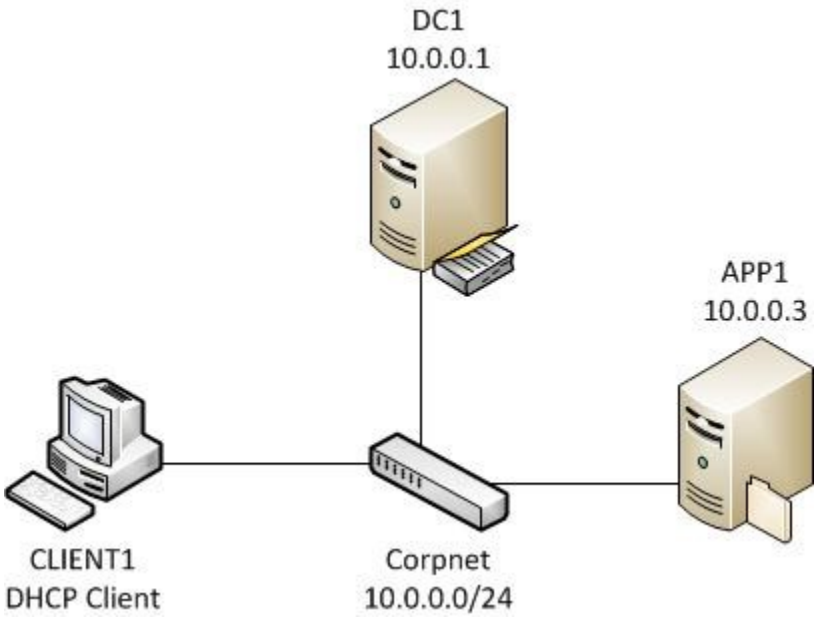

Hardware and software requirements

The following are required components for the test lab:

- The computer on which are you installing the Quick Start must have access to join computers to the domain. For the purposes of this guide, APP1 computer object must have permissions to join computers to the CORP domain.
- The virtual desktop template (a virtual hard disk) must be generalized by using Sysprep and then shut down.
- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- A physical computer that meets the minimum hardware requirements for Windows Server 2012.

Note

If you choose the VDI scenario-based installation, a virtual switch named RDS Virtual is automatically created.

Steps for configuring the test lab

There are four steps to follow when you set up a VDI Quick Start deployment test lab that is based on Test Lab Guide: Base Configuration.

Step 1: Complete the base configuration test lab The VDI Quick Start deployment test lab requires using the base configuration test lab as its starting point.

Step 2: Configure APP1 APP1 is a computer running Windows Server 2012 that is configured in the base configuration, and it will host the VDI Quick Start.

Important

You must have hardware-assisted virtualization on APP1 to successfully complete this test lab. Since all the role services will be installed on APP1 in this Quick Start, APP1 cannot run in Hyper-V. You can run all Remote Desktop Services roles on Hyper-V virtual machines except for the RD Virtualization Host role service.

Step 3: Test Remote Desktop Services connectivity CLIENT1 is a computer running Windows 8 that is used to test that you can successfully connect to the VDI Quick Start deployment virtual desktop collection.

Step 4: Take a snapshot of the configuration At the completion of the lab, a snapshot of the configuration will enable you to later return to a working VDI Quick Start deployment test lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks that are described in this guide. If you cannot complete a task when you are logged on with an account that is a member of the Administrators group, try performing the task when you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the base configuration test lab

This test lab guide uses the base configuration test lab as a starting place. Before you proceed with the remainder of the steps in this guide, complete the steps in the Test Lab Guide: Base Configuration. If you have already completed the steps in Test Lab Guide: Base Configuration, and you saved a disk image or a virtual machine snapshot of the base configuration, you can restore the base configuration in Hyper-V Manager, and then proceed to the next step.

Step 2: Configure APP1

The APP1 configuration for the VDI Quick Start deployment test lab consists of deploying VDI Quick Start.

To deploy the VDI Quick Start

1. Log on to the APP1 server with the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Before you begin** page of the **Add Roles and Features Wizard**, click **Next**.
5. On the **Select installation type** page, click **Remote Desktop Services installation**, and then click **Next**.
6. On the **Select deployment type** page, click **Quick Start**, and then click **Next**.
7. On the **Select deployment scenario** page, click **Virtual machine-based desktop deployment**, and then click **Next**.
8. On the **Select a server** page, accept the default selections, and then click **Next**.
9. On the **Select a virtual desktop template** page, in the **Virtual hard disk location** box, click **Browse**.

   Navigate to the location of the virtual desktop template, click **Open**, and then click **Next**.
10. On the **Confirm selections** page, select the **Restart the destination server automatically if required** check box, and then click **Deploy**.

    Note

    The computer will restart after the RD Virtualization Host role service is installed. After the restart is complete and you log on as CORP\Administrator, VDI Quick Start will finish the installation.

Step 3: Test Remote Desktop Services connectivity

You can ensure that VDI Quick Start deployed successfully by connecting to the virtual desktop collection that was created during installation. You log on to the RD Web Access server on the APP1 computer and connect to the QuickVMCollection virtual desktop collection.

Note

If you plan to connect to an RD Web Access server website from a server, you must turn off Internet Explorer Enhanced Security Configuration by using Server Manager.

To connect to the virtual desktop collection

1. Log on to the CLIENT1 computer with the CORP\User1 user account.
2. Open Internet Explorer. (To open Internet Explorer, click **Start**, and then click **Internet Explorer.**)
3. In the Internet Explorer address bar, type **https://app1.corp.contoso.com/RDWeb**, and then press ENTER.
4. Click **Continue to this website (not recommended)**.
5. In the **Domain\user name** box, type **CORP\Administrator**.
6. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
7. Click **QuickVMCollection**, and then click **Connect**.

Step 4: Take a snapshot of the configuration

This completes the VDI Quick Start deployment. To save this configuration so that you can quickly return to a working VDI Quick Start deployment configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then turn them off.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots. If your lab uses physical computers, create disk images to save the VDI Quick Start deployment test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

You Are Here -> Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

A Virtual Desktop Infrastructure (VDI) standard deployment enables you to install the appropriate role services on separate computers. Unlike the VDI Quick Start deployment, a standard deployment provides more precise control of virtual desktops and virtual desktop collections by not creating them automatically.

In this guide

This guide contains instructions for setting up a test lab which deploys VDI by using four servers. For detailed information about how the lab is set up, see Test Lab Guide: Base Configuration. The resulting test lab demonstrates how to successfully install and configure a VDI standard deployment.

Important
The following instructions explain how to configure a VDI standard deployment test lab that uses a minimum number of computers. Individual computers are needed to separate the services that are provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, the VDI standard deployment is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named RDVH1 that is configured as the Remote Desktop Virtualization Host (RD Virtualization Host) server.
- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) server.

The VDI standard deployment test lab consists of one subnet (10.0.0.0/24) that simulates an intranet named Corpnet.

Computers are connected to each other by using a Hyper-V virtual switch as shown the following figure.

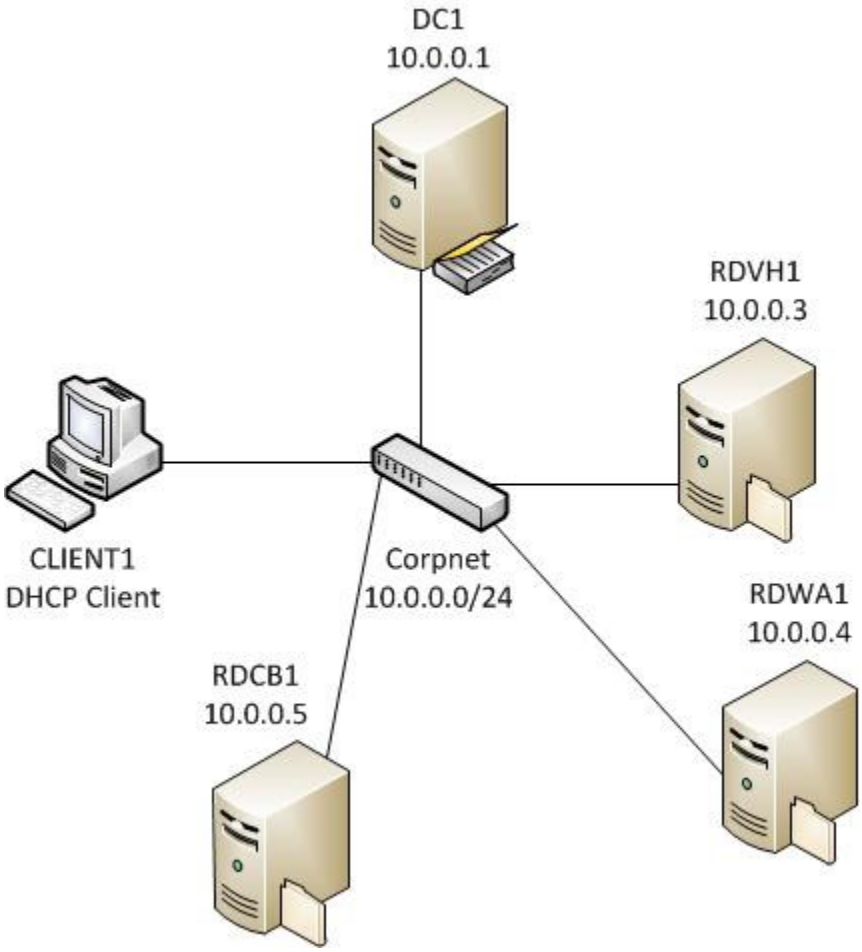

Hardware and software requirements

The following are required components of the test lab:

- You must have a network shared folder that is configured to store user profile disks. All RD Virtualization Host computer accounts must have the ability to read and write to this location.
- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- A physical computer that meets the minimum hardware requirements for Windows Server 2012.

Note

If you choose the VDI scenario-based installation, a virtual switch named RDS Virtual is automatically created.

Steps for configuring the test lab

There are seven steps to follow when you set up the VDI standard deployment test lab based on the Test Lab Guide: Base Configuration

- **Step 1: Complete the Base Configuration Test Lab** The VDI standard deployment test lab requires the Base Configuration test lab as its starting point.
- **Step 2: Install and configure RDVH1** RDVH1 is a computer running Windows Server 2012 that is used as the RD Virtualization Host computer.
- **Step 3: Install and configure RDWA1** RDWA1 is a computer running Windows Server 2012 that is used as the RD Web Access computer.
- **Step 4: Install and configure RDCB1** RDCB1 is a computer running Windows Server 2012 that is used as the RD Connection Broker computer.
- **Step 5: Deploy the Virtual Desktop Infrastructure standard deployment** The Virtual Desktop Infrastructure standard deployment is deployed by using Server Manager.
- **Step 6: Test Remote Desktop Services connectivity** RDCB1 is used to test that the VDI standard deployment completed successfully.
- **Step 7: Take a snapshot of the configuration** At the completion of the lab, take a snapshot of the configuration so that you can later return to a working VDI standard deployment test lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the base configuration test lab

This test lab guide uses the base configuration test lab as a starting place. Before you proceed with the remainder of the steps in this guide, complete the steps in the Test Lab Guide: Base Configuration. If you have already completed the steps in the Test Lab Guide: Base Configuration and saved a disk image or a virtual machine snapshot of the base configuration, you can restore the base configuration and proceed to the next step.

Step 2: Install and Configure RDVH1

Installing and configuring the server used for the RD Session Host consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDVH1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDVH1

First, install Windows Server 2012.

To install the operating system on RDVH1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDVH1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.3**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDVH1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.3 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDVH1 to the CORP domain

To join RDVH1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDVH1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDVH1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 3: Install and Configure RDWA1

Installing and configuring the server used for RD Web Access in the Session Virtualization standard deployment test lab consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDWA1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDWA1

First, install Windows Server 2012.

To install the operating system on RDWA1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDWA1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.4**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDWA1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.4 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDWA1 to the CORP domain

To join RDWA1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDWA1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDWA1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 4: Install and Configure RDCB1

Installing and configuring the server used for the RD Connection Broker consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDCB1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDCB1

First, install Windows Server 2012.

To install the operating system on RDCB1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDCB1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.5**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDCB1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.5 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDCB1 to the CORP domain

To join RDCB1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDCB1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDCB1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 5: Deploy the Virtual Desktop Infrastructure standard deployment

Deploying the required role services for the VDI standard deployment consists of the following procedures:

- Add RDVH1, RDWA1, and RDCB1 to the server pool
- Deploy the VDI standard deployment

To add RDVH1, RDWA1, and RDCB1 to the server pool

1. Log on to the RDCB1 server by using the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add other servers to manage**.

4. In the **Name (CN):** box, type **RDWA1**, and then click **Find Now**.
5. Click **RDWA1**, and then click the right arrow.
6. Repeat steps 4- 5 for the RDVH1 server.
7. Click **OK**.

To deploy the VDI standard deployment

1. Log on to the RDCB1 server by using the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
5. On the **Select Installation Type** page, click **Remote Desktop Services installation**, and then click **Next**.
6. On the **Select deployment type** page, click **Standard deployment**, and then click **Next**.
7. On the **Select deployment scenario** page, click **Virtual machine-based desktop deployment**, and then click **Next**.
8. On the **Review role services** page, click **Next**.
9. On the **Specify RD Connection Broker server** page, click **RDCB1.corp.contoso.com**, click the right arrow, and then click **Next**.
10. On the **Specify RD Web Access server** page, click **RDWA1.corp.contoso.com**, click the right arrow, and then click **Next**.
11. On the **Specify RD Virtualization Host server** page, click **RDVH1.corp.contoso.com**, click the right arrow, and then click **Next**.
12. On the **Confirm selections** page, select the **Restart the destination server automatically if required** check box, and then click **Deploy**.
13. After the installation is complete, click **Close**.

Step 6: Test the VDI standard deployment connectivity

You can ensure that VDI standard deployment deployed successfully by using Server Manager to check the Remote Desktop Services deployment overview.

Use the following procedure to connect to the VDI standard deployment.

To check the deployment overview

1. Log on to the RDCB1 server by using the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Remote Desktop Services**, and then click **Overview**.
4. In the **DEPLOYMENT OVERVIEW** section, ensure that the RD Web Access, RD Connection Broker, and RD Virtualization Host role services are installed. If there is an icon and not a green plus sign (+) next to the role service name, the role service is

installed and part of the deployment.

Step 7: Take a snapshot of the configuration

This completes the VDI standard deployment. To save this configuration so that you can quickly return to a working VDI standard deployment configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then turn them off.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots. If your lab uses physical computers, create disk images to save the VDI standard deployment test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

You Are Here -> Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Managed Pooled Virtual Desktop Collections

In Remote Desktop Services, there are two types of virtual desktop collections available: personal and pooled. You have the option to let Remote Desktop Services automatically manage pooled virtual desktops in a collection, or you can manually manage them. This topic addresses automatically managed pooled virtual desktop collections.

A managed pooled virtual desktop collection offers the following capabilities:

- Automatically create pooled virtual desktops based on a virtual desktop template.
- Automatically install security updates and applications based on a virtual desktop template.
- Live migration with local caching.

- User profile disk support. A user profile disk stores user profile information in a separate virtual hard disk so that user profile settings are persistent across pooled virtual desktops.

This test lab walks you through the process of creating a managed pooled virtual desktop collection, which completes the following tasks:

- Creates a network share that is used to store the user profile disks.
- Adds the virtual desktop template as a Hyper-V virtual machine on the RDVH1.
- Creates a managed pooled virtual desktop collection by using Server Manager.

In this guide

This topic contains instructions for setting up a test lab that is based on information in Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment. You will deploy a managed pooled virtual desktop collection that uses four servers and one client computer. The resulting test lab represents a successfully created managed pooled virtual desktop collection.

Important

The following instructions are for configuring a managed pooled virtual desktop collection test lab that uses the minimum number of computers. Individual computers are needed to separate the services that are provided on the network and to clearly show the desired functionality. This configuration is not designed to reflect best practices or to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed to work only on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named RDVH1 that is configured as the Remote Desktop Virtualization Host (RD Virtualization Host) server.
- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) server.
- One client computer running Windows 8 named CLIENT1 that is used to connect to the virtual desktop collection.

In addition, the managed pooled virtual desktop collection test lab includes:

One subnet (10.0.0.0/24) that simulates an intranet named Corpnet.

Computers connected to each other by using a Hyper-V virtual switch as shown in the following figure.

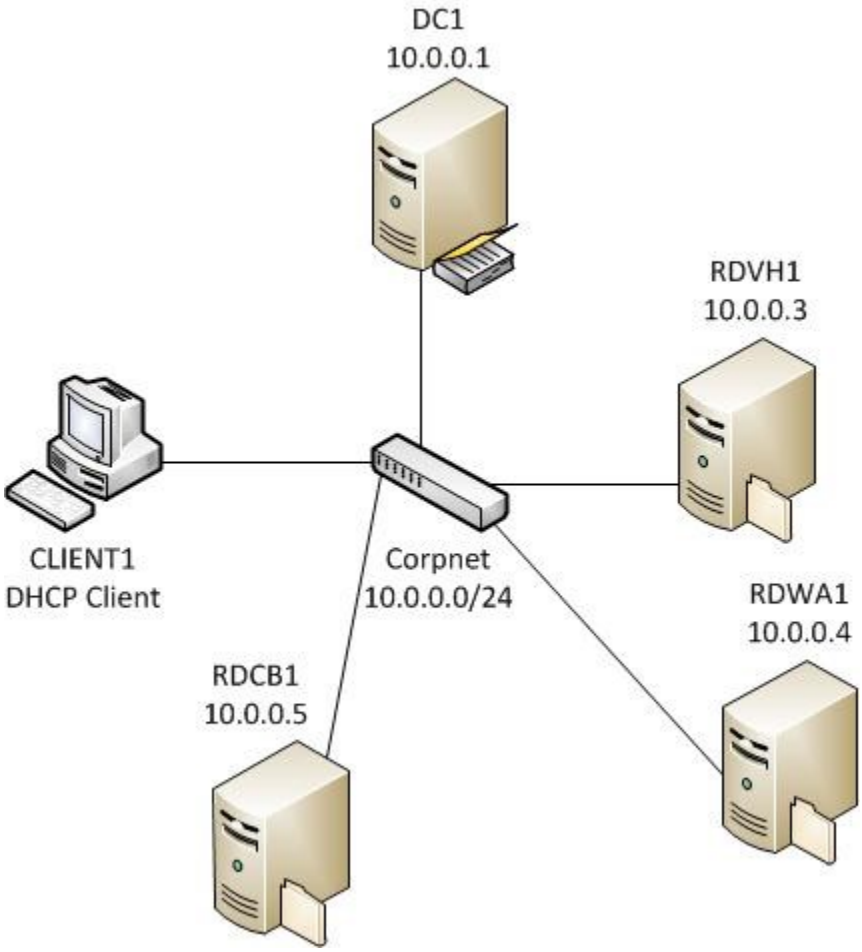

Hardware and software requirements

The following are required components of the test lab:

- A network share must exist that is used to store the user profile disks. The RD Virtualization Host computer accounts must have write access to this location.
- When creating the virtual desktop collection, you must be logged on to the computer with a user account with access to join computers to the domain. This is required so the pooled virtual desktops in the collection are added to the domain.
- The virtual desktop template must be added as a Hyper-V virtual machine. The virtual machine must be generalized by using Sysprep and shut down.

  The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- A physical computer that meets the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are six steps to follow when you set up a managed pooled virtual desktop collection test lab based on the Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment.

Step 1: Complete the VDI standard deployment test lab The managed pooled virtual desktop collection test lab requires the VDI standard deployment test lab as its starting point.

Step 2: Configure RDCB1 RDCB1 stores the user profile disks that are created for the managed pooled virtual desktop collection.

Step 3: Configure RDVH1 The virtual desktop template is added as a Hyper-V virtual machine on RDVH1.

Step 4: Create the managed pooled virtual desktop collection The managed pooled virtual desktop collection is created by running the Create Collection Wizard in Server Manager on RDCB1.

Step 5: Test Remote Desktop Services connectivity CLIENT1 is a computer running Windows 8, which is used to test that you can successfully connect to the virtual desktop collection.

Step 6: Take a snapshot of the configuration At the completion of the lab, take a snapshot of the configuration so that you can later return to a working managed pooled virtual desktop collection test lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks that are described in this guide. If you cannot complete a task when you are logged on with an account that is a member of the Administrators group, try performing the task when you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the VDI standard deployment test lab

This test lab guide uses the VDI standard deployment test lab as a starting place. Complete the steps in Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment before you proceed with the remainder of the steps in this guide. If you have already completed the steps in the Test Lab Guide: Virtual Desktop Infrastructure standard deployment, and you saved a disk image or a virtual machine snapshot of the VDI standard deployment, you can restore the VDI standard deployment in Hyper-V Manager, and then proceed to the next step.

Step 2: Configure RDCB1

You must create a network share on a computer in the CORP domain to store the user profile disks. In this test lab, we will create the network share on the RDCB1 computer.

Use the following procedures to connect to the virtual desktop collection:

- Create the user profile disk network share
- Adjust permissions on the network share

To create the user profile disk network share

1. Log on to the RDCB1 computer by using the CORP\Administrator user account.
2. Open Windows Explorer.
3. Click **Computer**, and then double-click **Local Disk (C:)**.
4. Click **Home**, click **New Folder**, type **UPD** and then press ENTER.
5. Right-click the **UPD** folder, and then click **Properties**.
6. Click **Sharing**, and then click **Advanced Sharing**.
7. Select the **Share this folder** check box.
8. Click **Permissions**, and then grant Full Control permissions to the Everyone group.
9. Click **OK** twice, and then click **Close**.

To adjust permissions on the network share

1. Right-click the **UPD** folder, and then click **Properties**.
2. Click **Security**, and then click **Edit**.
3. Click **Add**.
4. Click **Object Types**, select the **Computers** check box, and then click **OK**.
5. In the **Enter the object names to select** box, type **rdvh1.corp.contoso.com**, and then click **OK**.

   Important
   Each RD Virtualization Host computer account in your deployment must have Modify permissions on the user profile disk network share.
6. Click **RDVH1**, and then select the **Allow** check box next to **Modify**.
7. Click **OK** two times.

Step 3: Configure RDVH1

You must add the virtual desktop template to Hyper-V so you can assign it to the pooled virtual desktop collection.

To add the virtual desktop template

1. Log on to the RDVH1 computer as a CORP\Administrator user account.
2. Click **Start**, and then click **Hyper-V Manager**.
3. Right-click **RDVH1**, point to **New**, and then click **Virtual Machine**.
4. On the **Before You Begin** page, click **Next**.
5. On the **Specify Name and Location** page, in the **Name** box, type **Virtual Desktop Template**, and then click **Next**.
6. On the **Assign Memory** page, in the **Startup memory** box, type **1024**, and then click **Next**.

7. On the **Configure Networking** page, in the **Connection** box, click **RDS Virtual**, and then click **Next**.
8. On the **Connect Virtual Hard Disk** page, click the **Use an existing virtual hard disk** option.
9. Click **Browse**, navigate to the virtual hard disk that should be used as the virtual desktop template, and then click **Open**. Click **Next**.
10. On the **Summary** page, click **Finish**.

Step 4: Create the managed pooled virtual desktop collection

Create the managed pooled virtual desktop collection so that users can connect to desktops in the collection.

To create the managed pooled virtual desktop collection

1. Log on to the RDCB1 server as a CORP\Administrator user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the left pane, click **Remote Desktop Services**, and then click **Collections**.
4. Click **Tasks**, and then click **Create Virtual Desktop Collection**.
5. On the **Before you begin** page, click **Next**.
6. On the **Name the collection** page, in the **Name** box, type **Contoso Managed Pool**, and then click **Next**.
7. On the **Specify the collection type** page, click the **Pooled virtual desktop collection** option, ensure that the **Automatically create and manage virtual desktops** check box is selected, and then click **Next**.
8. On the **Specify the virtual desktop template** page, click **Virtual Desktop Template**, and then click **Next**.
9. On the **Specify the virtual desktop settings** page, click **Provide unattended settings**, and then click **Next**. In this step of the wizard, you can also choose to provide an answer file. For more information, see Walkthrough: Build a Simple Answer File on the Microsoft TechNet website.
10. On the **Specify the unattended settings** page, enter the following information and retain the default settings for the options that are not specified, and then click **Next**.
    - In the **Local Administrator account password** and **Confirm password** boxes, type the same strong password.
    - In the **Time zone** box, click the time zone that is appropriate for your location.
11. On the **Specify users and collection size** page, accept the default selections, and then click **Next**.
12. On the **Specify virtual desktop allocation** page, accept the default selections, and then click **Next**.
13. On the **Specify virtual desktop storage** page, accept the default selections, and then click **Next**.

14. On the **Specify user profile disks** page, in the **Location user profile disks** box, type **\\rdcb1\upd**, and then click **Next**. Make sure that the RD Virtualization Host computer accounts have read and write access to this location.
15. On the **Confirm selections** page, click **Create**.

    Note
    This process will take some time to complete.

Step 5: Test Remote Desktop Services connectivity

You can ensure the managed pooled virtual desktop collection was created successfully by connecting to the RD Web Access server and then connecting to the virtual desktop in the Contoso Managed Pool collection.

Note
If you plan to connect to an RD Web Access server website from a server, you must turn off Internet Explorer Enhanced Security Configuration by using Server Manager.

To connect to the virtual desktop collection

1. Open Internet Explorer. (To open Internet Explorer, click **Start**, and then click **Internet Explorer.**)
2. In the Internet Explorer address bar, type **https://rdwa1.corp.contoso.com/RDWeb**, and then press ENTER.
3. Click **Continue to this website (not recommended)**.
4. In the **Domain\user name** box, type **CORP\Administrator**.
5. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
6. Click **Contoso Managed Pool**, and then click **Connect**.

Step 6: Take a snapshot of the configuration

This completes the managed pooled virtual desktop collection test lab guide. To save this configuration so that you can quickly return to a working managed pooled virtual desktop collection configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then turn them off.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots. If your lab uses physical computers, create disk images to save the managed pooled virtual desktop collection test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

You Are Here -> Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

In Remote Desktop Services, there are two types of virtual desktop collections available: personal and pooled. You have the option to let Remote Desktop Services automatically manage virtual desktops in a virtual desktop collection, or you can manually manage them.

An unmanaged pooled virtual desktop collection offers the following capability:

- User profile disk support. A user profile disk stores user profile information in a separate virtual hard disk so that user profile settings are persistent across pooled virtual desktops.

This test lab walks you through the process of creating an unmanaged pooled virtual desktop collection, which completes the following tasks:

- Creates a network share used to store the user profile disks.
- Creates an unmanaged pooled virtual desktop collection by using Server Manager.

In this guide

This guide contains instructions for setting up a test lab that is based on information in Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment. You will deploy an unmanaged pooled virtual desktop collection that uses four servers and one client computer. The resulting lab demonstrates how to successfully create an unmanaged pooled virtual desktop collection.

Important

The following instructions are for configuring an unmanaged pooled virtual desktop collection test lab that uses the minimum number of computers. Individual computers are needed to separate the services that are provided on the network and to clearly show the

desired functionality. This configuration is not designed to reflect best practices or to reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named RDVH1 that is configured as the Remote Desktop Virtualization Host (RD Virtualization Host) server.
- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) server.
- One client computer running Windows 8 named CLIENT1 that is used to connect to the virtual desktop collection.

In addition, the unmanaged pooled virtual desktop collection test lab consists of one subnet (10.0.0.0/24) that simulates an intranet named Corpnet.

Computers are connected to each other by using a Hyper-V virtual switch as shown in the following figure.

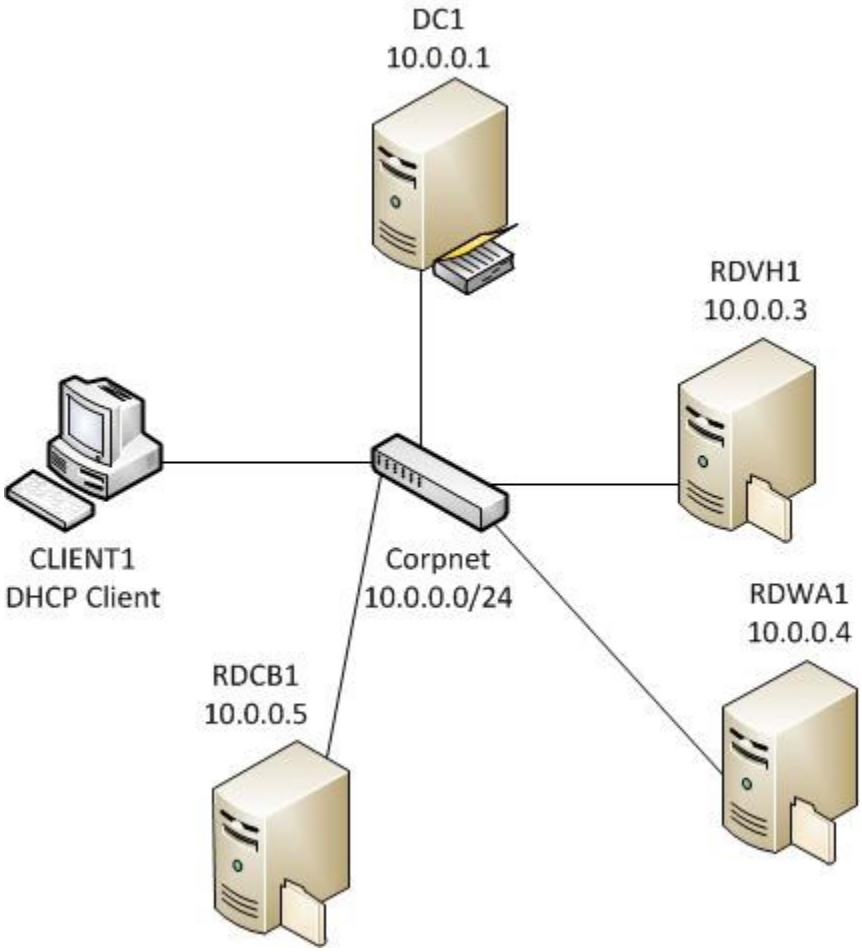

Hardware and software requirements

The following are required components of the test lab:

- Two identical client virtual machines must exist. These virtual machines will become part of the unmanaged pooled virtual desktop collection.
- A network share must exist that is used to store the user profile disks. The RD Virtualization Host computer accounts must have Write access to this location.
- The product disc or files for Windows 8.
- A physical computer that meets the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are four steps to follow when setting up an unmanaged pooled virtual desktop collection test lab based on the Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment.

- **Step 1: Complete the VDI standard deployment test lab** The unmanaged pooled virtual desktop collection test lab requires the VDI standard deployment test lab as its starting point.
- **Step 2: Configure RDCB1** RDCB1 stores the user profile disks that are created for the unmanaged pooled virtual desktop collection.
- **Step 3: Create the unmanaged pooled virtual desktop collection** The unmanaged pooled virtual desktop collection is created by running the Server Manager Create Collection Wizard on RDCB1.
- **Step 4: Test Remote Desktop Services connectivity** CLIENT1 is a Windows 8 computer that is used to test that you can successfully connect to the virtual desktop collection.
- **Step 5: Take a snapshot of the configuration** At the completion of the lab, take a snapshot of the configuration so that you can later return to a working unmanaged pooled virtual desktop collection test lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks that are described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the VDI standard deployment test lab

This test lab guide uses the VDI standard deployment test lab as a starting place. Complete the steps in Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment before you proceed with the remainder of the steps in this guide. If you have already completed the steps in the Test Lab Guide: Virtual Desktop Infrastructure standard Deployment, and you saved a disk image or a virtual machine snapshot of the VDI standard deployment, you can restore the VDI standard deployment in Hyper-V Manager, and then proceed to the next step.

Step 2: Configure RDCB1

You must create a network share on a computer in the CORP domain to store the user profile disks. In this test lab, we will create the network share on the RDCB1 computer.

Use the following procedures to connect to the virtual desktop collection:

- Create the user profile disk network share
- Adjust permissions on the network share

To create the user profile disk network share

1. Log on to the RDCB1 computer by using the CORP\Administrator user account.
2. Open Windows Explorer.
3. Click **Computer**, and then double-click **Local Disk (C:)**.
4. Click **New Folder**, type **UPD** and then press ENTER.
5. Right-click the **UPD** folder, and then click **Properties**.

6. Click **Sharing**, and then click **Advanced Sharing**.
7. Select the **Share this folder** check box.
8. Click **Permissions**, and then grant Full Control permissions to the Everyone group.
9. Click **OK** twice, and then click **Close**.

To adjust permissions on the network share

1. Right-click the **UPD** folder, and then click **Properties**.
2. Click **Security**, and then click **Edit**.
3. Click **Add**.
4. Click **Object Types**, select the **Computers** check box, and then click **OK**.
5. In the **Enter the object names to select** box, type **rdvh1.corp.contoso.com** and then click **OK**.

   Important
   Each RD Virtualization Host computer account in your deployment must have Modify permissions on the user profile disk network share.
6. Click **RDVH1**, and then select the **Allow** check box next to **Modify**.
7. Click **OK** twice.

Step 3: Create the unmanaged pooled virtual desktop collection

Use the following procedure to create an unmanaged pooled virtual desktop collection:

To create the unmanaged pooled virtual desktop collection

1. Log on to the RDCB1 server by using the CORP\Administrator user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the left pane, click **Remote Desktop Services**.
4. Click **Collections**.
5. Click **Tasks**, and then click **Create Virtual Desktop Collection**.
6. On the **Before you begin** page, click **Next**.
7. On the **Name the collection** page, in the **Name** box, type **Contoso Unmanaged Pool**, and then click **Next**.
8. On the **Specify the collection Type** page, click the **Pooled virtual desktop collection** option, clear the **Automatically create and manage virtual desktops** check box, and then click **Next**.
9. On the **Specify the existing virtual desktops** page, click the name of the virtual desktop, and then click **Add**.

   Note
   If the virtual desktop does not show up, you should close this wizard, add the

virtual machines by using Hyper-V Manager, and then create the collection again.

10. Click the name of another virtual desktop, click **Add**, and then click **Next**.
11. On the **Specify user groups** page, accept the default selections, and then click **Next**.
12. On the **Specify user profile disks** page, in the **Location of user profile disks** box, type **\\rdcb1\upd**, and then click **Next**. Make sure that the RD Virtualization Host computer accounts have Read and Write access to this location.
13. On the **Confirm selections** page, click **Create**.

Note
This procedure will take a while to complete.

Step 4: Test Remote Desktop Services connectivity

You can ensure that the unmanaged pooled virtual desktop collection was created successfully by connecting to the RD Web Access server on the CLIENT1 computer, and then connect to a pooled virtual desktop in the Contoso Unmanaged Pool virtual desktop collection.

Note
If you plan to connect to an RD Web Access server website from a server, you must use Server Manager to turn off Internet Explorer Enhanced Security Configuration.

Use the following procedure to connect to the virtual desktop collection.

To connect to the virtual desktop collection

1. Open Internet Explorer. To open Internet Explorer, click **Start**, and then click **Internet Explorer.**
2. In the Internet Explorer address bar, type **https://rdwa1.corp.contoso.com/RDWeb**, and then press Enter.
3. Click **Continue to this website (not recommended)**.
4. In the **Domain\user name** box, type **CORP\Administrator**.
5. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
6. Click **Contoso Unmanaged Pool**, and then click **Connect**.

Step 5: Take a snapshot of the configuration

This completes the unmanaged pooled virtual desktop collection test lab guide. To save this configuration so that you can quickly return to a working unmanaged pooled virtual desktop collection configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then turn them off.

2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots. If your lab uses physical computers, create disk images to save the unmanaged pooled virtual desktop collection test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

You Are Here -> Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Introduction

In Windows Server 2012 the Remote Desktop Services Session Virtualization standard deployment is a scenario-based installation within Server Manager that allows you to install, configure, and manage RD Session Host servers from a central location. With the Session Virtualization standard deployment scenario, you're presented with two deployment types:

- **Standard Deployment:** Allows you to flexibly deploy the various Remote Desktop Services role services to different servers.
- **Quick Start:** Installs all the necessary Remote Desktop Services' role services on one computer to let you install and configure in a test environment. . To use the standard deployment instead of the Quick Start, see: Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment.

Session Virtualization is a scenario-based installation within Server Manager that allows you to install, configure, and manage RD Session Host servers from a central location.

This test lab walks you through the process of creating a **Session Virtualization Quick Start** by doing the following:

- Installs the RD Session Host, RD Connection Broker, and RD Web Access role services on a single computer.
- Creates a session collection.
- Publishes a session-based desktop for each RD Session Host server in the collection.

- Publishes applications as RemoteApp programs.

In this guide

This topic contains instructions for setting up a test lab based on the Test Lab Guide: Base Configuration and deploying Remote Desktop Services Session Virtualization Quick Start deployment using two server computers and one client computers. The resulting Session Virtualization Quick Start deployment test lab demonstrates how to successfully deploy a Session Virtualization Quick Start deployment.

Important
The following instructions are for configuring a Session Virtualization Quick Start deployment test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services Session Virtualization is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named APP1 that will be configured as the Session Virtualization Quick Start deployment server.
- One roaming member client computer running Windows 8 named CLIENT1 that is configured to the Session Virtualization Quick Start deployment.

The Remote Desktop Services Session Virtualization test lab consists of one subnet that simulates an intranet named Corpnet (10.0.0.0/24).

Computers on each subnet connect using a Hyper-V virtual switch. See the following figure.

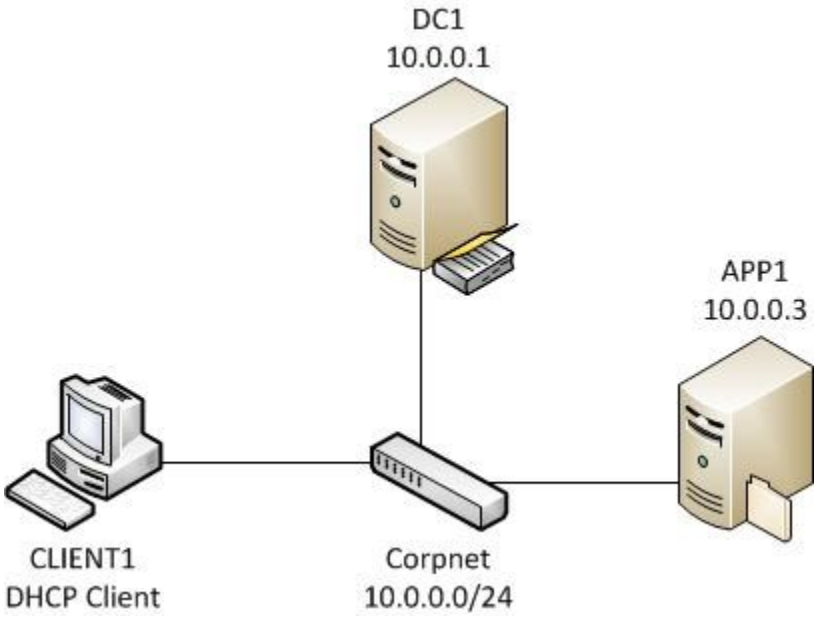

Hardware and software requirements

The following are required components of the test lab:

- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- A physical computer that meet the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are 4 steps to follow when setting up Session Virtualization Quick Start deployment test lab based on the Base Configuration Test Lab Guide.

Note
You must be logged on as a member of the Domain Admins group or as a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

- **Step 1: Complete the Base Configuration Test Lab** - The Session Virtualization Quick Start deployment test lab requires the Base Configuration test lab as its starting point.
- **Step 2: Configure APP1.** APP1 is a Windows Server 2012 computer that will host the Session Virtualization Quick Start.

- **Step 3: Test Remote Desktop Services Connectivity.** CLIENT1 is a Windows 8 computer that is used to test that you can successfully connect to the Session Virtualization Quick Start deployment session collection.
- **Step 4: Snapshot the configuration.** At the completion of the lab, snapshot the configuration so that you can later return to a working Session Virtualization Quick Start deployment Test Lab.

Step 1: Complete the Base Configuration Test Lab

This Test Lab Guide uses the Base Configuration network as a starting place. Please complete the steps in Test Lab Guide: Base Configuration before proceeding with the remainder of the steps in this guide. If you have already completed the steps in the Base Configuration Test Lab Guide and saved a disk image or a virtual machine snapshot of the Base Configuration, then you can restore the Base Configuration and proceed to the next step.

Step 2: Configure APP1

APP1 configuration for the Session Virtualization Quick Start deployment test lab consists of the following procedures:

- Deploy Session Virtualization Quick Start

The following sections explain these procedures in detail.

Deploy the Session Virtualization Quick Start

The Session Virtualization Quick Start is deployed by using Server Manager.

To deploy the Session Virtualization Quick Start

1. Log on to the APP1 server as the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Before You Begin** page of the Add Roles and Features Wizard, click **Next**.
5. On the **Select Installation Type** page, click **Remote Desktop Services installation**, and then click **Next**.
6. On the **Select deployment type** page, click **Quick Start**, and then click **Next**.
7. On the **Select deployment scenario** page, click **Session-based desktop deployment**, and then click **Next**.
8. On the **Server Selection** page, accept the default selections, and then click **Next**.
9. On the **Confirm selections** page, select the **Restart the destination server automatically if required** check box, and then click **Deploy**.
10. Click **Close**.

Note

The computer will restart after the RD Session Host role service is installed. After the restart is complete and you log on as CORP\User1, the Session Virtualization Quick Start deployment will finish the installation.

Step 3: Test Remote Desktop Service Connectivity

Verify that the Session Virtualization Quick Start deployment deployed successfully by connecting to the session collection that was created during installation.

Note

If you plan on connecting to an RD Web Access server website from a server, you must turn off Internet Explorer Enhanced Security Configuration by using Server Manager.

Use the following procedures to connect to the Session Collection:

- Connect to the virtual desktop collection.

The following sections explain these procedures in detail.

Connect to the Session Collection

To test that the Session Virtualization Quick Start deployment deployed successfully, you log on to the RD Web Access server on the APP1 computer and connect to the session collection.

To connect to the session collection

1. Log on to the CLIENT1 computer as the CORP\User1 user account.
2. Open Internet Explorer. Click **Start**, point to **All Programs**, and then click **Internet Explorer**.
3. In the Internet Explorer address bar, type **https://app1.corp.contoso.com/RDWeb** and then press **Enter**.
4. Click **Continue to this website (not recommended).**
5. In the **Domain/user name** box, type **CORP\Administrator**.
6. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
7. Click **No**, and then click **Allow**.

Step 4: Snapshot the Configuration

This completes the Session Virtualization Quick Start deployment test lab. To save this configuration so that you can quickly return to a working Session Virtualization Quick Start deployment configuration from which you can test other modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots Session Virtualization Quick Start deployment. If your lab uses physical

computers, create disk images to save the Session Virtualization Quick Start deployment test lab configuration.

Additional Resources

- What's New in Remote Desktop Services in Windows Server 2012
- Test Lab Guide: Virtual Desktop Infrastructure Quick Start
- Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment
- Test Lab Guide: Managed Pooled Virtual Desktop Collections
- Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections
- **You Are Here ->** Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start
- Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment
- Test Lab Guide: Remote Desktop Services Publishing
- Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

In Windows Server 2012 the Remote Desktop Services Session Virtualization standard deployment is a scenario-based installation within Server Manager that allows you to install, configure, and manage RD Session Host servers from a central location. With the Session Virtualization standard deployment scenario, you're presented with two deployment types:

- **Standard Deployment:** Allows you to flexibly deploy the various Remote Desktop Services role services to different servers.
- **Quick Start:** Installs all the necessary Remote Desktop Services' role services on one computer to let you install and configure in a test environment. To use the Quick Start instead of the standard deployment, see: Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

This test lab walks you through the process of creating a Session Virtualization **standard deployment** by doing the following:

- Installs the RD Connection Broker, RD Session Host, and RD Web Access role services on separate computers.
- Creates a session collection.
- Publishes a session-based desktop for each RD Session Host server in the collection.
- Publishes applications as RemoteApp programs.

A Remote Desktop Services Session Virtualization standard deployment allows you to install the appropriate role services on separate computers. Unlike the Session Virtualization Quick Start deployment, a standard deployment provides more precise control of session-based desktops and session virtualization collections by not creating them automatically.

In this guide

This topic contains instructions for setting up a test lab based on the Test Lab Guide: Base Configuration and deploying Session Virtualization using three server computers and one client computer. The resulting Session Virtualization standard deployment test lab demonstrates how to successfully install and configure a Session Virtualization standard deployment.

Important
The following instructions describe how to configure a Session Virtualization standard deployment test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices, nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, the Session Virtualization standard deployment is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, and Dynamic Host Configuration Protocol (DHCP) server.
- One intranet member server running Windows Server 2012 named RDSH1 that is configured as the Remote Desktop Session Host (RD Session Host) server.
- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) server.
- One roaming member client computer running Windows 8 named CLIENT1 that is used to connect to the Session Virtualization deployment.

The Session Virtualization standard deployment test lab consists of one subnet that simulate the following:

- An intranet, referred to as the Corpnet subnet (10.0.0.0/24).

Computers on the subnet connect using a physical hub, switch, or virtual switch. See the following figure for the configuration of the Session Virtualization standard deployment test lab.

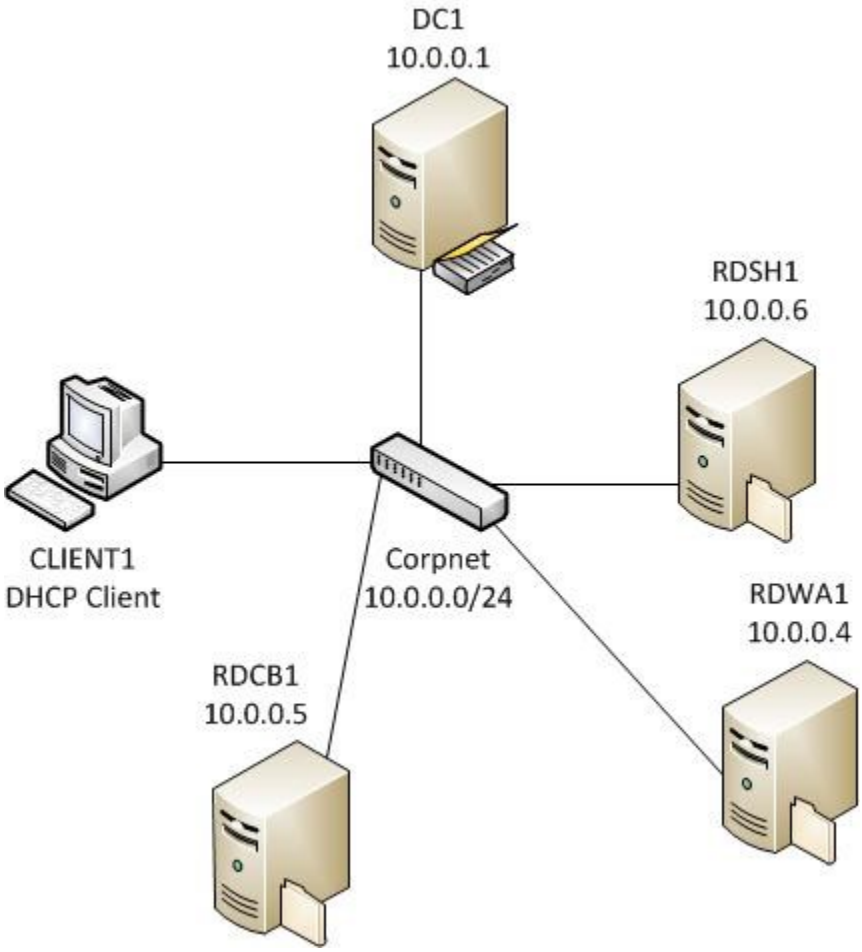

Hardware and software requirements

The following are required components of the test lab:

- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- Three physical computers that meet the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are 8 steps to follow when setting up the Session Virtualization standard deployment test lab based on the Test Lab Guide: Base Configuration.

- **Step 1: Complete the Base Configuration Test Lab** -- The Session Virtualization standard deployment test lab requires the Base Configuration test lab as its starting point.
- **Step 2: Install and Configure RDSH1** – RDSH1 is a Windows Server 2012 computer that is used as the RD Session Host computer.

- **Step 3: Install and Configure RDWA1** – RDWA1 is a Windows Server 2012 computer that is used as the RD Web Access computer.
- **Step 4: Install and Configure RDCB1** – RDCB1 is a Windows Server 2012 computer that is used as the RD Connection Broker computer.
- **Step 5: Deploy Session Virtualization Standard Deployment** – The Session Virtualization standard deployment is deployed by using Server Manager.
- **Step 6: Create Session Collection** – A Session Collection is created by using Server Manager.
- **Step 7: Test Remote Desktop Services Connectivity** -- CLIENT1 is a Windows 8 computer that is used to test that you can successfully connect to the Session Virtualization standard deployment.
- **Step 8: Snapshot the configuration** -- At the completion of the lab, snapshot the configuration so that you can later return to a working Session Virtualization standard deployment Test Lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the base configuration test lab

This test lab guide uses the base configuration test lab as a starting place. Before you proceed with the remainder of the steps in this guide, complete the steps in the Test Lab Guide: Base Configuration. If you have already completed the steps in the Test Lab Guide: Base Configuration and saved a disk image or a virtual machine snapshot of the base configuration, you can restore the base configuration and proceed to the next step.

Step 2: Install and Configure RDSH1

Installing and configuring the server used for the RD Session Host consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDSH1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDSH1

First, install Windows Server 2012.

To install the operating system on RDSH1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDSH1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.6**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDSH1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.6 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDSH1 to the CORP domain

To join RDSH1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDSH1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDSH1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 3: Install and Configure RDWA1

Installing and configuring the server used for RD Web Access in the Session Virtualization standard deployment test lab consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDWA1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDWA1

First, install Windows Server 2012.

To install the operating system on RDWA1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDWA1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.4**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDWA1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.4 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDWA1 to the CORP domain

To join RDWA1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDWA1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDWA1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 4: Install and Configure RDCB1

Installing and configuring the server used for the RD Connection Broker consists of the following procedures:

- Install the operating system.
- Configure TCP/IP properties.
- Join RDCB1 to the CORP domain.

The following sections explain these procedures in detail.

Install the operating system on RDCB1

First, install Windows Server 2012.

To install the operating system on RDCB1

1. Start the installation of Windows Server 2012.
2. Follow the instructions to complete the installation, specifying a strong password for the local Administrator account. Log on using the local Administrator account.
3. Connect RDCB1 to a network that has Internet access and run Windows Update to install the latest updates for Windows Server 2012.

Configure TCP/IP properties

Next, configure TCP/IP.

To configure TCP/IP properties

1. In Server Manager, click **Local Server** in the console tree. Click the link next to **Wired Ethernet Connection** in the Properties tile.
2. In **Network Connections**, right-click **Wired Ethernet Connection**, and then click **Properties**. Note that the "Wired Ethernet Connection" interface name may be different on your computer.
3. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Select **Use the following IP address**. In **IP address**, type **10.0.0.5**. In **Subnet mask**, type **255.255.255.0**.
5. Select **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close**. Close the **Network Connections** window.
7. From the Start screen, type **command**, and then click **Command Prompt**.
8. To check name resolution and network communication between RDCB1 and DC1, type **ping dc1.corp.contoso.com** in the command prompt window and hit **ENTER**.
9. Verify that there are four replies from 10.0.0.1.
10. Close the Command Prompt window.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Use the **ipconfig /all** command to list all the interfaces.

```
New-NetIPAddress -InterfaceAlias "Wired Ethernet Connection" -IPv4Address 10.0.0.5 -
PrefixLength 24

Set-DnsClientServerAddress -InterfaceAlias "Wired Ethernet Connection" -ServerAddresses
10.0.0.1
```

Join RDCB1 to the CORP domain

To join RDCB1 to the CORP domain

1. Log on to the computer with the local Administrator user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Local Server**.
4. Click the name of the computer next to **Computer name**, and then click **Change**.
5. In the **Computer name** box, type **RDCB1**. In the **Member of** box, click the **Domain** option, type **corp.contoso.com** and then click **OK**.
6. When you are prompted for a user name and password, provide the CORP\Administrator credentials, and then click **OK**.
7. When you see a dialog box welcoming you to the corp.contoso.com domain, click **OK**.
8. When you are prompted that you must restart the computer, click **OK**.
9. On the **System Properties** dialog box, click **Close**, and then click **Restart Now**.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.. Note that you must supply domain credentials after entering the Add-Computer command below.

```
Rename-Computer -NewName RDCB1

Restart-Computer

Add-Computer -DomainName corp.contoso.com

Restart-Computer
```

Step 5: Deploy Session Virtualization Standard

Deploying the required role services for the Session Virtualization standard deployment consists of the following procedures:

- Add RDSH1, RDWA1, and RDCB1 to the server pool.
- Deploy Session Virtualization standard deployment.

The following sections explain these procedures in detail.

Add RDSH1, RDWA1, and RDCB1 to the server pool

To add RDSH1, RDWA1, and RDCB1 to the server pool

1. Log on to the RDCB1 server as the CORP\User1 user account.
2. Server Manager will start automatically. If it does not start automatically, click **Start**, type

servermanager.exe, and then click **Server Manager**.

3. In the **QUICK START** section, click **Add other servers to manage**.
4. In the **Name (CN):** box, type **RDWA1** and then click **Find Now**.
5. Click **RDWA1** and then click the right arrow.
6. Repeat steps 4 – 5 for the RDSH1 server.
7. Click **OK**.

Deploy the Session Virtualization Standard deployment

The Session Virtualization standard deployment is deployed by using Server Manager.

To deploy Session Virtualization Standard

1. Log on to the RDCB1 server as the CORP\User1 user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. In the **QUICK START** section, click **Add roles and features**.
4. On the **Before You Begin** page of the **Add Roles and Features Wizard**, click **Next**.
5. On the **Select Installation Type** page, click **Remote Desktop Services installation**, and then click **Next**.
6. On the **Select deployment type** page, click **Standard deployment**, and then click **Next**.
7. On the **Select deployment scenario** page, click **Session based desktop deployment**, and then click **Next**.
8. On the **Review role services** page, click **Next**.
9. On the **Specify RD Connection Broker server** page, click **RDCB1.corp.contoso.com**, click the right arrow, and then click **Next**.
10. On the **Specify RD Web Access server** page, click **RDWA1.corp.contoso.com**, click the right arrow, and then click **Next**.
11. On the **Specify RD Session Host servers** page, click **RDSH1.corp.contoso.com**, click the right arrow, and then click **Next**.
12. On the **Confirm selections** page, select the **Restart the destination server automatically if required** check box, and then click **Deploy**.
13. After the installation is complete, click **Close**.

Step 6: Create Session Collection

A Session Collection is created by using Server Manager.

Use the following procedures to create a Session Collection:

- Create a Session Collection

The following sections explain these procedures in detail.

Create a session collection

To create a session collection

1. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
2. On the left side of the window, click **Remote Desktop Services**.
3. Click **Collections**.
4. Click **Tasks**, and then click **Create Session Collection**.
5. On the **Before You Begin** page of the Create Collection wizard, click **Next**.
6. On the **Name the collection** page, type **SessionCollection** in the **Name** box, and then click **Next**.
7. On the **Specify RD Session Host servers** page, select **RDSH1**, click the right arrow, and then click **Next**.
8. On the **Specify user groups** page, accept the default selections, and then click **Next**.
9. On the **Specify user profile disks** page, clear the **Enable user profile disks** check box, and then click **Next**.
10. On the **Confirm selections** page, click **Create**.
11. After the session collection is created, click **Close**.

Step 7: Test Remote Desktop Services Connectivity

You can ensure Session Virtualization standard deployment deployed successfully by connecting to the session collection that was created.

Note
If you plan on connecting to an RD Web Access server website from a server, you must turn off Internet Explorer Enhanced Security Configuration by using Server Manager.

Use the following procedures to connect to the Session Virtualization deployment:

- Connect to the session collection.

The following sections explain these procedures in detail.

Connect to the session collection

To test that the Session Virtualization standard deployment deployed and the session collection was created successfully, you log on to the RD Web Access server on the RDWA1 computer and connect to a session in the SessionCollection session collection.

To connect to the session collection

1. Log on to the CLIENT1 computer as the CORP\User1 user account.
2. Open Internet Explorer. To open Internet Explorer, click **Start**, and then click **Internet Explorer**.

3. In the Internet Explorer address bar, type **https://rdwa1.corp.contoso.com/RDWeb** and then press Enter.
4. Click **Continue to this website (not recommended)**.
5. In the **Domain/user name** box, type **CORP\Administrator**.
6. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
7. Click **SessionCollection**, and then click **Connect**.
8. Verify that the session-based desktop appears correctly.

Step 8: Snapshot the Configuration

This completes the Session Virtualization standard deployment test lab. To save this configuration so that you can quickly return to a working VDI Quick Start configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots Session Virtualization standard deployment. If your lab uses physical computers, create disk images to save the Session Virtualization standard deployment test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

You Are Here -> Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Remote Desktop Services Publishing

Introduction

In Windows Server 2008 R2, publishing and managing applications on both pooled and personal virtual desktops is a very time-consuming and costly process. Because RemoteApp programs only partially integrate with the native Windows experience, they add to the management cost since there's no way to organize published RemoteApp programs to users.

Remote Desktop Services in Windows Server® 2012 enables you to publish and manage resources, such as RemoteApp programs, session-based desktops, and virtual desktops, from a centralized console. Using this new publishing feature, you get an historic view of resources assigned to end users, you can change published resources for any given collection, and you can edit properties of published resources.

Centralized resource publishing provides end users with an experience that can replace locally installed applications.

In this guide

This paper contains instructions for setting up a test lab based on the Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment and publishing and managing resources using four server computers and one client computer. The resulting centralized resource publishing test lab demonstrates how to successfully install and configure resource publishing.

Important

The following instructions are for configuring the Remote Desktop Services Publishing test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services Publishing is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named RDSH1 that is configured as the Remote Desktop Session Host (RD Session Host) server.

- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) server.
- One roaming member client computer running Windows 8 named CLIENT1 that is used to connect to the Session Virtualization deployment.

The Remote Desktop Services Publishing test lab consists of 1 subnet (10.0.0.0/24) that simulates an intranet named Corpnet.

Computers are connected to each other by using a Hyper-V virtual switch. See the following figure.

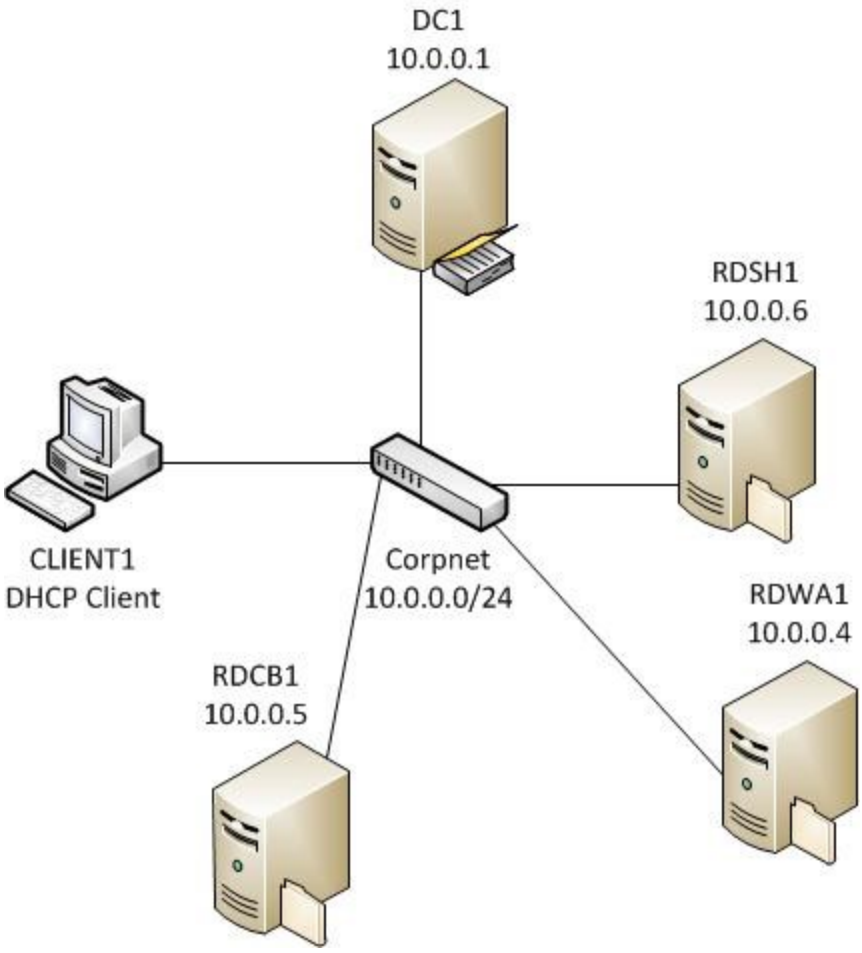

Hardware and software requirements

The following are required components of the test lab:

- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.

- A physical computer that meet the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are 6 steps to follow when setting up the Remote Desktop Services Publishing test lab based on the Test Lab Guide: Demonstrate Remote Desktop Services Session Virtualization Standard Deployment.

- **Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab** -- The Publishing test lab requires the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment test lab as its starting point.
- **Step 2: Publish a RemoteApp Program** – RemoteApp programs are published to either a session collection using Server Manager.
- **Step 3: Enable File Type Association on a RemoteApp Program** – Enable file type association to assign the file association on the client device to the associated RemoteApp program.
- **Step 4: Configure DNS Feed Lookup for RemoteApp and Desktop Connections** – Configure the Remote Desktop Services DNS feed lookup to automatically discover the RemoteApp and Desktop connection URL when users type in their email addresses.
- **Step 5: Test Published RemoteApp Program using RD Web Access** -- CLIENT1 is a Windows 8 computer that is used to test that you can successfully connect to the published RemoteApp program.
- **Step 6: Snapshot the configuration** -- At the completion of the lab, snapshot the configuration so that you can later return to a working Session Virtualization deployment Test Lab.

Note

You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab

This Test Lab Guide uses the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment test lab as a starting place. Please complete the steps in Test Lab Guide: Demonstrate Remote Desktop Services Session Virtualization Standard Deployment before proceeding with the remainder of the steps in this guide. If you have already completed the steps in the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab Guide and saved a disk image or a virtual machine snapshot of the Demonstrate Remote

Desktop Services Session Virtualization Standard Deployment, then you can restore the Configuration and proceed to the next step.

Step 2: Publish a RemoteApp Program

RemoteApp programs are published to either a session collection or virtual desktop collection using Server Manager. Publishing a RemoteApp program in the Session Virtualization deployment test lab consists of the following procedures:

- Publish a RemoteApp program

The following sections explain these procedures in detail.

Publish a RemoteApp program

To publish a RemoteApp program

1. Log on to the RDSH1 server as the CORP\Administrator user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. On the left side of the window, click **Remote Desktop Services**.
4. Under **Collections**, click **SessionCollection**.
5. In the **REMOTEAPP PROGRAMS** tile, click **Tasks**, and then click **Publish RemoteApp Programs**.
6. On the **Select RemoteApp Programs** page, select **Calculator** and **Wordpad**, and then click **Next**.
7. On the **Confirmation** page, click **Publish**.
8. When the RemoteApp program is published, click **Close**.

Step 3: Enable File Type Association on a RemoteApp Program

Installing and configuring the server that will be used for RD Web Access in the Session Virtualization deployment test lab consists of the following procedures:

- Enable the default connection URL using Group Policy
- Configure the file type association

The following sections explain these procedures in detail.

Enable the default connection URL using Group Policy

To enable the default connection URL using Group Policy

1. Log on to the domain controller as the CORP\Administrator user account.
2. On the domain controller, open Group Policy Management. Click **Start**, type **Group Policy Management**, and then click **Group Policy Management**..

3. Expand **Forest: corp.contoso.com**, and expand **Domains**, and then expand **corp.contoso.com**.
4. Right-click **Default Domain Policy**, and then click **Edit**.
5. Navigate to **User Configuration** -> **Policies** ->**Administrative Templates** -> **Windows Components** -> **Remote Desktop Services** -> **RemoteApp and Desktop Connections**.
6. Double-click **Specify Default Connection URL**, and then click **Enabled**.
7. In the **Default Connection URL** box, type **https://RDWA1.corp.contoso.com/RDWeb**.
8. Click **OK**.

Configure the File Type Association

To configure the file type association

1. Log on to the RDSH1 server as the CORP\Administrator user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. On the left side of the window, click **Remote Desktop Services**.
4. Under **Collections**, click **SessionCollection**.
5. Under the **REMOTEAPP PROGAMS** heading, right-click **WordPad**, and then click **Edit Properties**.
6. Click **File Type Association**.
7. Select the **.docx**, **.odt**, and **.rtf**
8. Click **OK**.

Step 4: Configure DNS Feed Lookup for RemoteApp and Desktop Connections

The Remote Desktop Services DNS feed lookup is used to automatically discover the RemoteApp and Desktop connection URL when users type in their email addresses. This is done by creating a DNS TXT record on the DNS server that contains the connection URL. Configure DNS Feed Lookup for RemoteApp and Desktop Connections to enable automatic discovery of the RemoteApp and Desktop connection URL consists of the following procedures:

- Configure DNS Feed Lookup for RemoteApp and Desktop Connections.

The following sections explain these procedures in detail.

Configure DNS Feed Lookup for RemoteApp and Desktop Connections.

To configure DNS feed lookup for RemoteApp and Desktop Connection

1. Log on to the domain controller as the CORP\Administrator user account.
2. On the domain controller, open DNS Manager. Click **Start**, type **DNS**, and then click

DNS.

3. Right-click the forward lookup zone where you want to create the TXT record, and then click **Other New Records**.
4. Click **Text (TXT)**, and then click **Create Record**.
5. In the **Record name (use parent domain if left blank)** box, type **_msradc**.
6. In the **Text** box, type **https://RDWA1.corp.contoso.com/RDWeb/Feed**.
7. Click **OK**.

Step 5: Test Published RemoteApp Program using RD Web Access

You can ensure the RemoteApp program published successfully by connecting to the WordPad application that was published.

- Test Published RemoteApp Program using RD Web Access

The following sections explain these procedures in detail.

Test Published RemoteApp Program using RD Web Access

To test published RemoteApp program using RD Web Access

1. Log on to the CLIENT1 computer as the CORP\Administrator user account.
2. Open Internet Explorer. To open Internet Explorer, click **Start**, and then click **Internet Explorer.**
3. In the Internet Explorer address bar, type **https://RDWA1.corp.contoso.com/RDWeb** and then press Enter.
4. Click **Continue to this website (not recommended)**.
5. In the **Domain/user name** box, type **CORP\Administrator**.
6. In the **Password** box, type the password for the CORP\Administrator user account, and then click **Sign in**.
7. Click **WordPad**, and then click **Connect**.
8. Verify that the WordPad application appears correctly.

Step 6: Snapshot the Configuration

This completes the Remote Desktop Services Virtual Desktop Infrastructure Quick Start test lab. To save this configuration so that you can quickly return to a working VDI Quick Start configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.

2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots Publish RemoteApp program. If your lab uses physical computers, create disk images to save the Remote Desktop Services Publishing test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

You Are Here -> Test Lab Guide: Remote Desktop Services Publishing

Test Lab Guide: Remote Desktop Services Licensing

Test Lab Guide: Remote Desktop Services Licensing

Introduction

Remote Desktop Licensing (RD Licensing) is a Remote Desktop Services role service that is included with Windows Server® 2012. RD Licensing manages the licenses that are required to connect to a session-based desktop or a virtual desktop. You can use RD Licensing to install, issue, and track the availability and usage of licenses.

In this guide

This paper contains instructions for setting up a test lab based on the Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment or Test Lab Guide: Managed Pooled Virtual Desktop Collections and installing and configuring licensing for Remote Desktop Services using four server computers. The resulting test lab demonstrates how to successfully install and configure licensing for Remote Desktop Services.

Important

The following instructions are for configuring the Remote Desktop Services Licensing test lab using the minimum number of computers. Individual computers are needed to separate the services provided on the network and to clearly show the desired functionality. This configuration is neither designed to reflect best practices nor does it reflect a desired or recommended configuration for a production network. The

configuration, including IP addresses and all other configuration parameters, is designed only to work on a separate test lab network.

Test lab overview

In this test lab, Remote Desktop Services Publishing is deployed with:

- One computer running Windows Server 2012 named DC1 that is configured as an intranet domain controller, Domain Name System (DNS) server, Dynamic Host Configuration Protocol (DHCP) server, and an enterprise root certification authority (CA).
- One intranet member server running Windows Server 2012 named RDVH1 that is configured as the Remote Desktop Virtualization Host (RD Virtualization Host) server.
- One intranet member server running Windows Server 2012 named RDWA1 that is configured as the Remote Desktop Web Access (RD Web Access) server.
- One intranet member server running Windows Server 2012 named RDCB1 that is configured as the Remote Desktop Connection Broker (RD Connection Broker) and the Remote Desktop Services Licensing server.

The Remote Desktop Services Licensing test lab consists of 1 subnet (10.0.0.0/24) that simulates an intranet named Corpnet.

Computers are connected to each other by using a Hyper-V virtual switch. See the following figure.

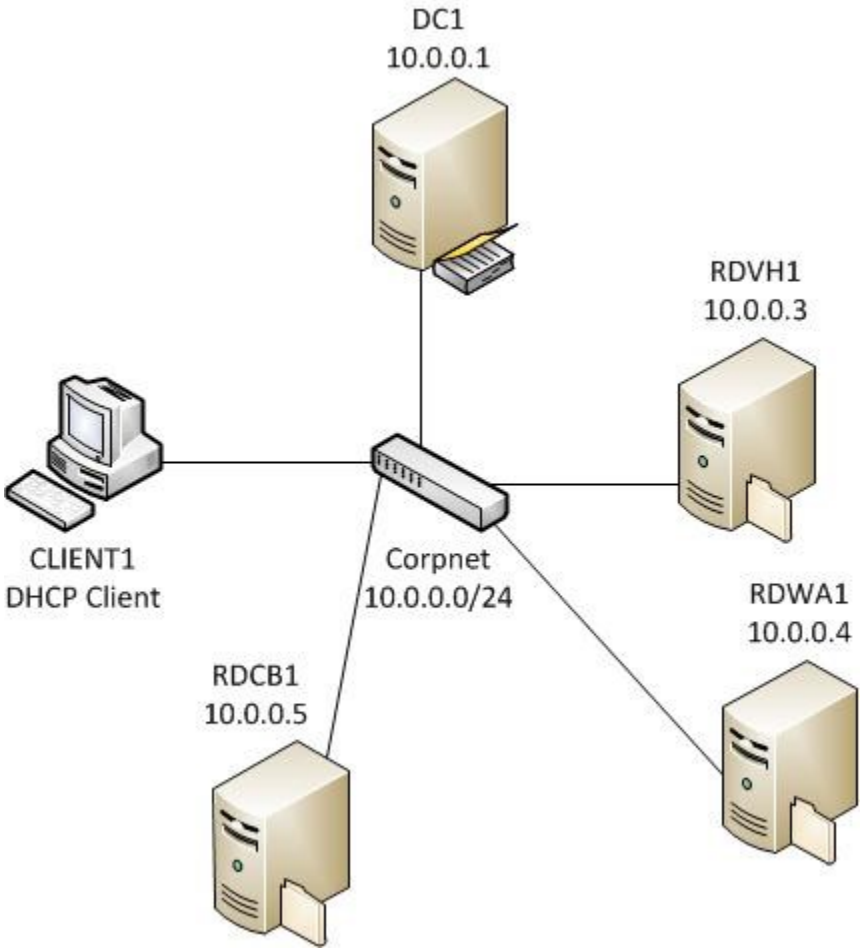

Hardware and software requirements

The following are required components of the test lab:

- You must have a network shared folder that is configured to store user profile disks. All RD Virtualization Host computer accounts must have the ability to read and write to this location.
- The product disc or files for Windows Server 2012.
- The product disc or files for Windows 8.
- A physical computer that meet the minimum hardware requirements for Windows Server 2012.

Steps for configuring the test lab

There are 7 steps to follow when you set up Remote Desktop Services Licensing test lab based on the Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment or Test Lab Guide: Managed Pooled Virtual Desktop Collections.

- **Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab** -- The Publishing test lab requires the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment test lab as its starting point.
- **Step 2: Install the RD Licensing role service** – Install the RD Licensing role service on RDCB1.
- **Step 3: Activate the Remote Desktop license server** – Activate the Remote Desktop license server to certify the license server.
- **Step 4: Request RDS Beta CALs** – Request Beta licenses for the Remote Desktop license server.
- **Step 5: Install RDS CALs on the Remote Desktop license server** -- Install RDS CALs to allow the license server to issue Remote Desktop Services licenses to clients.
- **Step 6: Verify licensing for the Remote Desktop Services** --
- **Step 7: Snapshot the configuration** -- At the completion of the lab, snapshot the configuration so that you can later return to a working Session Virtualization deployment Test Lab.

Note
You must be logged on as a member of the Domain Admins group or a member of the Administrators group on each computer to complete the tasks described in this guide. If you cannot complete a task while you are logged on with an account that is a member of the Administrators group, try performing the task while you are logged on with an account that is a member of the Domain Admins group.

Step 1: Complete the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab

This Test Lab Guide uses the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment test lab as a starting place. Please complete the steps in Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment or Test Lab Guide: Managed Pooled Virtual Desktop Collections before proceeding with the remainder of the steps in this guide. If you have already completed the steps in the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment Test Lab Guide and saved a disk image or a virtual machine snapshot of the Demonstrate Remote Desktop Services Session Virtualization Standard Deployment, then you can restore the Configuration and proceed to the next step.

Step 2: Install the RD Licensing role service

Install the RD Licensing role service on the RDCB1 server.

- Install the RD Licensing role service

The following sections explain these procedures in detail.

Install the RD Licensing role service

To install the RD Licensing role service

1. Log on to the RDCB1 server as the CORP\Administrator user account.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. On the left side of the window, click **Remote Desktop Services**, and then click **Overview**.
4. Under **Deployment Overview**, click the **+** icon above RD Licensing.
5. On the **Select a server** page, click **RDCB1.corp.contoso.com**, click the right arrow, and then click **Next**.
6. On the **Confirm selections** page, verify the RD Licensing role service will be installed, and then click **Add**.
7. On the **View progress** page, confirm the installation of the RD Licensing role service was successful, and then click **Close**.

Step 3: Activate the Remote Desktop License Server

Activate the Remote Desktop license server. A license server must be activated to certify the license server and prepare the license server to issue Remote Desktop Services client access licenses (RDS CALs), or VDI Suite licenses.

- Activate the Remote Desktop License Server

The following sections explain these procedures in detail.

Activate the Remote Desktop license server

To activate a Remote Desktop license server automatically

1. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
2. On the left side of the window, click **Remote Desktop Services**.
3. Click **Servers**.
4. Right-click **RDCB1**, and then click **RD Licensing Manager**.
5. Right-click **RDCB1**, and then click **Activate Server**. The Activate Server Wizard starts.
6. In the **Activate Server Wizard**, on the **Welcome to the Activate Server Wizard** page, click **Next**.
7. On the **Connection Method** page, in the **Connection method:** drop-down list, select **Automatic connection (recommended)**, and then click **Next**.
8. On the **Company Information** page, enter your **First Name:**, **Last Name:**, and **Company:**, select your **Country or Region:** from the drop-down list, and then click **Next**.
9. On the **Company Information** optional information page, specify any other information

that you want to provide, such as **Email:**, and **Company address:**, and then click **Next**.

10. On the **Completing the Activate Server Wizard** page, ensure that the **Start Install Licenses Wizard** now check box is cleared, and then click **Next**.

Step 4: Request RDS Beta CALs

Request RDS CALs over the Internet.

- Request RDS Beta CALs

The following sections explain these procedures in detail.

Request RDS Beta CALs

To request RDS Beta CALs

1. Open Internet Explorer. To open Internet Explorer, click **Start**, and then click **Internet Explorer**.
2. In the **Internet Explorer address bar**, type **https://licensecode.one.microsoft.com/** and then press Enter.
3. On the **Remote Desktop Services License Code Request page**, select the following information and then click **Next**:
   - **Country:**
   - **Source:** Select your beta/evaluation source.
   - **License Type:** You have the choice to select Per User CALs or Per Device CALs.
   - **eMail Address:** Select an email address to receive the RDS CALs.
4. On the **Confirm License Code Request** page, re-enter the **email address** and ensure that the information you provided in the confirm section is correct. Enter the capcha code and then click **Submit**.
5. You will receive an email from RDSLMIC (it may end up in your junk folder) with a **5x5 Beta evaluation license key**.

Step 5: Install RDS CALs on the Remote Desktop License Server

Install RDS CALs on the Remote Desktop license server, to allow the license server to issue Remote Desktop Services client access licenses (RDS CALs).

- Install RDS CALs on the Remote Desktop license server

The following sections explain these procedures in detail.

Install RDS CALs on the Remote Desktop license server

To install Remote Desktop Services client access licenses (RDS CALs) automatically

1. On the left side of the window, click **Remote Desktop Services**, and then click **Servers**, select the server **RDCB1**, and then click **RD Licensing Manager**.

2. Right-click the license server **RDCB1**, and then click **Install Licenses**.
3. In the **Install Licenses Wizard**, on the **Welcome to the Install Licenses Wizard** page, click **Next**.
4. On the **License Program** page, select **License Pack (Retail Purchase)**, and then click **Next**.
5. Enter the 5x5 Beta evaluation license key that you received by email, click **Add**, and then click **Next**.
6. Click **Finish**, to complete installation of the licenses. The license server can now issue RDS CALs to clients that connect to a Remote Desktop Services deployment.

Step 6: Verify Licensing for the Remote Desktop Services

To verify the functionality of the deployment, use Remote Desktop Connection (RDC) to connect to the remote host. The use of the license will be verified on the Remote Desktop license server.

- Verify client connectivity:
  - To verify client connectivity in a session-based desktop deployment: Establish a client connection in a session-based desktop deployment using a valid license
  - To verify client connectivity in a virtual machine-based desktop deployment: Establish a client connection in a virtual machine-based desktop deployment using a valid license
- Verify the RDS CAL by using Remote Desktop Licensing Manager

The following sections explain these procedures in detail.

Establish a client connection in a session-based desktop deployment using a valid license

To establish a client connection in a session-based desktop deployment using a valid license

1. Use the appropriate deployment guide to verify client connectivity. After client connectivity has been established, verify the RDS CAL by using Remote Desktop Licensing Manager.
   - Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start
   - Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Establish a client connection in a virtual machine-based desktop deployment using a valid license

To establish a client connection in a virtual machine-based desktop deployment using a valid license

1. Use the appropriate deployment guide to verify client connectivity. After client connectivity has been established, verify the RDS CAL by using Remote Desktop

Licensing Manager.

- Test Lab Guide: Managed Pooled Virtual Desktop Collections

Verify the RDS CAL by using Remote Desktop Licensing Manager

To verify the RDS CAL by using Remote Desktop Licensing Manager

1. On the left side of the window, click **Remote Desktop Services**, and then click **Servers**, select the server **RDCB1**, and then click **RD Licensing Manager**.
2. In the console tree, expand **All servers**, and then click **RDCB1**.
3. In the results pane, verify that a license is used in the **Issued** column.

Step 7: Snapshot the Configuration

This completes the Remote Desktop Services Licensing test lab. To save this configuration so that you can quickly return to a working Remote Desktop Services Licensing configuration from which you can test other Remote Desktop Services modular test lab guides (TLGs), TLG extensions, or for your own experimentation and learning, do the following:

1. On all physical computers or virtual machines in the test lab, close all windows and then perform a graceful shutdown.
2. If your lab is based on virtual machines, save a snapshot of each virtual machine and name the snapshots Remote Desktop Services Licensing. If your lab uses physical computers, create disk images to save the Remote Desktop Services Licensing test lab configuration.

Additional Resources

What's New in Remote Desktop Services in Windows Server 2012

Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment

Test Lab Guide: Managed Pooled Virtual Desktop Collections

Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections

Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start

Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment

Test Lab Guide: Remote Desktop Services Publishing

You Are Here -> Test Lab Guide: Remote Desktop Services Licensing

Microsoft Remote Desktop Clients

You can use the Microsoft Remote Desktop Client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a remote desktop client

designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps using the RD client.

This page lists the platforms that are supported by Remote Desktop client. Also included are links to resources that will help you get started with using Remote Desktop on your mobile device.

Remote Desktop Client for Android

- Getting Started with Remote Desktop Client on Android
- Remote Desktop Client on Android: FAQ
- Microsoft Remote Desktop for Android

Remote Desktop Client for iOS

- Getting Started with Remote Desktop Client on iOS
- Remote Desktop Client on iOS: FAQ
- Microsoft Remote Desktop for iOS

Remote Desktop Client for Mac

- Getting Started with Remote Desktop Client on Mac
- Remote Desktop Client on Mac: FAQ
- Microsoft Remote Desktop for Mac

Remote Desktop Client for Windows

- Getting Started with Remote Desktop Client on Windows
- Remote Desktop Connection: FAQ

See also

The following table provides additional resources for evaluating Remote Desktop Services.

Content type	References
Product evaluation	• Test Lab Guide: Remote Desktop Services Session Virtualization Quick Start • Test Lab Guide: Remote Desktop Services Session Virtualization Standard Deployment • Test Lab Guide: Virtual Desktop Infrastructure Quick Start

Content type	References
	• Test Lab Guide: Virtual Desktop Infrastructure Standard Deployment • Test Lab Guide: Managed Pooled Virtual Desktop Collections • Test Lab Guide: Unmanaged Pooled Virtual Desktop Collections • Test Lab Guide: Remote Desktop Services Publishing • Test Lab Guide: Remote Desktop Services Licensing • Remote desktop experience in Windows 8 • Test Lab Guide: Demonstrate Remote Desktop Services in Windows Server "8" Beta • Test Lab Guide: Demonstrate Remote Desktop Services Desktop Virtualization in Windows Server "8" Beta • Test Lab Guide: Deploying Remote Desktop Licensing in Windows Server "8" Beta • Understand and Troubleshoot Remote Desktop Services Desktop Virtualization in Windows Server “8” Beta • Understand and Troubleshoot Remote Desktop Services in Windows Server “8” Beta

Community resources	• Remote Desktop Services (Terminal Services) Team Blog
	• Remote Desktop Services Web Forum
	• Remote Desktop Services TechNet Wiki
	• Windows Server 2008 R2 Virtualization Web Forum
	• Windows Virtualization Team Blog
	• Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles
	  • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (01) Introduction to Desktop Virtualization
	  • Microsoft Virtual Academy: Using

Content type	References
	Microsoft VDI to Enable New Workstyles: (02) Optimizing the User Experience • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (03) Deploying Virtual Desktops with Windows Server 2012 and RDS • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (04) Hyper-V for VDI • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (05) Capacity Planning and Architecture • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (06) Leveraging Citrix • Microsoft Virtual Academy: Using Microsoft VDI to Enable New Workstyles: (07) Microsoft VDI Licensing
Related technologies	• Terminal Services in Windows Server 2008 • Remote Desktop Services in Windows Server 2008 R2 • Remote Desktop Services resources (all versions)

Getting Started with Remote Desktop Client on Android

Devices running Android version 2.2 and newer are supported.

Help contents for the Microsoft Remote Desktop Client on Android

You can use the Microsoft Remote Desktop Client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a Remote Desktop client

designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then continue working on those same apps using the Remote Desktop Client.

Here's how to get started with the Remote Desktop Client:

Work with files and apps on another PC. If you want to work on files and access apps on another PC as if you were sitting right in front of it, the Remote Desktop Client is the way to go.

1. Download the Remote Desktop Client from Google Play.
2. Set up your remote PC to allow remote connections.
   - Set up your computer running Windows Vista.
   - Set up your computer running Windows 7.
   - Set up your computer running Windows 8.
3. **Create a Remote Desktop**, or add a **Remote Resource** using the Remote Desktop Client.

A remote desktop connection summarizes all the settings for the remote desktop you want to connect to.

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections. This feature is typically available in corporate environments.

Important
For a list of supported Windows versions, see Remote Desktop Client on Android: FAQ

Getting to know the Connection Center

After starting the Remote Desktop Client for Android the Connection Center opens. You can use the Connection Center to:

- Connect to a Windows PC or server
- Access work resources published using RemoteApp and Desktop Connections
- Manage remote desktops
- Modify settings

Create / Edit Remote Desktops

A remote desktop summarizes all the settings for the remote desktop connection.

Create a Remote Desktop

1. On the Connection Center screen, tap **Remote Desktops**.
2. Tap the add icon.
3. Enter information for the remote desktop connection:

- Remote Desktop
  - Connection name (Optional): The name of the remote desktop in the connection center. This can be any string. If this is field is empty, the PC name will be displayed.
  - PC name: The name of the PC you want to connect to. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally append port information to the PC name, for example: MyDesktop:3389 or 10.0.0.1:3389.
  - Gateway (Optional): The Remote Desktop Gateway enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device. You can either select a gateway that is already configured, or add a new one. No gateway configured is displayed if no Gateway is selected. How to configure a gateway is explained in the section on the **Remote Desktop Gateway**. See your administrator for the name of your Remote Desktop Gateway.
  - Credentials: The user name and password to be used for the remote PC you are connecting to. If no credentials are displayed, the log on screen will be displayed after clicking the remote desktop. The user name can be entered in the following formats: user_name, domain\user_name, and user_name@domain.com.
- Advanced
  - Swap mouse buttons: Whenever a mouse gesture would send a command with the left mouse button, it sends the same command with the right mouse button instead. This is necessary if the remote PC is configured for left-handed mouse mode.
  - Connect to console session: Connect to the console session on a server running Windows Server 2003 and newer.
  - Redirect SD card: Mount your SD card as remote file system on the host.
  - Sound: Select the device to play audio during the remote session. You can configure sound to be played on the local device, the remote PC or not at all.

4. Tap **Done** to create the new remote desktop.

Edit Remote Desktop Settings

To change the settings of a specific remote desktop tap the edit icon available next to the connection bar or within the context menu. When adding or modifying a remote desktop you can modify the following settings:

- Remote Desktop
- Advanced

Copy a Remote Desktop

1. On the Connection Center screen, tap **Remote Desktops**.
2. Long-tap a remote desktop in the list.
3. You can select additional remote desktops to copy.
4. Tap the plus sign to copy the selected remote desktops.

Delete a Remote Desktop

1. On the Connection Center screen, tap **Remote Desktops**.

2. Long-tap a remote desktop in the list.
3. You can select additional remote desktops to delete.
4. Tap the trash can to delete the selected remote desktops.
5. In the **Delete Remote Desktop** window, tap **OK**.

Manage Favorites

Remote desktop connections that are frequently used can be marked as Favorites, which will add the remote desktop connection to the favorites screen.

Add a Remote Desktop to Favorites

1. On the Connection Center screen, tap **Remote Desktops**.
2. Long-tap a remote desktop in the list.
3. You can select additional remote desktops to mark as Favorite.
4. Tap the star to mark the selected remote desktops as a Favorite.
5. On the Connection Center screen, click **Favorites** to access the remote desktop.

Remove a Remote Desktop from Favorites

1. On the Connection Center screen, tap **Remote Desktops**, and then click **Favorites**.
2. Long-tap a remote desktop in the list.
3. You can select additional remote desktops.
4. Tap the star to deselect the remote desktops and remove from favorites.

Search Remote Desktops

Search can be a useful tool when you have to manage many remote desktops.

Search for a Remote Desktop

1. On the Connection Center screen, tap **Remote Desktops**.
2. Tap the magnifier to activate the search bar.
3. Enter text included in the PC name, or the connection name of the remote desktop to filter the remote desktops list.

Manage Settings

These settings apply for all connections and session settings:

Update Settings: Tap menu, and then tap **Settings**.

- Session:
    - Hardware keyboard mode: This setting only affects devices with a hardware keyboard. If your device features an on-device hardware keyboard, you can select the Native mode to use the default keyboard behavior. In Native mode, no emulation is happening, modifiers, such as the Shift, Alt, and Sym keys, have their default behavior. In Desktop mode the default, keystrokes are sent as individual keys to the server.

- Graphics acceleration: Use hardware accelerated rendering to improve performance. Hardware accelerated rendering is available on Android 3.0 and newer, and enabled by default for Android 4.1 and newer. If rendering artifacts are visible on versions earlier than Android 4.1, select the option for software rendering.

- Support:

  Tap **Send log via email** to open the standard email client on your device and create an email that contains log details. Enter the address for the recipient and send the email, to email a copy of the log.

  Warning
  Do not send this log to untrusted people.

Remote Resources

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections.

- The URL displays the link to the RD Web Access server enabling access to RemoteApp and Desktop Connections.
- The configured RemoteApp and Desktop Connections are listed.

You can use the edit icon to change the settings of your remote resource URL.

Add Remote Resources

1. On the Connection Center screen, tap **Remote Resources**, and then tap the add icon.
2. Enter information for the Remote Resource, and then tap **Done**:
   - URL: The URL of the RD Web Access server.
   - User name: The user name to be used for the RD Web Access server you are connecting to.
   - Password: The password to be used for the RD Web Access server you are connecting to.
3. Tap **Done** to add the remote resources.

Refresh Remote Resources

1. On the Connection Center screen, tap **Remote Resources**.
2. Tap refresh to the right of the URL to refresh the selected remote resources.

Delete Remote Resources

1. On the Connection Center screen, tap **Remote Resources**.
2. Tap the trash icon to the right of the URL to delete the selected remote resources.
3. In the **Delete Remote Resources** window, tap **OK**.

Drive Redirection for Remote Resources

Drive redirection is supported for remote resources. Save files created with a remote application locally to your Android device. The redirected folder is displayed as network drive in the remote session. Look for a network drive called **Android on …**.

Note

In order to use this feature the administrator needs to set the appropriate settings on the server.

Remote Desktop Gateway

A Remote Desktop Gateway (RD Gateway) enables authorized users to connect to a remote computer on a corporate network from anywhere on the Internet. You can create and manage your gateways using the Remote Desktop Client.

Set up a New Gateway

1. On the Connection Center screen, tap **Gateways**, and then tap the plus sign.
2. Enter information for the Remote Desktop Gateway, and then tap **Done**:
   - Gateway name: The name of the gateway to be displayed. This can be any string. If this is field is empty, the server will be displayed.
   - Server: The name of the computer you want to use as a gateway. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally append port information to the server name, for example: RDGateway:443 or 10.0.0.1:443.
   - Credentials: The user name and password to be used for the Remote Desktop Gateway you are connecting to. If credentials are not entered, the credentials for the remote desktop connection will be used.

Delete a Gateway

1. On the Connection Center screen, tap **Gateways**.
2. Long-tap a gateway in the list.
3. You can select additional gateway to delete.
4. Tap the trash can to delete the selected gateway.
5. In the **Delete gateway** window, tap **OK**.

Certificates

The Remote Desktop Client will try to validate the certificate when connecting to a remote PC. If the remote PC certificate cannot be validated, you will be presented with a choice of not connecting, or allowing the connection. If the remote PC certificate is permanently accepted, it is stored in the Certificate list, so a user can view the certificate or remove it.

Delete a Certificate

1. On the Connection Center screen, tap **Certificates**.
2. Long-tap a certificate in the list.
3. You can select additional certificates to delete.
4. Tap the trash can to delete the selected certificates.
5. In the **Delete certificate** window, tap **OK**.

Trusted Hosts

Displays a list of all low security remote PCs you have accepted to always trust. If you delete a remote PC as a trusted host, you will be prompted if to trust the remote PC the next time you connect.

Delete a Host

1. On the Connection Center screen, tap **Trusted Hosts**.
2. Long-tap a trusted host in the list.
3. You can select additional certificates to delete.
4. Tap the trash can to delete the selected trusted hosts.
5. In the **Delete host** window, tap **OK**.

Navigating the Remote Desktop Session

Once you start a remote desktop connection there are tools available that you can use to navigate the session.

Start a Remote Desktop Connection

1. On the Connection Center screen, tap **Remote Desktop**.
2. Tap the remote desktop connection, to start the remote desktop session.
3. If the **Untrusted Certificate** window is displayed, tap **Trust Once**, or **Trust Always** if you trust this remote PC.

Once the remote desktop connection has started, the connection bar, a small label with the name of the connection appears at the top of the screen. The connection bar includes tools you can use to navigate the session.

Connection Bar

- Pan Control: The pan control enables the screen to be enlarged and moved around.

  Note
  The pan control is available using direct touch only.
  - Enable / Disable the pan control: Tap the pan icon in the connection bar to display the pan control and zoom the screen. Tap the pan icon in the connection bar again to hide the control and return the screen to its original resolution.
  - Use the pan control: Tap and hold the pan control and then drag in the direction you want to move the screen.
  - Move the pan control: Double tap and hold the pan control to move the control on the screen.
- Connection name / Application name: The current connection name is displayed. Tap the connection name to display the utility bar.
- Keyboard: The keyboard can be displayed, or hidden by tapping the keyboard icon. The pan control is displayed automatically when the keyboard is displayed.

- Move the connection bar: Tap and hold the connection bar, and then drag and drop to a new location at the top of the screen.

Utility Bar

Tap the connection bar to display the utility bar above the keyboard or at the bottom of the screen. The utility bar contains modifier keys, and enables switching between mouse modes, and additional keyboards. There are additional buttons for switching between active applications or starting a new application when connected to remote resources.

Mouse Modes

The mouse modes available are defined in the table.

Note

Interacting with Windows 8 or newer the native touch gestures are supported in Direct Touch mode. For more information on Windows 8 gestures see Touch: Swipe, tap, and beyond.

Mouse mode	Mouse operation	Gesture
Direct Touch	Move pointer	Your finger is the pointer
Direct Touch	Left click	1 finger tap
Direct Touch	Left click and Drag	1 finger tap and drag
Direct Touch	Right click	1 finger tap and hold
Direct Touch	Right click and Drag	1 finger tap and hold, then drag
Direct Touch	Zoom	Use two fingers and pinch to zoom in or move fingers away from each other to zoom out
Mouse Pointer	Move pointer	1 finger tap and drag
Mouse Pointer	Left click	1 finger tap
Mouse Pointer	Left click and drag	1 finger double tap and hold, then drag
Mouse Pointer	Right click	2 finger tap
Mouse Pointer	Right click and drag	2 finger double tap and hold, then drag
Mouse Pointer	Mouse wheel	2 finger tap and hold, then drag up/down

On-screen Keyboards

You can use the utility bar to switch between the standard and additional keyboard.

External Mouse and Keyboard

The right mouse button is supported on Android version 2.3 and newer. Mouse wheel is supported on Android version 3.2 and newer. External keyboards are supported. If your device has an on-device hardware keyboard, review the setting Hardware keyboard mode.

Add a Remote Desktop Widget to the Home Screen

The Remote Desktop Client offers a widget that enables you to create a shortcut for remote desktops on your home screen.

Add a Remote Desktop Widget

1. Tap the All Apps button on the Home screen, and then tap **Widgets**.
2. Scroll through the widgets list to find the Remote Desktop Shortcut widget and then long-press **Remote Desktop Shortcut**.
3. Drag and drop the Remote Desktop Shortcut widget to a Home screen panel.
4. In the Remote Desktop Shortcut list, tap the remote desktop you want to add to your home screen.
5. In the **Shortcut Label** window, enter a label for the shortcut and then click **OK**. The shortcut for the remote desktop connection will be displayed on the window.

You can tap the remote desktop shortcut to start the remote desktop connection.

See also

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Remote Desktop Client on Android: FAQ
- **iPad or iOS:** Remote Desktop Client on iOS: FAQ
- **Mac:** Remote Desktop Client on Mac: FAQ
- **Windows 8:** Remote Desktop Connection: FAQ

Getting Started with Remote Desktop Client on iOS

Devices running iOS 6.X or newer are supported.

Tip

Remote Desktop Client on iOS is enhanced in Windows Server 2012 R2. Download the free Windows Server 2012 R2 trial, learn about the benefits of Windows Server 2012 R2, and get free Windows Server online technical training.

Help contents for the Microsoft Remote Desktop Client on iOS

You can use the Microsoft Remote Desktop Client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a Remote Desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then continue working on those same apps using the Remote Desktop Client.

Here's how to get started with the Remote Desktop Client:

Work with files and apps on another PC. If you want to work on files and access apps on another PC as if you were sitting right in front of it, the Remote Desktop Client is the way to go.

1. Download the Microsoft Remote Desktop Client from iTunes.
2. Set up your remote PC to allow remote connections.
    - Set up your computer running Windows Vista.
    - Set up your computer running Windows 7.
    - Set up your computer running Windows 8.
3. **Create a Remote Desktop**, or add a **Remote Resource** using the Remote Desktop Client.

A remote desktop connection summarizes all the settings for the remote desktop you want to connect to.

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections. This feature is typically available in corporate environments.

Important
For a list of supported Windows versions, see Remote Desktop Client on iOS: FAQ

Getting to know the Connection Center

After starting the Remote Desktop Client for iOS the Connection Center opens. You can use the Connection Center to:

- Connect to a Windows PC or server
- Access work resources published using RemoteApp and Desktop Connections
- Manage remote desktops
- Modify settings

Create / Edit Remote Desktops

A remote desktop summarizes all the settings for the remote desktop connection.

Create a Remote Desktop

1. On the Connection Center screen, tap **New Remote Desktop**.
2. In the **Remote Desktop** window, enter information for the remote desktop connection:
   - REMOTE PC
     - Connection name (Optional): The name of the remote desktop in the connection center. This can be any string. If this is field is empty, the PC name will be displayed.
     - PC name: The name of the PC you want to connect to. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally append port information to the PC name, for example: MyDesktop:3389 or 10.0.0.1:3389.
     - User name: The user name to be used for the remote PC you are connecting to. The user name can be entered in the following formats: user_name, domain\user_name, and user_name@domain.com.
     - Gateway (Optional): The Remote Desktop Gateway enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device. How to configure a gateway is explained in the section on the Remote Desktop Gateway. See your administrator for the name of your Remote Desktop Gateway.
   - SETTINGS
     - Console: Connect to the console session on a server running Windows Server 2003 and newer.
     - Sound: Select the device to play audio during the remote session. You can configure sound to be played on the local device, the remote PC or not at all.
     - Swap mouse buttons: Whenever a mouse gesture would send a command with the left mouse button, it sends the same command with the right mouse button instead. This is necessary if the remote PC is configured for left-handed mouse mode.
3. Tap **Save**.

Edit Remote Desktop Settings

To change the settings of a specific remote desktop tap the edit icon. When adding or modifying a remote desktop you can manage the following categories:

- Remote PC
- Settings

Delete a Remote Desktop

1. On the Connection Center screen, swipe a row from right to left to delete the remote desktop.
2. Tap **Delete** to delete the remote desktop.

Rearrange Remote Desktops

1. On the Connection Center screen, tap the rearrange icon.
2. Tap and hold a remote desktop connection, and then drag to a spot on the list.

3. Tap the rearrange icon when finished.

Search Remote Desktops

Search can be a useful tool when you have to manage many remote desktops.

Search for a Remote Desktop

1. On the Connection Center screen, tap the search text box.
2. Enter text included in the PC name, or the connection name of the remote desktop to filter remote desktops list.

Manage Configuration

These settings apply for all connections and session settings:

- Settings

  Tap **Settings** to manage the session settings.

  - Certificates: The Remote Desktop Client will try to validate the certificate when connecting to a remote PC. If the remote PC certificate cannot be validated, you will be presented with a choice of not connecting, or allowing the connection. If the remote PC certificate is permanently accepted, it is stored in the Certificate list, so a user can view the certificate.
  - Gateway: Display the list of available Remote Desktop Gateways. Configure, or add an RD Gateway.
  - Remote Desktop Licensing: Display the list of client licenses that are used by the remote desktop client.
  - Send log via email:

    Tap **Send log via email** to open the standard email client on your device and create an email that contains log details. Enter the address for the recipient and send the email, to email a copy of the log.

    Warning

    Do not send this log to untrusted people.

- Remote Resources

  Tap **Remote Resources** to configure Remote Resources.

  Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections.

- Help

  Tap **Help** to view web-based help.

- About

  Tap **About** to display information about the Remote Desktop Client.

Certificates

The Remote Desktop Client will try to validate the certificate when connecting to a remote PC. If the remote PC certificate cannot be validated, you will be presented with a choice of not connecting, or allowing the connection. If the remote PC certificate is permanently accepted, it is stored in the Certificate list, so a user can view the certificate.

View Details of a Certificate

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Certificates**.
2. In the **Certificates** window, tap the certificate to view details.

Delete a Certificate

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Certificates**.
2. In the **Certificates** window, swipe a row from right to left to delete the certificate.
3. Tap **Delete** to delete the certificate.

Lower security PCs are displayed. Lower security PCs are remote PCs that could not be connected to using a secure channel. Swipe a remote PC, and then tap delete to delete a trusted host from this list. If you delete a host, you will be prompted if you would like to trust this connection next time.

Gateway

A Remote Desktop Gateway (RD Gateway) enables authorized users to connect to a remote computer on a corporate network from anywhere on the Internet. You can create and manage your gateways using the Remote Desktop Client.

Set up a New Gateway

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Gateway**.
2. In the **Gateway** window tap **Add gateway**, and then tap **No gateway configured**.
3. In the **No gateway configured** window, enter information for the Remote Desktop Gateway:
   - Gateway name (Optional): The name of the gateway to be displayed. This can be any string. If this is field is empty, the server will be displayed.
   - Server: The name of the computer you want to use as the gateway. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally append port information to the server name, for example: RDGateway:443 or 10.0.0.1:443.
   - Credentials: The user name and password to be used for the Remote Desktop Gateway you are connecting to. If credentials are not entered, the credentials for the remote desktop connection will be used.
4. In the **No gateway configured** window tap **Gateway**, and then tap **Settings** to return to the Settings window.
5. In the **Settings** window tap **Done** to save the new configuration.

Delete a Gateway

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Gateway**.
2. In the **Gateway** window, swipe a row from right to left to delete the RD Gateway.
3. Tap **Delete** to delete the gateway.

Remote Desktop licensing

A Remote Desktop Session Host server determines if a Remote Desktop Services client access licenses (RDS CALs) is required. If an RDS CAL is required a list of the licenses currently issued to the Remote Desktop Client will be displayed.

View Details of a License

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Remote Desktop licensing**.
2. In the **Remote Desktop licensing** window, tap the license to view details.

Delete a License

1. On the Connection Center screen, under **Configuration** tap **Settings**, and then tap **Remote Desktop licensing**.
2. In the **Remote Desktop licensing** window, swipe a row from right to left to delete the license.
3. Tap **Delete** to delete the license.

Remote Resources

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections.

- The URL displays the link to the RD Web Access server enabling access to RemoteApp and Desktop Connections.
- The configured RemoteApp and Desktop Connections are listed.

Add Remote Resources

1. On the Connection Center screen, under **Configuration** tap **Remote Resources**.
2. In the **Remote Resources** window, enter information for the Remote Resource:
    - URL: The URL of the RD Web Access server.
    - User name: The user name to be used for the RD Web Access server you are connecting to.
    - Password: The password to be used for the RD Web Access server you are connecting to.
3. Tap **Save**. The remote resources will be displayed in the Connection Center.

Refresh Remote Resources

- On the Connection Center screen, tap the <refresh> icon to the right of the remote resource.

Delete Remote Resources

1. On the Connection Center screen, under **Configuration** tap **Remote Resources**.
2. In the **Remote Resources** window, tap **URL**.
3. In the URL window, tap the **x** button, and then tap **Save**.
4. Tap **Save**, to finalize deleting the Remote Resources.

Navigating the Remote Desktop Session

Once you start a remote desktop connection there are tools available that you can use to navigate the session.

Start a Remote Desktop Connection

1. Tap the remote desktop connection, to start the remote desktop session.
2. If you are asked to verify the certificate for the remote PC, tap **Connect once**, or **Connect always** if you trust this remote PC.

Connection Bar

- Pan Control: The pan control enables the screen to be enlarged and moved around.

  Note

  The pan control is available using direct touch only.

  - Enable / Disable the pan control: Tap the pan icon in the connection bar to display the pan control and zoom the screen. Tap the pan icon in the connection bar again to hide the control and return the screen to its original resolution.
  - Use the pan control: Tap and hold the pan control and then drag in the direction you want to move the screen.
  - Move the pan control: Double tap and hold the pan control to move the control on the screen.
- Connection name / Application name: The current connection name is displayed. Tap the connection name to display the utility bar.
- Keyboard: The keyboard can be displayed, or hidden by tapping the keyboard icon. The pan control is displayed automatically when the keyboard is displayed.
- Move the connection bar: Tap and hold the connection bar, and then drag and drop to a new location at the top of the screen.

Utility Bar

Tap the connection bar to display the utility bar at the top of the screen. The utility bar enables switching between mouse modes, and disconnects the session. There are additional buttons for switching between active applications or starting a new application when connected to remote resources.

Mouse Modes

The mouse modes available are defined in the table.

Note

Interacting with Windows 8 or newer the native touch gestures are supported in Direct Touch mode. For more information on Windows 8 gestures see Touch: Swipe, tap, and beyond.

Mouse mode	Mouse operation	Gesture
Direct Touch	Move pointer	Your finger is the pointer
Direct Touch	Left click	1 finger tap
Direct Touch	Left click and Drag	1 finger tap and drag
Direct Touch	Right click	1 finger tap and hold
Direct Touch	Right click and Drag	1 finger tap and hold, then drag
Direct Touch	Zoom	Use two fingers and pinch to zoom in or move fingers away from each other to zoom out
Mouse Pointer	Move pointer	1 finger tap and drag
Mouse Pointer	Left click	1 finger tap
Mouse Pointer	Left click and drag	1 finger double tap and hold, then drag
Mouse Pointer	Right click	2 finger tap
Mouse Pointer	Right click and drag	2 finger double tap and hold, then drag
Mouse Pointer	Mouse wheel	2 finger tap and hold, then drag up/down

On-screen Keyboards

On the iPad you can use the buttons on the edges of the bar above the keyboard to switch between the standard and additional keyboard.

On the iPhone and iPod touch devices you can switch between four different keyboard types depending on the keys that are needed.

- Letter keyboard
- Numeric keypad
- Special keys
- Function keys

Bluetooth Keyboard

If Bluetooth is enabled, a Bluetooth keyboard should be automatically detected.

Due to limitations on the iPhone OS special keys such as Ctrl, Option, and Function will not work as expected. The following keys currently work:

- Alphanumeric keys
- Cursor keys
- Tab: Tab works, but Shift+Tab does not work
- Home / Pos1: Alt+Left = Home
- End: Alt+Right = End
- Page Up: Alt+Up = Page Up
- Page Down: Alt+Down = Page Down
- Select All: Command+A = Ctrl+A (Select all in most programs)
- Cut: Command+X = Ctrl+X (Cut in most programs)
- Copy: Command+C = Ctrl+C (Copy in most programs)
- Paste: Command+V = Ctrl+V (Paste in most programs)
- Symbols: Alt+Alphanumeric keys will produce different symbols depending on the language configured

See also

The following list provides additional resources for using Remote Desktop Clients.

- **Android:** Remote Desktop Client on Android: FAQ
- **iPad or iOS:** Remote Desktop Client on iOS: FAQ
- **Mac:** Remote Desktop Client on Mac: FAQ
- **Windows 8:** Remote Desktop Connection: FAQ

Getting Started with Remote Desktop Client on Mac

Devices running OS 10.6.8 and newer are supported.

Help contents for the Microsoft Remote Desktop Client on Mac

You can use the Microsoft Remote Desktop Client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a Remote Desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then continue working on those same apps using the Remote Desktop Client.

Here's how to get started with the Remote Desktop Client:

Work with files and apps on another PC. If you want to work on files and access apps on another PC as if you were sitting right in front of it, the Remote Desktop Client is the way to go.

1. Download the Remote Desktop Client from the Mac App Store.
2. Set up your remote PC to allow remote connections.
   - Set up your computer running Windows Vista.
   - Set up your computer running Windows 7.
   - Set up your computer running Windows 8.
3. **Create a Remote Desktop**, or add a **Remote Resource** using the Remote Desktop Client.

A remote desktop connection summarizes all the settings for the remote desktop you want to connect to.

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections. This feature is typically available in corporate environments.

Warning
For a list of supported Windows versions, see Remote Desktop Client on Mac: FAQ

Getting to know the Connection Center

After starting the Remote Desktop Client for Mac the Connection Center opens. You can use the Connection Center to:

- Connect to a Windows PC or server
- Access work resources published using RemoteApp and Desktop Connections
- Manage remote desktops
- Modify settings

Create / Edit Remote Desktops

A remote desktop summarizes all the settings for the remote desktop connection.

Create a Remote Desktop

1. On the Connection Center screen, click **New**.
2. In the **Edit Remote Desktops** window, enter information for the remote desktop connection:
   - General
     - Connection name (Optional): The name of the remote desktop in the connection center. This can be any string. If this is field is empty, the PC name will be displayed.
     - PC name: The name of the PC you want to connect to. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally

append port information to the PC name, for example: MyDesktop:3389 or 10.0.0.1:3389.

- Gateway: The Remote Desktop Gateway enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device. You can either select a gateway that is already configured, or add a new one. No gateway configured is displayed if no Gateway is selected. How to configure a gateway is explained in the section on the Remote Desktop Gateway. See your administrator for the name of your Remote Desktop Gateway.
- Credentials: The user name and password to be used for the remote PC you are connecting to. If no credentials are displayed, the log on screen will be displayed after clicking the remote desktop. The user name can be entered in the following formats: user_name, domain\user_name, and user_name@domain.com.
- Resolution: Select the display resolution for the remote session.
- Colors: Select the color depth for the remote session.
- Full screen mode: Choose between the OS X native full screen feature (Mac OS X 10.7+) or the custom full screen mode which enables the multi-monitor feature to provide a better experience for users of Mac OS X versions 10.7 and 10.8..
- Start session in full screen: The remote session will start using the full screen. Disable this option if you prefer to start a remote desktop connection in windowed mode.
- Scale content: Enable the aspect ratio to be fixed for the session content.
- Use all monitors: Enable use of all monitors for the remote session.

- Session
  - Sound: Select the device to play audio during the remote session. You can configure sound to be played on the local device, the remote PC or not at all.
  - Connect to console session: Connect to the console session on a server running Windows Server 2003 and newer.
  - Forward printing devices: Locally installed printers are available during the remote session.
  - Swap mouse buttons: Whenever a mouse gesture would send a command with the left mouse button, it sends the same command with the right mouse button instead. This is necessary if the remote PC is configured for left-handed mouse mode.
- Redirection

  Folder redirection enables access local folders during the remote session. Click the + button at the bottom of the dialog and choose a folder you want to have redirected.

  Configure Folder Redirection

  i. Click the **+** button.

  ii. In the Add Local Folder window, enter the following information:
    - Name: Set a name for the folder to be available during the remote session.
    - Path: Select the path to the folder to be available during the remote session.

3. Click the close button to save the remote desktop.

Edit Remote Desktop Settings

To change the settings of a specific remote desktop, right-click the remote desktop and then click edit. When adding or modifying a remote desktop the following tabs are available:

- General
- Session
- Redirection

Delete a Remote Desktop

1. On the Connection Center screen, tap the right-click the remote desktop to delete.
2. Click **Delete**, and then click **Yes**.

Copy a Remote Desktop

1. On the Connection Center screen, tap the right-click the remote desktop to copy.
2. Click **Duplicate**, to copy the remote desktop.

Export Remote Desktops

Remote desktops can be exported. Remote desktops are saved in separate files with an RDP extension.

Export a Remote Desktop

1. On the Connection Center screen, tap the right-click the remote desktop to export.
2. Click **Export**.
3. In the **Export Options** window, click **Browse** to select a path for the file.
4. Click **OK** to export the remote desktop.

Import Remote Desktops

Exported remote desktop files can be imported.

Import a Remote Desktop

1. On the menu bar click **File**, and then click **Import**.
2. In the **Import Bookmarks** window, browse and select the remote desktop file.
3. Click **Open** to import the remote desktop connection.

Search Remote Desktops

Search can be a useful tool when you have to manage many remote desktops.

Search for a Remote Desktop

1. On the Connection Center screen, tap the search text box.
2. Enter text included in the PC name, or the connection name of the remote desktop to filter remote desktops list.

Rearrange Remote Desktops

Depending on your environment, the list of remote desktops can get quite long. The remote desktops in this list can be rearranged to bring frequently used remote desktops to the top of the list and remote desktops that are used less frequently can be moved towards the bottom of the list.

Rearrange Remote Desktop in Connection Center

1. Click and hold a remote desktop.
2. Drag and drop the remote desktop to a new position in the list.

Manage Preferences

These settings apply for all connections and session preferences:

Update Preferences: In the Connection Center, click **Preferences**.

- Gateway: Display the list of available Remote Desktop Gateways. Select, configure, or add an RD Gateway.
- Resolution: Select the display resolution for remote desktop sessions.

Remote Desktop Gateway

A Remote Desktop Gateway (RD Gateway) enables authorized users to connect to a remote computer on a corporate network from anywhere on the Internet. The Remote Desktop Client for Mac supports RD Gateway. You can create and manage your gateways using the Remote Desktop Client.

Set up a New Gateway

1. On the Connection Center screen, click **Preferences**.
2. In the **Preferences** window click **Gateway**. Enter information for the Remote Desktop Gateway:
    - Gateway name: The name of the gateway to be displayed. This can be any string. If this is field is empty, the server will be displayed.
    - Server: The name of the computer you want to use as gateway. This can be a Windows computer name, an Internet domain name or an IP address. You can optionally append port information to the server name, for example: RDGateway:443 or 10.0.0.1:443.
    - User name: The user name to be used for the RD Gateway you are connecting to. If a user name and password are not entered, the credentials for the remote desktop connection will be used.
    - Password: The password to be used for the RD Gateway you are connecting to.
3. Click **+** to add the RD Gateway, and then click close.

Delete a Gateway

1. On the Connection Center screen, click **Preferences**.
2. In the **Preferences** window click **Gateway**.
3. Click the RD Gateway to delete, click **-**, and then click close.

Resolution

The resolution setting allows you to select the display resolution for the remote desktop session.

Add a Resolution

1. On the Connection Center screen, click **Preferences**.
2. In the **Preferences** window click **Resolution**.
3. Click **+**.
4. In the **Add Resolution** window, enter a resolution width and height and then click **OK**.

Delete a Resolution

1. On the Connection Center screen, click **Preferences**.
2. In the **Preferences** window click **Resolution**.
3. Click the resolution to delete, click **-**, and then click close.

Displays have separate Spaces

Mac OS X 10.9 users who have disabled **Displays have separate Spaces** in Mavericks (System Preferences -> Mission Control) need to configure this setting in the Microsoft Remote Desktop client using the same option.

Remote Resources

Remote Resources enables access to RemoteApp programs, session-based desktops, and virtual desktops published using RemoteApp and Desktop Connections.

- The URL displays the link to the RD Web Access server enabling access to RemoteApp and Desktop Connections.
- The configured RemoteApp and Desktop Connections are listed.

Add Remote Resources

1. On the Connection Center screen, tap **Remote Resources**.
2. Enter information for the Remote Resource:
    - URL: The URL of the RD Web Access server.
    - User name: The user name to be used for the RD Web Access server you are connecting to.
    - Password: The password to be used for the RD Web Access server you are connecting to.
3. Click **Refresh** and then click close.

Refresh Remote Resources

1. On the Connection Center screen, tap **Remote Resources**.
2. Click **Refresh** and then click close.

Delete Remote Resources

1. On the Connection Center screen, tap **Remote Resources**.
2. Delete the URL.
3. Click **Refresh** and then click close.

Drive Redirection for Remote Resources

Drive redirection is supported for remote resources. Save files created with a remote application locally to your Mac. The redirected folder is always your home directory displayed as a network drive in the remote session.

Note

In order to use this feature the administrator needs to set the appropriate settings on the server.

Navigating the Remote Desktop Session

Once you start a remote desktop connection there are tools available that you can use to navigate the session.

Start a Remote Desktop Connection: Double-click the remote desktop connection.

If you are in full screen mode, the menu bar at the top of the screen will be hidden. This remote connection menu bar includes tools you can use to navigate the session.

Remote Desktop Menu Bar

- Microsoft Remote Desktop: This menu enables you to manage the remote desktop session.
- Window: Enables you to manage the windows.
- Help: View web-based help.

Start a Remote Resource: Double-click the remote application.

Remote Resource Menu Bar

- Microsoft Remote Desktop: This menu enables you to manage the remote desktop session.
- Remote Desktop Client: Right-click the remote desktop client icon in the menu bar, to switch between active applications. You can also right-click the Remote Desktop Client icon in the dock to switch between applications.

Mac and Windows Keyboard Layouts

Mac keyboard layouts differ from the Windows keyboard layouts. The OS keyboard locale, is suggested to the remote PC. If the OS keyboard locale is not used, we recommend checking the keyboard setting on the remote PC and changing the setting manually.

- The Command key on the Mac keyboard equals the Windows key.
- All actions, which you perform using the Command button on the Mac, you need to use the control button in Windows (e.g.: Copy = Ctrl + C).
- The function keys can be activated in the session by pressing additionally the FN key (e.g.: FN + F1).
- The Alt key to the right of the space bar on the Mac keyboard equals the Alt Gr/right Alt key in Windows.

See also

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Remote Desktop Client on Android: FAQ
- **iPad or iOS:** Remote Desktop Client on iOS: FAQ
- **Mac:** Remote Desktop Client on Mac: FAQ
- **Windows 8:** Remote Desktop Connection: FAQ

Remote Desktop Client on Android: FAQ

You can use the Remote Desktop client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a remote desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps using the RD client.

Remote Desktop Client: FAQ

You can use the RD Client to sit at a PC and connect to another PC in a different location (the remote PC). For example, you can connect to your work PC from your home PC and use all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps on your home PC.

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Getting Started with Remote Desktop Client on Android
- **iPad or iOS:** Getting Started with Remote Desktop Client on iOS
- **Mac:** Getting Started with Remote Desktop Client on Mac
- Set up your computer running Windows Vista
- Set up your computer running Windows 7
- Set up your computer running Windows 8

How do I set up a PC for Remote Desktop Connection?

On the PC you plan to connect to remotely, do the following:

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
3. In the **System Properties** dialog box, under **Remote Desktop**, select one of the two options. For more info about these options, see What types of Remote Desktop connections should I allow?
4. Tap or click **Select Users**, and in the **Remote Desktop Users** dialog box, tap or click **Add**.
5. In the **Select Users or Groups** dialog box, do the following:

6. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
    - To select the search location, tap or click **Locations**, and then select the location you want to search.
    - In **Enter the object names to select**, enter the name of the person that you want to add, and then tap or click **OK**. The name will be added to the list of people who can connect to this PC.

Note

You can't connect to a PC that's asleep or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

When you want to connect to a remote PC, you'll need to know the full name of that PC.

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. The full PC name is listed under **Computer name, domain, and workgroup settings**.

Which PCs can I connect to using Remote Desktop Connection?

You can only connect to PCs that are running these Windows operating systems:

- Windows 8 Enterprise
- Windows 8 Pro
- Windows 7 Professional
- Windows 7 Enterprise
- Windows 7 Ultimate
- Windows Vista Business
- Windows Vista Ultimate
- Windows Vista Enterprise
- Windows XP Professional
- Windows Server 2000
- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Multipoint Server 2011
- Windows Multipoint Server 2012
- Windows Small Business Server 2003

- Windows Small Business Server 2008
- Windows Small Business Server 2011

These Windows operating systems can be used as RD Gateway servers:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Small Business Server 2011

These Windows operating systems can be used to serve as RD Web Access or RemoteApp servers:

- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Why can’t I connect using Remote Desktop Connection?

Here are some possible solutions to common problems you might encounter when trying to connect to a remote PC. If these solutions don’t work, you can find more help on the Microsoft Community website.

- **The remote PC can’t be found.** If you see this error, talk to the system admin of the remote PC to make sure you have the right PC name, and then check to see if you entered the name correctly. If you still can't connect, try entering the IP address of the remote PC instead of the PC name.
- **There’s a problem with the network.** Make sure that:
  - Your router is turned on (home networks only).
  - The Ethernet cable is plugged into your network adapter (wired networks only).
  - Your PC’s wireless switch is turned on (laptops on wireless networks only).
  - Your network adapter is working.

  If the problem keeps happening, talk to the person who set up your network.
- **The Remote Desktop port might be blocked by a firewall.** If you're using Windows Firewall, follow these steps:
  a. Open Windows Firewall by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **firewall** in the search box, tapping or clicking **Settings**, and then tapping or clicking **Windows Firewall**.
  b. Tap or click **Allow an app or feature through Windows Firewall**.
  c. Tap or click **Change settings**. You might be asked for an admin password or to confirm your choice.
  d. Under **Allowed apps and features**, select **Remote Desktop**, and then tap or click **OK**.

  If you're using a different firewall, make sure the port for Remote Desktop (usually 3389) is open.

- **Remote connections might not be set up on the remote PC.** To fix this, see the "How do I set up a PC for Remote Desktop Connection?" question in this topic.
- **The remote PC might only allow PCs to connect that have Network Level Authentication set up.** For more info, see What types of Remote Desktop connections should I allow?
- **The remote PC might be turned off.** You can't connect to a PC that's turned off, asleep, or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

Why can't I sign in to a remote PC?

If you can see the sign-in screen of the remote PC but you can't sign in, you might not have been added to the Remote Desktop Users Group or to any group with administrator rights on the remote PC. Ask your system admin to do this for you.

Why are the characters mixed up in the session?

This issue can occur if the keyboard language has been switched on the remote PC while running a remote desktop session. In the case you would like to switch to a different keyboard language follow these instructions.

Important
The remote PC selects the system language of the device running the remote desktop client and not the keyboard language. When switching the keyboard language on your device, you need to manually switch the keyboard language on the remote PC.

1. Open settings on your device.
2. Tap **General** and then tap **International**.
3. Tap **Language** and select the language running on the remote PC.
4. Open the RD Client and connect to the remote PC and check if the Windows keyboard language matches the language that you selected.

I cannot find or connect to my computer

Verify that your computer is on. You can either enter your computers IP address or host name.

Important
Using the host name requires your network to resolve the host name properly through DNS. In many home network configurations only the IP address will work properly.

Verify that you are connecting to a supported Windows version.

Note
Windows XP Home, Windows Media Center Edition, Windows Vista Home and Windows 7 Home or Starter are not supported without 3rd party software.

VPN does not work

These issues can have several reasons. For all reasons it is important to note, that if you can, please check that the VPN works on the same network from your PC or Mac. If you can't test from a PC or Mac you can try to access some company intranet webpage with your device's web browser and see if this works for at least a couple of minutes.

- **The 3G network blocks or corrupts VPN:** There are several 3G providers in the world who seem to block or corrupt 3G traffic. Verify VPN connectivity works correctly for over a minute.
- **L2TP or PPTP VPNs:** If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.
- **VPN is misconfigured:** A misconfigured VPN server can be the reason why the VPN connections never worked or stopped working after some time. Ensure testing with the iOS device's web browser or a PC or Mac on the same network if this happens.

Which connection methods are supported for company networks?

If you want to access your office desktop from outside your company network, your company must provide you with a means of remote access. The RD Client currently supports the following:

- Terminal Server Gateway or Remote Desktop Gateway
- Remote Desktop Web Access
- VPN (through iOS built-in VPN options

Is bi-directional sound supported?

Sound upstream (from client to server, for microphones) is not supported by the Remote Desktop Client.

Is the Remote Desktop Client compatible with RDP 7.1?

Yes, the Remote Desktop client supports RDP 7.1 and all previous versions including: RDP 7, RDP 6.1, RDP 6, RDP 5 and RDP 4.

Note

Support for a specific version of RDP 7 does not include all virtual channels. RemoteFX codecs are not supported in RDP 7.1.

How to configure L2TP or PPTP VPN connections

If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.

Error: Insufficient privileges

You are not allowed to access the session you want to connect to. The most likely cause is that you are trying to connect to the console session. Only administrators are allowed to connect to the console. Verify that the console switch is off in the advanced settings of the remote desktop. If this is not the source of the problem, please contact your system administrator for further assistance.

Error: Failed to parse NTLM challenge

This error can be caused by a misconfiguration on the remote PC. Verify the RDP security level setting in the remote PC options is set to Client Compatible.

What can I do if the sound does not play?

If sound doesn't play, log out of the session, not just disconnect, and log in again.

Unsupported Windows Versions and Editions

The Remote Desktop Client will not connect to these Windows Versions and Editions:

- Windows XP Starter
- Windows XP Home
- Windows XP Tablet PC
- Windows Media Center
- Windows Vista Starter
- Windows Vista Home
- Windows 7 Starter
- Windows 7 Home

If you want to access computers that have one of these Windows versions installed, we recommend you upgrade to a Windows version that supports RDP.

RD Gateway messaging is not supported

Remote Desktop Client does not support RD Gateway messaging. Verify that the Remote Desktop Resource Access Policy (RD RAP) for your RD Gateway server does not specify **Only allow computers with support for RD Gateway Messaging** or you will not be able to connect.

ERROR: TS_RAP "You are not allowed to connect to the given host ..."

The error message “You are not allowed to connect to the given host through the gateway server because of a Resource Authorization Policy (TS_RAP)” means that a Resource Authorization Policy on the Gateway server prevents your username from connecting to the remote PC. This could be the case:

- You have provided the same name for the remote PC as for the gateway.
- You are then trying to connect to the RD Gateway server itself which is probably not allowed for your username.
- Your Windows user account is not a member of the user groups for remote access on the Remote Desktop server.

If you indeed want to connect to the gateway server itself via the gateway, do not use the external gateway name as PC name, use “localhost” or “127.0.0.1” or the internal server name for the remote desktop. If the entries are correct, ask your system administrator to view the event log on the Terminal server and to check if your user account is a member of the required user groups.

Error message that there is no CAL

When a remote desktop client connects to a Remote Desktop server, the server issues a Remote Desktop Services Client Access License (RDS CAL) stored by the client. Whenever the client connects again it will use its RDS CAL and the server will not issue another license. The server will issue another license if the RDS CAL on the device is missing or corrupt. When the maximum number of licensed devices is reached the server will not issue new RDS CALs. Contact your network administrator for assistance.

Error: Access Denied

The "Access Denied" error is a generated by the Remote Desktop Gateway and the result of incorrect credentials during the connection attempt. Verify your username and password. If the connection worked before and the error occurred recently, you possibly changed your Windows user account password and haven't updated it yet in the remote desktop settings.

How can I test if VPN is working properly?

Verify that VPN is enabled on your device. You can test your VPN connection by going to a webpage on your internal network or using a web service which is only available via the VPN.

Error: RPC Error 23014 or Error 0x59e6

In case of an **RPC error 23014** or **Error 0x59E6 try again after waiting a few minutes**, the RD Gateway server has reached the maximum number of active connections. Depending on the Windows version running on the RD Gateway the maximum number of connections differs: The Windows Server 2008 R2 Standard implementation limits the number of connections to 250. The Windows Server 2008 R2 Foundation implementation limits the number of connections to 50. All other Windows implementations allow an unlimited number of connections.

See also

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Getting Started with Remote Desktop Client on Android

- **iPad or iOS:** Getting Started with Remote Desktop Client on iOS
- **Mac:** Getting Started with Remote Desktop Client on Mac
- **Windows 8:** Remote Desktop Connection: FAQ

Remote Desktop Client on iOS: FAQ

You can use the Remote Desktop client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a remote desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps using the RD client.

Remote Desktop Client: FAQ

You can use the RD Client to sit at a PC and connect to another PC in a different location (the remote PC). For example, you can connect to your work PC from your home PC and use all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps on your home PC.

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Getting Started with Remote Desktop Client on Android
- **iPad or iOS:** Getting Started with Remote Desktop Client on iOS
- **Mac:** Getting Started with Remote Desktop Client on Mac
- Set up your computer running Windows Vista
- Set up your computer running Windows 7
- Set up your computer running Windows 8

How do I set up a PC for Remote Desktop Connection?

On the PC you plan to connect to remotely, do the following:

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
3. In the **System Properties** dialog box, under **Remote Desktop**, select one of the two options. For more info about these options, see What types of Remote Desktop connections should I allow?
4. Tap or click **Select Users**, and in the **Remote Desktop Users** dialog box, tap or click **Add**.
5. In the **Select Users or Groups** dialog box, do the following:

6. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
    - To select the search location, tap or click **Locations**, and then select the location you want to search.
    - In **Enter the object names to select**, enter the name of the person that you want to add, and then tap or click **OK**. The name will be added to the list of people who can connect to this PC.

Note

You can't connect to a PC that's asleep or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

When you want to connect to a remote PC, you'll need to know the full name of that PC.

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. The full PC name is listed under **Computer name, domain, and workgroup settings**.

Which PCs can I connect to using Remote Desktop Connection?

You can only connect to PCs that are running these Windows operating systems:

- Windows 8 Enterprise
- Windows 8 Pro
- Windows 7 Professional
- Windows 7 Enterprise
- Windows 7 Ultimate
- Windows Vista Business
- Windows Vista Ultimate
- Windows Vista Enterprise
- Windows XP Professional
- Windows Server 2000
- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Multipoint Server 2011
- Windows Multipoint Server 2012
- Windows Small Business Server 2003

- Windows Small Business Server 2008
- Windows Small Business Server 2011

These Windows operating systems can be used as RD Gateway servers:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Small Business Server 2011

These Windows operating systems can be used to serve as RD Web Access or RemoteApp servers:

- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Why can’t I connect using Remote Desktop Connection?

Here are some possible solutions to common problems you might encounter when trying to connect to a remote PC. If these solutions don’t work, you can find more help on the Microsoft Community website.

- **The remote PC can’t be found.** If you see this error, talk to the system admin of the remote PC to make sure you have the right PC name, and then check to see if you entered the name correctly. If you still can't connect, try entering the IP address of the remote PC instead of the PC name.
- **There’s a problem with the network.** Make sure that:
  - Your router is turned on (home networks only).
  - The Ethernet cable is plugged into your network adapter (wired networks only).
  - Your PC’s wireless switch is turned on (laptops on wireless networks only).
  - Your network adapter is working.

  If the problem keeps happening, talk to the person who set up your network.
- **The Remote Desktop port might be blocked by a firewall.** If you're using Windows Firewall, follow these steps:
  a. Open Windows Firewall by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **firewall** in the search box, tapping or clicking **Settings**, and then tapping or clicking **Windows Firewall**.
  b. Tap or click **Allow an app or feature through Windows Firewall**.
  c. Tap or click **Change settings**. You might be asked for an admin password or to confirm your choice.
  d. Under **Allowed apps and features**, select **Remote Desktop**, and then tap or click **OK**.

  If you're using a different firewall, make sure the port for Remote Desktop (port 3389) is open.

- **Remote connections might not be set up on the remote PC.** To fix this, see the "How do I set up a PC for Remote Desktop Connection?" question in this topic.
- **The remote PC might only allow PCs to connect that have Network Level Authentication set up.** For more info, see What types of Remote Desktop connections should I allow?
- **The remote PC might be turned off.** You can't connect to a PC that's turned off, asleep, or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

Why can't I sign in to a remote PC?

If you can see the sign-in screen of the remote PC but you can't sign in, you might not have been added to the Remote Desktop Users Group or to any group with administrator rights on the remote PC. Ask your system admin to do this for you.

Error: Insufficient privileges

You are not allowed to access the session you want to connect to. The most likely cause is that you are trying to connect to the console session. Only administrators are allowed to connect to the console. Verify that the console switch is off in the settings for the remote desktop. If this is not the source of the problem, please contact your system administrator for further assistance.

How to configure L2TP or PPTP VPN connections

If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.

What do I need to connect to an office computer using Remote Desktop Gateway?

To connect using a Remote Desktop Gateway you need three things:

1. Server name of the Remote Desktop Gateway: rdgw.contoso.com.
2. Name of your computer/workstation at the office.
3. Username, domain and password

I cannot find or connect to my computer

Verify that your computer is on. You can either enter your computers IP address or host name.

Important
Using the host name requires your network to resolve the host name properly through DNS. In many home network configurations only the IP address will work properly.

Verify that you are connecting to a supported Windows version.

Note

Windows XP Home, Windows Media Center Edition, Windows Vista Home and Windows 7 Home or Starter, and Windows 8 do not support RDP.

VPN does not work

The VPN not working can be due to several issues. Verify the VPN works on the same network from your PC. If you are unable to test from a PC you can try to access some company intranet webpage with your device's web browser and see if this works for at least a couple of minutes.

- **The 3G network blocks or corrupts VPN:** There are several 3G providers in the world who seem to block or corrupt 3G traffic. Verify VPN connectivity works correctly for over a minute.
- **L2TP or PPTP VPNs:** If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.
- **VPN is misconfigured:** A misconfigured VPN server can cause the VPN connections not to work or stop working after a period of time. Verify testing with the iOS device's web browser or a PC or Mac on the same network to test this issue.

What can I do when the app does not start anymore?

To resolve this issue please open the multi-tasking menu on your iOS device by double-pressing the home button. Then navigate to the RD Client icon and tap and hold your finger on the icon until the icon shows a red circle. Press this red circle to kill the app and then try to restart it.

If this does not work, it could be that settings have been corrupted which prevent the Remote Desktop app for iOS to start properly. To resolve this, delete the application and reinstall.

Which connection methods are supported for company networks?

If you want to access your office desktop from outside your company network, your company must provide you with a means of remote access. The RD Client currently supports the following:

- Terminal Server Gateway or Remote Desktop Gateway
- Remote Desktop Web Access
- VPN (through iOS built-in VPN options

Error: STOP error in Windows XP

This is a bug in Windows XP SP2 and SP3 and has been addressed by Microsoft. You need to open a support ticket with Microsoft to officially download the patch.

Is AirPrint supported?

The RD Client does not support AirPrint.

Is bi-directional sound supported?

Sound upstream (from client to server, for microphones) is not supported by the RD Client.

Why are the characters mixed up in the session?

This issue can occur if the keyboard language has been switched on the remote PC while running a remote desktop session. In the case you would like to switch to a different keyboard language follow these instructions.

Important
The remote PC selects the system language of the device running the remote desktop client and not the keyboard language. When switching the keyboard language on your device, you need to manually switch the keyboard language on the remote PC.

1. Open settings on your device.
2. Tap **General** and then tap **International**.
3. Tap **Language** and select the language running on the remote PC.
4. Open the RD Client and connect to the remote PC and check if the Windows keyboard language matches the language that you selected.

Why don’t special keys work on a Bluetooth keyboard?

On Bluetooth keyboards modifier keys such as Ctrl, Alt, Shift F-keys do not work as expected. This is a limitation of iOS. You can work around this by using the on-screen keyboard within the remote desktop session that you can activate by a tap on the keyboard icon in the connection bar.

Is the Remote Desktop Client compatible with RDP 7.1?

Yes, the Remote Desktop client supports RDP 7.1 and all previous versions including: RDP 7, RDP 6.1, RDP 6, RDP 5 and RDP 4.

Note
Support for a specific version of RDP 7 does not include all virtual channels. RemoteFX codecs are not supported in RDP 7.1.

Error: Failed to parse NTLM challenge

This error can be caused by a misconfiguration on the remote PC. Verify the RDP security level setting in the remote PC options is set to Client Compatible.

What can I do if the sound does not play?

If sound doesn't play, log out of the session, not just disconnect, and log in again.

Unsupported Windows Versions and Editions

The Remote Desktop Client will not connect to these Windows Versions and Editions:

- Windows XP Starter
- Windows XP Home
- Windows XP Tablet PC
- Windows Media Center
- Windows Vista Starter
- Windows Vista Home
- Windows 7 Starter
- Windows 7 Home

If you want to access computers that have one of these Windows versions installed, we recommend you upgrade to a Windows version that supports RDP.

RD Gateway messaging is not supported

Remote Desktop Client does not support RD Gateway messaging. Verify that the Remote Desktop Resource Access Policy (RD RAP) for your RD Gateway server does not specify **Only allow computers with support for RD Gateway Messaging** or you will not be able to connect.

ERROR: TS_RAP "You are not allowed to connect to the given host ..."

The error message “You are not allowed to connect to the given host through the gateway server because of a Resource Authorization Policy (TS_RAP)” means that a Resource Authorization Policy on the Gateway server prevents your username from connecting to the remote PC. This could be the case:

- You have provided the same name for the remote PC as for the gateway.
- You are then trying to connect to the RD Gateway server itself which is probably not allowed for your username.
- Your Windows user account is not a member of the user groups for remote access on the Remote Desktop server.

If you indeed want to connect to the gateway server itself via the gateway, do not use the external gateway name as PC name, use “localhost” or “127.0.0.1” or the internal server name for the remote desktop. If the entries are correct, ask your system administrator to view the event log on the Terminal server and to check if your user account is a member of the required user groups.

Error message that there is no CAL

When a remote desktop client connects to a Remote Desktop server, the server issues a Remote Desktop Services Client Access License (RDS CAL) stored by the client. Whenever the client connects again it will use its RDS CAL and the server will not issue another license. The server will issue another license if the RDS CAL on the device is missing or corrupt. When the maximum

number of licensed devices is reached the server will not issue new RDS CALs. Contact your network administrator for assistance.

Error: Access Denied

The "Access Denied" error is a generated by the Remote Desktop Gateway and the result of incorrect credentials during the connection attempt. Verify your username and password. If the connection worked before and the error occurred recently, you possibly changed your Windows user account password and haven't updated it yet in the remote desktop settings.

How can I test if VPN is working properly?

Verify that VPN is enabled on your device. You can test your VPN connection by going to a webpage on your internal network or using a web service which is only available via the VPN.

Error: RPC Error 23014 or Error 0x59e6

In case of an **RPC error 23014** or **Error 0x59E6 try again after waiting a few minutes**, the RD Gateway server has reached the maximum number of active connections. Depending on the Windows version running on the RD Gateway the maximum number of connections differs: The Windows Server 2008 R2 Standard implementation limits the number of connections to 250. The Windows Server 2008 R2 Foundation implementation limits the number of connections to 50. All other Windows implementations allow an unlimited number of connections.

See also

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Remote Desktop Client on Android: FAQ
- **iPad or iOS:** Remote Desktop Client on iOS: FAQ
- **Mac:** Remote Desktop Client on Mac: FAQ
- **Windows 8:** Remote Desktop Connection: FAQ

Remote Desktop Client on Mac: FAQ

You can use the Remote Desktop client to connect to a remote PC and your work resources from almost anywhere. Experience rich interactivity using a remote desktop client designed to help you get your work done wherever you are. For example, you can connect to your work PC and have access to all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps using the RD client.

Remote Desktop Client: FAQ

You can use the RD Client to sit at a PC and connect to another PC in a different location (the remote PC). For example, you can connect to your work PC from your home PC and use all of your apps, files, and network resources as if you were sitting right in front of your work PC. You can leave apps open at work and then see those same apps on your home PC.

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Getting Started with Remote Desktop Client on Android
- **iPad or iOS:** Getting Started with Remote Desktop Client on iOS
- **Mac:** Getting Started with Remote Desktop Client on Mac
- Set up your computer running Windows Vista
- Set up your computer running Windows 7
- Set up your computer running Windows 8

How do I set up a PC for Remote Desktop Connection?

On the PC you plan to connect to remotely, do the following:

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
3. In the **System Properties** dialog box, under **Remote Desktop**, select one of the two options. For more info about these options, see What types of Remote Desktop connections should I allow?
4. Tap or click **Select Users**, and in the **Remote Desktop Users** dialog box, tap or click **Add**.
5. In the **Select Users or Groups** dialog box, do the following:
6. Tap or click **Remote settings**. You might be asked for an admin password or to confirm your choice.
    - To select the search location, tap or click **Locations**, and then select the location you want to search.
    - In **Enter the object names to select**, enter the name of the person that you want to add, and then tap or click **OK**. The name will be added to the list of people who can connect to this PC.

Note

You can’t connect to a PC that's asleep or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

When you want to connect to a remote PC, you'll need to know the full name of that PC.

1. Open System by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **System** in the search box, tapping or clicking **Settings**, and then tapping or clicking **System**.
2. The full PC name is listed under **Computer name, domain, and workgroup settings**.

Which PCs can I connect to using Remote Desktop Connection?

You can only connect to PCs that are running these Windows operating systems:

- Windows 8 Enterprise
- Windows 8 Pro
- Windows 7 Professional
- Windows 7 Enterprise
- Windows 7 Ultimate
- Windows Vista Business
- Windows Vista Ultimate
- Windows Vista Enterprise
- Windows XP Professional
- Windows Server 2000
- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Multipoint Server 2011
- Windows Multipoint Server 2012
- Windows Small Business Server 2003
- Windows Small Business Server 2008
- Windows Small Business Server 2011

These Windows operating systems can be used as RD Gateway servers:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Small Business Server 2011

These Windows operating systems can be used to serve as RD Web Access or RemoteApp servers:

- Windows Server 2008 R2

- Windows Server 2012
- Windows Server 2012 R2

Why can’t I connect using Remote Desktop Connection?

Here are some possible solutions to common problems you might encounter when trying to connect to a remote PC. If these solutions don’t work, you can find more help on the Microsoft Community website.

- **The remote PC can’t be found.** If you see this error, talk to the system admin of the remote PC to make sure you have the right PC name, and then check to see if you entered the name correctly. If you still can't connect, try entering the IP address of the remote PC instead of the PC name.
- **There’s a problem with the network.** Make sure that:
  - Your router is turned on (home networks only).
  - The Ethernet cable is plugged into your network adapter (wired networks only).
  - Your PC’s wireless switch is turned on (laptops on wireless networks only).
  - Your network adapter is working.

  If the problem keeps happening, talk to the person who set up your network.
- **The Remote Desktop port might be blocked by a firewall.** If you're using Windows Firewall, follow these steps:
  a. Open Windows Firewall by swiping in from the right edge of the screen, tapping **Search** (or if you're using a mouse, pointing to the upper-right corner of the screen, moving the mouse pointer down, and then clicking **Search**), entering **firewall** in the search box, tapping or clicking **Settings**, and then tapping or clicking **Windows Firewall**.
  b. Tap or click **Allow an app or feature through Windows Firewall**.
  c. Tap or click **Change settings**. You might be asked for an admin password or to confirm your choice.
  d. Under **Allowed apps and features**, select **Remote Desktop**, and then tap or click **OK**.

  If you're using a different firewall, make sure the port for Remote Desktop (usually 3389) is open.
- **Remote connections might not be set up on the remote PC.** To fix this, see the "How do I set up a PC for Remote Desktop Connection?" question in this topic.
- **The remote PC might only allow PCs to connect that have Network Level Authentication set up.** For more info, see What types of Remote Desktop connections should I allow?
- **The remote PC might be turned off.** You can’t connect to a PC that's turned off, asleep, or hibernating, so make sure the settings for sleep and hibernation on the remote PC are set to Never (hibernation isn't available on all PCs.). For more info, see Power plans: Frequently asked questions.

Why can’t I sign in to a remote PC?

If you can see the sign-in screen of the remote PC but you can't sign in, you might not have been added to the Remote Desktop Users Group or to any group with administrator rights on the remote PC. Ask your system admin to do this for you.

Is retina resolution supported?

Retina resolution is supported.

How do I enable secondary right-click on your Mac?

In order to make use of the right-click inside an open session you have three options:

- Standard PC two button USB mouse
- Apple Magic Mouse: To enable right-click, click**System Preferences** in the dock, click **Mouse**, and then enable **Secondary click**.
- Apple Magic Trackpad or MacBook Trackpad: To enable right-click, click **System Preferences** in the dock, click **Mouse**, and then enable **Secondary click**.

How do I make use of all of my monitors?

To enable the usage of two or more screens connected to your Mac follow these steps:

- Right-click the Remote Desktop you want to enable multiple screens for, and then click **Edit**.
- Enable the options **Use all monitors** and **Full screen**.

How do I type various symbols on a Mac keyboard?

The location of many symbols on your keyboard depends on the used keyboard layout and language setting.

I cannot find or connect to my computer

Verify that your computer is on. You can either enter your computers IP address or host name.

Important
Using the host name requires your network to resolve the host name properly through DNS. In many home network configurations only the IP address will work properly.

Verify that you are connecting to a supported Windows version.

Note
Windows XP Home, Windows Media Center Edition, Windows Vista Home and Windows 7 Home or Starter are not supported without 3rd party software.

VPN does not work

These issues can have several reasons. For all reasons it is important to note, that if you can, please check that the VPN works on the same network from your PC or Mac. If you can't test from a PC or Mac you can try to access some company intranet webpage with your device's web browser and see if this works for at least a couple of minutes.

- **L2TP or PPTP VPNs:** If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.
- **VPN is misconfigured:** A misconfigured VPN server can be the reason why the VPN connections never worked or stopped working after some time. Ensure testing with the iOS device's web browser or a PC or Mac on the same network if this happens.

Which connection methods are supported for company networks?

If you want to access your office desktop from outside your company network, your company must provide you with a means of remote access. The RD Client currently supports the following:

- Terminal Server Gateway or Remote Desktop Gateway
- Remote Desktop Web Access
- VPN (through iOS built-in VPN options

Is bi-directional sound supported?

Sound upstream (from client to server, for microphones) is not supported by the Remote Desktop Client.

Why are the characters mixed up in the session?

This issue can occur if the keyboard language has been switched on the remote PC while running a remote desktop session. In the case you would like to switch to a different keyboard language follow these instructions.

Important
The remote PC selects the system language of the device running the remote desktop client and not the keyboard language. When switching the keyboard language on your device, you need to manually switch the keyboard language on the remote PC.

1. Open settings on your device.
2. Tap **General** and then tap **International**.
3. Tap **Language** and select the language running on the remote PC.
4. Open the RD Client and connect to the remote PC and check if the Windows keyboard language matches the language that you selected.

Is the Remote Desktop Client compatible with RDP 7.1?

Yes, the Remote Desktop client supports RDP 7.1 and all previous versions including: RDP 7, RDP 6.1, RDP 6, RDP 5 and RDP 4.

Note
Support for a specific version of RDP 7 does not include all virtual channels. RemoteFX codecs are not supported in RDP 7.1.

How to configure L2TP or PPTP VPN connections

If you are using L2TP or PPTP in your VPN, please set **Send All Traffic** to **ON** in the VPN configuration.

Error: Insufficient privileges

You are not allowed to access the session you want to connect to. The most likely cause is that you are trying to connect to the console session. Only administrators are allowed to connect to the console. Verify that the console switch is off in the settings for the remote desktop. If this is not the source of the problem, please contact your system administrator for further assistance.

Error: Failed to parse NTLM challenge

This error can be caused by a misconfiguration on the remote PC. Verify the RDP security level setting in the remote PC options is set to Client Compatible.

What can I do if the sound does not play?

If sound doesn't play, log out of the session, not just disconnect, and log in again.

Unsupported Windows Versions and Editions

The Remote Desktop Client will not connect to these Windows Versions and Editions:

- Windows XP Starter
- Windows XP Home
- Windows XP Tablet PC
- Windows Media Center
- Windows Vista Starter
- Windows Vista Home
- Windows 7 Starter
- Windows 7 Home

If you want to access computers that have one of these Windows versions installed, we recommend you upgrade to a Windows version that supports RDP.

RD Gateway messaging is not supported

Remote Desktop Client does not support RD Gateway messaging. Verify that the Remote Desktop Resource Access Policy (RD RAP) for your RD Gateway server does not specify **Only allow computers with support for RD Gateway Messaging** or you will not be able to connect.

ERROR: TS_RAP "You are not allowed to connect to the given host ..."

The error message “You are not allowed to connect to the given host through the gateway server because of a Resource Authorization Policy (TS_RAP)” means that a Resource Authorization Policy on the Gateway server prevents your username from connecting to the remote PC. This could be the case:

- You have provided the same name for the remote PC as for the gateway.
- You are then trying to connect to the RD Gateway server itself which is probably not allowed for your username.
- Your Windows user account is not a member of the user groups for remote access on the Remote Desktop server.

If you indeed want to connect to the gateway server itself via the gateway, do not use the external gateway name as PC name, use “localhost” or “127.0.0.1” or the internal server name for the remote desktop. If the entries are correct, ask your system administrator to view the event log on the Terminal server and to check if your user account is a member of the required user groups.

Error message that there is no CAL

When a remote desktop client connects to a Remote Desktop server, the server issues a Remote Desktop Services Client Access License (RDS CAL) stored by the client. Whenever the client connects again it will use its RDS CAL and the server will not issue another license. The server will issue another license if the RDS CAL on the device is missing or corrupt. When the maximum number of licensed devices is reached the server will not issue new RDS CALs. Contact your network administrator for assistance.

Error: Access Denied

The "Access Denied" error is a generated by the Remote Desktop Gateway and the result of incorrect credentials during the connection attempt. Verify your username and password. If the connection worked before and the error occurred recently, you possibly changed your Windows user account password and haven't updated it yet in the remote desktop settings.

How can I test if VPN is working properly?

Verify that VPN is enabled on your device. You can test your VPN connection by going to a webpage on your internal network or using a web service which is only available via the VPN.

Error: RPC Error 23014 or Error 0x59e6

In case of an **RPC error 23014** or **Error 0x59E6 try again after waiting a few minutes**, the RD Gateway server has reached the maximum number of active connections. Depending on the Windows version running on the RD Gateway the maximum number of connections differs: The Windows Server 2008 R2 Standard implementation limits the number of connections to 250. The Windows Server 2008 R2 Foundation implementation limits the number of connections to 50. All other Windows implementations allow an unlimited number of connections.

See also

The following list provides additional resources for using Remote Desktop mobile clients.

- **Android:** Remote Desktop Client on Android: FAQ
- **iPad or iOS:** Remote Desktop Client on iOS: FAQ
- **Mac:** Remote Desktop Client on Mac: FAQ
- **Windows 8:** Remote Desktop Connection: FAQ

Security and Protection

This collection contains descriptions and links to information about changes in security technologies in Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8.

The following table provides links to available information for the IT professional about security technologies and features for Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8. More technologies and features will be added to this table as content becomes available.

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
Access Control	Access Control Overview Access control helps protect files, applications, and other resources from unauthorized use.	The Protected Users security group and Authentication Policy Silos add more credentials protection. They are administered through Active Directory Domain Services. A restricted administration mode is available in the	Added the ability to use dynamic rules-based policies to protect shared folders and files. For more information, see Dynamic Access Control: Scenario Overview Redesigned the Access Control List Editor (ACL editor) to more clearly present key information

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
		Remote Desktop Services (RDS) client. For more information see, Credentials Protection and Management.	needed to assess and manage access control. For more information, see **Enhanced ACL Editor**.
AppLocker	AppLocker Overview AppLocker provides policy-based access control management for applications.	To assist you in process analysis, AppLocker captures command information for each process at runtime, and writes that data to the security log and states, "The system is attempting to launch a process with the following attributes:"	Added functionality to set rules on app packages, which helps manage Windows Store apps. For more information, see Packaged Apps and Packaged App Installer Rules in AppLocker.
BitLocker	BitLocker Overview BitLocker Drive Encryption enables you to encrypt all data that is stored on the operating system volume and configured data volumes for computers running supported versions of Windows. By using a Trusted Platform Module (TPM), it can help ensure the integrity of early startup components.	Broadening support for additional platforms. Recovery password now FIPS-compliant. For more information, see **What's New in BitLocker for Windows 8.1 and Windows Server 2012 R2**.	Added improvements for provisioning and encryption methods, the ability for standard users to change their PINs, support for encrypted hard drives, and a network unlock feature. For more information, see What's New in BitLocker for Windows 8 and Windows Server 2012.
Credential Locker	**Credential Locker Overview** Credential Locker is	Enhancement of credential storage through web	Added ability to program Windows Store apps to use Credential Locker,

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
	managed through the Control Panel by Credential Manager, and supports mostly consumer scenarios.	authentication broker-capable apps, and ability to select a default credential for each site	and improvements to credential roaming (which is set to be disabled for domain-joined computers. For more information, see **New and changed functionality**.
Credentials protection	**Credentials protection and management**. New techniques and features to manage and protect credentials during authentication.	Added additional LSA protection configuration options, new security group, new ways to group users and apply specific authentication policies, For more information, see Credentials Protection and Management	Not available
Encrypted Hard Drive	Encrypted Hard Drive Encrypted Hard Drive is a feature that is provided with BitLocker to enhance data security and management.	Device encryption is available on most editions of Windows. For more information, see **Device encryption**.	Introduced in Windows Server 2012 and Windows 8. For more information, see Support for Encrypted Hard Drives for Windows.
Exchange ActiveSync Policy Engine	Exchange ActiveSync Policy Engine Overview Set of APIs that enable apps to apply EAS policies on desktops, laptops, and tablets to protect data that is synchronized from the cloud, such as data	In certain cases, biometrics sign-in methods are not disabled when the failed-attempts limit is exceeded. For more information, see New and changed functionality.	Introduced in Windows Server 2012.

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
	from Exchange Server.		
Group Managed Service Accounts	Group Managed Service Accounts Overview The group Managed Service Account provides the same functionality as the standalone Managed Service Account within the domain, and it extends that functionality over multiple servers.	No changes.	Added the group Managed Service Account. For more information, see What's New for Managed Service Accounts.
Kerberos	Kerberos Authentication Overview Kerberos protocol is an authentication mechanism that verifies the identity of a user or host.	Change of behavior when the account is in the Protected User security group. For more information, see **Credentials protection and management**.	Reduced authentication failures due to larger service tickets, added changes for developers and IT professionals, changes to smart card sign on KDC validation defaults, and added configuration and maintenance improvements. **Important** For domain-joined devices, the default changed so that now smart card sign-on requires that the KDC certificate chains to a CA in the NTAuth

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
			store. For more information, see What's New in Kerberos Authentication
Local Computer Policy Settings	Security Policy Settings Overview Security policy settings are the configurable rules that the operating system follows when it determines the permissions to grant in response to a request for access to resources. Group Policy Administrative Templates can also be used for security management.	The policy setting System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signingwas changed to reflect changes in the BitLocker recovery password process. For improved process auditing, **Audit Process Creation** was added to the **System** node of **Administrative Templates** under **Computer Configuration.**	Added new security policies to improve security management. For more information, see New and changed functionality.
NTLM	NTLM Overview The NTLM authentication protocols are based on a challenge-and-response mechanism that proves to a server or domain controller that a user knows the password associated with an account.	Change of behavior when the account is in the Protected User security group. For more information, see Protected Users Security Group.	No changes.
Passwords	Passwords Overview The most common	No changes. Microsoft offers other	No changes.

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
	method for authenticating a user's identity is to use a secret passphrase or password as part of the sign-in process.	means for proving identity. For more information, see Smart Card Overview and **Virtual smart cards**.	
Security Auditing	Security Auditing Overview Security auditing can help identify attacks (successful or not) that pose a threat to your network, or attacks against resources that you have determined are of value through a risk assessment.	No changes.	Added expression-based audit policies, and improvements in the ability to audit new types of securable objects and removable storage devices. For more information, see What's New in Security Auditing.
Security Configuration Wizard	Security Configuration Wizard The Security Configuration Wizard is an attack-surface reduction tool that helps administrators create security policies that are based on the minimum functionality required for a server's roles.	No changes.	No changes.
Smart Cards	Smart Card Overview Smart cards provide a tamper-resistant and portable security solution for tasks such as	The process to enroll TPM-enabled devices as a virtual smart card device has improved. APIs are added to simplify the enrollment	Changed the smart card sign-in experience, service start and stop behavior, and smart card transactions, by adding support for

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
	authenticating clients, signing in to domains, signing code, and securing email.	process, making it easier to enroll a device with a virtual smart card regardless of whether they are domain joined and regardless of the hardware.	Windows RT devices and Windows 8 applications. For more information, see What's New in Smart Cards.
Software Restriction Policies	Software Restriction Policies Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run.	No changes.	No changes. Added greater flexibility for AppLocker to control programs in your enterprise. For more information, see AppLocker Overview.
TLS/SSL (Schannel SSP)	TLS/SSL (Schannel SSP) Overview Schannel is a Security Support Provider (SSP) that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols.	Supports server-side "TLS/SSL Session Resumption without Server-Side State extension" (also known as RFC 5077). Addition of the client-side Application Protocol Negotiation For more information, see New and changed functionality in Windows Server 2012 R2.	Changed how trusted issuers for client authentication can be managed, added TLS support for Server Name Indicator (SNI) Extensions, and added Datagram Transport Layer Security (DTLS) for the provider. For more information, see New and changed functionality in Windows Server 2012.
Trusted Platform Module (TPM)	Trusted Platform Module Technology Overview	Improvements to the TPM Key Storage Provider for platform	Improved administration and functionality, including automated provisioning and

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
	Trusted Platform Module (TPM) technology is designed to provide hardware-based, security-related functions.	and key attestation. For more information, see **Malware resistance** and **What's New for the TPM in Windows 8.1**.	management, Measured Boot with support for attestation, TPM-based Virtual Smart Card, and secure storage for critical elements. For more information, see New and changed functionality.
User Account Control (UAC)	User Account Control Overview UAC helps mitigate the impact of malicious programs.	No changes.	Refined to allow easier administration of UAC configuration and messages. For more information, see New and changed functionality.
Virtual Smart Card	Virtual smart cards offer multifactor authentication and compatibility with many smart card infrastructures, and offer users the convenience of not having to carry a physical card, so users are more likely to follow their organization's security guidelines rather than working around them. Understanding and Evaluating Virtual Smart Cards	The process to enroll TPM-enabled devices as a virtual smart card device has improved. APIs are added to simplify the enrollment process, making it easier to enroll a device with a virtual smart card regardless of whether they are domain joined and regardless of the hardware. For more information, see **Virtual smart cards**	Introduced in Windows Server 2012.
Windows Biometric Framework and **Windows**	**Windows Biometric Framework Overview**	Enhanced the client and associated APIs. For more information,	Better integration of fingerprint readers with Fast User Switching,

Feature or Technology	Overview	What changed in Windows Server 2012 R2	What changed in Windows Server 2012
Biometrics	The Windows Biometric Framework (WBF) is a set of services and interfaces that permit consistent development and management of biometric devices, such as fingerprint readers. WBF improves the reliability and compatibility with biometric services and drivers.	see **Fingerprint biometrics**.	and synchronization of passwords with fingerprints. For more information, see **New and changed functionality**
Windows Defender	Windows Defender is a full-featured antimalware solution that is capable of detecting and stopping a wider range of potentially malicious software, including viruses.	Available and enabled by default on Server Core installation options and Core System Server (without the user interface). For more information, see **Windows Defender**.	Upgraded from antispyware to a full-featured antimalware solution that is capable of detecting and stopping a wider range of potentially malicious software, including viruses.

See also

Secure Windows Server 2012 R2 and Windows Server 2012

Secure Windows 8

What's Changed in Security Technologies in Windows 8.1

Access Control Overview

This topic for the IT professional describes access control in Windows, which is the process of authorizing users, groups, and computers to access objects on the network or computer. Key concepts that make up access control are permissions, ownership of objects, inheritance of permissions, user rights, and object auditing.

Feature description

Computers that are running a supported version of Windows can control the use of system and network resources through the interrelated mechanisms of authentication and authorization. After a user is authenticated, the Windows operating system uses built-in authorization and access control technologies to implement the second phase of protecting resources: determining if an authenticated user has the correct permissions to access a resource.

Shared resources are available to users and groups other than the resource's owner, and they need to be protected from unauthorized use. In the access control model, users and groups (also referred to as security principals) are represented by unique security identifiers (SIDs). They are assigned rights and permissions that inform the operating system what each user and group can do. Each resource has an owner who grants permissions to security principals. During the access control check, these permissions are examined to determine which security principals can access the resource and how they can access it.

Security principals perform actions (which include Read, Write, Modify, or Full control) on objects. Objects include files, folders, printers, registry keys, and Active Directory Domain Services (AD DS) objects. Shared resources use access control lists (ACLs) to assign permissions. This enables resource managers to enforce access control in the following ways:

- Deny access to unauthorized users and groups
- Set well-defined limits on the access that is provided to authorized users and groups

Object owners generally grant permissions to security groups rather than to individual users. Users and computers that are added to existing groups assume the permissions of that group. If an object (such as a folder) can hold other objects (such as subfolders and files), it is called a container. In a hierarchy of objects, the relationship between a container and its content is expressed by referring to the container as the parent. An object in the container is referred to as the child, and the child inherits the access control settings of the parent. Object owners often define permissions for container objects, rather than individual child objects, to ease access control management.

Practical applications

Administrators who use the supported version of Windows can refine the application and management of access control to objects and subjects to provide the following security:

- Protect a greater number and variety of network resources from misuse.

- Provision users to access resources in a manner that is consistent with organizational policies and the requirements of their jobs.
- Enable users to access resources from a variety of devices in numerous locations.
- Update users' ability to access resources on a regular basis as an organization's policies change or as users' jobs change.
- Account for a growing number of use scenarios (such as access from remote locations or from a rapidly expanding variety of devices, such as tablet computers and mobile phones).
- Identify and resolve access issues when legitimate users are unable to access resources that they need to perform their jobs.

Permissions

Permissions define the type of access that is granted to a user or group for an object or object property. For example, the Finance group can be granted Read and Write permissions for a file named Payroll.dat.

By using the access control user interface, you can set NTFS permissions for objects such as files, Active Directory objects, registry objects, or system objects such as processes. Permissions can be granted to any user, group, or computer. It is a good practice to assign permissions to groups because it improves system performance when verifying access to an object.

For any object, you can grant permissions to:

- Groups, users, and other objects with security identifiers in the domain.
- Groups and users in that domain and any trusted domains.
- Local groups and users on the computer where the object resides.

The permissions attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. Some permissions, however, are common to most types of objects. These common permissions are:

- Read
- Modify
- Change owner
- Delete

When you set permissions, you specify the level of access for groups and users. For example, you can let one user read the contents of a file, let another user make changes to the file, and prevent all other users from accessing the file. You can set similar permissions on printers so that certain users can configure the printer and other users can only print.

When you need to change the permissions on a file, you can run Windows Explorer, right-click the file name, and click **Properties**. On the **Security** tab, you can change permissions on the file. For more information, see Managing Permissions.

Note

Another kind of permissions, called share permissions, is set on the Sharing tab of a folder's **Properties** page or by using the Shared Folder Wizard. For more information see Share and NTFS Permissions on a File Server.

Ownership of objects

An owner is assigned to an object when that object is created. By default, the owner is the creator of the object. No matter what permissions are set on an object, the owner of the object can always change the permissions. For more information, see Manage Object Ownership.

Inheritance of permissions

Inheritance allows administrators to easily assign and manage permissions. This feature automatically causes objects within a container to inherit all the inheritable permissions of that container. For example, the files within a folder inherit the permissions of the folder. Only permissions marked to be inherited will be inherited.

User rights

User rights grant specific privileges and sign-in rights to users and groups in your computing environment. Administrators can assign specific rights to group accounts or to individual user accounts. These rights authorize users to perform specific actions, such as signing in to a system interactively or backing up files and directories.

User rights are different from permissions because user rights apply to user accounts, and permissions are associated with objects. Although user rights can apply to individual user accounts, user rights are best administered on a group account basis. There is no support in the access control user interface to grant user rights. However, user rights assignment can be administered through **Local Security Settings**.

For more information about user rights, see User Rights Assignment.

Object auditing

With administrator's rights, you can audit users' successful or failed access to objects. You can select which object access to audit by using the access control user interface, but first you must enable the audit policy by selecting **Audit object access** under **Local Policies** in **Local Security Settings**. You can then view these security-related events in the Security log in Event Viewer.

For more information about auditing, see Security Auditing Overview.

Additional resources

- For more information about authorization and access control, see Windows Security Collection.

- For information about authorization strategy, see Designing a Resource Authorization Strategy.

Dynamic Access Control Overview

This overview topic for the IT professional describes Dynamic Access Control and its associated elements, which were introduced in Windows Server 2012 and Windows 8.

Domain-based Dynamic Access Control enables administrators to apply access-control permissions and restrictions based on well-defined rules that can include the sensitivity of the resources, the job or role of the user, and the configuration of the device that is used to access these resources.

For example, a user might have different permissions when they access a resource from their office computer versus when they are using a portable computer over a virtual private network. Or access may be allowed only if a device meets the security requirements that are defined by the network administrators. When Dynamic Access Control is used, a user's permissions change dynamically without additional administrator intervention if the user's job or role changes (resulting in changes to the user's account attributes in AD DS).

Dynamic Access Control is not supported in Windows operating systems prior to Windows Server 2012 and Windows 8. When Dynamic Access Control is configured in environments with supported and non-supported versions of Windows, only the supported versions will implement the changes.

Features and concepts associated with Dynamic Access Control include:

- Central access rules
- Central access policies
- Claims
- Expressions
- Proposed permissions

Central access rules

A central access rule is an expression of authorization rules that can include one or more conditions involving user groups, user claims, device claims, and resource properties. Multiple central access rules can be combined into a central access policy.

If one or more central access rules have been defined for a domain, file share administrators can match specific rules to specific resources and business requirements.

Central access policies

Central access policies are authorization policies that include conditional expressions. For example, let's say an organization has a business requirement to restrict access to personally

identifiable information (PII) in files to only the file owner and members of the human resources (HR) department who are allowed to view PII information. This represents an organization-wide policy that applies to PII files wherever they are located on file servers across the organization. To implement this policy, an organization needs to be able to:

- Identify and mark the files that contain the PII.
- Identify the group of HR members who are allowed to view the PII information.
- Add the central access policy to a central access rule, and apply the central access rule to all files that contain the PII, wherever they are located amongst the file servers across the organization.

Central access policies act as security umbrellas that an organization applies across its servers. These policies are in addition to (but do not replace) the local access policies or discretionary access control lists (DACLs) that are applied to files and folders.

Claims

A claim is a unique piece of information about a user, device, or resource that has been published by a domain controller. The user's title, the department classification of a file, or the health state of a computer are valid examples of a claim. An entity can involve more than one claim, and any combination of claims can be used to authorize access to resources. The following types of claims are available in the supported versions of Windows:

- **User claims** Active Directory attributes that are associated with a specific user.
- **Device claims** Active Directory attributes that are associated with a specific computer object.
- **Resource attributes** Global resource properties that are marked for use in authorization decisions and published in Active Directory.

Claims make it possible for administrators to make precise organization- or enterprise-wide statements about users, devices, and resources that can be incorporated in expressions, rules, and policies.

Expressions

Conditional expressions are an enhancement to access control management that allow or deny access to resources only when certain conditions are met, for example, group membership, location, or the security state of the device. Expressions are managed through the Advanced Security Settings dialog box of the ACL Editor or the Central Access Rule Editor in the Active Directory Administrative Center (ADAC).

Expressions help administrators manage access to sensitive resources with flexible conditions in increasingly complex business environments.

Proposed permissions

Proposed permissions enable an administrator to more accurately model the impact of potential changes to access control settings without actually changing them.

Predicting the effective access to a resource helps you plan and configure permissions for those resources before implementing those changes.

Additional changes

Additional enhancements in the supported versions of Windows that support Dynamic Access Control include:

Support in the Kerberos authentication protocol to reliably provide user claims, device claims, and device groups.

By default, devices running any of the supported versions of Windows are able to process Dynamic Access Control-related Kerberos tickets, which include data needed for compound authentication. Domain controllers are able to issue and respond to Kerberos tickets with compound authentication-related information. When a domain is configured to recognize Dynamic Access Control, devices receive claims from domain controllers during initial authentication, and they receive compound authentication tickets when submitting service ticket requests. Compound authentication results in an access token that includes the identity of the user and the device on the resources that recognize Dynamic Access Control.

Support for using the Key Distribution Center (KDC) Group Policy setting to enable Dynamic Access Control for a domain.

Every domain controller needs to have the same Administrative Template policy setting, which is located at **Computer Configuration\Policies\Administrative Templates\System\KDC\Support Dynamic Access Control and Kerberos armoring**.

Support for using the Key Distribution Center (KDC) Group Policy setting to enable Dynamic Access Control for a domain.

Every domain controller needs to have the same Administrative Template policy setting, which is located at **Computer Configuration\Policies\Administrative Templates\System\KDC\Support Dynamic Access Control and Kerberos armoring**.

Support in Active Directory to store user and device claims, resource properties, and central access policy objects.

Support for using Group Policy to deploy central access policy objects.

The following Group Policy setting enables you to deploy central access policy objects to file servers in your organization: **Computer Configuration\Policies\ Windows Settings\Security Settings\File System\Central Access Policy**.

Support for claims-based file authorization and auditing for file systems by using Group Policy and Global Object Access Auditing

You must enable staged central access policy auditing to audit the effective access of central access policy by using proposed permissions. You configure this setting for the computer under **Advanced Audit Policy Configuration** in the **Security Settings** of a Group Policy Object (GPO). After you configure the security setting in the GPO, you can deploy the GPO to computers in your network.

Support for transforming or filtering claim policy objects that traverse Active Directory forest trusts

You can filter or transform incoming and outgoing claims that traverse a forest trust. There are three basic scenarios for filtering and transforming claims:

- **Value-based filtering** Filters can be based on the value of a claim. This allows the trusted forest to prevent claims with certain values from being sent to the trusting forest. Domain controllers in trusting forests can use value-based filtering to guard against an elevation-of-privilege attack by filtering the incoming claims with specific values from the trusted forest.
- **Claim type-based filtering** Filters are based on the type of claim, rather than the value of the claim. You identify the claim type by the name of the claim. You use claim type-based filtering in the trusted forest, and it prevents Windows from sending claims that disclose information to the trusting forest.
- **Claim type-based transformation** Manipulates a claim before sending it to the intended target. You use claim type-based transformation in the trusted forest to generalize a known claim that contains specific information. You can use transformations to generalize the claim-type, the claim value, or both.

Software requirements

Because claims and compound authentication for Dynamic Access Control require Kerberos authentication extensions, any domain that supports Dynamic Access Control must have enough domain controllers running the supported versions of Windows to support authentication from Dynamic Access Control-aware Kerberos clients. By default, devices must use domain controllers in other sites. If no such domain controllers are available, authentication will fail. Therefore, you must support one of the following conditions:

- Every domain that supports Dynamic Access Control must have enough domain controllers running the supported versions of Windows Server to support authentication from all devices running the supported versions of Windows or Windows Server.
- Devices running the supported versions of Windows or that do not protect resources by using claims or compound identity, should disable Kerberos protocol support for Dynamic Access Control.

For domains that support user claims, every domain controller running the supported versions of Windows server must be configured with the appropriate setting to support claims and compound

authentication, and to provide Kerberos armoring. Configure settings in the KDC Administrative Template policy as follows:

- **Always provide claims** Use this setting if all domain controllers are running the supported versions of Windows Server. In addition, set the domain functional level to Windows Server 2012 or higher.
- **Supported** When you use this setting, monitor domain controllers to ensure that the number of domain controllers running the supported versions of Windows Server is sufficient for the number of client computers that need to access resources protected by Dynamic Access Control.

If the user domain and file server domain are in different forests, all domain controllers in the file server's forest root must be set at the Windows Server 2012 or higher functional level.

If clients do not recognize Dynamic Access Control, there must be a two-way trust relationship between the two forests.

If claims are transformed when they leave a forest, all domain controllers in the user's forest root must be set at the Windows Server 2012 or higher functional level.

A file server running Windows Server 2012 or Windows Server 2012 R2 must have a Group Policy setting that specifies whether it needs to get user claims for user tokens that do not carry claims. This setting is set by default to **Automatic**, which results in this Group Policy setting to be turned **On** if there is a central policy that contains user or device claims for that file server. If the file server contains discretionary ACLs that include user claims, you need to set this Group Policy to **On** so that the server knows to request claims on behalf of users that do not provide claims when they access the server.

Additional resource

For information about implementing solutions based on this technology, see Dynamic Access Control: Scenario Overview.

Security Principals Technical Overview

This reference topic for the IT professional describes security principals in regards to Windows accounts and security groups, in addition to security technologies that are related to security principals.

What are security principals?

Security principals are any entity that can be authenticated by the operating system, such as a user account, a computer account, or a thread or process that runs in the security context of a user or computer account, or the security groups for these accounts. Security principals have long been a foundation for controlling access to securable resources on Windows computers.

Each security principal is represented in the operating system by a unique security identifier (SID).

The following content applies to the versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

How security principals work

Security principals that are created in an Active Directory domain are Active Directory objects, which can be used to manage access to domain resources. Each security principal is assigned a unique identifier, which it retains for its entire lifetime. Local user accounts and security groups are created on a local computer, and they can be used to manage access to resources on that computer. Local user accounts and security groups are managed by the Security Accounts Manager (SAM) on the local computer.

Authorization and access control components

The following diagram illustrates the Windows authorization and access control process. In this diagram, the subject (a process that is initiated by a user) attempts to access an object, such as a shared folder. The information in the user's access token is compared to the access control entries (ACEs) in the object's security descriptor, and the access decision is made. The SIDs of security principals are used in the user's access token and in the ACEs in the object's security descriptor.

Authorization and access control process

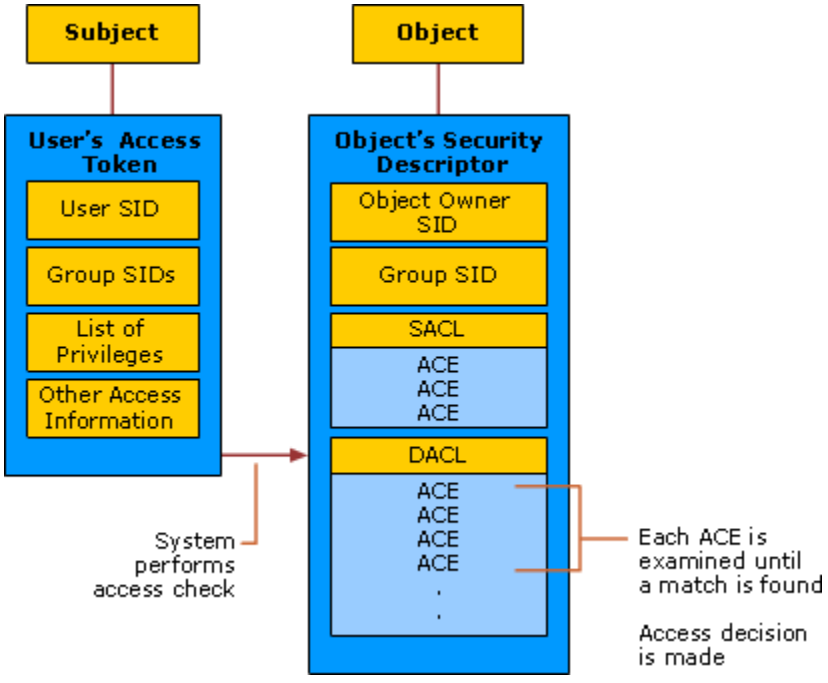

Security principals are closely related to the following components and technologies:

- Security identifiers
- Access tokens
- Security descriptors and access control lists
- Permissions

Security identifiers

Security identifiers (SIDs) provide a fundamental building block of the Windows security model. They work with specific components of the authorization and access control technologies in the security infrastructure of the Windows Server operating systems. This helps protect access to network resources and provides a more secure computing environment.

A SID is a value of variable length that is used to uniquely identify a security principal that represents any entity that can be authenticated by the system. These entities include a user account, a computer account, or a thread or process that runs in the security context of a user or computer account. Each security principal is automatically assigned a SID when it is created. The SID is stored in a security database. When a SID is used as the unique identifier for a user or group, it can never be used to identify another user or group.

Each time a user signs in, the system creates an access token for that user. The access token contains the user's SID, user rights, and the SIDs for groups that the user belongs to. This token provides the security context for whatever actions the user performs on that computer.

In addition to the uniquely created, domain-specific SIDs that are assigned to specific users and groups, there are well-known SIDs that identify generic groups and generic users. For example, the Everyone and the World SIDs identify groups that includes all users. Well-known SIDs have values that remain constant across all operating systems.

Access tokens

An access token is a protected object that contains information about the identity and user rights that are associated with a user account.

When a user signs in interactively or tries to make a network connection to a computer running Windows, the sign-in process authenticates the user's credentials. If authentication is successful, the process returns a SID for the user and a list of SIDs for the user's security groups. The Local Security Authority (LSA) on the computer uses this information to create an access token (in this case, the primary access token). This includes the SIDs that are returned by the sign-in process and a list of user rights that are assigned by the local security policy to the user and to the user's security groups.

After the LSA creates the primary access token, a copy of the access token is attached to every thread and process that executes on the user's behalf. Whenever a thread or process interacts with a securable object or tries to perform a system task that requires user rights, the operating system checks the access token that is associated with the thread to determine the level of authorization.

There are two kinds of access tokens, primary and impersonation. Every process has a primary token that describes the security context of the user account that is associated with the process.

A primary access token is typically assigned to a process to represent the default security information for that process. Impersonation tokens, on the other hand, are usually used for client and server scenarios. Impersonation tokens enable a thread to run in a security context that differs from the security context of the process that owns the thread.

Security descriptors and access control lists

A security descriptor is a data structure that is associated with each securable object. All objects in Active Directory and all securable objects on a local computer or on the network have security descriptors to help control access to the objects. Security descriptors include information about who owns an object, who can access it and in what way, and what types of access are audited. Security descriptors contain the access control list (ACL) of an object, which includes all of the security permissions that apply to that object. An object's security descriptor can contain two types of ACLs:

- A discretionary access control list (DACL), which identifies the users and groups who are allowed or denied access
- A system access control list (SACL), which controls how access is audited

You can use this access control model to individually secure objects and attributes such as files and folders, Active Directory objects, registry keys, printers, devices, ports, services, processes, and threads. Because of this individual control, you can adjust the security of objects to meet the needs of your organization, delegate authority over objects or attributes, and create custom objects or attributes that require unique security protections to be defined.

Permissions

Permissions enable the owner of each securable object, such as a file, Active Directory object, or registry key, to control who can perform an operation or a set of operations on the object or object property. Permissions are expressed in the security architecture as access control entries (ACEs). Because access to an object is at the discretion of the object's owner, the type of access control that is used in Windows is called discretionary access control.

Permissions are different from user rights in that permissions are attached to objects, and user rights apply to user accounts. Administrators can assign user rights to groups or users. These rights authorize users to perform specific actions, such as signing in to a system interactively or backing up files and directories.

On computers, user rights enable administrators to control who has the authority to perform operations that affect an entire computer, rather than a particular object. Administrators assign user rights to individual users or groups as part of the security settings for the computer. Although user rights can be managed centrally through Group Policy, they are applied locally. Users can (and usually do) have different user rights on different computers.

For information about which user rights are available and how they can be implemented, see User Rights Assignment.

Security context in authentication

A user account enables a user to sign in to computers, networks, and domains with an identity that can be authenticated by the computer, network, or domain.

In Windows, any user, service, group, or computer that can initiate action is a security principal. Security principals have accounts, which can be local to a computer or domain-based. For example, domain-joined Windows client computers can participate in a network domain by communicating with a domain controller, even when no user is signed in.

To initiate communications, the computer must have an active account in the domain. Before accepting communications from the computer, the Local Security Authority on the domain controller authenticates the computer's identity and then defines the computer's security context just as it would for a user's security principal.

This security context defines the identity and capabilities of a user or service on a particular computer, or of a user, service, group or computer on a network. For example, it defines the resources (such as a file share or printer) that can be accessed and the actions (such as Read, Write, or Modify) that can be performed by a user, service, or computer on that resource.

The security context of a user or computer can vary from one computer to another, such as when a user authenticates to a server or a workstation other than the user's primary workstation. It can also vary from one session to another, such as when an administrator modifies the user's rights and permissions. In addition, the security context is usually different when a user or computer is operating on a stand-alone basis, in a mixed network domain, or as part of an Active Directory domain.

Accounts and security groups

Accounts and security groups that are created in an Active Directory domain are stored in the Active Directory database and managed by using Active Directory tools. These security principals are directory objects, and they can be used to manage access to domain resources.

Local user accounts and security groups are created on a local computer, and they can be used to manage access to resources on that computer. Local user accounts and security groups are stored in and managed by the Security Accounts Manager (SAM) on the local computer.

User accounts

A user account uniquely identifies a person who is using a computer system. The account signals the system to enforce the appropriate authorization to allow or deny that user access to resources. User accounts can be created in Active Directory and on local computers, and administrators use them to:

- Represent, identify, and authenticate the identity of a user. A user account enables a user to sign in to computers, networks, and domains with a unique identifier that can be authenticated by the computer, network, or domain.

- Authorize (grant or deny) access to resources. After a user has been authenticated, the user is authorized access to resources based on the permissions that are assigned to that user for the resource.
- Audit the actions that are carried out on a user account.

Windows and the Windows Server operating systems have built-in user accounts, or you can create user accounts to meet the requirements of your organization.

Security groups

A security group is a collection of user accounts, computer accounts, and other groups of accounts that can be managed as a single unit from a security perspective. In Windows operating systems, there are several built-in security groups that are preconfigured with the appropriate rights and permissions for performing specific tasks. Additionally, you can (and, typically, will) create a security group for each unique combination of security requirements that applies to multiple users in your organization.

Groups can be Active Directory-based or local to a particular computer:

- Active Directory security groups are used to manage rights and permissions to domain resources.
- Local groups exist in the SAM database on local computers (on all Windows-based computers) except domain controllers. You use local groups to manage rights and permissions only to resources on the local computer.

By using security groups to manage access control, you can:

- Simplify administration. You can assign a common set of rights, a common set of permissions, or both to many accounts at one time, rather than assigning them to each account individually. Also, when users transfer jobs or leave the organization, permissions are not tied to their user accounts, making permission reassignment or removal easier.
- Implement a role-based access-control model. You can use this model to grant permissions by using groups with different scopes for appropriate purposes. Scopes that are available in Windows include local, global, domain local, and universal.
- Minimize the size of access control lists (ACLs) and speed security checking. A security group has its own SID; therefore, the group SID can be used to specify permissions for a resource. In an environment with more than a few thousand users, if the SIDs of individual user accounts are used to specify access to a resource, the ACL of that resource can become unmanageably large, and the time that is needed for the system to check permissions to the resource can become unacceptable.

For descriptions and settings information about the domain security groups that are defined in Active Directory, see Active Directory Security Groups.

For descriptions and settings information about the Special Identities group, see Special Identities.

Service Accounts

This topic for the IT professional explains group and standalone managed service accounts, and the computer-specific virtual computer account, and it points to resources about these service accounts.

Overview

A service account is a user account that is created explicitly to provide a security context for services running on Windows Server operating systems. The security context determines the service's ability to access local and network resources. The Windows operating systems rely on services to run various features. These services can be configured through the applications, the Services snap-in, or Task Manager, or by using Windows PowerShell.

This topic contains information about the following types of service accounts:

- Standalone managed service accounts
- Group managed service accounts
- Virtual accounts

Standalone managed service accounts

A managed service account is designed to isolate domain accounts in crucial applications, such as Internet Information Services (IIS), and eliminate the need for an administrator to manually administer the service principal name (SPN) and credentials for the accounts.

To use managed service accounts, the server on which the application or service is installed must be running at least Windows Server 2008 R2. One managed service account can be used for services on a single computer. Managed service accounts cannot be shared between multiple computers, and they cannot be used in server clusters where a service is replicated on multiple cluster nodes. For this scenario, you must use a group managed service account. For more information, see Group Managed Service Accounts Overview.

In addition to the enhanced security that is provided by having individual accounts for critical services, there are four important administrative benefits associated with managed service accounts:

- You can create a class of domain accounts that can be used to manage and maintain services on local computers.
- Unlike domain accounts in which administrators must reset manually passwords, the network passwords for these accounts are automatically reset.
- You do not have to complete complex SPN management tasks to use managed service accounts.
- Administrative tasks for managed service accounts can be delegated to non-administrators.

Software requirements

Managed service accounts apply to the Windows operating systems that are designated in the **Applies To** list at the beginning of this topic.

Group managed service accounts

Group managed service accounts are an extension of the standalone managed service accounts, which were introduced in Windows Server 2008 R2. These are managed domain accounts that provide automatic password management and simplified service principal name (SPN) management, including delegation of management to other administrators.

The group managed service account provides the same functionality as a standalone managed service account within the domain, but it extends that functionality over multiple servers. When connecting to a service that is hosted on a server farm, such as Network Load Balancing, the authentication protocols that support mutual authentication require all instances of the services to use the same principal. When group managed service accounts are used as service principals, the Windows Server operating system manages the password for the account instead of relying on the administrator to manage the password.

The Microsoft Key Distribution Service (kdssvc.dll) provides the mechanism to securely obtain the latest key or a specific key with a key identifier for an Active Directory account. This service was introduced in Windows Server 2012, and it does not run on previous versions of the Windows Server operating system. The Key Distribution Service shares a secret, which is used to create keys for the account. These keys are periodically changed. For a group managed service account, the domain controller computes the password on the key that is provided by the Key Distribution Services, in addition to other attributes of the group managed service account.

Practical applications

Group managed service accounts provide a single identity solution for services running on a server farm, or on systems that use Network Load Balancing. By providing a group managed service account solution, services can be configured for the group managed service account principal, and the password management is handled by the operating system.

By using a group managed service account, services or service administrators do not need to manage password synchronization between service instances. The group managed service account supports hosts that are kept offline for an extended time period and the management of member hosts for all instances of a service. This means that you can deploy a server farm that supports a single identity to which existing client computers can authenticate without knowing the instance of the service to which they are connecting.

Failover clusters do not support group managed service account s. However, services that run on top of the Cluster service can use a group managed service account or a standalone managed service account if they are a Windows service, an App pool, a scheduled task, or if they natively support group managed service account or standalone managed service accounts.

Software requirements

Group managed service accounts can only be configured and administered on computers running at least Windows Server 2012, but they can be deployed as a single service identity solution in domains that still have domain controllers running operating systems earlier than Windows Server 2012. There are no domain or forest functional level requirements.

A 64-bit architecture is required to run the Windows PowerShell commands that are used to administer group managed service accounts.

A managed service account is dependent on encryption types supported by Kerberos. When a client computer authenticates to a server by using Kerberos protocol, the domain controller creates a Kerberos service ticket that is protected with encryption that the domain controller and the server support. The domain controller uses the account's **msDS-SupportedEncryptionTypes** attribute to determine what encryption the server supports, and if there is no attribute, it assumes that the client computer does not support stronger encryption types. The Advanced Encryption Standard (AES) should always be explicitly configured for managed service accounts. If computers that host the managed service account are configured to not support RC4, authentication will always fail.

Note

Introduced in Windows Server 2008 R2, the Data Encryption Standard (DES) is disabled by default. For more information about supported encryption types, see Changes in Kerberos Authentication.

Group managed service accounts are not applicable in Windows operating systems prior to Windows Server 2012.

Virtual accounts

Virtual accounts were introduced in Windows Server 2008 R2 and Windows 7, and are managed local accounts that provide the following features to simplify service administration:

- The virtual account is automatically managed.
- The virtual account can access the network in a domain environment.
- No password management is required. For example, if the default value is used for the service accounts during SQL Server setup on Windows Server 2008 R2, a virtual account that uses the instance name as the service name is established in the format NT SERVICE\<SERVICENAME>.

Services that run as virtual accounts access network resources by using the credentials of the computer account in the format <domain_name>\<computer_name>$.

For information about how to configure and use virtual service accounts, see Service Accounts Step-by-Step Guide.

Software requirements

Virtual accounts apply to the Windows operating systems that are designated in the **Applies To** list at the beginning of this topic.

See also

The following table provides links to additional resources that are related to standalone managed service accounts, group managed service accounts, and virtual accounts.

Content type	References
Product evaluation	What's New for Managed Service Accounts Managed Service Accounts Documentation for Windows 7 and Windows Server 2008 R2 Service Accounts Step-by-Step Guide Getting Started with Group Managed Service Accounts
Deployment	Windows Server 2012: Group Managed Service Accounts - Ask Premier Field Engineering (PFE) Platforms - Site Home - TechNet Blogs
Operations	Managed Service Accounts in Active Directory
Tools and settings	Managed Service Accounts in Active Directory Domain Services
Community resources	Managed Service Accounts: Understanding, Implementing, Best Practices, and Troubleshooting
Related technologies	Security Principals Technical Overview Active Directory Domain Services Overview

Special Identities

This reference topic for the IT professional describes the special identity groups (which are sometimes referred to as security groups) that are used in Windows access control.

Special identity groups are similar to Active Directory security groups as listed in the users and built-in containers. Special identity groups can provide an efficient way to assign access to resources in your network. By using special identity groups, you can:

- Assign user rights to security groups in Active Directory.
- Assign permissions to security groups for the purpose of accessing resources.

Servers that are running the supported Windows Server operating systems designated in the **Applies To** list at the beginning of this topic include several special identity groups. These special

identity groups do not have specific memberships that can be modified, but they can represent different users at different times, depending on the circumstances.

Although the special identity groups can be assigned rights and permissions to resources, the memberships cannot be modified or viewed. Group scopes do not apply to special identity groups. Users are automatically assigned to these special identity groups whenever they sign in or access a particular resource.

For information about security groups and group scope, see Active Directory Security Groups.

The special identity groups are described in the following tables.

- Anonymous Logon
- Authenticated Users
- Batch
- Creator Group
- Creator Owner
- Dialup
- Digest Authentication
- Enterprise Domain Controllers
- Everyone
- Interactive
- Local Service
- LocalSystem
- Network
- Network Service
- NTLM Authentication
- Other Organization
- Principal Self
- Remote Interactive Logon
- Restricted
- SChannel Authentication
- Service
- Terminal Server User
- This Organization
- Window Manager\Window Manager Group

Anonymous Logon

Any user who accesses the system through an anonymous logon has the Anonymous Logon identity. This identity allows anonymous access to resources, such as a web page that is published on corporate servers. The Anonymous Logon group is not a member of the Everyone group by default.

Attribute	Value
Well-Known SID/RID	S-1-5-7
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Authenticated Users

Any user who accesses the system through a sign-in process has the Authenticated Users identity. This identity allows access to shared resources within the domain, such as files in a shared folder that should be accessible to all the workers in the organization. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-11
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	Access this computer from the network: SeNetworkLogonRight Add workstations to domain: SeMachineAccountPrivilege Bypass traverse checking: SeChangeNotifyPrivilege

Batch

Any user or process that accesses the system as a batch job (or through the batch queue) has the Batch identity. This identity allows batch jobs to run scheduled tasks, such as a nightly cleanup job that deletes temporary files. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-3
Object Class	Foreign Security Principal

Attribute	Value
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Creator Group

The person who created the file or the directory is a member of this special identity group. Windows Server operating systems use this identity to automatically grant access permissions to the creator of a file or directory.

A placeholder security identifier (SID) is created in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces this SID with the SID for the primary group of the object's current owner. The primary group is used only by the Portable Operating System Interface for UNIX (POSIX) subsystem.

Attribute	Value
Well-Known SID/RID	S-1-3-1
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Creator Owner

The person who created the file or the directory is a member of this special identity group. Windows Server operating systems use this identity to automatically grant access permissions to the creator of a file or directory. A placeholder SID is created in an inheritable ACE. When the ACE is inherited, the system replaces this SID with the SID for the object's current owner.

Attribute	Value
Well-Known SID/RID	S-1-3-0
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Dialup

Any user who accesses the system through a dial-up connection has the Dial-Up identity. This identity distinguishes dial-up users from other types of authenticated users.

Attribute	Value
Well-Known SID/RID	S-1-5-1
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Digest Authentication

Attribute	Value
Well-Known SID/RID	S-1-5-64-21
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Enterprise Domain Controllers

This group includes all domain controllers in an Active Directory forest. Domain controllers with enterprise-wide roles and responsibilities have the Enterprise Domain Controllers identity. This identity allows them to perform certain tasks in the enterprise by using transitive trusts. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-9
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights Assignment	Access this computer from the network:

Attribute	Value
	SeNetworkLogonRight Allow log on locally: SeInteractiveLogonRight

Everyone

All interactive, network, dial-up, and authenticated users are members of the Everyone group. This special identity group gives wide access to system resources. Whenever a user logs on to the network, the user is automatically added to the Everyone group.

On computers running Windows 2000 and earlier, the Everyone group included the Anonymous Logon group as a default member, but as of Windows Server 2003, the Everyone group contains only Authenticated Users and Guest; and it no longer includes Anonymous Logon by default (although this can be changed).

Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-1-0
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	Access this computer from the network: SeNetworkLogonRight Act as part of the operating system: SeTcbPrivilege Bypass traverse checking: SeChangeNotifyPrivilege

Interactive

Any user who is logged on to the local system has the Interactive identity. This identity allows only local users to access a resource. Whenever a user accesses a given resource on the computer to which they are currently logged on, the user is automatically added to the Interactive group. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-4

Attribute	Value
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Local Service

The Local Service account is similar to an Authenticated User account. The Local Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Local Service account access network resources as a null session with anonymous credentials. The name of the account is NT AUTHORITY\LocalService. This account does not have a password.

Attribute	Value
Well-Known SID/RID	S-1-5-19
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default user rights	Adjust memory quotas for a process: SeIncreaseQuotaPrivilege Bypass traverse checking: SeChangeNotifyPrivilege Change the system time: SeSystemtimePrivilege Change the time zone: SeTimeZonePrivilege Create global objects: SeCreateGlobalPrivilege Generate security audits: SeAuditPrivilege Impersonate a client after authentication: SeImpersonatePrivilege Replace a process level token: SeAssignPrimaryTokenPrivilege

LocalSystem

This is a service account that is used by the operating system. The LocalSystem account is a powerful account that has full access to the system and acts as the computer on the network. If a service logs on to the LocalSystem account on a domain controller, that service has access to the entire domain. Some services are configured by default to log on to the LocalSystem account. Do not change the default service setting. The name of the account is LocalSystem. This account does not have a password.

Attribute	Value
Well-Known SID/RID	S-1-5-18
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Network

This group implicitly includes all users who are logged on through a network connection. Any user who accesses the system through a network has the Network identity. This identity allows only remote users to access a resource. Whenever a user accesses a given resource over the network, the user is automatically added to the Network group. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-2
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Network Service

The Network Service account is similar to an Authenticated User account. The Network Service account has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard your system if individual services or processes are compromised. Services that run as the Network Service account access network resources by

using the credentials of the computer account. The name of the account is NT AUTHORITY\NetworkService. This account does not have a password.

Attribute	Value
Well-Known SID/RID	S-1-5-20
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	Adjust memory quotas for a process: SeIncreaseQuotaPrivilege Bypass traverse checking: SeChangeNotifyPrivilege Create global objects: SeCreateGlobalPrivilege Generate security audits: SeAuditPrivilege Impersonate a client after authentication: SeImpersonatePrivilege Restore files and directories: SeAssignPrimaryTokenPrivilege Replace a process level token: SeAssignPrimaryTokenPrivilege

NTLM Authentication

Attribute	Value
Well-Known SID/RID	S-1-5-64-10
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Other Organization

This group implicitly includes all users who are logged on to the system through a dial-up connection. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-1000
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Principal Self

This identify is a placeholder in an ACE on a user, group, or computer object in Active Directory. When you grant permissions to Principal Self, you grant them to the security principal that is represented by the object. During an access check, the operating system replaces the SID for Principal Self with the SID for the security principal that is represented by the object.

Attribute	Value
Well-Known SID/RID	S-1-5-10
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Remote Interactive Logon

This identity represents all users who are currently logged on to a computer by using a Remote Desktop connection. This group is a subset of the Interactive group. Access tokens that contain the Remote Interactive Logon SID also contain the Interactive SID.

Attribute	Value
Well-Known SID/RID	S-1-5-14
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Restricted

Users and computers with restricted capabilities have the Restricted identity. This identity group is used by a process that is running in a restricted security context, such as running an application with the RunAs service. When code runs at the Restricted security level, the Restricted SID is added to the user's access token.

Attribute	Value
Well-Known SID/RID	S-1-5-12
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

SChannel Authentication

Attribute	Value
Well-Known SID/RID	S-1-5-64-14
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Service

Any service that accesses the system has the Service identity. This identity group includes all security principals that are signed in as a service. This identity grants access to processes that are being run by Windows Server services. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-6
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>

Attribute	Value
Default User Rights	Create global objects: SeCreateGlobalPrivilege Impersonate a client after authentication: SeImpersonatePrivilege

Terminal Server User

Any user accessing the system through Terminal Services has the Terminal Server User identity. This identity allows users to access Terminal Server applications and to perform other necessary tasks with Terminal Server services. Membership is controlled by the operating system.

Attribute	Value
Well-Known SID/RID	S-1-5-13
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

This Organization

Attribute	Value
Well-Known SID/RID	S-1-5-15
Object Class	Foreign Security Principal
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	None

Window Manager\Window Manager Group

Attribute	Value
Well-Known SID/RID	
Object Class	

Attribute	Value
Default Location in Active Directory	cn=WellKnown Security Principals, cn=Configuration, dc=<forestRootDomain>
Default User Rights	Bypass traverse checking: SeChangeNotifyPrivilege Increase a process working set: SeIncreaseWorkingSetPrivilege

See also

Active Directory Security Groups

Security Principals Technical Overview

Microsoft Accounts

This topic for the IT professional explains how a Microsoft account works to enhance security and privacy for users, and how you can manage this consumer account type in your organization.

Microsoft sites, services, and properties such as Windows Live, MSN, Xbox LIVE, Zune, Windows Phone, and computers running Windows 8.1, Windows 8, and Windows RT use a Microsoft account as a mean of identifying users. Microsoft account is the name for what was previously called Windows Live ID. It has user-defined secrets associated with it, and it consists of a unique email address and a password.

There are some benefits and considerations when using Microsoft accounts in the enterprise. For more information, see The Microsoft account in the enterprise later in this topic.

When a user signs in with a Microsoft account, their device is connected to cloud services, and many of the settings, preferences, and apps associated with that user account can roam between devices.

Note
This content applies to the operating system versions that are designated in the **Applies To** list at the beginning of this topic.

How a Microsoft account works

The Microsoft account allows users to sign in to websites that support this service by using a single set of credentials. Users' credentials are validated by a Microsoft account authentication server that is associated with a website. The Windows Store is an example of this association. When new users sign in to websites that are enabled to use Microsoft accounts, they are redirected to the nearest authentication server, which asks for a user name and password.

Windows uses the Schannel Security Support Provider to open a Transport Level Security/Secure Sockets Layer (TLS/SSL) connection for this function. Users then have the option to use Credential Manager to store their credentials.

When users sign in to websites that are enabled to use a Microsoft account, a time-limited cookie is installed on their computers, which includes a triple DES encrypted ID tag. This encrypted ID tag has been agreed upon between the authentication server and the website. This ID tag is sent to the website, and the website plants another time-limited encrypted HTTP cookie on the user's computer. When these cookies are valid, users are not required to supply a user name and password. If a user actively signs out of their Microsoft account, these cookies are removed.

Important

Local Windows account functionality has not been removed, and it is still an option to use in managed environments.

How Microsoft accounts are created

To prevent fraud, the Microsoft system verifies the IP address when a user creates an account. If a user tries to create multiple Microsoft accounts with the same IP address, they are stopped.

Microsoft accounts are not designed to be created in batches, for example, for a group of domain users within your enterprise.

There are two methods for creating a Microsoft account:

- **Use an existing email address**.

  Users are able to use their valid email addresses to sign up for Microsoft accounts. The service turns the requesting user's email address into a Microsoft account. Users can also choose their personal password.

- **Sign up for a Microsoft email address**.

  Users can sign up for an email account with Microsoft's webmail services. This account can be used to sign in to websites that are enabled to use Microsoft accounts.

How the Microsoft account information is safeguarded

Credential information is encrypted twice. The first encryption is based on the account's password. Credentials are encrypted again when they are sent across the Internet. The data that is stored is not available to other Microsoft or non-Microsoft services.

- **Strong password is required**.

  Blank passwords are not allowed.

  For more information, see Microsoft Account Security Overview.

- **Secondary proof of identity is required**.

  Before user profile information and settings can be accessed on a second supported Windows computer for the first time, trust must established for that device by providing secondary proof of identity. This can be accomplished by providing Windows with a code that

is sent to a mobile phone number or by following the instructions that are sent to an alternate email address that a user specifies in the account settings.

- **All user profile data is encrypted on the client before it is transmitted to the cloud**.

  User data does not roam over a wireless wide area network (WWAN) by default, thereby protecting profile data. All data and settings that leave a device are transmitted through the TLS/SSL protocol.

Microsoft account security information is added.

Users can add security information to their Microsoft accounts through the **Accounts** interface on computers running the supported versions of Windows. This feature allows the user to update the security information that they provided when they created their accounts. This security information includes an alternate email address or phone number so if their password is compromised or forgotten, a verification code can be sent to verify their identity. Users can potentially use their Microsoft accounts to store corporate data on a personal SkyDrive or email app, so it is safe practice for the account owner to keep this security information up-to-date. To see how this works, see Microsoft account security info: FAQ.

The Microsoft account in the enterprise

Although the Microsoft account was designed to serve consumers, you might find situations where your domain users can benefit by using their personal Microsoft account in your enterprise. The following list describes some advantages.

- **Download Windows Store apps**:

  If your enterprise chooses to distribute software through the Windows Store, your users can use their Microsoft accounts to download and use them on up to five devices running any version of Windows 8.1, Windows 8, or Windows RT.

- **Single sign-on**:

  Your users can use Microsoft account credentials to sign in to devices running Windows 8.1, Windows 8 or Windows RT. When they do this, Windows works with your Windows Store app to provide authenticated experiences for them. Users can associate a Microsoft account with their sign-in credentials for Windows Store apps or websites, so that these credentials roam across any devices running these supported versions.

- **Personalized settings synchronization**:

  Users can associate their most commonly used operating-system settings with a Microsoft account. These settings are available whenever a user signs in with that account on any device that is running a supported version of Windows and is connected to the cloud. After a user signs in, the device automatically attempts to get the user's settings from the cloud and apply them to the device.

- **App synchronization**:

  Windows Store apps can store user-specific settings so that these settings are available to any device. As with operating system settings, these user-specific app settings are available whenever the user signs in with the same Microsoft account on any device that is running a

supported version of Windows and is connected to the cloud. After the user signs in, that device automatically downloads the settings from the cloud and applies them when the app is installed.

- **Integrated social media services**:

  Contact information and status for your users' friends and associates automatically stay up-to-date from sites such as Hotmail, Outlook, Facebook, Twitter, and LinkedIn. Users can also access and share photos, documents, and other files from sites such as SkyDrive, Facebook, and Flickr.

Managing the Microsoft account in the domain

Depending on your IT and business models, introducing Microsoft accounts into your enterprise might add complexity or it might provide solutions. You should address the following considerations before you allow the use of these account types in your enterprise:

- Restrict the use of the Microsoft account
- Configure connected accounts
- Provision Microsoft accounts in the enterprise
- Audit account activity
- Perform password resets
- Restrict app installation and usage

Restrict the use of the Microsoft account

If employees are allowed to join the domain with their personal devices, they might expect to connect to enterprise resources by using their Microsoft accounts. If you want to prevent any use of Microsoft accounts within your enterprise, you can configure the local security policy setting Accounts: Block Microsoft accounts. However, this setting can prevent the users from signing in to their Windows devices with their Microsoft accounts (if they had set them up to do so) when they are joined to the domain.

The default for this setting is **Disabled**, which enables users to use their Microsoft accounts on devices that are joined to your domain. Other options in the setting can:

1. Prevent users from creating new Microsoft accounts on a computer, switch a local account to a Microsoft account, or connect a domain account to a Microsoft account. This is the preferred option if you need to limit the use of Microsoft accounts in your enterprise.
2. Prevent users with an existing Microsoft account from signing in to Windows. Selecting this option might make it impossible for an existing administrator to sign in to a computer and manage the system.

Configure connected accounts

Users can connect a Microsoft account to their domain account and synchronize the settings and preferences between them. This enables users to see the same desktop background, app

settings, browser history and favorites, and other Microsoft account settings on their other devices.

Users can disconnect a Microsoft account from their domain account at any time as follows: In **PC settings**, tap or click **Users**, tap or click **Disconnect**, and then tap or click **Finish**.

Note
Connecting Microsoft accounts with domain accounts can limit access to some high-privileged tasks in Windows. For example, Task Scheduler will evaluate the connected Microsoft account for access and fail. In these situations, the account owner should disconnect the account.

Provision Microsoft accounts in the enterprise

Microsoft accounts are private user accounts. There are no methods provided by Microsoft to provision Microsoft accounts for an enterprise. Enterprises should use domain accounts.

Audit account activity

Because Microsoft accounts are Internet-based, Windows does not have a mechanism to audit their use until the account is associated with a domain account. But this association does not restrict the user from disconnecting the account or disjoining from the domain. It is not possible to audit the activity of accounts that are not associated with your domain.

Perform password resets

Only the owner of the Microsoft account can change the password. Passwords can be changed in the Microsoft account sign-in portal.

Restrict app installation and usage

Within your organization, you can set application control policies to regulate app installation and usage for Microsoft accounts. For more information, see AppLocker Overview and Packaged Apps and Packaged App Installer Rules in AppLocker.

See also

Managing Privacy: Using a Microsoft Account to Logon and Resulting Internet Communication

Active Directory Security Groups

This reference topic for the IT professional describes the default Active Directory security groups.

There are two forms of common security principals in Active Directory: user accounts and computer accounts. These accounts represent a physical entity (a person or a computer). User

accounts can also be used as dedicated service accounts for some applications. Security groups are used to collect user accounts, computer accounts, and other groups into manageable units.

In the Windows Server operating system, there are several built-in accounts and security groups that are preconfigured with the appropriate rights and permissions to perform specific tasks. For Active Directory, there are two types of administrative responsibilities:

- **Service administrators** Responsible for maintaining and delivering Active Directory Domain Services (AD DS), including managing domain controllers and configuring the AD DS.
- **Data administrators** Responsible for maintaining the data that is stored in AD DS and on domain member servers and workstations.

About Active Directory groups

Groups are used to collect user accounts, computer accounts, and other groups into manageable units. Working with groups instead of with individual users helps simplify network maintenance and administration.

There are two types of groups in Active Directory:

- **Distribution groups** Used to create email distribution lists.
- **Security groups** Used to assign permissions to shared resources.

Distribution groups

Distribution groups can be used only with email applications (such as Exchange Server) to send email to collections of users. Distribution groups are not security enabled, which means that they cannot be listed in discretionary access control lists (DACLs).

Security groups

Security groups can provide an efficient way to assign access to resources on your network. By using security groups, you can:

- Assign user rights to security groups in Active Directory.

  User rights are assigned to a security group to determine what members of that group can do within the scope of a domain or forest. User rights are automatically assigned to some security groups when Active Directory is installed to help administrators define a person's administrative role in the domain.

  For example, a user who is added to the Backup Operators group in Active Directory has the ability to back up and restore files and directories that are located on each domain controller in the domain. This is possible because, by default, the user rights **Backup files and directories** and **Restore files and directories** are automatically assigned to the Backup Operators group. Therefore, members of this group inherit the user rights that are assigned to that group.

You can use Group Policy to assign user rights to security groups to delegate specific tasks. For more information about using Group Policy, see User Rights Assignment.

- Assign permissions to security groups for resources.

  Permissions are different than user rights. Permissions are assigned to the security group for the shared resource. Permissions determine who can access the resource and the level of access, such as Full Control. Some permissions that are set on domain objects are automatically assigned to allow various levels of access to default security groups, such as the Account Operators group or the Domain Admins group.

  Security groups are listed in DACLs that define permissions on resources and objects. When assigning permissions for resources (file shares, printers, and so on), administrators should assign those permissions to a security group rather than to individual users. The permissions are assigned once to the group, instead of several times to each individual user. Each account that is added to a group receives the rights that are assigned to that group in Active Directory, and the user receives the permissions that are defined for that group.

Like distribution groups, security groups can be used as an email entity. Sending an email message to the group sends the message to all the members of the group.

Group scope

Groups are characterized by a scope that identifies the extent to which the group is applied in the domain tree or forest. The scope of the group defines where the group can be granted permissions. The following three group scopes are defined by Active Directory:

- Universal
- Global
- Domain Local

Note
In addition to these three scopes, the default groups in the **Builtin** container have a group scope of Builtin Local. This group scope and group type cannot be changed.

The following table lists the three group scopes and more information about each scope for a security group.

Group scopes

Scope	Possible Members	Scope Conversion	Can Grant Permissions	Possible Member of
Universal	Accounts from any domain in the same forest Global groups from any domain in the same	Can be converted to Domain Local scope Can be converted to Global scope if the group does	On any domain in the same forest or trusting forests	Other Universal groups in the same forest Domain Local groups in the same forest or

Scope	Possible Members	Scope Conversion	Can Grant Permissions	Possible Member of
	forest Other Universal groups from any domain in the same forest	not contain any other Universal groups		trusting forests Local groups on computers in the same forest or trusting forests
Global	Accounts from the same domain Other Global groups from the same domain	Can be converted to Universal scope if the group is not a member of any other global group	On any domain in the same forest, or trusting domains or forests	Universal groups from any domain in the same forest Other Global groups from the same domain Domain Local groups from any domain in the same forest, or from any trusting domain
Domain Local	Accounts from any domain or any trusted domain Global groups from any domain or any trusted domain Universal groups from any domain in the same forest Other Domain Local groups from the same domain Accounts, Global groups, and Universal groups from other	Can be converted to Universal scope if the group does not contain any other Domain Local groups	Within the same domain	Other Domain Local groups from the same domain Local groups on computers in the same domain, excluding built-in groups that have well-known SIDs

Scope	Possible Members	Scope Conversion	Can Grant Permissions	Possible Member of
	forests and from external domains			

Special identity groups

Special identities are generally referred to as groups. Special identity groups do not have specific memberships that can be modified, but they can represent different users at different times, depending on the circumstances. Some of these groups include Creator Owner, Batch, and Authenticated User.

For information about all the special identity groups, see Special Identities.

Default security groups

Default groups, such as the Domain Admins group, are security groups that are created automatically when you create an Active Directory domain. You can use these predefined groups to help control access to shared resources and to delegate specific domain-wide administrative roles.

Many default groups are automatically assigned a set of user rights that authorize members of the group to perform specific actions in a domain, such as logging on to a local system or backing up files and folders. For example, a member of the Backup Operators group has the right to perform backup operations for all domain controllers in the domain.

When you add a user to a group, the user receives all the user rights that are assigned to the group and all the permissions that are assigned to the group for any shared resources.

Default groups are located in the **Builtin** container and in the **Users** container in Active Directory Users and Computers. The **Builtin** container includes groups that are defined with the Domain Local scope. The **Users** includes contains groups that are defined with Global scope and groups that are defined with Domain Local scope. You can move groups that are located in these containers to other groups or organizational units (OU) within the domain, but you cannot move them to other domains.

Some of the administrative groups that are listed in this topic and all members of these groups are protected by a background process that periodically checks for and applies a specific security descriptor. This descriptor is a data structure that contains security information associated with a protected object. This process ensures that any successful unauthorized attempt to modify the security descriptor on one of the administrative accounts or groups will be overwritten with the protected settings.

The security descriptor is present on the **AdminSDHolder** object. This means that if you want to modify the permissions on one of the service administrator groups or on any of its member accounts, you must modify the security descriptor on the **AdminSDHolder** object so that it will be

applied consistently. Be careful when you make these modifications because you are also changing the default settings that will be applied to all of your protected administrative accounts.

Active Directory default security groups by operating system version

The following tables provide descriptions of the default groups that are located in the **Builtin** and **Users** containers in each operating system.

Default Security Group	Windows Server 2012 R2	Windows Server 2012	Windows Server 2008 R2	Windows Server 2008
Access Control Assistance Operators	Yes	Yes		
Account Operators	Yes	Yes	Yes	Yes
Administrators	Yes	Yes	Yes	Yes
Allowed RODC Password Replication Group	Yes	Yes	Yes	Yes
Backup Operators	Yes	Yes	Yes	Yes
Certificate Service DCOM Access	Yes	Yes	Yes	Yes
Cert Publishers	Yes	Yes	Yes	Yes
Cloneable Domain Controllers	Yes	Yes		
Cryptographic Operators	Yes	Yes	Yes	Yes
Denied RODC Password Replication Group	Yes	Yes	Yes	Yes
Distributed COM Users	Yes	Yes	Yes	Yes
DnsUpdateProxy	Yes	Yes	Yes	Yes
DnsAdmins	Yes	Yes	Yes	Yes
Domain Admins	Yes	Yes	Yes	Yes
Domain Computers	Yes	Yes	Yes	Yes
Domain Controllers	Yes	Yes	Yes	Yes
Domain Guests	Yes	Yes	Yes	Yes

Default Security Group	Windows Server 2012 R2	Windows Server 2012	Windows Server 2008 R2	Windows Server 2008
Domain Users	Yes	Yes	Yes	Yes
Enterprise Admins	Yes	Yes	Yes	Yes
Enterprise Read-Only Domain Controllers	Yes	Yes	Yes	Yes
Event Log Readers	Yes	Yes	Yes	Yes
Group Policy Creators Owners	Yes	Yes	Yes	Yes
Guests	Yes	Yes	Yes	Yes
Hyper-V Administrators	Yes	Yes		
IIS_IUSRS	Yes	Yes	Yes	Yes
Incoming Forest Trust Builders	Yes	Yes	Yes	Yes
Network Configuration Operators	Yes	Yes	Yes	Yes
Performance Log Users	Yes	Yes	Yes	Yes
Performance Monitor Users	Yes	Yes	Yes	Yes
Pre–Windows 2000 Compatible Access	Yes	Yes	Yes	Yes
Print Operators	Yes	Yes	Yes	Yes
Protected Users	Yes			
RAS and IAS Servers	Yes	Yes	Yes	Yes
RDS Endpoint Servers	Yes	Yes		
RDS Management Servers	Yes	Yes		
RDS Remote Access Servers	Yes	Yes		
Read-Only Domain Controllers	Yes	Yes	Yes	Yes
Remote Desktop Users	Yes	Yes	Yes	Yes
Remote Management Users	Yes	Yes		

Default Security Group	Windows Server 2012 R2	Windows Server 2012	Windows Server 2008 R2	Windows Server 2008
Replicator	Yes	Yes	Yes	Yes
Schema Admins	Yes	Yes	Yes	Yes
Server Operators	Yes	Yes	Yes	Yes
Terminal Server License Servers	Yes	Yes	Yes	Yes
Users	Yes	Yes	Yes	Yes
Windows Authorization Access Group	Yes	Yes	Yes	Yes
WinRMRemoteWMIUsers_	Yes	Yes		

Access Control Assistance Operators

Members of this group can remotely query authorization attributes and permissions for resources on the computer.

The Access Control Assistance Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-579
Type	BuiltIn Local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Attribute	Value

Account Operators

The Account Operators group grants limited account creation privileges to a user. Members of this group can create and modify most types of accounts, including those of users, local groups, and global groups, and members can log in locally to domain controllers.

Members of the Account Operators group cannot manage the Administrator user account, the user accounts of administrators, or the Administrators, Server Operators, Account Operators, Backup Operators, or Print Operators groups. Members of this group cannot modify user rights.

The Account Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note

By default, this built-in group has no members, and it can create and manage users and groups in the domain, including its own membership and that of the Server Operators group. This group is considered a service administrator group because it can modify Server Operators, which in turn can modify domain controller settings. As a best practice, leave the membership of this group empty, and do not use it for any delegated administration. This group cannot be renamed, deleted, or moved.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-548
Type	BuiltIn Local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Allow log on locally: SeInteractiveLogonRight

Administrators

Members of the Administrators group have complete and unrestricted access to the computer, or if the computer is promoted to a domain controller, members have unrestricted access to the domain.

The Administrators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Notes

The Administrators group has built-in capabilities that give its members full control over the system. This group cannot be renamed, deleted, or moved. This built-in group controls access to all the domain controllers in its domain, and it can change the membership of all administrative groups.

Membership can be modified by members of the following groups: the default service Administrators, Domain Admins in the domain, or Enterprise Admins. This group has the special privilege to take ownership of any object in the directory or any resource on a domain controller. This account is considered a service administrator group because its members have full access to the domain controllers in the domain.

This security group includes the following changes since Windows Server 2008:

- Default user rights changes: **Allow log on through Terminal Services** existed in Windows Server 2008, and it was replaced by Allow log on through Remote Desktop Services.
- Remove computer from docking station was removed in Windows Server 2012 R2.

Attribute	Value
Well-Known SID/RID	S-1-5-32-544
Type	BuiltIn Local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	Administrator, Domain Admins, Enterprise Admins
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Adjust memory quotas for a process: SeIncreaseQuotaPrivilege Access this computer from the network: SeNetworkLogonRight

Attribute	Value
	Allow log on locally: SeInteractiveLogonRight Allow log on through Remote Desktop Services: SeRemoteInteractiveLogonRight Back up files and directories: SeBackupPrivilege Bypass traverse checking: SeChangeNotifyPrivilege Change the system time: SeSystemTimePrivilege Change the time zone: SeTimeZonePrivilege Create a pagefile: SeCreatePagefilePrivilege Create global objects: SeCreateGlobalPrivilege Create symbolic links: SeCreateSymbolicLinkPrivilege Debug programs: SeDebugPrivilege Enable computer and user accounts to be trusted for delegation: SeEnableDelegationPrivilege Force shutdown from a remote system: SeRemoteShutdownPrivilege Impersonate a client after authentication: SeImpersonatePrivilege Increase scheduling priority: SeIncreaseBasePriorityPrivilege Load and unload device drivers: SeLoadDriverPrivilege Log on as a batch job: SeBatchLogonRight Manage auditing and security log: SeSecurityPrivilege Modify firmware environment values: SeSystemEnvironmentPrivilege Perform volume maintenance tasks: SeManageVolumePrivilege Profile system performance: SeSystemProfilePrivilege Profile single process: SeProfileSingleProcessPrivilege

Attribute	Value
	Remove computer from docking station: SeUndockPrivilege Restore files and directories: SeRestorePrivilege Shut down the system: SeShutdownPrivilege Take ownership of files or other objects: SeTakeOwnershipPrivilege

Allowed RODC Password Replication Group

The purpose of this security group is to manage a RODC password replication policy. This group has no members by default, and it results in the condition that new Read-only domain controllers do not cache user credentials. The Denied RODC Password Replication Group group contains a variety of high-privilege accounts and security groups. The Denied RODC Password Replication group supersedes the Allowed RODC Password Replication group.

The Allowed RODC Password Replication group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-571
Type	Domain local
Default container	CN=Users DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Backup Operators

Members of the Backup Operators group can back up and restore all files on a computer, regardless of the permissions that protect those files. Backup Operators also can log on to and shut down the computer. This group cannot be renamed, deleted, or moved. By default, this built-in group has no members, and it can perform backup and restore operations on domain controllers. Its membership can be modified by the following groups: default service Administrators, Domain Admins in the domain, or Enterprise Admins. It cannot modify the membership of any administrative groups. While members of this group cannot change server settings or modify the configuration of the directory, they do have the permissions needed to replace files (including operating system files) on domain controllers. Because of this, members of this group are considered service administrators.

The Backup Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-551
Type	Builtin local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Allow log on locally: SeInteractiveLogonRight Back up files and directories: SeBackupPrivilege Log on as a batch job: SeBatchLogonRight Restore files and directories: SeRestorePrivilege Shut down the system: SeShutdownPrivilege

Certificate Service DCOM Access

Members of this group are allowed to connect to certification authorities in the enterprise.

The Certificate Service DCOM Access group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-<domain>-574
Type	Domain Local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Cert Publishers

Members of the Cert Publishers group are authorized to publish certificates for User objects in Active Directory.

The Cert Publishers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-517
Type	Domain Local
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to	No

Attribute	Value
non-Service admins?	
Default User Rights	None

Cloneable Domain Controllers

Members of the Cloneable Domain Controllers group that are domain controllers may be cloned. In Windows Server 2012 R2 and Windows Server 2012, you can deploy domain controllers by copying an existing virtual domain controller. In a virtual environment, you no longer have to repeatedly deploy a server image that is prepared by using sysprep.exe, promote the server to a domain controller, and then complete additional configuration requirements for deploying each domain controller (including adding the virtual domain controller to this security group).

For more information, see Introduction to Active Directory Domain Services (AD DS) Virtualization (Level 100).

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-522
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Cryptographic Operators

Members of this group are authorized to perform cryptographic operations. This security group was added in Windows Vista Service Pack 1 (SP1) to configure Windows Firewall for IPsec in Common Criteria mode.

The Cryptographic Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group was introduced in Windows Vista Service Pack 1, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-32-569
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Denied RODC Password Replication Group

Members of the Denied RODC Password Replication group cannot have their passwords replicated to any Read-only domain controller.

The purpose of this security group is to manage a RODC password replication policy. This group contains a variety of high-privilege accounts and security groups. The Denied RODC Password Replication Group supersedes the Allowed RODC Password Replication Group.

This security group includes the following changes since Windows Server 2008:

- Windows Server 2012 changed the default members to include Cert Publishers.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-572
Type	Domain local
Default container	CN=Users, DC=<domain>, DC=
Default members	Cert Publishers Domain Admins Domain Controllers Enterprise Admins Group Policy Creator Owners

Attribute	Value
	krbtgt Read-Only Domain Controllers Schema Admins
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Distributed COM Users

Members of the Distributed COM Users group are allowed to launch, activate, and use Distributed COM objects on the computer. Microsoft Component Object Model (COM) is a platform-independent, distributed, object-oriented system for creating binary software components that can interact. Distributed Component Object Model (DCOM) allows applications to be distributed across locations that make the most sense to you and to the application. This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Distributed COM Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-562
Type	Builtin Local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	

Attribute	Value
Default User Rights	None

DnsUpdateProxy

Members of the DnsUpdateProxy group are DNS clients. They are permitted to perform dynamic updates on behalf of other clients (such as DHCP servers). A DNS server can develop stale resource records when a DHCP server is configured to dynamically register host (A) and pointer (PTR) resource records on behalf of DHCP clients by using dynamic update. Adding clients to this security group mitigates this scenario.

However, to protect against unsecured records or to permit members of the DnsUpdateProxy group to register records in zones that allow only secured dynamic updates, you must create a dedicated user account and configure DHCP servers to perform DNS dynamic updates by using the credentials of this account (user name, password, and domain). Multiple DHCP servers can use the credentials of one dedicated user account.

For information, see DNS Record Ownership and the DnsUpdateProxy Group.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-1103
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

DnsAdmins

Members of DNSAdmins group have access to network DNS information. The default permissions are as follows: Allow: Read, Write, Create All Child objects, Delete Child objects, Special Permissions.

For information about other means to secure the DNS server service, see Securing the DNS Server Service.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-1102
Type	Domain local
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Domain Admins

Members of the Domain Admins security group are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain, including the domain controllers. The Domain Admins group is the default owner of any object that is created in Active Directory for the domain by any member of the group. If members of the group create other objects, such as files, the default owner is the Administrators group.

The Domain Admins group controls access to all domain controllers in a domain, and it can modify the membership of all administrative accounts in the domain. Membership can be modified by members of the service administrator groups in its domain (Administrators and Domain Admins), and by members of the Enterprise Admins group. This is considered a service administrator account because its members have full access to the domain controllers in a domain.

The Domain Admins group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-512

Attribute	Value
Type	Domain Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Administrator
Default member of	Administrators Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Administrators See Denied RODC Password Replication Group

Domain Computers

This group can include all computers and servers that have joined the domain, excluding domain controllers. By default, any computer account that is created automatically becomes a member of this group.

The Domain Computers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-515
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	All computers joined to the domain, excluding domain controllers
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes (but not required)
Safe to delegate management of this group to	Yes

Attribute	Value
non-Service admins?	
Default User Rights	None

Domain Controllers

The Domain Controllers group can include all domain controllers in the domain. New domain controllers are automatically added to this group.

The Domain Controllers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-516
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Computer accounts for all domain controllers of the domain
Default member of	Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	No
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	None

Domain Guests

The Domain Guests group includes the domain's built-in Guest account. When members of this group sign in as local guests on a domain-joined computer, a domain profile is created on the local computer.

The Domain Guests group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-514
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Guest
Default member of	Guests
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Can be moved out but it is not recommended
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Guests

Domain Users

The Domain Users group includes all user accounts in a domain. When you create a user account in a domain, it is automatically added to this group.

By default, any user account that is created in the domain automatically becomes a member of this group. This group can be used to represent all users in the domain. For example, if you want all domain users to have access to a printer, you can assign permissions for the printer to this group (or add the Domain Users group to a local group on the print server that has permissions for the printer).

The Domain Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-513
Type	Domain Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Administrator krbtgt
Default member of	Users
Protected by ADMINSDHOLDER?	No

Attribute	Value
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Users

Enterprise Admins

The Enterprise Admins group exists only in the root domain of an Active Directory forest of domains. It is a Universal group if the domain is in native mode; it is a Global group if the domain is in mixed mode. Members of this group are authorized to make forest-wide changes in Active Directory, such as adding child domains.

By default, the only member of the group is the Administrator account for the forest root domain. This group is automatically added to the Administrators group in every domain in the forest, and it provides complete access for configuring all domain controllers. Members in this group can modify the membership of all administrative groups. Membership can be modified only by the default service administrator groups in the root domain. This is considered a service administrator account.

The Enterprise Admins group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<root domain>-519
Type	Universal (if Domain is in Native-Mode) else Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Administrator
Default member of	Administrators Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Administrators

Attribute	Value
	See Denied RODC Password Replication Group

Enterprise Read-Only Domain Controllers

Members of this group are Read-Only Domain Controllers in the enterprise. Except for account passwords, a Read-only domain controller holds all the Active Directory objects and attributes that a writable domain controller holds. However, changes cannot be made to the database that is stored on the Read-only domain controller. Changes must be made on a writable domain controller and then replicated to the Read-only domain controller.

Read-only domain controllers address some of the issues that are commonly found in branch offices. These locations might not have a domain controller. Or, they might have a writable domain controller, but not the physical security, network bandwidth, or local expertise to support it.

For more information, see AD DS: Read-Only Domain Controllers.

The Enterprise Read-Only Domain Controllers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-498
Type	Universal
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Event Log Readers

Members of this group can read event logs from local computers. The group is created when the server is promoted to a domain controller.

The Event Log Readers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-573
Type	Builtin local
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Group Policy Creators Owners

This group is authorized to create, edit, or delete Group Policy Objects in the domain. By default, the only member of the group is Administrator.

For information about other features you can use with this security group, see Group Policy Planning and Deployment Guide.

The Group Policy Creators Owners group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<domain>-520
Type	Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Administrator

Attribute	Value
Default member of	Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	No
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Denied RODC Password Replication Group

Guests

Members of the Guests group have the same access as members of the Users group by default, except that the Guest account has further restrictions. By default, the only member is the Guest account. The Guests group allows occasional or one-time users to sign in with limited privileges to a computer's built-in Guest account.

When a member of the Guests group signs out, the entire profile is deleted. This includes everything that is stored in the **%userprofile%** directory, including the user's registry hive information, custom desktop icons, and other user-specific settings. This implies that a guest must use a temporary profile to sign in to the system. This security group interacts with the Group Policy setting **Do not logon users with temporary profiles** when it is enabled. This setting is located under the following path:

Computer Configuration\Administrative Templates\System\User Profiles

Notes

A Guest account is a default member of the Guests security group. People who do not have an actual account in the domain can use the Guest account. A user whose account is disabled (but not deleted) can also use the Guest account.

The Guest account does not require a password. You can set rights and permissions for the Guest account as in any user account. By default, the Guest account is a member of the built-in Guests group and the Domain Guests global group, which allows a user to sign in to a domain. The Guest account is disabled by default, and we recommend that it stay disabled.

The Guests group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-546

Attribute	Value
Type	Builtin Local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	Guest
Default member of	Domain Guests Guest
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	None

Hyper-V Administrators

Members of the Hyper-V Administrators group have complete and unrestricted access to all the features in Hyper-V. Adding members to this group helps reduce the number of members required in the Administrators group, and further separates access.

Note
Prior to Windows Server 2012, access to features in Hyper-V was controlled in part by membership in the Administrators group.

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-32-578
Type	Builtin local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	None
Default member of	No
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	

Attribute	Value
Default User Rights	None

IIS_IUSRS

IIS_IUSRS is a built-in group that is used by Internet Information Services beginning with IIS 7.0. A built-in account and group are guaranteed by the operating system to always have a unique SID. IIS 7.0 replaces the IUSR_MachineName account and the IIS_WPG group with the IIS_IUSRS group to ensure that the actual names that are used by the new account and group will never be localized. For example, regardless of the language of the Windows operating system that you install, the IIS account name will always be IUSR, and the group name will be IIS_IUSRS.

For more information, see Understanding Built-In User and Group Accounts in IIS 7.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-568
Type	BuiltIn Local
Default container	CN=BuiltIn, DC=<domain>, DC=
Default members	IUSR
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Incoming Forest Trust Builders

Members of the Incoming Forest Trust Builders group can create incoming, one-way trusts to this forest. Active Directory provides security across multiple domains or forests through domain and forest trust relationships. Before authentication can occur across trusts, Windows must determine whether the domain being requested by a user, computer, or service has a trust relationship with the logon domain of the requesting account.

To make this determination, the Windows security system computes a trust path between the domain controller for the server that receives the request and a domain controller in the domain of

the requesting account. A secured channel extends to other Active Directory domains through interdomain trust relationships. This secured channel is used to obtain and verify security information, including security identifiers (SIDs) for users and groups.

Note
This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

For more information, see How Domain and Forest Trusts Work: Domain and Forest Trusts.

The Incoming Forest Trust Builders group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note
This group cannot be renamed, deleted, or moved.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-557
Type	BuiltIn local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	None

Network Configuration Operators

Members of the Network Configuration Operators group can have the following administrative privileges to manage configuration of networking features:

- Modify the Transmission Control Protocol/Internet Protocol (TCP/IP) properties for a local area network (LAN) connection, which includes the IP address, the subnet mask, the default gateway, and the name servers.
- Rename the LAN connections or remote access connections that are available to all the users.
- Enable or disable a LAN connection.

- Modify the properties of all of remote access connections of users.
- Delete all the remote access connections of users.
- Rename all the remote access connections of users.
- Issue **ipconfig**, **ipconfig /release**, or **ipconfig /renew** commands.
- Enter the PIN unblock key (PUK) for mobile broadband devices that support a SIM card.

Note
This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Network Configuration Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note
This group cannot be renamed, deleted, or moved.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-556
Type	BuiltIn local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	None

Performance Log Users

Members of the Performance Log Users group can manage performance counters, logs, and alerts locally on the server and from remote clients without being a member of the Administrators group. Specifically, members of this security group:

- Can use all the features that are available to the Performance Monitor Users group.
- Can create and modify Data Collector Sets after the group is assigned the Log on as a batch job user right.

Warning

If you are a member of the Performance Log Users group, you must configure Data Collector Sets that you create to run under your credentials.

- Cannot use the Windows Kernel Trace event provider in Data Collector Sets.

For members of the Performance Log Users group to initiate data logging or modify Data Collector Sets, the group must first be assigned the Log on as a batch job user right. To assign this user right, use the Local Security Policy snap-in in Microsoft Management Console.

Note

This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Performance Log Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note

This account cannot be renamed, deleted, or moved.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-559
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	Log on as a batch job: SeBatchLogonRight

Performance Monitor Users

Members of this group can monitor performance counters on domain controllers in the domain, locally and from remote clients, without being a member of the Administrators or Performance Log Users groups. The Windows Performance Monitor is a Microsoft Management Console (MMC) snap-in that provides tools for analyzing system performance. From a single console, you

can monitor application and hardware performance, customize what data you want to collect in logs, define thresholds for alerts and automatic actions, generate reports, and view past performance data in a variety of ways.

Specifically, members of this security group:

- Can use all the features that are available to the Users group.
- Can view real-time performance data in Performance Monitor.

  Can change the Performance Monitor display properties while viewing data.
- Cannot create or modify Data Collector Sets.

> **Warning**
> You cannot configure a Data Collector Set to run as a member of the Performance Monitor Users group.

> **Note**
> This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO). This group cannot be renamed, deleted, or moved.

The Performance Monitor Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-558
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	None

Pre–Windows 2000 Compatible Access

Members of the Pre–Windows 2000 Compatible Access group have Read access for all users and groups in the domain. This group is provided for backward compatibility for computers

running Windows NT 4.0 and earlier. By default, the special identity group, Everyone, is a member of this group. Add users to this group only if they are running Windows NT 4.0 or earlier.

Warning
This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Pre–Windows 2000 Compatible Access group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-554
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	If you choose the Pre–Windows 2000 Compatible Permissions mode, Everyone and Anonymous are members, and if you choose the Windows 2000-only permissions mode, Authenticated Users are members.
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Access this computer from the network: SeNetworkLogonRight Bypass traverse checking: SeChangeNotifyPrivilege

Print Operators

Members of this group can manage, create, share, and delete printers that are connected to domain controllers in the domain. They can also manage Active Directory printer objects in the domain. Members of this group can locally sign in to and shut down domain controllers in the domain.

This group has no default members. Because members of this group can load and unload device drivers on all domain controllers in the domain, add users with caution. This group cannot be renamed, deleted, or moved.

The Print Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008. However, in Windows Server 2008 R2, functionality was added to manage print administration. For more information, see Assigning Delegated Print Administrator and Printer Permission Settings in Windows Server 2008 R2.

Attribute	Value
Well-Known SID/RID	S-1-5-32-550
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Allow log on locally: SeInteractiveLogonRight Load and unload device drivers: SeLoadDriverPrivilege Shut down the system: SeShutdownPrivilege

Protected Users

Members of the Protected Users group are afforded additional protection against the compromise of credentials during authentication processes.

This security group is designed as part of a strategy to effectively protect and manage credentials within the enterprise. Members of this group automatically have non-configurable protection applied to their accounts. Membership in the Protected Users group is meant to be restrictive and proactively secure by default. The only method to modify the protection for an account is to remove the account from the security group.

This domain-related, global group triggers non-configurable protection on devices and host computers running Windows Server 2012 R2 and Windows 8.1, and on domain controllers in domains with a primary domain controller running Windows Server 2012 R2. This greatly reduces

the memory footprint of credentials when users sign in to computers on the network from a non-compromised computer.

Depending on the account's domain functional level, members of the Protected Users group are further protected due to behavior changes in the authentication methods that are supported in Windows.

- Members of the Protected Users group cannot authenticate by using the following Security Support Providers (SSPs): NTLM, Digest Authentication, or CredSSP. Passwords are not cached on a device running Windows 8.1, so the device fails to authenticate to a domain when the account is a member of the Protected User group.
- The Kerberos protocol will not use the weaker DES or RC4 encryption types in the preauthentication process. This means that the domain must be configured to support at least the AES cipher suite.
- The user's account cannot be delegated with Kerberos constrained or unconstrained delegation. This means that former connections to other systems may fail if the user is a member of the Protected Users group.
- The default Kerberos ticket-granting tickets (TGTs) lifetime setting of four hours is configurable by using Authentication Policies and Silos, which can be accessed through the Active Directory Administrative Center. This means that when four hours has passed, the user must authenticate again.

The Protected Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This group was introduced in Windows Server 2012 R2. For more information about how this group works, see Protected Users Security Group.

The following table specifies the properties of the Protected Users group.

Attribute	Value
Well-known SID/RID	S-1-5-21-<domain>-525
Type	Domain Global
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-service admins?	No
Default user rights	None

RAS and IAS Servers

Computers that are members of the RAS and IAS Servers group, when properly configured, are allowed to use remote access services. By default, this group has no members. Computers that are running the Routing and Remote Access service are added to the group automatically, such as IAS servers and Network Policy Servers. Members of this group have access to certain properties of User objects, such as Read Account Restrictions, Read Logon Information, and Read Remote Access Information.

The RAS and IAS Servers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-553
Type	Domain local
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	None

RDS Endpoint Servers

Servers that are members in the RDS Endpoint Servers group can run virtual machines and host sessions where user RemoteApp programs and personal virtual desktops run. This group needs to be populated on servers running RD Connection Broker. Session Host servers and RD Virtualization Host servers used in the deployment need to be in this group.

For information about Remote Desktop Services, see Remote Desktop Services Design Guide.

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-553

Attribute	Value
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

RDS Management Servers

Servers that are members in the RDS Management Servers group can be used to perform routine administrative actions on servers running Remote Desktop Services. This group needs to be populated on all servers in a Remote Desktop Services deployment. The servers running the RDS Central Management service must be included in this group.

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-32-577
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

RDS Remote Access Servers

Servers in the RDS Remote Access Servers group provide users with access to RemoteApp programs and personal virtual desktops. In Internet facing deployments, these servers are typically deployed in an edge network. This group needs to be populated on servers running RD Connection Broker. RD Gateway servers and RD Web Access servers that are used in the deployment need to be in this group.

For information about RemoteApp programs, see Overview of RemoteApp

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-32-575
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

Remote Desktop Users

The Remote Desktop Users group on an RD Session Host server is used to grant users and groups permissions to remotely connect to an RD Session Host server. This group cannot be renamed, deleted, or moved. It appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Remote Desktop Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-555

Attribute	Value
Type	Builtin Local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	None

Read-Only Domain Controllers

This group is comprised of the Read-only domain controllers in the domain. A Read-only domain controller makes it possible for organizations to easily deploy a domain controller in scenarios where physical security cannot be guaranteed, such as branch office locations, or in scenarios where local storage of all domain passwords is considered a primary threat, such as in an extranet or in an application-facing role.

Because administration of a Read-only domain controller can be delegated to a domain user or security group, an Read-only domain controller is well suited for a site that should not have a user who is a member of the Domain Admins group. A Read-only domain controller encompasses the following functionality:

- Read-only AD DS database
- Unidirectional replication
- Credential caching
- Administrator role separation
- Read-only Domain Name System (DNS)

For information about deploying a Read-only domain controller, see Read-Only Domain Controllers Step-by-Step Guide.

This security group was introduced in Windows Server 2008, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-521
Type	

Attribute	Value
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	
Default User Rights	See Denied RODC Password Replication Group

Remote Management Users

Members of the Remote Management Users group can access WMI resources over management protocols (such as WS-Management via the Windows Remote Management service). This applies only to WMI namespaces that grant access to the user.

The Remote Management Users group is generally used to allow users to manage servers through the Server Manager console, whereas the WinRMRemoteWMIUsers_ group is allows remotely running Windows PowerShell commands.

For more information, see WS-Management Protocol (Windows) and About WMI (Windows).

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-32-580
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	

Attribute	Value
Default User Rights	None

Replicator

Computers that are members of the Replicator group support file replication in a domain. Windows Server operating systems use the File Replication service (FRS) to replicate system policies and logon scripts stored in the System Volume (SYSVOL). Each domain controller keeps a copy of SYSVOL for network clients to access. FRS can also replicate data for the Distributed File System (DFS), synchronizing the content of each member in a replica set as defined by DFS. FRS can copy and maintain shared files and folders on multiple servers simultaneously. When changes occur, content is synchronized immediately within sites and by a schedule between sites.

Important

In Windows Server 2008 R2, FRS cannot be used for replicating DFS folders or custom (non-SYSVOL) data. A Windows Server 2008 R2 domain controller can still use FRS to replicate the contents of a SYSVOL shared resource in a domain that uses FRS for replicating the SYSVOL shared resource between domain controllers.

However, Windows Server 2008 R2 servers cannot use FRS to replicate the contents of any replica set apart from the SYSVOL shared resource. The DFS Replication service is a replacement for FRS, and it can be used to replicate the contents of a SYSVOL shared resource, DFS folders, and other custom (non-SYSVOL) data. You should migrate all non-SYSVOL FRS replica sets to DFS Replication. For more information, see File Replication Service (FRS) Is Deprecated in Windows Server 2008 R2 (Windows).

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-552
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	

Attribute	Value
Default User Rights	None

Schema Admins

Members of the Schema Admins group can modify the Active Directory schema. This group exists only in the root domain of an Active Directory forest of domains. It is a Universal group if the domain is in native mode; it is a Global group if the domain is in mixed mode.

The group is authorized to make schema changes in Active Directory. By default, the only member of the group is the Administrator account for the forest root domain. This group has full administrative access to the schema.

The membership of this group can be modified by any of the service administrator groups in the root domain. This is considered a service administrator account because its members can modify the schema, which governs the structure and content of the entire directory.

For more information, see What Is the Active Directory Schema?: Active Directory.

The Schema Admins group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-<root domain>-518
Type	Universal (if Domain is in Native-Mode) else Global
Default container	CN=Users, DC=<domain>, DC=
Default members	Administrator
Default member of	Denied RODC Password Replication Group
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	See Denied RODC Password Replication Group

Server Operators

Members in the Server Operators group can administer domain servers. This group exists only on domain controllers. By default, the group has no members. Memebers of the Server Operators group can sign in to a server interactively, create and delete network shared resources, start and stop services, back up and restore files, format the hard disk drive of the computer, and shut down the computer. This group cannot be renamed, deleted, or moved.

By default, this built-in group has no members, and it has access to server configuration options on domain controllers. Its membership is controlled by the service administrator groups, Administrators and Domain Admins, in the domain, and the Enterprise Admins group. Members in this group cannot change any administrative group memberships. This is considered a service administrator account because its members have physical access to domain controllers, they can perform maintenance tasks (such as backup and restore), and they have the ability to change binaries that are installed on the domain controllers. Note the default user rights in the following table.

The Server Operators group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-549
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	Yes
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	Allow log on locally: SeInteractiveLogonRight Back up files and directories: SeBackupPrivilege Change the system time: SeSystemTimePrivilege Change the time zone: SeTimeZonePrivilege Force shutdown from a remote system: SeRemoteShutdownPrivilege Restore files and directories: Restore files and

Attribute	Value
	directories SeRestorePrivilege Shut down the system: SeShutdownPrivilege

Terminal Server License Servers

Members of the Terminal Server License Servers group can update user accounts in Active Directory with information about license issuance. This is used to track and report TS Per User CAL usage. A TS Per User CAL gives one user the right to access a Terminal Server from an unlimited number of client computers or devices. This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

For more information about this security group, see Terminal Services License Server Security Group Configuration.

The Terminal Server License Servers group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note

This group cannot be renamed, deleted, or moved.

This security group only applies to Windows Server 2003 and Windows Server 2008 because Terminal Services was replaced by Remote Desktop Services in Windows Server 2008 R2.

Attribute	Value
Well-Known SID/RID	S-1-5-32-561
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	None
Default member of	None
Safe to move out of default container?	Cannot be moved
Protected by ADMINSDHOLDER?	No
Safe to delegate management of this group to non-Service admins?	Yes
Default User Rights	None

Users

Members of the Users group are prevented from making accidental or intentional system-wide changes, and they can run most applications. After the initial installation of the operating system, the only member is the Authenticated Users group. When a computer joins a domain, the Domain Users group is added to the Users group on the computer.

Users can perform tasks such as running applications, using local and network printers, shutting down the computer, and locking the computer. Users can install applications that only they are allowed to use if the installation program of the application supports per-user installation. This group cannot be renamed, deleted, or moved.

The Users group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

This security group includes the following changes since Windows Server 2008:

- In Windows Server 2008 R2, INTERACTIVE was added to the default members list.
- In Windows Server 2012, the default **Member Of** list changed from Domain Users to none.

Attribute	Value
Well-Known SID/RID	S-1-5-32-545
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	Authenticated Users Domain Users INTERACTIVE
Default member of	Domain Users (this membership is due to the fact that the Primary Group ID of all user accounts is Domain Users.)
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	No
Default User Rights	None

Windows Authorization Access Group

Members of this group have access to the computed token GroupsGlobalAndUniversal attribute on User objects. Some applications have features that read the token-groups-global-and-universal (TGGAU) attribute on user account objects or on computer account objects in Active

Directory Domain Services. Some Win32 functions make it easier to read the TGGAU attribute. Applications that read this attribute or that call an API (referred to as a function) that reads this attribute do not succeed if the calling security context does not have access to the attribute. This group appears as a SID until the domain controller is made the primary domain controller and it holds the operations master role (also known as flexible single master operations or FSMO).

The Windows Authorization Access group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

Note

This group cannot be renamed, deleted, or moved.

This security group has not changed since Windows Server 2008.

Attribute	Value
Well-Known SID/RID	S-1-5-32-560
Type	Builtin local
Default container	CN=Builtin, DC=<domain>, DC=
Default members	Enterprise Domain Controllers
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Cannot be moved
Safe to delegate management of this group to non-Service admins?	Yes
Default user rights	None

WinRMRemoteWMIUsers_

In Windows 8 and in Windows Server 2012, a **Share** tab was added to the Advanced Security Settings user interface. This tab displays the security properties of a remote file share. To view this information, you must have the following permissions and memberships, as appropriate for the version of Windows Server that the file server is running.

The WinRMRemoteWMIUsers_ group applies to versions of the Windows Server operating system listed in the Active Directory default security groups by operating system version.

- If the file share is hosted on a server that is running a supported version of the operating system:
  - You must be a member of the WinRMRemoteWMIUsers__ group or the BUILTIN\Administrators group.
  - You must have Read permissions to the file share.

- If the file share is hosted on a server that is running a version of Windows Server that is earlier than Windows Server 2012:
  - You must be a member of the BUILTIN\Administrators group.
  - You must have Read permissions to the file share.

In Windows Server 2012, the Access Denied Assistance functionality adds the Authenticated Users group to the local WinRMRemoteWMIUsers__ group. Therefore, when the Access Denied Assistance functionality is enabled, all authenticated users who have Read permissions to the file share can view the file share permissions.

Note
The WinRMRemoteWMIUsers_ group allows running Windows PowerShell commands remotely whereas the Remote Management Users group is generally used to allow users to manage servers by using the Server Manager console.

This security group was introduced in Windows Server 2012, and it has not changed in subsequent versions.

Attribute	Value
Well-Known SID/RID	S-1-5-21-<domain>-1000
Type	Domain local
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-Service admins?	
Default User Rights	None

See also

Security Principals Technical Overview

Special Identities

AppLocker Overview

This technical overview for the IT professional provides a description of AppLocker. This information can help you decide if your organization can benefit from deploying AppLocker application control policies. AppLocker helps administrators control which applications and files users can run. These include executable files, scripts, Windows Installer files, dynamic-link libraries (DLLs), packaged apps, and packaged app installers.

Tip

To digitally save or print pages from this library, click **Export** (in the upper-right corner of the page), and then follow the instructions.

What does AppLocker do?

By using AppLocker, you can:

- Define rules based on file attributes that persist across application updates, such as the publisher name (derived from the digital signature), product name, file name, and file version. You can also create rules based on the file path and hash.
- Assign a rule to a security group or an individual user.
- Create exceptions to rules. For example, you can create a rule that allows all users to run all Windows binaries except the Registry Editor (Regedit.exe).
- Use audit-only mode to deploy the policy and understand its impact before enforcing it.
- Create rules on a staging server, test them, then export them to your production environment and import them into a Group Policy Object.
- Simplify creating and managing AppLocker rules by using Windows PowerShell cmdlets for AppLocker.

AppLocker helps reduce administrative overhead and helps reduce the organization's cost of managing computing resources by decreasing the number of Help Desk calls that result from users running unapproved applications. AppLocker addresses the following application security scenarios:

- Application inventory

  AppLocker has the ability to enforce its policy in an audit-only mode where all application access activity is registered in event logs. These events can be collected for further analysis. Windows PowerShell cmdlets also help you analyze this data programmatically.

- Protection against unwanted software

  AppLocker has the ability to deny applications from running when you exclude them from the list of allowed applications. When AppLocker rules are enforced in the production environment, any application that is not included in the allowed rules is blocked from running.

- Licensing conformance

  AppLocker can help you create rules that preclude unlicensed software from running and restrict licensed software to authorized users.

- Software standardization

  AppLocker policies can be configured to allow only supported or approved applications to run on computers within a business group. This permits a more uniform application deployment.
- Manageability improvement

  AppLocker includes a number of improvements in manageability as compared to its predecessor Software Restriction Policies. Importing and exporting policies, automatic generation of rules from multiple files, audit-only mode deployment, and Windows PowerShell cmdlets are a few of the improvements over Software Restriction Policies.

When to use AppLocker

In many organizations, information is the most valuable asset, and ensuring that only approved users have access to that information is imperative. Access control technologies such as Active Directory Rights Management Services (AD RMS) and access control lists (ACLs) help control what users are allowed to access.

However, when a user runs a process, that process has the same level of access to data that the user has. As a result, sensitive information could easily be deleted or transmitted out of the organization if a user knowingly or unknowingly runs malicious software. AppLocker can help mitigate these types of security breaches by restricting the files that users or groups are allowed to run.

Software publishers are beginning to create more applications that can be installed by non-administrative users. This could jeopardize an organization's written security policy and circumvent traditional application control solutions that rely on the inability of users to install applications. By allowing administrators to create an allowed list of approved files and applications, AppLocker helps prevent such per-user applications from running. Because AppLocker can control DLLs, it is also useful to control who can install and run ActiveX controls.

AppLocker is ideal for organizations that currently use Group Policy to manage their Windows-based computers. Because AppLocker relies on Group Policy for authoring and deployment, experience with Group Policy is helpful if you plan to use AppLocker.

The following are examples of scenarios in which AppLocker can be used:

- Your organization's security policy dictates the use of only licensed software, so you need to prevent users from running unlicensed software and also restrict the use of licensed software to authorized users.
- An application is no longer supported by your organization, so you need to prevent it from being used by everyone.
- The potential that unwanted software can be introduced in your environment is high, so you need to reduce this threat.
- The license to an application has been revoked or it is expired in your organization, so you need to prevent it from being used by everyone.

- A new application or a new version of an application is deployed, and you need to prevent users from running the old version.
- Specific software tools are not allowed within the organization, or only specific users should have access to those tools.
- A single user or small group of users needs to use a specific application that is denied for all others.
- Some computers in your organization are shared by people who have different software usage needs, and you need to protect specific applications.
- In addition to other measures, you need to control the access to sensitive data through application usage.

AppLocker can help you protect the digital assets within your organization, reduce the threat of malicious software being introduced into your environment, and improve the management of application control and the maintenance of application control policies.

Versions, interoperability, and differences in functionality

Supported versions and interoperability considerations

AppLocker policies can only be configured on and applied to computers that are running the support versions and editions Windows operating system. For more information see Requirements to Use AppLocker.

Differences in functionality between versions

The following table lists the differences by operating system version for each of the major features in or functions of AppLocker:

Feature or function	Windows Server 2008 R2 and Windows 7	Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8
Ability to set rules for packaged apps and packaged app installers.	No	Yes
AppLocker policies are maintained through Group Policy, and only the administrator of the computer can update an AppLocker policy.	Yes	Yes
AppLocker permits customizing error messages so that administrators can direct users to a web page for help.	Yes	Yes

Feature or function	Windows Server 2008 R2 and Windows 7	Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8
Ability to work in conjunction with Software Restriction Policies (by using separate GPOs).	Yes	Yes
AppLocker supports a small set of Windows PowerShell cmdlets to aid administration and maintenance.	Yes	Yes
AppLocker rules can control the listed file formats.	• .exe • .com • .ps1 • .bat • .cmd • .vbs • .js • .msi • .msp • .dll • .ocx	• .exe • .com • .ps1 • .bat • .cmd • .vbs • .js • .msi • .msp • .mst • .dll • .ocx • .appx

For information about comparing the application control functions in Software Restriction Policies and AppLocker, and about using the two features together, see Use AppLocker and Software Restriction Policies in the Same Domain.

System requirements

AppLocker policies can only be configured on and applied to computers that are running on the supported versions and editions of the Windows operating system. Group Policy is required to distribute Group Policy Objects that contain AppLocker policies. For more information, see Requirements to Use AppLocker.

AppLocker rules can be created on domain controllers.

Note
The ability to author or enforce rules for packaged apps and packaged app installers is not available in Windows Server 2008 R2 and Windows 7.

Installing AppLocker

AppLocker is included with enterprise-level editions of Windows. You can author AppLocker rules for a single computer or for a group of computers. For a single computer, you can author the rules by using the Local Security Policy editor (secpol.msc). For a group of computers, you can author the rules within a Group Policy Object by using the Group Policy Management Console (GPMC).

Note

The GPMC is available in client computers running Windows only by installing the Remote Server Administration Tools. On computer running Windows Server, you must install the Group Policy Management feature.

Using AppLocker on Server Core

Windows PowerShell can be used to manage AppLocker on Server Core installations by using the AppLocker cmdlets, and if administered within a GPO, the Group Policy PowerShell cmdlets. For more information, see the AppLocker PowerShell Command Reference.

Virtualization considerations

You can administer AppLocker policies by using a virtualized instance of Windows provided it meets all the system requirements listed previously. You can also run Group Policy in a virtualized instance. However, you do risk losing the policies that you created and maintain if the virtualized instance is removed or fails.

Security considerations

Application control policies specify which programs are allowed to run on the local computer.

The variety of forms that malicious software can take make it difficult for users to know what is safe to run. When activated, malicious software can damage content on a hard disk drive, flood a network with requests to cause a denial-of-service (DoS) attack, send confidential information to the Internet, or compromise the security of a computer.

The countermeasure is to create a sound design for your application control policies on end-user computers in your organization, and then thoroughly test the policies in a lab environment before you deploy them in a production environment. AppLocker can be part of your application control strategy because you can control what software is allowed to run on your computers.

A flawed application control policy implementation can disable necessary applications or allow malicious or unintended software to run. Therefore, it is important that organizations dedicate sufficient resources to manage and troubleshoot the implementation of such policies.

For additional information about specific security issues, see Security Considerations for AppLocker.

When you use AppLocker to create application control policies, you should be aware of the following security considerations:

- Who has the rights to set AppLocker policies?

- How do you validate that the policies are enforced?
- What events should you audit?

For reference in your security planning, the following table identifies the baseline settings for a client computer with the AppLocker feature installed:

Setting	Default value
Accounts created	None
Authentication method	Not applicable
Management interfaces	AppLocker can be managed by using a Microsoft Management Console snap-in, Group Policy Management, and Windows PowerShell
Ports opened	None
Minimum privileges required	Administrator on the local computer; Domain Admin, or any set of rights that allow you to create, edit and distribute Group Policy Objects.
Protocols used	Not applicable
Scheduled Tasks	Appidpolicyconverter.exe is put in a scheduled task to be run on demand.
Security Policies	None required. AppLocker creates security policies.
System Services required	Application Identity service (appidsvc) runs under LocalServiceAndNoImpersonation.
Storage of credentials	None

Maintaining AppLocker policies

Information about Applocker policy maintenance is provided in the following topics:

- Administer AppLocker
- Monitor Application Usage with AppLocker
- Optimize AppLocker Performance
- Use AppLocker and Software Restriction Policies in the Same Domain
- Manage Packaged Apps with AppLocker

See also

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8
Product evaluation	Frequently Asked Questions AppLocker Step-by-Step Guide	Frequently Asked Questions AppLocker Step-by-Step Guide
Procedures	AppLocker Operations Guide	Administer AppLocker Manage Packaged Apps with AppLocker
Scripting	Using the AppLocker Windows PowerShell Cmdlets	Use the AppLocker Windows PowerShell Cmdlets
Technical content	AppLocker Technical Reference	AppLocker Technical Reference
Design, planning, and deployment	AppLocker Policies Design Guide AppLocker Policies Deployment Guide	AppLocker Policies Design Guide AppLocker Policies Deployment Guide

Administer AppLocker

This topic provides links to specific procedures to use when administering AppLocker policies and rules in those operating system versions designated in the **Applies To** list at the beginning of this topic.

AppLocker helps administrators control how users can access and use files, such as executable files, packaged apps, scripts, Windows Installer files, and DLLs. Using AppLocker, you can:

- Define rules based on file attributes derived from the digital signature, including the publisher, product name, file name, and file version. For example, you can create rules based on the publisher attribute that is persistent through updates, or you can create rules for a specific version of a file.
- Assign a rule to a security group or an individual user.
- Create exceptions to rules. For example, you can create a rule that allows all Windows processes to run except Registry Editor (Regedit.exe).
- Use audit-only mode to deploy the policy and understand its impact before enforcing it.
- Import and export rules. The import and export affects the entire policy. For example, if you export a policy, all of the rules from all of the rule collections are exported, including the enforcement settings for the rule collections. If you import a policy, the existing policy is overwritten.
- Simplify creating and managing AppLocker rules by using AppLocker PowerShell cmdlets.

Note

For more information about enhanced capabilities of AppLocker to control Windows apps, see Packaged Apps and Packaged App Installer Rules in AppLocker.

The following topics are included to administer AppLocker:

Maintain AppLocker Policies

Edit an AppLocker Policy

Test and Update an AppLocker Policy

Deploy AppLocker Policies by Using the Enforce Rules Setting

Use the AppLocker Windows PowerShell Cmdlets

Optimize AppLocker Performance

Monitor Application Usage with AppLocker

Use AppLocker and Software Restriction Policies in the Same Domain

Manage Packaged Apps with AppLocker

How to work with policies

- Configure an AppLocker Policy for Audit Only
- Configure an AppLocker Policy for Enforce Rules
- Display a Custom URL Message When Users Try to Run a Blocked Application
- Export an AppLocker Policy from a GPO
- Export an AppLocker Policy to an XML File
- Import an AppLocker Policy from Another Computer
- Import an AppLocker Policy into a GPO
- Merge AppLocker Policies by Using Set-ApplockerPolicy
- Merge AppLocker Policies Manually
- Refresh an AppLocker Policy
- Test an AppLocker Policy by Using Test-AppLockerPolicy

How to work with rules

- Create a Rule That Uses a File Hash Condition
- Create a Rule That Uses a Path Condition
- Create a Rule That Uses a Publisher Condition
- Create AppLocker Default Rules
- Configure Exceptions for an AppLocker Rule
- Create a Rule for Packaged Apps
- Delete an AppLocker Rule
- Edit AppLocker Rules
- Enable the DLL Rule Collection
- Enforce AppLocker Rules
- Run the Automatically Generate Rules Wizard

Using the MMC snap-ins to administer AppLocker

You can administer AppLocker policies by using the Group Policy Management Console to create or edit a Group Policy Object (GPO), or to create or edit an AppLocker policy on a local computer by using the Local Group Policy Editor snap-in or the Local Security Policy snap-in.

Administer Applocker using Group Policy

You must have Edit Setting permission to edit a GPO. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission. Also, the Group Policy Management feature must be installed on the computer.

1. On the **Start** screen, type **gpmc.msc** or open the Group Policy Management Console (GPMC).
2. Locate the GPO that contains the AppLocker policy to modify, right-click the GPO, and click **Edit**.
3. In the console tree, double-click **Application Control Policies**, double-click **AppLocker**, and then click the rule collection that you want to create the rule for.

Administer AppLocker on the local computer

1. On the **Start** screen, type **secpol.msc** or **gpedit.msc**.
2. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. In the console tree of the snap-in, double-click **Application Control Policies**, double-click **AppLocker**, and then click the rule collection that you want to create the rule for.

Using Windows PowerShell to administer AppLocker

For how-to information about administering AppLocker with Windows PowerShell, see Use the AppLocker Windows PowerShell Cmdlets. For reference information and examples how to administer AppLocker with Windows PowerShell, see the AppLocker PowerShell Command Reference.

Maintain AppLocker Policies

This topic describes how to maintain rules within AppLocker policies in Windows Server 2012 and Windows 8.

Common AppLocker maintenance scenarios include:

- A new application is deployed, and you need to update an AppLocker policy.

- A new version of an application is deployed, and you need to either update an AppLocker policy or create a new rule to update the policy.
- An application is no longer supported by your organization, so you need to prevent it from being used.
- An application appears to be blocked but should be allowed.
- An application appears to be allowed but should be blocked.
- A single user or small subset of users needs to use a specific application that is blocked.

There are two methods you can use to maintain AppLocker policies:

- Maintaining AppLocker policies by using Group Policy
- Maintaining AppLocker policies by using the Local Security Policy snap-in

As new applications are deployed or existing applications are removed by your organization or updated by the software publisher, you might need to make revisions to your rules and update the Group Policy Object (GPO) to ensure that your policy is current.

You can edit an AppLocker policy by adding, changing, or removing rules. However, you cannot specify a version for the AppLocker policy by importing additional rules. To ensure version control when modifying an AppLocker policy, use Group Policy management software that allows you to create versions of GPOs, such as Microsoft Advanced Group Policy Management (AGPM). For more information about AGPM, see Advanced Group Policy Management Overview (http://go.microsoft.com/fwlink/?LinkId=145013).

Caution
You should not edit an AppLocker rule collection while it is being enforced in Group Policy. Because AppLocker controls what files are allowed to run, making changes to a live policy can create unexpected behavior.

Maintaining AppLocker policies by using Group Policy

For every scenario, the steps to maintain an AppLocker policy distributed by Group Policy include the following tasks.

Step 1: Understand the current behavior of the policy

Before modifying a policy, evaluate how the policy is currently implemented. For example, if a new version of the application is deployed, you can use **Test-AppLockerPolicy** to verify the effectiveness of your current policy for that application. To read the procedures necessary to understand the current behavior of the policy, see Discovering the Effect of an AppLocker Policy. Updating your AppLocker planning document will help you track your findings. For information about creating this document, see Creating Your AppLocker Planning Document. For information about **Test-AppLockerPolicy** and examples of how to use it, see Test-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169000).

Note
These documents apply to Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7 operating systems.

Step 2: Export the AppLocker policy from the GPO

Updating an AppLocker policy that is currently enforced in your production environment can have unintended results. Therefore, export the policy from the GPO and update the rule or rules by using AppLocker on your AppLocker reference or test computer. To prepare an AppLocker policy for modification, see Export an AppLocker Policy from a GPO

Step 3: Update the AppLocker policy by editing the appropriate AppLocker rule

After the AppLocker policy has been exported from the GPO into the AppLocker reference or test computer, or has been accessed on the local computer, the specific rules can be modified as required.

To modify AppLocker rules, see the following:

- Edit AppLocker Rules
- Merge AppLocker Policies by Using Set-ApplockerPolicy or Merge AppLocker Policies Manually
- Delete an AppLocker Rule
- Enforce AppLocker Rules

Step 4: Test the AppLocker policy

You should test each collection of rules to ensure that the rules perform as intended. (Because AppLocker rules are inherited from linked GPOs, you should deploy all rules for simultaneous testing in all test GPOs.) For steps to perform this testing, see Test and Update an AppLocker Policy.

Step 5: Import the AppLocker policy into the GPO

After testing, import the AppLocker policy back into the GPO for implementation. To update the GPO with a modified AppLocker policy, see Import an AppLocker Policy into a GPO.

Step 6: Monitor the resulting policy behavior

After deploying a policy, evaluate the policy's effectiveness. For steps to understand the new behavior of the policy, see Discovering the Effect of an AppLocker Policy.

Maintaining AppLocker policies by using the Local Security Policy snap-in

For every scenario, the steps to maintain an AppLocker policy by using the Local Group Policy Editor or the Local Security Policy snap-in include the following tasks.

Step 1: Understand the current behavior of the policy

Before modifying a policy, evaluate how the policy is currently implemented. To read the procedures necessary to understand the current behavior of the policy, see Monitor Application Usage with AppLocker. Updating your AppLocker planning document will help you track your findings. For information about creating this document, see Creating Your AppLocker Planning Document.

Note
The AppLocker Planning Guide applies to Windows Server 2012, Windows Server 2008 R2, Windows 8 and Windows 7.

Step 2: Update the AppLocker policy by modifying the appropriate AppLocker rule

Rules are grouped into a collection, which can have the policy enforcement setting applied to it. By default, AppLocker rules do not allow users to open or run any files that are not specifically allowed.

To modify AppLocker rules, see the appropriate topic listed on Administer AppLocker.

Step 3: Test the AppLocker policy

You should test each collection of rules to ensure that the rules perform as intended. For steps to perform this testing, see Test and Update an AppLocker Policy.

Step 4: Deploy the policy with the modified rule

You can export and then import AppLocker policies to deploy the policy to other computers running Windows 8 or Windows Server 2012. To perform this task, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy from Another Computer.

Step 5: Monitor the resulting policy behavior

After deploying a policy, evaluate the policy's effectiveness. For steps to understand the new behavior of the policy, see Discovering the Effect of an AppLocker Policy in Discovering the Effect of an AppLocker Policy.

Note

This topic applies to Windows Server 2012, Windows Server 2008 R2, Windows 8 and Windows 7.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Edit an AppLocker Policy

This topic describes the steps you need to perform to modify an AppLocker policy in Windows Server 2012 and Windows 8.

You can edit an AppLocker policy by adding, changing, or removing rules. However, you cannot create a new version of the policy by importing additional rules. To modify an AppLocker policy that is in production, you should use Group Policy management software that allows you to version Group Policy Objects (GPOs). If you have created multiple AppLocker policies and need to merge them to create one AppLocker policy, you can either manually merge the policies or use the Windows PowerShell cmdlets for AppLocker. You cannot automatically merge policies by using the AppLocker snap-in. You must create one rule collection from two or more policies. The AppLocker policy is saved in XML format, and the exported policy can be edited with any text or XML editor. For information about merging policies, see Merge AppLocker Policies Manually or Merge AppLocker Policies by Using Set-ApplockerPolicy.

There are two methods you can use to edit an AppLocker policy:

- Editing an AppLocker policy by using Group Policy
- Editing an AppLocker policy by using the Local Security Policy snap-in

Editing an AppLocker policy by using Group Policy

The steps to edit an AppLocker policy distributed by Group Policy include the following:

Step 1: Use Group Policy management software to export the AppLocker policy from the GPO

AppLocker provides a feature to export and import AppLocker policies as an XML file. This allows you to modify an AppLocker policy outside your production environment. Because updating an AppLocker policy in a deployed GPO could have unintended consequences, you should first

export the AppLocker policy to an XML file. For the procedure to do this, see Export an AppLocker Policy from a GPO.

Step 2: Import the AppLocker policy into the AppLocker reference computer or the computer you use for policy maintenance

After exporting the AppLocker policy to an XML file, you should import the XML file onto a reference computer so that you can edit the policy. For the procedure to import an AppLocker policy, see Import an AppLocker Policy from Another Computer.

Caution
Importing a policy onto another computer will overwrite the existing policy on that computer.

Step 3: Use AppLocker to modify and test the rule

AppLocker provides ways to modify, delete, or add rules to a policy by modifying the rules within the collection.

- For the procedure to modify a rule, see Edit AppLocker Rules.
- For the procedure to delete a rule, see Delete an AppLocker Rule.
- For procedures to create rules, see:
  - Create a Rule That Uses a Publisher Condition
  - Create a Rule That Uses a Path Condition
  - Create a Rule That Uses a File Hash Condition
  - Enable the DLL Rule Collection
- For steps to test an AppLocker policy, see Test and Update an AppLocker Policy.
- For procedures to export the updated policy from the reference computer back into the GPO, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy into a GPO.

Step 4: Use AppLocker and Group Policy to import the AppLocker policy back into the GPO

For procedures to export the updated policy from the reference computer back into the GPO, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy into a GPO.

Caution
You should never edit an AppLocker rule collection while it is being enforced in Group Policy. Because AppLocker controls what files are allowed run, making changes to a live policy can create unexpected behavior. For information about testing policies, see Test and Update an AppLocker Policy.

Note

If you are performing these steps by using Microsoft Advanced Group Policy Management (AGPM), check out the GPO before exporting the policy.

Editing an AppLocker policy by using the Local Security Policy snap-in

The steps to edit an AppLocker policy distributed by using the Local Security Policy snap-in include the following tasks.

Step 1: Import the AppLocker policy

On the computer where you maintain policies, open the AppLocker snap-in from the Local Security Policy snap-in. If you exported the AppLocker policy from another computer, use AppLocker to import it onto the computer.

After exporting the AppLocker policy to an XML file, you should import the XML file onto a reference computer so that you can edit the policy. For the procedure to import an AppLocker policy, see Import an AppLocker Policy from Another Computer.

Caution
Importing a policy onto another computer will overwrite the existing policy on that computer.

Step 2: Identify and modify the rule to change, delete, or add

AppLocker provides ways to modify, delete, or add rules to a policy by modifying the rules within the collection.

- For the procedure to modify a rule, see Edit AppLocker Rules.
- For the procedure to delete a rule, see Delete an AppLocker Rule.
- For procedures to create rules, see:
  - Create a Rule That Uses a Publisher Condition
  - Create a Rule That Uses a Path Condition
  - Create a Rule That Uses a File Hash Condition
  - Enable the DLL Rule Collection

Step 3: Test the effect of the policy

For steps to test an AppLocker policy, see Test and Update an AppLocker Policy.

Step 4: Export the policy to an XML file and propagate it to all targeted computers

For procedures to export the updated policy from the reference computer to targeted computers, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy from Another Computer.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Test and Update an AppLocker Policy

This topic discusses the steps required to test an AppLocker policy prior to deployment in Windows Server 2012 and Windows 8.

You should test each set of rules to ensure that the rules perform as intended. If you use Group Policy to manage AppLocker policies, complete the following steps for each Group Policy Object (GPO) where you have created AppLocker rules. Because AppLocker rules are inherited from linked GPOs, you should deploy all of the rules for simultaneous testing in all of your test GPOs.

Step 1: Enable the Audit only enforcement setting

By using the **Audit only** enforcement setting, you can ensure that the AppLocker rules that you have created are properly configured for your organization. This setting can be enabled on the **Enforcement** tab of the **AppLocker Properties** dialog box. For the procedure to do this, see Configure an AppLocker Policy for Audit Only.

Step 2: Configure the Application Identity service to start automatically

Because AppLocker uses the Application Identity service to verify the attributes of a file, you must configure it to start automatically in any one GPO that applies AppLocker rules. For the procedure to do this, see Configure the Application Identity Service. For AppLocker policies that are not managed by a GPO, you must ensure that the service is running on each computer in order for the policies to be applied.

Step 3: Test the policy

Test the AppLocker policy to determine if your rule collection needs to be modified. Because you have created AppLocker rules, enabled the Application Identity service, and enabled the **Audit only** enforcement setting, the AppLocker policy should be present on all client computers that are configured to receive your AppLocker policy.

The **Test-AppLockerPolicy** Windows PowerShell cmdlet can be used to determine whether any of the rules in your rule collection will be blocked on your reference computers. For the procedure to do this, see Test an AppLocker Policy by Using Test-AppLockerPolicy.

Step 4: Analyze AppLocker events

You can either manually analyze AppLocker events or use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to automate the analysis.

To manually analyze AppLocker events

You can view the events either in Event Viewer or a text editor and then sort those events to perform an analysis, such as looking for patterns in application usage events, access frequencies, or access by user groups. If you have not configured an event subscription, then you will have to review the logs on a sampling of computers in your organization. For more information about using Event Viewer, see View the AppLocker Log in Event Viewer.

To analyze AppLocker events by using Get-AppLockerFileInformation

You can use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to analyze AppLocker events from a remote computer. If an application is being blocked and should be allowed, you can use the AppLocker cmdlets to help troubleshoot the problem.

For both event subscriptions and local events, you can use the **Get-AppLockerFileInformation** cmdlet to determine which files have been blocked or would have been blocked (if you are using the **Audit only** enforcement mode) and how many times the event has occurred for each file. For the procedure to do this, see Review AppLocker Events with Get-AppLockerFileInformation.

After using **Get-AppLockerFileInformation** to determine how many times that a file would have been blocked from running, you should review your rule list to determine whether a new rule should be created for the blocked file or whether an existing rule is too strictly defined. Ensure that you check which GPO is currently preventing the file from running. To determine this, you can use the Group Policy Results Wizard to view rule names.

Step 5: Modify the AppLocker policy

After you have identified which rules need to be edited or added to the policy, you can use the Group Policy Management Console to modify the AppLocker rules in the relevant GPOs. For AppLocker policies that are not managed by a GPO, you can use the Local Security Policy snap-in. For information how to modify an AppLocker policy, see, Edit an AppLocker Policy.

Step 6: Repeat policy testing, analysis, and policy modification

Repeat the previous steps 3–5 until all the rules perform as intended before applying enforcement.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Deploy AppLocker Policies by Using the Enforce Rules Setting

This topic describes the steps to deploy AppLocker policies by using the enforcement setting method in Windows Server 2012 and Windows 8.

Background and prerequisites

These procedures assume that you have already deployed AppLocker policies with the enforcement set to **Audit only**, and you have been collecting data through the AppLocker event logs and other channels to determine what effect these policies have on your environment and the policy's adherence to your application control design.

For information about the AppLocker policy enforcement setting, see Understanding AppLocker Enforcement Settings.

For information about how to plan an AppLocker policy deployment, see AppLocker Policies Design Guide.

Step 1: Retrieve the AppLocker policy

Updating an AppLocker policy that is currently enforced in your production environment can have unintended results. Using Group Policy, you can export the policy from the Group Policy Object (GPO) and then update the rule or rules by using AppLocker on your AppLocker reference or test computer. For the procedure to do this, see Export an AppLocker Policy from a GPO and Import an AppLocker Policy into a GPO. For local AppLocker policies, you can update the rule or rules by using the Local Security policy snap-in on your AppLocker reference or test computer. For the procedures to do this, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy from Another Computer.

Step 2: Alter the enforcement setting

Rule enforcement is applied only to a collection of rules, not to individual rules. AppLocker divides the rules into collections: executable files, Windows Installer files, packaged apps, scripts, and DLL files. By default, if enforcement is not configured and rules are present in a rule collection, those rules are enforced. For information about the enforcement setting, see Understanding AppLocker Enforcement Settings. For the procedure to alter the enforcement setting, see Configure an AppLocker Policy for Audit Only.

Step 3: Update the policy

You can edit an AppLocker policy by adding, changing, or removing rules. However, you cannot specify a version for the AppLocker policy by importing additional rules. To ensure version control when modifying an AppLocker policy, use Group Policy management software that allows you to create versions of GPOs. An example of this type of software is the Advanced Group Policy Management feature from the Microsoft Desktop Optimization Pack. For more information about Advanced Group Policy Management, see Advanced Group Policy Management Overview (http://go.microsoft.com/fwlink/?LinkId=145013).

Caution
You should not edit an AppLocker rule collection while it is being enforced in Group Policy. Because AppLocker controls what files are allowed to run, making changes to a live policy can create unexpected behavior.

For the procedure to update the GPO, see Import an AppLocker Policy into a GPO.

For the procedures to distribute policies for local computers by using the Local Security Policy snap-in, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy from Another Computer.

Step 4: Monitor the effect of the policy

When a policy is deployed, it is important to monitor the actual implementation of that policy. You can do this by monitoring your support organization's application access request activity and reviewing the AppLocker event logs. To monitor the effect of the policy, see View the AppLocker Log in Event Viewer and Review AppLocker Events with Get-AppLockerFileInformation.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Use the AppLocker Windows PowerShell Cmdlets

This topic describes how each AppLocker Windows PowerShell cmdlet can help you administer your AppLocker application control policies in Windows Server 2012 and Windows 8.

AppLocker Windows PowerShell cmdlets

The five AppLocker cmdlets are designed to streamline the administration of an AppLocker policy. They can be used to help create, test, maintain, and troubleshoot an AppLocker policy. The cmdlets are intended to be used in conjunction with the AppLocker user interface that is accessed through the Microsoft Management Console (MMC) snap-in extension to the Local Security Policy snap-in and Group Policy Management Console.

To edit or update a Group Policy Object (GPO) by using the AppLocker cmdlets, you must have Edit Setting permission. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission. To perform tasks by using the Local Security policy snap-in, you must be a member of the local **Administrators** group, or equivalent, on the computer.

Import the AppLocker PowerShell cmdlet module

To use the AppLocker cmdlets, you must first import the AppLocker module by using the following command at the Windows PowerShell command prompt: `C:\PS> Import-Module AppLocker`. Scripting must be enabled on the computer.

Retrieve application information

The Get-AppLockerFileInformation (http://go.microsoft.com/fwlink/?LinkId=169154) cmdlet retrieves the AppLocker file information from a list of files or from an event log. File information that is retrieved can include publisher information, file hash information, and file path information. File information from an event log may not contain all of these fields. Files that are not signed do not have any publisher information.

Set AppLocker policy

The Set-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169167) cmdlet sets the specified GPO to contain the specified AppLocker policy. If no Lightweight Directory Access Protocol (LDAP) is specified, the local GPO is the default.

Retrieve an AppLocker policy

The Get-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169165) cmdlet gets the AppLocker policy from the local GPO, from a specified GPO, or from the effective AppLocker

policy on the computer. The output of the AppLocker policy is an AppLockerPolicy object or an XML-formatted string.

Generate rules for a given user or group

The New-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169166) cmdlet uses a list of file information to automatically generate rules for a given user or group. It can generate rules based on publisher, hash, or path information. Use **Get-AppLockerFileInformation** to create the list of file information.

Test the AppLocker Policy against a file set

The Test-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkID=169000) cmdlet uses the specified AppLocker policy to test whether a specified list of files are allowed to run or not on the local computer for a specific user.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Use AppLocker and Software Restriction Policies in the Same Domain

This topic for the IT professional describes concepts and procedures to help you manage your application control strategy using Software Restriction Policies and AppLocker.

Using AppLocker and Software Restriction Policies in the same domain

AppLocker is supported on systems running Windows 7 and above. Software Restriction Policies (SRP) is supported on systems running Windows Vista or earlier. You can continue to use SRP for application control on your pre-Windows 7 computers, but use AppLocker for computers running Windows Server 2008 R2, Windows 7 and later. It is recommended that you author AppLocker and SRP rules in separate GPOs and target the GPO with SRP policies to systems running Windows Vista or earlier. When both SRP and AppLocker policies are applied to computers running Windows Server 2008 R2, Windows 7 and later, the SRP policies are ignored.

The following table compares the features and functions of Software Restriction Policies (SRP) and AppLocker.

Application control function	SRP	AppLocker
Scope	SRP policies can be applied to all Windows operating systems beginning with Windows XP and Windows Server 2003.	AppLocker policies apply only to Windows Server 2008 R2, Windows Server 2012, Windows 7, and Windows 8.
Policy creation	SRP policies are maintained through Group Policy and only the administrator of the GPO can update the SRP policy. The administrator on the local computer can modify the SRP policies defined in the local GPO.	AppLocker policies are maintained through Group Policy and only the administrator of the GPO can update the policy. The administrator on the local computer can modify the AppLocker policies defined in the local GPO. AppLocker permits customization of error messages to direct users to a Web page for help.
Policy maintenance	SRP policies must be updated by using the Local Security Policy snap-in (if the policies are created locally) or the Group Policy Management Console (GPMC).	AppLocker policies can be updated by using the Local Security Policy snap-in (if the policies are created locally), or the GPMC, or the Windows PowerShell AppLocker cmdlets.
Policy application	SRP policies are distributed through Group Policy.	AppLocker policies are distributed through Group Policy.
Enforcement mode	SRP works in the “deny list mode” where administrators can create rules for files that they do not want to allow in this Enterprise whereas the rest of the file are allowed to run by default. SRP can also be configured in the “allow list mode” such that the by default all files are blocked and administrators	AppLocker by default works in the “allow list mode” where only those files are allowed to run for which there is a matching allow rule.

Application control function	SRP	AppLocker
	need to create allow rules for files that they want to allow.	
File types that can be controlled	SRP can control the following file types: • Executables • Dlls • Scripts • Windows Installers SRP cannot control each file type separately. All SRP rules are in a single rule collection.	AppLocker can control the following file types on Windows Server 2008 R2 and Windows 7 and later versions: • Executables • Dlls • Scripts • Windows Installers Packaged apps and installersWere added beginning withWindows Server 2012 and Windows 8 AppLocker maintains a separate rule collection for each of the five file types.
Designated file types	SRP supports an extensible list of file types that are considered executable. Administrators can add extensions for files that should be considered executable.	AppLocker does not support this. AppLocker currently supports the following file extensions: • Executables (.exe, .com) • Dlls (.ocx, .dll) • Scripts (.vbs, .js, .ps1, .cmd, .bat) • Windows Installers (.msi, .mst, .msp) • Packaged app installers (.appx)
Rule types	SRP supports four types of rules: • Hash • Path • Signature • Internet zone	AppLocker supports three types of rules: • Hash • Path • Publisher
Editing the hash value	In Windows XP, SRP allows administrators to provide	AppLocker computes the hash value itself. Internally it uses the SHA2 Authenticode hash

Application control function	SRP	AppLocker
	custom hash values. Beginning with Windows 7, and Windows Server 2008 R2, you can only select the file to hash, not provide the hash value.	for Portable Executables (Exe and Dll) and Windows Installers and a SHA2 flat file hash for the rest.
Support for different security levels	With SRP administrators can specify the permissions with which an app can run. So, an administrator can configure a rule such that notepad always runs with restricted permissions and never with administrative privileges. SRP on Windows Vista and earlier supported multiple security levels. On Windows 7 that list was restricted to just two levels: Disallowed and Unrestricted (Basic User translates to Disallowed).	AppLocker does not support security levels.
Manage Packaged apps and Packaged app installers.	Unable	.appx is a valid file type which AppLocker can manage.
Targeting a rule to a user or a group of users	SRP rules apply to all users on a particular computer.	AppLocker rules can be targeted to a specific user or a group of users.
Support for rule exceptions	SRP does not support rule exceptions	AppLocker rules can have exceptions which allows administrators to create rules such as “Allow everything from Windows except for Regedit.exe”.
Support for audit mode	SRP does not support audit mode. The only way to test SRP policies is to set up a test environment and run a few experiments.	AppLocker supports audit mode which allows administrators to test the effect of their policy in the real production environment without impacting the user experience. Once you are satisfied with the results, you can start enforcing

Application control function	SRP	AppLocker
		the policy.
Support for exporting and importing policies	SRP does not support policy import/export.	AppLocker supports the importing and exporting of policies. This allows you to create AppLocker policy on a sample computer, test it out and then export that policy and import it back into the desired GPO.
Rule enforcement	Internally, SRP rules enforcement happens in the user-mode which is less secure.	Internally, AppLocker rules for Exes and Dlls are enforced in the kernel-mode which is more secure than enforcing them in the user-mode.

See Also

AppLocker Overview

Optimize AppLocker Performance

This topic describes how to optimize AppLocker policy enforcement in Windows Server 2012 and Windows 8.

Optimization of Group Policy

AppLocker policies can be implemented by organization unit (OU) using Group Policy. If so, your Group Policy infrastructure should be optimized and retested for performance when AppLocker policies are added to existing Group Policy Objects (GPOs) or new GPOs are created, as you do with adding any policies to your GPOs.

For additional information, see the Optimizing Group Policy Performance article in TechNet Magazine (http://go.microsoft.com/fwlink/?LinkId=163238).

AppLocker rule limitations

The more rules per GPO, the longer AppLocker requires for evaluation. There is no set limitation on the number of rules per GPO, but the number of rules that can fit into a 100 MB GPO varies based on the complexity of the rule, such as the number of file hashes included in a single file hash condition.

Using the DLL rule collection

When the DLL rule collection is enabled, AppLocker must check each DLL that an application loads. The more DLLs, the longer AppLocker requires to complete the evaluation.

See Also

AppLocker Overview

Monitor Application Usage with AppLocker

This topic describes how to monitor application usage when AppLocker policies are applied in Windows Server® 2012 and Windows® 8.

Once you set rules and deploy the AppLocker policies, it is good practice to determine if the policy implementation is what you expected.

Discover the effect of an AppLocker policy

You can evaluate how the AppLocker policy is currently implemented for documentation or audit purposes, or before you modify the policy. Updating your AppLocker Policy Deployment Planning document will help you track your findings. For information about creating this document, see Creating Your AppLocker Planning Document. You can perform one or more of the following steps to understand what application controls are currently enforced through AppLocker rules.

- **Analyze the AppLocker logs in Event Viewer**

  When AppLocker policy enforcement is set to **Enforce rules**, rules are enforced for the rule collection and all events are audited. When AppLocker policy enforcement is set to **Audit only**, rules are not enforced but are still evaluated to generate audit event data that is written to the AppLocker logs.

  For the procedure to access the log, see View the AppLocker Log in Event Viewer.

- **Enable the Audit only AppLocker enforcement setting**

  By using the **Audit only** enforcement setting, you can ensure that the AppLocker rules are properly configured for your organization. When AppLocker policy enforcement is set to **Audit only**, rules are only evaluated but all events generated from that evaluation are written to the AppLocker log.

  For the procedure to do this, see Configure an AppLocker Policy for Audit Only.

- **Review AppLocker events with Get-AppLockerFileInformation**

  For both event subscriptions and local events, you can use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to determine which files have been blocked or would have been blocked (if you are using the audit-only enforcement mode) and how many times the event has occurred for each file.

For the procedure to do this, see Review AppLocker events with Get-AppLockerFileInformation.

- **Review AppLocker events with Test-AppLockerPolicy**

  You can use the **Test-AppLockerPolicy** Windows PowerShell cmdlet to determine determine whether any of the rules in your rule collections will be blocked on your reference computer or the computer on which you maintain policies.

  For the procedure to do this, see Test an AppLocker Policy by Using Test-AppLockerPolicy.

Review AppLocker events with Get-AppLockerFileInformation

For both event subscriptions and local events, you can use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to determine which files have been blocked or would have been blocked (if the **Audit only** enforcement setting is applied) and how many times the event has occurred for each file.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

Note
If the AppLocker logs are not on the local computer, you will need permission to view the logs. If the output is saved to a file, you will need permission to read that file.

To review AppLocker events with Get-AppLockerFileInformation

1. Open a Command Prompt window.
2. At the command prompt, type **PowerShell**, and then press ENTER.
3. Run the following command to review how many times a file would have been blocked from running if rules were enforced:

   ```
   Get-AppLockerFileInformation -EventLog -Logname "Microsoft-Windows-AppLocker\EXE and DLL" -EventType Audited -Statistics
   ```

   Note
   For an event subscription, specify the path to the forwarded event log for the Logname parameter.
4. Run the following command to review how many times a file has been allowed to run or prevented from running:

   ```
   Get-AppLockerFileInformation -EventLog -Logname "Microsoft-Windows-AppLocker\EXE and DLL" -EventType Allowed -Statistics
   ```

View the AppLocker Log in Event Viewer

When AppLocker policy enforcement is set to **Enforce rules**, rules are enforced for the rule collection and all events are audited. When AppLocker policy enforcement is set to **Audit only**,

rules are only evaluated but all events generated from that evaluation are written to the AppLocker log.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To view events in the AppLocker log by using Event Viewer

1. Open Event Viewer. To do this, click **Start**, type **eventvwr.msc** in the **Search programs and files** box, and then press ENTER.
2. In the console tree under **Application and Services Logs\Microsoft\Windows**, double-click **AppLocker**.

AppLocker events are listed in either the **EXE and DLL** log, the **MSI and Script** log, or the **Packaged app-Deployment** or **Packaged app-Execution** log. Event information includes the enforcement setting, file name, date and time, and user name. The logs can be exported to other file formats for further analysis.

See Also

AppLocker Overview

Manage Packaged Apps with AppLocker

This topic for the IT professional describes concepts and lists procedures to help you manage Packaged apps with AppLocker as part of your overall application control strategy in Windows Server 2012 and Windows 8.

Understanding Packaged apps and Packaged app installers for AppLocker

Packaged apps (also known as Windows apps) in Windows Server 2012 and Windows 8 are based on a model that ensures all the files within an app package share the same identity. With classic Win32 applications, each file within the application could have a unique identity. With Packaged apps, it is possible to control the entire application by using a single AppLocker rule.

Note
AppLocker supports only publisher rules for Packaged apps. All Packaged apps must be signed by the software publisher because Windows does not support unsigned Packaged apps.

Typically, an app consists of multiple components: the installer that is used to install the app, and one or more Exes, Dlls or Scripts. With classic Win32 applications, not all these components always share common attributes such as the software's publisher name, product name, and product version. Therefore, AppLocker controls each of these components separately through

different rule collections, such as Exe, Dll, Script, and Windows Installer rules. In contrast, all the components of a Packaged app share the same publisher name, package name, and package version attributes. Therefore, you can control an entire app with a single rule.

AppLocker enforces rules for Packaged apps separately from classic applications. A single AppLocker rule for a packaged app can control installing and running the app.

Comparing classic apps and Packaged apps

AppLocker policies for Packaged apps can only be applied to applications installed on computers running Windows Server 2012 or Windows 8, but classic apps can be controlled on Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7. The rules for classic apps and Packaged apps can be enforced in tandem. The differences between Packaged apps and classic apps that you should consider include:

- **Installing the apps** All Packaged apps can be installed by a standard user, whereas a number of classic apps require administrative privileges to install. In an environment where most of the users are standard users, you might not have numerous Exe rules (because classic apps require administrative privileges to install), but you might want to have more explicit policies for Packaged apps.
- **Changing the system state** Classic apps can be written to change the system state if they are run with administrative privileges. Most Packaged apps cannot change the system state because they run with limited privileges. When you design your AppLocker policies, it is important to understand whether an app that you are allowing can make system-wide changes.
- **Acquiring the apps** Windows apps can be acquired through the Microsoft AppStore or by loading using Windows PowerShell cmdlets (which requires a special Enterprise license). Classic apps can be acquired through traditional means.

AppLocker uses different rule collections to control Packaged apps and classic apps. You have the choice to control one type, the other type, or both.

For information about controlling classic apps, see Administer AppLocker.

For more information about Packaged apps, see Packaged Apps and Packaged App Installer Rules in AppLocker.

Design and deployment decisions

You can use two methods to create an inventory of Packaged apps on a computer: the AppLocker UI or the **Get-AppxPackage** Windows PowerShell cmdlet.

Note

Not all Packaged apps are listed in AppLocker's application inventory wizard. Certain application packages are framework packages that are leveraged by other apps. By themselves, these packages cannot do anything, but blocking such packages can inadvertently cause failure for apps that you want to allow. Instead, you can create Allow or Deny rules for the Packaged apps that use these framework packages. The

AppLocker user interface deliberately filters out all the packages that are registered as framework packages.For information about how to create an inventory list, see Creating the List of Applications Deployed to Each Business Group.

For information about how to use the **Get-AppxPackage** Windows PowerShell cmdlet, see the AppLocker PowerShell Command Reference.

For information about creating rules for Packaged apps, see Create a Rule for Packaged Apps.

Consider the following information when you are designing and deploying apps:

- Because AppLocker supports only publisher rules for Packaged apps, collecting the installation path information for Packaged apps is not necessary.
- You cannot create hash- or path-based rules for Packaged apps because all Packaged apps and Packaged app installers are signed by the software publisher of the package. Classic apps were not always consistently signed; therefore, AppLocker has to support hash- or path-based rules.
- By default, if there are no rules in a particular rule collection, AppLocker allows every file that is included in that rule collection. For example, if there are no Windows Installer rules, AppLocker allows all .msi, .msp, and .mst files to run. An existing AppLocker policy that was targeted at computers running Windows Server 2008 R2 and Windows 7 would not have rules for Packaged apps. Therefore, when a computer running Windows Server 2012 or Windows 8 joins a domain where an AppLocker policy is already configured, users would be allowed to run any Packaged app. This might be contrary to your design.

  To prevent all Packaged apps from running on a newly domain-joined computer, by default AppLocker blocks all Packaged apps on a computer running Windows Server 2012 or Windows 8 if the existing domain policy has rules configured in the Exe rule collection. You must take explicit action to allow Packaged apps in your enterprise. You can allow only a select set of Packaged apps. Or if you want to allow all Packaged apps, you can create a default rule for the Packaged apps collection.

Using AppLocker to manage Packaged apps

Just as there are differences in managing each rule collection, you need to manage the Packaged apps with the following strategy:

1. Gather information about which Packaged apps are running in your environment. For information about how to do this, see Creating the List of Applications Deployed to Each Business Group.
2. Create AppLocker rules for specific Packaged apps based on your policy strategies. For more information, see Create a Rule for Packaged Apps and Packaged Apps Default Rules in AppLocker.
3. Continue to update the AppLocker policies as new Package apps are introduced into your environment. To do this, see Add Rules for Packaged Apps to Existing AppLocker Rule-set.
4. Continue to monitor your environment to verify the effectiveness of the rules that are deployed in AppLocker policies. To do this, see Monitor Application Usage with AppLocker.

See Also

Administer AppLocker

Working with AppLocker Rules

This topic describes AppLocker rule types and how to work with them for your application control policies using Windows Server® 2012 and Windows® 8.

Procedures

- Create a Rule That Uses a File Hash Condition
- Create a Rule That Uses a Path Condition
- Create a Rule That Uses a Publisher Condition
- Create a Rule for Packaged Apps
- Create AppLocker Default Rules
- Configure Exceptions for an AppLocker Rule
- Delete an AppLocker Rule
- Edit AppLocker Rules
- Enable the DLL Rule Collection
- Enforce AppLocker Rules
- Run the Automatically Generate Rules Wizard

The three AppLocker enforcement modes are described in the following table. The enforcement mode setting defined here can be overwritten by the setting derived from a linked Group Policy Object (GPO) with a higher precedence.

Enforcement mode	Description
Not configured	This is the default setting which means that the rules defined here will be enforced unless a linked GPO with a higher precedence has a different value for this setting.
Enforce rules	Rules are enforced.
Audit only	Rules are audited but not enforced. When a user runs an application that is affected by an AppLocker rule, the application is allowed to run and the information about the application is added to the AppLocker event log. The Audit-only enforcement mode helps you determine

Enforcement mode	Description
	which applications will be affected by the policy before the policy is enforced. When the AppLocker policy for a rule collection is set to **Audit only**, rules for that rule collection are not enforced

When AppLocker policies from various GPOs are merged, the rules from all the GPOs are merged and the enforcement mode setting of the winning GPO is applied.

For information about GPOs and Group Policy inheritance, see the Group Policy Planning and Deployment Guide http://go.microsoft.com/fwlink/p/?linkid=143138.

Rule collections

The AppLocker user interface is accessed through the Microsoft Management Console (MMC), and it is organized into rule collections, which are Executable files, Scripts, Windows Installer files, Packaged apps and packaged app installers, and DLL files. These collections give the administrator an easy way to differentiate the rules for different types of applications. The following table lists the file formats that are included in each rule collection.

Rule collection	Associated file formats
Executable files	.exe .com
Scripts	.ps1 .bat .cmd .vbs .js
Windows Installer files	.msi .msp .mst
Packaged apps and packaged app installers	.appx
DLL files	.dll .ocx

Important

If you use DLL rules, you need to create an allow rule for each DLL that is used by all of the allowed applications.

When DLL rules are used, AppLocker must check each DLL that an application loads. Therefore, users may experience a reduction in performance if DLL rules are used.

The DLL rule collection is not enabled by default. To learn how to enable the DLL rule collection, see DLL rule collection.

Rule conditions

Rule conditions are criteria that help AppLocker identify the applications to which the rule applies. The three primary rule conditions are publisher, path, and file hash.

- Publisher: Identifies an application based on its digital signature
- Path: Identifies an application by its location in the file system of the computer or on the network
- File hash: Represents the system computed cryptographic hash of the identified file

Publisher

This condition identifies an application based on its digital signature and extended attributes when available. The digital signature contains information about the company that created the application (the publisher). Executable files, Dlls, Windows installers, packaged apps and packaged app installers also have extended attributes, which are obtained from the binary resource. In case of Executable files, Dlls and Windows installers these attributes contain the name of the product that the file is a part of, the original name of the file as supplied by the publisher and the version number of the file. In case of packaged apps and packaged app installers these extended attributes contain the name and the version of the application package.

Note

Rules created in the packaged apps and packaged app installers rule collection can only have publisher conditions since Windows does not support unsigned packaged apps and packaged app installers.

Note

Use a publisher rule condition when possible because they can survive application updates as well as a change in the location of files.

When you select a reference file for a publisher condition, the wizard creates a rule that specifies the publisher, product, file name, and version number. You can make the rule more generic by moving the slider up or by using a wildcard character (*) in the product, file name, or version number fields.

Note

To enter custom values for any of the fields of a publisher rule condition in the Create Rules Wizard, you must select the **Use custom values** check box. When this check box is selected, you cannot use the slider.

The **File version** and **Package version** control whether a user can run a specific version, earlier versions, or later versions of the application. You can choose a version number and then configure the following options:

- **Exactly.** The rule applies only to this version of the application.
- **And above.** The rule applies to this version and all later versions.
- **And below.** The rule applies to this version and all earlier versions.

The following table describes how a publisher condition is applied.

Option		The publisher condition allows or denies…
All signed files		All files that are signed by any publisher.
Publisher only		All files that are signed by the named publisher.
Publisher and product name		All files for the specified product that are signed by the named publisher.
Publisher and product name, and file name		Any version of the named file or package for the named product that are signed by the publisher.
Publisher, product name, file name, and file version	**Exactly**	The specified version of the named file or package for the named product that are signed by the publisher.
Publisher, product name, file name, and file version	**And above**	The specified version of the named file or package and any new releases for the product that are signed by the publisher.
Publisher, product name, file name, and file version	**And below**	The specified version of the named file or package and any earlier versions for the product that are signed by the publisher.

Option	The publisher condition allows or denies…
Custom	You can edit the **Publisher**, **Product name**, **File name**, **Version Package name**, and **Package version** fields to create a custom rule.

Path

This rule condition identifies an application by its location in the file system of the computer or on the network.

AppLocker uses custom path variables for well-known paths, such as Program Files and Windows.

The following table details these path variables.

Windows directory or disk	AppLocker path variable	Windows environment variable
Windows	%WINDIR%	%SystemRoot%
System32	%SYSTEM32%	%SystemDirectory%
Windows installation directory	%OSDRIVE%	%SystemDrive%
Program Files	%PROGRAMFILES%	%ProgramFiles% and %ProgramFiles(x86)%
Removable media (for example, a CD or DVD)	%REMOVABLE%	
Removable storage device (for example, a USB flash drive)	%HOT%	

Important

Because a path rule condition can be configured to include a large number of folders and files, path conditions should be carefully planned. For example, if an allow rule with a path condition includes a folder location that non-administrators are allowed to write data into, a user can copy unapproved files into that location and run the files. For this reason, it is a best practice to not create path conditions for standard user writable locations, such as a user profile.

File hash

When you choose the file hash rule condition, the system computes a cryptographic hash of the identified file. The advantage of this rule condition is that because each file has a unique hash, a file hash rule condition applies to only one file. The disadvantage is that each time the file is updated (such as a security update or upgrade) the file's hash will change. As a result, you must manually update file hash rules.

AppLocker default rules

AppLocker allows you to generate default rules for each rule collection.

Executable default rule types include:

- Allow members of the local **Administrators** group to run all applications.
- Allow members of the **Everyone** group to run applications that are located in the Windows folder.
- Allow members of the **Everyone** group to run applications that are located in the Program Files folder.

Script default rule types include:

- Allow members of the local **Administrators** group to run all scripts.
- Allow members of the **Everyone** group to run scripts that are located in the Program Files folder.
- Allow members of the **Everyone** group to run scripts that are located in the Windows folder.

Windows Installer default rule types include:

- Allow members of the local **Administrators** group to run all Windows Installer files.
- Allow members of the **Everyone** group to run all digitally signed Windows Installer files.
- Allow members of the **Everyone** group to run all Windows Installer files that are located in the Windows\Installer folder.

DLL default rule types:

- Allow members of the local **Administrators** group to run all DLLs.
- Allow members of the **Everyone** group to run DLLs that are located in the Program Files folder.
- Allow members of the **Everyone** group to run DLLs that are located in the Windows folder.

Packaged apps default rule types:

- Allow members of the **Everyone** group to install and run all signed packaged apps and packaged app installers.

AppLocker rule behavior

If no AppLocker rules for a specific rule collection exist, all files with that file format are allowed to run. However, when an AppLocker rule for a specific rule collection is created, only the files explicitly allowed in a rule are permitted to run. For example, if you create an executable rule that

allows .exe files in *%SystemDrive%\FilePath* to run, only executable files located in that path are allowed to run.

A rule can be configured to use allow or deny actions:

- **Allow.** You can specify which files are allowed to run in your environment, and for which users or groups of users. You can also configure exceptions to identify files that are excluded from the rule.
- **Deny.** You can specify which files are *not* allowed to run in your environment, and for which users or groups of users. You can also configure exceptions to identify files that are excluded from the rule.

Important
For a best practice, use allow actions with exceptions. You can use a combination of allow and deny actions but understand that deny actions override allow actions in all cases, and can be circumvented.

Important
If you join a computer running Windows Server 2012 or Windows 8 to a domain that already enforces AppLocker rules for Executables, users will not be able to run any packaged apps unless you also create rules for packaged apps. If you want to allow any packaged apps in your environment while continuing to control Executables, you should create the default rules for packaged apps and set the enforcement mode to Audit-only for the packaged apps rule collection.

Rule exceptions

You can apply AppLocker rules to individual users or to a group of users. If you apply a rule to a group of users, all users in that group are affected by that rule. If you need to allow a subset of a user group to use an application, you can create a special rule for that subset. For example, the rule "Allow Everyone to run Windows except Registry Editor" allows everyone in the organization to run the Windows operating system, but it does not allow anyone to run Registry Editor.

The effect of this rule would prevent users such as Help Desk personnel from running a program that is necessary for their support tasks. To resolve this problem, create a second rule that applies to the Help Desk user group: "Allow Help Desk to run Registry Editor." If you create a deny rule that does not allow any users to run Registry Editor, the deny rule will override the second rule that allows the Help Desk user group to run Registry Editor.

DLL rule collection

Because the DLL rule collection is not enabled by default, you must perform the following procedure before you can create and enforce DLL rules.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To enable the DLL rule collection

1. On the **Start** screen, type **secpol.msc** in the **Search programs and files** box, and then press ENTER.
2. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. In the console tree, double-click **Application Control Policies**, right-click **AppLocker**, and then click **Properties**.
4. Click the **Advanced** tab, select the **Enable the DLL rule collection** check box, and then click **OK**.

   Important
   Before you enforce DLL rules, make sure that there are allow rules for each DLL that is used by any of the allowed applications.

AppLocker wizards

You can create rules by using two AppLocker wizards:

1. The Create Rules Wizard enables you to create one rule at a time.
2. The Automatically Generate Rules Wizard allows you to create multiple rules at one time. You can either select a folder and let the wizard create rules for the relevant files within that folder or in case of packaged apps let the wizard create rules for all packaged apps installed on the computer. You can also specify the user or group to which to apply the rules. This wizard automatically generates allow rules only.

Additional considerations

- By default, AppLocker rules do not allow users to open or run any files that are not specifically allowed. Administrators should maintain an up-to-date list of allowed applications.
- There are two types of AppLocker conditions that do not persist following an update of an application:
  - **File hash condition.** File hash rule conditions can be used with any application because a cryptographic hash value of the application is generated at the time the rule is created. However, the hash value is specific to that exact version of the application. If there are several versions of the application in use within the organization, you need to create file hash conditions for each version in use and for any new versions that are released.
  - **A publisher condition with a specific product version set.** If you create a publisher rule condition that uses the **Exactly** version option, the rule cannot persist if a new version of the application is installed. A new publisher condition must be created, or the version must be edited in the rule to be made less specific.
- If an application is not digitally signed, you cannot use a publisher rule condition for that application.

- AppLocker rules cannot be used to manage computers running a Windows operating system earlier than Windows Server 2008 R2 or Windows 7. Software Restriction Policies must be used instead. If AppLocker rules are defined in a Group Policy Object (GPO), only those rules are applied. To ensure interoperability between Software Restriction Policies rules and AppLocker rules, define Software Restriction Policies rules and AppLocker rules in different GPOs.
- The packaged apps and packaged apps installer rule collection is available only on Windows Server 2012 and Windows 8.
- When the rules for the Executable rule collection are enforced and the packaged apps and packaged app installers rule collection does not contain any rules, no packaged apps and packaged app installers are allowed to run. In order to allow any packaged apps and packaged app installers you must create rules for the packaged apps and packaged app installers rule collection.
- When an AppLocker rule collection is set to **Audit only**, the rules are not enforced. When a user runs an application that is included in the rule, the application is opened and runs normally, and information about that application is added to the AppLocker event log.
- A custom configured URL can be included in the message that is displayed when an application is blocked.
- Expect an increase in the number of Help Desk calls initially because of blocked applications until users understand that they cannot run applications that are not allowed.

See Also

AppLocker Overview

Create a Rule That Uses a File Hash Condition

This topic shows how to create an AppLocker rule with a file hash condition in Windows Server 2012 and Windows 8.

File hash rules use a system-computed cryptographic hash of the identified file.

For information about the file hash condition, see Understanding the File Hash Rule Condition in AppLocker.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To create a new rule with a file hash condition

1. Open the snap-in, and in the console tree, double-click **Application Control Policies**, double-click **AppLocker**, and then click the rule collection that you want to create the rule for.
2. On the **Action** menu, click **Create New Rule**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Permissions** page, select the action (allow or deny) and the user or group that the rule should apply to, and then click **Next**.
5. On the **Conditions** page, select the **File hash** rule condition, and then click **Browse Files** to locate the targeted application file.

   Note
   You can also click **Browse Folders** which calculates the hash for all the appropriate files relative to the rule collection. To remove hashes individually, click the **Remove** button.
6. Click **Next**.
7. On the **Name and Description** page, either accept the automatically generated rule name or type a new rule name, and then click **Create**.

Create a Rule That Uses a Path Condition

This topic shows how to create an AppLocker rule with a path condition in Windows Server 2012 and Windows 8.

The path condition identifies an application by its location in the file system of the computer or on the network.

Important
When creating a rule that uses a deny action, path conditions are less secure for preventing access to a file because a user could easily copy the file to a different location than what is specified in the rule. Because path rules correspond to locations within the file system, you should ensure that there are no subdirectories that are writable by non-administrators. For example, if you create a path rule for C:\ with the allow action, any file within C:\ will be allowed to run, including users' profiles.

For information about the path condition, see Understanding the Path Rule Condition in AppLocker.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To create a new rule with a path condition

1. In the console tree of the snap-in, double-click **Application Control Policies**, double-click **AppLocker**, and then click the rule collection that you want to create the rule for.
2. On the **Action** menu, click **Create New Rule**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Permissions** page, select the action (allow or deny) and the user or group that the rule should apply to, and then click **Next**.
5. On the **Conditions** page, select the **Path** rule condition, and then click **Browse Files** to locate the targeted folder for the application.

   Note
   When you browse to a file or folder location, the wizard automatically converts absolute file paths to use AppLocker path variables. You may edit the path after browsing to specify an absolute path, or you may type the path directly into the **Path** box. To learn more about AppLocker path variables, see Understanding the Path Rule Condition in AppLocker.
6. Click **Next**.
7. (Optional) On the **Exceptions** page, specify conditions by which to exclude files from being affected by the rule. Click **Next**.
8. On the **Name and Description** page, either accept the automatically generated rule name or type a new rule name, and then click **Create**.

Create a Rule That Uses a Publisher Condition

This topic shows how to create an AppLocker rule with a publisher condition in Windows Server 2012 and Windows 8.

You can use publisher conditions only for files that are digitally signed; the publisher condition identifies an application based on its digital signature and extended attributes. The digital signature contains information about the company that created the application (the publisher). The extended attributes, which are obtained from the binary resource, contain the name of the product that the file is part of and the version number of the application. The publisher may be a software development company, such as Microsoft, or the information technology department of your organization.

Packaged app rules are by definition rules that use publisher conditions. For information about creating a packaged app rule, see Create a Rule for Packaged Apps.

For information about the publisher condition, see Understanding the Publisher Rule Condition in AppLocker.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an

AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To create a new rule with a publisher condition

1. In the console tree of the snap-in, double-click **Application Control Policies**, double-click **AppLocker**, and then click the rule collection that you want to create the rule for.
2. On the **Action** menu, click **Create New Rule**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Permissions** page, select the action (allow or deny) and the user or group that the rule should apply to, and then click **Next**.
5. On the **Conditions** page, select the **Publisher** rule condition, and then click **Next**.
6. On the **Publisher** page, click **Browse** to select a signed file, and then use the slider to specify the scope of the rule. To use custom values in any of the fields or to specify a specific file version, select the **Use custom values** check box. For example, you can use the asterisk (*) wildcard character within a publisher rule to specify that any value should be matched.
7. Click **Next**.
8. (Optional) On the **Exceptions** page, specify conditions by which to exclude files from being affected by the rule. Click **Next**.
9. On the **Name and Description** page, either accept the automatically generated rule name or type a new rule name, and then click **Create**.

Create AppLocker Default Rules

This topic describes the steps to create a standard set of AppLocker rules that will allow Windows system files to run in Windows Server 2012 and Windows 8.

AppLocker includes default rules for each rule collection. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed to run.

Important

You can use the default rules as a template when creating your own rules to allow files within the Windows folders to run. However, these rules are only meant to function as a starter policy when you are first testing AppLocker rules. The default rules can be modified in the same way as other AppLocker rule types.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To create default rules

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Right-click the appropriate rule type for which you want to automatically generate default rules. You can automatically generate rules for executable, Windows Installer, script rules and Packaged app rules.
3. Click **Create Default Rules**.

Configure Exceptions for an AppLocker Rule

This topic describes the steps to specify which applications can or cannot run as exceptions to an AppLocker rule in Windows Server 2012 and Windows 8.

Rule exceptions allow you to specify files or folders to exclude from the rule. For more information about exceptions, see Understanding AppLocker Rule Exceptions.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To configure exceptions for a rule

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Expand the rule collection, right-click the rule that you want to configure exceptions for, and then click **Properties**.
3. Click the **Exceptions** tab.
4. In the **Add exception** box, select the rule type that you want to create, and then click **Add**.
   - For a publisher exception, click **Browse**, select the file that contains the publisher to exclude, and then click **OK**.
   - For a path exception, choose the file or folder path to exclude, and then click **OK**.
   - For a file hash exception, edit the file hash rule, and click **Remove**.
   - For Packaged apps exception, click **Add** to create the exceptions based on reference app and rule scope.

Create a Rule for Packaged Apps

This topic shows how to create an AppLocker rule for packaged apps with a publisher condition in Windows Server 2012 and Windows 8.

Packaged apps (also known as Windows 8 apps) are new to Windows Server 2012 and Windows 8. They are based on the new app model that ensures that all the files within an app package share the same identity. Therefore, it is possible to control the entire application using a single AppLocker rule as opposed to the non-packaged apps where each file within the app could have a unique identity. Windows does not support unsigned packaged apps which implies all packaged apps must be signed. AppLocker supports only publisher rules for Packaged apps. A publisher rule for a Packaged app is based on the following information:

- Publisher of the package
- Package name
- Package version

All the files within a package as well as the package installer share these attributes. Therefore, an AppLocker rule for a Packaged app controls both the installation as well as the running of the app. Otherwise, the publisher rules for Packaged apps are no different than the rest of the rule collections; they support exceptions, can be increased or decreased in scope, and can be assigned to users and groups.

For information about the publisher condition, see Understanding the Publisher Rule Condition in AppLocker.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To create a Packaged app rule

1. In the console tree of the snap-in, double-click **Application Control Policies**, double-click **AppLocker**, and then click **Packaged app Rules**.
2. On the **Action** menu, or by right-clicking on **Packaged app Rules**, click **Create New Rule**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Permissions** page, select the action (allow or deny) and the user or group that the rule should apply to, and then click **Next**.
5. On the **Publisher** page, you can select a specific reference for the Packaged app rule and set the scope for the rule. The following table describes the reference options.

Selection	Description	Example
Use an installed packaged app as a	If selected, AppLocker requires you to choose an	You want the Sales group only to use the app named

reference	app that is already installed on which to base your new rule. AppLocker uses the publisher, package name and package version to define the rule.	Microsoft.BingMaps for its outside sales calls. The Microsoft.BingMaps app is already installed on the computer where you are creating the rule, so you choose this option, and select the app from the list of apps installed on the computer and create the rule using this app as a reference.
Use a packaged app installer as a reference	If selected, AppLocker requires you to choose an app installer on which to base your new rule. A packaged app installer has the .appx extension. AppLocker uses the publisher, package name and package version of the installer to define the rule.	Your company has developed a number of internal LoB packaged apps. The app installers are stored on a common file share. Employees can install the required apps from that file share. You want to allow all your employees to install the Payroll app from this share. So you choose this option from the wizard, browse to the file share and choose the installer for the Payroll app as a reference to create your rule.

The following table describes setting the scope for the Package app rule.

Selection	Description	Example
Applies to **Any publisher**	This is the least restrictive scope condition for an **Allow** rule. It permits every packaged app to run or install. Conversely, if this is a **Deny** rule, then this option is the most restrictive because it denies all apps from installing or running.	You want the Sales group to use any packaged app from any signed publisher. You set the Permissions to allow Salesgroup to be able to run any app.
Applies to a specific	This scopes the rule to all	You want to allow all your users

Publisher	apps published by a particular publisher.	to install apps published by the publisher of Microsoft.BingMaps. You could select Microsoft.BingMaps as a reference and choose this rule scope.
Applies to a **Package name**	This scopes the rule to all packages that share the Publisher name and Package name as the reference file.	You want to allow your Sales group to install any version of the Microsoft.BingMaps app. You could select the Microsoft.BingMaps app as a reference and choose this rule scope.
Applies to a **Package version**	This scopes the rule to a particular version of the Package.	You want to be very selective in what you allow. You do not want to implicitly trust all future updates of the Microsoft.BingMaps app. You can limit the scope of your rule to the version of the app currently installed on your reference computer.
Applying custom values to the rule	Selecting Use custom values allows you to adjust the scope fields for your particular circumstance.	You want to allow users to install all Microsoft.Bing* applications which include Microsoft.BingMaps, Microsoft.BingWeather, Microsoft.BingMoney. You can choose the Microsoft.BingMaps as a reference, choose the **Use custom values** scope option and manually edit the Package name field by adding “Microsoft.Bing*” as the Package name.

6. Click **Next**.
7. (Optional) On the **Exceptions** page, specify conditions by which to exclude files from being affected by the rule. This allows you to add exceptions based on the same rule reference and rule scope as you set before. Click **Next**.
8. On the **Name and Description** page, either accept the automatically generated rule name or type a new rule name, and then click **Create**.

Delete an AppLocker Rule

This topic describes the steps to delete an AppLocker rule, which will prevent the application that is specified in the rule from running in Windows Server 2012 and Windows 8.

As older applications are retired and new applications are deployed in your organization, it will be necessary to modify the application control policies. If an application becomes unsupported by the IT department or is no longer allowed due to the organization's security policy, then deleting the rule or rules associated with that application will prevent the application from running.

For information about testing an AppLocker policy to see what rules affect which files or applications, see Test an AppLocker Policy by Using Test-AppLockerPolicy.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To delete a rule in an AppLocker policy

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Click the appropriate rule collection for which you want to delete the rule.
3. In the details pane, right-click the rule to delete, click **Delete**, and then click **Yes**.

Notes

When using Group Policy, for the rule deletion to take effect on computers within the domain, the GPO must be distributed or refreshed.

When this procedure is performed on the local computer, the AppLocker policy takes effect immediately.

Edit AppLocker Rules

This topic describes the steps to edit a publisher rule, path rule, and file hash rule in AppLocker in Windows Server 2012 and Windows 8.

For more information about these rule types, see Understanding AppLocker Rule Condition Types.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To edit a publisher rule

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Click the appropriate rule collection.
3. In the **Action** pane, right-click the publisher rule, and then click **Properties**.
4. Click the appropriate tab to edit the rule properties.
   - Click the **General** tab to change the rule name, add a rule description, configure whether the rule is used to allow or deny applications, and set the security group for which this rule should apply.
   - Click the **Publisher** tab to configure the certificate's common name, the product name, the file name, or file version of the publisher.
   - Click the **Exceptions** tab to create or edit exceptions.
   - When you finish updating the rule, click **OK**.

To edit a file hash rule

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Choose the appropriate rule collection.
3. In the **Action** pane, right-click the file hash rule, and then click **Properties**.
4. Click the appropriate tab to edit the rule properties.
   - Click the **General** tab to change the rule name, add a rule description, configure whether the rule is used to allow or deny applications, and set the security group in which this rule should apply.
   - Click the **File Hash** tab to configure the files that should be used to enforce the rule. You can click **Browse Files** to add a specific file or click **Browse Folders** to add all files in a specified folder. To remove hashes individually, click **Remove**.
   - When you finish updating the rule, click **OK**.

To edit a path rule

1. In the console tree in the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Choose the appropriate rule collection.
3. In the **Action** pane, right-click the path rule, and then click **Properties**.
4. Click the appropriate tab to edit the rule properties.
   - Click the **General** tab to change the rule name, add a rule description, configure whether the rule is used to allow or deny applications, and set the security group in which this rule should apply.
   - Click the **Path** tab to configure the path on the computer in which the rule should be enforced.

- Click the **Exceptions** tab to create exceptions for specific files in a folder.
- When you finish updating the rule, click **OK**.

Enable the DLL Rule Collection

This topic describes the steps to enable the DLL rule collection feature for AppLocker in Windows Server 2012 and Windows 8.

The DLL rule collection includes the .dll and .ocx file formats.

For information about these rules, see DLL Rules in AppLocker.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To enable the DLL rule collection

1. In the console tree in the snap-in, double-click **Application Control Policies**, right-click **AppLocker**, and then click **Properties**.
2. Click the **Advanced** tab, select the **Enable the DLL rule collection** check box, and then click **OK**.

   Important
   Before you enforce DLL rules, make sure that there are allow rules for each DLL that is used by any of the allowed applications.

Enforce AppLocker Rules

This topic describes how to enforce application control rules by using AppLocker in Windows Server 2012 and Windows 8.

After AppLocker rules are created within the rule collection, you can configure the enforcement setting to **Enforce rules** or **Audit only** on the rule collection.

When AppLocker policy enforcement is set to **Enforce rules**, rules are enforced for the rule collection and all events are audited. When AppLocker policy enforcement is set to **Audit only**, rules are only evaluated but all events generated from that evaluation are written to the AppLocker log.

There is no audit mode for the DLL rule collection. DLL rules affect specific applications. Therefore, test the impact of these rules first before deploying them to production.

To enforce AppLocker rules by configuring an AppLocker policy to **Enforce rules**, see Configure an AppLocker Policy for Enforce Rules.

Caution
AppLocker rules will be enforced immediately on the local computer or when the Group Policy object (GPO) is updated by performing this procedure. If you want to see the effect of applying an AppLocker policy before setting the enforcement setting to **Enforce rules**, configure the policy to **Audit only**. For information about how to do this, see Configure an AppLocker Policy for Audit Onlyor Test an AppLocker Policy by Using Test-AppLockerPolicy.

Run the Automatically Generate Rules Wizard

This topic describes steps to run the wizard to create AppLocker rules on a reference computer in Windows Server 2012 and Windows 8.

AppLocker allows you to automatically generate rules for all files within a folder. It will scan the specified folder and create the condition types that you choose for each file in that folder.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To automatically generate rules

1. In the console tree of the snap-in, double-click **Application Control Policies**, and then double-click **AppLocker**.
2. Right-click the appropriate rule type for which you want to automatically generate rules. You can automatically generate rules for executable, Windows Installer, script and Packaged app rules.
3. Click **Automatically Generate Rules**.
4. On the **Folder and Permissions** page, click **Browse** to choose the folder to be analyzed. By default, this is the Program Files folder.
5. Click **Select** to choose the security group in which the default rules should be applied. By default, this is the **Everyone** group.
6. The wizard provides a name in the **Name to identify this set of rules** box based on the name of the folder that you have selected. Accept the provided name or type a different name, and then click **Next**.
7. On the **Rule Preferences** page, choose the conditions that you want the wizard to use while creating rules, and then click **Next**. For more information about rule conditions, see Understanding AppLocker Rule Condition Types.

Notes

The **Reduce the number of rules created by grouping similar files** check box is selected by default. This helps you organize AppLocker rules and reduce the number of rules that you create by performing the following operations for the rule condition that you select:

- One publisher condition is created for all files that have the same publisher and product name.
- One path condition is created for the folder that you select. For example, if you select *C:\Program Files\ProgramName* and the files in that folder are not signed, the wizard creates a rule for *%programfiles%\ProgramName**.
- One file hash condition is created that contains all of the file hashes. When rule grouping is disabled, the wizard creates a file hash rule for each file.

8. Review the files that were analyzed and the rules that will be automatically created. To make changes, click **Previous** to return to the page where you can change your selections. After reviewing the rules, click **Create**.

Note

If you are running the wizard to create your first rules for a GPO, you will be prompted to create the default rules, which allow critical system files to run, after completing the wizard. You may edit the default rules at any time. If your organization has decided to edit the default rules or create custom rules to allow the Windows system files to run, ensure that you delete the default rules after replacing them with your custom rules.

Working with AppLocker Policies

This topic provides links to procedural topics about creating, maintaining, and testing AppLocker policies in Windows Server® 2012 and Windows® 8.

- Configure the Application Identity Service
- Configure an AppLocker Policy for Audit Only
- Configure an AppLocker Policy for Enforce Rules
- Display a Custom URL Message When Users Try to Run a Blocked Application
- Export an AppLocker Policy from a GPO
- Export an AppLocker Policy to an XML File
- Import an AppLocker Policy from Another Computer
- Import an AppLocker Policy into a GPO
- Add Rules for Packaged Apps to Existing AppLocker Rule-set
- Merge AppLocker Policies by Using Set-ApplockerPolicy
- Merge AppLocker Policies Manually
- Refresh an AppLocker Policy
- Test an AppLocker Policy by Using Test-AppLockerPolicy

What does AppLocker do?

AppLocker helps reduce administrative overhead and helps reduce the organization's cost of managing computing resources by decreasing the number of help desk calls that result from users running unapproved applications. AppLocker addresses the following application security scenarios:

- Application inventory

  AppLocker has the ability to enforce its policy in an audit-only mode where all application access activity is registered in event logs. These events can be collected for further analysis. Windows PowerShell cmdlets also help you analyze this data programmatically.

- Protection against unwanted software

  AppLocker has the ability to deny applications from running when you exclude them from the list of allowed applications. Once AppLocker rules are enforced in the production environment any application that is not covered by the allow rules is blocked from executing.

- Licensing conformance

  AppLocker can help you create rules that preclude unlicensed software from running and restricting licensed software to authorized users.

- Software standardization

  AppLocker policies can be configured to allow only supported or approved applications to run on computers within a business group. This permits a more uniform application deployment.

- Manageability improvement

  AppLocker includes a number of improvements in manageability as compared to its predecessor Software Restriction Policies (SRP). Importing and exporting policies, automatic generation of rules from multiple files, audit-only mode deployment and PowerShell cmdlets are a few of the improvements over SRP.

For more information about how AppLocker policies work, see the AppLocker Overview.

For information about Software Restriction Policies and AppLocker policies, see Use AppLocker and Software Restriction Policies in the Same Domain.

See Also

AppLocker Overview

Configure the Application Identity Service

This topic shows how to configure the Application Identity service to start automatically or manually in Windows Server 2012 and Windows 8.

The Application Identity service determines and verifies the identity of an application. Stopping this service will prevent AppLocker policies from being enforced.

Important

When using Group Policy, you must configure it to start automatically in at least one Group Policy Object (GPO) that applies AppLocker rules. This is because AppLocker uses this service to verify the attributes of a file.

To start the Application Identity service automatically using Group Policy

1. On the **Start** screen, type **gpmc.msc** to open the Group Policy Management Console (GPMC).
2. Locate the GPO to edit, right-click the GPO, and click **Edit**.
3. In the console tree under **Computer Configuration\Windows Settings\Security Settings**, click **System Services**.
4. In the details pane, double-click **Application Identity**.
5. In **Application Identity Properties**, configure the service to start automatically.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To start the Application Identity service manually

1. Right-click the taskbar, and click **Start Task Manager**.
2. Click the **Services** tab, right-click **AppIDSvc**, and then click **Start Service**.
3. Verify that the status for the Application Identity service is **Running**.

Configure an AppLocker Policy for Audit Only

This topic describes how to set AppLocker policies to **Audit only** within your IT environment by using AppLocker in Windows Server 2012 and Windows 8.

After AppLocker rules are created within the rule collection, you can configure the enforcement setting to **Enforce rules** or **Audit only**.

When AppLocker policy enforcement is set to **Enforce rules**, rules are enforced for the rule collection and all events are audited. When AppLocker policy enforcement is set to **Audit only**, rules are only evaluated but all events generated from that evaluation are written to the AppLocker log.

Note

There is no audit mode for the DLL rule collection. DLL rules affect specific applications. Therefore, test the impact of these rules first before deploying them to production. To enable the DLL rule collection, see Enable the DLL Rule Collection.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To audit rule collections

1. In the console tree of the snap-in, double-click **Application Control Policies**, right-click **AppLocker**, and then click **Properties**.
2. On the **Enforcement** tab, select the **Configured** check box for the rule collection that you want to enforce, and then verify that **Audit only** is selected in the list for that rule collection.
3. Repeat the above step to configure the enforcement setting to **Audit only** for additional rule collections.
4. Click **OK**.

Configure an AppLocker Policy for Enforce Rules

This topic describes the steps to enable the AppLocker policy enforcement setting in Windows Server 2012 and Windows 8.

Note
When AppLocker policy enforcement is set to **Enforce rules**, rules are enforced for the rule collection and all events are audited.

For information about how AppLocker policies are applied within a GPO structure, see Understanding AppLocker Rules and Enforcement Setting Inheritance in Group Policy.

You can perform this task by using the Group Policy Management Console for an AppLocker policy in a Group Policy Object (GPO) or by using the Local Security Policy snap-in for an AppLocker policy on a local computer or in a security template. For information how to use these MMC snap-ins to administer AppLocker, see Using the MMC snap-ins to administer AppLocker.

To enable the Enforce rules enforcement setting

1. In the console tree of the snap-in under **Computer Configuration\Windows Settings\Security Settings\Application Control Policies**, right-click **AppLocker**, and then click **Properties**.
2. On the **Enforcement** tab of the **AppLocker Properties** dialog box, select the **Configured** check box for the rule collection that you are editing, and then verify that **Enforce rules** is selected.
3. Click **OK**.

For information about viewing the events generated from rules enforcement, see View the AppLocker Log in Event Viewer.

Display a Custom URL Message When Users Try to Run a Blocked Application

This topic describes the steps for displaying a customized message to users when an AppLocker policy denies access to an application in Windows Server 2012 and Windows 8.

Using Group Policy, AppLocker can be configured to display a message with a custom URL. You can use this URL to redirect users to a support site that contains information about why the user received the error and which applications are allowed. If you do not display a custom message when an application is blocked, the default access denied message is displayed.

To complete this procedure, you must have Edit Setting permission to edit a GPO. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission.

To display a custom URL message when users try to run a blocked application

1. On the **Start** screen, type **gpmc.msc** to open the Group Policy Management Console (GPMC).
2. Navigate to the Group Policy Object (GPO) that you want to edit.
3. Right-click the GPO, and then click **Edit**.
4. In the console tree under **Policies\Administrative Templates\Windows Components**, click **Windows Explorer**.
5. In the details pane, double-click **Set a support web page link**.
6. Click **Enabled**, and then type the URL of the custom Web page in the **Support Web page URL** box.
7. Click **OK** to apply the setting.

Export an AppLocker Policy from a GPO

This topic describes the steps to export an AppLocker policy from a Group Policy Object (GPO) so that it can be modified in Windows Server 2012 and Windows 8.

Updating an AppLocker policy that is currently enforced in your production environment can have unintended results. Therefore, export the policy from the GPO and update the rule or rules by using AppLocker on your AppLocker reference or test computer.

To complete this procedure, you must have Edit Setting permission to edit a GPO. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission.

Export the policy from the GPO

1. In the Group Policy Management Console (GPMC), open the GPO that you want to edit.
2. In the console tree under **Computer Configuration\Policies\Windows Settings\Security Settings\Application Control Policies**, click **AppLocker**.
3. Right-click **AppLocker**, and then click **Export Policy**.
4. In the **Export Policy** dialog box, type a name for the exported policy (for example, the name of the GPO), select a location to save the policy, and then click **Save**.
5. The **AppLocker** dialog box will notify you of how many rules were exported. Click **OK**.

Export an AppLocker Policy to an XML File

This topic describes the steps to export an AppLocker policy to an XML file for review or testing in Windows Server 2012 and Windows 8.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To export an AppLocker policy to an XML file

1. On the **Start** screen, type **secpol.msc**.
2. In the console tree, expand **Application Control Policies**, right-click **AppLocker**, and then click **Export Policy**.
3. Browse to the location where you want to save the XML file.
4. In the **File name** box, type a file name for the XML file, and then click **Save**.

Import an AppLocker Policy from Another Computer

This topic describes how to import an AppLocker policy in Windows Server 2012 and Windows 8.

Before completing this procedure, you should have exported an AppLocker policy. For more information, see Export an AppLocker Policy to an XML File.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

Caution
Importing a policy will overwrite the existing policy on that computer.

To import an AppLocker policy

1. On the computer where you want to edit or apply the policy, on the **Start** screen, type **secpol.msc** to open the Local Security Policy snap-in.
2. In the console tree, expand **Application Control Policies**, right-click **AppLocker**, and then click **Import Policy**.
3. In the **Import Policy** dialog box, locate the file that you exported, and then click **Open**.
4. The **Import Policy** dialog box will warn you that importing a policy will overwrite the existing rules and enforcement settings. If acceptable, click **OK** to import and overwrite the policy.
5. The **AppLocker** dialog box will notify you of how many rules were overwritten and imported. Click **OK**.

Import an AppLocker Policy into a GPO

This topic describes the steps to import an AppLocker policy into a Group Policy Object (GPO) in Windows Server 2012 and Windows 8.

AppLocker policies can be created as local security policies and modified like any other local security policy, or they can be created as part of a GPO and managed by using Group Policy. You can create AppLocker policies on any supported computer. For information about which Windows editions are supported, see Requirements to Use AppLocker.

Important
Follow your organization's standard procedures for updating GPOs. For information about specific steps to follow for AppLocker policies, see Maintain AppLocker Policies.

To complete this procedure, you must have Edit Setting permission to edit a GPO. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission.

To import an AppLocker policy into a GPO

1. In the Group Policy Management Console (GPMC), open the GPO that you want to edit.
2. In the console tree under **Computer Configuration\Policies\Windows Settings\Security Settings\Application Control Policies**, click **AppLocker**.
3. Right-click **AppLocker**, and then click **Import Policy**.
4. In the **Import Policy** dialog box, locate the XML policy file, and click **Open**.
5. The **AppLocker** dialog box will notify you of how many rules were imported. Click **OK**.

Add Rules for Packaged Apps to Existing AppLocker Rule-set

This topic describes how to update your existing AppLocker policies for Packaged apps using the Remote Server Administration Toolkit (RSAT) in Windows Server® 2012 and Windows® 8.

You can create Packaged app rules for the computers running Windows Server 2012 or Windows 8 in your domain by updating your existing AppLocker rule set. All you need is a computer running Windows 8. Download and install the Remote Server Administration Toolkit (RSAT) from the Microsoft Download Center.

RSAT comes with the Group Policy Management Console which allows you to edit the GPO or GPOs where your existing AppLocker policy are authored. RSAT for Windows 8 has the necessary files required to author Packaged app rules. Packaged app rules will be ignored on computers running Windows 7 but will be enforced on those computers in your domain running Windows Server 2012 and Windows 8.

See Also

Administer AppLocker

Merge AppLocker Policies by Using Set-ApplockerPolicy

This topic describes the steps to merge AppLocker policies by using Windows PowerShell in Windows Server 2012 and Windows 8.

The **Set-AppLockerPolicy** cmdlet sets the specified Group Policy Object (GPO) to contain the specified AppLocker policy. If no Lightweight Directory Access Protocol (LDAP) is specified, the local GPO is the default. When the Merge parameter is used, rules in the specified AppLocker policy will be merged with the AppLocker rules in the target GPO specified in the LDAP path. The merging of policies will remove rules with duplicate rule IDs, and the enforcement setting specified by the AppLocker policy in the target GPO will be preserved. If the Merge parameter is not specified, then the new policy will overwrite the existing policy.

For information about using **Set-AppLockerPolicy**, including syntax descriptions and parameters, see Set-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169167).

For information about using Windows PowerShell for AppLocker, including how to import the AppLocker cmdlets into Windows PowerShell, see Use the AppLocker Windows PowerShell Cmdlets.

You can also manually merge AppLocker policies. For the procedure to do this, see Merge AppLocker Policies Manually.

To merge a local AppLocker policy with another AppLocker policy by using LDAP paths

1. Open the PowerShell command window. For information about performing Windows PowerShell commands for AppLocker, see Use the AppLocker Windows PowerShell Cmdlets.
2. At the command prompt, type **C:\PS>Get-AppLockerPolicy -Local | Set-AppLockerPolicy -LDAP "LDAP:** *//<string>***" -Merge** where *<string>* specifies the LDAP path of the unique GPO.

Example

Description

Gets the local AppLocker policy, and then merges the policy with the existing AppLocker policy in the GPO specified in the LDAP path.

Code

```
C:\PS>Get-AppLockerPolicy -Local | Set-AppLockerPolicy -LDAP
"LDAP://DC13.Contoso.com/CN={31B2F340-016D-11D2-945F-
00C044FB984F9},CN=Policies,CN=System,DC=Contoso,DC=com" -Merge
```

Merge AppLocker Policies Manually

This procedural topic describes the steps to manually merge AppLocker policies to update the Group Policy Object (GPO) in Windows Server 2012 and Windows 8.

If you have created multiple AppLocker policies and need to merge them to create one AppLocker policy, you can either manually merge the policies or use the Windows PowerShell cmdlets for AppLocker. You cannot automatically merge policies by using the AppLocker snap-in. You must create one rule collection from two or more policies. For information about merging policies by using the cmdlet, see Merge AppLocker Policies by Using Set-ApplockerPolicy.

The AppLocker policy is saved in XML format, and the exported policy can be edited with any text or XML editor. Rule collections are specified within the **RuleCollection Type** element. The XML schema includes five attributes for the different rule collections, as shown in the following table.

Rule collection	RuleCollection Type element
Executable rules	Exe
Windows Installer rules	.msi file
Script rules	Script
DLL rules	Dll
Packaged apps and packaged app installers	Appx

Rule enforcement is specified with the **EnforcementMode** element. The three enforcement modes in the XML correspond to the three enforcement modes in the AppLocker snap-in, as shown in the following table.

XML enforcement mode	Enforcement mode in Group Policy
NotConfigured	Not configured (rules are enforced)
AuditOnly	Audit only
Enabled	Enforce rules

Each of the three condition types use specific elements. For XML examples of the different rule types, see Merge AppLocker Policies Manually in the Windows Server 2008 R2 Technical Library.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To merge two or more AppLocker policies

1. Open an XML policy file in a text editor or XML editor, such as Notepad.
2. Select the rule collection where you want to copy rules from.
3. Select the rules that you want to add to another policy file, and then copy the text.

4. Open the policy where you want to add the copied rules.
5. Select and expand the rule collection where you want to add the rules.
6. At the bottom of the rule list for the collection, after the closing element, paste the rules that you copied from the first policy file. Verify that the opening and closing elements are intact, and then save the policy.
7. Upload the policy to a reference computer to ensure that it is functioning properly within the GPO.

Refresh an AppLocker Policy

This topic describes the steps to force an update for an AppLocker policy in Windows Server 2012 and Windows 8.

If you update the rule collection on a local computer by using the Local Security Policy snap-in, the policy will take effect immediately. If Group Policy is used to distribute the AppLocker policy and you want to immediately implement the policy, you must manually refresh the policy. The Group Policy refresh might take several minutes, depending upon the number of policies within the Group Policy Object (GPO) and the number of target computers.

To use Group Policy to distribute the AppLocker policy change, you need to retrieve the deployed AppLocker policy first. To prepare for the update and subsequent refresh, see Edit an AppLocker Policy and Use the AppLocker Windows PowerShell Cmdlets.

To complete this procedure, you must have Edit Setting permission to edit a GPO. By default, members of the **Domain Admins** group, the **Enterprise Admins** group, and the **Group Policy Creator Owners** group have this permission.

To manually refresh the AppLocker policy by using Group Policy

1. Open a Command Prompt window.
2. At the command prompt, type **gpupdate /force**, and then press ENTER.
3. When the command finishes, close the Command Prompt window, and then verify that the intended rule behavior is correct. You can do this by checking the AppLocker event logs for events that include "policy applied."

To change a policy on an individual computer, or to implement that policy on other computers, without using Group Policy, you first need to update the rule within the rule collection. For information about updating existing rules, see Edit AppLocker Rules. For information about creating a new rule for an existing policy, see:

- Create a Rule That Uses a Publisher Condition

- Create a Rule That Uses a File Hash Condition
- Create a Rule That Uses a Path Condition

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To refresh the AppLocker policy on the local computer

- Update the rule collection by using the Local Security Policy snap-in with one of the following procedures:
  - Edit AppLocker Rules
  - Delete an AppLocker Rule
  - Configure Exceptions for an AppLocker Rule

When finished, the policy is in effect.

To make the same change on another computer, you can use any of the following methods:

- From the computer that you made the change on, export the AppLocker policy, and then import the policy onto the other computer. To do this, use the AppLocker **Export Policy** and **Import Policy** features to copy the rules from the changed computer.

  Caution
  When importing rules from another computer, all the rules will be applied, not just the one that was updated. Merging policies allows both existing and updated (or new) rules to be applied.

- Merge AppLocker policies. For procedures to do this, see Merge AppLocker Policies Manually and Merge AppLocker Policies by Using Set-ApplockerPolicy.

Test an AppLocker Policy by Using Test-AppLockerPolicy

This procedural topic describes the steps to test an AppLocker policy prior to importing it into a Group Policy Object (GPO) or another computer in Windows Server 2012 and Windows 8.

The **Test-AppLockerPolicy** Windows PowerShell cmdlet can be used to determine whether any of the rules in your rule collections will be blocked on your reference computer or the computer on which you maintain policies. Perform the following steps on any computer where the AppLocker policies are applied.

Any user account can be used to complete this procedure.

To test an AppLocker policy by using Test-AppLockerPolicy

1. Export the effective AppLocker policy. To do this, you must use the **Get-**

AppLockerPolicy Windows PowerShell cmdlet.

a. Open a Windows PowerShell Prompt window as an administrator.

b. Before you can use the AppLocker cmdlets, you must import them into Windows PowerShell. To do this, run the following command: `import-module AppLocker`.

c. Use the **Get-AppLockerPolicy** cmdlet to export the effective AppLocker policy to an XML file:

```
Get-AppLockerPolicy –Effective –XML > <PathofFiletoExport.XML>
```

2. Use the **Get-ChildItem** cmdlet to specify the directory that you want to test, specify the **Test-AppLockerPolicy** cmdlet with the XML file from the previous step to test the policy, and use the **Export-CSV** cmdlet to export the results to a file to be analyzed:

```
Get-ChildItem <DirectoryPathtoReview> -Filter <FileExtensionFilter> -Recurse |
Convert-Path | Test-AppLockerPolicy –XMLPolicy <PathToExportedPolicyFile> -User
<domain\username> -Filter <TypeofRuletoFilterFor> | Export-CSV
<PathToExportResultsTo.CSV>
```

The following shows example input for **Test-AppLockerPolicy**:

```
PS C:\ Get-AppLockerPolicy –Effective –XML > C:\Effective.xml

PS C:\ Get-ChildItem 'C:\Program Files\Microsoft Office\' –filter *.exe –Recurse |
Convert-Path | Test-AppLockerPolicy –XMLPolicy C:\Effective.xml –User contoso\zwie –
Filter Denied,DeniedByDefault | Export-CSV C:\BlockedFiles.csv
```

In the example, the effective AppLocker policy is exported to the file C:\Effective.xml. The **Get-ChildItem** cmdlet is used to recursively gather path names for the .exe files in C:\Program Files\Microsoft Office\. The XMLPolicy parameter specifies that the C:\Effective.xml file is an XML AppLocker policy file. By specifying the User parameter, you can test the rules for specific users, and the **Export-CSV** cmdlet allows the results to be exported to a comma-separated file. In the example, `-FilterDenied,DeniedByDefault` displays only those files that will be blocked for the user under the policy.

Packaged Apps and Packaged App Installer Rules in AppLocker

This topic explains the AppLocker rule collection for packaged app installers and packaged apps introduced in Windows Server 2012 and Windows 8.

Commonly known as Windows apps, packaged apps can be installed through the Microsoft AppStore or can be side loaded using the Windows PowerShell cmdlets if you have an Enterprise license. Packaged apps can be installed by a standard user unlike some desktop applications that sometimes require administrative privileges for installation. In this topic, desktop applications refer to Win32 apps that run on the classic user desktop.

Typically, an app consists of multiple components – the installer used to install the app and one or more Exes, Dlls or Scripts. With desktop applications, not all those components always share

common attributes such as the publisher name, product name and product version. Therefore, AppLocker has to control each of these components separately through different rule collections – Exe, Dll, Script and Windows Installers. In contrast, all the components of a packaged app share the same attributes: Publisher name, Package name and Package version. It is therefore possible to control an entire app with a single rule.

In Windows Server 2012 and Windows 8, AppLocker enforces rules for packaged apps separately from desktop applications. A single AppLocker rule for a packaged app can control both the installation and the running of an app. Because all packaged apps are signed, AppLocker supports only publisher rules for packaged apps. A publisher rule for a packaged app is based on the following attributes of the app:

- Publisher name
- Package name
- Package version

In summary, including AppLocker rules for packaged apps in your policy design provides:

- The ability to control the installation and the running the packaged app
- The ability to control all the components of the app with a single rule rather than controlling individual binaries within the app
- The ability to create application control policies that survive application updates
- Management of packaged apps through Group Policy.

See also

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012 and Windows 8
Product evaluation	Frequently Asked Questions AppLocker Step-by-Step Guide	AppLocker Overview
Procedures	AppLocker Operations Guide	Administer AppLocker Manage Packaged Apps with AppLocker
Scripting	Using the AppLocker Windows PowerShell Cmdlets	Using the AppLocker Windows PowerShell Cmdlets
Technical content	AppLocker Technical Reference	AppLocker Technical Reference
Design, planning and deployment	AppLocker Policies Design Guide AppLocker Policies Deployment Guide	AppLocker Policies Design Guide AppLocker Policies Deployment Guide

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012 and Windows 8
General information and additional resources	AppLocker Documentation for Windows 7 and Windows Server 2008 R2	**AppLocker Navigation Page**

AppLocker Policies Design Guide

This topic for the IT professional introduces the design and planning steps required to deploy application control policies by using AppLocker which was introduced in Windows Server® 2008 R2 and Windows® 7.

Purpose of this guide

This guide provides important designing and planning information for deploying application control policies by using AppLocker. It is intended for security architects, security administrators, and system administrators. Through a sequential and iterative process, you can create an AppLocker policy deployment plan for your organization that will address your specific application control requirements by department, organizational unit, or business group.

This guide does not cover the deployment of application control policies by using Software Restriction Policies (SRP). However, SRP is discussed as a deployment option in conjunction with AppLocker policies. For information about these options, see Determine Your Application Control Objectives.

To understand if AppLocker is the correct application control solution for your organization, see Understand AppLocker Policy Design Decisions.

Tip

To digitally save or print pages from this library, click **Export** (in the upper-right corner of the page), and then follow the instructions.

Contents of this guide

This guide contains the following topics:

- Understand the AppLocker Policy Deployment Process
- Understand AppLocker Policy Design Decisions
- Determine Your Application Control Objectives
- Create List of Applications Deployed to Each Business Group
- Select Types of Rules to Create
- Determine Group Policy Structure and Rule Enforcement

- Plan for AppLocker Policy Management
- Create Your AppLocker Planning Document

After careful design and detailed planning, the next step is to deploy AppLocker policies. AppLocker Policies Deployment Guide covers the creation and testing of policies, deploying the enforcement setting, and managing and maintaining the policies.

Understand the AppLocker Policy Deployment Process

This planning and deployment topic for the IT professional describes the process for using AppLocker when deploying application control policies.

To successfully deploy AppLocker policies, you need to identify your application control objectives and construct the policies for those objectives. The key to the process is taking an accurate inventory of your organization's applications, which requires investigation of all the targeted business groups. With an accurate inventory, you can create rules and set enforcement criteria that will allow the organization to use the required applications and allow the IT department to manage a controlled set of applications.

The following diagram shows the main points in the design, planning, and deployment process for AppLocker.

Planning an AppLocker Deployment

1 Determine how to implement AppLocker
- Determine whether interoperability with Software Restriction Policies (SRP) is needed.
- Document where to use AppLocker.
- Select whether to use allow actions only or allow and deny actions.

2 Create a list of applications
- Create an inventory of applications installed in different OUs.

3 Select the types of rules to create
- Select which rule collections to use.
- Determine how to allow system files to run.

4 Define the Group Policy structure and rule enforcement
- Select enforcement settings for each OU.
- Determine rule and enforcement setting inheritance in Group Policy.

5 Create a process for managing rules
- Document an end-user support policy for blocked applications.
- Determine whether to use event forwarding.
- Determine how to maintain the AppLocker policy.

6 Document your AppLocker plan
- Document your AppLocker plan to use when deploying AppLocker and for future reference.

Deploying AppLocker

1 Author rules in a test environment
Using your documented AppLocker plan:
- Create a reference computer.
- Automatically generate rules.
- Modify the rule set.

2 Test and update the policy
- Enable **Audit only** enforcement.
- Start the Application Identity service.
- Test the policy.
- Analyze events.
- Modify the policy.
- Repeat the testing process until you are satisfied with the results.

3 Deploy the production enforcement setting
- Deploy the AppLocker enforcement settings to be used in production.

4 Maintain the policy
- Add new rules for new applications.
- Edit existing rules for revisions to organization policy.
- Provide end-user support.

Resources to support the deployment process

The following documentation contains information about designing, planning, deploying, and maintaining AppLocker policies:

- For information about the AppLocker policy design and planning requirements and process, see AppLocker Policies Design Guide.
- For information about the AppLocker policy deployment requirements and process, see AppLocker Policies Deployment Guide.
- For information about AppLocker policy maintenance and monitoring for Windows Server 2008 R2 and Windows 7, see AppLocker Operations Guide.

  For information about AppLocker policy maintenance and monitoring in versions of Windows later than Windows Server 2008 R2 and Windows 7, see Administer AppLocker.
- For information about AppLocker policy architecture, components, and processing, see AppLocker Technical Reference.

Understand AppLocker Policy Design Decisions

This topic for the IT professional lists the design questions, possible answers, and ramifications of the decisions when you plan a deployment of application control policies by using AppLocker within a Windows operating system environment.

When you begin the design and planning process, you should consider the ramifications of your design choices. The resulting decisions will affect your policy deployment scheme and subsequent application control policy maintenance.

You should consider using AppLocker as part of your organization's application control policies if all the following are true:

- You have deployed or plan to deploy the supported versions of Windows in your organization. For specific operating system version requirements, see Requirements to Use AppLocker.
- You need improved control over the access to your organization's applications and the data your users access.
- The number of applications in your organization is known and manageable.
- You have resources to test policies against the organization's requirements.
- You have resources to involve Help Desk or to build a self-help process for end-user application access issues.
- The group's requirements for productivity, manageability, and security can be controlled by restrictive policies.

The following questions are not in priority or sequential order. They should be considered when you deploy application control policies (as appropriate for your targeted environment).

Which applications do you need to control in your organization?

You might need to control a limited number of applications because they access sensitive data, or you might have to exclude all applications except those that are sanctioned for business purposes. There might be certain business groups that require strict control, and others that promote independent application usage.

Possible answers	Design considerations
Control all applications	AppLocker policies control applications by creating an allowed list of applications by file type. Exceptions are also possible. AppLocker policies can only be applied to applications installed on computers running one of the supported versions of Windows. For specific operating system version requirements, see Requirements to Use AppLocker. If versions of Windows operating systems earlier than Windows Server 2008 R2 or Windows 7 are also deployed, you can use Software Restriction Policies (SRP) with AppLocker.
Control specific applications	When you create AppLocker rules, a list of allowed applications is created. All applications on that list will be allowed to run (except those on the exception list). Applications that are not on the list will be prevented from running. AppLocker policies can only be applied to applications installed on computers running any of the supported versions of Windows. For specific operating system version requirements, see Requirements to Use AppLocker.
Control only classic apps, only Windows Store apps, or both	AppLocker policies control applications by creating an allowed list of applications by file type. Because Windows Store apps are categorized under the Publisher condition, desktop apps and Windows Store apps can be controlled together. AppLocker policies for Windows Store apps can be applied only to applications that are installed on computers that support the Windows Store, but desktop applications can be controlled with AppLocker on all supported versions of Windows. The rules you currently have configured for classic

Possible answers	Design considerations
	desktop applications can remain, and you can create new ones for Windows Store apps. For a comparison of desktop applications and Windows Store apps, see Comparing classic desktop applications and Windows Store apps for AppLocker policy design decisions in this topic.
Control applications by business group and user	Similar to SRP, AppLocker policies can be applied through a Group Policy Object (GPO) to computer objects within an organizational unit (OU). Individual AppLocker rules can be applied to individual users or to groups of users. SRP policies can be used without modification for Windows operating systems earlier than Windows Server 2008 R2 and Windows 7, but there are special considerations when they are used in the same GPO as AppLocker policies.
Control applications by computer, not user	AppLocker is a computer-based policy implementation. If your domain or site organizational structure is not based on a logical user structure, such as an OU, you might want to set up that structure before you begin your AppLocker planning. Otherwise, you will have to identify users, their computers, and their application access requirements.
Understand application usage, but there is no need to control any applications yet	AppLocker policies can be set to audit application usage to help you track which applications are used in your organization. You can then use the AppLocker event log to create AppLocker policies.

Important

The following list contains files or types of files that cannot be managed by AppLocker:

- AppLocker does not protect against running 16-bit DOS binaries in a NT Virtual DOS Machine (NTVDM). This technology allows running legacy DOS and 16-bit Windows programs on computers that are using Intel 80386 or higher when there is already another operating system running and controlling the hardware. The result is that 16-bit binaries can still run on Windows Server 2008 R2 and Windows 7 when AppLocker is configured to

otherwise block binaries and libraries. If it is a requirement to prevent 16-bit applications from running, you must configure the Deny rule in the Executable rule collection for NTVDM.exe.

- You cannot use AppLocker (or Software Restriction Policies) to prevent code from running outside the Win32 subsystem. In particular, this applies to the (POSIX) subsystem in Windows NT. If it is a requirement to prevent applications from running in the POSIX subsystem, you must disable the subsystem. For more information, see Disabling the POSIX Subsystem.
- AppLocker can only control VBScript, JScript, .bat files, .cmd files and Windows PowerShell scripts. It does not control all interpreted code that runs within a host process, for example Perl scripts and macros. Interpreted code is a form of executable code that runs within a host process. For example, Windows batch files (*.bat) run within the context of the Windows Command Host (cmd.exe). To use AppLocker to control interpreted code, the host process must call AppLocker before it runs the interpreted code, and then enforce the decision that is returned by AppLocker. Not all host processes call into AppLocker. Therefore, AppLocker cannot control every kind of interpreted code, for example Microsoft Office macros.

  Important
  You should configure the appropriate security settings of these host processes if you must allow them to run. For example, configure the security settings in Microsoft Office to ensure that only signed and trusted macros are loaded.

- AppLocker rules allow or prevent an application from launching. AppLocker does not control the behavior of applications after they are launched. Applications could contain flags that are passed to functions that signal AppLocker to circumvent the rules and allow another .exe or .dll file to be loaded. In practice, an application that is allowed by AppLocker could use these flags to bypass AppLocker rules and launch child processes. You must follow a process that best suits your needs to thoroughly vet each application before allowing them to run using AppLocker rules.

  For more information, see Security Considerations for AppLocker.

Comparing classic desktop applications and Windows Store apps for AppLocker policy design decisions

AppLocker policies for Windows Store apps can only be applied to applications that are installed on computers running Windows operating systems that support Windows Store apps. However, desktop applications can be controlled in Windows Server 2008 R2 and Windows 7, in addition to those computers that support Windows Store apps. The rules for desktop applications and Windows Store apps can be enforced together. The differences you should consider for Windows Store apps are:

- All Windows Store apps can be installed by a standard user, whereas a number of classic desktop applications require administrative credentials to install. So in an environment where most of the users are standard users, you might not need numerous Exe rules, but you might want more explicit policies for packaged apps.
- Desktop applications can be written to change the system state if they run with administrative credentials. Most Windows Store apps cannot change the system state because they run

with limited permissions. When you design your AppLocker policies, it is important to understand whether an app that you are allowing can make system-wide changes.

- Windows Store apps can be acquired through the Microsoft AppStore, or they can be side-loaded by using Windows PowerShell cmdlets. If you use Windows PowerShell cmdlets, a special Enterprise license is required to acquire Windows Store apps. Desktop applications can be acquired through traditional means, such as through software vendors or retail distribution.

AppLocker controls Windows Store apps and desktop applications by using different rule collections. You have the choice to control the Windows Store app, the desktop applications, or both.

For more information, see Packaged Apps and Packaged App Installer Rules in AppLocker.

How do you currently control application usage in your organization?

Most organizations have evolved application control policies and methods over time. With heightened security concerns and an emphasis on tighter IT control over desktop use, your organization might decide to consolidate application control practices or design a comprehensive application control scheme. AppLocker includes improvements over SRP in the architecture and management of application control policies.

Possible answers	Design considerations
Security polices (locally set or through Group Policy)	Using AppLocker requires increased effort in planning to create correct policies, but this results in a simpler distribution method.
Specific SRP polices	Using AppLocker does not require rewriting SRP rules for computers running versions of Windows operating systems earlier than Windows Server 2008 R2 or Windows 7. However, you must evaluate existing SRP rules if you use AppLocker with SRP in the supported versions of Windows. For specific operating system version requirements, see Requirements to Use AppLocker.
Non-Microsoft application control software	Using AppLocker requires a complete application control policy evaluation and implementation.
Managed usage by group or OU	Using AppLocker requires a complete application control policy evaluation and implementation.
Authorization Manager or other role-based	Using AppLocker requires a complete

Possible answers	Design considerations
access technologies	application control policy evaluation and implementation.
Other	Using AppLocker requires a complete application control policy evaluation and implementation.

Which Windows desktop and server operating systems are running in your organization?

If your organization supports multiple Windows operating systems, application control policy planning becomes more complex. Your initial design decisions should consider the security and management priorities of applications that are installed on each version of the operating system.

Possible answers	Design considerations
Your organization's computers are running a combination of the following operating systems: • Windows 8 • Windows 7 • Windows Vista • Windows XP • Windows Server 2012 • Windows Server 2008 R2 • Windows Server 2008 • Windows Server 2003	AppLocker rules are only applied to computers running the supported versions of Windows, but SRP rules can be applied to all versions of Windows beginning with Windows XP and Windows Server 2003. For specific operating system version requirements, see Requirements to Use AppLocker. **Note** If you are using the Basic User security level as assigned in SRP, those privileges are not supported on computers running that support AppLocker. AppLocker policies as applied through a GPO take precedence over SRP policies in the same or linked GPO. SRP policies can be created and maintained the same way.
Your organization's computers are running only the following operating systems: • Windows Server 2012 R2 • Windows Server 2012 • Windows Server 2008 R2 • Windows 8.1	Use AppLocker to create your application control policies.

Possible answers	Design considerations
• Windows 8 • Windows 7	

Are there specific groups in your organization that need customized application control policies?

Most business groups or departments have specific security requirements that pertain to data access and the applications used to access that data. You should consider the scope of the project for each group and the group's priorities before you deploy application control policies for the entire organization.

Possible answers	Design considerations
Yes	For each group, you need to create a list that includes their application control requirements. Although this may increase the planning time, it will most likely result in a more effective deployment. If your GPO structure is not currently configured so that you can apply different policies to specific groups, you can alternatively apply AppLocker rules in a GPO to specific user groups.
No	AppLocker policies can be applied globally to applications that are installed on computers running the supported versions of Windows as listed in Requirements to Use AppLocker. Depending on the number of applications you need to control, managing all the rules and exceptions might be challenging.

Does your IT department have resources to analyze application usage, and to design and manage the policies?

The time and resources that are available to you to perform the research and analysis can affect the detail of your plan and processes for continuing policy management and maintenance.

Possible answers	Design considerations
Yes	Invest the time to analyze your organization's

Possible answers	Design considerations
	application control requirements, and plan a complete deployment that uses rules that are as simply constructed as possible.
No	Consider a focused and phased deployment for specific groups by using a small number of rules. As you apply controls to applications in a specific group, learn from that deployment to plan your next deployment.

Does your organization have Help Desk support?

Preventing your users from accessing known, deployed, or personal applications will initially cause an increase in end-user support. It will be necessary to address the various support issues in your organization so security policies are followed and business workflow is not hampered.

Possible answers	Design considerations
Yes	Involve the support department early in the planning phase because your users may inadvertently be blocked from using their applications, or they may seek exceptions to use specific applications.
No	Invest time in developing online support processes and documentation before deployment.

Do you know what applications require restrictive policies?

Any successful application control policy implementation is based on your knowledge and understanding of application usage within the organization or business group. In addition, the application control design is dependent on the security requirements for data and the applications that access that data.

Possible answers	Design considerations
Yes	You should determine the application control priorities for a business group and then attempt to design the simplest scheme for their application control policies.

Possible answers	Design considerations
No	You will have to perform an audit and requirements gathering project to discover the application usage. AppLocker provides the means to deploy policies in **Audit only** mode, and tools to view the event logs.

How do you deploy or sanction applications (upgraded or new) in your organization?

Implementing a successful application control policy is based on your knowledge and understanding of application usage within the organization or business group. In addition, the application control design is dependent on the security requirements for data and the applications that access that data. Understanding the upgrade and deployment policy will help shape the construction of the application control policies.

Possible answers	Design considerations
Ad hoc	You need to gather requirements from each group. Some groups might want unrestricted access or installation, while other groups might want strict controls.
Strict written policy or guidelines to follow	You need to develop AppLocker rules that reflect those policies, and then test and maintain the rules.
No process in place	You need to determine if you have the resources to develop an application control policy, and for which groups.

Does your organization already have SRP deployed?

Although SRP and AppLocker have the same goal, AppLocker is a major revision of SRP.

Possible answers	Design considerations
Yes	You cannot use AppLocker to manage SRP settings, but you can use SRP to manage application control policies on computers running on any of the supported operating systems listed in Requirements to Use AppLocker. In addition, if AppLocker and SRP

Possible answers	Design considerations
	settings are configured in the same GPO, only the AppLocker settings will be enforced on computers running those supported operating systems. **Note** If you are using the Basic User security level as assigned in SRP, those permissions are not supported on computers running the supported operating systems.
No	Policies that are configured for AppLocker can only be applied to computers running the supported operating systems, but SRP is also available on those operating systems.

What are your organization's priorities when implementing application control policies?

Some organizations will benefit from application control policies as shown by an increase in productivity or conformance, while others will be hindered in performing their duties. Prioritize these aspects for each group to allow you to evaluate the effectiveness of AppLocker.

Possible answers	Design considerations
Productivity: The organization assures that tools work and required applications can be installed.	To meet innovation and productivity goals, some groups require the ability to install and run a variety of software from different sources, including software that they developed. Therefore, if innovation and productivity is a high priority, managing application control policies through an allowed list might be time consuming and an impediment to progress.
Management: The organization is aware of and controls the applications it supports.	In some business groups, application usage can be managed from a central point of control. AppLocker policies can be built into a GPO for that purpose. This shifts the burden of application access to the IT department, but it also has the benefit of controlling the number of applications that can be run and controlling the

Possible answers	Design considerations
	versions of those applications.
Security: The organization must protect data in part by ensuring that only approved applications are used.	AppLocker can help protect data by allowing a defined set of users access to applications that access the data. If security is the top priority, the application control policies will be the most restrictive.

How are applications currently accessed in your organization?

AppLocker is very effective for organizations that have application restriction requirements if they have environments with a simple topography and application control policy goals that are straightforward. For example, AppLocker can benefit an environment where non-employees have access to computers that are connected to the organizational network, such as a school or library. Large organizations also benefit from AppLocker policy deployment when the goal is to achieve a detailed level of control on the desktop computers with a relatively small number of applications to manage, or when the applications are manageable with a small number of rules.

Possible answers	Design considerations
Users run without administrative rights. Applications are installed by using an installation deployment technology.	AppLocker can help reduce the total cost of ownership for business groups that typically use a finite set of applications, such as human resources and finance departments. At the same time, these departments access highly sensitive information, much of which contains confidential and proprietary information. By using AppLocker to create rules for specific applications that are allowed to run, you can help limit unauthorized applications from accessing this information. **Note** AppLocker can also be effective in helping create standardized desktops in organizations where users run as administrators. However, it is important to note that users with administrativecredentials can add new rules to the local AppLocker policy.
Users must be able to install applications as	Enforcing AppLocker rules is not suited for business groups that must be able to install

Possible answers	Design considerations
needed. Users currently have administrator access, and it would be difficult to change this.	applications as needed and without approval from the IT department. If one or more OUs in your organization has this requirement, you can choose not to enforce application rules in those OUs by using AppLocker or to implement the **Audit only** enforcement setting through AppLocker.

Is the structure in Active Directory Domain Services based on the organization's hierarchy?

Designing application control policies based on an organizational structure that is already built into Active Directory Domain Services (AD DS) is easier than converting the existing structure to an organizational structure. Because the effectiveness of application control policies is dependent on the ability to update policies, consider what organizational work needs to be accomplished before deployment begins.

Possible answers	Design considerations
Yes	AppLocker rules can be developed and implemented through Group Policy, based on your AD DS structure.
No	The IT department must create a scheme to identify how application control policies can be applied to the correct user or computer.

Record your findings

The next step in the process is to record and analyze your answers to the preceding questions. If AppLocker is the right solution for your goals, tyou can set your application control policy objectives and plan your AppLocker rules. This process culminates in creating your planning document.

- For information about setting your policy goals, see Determine Your Application Control Objectives.
- For information about creating your planning document, see Create Your AppLocker Planning Document.

Determine Your Application Control Objectives

This topic helps you with the decisions you need to make to determine what applications to control and how to control them by comparing Software Restriction Policies and AppLocker.

AppLocker is very effective for organizations with application restriction requirements whose environments have a simple topography and the application control policy goals are straightforward. For example, AppLocker can benefit an environment where non-employees have access to computers connected to the organizational network, such as a school or library. Large organizations also benefit from AppLocker policy deployment when the goal is to achieve a detailed level of control on the desktop computers that they manage for a relatively small number of applications.

There are management and maintenance costs associated with a list of allowed applications. In addition, the purpose of application control policies is to allow or prevent employees from using applications that might actually be productivity tools. Keeping employees or users productive while implementing the policies can cost time and effort. Lastly, creating user support processes and network support processes to keep the organization productive are also concerns.

Use the following table to develop your own objectives and determine which application control feature best addresses those objectives.

Application control function	SRP	AppLocker
Scope	SRP policies can be applied to all Windows operating systems beginning with Windows XP and Windows Server 2003.	AppLocker policies apply only to the support versions of Windows listed in Requirements to Use AppLocker.
Policy creation	SRP policies are maintained through Group Policy and only the administrator of the GPO can update the SRP policy. The administrator on the local computer can modify the SRP policies defined in the local GPO.	AppLocker policies are maintained through Group Policy and only the administrator of the GPO can update the policy. The administrator on the local computer can modify the AppLocker policies defined in the local GPO. AppLocker permits customization of error messages to direct users to a Web page for help.

Application control function	SRP	AppLocker
Policy maintenance	SRP policies must be updated by using the Local Security Policy snap-in (if the policies are created locally) or the Group Policy Management Console (GPMC).	AppLocker policies can be updated by using the Local Security Policy snap-in (if the policies are created locally), or the GPMC, or the Windows PowerShell AppLocker cmdlets.
Policy application	SRP policies are distributed through Group Policy.	AppLocker policies are distributed through Group Policy.
Enforcement mode	SRP works in the “deny list mode” where administrators can create rules for files that they do not want to allow in this Enterprise whereas the rest of the file are allowed to run by default. SRP can also be configured in the “allow list mode” such that the by default all files are blocked and administrators need to create allow rules for files that they want to allow.	AppLocker by default works in the “allow list mode” where only those files are allowed to run for which there is a matching allow rule.
File types that can be controlled	SRP can control the following file types: • Executables • Dlls • Scripts • Windows Installers SRP cannot control each file type separately. All SRP rules are in a single rule collection.	AppLocker can control the following file types: • Executables • Dlls • Scripts • Windows Installers • Packaged apps and installers (but not in Windows Server 2008 R2 and Windows 7). AppLocker maintains a separate rule collection for each of the five file types.
Designated file types	SRP supports an extensible list of file types that are considered executable. Administrators can	AppLocker does not support this. AppLocker currently supports the following file

Application control function	SRP	AppLocker
	add extensions for files that should be considered executable.	extensions: • Executables (.exe, .com) • Dlls (.ocx, .dll) • Scripts (.vbs, .js, .ps1, .cmd, .bat) • Windows Installers (.msi, .mst, .msp) • Packaged app installers (.appx)
Rule types	SRP supports four types of rules: • Hash • Path • Signature • Internet zone	AppLocker supports three types of rules: • Hash • Path • Publisher
Editing the hash value	SRP allows administrators to select a file to hash.	AppLocker computes the hash value itself. Internally it uses the SHA2 Authenticode hash for Portable Executables (Exe and Dll) and Windows Installers and a SHA2 flat file hash for the rest.
Support for different security levels	With SRP administrators can specify the permissions with which an app can run. So, an administrator can configure a rule such that notepad always runs with restricted permissions and never with administrative privileges. SRP on Windows Vista and earlier supported multiple security levels. On Windows 7 that list was restricted to just two levels: Disallowed and Unrestricted (Basic User translates to Disallowed).	AppLocker does not support security levels.
Manage Packaged apps and	Unable	.appx is a valid file type which

Application control function	SRP	AppLocker
Packaged app installers.		AppLocker can manage.
Targeting a rule to a user or a group of users	SRP rules apply to all users on a particular computer.	AppLocker rules can be targeted to a specific user or a group of users.
Support for rule exceptions	SRP does not support rule exceptions	AppLocker rules can have exceptions which allow administrators to create rules such as “Allow everything from Windows except for Regedit.exe”.
Support for audit mode	SRP does not support audit mode. The only way to test SRP policies is to set up a test environment and run a few experiments.	AppLocker supports audit mode which allows administrators to test the effect of their policy in the real production environment without impacting the user experience. Once you are satisfied with the results, you can start enforcing the policy.
Support for exporting and importing policies	SRP does not support policy import/export.	AppLocker supports the importing and exporting of policies. This allows you to create AppLocker policy on a sample computer, test it out and then export that policy and import it back into the desired GPO.
Rule enforcement	Internally, SRP rules enforcement happens in the user-mode which is less secure.	Internally, AppLocker rules for Exes and Dlls are enforced in the kernel-mode which is more secure than enforcing them in the user-mode.

For information about the capabilities of Software Restriction Policies, see Software Restriction Policies Technical Overview.

For information about the capabilities of AppLocker, see AppLocker Overview.

Create List of Applications Deployed to Each Business Group

This topic describes the process of gathering application usage requirements from each business group in order to implement application control policies by using AppLocker.

Determining application usage

For each business group, determine the following:

- The complete list of applications used, including different versions of an application.
- The full installation path of the application.
- The publisher and signed status of each application.
- The type of requirement the business groups set for each application, such as business critical, business productivity, optional, or personal. It might also be helpful during this effort to identify which applications are supported or unsupported by your IT department, or supported by others outside your control.
- A list of files or applications that require administrative credentials to install or run. If the file requires administrative credentials to install or run, users who cannot provide administrative credentials will be prevented from running the file even if the file is explicitly allowed by an AppLocker policy. Even with AppLocker policies enforced, only members of the Administrators group can install or run files that require administrative credentials.

How to perform the application usage assessment

Although you might already have a method in place to understand application usage for each business group, you will need to use this information to help create your AppLocker rule collection. AppLocker includes the Automatically Generate Rules wizard and the **Audit only** enforcement configuration to assist you with planning and creating your rule collection.

Application inventory methods

Using the Automatically Generate Rules wizard quickly creates rules for the applications you specify. The wizard is designed specifically to build a rule collection. You can use the Local Security Policy snap-in to view and edit the rules. This method is very useful when creating rules from a reference computer, and when creating and evaluating AppLocker policies in a testing environment. However, it does require that the files be accessible on the reference computer or through a network drive. This might mean additional work in setting up the reference computer and determining a maintenance policy for that computer.

Using the **Audit only** enforcement method permits you to view the logs because it collects information about every process on the computers receiving the Group Policy Object (GPO). Therefore, you can see what the enforcement will be on the computers in a business group. AppLocker includes Windows PowerShell cmdlets that you can use to analyze the events from the event log and cmdlets to create rules. However, when you use Group Policy to deploy to several computers, a means to collect events in a central location is very important for

manageability. Because AppLocker logs information about files that users or other processes start on a computer, you could miss creating some rules initially. Therefore, you should continue your evaluation until you can verify that all required applications that are allowed to run are accessed successfully.

Tip

If you run Application Verifier against a custom application with any AppLocker policies enabled, it might prevent the application from running. You should either disable Application Verifier or AppLocker.

You can create an inventory of Windows 8 apps on a computer by using two methods: the **Get-AppxPackage** Windows PowerShell cmdlet or the AppLocker UI.

The following topics in the AppLocker Step-by-Step Guide describe how to perform each method:

- Automatically Generating Executable Rules from a Reference Computer
- Using Auditing to Track Which Applications Are Used

Prerequisites to completing the inventory

Identify the business group and each organizational unit (OU) within that group to which you will apply application control policies. In addition, you should have identified whether or not AppLocker is the most appropriate solution for these policies. For information about these steps, see the following topics:

- Understand AppLocker Policy Design Decisions
- Determine Your Application Control Objectives

Next steps

Identify and develop the list of applications. Record the name of the application, whether it is signed or not as indicated by the publisher's name, and whether or not it is a mission critical, business productivity, optional, or personal application. Record the installation path of the applications. For information about how to do this, see Document Your Application List.

After you have created the list of applications, the next step is to identify the rule collections, which will become the policies. This information can be added to the table under columns labeled:

- Use default rule or define new rule condition
- Allow or deny
- GPO name

To do this, see the following topics:

- Select Types of Rules to Create
- Determine Group Policy Structure and Rule Enforcement

Document Your Application List

This planning topic describes the application information that you should document when you create a list of applications for AppLocker policies.

Record your findings

Applications

Record the name of the application, whether it is signed as indicated by the publisher's name, and whether it is a mission critical, business productivity, optional, or personal application. Later, as you manage your rules, AppLocker displays this information in the format shown in the following example: *MICROSOFT OFFICE INFOPATH signed by O=MICROSOFT CORPORATION, L=REDMOND, S=WASHINGTON, C=US*.

Installation path

Record the installation path of the applications. For example, Microsoft Office 2007 installs files to *%programfiles%\Microsoft Office\Office12*, which is *C:\Program Files\Microsoft Office\Office12* on most computers.

The following table provides an example of how to list applications for each business group at the early stage of designing your application control policies. Eventually, as more planning information is added to the list, the information can be used to build AppLocker rules.

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path
Bank Tellers	Teller-East and Teller-West	Yes	Teller Software	C:\Program Files\Woodgrove\Teller.exe
			Windows files	C:\Windows
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe
			Internet Explorer 7	C:\Program Files\Internet Explorer\
			Windows files	C:\Windows

Note

AppLocker only supports publisher rules for Windows Store apps. Therefore, collecting the installation path information for Windows Store apps is not necessary.

Event processing

As you create your list of applications, you need to consider how to manage the events that are generated by user access, or you need to deny running those applications to make your users as productive as possible. The following list is an example of what to consider and what to record:

- Will event forwarding be implemented for AppLocker events?
- What is the location of the AppLocker event collection?
- Should an event archival policy be implemented?
- Will the events be analyzed and how often?
- Should a security policy be in place for event collection?

Policy maintenance

As you create your list of applications, you need to consider how to manage and maintain the policies that you will eventually create. The following list is an example of what to consider and what to record:

- How will rules be updated for emergency application access and permanent access?
- How will applications be removed?
- How many older versions of the same application will be maintained?
- How will new applications be introduced?

Next steps

After you have created the list of applications, the next step is to identify the rule collections, which will become the application control policies. This information can be added to the table under the following columns:

- Use default rule or define new rule condition
- Allow or deny
- GPO name

To identify the rule collections, see the following topics:

- Select Types of Rules to Create
- Determine Group Policy Structure and Rule Enforcement

Select Types of Rules to Create

This topic lists resources you can use when selecting your application control policy rules by using AppLocker.

When determining what types of rules to create for each of your groups, you should also determine what enforcement setting to use for each group. Different rule types are more applicable for some applications, depending on the way that the applications are deployed in a specific business group.

The following topics provide additional information about AppLocker rules that can help you decide what rules to use for your applications:

- Understanding AppLocker Rule Behavior
- Understanding AppLocker Rule Exceptions
- Understanding AppLocker Rule Collections
- Understanding AppLocker Allow and Deny Actions on Rules
- Understanding AppLocker Rule Condition TypesUnderstanding AppLocker Rule Condition Types
- Understanding AppLocker Default Rules

Select the rule collection

The rules you create will be in one of the following rule collections:

- Executable files: .exe and .com
- Windows Installer files: .msi, .msp, and .mst
- Scripts: .ps1, .bat, .cmd, .vbs, and .js
- Packaged apps and packaged app installers: .appx
- DLLs: .dll and .ocx

Note

.appx and .mst file types are not applicable to AppLocker running on Windows Server 2008 R2 and Windows 7.

By default, the rules will allow a file to run based upon user or group privilege. If you use DLL rules, a DLL allow rule has to be created for each DLL that is used by all of the allowed applications. The DLL rule collection is not enabled by default.

In the Woodgrove Bank example, the line-of-business application for the Bank Tellers business group is C:\Program Files\Woodgrove\Teller.exe, and this application needs to be included in a rule. In addition, because this rule is part of a list of allowed applications, all the Windows files under C:\Windows must be included as well.

Determine the rule condition

A rule condition is criteria upon which an AppLocker rule is based and can only be one of the rule conditions in the following table.

Rule condition	Usage scenario	Resources
Publisher	To use a publisher condition,	For more information about this

Rule condition	Usage scenario	Resources
	the files must be digitally signed by the software publisher, or you must do so by using an internal certificate. Rules that are specified to the version level might have to be updated when a new version of the file is released.	rule condition, see Understanding the Publisher Rule Condition in AppLocker.
Path	Any file can be assigned this rule condition; however, because path rules specify locations within the file system, any subdirectory will also be affected by the rule (unless explicitly exempted).	For more information about this rule condition, see Understanding the Path Rule Condition in AppLocker.
File hash	Any file can be assigned this rule condition; however, the rule must be updated each time a new version of the file is released because the hash value is based in part upon the version.	For more information about this rule condition, see Understanding the File Hash Rule Condition in AppLocker.

In the Woodgrove Bank example, the line-of-business application for the Bank Tellers business group is signed and is located at C:\Program Files\Woodgrove\Teller.exe. Therefore, the rule can be defined with a publisher condition. If the rule is defined to a specific version and above (for example, Teller.exe version 8.0 and above), then this will allow any updates to this application to occur without interruption of access to the users if the application's name and signed attributes stay the same.

Determine how to allow system files to run

Because AppLocker rules build a list of allowed applications, a rule or rules must be created to allow all Windows files to run. AppLocker provides a means to ensure system files are properly considered in your rule collection by generating the default rules for each rule collection. You can use the default rules as a template when creating your own rules. However, these rules are only meant to function as a starter policy when you are first testing AppLocker rules so that the system files in the Windows folders will be allowed to run. When a default rule is created, it is denoted with "(Default rule)" in its name as it appears in the rule collection.

You can also create a rule for the system files based on the path condition. In the preceding example, for the Bank Tellers group, all Windows files reside under C:\Windows and can be defined with the path rule condition type. This will permit access to these files whenever updates are applied and the files change. If you require additional application security, you might need to modify the rules created from the built-in default rule collection. For example, the default rule to allow all users to run .exe files in the Windows folder is based on a path condition that allows all files within the Windows folder to run. The Windows folder contains a Temp subfolder to which the Users group is given the following permissions:

- Traverse Folder/Execute File
- Create Files/Write Data
- Create Folders/Append Data

These permissions settings are applied to this folder for application compatibility. However, because any user can create files in this location, allowing applications to be run from this location might conflict with your organization's security policy.

Next steps

After you have selected the types of rules to create, record your findings as explained in Document Your AppLocker Rules.

After recording your findings for the AppLocker rules to create, you will need to consider how to enforce the rules. For information about how to do this, see Determine Group Policy Structure and Rule Enforcement.

Document Your AppLocker Rules

This topic describes what rule conditions to associate with each file, how to associate the rule conditions with each file, the source of the rule, and whether the file should be included or excluded.

Record your findings

To complete this AppLocker planning document, you should first complete the following steps:

1. Determine Your Application Control Objectives
2. Create List of Applications Deployed to Each Business Group
3. Select Types of Rules to Create

Document the following items for each business group or organizational unit:

- Whether your organization will use the built-in default AppLocker rules to allow system files to run.
- The types of rule conditions that you will use to create rules, stated in order of preference.

The following table details sample data for documenting rule type and rule condition findings. In addition, you should now consider whether to allow an application to run or deny permission for it to run. For information about these settings, see Understanding AppLocker Allow and Deny Actions on Rules.

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny
Bank Tellers	Teller-East and Teller-West	Yes	Teller Software	C:\Program Files\Woodgrove\Teller.exe	File is signed; create a publisher condition	
			Windows files	C:\Windows	Create a path exception to the default rule to exclude \Windows\Temp	
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe	File is signed; create a publisher condition	
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe	File is not signed; create a file hash condition	
			Internet Explorer 7	C:\Program Files\Internet Explorer\	File is signed; create a publisher condition	
			Windows files	C:\Windows	Use the default rule for the	

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny
					Windows path	

Next steps

For each rule, determine whether to use the allow or deny option. Then, three tasks remain:

- Determine Group Policy Structure and Rule Enforcement
- Plan for AppLocker Policy Management
- Create Your AppLocker Planning Document

Determine Group Policy Structure and Rule Enforcement

This overview topic describes the process to follow when you are planning to deploy AppLocker rules.

You should review the following topics to learn how to structure AppLocker rules for the targeted business groups in your organization:

- Understand AppLocker Enforcement Settings

  This topic describes the AppLocker enforcement settings for rule collections.
- Understand AppLocker Rules and Enforcement Setting Inheritance in Group Policy

  This topic describes what you need to investigate, determine, and record in your application control policies plan.

When you are determining how many Group Policy Objects (GPOs) to create when you apply an AppLocker policy in your organization, you should consider the following:

- Whether you are creating new GPOs or using existing GPOs
- Whether you are implementing Software Restriction Policies (SRP) policies and AppLocker policies in the same GPO
- GPO naming conventions
- GPO size limits

Note

There is no default limit on the number of AppLocker rules that you can create. However, in Windows Server 2008 R2, GPOs have a 2 MB size limit for performance. In subsequent versions, that limit is raised to 100 MB.

After you have determined your Group Policy structure and rule enforcement, record your findings as explained in Document Group Policy Structure and AppLocker Rule Enforcement.

Understand AppLocker Enforcement Settings

This topic describes the AppLocker enforcement settings for rule collections.

Rule enforcement is applied only to a collection of rules, not to individual rules. AppLocker divides the rules into four collections: executable files, Windows Installer files, scripts, and DLL files. For more information about rule collections, see Understanding AppLocker Rule Collections. By default, if enforcement is not configured and rules are present in a rule collection, those rules are enforced. The following table details the three AppLocker rule enforcement settings in Group Policy for each rule collection.

Enforcement setting	Description
Not configured	By default, enforcement is not configured in a rule collection. If rules are present in the corresponding rule collection, they are enforced. If rule enforcement is configured in a higher-level linked Group Policy object (GPO), that enforcement value overrides the **Not configured** value.
Enforce rules	Rules are enforced for the rule collection, and all rule events are audited.
Audit only	Rule events are audited only. Use this value when planning and testing AppLocker rules.

For the AppLocker policy to be enforced on a computer, the Application Identity service must be running. For more information about the Application Identity service, see Configure the Application Identity Service.

When AppLocker policies from various GPOs are merged, the enforcement modes are merged by using the standard Group Policy order of inheritance, which is local, domain, site, and organizational unit (OU). The Group Policy setting that was last written or applied by order of inheritance is used for the enforcement mode, and all rules from linked GPOs are applied.

Understand AppLocker Rules and Enforcement Setting Inheritance in Group Policy

This topic for the IT professional describes how application control policies configured in AppLocker are applied through Group Policy.

Rule enforcement is applied only to collections of rules, not individual rules. AppLocker divides the rules into the following collections: executable files, Windows Installer files, scripts, Packaged apps and Packaged app installers, and DLL files. The options for rule enforcement are **Not configured**, **Enforce rules**, or **Audit only**. Together, all AppLocker rule collections compose the application control policy, or AppLocker policy.

Group Policy merges AppLocker policy in two ways:

- **Rules.** Group Policy does not overwrite or replace rules that are already present in a linked Group Policy Object (GPO). For example, if the current GPO has 12 rules and a linked GPO has 50 rules, 62 rules are applied to all computers that receive the AppLocker policy.

  Important
  When determining whether a file is permitted to run, AppLocker processes rules in the following order:

  a. **Explicit deny.** An administrator created a rule to deny a file.
  b. **Explicit allow.** An administrator created a rule to allow a file.
  c. **Implicit deny.** This is also called the default deny because all files that are not affected by an allow rule are automatically blocked.

- **Enforcement settings.** The last write to the policy is applied. For example, if a higher-level GPO has the enforcement setting configured to **Enforce rules** and the closest GPO has the setting configured to **Audit only**, **Audit only** is enforced. If enforcement is not configured on the closest GPO, the setting from the closest linked GPO will be enforced.

Because a computer's effective policy includes rules from each linked GPO, duplicate rules or conflicting rules could be enforced on a user's computer. Therefore, you should carefully plan your deployment to ensure that only rules that are necessary are present in a GPO.

The following figure demonstrates how AppLocker rule enforcement is applied through linked GPOs.

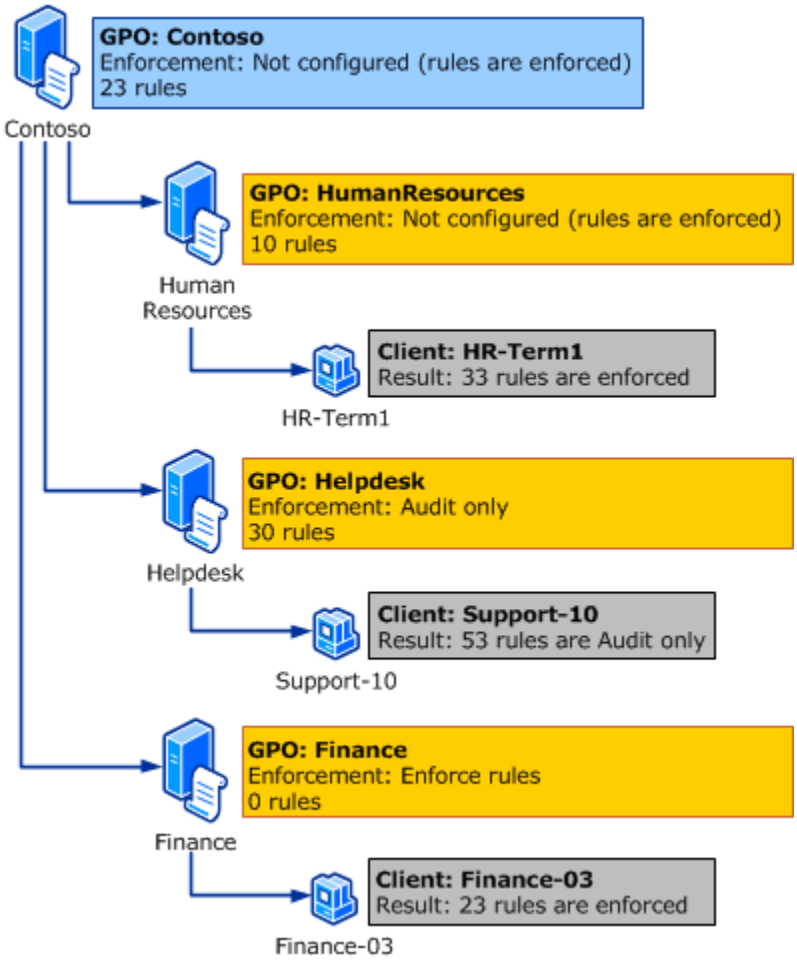

In the preceding illustration, note that all GPOs linked to Contoso are applied in order as configured. The rules that are not configured are also applied. For example, the result of the Contoso and Human Resources GPOs is 33 rules enforced, as shown in the client HR-Term1. The Human Resources GPO contains 10 non-configured rules. When the rule collection is configured for **Audit only**, no rules are enforced.

When constructing the Group Policy architecture for applying AppLocker policies, it is important to remember:

- Rule collections that are not configured will be enforced.
- Group Policy does not overwrite or replace rules that are already present in a linked GPO.
- AppLocker processes the explicit deny rule configuration before the allow rule configuration.
- For rule enforcement, the last write to the GPO is applied.

Document Group Policy Structure and AppLocker Rule Enforcement

This planning topic describes what you need to investigate, determine, and record in your application control policies plan when you use AppLocker.

Record your findings

To complete this AppLocker planning document, you should first complete the following steps:

1. Determine Your Application Control Objectives
2. Create List of Applications Deployed to Each Business Group
3. Select Types of Rules to Create
4. Determine Group Policy Structure and Rule Enforcement

After you determine how to structure your Group Policy Objects (GPOs) so that you can apply AppLocker policies, you should record your findings. You can use the following table to determine how many GPOs to create (or edit) and which objects they are linked to. If you decided to create custom rules to allow system files to run, note the high-level rule configuration in the **Use default rule or define new rule condition** column.

The following table includes the sample data that was collected when you determined your enforcement settings and the GPO structure for your AppLocker policies.

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name
Bank Tellers	Teller-East and Teller-West	Yes	Teller Software	C:\Program Files\Woodgrove\Teller.exe	File is signed; create a publisher condition	Allow	Tellers-AppLockerTellerRules
			Windows files	C:\Windows	Create a path exception to the default rule to exclude \Windows\	Allow	

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name
					Temp		
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe	File is signed; create a publisher condition	Allow	HR-AppLockerHR Rules
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe	File is not signed; create a file hash condition	Allow	
			Internet Explorer 7	C:\Program Files\Internet Explorer\	File is signed; create a publisher condition	Deny	
			Windows files	C:\Windows	Use a default rule for the Windows path	Allow	

Note
Excluding Windows Server 2008 R2 and Windows 7, AppLocker can manage Windows Store apps. For information about how to add rules for these apps to your existing GPO, see Add Rules for Packaged Apps to Existing AppLocker Rule-set.

Next steps

After you have determined the Group Policy structure and rule enforcement strategy for each business group's applications, the following tasks remain:

- Plan for AppLocker Policy Management

- Create Your AppLocker Planning Document

Plan for AppLocker Policy Management

This topic for describes the decisions you need to make to establish the processes for managing and maintaining AppLocker policies.

Policy management

Before you begin the deployment process, consider how the AppLocker rules will be managed. Developing a process for managing AppLocker rules helps assure that AppLocker continues to effectively control how applications are allowed to run in your organization.

Application and user support policy

Developing a process for managing AppLocker rules helps assure that AppLocker continues to effectively control how applications are allowed to run in your organization. Considerations include:

- What type of end-user support is provided for blocked applications?
- How are new rules added to the policy?
- How are existing rules updated?
- Are events forwarded for review?

Help desk support

If your organization has an established help desk support department in place, consider the following when deploying AppLocker policies:

- What documentation does your support department require for new policy deployments?
- What are the critical processes in each business group both in work flow and timing that will be affected by application control policies and how could they affect your support department's workload?
- Who are the contacts in the support department?
- How will the support department resolve application control issues between the end user and those who maintain the AppLocker rules?

End-user support

Because AppLocker is preventing unapproved applications from running, it is important that your organization carefully plan how to provide end-user support. Considerations include:

- Do you want to use an intranet site as a first line of support for users who have tried to run a blocked application?
- How do you want to support exceptions to the policy? Will you allow users to run a script to temporarily allow access to a blocked application?

Using an intranet site

AppLocker can be configured to display the default message but with a custom URL. You can use this URL to redirect users to a support site that contains information about why the user received the error and which applications are allowed. If you do not display a custom URL for the message when an application is blocked, the default URL is used.

The following image shows an example of the error message for a blocked application. You can use the **Set a support web link** policy setting to customize the **More information** link.

Blocked application error message

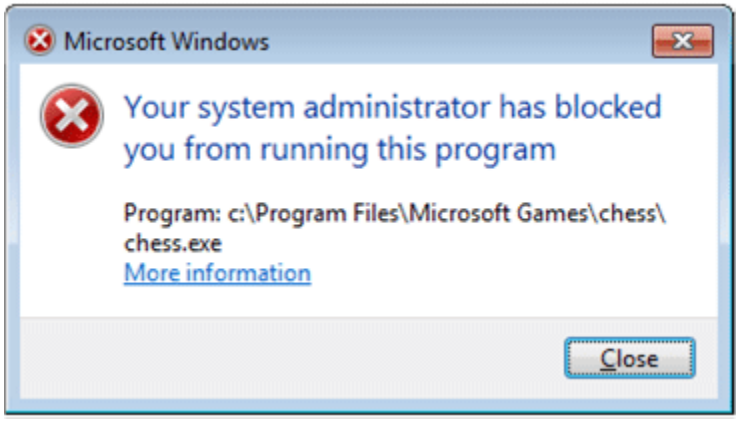

For steps to display a custom URL for the message, see Display a Custom URL Message When Users Try to Run a Blocked Application (http://go.microsoft.com/fwlink/?LinkId=160265).

AppLocker event management

Each time that a process requests permission to run, AppLocker creates an event in the AppLocker event log. The event details which file tried to run, the attributes of that file, the user that initiated the request, and the rule GUID that was used to make the AppLocker execution decision. The AppLocker event log is located in the following path: **Applications and Services Logs\Microsoft\Windows\AppLocker**. The AppLocker log includes three logs:

1. **EXE and DLL**. Contains events for all files affected by the executable and DLL rule collections (.exe, .com, .dll, and .ocx).
2. **MSI and Script**. Contains events for all files affected by the Windows Installer and script rule collections (.msi, .msp, .ps1, .bat, .cmd, .vbs, and .js).
3. **Packaged app-Deployment** or **Packaged app-Execution**, contains events for all Windows 8 apps affected by the Packaged app and Packed app installer rule collection (.appx).

Collecting these events in a central location can help you maintain your AppLocker policy and troubleshoot rule configuration problems. Event collection technologies such as those available in Windows allow administrators to subscribe to specific event channels and have the events from source computers aggregated into a forwarded event log on a Windows Server operating system collector. For more information about setting up an event subscription, see Configure Computers to Collect and Forward Events (http://go.microsoft.com/fwlink/?LinkId=145012).

Policy maintenance

As new applications are deployed or existing applications are updated by the software publisher, you will need to make revisions to your rule collections to ensure that the policy is current.

You can edit an AppLocker policy by adding, changing, or removing rules. However, you cannot specify a version for the policy by importing additional rules. To ensure version control when modifying an AppLocker policy, use Group Policy management software that allows you to create versions of Group Policy Objects (GPOs). An example of this type of software is the Advanced Group Policy Management feature from the Microsoft Desktop Optimization Pack. For more information about Advanced Group Policy Management, see Advanced Group Policy Management Overview (http://go.microsoft.com/fwlink/?LinkId=145013).

Caution

You should not edit an AppLocker rule collection while it is being enforced in Group Policy. Because AppLocker controls what files are allowed to run, making changes to a live policy can create unexpected behavior.

New version of a supported application

When a new version of an application is deployed in the organization, you need to determine whether to continue to support the previous version of that application. To add the new version, you might only need to create a new rule for each file that is associated with the application. If you are using publisher conditions and the version is not specified, then the existing rule or rules might be sufficient to allow the updated file to run. You must ensure, however, that the updated application has not altered the file names or added files to support new functionality. If so, then you must modify the existing rules or create new rules. To continue to reuse a publisher-based rule without a specific file version, you must also ensure that the file's digital signature is still identical to the previous version—the publisher, product name, and file name (if configured in your rule) must all match for the rule to be correctly applied.

To determine whether a file has been modified during an application update, review the publisher's release details provided with the update package. You can also review the publisher's Web page to retrieve this information. Each file can also be inspected to determine the version.

For files that are allowed or denied with file hash conditions, you must retrieve the new file hash. To add support for a new version and maintain support for the older version, you can either create a new file hash rule for the new version or edit the existing rule and add the new file hash to the list of conditions.

For files with path conditions, you should verify that the installation path has not changed from what is stated in the rule. If the path has changed, you need to update the rule before installing the new version of the application.

Recently deployed application

To support a new application, you must add one or more rules to the existing AppLocker policy.

Application is no longer supported

If your organization has determined that it will no longer support an application that has AppLocker rules associated with it, the easiest way to prevent users from running the application is to delete these rules.

Application is blocked but should be allowed

A file could be blocked for three reasons:

- The most common reason is that no rule exists to allow the application to run.
- There may be an existing rule that was created for the file that is too restrictive.
- A deny rule, which cannot be overridden, is explicitly blocking the file.

Before editing the rule collection, first determine what rule is preventing the file from running. You can troubleshoot the problem by using the **Test-AppLockerPolicy** Windows PowerShell cmdlet. For more information about troubleshooting an AppLockerpolicy, see Testing and Updating an AppLocker Policy (http://go.microsoft.com/fwlink/?LinkId=160269).

Next steps

After deciding how your organization will manage your AppLocker policy, record your findings.

- **End-user support policy.** Document the process that you will use for handling calls from users who have attempted to run a blocked application, and ensure that support personnel have clear escalation steps so that the administrator can update the AppLocker policy, if necessary.
- **Event processing.** Document whether events will be collected in a central location called a store, how that store will be archived, and whether the events will be processed for analysis.
- **Policy maintenance.** Detail how rules will be added to the policy and in which GPO the rules are defined.

For information and steps how to document your processes, see Document Your Application Control Management Processes.

Document Your Application Control Management Processes

This planning topic describes the AppLocker policy maintenance information to record for your design document.

Record your findings

To complete this AppLocker planning document, you should first complete the following steps:

1. Determine Your Application Control Objectives
2. Create List of Applications Deployed to Each Business Group
3. Select Types of Rules to Create
4. Determine Group Policy Structure and Rule Enforcement

5. Plan for AppLocker Policy Management

The three key areas to determine for AppLocker policy management are:

1. Support policy

   Document the process that you will use for handling calls from users who have attempted to run a blocked application, and ensure that support personnel know recommended troubleshooting steps and escalation points for your policy.

2. Event processing

   Document whether events will be collected in a central location, how that store will be archived, and whether the events will be processed for analysis.

3. Policy maintenance

   Detail how rules will be added to the policy, in which Group Policy Object (GPO) the rules should be defined, and how to modify rules when applications are retired, updated, or added.

The following table contains the added sample data that was collected when determining how to maintain and manage AppLocker policies.

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy
Bank Tellers	Teller-East and Teller-West	Yes	Teller Software	C:\Program Files\Woodgrove\Teller.exe	File is signed; create a publisher condition	Allow	Tellers-AppLockerTellerRules	Web help
			Windows files	C:\Windows	Create a path exception to the default rule to exclude \Windows\Temp	Allow		Help desk
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe	File is signed; create a publisher condition	Allow	HR-AppLockerHRRules	Web help

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe	File is not signed; create a file hash condition	Allow		Web help
			Internet Explorer 7	C:\Program Files\Internet Explorer\	File is signed; create a publisher condition	Deny		Web help
			Windows files	C:\Windows	Use the default rule for the Windows path	Allow		Help desk

The following two tables illustrate examples of documenting considerations to maintain and manage AppLocker policies.

Event processing policy

One discovery method for application usage is to set the AppLocker enforcement mode to **Audit only**. This will write events to the AppLocker logs, which can be managed and analyzed like other Windows logs. After applications have been identified, you can begin to develop policies regarding the processing and access to AppLocker events.

The following table is an example of what to consider and record.

Business group	AppLocker event collection location	Archival policy	Analyzed?	Security policy
Bank Tellers	Forwarded to: AppLocker Event Repository on srvBT093	Standard	None	Standard

Business group	AppLocker event collection location	Archival policy	Analyzed?	Security policy
Human Resources	DO NOT FORWARD. srvHR004	60 months	Yes, summary reports monthly to managers	Standard

Policy maintenance policy

When applications are identified and policies are created for application control, then you can begin documenting how you intend to update those policies.

The following table is an example of what to consider and record.

Business group	Rule update policy	Application decommission policy	Application version policy	Application deployment policy
Bank Tellers	Planned: Monthly through business office triage Emergency: Request through help desk	Through business office triage 30-day notice required	General policy: Keep past versions for 12 months List policies for each application	Coordinated through business office 30-day notice required
Human Resources	Planned: Monthly through HR triage Emergency: Request through help desk	Through HR triage 30-day notice required	General policy: Keep past versions for 60 months List policies for each application	Coordinated through HR 30-day notice required

Next steps

After you have determined your application control management strategy for each of the business group's applications, the following task remains:

- Create Your AppLocker Planning Document

Create Your AppLocker Planning Document

This planning topic for the IT professional summarizes the information you need to research and include in your AppLocker planning document.

The AppLocker deployment design

The design process and the planning document help you investigate application usage in your organization and record your findings so you can effectively deploy and maintain application control policies by using AppLocker.

You should have completed these steps in the design and planning process:

1. Determine Your Application Control Objectives
2. Create List of Applications Deployed to Each Business Group
3. Select Types of Rules to Create
4. Determine Group Policy Structure and Rule Enforcement
5. Plan for AppLocker Policy Management

AppLocker planning document contents

Your planning document should contain:

- A list of business groups that will participate in the application control policy project, their requirements, a description of their business processes, and contact information.
- Application control policy project target dates, both for planning and deployment.
- A complete list of applications used by each business group (or organizational unit), including version information and installation paths.
- What condition to apply to rules governing each application (or whether to use the default set provided by AppLocker).
- A strategy for using Group Policy to deploy the AppLocker policies.
- A strategy in processing the application usage events generated by AppLocker.
- A strategy to maintain and manage AppLocker polices after deployment.

Sample template for an AppLocker planning document

You can use the following form to construct your own AppLocker planning document.

Business group:

Operating system environment: (Windows and non-Windows)

Contacts	Business contact:	Technical contact:
Other departments	In this business group:	Affected by this project:
Security policies	Internal:	Regulatory/compliance:
Business goals	Primary:	Secondary:
Project target dates	Design signoff date:	Policy deployment date:

Rules

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy

Event processing

Business group	AppLocker event collection location	Archival policy	Analyzed?	Security policy

Policy maintenance

Business group	Rule update policy	Application decommission policy	Application version policy	Application deployment policy
	Planned: Emergency:			

Example of an AppLocker planning document

Rules

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy
Bank Teller	Teller-East	Yes	Teller Softwa	C:\Program Files\Woodgrove\Tell	File is signed;	All	Tellers-AppLockerT	Web

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy
s	and Teller-West		re	er.exe	create a publisher condition	ow	ellerRules	help
			Windows files	C:\Windows	Create a path exception to the default rule to exclude \Windows\Temp	Allow		Helpdesk
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe	File is signed; create a publisher condition	Allow	HR-AppLockerHRRules	Webhelp
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe	File is not signed; create a file hash condition	Allow		Webhelp
			Internet Explorer 7	C:\Program Files\Internet Explorer\	File is signed; create a publisher condition	Deny		Webhelp
			Windows files	C:\Windows	Use the default rule for the Windows path	Allow		Helpdesk

Event processing

Business group	AppLocker event collection location	Archival policy	Analyzed?	Security policy
Bank Tellers	Forwarded to: AppLocker Event Repository on srvBT093	Standard	None	Standard
Human Resources	DO NOT FORWARD. srvHR004	60 months	Yes, summary reports monthly to managers	Standard

Policy maintenance

Business group	Rule update policy	Application decommission policy	Application version policy	Application deployment policy
Bank Tellers	Planned: Monthly through business office triage Emergency: Request through help desk	Through business office triage 30-day notice required	General policy: Keep past versions for 12 months List policies for each application	Coordinated through business office 30-day notice required
Human Resources	Planned: Monthly through HR triage Emergency: Request through help desk	Through HR triage 30-day notice required	General policy: Keep past versions for 60 months List policies for each application	Coordinated through HR 30-day notice required

Additional resources

- The AppLocker Policies Design Guide is the predecessor to the AppLocker Policies Deployment Guide. When planning is complete, see the AppLocker Policies Deployment Guide.
- For a complete list of AppLocker documentation and other resources, see **Windows AppLocker Overview** in the Windows Server Technical Library.

AppLocker Policies Deployment Guide

This topic for the IT professional introduces the concepts and describes the steps required to deploy AppLocker™ policies introduced in Windows Server 2008 R2 and Windows 7.

Purpose of this guide

This guide provides steps based on your design and planning investigation for deploying application control policies by using AppLocker. It is intended for security architects, security administrators, and system administrators. Through a sequential and iterative deployment process, you can create application control policies, test and adjust the policies, and implement a method for maintaining those policies as the needs in your organization change.

This guide covers the use of Software Restriction Policies (SRP) in conjunction with AppLocker policies to control application usage. For a comparison of SRP and AppLocker, see Using Software Restriction Policies and AppLocker Policies in this guide. To understand if AppLocker is the correct application control solution for you, see Understand AppLocker Policy Design Decisions.

Tip
To digitally save or print pages from this library, click **Export** (in the upper-right corner of the page), and then follow the instructions.

Prerequisites to deploying AppLocker policies

The following are prerequisites or recommendations to deploying policies:

- Understand the capabilities of AppLocker:
  - **AppLocker Technical Overview**
  - AppLocker Step-by-Step Guide
- Document your application control policy deployment plan by addressing these tasks:
  - Understand the AppLocker Policy Deployment Process
  - Understand AppLocker Policy Design Decisions
  - Determine Your Application Control Objectives
  - Create List of Applications Deployed to Each Business Group
  - Select Types of Rules to Create
  - Determine Group Policy Structure and Rule Enforcement
  - Plan for AppLocker Policy Management
  - Create Your AppLocker Planning Document

Contents of this guide

This guide provides steps based on your design and planning investigation for deploying application control policies created and maintained by AppLocker for computers running any of the supported versions of Windows listed in Requirements to Use AppLocker. It contains the following topics:

- Understand the AppLocker Policy Deployment Process
- Requirements for Deploying AppLocker Policies
- Using Software Restriction Policies and AppLocker Policies
- Create Your AppLocker Policies
- Deploy the AppLocker Policy into Production
- Maintain AppLocker Policies

Additional resources

Using Software Restriction Policies to Protect Against Unauthorized Software

(http://go.microsoft.com/fwlink/?LinkID=155634)

This TechNet article is about SRP in Windows XP and Windows Server 2003 and is also applicable to Windows Vista® and Windows Server 2008. It provides an in-depth look at how software restriction policies can be used to fight viruses, regulate which ActiveX controls can be downloaded, run only digitally signed scripts, and enforce that only approved software is installed on system computers.

Software Restriction Policies

Windows Server 2003 product help **Software Restriction Policies**. This collection of topics describes the concepts to understand and the steps to implement and maintain SRP.

More recent guidance to Administer Software Restriction Policies. This collection contains procedures how to administer application control policies using Software Restriction Policies (SRP).

AppLocker Overview

This topic lists AppLocker documentation resources for the IT professional.

Understand the AppLocker Policy Deployment Process

This planning and deployment topic for the IT professional describes the process for using AppLocker when deploying application control policies.

To successfully deploy AppLocker policies, you need to identify your application control objectives and construct the policies for those objectives. The key to the process is taking an accurate inventory of your organization's applications, which requires investigation of all the

targeted business groups. With an accurate inventory, you can create rules and set enforcement criteria that will allow the organization to use the required applications and allow the IT department to manage a controlled set of applications.

The following diagram shows the main points in the design, planning, and deployment process for AppLocker.

Planning an AppLocker Deployment

1 Determine how to implement AppLocker
- Determine whether interoperability with Software Restriction Policies (SRP) is needed.
- Document where to use AppLocker.
- Select whether to use allow actions only or allow and deny actions.

2 Create a list of applications
- Create an inventory of applications installed in different OUs.

3 Select the types of rules to create
- Select which rule collections to use.
- Determine how to allow system files to run.

4 Define the Group Policy structure and rule enforcement
- Select enforcement settings for each OU.
- Determine rule and enforcement setting inheritance in Group Policy.

5 Create a process for managing rules
- Document an end-user support policy for blocked applications.
- Determine whether to use event forwarding.
- Determine how to maintain the AppLocker policy.

6 Document your AppLocker plan
- Document your AppLocker plan to use when deploying AppLocker and for future reference.

Deploying AppLocker

1 Author rules in a test environment
Using your documented AppLocker plan:
- Create a reference computer.
- Automatically generate rules.
- Modify the rule set.

2 Test and update the policy
- Enable **Audit only** enforcement.
- Start the Application Identity service.
- Test the policy.
- Analyze events.
- Modify the policy.
- Repeat the testing process until you are satisfied with the results.

3 Deploy the production enforcement setting
- Deploy the AppLocker enforcement settings to be used in production.

4 Maintain the policy
- Add new rules for new applications.
- Edit existing rules for revisions to organization policy.
- Provide end-user support.

Resources to support the deployment process

The following documentation contains information about designing, planning, deploying, and maintaining AppLocker policies:

- For information about the AppLocker policy design and planning requirements and process, see AppLocker Policies Design Guide.
- For information about the AppLocker policy deployment requirements and process, see AppLocker Policies Deployment Guide.
- For information about AppLocker policy maintenance and monitoring for Windows Server 2008 R2 and Windows 7, see AppLocker Operations Guide.

  For information about AppLocker policy maintenance and monitoring in versions of Windows later than Windows Server 2008 R2 and Windows 7, see Administer AppLocker.
- For information about AppLocker policy architecture, components, and processing, see AppLocker Technical Reference.

Requirements for Deploying AppLocker Policies

This deployment topic for the IT professional lists the requirements that you need to consider before you deploy AppLocker policies.

The following requirements must be met or addressed before you deploy your AppLocker policies:

- Deployment plan
- Supported operating systems
- Policy distribution mechanism
- Event collection and analysis system

Deployment plan

An AppLocker policy deployment plan is the result of investigating which applications are required and necessary in your organization, which applications are optional, and which applications are forbidden. To develop this plan, see AppLocker Policies Design Guide. The following table is an example of the data you need to collect and the decisions you need to make to successfully deploy AppLocker policies on the supported operating systems (as listed in Requirements to Use AppLocker).

Business group	Organizational unit	Implement AppLocker?	Applications	Installation path	Use default rule or define new rule condition	Allow or deny	GPO name	Support policy
Bank Tellers	Teller-East and Teller-West	Yes	Teller software	C:\Program Files\Woodgrove\Teller.exe	File is signed; create a publisher condition	Allow	Tellers	Web help
			Windows files	C:\Windows	Create a path exception to the default rule to exclude \Windows\Temp	Allow		Help Desk
			Time Sheet Organizer	C:\Program Files\Woodgrove\HR\Timesheet.exe	File is not signed; create a file hash condition	Allow		Web help
Human Resources	HR-All	Yes	Check Payout	C:\Program Files\Woodgrove\HR\Checkcut.exe	File is signed; create a publisher condition	Allow	HR	Web help
			Internet Explorer 7	C:\Program Files\Internet Explorer\	File is signed; create a publisher condition	Deny		Help Desk
			Windows files	C:\Windows	Use the default rule for the Windows path	Allow		Help Desk

Event processing policy

Business group	AppLocker event collection location	Archival policy	Analyzed?	Security policy
Bank Tellers	Forwarded to: srvBT093	Standard	None	Standard
Human Resources	Do not forward	60 months	Yes; summary reports monthly to managers	Standard

Policy maintenance policy

Business group	Rule update policy	Application decommission policy	Application version policy	Application deployment policy
Bank Tellers	Planned: Monthly through business office triage Emergency: Request through Help Desk	Through business office triage; 30-day notice required	General policy: Keep past versions for 12 months List policies for each application	Coordinated through business office; 30-day notice required
Human Resources	Planned: Through HR triage Emergency: Request through Help Desk	Through HR triage; 30-day notice required	General policy: Keep past versions for 60 months List policies for each application	Coordinated through HR; 30-day notice required

Supported operating systems

AppLocker is supported only on certain operating systems. Some features are not available on all operating systems. For more information, see Requirements to Use AppLocker.

Software Restriction Policies are supported in Windows XP and Windows Server 2003, in addition to the versions listed in Requirements to Use AppLocker.

Policy distribution mechanism

You need a way to distribute the AppLocker policies throughout the targeted business groups. AppLocker uses Group Policy management architecture to effectively distribute application control policies. AppLocker policies can also be configured on individual computers by using the Local Security Policy snap-in.

Event collection and analysis system

Event processing is important to understand application usage. You must have a process in place to collect and analyze AppLocker events so that application usage is appropriately restricted and understood. For procedures to monitor AppLocker events, see:

- Configure an AppLocker Policy for Audit Only
- Configure an AppLocker Policy for Enforce Rules
- Monitor Application Usage with AppLocker

See also

AppLocker Policies Deployment Guide

Using Software Restriction Policies and AppLocker Policies

This topic for the IT professional describes how to use Software Restriction Policies (SRP) and AppLocker policies in the same Windows deployment.

Understand the difference between SRP and AppLocker

You might want to deploy application control policies in Windows operating systems earlier than Windows Server 2008 R2 or Windows 7. You can use AppLocker policies only on the supported versions and editions of Windows as listed in Requirements to Use AppLocker. However, you can use SRP on those supported editions of Windows plus Windows Server 2003 and Windows XP. To compare features and functions in SRP and AppLocker so that you can determine when to use each technology to meet your application control objectives, see Determine Your Application Control Objectives.

Use SRP and AppLocker in the same domain

SRP and AppLocker use Group Policy for domain management. However, when policies are generated by SRP and AppLocker exist in the same domain, and they are applied through Group Policy, AppLocker policies take precedence over policies generated by SRP on computers that

are running an operating system that supports AppLocker. For information about how inheritance in Group Policy applies to AppLocker policies and policies generated by SRP, see Understand AppLocker Rules and Enforcement Setting Inheritance in Group Policy.

Important

As a best practice, use separate Group Policy Objects to implement your SRP and AppLocker policies. To reduce troubleshooting issues, do not combine them in the same GPO.

The following scenario provides an example of how each type of policy would affect a bank teller software application, where the application is deployed on different Windows desktop operating systems and managed by the Tellers GPO.

Operating system	Tellers GPO with AppLocker policy	Tellers GPO with SRP	Tellers GPO with AppLocker policy and SRP
Windows 8.1, Windows 8,and Windows 7	AppLocker policies in the GPO are applied, and they supersede any local AppLocker policies.	Local AppLocker policies supersede policies generated by SRP that are applied through the GPO.	AppLocker policies in the GPO are applied, and they supersede the policies generated by SRP in the GPO and local AppLocker policies or policies generated by SRP.
Windows Vista	AppLocker policies are not applied.	Policies generated by SRP in the GPO are applied, and they supersede local policies generated by SRP.AppLocker policies are not applied.	Policies generated by SRP in the GPO are applied, and they supersede local policies generated by SRP. AppLocker policies not applied.
Windows XP	AppLocker policies are not applied.	Policies generated by SRP in the GPO are applied, and they supersede local policies generated by SRP. AppLocker policies are not applied.	Policies generated by SRP in the GPO are applied, and they supersede local policies generated by SRP. AppLocker policies not applied.

Note

For information about supported versions and editions of the Windows operating system, see Requirements to Use AppLocker.

Test and validate SRPs and AppLocker policies that are deployed in the same environment

Because SRPs and AppLocker policies function differently, they should not be implemented in the same GPO. This makes testing the result of the policy straightforward, which is critical to successfully controlling application usage in the organization. Configuring a testing and policy distribution system can help you understand the result of a policy. The effects of policies generated by SRP and AppLocker policies need to be tested separately and by using different tools.

Step 1: Test the effect of SRPs

You can use the Group Policy Management Console (GPMC) or the Resultant Set of Policy (RSoP) snap-in to determine the effect of applying SRPs by using GPOs. For information about using the GPMC, see **Group Policy Management Console**. For information about using RSoP, see **Resultant Set of Policy**.

Step 2: Test the effect of AppLocker policies

You can test AppLocker policies by using Windows PowerShell cmdlets. For information about investigating the result of a policy, see:

- Test an AppLocker Policy by Using Test-AppLockerPolicy
- Monitor Application Usage with AppLocker

Another method to use when determining the result of a policy is to set the enforcement mode to **Audit only**. When the policy is deployed, events will be written to the AppLocker logs as if the policy was enforced. For information about using the **Audit only** mode, see:

Understand AppLocker Enforcement Settings

Configure an AppLocker Policy for Audit Only

See also

AppLocker Policies Deployment Guide

Create Your AppLocker Policies

This overview topic for the IT professional describes the steps to create an AppLocker policy and prepare it for deployment.

Creating effective application control policies with AppLocker starts by creating the rules for each application. Rules are grouped into one of five rule collections. The rule collection can be configured to be enforced or to run in **Audit only** mode. An AppLocker policy includes the rules in the five rule collections and the enforcement settings for each rule collection.

Step 1: Use your plan

You can develop an application control policy plan to guide you in making successful deployment decisions. For more information about how to do this and what you should consider, see the AppLocker Policies Design Guide. This guide is intended for security architects, security administrators, and system administrators. It contains the following topics to help you create an AppLocker policy deployment plan for your organization that will address your specific application control requirements by department, organizational unit, or business group:

1. Understand the AppLocker Policy Deployment Process
2. Understand AppLocker Policy Design Decisions
3. Determine Your Application Control Objectives
4. Create List of Applications Deployed to Each Business Group
5. Select Types of Rules to Create
6. Determine Group Policy Structure and Rule Enforcement
7. Plan for AppLocker Policy Management
8. Create Your AppLocker Planning Document

Step 2: Create your rules and rule collections

Each rule applies to one or more applications, and it imposes a specific rule condition on them. Rules can be created individually or they can be generated by the Automatically Generate Rules Wizard. For the steps to create the rules, see Create Your AppLocker Rules.

Step 3: Configure the enforcement setting

An AppLocker policy is a set of rule collections that are configured with a rule enforcement setting. The enforcement setting can be **Enforce rules**, **Audit only**, or **Not configured**. If an AppLocker policy has at least one rule, and it is set to **Not configured**, all the rules in that policy will be enforced. For information about configuring the rule enforcement setting, see Configure an AppLocker Policy for Audit Only and Configure an AppLocker Policy for Enforce Rules.

Step 4: Update the GPO

AppLocker policies can be defined locally on a computer or applied through Group Policy. To use Group Policy to apply AppLocker policies, you must create a new Group Policy Object (GPO) or you must update an existing GPO. You can create or modify AppLocker policies by using the

Group Policy Management Console (GPMC), or you can import an AppLocker policy into a GPO. For the procedure to do this, see Import an AppLocker Policy into a GPO.

Step 5: Test the effect of the policy

In a test environment or with the enforcement setting set at **Audit only**, verify that the results of the policy are what you intended. For information about testing a policy, see Test and Update an AppLocker Policy.

Step 6: Implement the policy

Depending on your deployment method, import the AppLocker policy to the GPO in your production environment, or if the policy is already deployed, change the enforcement setting to your production environment value—**Enforce rules** or **Audit only**.

Step 7: Test the effect of the policy and adjust

Validate the effect of the policy by analyzing the AppLocker logs for application usage, and then modify the policy as necessary. To do this, see Monitor Application Usage with AppLocker.

Next steps

Follow the steps described in the following topics to continue the deployment process:

1. Create Your AppLocker Rules
2. **Testing and Updating an AppLocker Policy**
3. Deploy the AppLocker Policy into Production

See also

AppLocker Policies Deployment Guide

Create Your AppLocker Rules

This topic for the IT professional describes what you need to know about AppLocker rules and the methods that you can to create rules.

Creating AppLocker rules

AppLocker rules apply to the targeted application, and they are the components that make up the AppLocker policy. Depending on your IT environment and the business group that requires application control policies, setting these access rules for each application can be time-

consuming and prone to error. With AppLocker, you can generate rules automatically or create rules individually. Creating rules that are derived from your planning document can help you avoid unintended results. For information about this planning document and other planning activities, see AppLocker Policies Design Guide.

Automatically generate your rules

You can use a reference computer to automatically create a set of default rules for each of the installed applications, test and modify each rule as necessary, and deploy the policies. Creating most of the rules for all the installed applications gives you a starting point to build and test your policies. For information about performing this task, see the following topics:

- Configure the AppLocker Reference Computer
- Run the Automatically Generate Rules Wizard
- Create AppLocker Default Rules
- Edit AppLocker Rules
- Configure Exceptions for an AppLocker Rule

Create your rules individually

You can create rules and set the mode to **Audit only** for each installed application, test and update each rule as necessary, and then deploy the policies. Creating rules individually might be best when you are targeting a small number of applications within a business group.

Note
AppLocker includes default rules for each rule collection. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed in an AppLocker rule collection. You can also edit the default rules. For information about creating the default rules for the Windows operating system, see Create AppLocker Default Rules.

For information about performing this task, see:

1. Create a Rule That Uses a Publisher Condition
2. Create a Rule That Uses a Path Condition
3. Create a Rule That Uses a File Hash Condition
4. Edit AppLocker Rules
5. Enforce AppLocker Rules
6. Configure an AppLocker Policy for Audit Only

About selecting rules

AppLocker policies are composed of distinct rules for specific applications. These rules are grouped by collection, and they are implemented through an AppLocker policy definition. AppLocker policies are managed by using Group Policy or by using the Local Security Policy snap-in for a single computer.

When you determine what types of rules to create for each of your business groups or organizational units (OUs), you should also determine what enforcement setting to use for each group. Certain rule types are more applicable for some applications, depending on how the applications are deployed in a specific business group.

For information about how to determine and document your AppLocker rules, see AppLocker Policies Design Guide.

For information about AppLocker rules and AppLocker policies, see the following topics:

- Understanding AppLocker Rule Behavior
- Understanding AppLocker Rule Exceptions
- Understanding AppLocker Rule Collections
- Understanding AppLocker Allow and Deny Actions on Rules
- Understanding AppLocker Rule Condition Types
- Understanding AppLocker Default Rules

Next steps

1. Import an AppLocker Policy into a GPO
2. Import an AppLocker Policy from Another Computer
3. Test and Update an AppLocker Policy
4. Deploy the AppLocker Policy into Production

See Also

Create Your AppLocker Policies

Test and Update an AppLocker Policy

This topic discusses the steps required to test an AppLocker policy prior to deployment in Windows Server 2012 and Windows 8.

You should test each set of rules to ensure that the rules perform as intended. If you use Group Policy to manage AppLocker policies, complete the following steps for each Group Policy Object (GPO) where you have created AppLocker rules. Because AppLocker rules are inherited from linked GPOs, you should deploy all of the rules for simultaneous testing in all of your test GPOs.

Step 1: Enable the Audit only enforcement setting

By using the **Audit only** enforcement setting, you can ensure that the AppLocker rules that you have created are properly configured for your organization. This setting can be enabled on the **Enforcement** tab of the **AppLocker Properties** dialog box. For the procedure to do this, see Configure an AppLocker Policy for Audit Only.

Step 2: Configure the Application Identity service to start automatically

Because AppLocker uses the Application Identity service to verify the attributes of a file, you must configure it to start automatically in any one GPO that applies AppLocker rules. For the procedure to do this, see Configure the Application Identity Service. For AppLocker policies that are not managed by a GPO, you must ensure that the service is running on each computer in order for the policies to be applied.

Step 3: Test the policy

Test the AppLocker policy to determine if your rule collection needs to be modified. Because you have created AppLocker rules, enabled the Application Identity service, and enabled the **Audit only** enforcement setting, the AppLocker policy should be present on all client computers that are configured to receive your AppLocker policy.

The **Test-AppLockerPolicy** Windows PowerShell cmdlet can be used to determine whether any of the rules in your rule collection will be blocked on your reference computers. For the procedure to do this, see Test an AppLocker Policy by Using Test-AppLockerPolicy.

Step 4: Analyze AppLocker events

You can either manually analyze AppLocker events or use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to automate the analysis.

To manually analyze AppLocker events

You can view the events either in Event Viewer or a text editor and then sort those events to perform an analysis, such as looking for patterns in application usage events, access frequencies, or access by user groups. If you have not configured an event subscription, then you will have to review the logs on a sampling of computers in your organization. For more information about using Event Viewer, see View the AppLocker Log in Event Viewer.

To analyze AppLocker events by using Get-AppLockerFileInformation

You can use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet to analyze AppLocker events from a remote computer. If an application is being blocked and should be allowed, you can use the AppLocker cmdlets to help troubleshoot the problem.

For both event subscriptions and local events, you can use the **Get-AppLockerFileInformation** cmdlet to determine which files have been blocked or would have been blocked (if you are using the **Audit only** enforcement mode) and how many times the event has occurred for each file. For the procedure to do this, see Review AppLocker Events with Get-AppLockerFileInformation.

After using **Get-AppLockerFileInformation** to determine how many times that a file would have been blocked from running, you should review your rule list to determine whether a new rule should be created for the blocked file or whether an existing rule is too strictly defined. Ensure that you check which GPO is currently preventing the file from running. To determine this, you can use the Group Policy Results Wizard to view rule names.

Step 5: Modify the AppLocker policy

After you have identified which rules need to be edited or added to the policy, you can use the Group Policy Management Console to modify the AppLocker rules in the relevant GPOs. For AppLocker policies that are not managed by a GPO, you can use the Local Security Policy snap-in. For information how to modify an AppLocker policy, see, Edit an AppLocker Policy.

Step 6: Repeat policy testing, analysis, and policy modification

Repeat the previous steps 3–5 until all the rules perform as intended before applying enforcement.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

Deploy the AppLocker Policy into Production

This topic for the IT professional describes the tasks that should be completed before you deploy AppLocker application control settings.

After successfully testing and modifying the AppLocker policy for each Group Policy Object (GPO), you are ready to deploy the enforcement settings into production. For most organizations, this means switching the AppLocker enforcement setting from **Audit only** to **Enforce rules**. However, it is important to follow the deployment plan that you created earlier. For more information, see the AppLocker Policies Design Guide. Depending on the needs of different business groups in your organization, you might deploy different enforcement settings for linked GPOs.

Understand your design decisions

Before you deploy an AppLocker policy, you should determine:

- For each business group, which applications will be controlled and in what manner. For more information, see Create List of Applications Deployed to Each Business Group.

- How to handle requests for application access. For information about what to consider when developing your support policies, see Plan for AppLocker Policy Management.
- How to manage events, including forwarding events. For information about event management in AppLocker, see Monitor Application Usage with AppLocker.
- Your GPO structure, including how to include policies generated by Software Restriction Policies and AppLocker policies. For more information, see Determine Group Policy Structure and Rule Enforcement.

For information about how AppLocker deployment is dependent on design decisions, see Understand AppLocker Policy Design Decisions.

AppLocker deployment methods

If you have configured a reference computer, you can create and update your AppLocker policies on this computer, test the policies, and then export the policies to the appropriate GPO for distribution. Another method is to create the policies and set the enforcement setting on **Audit only**, then observe the events that are generated.

- Use a Reference Computer to Create and Maintain AppLocker Policies

  This topic describes the steps to use an AppLocker reference computer to prepare application control policies for deployment by using Group Policy or other means.
- Deploy AppLocker Policies by Using the Enforce Rules Setting

  This topic describes the steps to deploy the AppLocker policy by changing the enforcement setting to **Audit only** or to **Enforce rules**.

See also

AppLocker Policies Deployment Guide

Use a Reference Computer to Create and Maintain AppLocker Policies

This topic for the IT professional describes the steps to create and maintain AppLocker policies by using a reference computer.

Background and prerequisites

An AppLocker reference computer is a baseline computer you can use to configure policies and can subsequently be used to maintain AppLocker policies. For the procedure to configure a reference computer, see Configure the AppLocker Reference Computer.

An AppLocker reference computer that is used to create and maintain AppLocker policies should contain the corresponding applications for each organizational unit (OU) to mimic your production environment.

Important

The reference computer must be running one of the supported editions of Windows. For information about operating system requirements for AppLocker, see Requirements to Use AppLocker.

You can perform AppLocker policy testing on the reference computer by using the **Audit only** enforcement setting or Windows PowerShell cmdlets. You can also use the reference computer as part of a testing configuration that includes policies that are created by using Software Restriction Policies.

Step 1: Automatically generate rules on the reference computer

With AppLocker, you can automatically generate rules for all files within a folder. AppLocker scans the specified folder and creates the condition types that you choose for each file in that folder. For the procedure to do this, see Run the Automatically Generate Rules Wizard.

Note

If you run this wizard to create your first rules for a Group Policy Object (GPO), after you complete the wizard, you will be prompted to create the default rules, which allow critical system files to run. You can edit the default rules at any time. If your organization has decided to edit the default rules or create custom rules to allow the Windows system files to run, ensure that you delete the default rules after you replace them with your custom rules.

Step 2: Create the default rules on the reference computer

AppLocker includes default rules for each rule collection. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed in an AppLocker rule collection. You must run the default rules for each rule collection. For information about default rules and considerations for using them, see Understanding AppLocker Default Rules. For the procedure to create default rules, see Create AppLocker Default Rules.

Important

You can use the default rules as a template when you create your own rules. This allows files within the Windows directory to run. However, these rules are only meant to function as a starter policy when you are first testing AppLocker rules.

Step 3: Modify rules and the rule collection on the reference computer

If AppLocker policies are currently running in your production environment, export the policies from the corresponding GPOs and save them to the reference computer. For the procedure to do this, see Export an AppLocker Policy from a GPO. If no AppLocker policies have been deployed, create the rules and develop the policies by using the following procedures:

- Create a Rule That Uses a Publisher Condition
- Create a Rule That Uses a File Hash Condition
- Create a Rule That Uses a Path Condition
- Edit AppLocker Rules
- Configure Exceptions for an AppLocker Rule
- Delete an AppLocker Rule
- Enable the DLL Rule Collection
- Enforce AppLocker Rules

Step 4: Test and update AppLocker policy on the reference computer

You should test each set of rules to ensure that they perform as intended. The **Test-AppLockerPolicy** Windows PowerShell cmdlet can be used to determine whether any of the rules in your rule collection will be blocked on your reference computer. Perform the steps on each reference computer that you used to define the AppLocker policy. Ensure that the reference computer is joined to the domain and that it is receiving the AppLocker policy from the appropriate GPO. Because AppLocker rules are inherited from linked GPOs, you should deploy all of the rules to simultaneously test all of your test GPOs. Use the following procedures to complete this step:

- Test an AppLocker Policy with Test-AppLockerPolicy
- Discover the Effect of an AppLocker Policy

Caution

If you have set the enforcement setting on the rule collection to **Enforce rules** or you have not configured the rule collection, the policy will be implemented when the GPO is updated in the next step. If you have set the enforcement setting on the rule collection to **Audit only**, application access events are written to the AppLocker log, and the policy will not take effect.

Step 5: Export and import the policy into production

When the AppLocker policy has been tested successfully, it can be imported into the GPO (or imported into individual computers that are not managed by Group Policy) and checked for its intended effectiveness. To do this, perform the following procedures:

- Export an AppLocker Policy to an XML File
- Import an AppLocker Policy into a GPO or
- Discover the Effect of an AppLocker Policy

If the AppLocker policy enforcement setting is **Audit only** and you are satisfied that the policy is fulfilling your intent, you can change it to **Enforce rules**. For information about how to change the enforcement setting, see Configure an AppLocker Policy for Enforce Rules.

Step 6: Monitor the effect of the policy in production

If additional refinements or updates are necessary after a policy is deployed, use the appropriate following procedures to monitor and update the policy:

- Monitor Application Usage with AppLocker
- Edit an AppLocker Policy
- Refresh an AppLocker Policy

See also

Deploy the AppLocker Policy into Production

Determine Which Applications Are Digitally Signed on a Reference Computer

This topic for the IT professional describes how to use AppLocker logs and tools to determine which applications are digitally signed.

The Windows PowerShell cmdlet **Get-AppLockerFileInformation** can be used to determine which applications installed on your reference computers are digitally signed. Perform the following steps on each reference computer that you used to define the AppLocker policy. The computer does not need to be joined to the domain.

Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.

To determine which applications are digitally signed on a reference computer

1. From the command line on the reference computer, run **Get-AppLockerFileInformation** with the appropriate parameters.

   The **Get-AppLockerFileInformation** cmdlet retrieves the AppLocker file information from a list of files or from an event log. File information that is retrieved can include publisher information, file hash information, and file path information. File information from an event log may not contain all of these fields. Files that are not signed do not have any publisher information.
2. Analyze the publisher's name and digital signature status from the output of the command.

For command parameters, syntax, and examples, see Get-AppLockerFileInformation.

See Also

Use a Reference Computer to Create and Maintain AppLocker Policies

Configure the AppLocker Reference Computer

This topic for the IT professional describes the steps to create an AppLocker policy platform structure on a reference computer.

An AppLocker reference computer that is used for the development and deployment of AppLocker policies should mimic the directory structure and corresponding applications in the organizational unit (OU) or business group for the production environment. On a reference computer, you can:

- Maintain an application list for each business group.
- Develop AppLocker policies by creating individual rules or by creating a policy by automatically generating rules.
- Create the default rules to allow the Windows system files to run properly.
- Run tests and analyze the event logs to determine the affect of the policies that you intend to deploy.

The reference computer does not need to be joined to a domain, but it must be able to import and export AppLocker policies in XML format. The reference computer must be running one of the supported editions of Windows as listed in Requirements to Use AppLocker.

To configure a reference computer

1. If the operating system is not already installed, install one of the supported editions of Windows on the computer.

Note

If you have the Group Policy Management Console (GPMC) installed on another computer to test your implementation of AppLocker policies, you can export the policies to that computer.

2. Configure the administrator account.

   To update local policies, you must be a member of the local Administrators group. To update domain policies, you must be a member of the Domain Admins group or have been delegated privileges to use Group Policy to update a Group Policy Object (GPO).

3. Install all applications that run in the targeted business group or OU by using the same directory structure.

   The reference computer should be configured to mimic the structure of your production environment. It depends on having the same applications in the same directories to accurately create the rules.

4. Import the AppLocker Windows PowerShell cmdlet module.

   To use the AppLocker cmdlets, you must first import the AppLocker module by using the following command at the Windows PowerShell command prompt: `C:\PS> Import-Module AppLocker`. Scripting must be enabled on the computer. For information about Windows PowerShell, see the Windows PowerShell Help file (WindowsPowerShellHelp.chm). For information about using the cmdlets, see Use the AppLocker Windows PowerShell Cmdlets.

See also

- After you configure the reference computer, you can create the AppLocker rule collections. You can build, import, or automatically generate the rules. For procedures to do this, see Working with AppLocker Rules.
- Use a Reference Computer to Create and Maintain AppLocker Policies

Maintain AppLocker Policies

This topic describes how to maintain rules within AppLocker policies in Windows Server 2012 and Windows 8.

Common AppLocker maintenance scenarios include:

- A new application is deployed, and you need to update an AppLocker policy.
- A new version of an application is deployed, and you need to either update an AppLocker policy or create a new rule to update the policy.
- An application is no longer supported by your organization, so you need to prevent it from being used.
- An application appears to be blocked but should be allowed.
- An application appears to be allowed but should be blocked.

- A single user or small subset of users needs to use a specific application that is blocked.

There are two methods you can use to maintain AppLocker policies:

- Maintaining AppLocker policies by using Group Policy
- Maintaining AppLocker policies by using the Local Security Policy snap-in

As new applications are deployed or existing applications are removed by your organization or updated by the software publisher, you might need to make revisions to your rules and update the Group Policy Object (GPO) to ensure that your policy is current.

You can edit an AppLocker policy by adding, changing, or removing rules. However, you cannot specify a version for the AppLocker policy by importing additional rules. To ensure version control when modifying an AppLocker policy, use Group Policy management software that allows you to create versions of GPOs, such as Microsoft Advanced Group Policy Management (AGPM). For more information about AGPM, see Advanced Group Policy Management Overview (http://go.microsoft.com/fwlink/?LinkId=145013).

Caution

You should not edit an AppLocker rule collection while it is being enforced in Group Policy. Because AppLocker controls what files are allowed to run, making changes to a live policy can create unexpected behavior.

Maintaining AppLocker policies by using Group Policy

For every scenario, the steps to maintain an AppLocker policy distributed by Group Policy include the following tasks.

Step 1: Understand the current behavior of the policy

Before modifying a policy, evaluate how the policy is currently implemented. For example, if a new version of the application is deployed, you can use **Test-AppLockerPolicy** to verify the effectiveness of your current policy for that application. To read the procedures necessary to understand the current behavior of the policy, see Discovering the Effect of an AppLocker Policy. Updating your AppLocker planning document will help you track your findings. For information about creating this document, see Creating Your AppLocker Planning Document. For information about **Test-AppLockerPolicy** and examples of how to use it, see Test-AppLockerPolicy (http://go.microsoft.com/fwlink/?LinkId=169000).

Note

These documents apply to Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7 operating systems.

Step 2: Export the AppLocker policy from the GPO

Updating an AppLocker policy that is currently enforced in your production environment can have unintended results. Therefore, export the policy from the GPO and update the rule or rules by using AppLocker on your AppLocker reference or test computer. To prepare an AppLocker policy for modification, see Export an AppLocker Policy from a GPO

Step 3: Update the AppLocker policy by editing the appropriate AppLocker rule

After the AppLocker policy has been exported from the GPO into the AppLocker reference or test computer, or has been accessed on the local computer, the specific rules can be modified as required.

To modify AppLocker rules, see the following:

- Edit AppLocker Rules
- Merge AppLocker Policies by Using Set-ApplockerPolicy or Merge AppLocker Policies Manually
- Delete an AppLocker Rule
- Enforce AppLocker Rules

Step 4: Test the AppLocker policy

You should test each collection of rules to ensure that the rules perform as intended. (Because AppLocker rules are inherited from linked GPOs, you should deploy all rules for simultaneous testing in all test GPOs.) For steps to perform this testing, see Test and Update an AppLocker Policy.

Step 5: Import the AppLocker policy into the GPO

After testing, import the AppLocker policy back into the GPO for implementation. To update the GPO with a modified AppLocker policy, see Import an AppLocker Policy into a GPO.

Step 6: Monitor the resulting policy behavior

After deploying a policy, evaluate the policy's effectiveness. For steps to understand the new behavior of the policy, see Discovering the Effect of an AppLocker Policy.

Maintaining AppLocker policies by using the Local Security Policy snap-in

For every scenario, the steps to maintain an AppLocker policy by using the Local Group Policy Editor or the Local Security Policy snap-in include the following tasks.

Step 1: Understand the current behavior of the policy

Before modifying a policy, evaluate how the policy is currently implemented. To read the procedures necessary to understand the current behavior of the policy, see Monitor Application Usage with AppLocker. Updating your AppLocker planning document will help you track your findings. For information about creating this document, see Creating Your AppLocker Planning Document.

Note
The AppLocker Planning Guide applies to Windows Server 2012, Windows Server 2008 R2, Windows 8 and Windows 7.

Step 2: Update the AppLocker policy by modifying the appropriate AppLocker rule

Rules are grouped into a collection, which can have the policy enforcement setting applied to it. By default, AppLocker rules do not allow users to open or run any files that are not specifically allowed.

To modify AppLocker rules, see the appropriate topic listed on Administer AppLocker.

Step 3: Test the AppLocker policy

You should test each collection of rules to ensure that the rules perform as intended. For steps to perform this testing, see Test and Update an AppLocker Policy.

Step 4: Deploy the policy with the modified rule

You can export and then import AppLocker policies to deploy the policy to other computers running Windows 8 or Windows Server 2012. To perform this task, see Export an AppLocker Policy to an XML File and Import an AppLocker Policy from Another Computer.

Step 5: Monitor the resulting policy behavior

After deploying a policy, evaluate the policy's effectiveness. For steps to understand the new behavior of the policy, see Discovering the Effect of an AppLocker Policy in Discovering the Effect of an AppLocker Policy.

Note
This topic applies to Windows Server 2012, Windows Server 2008 R2, Windows 8 and Windows 7.

Additional resources

- For steps to perform other AppLocker policy tasks, see Administer AppLocker.

See Also

AppLocker Overview

AppLocker Technical Reference

This overview topic for the IT professional explains the AppLocker application control feature and provides links to the topics in the technical reference.

AppLocker was introduced in Windows Server 2008 R2 and Windows 7 that advances the application control features and functionality of Software Restriction Policies. AppLocker contains new capabilities and extensions that allow you to create rules to allow or deny applications from running based on unique identities of files and to specify which users or groups can run those applications. These capabilities have been enhanced in Windows Server 2012 and Windows 8.

The AppLocker Technical Reference contains the following subjects:

What Is AppLocker?

This topic describes what AppLocker is and how its features differ from Software Restriction Policies.

Requirements to Use AppLocker

This topic lists software requirements that allow you to use AppLocker.

When to Use AppLocker

This topic describes the scenarios and applications in which AppLocker is useful.

AppLocker Policy Use Scenarios

This topic lists the various application control scenarios in which AppLocker policies can be effectively implemented.

How AppLocker Works

This topic explains how the application identity service and the AppLocker engine work to set and enforce policies.

AppLocker Architecture and Components

This topic describes the components that AppLocker uses for creating and applying application control policies.

AppLocker Processes and Interactions

This topic explains how AppLocker processes policies for both auditing and enforcing each rule collection type.

AppLocker Functions

This topic lists the functions and security levels for the Software Restriction Policies and AppLocker policies beginning with Windows Server 2003.

Security Considerations for AppLocker

This topic describes the threats and vulnerabilities to AppLocker and how you can mitigate them.

Tools to Use with AppLocker

This topic lists and describes the tools available to evaluate, review, create, modify, and test AppLocker policies.

AppLocker Settings

This topic lists the relevant Windows operating system settings for AppLocker.

AppLocker PowerShell Command Reference

This topic introduces and provides links to the Windows PowerShell cmdlets that can be used to help manage AppLocker.

Tip

To digitally save or print pages from this library, click **Export** (in the upper-right corner of the page), and then follow the instructions.

What Is AppLocker?

This topic for the IT professional describes what AppLocker is and how its features differ from Software Restriction Policies.

AppLocker was introduced in Windows Server 2008 R2 and Windows 7 that advances the application control features and functionality of Software Restriction Policies. AppLocker contains new capabilities and extensions that allow you to create rules to allow or deny applications from running based on unique identities of files and to specify which users or groups can run those applications.

Using AppLocker, you can:

- Control the following types of applications: executable files (.exe and .com), scripts (.js, .ps1, .vbs, .cmd, and .bat), Windows Installer files (.mst, .msi and .msp), and DLL files (.dll and .ocx), and packaged apps and packaged app installers (appx).
- Define rules based on file attributes derived from the digital signature, including the publisher, product name, file name, and file version. For example, you can create rules based on the publisher attribute that is persistent through updates, or you can create rules for a specific version of a file.
- Assign a rule to a security group or an individual user.
- Create exceptions to rules. For example, you can create a rule that allows all Windows processes to run except Registry Editor (Regedit.exe).
- Use audit-only mode to deploy the policy and understand its impact before enforcing it.
- Import and export rules. The import and export affects the entire policy. For example, if you export a policy, all of the rules from all of the rule collections are exported, including the enforcement settings for the rule collections. If you import a policy, all criteria in the existing policy are overwritten.

- Streamline creating and managing AppLocker rules by using Windows PowerShell cmdlets.

AppLocker helps reduce administrative overhead and helps reduce the organization's cost of managing computing resources by decreasing the number of help desk calls that result from users running unapproved applications.

For information about the application control scenarios that AppLocker addresses, see AppLocker Policy Use Scenarios.

What features are different between Software Restriction Policies and AppLocker?

Feature differences

The following table compares AppLocker to Software Restriction Policies.

Feature	Software Restriction Policies	AppLocker
Rule scope	All users	Specific user or group
Rule conditions provided	File hash, path, certificate, registry path, and Internet zone	File hash, path, and publisher
Rule types provided	Defined by the security levels: • Disallowed • Basic User • Unrestricted	Allow and deny
Default rule action	Unrestricted	Implicit deny
Audit-only mode	No	Yes
Wizard to create multiple rules at one time	No	Yes
Policy import or export	No	Yes
Rule collection	No	Yes
Windows PowerShell support	No	Yes
Custom error messages	No	Yes

Application control function differences

The following table compares the application control functions of Software Restriction Policies (SRP) and AppLocker.

Application control function	SRP	AppLocker
Operating system scope	SRP policies can be applied to all Windows operating systems beginning with Windows XP and Windows Server 2003.	AppLocker policies apply only to those supported operating system versions and editions listed in Requirements to Use AppLocker. But these systems can also use SRP. **Note** Use different GPOs for SRP and AppLocker rules.
User support	SRP allows users to install applications as an administrator.	AppLocker policies are maintained through Group Policy, and only the administrator of the computer can update an AppLocker policy. AppLocker permits customization of error messages to direct users to a Web page for help.
Policy maintenance	SRP policies are updated by using the Local Security Policy snap-in or the Group Policy Management Console (GPMC).	AppLocker policies are updated by using the Local Security Policy snap-in or the GPMC. AppLocker supports a small set of PowerShell cmdlets to aid in administration and maintenance.
Policy management infrastructure	To manage SRP policies, SRP uses Group Policy within a domain and the Local Security Policy snap-in for a local computer.	To manage AppLocker policies, AppLocker uses Group Policy within a domain and the Local Security Policy snap-in for a local computer.
Block malicious scripts	Rules for blocking malicious scripts prevents all scripts associated with the Windows Script Host from running, except those that are digitally signed by your organization.	AppLocker rules can control the following file formats: .ps1, .bat, .cmd, .vbs, and .js. In addition, you can set exceptions to allow specific files to run.
Manage software installation	SRP can prevent all Windows Installer packages from	The Windows Installer rule collection is a set of rules

Application control function	SRP	AppLocker
	installing. It allows .msi files that are digitally signed by your organization to be installed.	created for Windows Installer file types (.mst, .msi and .msp) to allow you to control the installation of files on client computers and servers.
Manage all software on the computer	All software is managed in one rule set. By default, the policy for managing all software on a computer disallows all software on the user's computer, except software that is installed in the Windows folder, Program Files folder, or subfolders.	Unlike SRP, each AppLocker rule collection functions as an allowed list of files. Only the files that are listed within the rule collection will be allowed to run. This configuration makes it easier for administrators to determine what will occur when an AppLocker rule is applied.
Different policies for different users	Rules are applied uniformly to all users on a particular computer.	On a computer that is shared by multiple users, an administrator can specify the groups of users who can access the installed software. Using AppLocker, an administrator can specify the user to whom a specific rule should apply.

See Also

AppLocker Technical Reference

Requirements to Use AppLocker

This topic for the IT professional lists software requirements to use AppLocker on the supported Windows operating systems.

General requirements

To use AppLocker, you need:

- A computer running a supported operating system to create the rules. The computer can be a domain controller.

- For Group Policy deployment, at least one computer with the Group Policy Management Console (GPMC) or Remote Server Administration Tools (RSAT) installed to host the AppLocker rules.
- Computers running a supported operating system to enforce the AppLocker rules that you create.

Note
You can use Software Restriction Policies with AppLocker, but with some limitations. For more information, see Use AppLocker and Software Restriction Policies in the Same Domain.

Operating system requirements

The following table show the on which operating systems AppLocker features are supported.

Version	Can be configured	Can be enforced	Available rules	Notes
Windows Server 2012 R2	Yes	Yes	Packaged apps Executable Windows Installer Script DLL	
Windows 8.1	Yes	Yes	Packaged apps Executable Windows Installer Script DLL	Only the Enterprise edition supports AppLocker
Windows RT 8.1	No	No	NA	
Windows Server 2012 Standard	Yes	Yes	Packaged apps Executable Windows Installer Script DLL	
Windows Server 2012 Datacenter	Yes	Yes	Packaged apps Executable Windows Installer Script DLL	
Windows 8 Pro	No	No	NA	
Windows 8	Yes	Yes	Packaged apps	

Version	Can be configured	Can be enforced	Available rules	Notes
Enterprise			Executable Windows Installer Script DLL	
Windows RT	No	No	NA	
Windows Server 2008 R2 Standard	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows Server 2008 R2 Enterprise	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows Server 2008 R2 Datacenter	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows Server 2008 R2 for Itanium-Based Systems	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows 7 Ultimate	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows 7 Enterprise	Yes	Yes	Executable Windows Installer Script DLL	Packaged app rules will not be enforced.
Windows 7 Professional	Yes	No	Executable Windows Installer Script DLL	No AppLocker rules are enforced.

AppLocker is not supported on versions of the Windows operating system not listed above. Software Restriction Policies can be used with those versions. However, the SRP Basic User feature is not supported on the above operating systems.

See also

Administer AppLocker

Monitor Application Usage with AppLocker

Optimize AppLocker Performance

Use AppLocker and Software Restriction Policies in the Same Domain

Manage Packaged Apps with AppLocker

AppLocker Policies Design Guide

AppLocker Policy Use Scenarios

This topic for the IT professional lists the various application control scenarios in which AppLocker policies can be effectively implemented.

AppLocker can help you improve the management of application control and the maintenance of application control policies. Application control scenarios addressed by AppLocker can be categorized as follows:

1. **Application inventory**

   AppLocker has the ability to enforce its policy in an audit-only mode where all application access activity is collected in event logs for further analysis. Windows PowerShell cmdlets are also available to help you understand application usage and access.

2. **Protection against unwanted software**

   AppLocker has the ability to deny applications from running simply by excluding them from the list of allowed applications per business group or user. If an application is not specifically identified by its publisher, installation path, or file hash, the attempt to run the application fails.

3. **Licensing conformance**

   AppLocker can provide an inventory of software usage within your organization, so you can identify the software that corresponds to your software licensing agreements and restrict application usage based on licensing agreements.

4. **Software standardization**

   AppLocker policies can be configured to allow only supported or approved applications to run on computers within a business group. This permits a more uniform application deployment.

5. **Manageability improvement**

   AppLocker policies can be modified and deployed through your existing Group Policy infrastructure and can work in conjunction with policies created by using Software Restriction Policies. As you manage ongoing change in your support of a business group's applications, you can modify policies and use the AppLocker cmdlets to test the policies for the expected

results. You can also design application control policies for situations in which users share computers.

Use scenarios

The following are examples of scenarios in which AppLocker can be used:

- Your organization implements a policy to standardize the applications used within each business group, so you need to determine the expected usage compared to the actual usage.
- The security policy for application usage has changed, and you need to evaluate where and when those deployed applications are being accessed.
- Your organization's security policy dictates the use of only licensed software, so you need to determine which applications are not licensed or prevent unauthorized users from running licensed software.
- An application is no longer supported by your organization, so you need to prevent it from being used by everyone.
- Your organization needs to restrict the use of Windows Store apps to just those your organization approves of or develops.
- The potential that unwanted software can be introduced in your environment is high, so you need to reduce this threat.
- The license to an application has been revoked or is expired in your organization, so you need to prevent it from being used by everyone.
- A new application or a new version of an application is deployed, and you need to allow certain groups to use it.
- Specific software tools are not allowed within the organization, or only specific users have access to those tools.
- A single user or small group of users needs to use a specific application that is denied for all others.
- Some computers in your organization are shared by people who have different software usage needs.
- In addition to other measures, you need to control the access to sensitive data through application usage.

See also

To explore specific use scenarios, see AppLocker Step-by-Step Scenarios.

See Also

AppLocker Technical Reference

How AppLocker Works

This topic for the IT professional provides links to topics about AppLocker architecture and components, processes and interactions, rules and policies..

The following topics explain how AppLocker policies for each of the rule condition types are evaluated:

- AppLocker Architecture and Components
- AppLocker Processes and Interactions

The following topics explain how AppLocker rules and policies work:

- Understanding AppLocker Rule Behavior
- Understanding AppLocker Rule Exceptions
- Understanding AppLocker Rule Collections
- Understanding AppLocker Allow and Deny Actions on Rules
- Understanding AppLocker Rule Condition Types
  - Understanding the Publisher Rule Condition in AppLocker
  - Understanding the Path Rule Condition in AppLocker
  - Understanding the File Hash Rule Condition in AppLocker
- Understanding AppLocker Default Rules
  - Executable Rules in AppLocker
  - Windows Installer Rules in AppLocker
  - Script Rules in AppLocker
  - DLL Rules in AppLocker
  - Packaged Apps and Packaged App Installer Rules in AppLocker

Additional resources

- AppLocker Policies Design Guide
- AppLocker Policies Deployment Guide
- Administer AppLocker

See Also

AppLocker Technical Architecture

Understanding AppLocker Rule Behavior

This topic describes how AppLocker rules are enforced by using the allow and deny options in AppLocker.

If no AppLocker rules for a specific rule collection exist, all files with that file format are allowed to run. However, when an AppLocker rule for a specific rule collection is created, only the files explicitly allowed in a rule are permitted to run. For example, if you create an executable rule that allows .exe files in *%SystemDrive%\FilePath* to run, only executable files located in that path are allowed to run.

A rule can be configured to use either an allow or deny action:

- **Allow**. You can specify which files are allowed to run in your environment and for which users or groups of users. You can also configure exceptions to identify files that are excluded from the rule.
- **Deny**. You can specify which files are *not* allowed to run in your environment and for which users or groups of users. You can also configure exceptions to identify files that are excluded from the rule.

Important
You can use a combination of allow actions and deny actions. However, we recommend using allow actions with exceptions because deny actions override allow actions in all cases. Deny actions can also be circumvented. For example, if you configure a deny action for a file or folder path, the user can still run the file from any other path.

See Also

How AppLocker Works

Understanding AppLocker Rule Exceptions

This topic describes the result of applying AppLocker rule exceptions to rule collections.

You can apply AppLocker rules to individual users or a group of users. If you apply a rule to a group of users, all users in that group are affected by that rule. If you need to allow a subset of a user group to use an application, you can create a special rule for that subset.

For example, the rule "Allow Everyone to run Windows except Registry Editor" allows everyone in the organization to run Windows but does not allow anyone to run Registry Editor. The effect of this rule would prevent users such as help desk personnel from running a program that is necessary for their support tasks. To resolve this problem, create a second rule that applies to the Helpdesk user group: "Allow Helpdesk to run Registry Editor." If you create a deny rule that does not allow any users to run Registry Editor, the deny rule will override the second rule that allows the Helpdesk user group to run Registry Editor.

See Also

How AppLocker Works

Understanding AppLocker Rule Collections

This topic explains the five different types of AppLocker rules used to enforce AppLocker policies.

An AppLocker rule collection is a set of rules that apply to one of five types:

- Executable files: .exe and .com
- Windows Installer files: .msi, mst and .msp
- Scripts: .ps1, .bat, .cmd, .vbs, and .js
- DLLs: .dll and .ocx
- Packaged apps and Packaged app installers: .appx

If you use DLL rules, a DLL allow rule has to be created for each DLL that is used by all of the allowed applications.

Important
Each application can load several DLLs, and AppLocker must check each DLL before it is allowed to run. Therefore, creating DLL rules might cause performance problems on some computers. Denying some DLLs from running can also create application compatibility problems. As a result, the DLL rule collection is not enabled by default.

For information about how to enable the DLL rule collection, see Enable the DLL Rule Collection.

See Also

How AppLocker Works

Understanding AppLocker Allow and Deny Actions on Rules

This topic explains the differences between allow and deny actions on AppLocker rules.

Allow action versus deny action on rules

Unlike Software Restriction Policies (SRP), each AppLocker rule collection functions as an allowed list of files. Only the files that are listed within the rule collection are allowed to run. This configuration makes it easier to determine what will occur when an AppLocker rule is applied.

You can also create rules that use the deny action. When applying rules, AppLocker first checks whether any explicit deny actions are specified in the rule list. If you have denied a file from running in a rule collection, the deny action will take precedence over any allow action, regardless of which Group Policy Object (GPO) the rule was originally applied in. Because AppLocker functions as an allowed list by default, if no rule explicitly allows or denies a file from running, AppLocker's default deny action will block the file.

Deny rule considerations

Although you can use AppLocker to create a rule to allow all files to run and then use rules to deny specific files, this configuration is not recommended. The deny action is generally less secure than the allow action because a malicious user could modify the file to invalidate the rule. Deny actions can also be circumvented. For example, if you configure a deny action for a file or folder path, the user can still run the file from any other path. The following table details security concerns for different rule conditions with deny actions.

Rule condition	Security concern with deny action
Publisher	A user could modify the properties of a file (for example, re-signing the file with a different certificate).
File hash	A user could modify the hash for a file.
Path	A user could move the denied file to a different location and run it from there.

Important
If you choose to use the deny action on rules, you must ensure that you first create rules that allow the Windows system files to run. AppLocker enforces rules for allowed applications by default, so after one or more rules have been created for a rule collection (affecting the Windows system files), only the applications that are listed as being allowed will be permitted to run. Therefore, creating a single rule in a rule collection to deny a malicious file from running will also deny all other files on the computer from running.

See Also

How AppLocker Works

Understanding AppLocker Rule Condition Types

This topic for the IT professional describes the three types of AppLocker rule conditions.

Rule conditions are criteria that the AppLocker rule is based on. Primary conditions are required to create an AppLocker rule. The three primary rule conditions are publisher, path, and file hash.

Publisher

To use a publisher condition, the files must be digitally signed by the software publisher, or you must do so by using an internal certificate. Rules that are specified to the version level might have

to be updated when a new version of the file is released. For more information about this rule condition, see Understanding the Publisher Rule Condition in AppLocker.

Path

Any file can be assigned this rule condition; however, because path rules specify locations within the file system, any subdirectory will also be affected by the rule (unless explicitly exempted). For more information about this rule condition, see Understanding the Path Rule Condition in AppLocker.

File hash

Any file can be assigned this rule condition; however, the rule must be updated each time a new version of the file is released because the hash value is unique to that the version of the file. For more information about this rule condition, see Understanding the File Hash Rule Condition in AppLocker.

Considerations

Selecting the appropriate condition for each rule depends on the overall application control policy goals of the organization, the AppLocker rule maintenance goals, and the condition of the existing (or planned) application deployment. The following questions can help you decide which rule condition to use.

1. Is the file digitally signed by a software publisher?

   If the file is signed by a software publisher, we recommend that you create rules with publisher conditions. You may still create file hash and path conditions for signed files. However, if the file is not digitally signed by a software publisher, you can:

   - Sign the file by using an internal certificate.
   - Create a rule by using a file hash condition.
   - Create a rule by using a path condition.

   > **Note**
   > To determine how many applications on a reference computer are digitally signed, you can use the **Get-AppLockerFileInformation** Windows PowerShell cmdlet for a directory of files. For example, `Get-AppLockerFileInformation -Directory C:\Windows\ -FileType EXE -recurse` displays the properties for all .exe and .com files within the Windows directory.

2. What rule condition type does your organization prefer?

   If your organization is already using Software Restriction Policies (SRP) to restrict what files users can run, rules using file hash or path conditions are probably already in place.

   > **Note**
   > For a list of supported operating system versions and editions to which SRP and AppLocker rules can be applied, see Requirements to Use AppLocker.

See Also

How AppLocker Works

Understanding the Publisher Rule Condition in AppLocker

This topic explains the AppLocker publisher rule condition, what controls are available, and how it is applied.

Publisher conditions can be made only for files that are digitally signed; this condition identifies an application based on its digital signature and extended attributes. The digital signature contains information about the company that created the application (the publisher). The extended attributes, which are obtained from the binary resource, contain the name of the product that the application is part of and the version number of the application. The publisher may be a software development company, such as Microsoft, or the Information Technology department of your organization.

Publisher conditions are easier to maintain than file hash conditions and are generally more secure than path conditions. Rules that are specified to the version level might have to be updated when a new version of the file is released. The following table describes the advantages and disadvantages of the publisher condition.

Publisher condition advantages	Publisher condition disadvantages
• Frequent updating is not required. • You can apply different values within a certificate. • A single rule can be used to allow an entire product suite. • You can use the asterisk (*) wildcard character within a publisher rule to specify that any value should be matched.	• The file must be signed. • Although a single rule can be used to allow an entire product suite, all files in the suite must be signed uniformly.

Wildcard characters can be used as values in the publisher rule fields according to the following specifications:

- **Publisher**

  The asterisk (*) character used by itself represents any publisher. When combined with any string value, the rule is limited to the publisher with a value in the signed certificate that matches the character string. In other words, the asterisk is not treated as a wildcard character if used with other characters in this field. For example, using the characters "M*" limits the publisher name to only a publisher with the name "M*." Using the characters "*x*"

limits the publisher name only to the name “*x*”. A question mark (?) is not a valid wildcard character in this field.

- **Product name**

  The asterisk (*) character used by itself represents any product name. When combined with any string value, the rule is limited to the product of the publisher with a value in the signed certificate that matches the character string. In other words, the asterisk is not treated as a wildcard character if used with other characters in this field. A question mark (?) is not a valid wildcard character in this field.

- **File name**

  Either the asterisk (*) or question mark (?) characters used by themselves represent any and all file names. When combined with any string value, the string is matched with any file name containing that string.

- **File version**

  The asterisk (*) character used by itself represents any file version. If you want to limit the file version to a specific version or as a starting point, you can state the file version and then use the following options to apply limits:

  - **Exactly**. The rule applies only to this version of the application.
  - **And above**. The rule applies to this version and all later versions.
  - **And Below**. The rule applies to this version and all earlier versions.

The following table describes how a publisher condition is applied.

Option		The publisher condition allows or denies…
All signed files		All files that are signed by a publisher.
Publisher only		All files that are signed by the named publisher.
Publisher and product name		All files for the specified product that are signed by the named publisher.
Publisher, product name, and file name		Any version of the named file for the named product that is signed by the publisher.
Publisher, product name, file name, and file version	**Exactly**	The specified version of the named file for the named product that is signed by the publisher.
Publisher, product name, file	**And above**	The specified version of the

Option		The publisher condition allows or denies…
name, and file version		named file and any new releases for the product that are signed by the publisher.
Publisher, product name, file name, and file version	**And below**	The specified version of the named file and any older versions for the product that are signed by the publisher.
Custom		You can edit the **Publisher**, **Product name**, **File name**, and **Version** fields to create a custom rule.

For an overview of the three types of AppLocker rule conditions and explanations of the advantages and disadvantages of each, see Understanding AppLocker Rule Condition Types.

See Also

How AppLocker Works

Understanding the Path Rule Condition in AppLocker

This topic explains the AppLocker path rule condition, the advantages and disadvantages, and how it is applied.

The path condition identifies an application by its location in the file system of the computer or on the network.

When creating a rule that uses a deny action, path conditions are less secure than publisher and file hash conditions for preventing access to a file because a user could easily copy the file to a different location than the location specified in the rule. Because path rules specify locations within the file system, you should ensure that there are no subdirectories that are writable by non-administrators. For example, if you create a path rule for C:\ with the allow action, any file under that location will be allowed to run, including within users' profiles. The following table describes the advantages and disadvantages of the path condition.

Path condition advantages	Path condition disadvantages
• You can easily control many directories or a single file. • You can use the asterisk (*) as a wildcard character within path rules.	• It might be less secure if a rule that is configured to use a directory path contains subdirectories that are writable by non-administrators. • You must specify the full path to a file or folder when creating path rules so that the rule will be properly enforced.

AppLocker does not enforce rules that specify paths with short names. You should always specify the full path to a file or folder when creating path rules so that the rule will be properly enforced.

The asterisk (*) wildcard character can be used within **Path** field. The asterisk (*) character used by itself represents any path. When combined with any string value, the rule is limited to the path of the file and all the files under that path. For example, %ProgramFiles%\Internet Explorer* indicates that all files and subfolders within the Internet Explorer folder will be affected by the rule.

AppLocker uses path variables for well-known directories in Windows. Path variables are not environment variables. The AppLocker engine can only interpret AppLocker path variables. The following table details these path variables.

Windows directory or drive	AppLocker path variable	Windows environment variable
Windows	%WINDIR%	%SystemRoot%
System32	%SYSTEM32%	%SystemDirectory%
Windows installation directory	%OSDRIVE%	%SystemDrive%
Program Files	%PROGRAMFILES%	%ProgramFiles% and %ProgramFiles(x86)%
Removable media (for example, CD or DVD)	%REMOVABLE%	
Removable storage device (for example, USB flash drive)	%HOT%	

For an overview of the three types of AppLocker rule conditions and explanations of the advantages and disadvantages of each, see Understanding AppLocker Rule Condition Types.

See Also

How AppLocker Works

Understanding the File Hash Rule Condition in AppLocker

This topic explains the AppLocker file hash rule condition, the advantages and disadvantages, and how it is applied.

File hash rules use a system-computed cryptographic hash of the identified file. For files that are not digitally signed, file hash rules are more secure than path rules. The following table describes the advantages and disadvantages of the file hash condition.

File hash condition advantages	File hash condition disadvantages
Because each file has a unique hash, a file hash condition applies to only one file.	Each time that the file is updated (such as a security update or upgrade), the file's hash will change. As a result, you must manually update file hash rules.

For an overview of the three types of AppLocker rule conditions and explanations of the advantages and disadvantages of each, see Understanding AppLocker Rule Condition Types.

See Also

How AppLocker Works

Understanding AppLocker Default Rules

This topic for IT professional describes the set of rules that can be used to ensure that required Windows system files are allowed to run when the policy is applied.

AppLocker includes default rules for each rule collection. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed in an AppLocker rule collection.

Important
You can use the default rules as a template when creating your own rules. However, these rules are only meant to function as a starter policy when you are first testing AppLocker rules so that the system files in the Windows folders will be allowed to run.

If you require additional application security, you might need to modify the rules created from the built-in default rule collection. For example, the default rule to allow all users to run .exe files in the Windows folder is based on a path condition that allows all files within the Windows folder to

run. The Windows folder contains a Temp subfolder to which the Users group is given the following permissions:

- Traverse Folder/Execute File
- Create Files/Write Data
- Create Folders/Append Data

These permissions settings are applied to this folder for application compatibility. However, because any user can create files in this location, allowing applications to be run from this location might conflict with your organization's security policy.

For information about each default rule in AppLocker, see:

- Executable Rules in AppLocker
- Windows Installer Rules in AppLocker
- Script Rules in AppLocker
- DLL Rules in AppLocker
- Packaged Apps and Packaged App Installer Rules in AppLocker

See Also

How AppLocker Works

Executable Rules in AppLocker

This topic describes the file formats and available default rules for the executable rule collection.

AppLocker defines executable rules as any files with the .exe and .com extensions that are associated with an application. Because all of the default rules for the executable rule collection are based on folder paths, all files under those paths will be allowed. The following table lists the default rules that are available for the executable rule collection.

Purpose	Name	User	Rule condition type
Allow members of the local Administrators group access to run all executable files	(Default Rule) All files	BUILTIN\Administrators	Path: *
Allow all users to run executable files in the Windows folder	(Default Rule) All files located in the Windows folder	Everyone	Path: %windir%*
Allow all users to run	(Default Rule) All	Everyone	Path: %programfiles%*

Purpose	Name	User	Rule condition type
executable files in the Program Files folder	files located in the Program Files folder		

See Also

Understanding AppLocker Default Rules

Windows Installer Rules in AppLocker

This topic describes the file formats and available default rules for the Windows Installer rule collection.

AppLocker defines Windows Installer rules to include only the following file formats:

- .msi
- .msp
- .mst

The purpose of this collection is to allow you to control the installation of files on client computers and servers through Group Policy or the Local Security Policy snap-in. The following table lists the default rules that are available for the Windows Installer rule collection.

Purpose	Name	User	Rule condition type
Allow members of the local Administrators group to run all Windows Installer files	(Default Rule) All Windows Installer files	BUILTIN\Administrators	Path: *
Allow all users to run Windows Installer files that are digitally signed	(Default Rule) All digitally signed Windows Installer files	Everyone	Publisher: * (all signed files)
Allow all users to run	(Default Rule) All Windows Installer files in	Everyone	Path: %windir%\Installer\

Purpose	Name	User	Rule condition type
Windows Installer files that are located in the Windows Installer folder	%systemdrive%\Windows\Installer		*

See Also

Understanding AppLocker Default Rules

Script Rules in AppLocker

This topic describes the file formats and available default rules for the script rule collection.

AppLocker defines script rules to include only the following file formats:

- .ps1
- .bat
- .cmd
- .vbs
- .js

The following table lists the default rules that are available for the script rule collection.

Purpose	Name	User	Rule condition type
Allows members of the local Administrators group to run all scripts	(Default Rule) All scripts	BUILTIN\Administrators	Path: *
Allow all users to run scripts in the Windows folder	(Default Rule) All scripts located in the Windows folder	Everyone	Path: %windir%*
Allow all users to run scripts in the Program Files folder	(Default Rule) All scripts located in the Program Files folder	Everyone	Path: %programfiles%*

See Also

Understanding AppLocker Default Rules

DLL Rules in AppLocker

This topic describes the file formats and available default rules for the DLL rule collection.

AppLocker defines DLL rules to include only the following file formats:

- .dll
- .ocx

The following table lists the default rules that are available for the DLL rule collection.

Purpose	Name	User	Rule condition type
Allows members of the local Administrators group to run all DLLs	(Default Rule) All DLLs	BUILTIN\Administrators	Path: *
Allow all users to run DLLs in the Windows folder	(Default Rule) Microsoft Windows DLLs	Everyone	Path: %windir%*
Allow all users to run DLLs in the Program Files folder	(Default Rule) All DLLs located in the Program Files folder	Everyone	Path: %programfiles%*

Important
If you use DLL rules, a DLL allow rule has to be created for each DLL that is used by all of the allowed applications.

Caution
When DLL rules are used, AppLocker must check each DLL that an application loads. Therefore, users may experience a reduction in performance if DLL rules are used.

See Also

Understanding AppLocker Default Rules

Packaged Apps and Packaged App Installer Rules in AppLocker

This topic explains the AppLocker rule collection for packaged app installers and packaged apps introduced in Windows Server 2012 and Windows 8.

Commonly known as Windows apps, packaged apps can be installed through the Microsoft AppStore or can be side loaded using the Windows PowerShell cmdlets if you have an Enterprise license. Packaged apps can be installed by a standard user unlike some desktop applications that sometimes require administrative privileges for installation. In this topic, desktop applications refer to Win32 apps that run on the classic user desktop.

Typically, an app consists of multiple components – the installer used to install the app and one or more Exes, Dlls or Scripts. With desktop applications, not all those components always share common attributes such as the publisher name, product name and product version. Therefore, AppLocker has to control each of these components separately through different rule collections – Exe, Dll, Script and Windows Installers. In contrast, all the components of a packaged app share the same attributes: Publisher name, Package name and Package version. It is therefore possible to control an entire app with a single rule.

In Windows Server 2012 and Windows 8, AppLocker enforces rules for packaged apps separately from desktop applications. A single AppLocker rule for a packaged app can control both the installation and the running of an app. Because all packaged apps are signed, AppLocker supports only publisher rules for packaged apps. A publisher rule for a packaged app is based on the following attributes of the app:

- Publisher name
- Package name
- Package version

In summary, including AppLocker rules for packaged apps in your policy design provides:

- The ability to control the installation and the running the packaged app
- The ability to control all the components of the app with a single rule rather than controlling individual binaries within the app
- The ability to create application control policies that survive application updates
- Management of packaged apps through Group Policy.

See also

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012 and Windows 8
Product evaluation	Frequently Asked Questions AppLocker Step-by-Step Guide	AppLocker Overview

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012 and Windows 8
Procedures	AppLocker Operations Guide	Administer AppLocker Manage Packaged Apps with AppLocker
Scripting	Using the AppLocker Windows PowerShell Cmdlets	Using the AppLocker Windows PowerShell Cmdlets
Technical content	AppLocker Technical Reference	AppLocker Technical Reference
Design, planning and deployment	AppLocker Policies Design Guide AppLocker Policies Deployment Guide	AppLocker Policies Design Guide AppLocker Policies Deployment Guide
General information and additional resources	AppLocker Documentation for Windows 7 and Windows Server 2008 R2	**AppLocker Navigation Page**

AppLocker Architecture and Components

This topic for IT professional describes AppLocker's basic architecture and its major components.

AppLocker relies on the Application Identity service to provide attributes for a file and to evaluate the AppLocker policy for the file. AppLocker policies are conditional access control entries (ACEs), and policies are evaluated by using the attribute-based access control **SeAccessCheckWithSecurityAttributes** or **AuthzAccessCheck** functions.

AppLocker provides three ways to intercept and validate if a file is allowed to execute according to an AppLocker policy.

A new process is created

When a new process is created, such as an executable file or a packaged app is run, AppLocker invokes the Application Identity component to calculate the attributes of the main executable file used to create a new process. It then updates the new process's token with these attributes and checks the AppLocker policy to verify that the executable file is allowed to run.

A DLL is loaded

When a new DLL loads, a notification is sent to AppLocker to verify that the DLL is allowed to load. AppLocker calls the Application Identity component to calculate the file attributes. It duplicates the existing process token and replaces those Application Identity attributes in the duplicated token with attributes of the loaded DLL. AppLocker then evaluates the policy for this

DLL, and the duplicated token is discarded. Depending on the result of this check, the system either continues to load the DLL or stops the process.

A script is run

Before a script file is run, the script host (for example. for .ps1 files the script host is PowerShell) invokes AppLocker to verify the script. AppLocker invokes the Application Identity component in user-mode with the file name or file handle to calculate the file properties. The script file then is evaluated against the AppLocker policy to verify that it is allowed to run. In each case, the actions taken by AppLocker are written to the event log.

See Also

AppLocker Technical Reference

AppLocker Processes and Interactions

This topic for the IT professional describes the process dependencies and interactions when AppLocker evaluates and enforces rules.

How policies are implemented by AppLocker

AppLocker policies are collections of AppLocker rules that might contain any one of the enforcement settings configured. When applied, each rule is evaluated within the policy and the collection of rules is applied according to the enforcement setting and according to your Group Policy structure.

The AppLocker policy is enforced on a computer through the Application Identity service, which is the engine that evaluates the policies. If the service is not running, policies will not be enforced. The Application Identity service returns the information from the binary—even if product or binary names are empty—to the results pane of the Local Security Policy snap-in.

AppLocker policies are stored in a security descriptor format according to Application Identity service requirements. It uses file path, hash, or fully qualified binary name attributes to form allow or deny actions on a rule. Each rule is stored as an access control entry (ACE) in the security descriptor and contains the following information:

- Either an allow or a deny ACE ("XA" or "XD" in security descriptor definition language (SDDL) form).
- The user security identifier (SID) that this rule is applicable to. (The default is the authenticated user SID, or "AU" in SDDL.)
- The rule condition containing the **appid** attributes.

For example, an SDDL for a rule that allows all files in the %windir% directory to run uses the following format: XA;;FX;;;AU;(APPID://PATH == "%windir%*").

An AppLocker policy for DLLs and executable files is read and cached by kernel mode code, which is part of appid.sys. Whenever a new policy is applied, appid.sys is notified by a policy

converter task. For other file types, the AppLocker policy is read every time a **SaferIdentifyLevel** call is made.

Understanding AppLocker rules

An AppLocker rule is a control placed on a file to govern whether or not it is allowed to run for a specific user or group. Rules apply to five different types, or collections, of files:

- An executable rule controls whether a user or group can run an executable file. Executable files most often have the .exe or .com file name extensions and apply to applications.
- A script rule controls whether a user or group can run scripts with a file name extension of .ps1, .bat, .cmd, .vbs, and .js.
- A Windows Installer rule controls whether a user or group can run files with a file name extension of .msi, mst and .msp (Windows Installer patch).
- A DLL rule controls whether a user or group can run files with a file name extension of .dll and .ocx.
- A Packaged app and Packaged app installer rule controls whether a user or group can run or install a Packaged app. A Packaged app installer has the .appx extension.

There are three different types of conditions that can be applied to rules:

- A publisher condition on a rule controls whether a user or group can run files from a specific software publisher. The file must be signed.
- A path condition on a rule controls whether a user or group can run files from within a specific directory or its subdirectories.
- A file hash condition on a rule controls whether a user or group can run files with matching encrypted hashes.
- Understanding AppLocker Rule Collections

  An AppLocker rule collection is a set of rules that apply to one of the following types: executable files, Windows Installer files, scripts, DLLs, and Packaged apps.
- Understanding AppLocker Rule Condition Types

  Rule conditions are criteria that the AppLocker rule is based on. Primary conditions are required to create an AppLocker rule. The three primary rule conditions are publisher, path, and file hash.
  - Understanding the Publisher Rule Condition in AppLocker
  - Understanding the Path Rule Condition in AppLocker
  - Understanding the File Hash Rule Condition in AppLocker
- Understanding AppLocker Default Rules

  AppLocker includes default rules for each rule collection. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed in an AppLocker rule collection.
  - Executable Rules in AppLocker
  - Windows Installer Rules in AppLocker
  - Script Rules in AppLocker

    - DLL Rules in AppLocker
    - Packaged Apps and Packaged App Installer Rules in AppLocker
- Understanding AppLocker Rule Exceptions

  You can apply AppLocker rules to individual users or a group of users. If you apply a rule to a group of users, all users in that group are affected by that rule. If you need to allow only a subset of a user group to use an application, you can create a special rule for that subset.
- Understanding AppLocker Rule Behavior and Understanding AppLocker Allow and Deny Actions on Rules

  Each AppLocker rule collection functions as an allowed list of files.

Understanding AppLocker policies

An AppLocker policy is a set of rule collections and their corresponding configured enforcement settings that have been applied to one or more computers.

- Understand AppLocker Enforcement Settings

  Rule enforcement is applied only to collections of rules, not individual rules. AppLocker divides the rules into four collections: executable files, Windows Installer files, scripts, and DLL files. The options for rule enforcement are **Not configured**, **Enforce rules**, or **Audit only**. Together, all AppLocker rule collections compose the application control policy, or AppLocker policy. By default, if enforcement is not configured and rules are present in a rule collection, those rules are enforced.

Understanding AppLocker and Group Policy

Group Policy can be used to create, modify, and distribute AppLocker policies in separate objects or in combination with other policies.

- Understand AppLocker Rules and Enforcement Setting Inheritance in Group Policy

  When Group Policy is used to distribute AppLocker policies, rule collections that are not configured will be enforced. Group Policy does not overwrite or replace rules that are already present in a linked Group Policy Object (GPO) and applies the AppLocker rules in addition to existing rules. AppLocker processes the explicit deny rule configuration before the allow rule configuration, and for rule enforcement, the last write to the GPO is applied.

See Also

AppLocker Technical Reference

AppLocker Functions

This topic for the IT professional lists the functions and security levels for the Software Restriction Policies (SRP) and AppLocker features.

Functions

The following list includes the SRP functions beginning with Windows Server 2003 and AppLocker functions beginning with Windows Server 2008 R2 and links to current documentation on MSDN:

- SaferGetPolicyInformation Function (http://go.microsoft.com/fwlink/?LinkId=159781)
- SaferCreateLevel Function (http://go.microsoft.com/fwlink/?LinkId=159782)
- SaferCloseLevel Function (http://go.microsoft.com/fwlink/?LinkId=159783)
- SaferIdentifyLevel Function (http://go.microsoft.com/fwlink/?LinkId=159784)
- SaferComputeTokenFromLevel Function (http://go.microsoft.com/fwlink/?LinkId=159785)
- SaferGetLevelInformation Function (http://go.microsoft.com/fwlink/?LinkId=159787)
- SaferRecordEventLogEntry Function (http://go.microsoft.com/fwlink/?LinkId=159789)
- SaferIsExecutableFileType Function (http://go.microsoft.com/fwlink/?LinkId=159790)

Security level ID

AppLocker and SRP use the security level IDs to stipulate the access requirements to files listed in policies. The following table shows those security levels supported in SRP and AppLocker.

Security level ID	SRP	AppLocker
SAFER_LEVELID_FULLYTRUSTED	Supported	Supported
SAFER_LEVELID_NORMALUSER	Supported	Not supported
SAFER_LEVELID_CONSTRAINED	Supported	Not supported
SAFER_LEVELID_UNTRUSTED	Supported	Not supported
SAFER_LEVELID_DISALLOWED	Supported	Supported

In addition, URL zone ID is not supported in AppLocker.

See Also

AppLocker Technical Reference

Security Considerations for AppLocker

This topic for the IT professional describes the security considerations you need to address when implementing AppLocker.

The purpose of AppLocker is to restrict the access to software, and therefore, the data accessed by the software, to a specific group of users or within a defined business group. The following are security considerations for AppLocker:

AppLocker is deployed within an enterprise and administered centrally by those in IT with trusted credentials. This makes its policy creation and deployment conform to similar policy deployment processes and security restrictions.

AppLocker policies are distributed through known processes and by known means within the domain through Group Policy. But AppLocker policies can also be set on individual computers if the person has administrator privileges, and those policies might be contrary to the organization's written security policy. The enforcement settings for local policies are overridden by the same AppLocker policies in a Group Policy Object (GPO). However, because AppLocker rules are additive, a local policy that is not in a GPO will still be evaluated for that computer.

Microsoft does not provide a way to develop any extensions to AppLocker. The interfaces are not public. A user with administrator credentials can automate some AppLocker processes by using Windows PowerShell cmdlets. For information about the Windows PowerShell cmdlets for AppLocker, see the **AppLocker PowerShell Command Reference**.

AppLocker runs in the context of Administrator or LocalSystem, which is the highest privilege set. This security context has the potential of misuse. If a user with administrative credentials makes changes to an AppLocker policy on a local computer that is joined to a domain, those changes could be overwritten or disallowed by the GPO that contains the AppLocker rule for the same file (or path) that was changed on the local computer. However, because AppLocker rules are additive, a local policy that is not in a GPO will still be evaluated for that computer. If the local computer is not joined to a domain and is not administered by Group Policy, a person with administrative credentials can alter the AppLocker policy.

When securing files in a directory with a rule of the path condition type, whether using the allow or deny action on the rule, it is still necessary and good practice to restrict access to those files by setting the access control lists (ACLs) according to your security policy.

AppLocker does not protect against running 16-bit DOS binaries in the Virtual DOS Machine (NTVDM). This technology allows running legacy DOS and 16-bit Windows programs on computers that are using Intel 80386 or later when there is already another operating system running and controlling the hardware. The result is that 16-bit binaries can still run on Windows Server 2008 R2 and Windows 7 when AppLocker is configured to otherwise block binaries and libraries. If it is a requirement to prevent 16-bit applications from running, you must configure the Deny rule in the Executable rule collection for NTVDM.exe.

You cannot use AppLocker (or Software Restriction Policies) to prevent code from running outside the Win32 subsystem. In particular, this applies to the (POSIX) subsystem in Windows NT. If it is a requirement to prevent applications from running in the POSIX subsystem, you must disable the subsystem. For more information, see article 101270 in the Microsoft Knowledge Base.

AppLocker can only control VBScript, JScript, .bat files, .cmd files, and Windows PowerShell scripts. It does not control all interpreted code that runs within a host process, for example,

Perl scripts and macros. Interpreted code is a form of executable code that runs within a host process. For example, Windows batch files (*.bat) run within the context of the Windows Command Host (cmd.exe). To control interpreted code by using AppLocker, the host process must call AppLocker before it runs the interpreted code, and then enforce the decision returned by AppLocker. Not all host processes call into AppLocker and, therefore, AppLocker cannot control every kind of interpreted code, such as Microsoft Office macros.

Important

You should configure the appropriate security settings of these host processes if you must allow them to run. For example, configure the security settings in Microsoft Office to ensure that only signed and trusted macros are loaded.

AppLocker rules either allow or prevent an application from launching. AppLocker does not control the behavior of applications after they are launched. Applications could contain flags passed to functions that signal AppLocker to circumvent the rules and allow another .exe or .dll to be loaded. In practice, an application that is allowed by AppLocker could use these flags to bypass AppLocker rules and launch child processes. You must thoroughly examine each application before allowing them to run by using AppLocker rules.

Note

Two flags that illustrate this condition are `SANDBOX_INERT`, which can be passed to `CreateRestrictedToken`, and `LOAD_IGNORE_CODE_AUTHZ_LEVEL`, which can be passed to `LoadLibraryEx`. Both of these flags signal AppLocker to circumvent the rules and allow a child .exe or .dll to be loaded.

In Windows 7, administrators have the option of restricting standard users from bypassing AppLocker policies when logging into their local computers in safe mode. For more information about this setting, see article 977542 in the Microsoft Knowledge Base.

See Also

AppLocker Technical Reference

Tools to Use with AppLocker

This topic for the IT professional describes the tools available to create and administer AppLocker policies.

The following tools can help you administer the application control policies created by using AppLocker on the local computer or by using Group Policy. For information about the basic requirements for using AppLocker, see Requirements to Use AppLocker.

- **AppLocker Local Security Policy MMC snap-in**

  The AppLocker rules can be maintained by using the Local Security Policy snap-in (secpol.msc) of the Microsoft Management Console (MMC). For procedures to create, modify, and delete AppLocker rules, see Working with AppLocker Rules.

- **Generate Default Rules tool**

  AppLocker includes default rules for each rule collection accessed through the Local Security Policy snap-in. These rules are intended to help ensure that the files that are required for Windows to operate properly are allowed in an AppLocker rule collection. For information about how to use this tool, see Create AppLocker Default Rules.

- **Automatically Generate AppLocker Rules wizard**

  By using the Local Security Policy snap-in, you can automatically generate rules for all files within a folder. The wizard will scan the specified folder and create the condition types that you choose for each file in that folder. For information about how to use this wizard, see Run the Automatically Generate Rules Wizard.

- **Group Policy**

  You can edit an AppLocker policy by adding, changing, or removing rules by using the Group Policy Management Console (GPMC).

  If you want additional features to manage AppLocker policies, such as version control, use Group Policy management software that allows you to create versions of Group Policy Objects (GPOs). An example of this type of software is the Advanced Group Policy Management feature from the Microsoft Desktop Optimization Pack. For more information about Advanced Group Policy Management, see Advanced Group Policy Management Overview (http://go.microsoft.com/fwlink/?LinkId=145013).

- **Remote Server Administration Tools (RSAT)**

  You can use a computer with a supported operating system that has the Remote Server Administration Tools (RSAT) installed to create and maintain AppLocker policies. To download RSAT, see Remote Server Administration Tools for Windows 7 (http://go.microsoft.com/fwlink/?LinkID=153874) or Remote Server Administration Tools for Windows 8.

- **Event Viewer**

  The AppLocker log contains information about applications that are affected by AppLocker rules. For information about using Event Viewer to review the AppLocker logs, see Using Event Viewer with AppLocker, and View the AppLocker Log in Event Viewer.

- **AppLocker PowerShell cmdlets**

  The AppLocker Windows PowerShell cmdlets are designed to streamline the administration of AppLocker policy. They can be used to help create, test, maintain, and troubleshoot an AppLocker policy. The cmdlets are intended to be used in conjunction with the AppLocker user interface that is accessed through the Local Security Policy snap-in and the GPMC. For information about the cmdlets, see the AppLocker PowerShell Command Reference.

See Also

AppLocker Technical Reference

Using Event Viewer with AppLocker

This topic lists AppLocker events and describes how to use Event Viewer with AppLocker.

The AppLocker log contains information about applications that are affected by AppLocker rules. Each event in the log contains detailed information about:

- Which file is affected and the path of that file
- Which Packaged app is affected and the package identifier of the app
- Whether the file or Packaged app is allowed or blocked
- The rule type (path, file hash, or publisher)
- The rule name
- The security identifier (SID) for the user or group identified in the rule

Review the entries in the Event Viewer to determine if any applications are not included in the rules that you automatically generated. For instance, some line-of-business applications are installed to non-standard locations, such as the root of the active drive (for example: %SystemDrive%).

For information about what to look for in the AppLocker Event logs for Windows Server 2012 and Windows 8, see Monitor Application Usage with AppLocker.

To review the AppLocker log in Event Viewer

1. Open Event Viewer. To do this, click **Start**, type **eventvwr.msc** in the **Search programs and files** box, and then press ENTER. On Windows Server 2012 and Windows 8, on the **Start** screen, type **eventvwr.msc**.
2. In the console tree under **Application and Services Logs\Microsoft\Windows**, click **AppLocker**.

The following table contains information about the events that you can use to determine which applications are affected by AppLocker rules.

Event ID	Level	Event message	Description
8000	Error	Application Identity Policy conversion failed. Status *<%1>*	Indicates that the policy was not applied correctly to the computer. The status message is provided for troubleshooting purposes.
8001	Information	The AppLocker policy was applied successfully to this computer.	Indicates that the AppLocker policy was successfully applied to the computer.

Event ID	Level	Event message	Description
8002	Information	<*File name*> was allowed to run.	Specifies that the .exe or .dll file is allowed by an AppLocker rule.
8003	Warning	<*File name*> was allowed to run but would have been prevented from running if the AppLocker policy were enforced.	Applied only when the **Audit only** enforcement mode is enabled. Specifies that the .exe or .dll file would be blocked if the **Enforce rules** enforcement mode were enabled.
8004	Error	<*File name*> was not allowed to run.	Access to <*file name*> is restricted by the administrator. Applied only when the **Enforce rules** enforcement mode is set either directly or indirectly through Group Policy inheritance. The .exe or .dll file cannot run.
8005	Information	<*File name*> was allowed to run.	Specifies that the script or .msi file is allowed by an AppLocker rule.
8006	Warning	<*File name*> was allowed to run but would have been prevented from running if the AppLocker policy were enforced.	Applied only when the **Audit only** enforcement mode is enabled. Specifies that the script or .msi file would be blocked if the **Enforce rules** enforcement mode were enabled.
8007	Error	<*File name*> was not allowed to run.	Access to <*file name*> is restricted by the administrator. Applied only when the **Enforce rules** enforcement mode is set either directly or indirectly through Group Policy inheritance. The

Event ID	Level	Event message	Description
			script or .msi file cannot run.
8007	Error	AppLocker disabled on the SKU.	Added in Windows Server 2012 and Windows 8.
8020	Information	Packaged app allowed.	Added in Windows Server 2012 and Windows 8.
8021	Information	Packaged app audited.	Added in Windows Server 2012 and Windows 8.
8022	Information	Packaged app disabled.	Added in Windows Server 2012 and Windows 8.
8023	Information	Packaged app installation allowed.	Added in Windows Server 2012 and Windows 8.
8024	Information	Packaged app installation audited.	Added in Windows Server 2012 and Windows 8.
8025	Warning	Packaged app installation disabled.	Added in Windows Server 2012 and Windows 8.
8027	Warning	No Packaged app rule configured.	Added in Windows Server 2012 and Windows 8.

See Also

Tools to Use with AppLocker

AppLocker Settings

This topic for the IT professional lists the settings used by AppLocker.

The following table describes the settings and values used by AppLocker.

Setting	Value
Registry path	Policies are stored in **\HKEY_LOCAL_Machine\Software\Policies\Microsoft\Windows\SrpV2**
Firewall ports	Not applicable
Security policies	Custom created, no default
Group Policy settings	Custom created, no default
Network ports	Not applicable
Service accounts	Not applicable
Performance counters	Not applicable

See Also

AppLocker Technical Reference

BitLocker Overview

This topic provides a high-level overview of BitLocker including a list of new and changed features, system requirements, practical applications, and deprecated features. It provides links to additional content that will help you to learn more about working with BitLocker.

Did you mean…

- AppLocker
- EFS
- BitLocker Drive Encryption Provider

Feature description

BitLocker Drive Encryption is a data protection feature of the operating system that was first made available in Windows Vista. Subsequent operating system releases have continued to improve the security offered by BitLocker protection to allow the operating system to provide BitLocker protection to more drives and devices. Having BitLocker integrated with the operating system addresses the threats of data theft or exposure from lost, stolen, or inappropriately decommissioned computers. Manage-bde is the command-line tool that can also be used to perform tasks on the computer related to BitLocker. When installing the BitLocker optional component on a server you will also need to install the Enhanced Storage feature. This feature is used to support hardware encrypted drives. On servers there is also an additional BitLocker feature that can be installed, BitLocker Network Unlock. Computers running Windows RT,

Windows RT 8.1 or Windows 8.1 can be protected using Device Encryption, which is a customized version of BitLocker for

BitLocker provides the most protection when used with a Trusted Platform Module (TPM) version 1.2 or later. The TPM is a hardware component installed in many newer computers by the computer manufacturers. It works with BitLocker to help protect user data and to ensure that a computer has not been tampered with while the system was offline.

On computers that do not have a TPM version 1.2 or later, you can still use BitLocker to encrypt the Windows operating system drive. However, this implementation will require the user to insert a USB startup key to start the computer or resume from hibernation. In Windows 8 using an operating system volume password is another option to protect the operating system volume on a computer without TPM. Both options do not provide the pre-startup system integrity verification offered by BitLocker with a TPM. In addition to the TPM, BitLocker offers the option to lock the normal startup process until the user supplies a personal identification number (PIN) or inserts a removable device, such as a USB flash drive, that contains a startup key. These additional security measures provide multifactor authentication and assurance that the computer will not start or resume from hibernation until the correct PIN or startup key is presented.

Practical applications

Data on a lost or stolen computer is vulnerable to unauthorized access, either by running a software-attack tool against it or by transferring the computer's hard disk to a different computer. BitLocker helps mitigate unauthorized data access by enhancing file and system protections. BitLocker also helps render data inaccessible when BitLocker-protected computers are decommissioned or recycled.

There are two additional tools in the Remote Server Administration Tools, which you can use to manage BitLocker.

- **BitLocker Recovery Password Viewer**. The BitLocker Recovery Password Viewer enables you to locate and view BitLocker Drive Encryption recovery passwords that have been backed up to Active Directory Domain Services (AD DS). You can use this tool to help recover data that is stored on a drive that has been encrypted by using BitLocker. The BitLocker Recovery Password Viewer tool is an extension for the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in.

  By using this tool, you can examine a computer object's **Properties** dialog box to view the corresponding BitLocker recovery passwords. Additionally, you can right-click a domain container and then search for a BitLocker recovery password across all the domains in the Active Directory forest. To view recovery passwords, you must be a domain administrator, or you must have been delegated permissions by a domain administrator.
- **BitLocker Drive Encryption Tools**. BitLocker Drive Encryption Tools include the command-line tools manage-bde and repair-bde and the BitLocker cmdlets for Windows PowerShell. Both manage-bde and the BitLocker cmdlets can be used to perform any task that can be accomplished through the BitLocker control panel, and they are appropriate to use for automated deployments and other scripting scenarios. Repair-bde is provided for disaster

recovery scenarios in which a BitLocker protected drive cannot be unlocked normally or by using the recovery console.

New and changed functionality

The following table identifies the new and changed functionality for BitLocker in Windows 8 and Windows Server 2012. Check out **What's New in BitLocker for Windows 8.1 and Windows Server 2012 R2**.

Feature/functionality	Windows 7	Windows 8 and Windows Server 2012
Reset the BitLocker PIN or password	Administrator privileges are required to reset the BitLocker PIN on an operating system drive and the password on a fixed or removable data drive.	Standard users can reset the BitLocker PIN and password on operating system drives, fixed data drives, and removable data drives.
Disk encryption	The entire disk is encrypted when BitLocker is enabled.	You can choose to encrypt the entire disk or just the used space on the disk when BitLocker is enabled. As disk space is consumed the disk will be encrypted.
Hardware Encrypted Drive support	Not natively supported by BitLocker.	BitLocker can support hard drives with the Windows logo that come pre-encrypted from the manufacturer.
Unlocking using a network based key to provide dual-factor authentication	Not available. Physical presence at the computer was required to have dual-factor authentication.	A new type of key protector allows for a special network key to be used to unlock and skip the PIN entry prompt in situations where computers are rebooted on trusted wired networks. This allows for remote maintenance on computers that are protected by using a PIN during non-work hours, and it provides dual-factor authentication without requiring physical presence at the computer, while still ensuring user

Feature/functionality	Windows 7	Windows 8 and Windows Server 2012
		authentication is required when the computer is not connected to the trusted network.
Protection for clusters	Not available.	Windows Server 2012 includes BitLocker support for Windows Cluster Shared Volumes and Windows Failover Clusters that are running in domain established by a Windows Server 2012 domain controller that has enabled the Kerberos Key Distribution Center Service.
Linking a BitLocker key protector to an Active Directory account	Not available.	Allows a BitLocker key protector to be tied to a user, group, or computer account in Active Directory. That key protector can be used to unlock BitLocker protected data volumes when a user is signed in with the correct credentials or signed on to a computer with the correct credentials.

Removed or deprecated functionality

The Diffuser option is no longer available to be added to the Advanced Encryption Standard (AES) encryption algorithm.

The "Configure TPM validation profile" Group Policy setting is deprecated in Windows 8 and Windows Server 2012. It has been replaced with system specific policies for BIOS-based and UEFI-based computers.

The **–tpm** option is no longer supported by manage-bde.

System requirements

BitLocker has the following hardware requirements:

For BitLocker to use the system integrity check provided by a Trusted Platform Module (TPM), the computer must have TPM 1.2 or TPM 2.0. If your computer does not have a TPM, enabling BitLocker requires that you save a startup key on a removable device, such as a USB flash drive.

A computer with a TPM must also have a Trusted Computing Group (TCG)-compliant BIOS or UEFI firmware. The BIOS or UEFI firmware establishes a chain of trust for the pre-operating system startup, and it must include support for TCG-specified Static Root of Trust Measurement. A computer without a TPM does not require TCG-compliant firmware.

The system BIOS or UEFI firmware (for TPM and non-TPM computers) must support the USB mass storage device class, including reading small files on a USB flash drive in the pre-operating system environment. For more information about USB, see the USB Mass Storage Bulk-Only and the Mass Storage UFI Command specifications on the USB Web site.

The hard disk must be partitioned with at least two drives:

- The operating system drive (or boot drive) contains the operating system and its support files. It must be formatted with the NTFS file system.
- The system drive contains the files that are needed to load Windows after the firmware has prepared the system hardware. BitLocker is not enabled on this drive. For BitLocker to work, the system drive must not be encrypted, must differ from the operating system drive, and must be formatted with the FAT32 file system on computers that use UEFI-based firmware or with the NTFS file system on computers that use BIOS firmware. We recommend that system drive be approximately 350 MB in size. After BitLocker is turned on it should have approximately 250 MB of free space.

When installed on a new computer, Windows will automatically create the partitions that are required for BitLocker.

See also

- BitLocker Frequently Asked Questions
- Prepare your organization for BitLocker: Planning and Policies
- BitLocker Basic Deployment
- BitLocker Group Policy Settings
- BCD Settings and BitLocker
- BitLocker Recovery Guide
- BitLocker: How to enable Network Unlock
- BitLocker: How to deploy on Windows Server 2012
- Protecting cluster shared volumes and storage area networks with BitLocker
- BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker
- BitLocker: Use BitLocker Recovery Password Viewer
- Encrypted Hard Drive

- BitLocker Cmdlets in Windows PowerShell
- TechNet Wiki ArticlesTechNet Wiki Articles

What's New in BitLocker for Windows 8 and Windows Server 2012

The following is a list of new functionality in BitLocker for Windows 8 and Windows Server 2012:

- New functionality in Windows 8 and Windows Server 2012:
  - BitLocker provisioning

    BitLocker can be used to deploy drives to an encrypted state during installation prior to calling setup.
  - Used Disk Space Only encryption

    BitLocker now offers two encryption methods, Used Disk Space Only and Full volume encryption. Used Disk Space Only allows for a much quicker encryption experience by only encrypting used blocks on the targeted volume.
  - Standard User PIN and password change

    Allows a standard user to change the BitLocker PIN or password on operating system volumes and the BitLocker password on data volumes, reducing internal help desk call volume.
  - Network Unlock

    Enables a BitLocker system on a wired network to automatically unlock the system volume during boot (on capable Windows Server 2012 networks), reducing internal help desk call volumes for lost PINs.
  - Support for Encrypted Hard Drives for Windows

    BitLocker support for Encrypted Hard Drives provides users a familiar method for managing drive encryption along with the benefit of using hardware-based encryption.

BitLocker provisioning

In Windows Vista and Windows 7, BitLocker is provisioned post installation for system and data volumes through either the manage-bde command line interface or the Control Panel user interface. In Windows 8 and Windows 8.1, BitLocker can also be easily provisioned before the operating system is installed.

With this improvement administrators can enable BitLocker from the Windows Preinstallation Environment (WinPE) prior to operating system deployment. This is done with a randomly generated clear protector applied to the formatted volume and encrypting the volume prior to running the Windows setup process. If the encryption uses the Used Disk Space Only option described in the next section, this step takes only a few seconds and so incorporates well into regular deployment processes.

To check the BitLocker status of a particular volume, administrators can look at the status of the drive in the BitLocker control panel applet or Windows Explorer. When a drive is pre-provisioned for BitLocker, a status of "Waiting For Activation" is displayed with a yellow exclamation icon in the BitLocker Control Panel. This status means that there was only a clear protector used when encrypting the volume. In this case, the volume is not protected and needs to have a secure key added to the volume before the drive is considered fully protected. You can use the control panel, manage-bde tool or WMI APIs to add an appropriate key protector and the volume status will be updated. The following table shows the appropriate key protectors that can be added to drives that have been pre-provisioned with BitLocker protection:

Drive Type	Key protector
Operating System	TPM TPM+PIN Startup Key (for systems without a TPM) Password (for systems without a TPM)
Fixed data drive	Automatic unlock Password Smart card
Removable data drive	Password Smart card

Used Disk Space Only encryption

In Windows 7, BitLocker requires that all data and free space on the drive are encrypted. The encryption process can take a very long time on larger volumes. In Windows 8 and Windows 8.1, administrators can choose to encrypt the entire volume or the used space only. When you choose the Used Disk Space Only encryption option, only the portion of the drive that has data will be encrypted. Free disk space will not be encrypted. Used Disk Space Only encryption allows encryption to complete much faster on empty or partly empty drives than previous implementations of BitLocker. When provisioning BitLocker during Windows deployments, Used Disk Space Only encryption allows BitLocker to encrypt a drive in a short amount of time before installing the operating system. Full Encryption encrypts both data and free space on the volume, similar to the way BitLocker works in Windows 7 and Windows Vista.

New Group Policy settings for encryption type

You can use Group Policy settings to enforce that either Used Disk Space Only or Full Encryption is used when BitLocker is enabled on a drive. Group Policy settings for BitLocker Drive Encryption are located under the **\Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption** path of the Local Group Policy Editor

The following new Group Policies are available:

- **Fixed Data Drives\Enforce drive encryption type on fixed data drives**
- **Operating System Drives\Enforce drive encryption type on operating system drives**
- **Removable Data Drives\Enforce drive encryption type on removable data drives**

For each of these policies, once they are enabled you can then specify which type of encryption is required to be used on which drive type. If the policy is not configured the user will be able to choose the encryption method when they turn on BitLocker.

Standard User PIN and password change

Administrative privileges are required to configure BitLocker for operating system drives. In an organization where computers are managed by IT professionals and users are not normally granted administrative privileges, deploying the TPM + PIN protection option to large numbers of computers can be challenging. In Windows 8, administrative privileges are still required to configure BitLocker, however standard users are allowed to change the BitLocker PIN or password for the operating system volume or the BitLocker password for fixed data volumes by default. This gives users the ability to choose PINs and passwords that correspond to a personal mnemonic instead of requiring the user remember a randomly generated character set and allows IT professionals to use the same initial PIN or password setting for all computer images. This also presents the opportunity for users to choose passwords and PINs that are more susceptible to password guessing, dictionary attacks, and social engineering attacks and gives users the ability unlock any computer that still uses the original PIN or password assignment. Requiring password complexity and PIN complexity by Group Policy is recommended to help ensure that users take appropriate care when setting passwords and PINs.

Standard users are required to enter the current PIN or password for the drive to change the BitLocker PIN or BitLocker password. If a user enters an incorrect current PIN or password, the default tolerance for retry attempts is set to 5. Once the retry limit is reached, a standard user will not be able to change the BitLocker PIN or BitLocker password. The retry counter is set to zero when the computer is restarted or when an administrator resets the BitLocker PIN or BitLocker password.

You can disable the option to allow standard users to change PINs and passwords using the Group Policy setting **Disallow standard users from changing the PIN** located in the **\Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives** section of Local Group Policy Editor.

Network Unlock

Windows Server 2012 has added a new BitLocker protector option for Operating System Volumes called Network Unlock. Network Unlock will enable easier management for BitLocker enabled desktops and servers in a domain environment by providing automatic unlock of operating system volumes at system reboot when connected to a trusted wired corporate

network. This feature requires the client hardware to have a DHCP driver implemented in its UEFI firmware.

Operating system volumes protected by TPM+PIN protectors require a PIN to be entered when a machine reboots or resumes from hibernation (for example, when configured for Wake on LAN). The requirement to enter a PIN can make it difficult for enterprises to install software patches to unattended desktops and servers. Network Unlock provides a method by which computers that are configured to use a TPM+PIN key protector can start Windows without user intervention. Network Unlock works in a similar fashion to the TPM+StartupKey. Rather than needing to read the StartupKey from USB media, however, the key for Network Unlock is composed from a key stored in the TPM and an encrypted network key that is sent to the server, decrypted and returned to the client in a secure session. The network key is stored on the system drive along with an AES 256 session key, and encrypted with the 2048-bit RSA public key of the unlock server's certificate. The network key is decrypted with the help of a provider on a Windows Server 2012 WDS server and returned encrypted with its corresponding session key. In instances where the Network Unlock provider is unavailable, the standard TPM+PIN unlock screen is presented to unlock the drive. The server side configuration to enable Network Unlock also requires provisioning a 2048 bit RSA public/private key pair in the form of an X.509 certificate, and for the public key certificate to be distributed to the clients. This certificate must be managed and deployed through the Group Policy Management Console directly on Windows Server 2012 Domain Controller. For more information see BitLocker: How to enable Network Unlock

Support for Encrypted Hard Drives for Windows

BitLocker provides Full Volume Encryption (FVE) of Windows operating system and data volumes using software-based encryption. In Windows 8 BitLocker also provide support for a new enhanced storage device type, the Encrypted Hard Drive, that is becoming a more common option in new servers and computers. Encrypted Hard Drives offer Full Disk Encryption (FDE), which means encryption occurs on each block of the physical drive. Encryption operations are more efficient on Encrypted Hard Drives because the encryption process is offloaded to the storage controller on the drive (also known as hardware-based encryption).

Windows 8 supports Encrypted Hard Drives natively in the operating system through the following mechanisms:

- Identification: Windows 8 will be able to identify that the drive is a Encrypted Hard Drive device type
- Activation: Windows 8 disk management will activate, create and map volumes to ranges/bands as appropriate
- Configuration: Windows 8 will create and map volumes to ranges/bands as appropriate
- API: Windows 8 provides API support for applications to manage Encrypted Hard Drives independently of BitLocker Drive Encryption
- BitLocker: BitLocker Control Panel will enable users to manage Encrypted Hard Drives in the same manner as full volume encrypted drives.

For more information about the system requirements and usage see Encrypted Hard Drive.

Additional information

Prepare your organization for BitLocker: Planning and Policies

BitLocker Group Policy Settings

BCD Settings and BitLocker

BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker

BitLocker: Use BitLocker Recovery Password Viewer

BitLocker: How to deploy on Windows Server 2012

BitLocker Frequently Asked Questions (FAQ)

BitLocker Drive Encryption is a data protection feature available in Windows 8 Pro, Windows 8 Enterprise, and in all editions of Windows Server 2012. This topic includes frequently asked questions about BitLocker in Windows 8.

- Overview and requirements
  - What is BitLocker? How does it work?
  - Does BitLocker support multifactor authentication?
  - What are the BitLocker hardware and software requirements?
  - Why are two partitions required? Why does the system drive have to be so large?
  - Which Trusted Platform Modules (TPMs) does BitLocker support?
  - How can I tell if a TPM is on my computer?
  - Can I use BitLocker on an operating system drive without a TPM?
  - How do I obtain BIOS support for the TPM on my computer?
  - What credentials are required to use BitLocker?
  - What is the recommended boot order for computers that are going to be BitLocker-protected?
- Upgrading
  - Can I upgrade my Windows 7–based computer to Windows 8 with BitLocker enabled?
  - What is the difference between suspending and decrypting BitLocker?
  - Do I have to decrypt my BitLocker-protected drive to download and install system updates and upgrades?
- Deployment and administration
  - Can BitLocker deployment be automated in an enterprise environment?
  - Can BitLocker encrypt more than just the operating system drive?
  - Is there a noticeable performance impact when BitLocker is enabled on a Windows 8–based computer?
  - Approximately how long will initial encryption take when BitLocker is turned on?
  - What happens if the computer is turned off during encryption or decryption?

    - Does BitLocker encrypt and decrypt the entire drive all at once when reading and writing data?
    - How can I prevent users on a network from storing data on an unencrypted drive?
    - What system changes would cause the integrity check on my operating system drive to fail?
    - What causes BitLocker to start into recovery mode when attempting to start the operating system drive?
    - Can I swap hard disks on the same computer if BitLocker is enabled on the operating system drive?
    - Can I access my BitLocker-protected drive if I insert the hard disk into a different computer?
    - Why is "Turn BitLocker on" not available when I right-click a drive?
    - What type of disk configurations are supported by BitLocker?
    - What if my disk configuration is not listed?
- Key management
    - What is the difference between a TPM owner password, recovery password, recovery key, password, PIN, enhanced PIN, and startup key?
    - How can the recovery password and recovery key be stored?
    - Is it possible to add an additional method of authentication without decrypting the drive if I only have the TPM authentication method enabled?
    - If I lose my recovery information, will the BitLocker-protected data be unrecoverable?
    - Can the USB flash drive that is used as the startup key also be used to store the recovery key?
    - Can I save the startup key on multiple USB flash drives?
    - Can I save multiple (different) startup keys on the same USB flash drive?
    - Can I generate multiple (different) startup keys for the same computer?
    - Can I generate multiple PIN combinations?
    - What encryption keys are used in BitLocker? How do they work together?
    - Where are the encryption keys stored?
    - Why do I have to use the function keys to enter the PIN or the 48-character recovery password?
    - How does BitLocker help prevent an attacker from discovering the PIN that unlocks my operating system drive?
    - How can I determine the manufacturer of my TPM?
    - How can I evaluate a TPM's dictionary attack mitigation mechanism?
    - Can PIN length and complexity be managed with Group Policy?
- BitLocker To Go
    - What is BitLocker To Go?
    - How can I authenticate or unlock my removable data drive?

    - What happens if I try to open a BitLocker-protected, NTFS-formatted removable drive by using a computer running Windows XP or Windows Vista?
    - Is there a way to ensure the BitLocker To Go Reader is not installed on FAT-formatted drives?
    - Can I save files to my BitLocker-protected removable drive when I am using Windows XP or Windows Vista?
    - Can I download a copy of the BitLocker To Go Reader?
    - Why am I unable to access my removable drive on computers running Windows XP or Windows Vista when using the BitLocker To Go Reader?
- Active Directory Domain Services (AD DS)
    - Does BitLocker require a schema extension to store recovery information in AD DS?
    - What type of information is stored in AD DS?
    - Does BitLocker encrypt recovery information as it is sent to AD DS?
    - Is the BitLocker recovery information stored in plaintext in AD DS?
    - What if BitLocker is enabled on a computer before the computer has joined the domain?
    - Is there an event log entry recorded on the client computer to indicate the success or failure of the Active Directory backup?
    - If I change the BitLocker recovery password on my computer and store the new password in AD DS, will AD DS overwrite the old password?
    - What happens if the backup initially fails? Will BitLocker retry the backup?
- Security
    - What form of encryption does BitLocker use? Is it configurable?
    - What is the best practice for using BitLocker on an operating system drive?
    - What are the implications of using the sleep or hibernate power management options?
    - What are the advantages of a TPM?
    - Is Microsoft pursuing any security certification for BitLocker?
- BitLocker Network Unlock
    - What is BitLocker Network Unlock?
    - What is required to use BitLocker Network Unlock?
    - How is BitLocker Network Unlock different from automatic unlock?
    - How do I unlock my computer when not connected to the network?
- Other questions
    - Can I use EFS with BitLocker?
    - Can I run a kernel debugger with BitLocker?
    - How does BitLocker handle memory dumps?
    - Can BitLocker support smart cards for pre-boot authentication?
    - Can I use a non-Microsoft TPM driver?
    - Can other tools that manage or modify the master boot record work with BitLocker?
    - Why is the system check failing when I am encrypting my operating system drive?

- What can I do if the recovery key on my USB flash drive cannot be read?
- Why am I unable to save my recovery key to my USB flash drive?
- Why am I unable to automatically unlock my drive?
- Can I use BitLocker in Safe Mode?
- How do I "lock" a data drive?
- Can I use BitLocker with the Volume Shadow Copy Service?
- Does BitLocker support virtual hard disks (VHDs)?

Overview and requirements

What is BitLocker? How does it work?

BitLocker encrypts the hard drives on your computer to provide enhanced protection against data theft or exposure on computers and removable drives that are lost or stolen, and more secure data deletion when BitLocker-protected computers are decommissioned as it is much more difficult to recover deleted data from an encrypted drive than from a non-encrypted drive.

How BitLocker works with operating system drives

Data on a lost or stolen computer is vulnerable to unauthorized access, either by running a software attack tool against it or by transferring the computer's hard disk to a different computer. BitLocker helps mitigate unauthorized data access on lost or stolen computers by:

- **Encrypting the entire Windows operating system drive on the hard disk.** BitLocker encrypts all user files and system files on the operating system drive, including the swap files and hibernation files.
- **Checking the integrity of early boot components and boot configuration data.** On computers that have a Trusted Platform Module (TPM) version 1.2 or 2.0, BitLocker uses the enhanced security capabilities of the TPM to help ensure that your data is accessible only if the computer's boot components appear unaltered and the encrypted disk is located in the original computer.

BitLocker is integrated into Windows 8 and provides enterprises with enhanced data protection that is easy to manage and configure. For example, BitLocker can use an existing Active Directory Domain Services (AD DS) infrastructure to remotely store BitLocker recovery keys.

How BitLocker works with fixed and removable data drives

BitLocker can also be used to protect fixed and removable data drives. When used with data drives, BitLocker encrypts the entire contents of the drive and can be configured by using Group Policy to require that BitLocker be enabled on a drive before the computer can write data to the drive. BitLocker can be configured with the following unlock methods for data drives:

- **Automatic unlock.** Fixed data drives can be set to automatically unlock on a computer where the operating system drive is encrypted. Removable data drives can be set to automatically unlock on a computer running Windows 8 after the password or smart card is initially used to unlock the drive. However, removable data drives must always have either a password or smart card unlock method in addition to the automatic unlock method.

- **Password.** When users attempt to open a drive, they are prompted to enter their password before the drive will be unlocked. This method can be used with the BitLocker To Go Reader on computers running Windows Vista or Windows XP, to open BitLocker-protected drives as read-only.
- **Smart card.** When users attempt to open a drive, they are prompted to insert their smart card before the drive will be unlocked.
- **Active Directory Account or Group.** A key can be assigned to an Active Directory user, group, or computer account and when those credentials are presented the drive will be unlocked. Using this key protector requires using manage-mde or the Windows PowerShell cmdlets for BitLocker to manually add the **ADAccountorGroup** key protector. For more information on the cmdlet, see Add-BitLockerKeyProtector. For more information on the command-line syntax, see Manage-bde: protectors.

A drive can support multiple unlock methods. For example, a removable data drive can be configured to be automatically unlocked on your primary work computer but query you for a password if used with another computer.

Does BitLocker support multifactor authentication?

Yes, BitLocker supports multifactor authentication for operating system drives. If you enable BitLocker on a computer that has a TPM version 1.2 or 2.0, you can use additional forms of authentication with the TPM protection. BitLocker offers the option to lock the normal boot process until the user supplies a password (which can be either a personal identification number (PIN), a passphrase, or a password) or inserts a USB device (such as a flash drive) that contains a BitLocker startup key, or both the password and the USB device can be required. These additional security measures provide multifactor authentication and help ensure that the computer will not start or resume from hibernation until the correct authentication method is presented.

Note

Use of both the USB and password along with the TPM must be configured by using the Manage-bde command-line tool. This protection method cannot be specified by using the BitLocker setup wizard.

What are the BitLocker hardware and software requirements?

To use all BitLocker features, your computer must meet the hardware and software requirements listed in the following table.

Note

Dynamic disks are not supported by BitLocker. Dynamic data volumes will not be displayed in the Control Panel. Although the operating system volume will always be displayed in the Control Panel, regardless of whether it is a Dynamic disk, if it is a dynamic disk it is cannot be protected by BitLocker. For a list of supported disk configurations, see What type of disk configurations are supported by BitLocker?

BitLocker hardware and software requirements for operating system drives

Requirement	Description
Hardware configuration	The computer must meet the minimum requirements for Windows 8. For more information about Windows 8 requirements, see the Windows 8 Web site.
Operating system	Windows 8 or Windows Server 2012 **Note** BitLocker is an optional feature of Windows Server 2012. Use Server Manager to install BitLocker on a computer running Windows Server 2012.
Hardware TPM	TPM version 1.2 or 2.0 A TPM is not required for BitLocker; however, only a computer with a TPM can provide the additional security of pre-startup system integrity verification and multifactor authentication.
BIOS configuration	• A Trusted Computing Group (TCG)-compliant BIOS or UEFI firmware. • The boot order must be set to start first from the hard disk, and not the USB or CD drives. • The firmware must be able to read from a USB flash drive during startup.
File system	For computers that boot natively with UEFI firmware at least one FAT32 partition for the system drive and one NTFS partition for the operating system drive. For computers with legacy BIOS firmware, at least two NTFS disk partitions, one for the system drive and one for the operating system drive. For either firmware, the system drive partition must be at least 350 megabytes (MB) and set as the active partition.
Hardware encrypted drive prerequisites	To use a hardware encrypted drive as the boot drive, the drive must be in the uninitialized state

Requirement	Description
(optional)	and in the security inactive state. In addition, the system must always boot with native UEFI version 2.3.1 or higher and the CSM (if any) disabled.

BitLocker hardware and software requirements for data drives

Requirement	Description
File system	For a fixed or removable data drive to be BitLocker-protected, it must be formatted by using the exFAT, FAT16, FAT32, or NTFS file system. **Note** To use the BitLocker To Go Reader to read data on a removable data drive, the drive must be formatted by using the exFAT, FAT16, or FAT32 file system. If the drive is NTFS formatted, it can only be unlocked on a computer running Windows Server 2008 R2 or Windows 7 or later. Prior versions of the Windows operating system will not recognize the drive and will prompt you to format the drive.
Drive size	The drive must be least 64 MB in size.

Why are two partitions required? Why does the system drive have to be so large?

Two partitions are required to run BitLocker because pre-startup authentication and system integrity verification must occur on a separate partition from the encrypted operating system drive. This configuration helps protect the operating system and the information in the encrypted drive. The system drive may also be used to store the Windows Recovery Environment (Windows RE) and other files that may be specific to setup or upgrade programs. Computer manufacturers and enterprise customers can also store system tools or other recovery tools on this drive, which will increase the required size of the system drive. For example, using the system drive to store Windows RE along with the BitLocker startup file will increase the size of the system drive to 350 MB. The system drive is hidden by default and is not assigned a drive letter. The system drive is created automatically when Windows 8 is installed.

Which Trusted Platform Modules (TPMs) does BitLocker support?

BitLocker supports TPM version 1.2 and 2.0. BitLocker does not support previous versions of TPMs. Version 1.2 and later TPMs provide increased standardization, security enhancement, and improved functionality over previous versions. In addition, you must use a Microsoft-provided TPM driver. To check the TPM driver provider, type **Windows logo key + R**, type **devmgmt.msc** in the **Open** box, and then press ENTER to open Device Manager. Expand **Security Devices**, right-click the TPM, and click **Properties**. Click the **Driver** tab, and verify that the **Driver Provider** field displays **Microsoft**.

Important

When using BitLocker with a TPM, it is recommended that BitLocker be turned on immediately after the computer has been restarted. If the computer has resumed from sleep prior to turning on BitLocker, the TPM may incorrectly measure the pre-boot components on the computer. In this situation, when the user subsequently attempts to unlock the computer, the TPM verification check will fail and the computer will enter BitLocker recovery mode and prompt the user to provide recovery information before unlocking the drive.

How can I tell if a TPM is on my computer?

Type **Windows logo key + Q** to open **Apps**. In the Search panel, click **Settings** , type **BitLocker**, and then click **Turn on BitLocker**. If your computer does not have a compatible TPM or the BIOS is not compatible with the TPM, you will receive the following error message informing you that a TPM was not found.

If you receive this error message on a computer that has a TPM, check if either of the following situations applies to your computer:

- Some computers have TPMs that do not appear in the Windows 8 TPM Microsoft Management Console snap-in (tpm.msc) due to a BIOS or UEFI setting that hides the TPM by default and does not make the TPM available unless it is first enabled in the BIOS or UEFI firmware. If your TPM might be hidden in the BIOS or UEFI, consult the manufacturer's documentation for instructions to display or enable the TPM.
- Some computers might have an earlier version of the TPM or an earlier version of the system BIOS that is not compatible with BitLocker. Contact the computer manufacturer to verify that the computer has a TPM version 1.2 or to get a BIOS update.

Can I use BitLocker on an operating system drive without a TPM?

Yes, you can enable BitLocker on an operating system drive without a TPM version 1.2 or 2.0, if the BIOS or UEFI firmware has the ability to read from a USB flash drive in the boot environment. This is because BitLocker will not unlock the protected drive until BitLocker's own volume master key is first released by either the computer's TPM or by a USB flash drive containing the

BitLocker startup key for that computer. However, computers without TPMs will not be able to use the system integrity verification that BitLocker can also provide.

To help determine whether a computer can read from a USB device during the boot process, use the BitLocker system check as part of the BitLocker setup process. This system check performs tests to confirm that the computer can properly read from the USB devices at the appropriate time and that the computer meets other BitLocker requirements.

To enable BitLocker on a computer without a TPM, you must enable the **Require additional authentication at setup** Group Policy setting, which is located in **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives**. You must select the **Allow BitLocker without a compatible TPM** check box. After this setting is applied to the local computer, the non-TPM settings appear in the BitLocker setup wizard.

How do I obtain BIOS support for the TPM on my computer?

Contact the computer manufacturer to request a Trusted Computing Group (TCG)-compliant BIOS or UEFI boot firmware that meets the following requirements:

- It is compatible with Windows 8 and has passed the Windows 8 logo tests.
- It is compliant with the TCG standards for a client computer.
- It has a secure update mechanism to help prevent a malicious BIOS or boot firmware from being installed on the computer.

What credentials are required to use BitLocker?

To turn on, turn off, or change configurations of BitLocker on operating system and fixed data drives, membership in the local **Administrators** group is required. Standard users can turn on, turn off, or change configurations of BitLocker on removable data drives. Disable the **Control use of BitLocker on removable drives** policy setting (located in **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives**) to restrict standard users from turning on or turning off BitLocker on removable data drives. In Windows 8 and Windows Server 2012, standard users can also change the PIN or password on operating system drives and fixed data drives after they provide the current PIN or password.

What is the recommended boot order for computers that are going to be BitLocker-protected?

You should configure the startup options of your computer to have the hard disk drive first in the boot order, before any other drives such ach as CD/DVD drives or USB drives. If the hard disk is not first and you typically boot from hard disk, then a boot order change may be detected or assumed when removable media is found during boot. The boot order typically affects the system measurement that is verified by BitLocker and a change in boot order will cause you to be prompted for your BitLocker recovery key. For the same reason, if you have a laptop with a

docking station, ensure that the hard disk drive is first in the boot order both when docked and undocked.

Upgrading

Can I upgrade my Windows 7–based computer to Windows 8 with BitLocker enabled?

Yes. To upgrade from Windows 7 to Windows 8 without decrypting the operating system drive, open the **BitLocker Drive Encryption** Control Panel item in Windows 7, click **Manage BitLocker**, and then and click **Suspend**. Suspending protection does not decrypt the drive; it disables the authentication mechanisms used by BitLocker and uses a clear key on the drive to enable access. Proceed with the upgrade process by using your Windows 8 DVD. After the upgrade has completed, open Windows Explorer, right-click the drive, and then click **Resume Protection**. This reapplies the BitLocker authentication methods and deletes the clear key.

What is the difference between suspending and decrypting BitLocker?

Decrypt completely removes BitLocker protection and fully decrypts the drive.

When BitLocker is suspended, BitLocker keeps the data encrypted but encrypts the BitLocker volume master key with a clear key. The clear key is a cryptographic key stored unencrypted and unprotected on the disk drive. By storing this key unencrypted, the **Suspend** option allows for changes or upgrades to the computer without the time and cost of decrypting and re-encrypting the entire drive. After the changes are made and BitLocker is again enabled, BitLocker will reseal the encryption key to the new values of the measured components that changed as a part of the upgrade, the volume master key is changed, the protectors are updated to match and the clear key is erased.

Do I have to decrypt my BitLocker-protected drive to download and install system updates and upgrades?

Operating system upgrades from Windows Anytime Upgrade require that the operating system drive be decrypted prior to installation. If you upgrade from Windows 7 to Windows 8 or install other non-Microsoft updates, you might need to disable or suspend BitLocker so that a new measurement of the system can be taken after the upgrade or update has been applied. Software and operating system updates from Microsoft Update do not require drive decryption or that you disable or suspend BitLocker.

Please refer to the following table to determine whether you must disable or suspend BitLocker or decrypt your drive before you perform an upgrade or update installation.

Type of update	Action
Windows Anytime Upgrade	Decrypt
Upgrade from Windows 7 to Windows 8	Suspend
Non-Microsoft software updates, such as: • Computer manufacturer firmware updates • TPM firmware updates • Non-Microsoft application updates that modify boot components	Suspend

If you suspended BitLocker, you can resume BitLocker protection after you have installed the upgrade or update. Upon resuming protection, BitLocker will reseal the encryption key to the new values of the measured components that changed as a part of the upgrade or update. If these types of upgrades or updates are applied without suspending BitLocker, your computer will enter recovery mode when restarting and will require a recovery key or password to access the computer.

Deployment and administration

Can BitLocker deployment be automated in an enterprise environment?

Yes, you can automate the deployment and configuration of BitLocker and the TPM using either WMI or Windows PowerShell scripts. How you choose to implement the scripts depends on your environment. You can also use the BitLocker command-line tool, Manage-bde.exe, to locally or remotely configure BitLocker. For additional information about writing scripts that use the BitLocker WMI providers, see the MSDN topic BitLocker Drive Encryption Provider. For additional information about using Windows PowerShell cmdlets with BitLocker Drive Encryption see BitLocker Cmdlets in Windows PowerShell.

Can BitLocker encrypt more than just the operating system drive?

Yes. In Windows Vista, BitLocker could only encrypt operating system drives. Windows Vista SP1 and Windows Server 2008 added support for encrypting fixed data drives. In Windows 8, Windows Server 2012, Windows 7, and Windows Server 2008 R2, BitLocker can encrypt operating system drives, fixed data drives, and removable data drives.

Is there a noticeable performance impact when BitLocker is enabled on a Windows 8–based computer?

Generally it imposes a single-digit percentage performance overhead.

Approximately how long will initial encryption take when BitLocker is turned on?

BitLocker encryption occurs in the background while you continue to work, and the system remains usable, but encryption times vary depending on the type of drive that is being encrypted, the size of the drive, and the speed of the drive. If you are encrypting very large drives, you may want to set encryption to occur during times when you will not be using the drive.

In Windows 8 and Windows Server 2012, you can choose whether or not BitLocker should encrypt the entire drive or just the used space on the drive when you turn on BitLocker. On a new hard drive, encrypting just the used spaced can be considerably faster than encrypting the entire drive. When this encryption option is selected, BitLocker automatically encrypts data as it is saved, ensuring that no data is stored unencrypted.

What happens if the computer is turned off during encryption or decryption?

If the computer is turned off or goes into hibernation, the BitLocker encryption and decryption process will resume where it stopped the next time Windows starts. This is true even if the power is suddenly unavailable.

Does BitLocker encrypt and decrypt the entire drive all at once when reading and writing data?

No, BitLocker does not encrypt and decrypt the entire drive when reading and writing data. The encrypted sectors in the BitLocker-protected drive are decrypted only as they are requested from system read operations. Blocks that are written to the drive are encrypted before the system writes them to the physical disk. No unencrypted data is ever stored on a BitLocker-protected drive.

How can I prevent users on a network from storing data on an unencrypted drive?

In Windows 8, you can enable Group Policy settings to require that data drives be BitLocker-protected before a BitLocker-protected computer can write data to them. The policy settings you use for this are:

- **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives\Deny write access to fixed drives not protected by BitLocker**
- **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives\Deny write access to removable drives not protected by BitLocker**

When these policy settings are enabled, the BitLocker-protected operating system will mount any data drives that are not protected by BitLocker as read-only.

If you are concerned that your users might inadvertently store data in an unencrypted drives while using a computer that does not have BitLocker enabled, use access control lists (ACLs) and Group Policy to configure access control for the drives or hide the drive letter.

For additional information about how to hide drive letters, see article 231289 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=83219).

What system changes would cause the integrity check on my operating system drive to fail?

The following types of system changes can cause an integrity check failure and prevent the TPM from releasing the BitLocker key to decrypt the protected operating system drive:

- Moving the BitLocker-protected drive into a new computer.
- Installing a new motherboard with a new TPM.
- Turning off, disabling, or clearing the TPM.
- Changing any boot configuration settings.
- Changing the BIOS, UEFI firmware, master boot record, boot sector, boot manager, option ROM, or other early boot components or boot configuration data.

This functionality is by design; BitLocker treats unauthorized modification of any of the early boot components as a potential attack and will place the system into recovery mode. Authorized administrators can update boot components without entering recovery mode by disabling BitLocker beforehand.

What causes BitLocker to start into recovery mode when attempting to start the operating system drive?

The following list provides examples of specific events that will cause BitLocker to enter recovery mode when attempting to start the operating system drive:

- Changing the boot order to boot another drive in advance of the hard drive.

  Notes

  If you have configured your computers for Wake On LAN you should verify that the BIOS option under System Configuration\Built In Device Options\Wake On LAN is set to “Follow Boot Order”.

  If the option “Boot to Network” is selected, when the computer is awoken from the LAN, it will boot to the network, potentially fail and then boot to the hard drive by following the normal boot order. This results in a different measurement on a conventional BIOS system than following the normal boot order alone.

  The different measurement results in the message from BitLocker that the system boot information has changed and BitLocker goes in recovery mode (on reboot). Configuring the setting to “Follow Boot Order” will cause the same boot sequence BitLocker used when it was turned on. This will cause BitLocker to boot normally.

- Having the CD or DVD drive before the hard drive in the BIOS boot order and then inserting or removing a CD or DVD.
- Failing to boot from a network drive before booting from the hard drive.
- Docking or undocking a portable computer. In some instances (depending on the computer manufacturer and the BIOS), the docking condition of the portable computer is part of the system measurement and must be consistent to validate the system status and unlock BitLocker. This means that if a portable computer is connected to its docking station when BitLocker is turned on, then it might also need to be connected to the docking station when it is unlocked. Conversely, if a portable computer is not connected to its docking station when BitLocker is turned on, then it might need to be disconnected from the docking station when it is unlocked.
- Changes to the NTFS partition table on the disk including creating, deleting, or resizing a primary partition.
- Entering the personal identification number (PIN) incorrectly too many times so that the anti-hammering logic of the TPM is activated. Anti-hammering logic is software or hardware methods that increase the difficulty and cost of a brute force attack on a PIN by not accepting PIN entries until after a certain amount of time has passed.
- Turning off the support for reading the USB device in the pre-boot environment from the BIOS or UEFI firmware if you are using USB-based keys instead of a TPM.
- Turning off, disabling, deactivating, or clearing the TPM.
- Upgrading critical early startup components, such as a BIOS or UEFI firmware upgrade, causing the related boot measurements to change.
- Forgetting the PIN when PIN authentication has been enabled.
- Updating option ROM firmware.
- Upgrading TPM firmware.
- Adding or removing hardware. For example, inserting a new card in the computer, including some PCMIA wireless cards.
- Removing, inserting, or completely depleting the charge on a smart battery on a portable computer.
- Changes to the master boot record on the disk.
- Changes to the boot manager on the disk.
- Hiding the TPM from the operating system. Some BIOS or UEFI settings can be used to prevent the enumeration of the TPM to the operating system. When implemented, this option can make the TPM hidden from the operating system. When the TPM is hidden, BIOS and UEFI secure startup are disabled, and the TPM does not respond to commands from any software.
- Using a different keyboard that does not correctly enter the PIN or whose keyboard map does not match the keyboard map assumed by the pre-boot environment. This can prevent the entry of enhanced PINs.
- Modifying the Platform Configuration Registers (PCRs) used by the TPM validation profile. For example, including **PCR[1]** would result in BitLocker measuring most changes to BIOS settings, causing BitLocker to enter recovery mode even when non-boot critical BIOS settings change.

Note
Some computers have BIOS settings that skip measurements to certain PCRs, such as **PCR[2]**. Changing this setting in the BIOS would cause BitLocker to enter recovery mode because the PCR measurement will be different.

- Moving the BitLocker-protected drive into a new computer.
- Upgrading the motherboard to a new one with a new TPM.
- Losing the USB flash drive containing the startup key when startup key authentication has been enabled.
- Failing the TPM self-test.
- Having a BIOS, UEFI firmware, or an option ROM component that is not compliant with the relevant Trusted Computing Group standards for a client computer. For example, a non-compliant implementation may record volatile data (such as time) in the TPM measurements, causing different measurements on each startup and causing BitLocker to start in recovery mode.
- Changing the usage authorization for the storage root key of the TPM to a non-zero value.

Note
The BitLocker TPM initialization process sets the usage authorization value to zero, so another user or process must explicitly have changed this value.

- Disabling the code integrity check or enabling test signing on Windows Boot Manager (Bootmgr).
- Pressing the F8 or F10 key during the boot process.
- Adding or removing add-in cards (such as video or network cards), or upgrading firmware on add-in cards.
- Using a BIOS hot key during the boot process to change the boot order to something other than the hard drive.

Can I swap hard disks on the same computer if BitLocker is enabled on the operating system drive?

Yes, you can swap multiple hard disks on the same computer if BitLocker is enabled, but only if the hard disks were BitLocker-protected on the same computer. The BitLocker keys are unique to the TPM and operating system drive, so if you want to prepare a backup operating system or data drive for use in case of disk failure, you need to make sure that they were matched with the correct TPM. You can also configure different hard drives for different operating systems and then enable BitLocker on each one with different authentication methods (such as one with TPM-only and one with TPM+PIN) without any conflicts.

Can I access my BitLocker-protected drive if I insert the hard disk into a different computer?

Yes, if the drive is a data drive, you can unlock it from the **BitLocker Drive Encryption** Control Panel item just as you would any other data drive by using a password or smart card. If the data

drive was configured for automatic unlock only, you will have to unlock it by using the recovery key. If it is an operating system drive mounted on another computer running Windows 7 or later, the encrypted hard disk can be unlocked by a data recovery agent (if one was configured) or it can be unlocked by using the recovery key.

Note
Mounting the hard disk on another computer running Windows 8 is a quick and straightforward way to recover information from a damaged computer that has a BitLocker-protected drive on the hard disk.

Why is "Turn BitLocker on" not available when I right-click a drive?

Some drives cannot be encrypted with BitLocker. Reasons a drive cannot be encrypted include insufficient disk size, an incompatible file system, if the drive is a dynamic disk, or a drive is designated as the system partition. By default, the system drive (or system partition) is hidden from display in the Computer window. However, if it is not created as a hidden drive when the operating system was installed due to a custom installation process, that drive might be displayed but cannot be encrypted.

What type of disk configurations are supported by BitLocker?

In Windows Server 2012, Windows Server 2008, and Windows 8, any number of internal, fixed data drives can be protected with BitLocker. ATA and SATA-based, direct-attached storage devices are also supported. The following table details which disk configurations are supported and not supported by BitLocker.

Drive configuration	Supported	Not supported
Network	None	Network file system (NFS) Distributed File System (DFS)
Optical media	None	CD file system (CDFS) Live File System Universal Disk Format (UDF)
Software	Basic volumes	Software-based RAID systems Bootable and non-bootable virtual hard disks (VHDs) Dynamic volumes RAM disks
File system	NTFS FAT16	CD File system

Drive configuration	Supported	Not supported
	FAT32 ExFAT	
Drive connection	USB Firewire SATA SAS ATA IDE SCSI eSATA iSCSI (Windows 8 and Windows Server 2012 only) Fiber Channel (Windows 8 and Windows Server 2012 only)	Bluetooth
Device type	Solid state drives, such as USB flash drives Hardware-based RAID systems Hard disk drive	None

What if my disk configuration is not listed?

If your disk configuration is not listed in the previous question, it is a configuration that has not been fully tested by Microsoft.

Key management

What is the difference between a TPM owner password, recovery password, recovery key, password, PIN, enhanced PIN, and startup key?

There are multiple keys that can be generated and used by BitLocker. Some keys are required a nd some are optional protectors you can choose to use depending on the level of security you require.

TPM owner password

Prior to enabling BitLocker on a computer with a TPM version 1.2 or later, you must initialize the TPM. The initialization process generates a TPM owner password, which is a password set on the TPM. You must be able to supply the TPM owner password to change the state of the TPM, such as when enabling or disabling the TPM or resetting the TPM lockout.

Recovery password and recovery key

When you set up BitLocker, you must choose how access to BitLocker-protected drives can be recovered in the event that the specified unlock method cannot be used (such as if the TPM cannot validate the boot components, the personal identification number (PIN) is forgotten, or the password is forgotten). In these situations, you must be able to supply either the recovery key or the recovery password to unlock the encrypted data on the drive. In BitLocker user interface, the term "recovery key" is used generically to refer to both the recovery key file and the recovery password. When you supply the recovery information, you can use either of the following formats:

- A recovery password consisting of 48 digits divided into eight groups. During recovery, you need to type this password into the BitLocker recovery console by using the function keys on your keyboard.
- A key file on a USB flash drive that is read directly by the BitLocker recovery console. During recovery, you need to insert this USB device.

Password

A password can be used to protect fixed and removable data drives as well as operating system drives. When used with operating system drives, it can be used as an alternative key to the USB key for using BitLocker on computers that do not have a TPM. The password can consist of 8 to 255 characters as specified by the **Configure use of passwords for operating system drives**, **Configure use of passwords for removable data drives** and **Configure use of passwords for fixed data drives** Group Policy settings and is stored internally as a 256-bit hash of the entered characters. This value is never displayed to the user.

Warning

Passwords cannot be used if FIPS-only compliance is enabled for computers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows 8, Windows 7, or Windows Vista.

Introduced in Windows Server 2012 R2 and Windows 8.1, you can use passwords when FIPS compliant algorithms are used.

The **System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing** policy setting in **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options** specifies whether FIPS-only compliance is enabled.

PIN and enhanced PIN

For a higher level of security with the TPM, you can configure BitLocker with a personal identification number (PIN) has been typed. The PIN is a user-created value that must be entered each time the computer starts or resumes from hibernation.

The PIN can consist of 4 to 20 digits as specified by the **Configure minimum PIN length for startup** Group Policy setting and is stored internally as a 256-bit hash of the entered Unicode

characters. This value is never displayed to the user. The PIN is used to provide another factor of authentication in conjunction with TPM authentication.

For an even higher level of security with the TPM, you can configure BitLocker to use enhanced PINs. Enhanced PINs are PINs that use the full keyboard character set in addition to the numeric set to allow for more possible PIN combinations and are between 4 and 20 characters in length. To use enhanced PINs, you must enable the **Allow enhanced PINs for startup** Group Policy setting before adding the PIN to the drive. By enabling this policy, all PINs created can utilize full keyboard characters.

Note

To use enhanced PINs, your computer's BIOS or UEFI firmware must support using the full keyboard in the pre-boot environment. Users can run the optional system check during the BitLocker setup process to ensure the PIN can be entered correctly in the pre-boot environment. You should verify that the computers in your organization are compatible before making the use of enhanced PINs an organizational requirement.

When setting a BitLocker PIN by using the BitLocker setup wizard, the Manage-bde command-line tool, or through Windows Management Instrumentation (WMI) remote administration, you can use the wide character set. However, system firmware, either BIOS or Unified Extensible Firmware Interface (UEFI), may only support a standard EN-US keyboard and keymap during system startup. Additionally, BIOS-based systems are limited to 7-bit ASCII input during PIN entry. Thus, the use of either non-English characters or keys that differ in position from the EN-US keymap, such as QWERTZ and AZERTY keyboards, may cause boot-time PIN entry to fail. If your computer is affected by this limitation, it should be identified during the system check run by the BitLocker setup wizard. If it is not identified during the system check and the PIN is not able to be entered, you will need to supply the recovery key to unlock the drive.

We recommend that users set their keyboard layout to EN-US during enhanced PIN entry to avoid PIN entry failure in the pre-boot environment. If you are unable to enter an enhanced PIN from your keyboard even after setting the keyboard layout to EN-US, you must use a numeric-only PIN.

The following list identifies characters that are not currently supported by system firmware:

- Roman characters on keyboards with a non-EN-US keymap. For example, "Z" and "Y" on German keyboards and "Q" and "A" on French keyboards.
- Characters that are not available in 7-bit ASCII. For example, characters with umlauts, grave accents, and tildes.
- Symbols that are not available in 7-bit ASCII. For example, squared superscript, fractions, copyright, trademark, and international currency symbols.

Startup key

Configuring a startup key is another method to enable a higher level of security with the TPM. The startup key is a key stored on a USB flash drive, and the USB flash drive must be inserted every time the computer starts. The startup key is used to provide another factor of authentication in conjunction with TPM authentication. To use a USB flash drive as a startup key, the USB flash drive must be formatted by using the NTFS, FAT, or FAT32 file system.

How can the recovery password and recovery key be stored?

The recovery password and recovery key for an operating system drive or a fixed data drive can be saved to a folder, saved to one or more USB devices, saved to your Microsoft account online, or printed.

For removable data drives, the recovery password and recovery key can be saved to a folder, saved to your Microsoft account online, or printed. By default, you cannot store a recovery key for a removable drive on a removable drive.

A domain administrator can additionally configure Group Policy to automatically generate recovery passwords and store them in Active Directory Domain Services (AD DS) for any BitLocker-protected drive.

Is it possible to add an additional method of authentication without decrypting the drive if I only have the TPM authentication method enabled?

You can use the Manage-bde.exe command-line tool to replace your TPM-only authentication mode with a multifactor authentication mode. For example, if BitLocker is enabled with TPM authentication only and you want to add PIN authentication, use the following commands from an elevated command prompt, replacing *<4-20 digit numeric PIN>* with the numeric PIN you want to use:

manage-bde –protectors –delete %systemdrive% -type tpm

manage-bde –protectors –add %systemdrive% -tpmandpin *<4-20 digit numeric PIN>*

If I lose my recovery information, will the BitLocker-protected data be unrecoverable?

BitLocker is designed to make the encrypted drive unrecoverable without the required authentication. When in recovery mode, the user needs the recovery password or recovery key to unlock the encrypted drive. Therefore, we highly recommend that you either store the recovery information in AD DS, along with your Microsoft account online, or another safe location.

Can the USB flash drive that is used as the startup key also be used to store the recovery key?

While this is technically possible, it is not a best practice to use one USB flash drive to store both keys. If the USB flash drive that contains your startup key is lost or stolen, you also lose access to your recovery key. In addition, inserting this key would cause your computer to automatically boot from the recovery key even if TPM-measured files have changed, which circumvents the TPM's system integrity check.

Can I save the startup key on multiple USB flash drives?

Yes, you can save a computer's startup key on multiple USB flash drives. Right-clicking a BitLocker-protected drive and selecting **Manage BitLocker** will provide you the options to duplicate the recovery keys as needed.

Can I save multiple (different) startup keys on the same USB flash drive?

Yes, you can save BitLocker startup keys for different computers on the same USB flash drive.

Can I generate multiple (different) startup keys for the same computer?

You can generate different startup keys for the same computer through scripting. However, for computers that have a TPM, creating different startup keys prevents BitLocker from using the TPM's system integrity check.

Can I generate multiple PIN combinations?

It is not possible to generate multiple PIN combinations.

What encryption keys are used in BitLocker? How do they work together?

Raw data is encrypted with the full volume encryption key, which is then encrypted with the volume master key. The volume master key is in turn encrypted by one of several possible methods depending on your authentication (that is, key protectors or TPM) and recovery scenarios.

Where are the encryption keys stored?

The full volume encryption key is encrypted by the volume master key and stored in the encrypted drive. The volume master key is encrypted by the appropriate key protector and stored in the encrypted drive. If BitLocker has been suspended, the clear key that is used to encrypt the volume master key is also stored in the encrypted drive, along with the encrypted volume master key.

This storage process ensures that the volume master key is never stored unencrypted and is protected unless you disable BitLocker. The keys are also saved to two additional locations on the drive for redundancy. The keys can be read and processed by the boot manager.

Why do I have to use the function keys to enter the PIN or the 48-character recovery password?

The F1 through F10 keys are universally mapped scan codes available in the pre-boot environment on all computers and in all languages. The numeric keys 0 through 9 are not usable in the pre-boot environment on all keyboards.

When using an enhanced PIN, users should run the optional system check during the BitLocker setup process to ensure that the PIN can be entered correctly in the pre-boot environment. For more information about enhanced PINs see What is the difference between a TPM owner password, recovery password, recovery key, password, PIN, enhanced PIN, and startup key?

How does BitLocker help prevent an attacker from discovering the PIN that unlocks my operating system drive?

It is possible that a personal identification number (PIN) can be discovered by an attacker performing a brute force attack. A brute force attack occurs when an attacker uses an automated tool to try different PIN combinations until the correct one is discovered. For BitLocker-protected computers, this type of attack, also known as a dictionary attack, requires that the attacker have physical access to the computer.

The TPM has the built-in ability to detect and react to these types of attacks. Because different manufacturers' TPMs may support different PIN and attack mitigations, contact your TPM's manufacturer to determine how your computer's TPM mitigates PIN brute force attacks.

After you have determined your TPM's manufacturer (see How can I determine the manufacturer of my TPM?), contact the manufacturer to gather the TPM's vendor-specific information. Most manufacturers use the PIN authentication failure count to exponentially increase lockout time to the PIN interface. However, each manufacturer has different policies regarding when and how the failure counter is decreased or reset.

How can I determine the manufacturer of my TPM?

To determine your TPM manufacturer, use the following procedure.

To determine the TPM manufacturer

1. On the **Start** screen, type **tpm.msc**.
2. The TPM manufacturer is listed in the main pane, under **TPM Manufacturer Information**.

Note

The **Manufacturer Name** field in the **TPM Manufacturer Information** listing is information provided by the TPM and is often an abbreviation (such as ATML for Atmel, BRCM for Broadcomm, or IFX for Infineon).

How can I evaluate a TPM's dictionary attack mitigation mechanism?

The following questions can assist you when asking a TPM manufacturer about the design of a dictionary attack mitigation mechanism:

- How many failed authorization attempts can occur before lockout?
- What is the algorithm for determining the duration of a lockout based on the number of failed attempts and any other relevant parameters?
- What actions can cause the failure count and lockout duration to be decreased or reset?

Can PIN length and complexity be managed with Group Policy?

Yes and No. You can configure the minimum personal identification number (PIN) length by using the **Configure minimum PIN length for startup** Group Policy setting and allow the use of alphanumeric PINs by enabling the **Allow enhanced PINs for startup** Group Policy setting. However, you cannot require PIN complexity by Group Policy.

BitLocker To Go

What is BitLocker To Go?

BitLocker To Go is BitLocker Drive Encryption on removable data drives. This includes the encryption of USB flash drives, SD cards, external hard disk drives, and other drives formatted by using the NTFS, FAT16, FAT32, or exFAT file systems.

How can I authenticate or unlock my removable data drive?

You can unlock removable data drives by using a password or a smart card. After you start encryption, the drive can also be automatically unlocked on a specific computer for a specific user account. System administrators can configure which options are available for users, as well as password complexity and minimum length requirements.

What happens if I try to open a BitLocker-protected, NTFS-formatted removable drive by using a computer running Windows XP or Windows Vista?

In most cases, Windows XP and Windows Vista will not be able to recognize a BitLocker-protected, NTFS-formatted removable drive. In many situations, the user will be prompted to format the drive. Because of this, it is recommended that removable drives be formatted by using the FAT, FAT32, or exFAT file system when using BitLocker.

Is there a way to ensure the BitLocker To Go Reader is not installed on FAT-formatted drives?

Yes. Group Policy can prevent the application from being installed on the drives. The first option is to disable the policy settings that allow computers running Windows Vista, Windows XP with Service Pack 3 (SP3), or Windows XP with Service Pack 2 (SP2) to open BitLocker-protected data drives. These policy settings are located in the following locations:

- **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives\Allow access to BitLocker-protected removable data drives from earlier versions of Windows**
- **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives\Allow access to BitLocker-protected fixed data drives from earlier versions of Windows**

  To allow computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2 to open BitLocker-protected data drives but prevent the BitLocker To Go Reader application from being installed on the drive, enable these policy settings and select the **Do not install BitLocker To Go Reader on FAT formatted removable drives** and **Do not install BitLocker To Go Reader on FAT formatted fixed drives** check boxes on the respective policy settings.

Note

Before deleting the BitLocker To Go Reader from a drive, BitLocker checks that the identification field of the drive is either blank or matches the identification field for your organization.

Can I save files to my BitLocker-protected removable drive when I am using Windows XP or Windows Vista?

No. The BitLocker To Go Reader provides read-only access to BitLocker-protected removable drives.

Can I download a copy of the BitLocker To Go Reader?

Yes. To download the BitLocker To Go Reader, see Microsoft Download Center: BitLocker To Go Reader.

Why am I unable to access my removable drive on computers running Windows XP or Windows Vista when using the BitLocker To Go Reader?

The most common reason for this situation is that the drive is not formatted by using the FAT, FAT32, or exFAT file systems. To check for this, insert the drive in a computer running Windows 7 or later, right-click the drive, and then click **Properties** to see the file format of the drive. Another reason could be that the system administrator has disabled access to removable drives

from previous versions of Windows by using the BitLocker Group Policy settings. To check for this, attempt to access the drive from a computer running Windows XP or Windows Vista that is not joined to the domain.

Important
You should ensure that BitLocker has finished the encryption process on your drive before attempting to view the drive by using the BitLocker To Go Reader.

Active Directory Domain Services (AD DS)

Does BitLocker require a schema extension to store recovery information in AD DS?

This depends on the operating system and AD DS implementation.

Windows Server 2003 with Service Pack 1 (SP1)

In Windows Server 2003 with SP1, the schema must be extended to support storing BitLocker and TPM recovery and password information.

Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012

For these servers the schema already includes the required attributes.

What type of information is stored in AD DS?

Three primary pieces of information are stored in AD DS. The following table details this information.

Stored information	Description
Hash of the TPM owner password	The password hash can be stored only if the TPM is owned and the ownership was taken by using Windows features, such as the BitLocker Setup Wizard or the TPM snap-in.
BitLocker recovery password	The recovery password allows you to unlock and access the drive in the event of a recovery incident. Domain administrators can view the BitLocker recovery password by using the BitLocker Recovery Password Viewer.
BitLocker key package	The key package helps to repair damage to the hard disk that would otherwise prevent standard recovery. Using the key package for recovery requires the BitLocker Repair Tool, Repair-bde.

Does BitLocker encrypt recovery information as it is sent to AD DS?

Yes, the transmission of recovery information from a computer that is running Windows 7 or later to AD DS is protected by using the Kerberos authentication protocol. Specifically, the connection uses the authentication flags ADS_SECURE_AUTHENTICATION, ADS_USE_SEALING, and ADS_USE_SIGNING.

For more information about Active Directory authentication flags, see ADS_AUTHENTICATION_ENUM Enumeration (http://go.microsoft.com/fwlink/?LinkId=79643).

Notes

After recovery information is transmitted, AD DS does not store the BitLocker and TPM recovery information in an encrypted format. However, access control permissions are set so that only domain administrators or appropriate delegates can read the stored information when the server is online. Enterprises concerned about offline attacks on branch office servers should consider enabling BitLocker on those servers. We also recommended configuring your domain controllers to support encryption sealing and that any recovery retrieval application used in your organization use sealing as well.

For more information about developing applications that exchange encrypted data over a network, see the following articles on MSDN:

- Binding with Encryption (http://go.microsoft.com/fwlink/?LinkId=151844)
- Using ldap_init (http://go.microsoft.com/fwlink/?LinkId=151845)

For more information about configuring servers to support encryption sealing, see the following articles:

- Using SASL (http://go.microsoft.com/fwlink/?LinkId=151846)
- Using SSL/TLS (http://go.microsoft.com/fwlink/?LinkId=151847)
- How to enable LDAP signing in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=151848)

Is the BitLocker recovery information stored in plaintext in AD DS?

Yes, the recovery information is stored unencrypted in AD DS, but the entries have access control lists (ACLs) that limit access to only domain administrators.

If an attacker gains full access to AD DS, all computers in the domain, including BitLocker-protected computers, can be compromised. For more information about securing access to AD DS, see Securing Active Directory Administrative Groups and Accounts (http://go.microsoft.com/fwlink/?LinkId=83266).

What if BitLocker is enabled on a computer before the computer has joined the domain?

If BitLocker is enabled on a drive before Group Policy has been applied to enforce backup, the recovery information will not be automatically backed up to AD DS when the computer joins the domain or when Group Policy is subsequently applied. However, in Windows 8 you can use the **Choose how BitLocker-protected operating system drives can be recovered**, **Choose how BitLocker-protected fixed drives can be recovered** and **Choose how BitLocker-protected removable drives can be recovered** Group Policy settings to require that the computer be connected to a domain before BitLocker can be enabled to help ensure that recovery information for BitLocker-protected drives in your organization is backed up to AD DS.

The BitLocker Windows Management Instrumentation (WMI) interface does allow administrators to write a script to back up or synchronize an online client's existing recovery information; however, BitLocker does not automatically manage this process. The Manage-bde command-line tool can also be used to manually back up recovery information to AD DS. For example, to back up all of the recovery information for the C: drive to AD DS, you would use the following command from an elevated command prompt: **manage-bde -protectors -adbackup C:**.

Important

Joining a computer to the domain should be the first step for new computers within an organization. After computers are joined to a domain, storing the BitLocker recovery key to AD DS is automatic (when enabled in Group Policy).

Is there an event log entry recorded on the client computer to indicate the success or failure of the Active Directory backup?

Yes, an event log entry that indicates the success or failure of an Active Directory backup is recorded on the client computer. However, even if an event log entry says "Success," the information could have been subsequently removed from AD DS, or BitLocker could have been reconfigured in such a way that the Active Directory information can no longer unlock the drive (such as by removing the recovery password key protector). In addition, it is also possible that the log entry could be spoofed.

Ultimately, determining whether a legitimate backup exists in AD DS requires querying AD DS with domain administrator credentials by using the BitLocker password viewer tool.

If I change the BitLocker recovery password on my computer and store the new password in AD DS, will AD DS overwrite the old password?

No. By design, BitLocker recovery password entries do not get deleted from AD DS; therefore, you might see multiple passwords for each drive. To identify the latest password, check the date on the object.

What happens if the backup initially fails? Will BitLocker retry the backup?

If the backup initially fails, such as when a domain controller is unreachable at the time when the BitLocker setup wizard is run, BitLocker does not try again to back up the recovery information to AD DS.

When an administrator selects the **Require BitLocker backup to AD DS** check box of the **Store BitLocker recovery information in Active Directory Domain Service (Windows 2008 and Windows Vista)** policy setting, or the equivalent **Do not enable BitLocker until recovery information is stored in AD DS for (operating system | fixed data | removable data) drives** check box in any of the **Choose how BitLocker-protected operating system drives can be recovered**, **Choose how BitLocker-protected fixed data drives can be recovered**, **Choose how BitLocker-protected removable data drives can be recovered** policy settings, this prevents users from enabling BitLocker unless the computer is connected to the domain and the backup of BitLocker recovery information to AD DS succeeds. With these settings configured if the backup fails, BitLocker cannot be enabled, ensuring that administrators will be able to recover BitLocker-protected drives in the organization.

When an administrator clears these check boxes, the administrator is allowing a drive to be BitLocker-protected without having the recovery information successfully backed up to AD DS; however, BitLocker will not automatically retry the backup if it fails. Instead, administrators can create a script for the backup, as described earlier in **What if BitLocker is enabled on a computer before the computer has joined the domain?** to capture the information after connectivity is restored.

Security

What form of encryption does BitLocker use? Is it configurable?

BitLocker uses Advanced Encryption Standard (AES) as its encryption algorithm with configurable key lengths of 128 or 256 bits. The default encryption setting is AES-128, but the options are configurable by using Group Policy.

What is the best practice for using BitLocker on an operating system drive?

The recommended practice for BitLocker configuration on an operating system drive is to implement BitLocker on a computer with a TPM version 1.2 or 2.0 and a Trusted Computing Group (TCG)-compliant BIOS or UEFI firmware implementation, plus a PIN. By requiring a PIN that was set by the user in addition to the TPM validation, a malicious user that has physical access to the computer cannot simply start the computer.

What are the implications of using the sleep or hibernate power management options?

BitLocker on operating system drives in its basic configuration (with a TPM but without advanced authentication) provides additional security for the hibernate mode. However, BitLocker provides greater security when it is configured to use an advanced authentication mode (TPM+PIN, TPM+USB, or TPM+PIN+USB) with the hibernate mode. This method is more secure because returning from hibernation requires BitLocker authentication. As a best practice, we recommend that sleep mode be disabled and that you use TPM+PIN for the authentication method.

What are the advantages of a TPM?

Most operating systems use a shared memory space and rely on the operating system to manage physical memory. A TPM is a hardware component that uses its own internal firmware and logic circuits for processing instructions, thus shielding it from external software vulnerabilities. Attacking the TPM requires physical access to the computer. Additionally, the tools and skills necessary to attack hardware are often more expensive, and usually are not as available as the ones used to attack software. And because each TPM is unique to the computer that contains it, attacking multiple TPM computers would be difficult and time-consuming.

Note

Configuring BitLocker with an additional factor of authentication provides even more protection against TPM hardware attacks.

Is Microsoft pursuing any security certification for BitLocker?

All of the versions of BitLocker that have been included with the operating system have obtained the Federal Information Processing Standard (FIPS) 140-2 certification, and have been Common Criteria certified EAL4+. These certifications have also been completed for Windows 8, and Windows Server 2012, and Windows 8.1 and Windows Server 2012 R2 are in process.

BitLocker Network Unlock

What is BitLocker Network Unlock?

BitLocker Network Unlock enables easier management for BitLocker-enabled desktops and servers that use the TPM+PIN protection method in a domain environment. When a computer that is connected to a wired corporate network is rebooted, Network Unlock allows the PIN entry prompt to be bypassed. It automatically unlocks BitLocker-protected operating system volumes by using a trusted key that is provided by the Windows Deployment Services server as its secondary authentication method.

What is required to use BitLocker Network Unlock?

BitLocker Network Unlock has the following software and hardware requirements that must be met before you can use it:

Client computer requirements

- A DHCP driver that is implemented in the UEFI firmware
- Trusted Platform Module (TPM) 1.2 or TPM 2.0
- BitLocker enabled on the operating system volume

Windows Deployment Services server requirements

- BitLocker Network Unlock feature installed
- 2,048-bit RSA public/private key pair X.509 certificate present in FVENKP certificate store

Domain controller requirements

- Copy of the BitLocker Network Unlock Certificate from the Windows Deployment Services server on the domain controller to set Group Policy settings for Network Unlock.

How is BitLocker Network Unlock different from automatic unlock?

Automatic unlock uses a single protector, the one stored in the TPM. Network Unlock uses two protectors, the TPM protector and the one provided by the network or by your PIN. If the computer is joined to a network without the key protector it will prompt you to enter your PIN. If the PIN is not available you will need to use the recovery key to unlock the computer if it can ot be connected to the network.

How do I unlock my computer when not connected to the network?

To use Network Unlock you must also have a PIN configured for your computer. When your computer is not connected to the network you will need to provide the PIN to unlock it.

Other questions

Can I use EFS with BitLocker?

Yes, you can use Encrypting File System (EFS) to encrypt files on a BitLocker-protected drive. BitLocker helps protect the entire operating system drive against offline attacks, whereas EFS can provide additional user-based file level encryption for security separation between multiple users of the same computer. You can also use EFS to encrypt files on other drives that are not encrypted by BitLocker. The root secrets of EFS are stored by default on the operating system drive; therefore, if BitLocker is enabled for the operating system drive, data that is encrypted by EFS on other drives is also indirectly protected by BitLocker.

Can I run a kernel debugger with BitLocker?

Yes. However, the debugger should be turned on before enabling BitLocker. Turning on the debugger ensures that the correct measurements are calculated when sealing to the TPM, allowing the computer to start properly. If you need to turn debugging on or off when using BitLocker, be sure to suspend BitLocker first to avoid putting your computer into recovery mode.

How does BitLocker handle memory dumps?

BitLocker has a storage driver stack that ensures emory dumps are encrypted when BitLocker is enabled.

Can BitLocker support smart cards for pre-boot authentication?

BitLocker does not support smart cards for pre-boot authentication. There is no single industry standard for smart card support in the firmware, and most computers either do not implement firmware support for smart cards, or only support specific smart cards and readers. This lack of standardization makes supporting them very difficult.

Can I use a non-Microsoft TPM driver?

Microsoft does not support non-Microsoft TPM drivers and strongly recommends against using them with BitLocker. Attempting to use a non-Microsoft TPM driver with BitLocker may cause BitLocker to report that a TPM is not present on the computer and not allow the TPM to be used with BitLocker.

Can other tools that manage or modify the master boot record work with BitLocker?

We do not recommend modifying the master boot record on computers whose operating system drives are BitLocker-protected for a number of security, reliability, and product support reasons. Changes to the master boot record (MBR) could change the security environment and prevent the computer from starting normally, as well as complicate any efforts to recover from a corrupted MBR. Changes made to the MBR by anything other than Windows might force the computer into recovery mode or prevent it from booting entirely.

Why is the system check failing when I am encrypting my operating system drive?

The system check is designed to ensure your computer's BIOS or UEFI firmware is compatible with BitLocker and that the TPM is working correctly. The system check can fail for several reasons:

- The computer's BIOS or UEFI firmware cannot read USB flash drives.
- The computer's BIOS, uEFI firmware, or boot menu does not have reading USB flash drives enabled.

- There are multiple USB flash drives inserted into the computer.
- The PIN was not entered correctly.
- The computer's BIOS or UEFI firmware only supports using the function keys (F1–F10) to enter numerals in the pre-boot environment.
- The startup key was removed before the computer finished rebooting.
- The TPM has malfunctioned and fails to unseal the keys.

What can I do if the recovery key on my USB flash drive cannot be read?

Some computers cannot read USB flash drives in the pre-boot environment. First, check your BIOS or UEFI firmware and boot settings to ensure that the use of USB drives is enabled. If it is not enabled, enable the use of USB drives in the BIOS or UEFI firmware and boot settings and then try to read the recovery key from the USB flash drive again. If it still cannot be read, you will have to mount the hard drive as a data drive on another computer so that there is an operating system to attempt to read the recovery key from the USB flash drive. If the USB flash drive has been corrupted or damaged, you may need to supply a recovery password or use the recovery information that was backed up to AD DS. Also, if you are using the recovery key in the pre-boot environment, ensure that the drive is formatted by using the NTFS, FAT16, or FAT32 file system.

Why am I unable to save my recovery key to my USB flash drive?

The **Save to USB** option is not shown by default for removable drives. If the option is unavailable, it means that a system administrator has disallowed the use of recovery keys.

Why am I unable to automatically unlock my drive?

Automatic unlocking for fixed data drives requires that the operating system drive also be protected by BitLocker. If you are using a computer that does not have a BitLocker-protected operating system drive, the drive cannot be automatically unlocked. For removable data drives, you can add automatic unlocking by right-clicking the drive in Windows Explorer and clicking **Manage BitLocker**. You will still be able to use the password or smart card credentials you supplied when you turned on BitLocker to unlock the removable drive on other computers.

Can I use BitLocker in Safe Mode?

Limited BitLocker functionality is available in Safe Mode. BitLocker-protected drives can be unlocked and decrypted by using the **BitLocker Drive Encryption** Control Panel item. Right-clicking to access BitLocker options from Windows Explorer is not available in Safe Mode.

How do I "lock" a data drive?

Both fixed and removable data drives can be locked by using the Manage-bde command-line tool and the –lock command.

Note
Ensure all data is saved to the drive before locking it. Once locked, the drive will become inaccessible.

The syntax of this command is:

manage-bde <*driveletter*> **-lock**

Outside of using this command, data drives will be locked on shutdown and restart of the operating system. A removable data drive will also be locked automatically when the drive is removed from the computer.

Can I use BitLocker with the Volume Shadow Copy Service?

Yes. However, shadow copies made prior to enabling BitLocker will be automatically deleted when BitLocker is enabled on software-encrypted drives. If you are using a hardware encrypted drive, the shadow copies are retained.

Does BitLocker support virtual hard disks (VHDs)?

BitLocker is not supported on bootable VHDs, but BitLocker is supported on data volume VHDs, such as those used by clusters, if you are running Windows 8, Windows 8.1, Windows Server 2012 or Windows Server 2012 R2.

More information

- Prepare your organization for BitLocker: Planning and Policies
- BitLocker Group Policy Settings
- BCD Settings and BitLocker
- BitLocker: How to enable Network Unlock
- BitLocker: How to deploy on Windows Server 2012
- BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker
- BitLocker: Use BitLocker Recovery Password Viewer
- Encrypted Hard Drive
- BitLocker Cmdlets in Windows PowerShell

If you have a question about using BitLocker that hasn't been answered in this FAQ, please post it in the Windows 8 IT Pro Security Forum.

BitLocker Basic Deployment

BitLocker is included with the operating system for computers running Windows 8.1 Enterprise or Ultimate editions. All versions of Windows 8.1 support Device Encryption, which is based off BitLocker Drive Encryption, but is not as customizable. For all Windows Server editions, BitLocker is installed as an optional feature using Server Manager. In Windows Vista and Windows 7, BitLocker is provisioned post installation for system and data volumes through either the manage-bde command line interface or the Control Panel user interface. For computers that will be running Windows 8, Windows Server 2012 or later operating systems BitLocker can be provisioned before the operating system is installed.

The following sections provide information that will help you put together your basic deployment plan for implementing BitLocker in your organization:

1. Using BitLocker to encrypt volumes
2. Down-level compatibility
3. Using manage-bde to encrypt volumes with BitLocker
4. Using PowerShell to encrypt volumes with BitLocker

Using BitLocker to encrypt volumes

BitLocker provides full volume encryption (FVE) for operating system volumes, as well as fixed and removable data volumes. To support fully encrypted operating system volumes, BitLocker uses an unencrypted system volume for the files required to boot, decrypt, and load the operating system. This volume is automatically created during a new installation of both client and server operating systems.

In the event that the disk was prepared as a single contiguous space, BitLocker requires a new volume to hold the boot files. The BdeHdCfg.exe tool shipped with Windows can create these volumes.

Notes

For more information about using BdeHdCfg.exe, see:

http://technet.microsoft.com/en-us/library/ff829850(WS.10).aspx

BitLocker encryption can be done using the following methods:

- BitLocker control panel
- Windows Explorer
- manage-bde command line interface
- BitLocker Windows PowerShell cmdlets

Encrypting volumes using the BitLocker control panel

Encrypting volumes with the BitLocker control panel is how many users will utilize BitLocker. The name of the BitLocker control panel is BitLocker Drive Encryption. The BitLocker control panel

supports encrypting operating system, fixed data and removable data volumes. The BitLocker control panel will organize available drives in the appropriate category based on how the device reports itself to Windows. Only formatted volumes with assigned drive letters will appear properly in the BitLocker control panel applet.

To start encryption for a volume, select **Turn on BitLocker** for the appropriate drive to initialize the BitLocker Drive Encryption Wizard. BitLocker Drive Encryption Wizard options vary based on volume type (Operating system volume or Data volume).

Operating system volume

Upon launch, the BitLocker Drive Encryption Wizard verifies the computer meets the BitLocker system requirements for encrypting an operating system volume. By default, the system requirements are:

Requirement	Description
Hardware configuration	The computer must meet the minimum requirements for Windows 8. For more information about Windows 8 requirements, see the Windows 8 Web site.
Operating system	Windows 8, Windows Server 2012 **Note** BitLocker is an optional feature of Windows Server 2012. Use Server Manager to install BitLocker on a computer running Windows Server 2012.
Hardware TPM	TPM version 1.2 or 2.0 A TPM is not required for BitLocker; however, only a computer with a TPM can provide the additional security of pre-startup system integrity verification and multifactor authentication.
BIOS configuration	• A Trusted Computing Group (TCG)-compliant BIOS or UEFI firmware. • The boot order must be set to start first from the hard disk, and not the USB or CD drives. • The firmware must be able to read from a USB flash drive during startup.
File system	For computers that boot natively with UEFI

Requirement	Description
	firmware, at least one FAT32 partition for the system drive and one NTFS partition for the operating system drive. For computers with legacy BIOS firmware, at least two NTFS disk partitions, one for the system drive and one for the operating system drive. For either firmware, the system drive partition must be at least 350 megabytes (MB) and set as the active partition.
Hardware encrypted drive prerequisites (optional)	To use a hardware encrypted drive as the boot drive, the drive must be in the uninitialized state and in the security inactive state. In addition, the system must always boot with native UEFI version 2.3.1 or higher and the CSM (if any) disabled.

Upon passing the initial configuration, users are required to enter a password for the volume. If the volume does not pass the initial configuration for BitLocker, the user is presented with an error dialog describing the appropriate actions to be taken.

Notes

For more information on creating strong passwords, see:

http://www.microsoft.com/security/online-privacy/passwords-create.aspx

Once a strong password has been created for the volume, a recovery key will be generated. The BitLocker Drive Encryption Wizard will prompt for a location to save this key. A BitLocker recovery key is a special key that you can create when you turn on BitLocker Drive Encryption for the first time on each drive that you encrypt. You can use the recovery key to gain access to your computer if the drive that Windows is installed on (the operating system drive) is encrypted using BitLocker Drive Encryption and BitLocker detects a condition that prevents it from unlocking the drive when the computer is starting up. A recovery key can also be used to gain access to your files and folders on a removable data drive (such as an external hard drive or USB flash drive) that is encrypted using BitLocker To Go, if for some reason you forget the password or your computer cannot access the drive.

You should store the recovery key by printing it, saving it on removable media, or saving it as a file in a network folder or on your SkyDrive, or on another drive of your computer that you are not encrypting. You cannot save the recovery key to the root directory of a non-removable drive and cannot be stored on the encrypted volume. You cannot save the recovery key for a removable data drive (such as a USB flash drive) on removable media. Ideally, you should store the

recovery key separate from your computer. After you create a recovery key, you can use the BitLocker control panel to make additional copies.

When the recovery key has been properly stored, the BitLocker Drive Encryption Wizard will prompt the user to choose how to encrypt the drive. There are two options:

- Encrypt used disk space only - Encrypts only disk space that contains data
- Encrypt entire drive - Encrypts the entire volume including free space

It is recommended that drives with little to no data utilize the **used disk space only** encryption option and that drives with data or an operating system utilize the **encrypt entire drive** option.

Note

Deleted files appear as free space to the file system, which is not encrypted by **used disk space only**. Until they are wiped or overwritten, deleted files hold information that could be recovered with common data forensic tools.

Selecting an encryption type and choosing **Next** will give the user the option of running a BitLocker system check (selected by default) which will ensure that BitLocker can properly access the recovery and encryption keys before the volume encryption begins. It is recommended to run this system check before starting the encryption process. If the system check is not run and a problem is encountered when the operating system attempts to start, the user will need to provide the recovery key to start Windows.

After completing the system check (if selected), the BitLocker Drive Encryption Wizard will restart the computer to begin encryption. Upon reboot, users are required to enter the password chosen to boot into the operating system volume. Users can check encryption status by checking the system notification area or the BitLocker control panel.

Until encryption is completed, the only available options for managing BitLocker involve manipulation of the password protecting the operating system volume, backing up the recovery key, and turning BitLocker off.

Data volume

Encrypting data volumes using the BitLocker control panel interface works in a similar fashion to encryption of the operating system volumes. Users select **Turn on BitLocker** within the control panel to begin the BitLocker Drive Encryption wizard.

Unlike for operating system volumes, data volumes are not required to pass any configuration tests for the wizard to proceed. Upon launching the wizard, a choice of authentication methods to unlock the drive appears. The available options are **password** and **smart card** and **automatically unlock this drive on this computer**. Disabled by default, the latter option will unlock the data volume without user input when the operating system volume is unlocked.

After selecting the desired authentication method and choosing **Next**, the wizard presents options for storage of the recovery key. These options are the same as for operating system volumes.

With the recovery key saved, selecting **Next** in the wizard will show available options for encryption. These options are the same as for operating system volumes; **used disk space only**

and **full drive encryption**. If the volume being encrypted is new or empty, it is recommended that used space only encryption is selected.

With an encryption method chosen, a final confirmation screen displays before beginning the encryption process. Selecting **Start encrypting** will begin encryption.

Encryption status displays in the notification area or within the BitLocker control panel.

Windows SkyDrive option

There is a new option for storing the BitLocker recovery key using the Windows SkyDrive. This option requires that computers are not members of a domain and that the user is using a Microsoft Account. A Microsoft Account is an account tied to a user's Windows Live ID. Local accounts do not give the option to utilize SkyDrive. Using the SkyDrive option is the default, recommended recovery key storage method for computers that are not joined to a domain.

Users can verify the recovery key was saved properly by checking their SkyDrive for the BitLocker folder which is created automatically during the save process. The folder will contain two files, a readme.txt and the recovery key. For users storing more than one recovery password on their SkyDrive, they can identify the required recovery key by looking at the file name. The recovery key ID is appended to the end of the file name.

Using BitLocker within Windows Explorer

Windows Explorer allows users to launch the BitLocker Drive Encryption wizard by right clicking on a volume and selecting **Turn On BitLocker**. This option is available on client computers by default. On servers, you must first install the BitLocker and Desktop-Experience features for this option to be available. After selecting **Turn on BitLocker**, the wizard works exactly as it does when launched using the BitLocker control panel.

Down-level compatibility

The following table shows the compatibility matrix for systems that have been BitLocker enabled then presented to a different version of Windows.

Table 1: Cross compatibility for Windows 8.1, Windows 8, and Windows 7 encrypted volumes

Encryption Type	Windows 8.1	Windows 8	Windows 7
Fully encrypted on Windows 8	Presents as fully encrypted	N/A	Presented as fully encrypted
Used Disk Space Only encrypted on Windows 8	Presents as encrypt on write	N/A	Presented as fully encrypted
Fully encrypted volume from Windows	Presents as fully encrypted	Presented as fully encrypted	N/A

7			
Partially encrypted volume from Windows 7	Windows 8.1 will complete encryption regardless of policy	Windows 8 will complete encryption regardless of policy	N/A

Encrypting volumes using the manage-bde command line interface

Manage-bde is a command-line utility that can be used for scripting BitLocker operations. Manage-bde offers additional options not displayed in the BitLocker control panel. For a complete list of the options, see Manage-bde.

Manage-bde offers a multitude of wider options for configuring BitLocker. This means that using the command syntax may require care and possibly later customization by the user. For example, using just the `manage-bde -on` command on a data volume will fully encrypt the volume without any authenticating protectors. A volume encrypted in this manner still requires user interaction to turn on BitLocker protection, even though the command successfully completed because an authentication method needs to be added to the volume for it to be fully protected.

Command line users need to determine the appropriate syntax for a given situation. The following section covers general encryption for operating system volumes and data volumes.

Operating system volume

Listed below are examples of basic valid commands for operating system volumes. In general, using only the `manage-bde -on <drive letter>` command will encrypt the operating system volume with a TPM-only protector and no recovery key. However, many environments require more secure protectors such as passwords or PIN and expect to be able to recover information with a recovery key.

Determining volume status

A good practice when using manage-bde is to determine the volume status on the target system. Use the following command to determine volume status:

```
manage-bde -status
```

This command returns the volumes on the target, current encryption status and volume type (operating system or data) for each volume. Using this information, users can determine the best encryption method for their environment.

Enabling BitLocker without a TPM

For example, suppose that you want to enable BitLocker on a computer without a TPM chip. To properly enable BitLocker for the operating system volume, you will need to use a USB flash drive as a startup key to boot (in this example, the drive letter E). You would first create the startup key needed for BitLocker using the –protectors option and save it to the USB drive on E: and then begin the encryption process. You will need to reboot the computer when prompted to complete the encryption process.

```
manage-bde –protectors -add C: -startupkey E:

manage-bde -on C:
```

Enabling BitLocker with a TPM only

It is possible to encrypt the operating system volume without any defined protectors using manage-bde. The command to do this is:

```
manage-bde -on C:
```

This will encrypt the drive using the TPM as the protector. If a user is unsure of the protector for a volume, they can use the -protectors option in manage-bde to list this information with the command:

```
 manage-bde -protectors -get <volume>
```

Provisioning BitLocker with two protectors

Another example is a user on non-TPM hardware who wishes to add a password and SID-based protector to the operating system volume. In this instance, the user adds the protectors first. This is done with the command:

```
manage-bde -protectors -add C: -pw -sid <user or group>
```

This command will require the user to enter and then confirm the password protector before adding them to the volume. With the protectors enabled on the volume, the user just needs to turn BitLocker on.

Data volume

Data volumes use the same syntax for encryption as operating system volumes but they do not require protectors for the operation to complete. Encrypting data volumes can be done using the base command: `manage-bde -on <drive letter>` or users can choose to add protectors to the volume. It is recommended that at least one primary protector and a recovery protector be added to a data volume.

Enabling BitLocker with a password

A common protector for a data volume is the password protector. In the example below, we add a password protector to the volume and turn BitLocker on.

```
manage-bde -protectors -add -pw C:

manage-bde -on C:
```

Using manage-bde to encrypt volumes with BitLocker

Encrypting volumes using the BitLocker Windows PowerShell cmdlets

Windows PowerShell cmdlets provide an alternative way to work with BitLocker. Using Windows PowerShell's scripting capabilities, administrators can integrate BitLocker options into existing scripts with ease. The list below displays the available BitLocker cmdlets.

Name	Parameters
Add-BitLockerKeyProtector	-ADAccountOrGroup -ADAccountOrGroupProtector -Confirm -MountPoint -Password -PasswordProtector -Pin -RecoveryKeyPath -RecoveryKeyProtector -RecoveryPassword -RecoveryPasswordProtector -Service -StartupKeyPath -StartupKeyProtector -TpmAndPinAndStartupKeyProtector -TpmAndPinProtector -TpmAndStartupKeyProtector -TpmProtector -WhatIf
Backup-BitLockerKeyProtector	-Confirm -KeyProtectorId -MountPoint -WhatIf
Disable-BitLocker	-Confirm -MountPoint

	-WhatIf
Disable-BitLockerAutoUnlock	-Confirm -MountPoint -WhatIf
Enable-BitLocker	-AdAccountOrGroup -AdAccountOrGroupProtector -Confirm -EncryptionMethod -HardwareEncryption -Password -PasswordProtector -Pin -RecoveryKeyPath -RecoveryKeyProtector -RecoveryPassword -RecoveryPasswordProtector -Service -SkipHardwareTest -StartupKeyPath -StartupKeyProtector -TpmAndPinAndStartupKeyProtector -TpmAndPinProtector -TpmAndStartupKeyProtector -TpmProtector -UsedSpaceOnly -WhatIf
Enable-BitLockerAutoUnlock	-Confirm -MountPoint -WhatIf
Get-BitLockerVolume	-MountPoint
Lock-BitLocker	-Confirm -ForceDismount -MountPoint -WhatIf

Remove-BitLockerKeyProtector	-Confirm -KeyProtectorId -MountPoint -WhatIf
Resume-BitLocker	-Confirm -MountPoint -WhatIf
Suspend-BitLocker	-Confirm -MountPoint -RebootCount -WhatIf
Unlock-BitLocker	-AdAccountOrGroup -Confirm -MountPoint -Password -RecoveryKeyPath -RecoveryPassword -RecoveryPassword -WhatIf

Similar to manage-bde, the Windows PowerShell cmdlets allow configuration beyond the options offered in the control panel. As with manage-bde, users need to consider the specific needs of the volume they are encrypting prior to running Windows PowerShell cmdlets.

A good initial step is to determine the current state of the volume(s) on the computer. You can do this using the `Get-BitLocker` volume cmdlet.

This output gives information on the volume type, protectors, protection status, and other useful information. This output shows that currently none of the volumes on this computer is BitLocker enabled, though two have existing key protectors.

Occasionally, all protectors may not be shown when using **Get-BitLockerVolume** due to lack of space in the output display. If you do not see all of the protectors for a volume, you can use the Windows PowerShell pipe command (|) to format a listing of the protectors.

Note

In the event that there are more than four protectors for a volume, the pipe command may run out of display space. For volumes with more than four protectors, use the method described in the section below to generate a listing of all protectors with protector ID.

```
Get-BitLockerVolume C: | fl
```

If you wanted to remove the existing protectors prior to provisioning BitLocker on the volume, you can utilize the `Remove-BitLockerKeyProtector` cmdlet. Accomplishing this requires the GUID associated with the protector to be removed.

A simple script can pipe the values of each **Get-BitLockerVolume** return out to another variable as seen below:

```
$vol = Get-BitLockerVolume

$keyprotectors = $vol.KeyProtector
```

Using this, we can display the information in the **$keyprotectors** variable to determine the GUID for each protector.

Using this information, we can then remove the key protector for a specific volume using the command:

```
Remove-BitLockerKeyProtector <volume>: -KeyProtectorID "{GUID}"
```

Note

The BitLocker cmdlet requires the key protector GUID enclosed in quotation marks to execute. Ensure the entire GUID, with braces, is included in the command.

Operating system volume

Using the BitLocker Windows PowerShell cmdlets is similar to working with the manage-bde tool for encrypting operating system volumes. Windows PowerShell offers users a lot of flexibility. For example, users can add the desired protector as part command for encrypting the volume. Below are examples of common user scenarios and steps to accomplish them using the BitLocker cmdlets for Windows PowerShell.

To enable BitLocker with just the TPM protector. This can be done using the command:

```
Enable-BitLocker C:
```

The example below adds one additional protector, the StartupKey protectors, and chooses to skip the BitLocker hardware test. In this example, encryption starts immediately without the need for a reboot.

```
Enable-BitLocker C: -StartupKeyProtector -StartupKeyPath <path> -SkipHardwareTest
```

Data volume

Data volume encryption using Windows PowerShell is the same as for operating system volumes. You should add the desired protectors prior to encrypting the volume. The following example adds a password protector to the E: volume using the variable $pw as the password. The $pw variable is held as a SecureString value to store the user defined password. Last, encryption begins.

```
$pw = Read-Host -AsSecureString

<user inputs password>
```

```
Enable-BitLockerKeyProtector E: -PasswordProtector -Password $pw
```

Using a SID based protector in Windows PowerShell

The ADAccountOrGroup protector is an Active Directory SID-based protector. This protector can be added to both operating system and data volumes, although it does not unlock operating system volumes in the pre-boot environment. The protector requires the SID for the domain account or group to link with the protector. BitLocker can protect a cluster-aware disk by adding a SID-based protector for the Cluster Name Object (CNO) that lets the disk properly failover and be unlocked to any member computer of the cluster.

Warning

The SID-based protector requires the use of an additional protector (such as TPM, PIN, recovery key, etc.) when used on operating system volumes.

To add an ADAccountOrGroup protector to a volume requires either the actual domain SID or the group name preceded by the domain and a backslash. In the example below, the CONTOSO\Administrator account is added as a protector to the data volume G.

```
Enable-BitLocker G: -AdAccountOrGroupProtector -AdAccountOrGroup <account>
```

For users who wish to use the SID for the account or group, the first step is to determine the SID associated with the account. To get the specific SID for a user account in Windows PowerShell, use the following command:

```
get-aduser -filter {samaccountname -eq "administrator"}
```

Note

Use of this command requires the RSAT-AD-PowerShell feature.

Tip

In addition to the PowerShell command above, information about the locally logged on user and group membership can be found using: WHOAMI /ALL. This does not require the use of additional features.

In the example below, the user wishes to add a domain SID based protector to the previously encrypted operating system volume. The user knows the SID for the user account or group they wish to add and uses the following command:

```
Add-BitLockerKeyProtector C: -ADAccountOrGroupProtector -ADAccountOrGroup "<SID>"
```

Note

Active Directory-based protectors are normally used to unlock Failover Cluster enabled volumes.

Using PowerShell to encrypt volumes with BitLocker

Checking BitLocker Status

To check the BitLocker status of a particular volume, administrators can look at the status of the drive in the BitLocker control panel applet, Windows Explorer, manage-bde command line tool, or Windows PowerShell cmdlets. Each option offers different levels of detail and ease of use. We will look at each of the available methods in the following section.

Checking BitLocker status with the control panel

Checking BitLocker status with the control panel is the most common method used by most users. Once opened, the status for each volume will display next to the volume description and drive letter. Available status return values with the control panel include:

Status	Description
On	BitLocker is enabled for the volume
Off	BitLocker is not enabled for the volume
Suspended	BitLocker is suspended and not actively protecting the volume
Waiting for Activation	BitLocker is enabled with a clear protector key and requires further action to be fully protected

If a drive is pre-provisioned with BitLocker, a status of "Waiting for Activation" displays with a yellow exclamation icon on volume E. This status means that there was only a clear protector used when encrypting the volume. In this case, the volume is not in a protected state and needs to have a secure key added to the volume before the drive is fully protected. Administrators can use the control panel, manage-bde tool, or WMI APIs to add an appropriate key protector. Once complete, the control panel will update to reflect the new status.

Using the control panel, administrators can choose **Turn on BitLocker** to start the BitLocker Drive Encryption wizard and add a protector, like PIN for an operating system volume (or password if no TPM exists), or a password or smart card protector to a data volume.

The drive security window displays prior to changing the volume status. Selecting **Activate BitLocker** will complete the encryption process.

Once BitLocker protector activation is completed, the completion notice is displayed.

Checking BitLocker status with manage-bde

Administrators who prefer a command line interface can utilize manage-bde to check volume status. Manage-bde is capable of returning more information about the volume than the graphical

user interface tools in the control panel. For example, manage-bde can display the BitLocker version in use, the encryption type, and the protectors associated with a volume.

To check the status of a volume using manage-bde, use the following command:

```
manage-bde -status <volume>
```

Note

If no volume letter is associated with the -status command, all volumes on the computer display their status.

Checking BitLocker status with Windows PowerShell

Windows PowerShell commands offer another way to query BitLocker status for volumes. Like manage-bde, Windows PowerShell includes the advantage of being able to check the status of a volume on a remote computer.

Using the Get-BitLockerVolume cmdlet, each volume on the system will display its current BitLocker status. To get information that is more detailed on a specific volume, use the following command:

```
Get-BitLockerVolume <volume> -Verbose | fl
```

This command will display information about the encryption method, volume type, key protectors, etc.

Provisioning BitLocker during operating system deployment

Administrators can enable BitLocker prior to operating system deployment from the Windows Pre-installation Environment. This is done with a randomly generated clear key protector applied to the formatted volume and encrypting the volume prior to running the Windows setup process. If the encryption uses the Used Disk Space Only option described later in this document, this step takes only a few seconds and incorporates well into regular deployment processes.

Decrypting BitLocker volumes

Decrypting volumes removes BitLocker and any associated protectors from the volumes. Decryption should occur when protection is no longer required. BitLocker decryption should not occur as a troubleshooting step. BitLocker can be removed from a volume using the BitLocker control panel applet, manage-bde, or Windows PowerShell cmdlets. We will discuss each method further below.

Decrypting volumes using the BitLocker control panel applet

BitLocker decryption using the control panel is done using a Wizard. The control panel can be called from Windows Explorer or by opening the directly. After opening the BitLocker control panel, users will select the Turn off BitLocker option to begin the process.

Once selected, the user chooses to continue by clicking the confirmation dialog. With Turn off BitLocker confirmed, the drive decryption process will begin and report status to the control panel.

The control panel does not report decryption progress but displays it in the notification area of the task bar. Selecting the notification area icon will open a modal dialog with progress.

Once decryption is complete, the drive will update its status in the control panel and is available for encryption.

Decrypting volumes using the manage-bde command line interface

Decrypting volumes using manage-bde is very straightforward. Decryption with manage-bde offers the advantage of not requiring user confirmation to start the process. Manage-bde uses the -off command to start the decryption process. A sample command for decryption is:

```
manage-bde -off C:
```

This command disables protectors while it decrypts the volume and removes all protectors when decryption is complete. If a user wishes to check the status of the decryption, they can use the following command:

```
manage-bde -status C:
```

Decrypting volumes using the BitLocker Windows PowerShell cmdlets

Decryption with Windows PowerShell cmdlets is straightforward, similar to manage-bde. The additional advantage Windows PowerShell offers is the ability to decrypt multiple drives in one pass. In the example below, the user has three encrypted volumes, which they wish to decrypt.

Using the Disable-BitLocker command, they can remove all protectors and encryption at the same time without the need for additional commands. An example of this command is:

```
DisableBitLocker
```

If a user did not want to input each mount point individually, using the `-MountPoint` parameter in an array can sequence the same command into one line without requiring additional user input. An example command is:

```
Disable-BitLocker -MountPoint E:,F:,G:
```

More information

BitLocker Overview

Encrypted Hard Drive

BitLocker Frequently Asked Questions (FAQ)

What's New in BitLocker for Windows 8 and Windows Server 2012

What's New in BitLocker for Windows 8.1 and Windows Server 2012 R2

Prepare your organization for BitLocker: Planning and Policies

BitLocker Recovery Guide

BitLocker: How to enable Network Unlock

Prepare your organization for BitLocker: Planning and Policies

When you design your BitLocker deployment strategy, define the appropriate policies and configuration requirements based on the business requirements of your organization. The following topics will help you collect information that you can use to frame your decision-making process about deploying and managing BitLocker systems.

- Audit your environment
- Encryption keys and authentication
- TPM hardware configurations
- Non-TPM hardware configurations
- Disk configuration considerations
- BitLocker provisioning
- Used Disk Space Only encryption
- Active Directory Domain Services considerations
- FIPS support for recovery password protector
- BitLocker Group Policy Settings

Audit your environment

To plan your enterprise deployment of BitLocker, you must first understand your current environment. Conduct an informal audit to define your current policies, procedures, and hardware environment. Begin by reviewing your existing corporate security policies as they relate to disk encryption software. If your organization is not currently using disk encryption software, none of these policies will exist. If you are using disk encryption software, then you might need to modify your organization's policies to address the capabilities of BitLocker.

Use the following questions to help you document your organization's current disk encryption security policies:

1. Are there policies to address which computers will use BitLocker and which computers will not use BitLocker?
2. What policies exist to control recovery password and recovery key storage?
3. What are the policies for validating the identity of users that need to perform BitLocker recovery?
4. What policies exist to control who in the organization has access to recovery data?
5. What policies exist to control computer decommissioning or retirement?

Encryption keys and authentication

BitLocker helps prevent unauthorized access to data on lost or stolen computers by:

- Encrypting the entire Windows operating system volume on the hard disk.

- Verifying the boot process integrity.

The trusted platform module (TPM)is a hardware component installed in many newer computers by the computer manufacturers. It works with BitLocker to help protect user data and to ensure that a computer has not been tampered with while the system was offline.

In addition, BitLocker offers the option to lock the normal startup process until the user supplies a personal identification number (PIN) or inserts a removable USB device, such as a flash drive, that contains a startup key. These additional security measures provide multifactor authentication and assurance that the computer will not start or resume from hibernation until the correct PIN or startup key is presented.

On computers that do not have a TPM version 1.2 or greater, you can still use BitLocker to encrypt the Windows operating system volume. However, this implementation will require the user to insert a USB startup key to start the computer or resume from hibernation, and does not provide the pre-startup system integrity verification offered by BitLocker working with a TPM.

BitLocker key protectors

Key protector	Description
TPM	A hardware device used to help establish a secure root-of-trust. BitLocker only supports TPM version 1.2 or greater.
PIN	A user-entered numeric key protector that can only be used in addition to the TPM.
Enhanced PIN	A user-entered alphanumeric key protector that can only be used in addition to the TPM.
Startup key	An encryption key that can be stored on most removable media. This key protector can be used alone on non-TPM computers, or in conjunction with a TPM for added security.
Recovery password	A 48-digit number used to unlock a volume when it is in recovery mode. Numbers can often be typed on a regular keyboard, if the numbers on the normal keyboard are not responding you can always use the function keys (F1-F10) to input the numbers.
Recovery key	An encryption key stored on removable media that can be used for recovering data encrypted on a BitLocker volume.

BitLocker authentication methods

Authentication method	Requires user interaction	Description
TPM only	No	TPM validates early boot components.
TPM + PIN	Yes	TPM validates early boot components. The user must enter the correct PIN before the start-up process can continue, and before the drive can be unlocked. The TPM will enter lockout if the incorrect PIN is entered repeatedly to protect the PIN from brute force attacks. The number of repeated attempts that will trigger a lockout is variable.
TPM + Network key	No	The TPM successfully validates early boot components, and a valid encrypted network key has been provided from the WDS server. This authentication method provides automatic unlock of operating system volumes at system reboot while still maintaining multifactor authentication.
TPM + startup key	Yes	The TPM successfully validates early boot components, and a USB flash drive containing the startup key has been inserted.
Startup key only	Yes	The user is prompted to insert the USB flash drive that holds the recovery key and/or startup key and reboot the computer.

Will you support computers without TPM version 1.2 or greater?

Determine whether you will support computers that do not have a TPM version 1.2 or greater in your environment. If you choose to support BitLocker on this type of computer, a user must use a

USB startup key to boot the system. This requires additional support processes similar to multifactor authentication.

What areas of your organization need a baseline level of data protection?

The TPM-only authentication method will provide the most transparent user experience for organizations that need a baseline level of data protection to meet security policies. It has the lowest total cost of ownership. TPM-only might also be more appropriate for computers that are unattended or that must reboot unattended.

However, TPM-only authentication method offers the lowest level of data protection. This authentication method protects against attacks that modify early boot components, but the level of protection can be affected by potential weaknesses in hardware or in the early boot components. BitLocker's multifactor authentication methods significantly increase the overall level of data protection.

What areas of your organization need a more secure level of data protection?

If there are areas of your organization where data residing on user computers is considered highly-sensitive, consider the best practice of deploying BitLocker with multifactor authentication on those systems. Requiring the user to input a PIN significantly increases the level of protection for the system. You can also use BitLocker Network Unlock to allow these computers to automatically unlock when connected to a trusted wired network that can provide the Network Unlock key.

What multifactor authentication method does your organization prefer?

The protection differences provided by multifactor authentication methods cannot be easily quantified. Consider each authentication method's impact on Helpdesk support, user education, user productivity, and automated systems management processes.

TPM hardware configurations

In your deployment plan, identify what TPM-based hardware platforms will be supported. Document the hardware models from an OEM of your choice, so that their configurations can be tested and supported. TPM hardware requires special consideration during all aspects of planning and deployment.

TPM states of existence

For each of the TPM states of existence, the TPM can transition into another state (for example, moving from disabled to enabled). The states are not exclusive.

State	Description
Enabled	Most features of the TPM are available. The TPM may be enabled and disabled multiple times within a boot period, if ownership is taken.

State	Description
Disabled	The TPM restricts most operations. Exceptions include the ability to report TPM capabilities, extend and reset Platform Configuration Register (PCR) functions, and to perform hashing and basic initialization. The TPM may be enabled and disabled multiple times within a boot period.
Activated	Most features of the TPM are available. The TPM may be activated and deactivated only through physical presence which requires a reboot.
Deactivated	Similar to disabled, with the exception that ownership can be taken while deactivated and enabled. The TPM may be activated and deactivated only through physical presence which requires a reboot.
Owned	Most features of the TPM are available. The TPM has an endorsement key and storage root key, and the owner knows information about owner authorization data.
Un-owned	The TPM does not have a storage root key and may or may not have an endorsement key.

Important
BitLocker cannot use the TPM until it is in the following state: enabled, activated, and owned. When the TPM is in this state and only when it is in this state, all operations are available.

The state of the TPM exists independent of the computer's operating system. Once the TPM is enabled, activated, and owned, the state of the TPM is preserved if the operating system is reinstalled.

Endorsement keys

For a TPM to be usable by BitLocker, it must contain an endorsement key, which is an RSA key pair. The private half of the key pair is held inside the TPM and is never revealed or accessible outside the TPM. If the TPM does not contain an endorsement key, BitLocker will force the TPM to generate one automatically as part of BitLocker setup.

An endorsement key can be created at various points in the TPM's lifecycle, but needs to be created only once for the lifetime of the TPM. If an endorsement key does not exist for the TPM, it must be created before TPM ownership can be taken.

For more information about the TPM and the TCG, see the Trusted Computing Group: Trusted Platform Module (TPM) Specifications (http://go.microsoft.com/fwlink/?LinkID=69584).

Non-TPM hardware configurations

Devices that do not include a TPM can still be protected by drive encryption. Windows To Go workspaces can be BitLocker protected using a startup password and PCs without a TPM can use a startup key.

Use the following questions to identify issues that might affect your deployment in a non-TPM configuration:

- Are password complexity rules in place?
- Do you have budget for USB flash drives for each of these computers?
- Do your existing non-TPM devices support USB devices at boot time?

Test your individual hardware platforms with the BitLocker system check option while you are enabling BitLocker. The system check will ensure that BitLocker can read the recovery information from a USB device and encryption keys correctly before it encrypts the volume. CD and DVD drives cannot act as a block storage device and cannot be used to store the BitLocker recovery material.

Disk configuration considerations

To function correctly, BitLocker requires a specific disk configuration. BitLocker requires two partitions that meet the following requirements:

- The operating system partition contains the operating system and its support files; it must be formatted with the NTFS file system
- The system partition (or boot partition) contains the files that are needed to load Windows after the BIOS or UEFI firware has prepared the system hardware. BitLocker is not enabled on this partition. For BitLocker to work, the system partition must not be encrypted and must be on a different partition than the operating system. On UEFI platforms the system partition must be formatted with the FAT 32 file system. On BIOS platforms the system partition must be formatted with the NTFS file system. It should be at least 350 MB in size

Windows setup will automatically configure the disk drives of your computer to support BitLocker encryption.

Windows Recovery Environment (Windows RE) is an extensible recovery platform that is based on Windows Pre-installation Environment (Windows PE). When the computer fails to start, Windows automatically transitions into this environment, and the Startup Repair tool in Windows RE automates the diagnosis and repair of an unbootable Windows installation. Windows RE also contains the drivers and tools that are needed to unlock a volume protected by BitLocker by providing a recovery key or recovery password. To use Windows RE in conjunction with

BitLocker, the Windows RE boot image must reside on a volume that is not protected by BitLocker.

Windows RE can also be used from boot media other than the local hard disk. If you choose not to install Windows RE on the local hard disk of BitLocker-enabled computers, you can use alternate boot methods, such as Windows Deployment Services, CD-ROM, or USB flash drive, for recovery.

BitLocker provisioning

In Windows Vista and Windows 7, BitLocker was provisioned post installation for system and data volumes through either the manage-bde command line interface or the Control Panel user interface. With newer operating systems, BitLocker can be easily provisioned before the operating system is installed. Preprovisioning requires that the computer have a TPM.

To check the BitLocker status of a particular volume, administrators can look at the status of the drive in the BitLocker control panel applet or Windows Explorer. A status of "Waiting For Activation" with a yellow exclamation icon means that the drive was preprovisioned for BitLocker. This status means that there was only a clear protector used when encrypting the volume. In this case, the volume is not protected and needs to have a secure key added to the volume before the drive is considered fully protected. Administrators can use the control panel options, manage-bde tool or WMI APIs to add an appropriate key protector and the volume status will be updated.

When using the control panel options, administrators can choose to **Turn on BitLocker** and follow the steps in the wizard to add a protector, such as a PIN for an operating system volume (or a password if no TPM exists), or a password or smart card protector to a data volume. Then the drive security window is presented prior to changing the volume status.

Administrators can enable BitLocker prior to operating system deployment from the Windows Pre-installation Environment (WinPE). This is done with a randomly generated clear key protector applied to the formatted volume and encrypting the volume prior to running the Windows setup process. If the encryption uses the Used Disk Space Only option this step takes only a few seconds and so incorporates well into regular deployment processes.

Used Disk Space Only encryption

The BitLocker Setup wizard provides administrators the ability to choose the Used Disk Space Only or Full encryption method when enabling BitLocker for a volume. Administrators can use the new BitLocker Group Policy setting to enforce either Used Disk Space Only or Full disk encryption.

Launching the BitLocker Setup wizard prompts for the authentication method to be used (password and smart card are available for data volumes). Once the method is chosen and the recovery key is saved, you are asked to choose the drive encryption type, either Used Disk Space Only or Full drive encryption.

Used Disk Space Only means that only the portion of the drive that contains data will be encrypted, unused space will remain unencrypted. This causes the encryption process to be

much faster, especially for new PCs and data drives. When BitLocker is enabled with this method as data is added to the drive the portion of the drive used will be encrypted, so there is never unencrypted data stored on the drive.

Full drive encryption means that the entire drive will be encrypted, regardless of whether data is stored on it or not. This is useful for drives that have been repurposed and may contain data remnants from their previous use.

Active Directory Domain Services considerations

BitLocker integrates with Active Directory Domain Services (AD DS) to provide centralized key management. By default, no recovery information is backed up to Active Directory. Administrators can configure Group Policy settings to enable backup of BitLocker or TPM recovery information. Before configuring these settings verify that access permissions have been granted to perform the backup.

By default, domain administrators are the only users that will have access to BitLocker recovery information. When you plan your support process, define what parts of your organization need access to BitLocker recovery information. Use this information to define how the appropriate rights will be delegated in your AD DS environment.

It is a best practice to require backup of recovery information for both the TPM and BitLocker to AD DS. You can implement this practice by configuring the Group Policy settings below for your BitLocker-protected computers.

BitLocker Group Policy setting	Configuration
BitLocker Drive Encryption: Turn on BitLocker backup to Active Directory Domain Services	Require BitLocker backup to AD DS (Passwords and key packages)
Trusted Platform Module Services: Turn on TPM backup to Active Directory Domain Services	Require TPM backup to AD DS

The following recovery data will be saved for each computer object:

- **Recovery password**

  A 48-digit recovery password used to recover a BitLocker-protected volume. Users enter this password to unlock a volume when BitLocker enters recovery mode.
- **Key package data**

  With this key package and the recovery password, you will be able decrypt portions of a BitLocker-protected volume if the disk is severely damaged. Each key package will only work with the volume it was created on, which can be identified by the corresponding volume ID.
- **TPM owner authorization password hash**

  When ownership of the TPM is taken a hash of the ownership password can be taken and stored in AD DS. This information can then be used to reset ownership of the TPM.

For Windows 8 a change to how the TPM owner authorization value is stored in AD DS was implemented in the AD DS schema. The TPM owner authorization value is now stored in a separate object which is linked to the Computer object. This value was stored as a property in the Computer object itself for the default Windows Server 2008 R2 schemas.

To take advantage of this integration, you must upgrade your domain controllers to Windows Server 2012 or extend the Active Directory schema and configure BitLocker-specific Group Policy objects.

Note

The account that you use to update the Active Directory schema must be a member of the Schema Admins group.

Windows Server 2012 domain controllers have the default schema to backup TPM owner authorization information in the separate object. If you are not upgrading your domain controller to Windows Server 2012 you need to extend the schema to support this change.

To support Windows 8 computers that are managed by a Windows Server 2003 or Windows 2008 domain controller

There are two schema extensions that you can copy down and add to your AD DS schema:

- **TpmSchemaExtension.ldf**

  This schema extension brings parity with the Windows Server 2012 schema. With this change, the TPM owner authorization information is stored in a separate TPM object linked to the corresponding computer object. Only the Computer object that has created the TPM object can update it. This means that any subsequent updates to the TPM objects will not succeed in dual boot scenarios or scenarios where the computer is reimaged resulting in a new AD computer object being created. To support such scenarios, an update to the schema was created.

- **TpmSchemaExtensionACLChanges.ldf**

  This schema update modifies the ACLs on the TPM object to be less restrictive so that any subsequent operating system which takes ownership of the computer object can update the owner authorization value in AD DS. However, this is less secure as any computer in the domain can now update the OwnerAuth of the TPM object (although it cannot read the OwnerAuth) and DOS attacks can be made from within the enterprise. The recommended mitigation in such a scenario is to do regular backup of TPM objects and enable auditing to track changes for these objects.

To download the schema extensions, see Schema Extensions for Windows Server 2008 R2 to support AD DS backup of TPM information from Windows 8 clients.

If you have a Windows Server 2012 domain controller in your environment, the schema extensions are already in place and do not need to be updated.

Caution

To configure Group Policy objects to backup TPM and BitLocker information in AD DS at least one of the domain controllers in your forest must be running Windows Server 2008 R2 or Windows Server 2012.

If Active Directory backup of the TPM owner authorization value is enabled in an environment without the required schema extensions, the TPM provisioning will fail and the TPM will remain in a Not Ready state for computers running Windows 8.

Setting the correct permissions in AD DS

To initialize the TPM successfully so that you can turn on BitLocker requires that the correct permissions for the SELF account in be set in AD DS for the **ms-TPMOwnerInformation** attribute. The following steps detail setting these permissions as required by BitLocker:

1. Open **Active Directory Users and Computers**.
2. Select the organizational unit (OU) which contains the computer accounts that will have BitLocker turned on.
3. Right-click the OU and click **Delegate Control** to open the **Delegation of Control** wizard.
4. Click **Next** to go to the **Users or Groups** page and then click **Add**.
5. In the **Select Users, Computers, or Groups** dialog box, type **SELF** as the object name and then click **OK** Once the object has been validated you will be returned to the **Users or Groups** wizard page and the SELF account will be listed. Click **Next**.
6. On the **Tasks to Delegate** page, choose **Create a custom task to delegate** and then click **Next**.
7. On the **Active Directory Object Type** page, choose **Only the following objects in the folder** and then check **Computer Objects** and then click **Next**.
8. On the **Permissions** page, for **Show these permissions**, check **General**, **Property-specific**, and **Creation/deletion of specific child objects**. Scroll down the **Permissions** list and check both **Write msTPM-OwnerInformation** and **Write msTPM-TpmInformationForComputer** then click **Next**.
9. Click **Finish** to apply the permissions settings.

FIPS support for recovery password protector

Functionality introduced in Windows Server 2012 R2 and Windows 8.1, allows BitLocker to be fully functional in FIPS mode.

Note

The United States Federal Information Processing Standard (FIPS) defines security and interoperability requirements for computer systems that are used by the U.S. federal government. The FIPS 140 standard defines approved cryptographic algorithms. The FIPS 140 standard also sets forth requirements for key generation and for key management. The National Institute of Standards and Technology (NIST) uses the Cryptographic Module Validation Program (CMVP) to determine whether a particular implementation of a cryptographic algorithm is compliant with the FIPS 140 standard. An implementation of a cryptographic algorithm is considered FIPS 140-compliant only if it has been submitted for and has passed NIST validation. An algorithm that has not been submitted cannot be considered FIPS-compliant even if the implementation produces identical data as a validated implementation of the same algorithm.

Prior to these supported versions of Windows, when Windows was in FIPS mode, BitLocker prevented the creation or use of recovery passwords and instead forced the user to use recovery keys. For more information about these issues, see the support article kb947249.

But on computers running these supported systems with BitLocker enabled:

- FIPS-compliant recovery password protectors can be created when Windows is in FIPS mode. These protectors use the FIPS 140 NIST SP800-132 algorithm.
- Recovery passwords created in FIPS mode on Windows 8.1 can be distinguished from recovery passwords created on other systems.
- Recovery unlock using the FIPS-compliant algorithm based recovery password protector work in all cases that currently work for recovery passwords.
- When FIPS-compliant recovery passwords unlock volumes, the volume is unlocked to allow read/write access even while in FIPS mode.
- FIPS-compliant recovery password protectors can be exported and stored in AD a while in FIPS mode.

To set FIPS-compliant mode for Windows, see System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing.

The BitLocker Group Policy settings for recovery passwords work the same for all Windows versions that support BitLocker, whether in FIPs mode or not.

However, you cannot use recovery passwords generated on a system in FIPS mode for systems earlier than Windows Server 2012 R2 and Windows 8.1. Recovery passwords created on Windows Server 2012 R2 and Windows 8.1 are incompatible with BitLocker on operating systems prior to Windows Server 2012 R2 and Windows 8.1; so recovery keys should be used instead.

More information

Trusted Platform Module Technology Overview

Windows 8 TPM Group Policy Settings

BitLocker Frequently Asked Questions (FAQ)

BitLocker Overview

BitLocker Group Policy Settings

BitLocker Basic Deployment

BitLocker Group Policy Settings

This reference topic for the IT professional describes the function, location, and effect of each Group Policy setting that is used to manage BitLocker Drive Encryption.

Overview

To control what drive encryption tasks the user can perform from the Windows Control Panel or to modify other configuration options, you can use Group Policy administrative templates or local computer policy settings. How you configure these policy settings depends on how you implement BitLocker and what level of user interaction will be allowed.

Note

A separate set of Group Policy settings supports the use of the Trusted Platform Module (TPM). For details about those settings, see **Trusted Platform Module Services Group Policy Settings**.

BitLocker Group Policy settings can be accessed using the Local Group Policy Editor and the Group Policy Management Console (GPMC) under **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption**.

Most of the BitLocker Group Policy settings are applied when BitLocker is initially turned on for a drive. If a computer is not compliant with existing Group Policy settings, BitLocker may not be turned on or modified until the computer is in a compliant state. When a drive is out of compliance with Group Policy settings (for example, if a Group Policy setting was changed after the initial BitLocker deployment in your organization, and then the setting was applied to previously encrypted drives), no change can be made to the BitLocker configuration of that drive except a change that will bring it into compliance.

If multiple changes are necessary to bring the drive into compliance, you must suspend BitLocker protection, make the necessary changes, and then resume protection. This situation could occur, for example, if a removable drive was initially configured to be unlocked with a password and then Group Policy settings are changed to disallow passwords and require smart cards. In this situation, you need to suspend BitLocker protection by using the **Manage-bde** command-line tool, delete the password unlock method, and add the smart card method. After this is complete, BitLocker is compliant with the Group Policy setting and BitLocker protection on the drive can be resumed.

BitLocker Group Policy settings

The following sections provide a comprehensive list of BitLocker Group Policy settings that are organized by usage. BitLocker Group Policy settings include settings for specific drive types (operating system drives, fixed data drives, and removable data drives) and settings that are applied to all drives.

The following policy settings can be used to determine how a BitLocker-protected drive can be unlocked.

- Allow network unlock at startup
- Require additional authentication at startup
- Allow enhanced PINs for startup
- Configure minimum PIN length for startup

- Disallow standard users from changing the PIN or password
- Configure use of passwords for operating system drives
- Require additional authentication at startup (Windows Server 2008 and Windows Vista)
- Configure use of smart cards on fixed data drives
- Configure use of passwords on fixed data drives
- Configure use of smart cards on removable data drives
- Configure use of passwords on removable data drives
- Validate smart card certificate usage rule compliance
- Enable use of BitLocker authentication requiring preboot keyboard input on slates

The following policy settings are used to control how users can access drives and how they can use BitLocker on their computers.

- Deny write access to fixed drives not protected by BitLocker
- Deny write access to removable drives not protected by BitLocker
- Control use of BitLocker on removable drives

The following policy settings determine the encryption methods and encryption types that are used with BitLocker.

- Choose drive encryption method and cipher strength
- Configure use of hardware-based encryption for fixed data drives
- Configure use of hardware-based encryption for operating system drives
- Configure use of hardware-based encryption for removable data drives
- Enforce drive encryption type on fixed data drives
- Enforce drive encryption type on operating system drives
- Enforce drive encryption type on removable data drives

The following policy settings define the recovery methods that can be used to restore access to a BitLocker-protected drive if an authentication method fails or is unable to be used.

- Choose how BitLocker-protected operating system drives can be recovered
- Choose how users can recover BitLocker-protected drives (Windows Server 2008 and Windows Vista)
- Store BitLocker recovery information in Active Directory Domain Services (Windows Server 2008 and Windows Vista)
- Choose default folder for recovery password
- Choose how BitLocker-protected fixed drives can be recovered
- Choose how BitLocker-protected removable drives can be recovered

The following policies are used to support customized deployment scenarios in your organization.

- Allow Secure Boot for integrity validation
- Provide the unique identifiers for your organization
- Prevent memory overwrite on restart
- Configure TPM platform validation profile for BIOS-based firmware configurations

- Configure TPM platform validation profile (Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2)
- Configure TPM platform validation profile for native UEFI firmware configurations
- Reset platform validation data after BitLocker recovery
- Use enhanced Boot Configuration Data validation profile
- Allow access to BitLocker-protected fixed data drives from earlier versions of Windows
- Allow access to BitLocker-protected removable data drives from earlier versions of Windows

Allow network unlock at startup

This policy controls a portion of the behavior of the Network Unlock feature in BitLocker. This policy is required to enable BitLocker Network Unlock on a network because it allows clients running BitLocker to create the necessary network key protector during encryption. This policy is used in addition to the BitLocker Drive Encryption Network Unlock Certificate security policy (located in the **Public Key Policies** folder of Local Computer Policy) to allow systems that are connected to a trusted network to properly utilize the Network Unlock feature.

Policy description	With this policy setting, you can control whether a BitLocker-protected computer that is connected to a trusted local area network and joined to a domain can create and use network key protectors on TPM-enabled computers to automatically unlock the operating system drive when the computer is started.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	Clients configured with a BitLocker Network Unlock certificate can create and use Network Key Protectors.
When disabled or not configured	Clients cannot create and use Network Key Protectors

Reference

To use a network key protector to unlock the computer, the computer and the server that hosts BitLocker Drive Encryption Network Unlock must be provisioned with a Network Unlock certificate. The Network Unlock certificate is used to create a network key protector and to protect

the information exchange with the server to unlock the computer. You can use the Group Policy setting **Computer Configuration\Windows Settings\Security Settings\Public Key Policies\BitLocker Drive Encryption Network Unlock Certificate** on the domain controller to distribute this certificate to computers in your organization. This unlock method uses the TPM on the computer, so computers that do not have a TPM cannot create network key protectors to automatically unlock by using Network Unlock.

Note
For reliability and security, computers should also have a TPM startup PIN that can be used when the computer is disconnected from the wired network or cannot connect to the domain controller at startup.

For more information about Network Unlock, see BitLocker: How to enable Network Unlock.

Require additional authentication at startup

This policy setting is used to control which unlock options are available for operating system drives.

Policy description	With this policy setting, you can configure whether BitLocker requires additional authentication each time the computer starts and whether you are using BitLocker with a Trusted Platform Module (TPM). This policy setting is applied when you turn on BitLocker.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	If one authentication method is required, the other methods cannot be allowed. Use of BitLocker with a TPM startup key or with a TPM startup key and a PIN must be disallowed if the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled.
When enabled	Users can configure advanced startup options in the BitLocker Setup Wizard.
When disabled or not configured	Users can configure only basic options on computers with a TPM.

	Only one of the additional authentication options can be required at startup; otherwise, a policy error occurs.

Reference

If you want to use BitLocker on a computer without a TPM, select the **Allow BitLocker without a compatible TPM** check box. In this mode, a USB drive is required for startup. Key information that is used to encrypt the drive is stored on the USB drive, which creates a USB key. When the USB key is inserted, access to the drive is authenticated and the drive is accessible. If the USB key is lost or unavailable, you need to use one of the BitLocker recovery options to access the drive.

On a computer with a compatible TPM, four types of authentication methods can be used at startup to provide added protection for encrypted data. When the computer starts, it can use:

- only the TPM for authentication
- insertion of a USB flash drive containing the startup key
- the entry of a 4-digit to 20-digit personal identification number (PIN)
- a combination of the PIN and the USB flash drive

There are four options for TPM-enabled computers or devices:

- Configure TPM startup
  - Allow TPM
  - Require TPM
  - Do not allow TPM
- Configure TPM startup PIN
  - Allow startup PIN with TPM
  - Require startup PIN with TPM
  - Do not allow startup PIN with TPM
- Configure TPM startup key
  - Allow startup key with TPM
  - Require startup key with TPM
  - Do not allow startup key with TPM
- Configure TPM startup key and PIN
  - Allow TPM startup key with PIN
  - Require startup key and PIN with TPM
  - Do not allow TPM startup key with PIN

Allow enhanced PINs for startup

This policy setting permits the use of enhanced PINs when you use an unlock method that includes a PIN.

Policy description	With this policy setting, you can configure whether enhanced startup PINs are used with BitLocker.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	All new BitLocker startup PINs that are set will be enhanced PINs. Existing drives that were protected by using standard startup PINs are not affected.
When disabled or not configured	Enhanced PINs will not be used.

Reference

Enhanced startup PINs permit the use of characters (including uppercase and lowercase letters, symbols, numbers, and spaces). This policy setting is applied when you turn on BitLocker.

Important

Not all computers support enhanced PIN characters in the preboot environment. It is strongly recommended that users perform a system check during the BitLocker setup to verify that enhanced PIN characters can be used.

Configure minimum PIN length for startup

This policy setting is used to set a minimum PIN length when you use an unlock method that includes a PIN.

Policy description	With this policy setting, you can configure a minimum length for a TPM startup PIN. This policy setting is applied when you turn on BitLocker. The startup PIN must have a minimum length of 4 digits, and it can have a maximum length of 20 digits.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative

	Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	You can require that users enter a minimum number of digits to when setting their startup PINs.
When disabled or not configured	Users can configure a startup PIN of any length between 4 and 20 digits.

Reference

This policy setting is applied when you turn on BitLocker. The startup PIN must have a minimum length of 4 digits and can have a maximum length of 20 digits.

Disallow standard users from changing the PIN or password

This policy setting allows you to configure whether standard users are allowed to change the PIN or password that is used to protect the operating system drive.

Policy description	With this policy setting, you can configure whether standard users are allowed to change the PIN or password used to protect the operating system drive.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	Standard users are not allowed to change BitLocker PINs or passwords.
When disabled or not configured	Standard users are permitted to change BitLocker PINs or passwords.

Reference

To change the PIN or password, the user must be able to provide the current PIN or password. This policy setting is applied when you turn on BitLocker.

Configure use of passwords for operating system drives

This policy controls how non-TPM based systems utilize the password protector. Used in conjunction with the **Password must meet complexity requirements** policy, this policy allows administrators to require password length and complexity for using the password protector. By default, passwords must be eight characters in length. Complexity configuration options determine how important domain connectivity is for the client. For the strongest password security, administrators should choose **Require password complexity** because it requires domain connectivity, and it requires that the BitLocker password meets the same password complexity requirements as domain sign-in passwords.

For more information, see Password must meet complexity requirements.

Policy description	With this policy setting, you can specify the constraints for passwords that are used to unlock operating system drives that are protected with BitLocker.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	Passwords cannot be used if FIPS-compliance is enabled. **Note** The **System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing** policy setting, which is located at **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options** specifies whether FIPS-compliance is enabled.
When enabled	Users can configure a password that meets the requirements you define. To enforce complexity requirements for the password, select **Require complexity**.
When disabled or not configured	The default length constraint of 8 characters will apply to operating system drive passwords and no complexity checks will occur.

Reference

If non-TPM protectors are allowed on operating system drives, you can provision a password, enforce complexity requirements on the password, and configure a minimum length for the password. For the complexity requirement setting to be effective, the Group Policy setting **Password must meet complexity requirements**, which is located at **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy** must be also enabled.

For more information, see Password must meet complexity requirements.

Note
These settings are enforced when turning on BitLocker, not when unlocking a volume. BitLocker allows unlocking a drive with any of the protectors that are available on the drive.

When set to **Require complexity**, a connection to a domain controller is necessary when BitLocker is enabled to validate the complexity the password. When set to **Allow complexity**, a connection to a domain controller is attempted to validate that the complexity adheres to the rules set by the policy. If no domain controllers are found, the password will be accepted regardless of actual password complexity, and the drive will be encrypted by using that password as a protector. When set to **Do not allow complexity**, there is no password complexity validation.

Passwords must be at least 8 characters. To configure a greater minimum length for the password, enter the desired number of characters in the **Minimum password length** box.

When this policy setting is enabled, you can set the option **Configure password complexity for operating system drives** to:

- Allow password complexity
- Do not allow password complexity
- Require password complexity

Require additional authentication at startup (Windows Server 2008 and Windows Vista)

This policy setting is used to control what unlock options are available for computers running Windows Server 2008 or Windows Vista.

Policy description	With this policy setting, you can control whether the BitLocker Setup Wizard on computers running Windows Vista or Windows Server 2008 can set up an additional authentication method that is required each time the computer starts.
Introduced	Windows Server 2008 and Windows Vista
Drive type	Operating system drives (Windows Server 2008

	and Windows Vista)
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	If you choose to require an additional authentication method, other authentication methods cannot be allowed.
When enabled	The BitLocker Setup Wizard displays the page that allows the user to configure advanced startup options for BitLocker. You can further configure setting options for computers with or without a TPM.
When disabled or not configured	The BitLocker Setup Wizard displays basic steps that allow users to enable BitLocker on computers with a TPM. In this basic wizard, no additional startup key or startup PIN can be configured.

Reference

On a computer with a compatible TPM, two authentication methods can be used at startup to provide added protection for encrypted data. When the computer starts, it can require users to insert a USB drive that contains a startup key. It can also require users to enter a 4-digit to 20-digit startup PIN.

A USB drive that contains a startup key is needed on computers without a compatible TPM. Without a TPM, BitLocker-encrypted data is protected solely by the key material that is on this USB drive.

There are two options for TPM-enabled computers or devices:

- Configure TPM startup PIN
  - Allow startup PIN with TPM
  - Require startup PIN with TPM
  - Do not allow startup PIN with TPM
- Configure TPM startup key
  - Allow startup key with TPM
  - Require startup key with TPM
  - Do not allow startup key with TPM

These options are mutually exclusive. If you require the startup key, you must not allow the startup PIN. If you require the startup PIN, you must not allow the startup key. Otherwise, a policy error will occur.

To hide the advanced page on a TPM-enabled computer or device, set these options to **Do not allow** for the startup key and for the startup PIN.

Configure use of smart cards on fixed data drives

This policy setting is used to require, allow, or deny the use of smart cards with fixed data drives.

Policy description	With this policy setting, you can specify whether smart cards can be used to authenticate user access to the BitLocker-protected fixed data drives on a computer.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	To use smart cards with BitLocker, you may also need to modify the object identifier setting in the **Computer Configuration\Administrative Templates\BitLocker Drive Encryption\Validate smart card certificate usage rule compliance** policy setting to match the object identifier of your smart card certificates.
When enabled	Smart cards can be used to authenticate user access to the drive. You can require smart card authentication by selecting the **Require use of smart cards on fixed data drives** check box.
When disabled	Users cannot use smart cards to authenticate their access to BitLocker-protected fixed data drives.
When not configured	Smart cards can be used to authenticate user access to a BitLocker-protected drive.

Reference

Note

These settings are enforced when turning on BitLocker, not when unlocking a drive. BitLocker allows unlocking a drive by using any of the protectors that are available on the drive.

Configure use of passwords on fixed data drives

This policy setting is used to require, allow, or deny the use of passwords with fixed data drives.

Policy description	With this policy setting, you can specify whether a password is required to unlock BitLocker-protected fixed data drives.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	To use password complexity, the **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\Password must meet complexity requirements** policy setting must also be enabled.
When enabled	Users can configure a password that meets the requirements you define. To require the use of a password, select **Require password for fixed data drive**. To enforce complexity requirements on the password, select **Require complexity**.
When disabled	The user is not allowed to use a password.
When not configured	Passwords are supported with the default settings, which do not include password complexity requirements and require only 8 characters.

Reference

When set to **Require complexity**, a connection to a domain controller is necessary to validate the complexity of the password when BitLocker is enabled.

When set to **Allow complexity**, a connection to a domain controller is attempted to validate that the complexity adheres to the rules set by the policy. However, if no domain controllers are found,

the password is accepted regardless of the actual password complexity, and the drive is encrypted by using that password as a protector.

When set to **Do not allow complexity**, no password complexity validation is performed.

Passwords must be at least 8 characters. To configure a greater minimum length for the password, enter the desired number of characters in the **Minimum password length** box.

Note
These settings are enforced when turning on BitLocker, not when unlocking a drive. BitLocker allows unlocking a drive with any of the protectors that are available on the drive.

For the complexity requirement setting to be effective, the Group Policy setting **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy\Password must meet complexity requirements** must also be enabled.

This policy setting is configured on a per-computer basis. This means that it applies to local user accounts and domain user accounts. Because the password filter that is used to validate password complexity is located on the domain controllers, local user accounts cannot access the password filter because they are not authenticated for domain access. When this policy setting is enabled, if you sign in with a local user account, and you attempt to encrypt a drive or change a password on an existing BitLocker-protected drive, an "Access denied" error message is displayed. In this situation, the password key protector cannot be added to the drive.

Enabling this policy setting requires that connectivity to a domain be established before adding a password key protector to a BitLocker-protected drive. Users who work remotely and have periods of time in which they cannot connect to the domain should be made aware of this requirement so that they can schedule a time when they will be connected to the domain to turn on BitLocker or to change a password on a BitLocker-protected data drive.

Important
Passwords cannot be used if FIPS compliance is enabled. The **System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing** policy setting in **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options** specifies whether FIPS compliance is enabled.

Configure use of smart cards on removable data drives

This policy setting is used to require, allow, or deny the use of smart cards with removable data drives.

Policy description	With this policy setting, you can specify whether smart cards can be used to authenticate user access to BitLocker-protected removable data drives on a computer.
Introduced	Windows Server 2008 R2 and Windows 7

Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	To use smart cards with BitLocker, you may also need to modify the object identifier setting in the **Computer Configuration\Administrative Templates\BitLocker Drive Encryption\Validate smart card certificate usage rule compliance** policy setting to match the object identifier of your smart card certificates.
When enabled	Smart cards can be used to authenticate user access to the drive. You can require smart card authentication by selecting the **Require use of smart cards on removable data drives** check box.
When disabled or not configured	Users are not allowed to use smart cards to authenticate their access to BitLocker-protected removable data drives.
When not configured	Smart cards are available to authenticate user access to a BitLocker-protected removable data drive.

Reference

Note
These settings are enforced when turning on BitLocker, not when unlocking a drive. BitLocker allows unlocking a drive with any of the protectors that are available on the drive.

Configure use of passwords on removable data drives

This policy setting is used to require, allow, or deny the use of passwords with removable data drives.

Policy description	With this policy setting, you can specify whether a password is required to unlock BitLocker-protected removable data drives.

Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	To use password complexity, the **Password must meet complexity requirements** policy setting, which is located at **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy** must also be enabled.
When enabled	Users can configure a password that meets the requirements you define. To require the use of a password, select **Require password for removable data drive**. To enforce complexity requirements on the password, select **Require complexity**.
When disabled	The user is not allowed to use a password.
When not configured	Passwords are supported with the default settings, which do not include password complexity requirements and require only 8 characters.

Reference

If you choose to allow the use of a password, you can require a password to be used, enforce complexity requirements, and configure a minimum length. For the complexity requirement setting to be effective, the Group Policy setting **Password must meet complexity requirements**, which is located at **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy** must also be enabled.

For information about this setting, see Password must meet complexity requirements.

Note
These settings are enforced when turning on BitLocker, not when unlocking a drive. BitLocker allows unlocking a drive with any of the protectors that are available on the drive.

Passwords must be at least 8 characters. To configure a greater minimum length for the password, enter the desired number of characters in the **Minimum password length** box.

When set to **Require complexity**, a connection to a domain controller is necessary when BitLocker is enabled to validate the complexity the password.

When set to **Allow complexity**, a connection to a domain controller will be attempted to validate that the complexity adheres to the rules set by the policy. However, if no domain controllers are found, the password will still be accepted regardless of actual password complexity and the drive will be encrypted by using that password as a protector.

When set to **Do not allow complexity**, no password complexity validation will be done.

Note
Passwords cannot be used if FIPS compliance is enabled. The **System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing** policy setting in **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options** specifies whether FIPS compliance is enabled.

For information about this setting, see System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing.

Validate smart card certificate usage rule compliance

This policy setting is used to determine what certificate to use with BitLocker.

Policy description	With this policy setting, you can associate an object identifier from a smart card certificate to a BitLocker-protected drive.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed and removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	None
When enabled	The object identifier that is specified in the **Object identifier** setting must match the object identifier in the smart card certificate.
When disabled or not configured	The default object identifier is used.

Reference

This policy setting is applied when you turn on BitLocker.

The object identifier is specified in the enhanced key usage (EKU) of a certificate. BitLocker can identify which certificates can be used to authenticate a user certificate to a BitLocker-protected drive by matching the object identifier in the certificate with the object identifier that is defined by this policy setting.

The default object identifier is 1.3.6.1.4.1.311.67.1.1.

Note
BitLocker does not require that a certificate have an EKU attribute; however, if one is configured for the certificate, it must be set to an object identifier that matches the object identifier configured for BitLocker.

Enable use of BitLocker authentication requiring preboot keyboard input on slates

This policy setting allows users to enable authentication options that require user input from the preboot environment even if the platform indicates a lack of preboot input capability.

Policy description	With this policy setting, you can allow users to enable authentication options that require user input from the preboot environment, even if the platform indicates a lack of preboot input capability.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drive
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drive
Conflicts	None
When enabled	Devices must have an alternative means of preboot input (such as an attached USB keyboard).
When disabled or not configured	The Windows Recovery Environment must be enabled on tablets to support entering the BitLocker recovery password.

Reference

The Windows touch keyboard (such as used by tablets) is not available in the preboot environment where BitLocker requires additional information, such as a PIN or password.

It is recommended that administrators enable this policy only for devices that are verified to have an alternative means of preboot input, such as attaching a USB keyboard.

When the Windows Recovery Environment is not enabled and this policy is not enabled, you cannot turn on BitLocker on a device that uses the Windows touch keyboard.

If you do not enable this policy setting, the following options in the **Require additional authentication at startup** policy might not be available:

- Configure TPM startup PIN: Required and Allowed

- Configure TPM startup key and PIN: Required and Allowed
- Configure use of passwords for operating system drives

Deny write access to fixed drives not protected by BitLocker

This policy setting is used to require encryption of fixed drives prior to granting Write access.

Policy description	With this policy setting, you can set whether BitLocker protection is required for fixed data drives to be writable on a computer.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	See the Reference section for a description of conflicts.
When enabled	All fixed data drives that are not BitLocker-protected are mounted as Read-only. If the drive is protected by BitLocker, it is mounted with Read and Write access.
When disabled or not configured	All fixed data drives on the computer are mounted with Read and Write access.

Reference

This policy setting is applied when you turn on BitLocker.

Conflict considerations include:

1. When this policy setting is enabled, users receive "Access denied" error messages when they try to save data to unencrypted fixed data drives. See the Reference section for additional conflicts.
2. If BdeHdCfg.exe is run on a computer when this policy setting is enabled, you could encounter the following issues:
    - If you attempted to shrink the drive and create the system drive, the drive size is successfully reduced and a raw partition is created. However, the raw partition is not formatted. The following error message is displayed: "The new active drive cannot be formatted. You may need to manually prepare your drive for BitLocker."
    - If you attempt to use unallocated space to create the system drive, a raw partition will be created. However, the raw partition will not be formatted. The following error message is displayed: "The new active drive cannot be formatted. You may need to manually prepare your drive for BitLocker."

- If you attempt to merge an existing drive into the system drive, the tool fails to copy the required boot file onto the target drive to create the system drive. The following error message is displayed: "BitLocker setup failed to copy boot files. You may need to manually prepare your drive for BitLocker."

3. If this policy setting is enforced, a hard drive cannot be repartitioned because the drive is protected. If you are upgrading computers in your organization from a previous version of Windows, and those computers were configured with a single partition, you should create the required BitLocker system partition before you apply this policy setting to the computers.

Deny write access to removable drives not protected by BitLocker

This policy setting is used to require that removable drives are encrypted prior to granting Write access, and to control whether BitLocker-protected removable drives that were configured in another organization can be opened with Write access.

Policy description	With this policy setting, you can configure whether BitLocker protection is required for a computer to be able to write data to a removable data drive.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	See the Reference section for a description of conflicts.
When enabled	All removable data drives that are not BitLocker-protected are mounted as Read-only. If the drive is protected by BitLocker, it is mounted with Read and Write access.
When disabled or not configured	All removable data drives on the computer are mounted with Read and Write access.

Reference

If the **Deny write access to devices configured in another organization** option is selected, only drives with identification fields that match the computer's identification fields are given Write access. When a removable data drive is accessed, it is checked for a valid identification field and allowed identification fields. These fields are defined by the **Provide the unique identifiers for your organization** policy setting.

Note

You can override this policy setting with the policy settings under **User Configuration\Administrative Templates\System\Removable Storage Access**. If the **Removable Disks: Deny write access** policy setting is enabled, this policy setting will be ignored.

Conflict considerations include:

1. Use of BitLocker with the TPM plus a startup key or with the TPM plus a PIN and startup key must be disallowed if the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled.
2. Use of recovery keys must be disallowed if the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled.
3. You must enable the **Provide the unique identifiers for your organization** policy setting if you want to deny Write access to drives that were configured in another organization.

Control use of BitLocker on removable drives

This policy setting is used to prevent users from turning BitLocker on or off on removable data drives.

Policy description	With this policy setting, you can control the use of BitLocker on removable data drives.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	None
When enabled	You can select property settings that control how users can configure BitLocker.
When disabled	Users cannot use BitLocker on removable data drives.
When not configured	Users can use BitLocker on removable data drives.

Reference

This policy setting is applied when you turn on BitLocker.

For information about suspending BitLocker protection, see BitLocker Basic Deployment.

The options for choosing property settings that control how users can configure BitLocker are:

- **Allow users to apply BitLocker protection on removable data drives** Enables the user to run the BitLocker Setup Wizard on a removable data drive.
- **Allow users to suspend and decrypt BitLocker on removable data drives** Enables the user to remove BitLocker from the drive or to suspend the encryption while performing maintenance.

Choose drive encryption method and cipher strength

This policy setting is used to control the encryption method and cipher strength.

Policy description	With this policy setting, you can control the encryption method and strength for drives.
Introduced	Windows Server 2012 and Windows 8
Drive type	All drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	None
When enabled	You can choose an encryption algorithm and key cipher strength for BitLocker to use to encrypt drives.
When disabled or not configured	BitLocker uses the default encryption method of AES 128-bit or the encryption method that is specified by the setup script.

Reference

By default, BitLocker uses AES 128-bit encryption. Available options are AES-128 and AES-256. The values of this policy determine the strength of the cipher that BitLocker uses for encryption. Enterprises may want to control the encryption level for increased security (AES-256 is stronger than AES-128).

Changing the encryption method has no effect if the drive is already encrypted or if encryption is in progress. In these cases, this policy setting is ignored.

Warning

This policy does not apply to encrypted drives. Encrypted drives utilize their own algorithm, which is set by the drive during partitioning.

When this policy setting is disabled, BitLocker uses AES with the same bit strength (128-bit or 256-bit) as specified in the policy setting **Choose drive encryption method and cipher strength (Windows Vista, Windows Server 2008, Windows 7)**. If neither policy is set,

BitLocker uses the default encryption method, AES-128, or the encryption method that is specified in the setup script.

Configure use of hardware-based encryption for fixed data drives

This policy controls how BitLocker reacts to systems that are equipped with encrypted drives when they are used as fixed data volumes. Using hardware-based encryption can improve the performance of drive operations that involve frequent reading or writing of data to the drive.

Policy description	With this policy setting, you can manage BitLocker's use of hardware-based encryption on fixed data drives and to specify which encryption algorithms BitLocker can use with hardware-based encryption.
Introduced	Windows Server 2012 and Windows 8
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	None
When enabled	You can specify additional options that control whether BitLocker software-based encryption is used instead of hardware-based encryption on computers that do not support hardware-based encryption. You can also specify whether you want to restrict the encryption algorithms and cipher suites that are used with hardware-based encryption.
When disabled	BitLocker cannot use hardware-based encryption with fixed data drives, and BitLocker software-based encryption is used by default when the drive in encrypted.
When not configured	BitLocker uses hardware-based encryption with the encryption algorithm that is set for the drive. If hardware-based encryption is not available, BitLocker software-based encryption is used instead.

Reference

Note

The **Choose drive encryption method and cipher strength** policy setting does not apply to hardware-based encryption.

The encryption algorithm that is used by hardware-based encryption is set when the drive is partitioned. By default, BitLocker uses the algorithm that is configured on the drive to encrypt the drive. The **Restrict encryption algorithms and cipher suites allowed for hardware-based encryption** option of this setting enables you to restrict the encryption algorithms that BitLocker can use with hardware encryption. If the algorithm that is set for the drive is not available, BitLocker disables the use of hardware-based encryption. Encryption algorithms are specified by object identifiers (OID), for example:

- Advanced Encryption Standard (AES) 128 in Cipher Block Chaining (CBC) mode OID: 2.16.840.1.101.3.4.1.2
- AES 256 in CBC mode OID: 2.16.840.1.101.3.4.1.42

Configure use of hardware-based encryption for operating system drives

This policy controls how BitLocker reacts when encrypted drives are used as operating system drives. Using hardware-based encryption can improve the performance of drive operations that involve frequent reading or writing of data to the drive.

Policy description	With this policy setting, you can manage BitLocker's use of hardware-based encryption on operating system drives and specify which encryption algorithms it can use with hardware-based encryption.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	You can specify additional options that control whether BitLocker software-based encryption is used instead of hardware-based encryption on computers that do not support hardware-based encryption. You can also specify whether you want to restrict the encryption algorithms and cipher suites that are used with hardware-based encryption.

When disabled	BitLocker cannot use hardware-based encryption with operating system drives, and BitLocker software-based encryption is used by default when the drive in encrypted.
When not configured	BitLocker uses hardware-based encryption with the encryption algorithm that is set for the drive. If hardware-based encryption is not available, BitLocker software-based encryption is used instead.

Reference

If hardware-based encryption is not available, BitLocker software-based encryption is used instead.

Note

The **Choose drive encryption method and cipher strength** policy setting does not apply to hardware-based encryption.

The encryption algorithm that is used by hardware-based encryption is set when the drive is partitioned. By default, BitLocker uses the algorithm that is configured on the drive to encrypt the drive. The **Restrict encryption algorithms and cipher suites allowed for hardware-based encryption** option of this setting enables you to restrict the encryption algorithms that BitLocker can use with hardware encryption. If the algorithm that is set for the drive is not available, BitLocker disables the use of hardware-based encryption. Encryption algorithms are specified by object identifiers (OID), for example:

- Advanced Encryption Standard (AES) 128 in Cipher Block Chaining (CBC) mode OID: 2.16.840.1.101.3.4.1.2
- AES 256 in CBC mode OID: 2.16.840.1.101.3.4.1.42

Configure use of hardware-based encryption for removable data drives

This policy controls how BitLocker reacts to encrypted drives when they are used as removable data drives. Using hardware-based encryption can improve the performance of drive operations that involve frequent reading or writing of data to the drive.

Policy description	With this policy setting, you can manage BitLocker's use of hardware-based encryption on removable data drives and specify which encryption algorithms it can use with hardware-based encryption.
Introduced	Windows Server 2012 and Windows 8

Drive type	Removable data drive
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	None
When enabled	You can specify additional options that control whether BitLocker software-based encryption is used instead of hardware-based encryption on computers that do not support hardware-based encryption. You can also specify whether you want to restrict the encryption algorithms and cipher suites that are used with hardware-based encryption.
When disabled	BitLocker cannot use hardware-based encryption with removable data drives, and BitLocker software-based encryption is used by default when the drive in encrypted.
When not configured	BitLocker uses hardware-based encryption with the encryption algorithm that is set for the drive. If hardware-based encryption is not available, BitLocker software-based encryption is used instead.

Reference

If hardware-based encryption is not available, BitLocker software-based encryption is used instead.

Note

The **Choose drive encryption method and cipher strength** policy setting does not apply to hardware-based encryption.

The encryption algorithm that is used by hardware-based encryption is set when the drive is partitioned. By default, BitLocker uses the algorithm that is configured on the drive to encrypt the drive. The **Restrict encryption algorithms and cipher suites allowed for hardware-based encryption** option of this setting enables you to restrict the encryption algorithms that BitLocker can use with hardware encryption. If the algorithm that is set for the drive is not available, BitLocker disables the use of hardware-based encryption. Encryption algorithms are specified by object identifiers (OID), for example:

- Advanced Encryption Standard (AES) 128 in Cipher Block Chaining (CBC) mode OID: 2.16.840.1.101.3.4.1.2
- AES 256 in CBC mode OID: 2.16.840.1.101.3.4.1.42

Enforce drive encryption type on fixed data drives

This policy controls whether fixed data drives utilize Used Space Only encryption or Full encryption. Setting this policy also causes the BitLocker Setup Wizard to skip the encryption options page so no encryption selection displays to the user.

Policy description	With this policy setting, you can configure the encryption type that is used by BitLocker.
Introduced	Windows Server 2012 and Windows 8
Drive type	Fixed data drive
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	None
When enabled	This policy defines the encryption type that BitLocker uses to encrypt drives, and the encryption type option is not presented in the BitLocker Setup Wizard.
When disabled or not configured	The BitLocker Setup Wizard asks the user to select the encryption type before turning on BitLocker.

Reference

This policy setting is applied when you turn on BitLocker. Changing the encryption type has no effect if the drive is already encrypted or if encryption is in progress. Choose Full encryption to require that the entire drive be encrypted when BitLocker is turned on. Choose Used Space Only encryption to require that only the portion of the drive that is used to store data is encrypted when BitLocker is turned on.

Note

This policy is ignored when you are shrinking or expanding a volume and the BitLocker driver uses the current encryption method. For example, when a drive that is using Used Space Only encryption is expanded, the new free space is not wiped as it would be for a drive that is using Full encryption. The user could wipe the free space on a Used Space Only drive by using the following command: **manage-bde -w**. If the volume is shrunk, no action is taken for the new free space.

For more information about the tool to manage BitLocker, see **Manage-bde**.

Enforce drive encryption type on operating system drives

This policy controls whether operating system drives utilize Full encryption or Used Space Only encryption. Setting this policy also causes the BitLocker Setup Wizard to skip the encryption options page, so no encryption selection displays to the user.

Policy description	With this policy setting, you can configure the encryption type that is used by BitLocker.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drive
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	The encryption type that BitLocker uses to encrypt drives is defined by this policy, and the encryption type option is not presented in the BitLocker Setup Wizard.
When disabled or not configured	The BitLocker Setup Wizard asks the user to select the encryption type before turning on BitLocker.

Reference

This policy setting is applied when you turn on BitLocker. Changing the encryption type has no effect if the drive is already encrypted or if encryption is in progress. Choose Full encryption to require that the entire drive be encrypted when BitLocker is turned on. Choose Used Space Only encryption to require that only the portion of the drive that is used to store data is encrypted when BitLocker is turned on.

Note

This policy is ignored when shrinking or expanding a volume, and the BitLocker driver uses the current encryption method. For example, when a drive that is using Used Space Only encryption is expanded, the new free space is not wiped as it would be for a drive that uses Full encryption. The user could wipe the free space on a Used Space Only drive by using the following command: **manage-bde -w**. If the volume is shrunk, no action is taken for the new free space.

For more information about the tool to manage BitLocker, see **Manage-bde**.

Enforce drive encryption type on removable data drives

This policy controls whether fixed data drives utilize Full encryption or Used Space Only encryption. Setting this policy also causes the BitLocker Setup Wizard to skip the encryption options page, so no encryption selection displays to the user.

Policy description	With this policy setting, you can configure the encryption type that is used by BitLocker.
Introduced	Windows Server 2012 and Windows 8
Drive type	Removable data drive
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	None
When enabled	The encryption type that BitLocker uses to encrypt drives is defined by this policy, and the encryption type option is not presented in the BitLocker Setup Wizard.
When disabled or not configured	The BitLocker Setup Wizard asks the user to select the encryption type before turning on BitLocker.

Reference

This policy setting is applied when you turn on BitLocker. Changing the encryption type has no effect if the drive is already encrypted or if encryption is in progress. Choose Full encryption to require that the entire drive be encrypted when BitLocker is turned on. Choose Used Space Only encryption to require that only the portion of the drive that is used to store data is encrypted when BitLocker is turned on.

Note

This policy is ignored when shrinking or expanding a volume, and the BitLocker driver uses the current encryption method. For example, when a drive that is using Used Space Only encryption is expanded, the new free space is not wiped as it would be for a drive that is using Full Encryption. The user could wipe the free space on a Used Space Only drive by using the following command: **manage-bde -w**. If the volume is shrunk, no action is taken for the new free space.

For more information about the tool to manage BitLocker, see **Manage-bde**.

Choose how BitLocker-protected operating system drives can be recovered

This policy setting is used to configure recovery methods for operating system drives.

Policy description	With this policy setting, you can control how BitLocker-protected operating system drives are recovered in the absence of the required startup key information.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	You must disallow the use of recovery keys if the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled. When using data recovery agents, you must enable the **Provide the unique identifiers for your organization** policy setting.
When enabled	You can control the methods that are available to users to recover data from BitLocker-protected operating system drives.
When disabled or not configured	The default recovery options are supported for BitLocker recovery. By default, a data recovery agent is allowed, the recovery options can be specified by the user (including the recovery password and recovery key), and recovery information is not backed up to AD DS.

Reference

This policy setting is applied when you turn on BitLocker.

The **Allow data recovery agent** check box is used to specify whether a data recovery agent can be used with BitLocker-protected operating system drives. Before a data recovery agent can be used, it must be added from **Public Key Policies**, which is located in the Group Policy Management Console (GPMC) or in the Local Group Policy Editor.

For more information about adding data recovery agents, see BitLocker Basic Deployment.

In **Configure user storage of BitLocker recovery information**, select whether users are allowed, required, or not allowed to generate a 48-digit recovery password.

Select **Omit recovery options from the BitLocker setup wizard** to prevent users from specifying recovery options when they enable BitLocker on a drive. This means that you will not be able to specify which recovery option to use when you enable BitLocker. Instead, BitLocker recovery options for the drive are determined by the policy setting.

In **Save BitLocker recovery information to Active Directory Domain Services**, choose which BitLocker recovery information to store in Active Directory Domain Services (AD DS) for operating system drives. If you select **Store recovery password and key packages**, the BitLocker recovery password and the key package are stored in AD DS. Storing the key package supports recovering data from a drive that is physically corrupted. If you select **Store recovery password only**, only the recovery password is stored in AD DS.

Select the **Do not enable BitLocker until recovery information is stored in AD DS for operating system drives** check box if you want to prevent users from enabling BitLocker unless the computer is connected to the domain and the backup of BitLocker recovery information to AD DS succeeds.

Note
If the **Do not enable BitLocker until recovery information is stored in AD DS for operating system drives** check box is selected, a recovery password is automatically generated.

Choose how users can recover BitLocker-protected drives (Windows Server 2008 and Windows Vista)

This policy setting is used to configure recovery methods for BitLocker-protected drives on computers running Windows Server 2008 or Windows Vista.

Policy description	With this policy setting, you can control whether the BitLocker Setup Wizard can display and specify BitLocker recovery options.
Introduced	Windows Server 2008 and Windows Vista
Drive type	Operating system drives and fixed data drives on computers running Windows Server 2008 and Windows Vista
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	This policy setting provides an administrative method of recovering data that is encrypted by BitLocker to prevent data loss due to lack of

	key information. If you choose the **Do not allow** option for both user recovery options, you must enable the **Store BitLocker recovery information in Active Directory Domain Services (Windows Server 2008 and Windows Vista)** policy setting to prevent a policy error.
When enabled	You can configure the options that the Bitlocker Setup Wizard displays to users for recovering BitLocker encrypted data.
When disabled or not configured	The BitLocker Setup Wizard presents users with ways to store recovery options.

Reference

This policy is only applicable to computers running Windows Server 2008 or Windows Vista. This policy setting is applied when you turn on BitLocker.

Two recovery options can be used to unlock BitLocker-encrypted data in the absence of the required startup key information. Users can type a 48-digit numerical recovery password, or they can insert a USB drive that contains a 256-bit recovery key.

Saving the recovery password to a USB drive stores the 48-digit recovery password as a text file and the 256-bit recovery key as a hidden file. Saving it to a folder stores the 48-digit recovery password as a text file. Printing it sends the 48-digit recovery password to the default printer. For example, not allowing the 48-digit recovery password prevents users from printing or saving recovery information to a folder.

Important

If TPM initialization is performed during the BitLocker setup, TPM owner information is saved or printed with the BitLocker recovery information.

The 48-digit recovery password is not available in FIPS-compliance mode.

Important

To prevent data loss, you must have a way to recover BitLocker encryption keys. If you do not allow both recovery options, you must enable the backup of BitLocker recovery information to AD DS. Otherwise, a policy error occurs.

Store BitLocker recovery information in Active Directory Domain Services (Windows Server 2008 and Windows Vista)

This policy setting is used to configure the storage of BitLocker recovery information in AD DS. This provides an administrative method of recovering data that is encrypted by BitLocker to prevent data loss due to lack of key information.

Policy description	With this policy setting, you can manage the AD DS backup of BitLocker Drive Encryption recovery information.
Introduced	Windows Server 2008 and Windows Vista
Drive type	Operating system drives and fixed data drives on computers running Windows Server 2008 and Windows Vista.
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	None
When enabled	BitLocker recovery information is automatically and silently backed up to AD DS when BitLocker is turned on for a computer.
When disabled or not configured	BitLocker recovery information is not backed up to AD DS.

Reference

This policy is only applicable to computers running Windows Server 2008 or Windows Vista.

This policy setting is applied when you turn on BitLocker.

BitLocker recovery information includes the recovery password and unique identifier data. You can also include a package that contains an encryption key for a BitLocker-protected drive. This key package is secured by one or more recovery passwords, and it can help perform specialized recovery when the disk is damaged or corrupted.

If you select **Require BitLocker backup to AD DS**, BitLocker cannot be turned on unless the computer is connected to the domain and the backup of BitLocker recovery information to AD DS succeeds. This option is selected by default to help ensure that BitLocker recovery is possible.

A recovery password is a 48-digit number that unlocks access to a BitLocker-protected drive. A key package contains a drive's BitLocker encryption key, which is secured by one or more recovery passwords. Key packages may help perform specialized recovery when the disk is damaged or corrupted.

If the **Require BitLocker backup to AD DS** option is not selected, AD DS backup is attempted, but network or other backup failures do not prevent the BitLocker setup. The Backup process is not automatically retried, and the recovery password might not be stored in AD DS during BitLocker setup.

TPM initialization might be needed during the BitLocker setup. Enable the **Turn on TPM backup to Active Directory Domain Services** policy setting in **Computer**

Configuration\Administrative Templates\System\Trusted Platform Module Services to ensure that TPM information is also backed up.

For more information about this setting, see **Turn on TPM backup to Active Directory Domain Services**.

If you are using domain controllers running Windows Server 2003 with Service Pack 1, you must first set up appropriate schema extensions and access control settings on the domain before a backup to AD DS can succeed. For more information, see **Backing Up BitLocker and TPM Recovery Information to AD DS**.

Choose default folder for recovery password

This policy setting is used to configure the default folder for recovery passwords.

Policy description	With this policy setting, you can specify the default path that is displayed when the BitLocker Setup Wizard prompts the user to enter the location of a folder in which to save the recovery password.
Introduced	Windows Vista
Drive type	All drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	None
When enabled	You can specify the path that will be used as the default folder location when the user chooses the option to save the recovery password in a folder. You can specify a fully qualified path or include the target computer's environment variables in the path. If the path is not valid, the BitLocker Setup Wizard displays the computer's top-level folder view.
When disabled or not configured	The BitLocker Setup Wizard displays the computer's top-level folder view when the user chooses the option to save the recovery password in a folder.

Reference

This policy setting is applied when you turn on BitLocker.

Note
This policy setting does not prevent the user from saving the recovery password in another folder.

Choose how BitLocker-protected fixed drives can be recovered

This policy setting is used to configure recovery methods for fixed data drives.

Policy description	With this policy setting, you can control how BitLocker-protected fixed data drives are recovered in the absence of the required credentials.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	You must disallow the use of recovery keys if the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled. When using data recovery agents, you must enable and configure the **Provide the unique identifiers for your organization** policy setting.
When enabled	You can control the methods that are available to users to recover data from BitLocker-protected fixed data drives.
When disabled or not configured	The default recovery options are supported for BitLocker recovery. By default, a data recovery agent is allowed, the recovery options can be specified by the user (including the recovery password and recovery key), and recovery information is not backed up to AD DS.

Reference

This policy setting is applied when you turn on BitLocker.

The **Allow data recovery agent** check box is used to specify whether a data recovery agent can be used with BitLocker-protected fixed data drives. Before a data recovery agent can be used, it

must be added from **Public Key Policies**, which is located in the Group Policy Management Console (GPMC) or in the Local Group Policy Editor.

In **Configure user storage of BitLocker recovery information**, select whether users are allowed, required, or not allowed to generate a 48-digit recovery password or a 256-bit recovery key.

Select **Omit recovery options from the BitLocker setup wizard** to prevent users from specifying recovery options when they enable BitLocker on a drive. This means that you cannot specify which recovery option to use when you enable BitLocker. Instead, BitLocker recovery options for the drive are determined by the policy setting.

In **Save BitLocker recovery information to Active Directory Doman Services**, choose which BitLocker recovery information to store in AD DS for fixed data drives. If you select **Backup recovery password and key package**, the BitLocker recovery password and the key package are stored in AD DS. Storing the key package supports recovering data from a drive that has been physically corrupted. To recover this data, you can use the **Repair-bde** command-line tool. If you select **Backup recovery password only**, only the recovery password is stored in AD DS.

For more information about the BitLocker repair tool, see **Repair-bde**.

Select the **Do not enable BitLocker until recovery information is stored in AD DS for fixed data drives** check box if you want to prevent users from enabling BitLocker unless the computer is connected to the domain and the backup of BitLocker recovery information to AD DS succeeds.

Note
If the **Do not enable BitLocker until recovery information is stored in AD DS for fixed data drives** check box is selected, a recovery password is automatically generated.

Choose how BitLocker-protected removable drives can be recovered

This policy setting is used to configure recovery methods for removable data drives.

Policy description	With this policy setting, you can control how BitLocker-protected removable data drives are recovered in the absence of the required credentials.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	You must disallow the use of recovery keys if

	the **Deny write access to removable drives not protected by BitLocker** policy setting is enabled. When using data recovery agents, you must enable and configure the **Provide the unique identifiers for your organization** policy setting.
When enabled	You can control the methods that are available to users to recover data from BitLocker-protected removable data drives.
When disabled or not configured	The default recovery options are supported for BitLocker recovery. By default, a data recovery agent is allowed, the recovery options can be specified by the user (including the recovery password and recovery key), and recovery information is not backed up to AD DS.

Reference

This policy setting is applied when you turn on BitLocker.

The **Allow data recovery agent** check box is used to specify whether a data recovery agent can be used with BitLocker-protected removable data drives. Before a data recovery agent can be used, it must be added from **Public Key Policies** , which is accessed using the GPMC or the Local Group Policy Editor.

In **Configure user storage of BitLocker recovery information**, select whether users are allowed, required, or not allowed to generate a 48-digit recovery password.

Select **Omit recovery options from the BitLocker setup wizard** to prevent users from specifying recovery options when they enable BitLocker on a drive. This means that you cannot specify which recovery option to use when you enable BitLocker. Instead, BitLocker recovery options for the drive are determined by the policy setting.

In **Save BitLocker recovery information to Active Directory Domain Services**, choose which BitLocker recovery information to store in AD DS for removable data drives. If you select **Backup recovery password and key package**, the BitLocker recovery password and the key package are stored in AD DS. If you select **Backup recovery password only**, only the recovery password is stored in AD DS.

Select the **Do not enable BitLocker until recovery information is stored in AD DS for removable data drives** check box if you want to prevent users from enabling BitLocker unless the computer is connected to the domain and the backup of BitLocker recovery information to AD DS succeeds.

Note

If the **Do not enable BitLocker until recovery information is stored in AD DS for fixed data drives** check box is selected, a recovery password is automatically generated.

Allow Secure Boot for integrity validation

This policy controls how BitLocker-enabled system volumes are handled in conjunction with the Secure Boot feature. Enabling this feature forces Secure Boot validation during the boot process and verifies Boot Configuration Data (BCD) settings according to the Secure Boot policy.

Policy description	With this policy setting, you can configure whether Secure Boot will be allowed as the platform integrity provider for BitLocker operating system drives.
Introduced	Windows Server 2012 and Windows 8
Drive type	All drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	If the **Configure TPM platform validation profile for native UEFI firmware configurations** Group Policy setting is enabled and PCR 7 is omitted, BitLocker is prevented from using Secure Boot for platform or BCD integrity validation. For more information about PCR 7, see About the Platform Configuration Register (PCR) in this topic.
When enabled	BitLocker uses Secure Boot for platform integrity if the platform is capable of Secure Boot-based integrity validation.
When disabled or not configured	BitLocker uses legacy platform integrity validation, even on systems that are capable of Secure Boot-based integrity validation.

Reference

Secure Boot ensures that the computer's preboot environment loads only firmware that is digitally signed by authorized software publishers. Secure Boot also provides more flexibility for managing preboot configurations than BitLocker integrity checks prior to Windows Server 2012 and Windows 8.

When this policy is enabled and the hardware is capable of using Secure Boot for BitLocker scenarios, the **Use enhanced Boot Configuration Data validation profile** Group Policy setting is ignored, and Secure Boot verifies BCD settings according to the Secure Boot policy setting, which is configured separately from BitLocker.

Warning
Enabling this policy might result in BitLocker recovery when manufacturer-specific firmware is updated. If you disable this policy, suspend BitLocker prior to applying firmware updates.

Provide the unique identifiers for your organization

This policy setting is used to establish an identifier that is applied to all drives that are encrypted in your organization.

Policy description	With this policy setting, you can associate unique organizational identifiers to a new drive that is enabled with BitLocker.
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	All drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	Identification fields are required to manage certificate-based data recovery agents on BitLocker-protected drives. BitLocker manages and updates certificate-based data recovery agents only when the identification field is present on a drive and it is identical to the value that is configured on the computer.
When enabled	You can configure the identification field on the BitLocker-protected drive and any allowed identification field that is used by your organization.
When disabled or not configured	The identification field is not required.

Reference

These identifiers are stored as the identification field and the allowed identification field. The identification field allows you to associate a unique organizational identifier to BitLocker-protected

drives. This identifier is automatically added to new BitLocker-protected drives, and it can be updated on existing BitLocker-protected drives by using the **Manage-bde** command-line tool.

An identification field is required to manage certificate-based data recovery agents on BitLocker-protected drives and for potential updates to the BitLocker To Go Reader. BitLocker manages and updates data recovery agents only when the identification field on the drive matches the value that is configured in the identification field. In a similar manner, BitLocker updates the BitLocker To Go Reader only when the identification field on the drive matches the value that is configured for the identification field.

For more information about the tool to manage BitLocker, see **Manage-bde**.

The allowed identification field is used in combination with the **Deny write access to removable drives not protected by BitLocker** policy setting to help control the use of removable drives in your organization. It is a comma-separated list of identification fields from your organization or external organizations.

You can configure the identification fields on existing drives by using the **Manage-bde** command-line tool.

When a BitLocker-protected drive is mounted on another BitLocker-enabled computer, the identification field and the allowed identification field are used to determine whether the drive is from an outside organization.

Multiple values separated by commas can be entered in the identification and allowed identification fields. The identification field can be any value up to 260 characters.

Prevent memory overwrite on restart

This policy setting is used to control whether the computer's memory will be overwritten the next time the computer is restarted.

Policy description	With this policy setting, you can control computer restart performance at the risk of exposing BitLocker secrets.
Introduced	Windows Vista
Drive type	All drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption
Conflicts	None
When enabled	The computer will not overwrite memory when it restarts. Preventing memory overwrite may improve restart performance, but it increases the risk of exposing BitLocker secrets.

When disabled or not configured	BitLocker secrets are removed from memory when the computer restarts.

Reference

This policy setting is applied when you turn on BitLocker. BitLocker secrets include key material that is used to encrypt data. This policy setting applies only when BitLocker protection is enabled.

Configure TPM platform validation profile for BIOS-based firmware configurations

This policy setting determines what values the TPM measures when it validates early boot components before it unlocks an operating system drive on a computer with a BIOS configuration or with UEFI firmware that has the Compatibility Support Module (CSM) enabled.

Policy description	With this policy setting, you can configure how the computer's TPM security hardware secures the BitLocker encryption key.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	You can configure the boot components that the TPM validates before unlocking access to the BitLocker-encrypted operating system drive. If any of these components change while BitLocker protection is in effect, the TPM does not release the encryption key to unlock the drive. Instead, the computer displays the BitLocker Recovery console and requires that the recovery password or the recovery key is provided to unlock the drive.
When disabled or not configured	The TPM uses the default platform validation profile or the platform validation profile that is specified by the setup script.

Reference

This policy setting does not apply if the computer does not have a compatible TPM or if BitLocker has already been turned on with TPM protection.

Important

This Group Policy setting only applies to computers with BIOS configurations or to computers with UEFI firmware with the CSM enabled. Computers that use a native UEFI firmware configuration store different values in the Platform Configuration Registers (PCRs). Use the **Configure TPM platform validation profile for native UEFI firmware configurations** Group Policy setting to configure the TPM PCR profile for computers that use native UEFI firmware.

A platform validation profile consists of a set of PCR indices that range from 0 to 23. The default platform validation profile secures the encryption key against changes to the following:

- Core Root of Trust of Measurement (CRTM), BIOS, and Platform Extensions (PCR 0)
- Option ROM Code (PCR 2)
- Master Boot Record (MBR) Code (PCR 4)
- NTFS Boot Sector (PCR 8)
- NTFS Boot Block (PCR 9)
- Boot Manager (PCR 10)
- BitLocker Access Control (PCR 11)

Note

Changing from the default platform validation profile affects the security and manageability of your computer. BitLocker's sensitivity to platform modifications (malicious or authorized) is increased or decreased depending on inclusion or exclusion (respectively) of the PCRs.

The following list identifies all of the PCRs available:

- PCR 0: Core root-of-trust for measurement, BIOS, and Platform extensions
- PCR 1: Platform and motherboard configuration and data.
- PCR 2: Option ROM code
- PCR 3: Option ROM data and configuration
- PCR 4: Master Boot Record (MBR) code
- PCR 5: Master Boot Record (MBR) partition table
- PCR 6: State transition and wake events
- PCR 7: Computer manufacturer-specific
- PCR 8: NTFS boot sector
- PCR 9: NTFS boot block
- PCR 10: Boot manager
- PCR 11: BitLocker access control
- PCR 12-23: Reserved for future use

Configure TPM platform validation profile (Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2)

This policy setting determines what values the TPM measures when it validates early boot components before unlocking a drive on a computer running Windows Vista, Windows Server 2008, or Windows 7.

Policy description	With this policy setting, you can configure how the computer's TPM security hardware secures the BitLocker encryption key.
Introduced	Windows Server 2008 and Windows Vista
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	You can configure the boot components that the TPM validates before unlocking access to the BitLocker-encrypted operating system drive. If any of these components change while BitLocker protection is in effect, the TPM does not release the encryption key to unlock the drive. Instead, the computer displays the BitLocker Recovery console and requires that the recovery password or the recovery key is provided to unlock the drive.
When disabled or not configured	The TPM uses the default platform validation profile or the platform validation profile that is specified by the setup script.

Reference

This policy setting does not apply if the computer does not have a compatible TPM or if BitLocker is already turned on with TPM protection.

A platform validation profile consists of a set of PCR indices that range from 0 to 23. The default platform validation profile secures the encryption key against changes to the following:

- Core Root of Trust of Measurement (CRTM), BIOS, and Platform Extensions (PCR 0)
- Option ROM Code (PCR 2)
- Master Boot Record (MBR) Code (PCR 4)
- NTFS Boot Sector (PCR 8)

- NTFS Boot Block (PCR 9)
- Boot Manager (PCR 10)
- BitLocker Access Control (PCR 11)

Note
The default TPM validation profile PCR settings for computers that use an Extensible Firmware Interface (EFI) are the PCRs 0, 2, 4, and 11 only.

The following list identifies all of the PCRs available:

- PCR 0: Core root-of-trust for measurement, EFI boot and run-time services, EFI drivers embedded in system ROM, ACPI static tables, embedded SMM code, and BIOS code
- PCR 1: Platform and motherboard configuration and data. Hand-off tables and EFI variables that affect system configuration
- PCR 2: Option ROM code
- PCR 3: Option ROM data and configuration
- PCR 4: Master Boot Record (MBR) code or code from other boot devices
- PCR 5: Master Boot Record (MBR) partition table. Various EFI variables and the GPT table
- PCR 6: State transition and wake events
- PCR 7: Computer manufacturer-specific
- PCR 8: NTFS boot sector
- PCR 9: NTFS boot block
- PCR 10: Boot manager
- PCR 11: BitLocker access control
- PCR 12 - 23: Reserved for future use

Warning
Changing from the default platform validation profile affects the security and manageability of your computer. BitLocker's sensitivity to platform modifications (malicious or authorized) is increased or decreased depending on inclusion or exclusion (respectively) of the PCRs.

Configure TPM platform validation profile for native UEFI firmware configurations

This policy setting determines what values the TPM measures when it validates early boot components before unlocking an operating system drive on a computer with native UEFI firmware configurations.

Policy description	With this policy setting, you can configure how the computer's Trusted Platform Module (TPM) security hardware secures the BitLocker encryption key.

Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	Setting this policy with PCR 7 omitted, overrides the **Allow Secure Boot for integrity validation** Group Policy setting, and it prevents BitLocker from using Secure Boot for platform or Boot Configuration Data (BCD) integrity validation. If your environments use TPM and Secure Boot for platform integrity checks, this policy should not be configured. For more information about PCR 7, see About the Platform Configuration Register (PCR) in this topic.
When enabled	Before you turn on BitLocker, you can configure the boot components that the TPM validates before it unlocks access to the BitLocker-encrypted operating system drive. If any of these components change while BitLocker protection is in effect, the TPM does not release the encryption key to unlock the drive. Instead, the computer displays the BitLocker Recovery console and requires that the recovery password or the recovery key is provided to unlock the drive.
When disabled or not configured	BitLocker uses the default platform validation profile or the platform validation profile that is specified by the setup script.

Reference

This policy setting does not apply if the computer does not have a compatible TPM or if BitLocker is already turned on with TPM protection.

Important

This Group Policy setting only applies to computers with a native UEFI firmware configuration. Computers with BIOS or UEFI firmware with a Compatibility Support Module (CSM) enabled store different values in the Platform Configuration Registers

(PCRs). Use the **Configure TPM platform validation profile for BIOS-based firmware configurations** Group Policy setting to configure the TPM PCR profile for computers with BIOS configurations or for computers with UEFI firmware with a CSM enabled.

A platform validation profile consists of a set of Platform Configuration Register (PCR) indices ranging from 0 to 23. The default platform validation profile secures the encryption key against changes to the core system firmware executable code (PCR 0), extended or pluggable executable code (PCR 2), boot manager (PCR 4), and the BitLocker access control (PCR 11).

The following list identifies all of the PCRs available:

- PCR 0: Core System Firmware executable code
- PCR 1: Core System Firmware data
- PCR 2: Extended or pluggable executable code
- PCR 3: Extended or pluggable firmware data
- PCR 4: Boot Manager
- PCR 5: GPT/Partition Table
- PCR 6: Resume from S4 and S5 Power State Events
- PCR 7: Secure Boot State

  For more information about this PCR, see About the Platform Configuration Register (PCR) in this topic.
- PCR 8: Initialized to 0 with no Extends (reserved for future use)
- PCR 9: Initialized to 0 with no Extends (reserved for future use)
- PCR 10: Initialized to 0 with no Extends (reserved for future use)
- PCR 11: BitLocker access control
- PCR 12: Data events and highly volatile events
- PCR 13: Boot Module Details
- PCR 14: Boot Authorities
- PCR 15 – 23: Reserved for future use

Warning

Changing from the default platform validation profile affects the security and manageability of your computer. BitLocker's sensitivity to platform modifications (malicious or authorized) is increased or decreased depending on inclusion or exclusion (respectively) of the PCRs.

Reset platform validation data after BitLocker recovery

This policy setting determines if you want platform validation data to refresh when Windows is started following a BitLocker recovery. A platform validation data profile consists of the values in a set of Platform Configuration Register (PCR) indices that range from 0 to 23.

Policy description	With this policy setting, you can control whether platform validation data is refreshed when

	Windows is started following a BitLocker recovery.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	None
When enabled	Platform validation data is refreshed when Windows is started following a BitLocker recovery.
When disabled	Platform validation data is not refreshed when Windows is started following a BitLocker recovery.
When not configured	Platform validation data is refreshed when Windows is started following a BitLocker recovery.

Reference

For more information about the recovery process, see the BitLocker Recovery Guide.

Use enhanced Boot Configuration Data validation profile

This policy setting determines specific Boot Configuration Data (BCD) settings to verify during platform validation. A platform validation uses the data in the platform validation profile, which consists of a set of Platform Configuration Register (PCR) indices that range from 0 to 23.

Policy description	With this policy setting, you can specify Boot Configuration Data (BCD) settings to verify during platform validation.
Introduced	Windows Server 2012 and Windows 8
Drive type	Operating system drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives
Conflicts	When BitLocker is using Secure Boot for platform and Boot Configuration Data integrity

	validation, the **Use enhanced Boot Configuration Data validation profile** Group Policy setting is ignored (as defined by the **Allow Secure Boot for integrity validation** Group Policy setting).
When enabled	You can add additional BCD settings, exclude the BCD settings you specify, or combine inclusion and exclusion lists to create a customized BCD validation profile, which gives you the ability to verify those BCD settings.
When disabled	The computer reverts to a BCD profile validation similar to the default BCD profile that is used by Windows 7.
When not configured	The computer verifies the default BCD settings in Windows.

Reference

Note
The setting that controls boot debugging (0x16000010) is always validated, and it has no effect if it is included in the inclusion or the exclusion list.

Allow access to BitLocker-protected fixed data drives from earlier versions of Windows

This policy setting is used to control whether access to drives is allowed by using the BitLocker To Go Reader, and if the application is installed on the drive.

Policy description	With this policy setting, you can configure whether fixed data drives that are formatted with the FAT file system can be unlocked and viewed on computers running Windows Vista, Windows XP with Service Pack 3 (SP3), or Windows XP with Service Pack 2 (SP2).
Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Fixed data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Fixed Data Drives
Conflicts	None

When enabled and **When not configured**	Fixed data drives that are formatted with the FAT file system can be unlocked on computers running Windows Server 2008, Windows Vista, Windows XP with SP3, or Windows XP with SP2, and their content can be viewed. These operating systems have Read-only access to BitLocker-protected drives.
When disabled	Fixed data drives that are formatted with the FAT file system and are BitLocker-protected cannot be unlocked on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2. BitLocker To Go Reader (bitlockertogo.exe) is not installed.

Reference

Note

This policy setting does not apply to drives that are formatted with the NTFS file system.

When this policy setting is enabled, select the **Do not install BitLocker To Go Reader on FAT formatted fixed drives** check box to help prevent users from running BitLocker To Go Reader from their fixed drives. If BitLocker To Go Reader (bitlockertogo.exe) is present on a drive that does not have an identification field specified, or if the drive has the same identification field as specified in the **Provide unique identifiers for your organization** policy setting, the user is prompted to update BitLocker, and BitLocker To Go Reader is deleted from the drive. In this situation, for the fixed drive to be unlocked on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2, BitLocker To Go Reader must be installed on the computer. If this check box is not selected, BitLocker To Go Reader will be installed on the fixed drive to enable users to unlock the drive on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2.

Allow access to BitLocker-protected removable data drives from earlier versions of Windows

This policy setting controls access to removable data drives that are using the BitLocker To Go Reader and whether the BitLocker To Go Reader can be installed on the drive.

Policy description	With this policy setting, you can configure whether removable data drives that are formatted with the FAT file system can be unlocked and viewed on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2.

Introduced	Windows Server 2008 R2 and Windows 7
Drive type	Removable data drives
Policy path	Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Removable Data Drives
Conflicts	None
When enabled and **When not configured**	Removable data drives that are formatted with the FAT file system can be unlocked on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2, and their content can be viewed. These operating systems have Read-only access to BitLocker-protected drives.
When disabled	Removable data drives that are formatted with the FAT file system that are BitLocker-protected cannot be unlocked on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2. BitLocker To Go Reader (bitlockertogo.exe) is not installed.

Reference

Note

This policy setting does not apply to drives that are formatted with the NTFS file system.

When this policy setting is enabled, select the **Do not install BitLocker To Go Reader on FAT formatted removable drives** check box to help prevent users from running BitLocker To Go Reader from their removable drives. If BitLocker To Go Reader (bitlockertogo.exe) is present on a drive that does not have an identification field specified, or if the drive has the same identification field as specified in the **Provide unique identifiers for your organization** policy setting, the user will be prompted to update BitLocker, and BitLocker To Go Reader is deleted from the drive. In this situation, for the removable drive to be unlocked on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2, BitLocker To Go Reader must be installed on the computer. If this check box is not selected, BitLocker To Go Reader will be installed on the removable drive to enable users to unlock the drive on computers running Windows Vista, Windows XP with SP3, or Windows XP with SP2 that do not have BitLocker To Go Reader installed.

FIPS setting

You can configure the Federal Information Processing Standard (FIPS) setting for FIPS compliance. As an effect of FIPS compliance, users cannot create or save a BitLocker password for recovery or as a key protector. The use of a recovery key is permitted.

Policy description	Notes
Introduced	Windows Server 2003 with SP1
Drive type	System-wide
Policy path	Local Policies\Security Options\System cryptography: **Use FIPS compliant algorithms for encryption, hashing, and signing**
Conflicts	Some applications, such as Terminal Services, do not support FIPS-140 on all operating systems.
When enabled	Users will be unable to save a recovery password to any location. This includes AD DS and network folders. In addition, you cannot use WMI or the BitLocker Drive Encryption Setup izard to create a recovery password.
When disabled or not configured	No BitLocker encryption key is generated

Reference

This policy needs to be enabled before any encryption key is generated for BitLocker. Note that when this policy is enabled, BitLocker prevents creating or using recovery passwords, so recovery keys should be used instead.

You can save the optional recovery key to a USB drive. Because recovery passwords cannot be saved to AD DS when FIPS is enabled, an error is caused if AD DS backup is required by Group Policy.

You can edit the FIPS setting by using the Security Policy Editor (Secpol.msc) or by editing the Windows registry. You must be an administrator to perform these procedures.

For more information about setting this policy, see System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing.

Power management Group Policy settings: Sleep and Hibernate

PCs default power settings for a computer will cause the computer to enter Sleep mode frequently to conserve power when idle and to help extend the system's battery life. When a

computer transitions to Sleep, open programs and documents are persisted in memory. When a computer resumes from Sleep, users are not required to re-authenticate with a PIN or USB startup key to access encrypted data. This might lead to conditions where data security is compromised.

However, when a computer hibernates the drive is locked, and when it resumes from hibernation the drive is unlocked, which means that users will need to provide a PIN or a startup key if using multifactor authentication with BitLocker. Therefore, organizations that use BitLocker may want to use Hibernate instead of Sleep for improved security. This setting does not have an impact on TPM-only mode, because it provides a transparent user experience at startup and when resuming from the Hibernate states.

You can use disable the following Group Policy settings, which are located in **Computer Configuration\Administrative Templates\System\Power Management** to disable all available sleep states:

- Allow Standby States (S1-S3) When Sleeping (Plugged In)
- Allow Standby States (S1-S3) When Sleeping (Battery)

About the Platform Configuration Register (PCR)

A platform validation profile consists of a set of PCR indices that range from 0 to 23. The scope of the values can be specific to the version of the operating system.

Changing from the default platform validation profile affects the security and manageability of your computer. BitLocker's sensitivity to platform modifications (malicious or authorized) is increased or decreased depending on inclusion or exclusion (respectively) of the PCRs.

About PCR 7

PCR 7 measures the state of Secure Boot. With PCR 7, BitLocker can leverage Secure Boot for integrity validation. Secure Boot ensures that the computer's preboot environment loads only firmware that is digitally signed by authorized software publishers. PCR 7 measurements indicate whether Secure Boot is on and which keys are trusted on the platform. If Secure Boot is on and the firmware measures PCR 7 correctly per the UEFI specification, BitLocker can bind to this information rather than to PCRs 0, 2, and 4 which have the measurements of the exact firmware and Bootmgr images loaded. This reduces the likelihood of BitLocker starting in recovery mode as a result of firmware and image updates, and it provides you with greater flexibility to manage the preboot configuration.

PCR 7 measurements must follow the guidance that is described in Appendix A Trusted Execution Environment EFI Protocol.

PCR 7 measurements are a mandatory logo requirement for systems that support InstantGo (also known as Always On, Always Connected PCs), such as the Microsoft Surface RT. On such systems, if the TPM with PCR 7 measurement and Secure Boot are correctly configured, BitLocker binds to PCR 7 and PCR 11 by default.

See also

Trusted Platform Module Technology Overview

Windows 8 TPM Group Policy Settings

BitLocker Frequently Asked Questions (FAQ)

BitLocker Overview

Prepare your organization for BitLocker: Planning and Policies

BitLocker: How to deploy on Windows Server 2012

For all Windows Server editions, BitLocker must be installed using Server Manager. However, you can still provision BitLocker before the server operating system is installed as part of your deployment.

Installing BitLocker

BitLocker requires administrator privileges on the server to install. You can install BitLocker either by using Server Manager or Windows PowerShell cmdlets.

- To install BitLocker using Server Manager
- To install BitLocker using Windows PowerShell

To install BitLocker using Server Manager

1. Open Server Manager by selecting the Server Manager icon or running servermanager.exe.
2. Select **Manage** from the **Server Manager Navigation** bar and select **Add Roles and Features** to start the **Add Roles and Features Wizard.**
3. With the **Add Roles and Features Wizard** open, select **Next** at the **Before you begin** pane (if shown).
4. Select **Role-based or feature-based installation** on the **Installation type** pane of the **Add Roles and Features Wizard** pane and select **Next** to continue.
5. Select the **Select a server from the server pool option** in the **Server Selection** pane and confirm the server for the BitLocker feature install.
6. Server roles and features install using the same wizard in Server Manager. Select **Next** on the **Server Roles** pane of the **Add Roles and Features** wizard to proceed to the **Features** pane.
7. Select the check box next to **BitLocker Drive Encryption** within the **Features** pane of the **Add Roles and Features Wizard**. The wizard will show the additional management features available for BitLocker. If you do not want to install these features, deselect the **Include**

management tools option and select **Add Features**. Once optional features selection is complete, select **Next** to proceed in the wizard.

> **Note**
> The **Enhanced Storage** feature is a required feature for enabling BitLocker. This feature enables support for Encrypted Hard Drives on capable systems. For more information on Encrypted Hard Drive support, please see Encrypted Hard Drive.

8. Select **Install** on the **Confirmation** pane of the **Add Roles and Features Wizard** to begin BitLocker feature installation. The BitLocker feature requires a restart to complete. Selecting the **Restart the destination server automatically if required** option in the **Confirmation** pane will force a restart of the computer after installation is complete.
9. If the **Restart the destination server automatically if required** check box is not selected, the **Results pane** of the **Add Roles and Features Wizard** will display the success or failure of the BitLocker feature installation. If required, a notification of additional action necessary to complete the feature installation, such as the restart of the computer, will be displayed in the results text.

To install BitLocker using Windows PowerShell

Windows PowerShell offers administrators another option for BitLocker feature installation. Windows PowerShell installs features using the `servermanager` or `dism` module; however, the `servermanager` and `dism` modules do not always share feature name parity. Because of this, it is advisable to confirm the feature or role name prior to installation.

> **Note**
> The BitLocker feature requires the server to be restarted after installation to fully install.

Using the servermanager module to install BitLocker

The `servermanager` Windows PowerShell module can use either the `Install-WindowsFeature` or `Add-WindowsFeature` to install the BitLocker feature. The `Add-WindowsFeature` cmdlet is merely a stub to the `Install-WindowsFeature`. This example uses the `Install-WindowsFeature` cmdlet. The feature name for BitLocker in the `servermanager` module is `BitLocker`. This can be determined using the `Get-WindowsFeature` cmdlet with a query such as:

```
Get-WindowsFeature Bit
```

The results of this command displays a table of all of the feature names beginning with “Bit” as their prefix. This allows you to confirm that the feature name is `BitLocker` for the BitLocker feature.

By default, installation of features in Windows PowerShell does not include optional sub-features or management tools as part of the install process. This can be seen using the `-WhatIf` option in Windows PowerShell.

```
Install-WindowsFeature BitLocker -WhatIf
```

The results of this command show that only the BitLocker Drive Encryption feature installs using this command.

To see what would be installed with the BitLocker feature including all available management tools and sub-features, use the following command:

```
Install-WindowsFeature BitLocker -IncludeAllSubFeature -IncludeManagementTools -WhatIf | fl
```

The result of this command displays the following list of all the administration tools for BitLocker that would be installed along with the feature, including tools for use with Active Directory Domain Services (AD DS) and Active Directory Lightweight Directory Services (AD LDS).

- BitLocker Drive Encryption
- BitLocker Drive Encryption Tools
- BitLocker Drive Encryption Administration Utilities
- BitLocker Recovery Password Viewer
- AD DS Snap-Ins and Command-Line Tools
- AD DS Tools
- AD DS and AD LDS Tools

The command to complete a full installation of the BitLocker feature with all available features and then rebooting the server at completion is:

```
Install-WindowsFeature BitLocker -IncludeAllSubFeature -IncludeManagementTools -Restart
```

Important
Installing the BitLocker feature using Windows PowerShell does not install the Enhanced Storage feature. Administrators wishing to support Encrypted Hard Drives in their environment will need to install the Enhanced Storage feature separately.

Using the dism module to install BitLocker

The `dism` Windows PowerShell module uses the `Enable-WindowsOptionalFeature` cmdlet to install features. The BitLocker feature name for BitLocker is `BitLocker`. The `dism` module does not support wildcards when searching for feature names. To list feature names for the `dism` module, use the `Get-WindowsOptionalFeatures` cmdlet. The following command will list all of the optional features in an online (running) operating system.

```
Get-WindowsOptionalFeature -Online | ft
```

From this output, we can see that there are three BitLocker related optional feature names: BitLocker, BitLocker-Utilities and BitLocker-NetworkUnlock. To install the BitLocker feature, the BitLocker and BitLocker-Utilities features are the only required items.

To install BitLocker using the `dism` module, use the following command:

```
Enable-WindowsOptionalFeature -Online -FeatureName BitLocker -All
```

This command will prompt the user for a reboot. The Enable-WindowsOptionalFeature cmdlet does not offer support for forcing a reboot of the computer. This command does not include installation of the management tools for BitLocker. For a complete installation of BitLocker and all available management tools, use the following command:

```
Enable-WindowsOptionalFeature -Online -FeatureName BitLocker, BitLocker-Utilities -All
```

More information

BitLocker Overview

Encrypted Hard Drive

BitLocker Frequently Asked Questions (FAQ)

What's New in BitLocker for Windows 8 and Windows Server 2012

What's New in BitLocker for Windows 8.1 and Windows Server 2012 R2

Prepare your organization for BitLocker: Planning and Policies

BitLocker: How to enable Network Unlock

Protecting cluster shared volumes and storage area networks with BitLocker

Windows Server 2012 introduces the ability to use BitLocker encrypted volumes on a failover cluster. BitLocker will have the ability to protect both physical disk resources and cluster shared volumes version 2.0 (CSV2.0). BitLocker on clustered volumes allows for an additional layer of protection for administrators wishing to protect sensitive, highly available data. By adding additional protectors to the clustered volume, administrators can also add an additional barrier of security to resources within an organization by allowing only certain user accounts access to unlock the BitLocker volume.

Configuring BitLocker on Cluster Shared Volumes

Using BitLocker with Clustered Volumes

BitLocker on volumes within a cluster are managed based on how the cluster service "views" the volume to be protected. The volume can be a physical disk resource such as a logical unit number (LUN) on a storage area network (SAN) or network attached storage (NAS).

Important
SANs used with BitLocker must have obtained Windows Hardware Certification. For more information see the Windows Hardware Certification Program FAQ.

Alternatively, the volume can be a cluster-shared volume, a shared namespace, within the cluster. Windows Server 2012 has expanded the CSV architecture, now known as CSV2.0, to enable support for BitLocker. When using BitLocker with volumes designated for a cluster, the volume will need to turn on BitLocker before its addition to the storage pool within cluster or put the resource into maintenance mode before BitLocker operations will complete.

Windows PowerShell or the manage-bde command line interface is the preferred method to manage BitLocker on CSV2.0 volumes. This is recommended over the BitLocker Control Panel item because CSV2.0 volumes are mount points. Mount points are an NTFS object that is used to

provide an entry point to other volumes. Mount points do not require the use of a drive letter. Volumes that lack drive letters do not appear in the BitLocker Control Panel item. Additionally, the new Active Directory-based protector option required for cluster disk resource or CSV2.0 resources is not available in the Control Panel item.

Note

Mount points can be used to support remote mount points on SMB based network shares. This type of share is not supported for BitLocker encryption.

For thinly provisioned storage, such as a Dynamic Virtual Hard Disk (VHD), BitLocker runs in Used Disk Space Only encryption mode. You cannot use the **manage-bde –WipeFreeSpace** command to transition the volume to full-volume encryption on these types of volumes. This occurs because Full Encryption requires an end marker for the volume and dynamically expanding VHDs do not have a static end of volume marker.

Active Directory based Protector

A new BitLocker protector is available for protecting clustered volumes held within the Active Directory Domain Services (AD DS) infrastructure. The **ADAccountOrGroup** protector is a domain security identifier (SID) based protector that can be bound to a user account, machine account or group. When an unlock request is made for a protected volume, the BitLocker service interrupts the request and uses the BitLocker protect/unprotect APIs to unlock or deny the request. BitLocker will unlock protected volumes without user intervention by attempting protectors in the following order:

1. Clear key
2. Driver-based auto-unlock key
3. ADAccountOrGroup protector
    a. Service context protector
    b. User protector
4. Registry-based auto-unlock key

Note

A Windows Server 2012 domain controller is required for this feature to work properly.

Turning on BitLocker before adding disks to a cluster using Windows PowerShell

BitLocker encryption is available for disks before or after addition to a cluster storage pool. The advantage of encrypting volumes prior to adding them to a cluster is that the disk resource does not require suspending the resource to complete the operation. To turn on BitLocker for a disk before adding it to a cluster, do the following:

1. Install the BitLocker Drive Encryption feature if it is not already installed
2. Ensure the disk is formatted NTFS and has a drive letter assigned to it

3. Enable BitLocker on the volume using your choice of protector. A password protector is used in the Windows PowerShell script example below.

```
Enable-BitLocker E: -PasswordProtector -Password $pw
```

4. Identify the name of the cluster with Windows PowerShell

```
Get-Cluster
```

5. Add an **ADAccountOrGroup**protector to the volume using the cluster name using a command such as:

```
Add-BitLockerProtector E: -ADAccountOrGroupProtector -
ADAccountOrGroup CLUSTER$
```

Warning
You must add an **ADAccountOrGroup** protector using the cluster CNO for a BitLocker enabled volume to either be shared in a Cluster Shared Volume or to failover properly in a traditional failover cluster.

6. Repeat steps 1-6 for each disk in the cluster
7. Add the volume(s) to the cluster

Turning on BitLocker for a clustered disk using Windows PowerShell

When the cluster service owns a disk resource already, it needs to be set into maintenance mode before BitLocker can be enabled. Use the following steps for turning BitLocker on for a clustered disk:

1. Install the BitLocker Drive Encryption feature if it is not already installed
2. Check the status of the cluster disk using Windows PowerShell

```
Get-ClusterResource "Cluster Disk 1"
```

3. Put the physical disk resource into maintenance mode using Windows PowerShell

```
Get-ClusterResource "Cluster Disk 1" | Suspend-ClusterResource
```

4. Enable BitLocker on the volume using your choice of protector. A password protector is used in the example below.

```
Enable-BitLocker E: -PasswordProtector -Password $pw
```

5. Identify the name of the cluster with Windows PowerShell

```
Get-Cluster
```

6. Add an **ADAccountOrGroup** protector with the Cluster Name Object (CNO) to the volume using a command such as:

```
Add-BitLockerProtector E: -ADAccountOrGroupProtector -
ADAccountOrGroup CLUSTER$
```

Warning

You must add an **ADAccountOrGroup** protector using the cluster CNO for a BitLocker enabled volume to either be shared in a Cluster Shared Volume or to failover properly in a traditional failover cluster.

7. Repeat steps 1-6 for each disk in the cluster
8. Add the volume(s) to the cluster

Adding BitLocker encrypted volumes to a cluster using manage-bde

You can also use manage-bde to enable BitLocker on clustered volumes. The steps needed to add a physical disk resource or CSV2.0 volume to an existing cluster includes the following:

1. Verify the BitLocker Drive Encryption feature is installed on the computer
2. Ensure new storage is formatted as NTFS
3. Encrypt the volume, add a recovery key and add the cluster administrator as a protector key using the manage-bde command line interface (see example):
    a. `Manage-bde -on -used <drive letter> -RP -sid domain\CNO$ -sync`
        i. BitLocker will check to see if the disk is already part of a cluster. If it is, administrators will encounter a hard block. Otherwise, the encryption will continue.
        ii. Using the -sync parameter is optional. Using it ensures the command waits until the encryption for the volume is completed before releasing the volume for use in the cluster storage pool.
4. Open the Failover Cluster Manager snap-in or cluster PowerShell cmdlets to enable the disk to be clustered
    a. Once the disk is clustered it can also be enabled for CSV.
5. During the resource online operation, cluster will check to see if the disk is BitLocker encrypted.
    a. If the volume is not BitLocker enabled, traditional cluster online operations occur.
    b. If the volume is BitLocker enabled, the following check occurs:
        i. If volume is **locked**, BitLocker will impersonate the CNO and unlock the volume using the CNO protector. If this operation fails an event will be logged that the volume could not be unlocked and the online operation will fail.
6. Once the disk is online in the storage pool, it can be added to a CSV by right clicking on the disk resource and choosing "**Add to cluster shared volumes**".

CSVs can include both encrypted and unencrypted volumes. To check the status of a particular volume for BitLocker encryption, administrators can utilize the manage-bde -status command with a path to the volume inside the CSV namespace as seen in the example command line below.

```
manage-bde -status "C:\ClusterStorage\volume1"
```

Physical Disk Resources

Unlike CSV2.0 volumes, physical disk resources can only be accessed by one cluster node at a time. This means that operations such as encrypting, decrypting, locking or unlocking volumes

require context to perform. For example, you cannot unlock or decrypt a physical disk resource if you are not administering the cluster node that owns the disk resource because the disk resource is not available.

Restrictions on BitLocker actions with cluster volumes

The following table contains information about both Physical Disk Resources (i.e. traditional failover cluster volumes) and Cluster Shared Volumes (CSV) and the actions that are allowed by BitLocker in each situation.

Action	On owner node of failover volume	On Metadata Server (MDS) of CSV	On (Data Server) DS of CSV	Maintenance Mode
Manage-bde –on	Blocked	Blocked	Blocked	Allowed
Manage-bde –off	Blocked	Blocked	Blocked	Allowed
Manage-bde Pause/Resume	Blocked	Blocked**	Blocked	Allowed
Manage-bde –lock	Blocked	Blocked	Blocked	Allowed
manage-bde –wipe	Blocked	Blocked	Blocked	Allowed
Unlock	Automatic via cluster service	Automatic via cluster service	Automatic via cluster service	Allowed
manage-bde –protector –add	Allowed	Allowed	Blocked	Allowed
manage-bde -protector -delete	Allowed	Allowed	Blocked	Allowed
manage-bde –autounlock	Allowed (not recommended)	Allowed (not recommended)	Blocked	Allowed (not recommended)
Manage-bde -upgrade	Allowed	Allowed	Blocked	Allowed
Shrink	Allowed	Allowed	Blocked	Allowed
Extend	Allowed	Allowed	Blocked	Allowed

Note

Although the manage-bde -pause command is Blocked in clusters, the cluster service will automatically resume a paused encryption or decryption from the MDS node

In the case where a physical disk resource experiences a failover event during conversion, the new owning node will detect the conversion is not complete and will complete the conversion process.

Other considerations when using BitLocker on CSV 2.0

Some other considerations to take into account for BitLocker on clustered storage include the following:

- BitLocker volumes have to be initialized and beginning encryption before they are available to add to a CSV2.0 volume.
- If an administrator needs to decrypt a CSV volume, remove the volume from the cluster or put into disk maintenance mode. You can add the CSV back to the cluster while waiting for decryption to complete.
- If an administrator needs to start encrypting a CSV volume, remove the volume from the cluster or put it in maintenance mode.
- If conversion is paused with encryption in progress and the CSV volume is offline from the cluster, the cluster thread (health check) will automatically resume conversion when the volume is online to the cluster.
- If conversion is paused with encryption in progress and a physical disk resource volume is offline from the cluster, the BitLocker driver will automatically resume conversion when the volume is online to the cluster.
- If conversion is paused with encryption in progress, while the CSV volume is in maintenance mode, the cluster thread (health check) will automatically resume conversion when moving the volume back from maintenance.
- If conversion is paused with encryption in progress, while the disk resource volume is in maintenance mode, the BitLocker driver will automatically resume conversion when the volume is moved back from maintenance mode.

BitLocker: How to enable Network Unlock

Windows 8 and Windows Server 2012 has added a new BitLocker protector option for operating system volumes called Network Unlock. Network Unlock enables easier management for BitLocker enabled desktops and servers in a domain environment by providing automatic unlock of operating system volumes at system reboot when connected to a wired corporate network. This feature requires the client hardware to have a DHCP driver implemented in its UEFI firmware.

Without Network Unlock, operating system volumes protected by TPM+PIN protectors require a PIN to be entered when a machine reboots or resumes from hibernation (for example, by Wake on LAN). This can make it difficult to enterprises to roll out software patches to unattended desktops and remotely administered servers.

Network Unlock allows BitLocker-enabled systems with TPM+PIN and that meet the hardware requirements to boot into Windows without user intervention. Network Unlock works in a similar fashion to the TPM+StartupKey at boot. Rather than needing to read the StartupKey from USB media, however, the key for Network Unlock is composed from a key stored in the TPM and an encrypted network key that is sent to the server, decrypted and returned to the client in a secure session.

Network Unlock Core requirements

Network Unlock must meet mandatory hardware and software requirements before the feature can automatically unlock domain joined systems. These requirements include:

- Computers running Windows 8 or Windows Server 2012 with UEFI DHCP drivers can be Network Unlock clients.
- BitLocker Network Unlock optional feature installed on Windows Server 2012.
- A separate Windows Server 2012 server running the Windows Deployment Services (WDS) role.
- A DHCP server, separate from the WDS server.
- Properly configured public/private key pairing.
- Network Unlock Group Policy settings configured.

Notes

To properly support DHCP within UEFI, the UEFI-based system should be in native mode without a compatibility support module (CSM) enabled.

For Network Unlock to work reliably, the first network interface card (NIC) on the computer, usually the onboard NIC, must be configured to support DHCP and used for Network Unlock. This is especially worth noting when you have multiple NICs, and you wish to configure one without DHCP, such as for a lights-out management protocol. This configuration is necessary because Network Unlock will stop enumerating NICs when it reaches one with a DHCP port failure for any reason. Thus, if the first enumerated NIC does not support DHCP, is not plugged into the network, or fails to report availability of the DHCP port for any reason, then Network Unlock will fail.

The Network Unlock server component installs on Windows Server 2012 as a Windows feature using Server Manager or Windows PowerShell cmdlets. The feature name is BitLocker Network Unlock in Server Manager and BitLocker-NetworkUnlock in Windows PowerShell. This feature is a core requirement.

Network Unlock requires a Windows Server 2012 running Windows Deployment Services (WDS) in the environment where the feature will be utilized. Configuration of the WDS installation is not required; however, the WDS service needs to be running on the server.

The network key is stored on the system drive along with an AES 256 session key, and encrypted with the 2048-bit RSA public key of the unlock server's certificate. The network key is decrypted with the help of a provider on a Windows Server 2012 WDS server and returned encrypted with its corresponding session key.

Network Unlock sequence

The unlock sequence starts on the client side, when the Windows boot manager detects the existence of Network Unlock protector. It leverages the DHCP driver in UEFI to obtain an IP address for IPv4 and then broadcasts a vendor-specific DHCP request that contains the network key and a session key for the reply, all encrypted by the server's Network Unlock certificate, as described above. The Network Unlock provider on the Windows Server 2012 WDS server recognizes the vendor-specific request, decrypts it with the RSA private key, and returns the network key encrypted with the session key via its own vendor-specific DHCP reply.

On the server side, the WDS server role has an optional plugin component, like a PXE provider, which is what handles the incoming Network Unlock requests. The provider can also be configured with subnet restrictions, which would require that the IP address provided by the client in the Network Unlock request belong to a permitted subnet in order to release the network key to the client. In instances where the Network Unlock provider is unavailable, BitLocker fails over to the next available protector to unlock the drive. In a typical configuration, this means the standard TPM+PIN unlock screen is presented to unlock the drive.

The server side configuration to enable Network Unlock also requires provisioning a 2048-bit RSA public/private key pair in the form of an X.509 certificate, and for the public key certificate to be distributed to the clients. This certificate must be managed and deployed through the Group Policy editor directly on Windows Server 2012 domain controller. This certificate is the public key that encrypts the intermediate network key (which is one of the two secrets required to unlock the drive; the other secret is stored in the TPM).

Configuring Network Unlock

The following steps allow an administrator to configure Network Unlock in a domain.

Step One: Install the WDS Server role

The BitLocker Network Unlock feature will install the WDS role if it is not already installed. If you want to install it separately before you install BitLocker Network Unlock you can use Server Manager or Windows PowerShell. To install the role using Server Manager, select the **Windows Deployment Services** role in Server Manager.

To install the role using Windows PowerShell, use the following command:

```
Install-WindowsFeature WDS-Deployment
```

Step Two: Confirm the WDS Service is running

To confirm the WDS service is running, use the Services Management Console or Windows PowerShell. To confirm the service is running in Services Management Console, open the console using **services.msc** and check the status of the Windows Deployment Services service.

To confirm the service is running using Windows PowerShell, use the following command:

```
Get-Service WDSServer
```

Step Three: Install the Network Unlock feature

To install the Network Unlock feature, use Server Manager or Windows PowerShell. To install the feature using Server Manager, select the **BitLocker Network Unlock** feature in the Server Manager console.

To install the feature using Windows PowerShell, use the following command:

```
Install-WindowsFeature BitLocker-NetworkUnlock
```

Step Four: Create the Network Unlock certificate

Network Unlock can use imported certificates from an existing PKI infrastructure, or you can use a self-signed certificate.

To enroll a certificate from an existing certification authority (CA), do the following:

1. Open Certificate Manager on the WDS server using **certmgr.msc**
2. Under the Certificates - Current User item, right-click Personal
3. Select All Tasks, then **Request New Certificate**
4. Select **Next** when the Certificate Enrollment wizard opens
5. Select Active Directory Enrollment Policy
6. Choose the certificate template created for Network Unlock on the Domain controller and select **Enroll**. When prompted for more information, add the following attribute to the certificate:
    - Select the **Subject Name** pane and provide a friendly name value. It is suggested that this friendly name include information for the domain or organizational unit for the certificate. For example "BitLocker Network Unlock Certificate for Contoso domain"
7. Create the certificate. Ensure the certificate appears in the Personal folder.
8. Export the public key certificate for Network Unlock
    a. Create a .cer file by right-clicking the previously created certificate, choosing **All Tasks**, then **Export**.
    b. Select **No, do not export the private key**.
    c. Select **DER encoded binary X.509** and complete exporting the certificate to a file.
    d. Give the file a name such as BitLocker-NetworkUnlock.cer.
9. Export the public key with a private key for Network Unlock
    a. Create a .pfx file by right-clicking the previously created certificate, choosing **All Tasks**, then **Export**.
    b. Select **Yes, export the private key**.
    c. Complete the wizard to create the .pfx file.

To create a self-signed certificate, do the following:

1. Create a text file with an .inf extension. For example, notepad.exe BitLocker-NetworkUnlock.inf
2. Add the following contents to the previously created file:

```
[NewRequest]
```

```
Subject="CN=BitLocker Network Unlock certificate"
Exportable=true
RequestType=Cert
KeyLength=2048

[Extensions]
1.3.6.1.4.1.311.21.10 = "{text}"
_continue_ = "OID=1.3.6.1.4.1.311.67.1.1"

2.5.29.37 = "{text}"
_continue_= "1.3.6.1.4.1.311.67.1.1"
```

3. Open an elevated command prompt and use the certreq tool to create a new certificate using the following command:

```
certreq -new BitLocker-NetworkUnlock.inf BitLocker-NetworkUnlock.cer
```

4. Verify the previous command properly created the certificate by confirming the .cer file exists
5. Launch the Certificate Manager by running **certmgr.msc**
6. Create a .pfx file by navigating to the **Certificates – Current User** pane, right-clicking the previously imported certificate, selecting **All Tasks**, then **Export**. Follow through the wizard to create the .pfx file.

Step Five: Deploy the private key and certificate to the WDS server

With the certificate and key created, deploy them to the infrastructure to properly unlock systems. To deploy the certificates, do the following:

1. On the WDS server, open a new MMC and add the certificates snap-in. Select the computer account and local computer when given the options.
2. Right-click the Certificates (Local Computer) - BitLocker Drive Encryption Network Unlock item, choose All Tasks, then **Import**
3. In the **File to Import** dialog, choose the .pfx file created previously.
4. Enter the password used to create the .pfx and complete the wizard.

Step Six: Configure Group Policy settings for Network Unlock

With certificate and key deployed to the WDS server for Network Unlock, the final step is to use Group Policy settings to deploy the public key certificate to computers that you want to be able to unlock using the Network Unlock key. The following steps describe how to deploy the required Group Policy setting:

1. Copy the .cer file created for Network Unlock to the domain controller

2. On the domain controller, launch Group Policy Management Console (gpmc.msc)
3. Create a new Group Policy Object or modify an existing object to enable the **Allow network unlock at startup** setting.
4. Deploy the public certificate to clients
    a. Within Group Policy Management Console, navigate to the following location: **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies\BitLocker Drive Encryption Network Unlock Certificate**
    b. Right-click the folder and choose **Add Network Unlock Certificate**
    c. Follow the wizard steps and import the .cer file that was copied earlier.

Note

Only one network unlock certificate can be available at a time. If a new certificate is required, delete the current certificate before deploying a new one. The Network Unlock certificate is located in the **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\SystemCertificates\FVE_NKP** key on the client computer.

Step Seven: Require TPM+PIN protectors at startup (recommended)

An additional recommendation is for enterprises to use TPM+PIN protectors for an extra level of security. To require TPM+PIN protectors in an environment, do the following:

1. Open Group Policy Management Console (gpmc.msc)
2. Enable the policy **Require additional authentication at startup** and select the **Require startup PIN with TPM** option
3. Turn on BitLocker with TPM+PIN protectors on all domain-joined computers

Creating the certificate template for Network Unlock

The following steps detail how to create a certificate template for use with BitLocker Network Unlock. A properly configured Active Directory Services Certification Authority can use this certificate to create and issue Network Unlock certificates.

1. Open the Certificates Template snap-in (certtmpl.msc).
2. Locate the User template. Right-click the template name and select **Duplicate Template**
3. On the **Compatibility** tab, change the **Certification Authority** and **Certificate recipient** fields to Windows Server 2012 and Windows 8respectively. Ensure the **Show resulting changes** dialog box is selected.
4. Select the **General** tab of the template. The **Template display name** and **Template name** should clearly identify that the template will be used for Network Unlock. Clear the checkbox for the **Publish certificate in Active Directory** option.
5. Select the **Request Handling** tab. Select **Encryption** from the **Purpose** drop down menu. Ensure the **Allow private key to be exported** option is selected.

6. Select the **Cryptography** tab. Set the **Minimum key size** to 2048. (Any Microsoft cryptographic provider that supports RSA can be used for this template, but for simplicity and forward compatibility we recommend using the **Microsoft Software Key Storage Provider**.)
7. Select the **Requests must use one of the following providers** option and clear all options except for the cryptography provider you selected, such as the **Microsoft Software Key Storage Provider**.
8. Select the **Subject Name** tab. Select **Supply in the request**. Select **OK** if the certificate templates pop-up dialog appears.
9. Select the **Issuance Requirements** tab. Select both **CA certificate manager approval** and **Valid existing certificate** options.
10. Select the **Extensions** tab. Select **Application Policies** and choose **Edit…**.
11. In the **Edit Application Policies Extension** options dialog box, select **Client Authentication**, **Encrypting File System**, **and Secure Email** and choose **Remove**.
12. On the **Edit Application Policies Extension** dialog box, select **Add**.
13. On the **Add Application Policy** dialog box, select **New**. In the **New Application Policy** dialog box enter the following information in the space provided and then click **OK** to create the BitLocker Network Unlock application policy:
    - **Name: BitLocker Network Unlock**
    - **Object Identifier: 1.3.6.1.4.1.311.67.1.1**
14. Select the newly created **BitLocker Network Unlock** application policy and select **OK**
15. With the **Extensions** tab still open, select the **Edit Key Usage Extension** dialog, select the **Allow key exchange only with key encryption (key encipherment)** option. Select the **Make this extension critical** option.
16. Select the **Security** tab. Confirm that the **Domain Admins** group has been granted **Enroll** permission
17. Select **OK** to complete configuration of the template.

To add the Network Unlock template to the Certification Authority, open the Certification Authority snap-in (certsrv.msc). Right-click the **Certificate Templates** item and choose **New, Certificate Template to issue**. Select the previously created BitLocker Network Unlock certificate.

After adding the Network Unlock template to the Certification Authority, this certificate can be used to configure BitLocker Network Unlock.

Subnet policy configuration files on WDS Server (Optional)

By default, all clients with the correct Network Unlock Certificate and valid Network Unlock protectors that have wired access to a Network Unlock-enabled WDS server via DHCP are unlocked by the server. A subnet policy configuration file on the WDS server can be created to limit which subnet(s) Network Unlock clients can use to unlock.

The configuration file, called bde-network-unlock.ini, must be located in the same directory as the Network Unlock provider DLL and it applies to both IPv6 and IPv4 DHCP implementations. If the subnet configuration policy becomes corrupted, the provider will fail and stop responding to requests.

The subnet policy configuration file must use a “[SUBNETS]” section to identify the specific subnets. The named subnets may then be used to specify restrictions in certificate subsections. Subnets are defined as simple name-value pairs, in the common INI format, where each subnet has its own line, with the name on the left of the equals sign, and the subnet identified on the right of the equal sign as a Classless Inter-Domain Routing (CIDR) address or range. The key word “ENABLED” is disallowed for subnet names.

```
 [SUBNETS]

SUBNET1=10.185.250.0/24 ; comment about this subrange could be here, after the semi-colon

SUBNET2=10.185.252.200/28

SUBNET3= 2001:4898:a:2::/64 ; an IPv6 subnet

SUBNET4=2001:4898:a:3::/64; in production, the admin would likely give more useful names,
like BUILDING9-EXCEPT-RECEP.
```

Following the [SUBNETS] section, there can be sections for each Network Unlock certificate, identified by the certificate thumbprint formatted without any spaces, which define subnets clients can be unlocked from with that certificate.

Note

When specifying the certificate thumbprint, do not include any spaces. If spaces are included in the thumbprint the subnet configuration will fail because the thumbprint will not be recognized as valid.

Subnet restrictions are defined within each certificate section by denoting the allowed list of permitted subnets. If any subnet is listed in a certificate section, then only those subnets listed are permitted for that certificate. If no subnet is listed in a certificate section, then all subnets are permitted for that certificate. If a certificate does not have a section in the subnet policy configuration file, then no subnet restrictions are applied for unlocking with that certificate. This means for restrictions to apply to every certificate, there must be a certificate section for every Network Unlock certificate on the server, and an explicit allowed list set for each certificate section.

Subnet lists are created by putting the name of a subnet from the [SUBNETS] section on its own line below the certificate section header. Then, the server will only unlock clients with this certificate on the subnet(s) specified as in the list. For troubleshooting, a subnet can be quickly excluded without deleting it from the section by simply commenting it out with a prepended semi-colon.

```
 [2158a767e1c14e88e27a4c0aee111d2de2eafe60]

;Comments could be added here to indicate when the cert was issued, which Group Policy
should get it, and so on.

;This list shows this cert is only allowed to unlock clients on SUBNET1 and SUBNET3
subnets. In this example, SUBNET2 is commented out.

SUBNET1

;SUBNET2
```

```
SUBNET3
```

To disallow the use of a certificate altogether, its subnet list may contain the line “DISABLED".

Turning off Network Unlock

To turn off the unlock server, the PXE provider can be unregistered from the WDS server or uninstalled altogether. However, to stop clients from creating Network Unlock protectors the **Allow Network Unlock at startup** Group Policy setting should be disabled. When this policy setting is updated to disabled on client computers any Network Unlock key protectors on the computer will be deleted. Alternatively, the BitLocker Network Unlock certificate policy can be deleted on the domain controller to accomplish the same task for an entire domain.

Note

Removing the FVENKP certificate store that contains the Network Unlock certificate and key on the WDS server will also effectively disable the server’s ability to respond to unlock requests for that certificate. However, this is seen as an error condition and is not a supported or recommended method for turning off the Network Unlock server.

Updating Network Unlock certificates

To update the certificates used by Network Unlock, administrators need to import or generate the new certificate for the server and then update the Network Unlock certificate Group Policy setting on the domain controller.

Troubleshooting Network Unlock

Troubleshooting Network Unlock issues begins by verifying the environment. Many times, a small configuration issue will be the root cause of the failure. Items to verify include:

- Verify client hardware is UEFI-based and is on firmware version is 2.3.1 and that the UEFI firmware is in native mode without a Compatibility Support Module (CSM) for BIOS mode enabled. Do this by checking that the firmware does not have an option enabled such as "Legacy mode" or "Compatibility mode" or that the firmware does not appear to be in a BIOS-like mode.
- All required roles and services are installed and started
- Public and private certificates have been published and are in the proper certificate containers. The presence of the Network Unlock certificate can be verified in the Microsoft Management Console (MMC.exe) on the WDS server with the certificate snap-ins for the local computer enabled. The client certificate can be verified by checking the registry key **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\SystemCertificates\FVE_NKP** on the client computer.
- Group policy for Network Unlock is enabled and linked to the appropriate domains
- Verify group policy is reaching the clients properly. This can be done using the GPRESULT.exe or RSOP.msc utilities.

- Verify the **Network (Certificate Based)** protector is listed on the client. This can be done using either manage-bde or Windows PowerShell cmdlets. For example the following command will list the key protectors currently configured on the C: drive of the lcoal computer:

  ```
  Manage-bde –protectors –get C:
  ```

Note

Use the output of manage-bde along with the WDS debug log to determine if the proper certificate thumbprint is being used for Network Unlock

Files to gather when troubleshooting BitLocker Network Unlock include:

1. The Windows event logs. Specifically the BitLocker event logs and the Microsoft-Windows-Deployment-Services-Diagnostics-Debug log

   Debug logging is turned off by default for the WDS server role, so you will need to enable it first. You can use either of the following two methods to turn on WDS debug logging.

   a. Start an elevated command prompt and run the following command:

      ```
      wevtutil sl Microsoft-Windows-Deployment-Services-
      Diagnostics/Debug /e:true
      ```

   b. Open Event Viewer on the WDS server.

      In the left pane, click **Applications and Services Logs**, click **Microsoft**, click **Windows**, click **Deployment-Services-Diagnostics**, and then click **Debug**.

      In the right pane, click **Enable Log**.

2. The DHCP subnet configuration file (if one exists).
3. The output of the BitLocker status on the volume, this can be gathered into a text file using **manage-bde -status** or **Get-BitLockerVolume** in Windows PowerShell
4. Network Monitor capture on the server hosting the WDS role, filtered by client IP address

More information

BitLocker Overview

Encrypted Hard Drive

BitLocker Frequently Asked Questions (FAQ)

What's New in BitLocker for Windows 8 and Windows Server 2012

Prepare your organization for BitLocker: Planning and Policies

BitLocker Recovery Guide

Organizations can use BitLocker recovery information saved in Active Directory Domain Services (AD DS) to access BitLocker-protected data. Creating a recovery model for BitLocker while you are planning your BitLocker deployment is recommended.

This article assumes that you understand how to set up AD DS to back up BitLocker recovery information automatically, and what types of recovery information are saved to AD DS.

This article does not detail how to configure AD DS to store the BitLocker recovery information.

This article contains the following topics:

- What is BitLocker recovery?
- Testing recovery
- Planning your recovery process
- Using additional recovery information
- Resetting recovery passwords
- Retrieving the BitLocker key package

What is BitLocker recovery?

BitLocker recovery is the process by which you can restore access to a BitLocker-protected drive in the event that you cannot unlock the drive normally. In a recovery scenario you have the following options to restore access to the drive:

- The user can supply the recovery password. If your organization allows users to print or store recovery passwords, the user can type in the 48-digit recovery password that they printed or stored on a USB drive or with your Microsoft Account online. (Saving a recovery password with your Microsoft Account online is only allowed when BitLocker is used on a PC that is not a member of a domain).
- A data recovery agent can use their credentials to unlock the drive. If the drive is an operating system drive, the drive must be mounted as a data drive on another computer for the data recovery agent to unlock it.
- A domain administrator can obtain the recovery password from AD DS and use it to unlock the drive. Storing recovery passwords in AD DS is recommended to provide a way for IT professionals to be able to obtain recovery passwords for drives in their organization if needed. This method requires that you have enabled this recovery method in the BitLocker Group Policy setting **Choose how BitLocker-protected operating system drives can be recovered** located at **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives** in the Local Group Policy Editor. For more information, see BitLocker Group Policy Settings.

What causes BitLocker recovery?

The following list provides examples of specific events that will cause BitLocker to enter recovery mode when attempting to start the operating system drive:

- Device Lockout feature in Windows 8, Windows RT, Windows 8.1 andWindows RT 8.1. On PCs that use either BitLocker or Device Encryption when an attack is detected the device will immediately reboot and enter into BitLocker recovery mode. To take advantage of this functionality Administrators can set the **Interactive logon: Machine account lockout threshold** Group Policy setting located in **\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options** in the Local Group Policy Editor, or use the **MaxFailedPasswordAttempts** policy of Exchange ActiveSync (also configurable through Windows Intune), to limit the number of failed password attempts before the device goes into Device Lockout.
- Changing the boot order to boot another drive in advance of the hard drive.
- Having the CD or DVD drive before the hard drive in the BIOS boot order and then inserting or removing a CD or DVD.
- Failing to boot from a network drive before booting from the hard drive.
- Docking or undocking a portable computer. In some instances (depending on the computer manufacturer and the BIOS), the docking condition of the portable computer is part of the system measurement and must be consistent to validate the system status and unlock BitLocker. This means that if a portable computer is connected to its docking station when BitLocker is turned on, then it might also need to be connected to the docking station when it is unlocked. Conversely, if a portable computer is not connected to its docking station when BitLocker is turned on, then it might need to be disconnected from the docking station when it is unlocked.
- Changes to the NTFS partition table on the disk including creating, deleting, or resizing a primary partition.
- Entering the personal identification number (PIN) incorrectly too many times so that the anti-hammering logic of the TPM is activated. Anti-hammering logic is software or hardware methods that increase the difficulty and cost of a brute force attack on a PIN by not accepting PIN entries until after a certain amount of time has passed.
- Turning off the support for reading the USB device in the pre-boot environment from the BIOS or UEFI firmware if you are using USB-based keys instead of a TPM.
- Turning off, disabling, deactivating, or clearing the TPM.
- Upgrading critical early startup components, such as a BIOS or UEFI firmware upgrade, causing the related boot measurements to change.
- Forgetting the PIN when PIN authentication has been enabled.
- Updating option ROM firmware.
- Upgrading TPM firmware.
- Adding or removing hardware; for example, inserting a new card in the computer, including some PCMIA wireless cards.
- Removing, inserting, or completely depleting the charge on a smart battery on a portable computer.
- Changes to the master boot record on the disk.
- Changes to the boot manager on the disk.
- Hiding the TPM from the operating system. Some BIOS or UEFI settings can be used to prevent the enumeration of the TPM to the operating system. When implemented, this option

can make the TPM hidden from the operating system. When the TPM is hidden, BIOS and UEFI secure startup are disabled, and the TPM does not respond to commands from any software.

- Using a different keyboard that does not correctly enter the PIN or whose keyboard map does not match the keyboard map assumed by the pre-boot environment. This can prevent the entry of enhanced PINs.
- Modifying the Platform Configuration Registers (PCRs) used by the TPM validation profile. For example, including **PCR[1]** would result in BitLocker measuring most changes to BIOS settings, causing BitLocker to enter recovery mode even when non-boot critical BIOS settings change.

  Note
  Some computers have BIOS settings that skip measurements to certain PCRs, such as **PCR[2]**. Changing this setting in the BIOS would cause BitLocker to enter recovery mode because the PCR measurement will be different.
- Moving the BitLocker-protected drive into a new computer.
- Upgrading the motherboard to a new one with a new TPM.
- Losing the USB flash drive containing the startup key when startup key authentication has been enabled.
- Failing the TPM self-test.
- Having a BIOS, UEFI firmware, or an option ROM component that is not compliant with the relevant Trusted Computing Group standards for a client computer. For example, a non-compliant implementation may record volatile data (such as time) in the TPM measurements, causing different measurements on each startup and causing BitLocker to start in recovery mode.
- Changing the usage authorization for the storage root key of the TPM to a non-zero value.

  Note
  The BitLocker TPM initialization process sets the usage authorization value to zero, so another user or process must explicitly have changed this value.
- Disabling the code integrity check or enabling test signing on Windows Boot Manager (Bootmgr).
- Pressing the F8 or F10 key during the boot process.
- Adding or removing add-in cards (such as video or network cards), or upgrading firmware on add-in cards.
- Using a BIOS hot key during the boot process to change the boot order to something other than the hard drive.

Note
Before you begin recovery, we recommend that you determine what caused recovery. This might help prevent the problem from occurring again in the future. For instance, if you determine that an attacker has modified your computer by obtaining physical access, you can create new security policies for tracking who has physical presence. After the

recovery password has been used to recover access to the PC, BitLocker will reseal the encryption key to the current values of the measured components.

For planned scenarios, such as a known hardware or firmware upgrades, you can avoid initiating recovery by temporarily suspending BitLocker protection. Because suspending BitLocker leaves the drive fully encrypted, the administrator can quickly resume BitLocker protection after the planned task has been completed. Using suspend and resume also reseals the encryption key without requiring the entry of the recovery key.

Notes

If suspended BitLocker will automatically resume protection when the PC is rebooted, unless a reboot count is specified using the manage-bde command line tool.

If software maintenance requires the computer be restarted and you are using two-factor authentication, you can enable BitLocker Network Unlock to provide the secondary authentication factor when the computers do not have an on-premise user to provide the additional authentication method.

Recovery has been described within the context of unplanned or undesired behavior, but you can also cause recovery as an intended production scenario, in order to manage access control. For example, when you redeploy desktop or laptop computers to other departments or employees in your enterprise, you can force BitLocker into recovery before the computer is given to a new user.

Testing recovery

Before you create a thorough BitLocker recovery process, we recommend that you test how the recovery process works for both end users (people who call your helpdesk for the recovery password) and administrators (people who help the end user get the recovery password). The –forcerecovery command of manage-bde is an easy way for you to step through the recovery process before your users encounter a recovery situation.

To force a recovery for the local computer

1. Click the **Start** button, type **cmd** in the **Start Search** box, right-click **cmd.exe**, and then click **Run as administrator**.
2. At the command prompt, type the following command and then press ENTER:

   manage-bde -forcerecovery <*Volume*>

To force recovery for a remote computer

1. On the **Start** screen, type **cmd.exe**, and then click **Run as administrator**.
2. At the command prompt, type the following command and then press ENTER:

   manage-bde. -ComputerName <*ComputerName*> **-forcerecovery** <*Volume*>

Note

<*ComputerName*> represents the name of the remote computer. <*Volume*> represents the volume on the remote computer that is protected with BitLocker.

Planning your recovery process

When planning the BitLocker recovery process, first consult your organization's current best practices for recovering sensitive information. For example: How does your enterprise handle lost Windows passwords? How does your organization perform smart card PIN resets? You can use these best practices and related resources (people and tools) to help formulate a BitLocker recovery model.

Organizations that rely on BitLocker Drive Encryption and BitLocker To Go to protect data on a large number of computers and removable drives running the Windows 8 or Windows 7 operating systems and Windows to Go should consider using the Microsoft BitLocker Administration and Monitoring (MBAM) Tool version 2.0, which is included in the Microsoft Desktop Optimization Pack (MDOP) for Microsoft Software Assurance. MBAM makes BitLocker implementations easier to deploy and manage and allows administrators to provision and monitor encryption for operating system and fixed drives. MBAM prompts the user before encrypting fixed drives. MBAM also manages recovery keys for fixed and removable drives, making recovery easier to manage. MBAM can be used as part of a Microsoft System Center deployment or as a stand-alone solution. For more information, see Microsoft BitLocker Administration and Monitoring.

After a BitLocker recovery has been initiated, users can use a recovery password to unlock access to encrypted data. You must consider both self-recovery and recovery password retrieval methods for your organization.

When you determine your recovery process, you should:

- Become familiar with how you can retrieve the recovery password. See:
  - Self-recovery
  - Recovery password retrieval
- Determine a series of steps for post-recovery, including analyzing why the recovery occurred and resetting the recovery password. See:
  - Post-recovery analysis

Self-recovery

In some cases, users might have the recovery password in a printout or a USB flash drive and can perform self-recovery. We recommend that your organization create a policy for self-recovery. If self-recovery includes using a password or recovery key stored on a USB flash drive, the users should be warned not to store the USB flash drive in the same place as the PC, especially during travel, for example if both the PC and the recovery items are in the same bag it would be very easy for access to be gained to the PC by an unauthorized user. Another policy to consider is having users contact the Helpdesk before or after performing self-recovery so that the root cause can be identified.

Recovery password retrieval

If the user does not have a recovery password in a printout or on a USB flash drive, the user will need to be able to retrieve the recovery password from an online source. If the PC is a member of a domain the recovery password can be backed up to AD DS. However, this does not happen by default, you must have configured the appropriate Group Policy settings before BitLocker was enabled on the PC. BitLocker Group Policy settings can be found in the Local Group Policy Editor or the Group Policy Management Console (GPMC) under **Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption**. The following policy settings define the recovery methods that can be used to restore access to a BitLocker-protected drive if an authentication method fails or is unable to be used.

- **Choose how BitLocker-protected operating system drives can be recovered**
- **Choose how BitLocker-protected fixed drives can be recovered**
- **Choose how BitLocker-protected removable drives can be recovered**

In each of these policies select **Save BitLocker recovery information to Active Directory Domain Services** and then choose which BitLocker recovery information to store in Active Directory Domain Services (AD DS). Select the **Do not enable BitLocker until recovery information is stored in AD DS** check box if you want to prevent users from enabling BitLocker unless the computer is connected to the domain and the backup of BitLocker recovery information for the drive to AD DS succeeds.

Note

If the PCs are part of a workgroup, users should be advised to save their BitLocker recovery password with their Microsoft Account online. Having an online copy of your BitLocker recovery password is recommended to help ensure that you do not lose access to your data in the event that recovery is required.

The BitLocker Recovery Password Viewer for Active Directory Users and Computers tool allows domain administrators to view BitLocker recovery passwords for specific computer objects in Active Directory.

You can use the following list as a template for creating your own recovery process for recovery password retrieval. This sample process uses the BitLocker Recovery Password Viewer for Active Directory Users and Computers tool.

- Record the name of the user's computer
- Verify the user's identity
- Locate the recovery password in AD DS
- Gather information to determine why recovery occurred
- Give the user the recovery password

Record the name of the user's computer

You can use the name of the user's computer to locate the recovery password in AD DS. If the user does not know the name of the computer, ask the user to read the first word of the **Drive**

Label in the **BitLocker Drive Encryption Password Entry** user interface. This is the computer name when BitLocker was enabled and is probably the current name of the computer.

Verify the user's identity

You should verify that the person that is asking for the recovery password is truly the authorized user of that computer. You may also wish to verify that the computer with the name the user provided belongs to the user.

Locate the recovery password in AD DS

Locate the Computer object with the matching name in AD DS. Because Computer object names are listed in the AD DS global catalog, you should be able to locate the object even if you have a multi-domain forest.

Multiple recovery passwords

If multiple recovery passwords are stored under a computer object in AD DS, the name of the BitLocker recovery information object includes the date that the password was created.

If at any time you are unsure what password to provide, or if you think you might be providing the incorrect password, ask the user to read the eight character password ID that is displayed in the recovery console.

Since the password ID is a unique value that is associated with each recovery password stored in AD DS, running a query using this ID will find the correct password to unlock the encrypted volume.

Gather information to determine why recovery occurred

Before you give the user the recovery password, you should gather any information that will help determine why the recovery was needed, in order to analyze the root cause during the post-recovery analysis. For more information about post-recovery analysis, see Post-recovery analysis.

Give the user the recovery password

Because the recovery password is 48 digits long the user may need to record the password by writing it down or typing it on a different computer. If you are using MBAM, the recovery password will be regenerated after it is recovered from the MBAM database to avoid the security risks associated with an uncontrolled password.

Note

Because the 48-digit recovery password is long and contains a combination of digits, the user might mishear or mistype the password. The boot-time recovery console uses built-in checksum numbers to detect input errors in each 6-digit block of the 48-digit recovery password, and offers the user the opportunity to correct such errors.

Post-recovery analysis

When a volume is unlocked using a recovery password, an event is written to the event log and the platform validation measurements are reset in the TPM to match the current configuration. Unlocking the volume means that the encryption key has been released and is ready for on-the-fly encryption when data is written to the volume, and on-the-fly decryption when data is read from the volume. After the volume is unlocked, BitLocker behaves the same way, regardless of how the access was granted.

If you notice that a computer is having repeated recovery password unlocks, you might want to have an administrator can perform post-recovery analysis to determine the root cause of the recovery and refresh BitLocker platform validation so that the user no longer needs to enter a recovery password each time that the computer starts up. See:

- Determine the root cause of the recovery
- Resolve the root cause

Determine the root cause of the recovery

If a user needed to recover the drive, it is important to determine the root cause that initiated the recovery as soon as possible. Properly analyzing the state of the computer and detecting tampering may reveal threats that have broader implications for enterprise security.

While an administrator can remotely investigate the cause of recovery in some cases, the end user might need to bring the computer that contains the recovered drive on site to analyze the root cause further.

Review and answer the following questions for your organization:

1. What BitLocker protection mode is in effect (TPM, TPM + PIN, TPM + startup key, startup key only)? Which PCR profile is in use on the PC?
2. Did the user merely forget the PIN or lose the startup key? If a token was lost, where might the token be?
3. If TPM mode was in effect, was recovery caused by a boot file change?
4. If recovery was caused by a boot file change, is this due to an intended user action (for example, BIOS upgrade), or to malicious software?
5. When was the user last able to start the computer successfully, and what might have happened to the computer since then?
6. Might the user have encountered malicious software or left the computer unattended since the last successful startup?

To help you answer these questions, use the BitLocker command-line tool to view the current configuration and protection mode (for example, **manage-bde -status**). Scan the event log to find events that help indicate why recovery was initiated (for example, if boot file change occurred). Both of these capabilities can be performed remotely.

Resolve the root cause

After you have identified what caused recovery, you can reset BitLocker protection and avoid recovery on every startup.

The details of this reset can vary according to the root cause of the recovery. If you cannot determine the root cause, or if malicious software or a rootkit might have infected the computer, Helpdesk should apply best-practice virus policies to react appropriately.

Note
You can perform a BitLocker validation profile reset by suspending and resuming BitLocker.

- Unknown PIN
- Lost startup key
- Changes to boot files

Unknown PIN

If a user has forgotten the PIN, you must reset the PIN while you are logged on to the computer in order to prevent BitLocker from initiating recovery each time the computer is restarted.

To prevent continued recovery due to an unknown PIN

1. Unlock the computer using the recovery password.
2. Reset the PIN:
   a.
   b. Right-click the drive and then click **Change PIN**
   c. In the BitLocker Drive Encryption dialog, click **Reset a forgotten PIN**. If you are not logged in with an administrator account you must provide administrative credentials at this time.
   d. In the PIN reset dialog, provide and confirm the new PIN to use and then click **Finish**.
3. You will use the new PIN the next time you unlock the drive.

Lost startup key

If you have lost the USB flash drive that contains the startup key, then you must unlock the drive by using the recovery key and then create a new startup key.

To prevent continued recovery due to a lost startup key

1. Log on as an administrator to the computer that has the lost startup key.
2. From the Start screen open Manage BitLocker
3. Click **Duplicate start up key**, insert the clean USB drive on which you are going to write the key and then click **Save**.

Changes to boot files

This error might occur if you updated the firmware. As a best practice you should suspend BitLocker before making changes the firmware and then resume protection after the update has completed. This prevents the computer from going into recovery mode. However if changes were made when BitLocker protection was on you can simply log on to the computer using the recovery password and the platform validation profile will be updated so that recovery will not occur the next time.

Windows RE and BitLocker

Windows Recovery Environment (RE) can be used to recover access to a drive protected by BitLocker or by Device Encryption. If a PC is unable to boot after two failures, Startup Repair will automatically start. When Startup Repair is launched automatically due to boot failures, it will only execute operating system and driver file repairs, provided that the boot logs or any available crash dump point to a specific corrupted file. In Windows 8.1 on systems that include firmware to support specific TPM measurements for PCR[7] the TPM can validate that Windows RE is a trusted operating environment and will unlock any BitLocker-protected drives if Windows RE has not been modified. If the Windows RE environment has been modified, for example the TPM has been disabled, the drives will stay locked until the BitLocker recovery key is provided. If Startup Repair is not able to be run automatically from the PC and instead Windows RE is manually started from a repair disk, the BitLocker recovery key must be provided to unlock the BitLocker–protected drives.

Using additional recovery information

Besides the 48-digit BitLocker recovery password, other types of recovery information are stored in Active Directory. This section describes how this additional information can be used.

BitLocker key package

If the recovery methods discussed earlier in this document do not unlock the volume, you can use the BitLocker Repair tool to decrypt the volume at the block level. The tool uses the BitLocker key package to help recover encrypted data from severely damaged drives. You can then use this recovered data to salvage encrypted data, even after the correct recovery password has failed to unlock the damaged volume. We recommend that you still save the recovery password. A key package cannot be used without the corresponding recovery password.

Note

You must use the BitLocker Repair tool **repair-bde** to use the BitLocker key package.

The BitLocker key package is not saved by default. To save the package along with the recovery password in AD DS you must select the **Backup recovery password and key package** option in the Group Policy settings that control the recovery method. You can also export the key package from a working volume. For more details on how to export key packages, see Retrieving the BitLocker key package.

Resetting recovery passwords

You should invalidate a recovery password after it has been provided and used. It should also be done when you intentionally want to invalidate an existing recovery password for any reason.

You can reset the recovery password in two ways:

- **Use manage-bde** You can use manage-bde to remove the old recovery password and add a new recovery password. The procedure identifies the command and the syntax for this method.
- **Run a script** You can run a script to reset the password without decrypting the volume. The sample script in the procedure illustrates this functionality. The sample script creates a new recovery password and invalidates all other passwords.

To reset a recovery password using manage-bde

1. Remove the previous recovery password

   ```
   Manage-bde -protectors -delete C: -type RecoveryPassword
   ```

2. Add the new recovery password

   ```
   Manage-bde -protectors -add C: -RecoveryPassword
   ```

3. Get the ID of the new recovery password. From the screen copy the ID of the recovery password.

   ```
   Manage-bde -protectors -get C: -Type RecoveryPassword
   ```

4. Backup the new recovery password to AD DS

   ```
   Manage-bde -protectors -adbackup C: -id {EXAMPLE6-5507-4924-
   AA9E-AFB2EB003692}
   ```

   Warning
   You must include the braces in the ID string.

To run the sample recovery password script

1. Save the following sample script in a VBScript file. For example: ResetPassword.vbs.
2. At the command prompt, type a command similar to the following:

   cscript ResetPassword.vbs

Important

This sample script is configured to work only for the C volume. You must customize the script to match the volume where you want to test password reset.

Note

To manage a remote computer, you can specify the remote computer name rather than the local computer name.

You can use the following sample script to create a VBScript file to reset the recovery passwords.

```
' Target drive letter
```

```
strDriveLetter = "c:"

' Target computer name
' Use "." to connect to the local computer
strComputerName = "."


' ------------------------------------------------------------------------------
' Connect to the BitLocker WMI provider class
' ------------------------------------------------------------------------------

strConnectionStr = "winmgmts:" _
                 & "{impersonationLevel=impersonate,authenticationLevel=pktPrivacy}!\\" _
                 & strComputerName _
                 & "\root\cimv2\Security\MicrosoftVolumeEncryption"


On Error Resume Next 'handle permission errors

Set objWMIService = GetObject(strConnectionStr)


If Err.Number <> 0 Then
     WScript.Echo "Failed to connect to the BitLocker interface (Error 0x" &
Hex(Err.Number) & ")."
     Wscript.Echo "Ensure that you are running with administrative privileges."
     WScript.Quit -1
End If

On Error GoTo 0

strQuery = "Select * from Win32_EncryptableVolume where DriveLetter='" & strDriveLetter &
"'"
Set colTargetVolumes = objWMIService.ExecQuery(strQuery)
```

```
If colTargetVolumes.Count = 0 Then
    WScript.Echo "FAILURE: Unable to find BitLocker-capable drive " &  strDriveLetter & "
on computer " & strComputerName & "."
    WScript.Quit -1
End If


' there should only be one volume found
For Each objFoundVolume in colTargetVolumes
    set objVolume = objFoundVolume
Next


' objVolume is now our found BitLocker-capable disk volume


' ----------------------------------------------------------------------------
' Perform BitLocker WMI provider functionality
' ----------------------------------------------------------------------------


' Add a new recovery password, keeping the ID around so it doesn't get deleted later
' ----------------------------------------------------------------------------

nRC = objVolume.ProtectKeyWithNumericalPassword("Recovery Password Refreshed By Script",
, sNewKeyProtectorID)

If nRC <> 0 Then
WScript.Echo "FAILURE: ProtectKeyWithNumericalPassword failed with return code 0x" &
Hex(nRC)
WScript.Quit -1
End If
```

```
' Removes the other, "stale", recovery passwords
' -----------------------------------------------------------------------------

nKeyProtectorTypeIn = 3 ' type associated with "Numerical Password" protector

nRC = objVolume.GetKeyProtectors(nKeyProtectorTypeIn, aKeyProtectorIDs)

If nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyProtectors failed with return code 0x" & Hex(nRC)
WScript.Quit -1
End If

' Delete those key protectors other than the one we just added.

For Each sKeyProtectorID In aKeyProtectorIDs

If sKeyProtectorID <> sNewKeyProtectorID Then
nRC = objVolume.DeleteKeyProtector(sKeyProtectorID)

If nRC <> 0 Then
WScript.Echo "FAILURE: DeleteKeyProtector on ID " & sKeyProtectorID & " failed with
return code 0x" & Hex(nRC)
WScript.Quit -1
Else
' no output
'WScript.Echo "SUCCESS: Key protector with ID " & sKeyProtectorID & " deleted"
End If
End If

Next

WScript.Echo "A new recovery password has been added. Old passwords have been removed."
```

```
' - some advanced output (hidden)
'WScript.Echo ""
'WScript.Echo "Type ""manage-bde -protectors -get " & strDriveLetter & " -type
recoverypassword"" to view existing passwords."
```

Retrieving the BitLocker key package

You can use two methods to retrieve the key package, as described in Using additional recovery information:

- **Export a previously-saved key package from AD DS.** You must have Read access to BitLocker recovery passwords that are stored in AD DS.
- **Export a new key package from an unlocked, BitLocker-protected volume.** You must have local administrator access to the working volume, before any damage has occurred.

The following sample script exports all previously-saved key packages from AD DS.

To run the sample key package retrieval script

1. Save the following sample script in a VBScript file. For example: GetBitLockerKeyPackageADDS.vbs.
2. At the command prompt, type a command similar to the following:

   cscript GetBitLockerKeyPackageADDS.vbs -?

You can use the following sample script to create a VBScript file to retrieve the BitLocker key package from AD DS.

```
' ----------------------------------------------------------------------------
' Usage
' ----------------------------------------------------------------------------


Sub ShowUsage
   Wscript.Echo "USAGE: GetBitLockerKeyPackageAD [Path To Saved Key Package] [Optional
Computer Name]"
   Wscript.Echo "If no computer name is specified, the local computer is assumed."
   Wscript.Echo
   Wscript.Echo "Example: GetBitLockerKeyPackageAD E:\bitlocker-ad-key-package
mycomputer"
   WScript.Quit
End Sub
```

```
' ----------------------------------------------------------------------------
' Parse Arguments
' ----------------------------------------------------------------------------

Set args = WScript.Arguments

Select Case args.Count
  Case 1
    If args(0) = "/?" Or args(0) = "-?" Then
    ShowUsage
    Else
      strFilePath = args(0)
      ' Get the name of the local computer
      Set objNetwork = CreateObject("WScript.Network")
      strComputerName = objNetwork.ComputerName
    End If

  Case 2
    If args(0) = "/?" Or args(0) = "-?" Then
      ShowUsage
    Else
      strFilePath = args(0)
      strComputerName = args(1)
    End If
  Case Else
    ShowUsage

End Select




' ----------------------------------------------------------------------------
' Get path to Active Directory computer object associated with the computer name
' ----------------------------------------------------------------------------
```

```
Function GetStrPathToComputer(strComputerName)

    ' Uses the global catalog to find the computer in the forest
    ' Search also includes deleted computers in the tombstone

    Set objRootLDAP = GetObject("LDAP://rootDSE")
    namingContext = objRootLDAP.Get("defaultNamingContext") ' e.g. string
dc=fabrikam,dc=com

    strBase = "<GC://" & namingContext & ">"

    Set objConnection = CreateObject("ADODB.Connection")
    Set objCommand = CreateObject("ADODB.Command")
    objConnection.Provider = "ADsDSOOBject"
    objConnection.Open "Active Directory Provider"
    Set objCommand.ActiveConnection = objConnection

    strFilter = "(&(objectCategory=Computer)(cn=" &  strComputerName & "))"
    strQuery = strBase & ";" & strFilter  & ";distinguishedName;subtree"

    objCommand.CommandText = strQuery
    objCommand.Properties("Page Size") = 100
    objCommand.Properties("Timeout") = 100
    objCommand.Properties("Cache Results") = False

    ' Enumerate all objects found.

    Set objRecordSet = objCommand.Execute
    If objRecordSet.EOF Then
      WScript.echo "The computer name '" &  strComputerName & "' cannot be found."
      WScript.Quit 1
    End If
```

```
    ' Found object matching name

    Do Until objRecordSet.EOF
      dnFound = objRecordSet.Fields("distinguishedName")
      GetStrPathToComputer = "LDAP://" & dnFound
      objRecordSet.MoveNext
    Loop


    ' Clean up.
    Set objConnection = Nothing
    Set objCommand = Nothing
    Set objRecordSet = Nothing

End Function


' ----------------------------------------------------------------------------
' Securely access the Active Directory computer object using Kerberos
' ----------------------------------------------------------------------------


Set objDSO = GetObject("LDAP:")
strPathToComputer = GetStrPathToComputer(strComputerName)

WScript.Echo "Accessing object: " + strPathToComputer

Const ADS_SECURE_AUTHENTICATION = 1
Const ADS_USE_SEALING = 64 '0x40
Const ADS_USE_SIGNING = 128 '0x80


' ----------------------------------------------------------------------------
' Get all BitLocker recovery information from the Active Directory computer object
```

```
' ----------------------------------------------------------------------------

' Get all the recovery information child objects of the computer object

Set objFveInfos = objDSO.OpenDSObject(strPathToComputer, vbNullString, vbNullString, _
                                      ADS_SECURE_AUTHENTICATION + ADS_USE_SEALING +
ADS_USE_SIGNING)

objFveInfos.Filter = Array("msFVE-RecoveryInformation")

' Iterate through each recovery information object and saves any existing key packages

nCount = 1
strFilePathCurrent = strFilePath & nCount

For Each objFveInfo in objFveInfos

   strName = objFveInfo.Get("name")

   strRecoveryPassword = objFveInfo.Get("msFVE-RecoveryPassword")
   strKeyPackage = objFveInfo.Get("msFVE-KeyPackage")

   WScript.echo
   WScript.echo "Recovery Object Name: " + strName
   WScript.echo "Recovery Password: " + strRecoveryPassword

   ' Validate file path
   Set fso = CreateObject("Scripting.FileSystemObject")

   If (fso.FileExists(strFilePathCurrent)) Then
 WScript.Echo "The file " & strFilePathCurrent & " already exists. Please use a different
path."
WScript.Quit -1
   End If
```

```
    ' Save binary data to the file
    SaveBinaryDataText strFilePathCurrent, strKeyPackage

    WScript.echo "Related key package successfully saved to " + strFilePathCurrent


    ' Update next file path using base name
    nCount = nCount + 1
    strFilePathCurrent = strFilePath & nCount

Next


'-----------------------------------------------------------------------------------------
' Utility functions to save binary data
'-----------------------------------------------------------------------------------------

Function SaveBinaryDataText(FileName, ByteArray)
  'Create FileSystemObject object
  Dim FS: Set FS = CreateObject("Scripting.FileSystemObject")

  'Create text stream object
  Dim TextStream
  Set TextStream = FS.CreateTextFile(FileName)

  'Convert binary data To text And write them To the file
  TextStream.Write BinaryToString(ByteArray)
End Function

Function BinaryToString(Binary)
  Dim I, S
  For I = 1 To LenB(Binary)
    S = S & Chr(AscB(MidB(Binary, I, 1)))
```

```
  Next
  BinaryToString = S
End Function

WScript.Quit
```

The following sample script exports a new key package from an unlocked, encrypted volume.

To run this script, start by saving the code into a VBS file (for example, GetBitLockerKeyPackage.vbs). Then, open an administrator command prompt and use “cscript” to run the saved file (for example, type "cscript GetBitLockerKeyPackage.vbs -?").

```
' ----------------------------------------------------------------------------
' Usage
' ----------------------------------------------------------------------------

Sub ShowUsage
   Wscript.Echo "USAGE: GetBitLockerKeyPackage [VolumeLetter/DriveLetter:] [Path To Saved Key Package]"
   Wscript.Echo
   Wscript.Echo "Example: GetBitLockerKeyPackage C: E:\bitlocker-backup-key-package"
   WScript.Quit
End Sub

' ----------------------------------------------------------------------------
' Parse Arguments
' ----------------------------------------------------------------------------

Set args = WScript.Arguments

Select Case args.Count
  Case 2
```

```
    If args(0) = "/?" Or args(0) = "-?" Then
      ShowUsage
    Else
      strDriveLetter = args(0)
      strFilePath = args(1)
    End If
  Case Else
    ShowUsage

End Select

' ----------------------------------------------------------------------------
' Other Inputs
' ----------------------------------------------------------------------------

' Target computer name
' Use "." to connect to the local computer
strComputerName = "."

' Default key protector ID to use. Specify "" to let the script choose.

strDefaultKeyProtectorID = ""

' strDefaultKeyProtectorID = "{001298E0-870E-4BA0-A2FF-FC74758D5720}"  ' sample



' ----------------------------------------------------------------------------
' Connect to the BitLocker WMI provider class
' ----------------------------------------------------------------------------

strConnectionStr = "winmgmts:" _
                 & "{impersonationLevel=impersonate,authenticationLevel=pktPrivacy}!\\" _
                 & strComputerName _
                 & "\root\cimv2\Security\MicrosoftVolumeEncryption"
```

```
On Error Resume Next 'handle permission errors

Set objWMIService = GetObject(strConnectionStr)



If Err.Number <> 0 Then
     WScript.Echo "Failed to connect to the BitLocker interface (Error 0x" &
Hex(Err.Number) & ")."
     Wscript.Echo "Ensure that you are running with administrative privileges."
     WScript.Quit -1
End If

On Error GoTo 0

strQuery = "Select * from Win32_EncryptableVolume where DriveLetter='" & strDriveLetter &
"'"
Set colTargetVolumes = objWMIService.ExecQuery(strQuery)



If colTargetVolumes.Count = 0 Then
    WScript.Echo "FAILURE: Unable to find BitLocker-capable drive " &  strDriveLetter & "
on computer " & strComputerName & "."
    WScript.Quit -1
End If



' there should only be one volume found
For Each objFoundVolume in colTargetVolumes
    set objVolume = objFoundVolume
Next
```

```
' objVolume is now our found BitLocker-capable disk volume


' ------------------------------------------------------------------------------
' Perform BitLocker WMI provider functionality
' ------------------------------------------------------------------------------


' Collect all possible valid key protector ID's that can be used to get the package
' ------------------------------------------------------------------------------

nNumericalKeyProtectorType = 3 ' type associated with "Numerical Password" protector

nRC = objVolume.GetKeyProtectors(nNumericalKeyProtectorType, aNumericalKeyProtectorIDs)
If nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyProtectors failed with return code 0x" & Hex(nRC)
WScript.Quit -1
End If

nExternalKeyProtectorType = 2 ' type associated with "External Key" protector

nRC = objVolume.GetKeyProtectors(nExternalKeyProtectorType, aExternalKeyProtectorIDs)
If nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyProtectors failed with return code 0x" & Hex(nRC)
WScript.Quit -1
End If


' Get first key protector of the type "Numerical Password" or "External Key", if any
' ------------------------------------------------------------------------------

if strDefaultKeyProtectorID = "" Then

' Save first numerical password, if exists
```

```
If UBound(aNumericalKeyProtectorIDs) <> -1 Then
strDefaultKeyProtectorID = aNumericalKeyProtectorIDs(0)
End If

' No numerical passwords exist, save the first external key
If strDefaultKeyProtectorID = "" and UBound(aExternalKeyProtectorIDs) <> -1 Then
strDefaultKeyProtectorID = aExternalKeyProtectorIDs(0)
End If

' Fail case: no recovery key protectors exist.
If strDefaultKeyProtectorID = "" Then
WScript.Echo "FAILURE: Cannot create backup key package because no recovery passwords or
recovery keys exist. Check that BitLocker protection is on for this drive."
WScript.Echo "For help adding recovery passwords or recovery keys, type ""manage-bde -
protectors -add -?""."
WScript.Quit -1
End If

End If

' Get some information about the chosen key protector ID
' ------------------------------------------------------------------------------

' is the type valid?

nRC = objVolume.GetKeyProtectorType(strDefaultKeyProtectorID, nDefaultKeyProtectorType)

If Hex(nRC) = "80070057" Then
WScript.Echo "The key protector ID " & strDefaultKeyProtectorID & " is not valid."
WScript.Echo "This ID value may have been provided by the script writer."
ElseIf nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyProtectorType failed with return code 0x" & Hex(nRC)
WScript.Quit -1
End If
```

```
' what's a string that can be used to describe it?

strDefaultKeyProtectorType = ""

Select Case nDefaultKeyProtectorType

  Case nNumericalKeyProtectorType
      strDefaultKeyProtectorType = "recovery password"

  Case nExternalKeyProtectorType
      strDefaultKeyProtectorType = "recovery key"

  Case Else
      WScript.Echo "The key protector ID " & strDefaultKeyProtectorID & " does not refer
to a valid recovery password or recovery key."
      WScript.Echo "This ID value may have been provided by the script writer."

End Select


' Save the backup key package using the chosen key protector ID
' ------------------------------------------------------------------------------

nRC = objVolume.GetKeyPackage(strDefaultKeyProtectorID, oKeyPackage)
If nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyPackage failed with return code 0x" & Hex(nRC)
WScript.Quit -1
End If

' Validate file path
Set fso = CreateObject("Scripting.FileSystemObject")
If (fso.FileExists(strFilePath)) Then
WScript.Echo "The file " & strFilePath & " already exists. Please use a different path."
```

```
WScript.Quit -1
End If

Dim oKeyPackageByte, bKeyPackage
For Each oKeyPackageByte in oKeyPackage
  'WScript.echo "key package byte: " & oKeyPackageByte
  bKeyPackage = bKeyPackage & ChrB(oKeyPackageByte)
Next

' Save binary data to the file
SaveBinaryDataText strFilePath, bKeyPackage

' Display helpful information
' ----------------------------------------------------------------------------

WScript.Echo "The backup key package has been saved to " & strFilePath & "."

WScript.Echo "IMPORTANT: To use this key package, the " & strDefaultKeyProtectorType & "
must also be saved."

' Display the recovery password or a note about saving the recovery key file

If nDefaultKeyProtectorType = nNumericalKeyProtectorType Then

nRC = objVolume.GetKeyProtectorNumericalPassword(strDefaultKeyProtectorID,
sNumericalPassword)
If nRC <> 0 Then
WScript.Echo "FAILURE: GetKeyProtectorNumericalPassword failed with return code 0x" &
Hex(nRC)
WScript.Quit -1
End If
WScript.Echo "Save this recovery password: " & sNumericalPassword

ElseIf nDefaultKeyProtectorType = nExternalKeyProtectorType Then
```

```
WScript.Echo "The saved key file is named " & strDefaultKeyProtectorID & ".BEK"
WScript.Echo "For help re-saving this external key file, type ""manage-bde -protectors -get -?"""
End If


'-----------------------------------------------------------------------------------------
' Utility functions to save binary data
'-----------------------------------------------------------------------------------------

Function SaveBinaryDataText(FileName, ByteArray)
  'Create FileSystemObject object
  Dim FS: Set FS = CreateObject("Scripting.FileSystemObject")

  'Create text stream object
  Dim TextStream
  Set TextStream = FS.CreateTextFile(FileName)

  'Convert binary data To text And write them To the file
  TextStream.Write BinaryToString(ByteArray)
End Function

Function BinaryToString(Binary)
  Dim I, S
  For I = 1 To LenB(Binary)
    S = S & Chr(AscB(MidB(Binary, I, 1)))
  Next
  BinaryToString = S
End Function
```

See also

- BitLocker Overview

BCD Settings and BitLocker

When protecting data at rest on an operating system volume, during the boot process BitLocker verifies that the security sensitive boot configuration data (BCD) settings have not changed since BitLocker was last enabled, resumed, or recovered.

BitLocker and BCD Settings

In Windows 7 and Windows Server 2008 R2, BitLocker validated nearly all BCD settings with the winload, winresume, and memtest prefixes. However, this high degree of validation caused BitLocker to go into recovery mode for benign setting changes, for example, when applying a language pack BitLocker would enter recovery.

In Windows 8, Windows Server 2012, and later operating systems BitLocker narrows the set of BCD settings validated to reduce the chance of benign changes causing a BCD validation problem. If you believe that there is a risk in excluding a particular BCD setting from the validation profile, you can increase BCD validation coverage to suit your validation preferences. Alternatively, if a default BCD setting is persistently triggering recovery for benign changes, then you can exclude that BCD setting from the validation profile.

When secure boot is enabled

Computers with UEFI firmware can use Secure Boot to provide enhanced boot security. When BitLocker is able to use Secure Boot for platform and BCD integrity validation, as defined by the **Allow Secure Boot for integrity validation** group policy setting, the **Use enhanced Boot Configuration Data validation profile** group policy is ignored.

One of the benefits of using Secure Boot is that it can correct BCD settings during boot without triggering recovery events. Secure Boot enforces the same BCD settings as BitLocker. Secure Boot BCD enforcement is not configurable from within the operating system.

Customizing BCD validation settings

To modify the BCD settings BitLocker validates the IT Pro will add or exclude BCD settings from the platform validation profile by enabling and configuring the **Use enhanced Boot Configuration Data validation profile** Group Policy setting.

For the purposes of BitLocker validation, BCD settings are associated with a specific set of Microsoft boot applications. BCD settings are either associated with a specific boot application or can apply to all boot applications by associating a prefix to the BCD setting entered in the Group Policy setting. Prefix values include:

- winload
- winresume
- memtest
- all

All BCD settings are specified by combining the prefix value with either a hexadecimal (hex) value or a “friendly name.”

The BCD setting hex value is reported when BitLocker enters recovery mode and is stored in the event log (event ID 523). The hex value uniquely identifies which BCD setting caused the recovery event.

You can quickly obtain the friendly name for the BCD settings on your computer by using the command “`bcdedit.exe /enum all`”.

Not all BCD settings have friendly names, for those settings the hex value is the only way to configure an exclusion policy.

When specifying BCD values in the **Use enhanced Boot Configuration Data validation profile** Group Policy setting, use the following syntax:

- Prefix the setting with the boot application prefix
- Append a colon ‘:’
- Append either the hex value or the friendly name
- If entering more than one BCD setting, you will need to enter each BCD setting on a new line

For example, either “`winload:hypervisordebugport`” or “`winload:0x250000f4`” yield the same value.

Setting that applies to all boot applications may be applied only to an individual application, however the reverse is not true. For example, one can specify either: “`all:locale`” or “`winresume:locale`”, but as the bcd setting “`win-pe`” does not apply to all boot applications, “`winload:winpe`” is valid, but “`all:winpe`” is not valid. The setting that controls boot debugging (“`bootdebug`” or 0x16000010) will always be validated and will have no effect if it is included in the provided fields.

Note

Take care when configuring BCD entries in the Group Policy setting. The Local Group Policy Editor does not validate the correctness of the BCD entry. BitLocker will fail to be enabled if the Group Policy setting specified is invalid.

Default BCD validation profile

The following table contains the default BCD validation profile used by BitLocker in Windows 8, Windows Server 2012, and later operating systems:

Hex Value	Prefix	Friendly Name
0x11000001	all	device
0x12000002	all	path
0x12000030	all	loadoptions
0x16000010	all	bootdebug
0x16000040	all	advancedoptions

Hex Value	Prefix	Friendly Name
0x16000041	all	optionsedit
0x16000048	all	nointegritychecks
0x16000049	all	testsigning
0x16000060	all	isolatedcontext
0x1600007b	all	forcefipscrypto
0x22000002	winload	systemroot
0x22000011	winload	kernel
0x22000012	winload	hal
0x22000053	winload	evstore
0x25000020	winload	nx
0x25000052	winload	restrictapiccluster
0x26000022	winload	winpe
0x26000025	winload	lastknowngood
0x26000081	winload	safebootalternateshell
0x260000a0	winload	debug
0x260000f2	winload	hypervisordebug
0x26000116	winload	hypervisorusevapic
0x21000001	winresume	filedevice
0x22000002	winresume	filepath
0x26000006	winresume	debugoptionenabled

Full list of friendly names for ignored BCD settings

This following is a full list of BCD settings with friendly names which are ignored by default. These settings are not part of the default BitLocker validation profile, but can be added if you see a need to validate any of these settings before allowing a BitLocker–protected operating system drive to be unlocked.

Note

Additional BCD settings exist that have hex values but do not have friendly names. These settings are not included in this list.

Hex Value	Prefix	Friendly Name
0x12000004	all	description
0x12000005	all	locale
0x12000016	all	targetname
0x12000019	all	busparams
0x1200001d	all	key
0x1200004a	all	fontpath
0x14000006	all	inherit
0x14000008	all	recoverysequence
0x15000007	all	truncatememory
0x1500000c	all	firstmegabytepolicy
0x1500000d	all	relocatephysical
0x1500000e	all	avoidlowmemory
0x15000011	all	debugtype
0x15000012	all	debugaddress
0x15000013	all	debugport
0x15000014	all	baudrate
0x15000015	all	channel
0x15000018	all	debugstart
0x1500001a	all	hostip
0x1500001b	all	port
0x15000022	all	emsport
0x15000023	all	emsbaudrate
0x15000042	all	keyringaddress
0x15000047	all	configaccesspolicy
0x1500004b	all	integrityservices
0x1500004c	all	volumebandid
0x15000051	all	initialconsoleinput
0x15000052	all	graphicsresolution

Hex Value	Prefix	Friendly Name
0x15000065	all	displaymessage
0x15000066	all	displaymessageoverride
0x16000009	all	recoveryenabled
0x1600000b	all	badmemoryaccess
0x1600000f	all	traditionalkseg
0x16000017	all	noumex
0x1600001c	all	dhcp
0x1600001e	all	vm
0x16000020	all	bootems
0x16000046	all	graphicsmodedisabled
0x16000050	all	extendedinput
0x16000053	all	restartonfailure
0x16000054	all	highestmode
0x1600006c	all	bootuxdisabled
0x16000072	all	nokeyboard
0x16000074	all	bootshutdowndisabled
0x1700000a	all	badmemorylist
0x17000077	all	allowedinmemorysettings
0x31000003	all	ramdisksdidevice
0x32000004	all	ramdisksdipath
0x35000001	all	ramdiskimageoffset
0x35000002	all	ramdisktftpclientport
0x35000005	all	ramdiskimagelength
0x35000007	all	ramdisktftpblocksize
0x35000008	all	ramdisktftpwindowsize
0x36000006	all	exportascd
0x36000009	all	ramdiskmcenabled
0x3600000a	all	ramdiskmctftpfallback

Hex Value	Prefix	Friendly Name
0x3600000b	all	ramdisktftpvarwindow
0x21000001	winload	osdevice
0x22000013	winload	dbgtransport
0x220000f9	winload	hypervisorbusparams
0x22000110	winload	hypervisorusekey
0x23000003	winload	resumeobject
0x25000021	winload	pae
0x25000031	winload	removememory
0x25000032	winload	increaseuserva
0x25000033	winload	perfmem
0x25000050	winload	clustermodeaddressing
0x25000055	winload	x2apicpolicy
0x25000061	winload	numproc
0x25000063	winload	configflags
0x25000066	winload	groupsize
0x25000071	winload	msi
0x25000072	winload	pciexpress
0x25000080	winload	safeboot
0x250000a6	winload	tscsyncpolicy
0x250000c1	winload	driverloadfailurepolicy
0x250000c2	winload	bootmenupolicy
0x250000e0	winload	bootstatuspolicy
0x250000f0	winload	hypervisorlaunchtype
0x250000f3	winload	hypervisordebugtype
0x250000f4	winload	hypervisordebugport
0x250000f5	winload	hypervisorbaudrate
0x250000f6	winload	hypervisorchannel
0x250000f7	winload	bootux

Hex Value	Prefix	Friendly Name
0x250000fa	winload	hypervisornumproc
0x250000fb	winload	hypervisorrootprocpernode
0x250000fd	winload	hypervisorhostip
0x250000fe	winload	hypervisorhostport
0x25000100	winload	tpmbootentropy
0x25000113	winload	hypervisorrootproc
0x25000115	winload	hypervisoriommupolicy
0x25000120	winload	xsavepolicy
0x25000121	winload	xsaveaddfeature0
0x25000122	winload	xsaveaddfeature1
0x25000123	winload	xsaveaddfeature2
0x25000124	winload	xsaveaddfeature3
0x25000125	winload	xsaveaddfeature4
0x25000126	winload	xsaveaddfeature5
0x25000127	winload	xsaveaddfeature6
0x25000128	winload	xsaveaddfeature7
0x25000129	winload	xsaveremovefeature
0x2500012a	winload	xsaveprocessorsmask
0x2500012b	winload	xsavedisable
0x25000130	winload	claimedtpmcounter
0x26000004	winload	stampdisks
0x26000010	winload	detecthal
0x26000024	winload	nocrashautoreboot
0x26000030	winload	nolowmem
0x26000040	winload	vga
0x26000041	winload	quietboot
0x26000042	winload	novesa
0x26000043	winload	novga

Hex Value	Prefix	Friendly Name
0x26000051	winload	usephysicaldestination
0x26000054	winload	uselegacyapicmode
0x26000060	winload	onecpu
0x26000062	winload	maxproc
0x26000064	winload	maxgroup
0x26000065	winload	groupaware
0x26000070	winload	usefirmwarepcisettings
0x26000090	winload	bootlog
0x26000091	winload	sos
0x260000a1	winload	halbreakpoint
0x260000a2	winload	useplatformclock
0x260000a3	winload	forcelegacyplatform
0x260000a4	winload	useplatformtick
0x260000a5	winload	disabledynamictick
0x260000b0	winload	ems
0x260000c3	winload	onetimeadvancedoptions
0x260000c4	winload	onetimeoptionsedit
0x260000e1	winload	disableelamdrivers
0x260000f8	winload	hypervisordisableslat
0x260000fc	winload	hypervisoruselargevtlb
0x26000114	winload	hypervisordhcp
0x21000005	winresume	associatedosdevice
0x25000007	winresume	bootux
0x25000008	winresume	bootmenupolicy
0x26000003	winresume	customsettings
0x26000004	winresume	pae
0x25000001	memtest	passcount
0x25000002	memtest	testmix

Hex Value	Prefix	Friendly Name
0x25000005	memtest	stridefailcount
0x25000006	memtest	invcfailcount
0x25000007	memtest	matsfailcount
0x25000008	memtest	randfailcount
0x25000009	memtest	chckrfailcount
0x26000003	memtest	cacheenable
0x26000004	memtest	failuresenabled

BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker

BitLocker Drive Encryption Tools include the command line tools manage-bde and repair-bde and the BitLocker cmdlets for Windows PowerShell.

Both manage-bde and the BitLocker cmdlets can be used to perform any task that can be accomplished through the BitLocker control panel and are appropriate to use for automated deployments and other scripting scenarios.

Repair-bde is a special circumstance tool that is provided for disaster recovery scenarios in which a BitLocker protected drive cannot be unlocked normally or using the recovery console.

1. Manage-bde
2. Repair-bde
3. BitLocker cmdlets for Windows PowerShell

Manage-bde

Manage-bde is a command-line tool that can be used for scripting BitLocker operations. Manage-bde offers additional options not displayed in the BitLocker control panel. For a complete list of the manage-bde options, see the **Manage-bde** command-line reference.

Manage-bde includes less default settings and requires greater customization for configuring BitLocker. For example, using just the `manage-bde -on` command on a data volume will fully encrypt the volume without any authenticating protectors. A volume encrypted in this manner still requires user interaction to turn on BitLocker protection, even though the command successfully completed because an authentication method needs to be added to the volume for it to be fully protected. The following sections provide examples of common usage scenarios for manage-bde.

Using manage-bde with operating system volumes

Listed below are examples of basic valid commands for operating system volumes. In general, using only the `manage-bde -on <drive letter>` command will encrypt the operating system volume with a TPM-only protector and no recovery key. However, many environments require more secure protectors such as passwords or PIN and expect to be able to recover information with a recovery key. It is recommended that at least one primary protector and a recovery protector be added to an operating system volume.

A good practice when using manage-bde is to determine the volume status on the target system. Use the following command to determine volume status:

```
manage-bde -status
```

This command returns the volumes on the target, current encryption status and volume type (operating system or data) for each volume.

The following example illustrates enabling BitLocker on a computer without a TPM chip. Before beginning the encryption process you must create the startup key needed for BitLocker and save it to the USB drive. When BitLocker is enabled for the operating system volume, the BitLocker will need to access the USB flash drive to obtain the encryption key (in this example, the drive letter E represents the USB drive). You will be prompted to reboot to complete the encryption process.

```
manage-bde –protectors -add C: -startupkey E:

manage-bde -on C:
```

Note

After the encryption is completed, the USB startup key must be inserted before the operating system can be started.

An alternative to the startup key protector on non-TPM hardware is to use a password and an **ADaccountorgroup** protector to protect the operating system volume. In this scenario, you would add the protectors first. This is done with the command:

```
manage-bde -protectors -add C: -pw -sid <user or group>
```

This command will require you to enter and then confirm the password protector before adding them to the volume. With the protectors enabled on the volume, you can then turn BitLocker on.

On computers with a TPM it is possible to encrypt the operating system volume without any defined protectors using manage-bde. The command to do this is:

```
manage-bde -on C:
```

This will encrypt the drive using the TPM as the default protector. If you are not sure if a TPM protector is available, to list the protectors available for a volume, run the following command:

```
manage-bde -protectors -get <volume>
```

Using manage-bde with data volumes

Data volumes use the same syntax for encryption as operating system volumes but they do not require protectors for the operation to complete. Encrypting data volumes can be done using the base command: `manage-bde -on <drive letter>` or you can choose to add additional protectors to the volume first. It is recommended that at least one primary protector and a recovery protector be added to a data volume.

A common protector for a data volume is the password protector. In the example below, we add a password protector to the volume and turn BitLocker on.

```
manage-bde -protectors -add -pw C:

manage-bde -on C:
```

Note

For more information about the types of protectors available and the volume types they are compatible with, see.

Repair-bde

You may experience a problem that damages an area of a hard disk on which BitLocker stores critical information. This kind of problem may be caused by a hard disk failure or if Windows exits unexpectedly.

The BitLocker Repair Tool (Repair-bde) can be used to access encrypted data on a severely damaged hard disk if the drive was encrypted by using BitLocker. Repair-bde can reconstruct critical parts of the drive and salvage recoverable data as long as a valid recovery password or recovery key is used to decrypt the data. If the BitLocker metadata data on the drive has become corrupt, you must be able to supply a backup key package in addition to the recovery password or recovery key. This key package is backed up in Active Directory Domain Services (AD DS) if you used the default setting for AD DS backup. With this key package and either the recovery password or recovery key, you can decrypt portions of a BitLocker-protected drive if the disk is corrupted. Each key package will work only for a drive that has the corresponding drive identifier. You can use the BitLocker Recovery Password Viewer to obtain this key package from AD DS.

Tip

If you are not backing up recovery information to AD DS or if you want to save key packages alternatively, you can use the command `manage-bde -KeyPackage` to generate a key package for a volume.

The Repair-bde command-line tool is intended for use when the operating system does not start or when you cannot start the BitLocker Recovery Console. You should use Repair-bde if the following conditions are true:

1. You have encrypted the drive by using BitLocker Drive Encryption.
2. Windows does not start, or you cannot start the BitLocker recovery console.
3. You do not have a copy of the data that is contained on the encrypted drive.

Note

Damage to the drive may not be related to BitLocker. Therefore, we recommend that you try other tools to help diagnose and resolve the problem with the drive before you use the BitLocker Repair Tool. The Windows Recovery Environment (Windows RE) provides additional options to repair computers.

The following limitations exist for Repair-bde:

- The Repair-bde command-line tool cannot repair a drive that failed during the encryption or decryption process.
- The Repair-bde command-line tool assumes that if the drive has any encryption, then the drive has been fully encrypted.

For more information about using repair-bde see **Repair-bde**

BitLocker cmdlets for Windows PowerShell

Windows PowerShell cmdlets provide a new way for administrators to use when working with BitLocker. Using Windows PowerShell's scripting capabilities, administrators can integrate BitLocker options into existing scripts with ease. The list below displays the available BitLocker cmdlets.

Name	Parameters
Add-BitLockerKeyProtector	-ADAccountOrGroup -ADAccountOrGroupProtector -Confirm -MountPoint -Password -PasswordProtector -Pin -RecoveryKeyPath -RecoveryKeyProtector -RecoveryPassword -RecoveryPasswordProtector -Service -StartupKeyPath -StartupKeyProtector -TpmAndPinAndStartupKeyProtector -TpmAndPinProtector -TpmAndStartupKeyProtector -TpmProtector

	-WhatIf
Backup-BitLockerKeyProtector	-Confirm -KeyProtectorId -MountPoint -WhatIf
Disable-BitLocker	-Confirm -MountPoint -WhatIf
Disable-BitLockerAutoUnlock	-Confirm -MountPoint -WhatIf
Enable-BitLocker	-AdAccountOrGroup -AdAccountOrGroupProtector -Confirm -EncryptionMethod -HardwareEncryption -Password -PasswordProtector -Pin -RecoveryKeyPath -RecoveryKeyProtector -RecoveryPassword -RecoveryPasswordProtector -Service -SkipHardwareTest -StartupKeyPath -StartupKeyProtector -TpmAndPinAndStartupKeyProtector -TpmAndPinProtector -TpmAndStartupKeyProtector -TpmProtector -UsedSpaceOnly -WhatIf
Enable-BitLockerAutoUnlock	-Confirm

	-MountPoint -WhatIf
Get-BitLockerVolume	-MountPoint
Lock-BitLocker	-Confirm -ForceDismount -MountPoint -WhatIf
Remove-BitLockerKeyProtector	-Confirm -KeyProtectorId -MountPoint -WhatIf
Resume-BitLocker	-Confirm -MountPoint -WhatIf
Suspend-BitLocker	-Confirm -MountPoint -RebootCount -WhatIf
Unlock-BitLocker	-AdAccountOrGroup -Confirm -MountPoint -Password -RecoveryKeyPath -RecoveryPassword -RecoveryPassword -WhatIf

Similar to manage-bde, the Windows PowerShell cmdlets allow configuration beyond the options offered in the control panel. As with manage-bde, users need to consider the specific needs of the volume they are encrypting prior to running Windows PowerShell cmdlets.

A good initial step is to determine the current state of the volume(s) on the computer. You can do this using the `Get-BitLockerVolume` cmdlet.

The `Get-BitLockerVolume` cmdlet output gives information on the volume type, protectors, protection status and other details.

Tip

Occasionally, all protectors may not be shown when using `Get-BitLockerVolume` due to lack of space in the output display. If you do not see all of the protectors for a volume, you can use the Windows PowerShell pipe command (|) to format a full listing of the protectors.

```
Get-BitLockerVolume C: | fl
```

If you want to remove the existing protectors prior to provisioning BitLocker on the volume, you could use the `Remove-BitLockerKeyProtector` cmdlet. Accomplishing this requires the GUID associated with the protector to be removed.

A simple script can pipe the values of each Get-BitLockerVolume return out to another variable as seen below:

```
$vol = Get-BitLockerVolume

$keyprotectors = $vol.KeyProtector
```

Using this, you can display the information in the $keyprotectors variable to determine the GUID for each protector.

Using this information, you can then remove the key protector for a specific volume using the command:

```
Remove-BitLockerKeyProtector <volume>: -KeyProtectorID "{GUID}"
```

Note

The BitLocker cmdlet requires the key protector GUID enclosed in quotation marks to execute. Ensure the entire GUID, with braces, is included in the command.

Using the BitLocker Windows PowerShell cmdlets with operating system volumes

Using the BitLocker Windows PowerShell cmdlets is similar to working with the manage-bde tool for encrypting operating system volumes. Windows PowerShell offers users a lot of flexibility. For example, users can add the desired protector as part command for encrypting the volume. Below are examples of common user scenarios and steps to accomplish them in BitLocker Windows PowerShell.

The following example shows how to enable BitLocker on an operating system drive using only the TPM protector:

```
Enable-BitLocker C:
```

In the example below, adds one additional protector, the StartupKey protector and chooses to skip the BitLocker hardware test. In this example, encryption starts immediately without the need for a reboot.

```
Enable-BitLocker C: -StartupKeyProtector -StartupKeyPath <path> -SkipHardwareTest
```

Using the BitLocker Windows PowerShell cmdlets with data volumes

Data volume encryption using Windows PowerShell is the same as for operating system volumes. You should add the desired protectors prior to encrypting the volume. The following example adds a password protector to the E: volume using the variable $pw as the password. The $pw variable is held as a SecureString value to store the user defined password.

```
$pw = Read-Host -AsSecureString

<user inputs password>

Enable-BitLockerKeyProtector E: -PasswordProtector -Password $pw
```

Using an AD Account or Group protector in Windows PowerShell

A new protector in available in Windows 8 and **win8_server_2** is the **ADAccountOrGroup** protector, an Active Directory SID-based protector. This protector can be added to both operating system and data volumes, although it does not unlock operating system volumes in the pre-boot environment. The protector requires the SID for the domain account or group to link with the protector. BitLocker can protect a cluster-aware disk by adding a SID-based protector for the Cluster Name Object (CNO) that lets the disk properly failover to and be unlocked by any member computer of the cluster.

Warning

The **ADAccountOrGroup** protector requires the use of an additional protector for use (such as TPM, PIN, or recovery key) when used on operating system volumes

To add an **ADAccountOrGroup** protector to a volume requires either the actual domain SID or the group name preceded by the domain and a backslash. In the example below, the CONTOSO\Administrator account is added as a protector to the data volume G.

```
Enable-BitLocker G: -AdAccountOrGroupProtector -AdAccountOrGroup CONTOSO\Administrator
```

For users who wish to use the SID for the account or group, the first step is to determine the SID associated with the account. To get the specific SID for a user account in Windows PowerShell, use the following command:

Note

Use of this command requires the RSAT-AD-PowerShell feature.

```
get-aduser -filter {samaccountname -eq "administrator"}
```

Tip

In addition to the PowerShell command above, information about the locally logged on user and group membership can be found using: WHOAMI /ALL. This does not require the use of additional features.

The following example adds an **ADAccountOrGroup** protector to the previously encrypted operating system volume using the SID of the account:

```
Add-BitLockerKeyProtector C: -ADAccountOrGroupProtector -ADAccountOrGroup S-1-5-21-
3651336348-8937238915-291003330-500
```

Note
Active Directory-based protectors are normally used to unlock Failover Cluster enabled volumes.

More information

BitLocker Overview

Encrypted Hard Drive

BitLocker Frequently Asked Questions (FAQ)

What's New in BitLocker for Windows 8 and Windows Server 2012

Prepare your organization for BitLocker: Planning and Policies

BitLocker: How to enable Network Unlock

BitLocker: How to deploy on Windows Server 2012

BitLocker: Use BitLocker Recovery Password Viewer

The BitLocker Recovery Password Viewer tool is an optional tool included with the Remote Server Administration Tools (RSAT) for Windows Server 2012 that are available to install when you install the BitLocker feature. This tool lets you locate and view BitLocker recovery passwords that are stored in Active Directory Domain Services (AD DS). You can use this tool to help recover data that is stored on a drive that has been encrypted by using BitLocker. The BitLocker Active Directory Recovery Password Viewer tool is an extension for the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in. Using this tool, you can examine a computer object's **Properties** dialog box to view the corresponding BitLocker recovery passwords. Additionally, you can right-click a domain container and then search for a BitLocker recovery password across all the domains in the Active Directory forest. You can also search for a password by password identifier (ID).

Before you start

To complete the procedures in this scenario:

- You must have domain administrator credentials.
- Your test computers must be joined to the domain.
- On the test computers, BitLocker must have been turned on after joining the domain.

The following procedures describe the most common tasks performed by using the BitLocker Recovery Password Viewer.

To view the recovery passwords for a computer

1. In **Active Directory Users and Computers**, locate and then click the container in which the computer is located.
2. Right-click the computer object, and then click **Properties**.
3. In the **Properties** dialog box, click the **BitLocker Recovery** tab to view the BitLocker recovery passwords that are associated with the computer.

To copy the recovery passwords for a computer

1. Follow the steps in the previous procedure to view the BitLocker recovery passwords.
2. On the **BitLocker Recovery** tab of the **Properties** dialog box, right-click the BitLocker recovery password that you want to copy, and then click **Copy Details**.
3. Press CTRL+V to paste the copied text to a destination location, such as a text file or spreadsheet.

To locate a recovery password by using a password ID

1. In Active Directory Users and Computers, right-click the domain container, and then click **Find BitLocker Recovery Password**.
2. In the **Find BitLocker Recovery Password** dialog box, type the first eight characters of the recovery password in the **Password ID (first 8 characters)** box, and then click **Search**.

By completing the procedures in this scenario, you have viewed and copied the recovery passwords for a computer and used a password ID to locate a recovery password.

More information

BitLocker Overview

Encrypted Hard Drive

BitLocker Frequently Asked Questions (FAQ)

What's New in BitLocker for Windows 8 and Windows Server 2012

Prepare your organization for BitLocker: Planning and Policies

BitLocker: How to deploy on Windows Server 2012

BitLocker: Use BitLocker Drive Encryption Tools to manage BitLocker

Credentials Protection and Management

This topic for the IT professional discusses features and methods introduced in Windows Server 2012 R2 and Windows 8.1 for credential protection and domain authentication controls to reduce credential theft.

Restricted Admin mode for Remote Desktop Connection

Restricted Admin mode provides a method of interactively logging on to a remote host server without transmitting your credentials to the server. This prevents your credentials from being harvested during the initial connection process if the server has been compromised.

Using this mode with administrator credentials, the remote desktop client attempts to interactively logon to a host that also supports this mode without sending credentials. When the host verifies that the user account connecting to it has administrator rights and supports Restricted Admin mode, the connection succeeds. Otherwise, the connection attempt fails. Restricted Admin mode does not at any point send plain text or other re-usable forms of credentials to remote computers.

For more information, see What's New in Remote Desktop Services in Windows Server 2012 R2.

LSA protection

The Local Security Authority (LSA), which resides within the Local Security Authority Security Service (LSASS) process, validates users for local and remote sign-ins and enforces local security policies. The Windows 8.1 operating system provides additional protection for the LSA to prevent code injection by non-protected processes. This provides added security for the credentials that the LSA stores and manages. This protected process setting for LSA can be configured in Windows 8.1 but is on by default in Windows RT 8.1 and cannot be changed.

For information about configuring LSA protection, see Configuring Additional LSA Protection.

Protected Users security group

This new domain global group triggers new non-configurable protection on devices and host computers running Windows Server 2012 R2 and Windows 8.1. The Protected Users group enables additional protections for domain controllers and domains in Windows Server 2012 R2 domains. This greatly reduces the types of credentials available when users are signed in to computers on the network from a non-compromised computer.

Members of the Protected Users group are limited further by the following methods of authentication:

- A member of the Protected Users group can only sign on using the Kerberos protocol. The account cannot authenticate using NTLM, Digest Authentication, or CredSSP. On a device running Windows 8.1, passwords are not cached, so the device that uses any one of these Security Support Providers (SSPs) will fail to authenticate to a domain when the account is a member of the Protected User group.
- The Kerberos protocol will not use the weaker DES or RC4 encryption types in the preauthentication process. This means that the domain must be configured to support at least the AES cypher suite.
- The user's account cannot be delegated with Kerberos constrained or unconstrained delegation. This means that former connections to other systems may fail if the user is a member of the Protected Users group.
- The default Kerberos Ticket Granting Tickets (TGTs) lifetime setting of four hours is configurable using Authentication Policies and Silos accessed through the Active Directory

Administrative Center (ADAC). This means that when four hours has passed, the user must authenticate again.

For more information about this group, see Protected Users Security Group.

Authentication Policy and Authentication Policy Silos

Forest-based Active Directory policies are introduced, and they can be applied to accounts in a domain with a Windows Server 2012 R2 domain functional level. These authentication policies can control which hosts a user can use to sign in. They work in conjunction with the Protect Users security group, and admins can apply access control conditions for authentication to the accounts. These authentication policies isolate related accounts to constrain the scope of a network.

The new Active Directory object class, Authentication Policy, allows you to apply authentication configuration to account classes in domains with a Windows Server 2012 R2 domain functional level. Authentication policies are enforced during the Kerberos AS or the TGS exchange. Active Directory account classes are:

- User
- Computer
- Managed Service Account
- Group Managed Service Account

For more information, see Authentication Policies and Authentication Policy Silos.

For more information how to configure protected accounts, see How to Configure Protected Accounts.

See also

For more information about the LSA and the LSASS, see the Windows Logon and Authentication Technical Overview.

For more information about how Windows manages credentials, see Cached and Stored Credentials Technical Overview.

Configuring Additional LSA Protection

This topic for the IT professional explains how to configure additional protection for the Local Security Authority (LSA) process to prevent code injection that could compromise credentials.

The LSA, which includes the Local Security Authority Server Service (LSASS) process, validates users for local and remote sign-ins and enforces local security policies. The Windows 8.1 operating system provides additional protection for the LSA to prevent reading memory and code injection by non-protected processes. This provides added security for the credentials that the LSA stores and manages. The protected process setting for LSA can be configured in Windows 8.1, but it cannot be configured in Windows RT 8.1. When this setting is used in conjunction with

Secure Boot, additional protection is achieved because disabling the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa registry key has no effect.

Protected process requirements for plug-ins or drivers

For an LSA plug-in or driver to successfully load as a protected process, it must meet the following criteria:

1. Signature verification

   Protected mode requires that any plug-in that is loaded into the LSA is digitally signed with a Microsoft signature. Therefore, any plug-ins that are unsigned or are not signed with a Microsoft signature will fail to load in LSA. Examples of these plug-ins are smart card drivers, cryptographic plug-ins, and password filters.

   LSA plug-ins that are drivers, such as smart card drivers, need to be signed by using the WHQL Certification. For more information, see WHQL Release Signature (Windows Drivers).

   Warning
   Currently, WHQL is the only supported method to sign plug-ins and drivers with a Microsoft signature.

2. Adherence to the Microsoft Security Development Lifecycle (SDL) process guidance

   All of the plug-ins must conform to the applicable SDL process guidance. For more information, see the Microsoft Security Development Lifecycle (SDL) Appendix.

   Even if the plug-ins are properly signed with a Microsoft signature, non-compliance with the SDL process can result in failure to load a plug-in.

Recommended practices

Use the following list to thoroughly test that LSA protection is enabled before you broadly deploy the feature:

- Identify all of the LSA plug-ins and drivers that are in use within your organization. This includes non-Microsoft drivers or plug-ins such as smart card drivers and cryptographic plug-ins, and any internally developed software that is used to enforce password filters or password change notifications.
- Ensure that all of the LSA plug-ins are digitally signed with a Microsoft certificate so that the plug-in will not fail to load.
- Ensure that all of the correctly signed plug-ins can successfully load into LSA and that they perform as expected.
- Use the audit logs to identify LSA plug-ins and drivers that fail to run as a protected process.

How to identify LSA plug-ins and drivers that fail to run as a protected process

The events described in this section are located in the Operational log under Applications and Services Logs\Microsoft\Windows\CodeIntegrity. They can help you identify LSA plug-ins and

drivers that are failing to load due to signing reasons. To manage these events, you can use the **wevtutil** command-line tool. For information about this tool, see **Wevtutil**.

Before opting in: How to identify plug-ins and drivers loaded by the lsass.exe

You can use the audit mode to identify LSA plug-ins and drivers that will fail to load in LSA Protection mode. While in the audit mode, the system will generate event logs, identifying all of the plug-ins and drivers that will fail to load under LSA if LSA Protection is enabled. The messages are logged without blocking the plug-ins or drivers.

To enable the audit mode for Lsass.exe on a single computer by editing the Registry

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe.
2. Set the value of the registry key to **AuditLevel=dword:00000008**.
3. Restart the computer.

Analyze the results of event 3065 and event 3066.

- **Event 3065**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a particular driver that did not meet the security requirements for Shared Sections. However, due to the system policy that is set, the image was allowed to load.
- **Event 3066**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a particular driver that did not meet the Microsoft signing level requirements. However, due to the system policy that is set, the image was allowed to load.

Important

These operational events are not generated when a kernel debugger is attached and enabled on a system.

If a plug-in or driver contains Shared Sections, Event 3066 is logged with Event 3065. Removing the Shared Sections should prevent both the events from occurring unless the plug-in does not meet the Microsoft signing level requirements.

To enable audit mode for multiple computers in a domain, you can use the Registry Client-Side Extension for Group Policy to deploy the Lsass.exe audit-level registry value. You need to modify HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe registry key.

To create the AuditLevel value setting in a GPO

1. Open the Group Policy Management Console (GPMC).
2. Create a new Group Policy Object (GPO) that is linked at the domain level or that is linked to the organizational unit that contains your computer accounts. Or you can select

a GPO that is already deployed.

3. Right-click the GPO, and then click **Edit** to open the Group Policy Management Editor.
4. Expand **Computer Configuration**, expand **Preferences**, and then expand **Windows Settings**.
5. Right-click **Registry**, point to **New**, and then click **Registry Item**. The **New Registry Properties** dialog box appears.
6. In the **Hive** list, click **HKEY_LOCAL_MACHINE.**
7. In the **Key Path** list, browse to **SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe**.
8. In the **Value name** box, type **AuditLevel**.
9. In the **Value type** box, click to select the **REG_DWORD**.
10. In the **Value data** box, type **00000008**.
11. Click **OK**.

Note

For the GPO take effect, the GPO change must be replicated to all domain controllers in the domain.

To opt-in for additional LSA protection on multiple computers, you can use the Registry Client-Side Extension for Group Policy by modifying HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa. For steps about how to do this, see How to configure additional LSA protection of credentials in this topic.

After opting in: How to identify plug-ins and drivers loaded by the lsass.exe

You can use the event log to identify LSA plug-ins and drivers that failed to load in LSA Protection mode. When the LSA protected process is enabled, the system generates event logs that identify all of the plug-ins and drivers that failed to load under LSA.

Analyze the results of Event 3033 and Event 3063.

- **Event 3033**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a driver that did not meet the Microsoft signing level requirements.
- **Event 3063**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a driver that did not meet the security requirements for Shared Sections.

Shared Sections are typically the result of programming techniques that allow instance data to interact with other processes that use the same security context. This can create security vulnerabilities.

How to configure additional LSA protection of credentials

On devices running Windows 8.1 (with or without Secure Boot or UEFI), configuration is possible by performing the procedures described in this section. For devices running Windows RT 8.1, lsass.exe protection is always enabled, and it cannot be turned off.

- For information about changes in Secure Boot in Windows 8 and Windows 8.1, see **Secure Boot**.
- For information about UEFI in Windows 8 and Windows 8.1, see **What's Changed in Security Technologies in Windows 8.1**.

On x86-based or x64-based devices using Secure Boot and UEFI or not

On x86-based or x64-based devices that use Secure Boot and UEFI, a UEFI variable is set in the UEFI firmware when LSA protection is enabled by using the registry key. When the setting is stored in the firmware, the UEFI variable cannot be deleted or changed in the registry key. The UEFI variable must be reset.

x86-based or x64-based devices that do not support UEFI or Secure Boot are disabled, cannot store the configuration for LSA protection in the firmware, and rely solely on the presence of the registry key. In this scenario, it is possible to disable LSA protection by using remote access to the device.

You can use the following procedures to enable or disable LSA protection:

To enable LSA protection on a single computer

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa.
2. Set the value of the registry key to: "RunAsPPL"=dword:00000001.
3. Restart the computer.

To enable LSA protection using Group Policy

1. Open the Group Policy Management Console (GPMC).
2. Create a new GPO that is linked at the domain level or that is linked to the organizational unit that contains your computer accounts. Or you can select a GPO that is already deployed.
3. Right-click the GPO, and then click **Edit** to open the Group Policy Management Editor.
4. Expand **Computer Configuration**, expand **Preferences**, and then expand **Windows Settings**.
5. Right-click **Registry**, point to **New**, and then click **Registry Item**. The **New Registry Properties** dialog box appears.
6. In the **Hive** list, click **HKEY_LOCAL_MACHINE**.

7. In the **Key Path** list, browse to **SYSTEM\CurrentControlSet\Control\Lsa**.
8. In the **Value name** box, type **RunAsPPL**.
9. In the **Value type** box, click the **REG_DWORD**.
10. In the **Value data** box, type **00000001**.
11. Click **OK**.

To disable LSA protection

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa.
2. Delete the following value from the registry key: "RunAsPPL"=dword:00000001.
3. Use the Local Security Authority (LSA) Protected Process Opt-out tool to delete the UFEI variable if the device is using Secure Boot.

   For more information about the opt-out tool, see Download Local Security Authority (LSA) Protected Process Opt-out from Official Microsoft Download Center.

   For more information about managing Secure Boot, see UEFI Firmware.

   Warning
   When Secure Boot is turned off, all the Secure Boot and UEFI-related configurations are reset. You should turn off Secure Boot only when all other means to disable LSA protection have failed.

Verifying LSA protection

To discover if LSA was started in protected mode when Windows started, search for the following WinInit event in the **System** log under **Windows Logs**:

- 12: LSASS.exe was started as a protected process with level: 4

Additional resources

Credentials Protection and Management

Cached and Stored Credentials Technical Overview

This topic for the IT professional describes how credentials are formed in Windows and how the operating system manages them.

Introduction

When a user or service wants to access a computing resource, they must provide information that proves their identity. Their identity is typically in the form of their account's user name. This might be the user name that is the Security Accounts Manager (SAM) account name or the User Principal Name (UPN). But to prove their identity, they must provide secret information, which is called the *authenticator*. An authenticator can take various forms depending on the authentication protocol and method. The combination of an identity and an authenticator is called an *authentication credential*.

The process of creating, submitting, and verifying credentials is described simply as authentication, which is implemented through various authentication protocols, such as the Kerberos protocol. Authentication establishes the identity of the user, but not necessarily the user's permission to access or change a specific computing resource. That process is known as *authorization*.

Credentials are typically created or converted to a form that is required by the authentication protocols that are available on a computer. Credentials can be stored in the Local Security Authority Subsystem Service (LSASS) process memory for use by the account during a session. Credentials must also be stored on a hard disk drive in authoritative databases, such as the SAM database and in the database that is used by Active Directory Domain Services (AD DS).

For more information about storage, see Credentials storage in this topic.

Credential authenticators

The authenticator types used in the Windows operating system are as follows:

Plaintext credentials

When a user signs in to a computer running Windows and provides a user name and credentials (such as a password or PIN), the information is provided to the computer in plaintext. This plaintext password is used to authenticate the user's identity by converting it into the form that is required by the authentication protocol. Some versions of Windows also retain an encrypted copy of this password that can be unencrypted to plaintext for use with authentication methods such as Digest authentication.

Note

Windows operating systems never store any plaintext credentials in memory or on the hard disk drive. Only reversibly encrypted credentials are stored there. When later access to the plaintext forms of the credentials is required, Windows stores the passwords in encrypted form that can only be decrypted by the operating system to provide access in authorized circumstances.

These protections, however, cannot prevent a malicious user with system-level access from illicitly extracting them in the same manner that the operating system would for legitimate use.

NT hash

The NT hash of the password is calculated by using an unsalted MD4 hash algorithm. MD4 is a cryptographic one-way function that produces a mathematical representation of a password. This hashing function is designed to always produce the same result from the same password input, and to minimize collisions where two different passwords can produce the same result. This hash is always the same length and cannot be directly decrypted to reveal the plaintext password. Because the NT hash only changes when the password changes, an NT hash is valid for authentication until a user's password is changed.

Notes

To protect against brute-force attacks on the NT hashes or online systems, users who authenticate with passwords should set strong passwords or passphrases that include characters from multiple sets and are as long as the user can easily remember.

For password complexity guidelines, see the **Strong passwords** section in the Passwords Technical Overview.

LM hash

LAN Manager (LM) hashes are derived from the user password. Legacy support for LM hashes and the LAN Manager authentication protocol remains in the NTLM protocol suite. Default configurations in Windows and Microsoft security guidance have discouraged its use.

LM hashes inherently are more vulnerable to attacks because:

- LM hashes require a password to be less than 15 characters long and they contain only ASCII characters.
- LM hashes do not differentiate between uppercase and lowercase letters.

Windows logon cached password verifiers

These verifiers are not credentials because they cannot be presented to another computer for authentication, and they can only be used to locally verify a credential. They are stored in the registry on the local computer and provide credentials validation when a domain-joined computer cannot connect to AD DS during a user's logon. These "cached logons" or more specifically, cached domain account information, can be managed using the security policy setting Interactive logon: Number of previous logons to cache (in case domain controller is not available).

Credentials storage

The following sections describe where credentials are stored in Windows operating systems. Windows credentials are composed of a combination of an account name and the authenticator. These are stored and retrieved from the following locations depending on the status of the user's session, which might be active or inactive, and local or networked.

Security Accounts Manager database

The SAM database is stored as a file on the local hard disk drive, and it is the authoritative credential store for local accounts on each Windows computer. This database contains all the credentials that are local to that specific computer, including the built-in local Administrator account and any other local accounts for that computer.

The SAM database stores information on each account, including the user name and the NT password hash. By default, the SAM database does not store LM hashes on current versions of Windows. No password is ever stored in a SAM database—only the password hashes. The NT password hash is an unsalted MD4 hash of the account's password. This means that if two accounts use an identical password, they will also have an identical NT password hash.

LSASS process memory

The Local Security Authority Subsystem Service (LSASS) stores credentials in memory on behalf of users with active Windows sessions. This allows users to seamlessly access network resources, such as file shares, Exchange Server mailboxes, and SharePoint sites, without re-entering their credentials for each remote service.

LSASS can store credentials in multiple forms, including:

- Reversibly encrypted plaintext
- Kerberos tickets (TGTs, service tickets)
- NT hash
- LM hash

If the user logs on to Windows by using a smart card, LSASS will not store a plaintext password, but it will store the corresponding NT hash value for the account and the plaintext PIN for the smart card. If the account attribute is enabled for a smart card that is required for interactive logon, a random NT hash value is automatically generated for the account instead of the original password hash. The password hash that is automatically generated when the attribute is set does not change.

If a user logs on to Windows with a password that is compatible with LM hashes, this authenticator will be present in memory.

The storage of plaintext credentials in memory cannot be disabled, even if the credential providers that require them are disabled.

The stored credentials are directly associated with the LSASS logon sessions that have been started since the last restart and have not been closed. For example, LSA sessions with stored LSA credentials are created when a user does any of the following:

- Logs on to a local session or RDP session on the computer
- Runs a task by using the **RunAs** option
- Runs an active Windows service on the computer
- Runs a scheduled task or batch job
- Runs a task on the local computer by using a remote administration tool

LSA secrets on the hard disk drive

A Local Security Authority (LSA) secret is a secret piece of data that is accessible only to SYSTEM account processes. Some of these secrets are credentials that must persist after reboot, and they are stored in encrypted form on the hard disk drive. Credentials stored as LSA secrets might include:

- Account password for the computer's AD DS account
- Account passwords for Windows services that are configured on the computer
- Account passwords for configured scheduled tasks
- Account passwords for IIS application pools and websites

AD DS database (NTDS.DIT)

The Active Directory Domain Services (AD DS) database is the authoritative store of credentials for all user and computer accounts in an AD DS domain. The two types of domain controllers in AD DS that manage credentials differently are:

Writable Each writable domain controller in the domain contains a full copy of the domain's AD DS database, including account credentials for all accounts in the domain.

Read-only Read-only domain controllers (RODCs) house a partial local replica with credentials for a select subset of the accounts in the domain. By default, RODCs do not have a copy of privileged domain accounts.

The database stores a number of attributes for each account, which includes user names types and the following:

- NT hash for the current password
- NT hashes for password history (if configured)

NT hash values are also retained in AD DS for previous passwords to enforce password history during password change operations. The number of password history NT hash values retained is equal to the number of passwords configured in the password history enforcement policy.

LM hashes may also be stored in the AD DS database depending on the domain controller operating system version, configuration settings, and password change frequency.

Credential Manager store

Users may choose to save passwords in Windows by using an application or through the Credential Manager Control Panel applet. These credentials are stored on the hard disk drive and protected by using the Data Protection Application Programming Interface (DPAPI). Any program running as that user will be able to access credentials in this store.

Credential Manager can obtain its information in two ways:

Explicit creation When users enter a user name and password for a target computer or domain, that information is stored and used when the users attempt to log on to an appropriate computer. If no stored information is available and users supply a user name and password, they

can save the information. If the user decides to save the information, Credential Manager receives and stores it.

System population When the operating system attempts to connect to a new computer on the network, it supplies the current user name and password to the computer. If this is not sufficient to provide access, Credential Manager attempts to supply the necessary user name and password. All stored user names and passwords are examined, from most specific to least specific as appropriate to the resource, and the connection is attempted in the order of those user names and passwords. Because user names and passwords are read and applied in order, from most to least specific, no more than one user name and password can be stored for each individual target or domain.

Credential Manager uses the Credential Locker, formerly known as Windows Vault, for secure storage of user names and passwords.

Related resources

- Passwords Technical Overview
- Windows Logon and Authentication Technical Overview
- Credential Locker Overview

Protected Users Security Group

This topic for the IT professional describes the Active Directory security group Protected Users, and explains how it works. This group was introduced in Windows Server 2012 R2.

Members of this group are afforded additional protections against the compromise of credentials during authentication processes.

This security group is designed as part of a strategy to effectively protect and manage credentials within the enterprise. Members of this group automatically have non-configurable protections applied to their accounts. Membership in the Protected Users group is meant to be restrictive and proactively secure by default. The only method to modify these protections for an account is to remove the account from the security group.

This domain-related, global group triggers non-configurable protection on devices and host computers running Windows Server 2012 R2 and Windows 8.1, and on domain controllers in domains with a primary domain controller running Windows Server 2012 R2. This greatly reduces the memory footprint of credentials when users sign in to computers on the network from a non-compromised computer.

Depending on the account's domain functional level, members of the Protected Users group are further protected due to behavior changes in the authentication methods that are supported in Windows.

- The member of the Protected Users group cannot authenticate by using NTLM, Digest Authentication, or CredSSP. On a device running Windows 8.1, passwords are not cached,

so the device that uses any one of these Security Support Providers (SSPs) will fail to authenticate to a domain when the account is a member of the Protected User group.

- The Kerberos protocol will not use the weaker DES or RC4 encryption types in the pre-authentication process. This means that the domain must be configured to support at least the AES cipher suite.
- The user's account cannot be delegated with Kerberos constrained or unconstrained delegation. This means that former connections to other systems may fail if the user is a member of the Protected Users group.
- The default Kerberos Ticket Granting Tickets (TGTs) lifetime setting of four hours is configurable by using Authentication Policies and Silos, which can be accessed through the Active Directory Administrative Center (ADAC). This means that when four hours has passed, the user must authenticate again.

For more information, see How the Protected Users group works in this topic.

The following table specifies the properties of the Protected Users group.

Attribute	Value
Well-known SID/RID	S-1-5-21-<domain>-525
Type	Domain Global
Default container	CN=Users, DC=<domain>, DC=
Default members	None
Default member of	None
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Yes
Safe to delegate management of this group to non-service admins?	No
Default user rights	No default user rights

How the Protected Users group works

This section explains how the Protected Users group works when:

- Windows 8.1 devices are connecting to Windows Server 2012 R2 hosts.
- The account is located at the Windows Server 2012 R2 domain functional level.

When Windows 8.1 devices are connecting to Windows Server 2012 R2 hosts

When the Protected Users' group account is upgraded to the Windows Server 2012 R2 domain functional level, domain controller-based protections are automatically applied. Members of the Protected Users group who authenticate to a Windows Server 2012 R2 domain can no longer authenticate by using:

- Default credential delegation (CredSSP). Plain text credentials are not cached even when the **Allow delegating default credentials** Group Policy setting is enabled.
- Windows Digest. Plain text credentials are not cached even when Windows Digest is enabled.
- NTLM. The result of the NT one-way function, NTOWF, is not cached.
- Kerberos long-term keys. The keys from Kerberos initial TGT requests are typically cached so the authentication requests are not interrupted. For accounts in this group, Kerberos protocol verifies authentication at each request..
- Sign-in offline. A cached verifier is not created at sign-in.

Non-configurable settings to the TGTs expiration are established for every account in the Protected Users group. Normally, the domain controller sets the TGTs lifetime and renewal, based on the domain policies, **Maximum lifetime for user ticket** and **Maximum lifetime for user ticket renewal**. For the Protected Users group, 600 minutes is set for these domain policies.

After the user account is added to the Protected Users group, protection is already in place when the user signs in to the domain.

When the account is located at the Windows Server 2012 R2 domain functional level

The Protected Users group can be applied to domains that are set to a domain functional level for an operating system earlier than Windows Server 2012 R2. This allows the added security that is achieved by using the Protected Users group to be applied throughout the domain. To do this, promote the primary domain controller (PDC) to Windows Server 2012 R2, and then allow the upgraded PDC to replicate the Protected Users group to other domain controllers. When the replication completes, the PDC can be set back to any available domain functional level (if desired), and the domain controller-based protections are automatically applied.

Accounts that are members of the Protected Users group that authenticate to a Windows Server 2012 R2 domain are unable to:

- Authenticate with NTLM authentication.
- Use DES or RC4 encryption types in Kerberos pre-authentication.
- Be delegated with unconstrained or constrained delegation.
- Renew the Kerberos TGTs beyond the initial four-hour lifetime.

Warning

Accounts for services and computers should not be members of the Protected Users group. This group provides no local protection because the password or certificate is always available on the host.

Event log information

Two operational administrative logs are available to help troubleshoot events that are related to Protected Users. These new logs are located in Event Viewer and are disabled by default, and are located under **Applications and Services Logs\Microsoft\Windows\Microsoft\Authentication**.

Event ID and Log	Description
104 **ProtectedUser-Client**	Reason: The security package on the client does not contain the credentials. The error is logged in the client computer when the account is a member of the Protected Users security group. This event indicates that the security package does not cache the credentials that are needed to authenticate to the server. Displays the package name, user name, domain name, and server name.
304 **ProtectedUser-Client**	Reason: The security package does not store the Protected User's credentials. An informational event is logged in the client to indicate that the security package does not cache the user's sign-in credentials. It is expected that Digest (WDigest), Credential Delegation (CredSSP), and NTLM fail to have sign-on credentials for Protected Users. Applications can still succeed if they prompt for credentials. Displays the package name, user name, and domain name.
100 **ProtectedUserFailures-DomainController**	Reason: An NTLM sign-in failure occurs for an account that is in the Protected Users security group. An error is logged in the domain controller to indicate that NTLM authentication failed because the account was a member of the Protected Users security group. Displays the account name and device name.
104 **ProtectedUserFailures-DomainController**	Reason: DES or RC4 encryption types are used for Kerberos authentication and a sign-in failure occurs for a user in the Protected User security group. Kerberos preauthentication failed because DES and RC4 encryption types cannot be used when the account is a member of the Protected Users security group.

Event ID and Log	Description
	(AES is acceptable.)
303 **ProtectedUserSuccesses-DomainController**	Reason: A Kerberos ticket-granting-ticket (TGT) was successfully issued for a member of the Protected User group.

Deployment requirements

Requirements to provide client-side protection for members of the Protected Users group include:

- The Protected Users global security group is replicated to all domain controllers in the account domain.
- Devices and hosts are running Windows 8.1 or Windows Server 2012 R2.

Requirements to provide domain controller protection for members of the Protected Users group include:

- The domain functional level in the account domains is set to Windows Server 2012 R2.

To enable Windows Server 2012 R2 and Windows 8.1 protection for clients on domains with pre-Windows Server 2012 R2 domain functional levels, after the Protected Users group has replicated throughout the domain, a user signs in with an account that is a member of a Protected Users group.

Additional resources

- Credentials Protection and Management
- Authentication Policies and Authentication Policy Silos
- How to Configure Protected Accounts

Authentication Policies and Authentication Policy Silos

This topic for the IT professional describes authentication policy silos and the policies that can restrict accounts to those silos. It also explains how authentication policies can be used to restrict the scope of accounts.

Authentication policy silos and the accompanying policies provide a way to contain high-privilege credentials to systems that are only pertinent to selected users, computers, or services. Silos can be defined and managed in Active Directory Domain Services (AD DS) by using the Active Directory Administrative Center and the Active Directory Windows PowerShell cmdlets.

Authentication policy silos are containers to which administrators can assign user accounts, computer accounts, and service accounts. Sets of accounts can then be managed by the

authentication policies that have been applied to that container. This reduces the need for the administrator to track access to resources for individual accounts, and helps prevent malicious users from accessing other resources through credential theft.

Capabilities introduced in Windows Server 2012 R2, allow you to create authentication policy silos, which host a set of high-privilege users. You can then assign authentication policies for this container to limit where privileged accounts can be used in the domain. When accounts are in the Protected Users security group, additional controls are applied, such as the exclusive use of the Kerberos protocol.

With these capabilities, you can limit high-value account usage to high-value hosts. For example, you could create a new Forest Administrators silo that contains enterprise, schema, and domain administrators. Then you could configure the silo with an authentication policy so that password and smartcard-based authentication from systems other than domain controllers and domain administrator consoles would fail.

For information about configuring authentication policy silos and authentication policies, see How to Configure Protected Accounts.

About authentication policy silos

An authentication policy silo controls which accounts can be restricted by the silo and defines the authentication policies to apply to the members. You can create the silo based on the requirements of your organization. The silos are Active Directory objects for users, computers, and services as defined by the schema in the following table.

Active Directory schema for authentication policy silos

Display Name	Description
Authentication Policy Silo	An instance of this class defines authentication policies and related behaviors for assigned users, computers, and services.
Authentication Policy Silos	A container of this class can contain authentication policy silo objects.
Authentication Policy Silo Enforced	Specifies whether the authentication policy silo is enforced. When not enforced, the policy by default is in audit mode. Events that indicate potential successes and failures are generated, but protections are not applied to the system.
Assigned Authentication Policy Silo Backlink	This attribute is the back link for msDS-AssignedAuthNPolicySilo.
Authentication Policy Silo Members	Specifies which principals are assigned to the AuthNPolicySilo.

Display Name	Description
Authentication Policy Silo Members Backlink	This attribute is the back link for msDS-AuthNPolicySiloMembers.

Authentication policy silos can be configured by using the Active Directory Administrative Console or Windows PowerShell. For more information, see How to Configure Protected Accounts.

About authentication policies

An authentication policy defines the Kerberos protocol ticket-granting ticket (TGT) lifetime properties and authentication access control conditions for an account type. The policy is built on and controls the AD DS container known as the authentication policy silo.

Authentication policies control the following:

- The TGT lifetime for the account, which is set to be non-renewable.
- The criteria that device accounts need to meet to sign in with a password or a certificate.
- The criteria that users and devices need to meet to authenticate to services running as part of the account.

The Active Directory account type determines the caller's role as one of the following:

- **User**

  Users should always be members of the Protected Users security group, which by default rejects attempts to authentication using NTLM.

  Policies can be configured to set the TGT lifetime of a user account to a shorter value or restrict the devices to which a user account can sign in. Rich expressions can be configured in the authentication policy to control the criteria that the users and their devices need to meet to authenticate to the service.

  For more information see Protected Users Security Group.

- **Service**

  Standalone managed service accounts, group managed service accounts, or a custom account object that is derived from these two types of service accounts are used. Policies can set a device's access control conditions, which are used to restrict managed service account credentials to specific devices with an Active Directory identity. Services should never be members of the Protected Users security group because all incoming authentication will fail.

- **Computer**

  The computer account object or the custom account object that is derived from the computer account object is used. Policies can set the access control conditions that are required to allow authentication to the account based on user and device properties. Computers should never be members of the Protected Users security group because all incoming authentication will fail. By default, attempts to use NTLM authentication are rejected. A TGT lifetime should not be configured for computer accounts.

Note

It is possible to set an authentication policy on a set of accounts without associating the policy to an authentication policy silo. You can use this strategy when you have a single account to protect.

Active Directory schema for authentication policies

The policies for the Active Directory objects for users, computers, and services are defined by the schema in the following table.

Type	Display Name	Description
Policy	Authentication Policy	An instance of this class defines authentication policy behaviors for assigned principals.
Policy	Authentication Policies	A container of this class can contain authentication policy objects.
Policy	Authentication Policy Enforced	Specifies whether the authentication policy is enforced. When not enforced, the policy by default is in audit mode, and events that indicate potential successes and failures are generated, but protections are not applied to the system.
Policy	Assigned Authentication Policy Backlink	This attribute is the back link for msDS-AssignedAuthNPolicy.
Policy	Assigned Authentication Policy	Specifies which AuthNPolicy should be applied to this principal.
User	User Authentication Policy	Specifies which AuthNPolicy should be applied to users who are assigned to this silo object.
User	User Authentication Policy Backlink	This attribute is the back link for msDS-UserAuthNPolicy.
User	ms-DS-User-Allowed-To-Authenticate-To	This attribute is used to determine the set of principals allowed to authenticate to a service running under the user account.
User	ms-DS-User-Allowed-To-Authenticate-From	This attribute is used to determine the set of devices to which a user account has permission to sign in.

Type	Display Name	Description
User	User TGT Lifetime	Specifies the maximum age of a Kerberos TGT that is issued to a user (expressed in seconds). Resultant TGTs are non-renewable.
Computer	Computer Authentication Policy	Specifies which AuthNPolicy should be applied to computers that are assigned to this silo object.
Computer	Computer Authentication Policy Backlink	This attribute is the back link for msDS-ComputerAuthNPolicy.
Computer	ms-DS-Computer-Allowed-To-Authenticate-To	This attribute is used to determine the set of principals that are allowed to authenticate to a service running under the computer account.
Computer	Computer TGT Lifetime	Specifies the maximum age of a Kerberos TGT that is issued to a computer (expressed in seconds). It is not recommended to change this setting.
Service	Service Authentication Policy	Specifies which AuthNPolicy should be applied to services that are assigned to this silo object.
Service	Service Authentication Policy Backlink	This attribute is the back link for msDS-ServiceAuthNPolicy.
Service	ms-DS-Service-Allowed-To-Authenticate-To	This attribute is used to determine the set of principals that are allowed to authenticate to a service running under the service account.
Service	ms-DS-Service-Allowed-To-Authenticate-From	This attribute is used to determine the set of devices to which a service account has permission to sign in.
Service	Service TGT Lifetime	Specifies the maximum age of a Kerberos TGT that is issued to a

Type	Display Name	Description
		service (expressed in seconds).

Authentication policies can be configured for each silo by using the Active Directory Administrative Console or Windows PowerShell. For more information, see How to Configure Protected Accounts.

How it works

This section describes how authentication policy silos and authentication policies work in conjunction with the Protected Users security group and implementation of the Kerberos protocol in Windows.

- How the Kerberos protocol is used with authentication policy silos and policies
- How restricting a user sign-in works
- How restricting service ticket issuance works

Protected accounts

The Protected Users security group triggers non-configurable protection on devices and host computers running Windows Server 2012 R2 and Windows 8.1, and on domain controllers in domains with a primary domain controller running Windows Server 2012 R2. Depending on the domain functional level of the account, members of the Protected Users security group are further protected because of changes in the authentication methods that are supported in Windows.

- The member of the Protected Users security group cannot authenticate by using NTLM, Digest Authentication, or CredSSP default credential delegation. On a device running Windows 8.1 that uses any one of these Security Support Providers (SSPs), authentication to a domain will fail when the account is a member of the Protected Users security group.
- The Kerberos protocol will not use the weaker DES or RC4 encryption types in the preauthentication process. This means that the domain must be configured to support at least the AES encryption type.
- The user's account cannot be delegated with Kerberos constrained or unconstrained delegation. This means that former connections to other systems may fail if the user is a member of the Protected Users security group.
- The default Kerberos TGTs lifetime setting of four hours is configurable by using authentication policies and silos, which can be accessed through the Active Directory Administrative Center. This means that when four hours has passed, the user must authenticate again.

For more information about this security group, see How the Protected Users group works.

Silos and authentication policies

Authentication policy silos and authentication policies leverage the existing Windows authentication infrastructure. The use of the NTLM protocol is rejected, and the Kerberos protocol with newer encryption types is used. Authentication policies complement the Protected Users security group by providing a way to apply configurable restrictions to accounts, in addition to

providing restrictions for accounts for services and computers. Authentication policies are enforced during the Kerberos protocol authentication service (AS) or ticket-granting service (TGS) exchange. For more information about how Windows uses the Kerberos protocol, and what changes have been made to support authentication policy silos and authentication policies, see:

- How the Kerberos Version 5 Authentication Protocol Works
- What's New in Kerberos Authentication (Windows Server 2012)
- Changes in Kerberos Authentication (Windows Server 2008 R2 and Windows 7)

How the Kerberos protocol is used with authentication policy silos and policies

When a domain account is linked to an authentication policy silo, and the user signs in, the Security Accounts Manager adds the claim type of Authentication Policy Silo that includes the silo as the value. This claim on the account provides the access to the targeted silo.

When an authentication policy is enforced and the authentication service request for a domain account is received on the domain controller, the domain controller returns a non-renewable TGT with the configured lifetime (unless the domain TGT lifetime is shorter).

Note
The domain account must have a configured TGT lifetime and must be either directly linked to the policy or indirectly linked through the silo membership.

When an authentication policy is in audit mode and the authentication service request for a domain account is received on the domain controller, the domain controller checks if authentication is allowed for the device so that it can log a warning if there is a failure. An audited authentication policy does not alter the process, so authentication requests will not fail if they do not meet the requirements of the policy.

Note
The domain account must be either directly linked to the policy or indirectly linked through the silo membership.

When an authentication policy is enforced and the authentication service is armored, the authentication service request for a domain account is received on the domain controller, the domain controller checks if authentication is allowed for the device. If it fails, the domain controller returns an error message and logs an event. For information about Kerberos armoring, see Support for claims, compound authentication, and Kerberos armoring.

Note
The domain account must be either directly linked to the policy or indirectly linked through the silo membership.

When an authentication policy is in audit mode and a ticket-granting service request is received by the domain controller for a domain account, the domain controller checks if authentication is allowed based on the request's ticket Privilege Attribute Certificate (PAC) data, and it logs a warning message if it fails. The PAC contains various types of authorization data, including

groups that the user is a member of, rights the user has, and what policies apply to the user. This information is used to generate the user's access token. If it is an enforced authentication policy which allows authentication to a user, device, or service, the domain controller checks if authentication is allowed based on the request's ticket PAC data. If it fails, the domain controller returns an error message and logs an event.

Note

The domain account must be either directly linked or linked through silo membership to an audited authentication policy which allows authentication to a user, device or service,

You can use a single authentication policy for all members of a silo, or you can use separate policies for users, computers, and managed service accounts.

Authentication policies can be configured for each silo by using the Active Directory Administrative Console or Windows PowerShell. For more information, see How to Configure Protected Accounts.

How restricting a user sign-in works

Because these authentication policies are applied to an account, it also applies to accounts that are used by services. If you want to limit the usage of a password for a service to specific hosts, this setting is useful. For example, group managed service accounts are configured where the hosts are allowed to retrieve the password from Active Directory Domain Services. However, that password can be used from any host for initial authentication. By applying an access control condition, an additional layer of protection can be achieved by limiting the password to only the set of hosts that can retrieve the password.

When services that run as system, network service, or other local service identity connect to network services, they use the host's computer account. Computer accounts cannot be restricted. So even if the service is using a computer account that is not for a Windows host, it cannot be restricted.

Restricting user sign-in to specific hosts requires the domain controller to validate the host's identity. When using Kerberos authentication with Kerberos armoring (which is part of Dynamic Access Control), the Key Distribution Center is provided with the TGT of the host from which the user is authenticating. The content of this armored TGT is used to complete an access check to determine if the host is allowed.

When a user signs in to Windows or enters their domain credentials in a credential prompt for an application, by default, Windows sends an unarmored AS-REQ to the domain controller. If the user is sending the request from a computer that does not support armoring, such as computers running Windows 7 or Windows Vista, the request fails.

The following list describes the process:

- The domain controller in a domain running Windows Server 2012 R2 queries for the user account and determines if it is configured with an authentication policy that restricts initial authentication that requires armored requests.
- The domain controller will fail the request.

- Because armoring is required, the user can attempt to sign in by using a computer running Windows 8.1 or Windows 8, which is enabled to support Kerberos armoring to retry the sign-in process.
- Windows detects that the domain supports Kerberos armoring and sends an armored AS-REQ to retry the sign-in request.
- The domain controller performs an access check by using the configured access control conditions and the client operating system's identity information in the TGT that was used to armor the request.
- If the access check fails, the domain controller rejects the request.

Even when operating systems support Kerberos armoring, access control requirements can be applied and must be met before access is granted. Users sign in to Windows or enter their domain credentials in a credential prompt for an application. By default, Windows sends an unarmored AS-REQ to the domain controller. If the user is sending the request from a computer that supports armoring, such as Windows 8.1 or Windows 8, authentication policies are evaluated as follows:

1. The domain controller in a domain running Windows Server 2012 R2 queries for the user account and determines if it is configured with an authentication policy that restricts initial authentication that requires armored requests.
2. The domain controller performs an access check by using the configured access control conditions and the system's identity information in the TGT that is used to armor the request. The access check succeeds.

   Note

   If legacy workgroup restrictions are configured, those also need to be met.
3. The domain controller replies with an armored reply (AS-REP), and the authentication continues.

How restricting service ticket issuance works

When an account is not allowed and a user who has a TGT attempts to connect to the service (such as by opening an application that requires authentication to a service that is identified by the service's service principal name (SPN), the following sequence occurs:

1. In an attempt to connect to SPN1 from SPN, Windows sends a TGS-REQ to the domain controller that is requesting a service ticket to SPN1.
2. The domain controller in a domain running Windows Server 2012 R2 looks up SPN1 to find the Active Directory Domain Services account for the service and determines that the account is configured with an authentication policy that restricts service ticket issuance.
3. The domain controller performs an access check by using the configured access control conditions and the user's identity information in the TGT. The access check fails.
4. The domain controller rejects the request.

When an account is allowed because the account meets the access control conditions that are set by the authentication policy, and a user who has a TGT attempts to connect to the service

(such as by opening an application that requires authentication to a service that is identified by the service's SPN), the following sequence occurs:

1. In an attempt to connect to SPN1, Windows sends a TGS-REQ to the domain controller that is requesting a service ticket to SPN1.
2. The domain controller in a domain running Windows Server 2012 R2 looks up SPN1 to find the Active Directory Domain Services account for the service and determines that the account is configured with an authentication policy that restricts service ticket issuance.
3. The domain controller performs an access check by using the configured access control conditions and the user's identity information in the TGT. The access check succeeds.
4. The domain controller replies to the request with a ticket-granting service reply (TGS-REP).

Associated error and informational event messages

The following table describes the events that are associated with Protected Users security group and the authentication policies that are applied to authentication policy silos.

The events are recorded in the Applications and Services Logs at **Microsoft\Windows\Authentication**.

For troubleshooting steps that use these events, see Troubleshoot Authentication Policies and Troubleshoot events related to Protected Users.

Event ID and Log	Description
101 **AuthenticationPolicyFailures-DomainController**	Reason: An NTLM sign-in failure occurs because the authentication policy is configured. An event is logged in the domain controller to indicate that NTLM authentication failed because access control restrictions are required, and those restrictions cannot be applied to NTLM. Displays the account, device, policy, and silo names.
105 **AuthenticationPolicyFailures-DomainController**	Reason: A Kerberos restriction failure occurs because the authentication from a particular device was not permitted. An event is logged in the domain controller to indicate that a Kerberos TGT was denied because the device did not meet the enforced access control restrictions. Displays the account, device, policy, silo names, and TGT lifetime.

Event ID and Log	Description
305 **AuthenticationPolicyFailures-DomainController**	Reason: A potential Kerberos restriction failure might occur because the authentication from a particular device was not permitted. In audit mode, an informational event is logged in the domain controller to determine if a Kerberos TGT will be denied because the device did not meet the access control restrictions. Displays the account, device, policy, silo names, and TGT lifetime.
106 **AuthenticationPolicyFailures-DomainController**	Reason: A Kerberos restriction failure occurs because the user or device was not allowed to authenticate to the server. An event is logged in the domain controller to indicate that a Kerberos service ticket was denied because the user, device, or both do not meet the enforced access control restrictions. Displays the device, policy, and silo names.
306 **AuthenticationPolicyFailures-DomainController**	Reason: A Kerberos restriction failure might occur because the user or device was not allowed to authenticate to the server. In audit mode, an informational event is logged on the domain controller to indicate that a Kerberos service ticket will be denied because the user, device, or both do not meet the access control restrictions. Displays the device, policy, and silo names.

See also

How to Configure Protected Accounts

Credentials Protection and Management

Protected Users Security Group

Encrypted Hard Drive

The Encrypted Hard Drive feature in Windows 8 and Windows Server 2012 uses the rapid encryption that is provided by BitLocker Drive Encryption to enhance data security and management. By offloading the cryptographic operations to hardware, Encrypted Hard Drives increase BitLocker performance and reduce CPU usage and power consumption. Because Encrypted Hard Drives encrypt data quickly, enterprise clients can expand BitLocker deployment with minimal impact on productivity.

Encrypted Hard Drives are a new class of hard drives that are self-encrypting at a hardware level and allow for full disk hardware encryption. Windows 8 and Windows Server 2012 support installing to these devices without additional modification.

Some of the benefits of Encrypted Hard Drives include:

- **Better performance**: Encryption hardware, integrated into the drive controller, allows the drive to operate at full data rate with no performance degradation.
- **Strong security based in hardware**: Encryption is always "on" and the keys for encryption never leave the hard drive. User authentication is performed by the drive before it will unlock, independently of the operating system
- **Ease of use**: Encryption is transparent to the user because it is on by default. There is no user interaction needed to enable encryption. Encrypted Hard Drives are easily erased using on-board encryption key; there is no need to re-encrypt data on the drive.
- **Lower cost of ownership**: There is no need for new infrastructure to manage encryption keys, since BitLocker leverages your Active Directory Domain Services infrastructure to store recovery information. Your computer operates more efficiently because processor cycles do not need to be used for the encryption process.

Windows 8 and Windows Server 2012 support Encrypted Hard Drives natively in the operating system through the following mechanisms:

- **Identification**: The operating system can identify that the drive is an Encrypted Hard Drive device type
- **Activation**: The operating system disk management utility can activate, create and map volumes to ranges/bands as appropriate
- **Configuration**: The operating system can create and map volumes to ranges/bands as appropriate
- **API**: Windows 8 and Windows Server 2012 provides API support for applications to manage Encrypted Hard Drives independently of BitLocker Drive Encryption (BDE)
- **BitLocker support**: Integration with the BitLocker Control Panel provides a seamless BitLocker end user experience.

Warning

Self-Encrypting Hard Drives and Encrypted Hard Drives for Windows are not the same type of device. Encrypted Hard Drives for Windows require compliance for specific TCG protocols as well as IEEE 1667 compliance; Self-Encrypting Hard Drives do not have

these requirements. It is important to confirm the device type is an Encrypted Hard Drive for Windows when planning for deployment.

System Requirements

To use an Encrypted Hard Drive on Windows 8 or Windows Server 2012., the following system requirements apply:

For Encrypted Hard Drives used as **data drives**:

- The drive must be in an uninitialized state.
- The drive must be in a security inactive state.

For Encrypted Hard Drives used as **startup drives**:

- The drive must be in an uninitialized state.
- The drive must be in a security inactive state.
- The computer must be UEFI 2.3.1 based and have the EFI_STORAGE_SECURITY_COMMAND_PROTOCOL defined. (This protocol is used to allow programs running in the EFI boot services environment to send security protocol commands to the drive).
- The computer must have the Compatibility Support Module (CSM) disabled in UEFI.
- The computer must always boot natively from UEFI.

Warning
All Encrypted Hard Drives must be attached to non-RAID controllers to function properly in Windows 8 or Windows Server 2012.

Technical overview

Rapid encryption in BitLocker directly addresses the security needs of enterprises while offering significantly improved performance. In earlier versions of Windows, BitLocker required a two-step process to complete read/write requests, but in Windows Server 2012, Encrypted Hard Drives offload the cryptographic operations to the drive controller for much greater efficiency. When Windows Server 2012 initializes an Encrypted Hard Drive, it activates the security mode. This activation lets the drive controller generate a media key for every volume that the host computer creates. This media key, which is never exposed outside the disk, is used to rapidly encrypt or decrypt every byte of data that is sent or received from the disk.

Configuring Encrypted Hard Drives as Startup drives

Configuration of Encrypted Hard Drives as startup drives is done using the same methods as standard hard drives. These methods include:

- **Deploy from media**: This deployment method involves installing Windows 8 or Windows Server 2012. from DVD media. Configuration of Encrypted Hard Drives happens automatically through the installation process.
- **Deploy from network**: This deployment method involves booting a Windows PE environment and using imaging tools to apply a Windows image from a network share. Using this method, the Enhanced Storage optional component needs to be included in the Windows PE image. You can enable this component using Server Manager, Windows PowerShell or the DISM command line tool. If this component is not present, configuration of Encrypted Hard Drives will not work.
- **Deploy from server**: This deployment method involves PXE booting a client with Encrypted Hard Drives present. Configuration of Encrypted Hard Drives happens automatically in this environment when the Enhanced Storage component is added to the PXE boot image. During deployment, the `TCGSecurityActivationDisabled` setting in unattend.xml controls the encryption behavior of Encrypted Hard Drives. For further information, see the Windows Assessment and Deployment Kit (ADK)and related release notes.
- **Disk Duplication**: This deployment method involves use of a previously configured Windows 8 or Windows Server 2012. image and disk duplication tools to apply a Windows image to an Encrypted Hard Drive. Disks must be partitioned using Windows 8 or Windows Server 2012.setup tools for this configuration to work. Images made using disk duplicators will not work.

Encrypted Hard Drive Architecture

Encrypted Hard Drives utilize two encryption keys on the device to control the locking and unlocking of data on the drive. These are the Data Encryption Key (DEK) and the Authentication Key (AK).

The Data Encryption Key is the key used to encrypt all of the data on the drive. The drive generates the DEK and it never leaves the device. It is stored in an encrypted format at a random location on the drive. If the DEK is changed or erased, data encrypted using the DEK is irrecoverable.

The Authentication Key is the key used to unlock data on the drive. A hash of the key is stored on drive and requires confirmation to decrypt the DEK.

When a computer with an Encrypted Hard Drive is in a powered off state, the drive locks automatically. As a computer powers on, the device remains in a locked state and is only unlocked after the Authentication Key decrypts the Data Encryption Key. Once the Authentication Key decrypts the Data Encryption Key, read-write operations can take place on the device.

When writing data to the drive, it passes through an encryption engine before the write operation completes. Likewise, reading data from the drive requires the encryption engine to decrypt the data before passing that data back to the user. In the event that the DEK needs to be changed or erased, the data on the drive does not need to be re-encrypted. A new Authentication Key needs to be created and it will re-encrypt the DEK. Once completed, the DEK can now be unlocked using the new AK and read-writes to the volume can continue.

Re-configuring Encrypted Hard Drives

Many Encrypted Hard Drive devices come pre-configured for use. If reconfiguration of the drive is required, use the following procedure after removing all available volumes and reverting the drive to an uninitialized state:

1. Open Disk Management (diskmgmt.msc)
2. Initialize the disk and select the appropriate partition style (MBR or GPT)
3. Create one or more volumes on the disk.
4. Use the BitLocker setup wizard to enable BitLocker on the volume.

Exchange ActiveSync Policy Engine Overview

This topic for the IT professional describes the Exchange ActiveSync policy engine and lists resources for using it.

Did you mean…

- Exchange ActiveSync
- Security Policy Settings

Feature description

The Exchange ActiveSync (EAS) policy engine was introduced in Windows Server 2012, Windows 8, and Windows RT to enable apps to apply EAS policies on desktops, laptops, and tablets to protect data that is synchronized from the cloud, such as data from an Exchange Server. It supports a core set of Windows security primitives.

Practical applications

The EAS policy engine contains a set of WinRT APIs that enable Windows Store apps to manage security primitives on devices. The policies supported by the EAS policy engine include password requirements, inactivity timers, sign-in methods, and disk encryption status. The policy engine enables you to check device status and if the status conforms to your policies. Windows Store apps can leverage the EAS policy engine APIs to verify and enforce these policies.

The Exchange ActiveSync (EAS) protocol is an XML-based protocol designed to synchronize e-mail, contacts, calendar, tasks, notes, and policies between Exchange Server and a client device. The EAS policy engine can enforce a subset of the policies defined in the EAS protocol on devices running any of the supported versions of the Windows operating system (listed in the **Applies To** list at the beginning of this topic).

How it works

Supported devices

Devices, including servers, desktops, laptops and tablets, that are running the supported versions of Windows and are capable of installing a mail app from the Windows Store can enforce EAS policies.

Available device information

A mail app that is using the EAS policy engine also has the ability to read device information and report it to the Exchange Server. Following is a list of device information that is available:

- A unique device identifier called Device ID or ID
- Name of the computer
- Operating system running on the device
- System manufacturer
- System product name
- System SKU

Policies supported by a mail app and EAS policy engine

To synchronize data from an Exchange Server, the mail app first applies security policies on the client device. A mail app can apply a core set of these policies by using WinRT APIs in the EAS policy engine.

The EAS policy engine has the ability to check for policies applied on a device and check if the accounts on that device conform to these policies. It can also enforce policies related to sign in methods, password, and device inactivity.

The EAS policy engine does not support all of the policies specified by the EAS protocol in MS-ASPROV. Particularly, policies that pertain to storage card access, mail retention, and S/MIME are not supported. For more information see, [MS-ASPROV]: Exchange ActiveSync: Provisioning Protocol in the MSDN Library. Following is a list of all the supported policies.

EAS Policy Name	Description	Correlation to Group Policies
AllowSimpleDevicePassword	Specifies whether the user is allowed to use convenience sign-in methods like pin, picture	There are three sets of Group Policies that control pin, picture password, and biometrics. These policies must be set to allow non-administrator accounts to become compliant without requiring an administrator action. **Pin** Policy setting: **Turn on PIN sign-in** Location in the Local Security Policy snap-in: Computer Configuration/Administrative Templates/System/Logon/

EAS Policy Name	Description	Correlation to Group Policies
	password, or biometrics.	**Picture** Policy setting: **Turn off picture password sign-in** Location in the Local Security Policy snap-in: Computer Configuration/Administrative Templates/System/Logon/ **Biometrics** Policy settings: **Allow the use of biometrics** **Allow users to log on using biometrics** **Note** Recommendation is to set to ensure that all non-administrator users can be compliant. **Allow domain users to log on using biometrics** Location in the Local Security Policy snap-in: Computer Configuration/Administrative Templates/Windows Components/Biometrics/ **Note** For client devices running the supported versions of Windows, if the Pin or Picture Group Policy setting is applied to bring non-administrator accounts into compliance, it is necessary to set the registry key that corresponds to DisallowConvenienceLogon, which is HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\System\DisallowConvenienceLogon.
MaxInactivityTimeDeviceLock	Specifies the length of time a device can go without user input before it is locked.	Policy setting: **Interactive logon: Machine inactivity limit** Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options/
MaxDevicePasswordFailedAttempts	Specifies how many	Policy setting: **Interactive logon: Machine account lockout threshold**

EAS Policy Name	Description	Correlation to Group Policies
	times an incorrect password can be entered before the device is rebooted. The device may be put into lock mode, which requires the recovery key to unlock the device in the presence of disk encryption.	Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options/
MinDevicePasswordComplexCharacters	Specifies the minimum number of complex characters that are required for a password.	Policy setting: **Password must meet complexity requirements** Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy/
MinDevicePasswordLength	Specifies the minimum	Policy setting: **Minimum password length** Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security

EAS Policy Name	Description	Correlation to Group Policies
	password length.	Settings/Account Policies/Password Policy/
DevicePasswordExpiration	Specifies the length of time after which a user password must be changed.	Policy setting: **Maximum password age** Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy/
DevicePasswordHistory	Specifies the number of previous passwords that cannot be reused.	Policy setting: **Enforce password history** Location in the Local Security Policy snap-in: Computer Configuration/Windows Settings/Security Settings/Account Policies/Password Policy/
RequireDeviceEncryption	Specifies whether device encryption is required. If set to True, the device must be able to support and implement encryption to become compliant. **Note** E	There is no Local Security Policy or Group Policy Administrative Template that corresponds to this EAS policy. An explicit action is required to encrypt a device if this policy is required.

EAS Policy Name	Description	Correlation to Group Policies
	nc ry pti o n ca n n ot b e e n a bl e d by th e E A S p oli cy e n gi n e, a n d a n ex pli cit us er	

EAS Policy Name	Description	Correlation to Group Policies
	action is required to turn on encryption, if it is not already running	

For more information about each policy, see Use Exchange ActiveSync Policies for Device Management.

About password policies and accounts

Password policies help prevent a malicious user from accessing content on the device by requiring a password. This is also true for computers or devices that have one or more administrator accounts. Because administrators can potentially access data for other accounts, is it necessary that all administrator accounts adhere to password policies, even if EAS policies are not applied.

If password policies are applied, all administrator and all control user accounts are required to conform to password requirements at the next sign in or unlock action. In addition, if an administrator account has a blank password, a compliant password must be applied before the computer or device can be compliant to synchronize with Exchange data.

EAS protocol defines DevicePasswordEnabled, which is not directly supported in the operating system. If an Exchange Server sets DevicePasswordEnabled, apps that apply these policies are expected to set some of the underlying password policies with default values. These policies include length, complexity, history, expiration, and failed attempts.

MaxDevicePasswordFailedAttempts is set with a default value of 4. In the presence of BitLocker or Device Encryption, incorrect password attempts will result in a restart followed by device lock. If the disk is not encrypted, the device will restart to slow down casual brute-force attacks.

Administrator or user accounts that are disabled are set to change the password at the next sign in. But because they are disabled, they are considered to be compliant.

Password history and expiration apply only at the time a local user account password is evaluated or changed. They are only required to be locally enforced. These policies are not enforced for domain or Microsoft accounts. Microsoft accounts and Active Directory domain accounts have varying policies that are enforced on the server; therefore, they will not be bound by the policies from an EAS policy.

Password length and complexity supported by account types

Password length and complexity policies are evaluated and applied on different account types differently.

Local accounts

Current local account and all administrator accounts are required to have a password, if the EAS policies set minimum password length or complexity, or both. Failing to meet this requirement results in non-compliance. When password length and complexity rules are applied, all the control user and administrator accounts are marked to change the password at the next sign in to ensure complexity requirements are met.

Local accounts can support the full password length policy, but they can only support three character sets, not the full four that EAS protocol can specify. If an EAS policy is set to require four character sets, all local accounts will become non-compliant. This is because the Windows operating system does not explicitly support choosing the number of complex characters in a password; rather, it requires that passwords meet a certain complexity level. This complexity translates to three complex characters. Therefore, an EAS policy that requires

MinDevicePasswordComplexCharacters greater than 3 cannot be supported by Windows accounts.

Local accounts default to a minimum password length and complexity policy ratio of 6 and 3 when any password policy is set. Complexity requirements are enforced when passwords are changed or created. All security threats over the network on accounts with a blank password are mitigated. However, when a user sets a password for a local account, the account is immediately vulnerable to password guessing or other password attacks over the network. So, to mitigate threats through network access, minimum requirements are set for local accounts, which provides a reasonable level of security.

Domain accounts

Domain accounts are not evaluated locally for password policies that are set by EAS because it is assumed that the EAS policies and the domain account policies belong to the same account authority. These policies include complexity, length, expiration, and history settings.

Microsoft accounts

Note

Microsoft accounts were formerly known as Windows Live ID accounts.

Much like a domain account, a Microsoft account is governed by a policy authority that is not related to a local device. Properties including password complexity, length, expiration, and history are part of the Microsoft account.

Microsoft accounts enforce a minimum password length of 8 characters and 2 character sets in a password. Therefore, Microsoft accounts can comply if the MinDevicePasswordLength policy is set at less than or equal to 8 characters and the MinDevicePasswordComplexCharacters policy is set at less than or equal to 2.

A password can still conform to the EAS policy rules, but nothing can prevent a user from creating a less-complex password for the Microsoft account. Non-compliance results if EAS policies are stricter than those that Microsoft accounts can enforce.

Expiration and history are not evaluated locally for Microsoft accounts.

Policy application on administrator and standard user accounts

EAS policies are applied to all administrator accounts, regardless of whether they have an app configured to use EAS policies.

EAS policies are also applied to any standard (non-administrator) account that has an app configured to use EAS policies. These accounts are called control user accounts. The EAS policy, MaxInactivityTimeDeviceLock is the exception because it is not applied to accounts, but rather to the device.

Policies specified by different sources

Given a set of policies that are enforced by Exchange ActiveSync, Group Policies, a Microsoft account, or local policies, the Windows operating system always enforces the strictest policy out of the set of governing policies.

Multiuser support

Windows provides for multiple users on a single device. Windows Live Mail also permits multiple EAS accounts for each user. When there are multiple EAS accounts with policies set on a device running any supported version of Windows, the policies are merged into the most restrictive resultant set.

EAS policies do not apply to all users on a device running Windows. Windows restricts standard user accounts in their ability to access data in other user profiles or in privileged locations. For this reason, EAS policies do not apply uniformly to standard users. The policies apply only to standard users who have a configured Exchange account that requires an EAS policy.

Accounts for users with Administrator rights always have the EAS policies applied.

Windows only provides a mechanism to apply a single instance of the EAS policies. Any account that is subject to an EAS policy is controlled by the strictest policy applied to the device.

Policy reset

To prevent an app from reducing security policies and posing a risk to the device, a policy cannot be reduced, even if the policy no longer exists on the server. A user must take action to reset policies in the event of policy relaxation, policy removal, account removal, or an app removal.

Policies can be reset by using the Control Panel. Click **User Accounts and Family Safety**, click **User Accounts**, and click **Reset Security Policies**. Users can also use **Reset Security Policies** to reset an EAS policy that causes an email delivery failure.

Note

The option to reset security policies is present only if policies are applied by the EAS policy engine.

EAS policies and provisioning refresh

Exchange servers can force users to provision devices and reapply policies after a certain period of time. This ensures that if the device is no longer compliant with EAS policies, the policies are reapplied or the device is deemed non-compliant.

The Windows Mail client does not enforce the provisioning refresh, so it is the responsibility of the EAS administrator to make sure that if a policy change occurs, the provisioning refresh is triggered within a given time frame.

A provisioning refresh can be controlled by setting the **Refresh Interval** policy in the Exchange ActiveSync Mailbox Policy settings. For more information about setting the refresh interval, see View or Configure Exchange ActiveSync Mailbox Policy Properties.

Device lock

The supported Windows operating systems provide data protection through BitLocker Drive Encryption. When combined with password policies and the MaxDevicePasswordFailedAttempts policy, Windows Live Mail provides a robust data protection scheme to limit the risks of data loss due to device loss or theft.

Although many mobile devices support a complete removal of device content when a remote instruction is received or when the MaxDevicePasswordFailedAttempts threshold is reached by a user, the supported Windows operating systems do not.

The device lock feature in the supported Windows operating systems enables devices that are encrypted with BitLocker to cryptographically lock all encrypted volumes and reboot the device into the recovery console. A lost or stolen device is not readable unless the device recovery key is available to the possessor of the device.

The device lock feature provides protection for lost or stolen devices and provides a means for legitimate users who accidentally enter the device lock state to recover their device and continue using it.

Autologon behavior

Implementation of EAS policies prevents users from using the Autologon feature. Autologon permits the credentials of the signed on account to be encrypted and stored in the registry so when the computer is restarted the user's account is automatically signed in using those stored credentials. Autologon is useful when administrators are required to sign in to a computer multiple times during configuration, or when a user has secure possession of their personal computer and desires an easier ability to sign in.

When EAS policies are implemented, autologon does not function by default and the user who is attempting to sign in must enter their account's credentials. If autologon behavior is desired, the EAS policies must be disabled.

Warning
Disabling EAS policies will reduce security effectiveness.

For more information about this downloadable tool, see, Autologon.

New and changed functionality

The EAS policy engine was introduced in Windows Server 2012 and Windows 8.

In Windows Server 2012 and Windows 8, Biometrics sign-in methods do not count towards the limits set in the MaxDevicePasswordFailedAttempts policy. In Windows Server 2012 R2 and

Windows 8.1, if the device is encrypted using BitLocker or any other disk encryption software, Biometrics sign-in methods are not disabled when the limit is exceeded if the DisallowConvenienceLogon policy is set.

Software requirements

The EAS policy engine and its policy implementation is supported only on versions of the Windows operating system designated in the **Applies To** list at the beginning of this topic.

Additional resources

The following table provides additional and related information about the Exchange ActiveSync policy engine, including policies and APIs.

Content type	References
Product evaluation	This topic
Planning	None
Deployment	None
Operations	Use Exchange ActiveSync Policies for Device Management
Troubleshooting	None
Security	None
Tools and settings	Security policy settings reference: Password Policy
Community resources	None
Related technologies	Managing Exchange ActiveSync: Exchange 2010 Help Exchange ActiveSync FAQ [MS-ASPROV]: Exchange ActiveSync: Provisioning Protocol

Use Exchange ActiveSync Policies for Device Management

This topic for the IT professional and mail app developer describes the policies available to use the Exchange ActiveSync (EAS) policy engine to comply with security policies, and it describes the details of each applicable policy.

The EAS policy engine was introduced in Windows Server 2012, Windows 8, and Windows RT to enable Windows store apps to apply EAS policies on desktops, laptops, and tablets to protect data that is synchronized from the cloud, such as data from an Exchange Server. It supports a core set of Windows security primitives.

For more information about the EAS policy engine and how it works, see Exchange ActiveSync Policy Engine Overview. This topic contains:

- Addressing policy compliance
- Supported EAS policy engine policies details
  - AllowSimpleDevicePassword
  - DisallowConvenienceLogon
  - MaxInactivityTimeDeviceLock
  - MaxDevicePasswordFailedAttempts
  - MinDevicePasswordComplexCharacters
  - MinDevicePasswordLength
  - DevicePasswordExpiration
  - DevicePasswordHistory
  - RequireDeviceEncryption

Addressing policy compliance

Most of the policies listed in the previous section can be applied by the user if the user is an administrator. However, there are certain cases where the user must take action first. These are:

- When the server requires encryption
  - The Windows client operating system supports BitLocker Drive Encryption, but it is not enabled. The user must manually enable BitLocker.
  - The Windows client operating system runs on a Windows RT device that supports device encryption, but it is suspended. The user must restart the device.
  - The Windows client operating system runs on a Windows RT device that supports device encryption, but it isn't enabled. The user must sign into Windows with a Microsoft account.
- When an administrator on a computer does not have a strong password. All administrator accounts must have a strong password.
- When the user's account does not have a strong password. Users must set strong passwords.

A device that is shared between an administrator and a standard user, when the user is configuring Exchange Server, requires that user and administrator are compliant with the password policy.

Microsoft accounts have a password length and complexity of ratio of 8 and 2, respectively. This means that connected accounts cannot be enforced with a length/complexity of greater than 8 and 2. Users who are blocked by an Exchange Server that has an EAS policy complexity ratio greater than what is supported by the connected accounts should create a local account with a higher complexity and ensure that their connected accounts are non-administrator accounts.

Users with non-administrator accounts cannot apply policies that require a state change to the computer or that impact other accounts. If a user with a non-administrator account attempts to apply EAS policies, the application might fail. To mitigate this problem, the user can take the following steps:

1. Change the account from a standard user to an administrator.
2. Use an administrator account to accept and apply policies. When the policies are applied, all the administrator and control user accounts can be made compliant without further administrator action. If more stringent policies are needed in the future, the administrator might have to accept the new policies before non-administrator accounts will become compliant.
3. An administrator can apply local policies and Group Policies equal to or stricter than the EAS policies. This ensures that all non-administrator accounts can be compliant because they conform to policy requirements that are equal to or stricter than what the EAS policy engine expects.

Supported EAS policy engine policies details

A policy that is applied using the EAS policy engine prompts the user for consent before a change is made to the system. Policy changes that require a change in the computer's state or impact other accounts require administrator consent. Policies that are applied using this mechanism are always more stringent.

Note
Windows Store apps (for example, the Mail app) use the WinRT APIs that are provided in the EasClientSecurityPolicy class to check and apply security policies. The following list shows the recommended mapping between this policy and the API properties supported by the operating system. For more information, see EasClientSecurityPolicy class (Windows).

AllowSimpleDevicePassword

Originally, the AllowSimpleDevicePassword policy was used to allow or deny setting PINs and passwords that consisted of repeating or sequential numbers or characters. Windows 8 introduced DisallowConvenienceLogon to control sign-in methods. Because different types of sign-in methods exist for the user's convenience, AllowSimplePassword is no longer directly

exposed as a policy setting in the operating system. AllowSimpleDevicePassword is mapped to DisallowConvenienceLogon by the Mail app.

DisallowConvenienceLogon

The DisallowConvenienceLogon policy can be used to disable convenience sign-in methods like PINs, pictures, and biometrics. It prevents users from using sign-in methods that are less secure in certain configurations.

PIN and picture sign-in methods allow users to attempt five incorrect sign-ins. These incorrect sign-in attempts also contribute toward the MaxDevicePasswordFailedAttempts policy. When a certain threshold is reached, the computer will restart, and users are possibly locked out of their devices.

In the supported versions of Windows (designated in the **Applies To** list at the top of this page), biometric sign-in methods, do not count toward the MaxDevicePasswordFailedAttempts policy.

In Windows Server 2012 R2 and Windows 8.1, if the device is encrypted with BitLocker or any other disk encryption software, biometrics is not disabled if the DisallowConvenienceLogon policy is set.

The DisallowConvenienceLogon policy applies to all administrator and control user accounts. It is recommended to allow the use of simple passwords and to set adequate password lengths and other related policies.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, local, or domain policy applies.
Domain accounts	The strictest policy of either EAS, local, or domain policy applies.
Connected accounts (for example, Microsoft account)	The strictest policy of either EAS, local, or domain policy applies.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: DisallowConvenienceLogon
- Relationship: Inverse

MaxInactivityTimeDeviceLock

The MaxInactivityTimeDeviceLock policy ensures that unattended computers and devices are locked after a certain amount of time. However, locking the device alone is not enough protection from a security perspective. Therefore, when this policy is applied, it also forces the system to require a password when the computer is activated.

A user can configure an inactivity time that is shorter than the time applied in the policy. The shorter of the two will become the effective inactivity timer.

MaxInactivityTimeDeviceLock is a device global timer that applies to all administrator and all user accounts. This is unique relative to other policies because they apply only to administrative and controlled user accounts.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, local, or domain policy applies.
Domain accounts	The strictest policy of either EAS, local, or domain policy applies.
Connected accounts (for example, Microsoft account)	The strictest policy of either EAS, local, or domain policy applies.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: MaxInactivityTimeLock
- Relationship: One-to-One

MaxDevicePasswordFailedAttempts

The MaxDevicePasswordFailedAttempts policy ensures that there is a limit to incorrect sign-in attempts that are made to a local account. When specified, the policy causes the computer or device to restart after the specified threshold is reached. If the system is protected with BitLocker or Device Encryption, it will be put into device lock mode. Restarting a device prevents lightweight attacks such as password guessing. However, it cannot prevent against motivated attacks and data theft if the disk is not encrypted. When coupled with disk encryption, there is a better stance against data theft if a device is lost or stolen.

Failed password attempt counts persist across restarts, and they are maintained on a per user basis to ensure that the computer is put into device lock mode, even if incorrect password attempts are made across restarts.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, local, or domain policy applies.
Domain accounts	The strictest policy of either EAS, local, or domain policy applies.
Connected accounts (for example, Microsoft	The strictest policy of either EAS, local, or

Account types	Evaluation scope
account)	domain policy applies.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: MaxPasswordFailedAttempts
- Relationship: One-to-One

MinDevicePasswordComplexCharacters

The MinDevicePasswordComplexCharacters policy defines the number of minimum character types that should be present in an account password to meet security requirements. The character types are defined as:

- Lower-case alphabetical
- Upper-case alphabetical
- Numbers
- Non-alphanumeric

If MinDevicePasswordComplexCharacters is greater than 1, DisallowConvenienceLogon is ignored. The password cannot contain the user's account name or parts of the user's full name that exceed two consecutive characters. It also needs to be at least six characters in length.

MinDevicePasswordComplexCharacters applies to all administrator and control user accounts.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, local, or domain policy applies.
Domain accounts	Not evaluated
Connected accounts (for example, Microsoft account)	EAS Policy set at less than or equal to 2 is always compliant and a value greater than 2 results in non-compliance.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: MinPasswordComplexCharacters
- Relationship: One-to-One

MinDevicePasswordLength

The MinDevicePasswordLength policy defines the minimum length of a password.

MinDevicePasswordLength applies to all administrator and control user accounts.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, and local policy applies.
Domain accounts	Not evaluated
Connected accounts (for example, Microsoft account)	EAS Policy less than or equal to 8 is always compliant and a value greater than 8 results in non-compliance.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: MinPasswordLength
- Relationship: One-to-One

DevicePasswordExpiration

The DevicePasswordExpiration policy determines the period of time (in days) that a password can be used before the system requires the user to change it.

DevicePasswordExpiration applies to all administrator and control user accounts and only governs local accounts.

The following table describes policy evaluation for account types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, and local policy applies.
Domain accounts	Not evaluated
Connected accounts (for example, Microsoft account)	Not evaluated

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: PasswordExpiration
- Relationship: One-to-One

DevicePasswordHistory

The DevicePasswordHistory policy allows administrators to enhance security by ensuring that old passwords are not reused. It is recommended that anyone setting this policy also sets the **Minimum password age** policy (by using Group Policy) to prevent users from rolling over passwords to get back to a favorite password.

DevicePasswordHistory applies to all administrator and control user accounts and only governs local accounts.

The following table describes policy evaluation for accounts types.

Account types	Evaluation scope
Local accounts	The strictest policy of either EAS, and local policy applies.
Domain accounts	Not evaluated
Connected accounts (for example, Microsoft account)	Not evaluated

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: PasswordHistory
- Relationship: One-to-One

RequireDeviceEncryption

The RequireDeviceEncryption policy requires that all fixed data partitions that can contain user data must be encrypted. This includes the operating system partition and any other data partitions that are created by the user.

The EAS policy engine in Windows only checks for compliance related to this policy. If a computer or device is found to not be in compliance, a user action is required to encrypt the hard disk drive before the computer or device can be compliant so that data can be synchronized.

On Windows RT devices, encryption starts as soon as an account is connected so the recovery key is automatically bound to the corresponding Microsoft account.

The following table describes policy evaluation for accounts types.

Account types	Evaluation scope
Local accounts	All non-conflicting EAS policies apply.
Domain accounts	All non-conflicting EAS policies apply.
Connected accounts (for example, Microsoft account)	All non-conflicting EAS policies apply.

EAS policy mapping at the WinRT API layer

- WinRT API Policy Object: RequireEncryption
- Relationship: One-to-One

Related resources

For more information about the EAS policy engine and how it works, see Exchange ActiveSync Policy Engine Overview.

Group Managed Service Accounts Overview

This topic for the IT professional introduces the group Managed Service Account by describing practical applications, changes in Microsoft's implementation, hardware and software requirements, plus additional resources for Windows Server 2012.

Did you mean…

- Service accounts
- Managed Service Accounts (standalone)

Feature description

Standalone Managed Service Accounts, which were introduced in Windows Server 2008 R2 and Windows 7, are managed domain accounts that provide automatic password management and simplified SPN management, including delegation of management to other administrators.

The group Managed Service Account provides the same functionality within the domain but also extends that functionality over multiple servers. When connecting to a service hosted on a server farm, such as Network Load Balance, the authentication protocols supporting mutual authentication require that all instances of the services use the same principal. When group Managed Service Account are used as service principals, the Windows operating system manages the password for the account instead of relying on the administrator to manage the password.

The Microsoft Key Distribution Service (kdssvc.dll) provides the mechanism to securely obtain the latest key or a specific key with a key identifier for an Active Directory account. This service is new to Windows Server 2012 and does not run on previous versions of the Windows Server operating system. The Key Distribution Service shares a secret which is used to create keys for the account. These keys are periodically changed. For a group Managed Service Account the Windows Server 2012 domain controller computes the password on the key provided by the Key Distribution Services in addition to other attributes of the group Managed Service Account. Windows Server 2012 and Windows 8 member hosts can obtain the current and preceding password values by contacting a Windows Server 2012 domain controller.

Practical applications

Group Managed Service Accounts provide a single identity solution for services running on a server farm, or on systems behind Network Load Balance. By providing a group MSA solution,

services can be configured for the new group MSA principal and the password management is handled by Windows.

Using a group Managed Service Account, services or service administrators do not need to manage password synchronization between service instances. The group Managed Service Account supports hosts that are kept offline for an extended time period, and management of member hosts for all instances of a service. This means you can deploy a server farm that supports a single identity to which existing client computers can authenticate without knowing the instance of the service to which they are connecting.

Failover clusters do not support gMSAs. However, services that run on top of the Cluster service can use a gMSA or a sMSA if they are a Windows service, an App pool, a scheduled task, or natively support gMSA or sMSA.

New and changed functionality

The following table notes the changes to the MSA feature.

Feature/functionality	Windows Server 2008 R2	Windows Server 2012
Virtual Computer Accounts	X	X
Managed Service Accounts	X	X
Group Managed Service Accounts		X
Windows PowerShell cmdlets	X	X

For information about these changes in functionality for MSA, see What's New for Managed Service Accounts.

Deprecated functionality

For Windows Server 2012, the Windows PowerShell cmdlets default to managing the group Managed Service Accounts instead of the original standalone Managed Service Accounts.

Software requirements

Managed Service Accounts (and Virtual Computer Accounts) apply to both Windows Server 2008 R2 and Windows Server 2012. Group Managed Service Accounts can only be configured and administered on computers running Windows Server 2012 but can be deployed as a single service identity solution in domains that still have some DCs running operating systems earlier than Windows Server 2012. There are no domain or forest functional level requirements.

A 64-bit architecture is required to run the Windows PowerShell commands which are used to administer group Managed Service Accounts.

A managed service account is dependent upon Kerberos supported encryption types.When a client computer authenticates to a server using Kerberos the DC creates a Kerberos service ticket protected with encryption both the DC and server supports. The DC uses the account's msDS-SupportedEncryptionTypes attribute to determine what encryption the server supports and, if there is no attribute, it assumes the client computer does not support stronger encryption types. If the Windows Server 2012 host is configured to not support RC4 then authentication will always fail. For this reason, AES should always be explicitly configured for MSAs.

Note
Beginning with Windows Server 2008 R2, DES is disabled by default. For more information about supported encryption types, see Changes in Kerberos Authentication.

Group Managed Service Accounts are not applicable to Windows operating systems prior to Windows Server 2012.

Server Manager information

There are no configuration steps necessary to implement MSA and group MSA using Server Manager or the Install-WindowsFeature cmdlet.

See also

The following table provides links to additional resources related to Managed Service Accounts and group Managed Service Accounts.

Content type	References
Product evaluation	What's New for Managed Service Accounts Managed Service Accounts Documentation for Windows 7 and Windows Server 2008 R2 Service Accounts Step-by-Step Guide
Planning	Not yet available
Deployment	Not yet available
Operations	Managed Service Accounts in Active Directory
Troubleshooting	Not yet available
Evaluation	Getting Started with Group Managed Service Accounts
Tools and settings	Managed Service Accounts in Active Directory Domain Services

Content type	References
Community resources	Managed Service Accounts: Understanding, Implementing, Best Practices, and Troubleshooting
Related technologies	Active Directory Domain Services Overview

What's New for Managed Service Accounts

This topic for the IT professional describes the changes in functionality for Managed Service Accounts with the introduction of the group Managed Service Account (gMSA) in Windows Server® 2012 and Windows® 8.

Managed Service Accounts description

The managed service account is designed to provide services and tasks such as Windows services and IIS application pools to share their own domain accounts, while eliminating the need for an administrator to manually administer passwords for these accounts. It is a managed domain account that provides automatic password management.

New and changed functionality

The following describes what changes in functionality were made to MSA in Windows Server 2012 and Windows 8.

Group Managed Service Accounts

When a domain account is configured for a server in a domain, the client computer can authenticate and connect to that service. Previously, only two account types have provided identity without requiring password management. But these account types have limitations:

- Computer account is limited to one domain server and the passwords are managed by the computer
- Managed Service Account is limited to one domain server and the passwords are managed by the computer.

These accounts cannot be shared across multiple systems. Therefore, you must regularly maintain the account for each service on each system to prevent unwanted password expiration.

What value does this change add?

The group Managed Service Account solves this problem because the account password is managed by Windows Server 2012 domain controllers and can be retrieved by multiple Windows

Server 2012 systems. This minimizes the administrative overhead of a service account by allowing Windows to handle password management for these accounts.

What works differently?

On computers running Windows Server 2012 or Windows 8, a group MSA can be created and managed through the Service Control Manager so that numerous instances of the service, such as deployed over a server farm, can be managed from one server. Tools and utilities that you used to administer Managed Service Accounts, such as IIS Application Pool Manager, can be used with group Managed Service Accounts. Domain administrators can delegate service management to service administrators, who can manage the entire lifecycle of a Managed Service Account or the group Managed Service Account. Existing client computers will be able to authenticate to any such service without knowing which service instance they are authenticating to.

Removed or deprecated functionality

For Windows Server 2012, the Windows PowerShell cmdlets default to managing the group Managed Service Accounts instead of the server Managed Service Accounts.

See also

The following table provides additional resources for evaluating Managed Service Accounts and the group Managed Service Account.

Content type	References
Product evaluation	Group Managed Service Accounts Overview
Community resources	Managed Service Accounts: Understanding, Implementing, Best Practices, and Troubleshooting
Related technologies	**Active Directory Domain Services Overview**

Getting Started with Group Managed Service Accounts

This guide provides step-by-step instructions and background information for enabling and using group Managed Service Accounts in Windows Server 2012.

In this document

- Prerequisites
- Introduction

- Deploying a new server farm
- Adding member hosts to an existing server farm
- Updating the group Managed Service Account properties
- Decommissioning member hosts from an existing server farm
- Decommissioning an existing server farm

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Prerequisites

See the section in this topic on Requirements for group Managed Service Accounts.

Introduction

When a client computer connects to a service which is hosted on a server farm using network load balancing (NLB) or some other method where all the servers appear to be the same service to the client, then authentication protocols supporting mutual authentication such as Kerberos cannot be used unless all the instances of the services use the same principal. This means that each service has to use the same passwords/keys to prove their identity.

Note
Failover clusters do not support gMSAs. However, services that run on top of the Cluster service can use a gMSA or a sMSA if they are a Windows service, an App pool, a scheduled task, or natively support gMSA or sMSA.

Services have the following principals from which to choose, and each has certain limitations.

Principals	Scope	Services supported	Password management
Computer Account of Windows system	Domain	Limited to one domain joined server	Computer manages
Computer Account without Windows system	Domain	Any domain joined server	None
Virtual Account	Local	Limited to one server	Computer manages
Windows 7 standalone Managed Service Account	Domain	Limited to one domain joined server	Computer manages
User Account	Domain	Any domain joined server	None

Principals	Scope	Services supported	Password management
Group Managed Service Account	Domain	Any Windows Server 2012 domain-joined server	The domain controller manages, and the host retrieves

A Windows computer account, or a Windows 7 standalone Managed Service Account (sMSA), or virtual accounts cannot be shared across multiple systems. If you configure one account for services on server farms to share, you would have to choose a user account or a computer account apart from a Windows system. Either way, these accounts do not have the capability of single-point-of-control password management. This creates problem where each organization needs to create an expensive solution to update keys for the service in Active Directory and then distribute the keys to all instances of those services.

With Windows Server 2012, services or service administrators do not need to manage password synchronization between service instances when using group Managed Service Accounts (gMSA). You provision the gMSA in AD and then configure the service which supports Managed Service Accounts. You can provision a gMSA using the *-ADServiceAccount cmdlets which are part of the Active Directory module. Service identity configuration on the host is supported by:

- Same APIs as sMSA, so products which support sMSA will support gMSA
- Services which use Service Control Manager to configure logon identity
- Services which use the IIS manager for application pools to configure identity
- Tasks using Task Scheduler.

For more information about standalone Managed Service Accounts, see Managed Service Accounts. For more information about group Managed Service Accounts, see Group Managed Service Accounts Overview.

Requirements for group Managed Service Accounts

The following table lists the operating system requirements for Kerberos authentication to work with services using gMSA. The Active Directory requirements are listed after the table.

A 64-bit architecture is required to run the Windows PowerShell commands used to administer group Managed Service Accounts.

Operating system requirements

Element	Requirement	Operating system
Client Application host	RFC compliant Kerberos client	At least Windows XP
User account's domain DCs	RFC compliant KDC	At least Windows Server 2003
Shared service member hosts		Windows Server 2012
Member host's domain DCs	RFC compliant KDC	At least Windows Server 2003

Element	Requirement	Operating system
gMSA account's domain DCs	Windows Server 2012 DCs available for host to retrieve the password	Domain with Windows Server 2012 which can have some systems earlier than Windows Server 2012
Backend service host	RFC compliant Kerberos application server	At least Windows Server 2003
Backend service account's domain DCs	RFC compliant KDC	At least Windows Server 2003
Windows PowerShell for Active Directory	Windows PowerShell for Active Directory installed locally on a computer supporting a 64-bit architecture or on your remote management computer (for example, using the Remote Server Administration Toolkit)	Windows Server 2012

Active Directory Domain Service requirements

- The Active Directory schema in the gMSA domain's forest needs to be updated to Windows Server 2012 to create a gMSA.

  You can update the schema by installing a domain controller that runs Windows Server 2012 or by running the version of adprep.exe from a computer running Windows Server 2012. The object-version attribute value for the object CN=Schema,CN=Configuration,DC=Contoso,DC=Com must be 52.
- New gMSA account provisioned
- If you are managing the service host permission to use gMSA by group, then new or existing security group
- If managing service access control by group, then new or existing security group
- If the first master root key for Active Directory is not deployed in the domain or has not been created, then create it. The result of its creation can be verified in the KdsSvc Operational log, Event ID 4004.

For instructions how to create the key, see Create the Key Distribution Services KDS Root Key. Microsoft Key Distribution Service (kdssvc.dll) the root key for AD.

Lifecycle

The lifecycle of a server farm using the gMSA feature typically involves the following tasks:

- Deploying a new server farm
- Adding member hosts to an existing server farm
- Decommissioning member hosts from an existing server farm
- Decommissioning an existing server farm

- Removing a compromised member host from a server farm if required.

Deploying a new server farm

When deploying a new server farm, the service administrator will need to determine:

- If the service supports using gMSAs
- If the service requires inbound or outbound authenticated connections
- The computer account names for the member hosts for the service using the gMSA
- The NetBIOS name for the service
- The DNS host name for the service
- The Service Principal Names (SPNs) for the service
- The password change interval (default is 30 days).

Step 1: Provisioning group Managed Service Accounts

You can create a gMSA only if the forest schema has been updated to Windows Server 2012, the master root key for Active Directory has been deployed, and there is at least one Windows Server 2012 DC in the domain in which the gMSA will be created.

Membership in **Domain Admins**, **Account Operators** or ability to create msDS-GroupManagedServiceAccount objects, is the minimum required to complete the following procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To create a gMSA using the New-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell, type the following commands, and then press ENTER. (The Active Directory module will load automatically.)

   New-ADServiceAccount [-Name] <string> -DNSHostName <string> [-KerberosEncryptionType <ADKerberosEncryptionType>] [-ManagedPasswordIntervalInDays <Nullable[Int32]>] [-PrincipalsAllowedToRetrieveManagedPassword <ADPrincipal[]>] -SamAccountName <string> -ServicePrincipalNames <string[]>

Parameter	String	Example
Name	Name of the account	ITFarm1
DNSHostName	DNS host name of	ITFarm1.contoso.com

	service	
KerberosEncryptionType	Any encryption types supported by the host servers	RC4, AES128, AES256
ManagedPasswordIntervalInDays	Password change interval in days (default is 30 days if not provided)	90
PrincipalsAllowedToRetrieveManagedPassword	The computer accounts of the member hosts or the security group that the member hosts are a member of	ITFarmHosts
SamAccountName	NetBIOS name for the service if not same as	ITFarm1

	Name	
ServicePrincipalNames	Service Principal Names (SPNs) for the service	http/ITFarm1.contoso.com/contoso.com, http/ITFarm1.contoso.com/contoso, http/ITFarm1/contoso.com, http/ITFarm1/contoso

Important
The password change interval can only be set during creation. If you need to change the interval, you must create a new gMSA and set it at creation time.

Example

Enter the command on a single line, even though they might appear word-wrapped across several lines here because of formatting constraints.

```
New-ADServiceAccount ITFarm1 -DNSHostName ITFarm1.contoso.com
-PrincipalsAllowedToRetrieveManagedPassword ITFarmHosts -
KerberosEncryptionType RC4, AES128, AES256 -
ServicePrincipalNames http/ITFarm1.contoso.com/contoso.com,
http/ITFarm1.contoso.com/contoso, http/ITFarm1/contoso.com,
http/ITFarm1/contoso
```

Membership in **Domain Admins**, **Account Operators**, or ability to create msDS-GroupManagedServiceAccount objects, is the minimum required to complete this procedure. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To create a gMSA for outbound authentication only using the New-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   New-ADServiceAccount [-Name] <string> -RestrictToOutboundAuthenticationOnly [-ManagedPasswordIntervalInDays <Nullable[Int32]>] [-PrincipalsAllowedToRetrieveManagedPassword <ADPrincipal[]>]

Parameter	String	Example

Name	Name the account	ITFarm1
ManagedPasswordIntervalInDays	Password change interval in days (default is 30 days if not provided)	75
PrincipalsAllowedToRetrieveManagedPassword	The computer accounts of the member hosts or the security group that the member hosts are a member of	ITFarmHosts

Important
The password change interval can only be set during creation. If you need to change the interval, you must create a new gMSA and set it at creation time.

Example

```
New-ADServiceAccount ITFarm1 -RestrictToOutboundAuthenticationOnly -
PrincipalsAllowedToRetrieveManagedPassword ITFarmHosts
```

Step 2: Configuring service identity application service

To configure the services in Windows Server 2012, see the following feature documentation:

- IIS application pool

  For more information, see Specify an Identity for an Application Pool (IIS 7).
- Windows Services

  For more information, see Services.
- Tasks

  For more information, see the Task Scheduler Overview.

Other services could support gMSA. See the appropriate product documentation for details on how to configure those services.

Adding member hosts to an existing server farm

If using security groups for managing member hosts, add the computer account for the new member host to the security group (that the gMSA's member hosts are a member of) using one of the following methods.

Membership in **Domain Admins**, or the ability to add members to the security group object, is the minimum required to complete these procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

- Method 1: Active Directory Users and Computers

  For procedures how to use this method, see Add a computer account to a group using the Windows interface, and Manage Different Domains in Active Directory Administrative Center.
- Method 2: dsmod

  For procedures how to use this method, see Add a computer account to a group using the command line.
- Method 3: Windows PowerShell Active Directory cmdlet Add-ADPrincipalGroupMembership

  For procedures how to use this method, see Add-ADPrincipalGroupMembership.

If using computer accounts, find the existing accounts and then add the new computer account.

Membership in **Domain Admins**, **Account Operators**, or ability to manage msDS-GroupManagedServiceAccount objects, is the minimum required to complete this procedure. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To add member hosts using the Set-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Get-ADServiceAccount [-Name] <string> -PrincipalsAllowedToRetrieveManagedPassword
3. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Set-ADServiceAccount [-Name] <string> -PrincipalsAllowedToRetrieveManagedPassword <ADPrincipal[]>

Parameter	String	Example
Name	Name the account	ITFarm1
PrincipalsAllowedToRetrieveManagedPassword	The computer accounts of the	Host1, Host2, Host3

Parameter	String	Example
	member hosts or the security group that the member hosts are a member of	

Example

For example, to add member hosts type the following commands, and then press ENTER.

```
Get-ADServiceAccount [-Name] ITFarm1 -PrincipalsAllowedToRetrieveManagedPassword
```

```
Set-ADServiceAccount [-Name] ITFarm1-PrincipalsAllowedToRetrieveManagedPassword Host1
Host2 Host3
```

Updating the group Managed Service Account properties

Membership in **Domain Admins**, **Account Operators**, or the ability to write to msDS-GroupManagedServiceAccount objects, is the minimum required to complete these procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

Open the Active Directory Module for Windows PowerShell, and set any property by using the Set-ADServiceAccount cmdlet.

For detailed information how to set these properties, see Set-ADServiceAccount in the TechNet Library or by typing **Get-Help Set-ADServiceAccount** at the Active Directory module for Windows PowerShell command prompt and pressing ENTER.

Decommissioning member hosts from an existing server farm

Membership in **Domain Admins**, or ability to remove members from the security group object, is the minimum required to complete these procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

Step 1: Remove member host from gMSA

If using security groups for managing member hosts, remove the computer account for the decommissioned member host from the security group that the gMSA's member hosts are a member of using either of the following methods.

- Method 1: Active Directory Users and Computers

  For procedures how to use this method, see Delete a Computer Account using the Windows interface, and Manage Different Domains in Active Directory Administrative Center.
- Method 2: drsm

  For procedures how to use this method, see Delete a Computer Account using the command line.
- Method 3: Windows PowerShell Active Directory cmdlet Remove-ADPrincipalGroupMembership

  For detailed information how to do this, see Remove-ADPrincipalGroupMembership in the TechNet Library or by typing **Get-Help Remove-ADPrincipalGroupMembership** at the Active Directory module for Windows PowerShell command prompt and pressing ENTER.

If listing computer accounts, retrieve the existing accounts and then add all but the removed computer account.

Membership in **Domain Admins**, **Account Operators**, or ability to manage msDS-GroupManagedServiceAccount objects, is the minimum required to complete this procedure. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To remove member hosts using the Set-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Get-ADServiceAccount [-Name] <string> -PrincipalsAllowedToRetrieveManagedPassword
3. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Set-ADServiceAccount [-Name] <string> -PrincipalsAllowedToRetrieveManagedPassword <ADPrincipal[]>

Parameter	String	Example
Name	Name the account	ITFarm1
PrincipalsAllowedToRetrieveManagedPassword	The computer accounts of the member hosts or the security group that the member hosts are a member of	Host1, Host3

Example

For example, to remove member hosts type the following commands, and then press ENTER.

```
Get-ADServiceAccount [-Name] ITFarm1 -PrincipalsAllowedToRetrieveManagedPassword
```

```
Set-ADServiceAccount [-Name] ITFarm1 -PrincipalsAllowedToRetrieveManagedPassword Host1
Host3
```

Step 2: Removing a group Managed Service Account from the system

Remove the cached gMSA credentials from the member host using Uninstall-ADServiceAccount or the NetRemoveServiceAccount API on the host system.

Membership in **Administrators**, or equivalent, is the minimum required to complete these procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To remove a gMSA using the Uninstall-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Uninstall-ADServiceAccount < ADServiceAccount>

   Example

   For example, to remove the cached credentials for a gMSA named ITFarm1 type the following command, and then press ENTER:

   ```
   Uninstall-ADServiceAccount ITFarm1
   ```

For more information about the Uninstall-ADServiceAccount cmdlet, at the Active Directory module for Windows PowerShell command prompt, type **Get-Help Uninstall-ADServiceAccount**, and then press ENTER, or see the information on the TechNet web at Uninstall-ADServiceAccount.

Decommissioning an existing server farm

When decommissioning an existing server farm, you should remove the following objects from AD:

- If the gMSA is the only member, the security group that the gMSA is a member of which is used for access control
- If the security group is only used for member hosts, the security group that the member hosts are a member of

- gMSA.

To delete a security group, use Active Directory Users and Computers, dsrm or Remove-ADGroup. To delete a gMSA, use Active Directory Users and Computers or Remove-ADServiceAccount.

Step 1: Deleting Active Directory Objects

Membership in **Domain Admins**, **Account Operators**, or ability to delete msDS-GroupManagedServiceAccount and security group objects, is the minimum required to complete these procedures. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

To remove a gMSA using the Remove-ADServiceAccount cmdlet

1. On the Windows Server 2012 domain controller, run Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Remove-ADServiceAccount < ADServiceAccount>

   Example

   For example, to delete a gMSA named ITFarm1 type the following command, and then press ENTER:

   ```
   Remove-ADServiceAccount ITFarm1
   ```

For more information about the Remove-ADServiceAccount cmdlet, at the Active Directory module for Windows PowerShell command prompt, type **Get-Help Remove-ADServiceAccount**, and then press ENTER, or see the information on the TechNet web at Remove-ADServiceAccount.

Step 2: Removing a group Managed Service Account from the system

Use the procedures described in Step 2: Removing a group Managed Service Account from the system in this topic.

See also

- Group Managed Service Accounts Overview
- Managed Service Accounts

Create the Key Distribution Services KDS Root Key

This topic for the IT professional describes how to create a Microsoft Key Distribution Service (kdssvc.dll) root key on the domain controller using Windows PowerShell to generate group Managed Service Account passwords in Windows Server® 2012.

Windows Server 2012 domain controllers (DC) require a root key to begin generating gMSA passwords. The domain controllers will wait up to 10 hours from time of creation to allow all domain controllers to converge their AD replication before allowing the creation of a gMSA. The 10 hours is a safety measure to prevent password generation from occurring before all DCs in the environment are capable of answering gMSA requests. If you try to use a gMSA too soon the key might not have been replicated to all Windows Server 2012 DCs and therefore password retrieval might fail when the gMSA host attempts to retrieve the password. gMSA password retrieval failures can also occur when using DCs with limited replication schedules or if there is a replication issue.

Membership in the **Domain Admins** or **Enterprise Admins** groups, or equivalent, is the minimum required to complete this procedure. For detailed information about using the appropriate accounts and group memberships, see Local and Domain Default Groups.

Note

A 64-bit architecture is required to run the Windows PowerShell commands which are used to administer group Managed Service Accounts.

To create the KDS root key using the New-KdsRootKey cmdlet

1. On the Windows Server 2012 domain controller, run the Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   Add-KdsRootKey –EffectiveImmediately

   Tip

   The Effective time parameter can be used to give time for keys to be propagated to all DCs before use. Using Add-KdsRootKey –EffectiveImmediately will add a root key to the target DC which will be used by the KDS service immediately. However, other Windows Server 2012 DCs will not be able to use the root key until replication is successful.

For test environments with only one DC, you can create a KDS root key and set the start time in the past to avoid the interval wait for key generation by using the following procedure. Validate that a 4004 event has been logged in the kds event log.

To create the KDS root key in a test environment for immediate effectiveness

1. On the Windows Server 2012 domain controller, run the Windows PowerShell from the Taskbar.
2. At the command prompt for the Windows PowerShell Active Directory module, type the following commands, and then press ENTER:

   $a=Get-Date

   $b=$a.AddHours(-10)

   Add-KdsRootKey –EffectiveTime $b

   Or use a single command

   Add-KdsRootKey –EffectiveTime ((get-date).addhours(-10))

See Also

Getting Started with Group Managed Service Accounts

Kerberos Authentication Overview

Kerberos is an authentication protocol that is used to verify the identity of a user or host. This topic contains information about Kerberos authentication in Windows Server 2012 and Windows 8.

Feature description

The Windows Server operating systems implement the Kerberos version 5 authentication protocol and extensions for public key authentication, transporting authorization data, and delegation. The Kerberos authentication client is implemented as a security support provider (SSP), and it can be accessed through the Security Support Provider Interface (SSPI). Initial user authentication is integrated with the Winlogon single sign-on architecture.

The Kerberos Key Distribution Center (KDC) is integrated with other Windows Server security services that run on the domain controller. The KDC uses the domain's Active Directory Domain Services database as its security account database. Active Directory Domain Services is required for default Kerberos implementations within the domain or forest.

Practical applications

The benefits gained by using Kerberos for domain-based authentication are:

- **Delegated authentication.**

  Services that run on Windows operating systems can impersonate a client computer when accessing resources on the client's behalf. In many cases, a service can complete its work for the client by accessing resources on the local computer. When a client computer authenticates to the service, NTLM and Kerberos protocol provide the authorization

information that a service needs to impersonate the client computer locally. However, some distributed applications are designed so that a front-end service must use the client computer's identity when it connects to back-end services on other computers. Kerberos authentication supports a delegation mechanism that enables a service to act on behalf of its client when connecting to other services.

- **Single sign on.**

  Using Kerberos authentication within a domain or in a forest allows the user or service access to resources permitted by administrators without multiple requests for credentials. After initial domain sign on through Winlogon, Kerberos manages the credentials throughout the forest whenever access to resources is attempted.

- **Interoperability.**

  The implementation of the Kerberos V5 protocol by Microsoft is based on standards-track specifications that are recommended to the Internet Engineering Task Force (IETF). As a result, in Windows operating systems, the Kerberos protocol lays a foundation for interoperability with other networks in which the Kerberos protocol is used for authentication. In addition, Microsoft publishes Windows Protocols documentation for implementing the Kerberos protocol. The documentation contains the technical requirements, limitations, dependencies, and Windows-specific protocol behavior for Microsoft's implementation of the Kerberos protocol.

- **More efficient authentication to servers.**

  Before Kerberos, NTLM authentication could be used, which requires an application server to connect to a domain controller to authenticate every client computer or service. With the Kerberos protocol, renewable session tickets replace pass-through authentication. The server is not required to go to a domain controller (unless it needs to validate a Privilege Attribute Certificate (PAC)). Instead, the server can authenticate the client computer by examining credentials presented by the client. Client computers can obtain credentials for a particular server once and then reuse those credentials throughout a network logon session.

- **Mutual authentication.**

  By using the Kerberos protocol, a party at either end of a network connection can verify that the party on the other end is the entity it claims to be. NTLM does not enable clients to verify a server's identity or enable one server to verify the identity of another. NTLM authentication was designed for a network environment in which servers were assumed to be genuine. The Kerberos protocol makes no such assumption.

New and changed functionality

For a description of changes made to the Kerberos implementation in Windows, see What's New in Kerberos Authentication.

See also

Content type	References
Product evaluation	Kerberos Constrained Delegation Overview Dynamic Access Control: Scenario Overview Access Control and Authorization Overview **New and changed functionality** What's New in Active Directory Domain Services (AD DS)
Planning	What's New in Kerberos Authentication Kerberos Documentation for Windows 7, Windows Vista, Windows Server 2008 R2, and Windows Server 2008
Deployment	Not yet available
Operations	Not yet available
Troubleshooting	Not yet available
Security	Not yet available
Tools and settings	[MS-KILE]: Kerberos Protocol Extensions [MS-KKDCP]: Kerberos Key Distribution Center (KDC) Proxy Protocol Specification
Community resources	Kerberos Survival Guide
Related technologies	Active Directory Domain Services Overview

See Also

Windows Authentication Overview

What's New in Kerberos Authentication

This topic for the IT professional describes new capabilities and improvements to Windows implementation of the Kerberos authentication protocol in Windows Server 2012 and Windows 8.

Feature description

The Windows Server operating systems implement the Kerberos version 5 authentication protocol and extensions for public key and password-based authentication. The Kerberos authentication client is implemented as a security support provider (SSP), which can be accessed

through the Security Support Provider Interface (SSPI). Initial user authentication is integrated with the Winlogon service single sign-on architecture. The Kerberos Key Distribution Center (KDC) is integrated in the domain controller with other security services in Windows Server. The KDC uses the domain's Active Directory Domain Services (AD DS) as its security account database. AD DS is required for default Kerberos implementations within the domain or forest. In Windows Server 2012 and Windows 8, Kerberos authentication can be leveraged to address lack of connectivity to the domain controller from outside the corporate firewall. To do so, it proxies Kerberos authentication and password change messages for users who are requesting domain access through DirectAccess or Remote Desktop Services.

Improvements to reduce authentication failures due to large service tickets

As sign-in scenarios become more complex, the amount of authorization information that is contained in the authentication messages increases, which in some cases leads to an authentication failure. In Windows Server 2012 and Windows 8, improvements have been made to accommodate this increase.

KDC resource group compression

Within Kerberos authentication processes, authorization data is retrieved by the client from the domain controller and then sent to the resource. As the complexity of organizations increases, the number of groups that a user belongs to can increase significantly. When users are members of many groups in the resource domain, universal groups, and groups associated with security identifier history, the size of the authorization data in the service ticket can increase. The authentication message can exceed the default context token buffer size that Kerberos reports to applications, and this can cause authentication failures. In addition, when authorization data contains claims and device information, the size of that data increases.

What value does this change add?

This compression makes deployments easier when organizations have many users with high resource domain membership. In Windows Server 2012 and Windows 8, the KDC automatically compresses the groups in the resource domain, which can reduce the size of the service ticket and reduce application authentication failures. After the Windows Server 2012 KDC performs the compression, the supported operating systems Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 8, Windows 7, Windows Vista and Windows XP can use the more compact Kerberos service ticket.

What works differently?

KDC resource group compression is a default feature of the implementation of the Kerberos protocol in the Windows operating systems.

Increase in the Kerberos SSPI context token buffer size

When applications query to determine the maximum size of the authentication context token buffer so they can preallocate memory, they can fail authentication when the Kerberos or negotiate message returned is larger than expected.

What value does this change add?

The increase in the default size of this buffer prevents a greater number of application authentication requests from failing due to a limit in space.

What works differently?

In Windows Server 2012 and Windows 8, the default Kerberos SSPI context token buffer size has been increased to advised to set this value to more than 48,000 bytes.

Group Policy to set a maximum for the Kerberos SSPI context token buffer size

The new **Set maximum Kerberos SSPI context token buffer size** Kerberos administrative template policy setting enables you to set the value that is returned to applications that request the maximum size of the authentication context token buffer. This Group Policy setting replaces the need to manually set the MaxTokenSize registry value in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters.

What value does this change add?

This Group Policy setting enables you to finely control the context token buffer size to meet the authentication requirements of the applications in your environment.

What works differently?

If you do not already configure the MaxTokenSize registry value, you should consider applying this policy with a setting of 48,000 bytes. The new warning for large Kerberos tickets can be used to determine if your organization has or is approaching a large authentication context token buffer issue.

KDC with warning events for large Kerberos tickets

The new **Warning events for large Kerberos tickets** KDC administrative template policy setting enables you to control the system warning events that are logged by a KDC when tickets issued during Kerberos authentication are close to or greater than a configured threshold value size. The ticket size warnings are Event ID 31 in the System log.

What works differently?

But there are some considerations for setting this warning threshold. If it is set too high, warnings are not produced before authentication failures due to larger than expected authentication context token buffer size. If it is set too low, log entries increase, which makes analysis more difficult. The size should be set to one of the following:

- Your custom MaxTokenSize registry value

- The size deployed by the **Set maximum Kerberos SSPI context token buffer size** Group Policy setting

If the **Warning events for large Kerberos tickets** setting is not enabled, the threshold value defaults to 12,000 bytes, which is the default setting for MaxTokenSize for Kerberos tickets in Windows 7, Windows Server 2008 R2, and prior versions.

New and changed functionality for developers

The following features are introduced for Kerberos authentication in Windows Server 2012 and Windows 8 to enable you to build more robust applications.

Extended protection for untrusted targets

Extended Protection for Authentication is designed to protect against authentication relay attacks. It uses a combination of channel binding and service binding techniques. Extended Protection was introduced in Windows 7 and made available on a number of previously released operating systems.

Extended Protection relies on a collaboration of client-side and server-side technologies and a specific set of rules. For example, one rule is that the target of authentication, the Service Principal Name (SPN), must be intentional—even when a user is lured into authentication. The target SPN must be derived from user intent, and not from an untrusted source.

What works differently?

In Windows Server 2012, InitializeSecurityContext () adds support for the ISC_REQ_UNVERIFIED_TARGET_NAME flag for applications to use when they provide target names from an untrusted source. This allows services that are configured for Extended Protection to handle authentication to an untrusted source securely. Because NTLM and Kerberos clients that are attempting to authenticate using an untrusted SPN state that explicitly in their authentication messages, the services then will allow or fail the authentication request based on its Extended Protection configuration:

- Disabled: allow
- Enabled: fail

Kernel support for EncryptMessage and DecryptMessage

The Kerberos SSP now supports EncryptMessage/DecryptMessage from kernel applications.

Device sign-in with certificates

The Kerberos SSP now supports device sign-in by using a certificate.

New and changed functionality for IT professionals

Branch office support for authentication to resources outside the branch office

In Windows 7, when the user signed in by using a smart card with a mapped group, the group was lost when the user connected to resources outside the branch office. A device running Windows 8 in a branch office uses a hub domain controller running Windows Server 2012 to get a ticket-granting ticket (TGT). It uses the TGT to request service tickets for resources outside the branch office. There is no configuration required to use this feature.

Support for claims, compound authentication, and Kerberos armoring

If you want to create access control based on claims and compound authentication, you need to deploy Dynamic Access Control. This requires that you upgrade to Kerberos clients and use the KDC, which support these new authorization types. With Windows Server 2012, you do not have to wait until all the domain controllers and the domain functional level are upgraded to take advantage of new access control options.

For information about Dynamic Access Control in Windows Server 2012, see Dynamic Access Control: Scenario Overview and the related topics.

What value does this change add?

You can create access control based on claims and compound authentication.

What works differently?

The new **KDC support for claims, compound authentication, and Kerberos armoring** KDC administrative template policy setting allows you to configure a domain controller to support claims and compound authentication for Dynamic Access Control and Kerberos armoring by using Kerberos authentication. This policy setting is configured on the domain controller OU.

When the supported (or higher) setting is configured, domain controllers running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 will advertise domain support for claims and compound authentication for Dynamic Access Control and Kerberos armoring. No domain controllers running Windows Server 2003 can be in a domain that support claims and compound authentication for Dynamic Access Control and Kerberos armoring.

Additionally, the **Kerberos client support for claims, compound authentication and Kerberos armoring** administrative template policy setting enables you to configure devices running Windows 8 to support claims and compound authentication for Dynamic Access Control and Kerberos armoring by using Kerberos authentication. Devices running Windows 8 will fail authentication if they cannot find a domain controller running Windows Server 2012 . It is

important to ensure that there are sufficient domain controllers running Windows Server 2012 for any account, referral, and resource domains that are supported.

The following table lists the four configurations that are available in **KDC support for claims, compound authentication, and Kerberos armoring**.

Configuration	Results	Domain controller behavior in Windows Server 2012
Not supported (default)	No minimum requirement for domain controllers running Windows Server 2012	Claims not provided Compound authentication not supported Kerberos armoring not supported
Supported	All domain controllers advertise support for claims and compound authentication for Dynamic Access Control and Kerberos armoring Requires sufficient domain controllers running Windows Server 2012 to handle the authentication requests for devices running Windows 8 in the domain	Claims provided on request Compound authentication provided on request when resource supports it Kerberos armoring supported
Always provide claims	All domain controllers advertise support for claims and compound authentication for Dynamic Access Control and Kerberos armoring Requires Windows Server 2012 domain functional level	Claims always provided Compound authentication provided on request when resource supports it Kerberos armoring supported and Flexible Authentication via Secure Tunneling (RFC FAST) behavior supported
Fail unarmored authentication requests	All domain controllers advertise support for claims and compound authentication for Dynamic Access Control and Kerberos armoring Requires Windows Server 2012 domain functional level Requires that all FAST-aware	Claims always provided Compound authentication provided on request when resource supports it Rejects unarmored Kerberos messages and supports the Flexible Authentication via Secure Tunneling (RFC FAST)

Configuration	Results	Domain controller behavior in Windows Server 2012
	devices request authentication	behavior

Always provide claims and **Fail unarmored authentication requests** options cause intermittent authentication or access control failures if there are any domain controllers not running -Windows Server 2012 in the domain. So neither of these options will take effect until the domain is set at the Windows Server 2012 functional level. Until then, domain controllers running Windows Server 2012 will behave as if the **Supported** option is configured.

Supporting claims and compound authentication for Dynamic Access Control and Kerberos armoring will impact the domain controller because:

1. Secure Kerberos domain capability discovery is required, which results in additional message exchanges. Resource servers running Windows Server 2012 send protocol transition requests for non-Windows 8 devices, which results in additional messages exchanges to account domains and referral domains. These requests will go outside the site when no domain controller is available in the site.
2. Claims and compound authentication for Dynamic Access Control increases the size and complexity of the data in the message, which results in more processing time and greater Kerberos service ticket size.
3. Kerberos armoring fully encrypts Kerberos messages and signs Kerberos errors. This results in increased processing time, but it does not change the Kerberos service ticket size.

Supporting claims and compound authentication for Dynamic Access Control and Kerberos armoring (when the Supported option or higher is selected), could impact connectivity and file access because:

1. The setting configures a domain for the new access control model, which new services can choose to support.
2. Any device that supports claims, compound authentication, and Kerberos armoring cannot authenticate unless it can find a domain controller that supports claims, compound authentication and Kerberos armoring. This ensures that the client device receives complete domain authorization information during authentication for all access control checks (local and remote).
3. Devices running Windows 8 that do not support claims, compound authentication, and Kerberos armoring should not be configured with claims-based access controls for startup or sign-on services or with claims-based access controls for local files. Access will be denied, and claims will not be available until the domain is configured to the Windows Server 2012 domain functional level and configured to always provide claims.

The **KDC support for claims, compound authentication, and Kerberos armoring** administrative template policy setting also impacts network load and traffic patterns as follows:

1. Secure Kerberos domain capability discovery is required, which results in additional message exchanges.

2. Requests will go outside of the site when no domain controllers running Windows Server 2012 are available on the site for devices running Windows 8 that send authentication requests, and resource servers running Windows Server 2012 that send protocol transition requests.

Claims

Claims are new authorization data that are provided by Active Directory. When claims are provisioned, Windows Server 2012 KDCs can create service tickets with a principal's claims. Access tokens that are created from these service tickets include claims that can be used for access control.

What value does this change add?

This enables Windows Server 2012 to control access to resources with expressions in terms of:

- User claims
- Device claims, if the resource's domain supports claims and compound authentication for Dynamic Access Control and Kerberos armoring

What works differently?

To support resources that use claims-based access control, the principal's domains will need to be running one of the following:

- All Windows Server 2012 domain controllers
- Sufficient Windows Server 2012 domain controllers to handle all the Windows 8 device authentication requests
- Sufficient Windows Server 2012 domain controllers to handle all the Windows Server 2012 resource protocol transition requests to support non-Windows 8 devices

Additionally, you will need to configure on the domain controller OU, the new **KDC support for claims, compound authentication and Kerberos armoring** policy with "Supported" or higher setting, and if the client is retrieving claims, configure the new **Kerberos client support for claims, compound authentication and Kerberos armoring** policy on each client.

To support access control across forests, the forest root domains need the following:

- All Windows Server 2012 domain controllers. This helps ensure that claims are not lost from trusted forests.
- If users across forests sign in to devices in child domains, you must apply the QFE NetBIOS domain name\username format cannot be used with the Kerberos referral mechanism to log on to a computer in an across forest environment to down-level global catalogs.

To support access control for branch offices, you need the following:

- If groups or claims are generated from certificate-based sign-ins, sufficient hub domain controllers running Windows Server 2012 for the corresponding branch office read-only domain controllers (RODC) to handle all the Windows 8 device authentication requests for resources outside the branch office.
- For resources that do not sign in to an RODC, sufficient hub domain controllers running Windows Server 2012 for the corresponding branch office user's RODC. These domain

controllers handle all the Windows Server 2012 resource protocol transition requests to support devices that are not running Windows 8.

Compound authentication

Compound authentication is an extension to Flexible Authentication Secure Tunneling (FAST), which allows Kerberos clients to provide the device's TGT. This allows Windows Server 2012 KDCs to create service tickets with device authorization data for services that are hosted on Windows 8 and configured to support device authorization data. Access tokens that are created from these service tickets include the device's groups and claims, which can be used for access control.

What value does this change add?

This enables Windows Server 2012 to control access to resources with expressions in terms of:

- Devices
- Device groups
- Device claims, if the device's domain supports claims and compound authentication for Dynamic Access Control and Kerberos armoring

What works differently?

To support resources by using compound authentication for access control, the resource domains need one of the following:

- All Windows Server 2012 domain controllers
- Sufficient Windows Server 2012 domain controllers to handle all the Windows 8 device authentication requests

Additionally, you will need to configure on the domain controller OU, the new **KDC support for claims, compound authentication and Kerberos armoring** policy with "Supported" or higher setting, and if the client is retrieving claims, configure the new **Kerberos client support for claims, compound authentication and Kerberos armoring** policy on each client.

Access control based on device claims also require the device's domain to be provisioned for claims. Access control based on device or device groups, does not.

To support across forest access control, the forest root domains need:

- All Windows Server 2012 domain controllers. This helps ensure that claims are not lost from trusted forests. If domain controllers that are running a Windows Server version earlier than Windows Server 2012 exist in the forest root domains, those domain controllers will discard the claims. This results in intermittent access control failures, and the claims data is not disclosed to the trusting forest.

To support access control for resources outside the branch office, sufficient hub domain controllers running Windows Server 2012 for the corresponding branch office domain controllers (RODC).

The new **Support compound authentication** Kerberos administrative template policy setting enables you to explicitly configure support for KDCs to provide compound authentication to services by using the device's identity. Generally, this is not needed because applications will configure support for compound authentication. There are three options:

- **Never**: KDC will not provide compound authentication.
- **Automatic**: When a Dynamic Access Control-aware application is installed, the KDC always provides compound authentication. After the last Dynamic Access Control-aware application is removed, the KDC will not provide compound authentication.
- **Always**: KDC always provides compound authentication.

For services that are using a managed service account identity, Windows PowerShell must be used so that the KDC can start providing compound authentication. Use **Set-ADServiceAccount** with the **CompoundIdentitySupported** parameter.

Kerberos Armoring (Flexible Authentication Secure Tunneling (FAST))

Flexible Authentication Secure Tunneling (FAST) provides a protected channel between the Kerberos client and the KDC. FAST is implemented as Kerberos armoring in Windows Server 2012, and it is only available for authentication service (AS) and ticket-granting service (TGS) exchanges.

What value does this change add?

This enables the following for domain-joined systems running Windows Server 2012 and Windows 8:

- Protection against offline dictionary attacks. Kerberos armoring protects the user's preauthentication data, which is vulnerable to offline dictionary attacks when it is generated from a password.
- Authenticated Kerberos errors. Kerberos armoring protects user Kerberos authentications from KDC Kerberos error spoofing, which can downgrade to NTLM or weaker cryptography.
- Compounded authentication

What works differently?

Kerberos armoring uses a ticket-granting ticket (TGT) for the device to protect authentication service exchanges with the KDC, so the computer's authentication service exchange is not armored. The user's TGT is used to protect its TGS exchanges with the KDC.

The new **Fail authentication requests when Kerberos armoring is not available** Kerberos administrative template policy setting controls whether a device requires that Kerberos message exchanges be armored when communicating with a domain controller. When this policy setting is enabled, a device running Windows 8 only allows the authentication service exchange to be unarmored and fail authentication when a domain does not support Kerberos armoring, or when a domain controller running Windows Server 2012 cannot be found for a domain that supports Kerberos armoring.

Important

This policy setting is affected by another setting. When a domain does not support Kerberos armoring by enabling **KDC support claims, compound authentication, and Kerberos armoring**, all authentication attempts for all its users will fail from computers that have this policy enabled.

Resource-based constrained delegation across domains and forest

In Windows Server 2012, the new resource-based Kerberos constrained delegation can be used to provide constrained delegation when the front-end services and the resource services are not in the same domain.

Kerberos constrained delegation provides a safer form of delegation to be used by services. It restricts the services to which a given server can act on the behalf of a user. However, this model limited constrained delegation to a single domain and required a domain administrator because it is configured on the front-end service's account. In today's enterprise, front-end services are not designed to limit themselves to integration with only services in their domain. In addition, because provisioning is required by domain administrators, deploying new services cannot be done by service administrators.

Currently, a service administrator has no useful way of knowing which front-end services delegate to the resource services they own. This introduces security risks because any front-end service that can delegate to a resource service represents a potential attack point. If a server that hosts a front-end service that delegates to a resource service is compromised, the resource service may be compromised.

What value does this change add?

In Windows Server 2012, service administrators are able to specify which service identities can impersonate users to their services. Configuring constrained delegation on the resource permits services such as ISA Server, OWA, and SharePoint to be configured to use constrained delegation to authenticate to servers in other domains or in other forests. This provides support for multihop service solutions by using an existing Kerberos infrastructure.

What works differently?

To configure a resource service to allow a front-end service access on the behalf of users, Windows PowerShell cmdlets are used to provide a list of principals that are allowed. Use the **New** and **Set** cmdlets for **ADComputer**, **ADServiceAccount**, and **ADUser** with the **PrincipalsAllowedToDelegateToAccount** parameter.

Note

- You must apply the QFE Resource-based constrained delegation KDC_ERR_POLICY failure in environments that have Windows Server 2008 R2-based domain controllers to all down-level domain controllers in user account domains between the front-end and back-end domains.

To retrieve the list of principals, use the **Get** cmdlet for **ADComputer**, **ADServiceAccount**, and **ADUser** with the **PrincipalsAllowedToDelegateToAccount** parameter.

Enforcement for forest boundary for Kerberos full delegation

When full delegation is enabled for Kerberos on a server, the server can use the delegated ticket-granting ticket (TGT) to connect as the user to any server, including those across a one way trust.

In Windows Server 2012, a trust across forests can be configured to enforce the security boundary by disallowing forwarding TGTs to enter other forests.

What works differently?

To disable Kerberos full delegation on an outbound forest trust, all domain controllers in the forest root must support this feature, and you must set **netdom trust** for **EnableTGTDelegation on the across forest trust** to **No**.

KDC proxy service

The KDC proxy service is installed with select DirectAccess or Remote Desktop solutions to provide a way for Kerberos authentication to be used by Internet clients. DirectAccess and Remote Desktop clients create a TLS/SSL secure channel to the KDC proxy service running on the DirectAccess server or Remote Desktop Gateway (RD Gateway). To obtain the service ticket for the DirectAccess server or the Remote Desktop server, Kerberos messages are sent to the corresponding KDC proxy service. The service sends the request to a domain controller in the corporate network, and then returns the response.

What value does this change add?

This benefits the administrator in that the same protocol used in support of Active Directory domains, Kerberos, is also used for authentication traffic and credential management, such as password change for DirectAccess or Remote Desktop users who are requesting domain access.

What works differently?

For currently deployed systems that use Kerberos authentication, there are no differences. For some deployments of DirectAccess and Remote Desktop Services solutions with Windows Server 2012, the KDC proxy service is available by default.

The **Disable revocation checking for the SSL certificate of KDC proxy servers** Kerberos administrative template policy setting allows you to create a test environment for DirectAccess or RD Gateway with a KDC proxy server without needing to configure a CRL/OCSP environment. When enabled, the Kerberos client will not enforce the revocation check for the KDC proxy service's TLS/SSL certificate.

Warning

Because this setting allows the client to connect to services with an invalid certificate, it should never be set on production servers.

Strict KDC Validation default changes

In Windows 8, the default behavior for domain-joined smart card sign on for KDC validation was changed so that validation using NTAuth store policy, or the root CA on the smart card, is required.

Notes

This behavior is not the same as enabling **Require strict KDC validation**. Strict KDC validation is a more restrictive set of criteria which ensures all of the following are met:

- The domain controller has the private key for the certificate provided.
- For domain-joined systems, the certification authority (CA) that issued the KDC's certificate is in the NTAuth store.
- For non-domain-joined systems, the root CA of the KDC's certificate is in the Third-Party Root CA or Smart Card Trusted Roots store.
- KDC's certificate has the KDC EKU.
- KDC certificate's DNSName field of the subjectAltName (SAN) extension matches the DNS name of the domain.

For non-domain-joined smart card sign on, strict KDC validation is required.

To disable this default behavior, disable the Group Policy setting **Require strict KDC validation**.

Configuration and maintenance improvements

KDC service servicing without a reboot

The Windows Server 2012 Kerberos Key Distribution Center (KDC) service (kdcsvc) can be patched without requiring a system reboot.

New KDC system events

The following table describes new KDC events, which are written to the System log on a domain controller.

Event	Text	Resolution
31	Warning service ticket issued that approaches the configured ticket threshold	This is controlled by the new "Warning events for large Kerberos tickets" KDC administrative template policy setting. The Explain text of this policy setting has a detailed description.
32	KDC certificate is missing a KDC EKU	This event is triggered by a certificate sign-in request. The KDC has a certificate that is using an old template that does not have the KDC EKU, and it should be replaced.
33	Windows Server 2012 domain controller failed to configure the domain to advertise support for claims and compound	To resolve this issue, try running **gpupdate /force** on a domain controller running Windows Server 2012.

Event	Text	Resolution
	authentication for Dynamic Access Control and Kerberos armoring.	
34	Windows Server 2012 domain controllers are configured for supporting claims and compound authentication for Dynamic Access Control and Kerberos armoring option which requires Windows Server 2012 domain functional level and the domain is not at this level.	To resolve this condition, update the domain functional level to Windows Server 2012 or set the domain support for claims and compound authentication for Dynamic Access Control and Kerberos armoring to the **Supported** option.

KDC operational log

The new KDC operational log is located in **Application and Services/Microsoft/Windows/Kerberos-Key-Distribution-Center/Operational**. It provides information that can help you determine the cause of a certificate logon failure or when the KDC service was stopped or started. In Windows Server 2012, the following events are logged when the operational log is enabled:

- 120: Error: KDC certificate validation failure
- 200: Warning: KDC certificate not found
- 300: Informational: KDC service was started
- 301: Informational: KDC service was stopped
- 302: Informational: KDC certificate is current

Kerberos operational log

The new Kerberos operational log is located in **Application and Services/Microsoft/Windows/Security-Kerberos/Operational**. It provides information that can help you determine the cause of a failure without collecting a trace and contacting Microsoft Product Support. In Windows Server 2012, the following events are logged when the operational log is enabled:

- 100: Error: Authentication failure due to unregistered SPN
- 101: Error: Authentication failure due to duplicate SPN
- 102: Error: Certificate based authentication failure due to KDC certificate validation failure
- 103: Error: Certificate based authentication failure due to client certificate validation failure
- 104: Error: Certificate based authentication failure due to KDC either not supported certificate based authentication or not provisioned with a KDC certificate

- 105: Error: Attempt to retrieve group managed service account password failed
- 106: Error: Certificate based authentication failure due to KDC certificate validation failure due to KDC EKU missing
- 107: Error: Certificate based authentication failure due to KDC certificate validation failure due to wrong domain name
- 108: Error: Authentication failure due to no domain controller available and inability to use Kerberos proxy messages
- 109: Error: Authentication failure due to no domain controller available and attempt to use Kerberos proxy failed because the authentication request did not have authentication credentials for HTTP proxy
- 200: Warning: Domain controller could not be located
- 300: Informational: Domain controller which was located
- 301: Informational: Authentication using a stored credential
- 302: Informational: Domain controller not responding
- 303: Informational: Group managed service account password retrieved

Performance monitor counters

In Windows Server 2008 R2, Kerberos had four performance monitor counters:

- Forwarded Kerberos Requests for read-only domain controllers
- KDC AS Requests for domain controllers
- KDC TGS Requests for domain controllers
- Kerberos Authentications for servers

In Windows Server 2012, the following Kerberos performance counters have been added:

- KDC claims-aware AS Requests

  Measures the number of claims-aware Authentication Service requests (AS-REQ) that are being received by a domain controller. The load of existing claims-aware devices can be determined without requiring that the domain support claims, compound authentication, or provide Kerberos armoring. However, because devices running Windows 8 are not claims-aware by default, this counter cannot be used to determine initial authentication requests unless the devices are configured to support claims, compound authentication, and provide Kerberos armoring.

  Note
  Windows Server 2012 domain controllers control the provisioning of claims authorization data, and only those domain controllers that support claims and compound authentication for Dynamic Access Control and Kerberos armoring will actually return AS-REP messages that contain claims.

- KDC AS Requests with FAST for domain controllers

  Measures the number of armored Authentication Service request messages (AS-REQ) that are being processed by a domain controller.

Note
Domain controllers that do not support claims and compound authentication for Dynamic Access Control and Kerberos armoring cannot process armored Kerberos messages.

- KDC armored TGS Requests

  Measures the number of armored ticket-granting service request messages (TGS-REQ) that are being processed by a domain controller.

  Note
  Domain controllers that do not support claims and compound authentication for Dynamic Access Control and Kerberos armoring cannot process armored Kerberos messages.

- KDC claims-aware TGS Requests

  Measures the number of claims-aware ticket-granting service request messages (TGS-REQ) that are being received by a domain controller. The load of existing claims-aware devices can be determined without requiring the domain to support claims, compound authentication, and provide Kerberos armoring. However, because devices running Windows 8 are not claims-aware by default, this counter cannot be used to determine service ticket requests unless the devices are configured to support claims, compound authentication, and provide Kerberos armoring.

- KDC claims-aware service asserted identity TGS Requests

  Measures the number of service asserted identity (S4U2Self) ticket-granting service request messages (TGS-REQ) that are explicitly asking for claims.

  Note
  Domain controllers running Windows Server 2012 control the provisioning of claims authorization data, and only those domain controllers that support claims and compound authentication for Dynamic Access Control and Kerberos armoring will actually return TGS-REP messages that contain claims.

- KDC classic type constrained delegation TGS Requests

  Measures the number of constrained delegation (S4U2Proxy) ticket-granting service request messages (TGS-REQ) that are being processed by a domain controller by using the SPNs that are configured on the middle-tier service account object.

- KDC resource type constrained delegation TGS Requests

  Measures the number of constrained delegation (S4U2Proxy) ticket-granting service request messages (TGS-REQ) that are being processed by a domain controller by using the resource's account object configuration.

Changes to SetSPN

In Windows Server 2012, SetSPN will no longer be able to register duplicate SPNs in a domain. When SetSPN –a is used, SetSPN will treat it as SetSPN –s.

For eveloping information about SetSPN, see Service Principal Names (SPNs) SetSPN Syntax (Setspn.exe) on the TechNet Wiki. For the command reference, see **Setspn** in the TechNet Library.

See also

The following table provides additional resources to understand Kerberos authentication in Windows Server 2012.

Content type	References
Product evaluation	Kerberos Authentication Overview Dynamic Access Control: Scenario Overview
Community resources	Not yet available
Related technologies	Dynamic Access Control: Scenario Overview Active Directory Domain Services Overview

Kerberos Constrained Delegation Overview

This overview topic for the IT professional describes new capabilities for Kerberos constrained delegation in Windows Server 2012 R2 and Windows Server 2012.

Did you mean…

- Kerberos Delegation
- Domain Name System (DNS) Delegation
- Active Directory Delegation of Control

Feature description

Kerberos constrained delegation was introduced in Windows Server 2003 to provide a safer form of delegation that could be used by services. When it is configured, constrained delegation restricts the services to which the specified server can act on the behalf of a user. This requires domain administrator privileges to configure a domain account for a service and is restricts the account to a single domain. In today's enterprise, front-end services are not designed to be limited to integration with only services in their domain.

In earlier operating systems where the domain administrator configured the service, the service administrator had no useful way to know which front-end services delegated to the resource services they owned. And any front-end service that could delegate to a resource service represented a potential attack point. If a server that hosted a front-end service was compromised,

and it was configured to delegate to resource services, the resource services could also be compromised.

In Windows Server 2012 R2 and Windows Server 2012, ability to configure constrained delegation for the service has been transferred from the domain administrator to the service administrator. In this way, the back-end service administrator can allow or deny front-end services.

For detailed information about constrained delegation as introduced in Windows Server 2003, see Kerberos Protocol Transition and Constrained Delegation.

The Windows Server 2012 R2 and Windows Server 2012 implementation of the Kerberos protocol includes extensions specifically for constrained delegation. Service for User to Proxy (S4U2Proxy) allows a service to use its Kerberos service ticket for a user to obtain a service ticket from the Key Distribution Center (KDC) to a back-end service. These extensions allow constrained delegation to be configured on the back-end service's account, which can be in another domain. For more information about these extensions, see [MS-SFU]: Kerberos Protocol Extensions: Service for User and Constrained Delegation Protocol Specification in the MSDN Library. For more information about these changes, see What works differently later in this topic.

Practical applications

Constrained delegation gives service administrators the ability to specify and enforce application trust boundaries by limiting the scope where application services can act on a user's behalf. Service administrators can configure which front-end service accounts can delegate to their back-end services.

By supporting constrained delegation across domains in Windows Server 2012 R2 and Windows Server 2012, front-end services such as Microsoft Internet Security and Acceleration (ISA) Server, Microsoft Forefront Threat Management Gateway, Microsoft Exchange Outlook Web Access (OWA), and Microsoft SharePoint Server can be configured to use constrained delegation to authenticate to servers in other domains. This provides support for across domains service solutions by using an existing Kerberos infrastructure. Kerberos constrained delegation can be managed by domain administrators or service administrators.

New and changed functionality

Resource-based constrained delegation across domains

Kerberos constrained delegation can be used to provide constrained delegation when the front-end service and the resource services are not in the same domain. Service administrators are able to configure the new delegation by specifying the domain accounts of the front-end services which can impersonate users on the account objects of the resource services.

What value does this change add?

By supporting constrained delegation across domains, services can be configured to use constrained delegation to authenticate to servers in other domains rather than using

unconstrained delegation. This provides authentication support for across domain service solutions by using an existing Kerberos infrastructure without needing to trust front-end services to delegate to any service.

What works differently?

A change in the underlying protocol allows constrained delegation across domains. The Windows Server 2012 R2 and Windows Server 2012 implementation of the Kerberos protocol includes extensions to Service for User to Proxy (S4U2Proxy) protocol. This is a set of extensions to the Kerberos protocol that allows a service to use its Kerberos service ticket for a user to obtain a service ticket from the Key Distribution Center (KDC) to a back-end service.

For implementation information about these extensions, see [MS-SFU]: Kerberos Protocol Extensions: Service for User and Constrained Delegation Protocol Specification in MSDN.

For more information about the basic message sequence for Kerberos delegation with a forwarded ticket-granting ticket (TGT) as compared to Service for User (S4U) extensions, see section 1.3.3 Protocol Overview in the [MS-SFU]: Kerberos Protocol Extensions: Service for User and Constrained Delegation Protocol Specification.

To configure a resource service to allow a front-end service access on the behalf of users, use Windows PowerShell cmdlets.

- To retrieve a list of principals, use the **Get-ADComputer**, **Get-ADServiceAccount**, and **Get-ADUser** cmdlets with the **–Properties PrincipalsAllowedToDelegateToAccount** parameter.
- To configure the resource service, use the **New-ADComputer**, **New-ADServiceAccount**, **New-ADUser**, **Set-ADComputer**, **Set-ADServiceAccount**, and **Set-ADUser** cmdlets with the **–PrincipalsAllowedToDelegateToAccount** parameter.

Software requirements

Resource-based constrained delegation can only be configured on a domain controller running Windows Server 2012 R2 and Windows Server 2012, but can be applied within a mixed-mode forest.

You must apply the following hotfix to all domain controllers running Windows Server 2012 in user account domains on the referral path between the front-end and back-end domains that are running operating systems earlier than Windows Server: Resource-based constrained delegation KDC_ERR_POLICY failure in environments that have Windows Server 2008 R2-based domain controllers. For more details or to download the hotfix, see article 2665790 in the Microsoft Knowledge Base.

See also

Content type	References
Product evaluation	What's New in Kerberos Authentication

Content type	References
Planning	Kerberos Protocol Transition and Constrained Delegation in Windows Server 2003. Kerberos Documentation for Windows 7, Windows Vista, Windows Server 2008 R2, and Windows Server 2008
Deployment	Not yet available
Operations	Not yet available
Troubleshooting	Not yet available
Security	Not yet available
Tools and settings	[MS-SFU]: Kerberos Protocol Extensions: Service for User and Constrained Delegation Protocol Specification Active Directory Cmdlets in Windows PowerShell
Community resources	Kerberos Survival Guide
Related technologies	Active Directory Domain Services Overview

NTLM Overview

This topic for the IT professional describes NTLM, any changes in functionality, and provides links to technical resources to Windows Authentication and NTLM for Windows Server 2012 and previous versions.

Did you mean…

- LMcompatibilitylevel

Feature description

NTLM authentication is a family of authentication protocols that are encompassed in the Windows Msv1_0.dll. The NTLM authentication protocols include LAN Manager version 1 and 2, and NTLM version 1 and 2. The NTLM authentication protocols authenticate users and computers based on a challenge/response mechanism that proves to a server or domain controller that a user knows the password associated with an account. When the NTLM protocol is used, a resource server must take one of the following actions to verify the identity of a computer or user whenever a new access token is needed:

- Contact a domain authentication service on the domain controller for the computer's or user's account domain, if the account is a domain account.
- Look up the computer's or user's account in the local account database, if the account is a local account.

Current applications

NTLM authentication is still supported and must be used for Windows authentication with systems configured as a member of a workgroup. NTLM authentication is also used for local logon authentication on non-domain controllers. Kerberos version 5 authentication is the preferred authentication method for Active Directory environments, but a non-Microsoft or Microsoft application might still use NTLM.

Reducing the usage of the NTLM protocol in an IT environment requires both the knowledge of deployed application requirements on NTLM and the strategies and steps necessary to configure computing environments to use other protocols. New tools and settings have been added to help you discover how NTLM is used in order to selectively restrict NTLM traffic. For information about how to analyze and restrict NTLM usage in your environments, see Introducing the Restriction of NTLM Authentication to access the Auditing and restricting NTLM usage guide.

New and changed functionality

There are no changes in functionality for NTLM for Windows Server 2012.

Removed or deprecated functionality

There is no removed or deprecated functionality for NTLM for Windows Server 2012.

Server Manager information

NTLM cannot be configured from Server Manager. You can use Security Policy settings or Group Policies to manage NTLM authentication usage between computer systems. In a domain, Kerberos is the default authentication protocol.

See also

The following table lists relevant resources for NTLM and other Windows authentication technologies.

Content type	References
Product evaluation	**Introducing the Restriction of NTLM Authentication** **Changes in NTLM Authentication**

Content type	References
Planning	IT Infrastructure Threat Modeling Guide Threats and Countermeasures: Security Settings in Windows Server 2003 and Windows XP **Threats and Countermeasures Guide: Security Settings in Windows Server 2008 and Windows Vista** **Threats and Countermeasures Guide: Security Settings in Windows Server 2008 R2 and Windows 7** **Designing Your Windows Authentication Strategy**
Deployment	Extended Protection for Authentication Auditing and restricting NTLM usage guide Ask the Directory Services Team : NTLM Blocking and You: Application Analysis and Auditing Methodologies in Windows 7 Windows Authentication Blog Configuring MaxConcurrentAPI for NTLM pass-through authentication
Development	Microsoft NTLM (Windows) [MS-NLMP]: NT LAN Manager (NTLM) Authentication Protocol Specification [MS-NNTP]: NT LAN Manager (NTLM) Authentication: Network News Transfer Protocol (NNTP) Extension [MS-NTHT]: NTLM Over HTTP Protocol Specification
Troubleshooting	Not yet available
Community resources	Is this horse dead yet: NTLM Bottlenecks and the RPC runtime
Related technologies	Windows Authentication Overview

Passwords Overview

This topic for the IT professional describes passwords as used in the Windows operating systems, and links to documentation and discussions about the use of passwords in a credential management strategy.

Did you mean…

- Personal Identifier Number (PIN) usage in Smart Cards
- Fine-grain password policy usage in Active Directory Domain Services
- Picture passwords - Signing in with a picture password
- Password usage in Managed Service Accounts
- Password usage in Group Managed Service Accounts Overview

Feature description

Operating systems and applications today are architected around passwords and even if you use smart cards or biometric systems, all accounts still have passwords and they can still be used in some circumstances. Some accounts, notably accounts used to run services, cannot even use smart cards and biometric tokens and therefore must use a password to authenticate. Windows protects passwords using cryptographic hashes.

For more information about Windows passwords, see the Passwords Technical Overview in the Windows Server 2008 TechNet Library.

Practical applications

In Windows and many other operating systems, the most common method for authenticating a user's identity is to use a secret passphrase or password. Securing your network environment requires that strong passwords be used by all users. This helps avoid the threat of a malicious user guessing a weak password, whether through manual methods or by using tools, to acquire the credentials of a compromised user account. This is especially true for administrative accounts. When you change a complex password regularly, it reduces the likelihood of a password attack compromising that account.

New and changed functionality

In Windows Server 2012 and Windows 8, picture passwords are new. Picture passwords are a combination of a user selected image coupled with a series of gestures. Picture password functionality is disabled on domain-joined computers. Links to more information about picture passwords are listed in See also below.

There has been no change to password functionality in Windows Server 2012 and Windows 8. No new Group Policy settings have been added. However, improvements and enhancements have been made in credential (and password) management, such as with picture passwords,

Credential Locker and signing in to Windows 8 with a Microsoft account, formerly known as a Windows Live ID.

Deprecated functionality

No password functionality has been deprecated in Windows Server 2012 and Windows 8.

Software requirements

In enterprise environments, passwords are typically managed with Active Directory Domain Services. Passwords can also be managed on the local computer using the settings in local Security Settings, Account Policies, Password Policy.

See also

This table lists additional resources for password features, technology and credential management.

Content type	References
Scenario documentation	Protecting your digital identity
Operations	Active Directory Users and Computers
Troubleshooting	Find out when your Password Expires - Active Directory PowerShell Blog
Security	Windows Server 2008 R2 and Windows 7 Threats and Countermeasures Guide: Account Policies Guidance to change and create strong passwords
Tools and settings	Group Policy Settings Reference for Windows and Windows Server on the Microsoft Download Center
Community resources	Protecting your digital identity Signing in to Windows 8 with a Windows Live ID Signing in with a picture password Optimizing picture password security
Related technologies	Active Directory Domain Services Overview

Content type	References
	Group Policy Overview **Credential Locker Overview**

Security Auditing Overview

This technical overview for the IT professional describes the security auditing features in Windows and how your organization can benefit from using these technologies to enhance the security and manageability of your network.

Did you mean…

Security Policy Settings Reference

Event Viewer

Feature description

Security auditing is a powerful tool to help maintain the security of an enterprise. Auditing can be used for a variety of purposes, including forensic analysis, regulatory compliance, monitoring user activity, and troubleshooting. Industry regulations in various countries or regions require enterprises to implement a strict set of rules related to data security and privacy. Security audits can help implement such policies and prove that these policies have been implemented. Also, security auditing can be used for forensic analysis, to help administrators detect anomalous behavior, to identify and mitigate gaps in security policies, and to deter irresponsible behavior by tracking critical user activities.

Practical applications

You can use Windows security and system logs to create a security events tracking system, to record and store network activities that are associated with potentially harmful behaviors, and to mitigate those risks. You can enable auditing based on categories of security events such as:

- Changes to user account and resource permissions.
- Failed attempts by users to log on.
- Failed attempts to access resources.
- Changes to system files.

In Windows Server 2008 R2 and Windows 7, the number of security audit policy settings was increased from nine to 53, and all auditing capabilities were integrated with Group Policy. This allows administrators to configure, deploy, and manage a wide range of settings in the Group Policy Management Console (GPMC) or the Local Security Policy snap-in for a domain, site, or organizational unit (OU). This makes it easier for IT professionals to track when precisely defined,

significant activities take place on the network. For more information, see Advanced Security Audit Policy Settings.

New and changed functionality

There are no new changes in functionality to security auditing in Windows Server 2012 R2.

In Windows Server 2012, changes to security auditing were introduced for the following purposes:

- **Reduce the volume of audits.** You can target audit policies to specific files and users based on resource attributes and user and device claims.
- **Improve the manageability of audit policies.** The introduction of Global Object Access Auditing provides an effective means for enforcing the application of security audit policies on resources. Combining Global Object Access Auditing with claims and Dynamic Access Control allows you to take this global enforcement mechanism and apply it to a more precise set of activities of potential interest.
- **Improve ability to locate critical security audit data.** Existing data access events can log additional information regarding user, computer, and resource claims. This makes it easier for event collection and analysis tools to be configured to get the most relevant event data quickly.
- **Enable security auditing of removable storage devices.** The growing popularity of removable storage devices makes their attempted use a significant security concern that needs to be monitored.

Dynamic claim-based auditing leads to more precise and easier-to-manage audit policies. It enables scenarios that have been impossible or too difficult to configure. In addition to these improvements, new audit events and categories for tracking changes to Dynamic Access Control (DAC) policy elements include:

- Resource attributes on files
- Central access policies associated with files
- User and device claims
- Resource property definitions
- Central access policy and central access rule definitions

The following are examples of audit policies that administrators can author:

- Anyone without a “High” security clearance who attempts to access documents classified as High Business Impact (HBI)—for example, Audit | Everyone | All-Access | Resource.BusinessImpact=HBI AND User.SecurityClearance!=High.
- Audit all vendors when they access documents related to projects that they are not working on—for example, Audit | Everyone | All-Access | User.EmploymentStatus=Vendor AND User.Project Not_AnyOf Resource.Project.

These policies help regulate the volume of audit events and limit them to only the most relevant data or users.

To provide a full view of events across the organization, Microsoft is working with partners to provide event collection and analysis tools, such as Microsoft System Center.

Managing security auditing

To use security auditing, you need to configure the system access control list (SACL) for an object, and apply the appropriate security audit policy to the user or computer. For more information, see Managing Security Auditing.

For information about managing Advanced Security Auditing, see Advanced Security Auditing Walkthrough.

For information about auditing Dynamic Access Control, see Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects.

Related resources

Content type	Resources
Product evaluation	Advanced Security Auditing FAQ Advanced Security Auditing Walkthrough What's New in Security Auditing for Windows Server 2012
Planning and Deployment	Planning and Deploying Advanced Security Audit Policies
Operations	Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects
Troubleshooting	Not yet available
Tools and settings	Advanced Security Audit Policy Settings
Community resources	Advanced Security Auditing in Windows 7 and Windows Server 2008 R2
Related technologies	Active Directory Domain Services Overview

Advanced Security Auditing FAQ

This topic for the IT professional lists questions and answers about understanding, deploying, and managing security audit policies.

- What is Windows security auditing and why might I want to use it?

- What is the difference between audit policies located in Local Policies\Audit Policy and audit policies located in Advanced Audit Policy Configuration?
- What is the interaction between basic audit policy settings and advanced audit policy settings?
- How are audit settings merged by Group Policy?
- What is the difference between an object DACL and an object SACL?
- Why are audit policies applied on a per-computer basis rather than per user?
- What are the differences in auditing functionality between versions of Windows?
- Can I use advanced audit policies from a domain controller running Windows Server 2003 or Windows 2000 Server?
- What is the difference between success and failure events? Is something wrong if I get a failure audit?
- How can I set an audit policy that affects all objects on a computer?
- How do I figure out why someone was able to access a resource?
- How do I know when changes are made to access control settings, by whom, and what the changes were?
- How can I roll back security audit policies from the advanced audit policy to the basic audit policy?
- How can I monitor if changes are made to audit policy settings?
- How can I minimize the number of events that are generated?
- What are the best tools to model and manage audit policies?
- Where can I find information about all the possible events that I might receive?
- Where can I find more detailed information?

What is Windows security auditing and why might I want to use it?

Security auditing is a methodical examination and review of activities that may affect the security of a system. In the Windows operating systems, security auditing is more narrowly defined as the features and services that enable an administrator to log and review events for specified security-related activities.

Hundreds of events occur as the Windows operating system and the applications that run on it perform their tasks. Monitoring these events can provide valuable information to help administrators troubleshoot and investigate security-related activities.

What is the difference between audit policies located in Local Policies\Audit Policy and audit policies located in Advanced Audit Policy Configuration?

The basic security audit policy settings in **Security Settings\Local Policies\Audit Policy** and the advanced security audit policy settings in **Security Settings\Advanced Audit Policy Configuration\System Audit Policies** appear to overlap, but they are recorded and applied differently. When you apply basic audit policy settings to the local computer by using the Local Security Policy snap-in, you are editing the effective audit policy, so changes made to basic audit policy settings will appear exactly as configured in Auditpol.exe.

There are a number of additional differences between the security audit policy settings in these two locations.

There are nine basic audit policy settings under **Security Settings\Local Policies\Audit Policy** and 53 settings under **Advanced Audit Policy Configuration**. The settings available in **Security Settings\Advanced Audit Policy Configuration** address similar issues as the nine basic settings in **Local Policies\Audit Policy**, but they allow administrators to be more selective in the number and types of events to audit. For example, the basic audit policy provides a single setting for account logon, and the advanced audit policy provides four. Enabling the single basic account logon setting would be the equivalent of setting all four advanced account logon settings. In comparison, setting a single advanced audit policy setting does not generate audit events for activities that you are not interested in tracking.

In addition, if you enable success auditing for the basic **Audit account logon events** setting, only success events will be logged for all account logon–related behaviors. In comparison, depending on the needs of your organization, you can configure success auditing for one advanced account logon setting, failure auditing for a second advanced account logon setting, success and failure auditing for a third advanced account logon setting, or no auditing.

The nine basic settings under **Security Settings\Local Policies\Audit Policy** were introduced in Windows 2000. Therefore, they are available in all versions of Windows released since then. The advanced audit policy settings were introduced in Windows Vista and Windows Server 2008. The advanced settings can only be used on computers running at least Windows 7, Windows Vista, Windows Server 2008 R2, or Windows Server 2008.

Notes

In Windows Vista and Windows Server 2008, advanced audit event settings were not integrated with Group Policy, and they could be deployed only by using logon scripts that were generated with the Auditpol.exe command-line tool. In Windows Server 2008 R2 and Windows 7, all auditing capabilities were integrated with Group Policy. This allows administrators to configure, deploy, and manage these settings in the Group Policy Management Console (GPMC) or Local Security Policy snap-in for a domain, site, or organizational unit (OU).

In Windows Server 2012, changes to security auditing were introduced to:

- **Reduce the volume of audits.** You can target audit policies to specific files and users based on resource attributes and user and device claims.
- **Improve the manageability of audit policies.** The introduction of Global Object Access Auditing in Windows Server 2008 R2 provides an effective means for enforcing the application of security audit policy on resources. Combining Global Object Access Auditing with claims and Dynamic Access Control allows you to apply this global enforcement mechanism to a more precise set of activities of potential interest.
- **Improve ability to locate critical security audit data.** Existing data access events can log additional information regarding user, computer, and resource claims. This makes it easier for event collection and analysis tools to be configured to get the most relevant event data quickly.
- **Enable security auditing of removable storage devices.** The growing popularity of removable storage devices makes their attempted use a significant security concern that needs to be monitored.

Dynamic claim-based auditing leads to more precise and easier-to-manage audit policies. It enables scenarios that have been impossible or too difficult to configure. In addition to these improvements, new audit events and categories for tracking changes to Dynamic Access Control (DAC) policy elements include or initiated:

- Changes to resource attributes on files
- Changes to central access policies associated with files
- User and device claims
- Changes to user and device claims and resource property definitions
- Changes to central access policy and central access rule definitions

What is the interaction between basic audit policy settings and advanced audit policy settings?

Basic audit policy settings are not compatible with advanced audit policy settings that are applied by using Group Policy. When advanced audit policy settings are applied by using Group Policy, the current computer's audit policy settings are cleared before the resulting advanced audit policy settings are applied. After you apply advanced audit policy settings by using Group Policy, you can only reliably set system audit policy for the computer by using the advanced audit policy settings.

Editing and applying the advanced audit policy settings in Local Security Policy modifies the local Group Policy Object (GPO), so changes made here may not be exactly reflected in Auditpol.exe if there are policies from other domain GPOs or logon scripts. Both types of policies can be edited and applied by using domain GPOs, and these settings will override any conflicting local audit policy settings. However, because the basic audit policy is recorded in the effective audit policy, that audit policy must be explicitly removed when a change is desired, or it will remain in the effective audit policy. Policy changes that are applied by using local or domain Group Policy settings are reflected as soon as the new policy is applied.

Important

Whether you apply advanced audit policies by using Group Policy or by using logon scripts, do not use both the basic audit policy settings under **Local Policies\Audit Policy** and the advanced settings under **Security Settings\Advanced Audit Policy Configuration**. Using both advanced and basic audit policy settings can cause unexpected results in audit reporting.

If you use Advanced Audit Policy Configuration settings or use logon scripts (for computers running Windows Vista or Windows Server 2008) to apply advanced audit policies, be sure to enable the **Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings** policy setting under **Local Policies\Security Options**. This will prevent conflicts between similar settings by forcing basic security auditing to be ignored.

How are audit settings merged by Group Policy?

By default, policy options that are set in GPOs and linked to higher levels of Active Directory sites, domains, and OUs are inherited by all OUs at lower levels. However, an inherited policy can be overridden by a GPO that is linked at a lower level.

For example, you might use a domain GPO to assign an organization-wide group of audit settings, but want a certain OU to get a defined group of additional settings. To accomplish this, you can link a second GPO to that specific lower-level OU. Therefore, a logon audit setting that is applied at the OU level will override a conflicting logon audit setting that is applied at the domain level (unless you have taken special steps to apply Group Policy loopback processing).

The rules that govern how Group Policy settings are applied propagate to the subcategory level of audit policy settings. This means that audit policy settings configured in different GPOs will be merged if no policy settings configured at a lower level exist. The following table illustrates this behavior.

Auditing subcategory	Setting configured in an OU GPO (higher priority)	Setting configured in a domain GPO (lower priority)	Resulting policy for the target computer
Detailed File Share Auditing	Success	Failure	Success
Process Creation Auditing	Disabled	Success	Disabled
Logon Auditing	Success	Failure	Failure

What is the difference between an object DACL and an object SACL?

All objects in Active Directory Domain Services (AD DS), and all securable objects on a local computer or on the network, have security descriptors to help control access to the objects. Security descriptors include information about who owns an object, who can access it and in what way, and what types of access are audited. Security descriptors contain the access control list (ACL) of an object, which includes all of the security permissions that apply to that object. An object's security descriptor can contain two types of ACLs:

- A discretionary access control list (DACL) that identifies the users and groups who are allowed or denied access
- A system access control list (SACL) that controls how access is audited

The access control model that is used in Windows is administered at the object level by setting different levels of access, or permissions, to objects. If permissions are configured for an object, its security descriptor contains a DACL with security identifiers (SIDs) for the users and groups that are allowed or denied access.

If auditing is configured for the object, its security descriptor also contains a SACL that controls how the security subsystem audits attempts to access the object. However, auditing is not completely configured unless a SACL has been configured for an object and a corresponding **Object Access** audit policy setting has been configured and applied.

Why are audit policies applied on a per-computer basis rather than per user?

In security auditing in Windows, the computer, objects on the computer, and related resources are the primary recipients of actions by clients including applications, other computers, and users. In a security breach, malicious users can use alternate credentials to hide their identity, or malicious applications can impersonate legitimate users to perform undesired tasks. Therefore, the most consistent way to apply an audit policy is to focus on the computer and the objects and resources on that computer.

In addition, because audit policy capabilities can vary between computers running different versions of Windows, the best way to ensure that the audit policy is applied correctly is to base these settings on the computer instead of the user.

However, in cases where you want audit settings to apply only to specified groups of users, you can accomplish this by configuring SACLs on the relevant objects to enable auditing for a security group that contains only the users you specify. For example, you can configure a SACL for a folder called Payroll Data on Accounting Server 1. This can audit attempts by members of the Payroll Processors OU to delete objects from this folder. The **Object Access\Audit File System** audit policy setting applies to Accounting Server 1, but because it requires a corresponding resource SACL, only actions by members of the Payroll Processors OU on the Payroll Data folder generates audit events.

What are the differences in auditing functionality between versions of Windows?

Basic audit policy settings are available in all versions of Windows since Windows 2000, and they can be applied locally or by using Group Policy. Advanced audit policy settings were introduced in Windows Vista and Windows Server 2008, but the settings can only be applied by using logon scripts in those versions. Advanced audit policy settings, which were introduced in Windows 7 and Windows Server 2008 R2, can be configured and applied by using local and domain Group Policy settings.

Can I use advanced audit policies from a domain controller running Windows Server 2003 or Windows 2000 Server?

To use advanced audit policy settings, your domain controller must be installed on a computer running Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 with Service Pack 2 (SP2). Windows 2000 Server is not supported.

What is the difference between success and failure events? Is something wrong if I get a failure audit?

A success audit event is triggered when a defined action, such as accessing a file share, is completed successfully.

A failure audit event is triggered when a defined action, such as a user logon, is not completed successfully.

The appearance of failure audit events in the event log does not necessarily mean that something is wrong with your system. For example, if you configure Audit Logon events, a failure event may simply mean that a user mistyped his or her password.

How can I set an audit policy that affects all objects on a computer?

System administrators and auditors increasingly want to verify that an auditing policy is applied to all objects on a system. This has been difficult to accomplish because the system access control lists (SACLs) that govern auditing are applied on a per-object basis. Thus, to verify that an audit policy has been applied to all objects, you would have to check every object to be sure that no changes have been made—even temporarily to a single SACL.

Introduced in Windows Server 2008 R2 and Windows 7, security auditing allows administrators to define global object access auditing policies for the entire file system or for the registry on a computer. The specified SACL is then automatically applied to every object of that type. This can be useful for verifying that all critical files, folders, and registry settings on a computer are protected, and for identifying when an issue with a system resource occurs. If a file or folder SACL and a global object access auditing policy (or a single registry setting SACL and a global object access auditing policy) are configured on a computer, the effective SACL is derived from combining the file or folder SACL and the global object access auditing policy. This means that an audit event is generated if an activity matches either the file or folder SACL or the global object access auditing policy.

How do I figure out why someone was able to access a resource?

Often it is not enough to know simply that an object such as a file or folder was accessed. You may also want to know why the user was able to access this resource. You can obtain this forensic data by configuring the **Audit Handle Manipulation** setting with the **Audit File System** or with the **Audit Registry** audit setting. For more information, see "Step 3: Creating and verifying an audit policy that provides the reason for object access" in the Advanced Security Auditing Walkthrough.

How do I know when changes are made to access control settings, by whom, and what the changes were?

To track access control changes on computers running Windows Server 2012 R2, Windows Server 2012 Windows 7, Windows Server 2008 R2, Windows Vista, or Windows Server 2008, you need to enable the following settings, which track changes to DACLs:

- **Audit File System** subcategory: Enable for success, failure, or success and failure
- **Audit Authorization Policy Change** setting: Enable for success, failure, or success and failure
- A SACL with **Write** and **Take ownership** permissions: Apply to the object that you want to monitor

In Windows XP and Windows Server 2003, you need to use the **Audit policy change** subcategory.

How can I roll back security audit policies from the advanced audit policy to the basic audit policy?

Applying advanced audit policy settings replaces any comparable basic security audit policy settings. If you subsequently change the advanced audit policy setting to **Not configured**, you need to complete the following steps to restore the original basic security audit policy settings:

1. Set all Advanced Audit Policy subcategories to **Not configured**.
2. Delete all audit.csv files from the %SYSVOL% folder on the domain controller.
3. Reconfigure and apply the basic audit policy settings.

Unless you complete all of these steps, the basic audit policy settings will not be restored.

How can I monitor if changes are made to audit policy settings?

Changes to security audit policies are critical security events. You can use the **Audit Audit Policy Change** setting to determine if the operating system generates audit events when the following types of activities take place:

- Permissions and audit settings on the audit policy object are changed
- The system audit policy is changed
- Security event sources are registered or unregistered
- Per-user audit settings are changed
- The value of **CrashOnAuditFail** is modified
- Audit settings on a file or registry key are changed
- A Special Groups list is changed

How can I minimize the number of events that are generated?

Finding the right balance between auditing enough network and computer activity and auditing too little network and computer activity can be challenging. You can achieve this balance by identifying the most important resources, critical activities, and users or groups of users. Then design a security audit policy that targets these resources, activities, and users. Useful guidelines and recommendations for developing an effective security auditing strategy can be found in Planning and Deploying Advanced Security Audit Policies.

What are the best tools to model and manage audit policies?

The integration of advanced audit policy settings with domain Group Policy, introduced in Windows 7 and Windows Server 2008 R2, is designed to simplify the management and implementation of security audit policies in an organization's network. As such, tools used to plan and deploy Group Policy Objects for a domain can also be used to plan and deploy security audit policies.

On an individual computer, the Auditpol command-line tool can be used to complete a number of important audit policy–related management tasks. For more information, see **Auditpol** Help.

In addition, there are a number of computer management products, such as the Audit Collection Services in the Microsoft System Center Operations Manager products, which can be used to collect and filter event data.

Where can I find information about all the possible events that I might receive?

Users who examine the security event log for the first time can be a bit overwhelmed by the number of audit events that are stored there (which can quickly number in the thousands) and by the structured information that is included for each audit event. Additional information about these events, and the settings used to generate them, can be obtained from the following resources:

- Windows 8 and Windows Server 2012 Security Event Details
- Security Audit Events for Windows 7 and Windows Server 2008 R2
- Security Audit Events for Windows Server 2008 and Windows Vista
- Security Audit Policy Reference
- Advanced Security Audit Policy Settings

Where can I find more detailed information?

To learn more about security audit policies, see the following resources:

- Planning and Deploying Advanced Security Audit Policies
- Security Monitoring and Attack Detection Planning Guide
- **Security Audit Policy Reference**
- Security Audit Events for Windows 7 and Windows Server 2008 R2
- Security Audit Events for Windows Server 2008 and Windows Vista
- Windows Security Logging and Other Esoterica

Advanced Security Auditing Walkthrough

This procedural topic for the IT professional describes the process of setting up an advanced security auditing policy infrastructure in a test environment.

About this guide

The security auditing enhancements introduced in Windows Server 2008 R2 and Windows 7 can help your organization audit compliance with important business-related and security-related rules by tracking precisely defined activities, for example:

- A group administrator has modified settings or data on servers that contain finance information.
- An employee within a defined group has accessed an important file.
- The correct system access control list (SACL) is applied to every file and folder or registry key on a computer or file share as a verifiable safeguard against undetected access.

The advanced security auditing settings allow you to select only the behaviors that you want to monitor and exclude audit results for behaviors that are of little or no concern to you, or behaviors that create an excessive number of log entries. Security audit policies can be applied by using domain Group Policy Objects, and audit policy settings can be modified, tested, and deployed to selected users and groups with relative simplicity.

This walkthrough demonstrates the process of setting up an advanced security auditing policy infrastructure in a test environment. It also guides you through the process of configuring some representative advanced security audit policy settings. When you have completed these initial tasks, you are encouraged to use the procedures in this guide to choose, configure, apply, and assess additional security audit policy settings.

During this process, you will create an Active Directory domain, install a supported version of the Windows server operating system on a member server, install a supported version of the Windows client operating system on a client computer, and configure new advanced security audit policy settings, including global object access auditing. In addition, this document examines the "reason for access" data that is available by using a number of audit policy settings.

Note

Because this topic applies to the supported Windows operating systems designated in the **Applies To** list at the beginning of this topic, combinations of operating systems can be configured to complete these walkthrough procedures. However, in this example, procedures relating to Windows Server 2008 R2 and Windows 7 are used.

After you complete the settings for this text environment, you can apply different sets of advanced security audit policy settings and assess how they might be used to enhance security in your organization.

As you complete the steps in this guide, you will be able to:

- Create and apply advanced audit policy settings to a defined group of computers in your organization.

- Verify that the audit policy settings are applied to a defined group of client computer in your organization.
- Use the "reason for access" security event data to identify the permissions that were used to determine whether a particular security event was triggered.
- Configure, apply, and analyze the impact of different audit policy settings to identify the settings that are important to your organization.
- Manage per-user auditing.

Deploying advanced audit policy settings in a test environment

After you complete this guide, you will have a working advanced security auditing infrastructure. You can then test it and learn about additional advanced security audit policy settings by logging on to CONTOSO-CLNT and ensuring that the correct audit policy is being applied on the computer.

Important
We recommend that you first use the procedures in this guide in a test lab environment. Walkthroughs are not meant to be used to deploy Windows features without additional deployment planning and documentation.

The test environment that is described in this guide includes three computers that are connected to a private network and use the following operating systems, applications, and services.

Note
The supported Windows operating systems are designated in the **Applies To** list at the beginning of this topic.

Computer name	Operating system	Applications and services
CONTOSO-DC	Supported versions of the Windows client operating system **Note** Supported versions of the Windows server operating system or Windows Server 2003 with Service Pack 2 (SP2) can also be used for the domain controller.	Active Directory Domain Services (AD DS) and Domain Name System (DNS)
CONTOSO-SRV	Supported versions of Windows server	Group Policy Management Console (GPMC)

Computer name	Operating system	Applications and services
CONTOSO-CLNT	Supported versions of Windows client **Note** Supported versions of Windows server can also be used as the client computer.	

Note
For more information about operating system compatibility and requirements, see Which Editions of Windows Support Advanced Audit Policy Configuration.

The computers form a private intranet and are connected through a common hub or Layer 2 switch. This configuration can be emulated in a virtual machine environment, if desired. This walkthrough uses private addresses throughout the test lab configuration. The private network ID 10.0.0.0/24 is used for the intranet. The domain controller for the domain named contoso.com is named CONTOSO-DC. The following figure shows the configuration of the test environment.

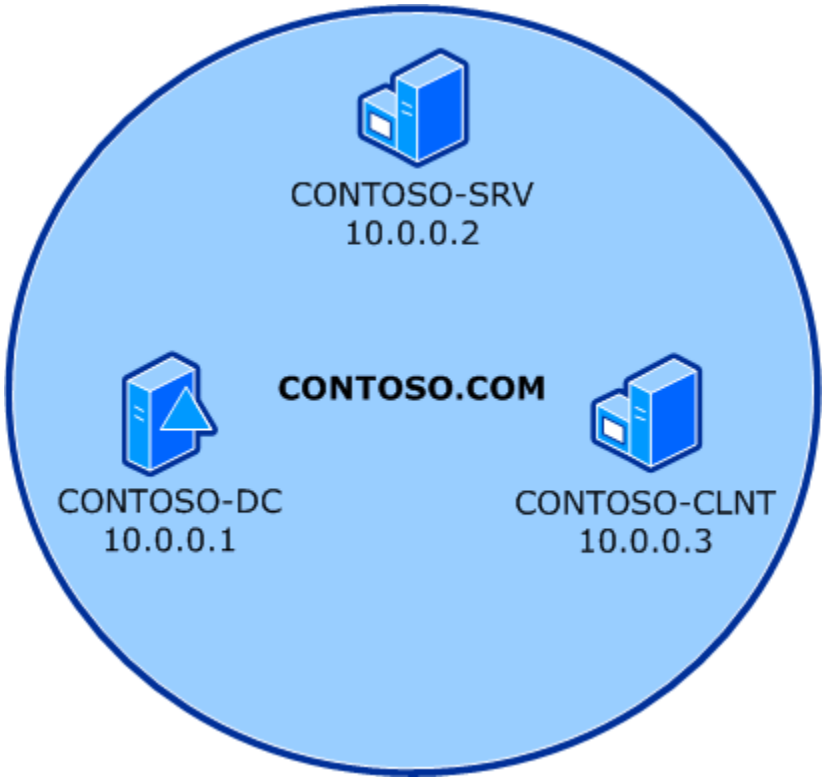

Steps for deploying advanced audit policies in a test environment

Complete the following steps to deploy advanced audit policy settings in a test environment.

Step 1: Set up the infrastructure

Step 2: Create and verify an advanced audit policy

Step 3: Create and verify an audit policy that provides the reason for object access

Step 4: Create and verify a global object access policy

Step 5: Create and verify additional advanced audit policies

Optional section: Roll back security audit policy from advanced audit policy to basic audit policy

Step 1: Set up the infrastructure

To prepare your test environment in the CONTOSO domain, you must complete the following tasks:

- Configure the domain controller (CONTOSO-DC)
- Configure the member server (CONTOSO-SRV)
- Configure the client computer (CONTOSO-CLNT)

Use the following table as a reference when you set up the appropriate computer names, operating systems, and network settings that are required to complete the steps in this guide.

Important
Before you configure your computers with static IP addresses, we recommend that you first complete two important tasks that require Internet connectivity: Complete Windows product activation and use Windows Update to obtain and install any available critical security updates.

Computer name	Operating system requirement	IP settings	DNS settings
CONTOSO-DC	Any supported Windows server operating systems, or, Windows Server 2008, or Windows Server 2003 with Service Pack 2 (SP2)	IP address: 10.0.0.1 Subnet mask: 255.255.255.0	Configured by DNS server role
CONTOSO-SRV	Windows Server 2008 R2	IP address: 10.0.0.2 Subnet mask: 255.255.255.0	Preferred: 10.0.0.1

Computer name	Operating system requirement	IP settings	DNS settings
CONTOSO-CLNT	Any of the supported Windows operating systems	IP address: 10.0.0.3 Subnet mask: 255.255.255.0	Preferred: 10.0.0.1

Configure the domain controller (CONTOSO-DC)

Depending on your environment, you may evaluate advanced audit policy settings in a domain running supported Windows server operating systems, Windows Server 2008, or Windows Server 2003. For this guide, the domain is running Windows Server 2008 R2.

To configure the domain controller, CONTOSO-DC, running Windows Server 2008 R2, you must:

- Install Windows Server 2008 R2
- Configure TCP/IP properties
- Install AD DS
- Create a Finance organizational unit (OU)

First, install Windows Server 2008 R2 on a stand-alone server.

To install Windows Server 2008 R2

1. Start your computer by using the Windows Server 2008 R2 product CD.
2. When prompted for a computer name, type **CONTOSO-DC**.
3. Follow the rest of the instructions that appear on your screen to finish the installation.

Next, configure TCP/IP properties so that CONTOSO-DC has an IPv4 static IP address of 10.0.0.1.

To configure TCP/IP properties

1. Log on to CONTOSO-DC with the CONTOSO-DC\Administrator account.
2. Click **Start**, click **Control Panel**, click **Network and Internet**, click **Network and Sharing Center**, click **Change Adapter Settings**, right-click **Local Area Connection**, and then click **Properties**.
3. On the **Networking** tab, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Click **Use the following IP address**.
   - In the **IP address** box, type **10.0.0.1**.
   - In the **Subnet mask** box, type **255.255.255.0**.
   - In the **Default gateway** box, type **10.0.0.1**.
5. In the **Preferred DNS server** box, type **10.0.0.1**, and then click **OK**.
6. On the **Networking** tab, clear the **Internet Protocol Version 6 (TCP/IPv6)** check box,

and then click **Close**.

Next, configure the computer as a domain controller running Windows Server 2008 R2 and create a Finance OU.

To configure CONTOSO-DC as a domain controller

1. Click **Start**, and then click **Run**. In the **Open** box, type **dcpromo**, and then click **OK**.
2. On the **Welcome to the Active Directory Domain Services Installation Wizard** page, click **Next**, and then click **Next** again.
3. Click **Create a new domain in a new forest**, and then click **Next**.
4. In the **FQDN of the forest root domain** box, type **contoso.com**, and then click **Next**.
5. Leave the default value in the **Domain NetBIOS name** box, and then click **Next**.
6. In the **Forest functional level** list, click **Windows Server 2008 R2**, and then click **Next**.
7. In the **Domain functional level** list, click **Windows Server 2008 R2**, and then click **Next**.
8. Ensure that the **DNS server** check box is selected, and then click **Next**.
9. Click **Yes**, confirming that you want to create a delegation for this DNS server.
10. On the **Location for Database, Log Files, and SYSVOL** page, click **Next**.
11. In the **Password** and **Confirm password** boxes, type a strong password, and then click **Next**.
12. On the **Summary** page, click **Next** to start the installation.
13. When the installation is complete, click **Finish**, and then click **Restart Now**.

    Note
    You must restart the computer after you complete this procedure.

To create a Finance OU in contoso.com

1. Log on to CONTOSO-DC with the CONTOSO-DC\Administrator account.
2. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
3. In the console tree, right-click **contoso.com**, point to **New**, and then click **Organizational Unit**.
4. Type the name of the new OU, **Finance**, and then click **OK**.

Configure the member server (CONTOSO-SRV) running Windows Server 2008 R2

To configure the member server, CONTOSO-SRV, you must:

- Install Windows Server 2008 R2
- Configure TCP/IP properties
- Join CONTOSO-SRV to the domain contoso.com
- Add CONTOSO-SRV to the Finance OU

- Install the GPMC

First, install Windows Server 2008 R2 as a stand-alone server.

To install Windows Server 2008 R2

1. Start your computer by using the Windows Server 2008 R2 product CD.
2. When prompted for a computer name, type **CONTOSO-SRV**.
3. Follow the remaining instructions that appear on your screen to finish the installation.

Next, configure TCP/IP properties so that CONTOSO-SRV has a static IP address of 10.0.0.2. In addition, configure the DNS server by using the IP address of CONTOSO-DC (10.0.0.1).

To configure TCP/IP properties

1. Log on to CONTOSO-SRV with the CONTOSO-SRV\Administrator account or another user account in the local **Administrators** group.
2. Click **Start**, click **Control Panel**, double-click **Network and Sharing Center**, click **Manage Network Connections**, right-click **Local Area Connection**, and then click **Properties**.
3. On the **Networking** tab, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. Click **Use the following IP address**.
   - In the **IP address** box, type **10.0.0.2**.
   - In the **Subnet mask** box, type **255.255.255.0**.
   - In the **Default gateway** box, type **10.0.0.1**.
5. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
6. Click **OK**, and then click **Close** to close the **Local Area Connection Properties** dialog box.

Next, join CONTOSO-SRV to the contoso.com domain.

To join CONTOSO-SRV to the contoso.com domain

1. Click **Start**, right-click **Computer**, and then click **Properties**.
2. Click **Change settings** (in the right pane under **Computer name, domain, and workgroup settings**), and then click **Change**.
3. In the **Computer Name/Domain Changes** dialog box, click **Domain**, and then type **contoso.com**.
4. Click **More**, and in the **Primary DNS suffix of this computer** box, type **contoso.com**.
5. Click **OK**, and then click **OK** again.
6. When a **Computer Name/Domain Changes** dialog box appears prompting you for administrative credentials, provide the credentials for CONTOSO\Administrator, and then click **OK**.

7. When a **Computer Name/Domain Changes** dialog box appears welcoming you to the contoso.com domain, click **OK**.
8. When a **Computer Name/Domain Changes** dialog box appears telling you that the computer must be restarted, click **OK**, and then click **Close**.
9. Click **Restart Now**.

After the computer restarts, add CONTOSO-SRV to the Finance OU.

To add a server to the Finance OU

1. Log on to CONTOSO-DC with the CONTOSO-DC\Administrator account.
2. Click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.
3. In the console tree, right-click **contoso.com**.
4. In the console tree, right-click the **Finance** OU, point to **New**, and then click **Group**.
5. Type the name of the new group, **Computers**, and then in **Group scope**, click **Domain local**, and in **Group type**, click **Security group**.
6. Right-click **Computers**, and then click **Properties**. On the **Members** tab, click **Add**.
7. In **Enter the object names to select**, type **CONTOSO-SRV**, and then click **OK**.

Finally, use Server Manager to install the GPMC on CONTOSO-SRV. This will be used to configure the advanced security audit policy settings.

To install the GPMC

1. Log on to CONTOSO-SRV as a member of the local **Administrators** group.
2. Click **Start**, point to **Administrative Tools**, and then click **Server Manager**.
3. Under **Feature Summary**, click **Add Features**.
4. Select the **Group Policy Management** check box, and then click **Install**.
5. Close Server Manager.

Configure the client computer (CONTOSO-CLNT)

To configure CONTOSO-CLNT, you must:

- Install Windows 7
- Configure TCP/IP properties
- Join CONTOSO-CLNT to the domain contoso.com

To install Windows 7

1. Start your computer by using the Windows 7 product CD.
2. Follow the instructions that appear on your screen, and when prompted for a computer name, type **CONTOSO-CLNT**.

Next, configure TCP/IP properties so that CONTOSO-CLNT has a static IP address of 10.0.0.3. In addition, configure the DNS server of CONTOSO-DC (10.0.0.1).

To configure TCP/IP properties

1. Log on to CONTOSO-CLNT with the CONTOSO-CLNT\Administrator account or another user account in the local **Administrators** group.
2. Click **Start**, click **Control Panel**, click **Network and Internet**, and then click **Network and Sharing Center**.
3. Click **Change adapter settings**, right-click **Local Area Connection**, and then click **Properties**.
4. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
5. On the **Networking** tab, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
6. Click **Use the following IP address**. In **IP address**, type **10.0.0.3**. In **Subnet mask**, type **255.255.255.0**.
7. Click **Use the following DNS server addresses**. In **Preferred DNS server**, type **10.0.0.1**.
8. Click **OK**, and then click **Close** to close the **Local Area Connection Properties** dialog box.

Next, join CONTOSO-CLNT to the contoso.com domain.

To join CONTOSO-CLNT to the contoso.com domain

1. Click **Start**, right-click **Computer**, and then click **Properties**.
2. Under **Computer name, domain, and workgroup settings**, click **Change settings**.
3. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
4. On the **Computer Name** tab, click **Change**.
5. In the **Computer Name/Domain Changes** dialog box, click **Domain**, and then type **contoso.com**.
6. Click **More**, and in the **Primary DNS suffix of this computer** box, type **contoso.com**.
7. Click **OK**, and then click **OK** again.
8. When a **Computer Name/Domain Changes** dialog box appears prompting you for administrative credentials, provide the credentials, and then click **OK**.
9. When a **Computer Name/Domain Changes** dialog box appears welcoming you to the contoso.com domain, click **OK**.
10. When a **Computer Name/Domain Changes** dialog box appears telling you that the computer must be restarted, click **OK**, and then click **Close**.
11. In the **System Settings Change** dialog box, click **Yes** to restart the computer.

Step 2: Create and verify an advanced audit policy

The nine basic audit policies under **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policy** allow you to configure security audit policy settings for broad sets of behaviors, some of which generate many audit events. An administrator has to review all events that are generated, whether they are of interest.

In any of the supported versions of the Windows operating systems, administrators can audit more specific aspects of client behavior on the computer or network, thus making it easier to identify the behaviors that are of greatest interest. For example, in **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policy**, there is only one policy setting for logon events, **Audit logon events**. In **Computer Configuration\Policies\Windows Settings\Security Settings\Advanced Audit Policy Configuration\System Audit Policies**, you can choose from eight policy settings in the **Logon/Logoff** category. This provides you with more detailed control of what aspects of logon and logoff you can track.

A default domain policy is automatically generated when a new domain is created. In this section, we will edit the default domain policy and add an advanced security audit policy setting that audits when a user successfully or unsuccessfully logs on to a computer in the CONTOSO domain.

To configure, apply, and validate an advanced domain logon audit policy setting, you must:

- Configure an advanced domain logon policy setting
- Ensure that Advanced Audit Policy Configuration settings are not overwritten
- Update Group Policy settings
- Verify that the advanced logon security audit policy settings were applied correctly

To configure an advanced domain logon audit policy setting

1. Log on to CONTOSO-SRV as a member of the local **Administrators** group.
2. Click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Forest: contoso.com**, double-click **Domains**, and then double-click **contoso.com**.
4. Right-click **Default Domain Policy**, and then click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **System Audit Policies**.
7. Double-click **Logon/Logoff**, and then double-click **Logon**.
8. Select the **Configure the following audit events** check box, select the **Success** check box, select the **Failure** check box, and then click **OK**.

When you use **Advanced Audit Policy Configuration** settings, you need to confirm that these settings are not overwritten by basic audit policy settings. The following procedure shows how to prevent conflicts by blocking the application of any basic audit policy settings.

To ensure that Advanced Audit Policy Configuration settings are not overwritten

1. On CONTOSO-SRV, click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.
2. In the console tree, double-click **Forest: contoso.com**, double-click **Domains**, and then double-click **contoso.com**.
3. Right-click **Default Domain Policy**, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
5. Double-click **Security Settings**, and then click **Security Options**.
6. Double-click **Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings**, and then click **Define this policy setting**.
7. Click **Enabled**, and then click **OK**.

Before you can verify the functionality of advanced security audit policy settings in the contoso.com domain, you will log on to CONTOSO-CLNT as the domain administrator of the contoso.com domain and ensure that the Group Policy settings have been applied.

To update Group Policy settings

1. Log on to CONTOSO-CLNT as CONTOSO\Administrator.
2. Click **Start**, point to **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.
3. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
4. Type **gpupdate**, and then press ENTER.

After the Group Policy settings have been applied, you can verify that the audit policy settings were applied correctly.

To verify that the advanced logon security audit policy settings were applied correctly

1. Log on to CONTOSO-CLNT as CONTOSO\Administrator.
2. Click **Start**, point to **All Programs**, click **Accessories**, right-click **Command Prompt**, and then click **Run as administrator**.
3. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
4. Type **auditpol.exe /get /category:***, and then press ENTER.
5. Verify that **Success**, **Failure**, or **Success and Failure** are shown to the right of **Logon**.

Step 3: Create and verify an audit policy that provides the reason for object access

One of the most common auditing needs is to track access to a particular file or folder. For example, you might need to identify an activity such as users writing to files that they should not have had access to. By enabling "reason for access" auditing, not only will you be able to track this type of activity, but you will also be able to identify the exact access control entry that allowed the undesired access, which can significantly simplify the task of modifying access control settings to prevent similar undesired object access in the future.

To configure, apply, and validate a reason for object access policy, you must:

- Configure the file system audit policy
- Enable auditing for a file or folder
- Enable the handle manipulation audit policy
- Ensure that Advanced Audit Policy Configuration settings are not overwritten
- Update Group Policy settings
- Review and verify the reason for access auditing data

To configure the file system audit policy

1. Log on to CONTOSO-SRV as a member of the local **Administrators** group.
2. Click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Forest: contoso.com**, double-click **Domains**, and then double-click **contoso.com**.
4. Right-click **Default Domain Policy**, and then click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **System Audit Policies**.
7. Double-click **Object Access**, and then double-click **File System**.
8. Select the **Configure the following events** check box, and then select the **Success**, **Failure**, or both **Success** and **Failure** check boxes.
9. Click **OK**.

The file system audit policy is only used to monitor objects for which auditing SACLs have been configured. The following procedure shows how to configure auditing for a file or folder.

To enable auditing for a file or folder

1. Log on to CONTOSO-CLNT as a member of the local **Administrators** group.
2. Create a new folder or .txt document.
3. Right-click the new object, click **Properties**, and click the **Security** tab.
4. Click **Advanced**, and then click the **Auditing** tab.

5. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
6. Click **Add**, type a user name or computer name in the format **contoso\user1**, and then click **OK**.
7. In the **Auditing Entries for** dialog box, select the permissions that you want to audit, such as **Full Control** or **Delete**.
8. Click **OK** four times to complete configuration of the object SACL.

In Windows 7 and Windows Server 2008 R2, the reason that someone is granted or denied access is added to the open handle event. This makes it possible for administrators to understand why someone was able to open a file, folder, or file share for a specific access. To enable this functionality, the handle manipulation audit policy also needs to be enabled to show **Success** access attempts that were allowed and **Failure** access attempts that were denied.

To enable the handle manipulation audit policy

1. Log on to CONTOSO-SRV as a member of the local **Administrators** group.
2. Click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Forest: contoso.com**, double-click **Domains**, and then double-click **contoso.com**.
4. Double-click the **Finance** OU, right-click **Finance Audit Policy**, and click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **System Audit Policies**.
7. Double-click **Object Access**, right-click **Handle Manipulation**, and click **Properties**.
8. Select the **Configure the following audit events** check box, select the **Success** and **Failure** check boxes, and then click **OK**.

After you have created this audit policy, confirm that these advanced audit policy settings cannot be overwritten. For more information, see the "To ensure that Advanced Audit Policy Configuration settings are not overwritten" procedure in the Step 2: Create and verify an advanced audit policy section.

Then apply the Group Policy updates by using the "To update Group Policy settings" procedure in the Step 2: Create and verify an advanced audit policy section.

After the updated Group Policy settings have been applied, be sure to log on to and log off from CONTOSO-CLNT and complete some tasks that will generate reason for object access events. After you have completed these steps, you can review the auditing data that provides the reason for access.

To review reason for access auditing data

1. On CONTOSO-CLNT, click **Start**, point to **Administrative Tools**, and then click **Event Viewer**.

2. Click **Windows Logs**, and then click **Security**.
3. In the **Actions** pane, click **Clear Log**.
4. Find the file or folder that you configured in the domain-level object access procedure, and modify the file or folder by using the permissions that you configured for the user account.
5. Go back to **Event Viewer**, and in the **Actions** pane, click **Refresh**.
6. In the **Event ID** column, click the event or events titled **4656**, scroll down to the **Access Request Information** section, and confirm the permissions that were used to perform the task.

Step 4: Create and verify a global object access policy

A global object access audit policy can be used to enforce object access audit policy for a computer, file share, or registry without having to configure and propagate conventional SACLs. Configuring and propagating SACLs is a more complex administrative task and is difficult to verify, particularly if you need to verify to an auditor that the security policy is being enforced. By using a global object access audit policy, you can enforce a security policy such as "Log all administrative Write activity on servers containing Finance information," and verify that critical assets are being protected.

In this case, you will be auditing any changes made to registry keys by members of a specified group rather than changes made to file system objects.

To configure, apply, and validate a global object access audit policy, you must:

- Configure a domain global object access audit policy
- Ensure that Advanced Audit Policy Configuration settings are not overwritten
- Update Group Policy settings
- Confirm that global object access auditing is taking place

To configure a domain global object access audit policy

1. Log on to CONTOSO-SRV as a member of the local **Administrators** group.
2. Click **Start**, point to **Administrative Tools**, and then click **Group Policy Management**.
3. In the console tree, double-click **Forest: contoso.com**, double-click **Domains**, and then double-click **contoso.com**.
4. Right-click **Default Domain Policy**, and then click **Edit**.
5. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
6. Double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, and then double-click **System Audit Policies**.
7. Double-click **Object Access**, and then double-click **Registry**.
8. Select the **Configure the following events** check box, select the **Success** and **Failure**

check boxes, and then click **OK**.

9. Double-click **Global Object Access Policies**, and then double-click **Registry**.
10. Select the **Define this policy setting** check box, and click **Configure**.
11. In the **Advanced Security Settings for Registry SACL** box, click **Add**.
12. Type a user name or computer name in the format **contoso\user1**, **user1@contoso.com**, or **CONTOSO-CLNT**, and click **OK**.
13. In the **Auditing Entry for Global Registry SACL** box, select the **Successful** or **Failed** activities for which you want to log audit entries—for example, **Create Subkey**, **Delete**, or **Read**.
14. Click **OK** three times to complete the audit policy configuration.

After you have created the audit policy, confirm that these advanced audit policy settings cannot be overwritten. For more information, see the "To ensure that Advanced Audit Policy Configuration settings are not overwritten" procedure in the Step 2: Create and verify an advanced audit policy section.

Then apply the Group Policy updates by using the "To update Group Policy settings" procedure in the Step 2: Create and verify an advanced audit policy section. After the updated Group Policy settings have been applied, log on to and log off from CONTOSO-CLNT.

To verify that the global object access policy has been applied

1. Open Registry Editor, and create and modify one or more registry settings.
2. Delete one or more of the registry settings that you created.
3. Open Event Viewer, and confirm that your activities resulted in audit events, even though you did not set explicit auditing SACLs on the registry settings that you created, modified, and deleted.

Step 5: Create and verify additional advanced audit policies

Now that you have created, applied, and validated the three basic types of advanced security audit policy settings, you can continue to identify and test additional advanced security audit policy settings by using the basic procedures outlined in the previous sections.

Additional information is available in **Computer Configuration\Policies\Windows Settings\Security Settings\Advanced Audit Policy Configuration\System Audit Policies** by right-clicking individual settings, clicking **Properties**, and clicking the **Explain** tab.

As you apply and test additional settings, consider how the audit event data that is generated can help you create a more secure network. In particular, consider the following:

- Is the information provided by these audit events useful?
- Is sufficient information provided by the audit data?
- Is too much information provided by the audit data?
- How can I adjust these audit policy settings to get only the information that I need?

Security auditing is a critical and essential tool to help you ensure that your network assets are secure. You should spend as much time as necessary to explore and understand the new advanced security audit policy settings in the supported Windows operating systems.

Managing per-user auditing

Security audit policy settings in the supported Windows operating systems can be configured and used only on a per-computer basis, not on a per-user basis. However, there are several ways to apply audit settings to specific users, including:

- Where available, configure the advanced security permissions on the object being audited so that the audit policy applies only to a specific group. For example, if you want the **Object Access** policy setting to apply to a file or folder, you can configure permissions on the file or folder so that object access is only tracked for the individuals or groups you specify. The procedure titled To enable auditing for a file or folder earlier in this document describes how to complete this task.
- Define and deploy per-user audit settings by using an audit policy text file, a logon script, and the Auditpol.exe command-line tool.

  Important
  Per-user auditing based on logon scripts can only be applied to individual users, not groups. You cannot use logon scripts to exclude subcategories or categories of audit policy settings for administrators.

The following procedure describes how to create an audit policy text file that can be deployed by using a logon script. For more information about using logon scripts to deploy an audit policy, see article 921469 in the Microsoft Knowledge Base.

To create an audit policy text file

1. At a command prompt, type **auditpol /set /user:***securityprincipalname***/category:"***subcategoryname***" /include /***Success or Failure***:enable** to add a per-user audit setting. Repeat this step for each audit policy subcategory and user or group that you want to add to your audit policy text file.

   Note
   To obtain a list of possible audit settings in report format, open a Command Prompt window, type **auditpol /list /subcategory:* /r**, and press ENTER. For more information about using Auditpol, see Auditpol set and Auditpol list.
2. At a command prompt, type **auditpol /backup /file:** *auditpolicyfilename***.txt** to export the policy.
3. Format your policy by opening *auditpolicyfilename.txt* and removing all lines except the first line of text and the per-user audit lines of text.

   Note
   Per-user audit policy text will be in the form: *ComputerName,S-1-XXXX,SubcategoryName,GUID,TextIncludeSettings,TextExcludeSettings,#*.

System settings will be in the form: *ComputerName,System,SubcategoryName,GUID,TextAuditSettings,#*. Also, be sure to remove the last six lines, which contain audit option settings.

4. When you have finished creating your file, on the **File** menu, click **Save As**, confirm that **ANSI** is selected in the **Encoding** list, and then click **OK**.
5. At a command prompt, type **auditpol /restore /file:** *auditpolicyfilename*.**txt**, and press ENTER to confirm that the desired audit settings are configured. Type **auditpol /list /user**, and press ENTER to list any users with per-user audit settings.
6. Copy the *auditpolicyfilename.txt* file to the Netlogon share of the domain controller that holds the primary domain controller (PDC) emulator role in the domain.

Important

Do not import audit policies that contain per-user auditing settings directly into a Group Policy Object (GPO). When per-user audit settings are deployed through Group Policy and not through logon scripts as described in this procedure, unexpected levels of failure events can appear in your security audit logs.

Optional section: Roll back security audit policy from advanced audit policy to basic audit policy

When you apply advanced audit policy settings, any comparable basic security audit policy settings are replaced. If you subsequently change the advanced audit policy setting to **Not configured**, you need to complete the following steps to restore the original basic security audit policy settings:

1. Set all Advanced Audit Policy sub-categories to **Not configured**.
2. Delete all audit.csv files from the %SYSVOL% folder on the domain controller.
3. Reconfigure and apply the basic audit policy settings.

Unless you complete each of these three steps, the basic audit policy settings will not be restored.

Which Editions of Windows Support Advanced Audit Policy Configuration

This reference topic for the IT professional describes which versions of the Windows operating systems support advanced security auditing policies.

All versions of Windows server operating systems and Windows client operating systems designated in the **Applies To** list at the beginning of this topic that can process Group Policy can be configured to use the advanced security auditing enhancements. Versions of the Windows operating system that cannot join a domain do not have access to these features. There is no difference in security auditing support between 32-bit and 64-bit versions.

Are there any special considerations?

In addition, the following special considerations apply to the various tasks associated with advanced security auditing enhancements:

- **Creating an audit policy.** To create an advanced security auditing policy, you must use a computer running any supported version of the Windows server operating system and Windows client operating system. You can use the Group Policy Management Console (GPMC) on a computer running a supported version of the Windows client operating system after installing the Remote Server Administration Tools.
- **Applying audit policy settings.** If you are using Group Policy to apply the advanced audit policy settings and global object access settings, client computers must be running any supported version of the Windows server operating system or Windows client operating system. In addition, only computers running any of these supported operating systems can provide "reason for access" reporting data.
- **Developing an audit policy model.** To plan advanced security audit settings and global object access settings, you must use the GPMC that targets a domain controller running a supported version of the Windows server operating system.
- **Distributing the audit policy.** After a Group Policy Object (GPO) that includes advanced security auditing settings is developed, it can be distributed by using domain controllers running any Windows Server operating system. However, if you cannot put client computers running a supported version of the Windows client operating system into a separate organizational unit (OU), you should use Windows Management Instrumentation (WMI) filtering to ensure that the advanced security auditing policy settings are applied only to client computers running a supported version of the Windows client operating system.

Note

Advanced security auditing policy settings can also be applied to client computers running Windows Vista. However, the audit policies for these computers must be created and applied separately by using Auditpol.exe logon scripts.

Important

Using both the basic auditing policy settings under **Local Policies\Audit Policy** and the advanced auditing policy settings under **Advanced Audit Policy Configuration** can cause unexpected results in audit reporting. Therefore, the two sets of audit policy settings should not be combined. If you use advanced audit policy configuration settings, you should enable the **Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings** policy setting under **Local Policies\Security Options**. This will prevent conflicts between similar settings by forcing basic security auditing to be ignored.

Planning and Deploying Advanced Security Audit Policies

This topic for the IT professional explains the options that security policy planners must consider and the tasks they must complete to deploy an effective security audit policy in a network that includes Advanced Security Audit policies.

Overview

Organizations invest a large portion of their information technology budgets on security applications and services, such as antivirus software, firewalls, and encryption. But no matter how much security hardware or software you deploy, how tightly you control the rights of users, or how carefully you configure security permissions on your data, you should not consider the job complete unless you have a well-defined, timely auditing strategy to track the effectiveness of your defenses and identify attempts to circumvent them.

To be well defined and timely, an auditing strategy must provide useful tracking data for an organization's most important resources, critical behaviors, and potential risks. In a growing number of organizations, it must also provide absolute proof that IT operations comply with corporate and regulatory requirements.

Unfortunately, no organization has unlimited resources to monitor every resource and activity on a network. If you do not plan well, you will likely have gaps in your auditing strategy. However, if you try to audit every resource and activity, you may find yourself with far too much monitoring data, including thousands of benign audit entries that an analyst needs to sift through to identify the narrow set of entries that warrant closer examination. This could cause delays or even prevent auditors from identifying suspicious activity. Thus, too much monitoring can leave an organization as vulnerable as not enough monitoring.

Auditing features introduced in Windows 7 and Windows Server 2008 R2 made it easier to audit important user, computer, and network activities in a focused, efficient manner. These features include:

- **Advanced audit policy settings** Allow administrators to more narrowly apply and manage detailed audit policy settings through their existing Group Policy framework.
- **"Reason for access" auditing** Enables administrators to specify and identify the permissions that were used to generate a particular object access security event.
- **Global object access auditing** Allows administrators to define system access control lists (SACLs) for an entire computer file system or registry.

To deploy these features and plan an effective security auditing strategy, you need to:

- Identify your most critical resources and the most important activities that need to be tracked.
- Identify the audit settings that can be used to track these activities.
- Assess the advantages and potential costs associated with each.
- Test these settings to validate your choices.

- Develop plans for deploying and managing your audit policy.

About this guide

This document will guide you through the steps needed to plan a security auditing policy that uses Windows auditing features. This policy must identify and address vital business needs, including:

- Network reliability
- Regulatory requirements
- Protection of the organization's data and intellectual property
- Users, including employees, contractors, partners, and customers
- Client computers and applications
- Servers and the applications and services running on those servers

The audit policy also must identify processes for managing audit data after it has been logged, including:

- Collecting, evaluating, and reviewing audit data
- Storing and (if required) disposing of audit data

By carefully planning, designing, testing, and deploying a solution based on your organization's business requirements, you can provide the standardized functionality, security, and management control that your organization needs.

Supported versions of Windows

Basic security audit policy settings

The basic security audit policy settings can be used to log audit events on any computer running any of the supported versions of Windows in addition to Windows Server 2008, Windows Server 2003, Windows 2000 Server, or a computer running Windows Vista, Windows XP, or Windows 2000 that can join a domain.

Note
Supported versions are designated in the **Applies To** list at the beginning of this topic,

Advanced audit policy settings

The advanced audit policy settings can be configured and used to log audit events on any computer running any of the supported versions of Windows.

In addition, the advanced audit policy settings can be applied by using a logon script on computers running Windows Vista or Windows Server 2008.

Terminology used in this guide

Understanding the following terms will help you apply the guidance in this document:

Group Policy. The Windows feature that allows administrators to specify options to manage configurations for groups of computers and users.

Audit policy. The nine security audit policy settings under **Security Settings\Local Policies\Audit Policy** and 53 audit policy settings under **Security Settings\Advanced Audit Policy Configuration** that determine the security events to be recorded in the security event log.

Success audit. An audit event that is triggered when a defined action, such as accessing a file share, is completed successfully.

Failure audit. An audit event that is triggered when a defined action, such as a user logon, is not completed successfully.

System access control lists (SACLs). A section of the security descriptor for objects, which is used to maintain per-object auditing information. For many types of objects, such as file system objects or registry settings, audit events are logged only if the relevant audit policy has been applied to the computer and a corresponding SACL has been configured on the object. However, you can ensure that a system-wide SACL is applied to all file system objects or all registry settings on a computer by using global object access auditing.

"Reason for access" audit. This list of access control entries (ACEs) provides the rights for which the decision to allow or deny access to the object was based. This can be useful for documenting the permissions, such as group memberships, that allow or prevent the occurrence of a particular auditable event. For more information about using "reason for access" auditing, see the Advanced Security Auditing Walkthrough.

Global object access auditing. In any of the supported versions of Windows , a computer-wide SACL can be applied to all objects in the file system or to all registry settings. The specified SACL is automatically applied to every object of that type, overriding any other SACLS that are configured for or applied to that object. For more information about using global object access auditing, see the Advanced Security Auditing Walkthrough.

Understanding the security audit policy design process

The process of designing and deploying a Windows security audit policy involves the following tasks, which are described in greater detail throughout this document:

- Identifying your Windows security audit policy deployment goals

  This section helps define the business objectives that will guide your Windows security audit policy. It also helps you define the resources, users, and computers that will be the focus of your security auditing.

- Mapping the security audit policy to groups of users, computers, and resources in your organization

  This section explains how to integrate security audit policy settings with domain Group Policy settings for different groups of users, computers, and resources. In addition, if your network

includes multiple versions of Windows client and server operating systems, it also explains when to use basic audit policy settings and when to use advanced security audit policy settings.

- Mapping your security auditing goals to a security audit policy configuration

  This section explains the categories of Windows security auditing settings that are available. It also identifies individual Windows security auditing policy settings that can be of particular value to address auditing scenarios.

- Planning for security audit monitoring and management

  This section helps you plan to collect, analyze, and store Windows audit data. Depending on the number of computers and types of activity that you want to audit, Windows event logs can fill up quickly. In addition, this section explains how auditors can access and aggregate event data from multiple servers and desktop computers. It also explains how to address storage requirements, including how much audit data to store and how it must be stored.

- Deploying the security audit policy

  This section provides recommendations and guidelines for the effective deployment of a Windows security audit policy. Configuring and deploying Windows audit policy settings in a test lab environment can help you confirm that the settings you have selected will produce the type of audit data you need. However, only a carefully staged pilot and incremental deployments based on your domain and organizational unit (OU) structure will enable you to confirm that the audit data you generate can be monitored and that it meets your organization's audit needs.

Identifying your Windows security audit policy deployment goals

A security audit policy must support and be a critical and integrated aspect of an organization's overall security design and framework.

Every organization has a unique set of data and network assets (such as customer and financial data and trade secrets), physical resources (such as desktop computers, portable computers, and servers), and users (which can include various internal groups such as finance and marketing, and external groups such as partners, customers, and anonymous users on the website). Not all of these assets, resources, and users justify the cost of an audit. Your task is to identify which assets, resources, and users provide the strongest justification for the focus of a security audit.

To create your Windows security audit plan, begin by identifying:

- The overall network environment, including the domains, OUs, and security groups.
- The resources on the network, the users of those resources, and how those resources are being used.
- Regulatory requirements.

Network environment

An organization's domain and OU structure provide a fundamental starting point for thinking about how to apply a security audit policy because it likely provides a foundation of Group Policy Objects (GPOs) and logical grouping of resources and activities that you can use to apply the audit settings that you choose. It is also likely that certain portions of your domain and OU structure already provide logical groups of users, resources, and activities that justify the time and resources needed to audit them. For information about how to integrate a security audit policy with your domain and OU structure, see Mapping the security audit policy to groups of users, computers, and resources in your organization later in this document.

In addition to your domain model, you should also find out whether your organization creates and maintains a systematic threat model. A good threat model can help you identify threats to key components in your infrastructure, so you can define and apply audit settings that enhance the organization's ability to identify and counter those threats.

Important
Including auditing within your organization's security plan also makes it possible to budget your resources on the areas where auditing can achieve the most positive results.

For additional details about how to complete each of these steps and how to prepare a detailed threat model, download the IT Infrastructure Threat Modeling Guide.

Data and resources

For data and resource auditing, you need to identify the most important types of data and resources (such as patient records, accounting data, or marketing plans) that can benefit from the closer monitoring that Windows auditing can provide. Some of these data resources might already be monitored through auditing features in products such as Microsoft SQL Server and Exchange Server. If so, you may want to consider how Windows auditing features can enhance the existing audit strategy. As with the domain and OU structure discussed previously, security auditing should focus on your most critical resources. You also must consider how much audit data you will be able to manage.

You can record if these resources have high business impact, medium business impact, or low business impact, the cost to the organization if these data resources are accessed by unauthorized users, and the risk that this access can pose to the organization. The type of access by users (such as Read, Modify, or Copy) can also pose different levels of risk to an organization.

Increasingly, data access and use is governed by regulations, and a breach can result in severe penalties and a loss in credibility for the organization. If regulatory compliance plays a role in how you manage your data, be sure to also document this information.

The following table provides an example of a resource analysis for an organization.

Resource class	Where stored	Organizational unit	Business impact	Security or regulatory requirements
Payroll data	Corp-Finance-1	Accounting: Read/Write on Corp-Finance-1 Departmental Payroll Managers: Write only on Corp-Finance-1	High	Financial integrity and employee privacy
Patient medical records	MedRec-2	Doctors and Nurses: Read/Write on Med/Rec-2 Lab Assistants: Write only on MedRec-2 Accounting: Read only on MedRec-2	High	Strict legal and regulatory standards
Consumer health information	Web-Ext-1	Public Relations Web Content Creators: Read/Write on Web-Ext-1 Public: Read only on Web-Ext-1	Low	Public education and corporate image

Users

Many organizations find it useful to classify the types of users they have and base permissions on this classification. This same classification can help you identify which user activities should be the subject of security auditing and the amount of audit data they will generate.

Organizations can create distinctions based on the type of rights and permissions needed by users to perform their jobs. For example, under the classification Administrators, larger organizations might assign local administrator responsibilities for a single computer, for specific applications such as Exchange Server or SQL Server, or for an entire domain. Under Users, permissions and Group Policy settings can apply to as many as all users in an organization or as few as a subset of the employees in a given department.

Also, if your organization is subject to regulatory requirements, user activities such as accessing medical records or financial data may need to be audited to verify that you are complying with these requirements.

To effectively audit user activity, begin by listing the different types of users in your organization and the types of data they need access to—in addition to the data they should not have access to.

Also, if external users can access any of your organization's data, be sure to identify them, including if they belong to a business partner, customer, or general user, the data they have access to, and the permissions they have to access that data.

The following table illustrates an analysis of users on a network. Although our example contains a single column titled "Possible auditing considerations," you may want to create additional columns to differentiate between different types of network activity, such as logon hours and permission use.

Groups	Data	Possible auditing considerations
Account administrators	User accounts and security groups	Account administrators have full privileges to create new user accounts, reset passwords, and modify security group memberships. We need a mechanism to monitor these changes.
Members of the Finance OU	Financial records	Users in Finance have Read/Write access to critical financial records, but no ability to change permissions on these resources. These financial records are subject to government regulatory compliance requirements.
External partners	Project Z	Employees of partner organizations have Read/Write access to certain project data and servers relating to Project Z, but not to other servers or data on the network.

Computers

Security and auditing requirements and audit event volume can vary considerably for different types of computers in an organization. These requirements can be based on:

- If the computers are servers, desktop computers, or portable computers.

- The important applications the computers run, such as Exchange Server, SQL Server, or Forefront Identity Manager.

  Note
  If the server applications (including Exchange Server and SQL Server) have audit settings. For more information about auditing in Exchange Server, see the Exchange 2010 Security Guide. For more information about auditing in SQL Server 2008, see Auditing (Database Engine). For SQL Server 2012, see SQL Server Audit (Database Engine).

- The operating system versions.

  Note
  The operating system version determines which auditing options are available and the volume of audit event data.

- The business value of the data.

For example, a web server that is accessed by external users requires different audit settings than a root certification authority (CA) that is never exposed to the public Internet or even to regular users on the organization's network.

The following table illustrates an analysis of computers in an organization.

Type of computer and applications	Operating system version	Where located
Servers hosting Exchange Server	Windows Server 2008 R2	ExchangeSrv OU
File servers	Windows Server 2012	Separate resource OUs by department and (in some cases) by location
Portable computers	Windows Vista and Windows 7	Separate portable computer OUs by department and (in some cases) by location
Web servers	Windows Server 2008 R2	WebSrv OU

Regulatory requirements

Many industries and locales have strict and specific requirements for network operations and how resources are protected. In the health care and financial industries, for example, there are strict guidelines for who has access to records and how they are used. Many countries have strict privacy rules. To identify regulatory requirements, work with your organization's legal department and other departments responsible for these requirements. Then consider the security configuration and auditing options that can be used to comply with and verify compliance with these regulations.

For more information, see the System Center Process Pack for IT GRC.

Mapping the security audit policy to groups of users, computers, and resources in your organization

By using Group Policy, you can apply your security audit policy to defined groups of users, computers, and resources. To map a security auditing policy to these defined groups in your organization, you should understand the following considerations for using Group Policy to apply security audit policy settings:

- The policy settings you identify can be applied by using one or more GPOs. To create and edit a GPO, use the Group Policy Management Console (GPMC). By using the GPMC to link a GPO to selected Active Directory sites, domains, and OUs, you apply the policy settings in the GPO to the users and computers in those Active Directory objects. An OU is the lowest-level Active Directory container to which you can assign Group Policy settings.
- For every policy setting that you select, you need to decide whether it should be enforced across the organization, or whether it should apply only to selected users or computers. You can then combine these audit policy settings into GPOs and link them to the appropriate Active Directory containers.
- By default, options set in GPOs that are linked to higher levels of Active Directory sites, domains, and OUs are inherited by all OUs at lower levels. However, a GPO that is linked at a lower level can overwrite inherited policies.

  For example, you might use a domain GPO to assign an organization-wide group of audit settings, but want a certain OU to get a defined group of additional settings. To accomplish this, you can link a second GPO to that specific lower-level OU. Therefore, a logon audit setting that is applied at the OU level will override a conflicting logon audit setting that is applied at the domain level (unless you have taken special steps to apply Group Policy loopback processing).
- Audit policies are computer policies. Therefore, they must be applied through GPOs that are applied to computer OUs, not to user OUs. However, in most cases you can apply audit settings for only specified resources and groups of users by configuring SACLs on the relevant objects. This enables auditing for a security group that contains only the users you specify.

  For example, you could configure a SACL for a folder called Payroll Data on Accounting Server 1. This can audit attempts by members of the Payroll Processors OU to delete objects from this folder. The **Object Access\Audit File System** audit policy setting applies to Accounting Server 1, but because it requires a corresponding resource SACL, only actions by members of the Payroll Processors OU on the Payroll Data folder generates audit events.
- Advanced security audit policy settings were introduced in Windows Server 2008 R2 or Windows 7 and can be applied to those operating systems designated in the **Applies To** list at the beginning of this topic. These advanced audit polices can only be applied by using Group Policy.

Note

These settings can also be applied to computers running Windows Vista or Windows Server 2008, but with these operating systems you must use logon scripts to apply advanced audit policies. Therefore, consider upgrading computers in high-security OUs first so that you can use Group Policy to apply and manage advanced audit policies to monitor high-value computers and their users.

Important

Whether you apply advanced audit policies by using Group Policy or by using logon scripts, do not use both the basic audit policy settings under **Local Policies\Audit Policy** and the advanced settings under **Security Settings\Advanced Audit Policy Configuration**. Using both basic and advanced audit policy settings can cause unexpected results in audit reporting.

If you use **Advanced Audit Policy Configuration** settings or use logon scripts (for computers running Windows Vista or Windows Server 2008) to apply advanced audit policies, be sure to enable the **Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings** policy setting under **Local Policies\Security Options**. This will prevent conflicts between similar settings by forcing basic security auditing to be ignored.

For more information, see article 921469 in the Microsoft Knowledge Base.

The following are examples of how audit policies can be applied to an organization's OU structure:

- Apply data activity settings to an OU that contains file servers. If your organization has servers that contain particularly sensitive data, consider putting them in a separate OU so that you can configure and apply a more precise audit policy to these servers.
- Apply user activity audit policies to an OU that contains all computers in the organization. If your organization places users in OUs based on the department they work in, consider configuring and applying more detailed security permissions on critical resources that are accessed by employees who work in more sensitive areas, such as network administrators or the legal department.
- Apply network and system activity audit policies to OUs that contain the organization's most critical servers, such as domain controllers, CAs, email servers, or database servers.

For more information about planning and deploying Group Policy, see the **Group Policy Planning and Deployment Guide**.

Mapping your security auditing goals to a security audit policy configuration

After you identify your security auditing goals, you can begin to map them to a security audit policy configuration. This audit policy configuration must address your most critical security auditing goals, but it also must address your organization's constraints, such as the number of computers that need to be monitored, the number of activities that you want to audit, the number

of audit events that your desired audit configuration will generate, and the number of administrators available to analyze and act upon audit data.

To create your audit policy configuration, you need to:

1. Explore all of the audit policy settings that can be used to address your needs.
2. Choose the audit settings that will most effectively address the audit requirements identified in the previous section.
3. Confirm that the settings you choose are compatible with the operating systems running on the computers that you want to monitor.
4. Decide which configuration options (Success, Failure, or both Success and Failure) you want to use for the audit settings.
5. Deploy the audit settings in a lab or test environment to verify that they meet your desired results in terms of volume, supportability, and comprehensiveness. Then deploy the audit settings in a pilot production environment to ensure that your estimates of how much audit data your audit plan will generate are realistic and that you can manage this data.

Exploring audit policy options

Security audit policy settings in the supported versions of Windows can be viewed and configured in the following locations:

- **Security Settings\Local Policies\Audit Policy**. For more information, see Audit Policy Settings Under Local Policies\Audit Policy.
- **Security Settings\Local Policies\Security Options**. For more information, see Audit Policy Settings Under Local Policies\Security Options.
- **Security Settings\Advanced Audit Policy Configuration**. For more information, see Advanced Security Audit Policy Settings.

Choosing audit settings to use

Depending on your goals, different sets of audit settings may be of particular value to you. For example, some settings under **Security Settings\Advanced Audit Policy Configuration** can be used to monitor the following types of activity:

- Data and resources
- Users
- Network

Important

These are only some of the settings that you should consider. For additional settings and their descriptions, see the Audit Policy Settings Under Local Policies\Security Options.

Settings that are described in the Reference might also provide valuable information about activity audited by another setting. For example, the settings used to monitor user activity and network activity have obvious relevance to protecting your data resources. Likewise, attempts to compromise data resources have huge implications for overall

network status, and potentially for how well you are managing the activities of users on the network.

Data and resource activity

For many organizations, compromising the organization's data resources can cause tremendous financial losses, in addition to lost prestige and legal liability. If your organization has critical data resources that need to be protected against any breach, the following settings can provide extremely valuable monitoring and forensic data:

- **Object Access\Audit File Share**. This policy setting allows you to track what content was accessed, the source (IP address and port) of the request, and the user account that was used for the access. The volume of event data generated by this setting will vary depending on the number of client computers that attempt to access the file share. On a file server or domain controller, volume may be high due to SYSVOL access by client computers for policy processing. If you do not need to record routine access by client computers that have permissions on the file share, you may want to log audit events only for failed attempts to access the file share.
- **Object Access\Audit File System**. This policy setting determines whether the operating system audits user attempts to access file system objects. Audit events are only generated for objects (such as files and folders) that have configured SACLs, and only if the type of access requested (such as Write, Read, or Modify) and the account that is making the request match the settings in the SACL.

  If success auditing is enabled, an audit entry is generated each time any account successfully accesses a file system object that has a matching SACL. If failure auditing is enabled, an audit entry is generated each time any user unsuccessfully attempts to access a file system object that has a matching SACL. The amount of audit data generated by the **Audit File System** policy setting can vary considerably, depending on the number of objects that have been configured to be monitored.

  Note

  To audit user attempts to access all file system objects on a computer, use the Global Object Access Auditing settings Registry (Global Object Access Auditing) or File System (Global Object Access Auditing).
- **Object Access\Audit Handle Manipulation**. This policy setting determines whether the operating system generates audit events when a handle to an object is opened or closed. Only objects with configured SACLs generate these events, and only if the attempted handle operation matches the SACL.

  Event volume can be high, depending on how SACLs are configured. When used together with the **Audit File System** or **Audit Registry** policy settings, the **Audit Handle Manipulation** policy setting can provide an administrator with useful "reason for access" audit data that details the precise permissions on which the audit event is based. For example, if a file is configured as a Read-only resource but a user attempts to save changes to the file, the audit event will log not only the event, but also the permissions that were used (or attempted to be used) to save the file changes.

- **Global Object Access Auditing**. A growing number of organizations are using security auditing to comply with regulatory requirements that govern data security and privacy. But demonstrating that strict controls are being enforced can be extremely difficult. To address this issue, the supported versions of Windows include two **Global Object Access Auditing** policy settings, one for the registry and one for the file system. When you configure these settings, they apply a global system access control SACL on all objects of that class on a system, which cannot be overridden or circumvented.

  Important
  The **Global Object Access Auditing** policy settings must be configured and applied in conjunction with the **Audit File System** and **Audit Registry** audit policy settings in the **Object Access** category. For more information about using the **Global Object Access Auditing** policy settings, see the Advanced Security Auditing Walkthrough.

User activity

The settings in the previous section relate to activity involving the files, folders, and network shares that are stored on a network, and the settings in this section focus on the users, including employees, partners, and customers, who may try to access those resources.

In the majority of cases, these attempts will be legitimate and a network needs to make vital data readily available to legitimate users. However in other cases, employees, partners, and others may attempt to access resources that they have no legitimate reason to access. Security auditing can be used to track a wide variety of user activities on a particular computer to diagnose and resolve problems for legitimate users and identify and address illegitimate activities. The following are a few important settings that you should evaluate to track user activity on your network:

- **Account Logon\Audit Credential Validation**. This is an extremely important policy setting because it enables you to track every successful and unsuccessful attempt to present credentials for a user logon. In particular, a pattern of unsuccessful attempts may indicate that a user or application is using credentials that are no longer valid, or attempting to use a variety of credentials in succession in hope that one of these attempts will eventually be successful. These events occur on the computer that is authoritative for the credentials. For domain accounts, the domain controller is authoritative. For local accounts, the local computer is authoritative.
- **Detailed Tracking\Audit Process Creation** and **Detailed Tracking\Audit Process Termination**. These policy settings can enable you to monitor the applications that a user opens and closes on a computer.
- **DS Access\Audit Directory Service Access** and **DS Access\Audit Directory Service Changes**. These policy settings provide a detailed audit trail of attempts to access create, modify, delete, move, or undelete objects in Active Directory Domain Services (AD DS). Only domain administrators have permissions to modify AD DS objects, so it is extremely important to identify malicious attempts to modify these objects. In addition, although domain administrators should be among an organization's most trusted employees, the use of **Audit Directory Service Access** and **Audit Directory Service Changes** settings allow you to monitor and verify that only approved changes are made to AD DS. These audit events are logged only on domain controllers.

- **Logon/Logoff\Audit Account Lockout**. Another common security scenario occurs when a user attempts to log on with an account that has been locked out. It is important to identify these events and to determine whether the attempt to use an account that has been locked out is malicious.
- **Logon/Logoff\Audit Logoff** and **Logon/Logoff\Audit Logon**. Logon and logoff events are essential to tracking user activity and detecting potential attacks. Logon events are related to the creation of logon sessions, and they occur on the computer that was accessed. For an interactive logon, events are generated on the computer that was logged on to. For network logon, such as accessing a shared resource, events are generated on the computer that hosts the resource that was accessed. Logoff events are generated when logon sessions are terminated.

  Note

  There is no failure event for logoff activity because failed logoffs (such as when a system abruptly shuts down) do not generate an audit record. Logoff events are not 100 percent reliable. For example, the computer can be turned off without a proper logoff and shutdown, and a logoff event is not generated.
- **Logon/Logoff\Audit Special Logon**. A special logon has administrator-equivalent rights and can be used to elevate a process to a higher level. It is recommended to track these types of logons. For more information about this feature, see article 947223 in the Microsoft Knowledge Base.
- **Object Access\Audit Certification Services**. This policy setting allows you to track and monitor a wide variety of activities on a computer that hosts Active Directory Certificate Services (AD CS) role services to ensure that only authorized users are performing or attempting to perform these tasks, and that only authorized or desired tasks are being performed.
- **Object Access\Audit File System** and **Object Access\Audit File Share**. These policy settings are described in the previous section.
- **Object Access\Audit Handle Manipulation**. This policy setting and its role in providing "reason for access" audit data is described in the previous section.
- **Object Access\Audit Registry**. Monitoring for changes to the registry is one of the most critical means that an administrator has to ensure malicious users do not make changes to essential computer settings. Audit events are only generated for objects that have configured SACLs, and only if the type of access that is requested (such as Write, Read, or Modify) and the account making the request match the settings in the SACL.

  Important

  On critical systems where all attempts to change registry settings need to be tracked, you can combine the **Audit Registry** policy setting with the **Global Object Access Auditing** policy settings to ensure that all attempts to modify registry settings on a computer are tracked.
- **Object Access\Audit SAM**. The Security Accounts Manager (SAM) is a database that is present on computers running Windows that stores user accounts and security descriptors for users on the local computer. Changes to user and group objects are tracked by the **Account Management** audit category. However, user accounts with the proper user rights could

potentially alter the files where the account and password information is stored in the system, bypassing any **Account Management** events.

- **Privilege Use\Audit Sensitive Privilege Use**. **Privilege Use** policy settings and audit events allow you to track the use of certain rights on one or more systems. If you configure this policy setting, an audit event is generated when sensitive rights requests are made.

Network activity

The following network activity policy settings allow you to monitor security-related issues that are not necessarily covered in the data or user activity categories, but that can be equally important for network status and protection.

- **Account Management**. The policy settings in this category can be used to track attempts to create, delete, or modify user or computer accounts, security groups, or distribution groups. Monitoring these activities complements the monitoring strategies you select in the user activity and data activity sections.
- **Account Logon\Audit Kerberos Authentication Service** and **Account Logon\Audit Kerberos Service Ticket Operations**. Audit policy settings in the **Account Logon** category monitor activities that relate to the use of domain account credentials. These policy settings complement the policy settings in the **Logon/Logoff** category. The **Audit Kerberos Authentication Service** policy setting allows you to monitor the status of and potential threats to the Kerberos service. The Audit **Kerberos Service Ticket Operations** policy setting allows you to monitor the use of Kerberos service tickets.

  Note
  Account Logon policy settings apply only to specific domain account activities, regardless of the computer that is accessed, whereas **Logon/Logoff** policy settings apply to the computer that hosts the resources being accessed.
- **Account Logon\Audit Other Account Logon Events**. This policy setting can be used to track a number of different network activities, including attempts to create Remote Desktop connections, wired network connections, and wireless connections.
- **DS Access**. Policy settings in this category allow you to monitor the AD DS role services, which provide account data, validate logons, maintain network access permissions, and provide other services that are critical to the secure and proper functioning of a network. Therefore, auditing the rights to access and modify the configuration of a domain controller can help an organization maintain a secure and reliable network. In addition, one of the key tasks performed by AD DS is the replication of data between domain controllers.
- **Logon/Logoff\Audit IPsec Extended Mode**, **Logon/Logoff\Audit IPsec Main Mode**, and **Logon/Logoff\Audit IPsec Quick Mode**. Many networks support large numbers of external users, including remote employees and partners. Because these users are outside the organization's network boundaries, IPsec is often used to help protect communications over the Internet by enabling network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and protection against replay attacks. You can use these settings to ensure that IPsec services are functioning properly.
- **Logon/Logoff\Audit Network Policy Server**. Organizations that use RADIUS (IAS) and Network Access Protection (NAP) to set and maintain security requirements for external

users can use this policy setting to monitor the effectiveness of these policies and to determine whether anyone is attempting to circumvent these protections.

- **Policy Change**. These policy settings and events allow you to track changes to important security policies on a local computer or network. Because policies are typically established by administrators to help secure network resources, any changes or attempts to change these policies can be an important aspect of security management for a network.
- **Policy Change\Audit Audit Policy Change**. This policy setting allows you to monitor changes to the audit policy. If malicious users obtain domain administrator credentials, they can temporarily disable essential security audit policy settings so that their other activities on the network cannot be detected.
- **Policy Change\Audit Filtering Platform Policy Change**. This policy setting can be used to monitor a large variety of changes to an organization's IPsec policies.
- **Policy Change\Audit MPSSVC Rule-Level Policy Change**. This policy setting determines if the operating system generates audit events when changes are made to policy rules for the Microsoft Protection Service (MPSSVC.exe), which is used by Windows Firewall. Changes to firewall rules are important for understanding the security state of the computer and how well it is protected against network attacks.

Confirm operating system version compatibility

Not all versions of Windows support advanced audit policy settings or the use of Group Policy to apply and manage these settings. Refer to the supported versions of Windows designated in the **Applies To** list at the beginning of this topic for version-applicable information that relates to the information in this guide. For more information, see Which Editions of Windows Support Advanced Audit Policy Configuration.

The audit policy settings under **Local Policies\Audit Policy** overlap with audit policy settings under **Security Settings\Advanced Audit Policy Configuration**. However, the advanced audit policy categories and subcategories make it possible to focus your auditing efforts on the most critical activities while reducing the amount of audit data that is less important to your organization.

For example, **Local Policies\Audit Policy** contains a single setting called Audit account logon events. When this setting is configured, it generates at least 10 types of audit events.

In comparison, the Account Logon category under **Security Settings\Advanced Audit Policy Configuration** provides the following advanced settings, which allow you to focus your auditing:

- Credential Validation
- Kerberos Authentication Service
- Kerberos Service Ticket Operations
- Other Account Logon Events

These settings allow you to exercise much tighter control over which activities or events generate event data. Some activities and events will be more important to your organization, so define the scope of your security audit policy as narrowly as possible.

Success, failure, or both

Whichever event settings you include in your plan, you also have to decide whether you want to log an event when the activity fails, when an activity succeeds, or both successes and failures. This is an important question, and the answer will be based on the criticality of the event and the implications of the decision on event volume.

For example, on a file server that is accessed frequently by legitimate users, you may be interested in logging an event only when an unsuccessful attempt to access data takes place, because this could be evidence of an unauthorized or malicious user. And in this instance, logging successful attempts to access the server would quickly fill the event log with benign events.

On the other hand, if the file share has extremely sensitive and valuable information, such as trade secrets, you may want to log every access attempt, whether successful or unsuccessful, so that you have an audit trail of every user who accessed the resource.

Planning for security audit monitoring and management

Networks can contain hundreds of servers running critical services or storing critical data, all of which need to be monitored. The number of client computers on the network can easily range into the tens or even hundreds of thousands. This may not be an issue if the ratio of servers or client computers per administrator is low. Even if an administrator who is responsible for auditing security and performance issues has relatively few computers to monitor, you need to decide how an administrator will obtain event data to review. Following are some options for obtaining the event data.

- Will you keep event data on a local computer until an administrator logs on to review this data? If so, then the administrator needs to have physical or remote access to the Event Viewer on each client computer or server, and the remote access and firewall settings on each client computer or server need to be configured to enable this access. In addition, you need to decide how often an administrator can visit each computer, and adjust the size of the audit log so that critical information is not deleted if the log reaches its maximum capacity.
- Will you collect event data so that it can be reviewed from a central console? If so, there are a number of computer management products, such as the Audit Collection Services in Microsoft System Center Operations Manager 2007 and 2012, which can be used to collect and filter event data. Presumably this solution enables a single administrator to review larger amounts of data than using the local storage option. But in some cases, this can make it more difficult to detect clusters of related events that can occur on a single computer. For more information about ACS pertaining to these versions, see Audit Collection Services (ACS) or Collecting Security Events Using Audit Collection Services in Operations Manager.

In addition, whether you choose to leave audit data on an individual computer or consolidate it at a central location, you need to decide how large the log file should be and what should happen when the log reaches its maximum size. To configure these options, open Event Viewer, expand

Windows Logs, right-click **Security**, and click **Properties**. You can configure the following properties:

- **Overwrite events as needed (oldest events first)**. This is the default option, which is an acceptable solution in most situations.
- **Archive the log when full, do not overwrite events**. This option can be used when all log data needs to be saved, but it also suggests that you may not be reviewing audit data frequently enough.
- **Do not overwrite events (Clear logs manually)**. This option stops the collection of audit data when the log file reaches its maximum size. Older data is retained at the expense of the most recent audit events. Use this option only if you do not want to lose any audit data, do not want to create an archive of the event log, and are committed to reviewing data before the maximum log size is reached.

You can also configure the audit log size and other key management options by using Group Policy settings. You can configure the event log settings in the following locations within the GPMC: **Computer Configuration\Administrative Templates\Windows Components\Event Log Service\Security**. These options include:

- **Maximum Log Size (KB)**. This policy setting specifies the maximum size of the log files. The user interfaces in the Local Group Policy Editor and the Microsoft Management Console (MMC) Event Viewer snap-in allow you to enter values as large as 2 terabytes. If this setting is not configured, event logs have a default maximum size of 20 megabytes.
- **Log Access**. This policy setting determines which user accounts have access to log files and what usage rights are granted.
- **Retain old events**. This policy setting controls event log behavior when the log file reaches its maximum size. When this policy setting is enabled and a log file reaches its maximum size, new events are not written to the log and are lost. When this policy setting is disabled and a log file reaches its maximum size, new events overwrite old events.
- **Backup log automatically when full**. This policy setting controls event log behavior when the log file reaches its maximum size and takes effect only if the **Retain old events** policy setting is enabled. If you enable these policy settings, the event log file is automatically closed and renamed when it is full. A new file is then started. If you disable or do not configure this policy setting and the **Retain old events** policy setting is enabled, new events are discarded and the old events are retained.

For more information about the potential vulnerabilities addressed by these settings and their potential impacts, see the **Threats and Countermeasures Guide: Security Settings in Windows Server 2008 and Windows Vista** or the Threats and Countermeasures Guide: Security Settings in Windows Server 2008 R2 and Windows 7. Also, security policy settings are described in the Security Policy Settings Reference.

In addition, a growing number of organizations are being required to store archived log files for a number of years. You should consult with regulatory compliance officers in your organization to determine whether such guidelines apply to your organization. For more information, see the IT Compliance Management Guide.

Deploying the security audit policy

Before deploying the audit policy in a production environment, it is critical that you determine the effects of the policy settings that you have configured.

The first step in assessing your audit policy deployment is to create a test environment in a lab and use it to simulate the various use scenarios that you have identified to confirm that the audit settings you have selected are configured correctly and generate the type of results you intend. For more information about setting up a lab environment for testing security audit policy, see the Advanced Security Auditing Walkthrough.

However, unless you are able to run fairly realistic simulations of network usage patterns, a lab setup cannot provide you with accurate information about the volume of audit data that the audit policy settings you selected will generate and how effective your plan for monitoring audit data will be. To provide this type of information, you need to conduct one or more pilot deployments. These pilot deployments could involve:

- A single OU that contains critical data servers or an OU that contains all desktop computers in a specified location.
- A limited set of security audit policy settings, such as **Logon/Logoff** and **Account Logon**.
- A combination of limited OUs and audit policy settings—for example, targeting servers in only the Accounting OU with **Object Access** policy settings.

After you have successfully completed one or more limited deployments, you should confirm that the audit data that is collected is manageable with your management tools and administrators. When you have confirmed that the pilot deployment is effective, you need to confirm that you have the necessary tools and staff to expand the deployment to include additional OUs and sets of audit policy settings until the production deployment is complete.

For more information about modeling and deploying Group Policy, see the **Group Policy Planning and Deployment Guide**.

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

This guide explains the process of setting up advanced security auditing capabilities that are made possible through settings and events that were introduced in Windows 8 and Windows Server 2012.

These procedures can be deployed with the advanced security auditing capabilities described in the following documents:

- Advanced Security Auditing Walkthrough
- Deploy Security Auditing with Central Audit Policies (Demonstration Steps)

Note

These auditing policies use the settings and events that were introduced in Windows 8 and Windows Server 2012. The contents of this guide apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

In this guide

Domain administrators can create and deploy expression-based security audit policies by using file classification information (resource attributes), user claims, and device claims to target specific users and resources to monitor potentially significant activities on one or more computers. These policies can be deployed centrally by using Group Policy, or directly on a computer, in a folder, or in individual files.

The procedures in this document describe how to:

- Monitor the Central Access Policies that Apply on a File Server
- Monitor the Central Access Policies Associated with Files and Folders
- Monitor the Resource Attributes on Files and Folders
- Monitor Claim Types
- Monitor User and Device Claims During Sign-in
- Monitor Central Access Policy and Rule Definitions
- Monitor Resource Attribute Definitions
- Monitor the Use of Removable Storage Devices

Important
This procedure can be configured on computers running any of the supported Windows operating systems. The other monitoring procedures can be configured only as part of a functioning dynamic access control deployment. If you have not yet deployed dynamic access control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Related resources

Security Auditing Overview

Monitor the Central Access Policies that Apply on a File Server

This topic for the IT professional describes how to monitor changes to the central access policies that apply to a file server when using advanced security auditing options to monitor dynamic

access control objects. Central access policies are created on a domain controller and then applied to file servers through Group Policy management.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to configure and verify security auditing settings that are used to monitor changes to the set of central access policies on a file server.The following procedures assume that you have configured and deployed dynamic access control, including central access policies, and claims in your network. If you have not yet deployed dynamic access control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor changes to central access policies

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the flexible access Group Policy Object, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, double-click **Policy Change**, and then double-click **Other Policy Change Events**.

   Note
   This policy setting monitors policy changes that might not be captured otherwise, such as central access policy changes or trusted platform module configuration changes.
5. Select the **Configure the following audit events** check box, select the **Success** check box (and the **Failure** check box, if desired), and then click **OK**.

After you modify the central access policies on the domain controller, verify that the changes have been applied to the file server and that the proper events are logged.

To verify changes to the central access policies

1. Sign in to your domain controller by using domain administrator credentials.
2. Open the Group Policy Management Console.
3. Right-click **Default domain policy**, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Policies**, and then double-click **Windows Settings**.
5. Double-click **Security Settings**, right-click **File system**, and then click **Manage CAPs**.
6. In the wizard that appears, follow the instructions to add a new central access policy

(CAP), and then click **OK**.

7. Use local administrator credentials to sign in to the server that hosts resources that are subject to the central access policies you changed.
8. Press the Windows key + R, then type **cmd** to open a Command Prompt window.

   Note
   If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
9. Type **gpupdate /force**, and press ENTER.
10. In Server Manager, click **Tools**, and then click **Event Viewer**.
11. Expand **Windows Logs**, and then click **Security**. Verify that event 4819 appears in the security log.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor the Use of Removable Storage Devices

This topic for the IT professional describes how to monitor attempts to use removable storage devices to access network resources. It describes how to use advanced security auditing options to monitor dynamic access control objects.

If you configure this policy setting, an audit event is generated each time a user attempts to copy, move, or save a resource to a removable storage device.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic. In these supported operating systems, administrators can set the Removable Storage Access policy to limit or deny users the ability to use removable storage devices. However, in earlier versions of the Windows and Windows Server operating systems, administrators could not track the use of removable storage devices.

Use the following procedures to monitor the use of removable storage devices and to verify that the devices are being monitored.

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor removable storage devices

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the flexible access Group Policy Object on the domain controller, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, double-click **Object Access**, and then double-click **Audit Removable Storage**.
5. Select the **Configure the following audit events** check box, select the **Success** check box (and the **Failure** check box, if desired), and then click **OK**.
6. If you selected the **Failure** check box, double-click **Audit Handle Manipulation**, select the **Configure the following audit events check box**, and then select **Failure**.
7. Click **OK**, and then close the Group Policy Management Editor.

After you configure the settings to monitor removable storage devices, use the following procedure to verify that the settings are active.

To verify that removable storage devices are monitored

1. Sign in to the computer that hosts the resources that you want to monitor. Press the Windows key + R, and then type **cmd** to open a Command Prompt window.

   Note
   If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. Type **gpupdate /force**, and press ENTER.
3. Connect a removable storage device to the targeted computer and attempt to copy a file that is protected with the Removable Storage Audit policy.
4. In Server Manager, click **Tools**, and then click **Event Viewer**.
5. Expand **Windows Logs**, and then click **Security**.
6. Look for event 4663, which logs successful attempts to write to or read from a removable storage device. Failures will log event 4656. Both events include **Task Category = Removable Storage device**.

   Key information to look for includes the name and account domain of the user who attempted to access the file, the object that the user is attempting to access, resource attributes of the resource, and the type of access that was attempted.

   Note
   We do not recommend that you enable this category on a file server that hosts file shares on a removable storage device. When Removable Storage Auditing is configured, any attempt to access the removable storage device will generate an audit event.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor Resource Attribute Definitions

This topic for the IT professional describes how to monitor changes to resource attribute definitions when you are using Advanced Security Auditing options to monitor dynamic access control objects.

Resource attribute definitions define the basic properties of resource attributes, such as what it means for a resource to be defined as “high business value.” Resource attribute definitions are stored in AD DS under the Resource Properties container. Changes to these definitions could significantly change the protections that govern a resource, even if the resource attributes that apply to the resource remain unchanged. Changes can be monitored like any other AD DS object.

For information about monitoring changes to the resource attributes that apply to files, see Monitor the Resource Attributes on Files and Folders.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to configure settings to monitor changes to resource attribute definitions in AD DS and to verify the changes. These procedures assume that you have configured and deployed Dynamic Access Control, including central access policies, claims, and other components, in your network. If you have not yet deployed Dynamic Access Control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor changes to resource attributes

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the Group Policy Object for the default domain controller, and then click **Edit**.
4. Double-click **Computer Configuration**, click **Security Settings**, expand **Advanced Audit Policy Configuration**, expand **System Audit Policies**, click **DS Access**, and then double-click **Audit directory service changes**.
5. Select the **Configure the following audit events** check box, select the **Success** check

box (and the **Failure** check box, if desired), and then click **OK**.

6. Close the Group Policy Management Editor.
7. Open the Active Directory Administrative Center.
8. Under **Dynamic Access Control**, right-click **Resource Properties**, and then click **Properties**.
9. Click the **Security** tab, click **Advanced** to open the **Advanced Security Settings** dialog box, and then click the **Auditing** tab.
10. Click **Add**, add a security auditing setting for the container, and then close all Security properties dialog boxes.

After you configure settings to monitor changes to resource attributes in AD DS, verify that the changes are being monitored.

To verify that changes to resource definitions are monitored

1. Sign in to your domain controller by using domain administrator credentials.
2. Open the Active Directory Administrative Center.
3. Under **Dynamic Access Control**, click **Resource Properties**, and then double-click a resource attribute.
4. Make changes to this resource attribute.
5. Click **OK**, and then close the Active Directory Administrative Center.
6. In Server Manager, click **Tools**, and then click **Event Viewer**.
7. Expand **Windows Logs**, and then click **Security**. Verify that event 5137 appears in the security log.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor Central Access Policy and Rule Definitions

This topic for the IT professional describes how to monitor changes to central access policy and central access rule definitions when you use Advanced Security Auditing options to monitor dynamic access control objects.

Central access policies and rules determine access permissions for multiple files on multiple file servers. Therefore, it is important to monitor changes to them. Like user claim and device claim definitions, central access policy and rule definitions reside in Active Directory Domain Services (AD DS), and they can be monitored just like any other object in Active Directory. Central access policies and rules are critical elements in a Dynamic Access Control deployment. These policies and rules are stored in AD DS, so they should be less likely to be tampered with than other

network objects. However, it is important to monitor these objects for potential changes in security auditing and to verify that policies are being enforced.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to configure settings to monitor changes to central access policy and central access rule definitions and to verify the changes. These procedures assume that you have configured and deployed Dynamic Access Control, including central access policies, claims, and other components, in your network. If you have not yet deployed Dynamic Access Control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor changes to central access policy and rule definitions

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the default domain controller Group Policy Object, and then click **Edit**.
4. Double-click **Computer Configuration**, click **Security Settings**, expand **Advanced Audit Policy Configuration**, expand **System Audit Policies**, click **DS Access**, and then double-click **Audit directory service changes**.
5. Select the **Configure the following audit events** check box, select the **Success** check box (and the **Failure** check box, if desired), and then click **OK**.
6. Close the Group Policy Management Editor.
7. Open the Active Directory Administrative Center.
8. Under Dynamic Access Control, right-click **Central Access Policies**, and then select **Properties**.
9. Click the **Security** tab, click **Advanced** to open the **Advanced Security Settings** dialog box, and then click the **Auditing** tab.
10. Click **Add**, add a security auditing setting for the container, and then close all Security properties dialog boxes.

After you configure settings to monitor changes to central access policy and central access rule definitions, verify that the changes are being monitored.

To verify that changes to central access policy and rule definitions are monitored

1. Sign in to your domain controller by using domain administrator credentials.

2. Open the Active Directory Administrative Center.
3. Under **Dynamic Access Control**, right-click **Central Access Policies**, and then click **Properties**.
4. Click the **Security** tab, click **Advanced** to open the **Advanced Security Settings** dialog box, and then click the **Auditing** tab.
5. Click **Add**, add a security auditing setting for the container, and then close all Security properties dialog boxes.
6. In the **Central Access Policies** container, add a new central access policy (or select one that exists), click **Properties** in the **Tasks** pane, and then change one or more attributes.
7. Click **OK**, and then close the Active Directory Administrative Center.
8. In Server Manager, click **Tools**, and then click **Event Viewer**.
9. Expand **Windows Logs**, and then click **Security**. Verify that event 4819 appears in the security log.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor User and Device Claims During Sign-in

This topic for the IT professional describes how to monitor user and device claims that are associated with a user's security token when you are using Advanced Security Auditing options to monitor dynamic access control objects.

Device claims are associated with the system that is used to access resources that are protected with Dynamic Access Control. User claims are attributes that are associated with a user. User claims and device claims are included in the user's security token used at sign-on. For example, information about Department, Company, Project, or Security clearances might be included in the token.

The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to monitor changes to user claims and device claims in the user's sign-on token and to verify the changes. These procedures assume that you have configured and deployed Dynamic Access Control, including central access policies, claims, and other components, in your network. If you have not yet deployed Dynamic Access Control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Note

Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To monitor user and device claims in user logon token

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the flexible access Group Policy Object, and then click **Edit**.
4. Double-click **Computer Configuration**, click **Security Settings**, expand **Advanced Audit Policy Configuration**, expand **System Audit Policies**, click **Logon/Logoff**, and then double-click **Audit User/Device claims**.
5. Select the **Configure the following audit events** check box, select the **Success** check box (and the **Failure** check box, if desired), and then click **OK**.
6. Close the Group Policy Management Editor.

After you configure settings to monitor user and device claims, verify that the changes are being monitored.

To verify that user and device claims in user logon token are monitored

1. With local administrator credentials, sign in to a file server that is subject to the flexible access Group Policy Object.
2. Open an elevated command prompt, and run the following command:

   Gpudate force
3. From a client computer, connect to a file share on the file server as a user who has access permissions to the file server.
4. On the file server, open Event Viewer, expand **Windows Logs**, and select the **Security** log. Look for event 4626, and confirm that it contains information about user claims and device claims.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor the Resource Attributes on Files and Folders

This topic for the IT professional describes how to monitor attempts to change settings to the resource attributes on files when you are using Advanced Security Auditing options to monitor dynamic access control objects.

The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

If your organization has a carefully thought out authorization configuration for resources, changes to these resource attributes can create potential security risks. Examples include:

- Changing files that have been marked as high business value to low business value.
- Changing the Retention attribute of files that have been marked for retention.
- Changing the Department attribute of files that are marked as belonging to a particular department.

Use the following procedures to configure settings to monitor changes to resource attributes on files and folders. These procedures assume that have configured and deployed central access policies in your network. For more information about how to configure and deploy central access policies, see Dynamic Access Control: Scenario Overview .

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To monitor changes to resource attributes on files

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the flexible access Group Policy Object, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, double-click **Policy Change**, and then double-click **Audit Authorization Policy Change**.
5. Select the **Configure the following audit events** check box, select the **Success** and **Failure** check boxes, and then click **OK**.

After you configure settings to monitor resource attributes on files, verify that the changes are being monitored.

To verify that changes to resource attributes on files are monitored

1. Use administrator credentials to sign in to the server that hosts the resource you want to monitor. Press the Windows key + R, and then type **cmd** to open a Command Prompt window.

   Note
   If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
2. Type **gpupdate /force**, and press ENTER.
3. Attempt to change resource properties on one or more files and folders.

4. In Server Manager, click **Tools**, and then click **Event Viewer**.
5. Expand **Windows Logs**, and then click **Security**.
6. Depending on which resource attributes you attempted to change, you should look for the following events:
   - Event 4911, which tracks changes to file attributes
   - Event 4913, which tracks changes to central access policies

   Key information to look for includes the name and account domain of the principal attempting to change the resource attribute, the object that the principal is attempting to modify, and information about the changes that are being attempted.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor the Central Access Policies Associated with Files and Folders

This topic for the IT professional describes how to monitor changes to the central access policies that are associated with files and folders when you are using Advanced Security Auditing options to monitor dynamic access control objects.

This security audit policy and the event that it records are generated when the central access policy that is associated with a file or folder is changed. This security audit policy is useful when an administrator wants to monitor potential changes on some, but not all, files and folders on a file server.

For information about monitoring potential central access policy changes for an entire file server, see Monitor the Central Access Policies that Apply on a File Server.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to configure settings to monitor central access policies that are associated with files. These procedures assume that you have configured and deployed Dynamic Access Control in your network. For more information about how to configure and deploy Dynamic Access Control, see Dynamic Access Control: Scenario Overview.

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor central access policies associated with files or folders

1. Sign in to your domain controller by using domain administrator credentials.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the flexible access Group Policy Object, and then click **Edit**.
4. Double-click **Computer Configuration**, double-click **Security Settings**, double-click **Advanced Audit Policy Configuration**, double-click **Policy Change**, and then double-click **Audit Authorization Policy Change**.
5. Select the **Configure the following audit events** check box, select the **Success** check box (and the **Failure** check box, if desired), and then click **OK**.
6. Enable auditing for a file or folder as described in the following procedure.

To enable auditing for a file or folder

1. Sign in as a member of the local administrators group on the computer that contains the files or folders that you want to audit.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Auditing** tab.

   If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
4. Click **Add**, type a user name or group name in the format **contoso\user1**, and then click **OK**.
5. In the **Auditing Entries for** dialog box, select the permissions that you want to audit, such as **Full Control** or **Delete**.
6. Click **OK** four times to complete the configuration of the object SACL.
7. Open a File Explorer window and select or create a file or folder to audit.
8. Open an elevated command prompt, and run the following command:

   gpupdate /force

After you configure settings to monitor changes to the central access policies that are associated with files and folders, verify that the changes are being monitored.

To verify that changes to central access policies associated with files and folders are monitored

1. Sign in as a member of the local administrators group on the computer that contains the files or folders that you want to audit.
2. Open a File Explorer window and select the file or folder that you configured for auditing in the previous procedure.
3. Right-click the file or folder, click **Properties**, click the **Security** tab, and then click **Advanced**.
4. Click the **Central Policy** tab, click **Change**, and select a different central access policy (if one is available) or select **No Central Access Policy**, and then click **OK** twice.

   Note

You must select a setting that is different than your original setting to generate the audit event.

5. In Server Manager, click **Tools**, and then click **Event Viewer**.
6. Expand **Windows Logs**, and then click **Security**.
7. Look for event 4913, which is generated when the central access policy that is associated with a file or folder is changed. This event includes the security identifiers (SIDs) of the old and new central access policies.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Monitor Claim Types

This topic for the IT professional describes how to monitor changes to claim types that are associated with dynamic access control when you are using Advanced Security Auditing options.

Claim types are one of the basic building blocks of Dynamic Access Control. Claim types can include attributes such as the departments in an organization or the levels of security clearance that apply to classes of users. You can use security auditing to track whether claims are added, modified, enabled, disabled, or deleted.

Note
The contents of this topic apply to the list of supported Windows operating systems designated in the **Applies To** list at the beginning of this topic.

Use the following procedures to configure settings to monitor changes to claim types in AD DS. These procedures assume that you have configured and deployed Dynamic Access Control, including central access policies, claims, and other components, in your network. If you have not yet deployed Dynamic Access Control in your network, see Deploy a Central Access Policy (Demonstration Steps).

Note
Your server might function differently based on the version and edition of the operating system that is installed, your account permissions, and your menu settings.

To configure settings to monitor changes to claim types

1. Sign in to your domain controller by using domain administrator credential.
2. In Server Manager, point to **Tools**, and then click **Group Policy Management**.
3. In the console tree, right-click the default domain controller Group Policy Object, and then

click **Edit**.

4. Double-click **Computer Configuration**, click **Security Settings**, expand **Advanced Audit Policy Configuration**, expand **System Audit Policies**, click **DS Access**, and then double-click **Audit directory service changes**.
5. Select the **Configure the following audit events** check box, select the **Success** check box (andthe **Failure** check box, if desired), and then click **OK**.

After you configure settings to monitor changes to claim types in AD DS, verify that the changes are being monitored.

To verify that changes to claim types are monitored

1. Sign in to your domain controller by using domain administrator credentials.
2. Open the Active Directory Administrative Center.
3. Under **Dynamic Access Control**, right-click **Claim Types**, and then click **Properties**.
4. Click the **Security** tab, click **Advanced** to open the **Advanced Security Settings** dialog box, and then click the **Auditing** tab.
5. Click **Add**, add a security auditing setting for the container, and then close all the Security properties dialog boxes.
6. In the **Claim Types** container, add a new claim type or select an existing claim type. In the **Tasks** pane, click **Properties**, and then change one or more attributes.

   Click **OK**, and then close the Active Directory Administrative Center.
7. Open Event Viewer on this domain controller, expand **Windows Logs**, and select the **Security** log.

   Look for event 5137. Key information to look for includes the name of the new attribute that was added, the type of claim that was created, and the user who created the claim.

Related resource

Using Advanced Security Auditing Options to Monitor Dynamic Access Control Objects

Advanced Security Audit Policy Settings

This reference for the IT professional provides information about the Advanced Audit policy settings that are available in Windows operating systems and the audit events that they generate.

The 53 security audit policy settings under **Security Settings\Advanced Audit Policy Configuration** can help your organization audit compliance with important business-related and security-related rules by tracking precisely defined activities, such as:

- A group administrator has modified settings or data on servers that contain finance information.
- An employee within a defined group has accessed an important file.

- The correct system access control list (SACL) is applied to every file and folder or registry key on a computer or file share as a verifiable safeguard against undetected access.

You can access these audit policy settings through the Local Security Policy snap-in (secpol.msc) on the local computer or by using Group Policy.

These Advanced Audit policy settings allow you to select only the behaviors that you want to monitor. You can exclude audit results for behaviors that are of little or no concern to you, or behaviors that create an excessive number of log entries. In addition, because security audit policies can be applied by using domain Group Policy Objects, audit policy settings can be modified, tested, and deployed to selected users and groups with relative simplicity.

When Advanced Security Audit policy settings are configured, events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Audit policy settings under **Security Settings\Advanced Audit Policy Configuration** are available in the following categories:

Account Logon

Configuring policy settings in this category can help you document attempts to authenticate account data on a domain controller or on a local Security Accounts Manager (SAM). Unlike Logon and Logoff policy settings and events, which track attempts to access a particular computer, settings and events in this category focus on the account database that is used. This category includes the following subcategories:

- Audit Credential Validation
- Audit Kerberos Authentication Service
- Audit Kerberos Service Ticket Operations
- Audit Other Logon/Logoff Events

Account Management

The security audit policy settings in this category can be used to monitor changes to user and computer accounts and groups. This category includes the following subcategories:

- Audit Application Group Management
- Audit Computer Account Management
- Audit Distribution Group Management
- Audit Other Account Management Events
- Audit Security Group Management
- Audit User Account Management

Detailed Tracking

Detailed Tracking security policy settings and audit events can be used to monitor the activities of individual applications and users on that computer, and to understand how a computer is being used. This category includes the following subcategories:

- Audit DPAPI Activity
- Audit Process Creation

- Audit Process Termination
- Audit RPC Events

DS Access

DS Access security audit policy settings provide a detailed audit trail of attempts to access and modify objects in Active Directory Domain Services (AD DS). These audit events are logged only on domain controllers. This category includes the following subcategories:

- Audit Detailed Directory Service Replication
- Audit Directory Service Access
- Audit Directory Service Changes
- Audit Directory Service Replication

Logon/Logoff

Logon/Logoff security policy settings and audit events allow you to track attempts to log on to a computer interactively or over a network. These events are particularly useful for tracking user activity and identifying potential attacks on network resources. This category includes the following subcategories:

- Audit Account Lockout
- Audit IPsec Extended Mode
- Audit IPsec Main Mode
- Audit IPsec Quick Mode
- Audit Logoff
- Audit Logon
- Audit Network Policy Server
- Audit Other Logon/Logoff Events
- Audit Special Logon

Object Access

Object Access policy settings and audit events allow you to track attempts to access specific objects or types of objects on a network or computer. To audit attempts to access a file, directory, registry key, or any other object, you must enable the appropriate Object Access auditing subcategory for success and/or failure events. For example, the File System subcategory needs to be enabled to audit file operations, and the Registry subcategory needs to be enabled to audit registry accesses.

Proving that these audit policies are in effect to an external auditor is more difficult. There is no easy way to verify that the proper SACLs are set on all inherited objects. To address this issue, see Global Object Access Auditing.

This category includes the following subcategories:

- Audit Application Generated
- Audit Certification Services
- Audit Detailed File Share
- Audit File Share

- Audit File System
- Audit Filtering Platform Connection
- Audit Filtering Platform Packet Drop
- Audit Handle Manipulation
- Audit Kernel Object
- Audit Other Object Access Events
- Audit Registry
- Audit SAM

Policy Change

Policy Change audit events allow you to track changes to important security policies on a local system or network. Because policies are typically established by administrators to help secure network resources, monitoring changes or attempts to change these policies can be an important aspect of security management for a network. This category includes the following subcategories:

- Audit Audit Policy Change
- Audit Authentication Policy Change
- Audit Authorization Policy Change
- Audit Filtering Platform Policy Change
- Audit MPSSVC Rule-Level Policy Change
- Audit Other Policy Change Events

Privilege Use

Permissions on a network are granted for users or computers to complete defined tasks. Privilege Use security policy settings and audit events allow you to track the use of certain permissions on one or more systems. This category includes the following subcategories:

- Audit Non-Sensitive Privilege Use
- Audit Sensitive Privilege Use
- Audit Other Privilege Use Events

System

System security policy settings and audit events allow you to track system-level changes to a computer that are not included in other categories and that have potential security implications. This category includes the following subcategories:

- Audit IPsec Driver
- Audit Other System Events
- Audit Security State Change
- Audit Security System Extension
- Audit System Integrity

Global Object Access

Global Object Access Auditing policy settings allow administrators to define computer system access control lists (SACLs) per object type for the file system or for the registry. The specified SACL is then automatically applied to every object of that type.

Auditors will be able to prove that every resource in the system is protected by an audit policy by viewing the contents of the Global Object Access Auditing policy settings. For example, if auditors see a policy setting called "Track all changes made by group administrators," they know that this policy is in effect.

Resource SACLs are also useful for diagnostic scenarios. For example, setting the Global Object Access Auditing policy to log all the activity for a specific user and enabling the policy to track "Access denied" events for the file system or registry can help administrators quickly identify which object in a system is denying a user access.

Note

If a file or folder SACL and a Global Object Access Auditing policy setting (or a single registry setting SACL and a Global Object Access Auditing policy setting) are configured on a computer, the effective SACL is derived from combining the file or folder SACL and the Global Object Access Auditing policy. This means that an audit event is generated if an activity matches the file or folder SACL or the Global Object Access Auditing policy.

This category includes the following subcategories:

- File System (Global Object Access Auditing)
- Registry (Global Object Access Auditing)

Audit Credential Validation

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Credential Validation**, which determines whether the operating system generates audit events on credentials that are submitted for a user account logon request.

These events occur on the computer that is authoritative for the credentials as follows:

- For domain accounts, the domain controller is authoritative.
- For local accounts, the local computer is authoritative.

Event volume: High on domain controllers

Because domain accounts are used much more frequently than local accounts in enterprise environments, most of the Account Logon events in a domain environment occur on the domain controllers that are authoritative for the domain accounts. However, these events can occur on any computer, and they may occur in conjunction with or on separate computers from Logon and Logoff events.

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4774	An account was mapped for logon.
4775	An account could not be mapped for logon.
4776	The domain controller attempted to validate the credentials for an account.
4777	The domain controller failed to validate the credentials for an account.

Related resource

Advanced Security Audit Policy Settings

Audit Kerberos Authentication Service

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Kerberos Authentication Service**, which determines whether to generate audit events for Kerberos authentication ticket-granting ticket (TGT) requests.

If you configure this policy setting, an audit event is generated after a Kerberos authentication TGT request. Success audits record successful attempts and Failure audits record unsuccessful attempts.

Event volume: High on Kerberos Key Distribution Center servers

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4768	A Kerberos authentication ticket (TGT) was

Event ID	Event message
	requested.
4771	Kerberos preauthentication failed.
4772	A Kerberos authentication ticket request failed.

Related resource

Advanced Security Audit Policy Settings

Audit Kerberos Service Ticket Operations

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Kerberos Service Ticket Operations**, which determines whether the operating system generates security audit events for Kerberos service ticket requests.

Events are generated every time Kerberos is used to authenticate a user who wants to access a protected network resource. Kerberos service ticket operation audit events can be used to track user activity.

Event volume:

- High on a domain controller that is in a Key Distribution Center (KDC)
- Low on domain members

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4769	A Kerberos service ticket was requested.
4770	A Kerberos service ticket was renewed.

Related resource

Advanced Security Audit Policy Settings

Audit Other Account Logon Events

This topic for the IT professional describes the advanced security audit policy setting, **Audit Other Account Logon Events**, which allows you to audit events generated by responses to credential requests submitted for a user account logon that are not credential validation or Kerberos tickets.

Examples can include the following:

- Remote Desktop session disconnections
- New Remote Desktop sessions
- Locking and unlocking a workstation
- Invoking a screen saver
- Dismissing a screen saver
- Detection of a Kerberos replay attack, in which a Kerberos request with identical information was received twice

  Note
  This condition could be caused by a network misconfiguration.
- Access to a wireless network granted to a user or computer account
- Access to a wired 802.1x network granted to a user or computer account

Event volume: Varies, depending on system use

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of Windows as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4649	A replay attack was detected.
4778	A session was reconnected to a Window Station.
4779	A session was disconnected from a Window Station.
4800	The workstation was locked.
4801	The workstation was unlocked.
4802	The screen saver was invoked.
4803	The screen saver was dismissed.
5378	The requested credentials delegation was disallowed by policy.

Event ID	Event message
5632	A request was made to authenticate to a wireless network.
5633	A request was made to authenticate to a wired network.

Related resource

Advanced Security Audit Policy Settings

Audit Application Group Management

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Application Group Management**, which determines whether the operating system generates audit events when application group management tasks are performed.

Application group management tasks include:

- An application group is created, changed, or deleted.
- A member is added to or removed from an application group.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4783	A basic application group was created.
4784	A basic application group was changed.
4785	A member was added to a basic application group.
4786	A member was removed from a basic application group.

Event ID	Event message
4787	A non-member was added to a basic application group.
4788	A non-member was removed from a basic application group.
4789	A basic application group was deleted.
4790	An LDAP query group was created.

Related resource

Advanced Security Audit Policy Settings

Audit Computer Account Management

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Computer Account Management**, which determines whether the operating system generates audit events when a computer account is created, changed, or deleted.

This policy setting is useful for tracking account-related changes to computers that are members of a domain.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4741	A computer account was created.
4742	A computer account was changed.
4743	A computer account was deleted.

Related resource

Advanced Security Audit Policy Settings

Audit Distribution Group Management

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Distribution Group Management**, which determines whether the operating system generates audit events for specific distribution-group management tasks.

Tasks for distribution-group management that can be audited include:

- A distribution group is created, changed, or deleted.
- A member is added to or removed from a distribution group.

This subcategory to which this policy belongs is logged only on domain controllers.

Note
Distribution groups cannot be used to manage access control permissions.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008.

Event ID	Event message
4744	A security-disabled local group was created.
4745	A security-disabled local group was changed.
4746	A member was added to a security-disabled local group.
4747	A member was removed from a security-disabled local group.
4748	A security-disabled local group was deleted.
4749	A security-disabled global group was created.
4750	A security-disabled global group was changed.
4751	A member was added to a security-disabled global group.
4752	A member was removed from a security-disabled global group.

Event ID	Event message
4753	A security-disabled global group was deleted.
4759	A security-disabled universal group was created.
4760	A security-disabled universal group was changed.
4761	A member was added to a security-disabled universal group.
4762	A member was removed from a security-disabled universal group.

Related resource

Advanced Security Audit Policy Settings

Audit Other Account Management Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Other Account Management Events**, which determines whether the operating system generates user account management audit events.

Events can be generated for user account management auditing when:

- The password hash of an account is accessed. This typically happens when the Active Directory Migration Tool (ADMT) is moving password data.
- The Password Policy Checking application programming interface (API) is called. Calls to this function could be part of an attack from a malicious application that is testing whether password complexity policy settings are being applied.
- Changes are made to domain policy under **Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy** or **Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy**.

Note
These events are logged when the domain policy is applied (on refresh or restart), not when settings are modified by an administrator.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event Message Summary
4782	The password hash for an account was accessed.
4793	The Password Policy Checking API was called.

Related resource

Advanced Security Audit Policy Settings

Audit Security Group Management

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Security Group Management**, which determines whether the operating system generates audit events when specific security group management tasks are performed.

Tasks for security group management include:

- A security group is created, changed, or deleted.
- A member is added to or removed from a security group.
- A group's type is changed.

Security groups can be used for access control permissions and also as distribution lists.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4727	A security-enabled global group was created.
4728	A member was added to a security-enabled global group.
4729	A member was removed from a security-enabled global group.
4730	A security-enabled global group was deleted.

Event ID	Event message
4731	A security-enabled local group was created.
4732	A member was added to a security-enabled local group.
4733	A member was removed from a security-enabled local group.
4734	A security-enabled local group was deleted.
4735	A security-enabled local group was changed.
4737	A security-enabled global group was changed.
4754	A security-enabled universal group was created.
4755	A security-enabled universal group was changed.
4756	A member was added to a security-enabled universal group.
4757	A member was removed from a security-enabled universal group.
4758	A security-enabled universal group was deleted.
4764	A group's type was changed.

Related resource

Advanced Security Audit Policy Settings

Audit User Account Management

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit User Account Management**, which determines whether the operating system generates audit events when specific user account management tasks are performed.

Tasks that are audited for user account management include:

- A user account is created, changed, deleted, renamed, disabled, enabled, locked out, or unlocked.
- A user account password is set or changed.

- Security identifier (SID) history is added to a user account.
- The Directory Services Restore Mode password is set.
- Permissions are changed on accounts that are members of administrator groups.
- Credential Manager credentials are backed up or restored.

This policy setting is essential for tracking events that involve provisioning and managing user accounts.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4720	A user account was created.
4722	A user account was enabled.
4723	An attempt was made to change an account's password.
4724	An attempt was made to reset an account's password.
4725	A user account was disabled.
4726	A user account was deleted.
4738	A user account was changed.
4740	A user account was locked out.
4765	SID History was added to an account.
4766	An attempt to add SID History to an account failed.
4767	A user account was unlocked.
4780	The ACL was set on accounts which are members of administrators groups.
4781	The name of an account was changed:
4794	An attempt was made to set the Directory Services Restore Mode.
5376	Credential Manager credentials were backed up.

Event ID	Event message
5377	Credential Manager credentials were restored from a backup.

Related resource

Advanced Security Audit Policy Settings

Audit DPAPI Activity

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit DPAPI Activity**, which determines whether the operating system generates audit events when encryption or decryption calls are made into the data protection application interface (DPAPI).

DPAPI is used to protect secret information such as stored passwords and key information. For more information about DPAPI, see Windows Data Protection (http://go.microsoft.com/fwlink/?LinkID=121720).

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4692	Backup of data protection master key was attempted.
4693	Recovery of data protection master key was attempted.
4694	Protection of auditable protected data was attempted.
4695	Unprotection of auditable protected data was attempted.

Related resource

Advanced Security Audit Policy Settings

Audit Process Creation

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Process Creation**, which determines whether the operating system generates audit events when a process is created (starts).

These audit events can help you track user activity and understand how a computer is being used. Information includes the name of the program or the user that created the process.

Event volume: Low to medium, depending on system usage

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4688	A new process has been created.
4696	A primary token was assigned to a process.

Related resource

Advanced Security Audit Policy Settings

Audit Process Termination

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Process Termination**, which determines whether the operating system generates audit events when an attempt is made to end a process.

Success audits record successful attempts and Failure audits record unsuccessful attempts.

If you do not configure this policy setting, no audit event is generated when a process ends.

This policy setting can help you track user activity and understand how the computer is used.

Event volume: Varies, depending on how the computer is used

Default: Not configured

If this policy setting is configured, the following event appears on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4689	A process has exited.

Related resource

Advanced Security Audit Policy Settings

Audit RPC Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit RPC Events**, which determines whether the operating system generates audit events when inbound remote procedure call (RPC) connections are made.

RPC is a technology for creating distributed client/server programs. RPC is an interprocess communication technique that enables client and server software to communicate. For more information, see What Is RPC?.

Event volume: High on RPC servers

Default: Not configured

If this policy setting is configured, the following event appears on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
5712	A Remote Procedure Call (RPC) was attempted.

Related resource

Advanced Security Audit Policy Settings

Audit Detailed Directory Service Replication

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Detailed Directory Service Replication**, which determines whether the operating system generates audit events that contain detailed tracking information about data that is replicated between domain controllers.

This audit subcategory can be useful to diagnose replication issues.

Event volume: These events can create a very high volume of event data.

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008.

Event ID	Event message
4928	An Active Directory replica source naming context was established.
4929	An Active Directory replica source naming context was removed.
4930	An Active Directory replica source naming context was modified.
4931	An Active Directory replica destination naming context was modified.
4934	Attributes of an Active Directory object were replicated.
4935	Replication failure begins.
4936	Replication failure ends.
4937	A lingering object was removed from a replica.

Related resource

Advanced Security Audit Policy Settings

Audit Directory Service Access

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Directory Service Access**, which determines whether the operating system generates audit events when an Active Directory Domain Services (AD DS) object is accessed.

These events are similar to the Directory Service Access events in previous versions of Windows Server operating systems.

Important
Audit events are generated only on objects with configured system access control lists (SACLs), and only when they are accessed in a manner that matches the SACL settings.

Event volume: High on servers running AD DS role services; none on client computers

Default: Not configured

If this policy setting is configured, the following event appears on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008.

Event ID	Event message
4662	An operation was performed on an object.

Related resource

Advanced Security Audit Policy Settings

Audit Directory Service Changes

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Directory Service Changes**, which determines whether the operating system generates audit events when changes are made to objects in Active Directory Domain Services (AD DS).

The types of changes that are reported are:

- Create
- Delete
- Modify
- Move
- Undelete

Directory Service Changes auditing, where appropriate, indicates the old and new values of the changed properties of the objects that were changed.

Important
Audit events are generated only for objects with configured system access control lists (SACLs), and only when they are accessed in a manner that matches their SACL settings. Some objects and properties do not cause audit events to be generated due to settings on the object class in the schema.

This subcategory only logs events on domain controllers. Changes to Active Directory objects are important events to track in order to understand the state of the network policy.

Event volume: High on domain controllers; none on client computers

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008.

Event ID	Event message
5136	A directory service object was modified.
5137	A directory service object was created.
5138	A directory service object was undeleted.
5139	A directory service object was moved.
5141	A directory service object was deleted.

Related resource

Advanced Security Audit Policy Settings

Audit Directory Service Replication

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Directory Service Replication**, which determines whether the operating system generates audit events when replication between two domain controllers begins and ends.

Event volume: Medium on domain controllers; none on client computers

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008.

Event ID	Event message
4932	Synchronization of a replica of an Active Directory naming context has begun.
4933	Synchronization of a replica of an Active Directory naming context has ended.

Related resource

Advanced Security Audit Policy Settings

Audit Account Lockout

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Account Lockout**, which enables you to audit security events that are generated by a failed attempt to log on to an account that is locked out.

If you configure this policy setting, an audit event is generated when an account cannot log on to a computer because the account is locked out. Success audits record successful attempts and failure audits record unsuccessful attempts.

Account lockout events are essential for understanding user activity and detecting potential attacks.

Event volume: Low

Default setting: Success

If this policy setting is configured, the following event appears on computers that run the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4625	An account failed to log on.

Related resource

Advanced Security Audit Policy Settings

Audit IPsec Extended Mode

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit IPsec Extended Mode**, which determines whether the operating system generates audit events for the results of the Internet Key Exchange (IKE) protocol and Authenticated Internet Protocol (AuthIP) during Extended Mode negotiations.

IKE is an Internet standard, defined in RFC 2409, that defines a mechanism to establish IPsec security associations (SAs). An SA is a combination of a mutually agreeable policy and keys that define the security services and mechanisms that help protect communication between IPsec peers.

AuthIP is an enhanced version of IKE that offers additional flexibility with support for user-based authentication, authentication with multiple credentials, improved authentication method negotiation, and asymmetric authentication. Like IKE, AuthIP supports main-mode and quick-mode negotiation. AuthIP also supports Extended Mode, a part of IPsec peer negotiation during which a second round of authentication can be performed. Extended Mode, which is optional, can

be used for multiple authentications. For example, with extended mode you can perform separate computer-based and user-based authentications.

Event volume: High

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4978	During Extended Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
4979	IPsec Main Mode and Extended Mode security associations were established. **Note** This event provides event data in the following categories: Main Mode Local Endpoint, Main Mode Remote Endpoint, Main Mode Cryptographic Information, Main Mode Security Association, Main Mode Additional Information, and Extended Mode Information.
4980	IPsec Main Mode and Extended Mode security associations were established. **Note** This event provides event audit data in the following categories: Main Mode Local Endpoint, Main Mode Remote Endpoint. Main Mode Cryptographic Information, Main Mode Security Association, Main Mode Additional Information, Extended Mode Local Endpoint, Extended Mode Remote Endpoint, and Extended Mode Additional Information:
4981	IPsec Main Mode and Extended Mode security

Event ID	Event message
	associations were established. **Note** This event provides event audit data in the following categories: Local Endpoint, Local Certificate, Remote Endpoint, Remote Certificate, Cryptographic Information, Security Association Information, Additional Information, and Extended Mode Information.
4982	IPsec Main Mode and Extended Mode security associations were established. **Note** This event provides event audit data in the following categories: Local Endpoint, Local Certificate, Remote Endpoint, Remote Certificate, Cryptographic Information, Security Association Information, Additional Information, Extended Mode Local Endpoint, Extended Mode Remote Endpoint, and Extended Mode Additional Information.
4983	An IPsec Extended Mode negotiation failed. The corresponding Main Mode security association has been deleted. **Note** This event provides event audit data in the following categories: Local Endpoint, Local Certificate, Remote Endpoint, Remote Certificate, and Failure Information.
4984	An IPsec Extended Mode negotiation failed. The corresponding Main Mode security association has been deleted. **Note** This event provides event audit data in

Event ID	Event message
	the following categories: Local Endpoint, Remote Endpoint, Additional Information, and Failure Information.

Related resource

Advanced Security Audit Policy Settings

Audit IPsec Main Mode

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit IPsec Main Mode**, which determines whether the operating system generates events for the results of the Internet Key Exchange (IKE) protocol and Authenticated Internet Protocol (AuthIP) during Main Mode negotiations.

IKE is an Internet standard, defined in RFC 2409, that defines a mechanism to establish IPsec security associations (SAs). An SA is a combination of a mutually agreeable policy and keys that define the security services and mechanisms that help protect communication between IPsec peers.

AuthIP is an enhanced version of IKE that offers additional flexibility with support for user-based authentication, authentication with multiple credentials, improved authentication method negotiation, and asymmetric authentication. Like IKE, AuthIP supports Main Mode and Quick Mode negotiation.

Main Mode Internet Key Exchange (IKE) negotiation establishes a secure channel, known as the Internet Security Association and Key Management Protocol (ISAKMP) security association (SA), between two computers. To establish the secure channel, Main Mode negotiation determines a set of cryptographic protection suites, exchanges keying material to establish the shared secret key, and authenticates computer identities.

Event volume: High

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4646	Security ID: %1
4650	An IPsec Main Mode security association was established. Extended Mode was not enabled.

Event ID	Event message
	Certificate authentication was not used.
4651	An IPsec Main Mode security association was established. Extended Mode was not enabled. A certificate was used for authentication.
4652	An IPsec Main Mode negotiation failed. **Note** This audit event returns detailed audit data in the following categories: Local Endpoint, Local Certificate, Remote Endpoint, Remote Certificate, Additional Information, and Failure Information.
4653	An IPsec Main Mode negotiation failed. **Note** This audit event returns detailed audit data in the following categories: Local Endpoint, Remote Endpoint, Additional Information, and Failure Information.
4655	An IPsec Main Mode security association ended.
4976	During Main Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
5049	An IPsec Security Association was deleted.
5453	An IPsec negotiation with a remote computer failed because the IKE and AuthIP IPsec Keying Modules (IKEEXT) service is not started.

Related resource

Advanced Security Audit Policy Settings

Audit IPsec Quick Mode

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit IPsec Quick Mode**, which determines whether the operating system generates audit events for the results of the Internet Key Exchange (IKE) protocol and Authenticated Internet Protocol (AuthIP) during Quick Mode negotiations.

IKE is an Internet standard, defined in RFC 2409, that defines a mechanism to establish IPsec security associations (SAs). An SA is a combination of a mutually agreeable policy and keys that define the security services and mechanisms that help protect communication between IPsec peers.

AuthIP is an enhanced version of IKE that offers additional flexibility with support for user-based authentication, authentication with multiple credentials, improved authentication method negotiation, and asymmetric authentication. Like IKE, AuthIP supports Main Mode and Quick Mode negotiation.

Quick Mode (also known as Phase 2) IKE negotiation establishes a secure channel between two computers to protect data. Because this phase involves the establishment of security associations (SAs) that are negotiated on behalf of the IPsec service, the SAs that are created during Quick Mode are called the IPsec SAs. During Quick Mode, keying material is refreshed or, if necessary, new keys are generated. A protection suite that protects specified IP traffic is also selected. A protection suite is a defined set of data integrity or data encryption settings. Quick Mode is not considered a complete exchange because it is dependent on a Main Mode exchange.

Event volume: High

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4977	During Quick Mode negotiation, IPsec received an invalid negotiation packet. If this problem persists, it could indicate a network issue or an attempt to modify or replay this negotiation.
5451	An IPsec Quick Mode security association was established.
5452	An IPsec Quick Mode security association ended.

Related resource

Advanced Security Audit Policy Settings

Audit Logoff

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Logoff**, which determines whether the operating system generates audit events when logon sessions are terminated.

These events occur on the computer that was accessed. In the case of an interactive logon, these events are generated on the computer that was logged on to.

Note
There is no failure event in this subcategory because failed logoffs (such as when a system abruptly shuts down) do not generate an audit record.

Logon events are essential to understanding user activity and detecting potential attacks. Logoff events are not 100 percent reliable. For example, the computer can be turned off without a proper logoff and shutdown; in this case, a logoff event is not generated.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4634	An account was logged off.
4647	User initiated logoff.

Related resource

Advanced Security Audit Policy Settings

Audit Logon

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Logon**, which determines whether the operating system generates audit events when a user attempts to log on to a computer.

These events are related to the creation of logon sessions and occur on the computer that was accessed. For an interactive logon, events are generated on the computer that was logged on to. For a network logon, such as accessing a share, events are generated on the computer that hosts the resource that was accessed.

The following events are recorded:

- Logon success and failure.
- Logon attempts by using explicit credentials. This event is generated when a process attempts to log on an account by explicitly specifying that account's credentials. This most commonly occurs in batch configurations such as scheduled tasks, or when using the Runas command.
- Security identifiers (SIDs) are filtered.

Logon events are essential to tracking user activity and detecting potential attacks.

Event volume: Low on a client computer; medium on a domain controller or network server

Default: Success for client computers; success and failure for servers

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4624	An account was successfully logged on.
4625	An account failed to log on.
4648	A logon was attempted using explicit credentials.
4675	SIDs were filtered.

Related resource

Advanced Security Audit Policy Settings

Audit Network Policy Server

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Network Policy Server**, which determines whether the operating system generates audit events for RADIUS (IAS) and Network Access Protection (NAP) activity on user access requests (Grant, Deny, Discard, Quarantine, Lock, and Unlock).

NAP events can be used to help understand the overall health of the network.

Event volume: Medium to high on servers that are running Network Policy Server (NPS); moderate on other servers or on client computers

Default: Success and failure

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008, Windows Vista with Service Pack 2 (SP2), or Windows Vista with Service Pack 1 (SP1).

Event ID	Event message
6272	Network Policy Server granted access to a user.
6273	Network Policy Server denied access to a user.
6274	Network Policy Server discarded the request for a user.
6275	Network Policy Server discarded the accounting request for a user.
6276	Network Policy Server quarantined a user.
6277	Network Policy Server granted access to a user but put it on probation because the host did not meet the defined health policy.
6278	Network Policy Server granted full access to a user because the host met the defined health policy.
6279	Network Policy Server locked the user account due to repeated failed authentication attempts.
6280	Network Policy Server unlocked the user account.

Related resource

Advanced Security Audit Policy Settings

Audit Other Logon/Logoff Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Other Logon/Logoff Events**, which determines whether Windows generates audit events for other logon or logoff events.

These other logon or logoff events include:

- A Remote Desktop session connects or disconnects.
- A workstation is locked or unlocked.
- A screen saver is invoked or dismissed.
- A replay attack is detected. This event indicates that a Kerberos request was received twice with identical information. This condition could also be caused by network misconfiguration.
- A user is granted access to a wireless network. It can either be a user account or the computer account.
- A user is granted access to a wired 802.1x network. It can either be a user account or the computer account.

Logon events are essential to understanding user activity and detecting potential attacks.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of Windows as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4649	A replay attack was detected.
4778	A session was reconnected to a Window Station.
4779	A session was disconnected from a Window Station.
4800	The workstation was locked.
4801	The workstation was unlocked.
4802	The screen saver was invoked.
4803	The screen saver was dismissed.
5378	The requested credentials delegation was disallowed by policy.
5632	A request was made to authenticate to a wireless network.

Event ID	Event message
5633	A request was made to authenticate to a wired network.

Related resource

Advanced Security Audit Policy Settings

Audit Special Logon

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Special Logon**, which determines whether the operating system generates audit events under special sign on (or log on) circumstances.

This security policy setting determines whether the operating system generates audit events when:

- A special logon is used. A special logon is a logon that has administrator-equivalent privileges and can be used to elevate a process to a higher level.
- A member of a special group logs on. Special Groups is a Windows feature that enables the administrator to find out when a member of a certain group has logged on. The administrator can set a list of group security identifiers (SIDs) in the registry. If any of these SIDs is added to a token during logon and this auditing subcategory is enabled, a security event is logged. For more information about this feature, see article 947223 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkID=120183).

Users holding special privileges can potentially make changes to the system. We recommend that you track their activity.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies To** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4964	Special groups have been assigned to a new logon.

Related resource

Advanced Security Audit Policy Settings

Audit Application Generated

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Application Generated**, which determines whether the operating system generates audit events when applications attempt to use the Windows Auditing application programming interfaces (APIs).

The following events can generate audit activity:

- Creation, deletion, or initialization of an application client context
- Application operations

Applications that are designed to use the Windows Auditing APIs can use this subcategory to log auditing events that are related to those APIs. The level, volume, relevance, and importance of these audit events depend on the application that generates them. The operating system logs the events as they are generated by the application.

Event volume: Depends on the installed application's use of Windows Auditing

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4665	An attempt was made to create an application client context.
4666	An application attempted an operation:
4667	An application client context was deleted.
4668	An application was initialized.

Related resource

Advanced Security Audit Policy Settings

Audit Certification Services

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Certification Services**, which determines whether the operating system generates events when Active Directory Certificate Services (AD CS) operations are performed.

Examples of AD CS operations include:

- AD CS starts, shuts down, is backed up, or is restored.

- Certificate revocation list (CRL)-related tasks are performed.
- Certificates are requested, issued, or revoked.
- Certificate manager settings for AD CS are changed.
- The configuration and properties of the certification authority (CA) are changed.
- AD CS templates are modified.
- Certificates are imported.
- A CA certificate is published to Active Directory Domain Services.
- Security permissions for AD CS role services are modified.
- Keys are archived, imported, or retrieved.
- The OCSP Responder Service is started or stopped.

Monitoring these operational events is important to ensure that AD CS role services are functioning properly.

Event volume: Low to medium on servers that host AD CS role services

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4868	The certificate manager denied a pending certificate request.
4869	Certificate Services received a resubmitted certificate request.
4870	Certificate Services revoked a certificate.
4871	Certificate Services received a request to publish the certificate revocation list (CRL).
4872	Certificate Services published the certificate revocation list (CRL).
4873	A certificate request extension changed.
4874	One or more certificate request attributes changed.
4875	Certificate Services received a request to shut down.
4876	Certificate Services backup started.
4877	Certificate Services backup completed.

Event ID	Event message
4878	Certificate Services restore started.
4879	Certificate Services restore completed.
4880	Certificate Services started.
4881	Certificate Services stopped.
4882	The security permissions for Certificate Services changed.
4883	Certificate Services retrieved an archived key.
4884	Certificate Services imported a certificate into its database.
4885	The audit filter for Certificate Services changed.
4886	Certificate Services received a certificate request.
4887	Certificate Services approved a certificate request and issued a certificate.
4888	Certificate Services denied a certificate request.
4889	Certificate Services set the status of a certificate request to pending.
4890	The certificate manager settings for Certificate Services changed.
4891	A configuration entry changed in Certificate Services.
4892	A property of Certificate Services changed.
4893	Certificate Services archived a key.
4894	Certificate Services imported and archived a key.
4895	Certificate Services published the CA certificate to Active Directory Domain Services.
4896	One or more rows have been deleted from the certificate database.
4897	Role separation enabled:
4898	Certificate Services loaded a template.

Related resource

Advanced Security Audit Policy Settings

Audit Detailed File Share

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Detailed File Share**, which allows you to audit attempts to access files and folders on a shared folder.

The Detailed File Share setting logs an event every time a file or folder is accessed, whereas the File Share setting only records one event for any connection established between a client computer and file share. Detailed File Share audit events include detailed information about the permissions or other criteria used to grant or deny access.

Note
There are no system access control lists (SACLs) for shared folders. If this policy setting is enabled, access to all shared files and folders on the system is audited.

Event volume: High on a file server or domain controller because of SYSVOL network access required by Group Policy

Default: Not configured

If this policy setting is configured, the following event appears on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
5145	A network share object was checked to see whether the client can be granted desired access.

See Also

Advanced Security Audit Policy Settings

Audit File Share

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit File Share**, which determines whether the operating system generates audit events when a file share is accessed.

Audit events are not generated when shares are created, deleted, or when share permissions change.

Note
There are no system access control lists (SACLs) for shares; therefore, after this setting is enabled, access to all shares on the system will be audited.

Combined with File System auditing, File Share auditing enables you to track what content was accessed, the source (IP address and port) of the request, and the user account that was used for the access.

Event volume: High on a file server or domain controller (due to SYSVOL access by client computers for policy processing)

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, unless noted below.

Event ID	Event message
5140	A network share object was accessed. **Note** This event is logged on computers running Windows Server 2008 R2, Windows Server 2008, Windows 7, or Windows Vista.
5142	A network share object was added.
5143	A network share object was modified.
5144	A network share object was deleted.
5168	SPN check for SMB/SMB2 failed.

Related resource

Advanced Security Audit Policy Settings

Audit File System

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit File System**, which determines whether the operating system generates audit events when users attempt to access file system objects.

Audit events are generated only for objects that have configured system access control lists (SACLs), and only if the type of access requested (such as Write, Read, or Modify) and the account making the request match the settings in the SACL.

If success auditing is enabled, an audit entry is generated each time any account successfully accesses a file system object that has a matching SACL. If failure auditing is enabled, an audit entry is generated each time any user unsuccessfully attempts to access a file system object that has a matching SACL.

These events are essential for tracking activity for file objects that are sensitive or valuable and require extra monitoring.

Event volume: Varies, depending on how file system SACLs are configured

No audit events are generated for the default file system SACLs.

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4664	An attempt was made to create a hard link.
4985	The state of a transaction has changed.
5051	A file was virtualized.

Related resource

Advanced Security Audit Policy Settings

Audit Filtering Platform Connection

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Filtering Platform Connection**, which determines whether the operating system generates audit events when connections are allowed or blocked by the Windows Filtering Platform.

Windows Filtering Platform (WFP) was introduced in Windows Server 2008 and Windows Vista to enable independent software vendors (ISVs) to filter and modify TCP/IP packets, monitor or authorize connections, filter Internet Protocol security (IPsec)-protected traffic, and filter remote procedure calls (RPCs).

This security policy enables you to audit the following types of actions:

- The Windows Firewall service blocks an application from accepting incoming connections on the network.
- The Windows Filtering Platform allows or blocks a connection.

- The Windows Filtering Platform permits or blocks a bind to a local port.
- The Windows Filtering Platform permits or blocks an application or service from listening for incoming connections on a port.

Event volume: High

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
5031	The Windows Firewall Service blocked an application from accepting incoming connections on the network.
5140	A network share object was accessed. **Note** This event is logged only on computers running the supported versions of the Windows operating system as designated in the **Applies To** list.
5150	The Windows Filtering Platform blocked a packet.
5151	A more restrictive Windows Filtering Platform filter has blocked a packet.
5154	The Windows Filtering Platform has permitted an application or service to listen on a port for incoming connections.
5155	The Windows Filtering Platform has blocked an application or service from listening on a port for incoming connections.
5156	The Windows Filtering Platform has allowed a connection.
5157	The Windows Filtering Platform has blocked a connection.
5158	The Windows Filtering Platform has permitted a bind to a local port.
5159	The Windows Filtering Platform has blocked a bind to a local port.

Related resource

Advanced Security Audit Policy Settings

Audit Filtering Platform Packet Drop

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Filtering Platform Packet Drop**, which determines whether the operating system generates audit events when packets are dropped by the Windows Filtering Platform.

Windows Filtering Platform (WFP) was introduced in Windows Server 2008 and Windows Vista to enable independent software vendors (ISVs) to filter and modify TCP/IP packets, monitor or authorize connections, filter Internet Protocol security (IPsec)-protected traffic, and filter remote procedure calls (RPCs).

A high rate of dropped packets may indicate that there have been attempts to gain unauthorized access to computers on your network.

Event volume: High

Default setting: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
5152	The Windows Filtering Platform blocked a packet.
5153	A more restrictive Windows Filtering Platform filter has blocked a packet.

Related resource

Advanced Security Audit Policy Settings

Audit Handle Manipulation

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Handle Manipulation**, which determines whether the operating system generates audit events when a handle to an object is opened or closed.

Only objects with configured system access control lists (SACLs) generate these events, and only if the attempted handle operation matches the SACL.

Important
Handle Manipulation events are generated only for object types where the corresponding File System or Registry Object Access subcategory is enabled. For more information, see Audit File System or Audit Registry.

Event volume: High, depending on how SACLs are configured

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4656	A handle to an object was requested.
4658	The handle to an object was closed.
4690	An attempt was made to duplicate a handle to an object.

Related resource

Advanced Security Audit Policy Settings

Audit Kernel Object

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Kernel Object**, which determines whether the operating system generates audit events when users attempt to access the system kernel, which includes mutexes and semaphores.

Only kernel objects with a matching system access control list (SACL) generate security audit events. The audits generated are usually useful only to developers.

Typically, kernel objects are given SACLs only if the **AuditBaseObjects** or **AuditBaseDirectories** auditing options are enabled.

Note
The **Audit: Audit the access of global system objects** policy setting controls the default SACL of kernel objects.

Event volume: High if you have enabled one of the Global Object Access Auditing settings

Default setting: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4659	A handle to an object was requested with intent to delete.
4660	An object was deleted.
4661	A handle to an object was requested.
4663	An attempt was made to access an object.

Related resource

Advanced Security Audit Policy Settings

Audit Other Object Access Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Other Object Access Events**, which determines whether the operating system generates audit events for the management of Task Scheduler jobs or COM+ objects.

For scheduler jobs, the following actions are audited:

- Job created.
- Job deleted.
- Job enabled.
- Job disabled.
- Job updated.

For COM+ objects, the following actions are audited:

- Catalog object added.
- Catalog object updated.
- Catalog object deleted.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4671	An application attempted to access a blocked ordinal through the TBS.
4691	Indirect access to an object was requested.
4698	A scheduled task was created.
4699	A scheduled task was deleted.
4700	A scheduled task was enabled.
4701	A scheduled task was disabled.
4702	A scheduled task was updated.
5148	The Windows Filtering Platform has detected a DoS attack and entered a defensive mode; packets associated with this attack will be discarded. **Note** This event is logged only on computers running the supported versions of the Windows operating system.
5149	The DoS attack has subsided and normal processing is being resumed. **Notes** This event is logged only on computers running the supported versions of the Windows operating system.
5888	An object in the COM+ Catalog was modified.
5889	An object was deleted from the COM+ Catalog.
5890	An object was added to the COM+ Catalog.

Related resource

Advanced Security Audit Policy Settings

Audit Registry

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Registry**, which determines whether the operating system generates audit events when users attempt to access registry objects.

Audit events are generated only for objects that have configured system access control lists (SACLs) specified, and only if the type of access requested (such as Write, Read, or Modify) and the account making the request match the settings in the SACL.

If success auditing is enabled, an audit entry is generated each time any account successfully accesses a registry object that has a matching SACL. If failure auditing is enabled, an audit entry is generated each time any user unsuccessfully attempts to access a registry object that has a matching SACL.

Event volume: Low to medium, depending on how registry SACLs are configured

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4657	A registry value was modified.
5039	A registry key was virtualized.

Related resource

Advanced Security Audit Policy Settings

Audit SAM

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit SAM**, which enables you to audit events that are generated by attempts to access Security Account Manager (SAM) objects.

The Security Account Manager (SAM) is a database that is present on computers running Windows operating systems that stores user accounts and security descriptors for users on the local computer.

SAM objects include the following:

- SAM_ALIAS: A local group
- SAM_GROUP: A group that is not a local group

- SAM_USER: A user account
- SAM_DOMAIN: A domain
- SAM_SERVER: A computer account

If you configure this policy setting, an audit event is generated when a SAM object is accessed. Success audits record successful attempts, and failure audits record unsuccessful attempts.

Note
Only the SACL for SAM_SERVER can be modified.

Changes to user and group objects are tracked by the Account Management audit category. However, user accounts with enough privileges could potentially alter the files in which the account and password information is stored in the system, bypassing any Account Management events.

Event volume: High on domain controllers

Note
For information about reducing the number of events generated in this subcategory, see article 841001 in the Microsoft Knowledge Base (http://go.microsoft.com/fwlink/?LinkId=121698).

Default setting: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4659	A handle to an object was requested with intent to delete.
4660	An object was deleted.
4661	A handle to an object was requested.
4663	An attempt was made to access an object.

Related resource

Advanced Security Audit Policy Settings

Audit Audit Policy Change

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Audit Policy Change**, which determines whether the operating system generates audit events when changes are made to audit policy.

Changes to audit policy that are audited include:

- Changing permissions and audit settings on the audit policy object (by using **auditpol /set /sd**).
- Changing the system audit policy.
- Registering and unregistering security event sources.
- Changing per-user audit settings.
- Changing the value of **CrashOnAuditFail**.
- Changing audit settings on an object (for example, modifying the system access control list (SACL) for a file or registry key).

  Note
  SACL change auditing is performed when a SACL for an object has changed and the Policy Change category is configured. Discretionary access control list (DACL) and owner change auditing are performed when Object Access auditing is configured and the object's SACL is set for auditing of the DACL or owner change.
- Changing anything in the Special Groups list.

Important
Changes to the audit policy are critical security events.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista, unless otherwise noted.

Event ID	Event message
4715	The audit policy (SACL) on an object was changed.
4719	System audit policy was changed.
4817	Auditing settings on an object were changed. **Note** This event is logged only on computers running the supported versions of the

Event ID	Event message
	Windows operating system.
4902	The Per-user audit policy table was created.
4904	An attempt was made to register a security event source.
4905	An attempt was made to unregister a security event source.
4906	The CrashOnAuditFail value has changed.
4907	Auditing settings on object were changed.
4908	Special Groups Logon table modified.
4912	Per User Audit Policy was changed.

Related resource

Advanced Security Audit Policy Settings

Audit Authentication Policy Change

This topic for the IT professional describes this Advanced Security Audit policy setting, **Audit Authentication Policy Change**, which determines whether the operating system generates audit events when changes are made to authentication policy.

Changes made to authentication policy include:

- Creation, modification, and removal of forest and domain trusts.
- Changes to Kerberos policy under **Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy**.

  Note
  The audit event is logged when the policy is applied, not when settings are modified by the administrator.
- When any of the following user rights is granted to a user or group:
  - **Access this computer from the network**
  - **Allow logon locally**
  - **Allow logon through Remote Desktop**
  - **Logon as a batch job**
  - **Logon as a service**

- Namespace collision, such as when an added trust collides with an existing namespace name.

This setting is useful for tracking changes in domain-level and forest-level trust and privileges that are granted to user accounts or groups.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4713	Kerberos policy was changed.
4716	Trusted domain information was modified.
4717	System security access was granted to an account.
4718	System security access was removed from an account.
4739	Domain Policy was changed.
4864	A namespace collision was detected.
4865	A trusted forest information entry was added.
4866	A trusted forest information entry was removed.
4867	A trusted forest information entry was modified.

Related resource

Advanced Security Audit Policy Settings

Audit Authorization Policy Change

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Authorization Policy Change**, which determines whether the operating system generates audit events when specific changes are made to the authorization policy.

Authorization policy changes that can be audited include:

- Assigning or removing user rights (privileges) such as **SeCreateTokenPrivilege**, except for the system access rights that are audited by using the Audit Authentication Policy Change subcategory.

- Changing the Encrypting File System (EFS) policy.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4704	A user right was assigned.
4705	A user right was removed.
4706	A new trust was created to a domain.
4707	A trust to a domain was removed.
4714	Encrypted data recovery policy was changed.

Related resource

Advanced Security Audit Policy Settings

Audit Filtering Platform Policy Change

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Filtering Platform Policy Change**, which determines whether the operating system generates audit events for certain IPsec and Windows Filtering Platform actions.

Windows Filtering Platform (WFP) enables independent software vendors (ISVs) to filter and modify TCP/IP packets, monitor or authorize connections, filter Internet Protocol security (IPsec)-protected traffic, and filter remote procedure calls (RPCs).

This security policy setting determines whether the operating system generates audit events for:

- IPsec services status.
- Changes to IPsec settings.
- Status and changes to the Windows Filtering Platform engine and providers.
- IPsec Policy Agent service activities.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4709	IPsec Services was started.
4710	IPsec Services was disabled.
4711	May contain any one of the following: • PAStore Engine applied locally cached copy of Active Directory storage IPsec policy on the computer. • PAStore Engine applied Active Directory storage IPsec policy on the computer. • PAStore Engine applied local registry storage IPsec policy on the computer. • PAStore Engine failed to apply locally cached copy of Active Directory storage IPsec policy on the computer. • PAStore Engine failed to apply Active Directory storage IPsec policy on the computer. • PAStore Engine failed to apply local registry storage IPsec policy on the computer. • PAStore Engine failed to apply some rules of the active IPsec policy on the computer. • PAStore Engine failed to load directory storage IPsec policy on the computer. • PAStore Engine loaded directory storage IPsec policy on the computer. • PAStore Engine failed to load local storage IPsec policy on the computer. • PAStore Engine loaded local storage IPsec policy on the computer. • PAStore Engine polled for changes to the active IPsec policy and detected no changes.
4712	IPsec Services encountered a potentially serious failure.
5040	A change has been made to IPsec settings. An Authentication Set was added.
5041	A change has been made to IPsec settings. An Authentication Set was modified.

Event ID	Event message
5042	A change has been made to IPsec settings. An Authentication Set was deleted.
5043	A change has been made to IPsec settings. A Connection Security Rule was added.
5044	A change has been made to IPsec settings. A Connection Security Rule was modified.
5045	A change has been made to IPsec settings. A Connection Security Rule was deleted.
5046	A change has been made to IPsec settings. A Crypto Set was added.
5047	A change has been made to IPsec settings. A Crypto Set was modified.
5048	A change has been made to IPsec settings. A Crypto Set was deleted.
5440	The following callout was present when the Windows Filtering Platform Base Filtering Engine started.
5441	The following filter was present when the Windows Filtering Platform Base Filtering Engine started.
5442	The following provider was present when the Windows Filtering Platform Base Filtering Engine started.
5443	The following provider context was present when the Windows Filtering Platform Base Filtering Engine started.
5444	The following sub-layer was present when the Windows Filtering Platform Base Filtering Engine started.
5446	A Windows Filtering Platform callout has been changed.
5448	A Windows Filtering Platform provider has been changed.
5449	A Windows Filtering Platform provider context

Event ID	Event message
	has been changed.
5450	A Windows Filtering Platform sub-layer has been changed.
5456	PAStore Engine applied Active Directory storage IPsec policy on the computer.
5457	PAStore Engine failed to apply Active Directory storage IPsec policy on the computer.
5458	PAStore Engine applied locally cached copy of Active Directory storage IPsec policy on the computer.
5459	PAStore Engine failed to apply locally cached copy of Active Directory storage IPsec policy on the computer.
5460	PAStore Engine applied local registry storage IPsec policy on the computer.
5461	PAStore Engine failed to apply local registry storage IPsec policy on the computer.
5462	PAStore Engine failed to apply some rules of the active IPsec policy on the computer. Use the IP Security Monitor snap-in to diagnose the problem.
5463	PAStore Engine polled for changes to the active IPsec policy and detected no changes.
5464	PAStore Engine polled for changes to the active IPsec policy, detected changes, and applied them to IPsec Services.
5465	PAStore Engine received a control for forced reloading of IPsec policy and processed the control successfully.
5466	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory cannot be reached, and will use the cached copy of the Active Directory IPsec policy instead. Any changes made to the Active Directory IPsec policy since the last poll could not be applied.

Event ID	Event message
5467	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory can be reached, and found no changes to the policy. The cached copy of the Active Directory IPsec policy is no longer being used.
5468	PAStore Engine polled for changes to the Active Directory IPsec policy, determined that Active Directory can be reached, found changes to the policy, and applied those changes. The cached copy of the Active Directory IPsec policy is no longer being used.
5471	PAStore Engine loaded local storage IPsec policy on the computer.
5472	PAStore Engine failed to load local storage IPsec policy on the computer.
5473	PAStore Engine loaded directory storage IPsec policy on the computer.
5474	PAStore Engine failed to load directory storage IPsec policy on the computer.
5477	PAStore Engine failed to add quick mode filter.

Related resource

Advanced Security Audit Policy Settings

Audit MPSSVC Rule-Level Policy Change

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit MPSSVC Rule-Level Policy Change**, which determines whether the operating system generates audit events when changes are made to policy rules for the Microsoft Protection Service (MPSSVC.exe).

The Microsoft Protection Service, which is used by Windows Firewall, is an integral part of the computer's threat protection against Internet-bound threats such as Trojans and spyware. The tracked activities include:

- Active policies when the Windows Firewall service starts.

- Changes to Windows Firewall rules.
- Changes to the Windows Firewall exception list.
- Changes to Windows Firewall settings.
- Rules ignored or not applied by the Windows Firewall service.
- Changes to Windows Firewall Group Policy settings.

Changes to firewall rules are important for understanding the security state of the computer and how well it is protected against network attacks.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4944	The following policy was active when the Windows Firewall started.
4945	A rule was listed when the Windows Firewall started.
4946	A change has been made to Windows Firewall exception list. A rule was added.
4947	A change has been made to Windows Firewall exception list. A rule was modified.
4948	A change has been made to Windows Firewall exception list. A rule was deleted.
4949	Windows Firewall settings were restored to the default values.
4950	A Windows Firewall setting has changed.
4951	A rule has been ignored because its major version number was not recognized by Windows Firewall.
4952	Parts of a rule have been ignored because its minor version number was not recognized by Windows Firewall. The other parts of the rule will be enforced.
4953	A rule has been ignored by Windows Firewall because it could not parse the rule.

Event ID	Event message
4954	Windows Firewall Group Policy settings have changed. The new settings have been applied.
4956	Windows Firewall has changed the active profile.
4957	Windows Firewall did not apply the following rule:
4958	Windows Firewall did not apply the following rule because the rule referred to items not configured on this computer:

Related resource

Advanced Security Audit Policy Settings

Audit Other Policy Change Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Other Policy Change Events**, which determines whether the operating system generates audit events for security policy changes that are not otherwise audited in the Policy Change category.

These other activities in the Policy Change category that can be audited include:

- Trusted Platform Module (TPM) configuration changes.
- Kernel-mode cryptographic self tests.
- Cryptographic provider operations.
- Cryptographic context operations or modifications.

Event volume: Low

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4670	Permissions on an object were changed.
4909	The local policy settings for the TBS were changed.

Event ID	Event message
4910	The group policy settings for the TBS were changed.
5063	A cryptographic provider operation was attempted.
5064	A cryptographic context operation was attempted.
5065	A cryptographic context modification was attempted.
5066	A cryptographic function operation was attempted.
5067	A cryptographic function modification was attempted.
5068	A cryptographic function provider operation was attempted.
5069	A cryptographic function property operation was attempted.
5070	A cryptographic function property modification was attempted.
5447	A Windows Filtering Platform filter has been changed.
6144	Security policy in the group policy objects has been applied successfully.
6145	One or more errors occurred while processing security policy in the group policy objects.

Related resource

Advanced Security Audit Policy Settings

Audit Sensitive Privilege Use

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Sensitive Privilege Use**, which determines whether the operating system generates audit events when sensitive privileges (user rights) are used.

Actions that can be audited include:

- A privileged service is called.
- One of the following privileges is called:

  Act as part of the operating system

  Back up files and directories

  Create a token object

  Debug programs

  Enable computer and user accounts to be trusted for delegation

  Generate security audits

  Impersonate a client after authentication

  Load and unload device drivers

  Manage auditing and security log

  Modify firmware environment values

  Replace a process-level token

  Restore files and directories

  Take ownership of files or other objects

If you configure this policy setting, an audit event is generated when sensitive privilege requests are made. Success audits record successful attempts, and failure audits record unsuccessful attempts.

Event volume: High

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4672	Special privileges assigned to new logon.
4673	A privileged service was called.
4674	An operation was attempted on a privileged object.

Related resource

Advanced Security Audit Policy Settings

Audit Non-Sensitive Privilege Use

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Non-Sensitive Privilege Use**, which determines whether the operating system generates audit events when non-sensitive privileges (user rights) are used.

The following privileges are non-sensitive:

- **Access Credential Manager as a trusted caller**
- **Access this computer from the network**
- **Add workstations to domain**
- **Adjust memory quotas for a process**
- **Allow log on locally**
- **Allow log on through Terminal Services**
- **Bypass traverse checking**
- **Change the system time**
- **Create a page file**
- **Create global objects**
- **Create permanent shared objects**
- **Create symbolic links**
- **Deny access to this computer from the network**
- **Deny log on as a batch job**
- **Deny log on as a service**
- **Deny log on locally**
- **Deny log on through Terminal Services**
- **Force shutdown from a remote system**
- **Increase a process working set**
- **Increase scheduling priority**
- **Lock pages in memory**
- **Log on as a batch job**
- **Log on as a service**
- **Modify an object label**
- **Perform volume maintenance tasks**
- **Profile single process**
- **Profile system performance**
- **Remove computer from docking station**
- **Shut down the system**
- **Synchronize directory service data**

If you configure this policy setting, an audit event is generated when a non-sensitive privilege is called. Success audits record successful attempts, and failure audits record unsuccessful attempts.

Event volume: Very high

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4672	Special privileges assigned to new logon.
4673	A privileged service was called.
4674	An operation was attempted on a privileged object.

Related resource

Advanced Security Audit Policy Settings

Audit Other Privilege Use Events

This security policy setting is not used.

Related resource

Advanced Security Audit Policy Settings

Audit IPsec Driver

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit IPsec Driver**, which determines whether the operating system generates audit events for the activities of the IPsec driver.

The IPsec driver, using the IP Filter List from the active IPsec policy, watches for outbound IP packets that must be secured and inbound IP packets that must be verified and decrypted. This security policy setting reports on the following activities of the IPsec driver:

- Startup and shutdown of IPsec services.
- Packets dropped due to integrity-check failure.

- Packets dropped due to replay-check failure.
- Packets dropped due to being in plaintext.
- Packets received with an incorrect Security Parameter Index (SPI). (This can indicate malfunctioning hardware or interoperability problems.)
- Failure to process IPsec filters.

A high rate of packet drops by the IPsec filter driver may indicate attempts to gain access to the network by unauthorized systems.

Failure to process IPsec filters poses a potential security risk because some network interfaces may not get the protection that is provided by the IPsec filter.

Event volume: Medium

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4960	IPsec dropped an inbound packet that failed an integrity check. If this problem persists, it could indicate a network issue or that packets are being modified in transit to this computer. Verify that the packets sent from the remote computer are the same as those received by this computer. This error might also indicate interoperability problems with other IPsec implementations.
4961	IPsec dropped an inbound packet that failed a replay check. If this problem persists, it could indicate a replay attack against this computer.
4962	IPsec dropped an inbound packet that failed a replay check. The inbound packet had too low a sequence number to ensure it was not a replay.
4963	IPsec dropped an inbound clear text packet that should have been secured. This is usually due to the remote computer changing its IPsec policy without informing this computer. This could also be a spoofing attack attempt.
4965	IPsec received a packet from a remote computer with an incorrect Security Parameter Index (SPI). This is usually caused by

Event ID	Event message
	malfunctioning hardware that is corrupting packets. If these errors persist, verify that the packets sent from the remote computer are the same as those received by this computer. This error may also indicate interoperability problems with other IPsec implementations. In that case, if connectivity is not impeded, then these events can be ignored.
5478	IPsec Services has started successfully.
5479	IPsec Services has been shut down successfully. The shutdown of IPsec Services can put the computer at greater risk of network attack or expose the computer to potential security risks.
5480	IPsec Services failed to get the complete list of network interfaces on the computer. This poses a potential security risk because some of the network interfaces may not get the protection provided by the applied IPsec filters. Use the IP Security Monitor snap-in to diagnose the problem.
5483	IPsec Services failed to initialize RPC server. IPsec Services could not be started.
5484	IPsec Services has experienced a critical failure and has been shut down. The shutdown of IPsec Services can put the computer at greater risk of network attack or expose the computer to potential security risks.
5485	IPsec Services failed to process some IPsec filters on a plug-and-play event for network interfaces. This poses a potential security risk because some of the network interfaces may not get the protection provided by the applied IPsec filters. Use the IP Security Monitor snap-in to diagnose the problem.

Related resource

Advanced Security Audit Policy Settings

Audit Other System Events

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Other System Events**, which determines whether the operating system audits various system events.

The system events in this category include:

- Startup and shutdown of the Windows Firewall service and driver.
- Security policy processing by the Windows Firewall service.
- Cryptography key file and migration operations.

Important
Failure to start the Windows Firewall service may result in a computer that is not fully protected against network threats.

Event volume: Low

Default: Success and failure

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista, unless otherwise noted.

Event ID	Event message
5024	The Windows Firewall Service has started successfully.
5025	The Windows Firewall Service has been stopped.
5027	The Windows Firewall Service was unable to retrieve the security policy from the local storage. The service will continue enforcing the current policy.
5028	The Windows Firewall Service was unable to parse the new security policy. The service will continue with currently enforced policy.
5029	The Windows Firewall Service failed to initialize the driver. The service will continue to enforce

Event ID	Event message
	the current policy.
5030	The Windows Firewall Service failed to start.
5032	Windows Firewall was unable to notify the user that it blocked an application from accepting incoming connections on the network.
5033	The Windows Firewall Driver has started successfully.
5034	The Windows Firewall Driver has been stopped.
5035	The Windows Firewall Driver failed to start.
5037	The Windows Firewall Driver detected critical runtime error. Terminating.
5058	Key file operation.
5059	Key migration operation.
6400	BranchCache: Received an incorrectly formatted response while discovering availability of content. **Notes** This event is logged only on computers running supported versions of the Windows operating system.
6401	BranchCache: Received invalid data from a peer. Data discarded. **Note** This event is logged only on computers running supported versions of the Windows operating system. .
6402	BranchCache: The message to the hosted cache offering it data is incorrectly formatted. **Note** This event is logged only on computers running supported versions of the Windows operating system.

Event ID	Event message
6403	BranchCache: The hosted cache sent an incorrectly formatted response to the client. **Note** This event is logged only on computers running supported versions of the Windows operating system.
6404	BranchCache: Hosted cache could not be authenticated using the provisioned SSL certificate. **Note** This event is logged only on computers running supported versions of the Windows operating system.
6405	BranchCache: %2 instance(s) of event id %1 occurred. **Note** This event is logged only on computers running supported versions of the Windows operating system.
6406	%1 registered to Windows Firewall to control filtering for the following: %2 **Note** This event is logged only on computers running supported versions of the Windows operating system.
6407	1% **Note** This event is logged only on computers running supported versions of the Windows operating system.
6408	Registered product %1 failed and Windows Firewall is now controlling the filtering for %2 **Notes** This event is logged only on computers running supported versions of the

Event ID	Event message
	Windows operating system.

Related resource

Advanced Security Audit Policy Settings

Audit Security State Change

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Security State Change**, which determines whether Windows generates audit events for changes in the security state of a system.

Changes in the security state of the operating system include:

- System startup and shutdown.
- Change of system time.
- System recovery from **CrashOnAuditFail**. This event is logged after a system reboots following **CrashOnAuditFail**.

Important
Some auditable activity may not be recorded when a system reboots due to **CrashOnAuditFail**.

System startup and shutdown events are important for understanding system usage.

Event volume: Low

Default: Success

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic in addition to Windows Server 2008 and Windows Vista.

Event ID	Event Message Summary	Minimum Requirement
4608	Windows is starting up.	Windows Vista, Windows Server 2008
4609	Windows is shutting down.	Windows Vista, Windows Server 2008
4616	The system time was changed.	Windows Vista, Windows Server 2008
4621	Administrator recovered system	Windows Vista, Windows

Event ID	Event Message Summary	Minimum Requirement
	from CrashOnAuditFail. Users who are not administrators will now be allowed to log on. Some auditable activity might not have been recorded.	Server 2008

Related resource

Advanced Security Audit Policy Settings

Audit Security System Extension

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit Security System Extension**, which determines whether the operating system generates audit events related to security system extensions.

Changes to security system extensions in the operating system include the following activities:

- A security extension code is loaded (such as an authentication, notification, or security package). A security extension code registers with the Local Security Authority and will be used and trusted to authenticate logon attempts, submit logon requests, and be notified of any account or password changes. Examples of this extension code are Security Support Providers, such as Kerberos and NTLM.
- A service is installed. An audit log is generated when a service is registered with the Service Control Manager. The audit log contains information about the service name, binary, type, start type, and service account.

Important
Attempts to install or load security system extensions or services are critical system events that could indicate a security breach.

Event volume: Low

These events are expected to appear more on a domain controller than on client computers or member servers.

Default: Not configured

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista.

Event ID	Event message
4610	An authentication package has been loaded by

Event ID	Event message
	the Local Security Authority.
4611	A trusted logon process has been registered with the Local Security Authority.
4614	A notification package has been loaded by the Security Account Manager.
4622	A security package has been loaded by the Local Security Authority.
4697	A service was installed in the system.

Related resource

Advanced Security Audit Policy Settings

Audit System Integrity

This topic for the IT professional describes the Advanced Security Audit policy setting, **Audit System Integrity**, which determines whether the operating system audits events that violate the integrity of the security subsystem.

Activities that violate the integrity of the security subsystem include the following:

- Audited events are lost due to a failure of the auditing system.
- A process uses an invalid local procedure call (LPC) port in an attempt to impersonate a client, reply to a client address space, read to a client address space, or write from a client address space.
- A remote procedure call (RPC) integrity violation is detected.
- A code integrity violation with an invalid hash value of an executable file is detected.
- Cryptographic tasks are performed.

Important
Violations of security subsystem integrity are critical and could indicate a potential security attack.

Event volume: Low

Default: Success and failure

If this policy setting is configured, the following events appear on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic, in addition to Windows Server 2008 and Windows Vista, unless otherwise noted.

Event ID	Event message
4612	Internal resources allocated for the queuing of audit messages have been exhausted, leading to the loss of some audits.
4615	Invalid use of LPC port.
4618	A monitored security event pattern has occurred.
4816	RPC detected an integrity violation while decrypting an incoming message.
5038	Code integrity determined that the image hash of a file is not valid. The file could be corrupt due to unauthorized modification or the invalid hash could indicate a potential disk device error.
5056	A cryptographic self-test was performed.
5057	A cryptographic primitive operation failed.
5060	Verification operation failed.
5061	Cryptographic operation.
5062	A kernel-mode cryptographic self-test was performed.
6281	Code Integrity determined that the page hashes of an image file are not valid. The file could be improperly signed without page hashes or corrupt due to unauthorized modification. The invalid hashes could indicate a potential disk device error. **Note** This event is logged only on computers running the supported versions of the Windows operating system.

Related resource

Advanced Security Audit Policy Settings

Registry (Global Object Access Auditing)

This topic for the IT professional describes the Advanced Security Audit policy setting, **Registry (Global Object Access Auditing)**, which enables you to configure a global system access control list (SACL) on the registry of a computer.

If you select the **Configure security** check box on this policy's property page, you can add a user or group to the global SACL. This enables you to define computer system access control lists (SACLs) per object type for the registry. The specified SACL is then automatically applied to every registry object type.

This policy setting must be used in combination with the **Registry** security policy setting under Object Access. For more information, see Audit Registry.

This policy setting is available on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic.

Related resource

Advanced Security Audit Policy Settings

File System (Global Object Access Auditing)

This topic for the IT professional describes the Advanced Security Audit policy setting, **File System (Global Object Access Auditing)**, which enables you to configure a global system access control list (SACL) on the file system for an entire computer.

If you select the **Configure security** check box on the policy's property page, you can add a user or group to the global SACL. This enables you to define computer system access control lists (SACLs) per object type for the file system. The specified SACL is then automatically applied to every file system object type.

If both a file or folder SACL and a global SACL are configured on a computer, the effective SACL is derived by combining the file or folder SACL and the global SACL. This means that an audit event is generated if an activity matches either the file or folder SACL or the global SACL.

This policy setting must be used in combination with the **File System** security policy setting under Object Access. For more information, see Audit File System.

This policy setting is available on computers running the supported versions of the Windows operating system as designated in the **Applies to** list at the beginning of this topic.

Related resource

Advanced Security Audit Policy Settings

Security Policy Settings Overview

This topic provides high-level information about security settings in, and links to more detailed reference topics, procedures, and guides about security policy settings.

Did you mean…

- Group Policy (see Group Policy Overview)
- Threats and Countermeasures Guide (see the resources on Windows Server Security)

Security policy settings description

Security settings incorporated into policies are rules that administrators configure on a computer or multiple computers for the purpose of protecting resources on a computer or network. The Security Settings extension of the Local Group Policy Object Editor snap-in allows you to define security configurations as part of a Group Policy Object (GPO). The GPOs are linked to Active Directory containers such as sites, domains, or organizational units, and enable administrators to manage security settings for multiple computers from any computer joined to the domain. Security settings are used as part of your overall security implementation to help secure domain controllers, servers, clients and other resources in your organization.

Practical applications

Security settings can control:

- User authentication to a network or computer.
- The resources that users are permitted to access.
- Whether to record a user's or group's actions in the Event log.
- Membership in a group.

New and changed functionality

There are no changes in functionality to the tools and snap-ins used to administer security policy settings on the local computer or throughout the domain using Group Policy.

There were no policy settings added in Windows Server 2012 R2 and Windows 8.1.

The following table lists the Security Settings new for Windows Server 2012 and Windows 8. For detailed descriptions of these and other settings, see Security Policy Settings Reference.

Security policy	Location	Description
Accounts: Block Microsoft accounts	Windows Settings/Security Settings/Local Policies/Security Options	This policy setting prevents users from adding new Microsoft accounts on this computer.

Security policy	Location	Description
Interactive logon: Machine account threshold.	Windows Settings/Security Settings/Local Policies/Security Options	The computer lockout policy is enforced only on those computers that have BitLocker enabled for protecting operating system volumes. Please ensure that appropriate recovery password backup policies are enabled.
Interactive logon: Machine inactivity limit.	Windows Settings/Security Settings/Local Policies/Security Options	Windows notices inactivity of a logon session, and if the amount of inactive time exceeds the inactivity limit, then the screen saver will run, locking the session.
Microsoft network server: Attempt S4U2Self to obtain claim information	Windows Settings/Security Settings/Local Policies/Security Options	This security setting is to support clients running a version of Windows prior to Windows 8 that are trying to access a file share that requires user claims. This setting determines whether the local file server will attempt to use Kerberos Service-For-User-To-Self (S4U2Self) functionality to obtain a network client principal's claims from the client's account domain
Packaged app Rules	Windows Settings/Security Settings/Application Control Policies/AppLocker	AppLocker rules that apply to files that share the same attributes: publisher name, package name, and package version.

Deprecated functionality

No functionality to set or maintain security settings has been deprecated in the versions designated in the **Applies to** list at the top of this page.

See also

For information about methods to administer security policies and policy settings, see Administer Security Policy Settings.

The following table provides links to additional resources that can help you understand and implement security policies using security settings in the versions designated in the **Applies to** list at the top of this page.

	References
Deployment	Group Policy Settings Reference for Windows and Windows Server from Official Microsoft Download Center
Operations	Administer Security Policy Settings
Troubleshooting	For Windows Server 2003 and Windows XP, see Security Settings Extension Technical Reference For Windows Server 2008 R2, Windows Server 2008, Windows 7, and Windows Vista, see Server Security Policy Management Technical Reference. For others, see Security Policy Settings Technical Overview.
Security	Windows Security Survival Guide Threats and Countermeasures Guide: Security Settings in Windows Server 2008 R2 and Windows 7 For Windows Server 2012 and Windows 8, see Microsoft Security Compliance Manager (SCM) in the TechNet Wiki.
Tools and settings	Security Policy Settings Technical Overview Security Policy Settings Reference For Windows Server 2003, see Security Settings Extension Technical Reference Download: Group Policy Settings Reference for Windows and Windows Server
Community resources	Windows Security Survival Guide
Related technologies	Group Policy Overview

Security Policy Settings Technical Overview

This reference topic describes the common scenarios, architecture, and processes for security settings.

Did you mean…

- Security Policy Settings Reference

Introduction

Security policy settings are rules that administrators configure on a computer or multiple computers for the purpose of protecting resources on a computer or network. The Security Settings extension of the Local Group Policy Editor snap-in allows you to define security configurations as part of a Group Policy Object (GPO). The GPOs are linked to Active Directory containers such as sites, domains, or organizational units, and they enable administrators to manage security settings for multiple computers from any computer joined to the domain. Security settings policies are used as part of your overall security implementation to help secure domain controllers, servers, clients, and other resources in your organization.

Security settings can control:

- User authentication to a network or computer.
- The resources that users are permitted to access.
- Whether to record a user's or group's actions in the Event log.
- Membership in a group.

To manage security configurations for multiple computers, you can use one of the following options:

- Edit specific security settings in a GPO.
- Use the Security Templates snap-in to create a security template that contains the security policies you want to apply, and then import the security template into a Group Policy Object. A security template is a file that represents a security configuration, and it can be imported to a GPO, applied to a local computer, or used to analyze security.

For more information about managing security configurations, see Administer Security Policy Settings.

The Security Settings extension of the Local Group Policy Editor includes the following types of security policies:

- **Account Policies.** These polices are defined on computers; they affect how user accounts can interact with the computer or domain. Account policies include the following types of policies:
  - **Password Policy.** These policies determine settings for passwords, such as enforcement and lifetimes. Password policies are used for domain accounts.
  - **Account Lockout Policy.** These policies determine the conditions and length of time that an account will be locked out of the system. Account lockout policies are used for domain or local user accounts.

    - **Kerberos Policy.** These policies are used for domain user accounts; they determine Kerberos-related settings, such as ticket lifetimes and enforcement.
- **Local Policies.** These policies apply to a computer and include the following types of policy settings:
    - **Audit Policy.** Specify security settings that control the logging of security events into the Security log on the computer, and specifies what types of security events to log (success, failure, or both).

      Note
      For computers running Windows 7, Windows Server 2008 R2, Windows 8 and Windows Server 2012, it is recommended to use the settings under Advanced Audit Policy Configuration rather than the Audit Policy settings under Local Policies.
    - **User Rights Assignment.** Specify the users or groups that have logon rights or privileges on a computer.
    - **Security Options.** Specify security settings for the computer, such as Administrator and Guest Account names; access to floppy disk drives and CD-ROM drives; installation of drivers; logon prompts; and so on.
- **Windows Firewall with Advanced Security.** Specify settings to protect the computers on your network by using a stateful firewall that allows you to determine which network traffic is permitted to pass between your computer and the network.
- **Network List Manager Policies.** Specify settings that you can use to configure different aspects of how networks are listed and displayed on one computer or on many computers.
- **Public Key Policies.** Specify settings to control Encrypting File System, Data Protection, and BitLocker Drive Encryption in addition to certain certificate paths and services settings.
- **Software Restriction Policies.** Specify settings to identify software and to control its ability to run on your local computer, organizational unit, domain, or site.
- **Application Control Policies.** Specify settings to control which users or groups can run particular applications in your organization based on unique identities of files.
- **IP Security Policies on Local Computer.** Specify settings to ensure private, secure communications over IP networks through the use of cryptographic security services. IPsec establishes trust and security from a source IP address to a destination IP address.
- **Advanced Audit Policy Configuration.** Specify settings that control the logging of security events into the Security log on the computer. The settings under Advanced Audit Policy Configuration provide finer control over which activities to monitor as opposed to the Audit Policy settings under Local Policies.

Policy-based Security Settings management

The Security Settings extension to Group Policy provides an integrated policy-based management infrastructure to help you manage and enforce your security policies.

You can define and apply security settings policies to users, groups, and network servers and clients through Group Policy and Active Directory Domain Services (AD DS). A group of servers

with the same functionality can be created (for example, a Microsoft Web (IIS) server, and then Group Policy Objects can be used to apply common security settings to the group. If more servers are added to this group later, many of the common security settings are automatically applied, reducing deployment and administrative labor.

Common scenarios for using Security Settings policies

Security Settings policies are used to manage the following aspects of security: accounts policy, local policy, user rights assignment, registry values, file and registry Access Control Lists (ACLs), service startup modes, and more.

As part of your security strategy, you can create GPOs with security settings policies configured specifically for the various computer roles in your organization, such as domain controllers, file servers, member servers, clients, and so on.

You can create an organizational unit (OU) structure that groups computers according to their roles. Using OUs is the best method for separating specific security requirements for the different computer roles in your network. This approach also allows you to apply customized security templates to each class of server or computer. After creating the security templates, you create a new GPO for each of the OUs, and then import the security template (.inf file) into the new GPO.

Importing a security template to a GPO ensures that any accounts to which the GPO is applied automatically receive the template's security settings when the Group Policy settings are refreshed. On a workstation or server, the security settings are refreshed at regular intervals (with a random offset of at most 30 minutes), and, on a domain controller, this process occurs every few minutes if changes have occurred in any of the GPO settings that apply. The settings are also refreshed every 16 hours, whether or not any changes have occurred.

Note
These refresh settings vary between versions of the operating system and can be configured.

By using Group Policy−based security configurations in conjunction with the delegation of administration, you can ensure that specific security settings, rights, and behavior are applied to all servers and computers within an OU. This approach makes it simple to update a number of servers with any additional changes required in the future.

Dependencies on other operating system technologies

For computers that are members of a Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 domain, Security Settings policies depend on the following technologies:

- **Active Directory Domain Services (AD DS)**

  The Windows-based directory service, AD DS, stores information about objects on a network and makes this information available to administrators and users. By using AD DS, you can view and manage network objects on the network from a single location, and users can access permitted network resources by using a single logon.
- **Group Policy**

The infrastructure within AD DS that enables directory-based configuration management of user and computer settings on computers running Windows Server operating systems. By using Group Policy, you can define configurations for groups of users and computers, including policy settings, registry-based policies, software installation, scripts, folder redirection, Remote Installation Services, Internet Explorer maintenance, and security.

- **Domain Name System (DNS)**

  A hierarchical naming system used for locating domain names on the Internet and on private TCP/IP networks. DNS provides a service for mapping DNS domain names to IP addresses, and IP addresses to domain names. This allows users, computers, and applications to query DNS to specify remote systems by fully qualified domain names rather than by IP addresses.

- **Winlogon**

  A part of the Windows operating system that provides interactive logon support. Winlogon is designed around an interactive logon model that consists of three components: the Winlogon executable, a credential provider, and any number of network providers.

- **Setup**

  Security configuration interacts with the operating system setup process during a clean installation or upgrade from earlier versions of a Windows Server operating system.

- **Security Accounts Manager (SAM)**

  A Windows service used during the logon process. SAM maintains user account information, including groups to which a user belongs.

- **Local Security Authority (LSA)**

  A protected subsystem that authenticates and logs users onto the local system. LSA also maintains information about all aspects of local security on a system, collectively known as the Local Security Policy of the system.

- **Windows Management Instrumentation (WMI)**

  A feature of the Microsoft Windows operating system, WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI provides access to information about objects in a managed environment. Through WMI and the WMI application programming interface (API), applications can query for and make changes to static information in the Common Information Model (CIM) repository and dynamic information maintained by the various types of providers.

- **Resultant Set of Policy (RSoP)**

  An enhanced Group Policy infrastructure that uses WMI in order to make it easier to plan and debug policy settings. RSoP provides public methods that expose what an extension to Group Policy would do in a what-if situation, and what the extension has done in an actual situation. This allows administrators to easily determine the combination of policy settings that apply to, or will apply to, a user or computer. In Windows Server 2012, this feature has tighter coupling with Group Policy's Infrastructure Status feature.

- **Service Control Manager (SCM)**

Used for configuration of service startup modes and security.

- **Registry**

  Used for configuration of registry values and security.

- **File system**

  Used for configuration of security.

- **File system conversions**

  Security is set when an administrator converts a file system from FAT to NTFS.

- **Microsoft Management Console (MMC)**

  The user interface for the Security Settings tool is an extension of the Local Group Policy Editor MMC snap-in.

Security Settings policies and Group Policy

The Security Settings extension of the Local Group Policy Editor is part of the Security Configuration Manager tool set. The following components are associated with Security Settings: a configuration engine; an analysis engine; a template and database interface layer; setup integration logic; and the Secedit.exe command-line tool. The security configuration engine is responsible for handling security configuration editor-related security requests for the system on which it runs. The analysis engine analyzes system security for a given configuration and saves the result. The template and database interface layer handles reading and writing requests from and to the template or database (for internal storage). The Security Settings extension of the Local Group Policy Editor handles Group Policy from a domain-based or local computer. The security configuration logic integrates with setup and manages system security for a clean installation or upgrade to a more recent Windows operating system. Security information is stored in templates (.inf files) or in the Secedit.sdb database.

The following diagram shows Security Settings and related features.

Security Settings Policies and Related Features

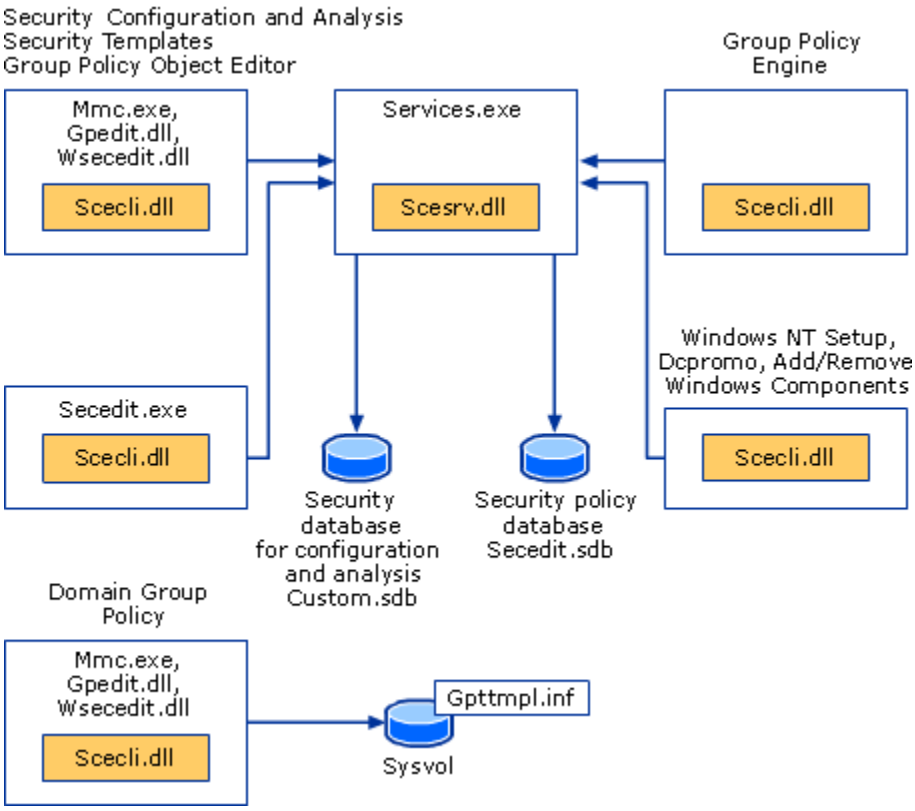

- **Scesrv.dll**

  Provides the core security engine functionality.
- **Scecli.dll**

  Provides the client-side interfaces to the security configuration engine and provides data to Resultant Set of Policy (RSoP).
- **Wsecedit.dll**

  The Security Settings extension of Local Group Policy Editor. Scecli.dll is loaded into Wsecedit.dll to support the Security Settings user interface.
- **Gpedit.dll**

  The Local Group Policy Editor MMC snap-in.

Security Settings extension architecture

The Security Settings extension of the Local Group Policy Editor is part of the Security Configuration Manager tools, as shown in the following diagram.

Security Settings Architecture

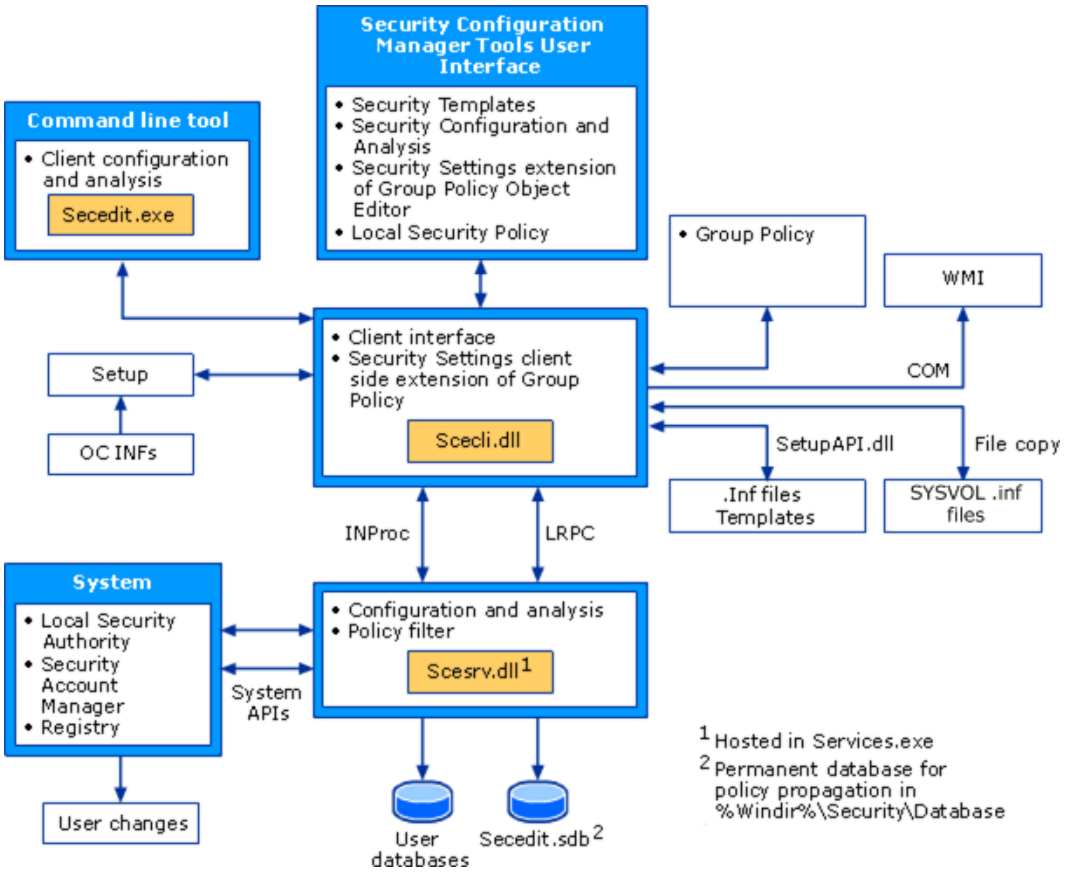

The security settings configuration and analysis tools include a security configuration engine, which provides local computer (non-domain member) and Group Policy−based configuration and analysis of security settings policies. The security configuration engine also supports the creation of security policy files. The primary features of the security configuration engine are Scecli.dll and Scesrv.dll.

The following list describes these primary features of the security configuration engine and other Security Settings−related features.

- **Scesrv.dll**

  This .dll is hosted in Services.exe and runs under local system context. Scesrv.dll provides core Security Configuration Manager functionality, such as import, configure, analyze, and policy propagation.

  Scesrv.dll performs configuration and analysis of various security-related system parameters by calling corresponding system APIs, including LSA, SAM, and the registry.

  Scesrv.dll exposes APIs such as import, export, configure, and analyze. It checks that the request is made over LRPC (Windows XP) and fails the call if it is not.

  Communication between parts of the Security Settings extension occurs by using the following methods:

- Component Object Model (COM) calls
- Local Remote Procedure Call (LRPC)
- Lightweight Directory Access Protocol (LDAP)
- Active Directory Service Interfaces (ADSI)
- Server Message Block (SMB)
- Win32 APIs
- Windows Management Instrumentation (WMI) calls

On domain controllers, Scesrv.dll receives notifications of changes made to SAM and the LSA that need to be synchronized across domain controllers. Scesrv.dll incorporates those changes into the Default Domain Controller Policy GPO by using in-process Scecli.dll template modification APIs.

Scesrv.dll also performs configuration and analysis operations.

- **Scecli.dll**

  This is the client-side interface or wrapper to Scesrv.dll. Scecli.dll is loaded into Wsecedit.dll to support MMC snap-ins. It is used by Setup to configure default system security and security of files, registry keys, and services installed by the Setup API .inf files.

  The command-line version of the security configuration and analysis user interfaces, Secedit.exe, uses Scecli.dll.

  Scecli.dll implements the client-side extension for Group Policy.

  Scesrv.dll uses Scecli.dll to download applicable Group Policy files from SYSVOL in order to apply Group Policy security settings to the local computer.

  Scecli.dll logs application of security policy into WMI (RSoP).

  Scesrv.dll policy filter uses Scecli.dll to update Default Domain Controller Policy GPO when changes are made to SAM and LSA.

- **Wsecedit.dll**

  The Security Settings extension of the Group Policy Object Editor snap-in. You use this tool to configure security settings in a Group Policy Object for a site, domain, or organizational unit. You can also use Security Settings to import security templates to a GPO.

- **Secedit.sdb**

  This is a permanent system database used for policy propagation including a table of persistent settings for rollback purposes.

- **User databases**

  A user database is any database other than the system database created by administrators for the purposes of configuration or analysis of security.

- **.Inf Templates**

  These are text files that contain declarative security settings. They are loaded into a database before configuration or analysis. Group Policy security policies are stored in .inf files on the SYSVOL folder of domain controllers, where they are downloaded (by using file copy) and merged into the system database during policy propagation.

Security Settings policy processes and interactions

For a domain-joined computer, where Group Policy is administered, Security Settings are processed in conjunction with Group Policy. Not all settings for computer security are configurable.

Group Policy processing

When a computer starts and a user logs on, computer policy and user policy are applied according to the following sequence:

1. The network starts. Remote Procedure Call System Service (RPCSS) and Multiple Universal Naming Convention Provider (MUP) start.
2. An ordered list of Group Policy Objects is obtained for the computer. The list might depend on these factors:
   - Whether the computer is part of a domain and, therefore, subject to Group Policy through Active Directory.
   - The location of the computer in Active Directory.
   - Whether the list of Group Policy Objects has changed. If the list of Group Policy Objects has not changed, no processing is done.
3. Computer policy is applied. These are the settings under Computer Configuration from the gathered list. This is a synchronous process by default and occurs in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while computer policies are processed.
4. Startup scripts run. This is hidden and synchronous by default; each script must complete or time out before the next one starts. The default time-out is 600 seconds. You can use several policy settings to modify this behavior.
5. The user presses CTRL+ALT+DEL to log on.
6. After the user is validated, the user profile loads; it is governed by the policy settings that are in effect.
7. An ordered list of Group Policy Objects is obtained for the user. The list might depend on these factors:
   - Whether the user is part of a domain and, therefore, subject to Group Policy through Active Directory.
   - Whether loopback policy processing is enabled, and if so, the state (Merge or Replace) of the loopback policy setting.
   - The location of the user in Active Directory.
   - Whether the list of Group Policy Objects has changed. If the list of Group Policy Objects has not changed, no processing is done.
8. User policy is applied. These are the settings under User Configuration from the gathered list. This is synchronous by default and in the following order: local, site, domain, organizational unit, child organizational unit, and so on. No user interface appears while user policies are processed.

9. Logon scripts run. Group Policy−based logon scripts are hidden and asynchronous by default. The user object script runs last.
10. The operating system user interface that is prescribed by Group Policy appears.

Group Policy Objects storage

A Group Policy Object (GPO) is a virtual object that is identified by a Globally Unique Identifier (GUID) and stored at the domain level. The policy setting information of a GPO is stored in the following two locations:

- **Group Policy containers in Active Directory.**

  The Group Policy container is an Active Directory container that contains GPO properties, such as version information, GPO status, plus a list of other component settings.

- **Group Policy templates in a domain's system volume folder (SYSVOL).**

  The Group Policy template is a file system folder that includes policy data specified by .admx files, security settings, script files, and information about applications that are available for installation. The Group Policy template is located in the SYSVOL folder in the domain\Policies subfolder.

The **GROUP_POLICY_OBJECT** structure provides information about a GPO in a GPO list, including the version number of the GPO, a pointer to a string that indicates the Active Directory portion of the GPO, and a pointer to a string that specifies the path to the file system portion of the GPO.

Group Policy processing order

Group Policy settings are processed in the following order:

1. **Local Group Policy Object.**

   Each computer running a Windows operating system beginning with Windows XP has exactly one Group Policy Object that is stored locally.

2. **Site.**

   Any Group Policy Objects that have been linked to the site are processed next. Processing is synchronous and in an order that is specified by the administrator.

3. **Domain.**

   Processing of multiple domain-linked Group Policy Objects is synchronous and in an order specified by the administrator.

4. **Organizational units.**

   Group Policy Objects that are linked to the organizational unit that is highest in the Active Directory hierarchy are processed first, then Group Policy Objects that are linked to its child organizational unit, and so on. Finally, the Group Policy Objects that are linked to the organizational unit that contains the user or computer are processed.

At the level of each organizational unit in the Active Directory hierarchy, one, many, or no Group Policy Objects can be linked. If several Group Policy Objects are linked to an organizational unit, their processing is synchronous and in an order that is specified by the administrator.

This order means that the local Group Policy Object is processed first, and Group Policy Objects that are linked to the organizational unit of which the computer or user is a direct member are processed last, which overwrites the earlier Group Policy Objects.

This is the default processing order and administrators can specify exceptions to this order. A Group Policy Object that is linked to a site, domain, or organizational unit (not a local Group Policy Object) can be set to **Enforced** with respect to that site, domain, or organizational unit, so that none of its policy settings can be overridden. At any site, domain, or organizational unit, you can mark Group Policy inheritance selectively as **Block Inheritance**. Group Policy Object links that are set to **Enforced** are always applied, however, and they cannot be blocked.

Security Settings policy processing

In the context of Group Policy processing, Security Settings policy is processed in the following order.

1. During Group Policy processing, the Group Policy engine determines which security settings policies to apply.
2. If security settings policies exist in a GPO, Group Policy invokes the Security Settings client-side extension.
3. The Security Settings extension downloads the policy from the appropriate location such as a specific domain controller.
4. The Security Settings extension merges all security settings policies according to precedence rules. The processing is according to the Group Policy processing order of local, site, domain, and organizational unit (OU), as described earlier in the “Group Policy processing order” section. If multiple GPOs are in effect for a given computer and there are no conflicting policies, then the policies are cumulative and are merged.

   This example uses the Active Directory structure shown in the following figure. A given computer is a member of OU2, to which the **GroupMembershipPolGPO** GPO is linked. This computer is also subject to the **UserRightsPolGPO** GPO, which is linked to OU1, higher in the hierarchy. In this case, no conflicting policies exist so the computer receives all of the policies contained in both the **UserRightsPolGPO** and the **GroupMembershipPolGPO** GPOs.

Multiple GPOs and Merging of Security Policy

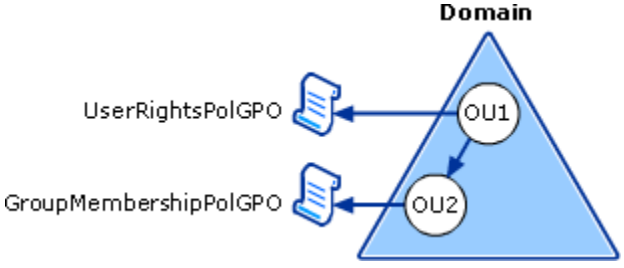

5. The resultant security policies are stored in Secedit.sdb, the security settings database. The security engine gets the security template files and imports them to Secedit.sdb.
6. The security settings policies are applied to computers.

The following figure illustrates the security settings policy processing.

Security Settings Policy Processing

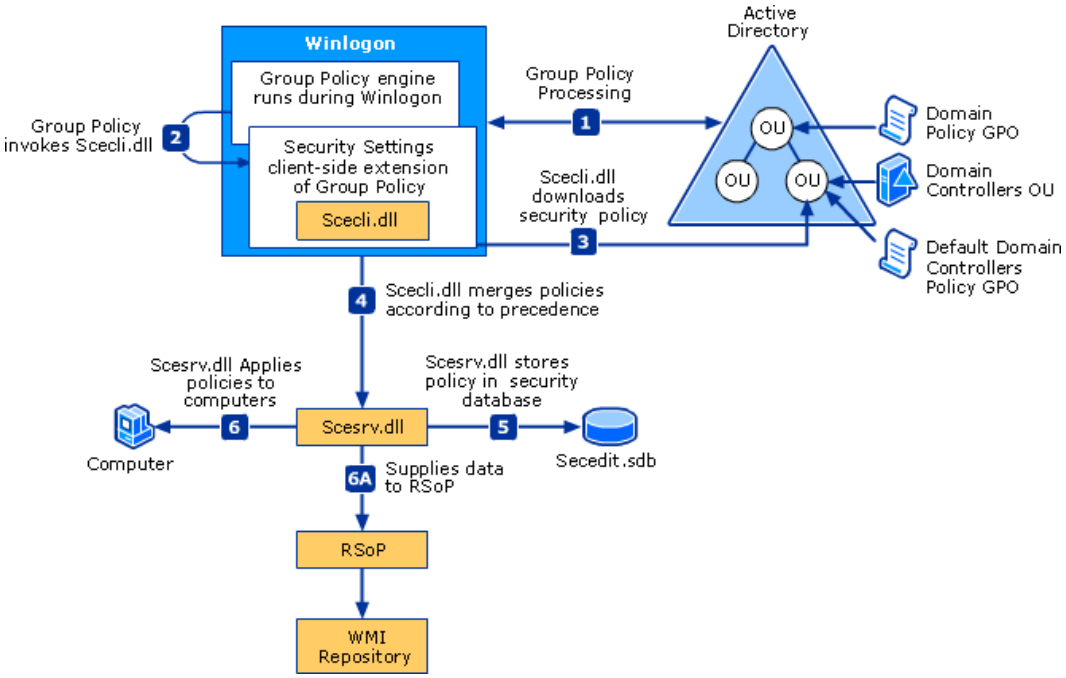

Merging of security policies on domain controllers

Password policies, Kerberos, and some security options are only merged from GPOs that are linked at the root level on the domain. This is done to keep those settings synchronized across all domain controllers in the domain. The following security options are merged:

- Network Security: Force logoff when logon hours expire
- Accounts: Administrator account status
- Accounts: Guest account status
- Accounts: Rename administrator account
- Accounts: Rename guest account

Another mechanism exists that allows security policy changes made by administrators by using net accounts to be merged into the Default Domain Policy GPO. User rights changes that are made by using Local Security Authority (LSA) APIs are filtered into the Default Domain Controllers Policy GPO.

Special considerations for domain controllers

If an application is installed on a primary domain controller (PDC) with operations master role (also known as flexible single master operations or FSMO) and the application makes changes to user rights or password policy, these changes must be communicated to ensure that synchronization across domain controllers occurs. Scesrv.dll receives a notification of any changes made to the security account manager (SAM) and LSA that need to be synchronized across domain controllers and then incorporates the changes into the Default Domain Controller Policy GPO by using Scecli.dll template modification APIs.

When security settings are applied

After you have edited the security settings policies, the settings are refreshed on the computers in the organizational unit linked to your Group Policy Object in the following instances:

- When a computer is restarted.
- Every 90 minutes on a workstation or server and every 5 minutes on a domain controller. This refresh interval is configurable.
- By default, Security Policy settings delivered by Group Policy are also applied every 16 hours (960 minutes) even if a GPO has not changed.

Persistence of Security Settings policy

Security Settings can persist even if a setting is no longer defined in the policy that originally applied it.

Security Settings might persist in the following cases:

- The setting has not been previously defined for the computer.
- The setting is for a registry security object.
- The settings are for a file system security object.

All settings applied through local policy or through a Group Policy Object are stored in a local database on your computer. Whenever a security setting is modified, the computer saves the security setting value to the local database, which retains a history of all the settings that have been applied to the computer. If a policy first defines a security setting and then no longer defines that setting, then the setting takes on the previous value in the database. If a previous value does not exist in the database then the setting does not revert to anything and remains defined as is. This behavior is sometimes referred to as “tattooing.”

Registry and file security settings will maintain the values applied through Group Policy until that setting is set to other values.

On domain controllers running Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003 all security settings persist.

Permissions required for policy to apply

Both Apply Group Policy and Read permissions are required to have the settings from a Group Policy Object apply to users or groups, and computers.

Filtering security policy

By default, all GPOs have Read and Apply Group Policy both Allowed for the Authenticated Users group. The Authenticated Users group includes both users and computers. Security settings policies are computer-based. To specify which client computers will or will not have a Group Policy Object applied to them, you can deny them either the Apply Group Policy or Read permission on that Group Policy Object. Changing these permissions allows you to limit the scope of the GPO to a specific set of computers within a site, domain, or OU.

Note

Do not use security policy filtering on a domain controller as this would prevent security policy from applying to it.

Migration of GPOs containing Security Settings

In some situations, a policy administrator might want to migrate GPOs from one domain environment to another environment. The two most common scenarios are test-to-production migration, and production-to-production migration. The GPO copying process has implications for some types of security settings.

Data for a single GPO is stored in multiple locations and in various formats; some data is contained in Active Directory and other data is stored on the SYSVOL share on the domain controllers. Certain policy data might be valid in one domain but might be invalid in the domain to which the GPO is being copied. For example, Security Identifiers (SIDs) stored in security policy settings are often domain-specific. So copying GPOs is not as simple as taking a folder and copying it from one computer to another.

The following security policies can contain security principals and might require some additional work to successfully move them from one domain to another.

- User rights assignment
- Restricted groups
- Services
- File system
- Registry
- The GPO DACL, if you choose to preserve it during a copy operation

To ensure that data is copied correctly, administrators can use Group Policy Management Console (GPMC). When migrating a GPO from one domain to another, GPMC ensures that all relevant data is properly copied. GPMC also offers migration tables, which can be used to update domain-specific data to new values as part of the migration process. GPMC hides much of the complexity involved in the migrating GPO operations, and it provides simple and reliable mechanisms for performing operations such as copy and backup of GPOs.

Additional resources

The following table lists relevant and supporting documentation for the Security Settings extensions.

Technology content	Resources
Group Policy	Group Policy Overview
Security Tools	Security Tools
Security settings	Security Policy Settings Reference Threats and Countermeasures Guide: Security Settings in Windows Server 2008 R2 and Windows 7 Threats and Countermeasures Guide: Security Settings in Windows Server 2008 and Windows Vista
Active Directory	Active Directory Domain Services Overview
Sign on technologies	Interactive Logon Technical Reference (2003)
Troubleshooting	Group Policy Infrastructure AD DS Windows Logon and Initialization TechNet Wiki: Troubleshooting Portal
Administer and configure	Administer Security Policy Settings Group Policy Management Console Local Group Policy Editor

See Also

Security Policy Settings Overview

Security Policy Settings Reference

This reference of security settings provides information about how to implement and manage security policies, including setting options and security considerations. Applicable operating system versions are listed on each policy setting description.

Overview

This reference focuses on those settings that are considered security settings. This reference examines only the settings and features in the Windows operating systems that can help organizations secure their enterprises against malicious software threats. Management features and those security features that are not configurable by administrators are not described in this reference.

Each policy setting described contains referential content such as a detailed explanation of the settings, best practices, default settings, differences between operating system versions, policy management considerations, and security considerations that include a discussion of vulnerability, countermeasures, and potential impact of those countermeasures.

For each grouping of policies, an overview topic describes the general use and implementation of those policies and provides links to the related topics.

The information that is provided within this reference should help you and members of your organization understand the implementation and countermeasures for security policies that are available in the named versions of the Windows operating system.

Contents of this guide

This guide consists of the following sections that contain descriptions of the settings that you should consider while planning the security policy for your organization.

Account Policies	This section discusses the settings that are applied at the domain level: password policies, account lockout policies, and Kerberos authentication protocol policies.
Audit Policy	This section provides information about the security audit policy settings under **Security Settings\Local Policies\Audit Policy** that provide broad security audit capabilities for client computers and servers that cannot use advanced security audit policy settings.
Security Options	This section provides guidance about how to configure the variety of settings specific to security on the local computer: • Accounts • Audit • DCOM • Devices • Domain controller • Domain member

	• Interactive logon • Microsoft network client • Microsoft network server • Network access • Network security • Recovery console • Shutdown • System cryptography • System objects • System settings • User Account Control
Advanced Audit Policy Configuration	This section provides information about advanced security audit policy settings, which allow administrators to use audit policies to help monitor and enforce business rules.
User Rights Assignment	This section provides information about the User Rights Assignment security policy settings that are available in the Windows operating system.

Additional resources

- Security Audit Policy Reference

  This reference for the IT professional provides information about the auditing settings available in the Windows operating system beginning with Windows Server 2003 and the audit events that they generate.

- Threats and Countermeasures Guide: Security Settings in Windows Server 2008 and Windows VistaThreats and Countermeasures Guide: Security Settings in Windows Server 2008 and Windows Vista

  This guide focuses on the security settings available in Windows Server 2008 and Windows Vista that can help you address specific computer security risks in an enterprise environment.

- Threats and Countermeasures Guide: Security Settings in Windows Server 2008 R2 and Windows 7

  This guide focuses on the security settings available in Windows Server 2008 R2 and Windows 7 that can help you address specific computer security risks in an enterprise environment.

- Security Policy Settings

  This reference documents security policy settings in Windows Server 2003.

- Security Compliance Manager

  The Security Baselines are part of the Microsoft Security Compliance Manager available as a downloadable tool that can help you plan, deploy, and monitor the security baselines of computers running the supported operating system versions designated in the **Applies To** list at the beginning of this topic.

How to Configure Security Policy Settings

This procedural topic for the IT professional describes steps to configure a security policy setting on the local computer, on a domain-joined computer, and on a domain controller.

This topic pertains to the versions of Windows designated in the **Applies To** list above. Some of the user interface elements that are described in this topic might differ from version to version.

You must have Administrators rights on the local computer, or you must have the appropriate permissions to update a Group Policy Object (GPO) on the domain controller to perform these procedures.

When a local setting is inaccessible, it indicates that a GPO currently controls that setting.

In this topic

- To configure a setting for your local computer
- To configure a setting for computer that is joined to a domain
- To configure a setting for a domain controller

To configure a setting for your local computer

1. To open Local Security Policy, on the **Start** screen, type, **secpol.msc**.
2. Navigate the console tree to Local Computer Policy\Windows Settings\Security Settings
3. Under **Security Settings** of the console tree, do one of the following:
   - Click **Account Policies** to edit the **Password Policy** or **Account Lockout Policy**.
   - Click **Local Policies** to edit an **Audit Policy**, a **User Rights Assignment**, or **Security Options**.
4. When you find the policy setting in the details pane, double-click the security policy that you want to modify.
5. Modify the security policy setting, and then click **OK**.

   Notes
   - Some security policy settings require that the computer be restarted before the setting takes effect.
   - Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

To configure a setting for computer that is joined to a domain

The following procedure describes how to configure a security policy setting for a Group Policy Object when you are on a workstation or server that is joined to a domain.

You must have the appropriate permissions to install and use the Microsoft Management Console (MMC), and to update a Group Policy Object (GPO) on the domain controller to perform these procedures.

1. To open the MMC and add the Group Policy Object Editor, on the **Start** screen, type **mmc.msc**.
2. On the **File** menu of the MMC, click **Add/Remove snap-in**, and then click **Add**.
3. In **Add Standalone Snap-in**, double-click **Group Policy Object Editor**.
4. In **Select Group Policy Object**, click **Browse**, browse to the GPO you would like to modify, and then click **Finish**.
5. Click **Close**, and then click **OK**.

   This procedure added the snap-in to the MMC.
6. In the console tree, locate *GroupPolicyObject [ComputerName]* Policy, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.
7. Do one of the following:
   - Click **Account Policies** to edit the **Password Policy** or **Account Lockout Policy**.
   - Click **Local Policies** to edit an **Audit Policy**, a **User Rights Assignment**, or **Security Options**.
   - Click **Event Log** to edit event log settings.
8. In the details pane, double-click the security policy setting that you want to modify.

   Note

   If this security policy has not yet been defined, select the **Define these policy settings** check box.
9. Modify the security policy setting and then click **OK**.

To configure a setting for a domain controller

The following procedure describes how to configure a security policy setting for only a domain controller (from the domain controller).

1. To open the domain controller security policy, in the console tree, locate *GroupPolicyObject [ComputerName]* Policy, click **Computer Configuration**, click **Windows Settings**, and then click **Security Settings**.
2. Do one of the following:
   - Double-click **Account Policies** to edit the **Password Policy**, **Account Lockout Policy**, or **Kerberos Policy**.
   - Click **Local Policies** to edit the **Audit Policy**, a **User Rights Assignment**, or **Security Options**.
   - Click **Event Log** to edit event log settings.

3. In the details pane, double-click the security policy that you want to modify.

   Note
   If this security policy has not yet been defined, select the **Define these policy settings** check box.

4. Modify the security policy setting, and then click **OK**.

Important

- Always test a newly created policy in a test organizational unit before you apply it to your network.
- When you change a security setting through a GPO and click **OK**, that setting will take effect the next time you refresh the settings.

See also

Security Policy Settings Reference

Account Policies

This security policy setting reference provides an overview of account policies in Windows and provides links to policy descriptions.

The account policies settings include the following:

Password Policy

Account Lockout Policy

Kerberos Policy

For information about setting security policies, see **How to Configure Security Policy Settings**.

All account policies settings applied by using Group Policy are applied at the domain level. Default values are present in the built-in default domain controller policy for Password Policy settings, Account Lockout Policy settings, and Kerberos Policy settings. The domain account policy becomes the default local account policy of any Windows-based computer that is a member of the domain. If these policies are set at any level below the domain level in Active Directory Domain Services (AD DS), they affect only local accounts on member servers.

Note
Each domain can have only one account policy. The account policy must be defined in the default domain policy or in a new policy that is linked to the root of the domain and given precedence over the default domain policy, which is enforced by the domain controllers in the domain. These domain-wide account policy settings (Password Policy, Account Lockout Policy, and Kerberos Policy) are enforced by the domain controllers in the domain; therefore, domain controllers always retrieve the values of these account policy settings from the default domain policy Group Policy Object (GPO).

The only exception is when another account policy is defined for an organizational unit (OU). The account policy settings for the OU affect the local policy on any computers that are contained in the OU. For example, if an OU policy defines a maximum password age that differs from the domain-level account policy, the OU policy will be applied and enforced only when users log on to the local computer. The default local computer policies apply only to computers that are in a workgroup or in a domain where neither an OU account policy nor a domain policy applies.

Password Policy

This security policy reference topic for the IT professional provides an overview of password policies for Windows and links to information for each policy setting.

In many operating systems, the most common method to authenticate a user's identity is to use a secret passphrase or password. A secure network environment requires all users to use strong passwords, which have at least eight characters and include a combination of letters, numbers, and symbols. These passwords help prevent the compromise of user accounts and administrative accounts by unauthorized users who use manual methods or automated tools to guess weak passwords. Strong passwords that are changed regularly reduce the likelihood of a successful password attack.

Introduced in Windows Server 2008 R2 and Windows Server 2008, Windows supports fine-grained password policies. This feature provides organizations with a way to define different password and account lockout policies for different sets of users in a domain. In Windows 2000 Server and Windows Server 2003 Active Directory domains, only one password policy and account lockout policy could be applied to all users in the domain. Fine-grained password policies apply only to user objects (or inetOrgPerson objects if they are used instead of user objects) and global security groups.

To apply a fine-grained password policy to users of an OU, you can use a shadow group. A shadow group is a global security group that is logically mapped to an OU to enforce a fine-grained password policy. You add users of the OU as members of the newly created shadow group and then apply the fine-grained password policy to this shadow group. You can create additional shadow groups for other OUs as needed. If you move a user from one OU to another, you must update the membership of the corresponding shadow groups.

Fine-grained password policies include attributes for all the settings that can be defined in the default domain policy (except Kerberos settings) in addition to account lockout settings. When you specify a fine-grained password policy, you must specify all of these settings. By default, only members of the Domain Admins group can set fine-grained password policies. However, you can also delegate the ability to set these policies to other users. The domain must be running at least Windows Server 2008 R2 or Windows Server 2008 to use fine-grained password policies. Fine-grained password policies cannot be applied to an organizational unit (OU) directly.

Fine-grained password policies do not interfere with custom password filters that you might use in the same domain. Organizations that have deployed custom password filters to domain controllers running Windows 2000 Server or Windows Server 2003 can continue to use those

password filters to enforce additional restrictions for passwords. For more information about fine-grained password policies, see AD DS: Fine-Grained Password Policies.

You can enforce the use of strong passwords through an appropriate password policy. There are password policy settings that control the complexity and lifetime of passwords, such as the **Passwords must meet complexity requirements** policy setting.

You can configure the password policy settings in the following location by using the Group Policy Management Console on your domain controller:

Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy

If individual groups require distinct password policies, these groups should be separated into another domain or forest, based on additional requirements.

For information about setting security policies, see How to Configure Security Policy Settings.

The following topics provide a discussion of password policy implementation and best practices considerations, policy location, default values for the server type or GPO, relevant differences in operating system versions, security considerations (including the possible vulnerabilities of each setting), countermeasures that you can take, and the potential impact for each setting.

- Enforce password history
- Maximum password age
- Minimum password age
- Minimum password length
- Password must meet complexity requirements
- Store passwords using reversible encryption

Enforce password history

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Enforce password history** policy setting determines the number of unique new passwords that must be associated with a user account before an old password can be reused.

Password reuse is an important concern in any organization. Many users want to reuse the same password for their account over a long period of time. The longer the same password is used for a particular account, the greater the chance that an attacker will be able to determine the password through brute force attacks. If users are required to change their password, but they can reuse an old password, the effectiveness of a good password policy is greatly reduced.

Specifying a low number for **Enforce password history** allows users to continually use the same small number of passwords repeatedly. If you do not also set **Minimum password age**, users can change their password as many times in a row as necessary to reuse their original password.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- User-specified number from 0 through 24
- Not defined

Best practices

- Set **Enforce password history** to 24. This will help mitigate vulnerabilities that are caused by password reuse.
- Set **Maximum password age** to expire passwords between 60 and 90 days. Try to expire the passwords between major business cycles to prevent work loss.
- Configure **Minimum password age** so that you do not allow passwords to be changed immediately.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default domain policy	24 passwords remembered
Default domain controller policy	Not defined
Stand-alone server default settings	0 passwords remembered
Domain controller effective default settings	24 passwords remembered
Member server effective default settings	24 passwords remembered
Effective GPO default settings on client computers	24 passwords remembered

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The longer a user uses the same password, the greater the chance that an attacker can determine the password through brute force attacks. Also, any accounts that may have been compromised remain exploitable for as long as the password is left unchanged. If password changes are required but password reuse is not prevented, or if users continually reuse a small number of passwords, the effectiveness of a good password policy is greatly reduced.

If you specify a low number for this policy setting, users can use the same small number of passwords repeatedly. If you do not also configure the **Minimum password age** policy setting, users might repeatedly change their passwords until they can reuse their original password.

Note

After an account has been compromised, a simple password reset might not be enough to restrict a malicious user because the malicious user might have modified the user's environment so that the password is changed back to a known value automatically at a certain time. If an account has been compromised, it is best to delete the account and assign the user a new account after all affected systems have been restored to normal operations and verified that they are no longer compromised.

Countermeasure

Configure the **Enforce password history** policy setting to 24 (the maximum setting) to help minimize the number of vulnerabilities that are caused by password reuse.

For this policy setting to be effective, you should also configure effective values for the **Minimum password age** and **Maximum password age** policy settings.

Potential impact

The major impact of configuring the **Enforce password history** setting to 24 is that users must create a new password every time they are required to change their old one. If users are required to change their passwords to new unique values, there is an increased risk of users who write their passwords somewhere so that they do not forget them. Another risk is that users may create passwords that change incrementally (for example, password01, password02, and so on) to facilitate memorization, but this makes them easier for an attacker to guess. Also, an excessively low value for the **Maximum password age** policy setting is likely to increase administrative overhead because users who forget their passwords might ask the Help Desk to reset them frequently.

See Also

Password Policy

Maximum password age

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Maximum password age** policy setting determines the period of time (in days) that a password can be used before the system requires the user to change it. You can set passwords to expire after a number of days between 1 and 999, or you can specify that passwords never expire by setting the number of days to 0. If **Maximum password age** is between 1 and 999 days, the minimum password age must be less than the maximum password age. If **Maximum password age** is set to 0, **Minimum password age** can be any value between 0 and 998 days.

Note
Setting **Maximum password age** to -1 is equivalent to 0, which means it never expires. Setting it to any other negative number is equivalent to setting it to **Not Defined**.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- User-specified number of days between 0 and 999
- Not defined

Best practices

Set **Maximum password age** to a value between 30 and 90 days, depending on your environment. This way, an attacker has a limited amount of time in which to compromise a user's password and have access to your network resources.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	42 days
Default domain controller policy	Not defined
Stand-alone server default settings	42 days
Domain controller effective default settings	42 days
Member server effective default settings	42 days
Effective GPO default settings on client computers	42 days

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The longer a password exists, the higher the likelihood that it will be compromised by a brute force attack, by an attacker gaining general knowledge about the user, or by the user sharing the password. Configuring the **Maximum password age** policy setting to 0 so that users are never required to change their passwords is a major security risk because that allows a compromised password to be used by the malicious user for as long as the valid user is authorized access.

Countermeasure

Configure the **Maximum password age** policy setting to a value that is suitable for your organization's business requirements.

Potential impact

If the **Maximum password age** policy setting is too low, users are required to change their passwords very often. Such a configuration can reduce security in the organization because users might keep their passwords in an unsecured location or lose them. If the value for this policy setting is too high, the level of security within an organization is reduced because it allows potential attackers more time in which to discover user passwords or to use compromised accounts.

See Also

Password Policy

Minimum password age

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Minimum password age** policy setting determines the period of time (in days) that a password can be used before the system requires the user to change it. You can set passwords to expire after a number of days between 1 and 999, or you can specify that passwords never expire by setting the number of days to 0. If **Maximum password age** is between 1 and 999

days, the minimum password age must be less than the maximum password age. If **Maximum password age** is set to 0, **Minimum password age** can be any value between 0 and 998 days.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- User-specified number of days between 0 and 998
- Not defined

Best practices

Set **Minimum password age** to a value of 2 days. Setting the number of days to 0 allows immediate password changes, which is not recommended.

If an administrator sets a password for a user and wants that user to change the administrator-defined password, the administrator must select the **User must change password at next logon** check box. Otherwise, the user will not be able to change the password until the number of days specified by **Minimum password age**.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	1 day
Default domain controller policy	Not defined
Stand-alone server default settings	0 days
Domain controller effective default settings	1 day
Member server effective default settings	1 day
Effective GPO default settings on client computers	1 day

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users may have favorite passwords that they like to use because they are easy to remember and they believe that their password choice is secure from compromise. Unfortunately, passwords can be compromised and if an attacker is targeting a specific individual user account, with knowledge of data about that user, reuse of old passwords can cause a security breach.

To address password reuse, you must use a combination of security settings. Using this policy setting with the **Enforce password history** policy setting prevents the easy reuse of old passwords. For example, if you configure the **Enforce password history** policy setting to ensure that users cannot reuse any of their last 12 passwords, but you do not configure the **Minimum password age** policy setting to a number that is greater than 0, users could change their password 13 times in a few minutes and reuse their original password. You must configure this policy setting to a number that is greater than 0 for the **Enforce password history** policy setting to be effective.

Countermeasure

Configure the **Minimum password age** policy setting to a value of at least 2 days. Users should know about this limitation and contact the Help Desk if they need to change their password during that two-day period. If you configure the number of days to 0, immediate password changes would be allowed, which we do not recommend.

Potential impact

If an administrator sets a password for a user but wants that user to change the password when the user first logs on, the administrator must select the **User must change password at next logon** check box, or the user cannot change the password until the next day.

See Also

Password Policy

Minimum password length

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Minimum password length** policy setting determines the least number of characters that can make up a password for a user account. You can set a value of between 1 and 14 characters, or you can establish that no password is required by setting the number of characters to 0.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- User-specified number of characters between 0 and 14
- Not defined

Best practices

Set **Minimum password length** to at least a value of 8. If the number of characters is set to 0, no password is required. In most environments, an eight-character password is recommended because it is long enough to provide adequate security and still short enough for users to easily remember. This value will help provide adequate defense against a brute force attack. Adding complexity requirements will help reduce the possibility of a dictionary attack. For more information, see Password must meet complexity requirements.

Permitting short passwords reduces security because short passwords can be easily broken with tools that perform dictionary or brute force attacks against the passwords. Requiring very long passwords can result in mistyped passwords that might cause an account lockout and subsequently increase the volume of Help Desk calls.

In addition, requiring extremely long passwords can actually decrease the security of an organization because users might be more likely to write down their passwords to avoid forgetting

them. However, if users are taught that they can use passphrases (sentences such as "I want to drink a $5 milkshake"), they should be much more likely to remember.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	7 characters
Default domain controller policy	Not defined
Stand-alone server default settings	0 characters
Domain controller effective default settings	7 characters
Member server effective default settings	7 characters
Effective GPO default settings on client computers	0 characters

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Types of password attacks include dictionary attacks (which attempt to use common words and phrases) and brute force attacks (which try every possible combination of characters). Also, attackers sometimes try to obtain the account database so they can use tools to discover the accounts and passwords.

Countermeasure

Configure the **Minimum password length** policy setting to a value of 8 or more. If the number of characters is set to 0, no password will be required.

In most environments, we recommend an eight-character password because it is long enough to provide adequate security, but not too difficult for users to easily remember. This configuration provides adequate defense against a brute force attack. Using the **Passwords must meet complexity requirements** policy setting in addition to the **Minimum password length** setting helps reduce the possibility of a dictionary attack.

Note

Some jurisdictions have established legal requirements for password length as part of establishing computer security regulations.

Potential impact

Requirements for extremely long passwords can actually decrease the security of an organization because users might leave the information in an unsecured location or lose it. If very long passwords are required, mistyped passwords could cause account lockouts and increase the volume of Help Desk calls. If your organization has issues with forgotten passwords due to password length requirements, consider teaching your users about passphrases, which are often easier to remember and, due to the larger number of character combinations, much harder to discover.

Note

Older versions of Windows (such as Windows 98 and Windows NT 4.0) do not support passwords that are longer than 14 characters. Computers that run these older operating systems are unable to authenticate with computers or domains that use accounts that require long passwords.

See Also

Password Policy

Password must meet complexity requirements

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Passwords must meet complexity requirements** policy setting determines whether passwords must meet a series of guidelines that are considered important for a strong password. Enabling this policy setting requires passwords to meet the following requirements:

1. Passwords may not contain the user's samAccountName (Account Name) value or entire displayName (Full Name value). Both checks are not case sensitive.

   The samAccountName is checked in its entirety only to determine whether it is part of the password. If the samAccountName is less than three characters long, this check is skipped.

   The displayName is parsed for delimiters: commas, periods, dashes or hyphens, underscores, spaces, pound signs, and tabs. If any of these delimiters are found, the displayName is split and all parsed sections (tokens) are confirmed to not be included in the password. Tokens that are less than three characters are ignored, and substrings of the tokens are not checked. For example, the name "Erin M. Hagens" is split into three tokens: "Erin", "M", and "Hagens". Because the second token is only one character long, it is ignored. Therefore, this user could not have a password that included either "erin" or "hagens" as a substring anywhere in the password.
2. The password contains characters from three of the following categories:
   - Uppercase letters of European languages (A through Z, with diacritic marks, Greek and Cyrillic characters)
   - Lowercase letters of European languages (a through z, sharp-s, with diacritic marks, Greek and Cyrillic characters)
   - Base 10 digits (0 through 9)
   - Non-alphanumeric characters (special characters) (for example, !, $, #, %)
   - Any Unicode character that is categorized as an alphabetic character but is not uppercase or lowercase. This includes Unicode characters from Asian languages.

Complexity requirements are enforced when passwords are changed or created.

The rules that are included in the Windows Server password complexity requirements are part of Passfilt.dll, and they cannot be directly modified.

Enabling the default Passfilt.dll may cause some additional Help Desk calls for locked-out accounts because users might not be used to having passwords that contain characters other than those found in the alphabet. However, this policy setting is liberal enough that all users should be able to abide by the requirements with a minor learning curve.

Additional settings that can be included in a custom Passfilt.dll are the use of non–upper-row characters. Upper-row characters are those that are typed by holding down the SHIFT key and typing any of the digits from 1 through 10.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

Set **Passwords must meet complexity requirements** to Enabled. This policy setting, combined with a minimum password length of 8, ensures that there are at least 218,340,105,584,896 different possibilities for a single password. This makes a brute force attack difficult, but still not impossible.

The use of ALT key character combinations can greatly enhance the complexity of a password. However, requiring all users in an organization to adhere to such stringent password requirements can result in unhappy users and an extremely busy Help Desk. Consider implementing a requirement in your organization to use ALT characters in the range from 0128 through 0159 as part of all administrator passwords. (ALT characters outside of this range can represent standard alphanumeric characters that do not add additional complexity to the password.)

Passwords that contain only alphanumeric characters are easy to compromise by using publicly available tools. To prevent this, passwords should contain additional characters and meet complexity requirements.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Enabled
Default domain controller policy	Enabled

Server type or Group Policy Object (GPO)	Default value
Stand-alone server default settings	Disabled
Domain controller effective default settings	Enabled
Member server effective default settings	Enabled
Effective GPO default settings on client computers	Disabled

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Passwords that contain only alphanumeric characters are extremely easy to discover with several publicly available tools.

Countermeasure

Configure the **Passwords must meet complexity requirements** policy setting to Enabled and advise users to use a variety of characters in their passwords.

When combined with a **Minimum password length** of 8, this policy setting ensures that the number of different possibilities for a single password is so great that it is difficult (but not impossible) for a brute force attack to succeed. (If the **Minimum password length** policy setting is increased, the average amount of time necessary for a successful attack also increases.)

Potential impact

If the default password complexity configuration is retained, additional Help Desk calls for locked-out accounts could occur because users might not be accustomed to passwords that contain non-alphabetical characters, or they might have problems entering passwords that contain accented characters or symbols on keyboards with different layouts. However, all users should be able to comply with the complexity requirement with minimal difficulty.

If your organization has more stringent security requirements, you can create a custom version of the Passfilt.dll file that allows the use of arbitrarily complex password strength rules. For example,

a custom password filter might require the use of non-upper-row symbols. (Upper-row symbols are those that require you to press and hold the SHIFT key and then press any of the digits between 1 and 0.) A custom password filter might also perform a dictionary check to verify that the proposed password does not contain common dictionary words or fragments.

The use of ALT key character combinations can greatly enhance the complexity of a password. However, such stringent password requirements can result in additional Help Desk requests. Alternatively, your organization could consider a requirement for all administrator passwords to use ALT characters in the 0128–0159 range. (ALT characters outside of this range can represent standard alphanumeric characters that would not add additional complexity to the password.)

See Also

Password Policy

Store passwords using reversible encryption

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Store password using reversible encryption** policy setting provides support for applications that use protocols that require the user's password for authentication. Storing encrypted passwords in a way that is reversible means that the encrypted passwords can be decrypted. A knowledgeable attacker who is able to break this encryption can then log on to network resources by using the compromised account. For this reason, never enable **Store password using reversible encryption** for all users in the domain unless application requirements outweigh the need to protect password information.

If you use the Challenge Handshake Authentication Protocol (CHAP) through remote access or Internet Authentication Services (IAS), you must enable this policy setting. CHAP is an authentication protocol that is used by remote access and network connections. Digest Authentication in Internet Information Services (IIS) also requires that you enable this policy setting.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

Set the value for **Store password using reversible encryption** to Disabled. If you use CHAP through remote access or IAS, or Digest Authentication in IIS, you must set this value to Enabled. This presents a security risk when you apply the setting by using Group Policy on a user-by-user basis because it requires opening the appropriate user account object in Active Directory Users and Computers.

Note
Do not enable this policy setting unless business requirements outweigh the need to protect password information.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Disabled
Default domain controller policy	Disabled
Stand-alone server default settings	Disabled
Domain controller effective default settings	Disabled
Member server effective default settings	Disabled
Effective GPO default settings on client computers	Disabled

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Enabling this policy setting allows the operating system to store passwords in a format that can weaken your overall security.

Countermeasure

Disable the **Store password using reversible encryption** policy setting.

Potential impact

If your organization uses CHAP through remote access or IAS, or Digest Authentication in IIS, you must configure this policy setting to Enabled. This presents a security risk when you apply the setting through Group Policy on a user-by-user basis because it requires the appropriate user account object to be opened in Active Directory Users and Computers.

See Also

Password Policy

Account Lockout Policy

This topic for the IT professional describes the Account Lockout Policy settings and links to information about each policy setting.

Overview of Account Lockout Policy settings

Someone who attempts to use more than a few unsuccessful passwords while trying to log on to your system might be a malicious user who is attempting to determine an account password by trial and error. Beginning with Windows Server 2003, Windows domain controllers keep track of logon attempts, and domain controllers can be configured to respond to this type of potential attack by disabling the account for a preset period of time. Account Lockout Policy settings control the threshold for this response and the actions to be taken after the threshold is reached. The Account Lockout Policy settings can be configured in the following location in the Group Policy Management Console: **Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy**.

For information about setting security policies, see How to Configure Security Policy Settings.

The following topics provide a discussion of each policy setting's implementation and best practices considerations, policy location, default values for the server type or Group Policy Object (GPO), relevant differences in operating system versions, and security considerations (including the possible vulnerabilities of each policy setting), countermeasures that you can implement, and the potential impact of implementing the countermeasures.

- Account lockout duration

- Account lockout threshold
- Reset account lockout counter after

Account lockout duration

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Account lockout duration** policy setting determines the number of minutes that a locked-out account remains locked out before automatically becoming unlocked. The available range is from 1 through 99,999 minutes. A value of 0 specifies that the account will be locked out until an administrator explicitly unlocks it. If **Account lockout threshold** is set to a number greater than zero, **Account lockout duration** must be greater than or equal to the value of **Reset account lockout counter after**.

This policy setting is dependent on the **Account lockout threshold** policy setting that is defined, and it must be greater than or equal to the value specified for the **Reset lockout counter after** policy setting.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- A user-defined number of minutes from 0 through 99,999
- Not defined

If **Account lockout threshold** is configured, after the specified number of failed attempts, the account will be locked out. If th **Account lockout duration** is set to 0, the account will remain locked until an administrator unlocks it manually.

It is advisable to set **Account lockout duration** to approximately 30 minutes. To specify that the account will never be locked out, set the value to 0. To configure the value for this policy setting so that it never automatically unlocks the account might seem like a good idea; however, doing so can increase the number of requests that your organization's Help Desk receives to unlock accounts that were locked by mistake.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not applicable
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Security considerations

More than a few unsuccessful password submissions during an attempt to log on to a computer might represent an attacker's attempts to determine an account password by trial and error. The Windows and Windows Server operating systems can track logon attempts, and you can configure the operating system to disable the account for a preset period of time after a specified number of failed attempts. Account lockout policy settings control the threshold for this response and what action to take after the threshold is reached.

Vulnerability

A denial-of-service (DoS) condition can be created if an attacker abuses the **Account lockout threshold** policy setting and repeatedly attempts to log on with a specific account. After you configure the **Account lockout threshold** policy setting, the account will be locked out after the specified number of failed attempts. If you configure the **Account lockout duration** policy setting to 0, the account remains locked until an administrator unlocks it manually.

Countermeasure

Configure the **Account lockout duration** policy setting to an appropriate value for your environment. To specify that the account will remain locked until an administrator manually unlocks it, configure the value to 0. When the **Account lockout duration** policy setting is configured to a nonzero value, automated attempts to guess account passwords are delayed for

this interval before resuming attempts against a specific account. Using this setting in combination with the **Account lockout threshold** policy setting makes automated password guessing attempts more difficult.

Potential impact

Configuring the **Account lockout duration** policy setting to 0 so that accounts cannot be automatically unlocked can increase the number of requests that your organization's Help Desk receives to unlock accounts that were locked by mistake.

See Also

Account Lockout Policy

Account lockout threshold

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Account lockout threshold** policy setting determines the number of failed logon attempts that will cause a user account to be locked out. A locked-out account cannot be used until it is reset by an administrator or until the number of minutes specified by **Account lockout duration** expires. You can set a value from 1 through 999 failed logon attempts, or you can specify that the account will never be locked out by setting the value to 0. If **Account lockout threshold** is set to a number greater than zero, **Account lockout duration** must be greater than or equal to the value of **Reset account lockout counter after**.

Failed password attempts on workstations or member servers that have been locked by using CTRL+ALT+DELETE or password-protected screen savers do not count as failed logon attempts unless **Interactive logon: Require Domain Controller authentication to unlock workstation** is set to Enabled. If **Interactive logon: Require Domain Controller authentication to unlock workstation** is enabled, repeated failed password attempts to unlock the workstation will count against the account lockout threshold.

Brute force password attacks can be automated to try thousands or even millions of password combinations for any or all user accounts. Limiting the number of failed logons that can be performed nearly eliminates the effectiveness of such attacks.

However, it is important to note that a denial-of-service attack could be performed on a domain that has an account lockout threshold configured. A malicious user could programmatically attempt a series of password attacks against all users in the organization. If the number of attempts is greater than the value of **Account lockout threshold**, the attacker could potentially lock out every account.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- A user-defined number from 0 through 999
- Not defined

Because vulnerabilities can exist when this value is configured and when it is not, any organization should weigh their identified threats and the risks that they are trying to mitigate. There are two options to consider for this policy setting:

- Set **Account lockout threshold** to 0. This ensures that accounts will not be locked out. This setting will prevent a denial-of-service attack that intentionally locks out all or some accounts. In addition, this setting helps reduce Help Desk calls because users cannot accidentally lock themselves out of their accounts.

  Because it will not prevent a brute force attack, a value of 0 should only be chosen if both of the following criteria are explicitly met:

  - Password Policy settings force all users to have complex passwords made up of eight or more characters.
  - A robust auditing mechanism is in place to alert administrators when a series of failed logons are occurring in the environment.

  If these criteria cannot be met, set **Account lockout threshold** to a high enough value that users can accidentally mistype their password several times before they are locked out of their account, but ensure that a brute-force password attack would still lock out the account. It is advisable to specify a value of 50 invalid logon attempts. Keep in mind, however, that although this setting can reduce the number of Help Desk calls by reducing the number of user lockouts, it cannot prevent a denial-of-service attack.

Best practices

The threshold that you select is a balance between operational efficiency and security, and it depends on your organization's risk level. To allow for user error and thwart brute force attacks, a setting above 4 and below 10 could be an acceptable starting point for your organization.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	0 invalid login attempts
Default domain controller policy	Not defined
Stand-alone server default settings	0 invalid login attempts
Domain controller effective default settings	0 invalid login attempts
Member server effective default settings	0 invalid login attempts
Effective GPO default settings on client computers	0 invalid login attempts

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Online brute force password attacks can use automated methods to try millions of password combinations for any user account. The effectiveness of such attacks can be almost eliminated if you limit the number of failed logons that can be performed.

However, a denial-of-service (DoS) attack could be performed on a domain that has an account lockout threshold configured. An attacker could programmatically attempt a series of password attacks against all users in the organization. If the number of attempts is greater than the account lockout threshold, the attacker might be able to lock out every account without needing any special privileges or being authenticated in the network.

Note
Offline password attacks are not countered by this policy setting.

Countermeasure

Because vulnerabilities can exist when this value is configured and when it is not configured, two distinct countermeasures are defined. Any organization should weigh the choice between the two, based on their identified threats and the risks that they want to mitigate. The two countermeasure options are:

- Configure the **Account lockout threshold** setting to 0. This configuration ensures that accounts will not be locked out, and it will prevent a DoS attack that intentionally attempts to lock out accounts. This configuration also helps reduce Help Desk calls because users cannot accidentally lock themselves out of their accounts. Because it does not prevent a brute force attack, this configuration should be chosen only if both of the following criteria are explicitly met:
  - The password policy requires all users to have complex passwords of 8 or more characters.
  - A robust audit mechanism is in place to alert administrators when a series of failed logons occur in the environment.
- Configure the **Account lockout threshold** policy setting to a sufficiently high value to provide users with the ability to accidentally mistype their password several times before the account is locked, but ensure that a brute force password attack still locks the account. A good recommendation for such a configuration is 50 invalid logon attempts, which prevents accidental account lockouts and reduces the number of Help Desk calls, but does not prevent a DoS attack. We recommend this option if your organization cannot implement complex password requirements and an audit policy that alerts administrators to a series of failed logon attempts. Using this type of policy must be accompanied by a process to unlock locked accounts that can be implemented whenever needed to help mitigate massive lockouts caused by an attack on your systems.

Potential impact

If this policy setting is enabled, a locked-out account is not usable until it is reset by an administrator or until the account lockout duration expires. Enabling this setting will likely generate a number of additional Help Desk calls.

If you configure the **Account lockout threshold** policy setting to 0, there is a possibility that an attacker's attempt to discover passwords with a brute force password attack might go undetected if a robust audit mechanism is not in place. If you configure this policy setting to a number greater than 0, an attacker can easily lock out any accounts for which the account name is known. This is especially dangerous considering that no privileges other than access to the network are necessary to lock the accounts.

See Also

Account Lockout Policy

Reset account lockout counter after

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Reset account lockout counter after** policy setting determines the number of minutes that must elapse from the time a user fails to log on before the failed logon attempt counter is reset to 0. If **Account lockout threshold** is set to a number greater than zero, this reset time must be less than or equal to the value of **Account lockout duration**.

A disadvantage to setting this too high is that users lock themselves out for an inconveniently long period if they exceed the account lockout threshold through logon errors. Users may make excessive Help Desk calls.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Possible values

- A user-defined number of minutes from 1 through 99,999
- Not defined

Best practices

- You need to determine the threat level for your organization and balance that against the cost of your Help Desk support for password resets. Each organization will have specific requirements.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Account Lockout Policy**

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not applicable
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users can accidentally lock themselves out of their accounts if they mistype their password multiple times.

Countermeasure

Configure the **Reset account lockout counter after** policy setting to 30.

Potential impact

If you do not configure this policy setting or if the value is configured to an interval that is too long, an attacker could attempt to log on to each user's account numerous times and lock out their accounts, a denial-of-service (DoS) attack might succeed, or administrators might have to manually unlock all locked-out accounts. If you configure this policy setting to a reasonable value, users can perform new attempts to log on after a failed logon within a reasonable time, without making brute force attacks feasible at high speeds. Be sure that you notify users of the values that are used for this policy setting so that they wait for the lockout timer to expire before they call the Help Desk.

See Also

Account Lockout Policy

Kerberos Policy

This topic for the IT professional describes the Kerberos Policy settings and provides links to policy setting descriptions.

The Kerberos version 5 authentication protocol provides the default mechanism for authentication services and the authorization data necessary for a user to access a resource and perform a task on that resource. By reducing the lifetime of Kerberos tickets, you reduce the risk of a legitimate user's credentials being stolen and successfully used by an attacker. However, this also

increases the authorization overhead. In most environments, these settings should not need to be changed.

These policy settings are located in *GPO_name***Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy** and can be set on a domain controller.

For information about setting security policies, see How to Configure Security Policy Settings.

The following topics provide a discussion of implementation and best practices considerations, policy location, default values for the server type or GPO, relevant differences in operating system versions, security considerations (including the possible settings vulnerabilities of each setting), countermeasures you can take, and the potential impact for each setting.

- Enforce user logon restrictions

  This policy setting determines whether the Kerberos V5 Key Distribution Center (KDC) validates every request for a session ticket against the user rights policy of the user account.
- Maximum lifetime for service ticket

  This policy setting determines the maximum number of minutes that a granted session ticket can be used to access a particular service.
- Maximum lifetime for user ticket

  This policy setting determines the maximum amount of time (in hours) that a user's ticket-granting ticket can be used.
- Maximum lifetime for user ticket renewal

  This policy setting determines the period of time (in days) during which a user's ticket-granting ticket can be renewed.
- Maximum tolerance for computer clock synchronization

  This policy setting determines the maximum time difference (in minutes) that the Kerberos V5 protocol tolerates between the time on the client clock and the time on the domain controller that provides Kerberos authentication.

Enforce user logon restrictions

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Enforce user logon restrictions** policy setting determines whether the Kerberos V5 Key Distribution Center (KDC) validates every request for a session ticket against the user rights policy of the user account. Validating each request for a session ticket is optional because the extra step takes time, and that can slow network access to services.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

The possible values for this Group Policy setting are:

- Enabled
- Disabled
- Not defined

Best practices

1. If this policy setting is disabled, users might be granted session tickets for services that they do not have the right to use.

   It is advisable to set **Enforce user logon restrictions** to Enabled.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

Default Values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server Type or GPO	Default Value
Default Domain Policy	Enabled
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not applicable
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Not applicable
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Group Policy

Client computers will get the new setting during the next scheduled and successful Group Policy refresh. But for domain controllers to assign these new settings immediately, a Gpupdate /force is required. On the local computer, the Security Configuration Engine will refresh this setting in about five minutes.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you disable this policy setting, users could receive session tickets for services that they no longer have the right to use because the right was removed after they logged on.

Countermeasure

Enable the **Enforce user logon restrictions** setting.

Potential impact

None. This is the default configuration.

See Also

Kerberos Policy

Maximum lifetime for service ticket

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Maximum lifetime for service ticket** policy setting determines the maximum number of minutes that a granted session ticket can be used to access a particular service. The value must be 10 minutes or greater, and it must be less than or equal to the value of the **Maximum lifetime for service ticket** policy setting.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 10 through 99,999, or 0 (in which case service tickets do not expire).
- Not defined.

If a client presents an expired session ticket when it requests a connection to a server, the server returns an error message. The client must request a new session ticket from the Kerberos V5 KDC. After a connection is authenticated, however, it no longer matters whether the session ticket remains valid. Session tickets are used only to authenticate new connections with servers. Ongoing operations are not interrupted if the session ticket that authenticated the connection expires during the connection.

If the value for this policy setting is too high, users might be able to access network resources outside of their logon hours. In addition, users whose accounts have been disabled might be able to continue accessing network services by using valid service tickets that were issued before their account was disabled. If the value is set to 0, service tickets never expire.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Best practices

1. It is advisable to set **Maximum lifetime for service ticket** to **600** minutes.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

Default Values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server Type or GPO	Default Value
Default Domain Policy	600 minutes
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not applicable
DC Effective Default Settings	600 minutes

Server Type or GPO	Default Value
Member Server Effective Default Settings	Not applicable
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

This policy setting is configured on the domain controller.

Group Policy

Client computers will get the new setting during the next scheduled and successful Group Policy refresh. But for domain controllers to assign these new settings immediately, a Gpupdate /force is required. On the local computer, the Security Configuration Engine will refresh this setting in about five minutes.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you configure the value for the **Maximum lifetime for service ticket** setting too high, users might be able to access network resources outside of their logon hours. Also, users whose accounts were disabled might continue to have access to network services with valid service tickets that were issued before their accounts were disabled.

Countermeasure

Configure the **Maximum lifetime for service ticket** setting to 600 minutes.

Potential impact

None. This is the default configuration.

See Also

Kerberos Policy

Maximum lifetime for user ticket

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Maximum lifetime for user ticket** policy setting determines the maximum amount of time (in hours) that a user's ticket-granting ticket can be used. When a user's ticket-granting ticket expires, a new one must be requested or the existing one must be renewed.

The possible values for this Group Policy setting are:

- A user-defined number of hours from 0 through 99,999
- Not defined

If the value for this policy setting is too high, users might be able to access network resources outside of their logon hours, or users whose accounts have been disabled might be able to continue to access network services by using valid service tickets that were issued before their account was disabled. If the value is set to 0, ticket-granting tickets never expire.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Best practices

1. It is advisable to set **Maximum lifetime for user ticket** to 10 hours.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

Default Values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server Type or GPO	Default Value
Default Domain Policy	10 hours
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not applicable
Domain Controller Effective Default Settings	10 hours
Member Server Effective Default Settings	Not applicable
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

This policy setting is configured on the domain controller.

Group Policy

Client computers will get the new setting during the next scheduled and successful Group Policy refresh. But for domain controllers to assign these new settings immediately, a Gpupdate /force is required. On the local computer, the Security Configuration Engine will refresh this setting in about five minutes.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you configure the value for the **Maximum lifetime for user ticket** setting too high, users might be able to access network resources outside of their logon hours. Also, users whose accounts were disabled might continue to have access to network services with valid user tickets that were issued before their accounts were disabled. If you configure this value too low, ticket requests to the KDC may affect the performance of your KDC and present an opportunity for a DoS attack.

Countermeasure

Configure the **Maximum lifetime for user ticket** setting with a value between 4 and 10 hours.

Potential impact

Reducing this setting from the default value reduces the likelihood that the ticket-granting ticket will be used to access resources that the user does not have rights to. However, it requires more frequent requests to the KDC for ticket-granting tickets on behalf of users. Most KDCs can support a value of four hours without too much additional burden.

See Also

Kerberos Policy

Maximum lifetime for user ticket renewal

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Maximum lifetime for user ticket renewal** policy setting determines the period of time (in days) during which a user's ticket-granting ticket can be renewed.

The possible values for this Group Policy setting are:

- A user-defined number of days from 0 through 99,999
- Not defined

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Best practices

1. If the value for this policy setting is too high, users may be able to renew very old user ticket-granting tickets. If the value is 0, ticket-granting tickets never expire.

   It is advisable to set **Maximum lifetime for user ticket renewal** to **7** days.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	7 days
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not applicable
Domain Controller Effective Default Settings	7 days
Member Server Effective Default Settings	Not applicable
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

This policy setting is configured on the domain controller.

Group Policy

Client computers will get the new setting during the next scheduled and successful Group Policy refresh. But for domain controllers to assign these new settings immediately, a Gpupdate /force is required. On the local computer, the Security Configuration Engine will refresh this setting in about five minutes.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If the value for the **Maximum lifetime for user ticket renewal** setting is too high, users might be able to renew very old user tickets.

Countermeasure

Configure the **Maximum lifetime for user ticket renewal** setting to 7 days.

Potential impact

None. This is the default configuration.

See Also

Kerberos Policy

Maximum tolerance for computer clock synchronization

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines the maximum time difference (in minutes) that Kerberos V5 tolerates between the time on the client clock and the time on the domain controller that provides Kerberos authentication.

To prevent "replay attacks," the Kerberos v5 protocol uses time stamps as part of its protocol definition. For time stamps to work properly, the clocks of the client and the domain controller need to be in sync as much as possible. In other words, both computers must be set to the same time and date. Because the clocks of two computers are often out of sync, administrators can use this policy setting to establish the maximum acceptable difference to the Kerberos protocol between a client clock and domain controller clock. If the difference between a client computer clock and the domain controller clock is less than the maximum time difference that is specified in this policy, any time stamp that is used in a session between the two computers is considered to be authentic.

Important

Changes to this setting are not persistent on computers running Windows Server 2003 and Windows XP. If you configure this setting and then restart the computer, this setting reverts to the default value. Beginning with Windows Server 2008 and Windows Vista, you can persist changes if you select the option **Define this policy in the database**. Selecting this option does not change current computer time settings.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

The possible values for this Group Policy setting are:

- A user-defined number of minutes from 1 through 99,999
- Not defined

To prevent replay attacks, the Kerberos protocol uses time stamps as part of its definition. For time stamps to work properly, the clocks of the client computer and the domain controller need to be synchronized as closely as possible. Because the clocks of two computers are often out of sync, administrators can use this policy setting to establish the maximum acceptable difference to the Kerberos protocol between a client computer clock and domain controller clock. If the difference between a client computer clock and the domain controller clock is less than **Maximum tolerance for computer clock synchronization**, any time stamp that is used in a session between the two computers is considered authentic.

Best practices

1. It is advisable to set **Maximum tolerance for computer clock synchronization** to a value of 5 minutes.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Account Policies\Kerberos Policy\

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	5 minutes
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not applicable
Domain Controller Effective Default Settings	5 minutes
Member Server Effective Default Settings	Not applicable
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in the way this policy setting works between supported versions of Windows.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

This policy setting is configured on the domain controller.

Group Policy

Client computers will get the new setting during the next scheduled and successful Group Policy refresh. But for domain controllers to assign these new settings immediately, a Gpupdate /force is required. On the local computer, the Security Configuration Engine will refresh this setting in about five minutes.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

To prevent "replay attacks" (which are attacks in which an authentication credential is resubmitted by a malicious user or program to gain access to a protected resource), the Kerberos protocol uses time stamps as part of its definition. For time stamps to work properly, the clocks of the client computer and the domain controller need to be closely synchronized. Because the clocks of two computers are often not synchronized, administrators can use this policy to establish the maximum acceptable difference to the Kerberos protocol between a client computer clock and a domain controller clock. If the difference between the client computer clock and the domain controller clock is less than the maximum time difference specified in this setting, any time stamp that is used in a session between the two computers is considered to be authentic.

Countermeasure

Configure the **Maximum tolerance for computer clock synchronization** setting to 5 minutes.

Potential impact

None. This is the default configuration.

Additional references

The following links provide additional information about topics that relate to securing domain controllers:

- Modify Default Security Policies on Windows Server 2008-Based Domain Controllers
- Best Practice Guide for Securing Windows Server Active Directory Installations

See Also

Kerberos Policy

Audit Policy

This reference topic for the IT professional provides information about one collection of auditing settings that is available in Windows and links to information about each setting.

The security audit policy settings under **Security Settings\Local Policies\Audit Policy** provide broad security audit capabilities for client computers and servers that cannot use advanced security audit policy settings.

For information about setting security policies, see How to Configure Security Policy Settings.

The basic audit policy settings under **Security Settings\Local Policies\Audit Policy** are:

Audit account logon events

Audit account management

Audit directory service access

Audit logon events

Audit object access

Audit policy change

Audit privilege use

Audit process tracking

Audit system events

Note
The descriptions of these policies are not located in this version of the TechNet Library.

Security Options

This reference topic for the IT professional provides an introduction to the settings under **Security Options** of the local security policies and links to information about each setting.

The **Security Options** contain the following groupings of security policy settings that allow you to configure the behavior of the local computer. Some of these policies can be included in a Group Policy Object and distributed over your enterprise.

If you edit policy settings locally on a computer, you will affect the settings on only that one computer. If you configure the settings in a Group Policy Object (GPO) hosted in an Active Directory domain, the settings apply to all computers that are subject to that GPO. For more information about Group Policy in an Active Directory domain, see Group Policy(http://go.microsoft.com/fwlink/?LinkId=55625).

For information about how the security policy snap-in and related technologies work, see Security Policy Settings Technical Overview.

Open the Local Security Policy snap-in (secpol.msc) and navigate to **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**.

Local computer permissions: Membership in the local Administrators group, or equivalent, is the minimum required to modify these policy settings.

For information about setting security policies, see How to Configure Security Policy Settings.

Grouping	Security Policy setting
Accounts	• Accounts: Administrator account status • Accounts: Block Microsoft accounts • Accounts: Guest account status • Accounts: Limit local account use of blank passwords to console logon only • Accounts: Rename administrator account

Grouping	Security Policy setting
	• Accounts: Rename guest account
Audit	• Audit: Audit the access of global system objects • Audit: Audit the use of Backup and Restore privilege • Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings • Audit: Shut down system immediately if unable to log security audits
DCOM	• DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax • DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax
Devices	• Devices: Allow undock without having to log on • Devices: Allowed to format and eject removable media • Devices: Prevent users from installing printer drivers • Devices: Restrict CD-ROM access to locally logged-on user only • Devices: Restrict floppy access to locally logged-on user only
Domain controller	• Domain controller: Allow server operators to schedule tasks • Domain controller: LDAP server signing requirements • Domain controller: Refuse machine account password changes
Domain member	• Domain member: Digitally encrypt or sign secure channel data (always) • Domain member: Digitally encrypt secure channel data (when possible) • Domain member: Digitally sign secure channel data (when possible) • Domain member: Disable machine account

Grouping	Security Policy setting
	password changes • Domain member: Maximum machine account password age • Domain member: Require strong (Windows 2000 or later) session key
Interactive logon	• Interactive logon: Display user information when the session is locked • Interactive logon: Do not display last user name • Interactive logon: Do not require CTRL+ALT+DEL • Interactive logon: Machine account lockout threshold • Interactive logon: Machine inactivity limit • Interactive logon: Message text for users attempting to log on • Interactive logon: Message title for users attempting to log on • Interactive logon: Number of previous logons to cache (in case domain controller is not available) • Interactive logon: Prompt user to change password before expiration • Interactive logon: Require Domain Controller authentication to unlock workstation • Interactive logon: Require smart card • Interactive logon: Smart card removal behavior
Microsoft network client	• Microsoft network client: Digitally sign communications (always) • Microsoft network client: Digitally sign communications (if server agrees) • Microsoft network client: Send unencrypted password to third-party SMB servers
Microsoft network server	• Microsoft network server: Amount of idle time required before suspending session • Microsoft network server: Attempt S4U2Self to obtain claim information

Grouping	Security Policy setting
	• Microsoft network server: Digitally sign communications (always) • Microsoft network server: Digitally sign communications (if client agrees) • Microsoft network server: Disconnect clients when logon hours expire • Microsoft network server: Server SPN target name validation level
Network access	• Network access: Allow anonymous SID/Name translation • Network access: Do not allow anonymous enumeration of SAM accounts • Network access: Do not allow anonymous enumeration of SAM accounts and shares • Network access: Do not allow storage of passwords and credentials for network authentication • Network access: Let Everyone permissions apply to anonymous users • Network access: Named Pipes that can be accessed anonymously • Network access: Remotely accessible registry paths • Network access: Remotely accessible registry paths and subpaths • Network access: Restrict anonymous access to Named Pipes and Shares • Network access: Shares that can be accessed anonymously • Network access: Sharing and security model for local accounts
Network security	• Network security: Allow Local System to use computer identity for NTLM • Network security: Allow LocalSystem NULL session fallback • Network Security: Allow PKU2U authentication requests to this computer to use online identities • Network security: Configure encryption types allowed for Kerberos

Grouping	Security Policy setting
	• Network security: Do not store LAN Manager hash value on next password change • Network security: Force logoff when logon hours expire • Network security: LAN Manager authentication level • Network security: LDAP client signing requirements • Network security: Minimum session security for NTLM SSP based (including secure RPC) clients • Network security: Minimum session security for NTLM SSP based (including secure RPC) servers • Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication • Network security: Restrict NTLM: Add server exceptions in this domain • Network Security: Restrict NTLM: Incoming NTLM Traffic • Network Security: Restrict NTLM: NTLM authentication in this domain • Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers • Network Security: Restrict NTLM: Audit Incoming NTLM Traffic • Network Security: Restrict NTLM: Audit NTLM authentication in this domain
Recovery console	• Recovery console: Allow automatic administrative logon • Recovery console: Allow floppy copy and access to all drives and folders
Shutdown	• Shutdown: Allow system to be shut down without having to log on • Shutdown: Clear virtual memory pagefile
System cryptography	• System cryptography: Force strong key protection for user keys stored on the computer

Grouping	Security Policy setting
	• System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing
System objects	• System objects: Require case insensitivity for non-Windows subsystems • System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links)
System settings	• System settings: Optional subsystems • System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies
User Account Control	• User Account Control: Admin Approval Mode for the Built-in Administrator account • User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop • User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode • User Account Control: Behavior of the elevation prompt for standard users • User Account Control: Detect application installations and prompt for elevation • User Account Control: Only elevate executables that are signed and validated • User Account Control: Only elevate UIAccess applications that are installed in secure locations • User Account Control: Run all administrators in Admin Approval Mode • User Account Control: Switch to the secure desktop when prompting for elevation • User Account Control: Virtualize file and registry write failures to per-user locations

Accounts: Administrator account status

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This security setting determines whether the local Administrator account is enabled or disabled.

If you try to enable the Administrator account after it has been disabled, and if the current Administrator password does not meet the password requirements, you cannot enable the account. In this case, an alternative member of the Administrators group must reset the password on the Administrator account.

If you disable this policy setting, and one of the following conditions exists on the computer, the Administrator account is not disabled.

1. No other local Administrator account exists
2. The Administrator account is currently in use
3. All other local Administrator accounts are:
    a. Disabled
    b. Listed in the Deny log on locally User Rights Assignment

If the current Administrator password does not meet the password requirements, you will not be able to enable the Administrator account again after it has been disabled. In this case, another member of the Administrators group must set the password on the Administrator account by using the Local Users and Groups tool.

Possible values

- Enabled
- Disabled
- Not defined

By default, this setting is **Not defined** on domain controllers and **Enabled** on stand-alone servers.

Best practices

1. Disabling the Administrator account can become a maintenance issue under certain circumstances. For example, in a domain environment, if the secure channel that constitutes your connection fails for any reason, and there is no other local Administrator account, you must restart the computer in safe mode to fix the problem that broke your connection status.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy setting between operating systems for the versions designated in the **Applies to** list at the beginning of this topic.

Policy management

Disabling the Administrator account can become a maintenance issue under certain circumstances. Reasons that an organization might consider disabling the built-in Administrator account include:

- For some organizations, periodically changing the passwords for local accounts can be a daunting management challenge.
- By default, the Administrator account cannot be locked—no matter how many failed attempts to sign in a user accrues. This makes it a prime target for brute-force, password-guessing attacks.
- This account has a well-known security identifier (SID). Some non-Microsoft tools allow you to authenticate over the network by specifying the SID rather than the account name. This means that even if you rename the Administrator account, a malicious user could start a brute-force attack by using the SID.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Safe mode considerations

When you start a computer in safe mode, the disabled Administrator account is enabled only if the computer is non-domain joined and there are no other active local Administrator accounts. If the computer is joined to a domain, the disabled Administrator account is not enabled.

If the Administrator account is disabled, you can still access the computer by using safe mode with the current administrative credentials. For example, if a failure occurs using a secure channel with a domain-joined computer, and there is no other local Administrator account, you must restart the computer in safe mode to fix the failure.

How to access a disabled Administrator account

Note
The following conditions and remedies apply to computers running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008.

You can use the following methods to access a disabled Administrator account:

- When there is only one local Administrator accountthat is disabled, start the computer in safe mode (locally or over a network), and sign in by using the credentials for the Administrator account on that computer.
- When there are local Administrator accounts in addition to the built-in account, start the computer in safe mode (locally or over a network), and sign in by using the credentials for the Administrator account on that computer. An alternate method is to sign in to Windows by using another local Administrator account that was created.
- When multiple domain-joined servers have a disabled local Administrator account that can be accessed in safe mode, you can remotely run the Sysinternals tool, PsExec, by using the following command: **net user administrator /active: no**.
- You can update a Group Policy Object (GPO) by using the **Group Policy Management Editor** as follows:

To access a disabled Administrator account

1. Under **Computer Configuration**, click **Preferences**, and then click **Control Panel Settings**.
2. Right-click **Local Users and Groups**, select **New**, and then click **Local User**.
3. In the **User name** field, type **Administrator**.
4. Enter and remember an acceptable password, if required.
5. Clear the **User must change password at next logon** check box.
6. Select **Account is disabled.**

7. Click **OK**.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The built-in Administrator account cannot be locked out no matter how many failed logons it accrues, which makes it a prime target for brute-force attacks that attempt to guess passwords. Also, this account has a well-known security identifier (SID), and there are non-Microsoft tools that allow authentication by using the SID rather than the account name. Therefore, even if you rename the Administrator account, an attacker could launch a brute-force attack by using the SID to log on. All other accounts that are members of the Administrator's group have the safeguard of locking out the account if the number of failed logons exceeds its configured maximum.

Countermeasure

Disable the **Accounts: Administrator account status** setting so that the built-in Administrator account cannot be used in a normal system startup.

If it is very difficult to maintain a regular schedule for periodic password changes for local accounts, you can disable the built-in Administrator account instead of relying on regular password changes to protect it from attack.

Potential impact

Maintenance issues can arise under certain circumstances if you disable the Administrator account. For example, if the secure channel between a member computer and the domain controller fails in a domain environment for any reason and there is no other local Administrator account, you must restart in safe mode to fix the problem that caused the secure channel to fail.

If the current Administrator password does not meet the password requirements, you cannot enable the Administrator account after it is disabled. If this situation occurs, another member of the Administrators group must set the password on the Administrator account with the Local Users and Groups tool.

Accounts: Block Microsoft accounts

This security policy reference topic for the IT professional describes the best practices, location, values, management, and security considerations for this policy setting.

Reference

This policy setting prevents users from adding new Microsoft accounts on a computer.

If you click the **Users can't add Microsoft accounts** setting option, users will not be able to create new Microsoft accounts on a computer, switch a local account to a Microsoft account, or connect a domain account to a Microsoft account. This is the preferred option if you need to limit the use of Microsoft accounts in your enterprise.

If you click the **Users can't add or log on with Microsoft accounts** setting option, existing Microsoft account users will not be able to log on to Windows. Selecting this option might make it impossible for an existing administrator to log on to a computer and manage the system.

If you disable or do not configure this policy (recommended), users will be able to use Microsoft accounts with Windows.

Possible values

- This policy is disabled
- Users can't add Microsoft accounts
- Users can't add or log on with Microsoft accounts

By default, this setting is not defined on domain controllers and disabled on stand-alone servers.

Best practices

- By disabling or not configuring this policy setting on the client computer, users will be able to use their Microsoft account, local account, or domain account for their sign-in session to Windows. It also enables the user to connect a local or domain account to a Microsoft account. This provides a convenient option for your users.
- If you need to limit the use of Microsoft accounts in your organization, click the **Users can't add Microsoft accounts** setting option so that users will not be able to create new Microsoft accounts on a computer, switch a local account to a Microsoft account, or connect a domain account to a Microsoft account.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

This policy setting was introduced in Windows Server 2012 and Windows 8.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Because this policy setting was introduced to control sign-in options for Windows Server 2012 and Windows 8, it has no affect on earlier Windows operating systems.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of the countermeasure implementation.

Vulnerability

Although Microsoft accounts are password-protected, they also have the potential of greater exposure outside of the enterprise. Additionally, if the owner of a Microsoft account is not easily distinguishable, auditing and forensics become more difficult on computers running Windows Server 2012 and Windows 8.

Countermeasure

Require only domain accounts in your enterprise by limiting the use of Microsoft accounts. Click the **Users can't add Microsoft accounts** setting option so that users will not be able to create new Microsoft accounts on a computer, switch a local account to a Microsoft account, or connect a domain account to a Microsoft account.

Potential impact

Establishing greater control over accounts in your organization can give you more secure management capabilities, including procedures around password resets.

See Also

Security Options

Accounts: Guest account status

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting

Reference

The **Accounts: Guest account status** policy setting determines whether the Guest account is enabled or disabled.

This account allows unauthenticated network users to gain access to the system by logging on as a Guest with no password. Unauthorized users can access any resources that are accessible to the Guest account over the network. This means that any network shared folders with permissions that allow access to the Guest account, the Guests group, or the Everyone group will be accessible over the network. This can lead to the exposure or corruption of data.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

Set **Accounts: Guest account status** to Disabled so that the built-in Guest account is no longer usable. All network users will have to authenticate before they can access shared resources on the system. If the Guest account is disabled and **Network access: Sharing and security model**

for local accounts is set to **Guest only**, network logons—such as those performed by the SMB Service—will fail.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The default Guest account allows unauthenticated network users to log on as a Guest with no password. These unauthorized users could access any resources that are accessible to the Guest account over the network. This capability means that any shared folders with permissions that allow access to the Guest account, the Guests group, or the Everyone group are accessible over the network, which could lead to the exposure or corruption of data.

Countermeasure

Disable the **Accounts: Guest account status** setting so that the built-in Guest account cannot be used.

Potential impact

All network users must be authenticated before they can access shared resources. If you disable the Guest account and the **Network Access: Sharing and Security Model** option is set to **Guest Only**, network logons, such as those performed by the Microsoft Network Server (SMB Service), fail. This policy setting should have little impact on most organizations because it is the default setting in the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 7, Windows Vista, Windows XP, and Microsoft Windows 2000 operating systems.

Accounts: Limit local account use of blank passwords to console logon only

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Accounts: Limit local account use of blank passwords to console logon only** policy setting determines whether remote interactive logons by network services such as Remote Desktop Services (formerly Terminal Services), Telnet, and File Transfer Protocol (FTP) are allowed for local accounts that have blank passwords. If this policy setting is enabled, a local account must have a nonblank password to be used to perform an interactive or network logon from a remote client.

This policy setting does not affect interactive logons that are performed physically at the console or logons that use domain accounts. It is possible for non-Microsoft applications that use remote interactive logons to bypass this policy setting.

Blank passwords are a serious threat to computer security and they should be forbidden through both corporate policy and suitable technical measures. Nevertheless, if a user with the ability to create new accounts creates one that has bypassed your domain-based password policy settings, that account might have a blank password. For example, a user could build a stand-alone system, create one or more accounts with blank passwords, and then join the computer to the domain. The local accounts with blank passwords would still function. Anyone who knows the account name can then use accounts with blank passwords to log on to systems.

Computers that are not in physically secure locations should always enforce strong password policies for all local user accounts. Otherwise, anyone with physical access to the computer can log on by using a user account that does not have a password. This is especially important for portable computers.

If you apply this security policy to the Everyone group, no one will be able to log on through Remote Desktop Services.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. It is advisable to set **Accounts: Limit local account use of blank passwords to console logon only** to Enabled.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

The policy as distributed through the GPO takes precedence over the locally configured policy setting on a computer joined to a domain. On the domain controller, use ADSI Edit or the dsquery command to determine effective minimum password length.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Blank passwords are a serious threat to computer security, and they should be forbidden through organizational policy and suitable technical measures.The default settings for Active Directory® domains in Windows Server 2003 require complex passwords of at least seven characters, and eight characters for Windows Server 2008 R2 and Windows Server 2008. However, if users with the ability to create new accounts bypass your domain-based password policies, they could create accounts with blank passwords. For example, a user could build a stand-alone computer, create one or more accounts with blank passwords, and then join the computer to the domain. The local accounts with blank passwords would still function. Anyone who knows the name of one of these unprotected accounts could then use it to log on.

Countermeasure

Enable the **Accounts: Limit local account use of blank passwords to console logon only** setting.

Potential impact

None. This is the default configuration.

Accounts: Rename administrator account

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Accounts: Rename administrator account** policy setting determines whether a different account name is associated with the security identifier (SID) for the Administrator account.

Because the Administrator account exists on all Windows server versions, renaming the account makes it slightly more difficult for attackers to guess this user name and password combination. By default, the built-in Administrator account cannot be locked out no matter how many times a malicious user might use a bad password. This makes the Administrator account a popular target for brute-force password-guessing attacks. The value of this countermeasure is lessened because this account has a well-known SID and there are non-Microsoft tools that allow you to initiate a brute-force attack over the network by specifying the SID rather than the account name. This means that even if you have renamed the Administrator account, a malicious user could start a brute-force attack by using the SID.

Rename the Administrator account by specifying a value for the **Accounts: Rename administrator account** policy setting.

Possible values

- User-defined text
- Not defined

Best practices

1. Be sure to inform users who are authorized to use this account of the new account name.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Administrator
DC Effective Default Settings	Administrator
Member Server Effective Default Settings	Administrator
Client Computer Effective Default Settings	Administrator

Operating system version differences

In Windows Server 2003, this policy setting is not configured in the security templates, nor is a new user name for the account suggested there. This is so that numerous organizations reading this guide will not implement the same new user name in their environments. In later versions of the Windows operating system, the preconfigured security templates were removed.

Beginning with Windows Vista, the person who installs the operating system specifies an account that is the first member of the Administrator group and has full rights to configure the computer.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The Administrator account exists on all computers that run the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server, or Windows XP Professional operating systems. If you rename this account, it is slightly more difficult for unauthorized persons to guess this privileged user name and password combination. Beginning with Windows Vista, the person who installs the operating system specifies an account that is the first member of the Administrator group and has full rights to configure the computer. The account cannot have the name Administrator, so this countermeasure is applied by default on new Windows Vista and Windows 7 installations. If a computer is upgraded from a previous version of Windows to Windows Vista or Windows 7, the account with the name Administrator is retained with all the rights and privileges that were defined for the account in the previous installation.

The built-in Administrator account cannot be locked out, regardless of how many times an attacker might use a bad password. This capability makes the Administrator account a popular target for brute-force attacks that attempt to guess passwords. The value of this countermeasure is lessened because this account has a well-known SID, and there are non-Microsoft tools that allow authentication by using the SID rather than the account name. Therefore, even if you rename the Administrator account, an attacker could launch a brute-force attack by using the SID to log on.

Countermeasure

Specify a new name in the **Accounts: Rename administrator account** setting to rename the Administrator account.

Potential impact

You must provide users who are authorized to use this account with the new account name. (The guidance for this setting assumes that the Administrator account was not disabled.)

Accounts: Rename guest account

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Accounts: Rename guest account** policy setting determines whether a different account name is associated with the security identifier (SID) for the Guest account.

The Guest account exists in the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 7, Windows Vista, and Windows XP Professional operating systems. Renaming the Guest account makes it slightly more difficult for unauthorized persons to guess this privileged user name and password combination.

Rename the Guest account by specifying a value for the **Accounts: Rename guest account** policy setting.

The Guest account is disabled by default in Windows Server 2003, Windows 2000 Server, and Windows XP, but it is set as **Guest** in Windows Server 2008 R2, Windows Server 2008, Windows 7, and Windows Vista.

Possible values

- *User-defined text*
- Guest

Best practices

1. For computers in unsecured locations, renaming the account makes it more difficult for unauthorized users to guess it.
2. For computers in secured or trusted locations, keeping the name of the account as Guest provides consistency among computers.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Guest
Default Domain Controller Policy	Guest
Stand-Alone Server Default Settings	Guest
DC Effective Default Settings	Guest
Member Server Effective Default Settings	Guest
Client Computer Effective Default Settings	*User-defined text*

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The Guest account exists in all Windows server and client operating system versions beginning with Windows Server 2003 and Windows XP Professional. Because the account name is well known, it provides a vector for a malicious user to get access to network resources and attempt to elevate privileges or install software that could be used for a later attack on your system.

Countermeasure

Specify a new name in the **Accounts: Rename guest account** setting to rename the Guest account. If you rename this account, it is slightly more difficult for unauthorized persons to guess this privileged user name and password combination.

Potential impact

There should be little impact because the Guest account is disabled by default in Windows 2000 Server, Windows Server 2003, and Windows XP. For later operating systems, the policy is enabled with **Guest** as the default.

Audit: Audit the access of global system objects

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

If you enable this policy setting, a default system access control list (SACL) is applied when the computer creates system objects such as mutexes, events, semaphores, and MS-DOS® devices. If you also enable the **Audit object access** audit setting, access to these system objects is audited.

Global system objects, also known as "base system objects" or "base named objects," are temporary kernel objects that have had names assigned to them by the application or system component that created them. These objects are most commonly used to synchronize multiple applications or multiple parts of a complex application. Because they have names, these objects are global in scope and, therefore, visible to all processes on the computer. These objects all have a security descriptor; but typically, they do not have a NULL SACL. If you enable this policy setting and it takes effect at startup time, the kernel assigns a SACL to these objects when they are created.

The threat is that a globally visible named object, if incorrectly secured, might be acted on by a malicious program that knows the name of the object. For instance, if a synchronization object such as a mutex has a poorly constructed discretionary access control list (DACL), a malicious program can access that mutex by name and cause the program that created it to malfunction. However, the risk of this occurring is very low.

Enabling this policy setting can generate a large number of security events, especially on busy domain controllers and application servers. This might cause servers to respond slowly and force the security log to record numerous events of little significance. Auditing for access to global system objects is an all-or-nothing affair; there is no way to filter which events get recorded and which do not. Even if an organization has the resources to analyze events generated when this policy setting is enabled, it is unlikely to have the source code or a description of what each named object is used for; therefore, it is unlikely that many organizations could benefit from enabling this policy setting.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Use the advanced security audit policy option, **Audit Kernel Object** in Advanced Security Audit Policy Settings\Object Access, to reduce the number of unrelated audit events that you generate.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008 and Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

A restart of the computer is required before this policy will be effective when changes to this policy are saved locally or distributed through Group Policy.

Group Policy

This policy can be applied to Windows beginning with Windows Server 2003 and Windows XP.

Auditing

To audit attempts to access global system objects, you can use one of two security audit policy settings:

- **Audit Kernel Object** in Advanced Security Audit Policy Settings\Object Access
- **Audit Object Access** under Security Settings\Local Policies\Audit Policy

If possible, use the Advanced Security Audit Policy option to reduce the number of unrelated audit events that you generate. For more information, see **Which Versions of Windows Support Advanced Audit Policy Configuration?**

If the **Audit Kernel Object** setting is configured, the following events are generated. The events appear on computers running Windows Server 2008 R2, Windows Server 2008, Windows 7, or Windows Vista.

Event ID	Event message
4659	A handle to an object was requested with intent to delete.
4660	An object was deleted.
4661	A handle to an object was requested.
4663	An attempt was made to access an object.

If the **Audit Object Access** setting is configured, the following events are generated.

Event ID	Event message
560	Access was granted to an already existing object.
562	A handle to an object was closed.
563	An attempt was made to open an object with the intent to delete it. **Note** This is used by file systems when the FILE_DELETE_ON_CLOSE flag is specified in Createfile().
564	A protected object was deleted.
565	Access was granted to an already existing object type.
567	A permission associated with a handle was used. **Note**

Event ID	Event message
	A handle is created with certain granted permissions (Read, Write, and so on). When the handle is used, up to one audit is generated for each of the permissions that was used.
569	The resource manager in Authorization Manager attempted to create a client context.
570	A client attempted to access an object. **Note** An event will be generated for every attempted operation on the object.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A globally visible named object, if incorrectly secured, could be acted upon by malicious software by using the name of the object. For instance, if a synchronization object such as a mutex had a poorly chosen discretionary access control list (DACL), malicious software could access that mutex by name and cause the program that created it to malfunction. However, the risk of such an occurrence is very low.

Countermeasure

Enable the **Audit: Audit the access of global system objects** setting.

Potential impact

If you enable the **Audit: Audit the access of global system objects** setting, a large number of security events could be generated, especially on busy domain controllers and application servers. Such an occurrence could cause servers to respond slowly and force the Security log to record numerous events of little significance. This policy setting can only be enabled or disabled, and there is no way to choose which events are recorded from this setting. Even organizations that have the resources to analyze events that are generated by this policy setting are not likely to have the source code or a description of what each named object is used for. Therefore, it is unlikely that most organizations would benefit by enabling this policy setting.

To reduce the number of audit events generated, use the advanced audit policy. There are some restrictions that apply. For more information, see **Which Versions of Windows Support Advanced Audit Policy Configuration?**

Additional resources

For the advanced audit policy options, see **Audit Kernel Object**.

For basic audit policy, see **Audit object access**.

Audit: Audit the use of Backup and Restore privilege

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Audit: Audit the use of Backup and Restore privilege** policy setting determines whether to audit the use of all user rights, including Backup and Restore, when the **Audit privilege use** policy setting is configured. Enabling both policy settings generates an audit event for every file that is backed up or restored.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set **Audit: Audit the use of Backup and Restore privilege** to Disabled. Enabling this policy setting can generate a large number of security events, which might cause servers to respond slowly and force the security log to record numerous events of little significance.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operation systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Auditing

Enabling this policy setting in conjunction with the **Audit privilege use** policy setting records any instance of user rights that are being exercised in the security log. If **Audit privilege use** is enabled but **Audit: Audit the use of Backup and Restore privilege** is disabled, when users use backup or restore user rights, those events will not be audited.

Enabling this policy setting when the **Audit privilege use** policy setting is also enabled generates an audit event for every file that is backed up or restored. This can help you to track down an administrator who is accidentally or maliciously restoring data in an unauthorized manner.

Alternately, you can use the advanced audit policy, **Audit Sensitive Privilege Use**, which can help you manage the number of events generated. For information about this policy, see **Audit Sensitive Privilege Use**.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When the backup and restore function is used, it creates a copy of the file system that is identical to the target of the backup. Making regular backup and restore volumes is an important part of your incident response plan. However, a malicious user could use a legitimate backup copy to gain access to information or to impersonate a legitimate network resource to compromise your enterprise.

Countermeasure

Enable the **Audit: Audit the use of Backup and Restore privilege** setting. Alternatively, implement automatic log backup by configuring the **AutoBackupLogFiles** registry key. If you enable this option when the **Audit privilege use** setting is also enabled, an audit event is generated for every file that is backed up or restored. This information could help you to identify an account that was used to accidentally or maliciously restore data in an unauthorized manner.

For more information about configuring this key, see article 100879 in the Microsoft Knowledge Base.

Potential impact

If you enable this policy setting, a large number of security events could be generated, which could cause servers to respond slowly and force the Security event log to record numerous events of little significance. If you increase the Security log size to reduce the chances of a system shutdown, an excessively large log file may affect system performance.

Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The applicable Windows operating systems allow audit policy to be managed in a more precise way by using audit policy subcategories.

There are over 40 auditing subcategories that provide precise details about activities on a computer. For information about these subcategories, see the **Advanced Security Audit Policy Settings**.

Possible values

- Enabled
- Disabled

Best practices

1. Leave the setting enabled. This provides the ability to audit events at the category level without revising a policy.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy for operating systems beginning with Windows Vista and Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

In Windows Server 2008 R2 and Windows 7, all auditing capabilities are integrated in Group Policy. This allows administrators to configure, deploy, and manage these settings in the Group Policy Management Console (GPMC) or Local Security Policy snap-in for a domain, site, or organizational unit (OU).

Auditing

To manage an audit policy by using subcategories without requiring a change to Group Policy, the SCENoApplyLegacyAuditPolicy registry value , prevents the application of category-level audit policy from Group Policy and from the Local Security Policy administrative tool. This registry value is available in Windows Vista and Windows 7.

If the category level audit policy that is set here is not consistent with the events that are currently being generated, the cause might be that this registry key is set.

Command-line tools

For information about displaying and manipulating audit policies from the command line, see **Auditpol**.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Prior to the introduction of auditing subcategories in Windows Vista, it was difficult to track events at a per-system or per-user level. The larger event categories created too many events, and the key information that needed to be audited was difficult to find.

Countermeasure

Enable audit policy subcategories as needed to track specific events.

Potential impacts

The individual audit policy subcategories available in Windows Vista and Windows Server 2008 are not exposed in the interface of Group Policy tools. Administrators can deploy a custom audit

policy that applies detailed security audit settings to Windows Vista–based or Windows Server 2008 computers in a Windows Server 2008, Windows Server 2003, or Windows 2000 domain. But the advanced security audit policies are available in the interface of Group Policy in Windows Server 2008 R2.

If you attempt to modify an audit setting by using Group Policy after enabling this setting through the command-line tools, the Group Policy audit setting is ignored in favor of the custom policy setting. To modify audit settings by using Group Policy, you must first disable the **SCENoApplyLegacyAuditPolicy** key.

Important
Be very cautious about audit settings that can generate a large volume of traffic. For example, if you enable success or failure auditing for all of the Privilege Use subcategories, the high volume of audit events that are generated can make it difficult to find other types of entries in the Security log. Such a configuration could also have a significant impact on system performance.

Audit: Shut down system immediately if unable to log security audits

This security policy reference topic for the IT professional describes the best practices, location, values, management practices, and security considerations for this policy setting.

Reference

The **Audit: Shut down system immediately if unable to log security audits** policy setting determines whether the system shuts down if it is unable to log security events. This policy setting is a requirement for Trusted Computer System Evaluation Criteria (TCSEC)-C2 and Common Criteria certification to prevent auditable events from occurring if the audit system is unable to log those events. Microsoft has chosen to meet this requirement by halting the system and displaying a Stop message in the case of a failure of the auditing system. Enabling this policy setting stops the system if a security audit cannot be logged for any reason. Typically, an event fails to be logged when the security audit log is full and the value of **Retention method for security log** is **Do not overwrite events (clear log manually)** or **Overwrite events by days**.

With **Audit: Shut down system immediately if unable to log security audits** set to **Enabled**, if the security log is full and an existing entry cannot be overwritten, the following Stop message appears:

```
STOP: C0000244 {Audit Failed}
An attempt to generate a security audit failed.
```

To recover, an administrator must log on, archive the log (optional), clear the log, and reset this option as desired.

If the computer is unable to record events to the security log, critical evidence or important troubleshooting information might not be available for review after a security incident.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Depending on your security audit requirements, you can enable the **Audit: Shut down system immediately if unable to log security audits** setting to ensure that security auditing information is captured for review. However, enabling this setting will increase the number of events logged.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between Windows operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

The administrative burden of enabling this policy setting can be very high, especially if you also set the **Retention method for security log** to **Do not overwrite events (clear log manually)**. This setting turns a repudiation threat (a backup operator could deny that they backed up or restored data) into a denial-of-service threat, because a server can be forced to shut down if it is overwhelmed with logon events and other security events that are written to the security log. Additionally, because the shutdown is not graceful, it is possible that irreparable damage to the operating system, applications, or data could result. Although the NTFS file system will guarantee that the file system's integrity will be maintained during a sudden system shutdown, it cannot guarantee that every data file for every application will still be in a usable form when the system is restarted.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Modifying this setting may affect compatibility with clients, services, and applications.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If the computer is unable to record events to the Security log, critical evidence or important troubleshooting information may not be available for review after a security incident. Also, an attacker could potentially generate a large volume of Security log events to purposely force a computer shutdown.

Countermeasure

Enable the **Audit: Shut down system immediately if unable to log security audits** setting to ensure that security auditing information is captured for review.

Potential impact

If you enable this policy setting, the administrative burden can be significant, especially if you also configure the **Retention method for the Security log** to **Do not overwrite events** (clear log

manually). This configuration causes a repudiation threat (a backup operator could deny that they backed up or restored data) to become a denial of service (DoS) vulnerability because a server could be forced to shut down if it is overwhelmed with logon events and other security events that are written to the Security log. Also, because the shutdown is abrupt, it is possible that irreparable damage to the operating system, applications, or data could result. Although the NTFS file system maintains its integrity when this type of computer shutdown occurs, there is no guarantee that every data file for every application will still be in a usable form when the computer restarts.

DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting allows administrators to define additional computer-wide controls that govern access to all Distributed Component Object Model (DCOM)–based applications on a computer. These controls restrict call, activation, or launch requests on the computer. A simple way to think about these access controls is as an additional access check that is performed against a computer-wide access control list (ACL) on each call, activation, or launch of any COM-based server. If the access check fails, the call, activation, or launch request is denied. (This check is in addition to any access check that is run against the server-specific ACLs.) In effect, it provides a minimum authorization standard that must be passed to access any COM-based server. This policy setting controls access permissions to cover call rights.

These computer-wide ACLs provide a way to override weak security settings that are specified by an application through the CoInitializeSecurity function or application-specific security settings. They provide a minimum security standard that must be passed, regardless of the settings of the specific server.

These ACLs also provide a centralized location for an administrator to set a general authorization policy that applies to all COM-based servers on the computer.

This policy setting allows you to specify an ACL in two different ways. You can type the security descriptor in SDDL, or you can grant or deny Local Access and Remote Access permissions to users and groups. We recommend that you use the built-in user interface to specify the ACL contents that you want to apply with this setting. The default ACL settings vary, depending on the version of Windows you are running.

Possible values

- *User-defined input* of the SDDL representation of the groups and privileges

When you specify the users or groups that are to be given permissions, the security descriptor field is populated with the Security Descriptor Definition Language representation of those groups and privileges. Users and groups can be given explicit Allow or Deny privileges for local access and remote access.

- Blank

  This represents how the local security policy deletes the policy enforcement key. This value deletes the policy and then sets it as Not defined. The Blank value is set by using the ACL editor to empty the list, and then pressing OK.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Blank
Default Domain Controller Policy	Blank
Stand-Alone Server Default Settings	Blank
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

The registry settings that are created as a result of enabling the **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** policy setting take precedence over the previous registry settings when this olicy setting was configured. The Remote Procedure Call (RPC) service checks the new registry keys in the Policies section for the computer restrictions, and these registry entries take precedence over the existing registry keys under OLE. This means that previously existing registry settings are no longer effective, and if you make changes to the existing settings, computer access permissions for users are not changed. Use care in configuring the list of users and groups.

If the administrator is denied permission to access DCOM applications due to the changes made to DCOM in the Windows operating system, the administrator can use the **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** policy setting to manage DCOM access to the computer. The administrator can use this setting to specify which users and groups can access the DCOM application on the computer locally and remotely. This will restore control of the DCOM application to the administrator and users. To do this, open the **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** setting, and click **Edit Security**. Specify the users or groups you want to include and the computer access permissions for those users or groups. This defines the setting and sets the appropriate SDDL value.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Many COM applications include some security-specific code (for example, to call CoInitializeSecurity), but they use weak settings that allow unauthenticated access to the process. Administrators cannot override these settings to force stronger security in earlier versions of Windows without modifying the application. An attacker could attempt to exploit weak security in an individual application by attacking it through COM calls.

Also, the COM infrastructure includes the Remote Procedure Call Services (RPCSS), a system service that runs during and after computer startup. This service manages activation of COM objects and the running object table and provides helper services to DCOM remoting. It exposes RPC interfaces that can be called remotely. Because some COM-based servers allow unauthenticated remote access, these interfaces can be called by anyone, including unauthenticated users. As a result, RPCSS can be attacked by malicious users who use remote, unauthenticated computers.

Countermeasure

To protect individual COM-based applications or services, set the **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** setting to an appropriate computer-wide ACL.

Potential impact

Windows operating systems implement default COM ACLs when they are installed. Modifying these ACLs from the default may cause some applications or components that communicate by using DCOM to fail. If you implement a COM-based server and you override the default security settings, confirm that the application-specific call permissions that ACL assigns are the correct permissions for appropriate users. If it does not, you must change your application-specific permission ACL to provide appropriate users with activation rights so that applications and Windows components that use DCOM do not fail.

DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting is similar to the **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** setting in that it allows administrators to define additional computer-wide controls that govern access to all DCOM–based applications on a computer. However, the ACLs that are specified in this policy setting control local and remote COM launch requests (not access requests) on the computer. A simple way to think about this access control is as an additional access check that is performed against a computer-wide ACL on each launch of any COM-based server. If the access check fails, the call, activation, or launch request is denied. (This check is in addition to any access check that is run against the server-specific ACLs.) In effect, it provides a minimum authorization standard that must be passed to launch any COM-based server. The **DCOM: Machine Access Restrictions in Security Descriptor Definition Language (SDDL) syntax** policy setting differs in that it provides a minimum access check that is applied to attempts to access an already launched COM-based server.

These computer-wide ACLs provide a way to override weak security settings that are specified by an application through CoInitializeSecurity or application-specific security settings. They provide a minimum security standard that must be passed, regardless of the settings of the specific COM-

based server. These ACLs provide a centralized location for an administrator to set a general authorization policy that applies to all COM-based servers.

The **DCOM: Machine Launch Restrictions in the Security Descriptor Definition Language (SDDL) syntax** setting allows you to specify an ACL in two ways. You can type the security descriptor in SDDL, or you can grant or deny Local Access and Remote Access permissions to users and groups. We recommend that you use the built-in user interface to specify the ACL contents that you want to apply with this setting. The default ACL settings vary, depending on the version of Windows you are running.

Possible values

- Blank

  This represents how the local security policy deletes the policy enforcement key. This value deletes the policy and then sets it to Not defined. The Blank value is set by using the ACL editor to empty the list, and then pressing OK.

- *User-defined input* of the SDDL representation of the groups and privileges

  When you specify the users or groups that are to be given permission, the security descriptor field is populated with the Security Descriptor Definition Language representation of those groups and privileges. Users and groups can be given explicit Allow or Deny privileges on both local access and remote access.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Blank
Default Domain Controller Policy	Blank
Stand-Alone Server Default Settings	Blank
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

The registry settings that are created as a result of this policy take precedence over the previous registry settings in this area. The Remote Procedure Call (RPC) service (RpcSs) checks the new registry keys in the Policies section for the computer restrictions; these entries take precedence over the existing registry keys under OLE.

If the administrator is denied access to activate and launch DCOM applications due to the changes made to DCOM in the Windows operating system, this policy setting can be used to control the DCOM activation and launch to the computer.

The administrator can specify which users and groups can launch and activate DCOM applications on the computer locally and remotely by using the **DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax** policy setting. This restores control of the DCOM application to the administrator and specified users. To do this, open the **DCOM: Machine Launch Restrictions in Security Descriptor Definition Language (SDDL) syntax** setting, and click **Edit Security**. Specify the groups that you want to include and the computer launch permissions for those groups. This defines the setting and sets the appropriate SDDL value.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Many COM applications include some security-specific code (for example, to call CoInitializeSecurity), but they use weak settings that allow unauthenticated access to the process. Administrators cannot override these settings to force stronger security in earlier versions of Windows without modifying the application. An attacker could attempt to exploit weak security in an individual application by attacking it through COM calls.

Also, the COM infrastructure includes the Remote Procedure Call Service (RPCSS), a system service that runs during computer startup and always runs after that. This service manages activation of COM objects and the running object table and provides helper services to DCOM remoting. It exposes RPC interfaces that can be called remotely. Because some COM-based servers allow unauthenticated remote component activation, these interfaces can be called by anyone, including unauthenticated users. As a result, RPCSS can be attacked by malicious users using remote, unauthenticated computers.

Countermeasure

To protect individual COM-based applications or services, set this policy setting to an appropriate computer-wide ACL.

Potential impact

Windows operating systems implement default COM ACLs when they are installed. Modifying these ACLs from the default may cause some applications or components that communicate by using DCOM to fail. If you implement a COM-based server and you override the default security settings, confirm that the application-specific launch permissions ACL assigns include activation permissions to appropriate users. If it does not, you must change your application-specific launch permission ACL to provide appropriate users with activation rights so that applications and Windows components that use DCOM do not fail.

Devices: Allow undock without having to log on

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting enables or disables the ability of a user to remove a portable computer from a docking station without logging on. If you enable this policy setting, users can press a docked portable computer's physical eject button to safely undock the computer. If you disable this policy setting, the user must log on to receive permission to undock the computer. Only users who have the **Remove Computer from Docking Station** privilege can obtain this permission.

Note
Disabling this policy setting only reduces theft risk for portable computers that cannot be mechanically undocked. Computers that can be mechanically undocked can be physically removed by the user whether or not they use the Windows undocking functionality.

Enabling this policy setting means that anyone with physical access to a computer that has been placed in its docking station can remove the computer and possibly tamper with it. For computers that do not have docking stations, this policy setting has no impact. However, for users with a mobile computer that is normally docked while they are in the office, this policy setting will help lower the risk of equipment theft or a malicious user gaining physical access to these computers.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

It is advisable to disable the **Devices: Allow undock without having to log on** policy setting. Users who have docked their computers will have to log on to the local console before they can undock their systems.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy setting between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If this policy setting is enabled, anyone with physical access to portable computers in docking stations could remove them and possibly tamper with them.

Countermeasure

Disable the **Devices: Allow undock without having to log on** setting.

Potential impact

Users who have docked their computers must log on to the local console before they can undock their computers. For computers that do not have docking stations, this policy setting has no impact.

Devices: Allowed to format and eject removable media

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines who is allowed to format and eject removable media.

Users can move removable disks to a different computer where they have administrative user rights and then take ownership of any file, assign themselves full control, and view or modify any file. The advantage of configuring this policy setting is diminished by the fact that most removable storage devices will eject media with the press of a button.

Possible values

- Administrators
- Administrators and Power Users
- Administrators and Interactive Users (not applicable to Windows Server 2008 R2 or Windows 7)
- Not defined

Best practices

1. It is advisable to set **Allowed to format and eject removable media** to **Administrators**. Only administrators will be able to eject NTFS-formatted removable media.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Administrators
DC Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Not defined

Operating system version differences

The option Administrators and Interactive Users is not applicable to Windows Server 2008 R2 or Windows 7, but it can be used for backward compatibility.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users could move data on removable disks to a different computer where they have administrative privileges. The user could then take ownership of any file, grant themselves full control, and view or modify any file. The fact that most removable storage devices eject media when a mechanical button is pressed diminishes the advantage of this policy setting.

Countermeasure

Configure the **Devices: Allowed to format and eject removable media** setting to **Administrators**.

Potential impact

Only administrators can format and eject removable media. If users are in the habit of using removable media for file transfers and storage, they must be informed of the change in policy.

Devices: Prevent users from installing printer drivers

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

For a computer to print to a network printer, the driver for that network printer must be installed on the local computer. The **Devices: Prevent users from installing printer drivers** policy setting determines who can install a printer driver as part of adding a network printer. When you set the value to **Enabled**, only Administrators and Power Users can install a printer driver as part of adding a network printer. Setting the value to **Disabled** allows any user to install a printer driver as part of adding a network printer. This setting prevents unprivileged users from downloading and installing an untrusted printer driver.

This setting has no impact if an administrator has configured a trusted path for downloading drivers. When using trusted paths, the print subsystem attempts to use the trusted path to download the driver. If the trusted path download succeeds, the driver is installed on behalf of any user. If the trusted path download fails, the driver is not installed and the network printer is not added.

Although it might be appropriate in some organizations to allow users to install printer drivers on their own workstations, this is not suitable for servers. Installing a printer driver on a server can cause the system to become less stable. Only administrators should have this user right on servers. A malicious user might deliberately try to damage the system by installing inappropriate printer drivers.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. It is advisable to set **Devices: Prevent users from installing printer drivers** to Enabled. Only users in the Administrative, Power User, or Server Operator groups will be able to install printers on servers. If this policy setting is enabled, but the driver for a network printer already exists on the local computer, users can still add the network printer. This policy setting does not affect a user's ability to add a local printer.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled

Server type or GPO	Default value
Client Computer Effective Default Settings	Disabled

Operating system version differences

None.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

It may be appropriate in some organizations to allow users to install printer drivers on their own workstations. However, you should allow only administrators, not users, to do so on servers because printer driver installation on a server may unintentionally cause the computer to become less stable. A malicious user could install inappropriate printer drivers in a deliberate attempt to damage the computer, or a user might accidentally install malicious software that masquerades as a printer driver.

Countermeasure

Enable the **Devices: Prevent users from installing printer drivers** setting.

Potential impact

Only members of the Administrator, Power Users, or Server Operator groups can install printers on the servers. If this policy setting is enabled but the driver for a network printer already exists on the local computer, users can still add the network printer.

Devices: Restrict CD-ROM access to locally logged-on user only

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines whether a CD is accessible to local and remote users simultaneously. If you enable this policy setting, only the interactively logged-on user is allowed to access removable CDs. If this policy setting is enabled and no one is logged on interactively, the CD can be accessed over the network.

The security benefit of enabling this policy setting is small because it only prevents network users from accessing the drive when someone is logged on to the local console of the system at the same time. Additionally, CD drives are not automatically made available as network shared drives; administrators must deliberately choose to share the drive. This is important when administrators are installing software or copying data from a CD-ROM, and they do not want network users to be able to execute the applications or view the data.

If this policy setting is enabled, users who connect to the server over the network will not be able to use any CD drives that are installed on the server when anyone is logged on to the local console of the server. Enabling this policy setting is not suitable for a system that serves as a CD jukebox for network users.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

- Best practices are dependent on your security and user accessibility requirements for CD drives.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A remote user could potentially access a mounted CD that contains sensitive information. This risk is small because CD drives are not automatically made available as shared drives; administrators must deliberately choose to share the drive. However, administrators can deny network users the ability to view data or run applications from removable media on the server.

Countermeasure

Enable the **Devices: Restrict CD-ROM drive access to locally logged-on user only** setting.

Potential impact

Users who connect to the server over the network cannot use any CD drives that are installed on the server when anyone is logged on to the local console of the server. System tools that require access to the CD drive will fail. For example, the Volume Shadow Copy service attempts to access all CD and floppy disk drives that are present on the computer when it initializes, and if the service cannot access one of these drives, it fails. This condition causes the Windows Backup tool to fail if volume shadow copies were specified for the backup job. Any non-Microsoft backup products that use volume shadow copies also fail. This policy setting would not be suitable for a computer that serves as a CD jukebox for network users.

Devices: Restrict floppy access to locally logged-on user only

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines whether removable floppy disks are accessible to local and remote users simultaneously. Enabling this policy setting allows only the interactively logged-on user to access removable floppy disks. If this policy setting is enabled and no one is logged on interactively, the floppy disk can be accessed over the network.

The security benefit of enabling this policy setting is small because it only prevents network users from accessing the floppy disk drive when someone is logged on to the local console of the system at the same time. Additionally, floppy disk drives are not automatically made available as network shared drives; administrators must deliberately choose to share the drive. This becomes important when administrators are installing software or copying data from a floppy disk and they do not want network users to be able to execute the applications or view the data.

If this policy setting is enabled, users who connect to the server over the network will not be able to use any floppy disk drives that are installed on the server when anyone is logged on to the local console of the server.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

- Best practices are dependent on your security and user accessibility requirements for CD drives.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disdabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A remote user could potentially access a mounted floppy disk that contains sensitive information. This risk is small because floppy disk drives are not automatically shared; administrators must deliberately choose to share the drive. However, administrators can deny network users the ability to view data or run applications from removable media on the server.

Countermeasure

Enable the **Devices: Restrict floppy access to locally logged-on user only** setting.

Potential impact

Users who connect to the server over the network cannot use any floppy disk drives that are installed on the server when anyone is logged on to the local console of the server. System tools that require access to floppy disk drives fail. For example, the Volume Shadow Copy service attempts to access all CD-ROM and floppy disk drives that are present on the computer when it initializes, and if the service cannot access one of these drives, it fails. This condition causes the Windows Backup tool to fail if volume shadow copies were specified for the backup job. Any non-Microsoft backup products that use volume shadow copies also fail.

Domain controller: Allow server operators to schedule tasks

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines whether server operators are allowed to submit jobs by means of the **at** command. If you enable this policy setting, jobs that are created by server operators by means of the **at** command run in the context of the account that runs the Task Scheduler service. By default, that is the Local System account.

Note
This security option setting affects only the scheduler tool for the **at** command. It does not affect the Task Scheduler tool.

Enabling this policy setting means jobs that are created by server operators through the **at** command will be executed in the context of the account that is running that service—by default, that is the Local System account. This means that server operators can perform tasks that the Local System account is able to do, but server operators would normally not be able to do, such as add their account to the local Administrators group.

The impact of enabling this policy setting should be small for most organizations. Users, including those in the Server Operators group, will still be able to create jobs by using the Task Scheduler Wizard, but those jobs will run in the context of the account that the user authenticates with when setting up the job.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

- Best practices for this policy are dependent on your security and operational requirements for task scheduling.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Command-line tools

The **at** command schedules commands and programs to run on a computer at a specified time and date. The Schedule service must be running to use the **at** command. For more information about the **at** command, see **AT**.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Tasks that run under the context of the Local System account can affect resources that are at a higher privilege level than the user account that scheduled the task.

Countermeasure

Disable the **Domain controller: Allow server operators to schedule tasks** setting.

Potential impact

The impact should be small for most organizations. Users (including those in the Server Operators group) can still create jobs by means of the Task Scheduler snap-in. However, those jobs run in the context of the account that the user authenticates with when setting up the job.

Domain controller: LDAP server signing requirements

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines whether the Lightweight Directory Access Protocol (LDAP) server requires LDAP clients to negotiate data signing.

Unsigned network traffic is susceptible to man-in-the-middle attacks, where an intruder captures packets between the server and the client computer and modifies them before forwarding them to the client computer. In the case of an LDAP server, this means that a malicious user can cause a client computer to make decisions based on false records from the LDAP directory. You can lower the risk of a malicious user accomplishing this in a corporate network by implementing strong physical security measures to protect the network infrastructure. Furthermore, implementing Internet Protocol security (IPsec) Authentication Header mode, which provides mutual authentication and packet integrity for IP traffic, can make all types of man-in-the-middle attacks extremely difficult.

This setting does not have any impact on LDAP simple bind or LDAP simple bind through SSL.

If signing is required, then LDAP simple bind and LDAP simple bind through SSL requests are rejected. No LDAP client computers running Windows XP Professional or servers running Windows Server 2003 use LDAP simple bind or LDAP simple bind through SSL to bind to directory service.

Caution

If you set the server to Require signature, you must also set the client computer. Not setting the client computer results in loss of connection with the server.

Possible values

- None. Data signatures are not required to bind with the server. If the client computer requests data signing, the server supports it.
- Require signature. The LDAP data-signing option must be negotiated unless Transport Layer Security/Secure Sockets Layer (TLS/SSL) is in use.
- Not defined.

Best practices

1. It is advisable to set **Domain controller: LDAP server signing requirements** to **Require signature**. Clients that do not support LDAP signing will be unable to execute LDAP queries against the domain controllers.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	None
Member Server Effective Default Settings	None
Client Computer Effective Default Settings	None

Operating system version differences

Modifying this setting may affect compatibility with clients, services, and applications. For more information about these compatibility issues, see article 823659 in the Microsoft Knowledge database.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Unsigned network traffic is susceptible to man-in-the-middle attacks. In such attacks, an intruder captures packets between the server and the client computer, modifies them, and then forwards them to the client computer. Where LDAP servers are concerned, an attacker could cause a client computer to make decisions that are based on false records from the LDAP directory. To lower the risk of such an intrusion in an organization's network, you can implement strong physical security measures to protect the network infrastructure. You could also implement Internet Protocol security (IPsec) Authentication Header mode, which performs mutual authentication and packet integrity for IP traffic to make all types of man-in-the-middle attacks extremely difficult.

Countermeasure

Configure the **Domain controller: LDAP server signing requirements** setting to **Require signature**.

Potential impact

Client computers that do not support LDAP signing cannot run LDAP queries against the domain controllers.

Domain controller: Refuse machine account password changes

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting enables or disables blocking a domain controller from accepting password change requests for computer accounts. By default, member computers change their computer account passwords every 30 days. If enabled, the domain controller will refuse computer account password change requests.

Possible values

- Enabled

  When enabled, this setting does not allow a domain controller to accept any changes to a computer account's password.
- Disabled

  When disabled, this setting allows a domain controller to accept any changes to a computer account's password.

- Not defined

  Same as Disabled.

Best practices

1. Enabling this policy setting on all domain controllers in a domain prevents domain members from changing their computer account passwords. This, in turn, leaves those passwords susceptible to attack. Make sure that this conforms to your overall security policy for the domain.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Not applicable

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you enable this policy setting on all domain controllers in a domain, domain members cannot change their computer account passwords, and those passwords are more susceptible to attack.

Countermeasure

Disable the **Domain controller: Refuse machine account password changes** setting.

Potential impact

None. This is the default configuration.

Domain member: Digitally encrypt or sign secure channel data (always)

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This setting determines whether all secure channel traffic that is initiated by the domain member meets minimum security requirements. Specifically, it determines whether all secure channel traffic that is initiated by the domain member must be signed or encrypted. Logon information that is transmitted over the secure channel is always encrypted regardless of whether the encryption of all other secure channel traffic is negotiated.

The following policy settings determine whether a secure channel can be established with a domain controller that is not capable of signing or encrypting secure channel traffic:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally encrypt secure channel data (when possible)
- Domain member: Digitally sign secure channel data (when possible)

Setting **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled** prevents establishing a secure channel with any domain controller that cannot sign or encrypt all secure channel data.

To protect authentication traffic from man-in-the-middle, replay, and other types of network attacks, Windows-based computers create a communication channel through NetLogon called

secure channels. These channels authenticate computer accounts. They also authenticate user accounts when a remote user connects to a network resource and the user account exists in a trusted domain. This is called pass-through authentication, and it allows a computer running the Windows operating system that has joined a domain to have access to the user account database in its domain and in any trusted domains.

To enable the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting on a member workstation or server, all domain controllers in the domain that the member belongs to must be capable of signing or encrypting all secure-channel data.

Enabling the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting automatically enables the **Domain member: Digitally sign secure channel data (when possible)** policy setting.

When a computer running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 joins a domain, a computer account is created. After joining the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. This secure channel is used to perform operations such as NTLM pass-through authentication and LSA SID/name Lookup. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the integrity of the channel is not checked, and not all information is encrypted. If a system is set to always encrypt or sign secure channel data, a secure channel cannot be established with a domain controller that is not capable of signing or encrypting all secure channel traffic. If the computer is configured to encrypt or sign secure channel data when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Possible values

- Enabled

  The policy **Domain member: Digitally sign secure channel data (when possible)** is assumed to be enabled regardless of its current setting. This ensures that the domain member attempts to negotiate at least signing of the secure channel traffic.

- Disabled

  The encryption and signing of all secure channel traffic is negotiated with the domain controller, in which case the level of signing and encryption depends on the version of the domain controller and the settings of the following policies:

  a. **Domain member: Digitally encrypt secure channel data (when possible)**
  b. **Domain member: Digitally sign secure channel data (when possible)**

- Not defined

Best practices

- Set **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled**.

- Set **Domain member: Digitally encrypt secure channel data (when possible)** to **Enabled**.
- Set **Domain member: Digitally sign secure channel data (when possible)** to **Enabled**.

You can enable this policy setting after you eliminate all Windows 95, Windows 98, and Windows Millennium Edition clients from the domain and upgrade all Windows NT 4.0 servers and domain controllers from trusted or trusting domains to Windows NT 4.0 with SP6a.

Note
You can enable the policy settings **Domain member: Digitally encrypt secure channel data (when possible)** and **Domain member: Digitally sign secure channel data (when possible)** on all computers in the domain that support these policy settings without affecting earlier-version clients and applications.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Enabled
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

Modifying this setting may affect compatibility with clients, services, and applications, including those running in operating systems earlier than Windows 2000 Server.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distribution of this policy through Group Policy overrides the Local Security Policy setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When a computer joins a domain, a computer account is created. After it joins the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the channel is not integrity-checked, and not all information is encrypted. If a computer is configured to always encrypt or sign secure channel data but the domain controller cannot sign or encrypt any portion of the secure channel data, the computer and domain controller cannot establish a secure channel. If the computer is configured to encrypt or sign secure channel data, when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Countermeasure

Select one of the following settings as appropriate for your environment to configure the computers in your domain to encrypt or sign secure channel data.

- **Domain member: Digitally encrypt or sign secure channel data (always)**
- **Domain member: Digitally encrypt secure channel data (when possible)**
- **Domain member: Digitally sign secure channel data (when possible)**

Potential impact

Digital encryption and signing of the secure channel is a good idea where it is supported. The secure channel protects domain credentials as they are sent to the domain controller. However, operating systems earlier than Windows NT 4.0 with SP6a support digital encryption and signing of the secure channel. Client computers running Windows 98 Second Edition do not support it unless they have the Active Directory Client Extension installed). Therefore, you cannot enable the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting on domain controllers that support client computers running Windows 98 as members of the domain. The potential impact of such a setting can include making it impossible to create or delete trust

relationships with earlier versions, disabling logons from earlier-version clients, and not being able to authenticate other domains' users from a trusted domain that is running an earlier version of the Windows operating system.

- The ability to create or delete trust relationships with client computers running versions of Windows earlier than Windows NT 4.0 with SP6a is disabled.
- Logons from client computers running versions of Windows earlier than Windows NT 4.0 with SP6a are disabled.
- The ability to authenticate other domains' users from a domain controller running a version of Windows earlier than Windows NT 4.0 with SP6a in a trusted domain is disabled.

You can enable this policy setting after you eliminate all client computers running Windows 95 or Windows 98 from the domain and upgrade all servers running Windows NT 4.0 and domain controllers from trusted/trusting domains to Windows NT 4.0 with SP6a. You can enable the policy settings **Domain member: Digitally encrypt secure channel data (when possible)** and **Domain member: Digitally encrypt sign channel data (when possible)** on all computers in the domain that support them and client computers running versions of the Windows operating system earlier than Windows NT 4.0 with SP6a. Applications that run on these versions of the Windows operating system are not affected.

Domain member: Digitally encrypt secure channel data (when possible)

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This setting determines whether all secure channel traffic that is initiated by the domain member meets minimum security requirements. Specifically, it determines whether all secure channel traffic that is initiated by the domain member must be encrypted. Logon information that is transmitted over the secure channel is always encrypted regardless of whether the encryption of all other secure channel traffic is negotiated.

In addition to this policy setting, the following policy settings determine whether a secure channel can be established with a domain controller that is not capable of signing or encrypting secure channel traffic:

- Domain member: Digitally encrypt or sign secure channel data (always)
- Domain member: Digitally sign secure channel data (when possible)

Setting **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled** prevents establishing a secure channel with any domain controller that cannot sign or encrypt all secure channel data.

To protect authentication traffic from man-in-the-middle, replay, and other types of network attacks, Windows-based computers create a communication channel through NetLogon called secure channels. These channels authenticate computer accounts. They also authenticate user accounts when a remote user connects to a network resource and the user account exists in a trusted domain. This is called pass-through authentication, and it allows a computer running the Windows operating system that has joined a domain to have access to the user account database in its domain and in any trusted domains.

Enabling the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting automatically enables the **Domain member: Digitally sign secure channel data (when possible)** policy setting.

When computers running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 joins a domain, a computer account is created. After joining the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. This secure channel is used to perform operations such as NTLM pass through authentication and LSA SID/name Lookup. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the integrity of the channel is not checked, and not all information is encrypted. If a system is set to always encrypt or sign secure channel data, a secure channel cannot be established with a domain controller that is not capable of signing or encrypting all secure channel traffic. If the computer is configured to encrypt or sign secure channel data when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Possible values

- Enabled

  The domain member will request encryption of all secure channel traffic. If the domain controller supports encryption of all secure channel traffic, then all secure channel traffic will be encrypted. Otherwise, only logon information that is transmitted over the secure channel will be encrypted.

- Disabled

  The domain member will not attempt to negotiate secure channel encryption.

  Note
  If the security policy setting **Domain member: Digitally encrypt or sign secure channel data (always)** is enabled, this setting will be overwritten.

- Not defined

Best practices

- Set **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled**.
- Set **Domain member: Digitally encrypt secure channel data (when possible)** to **Enabled**.

- Set **Domain member: Digitally sign secure channel data (when possible)** to **Enabled**.

Note
You can enable this policy setting on all computers in the domain that support these policy settings without affecting applications or client computers that are running earlier versions of the Windows operating system.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Enabled
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distribution of this policy through Group Policy does not override the Local Security Policy setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When a computer joins a domain, a computer account is created. After it joins the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the channel is not integrity-checked, and not all information is encrypted. If a computer is configured to always encrypt or sign secure channel data but the domain controller cannot sign or encrypt any portion of the secure channel data, the computer and domain controller cannot establish a secure channel. If the computer is configured to encrypt or sign secure channel data when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Countermeasure

Select one of the following settings as appropriate for your environment to configure the computers in your domain to encrypt or sign secure channel data:

- **Domain member: Digitally encrypt or sign secure channel data (always)**
- **Domain member: Digitally encrypt secure channel data (when possible)**
- **Domain member: Digitally sign secure channel data (when possible)**

Potential impact

Digital signing of the secure channel is a good idea where it is supported. The secure channel protects domain credentials as they are sent to the domain controller. However, operating systems earlier than Windows NT 4.0 with SP6a do not support digital encryption and signing of the secure channel. Client computers running Microsoft Windows 98 Second Edition do not support it unless they have the Active Directory Client Extension installed. Therefore, enabling this policy where it is unsupported on these systems will have no affect.

Domain member: Digitally sign secure channel data (when possible)

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This setting determines whether all secure channel traffic that is initiated by the domain member meets minimum security requirements. Specifically, it determines whether all secure channel traffic that is initiated by the domain member must be signed. Logon information that is transmitted over the secure channel is always encrypted regardless of whether the encryption of all other secure channel traffic is negotiated.

The following policy settings determine whether a secure channel can be established with a domain controller that is not capable of signing or encrypting secure channel traffic:

- **Domain member: Digitally encrypt or sign secure channel data (always)**
- **Domain member: Digitally encrypt secure channel data (when possible)**
- **Domain member: Digitally sign secure channel data (when possible)**

Setting **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled** prevents establishing a secure channel with any domain controller that cannot sign or encrypt all secure channel data.

To protect authentication traffic from man-in-the-middle, replay, and other types of network attacks, Windows-based computers create a communication channel through NetLogon called secure channels. These channels authenticate computer accounts. They also authenticate user accounts when a remote user connects to a network resource and the user account exists in a trusted domain. This is called pass-through authentication, and it allows a computer running the Windows operating system that has joined a domain to have access to the user account database in its domain and in any trusted domains.

Enabling the **Domain member: Digitally encrypt or sign secure channel data (always)** policy setting automatically enables the **Domain member: Digitally sign secure channel data (when possible)** policy setting.

When a computer running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 joins a domain, a computer account is created. After joining the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. This secure channel is used to perform operations such as NTLM pass through authentication and LSA SID/name Lookup. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the integrity of the channel is not checked, and not all information is encrypted. If a system is set to always encrypt or sign secure channel data, a secure channel cannot be established with a domain controller that is not capable of signing or encrypting all secure channel traffic. If the computer is configured to encrypt or sign secure channel data when

possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Possible values

- Enabled

  The domain member will request signing of all secure channel traffic. If the Domain Controller supports signing of all secure channel traffic, then all secure channel traffic will be signed which ensures that it cannot be tampered with in transit.
- Disabled

  Signing will not be negotiated unless the policy **Domain member: Digitally encrypt or sign secure channel data (always)** is enabled.
- Not defined

Best practices

- Set **Domain member: Digitally encrypt or sign secure channel data (always)** to **Enabled**.
- Set **Domain member: Digitally encrypt secure channel data (when possible)** to **Enabled**.
- Set **Domain member: Digitally sign secure channel data (when possible)** to **Enabled**.

Note

You can enable the other two policy settings, **Domain member: Digitally encrypt secure channel data (when possible)** and **Domain member: Digitally sign secure channel data (when possible)**, on all computers in the domain that support these policy settings without affecting earlier-version clients and applications.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Enabled
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled

Server type or GPO	Default value
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

Modifying this setting may affect compatibility with clients, services, and applications, including those running in operating systems earlier than Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distribution of this policy through Group Policy does not override the Local Security Policy setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When a computer joins a domain, a computer account is created. After it joins the domain, the computer uses the password for that account to create a secure channel with the domain controller for its domain every time it restarts. Requests that are sent on the secure channel are authenticated—and sensitive information such as passwords are encrypted—but the channel is not integrity-checked, and not all information is encrypted. If a computer is configured to always encrypt or sign secure channel data but the domain controller cannot sign or encrypt any portion of the secure channel data, the computer and domain controller cannot establish a secure channel. If the computer is configured to encrypt or sign secure channel data when possible, a secure channel can be established, but the level of encryption and signing is negotiated.

Countermeasure

Because these policies are closely related and useful depending on your environment, select one of the following settings as appropriate to configure the computers in your domain to encrypt or sign secure channel data when possible.

- **Domain member: Digitally encrypt or sign secure channel data (always)**
- **Domain member: Digitally encrypt secure channel data (when possible)**
- **Domain member: Digitally sign secure channel data (when possible)**

Potential impact

Digital signing of the secure channel is a good idea where it is supported. The secure channel protects domain credentials as they are sent to the domain controller. However, operating systems earlier than Windows NT 4.0 with SP6a do not support digital encryption and signing of the secure channel. Client computers running Windows 98 Second Edition do not support it unless they have the Active Directory Client Extension installed. Therefore, enabling this policy where it is unsupported on these systems will have no affect.

Domain member: Disable machine account password changes

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Domain member: Disable machine account password changes** policy setting determines whether a domain member periodically changes its computer account password. Setting its value to **Enabled** prevents the domain member from changing the computer account password. Setting it to **Disabled** allows the domain member to change the computer account password as specified by the value of the **Domain member: Maximum machine account password age** policy setting, which is every 30 days by default.

The default configuration for computers running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 that belong to a domain is that they are automatically required to change the passwords for their accounts every 30 days. Disabling this feature causes computers that are running these operating systems to retain the same passwords as their computer accounts. Computers that are no longer able to automatically change their account password are at risk of a malicious user determining the password for the system's domain account.

Verify that the **Domain member: Disable machine account password changes** option is set to **Disabled**.

Possible values

- Enabled
- Disabled

Best practices

1. Do not enable this policy setting. Computer account passwords are used to establish secure channel communications between members and domain controllers and between the domain controllers within the domain. After it is established, the secure channel transmits sensitive information that is necessary for making authentication and authorization decisions.
2. Do not use this policy setting in an attempt to support dual-boot scenarios that use the same computer account. If you want to dual-boot installations that are joined to the same domain, give the two installations different computer names. This policy setting was added to the Windows operating system to make it easier for organizations that stockpile prebuilt computers that are put into production months later; those computers do not have to be rejoined to the domain.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Disabled
Default Domain Controller Policy	Disabled
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

In Windows Server 2003, the default setting for Domain Policy and Domain Controller Policy was not defined. In Windows Server 2008 R2 and Windows Server 2008, it is **Disabled**. There is no difference in the effective setting between these two, however.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The default configuration for Windows Server operating systems that belong to a domain automatically require them to change the passwords for their accounts every certain number of days, typically 30. If you disable this policy setting, computers that run Windows Server operating systems retain the same passwords as their computer accounts. Computers that cannot automatically change their account password are at risk from an attacker who could determine the password for the computer's domain account.

Countermeasure

Verify that the **Domain member: Disable machine account password changes** setting is configured to **Disabled**.

Potential impact

None. This is the default configuration.

Domain member: Maximum machine account password age

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Domain member: Maximum machine account password age** policy setting determines the maximum allowable age for a computer account password.

In Active Directory–based domains, each computer has an account and password, just like every user. By default, the domain members automatically change their domain password every 30 days. Increasing this interval significantly, or setting it to **0** so that the computers no longer change their passwords, gives a malicious user more time to undertake a brute-force password-guessing attack against one of the computer accounts.

Possible values

- User-defined number of days between 0 and 999
- Not defined.

Best practices

1. It is often advisable to set **Domain member: Maximum machine account password age** to about 30 days.
2. Some organizations prebuild computers and then store them for later use or ship them to remote locations. If the computer's account has expired, it will no longer be able to authenticate with the domain. Computers that cannot authenticate with the domain must be removed from the domain and rejoined to it. For this reason, some organizations might want to create a special organizational unit (OU) for computers that are prebuilt, and configure the value for this policy setting to a larger number of days.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	30 days
DC Effective Default Settings	30 days
Member Server Effective Default Settings	30 days
Client Computer Effective Default Settings	30 days

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003. This setting also applies to Windows 2000 computers, but it is not available through the Security Configuration Manager tools on Windows 2000 computers.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

In Active Directory–based domains, each computer has an account and password, just as every user does. By default, the domain members automatically change their domain password every 30 days. If you increase this interval significantly, or set it to 0 so that the computers no longer change their passwords, an attacker has more time to undertake a brute-force attack to guess the password of one or more computer accounts.

Countermeasure

Configure the **Domain member: Maximum machine account password age** setting to 30 days.

Potential impact

None. This is the default configuration.

Domain member: Require strong (Windows 2000 or later) session key

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The **Domain member: Require strong (Windows 2000 or later) session key** policy setting determines whether a secure channel can be established with a domain controller that is not capable of encrypting secure channel traffic with a strong, 128-bit session key. Enabling this policy setting prevents establishing a secure channel with any domain controller that cannot encrypt secure channel data with a strong key. Disabling this policy setting allows 64-bit session keys.

Whenever possible, you should take advantage of these stronger session keys to help protect secure channel communications from eavesdropping and session-hijacking network attacks. Eavesdropping is a form of hacking in which network data is read or altered in transit. The data can be modified to hide or change the name of the sender, or it can be redirected.

Possible values

- Enabled

  When enabled on a member workstation or server, all domain controllers in the domain that the member belongs to must be capable of encrypting secure channel data with a strong, 128-bit key. This means that all such domain controllers must be running the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, or Windows 2000 operating system.
- Disabled

  Allows 64-bit session keys to be used.
- Not defined.

Best practices

1. It is advisable to set **Domain member: Require strong (Windows 2000 or later) session key** to Enabled. Enabling this policy setting ensures that all outgoing secure channel traffic will require a strong encryption key. Disabling this policy setting requires that key strength be negotiated. Only enable this option if the domain controllers in all trusted domains support strong keys. By default, this value is disabled.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

Modifying this setting may affect compatibility with clients, services, and applications running on Windows operating systems beginning with Windows 2000 Server.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

You will not be able to join computers that have this policy setting enabled to Windows NT 4.0 domains, nor will you be able to join computers that do not support this policy setting to domains where the domain controllers have this policy setting enabled.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Session keys that are used to establish secure channel communications between domain controllers and member computers are much stronger in Windows Server 2008 R2, Windows

Server 2008, Windows Server 2003, and Windows 2000 than they were in previous Windows operating systems.

Whenever possible, you should take advantage of these stronger session keys to help protect secure channel communications from attacks that attempt to hijack network sessions and eavesdrop. (Eavesdropping is a form of hacking in which network data is read or altered in transit. The data can be modified to hide or change the sender, or be redirected.)

Countermeasure

Enable the **Domain member: Require strong (Windows 2000 or later) session key** setting.

If you enable this policy setting, all outgoing secure channel traffic requires a strong encryption key. If you disable this policy setting, the key strength is negotiated. You should enable this policy setting only if the domain controllers in all trusted domains support strong keys. By default, this policy setting is disabled.

Potential impact

Computers that have this policy setting enabled cannot join Windows NT 4.0 domains, and trusts between Active Directory domains and Windows NT domains may not work properly. Also, computers that do not support this policy setting cannot join domains in which the domain controllers have this policy setting enabled.

Interactive logon: Display user information when the session is locked

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

When a session is locked in a Windows operating system (meaning the user at the computer pressed CTRL+ALT+DEL and the Secure Desktop is displayed), user information is displayed. By default, this information is in the form of **<user name> is logged on**. The displayed user name is the user's full name as set on the Properties page for that user. These settings do not apply to the logon tiles, which are displayed on the desktop after using the **Switch User** feature. The information that is displayed can be changed to meet your security requirements using the following possible values.

Possible values

- **User display name, domain and user names**

If this is a local logon, the user's full name is displayed on the Secure Desktop. If it is a domain logon, the user's domain and user's account name is displayed.

- **User display name only**

  The name of the user who locked the session is displayed on the Secure Desktop as the user's full name.

- **Do not display user information**

  No names are displayed on the Secure Desktop, but user's full names will be displayed on the **Switch user** desktop.

- Blank.

  Default setting. This translates to "Not defined," but it will display the user's full name in the same manner as the **User display name, domain and user names** option. When an option is set, you cannot reset this policy to blank, or not defined.

Best practices

Your implementation of this policy depends on your security requirements for displayed logon information. If you have running computers that store sensitive data, with monitors displayed in unsecured locations, or if you have computers with sensitive data that are remotely accessed, revealing logged on user's full names or domain account names might contradict your overall security policy.

Depending on your security policy, you might also want to enable the **Interactive logon: Do not display last user name** policy, which will prevent the Windows operating system from displaying the logon name and logon tile of the last user to logon.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or Group Policy object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	**User display name, domain and user names**
Member server effective default settings	**User display name, domain and user names**
Effective GPO default settings on client computers	**User display name, domain and user names**

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When a computer displays the Secure Desktop in an unsecured area, certain user information can be readily available to anyone looking at the monitor, either physically or through a remote connection. The displayed user information could include the domain user account name or the full name of the user who locked the session or who had logged on last.

Countermeasure

Enabling this policy setting allows the operating system to hide certain user information from being displayed on the Secure Desktop (after the computer has been booted or when the session has been locked by using CTRL+ALT+DEL). However, user information is displayed if the **Switch user** feature is used so that the logon tiles are displayed for each logged on user.

You might also want to enable the **Interactive logon: Do not display last user name** policy, which will prevent the Windows operating system from displaying the logon name and logon tile of the last user to logon.

Potential impact

If you do not enable this policy, the effect will be the same as enabling the policy and selecting the **User display name, domain and user names** option.

If the policy is enabled and set to **Do not display user information**, an observer cannot see who is logged onto the Secure Desktop, but the logon tile is still present if the **Interactive logon: Do not display last user name** policy is not enabled. Depending on how the logon tiles are configured, they could provide visual clues as to who is logged on. In addition, if the **Interactive logon: Do not display last user name** policy is not enabled, then the **Switch user** feature will show user information.

Interactive logon: Do not display last user name

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This security setting determines whether the name of the last user to log on to the computer is displayed on the Secure Desktop.

If this policy is enabled, the full name of the last user to successfully log on is not displayed on the Secure Desktop, nor is the user's logon tile displayed. Additionally, if the **Switch user** feature is used, the full name and logon tile are not displayed. The logon screen requests a qualified domain account name (or local user name) and password.

If this policy is disabled, the full name of the last user to log on is displayed, and the user's logon tile is displayed. This behavior is the same when the **Switch user** feature is used.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

Your implementation of this policy depends on your security requirements for displayed logon information. If you have running computers that store sensitive data, with monitors displayed in

unsecured locations, or if you have computers with sensitive data that are remotely accessed, revealing logged on user's full names or domain account names might contradict your overall security policy.

Depending on your security policy, you might also want to enable the **Interactive logon: Display user information when the session is locked** policy, which will prevent the Windows operating system from displaying the logon name when the session is locked or started.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or Group Policy object (GPO)	Default value
Default domain policy	Disabled
Default domain controller policy	Disabled
Stand-alone server default settings	Disabled
Domain controller effective default settings	Disabled
Member server effective default settings	Disabled
Effective GPO default settings on client computers	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008 and Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An attacker with access to the console (for example, someone with physical access or someone who can connect to the computer through Remote Desktop Session Host (RD Session Host)) could view the name of the last user who logged on. The attacker could then try to guess the password, use a dictionary, or use a brute-force attack to try to log on.

Countermeasure

Enable the **Interactive logon: Do not display last user name** setting.

Potential impact

Users must always type their user names and passwords when they log on locally or to the domain. The logon tiles of all logged on users are not displayed.

Interactive logon: Do not require CTRL+ALT+DEL

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This security setting determines whether pressing CTRL+ALT+DEL is required before a user can log on.

If this policy setting is enabled on a computer, a user is not required to press CTRL+ALT+DEL to log on. Not having to press CTRL+ALT+DEL leaves users susceptible to attacks that attempt to intercept the users' passwords. Requiring CTRL+ALT+DEL before users log on ensures that users are communicating by means of a trusted path when entering their passwords.

If this policy is disabled, any user is required to press CTRL+ALT+DEL before logging on to the Windows operating system (unless they are using a smart card for logon).

Important
The default behavior of smart card logon has changed from Windows Server 2003 and Windows XP. See Operating system version differences for more information.

Microsoft developed this feature to make it easier for users with certain types of physical impairments to log on to computers running the Windows operating system; however, not having to press the CTRL+ALT+DELETE key combination leaves users susceptible to attacks that attempt to intercept their passwords. Requiring CTRL+ALT+DELETE before users log on ensures that users are communicating by means of a trusted path when entering their passwords.

A malicious user might install a Trojan horse program that looks like the standard logon dialog box for the Windows operating system, and capture a user's password. The attacker can then log on to the compromised account with whatever level of user rights that user has.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. It is advisable to set **Disable CTRL+ALT+DEL requirement for logon** to **Disabled**. Unless they are using a smart card to log on, users will have to simultaneously press three keys before the logon dialog box appears.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled

Server type or GPO	Default value
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are differences in this policy between operating systems beginning with Windows Server 2003 and Windows XP.

In Windows Server 2003 and Windows XP, a smart card logon does not require the CTRL+ALT+DEL key combination if this policy is disabled. However, beginning with Windows Server 2008 and Windows Vista, the CTRL+ALT+DELETE key combination is required to authenticate if this policy is disabled.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

In Windows Server 2003 and Windows XP, a smart card logon does not require the CTRL+ALT+DEL key combination if this policy is disabled. However, beginning with Windows Server 2008 and Windows Vista, the CTRL+ALT+DELETE key combination is required to authenticate if this policy is disabled.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This setting makes it easier for users with certain types of physical impairments to log on to computers that run the Windows operating system. However, if users are not required to press CTRL+ALT+DEL, they are susceptible to attacks that attempt to intercept their passwords. If CTRL+ALT+DEL is required before logon, user passwords are communicated by means of a trusted path.

If this setting is enabled, an attacker could install a Trojan horse program that looks like the standard logon dialog box in the Windows operating system, and capture the user's password. The attacker would then be able to log on to the compromised account with whatever level of privilege that user has.

Countermeasure

Disable the **Interactive logon: Do not require CTRL+ALT+DEL** setting.

Potential impact

Unless they use a smart card to log on, users must simultaneously press the three keys before the logon dialog box is displayed.

Interactive logon: Machine inactivity limit

This security policy reference topic for the IT professional describes the best practices, location, values, management, and security considerations for this policy setting.

Reference

Beginning with Windows Server 2012 and Windows 8, Windows detects user-input inactivity of a sign-in (logon) session by using the security policy setting **Interactive logon: Machine inactivity limit**. If the amount of inactive time exceeds the inactivity limit set by this policy, then the user's session locks by invoking the screen saver. This policy setting allows you to control the locking time by using Group Policy.

In versions earlier than Windows Server 2012 and Windows 8, the desktop-locking mechanism was set on individual computers in Personalization in Control Panel.

Possible values

The automatic lock of the computer is set in elapsed seconds of inactivity, which can range from zero (0) to 599,940 seconds (166.65 hours).

If no value (blank) or zero (0) is present in the **Machine will be locked after** input field, then the policy setting is disabled and no action is taken on user-input inactivity for the session.

Best practices

Set the time for elapsed user-input inactivity based on the computer's usage and location requirements. For example, if the computer or device is in a public area, you might want to have the desktop automatically lock after a short period of inactivity to prevent unauthorized access. However, if the computer or device is used by an individual or group of trusted individuals, such as in a restricted manufacturing area, automatically locking the desktop might hinder productivity.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

This policy setting was introduced in Windows Server 2012 and Windows 8.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

Restart is required for changes to this policy to become effective when they are saved locally or distributed through Group Policy.

Group Policy

Because this policy setting was introduced in Windows Server 2012 and Windows 8, it can only be set locally on those computers that contain this policy setting, but it can be set and distributed through Group Policy to any computer running the Windows operating system that supports Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This policy setting helps you prevent unauthorized access to computers or devices under your control when the currently signed-in user leaves without deliberately locking the desktop. In versions earlier than Windows Server 2012 and Windows 8, the desktop-locking mechanism was set on individual computers in Personalization in Control Panel.

Countermeasure

Set the time for elapsed user-input inactivity time by using the security policy setting **Interactive logon: Machine inactivity limit** based on the computer's usage and location requirements.

Potential impact

This security policy setting can limit unauthorized access to unsecured computers; however, that requirement must be balanced with the productivity requirements of the intended computer user.

See Also

Security Options

Interactive logon: Machine account lockout threshold

This security policy reference topic for the IT professional describes the best practices, location, values, management, and security considerations for this policy setting.

Reference

Beginning with Windows Server 2012 and Windows 8, the **Interactive logon: Machine account threshold** security policy setting enforces the lockout policy on those computers that have BitLocker enabled to protect operating system volumes.

The security setting allows you to set a threshold for the number of failed logon attempts that causes the computer or device to be locked by using BitLocker. This means, if the specified maximum number of failed logon attempts is exceeded, the device will invalidate the Trusted Platform Module (TPM) protector and any other protector except the 48-digit recovery password, and then reboot. During Device Lockout mode, the computer or device only boots into the touch-enabled Windows Recovery Environment (WinRE) until an authorized user enters the recovery password to restore full access.

Failed password attempts on workstations or member servers that have been locked by using either Ctrl+Alt+Delete or password-protected screen savers count as failed logon attempts.

Possible values

You can set the **invalid logon attempts** value between 1 and 999. Values from 1 to 3 are interpreted as 4. If you set the value to 0, or leave blank, the computer or device will never be locked as a result of this policy setting.

Best practices

Use this policy setting in conjunction with your other failed account logon attempts policy. For example, if the **Account lockout threshold** policy setting is set at 4, then setting **Interactive logon: Machine account lockout threshold** at 6 allows the user to restore access to resources without having to restore access to the computer or device resulting from a BitLocker lock out.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled

Server type or GPO	Default value
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

This policy setting was introduced in Windows Server 2012 and Windows 8.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

A restart is required for changes to this policy to become effective when they are saved locally or distributed through Group Policy.

Group Policy

Because this policy setting was introduced in Windows Server 2012 and Windows 8, it can only be set locally on those computers that contain this policy setting, but it can be set and distributed through Group Policy to any computer running the Windows operating system that supports Group Policy and is BitLocker-enabled.

When setting this policy, consider the **Account lockout threshold** policy setting, which determines the number of failed logon attempts that will cause a user account to be locked out. For more information, see Best practices.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This policy setting helps protect a BitLocker-encrypted device from attackers attempting to brute-force guess the Windows sign-in password. If not set, then attackers can attempt innumerable passwords, if no other account protection mechanisms are in place.

Countermeasure

Use this policy setting in conjunction with your other failed account logon attempts policy. For example, if the **Account lockout threshold** policy setting is set at 4, then setting **Interactive logon: Machine account lockout threshold** at 6 allows the user to restore access to resources without having to restore access to the computer or device resulting from a BitLocker lock out.

Potential impact

If not set, the computer or device could be compromised by an attacker using brute-force password cracking software.

If set too low, productivity might be hindered because users who become locked out will be unable to access the computer or device without providing the 48-digit BitLocker recovery password.

See Also

Security Options

Interactive logon: Message text for users attempting to log on

This security policy reference topic for the IT professional describes the best practices, location, values, management, and security considerations for this policy setting.

Reference

The **Interactive logon: Message text for users attempting to log on** and **Interactive logon: Message title for users attempting to log on** policy settings are closely related. **Interactive logon: Message text for users attempting to log on** specifies a text message to be displayed to users when they log on. **Interactive logon: Message title for users attempting to log on** specifies a title to appear in the title bar of the window that contains the text message. This text is often used for legal reasons—for example, to warn users about the ramifications of misusing company information, or to warn them that their actions might be audited.

Not using this warning-message policy setting leaves your organization legally vulnerable to trespassers who unlawfully penetrate your network. Legal precedents have established that organizations that display warnings to users who connect to their servers over a network have a higher rate of successfully prosecuting trespassers.

When these policy settings are configured, users will see a dialog box before they can log on to the server console.

Possible values

The possible values for this Group Policy setting are:

- User-defined text
- Not defined

Best practices

1. It is advisable to set **Interactive logon: Message text for users attempting to log on** to a value similar to one of the following:
   a. IT IS AN OFFENSE TO CONTINUE WITHOUT PROPER AUTHORIZATION.
   b. This system is restricted to authorized users. Individuals who attempt unauthorized access will be prosecuted. If you are unauthorized, terminate access now. Click OK to indicate your acceptance of this information.

Important

Any warning that you display in the title or text should be approved by representatives from your organization's legal and human resources departments.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes different requirements to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Effect on Windows 2000 operating systems

In Windows 7, Windows Vista, and Windows XP Professional, configuring logon banners that exceed 512 characters in length and contain carriage–return, line-feed sequences is supported. However, client computers running Windows 2000 cannot interpret and display message text that is created by computers running Windows XP Professional or later. You must use a Windows 2000-based computer to create a logon message that applies to Windows 2000-based computers. If you inadvertently create a logon message by using a Windows XP Professional or later operating system, and you discover that it is not displayed properly on Windows 2000-based computers, do the following:

- Undefine the setting.
- Redefine the setting by using a computer running Windows 2000.

Note
Simply changing a logon message that is defined by a computer running Windows 7, Windows Vista, or Windows XP Professional by using a Windows 2000-based computer does not work. The setting must be undefined first.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

There are two policy settings that relate to logon displays:

- **Interactive logon: Message text for users attempting to log on**
- **Interactive logon: Message title for users attempting to log on**

The first policy setting specifies a text message that displays to users when they log on, and the second policy setting specifies a title for the title bar of the text message window. Many organizations use this text for legal purposes; for example, to warn users about the ramifications of misuse of company information, or to warn them that their actions may be audited.

Vulnerability

Users often do not understand the importance of security practices. However, the display of a warning message before logon may help prevent an attack by warning malicious or uninformed

users about the consequences of their misconduct before it happens. It may also help reinforce corporate policies by notifying employees of appropriate policies during the logon process.

Countermeasure

Configure the **Interactive logon: Message text for users attempting to log on** and **Interactive logon: Message title for users attempting to log on** settings to an appropriate value for your organization.

Note
Any warning message that displays should be approved by your organization's legal and human resources representatives.

Potential impact

Users see a message in a dialog box before they can log on to the server console.

Note
Client computers running Windows 2000 cannot interpret and display messages that exceed 512 characters in length and contain carriage-return, line-feed sequences. You must use a Windows 2000-based computer to create a logon message policy that applies to Windows 2000-based computers. If you do not create a logon message policy on a Windows 2000-based computer, and you discover that it does not display properly on Windows 2000-based computers, you must first change the policy setting to **Not Defined** and then reconfigure the setting by using a Windows 2000-based computer. If you do not do this, the changes do not take effect properly.

Interactive logon: Message title for users attempting to log on

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This security setting allows you to specifiy a title that appears in the title bar of the window that contains the **Interactive logon: Message text for users attempting to log on**. This text is often used for legal reasons—for example, to warn users about the ramifications of misusing company information, or to warn them that their actions might be audited.

The **Interactive logon: Message text for users attempting to log on** and **Interactive logon: Message title for users attempting to log on** policy settings are closely related. **Interactive**

logon: Message text for users attempting to log on specifies a text message to be displayed to users when they log on.

Not using this warning-message policy setting leaves your organization legally vulnerable to trespassers who unlawfully penetrate your network. Legal precedents have established that organizations that display warnings to users who connect to their servers over a network have a higher rate of successfully prosecuting trespassers.

When these policy settings are configured, users will see a dialog box before they can log on to the server console.

Possible values

- *User-defined title*
- Not defined

Best practices

1. It is advisable to set **Interactive logon: Message title for users attempting to log on** to a value similar to one the following:
   - RESTRICTED SYSTEM

     or
   - WARNING: This system is restricted to authorized users.
2. Set the policy **Interactive logon: Message text for users attempting to log on** to reinforce the meaning of the message's title.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined

Server type or GPO	Default value
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Effect on Windows 2000 operating systems

In Windows 7, Windows Vista, and Windows XP Professional, configuring logon banners that exceed 512 characters in length and contain carriage–return, line-feed sequences is supported. However, client computers running Windows 2000 cannot interpret and display message text that is created by computers running Windows XP Professional or later. You must use a Windows 2000-based computer to create a logon message that applies to Windows 2000-based computers. If you inadvertently create a logon message by using a Windows XP Professional or later operating system, and you discover that it is not displayed properly on Windows 2000-based computers, do the following:

- Undefine the setting.
- Redefine the setting by using a computer running Windows 2000.

Note
Simply changing a logon message that is defined by a computer running Windows 7, Windows Vista, or Windows XP Professional by using a Windows 2000-based computer does not work. The setting must be undefined first.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

There are two policy settings that relate to logon displays:

- **Interactive logon: Message text for users attempting to log on**

- **Interactive logon: Message title for users attempting to log on**

The first policy setting specifies a text message that displays to users when they log on, and the second policy setting specifies a title for the title bar of the text message window. Many organizations use this text for legal purposes; for example, to warn users about the ramifications of misuse of company information, or to warn them that their actions may be audited.

Vulnerability

Users often do not understand the importance of security practices. However, the display of a warning message with an appropriate title before logon may help prevent an attack by warning malicious or uninformed users about the consequences of their misconduct before it happens. It may also help reinforce corporate policies by notifying employees of appropriate policies during the logon process.

Countermeasure

Configure the **Interactive logon: Message text for users attempting to log on** and **Interactive logon: Message title for users attempting to log on** settings to an appropriate value for your organization.

Note

Any warning message that displays should be approved by your organization's legal and human resources representatives.

Potential impact

Users see a message in a dialog box before they can log on to the server console.

Note

Client computers running Windows 2000 cannot interpret and display messages that exceed 512 characters in length and contain carriage-return, line-feed sequences. You must use a Windows 2000-based computer to create a logon message policy that applies to Windows 2000-based computers. If you do not create a logon message policy on a Windows 2000-based computer, and you discover that it does not display properly on Windows 2000-based computers, you must first change the policy setting to **Not Defined** and then reconfigure the setting by using a Windows 2000-based computer. If you do not do this, the changes do not take effect properly.

Important

If you do not reconfigure this setting to **Not Defined** before you change the setting using a Windows 2000-based computer, the changes do not take effect properly.

Interactive logon: Number of previous logons to cache (in case domain controller is not available)

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The **Interactive logon: Number of previous logons to cache (in case domain controller is not available**) policy setting determines whether a user can log on to a Windows domain by using cached account information. Logon information for domain accounts can be cached locally so that, if a domain controller cannot be contacted on subsequent logons, a user can still log on. This policy setting determines the number of unique users whose logon information is cached locally.

If a domain controller is unavailable and a user's logon information is cached, the user is prompted with the following message:

A domain controller for your domain could not be contacted. You have been logged on using cached account information. Changes to your profile since you last logged on might not be available.

If a domain controller is unavailable and a user's logon information is not cached, the user is prompted with this message:

The system cannot log you on now because the domain *DOMAIN NAME* is not available.

The value of this policy setting indicates the number of users whose logon information the server caches locally. If the value is 10, the server caches logon information for 10 users. When an eleventh user logs on to the computer, the server overwrites the oldest cached logon session.

Users who access the server console will have their logon credentials cached on that server. A malicious user who is able to access the file system of the server can locate this cached information and use a brute-force attack to determine user passwords. Windows mitigates this type of attack by encrypting the information and keeping the cached credentials in the system's registries, which are spread across numerous physical locations.

Possible values

- A user-defined number from 0 through 50
- Not defined

Best practices

It is advisable to set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to 0. Setting this value to 0 disables the local caching of logon

information. Additional countermeasures include enforcing strong password policies and physically securing the computers. If the value is set to 0, users will be unable to log on to any computers if there is no domain controller available to authenticate them. Organizations might want to set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to 2 for end-user systems, especially for mobile users. Setting this value to 2 means that the user's logon information will still be in the cache even if a member of the IT department has recently logged on to their computer to perform system maintenance. This way, those users will be able to log on to their computers when they are not connected to the corporate network.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	10 logons
DC Effective Default Settings	10 logons
Member Server Effective Default Settings	10 logons
Client Computer Effective Default Settings	10 logons

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The number that is assigned to this policy setting indicates the number of users whose logon information is cache locally by the servers. If the number is set to 10, the server caches logon information for 10 users. When an eleventh user logs on to the computer, the server overwrites the oldest cached logon session.

Users who access the server console have their logon credentials cached on that server. An attacker who is able to access the file system of the server could locate this cached information and use a brute force attack to attempt to determine user passwords.

To mitigate this type of attack, Windows encrypts the information and obscures its physical location.

Countermeasure

Configure the **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** setting to 0, which disables the local caching of logon information. Additional countermeasures include enforcement of strong password policies and physically secure locations for the computers.

Potential impact

Users cannot log on to any computers if there is no domain controller available to authenticate them. Organizations can configure this value to 2 for end-user computers, especially for mobile users. A configuration value of 2 means that the user's logon information is still in the cache, even if a member of the IT department has recently logged on to the computer to perform system maintenance. This method allows users to log on to their computers when they are not connected to the organization's network.

Interactive logon: Prompt user to change password before expiration

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The **Interactive logon: Prompt user to change password before expiration** policy setting determines how many days in advance users are warned that their passwords are about to expire. With this advance warning, the user has time to construct a password that is sufficiently strong.

Possible values

- A user-defined number of days from 0 through 999.
- Not defined.

Best practices

1. Configure user passwords to expire periodically. Users will need warning that their passwords are going to expire, or they might inadvertently get locked out of the system. This could lead to confusion for users who access the network locally, or make it impossible for users who access the network through dial-up or virtual private network (VPN) connections to log on.
2. Set **Interactive logon: Prompt user to change password before expiration** to 5 days. When their password expiration date is 5 or fewer days away, users will see a dialog box each time they log on to the domain.
3. Do not set the value to 0, which results in displaying the password expiration warning every time the user logs on.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	14 days *
DC Effective Default Settings	14 days *
Member Server Effective Default Settings	14 days *
Client Computer Effective Default Settings	14 days *

Operating system version differences

Beginning with Windows Server 2008 R2 and Windows 7, the default value changed to 5 days.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If user passwords are configured to expire periodically in your organization, users need to be warned when this is about to happen, or they may be locked out of the computer inadvertently

when their passwords expire. This condition could lead to confusion for users who access the network locally, or make it impossible for users to access your organization's network through dial-up or virtual private network (VPN) connections.

Countermeasure

Configure the **Interactive logon: Prompt user to change password before expiration** setting to 14 days.

Potential impact

Users see a dialog-box prompt to change their password each time that they log on to the domain when their password is configured to expire in 14 or fewer days.

Interactive logon: Require Domain Controller authentication to unlock workstation

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

Unlocking a locked computer requires logon information. For domain accounts, the **Interactive logon: Require Domain Controller authentication to unlock workstation** policy setting determines whether it is necessary to contact a domain controller to unlock a computer. Enabling this policy setting requires a domain controller to authenticate the domain account that is being used to unlock the computer. Disabling this policy setting allows a user to unlock the computer without the computer verifying the logon information with a domain controller. However, if **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** is set to a value greater than zero, the user's cached credentials will be used to unlock the system.

The computer caches (locally in memory) the credentials of any users who have been authenticated. The computer uses these cached credentials to authenticate anyone who attempts to unlock the console.

When cached credentials are used, any changes that have recently been made to the account (such as user rights assignments, account lockout, or the account being disabled) are not considered or applied after this authentication process. This means not only that user rights are not updated, but more importantly that disabled accounts are still able to unlock the console of the system.

It is advisable to set **Interactive logon: Require Domain Controller authentication to unlock workstation** to Enabled and set **Interactive logon: Number of previous logons to cache (in**

case domain controller is not available) to 0. When the console of a computer is locked by a user or automatically by a screen saver time-out, the console can only be unlocked if the user is able to reauthenticate to the domain controller. If no domain controller is available, users cannot unlock their workstations.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set **Interactive logon: Require Domain Controller authentication to unlock workstation** to Enabled and set **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** to 0. When the console of a computer is locked by a user or automatically by a screen saver time-out, the console can only be unlocked if the user is able to reauthenticate to the domain controller. If no domain controller is available, users cannot unlock their workstations.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

By default, the computer caches locally in memory the credentials of any users who are authenticated. The computer uses these cached credentials to authenticate anyone who attempts to unlock the console. When cached credentials are used, any changes that have recently been made to the account—such as user rights assignments, account lockout, or the account being disabled—are not considered or applied after the account is authenticated. User privileges are not updated, and disabled accounts are still able to unlock the console of the computer.

Countermeasure

Configure the **Interactive logon: Require Domain Controller authentication to unlock workstation** setting to Enabled and configure the **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** setting to 0.

Potential impact

When the console on a computer is locked by a user or automatically by a screen-saver timeout, the console can be unlocked only if the user can reauthenticate to the domain controller. If no domain controller is available, users cannot unlock their workstations. If you configure the **Interactive logon: Number of previous logons to cache (in case domain controller is not available)** setting to 0, users whose domain controllers are unavailable (such as mobile or remote users) cannot log on.

Interactive logon: Require smart card

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The **Interactive logon: Require smart card** policy setting requires users to log on to a computer by using a smart card.

Requiring users to use long, complex passwords for authentication enhances network security, especially if the users must change their passwords regularly. This reduces the chance that a malicious user will be able to guess a user's password through a brute-force attack. Using smart cards rather than passwords for authentication dramatically increases security because, with today's technology, it is nearly impossible for a malicious user to impersonate another user. Smart cards that require personal identification numbers (PINs) provide two-factor authentication: the user who attempts to log on must possess the smart card and know its PIN. A malicious user who captures the authentication traffic between the user's computer and the domain controller will find it extremely difficult to decrypt the traffic: even if they do, the next time the user logs on to the network, a new session key will be generated for encrypting traffic between the user and the domain controller.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set **Interactive logon: Require smart card** to Enabled. All users will have to use smart cards to log on to the network. This means that the organization must have a reliable public key infrastructure (PKI) in place, and provide smart cards and smart card readers for all users.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

In Windows Server 2003, Microsoft implemented the PKI called Certificate Services. In later operating system versions, Certificate Services was renamed Active Directory® Certificate Services (AD CS), and it became a separate server role that is installed by using Server Manager.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

It can be difficult to make users choose strong passwords, and even strong passwords are vulnerable to brute-force attacks if an attacker has sufficient time and computing resources.

Countermeasure

For users with access to computers that contain sensitive data, issue smart cards to users and configure the **Interactive logon: Require smart card** setting to Enabled.

Potential impact

All users of a computer with this setting enabled must use smart cards to log on to the local computer. This means that the organization must have a reliable public key infrastructure (PKI) as well as smart cards and smart card readers for these users. These requirements are significant challenges because expertise and resources are required to plan for and deploy these technologies. Windows Server 2008 R2 and Windows Server 2008 include Active Directory Certificate Services (AD CS) to implement and manage certificates. When AD CS is used with Windows 7 or Windows Vista, features such as automatic user and computer enrollment and renewal become available.

Interactive logon: Smart card removal behavior

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines what happens when the smart card for a logged-on user is removed from the smart card reader.

If smart cards are used for authentication, the computer should automatically lock itself when the card is removed—that way, if users forget to manually lock their workstations when they are away from them, malicious users cannot gain access.

If you select **Force Logoff** in the property sheet for this policy setting, the user is automatically logged off when the smart card is removed. Users will have to reinsert their smart cards and reenter their PINs when they return to their workstations.

Note

Beginning with Windows Vista or Windows Server 2008, the Smart Card Removal Policy service must be started for this setting to work.

Possible values

- No Action
- Lock Workstation

  If you select this, the workstation is locked when the smart card is removed, allowing users to leave the area, take their smart card with them, and still maintain a protected session.
- Force Logoff

  If you select this, the user is automatically logged off when the smart card is removed.
- Disconnect if a remote Remote Desktop Services session

  If you select this, removal of the smart card disconnects the session without logging the user off. This allows the user to insert the smart card and resume the session later, or at another smart card reader-equipped computer, without having to log on again. If the session is local, this policy functions identically to Lock Workstation.
- Not Defined

Best practices

1. Set **Interactive logon: Smart card removal behavior** to **Lock Workstation**. If you select **Lock Workstation** in the property sheet for this policy setting, the workstation is locked when the smart card is removed. This allows users to leave the area, take their smart card with them, and still maintain a protected session.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	No Action
DC Effective Default Settings	No Action
Member Server Effective Default Settings	No Action
Client Computer Effective Default Settings	No Action

Operating system version differences

In Windows Server 2008, the option **Disconnect if a remote Terminal Services session** was added, and in Windows Server 2008 R2, it was changed to **Disconnect if a remote Remote Desktop Services session** to accommodate the renaming of that feature.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users sometimes forget to lock their workstations when they are away from them, allowing the possibility for malicious users to access their computers. If smart cards are used for authentication, the computer should automatically lock itself when the card is removed to ensure that only the user with the smart card is accessing resources by using those credentials.

Countermeasure

Configure the **Interactive logon: Smart card removal behavior** setting to **Lock Workstation**.

If you select **Lock Workstation** for this policy setting, the workstation locks when the smart card is removed. Users can leave the area, take their smart card with them, and still maintain a protected session. This behavior is similar to the setting that requires users to log on when resuming work on the computer after the screen saver has started.

If you select **Force Logoff** for this policy setting, the user is automatically logged off when the smart card is removed. This setting is useful when a computer is deployed as a public access point, such as a kiosk or other type of shared computer.

Potential impact

If you select **Force Logoff**, users must insert their smart cards and enter their PINs when they return to their workstations.

Microsoft network client: Digitally sign communications (always)

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The Server Message Block (SMB) protocol provides the basis for file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with the Server service is permitted.

Implementation of digital signatures in high-security networks helps prevent the impersonation of client computers and servers, which is known as "session hijacking." But misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

For this policy to take effect on computers running Windows 2000, client-side packet signing must also be enabled. To enable client-side SMB packet signing, set **Microsoft network client: Digitally sign communications (if server agrees)**. Computers that have this policy set will not be able to communicate with computers that do not have server-side packet signing enabled. By default, server-side packet signing is enabled only on domain controllers running Windows 2000, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2. Server-side packet signing can be enabled on these computers by setting **Microsoft network server: Digitally sign communications (if client agrees)**.

If server-side SMB signing is required, a client computer will not be able to establish a session with that server, unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers. Similarly, if client-side SMB signing is required, that client computer will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

If server-side SMB signing is enabled, SMB packet signing will be negotiated with client computers that have SMB signing enabled.

Using SMB packet signing can impose up to a 15 percent performance degradation on file service transactions.

There are three other policy settings that relate to packet-signing requirements for Server Message Block (SMB) communications:

- Microsoft network server: Digitally sign communications (always)
- Microsoft network client: Digitally sign communications (if server agrees)
- Microsoft network server: Digitally sign communications (if client agrees)

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Configure the following security policy settings as follows:
    - Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
    - Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
    - Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
    - Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.

2. Alternately, you can set all of these policy settings to Enabled, but enabling them can cause slower performance on client computers and prevent them from communicating with legacy SMB applications and operating systems.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy setting between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Session hijacking uses tools that allow attackers who have access to the same network as the client computer or server to interrupt, end, or steal a session in progress. Attackers can potentially intercept and modify unsigned Server Message Block (SMB) packets and then modify the traffic and forward it so that the server might perform objectionable actions. Alternatively, the attacker could pose as the server or client computer after legitimate authentication, and gain unauthorized access to data.

SMB is the resource-sharing protocol that is supported by many Windows operating systems. It is the basis of NetBIOS and many other protocols. SMB signatures authenticate users and the servers that host the data. If either side fails the authentication process, data transmission does not take place.

Countermeasure

Configure the settings as follows:

- Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
- Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
- Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
- Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.

In highly secure environments, we recommend that you configure all of these settings to Enabled. However, that configuration may cause slower performance on client computers and prevent communications with earlier SMB applications and operating systems.

Note

An alternative countermeasure that could protect all network traffic is to implement digital signatures with IPsec. There are hardware-based accelerators for IPsec encryption and signing that could be used to minimize the performance impact on the servers' CPUs. No such accelerators are available for SMB signing.

Potential impact

In the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server, Windows 7, Windows Vista, Windows XP Professional, and Windows 2000 Professional operating systems, implementations of the SMB file and print-sharing protocol support mutual authentication. This prevents session hijacking attacks and supports message authentication to prevent man-in-the-middle attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by the client and the server.

Implementation of SMB signing may negatively affect performance because each packet must be signed and verified. If these settings are enabled on a server that is performing multiple roles, such as a small business server that is serving as a domain controller, file server, print server, and application server, performance may be substantially slowed. Additionally, if you configure computers to ignore all unsigned SMB communications, older applications and operating systems

cannot connect. However, if you completely disable all SMB signing, computers are vulnerable to session-hijacking attacks.

Microsoft network client: Digitally sign communications (if server agrees)

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

The Server Message Block (SMB) protocol provides the basis for Microsoft file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with the Server service is permitted.

Implementation of digital signatures in high-security networks helps to prevent the impersonation of client computers and servers, which is known as "session hijacking." But misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

For this policy to take effect on computers running Windows 2000, client-side packet signing must also be enabled. To enable client-side SMB packet signing, set **Microsoft network client: Digitally sign communications (if server agrees)**. Computers that have this policy set will not be able to communicate with computers that do not have server-side packet signing enabled. By default, server-side packet signing is enabled only on domain controllers running Windows 2000, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2. Server-side packet signing can be enabled on these computers by setting **Microsoft network server: Digitally sign communications (if client agrees)**.

If server-side SMB signing is required, a client computer will not be able to establish a session with that server, unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers. Similarly, if client-side SMB signing is required, that client computer will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

If server-side SMB signing is enabled, SMB packet signing will be negotiated with client computers that have SMB signing enabled.

Using SMB packet signing can impose up to a 15 percent performance degradation on file service transactions.

There are three other policy settings that relate to packet-signing requirements for Server Message Block (SMB) communications:

- Microsoft network server: Digitally sign communications (always)
- Microsoft network client: Digitally sign communications (always)
- Microsoft network server: Digitally sign communications (if client agrees)

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Configure the following security policy settings as follows:
   - Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
   - Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
   - Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
   - Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.
2. Alternately, you can set all of these policy settings to Enabled, but enabling them can cause slower performance on client computers and prevent them from communicating with legacy SMB applications and operating systems.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy setting between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Session hijacking uses tools that allow attackers who have access to the same network as the client or server to interrupt, end, or steal a session in progress. Attackers can potentially intercept and modify unsigned Server Message Block (SMB) packets and then modify the traffic and forward it so that the server might perform objectionable actions. Alternatively, the attacker could pose as the server or client computer after legitimate authentication and gain unauthorized access to data.

SMB is the resource-sharing protocol that is supported by many Windows operating systems. It is the basis of NetBIOS and many other protocols. SMB signatures authenticate users and the servers that host the data. If either side fails the authentication process, data transmission does not take place.

Countermeasure

Configure the settings as follows:

- Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
- Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
- Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
- Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.

In highly secure environments we recommend that you configure all of these settings to Enabled. However, that configuration may cause slower performance on client computers and prevent communications with earlier SMB applications and operating systems.

Note

An alternative countermeasure that could protect all network traffic is to implement digital signatures with IPsec. There are hardware-based accelerators for IPsec encryption and signing that could be used to minimize the performance impact on the servers' CPUs. No such accelerators are available for SMB signing.

Potential impact

In the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server, Windows 7, Windows Vista, Windows XP Professional, and Windows 2000 Professional operating systems, implementations of the SMB file and print-sharing protocol support mutual authentication. This prevents session hijacking attacks and supports message authentication to prevent man-in-the-middle attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by the client and the server.

Implementation of SMB signing may negatively affect performance because each packet must be signed and verified. If these settings are enabled on a server that is performing multiple roles, such as a small business server that is serving as a domain controller, file server, print server, and application server, performance may be substantially slowed. Additionally, if you configure computers to ignore all unsigned SMB communications, older applications and operating systems cannot connect. However, if you completely disable all SMB signing, computers are vulnerable to session-hijacking attacks.

Microsoft network client: Send unencrypted password to third-party SMB servers

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The Server Message Block (SMB) protocol provides the basis for file and print sharing and many other networking operations, such as remote Windows administration. This policy setting allows or prevents the SMB redirector to send plaintext passwords to a non-Microsoft server service that does not support password encryption during authentication.

Possible values

- Enabled

  The Server Message Block (SMB) redirector is allowed to send plaintext passwords to a non-Microsoft server service that does not support password encryption during authentication.
- Disabled

The Server Message Block (SMB) redirector only sends encrypted passwords to non-Microsoft SMB server services. If those server services do not support password encryption, the authentication request will fail.

- Not defined

Best practices

1. It is advisable to set **Microsoft network client: Send unencrypted password to connect to third-party SMB servers** to Disabled. Some old applications and operating systems (such as MS-DOS, Windows for Workgroups 3.11, and Windows 95) might not be able to communicate with the servers in your organization through the SMB protocol.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you enable this policy setting, the server can transmit plaintext passwords across the network to other computers that offer SMB services. These other computers might not use any of the SMB security mechanisms that are included with Windows Server 2003 or later.

Countermeasure

Disable the **Microsoft network client: Send unencrypted password to connect to third-party SMB servers** setting.

Potential impact

Some old applications and operating systems (such as MS-DOS, Windows for Workgroups 3.11, and Microsoft Windows 95) may not be able to communicate with the servers in your organization by means of the SMB protocol.

Microsoft network server: Amount of idle time required before suspending session

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

Each Server Message Block (SMB) session consumes server resources. Establishing numerous null sessions will cause the server to slow down or possibly fail. A malicious user might repeatedly establish SMB sessions until the server stops responding; at this point, SMB services will become slow or unresponsive.

The **Microsoft network server: Amount of idle time required before suspending session** policy setting determines the amount of continuous idle time that must pass in an SMB session before the session is suspended due to inactivity. You can use this policy setting to control when

a computer suspends an inactive SMB session. The session is automatically reestablished when client computer activity resumes.

Possible values

- A user-defined number of minutes from 0 through 99,999

  For this policy setting, a value of 0 means to disconnect an idle session as quickly as is reasonably possible. The maximum value is 99999, which is 208 days. In effect, this value disables the policy.
- Not defined

Best practices

1. It is advisable to set this policy to 15 minutes. There will be little impact because SMB sessions will be reestablished automatically if the client resumes activity.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	15 minutes
DC Effective Default Settings	15 minutes
Member Server Effective Default Settings	15 minutes
Client Computer Effective Default Settings	15 minutes

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Each SMB session consumes server resources, and numerous null sessions slow the server or possibly cause it to fail. An attacker could repeatedly establish SMB sessions until the server's SMB services become slow or unresponsive.

Countermeasure

The default behavior on a server mitigates this threat by design in Windows Server 2003 and later.

Potential impact

There is little impact because SMB sessions are reestablished automatically if the client computer resumes activity.

Microsoft network server: Attempt S4U2Self to obtain claim information

This security policy reference topic for the IT professional describes the best practices, location, values, management, and security considerations for this policy setting.

Reference

This security setting supports client computers running a version of Windows prior to Windows 8 that are trying to access a file share that requires user claims. This setting determines whether the local file server will attempt to use Kerberos Service-for-User-to-Self (S4U2Self) functionality to obtain a network client principal’s claims from the client’s account domain. This setting should

only be enabled if the file server is using user claims to control access to files, and if the file server will support client principals whose accounts might be in a domain that has client computers and domain controllers running a version of Windows prior to Windows 8 or Windows Server 2012.

When enabled, this security setting causes the Windows file server to examine the access token of an authenticated network client principal and determines if claim information is present. If claims are not present, the file server will then use the Kerberos S4U2Self feature to attempt to contact a Windows Server 2012 domain controller in the client's account domain and obtain a claims-enabled access token for the client principal. A claims-enabled token might be needed to access files or folders that have claim-based access control policy applied.

If this setting is disabled, the Windows file server will not attempt to obtain a claim-enabled access token for the client principal.

Possible values

- **Default**

  The Windows file server will examine the access token of an authenticated network client principal and determine if claim information is present.
- **Enabled**

  Same as **Default**.
- **Disabled**
- **Not defined**

  Same as **Disabled**.

Best practices

This setting should be set to **Default** so that the file server can automatically evaluate whether claims are needed for the user. You should explicitly configure this setting to **Enabled** only if there are local file access policies that include user claims.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

This policy setting was introduced in Windows Server 2012 and Windows 8 to support client computers running a version of Windows prior to Windows 8 that are trying to access a file share that requires user claims.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

This setting should only be enabled if the file server is using user claims to control access to files, and if the file server will support client principals whose accounts might be in a domain that has client computers and domain controllers running a version of Windows prior to Windows 8 or Windows Server 2012.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

None. Enabling this policy setting allows you take advantage of features in Windows Server 2012 and Windows 8 for specific scenarios to use claims-enabled tokens to access files or folders that have claim-based access control policy applied on Windows operating systems prior to Windows Server 2012 and Windows 8. For more information, see Reference.

Countermeasure

Not applicable.

Potential impact

None.

See Also

Security Options

Microsoft network server: Digitally sign communications (always)

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The Server Message Block (SMB) protocol provides the basis for file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with the Server service is permitted.

Implementation of digital signatures in high-security networks helps to prevent the impersonation of client computers and servers, which is known as "session hijacking." But misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

For this policy to take effect on computers running Windows 2000, client-side packet signing must also be enabled. To enable client-side SMB packet signing, set **Microsoft network client: Digitally sign communications (if server agrees)**. Computers that have this policy set will not be able to communicate with computers that do not have server-side packet signing enabled. By default, server-side packet signing is enabled only on domain controllers running Windows 2000, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2. Server-side packet signing can be enabled on computers running Windows 2000 and later by setting **Microsoft network server: Digitally sign communications (if client agrees)**.

If server-side SMB signing is required, a client computer will not be able to establish a session with that server, unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers. Similarly, if client-side SMB signing is required, that client computer will not be able to establish a session with servers that do

not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

If server-side SMB signing is enabled, SMB packet signing will be negotiated with client computers that have SMB signing enabled.

Using SMB packet signing can impose up to a 15 percent performance degradation on file service transactions.

There are three other policy settings that relate to packet-signing requirements for Server Message Block (SMB) communications:

- Microsoft network client: Digitally sign communications (always)
- Microsoft network client: Digitally sign communications (if server agrees)
- Microsoft network server: Digitally sign communications (if client agrees)

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Configure the following security policy settings as follows:
   - Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
   - Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
   - Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
   - Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.
2. Alternately, you can set all of these policy settings to Enabled, but enabling them can cause slower performance on client computers and prevent them from communicating with legacy SMB applications and operating systems.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Enabled
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy setting between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Session hijacking uses tools that allow attackers who have access to the same network as the client computer or server to interrupt, end, or steal a session in progress. Attackers can potentially intercept and modify unsigned Server Message Block (SMB) packets and then modify the traffic and forward it so that the server might perform objectionable actions. Alternatively, the attacker could pose as the server or client computer after legitimate authentication and gain unauthorized access to data.

SMB is the resource-sharing protocol that is supported by many Windows operating systems. It is the basis of NetBIOS and many other protocols. SMB signatures authenticate users and the servers that host the data. If either side fails the authentication process, data transmission does not take place.

Countermeasure

Configure the settings as follows:

- Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
- Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
- Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
- Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.

In highly secure environments we recommend that you configure all of these settings to Enabled. However, that configuration may cause slower performance on client computers and prevent communications with earlier SMB applications and operating systems.

Note

An alternative countermeasure that could protect all network traffic is to implement digital signatures with IPsec. There are hardware-based accelerators for IPsec encryption and signing that could be used to minimize the performance impact on the servers' CPUs. No such accelerators are available for SMB signing.

Potential impact

In the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server, Windows 7, Windows Vista, Windows XP Professional, and Windows 2000 Professional operating systems, implementations of the SMB file and print-sharing protocol support mutual authentication. This prevents session hijacking attacks and supports message authentication to prevent man-in-the-middle attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by the client and the server.

Implementation of SMB signing may negatively affect performance because each packet must be signed and verified. If these settings are enabled on a server that is performing multiple roles, such as a small business server that is serving as a domain controller, file server, print server, and application server, performance may be substantially slowed. Additionally, if you configure computers to ignore all unsigned SMB communications, older applications and operating systems cannot connect. However, if you completely disable all SMB signing, computers are vulnerable to session-hijacking attacks.

Microsoft network server: Digitally sign communications (if client agrees)

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The Server Message Block (SMB) protocol provides the basis for file and print sharing and many other networking operations, such as remote Windows administration. To prevent man-in-the-middle attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. This policy setting determines whether SMB packet signing must be negotiated before further communication with the Server service is permitted.

Implementation of digital signatures in high-security networks helps to prevent the impersonation of client computers and servers, which is known as "session hijacking." But misuse of these policy settings is a common error that can cause data loss or problems with data access or security.

For this policy to take effect on computers running Windows 2000, client-side packet signing must also be enabled. To enable client-side SMB packet signing, set **Microsoft network client: Digitally sign communications (if server agrees)**. Computers that have this policy set will not be able to communicate with computers that do not have server-side packet signing enabled. By default, server-side packet signing is enabled only on domain controllers running Windows 2000, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2. Server-side packet signing can be enabled on these computers by setting**Microsoft network server: Digitally sign communications (if client agrees)**.

If server-side SMB signing is required, a client computer will not be able to establish a session with that server, unless it has client-side SMB signing enabled. By default, client-side SMB signing is enabled on workstations, servers, and domain controllers. Similarly, if client-side SMB signing is required, that client computer will not be able to establish a session with servers that do not have packet signing enabled. By default, server-side SMB signing is enabled only on domain controllers.

If server-side SMB signing is enabled, SMB packet signing will be negotiated with client computers that have SMB signing enabled.

Using SMB packet signing can impose up to a 15 percent performance degradation on file service transactions.

There are three other policy settings that relate to packet-signing requirements for Server Message Block (SMB) communications:

- Microsoft network server: Digitally sign communications (always)
- Microsoft network client: Digitally sign communications (if server agrees)
- Microsoft network client: Digitally sign communications (always)

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Configure the following security policy settings as follows:
   - Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
   - Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
   - Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
   - Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.
2. Alternately, you can set all of these policy settings to Enabled, but enabling them can cause slower performance on client computers and prevent them from communicating with legacy SMB applications and operating systems.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Enabled
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy setting between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Session hijacking uses tools that allow attackers who have access to the same network as the client computer or server to interrupt, end, or steal a session in progress. Attackers can potentially intercept and modify unsigned Server Message Block (SMB) packets and then modify the traffic and forward it so that the server might perform objectionable actions. Alternatively, the attacker could pose as the server or client computer after legitimate authentication and gain unauthorized access to data.

SMB is the resource-sharing protocol that is supported by many Windows operating systems. It is the basis of NetBIOS and many other protocols. SMB signatures authenticate users and the servers that host the data. If either side fails the authentication process, data transmission does not take place.

Countermeasure

Configure the settings as follows:

- Disable **Microsoft Network Client: Digitally Sign Communications (Always)**.
- Disable **Microsoft Network Server: Digitally Sign Communications (Always)**.
- Enable **Microsoft Network Client: Digitally Sign Communications (If Server Agrees)**.
- Enable **Microsoft Network Server: Digitally Sign Communications (If Client Agrees)**.

In highly secure environments we recommend that you configure all of these settings to Enabled. However, that configuration may cause slower performance on client computers and prevent communications with earlier SMB applications and operating systems.

Note

An alternative countermeasure that could protect all network traffic is to implement digital signatures with IPsec. There are hardware-based accelerators for IPsec encryption and signing that could be used to minimize the performance impact on the servers' CPUs. No such accelerators are available for SMB signing.

Potential impact

In the Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Windows 2000 Server, Windows 7, Windows Vista, Windows XP Professional, and Windows 2000 Professional operating systems, implementations of the SMB file and print-sharing protocol support mutual authentication. This prevents session hijacking attacks and supports message authentication to prevent man-in-the-middle attacks. SMB signing provides this authentication by placing a digital signature into each SMB, which is then verified by the client and the server.

Implementation of SMB signing may negatively affect performance because each packet must be signed and verified. If these settings are enabled on a server that is performing multiple roles, such as a small business server that is serving as a domain controller, file server, print server, and application server, performance may be substantially slowed. Additionally, if you configure computers to ignore all unsigned SMB communications, older applications and operating systems cannot connect. However, if you completely disable all SMB signing, computers are vulnerable to session-hijacking attacks.

Microsoft network server: Disconnect clients when logon hours expire

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting enables or disables the forced disconnection of users who are connected to the local computer outside their user account's valid logon hours. It affects the SMB component. If you enable this policy setting, client computer sessions with the SMB service are forcibly disconnected when the client's logon hours expire. If you disable this policy setting, established client computer sessions are maintained after the client computer's logon hours expire.

Possible values

- Enabled

  Client computer sessions with the SMB service are forcibly disconnected when the client computer's logon hours expire. If logon hours are not used in your organization, enabling this policy setting will have no impact.

- Disabled

  The system maintains an established client computer session after the client computer's logon hours have expired.

- Not defined

Best practices

1. If you enable this policy setting, you should also enable **Network security: Force logoff when logon hours expire**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

The default setting of this policy in operating systems previous to Windows Server 2008 and Windows Vista is Disabled for stand-alone servers and client computers.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

When using Group Policy to distribute this policy setting, the default is different for Windows operating systems prior to Windows Server 2008.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If your organization configures logon hours for users, it makes sense to enable this policy setting. Otherwise, users who should not have access to network resources outside of their logon hours can continue to use those resources with sessions that were established during allowed hours.

Countermeasure

Enable the **Microsoft network server: Disconnect clients when logon hours expire** setting.

Potential impact

If logon hours are not used in your organization, this policy setting has no impact. If logon hours are used, existing user sessions are forcibly terminated when their logon hours expire.

Microsoft network server: Server SPN target name validation level

This security policy reference topic for the IT professional describes the best practices, location, and values, policy management and security considerations for this policy setting.

Reference

This policy setting controls the level of validation that a server with shared folders or printers performs on the service principal name (SPN) that is provided by the client computer when the client computer establishes a session by using the Server Message Block (SMB) protocol. The level of validation can help prevent a class of attacks against SMB services (referred to as SMB relay attacks). This setting will affect both SMB1 and SMB2.

Servers that use SMB provide availability to their file systems and other resources, such as printers, to networked client computers. Most servers that use SMB validate user access to resources by using NT Domain authentication (NTLMv1 and NTLMv2) and the Kerberos protocol.

Windows Server 2008 and Windows Vista support SMB protocol (SMB 2.0 or SMB2). Windows Server 2008 R2 and Windows 7 supports SMB 2.0, SMB2 and SMB 2.1.

Possible values

The options for validation levels are:

- **Off**

  The SPN from a SMB client is not required or validated by the SMB server.
- **Accept if provided by client**

  The SMB server will accept and validate the SPN provided by the SMB client and allow a session to be established if it matches the SMB server's list of SPN's. If the SPN does not match, the session request for that SMB client will be denied.
- **Required from client**

  The SMB client must send a SPN name in session setup, and the SPN name provided must match the SMB server that is being requested to establish a connection. If no SPN is provided by the client computer, or the SPN provided does not match, the session is denied.

The default setting is Off.

Best practices

This setting affects the server SMB behavior, and its implementation should be carefully evaluated and tested to prevent disruptions to file and print serving capabilities.

Note
All Windows operating systems support a client-side SMB component and a server-side SMB component.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or Group Policy object (GPO)	Default value
Default domain policy	Off
Default domain controller policy	Off
Stand-alone server default settings	Off
Domain controller effective default settings	Validation level check not implemented
Member server effective default settings	Validation level check not implemented
Effective GPO default settings on client computers	Validation level check not implemented

Operating system version differences

This policy was introduced in Windows Server 2008 and Windows Vista, and it can be applied through Group Policy to all Windows server operating systems beginning with Windows Server 2003 and Windows XP.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

None.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This policy setting controls the level of validation that a server with shared folders or printers performs on the service principal name (SPN) that is provided by the client computer when the client computer establishes a session by using the SMB protocol. The level of validation can help prevent a class of attacks against SMB servers (referred to as SMB relay attacks). This setting will affect both SMB1 and SMB2.

Countermeasure

For countermeasures that are appropriate to your environment, see **Possible values** above.

Potential impact

All Windows operating systems support a client-side SMB component and a server-side SMB component. This setting affects the server SMB behavior, and its implementation should be carefully evaluated and tested to prevent disruptions to file and print serving capabilities.

Because the SMB protocol is widely deployed, setting the options to **Accept if provided by client** or **Required from client** will prevent some clients from successfully authenticating to some servers in your environment.

Network access: Allow anonymous SID/Name translation

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enables or disables the ability of an anonymous user to request security identifier (SID) attributes for another user.

If this policy setting is enabled, a user might use the well-known Administrators SID to get the real name of the built-in Administrator account, even if the account has been renamed. That person might then use the account name to initiate a brute-force password-guessing attack.

Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

Possible values

- Enabled

  An anonymous user can request the SID attribute for another user. An anonymous user with knowledge of an administrator's SID could contact a computer that has this policy enabled and use the SID to get the administrator's name. This setting affects the SID-to-name translation as well as the name-to-SID translation
- Disabled

  Prevents an anonymous user from requesting the SID attribute for another user.
- Not defined

Best practices

- Set this policy to Disabled. This is the default value on member computers; therefore, it will have no impact on them. The default value for domain controllers is Enabled. Disabling this policy setting on domain controllers means that legacy systems might be unable to communicate with Windows Server 2003-based domains.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Note defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

The default value of this setting has changed between operating systems as follows:

- The default on domain controllers running Windows Server 2003 R2 or earlier was set to Enabled.
- The default on domain controllers running Windows Server 2008 or Windows Server 2008 R2 is set to Disabled.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Modifying this setting may affect compatibility with client computers, services, and applications.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If this policy setting is enabled, a user with local access could use the well-known Administrator's SID to learn the real name of the built-in Administrator account, even if it has been renamed. That person could then use the account name to initiate a password-guessing attack.

Countermeasure

Disable the **Network access: Allow anonymous SID/Name translation** setting.

Potential impact

Disabled is the default configuration for this policy setting on member computers; therefore, it has no impact on them. The default configuration for domain controllers is Enabled. If you disable this policy setting on domain controllers, computers running versions of Windows earlier than Windows Server 2003 may not communicate with Windows Server 2003–based domains. For example, computers with the following configurations may not work:

- Windows NT 4.0–based Remote Access Service servers
- Servers that host Microsoft SQL Server® and run on Windows NT 3.x–based or Windows NT 4.0–based computers
- Servers that host Remote Access Service or Microsoft SQL Server that run on Windows 2000–based computers and are located in Windows NT domains

Network access: Do not allow anonymous enumeration of SAM accounts

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines which additional permissions will be assigned for anonymous connections to the computer. Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to give access to users in a trusted domain that does not maintain a reciprocal trust.

This policy setting has no impact on domain controllers.

Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

Possible values

- Enabled

  It will be impossible to establish trusts with Windows NT 4.0–based domains. This value will also cause problems with earlier-version client computers (such as those running Windows NT 3.51 and Windows 95) that are trying to use resources on the server.
- Disabled

  No additional permissions can be assigned by the administrator for anonymous connections to the computer. Anonymous connections will rely on default permissions.
- Not defined

Best practices

1. Best practices are dependent on your security goals and policies.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

Windows 2000 Server has a similar policy setting named **Additional Restrictions for Anonymous Connections** managed a registry entry **RestrictAnonymous**, located in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA key. In Windows Server 2003, the policy settings **Network access: Do not allow anonymous enumeration of SAM accounts** and **Network access: Do not allow anonymous enumeration of SAM accounts and shares** replace the Windows 2000 policy setting. They manage the registry entries **RestrictAnonymousSAM** and **RestrictAnonymous**, respectively, both located in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\ key.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflicts

Even with this policy setting enabled, anonymous users will have access to resources with permissions that explicitly include the built-in group, ANONYMOUS LOGON (on systems earlier than Windows Server 2008 and Windows Vista).

Group Policy

This policy has no impact on domain controllers.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An unauthorized user could anonymously list account names and use the information to perform social engineering attacks or attempt to guess passwords. Social engineering attackers try to deceive users in some way to obtain passwords or some form of security information.

Countermeasure

Enable the **Network access: Do not allow anonymous enumeration of SAM accounts** setting.

Potential impact

It is impossible to establish trusts with Windows NT 4.0–based domains. Also, client computers that run earlier versions of the Windows operating system, such as Windows NT 3.51 and Windows 95, experience problems when they try to use resources on the server.

Network access: Do not allow anonymous enumeration of SAM accounts and shares

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines which additional permissions will be assigned for anonymous connections to the computer. Windows allows anonymous users to perform certain activities, such as enumerating the names of domain accounts and network shares. This is convenient, for example, when an administrator wants to give access to users in a trusted domain that does not maintain a reciprocal trust. However, even with this policy setting enabled, anonymous users will have access to resources with permissions that explicitly include the built-in group, ANONYMOUS LOGON.

This policy setting has no impact on domain controllers.

Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

Possible values

- Enabled

  It will be impossible to establish trusts with Windows NT 4.0–based domains. This value will also cause problems with earlier-version client computers (such as those running Windows NT 3.51 and Windows 95) that are trying to use resources on the server.
- Disabled

  No additional permissions can be assigned by the administrator for anonymous connections to the computer. Anonymous connections will rely on default permissions. However, an unauthorized user could anonymously list account names and use the information to attempt to guess passwords or perform social-engineering attacks.
- Not defined

Best practices

1. Set this policy to Enabled. Note that when this is enabled, it will be impossible to establish trusts with Windows NT 4.0–based domains. This policy setting will also cause problems with

earlier-version clients, such as Windows NT 3.51 and Windows 95, that are trying to use resources on the server.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

Windows 2000 Server has a similar policy setting named **Additional Restrictions for Anonymous Connections** managed a registry entry **RestrictAnonymous**, located in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA key. In Windows Server 2003, the policy settings **Network access: Do not allow anonymous enumeration of SAM accounts** and **Network access: Do not allow anonymous enumeration of SAM accounts and shares** replace the Windows 2000 policy setting. They manage the registry entries **RestrictAnonymousSAM** and **RestrictAnonymous**, respectively, both located in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\ key.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflicts

Even with this policy setting enabled, anonymous users will have access to resources with permissions that explicitly include the built-in group, ANONYMOUS LOGON (on systems earlier than Windows Server 2008 and Windows Vista).

Group Policy

This policy has no impact on domain controllers.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An unauthorized user could anonymously list account names and shared resources and use the information to attempt to guess passwords or perform social-engineering attacks.

Countermeasure

Enable the **Network access: Do not allow anonymous enumeration of SAM accounts and shares** setting.

Potential impact

It is impossible to grant access to users of another domain across a one-way trust because administrators in the trusting domain are unable to enumerate lists of accounts in the other domain. Users who access file and print servers anonymously are unable to list the shared network resources on those servers; the users must be authenticated before they can view the lists of shared folders and printers.

Network access: Do not allow storage of passwords and credentials for network authentication

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This security setting determines whether Credential Manager saves passwords and credentials for later use when it gains domain authentication.

Possible values

- Enabled

  Credential Manager does not store passwords and credentials on the computer.
- Disabled

  Credential Manager will store passwords and credentials on this computer for later use for domain authentication.
- Not defined

Best practices

It is a recommended practice to disable the ability of the Windows operating system to cache credentials on any computer where credentials are not needed. Evaluate your servers and workstations to determine the requirements. Cached credentials are designed primarily to be used on laptops that require domain credentials when disconnected from the domain.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Disabled
Default domain controller policy	Disabled
Stand-alone server default settings	Disabled
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Effective GPO default settings on client computers	Not defined

Operating system version differences

This policy is present in Windows Server 2003 and Windows XP, and it is named **Network access: Do not allow storage of credentials or .NET Passports for network authentication**. The policy name was modified for Windows Server 2008 and Windows Vista. However, this policy can be applied to all Windows server operating systems through Group Policy.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

A restart of the computer is required before this policy will be effective when changes to this policy are saved locally or distributed through Group Policy.

Group Policy

Although the name of this policy was changed in Windows Server 2008 and Windows Vista, it can be applied to Windows Server 2003 and Windows XP.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Passwords that are cached can be accessed by the user when logged on to the computer. Although this information may sound obvious, a problem can arise if the user unknowingly runs malicious software that reads the passwords and forwards them to another, unauthorized user.

Note

The chances of success for this exploit and others that involve malicious software are reduced significantly for organizations that effectively implement and manage an enterprise antivirus solution combined with sensible software restriction policies. For more information, see **Software Restriction Policies**.

Regardless of what encryption algorithm is used to encrypt the password verifier, a password verifier can be overwritten so that an attacker can authenticate as the user to whom the verifier belongs. Therefore, the administrator's password may be overwritten. This procedure requires physical access to the computer. Utilities exist that can help overwrite the cached verifier. By using one of these utilities, an attacker can authenticate by using the overwritten value.

Overwriting the administrator's password does not help the attacker access data that is encrypted by using that password. Also, overwriting the password does not help the attacker access any Encrypting File System (EFS) data that belongs to other users on that computer. Overwriting the password does not help an attacker replace the verifier, because the base keying material is incorrect. Therefore, data that is encrypted by using Encrypting File System or by using the Data Protection API (DPAPI) will not decrypt.

Countermeasure

Enable the **Network access: Do not allow storage of passwords and credentials for network authentication** setting.

To limit the number of changed domain credentials that are stored on the computer, set the **cachedlogonscount** registry entry. By default, the operating system caches the verifier for each unique user's ten most recent valid logons. This value can be set to any value between 0 and 50. By default, all versions of the Windows operating system remember 10 cached logons, except Windows Server 2008 R2 and Windows Server 2008, which are set at 25.

When you try to log on to a domain from a Windows-based client computer, and a domain controller is unavailable, you do not receive an error message. Therefore, you may not notice that you logged on with cached domain credentials. You can set a notification of logon that uses cached domain credentials with the ReportDC registry entry.

Potential impact

Users are forced to type passwords whenever they log on to their Windows Live ID or other network resources that are not accessible to their domain account. This policy setting should have no impact on users who access network resources that are configured to allow access with their Active Directory–based domain account.

Additional resources

For information about how the Windows operating system stores and manages credentials, see Cached and Stored Credentials Technical Overview.

Network access: Let Everyone permissions apply to anonymous users

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines what additional permissions are granted for anonymous connections to the computer. If you enable this policy setting, anonymous users can enumerate the names of domain accounts and shared folders and perform certain other activities. This capability is convenient, for example, when an administrator wants to grant access to users in a trusted domain that does not maintain a reciprocal trust.

By default, the token that is created for anonymous connections does not include the Everyone SID. Therefore, permissions that are assigned to the Everyone group do not apply to anonymous users.

Possible values

- Enabled

  The Everyone SID is added to the token that is created for anonymous connections, and anonymous users can access any resource for which the Everyone group has been assigned permissions.
- Disabled

  The Everyone SID is removed from the token that is created for anonymous connections.
- Not defined

Best practices

1. Set this policy to **Disabled**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Polices\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled

Server type or GPO	Default value
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An unauthorized user could anonymously list account names and shared resources and use the information to attempt to guess passwords, perform social engineering attacks, or launch DoS attacks.

Countermeasure

Disable the **Network access: Let Everyone permissions apply to anonymous users** setting.

Potential impact

None. This is the default configuration.

Network access: Named Pipes that can be accessed anonymously

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines which communication sessions, or pipes, have attributes and permissions that allow anonymous access.

Restricting access over named pipes such as COMNAP and LOCATOR helps prevent unauthorized access to the network. The table in the Vulnerability section lists default named pipes and their purpose.

Possible values

- User-defined list of shared folders
- Not defined

Best practices

1. Set this policy to a null value; that is, enable the policy setting, but do not enter named pipes in the text box. This will disable null session access over named pipes, and applications that rely on this feature or on unauthenticated access to named pipes will no longer function. For example, with Microsoft Commercial Internet System 1.0, the Internet Mail Service runs under Inetinfo.exe. Inetinfo.exe starts in the context of the System account. When Internet Mail Service needs to query the Microsoft SQL Server database, it uses the System account, which uses null credentials to access an SQL pipe on the computer running SQL Server.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Netlogon, samr, lsarpc

Server type or GPO	Default value
Stand-Alone Server Default Settings	Null
DC Effective Default Settings	Netlogon, samr, lsarpc
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

The default settings for this policy changed from Windows Server 2003.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

For this policy setting to take effect, you must also enable the **Network access: Restrict anonymous access to named pipes and shares** setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

You can restrict access over named pipes such as COMNAP and LOCATOR to help prevent unauthorized access to the network. The following list describes available named pipes and their purpose. These pipes were granted anonymous access in earlier versions of Windows and some legacy applications may still use them.

Named pipe	Purpose
COMNAP	SNABase named pipe. Systems network Architecture (SNA) is a collection of network protocols that were originally developed for IBM

Named pipe	Purpose
	mainframe computers.
COMNODE	SNA Server named pipe.
SQL\QUERY	Default named pipe for SQL Server.
SPOOLSS	Named pipe for the Print Spooler service.
EPMAPPER	End Point Mapper named pipe.
LOCATOR	Remote Procedure Call Locator service named pipe.
TrlWks	Distributed Link Tracking Client named pipe.
TrkSvr	Distributed Link Tracking Server named pipe.

Countermeasure

Configure the **Network access: Named Pipes that can be accessed anonymously** setting to a null value (enable the setting but do not specify named pipes in the text box).

Potential impact

This configuration disables null-session access over named pipes, and applications that rely on this feature or on unauthenticated access to named pipes no longer function. This may break trust between Windows Server 2003 domains in a mixed mode environment.

Network access: Remotely accessible registry paths

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines which registry paths are accessible when an application or process references the WinReg key to determine access permissions.

The registry is a database for computer configuration information, much of which is sensitive. A malicious user can use the registry to facilitate unauthorized activities. To reduce the risk of this happening, suitable access control lists (ACLs) are assigned throughout the registry to help protect it from access by unauthorized users.

To allow remote access, you must also enable the Remote Registry service.

Possible values

- User-defined list of paths
- Not Defined

Best practices

1. Set this policy to a null value; that is, enable the policy setting but do not enter any paths in the text box. Remote management tools, such as the Microsoft Baseline Security Analyzer and Microsoft Systems Management Server (SMS), require remote access to the registry. Removing the default registry paths from the list of accessible paths might cause these and other management tools to fail.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	See the following registry key combination
DC Effective Default Settings	See the following registry key combination
Member Server Effective Default Settings	See the following registry key combination
Client Computer Effective Default Settings	See the following registry key combination

The combination of all the following registry keys apply to the previous settings:

1. System\CurrentControlSet\Control\ProductOptions
2. System\CurrentControlSet\Control\Server Applications
3. Software\Microsoft\Windows NT\CurrentVersion

Operating system version differences

This security setting is not available on Windows operating systems prior to Windows Server 2008. The security setting that appears on computers running Windows XP is **Network access: Remotely accessible registry paths**, and the setting in Windows Server 2003 is

Network access: Remotely accessible registry paths and subpaths. For more information, see Network access: Remotely accessible registry paths and subpaths.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An attacker could use information in the registry to facilitate unauthorized activities. To reduce the risk of such an attack, suitable ACLs are assigned throughout the registry to help protect it from access by unauthorized users.

Countermeasure

Configure the **Network access: Remotely accessible registry paths** setting to a null value (enable the setting, but do not enter any paths in the text box).

Potential impact

Remote management tools such as the Microsoft Baseline Security Analyzer (MBSA) and Microsoft Systems Management Server (SMS) require remote access to the registry to properly monitor and manage those computers. If you remove the default registry paths from the list of accessible ones, such remote management tools could fail.

Note
If you want to allow remote access, you must also enable the Remote Registry service.

Network access: Remotely accessible registry paths and subpaths

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy setting determines which registry paths and subpaths are accessible when an application or process references the WinReg key to determine access permissions.

The registry is a database for computer configuration information, much of which is sensitive. A malicious user can use it to facilitate unauthorized activities. The chance of this happening is reduced by the fact that the default ACLs that are assigned throughout the registry are fairly restrictive, and they help protect it from access by unauthorized users.

To allow remote access, you must also enable the Remote Registry service.

Possible values

- User-defined list of paths
- Not Defined

Best practices

1. Set this policy to a null value; that is, enable the policy setting, but do not enter any paths in the text box. Remote management tools, such as the Microsoft Baseline Security Analyzer and Microsoft Systems Management Server (SMS), require remote access to the registry. Removing the default registry paths from the list of accessible paths might cause these and other management tools to fail.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined

Server type or GPO	Default value
Stand-Alone Server Default Settings	See the following registry key combination
DC Effective Default Settings	See the following registry key combination
Member Server Effective Default Settings	See the following registry key combination
Client Computer Effective Default Settings	See the following registry key combination

The combination of all the following registry keys apply to the previous settings:

1. System\CurrentControlSet\Control\Print\Printers
2. System\CurrentControlSet\Services\Eventlog
3. Software\Microsoft\OLAP Server
4. Software\Microsoft\Windows NT\CurrentVersion\Print
5. Software\Microsoft\Windows NT\CurrentVersion\Windows
6. System\CurrentControlSet\Control\ContentIndex
7. System\CurrentControlSet\Control\Terminal Server
8. System\CurrentControlSet\Control\Terminal Server\UserConfig
9. System\CurrentControlSet\Control\Terminal Server\DefaultUserConfiguration
10. Software\Microsoft\Windows NT\CurrentVersion\Perflib
11. System\CurrentControlSet\Services\SysmonLog

Operating system version differences

This security setting is not available on Windows operating systems prior to Windows Server 2008. The security setting that appears on computers running Windows XP is **Network access: Remotely accessible registry paths**, and the setting in Windows Server 2003 is **Network access: Remotely accessible registry paths and subpaths**.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The registry contains sensitive computer configuration information that could be used by an attacker to facilitate unauthorized activities. The fact that the default ACLs that are assigned throughout the registry are fairly restrictive and help to protect the registry from access by unauthorized users reduces the risk of such an attack.

Countermeasure

Configure the **Network access: Remotely accessible registry paths and sub-paths** setting to a null value (enable the setting but do not enter any paths in the text box).

Potential impact

Remote management tools such as MBSA and SMS require remote access to the registry to properly monitor and manage those computers. If you remove the default registry paths from the list of accessible ones, such remote management tools could fail.

Note

If you want to allow remote access, you must also enable the Remote Registry service.

Network access: Restrict anonymous access to Named Pipes and Shares

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enables or disables the restriction of anonymous access to only those shared folders and pipes that are named in the **Network access: Named pipes that can be accessed anonymously** and **Network access: Shares that can be accessed anonymously** settings. The setting controls null session access to shared folders on your computers by adding RestrictNullSessAccess with the value 1 in the registry key **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanManServer\Parameters**. This registry value toggles null session shared folders on or off to control whether the Server service restricts unauthenticated clients' access to named resources.

Null sessions are a weakness that can be exploited through the various shared folders on the computers in your environment.

Possible values

- Enabled

- Disabled
- Not defined

Best practices

1. Set this policy to Enabled. Enabling this policy setting restricts null session access to unauthenticated users to all server pipes and shared folders except those listed in the **NullSessionPipes** and **NullSessionShares** registry entries.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this security policy between Windows operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Null sessions are a weakness that can be exploited through shared folders (including the default shared folders) on computers in your environment.

Countermeasure

Enable the **Network access: Restrict anonymous access to Named Pipes and Shares** setting.

Potential impact

You can enable this policy setting to restrict null-session access for unauthenticated users to all server pipes and shared folders except those that are listed in the NullSessionPipes and NullSessionShares entries.

If you choose to enable this setting, and you are supporting Windows NT 4.0 domains, determine whether any of the named pipes in the following list are required to maintain trust relationships between the domains and then add the pipe to the **Network access: Named pipes that can be accessed anonymously** setting:

- COMNAP–SNA session access
- COMNODE–SNA session access
- SQL\QUERY–SQL instance access
- SPOOLSS–Spooler service
- LLSRPC–License Logging service
- Netlogon–Net Logon service
- Lsarpc–LSA access
- Samr–Remote access to SAM objects
- browser–Computer Browser service

In operating systems earlier than Windows Server 2003 with Service Pack 1 (SP1), these named pipes were allowed anonymous access by default. In later operating systems, these pipes must be explicitly added if needed.

Network access: Shares that can be accessed anonymously

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines which shared folders can be accessed by anonymous users.

Possible values

- User-defined list of shared folders
- Not Defined

Best practices

1. Set this policy to a null value. There should be little impact because this is the default value. All users will have to be authenticated before they can access shared resources on the server.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

In Windows Server 2003, the default setting for stand-alone servers, member servers, and client computers was COMCFG, DFS$. In later versions of the Windows operating system, the default setting for these server types of GPOs is not defined (null).

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any shared folders that are listed can be accessed by any network user, which could lead to the exposure or corruption of sensitive data.

Countermeasure

Configure the **Network access: Shares that can be accessed anonymously** setting to a null value.

Potential impact

There should be little impact because this is the default configuration. Only authenticated users have access to shared resources on the server.

Network access: Sharing and security model for local accounts

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines how network logons that use local accounts are authenticated. If you configure this policy setting to Classic, network logons that use local account credentials authenticate with those credentials. If you configure this policy setting to Guest only, network logons that use local accounts are automatically mapped to the Guest account. The Classic model provides precise control over access to resources, and it enables you to grant different types of access to different users for the same resource. Conversely, the Guest only model treats all users equally, and they all receive the same level of access to a given resource, which can be either Read Only or Modify.

Notes

This policy setting does not affect network logons that use domain accounts. Nor does this policy setting affect interactive logons that are performed remotely through services such as Telnet or Remote Desktop Services.

When the computer is not joined to a domain, this policy setting also tailors the **Sharing** and **Security** tabs in Windows Explorer to correspond to the sharing and security model that is being used.

The default configuration for a stand-alone computer running Windows Vista or Windows XP Professional is Guest only. The default configuration for domain-joined computers running the Windows Vista or Windows XP Professional operating system is Classic. Computers that run the Windows Server 2008 or Windows Server 2003 operating system also use the Classic security model.

When the value of this policy setting is **Guest only - local users authenticate as Guest**, any user who can access your computer over the network does so with Guest user rights. This means that they will probably be unable to write to shared folders. Although this does increase security, it makes it impossible for authorized users to access shared resources on those systems. When the value is **Classic - local users authenticate as themselves**, local accounts must be password-protected; otherwise, anyone can use those user accounts to access shared system resources.

Possible values

- Classic - Local users authenticate as themselves
- Guest only - Local users authenticate as Guest
- Not defined

Best practices

1. For network servers, set this policy to **Classic - local users authenticate as themselves**.
2. On end-user systems, set this policy to **Guest only - local users authenticate as Guest**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Classic (local users authenticate as themselves)
DC Effective Default Settings	Classic (local users authenticate as themselves)
Member Server Effective Default Settings	Classic (local users authenticate as themselves)
Client Computer Effective Default Settings	Classic (local users authenticate as themselves)

Operating system version differences

There are no differences in this policy setting between Windows operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

This policy will have no impact on computers running Windows 2000.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

With the Guest only model, any user who can authenticate to your computer over the network does so with Guest privileges, which probably means that they do not have Write access to shared resources on that computer. Although this restriction does increase security, it makes it more difficult for authorized users to access shared resources on those computers because ACLs on those resources must include access control entries (ACEs) for the Guest account. With the Classic model, local accounts should be password protected. Otherwise, if Guest access is enabled, anyone can use those user accounts to access shared system resources.

Countermeasure

For network servers, configure the **Network access: Sharing and security model for local accounts setting** to **Classic – local users authenticate as themselves**. On end-user computers, configure this policy setting to **Guest only – local users authenticate as guest**.

Potential impact

None. This is the default configuration.

Network security: Allow Local System to use computer identity for NTLM

This security policy reference topic for the IT professional describes the location, values, policy management, and security considerations for this policy setting.

Reference

When services connect to computers that are running versions of the Windows operating system earlier than Windows Vista or Windows Server 2008, services that run as Local System and use SPNEGO (Negotiate) that revert to NTLM will authenticate anonymously. In Windows Server 2008 R2 and Windows 7, if a service connects to a computer running Windows Server 2008 or Windows Vista, the system service uses the computer identity.

When a service connects with the computer identity, signing and encryption are supported to provide data protection. (When a service connects anonymously, a system-generated session key is created, which provides no protection, but it allows applications to sign and encrypt data

without errors. Anonymous authentication uses a NULL session, which is a session with a server in which no user authentication is performed; and therefore, anonymous access is allowed.)

You can configure the **Network security: Allow Local System to use computer identity for NTLM** security policy setting to disallow Local System services that use Negotiate to use the computer identity when reverting to NTLM authentication.

Possible values

Setting	Windows Server 2008 and Windows Vista	At least Windows Server 2008 R2 and Windows 7
Enabled	Services running as Local System that use Negotiate will use the computer identity. This might cause some authentication requests between Windows operating systems to fail and log an error.	Services running as Local System that use Negotiate will use the computer identity. This is the default behavior.
Disabled	Services running as Local System that use Negotiate when reverting to NTLM authentication will authenticate anonymously. This is the default behavior.	Services running as Local System that use Negotiate when reverting to NTLM authentication will authenticate anonymously.
Neither	Services running as Local System that use Negotiate when reverting to NTLM authentication will authenticate anonymously.	Services running as Local System that use Negotiate will use the computer identity. This might cause some authentication requests between Windows operating systems to fail and log an error.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or Group Policy object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not applicable
Member server effective default settings	Not applicable
Effective GPO default settings on client computers	Not defined

Operating system version differences

This policy was introduced in Windows Server 2008 and Windows Vista, and it applies to authentication connectivity between these versions and later versions of the Windows operating system.

Beginning with Windows Server 2008 R2 and Windows 7, the operating system will by default not allow anonymous authentication because the system service uses the computer identity.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflict considerations

The policy **Network security: Allow LocalSystem NULL session fallback**, if enabled, will allow NTLM or Kerberos authentication to be used when a system service attempts authentication. This will increase the success of interoperability at the expense of security.

The anonymous authentication behavior is different for Windows Server 2008 and Windows Vista than later versions of Windows. Configuring and applying this policy setting on those systems might not produce the same results.

Group Policy

This policy setting can be configured by using the Group Policy Management Console (GPMC) to be distributed through Group Policy Objects (GPOs). If this policy is not contained in a distributed

GPO, this policy can be configured on the local computer by using the Local Security Policy snap-in.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When a service connects to computers running versions of Windows earlier than Windows Vista or Windows Server 2008, services that run as Local System and use SPNEGO (Negotiate) that revert to NTLM will use NULL session. In Windows Server 2008 R2 and Windows 7, if a service connects to a computer running Windows Server 2008 or Windows Vista, the system service uses the computer identity.

When a service connects with the computer identity, signing and encryption are supported to provide data protection. When a service connects with a NULL session, a system-generated session key is created, which provides no protection, but it allows applications to sign and encrypt data without errors.

Countermeasure

You can configure the **Network security: Allow Local System to use computer identity for NTLM** security policy setting to allow Local System services that use Negotiate to use the computer identity when reverting to NTLM authentication.

Potential impact

If you do not configure this policy setting on Windows Server 2008 and Windows Vista, services running as Local System that use the default credentials will use the NULL session and revert to NTLM authentication for Windows operating systems earlier than Windows Vista or Windows Server 2008.

Beginning with Windows Server 2008 R2 and Windows 7, the system allows Local System services that use Negotiate to use the computer identity when reverting to NTLM authentication.

Network security: Allow LocalSystem NULL session fallback

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This policy affects session security during the authentication process between computers running Windows Server 2008 R2 and Windows 7 and those computers running earlier versions of the Windows operating system. For computers running Windows Server 2008 R2 and Windows 7, services running as Local System require a service principal name (SPN) to generate the session key. However, if **Network security: Allow Local System to use computer identity for NTLM** is set to disabled, services running as Local System will fall back to using NULL session authentication when they transmit data to servers running versions of Windows earlier than Windows Vista or Windows Server 2008. NULL session does not establish a unique session key for each authentication; and thus, it cannot provide integrity or confidentiality protection. The setting **Network security: Allow LocalSystem NULL session fallback** determines whether services that request the use of session security are allowed to perform signature or encryption functions with a well-known key for application compatibility.

Possible values

- **Enabled**

  When a service running as Local System connects with a NULL session, a system-generated session key is created, which provides no protection but allows applications to sign and encrypt data without errors. This increases application compatibility, but it degrades the level of security.

- **Disabled**

  When a service running as Local System connects with a NULL session, session security will be unavailable. Calls seeking encryption or signing will fail. This setting is more secure, but at the risk of degrading application incompatibility. Calls that are using the computer identity instead of a NULL session will still have full use of session security.

- Not defined. When this policy is not defined, the default takes effect. This is Enabled for versions of the Windows operating system earlier than Windows Server 2008 R2 and Windows 7, and it is Disabled otherwise.

Best practices

When services connect with the computer identity, signing and encryption are supported to provide data protection. When services connect with a NULL session, this level of data protection is not provided. However, you will need to evaluate your environment to determine the Windows operating system versions that you support. If this policy is enabled, some services may not be able to authenticate.

This policy applies to Windows Server 2008 and Windows Vista (SP1 and later). When your environment no longer requires support for Windows NT 4, this policy should be disabled. By default, it is disabled in Windows 7 and Windows Server 2008 R2.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not applicable
Member server effective default settings	Not applicable
Effective GPO default settings on client computers	Not applicable

Operating system version differences

This policy was introduced in Windows Server 2008 R2 and Windows 7, and it can be applied to all computers running at least Windows Server 2008 with SP1 and Windows Vista with SP1.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If this setting is Enabled, when a service connects with a NULL session, a system-generated session key is created, which provides no protection but allows applications to sign and encrypt data without errors. Data that is intended to be protected might be exposed.

Countermeasure

You can configure the computer to use the computer identity for Local System with the policy **Network security: Allow Local System to use computer identity for NTLM**. If that is not possible, this policy can be used to prevent data from being exposed in transit if it was protected with a well-known key. To configure this policy, see Possible values for these options.

Potential impact

If you enable this policy, services that use NULL session with Local System could fail to authenticate because they will be prohibited from using signing and encryption.

Network Security: Allow PKU2U authentication requests to this computer to use online identities

This security policy reference topic for the IT professional describes the best practices, location, and values for this policy setting.

Reference

In Windows Server 2008 R2 and Windows 7 the Negotiate Security Support Provider (SSP) supports an extension SSP, Negoexts.dll. This extension SSP is treated as an authentication protocol by the Windows operating system, and it supports SSPs from Microsoft, including PKU2U. You can also develop or add other SSPs.

When computers are configured to accept authentication requests by using online IDs, Negoexts.dll calls the PKU2U SSP on the computer that is used to log on. The PKU2U SSP obtains a local certificate and exchanges the policy between the peer computers. When validated on the peer computer, the certificate within the metadata is sent to the logon peer for validation. It associates the user's certificate to a security token, and then the logon process completes.

Note

The ability to link online IDs can be performed by anyone with an account that has Standard User’s credentials through the **Credential Manager** under User Accounts from the Control Panel.

This policy is not configured by default on domain-joined computers. This would disallow the online identities to be able to authenticate to the domain-joined computers in Windows 7.

Possible values

- **Enabled**

  This will allow authentication to successfully complete between the two (or more) computers that have established a peer relationship through the use on online IDs. The PKU2U SSP obtains a local certificate and exchanges the policy between the peer computers. When validated on the peer computer, the certificate within the metadata is sent to the logon peer for validation. It associates the user's certificate to a security token, and then the logon process completes.
- **Disabled**

This will prevent online IDs from being used to authenticate the user to another computer in a peer-to-peer relationship.

- Not set. Not configuring this policy prevents online IDs from being used to authenticate the user. This is the default on domain-joined computers.

Best practices

Within a domain, domain accounts should be used for authentication. Set this policy to **Disabled** or do not configure this policy to exclude online identities from being used to authenticate.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Disabled
Member server effective default settings	Disabled
Effective GPO default settings on client computers	Disabled

Operating system version differences

This policy was introduced in Windows Server 2008 R2 and Windows 7, and it can be applied to Windows Server operating systems that support the Negotiate Extensions SSP (Negoexts.dll).

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Enabling this policy setting allows a user's account on one computer to be associated with an online identity, such as Windows Live ID, so that account can log on to a peer computer (if the peer computer is likewise configured) without the use of a Windows logon account (domain or local). Although this is beneficial for workgroups or home groups, using this feature in a domain-joined environment might circumvent your established security policies.

Countermeasure

Set this policy to Disabled or do not configure this security policy for domain-joined computers.

Potential impact

If you do not set or disable this policy, the PKU2U protocol will not be used to authenticate between peer computers, which forces users to follow domain defined access control policies. If you enable this policy, you will allow your users to authenticate by using local certificates between systems that are not part of a domain that uses PKU2U. This will allow users to share resources between computers.

Network security: Configure encryption types allowed for Kerberos

This security policy reference topic for the IT professional describes the best practices, location, values and security considerations for this policy setting.

Reference

This policy setting allows you to set the encryption types that the Kerberos protocol is allowed to use. If it is not selected, the encryption type will not be allowed. This setting might affect compatibility with client computers or services and applications. Multiple selections are permitted.

For more information, see article 977321 in the Microsoft Knowledge Base.

The following table lists and explains the allowed encryption types.

Encryption type	Description and version support
DES_CBC_CRC	Data Encryption Standard with Cipher Block Chaining using the Cyclic Redundancy Check function Supported in Windows 2000 Server, Windows XP, Windows Server 2003, Windows Vista, and Windows Server 2008. The

Encryption type	Description and version support
	Windows 7 and Windows Server 2008 R2 operating systems do not support DES by default.
DES_CBC_MD5	Data Encryption Standard with Cipher Block Chaining using the Message-Digest algorithm 5 checksum function Supported in Windows 2000 Server, Windows XP, Windows Server 2003, Windows Vista, and Windows Server 2008. The Windows 7 and Windows Server 2008 R2 operating systems do not support DES by default.
RC4_HMAC_MD5	Rivest Cipher 4 with Hashed Message Authentication Code using the Message-Digest algorithm 5 checksum function Supported in Windows 2000 Server, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2.
AES128_HMAC_SHA1	Advanced Encryption Standard in 128 bit cipher block with Hashed Message Authentication Code using the Secure Hash Algorithm (1). Not supported in Windows 2000 Server, Windows XP, or Windows Server 2003. Supported in Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2.
AES256_HMAC_SHA1	Advanced Encryption Standard in 256 bit cipher block with Hashed Message Authentication Code using the Secure Hash Algorithm (1). Not supported in Windows 2000 Server, Windows XP, or Windows Server 2003. Supported in Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2.
Future encryption types	Reserved by Microsoft for additional encryption types that might be implemented.

Possible values

The encryption type options include:

- DES_CBC_CRC
- DES_CBC_MD5
- RC4_HMAC_MD5
- AES128_HMAC_SHA1
- AES256_HMAC_SHA1
- Future encryption types

  As of the release of Windows 7 and Windows Server 2008 R2, this is reserved by Microsoft for additional encryption types that might be implemented.

Best practices

You must analyze your environment to determine which encryption types will be supported and then select those that meet that evaluation.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or Group Policy Object (GPO)	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	None of these encryption types that are available in this policy are allowed.
Member server effective default settings	None of these encryption types that are available in this policy are allowed.
Effective GPO default settings on client computers	None of these encryption types that are available in this policy are allowed.

Operating system version differences

This policy was introduced in Windows Server 2008 R2 and Windows 7, and it can be applied to Windows operating systems beginning with Windows 2000.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Windows Server 2008 R2 and Windows 7 do not support the DES cryptographic suites because stronger ones are available. To enable Kerberos interoperability with non-Windows versions of the Kerberos protocol, these suites can be enabled. However, doing so might open attack vectors on computers running Windows Server 2008 R2 and Windows 7. You can also disable DES for your computers running Windows Vista and Windows Server 2008.

Countermeasure

Do not configure this policy. This will force the computers running Windows Server 2008 R2 and Windows 7 to use the AES or RC4 cryptographic suites.

Potential impact

If you do not select any of the encryption types, computers running Windows Server 2008 R2 and Windows 7 might have Kerberos authentication failures when connecting with computers running non-Windows versions of the Kerberos protocol.

If you do select any encryption type, you will lower the effectiveness of encryption for Kerberos authentication but you will improve interoperability with computers running older versions of Windows.

Contemporary non-Windows implementations of the Kerberos protocol support RC4 and AES 128-bit and AES 256-bit encryption. Most implementations, including the MIT Kerberos protocol and the Windows Kerberos protocol, are deprecating DES encryption.

Network security: Do not store LAN Manager hash value on next password change

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting. This information applies to computers running at least the Windows Server 2008 operating system.

Reference

This policy setting determines whether LAN Manager is prevented from storing hash values for the new password the next time the password is changed. Hash values are a representation of

the password after the encryption algorithm is applied that corresponds to the format that is specified by the algorithm. To decrypt the hash value, the encryption algorithm must be determined and then reversed. The LAN Manager hash is relatively weak and prone to attack compared to the cryptographically stronger NTLM hash. Because the LM hash is stored on the local computer in the security database, the passwords can be compromised if the security database, Security Accounts Manager (SAM), is attacked.

By attacking the SAM file, attackers can potentially gain access to user names and password hashes. Attackers can use a password-cracking tool to determine what the password is. After they have access to this information, they can use it to gain access to resources on your network by impersonating users. Enabling this policy setting will not prevent these types of attacks, but it will make them much more difficult.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set **Network security: Do not store LAN Manager hash value on next password change** to **Enabled**.
2. Require all users to set new passwords the next time they log on to the domain so that LAN Manager hashes are removed. Legacy operating systems (such as Windows 95, Windows 98, and Windows Millennium Edition) will fail, as will some non-Microsoft applications.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled

Server type or GPO	Default value
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

On Windows operating systems before Windows Server 2008 and Windows Vista, the default for this policy was Disabled.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The SAM file can be targeted by attackers who seek access to user names and password hashes. Such attacks use special tools to discover passwords, which can then be used to impersonate users and gain access to resources on your network. These types of attacks are not prevented by enabling this policy setting because LAN Manager hashes are much weaker than NTLM hashes, but it is much more difficult for these attacks to succeed.

Countermeasure

Enable the **Network security: Do not store LAN Manager hash value on next password change** setting. Require all users to set new passwords the next time they log on to the domain so that LAN Manager hashes are removed.

Potential impact

Earlier operating systems such as Windows 95, Windows 98, and Windows Millennium Edition, as well as some non-Microsoft applications, cannot connect to the system.

Network security: Force logoff when logon hours expire

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting. This information applies to computers running at least the Windows Server 2008 operating system.

Reference

This security setting determines whether to disconnect users who are connected to the local computer outside their user account's valid logon hours. This setting affects the Server Message Block (SMB) component.

This policy setting does not apply to Administrator accounts, but it behaves as an account policy. For domain accounts, there can be only one account policy. The account policy must be defined in the Default Domain Policy, and it is enforced by the domain controllers that make up the domain. A domain controller always pulls the account policy from the Default Domain Policy Group Policy Object (GPO), even if there is a different account policy that is applied to the organizational unit that contains the domain controller. By default, workstations and servers that are joined to a domain (for example, member computers) also receive the same account policy for their local accounts. However, local account policies for member computers can be different from the domain account policy by defining an account policy for the organizational unit that contains the member computers. Kerberos settings are not applied to member computers.

Possible values

- Enabled

  When enabled, this policy causes client sessions with the SMB server to be forcibly disconnected when the client's logon hours expire.

- Disabled

  When disabled, this policy allows for the continuation of an established client session after the client's logon hours have expired.

- Not defined

Best practices

1. Set **Network security: Force logoff when logon hours expire** to Enabled. SMB sessions will be terminated on member servers when a user's logon time expires, and the user will be unable to log on to the system until their next scheduled access time begins.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Disabled
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If you disable this policy setting, users can remain connected to the computer outside of their allotted logon hours.

Countermeasure

Enable the **Network security: Force logoff when logon hours expire** setting. This policy setting does not apply to administrator accounts.

Potential impact

When a user's logon time expires, SMB sessions terminate. The user cannot log on to the computer until the next scheduled access time commences.

Network security: LAN Manager authentication level

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting. This information applies to computers running at least the Windows Server 2008 operating system.

Reference

This policy setting determines which challenge or response authentication protocol is used for network logons. LAN Manager (LM) includes client computer and server software from Microsoft that allows users to link personal computers together on a single network. Network capabilities include transparent file and print sharing, user security features, and network administration tools. In Active Directory domains, the Kerberos protocol is the default authentication protocol. However, if the Kerberos protocol is not negotiated for some reason, Active Directory uses LM, NTLM, or NTLM version 2 (NTLMv2).

LAN Manager authentication includes the LM, NTLM, and NTLMv2 variants, and it is the protocol that is used to authenticate all client computers running the Windows operating system when they perform the following operations:

- Join a domain
- Authenticate between Active Directory forests
- Authenticate to domains based on earlier versions of the Windows operating system
- Authenticate to computers that do not run Windows operating systems, beginning with Windows 2000
- Authenticate to computers that are not in the domain

Possible values

- Send LM & NTLM responses
- Send LM & NTLM - use NTLMv2 session security if negotiated
- Send NTLM responses only
- Send NTLMv2 responses only

- Send NTLMv2 responses only. Refuse LM
- Send NTLMv2 responses only. Refuse LM & NTLM
- Not Defined

The **Network security: LAN Manager authentication level** setting determines which challenge/response authentication protocol is used for network logons. This choice affects the authentication protocol level that clients use, the session security level that the computers negotiate, and the authentication level that servers accept. The following table identifies the policy settings, describes the setting, and identifies the security level used in the corresponding registry setting if you choose to use the registry to control this setting instead of the policy setting.

Setting	Description	Registry security level
Send LM & NTLM responses	Client computers use LM and NTLM authentication, and they never use NTLMv2 session security. Domain controllers accept LM, NTLM, and NTLMv2 authentication.	0
Send LM & NTLM – use NTLMv2 session security if negotiated	Client computers use LM and NTLM authentication, and they use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.	1
Send NTLM response only	Client computers use NTLMv1 authentication, and they use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.	2
Send NTLMv2 response only	Client computers use NTLMv2 authentication, and they use NTLMv2 session security if the server supports it. Domain controllers accept LM, NTLM, and NTLMv2 authentication.	3
Send NTLMv2 response only. Refuse LM	Client computers use NTLMv2 authentication, and they use NTLMv2 session security if the server supports it. Domain controllers refuse to accept LM	4

Setting	Description	Registry security level
	authentication, and they will accept only NTLM and NTLMv2 authentication.	
Send NTLMv2 response only. Refuse LM & NTLM	Client computers use NTLMv2 authentication, and they use NTLMv2 session security if the server supports it. Domain controllers refuse to accept LM and NTLM authentication, and they will accept only NTLMv2 authentication.	5

Best practices

1. Best practices are dependent on your specific security and authentication requirements.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Send NTLMv2 response only
DC Effective Default Settings	Send NTLMv2 response only
Member Server Effective Default Settings	Send NTLMv2 response only
Client Computer Effective Default Settings	Not defined

Operating system version differences

In Windows Server 2003, the Default Domain Controllers Policy was **Send NTLM response only**, which changed to **Not defined** in later versions.

In Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2, the default is **Send NTLMv2 response only**.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Modifying this setting may affect compatibility with client computers, services, and applications.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

In Windows 7 and Windows Vista, this setting is undefined. In Windows Server 2008 R2 and Windows Server 2008 this setting is configured to **Send NTLMv2 responses only**. In Windows 2000, Windows Server 2003, and Windows XP, client computers are configured by default to send LM and NTLM authentication responses (Windows 95-based and Windows 98-based client computers only send LM). The default setting on servers allows all client computers to authenticate with servers and use their resources. However, this means that LM responses—the weakest form of authentication response—are sent over the network, and it is potentially possible for attackers to intercept that traffic to reproduce the user's password more easily.

The Windows 95, Windows 98, and Windows NT operating systems cannot use Kerberos protocol version 5 for authentication. For this reason, in a Windows Server 2003 domain, these computers authenticate by default with both the LM and NTLM protocols for network authentication. You can enforce a more secure authentication protocol for Windows 95, Windows 98, and Windows NT by using NTLMv2. For the logon process, NTLMv2 uses a secure channel to protect the authentication process. Even if you use NTLMv2 for client computers and servers running these early versions of the Windows operating system, Windows-based client computers and servers that are members of the domain use the Kerberos protocol to authenticate with Windows Server 2003 domain controllers.

Countermeasure

Configure the **Network security: LAN Manager Authentication Level** setting to **Send NTLMv2 responses only**. Microsoft and a number of independent organizations strongly recommend this level of authentication when all client computers support NTLMv2.

For more information about how to enable NTLMv2 on older versions of the Windows operating system, see article 239869 in the Microsoft Knowledge Base. Windows NT 4.0 requires Service Pack 4 (SP4) to support NTLMv2, and Windows 95 and Windows 98 need the directory service client installed to support NTLMv2.

Potential impact

Client computers that do not support NTLMv2 authentication cannot authenticate in the domain and access domain resources by using LM and NTLM.

Note

For information about a hotfix to ensure that this setting works in networks that include Windows NT 4.0-based computers along with Windows 2000, Windows XP, and Windows Server 2003-based computers, see article 305379 in the Microsoft Knowledge Base.

Network security: LDAP client signing requirements

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting. This information applies to computers running at least the Windows Server 2008 operating system.

Reference

This policy setting determines the level of data signing that is requested on behalf of client computers that issue LDAP BIND requests. The levels of data signing are described in the following list:

- **None**. The LDAP BIND request is issued with the caller-specified options.
- **Negotiate signing**. If Transport Layer Security/Secure Sockets Layer (TLS/SSL) has not been started, the LDAP BIND request is initiated with the LDAP data signing option set in addition to the caller-specified options. If TLS/SSL has been started, the LDAP BIND request is initiated with the caller-specified options.
- **Require signing**. This level is the same as **Negotiate signing**. However, if the LDAP server's intermediate saslBindInProgress response does not indicate that LDAP traffic signing is required, the caller is returned a message that the LDAP BIND command request failed.

Note

This policy setting does not have any impact on ldap_simple_bind or ldap_simple_bind_s. LDAP client computers that are running Windows XP Professional or Windows Vista do not use ldap_simple_bind or ldap_simple_bind_s to communicate with a domain controller.

Misuse of this policy setting is a common error that can cause data loss or problems with data access or security.

Possible values

- None
- Negotiate signing
- Require signature
- Not Defined

Best practices

1. Set **Domain controller: LDAP server signing requirements** to **Require signature**. If you set the server to require LDAP signatures, you must also set the client computers to do so. Not setting the client computers will prevent client computers from communicating with the server. This can cause many features to fail, including user authentication, Group Policy, and logon scripts.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Negotiate signing
DC Effective Default Settings	Negotiate signing
Member Server Effective Default Settings	Negotiate signing
Client Computer Effective Default Settings	Negotiate signing

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Modifying this setting may affect compatibility with client computers, services, and applications.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Unsigned network traffic is susceptible to man-in-the-middle attacks in which an intruder captures the packets between the client computer and server, modifies them, and then forwards them to the server. For an LDAP server, this susceptibility means that an attacker could cause a server to make decisions that are based on false or altered data from the LDAP queries. To lower this risk in your network, you can implement strong physical security measures to protect the network infrastructure. Also, you can make all types of man-in-the-middle attacks extremely difficult if you require digital signatures on all network packets by means of IPsec authentication headers.

Countermeasure

Configure the **Network security: LDAP server signing requirements** setting to **Require signature**.

Potential impact

If you configure the server to require LDAP signatures, you must also configure the client computers. If you do not configure the client computers, they cannot communicate with the server, which could cause many features to fail, including user authentication, Group Policy, and logon scripts.

Network security: Minimum session security for NTLM SSP based (including secure RPC) clients

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting allows a client computer to require the negotiation of 128-bit encryption or NTLMv2 session security. These values are dependent on the **Network security: LAN Manager Authentication Level policy** setting value.

Possible values

- Require NTLMv2 session security

  The connection fails if strong encryption (128-bit) is not negotiated.
- Require 128-bit encryption

  The connection fails if the NTLMv2 protocol is not negotiated.

Best practices

Practices in setting this policy are dependent on your security requirements.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Require 128-bit encryption
DC Effective Default Settings	Require 128-bit encryption

Server type or GPO	Default value
Member Server Effective Default Settings	Require 128-bit encryption
Client Computer Effective Default Settings	Require 128-bit encryption

Operating system version differences

In Windows Server 2008 and earlier versions of the Windows operating system, no requirements were set. But in computers running Windows Server 2008 R2 **Require 128-bit encryption** is the default.

In Windows Server 2008 R2 and Windows 7, only 128-bit encryption or NTLMv2 session security options are available.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy conflicts

The settings for this security policy are dependent on the **Network security: LAN Manager Authentication Level policy** setting value. For information about this policy, see Network security: LAN Manager authentication level.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Network traffic that uses the NTLM Security Support Provider (NTLM SSP) could be exposed such that an attacker who has gained access to the network can create man-in-the-middle attacks.

Countermeasure

Enable all options that are available for the **Network security: Minimum session security for NTLM SSP based (including secure RPC) clients policy** setting.

Potential impact

Client computers that enforce these settings cannot communicate with older servers that do not support them.

Network security: Minimum session security for NTLM SSP based (including secure RPC) servers

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting allows a client computer to require the negotiation of 128-bit encryption or NTLMv2 session security. These values are dependent on the **Network security: LAN Manager Authentication Level policy** setting value.

Setting all of these values for this policy setting will help protect network traffic that uses the NTLM Security Support Provider (NTLM SSP) from being exposed or tampered with by a malicious user who has gained access to the same network. That is, these settings help protect against man-in-the-middle attacks.

Possible values

- Require 128-bit encryption. The connection fails if strong encryption (128-bit) is not negotiated.
- Require NTLMv2 session security. The connection fails if the NTLMv2 protocol is not negotiated.
- Not Defined.

Best practices

1. Enable all values that are available for this security policy. Legacy client computers that do not support these policy settings will be unable to communicate with the server.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Require 128-bit encryption
DC Effective Default Settings	Require 128-bit encryption
Member Server Effective Default Settings	Require 128-bit encryption
Client Computer Effective Default Settings	Require 128-bit encryption

Operating system version differences

In Windows Server 2008 and earlier versions, no requirements were set. But in later versions of the Windows operating system, **Require 128-bit encryption** is the default.

In Windows Server 2008 R2 and Windows 7, only 128-bit encryption or NTLMv2 session security options are available.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Policy dependencies

The settings for this security policy are dependent on the **Network security: LAN Manager Authentication Level policy** setting value. For information about this policy, see Network security: LAN Manager authentication level.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Network traffic that uses the NTLM Security Support Provider (NTLM SSP) could be exposed such that an attacker who has gained access to the network can create man-in-the-middle attacks.

Countermeasure

Enable all options that are available for the **Network security: Minimum session security for NTLM SSP based (including secure RPC) servers** policy setting.

Potential impact

Older client computers that do not support these security settings cannot communicate with the computer on which this policy is set.

Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network security: Restrict NTLM: Add server exceptions for NTLM authentication** policy setting allows you to create an exception list of remote servers to which client computers are allowed to use NTLM authentication if the **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** policy setting is configured.

If you configure this policy setting, you can define a list of remote servers to which client computers are allowed to use NTLM authentication.

If you do not configure this policy setting, no exceptions will be applied, and if **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** is enabled, NTLM authentication attempts from the client computers will fail.

List the NetBIOS server names that are used by the applications as the naming format, one per line. To ensure exceptions, the names that are used by all applications need to be in the list. A single asterisk (*) can be used anywhere in the string as a wildcard character.

Possible values

- User-defined list of remote servers

  When you enter a list of remote servers to which clients are allowed to use NTLM authentication, the policy is defined and enabled.

- Not defined

  If you do not configure this policy setting by defining a list of servers, the policy is undefined and no exceptions will be applied.

Best practices

1. First enforce the **Network Security: Restrict NTLM: Audit Incoming NTLM Traffic** or **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting and then review the Operational log to understand which servers are involved in these authentication attempts so you can decide which servers to exempt.
2. After you have set the server exception list, enforce the **Network Security: Restrict NTLM: Audit Incoming NTLM Traffic** or **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting and then review the Operational log again before setting the policies to block NTLM traffic.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes the features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy through Group Policy takes precedence over the setting on the local computer. If the Group Policy setting is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if your server exception list is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**.

There are no Security Audit policies that can be configured to view output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When it has been determined that the NTLM authentication protocol should not be used from a client computer to any remote servers because you are required to use a more secure protocol such as Kerberos, there might be some client applications that still use NTLM. If so, and you set **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** to any of the deny options, those applications will fail because the outbound NTLM authentication traffic from the client computer will be blocked.

If you define an exception list of servers to which client computers are allowed to use NTLM authentication, then NTLM authentication traffic will continue to flow between those client applications and servers. The servers then are vulnerable to any malicious attack that takes advantage of security weaknesses in NTLM.

Countermeasure

When you use **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** in audit-only mode, you can determine by reviewing which client applications are making NTLM authentication requests to the remote servers in your environment. When assessed, you will have to determine on a case-by-case basis if NTLM authentication still minimally meets your security requirements. If not, the client application has to be upgraded to use something other than NTLM authentication.

Potential impact

Defining a list of servers for this policy setting will enable NTLM authentication traffic from the client application that uses those servers, and this might result in a security vulnerability.

If this list is not defined and **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** is enabled, then client applications that use NTLM will fail to authenticate to those servers that they have previously used.

Network security: Restrict NTLM: Add server exceptions in this domain

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting allows you to create an exception list of servers in this domain to which client computers are allowed to use NTLM pass-through authentication if any of the deny options are set in the **Network Security: Restrict NTLM: NTLM authentication in this domain** policy setting.

If you configure this policy setting, you can define a list of servers in this domain to which client computers are allowed to use NTLM authentication.

If you do not configure this policy setting, no exceptions will be applied, and if **Network Security: Restrict NTLM: NTLM authentication in this domain** is enabled, all NTLM authentication attempts in the domain will fail.

List the NetBIOS server names as the naming format, one per line. A single asterisk (*) can be used anywhere in the string as a wildcard character.

Possible values

- User-defined list of servers

  When you enter a list of servers in this domain to which clients are allowed to use NTLM authentication, the policy is defined and enabled.
- Not defined

  If you do not configure this policy setting by defining a list of servers, the policy is undefined and no exceptions will be applied.

Best practices

1. First enforce the **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting, and then review the Operational log to understand what domain

controllers are involved in these authentication attempts so you can decide which servers to exempt.

2. After you have set the server exception list, enforce the **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting, and then review the Operational log again before setting the policies to block NTLM traffic.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy via Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if your server exception list is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**.

There are no Security Audit policies that can be configured to view output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

When it has been determined that the NTLM authentication protocol should not be used within a domain because you are required to use a more secure protocol such as Kerberos, there might be some NTLM authentication traffic that is still present in the domain. If so, and you set **Network Security: Restrict NTLM: NTLM authentication in this domain** to any of the deny options, any NTLM authentication request will fail because the pass-through member server will block the NTLM request.

If you define an exception list of servers in this domain to which client computers are allowed to use NTLM pass-through authentication, then NTLM authentication traffic will continue to flow between those servers, which make them vulnerable to any malicious attack that takes advantage of security weaknesses in NTLM.

Countermeasure

When you use **Network Security: Restrict NTLM: NTLM authentication in this domain** in audit-only mode, you can determine by reviewing which client applications are making NTLM authentication requests to the pass-through authentication servers. When assessed, you will have to determine on a case-by-case basis if NTLM authentication still minimally meets your security requirements.

Potential impact

Defining a list of servers for this policy setting will enable NTLM authentication traffic between those servers might result in a security vulnerability.

If this list is not defined and **Network Security: Restrict NTLM: NTLM authentication in this domain** is enabled, then NTLM authentication will fail on those pass-through servers in the domain that they have previously used

Network Security: Restrict NTLM: Incoming NTLM Traffic

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network Security: Restrict NTLM: Incoming NTLM Traffic** policy setting allows you to deny or allow incoming NTLM traffic from client computers, other member servers, or a domain controller.

Possible values

- **Allow all**

  The server will allow all NTLM authentication requests.
- **Deny all domain accounts**

  The server will deny NTLM authentication requests for domain logon, return an NTLM blocked error message to the client computer, and log the error, but the server will allow local account logon.
- **Deny all accounts**

  The server will deny NTLM authentication requests from all incoming traffic (whether domain account logon or local account logon), return an NTLM blocked error message to the client computer, and log the error.
- Not defined

  This is the same as **Allow all**, and the server will allow all NTLM authentication requests. See the Group Policy section for additional information.

Best practices

If you select **Deny all domain accounts** or **Deny all accounts**, incoming NTLM traffic to the member server will be restricted. It is better to set the **Network Security: Restrict NTLM: Audit Incoming NTLM Traffic** policy setting and then review the Operational log to understand what authentication attempts are made to the member servers, and subsequently what client applications are using NTLM.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if this policy is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**.

There are no Security Audit Event policies that can be configured to view event output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

NTLM and NTLMv2 authentication is vulnerable to a variety of malicious attacks, including SMB replay, man-in-the-middle attacks, and brute force attacks. Reducing and eliminating NTLM authentication from your environment forces the Windows operating system to use more secure protocols, such as the Kerberos version 5 protocol, or different authentication mechanisms, such as smart cards.

Vulnerability

Malicious attacks on NTLM authentication traffic that result in a compromised server can occur only if the server handles NTLM requests. If those requests are denied, brute force attacks on NTLM are eliminated.

Countermeasure

When it has been determined that the NTLM authentication protocol should not be used within a network because you are required to use a more secure protocol such as Kerberos, you can select one of several options that this security policy setting offers to restrict NTLM usage.

Potential impact

If you configure this policy setting, numerous NTLM authentication requests could fail within your network, which could degrade productivity. Before implementing this change through this policy setting, set **Network security: Restrict NTLM: Audit Incoming NTLM Traffic** to the same option so that you can review the log for the potential impact, perform an analysis of servers, and create an exception list of servers to exclude from this policy setting **Network security: Restrict NTLM: Add server exceptions in this domain**.

Network Security: Restrict NTLM: NTLM authentication in this domain

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network Security: Restrict NTLM: NTLM authentication in this domain** policy setting allows you to deny or allow NTLM authentication within a domain from this domain controller. This policy setting does not affect interactive logon to this domain controller.

Possible values

- **Disable**

  The domain controller will allow all NTLM pass-through authentication requests within the domain.

- **Deny for domain accounts to domain servers**

  The domain controller will deny all NTLM authentication logon attempts using accounts from this domain to all servers in the domain. The NTLM authentication attempts will be blocked and will return an NTLM blocked error unless the server name is on the exception list in the **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting.

  NTLM can be used if the users are connecting to other domains. This depends on if any Restrict NTLM policies have been set on those domains.

- **Deny for domain accounts**

  Only the domain controller will deny all NTLM authentication logon attempts from domain accounts and will return an NTLM blocked error unless the server name is on the exception list in the **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting.

- **Deny for domain servers**

  The domain controller will deny NTLM authentication requests to all servers in the domain and will return an NTLM blocked error unless the server name is on the exception list in the **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting. Servers that are not joined to the domain will not be affected if this policy setting is configured.

- **Deny all**

  The domain controller will deny all NTLM pass-through authentication requests from its servers and for its accounts and return an NTLM blocked error unless the server name is on the exception list in the **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting.

- Not defined

  The domain controller will allow all NTLM authentication requests in the domain where the policy is deployed.

Best practices

If you select any of the deny options, incoming NTLM traffic to the domain will be restricted. First, set the **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting, and then review the Operational log to understand what authentication attempts are made to the member servers. You can then add those member server names to a server exception list by using the **Network security: Restrict NTLM: Add server exceptions in this domain** policy setting.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not configured
Default domain controller policy	Not configured
Stand-alone server default settings	Not configured
Domain controller effective default settings	Not configured
Member server effective default settings	Not configured
Client computer effective default settings	Not configured

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if this policy is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**.

There are no Security Audit Event policies that can be configured to view output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

NTLM and NTLMv2 authentication is vulnerable to a variety of malicious attacks, including SMB replay, man-in-the-middle attacks, and brute force attacks. Reducing and eliminating NTLM authentication from your environment forces the Windows operating system to use more secure protocols, such as the Kerberos version 5 protocol, or different authentication mechanisms, such as smart cards.

Vulnerability

Malicious attacks on NTLM authentication traffic resulting in a compromised server or domain controller can occur only if the server or domain controller handles NTLM requests. If those requests are denied, this attack vector is eliminated.

Countermeasure

When it has been determined that the NTLM authentication protocol should not be used within a network because you are required to use a more secure protocol such as the Kerberos protocol, then you can select one of several options that this security policy setting offers to restrict NTLM usage within the domain.

Potential impact

If you configure this policy setting, numerous NTLM authentication requests could fail within the domain, which could degrade productivity. Before implementing this change through this policy setting, set **Network security: Restrict NTLM: Audit NTLM authentication in this domain** to the same option so that you can review the log for the potential impact, perform an analysis of servers, and create an exception list of servers to exclude from this policy setting by using **Network security: Restrict NTLM: Add server exceptions in this domain**.

Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network Security: Restrict NTLM: Outgoing NTLM traffic to remote servers** policy setting allows you to deny or audit outgoing NTLM traffic from a computer running Windows 7 or Windows Server 2008 R2 to any remote server running the Windows operating system.

Warning
Modifying this policy setting may affect compatibility with client computers, services, and applications.

Possible values

- **Allow all**

  The computer can authenticate identities to a remote server by using NTLM authentication because no restrictions exist.

- **Audit all**

  The computer that sends the NTLM authentication request to a remote server logs an event for each request. This allows you to identify those servers that receive NTLM authentication requests from the client computer.

- **Deny all**

  The computer cannot authenticate any identities to a remote server by using NTLM authentication. You can use the **Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication** policy setting to define a list of remote servers to which client computers are allowed to use NTLM authentication while denying others. This setting will also log an event on the computer that is making the authentication request.

- Not defined

  This is the same as **Allow all**, and the computer will allow all NTLM authentication requests when the policy is deployed. See the Group Policy section for additional information.

Best practices

If you select **Deny all**, the client computer cannot authenticate identities to a remote server by using NTLM authentication. First, select **Audit all** and then review the Operational log to understand which servers are involved in these authentication attempts. You can then add those server names to a server exception list by using the **Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication** policy setting.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if this policy is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**.

There are no Security Audit Event policies that can be configured to view event output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

NTLM and NTLMv2 authentication is vulnerable to a variety of malicious attacks, including SMB replay, man-in-the-middle attacks, and brute force attacks. Reducing and eliminating NTLM authentication from your environment forces the Windows operating system to use more secure protocols, such as the Kerberos version 5 protocol, or different authentication mechanisms, such as smart cards.

Vulnerability

Malicious attacks on NTLM authentication traffic that result in a compromised server or domain controller can occur only if the server or domain controller handles NTLM requests. If those requests are denied, this attack vector is eliminated.

Countermeasure

When it has been determined that the NTLM authentication protocol should not be used within a network because you are required to use a more secure protocol such as Kerberos, then you can select from several options to restrict NTLM usage to servers.

Potential impact

If you configure this policy setting to deny all requests, numerous NTLM authentication requests to remote servers could fail, which could degrade productivity. Before implementing this restriction through this policy setting, select **Audit all** so that you can review the log for the potential impact, perform an analysis of servers, and create an exception list of servers to exclude from this policy setting by using **Network security: Restrict NTLM: Add remote server exceptions for NTLM authentication**.

Network Security: Restrict NTLM: Audit Incoming NTLM Traffic

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting.

Reference

The **Network Security: Restrict NTLM: Audit Incoming NTLM Traffic** policy setting allows you to audit incoming NTLM traffic.

When this audit policy is enabled within Group Policy, it is enforced on any server where that Group Policy is distributed. The events will be recorded in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**. Using an audit event collection system can help you collect the events for analysis more efficiently.

When you enable this policy on a server, only authentication traffic to that server will be logged.

When you enable this audit policy, it functions in the same way as the **Network Security: Restrict NTLM: Incoming NTLM Traffic** policy, but it does not actually block any traffic. Therefore, you can use it effectively to understand the authentication traffic in your environment, and when you are ready to block that traffic, you can enable the **Network Security: Restrict NTLM: Incoming NTLM Traffic** policy setting and select **Deny all accounts** or **Deny all domain accounts**.

Possible values

- Disable

  The server on which this policy is set will not log events for incoming NTLM traffic.
- Enable auditing for domain accounts

  The server on which this policy is set will log events for NTLM pass-through authentication requests only for accounts in the domain that would be blocked when the **Network Security: Restrict NTLM: Incoming NTLM traffic** policy setting is set to **Deny all domain accounts**.
- Enable auditing for all accounts

  The server on which this policy is set will log events for all NTLM authentication requests that would be blocked when the **Network Security: Restrict NTLM: Incoming NTLM traffic** policy setting is set to **Deny all accounts**.
- Not defined

  This is the same as **Disable**, and it results in no auditing of NTLM traffic. See the Group Policy section for additional information.

Best practices

Depending on your environment and the duration of your testing, monitor the log size regularly.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined

Server type or GPO	Default value
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if this policy is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**. Using an audit event collection system can help you collect the events for analysis more efficiently.

There are no Security Audit Event policies that can be configured to view output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

NTLM and NTLMv2 authentication is vulnerable to a variety of malicious attacks, including SMB replay, man-in-the-middle attacks, and brute force attacks. Reducing and eliminating NTLM authentication from your environment forces the Windows operating system to use more secure protocols, such as the Kerberos version 5 protocol, or different authentication mechanisms, such as smart cards.

Vulnerability

Enabling this policy setting will reveal through logging which servers and client computers within your network or domain handle NTLM traffic. The identity of these computers can be used in malicious ways if NTLM authentication traffic is compromised. The policy setting does not prevent or mitigate any vulnerability because it is for audit purposes only.

Countermeasure

Restrict access to the log files when this policy setting is enabled in your production environment.

Potential impact

If you do not enable or configure this policy setting, no NTLM authentication traffic information will be logged. If you do enable this policy setting, only auditing functions will occur; no security enhancements will be implemented.

Network Security: Restrict NTLM: Audit NTLM authentication in this domain

This security policy reference topic for the IT professional describes the best practices, location, values, management aspects, and security considerations for this policy setting. This information applies to computers running at least the Windows Server 2008 R2 or Windows 7 operating systems.

Reference

The **Network Security: Restrict NTLM: Audit NTLM authentication in this domain** policy setting allows you to audit on the domain controller NTLM authentication in that domain.

When you enable this policy setting on the domain controller, only authentication traffic to that domain controller will be logged.

When you enable this audit policy, it functions in the same way as the **Network Security: Restrict NTLM: NTLM authentication in this domain** policy setting, but it does not actually block any traffic. Therefore, you can use it effectively to understand the authentication traffic to your domain controllers and when you are ready to block that traffic, you can enable the **Network Security: Restrict NTLM: NTLM authentication in this domain** policy setting and select **Deny for domain accounts to domain servers**, **Deny for domain servers**, or **Deny for domain accounts**.

Possible values

- **Disable**

  The domain controller on which this policy is set will not log events for incoming NTLM traffic.

- **Enable for domain accounts to domain servers**

  The domain controller on which this policy is set will log events for NTLM authentication logon attempts for accounts in the domain to domain servers when NTLM authentication would be denied because the **Network security: Restrict NTLM: NTLM authentication in this domain** policy setting is set to **Deny for domain accounts to domain servers**.
- **Enable for domain accounts**

  The domain controller will log events for NTLM authentication logon attempts that use domain accounts when NTLM authentication would be denied because the **Network security: Restrict NTLM: NTLM authentication in this domain** policy setting is set to **Deny for domain accounts**.
- Not defined

  This is the same as **Disable** and results in no auditing of NTLM traffic. See the Group Policy section for additional information.

Best practices

Depending on your environment and the duration of your testing, monitor the Operational log size regularly.

Location

*GPO_name***Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options**

Default values

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 R2 and Windows 7.

Policy management

This section describes different features and tools available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a restart when saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer. If the Group Policy is set to **Not Configured**, local settings will apply.

Auditing

View the Operational log to see if this policy is functioning as intended. Audit and block events are recorded on this computer in the Operational log located in **Applications and Services Log\Microsoft\Windows\NTLM**. Using an audit event collection system can help you collect the events for analysis more efficiently.

There are no Security Audit Event policies that can be configured to view output from this policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

NTLM and NTLMv2 authentication is vulnerable to a variety of malicious attacks, including SMB replay, man-in-the-middle attacks, and brute force attacks. Reducing and eliminating NTLM authentication from your environment forces the Windows operating system to use more secure protocols, such as the Kerberos version 5 protocol, or different authentication mechanisms, such as smart cards.

Vulnerability

Enabling this policy setting will reveal through logging which computers within your network or domain handle NTLM traffic. The identity of these computers can be used in malicious ways if NTLM authentication traffic is compromised. The policy setting does not prevent or mitigate any vulnerability because it is for audit purposes only.

Countermeasure

Restrict access to the log files when this policy setting is enabled in your production environment.

Potential impact

If you do not enable or configure this policy setting, no NTLM authentication traffic information will be logged. If you do enable this policy setting, only auditing functions will occur; no security enhancements will be implemented.

Recovery console: Allow automatic administrative logon

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting,

Reference

This policy setting determines whether the Administrator account password must be provided before access to the computer is granted. If you enable this setting, the Administrator account is automatically logged on to the computer at the Recovery Console; no password is required.

The Recovery Console can be very useful when troubleshooting and repairing systems that cannot be restarted. However, enabling this policy setting so a user can automatically log on to the console is dangerous. Anyone can walk up to the server, shut it down by disconnecting the power, reboot it, select **Recovery Console** from the **Restart** menu, and then assume full control of the server.

Possible values

- Enabled

  The Administrator account is automatically logged on to the computer at the Recovery Console; no password is required
- Disabled

  Automatic administrative logon is not allowed.
- Not defined

  Automatic administrative logon is not allowed.

Best practices

1. Set **Recovery Console: Allow automatic administrative logon** to **Disabled**. This requires a user to enter a user name and password to access the Recovery Console account.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are differences in this policy between Windows operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer.

Policy conflicts

None.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The Recovery Console can be very useful when you must troubleshoot and repair computers that do not start. However, allowing automatic logon to the Recovery Console can make it possible for someone to assume full control of the server.

Countermeasure

Disable the **Recovery console: Allow automatic administrative logon** setting.

Potential impact

Users must enter a user name and password to access the Recovery Console.

Recovery console: Allow floppy copy and access to all drives and folders

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enables or disables the Recovery Console SET command, which allows you to set the following Recovery Console environment variables.

- **AllowWildCards**. Enables wildcard support for some commands, such as the DEL command.
- **AllowAllPaths**. Allows access to all files and folders on the computer.
- **AllowRemovableMedia**. Allows files to be copied to removable media, such as a floppy disk.
- **NoCopyPrompt**. Suppresses the prompt that typically displays before an existing file is overwritten.

An authorized administrator might forget to remove removable media, such as CD or floppy disk, with sensitive data or applications that a malicious user could then steal. Or an authorized administrator could accidentally leave a startup disk in the computer after using the Recovery Console. If the computer is restarted for any reason and the BIOS has been configured to boot from the removable media before the hard disk drive, the server will start from the removable disk. This causes the server's network services to be unavailable.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set **Recovery Console: Allow floppy copy and access to drives and folders** to **Disabled**. Users who have started a server by using the Recovery Console and logged in with the built-in Administrator account will not be able to copy files and folders to a floppy disk.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Setting and deploying this policy using Group Policy takes precedence over the setting on the local computer.

Policy conflicts

None.

Command-line tools

Enabling this security option makes the Recovery Console SET command available, which allows you to set the following Recovery Console environment variables:

- AllowWildCards: Enable wildcard support for some commands (such as the DEL command).
- AllowAllPaths: Allow access to all files and folders on the computer.
- AllowRemovableMedia: Allow files to be copied to removable media, such as a floppy disk.
- NoCopyPrompt: Do not prompt when overwriting an existing file.

For more information about this command, see **Set**.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An attacker who can cause the system to restart into the Recovery Console could steal sensitive data and leave no audit or access trail.

Countermeasure

Disable the **Recovery console: Allow floppy copy and access to drives and folders** setting.

Potential impact

Users who have started a server through the Recovery Console and logged in with the built-in Administrator account cannot copy files and folders to a floppy disk.

Shutdown: Allow system to be shut down without having to log on

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines whether a computer can be shut down without having to log on to the Windows operating system. If you enable this policy setting, the **Shut Down** option is available on the logon screen in the Windows operating system. If you disable this policy setting, the **Shut Down** option is removed from the logon screen. This configuration requires that users are able to log on to the computer successfully and that they have the **Shut down the system** user right before they can perform a computer shutdown.

Users who can access the console locally can shut down the system. Attackers or misguided users can connect to the server by using Terminal Services or Remote Desktop Services, and then shut it down or restart it without having to identify themselves. A malicious user might also cause a temporary denial-of-service condition by walking up to the local console and restarting the server, or shutting down the server and thus rendering unavailable all its applications and services.

Possible values

- Enabled

  The Shut Down command is available on the logon screen in the Windows operating system.
- Disabled

  The Shut Down option is removed from the logon screen and users must have the **Shut down the system** user right before they can perform a computer shutdown.
- Not defined

Best practices

1. On servers, set this policy to **Disabled**. Administrators must log on to servers to shut them down or restart them.
2. On client computers, set this policy to **Enabled** and define the list of those with the right to shut them down or restart them with the User Rights Assignment policy **Shut down the system**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008 and Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

For information about the User Rights Assignment policy, **Shut down the system**, see the Shut down the system reference topic.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can access the console locally could shut down the computer.

Attackers who have access to the local console could restart the server, which would cause a temporary DoS condition. Attackers could also shut down the server and leave all of its applications and services unavailable.

Countermeasure

Disable the **Shutdown: Allow system to be shut down without having to log on** setting.

Potential impact

Operators must log on to servers to shut them down or restart them.

Shutdown: Clear virtual memory pagefile

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines whether the virtual memory paging file is cleared when the computer is shut down. Virtual memory support uses a system paging file to swap pages of memory to disk when they are not used. On a running computer, this paging file is opened exclusively by the operating system, and it is well protected. However, computers that are configured to allow other operating systems to start should verify that the system paging file is cleared as the computer shuts down. This confirmation ensures that sensitive information from process memory that might be placed in the paging file is not available to an unauthorized user who manages to directly access the paging file after shutdown.

Important information that is kept in real memory might be written periodically to the paging file. This helps computers running at least the Windows Server 2003 operating system handle multitasking functions. A malicious user who has physical access to a server that has been shut down can view the contents of the paging file. The attacker can move the system volume into a different computer and then analyze the contents of the paging file. This is a time-consuming process, but it can expose data that is cached from RAM to the paging file. A malicious user who has physical access to the server can bypass this countermeasure by simply unplugging the server from its power source.

Possible values

- Enabled

  The system paging file is cleared when the system shuts down normally. Also, this policy setting forces the computer to clear the hibernation file (hiberfil.sys) when hibernation is disabled on a portable computer.
- Disabled
- Not defined

Best practices

1. Set this policy to **Enabled**. This causes the Windows Server operating system to clear the paging file when the system is shut down. Depending on the size of the paging file, this process might take several minutes before the system completely shuts down. This delay in shutting down the server is especially noticeable on servers with large paging files. For a server with 2 gigabytes (GB) of RAM and a 2-GB paging file, this setting can add more than 30 minutes to the shutdown process. For some organizations, this downtime violates their internal service level agreements. Use caution when implementing this countermeasure in your environment.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Important information that is kept in real memory may be written periodically to the paging file to help the operating system handle multitasking functions. An attacker who has physical access to a server that has been shut down could view the contents of the paging file. The attacker could move the system volume into a different computer and then analyze the contents of the paging file. Although this process is time consuming, it could expose data that is cached from random access memory (RAM) to the paging file.

Caution

An attacker who has physical access to the computer could bypass this countermeasure by unplugging the computer from its power source.

Countermeasure

Enable the **Shutdown: Clear virtual memory page file when system shuts down** setting. This configuration causes the operating system to clear the paging file when the computer is shut down. The amount of time that is required to complete this process depends on the size of the page file. Because the process overwrites the storage area that is used by the page file several times, it could be several minutes before the computer completely shuts down.

Potential impact

It takes longer to shut down and restart the computer, especially on computers with large paging files. For a computer with 2 gigabytes (GB) of RAM and a 2-GB paging file, this policy setting could increase the shutdown process by more than 30 minutes. For some organizations this downtime violates their internal service level agreements. Therefore, use caution before you implement this countermeasure in your environment.

System cryptography: Force strong key protection for user keys stored on the computer - Explain text

---

Settings page: Options

User input is not required when keys are stored and used

User is prompted when the key is first used

User must enter a password each time they use a key

Explain text

System Cryptography: Force strong key protection for user keys stored on the computer

This security setting determines if users' private keys require a password to be used.

The options are:

User input is not required when new keys are stored and used

User is prompted when the key is first used

User must enter a password each time they use a key

For more information, see Public key infrastructure.

Default: This policy is not defined.

System cryptography: Force strong key protection for user keys stored on the computer

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines whether users can use private keys, such as their Secure/Multipurpose Internet Mail Extensions (S/MIME) key, without a password.

Configuring this policy setting so that users must provide a password every time they use a key (in addition to their domain password) makes it more difficult for a malicious user to access locally-stored user keys, even if the attacker takes control of the user's computer and determines their logon password.

Possible values

- **User input is not required when new keys are stored and used**
- **User is prompted when the key is first used**
- **User must enter a password each time they use a key**
- Not defined

Best practices

1. Set this policy to **User must enter a password each time they use a key**. Users must enter their password every time they access a key that is stored on their computer. For example, if users use an S/MIME certificate to digitally sign their email, they will be forced to enter the password for that certificate every time they send a signed email message. For some organizations, the overhead that is caused by using this value might be too high, but they should set the value at a minimum to **User is prompted when the key is first used**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
DC Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If a user's account is compromised or the user's computer is inadvertently left unsecured, the malicious user can use the keys that are stored for the user to access protected resources.

Countermeasure

Configure the **System cryptography: Force strong key protection for user keys stored on the computer** setting to **User must enter a password each time they use a key** so that users must provide a password that is distinct from their domain password every time they use a key. This configuration makes it more difficult for an attacker to access locally stored user keys, even if the attacker takes control of the user's computer and determines the logon password.

Potential impact

Users must type their password every time they access a key that is stored on their computer. For example, if users use an S/MIME certificate to digitally sign their email, they are forced to type the password for that certificate every time they send a signed email message. For some organizations, the overhead that is involved by using this configuration may be too high. At a minimum, this setting should be set to **User is prompted when the key is first used**.

System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

The Federal Information Processing Standard (FIPS) 140 is a security implementation that is designed for certifying cryptographic software. Windows implements these certified algorithms to meet the requirements and standards for cryptographic modules for use by departments and agencies of the United States federal government. For information about FIPS, see FIPS General Information.

TLS/SSL

This policy setting determines whether the TLS/SSL security provider supports only the FIPS-compliant strong cipher suite known as TLS_RSA_WITH_3DES_EDE_CBC_SHA, which means that the provider only supports the TLS protocol as a client computer and as a server, if applicable. It uses only the Triple Data Encryption Standard (3DES) encryption algorithm for the TLS traffic encryption, only the Rivest-Shamir-Adleman (RSA) public key algorithm for the TLS key exchange and authentication, and only the Secure Hash Algorithm version 1 (SHA-1) hashing algorithm for the TLS hashing requirements.

Encrypting File System (EFS)

For the EFS service, this policy setting supports the 3DES and Advanced Encryption Standard (AES) encryption algorithms for encrypting file data supported by the NTFS file system. To encrypt file data, by default EFS uses the Advanced Encryption Standard (AES) algorithm with a 256-bit key in the Windows Server 2003 and Windows Vista operating systems, and it uses a DESX algorithm in Windows XP.

Remote Desktop Services (RDS)

For encrypting Remote Desktop Services network communication, this policy setting supports only the Triple DES encryption algorithm.

BitLocker

For BitLocker, this policy setting needs to be enabled before any encryption key is generated.

Recovery passwords created on Windows Server 2012 R2 and Windows 8.1 when this policy is enabled are incompatible with BitLocker on operating systems prior to Windows Server 2012 R2 and Windows 8.1; BitLocker will prevent the creation or use of recovery passwords on these systems, so recovery keys should be used instead.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. For use with TLS, set this policy to **Enabled**. Client computers with this policy setting enabled will be unable to communicate through digitally encrypted or signed protocols with servers that do not support these algorithms. Client computers that are connected to the network and do not support these algorithms cannot use servers that require the algorithms for network communications. If you enable this policy setting, you must also configure Internet Explorer to use TLS.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

When this setting is enabled, the Encrypting File System (EFS) service supports only the Triple DES encryption algorithm for encrypting file data. By default, the Windows Vista and the Windows Server 2003 implementation of EFS uses the Advanced Encryption Standard (AES) with a 256-bit key. The Windows XP implementation uses DESX.

When this setting is enabled, BitLocker generates recovery password or recovery keys applicable to versions listed in the following:

Operating systems	Applicability
Windows Server 2012 R2 and Windows 8.1	When created on these operating systems, the recovery password cannot be used on other systems listed in this table.
Windows Server 2012 and Windows 8	When created on these operating systems, the recovery key can be used on other systems listed in this table as well.
Windows Server 2008 R2 and Windows 7	When created on these operating systems, the recovery key can be used on other systems listed in this table as well.
Windows Server 2008 and Windows Vista	When created on these operating systems, the recovery key can be used on other systems listed in this table as well.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Note

This setting is configured differently than in Windows Server 2003 and Windows XP in computers that are running at least Windows Vista or Windows Server 2008.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

You can enable this policy setting to ensure that the computer uses the most powerful algorithms that are available for digital encryption, hashing, and signing. Use of these algorithms minimize the risk of compromise of digitally encrypted or signed data by an unauthorized user.

Countermeasure

Enable the **System cryptography: Use FIPS compliant algorithms for encryption, hashing, and signing** setting.

Potential impact

Client computers that have this policy setting enabled cannot communicate by means of digitally encrypted or signed protocols with servers that do not support these algorithms. Network clients that do not support these algorithms cannot use servers that require them for network communications. For example, many Apache-based Web servers are not configured to support TLS. If you enable this setting, you must also configure Internet Explorer® to use TLS. This policy setting also affects the encryption level that is used for the Remote Desktop Protocol (RDP). The Remote Desktop Connection tool uses the RDP protocol to communicate with servers that run Terminal Services and client computers that are configured for remote control; RDP connections fail if both computers are not configured to use the same encryption algorithms.

To enable Internet Explorer to use TLS

1. On the Internet Explorer **Tools** menu, click **Internet Options**.
2. Click the **Advanced** tab.
3. Select the appropriate version for the **Use TLS** check box.

You can also configure this policy setting through Group Policy or by using the Internet Explorer Administrators Kit.

System objects: Require case insensitivity for non-Windows subsystems

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting

Reference

This policy setting determines whether case insensitivity is enforced for all subsystems. The Microsoft Win32 subsystem is not case sensitive; however, the kernel supports case sensitivity for other subsystems, such as Portable Operating System Interface for UNIX (POSIX). Enabling this policy setting enforces case insensitivity for all directory objects, symbolic links, and input/output (I/O) objects, including file objects. Disabling this policy setting does not allow the Win32 subsystem to become case sensitive.

Because Windows is case insensitive but the POSIX subsystem will support case sensitivity, if this policy setting is not enforced, it is possible for a user of that subsystem to create a file with the same name as another file but with a different mix of capital letters. That might confuse users when they try to access these files by using normal Win32 tools, because only one of the files will be available.

Possible values

- Enabled

  Case insensitivity is enforced for all directory objects, symbolic links, and IO objects, including file objects.
- Disabled

  Will not allow the Win32 subsystem to become case sensitive.
- Not defined

Best practices

1. Set this policy to **Enabled**. All subsystems will be forced to observe case insensitivity. However, this might confuse users who are familiar with one of the UNIX-based operating systems and are used to a case sensitive operating system.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy on operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Because Windows is case insensitive but the POSIX subsystem supports case sensitivity, failure to enable this policy setting makes it possible for a user of that subsystem to create a file with the same name as another file but with a different mix of uppercase and lowercase letters. Such a situation could potentially confuse users when they try to access such files from normal Win32 tools because only one of the files is available.

Countermeasure

Enable the **System objects: Require case insensitivity for non-Windows subsystems** setting.

Potential impact

All subsystems are forced to observe case insensitivity. This configuration may confuse users who are familiar with any UNIX-based operating systems that are case sensitive.

System objects: Strengthen default permissions of internal system objects (e.g. Symbolic Links)

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the strength of the default discretionary access control list (DACL) for objects. Windows Server 2003 maintains a global list of shared system resources such as MS-DOS device names, mutexes, and semaphores. By using this list, processes can locate and share objects. Each type of object is created with a default DACL that specifies who can access the objects with what permissions. Enabling this policy setting strengthens the default DACL and allows users who are not administrators to read, but not to modify, shared objects that they did not create.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. It is advisable to set this policy to **Enabled**.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\ Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This policy setting is enabled by default to protect against a known vulnerability that can be used with hard links or symbolic links. Hard links are actual directory entries in the file system. With hard links, the same data in a file system can be referred to by different file names. Symbolic links are text files that provide a pointer to the file that is interpreted and followed by the operating system as a path to another file or directory. Because symbolic links are a separate file, they can exist independently of the target location. If a symbolic link is deleted, its target location remains unaffected. When this setting is disabled, it is possible for a malicious user to destroy a data file by creating a link that looks like a temporary file that the system automatically creates, such as a sequentially named log file, but it points to the data file that the malicious user wants to eradicate. When the system writes the files with that name, the data is overwritten. Enabling **System objects: Strengthen default permissions of internal system objects (e.g., Symbolic Links)** prevents an attacker from exploiting programs that create files with predictable names by not allowing them to write to objects that they did not create.

Countermeasure

Enable the **System objects: Strengthen default permissions of global system objects (for example, Symbolic Links)** setting.

Potential impact

None. This is the default configuration.

System settings: Optional subsystems

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines which subsystems support your applications. You can use this security setting to specify as many subsystems as your environment demands.

The subsystem introduces a security risk that is related to processes that can potentially persist across logons. If a user starts a process and then logs out, the next user who logs on to the system might access the process that the previous user started. This is dangerous, because the process started by the first user can retain that user's system user rights; therefore, anything that the second user does using that process is performed with the user rights of the first user. This makes it difficult to trace who creates processes and objects, which is essential for post-security incident forensics.

Possible values

- User-defined list of subsystems
- Not defined

Best practices

1. Set this policy setting to a null value. The default value is **POSIX**, so applications that rely on the POSIX subsystem will no longer run. For example, Microsoft Services for UNIX 3.0 installs an updated version of the POSIX subsystem. Reset this policy setting in Group Policy for any servers that use Services for UNIX 3.0.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	POSIX
DC Effective Default Settings	POSIX
Member Server Effective Default Settings	POSIX
Client Computer Effective Default Settings	POSIX

Operating system version differences

There are no differences for this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The POSIX subsystem is an Institute of Electrical and Electronic Engineers (IEEE) standard that defines a set of operating system services. The POSIX subsystem is required if the server supports applications that use that subsystem.

The POSIX subsystem introduces a security risk that relates to processes that can potentially persist across logons. If a user starts a process and then logs out, there is a potential that the next user who logs on to the computer could access the previous user's process. This would allow the second user to take actions on the process by using the privileges of the first user.

Countermeasure

Configure the **System settings: Optional subsystems setting** to a null value. The default value is POSIX.

Potential impact

Applications that rely on the POSIX subsystem no longer operate. For example, Microsoft Services for UNIX (SFU) installs an updated version of the POSIX subsystem that is required, so you must reconfigure this setting in Group Policy for any servers that use SFU.

System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines whether digital certificates are processed when software restriction policies are enabled and a user or process attempts to run software with an .exe file name extension. This security setting enables or disables certificate rules (which are a type of software restriction policy). With a software restriction policy, you can create a certificate rule that allows or disallows Microsoft Authenticode®-signed software to run, based on the digital certificate that is associated with the software. For certificate rules to work in software restriction policies, you must enable this security setting.

Possible values

- Enabled
- Disabled
- Not defined

Best practices

1. Set this policy to **Enabled**. Enabling certificate rules results in software restriction policies checking a certificate revocation list (CRL) to make sure that the software's certificate and signature are valid. When you start signed programs, this setting can decrease system performance. You can disable CRLs by editing the software restriction policies in the desired GPO. In the **Trusted Publishers Properties** dialog box, clear the **Publisher** and **Timestamp** check boxes.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2003.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Without the use of software restriction policies, users and computers might be exposed to unauthorized software that could include malicious software such as viruses and Trojan horses.

Countermeasure

Enable the **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies** setting.

Potential impact

If you enable certificate rules, software restriction policies check a certificate revocation list (CRL) to verify that the software's certificate and signature are valid. This checking process may negatively affect performance when signed programs start. To disable this feature, you can edit the software restriction policies in the appropriate GPO. In the **Trusted Publishers Properties** dialog box, clear the **Publisher** and **Timestamp** check boxes.

User Account Control: Admin Approval Mode for the Built-in Administrator account

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the behavior of Admin Approval Mode for the built-in Administrator account.

When the Admin Approval Mode is enabled, the local Administrator account functions like a standard user account, but it has the ability to elevate privileges without logging on by using a different account. In this mode, any operation that requires elevation of privilege displays a prompt that allows the administrator to permit or deny the elevation of privilege. If Admin Approval Mode is not enabled, the built-in Administrator account logs on in Windows XP Mode, and it runs all applications by default with full administrative privileges. By default, this setting is set to **Disabled**.

Note
If a computer is upgraded from a previous version of the Windows operating system, and the Administrator account is the only account on the computer, the built-in Administrator account remains enabled, and this setting is also enabled.

Possible values

- Enabled

  The built-in Administrator account logs on in Admin Approval Mode so that any operation that requires elevation of privilege displays a prompt that provides the administrator the option to permit or deny the elevation of privilege.
- Disabled

  The built-in Administrator account logs on in Windows XP Mode, and it runs all applications by default with full administrative privileges.

Best practices

1. Do not enable the built-in Administrator account on the client computer, but use the Standard User account and User Account Control (UAC).

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

One of the risks of the User Account Control (UAC) feature that is included in computers running at least Windows Vista or Windows Server 2008 is that it is intended to mitigate malicious software running under elevated credentials without the user or administrator being aware of its activity. An attack vector for malicious programs is to discover the password of the Administrator account because that user account was created for all installations of the Windows operating system. To address this risk, the built-in Administrator account is disabled in computers running

at least Windows Vista. In computers running at least Windows Server 2008, the Administrator account is enabled, and the password must be changed the first time the Administrator logs on. In a default installation of a computer running at least Windows Vista, accounts with administrative control over the computer are initially set up in one of two ways:

- If the computer is not joined to a domain, the first user account you create has the equivalent permissions as a local administrator.
- If the computer is joined to a domain, no local administrator accounts are created. The enterprise or domain administrator must log on to the computer and create a local administrator account if one is warranted.

After Windows Vista or Windows 7 is installed, you can enable the built-in Administrator account, but we strongly recommend that this account remain disabled.

Countermeasure

Enable the **User Account Control: Admin Approval Mode for the Built-in Administrator account** setting if you have the built-in Administrator account enabled.

Potential impact

Users who log on by using the local Administrator account are prompted for consent whenever a program requests an elevation in privilege.

User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop

This security policy reference topic for the IT professional describes the best practices, location, values, and security considerations for this policy setting.

Reference

This security setting controls whether User Interface Accessibility (UIAccess or UIA) programs can automatically disable the secure desktop for elevation prompts that are used by a standard user.

Note
This setting does not change the behavior of the UAC elevation prompt for administrators.

Background

User Interface Privilege Isolation (UIPI) implements restrictions in the Windows subsystem that prevent lower-privilege applications from sending messages or installing hooks in higher-privilege

processes. Higher-privilege applications are permitted to send messages to lower-privilege processes. UIPI does not interfere with or change the behavior of messages between applications at the same privilege (or integrity) level.

Microsoft UI Automation is the current model to support accessibility requirements in the Windows operating systems. Applications that are designed to support an accessible user experience control the behavior of other Windows applications on behalf of the user. When all applications on the automation client computer and server are running as a standard user (that is, at a medium integrity level), the UIPI restrictions do not interfere with the Microsoft UI automation model.

However, there might be times when an administrative user runs an application with elevated privilege based on UAC in Admin Approval Mode. Microsoft UI Automation cannot drive the UI graphics of elevated applications on the desktop without the ability to bypass the restrictions that UIPI implements. The ability to bypass UIPI restrictions across privilege levels is available for UI automation programs by using UIAccess.

If an application presents a UIAccess attribute when it requests privileges, the application is stating a requirement to bypass UIPI restrictions for sending messages across privilege levels. Computers running at least the Windows Vista operating system implement the following policy checks before starting an application with UIAccess privilege.

1. The application must have a digital signature that can be verified by using a digital certificate that is associated with the Trusted Root Certification Authorities store on the local computer.
2. The application must be installed in a local folder that is writeable only by administrators, such as the Program Files directory. The allowed directories for UI automation applications are:
    a. %ProgramFiles% and its subdirectories.
    b. %WinDir% and its subdirectories, except a few subdirectories that are excluded because standard users have write access.

Resulting behavior

When this setting is enabled, UIAccess programs (including Windows Remote Assistance) can automatically disable the secure desktop for elevation prompts. Unless you have also disabled elevation prompts, the prompts appear on the interactive user's desktop instead of on the secure desktop. The prompts also appear on the remote administrator's view of the desktop during a Windows Remote Assistance session, and the remote administrator can provide the appropriate credentials for elevation.

If you disable this setting, the secure desktop can only be disabled by the user of the interactive desktop or by disabling the **User Account Control: Switch to the secure desktop when prompting for elevation** setting, which by default is enabled.

Possible values

- Enabled

  UIA programs can automatically disable the secure desktop for elevation prompts, and unless you have also disabled elevation prompts, the prompts appear on the interactive user's

desktop instead of on the secure desktop. Prompts will also appear on the remote administrator's view of the desktop during a Windows Remote Assistance session, and the remote administrator can provide the appropriate credentials for elevation.

- Disabled

  The secure desktop can be disabled only by the user of the interactive desktop or by disabling the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting.

Best practices

1. Best practices are dependent on your security policies and your remote operational requirements.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

This policy has no impact in Windows operating systems earlier than Windows Vista.

There are no differences in this security policy between operating systems beginning with Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

This policy has no impact in Windows operating systems earlier than Windows Vista.

Policy interactions

If you plan to enable this setting, you should also review the effect of the **User Account Control: Behavior of the elevation prompt for standard users** setting. If it is configured as **Automatically deny elevation requests**, elevation requests are not presented to the user. If you disable this setting, the secure desktop can only be disabled by the user of the interactive desktop or by disabling the **User Account Control: Switch to the secure desktop when prompting for elevation** setting, which by default is enabled.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

UIA programs are designed to interact with Windows and application programs on behalf of a user. This setting allows UIA programs to bypass the secure desktop to increase usability in certain cases, but it allows elevation requests to appear on the regular interactive desktop instead of on the secure desktop. This increases the risk that a malicious program could intercept data that is being transferred between the UI and the application. Because UIA programs must be able to respond to prompts regarding security issues, such as the UAC elevation prompt, UIA programs must be highly trusted. To be considered trusted, a UIA program must be digitally signed. By default, UIA programs can be run only from the following protected paths:

- ..\Program Files\ (and subfolders)
- ..\Program Files (x86)\ (and subfolders, in 64-bit versions of Windows only)
- ..\Windows\System32\

The requirement to be in a protected path can be disabled by the **User Account Control: Only elevate UIAccess applications that are installed in secure locations** setting. Although this setting applies to any UIA program, it is used primarily in certain Windows Remote Assistance scenarios. The Windows Remote Assistance program in Windows Vista and Windows 7 is a UIA program.

Countermeasure

Disable the **User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop** setting.

Potential impact

If a user requests remote assistance from an administrator and the remote assistance session is established, elevation prompts appear on the interactive user's secure desktop and the administrator's remote session is paused. To avoid pausing the remote administrator's session during elevation requests, the user can select the "Allow IT Expert to respond to User Account Control prompts" check box when setting up the remote assistance session. However, selecting this check box requires that the interactive user respond to an elevation prompt on the secure desktop. If the interactive user is a standard user, the user does not have the required credentials to allow elevation.

User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the behavior of the elevation prompt for accounts that have administrative credentials.

Possible values

- Elevate without prompting

  Assumes that the administrator will permit an operation that requires elevation, and additional consent or credentials are not required.

  Note
  Selecting **Elevate without prompting** minimizes the protection that is provided by the UAC feature. We do not recommend selecting this value unless administrator accounts are tightly controlled and the operating environment is highly secure.

- Prompt for credentials on the secure desktop

  When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a privileged user name and password. If the user enters valid credentials, the operation continues with the user's highest available privilege.

- Prompt for consent on the secure desktop

  When an operation requires elevation of privilege, the user is prompted on the secure desktop to select **Permit** or **Deny**. If the user selects **Permit**, the operation continues with the user's highest available privilege.
- Prompt for credentials

  An operation that requires elevation of privilege prompts the administrator to type the user name and password. If the administrator enters valid credentials, the operation continues with the applicable privilege.
- Prompt for consent

  An operation that requires elevation of privilege prompts the administrator to select **Permit** or **Deny**. If the administrator selects **Permit**, the operation continues with the administrator's highest available privilege.
- Prompt for consent for non-Windows binaries

  This is the default. When an operation for a non-Microsoft application requires elevation of privilege, the user is prompted on the secure desktop to select **Permit** or **Deny**. If the user selects **Permit**, the operation continues with the user's highest available privilege.

Best practices

1. Selecting the option **Elevate without prompting** minimizes the protection that is provided by the UAC feature. We do not recommend selecting this value unless administrator accounts are tightly controlled and the operating environment is highly secure.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Prompt for consent for non-Windows binaries
DC Effective Default Settings	Prompt for consent for non-Windows binaries
Member Server Effective Default Settings	Prompt for consent for non-Windows binaries
Client Computer Effective Default Settings	Prompt for consent for non-Windows binaries

Operating system version differences

The default in Windows Vista and Windows Server 2008 was **Prompt for consent**. In Windows Server 2008 R2 and Windows 7, the default was changed and three options were added:

- **Prompt for credentials on the secure desktop**
- **Prompt for consent on the secure desktop**
- **Prompt for consent for non-Windows binaries** (default)

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distributing this policy through Group Policy to computers running at least the Windows Vista or Windows Server 2008 operating systems will have no impact (because UAC was first introduced in Windows Vista).

Distributing this policy to computers running Windows Vista and Windows Server 2008 will change the default option to **Prompt for consent for non-Windows binaries**. You will need to test that the resultant behavior is as you expect.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

One of the risks that the UAC feature tries to mitigate is that of malicious software running under elevated credentials without the user or administrator being aware of its activity. This setting raises awareness to the administrator of elevated privilege operations, and it permits the administrator to prevent a malicious program from elevating its privilege when the program attempts to do so.

Countermeasure

Configure the **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** setting to **Prompt for consent**.

Potential impact

Administrators should be made aware that they will be prompted for consent when all binaries attempt to run.

User Account Control: Behavior of the elevation prompt for standard users

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the behavior of the elevation prompt for standard users.

Possible values

- **Automatically deny elevation requests**

  This option returns an “Access denied” error message to standard users when they try to perform an operation that requires elevation of privilege. Most enterprises that run desktops as standard users configure this policy to reduce Help Desk calls.

- **Prompt for credentials on the secure desktop**

  This is the default. When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a different user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.

- **Prompt for credentials**

  An operation that requires elevation of privilege prompts the user to type an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.

Best practices

1. Configure the **User Account Control: Behavior of the elevation prompt for standard users** to **Automatically deny elevation requests**. This setting requires the user to log on with an administrative account to run programs that require elevation of privilege.
2. As a security best practice, standard users should not have knowledge of administrative passwords. However, if your users have both standard and administrator-level accounts, set

Prompt for credentials so that the users do not choose to always log on with their administrator accounts, and they shift their behavior to use the standard user account.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Prompt for credentials on the secure desktop
DC Effective Default Settings	Prompt for credentials on the secure desktop
Member Server Effective Default Settings	Prompt for credentials on the secure desktop
Client Computer Effective Default Settings	Prompt for credentials on the secure desktop

Operating system version differences

The default in Windows Vista and Windows Server 2008 is **Prompt for credentials**. In Windows Server 2008 R2 and Windows 7, the default is changed to **Prompt for credentials on the secure desktop**.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distributing this policy through Group Policy to operating systems earlier than Windows Vista and Windows Server 2008 will have no impact (because UAC was first introduced in Windows Vista).

Distributing this policy to computers running Windows Vista and Windows Server 2008 will change the default option to **Prompt for credentials on the secure desktop**. You will need to test that the resultant behavior is as you expect.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

One of the risks that the UAC feature tries to mitigate is that of malicious programs running under elevated credentials without the user or administrator being aware of their activity. This setting raises awareness to the user that a program requires the use of elevated privilege operations, and it requires that the user supply administrative credentials for the program to run.

Countermeasure

Configure the **User Account Control: Behavior of the elevation prompt for standard users** to **Automatically deny elevation requests**. This setting requires the user to log on with an administrative account to run programs that require elevation of privilege. As a security best practice, standard users should not have knowledge of administrative passwords. However, if your users have both standard and administrator-level accounts, we recommend setting **Prompt for credentials** so that the users do not choose to always log on with their administrator accounts,and they shift their behavior to use the standard user account.

Potential impact

Users must provide administrative passwords to run programs with elevated privileges. This could cause an increased load on IT staff while the programs that are affected are identified and standard operating procedures are modified to support least privilege operations.

User Account Control: Detect application installations and prompt for elevation

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the behavior of application installation detection for the entire system.

Some software might attempt to install itself after being given permission to run. The user may give permission for the program to run because the program is trusted. Then the user is prompted to install an unknown component. This security policy provides another way to identify and stop these attempted software installations before they can do damage.

Possible values

- **Enabled**

  Application installation packages that require an elevation of privilege to install are detected and the user is prompted for administrative credentials.

- **Disabled**

  Application installation packages that require an elevation of privilege to install are not detected and the user is not prompted for administrative credentials.

Best practices

1. Installer detection is unnecessary when enterprises run standard user desktops that capitalize on delegated installation technologies like Group Policy Software Install (GPSI) or SMS. Therefore you can set this security policy to **Disabled**.
2. Enable the **User Account Control: Detect application installations and prompt for elevation** setting so standard users must provide administrative credentials before software is installed.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled

Server type or GPO	Default value
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Vista and Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Some malicious software might attempt to install itself after being given permission to run, for example, malicious software with a trusted application shell. The user may give permission for the program to run because the program is trusted. Then the user is prompted to install an unknown component. This policy provides another way to trap the software before it can do damage.

Countermeasure

Enable the **User Account Control: Detect application installations and prompt for elevation** setting.

Potential impact

Users must provide administrative passwords to install programs.

User Account Control: Only elevate executables that are signed and validated

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enforces public key infrastructure (PKI) signature checks on any interactive application that requests elevation of privilege. Enterprise administrators can control the applications that are allowed to run through the population of certificates in the local computer's Trusted Publishers store.

A trusted publisher is a certificate issuer that the computer's user has chosen to trust and that has certificate details that have been added to the store of trusted publishers.

The Windows operating system maintains certificates in certificate stores. These stores can be represented by containers in the file system or the registry, or they can be implemented as physical stores such as smart cards. Certificate stores are associated with the computer object or they are owned by a distinct user who has a security context and profile on that computer. In addition, services can have certificate stores. A certificate store will often contain numerous certificates, possibly issued from a number of different certification authorities (CAs).

When certificate path discovery is initiated, the Windows operating system attempts to locate the issuing CA for the certificates, and it builds a certificate path to the trusted root certificate. Intermediate certificates are included as part of the application protocol or are picked up from Group Policy or through URLs that are specified in the Authority Information Access (AIA) extension. When the path is built, each certificate in the path is verified for validity with respect to various parameters, such as name, time, signature, revocation status, and other constraints.

Possible values

- Enabled

  Enforces the PKI certificate chain validation of a given executable file before it is permitted to run.
- Disabled

  Does not enforce PKI certificate chain validation before a given executable file is permitted to run.

Best practices

1. Best practices are dependent on your security and performance goals.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Disabled
DC Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Vista and Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

Distributing this policy through Group Policy to operating systems earlier than Windows Vista and Windows Server 2008 will have no impact (because UAC was first introduced in Windows Vista).

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Intellectual property, personally identifiable information, and other confidential data are normally manipulated by applications on the computer, and elevated credentials are required to access the information. Users and administrators inherently trust applications that are used with these information sources, and they provide their credentials. If one of these applications is replaced by a rogue application that appears identical to the trusted application, the confidential data could be compromised and the user's administrative credentials would also be compromised.

Countermeasure

Enable the **User Account Control: Only elevate executables that are signed and validated**.

Potential impact

Enabling this setting requires that you have a PKI infrastructure and that your enterprise administrators have populated the Trusted Publishers store with the certificates for the allowed applications. Some older applications are not signed, and they cannot be used in an environment that is hardened with this setting. You should carefully test your applications in a preproduction environment before implementing this setting.

Control over the applications that are installed on the desktops and the hardware that joins your domain should provide similar protection from the vulnerability that is addressed by this setting. Additionally, the level of protection that is provided by this setting is not an assurance that all rogue applications will be found.

User Account Control: Only elevate UIAccess applications that are installed in secure locations

---

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enforces the requirement that applications that request running with a UIAccess integrity level (by means of a marking of UIAccess=true in their application manifest),

must reside in a secure location on the file system. Relatively secure locations are limited to the following directories:

- \Program Files\ including subdirectories
- \Windows\system32\
- \Program Files (x86)\ including subdirectories for 64-bit versions of Windows

Note
Windows enforces a PKI signature check on any interactive application that requests running with a UIAccess integrity level, regardless of the state of this security setting.

Background

User Interface Privilege Isolation (UIPI) implements restrictions in the Windows subsystem that prevent lower-privilege applications from sending messages or installing hooks in higher-privilege processes. Higher-privilege applications are permitted to send messages to lower-privilege processes. UIPI does not interfere with or change the behavior of messages between applications at the same privilege (or integrity) level.

Microsoft UI Automation is the current model to support accessibility requirements in the Windows operating systems. Applications that are designed to support an accessible user experience control the behavior of other Windows applications on behalf of the user. When all applications on the automation client computer and server are running as a standard user (that is, at a medium integrity level), the UIPI restrictions do not interfere with the Microsoft UI automation model.

However, there might be times when an administrative user runs an application with elevated privilege based on UAC in Admin Approval Mode. Microsoft UI Automation cannot drive the UI graphics of elevated applications on the desktop without the ability to bypass the restrictions that UIPI implements. The ability to bypass UIPI restrictions across privilege levels is available for UI automation programs by using UIAccess.

If an application presents a UIAccess attribute when it requests privileges, the application is stating a requirement to bypass UIPI restrictions for sending messages across privilege levels. Computers running at least the Windows Vista operating system implement the following policy checks before starting an application with UIAccess privilege.

1. The application must have a digital signature that can be verified by using a digital certificate that is associated with the Trusted Root Certification Authorities store on the local computer.
2. The application must be installed in a local folder that is writeable only by administrators, such as the Program Files directory. The allowed directories for UI automation applications are:
    a. %ProgramFiles% and its subdirectories.
    b. %WinDir% and its subdirectories, except a few subdirectories that are excluded because standard users have write access.

Possible values

- **Enabled**

An application can start with UIAccess integrity only if it resides in a secure location in the file system.

- **Disabled**

  An application can start with UIAccess integrity even if it does not reside in a secure location in the file system.

Best practices

1. Set this policy to **Enabled** to permit applications that are located in one of the designated secure directories to run with UIAccess integrity.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this policy between operating systems beginning with Windows Vista and Windows Server 2008.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they aresaved locally or distributed through Group Policy.

Group Policy

This policy has no impact in Windows operating systems earlier than Windows Vista.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

UIAccess integrity allows an application to bypass User Interface Privilege Isolation (UIPI) restrictions when an application is elevated in privilege from a standard user to an administrator. When this setting is enabled, an application that has the UIAccess flag set to true in its manifest can interchange information with applications that are running at a higher privilege level, such as logon prompts and privilege elevation prompts. This ability is required to support accessibility features such as screen readers that are transmitting user interfaces to alternative forms, but it is not required by most applications. A process that is started with UIAccess rights has the following abilities:

- Set the foreground window.
- Drive any application window by using the SendInput function.
- Use read input for all integrity levels by using low-level hooks, raw input, GetKeyState, GetAsyncKeyState, and GetKeyboardInput.
- Set journal hooks.
- Use AttachThreadInput to attach a thread to a higher integrity input queue.

Countermeasure

Enable the **User Account Control: Only elevate UIAccess applications that are installed in secure locations** setting.

Potential impact

If the application that requests UIAccess meets the UIAccess setting requirements, computers running at least the Windows Vista operating system start the application with the ability to bypass most of the UIPI restrictions. If the application does not meet the security restrictions, the application is started without UIAccess rights, and it can interact only with applications at the same or lower privilege level.

User Account Control: Run all administrators in Admin Approval Mode

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines the behavior of all User Account Control (UAC) policies for the entire system. This is the setting that turns UAC on or off.

Possible values

- **Enabled**

  Admin Approval Mode and all other UAC policies are dependent on this option being enabled. Changing this setting requires restarting the system.

- **Disabled**

  Admin Approval Mode and all related UAC policies are disabled.

  Note

  If this security setting is configured to **Disabled**, the Security Center notifies the user that the overall security of the operating system has been reduced.

Best practices

1. Enable this policy to allow all other UAC features and policies to function.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled

Server type or GPO	Default value
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this security policy between operating systems beginning with Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

A restart of the computer is required before this policy will be effective when changes to this policy are saved locally or distributed through Group Policy.

Group Policy

This policy has no impact in Windows operating systems earlier than Windows Vista.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

This is the setting that turns UAC on or off. If this setting is disabled, UAC is not used, and any security benefits and risk mitigations that are dependent on UAC are not present on the computer.

Countermeasure

Enable the **User Account Control: Run all users, including administrators, as standard users** setting.

Potential impact

Users and administrators must learn to work with UAC prompts and adjust their work habits to use least privilege operations.

User Account Control: Switch to the secure desktop when prompting for elevation

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines whether the elevation request prompts on the interactive user desktop or on the secure desktop.

The secure desktop presents the logon UI and restricts functionality and access to the system until the logon requirements are satisfied.

The secure desktop's primary difference from the user desktop is that only trusted processes running as SYSTEM are allowed to run here (that is, nothing is running at the user's privilege level). The path to get to the secure desktop from the user desktop must also be trusted through the entire chain.

Possible values

- **Enabled**

  All elevation requests by default go to the secure desktop.
- **Disabled**

  All elevation requests go to the interactive user desktop.

Best practices

1. Enable the **User Account Control: Switch to the secure desktop when prompting for elevation setting**. The secure desktop helps protect against input and output spoofing by presenting the credentials dialog box in a protected section of memory that is accessible only by trusted system processes.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this security policy between operating systems beginning with Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

This policy has no impact of Windows operating systems earlier than Windows Vista.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Elevation prompt dialog boxes can be spoofed, causing users to disclose their passwords to malicious software. Mouse cursors can be spoofed by hiding the real cursor and replacing it with an offset so the cursor is actually pointing to the **Allow** button.

Countermeasure

Enable the **User Account Control: Switch to the secure desktop when prompting for elevation setting**. The secure desktop helps protect against input and output spoofing by presenting the credentials dialog box in a protected section of memory that is accessible only by trusted system processes.

Potential impact

None. This is the default configuration.

User Account Control: Virtualize file and registry write failures to per-user locations

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting enables or disables the redirection of the write failures of earlier applications to defined locations in the registry and the file system. This feature mitigates applications that historically ran as administrator and wrote runtime application data to %ProgramFiles%, %Windir%, %Windir%\system32, or HKEY_LOCAL_MACHINE\Software\.

Virtualization facilitates running applications based on operating systems earlier than Windows Vista that historically failed to run with standard user privileges. This feature can be disabled for applications on computers running at least the Windows Vista operating system because it is unnecessary.

Possible values

- **Enabled**

  Setting this value facilitates the runtime redirection of application write failures to defined user locations for the file system and the registry.

- **Disabled**

  Applications that write data to protected locations fail, as they did in versions of the Windows operating system earlier than Windows Vista.

Best practices

1. If you run applications that are not Windows Vista-compliant, enable this security policy to prevent the possibility that these older applications could write data to unsecure locations.
2. If you only run at least Windows Vista–compliant applications, this feature is unnecessary so you can disable this policy.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options

Default values

The following table lists the actual and effective default values for this policy. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Enabled
DC Effective Default Settings	Enabled
Member Server Effective Default Settings	Enabled
Client Computer Effective Default Settings	Enabled

Operating system version differences

There are no differences in this security policy between operating systems beginning with Windows Vista.

Policy management

This section describes features and tools that are available to help you manage this policy.

Restart requirement

None. Changes to this policy become effective without a computer restart when they are saved locally or distributed through Group Policy.

Group Policy

This policy has no impact in Windows operating systems earlier than Windows Vista.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Earlier applications might not write data to secure locations.

Countermeasure

Enable the **User Account Control: Virtualize file and registry write failures to per-user locations** setting.

Potential impact

None. This is the default configuration.

Advanced Audit Policy Configuration

This reference for the IT professional provides information about one collection of auditing settings which were introduced in Windows Server 2008 R2 and Windows 7, and the audit events that they generate.

The security audit policy settings under **Security Settings\Advanced Audit Policy Configuration** can help your organization audit compliance with important business-related and security-related rules by tracking precisely defined activities, such as:

- A group administrator has modified settings or data on servers that contain finance information.
- An employee within a defined group has accessed an important file.
- The correct system access control list (SACL) is applied to every file and folder or registry key on a computer or file share as a verifiable safeguard against undetected access.

These settings allow you to select only the behaviors that you want to monitor and exclude audit results for other behaviors. In addition, because security audit policies can be applied by using domain Group Policy, audit policy settings can be modified, tested, and deployed to selected users and groups.

Audit policy settings under **Security Settings\Advanced Audit Policy Configuration** are available in the following categories:

Account Logon

Configuring policy settings in this category can help you document attempts to authenticate account data on a domain controller or on a local Security Accounts Manager (SAM). Unlike Logon and Logoff policy settings and events, which track attempts to access a particular

computer, settings and events in this category focus on the account database that is used. This category includes the following subcategories:

- Audit Credential Validation
- Audit Kerberos Authentication Service
- Audit Kerberos Service Ticket Operations
- Audit Other Logon/Logoff Events

Account Management

The security audit policy settings in this category can be used to monitor changes to user and computer accounts and groups. This category includes the following subcategories:

- Audit Application Group Management
- Audit Computer Account Management
- Audit Distribution Group Management
- Audit Other Account Management Events
- Audit Security Group Management
- Audit User Account Management

Detailed Tracking

Detailed Tracking security policy settings and audit events can be used to monitor the activities of individual applications and users on that computer, and to understand how a computer is being used. This category includes the following subcategories:

- Audit DPAPI Activity
- Audit Process Creation
- Audit Process Termination
- Audit RPC Events

DS Access

DS Access security audit policy settings provide a detailed audit trail of attempts to access and modify objects in Active Directory Domain Services (AD DS). These audit events are logged only on domain controllers. This category includes the following subcategories:

- Audit Detailed Directory Service Replication
- Audit Directory Service Access
- Audit Directory Service Changes
- Audit Directory Service Replication

Logon/Logoff

Logon/Logoff security policy settings and audit events allow you to track attempts to log on to a computer interactively or over a network. These events are particularly useful for tracking user activity and identifying potential attacks on network resources. This category includes the following subcategories:

- Audit Account Lockout
- Audit IPsec Extended Mode
- Audit IPsec Main Mode

- Audit IPsec Quick Mode
- Audit Logoff
- Audit Logon
- Audit Network Policy Server
- Audit Other Logon/Logoff Events
- Audit Special Logon

Object Access

Object Access policy settings and audit events allow you to track attempts to access specific objects or types of objects on a network or computer. To audit attempts to access a file, directory, registry key, or any other object, you must enable the appropriate Object Access auditing subcategory for success and/or failure events. For example, the File System subcategory needs to be enabled to audit file operations, and the Registry subcategory needs to be enabled to audit registry accesses.

Proving that these audit policies are in effect to an external auditor is more difficult. There is no easy way to verify that the proper SACLs are set on all inherited objects. To address this issue, see Global Object Access Auditing.

This category includes the following subcategories:

- Audit Application Generated
- Audit Certification Services
- Audit Detailed File Share
- Audit File Share
- Audit File System
- Audit Filtering Platform Connection
- Audit Filtering Platform Packet Drop
- Audit Handle Manipulation
- Audit Kernel Object
- Audit Other Object Access Events
- Audit Registry
- Audit SAM

Policy Change

Policy Change audit events allow you to track changes to important security policies on a local system or network. Because policies are typically established by administrators to help secure network resources, monitoring changes or attempts to change these policies can be an important aspect of security management for a network. This category includes the following subcategories:

- Audit Audit Policy Change
- Audit Authentication Policy Change
- Audit Authorization Policy Change
- Audit Filtering Platform Policy Change

- Audit MPSSVC Rule-Level Policy Change
- Audit Other Policy Change Events

Privilege Use

Permissions on a network are granted for users or computers to complete defined tasks. Privilege Use security policy settings and audit events allow you to track the use of certain permissions on one or more systems. This category includes the following subcategories:

- Audit Non-Sensitive Privilege Use
- Audit Sensitive Privilege Use
- Audit Other Privilege Use Events

System

System security policy settings and audit events allow you to track system-level changes to a computer that are not included in other categories and that have potential security implications. This category includes the following subcategories:

- Audit IPsec Driver
- Audit Other System Events
- Audit Security State Change
- Audit Security System Extension
- Audit System Integrity

Global Object Access

Global Object Access Auditing policy settings allow administrators to define computer system access control lists (SACLs) per object type for the file system or for the registry. The specified SACL is then automatically applied to every object of that type.

Auditors will be able to prove that every resource in the system is protected by an audit policy by viewing the contents of the Global Object Access Auditing policy settings. For example, if auditors see a policy setting called "Track all changes made by group administrators," they know that this policy is in effect.

Resource SACLs are also useful for diagnostic scenarios. For example, setting the Global Object Access Auditing policy to log all the activity for a specific user and enabling the policy to track "Access denied" events for the file system or registry can help administrators quickly identify which object in a system is denying a user access.

Note

If a file or folder SACL and a Global Object Access Auditing policy setting (or a single registry setting SACL and a Global Object Access Auditing policy setting) are configured on a computer, the effective SACL is derived from combining the file or folder SACL and the Global Object Access Auditing policy. This means that an audit event is generated if an activity matches the file or folder SACL or the Global Object Access Auditing policy.

This category includes the following subcategories:

- File System (Global Object Access Auditing)

- Registry (Global Object Access Auditing)

See Also

Which Editions of Windows Support Advanced Audit Policy Configuration

User Rights Assignment

This reference topic for the IT professional provides an overview and links to information about the User Rights Assignment security policy settings user rights that are available in the Windows operating system.

User rights govern the methods by which a user can log on to a system. User rights are applied at the local computer level, and they allow users to perform tasks on a computer or in a domain. User rights include logon rights and permissions. Logon rights control who is authorized to log on to a computer and how they can log on. User rights permissions control access to computer and domain resources, and they can override permissions that have been set on specific objects. User rights are managed in Group Policy under the **User Rights Assignment** item.

Each user right has a constant name and a Group Policy name associated with it. The constant names are used when referring to the user right in log events. You can configure the user rights assignment settings in the following location within the Group Policy Management Console (GPMC) under **Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment**, or on the local computer by using the Local Group Policy Editor (gpedit.msc).

For information about setting security policies, see How to Configure Security Policy Settings.

The following table links to each security policy setting and provides the constant name for each. Setting descriptions contain reference information, best practices for configuring the policy setting, default values, differences between operating system versions, and considerations for policy management and security.

Group Policy Setting	Constant Name
Access Credential Manager as a trusted caller	SeTrustedCredManAccessPrivilege
Access this computer from the network	SeNetworkLogonRight
Act as part of the operating system	SeTcbPrivilege
Add workstations to domain	SeMachineAccountPrivilege
Adjust memory quotas for a process	SeIncreaseQuotaPrivilege
Allow log on locally	SeInteractiveLogonRight
Allow log on through Remote Desktop Services	SeRemoteInteractiveLogonRight

Group Policy Setting	Constant Name
Back up files and directories	SeBackupPrivilege
Bypass traverse checking	SeChangeNotifyPrivilege
Change the system time	SeSystemtimePrivilege
Change the time zone	SeTimeZonePrivilege
Create a pagefile	SeCreatePagefilePrivilege
Create a token object	SeCreateTokenPrivilege
Create global objects	SeCreateGlobalPrivilege
Create permanent shared objects	SeCreatePermanentPrivilege
Create symbolic links	SeCreateSymbolicLinkPrivilege
Debug programs	SeDebugPrivilege
Deny access to this computer from the network	SeDenyNetworkLogonRight
Deny log on as a batch job	SeDenyBatchLogonRight
Deny log on as a service	SeDenyServiceLogonRight
Deny log on locally	SeDenyInteractiveLogonRight
Deny log on through Remote Desktop Services	SeDenyRemoteInteractiveLogonRight
Enable computer and user accounts to be trusted for delegation	SeEnableDelegationPrivilege
Force shutdown from a remote system	SeRemoteShutdownPrivilege
Generate security audits	SeAuditPrivilege
Impersonate a client after authentication	SeImpersonatePrivilege
Increase a process working set	SeIncreaseWorkingSetPrivilege
Increase scheduling priority	SeIncreaseBasePriorityPrivilege
Load and unload device drivers	SeLoadDriverPrivilege
Lock pages in memory	SeLockMemoryPrivilege
Log on as a batch job	SeBatchLogonRight
Log on as a service	SeServiceLogonRight
Manage auditing and security log	SeSecurityPrivilege
Modify an object label	SeRelabelPrivilege

Group Policy Setting	Constant Name
Modify firmware environment values	SeSystemEnvironmentPrivilege
Perform volume maintenance tasks	SeManageVolumePrivilege
Profile single process	SeProfileSingleProcessPrivilege
Profile system performance	SeSystemProfilePrivilege
Remove computer from docking station	SeUndockPrivilege
Replace a process level token	SeAssignPrimaryTokenPrivilege
Restore files and directories	SeRestorePrivilege
Shut down the system	SeShutdownPrivilege
Synchronize directory service data	SeSyncAgentPrivilege
Take ownership of files or other objects	SeTakeOwnershipPrivilege

Access Credential Manager as a trusted caller

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Access Credential Manager as a trusted caller** policy setting is used by Credential Manager during backup and restore. No accounts should have this privilege because it is assigned only to the Winlogon service. Saved credentials of users may be compromised if this privilege is given to other entities.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeTrustedCredManAccessPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

- Do not modify this policy setting from the default.

Location

\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined
Client computer effective default settings	Not defined

Operating system version differences

This setting was introduced with Windows Vista and Windows Server 2008. There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

If an account is given this user right, the user of the account may create an application that calls into Credential Manager and is returned the credentials for another user.

Countermeasure

Do not define the **Access Credential Manager as a trusted caller** policy setting for any accounts besides Credential Manager.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Access this computer from the network

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Access this computer from the network** policy setting determines which users can connect to the computer from the network. This capability is required by a number of network protocols, including Server Message Block (SMB)-based protocols, NetBIOS, Common Internet File System (CIFS), and Component Object Model Plus (COM+).

Users, computers, and service accounts gain or lose the **Access this computer from network** user right by being explicitly or implicitly added or removed from a security group that has been granted this user right. For example, a user account or a computer account may be explicitly added to a custom security group or a built-in security group by an administrator, or it may be implicitly added by the operating system to a computed security group such as Domain Users, Authenticated Users, or Enterprise Domain Controllers.

By default, user accounts and computer accounts are granted the **Access this computer from network** user right when computed groups such as Authenticated Users, and for domain

controllers, the Enterprise Domain Controllers group, are defined in the default domain controllers Group Policy Object (GPO).

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeNetworkLogonRight

Possible values

- User-defined list of accounts
- Not defined

Best practices

- On desktop computers or member servers, grant this right only to users and administrators.
- On domain controllers, grant this right only to authenticated users, enterprise domain controllers, and administrators.
- This setting includes the **Everyone** group to ensure backward compatibility. Upon Windows upgrade, after you have verified that all users and groups are correctly migrated, you should remove the **Everyone** group and use the **Authenticated Users** group instead.

Location

\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Everyone, Administrators, Authenticated Users, Enterprise Domain Controllers, Pre-Windows 2000 Compatible Access
Stand-alone server default settings	Everyone, Administrators, Users, Backup Operators
Domain controller effective default settings	Everyone, Administrators, Authenticated Users, Enterprise Domain Controllers, Pre-Windows 2000 Compatible Access
Member server effective default settings	Everyone, Administrators, Users, Backup

Server type or GPO	Default value
	Operators
Client computer effective default settings	Everyone, Administrators, Users, Backup Operators

Operating system version differences

This policy setting was introduced in Windows Server 2003 and Windows XP with Service Pack 2. This setting has not changed, and it includes the **Everyone** group to ensure backward compatibility. After you have verified that all users and groups are correctly migrated, you should remove the **Everyone** group and use the **Authenticated Users** group instead.

Policy management

When modifying this user right, the following actions might cause users and services to experience network access issues:

- Removing the Enterprise Domain Controllers security group
- Removing the Authenticated Users group or an explicit group that allows users, computers, and service accounts the user right to connect to computers over the network
- Removing all users and computers

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can connect from their computer to the network can access resources on target computers for which they have permission. For example, the **Access this computer from the network** user right is required for users to connect to shared printers and folders. If this user right is assigned to the **Everyone** group, anyone in the group can read the files in those shared folders. This situation is unlikely because the groups created by a default installation of Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7 do not include the **Everyone** group. However, if a computer is upgraded and the original computer includes the **Everyone** group as part of its defined users and groups, that group is transitioned as part of the upgrade process and is present on the computer.

Countermeasure

Restrict the **Access this computer from the network** user right to only those users and groups who require access to the computer. For example, if you configure this policy setting to the **Administrators** and **Users** groups, users who log on to the domain can access resources that are shared from servers in the domain if members of the **Domain Users** group are included in the local **Users** group.

Note

If you are using IPsec to help secure network communications in your organization, ensure that a group that includes computer accounts is given this right. This right is required for successful computer authentication. Assigning this right to **Authenticated Users** or **Domain Computers** meets this requirement.

Potential impact

If you remove the **Access this computer from the network** user right on domain controllers for all users, no one can log on to the domain or use network resources. If you remove this user right on member servers, users cannot connect to those servers through the network. If you have installed optional components such as ASP.NET or Internet Information Services (IIS), you may need to assign this user right to additional accounts that are required by those components. It is important to verify that authorized users are assigned this user right for the computers that they need to access the network.

See Also

User Rights Assignment

Act as part of the operating system

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The **Act as part of the operating system** policy setting determines whether a process can assume the identity of any user and thereby gain access to the resources that the user is authorized to access. Typically, only low-level authentication services require this user right. Potential access is not limited to what is associated with the user by default. The calling process may request that arbitrary additional privileges be added to the access token. The calling process may also build an access token that does not provide a primary identity for auditing in the system event logs.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeTcbPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

- Do not assign this right to any user accounts. Only assign this user right to trusted users.
- If a service requires this user right, configure the service to log on by using the local System account, which inherently includes this user right. Do not create a separate account and assign this user right to it.

Location

\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default domain policy	Not defined
Default domain controller policy	Not defined
Stand-alone server default settings	Not defined
Domain controller effective default settings	Not defined
Member server effective default settings	Not defined

Server type or GPO	Default value
Client computer effective default settings	Not defined

Operating system version differences

This setting was introduced with Windows Vista and Windows Server 2008. There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Act as part of the operating system** user right is extremely powerful. Users with this user right can take complete control of the computer and erase evidence of their activities.

Countermeasure

Restrict the **Act as part of the operating system** user right to as few accounts as possible—it should not even be assigned to the Administrators group under typical circumstances. When a service requires this user right, configure the service to log on with the Local System account,

which inherently inlcudes this privilege. Do not create a separate account and assign this user right to it.

Potential impact

There should be little or no impact because the **Act as part of the operating system** user right is rarely needed by any accounts other than the Local System account.

See Also

User Rights Assignment

Add workstations to domain

This security policy reference topic for the IT professional describes the best practices, location, values, policy management and security considerations for this policy setting.

Reference

This policy setting determines which users can add a computer to a specific domain. For it to take effect, it must be assigned so that it applies to at least one domain controller. A user who is assigned this user right can add up to ten workstations to the domain.

Adding a computer account to the domain allows the computer to participate in Active Directory-based networking.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeMachineAccountPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Configure this setting so that only authorized members of the IT team are allowed to add computers to the domain.

Location

Computer Configuration\Windows Settings\Security Settings\User Rights Assignment

Default values

By default, this setting allows access for Authenticated Users on domain controllers, and it is not defined on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Not Defined
Stand-Alone Server Default Settings	Not Defined
Domain Controller Effective Default Settings	Authenticated Users
Member Server Effective Default Settings	Not Defined
Client Computer Effective Default Settings	Not Defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

Users can also join a computer to a domain if they have the Create Computer Objects permission for an organizational unit (OU) or for the Computers container in the directory. Users who are assigned this permission can add an unlimited number of computers to the domain regardless of whether they have the **Add workstations to domain** user right.

Furthermore, computer accounts that are created by means of the **Add workstations to domain** user right have Domain Administrators as the owner of the computer account. Computer accounts that are created by means of permissions on the computer's container use the creator as the owner of the computer account. If a user has permissions on the container and also has the **Add workstation to domain** user right, the computer is added based on the computer container permissions rather than the user right.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This policy has the following security considerations:

Vulnerability

The **Add workstations to domain** user right presents a moderate vulnerability. Users with this right could add a computer to the domain that is configured in a way that violates organizational security policies. For example, if your organization does not want its users to have administrative privileges on their computers, users could install Windows on their computers and then add the computers to the domain. The user would know the password for the local administrator account, could log on with that account, and then add a personal domain account to the local Administrators group.

Countermeasure

Configure this setting so that only authorized members of the IT team are allowed to add computers to the domain.

Potential impact

For organizations that have never allowed users to set up their own computers and add them to the domain, this countermeasure has no impact. For those that have allowed some or all users to configure their own computers, this countermeasure forces the organization to establish a formal process for these procedures going forward. It does not affect existing computers unless they are removed from and then added to the domain.

See Also

User Rights Assignment

Adjust memory quotas for a process

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This privilege determines who can change the maximum memory that can be consumed by a process. This privilege is useful for system tuning on a group or user basis.

This user right is defined in the Default Domain Controller Group Policy Object (GPO) and in the local security policy of workstations and servers.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeIncreaseQuotaPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Restrict the **Adjust memory quotas for a process** user right to only users who require the ability to adjust memory quotas to perform their jobs.
2. If this user right is necessary for a user account, it can be assigned to a local computer account instead of to a domain account.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\User Rights Assignment\

Default values

By default, members of the Administrators, Local Service, and Network Service groups have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Administrators Local Service Network Service
Default Domain Controller Policy	Administrators Local Service Network Service

Server type or GPO	Default value
Stand-Alone Server Default Settings	Administrators Local Service Network Service
Domain Controller Effective Default Settings	Administrators Local Service Network Service
Member Server Effective Default Settings	Administrators Local Service Network Service
Client Computer Effective Default Settings	Administrators Local Service Network Service

Operating system version differences

There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A user with the **Adjust memory quotas for a process** privilege can reduce the amount of memory that is available to any process, which could cause business-critical network applications to become slow or to fail. This privilege could be used by a malicious user to start a denial-of-service (DoS) attack.

Countermeasure

Restrict the **Adjust memory quotas for a process** user right to users who require it to perform their jobs, such as application administrators who maintain database management systems or domain administrators who manage the organization's directory and its supporting infrastructure.

Potential impact

Organizations that have not restricted users to roles with limited privileges may find it difficult to impose this countermeasure. Also, if you have installed optional components such as ASP.NET or IIS, you may need to assign the **Adjust memory quotas for a process** user right to additional accounts that are required by those components. IIS requires that this privilege be explicitly assigned to the IWAM_<ComputerName>, Network Service, and Service accounts. Otherwise, this countermeasure should have no impact on most computers. If this user right is necessary for a user account, it can be assigned to a local computer account instead of to a domain account.

See Also

User Rights Assignment

Allow log on locally

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can start an interactive session on the computer. Users must have this user right to log on over a Remote Desktop Services or Terminal Services session that is running on a Windows-based member computer or domain controller.

Note
Users who do not have this right are still able to start a remote interactive session on the computer if they have the **Allow logon through Remote Desktop Services** right.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeInteractiveLogonRight

Possible values

- User-defined list of accounts
- Not Defined

By default, the members of the following groups have this right on workstations and servers:

- Administrators
- Backup Operators
- Users

By default, the members of the following groups have this right on domain controllers:

- Account Operators
- Administrators
- Backup Operators
- Print Operators
- Server Operators

Best practices

1. Restrict this user right to legitimate users who must log on to the console of the computer.
2. If you selectively remove default groups, you can limit the abilities of users who are assigned to specific administrative roles in your organization.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Account Operators

Server type or GPO	Default value
	Administrators Backup Operators Print Operators Server Operators
Stand-Alone Server Default Settings	Administrators Backup Operators Users
Domain Controller Effective Default Settings	Account Operators Administrators Backup Operators Print Operators Server Operators
Member Server Effective Default Settings	Administrators Backup Operators Users
Client Computer Effective Default Settings	Administrators Backup Operators Users

Operating system version differences

There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic. Remote Desktop Services was formerly known as Terminal Services.

Policy management

Restarting the computer is not required to implement this change.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Modifying this setting might affect compatibility with clients, services, and applications. Use caution when removing service accounts that are used by components and by programs on member computers and on domain controllers in the domain from the default domain controller's policy. Also use caution when removing users or security groups that log on to the console of member computers in the domain, or removing service accounts that are defined in the local Security Accounts Manager (SAM) database of member computers or of workgroup computers.

If you want to grant a user account the ability to log on locally to a domain controller, you must make that user a member of a group that already has the **Allowed logon locally** system right or grant the right to that user account.

The domain controllers in the domain share the Default Domain Controllers Group Policy Object (GPO). When you grant an account the **Allow logon locally** right, you are allowing that account to log on locally to all domain controllers in the domain.

If the Users group is listed in the **Allow log on locally** setting for a GPO, all domain users can log on locally. The Users built-in group contains Domain Users as a member.

Group Policy

Group Policy settings are applied through GPOs in the following order, which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any account with the **Allow log on locally** user right can log on to the console of the computer. If you do not restrict this user right to legitimate users who must log on to the console of the computer, unauthorized users could download and run malicious software to elevate their privileges.

Countermeasure

For domain controllers, assign the **Allow log on locally** user right only to the Administrators group. For other server roles, you may choose to add Backup Operators in addition to Administrators. For end-user computers, you should also assign this right to the Users group.

Alternatively, you can assign groups such as Account Operators, Server Operators, and Guests to the **Deny log on locally** user right.

Potential impact

If you remove these default groups, you could limit the abilities of users who are assigned to specific administrative roles in your environment. If you have installed optional components such

as ASP.NET or IIS, you may need to assign the **Allow log on locally** user right to additional accounts that are required by those components. IIS requires that this user right be assigned to the IUSR_<*ComputerName*> account. You should confirm that delegated activities are not adversely affected by any changes that you make to the **Allow log on locally** user rights assignments.

See Also

User Rights Assignment

Allow log on through Remote Desktop Services

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy.

Reference

This policy setting determines which users or groups can access the logon screen of a remote computer through a Remote Desktop Services connection. It is possible for a user to establish a Remote Desktop Services connection to a particular server but not be able to log on to the console of that same server.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeRemoteInteractiveLogonRight

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. To control who can open a Remote Desktop Services connection and log on to the computer, add users to or remove users from the Remote Desktop Users group.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, members of the Administrators group have this right on domain controllers, workstations, and servers. The Remote Desktops Users group also has this right on workstations and servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators Remote Desktop Users
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators Remote Desktop Users
Client Computer Effective Default Settings	Administrators Remote Desktop Users

Operating system version differences

There have been no changes to the settings and effects of this user right policy setting since it was introduced in Windows Server 2003 and Windows XP. However, the policy setting name was changed in Windows Server 2008 R2 and Windows 7 from **Allow log on through Terminal Services**.

Policy management

This section describes different features and tools available to help you manage this policy.

Group Policy

To use RDS to successfully log on to a remote computer, the user or group must be a member of the Remote Desktop Users or Administrators group and be granted the **Allow log on through Remote Desktop Services** right. It is possible for a user to establish an RDS session to a particular server, but not be able to log on to the console of that same server.

To exclude users or groups, you can assign the **Deny log on through Remote Desktop Services** user right to those users or groups. However, be careful when you use this method because you could create conflicts for legitimate users or groups that have been allowed access through the **Allow log on through Remote Desktop Services** user right.

For more information, see Deny log on through Remote Desktop Services.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy settings are applied through GPOs in the following order, which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any account with the **Allow log on through Remote Desktop Services** user right can log on to the remote console of the computer. If you do not restrict this user right to legitimate users who must log on to the console of the computer, unauthorized users could download and run malicious software to elevate their privileges.

Countermeasure

For domain controllers, assign the **Allow log on through Remote Desktop Services** user right only to the Administrators group. For other server roles and end-user computers, add the Remote Desktop Users group. For servers that have the Remote Desktop (RD) Session Host role service enabled and do not run in Application Server mode, ensure that only authorized IT personnel who must manage the computers remotely belong to these groups.

Caution
For RD Session Host servers that run in Application Server mode, ensure that only users who require access to the server have accounts that belong to the Remote Desktop Users group because this built-in group has this logon right by default.

Alternatively, you can assign the **Deny log on through Remote Desktop Services** user right to groups such as Account Operators, Server Operators, and Guests. However, be careful when you use this method because you could block access to legitimate administrators who also belong to a group that has the **Deny log on through Remote Desktop Services** user right.

Potential impact

Removal of the **Allow log on through Remote Desktop Services** user right from other groups (or membership changes in these default groups) could limit the abilities of users who perform specific administrative roles in your environment. You should confirm that delegated activities are not adversely affected.

See Also

User Rights Assignment

Back up files and directories

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This user right determines which users can bypass file and directory, registry, and other persistent object permissions for the purposes of backing up the system. This user right is effective only when an application attempts access through the NTFS backup application programming interface (API) through a backup tool such as NTBACKUP.EXE. Otherwise, standard file and directory permissions apply.

This user right is similar to granting the following permissions to the user or group you have selected on all files and folders on the system:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

Default on workstations and servers:

- Administrators
- Backup Operators

Default on domain controllers:

- Administrators
- Backup Operators
- Server Operators

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeBackupPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Restrict the **Back up files and directories** user right to members of the IT team who must back up organizational data as part of their daily job responsibilities. Because there is no way to be sure that a user is backing up data, stealing data, or copying data to be distributed, only assign this user right to trusted users.
2. If you are using backup software that runs under specific service accounts, only these accounts (and not the IT staff) should have the **Back up files and directories** user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this right is granted to Administrators and Backup Operators on workstations and servers. On domain controllers, Administrators, Backup Operators, and Server Operators have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators Backup Operators Server Operators
Stand-Alone Server Default Settings	Administrators Backup Operators
Domain Controller Effective Default Settings	Administrators Backup Operators Server Operators
Member Server Effective Default Settings	Administrators Backup Operators
Client Computer Effective Default Settings	Administrators

Server type or GPO	Default value
	Backup Operators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can back up data from a computer could take the backup media to a non-domain computer on which they have administrative privileges, and then restore the data. They could take ownership of the files and view any unencrypted data that is contained within the backup set.

Countermeasure

Restrict the **Back up files and directories** user right to members of the IT team who must back up organizational data as part of their daily job responsibilities. If you are using backup software that runs under specific service accounts, only these accounts (and not the IT staff) should have the **Back up files and directories** user right.

Potential impact

Changes in the membership of the groups that have the **Back up files and directories** user right could limit the abilities of users who are assigned to specific administrative roles in your environment. You should confirm that authorized backup administrators can still perform backup operations.

See Also

User Rights Assignment

Bypass traverse checking

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users (or a process that acts on behalf of the user's account) have permission to navigate an object path in the NTFS file system or in the registry without being checked for the Traverse Folder special access permission. This user right does not allow the user to list the contents of a folder. It only allows the user to traverse folders to access permitted files or subfolders.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeChangeNotifyPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Use access–based enumeration when you want to prevent users from seeing any folder or file to which they do not have access.
2. Use the default settings of this policy in most cases. If you change the settings, verify your intent through testing.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators Authenticated Users Everyone Local Service Network Service Pre-Windows 2000 Compatible Access
Stand-Alone Server Default Settings	Administrators Backup Operators Users Everyone Local Service Network Service
Domain Controller Effective Default Settings	Administrators Authenticated Users Everyone Local Service Network Service Pre-Windows 2000 Compatible Access
Member Server Effective Default Settings	Administrators Backup Operators Users Everyone Local Service Network Service
Client Computer Effective Default Settings	Administrators Backup Operators Users Everyone

Server type or GPO	Default value
	Local Service Network Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

Permissions to files and folders are controlled though the appropriate configuration of file system access control lists (ACLs).The ability to traverse the folder does not provide any Read or Write permissions to the user.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The default configuration for the **Bypass traverse checking** setting is to allow all users to bypass traverse checking. Permissions to files and folders are controlled though the appropriate configuration of file system access control lists (ACLs) because the ability to traverse the folder does not provide any Read or Write permissions to the user. The only scenario in which the default configuration could lead to a mishap would be if the administrator who configures

permissions does not understand how this policy setting works. For example, the administrator might expect that users who are unable to access a folder are unable to access the contents of any child folders. Such a situation is unlikely, and, therefore, this vulnerability presents little risk.

Countermeasure

Organizations that are extremely concerned about security may want to remove the Everyone group, and perhaps the Users group, from the list of groups that have the **Bypass traverse checking** user right. Taking explicit control over traversal assignments can be an effective way to limit access to sensitive information. Access–based enumeration can also be used. If you use access–based enumeration, users cannot see any folder or file to which they do not have access. For more information about this feature, see Access-based Enumeration.

Potential impact

The Windows operating systems and many applications were designed with the expectation that anyone who can legitimately access the computer will have this user right. Therefore, we recommend that you thoroughly test any changes to assignments of the **Bypass traverse checking** user right before you make such changes to production systems. In particular, IIS requires this user right to be assigned to the Network Service, Local Service, IIS_WPG, IUSR_<*ComputerName*>, and IWAM_<*ComputerName*> accounts. (It must also be assigned to the ASPNET account through its membership in the Users group.) We recommend that you leave this policy setting at its default configuration.

See Also

User Rights Assignment

Change the system time

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can adjust the time on the computer's internal clock. This right allows the computer user to change the date and time associated with records in the event logs, database transactions, and the file system. This right is also required by the process that performs time synchronization. This setting does not impact the user’s ability to change the time zone or other display characteristics of the system time. For information about assigning the right to change the time zone, see Change the time zone.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeSystemtimePrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Restrict the **Change the system time** user right to users with a legitimate need to change the system time, such as members of the IT team.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, members of the Administrators and Local Service groups have this right on workstations and servers. Members of the Administrators, Server Operators, and Local Service groups have this right on domain controllers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators Server Operators Local Service
Stand-Alone Server Default Settings	Administrators Local Service
DC Effective Default Settings	Administrators Server Operators Local Service
Member Server Effective Default Settings	Administrators Local Service
Client Computer Effective Default Settings	Administrators Local Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can change the time on a computer could cause several problems. For example:

- Time stamps on event log entries could be made inaccurate
- Time stamps on files and folders that are created or modified could be incorrect
- Computers that belong to a domain might not be able to authenticate themselves
- Users who try to log on to the domain from computers with inaccurate time might not be able to authenticate.

Also, because the Kerberos authentication protocol requires that the requester and authenticator have their clocks synchronized within an administrator-defined skew period, an attacker who changes a computer's time may cause that computer to be unable to obtain or grant Kerberos protocol tickets.

The risk from these types of events is mitigated on most domain controllers, member servers, and end-user computers because the Windows Time Service automatically synchronizes time with domain controllers in the following ways:

- All desktop client computers and member servers use the authenticating domain controller as their inbound time partner.
- All domain controllers in a domain nominate the primary domain controller (PDC) emulator operations master as their inbound time partner.
- All PDC emulator operations masters follow the hierarchy of domains in the selection of their inbound time partner.
- The PDC emulator operations master at the root of the domain is authoritative for the organization. Therefore, we recommend that you configure this computer to synchronize with a reliable external time server.

This vulnerability becomes much more serious if an attacker is able to change the system time and then stop the Windows Time Service or reconfigure it to synchronize with a time server that is not accurate.

Countermeasure

Restrict the **Change the system time** user right to users with a legitimate need to change the system time, such as members of the IT team.

Potential impact

There should be no impact because time synchronization for most organizations should be fully automated for all computers that belong to the domain. Computers that do not belong to the domain should be configured to synchronize with an external source, such as a web service.

See Also

User Rights Assignment

Change the time zone

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can adjust the time zone that is used by the computer for displaying the local time, which includes the computer's system time plus the time zone offset.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeTimeZonePrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

None.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators Users
Stand-Alone Server Default Settings	Administrators Users
Domain Controller Effective Default Settings	Administrators Users
Member Server Effective Default Settings	Administrators Users
Client Computer Effective Default Settings	Administrators Users

Operating system version differences

There are no differences in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the account for this user right assignment becomes effective the next time the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Changing the time zone represents little vulnerability because the system time is not affected. This setting merely enables users to display their preferred time zone while being synchronized with domain controllers in different time zones.

Countermeasure

Countermeasures are not required because system time is not affected by this setting.

Potential impact

None.

See Also

User Rights Assignment

Create a pagefile

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

The Windows operating system, by default, designates a section of the hard drive as virtual memory known as the page file, or more specifically, as pagefile.sys. It is used to supplement the computer's Random Access Memory (RAM) to improve performance for programs and data that are used frequently. Although the file is hidden from browsing, you can manage it using the system settings.

This policy setting determines which users can create and change the size of a page file. It determines whether users can specify a page file size for a particular drive in the **Performance Options** box located on the **Advanced** tab of the **System Properties** dialog box or through using internal application interfaces (APIs).

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeCreatePagefilePrivilege

Possible values

- User-defined list of accounts
- Administrators

Best practices

1. Restrict the **Create a pagefile** user right to Administrators, which is the default.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, members of the Administrators group have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Administrators
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators

Server type or GPO	Default value
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can change the page file size could make it extremely small or move the file to a highly fragmented storage volume, which could cause reduced computer performance.

Countermeasure

Restrict the **Create a pagefile** user right to members of the Administrators group.

Potential impact

None. Restricting this right to members of the Administrators group is the default configuration.

See Also

User Rights Assignment

Create a token object

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which accounts a process can use to create a token, and which accounts it can then use to gain access to local resources when the process uses NtCreateToken() or other token-creation APIs.

When a user logs on to the local computer or connects to a remote computer through a network, Windows builds the user's access token. Then the system examines the token to determine the level of the user's privileges. When you revoke a privilege, the change is immediately recorded, but the change is not reflected in the user's access token until the next time the user logs on or connects.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeCreateTokenPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. This user right is used internally by the operating system. Unless it is necessary, do not assign this user right to a user, group, or process other than Local System.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

This user right is used internally by the operating system. By default, it is not assigned to any user groups.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Not Defined
Stand-Alone Server Default Settings	Not Defined
Domain Controller Effective Default Settings	Local System
Member Server Effective Default Settings	Local System
Client Computer Effective Default Settings	Local System

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Caution
A user account that is given this user right has complete control over the system, and it can lead to the system being compromised. We highly recommend that you do not assign this right to any user accounts.

The operating system examines a user's access token to determine the level of the user's privileges. Access tokens are built when users log on to the local computer or connect to a remote computer over a network. When you revoke a privilege, the change is immediately recorded, but the change is not reflected in the user's access token until the next time the user logs on or connects. Users with the ability to create or modify tokens can change the level of access for any account on a computer if they are currently logged on. They could escalate their privileges or create a DoS condition.

Countermeasure

Do not assign the **Create a token object** user right to any users. Processes that require this user right should use the Local System account, which already includes it, instead of a separate user account that has this user right assigned.

Potential impact

None. Not Defined is the default configuration.

See Also

User Rights Assignment

Create global objects

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can create global objects that are available to all sessions. Users can still create objects that are specific to their own session if they do not have this user right.

A global object is an object that is created to be used by any number of processes or threads, even those not started within the user's session. Remote Desktop Services uses global objects in its processes to facilitate connections and access.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeCreateGlobalPrivilege

Possible values

- User-defined list of accounts
- Default accounts listed below

Best practices

1. Do not assign any user accounts this right.

Location

GPO_name\Computer Configuration\Windows Settings\Security SettingsLocal Policies\User Rights Assignment

Default values

By default, members of the Administrators group have this right, as do Local Service and Network Service accounts on the supported versions of Windows. Service is included for backwards compatibility with earlier versions of Windows.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Administrators Local Service Network Service Service
Stand-Alone Server Default Settings	Administrators Local Service Network Service Service
Domain Controller Effective Default Settings	Administrators Local Service Network Service Service

Server type or GPO	Default value
Member Server Effective Default Settings	Administrators Local Service Network Service Service
Client Computer Effective Default Settings	Administrators Local Service Network Service Service

Operating system version differences

Prior to Windows Server 2008 and Windows Vista, the default values of this security policy setting included Administrators, Service, and Local System. There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

A restart of the computer is not required for this policy setting to take effect.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Caution
A user account that is given this user right has complete control over the system, and it can lead to the system being compromised. We highly recommend that you do not assign this right to any user accounts.

The operating system examines a user's access token to determine the level of the user's privileges. Access tokens are built when users log on to the local computer or connect to a remote computer over a network. When you revoke a privilege, the change is immediately recorded, but the change is not reflected in the user's access token until the next time the user logs on or connects. Users with the ability to create or modify tokens can change the level of access for any currently logged on account. They could escalate their privileges or create a denial-of-service (DoS) condition.

Countermeasure

Do not assign the **Create a token object** user right to any users. Processes that require this user right should use the Local System account, which already includes it, instead of a separate user account with this user right assigned.

Potential impact

None. Not Defined is the default domain policy configuration.

See Also

User Rights Assignment

Create permanent shared objects

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This user right determines which accounts can be used by processes to create a directory object by using the object manager. Directory objects include Active Directory objects, files and folders, printers, registry keys, processes, and threads. Users who have this capability can create permanent shared objects, including devices, semaphores, and mutexes. This user right is useful to kernel-mode components that extend the object namespace. Because components that are running in kernel-mode inherently have this user right assigned to them, it is not necessary to specifically assign it.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeCreatePermanentPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Users who have the **Create permanent shared objects** user right could create new shared objects and expose sensitive data to the network. Therefore, do not assign this right to any users.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, **LocalSystem** is the only account that has this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Not Defined
Stand-Alone Server Default Settings	Not Defined
Domain Controller Effective Default Settings	**LocalSystem**
Member Server Effective Default Settings	**LocalSystem**
Client Computer Effective Default Settings	**LocalSystem**

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes different features and tools available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who have the **Create permanent shared objects** user right could create new shared objects and expose sensitive data to the network.

Countermeasure

Do not assign the **Create permanent shared objects** user right to any users. Processes that require this user right should use the System account, which already includes this user right, instead of a separate user account.

Potential impact

None. Not Defined is the default configuration.

See Also

User Rights Assignment

Create symbolic links

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This user right determines if users can create a symbolic link from the computer they are logged on to.

A symbolic link is a file-system object that points to another file-system object. The object that is pointed to is called the target. Symbolic links are transparent to users. The links appear as normal files or directories, and they can be acted upon by the user or application in exactly the same manner. Symbolic links are designed to aid in migration and application compatibility with UNIX operating systems. Microsoft has implemented symbolic links to function just like UNIX links.

Warning This privilege should only be given to trusted users. Symbolic links can expose security vulnerabilities in applications that aren't designed to handle them.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeCreateSymbolicLinkPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. This user right should only be given to trusted users. Symbolic links can expose security vulnerabilities in applications that are not designed to handle them.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, members of the Administrators group have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Not Defined
Stand-Alone Server Default Settings	Not Defined
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes different features and tools available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Command-line tools

This setting can be used in conjunction with a symbolic link file system setting that can be manipulated with the command-line tool to control the kinds of symlinks that are allowed on the computer. Type **fsutil behavior set symlinkevalution /?** at the command line to get more information about **fsutil** and symbolic links.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who have the **Create symbolic links** user right could inadvertently or maliciously expose your system to symbolic link attacks. Symbolic link attacks can be used to change the permissions on a file, to corrupt data, to destroy data, or as a DoS attack.

Countermeasure

Do not assign the **Create symbolic links** user right to standard users. Restrict this right to trusted administrators. You can use the **fsutil** command to establish a symbolic link file system setting that controls the kind of symbolic links that can be created on a computer.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Debug programs

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can attach to or open any process, even those they do not own. Developers who are debugging their own applications do not need to be assigned this user right. Developers who are debugging new system components need this user right. This user right provides access to sensitive and critical operating-system components.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDebugPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. Assign this user right only to trusted users to reduce security vulnerabilities.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, members of the Administrators group have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features and tools that are available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Debug programs** user right can be exploited to capture sensitive computer information from system memory or to access and modify kernel or application structures. Some attack tools exploit this user right to extract hashed passwords and other private security information or to insert rootkit code. By default, the **Debug programs** user right is assigned only to administrators, which helps mitigate risk from this vulnerability.

Countermeasure

Remove the accounts of all users and groups that do not require the **Debug programs** user right.

Potential impact

If you revoke this user right, no one can debug programs. However, typical circumstances rarely require this capability on production computers. If an issue arises that requires an application to be debugged on a production server, you can move the server to a different organizational unit (OU) temporarily and assign the **Debug programs** user right to a separate Group Policy for that OU.

See Also

User Rights Assignment

Deny access to this computer from the network

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines which users are prevented from accessing a computer over the network.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDenyNetworkLogonRight

Possible values

- User-defined list of accounts
- Guest

Best practices

1. Because all Active Directory Domain Services programs use a network logon for access, use caution when you assign this user right on domain controllers.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this setting is Guest on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Guest
Stand-Alone Server Default Settings	Guest
Domain Controller Effective Default Settings	Guest

Server type or GPO	Default value
Member Server Effective Default Settings	Guest
Client Computer Effective Default Settings	Guest

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features and tools available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

This policy setting supersedes the **Access this computer from the network** policy setting if a user account is subject to both policies.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users who can log on to the computer over the network can enumerate lists of account names, group names, and shared resources. Users with permission to access shared folders and files can connect over the network and possibly view or modify data.

Countermeasure

Assign the **Deny access to this computer from the network** user right to the following accounts:

- Anonymous logon
- Built-in local Administrator account
- Local Guest account
- All service accounts

An important exception to this list is any service accounts that are used to start services that must connect to the computer over the network. For example, let's say you have configured a shared folder for web servers to access, and you present content within that folder through a website. You may need to allow the account that runs IIS to log on to the server with the shared folder from the network. This user right is particularly effective when you must configure servers and workstations on which sensitive information is handled because of regulatory compliance concerns.

Potential impact

If you configure the **Deny access to this computer from the network** user right for other accounts, you could limit the abilities of users who are assigned to specific administrative roles in your environment. You should verify that delegated tasks are not negatively affected.

See Also

User Rights Assignment

Deny log on as a batch job

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which accounts are prevented from logging on by using a batch-queue tool to schedule and start jobs automatically in the future. The ability to log on by using a batch-queue tool is needed for any account that is used to start scheduled jobs by means of the Task Scheduler.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDenyBatchLogonRight

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. When you assign this user right, thoroughly test that the effect is what you intended.
2. Within a domain, modify this setting on the applicable Group Policy Object (GPO).
3. **Deny log on as a batch job** prevents administrators or operators from using their personal accounts to schedule tasks, which helps with business continuity when that person transitions to other positions or responsibilities.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Note
This setting is not compatible with computers running Windows 2000 Server with Service Pack 1 or earlier.

Policy management

This section describes features and tools available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

This policy setting might conflict with and negate the **Log on as a batch job** setting.

Group Policy

On a domain-joined computer, including the domain controller, this policy can be overwritten by a domain policy, which will prevent you from modifying the local policy setting.

For example, if you are trying to configure Task Scheduler on your domain controller, check the Settings tab of your two domain controller policy and domain policy GPOs in the Group Policy Management Console (GPMC). Verify the targeted account is not present in the **Deny log on as a batch job** User Rights Assignment and also correctly configured in the **Log on as a batch job** setting.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Accounts that have the **Deny log on as a batch job** user right could be used to schedule jobs that could consume excessive computer resources and cause a denial-of-service condition.

Countermeasure

Assign the **Deny log on as a batch job** user right to the local Guest account.

Potential impact

If you assign the **Deny log on as a batch job** user right to other accounts, you could deny the ability to perform required job activities to users who are assigned specific administrative roles. You should confirm that delegated tasks are not affected adversely.

See Also

User Rights Assignment

Deny log on as a service

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users are prevented from logging on to the service applications on a computer.

A service is an application type that runs in the system background without a user interface. It provides core operating system features, such as web serving, event logging, file serving, printing, cryptography, and error reporting.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDenyServiceLogonRight

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. When you assign this user right, thoroughly test that the effect is what you intended.
2. Within a domain, modify this setting on the applicable Group Policy Object (GPO).

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features and tools available to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

For information about configuring a service, see Configure How a Service Is Started.

Group Policy

On a domain-joined computer, including the domain controller, this policy can be overwritten by a domain policy, which will prevent you from modifying the local policy setting.

This policy setting might conflict with and negate the **Log on as a service** setting.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Accounts that can log on to a service application could be used to configure and start new unauthorized services, such as a keylogger or other malicious software. The benefit of the specified countermeasure is somewhat reduced by the fact that only users with administrative rights can install and configure services, and an attacker who has already attained that level of access could configure the service to run by using the System account.

Countermeasure

We recommend that you not assign the **Deny log on as a service** user right to any accounts. This is the default configuration. Organizations that are extremely concerned about security might assign this user right to groups and accounts when they are certain that they will never need to log on to a service application.

Potential impact

If you assign the **Deny log on as a service** user right to specific accounts, services may not start and a denial-of-service condition could result.

See Also

User Rights Assignment

Deny log on locally

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users are prevented from logging on directly at the computer's console.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDenyInteractiveLogonRight

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. Assign the **Deny log on locally** user right to the local guest account to restrict access by potentially unauthorized users.
2. Test your modifications to this policy setting in conjunction with the **Allow log on locally** policy setting to determine if the user account is subject to both policies.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Note
This setting is not compatible with computers running Windows 2000 Server with Service Pack 1 or earlier.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

If you apply this policy setting to the Everyone group, no one will be able to log on locally.

Group Policy

This policy setting supersedes the **Allow log on locally** policy setting if a user account is subject to both policies.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any account with the ability to log on locally could be used to log on at the console of the computer. If this user right is not restricted to legitimate users who must log on to the console of the computer, unauthorized users might download and run malicious software that elevates their user rights.

Countermeasure

Assign the **Deny log on locally** user right to the local Guest account. If you have installed optional components such as ASP.NET, you may want to assign this user right to additional accounts that are required by those components.

Potential impact

If you assign the **Deny log on locally** user right to additional accounts, you could limit the abilities of users who are assigned to specific roles in your environment. However, this user right should explicitly be assigned to the ASPNET account on computers that are configured with the Web Server role. You should confirm that delegated activities are not adversely affected.

See Also

User Rights Assignment

Deny log on through Remote Desktop Services

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users are prevented from logging on to the computer through a remote desktop connection through Remote Desktop Services. It is possible for a user to establish a remote desktop connection to a particular server, but not be able to log on to the console of that server.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeDenyRemoteInteractiveLogonRight

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. To control who can open a remote desktop connection and log on to the computer, add the user account to or remove user accounts from the Remote Desktop Users group.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Remote Desktop Services was formerly named Terminal Services. The naming difference does not impact how this policy setting works.

Note
This setting is not compatible with computers running Windows 2000 Server with Service Pack 1 or earlier.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

The **Remote System** property controls settings for Remote Desktop Services (**Allow or prevent remote connections to the computer**) and for Remote Assistance (**Allow Remote Assistance connections to this computer**).

Group Policy

This policy setting supersedes the **Allow log on through Remote Desktop Services** policy setting if a user account is subject to both policies.

Group Policy settings are applied in the following order. They overwrite settings on the local computer at the next Group Policy update.

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. Organizational unit policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any account with the right to log on through Remote Desktop Services could be used to log on to the remote console of the computer. If this user right is not restricted to legitimate users who need to log on to the console of the computer, malicious users might download and run software that elevates their user rights.

Countermeasure

Assign the **Deny log on through Remote Desktop Services** user right to the built-in local guest account and all service accounts. If you have installed optional components, such as ASP.NET, you may want to assign this user right to additional accounts that are required by those components.

Potential impact

If you assign the **Deny log on through Remote Desktop Services** user right to other groups, you could limit the abilities of users who are assigned to specific administrative roles in your environment. Accounts that have this user right cannot connect to the computer through Remote Desktop Services or Remote Assistance. You should confirm that delegated tasks are not negatively affected.

See Also

User Rights Assignment

Enable computer and user accounts to be trusted for delegation

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can set the **Trusted for Delegation** setting on a user or computer object.

Security account delegation provides the ability to connect to multiple servers, and each server change retains the authentication credentials of the original client. Delegation of authentication is a capability that client and server applications use when they have multiple tiers. It allows a public-facing service to use client credentials to authenticate to an application or database service. For this configuration to be possible, the client and the server must run under accounts that are trusted for delegation.

Only administrators who have the **Enable computer and user accounts to be trusted for delegation** credential can set up delegation. Domain admins and Enterprise admins have this credential. The procedure to allow a user to be trusted for delegation depends on the functionality level of the domain.

The user or computer object that is granted this right must have write access to the account control flags. A server process running on a computer (or under a user context) that is trusted for delegation can access resources on another computer by using the delegated credentials of a client. However, the client account must have Write access to the account control flags on the object.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeEnableDelegationPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. There is no reason to assign this user right to anyone on member servers and workstations that belong to a domain because it has no meaning in those contexts. It is only relevant on domain controllers and stand-alone computers.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic. However, configuring this policy setting might have different effects on servers if those servers have different domain functional levels.

Policy management

This section describes features, tools and guidance to help you manage this policy.

Modifying this setting might affect compatibility with clients, services, and applications.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

This user right is defined in the Default Domain Controller Group Policy Object (GPO) and in the local security policy of workstations and servers.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Misuse of the **Enable computer and user accounts to be trusted for delegation** user right could allow unauthorized users to impersonate other users on the network. An attacker could exploit this privilege to gain access to network resources and make it difficult to determine what has happened after a security incident.

Countermeasure

The **Enable computer and user accounts to be trusted for delegation** user right should be assigned only if there is a clear need for its functionality. When you assign this right, you should investigate the use of constrained delegation to control what the delegated accounts can do. On domain controllers, this right is assigned to the Administrators group by default.

Note

There is no reason to assign this user right to anyone on member servers and workstations that belong to a domain because it has no meaning in those contexts. It is only relevant on domain controllers and stand-alone computers.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Force shutdown from a remote system

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines which users are allowed to shut down a computer from a remote location on the network. This allows members of the Administrators group or specific users to manage computers (for tasks such as a restart) from a remote location.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeRemoteShutdownPrivilege

Possible values

- User-defined list of accounts

- Administrators

Best practices

1. Explicitly restrict this user right to members of the Administrators group or other specifically assigned roles that require this capability, such as nonadministrative operations staff.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators and Server Operators on domain controllers and Administrators on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators Server Operators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators Server Operators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

This setting was introduced with Windows Vista and Windows Server 2008. There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

This policy setting must be applied on the computer that is being accessed remotely.

Group Policy

This user right is defined in the Default Domain Controller Group Policy Object (GPO) and in the local security policy of workstations and servers.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any user who can shut down a computer could cause a denial-of-service condition to occur. Therefore, this user right should be tightly restricted.

Countermeasure

Restrict the **Force shutdown from a remote system** user right to members of the Administrators group or other specifically assigned roles that require this capability, such as nonadministrative operations staff.

Potential impact

On a domain controller, if you remove the **Force shutdown from a remote system** user right from the Server Operator group, you could limit the abilities of users who are assigned to specific administrative roles in your environment. You should confirm that delegated activities are not adversely affected.

See Also

User Rights Assignment

Generate security audits

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which accounts can be used by a process to generate audit records in the Security log. The Local Security Authority Subsystem Service writes events to the log. You can use the information in the Security log to trace unauthorized computer access.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeAuditPrivilege

Possible values

- User-defined list of accounts
- Local Service
- Network Service

Best practices

1. Because the audit log can potentially be an attack vector if an account is compromised, ensure that only the Local Service and Network Service accounts have the **Generate security audits** user right assigned to them.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this setting is Local Service and Network Service on domain controllers and stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Local Service Network Service

Server type or GPO	Default value
Stand-Alone Server Default Settings	Local Service Network Service
Domain Controller Effective Default Settings	Local Service Network Service
Member Server Effective Default Settings	Local Service Network Service
Client Computer Effective Default Settings	Local Service Network Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Misuse of this user right can result in the generation of many auditing events, potentially hiding evidence of an attack or causing a denial-of-service (DoS) if the **Audit: Shut down system immediately if unable to log security audits** security policy setting is enabled. For more information see **Audit: Shut down system immediately if unable to log security audits**.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A malicious user could use accounts that can write to the Security log to fill that log with meaningless events. If the computer is configured to overwrite events as needed, malicious users could use this method to remove evidence of their unauthorized activities. If the computer is configured to shut down when it is unable to write to the Security log, and it is not configured to automatically back up the log files, this method could be used to create a DoS condition.

Countermeasure

Ensure that only the Local Service and Network Service accounts have the **Generate security audits** user right assigned to them.

Potential impact

None. Restricting the **Generate security audits** user right to the Local Service and Network Service accounts is the default configuration.

See Also

User Rights Assignment

Impersonate a client after authentication

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which programs are allowed to impersonate a user or another specified account and act on behalf of the user. If this user right is required for this type of impersonation, an unauthorized user cannot cause a client to connect (for example, by remote procedure call (RPC) or named pipes) to a service that they have created to impersonate that client. (Such an action could elevate the unauthorized user's permissions to administrative or system levels.)

Impersonation is the ability of a thread to run in a security context that is different from the context of the process that owns the thread. Impersonation is designed to meet the security requirements

of client/server applications. When running in a client's security context, a service "is" the client, to some degree. One of the service's threads uses an access token representing the client's credentials to obtain access to the objects to which the client has access.

The primary reason for impersonation is to cause access checks to be performed against the client's identity. Using the client's identity for access checks can cause access to be either restricted or expanded, depending on what the client has permission to do.

Services that are started by the Service Control Manager have the built-in Service group added by default to their access tokens. COM servers that are started by the COM infrastructure and configured to run under a specific account also have the Service group added to their access tokens. As a result, these processes are assigned this user right when they are started.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeImpersonatePrivilege

Possible values

- User-defined list of accounts
- Default values
- Not defined

Best practices

1. A user can impersonate an access token if any of the following conditions exist:
   - The access token that is being impersonated is for this user.
   - The user in this session logged on to the network with explicit credentials to create the access token.
   - The requested level is less than Impersonate, such as Anonymous or Identify.

   Because of these factors, users do not usually need to have this user right assigned.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this setting is Administrators, Local Service, Network Service, and Service on domain controllers and stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy’s property page.

Server type or GPO	Default value
Default Domain Policy	Not eefined
Default Domain Controller Policy	Administrators Local Service Network Service Service
Stand-Alone Server Default Settings	Administrators Local Service Network Service Service
Domain Controller Effective Default Settings	Administrators Local Service Network Service Service
Member Server Effective Default Settings	Administrators Local Service Network Service Service
Client Computer Effective Default Settings	Administrators Local Service Network Service Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An attacker with the **Impersonate a client after authentication** user right could create a service, mislead a client into connecting to the service, and then impersonate that computer to elevate the attacker's level of access to that of the computer.

Countermeasure

On member servers, ensure that only the Administrators and Service groups (Local Service, Network Service, and Service) have the **Impersonate a client after authentication** user right assigned to them.

Potential impact

In most cases, this configuration has no impact. If you have installed optional components such as ASP.NET or IIS, you may need to assign the **Impersonate a client after authentication** user right to additional accounts that are required by those components, such as IUSR_<*ComputerName*>, IIS_WPG, ASP.NET, or IWAM_<*ComputerName*>.

See Also

User Rights Assignment

Increase a process working set

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can increase or decrease the size of the working set of a process. The working set of a process is the set of memory pages currently visible to the process in physical RAM. These pages are resident, and they are available for an application to use without triggering a page fault. The minimum and maximum working set sizes affect the virtual memory paging behavior of a process.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeIncreaseWorkingSetPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. You should make users aware that adverse performance issues may occur if they modify this security setting.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, standard users have this right.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not Defined
Default Domain Controller Policy	Users
Stand-Alone Server Default Settings	Users
Domain Controller Effective Default Settings	Users
Member Server Effective Default Settings	Users
Client Computer Effective Default Settings	Users

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Increasing the working set size for a process decreases the amount of physical memory that is available to the rest of the system.

Countermeasure

Increase user's awareness about the impact of increasing the working set of a process and how to recognize that their system is adversely affected if they change this setting.

Potential impact

None. Allowing standard users to increase the working set of a process is the default configuration.

See Also

User Rights Assignment

Increase scheduling priority

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which user accounts can increase the base priority class of a process. It is not a privileged operation to increase relative priority within a priority class. This user right is not required by administrative tools that are supplied with the operating system, but it might be required by software development tools.

Specifically, this security setting determines which accounts can use a process with Write Property access to another process to increase the run priority that is assigned to the other process. A user with this privilege can change the scheduling priority of a process through the Task Manager user interface.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeIncreaseBasePriorityPrivilege

Possible values

- User-defined list of accounts
- Not defined
- Administrators

Best practices

1. Allow the default value, Administrators, as the only account responsible for controlling process scheduling priorities.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A user who is assigned this user right could increase the scheduling priority of a process to Real-Time, which would leave little processing time for all other processes and could lead to a denial-of-service condition.

Countermeasure

Verify that only Administrators have the **Increase scheduling priority** user right assigned to them.

Potential impact

None. Restricting the **Increase scheduling priority** user right to members of the Administrators group is the default configuration.

See Also

User Rights Assignment

Load and unload device drivers

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can dynamically load and unload device drivers. This user right is not required if a signed driver for the new hardware already exists in the Driver.cab file on the computer. Device drivers run as highly privileged code.

Windows supports the Plug and Play specifications that define how a computer can detect and configure newly added hardware, and then automatically install the device driver. Prior to Plug and Play, users needed to manually configure devices before attaching them to the computer. This model allows a user to plug in the hardware, then Windows searches for an appropriate device driver package and automatically configures it to work without interfering with other devices.

Because device driver software runs as if it is a part of the operating system with unrestricted access to the entire computer, it is critical that only known and authorized device drivers be permitted.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeLoadDriverPrivilege

Possible values

- User-defined list of accounts
- Default values
- Not Defined

Best practices

1. Because of the potential security risk, do not assign this user right to any user, group, or process that you do not want to take over the system.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators and Print Operators on domain controllers and Administrators on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators Print Operators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators Print Operators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Device drivers run as highly privileged code. A user who has the **Load and unload device drivers** user right could unintentionally install malicious software that masquerades as a device driver. Administrators should exercise care and install only drivers with verified digital signatures.

Note

You must have this user right or be a member of the local Administrators group to install a new driver for a local printer or to manage a local printer and configure defaults for options such as duplex printing.

Countermeasure

Do not assign the **Load and unload device drivers** user right to any user or group other than Administrators on member servers. On domain controllers, do not assign this user right to any user or group other than Domain Admins.

Potential impact

If you remove the **Load and unload device drivers** user right from the Print Operators group or other accounts, you could limit the abilities of users who are assigned to specific administrative roles in your environment. You should ensure that delegated tasks are not negatively affected.

See Also

User Rights Assignment

Lock pages in memory

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which accounts can use a process to keep data in physical memory, which prevents the computer from paging the data to virtual memory on a disk.

Normally, an application running on Windows can negotiate for more physical memory, and in response to the request, the application begins to move the data from RAM (such as the data cache) to a disk. When the pageable memory is moved to a disk, more RAM is free for the operating system to use.

Enabling this policy setting for a specific account (a user account or a process account for an application) prevents paging of the data. Thereby, the amount of memory that Windows can reclaim under pressure is limited. This could lead to performance degradation.

Note

By configuring this policy setting, the performance of the Windows operating system will differ depending on if applications are running on 32-bit or 64-bit systems, and if they are virtualized images. Performance will also differ between earlier and later versions of the Windows operating system.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeLockMemoryPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

Best practices are dependent on the platform architecture and the applications running on those platforms.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic. However, recommendations for its implementation differ depending on if applications are running on 32-bit or 64-bit systems, and if they are virtualized images.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users with the **Lock pages in memory** user right could assign physical memory to several processes, which could leave little or no RAM for other processes and result in a denial-of-service condition.

Countermeasure

Do not assign the **Lock pages in memory** user right to any accounts.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Log on as a batch job

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which accounts can log on by using a batch-queue tool such as the Task Scheduler service. When an administrator uses the Add Scheduled Task Wizard to schedule a task to run under a particular user name and password, that user is automatically assigned the **Log on as a batch job** user right. When the scheduled time arrives, the Task Scheduler service logs on the user as a batch job instead of as an interactive user, and the task runs in the user's security context.

This policy setting is provides compatibility with older versions of Windows.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeBatchLogonRight

Possible values

- User-defined list of accounts
- Default values
- Not Defined

Best practices

1. Use discretion when assigning this right to specific users for security reasons. The default settings are sufficient in most cases.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this setting is for Administrators, Backup Operators, and Performance Log Users on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators Backup Operators Performance Log Users
Stand-Alone Server Default Settings	Administrators Backup Operators Performance Log Users
Domain Controller Effective Default Settings	Administrators Backup Operators Performance Log Users
Member Server Effective Default Settings	Administrators Backup Operators Performance Log Users
Client Computer Effective Default Settings	Administrators

Operating system version differences

In Windows Server 2008 R2, the Performance Log Users group was added as a default group for this policy setting. No other changes were made in the way this policy setting works between the supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Task Scheduler automatically grants this right when a user schedules a task. To override this behavior use the Deny log on as a batch job User Rights Assignment setting.

Group Policy settings are applied in the following order, which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Log on as a batch job** user right presents a low-risk vulnerability. For most organizations, the default settings are sufficient. Members of the local Administrators group have this right by default.

Countermeasure

You should allow the computer to manage this user right automatically if you want to allow scheduled tasks to run for specific user accounts. If you do not want to use the Task Scheduler in this manner, configure the **Log on as a batch job** user right for only the Local Service account.

For IIS servers, you should configure this policy locally instead of through domain–based Group Policy settings so that you can ensure the local IUSR_<*ComputerName*> and IWAM_<*ComputerName*> accounts have this user right.

Potential impact

If you configure the **Log on as a batch job** setting by using domain-based Group Policy settings, the computer cannot assign the user right to accounts that are used for scheduled jobs in the Task Scheduler. If you install optional components such as ASP.NET or IIS, you may need to assign this user right to additional accounts that are required by those components. For example, IIS requires assignment of this user right to the IIS_WPG group and the IUSR_<*ComputerName*>, ASPNET, and IWAM_<*ComputerName*> accounts. If this user right is not assigned to this group and these accounts, IIS cannot run some COM objects that are necessary for proper functionality.

See Also

User Rights Assignment

Log on as a service

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which service accounts can register a process as a service. Running a process under a service account circumvents the need for human intervention.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeServiceLogonRight

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Minimize the number of accounts that are granted this user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Network Service on domain controllers and Network Service on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Network Service
Member Server Effective Default Settings	Network Service
Client Computer Effective Default Settings	Network Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

The default changed in Windows Server 2008 R2 and Windows 7 from Not defined in that only the Network Service account has this right by default. Any service that runs under a separate user account must be assigned this user right.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

The policy setting **Deny logon as a service** supersedes this policy setting if a user account is subject to both policies.

Group Policy settings are applied in the following order, which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Log on as a service** user right allows accounts to start network services or services that run continuously on a computer, even when no one is logged on to the console. The risk is reduced by the fact that only users with administrative privileges can install and configure services. An attacker who has already attained that level of access could configure the service to run with the Local System account.

Countermeasure

By definition, the Network Service account has the **Log on as a service** user right. This right is not granted through the Group Policy setting. You should minimize the number of other accounts that are granted this user right.

Potential impact

On most computers, restricting the **Log on as a service** user right to the Local System, Local Service, and Network Service built-in accounts is the default configuration, and there is no negative impact. However, if you have installed optional components such as ASP.NET or IIS, you may need to assign the **Log on as a service** user right to additional accounts that are required by those components. IIS requires that this user right be explicitly granted to the ASPNET user account.

See Also

User Rights Assignment

Manage auditing and security log

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can specify object access audit options for individual resources such as files, Active Directory objects, and registry keys. These objects specify their system access control lists (SACL). A user who is assigned this user right can also view and clear the Security log in Event Viewer. For more information about the Object Access audit policy, see **Object Access**.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeSecurityPrivilege

Possible values

- User-defined list of accounts
- Administrators
- Not Defined

Best practices

1. Before removing this right from a group, investigate whether applications are dependent on this right.
2. Generally, assigning this user right to groups other than Administrators is not necessary.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators

Server type or GPO	Default value
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Audits for object access are not performed unless you enable them by using the Local Group Policy Editor, the Group Policy Management Console (GPMC), or the Auditpol command-line tool.

For more information about the Object Access audit policy, see **Object Access**. For information about Auditpol, see the Command-Line Reference topic **Auditpol**.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Anyone with the **Manage auditing and security log** user right can clear the Security log to erase important evidence of unauthorized activity.

Countermeasure

Ensure that only the local Administrators group has the **Manage auditing and security log** user right.

Potential impact

Restricting the **Manage auditing and security log** user right to the local Administrators group is the default configuration.

Warning

If groups other than the local Administrators group have been assigned this user right, removing this user right might cause performance issues with other applications. Before removing this right from a group, investigate whether applications are dependent on this right.

See Also

User Rights Assignment

Modify an object label

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This privilege determines which user accounts can modify the integrity label of objects, such as files, registry keys, or processes owned by other users. Processes running under a user account can modify the label of an object owned by that user to a lower level without this privilege.

The integrity label is used by the Windows Integrity Controls (WIC) feature, which was introduced in Windows Server 2008 and Windows Vista. WIC keeps lower integrity processes from modifying higher integrity processes by assigning one of six possible labels to objects on the system.

Although similar to NTFS file and folder permissions, which are discretionary controls on objects, the WIC integrity levels are mandatory controls that are put in place and enforced by the operating system. The following list describes the integrity levels from lowest to highest:

- **Untrusted** Default assignment for processes that are logged on anonymously.
- **Low** Default assignment for processes that interact with the Internet.
- **Medium** Default assignment for standard user accounts and any object that is not explicitly designated with a lower or higher integrity level.
- **High** Default assignment for administrator accounts and processes that request to run using administrative rights.
- **System** Default assignment for Windows kernel and core services.
- **Installer** Used by setup programs to install software. It is important that only trusted software is installed on computers because objects that are assigned the Installer integrity level can install, modify, and uninstall all other objects.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeRelabelPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Do not give any group this user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Not defined on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined

Server type or GPO	Default value
Domain Controller Effective Default Settings	Not defined
Member Server Effective Default Settings	Not defined
Client Computer Effective Default Settings	Not defined

Operating system version differences

This policy setting was introduced in Windows Server 2008 and Windows Vista and has no effect on earlier versions. There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Anyone with the **Modify an object label** user right can change the integrity level of a file or process so that it becomes elevated or decreased to a point where it can be deleted by lower

integrity processes. Either of these states effectively circumvents the protection that is offered by Windows Integrity Controls and makes your system vulnerable to attacks by malicious software.

If malicious software is set with an elevated integrity level such as Trusted Installer or System, administrator accounts do not have sufficient integrity levels to delete the program from the system. In that case, use of the **Modify an object label** right is mandated so that the object can be relabeled. However, the relabeling must occur by using a process that is at the same or a higher level of integrity than the object that you are attempting to relabel.

Countermeasure

Do not give any group this right. If necessary, implement it for a constrained period of time to a trusted individual to respond to a specific organizational need.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Modify firmware environment values

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines who can modify firmware environment values. Firmware environment values are settings that are stored in the nonvolatile RAM of non-x86-based computers. The effect of the setting depends on the processor.

On x86-based computers, the only firmware environment value that can be modified by assigning this user right is the **Last Known Good Configuration** setting, which should only be modified by the system.

On Itanium-based computers, boot information is stored in nonvolatile RAM. Users must be assigned this user right to run bootcfg.exe and to change the **Default Operating System** setting using the **Startup and Recovery** feature on the **Advanced** tab of **System Properties**.

The exact setting for firmware environment values is determined by the boot firmware. The location of these values is also specified by the firmware. For example, on a UEFI-based system, NVRAM contains firmware environment values that specify system boot settings.

On all computers, this user right is required to install or upgrade Windows.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeSystemEnvironmentPrivilege

Possible values

- User-defined list of accounts
- Administrators
- Not Defined

Best practices

1. Ensure that only the local Administrators group is assigned the **Modify firmware environment values** user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Adminstrators
Stand-Alone Server Default Settings	Adminstrators
Domain Controller Effective Default Settings	Adminstrators
Member Server Effective Default Settings	Adminstrators
Client Computer Effective Default Settings	Adminstrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

This security setting does not affect who can modify the system environment values and user environment values that are displayed on the **Advanced** tab of **System Properties**.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Anyone who is assigned the **Modify firmware environment values** user right could configure the settings of a hardware component to cause it to fail, which could lead to data corruption or a denial-of-service condition.

Countermeasure

Ensure that only the local Administrators group is assigned the **Modify firmware environment values** user right.

Potential impact

None. Restricting the **Modify firmware environment values** user right to the members of the local Administrators group is the default configuration.

See Also

User Rights Assignment

Perform volume maintenance tasks

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can perform volume or disk management tasks, such as defragmenting an existing volume, creating or removing volumes, and running the Disk Cleanup tool.

Use caution when assigning this user right. Users with this user right can explore disks and extend files in to memory that contains other data. When the extended files are opened, the user might be able to read and modify the acquired data.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeManageVolumePrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Ensure that only the local Administrators group is assigned the **Perform volume maintenance tasks** user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators

Server type or GPO	Default value
Stand-Alone Server Default Settings	Administrators
DC Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

A user who is assigned the **Perform volume maintenance tasks** user right could delete a volume, which could result in the loss of data or a denial-of- service condition. Also, disk maintenance tasks can be used to modify data on the disk, such as user rights assignments that might lead to escalation of privileges.

Countermeasure

Ensure that only the local Administrators group is assigned the **Perform volume maintenance tasks** user right.

Potential impact

None. Restricting the **Perform volume maintenance tasks** user right to the local Administrators group is the default configuration.

See Also

User Rights Assignment

Profile single process

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can view a sample performance of an application process. Typically, you do not need this user right to use the performance reporting tools included in the operating system. However, you do need this user right if the system's monitor components are configured to collect data through Windows Management Instrumentation (WMI).

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeProfileSingleProcessPrivilege

Possible values

- User-defined list of accounts
- Administrators
- Not Defined

Best practices

1. This right should not be granted to individual users. It should be granted only for trusted applications that monitor other programs.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic. Note that the Power Users group is assigned for backward compatibility to Windows XP.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Profile single process** user right presents a moderate vulnerability. Attackers with this user right could monitor a computer's performance to help identify critical processes that they might want to attack directly. Attackers may be able to determine what processes run on the computer so that they could identify countermeasures that they may need to avoid, such as antivirus software or an intrusion-detection system. They could also identify other users who are logged on to a computer.

Countermeasure

Ensure that only the local Administrators group is assigned the **Profile single process** user right.

Potential impact

If you remove the **Profile single process** user right from the Power Users group or other accounts, you could limit the abilities of users who are assigned to specific administrative roles in your environment. You should ensure that delegated tasks are not negatively affected.

See Also

User Rights Assignment

Profile system performance

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines which users can use Windows performance monitoring tools to monitor the performance of system processes.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeSystemProfilePrivilege

Possible values

- User-defined list of accounts
- Administrators
- Not defined

Best practices

1. Ensure that only the local Administrators group is assigned the **Profile system performance** user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Depending on your version of Windows and your environment, you might need to add this user right to the Local System account or the Local Service account if you encounter access errors when you use the Administrators account.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Profile system performance** user right poses a moderate vulnerability. Attackers with this user right could monitor a computer's performance to help identify critical processes that they might want to attack directly. Attackers might also be able to determine what processes are active on the computer so that they could identify countermeasures to avoid, such as antivirus software or an intrusion detection system.

Countermeasure

Ensure that only the local Administrators group is assigned the **Profile system performance** user right.

Potential impact

None. Restricting the **Profile system performance** user right to the local Administrators group is the default configuration.

See Also

User Rights Assignment

Remove computer from docking station

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines whether a user can undock a portable computer from its docking station without logging on. This policy setting only affects scenarios that involve a portable computer and its docking station.

If this user right is assigned to the user's account (or if the user is a member of the assigned group), the user must log on before removing the portable computer from its docking station. Otherwise, as a security measure, the user will not be able to log on after the computer is removed from the docking station. If this policy is not assigned, the user may remove the portable computer from its docking station without logging on, and then have the ability to start and log on to the computer afterwards in its undocked state.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeUndockPrivilege

Possible values

- User-defined list of accounts
- Not Defined

Best practices

1. Assign this user right to only those accounts that are permitted to use the portable computer.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

Although this portable computer scenario does not normally apply to servers, by default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic. However, the groups Users and Power Users were added as defaults for backwards compatibility to Windows Vista.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Anyone who has the **Remove computer from docking station** user right can log on and then remove a portable computer from its docking station. If this setting is not defined, it has the same effect as if everyone was granted this right. However, the value of implementing this countermeasure is reduced by the following factors:

- If attackers can restart the computer, they could remove it from the docking station after the BIOS starts but before the operating system starts.
- This setting does not affect servers because they typically are not installed in docking stations.
- An attacker could steal the computer and the docking station together.
- Computers that can be mechanically undocked can be physically removed by the user whether or not they use the Windows undocking functionality.

Countermeasure

Ensure that only the local Administrators group and the user account to which the computer is allocated are assigned the **Remove computer from docking station** user right.

Potential impact

By default, only members of the local Administrators group are granted this right. Other user accounts must be explicitly granted this user right as necessary. If your organization's users are not members of the local Administrators groups on their portable computers, they cannot remove their portable computers from their docking stations if they do not first shut down the computer. Therefore, you may want to assign the **Remove computer from docking station** privilege to the local Users group for portable computers.

See Also

User Rights Assignment

Replace a process level token

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which parent processes can replace the access token that is associated with a child process.

Specifically, the **Replace a process level token** setting determines which user accounts can call the CreateProcessAsUser() application programming interface (API) so that one service can start

another. An example of a process that uses this user right is Task Scheduler, where the user right is extended to any processes that can be managed by Task Scheduler.

An access token is an object that describes the security context of a process or thread. The information in a token includes the identity and privileges of the user account that is associated with the process or thread. With this user right, every child process that runs on behalf of this user account would have its access token replaced with the process level token.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeAssignPrimaryTokenPrivilege

Possible values

- User-defined list of accounts
- Defaults
- Not defined

Best practices

1. For member servers, ensure that only the Local Service and Network Service accounts have the **Replace a process level token** user right.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Network Service and Local Service on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Network Service Local Service
Stand-Alone Server Default Settings	Network Service Local Service
Domain Controller Effective Default Settings	Network Service Local Service

Server type or GPO	Default value
Member Server Effective Default Settings	Network Service Local Service
Client Computer Effective Default Settings	Network Service Local Service

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Users with the **Replace a process level token** user right can start processes as another user if they know the user's credentials.

Countermeasure

For member servers, ensure that only the Local Service and Network Service accounts have the **Replace a process level token** user right.

Potential impact

On most computers, restricting the **Replace a process level token** user right to the Local Service and the Network Service built-in accounts is the default configuration, and there is no negative impact. However, if you have installed optional components such as ASP.NET or IIS, you may need to assign the **Replace a process level token** user right to additional accounts. For example, IIS requires that the Service, Network Service, and IWAM_<*ComputerName*> accounts be explicitly granted this user right.

See Also

User Rights Assignment

Restore files and directories

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines which users can bypass file, directory, registry, and other persistent object permissions when they restore backed up files and directories, and it determines which users can set valid security principals as the owner of an object.

Granting this user right to an account is similar to granting the account the following permissions to all files and folders on the system:

- **Traverse folder / execute file**
- **Write**

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeRestorePrivilege

Possible values

- User-defined list of accounts
- Defaults
- Not Defined

Best practices

1. Users with this user right can overwrite registry settings, hide data, and gain ownership of system objects, so only assign this user right to trusted users.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default, this right is granted to the Administrators, Backup Operators, and Server Operators groups on domain controllers, and to the Administrators and Backup Operators groups on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	
Default Domain Controller Policy	Administrators Backup Operators Server Operators
Stand-Alone Server Default Settings	Administrators Backup Operators
Domain Controller Effective Default Settings	Administrators Backup Operators Server Operators
Member Server Effective Default Settings	Administrators Backup Operators
Client Computer Effective Default Settings	Administrators Backup Operators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

An attacker with the **Restore files and directories** user right could restore sensitive data to a computer and overwrite data that is more recent, which could lead to loss of important data, data corruption, or a denial-of-service condition. Attackers could overwrite executable files that are used by legitimate administrators or system services with versions that include malicious software to grant themselves elevated privileges, compromise data, or install programs that provide continued access to the computer.

Note

Even if the following countermeasure is configured, an attacker could restore data to a computer in a domain that is controlled by the attacker. Therefore, it is critical that organizations carefully protect the media that are used to back up data.

Countermeasure

Ensure that only the local Administrators group is assigned the **Restore files and directories** user right unless your organization has clearly defined roles for backup and for restore personnel.

Potential impact

If you remove the **Restore files and directories** user right from the Backup Operators group and other accounts, users who are not members of the local Administrators group cannot load data backups. If restoring backups is delegated to a subset of IT staff in your organization, you should verify that this change does not negatively affect the ability of your organization's personnel to do their jobs.

See Also

User Rights Assignment

Shut down the system

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This security setting determines if a user who is logged on locally to a computer can shut down the operating system by using the Shut Down command.

Shutting down domain controllers makes them unavailable to perform functions such as processing logon requests, processing Group Policy settings, and answering Lightweight Directory Access Protocol (LDAP) queries. Shutting down domain controllers that have been assigned operations master roles (also known as flexible single master operations or FSMO roles) can disable key domain functionality; for example, processing logon requests for new passwords, which is performed by the primary domain controller (PDC) emulator master.

The **Shut down the system** user right is required to enable hibernation support, to set the power management settings, and to cancela shutdown.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeShutdownPrivilege

Possible values

- A user-defined list of accounts
- Defaults
- Not defined

Best practices

1. Ensure that only Administrators and Backup Operators have the **Shut down the system** user right on member servers, and that only Administrators have the user right on domain controllers. Removing these default groups might limit the abilities of users who are assigned to specific administrative roles in your environment. Ensure that their delegated tasks will not be negatively affected.
2. The ability to shut down domain controllers should be limited to a very small number of trusted administrators. Even though a system shutdown requires the ability to log on to the server, you should be very careful about the accounts and groups that you allow to shut down a domain controller.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators, Backup Operators, Server Operators, and Print Operators on domain controllers, and Administrators and Backup Operators on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Administrators Backup Operators Server Operators Print Operators
Stand-Alone Server Default Settings	Administrators Backup Operators
Domain Controller Effective Default Settings	Administrators Backup Operators Server Operators Print Operators
Member Server Effective Default Settings	Administrators Backup Operators
Client Computer Effective Default Settings	Administrators

Server type or GPO	Default value
	Backup Operators Users

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic. However, the setting is not compatible with computers running Microsoft Windows 2000 Server with Service Pack 1 or earlier.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Do not apply this setting to computers running Microsoft Windows 2000 Server with Service Pack 1 or earlier.

This user right does not have the same effect as **Force shutdown from a remote system**. For more information, see Force shutdown from a remote system.

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The ability to shut down domain controllers should be limited to a very small number of trusted administrators. Although the **Shut down the system** user right requires the ability to log on to the server, you should be very careful about which accounts and groups you allow to shut down a domain controller.

When a domain controller is shut down, it is no longer available to process logon requests, process Group Policy settings, and answer Lightweight Directory Access Protocol (LDAP) queries. If you shut down domain controllers that possess operations master roles, you can disable key domain functionality, such as processing logon requests for new passwords, which is performed by the PDC master.

For other server roles, especially those where non-administrators have rights to log on to the server (such as RD Session Host servers), it is critical that this user right be removed from users that do not have a legitimate reason to restart the servers.

Countermeasure

Ensure that only the Administrators and Backup Operators groups are assigned the **Shut down the system** user right on member servers, and ensure that only the Administrators group is assigned the user right on domain controllers.

Potential impact

The impact of removing these default groups from the **Shut down the system** user right could limit the delegated abilities of assigned roles in your environment. You should confirm that delegated activities are not adversely affected.

See Also

User Rights Assignment

Synchronize directory service data

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users and groups have authority to synchronize all directory service data, regardless of the protection for objects and properties. This privilege is required to use LDAP directory synchronization (dirsync) services. Domain controllers have this user right inherently because the synchronization process runs in the context of the **System** account on domain controllers.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeSyncAgentPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. Ensure that no accounts are assigned the **Synchronize directory service data** user right. Only domain controllers need this privilege, which they inherently have.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is not defined on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined
Default Domain Controller Policy	Not defined
Stand-Alone Server Default Settings	Not defined
Domain Controller Effective Default Settings	Enabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

The **Synchronize directory service data** user right affects domain controllers (only domain controllers should be able to synchronize directory service data). Domain controllers have this user right inherently because the synchronization process runs in the context of the **System** account on domain controllers. Attackers who have this user right can view all information that is stored within the directory. They could then use some of that information to facilitate additional attacks or expose sensitive data, such as direct telephone numbers or physical addresses.

Countermeasure

Ensure that no accounts are assigned the **Synchronize directory service data** user right.

Potential impact

None. Not defined is the default configuration.

See Also

User Rights Assignment

Take ownership of files or other objects

This security policy reference topic for the IT professional describes the best practices, location, values, policy management, and security considerations for this policy setting.

Reference

This policy setting determines which users can take ownership of any securable object in the computer, including Active Directory objects, NTFS files and folders, printers, registry keys, services, processes, and threads.

Every object has an owner, whether the object resides in an NTFS volume or Active Directory database. The owner controls how permissions are set on the object and to whom permissions are granted.

By default, the owner is the person who or the process which created the object. Owners can always change permissions to objects, even when they are denied all access to the object.

This policy setting is supported on versions of Windows that are designated in the **Applies To** list at the beginning of this topic.

Constant: SeTakeOwnershipPrivilege

Possible values

- User-defined list of accounts
- Not defined

Best practices

1. Assigning this user right can be a security risk. Because owners of objects have full control of them, only assign this user right to trusted users.

Location

GPO_name\Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

Default values

By default this setting is Administrators on domain controllers and on stand-alone servers.

The following table lists the actual and effective default policy values for the most recent supported versions of Windows. Default values are also listed on the policy's property page.

Server type or GPO	Default value
Default Domain Policy	Not defined

Server type or GPO	Default value
Default Domain Controller Policy	Administrators
Stand-Alone Server Default Settings	Administrators
Domain Controller Effective Default Settings	Administrators
Member Server Effective Default Settings	Administrators
Client Computer Effective Default Settings	Administrators

Operating system version differences

There are no differences in the way this policy setting works between the supported versions of Windows that are designated in the**Applies To** list at the beginning of this topic.

Policy management

This section describes features, tools, and guidance to help you manage this policy.

A restart of the computer is not required for this policy setting to be effective.

Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Ownership can be taken by:

- An administrator. By default, the Administrators group is given the **Take ownership of files or other objects** user right.
- Anyone or any group who has the **Take ownership** user right on the object.
- A user who has the **Restore files and directories** user right.

Ownership can be transferred in the following ways:

- The current owner can grant the **Take ownership** user right to another user if that user is a member of a group defined in the current owner's access token. The user must take ownership to complete the transfer.
- An administrator can take ownership.
- A user who has the **Restore files and directories** user right can double-click **Other users and groups** and choose any user or group to assign ownership to.

Group Policy

Settings are applied in the following order through a Group Policy Object (GPO), which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings

4. OU policy settings

When a local setting is greyed out, it indicates that a GPO currently controls that setting.

Security considerations

This section describes how an attacker might exploit a feature or its configuration, how to implement the countermeasure, and the possible negative consequences of countermeasure implementation.

Vulnerability

Any users with the **Take ownership of files or other objects user right** can take control of any object, regardless of the permissions on that object, and then make any changes that they want to make to that object. Such changes could result in exposure of data, corruption of data, or a denial-of-service condition.

Countermeasure

Ensure that only the local Administrators group has the **Take ownership of files or other objects** user right.

Potential impact

None. Restricting the **Take ownership of files or other objects** user right to the local Administrators group is the default configuration.

See Also

User Rights Assignment

Administer Security Policy Settings

This topic for the IT professional discusses different methods to administer security policy settings on the local computer or throughout a small- or medium-sized IT organization by using Windows Server 2012 and Windows 8.

Introduction

Security policy settings should be used as part of your overall security implementation to help secure domain controllers, servers, client computers, and other resources in your organization.

Security settings policies are rules that you can configure on a computer, or multiple computers, for the purpose of protecting resources on a computer or network. The Security Settings extension of the Local Group Policy Editor snap-in (Gpedit.msc) allows you to define security

configurations as part of a Group Policy Object (GPO). The GPOs are linked to Active Directory containers such as sites, domains, and organizational units, and they enable administrators to manage security settings for multiple computers from any computer joined to the domain.

Security settings can control:

- User authentication to a network or computer.
- The resources that users are permitted to access.
- Whether to record a user's or group's actions in the Event log.
- Membership in a group.

For information about each setting, including descriptions, default settings, and management and security considerations, see the Security Policy Settings Reference in the Microsoft Download Center.

For information about how the security settings snap-ins work, see the Security Policy Settings Technical Overview.

To manage security configurations for multiple computers, you can use one of the following options:

- Edit specific security settings in a GPO.
- Use the Security Templates snap-in to create a security template that contains the security policies you want to apply, and then import the security template into a Group Policy Object. A security template is a file that represents a security configuration, and it can be imported to a GPO, or applied to a local computer, or it can be used to analyze security.

What's changed in how settings are administered?

Over time, new ways to manage security policy settings have been introduced, which include new operating system features and the addition of new settings. The following table lists different means by which security policy settings can be administered.

Tool or feature	Description and use
Using the Local Security Policy snap-in	Secpol.msc MMC snap-in designed to manage only security policy settings.
Using the Secedit command-line tool	Secedit.exe Configures and analyzes system security by comparing your current configuration to specified security templates.
Using the Security Compliance Manager	Tool download A Solution Accelerator that helps you plan,

Tool or feature	Description and use
	deploy, operate, and manage your security baselines for Windows client and server operating systems, and Microsoft applications.
Using the Security Configuration Wizard	Scw.exe SCW is a role-based tool available on servers only: You can use it to create a policy that enables services, firewall rules, and settings that are required for a selected server to perform specific roles.
Working with the Security Configuration Manager	This tool set allows you to create, apply, and edit the security for your local computer, organizational unit, or domain.
Working with Group Policy tools	Gpmc.msc and Gpedit.msc The Group Policy Management Console uses the Group Policy Object editor to expose the local Security options, which can then be incorporated into Group Policy Objects for distribution throughout the domain. The Local Group Policy Editor performs similar functions on the local computer.
Software Restriction Policies See Administer Software Restriction Policies in the TechNet Library.	Gpedit.msc Software Restriction Policies (SRP) is a Group Policy-based feature that identifies software programs running on computers in a domain, and it controls the ability of those programs to run.
AppLocker See Administer AppLocker in the TechNet Library.	Gpedit.msc Prevents malicious software (malware) and unsupported applications from affecting computers in your environment, and it prevents users in your organization from installing and using unauthorized applications.

Using the Local Security Policy snap-in

The Local Security Policy snap-in (Secpol.msc) restricts the view of local policy objects to the following policies and features:

- Account Policies
- Local Policies
- Windows Firewall with Advanced Security
- Network List Manager Policies
- Public Key Policies
- Software Restriction Policies
- Application Control Policies
- IP Security Policies on Local Computer
- Advanced Audit Policy Configuration

Policies set locally might be overwritten if the computer is joined to the domain.

The Local Security Policy snap-in is part of the Security Configuration Manager tool set. For information about other tools in this tool set, see Working with the Security Configuration Manager in this topic.

Using the Secedit command-line tool

The Secedit command-line tool works with security templates and provides six primary functions:

- The **Configure** parameter helps you resolve security discrepancies between servers by applying the correct security template to the errant server.
- The **Analyze** parameter compares the server's security configuration with the selected template.
- The **Import** parameter allows you to create a database from an existing template. The Security Configuration and Analysis tool does this also.
- The **Export** parameter allows you to export the settings from a database into a security settings template.
- The **Validate** parameter allows you to validate the syntax of each or any lines of text that you created or added to a security template. This ensures that if the template fails to apply syntax, the template will not be the issue.
- The **Generate Rollback** parameter saves the server's current security settings into a security template so it can be used to restore most of the server's security settings to a known state. The exceptions are that, when applied, the rollback template will not change access control list entries on files or registry entries that were changed by the most recently applied template.

For more information about the Secedit.exe tool, see **Secedit**.

Using the Security Compliance Manager

The Security Compliance Manager is a downloadable tool that helps you plan, deploy, operate, and manage your security baselines for Windows client and server operating systems, and for Microsoft applications. It contains a complete database of recommended security settings, methods to customize your baselines, and the option to implement those settings in multiple

formats—including XLS, GPOs, Desired Configuration Management (DCM) packs, or Security Content Automation Protocol (SCAP). The Security Compliance Manager is used to export the baselines to your environment to automate the security baseline deployment and compliance verification process.

To administer security policies by using the Security Compliance Manager

1. Download the most recent version of the Security Compliance Manager in the Microsoft Download Center.
2. Read the relevant security baseline documentation that is included in this tool.
3. Download and import the relevant security baselines. The installation process steps you through baseline selection.
4. Open the Help and follow instructions how to customize, compare, or merge your security baselines before deploying those baselines.

Using the Security Configuration Wizard

The Security Configuration Wizard (SCW) guides you through the process of creating, editing, applying, or rolling back a security policy. A security policy that you create with SCW is an .xml file that, when applied, configures services, network security, specific registry values, and audit policy. SCW is a role-based tool: You can use it to create a policy that enables services, firewall rules, and settings that are required for a selected server to perform specific roles. For example, a server might be a file server, a print server, or a domain controller.

The following are considerations for using SCW:

- SCW disables unnecessary services and provides Windows Firewall with Advanced Security support.
- Security policies that are created with SCW are not the same as security templates, which are files with an .inf extension. Security templates contain more security settings than those that can be set with SCW. However, it is possible to include a security template in an SCW security policy file.
- You can deploy security policies that you create with SCW by using Group Policy.
- SCW does not install or uninstall the features necessary for the server to perform a role. You can install role-specific features through Server Manager.
- SCW detects role dependencies. If you select a role, it automatically selects dependent roles.
- All applications that use the IP protocol and ports must be running on the server when you run SCW.
- In some cases, you must be connected to the Internet to use the links in the SCW Help.

Note

The SCW is available only on Windows Server and only applicable to server installations.

The SCW can be accessed through Server Manager or by running Scw.exe. The wizard steps you through server security configuration to:

- Create a security policy that can be applied to any server on your network.

- Edit an existing security policy.
- Apply an existing security policy.
- Roll back the last applied security policy.

The Security Policy Wizard configures services and network security based on the server's role, as well as configures auditing and registry settings.

For more information about SCW, including procedures, see the Security Configuration Wizard.

Working with the Security Configuration Manager

The Security Configuration Manager tool set allows you to create, apply, and edit the security for your local computer, organizational unit, or domain.

For procedures on how to use the Security Configuration Manager, see Security Configuration Manager.

The following table lists the features of the Security Configuration Manager.

Security Configuration Manager tools	Description
Security Configuration and Analysis	Defines a security policy in a template. These templates can be applied to Group Policy or to your local computer.
Security templates	Defines a security policy in a template. These templates can be applied to Group Policy or to your local computer.
Security Settings extension to Group Policy	Edits individual security settings on a domain, site, or organizational unit.
Local Security Policy	Edits individual security settings on your local computer.
Secedit	Automates security configuration tasks at a command prompt.

Security Configuration and Analysis

Security Configuration and Analysis is an MMC snap-in for analyzing and configuring local system security.

Security analysis

The state of the operating system and applications on a computer is dynamic. For example, you may need to temporarily change security levels so that you can immediately resolve an administration or network issue. However, this change can often go unreversed. This means that a computer may no longer meet the requirements for enterprise security.

Regular analysis enables an administrator to track and ensure an adequate level of security on each computer as part of an enterprise risk management program. An administrator can tune the security levels and, most importantly, detect any security flaws that may occur in the system over time.

Security Configuration and Analysis enables you to quickly review security analysis results. It presents recommendations alongside of current system settings and uses visual flags or remarks to highlight any areas where the current settings do not match the proposed level of security. Security Configuration and Analysis also offers the ability to resolve any discrepancies that analysis reveals.

Security configuration

Security Configuration and Analysis can also be used to directly configure local system security. Through its use of personal databases, you can import security templates that have been created with Security Templates and apply these templates to the local computer. This immediately configures the system security with the levels specified in the template.

Security templates

With the Security Templates snap-in for Microsoft Management Console, you can create a security policy for your computer or for your network. It is a single point of entry where the full range of system security can be taken into account. The Security Templates snap-in does not introduce new security parameters, it simply organizes all existing security attributes into one place to ease security administration.

Importing a security template to a Group Policy Object eases domain administration by configuring security for a domain or organizational unit at once.

To apply a security template to your local computer, you can use Security Configuration and Analysis or the Secedit command-line tool.

Security templates can be used to define:

- Account Policies
  - Password Policy
  - Account Lockout Policy
  - Kerberos Policy
- Local Policies
  - Audit Policy
  - User Rights Assignment
  - Security Options
- Event Log: Application, system, and security Event Log settings
- Restricted Groups: Membership of security-sensitive groups
- System Services: Startup and permissions for system services
- Registry: Permissions for registry keys

- File System: Permissions for folders and files

Each template is saved as a text-based .inf file. This enables you to copy, paste, import, or export some or all of the template attributes. With the exceptions of Internet Protocol security and public key policies, all security attributes can be contained in a security template.

Security Settings extension to Group Policy

Organizational units, domains, and sites are linked to Group Policy Objects. The Security Settings tool allows you change the security configuration of the Group Policy Object, in turn, affecting multiple computers. With Security Settings, you can modify the security settings of many computers, depending on the Group Policy Object you modify, from just one computer joined to a domain.

Security settings or security policies are rules that are configured on a computer or multiple computers for protecting resources on a computer or network. Security settings can control:

- How users are authenticated to a network or computer.
- What resources users are authorized to use.
- Whether or not a user's or group's actions are recorded in the event log.
- Group membership.

You can change the security configuration on multiple computers in two ways:

- Create a security policy by using a security template with Security Templates, and then import the template through Security Settings to a Group Policy Object.
- Change a few select settings with Security Settings.

Local Security Policy

A security policy is a combination of security settings that affect the security on a computer. You can use your local security policy to edit account policies and local policies on your local computer.

With the local security policy, you can control:

- Who accesses your computer.
- What resources users are authorized to use on your computer.
- Whether or not a user’s or group's actions are recorded in the event log.

If your local computer is joined to a domain, you are subject to obtaining a security policy from the domain's policy or from the policy of any organizational unit that you are a member of. If you are getting a policy from more than one source, conflicts are resolved in the following order of precedence.

1. Organizational unit policy
2. Domain policy
3. Site policy
4. Local computer policy

If you modify the security settings on your local computer by using the local security policy, then you are directly modifying the settings on your computer. Therefore, the settings take effect immediately, but this may only be temporary. The settings will actually remain in effect on your local computer until the next refresh of Group Policy security settings, when the security settings that are received from Group Policy will override your local settings wherever there are conflicts.

Using the Security Configuration Manager

For procedures on how to use the Security Configuration Manager, see Security Configuration Manager How To. This section contains information in this topic about:

- Applying security settings
- Importing and exporting security templates
- Analyzing security and viewing results
- Resolving security discrepancies
- Automating security configuration tasks

Applying security settings

Once you have edited the security settings, the settings are refreshed on the computers in the organizational unit linked to your Group Policy Object:

- When a computer is restarted, the settings on that computer will be refreshed.
- To force a computer to refresh its security settings as well as all Group Policy settings, see the **Gpupdate** command-line tool.

Precedence of a policy when more than one policy is applied to a computer

For security settings that are defined by more than one policy, the following order of precedence is observed:

1. Organizational Unit Policy
2. Domain Policy
3. Site Policy
4. Local computer Policy

For example, a workstation that is joined to a domain will have its local security settings overridden by the domain policy wherever there is a conflict. Likewise, if the same workstation is a member of an Organizational Unit, the settings applied from the Organizational Unit's policy will override both the domain and local settings. If the workstation is a member of more than one Organizational Unit, then the Organizational Unit that immediately contains the workstation has the highest order of precedence.

Notes

Use the **Gpresult** command-line tool to find out what policies are applied to a computer and in what order.

For domain accounts, there can be only one account policy that includes password policies, account lockout policies, and Kerberos policies.

Persistence in security settings

Security settings may still persist even if a setting is no longer defined in the policy that originally applied it.

Persistence in security settings occurs when:

- The setting has not been previously defined for the computer.
- The setting is for a registry object.
- The setting is for a file system object.

All settings applied through local policy or a Group Policy Object are stored in a local database on your computer. Whenever a security setting is modified, the computer saves the security setting value to the local database, which retains a history of all the settings that have been applied to the computer. If a policy first defines a security setting and then no longer defines that setting, then the setting takes on the previous value in the database. If a previous value does not exist in the database, then the setting does not revert to anything and remains defined as is. This behavior is sometimes called "tattooing."

Registry and file settings will maintain the values applied through policy until that setting is set to other values.

Filtering security settings based on group membership

You can also decide what users or groups will or will not have a Group Policy Object applied to them regardless of what computer they have logged onto by denying them either the Apply Group Policy or Read permission on that Group Policy Object. Both of these permissions are needed to apply Group Policy.

Importing and exporting security templates

Security Configuration and Analysis provides the ability to import and export security templates into or from a database.

If you have made any changes to the analysis database, you can save those settings by exporting them into a template. The export feature provides the ability to save the analysis database settings as a new template file. This template file can then be used to analyze or configure a system, or it can be imported to a Group Policy Object.

Analyzing security and viewing results

Security Configuration and Analysis performs security analysis by comparing the current state of system security against an *analysis database*. During creation, the analysis database uses at least one security template. If you choose to import more than one security template, the database will merge the various templates and create one composite template. It resolves conflicts in order of import; the last template that is imported takes precedence.

Security Configuration and Analysis displays the analysis results by security area, using visual flags to indicate problems. It displays the current system and base configuration settings for each security attribute in the security areas. To change the analysis database settings, right-click the entry, and then click **Properties**.

Visual flag	Meaning
Red X	The entry is defined in the analysis database and on the system, but the security setting values do not match.
Green check mark	The entry is defined in the analysis database and on the system and the setting values match.
Question mark	The entry is not defined in the analysis database and, therefore, was not analyzed. If an entry is not analyzed, it may be that it was not defined in the analysis database or that the user who is running the analysis may not have sufficient permission to perform analysis on a specific object or area.
Exclamation point	This item is defined in the analysis database, but does not exist on the actual system. For example, there may be a restricted group that is defined in the analysis database but does not actually exist on the analyzed system.
No highlight	The item is not defined in the analysis database or on the system.

If you choose to accept the current settings, the corresponding value in the base configuration is modified to match them. If you change the system setting to match the base configuration, the change will be reflected when you configure the system with Security Configuration and Analysis.

To avoid continued flagging of settings that you have investigated and determined to be reasonable, you can modify the base configuration. The changes are made to a copy of the template.

Resolving security discrepancies

You can resolve discrepancies between analysis database and system settings by:

- Accepting or changing some or all of the values that are flagged or not included in the configuration, if you determine that the local system security levels are valid due to the context (or role) of that computer. These attribute values are then updated in the database and applied to the system when you click **Configure Computer Now**.
- Configuring the system to the analysis database values, if you determine the system is not in compliance with valid security levels.
- Importing a more appropriate template for the role of that computer into the database as the new base configuration and applying it to the system.

Changes to the analysis database are made to the stored template in the database, not to the security template file. The security template file will only be modified if you either return to Security Templates and edit that template or export the stored configuration to the same template file.

You should use **Configure Computer Now** only to modify security areas *not* affected by Group Policy settings, such as security on local files and folders, registry keys, and system services. Otherwise, when the Group Policy settings are applied, it will take precedence over local settings—such as account policies. In general, do not use **Configure Computer Now** when you are analyzing security for domain-based clients, since you will have to configure each client individually. In this case, you should return to Security Templates, modify the template, and reapply it to the appropriate Group Policy Object.

Automating security configuration tasks

By calling the Secedit.exe tool at a command prompt from a batch file or automatic task scheduler, you can use it to automatically create and apply templates, and analyze system security. You can also run it dynamically from a command prompt.

Secedit.exe is useful when you have multiple computers on which security must be analyzed or configured, and you need to perform these tasks during off-hours.

Working with Group Policy tools

Group Policy is an infrastructure that allows you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences. For Group Policy settings that affect only a local computer or user, you can use the Local Group Policy Editor. You can manage Group Policy settings and Group Policy Preferences in an Active Directory Domain Services (AD DS) environment through the Group Policy Management Console (GPMC). Group Policy management tools also are included in the Remote Server Administration Tools pack to provide a way for you to administer Group Policy settings from your desktop.

To get started with managing security settings with Group Policy, see Group Policy Overview.

To begin using the Group Policy Management Console, see the Help for the Group Policy Management Console in the TechNet Library.

To begin using the Local Group Policy Editor, see the Help for the Local Group Policy Editor.

Network List Manager Policies

Network List Manager Policies are security settings that you can use to configure different aspects of how networks are listed and displayed on one computer or on many computers.

To configure Network List Manager Policies for one computer, you can use the Microsoft Management Console (MMC) with the Group Policy Object Editor snap-in, and edit the local computer policy. The Network List Manager Policies are located at the following path in Group Policy Object Editor:

Computer Configuration | Windows Settings | Security Settings | Network List Manager Policies

To configure Network List Manager Policies for many computers, such as for all of the Domain Computers in an Active Directory domain, follow Group Policy documentation to learn how to edit the policies for the object that you require. The path to the Network List Manager Policies is the same as the path listed above.

Policy settings for Network List Manager Policies

The following policy settings are provided for Network List Manager Policies. These policy settings are located in the details pane of the Group Policy Object Editor, in **Network Name**.

Unidentified Networks

This policy setting allows you to configure the **Network Location**, including the location type and the user permissions, for networks that Windows cannot identify due to a network issue or a lack of identifiable characters in the network information received by the operating system from the network. A network location identifies the type of network that a computer is connected to and automatically sets the appropriate firewall settings for that location. You can configure the following items for this policy setting:

- **Location type**. For this item, the following options are available:
  - **Not configured**. If you select this option, this policy setting does not apply a location type to unidentified network connections.
  - **Private**. If you select this option, this policy setting applies a location type of Private to unidentified network connections. A private network, such as a home or work network, is a location type that assumes that you trust the other computers on the network. Do not select this item if there is a possibility that an active, unidentified network is in a public place.
  - **Public**. If you select this option, this policy setting applies a location type of Public to unidentified network connections. A public network, such as a wireless network at an airport or coffee shop, is a location type that assumes that you do not trust the other computers on the network.
- **User permissions**. For this item, the following options are available:
  - **Not configured**. If you select this option, this policy setting does not specify whether users can change the location for unidentified network connections.
  - **User can change location**. If you select this option, this policy setting allows users to change an unidentified network connection location from Private to Public or from Public to Private.

- **User cannot change location**. If you select this option, this policy setting does not allow users to change the location of an unidentified network connection.

Identifying Networks

This policy setting allows you to configure the **Network Location** for networks that are in a temporary state while Windows works to identify the network and location type. A network location identifies the type of network that a computer is connected to and automatically sets the appropriate firewall settings for that location. You can configure the following items for this policy setting:

- **Location type**. For this item, the following options are available:
  - **Not configured**. If you select this option, this policy setting does not apply a location type to network connections that are in the process of being identified by Windows.
  - **Private**. If you select this option, this policy setting applies a location type of Private to network connections that are in the process of being identified. A private network, such as a home or work network, is a location type that assumes that you trust the other computers on the network. Do not select this item if there is a possibility that an active, unidentified network is in a public place.
  - **Public**. If you select this option, this policy setting applies a location type of Public to network connections that are in the process of being identified by Windows. A public network, such as a wireless network at an airport or coffee shop, is a location type that assumes that you do not trust the other computers on the network.

All Networks

This policy setting allows you to specify the **User Permissions** that control whether users can change the network name, location, or icon, for all networks to which the user connects. You can configure the following items for this policy setting:

- **Network name**. For this item, the following options are available:
  - **Not configured**. If you select this option, this policy setting does not specify whether users can change the network name for all network connections.
  - **User can change name**. If you select this option, users can change the network name for all networks to which they connect.
  - **User cannot change name**. If you select this option, users cannot change the network name for any networks to which they connect.
- **Network location**. For this item, the following options are available:
  - **Not configured**. If you select this option, this policy setting does not specify whether users can change the location for all network connections.
  - **User can change location**. If you select this option, this policy setting allows users to change all network locations from Private to Public or from Public to Private.
  - **User cannot change location**. If you select this option, this policy setting does not allow users to change the location for any networks to which they connect.
- **Network icon**. For this item, the following options are available:

- **Not configured**. If you select this option, this policy setting does not specify whether users can change the network icon for all network connections.
- **User can change icon**. If you select this option, this policy setting allows users to change the network icon for all networks to which the user connects.
- **User cannot change icon**. If you select this option, this policy setting does not allow users to change the network icon for any networks to which the user connects.

Smart Card Overview

This topic lists documentation resources for smart cards that include product evaluation, getting started guides, procedures, design and deployment guides, technical references, and command references.

Did you mean…

- Smart Cards in Windows 7

Feature description

Smart cards are a tamper-resistant and portable way to provide security solutions for tasks such as client authentication, logging on to domains, code signing, and securing e-mail.

Practical applications

Smart cards and their associated personal identification numbers (PINs) are an increasingly popular, reliable, and cost-effective form of two-factor authentication. With the right controls in place, the user must have the smart card and know the PIN to gain access to network resources. The two-factor requirement significantly reduces the likelihood of unauthorized access to an organization's network. Smart cards provide particularly effective security control in two scenarios: to secure administrator accounts and to secure remote access.

See also

The following table provides links to relevant resources in understanding and using smart cards.

Content type	References
Product evaluation	What's New in Smart Cards | Understanding and Evaluating Virtual Smart Cards
Scenario documentation	Strong Authentication with One-Time Passwords in Windows 7 and Windows Server 2008 R2

Content type	References
Operations	Not yet available
Troubleshooting	Smart Card Troubleshooting Guide
Security	Windows Smart Card Technical Reference
Community resources	Windows Security Survival Guide

What's New in Smart Cards

This document describes new smart card−related features on the Windows Server 2012, Windows 8, and Windows RT operating systems.

Feature description

Smart cards together with personal identification numbers (PINs) are an increasingly popular, reliable, and cost-effective form of two-factor authentication. With the right controls in place, a user must have the smart card and know the PIN to gain access to network resources. The two-factor requirement significantly reduces the likelihood of unauthorized access to an organization's network. Smart cards provide particularly effective security control for:

- authentication for scenarios such as remote access
- data integrity for scenarios such as document signing
- data confidentiality for scenarios that require encryption.

Their use in additional scenarios, such as secure access to high-value applications, is likely to grow as organizations deploy a new generation of secure applications.

New and changed functionality

The following changes have been made to smart card support in Windows Server 2012, Windows 8, and Windows RT:

- Virtual smart cards
- Changes to the smart card sign-in experience
- Smart Card Service start and stop behavior
- Smart card transactions
- Smart card support on Windows RT
- Smart card support in Windows 8 applications

Virtual smart cards

Virtual smart cards emulate the functionality of traditional smart cards, but they use the Trusted Platform Module (TPM) chip available on many organizations' computers rather than requiring the use of a separate physical smart card and reader. Virtual smart cards involve technical, functional, security, and cost differences with conventional smart cards.

To the end user, the virtual smart card is a smart card that is always available on the computer. If a user needs to use more than one computer, a new virtual smart card must be issued to the user for each computer. Also, a computer that is shared among multiple users can host multiple virtual smart cards, one for each user.

Conventional smart cards and TPM virtual smart cards offer comparable levels of security. TPM virtual smart cards can be deployed with no additional material cost, as long as employees have computers with built-in TPMs. For more information, see Understanding and Evaluating Virtual Smart Cards.

Changes to the smart card sign-in experience

For end users, the sign-in experience on Windows Server 2012 and Windows 8 has improved detection of whether a smart card reader was installed and whether a smart card or a password was used to sign in or unlock the computer the last time. If a smart card was not installed previously, and the user selects the smart card sign-in icon, a message appears telling the user to connect a smart card. After a card is connected, the smart card PIN dialog box appears. If the user does not want to use the sign-in option that automatically appears (if their smart card is not readily available, for example), a second message allows the user to select from different sign-in options.

Smart Card Service start and stop behavior

Smart card reader detection logic has been added so that the Smart Card Service runs only when appropriate. On Windows Server 2012 and Windows 8, the Smart Card Service (scardsvr) automatically starts when the user connects a smart card reader and automatically stops when a user removes a smart card reader and no other smart card reader is connected to the computer. On startup, the Smart Card Service automatically starts if a reader was previously connected to the computer but a reader is not currently connected to the system. If no smart card readers are connected to the computer, the service will automatically shut down one minute after the last API call into the Smart Card Service. If a reader was never previously connected to the computer, the service will not start automatically.

Smart card transactions

On Windows Server 2012, Windows 8, and Windows RT, if a transaction is held on the card for more than 5 seconds with no operations happening on the card, the card is reset. This is a change from the behavior in previous releases.

For more information about this behavior, see SCardBeginTransaction function.

Smart card support on Windows RT

Smart card support for Windows RT includes the following:

- **Smart card readers**

  Only smart card readers that connect over USB and support the USB Chip/Smart Card Interface Devices (CCID) specification are allowed on Windows RT. Such smart card readers must use the USB CCID specification smart card reader class driver that comes with Windows RT.

- **Smart cards**

  Only smart cards that support the Generic Identity Device Specification (GIDS) or the Personal Identity Verification (PIV) standard are supported on Windows RT. The class drivers for cards based upon these specifications are included within Windows.

Smart card support in Windows 8 applications

Windows 8 supports a number of new types of desktop applications. Developers who create applications that need the security benefits of smart cards must address the following requirements; otherwise, these applications cannot automatically use smart cards to support their functionality.

- To use a smart card, applications running in AppContainer must have the SharedUserCertificates capability in their application manifest. Without this capability, the application will not be permitted to use a smart card for authentication, signing, or encryption. For information about this capability and how to include it in the manifest, see Setting certificate store capabilities.
- For Windows RT applications, smart card support is limited to SSL client authentication. For a sample application that demonstrates the use of smart cards for SSL client authentication, see Windows Store app for banking: code walkthrough.

See also

Understanding and Evaluating Virtual Smart Cards

Software Restriction Policies

This topic for the IT professional describes Software Restriction Policies (SRP) in Windows Server 2012 and Windows 8, and provides links to technical information about SRP beginning with Windows Server 2003.

For procedures and troubleshooting tips, see Administer Software Restriction Policies and Troubleshoot Software Restriction Policies.

Did you mean…

- Security policies (see Security Policy Settings Overview)
- AppLocker (see **AppLocker Overview**)

Software Restriction Policies description

Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. Software restriction policies are part of the Microsoft security and management strategy to assist enterprises in increasing the reliability, integrity, and manageability of their computers.

You can also use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. Software restriction policies are integrated with Microsoft Active Directory and Group Policy. You can also create software restriction policies on stand-alone computers. Software restriction policies are trust policies, which are regulations set by an administrator to restrict scripts and other code that is not fully trusted from running.

You can define these policies through the Software Restriction Policies extension of the Local Group Policy Editor or the Local Security Policies snap-in to the Microsoft Management Console (MMC).

For in-depth information about SRP, see the Software Restriction Policies Technical Overview.

Practical applications

Administrators can use software restriction policies for the following tasks:

- Define what is trusted code
- Design a flexible Group Policy for regulating scripts, executable files, and ActiveX controls

Software restriction policies are enforced by the operating system and by applications (such as scripting applications) that comply with software restriction policies.

Specifically, administrators can use software restriction policies for the following purposes:

- Specify which software (executable files) can run on clients
- Prevent users from running specific programs on shared computers
- Specify who can add trusted publishers to clients
- Set the scope of the software restriction policies (specify whether policies affect all users or a subset of users on clients)
- Prevent executable files from running on the local computer, organizational unit (OU), site, or domain. This would be appropriate in cases when you are not using software restriction policies to address potential issues with malicious users.

New and changed functionality

There are no changes in functionality for Software Restriction Policies.

Removed or deprecated functionality

There is no removed or deprecated functionality for Software Restriction Policies.

Software requirements

The Software Restriction Policies extension to the Local Group Policy Editor can be accessed through the MMC.

The following features are required to create and maintain software restriction policies on the local computer:

- Local Group Policy Editor
- Windows Installer
- Authenticode and WinVerifyTrust

If your design calls for domain deployment of these policies, in addition to the above list, the following features are required:

- Active Directory Domain Services
- Group Policy

Server Manager information

Software Restriction Policies is an extension of the Local Group Policy Editor and is not installed through Server Manager, Add Roles and Features.

See also

The following table provides links to relevant resources in understanding and using SRP.

Content type	References
Product evaluation	Application Lockdown with Software Restriction Policies
Planning	Software Restriction Policies Technical Overview (Windows Server 2012) Software Restriction Policies Technical Reference (Windows Server 2003)
Deployment	No resources available.
Operations	Administer Software Restriction Policies (Windows Server 2012) Software Restriction Policies Product Help (Windows Server 2003)
Troubleshooting	Troubleshoot Software Restriction Policies (Windows Server 2012) Software Restriction Policies Troubleshooting (Windows Server 2003)

Content type	References
Security	Threats and Countermeasures for Software Restriction Polices (Windows Server 2008) Threats and Countermeasures for Software Restriction Polices (Windows Server 2008 R2)
Tools and settings	Software Restriction Policies Tools and Settings (Windows Server 2003)
Community resources	Application Lockdown with Software Restriction Policies
Related technologies	**AppLocker Overview**

Software Restriction Policies Technical Overview

This topic describes software restriction policies, when and how to use the feature, what changes have been implemented in past releases, and provides links to additional resources to help you create and deploy software restriction policies beginning with Windows Server 2008 and Windows Vista.

Introduction

Software restriction policies provide administrators with a Group Policy-driven mechanism to identify software and control its ability to run on the local computer. These policies can be used to protect computers running Microsoft Windows operating systems (beginning with Windows Server 2003 and Windows XP Professional) against known conflicts and safeguard the computers against security threats such as malicious viruses and Trojan horse programs. You can also use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. Software restriction policies are integrated with Microsoft Active Directory and Group Policy. You can also create software restriction policies on stand-alone computers.

Software restriction policies are trust policies, which are regulations set by an administrator to restrict scripts and other code that is not fully trusted from running. The Software Restriction Policies extension to the Local Group Policy Editor provides a single user interface through which the settings for restricting the use of applications can be managed on the local computer or throughout a domain.

Procedures

- Administer Software Restriction Policies
  - Determine Allow/Deny List and Application Inventory for Software Restriction Policies
  - Work with Software Restriction Policies Rules
  - Use Software Restriction Policies to Help Protect Your Computer Against an Email Virus
- Troubleshoot Software Restriction Policies

Software restriction policy usage scenarios

Business users collaborate by using e-mail, instant messaging, and peer-to-peer applications. As these collaborations increase, especially with the use of the Internet in business computing, so do the threats from malicious code, such as worms, viruses, and malicious user or attacker threats.

Users might receive hostile code in many forms, ranging from native Windows executable files (.exe files), to macros in documents (such as .doc files), to scripts (such as .vbs files). Malicious users or attackers often use social engineering methods to get users to run code containing viruses and worms. (Social engineering is a term for tricking people into revealing their password or some form of security information.) If such code is activated, it can generate denial-of-service attacks on the network, send sensitive or private data to the Internet, put the security of the computer at risk, or damage the contents of the hard disk drive.

IT organizations and users must be able to determine which software is safe to run and which is not. With the large numbers and forms that hostile code can take, this becomes a difficult task.

To help protect their network computers from both hostile code and unknown or unsupported software, organizations can implement software restriction policies as part of their overall security strategy.

Administrators can use software restriction policies for the following tasks:

- Define what is trusted code
- Design a flexible Group Policy for regulating scripts, executable files, and ActiveX controls

Software restriction policies are enforced by the operating system and by applications (such as scripting applications) that comply with software restriction policies.

Specifically, administrators can use software restriction policies for the following purposes:

- Specify which software (executable files) can run on client computers
- Prevent users from running specific programs on shared computers
- Specify who can add trusted publishers to client computers
- Set the scope of the software restriction policies (specify whether policies affect all users or a subset of users on client computers)
- Prevent executable files from running on the local computer, organizational unit (OU), site, or domain. This would be appropriate in cases when you are not using software restriction policies to address potential issues with malicious users.

Differences and changes in functionality

There are no changes in functionality in SRP for Windows Server 2012 and Windows 8.

Supported versions

Software Restriction Policies can only be configured on and applied to computers running at least Windows Server 2003, including Windows Server 2012, and at least Windows XP, including Windows 8.

Note

Certain editions of the Windows client operating system beginning with Windows Vista do not have Software Restrictions Policies. Computers not administered in a domain by Group Policy might not receive distributed policies.

Comparing application control functions in Software Restriction Policies and AppLocker

The following table compares the features and functions of the Software Restriction Policies (SRP) feature and AppLocker.

Application control function	SRP	AppLocker
Scope	SRP policies can be applied to all Windows operating systems beginning with Windows XP and Windows Server 2003.	AppLocker policies apply only to Windows Server 2008 R2, Windows Server 2012, Windows 7, and Windows 8.
Policy creation	SRP policies are maintained through Group Policy and only the administrator of the GPO can update the SRP policy. The administrator on the local computer can modify the SRP policies defined in the local GPO.	AppLocker policies are maintained through Group Policy and only the administrator of the GPO can update the policy. The administrator on the local computer can modify the AppLocker policies defined in the local GPO. AppLocker permits customization of error messages to direct users to a Web page for help.
Policy maintenance	SRP policies must be updated by using the Local Security Policy snap-in (if the policies are created locally) or the Group Policy Management Console (GPMC).	AppLocker policies can be updated by using the Local Security Policy snap-in (if the policies are created locally), or the GPMC, or the Windows PowerShell AppLocker cmdlets.

Application control function	SRP	AppLocker
Policy application	SRP policies are distributed through Group Policy.	AppLocker policies are distributed through Group Policy.
Enforcement mode	SRP works in the “deny list mode” where administrators can create rules for files that they do not want to allow in this Enterprise whereas the rest of the file are allowed to run by default. SRP can also be configured in the “allow list mode” such that the by default all files are blocked and administrators need to create allow rules for files that they want to allow.	AppLocker by default works in the “allow list mode” where only those files are allowed to run for which there is a matching allow rule.
File types that can be controlled	SRP can control the following file types: • Executables • Dlls • Scripts • Windows Installers SRP cannot control each file type separately. All SRP rules are in a single rule collection.	AppLocker can control the following file types: • Executables • Dlls • Scripts • Windows Installers • Packaged apps and installers (Windows Server 2012 and Windows 8) AppLocker maintains a separate rule collection for each of the five file types.
Designated file types	SRP supports an extensible list of file types that are considered executable. Administrators can add extensions for files that should be considered executable.	AppLocker does not support this. AppLocker currently supports the following file extensions: • Executables (.exe, .com) • Dlls (.ocx, .dll) • Scripts (.vbs, .js, .ps1, .cmd, .bat) • Windows Installers (.msi, .mst, .msp)

Application control function	SRP	AppLocker
		• Packaged app installers (.appx)
Rule types	SRP supports four types of rules: • Hash • Path • Signature • Internet zone	AppLocker supports three types of rules: • Hash • Path • Publisher
Editing the hash value	SRP allows administrators to provide custom hash values.	AppLocker computes the hash value itself. Internally it uses the SHA1 Authenticode hash for Portable Executables (Exe and Dll) and Windows Installers and a SHA1 flat file hash for the rest.
Support for different security levels	With SRP administrators can specify the permissions with which an app can run. So, an administrator can configure a rule such that notepad always runs with restricted permissions and never with administrative privileges. SRP on Windows Vista and earlier supported multiple security levels. On Windows 7 that list was restricted to just two levels: Disallowed and Unrestricted (Basic User translates to Disallowed).	AppLocker does not support security levels.
Manage Packaged apps and Packaged app installers	Unable	.appx is a valid file type which AppLocker can manage.
Targeting a rule to a user or a group of users	SRP rules apply to all users on a particular computer.	AppLocker rules can be targeted to a specific user or a group of users.
Support for rule exceptions	SRP does not support rule exceptions	AppLocker rules can have exceptions which allow administrators to create rules

Application control function	SRP	AppLocker
		such as “Allow everything from Windows except for Regedit.exe”.
Support for audit mode	SRP does not support audit mode. The only way to test SRP policies is to set up a test environment and run a few experiments.	AppLocker supports audit mode which allows administrators to test the effect of their policy in the real production environment without impacting the user experience. Once you are satisfied with the results, you can start enforcing the policy.
Support for exporting and importing policies	SRP does not support policy import/export.	AppLocker supports the importing and exporting of policies. This allows you to create AppLocker policy on a sample computer, test it out and then export that policy and import it back into the desired GPO.
Rule enforcement	Internally, SRP rules enforcement happens in the user-mode which is less secure.	Internally, AppLocker rules for Exes and Dlls are enforced in the kernel-mode which is more secure than enforcing them in the user-mode.

System requirements

Software restriction policies can only be configured on and applied to computers running at least Windows Server 2003, and at least Windows XP. Group Policy is required to distribute Group Policy Objects that contain software restriction policies.

Software restriction policies components and architecture

Software restriction policies provide a mechanism for the operating system and applications compliant with software restriction policies to restrict the runtime execution of software programs.

At a high level, software restriction policies consist of the following components:

- Software restriction policies API. The Application Programming Interfaces (APIs) are used to create and configure the rules that constitute the software restriction policy. There also are software restriction policies APIs for querying, processing, and enforcing software restriction policies.
- A software restriction policies management tool. This consists of the **Software Restriction Policies** extension of the **Local Group Policy Object Editor** snap-in, which administrators use to create and edit the software restriction policies.
- A set of operating system APIs and applications that call the software restriction policies APIs to provide enforcement of the software restriction policies at runtime.
- Active Directory and Group Policy. Software restriction policies depend on the Group Policy infrastructure to propagate the software restriction policies from the Active Directory to the appropriate clients, and for scoping and filtering the application of these policies to the appropriate target computers.
- Authenticode and WinVerify Trust APIs which are used to process signed executable files.
- Event Viewer. The functions used by software restriction policies log events to the Event Viewer logs.
- Resultant Set of Policies (RSoP), which can aid in the diagnosing of the effective policy that will be applied to a client.

For more information about SRP architecture, how SRP manages rules, processes and interactions, see How Software Restriction Policies Work in the Windows Server 2003 Technical Library.

Best practices

Do not modify the default domain policy.

- If you do not edit the default domain policy, you always have the option of reapplying the default domain policy if something goes wrong with your customized domain policy.

Create a separate Group Policy Object for software restriction policies.

- If you create a separate Group Policy Object (GPO) for software restriction policies, you can disable software restriction policies in an emergency without disabling the rest of your domain policy.

If you experience problems with applied policy settings, restart Windows in Safe Mode.

- Software restriction policies do not apply when Windows is started in Safe Mode. If you accidentally lock down a workstation with software restriction policies, restart the computer in Safe Mode, log on as a local administrator, modify the policy, run **gpupdate**, restart the computer, and then log on normally.

Use caution when defining a default setting of Disallowed.

- When you define a default setting of **Disallowed**, all software is disallowed except for software that has been explicitly allowed. Any file that you want to open has to have a software restriction policies rule that allows it to open.
- To protect administrators from locking themselves out of the system, when the default security level is set to **Disallowed**, four registry path rules are automatically created. You can delete or modify these registry path rules; however, this is not recommended.

For best security, use access control lists in conjunction with software restriction policies.

- Users might try to circumvent software restriction policies by renaming or moving disallowed files or by overwriting unrestricted files. As a result, it is recommended that you use access control lists (ACLs) to deny users the access necessary to perform these tasks.

Test new policy settings thoroughly in test environments before applying the policy settings to your domain.

- New policy settings might act differently than originally expected. Testing diminishes the chance of encountering a problem when you deploy policy settings across your network.
- You can set up a test domain, separate from your organization's domain, in which to test new policy settings. You can also test the policy settings by creating a test GPO and linking it to a test organizational unit. When you have thoroughly tested the policy settings with test users, you can link the test GPO to your domain.
- Do not set programs or files to **Disallowed** without testing to see what the effect may be. Restrictions on certain files can seriously affect the operation of your computer or network.
- Information that is entered incorrectly or typing mistakes can result in a policy setting that does not perform as expected. Testing new policy settings before applying them can prevent unexpected behavior.

Filter user policy settings based on membership in security groups.

- You can specify users or groups for which you do not want a policy setting to apply by clearing the **Apply Group Policy** and **Read** check boxes, which are located on the **Security** tab of the properties dialog box for the GPO.
- When the Read permission is denied, the policy setting is not downloaded by the computer. As a result, less bandwidth is consumed by downloading unnecessary policy settings, which enables the network to function more quickly. To deny the Read permission, select **Deny** for the **Read** check box, which is located on the **Security** tab of the properties dialog box for the GPO.

Do not link to a GPO in another domain or site.

- Linking to a GPO in another domain or site can result in poor performance.

Additional resources

Content type	References
Planning	Software Restriction Policies Technical Reference
Operations	Administer Software Restriction Policies
Troubleshooting	Software Restriction Policies Troubleshooting (2003)
Security	Threats and Countermeasures for Software Restriction Polices (2008) Threats and Countermeasures for Software Restriction Polices (2008 R2)
Tools and settings	Software Restriction Policies Tools and Settings (2003)
Community resources	Application Lockdown with Software Restriction Policies
Related technologies	**AppLocker Overview**

Administer Software Restriction Policies

This topic for the IT professional contains procedures how to administer application control policies using Software Restriction Policies (SRP) beginning with Windows Server 2008 and Windows Vista.

Introduction

Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. You use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. These are integrated with Microsoft Active Directory Domain Services and Group Policy but can also be configured on stand-alone computers. For more information about SRP, see the Software Restriction Policies.

Beginning with Windows Server 2008 R2 and Windows 7, Windows AppLocker can be used instead of or in concert with SRP for a portion of your application control strategy. For information about AppLocker in Windows Server 2012 and Windows 8, see the **AppLocker Overview**.

This topic contains:

- To open Software Restriction Policies
- To create new software restriction policies
- To add or delete a designated file type
- To prevent software restriction policies from applying to local administrators
- To change the default security level of software restriction policies
- To apply software restriction policies to DLLs

For information about how to accomplish specific tasks using SRP, see the following:

- Determine Allow/Deny List and Application Inventory for Software Restriction Policies
- Work with Software Restriction Policies Rules
- Use Software Restriction Policies to Help Protect Your Computer Against an Email Virus

To open Software Restriction Policies

- For your local computer
- For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain
- For a domain or organizational unit, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed
- For a site, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

For your local computer

1. Open Local Security Settings.
2. In the console tree, click **Software Restriction Policies**.

   Where?

   - Security Settings/Software Restriction Policies

Note
To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **Local Group Policy Object Editor**, and then click **Add**.
4. In **Select Group Policy Object**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy Object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, and then click **Finish**.
6. Click **Close**, and then click **OK**.

7. In the console tree, click **Software Restriction Policies**.

   Where?

   - *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or User Configuration/Windows Settings/Security Settings/Software Restriction Policies

Note
To perform this procedure, you must be a member of the Domain Admins group.

For a domain or organizational unit, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

1. Open Group Policy Management Console.
2. In the console tree, right-click the Group Policy Object (GPO) that you want to open software restriction policies for.
3. Click **Edit** to open the GPO that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
4. In the console tree, click **Software Restriction Policies**.

   Where?

   - *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or User Configuration/Windows Settings/Security Settings/Software Restriction Policies

Note
To perform this procedure, you must be a member of the Domain Admins group.

For a site, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

1. Open Group Policy Management Console.
2. In the console tree, right-click the site that you want to set Group Policy for.

   Where?

   - Active Directory Sites and Services [*Domain_Controller_Name.Domain_Name*]/Sites/Site
3. Click an entry in **Group Policy Object Links** to select an existing Group Policy Object (GPO), and then click **Edit**. You can also click **New** to create a new GPO, and then click **Edit**.
4. In the console tree, click **Software Restriction Policies**.

   Where

   - *Group Policy Object* [*ComputerName*] Policy/Computer Configuration or User Configuration/Windows Settings/Security Settings/Software Restriction Policies

Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- To set policy settings that will be applied to computers, regardless of which users log on to them, click **Computer Configuration**.
- To set policy settings that will be applied to users, regardless of which computer they log on to, click **User Configuration**.

To create new software restriction policies

1. Open Software Restriction Policies.
2. On the **Action** menu, click **New Software Restriction Policies**.

Warning

- Different administrative credentials are required to perform this procedure, depending on your environment:
  - If you create new software restriction policies for your local computer: Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.
  - If you create new software restriction policies for a computer that is joined to a domain, members of the Domain Admins group can perform this procedure.
- If software restriction policies have already been created for a Group Policy Object (GPO), the **New Software Restriction Policies** command does not appear on the **Action** menu. To delete the software restriction policies that are applied to a GPO, in the console tree, right-click **Software Restriction Policies**, and then click **Delete Software Restriction Policies**. When you delete software restriction policies for a GPO, you also delete all software restriction policies rules for that GPO. After you delete software restriction policies, you can create new software restriction policies for that GPO.

To add or delete a designated file type

1. Open Software Restriction Policies.
2. In the details pane, double-click **Designated File Types**.
3. Do one of the following:
   - To add a file type, in **File name extension**, type the file name extension, and then click **Add**.
   - To delete a file type, in **Designated file types**, click the file type, and then click **Remove**.

Notes

- Different administrative credentials are required to perform this procedure, depending on the environment in which you add or delete a designated file type:

    - If you add or delete a designated file type for your local computer: Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.
    - If you create new software restriction policies for a computer that is joined to a domain, members of the Domain Admins group can perform this procedure.
- It may be necessary to create a new software restriction policy setting for the Group Policy Object (GPO) if you have not already done so.
- The list of designated file types is shared by all rules for both Computer Configuration and User Configuration for a GPO.

To prevent software restriction policies from applying to local administrators

1. Open Software Restriction Policies.
2. In the details pane, double-click **Enforcement**.
3. Under **Apply software restriction policies to the following users**, click **All users except local administrators**.

Warning

- Membership in the local **Administrators** group, or equivalent, is the minimum required to complete this procedure.
- It may be necessary to create a new software restriction policy setting for the Group Policy Object (GPO) if you have not already done so.
- If it is common for users to be members of the local Administrators group on their computers in your organization, you may not want to enable this option.
- If you are defining a software restriction policy setting for your local computer, use this procedure to prevent local administrators from having software restriction policies applied to them. If you are defining a software restriction policy setting for your network, filter user policy settings based on membership in security groups through Group Policy.

To change the default security level of software restriction policies

1. Open Software Restriction Policies.
2. In the details pane, double-click **Security Levels**.
3. Right-click the security level that you want to set as the default, and then click **Set as default**.

Caution

In certain directories, setting the default security level to **Disallowed** can adversely affect your operating system.

Notes

- Different administrative credentials are required to perform this procedure, depending on the environment for which you change the default security level of software restriction policies.
- It may be necessary to create a new software restriction policy setting for this Group Policy Object (GPO) if you have not already done so.
- In the details pane, the current default security level is indicated by a black circle with a check mark in it. If you right-click the current default security level, the **Set as default** command does not appear in the menu.
- Software restriction policies rules are created to specify exceptions to the default security level. When the default security level is set to **Unrestricted**, rules can specify software that is not allowed to run. When the default security level is set to **Disallowed**, rules can specify software that is allowed to run.
- At installation, the default security level of software restriction policies on all files on your system is set to **Unrestricted**.

To apply software restriction policies to DLLs

1. Open Software Restriction Policies.
2. In the details pane, double-click **Enforcement**.
3. Under **Apply software restriction policies to the following**, click **All software files**.

Notes

- To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure.
- By default, software restriction policies do not check dynamic-link libraries (DLLs). Checking DLLs can decrease system performance, because software restriction policies must be evaluated every time a DLL is loaded. However, you may decide to check DLLs if you are concerned about receiving a virus that targets DLLs. If the default security level is set to **Disallowed**, and you enable DLL checking, you must create software restriction policies rules that allow each DLL to run.

Determine Allow/Deny List and Application Inventory for Software Restriction Policies

This topic for the IT professional gives guidance how to create an allow and deny list for applications to be managed by Software Restriction Policies (SRP) beginning with Windows Server 2008 and Windows Vista.

Introduction

Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. You use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. These are integrated with Microsoft Active Directory Domain Services and Group Policy but can also be configured on stand-alone computers. For a starting point for SRP, see the Software Restriction Policies.

Beginning with Windows Server® 2008 R2 and Windows® 7, Windows AppLocker can be used instead of or in concert with SRP for a portion of your application control strategy. For information about AppLocker in Windows Server 2012 and Windows 8, see the **AppLocker Overview**.

For information about how to accomplish specific tasks using SRP, see the following:

- Work with Software Restriction Policies Rules
- Use Software Restriction Policies to Help Protect Your Computer Against an Email Virus

What default rule to choose: Allow or Deny

Software restriction policies can be deployed in one of two modes that are the basis of your default rule: Allow List or Deny List. You can create a policy that identifies every application that is allowed to run in your environment; the default rule within your policy is Restricted and will block all applications that you do not explicitly allow to run. Or you can create a policy that identifies every application that cannot run; the default rule is Unrestricted and restricts only the applications that you have explicitly listed.

Important
The Deny List mode might be a high-maintenance strategy for your organization regarding application control. Creating and maintaining an evolving list that prohibits all malware and other problematic applications would be time consuming and susceptible to mistakes.

Create an inventory of your applications for the Allow list

To effectively use the Allow default rule, you need to determine exactly which applications are required in your organization. There are tools designed to produce an application inventory, such as the Inventory Collector in the Microsoft Application Compatibility Toolkit. But SRP has an advanced logging feature to help you understand exactly what applications are running in your environment.

To discover which applications to allow

1. In a test environment, deploy Software Restriction Policy with the default rule set to Unrestricted and remove any additional rules. If you enable SRP without forcing it to restrict any applications, SPR will be able to monitor what applications are being run.
2. Create the following registry value in order to enable the advanced logging feature and

set the path to where the log file should be written.

"HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\Safer\CodeIdentifiers"

String Value: *NameLogFile path to NameLogFile*

Because SRP is evaluating all applications when they run, an entry is written to the log file *NameLogFile* each time that application is run.

3. Evaluate the log file

   Each log entry states:

   - the caller of the software restriction policy and the process ID (PID) of the calling process
   - the target being evaluated
   - the SRP rule that was encountered when that application ran
   - an identifier for the SRP rule.

   An example of the output written to a log file:

```
explorer.exe (PID = 4728) identified

C:\Windows\system32\onenote.exe as Unrestricted using

path rule, Guid =

{320bd852-aa7c-4674-82c5-9a80321670a3}
```

   All applications and associated code that SRP checks and set to block will be noted in the log file, which you then can use to determine which executables should be considered for your Allowed list.

Work with Software Restriction Policies Rules

This topic describes procedures working with certificate, path, internet zone and hash rules using Software Restriction Policies.

Introduction

With software restriction policies, you can protect your computing environment from untrusted software by identifying and specifying what software is allowed to run. You can define a default security level of **Unrestricted** or **Disallowed** for a Group Policy Object (GPO) so that software is either allowed or not allowed to run by default. You can make exceptions to this default security

level by creating software restriction policies rules for specific software. For example, if the default security level is set to **Disallowed**, you can create rules that allow specific software to run. The types of rules are as follows:

- **Certificate rules**

  For procedures, see Working with certificate rules.
- **Hash rules**

  For procedures, see Working with hash rules.
- **Internet zone rules**

  For procedures, see Working with Internet Zone rules.
- **Path rules**

  For procedures, see Working with path rules.

For information about other tasks to manage Software Restriction Policies, see Administer Software Restriction Policies.

To learn how to work with rules using AppLocker, see Administer AppLocker.

Working with certificate rules

Software restriction policies can also identify software by its signing certificate. You can create a certificate rule that identifies software and then allows or does not allow the software to run, depending on the security level. For example, you can use certificate rules to automatically trust software from a trusted source in a domain without prompting the user. You can also use certificate rules to run files in disallowed areas of your operating system. Certificate rules are not enabled by default.

When rules are created for the domain using Group Policy, you must have permissions to create or modify a Group Policy Object. If you are creating rules for the local computer, you must have administrative credentials on that computer.

To create a certificate rule

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Certificate Rule**.
3. Click **Browse**, and then select a certificate or signed file.
4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

Notes

- It might be necessary to create a new software restriction policy setting for the Group Policy Object (GPO) if you have not already done so.
- Certificate rules are not enabled by default.

- The only file types that are affected by certificate rules are those that are listed in **Designated file types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts.

Enabling certificate rules

There are different procedures for enabling certificate rules depending on your environment:

- To enable certificate rules for your local computer
- To enable certificate rules For a Group Policy Object, and you are on a server that is joined to a domain
- To enable certificate rules for a Group Policy Object, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed
- To enable certificate rules for only domain controllers, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

To enable certificate rules for your local computer

1. Open Local Security Settings.
2. In the console tree, click **Security Options** located under Security Settings/Local Policies.
3. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.
4. Do one of the following, and then click **OK**:
   - To enable certificate rules, click **Enabled**.
   - To disable certificate rules, click **Disabled**.

To enable certificate rules For a Group Policy Object, and you are on a server that is joined to a domain

1. Open Microsoft Management Console (MMC).
2. On the **File** menu, click **Add/Remove snap-in**, and then click **Add**.
3. Click **Local Group Policy Object Editor**, and then click **Add**.
4. In **Select Group Policy Object**, click **Browse**.
5. In **Browse for a Group Policy Object**, select a Group Policy Object (GPO) in the appropriate domain, site, or organizational unit--or create a new one, and then click **Finish**.
6. Click **Close**, and then click **OK**.
7. In the console tree, click **Security Options** located under *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security

Settings/Local Policies/.

8. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.
9. If this policy setting has not yet been defined, select the **Define these policy settings** check box.
10. Do one of the following, and then click **OK**:
    - To enable certificate rules, click **Enabled**.
    - To disable certificate rules, click **Disabled**.

To enable certificate rules for a Group Policy Object, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

1. Open Active Directory Users and Computers.
2. In the console tree, right-click the Group Policy Object (GPO) for which you want to enable certificate rules.
3. Click **Properties**, and then click the **Group Policy** tab.
4. Click **Edit** to open the GPO that you want to edit. You can also click **New** to create a new GPO, and then click **Edit**.
5. In the console tree, click **Security Options** located under *GroupPolicyObject*[*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies.
6. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.
7. If this policy setting has not yet been defined, select the **Define these policy settings** check box.
8. Do one of the following, and then click **OK**:
    - To enable certificate rules, click **Enabled**.
    - To disable certificate rules, click **Disabled**.

To enable certificate rules for only domain controllers, and you are on a domain controller or on a workstation that has the Remote Server Administration Tools installed

1. Open Domain Controller Security Settings.
2. In the console tree, click **Security Options** located under *GroupPolicyObject* [*ComputerName*] Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies.
3. In the details pane, double-click **System settings: Use Certificate Rules on Windows Executables for Software Restriction Policies**.
4. If this policy setting has not yet been defined, select the **Define these policy settings** check box.
5. Do one of the following, and then click **OK**:
    - To enable certificate rules, click **Enabled**.

- To disable certificate rules, click **Disabled**.

Note

You must perform this procedure before certificate rules can take effect.

Set trusted publisher options

Software signing is being used by a growing number of software publishers and application developers to verify that their applications come from a trusted source. However, many users do not understand or pay little attention to the signing certificates associated with applications that they install.

The policy settings in the **Trusted Publishers** tab of the certificate path validation policy allows administrators to control which certificates can be accepted as coming from a trusted publisher.

To configure the trusted publishers policy settings for a local computer

1. On the **Start** screen, type **gpedit.msc** and then press ENTER.
2. In the console tree under **Local Computer Policy\Computer Configuration\Windows Settings\Security Settings**, click **Public Key Policies**.
3. Double-click **Certificate Path Validation Settings**, and then click the **Trusted Publishers** tab.
4. Select the **Define these policy settings** check box, select the policy settings that you want to apply, and then click **OK** to apply the new settings.

To configure the trusted publishers policy settings for a domain

1. Open **Group Policy Management**.
2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the **Default Domain Policy** Group Policy Object (GPO) that you want to edit.
3. Right-click the **Default Domain Policy** GPO, and then click **Edit**.
4. In the console tree under **Computer Configuration\Windows Settings\Security Settings**, click **Public Key Policies**.
5. Double-click **Certificate Path Validation Settings**, and then click the **Trusted Publishers** tab.
6. Select the **Define these policy settings** check box, select the policy settings that you want to apply, and then click **OK** to apply the new settings.

To allow only administrators to manage certificates used for code signing for a local computer

1. On the **Start** screen, type, **gpedit.msc** in the **Search programs and files** or in Windows 8, on the Desktop, and then press ENTER.
2. In the console tree under **Default Domain Policy** or **Local Computer Policy**, double-click **Computer Configuration**, **Windows Settings**, and **Security Settings**, and then click **Public Key Policies**.

3. Double-click **Certificate Path Validation Settings**, and then click the **Trusted Publishers** tab.
4. Select the **Define these policy settings** check box.
5. Under **Trusted publisher management**, click **Allow only all administrators to manage Trusted Publishers**, and then click **OK** to apply the new settings.

To allow only administrators to manage certificates used for code signing for a domain

1. Open **Group Policy Management**.
2. In the console tree, double-click **Group Policy Objects** in the forest and domain containing the **Default Domain Policy** GPO that you want to edit.
3. Right-click the **Default Domain Policy** GPO, and then click **Edit**.
4. In the console tree under **Computer Configuration\Windows Settings\Security Settings**, click **Public Key Policies**.
5. Double-click **Certificate Path Validation Settings**, and then click the **Trusted Publishers** tab.
6. Select the **Define these policy settings** check box, implement the changes you want, and then click **OK** to apply the new settings.

Working with hash rules

A hash is a series of bytes with a fixed length that uniquely identifies a software program or file. The hash is computed by a hash algorithm. When a hash rule is created for a software program, software restriction policies calculate a hash of the program. When a user tries to open a software program, a hash of the program is compared to existing hash rules for software restriction policies. The hash of a software program is always the same, regardless of where the program is located on the computer. However, if a software program is altered in any way, its hash also changes, and it no longer matches the hash in the hash rule for software restriction policies.

For example, you can create a hash rule and set the security level to **Disallowed** to prevent users from running a certain file. A file can be renamed or moved to another folder and still result in the same hash. However, any changes to the file itself also change its hash value and allow the file to bypass restrictions.

To create a hash rule

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Hash Rule**.
3. Click **Browse** to find a file.

   Note
   In Windows XP it is possible to paste a pre-calculated hash in **File hash**. In Windows Server 2008 R2, Windows 7 and later versions, this option is not

available.

4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

Notes

- It may be necessary to create a new software restriction policy setting for the Group Policy Object (GPO) if you have not already done so.
- A hash rule can be created for a virus or a Trojan horse to prevent them from running.
- If you want other people to use a hash rule so that a virus cannot run, calculate the hash of the virus by using software restriction policies, and then e-mail the hash value to the other people. Never e-mail the virus itself.
- If a virus has been sent through e-mail, you can also create a path rule to prevent execution of e-mail attachments.
- A file that is renamed or moved to another folder results in the same hash. Any change to the file itself results in a different hash.
- The only file types that are affected by hash rules are those that are listed in **Designated File Types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts.

Working with Internet Zone rules

Internet zone rules apply only to Windows Installer packages. A zone rule can identify software from a zone that is specified through Internet Explorer. These zones are Internet, Local intranet, Restricted sites, Trusted sites, and My Computer. An Internet Zone rule is designed to prevent users from downloading and installing software.

To create an Internet zone rule

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Internet Zone Rule**.
3. In **Internet zone**, click an Internet zone.
4. In **Security level**, click either **Disallowed** or **Unrestricted**, and then click **OK**.

Notes

- It may be necessary to create a new software restriction policy setting for the Group Policy Object (GPO) if you have not already done so.
- Zone rules only apply to files with an .msi file type, which are Windows Installer packages.

- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts.

Working with path rules

A path rule identifies software by its file path. For example, if you have a computer that has a default security level of **Disallowed**, you can still grant unrestricted access to a specific folder for each user. You can create a path rule by using the file path and setting the security level of the path rule to **Unrestricted**. Some common paths for this type of rule are %userprofile%, %windir%, %appdata%, %programfiles%, and %temp%. You can also create registry path rules that use the registry key of the software as its path.

Because these rules are specified by the path, if a software program is moved, the path rule no longer applies.

To create a path rule

1. Open Software Restriction Policies.
2. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Path Rule**.
3. In **Path**, type a path, or click **Browse** to find a file or folder.
4. In **Security level**, click either **Disallowed** or **Unrestricted**.
5. In **Description**, type a description for this rule, and then click **OK**.

Caution

- On certain folders, such as the Windows folder, setting the security level to **Disallowed** can adversely affect the operation of your operating system. Make sure that you do not disallow a crucial component of the operating system or one of its dependent programs.

Notes

- It may be necessary to create new software restriction policies for the Group Policy Object (GPO) if you have not already done so.
- If you create a path rule for software with a security level of **Disallowed**, users can still run the software by copying it to another location.
- The wildcard characters that are supported by the path rule are * and ?.
- You can use environment variables, such as %programfiles% or %systemroot%, in the path rule.
- If you want to create a path rule for software when you do not know where it is stored on a computer but you have its registry key, you can create a registry path rule.
- To prevent users from executing e-mail attachments, you can create a path rule for your e-mail program's attachment directory that prevents users from running e-mail attachments.

- The only file types that are affected by path rules are those that are listed in **Designated File Types** in the details pane for Software Restriction Policies. There is one list of designated file types that is shared by all rules.
- For software restriction policies to take effect, users must update policy settings by logging off from and logging on to their computers.
- When more than one software restriction policies rule is applied to policy settings, there is a precedence of rules for handling conflicts.

To create a registry path rule

1. On the **Start** screen, type regedit.
2. In the console tree, right-click the registry key that you want to create a rule for, and then click **Copy Key Name**. Note the value name in the details pane.
3. Open Software Restriction Policies.
4. In either the console tree or the details pane, right-click **Additional Rules**, and then click **New Path Rule**.
5. In **Path**, paste the registry key name, followed by the value name.
6. Enclose the registry path in percent signs (%), for example, .
7. In **Security level**, click either **Disallowed** or **Unrestricted**.
8. In **Description**, type a description for this rule, and then click **OK**.

Use Software Restriction Policies to Help Protect Your Computer Against an Email Virus

This topic provides information how to set application control polices using Software Restriction Policies (SRP) to help protect your computer against e-mail virus beginning with Windows Server 2008 and Windows Vista.

Introduction

Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run. You use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. These are integrated with Microsoft Active Directory Domain Services and Group Policy but can also be configured on stand-alone computers. For a starting point for SRP, see the Software Restriction Policies.

Beginning with Windows Server® 2008 R2 and Windows® 7, Windows AppLocker can be used instead of or in concert with SRP for a portion of your application control strategy. For information about AppLocker in Windows Server 2012 and Windows 8, see the **AppLocker Overview**.

Configure SRP to help protect against an e-mail virus

1. Review the best practices for software restriction policies to understand how SRP works.
   - Best practices
   - How Software Restriction Policies Work
2. Open Software Restriction Policies.
   - For your local computer
   - For a domain, site, or organizational unit, and you are on a member server or on a workstation that is joined to a domain
3. If you have not previously defined software restriction policies, create new software restriction policies.
   - To create new software restriction policies
4. Create a path rule for the folder that your e-mail program uses to run e-mail attachments, and then set the security level to **Disallowed**.
   - Working with path rules
5. Specify the file types to which the rule applies.
   - To add or delete a designated file type
6. Modify policy settings so that they apply to the users and groups that you want:
   - Specify users or groups to which you do not want the Group Policy Object's (GPO) policy settings to apply.
   - Exclude local administrators from the software restriction policies of a specific policy setting in Group Policy and still have the rest of Group Policy apply to the administrators.
     - To prevent software restriction policies from applying to local administrators
7. Test the policy.

Troubleshoot Software Restriction Policies

This topic describes common problems and their solutions when troubleshooting Software Restriction Policies (SRP) beginning with Windows Server 2008 and Windows Vista.

Introduction

Software Restriction Policies (SRP) is Group Policy-based feature that identifies software programs running on computers in a domain, and controls the ability of those programs to run.

You use software restriction policies to create a highly restricted configuration for computers, in which you allow only specifically identified applications to run. These are integrated with Microsoft Active Directory Domain Services and Group Policy but can also be configured on stand-alone computers. For more information about SRP, see the Software Restriction Policies.

Beginning with Windows Server® 2008 R2 and Windows® 7, Windows AppLocker can be used instead of or in concert with SRP for a portion of your application control strategy. For information about AppLocker in Windows Server 2012 and Windows 8, see the **AppLocker Overview**.

Windows cannot open a program

Users receive a message that says "Windows cannot open this program because it has been prevented by a software restriction policy. For more information, open Event Viewer or contact your system administrator." Or, on the command line, a message says "The system cannot execute the specified program."

Cause: The default security level (or a rule) was created so that the software program is set as **Disallowed**, and as a result it will not start.

Solution: Look in the event log for an in-depth description of the message. The event log message indicates what software program is set as **Disallowed** and what rule is applied to the program.

Modified software restriction policies are not taking effect

Cause: Software restriction policies that are specified in a domain through Group Policy override any policy settings that are configured locally. This might imply that there is a policy setting from the domain that is overriding your policy setting.

Cause: Group Policy might not have refreshed its policy settings. Group Policy applies changes to policy settings periodically; therefore, it is likely that the policy changes that were made in the directory have not yet been refreshed.

Solutions:

1. The computer on which you modify software restriction policies for the network must be able to contact a domain controller. Ensure the computer can contact a domain controller.
2. Refresh policy by logging off of the network and then logging on to the network again. If any policy is applied through Group Policy, logging back in will refresh those policies.
3. You can refresh policy settings with the command-line utility gpupdate or by logging off from and then logging back on to your computer. For best results, run gpupdate, and then log off from and log back on to your computer. Generally, the security settings are refreshed every 90 minutes on a workstation or server and every 5 minutes on a domain controller. The settings are also refreshed every 16 hours, whether or not there are any changes. These are configurable settings so refresh intervals might be different in each domain.
4. Check which policies apply. Check domain level policies for **No Override** settings.
5. Software restriction policies that are specified in a domain through Group Policy override any policies that are configured locally. Use Gpresult command-line tool to determine what the

net effect of the policy is. This might imply that there is a policy from the domain that is overriding your local setting.

6. If SRP and AppLocker policy settings are in the same GPO, AppLocker settings will take precedence on Windows 7, Windows Server 2008 R2, and later. It is recommended to put SRP and AppLocker policy settings in different GPOs.

After adding a rule through SRP, you cannot log on to your computer

Cause: Your computer accesses many programs and files when it starts. You might have inadvertently set one of these programs or files to **Disallowed**. Because the computer cannot access the program or file, it cannot start properly.

Solution: Start the computer in Safe Mode, log on as a local administrator, and then change software restriction policies to allow the program or file to run.

A new policy setting is not applying to a specific file name extension

Cause: The filename extension is not in the list of supported file types.

Solution: Add the filename extension to the list of file types supported by SRP.

Software restriction policies address the problem of regulating unknown or untrusted code. Software restriction policies are security settings to identify software and control its ability to run on a local computer, in a site, domain, or OU and can be implemented through a GPO.

A default rule is not restricting as expected

Cause: Rules which are applied in a particular order which can cause default rules to be overridden by specific rules. SRP applies rules in the following order (most specific to general):

1. Hash rules
2. Certificate rules
3. Path rules
4. Internet Zone rules
5. Default rules

Solution: Evaluate the rules restricting the application and, if appropriate, remove all but the Default rule.

Unable to discover which restrictions are applied

Cause: There is no apparent cause for the unexpected behavior, and GPO refresh has not solved the issue so further investigation is necessary.

Solutions:

1. Investigate the System Event Log, filtering on source of “Software Restriction Policy.” The entries explicitly state which rule is implemented for each application.

2. Enable advanced logging. See Determine Allow/Deny List and Application Inventory for Software Restriction Policies for more information.

TLS/SSL (Schannel SSP) Overview

This topic for the IT professional introduces the TLS/SSL implementation in Windows using the Schannel Security Service Provider (SSP) by describing practical applications, changes in Microsoft's implementation, and software requirements, plus additional resources for Windows Server 2012 and Windows 8.

Did you mean…

- The Schannel security package
- Secure Channel
- Transport Layer Security Protocol

TLS/SSL (Schannel) description

Schannel is a Security Support Provider (SSP) that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) Internet standard authentication protocols.

The Security Support Provider Interface (SSPI) is an API used by Windows systems to perform security-related functions including authentication. The SSPI functions as a common interface to several Security Support Providers (SSPs), including the Schannel SSP.

The Transport Layer Security (TLS) protocol versions 1.0, 1.1, and 1.2, Secure Sockets Layer (SSL) protocol, versions 2.0 and 3.0, Datagram Transport Layer Security (DTLS) version 1.0, and the Private Communications Transport (PCT) protocol are based on public key cryptography. The Security Channel (Schannel) authentication protocol suite provides these protocols. All Schannel protocols use a client/server model.

Practical applications

One problem when you administer a network is securing data that is being sent between applications across an untrusted network. You can use TLS/SSL to authenticate servers and client computers and then use the protocol to encrypt messages between the authenticated parties.

For example, you can use TLS/SSL for:

- SSL-secured transactions with an e-commerce website
- Authenticated client access to an SSL-secured website
- Remote access
- SQL access
- E-mail

New and changed functionality

The following table notes the changes to the Schannel SSP to allow for the Datagram Transport Layer Security (DTLS) (RFC 4347), Server Name Indicator, configurable hints for certificates, and ability to configure the Schannel provider to use specific trusted root stores.

Feature/functionality	Windows Server 2008 R2	Windows Server 2012
Inclusion of TLS 1.2	X	X
Inclusion of DTLS		X
TLS support for Server Name Indicator (SNI) extensions		X
Manageability improvements to configure a list of certificate roots to be used by a website as the trust anchors for the purposes of validating a client certificate		X
Manageability improvements to configure a list of certificates hints for eventual selection by the client computer		X

For information about these changes in functionality for TLS/SSL, see What's New in TLS/SSL (Schannel SSP).

Deprecated functionality

In the Schannel SSP for Windows Server 2012 and Windows 8, there are no deprecated features or functionality. The Private Communications Transport (PCT) protocol is disabled by default, as it was in the past version.

Software requirements

The TLS/SSL protocol use a client/server model and are based on certificate authentication, which requires a Public Key Infrastructure.

Server Manager information

There are no configuration steps necessary to implement TLS, SSL or Schannel using Server Manager or the Add Features feature.

See also

The following table provides links to additional resources related to TLS, SSL and the Schannel SSP.

Content type	References
Product evaluation	What's New in TLS/SSL (Schannel SSP)
Planning	TLS/SSL Technical Reference (2003)
Deployment	Not yet available
Operations	Not yet available
Troubleshooting	Not yet available
Security	Not yet available
Tools and settings	Not yet available
Community resources	Private Cloud Security Model - Wrapper Functionality
Related technologies	Active Directory Certificate Services Overview

What's New in TLS/SSL (Schannel SSP)

This topic for the IT professional describes the changes in functionality in the Schannel Security Support Provider (SSP), which includes the Transport Layer Security (TLS), the Secure Sockets Layer (SSL), and the Datagram Transport Layer Security (DTLS) authentication protocols, for Windows Server 2012 R2, Windows Server 2012, Windows 8.1, and Windows 8.

TLS/SSL (Schannel SSP) description

Schannel is a Security Support Provider (SSP) that implements the SSL, TLS and DTLS internet standard authentication protocols.

The Security Support Provider Interface (SSPI) is an API used by Windows systems to perform security-related functions including authentication. The SSPI functions as a common interface to several Security Support Providers (SSPs), including the Schannel SSP.

For more information about Microsoft's implementation of TLS and SSL in the Schannel SSP, see the TLS/SSL Technical Reference (2003).

New and changed functionality in Windows Server 2012 R2

The following describes what changes in functionality were made to TLS in the Schannel SSP.

TLS session resumption

The Transport Layer Security (TLS) protocol, a component of the Schannel Security Support Provider, is used to secure data that is sent between applications across an untrusted network. TLS/SSL can be used to authenticate servers and client computers, and also to encrypt messages between the authenticated parties.

Devices that connect TLS to servers frequently need to reconnect due to session expiration. Windows 8.1 and Windows Server 2012 R2 now support RFC 5077 (TLS Session Resumption without Server-Side State). This modification provides Windows Phone and Windows RT devices with:

- Reduced resource usage on the server
- Reduced bandwidth, which improves the efficiency of client connections
- Reduced time spent for the TLS handshake due to resumptions of the connection.

Note

The client-side implementation of RFC 5077 was added in Windows 8.

For information about stateless TLS session resumption, see the IETF document RFC 5077.

Application Protocol Negotiation

Windows Server 2012 R2 and Windows 8.1 support client-side TLS application protocol negotiation so applications can leverage protocols as part of the HTTP 2.0 standard development and users can access online services such as Google and Twitter using apps running the SPDY protocol.

How it works

Client and server applications enable application protocol negotiation extension by supplying lists of supported application protocol IDs, in descending order of preference. The TLS client indicates that it supports application protocol negotiation by including the Application Layer Protocol Negotiation (ALPN) extension with a list of client-supported protocols in the ClientHello message.

When the TLS client makes the request to the server, the TLS server reads its supported protocol list for the most-preferred application protocol which the client also supports. If such a protocol is found, the server responds with the selected protocol ID and continues with the handshake as usual. If there is no common application protocol, the server sends a fatal handshake failure alert.

New and changed functionality in Windows Server 2012

The following describes what changes in functionality were made to the Schannel.

- Management of trusted issuers for client authentication
- TLS support for Server Name Indicator (SNI) Extensions
- Datagram Transport Layer Security (DTLS)
- Deprecated functionality

Management of trusted issuers for client authentication

When authentication of the client computer is required using SSL or TLS, the server can be configured to send a list of trusted certificate issuers. This list contains the set of certificate issuers which the server will trust and is a hint to the client computer as to which client certificate to select if there are multiple certificates present. In addition, the certificate chain the client computer sends to the server must be validated against the configured trusted issuers list.

Prior to Windows Server 2012 and Windows 8, applications or processes that used the Schannel SSP (including HTTP.sys and IIS) could provide a list of the trusted issuers that they supported for Client Authentication through a Certificate Trust List (CTL).

In Windows Server 2012 and Windows 8, changes were made to the underlying authentication process so that:

- CTL-based trusted issuer list management is no longer supported
- The behavior to send the Trusted Issuer List by default is off: Default value of the SendTrustedIssuerList registry key is now 0 (off by default) instead of 1
- Compatibility to previous versions of Windows operating systems is preserved

What value does this change add?

Beginning with Windows Server 2012, the use of the CTL has been replaced with a certificate store-based implementation. This allows for more familiar manageability through the existing certificate management commandlets of the PowerShell provider, as well as command line tools such as certutil.exe.

Although the maximum size of the trusted certification authorities list that the Schannel SSP supports (16 KB) remains the same as in Windows Server 2008 R2, in Windows Server 2012 there is a new dedicated certificate store for client authentication issuers so that unrelated certificates are not included in the message.

What works differently?

In Windows Server 2012, the trusted issuers list is configured using certificate stores; one default global computer certificate store and one that is optional per site. The source of the list will be determined as follows:

- If there is a specific credential store configured for the site, it will be used as the source
- If no certificates exist in the application-defined store, then Schannel checks the **Client Authentication Issuers** store on the local computer and, if certificates are present, uses that

store as the source. If no certificate is found in either store, the Trusted Roots store is checked.

- If neither the global or local stores contain certificates, the Schannel provider will use the **Trusted Root Certifictation Authorities** store as the source of trusted issuers list. (This is the behavior for Windows Server 2008 R2.)

If the **Trusted Root Certifictation Authorities** store that was used contains a mix of Root (self-signed) and certification authority (CA) Issuer certificates, only the CA Issuer certificates will be sent to the server by default.

How to configure Schannel to use the certificate store for Trusted Issuers

The Schannel SSP architecture in Windows Server 2012 will by default use the stores as described above to manage the Trusted Issuers list. You can still use the existing certificate management commandlets of the PowerShell provider, as well as command line tools such as Certutil to manage certificates.

For information about managing certificates using the PowerShell provider, see AD CS Administration Cmdlets in Windows.

For information about managing certificates using the certificate utility, see certutil.exe.

For information about what data, including the application-defined store, is defined for an Schannel credential, see SCHANNEL_CRED structure (Windows).

Defaults for Trust Modes

There are three Client Authentication Trust Modes supported by the Schannel provider. The trust mode controls how validation of the client's certificate chain is performed and is a system-wide setting controlled by the REG_DWORD "ClientAuthTrustMode" under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\Schannel.

Value	Trust Mode	Description
0	Machine Trust (default)	Requires that the client certificate is issued by a certificate in the Trusted Issuers list.
1	Exclusive Root Trust	Requires that a client certificate chains to a root certificate contained in the caller-specified trusted issuer store. The certificate must also be issued by an issuer in the Trusted Issuers list
2	Exclusive CA Trust	Requires that a client certificate chain to either an intermediate CA certificate or root certificate in the caller-specified trusted

Value	Trust Mode	Description
		issuer store.

For information about authentication failures due to trusted issuers configuration issues, see Knowledge Base article 280256.

TLS support for Server Name Indicator (SNI) Extensions

Server Name Indication feature extends the SSL and TLS protocols to allow proper identification of the server when numerous virtual images are running on a single server. To properly secure the communication between a client computer and a server, the client computer requests a digital certificate from the server. After the server responds to the request and sends the certificate, the client computer examines it, uses it to encrypt the communication, and proceeds with the normal request-response exchange. However, in a virtual hosting scenario, several domains, each with its own potentially distinct certificate, are hosted on one server. In this case, the server has no way of knowing beforehand which certificate to send to the client computer. SNI allows the client computer to inform the target domain earlier in the protocol, and this allows the server to correctly select the proper certificate.

What value does this change add?

This additional functionality:

- Allows you to host multiple SSL websites on a single IP and port combination
- Reduces the memory usage when multiple SSL websites are hosted on a single web server
- Allows more users to connect to my SSL websites simultaneously
- Permits you to provide hints to end users through the computer interface for selecting the correct certificate during a client authentication process.

What works differently?

The Schannel SSP maintains an in-memory cache of client connection states permitted for clients. This allows client computers to reconnect quickly to the SSL server without subject to a full SSL handshake on subsequent visits. This efficient use of certificate management permits more sites to be hosted on a single Windows Server 2012 compared to previous operating system versions.

Certificate selection by the end user has been enhanced by allowing you to construct a list of probable certificate issuer names that provide the end user with hints on which one to choose. This list is configurable using Group Policy.

Datagram Transport Layer Security (DTLS)

The DTLS version 1.0 protocol has been added to the Schannel Security Support Provider. The DTLS protocol provides communications privacy for datagram protocols. The protocol allows client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, or message forgery. The DTLS protocol is based on the Transport Layer Security

(TLS) protocol and provides equivalent security guarantees, reducing the need to use IPsec or designing a custom application layer security protocol.

What value does this change add?

Datagrams are common in streaming media – such as gaming or secured video conferencing. Adding the DTLS protocol to the Schannel provider gives you the ability to use the familiar Windows SSPI model in securing the communication between client computers and servers. DTLS is deliberately designed to be as similar to TLS as possible, both to minimize new security invention and to maximize the amount of code and infrastructure reuse.

What works differently?

Applications that use DTLS over UDP can use the SSPI model in Windows Server 2012 and Windows 8. Certain cipher suites are available for configuration, similar to how you can configure TLS. Schannel continues to use the CNG cryptographic provider which takes advantage of FIPS 140 certification, which was introduced in Windows Vista.

Deprecated functionality

In the Schannel SSP for Windows Server 2012 and Windows 8, there are no deprecated features or functionality.

See also

The table provides additional resources for evaluating the Schannel Security Support Provider, TLS and SSL.

Content type	References
Product evaluation	TLS/SSL (Schannel SSP) Overview
Community resources	Private Cloud Security Model - Wrapper Functionality
Related technologies	Active Directory Certificate Services Overview

Trusted Platform Module Technology Overview

This topic for the IT professional describes the Trusted Platform Module (TPM) and how Windows uses it for access control and authentication. The topic provides links to other resources about the TPM.

Feature description

Trusted Platform Module (TPM) technology is designed to provide hardware-based, security-related functions. A TPM chip is a secure crypto-processor that is designed to carry out cryptographic operations. The chip includes multiple physical security mechanisms to make it tamper resistant, and malicious software is unable to tamper with the security functions of the TPM. Some of the key advantages of using TPM technology are that you can:

- Generate, store, and limit the use of cryptographic keys.
- Use TPM technology for platform device authentication by using the TPM's unique RSA key, which is burned into itself.
- Help ensure platform integrity by taking and storing security measurements.

The most common TPM functions are used for system integrity measurements and for key creation and use. During the boot process of a system, the boot code that is loaded (including firmware and the operating system components) can be measured and recorded in the TPM. The integrity measurements can be used as evidence for how a system started and to make sure that a TPM-based key was used only when the correct software was used to boot the system.

TPM-based keys can be configured in a variety of ways. One option is to make a TPM-based key unavailable outside the TPM. This is good to mitigate phishing attacks because it prevents the key from being copied and used without the TPM. TPM-based keys can also be configured to require an authorization value to use them. If too many incorrect authorization guesses occur, the TPM will activate its dictionary attack logic and prevent further authorization value guesses.

Different versions of the TPM are defined in specifications by the Trusted Computing Group (TCG). For more information, consult the TCG Web site (http://www.trustedcomputinggroup.org/developers/trusted_platform_module).

The Supported versions of the Windows operating system can automatically provision and manage the TPM. Group Policy settings can be configured to control whether the TPM owner authorization value is backed up in Active Directory. Because the TPM state persists across operating system installations, TPM information is stored in a location in Active Directory that is separate from computer objects. Depending on an enterprise's security goals, Group Policy can be configured to allow or prevent local administrators from resetting the TPM's dictionary attack logic. Standard users can use the TPM, but Group Policy controls limit how many authorization failures standard users can attempt so that one user is unable to prevent other users or the administrator from using the TPM. TPM technology can also be used as a virtual smart card and for secure certificate storage. With BitLocker Network Unlock, domain-joined computers are not prompted for a BitLocker PIN.

Practical applications

Certificates can be installed or created on computers that are using the TPM. After a computer is provisioned, the RSA private key for a certificate is bound to the TPM and cannot be exported. The TPM can also be used as a replacement for smart cards, which reduces the costs associated with creating and disbursing smart cards.

Automated provisioning in the TPM reduces the cost of TPM deployment in an enterprise. New APIs for TPM management can determine if TPM provisioning actions require physical presence of a service technician to approve TPM state change requests during the boot process.

Antimalware software can use the boot measurements of the operating system start state to prove the integrity of a computer running Windows 8.1, Windows 8 Windows Server 2012 R2 or Windows Server 2012. These measurements include the launch of Hyper-V to test that datacenters using virtualization are not running untrusted hypervisors. With BitLocker Network Unlock, IT administrators can push an update without concerns that a computer is waiting for PIN entry.

The TPM has several Group Policy settings that can be used to manage how it is used. These settings can be used to manage the owner authorization value, the blocked TPM commands, the standard user lockout, and the backup of the TPM to AD DS. For additional details, see **TPM Group Policy Settings**.

New and changed functionality

For information about new and changed functionality in the TPM for Windows Server 2012 R2 and Windows 8.1, see **What's New for the TPM in Windows 8.1**.

The following sections describe the new and changed functionalities in the TPM for Windows Server 2012 and Windows 8:

- Automated provisioning and management of the TPM
- Measured Boot with support for attestation
- TPM-based Virtual Smart Card
- TPM Owner Authorization Value
- TPM Cmdlets

Automated provisioning and management of the TPM

Windows Server 2012 and Windows 8 deliver a number of enhancements that streamline TPM provisioning, making it easier to deploy systems that are ready for BitLocker and other TPM-dependent features. These enhancements include simplifying the TPM state model to report **Ready**, **Ready with reduced functionality**, or **Not ready**. Windows Server 2012 and Windows 8 also include automatic provisioning of TPMs in the **Ready** state, remote provisioning to remove the requirement for the physical presence of a technician for the initial deployment. In addition, the TPM stack is available in the Windows Preinstallation Environment (Windows PE).

A number of management settings have been added for easier management and configuration of the TPM through Group Policy. The primary new settings include Active Directory-based backup of TPM owner authentication, the level of owner authentication that should be stored locally on the TPM, and the software-based TPM lockout settings for standard users. For more information about backing up owner authentication to Windows Server 2008 R2 AD DS domains, see

Schema Extensions for Windows Server 2008 R2 to support AD DS backup of TPM information from Windows 8 clients

Measured Boot with support for attestation

The Measured Boot feature provides antimalware software with a trusted (resistant to spoofing and tampering) log of all boot components. Antimalware software can use the log to determine whether components that ran before it are trustworthy versus infected with malware. It can also send the Measured Boot logs to a remote server for evaluation. The remote server can initiate remediation actions by interacting with software on the client or through out-of-band mechanisms, as appropriate.

TPM-based Virtual Smart Card

The Virtual Smart Card emulates the functionality of traditional smart cards, but Virtual Smart Cards use the TPM chip that is available on an organization's computers, rather than requiring the use of a separate physical smart card and reader. This greatly reduces the management and deployment cost of smart cards in an enterprise. To the end user, the Virtual Smart Card is always available on the computer. If a user needs to use more than one computer, a Virtual Smart Card must be issued to the user for each computer. A computer that is shared among multiple users can host multiple Virtual Smart Cards, one for each user.

For more information, see **Virtual Smart Card Overview**.

TPM-based certificate storage

The TPM can be used to protect certificates and RSA keys. Windows Server 2012 includes a new key storage provider (KSP), which provides easy, convenient use of the TPM as a way of strongly protecting private keys. The TPM KSP can be used to generate keys when an organization enrolls for certificates, and the KSP is managed by templates in the UI. The TPM can also be used to protect certificates that are imported from an outside source. TPM-based certificates can be used exactly as standard certificates with the added functionality that the certificate can never leave the TPM from which the keys were generated. The TPM can now be used for crypto-operations through Cryptography API: Next Generation (CNG). For more information, see Cryptography API: Next Generation.

TPM Owner Authorization Value

For Windows 8 a change to how the TPM owner authorization value is stored in AD DS was implemented in the AD DS schema. The TPM owner authorization value is now stored in a separate object which is linked to the Computer object. This value was stored as a property in the Computer object itself for the default Windows Server 2008 R2 schemas. Windows Server 2012 domain controllers have the default schema to backup TPM owner authorization information in the separate object. If you are not upgrading your domain controller to Windows Server 2012 you need to extend the schema to support this change. If Active Directory backup of the TPM owner

authorization value is enabled in a Windows Server 2008 R2 environment without extending the schema, the TPM provisioning will fail and the TPM will remain in a Not Ready state for computers running Windows 8.

If your computer is not being joined to a domain the TPM owner authorization value will be stored in the local computer registry. Using BitLocker to encrypt the operating system drive will protect the owner authorization value from being disclosed when the computer is at rest, but there is a risk that a malicious user could obtain the TPM owner authorization value when the computer is unlocked. Therefore, we recommend that in this situation you configure your computer to automatically lock after 30 seconds of inactivity. If automatic locking is not used, then you should consider removing full owner authorization from the computer registry.

Registry information

Registry key: HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\TPM

DWORD: OSManagedAuthLevel

Value Data	Setting
0	None
2	Delegated
4	Full

Note

If the operating system managed TPM authentication setting is changed from "Full" to "Delegated" the full TPM owner authorization value will be regenerated and any copies of the original TPM owner authorization value will be invalid. If you are backing up the TPM owner authorization value to AD DS, the new owner authorization value will be automatically backed up to AD DS when it is changed.

TPM Cmdlets

If you are using PowerShell to script and manage your computers, you can now manage the TPM using Windows PowerShell as well. To install the TPM cmdlets use the following command:

dism /online /enable-feature /FeatureName:tpm-psh-cmdlets

For details about the individual cmdlets, see TPM Cmdlets in Windows PowerShell

Supported versions

TPM version	Windows Server 2012 R2, Windows 8.1 and Windows RT	Windows Server 2012, Windows 8 and Windows RT	Windows Server 2008 R2 and Windows 7
TPM 1.2	X	X	X
TPM 2.0	X	X	

Additional Resources

TPM Fundamentals

How the TPM mitigates dictionary attacks

Trusted Platform Module Services Group Policy Settings

TPM Cmdlets in Windows PowerShell

Schema Extensions for Windows Server 2008 R2 to support AD DS backup of TPM information from Windows 8 clients

Prepare your organization for BitLocker: Planning and Policies - TPM configurations

What's New in Smart Cards

Schema Extensions for Windows Server 2008 R2 to support AD DS backup of TPM information from Windows 8 clients

For Windows 8 a change to how the TPM owner authorization value is stored in AD DS was implemented in the AD DS schema. This topic provides more details about this change and provides template schema extensions that you can incorporate into your organization as appropriate:

Why a schema extension is needed

The TPM owner authorization value is now stored in a separate object which is linked to the Computer object. This value was stored as a property in the Computer object itself for the default Windows Server 2008 R2 schemas. Windows Server 2012 domain controllers have the default schema to backup TPM owner authorization information in the separate object. If you are not upgrading your domain controller to Windows Server 2012 you need to extend the schema to support this change. If Active Directory backup of the TPM owner authorization value is enabled in a Windows Server 2008 R2 environment without extending the schema, the TPM provisioning will fail and the TPM will remain in a Not Ready state for computers running Windows 8. The

following are the two schema extensions that you can use to bring your Windows Server 2008 R2 domain to parity with Windows Server 2012:

TpmSchemaExtension.ldf

This schema extension brings parity with the Windows Server 2012 schema and is required if you want to store the TPM owner authorization value for a computer running Windows 8 in a Windows Server 2008 R2 AD DS domain. With this extension the TPM owner authorization information will be stored in a separate TPM object linked to the corresponding computer object.

```
#============================================================================

#

# Active Directory Domain Services schema extension for

# BitLocker Drive Encryption and Trusted Platform Module (TPM) recovery

#

# This file contains attributes and class objects that enable Windows Server

# 2008 and Windows Server 2008 R2 domain controllers to store TPM recovery

# information in a new, TPM-specific location.

#

# Change History:

#   07/2010 - Created

#

# To extend the schema, use the LDIFDE tool on the schema master of the forest.

#

# Sample command:

#   ldifde -i -v -f TPMSchemaExtension.ldf -c "DC=X" "DC=nttest,dc=microsoft,dc=com" -k -

j .

#

# For more information on LDIFDE tool, see

# http://support.microsoft.com/default.aspx?scid=kb;en-us;237677

#

#============================================================================


#============================================================================

# New schema attributes

#============================================================================
```

```
#
# ms-TPM-Srk-Pub-Thumbprint
# GUID: 19d706eb-4d76-44a2-85d6-1c342be3be37
#
dn: CN=ms-TPM-Srk-Pub-Thumbprint,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: attributeSchema
ldapDisplayName: msTPM-SrkPubThumbprint
adminDisplayName: TPM-SrkPubThumbprint
adminDescription: This attribute contains the thumbprint of the SrkPub corresponding to a
particular TPM. This helps to index the TPM devices in the directory.
attributeId: 1.2.840.113556.1.4.2107
attributeSyntax: 2.5.5.10
omSyntax: 4
isSingleValued: TRUE
searchFlags: 11
schemaIdGuid:: 6wbXGXZNokSF1hw0K+O+Nw==
showInAdvancedViewOnly: TRUE
isMemberOfPartialAttributeSet: FALSE
rangeUpper: 20

#
# ms-TPM-Owner-Information-Temp
# GUID: c894809d-b513-4ff8-8811-f4f43f5ac7bc
#
dn: CN=ms-TPM-Owner-Information-Temp,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: attributeSchema
ldapDisplayName: msTPM-OwnerInformationTemp
adminDisplayName: TPM-OwnerInformationTemp
adminDescription: This attribute contains temporary owner information for a particular
TPM.
attributeId: 1.2.840.113556.1.4.2108
attributeSyntax: 2.5.5.12
```

```
omSyntax: 64
isSingleValued: TRUE
searchFlags: 640
rangeUpper: 128
schemaIdGuid:: nYCUyBO1+E+IEfT0P1rHvA==
showInAdvancedViewOnly: TRUE
isMemberOfPartialAttributeSet: FALSE

#
# ms-TPM-Tpm-Information-For-Computer
# GUID: ea1b7b93-5e48-46d5-bc6c-4df4fda78a35
#
dn: CN=ms-TPM-Tpm-Information-For-Computer,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: attributeSchema
ldapDisplayName: msTPM-TpmInformationForComputer
adminDisplayName: TPM-TpmInformationForComputer
adminDescription: This attribute links a Computer object to a TPM object.
attributeId: 1.2.840.113556.1.4.2109
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: TRUE
searchFlags: 16
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: k3sb6khe1Ua8bE30/aeKNQ==
showInAdvancedViewOnly: TRUE
isMemberOfPartialAttributeSet: FALSE
linkId: 2182

#
# ms-TPM-TpmInformation-For-Computer-BL
# GUID: 14fa84c9-8ecd-4348-bc91-6d3ced472ab7
#
dn: CN=ms-TPM-Tpm-Information-For-Computer-BL,CN=Schema,CN=Configuration,DC=X
```

```
changetype: add
objectClass: attributeSchema
ldapDisplayName: msTPM-TpmInformationForComputerBL
adminDisplayName: TPM-TpmInformationForComputerBL
adminDescription: This attribute links a TPM object to the Computer objects associated
with it.
attributeId: 1.2.840.113556.1.4.2110
attributeSyntax: 2.5.5.1
omSyntax: 127
isSingleValued: FALSE
searchFlags: 0
omObjectClass:: KwwCh3McAIVK
schemaIdGuid:: yYT6FM2OSEO8kW087Ucqtw==
showInAdvancedViewOnly: TRUE
systemOnly: TRUE
linkId: 2183

#
# Commit the new attributes
#

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

#
# Modify the Computer schema to support the TPM link
#

dn: CN=computer,CN=Schema,CN=Configuration,DC=X
changetype: modify
add: mayContain
```

```
mayContain: msTPM-TpmInformationForComputer
-

#
# Commit the modification to the computer class
#

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-

#============================================================================
# New schema classes
#============================================================================

#
# ms-TPM-Information-Objects-Container
# GUID: e027a8bd-6456-45de-90a3-38593877ee74
#
dn: CN=ms-TPM-Information-Objects-Container,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: classSchema
ldapDisplayName: msTPM-InformationObjectsContainer
adminDisplayName: TPM-InformationObjectsContainer
adminDescription: Container for TPM objects.
governsID: 1.2.840.113556.1.5.276
objectClassCategory: 1
subClassOf: top
systemMustContain: cn
systemPossSuperiors: domain
systemPossSuperiors: domainDNS
schemaIdGUID:: vagn4FZk3kWQozhZOHfudA==
```

```
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;LOLCCCRP;;;
DC)
defaultHidingValue: TRUE
defaultObjectCategory: CN=ms-TPM-Information-Objects-
Container,CN=Schema,CN=Configuration,DC=X

#
# ms-TPM-Information-Object
# GUID: 85045b6a-47a6-4243-a7cc-6890701f662c
#
# NOTE: If the 'defaultSecurityDescriptor' value below is changed,
#   also change the other '.ldf' files in this directory, as appropriate.
#
dn: CN=ms-TPM-Information-Object,CN=Schema,CN=Configuration,DC=X
changetype: add
objectClass: classSchema
ldapDisplayName: msTPM-InformationObject
adminDisplayName: TPM-InformationObject
adminDescription: This class contains recovery information for a Trusted Platform Module
(TPM) device.
governsID: 1.2.840.113556.1.5.275
objectClassCategory: 1
subClassOf: top
systemMustContain: msTPM-OwnerInformation
systemMayContain: msTPM-SrkPubThumbprint
systemMayContain: msTPM-OwnerInformationTemp
systemPossSuperiors: 1.2.840.113556.1.5.276
schemaIdGUID:: alsEhaZHQ0KnzGiQcB9mLA==
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPLO;;;DC)(
A;;WP;;;CO)
defaultHidingValue: TRUE
defaultObjectCategory: CN=ms-TPM-Information-Object,CN=Schema,CN=Configuration,DC=X
#
```

```
# NOTE: If the 'defaultSecurityDescriptor' value above is changed,
#   also change the other '.ldf' files in this directory, as appropriate.
#


#
# Commit the new TPM object class
#


dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-


#=============================================================================
# New objects
#=============================================================================


#
# Add the TPM container to its location in the directory
#
dn: CN=TPM Devices,DC=X
changetype: add
objectClass: msTPM-InformationObjectsContainer
```

You should be aware that only the Computer object that has created the TPM object can update it. This means that any subsequent updates to the TPM objects will not succeed in dual boot scenarios or scenarios where the computer is reimaged resulting in a new AD computer object being created. If you are planning to support such scenarios, you will need to update the schema further as shown in the schema extension example, TpmSchemaExtensionACLChanges.ldf.

TpmSchemaExtensionACLChanges.ldf

This schema update modifies the ACLs on the TPM object to be less restrictive so that any subsequent operating system which takes ownership of the computer object can update the owner authorization value in AD DS.

Important

After implementing this schema update, any computer in the domain can update the OwnerAuth of the TPM object (although it cannot read the OwnerAuth). When using this extension, perform a regular backup of the TPM objects and enable auditing to track the changes for these objects.

```
#============================================================================

#

# Active Directory Domain Services schema extension for

# BitLocker Drive Encryption and Trusted Platform Module (TPM) recovery

#

# This file modifies a class object that enables Windows Server 2008

# and Windows Server 2008 R2 domain controllers to store TPM recovery

# information in a new, TPM-specific location.

#

# This file converts the standard schema extension in which only the creator

# of an 'ms-TPM-Information-Object' can write to the object to the Open

# schema extension in which any Domain Computer can write to the object.

#

# This conversion does not apply to any 'ms-TPM-Information-Object' that

# was created before the conversion.


#

# Change History:

#   12/2011 - Created

#

# To change the schema, use the LDIFDE tool on the schema master of the forest.

#

# Sample command:

#   ldifde -i -v -f TpmSchemaExtensionACLChanges.ldf

#      -c "DC=X" "DC=nttest,dc=microsoft,dc=com" -k -j .

#

# For more information on LDIFDE tool, see

# http://support.microsoft.com/default.aspx?scid=kb;en-us;237677

#

#============================================================================
```

```
#
# Modify the TPM-Information-Object class schema 'defaultSecurityDescriptor' to
# allow any Domain Computer to write its properties (including the TPM OwnerAuth
# value) from allowing only the creating Computer object to write its properties
#
# NOTE: Keep any changes to the 'defaultSecurityDescriptor' value in synchronization
#   with the value in the TPM-Information-Object class description in the
#   'TpmSchemaExtension.ldf' file
#

dn: CN=ms-TPM-Information-Object,CN=Schema,CN=Configuration,DC=X
changetype: modify
replace: defaultSecurityDescriptor
defaultSecurityDescriptor:
D:(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;DA)(A;;RPWPCRCCDCLCLORCWOWDSDDTSW;;;SY)(A;;RPWPLO;;;DC
)
-

#
# Commit the modification to the TPM-Information-Object schema
#

dn:
changetype: modify
add: schemaUpdateNow
schemaUpdateNow: 1
-
```

User Account Control Overview

User Account Control (UAC) is a fundamental component of Microsoft's overall security vision. UAC helps mitigate the impact of a malicious program.

Feature description

UAC allows all users to log on to their computers using a standard user account. Processes launched using a standard user token may perform tasks using access rights granted to a standard user. For instance, Windows Explorer automatically inherits standard user level permissions. Additionally, any programs that are executed using Windows Explorer (for example, by double-clicking an application shortcut) also run with the standard set of user permissions. Many applications, including those that are included with the operating system itself, are designed to work properly in this way.

Other applications, especially those that were not specifically designed with security settings in mind, often require additional permissions to run successfully. These types of programs are referred to as legacy applications. Additionally, actions such as installing new software and making configuration changes to programs such as Windows Firewall, require more permissions than what is available to a standard user account.

When an applications needs to run with more than standard user rights, UAC can restore additional user groups to the token. This enables the user to have explicit control of programs that are making system level changes to their computer or device.

Practical applications

Admin Approval Mode in UAC helps prevent malicious programs from silently installing without an administrator's knowledge. It also helps protect from inadvertent system-wide changes. Lastly, it can be used to enforce a higher level of compliance where administrators must actively consent or provide credentials for each administrative process.

New and changed functionality

In Windows Server 2012 UAC functionality is improved to:

- Allow a user with administrator privileges to configure the UAC experience in the Control Panel.
- Provide additional local security policies that enable a local administrator to change the behavior of the UAC messages for local administrators in Admin Approval Mode.
- Provide additional local security policies that enable a local administrator to change the behavior of the UAC messages for standard users.

How User Account Control Works

User Account Control (UAC) helps prevent malicious programs (also called malware) from damaging a computer and helps organizations deploy a better-managed desktop. With UAC, applications and tasks always run in the security context of a non-administrator account, unless an administrator specifically authorizes administrator-level access to the system. UAC can block

the automatic installation of unauthorized applications and prevent inadvertent changes to system settings.

UAC Process and Interactions

Each application that requires the administrator access token must prompt the administrator for consent. The one exception is the relationship that exists between parent and child processes. Child processes inherit the user access token from the parent process. Both the parent and child processes, however, must have the same integrity level. Windows Server 2012 protects processes by marking their integrity levels. Integrity levels are measurements of trust. A "high" integrity application is one that performs tasks that modify system data, such as a disk partitioning application, while a "low" integrity application is one that performs tasks that could potentially compromise the operating system, such as a Web browser. Applications with lower integrity levels cannot modify data in applications with higher integrity levels. When a standard user attempts to run an application that requires an administrator access token, UAC requires that the user provide valid administrator credentials.

In order to better understand how this process happens it is important to review the details of the Windows Server 2012 logon process.

Windows Server® 2012 Logon Process

The following illustration demonstrates how the logon process for an administrator differs from the logon process for a standard user.

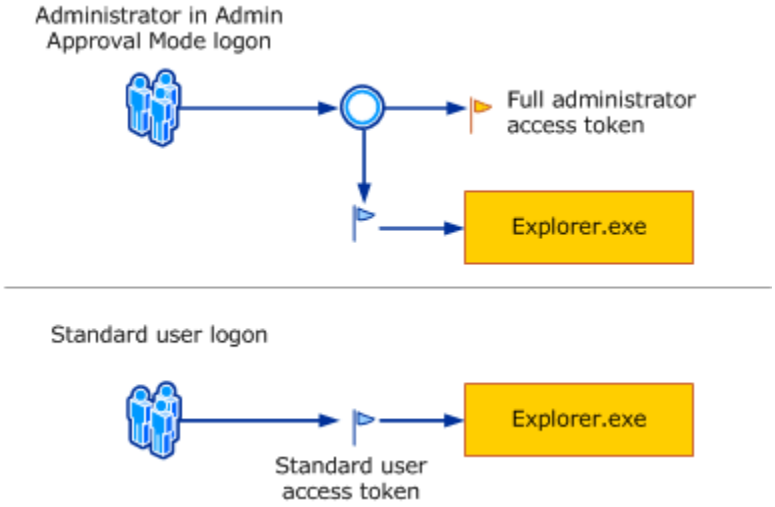

By default, standard users and administrators access resources and run applications in the security context of standard users. When a user logs on to a computer, the system creates an access token for that user. The access token contains information about the level of access that the user is granted, including specific security identifiers (SIDs) and Windows privileges.

When an administrator logs on, two separate access tokens are created for the user: a standard user access token and an administrator access token. The standard user access token contains

the same user-specific information as the administrator access token, but the administrative Windows privileges and SIDs are removed. The standard user access token is used to start applications that do not perform administrative tasks (standard user applications). The standard user access token is then used to display the desktop (Explorer.exe). Explorer.exe is the parent process from which all other user-initiated processes inherit their access token. As a result, all applications run as a standard user unless a user provides consent or credentials to approve an application to use a full administrative access token.

A user that is a member of the Administrators group can log on, browse the Web, and read e-mail while using a standard user access token. When the administrator needs to perform a task that requires the administrator access token, Windows Server 2012 automatically prompts the user for approval. This prompt is called an elevation prompt, and its behavior can be configured by using the Local Security Policy snap-in (Secpol.msc) or Group Policy.

Note

The term "elevate" is used to refer to the process in Windows Server 2012 that prompts the user for consent or credentials to use a full administrator access token.

The UAC User Experience

When UAC is enabled, the user experience for standard users is different from that of administrators in Admin Approval Mode. The recommended and more secure method of running Windows Server 2012 is to make your primary user account a standard user account. Running as a standard user helps to maximize security for a managed environment. With the built-in UAC elevation component, standard users can easily perform an administrative task by entering valid credentials for a local administrator account. The default, built-in UAC elevation component for standard users is the credential prompt.

The alternative to running as a standard user is to run as an administrator in Admin Approval Mode. With the built-in UAC elevation component, members of the local Administrators group can easily perform an administrative task by providing approval. The default, built-in UAC elevation component for an administrator account in Admin Approval Mode is called the consent prompt. The UAC elevation prompting behavior can be configured by using the Local Security Policy snap-in (Secpol.msc) or Group Policy.

The consent and credential prompts

With UAC enabled, Windows Server 2012 prompts for consent or prompts for credentials of a valid local administrator account before starting a program or task that requires a full administrator access token. This prompt ensures that no malicious software can be silently installed.

The consent prompt

The consent prompt is presented when a user attempts to perform a task that requires a user's administrative access token. The following is a screen shot of the UAC consent prompt.

The credential prompt

The credential prompt is presented when a standard user attempts to perform a task that requires a user's administrative access token. This standard user default prompt behavior can be configured by using the Local Security Policy snap-in (Secpol.msc) or Group Policy. Administrators can also be required to provide their credentials by setting the User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode policy setting value to Prompt for credentials.

The following screen shot is an example of the UAC credential prompt.

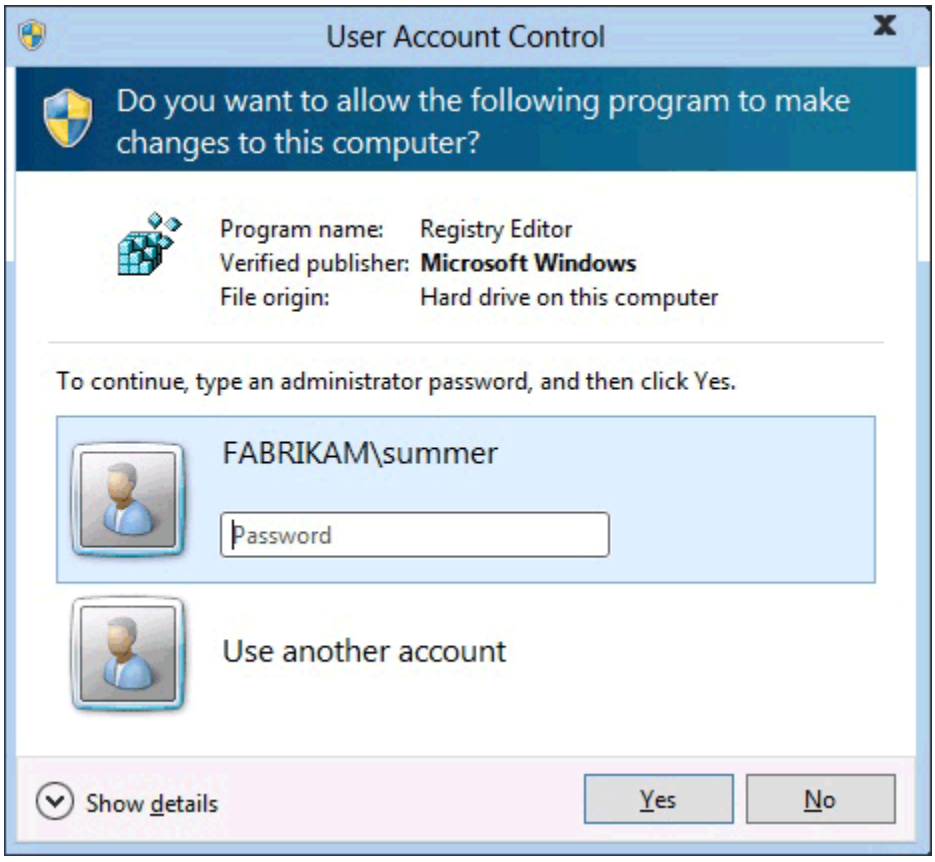

UAC elevation prompts

The UAC elevation prompts are color-coded to be application-specific, enabling for immediate identification of an application's potential security risk. When an application attempts to run with an administrator's full access token, Windows Server 2012 first analyzes the executable file to determine its publisher. Applications are first separated into three categories based on the executable file's publisher: Windows Server 2012, publisher verified (signed), and publisher not verified (unsigned). The following diagram illustrates how Windows Server 2012 determines which color elevation prompt to present to the user.

The elevation prompt color-coding is as follows:

- Red background with a red shield icon: The application is blocked by Group Policy or is from a publisher that is blocked.
- Blue background with a blue and gold shield icon: The application is a Windows Server 2012 administrative application, such as a Control Panel item.
- Blue background with a blue shield icon: The application is signed by using Authenticode and is trusted by the local computer.
- Yellow background with a yellow shield icon: The application is unsigned or signed but is not yet trusted by the local computer.

Shield icon

Some Control Panel items, such as **Date and Time Properties**, contain a combination of administrator and standard user operations. Standard users can view the clock and change the time zone, but a full administrator access token is required to change the local system time. The following is a screen shot of the **Date and Time Properties** Control Panel item.

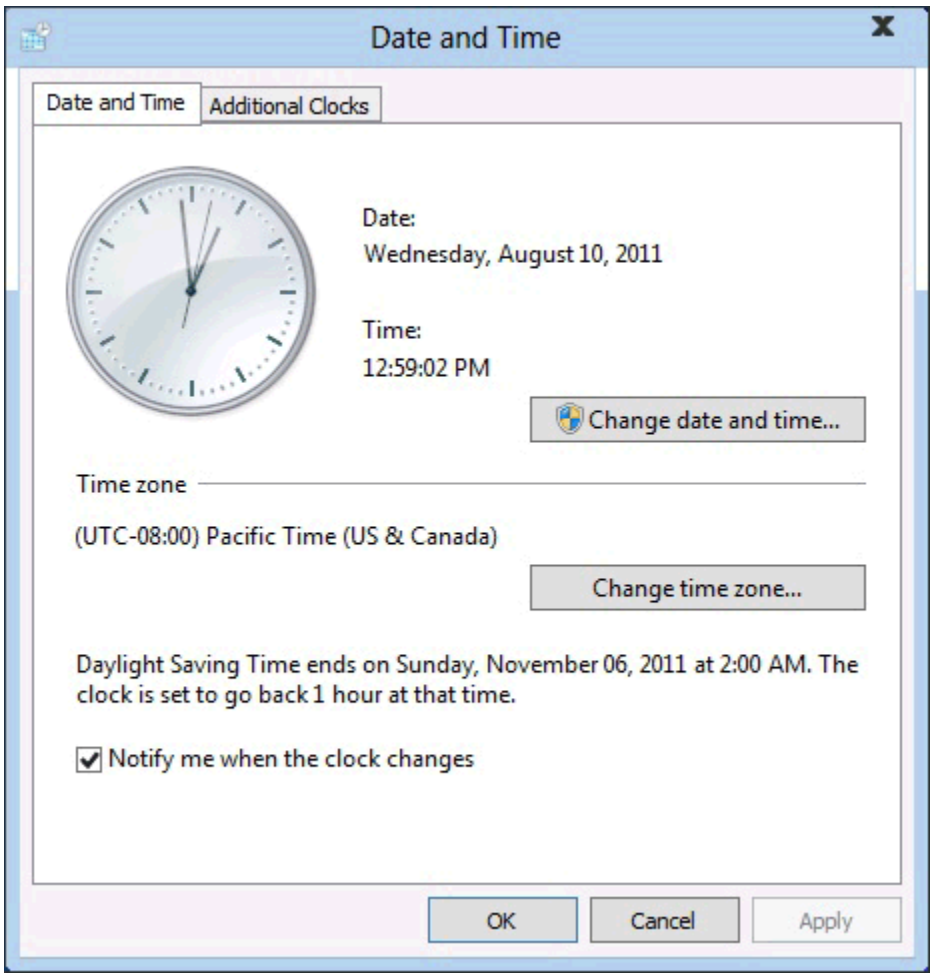

The shield icon on the **Change date and time** button indicates that the process requires a full administrator access token and will display a UAC elevation prompt.

Securing the elevation prompt

The elevation process is further secured by directing the prompt to the secure desktop. The consent and credential prompts are displayed on the secure desktop by default in Windows Server 2012. Only Windows processes can access the secure desktop. For higher levels of security, we recommend keeping the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting enabled.

When an executable file requests elevation, the interactive desktop, also called the user desktop, is switched to the secure desktop. The secure desktop dims the user desktop and displays an

elevation prompt that must be responded to before continuing. When the user clicks Yes or No, the desktop switches back to the user desktop.

Malware can present an imitation of the secure desktop, but when the User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode policy setting is set to Prompt for consent, the malware does not gain elevation if the user clicks Yes on the imitation. If the policy setting is set to Prompt for credentials, malware imitating the credential prompt may be able to gather the credentials from the user. However, the malware does not gain elevated privilege and the system has other protections that mitigate malware from taking control of the user interface even with a harvested password.

While malware could present an imitation of the secure desktop, this issue cannot occur unless a user previously installed the malware on the computer. Because processes requiring an administrator access token cannot silently install when UAC is enabled, the user must explicitly provide consent by clicking **Yes** or by providing administrator credentials. The specific behavior of the UAC elevation prompt is dependent upon Group Policy.

UAC Architecture

The following diagram details the UAC architecture.

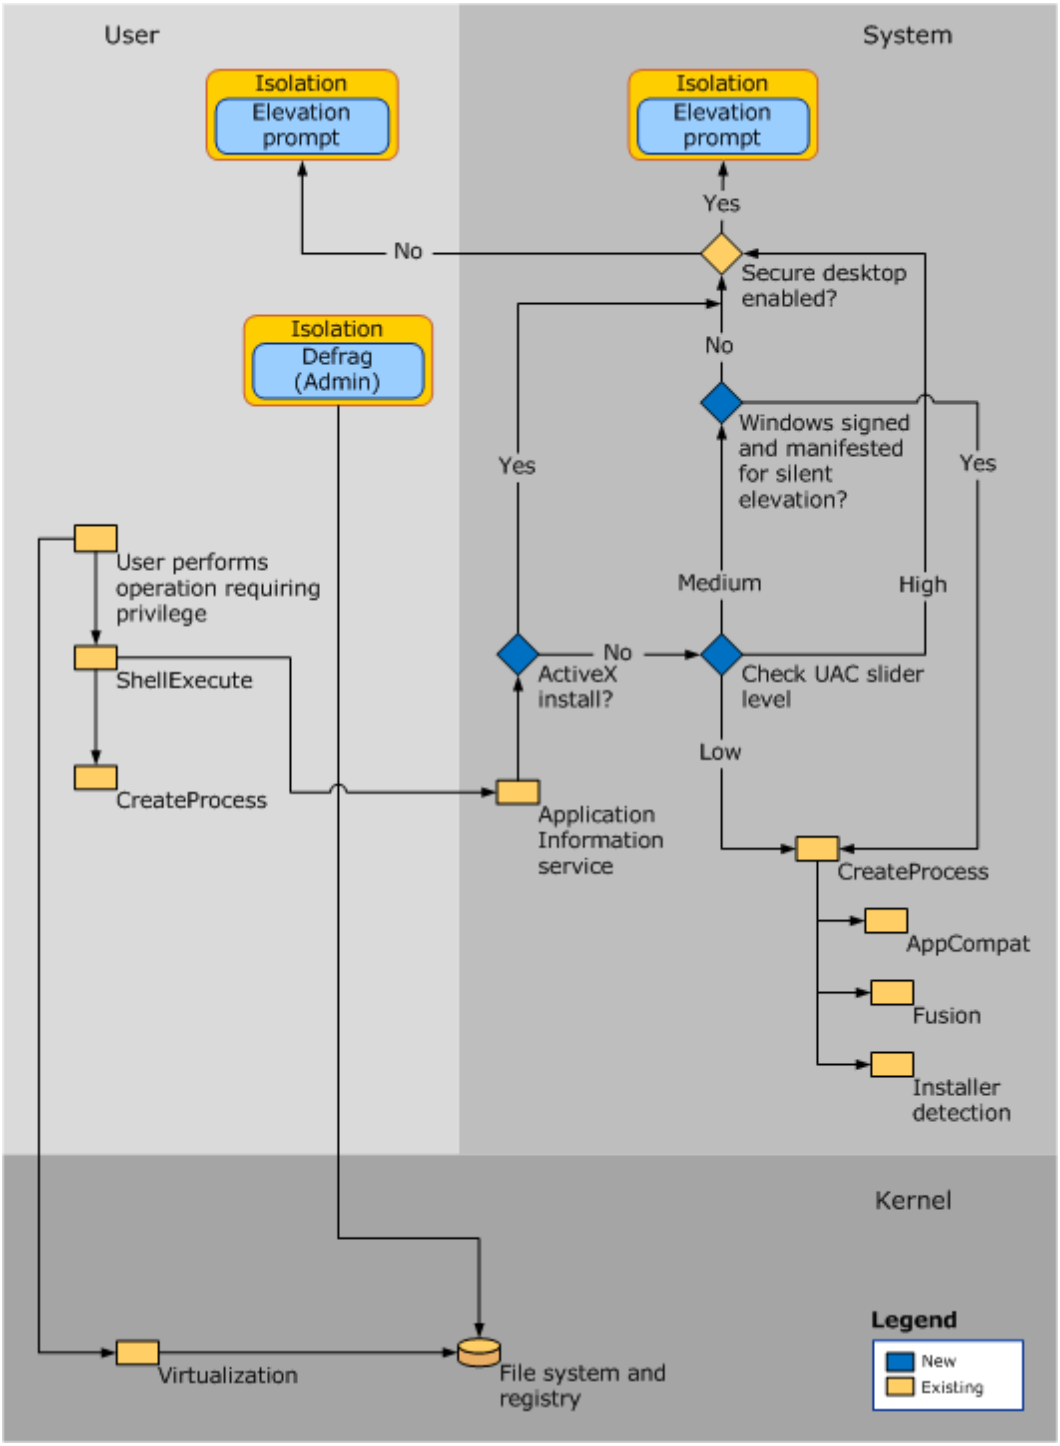

To better understand each component, review the table below:

Component	Description
User	
User performs operation requiring privilege	If the operation changes the file system or registry, Virtualization is called. All other operations call ShellExecute.
ShellExecute	ShellExecute calls CreateProcess. ShellExecute looks for the ERROR_ELEVATION_REQUIRED error from CreateProcess. If it receives the error, ShellExecute calls the Application Information service to attempt to perform the requested task with the elevated prompt.
CreateProcess	If the application requires elevation, CreateProcess rejects the call with ERROR_ELEVATION_REQUIRED.
System	
Application Information service	A system service that helps start applications that require one or more elevated privileges or user rights to run, such as local administrative tasks, and applications that require higher integrity levels. The Application Information service helps start such applications by creating a new process for the application with an administrative user's full access token when elevation is required and (depending on Group Policy) consent is given by the user to do so.
Elevating an ActiveX install	If ActiveX is not installed, the system checks the UAC slider level. If ActiveX is installed, the **User Account Control: Switch to the secure desktop when prompting for elevation** Group Policy setting is checked.
Check UAC slider level	UAC now has four levels of notification to choose from and a slider to use to select the notification level: • High If the slider is set to **Always notify**, the system checks whether the secure desktop is enabled.

Component	Description

	• Medium
	If the slider is set to **Default-Notify me only when programs try to make changes to my computer**, the **User Account Control: Only elevate executable files that are signed and validated** policy setting is checked:
	  • If the policy setting is enabled, the public key infrastructure (PKI) certification path validation is enforced for a given executable file before it is permitted to run.
	  • If the policy setting is not enabled (default), the PKI certification path validation is not enforced before a given executable file is permitted to run. The **User Account Control: Switch to the secure desktop when prompting for elevation** Group Policy setting is checked.
	• Low
	If the slider is set to **Notify me only when programs try to make changes to my computer (do not dim by desktop)**, the CreateProcess is called.
	• Never Notify
	If the slider is set to **Never notify me when**, UAC prompt will never notify when a program is trying to install or trying to make any change on the computer.
	Important
	This setting is not recommended. This setting is the same as setting the **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting to **Elevate without prompting**.

Secure desktop enabled	The **User Account Control: Switch to the secure desktop when prompting for**

Component	Description
	elevation policy setting is checked: • If the secure desktop is enabled, all elevation requests go to the secure desktop regardless of prompt behavior policy settings for administrators and standard users. • If the secure desktop is not enabled, all elevation requests go to the interactive user's desktop, and the per-user settings for administrators and standard users are used.
CreateProcess	CreateProcess calls AppCompat, Fusion, and Installer detection to assess if the application requires elevation. The executable file is then inspected to determine its requested execution level, which is stored in the application manifest for the executable file. CreateProcess fails if the requested execution level specified in the manifest does not match the access token and returns an error (ERROR_ELEVATION_REQUIRED) to ShellExecute. For more information, see **Requested execution levels** later in this section.
AppCompat	The AppCompat database stores information in the application compatibility fix entries for an application.
Fusion	The Fusion database stores information from application manifests that describe the applications. The manifest schema is updated to add a new requested execution level field.
Installer detection	Installer detection detects setup executable files, which helps prevent installations from being run without the user's knowledge and consent. For more information, see **Installer detection technology** later in this section.
Kernel	
Virtualization	Virtualization technology ensures that non-compliant applications do not silently fail to run

Component	Description
	or fail in a way that the cause cannot be determined. UAC also provides file and registry virtualization and logging for applications that write to protected areas. For more information, see **Virtualization** later in this section.
File system and registry	The per-user file and registry virtualization redirects per-computer registry and file write requests to equivalent per-user locations. Read requests are redirected to the virtualized per-user location first and to the per-computer location second.

There is a change on Windows Server 2012 UAC from previous Windows versions. The new slider will never turn UAC completely off. The new setting will:

- Keep the UAC service running.
- Cause all elevation request initiated by administrators to be auto-approved without showing a UAC prompt.
- Automatically deny all elevation requests for standard users.

Important
In order to fully disable UAC you must disable the policy **User Account Control: Run all administrators in Admin Approval Mode**.

Warning
Tailored Applications will not work on Windows Server 2012 when UAC is disabled.

Virtualization

Because system administrators in enterprise environments attempt to secure systems, many line-of-business (LOB) applications are designed to use only a standard user access token. As a result, IT administrators do not need to replace the majority of applications when running Windows Server 2012 with UAC enabled.

Windows Server 2012 includes file and registry virtualization technology for applications that are not UAC compliant and that require an administrator's access token to run correctly. Virtualization ensures that even applications that are not UAC compliant are compatible with Windows Server 2012 . When an administrative application that is not UAC compliant attempts to write to a protected directory, such as Program Files, UAC gives the application its own virtualized view of the resource it is attempting to change. The virtualized copy is maintained in the user's profile. This strategy creates a separate copy of the virtualized file for each user that runs the non-compliant application.

Most application tasks operate properly by using virtualization features. Although virtualization allows a majority of applications to run, it is a short-term fix and not a long-term solution. Application developers should modify their applications to be compliant with the Windows Server 2012 logo program as soon as possible, rather than relying on file, folder, and registry virtualization.

Virtualization is not in option in the following scenarios:

1. Virtualization does not apply to applications that are elevated and run with a full administrative access token.
2. Virtualization supports only 32-bit applications. Non-elevated 64-bit applications simply receive an access denied message when they attempt to acquire a handle (a unique identifier) to a Windows object. Native Windows 64-bit applications are required to be compatible with UAC and to write data into the correct locations.
3. Virtualization is disabled for an application if the application includes an application manifest with a requested execution level attribute.

Request Execution Levels

An application manifest is an XML file that describes and identifies the shared and private side-by-side assemblies that an application should bind to at run time. In Windows Server 2012, the application manifest includes entries for UAC application compatibility purposes. Administrative applications that include an entry in the application manifest prompt the user for permission to access the user's access token. Although they lack an entry in the application manifest, most administrative applications can run without modification by using application compatibility fixes. Application compatibility fixes are database entries that enable applications that are not UAC compliant to work properly with Windows Server 2012.

All UAC-compliant applications should have a requested execution level added to the application manifest. If the application requires administrative access to the system, then marking the application with a requested execution level of "require administrator" ensures that the system identifies this program as an administrative application and performs the necessary elevation steps. Requested execution levels specify the privileges required for an application.

Installer Detection Technology

Installation programs are applications designed to deploy software. Most installation programs write to system directories and registry keys. These protected system locations are typically writeable only by an administrator in Installer detection technology, which means that standard users do not have sufficient access to install programs. Windows Server 2012 heuristically detects installation programs and requests administrator credentials or approval from the administrator user in order to run with access privileges. Windows Server 2012 also heuristically detects updates and programs that uninstall applications. One of the design goals of UAC is to prevent installations from being run without the user's knowledge and consent because installation programs write to protected areas of the file system and registry.

Installer detection only applies to:

- 32-bit executable files.
- Applications without a requested execution level attribute.
- Interactive processes running as a standard user with UAC enabled.

Before a 32-bit process is created, the following attributes are checked to determine whether it is an installer:

- The file name includes keywords such as "install," "setup," or "update."
- Versioning Resource fields contain the following keywords: Vendor, Company Name, Product Name, File Description, Original Filename, Internal Name, and Export Name.
- Keywords in the side-by-side manifest are embedded in the executable file.
- Keywords in specific StringTable entries are linked in the executable file.
- Key attributes in the resource script data are linked in the executable file.
- There are targeted sequences of bytes within the executable file.

Note
The keywords and sequences of bytes were derived from common characteristics observed from various installer technologies.

Note
The User Account Control: Detect application installations and prompt for elevation policy setting must be enabled for installer detection to detect installation programs. This setting is enabled by default and can be configured locally by using the Local Security Policy snap-in (Secpol.msc) or configured for the domain, OU, or specific groups by Group Policy (Gpedit.msc).

Windows Authentication Overview

This navigation topic for the IT professional lists documentation resources for Windows authentication and logon technologies that include product evaluation, getting started guides, procedures, design and deployment guides, technical references, and command references for Windows Server 2012.

Feature description

Authentication is a process for verifying the identity of an object, service or person. When you authenticate an object, the goal is to verify that the object is genuine. When you authenticate a service or person, the goal is to verify that the credentials presented are authentic.

In a networking context, authentication is the act of proving identity to a network application or resource. Typically, identity is proven by a cryptographic operation that uses either a key only the user knows — as with public key cryptography — or a shared key. The server side of the authentication exchange compares the signed data with a known cryptographic key to validate the authentication attempt.

Storing the cryptographic keys in a secure central location makes the authentication process scalable and maintainable. Active Directory Domain Services is the recommended and default technology for storing identity information (including the cryptographic keys that are the user's' credentials). Active Directory is required for default NTLM and Kerberos implementations.

Authentication techniques range from a simple logon, which identifies users based on something that only the user knows — like a password, to more powerful security mechanisms that use something that the user has — like tokens, public key certificates, and biometrics. In a business environment, services or users might access multiple applications or resources on many types of servers within a single location or across multiple locations. For these reasons, authentication must support environments for other platforms and for other Windows operating systems.

The Windows operating system implements a default set of authentication protocols, including Kerberos, NTLM, Transport Layer Security/Secure Sockets Layer (TLS/SSL), and Digest, as part of an extensible architecture. In addition, some protocols are combined into authentication packages such as Negotiate and the Credential Security Support Provider. These protocols and packages enable authentication of users, computers, and services; the authentication process, in turn, enables authorized users and services to access resources in a secure manner.

Practical applications

Windows Authentication is used to verify that the information comes from a trusted source, whether from a person or computer object, such as another computer. Windows provides many different methods to achieve this goal as described below.

To…	Feature	Description
Authenticate within an Active Directory domain	Kerberos	The Microsoft Windows Server operating systems implement the Kerberos version 5 authentication protocol and extensions for public key authentication. The Kerberos authentication client is implemented as a security support provider (SSP) and can be accessed through the Security Support Provider Interface (SSPI). Initial user authentication is integrated with the Winlogon single sign-on architecture. The Kerberos Key Distribution Center (KDC) is integrated with other Windows Server security services running

To…	Feature	Description
		on the domain controller. The KDC uses the domain's Active Directory directory service database as its security account database. Active Directory is required for default Kerberos implementations. For additional resources, see Kerberos Authentication Overview.
Secure authentication on the web	TLS/SSL as implemented in the Schannel Security Support Provider	The Transport Layer Security (TLS) protocol versions 1.0, 1.1, and 1.2, Secure Sockets Layer (SSL) protocol, versions 2.0 and 3.0, and the Private Communications Transport (PCT) protocol, version 1.0, are based on public key cryptography. The Secure Channel (Schannel) provider authentication protocol suite provides these protocols. All Schannel protocols use a client and server model. For additional resources, see TLS/SSL (Schannel SSP) Overview.
Authenticate to a web service or application	Integrated Windows Authentication Digest Authentication	For additional resources, see Integrated Windows Authentication and Digest Authentication, and Advanced Digest Authentication.
Authenticate to legacy applications	NTLM	NTLM is a challenge-response style authentication protocol.In addition to authentication, the NTLM protocol optionally provides for session security—specifically message integrity and confidentiality through

To…	Feature	Description
		signing and sealing functions in NTLM. For additional resources, see NTLM Overview.
Leverage multifactor authentication	Smart card support Biometric support	Smart cards are a tamper-resistant and portable way to provide security solutions for tasks such as client authentication, logging on to domains, code signing, and securing e-mail. Biometrics relies on measuring an unchanging physical characteristic of a person to uniquely identify that person. Fingerprints are one of the most frequently used biometric characteristics, with millions of fingerprint biometric devices that are embedded in personal computers and peripherals. For additional resources, see Smart Card Overview and **Windows Biometric Framework Overview**.
Provide local management, storage and reuse of credentials	Credentials management Local Security Authority Passwords	Credential management in Windows ensures that credentials are stored securely. Credentials are collected on the Secure Desktop (for local or domain access), through apps or through websites so that the correct credentials are presented every time a resource is accessed. For additional resources, see **Credential Locker Overview** and Passwords Overview.
Extend modern authentication	Extended Protection for	This feature enhances the

To…	Feature	Description
protection to legacy systems	Authentication	protection and handling of credentials when authenticating network connections by using Integrated Windows Authentication (IWA). For additional resources, see Extended Protection for Authentication.

New and changed functionality

Feature	Change summary	Resources
Kerberos	Intraforest cross-domain Constrained Delegation Troubleshooting improvements (log and tracing, integration with DirectAccess Support for compound ID using FAST Support for Claims authorization data	What's New in Kerberos Authentication
TLS/SSL as implemented in the Schannel Security Support Provider	TLS support for Server Name Indicator (SNI) Extensions Datagram Transport Layer Security (DTLS)	What's New in TLS/SSL (Schannel SSP)
Integrated Windows Authentication and NTLM	No changes in functionality	Included by default in Windows Server 2012 and Windows 8.
Smart card support	Virtual smart cards closely mimic the functionality of physical smart cards. The virtual smart card is essentially a smart card that is always available on the computer.	What's New in Smart Cards
Biometric support	Improvements to fast user switching for biometric devices	**New and changed functionality**

Feature	Change summary	Resources
	and credentials provider support	
Credentials management	Architecture changes to the former Windows Vault which now is called Credential Locker	**New and changed functionality**
Managed Service Accounts	Administering standalone Managed Service Accounts is now easier with the addition of group Managed Service Accounts that extends the functionality to groups of servers.	Group Managed Service Accounts Overview

Removed or deprecated functionality

For a list of deprecated features in Windows Server 2012, see Features Removed or Deprecated in Windows Server 2012.

Software requirements

Windows Authentication is designed to be compatible with previous versions of the Windows operating system. However, improvements with each release are not necessarily applicable to previous versions. Refer to documentation about specific features for more information.

Server Manager information

Many authentication features can be configured using Group Policy, which can be installed using Server Manager. The Windows Biometric Framework feature is installed using Server Manager. Other server roles which are dependent upon authentication methods, such as Web Server (IIS) and Active Directory Domain Services, can also be installed using Server Manager.

Windows Server 2012 R2 Software Inventory Logging Overview

This technical overview describes Software Inventory Logging, a feature in Windows Server® 2012 R2 that aggregates software identification on a per server basis.

Did you mean…

- User Access Logging Overview

- Microsoft Assessment and Planning Toolkit
- Volume Activation Management Tool

Feature description

Software Inventory Logging in Windows Server 2012 R2 is a feature with a simple set of interfaces that helps server administrators retrieve a list of the software that is installed on their servers. Updates and fixes from software vendors are also included in the list, as well as components of the server system commonly used for inventory. In the cases of some software, unique client request counts are summed and also made available with these interfaces (for more information, see User Access Logging Overview). Windows Server 2012 R2 Hyper-V host servers can retrieve this data from virtual machines, running on the local hypervisor, if the virtual machines are also running Windows Server 2012 R2. Software Inventory Logging can be configured by the server administrator, by enabling a daily task via PowerShell, to collect all data from these interfaces and forward it using HTTPS to a server running a web service (configured by the server owner), for aggregation.

Retrieving data from the Software Inventory Logging interfaces, or manually forwarding this data to an aggregation server, is done through PowerShell cmdlets. Managing the feature, primarily for daily collection and forwarding, is also done with PowerShell commands.

Note

An aggregation server running a web service can be separately configured, but is not deployed as part of Software Inventory Logging.

Important

None of the data collected by Software Inventory Logging is sent to Microsoft as part of the feature's functionality.

Practical applications

Software Inventory Logging interfaces are intended to reduce the operational costs of retrieving accurate information regarding software deployed locally on a server, but especially across many servers in an IT deployment (assuming the feature exists on these servers). The ability to forward this data daily to an aggregation server (if set up separately by an IT administrator), allows the data to be collected in one place by a uniform, automatic, process. While this can also be achieved by querying the interfaces directly, Software Inventory Logging provides this added functionality, and by employing a forwarding (over the network) architecture initiated at the host server, can overcome machine discovery challenges typical for many software inventory and asset management scenarios. SSL is used to encrypt data forwarded over HTTPS to an administrator's aggregation server. Having the data in one place (on one server) makes the data easier to analyze, manipulate and report on. It is important to note that none of this data is sent to

Microsoft as part of the feature functionality. Software Inventory Logging data and functionality is meant for the sole use of the server software's licensed owner and administrators.

Software Inventory Logging can assist server administrators in performing the following tasks:

- Retrieving software and server inventory information from Windows Server 2012 R2 servers, remotely and on-demand.
- Aggregating software and server inventory information collected daily from Windows Server 2012 R2 by enabling each server's Software Inventory Logging task and choosing a web server target URI.
- Retrieving, collecting, and aggregating, counts of unique client requests reflected by the User Access Logging feature. Since Software Inventory Logging relies on the User Access Logging service to collect unique client request counts, it supports the same server roles and applications for this functionality as User Access Logging. See the User Access Logging Overview for the complete listing.

Software requirements

Software Inventory Logging can be used on any computer running Windows Server 2012 R2.

See Also

User Access Logging Overview

Microsoft Assessment and Planning Toolkit

Volume Activation Overview

Volume Activation Management Tool

Manage Software Inventory Logging in Windows Server 2012 R2

This document describes how to manage Software Inventory Logging in Windows Server® 2012 R2.

The Software Inventory Logging feature in Windows Server 2012 R2 is intended to help datacenter administrators easily log inventory data from their deployments. This feature is built into the operating system but only exists on Windows Server 2012 R2 (it does not collect or log data from other operating systems). There are two primary methods for inventorying using this feature:

1. Enabling the SilLogging daily task to collect from all data sources and forward over the network to a specified target (URI)
2. Manually querying the data using Powershell or WMI at any interval.

Enabling the SilLogging daily task involves some planning and foresight, but has significant advantages compare to manually querying the data.

Enabling SilLogging has the following three primary advantages for data center administrators:

- An ongoing history, or log, can be aggregated over time empowering flexible and comprehensive reports from a single source.
- If Software Inventory Logging is enabled within base virtual machine images, or some other widespread deployment method or policy, computer discovery challenges typical with many inventory tools can be overcome.
- If Software Inventory Logging is enabled within base virtual machine images, or some other widespread deployment method or policy, trust boundary challenges and the need for elevated user privileges, which is typical with many inventory tools, can be overcome, yet all the while maintaining a level of security since the data is encrypted over HTTPS with SSL.

Software Inventory Logging is installed and enabled manually (not by default) in Windows Server 2012 R2 and collects data on a nearly real-time basis. All configuration of Software Inventory Logging is done with PowerShell cmdlets. There are only a few configuration options for Software Inventory Logging. This document describes these options and their intended purpose.

To learn more about the benefits of Software Inventory Logging, see: Windows Server 2012 R2 Software Inventory Logging Overview.

In this document

The configuration options covered in this document include:

- Software Inventory Logging PowerShell cmdlet
- Enabling and Disabling Software Inventory Logging
- Displaying Software Inventory Logging data
- Forwarding Software Inventory Logging data
- Deleting data logged by Software Inventory Logging
- Backing up and restoring data logged by Software Inventory Logging
- Configuring certificates for Software Inventory Logging
- Enabling Software Inventory using unattended installation
- Creating a master VHD image to deploy Software Inventory Logging to servers in an enterprise

Note

This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Starting and Stopping Software Inventory Logging

Software Inventory Logging daily collection and forward over the network must be enabled on a computer running Windows Server 2012 R2 to log software inventory.

Note

You can use the **Get-SilLogging** PowerShell cmdlet to retrieve information about the Software Inventory Logging Service including whether it is running or stopped

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

To Start Software Inventory Logging

1. Sign in to the server with an account that has local administrator privileges.
2. Open PowerShell.
3. At the PowerShell prompt, type **Start-SilLogging**

The following Windows PowerShell cmdlets disable Software Inventory Logging.

To Stop Software Inventory Logging

1. Sign in to the server with an account that has local administrator privileges.
2. Open PowerShell.
3. At the PowerShell prompt, type **Stop-SilLogging**

Displaying Software Inventory Logging data

In addition to the PowerShell cmdlets described in the previous section, 5 additional cmdlets in Windows Server 2012 R2 can be used to collect Software Inventory Logging data:

- **Get-SilComputer**: Displays the point in time values for specific server and operating system-related data.
- **Get-SilData**: Displays a point in time collection of all Software Inventory Logging data.
- **Get-SilSoftware**: Displays the point in time identity of all software installed on the computer
- **Get-SilUalAccess**: Displays the total number of unique client device requests and client user requests of the server from two days prior.
- **Get-SilWindowsUpdate:**: Displays the point in time list of all Windows updates installed on the computer.

A typical use case scenario for Software Inventory Logging cmdlets would be for an administrator to query Software Inventory Logging for a point in time collection of all Software Inventory Logging data using **Get-SilData**

```
PS C:\> Get-SilData


Item   : __PARAMETERS

VmGuid : 15F8DF27-183C-46BE-A037-2D8797E75C1C

Item   : MsftSil_UalAccess (RoleGuid = "7fb09bd3-7fe6-435e-8348-7d8aefb6cea3")

VmGuid : 15F8DF27-183C-46BE-A037-2D8797E75C1C

Item   : MsftSil_Computer (Name = "Contoso-01.Server.Test.Contoso...)

VmGuid : 15F8DF27-183C-46BE-A037-2D8797E75C1C
```

```
PSComputerName :
ID          : KB999114
InstallDate : 5/15/2013
ChassisSerialNumber       : 7ZQ7TJ1
CollectedDateTime         : 5/20/2013 3:20:09 PM
Model                     : Precision WorkStation T7400
Name                      : Server01.Test.Contoso.com
NumberOfCores             : 4
NumberOfLogicalProcessors : 4
NumberOfProcessors        : 1
OSName                    : Microsoft Windows Server 2012 R2 Datacenter Preview
OSSku                     : 8
OSSuite                   : 400
OSSuiteMask               : 400
OSVersion                 : 6.3.9405
ProcessorFamily           : 179
ProcessorManufacturer     : GenuineIntel
ProcessorName             : Intel(R) Xeon(R) CPU           E5440  @ 2.83GHz
SystemManufacturer        : Dell Inc.
VmGuid                    :
ProductName             : Windows Server 2012 R2 Datacenter Preview
RoleGuid                : 10a9226f-50ee-49d8-a393-9a501d47ce04
RoleName                : File Server
SampleDate              : 5/18/2013
UniqueDeviceAccessCount : 32
UniqueUserAccessCount   : 55
ProductName             : Windows Server 2012 R2 Datacenter Preview
RoleGuid                : 7fb09bd3-7fe6-435e-8348-7d8aefb6cea3
RoleName                : Print and Document Services
SampleDate              : 5/18/2013
UniqueDeviceAccessCount : 2
UniqueUserAccessCount   : 6
ID            : {6B6533BD-1680-4368-908E-D50977561A86}
InstallDate    : 5/14/2013 12:00:00 AM
```

```
Name            : Microsoft Office Professional Plus 2010

Vendor          : Microsoft Corporation

Version         : 14.0.6029.1000
```

Note
Output from this cmdlet is the same as all other **Get-Sil** cmdlets for this feature combined, but is provided to the console asynchronously so the order of the objects may not always be the same.

Note
If this cmdlet is run on a Windows Server 2012 R2 Hyper-V host, Software Inventory Logging data will be provided from the guest virtual machines if they are running Windows Server 2012 R2 and have Software Inventory Logging started. This data will be 'collapsed' on the console's view, identifiable by the virtual machine's assigned GUID.

Forwarding Software Inventory Logging data

To forward data to an aggregation server, there are three PowerShell cmdlets that you need to run:

1. Specify the target aggregation server using the **Set-SilLogging** cmdlet.
2. Enable Software Inventory Logging as described above using the **Start-SilLogging** cmdlet.
3. Set the certificate thumbprint using the **Set-SilLogging** cmdlet

Before Software Inventory Logging data can be forwarded over the network to an aggregation server, you must first specify the target server. The **Set-SilLogging** cmdlets is used to specify an aggregation server. If you do not specify an aggregation server and the server is a virtual machine that runs on a Hyper-V host, then the daily collection forwards the Software Inventory Logging data to a location that its Windows Server host can access. **Set-SilLogging** sets the time of day that the daily collection occurs, the uniform resource identifier (URI) of the aggregation server, and the certificate thumbprint to use for encryption.

Once the aggregation server has been specified, you can enable a daily task via PowerShell to collect all data from these interfaces and forward it using HTTPS to a server running a web service (configured by the server owner), for aggregation.

Note
You can use the **Get-SilLogging** PowerShell cmdlet to retrieve information about the Software Inventory Logging Service including whether it is running or stopped. To see en example of **Get-SilLogging** data output, refer back to the **Displaying Software Inventory Logging data** section above.

The **Publish-SilData** cmdlet initiates a point in time collection of all Software Inventory Logging data, and forwards the data over the network to an aggregation server. If Software Inventory Logging was started by an administrator, daily collection and forwarding of the data to the aggregation server commences. In cases where Software Inventory Logging does not receive an acknowledgement from the aggregation server that data forwarding was successful, the service stores that data locally and attempts to forward both the previous and current data the next day. Administrators can also use this cmdlet to manually forward all historical Software Inventory Logging data that was accumulated locally, if that data was not been successfully forwarded to the aggregation server for a period of one or more days. After Software Inventory Logging successfully forwards the data, it deletes all historical data from the client server.

Important

If for any reason a repair installation or upgrade is necessary of the operating system, all configuration settings will be lost. In these situations, simply set up Software Inventory logging as if a fresh installation of the operating system has been performed.

Deleting data logged by Software Inventory Logging

Software Inventory Logging is not intended to be a mission critical component. Its design is intended to impact local system operations as little as possible while maintaining a high level of reliability. This also allows the administrator to manually delete Software Inventory Logging database and supporting files (every file in \Windows\System32\LogFiles\SIL\ directory) to meet operational needs.

To delete data logged by Software Inventory Logging

1. In PowerShell, stop Software Inventory Logging with the **Stop-SilLogging** command.
2. Open Windows Explorer.
3. Go to **\Windows\System32\Logfiles\SIL**.
4. Delete all files in the folder.

Backing up and restoring data logged by Software Inventory Logging

Software Inventory Logging will temporarily store daily collections of data if it has been configured to forward data to an aggregator, BUT forwards have failed due to external issues, i.e. network connectivity. The log files are stored in the \Windows\System32\LogFiles\SIL\ directory. Backups of this Software Inventory Logging data can be made with your regularly scheduled server backups.

Important

If for any reason a repair installation or upgrade is necessary of the operating system, any log files stored locally will be lost. If this data is critical for operations, it is recommended to be backed up prior to operating system installation. After repair or upgrade, simply restore to the same location.

Configuring certificates for Software Inventory Logging

Certificate based authentication uses certificates to authenticate the user/client (in this case, to access SIL). The certificate is used in place of credentials.

By deploying certificate based authentication, administrators can control which computers can be queried for Software Inventory Logging Data.

When you use certificate based authentication, there will be added administrative overhead. You will either need to deploy your own internal Public Key Infrastructure (PKI) using Active Directory Certificate Services (AD CS, formerly Windows Server Certificate Services) or a 3rd-party PKI solution, or you will have to purchase certificates from a public certification authority (CA). This will not be a one-time added overhead. Certificates expire, and when a user's certificate expires, they will need a new one, requiring either time invested in getting the user a new certificate, or budget invested in purchasing one.

Prerequisites

You need access to a CA for client certificates. This can be a public CA solution, individual certificates from a vendor, or an Active Directory Certificate Services solution. Regardless of the method used, the following requirements must be met:

1. The certificate must be issued for client authentication. The default User template from an AD CS server will work in this scenario.
2. The **User Principal Name** (UPN) for each client's account must match the **Subject Name** field in the client's certificate.
3. All servers must trust the entire CA trust chain. This chain includes the root CA certificate and any intermediate CA certificates. These certificates should be installed on all servers where SIL will run.
4. The root CA certificate must be in the **Trusted Root Certification Authorities** store, and any intermediate CA certificates in the intermediate store on all of these systems. The root CA certificate and intermediate CA certificates must also be installed on the server where SIL will run.
5. The client's certificate must be associated with the client's account in Active Directory.

Enabling Software Inventory using unattended installation

You can configure an unattended Windows Server 2012 R2 installation with the following script.to enable Software Inventory Logging. This script schedules data collection 10 seconds after first logon, and writes the collected data to C:\Windows\System32\LogFiles\SIL folder.

```
<unattend xmlns="urn:schemas-microsoft-com:unattend"
xmlns:wcm="http://schemas.microsoft.com/WMIConfig/2002/State">
    <settings pass="oobeSystem">
        <component name="Microsoft-Windows-Shell-Setup" processorArchitecture="amd64"
publicKeyToken="31bf3856ad364e35" language="neutral" versionScope="nonSxS">
            ...
            <FirstLogonCommands>
                        <SynchronousCommand wcm:action="add">
                                <CommandLine>%windir%\System32\cmd.exe /c powershell -c
Start-SilLogging;Set-SilLogging -TimeOfDay (Get-Date).AddSeconds(10) -TargetURI
"http://testdcil/file.upload"</CommandLine>
                                <Description>Enable Software Inventory
Logging.</Description>
                                <Order>1</Order>
                        </SynchronousCommand>
                ...
            </FirstLogonCommands>
        </component>
    </settings>
    ...
</unattend>
```

Creating a master VHD image to deploy Software Inventory Logging to servers in an enterprise

Once your unattended installation completes, you can sysprep it and use the resulting image to stage other VMs.

With SIL enabled and configured, remove the unattend file from this computer. If you want to change the default data collection time to 30 minutes after login, you can run the following cmndlet in PowerShell:

```
Set-SilLogging -TimeOfDay (Get-Date).AddMinutes(30) -TargetURI
"http://testdcil/file.upload"
```

Then, to start the sysprep, run the following command at a command prompt

```
C:\Windows\System32\Sysprep\sysprep.exe /generalize /oobe /shutdown
```

The resulting image will:

1. Preserve the registry setting for Software Inventory Logging.
2. Preserve the Task Scheduler Settings and enable the daily collection task.
3. Write a new BMIL file to C:\Windows\System32\LogFiles\SIL at a pre-defined time.

See also

- Windows Server 2012 R2 Software Inventory Logging Overview

Telemetry Overview

You can help Microsoft improve the quality, reliability, and performance of its operating systems by participating in the Microsoft Customer Experience Improvement Program (CEIP) and Windows Error Reporting (WER). If you accept, Microsoft Corporation collects statistical information about your system configuration, the performance of some features of Windows, and certain types of events. Windows periodically uploads a small file to Microsoft that contains a summary of the information collected.

Did you mean…

- Windows 8 and Windows Server 2012 Privacy Highlights
- Customer Experience Improvement Program for Windows Server 2008 R2

Feature description

Windows Feedback Forwarder is a service that enables you to automatically send feedback to Microsoft by deploying a Group Policy setting to one or more organizational units. Windows Feedback Forwarder is available on all editions of Windows Server 2012. You can also install Windows Feedback Forwarder on a Server Core installation of Windows Server 2012 by using the Windows PowerShell cmdlets for Server Manager.

See also

Content type	References
What's New	What's New in Telemetry
Deployment	Deploy Windows Feedback Forwarder
Operations	*Content is not available yet.*

Deploy Windows Feedback Forwarder

Using the Windows Feedback Forwarder feature available in Windows Server 2012, up to 15,000 computers in one or more domain or organizational units can be automatically opted in to the Microsoft Customer Experience Improvement Program (CEIP) through Group Policy settings. By participating in CEIP, you can help Microsoft improve the quality, reliability, and performance of its operating systems. Microsoft collects statistical information about your system configuration, the performance of some Windows features, and certain types of events. Once you opt in to CEIP, Windows Server 2012 periodically uploads a small file to Microsoft that contains a summary of the information collected. The file is transferred securely using the Internet.

The uploaded data contains no information that identifies you or your company. Microsoft does not contact you, and it all happens automatically—you are never interrupted and you should experience no loss in performance. If your computer is not connected to the Internet, the data is eventually discarded after several days. Microsoft shares this information only with selected partners as outlined in the Windows Server 2012 privacy statement; it is used only by Microsoft in aggregate form for the purpose of improving our software for our customers.

In this guide

- Participate in the Customer Experience Improvement Program
- Participate in Windows Error Reporting
- Configure all servers to participate in CEIP and WER
- Install Windows Feedback Forwarder
- Configure Windows Feedback Forwarder

Participate in the Customer Experience Improvement Program

You can help Microsoft improve the quality, reliability, and performance of its operating systems by participating in the Microsoft Customer Experience Improvement Program. You can participate in the CEIP by creating a Group Policy Object (GPO), or using Server Manager.

To create a group policy object to deploy the CEIP

1. From the domain controller, open the Group Policy Management Console. (Press the

Windows logo key + R to open the **RUN** dialog box. Type **gpmc.msc** in the text box, and then click **OK**).

2. In the GPMC console tree, right-click **Group Policy Objects** in the forest and domain in which you want to create a GPO, and then click **New**.
3. In the **New GPO** dialog box, specify a name for the new GPO, and click **OK**.
4. Right-click the new GPO, and click **Edit** to open the Group Policy Management Editor.
5. In Computer Configuration\Preferences\Windows Settings, right-click **Registry**, click **New**, and then click **Registry Item**.
6. In the key path, enter 'HKLM\Software\Policies\Microsoft\SQMClient\Windows', and configure the following:
   - Name: CEIPEnable
   - Type: REG_DWORD
   - Value: 1

To participate in the Customer Experience Improvement Program by using Server Manager

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. In Server Manager, click **Local Server**, and then click **Participating** next to **Customer Experience Improvement Program**. The **Customer Experience Improvement Program Configuration** dialog box opens.
3. Select **Yes, I want to participate in the CEIP** and click **OK**.

If you decide not to continue participating in CEIP, you can be removed from CEIP using one of the following procedures:

To disable Customer Experience Improvement Program by using Server Manager

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. In Server Manager, click **Local Server**, and then click **Participating** next to **Customer Experience Improvement Program**. The **Customer Experience Improvement Program Configuration** dialog box opens.
3. Select **No, I do not want to participate** and click **OK**.

Participate in Windows Error Reporting

Windows Error Reporting (WER) is a set of Windows technologies that capture software crash and hang data from end users. Through the Winqual website, software and hardware vendors can access these reports in order to analyze, fix and respond to these problems. WER technologies were originally implemented in Windows XP/Windows Server 2003, and are still a part of current Windows releases. For more information, see Windows Error Reporting. You can enable WER by creating a GPO, or using Server Manager.

To create a group policy object to enable WER

1. From the domain controller, open the Group Policy Management Console. (On the **Start** screen, type **gpmmc**).
2. Right-click the domain, organizational unit, or security group where you want to create the GPO, and then click **New**.
3. In the **New GPO** dialog box, specify a name for the new GPO, and click **OK**.
4. Right-click the new GPO, and click **Edit** to open the Group Policy Management Editor.
5. In Computer Configuration\Preferences\Windows Settings, right-click **Registry**, click **New**, and then click **Registry Item**.
6. In the key path, enter 'HKLM\SOFTWARE\Policies\Microsoft\Windows\Windows Error Reporting', and then configure the following:
   - **Name**: Disabled
   - **Type**: REG_DWORD
   - **Value**: 0

   Next, in the key path, enter 'HKLM\SOFTWARE\Policies\Microsoft\Windows\Windows Error Reporting\Consent', and then configure the following:

   a. **Name**: DefaultConsent
   b. **Type**: REG_DWORD
   c. **Value**: 1-3 (**1**: query before send; **2**: send summary; **3**: send detailed report)

To enable Windows Error Reporting by using Server Manager

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. In Server Manager, click **Local Server**, and then click **Off** next to **Windows Error Reporting**. The **Windows Error Reporting Configuration** dialog box opens.
3. The following two options enable Windows Error Reporting. Each option provides a different level of detail in the information sent to Microsoft when failures occur.
   - **Yes, automatically send detailed reports. Notify me if there are possible solutions to the problem.**

     When this option is selected, if an application failure occurs, detailed reports are automatically sent to Microsoft. The contents of log files, and portions of the registry, may also be sent to Microsoft.

     Important
     Portions of documents you were working on at the moment a program failed, or data you provided to Web sites or forms may be included in these reports.
   - **Yes, automatically send summary reports. Notify me if there are possible solutions to the problem.**

     When this option is selected, if an application failure occurs, only non-personal data is sent to Microsoft. Data about your computer's configuration, the program in which you were working, and any operations that triggered the failure is sent to Microsoft.
4. The remaining two options are for disabling WER.

Disable Windows Error Reporting

Use Server Manager to disable WER.

To disable Windows Error Reporting by using Server Manager

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. In Server Manager, click **Local Server**, and then click **On** next to **Windows Error Reporting**. The **Windows Error Reporting Configuration** dialog box opens.
3. Use one of the following two options to disable Windows Error Reporting and then click **OK**.
   - **I don't want to participate, and don't ask me again**

     This option disables WER, and prevents it from prompting you to send information about application failures to Microsoft.
   - **Ask me about sending reports every time an error occurs**

     This option disables WER, but allows it to prompt you to send information about application failures to Microsoft whenever a failure occurs.

Configure all servers to participate in CEIP and WER

To opt-in multiple servers at the same time, use the following procedure.

To configure all servers to participate in CEIP and WER

1. In Server Manager, click **All Servers**.
2. Select all of the servers that you want to participate in CEIP and WER.
3. Right-click the servers as a group, and then click **Configure Windows Automatic Feedback**.
4. Under **Participate in Customer Experience Improvement Program?**, select **Yes, I want to participate**.
5. Click **Windows Error Reporting**, and then configure how you want to participate in WER.
6. When finished, click **OK**.

Install Windows Feedback Forwarder

The Windows Feedback Forwarder feature can be used on all editions of Windows Server 2012, including with Server Core installations. Use one of the following procedures to install Windows Feedback Forwarder.

Install Windows Feedback Forwarder by using Add Roles and Features Wizard

Using the Add Roles and Features Wizard, you can install roles, role services, and features on the local server, a remote server that has been added to Server Manager, or an offline VHD. The server on which you want to mount an offline VHD must be added to Server Manager, and you must have Administrator rights on that server.

To install Windows Feedback Forwarder by using Add Roles and Features Wizard

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page, verify that your destination server and network environment are prepared for the role and feature you want to install. Click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation** to install the Windows Feedback Forwarder feature. Click **Next**.
5. On the **Select destination server** page, select a server or a remote server from the server pool, or select an offline VHD. To select an offline VHD as your destination server, first select the server on which to mount the VHD, and then select the VHD file. For information about how to add servers to your server pool, see Add servers to Server Manager. After you have selected the destination server, click **Next**.
6. Select roles, select role services for the role if applicable, and then click **Next** to select features. On the **Features** page, select **Windows Feedback Forwarder**, and then click **Next**.
7. On the **Windows Feedback Forwarder Description** page, review the information, and then click **Next**.
8. On the **Windows Feedback Forwarder** page, configure the incoming port (the default is 53533) and port.
9. On the **Confirm installation selections** page, review your feature and server selections. If you are ready to install, click **Install**.

   Note
   When installing Windows Feedback Forwarder on a remote server, the Add Roles and Features Wizard writes the settings to a configuration file, and then deploys the file to the remote server.
10. After you click **Install**, the **Installation progress** page displays installation progress, results, and messages such as warnings, failures, or post-installation configuration steps that are required for the roles or features that you installed.
11. After installation is complete, use the **Configure Windows Feedback Forwarder** dialog to configure group policy settings. You can open **Configure Windows Feedback Forwarder** from the Servers page of Server Manager. For more information, see Configure Windows Feedback Forwarder.

To install the Windows Feedback Forwarder feature by using the Install-WindowsFeature cmdlet

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note

   If you are installing roles and features on a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the tile for Windows PowerShell, and then on the app bar, click **Run as Administrator**.
2. Type **Get-Help Install-WindowsFeature**, and then press **Enter** to view the syntax and accepted parameters for the `Install-WindowsFeature` cmdlet.
3. Type the following, and then press **Enter**, where *computer_name* represents a remote computer on which you want to install roles and features. The `Restart` parameter automatically restarts the destination server if required by the role or feature installation.

   ```
   Install-WindowsFeature –Name WFF -ComputerName
   <computer_name> -Restart
   ```

   To install roles and features on an offline VHD, add both the `ComputerName` parameter and the `VHD` parameter. If you do not add the `ComputerName` parameter, the cmdlet assumes that the local computer is configured to mount the VHD. The `ComputerName` parameter contains the name of the server on which to mount the VHD, and the `VHD` parameter contains the path to the VHD file on the specified server.

   Notes

   You must add the `ComputerName` parameter if you are running the cmdlet from a computer that is running Windows 8.

   If you enter the name of a VHD file from a network share, that share must grant **Read** and **Write** permissions to the computer account of the server that you selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.

   ```
   Install-WindowsFeature -Name <feature_name> –VHD <path> -
   ComputerName <computer_name> -Restart
   ```
4. When installation is finished, verify installation by opening the **All Servers** page in Server Manager, selecting a server on which you installed roles and features, and viewing the **Roles and Features** tile on the page for the selected server. You can also run the `Get-WindowsFeature` cmdlet targeted at the selected server (Get-WindowsFeature -ComputerName <*computer_name*>) to view a list of roles and features that are installed on the server.

Install Windows Feedback Forwarder on a Server Core installation

The Server Manager deployment cmdlets for Windows PowerShell function similarly to the GUI-based Add Roles and Features Wizard, with an important difference. In Windows PowerShell, unlike in the Add Roles and Features Wizard, management tools and snap-ins for a role are not included by default. To include management tools as part of a role installation, add the `IncludeManagementTools` parameter to the cmdlet. If you are installing roles and features on a server that is running the Server Core installation option of Windows Server 2012, you can add a role's management tools to an installation, but GUI-based management tools and snap-ins cannot be installed on servers that are running the Server Core installation option of Windows Server. Only command-line and Windows PowerShell management tools can be installed on the Server Core installation option of Windows Server 2012.

Configure Windows Feedback Forwarder

You can configure Windows Feedback Forwarder using either Server Manager or, with Server Core installations, Windows PowerShell. Use one of the following procedures to configure Windows Feedback Forwarder.

To configure Windows Feedback Forwarder by using Server Manager

1. On the Windows **Start** screen, click the **Server Manager** tile.
2. In Server Manager, click **All Servers**, right-click the server on which Windows Feedback Forwarder is run, and then click **Windows Feedback Forwarder Configuration** from the list of available tasks. Optionally, when configuring Windows Feedback Forwarder from the local computer, click **Local Server**, click **Tasks**, and then click **Windows Feedback Forwarder Configuration**.
3. In **Deploy Group Policy**, under **Domain Name**, enter the domain name on which you want to deploy a group policy for Windows Feedback Forwarder. Click **Find** to retrieve the domain's organizational units.
4. Under **Organizational Unit**, select the checkbox next to the organization unit on which you want to deploy group policy. Clear the checkbox next to an organizational unit that you no longer want to use with the Windows Feedback Forwarder group policy.
5. In **Forwarding Settings**, under **Internet Proxy**, you have the option to change the incoming port number from the default value of 53533. This port is used for transferring CEIP data from the client to Windows Feedback Forwarder. The port is configurable to accommodate firewall policies. If you change the port number, the firewall exception for Windows Feedback Forwarder is updated.

   The **Windows Feedback Forwarding URL** textbox displays the complete URL that is used by the Windows Feedback Forwarder group policy to transfer CEIP data.

   Under **Internet Proxy**, enter the appropriate settings for **Server/URL** and **Port**.
6. Click **Apply**.

To configure Windows Feedback Forwarder on a Server Core installation

1. Open Windows PowerShell on the computer that is running the Server Core installation.
2. Open Server Manager on the computer that is running Windows Server 2012, and then remotely connect to the computer that is running the Server Core installation. Using Server Manager, configure Windows Feedback Forwarder (following the steps in the previous procedure).

What's New in Telemetry

You can help Microsoft improve the quality, reliability, and performance of its operating systems by participating in the Microsoft Customer Experience Improvement Program (CEIP) and Windows Error Reporting (WER). Your participation allows Microsoft Corporation to collect statistical information about your system configuration, the performance of some features of Windows, and certain types of events.

Feature description

By participating in CEIP, Microsoft learns how you use Windows Server 2012, and then improves the product based on your feedback. The CEIP report collection and transfer process is lightweight to allow CEIP participants to get the most out of their servers, especially when it comes to server performance and network bandwidth. Windows Server 2012 records CEIP usage information using a high-speed tracing component called Event Tracing for Windows (ETW). ETW enables Windows Server 2012 to record CEIP usage data with no noticeable impact to server performance. CEIP usage information is then transferred to Microsoft in a two part process using the Consolidator and Uploader scheduled tasks. The Consolidator scheduled task exports CEIP data into a compressed binary format that is ready for transfer. The binary data typically has a size of less than 1 megabyte (MB) so that the transfer has minimal impact to network bandwidth. The Uploader scheduled task runs once every twenty-four hours and transfers the CEIP binary data to the Microsoft frontend servers using the Windows Telemetry Protocol.

The CEIP data collected by Microsoft consists of basic information about how servers are configured and used, the roles and features that are installed, and information about hardware. CEIP does not intentionally collect Personally Identifiable Information (PII), so CEIP reports do not contain personal contact information, such as names, addresses, or phone numbers. CEIP does not ask you to participate in surveys or to read email, and you will not be contacted in any way. The Microsoft Customer Experience Improvement Program privacy statement discusses, in detail, the data collected by CEIP and how we use it.

The following list outlines some of the ways CEIP data is used to improve products.

- **Increased server reliability**. In pre-release versions of Windows Server 2012, Reliability Analysis Component (RAC) features are enabled to determine the root order cause of why a Windows server or application crashes or stops working. RAC combines CEIP data with Windows Error Reporting (WER) data to reconstruct a full view of the system state at the time

of the crash or hang. By analyzing the combined CEIP and WER data, high occurrence issues can be identified and then triaged and fixed so that a product is improved with each release. To learn more about the data collected by WER, see the Microsoft Error Reporting Privacy Statement.

- **Improved programmability for server administration scripts**. For large-scale deployments, system administration is often performed using Windows PowerShell and WMI scripts because scripting simplifies manageability. When a Windows PowerShell cmdlet or WMI interface changes or is removed, rewriting scripts can be cumbersome when you need to accommodate the platform changes. In Windows Server 2012, CEIP addresses this issue by monitoring deprecated API usage and ensuring that APIs are not removed until they have minimal impact to you. For example, in Windows Server 2012, the Win32_ServerFeature WMI interface, which detects installed roles and features, is considered deprecated and is replaced with MSFT_ServerManagerDeploymentTasks. As part of the deprecation process, CEIP data can now record interface usage and, based on the latest Windows Server 2012 Beta CEIP data, it was found that 47% of customers are still using Win32_ServerFeature. Using this data from the usage of Win32_ServerFeature, Microsoft can make sure that Win32_ServerFeature is not formally removed from the product until migration to MSFT_ServerManagerDeploymentTasks is complete.
- **Diversity of Windows-Certified hardware**. One of the frequently asked questions Microsoft gets is *What CEIP data does Microsoft share with partners*? There are certain scenarios where a subset of CEIP data, except for Personally Identifiable Information (PII), is shared with independent hardware or software vendors as part of the hardware certification process. It is important that Windows server supports high-quality drivers for the diverse set of devices that are available. CEIP data is used to model hardware profiles and map the diversity of different devices to notify the independent hardware or software vendors of the requirements for certification. With this data, independent hardware or software vendors determine the range of drivers to certify (based on what is available) and prioritize which devices get certified first (based on popularity).
- **Improved product experience**. CEIP data is used on a day-by-day basis to understand a broad range of feature configurations, so that Microsoft can prioritize work according to how you use the product. For example, to reduce the cost of installing and configuring new servers, CEIP records what settings you use. This allows Microsoft to refine default settings by adjusting them to reflect most common usage patterns, which allows you to more quickly install and configure a new server.

New and changed functionality

The following table lists functionality in Telemetry that is new for this release or has been changed.

Feature/functionality	New or updated
Windows Feedback Forwarder	New
Windows Automatic Feedback	New

Feature/functionality	New or updated
Telemetry module for Windows PowerShell	New

Windows Feedback Forwarder

Windows Feedback Forwarder is a forwarding service that enables you to automatically and securely send feedback to Microsoft after setting up a Group Policy domain or organizational unit. After configuring Group Policy and joining the Customer Experience Improvement Program (CEIP), Windows periodically collects and uploads statistical information.

What value does this change add?

Using the Windows Feedback Forwarder feature, up to 15,000 computers in one or more domain or organizational units can be automatically opted in to the Microsoft Customer Experience Improvement Program (CEIP) through Group Policy settings.

Windows Automatic Feedback

Windows Automatic Feedback allows you to enable both CEIP and WER on multiple computers simultaneously using Server Manager.

What value does this change add?

Using Windows Automatic Feedback, you can quickly configure multiple computers to use CEIP and WER.

See also

The following table provides additional resources for evaluating Telemetry.

Content type	References
Product evaluation	Telemetry Overview
Deployment	Deploy Windows Feedback Forwarder
Operations	*Content is not available yet.*

User Access Logging Overview

This technical overview describes User Access Logging (UAL), a feature in Windows Server 2012 that aggregates client usage data by role and products on a local server.

Did you mean…

- Microsoft Assessment and Planning Toolkit
- Volume Activation Overview
- Volume Activation Management Tool

Feature description

User Access Logging (UAL) in Windows Server 2012 is a feature to help server administrators quantify requests from client computers for roles and services on a local server.

UAL is installed and enabled by default in Windows Server 2012, and collects data in nearly real-time. No administrator configuration is required, although UAL can be disabled or enabled. For more information, see Manage User Access Logging.T he User Access Logging service aggregates client usage data by roles and products into local database files. IT administrators can later use Windows Management Instrumentation (WMI) or Windows PowerShell cmdlets to retrieve quantities and instances by server role (or software product), by user, by device, by the local server, and by date.

Note
UAL supports the Microsoft Assessment and Planning Toolkit.

Practical applications

UAL aggregates unique client device and user request events that are logged on a computer running Windows Server 2012 into a local database. These records are then made available (through a query by a server administrator) to retrieve quantities and instances by server role, by user, by device, by the local server, and by date. In addition, UAL has been extended to enable non-Microsoft software developers to instrument their UAL events to be aggregated by Windows Server 2012.

This information can be useful to server administrators at all levels. UAL can assist server administrators in performing the following tasks:

- Quantify client user requests for local physical or virtual servers.
- Quantify client user requests for installed software products on a local physical or virtual server.
- Retrieve data on a local server running Hyper-V to identify periods of high and low demand on a Hyper-V virtual computer.
- Retrieve UAL data from multiple remote servers.

In addition, software developers can instrument UAL events that can then be aggregated by Windows Server 2012 and retrieved by using WMI and Windows PowerShell interfaces.

The following server roles and services can be supported by UAL:

- Active Directory Certificate Services (AD CS)
- Active Directory Rights Management Services (AD RMS)
- BranchCache
- Domain Name System (DNS)

Note
UAL collects DNS data every 24 hours, and there is a separate UAL cmdlet for this scenario.

- Dynamic Host Configuration Protocol (DHCP)
- Fax Server
- File Services
- File Transfer Protocol (FTP) Server
- Hyper-V

Note
UAL collects Hyper-V data every 24 hours, and there is a separate UAL cmdlet for this scenario.

- Web Server (IIS)

Warning
To use UAL with IIS, you must use iisual.exe. For more information, see Analyzing Client Usage Data with IIS User Access Logging.

- Microsoft Message Queue (MSMQ) Services
- Network Policy and Access Services
- Print and Document Services
- Routing and Remote Access Service (RRAS)
- Windows Deployment Services (WDS)
- Windows Server Update Services (WSUS)

Important
UAL is not recommended for use on servers that are connected directly to the Internet, such as web servers on an Internet-accessible address space, or in scenarios where extremely high performance is the primary function of the server (such as in HPC workload environments). UAL is primarily intended for small, medium, and enterprise intranet scenarios where high volume is expected, but not as high as many deployments of Windows Server 2012 that serve Internet-facing traffic volume on a regular basis.

Important functionality

The following table describes key functions of UAL and their potential value.

Functionality	Value
Collect and aggregate client request event data in near real-time.	Up to three years of data can be saved. **Important** Administrators need to enforce

Functionality	Value
	compliance of the data collected and data retention periods with the organization's privacy policy and local regulations.
Query UAL by using WMI or Windows PowerShell interfaces to retrieve client request data on a local or remote server.	UAL enables a single view of ongoing usage data. Server and enterprise administrators can retrieve this data and coordinate with business administrators to optimize use of their volume software licenses.
Enabled by default.	Server administrators do not need to configure or otherwise set up this feature for all core functionality to be available and working.

Data logged with UAL

The following user-related data is logged with UAL.

Data	Description
UserName	The user name on the client that accompanies the UAL entries from installed roles and products, if applicable.
ActivityCount	The number of times a particular user accessed a role or service.
FirstSeen	The date and time when a user first accesses a role or service.
LastSeen	The date and time when a user last accessed a role or service.
ProductName	The name of the software parent product, such as Windows, that is providing UAL data.
RoleGUID	The UAL assigned or registered GUID that represents the server role or installed product.
RoleName	The name of the role, component, or subproduct that is providing UAL data. This is also associated with a ProductName and a RoleGUID.

Data	Description
TenantIdentifier	A unique GUID for a tenant client of an installed role or product that accompanies the UAL data, if applicable.

The following device-related data is logged with UAL.

Data	Description
IPAddress	The IP address of a client device that is used to access a role or service.
ActivityCount	The number of times a particular device accessed the role or service.
FirstSeen	The date and time when an IP address was first used to access a role or service.
LastSeen	The date and time when an IP address was last used to access a role or service.
ProductName	The name of the software parent product, such as Windows, that is providing UAL data.
RoleGUID	The UAL-assigned or registered GUID that represents the server role or installed product.
RoleName	The name of the role, component, or subproduct that is providing UAL data. This is also associated with a ProductName and a RoleGUID.
TenantIdentifier	A unique GUID for a tenant client of an installed role or product that accompanies the UAL data, if applicable.

Software requirements

UAL can be used on any computer running Windows Server 2012.

See also

User Access Logging on MSDN.

Manage User Access Logging

Manage User Access Logging

This document describes how to manage User Access Logging (UAL).

UAL is a feature that can help server administrators quantify the number of unique client requests of roles and services on a local server.

UAL is installed and enabled by default and collects data on a nearly real-time basis. There are only a few configuration options for UAL. This document describes these options and their intended purpose.

To learn more about the benefits of UAL, see the User Access Logging Overview.

In this document

The configuration options covered in this document include:

- Disabling and enabling the UAL service
- Collecting and removing data
- Deleting data logged by UAL
- Managing UAL in high volume environments
- Recovering from a corrupt state
- Enable Work Folders usage license tracking

Note
This topic includes sample Windows PowerShell cmdlets that you can use to automate some of the procedures described. For more information, see Using Cmdlets.

Disabling and enabling the UAL service

UAL is enabled and runs by default when a computer running Windows Server 2012or later is installed and started for the first time. Administrators may want to turn off and disable UAL to comply with privacy requirements or other operational needs. You can turn off UAL using the Services console, from the command line, or using PowerShell cmdlets. However, to ensure that UAL does not run again the next time the computer is started you also need to disable the service. The following procedures describes both how to turn off and disable UAL.

Note
You can use the `Get-Service UALSVC` PowerShell cmdlet to retrieve information about the UAL Service including whether it is running or stopped and whether it is enabled or disabled.

To stop and disable the UAL service by using the Services console

1. Sign in to the server with an account that has local administrator privileges.
2. In Server Manager, point to **Tools**, and then click **Services**.
3. Scroll down and select **User Access Logging Service**.Click **Stop the service**.
4. Right-click the service name and select **Properties**. On the **General** tab, change the

Startup type to **Disabled**, and then click **OK**.

To stop and disable UAL from the command line

1. Sign in to the server with an account that has local administrator privileges.
2. Press the Windows logo + R, then type **cmd** to open a Command Prompt window.

   Important
   If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. Type **net stop ualsvc**, and then press ENTER.
4. Type **netsh ualsvc set opmode mode=disable**, and then press ENTER.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

You can also stop and disable UAL by using the Stop-service and Disable-Ual Windows PowerShell commands.

```
Stop-service ualsvc

Disable-ual
```

If at a later date you want to restart and reenable UAL you can do so with the following procedures.

To start and enable the UAL service by using the Services console

1. Sign in to the server an account that has local administrator privileges.
2. In Server Manager, point to **Tools**, and then click **Services**.
3. Scroll down and select **User Access Logging Service**.Click **Start the service**.
4. Right-click the service name and select **Properties**. On the **General** tab, change the **Startup type** to **Automatic**, and then click **OK**.

To start and enable UAL from the command line

1. Sign in to the server with local administrator credentials.
2. Press the Windows logo + R, then type **cmd** to open a Command Prompt window.

   Important
   If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
3. Type **net start ualsvc**, and then press ENTER.
4. Type **netsh ualsvc set opmode mode=enable**, and then press ENTER.

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

You can also start and reenable UAL by using the Start-service and Enable-Ual Windows PowerShell commands.

```
Enable-ual

Start-service ualsvc
```

Collecting UAL data

In addition to the PowerShell cmdlets described in the previous section, 12 additional cmdlets in Windows Server 2012 can be used to collect UAL data:

- **Get-UalOverview**: Provides UAL related details and history of installed products and roles.
- **Get-UalServerUser**: Provides client user access data for the local or targeted server.
- **Get-UalServerDevice**: Provides client device access data for the local or targeted server.
- **Get-UalUserAccess**: Provides client user access data for each role or product installed on the local or targeted server.
- **Get-UalDeviceAccess**: Provides client device access data for each role or product installed on the local or targeted server.
- **Get-UalDailyUserAccess**: Provides client user access data for each day of the year.
- **Get-UalDailyDeviceAccess**: Provides client device access data for each day of the year.
- **Get-UalDailyAccess**: Provides both client device and user access data for each day of the year.
- **Get-UalHyperV**: Provides virtual machine data relevant to the local or targeted server.
- **Get-UalDns**: Provides DNS client specific data of the local or targeted DNS server.
- **Get-UalSystemId**: Provides system specific data to uniquely identify the local or targeted server.

`Get-UalSystemId` is meant to provide a unique profile of a server for all other data from that server to be correlated with. If a server experiences any change in the in one of the parameters of `Get-UalSystemId` a new profile is created. `Get-UalOverview` is meant to provide the administrator with a list of roles installed and being used on the server.

Note

Basic features of Print and Document Services and File Services are installed by default in Windows Server 2012. Therefore administrators can expect to always see information on these displayed as if the full roles are installed. Separate UAL cmdlets are included for Hyper-V and DNS because of the unique data that UAL collects for these server roles.

A typical use case scenario for UAL cmdlets would be for an administrator to query UAL for unique client accesses over the course of a date range. This can be done in a variety of ways. The following is a recommended method to query unique device accesses over a date range.

```
PS C:\Windows\system32>Gwmi -Namespace "root\AccessLogging" -query "SELECT * FROM
MsftUal_DeviceAccess WHERE LastSeen >='1/01/2013' and LastSeen <='3/31/2013'"
```

This will return a verbose listing of all unique client devices, by IP address, that have made requests of the server in that date range.

'ActivityCount' for each unique client is limited to 65,535 per day. Also, calling into WMI from PowerShell is only required when you query by date. All other UAL cmdlet parameters can be used within PS queries as expected, as in the following example:

```
PS C:\Windows\system32> Get-UalDeviceAccess -IPAddress "10.36.206.112"


ActivityCount    : 1
FirstSeen        : 6/23/2012 5:06:50 AM
IPAddress        : 10.36.206.112
LastSeen         : 6/23/2012 5:06:50 AM
ProductName      : Windows Server 2012 Datacenter
RoleGuid         : 10a9226f-50ee-49d8-a393-9a501d47ce04
RoleName         : File Server
TenantIdentifier : 00000000-0000-0000-0000-000000000000
PSComputerName
```

UAL retains up to two years' worth of history. To allow retrieval of UAL data by an administrator when the service is running, UAL makes a copy of the active database file, current.mdb, to a file named *GUID.mdb* every 24 hours for the WMI provider's use.

On the first day of the year, UAL will create a new *GUID.mdb*. The old *GUID.mdb* is retained as an archive for the provider's use. After two years, the original *GUID.mdb* will be overwritten.

Important

The following procedure should be performed only by an advanced user and would commonly be used by a developer testing their own instrumentation of UAL application programming interfaces...

To adjust the default 24-hour interval to make data visible to the WMI provider

1. Sign in to the server with an account that has local administrator privileges.
2. Press the Windows logo + R, then type **cmd** to open a Command Prompt window.
3. Add the registry value:
   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\WMI\AutoLogger\Sum\PollingInterval (REG_DWORD).

   Warning

   Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

   The following example shows how to add a two-minute interval (not recommended as a

long term running state): **REG ADD HKLM\System\CurrentControlSet\Control\WMI\AutoLogger\Sum /v PollingInterval /t REG_DWORD /d 120000 /F**

Time values are in milliseconds. The minimum value is 60 seconds, the maximum is seven days, and the default is 24 hours.

4. Use the Services console to stop and restart the User Access Logging Service.

Deleting data logged by UAL

UAL is not intended to be a mission critical component. Its design is intended to impact local system operations as little as possible while maintaining a high level of reliability. This also allows the administrator to manually delete UAL database and supporting files (every file in \Windows\System32\LogFiles\SUM\ directory) to meet operational needs.

To delete data logged by UAL

1. Stop the User Access Logging Service.
2. Open Windows Explorer.
3. Go to **\Windows\System32\Logfiles\SUM**.
4. Delete all files in the folder.

Managing UAL in high volume environments

This section describes what an administrator can expect when UAL is used on a server with high client volume:

The maximum number of accesses that can be recorded with UAL is 65,535 per day. UAL is not recommended for use on servers that are connected directly to the Internet, such as web servers that are connected directly to the Internet, or in scenarios where extremely high performance is the primary function of the server (such as in HPC workload environments). UAL is primarily intended for small, medium, and enterprise intranet scenarios where high volume is expected, but not as high as many deployments of Windows Server 2012 that serve Internet-facing traffic volume on a regular basis.

UAL in Memory: Because UAL uses the Extensible Storage Engine (ESE), UAL's memory requirements will increase over time (or by quantity of client requests). But memory will be relinquished as the system requires it to minimize impact on system performance.

UAL on Disk: UAL's hard disk requirements are approximately as shown below:

- 0 unique client records: 22M
- 50,000 unique client records: 80 M
- 500,000 unique client records: 384M
- 1,000,000 unique client records: 729M

Recovering from a corrupt state

This section discusses UAL's use of the Extensible Storage Engine (ESE) at a high level and what an administrator can do if UAL data is corrupted or unrecoverable.

UAL uses ESE to optimize use of system resources and for its resistance to corruption. For more information about the benefits of ESE, see Extensible Storage Engine on MSDN.

Each time the UAL service starts ESE performs a soft recovery. For more information, see Extensible Storage Engine Files on MSDN.

If there is a problem with the soft recovery, ESE will perform a crash recovery. For more information, see JetInit Function on MSDN.

If UAL is still unable to start with the existing set of ESE files, it will delete all files in the \Windows\System32\LogFiles\SUM\ directory. After these files are deleted, the User Access Logging Service will restart and new files are created. The UAL service will then resume as if on a freshly installed computer.

Important

The files in UAL database directory should never be moved or modified. Doing so will initiate the recovery steps, including the cleanup routine described in this section. If backups of the\Windows\System32\LogFiles\SUM\ directory are needed, then every file in this directory must be backed up together in order for a restore operation to function as expected.

Enable Work Folders usage license tracking

Work Folders server can use UAL to report client usage. Unlike UAL, Work Folder logging is not turned on by default. You can enable it with the following regkey change:

```
Reg add HKLM\Software\Microsoft\Windows\CurrentVersion\SyncShareSrv /v
EnableWorkFoldersUAL /t REG_DWORD /d 1
```

After the regkey is added, you must restart the SyncShareSvc service on the server, to enable logging.

After Work Folder logging is enabled, there will be two informational events logged in the Windows Logs\Application channel each time a client connects to the server. The more devices that connect, the longer it takes the server to complete logging its two events for every device and then start the query over again. For example, if there are 1000 users each with two devices, it can take up to 70 minutes for the server to log the two events for each device and then start over again. This can make troubleshooting the server difficult since you may have to wait for extended periods of time before the event you are troubleshooting gets logged. Consider the number of clients that will connect to your server before enabling usage license tracking for Work Folders.

See also

- User Access Logging Overview

Volume Activation Overview

This technical overview for the IT professional describes the volume activation technologies in Windows operating systems and how your organization can benefit from using these technologies to deploy and manage volume licenses for a medium to large number of computers.

Did you mean…

- Volume Activation for Windows 7 and Windows Server 2008 R2
- Microsoft Product Activation (Information about end-user licensing and how to license individual copies of Windows that are purchased at retail locations)
- Remote Desktop Licensing Manager

Volume activation description

A number of technologies are available in Windows client and server operating systems to simplify the task of configuring the distribution and management of an organization's volume software licenses. These technologies include:

- **The Volume Activation Services server role** Volume Activation Services is a server role Windows Server starting with Windows Server 2012 that enables you to automate and simplify the issuance and management of Microsoft software volume licenses for a variety of scenarios and environments. With Volume Activation Services, you can install and configure the Key Management Service (KMS) and enable Active Directory-based Activation.
- **Key Management Service (KMS)** KMS is a role service that allows organizations to activate systems within their network from a server where a KMS host has been installed. With KMS, IT pros can complete activations on their local network, eliminating the need for individual computers to connect to Microsoft for product activation. KMS does not require a dedicated system, and it can be cohosted on a system that provides other services. By default, volume editions of Windows client and server operating systems connect to a system that hosts the KMS service to request activation. No action is required from the user.
- **Active Directory-based Activation** Active Directory-based Activation is a role service that allows you to use Active Directory Domain Services (AD DS) to store activation objects, which can further simplify the task of maintaining volume activation services for a network. With Active Directory-based Activation, no additional host server is needed, and activation requests are processed during computer startup.

  Any computers running a client or server operating system with a Generic Volume License Key (GVLK) that are connected to the domain will activate automatically and transparently. They will stay activated as long as they remain members of the domain and maintain periodic contact with a domain controller. Activation takes place after the licensing service starts. When this service starts, the computer running a Windows client or server operating system contacts AD DS automatically, receives the activation object, and activates without user intervention.

  Note

AD DS must be at the schema level of Windows Server 2012 or newer to store activation objects. Domain controllers running earlier versions of Windows Server can activate clients after their schemas have been updated using the Windows Server 2012 version of Adprep.exe. For more information, see Running Adprep.exe.

- **Volume Activation Tools console** The Volume Activation Tools snap-in is installed with KMS and Active-Directory–based Activation. You can use the Volume Activation Tools console to install, activate, and manage one or more volume activation license keys in AD DS or on a KMS host.

Other Microsoft volume activation technologies that can be used in conjunction with Windows server operating systems include:

- **Volume Activation Management Tool (VAMT)** The VAMT is a free downloadable MMC snap-in that is included in the Windows Assessment and Deployment Kit (ADK) that enables network administrators and other IT professionals to automate and centrally manage the Windows and Microsoft Office volume and retail activation process. VAMT can manage volume activation by using Multiple Activation Keys (MAKs) or the Windows Key Management Service (KMS).

  VAMT 3.0 includes all the major functionality of the previously released versions of VAMT, and it includes an improved interface, data storage in a SQL Server database for greater scalability and speed, new reporting options, the ability to use Windows PowerShell cmdlets, and new context-sensitive Help. In addition, VAMT 3.0 supports Active Directory-based Activation in addition to KMS and MAK activation. VAMT version 3.1 is included in the Windows ADK. For more information about VAMT, see Volume Activation Management Tool (VAMT) Technical Reference, and VAMT How-To Topics
- **Multiple Activation Key (MAK)** A MAK is a volume license key that is used for one-time activation with activation services that are hosted by Microsoft. There are two ways to use MAK to activate computers:
  - **MAK independent activation** Each computer must independently connect and be activated by Microsoft over the Internet or by telephone.
  - **MAK proxy activation** A computer that is acting as a MAK proxy gathers activation information from multiple computers on the network, and then sends a centralized activation request to Microsoft on their behalf. MAK proxy activation is configured by using the Volume Activation Management Tool (VAMT).

How does volume activation work?

Before you configure and use volume activation technologies, it may be helpful to understand how volume activation works.

In an environment that is using a KMS host, the volume activation process is as follows:

1. KMS is installed on a server.
2. The KMS host key is installed on the KMS host, and then activated by contacting the hosted activation services at Microsoft.
3. After the KMS is enabled, it registers an SRV record in Domain Name System (DNS) each time KMS is started or once per day.

4. A KMS client computer discovers the KMS host from a configured registry entry or through the KMS SRV record in DNS.
5. The client sends an RPC request to the KMS host on 1688/TCP (the default setting). This request includes an encrypted client computer ID. If the computer has never been activated and there is no response from the KMS host, the client sends a new request after two hours. If the computer has been activated, the client sends a new reactivation request after seven days.
6. The KMS host adds the client computer ID to a table and returns the activation count to the client.
7. The client evaluates the activation count against the license policy and activates if the activation threshold is met.

Note

KMS maintains an activation threshold, which requires that the organization activate at least five computers (physical or virtual machines) running a server operating system with Windows Server 2012 or later and at least 25 client computers running Windows 8 or later. The computers running Windows Server 2012 must connect within a 30-day period to active KMS client computers.

In an environment that is using Active Directory-based Activation, the volume activation process is as follows:

1. A user with Enterprise Administrator permissions installs the Active Directory-based Activation role service on a domain controller, including the KMS host key, and then activates the KMS host key with the activation services that are hosted by Microsoft.

Note

This installation can be completed from any computer where the Volume Activation Tools console is installed.

2. When a domain-joined computer running Windows with a GVLK starts, the licensing service on the client automatically queries the domain controller for licensing information.

Note

Active Directory-based Activation cannot be used to license computers that are not members of the domain.

3. If a valid activation object is found, activation proceeds silently without user intervention. The same renewal guidelines apply to Active Directory-based Activation as to KMS activation.

Note

If volume license information is not found in AD DS, clients running Windows look for a KMS host, and then attempt activation by following the KMS activation process.

If Active Directory-based Activation or KMS activation are not options (such as to license computers in remote locations), MAK keys can be distributed as part of an operating system image. This is done through WMI script and the change product key wizard, or by using VAMT. After a MAK is installed, each client needs to connect to Microsoft individually over the Internet or by telephone to complete activation.

A second MAK validation option is available for client computers that are not connected to the Internet, such as in remote or highly secured environments. In MAK Proxy activation, VAMT installs a MAK product key on a client computer, obtains the installation ID (IID) from the target computer, sends the IID to Microsoft on behalf of the client, and obtains a confirmation ID (CID). The tool then activates the client computer by installing the CID.

Practical applications

Validating the licenses and activating tens, hundreds, or thousands of computers in multiple locations is an important network administration task. By doing so, you verify that operating systems and applications are genuine, have not been compromised, and will be supported. Volume Activation Services in can help you accomplish these tasks.

The following are common licensing and activation-related questions.

What is activation?

Product activation is the process of validating software with the manufacturer. Activation confirms the genuine status of a product and that the product key has not been compromised. It establishes a relationship between the software's product key and a particular installation of that software on a device. In earlier versions of the Windows operating system, activation and validation (using the Windows Genuine tool) occurred separately. This caused confusion for users who thought the terms were interchangeable. In Windows operating systems starting with Windows Server 2012 or Windows 8, activation and validation occur at the same time.

What are my activation options?

For released versions of Windows client and server operating systems, licenses can be obtained through the following channels:

- **Retail** Microsoft software products that are acquired through a retail store are individually licensed and come with one unique product key (printed on the product packaging), which the user enters during product installation. The computer uses this product key to complete activation after the software installation is complete. This final activation step can be accomplished online or by telephone.
- **Original equipment manufacturer (OEM)** Most OEMs sell systems that include a standard build of Windows and other software. OEM activation is performed by associating the operating system to the firmware of the computer. This process occurs before the computer is sent to the customer so that the software is already activated and no additional actions are required by the user. OEM activation remains valid as long as the customer uses the OEM-provided image on the computer. If customers want to create customized images, they can use the image provided by the OEM as the basis for creating custom images.
- **Volume licensing** A volume license is a product key that allows an organization to activate a specific number of clients to use a particular piece of software. Microsoft Volume Licensing offers customized programs that are tailored to the size and purchasing preference of a

variety of organizations. To become a volume licensing customer, the organization needs to set up a volume license agreement with Microsoft.

What happens if systems are not activated?

Activation is designed to be relatively transparent to users. However, the operating system must be activated immediately upon installation. Failure to activate the Windows operating systems will prevent users from completing customization.

When should I use volume activation?

Customers can use one or all three of the volume activation scenarios described earlier (KMS, Active Directory-based activation, or MAK) to activate systems in their environments.

- **When to use the Key Management Service (KMS)**: KMS is the default key for volume activation clients, regardless of the Windows version they are running. If you have a network where multiple versions of Windows client and server operating systems are deployed, you will probably want to use KMS to activate them.

  Important
  Installations of KMS on Windows 7 SP1 or Windows Server 2008 R2 SP1 must be updated in order to support activation of clients running operating systems starting with Windows 8 or Windows Server 2012. For more information, see An update that adds support for Windows 8 and for Windows Server 2012 to Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2 KMS hosts is available.

  KMS enables organizations to activate systems within their network, eliminating the need for individual computers to connect to Microsoft for product activation.

  To use KMS activation, the computer must have the qualifying operating system license (often obtained through OEMs as part of a new computer purchase), and they must contain a Windows marker in the BIOS.

- **When to use Active Directory-based Activation**: Active Directory-based Activation simplifies the process of activating clients that are running an operating system starting with Windows 8 or Windows Server 2012. In addition, it requires the AD DS schema of Windows Server 2012 or later.

  Activation objects themselves are not editable. However, an administrator with proper permissions can use advanced AD DS tools to view each activation object and to configure security access control lists (ACLs) for the activation objects to restrict access as needed. Administrators can delete activation objects, if necessary. On a local client, a user can use the command line to perform these functions if they have Write access to the activation object.

- **When to use Multiple Activation Key (MAK)**: Not all of an organization's computers are conveniently located on the organization's networks. Computers may be laptops that are in the possession of employees who travel frequently, may be located in branch offices that do

not have good connectivity to the host network, or may be part of extremely secure networks that are disconnected from the main network. You can use MAK to activate these clients.

Note
For more information about configuring and using MAK activation, see the Volume Activation Deployment Guide

System requirements

Volume activation services can be installed and used on any computer running Windows Server 2012 or later operating system.

Active Directory-based Activation requires that AD DS be at the schema level of Windows Server 2012 or later in order to store activation objects. Domain controllers that are running earlier versions of Windows Server can activate clients after their schemas have been updated using the Windows Server 2012 version of Adprep.exe.

KMS can be installed by using the Server Manager on any computer running Windows Server 2012 or later operating system. It can also be installed from the command line and used on computers running earlier versions of Windows Server.

KMS can be used to activate computers running any version of Windows that can be a member of a domain.

Additional information

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012, Windows Server 2012 R2, Windows 8, and Windows 8.1
Product evaluation		Test Lab Guide: Demonstrate Volume Activation Services
Procedures	Volume Activation Operations Guide	
Technical content	Volume Activation Technical Reference Guide	
Design, planning and deployment	Volume Activation Planning Guide Volume Activation Deployment Guide	Plan for Volume Activation
Troubleshooting	Solutions to Common Volume Activation Errors	

Resource	Windows Server 2008 R2 and Windows 7	Windows Server 2012, Windows Server 2012 R2, Windows 8, and Windows 8.1
General information and additional resources	Volume Activation Management Tool (VAMT) Overview	Plan volume activation of Office 2013 Preview Volume activation methods in Office 2013 Preview

Test Lab Guide: Demonstrate Volume Activation Services

This test lab guide describes the steps for configuring the Volume Activation Services server role in a Windows Server 2012 demonstration or test environment. Although the same steps can be used to deploy Volume Activation Services in a production environment, additional planning would be needed to ensure that your Volume Activation Services deployment meets the needs of your environment.

Setting up Volume Activation Services in a test environment

The following steps are needed to set up Volume Activation Services in a Windows Server 2012 test lab:

1. Step 1: Set up the base test lab configuration
2. Step 2: Install the Volume Activation Services server role
3. Step 3: Configure Active Directory-based Activation
4. Step 4: Verify that Active Directory-based Activation works
5. Step 5: Configure Key Management Services (KMS)
6. Step 6: Verify that KMS volume activation works

Step 1: Set up the base test lab infrastructure

Test labs enable you to get valuable hands-on experience with new products and technologies. The following base test lab configuration uses a predefined and tested methodology that results in a working configuration. The following instructions define what servers to create, how to configure the operating systems and system services, and how to install and configure additional products or technologies. This enables you to see all of the front-end and back-end components that are required for a product or technology, such as Volume Activation Services, or for a multiproduct or technology solution.

To prepare a base test lab infrastructure

1. Create the base configuration test lab as described in the following document: Test Lab Guide: Base Configuration (http://go.microsoft.com/fwlink/?LinkId=236358). The base configuration test lab contains the following computers:
   - **DC1** A domain controller running Windows Server 2012.

     Note
     AD DS must be at the Windows Server 2012 schema level to store activation objects. Domain controllers running earlier versions of Windows Server can activate clients after their schemas have been updated using the Windows Server 2012 version of Adprep.exe. For more information, see **Running Adprep.exe**.A new installation of Windows Server 2012 and AD DS with the configuration settings that are described in the base configuration test lab guide will be at the Windows Server 2012 schema level.
   - **APP1** A computer running Windows Server 2012 that can be configured as a Key Management Service host.
   - **CLIENT1** A computer running Windows 8 or Windows Server 2012 to serve as a volume licensing client for Active Directory-based Activation.
   - **CLIENT2 (optional)** A computer running Windows 7 or Windows Server 2008 R2 that is added to the base configuration to test Active Directory-based Activation and KMS activation at the same time.

     Note
     Configuration of CLIENT2 is not covered in the base configuration test lab guide.
2. Complete all the network configuration tasks for the base test lab.

   Note
   Local administrator permissions are needed on all computers in this configuration to complete these setup steps. Enterprise admin permissions are needed to manage or modify Active Directory Domain Services on DC1.

Step 2: Install the Volume Activation Services server role

The Volume Activation server role can be installed by using Windows Server Manager, Windows PowerShell DISM cmdlets, or Windows PowerShell ServerManager cmdlets.

To install Volume Activation Services by using Server Manager:

1. On either server running Windows Server 2012, open **Windows Server Manager**
2. Click **Manage**, click **Add Roles and Features**, and then click **Next**.
3. On the **Select installation type** page, select **Role-based or Feature-based installation**, and then click **Next**.

4. On the **Select destination server** page, select DC1 from the **SERVER POOL**, and then click **Next**.
5. On the **Select server roles** page, select **Volume Activation Services**, and then click **Next** three times.
6. On the **Confirm installation selections** page, confirm that the information is correct, then click **Install**.
7. On the **Installation progress** page, wait until the progress indicator shows that the installation was successful. Because additional configuration is required, click the **Volume Activation Tools** link. At this point, move to Steps 3 or 5 later in this document to configure either KMS-based volume activation or Active Directory-based Activation.

To install the Volume Activation Services by using the DISM cmdlets

1. To import the DISM cmdlets, open a Command Prompt window, type **import-module dism**, and then press ENTER.
2. Open a Windows PowerShell console and at the command prompt, type the following command, then press ENTER: **enable-windowsoptionalfeature -online -featurename VolumeActivation-Full-Role**

To install the Volume Activation Services by using the ServerManager cmdlets

1. To import the ServerManager cmdlets, open a Command Prompt window and type: **import-module servermanager**
2. Open a Windows PowerShell console and at the command prompt, type the following command, then press ENTER: **Install-WindowsFeature VolumeActivation**

Step 3: Configure Active Directory-based Activation

To configure Active Directory-based Activation

1. On DC1, open Server Manager, select **Tools**, and then click **Volume Activation Tools**.
2. On the **Select Volume Activation Method** page, select **Active Directory-based Activation**. If the account you are currently using does not have enterprise administrator permissions, enter the credentials for an account that has permissions to create a new container on the domain controller, and then click **Next**.
3. Enter the KMS host key and an optional name for the Active Directory object, and then click **Next**.

   Note
   Although license data is presented in English, the optional name option supports the Unicode character set, which enables you to more easily identify each volume license key in any language.
4. Activate the key online or by phone, and then click **Commit**. If you have no other changes, click **Close**.

After the KMS host key has been activated, client computers that are joined to a domain running Windows Server 2012 or Windows 8 with generic volume license keys (GVLKs) installed will automatically activate without user intervention.

Step 4: Verify that Active Directory-based Activation works

All Active Directory-based Activation events are logged to the Windows Application event log under the Source: Microsoft-Windows-Security-SPP. Look at Event 12308 for confirming information.

With clients running Windows Server 2012 or Windows 8, activation should occur automatically the next time the computer is started and the user logs on.

Note

If you set up multiple domain controllers, replication between domain controllers must be complete before clients can access this data. Therefore, in a production environment, you might need to wait until replication is complete before you start the client system and the user logs on to complete Active Directory-based Activation.

To verify that Active Directory-based Activation works

1. Log on to CLIENT1.
2. Open Windows Explorer, right-click **Computer**, and select **Properties**.
3. Scroll down to the **Windows activation** section and verify that this client has been activated.

Important

Active Directory-based Activation will not work for operating systems earlier than Windows Server 2012 or Windows 8. It also will not work with Microsoft Office 2010. Use KMS volume activation to activate Windows clients and applications that do not support Active Directory-based Activation.

Step 5: Configure Key Management Services

In earlier versions of Windows, Key Management Services (KMS) could be installed only from the command line. The following procedure describes how to complete this task by using the Volume Activation Tools console in Windows Server 2012.

To configure KMS

1. On DC1, open the **Volume Activation Tools** console.

   If you did not open the console when you completed the server role installation in Step 2, open Server Manager, click **Volume Activation Services**, right-click the name of the server on which you installed the server role, and then select **Volume Activation Tools**.

When the Volume Activation Tools console is open, click **Next**.

2. Select **Key Management Service (KMS)** and select APP1 as the server that you want to install KMS on.
3. If the account you are using does not have administrator permissions on APP1, enter the User name and Password for an account that does in the boxes, and then click **Alternate credentials (optional)**.

   After you have completed both tasks, click **Next**.
4. On the **Manage KMS Host** page, enter your **KMS host key**, and then click **Commit**.
5. On the **Product Key Installation Succeeded** page, click **Activate Products**, and then click **Next**.
6. Activate your key online or by phone, and then click **Next**.
7. On the **Configure Key Management Service options** page, review the default configuration options, click **Commit**, and then click **Close** if you have no other changes.

   Note
   Administrators should validate the firewall configurations in their environment to ensure that the appropriate exceptions are enabled for KMS traffic. If you need to change the firewall or other default options at a later date, you can open the Volume Activation Tools console and return to this dialog box at any time to modify the configuration.

Step 6: Verify that KMS volume activation works

You can verify KMS volume activation from the KMS host server, APP1 and from the client computer, CLIENT2.

Note
If you set up and configured Active Directory-based Activation before configuring KMS activation, you must use a client computer that is not running Windows 8 or Windows Server 2012to verify that KMS activation works. Otherwise, the activation will be complete before your clients contact the server that is hosting KMS.

Important
KMS volume activation requires a minimum threshold of 25 computers before activation quests will be processed. The verification process described here will increment the activation count each time a client computer contacts the KMS host. However, unless the activation threshold is reached, the verification will take the form of an error message rather than a confirmation message.

To verify that KMS volume activation works

1. On APP1, open the event log and confirm that DNS publishing is successful.

   On CLIENT2, open a Command Prompt window, type **Slmgr.vbs /ato**, and then press

ENTER. The response should show the license state and detailed Windows version information.

2. On either CLIENT2 or APP1, open an elevated Command Prompt window, type **Slmgr /dlv**, and then press ENTER. The response should return an error that states the KMS activation count is too low. This confirms that KMS is functioning correctly, even though the client has not been activated.

Additional information

For more information about Volume Activation Services, see the Volume Activation Overview.

For more information that can help you prepare to deploy Volume Activation Services in your organization, see Plan for Volume Activation.

For more information about Microsoft test lab guides, see Test Lab Guides on the Microsoft TechNet Wiki.

To provide the authors of this guide with feedback or suggestions for improvement, send an email message to tlgfb@microsoft.com.

Plan for Volume Activation

This document describes critical planning considerations that you must address to deploy Volume Activation services in your organization when you are running Windows Server 2012 or later operating system. The planning steps described in this document are:

- Step 1: Review and Select Activation Methods
- Step 2: Evaluate Client Connectivity
- Step 3: Determine Activation Method and Product License Requirements
- Step 4: Determine Monitoring and Reporting Needs

After you have completed these planning steps, see Test Lab Guide: Demonstrate Volume Activation Services for procedures to deploy Volume Activation Services. For more information, see the Volume Activation Overview.

Step 1: Review and Select Activation Methods

The first step in planning for Volume Activation services in Windows starting with Windows Server 2012 is to evaluate the features and deployment models that are available for activating client computers. When you understand the differences between and advantages of these features and deployment models, move to Step 2: Evaluate Client Connectivity and use this information to assess the types of network and clients that you have. Then move to Step 3: Determine

Activation Method and Product License Requirements to identify which solution is best for each group of clients.

The Volume Activation server role in Windows Server provides three methods for implementing volume activation. You can use one, two, or all three methods of volume activation. The methods you choose will depend on the size, network infrastructure, connectivity, and security requirements of your organization. The following table highlights these methods.

Task	Description
1.1. Plan for Active Directory-based Activation	Identify and resolve key planning issues that are associated with Active Directory-based Activation.
1.2. Plan for Key Management Services (KMS) activation	Identify and resolve key planning issues that are associated with KMS activation.
1.3. Plan for Multiple Activation Key (MAK) activation	Identify and resolve key planning issues that are associated with MAK activation.

When you understand your volume activation options, you choose the right combination of methods to meet your organization's needs.

1.1 Plan for Active Directory-based Activation

Active Directory-based Activation is a role service that allows you to use Active Directory Domain Services (AD DS) to store activation objects, which can greatly simplify the task of maintaining volume activation services for a network.

With Active Directory-based Activation, any computers running operating systems starting with Windows 8 or Windows Server 2012 with a KMS client setup key (GVLK) that are connected to the domain will activate automatically and transparently during computer startup. These clients stay activated as long as they remain members of the domain and maintain periodic contact with a domain controller.

Activation takes place after the Software Protection service starts. When the Software Protection service starts, the computer contacts AD DS automatically, receives the activation object, and activates without user intervention.

The following are planning considerations when working with Active Directory-based Activation:

- With Active Directory-based Activation, you do not need an additional host server; your existing domain controllers can support activation clients, with the following limitations:
  - Active Directory-based Activation cannot be configured on read-only domain controllers.
  - Active Directory-based Activation cannot be used with non-Microsoft directory services.
  - AD DS must be at the Windows Server 2012 schema level to store activation objects. Domain controllers running earlier versions of Windows Server can activate clients after

their schemas have been updated using the Windows Server 2012 version of Adprep.exe. For more information, see **Running Adprep.exe**.

> **Note**
> For Windows Server 2012 R2, you do not have to run Adprep.exe again; the Windows Server 2012 schema level is adequate.

- You only need one activation object forest.

> **Note**
> When the client and server KMS host keys are activated within an environment, two activation objects will exist in the forest. In this case, the activation object on the server activates computers running Windows 8 or Windows Server 2012 within the forest and the client activation object is not used. There is no harm in keeping the client activation object, but it can also safely be deleted.

- Unlike using KMS, you do not need to meet any threshold limits before clients or servers can be activated.

The following are additional planning considerations for Active Directory-based Activation.

Using Active Directory-based Activation in mixed Windows activation environments

Many organizations have elaborate volume license infrastructures to support KMS and Office installations. To add Active Directory-based Activation to these environments, administrators must assess their current implementations and determine what role Active Directory-based Activation will play in the environment.

Some considerations include how to upgrade these operating systems and applications to versions that support Active Directory-based Activation.

For environments that will run Windows 8, Windows Server 2012, or newer client or server operating systems exclusively, Active Directory-based Activation is a suitable option for activating all clients and servers, and you may be able to remove any KMS hosts from your environment.

If an environment will continue to contain earlier volume-licensed operating systems and applications, administrators need a KMS host to maintain activation status for earlier volume-licensed editions of Windows and Office in addition to enabling Active Directory-based Activation for Windows 8 and Windows Server 2012 (or newer) clients.

Utilizing Active Directory-based Activation failover capability

There can also be advantages to maintaining more than one activation method for your environment. To understand why, consider what happens when clients need to be reactivated. The activation status of clients that are joined to the domain and activated with Active Directory-based Activation lasts for 180 days. When a reactivation event occurs on an installation, the client queries AD DS for the activation object. Client systems examine the activation object and

compare it to the local edition as defined by the GVLK. If the object and GVLK match, reactivation occurs.

If the AD DS object is not reachable, clients attempt to use activation methods in the following order:

1. Active Directory-based Activation
2. KMS activation
3. MAK activation

If an installation was previously activated, a successful activation event will reset the activation period to 180 days. If an activation attempt fails, activation will be attempted every seven days by default until reactivation is successful. If the full 180-day activation period expires, the activation fails, a notification is displayed to the user, and an event is logged in the Windows Application Event Log.

If Active Directory-based Activation is not available, computers that were originally activated by using Active Directory-based Activation attempt to use KMS activation, and they continue to do so until the activation is successful. The next time the computer attempts to reactivate, the client first attempts to use Active Directory-based Activation. If that attempt is successful, the client is again based on Active Directory-based Activation.

Computers that are no longer members of a domain will fail activation when the computer or Software Protection service are restarted.

1.2 Plan for Key Management Services activation

The following information outlines initial planning considerations that you need to review for Key Management Services (KMS) activation.

KMS uses a client-server model to active clients. KMS clients connect to a KMS server, called the KMS host, for activation. The KMS host must reside on your local network.

KMS hosts do not need to be dedicated servers, and KMS can be cohosted with other services. You can run a KMS host on any physical or virtual system that is running Windows Server 2012, Windows 8, Windows Server 2008 R2, Windows Server 2008, Windows 7, or Windows Vista SP1 or SP2.

A KMS host running on Windows 8, Windows 7, or Windows Vista can only activate computers running client operating systems.

The following table summarizes KMS host and client requirements for networks that include Windows Server 2012 and Windows 8 (or newer) clients.

Product key group	KMS can be hosted on	Windows editions activated by this KMS host
Volume license for Windows 8	Windows Vista (An update is required. See the note following this table.)	Windows 8 Professional Windows 8 Enterprise Windows 7 Professional

Product key group	KMS can be hosted on	Windows editions activated by this KMS host
	Windows 7 (An update is required. See the note following this table.) Windows 8	Windows 7 Enterprise Windows Vista Business Windows Vista Enterprise
Volume license for Windows Server 2012	Windows Server 2008 (An update is required. See the note following this table.) Windows Server 2008 R2 (An update is required. See the note following this table.) Windows Server 2012 Standard Windows Server 2012 Datacenter	Windows Server 2012 (all editions) Windows 8 Professional Windows 8 Enterprise Windows 2008 (all editions) Windows 2008 R2 (all editions) Windows 7 Professional Windows 7 Enterprise Windows Vista Business Windows Vista Enterprise

Note
Installations of KMS on Windows Vista, Windows 7, Windows Server 2008, or Windows Server 2008 R2 must be updated in order to support activation of clients running Windows 8, Windows Server 2012, or newer client or server operating systems. For more information, see Update adds support for Windows 8 and Windows Server 2012 to Windows Server 2008, Windows 7, and Windows Server 2008 R2 KMS hosts.
Installations of KMS On Windows Server 2003 cannot be updated to support activation of clients running Windows 8, Windows Server 2012, or newer operating systems.

A single KMS host can support an unlimited number of KMS clients. If you have more than 25 clients, we recommend that you have at least two KMS hosts in case one of your KMS hosts becomes unavailable. Most organizations can operate with as few as two KMS hosts for their entire infrastructure.

Addressing KMS operational requirements

KMS can activate physical and virtual computers, but to qualify for KMS activation, a network must have a minimum number of computers (called the activation threshold). KMS clients activate only after this threshold is met. To ensure that the activation threshold is met, a KMS host counts the number of computers that are requesting activation on the network.

KMS hosts count the most recent connections. When a client or server contacts the KMS host, the host adds the machine ID to its count and then returns the current count value in its response. The client or server will activate if the count is high enough. Clients will activate if the count is 25

or higher. Servers and volume editions of Microsoft Office products will activate if the count is five or greater. The KMS only counts unique connections from the past 30 days, and only stores the 50 most recent contacts.

KMS activations are valid for 180 days, a period known as the activation validity interval. KMS clients must renew their activation by connecting to the KMS host at least once every 180 days to stay activated. By default, KMS client computers attempt to renew their activation every seven days. After a client's activation is renewed, the activation validity interval begins again.

Addressing KMS functional requirements

KMS activation requires TCP/IP connectivity. KMS hosts and clients are configured by default to use Domain Name System (DNS). By default, KMS hosts use DNS dynamic update to automatically publish the information that KMS clients need to find and connect to them. You can accept these default settings, or if you have special network and security configuration requirements, you can manually configure KMS hosts and clients.

After the first KMS host is activated, the KMS key that is used on the first host can be used to activate up to five more KMS hosts on your network. After a KMS host is activated, administrators can reactivate the same host up to nine times with the same key.

If your organization needs more than six KMS hosts, you should request additional activations for your organization's KMS key—for example, if you have ten physical locations under one volume licensing agreement and you want each location to have a local KMS host.

Note

To request this exception, contact your Activation Call Center. For more information, see Microsoft Volume Licensing.

Computers that are running volume licensing editions of Windows 8.1, Windows Server 2012 R2, Windows 8, Windows Server 2012, Windows 7, Windows Server 2008 R2, Windows Vista, and Windows Server 2008 are, by default, KMS clients with no additional configuration needed.

If you are converting a computer from a KMS host, MAK, or retail edition of Windows to a KMS client, install the applicable KMS Client Setup Key. For more information, see Appendix A: KMS Client Setup Keys.

KMS clients can locate a KMS host automatically by querying DNS for SRV records that publish the KMS service. If your network environment uses a nonstandard DNS configuration or complex DNS zoning, an administrator must manually configure a KMS client to use a specific KMS host. For more information, see Configuring KMS Clients.

1.3 Plan for Multiple Activation Key activation

The following information outlines initial considerations that you need to review for Multiple Activation Key (MAK) activation.

MAK is used for a one-time activation with activation services that are hosted by Microsoft. Each MAK key has a predetermined number of allowed activations. This number is based on your

volume licensing agreements, and it does not match your organization's exact license count. Each activation that uses a MAK with activation services that are hosted by Microsoft counts towards the activation limit.

A MAK activation is recommended for computers that rarely or never connect to the corporate network and for environments where the number of physical computers that need activation does not meet the KMS activation threshold.

Note
Computers running Windows 8 or Windows Server 2012 or newer operating systems can be converted from Active Directory-based Activation or KMS activation to MAK activation at any time. For information, see Converting KMS Clients to MAK Activation.

There are two ways to activate computers by using MAK:

- **MAK Independent**: MAK Independent activation requires that each computer independently connect and activate with Microsoft over the Internet or by telephone. MAK Independent activation is best suited for computers within an organization that do not maintain a connection to the corporate network.
- **MAK Proxy**: MAK Proxy activation enables a centralized activation request on behalf of multiple computers with one connection to Microsoft. MAK Proxy activation is configured by using the Volume Activation Management Tool (VAMT). MAK Proxy activation is appropriate for environments where security concerns can restrict direct access to the Internet or the corporate network. It is also suited for development and test labs that lack this connectivity.

  Note
  VAMT is a standalone application that collects activation requests from several systems then sends them, in bulk, to Microsoft. For more information, see Volume Activation Management Tool (VAMT) Overview.

MAK can be used for individual computers or with an image that can be duplicated or provided for download by using Microsoft deployment solutions. MAK can also be used on a computer that was originally configured to use KMS activation, if that computer's activation is about to or has reached the end of its activation validity interval.

See also

- Step 2: Evaluate Client Connectivity
- Plan for Volume Activation
- Volume Activation Overview
- Test Lab Guide: Demonstrate Volume Activation Services

Step 2: Evaluate Client Connectivity

After you understand some of the requirements that are associated with each volume activation option as described in Step 1: Review and Select Activation Methods, the next step in planning

for Volume Activation Services is to evaluate the types of clients you have and their location and ability to connect to the central network. After you have completed this step, you can proceed to Step 3: Determine Activation Method and Product License Requirements.

Activation scenarios

Each Volume Activation Services method that is available in Windows Server is best suited to a particular type of client and network configuration.

To select the best activation method or methods for your organization, you need to assess your network environment to identify the types of clients you have and how groups of clients connect to the network.

Connectivity to the corporate network, Internet access, and the number of computers that regularly connect to your corporate network are some of the important characteristics to identify. Most medium to large organizations use a combination of activation methods because of varied client connectivity needs.

This section illustrates a few common volume activation scenarios in heterogeneous corporate environments that require more than one activation method. Each scenario has a recommended activation solution, but some environments may have infrastructure or policy requirements that are best suited to a different solution.

Task	Description
2.1. Identify scenario requirements for the central network	Evaluate and address activation planning issues associated with well-connected computers.
2.2. Identify scenario requirements for the branch office network	Evaluate and address activation planning issues associated with computers that are on branch office networks.
2.3. Identify high-security zone scenario requirements	Evaluate and address activation planning issues associated with computers located in high-security zones.
2.4. Identify disconnected computer scenario requirements	Evaluate and address activation planning issues associated with computers that are disconnected from any of the organization's networks.
2.5 Identify test and development lab requirements	Evaluate and address activation planning issues associated with computers in test or development labs.

2.1. Identify scenario requirements for the central network

The central network is characterized by well-connected computers on multiple network segments that also have a connection to the Internet. For these environments, a combination of Active Directory-based Activation and a centralized KMS solution are the recommended activation methods for computers on the central network.

In Figure 1, the central network has a domain controller and KMS host, which are activated directly over the Internet.

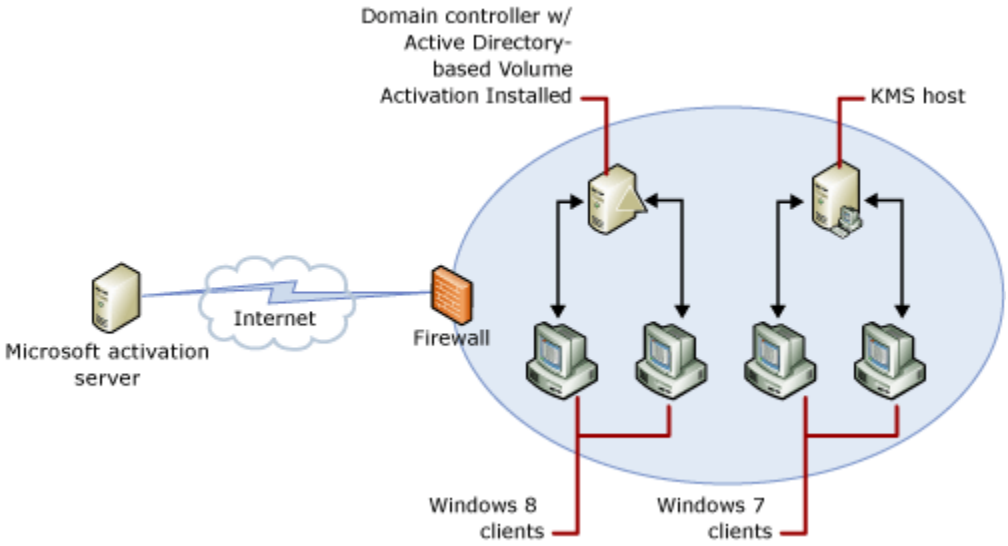

If physical computers are running Windows Server 2012 or Windows 8 or newer, use Active-Directory-based Activation. If you have more than 25 client computers running earlier versions of Windows, use KMS host activation.

Note

Clients running Windows Server 2012 or Windows 8 or newer can activate from the domain controller or from the KMS host.

If you have fewer than 25 client computers running earlier versions of Windows, use MAK activation (by telephone, Internet, or MAK Proxy).

Note

You can install a KMS host on a virtual computer. You should select a virtual computer that is unlikely to be moved to a different physical computer. If the KMS host on the virtual computer is moved to a different physical computer, the operating system detects the change in the underlying hardware, and the KMS host must be reactivated with Microsoft.

If you plan to use KMS activation and you have more than 50 clients, consider setting up at least two KMS hosts for failover if one host is unavailable.

2.2. Identify scenario requirements for the branch office network

Some networks are in a different physical location from the central network. You may need to use a combination of Active Directory-based Activation, KMS hosts, and MAK activation to meet the requirements of the branch office network scenario.

How your network's firewalls are configured is a significant factor in determining your branch office strategy.

If ports on firewalls can be opened between KMS clients and hosts, use KMS hosts that are located on the central network.

Figure 2 shows an enterprise network that supports client computers in three branch offices.

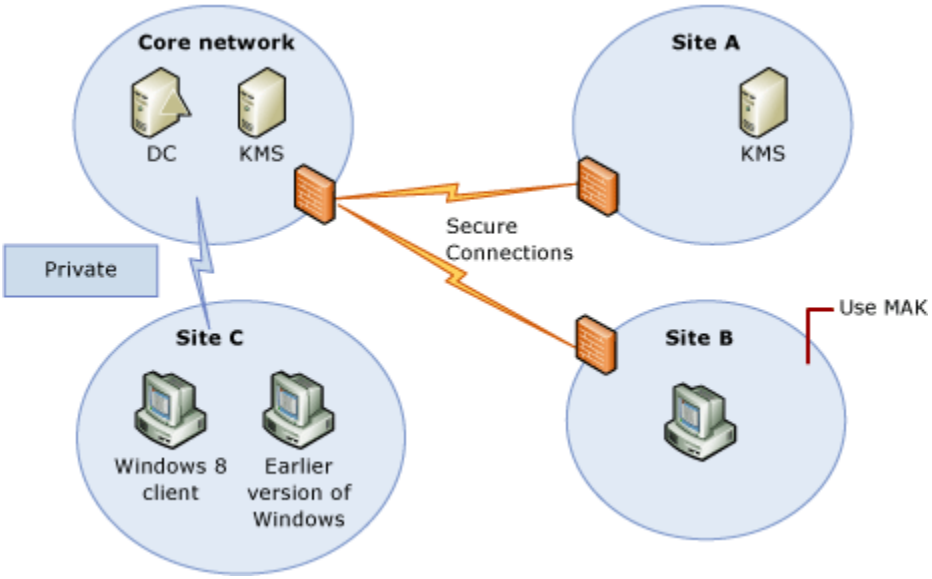

If you have more than 25 client computers and a site that has a secure TCP/IP connection to the central network, use a local KMS host for activation. Site A in the diagram is an example of this solution.

If you have remote sites with infrequent contact to the central network, use MAK activation. Site B is an example of the MAK solution.

If you have a site with a secure connection to the central network, you can use the Active Directory-based Activation for clients running Windows Server 2012 or Windows 8 or newer operating systems, and use KMS hosts from the central network to activate all other clients. Site C is an example of this solution.

Note

The firewall configuration on the host for RPC over TCP should allow inbound traffic on TCP port 1688.

2.3. Identify high-security zone scenario requirements

Many organizations have networks that are separated into multiple security zones, including a high-security zone that is isolated to protect sensitive information.

Figure 3 shows a high-security network environment that does not allow any traffic between computers in the high-security zone and the central network.

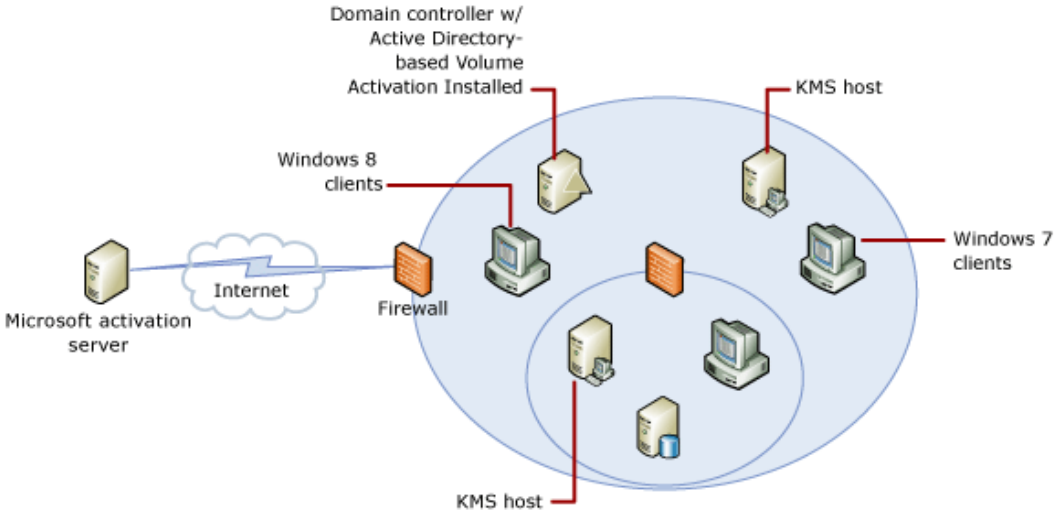

High-security zones are frequently separated by a firewall that limits communication to and from other networks.

If the high-security zone has only limited communication with the central network, but it has enough computers to meet the KMS activation threshold, the high-security zone has a local KMS host. This KMS host is activated by telephone.

If the high-security zone has only limited communication with the central network, but it does not have enough computers to meet the KMS activation threshold, use MAK activation. Each computer can be activated independently with Microsoft by telephone.

If the computers in a high-security zone are allowed access to the central network, you can activate the high-security zone computers by using Active Directory-based Activation or KMS hosts that are located on the central network.

Note

Firewalls between networks need to allow TCP 1688 RPC traffic to the KMS host from the client, and they must allow a stateful response from the KMS host to the client.

MAK Proxy activation enabled with VAMT is also possible in the high-security zone scenario. Because the computers in the high-security zone do not have Internet access, VAMT can discover them by using Active Directory, the computer name, an IP address, or membership in a workgroup. VAMT uses WMI to install MAK product keys and CIDs and to retrieve status on MAK clients. Because this traffic is not allowed through the firewall, you must have a local VAMT host in the high-security zone. For more information, see Volume Activation Management Tool (VAMT) Overview.

2.4. Identify disconnected computer scenario requirements

Some users in an organization may be in remote locations or may travel to many locations. This scenario is common for roaming clients, such as the computers of sales people or other users who work offsite, but not at branch locations. This scenario can also apply to remote branch office locations that have no connection or an intermittent connection to the central network.

Figure 4 illustrates the activation options that are available for disconnected computers.

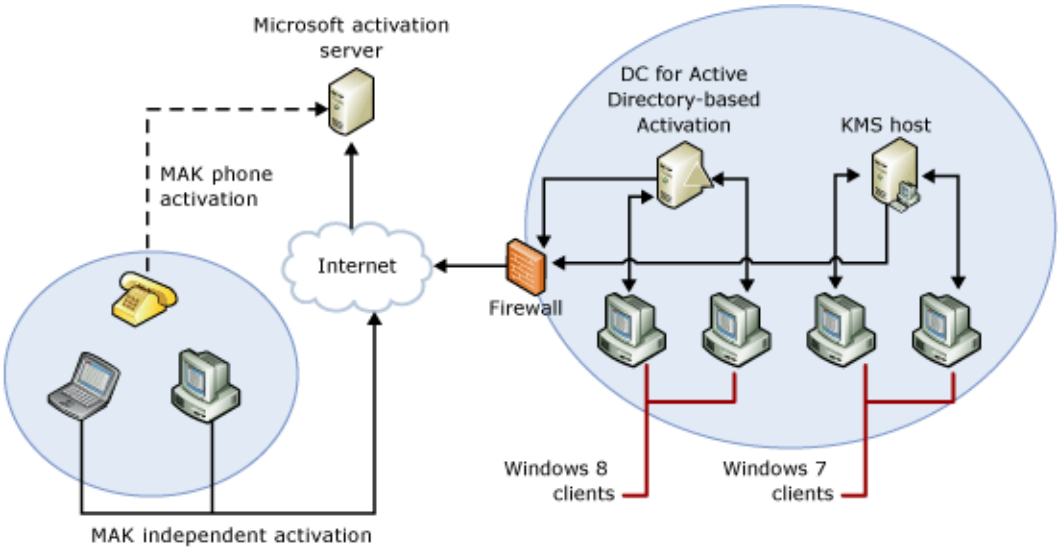

Disconnected computers can use Active Directory-based Activation, KMS activation, or MAK activation processes that connect to the central network (directly or through a VPN or secure channel) at least once every 180 days.

Use MAK Independent activation, by telephone or the Internet, for computers that rarely or never connect to the central network.

2.5 Identify test and development lab requirements

In a lab environment, computers are reconfigured often, and there are usually a large number of virtual computers. You should first determine whether the computers in test and development labs need activation.

If you do need activation, labs can use Active Directory-based Activation, KMS activation, or MAK activation, depending on the planning considerations described previously.

Use Active Directory-based Activation or KMS activation if the computers have connectivity to a central network where these activation options are available.

If the lab does not have connectivity to the central network and the number of physical computers in the lab meets the KMS activation threshold, you can also deploy a local KMS host in the lab. However, if your lab has a high turnover of computers and a small number of physical KMS clients, you need to monitor the KMS activation count to maintain a sufficient number of cached client machine identifications (CMIDs) on the KMS host. A KMS host caches activation requests from physical computers for 30 days.

If computers in your lab environment need activation, but they are not connected to the central network and do not meet the conditions for KMS activation, you can use MAK Independent activation, by telephone or Internet, if available.

MAK Proxy activation can also be used with VAMT in this scenario.

Note

If you install a local VAMT inside a lab that has no outside connectivity, you must manually update VAMT. You need to install VAMT in the isolated lab network and also in a network that has access to the Internet. VAMT in the isolated lab performs discovery, obtains status, installs a MAK product key, and obtains the IID of each computer in the lab. You can then export this information from VAMT, save it to removable media, and import the file to a computer running VAMT that has access to the Internet. VAMT sends the IIDs to Microsoft and obtains the corresponding CIDs needed to complete activation. After you export this data to removable media, you can take it to the isolated lab to import the CIDs so VAMT can complete the activations.

See also

- Step 3: Determine Activation Method and Product License Requirements
- Step 1: Review and Select Activation Methods
- Volume Activation Overview
- Volume Activation Overview

Step 3: Determine Activation Method and Product License Requirements

After selecting the activation methods that are appropriate for your organization in Step 2: Evaluate Client Connectivity, the next step is mapping computers that are using volume activation to activation methods. After you complete this step, you can complete your volume activation deployment plan by completing the steps in Step 4: Determine Monitoring and Reporting Needs.

You can use the following planning table to ensure that all computers are associated with an activation option.

Criteria	Activation Method	Number of Computers
Number of computers running Windows Server 2012 or Windows 8 or newer that will connect to the network at least once every 180 days (directly or through a VPN)	Active Directory-based Activation	

Criteria	Activation Method	Number of Computers
Number of computers not running Windows Server 2012 or Windows 8 or newer that will connect to the network at least once every 180 days (directly or through a VPN), and where the KMS activation threshold is met	KMS	
Number of computers that do not connect to network at least once every 180 days	MAK	
Number of computers running Windows Server 2012 or Windows 8 or newer in isolated networks that connect at least once every 180 days to the core network	Active Directory-based Activation or KMS	
Number of computers not running Windows Server 2012 or Windows 8 or newer in isolated networks that can contact a KMS host	KMS	
Number of computers in isolated networks where the KMS activation threshold is not met	MAK	
Number of computers in test and development labs that will not be activated	None	

See also

- Step 4: Determine Monitoring and Reporting Needs
- Step 2: Evaluate Client Connectivity
- Plan for Volume Activation
- Test Lab Guide: Demonstrate Volume Activation Services

Step 4: Determine Monitoring and Reporting Needs

After you have mapped computers that are using volume activation methods in Step 3: Determine Activation Method and Product License Requirements, you can complete your volume activation deployment plan by determining your organization's volume activation monitoring and reporting needs. When you have completed the planning steps in this section, you can deploy your volume activation infrastructure based on this plan.

Organizations that use the Volume Activation Services in Windows Server 2012need to track product key usage and the license conditions of activated computers. One tool that is available to volume licensing customers is the Volume Licensing Service Center. Volume licensing customers can log on to this website at any time to view their volume licensing key information. Several additional tools are available to help volume licensing customers manage activations and product key usage. This topic describes the available tools and how each assists volume licensing customers. The following sections list some initial considerations to review, in addition to hardware and software requirements for monitoring and reporting on volume activation.

The following table describes the leading management and reporting options for the Volume Activation services in versions starting with Windows Server 2012.

Task	Description
4.1 Managing Active Directory-based Activation objects	Use Active Directory management tools to monitor and manage activation issues.
4.2 Managing Activation with System Center Configuration Manager	Use System Center Configuration Manager to monitor and manage activation issues.
4.3 Managing Activation by Using Event Logs	Use event logs to monitor and manage activation issues.
4.4 Managing Activation by Using the Volume Activation Management Tool (VAMT)	Use VAMT to monitor and manage activation issues.

4.1 Managing Active Directory-based Activation objects

Active Directory-based Activation objects are held in the configuration partition of Active Directory. You can query the information about Active Directory-based Activation objects by using ADSIEDIT.msc or the SLMGR.vbs command-line interface.

ADSIEDIT allows members of the Enterprise Administrators group to change properties on the object or delete the object if they desire. Domain members have Read access. SLMGR allows an administrator to delete the object but not change its attributes.

ADSIEDIT or SLMGR allow administrators to gain other information about the Active Directory objects in their domain.

The following are a list of properties that can be reviewed by using ADSIEDIT:

- Partial product key
- KMS ProductID
- Active Directory object creation date
- Access control lists (ACLs)

SLMGR provides the following information about the Active Directory object:

- License state information, including activation status
- Ability to force immediate activation attempt
- Disable or force Active Directory-based Activation
- Install MAK product key (this disables Active Directory-based Activation)
- List all installed Active Directory objects (only available for accounts with rights to the Active Directory object container)
- Delete an Active Directory object (only available for accounts with rights to the Active Directory object container)

4.2 Managing Activation with System Center Configuration Manager

Data that is gathered during activation is accessible by using WMI. Several of the tools available use WMI to access volume activation data.

See Appendix 1 of the Volume Activation 2.0 Operations Guide for a list all of WMI methods, properties, registry keys, and event IDs for product activation. MAK and KMS can use System Center Configuration Manager to monitor the license conditions of their organization's computers. For a detailed description of the available license conditions, see Appendix 2: License Conditions in this guide.

4.3 Managing Activation by Using Event Logs

The Volume Activation service and KMS service record every action in the application logs of the client and host computers. A client records activation requests, renewals, and responses in the client's local application log by using Microsoft Windows Security Licensing (SLC) event numbers 12288 and 12289. The KMS host and domain controller log a separate entry for each request they receive from a client as SLC event number 12290. These entries are saved to the **Applications and Services Logs** folder. If you have multiple KMS hosts, each host keeps an individual log of activations. There is no replication of logs between KMS hosts.

4.4 Managing Activation by Using the Volume Activation Management Tool

Each MAK key has a predetermined number of allowed activations, based on an organization's volume licensing agreement. Every activation with activation services that are hosted by Microsoft reduces the MAK activation pool by one. MAK implementations should include how your organization plans to monitor the number of MAK activations that are left. MAK Independent and MAK Proxy activations can use VAMT to monitor this.

See also

- Step 3: Determine Activation Method and Product License Requirements
- Plan for Volume Activation
- Volume Activation Overview
- Volume Activation Management Tool (VAMT) Overview
- Test Lab Guide: Demonstrate Volume Activation Services

Appendix A: KMS Client Setup Keys

Computers that are running volume licensing editions of Windows 8, Windows Server 2012 R2, Windows 8, Windows Server 2012, Windows 7, Windows Server 2008 R2, Windows Vista, and Windows Server 2008 are, by default, KMS clients with no additional configuration needed.

If you are converting a computer from a KMS host, MAK, or retail edition of Windows to a KMS client, install the applicable setup key (GVLK) from the following tables.

Note
To install a client setup key, open an administrative command prompt on the client, type **slmgr /ipk** *<setup key>* and press ENTER.

Windows Server 2012 R2 and Windows 8.1 Client Setup Keys

Operating system edition	KMS Client Setup Key
Windows 8.1 Professional	GCRJD-8NW9H-F2CDX-CCM8D-9D6T9
Windows 8.1 Professional N	HMCNV-VVBFX-7HMBH-CTY9B-B4FXY
Windows 8.1 Enterprise	MHF9N-XY6XB-WVXMC-BTDCT-MKKG7
Windows 8.1 Enterprise N	TT4HM-HN7YT-62K67-RGRQJ-JFFXW

Operating system edition	KMS Client Setup Key
Windows Server 2012 R2 Server Standard	D2N9P-3P6X9-2R39C-7RTCD-MDVJX
Windows Server 2012 R2 Datacenter	W3GGN-FT8W3-Y4M27-J84CP-Q3VJ9
Windows Server 2012 R2 Essentials	KNC87-3J2TX-XB4WP-VCPJV-M4FWM

Windows Server 2012 and Windows 8 Client Setup Keys

Operating system edition	KMS Client Setup Key
Windows 8 Professional	NG4HW-VH26C-733KW-K6F98-J8CK4
Windows 8 Professional N	XCVCF-2NXM9-723PB-MHCB7-2RYQQ
Windows 8 Enterprise	32JNW-9KQ84-P47T8-D8GGY-CWCK7
Windows 8 Enterprise N	JMNMF-RHW7P-DMY6X-RF3DR-X2BQT
Windows Server 2012	BN3D2-R7TKB-3YPBD-8DRP2-27GG4
Windows Server 2012 N	8N2M2-HWPGY-7PGT9-HGDD8-GVGGY
Windows Server 2012 Single Language	2WN2H-YGCQR-KFX6K-CD6TF-84YXQ
Windows Server 2012 Country Specific	4K36P-JN4VD-GDC6V-KDT89-DYFKP
Windows Server 2012 Server Standard	XC9B7-NBPP2-83J2H-RHMBY-92BT4
Windows Server 2012 MultiPoint Standard	HM7DN-YVMH3-46JC3-XYTG7-CYQJJ
Windows Server 2012 MultiPoint Premium	XNH6W-2V9GX-RGJ4K-Y8X6F-QGJ2G
Windows Server 2012 Datacenter	48HP8-DN98B-MYWDG-T2DCC-8W83P

Windows 7 and Windows Server 2008 R2

Operating system edition	KMS Client Setup Key
Windows 7 Professional	FJ82H-XT6CR-J8D7P-XQJJ2-GPDD4
Windows 7 Professional N	MRPKT-YTG23-K7D7T-X2JMM-QY7MG
Windows 7 Professional E	W82YF-2Q76Y-63HXB-FGJG9-GF7QX
Windows 7 Enterprise	33PXH-7Y6KF-2VJC9-XBBR8-HVTHH

Operating system edition	KMS Client Setup Key
Windows 7 Enterprise N	YDRBP-3D83W-TY26F-D46B2-XCKRJ
Windows 7 Enterprise E	C29WB-22CC8-VJ326-GHFJW-H9DH4
Windows Server 2008 R2 Web	6TPJF-RBVHG-WBW2R-86QPH-6RTM4
Windows Server 2008 R2 HPC edition	TT8MH-CG224-D3D7Q-498W2-9QCTX
Windows Server 2008 R2 Standard	YC6KT-GKW9T-YTKYR-T4X34-R7VHC
Windows Server 2008 R2 Enterprise	489J6-VHDMP-X63PK-3K798-CPX3Y
Windows Server 2008 R2 Datacenter	74YFP-3QFB3-KQT8W-PMXWJ-7M648
Windows Server 2008 R2 for Itanium-based Systems	GT63C-RJFQ3-4GMB6-BRFB9-CB83V

Windows Vista and Windows Server 2008

Operating system edition	KMS Client Setup Key
Windows Vista Business	YFKBB-PQJJV-G996G-VWGXY-2V3X8
Windows Vista Business N	HMBQG-8H2RH-C77VX-27R82-VMQBT
Windows Vista Enterprise	VKK3X-68KWM-X2YGT-QR4M6-4BWMV
Windows Vista Enterprise N	VTC42-BM838-43QHV-84HX6-XJXKV
Windows Web Server 2008	WYR28-R7TFJ-3X2YQ-YCY4H-M249D
Windows Server 2008 Standard	TM24T-X9RMF-VWXK6-X8JC9-BFGM2
Windows Server 2008 Standard without Hyper-V	W7VD6-7JFBR-RX26B-YKQ3Y-6FFFJ
Windows Server 2008 Enterprise	YQGMW-MPWTJ-34KDK-48M3W-X4Q6V
Windows Server 2008 Enterprise without Hyper-V	39BXF-X8Q23-P2WWT-38T2F-G3FPG
Windows Server 2008 HPC	RCTX3-KWVHP-BR6TB-RB6DM-6X7HP
Windows Server 2008 Datacenter	7M67G-PC374-GR742-YH8V4-TCBY3
Windows Server 2008 Datacenter without Hyper-V	22XQ2-VRXRG-P8D42-K34TD-G3QQC
Windows Server 2008 for Itanium-Based Systems	4DWFP-JF3DJ-B7DTH-78FJB-PDRHK

Deploy Volume Activation

This guide describes Microsoft® Volume Activation deployment concepts. Volume Activation consists of three technologies—Active Directory-based Activation, Key Management Service (KMS), and Multiple Activation Key (MAK)—that allow Volume Licensing customers to activate Volume License editions of Windows operating systems.

Many organizations will want to deploy more than one or perhaps all three of these volume activation technologies. Before deploying Volume Activation Services, complete the planning steps in Plan for Volume Activation to determine which volume activation options are appropriate for your organization.

Note
This guide describes procedures that run scripts and make changes to the registry. These rights can be delegated to selected information technology (IT) implementers, and the rights to change product keys and perform activations can even be assigned to users, although Microsoft does not recommend this practice.

Methods of volume activation

Deploy Active-Directory-based Activation

Deploy KMS Activation

Deploy MAK Activation

Deploy Active-Directory-based Activation

Starting with Windows Server® 2012 and Windows® 8, you can avoid designating and configuring an activation server altogether and use Active Directory for activation instead. Whenever you activate a computer running these or later operating systems that is joined to a domain that is using the Windows 8 (or later) schema, an activation object is created and stored in Active Directory, instead of on the local computer. You can accomplish the initial activation of the first computer with Volume Activation Services in Server Manager or Slmgr.vbs Active Directory activation commands (see Slmgr.vbs Options for Volume Activation).

When any Windows Server 2012, Windows 8 (or later operating system) computers that are volume licensed join the domain, they access the activation object present in Active Directory and become activated automatically. They remain activated as long as they remain joined to the domain and in contact with it. At each renewal interval, any Key Management Service (KMS) clients, upon connection to the domain, will attempt to use Active Directory activation; if that is not available, they will continue to use KMS activation.

Using Active Directory-based activation

To use Active Directory-based activation, access any computer joined to the domain that is running Windows Server 2012, Windows 8 (or later) operating system and activate the computer with a volume product key (CSVLK) by using the UI or other activation methods.

Managing the activation object

You can use a one-time script (Appendix A in this topic) that runs at the Enterprise Administrator level, creates a Security Group, and designates an owner for it. Members that you add to that Security Group are authorized to create activation objects, allowing you to delegate licensing support. You cannot edit the activation objects, but you can use Volume Activation Services to control and view them.

To prepare Active Directory

1. If the forest is not already based on the Windows Server 2012 schema, run the Windows Server 2012 (or later) version of Adprep.exe to update the schema.
2. Optionally, using an account with administrative access to the Domain Root Directory (such as Enterprise Administrator), run the script available in Appendix A of this topic. This creates the Licensing Security Group.
3. Optionally, the Enterprise Administrator can add accounts to the Licensing Security Group to delegate the ability to store activation objects.

To activate the Active Directory forest

1. Obtain a customer-specific volume license key (CSVLK).
2. Log into any computer joined to the domain that is running Windows Server 2012, Windows 8 (or later operating system) with an account that is in the Licensing Security Group.
3. In the Volume Activation Tools of the Volume Activation service, start the CSVLK Installation wizard and complete the steps to install the CSVLK, choose **Activate AD**, and (optionally) create a friendly name for the activation object.

   Note
   If online activation is not possible, you can use phone activation to obtain a CID from the CSVLK. Then enter the CID in the CSVLK Installation wizard.
4. Alternately, instead of using the wizard, you can install the CSVLK with the command **slmgr /ad-activation-online <Product Key> [Activation Object name]**. (See Slmgr.vbs Options for Volume Activation.) You can also install the CSVLK with the phone activation option.

To activate Windows on other computers in the domain

Once the activation object has been created as in the previous procedure, any computer running Windows Server 2012, Windows 8 (or later operating system) that joins the domain becomes activated automatically.

To manage the activation object

1. From an account in the Licensing Security Group, start ADSI Edit.
2. Using ADSI Edit, you can view the following properties of the activation object:
   - Partial product key
   - The KMS product ID
   - Supported editions
   - Creation date of the activation object
   - ACLs applied to the activation object
3. Alternately, from the command line, you can install, view, or delete activation objects with the Active Directory-based activation configuration options of Slmgr.vbs. See Slmgr.vbs Options for Volume Activation.
4. To delete activation objects, use Volume Activation Services or Slmgr.vbs.

Appendix A: Administration delegation script

This script runs at the Enterprise Administrator level, creates the Licensing Security Group, and designates an owner for it. You can use it to delegate accounts the authority to view and manage the activation objects.

```
# Delegate activation script

# user_name


#....................................

# relax execution policy

#....................................

set-executionpolicy Unrestricted


#....................................

# import AD module

#....................................

import-module ActiveDirectory


#....................................
```

```
# Query domain Name
#.....................................
$domain = get-addomain

#.....................................
# Get the licensing admin
#.....................................
$Licensing_admin = Get-ADUser $args[0]
$licensing_admin_sid = New-Object System.Security.Principal.SecurityIdentifier
$Licensing_admin.SID
$licensing_admin_dn = $licensing_admin.DistinguishedName

#.....................................
# Get domain DN
#.....................................
$dn=$domain.DistinguishedName

#.....................................................................
# Create a new AD-Group: "ms-SPP-Licensing-Administrators" and let admin manage it
#.....................................................................
New-ADGroup -Name "MS SPP Licensing Administrators" -SamAccountName msSPPLicensing -
GroupCategory Security -GroupScope Universal -DisplayName "MS SPP Licensing
Administrators"
#Could possibly add error checking and warning if groups wasn't created and bail.

$licensing_group = Get-ADGroup -Identity msSPPLicensing
$licensing_group_sid = New-Object System.Security.Principal.SecurityIdentifier
$licensing_group.SID
$licensing_group_dn = $licensing_group.DistinguishedName

$grp_acl = Get-Acl -Path:("AD:\" + $licensing_group_dn)

#create ACE for write property (members)
$guid = New-Object System.Guid "bf9679c0-0de6-11d0-a285-00aa003049e2"
```

```
$grp_ace = New-Object System.DirectoryServices.ActiveDirectoryAccessRule
$licensing_admin_sid, "WriteProperty", "Allow", $guid
$grp_acl.AddAccessRule($grp_ace)

#create ACE for read property (members)
$grp_ace = New-Object System.DirectoryServices.ActiveDirectoryAccessRule
$licensing_admin_sid, "ReadProperty", "Allow", $guid
$grp_acl.AddAccessRule($grp_ace)

Set-Acl -ACLObject:$grp_acl -Path:("AD:\" + $licensing_group_dn)

#.....................................................................
# Get Activation container and make the licensing admin group in charge
#.....................................................................

# Path to the Activation Container
$Path_A="CN=Activation Objects,CN=Microsoft SPP,CN=Services,CN=Configuration,"+$dn
#Could possibly add error checking and warning if container wasn't found.

$grp_acl = Get-Acl -Path:("AD:\" + $Path_A)

#create ACE for FC
$grp_ace = New-Object System.DirectoryServices.ActiveDirectoryAccessRule
$licensing_group_sid, "GenericAll", "Allow"
$grp_acl.AddAccessRule($grp_ace)

Set-Acl -ACLObject:$grp_acl -Path:("AD:\" + $Path_A)

#.....................................................................
# Add the licensing admin to the "ms-SPP-Licensing-Administrators" group
#.....................................................................

Add-ADGroupMember $licensing_group -Members $licensing_admin_dn
```

Deploy KMS Activation

If your network environment supports the DNS dynamic update protocol and allows computers to publish services automatically, deploying a KMS host can require very little effort. If the organization has more than one KMS host or the network does not support dynamic update, additional configuration tasks may be necessary.

Caution
Some procedures in this section require changing the registry. Problems can occur if the registry is modified incorrectly by using Registry Editor or another method, and these problems might require reinstalling the operating system. Microsoft cannot guarantee that these problems can be resolved. IT pros modify the registry at their own risk.

The remainder of this section describes the following key tasks:

- **Configuring KMS hosts**
- **Configuring DNS**
- **Installing KMS hosts**
- **Configuring KMS clients**

Configuring KMS hosts

See Configuring KMS hosts.

Configuring DNS

The following sections describe concepts for configuring DNS to work with Volume Activation:

If more than one KMS host is used, see the section “Change the Default DNS Permissions for SRV Records.”

To enable KMS clients using different DNS servers to find KMS hosts, see the section “Publish to Multiple DNS Domains.”

To manually add SRV resource records for KMS hosts, see the sections “Manually Create SRV Records in DNS,” “Manually Create SRV Records in a BIND 8.2 or Higher DNS Server,” and “Disable Publishing of KMS SRV Records to DNS.”

Note
DNS changes might not be reflected until all DNS servers have been replicated.

Change the Default DNS Permissions for SRV Records

If you are using only one KMS host, you might not need to configure permissions in DNS. The default behavior is to allow a computer to create an SRV resource record and then update it. However, if you have more than one KMS host (the usual case), the other hosts will be unable to update the SRV resource record unless SRV default permissions are changed.

The following high-level procedure is an example from Microsoft's own environment. It does not give detailed steps, which might be different from one organization to another, and it is not the only way to achieve the desired result

Create a global security group in Active Directory® that will be used for your KMS hosts. An example is Key Management Service Group.Add each of your KMS hosts to this group. They must all be joined to the same domain.

Once the first KMS host is created, it will create the original SRV record. If the first KMS host is unable to create the SRV resource record, it may be because your organization has changed the default permissions. In this case, manually create the SRV resource record as the section "Manually Create SRV Records in DNS" describes.

Set the permissions for the SRV group to allow updates by members of the global security group.

Note
A domain administrator can delegate the ability to carry out the preceding steps to administrators in the organization. To do so, create a security group in Active Directory, give that group permission to change the SRV records, and then add the delegates.

Publish to Multiple DNS Domains

By default, the KMS host is registered only in the DNS domain to which the host belongs. If the network environment has only one DNS domain, no further action is required.

If there is more than one DNS domain name, a list of DNS domains can be created for a KMS host to use when publishing its SRV RR. Setting this registry value suspends the KMS host's default behavior of publishing only in the domain specified as the Primary DNS Suffix.Optionally, add priority and weight parameters to the **DnsDomainPublishList** registry value for KMS. This feature enables an administrator to establish KMS host priority groupings and weighting within each group to define which KMS host to try first and balance traffic among multiple KMS hosts.

Note
DNS changes might not be reflected until all DNS servers have been replicated. Changes made too frequently (time < replication time) can leave older records if the change is performed on a server that has not been replicated.

To automatically publish KMS in multiple DNS domains, add each DNS domain suffix to whichever KMS should publish to the multi-string registry value **DnsDomainPublishList** in **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform**. After changing the value, restart the Software Licensing Service to create the SRV RRs.

Note
This key has changed from the Windows Vista location of **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SL**

After configuring a KMS host to publish to multiple domains, export the registry key, and then import it in to the registry on additional KMS hosts. To verify that this procedure was successful,

check the Application event log on each KMS host. Event ID 12294 indicates that the KMS host successfully created the SRV RRs. Event ID 12293 indicates that the attempt to create the SRV RRs was unsuccessful. For a complete list of error codes, see the Volume Activation Operations Guide.

Manually Create SRV Records in DNS

If the environment does not support dynamic update, the SRV RRs must be manually created to publish the KMS host. Environments that do not support dynamic update should disable publishing on all KMS hosts to prevent event logs from collecting failed DNS publishing events. To disable auto-publishing, use the Slmgr.vbs script with the **/cdns** command-line option. See the "Configuring KMS" section for more information about the Slmgr.vbs script.

Note

Manually created SRV RRs can coexist with SRV RRs that KMS hosts automatically publish in other domains as long as all records are maintained to prevent conflicts.

Using DNS Manager, in the appropriate forwarding lookup zone, create a new SRV RR using the appropriate information for the location. By default, KMS listens on TCP port 1688, and the service is _VLMCS. Table 2 contains example settings for a SRV RR.

Name	Setting
Service	_VLMCS
Protocol	_TCP
Port number	1688
Host offering the service	FQDN of the KMS host

Manually Create SRV Records in a BIND 8.2 or higher DNS Server

If the organization uses a non-Microsoft DNS server, the needed SRV RRs can be created as long as the DNS server is compliant with Berkeley Internet Name Domain (BIND) 8.2 or higher. When creating the record, include the information shown in Table 3. The Priority and Weight settings shown in Table 3 are only used by Windows 7 and Windows Server 2008 R2.

Name	Setting
Name	_vlmcs._tcp
Type	SRV
Priority	0
Weight	0
Port	1688

Name	Setting
Host name	FQDN of the KMS host

To configure a BIND 8.2 or higher DNS server to support KMS auto-publishing, configure the BIND server to enable RR updates from KMS hosts. For example, add the following line to the zone definition in named.conf:`allow-update { any; };`

Note

An `allow-update` statement can also be added in named.conf options to allow dynamic update for all zones hosted on this server.

Disable Publishing of KMS SRV Records to DNS

KMS hosts automatically publish their existence by creating SRV RRs in DNS. To disable automatic DNS publishing by a KMS host, use the Slmgr.vbs script with the **/cdns** command-line option.

Using the Slmgr.vbs script to disable automatic DNS publishing is preferred, but you can also perform this task by creating a new DWORD value called **DisableDnsPublishing** in the registry, and set its value to 1. This value is at **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform** in the registry. To re-enable the default behavior for publishing of KMS SRV records to DNS, set the value to 0.

Installing KMS hosts

To enable KMS functionality, a KMS key is installed on a KMS host; then, the host is activated over the Internet or by phone using Microsoft's activation services. Computers running Windows 7 or Windows Server 2008 R2 can both serve as KMS hosts.Windows Vista, Windows Server 2003, and Windows Server 2008 can also serve as KMS hosts. The KMS clients that a KMS host can activate are dependent on the host key used to activate the KMS host. For more information about KMS host keys, see the Volume Activation Planning Guide.

Install and activate a KMS key on a Windows 7 or Windows Server 2008 R2 computer by using an elevated command prompt:

- To install a KMS key, type **slmgr.vbs /ipk <KmsKey>** at a command prompt.
- To activate online, type **slmgr.vbs /ato** at a command prompt.
- To activate by using the telephone, type **slui.exe 4** at a command prompt.

After activating the KMS key, restart the Software Protection Service.

Windows 7 and Windows Server 2008 R2 display the warning shown in [Figure 1] any time administrators install a KMS host key by using the UI (Users will not see this warning if they install a KMS host key by using the Slmgr.vbs script). This message prevents accidentally installing a KMS key on computers that administrators do not intend to be KMS hosts.

To verify that the KMS host is configured correctly, check the KMS count to see whether it is increasing. In the Command Prompt window on the KMS host, type slmgr.vbs /dli to display the

current KMS count. Administrators can also check the Key Management Service log in the Applications and Services Logs folder for event ID 12290. The Key Management Service log records activation requests from KMS clients. Each event displays the name of the computer and the time stamp of each activation request.

Configuring KMS clients

This section describes concepts for installing and configuring computers as KMS clients. By default, Volume License editions of Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 are KMS clients. If the computers the organization wants to activate by using KMS are using either of these operating systems and the network allows DNS auto-discovery, no further configuration is needed.

If a KMS client is configured to search for a KMS host using DNS but does not receive SRV records from DNS, Windows 7 and Windows Server 2008 R2 log the error in the event log.

Manually Specifying a KMS Host

Administrators can manually assign a KMS host to KMS clients by using KMS host caching. Manually assigning a KMS host disables auto-discovery of KMS on the KMS client. A KMS host is manually assigned to a KMS client by running:**/skms <Name[:Port] | : port> [Activation ID]** where <**Name**> is either the KMS_FQDN, IPv4Address, or NetbiosName of the KMS host and port is TCP port on the KMS host.

If the KMS host uses Internet Protocol version 6 (IPv6) only, the address must be specified in the format [hostname]:port (using the square brackets). IPv6 addresses contain colons (:), which will be parsed incorrectly by the Slmgr.vbs script.

Enable Auto-discovery for a KMS Client

By default, KMS clients automatically attempt to discover KMS hosts. Auto-discovery can be disabled by manually assigning a KMS host to a KMS client. This action also clears the KMS host name from the KMS client's cache. If auto-discovery is disabled, it can be re-enabled by running **slmgr.vbs /ckms** at a command prompt.

Adding Suffixed Entries to KMS Clients

By adding the address of a DNS server containing the SRV RR as a suffixed entry on KMS clients, administrators can advertise KMS hosts on one DNS server and allow KMS clients with other primary DNS servers to find it. For more information about configuring a domain suffix search list on KMS clients, see How to configure a domain suffix search list on the Domain Name System clients.

Deploy KMS Clients

The information in this section is for Volume Licensing customers using the Windows Automated Installation Kit (Windows AIK) to deploy and activate a Windows operating system. Prepare KMS clients for deployment by using the System Preparation Tool (Sysprep.exe).

Before capturing an image, run Sysprep.exe with the **/generalize** command-line option to reset the activation timer, security identifier (SID), and other important settings. Resetting the activation timer prevents the image's grace period from expiring before the image is deployed. Running

Sysprep.exe does not remove the installed product key, and administrators are not prompted for a new key during mini-setup. If no rearms are left, the Sysprep.exe operation completes but the activation timers are not changed and an error is returned that explains the situation.

When building demo virtual machines for internal use (e.g., building virtual machines for the organization's sales department or to set up a temporary training environment), running the Slmgr.vbs script with the **/rearm** command-line option extends the grace period another 30 days, which in turn resets the activation timer but makes no other changes to the computer. The activation timer can be reset three times for computers running Windows 7 or Windows Server 2008 R2.

Manually Activate a KMS Client

By default, KMS clients automatically attempt to activate themselves at preset intervals. To manually activate KMS clients (for example, disconnected clients) before distributing them to users, use the Control Panel System item, or run **slmgr.vbs /ato** at an elevated command prompt. The Slmgr.vbs script reports activation success or failure and provides a result code. To perform activation, the KMS client must have access to a KMS host on the organization's network.

Deploy MAK Activation

MAK activation is used for one-time activation through Microsoft's hosted activation services, with no renewals required. For background information about MAK activation, see the Volume Activation Planning Guide.

Installing a MAK During Operating System Installation

Administrators can convert a KMS client to a MAK client during the initial installation of Windows 7 or Windows Server 2008 R2 by including a MAK in an unattended setup (Unattend.xml) file. The Unattend.xml file can be used with Setup.exe or Windows Deployment Services. For more information, see the Unattended Windows Setup Reference help file in the Windows AIK at http://go.microsoft.com/fwlink/?LinkId=136976. For a sample unattended installation file, see "Appendix B: Sample Unattended Installation File."Note The MAK is stored in clear text in the Unattend.xml file. During an unattended installation, the file Unattend.xml or AutoUnattend.xml is copied to the %SystemRoot%\Panther folder of the target computer. However, at the end of the Setup process, the Setup program replaces it with "SENSITIVE*DATA*DELETED."

Installing a MAK After Operating System Installation

A volume edition of Windows 7 and Windows Server 2008 R2 can be configured to use MAK activation by using the Control Panel System item or by running the Slmgr.vbs script:

To install a MAK by using the System application, click the Change your product key link, and then type the MAK in the Change your product key for activation dialog box.

To install a MAK by using Slmgr.vbs, run the following command at a command prompt:**slmgr.vbs /ipk <MultipleActivationKey>**

where MultipleActivationKey is the MAK.

If users install a MAK by using the UI, the MAK client attempts to activate itself over the Internet one time. If the users install a MAK key by using the Slmgr.vbs script, the MAK client does not try to activate automatically.

Disabling Automatic Activation

To disable automatic activation on any MAK client, set the DWORD registry value **Manual** to 1. This value is in the registry subkey **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\Activation**.

Activating MAK Clients

The MAK client attempts to activate over the Internet at the next scheduled interval. Administrators can force immediate activation over the Internet, by telephone, or by using the Volume Activation Management Tool (VAMT).

To confirm activation, check the system tray for a notification that says, “Windows is activated.” Alternatively, type slmgr.vbs /dli at a command prompt to view the activation status of a computer.

Activating MAK Clients over the Internet

A MAK client can be activated over the Internet in either of two ways:

- Click the **Click here to activate Windows now** link in the Control Panel System item. Windows reports whether the activation was successful. If activation was unsuccessful, a wizard presents additional options.
- Run **slmgr.vbs /ato** at a command prompt. Additional options are not presented when using slmgr.vbs.

Activating MAK Clients Through a Proxy Server

Activation over the Internet may be blocked if the proxy server requires user authentication. In Microsoft Internet Security and Acceleration (ISA) Server, this setting is called Basic Authentication. Because activation requests do not present the user's credentials to the proxy server, Microsoft recommends not using Basic Authentication with ISA Server or other proxy servers. However, if Basic Authentication or a comparable mechanism must be used on the proxy server, add the following URLs to the Proxy Authentication exclusion list:

http://go.microsoft.com/*

https://sls.microsoft.com/*

https://sls.microsoft.com:443

http://crl.microsoft.com/pki/crl/products/MicrosoftRootAuthority.crl

http://crl.microsoft.com/pki/crl/products/MicrosoftProductSecureCommunications.crl

http://www.microsoft.com/pki/crl/products/MicrosoftProductSecureCommunications.crl

http://crl.microsoft.com/pki/crl/products/MicrosoftProductSecureServer.crl

http://www.microsoft.com/pki/crl/products/MicrosoftProductSecureServer.crl

Activating MAK Clients Using the Telephone

To activate computers that are connected to the organization's network but do not have Internet connectivity by using Slmgr.vbs. In the Command Prompt window, type:

slmgr.vbs TargetComputerName <Username> <Password> /dti

to display the information required to complete telephone activation. To obtain the telephone number for an Activation Call Center in your local region, run slui.exe 4. Use the Interactive Voice Response system to obtain the confirmation ID (CID); then, run:

slmgr.vbs TargetComputerName <UserName> <Password> /atp <CID>

to install the CID. If activation is performed frequently or multiple computers are activated, it may be more useful to automate the process using the Slmgr.vbs script.

Activating MAK Clients Using VAMT

VAMT allows automation of MAK deployment and activation over the network by distributing MAKs from a centralized console. VAMT queries Microsoft activation servers to get the number of remaining activations for a given MAK, then lists the activation status of all MAK-activated systems in the environment. This count is a snapshot in time, not a real-time count. VAMT version 3.1 is included in the Windows ADK. For more information about VAMT, see Volume Activation Management Tool (VAMT) Technical Reference, and VAMT How-To Topics.

Disable Automatic Activation

Automatic activation can be disabled on any KMS client by setting the existing DWORD value **Manual** to 1. This value is in the registry subkey **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\Activation**.

Integrating MAKs with Deployment Workbench

Microsoft Deployment Toolkit (MDT) also provides a solution for deploying MAKs. In Deployment Workbench, administrators configure the MAK in task sequences, which add the MAK to the Unattend.xml file used during installation. Administrators can prepare the reference image for KMS activation, then, during deployment, MDT activates the installation by using a MAK as long as it does not detect a KMS infrastructure. MDT applies the MAK after installing the image. For more information about MDT, see the Deployment TechCenter.

Volume Activation Technical Reference

These topics provide useful reference information related to Volume Activation.

Technical Reference for Volume Activation

- Slmgr.vbs Options for Volume Activation

- Registry Settings for Volume Activation
- KMS Activation Timing and Discovery for Volume Activation
- KMS and MAK Activation Scenarios for Volume Activation
- KMS Host Keys to Products Activated for Volume Activation
- Activation Policy Values for Volume Activation
- KMS Client Setup Keys for Volume Activation
- Activation Error Codes for Volume Activation
- WMI Properties and Methods for Volume Activation

See also

- Volume Activation Overview

Slmgr.vbs Options for Volume Activation

Because of WMI changes in Windows 7 and Windows Server 2008 R2, the Slmgr.vbs script is not intended to work across platforms. Using Slmgr.vbs to manage a Windows 7 or Windows Server 2008 R2 system from the Windows Vista® operating system is not supported. Attempting to manage an older system from Windows 7 or Windows Server 2008 R2 will generate a specific version mismatch error. For example, running cscript slmgr.vbs <vista_machine_name> /dlv produces the following output:

```
Microsoft (R) Windows Script Host Version 5.8

Copyright (C) Microsoft Corporation. All rights reserved.


The remote machine does not support this version of SLMgr.vbs
```

To manage remote clients, use the Volume Activation Management Tool (VAMT) version 1.2 or later, or create custom WMI scripts that are aware of the differences between platforms. For more information about WMI properties and methods for Volume Activation, see the section, “WMI Properties and Methods,” later in this reference guide.

The following describes the syntax of the Slmgr.vbs script, and Table 2 describes each command-line option:

```
slmgr.vbs [MachineName [User Password]] [<Option>]
```

Note
For support of other software using Volume Activation, see the documents specifically written for those applications.

Table 2. Slmgr.vbs command-line options

Option	Description
[MachineName]	Name of a remote computer (default is local computer)
[User]	Account with the required privilege on the remote computer
[Password]	Password for the account with required privileges on the remote computer
Global options	
/ipk <ProductKey>	Attempts to install a 5×5 product key. The product key provided by the parameter is confirmed valid and applicable to the installed operating system. If not, an error is returned. If the key is valid and applicable, the key is installed. If a key is already installed, it is silently replaced. To prevent instability in the license service, the system should be restarted or the Software Protection Service should be restarted. This operation must be run from an elevated command prompt, or the Standard User Operations registry value must be set to allow unprivileged users extra access to the Software Protection Service.
/ato [Activation ID]	For retail editions and volume systems with a KMS host key or a Multiple Activation Key (MAK) installed, **/ato** prompts Windows to attempt online activation. For systems with a Generic Volume License Key (GVLK) installed, this prompts an attempt at KMS activation. Systems that have been set to suspend automatic KMS activation attempts (/stao) still attempt KMS activation when /ato is run. **Note** Starting with Windows 8 (and Windows Server 2012), the /stao option is deprecated. Use the /act-type option instead. The parameter [Activation ID] expands /ato support to identify a Windows edition installed on the computer. Specifying the [Activation ID] parameter isolates the effects of the option to the edition associated with that Activation ID. Run Slmgr.vbs /dlv all to get the Activation IDs for the

Option	Description
	installed version of Windows. If you need to support other applications, see the guidance provided by that application for further instruction. KMS activation does not require elevated privileges. However, online activation does require elevation, or the Standard User Operations registry value must be set to allow unprivileged users extra access to the Software Protection Service.
/dli [Activation ID | All]	Display license information. By default, **/dli** displays the license information for the installed active Windows edition. Specifying the **[Activation ID]** parameter displays the license information for the specified edition associated with that Activation ID. Specifying **[All]** as the parameter will display all applicable installed products' license information. This operation does not require elevated privileges.
/dlv [Activation ID | All]	Display detailed license information. By default, **/dlv** displays the license information for the installed operating system. Specifying the **[Activation ID]** parameter displays the license information for the specified edition associated with that Activation ID. Specifying the **[All]** parameter displays all applicable installed products' license information. This operation does not require elevated privileges.
/xpr [Activation ID]	Display the activation expiration date for the product. By default, this refers to the current Windows edition and is primarily useful for KMS clients, because MAK and retail activation is perpetual. Specifying the **[Activation ID]** parameter displays the activation expiration date of the specified edition associated with that Activation ID.This operation does not require elevated privileges.
Advanced options	
/cpky	Some servicing operations require the product key to be available in the registry during Out-of-Box Experience (OOBE) operations. The **/cpky** option removes the product

Option	Description
	key from the registry to prevent this key from being stolen by malicious code. For retail installations that deploy keys, best practices recommend running this option. This option is not required for MAK and KMS host keys, because this is the default behavior for those keys. This option is only needed for other types of keys where the default behavior is not to clear the key from the registry. This operation must be run from an elevated command prompt.
/ilc <license_file>	This option installs the license file specified by the required parameter. These licenses may be installed as a troubleshooting measure, to support token-based activation, or as part of a manual installation of an on-boarded application. Licenses are not validated during this process: License validation is out of scope for Slmgr. Instead, validation is handled by the Software Protection Service at runtime. This operation must be run from an elevated command prompt, or the **Standard User Operations** registry value must be set to allow unprivileged users extra access to the Software Protection Service.
/rilc	This option reinstalls all licenses stored in %SystemRoot%\system32\oem and %SystemRoot%\System32\spp\tokens. These are “known-good” copies stored during installation. Any matching licenses in the Trusted Store are replaced. Any additional licenses—for example, Trusted Authority (TA) Issuance Licenses (ILs), licenses for applications—are not affected. This operation must be run from an elevated command prompt, or the **Standard User Operations** registry value must be set to allow unprivileged users extra access to the Software Protection Service.
/rearm	This option resets the activation timers. The **/rearm** process is also called by **sysprep /generalize**. This operation will do nothing if the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\SkipRearm

Option	Description
	registry subkey is equal to **1**. See the section “Registry Settings,” later in this guide, for details on this registry subkey. This operation must be run from an elevated command prompt, or the **Standard User Operations** registry value must be set to allow unprivileged users extra access to the Software Protection Service.
/rearm-app <Application ID>	Resets the licensing status of the specified app.
/rearm-sku <Application ID>	Resets the licensing status of the specified SKU.
/upk [ActivationID]	This option uninstalls the product key of the current Windows edition. After a restart, the system will be in an Unlicensed state unless a new product key is installed. Optionally, you can use the **[Activation ID]** parameter to specify a different installed product. This operation must be run from an elevated command prompt.
/dti [Activation ID]	Displays installation ID for offline activation.
/atp <Confirmation ID>	Activate product with user-provided confirmation ID.
Volume Licensing: KMS client options	
/skms <Name[:Port] | : port> [Activation ID]	This option specifies the name and, optionally, the port of the KMS host computer to contact. Setting this value disables auto-detection of the KMS host. If the KMS host uses Internet Protocol version 6 (IPv6) only, the address must be specified in the format [hostname]:port. IPv6 addresses contain colons (:), which will be parsed incorrectly by the Slmgr.vbs script. This operation must be run from an elevated command prompt.
/skms-domain <FQDN> [Activation ID]	Sets the specific DNS domain in which all KMS SRV records can be found. This setting has no effect if the specific single KMS host is set with the /skms option. Use this option, especially in disjoint namespace environments, to force KMS to ignore the DNS suffix search list and look for KMS host records in the specified DNS domain instead.

Option	Description
/ckms [Activation ID]	This option removes the specified KMS host name, address, and port information from the registry and restores KMS auto-discovery behavior. This operation must be run from an elevated command prompt.
/skhc	This option enables KMS host caching (default), which blocks the use of Domain Name System (DNS) priority and weight after the initial discovery of a working KMS host. If the system can no longer contact the working KMS host, discovery will be attempted again. This operation must be run from an elevated command prompt.
/ckhc	This option disables KMS host caching. This setting instructs the client to use DNS auto-discovery each time it attempts KMS activation (recommended when using priority and weight). This operation must be run from an elevated command prompt.
KMS host configuration	
/sai <interval>	This option sets the interval in minutes for unactivated clients to attempt KMS connection. The activation interval must be between 15 minutes and 30 days, although the default (2 hours) is recommended. The KMS client initially picks up this interval from registry but switches to the KMS setting after it receives the first KMS response. This operation must be run from an elevated command prompt.
/sri <interval>	This option sets the renewal interval in minutes for activated clients to attempt KMS connection. The renewal interval must be between 15 minutes and 30 days. This option is set initially on both the KMS server and client sides. The default is 10080 minutes (7 days). The KMS client initially picks up this interval from the registry but switches to the KMS setting after it receives the first KMS response.

Option	Description
	This operation must be run from an elevated command prompt.
/sprt <port>	This option sets the port on which the KMS host listens for client activation requests. The default TCP port is 1688. This operation must be run from an elevated command prompt.
/sdns	Enable DNS publishing by the KMS host (default). This operation must be run from an elevated command prompt.
/cdns	Disable DNS publishing by the KMS host. This operation must be run from an elevated command prompt.
/spri	Set the KMS priority to normal (default). This operation must be run from an elevated command prompt.
/cpri	Set the KMS priority to low. Use this option to minimize contention from KMS in a co-hosted environment. Note that this could lead to KMS starvation, depending on what other applications or server roles are active. Use with care. This operation must be run from an elevated command prompt.
/act-type [Activation-Type] [Activation ID]	This option sets a value in the registry that limits volume activation to a single type. Activation Type **1** limits activation to Active Directory only; **2** limits it to KMS activation; **3** to token-based activation. The **0** option allows any activation type and is the default value.
Token-based activation configuration	
/lil	List the installed token-based activation issuance licenses.
/ril <ILID> <ILvID>	Remove an installed token-based activation issuance license. This operation must be run from an elevated command prompt.
/stao	Set the **Token-based Activation Only** flag, disabling

Option	Description
	automatic KMS activation. This operation must be run from an elevated command prompt. This option has been removed beginning with Windows Server 2012 R2 and Windows 8.1. Use the **/act–type** option instead.
/ctao	Clear the **Token-based Activation Only** flag (default), enabling automatic KMS activation. This operation must be run from an elevated command prompt. This option has been removed beginning with Windows Server 2012 R2 and Windows 8.1. Use the **/act–type** option instead.
/ltc	List valid token-based activation certificates that can activate installed software.
/fta <Certificate Thumbprint> [<PIN>]	Force token-based activation using the identified certificate. The optional personal identification number (PIN) is provided to unlock the private key without a PIN prompt when using certificates that are protected by hardware (for example, smart cards).
Active Directory-based activation configuration	
/ad-activation-online <Product Key> [Activation Object name]	Collects Active Directory data and initiates Active Directory forest activation using the credentials that the command prompt is running. Local administrator access is not required, but read/write access to the activation object container in the root domain of the forest is required.
/ad-activation-get-IID <Product Key>	This option starts Active Directory forest activation in phone mode. The output is the installation ID (IID) that can be used to activate the forest over the telephone when Internet connectivity is not available. Upon providing the IID in the activation phone call, a CID is returned that is used to complete activation.
/ad-activation-apply-cid <Product Key> <Confirmation ID> [Activation Object name]	With this option, enter the CID provided from the activation phone call to complete activation

Option	Description
[/name: <AO_Name>]	Optionally, you can append the **/name** option to any of these commands to specify a name for the activation object stored in Active Directory. The name must not exceed 40 Unicode characters; use double quotation marks to explicitly define the name string. In Windows Server 2012 R2 and Windows 8.1, you can append the name directly after **/ad-activation-online <Produce Key>** and **/ad-activation-apply-cid** without having to use the /name option.
/ao-list	Displays all of the activation objects that are available to the local computer.
/del-ao <AO_DN> or **/del-ao <AO_RDN>**	Deletes the specified activation object from the forest.

See also

- Volume Activation Technical Reference
- Volume Activation Overview

Registry Settings for Volume Activation

Registry settings

Table 3 describes the registry settings that Volume Activation supports.

Table 3. Registry settings

Value	Type	Description	Scope
Registry subkey: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform			
Activation\AlternateURL	**REG_SZ**	Optionally contains an administrator-provided URL. If this key is populated, some activation dialog boxes will include a new link — **Learn About Activation Online** — that opens in the default browser.	All

Value	Type	Description	Scope
DisableDnsPublishing	**REG_DWORD**	Set this to a non-zero value to block auto-publishing to DNS by the KMS host.	KMS
DnsDomainPublishList	**REG_MULTI_SZ**	**Format:** *domain_name, <priority>, <weight>* Optionally contains a list of fully qualified domains that KMS uses to auto-publish its DNS service (SRV) record. The KMS home domain is always used, so it is not necessary to include it here. This depends on the **DisableDnsPublishing** value. The **priority** and **weight** parameters are optional. If not supplied, these are set to **0** in the SRV record.	KMS
EnableKmsLowPriority	**REG_DWORD**	Set this to a non-zero value to minimize contention from KMS in a co-hosted environment. Note that on a busy system, this can prevent the KMS host from maintaining a sufficient count.	KMS
KeyManagementServiceListeningPort	**REG_SZ**	Set this on the KMS host to force KMS to listen on the TCP port specified. The host uses TCP port 1688 if this value is not specified.	KMS
KeyManagementServiceName	**REG_SZ**	Set this value to force the use of a specific KMS host by the KMS client. No default. If you specify both App-ID and SKU-ID, this is created in a subkey under which the value is stored:..\SoftwareProtectionPlatform\APPID_GUID\SKUID_GUID\ KeyManagementServiceName If you specify only App-ID, SKU-ID is not allowed, and the resulting key is: ..\SoftwareProtection\APPID_GUID\ KeyManagementServiceName If you specify neither App-ID nor SKU-ID, the resulting key is: ..\SoftwareProtectionPlatform	KMS client
KeyManagementServicePort	**REG_SZ**	Set this to force the use of a specific TCP port by the KMS client when it communicates with a KMS. The client uses port 1688 if this value	KMS clien

Value	Type	Description	Scope
		is not specified.	t
KeyManagementServiceVersion	**REG_SZ**	Set this for Microsoft Operations Manager (MOM) 2005 automatic discovery of the KMS. Delete this value if the KMS is no longer functional on the computer.	KMS
UserOperations	**REG_DWORD**	This value does not exist by default, and the default is **0**, which requires elevation. Create this value and set it to **1** to enable standard users to install product keys and activate and rearm computers without requiring elevation. This setting can also be configured by GPOs, which are stored under the **\policies** registry subkey.	All (not just KMS)
VLActivationInterval	**REG_DWORD**	This is set initially on both KMS server and client sides. The default is 120 minutes (2 hours). The KMS client initially picks up this interval from the registry but switches to the KMS setting after it receives the first KMS response. The minimum value is 15 minutes, and the maximum is 43,200 minutes (30 days).	KMS Client and Server
VLRenewalInterval	**REG_DWORD**	This is set initially on both KMS server and client sides. The default is 10,080 minutes (7 days). The KMS client initially picks up this interval from the registry but switches to the KMS setting after it receives the first KMS response. As a result, the client-side setting will never be used. The minimum value is 15 minutes, and the maximum is 43,200 minutes (30 days).	KMS Client and Server
\(APP-ID)\(SKU-ID)\VLActivationType	**REG_DWORD**	Limits the type of activation to a single type. **1** = Active Directory activation only **2** = KMS activation only **3** = Token activation only **0, blank, or values other than 1-3** = all activation types allowed (default value)	[???]
Registry subkey: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows			

Value	Type	Description	Scope
NT\CurrentVersion\SoftwareProtectionPlatform\Activation			
Manual	**REG_DWORD**	**0** = Allow automatic online activation (default) **1** = Disable auto-activation	MAK and KMS host but not KMS client

See also

- Volume Activation Technical Reference
- Volume Activation Overview

KMS Activation Timing and Discovery for Volume Activation

On this page

KMS activation timing

KMS discovery search order

KMS activation timing

KMS activation timing is controlled by two settings: the Activation and Renewal timers:

- The *KMS Activation Interval* determines how often the KMS client attempts activation before it is activated. This value is set on the client by running **slmgr.vbs /sai** . See the sections “Slmgr.vbs Options” and “Registry Settings” for more information. Once a system contacts a KMS host, the KMS host setting will override the local client setting.
- The *KMS Renewal Interval* determines how often the KMS client attempts reactivation with KMS, once it has been activated. This value can be set using **slmgr /sri interval**. The

interval is set in minutes. The default value is 10,080 minutes (7 days). This value can be set on the client but will be overridden by the value that the KMS host provides.

In operation, the KMS client attempts to contact a KMS host on startup and again every Activation Interval minutes upon failure. By default, this happens every two hours. When KMS activation succeeds, the client receives and stores the KMS host's Renewal Interval, and Windows activates for 180 days. Activation renewal operates as a sliding window, similar to a Dynamic Host Configuration Protocol (DHCP) lease. The client is activated for 180 days but still attempts reactivation every Renewal Interval minutes—by default, 7 days. If it succeeds, the 180-day period is reset. If it fails, the system will watch for network changes or other qualifying events to trigger another attempt. After 15 minutes, the system stops monitoring for qualifying events but still attempts every Activation Interval minutes.

KMS discovery search order

When KMS auto-discovery is enabled (default) and KMS host caching is not disabled, KMS clients will use the following order to determine which KMS to use (they will try each step until an activated KMS host with the correct key responds):

- SKU-specific value in the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\AppID\SKUID\KeyManagementServiceName REG_SZ registry value
- AppID-specific value in the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\AppID\KeyManagementServiceName REG_SZ registry value
- Global value in the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\KeyManagementServiceName REG_SZ registry value
- SKU-specific cached KMS host (This is the cached identity of the host used in the last successful KMS activation.)
- DNS auto-discovery (**weight** and **priority** honored)

If a cached KMS does not respond, discovery is performed.

See also

- Volume Activation Technical Reference
- Volume Activation Overview

KMS and MAK Activation Scenarios for Volume Activation

The following sections describe key scenarios for KMS and MAK activation.

On this page

KMS scenarios

MAK scenarios

VAMT support for KMS activation

Converting KMS to MAK activation

KMS scenarios

KMS can support simple, single-site networks and global networks. The following scenarios show:

- The default implementation of KMS.
- The implementation expanded to support a global network.

Default KMS implementation for a single-site network

Contoso has 100 Windows 7 Enterprise clients and a mixed set of Windows Vista, Windows Server 2008, and Windows Server 2008 R2 systems. Contoso is a flat domain (Contoso.com). The DNS server runs Microsoft DNS in its default configuration. This configuration supports DNS dynamic update protocol and DNS record scavenging to remove stale records.

Contoso purchases a license agreement that provides a KMS key, which will activate all of its systems. The information technology (IT) administrator installs Contoso's KMS host key (CSVLK) on two KMS hosts running Windows Server 2008 R2 by using the following command run locally at an elevated command prompt:

```
Slmgr.vbs /ipk <KMS_host_key>
```

The IT administrator then creates a Security Group in Active Directory® Domain Services (AD DS) named KMS_Hosts. The administrator adds the servers KMS_1 and KMS_2 to the KMS_Hosts membership.

The host KMS_1 is activated against Microsoft via the Internet: **Slmgr.vbs /ato**. KMS_1 automatically publishes its SRV resource records (RRs) to DNS. The IT administrator accesses the DNS server, locates the RR for _vlmcs._tcp.contoso.com, and changes its permissions to give KMS_Hosts Read, Write, and Delete permission to the record. The host KMS_2 is now activated against Microsoft via the Internet: **Slmgr.vbs /ato**.

Finally, the administrator confirms that the KMS host exclusion is enabled in Windows Firewall. The Key Management Service firewall exception needs to be enabled.

KMS clients on the Contoso network query DNS and receive the SRV records for both KMS hosts. The clients pick one or the other host and are activated (as soon as the KMS count rises above the threshold). See the section, "Activation Policy Values," for more information about KMS count requirements.

KMS implementation in a complex, global network

Contoso has expanded into two domains, east.contoso.com and west.contoso.com. Network traffic can pass privately between the two networks through a firewalled wide area network (WAN) link. This link has limited bandwidth, so resources (including DNS) are replicated on both sides of the WAN link to reduce traffic when possible.

Contoso uses image-based deployment. Its client systems are standardized on Windows Vista, but the Windows Vista systems are being replaced by Windows 7 clients.

Users regularly travel geographically, and their network connection changes from one domain to the other. Clients' IP addressing is provided dynamically by DHCP, including specifying the local DNS host, local gateway address, and so on.

To provide activation support for the existing Windows Vista clients and the new Windows 7 clients, Contoso installs four KMS hosts (using its KMS B key). Two KMS hosts are configured in east.contoso.com, and two are configured in west.contoso.com.

The IT administrator configures the KMS hosts so that the DNS SRV records drive Windows 7 and Windows Server 2008 R2 clients in the east domain to KMS_E1, if available, or to KMS_E2. If neither is available, the clients will attempt KMS_W1 and, finally, KMS_W2. Likewise, DNS in the west domain is configured so that Windows 7 clients in the west domain will prefer contacting KMS_W1, then KMS_W2 and only when these fail will attempt activation with KMS_E1 and KMS_E2 on the far side of the WAN link.

The IT administrator accomplishes this by making the following configuration changes (see Table 5):

KMS_E* hosts are added to the security group KMS_E. This group is given Read, Write, and Delete rights to the record _vlmcs._tcp.east.contoso.com.

KMS_W* hosts are added to the security group KMS_W. This group is given Read, Write, and Delete rights to the record _vlmcs._tcp.west.contoso.com.

DHCP servers are configured to add east.contoso.com and west.contoso.com to the DNS suffix search list for all clients.

Firewalls between east and west domains are configured to allow RPC traffic to the KMS hosts on port 1688.

Table 5. KMS Host Configuration

Description	Configuration	Host
Configure east.contoso.com SRV priority and weight so that clients will contact KMS_E* hosts if available before trying KMS_W*	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform DnsDomainPublishList= KMS_E1, 10, 75 KMS_E2, 10, 25 KMS_W1, 90, 75 KMS_W2, 90, 25	KMS_E1 KMS_E2

Description	Configuration	Host
hosts. Traffic will be divided: 75% to the KMS_*1 hosts and 25% to the KMS_*2 hosts.		
Configure west.contoso.com SRV priority and weight so that clients will contact KMS_W* hosts if available before trying KMS_E* hosts. Traffic will be divided: 75% to the KMS_*1 hosts and 25% to the KMS_*2 hosts.	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform DnsDomainPublishList= KMS_W1, 10, 75 KMS_W2, 10, 25 KMS_E1, 90, 75 KMS_E2, 90, 25	KMS_W1 KMS_W2

Then, the administrator confirms that the Windows Firewall exceptions are set to allow KMS client traffic and configures client computers, as Table 6 describes.

Table 6. KMS Client Configuration

Description	Configuration
Disable KMS host caching	Slmgr /ckhc
A customer URL is set to direct users with activation issues to the Contoso help desk	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SoftwareProtectionPlatform\Activation AlternateURL = http://contoso.com/help

The client reference computer is KMS activated, then **sysprep /generalize** is run. The system is shut down and imaged using ImageX from the Windows 7 Windows Automated Installation Kit (Windows AIK). The Windows AIK is available to download at http://go.microsoft.com/fwlink/?LinkId=136976.

Top Of Page

MAK scenarios

The following sections describe scenarios for MAK activation. This section involves network and workgroup environments that are typical within large enterprises.

In the Core Network environment, all computers are within a common network managed by AD DS. The Secure Zone represents higher-security Core Network computers that have addition firewall protection.

The Isolated Lab environment is a workgroup that is physically separate from the Core Network, and its computers do not have Internet access. The network security policy states that no information that could identify a specific computer or user may be transferred out of the Isolated Lab.

MAK independent activation

The VAMT allows automation of MAK deployment and activation over the network by distributing MAKs from a centralized console. VAMT queries Microsoft activation servers to get the number of remaining activations for a given MAK, then lists the activation status of all MAK-activated systems in the environment. This count is a snapshot in time, not a real-time count. VAMT version 3.1 is included in the Windows ADK. For more information about VAMT, see Volume Activation Management Tool (VAMT) Technical Reference, and VAMT How-To Topics

In this scenario, the VAMT is deployed in the Core Network environment. The VAMT is installed on a central computer with network access to all client computers. Both the VAMT host and client computers have Internet access. The following instructions describe how to perform independent activation:

1. Install and launch the VAMT on a networked host computer:
    a. Install the Windows ADK on the host computer.
    b. Click Start, and then click **VAMT** to open the VAMT console.
2. Configure the Windows Management Instrumentation (WMI) firewall exception on target computers. Ensure that the WMI firewall exception has been enabled for all target computers.
3. Add computers to the computer information list (CIL):
    a. On the VAMT console, click **Action**, and then click **Add Computers** to display the **Add Computers** dialog box.
    b. Enter a Computer Group name such as Core Network Group to identify the group of computers you are activating.
    c. Click the drop-down list to select a search option. You can search for computers in a workgroup, in an AD DS domain, or by individual computer name or IP address.
    d. If you are searching by individual computer name or IP address, enter that information in the text box below the drop-down list.
    e. If you are searching a domain or a workgroup, select the domain or workgroup from the additional drop-down list displayed for those options. Use the **Filter by computer name** field to search for a specific computer within the domain or workgroup.
    f. Click **OK**.

    The VAMT queries Active Directory Domain Services (AD DS) via Lightweight Directory Access Protocol (LDAP) and adds the computers it finds to the CIL.
4. Collect status information from the discovered computers:

a. For the VAMT to perform an action on a computer, it must have current license status data for that computer. Collect status from individual computers by directly selecting one or more computers in the computer list view pane. To perform operations for an entire group, in the tree view pane, select the Status Unknown node or the User Defined Groups\Core Network Group node.

b. Right-click the group or the desired computers, and then click **Refresh Computer Status**.

c. If you are activating computers that require administrator credentials different from the ones you are currently using, select **Use Alternate Credentials**.

d. Click **OK**.

e. Provide a Domain Administrator account (to collect status from all computers in the domain) and password when prompted.

   The VAMT displays the **Collecting computer information** dialog box while it collects the status of all selected computers. When the process is finished, the refreshed status of each computer appears in the computer list view pane of the VAMT console.

   Note
   To retrieve the licensing status on the selected computers, the VAMT must have administrative permissions on the remote computers, and WMI must be accessible through Windows Firewall. In addition, for workgroup computers, a registry key must be created to enable remote administrative actions under User Account Control (UAC). For more information about configuring systems for VAMT remote management, see the VAMT help.

5. Add a MAK, and determine its remaining activation count:

   a. Click **Options**, and then click **Manage MAKs** to open the **Manage MAK Keys** dialog box.

   b. Click **Add** to enter a MAK.

   c. Enter the MAK, and then click **Validate**.

      Once validated, the **Edition** field is automatically populated.

   d. Provide a meaningful description, and then click **Add**.

      The MAK is now listed in the **Manage MAK Keys** dialog box.

   e. Click **Refresh remaining count** to retrieve the number of remaining activations for the listed MAKs from Microsoft. (This step requires Internet connectivity.)

   f. Click **Exit** to close the dialog box.

6. Install the MAK, and activate the client computers:

   a. Select the computers to be activated, either by clicking a group in the tree view pane or by selecting individual computers in the list view pane.

   b. Right-click the selected group or computers, and then click **MAK Independent Activate** to display the **MAK Independent Activate** dialog box.

   c. Select the appropriate MAK from the **Install MAK** list.

d. Select the Install **MAK (overwrite existing)** and **Activate Now** check boxes to instruct the selected computers to activate immediately by contacting Microsoft Windows Activation over the Internet.

   If an asterisk (*) appears next to the text for these check boxes, the action will apply only to applicable computers. For example, a computer installed with a Windows Vista retail edition cannot be activated using a MAK.

e. If you are activating computers that require administrator credentials different from the those you are currently using, select **Use Alternate Credentials**.

f. Click **OK**.

   VAMT displays the **Assigning Product Keys** dialog box, and then displays the **Activating Computers** dialog box until it completes the requested action. If you selected **Use Alternate Credentials**, you will be prompted to enter the credentials prior to these dialog boxes.

MAK proxy activation

In this scenario, the VAMT is used to activate workgroup computers within an Isolated Lab environment. For workgroups such as this, which are completely isolated from the larger network, you can perform MAK proxy activation by installing a second VAMT on a computer within the isolated workgroup and using removable media to transfer VAMT CILs between that computer and another VAMT host that has Internet access. For more information about configuring systems for VAMT remote management, see the VAMT help. The following steps describe how to use the VAMT tool for proxy activation:

1. Install a VAMT on a workgroup computer in the Isolated Lab:
   a. Install the Windows AIK on a host computer in the Isolated Lab workgroup.

      This computer can be running Windows XP with Service Pack 2 (SP2), Windows Server 2003, Windows Vista, Windows 7, Windows Server 2008, or Windows Server 2008 R2.

   b. Click Start, and then click **VAMT** to open the VAMT console.
   c. Configure the Windows Management Instrumentation (WMI) firewall exception on target computers. Ensure that the WMI firewall exception has been enabled for all target computers.
2. Add computers to the CIL:
   a. On the VAMT console, click **Action**, and then click **Add Computers** to open the **Add Computers** dialog box.
   b. Enter a Group name such as Isolated Lab Group to identify the group of computers you are activating.
   c. Select the **Workgroup** option, and specify the name of the workgroup, and then click **OK**.

      The VAMT searches the workgroup for computers.

      Note

For the VAMT to discover client computers in a workgroup, the network discovery Windows Firewall exception must be enabled.

3. Collect status from the discovered computers:
   a. For the VAMT to perform an action on a computer, it must have current license status data for that computer. Collect status from individual computers by directly selecting one or more computers in the computer list view pane. To perform operations for an entire group, in the left pane, select the Status Unknown node or the User Defined Groups\Isolated Lab Group node.
   b. Right-click the group or the desired computers, and then click **Refresh Computer Status**.
   c. If you are activating computers that require administrator credentials different from the those you are currently using, select **Use Alternate Credentials**.
   d. Click **OK**. When prompted, provide the credentials for an account that has local administrative rights on the selected workgroup computers.

      The VAMT displays the **Collecting computer information** dialog box while it collects the status of all selected computers. When the process is finished, the refreshed status of each computer appears in the computer list view pane of the VAMT console.

      Note
      To retrieve the licensing status on the selected computers, the VAMT must have administrative permissions on the remote computers, and WMI must be accessible through Windows Firewall. In addition, for workgroup computers, a registry key must be created to enable remote administrative actions under UAC.
4. Add a MAK:
   a. Click **Options**, and then click **Manage MAKs** to open the Manage **MAK Keys** dialog box.
   b. Click **Add** to enter a MAK.
   c. Enter the MAK, and then click **Validate**.

      Once validated, the **Edition** field is automatically populated.
   d. Provide a meaningful description, and then click **Add**.

      The MAK is now be listed in the **Manage MAK Keys** dialog box.
   e. Click **Exit** to close the dialog box.

      Note
      Because the VAMT is not installed on a computer with Internet access, the **Refresh Remaining Count** option will not operate. This feature requires Internet access.
5. Install the MAK on the Isolated Lab computers:
   a. Select the Isolated Lab group in the tree view pane.
   b. Right-click the selected group, and then click **MAK Proxy Activate** to display the **MAK Proxy Activate** dialog box.
   c. Select the appropriate MAK from the **Install MAK** list.

d. Select **Install MAK (overwrite existing)**.

   If an asterisk (*) appears next to the text for this checkbox, the action will apply only to applicable computers. For example, a computer installed with a Windows Vista retail edition cannot be activated using a MAK.

e. Clear the **Get Confirmation ID from Microsoft** check box, because this computer does not have Internet access.

f. Clear the **Apply Confirmation ID and Activate** check box, because confirmation IDs (CIDs) have not yet been requested.

g. If you are activating computers that require administrator credentials different from the those you are currently using, select **Use Alternate Credentials**, and then click **OK**.

   VAMT displays the Assigning Product Keys dialog box until it completes the requested action. If you selected Use Alternate Credentials, you will be prompted to enter the credentials prior to this dialog box.

   Note
   Selecting the **Install MAK (overwrite existing)** check box force-installs a MAK on a client computer. This must be done with care. If the pre-SP1 version of Windows Vista has been installed on the computer for more than 30 days, then its Initial Grace period has expired, and the computer will enter Reduced Functionality Mode (RFM) if activation is not completed successfully before the next logon. However, you can use MAK proxy activation to recover properly configured computers from RFM as long as the computers are accessible to the VAMT host. RFM only applies to the pre-SP1 version of Windows Vista. Windows Vista with SP1 or later, Windows 7, Windows Server 2008, or Windows Server 2008 R2 will not enter RFM.

6. Save the CIL:

   In this step, the full CIL is saved on the local (that is, workgroup) VAMT host. In MAK proxy activation, it is critical to retain this file, because the VAMT uses it to apply the CIDs to the proper computers.

   a. Using the VAMT tree view or list view pane, select the group or the individual computers that successfully received a MAK.

   b. From the **File** menu, click **Save** to display the Save the **Computer Information List** dialog box.

   c. Select a directory, and then enter a CIL file name such as *IsolatedLabGroup.CIL*.

   d. Click **Save**.

7. Export the CIL:

   As stated previously, the enterprise's security policy states that no information that could identify a specific computer or user may be transferred out of the Isolated Lab. Therefore, this type of data must be excluded from the CIL file to be transferred to the Core Network VAMT host.

   a. Using the VAMT tree view or list view pane, select the group or the individual computers that successfully received a MAK.

b. Right-click the selected group or computers, and then click **Export Computers** to display the **Save the Computer Information List** dialog box.

c. Select **Exclude any sensitive environment data**.

d. Click **Browse** to display the **Save As** dialog box.

e. Select a directory, and then enter a CIL file name such as *IsolatedLabGroupSecureExport.CIL*.

f. Click **Save**.

g. Copy the IsolatedLabGroupSecureExport.CIL file to removable media (disk drive, CD/DVD, or USB flash drive [UFD]).

Note

Selecting the **Exclude any sensitive environment data** check box excludes Personally Identifiable Information (PII) from being saved in the CIL. As a result, the CIL must be re-imported in this VAMT host and into the full-save CIL file so that the CIDs requested from Microsoft can be correctly assigned to the computers in the Isolated Lab group.

8. Import the CIL on a VAMT host with Internet access:

a. Copy IsolatedLabGroupSecureExport.CIL from the removable media to a drive on a VAMT host that has Internet access.

b. Open the VAMT console.

c. From the **File** menu, click **Import** to open the **Open Computer Information List** dialog box.

d. Locate and select **IsolatedLabGroupSecureExport.CIL** and then click **Open** to load the CIL into the VAMT.

e. Confirm that the computer, group, and workgroup names are not visible.

9. Request CIDs from Microsoft:

a. Right-click the imported computers, and then click **MAK Proxy Activate** to open the **MAK Proxy Activate** dialog box.

b. Clear the **Install MAK (overwrite existing)** check box.

c. Clear the **Get Confirmation ID from Microsoft** check box.

d. Clear the **Apply Confirmation ID and Activate** check box, because this VAMT host is not on the same network as the selected computers.

e. Click **OK**.

The VAMT displays the **Acquiring Confirmation ID online** dialog box while it contacts Microsoft and collects the CIDs.

10. Export the CIL from the Internet-connected VAMT host:

a. After the CIDs have been retrieved for all computers, save the CIL file. In the VAMT console, scroll to the right and confirm that the **PendingCID** column has values.

b. From the **File** menu, click **Save As**, and provide a new name for the file, such as *IsolatedLabGroupSecureExportwithCID.CIL*.

c. Copy the updated file IsolatedLabGroupSecureExportwithCID.CIL to removable media (disk drive, CD/DVD, or UFD).

11. Import the CIL on the original VAMT host within the Isolated Lab:
    a. Copy IsolatedLabGroupSecureExportwithCID.CIL from the removable media onto a drive on the VAMT host in the Isolated Lab.
    b. Open the VAMT console.
    c. From the **File** menu, click **Open** to display the Open **Computer Information List** dialog box.
    d. Locate and select IsolatedLabGroup.CIL, and then click **Open** to load the CIL into the VAMT.

       This displays the previously discovered computers and their license states in the computer list view pane.
    e. From the **File** menu, click **Import** and browse to the file that contains the CIDs (IsolatedLabGroupSecureExportwithCID.CIL).

       This starts the merge process that matches the CIDs to the computers' IIDs.
    f. From the **File** menu, click **Save**, and then enter a file name such as IsolatedLabGroupwithCID.CIL to save this file that associates the IIDs with the corresponding CIDs.

       Important
       Saving this file is critical for reimaging scenarios. It is strongly recommended that you back up this file and protect it from being overwritten.
12. Apply the CIDs, and activate the Isolated Lab computers:
    a. Right-click the Isolated Lab group in the tree view pane, and then click **MAK Proxy Activate** to open the **MAK Proxy Activate** dialog box.
    b. Clear the **Install MAK (overwrite existing)** check box, because the selected computers already have a MAK installed.
    c. Clear the **Get Confirmation ID from Microsoft** check box.
    d. Select the Apply Confirmation ID and Activate check box to instruct the VAMT to activate the selected computers by installing their CIDs.
    e. Click **OK**.

       The VAMT displays the **Assigning Confirmation IDs** dialog box while it installs the CIDs on the selected computers.
13. Optionally, reactivate reimaged computers in the Isolated Lab:

    If the computers in the Isolated Lab have been reimaged but the underlying hardware has not changed, the VAMT can reactivate them using the IsolatedLabGroupwithCID.CIL file.
    a. Redeploy Windows to each computer using the same computer names as before.
    b. Open the VAMT console on the local (workgroup) VAMT host.
    c. From the **File** menu, click **Open**, and then select the IsolatedLabGroupwithCID.CIL file.
    d. In the tree view pane, right-click the **Isolated Lab** group, and then click **Reapply Confirmation ID** to display the **MAK Proxy Reactivate** dialog box.

e. Leave the **Install MAK (overwrite existing)** check box selected, and ensure that the previously used MAK is selected from the list.

f. Leave the **Reapply Confirmation ID and Reactivate** check box selected to reapply the stored CIDs from the IsolatedLabGroupwithCID.CIL file.

g. Leave the Require exact IID match check box selected if the hardware has not changed since the initial MAK Proxy Activation and you are using the same MAK to reactivate the computer.

   Clear this check box if the hardware has changed but you would still like VAMT to attempt a reactivation using the previously stored CID. Note that Windows might not be successfully reactivated in this scenario.

h. If you are activating a computer that requires administrator credentials different from those you are currently using, select the **Use Alternate Credentials** check box.

i. Click **OK**.

   The VAMT displays the **Assigning Product Keys** dialog box while it applies the MAK to the Isolated Lab Group.

   The VAMT displays the **Assigning Confirmation IDs** dialog box while it installs the CIDs on the selected computers.

   Note
   Selecting the **Install MAK (overwrite existing)** check box force-installs a MAK. This must be done with care. If the pre-SP1 version of Windows Vista has been installed for more than 30 days, then its Initial Grace period has expired and it will enter RFM if activation is not completed successfully before the next logon. VAMT can be used to recover properly configured remote computers from RFM as long as they are accessible on the network. RFM only applies to the pre-SP1 version of Windows Vista. Windows Vista with SP1 or later, Windows 7, Windows Server 2008, and Windows Server 2008 R2 will not enter RFM.

Top Of Page

VAMT support for KMS activation

In this scenario, the VAMT is used to install and activate KMS client keys on either Core Network or Isolated Lab computers. The procedure described below assumes that the VAMT has been installed and the computers have been added to the CIL. These tasks are described in steps 1–4 in either Scenario 1 (for the Core Network) or Scenario 2 (for the Isolated Lab workgroup).

1. Open the VAMT console.
2. Select the computers to be activated, either by clicking a group in the tree view pane or by selecting individual computers in the list view pane.
3. Right-click the selected group or computers, and then click **Configure for KMS Activation** to display the **KMS Configuration** dialog box.
4. Select **Install KMS Client Key (overwrite existing)** to instruct the VAMT to install a KMS client key on a client computer.

5. Select the appropriate activation option:
   - **Auto-Discover KMS host using DNS.** The VAMT first clears any previously configured KMS host on the target computer, then instructs the computer to query DNS to locate a KMS host and attempt activation.
   - **Use specific KMS host and port.** The VAMT sets the specified KMS host name and port on the target computer, and then instruct the computer to attempt activation with that specific KMS host.
6. If you are activating computers that require administrator credentials different from those you are currently using, select the **Use Alternate Credentials** check box, and then click **OK**.

   The VAMT displays the **Assigning Product Keys** dialog box, and then displays the **Activating Computers** dialog box until it completes the requested action. If you selected the **Use Alternate Credentials** check box, you will be prompted to enter the credentials prior to these dialog boxes.

Top Of Page

Converting KMS to MAK activation

Windows 7 and Windows Server 2008 R2 install automatically as KMS clients. To convert a KMS client to MAK activation, install a MAK. A MAK can be installed during or anytime after operating system installation.

A MAK key can be installed on a reference image of Windows 7 and Windows Server 2008 R2 to make all installations from that image use MAK activation instead of the default KMS activation. Doing so alleviates the need to specify a MAK in an unattended installation file.

Installing a MAK during operating system installation

You can convert a KMS client to a MAK client during the initial installation of Windows 7 or Windows Server 2008 R2 by including a MAK in an unattended setup (Unattend.xml) file. The Unattend.xml file can be used with Setup.exe or Windows Deployment Services. For more information, see the Unattended Windows Setup Reference help file in the Windows AIK at http://go.microsoft.com/fwlink/?LinkId=136976.

Note

The MAK is stored in clear text in the Unattend.xml file. During an unattended installation, the file Unattend.xml or AutoUnattend.xml is copied to the %SystemRoot%\Panther folder of the target computer. However, at the end of the Setup process, the Setup program replaces it with “SENSITIVE*DATA*DELETED.”

Installing a MAK after operating system installation

You can configure a volume edition of Windows 7 or Windows Server 2008 R2 to use MAK activation by using the Control Panel System item or by running the Slmgr.vbs script:

- To install a MAK by using System Center, click the **Change your product key** link, and then type the MAK in the **Change your product key for activation** dialog box.

- To install a MAK by using Slmgr.vbs, run the following command at a command prompt:

  ```
  slmgr.vbs /ipk <MultipleActivationKey>
  ```

  where *MultipleActivationKey* is the MAK.

If users install a MAK by using the user interface (UI), the MAK client attempts to activate itself over the Internet one time. If the users install a MAK key by using the Slmgr.vbs script, the MAK client does not try to activate automatically. The user can use **Slmgr.vbs /ato** to manually activate the MAK over Internet.

Top Of Page

See also

- Volume Activation Technical Reference
- Volume Activation Overview

KMS Host Keys to Products Activated for Volume Activation

Table 7 describes the products that each KMS host can activate.

Note
This table is currently being updated for Windows 8.1 and Windows Server 2012 R2. Please check back for the updated version.

Table 7. KMS host keys matched to products they activate

Volume product group	Windows products	Key type	Products activated by key type
Windows Vista	Windows Vista Business Windows Vista Enterprise	KMS	Windows Vista Business Windows Vista Enterprise
Windows 7	Windows 7 Professional Windows 7 Enterprise	KMS	Windows Vista Business Windows Vista Enterprise Windows 7 Professional

Volume product group	Windows products	Key type	Products activated by key type
			Windows 7 Enterprise
Windows Server 2008 R2 Server Group A:	Firstref_server_7web Windows Server 2008 R2 HPC Edition Windows HPC Server 2008 R2	KMS_A	Windows Server 2008 R2 Group A Windows Server 2008 Group A Windows Vista Volume Editions Windows 7 Volume Editions
Windows Server 2008 R2 Server Group B	Windows Server 2008 R2 Standard Windows Server 2008 R2 Enterprise	KMS_B	Windows Server 2008 R2 Group A Windows Server 2008 Group A Windows Server 2008 R2 Group B Windows Server 2008 Group B Windows Vista Volume Editions Windows 7 Volume Editions
Windows Server 2008 R2 Server Group C	Windows Server 2008 R2 Datacenter Windows Server 2008 R2 for Itanium-Based Systems	KMS_C	Windows Server 2008 R2 Group A Windows Server 2008 Group A Windows Server 2008 R2 Group B Windows Server 2008 Group B Windows Server 2008 R2 Group

Volume product group	Windows products	Key type	Products activated by key type
			C Windows Server 2008 Group C Windows Vista Volume Editions Windows 7 Volume Editions

Windows Server 2008 R2 KMS keys follow the same hierarchy (groups A, B, C) as Windows Server 2008 with the addition of the Windows Server 2008 Datacenter without Hyper-V™, Windows Server 2008 Enterprise without Hyper-V, and Windows Server 2008 Standard without Hyper-V. The primary difference to note is that Windows Server 2008 R2 KMS keys can be used to activate a KMS host that will in turn activate Windows Server 2008, Windows 7, and Windows Vista.

Windows Server 2008 and Windows Vista keys cannot activate Windows 7 or Windows Server 2008 R2 systems.

Top Of Page

See also

- Volume Activation Technical Reference
- Volume Activation Overview

Activation Policy Values for Volume Activation

Table 8 lists policy values for Windows 8.1 and Windows Server 2012 R2.

Table 8. Windows 8.1 and Windows Server 2012 R2 policy values

Element	Client value	Server value	Range
Rearm limit	1000	1000	Not configurable
KMS Client Count Threshold	25	5	Not configurable

Element	Client value	Server value	Range
KMS Activation duration	180 days	180 days	Not configurable
KMS host listening port	1688	1688	Any open TCP port
OOB Grace	Not available	Not available	Not configurable
OOT Grace	Not available	Not available	Not configurable
KMS client reactivation (in minutes)	10,080 minutes (7 days—default)	10,080 minutes (7 days—default)	15–43,200 minutes (30 days—default)
Silent KMS retry interval—occurs on KMS activation failure	120 minutes	120 minutes	15–43,200 minutes (120 minutes—default)

See also

- Volume Activation Technical Reference
- Volume Activation Overview

KMS Client Setup Keys for Volume Activation

By default, the Windows 7 and Windows Server 2008 R2 and later operating systems use KMS for activation. In volume installations, the setup key is installed by default, which makes the system a KMS client. If you are converting a computer from a KMS host, MAK, or retail edition of Windows to a KMS client, install the applicable setup key (GVLK) from Appendix A: KMS Client Setup Keys by using **slmgr /ipk <setup key>**.

See also

- Volume Activation Technical Reference
- Volume Activation Overview

Activation Error Codes for Volume Activation

Table 10 provides troubleshooting help for problems that occur when activating volume editions of Windows 7 and Windows Server 2008 R2 operating systems.

Table 10. Error codes and descriptions

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
0xC004C001	The activation server determined the specified product key is invalid	MAK	An invalid MAK was entered.	Verify that the key is the MAK provided by Microsoft. Contact the Microsoft Activation Call Center to verify that the MAK is valid.
0xC004C003	The activation server determined the specified product key has been blocked	MAK	The MAK is blocked on the activation server.	Contact the Microsoft Activation Call Center to obtain a new MAK and install/activate the system.
0xC004C008	The activation server reported that the product key has exceeded its unlock limit.	KMS	The KMS key has exceeded the activation limit.	KMS host keys will activate up to 10 times on six different computers. If more activations are necessary, contact the Microsoft Activation Call Center.
0xC004C020	The activation server reported that the Multiple Activation Key has exceeded its limit.	MAK	The MAK has exceeded the activation limit.	MAKs by design have a limited number of activations. Contact the Microsoft Activation Call Center.
0xC004C021	The activation server reported that the Multiple Activation Key extension limit	MAK	The MAK has exceeded the activation limit.	MAKs by design have a limited number of activations. Contact the Microsoft Activation Call

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
	has been exceeded.			Center.
0xC004F009	The Software Protection Service reported that the grace period expired.	MAK	The grace period expired before the system was activated. Now, the system is in the Notifications state.	See the section “User Experience.”
0xC004F00F	The Software Licensing Server reported that the hardware ID binding is beyond level the of tolerance.	MAK/KMS client/KMS host	The hardware has changed or the drivers were updated on the system.	MAK: Reactivate the system during the OOT grace period using either online or phone activation. KMS: Restart, or run slmgr.vbs /ato.
0xC004F014	The Software Protection Service reported that the product key is not available	MAK/KMS client	No product keys are installed on the system.	Install a MAK product key, or install a KMS Setup key found in \sources\Product.ini on the installation media.
0xC004F02C	The Software Protection Service reported that the format for the offline activation data is incorrect.	MAK/KMS client	The system has detected that the data entered during phone activation is not valid.	Verify that the CID is correctly entered.
0xC004F038	The Software Protection Service reported that the computer could not be activated. The	KMS client	The count on the KMS host is not high enough. The KMS count must be ≥5 for Windows Server or ≥25 for Windows client.	More computers are needed in the KMS pool for KMS clients to activate. Run Slmgr.vbs /dli to get the current count on the KMS host.

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
	count reported by your Key Management Service (KMS) is insufficient. Please contact your system administrator.			
0xC004F039	The Software Protection Service reported that the computer could not be activated. The Key Management Service (KMS) is not enabled.	KMS client	This error occurs when a KMS request is not answered.	Troubleshoot the network connection between the KMS host and the client. Make sure that TCP port 1688 (default) is not blocked by a firewall or otherwise filtered.
0xC004F041	The Software Licensing Service determined that the Key Management Service (KMS) is not activated. KMS needs to be activated. Please contact system administrator.	KMS client	The KMS host is not activated.	Activate the KMS host with either online or phone activation.
0xC004F042	The Software Protection Service determined that the specified Key Management	KMS client	Mismatch between the KMS client and the KMS host.	This error occurs when a KMS client contacts a KMS host that cannot activate the client software. This can be common in mixed

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
	Service (KMS) cannot be used.			environments that contain application and operating system-specific KMS hosts, for example.
0xC004F050	The Software Protection Service reported that the product key is invalid.	KMS, KMS client, MAK	This can be caused by a typo in the KMS key or by typing in a Beta key on a Released version of the operating system.	Install the appropriate KMS key on the corresponding version of Windows. Check the spelling. If the key is being copied and pasted, make sure that em dashes have not been substituted for the dashes in the key.
0xC004F051	The Software Protection Service reported that the product key is blocked.	MAK/KMS	The product key on the activation server is blocked by Microsoft.	Obtain a new MAK/KMS key, install it on the system, and activate.
0xC004F074	The Software Protection Service reported that the computer could not be activated. No Key Management Service (KMS) could be contacted. Please see the Application Event Log for additional information.	KMS Client	All KMS host systems returned an error.	Troubleshoot errors from each event ID 12288 associated with the activation attempt.

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
0xC004F06C	The Software Protection Service reported that the computer could not be activated. The Key Management Service (KMS) determined that the request timestamp is invalid.	KMS client	The system time on the client computer is too different from the time on the KMS host.	Time sync is important to system and network security for a variety of reasons. Fix this issue by changing the system time on the client to sync with the KMS. Use of a Network Time Protocol (NTP) time source or Active Directory Domain Services for time synchronization is recommended. This issue uses UTP time and is independent of Time Zone selection.
0x80070005	Access denied. The requested action requires elevated privileges.	KMS client/MAK/KMS host	User Account Control (UAC) prohibits activation processes from running in a non-elevated command prompt.	Run slmgr.vbs from an elevated command prompt. Right-click cmd.exe, and then click Run as Administrator.
0x8007232A	DNS server failure.	KMS host	The system has network or DNS issues.	Troubleshoot network and DNS.
0x8007232B	DNS name does not exist.	KMS client	The KMS client cannot find KMS SRV RRs in DNS. If a KMS host does not exist on the network, a MAK should be installed.	Confirm that a KMS host has been installed and DNS publishing is enabled (default). If DNS is unavailable, point the KMS client to the

Error code	Error message	Activation type	Possible cause	Troubleshooting steps
				KMS host by using slmgr.vbs /skms <kms_host_name>. Optionally, obtain and install a MAK; then, activate the system. Finally, troubleshoot DNS.
0x800706BA	The RPC server is unavailable.	KMS client	Firewall settings are not configured on the KMS host, or DNS SRV records are stale.	Ensure the Key Management Service firewall exception is enabled on the KMS host computer. Ensure that SRV records point to a valid KMS host. Troubleshoot network connections.
0x8007251D	No records found for given DNS query.	KMS client	The KMS client cannot find KMS SRV RRs in DNS.	Troubleshoot network connections and DNS.

See also

- Volume Activation Technical Reference
- Volume Activation Overview

WMI Properties and Methods for Volume Activation

Table 12 lists WMI properties and methods for the Software Protection Platform. Software Protection Platform properties are defined in the file %WinDir%\System32\wbem\sppwmi.mof.

Table 12. Software protection platform WMI properties and methods

//properties SoftwareLicensingService Class	
Version	Version of the Software Protection Service
KeyManagementServiceMachine	The name of the KMS host. Returns null if **SetKeyManagementServiceMachine** has not been called.
KeyManagementServicePort	The TCP port used by clients to send KMS activation requests. Returns **0** if **SetKeyManagementServicePort** has not been called.
IsKeyManagementServiceMachine	Indicates whether KMS is enabled on the computer: **0** if false, **1** if true.
VLActivationInterval	The frequency, in minutes, of how often a client will contact the KMS host before the client is licensed.
VLRenewalInterval	The frequency, in minutes, of how often a client will contact the KMS host after the client is licensed.
KeyManagementServiceCurrentCount	The count of currently active KMS clients on the KMS host. **-1** indicates that the host is not enabled as a KMS or has not received any client licensing requests.
RequiredClientCount	The minimum number of clients required to connect to a KMS host in order to enable Volume Licensing.
PolicyCacheRefreshRequired	Indicates whether the licensing policy cache needs to be updated: **0**=not required, **1**=Refresh required.
ClientMachineID	The globally unique identifier (GUID) that identifies a KMS client to a KMS host. The client includes this in requests it sends to the KMS.
RemainingWindowsReArmCount	Remaining number of times the client can be rearmed successfully.

KeyManagementServiceListeningPort	The TCP port the KMS host uses to listen for activation requests.
KeyManagementServiceDnsPublishing	Indicates the DNS publishing status of a KMS host: **0**=Disabled, **1**=Auto-publish enabled (default).
KeyManagementServiceLowPriority	Indicates the thread priority status of KMS service: **0**=Normal Priority (default), **1**=Low priority.
KeyManagementServiceHostCaching	Indicates the caching status of KMS host name and port: **0**=Caching disabled, **1**=Caching enabled (default).
KeyManagementServiceUnlicensedRequests	The count of KMS requests from clients with License Status=0 (Unlicensed).
KeyManagementServiceLicensedRequests	The count of KMS requests from clients with License Status=1 (Licensed).
KeyManagementServiceOOBGraceRequests	The count of KMS requests from clients with License Status=2 (OOBGrace).
KeyManagementServiceOOTGraceRequests	The count of KMS requests from clients with License Status=3 (OOTGrace).
KeyManagementServiceNonGenuineGraceRequests	The count of KMS requests from clients with License Status=4 (NonGenuineGrace).
KeyManagementServiceNotificationRequests	The count of KMS requests from clients with License Status=5 (Notification).
KeyManagementServiceTotalRequests	The total count of valid KMS requests.
KeyManagementServiceFailedRequests	The count of invalid KMS requests.
KeyManagementServiceActivationDisabled	Indicates whether the volume activation through KMS is disabled.
//methods SoftwareLicensingService Class	
InstallProductKey	Installs a product key.

InstallLicense	Installs a license.
InstallLicensePackage	Installs a license package for the current product.
SetKeyManagementServiceMachine	Sets the KMS host name to use for volume activation.
ClearKeyManagementServiceMachine	Clears any previously configured KMS host name.
SetKeyManagementServicePort	Sets the TCP port used by a client to make requests of a KMS host. If not specified, port 1688 is used.
ClearKeyManagementServicePort	Clears any previously specified port number.
SetVLActivationInterval	The activation frequency, in minutes, of how often KMS clients will contact the KMS host before they become licensed. The frequency must be ≥15 and ≤43,200. An error is returned if the method is called and the computer is not a KMS.
SetVLRenewalInterval	The renewal frequency, in minutes, of how often KMS clients will contact the KMS host after they have become licensed. The frequency must be ≥15 and ≤43,200. An error is returned if the method is called and the computer is not a KMS.
ClearProductKeyFromRegistry	Clears product key from the registry.
AcquireGenuineTicket	Perform online Genuine Validation. Calling this method may change the Genuine status of the computer.
ReArmWindows	Resets the licensing status of the computer to OOB Grace (see **LicenseStatus**). **Note** You must restart the client for the changes to take effect.
RefreshLicenseStatus	Updates the licensing status of

	Windows so that applications have access to current licensing information.
SetKeyManagementServiceListeningPort	Sets the TCP port used by a KMS host to listen for activation requests. Applies to KMS hosts only. If not specified, port 1688 is used.
ClearKeyManagementServiceListeningPort	Clears any previously specified listening port. Applies to KMS hosts only.
DisableKeyManagementServiceDnsPublishing	Enable/disable DNS Publishing on a KMS host computer: **0**=Enable, **1**=Disable.
EnableKeyManagementServiceLowPriority	Enable/disable KMS service running with low priority: **0**=Disable, **1**=Enable.
DisableKeyManagementServiceHostCaching	Enable/disable the caching of the KMS host name and port on a volume activation client computer: **0**=Enable, **1**=Disable.
DisableKeyManagementServiceActivation	Enable/disable volume activation through a KMS computer: **0**=Enable, **1**=Disable.
DiscoveredKeyManagementServiceMachineName	Last discovered KMS host name through DNS.
DiscoveredKeyManagementServiceMachinePort	Last discovered KMS host port through DNS.
GenerateActiveDirectoryOfflineActivationId	Generates an installation ID for use in Active Directory activation using the specified product key.
DepositActiveDirectoryOfflineActivationConfirmation	Processes an installation ID, a confirmation ID, and ActivationObjectName to generate and publish the activation object in Active Directory.
DoActiveDirectoryOnlineActivation	Performs online activation of Active Directory using a specified product key.

SetVLActivationTypeEnabled	Sets the activation method.
ClearVLActivationTypeEnabled	Clears previously specified activation method.
//properties SoftwareLicensingProduct Class	
ID	Product identifier
Name	Product name
Description	Product description
ApplicationID	ID of current product's application
ProcessorURL	Software licensing server URL for the process certificate
MachineURL	Software licensing server URL for the binding certificate
ProductKeyURL	Software licensing server URL for the product certificate
UseLicenseURL	Software licensing server URL for the user license
LicenseStatus	License status of this product's application: **0**=Unlicensed, **1**=Licensed, **2**=OOBGrace, **3**=OOTGrace, **4**=NonGenuineGrace, **5**=Notification, **6**=ExtendedGrace
LicenseStatusReason	A diagnostic code that indicates why a computer is in a specific licensing state.
GracePeriodRemaining	Remaining time in minutes before the parent application goes into Notification mode. For volume clients, this is the remaining time before reactivation is required.
EvaluationEndDate	The expiration date of this product's application. After this date, the

	LicenseStatus will be **Unlicensed** and cannot be activated.
OfflineInstallationId	An identifier for this product's application that can be used for telephone or offline activation. Returns null if a product key is not installed.
PartialProductKey	Last five characters of this product's key. Returns null if a product key is not installed.
ProductKeyID	Product key ID. Returns null if a product key is not installed.
LicenseFamily	The family identifier for the SKU used to determine license relationships for add-ons.
LicenseDependsOn	The dependency identifier for the family of SKUs used to determine license relationships for add-ons.
LicenseIsAddon	Returns **True** if the product is identified as an add-on license.
VLActivationInterval	The frequency, in minutes, of how often a client will contact the KMS host before the product is licensed.
VLRenewalInterval	The frequency, in minutes, of how often a client will contact the KMS host after the product is licensed.
KeyManagementServiceProductKeyID	KMS product key ID. Returns null if not applicable.
KeyManagementServiceMachine	The name of the KMS host. Returns null if **SetKeyManagementServiceMachine** has not been called.
KeyManagementServicePort	The TCP port used by clients to send KMS activation requests. Returns **0** if **SetKeyManagementServicePort** has not been called.
DiscoveredKeyManagementServiceMachineName	Last discovered KMS host name

	through DNS.
DiscoveredKeyManagementServiceMachinePort	Last discovered KMS host port through DNS.
IsKeyManagementServiceMachine	Indicates if KMS is enabled on the computer: **1** if true, **0** if false.
KeyManagementServiceCurrentCount	The count of currently active KMS clients on the KMS host. **-1** indicates that the computer is not enabled as a KMS or has not received any client licensing requests.
RequiredClientCount	The minimum number of clients required to connect to a KMS host in order to enable Volume Licensing.
KeyManagementServiceUnlicensedRequests	The count of KMS requests from clients with License Status=0 (Unlicensed).
KeyManagementServiceLicensedRequests	The count of KMS requests from clients with License Status=1 (Licensed).
KeyManagementServiceOOBGraceRequests	The count of KMS requests from clients with License Status=2 (OOBGrace).
KeyManagementServiceOOTGraceRequests	The count of KMS requests from clients with License Status=3 (OOTGrace).
KeyManagementServiceNonGenuineGraceRequests	The count of KMS requests from clients with License Status=4 (NonGenuineGrace).
KeyManagementServiceTotalRequests	The total count of valid KMS requests.
KeyManagementServiceFailedRequests	The count of failed KMS requests.
KeyManagementServiceNotificationRequests	The count of KMS requests from clients with License Status=5 (Notification).
GenuineStatus	Genuine status for this product's application.
ExtendedGrace	Extended grace time in minutes

	before the parent application becomes unlicensed.
TrustedTime	Displays the currently stored trusted time in the system.
DisableKeyManagementServiceHostCaching	Enable/disable the caching of the KMS host name and port on a volume activation client computer: 0=Enable, 1=Disable.
VLActivationType	Defines the activation type used for the last successful client activation: 1=Active Directory, 2=KMS, 3=Token. If blank, the client was never activated.
VLActivationTypeEnabled	Activation type configured for the volume licensing client: 1=Active Directory, 2=KMS, 3=Token. If blank or 0, all activation types are configured.
ADActivationObjectName	Name of the activation object used for the last Active Directory activation.
ADActivationObjectDN	The fully qualified distinguished name of the activation object used for the last Active Directory activation.
ADActivationCsvlkPid	The CSVLK PID that is in the activation object used for the last Active Directory activation.
ADActivationCsvlkSkuId	The CSVLK SKU ID that is in the activation object used for the last Active Directory activation.
// methods SoftwareLicensingProduct Class	
UninstallProductKey	Uninstalls a product key for the current product.
Activate	Activates the current product.
DepositOfflineConfirmationId	Activates a product by depositing an Offline Confirmation Identifier for this product when performing telephone or offline activation.

GetPolicyInformationDWord	Gets license policy information of type DWORD.
GetPolicyInformationString	Gets policy information of type string.
SetKeyManagementServiceMachine	Sets the KMS host name to use for volume activation.
ClearKeyManagementServiceMachine	Clears any previously configured KMS host name.
SetKeyManagementServicePort	Sets the TCP port used by a client to make requests of a KMS host. If not specified, port 1688 is used.
ClearKeyManagementServicePort	Clears any previously specified port number.
SetVLActivationTypeEnabled	Sets the configured activation method for the volume licensing client.
ClearVLActivationTypeEnabled	Clears any previously specified activation method configured for the volume licensing client.

See also

- Volume Activation Technical Reference
- Volume Activation Overview

Web Server (IIS) Overview

This document introduces the Web Server (IIS) role of Windows Server® 2012, describes new IIS 8 features, and links to additional Microsoft and community information about IIS.

Did you mean…

- Application Server Overview
- Security and Protection
- Web Server (IIS 7)
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Role Description

The Web Server (IIS) role in Windows Server 2012 provides a secure, easy-to-manage, modular and extensible platform for reliably hosting websites, services, and applications. With IIS 8 you

can share information with users on the Internet, an intranet, or an extranet. IIS 8 is a unified web platform that integrates IIS, ASP.NET, FTP services, PHP, and Windows Communication Foundation (WCF).

The following list shows just some of the benefits you get by using IIS 8:

- Maximize web security through a reduced server foot print and automatic application isolation.
- Easily deploy and run ASP.NET, classic ASP, and PHP web applications on the same server.
- Achieve application isolation by giving worker processes a unique identity and sandboxed configuration by default, further reducing security risks.
- Easily add, remove, and even replace built-in IIS components with custom modules, suited for customer needs.
- Speed up your website through built-in dynamic caching and enhanced compression.

Practical Applications

Administrators can use the Web Server (IIS) role to set up and manage multiple websites, web applications, and FTP sites. Some of the specific features include the following:

- Use IIS Manager to configure IIS features and administer you websites.
- Use File Transfer Protocol (FTP) to allow website owners to upload and download files.
- Use website isolation to protect against one website from interfering with other sites on your server.
- Configure web applications that are written using various technologies, such as classic ASP, ASP.NET, and PHP.
- Use Windows PowerShell to automate management of most administration tasks for your web server.
- Configure multiple web servers into a server farm that you can manage using IIS.
- Take full advantage NUMA hardware, and get optimum performance from your NUMA-enabled server.

New and Changed Functionality

The following table shows some of the features that are new for IIS 8:

Feature/functionality	New or Updated	Summary
Centralized Certificates	New	Provides a single SSL certificate store for a server farm and simplifies the management of SSL bindings.
Dynamic IP Restrictions	New	Enables administrators to configure IIS 8 to block access for IP addresses that exceed

Feature/functionality	New or Updated	Summary
		the specified number of requests and to specify the behavior when an IP address is blocked.
FTP Logon Attempt Restrictions	New	Restricts the number of failed logon attempts that can be made to an FTP account in a specified time period.
Server Name Indication (SNI)	New	Extends the SSL and TSL protocols to allow a virtual domain name, or a hostname, to be used to identify the network end point.
Application Initialization	New	Enables web administrators to configure IIS 8 to initialize web applications, so the application is ready for the first request.
NUMA-aware scalability	New	Provides support for NUMA hardware, which allows 32-128 CPU cores. This support provides near optimal out-of-the-box performance on NUMA hardware.
IIS CPU Throttling	Updated	Limits CPU, memory, and bandwidth consumption by a single application pool in a multitenant deployment. IIS 8 includes additional throttling options.

Deprecated Functionality

The Internet Information Services (IIS) Manager 6.0 Microsoft Management Console (MMC) snap-in is deprecated in Windows Server 2012. It will be removed in a future release of Windows Server.

See Also

The following table contains links to resources related to this scenario.

Content type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center | IIS Media Services | What's New in IIS 8.0 for Windows 8?
Troubleshooting	IIS Troubleshooters
Security	Secure Windows Server 2012 R2 and Windows Server 2012 | Security and Protection
Tools and settings	Web Server (IIS) Administration Cmdlets in Windows PowerShell
Community resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related technologies	ASP.NET | ASP.NET Web Projects

Hosting-Friendly Web Server Platform (IIS): Scenario Overview

This document contains an overview for the hosting-friendly web server, IIS 8, by presenting a collection of scenarios for configuring IIS 8 to support various types of web applications.

Did you mean…

- Building Your Cloud Infrastructure: Scenario Overview
- Web Server (IIS) Overview

Scenario Description

The IIS 8 web server provides a secure, easy-to-manage, modular and extensible platform for reliably hosting websites, services, and applications. The web server scenarios provide end-to-end instructions on how to plan, install, and configure IIS to host various types of services and applications.

Each scenario is self-contained. You can read them in any order and concentrate on those scenarios that most closely meet your needs.

Web Server Scenarios

Scenario	Description
Build a Static Website on IIS	Shows how to install IIS with the minimum module configuration to server static websites. Describes how to add a website, configure anonymous authentication, set up a default document, set up static content compression, use IIS logging, and filter requests.
Build a Classic ASP Website on IIS	Shows how to install IIS with the modules required to server classic ASP websites. Explains how to add a websites and edit ASP application settings.
Build an ASP.NET Website on IIS	Guides you through planning, installing, and configuring an ASP.NET website on IIS. Describes what IIS modules to install. Shows how to configure ASP.NET settings, such as session state, pages and controls, application settings, .NET compilation and globalization settings. Explains how to configure data source settings. And shows how to improve security with application isolation, .NET trust levels, .NET authentication, machine key settings, and SSL communication.
Build a PHP Website on IIS	Guides you through planning, installing, and configuring a PHP website on IIS. Describes what IIS modules to install. Describes how to configure essential PHP settings. Explains how to configure WinCache and other PHP extensions. Describes how to improve security of a PHP application on IIS.
Build an FTP Site on IIS	Shows how to install FTP on an existing IIS web server. Describes how to add an FTP site, configure FTP site defaults, fire wall support, user isolation, and directory browsing options. Shows how to configure logon attempt restrictions, request filtering, FTP logging, and

Scenario	Description
	FTP messages.
Build a Web Farm with IIS Servers	Guides you through planning, installing, and configuring a simple server farm with IIS servers. Helps you plan and configure the farm infrastructure. Shows how set up load balancing by using Application Request Routing (ARR). Describes how to set up shared content and shared configuration. Shows how to set up a central certificate store for SSL certificates. Explains how to set up FTP and Web Deploy on the server farm.

Practical Applications

Whether you are an IT professional, a web developer, or you want to set up your own web server, this scenario can help you install IIS and configure it to serve your web application.

See Also

The following table contains links to resources related to IIS 8.

Content Type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center | IIS Media Services | What's New in IIS 8.0 for Windows 8?
Troubleshooting	IIS Troubleshooting
Security	Secure Windows Server 2012 R2 and Windows Server 2012 | Security and Protection
Tools and Settings	Web Server (IIS) Administration Cmdlets in Windows PowerShell
Community Resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related Technologies	ASP.NET | ASP.NET Web Projects

Build a Static Website on IIS

This document guides you through the process of installing an IIS web server and configuring it to serve static content. Static content is a web page (HTML) that is delivered to the user exactly as stored. By contrast, dynamic content is generated by a web application, such as an ASP.NET, classic ASP, or PHP application. Static content displays the same information for all users; dynamic content can display user-specific information, such as the user name.

A static-content web server is the most basic configuration of IIS for supporting HTML websites. You can use a static-content web server to host internal or external (public) websites. When you install IIS 8, the default installation provides all the IIS modules required to support a static-content web server. The default installation includes the ability to serve static HTML files, documents, and images. IIS 8 supports default documents, directory browsing, logging, and anonymous authentication for the static content server.

In this document

- Prerequisites
- Step 1: Install the IIS Web Server
- Step 2: Add a Website
- Step 3: Configure Anonymous Authentication
- Step 4: Configure the Default Documents
- Step 5: Configure Static Content Compression
- Next Steps

Prerequisites

To get the most from this tutorial, you must have access to a computer that is running one of the following operating systems:

- Windows Server® 2012
- Windows® 8

Step 1: Install the IIS Web Server

You can use the Web Platform Installer (Web PI) to install IIS and applications that run on IIS. The Web PI installs the latest versions of web platform offerings, including updates. To learn more about the Web PI, see Learn more and install the Web PI. If you use the Web PI to install IIS, you can skip to step 2.

You can also perform this procedure by using the Windows user interface (UI) or from a command line.

To install IIS on Windows Server 2012 by using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.

2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select the name of your computer, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, notice the preselected features, and then click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select Role Services** page, note the preselected role services, and then click **Next**.

   Note
   Install the IIS 8 default role services for a static-content web server.
10. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
11. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
12. To verify that IIS installed successfully, enter the following into a web browser:

    http://localhost

    You see the default IIS Welcome page.

To install IIS on Windows 8 by using the UI

1. On the **Start** page, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services**, and then click **OK**.

   This action installs the IIS 8 default features. Install only the default features for a static-content web server.
4. To verify that IIS installed successfully, enter the following into a web browser:

   http://localhost

   You see the default IIS Welcome page.

To Install IIS by using the command line

- Enter the following command at an elevated command prompt or into a script:

```
Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-
```

```
WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-
ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI
```

Step 2: Add a Website

This step tells you how to add a website to IIS by using the IIS Manager UI, or by running Appcmd.exe commands in a command-line window.

To add a website by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In the **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, right-click the **Sites** node in the tree, and then click **Add Website**.
3. In the **Add Website** dialog box, enter a friendly name for your website in the **Site name** box.
4. If you want to select a different application pool than the one listed in the **Application Pool** box, click **Select**. In the **Select Application Pool** dialog box, select an application pool from the **Application Pool** list and then click **OK**.
5. In the **Physical path** box, enter the physical path of the website's folder, or click the browse button (**...**) to navigate the file system to find the folder.
6. If the physical path that you entered in the previous step is to a remote share, click **Connect as** to specify credentials that have permission to access the path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option in the **Connect As** dialog box.
7. Select the protocol for the website from the **Type** list.
8. . If you must specify a static IP address for the website (by default, this is set to **All Unassigned**), enter the IP address in the **IP address** box.
9. Enter a port number in the **Port** text box.
10. Optionally, enter a host header name for the website in the **Host Header** box.
11. If you do not have to make any changes to the site, and you want the website to be immediately available, select the **Start Web site immediately** check box.
12. Click **OK**.

To add a website by using the command line

- Use the following syntax at an elevated command prompt or in a script:

Note

For this syntax to work, you either must be in the following directory, or have the directory in your path: `%windir%\system32\inetsrv`

`appcmd add site /name:string /id:`*`uint`* `/physicalPath:`*`string`* `/bindings:`*`string`*

The variable `name` is the name, and the variable `id` is a positive integer that you want to assign to the site. The variables `name` and `id` are the only variables that are required to add a site by using appcmd. However, if you add a site without specifying the values for the `bindings` and `physicalPath` attributes, the site will not be able to start.

The variable `physicalPath` is the absolute path of the site content in the file system.

The variable `bindings` contains information that is used to access the site. It must be in the form *protocol/IP_Address:port:hostheader*. For example, for a website, the binding `http/*:85:` configures the site to listen for HTTP requests on port 85 for all IP addresses and domain names (also known as host headers or host names). On the other hand, a binding of `http/*:85:`**marketing.contoso.com** configures a website to listen for HTTP requests on port 85 for all IP addresses and for the domain name **marketing.contoso.com**.

To add a website named **contoso** with an ID of 2 that has content in **c:\contoso**, and that listens for HTTP requests on port 85 for all IP addresses and a domain name of **marketing.contoso.com**, enter the following at the command prompt:

```
appcmd add site /name:contoso /id:2 /physicalPath:c:\contoso
/bindings:http/*:85:marketing.contoso.com
```

Step 3: Configure Anonymous Authentication

Anonymous authentication gives users access to the public areas of your website without prompting them for a user name or password. You can configure anonymous authentication by using the default anonymous user account (IUSR), or you can set up a local user account for anonymous users.

To configure anonymous authentication by using the UI

1. In **Features View** of IIS Manager, double-click **Authentication**.
2. On the **Authentication** page, select **Anonymous Authentication**.
3. In the **Actions** pane, click **Edit** to set the security principal (user credentials) under which anonymous users will connect to the site.
4. In the **Edit Anonymous Authentication Credentials** dialog box, select one of the following options:
   - If you want to configure a specific user account that IIS uses to access your site or application, select **Specific user**. Then click **Set** to open the **Set Credentials** dialog box, and enter a user name and password for the identity. Then click **OK**.
   - If you want IIS processes to run by using the account that is currently specified on the property page for the application pool, select **Application pool identity**. By default,

this identity is the IUSR account.

Important
If you use the IUSR account, you grant anonymous users all the internal network access associated with that account.

5. Click **OK** to close the **Edit Anonymous Authentication Credentials** dialog box.

To configure anonymous authentication by using the command line

- Use the following syntax to change the default account for anonymous access:

  `appcmd set config /section:anonymousAuthentication /userName:`*string* `/password:`*string*

  The variable `username` is the account that IIS uses for anonymous authentication, and the variable `password` is the password, which is encrypted in the configuration file by default. For example, to use an account named **Moe** and a password of **pssword1** for anonymous access, enter the following at the command prompt:

  ```
  appcmd set config /section:anonymousAuthentication /userName:Moe
  /password:pssword1
  ```

Step 4: Configure the Default Documents

When a client request to your website doesn't include a document name, IIS looks for a file whose name is defined as a default document. Typically, the default document name is Default.htm. You can define a list of default document names in order of precedence.

To configure the default document by using the UI

1. In **Features View** of IIS Manager, double-click **Default Document**.
2. In the **Actions** pane, click **Add**.
3. In the **Name** box, enter the file name that you want to add to the list of default documents and then click **OK**. This file name is added to the top of the default document list.
4. Optionally, select a default document in the list and in the **Actions** pane, click **Move Up** or **Move Down** to change the file's precedence.
5. Optionally, select a default document in the list, and in the **Actions** pane, click **Remove** to remove any file names that you do not want to use as default documents.

To configure the default document by using the command line

- To add a file name to the list of default documents, use the following syntax:

  `appcmd set config /section:defaultDocument /+files.[value='`*string*`']`

  The variable `string` is the file name that you want to add to the list. For example, to add a file named **home.html** to the default document list, enter the following at the command prompt:

  ```
  appcmd set config /section:defaultDocument /+files.[value='home.html']
  ```

To remove a file named **home.html** from the default document list, enter the following at the command prompt, and then press ENTER:

```
appcmd set config /section:defaultDocument /-files.[value='home.html']
```

Step 5: Configure Static Content Compression

You can optionally configure your web server to compress static content to user bandwidth more efficiently and to enhance the performance of your website.

To configure static content compression by using the UI

1. In **Features View** of IIS Manager, double-click **Compression**.
2. Select **Enable static content compression** to configure IIS to compress static content.
3. In the **Static Compression** box, configure the following settings:
   a. Optionally, select **Only compress files larger than (in bytes)** and enter the minimum file size that you want IIS to compress. The default size is 256 bytes.
   b. In the **Cache directory** text box, enter the path of a local directory or click the browse button (**…**) to locate a directory. After a static file is compressed, it is cached in this temporary directory until it expires, or until the content changes. The temporary directory must be on a local drive on an NTFS-formatted partition. The directory cannot be compressed, and should not be shared.
   c. Optionally, select **Per application pool disk space limit (in MB)** and enter the maximum amount of space per application pool, in megabytes, you want IIS to use when it compresses static content. For example, if there are 20 application pools on the server and the **Disk space limit** is set to 100, the maximum disk space will be 2GB. If you click the **Per application pool disk space limit (in MB)** option and enter a number in the text box under it, IIS automatically cleans up the temporary directory according to a least recently used rule when the set limit is reached. The default is 100 MB per application pool.
4. Click **Apply** in the **Actions** pane.

To configure static content compression by using the command line

- To enable HTTP compression of static content, at the command prompt, enter the following command, and then press Enter:

  ```
  appcmd set config /section:urlCompression /doStaticCompression:True
  ```

- To configure static-content compression settings, use the following syntax:

  `appcmd set config /section:urlCompression /minFileSizeforComp:`*int* `/directory:`*string* `/maxDiskSpace:`*int*

  The variable `minFileSizeforComp` sets the minimum number of bytes that a file must contain for it to be compressed. The default value is 256. The variable `directory` specifies the directory where compressed versions of static files are temporarily stored and cached. The default is the following folder:

`%SystemDrive%\inetpub\temp\IIS Temporary Compressed Files`

The variable `maxDiskSpace` sets the maximum amount of space per application pool, in megabytes, that you want IIS to use when it compresses static content. The default is 100 MB per application pool.

Next Steps

Test your website thoroughly to ensure that it functions as expected. Then consider configuring the following features.

- To help you troubleshoot or optimize the performance of your web server, set up IIS logging. For instructions see, Configure Logging in IIS.
- To improve the security of your web server, configure request filtering. For instructions see, Configure Request Filtering in IIS.

See also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build an ASP.NET Website on IIS
- Build a Classic ASP Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Configure Logging in IIS

You can configure logging on your web server or website that records information about HTTP requests and errors. The information in your log can help you troubleshoot or optimize your website.

Choose one or more of the following procedures to set up logging appropriately for your needs:

- Prerequisites
- Configure Logging at the Site Level
- Configure Per-site Logging at the Server Level
- Configure Per-server Logging at the Server Level
- Select W3C Fields to Log
- Configure Log File Rollover Options

Prerequisites

This guide was written for, and tested on, the following operating systems:

- Windows Server® 2012 R2

- Windows Server® 2012
- Windows® 8

Configure Logging at the Site Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure logging at the site level by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** tree view, select your website.
3. In **Features View**, double-click **Logging**.
4. On the **Logging** page, in the **Log file** section under **Format**, select one of the following log file formats:
   - **IIS**: to use the Microsoft IIS log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by commas, and time is recorded as local time. For more information about the IIS log file format, see IIS Log File Format (IIS 6.0).
   - **NCSA**: to use the National Center for Supercomputing Applications (NCSA) Common log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by spaces, and time is recorded as local time with the Coordinated Universal Time (UTC) offset. For more information about the NCSA log file format, see NCSA Common Log File Format (IIS 6.0).
   - **W3C**: to use the centralized W3C log file format to log information about all sites on the server. This format is handled by HTTP.sys, and is a *customizable* ASCII text-based format, which means that you specify the fields that are logged. Specify the fields that are logged on the **W3C Logging Fields** dialog box by clicking **Select Fields** on the **Logging** page. Fields are separated by spaces, and time is recorded in Coordinated Universal Time (UTC). For more information about the W3C log file format, see W3C Extended Log File Format (IIS 6.0).

     Note
     In Windows Server 2012 R2 and later, you can write log information to a log file, as Event Tracing for Windows (ETW) events, or both. For more

information on ETW, see Event Tracing.

- **Custom**: to use a custom format for a custom logging module. When you select this option, the **Logging** page becomes disabled, because custom logging cannot be configured in IIS Manager. For more information about how to use custom log file formats, see Custom Logging Modules (IIS 6.0).

5. Under **Directory**, specify the path where the log file should be stored. The default is *%SystemDrive%*\inetpub\logs\LogFiles

   Note
   As a best practice, store log files, such as failed request trace logs, in a directory other than *systemroot*.

6. In the **Log File Rollover** section, select one of the following options:
   - **Schedule**: to create new log file that is based on one of the following values:
   - **Hourly**: a new log file is created each hour.
   - **Daily**: a new log file is created each day.
   - **Weekly**: a new log file is created each week.
   - **Monthly**: a new log file is created each month.
   - **Maximum file size (in bytes)**: to create a log file when the file reaches a certain size (in bytes).The minimum file size is 1048576 bytes. If this attribute is set to a value less than 1048576 bytes, the default value is implicitly assumed as 1048576 bytes.
   - **Do not create a new log file**: there is a single log file that continues to grow as information is logged.
7. Select **Use local time for file naming and rollover** to specify that log file naming and time for log file rollover uses the local server time. When this option is not selected, Coordinated Universal Time (UTC) is used.

   Note
   Regardless of this setting, timestamps in the actual log file will use the time format for the log format that you select from the **Format** list. For example, NCSA and W3C log file formats use UTC time format for timestamps.

8. Click **Apply** in the **Actions** pane.

Configure Per-site Logging at the Server Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure per-site logging at the server level by using the UI

1. In the **Connections** tree view of IIS Manager, select your web server.
2. In **Features View**, double-click **Logging**.
3. On the **Logging** page under **One log file per site**, select **Site** from the drop-down list. By default, **Site** is selected.

4. Complete this procedure by following the site-level procedure by starting with step 4.

Configure Per-server Logging at the Server Level

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure per-server logging at the server level by using the UI

1. In the **Connections** tree view of IIS Manager, select your web server.
2. In **Features View**, double-click **Logging**.
3. On the **Logging** page, under **One log file per site**, select **Server** from the drop-down list. By default, **Site** is selected.
4. Complete this procedure by following the site-level procedure by starting with step 4.

Select W3C Fields to Log

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To select W3C fields to log by using the UI

1. In **Features View** of IIS Manager, double-click **Logging**.
2. On the **Logging** page, in the **Log file** section under **Format**, click **Select Fields**.
3. In the **W3C Logging Fields** dialog box, select one or more of the following options:
   - **Date (date)**: the date on which the request occurred.
   - **Time (time)**: the time, in Coordinated Universal Time (UTC), at which the request occurred.
   - **Client IP Address (c-ip)**: the IP address of the client that made the request.
   - **User Name (cs-username)**: the name of the authenticated user who accessed your server. Anonymous users are indicated by a hyphen.
   - **Service Name (s-sitename)**: the site instance number that fulfilled the request.
   - **Server Name (s-computername)**: the name of the server on which the log file entry was generated.
   - **Server IP Address (s-ip)**: the IP address of the server on which the log file entry was generated.
   - **Server Port (s-port)**: the server port number that is configured for the service.
   - **Method (cs-method)**: the requested action, for example, a GET method.
   - **URI Stem (cs-uri-stem)**: the Universal Resource Identifier, or target, of the action.
   - **URI Query (cs-uri-query)**: the query, if any, that the client was trying to perform. A Universal Resource Identifier (URI) query is necessary only for dynamic pages.
   - **Protocol Status (sc-status)**: the HTTP or FTP status code.

- **Protocol Sub-status (sc-substatus)**: the HTTP or FTP substatus code.
- **Win32 Status (sc-win32-status)**: the Windows status code.
- **Bytes Sent (sc-bytes)**: the number of bytes that the server sent.
- **Bytes Received (cs-bytes)**: the number of bytes that the server received.
- **Time Taken (time-taken)**: the length of time that the action took in milliseconds.
- **Protocol Version (cs-version)**: the protocol version, HTTP or FTP, that the client used.
- **Host (cs-host)**: the host name, if any.
- **User Agent (cs(UserAgent))**: the browser type that the client used.
- **Cookie (cs(Cookie))**: the content of the cookie sent or received, if any.
- **Referer (cs(Referer))**: the site that the user last visited. This site provided a link to the current site.

Note
In Windows Server 2012 R2 and later, you can specify additional custom fields to log from the HTTP request and response headers, and from server variables. To add custom fields, select **Add Field** in the **W3C Logging Fields** dialog box. Enhanced logging is available only for site-level logging - if you selected server-level logging, then **Add Field** is disabled. Note that if the total size of the custom fields that you define is greater than 64K bytes, the logged content is truncated to 64K bytes.

4. Click **Apply** in the **Actions** pane.

Configure Log File Rollover Options

You can perform this procedure by using the user interface (UI), or by editing configuration files directly.

To configure log file rollover options by using the UI

1. In **Features View** of IIS Manager, double-click **Logging**.
2. On the **Logging** page, in the **Log File Rollover** section, select one of the following options:
   - **Schedule**: to create new log file that is based on one of the following values:
   - **Hourly**: a new log file is created each hour.
   - **Daily**: a new log file is created each day.
   - **Weekly**: a new log file is created each week.
   - **Monthly**: a new log file is created each month.
   - **Maximum file size (in bytes)**: to create a log file when the file reaches a certain size (in bytes).The minimum file size is 1048576 bytes. If this attribute is set to a value less than 1048576 bytes, the default value is implicitly assumed as 1048576 bytes.
   - **Do not create a new log file**: This option means that there is a single log file that

continues to grow as information is logged. If you use a single log file for your site, it is helpful when you use log parsing utilities, but it also creates larger log files that could affect the overall performance of the server.

3. Select **Use local time for file naming and rollover** to specify that log file naming and time for log file rollover uses the local server time. When this option is not selected, Coordinated Universal Time (UTC) is used.

   Note
   Regardless of this setting, timestamps in the actual log file will use the time format for the log format that you select from the **Format** list. For example, NCSA and W3C log file formats use UTC time format for timestamps.

4. Click **Apply** in the **Actions** pane.

See Also

- Build an ASP.NET Website on IIS
- Build a Classic ASP Website on IIS
- Build a Static Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Configure Request Filtering in IIS

This document shows you how to use common request-filter settings to improve the security of your IIS 8 web server.

Request filters restrict the types of HTTP requests that IIS 8 processes. By blocking specific HTTP requests, request filters help prevent potentially harmful requests from reaching the server. The request filter module scans incoming requests and rejects requests that are unwanted based upon the rules that you set up.

By default, IIS rejects requests to browse critical code segments. It also rejects requests for some file name extensions.

You can configure a request filter at the server-wide level and then override the configuration at a website level.

- Prerequisites
- General Request Filter Settings
- File Name Extensions
- Filtering Rules
- Hidden Segments

- URL Filtering
- HTTP Verbs
- Header Size Limits
- Query Strings
- Request Filter Logging

Prerequisites

This guide was written for, and tested on, the following operating systems:

- Windows Server® 2012
- Windows® 8

General Request Filter Settings

The general settings include such settings as the following:

1. Whether to allow access to a file with an extension that is not listed for request filter.
2. Whether to allow requests that use HTTP verbs that are not listed.
3. Whether to allow requests that contain high-bit characters (non-ASCII).
4. Whether to allow requests that are double encoded.
5. Maximum length of the content requested.
6. Maximum length of the URL
7. Maximum size of a query string

To configure general request-filter options by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. In the **Actions** pane, click **Edit Feature Settings**.
4. In the **Edit Request Filtering Settings** dialog, edit the settings as desired, and then click **OK**.

To configure general request-filter options by using the command line

- To configure high-bit characters, use the following syntax:

  appcmd set config /section:requestfiltering /allowhighbitcharacters:true | false

  For example, to allow high-bit characters, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /allowhighbitcharacters:true
- To configure double escaping, use the following syntax:

  appcmd set config /section:requestfiltering /allowdoubleescaping:true | false

  For example, to enable double escaping, type the following at the command prompt, and

then press ENTER:

appcmd set config /section:requestfiltering /allowdoubleescaping:true

- To configure a value for the maximum allowed length of content, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxallowedcontentlength:*unit*

  The variable **requestlimits.maxallowedcontentlength** *unit* specifies the maximum length of content.

  For example, to specify 30000000 as the maximum length of content, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /requestlimits.maxallowedcontentlength:30000000

- To configure a value for the maximum allowed length of an incoming URL, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxurl:*unit*

  The variable **requestlimits.maxurl** *unit* specifies the maximum length of an incoming URL.

  For example, to specify 4096 as the maximum incoming URL length, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /requestlimits.maxurl:4096

- To configure a value for the maximum allowed length of an incoming query string, use the following syntax:

  appcmd set config /section:requestfiltering /requestlimits.maxquerystring:*unit*

  The variable **requestlimits.maxquertystring** *unit* specifies the maximum length of an incoming query string.

  For example, to specify 2048 as the maximum incoming query string, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /requestlimits.maxquerystring:2048

- To configure a size limit for a specific HTTP header, use the following syntax:

  appcmd set config /section:requestfiltering /+requestlimits.headerLimits.[header='*string***',sizelimit='***unit***']**

  The variable **header** *string* specifies the header this restriction applies to. The variable **sizelimit** *unit* specifies the maximum size of this header.

  For example, to specify a maximum size of 2048 for headers that include a value of contoso.com, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+requestlimits.headerLimits.[header='contoso.com',sizelimit='2048']

File Name Extensions

For each file name extension you add, you can indicate whether to allow or reject requests for that type of file.

To configure file name extensions by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **File Name Extensions** tab.
4. In the **Actions** pane, click either **Allow File Name Extension** or **Deny File Name Extension**.
5. Type the file name extension in the box, and then click **OK**.

To configure file name extensions by using the command line

- To configure how IIS deals with unlisted file name extensions, use the following syntax:

  appcmd set config /section:requestfiltering /fileExtensions.allowunlisted:true | false

  For example, to deny unlisted file name extensions, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /fileExtensions.allowunlisted:false
- To configure whether file name extensions apply to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /fileExtensions.applyToWebDAV:true | false

  For example, to configure IIS so that file name extensions do not apply to WebDAV requests, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /fileExtensions.applyToWebDAV:false
- To add a file name extension, use the following syntax:

  appcmd set config /section:requestfiltering /+fileExtensions.[fileextension='*.string***',allowed='true | false']**

  The variable **fileextension** *string* is the file name extension you want to allow or deny.

  For example, to add an allow rule for the file name extension .xxx, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+fileExtensions.[fileextension='*.xxx***',allowed='true']**
- To remove a rule for the file name extension .xxx, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /-fileExtensions.[fileextension='*.xxx***']**

Filtering Rules

IIS 8 permits you to define custom filter rules that apply to incoming requests. Using this feature, you can define filters that can do the following:

1. Scan the request URL.
2. Scan for query strings contained in the URL.
3. Scan for specific header fields.
4. Define what file name extensions the filter applies to.
5. Define strings you want to deny.

To configure a filtering rule by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Rules** tab.
4. In the **Actions** pane, double-click **Add Filtering Rule**.
5. In the **Name** box, type a name for the filtering rule.
6. If you want the URL scanned, select the **Scan Url** check box.
7. If you want the query string scanned, select the **Scan query string** check box.
8. Under **Scan Headers**, type one or more headers to scan for.
9. Under **Applies To**, type one or more file name extensions that the rule applies to.
10. Under the **Deny Strings**, type one or more strings you want to deny.
11. Click **OK**.

Hidden Segments

This feature allows you to reject requests that contain a URL segment (for example, a folder name).

To configure hidden segments by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Hidden Segments** tab.
4. In the **Actions** pane, click **Add Hidden Segment**.
5. Type the URL segment in the box, and then click **OK**.

To configure hidden segments by using the command line

- To configure whether hidden segments apply to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /hiddensegments.applyToWebDAV:true

| false

For example, to configure IIS so that hidden segments do not apply to WebDAV requests, type the following at the command prompt, and then press ENTER:

appcmd set config /section:requestfiltering /hiddensegments.applyToWebDAV:false

- To configure a hidden segment, use the following syntax:

  appcmd set config /section:requestfiltering /+hiddensegments.[segment='*string***']**

  The variable **segment** *string* specifies a URL segment that is hidden.

  For example, to specify that /bin is a hidden segment, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+hiddensegments.[segment='/bin']

URL Filtering

You can configure IIS to accept a specified URL. In addition, you can configure it to deny a specified URL sequence.

To configure URL filtering by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **URL** tab.
4. In the **Actions** pane, select either **Allow URL** or **Deny Sequence**.
5. Type the URL or the URL sequence in the box, and click **OK**.

To configure URL filtering by using the command line

- To deny a URL sequence, use the following syntax:

  appcmd set config /section:requestfiltering /+denyurlsequences.[sequence='*string***']**

  The variable **sequence** *string* specifies a sequence of characters in a URL that IIS is never allowed to parse.

  For example, to specify that IIS never parse URLs that contain two periods, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+denyurlsequences.[sequence='..']

HTTP Verbs

You can define a list of verbs that IIS 8 accepts as part of a request. When IIS rejects a request based on this feature, the error code logged is 404.6.

To configure unlisted HTTP verbs by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **HTTP Verbs** tab.
4. In the **Actions** pane, click either **Allow Verb** or **Deny Verb**.
5. Enter the verb in the box, and then click **OK**.

To configure unlisted HTTP verbs by using the command line

- To configure how IIS deals with unlisted verbs, use the following syntax:

  appcmd set config /section:requestfiltering /verbs.allowunlisted:true | false

  For example, to deny unlisted verbs, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /verbs.allowunlisted:false
- To configure whether verb filtering applies to WebDAV requests, use the following syntax:

  appcmd set config /section:requestfiltering /verbs.applyToWebDAV:true | false

  For example, to configure IIS so that verb filtering does not apply to WebDAV requests, type the following at the command prompt and then press ENTER:

  appcmd set config /section:requestfiltering /verbs.applyToWebDAV:false
- To configure a verb to filter, use the following syntax:

  appcmd set config /section:requestfiltering /+verbs.[verb='*string***',allowed='true | false']**

  The variable **verb** *string* specifies the verb this restriction applies to.

  For example, to specify GET is allowed, type the following at the command prompt, and then press ENTER:

  appcmd set config /section:requestfiltering /+verbs.[verb='GET',allowed='true']

Header Size Limits

You can limit the size of HTTP request headers to improve performance and security. Headers are name/value pairs that define the operating parameters of an HTTP transaction.

To configure header size limits by using the UI

1. Open IIS Manager and select the level for which you want to configure request filter.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Headers** tab, and click **Add Header**.
4. In the **Header** box, type the header field name.
5. In the **Size Limit** box, type a positive integer that represents the header size limit in bytes.

6. Click **OK**.

Query Strings

You can configure IIS 8 to allow or deny specific query strings contained in the requested URL. For example, if a denied query string is found in a request URL, the request is denied.

To configure query strings by using the UI

1. Open IIS Manager and select the level you want to configure request filter for.
2. In **Features View**, double-click **Request Filtering**.
3. Select the **Query Strings** tab, and click either **Allow Query String** or **Deny Query String**.
4. In the **Query string** box, type the query string.
5. Click **OK**.

Request Filter Logging

You can use IIS logging to evaluate and optimize your request filter configuration.

The following table shows the request filter error codes that you see in the log:

Error Description	Status Code
Request Filtering: URL Sequence denied	404.5
Request Filtering: Verb denied	404.6
Request Filtering: File name extension denied	404.7
Request Filtering: Denied by hidden segment	404.8
Request Filtering: Denied because request header is too long.	404.10
Request Filtering: Denied because URL doubled escaping	404.11
Request Filtering: Denied because of high bit characters	404.12
Request Filtering: Denied because content length too large	404.13
Request Filtering: Denied because URL too long	404.14
Request Filtering: Denied because query string	404.15

Error Description	Status Code
too long	

See Also

- Build an ASP.NET Website on IIS
- Build a Classic ASP Website on IIS
- Build a Static Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS ServersBuild a web farm using IIS servers
- Hosting-Friendly Web Server Platform (IIS): Scenario Overview

Build a Classic ASP Website on IIS

This document guides you through installing IIS and configuring a classic ASP website. Classic ASP is a server-side scripting environment that you can use to create and run dynamic web applications. With ASP, you can combine HTML pages, script commands, and COM components to create interactive web pages that are easy to develop and modify. Classic ASP is the predecessor to ASP.NET, but it is still in wide use today.

The Classic ASP server configuration adds IIS modules for ASP and ISAPI extensions to the default IIS installation.

In this document

- Prerequisites
- Step 1: Install the IIS Web Server
- Step 2: Add a Classic ASP Website
- Step 3: Edit ASP Application Settings
- Next Steps

Prerequisites

This guide was written for, and tested on, the following operating systems:

1. Windows Server® 2012
2. Windows® 8

Step 1: Install the IIS Web Server

You can use the Web Platform Installer (Web PI) to install IIS, and applications that run on IIS. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

You can also perform this procedure by using the Windows user interface (UI) or from a command line.

To install IIS on Windows Server 2012 by using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before You Begin** page, click **Next**.
4. On the **Select Installation Type** page, select **Role-based or Feature-based Installation** and click **Next**
5. On the **Select Destination Server** page, select **Select a server from the server pool**, select your server, and click **Next**.
6. On the **Select Server Roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Select Features** page, note the preselected features that are installed by default, and then select the following additional role services:
   - ASP
   - ISAPI Extensions
8. Click **Next**.
9. On the **Web Server Role (IIS)** page, click **Next**.
10. On the **Select Role Services** page, note the preselected role services that are installed by default, and then click **Next**.

    Note
    You only have to install the IIS 8 default role services for a static-content web server.
11. On the **Confirm Installation Selections** page, confirm your selections, and then click **Install**.
12. On the **Installation Progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
13. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    You should see the default IIS Welcome page.

To install IIS on Windows 8 by using the UI

1. On the **Start** page, type **Control Panel**, and then click the **Control Panel** icon in the search results.

2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services**, note the preselected features that are installed by default, and then select the following additional role services:
   - ASP
   - ISAPI Extensions
4. Click **OK**.
5. To verify that IIS installed successfully, type the following into a web browser:

   http://localhost

   You see the default IIS Welcome page.

To Install IIS by using the command line

- Type the following command at a command prompt or into a script:

  Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-ApplicationDevelopment;IIS-ASP;IIS-ISAPIExtensions;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI

Step 2: Add a Classic ASP Website

You can perform this procedure by using the user interface (UI), by running Appcmd.exe commands in a command-line window, by editing configuration files directly, or by writing WMI scripts.

To add a website by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. On the **Server Manager Dashboard**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, right-click the **Sites** node in the tree, and then click **Add Website**.
3. In the **Add Website** dialog box, type a friendly name for your website in the **Site name** box.

4. If you want to select a different application pool than the one listed in the **Application Pool** box, click **Select**. In the **Select Application Pool** dialog box, select an application pool from the **Application Pool** list and then click **OK**.
5. In the **Physical path** box, type the physical path of the Web site's folder, or click the browse button (**...**) to navigate the file system to find the folder.
6. If the physical path that you entered in step 5 is to a remote share, click **Connect as** to specify credentials that have permission to access the path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option in the **Connect As** dialog box.
7. Select the protocol for the Web site from the **Type** list.
8. The default value in the **IP address** box is **All Unassigned**. If you must specify a static IP address for the Web site, type the IP address in the **IP address** box.
9. Type a port number in the **Port** text box.
10. Optionally, type a host header name for the Web site in the **Host Header** box.
11. If you do not have to make any changes to the site, and you want the Web site to be immediately available, select the **Start Web site immediately** check box.
12. Click **OK**.

To add a website by using the command line

- Use the following syntax at the command prompt or in a script:

  Note
  For this syntax to work, you either must be in the following directory, or have the directory in your path: `%windir%\system32\inetsrv`

  appcmd add site /name:*string* **/id:***uint* **/physicalPath:***string* **/bindings:***string*

  The variable **name** *string* is the name, and the variable **id** *uint* is the unsigned integer that you want to assign to the site. The variables **name** *string* and **id** *uint* are the only variables that are required when you add a site in Appcmd.exe.

  Note
  When you add a site without specifying the values for the **bindings** and **physicalPath** attributes, the site will not be able to start.

  The variable **physicalPath** *string* is the path of the site content in the file system.

  The variable **bindings** *string* contains information that is used to access the site, and it should be in the form of *protocol*/*IP_address*:*port*:*host_header*. For example, a web site binding is the combination of protocol, IP address, port, and host header. A binding of **http/*:85:** enables a web site to listen for HTTP requests on port 85 for all IP addresses and domain names (also known as host headers or host names). On the other hand, a binding of **http/*:85:marketing.contoso.com** enables a web site to listen for HTTP requests on port 85 for all IP addresses and the domain name **marketing.contoso.com**.

  To add a web site named **contoso** with an ID of 2 that has content in **c:\contoso**, and that listens for HTTP requests on port 85 for all IP addresses and a domain name of

marketing.contoso.com, type the following at the command prompt, and then press ENTER:

appcmd add site /name:contoso /id:2 /physicalPath:c:\contoso /bindings:http/*:85:marketing.contoso.com

Step 3: Edit ASP Application Settings

IIS 8 provides default settings for ASP applications, but you can change those settings as needed. For example, you can enable client-side debugging on a test server to aide in troubleshooting issues during a test pass.

To edit ASP application settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **ASP**.
3. On the **ASP** page, edit settings as desired.
4. When finished, click **Apply** in the **Actions** pane.

To edit ASP application settings by using the command line

- **Specify default character set**

  To specify the default character set for an application, use the following syntax:

  appcmd set config /section:asp /codePage:*integerRange*

  The variable *integerRange* is the default character set. For example, to set the code page to a Latin character set used in American English and many European alphabets, type the following at the command prompt, and then press Enter:

  appcmd set config /section:asp /codePage:*1252*
- **Enable or disable buffering**

  To enable or disable buffering of ASP application output, use the following syntax:

  appcmd set config /section:asp /bufferingOn:True|False

  True enables buffering whereas **False** disables buffering. The default value is **True**.
- **Enable or disable HTTP 1.1 chunked transfer encoding**

  To enable HTTP 1.1 chunked transfer encoding for the World Wide Web publishing service, use the following syntax:

  appcmd set config /section:asp /enableChunkedEncoding:True|False

  True enables HTTP 1.1 chunked transfer encoding whereas **False** disables HTTP 1.1 chunked transfer encoding. The default value is **True**.
- **Enable or disable HTML fallback**

  To enable or disable HTML fallback, use the following syntax:

  appcmd set config /section:asp /enableASPHTMLFallback:True|False

True causes an .htm file that has the same name as the requested .asp file, if it exists, to be sent instead of the .asp file if the request is rejected because of a full request queue. The default value is **True**.

- **Enable or disable parent paths**

  To enable or disable paths relative to the current directory or above the current directory, use the following syntax:

  appcmd set config /section:asp /enableParentPaths:True|False

  True sets ASP pages to allow paths relative to the current directory or above the current directory. The default value is **True**.

- **Set client connection test interval**

  To set a time interval after which ASP will check to see whether the client is still connected before executing a request, use the following syntax:

  appcmd set config /section:asp /queueConnectionTestTime:*timeSpan*

  The variable *timeSpan* sets the time interval (hh:mm:ss) after which ASP checks to see whether the client is still connected before executing a request. The default value is 00:00:03.

- **Set maximum requesting entity body limit**

  To specify the maximum number of bytes allowed in the entity body of an ASP request, use the following syntax:

  appcmd set config /section:asp /maxRequestEntityAllowed:*int*

  The variable *int* represents the maximum number of bytes allowed in the body of an ASP request. The default value is 200000 bytes.

- **Set request queue length**

  To specify the maximum number of concurrent ASP requests allowed into the queue, use the following syntax:

  appcmd set config /section:asp /requestQueueMax:*int*

  The variable *int* represents the maximum number of concurrent ASP requests that are allowed into the request queue. The default value is 3000.

- **Set request queue time-out**

  To specify the period that an ASP request can wait in the request queue, use the following syntax:

  appcmd set config /section:asp /queueTimeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that an ASP request can wait in the request queue. The default value is 00:00:00.

- **Specify response buffering limit**

  To control the maximum number of bytes that an ASP page can write to the response buffer before a flush occurs, use the following syntax:

  appcmd set config /section:asp /bufferingLimit:*int*

  The variable *int* represents the maximum size, in bytes, of the ASP buffer. The default

value is 4194304 bytes.

- **Set script time-out**

  To specify the default length of time that ASP pages let a script run before terminating the script and writing an event to the Windows Event Log, use the following syntax:

  appcmd set config /section:asp /scriptTimeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that an ASP request can run before an event is written to the Windows Event Log. The default value is 00:01:30.

- **Specify threads per processor limit**

  To specify the maximum number of worker threads per processor that ASP can create, use the following syntax:

  appcmd set config /section:asp /processorThreadMax:*int*

  The variable *int* represents the maximum number of worker threads per processor that ASP can create. The default value is 25.

- **Specify default locale identifier**

  To define how dates, times, and currencies are formatted for an ASP application, use the following syntax:

  appcmd set config /section:asp /lcid:*int*

  The variable *int* represents the default locale identifier for an ASP application. The default value is 0.

- **Enable or disable automatic application restart**

  To enable or disable automatic restart of ASP applications whenever a configuration setting is changed, use the following syntax:

  appcmd set config /section:asp /enableApplicationRestart:True|False

  True enables ASP applications to be automatically restarted whenever a configuration setting is changed. The default value is **True**.

- **Enable or disable line number calculation**

  To enable or disable ASP to calculate and store the line number of each executed line of code to provide the number in an error report, use the following syntax:

  appcmd set config /section:asp /calLineNumber:True|False

  True enables line number calculation and storage. The default value is **True**.

- **Enable or disable COM component exception trapping**

  To enable or disable ASP pages to catch exceptions thrown by COM components, use the following syntax:

  appcmd set config /section:asp /exceptionCatchEnable:True|False

  True enables COM component exception trapping. If set to **False**, the Microsoft Script Debugger tool does not catch exceptions sent by the component that you are debugging. The default value is **True**.

- **Enable or disable client-side debugging**

  To enable or disable client-side debugging, use the following syntax:

  appcmd set config /section:asp /appAllowClientDebug:True|False

  True enables client-side debugging. The default value is **False**.
- **Enable or disable log error requests**

  To enable or disable the writing of ASP errors to the application section of the Windows event log, use the following syntax:

  appcmd set config /section:asp /logErrorRequests:True|False

  True enables log error requests. By default, ASP errors are written to the client browser and the IIS logs. The default value is **True**.
- **Enable or disable server-side debugging**

- **Enable or disable Windows event logging of ASP errors**

  To enable or disable ASP debugging on the server, use the following syntax:

  appcmd set config /section:asp /appAllowDebugging:True|False

  True enables server-side debugging for ASP applications. The default value is **False**.
- **Run On End Functions Anonymously**

  To enable or disable **SessionOnEnd** and **ApplicationOnEnd** global ASP functions to run as the anonymous user, use the following syntax:

  appcmd set config /section:asp /runOnEndAnonymously:True|False

  True enables **SessionOnEnd** and **ApplicationOnEnd** global ASP functions to run as the anonymous user. The default value is **True**.
- **Specify script error message**

  To specify the error message to send to the browser if specific debugging errors are not sent to the client, use the following syntax:

  appcmd set config /section:asp /scriptErrorMessage:*string*

  The variable *string* represents the error message that is sent to the browser when specific debugging errors are not sent to the client. The default value is "An error occurred on the server when processing the URL. Please contact the system administrator".
- **Enable or disable sending errors to the browser**

  To enable or disable the writing of debugging specifics (file name, error, line number, and description) to the client browser in addition to logging them to the Windows event log, use the following syntax:

  appcmd set config /section:asp /scriptErrorSentToBrowser:True|False

  True enables the writing of debugging specifics to the client browser. The default value is **False**.
- **Specify default script language**

To specify the default script language for all ASP applications that are running on the web server, use the following syntax:

appcmd set config /section:asp /scriptLanguage:*string*

The variable *string* represents the default script language. The default value is VBScript.

- **Specify cache directory path**

  To specify the name of the directory where ASP stores compiled ASP templates when the in-memory cache overflows, use the following syntax:

  appcmd set config /section:asp /diskTemplateCacheDirectory:*string*

  The variable *string* represents the cache directory path. The default value is *%windir%*\system32\inetsrv\ASP Compiled Templates.

- **Enable or disable type library caching**

  To enable or disable the caching of type libraries, use the following syntax:

  appcmd set config /section:asp /enableTypelibCache:True|False

  True enables the caching of type libraries. The default value is **True**.

- **Set maximum number of compiled ASP templates to store**

  To set the maximum number of compiled ASP templates that can be stored, use the following syntax:

  appcmd set config /section:asp /maxDiskTemplateCacheFiles:*int*

  The variable *int* represents the maximum number of compiled ASP templates to store. The default value is 2000.

- **Set maximum number of compiled ASP templates to store**

  To set the maximum number of precompiled script files to cache, use the following syntax:

  appcmd set config /section:asp /scriptFileCacheSize:*int*

  The variable *int* represents the number of precompiled script files to cache. If set to 0, no script files are cached. If set to 4294967295, all requested script files are cached. The default value is 500.

- **Set maximum number of scripting engines to cache**

  To set the maximum number of scripting engines that ASP pages keep cached in memory, use the following syntax:

  appcmd set config /section:asp /scriptEngineCacheMax:*int*

  The variable *int* represents the maximum number of scripting engines that are cached. The default value is 250.

- **Enable or disable COM+ side-by-side assemblies**

  To enable or disable side-by-side COM+ assemblies, which allow ASP applications to specify which version of a system DLL or classic COM component to use, use the following syntax:

  appcmd set config /section:asp /appServiceFlags:True|False

True enables COM+ side-by-side assemblies. The default value is **False**.

- **Enable or disable COM+ tracker**

  To enable or disable COM+ tracker, use the following syntax:

  appcmd set config /section:asp /enableTypelibCache:True|False

  True enables COM+ tracker, which lets administrators or developers debug ASP applications. The default value is **False**.

- **Enable or disable multithreaded environments**

  To enable or disable ASP to run in a multithreaded environment, use the following syntax:

  appcmd set config /section:asp /executeInMta:True|False

  True enables ASP to run in a multithreaded environment. The default value is **False**.

- **Enable or disable thread model checking**

  To enable or disable whether IIS checks the threading model of any component that your application creates, use the following syntax:

  appcmd set config /section:asp /trackThreadingModel:True|False

  True enables thread model checking. The default value is **False**.

- **Specify COM+ partition ID**

  To specify the Globally Unique Identifier (GUID) of the COM+ partition, use the following syntax:

  appcmd set config /section:asp /partitionID:*string*

  The variable *string* represents the GUID of the COM+ partition. The default value is 00000000-0000-0000-0000-000000000000.

  Note

  You must also set the **appServiceFlags** flag to **True**.

- **Specify COM+ application**

  To specify the name of the COM+ application, use the following syntax:

  appcmd set config /section:asp /sxsName:*string*

  The variable *string* represents name of the COM+ application.

  Note

  You must also set the **appSeviceFlags** flag to **True**.

- **Enable or disable COM+ partitioning**

  To enable or disable COM+ partitioning, use the following syntax:

  appcmd set config /section:asp /appServiceFlags:True|False

  True enables COM+ partitioning, which can be used to isolate applications in their own COM+ partition. The default value is **False**.

  Note

If set to **True**, you must also set a value for the **partitionID** property.

- **Enable or disable session state**

  To enable or disable session state persistence for an ASP application, use the following syntax:

  appcmd set config /section:asp /allowSessionState:True|False

  True enables session state persistence. The default value is **True**.

- **Set maximum number of concurrent sessions**

  To set the maximum number of concurrent sessions that ASP allows, use the following syntax:

  appcmd set config /section:asp /max:*int*

  The variable *int* represents the maximum number of concurrent sessions. The default value is -1.

- **Enable or disable secure session ID**

  To enable or disable the sending of a session ID as a secure cookie if assigned over a secure session channel, use the following syntax:

  appcmd set config /section:asp /keepSessionIdSecure:True|False

  True enables secure session ID. The default value is **True**.

- **Set session time-out**

  To specify the default time that a session object is maintained after the last request associated with the object is made, use the following syntax:

  appcmd set config /section:asp /timeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that a session object is maintained after the last request associated with the object is made. The default value is 00:20:00.

Next Steps

Test your website thoroughly to ensure that it functions as expected. Then consider configuring the following features.

- To help you troubleshoot or optimize the performance of your web server, set up IIS logging. For instructions see, Configure Logging in IIS.
- To improve the security of your web server, configure request filtering. For instructions see, Configure Request Filtering in IIS.

See also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build a Static Website on IIS
- Build an ASP.NET Website on IIS

- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Build an ASP.NET Website on IIS

This document contains an overview of the **Build an ASP.NET Website on IIS** scenario. It also contains links to additional information and community resources related to the scenario.

Did you mean…

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an FTP Site on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

Scenario Description

This scenario shows how to plan and configure an ASP.NET website on an IIS 8 webserver. It is divided into two phases: a plan and design phase, and an install and configure phase. In the plan and design phase, you are provided the information needed to make informed decisions about web server installation, ASP.NET settings, data source settings, and basic application security. In the install and configure phase, you are guided through the procedures required to install IIS, add an ASP.NET application, and configure IIS.

This scenario does not cover how to write an ASP.NET application.

In This Scenario

- Plan an ASP.NET Website on IIS
  - Step 1: Plan IIS Web Server and ASP.NET Modules Installation
  - Step 2: Plan ASP.NET Settings
  - Step 3: Plan Data Source Settings
  - Step 4: Plan Application Security
- Configure an ASP.NET Website on IIS
  - Step 1: Install IIS and ASP.NET Modules
  - Step 2: Configure ASP.NET Settings
  - Step 3: Configure Data Source Settings
  - Step 4: Configure Application Security

Practical Applications

Whether you are an IT professional, a web developer, or you just want to set up your own webserver, this scenario can help you install IIS and configure it to serve your ASP.NET web application.

Software Requirements

To get the most from this scenario, you must have access to a computer running one of the following operating systems:

- Windows Server® 2012
- Windows® 8

See Also

The following table contains links to resources related to this scenario.

Content type	References
Deployment	Deployment to a Hosting Provider | Web Deploy 2.0
Operations	IIS.NET | IIS Learning Center
Tools and Settings	Web Server (IIS) Administration Cmdlets in Windows PowerShell
Community Resources	IIS Blogs | IIS Forums | Robert McMurray's Blog | Scott Forsyth's Blog | Steve Schofield's Blog
Related Technologies	ASP.NET | ASP.NET Web Projects

Plan an ASP.NET Website on IIS

To develop a plan for installing an IIS web server and configuring it for ASP.NET web applications, follow the steps listed.

- Step 1: Plan IIS Web Server and ASP.NET Modules Installation
- Step 2: Plan ASP.NET Settings
- Step 3: Plan Data Source Settings
- Step 4: Plan Application Security

After you have reviewed these planning steps, see Configure an ASP.NET Website on IIS. For more information, see Build an ASP.NET Website on IIS.

Step 1: Plan IIS Web Server and ASP.NET Modules Installation

The first step in planning to build an ASP.NET website on IIS 8 is to install IIS along with the ASP.NET modules and to add your application files to IIS.

The following list shows the tasks required to complete this step:

1.1. Plan to Install IIS and ASP.NET Modules

1.2. Plan to Add the ASP.NET Application

When you are done with these tasks, record your design decisions before going on to Step 2: Plan ASP.NET Settings.

1.1. Plan to Install IIS and ASP.NET Modules

An ASP.NET Web server is an extended IIS server that has the ASP.NET runtime extensibility model integrated into the core server. The ASP.NET server configuration adds the following IIS modules to the default IIS web server installation:

- ASP.NET 4.5
- .NET Extensibility 4.5
- ISAPI Extensions
- ISAPI Filters

Important
To install IIS 8 on either Windows Server® 2012 or Windows® 8, you must sign on as a member of the Administrators group.

By default, Windows Server 2012 and Windows 8 installs .NET 4.5 only. If you want to run .NET 2 applications, install .NET 3.5, which supports web applications written using ASP.NET 2 to 3.5.

1.2. Plan to Add the ASP.NET Application

An application is a grouping of content at the root level of a website or a grouping of content in a separate folder under the website root directory. When you add an application in IIS 8, you designate a directory as the *application root*, or starting point, for the application. Then specify

properties specific to that particular application, such as the application pool that the application runs in.

To configure your ASP.NET application on an IIS website, provide the following information:

- **Alias**: The alias is used as part of the application root URL and should be short and descriptive. For example, the alias *marketing* added to Default Web Site on the local host computer, would produce the following URL: `//localhost/marketing`.
- **Application pool**: An application pool enables an application or a group of applications to run in isolation from one or more applications in another application pool.
- **Physical path**: The local path to the application files on the server.

See Also

- Step 2: Plan ASP.NET Settings
- Plan an ASP.NET Website on IIS
- Build an ASP.NET Website on IIS
- Step 1: Install IIS and ASP.NET Modules

Step 2: Plan ASP.NET Settings

In this phase of building your website, consider the following IIS server and website settings that support ASP.NET:

2.1. Session State Settings

2.2. Pages and Controls Settings

2.3. Application Settings

2.4. .NET Compilation Settings

2.5. .NET Globalization Settings

2.1. Session State Settings

When clients visit a site, they generally navigate from one page to another and frequently change some of the pages they visit. If you want to track pages that users browse and the changes they make during a visit to your website, configure session state.

When session state is enabled for your application, a user receives a unique session ID on their first request to a web page from your ASP.NET application. Session-state data is stored on the server in one of the following ways:

- **In-process**: Session state is stored in the worker process where the ASP.NET application runs.
- **State Server**: Session state is stored outside the worker process where the ASP.NET application runs.
- **SQL Server**: Session state is stored in a SQL Server database.

Note

You can also configure custom out-of-state storage for session-state data. However, it is beyond the scope of this tutorial. Visual Studio 11 project uses custom storage to support SQL Server, SQL Compact, and SQL Azure.

Session-state data can also be stored on the client in a cookie. A cookie is a text file that contains data used for storing information about a user, such as user preferences and user authentication information.

The following sections describe your session-state storage options in more detail.

- Store session state in process
- Store session state by using state server
- Store session state by using SQL server
- Cookie mode for session state

Store session state in process

The in-process session state mode stores session-state data for an ASP.NET application in the worker process where the application runs. This mode is the default for IIS 8.

The advantage of using in-process session state is that it provides the fastest access to session-state data. However, as you store more data in a session, you consume more memory, which can slow server performance.

Before you configure in-process session state, consider the effect of worker process recycling on session state data. If the worker process recycles, all session state data is lost. If your ASP.NET applications must preserve session-state data, and if speed of access to the data is not your primary objective, consider using an out-of-process session state mode for storing this data.

Important

The Windows state service (Aspnet_state.exe) must be running for in-process session state to take effect. By default, this service is installed when Windows Server® 2012 is installed and is configured for manual start. It is recommended that you change the start behavior to Automatic.

By default, the session expires when the user has not requested or refreshed a page in the ASP.NET application for 20 minutes. Because Session objects consume memory on the web server, consider decreasing the time-out value to conserve resources.

Important

Use caution when you adjust the session time-out value, because information stored in a user's Session object is lost when the user is not active on the website for the length of the time-out period.

If you decide to use in-process session-state storage, decide also if you also want to use cookies. For more information about cookies, see Cookie mode for session state.

Store session state by using state server

A *state server* maintains session-state date in memory that is outside the worker process. The advantage of this configuration is that session state is preserved when the application worker process recycles. Using a state server is recommended for medium-sized applications.

A state server depends on the Windows state service (Aspnet_state.exe) and requires a machine key to verify session state across the connection.

When a state server runs on the same web server that contains the applications for which it maintains state, a web garden configuration is supported. For increased protection of session state data, consider using a web farm configuration with a separate server that stores session state and shares it among all servers in the farm. Another approach is to use SQL server to maintain out-of-process session state.

Important

The Windows state service (Aspnet_state.exe) must be running for in-process session state to take effect. By default, this service is installed when Windows Server 2012 is installed and is configured for manual start. Change the start behavior to Automatic.

If you decide to store session state by using a state server, make the following design decisions:

- Define a connection string for the state server.
- Specify the number of seconds to wait before the connection time out.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie. For more information about cookies, see Cookie mode for session state.

Store session state by using SQL server

One type of out-of-process session state uses a SQL server to store session state data. The advantage of this configuration is that session state is preserved despite recycling of the application worker process, or if either the Windows state service or the web server goes down.

Note

This setting does not support SQL Azure.

When a SQL server runs on the same web server that has the applications for which it maintains state, it supports a web garden configuration, which increases web server scalability. When the SQL server runs on another server, it supports a web farm configuration, which greatly increases scalability across a group of servers.

Important

The Windows state service (Aspnet_state.exe) must be running for out-of-process session state to take effect. By default, this service is installed when Windows Server 2012 is installed and is configured for manual start. Change the start behavior to Automatic.

Important

Before you configure a SQL server for session state, run the InstallSqlState.sql script on the server. By default, this script is stored in %systemroot%\Microsoft.NET\Framework\V4.0.30319.

If you decide to store session state in a SQL Server database, make the following design decisions:

- Define a connection string for the database.
- Specify the number of seconds to wait before the connection time-out.
- Specify the number of seconds to wait before trying to reconnect.
- Decide whether to enable a custom database.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie. For more information about cookies, see Cookie mode for session state.

Cookie mode for session state

One way to track session state for clients that connect to a web server is to use cookies. You can configure a web server to use cookies, not to use cookies, or to select cookie behavior that depends on the browser that is used for the connection.

A session cookie associates session information with client information for the session. A session is the duration of a user's connection to a site. The cookie is passed together with all requests between a client and a web server in an HTTP header.

Using cookies to track session state is more efficient than any other method that does not use cookies, because ookies do not require any redirection. In addition, they allow users to bookmark web pages, and they retain state if a user leaves one site to visit another and then returns to the original site. The one drawback of user cookies is that users can disable cookies in their browser.

The *Use Device Profile* cookie mode causes the brower to use cookies if the it cookies; otherwise, no cookies are used. If the device profile indicates support for cookies, they are used regardless of whether the user has disabled cookie support.

Important

When you use the Use Device Profile cookie mode, set expired session IDs be regenerated. Doing so allows a web server to expire and regenerate tokens, which gives a potential attacker less time to capture a cookie and gain access to web server content.

The *Auto-Detect* cookie mode causes the mobile device to use cookies if its profile supports cookies; otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET tries to use cookies when cookie support is enabled in the browser. If cookie support is disabled, session state is stored in the URL.

Important

When you use the Auto-Detect cookie mode, set expired session IDs be regenerated. Doing so enables a web server to expire and regenerate tokens, which gives a potential

attacker less time to capture a cookie and gain access to web server content. Consider changing the time-out value to less than the 20-minute default.

You can configure session state without using cookies. When you use a Uniform Resource Identifier (URI) to handle session state, the session ID is embedded as a query string in the URI request, and then the URI is redirected to the originally requested URL. The changed URI request is used for the duration of the session, so that no cookie is necessary.

Important

When you use a URI, set expired session IDs be regenerated. Doing so enables a web server to expire and regenerate tokens, which gives a potential attacker less time to capture a cookie and gain access to web server content.

Using a URI to track session state can help you avoid the disadvantages of cookies, including browser support problems and the possibility that users disable cookies. However, using a URI has the following disadvantages:

- Cannot use absolute URLs without losing session state, which means that if a user goes to another application and then returns to the previous one, the user's input no longer exists on the page.
- Does not allow users to bookmark web pages, because session state is lost.

If you decide to use cookies to store session state, make the following design decisions:

- Select a cookie mode: auto detect, use cookies, use device profile, or use URI.
- Unless you selected use URI, specify the name of the cookie.
- Unless you selected use URI, specify the number of minutes before the cookie times out.
- Unless you selected use cookies, decide whether to regenerate an expired session ID.

2.2. Pages and Controls Settings

ASP.NET pages include extra elements that ASP.NET recognizes and processes when the page runs. ASP.NET pages can also contain custom, reusable controls. These custom controls are processed on the server. This lets you use server code to set ASP.NET web page properties.

Note

These settings apply only the ASP.NET Web Forms. They do not apply to ASP.NET MVC or ASP.NET Web Pages.

IIS 8 lets you configure the following ASP.NET page and user controls settings:

- Behavior settings: For example, whether the web page maintains its view state and the view state of any server controls it contains when the current page request ends.
- General settings: For example, namespaces that are included for all pages.
- Compilation settings: For example, whether pages are compiled or interpreted.
- Services: For example, whether session state is enabled.

IIS 8 provides default settings for ASP.NET pages and controls, but you can change those settings as needed. For example, you can set the master page file for a site or enable view state.

Web custom controls are compiled components that run on the server and that encapsulate user-interface and other related functionality into reusable packages. In IIS 8, you can specify the tag prefix and namespace mapping for a custom control that can be used in multiple pages in an application.

Add a custom control when you want to specify the tag prefix/namespace mapping for a custom control that is used in multiple pages in an application.

Note
Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

If you decide to configure ASP.NET custom controls, you need the following information for each control you want to configure:

- I must specify the tag prefix of the control.
- I must specify the .NET namespace of the control.
- I must specify the assembly the control is in.

2.3. Application Settings

Configure application settings when you want to store key/value pairs as part of your configuration in the Web.config file. Application settings provide quick and easy access to stored configuration data for your application.

To manage custom controls, you can view a list that contains all custom controls for a particular configuration level. You can sort this list by tag prefix, source or assembly, or scope (local or inherited). You can also group controls by scope to see which custom controls apply at the current configuration level and which custom controls are inherited from a parent level.

Add a custom control when you want to specify the tag prefix/namespace mapping for a custom control that is used in multiple pages in an application.

Note
Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

If you decide to configure application settings, you need the following information for each setting you want to configure:

1. Specify a name for the setting.
2. Specify a value for the setting.

2.4. .NET Compilation Settings

For application code to service requests by users, ASP.NET must first compile the code into one or more assemblies. Assemblies are files that have the file name extension .dll. Configure .NET compilation settings in IIS 8 when you want to control how ASP.NET code is compiled.

IIS lets you configure the following .NET compilation settings:

- Batch settings, such as the maximum file size that you can batch and the maximum number of pages that you can have per batched compilation.
- Behavior settings, such as the number of times resources are dynamically compiled before the application is restarted.
- General settings, such as the default programming language that is used in dynamic compilation files.

2.5. .NET Globalization Settings

Globalization is the process of internationalizing application code, then localizing the application to other languages and cultures. The internationalization process makes it possible to translate, store, retrieve, and present application content for any locale by using the same application code base whenever possible. Locale is the combination of both language and cultural environment. This includes date formats, times, currencies, telephone numbers, and so on. Localization means adapting your application to other locales by translating and formatting content according to culture, preferably without touching the code.

You can change globalization settings for ASP.NET applications at the web server level when you want them to apply to all ASP.NET applications on the server. You can also edit ASP.NET globalization settings for sites, applications, directories, and files.

IIS lets you configure the following globalization settings:

- Culture settings, such as the UI culture or UI language.
- Encoding settings, such as encoding for response headers.

Note

Editing a configuration setting changes the setting at the local level and for any child levels that inherit the setting.

Step 3: Plan Data Source Settings

In this phase of building your website, consider the data storage needs of your ASP.NET application. The following sections describe various data-source settings available in IIS:

3.1. Data source connection strings

3.2. ASP.NET providers

3.3. .NET profiles

3.4. .NET roles

3.5. .NET users

3.1. Data source connection strings

A connection string provides the information that an application or provider must have to communicate with a particular database. A connection string usually supplies the server or

location of the database server, the particular database to use, and the authentication information. A connection string enables you to connect to databases from managed code applications in a centralized manner.

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

To add a connection string to IIS, provide the following information:

- Specify a name for the connection string. This name must be the same name you reference in your application code to retrieve data from the database.
- Specify the server the database is on.
- Specify the name of the database.
- Provide the credentials, unless using Windows integrate security.

3.2. ASP.NET providers

ASP.NET 2.0 includes several services that store state in a database or other data store. A provider is a software module that implements a uniform interface between one of these services and a data source. In IIS 7, you can set the default provider for your application. You can also configure the provider properties. For example, **Users** is a provider-based feature where one provider stores the user data in SQL whereas another provider stores the user data in a text file.

Add a provider in IIS when you have an application that uses a provider-based service to store data in a database or other data store. For example, the session state service in ASP.NET is a provider-based service that manages per-user session state by storing it in process (in memory in the application domain of the host application), in memory in an external process (the "state server process"), or in a Microsoft SQL Server database.

Adding a configuration setting adds the setting at the local level and to any child levels that inherit the setting.

To add a provider for your application, provide the following configuration information:

- Select the IIS feature for the provider to provide: .NET Profiles, .NET Roles, or .NET Users.
- Select the provider type.
- Enter a name for the provider.
- If the feature selected is .NET Users, select any of the following provider behaviors that you want:
    - Enable password reset
    - Enable password retrieval
    - Requires question and answer
    - Requires unique email
    - Store password in secure format
- Provide the name of the connection string to the database.
- Enter the name of the application.

3.3. .NET profiles

The .NET profile feature associates information with an individual user and stores the information in a persistent format. .NET profiles let you manage user information without requiring you to create and maintain your own database.

In .NET profile, you can add properties or groups. A property stores information that is unique to a user, such as the user name. You can then use the information you have stored to present the user with a personalized version of your application. A group organizes related properties together. For example, the different properties of a user's address information can be grouped together in an Address group.

If you decide to add profile properties or groups for your ASP.NET application, provide the following information:

- For each profile property, provide the property name, data type (such as string or Boolean), a default value, a serialization option (string, XML, binary, or provider specific), whether it is read-only, and whether it is available for anonymous users.
- For each profile group, provide the group name.

3.4. .NET roles

Roles give you an easy way to manage access rules for groups of users. You create users and then assign the users to roles (in Windows, you assign users to groups). For example, you can create a set of pages that you want to restrict to certain users and store those pages in a folder by themselves. Then you can use IIS 8 to define rules that grant and deny access to restricted folders. If an unauthorized user tries to view a restricted page, the user either sees an error or is redirected to a page that you specify.

Roles do not work with anonymous users who access a site, application, or file.

Add roles when you have security settings that you want to apply to a group of users who differ from all existing groups.

If you decide to define roles for your application, provide a name for each role.

Important

To configure roles for your application, you must have previously configured a .NET roles provider.

3.5. .NET users

Associating user identities with an application helps you manage authentication, authorization, and other security-related operations for that application.

Add a user when you want to use IIS to define a user name, e-mail address, password, and security question for initiating an automatic account reset if a user loses or forgets their password. If the **.NET Roles** feature is enabled, you can add users to roles as you create them.

If you decide to configure users for your application, provide the following information for each user:

- User Name (must be unique).
- E-mail (must use standard format: name@domain.com).
- Password (must be a strong password).
- Question (enter a custom question or select from the list).
- Answer to the foregoing question.
- Roles the user is assigned to.

Important
To configure users for your application, you must have previously configured a .NET users provider.

Step 4: Plan Application Security

In this phase of building your website, consider the security needs of your ASP.NET application. The following sections describe application security settings available in IIS 8:

4.1. Isolate Web Applications

4.2. .NET Trust Levels

4.3. .NET Authentication

4.4. Machine Key Settings

4.5. TLS/SSL Communication

4.1. Isolate Web Applications

One of the most effective ways to improve security for your web application is to isolate it from other applications on your web server. An application pool has its own worker process, which processes requests and runs application code. The worker process has a security identifier (SID). And each application pool has a unique application-pool identity. By default, when you create a web application, a new application pool is also created with the same name as the application. If you keep web applications in separate application pools, you can isolate them from one another.

Web application isolation entails the following:

- Site isolation: Separate different applications into different sites with different application pools.
- Least privilege: Run your worker process as a low privileged identity (virtual application pool identity) that is unique per site.
- Temp isolation: Set up a separate ASP.NET temp folder per site and only give access to appropriate process identity.
- Content isolation: Make sure to set an ACL (access control list) on each site root to allow only access to the appropriate process identity.

Tip

It is a good idea to host your website and web application content on a drive other than your system drive (C:).

4.2. .NET Trust Levels

An application *trust level* determines the permissions that the ASP.NET code access security (CAS) policy grants. CAS defines two trust categories: full trust and partial trust. An application that has full trust permissions can access all resource types on a server and perform privileged operations. Applications with full trust are affected only by the security settings of the operating system.

Partial-trust web applications are applications that do not have full trust and have a restricted set of code access permissions. As a result, partial-trust applications are limited in their ability to access secured resources and perform other privileged operations. Certain permissions are denied to partial-trust applications, so resources that require those permissions cannot be directly accessed. Other permissions are granted in a restricted way, so resources that require those permissions might be accessible, but in a limited way. For example, restricted file IO permission give the application can access to the file system, but only in directories beneath the application's virtual directory root.

By configuring a web application or web service for partial trust, you can restrict the application ability to access crucial system resources or resources that belong to other web applications. By granting only the permissions that the application requires and no more, you can build least privileged web applications and limit potential damage if the web application is compromised by a code injection attack.

The following list shows the restrictions associated with each trust level:

- Full trust applications have unrestricted access to all resource types and can perform privileged operations.
- High, medium, low, or minimal trust applications are unable to call unmanaged code or serviced components, write to the event log, access Message Queuing queues, or access OLE DB data sources.
- High trust applications have unrestricted access to the file system.
- Medium trust applications have restricted file system access and can only access files in their own application-directory hierarchy.
- Low or minimal trust applications cannot access SQL Server databases.
- Minimal trust applications cannot access any resources.

4.3. .NET Authentication

Authentication helps you confirm the identity of clients who request access to your sites and applications. When authentication is enabled, IIS 8 uses the account credentials supplied by the user to determine what permissions the user has been granted and what resources the user can access.

This section describes the authentication modes that are specific to ASP.NET applications.

1. ASP.NET Forms Authentication
2. ASP.NET Impersonation Authentication

ASP.NET Forms Authentication

Forms authentication uses client-side redirection to forward unauthenticated users to an HTML form where they can enter their credentials, which are usually a user name and password. After the credentials are validated, users are redirected to the page they originally requested. Forms authentication often employs cookies to pass user credentials between the server and the client browser.

The following sections describe what you need to know to plan adding forms authentication to your site:

1. Forms authentication basics
2. Authentication cookies

Forms authentication basics

ASP.NET Forms-based authentication works well for sites or applications on public web servers that receive many requests. This authentication mode lets you manage client registration and authentication at the application level, instead of relying on the authentication mechanisms the operating system provides.

Important
Because Forms authentication sends the user name and password to the web server as plaintext, use Secure Sockets Layer (SSL) encryption for the logon page and for all other pages in your application except the home page. For information about SSL, see 4.5. TLS/SSL Communication.

Forms authentication lets users log on by using identities from an ASP.NET membership database. This authentication method uses redirection to an HTML logon page to confirm the identity of the user. You can configure Forms authentication at the site or application levels.

Forms authentication is convenient for the following reasons:

- It allows either a custom data store, such as a SQL server database, or Active Directory to be used for authentication.
- It integrates easily with a web user interface.
- Clients can use any browser.

If you want to use membership roles for authorization, use Forms authentication or a similar custom authentication method.

Important
If you select Forms authentication, you cannot use any of the challenge-based authentication methods at the same time.

By default, the login URL for Forms authentication is Login.aspx. You can create a unique login page for clients who visit a site or application. For example, you might want to collect specific information from visitors, or offer membership to selected pages on the site or selected applications.

The default time-out value for Forms authentication is 30 minutes. Consider changing the time-out value to a shorter period, to shorten the session lifetime and to reduce the chance of cookie replay attacks.

Authentication cookies

Authentication cookies are used as a token to verify that a client has access to some or all pages of an application. By contrast, personalization cookies contain user-specific settings that determine user experience on a specific site or application.

Important

Because authentication cookies are passed between client and server together with every request, always secure authentication cookies using Secure Sockets Layer (SSL). For information about SSL, see 4.5. TLS/SSL Communication.

Cookies are a more efficient way to track visitors to a site than query strings, because they do not require redirection. However, they are browser-dependent, and some browsers do not support their use. In addition, the use of cookie-based authentication is not always effective because some users disable cookie support in their browsers.

By default, the cookie name for ASP.NET applications is .ASPXAUTH. However, you can instead use a unique cookie name and path for each application. Doing so can prevent users who are authenticated for one application from being authenticated for other applications on a web server.

You can choose one of the following cookie modes for your site or application.

Mode	Description
Use cookies	Cookies are always used regardless of device.
Do not use cookies	Cookies are not used.
Auto Detect	Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET checks to determine whether cookies are enabled. This setting is the default.
Use device profile	Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. ASP.NET does not check to determine whether cookies are enabled on devices that support cookies. This setting is the default for IIS 8.

The cookie protection mode defines the function a Forms authentication cookie performs for a specific application. The following table shows the cookie protection modes that you can define:

Mode	Description
Encryption and validation	Specifies that the application use both data validation and encryption to help protect the cookie. This option uses the configured data validation algorithm (based on the machine key). If triple-DES (3DES) is available and if the key is long enough (48 bytes or more 3DES is used for encryption. This setting is the default (and recommended) value.
None	Specifies that both encryption and validation are disabled for sites that are using cookies only for personalization and have weaker security requirements. We do not recommend that you use cookies in this manner; however, it is the least resource-intensive way to enable personalization by using the .NET Framework.
Encryption	Specifies that the cookie is encrypted by using Triple-DES or DES, but data validation is not performed on the cookie. Cookies used in this manner might be subject to plaintext attacks.
Validation	Specifies that a validation scheme verifies that the contents of an encrypted cookie have not been changed in transit.

Important

For security reasons, consider keeping Encryption and Validation cookies separate from each other. The theft of encryption cookies would be a greater security exposure than the theft of validation cookies.

If an application contains objects that clients request frequently, improve application performance by caching those objects. If the user accesses the cached object before the authentication cookie times out, IIS 8 allows the cached object to remain in the cache, and the timer is reset. However, if the user does not access the cached object during that time, IIS 8 removes the cached object from the cache.

Consider enabling this setting under the following circumstances:

- You have a limited amount of memory available for caching.

- You have many objects to cache, because this setting allows only the most frequently requested objects to remain in the cache.

  Note
  You specify the number of minutes before an authentication cookie times out with **Authentication cookie time-out (in minutes)**.

ASP.NET Impersonation Authentication

Use ASP.NET impersonation when you want to run your ASP.NET application under a security context different from the default security context for ASP.NET applications.

If you enable impersonation for an ASP.NET application, that application can run in one of two different contexts: either as the user authenticated by IIS 8 or as an arbitrary account that you set up. For example, if you use Anonymous authentication and choose to run the ASP.NET application as the authenticated user, the application would run under an account that is set up for anonymous users (typically, IUSR). Likewise, if you chose to run the application under an arbitrary account, it would run under whatever security context was set up for that account.

By default, ASP.NET impersonation is disabled. If you enable impersonation, your ASP.NET application runs under the security context of the user authenticated by IIS 8.

4.4. Machine Key Settings

Machine keys help protect Forms authentication cookie data and page-level view state data. They also verify out-of-process session state identification. ASP.NET uses the following types of machine keys:

- A *validation key* computes a Message Authentication Code (MAC) to confirm the integrity of the data. This key is appended to either the Forms authentication cookie or the view state for a specific page.
- A *decryption key* is used to encrypt and decrypt Forms authentication tickets and view state.

IIS 8 enables you to configure validation and encryption settings and generate machine keys for use with ASP.NET application services, such as view state, forms authentication, membership, roles, and anonymous identification.

Before you generate machine keys for your application, make the following design decisions:

- Decide what validation method to use: AES, MD5, SHA1 (default), TripleDES, HMACSHA256, HMACSHA384, or HMACSHA512.
- Decide what encryption method to use: Auto (default), AES, TripleDES, or DES.
- Decide whether to generate the validation key at runtime automatically.
- Decide whether to generate a unique validation key for each application.
- Decide whether to generate the decryption key at runtime automatically.
- Decide whether to generate a unique decryption key for each application.

4.5. TLS/SSL Communication

Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL) are protocols that provide communication security your website. You can use TLS/SSL to authenticate servers and clients and then use it to encrypt messages between the authenticated parties.

In the authentication process, a TLS/SSL client sends a message to a TLS/SSL server, and the server responds with the information that the server needs to authenticate itself. The client and server perform an additional exchange of session keys, and the authentication dialog ends. When authentication is completed, SSL-secured communication can begin between the server and the client by using the symmetric encryption keys that are established during the authentication process.

To configure TSL/SSL for your website, do the following:

1. Obtain a server certificate from a certification authority (CA). See Server Certificates.
2. Add SSL binding to the site. See SSL Binding.
3. Set IIS to require SSL on the site. See Require SSL for Your Site.
4. Consider using client certificates for your site. See Client Certificates.

Server Certificates

You can obtain a server certificate from a certification authority (CA). Obtaining a server certificate from a certification authority is one step in configuring Secure Sockets Layer (SSL) or Transport Layer Security (TLS). You can obtain server certificates from a third-party CA. A third-party CA might require you to provide proof of identity before a certificate is issued. You can also issue your own server certificates by using an online CA, such as Microsoft Certificate Services.

Digital certificates are electronic files that work like an online password to verify the identity of a user or a computer. They are used to create the SSL encrypted channel that is used for client communications. A certificate is a digital statement that is issued by a certification authority (CA) that vouches for the identity of the certificate holder and enables the parties to communicate in a secure manner by using encryption.

Digital certificates do the following:

- They authenticate that their holders—people, web sites, and even network resources such as routers—are truly who or what they claim to be.
- They protect data that is exchanged online from theft or tampering.

Digital certificates are issued by a trusted third-party CA or a Microsoft Windows public key infrastructure (PKI) using Certificate Services, or they can be self-signed. Each type of certificate has advantages and disadvantages. Each type of digital certificate is tamper-proof and can't be forged.

Certificates can be issued for several uses. These uses include web user authentication, web server authentication, Secure/Multipurpose Internet Mail Extensions (S/MIME), Internet Protocol security (IPsec), Transport Layer Security (TLS), and code signing.

A certificate contains a public key and attaches that public key to the identity of a person, computer, or service that holds the corresponding private key. The public and private keys are

used by the client and the server to encrypt the data before it is transmitted. For Windows-based users, computers, and services, trust in a CA is established when there is a copy of the root certificate in the trusted root certificate store and the certificate contains a valid certification path. For the certificate to be valid, the certificate must not have been revoked and the validity period must not have expired.

SSL Binding

You can assign multiple bindings to a site when you have site content that serves different purposes or for which you must use a different protocol. For example, a commerce site might have an application that requires that users log on to an account to purchase merchandise. The company hosts the site over HTTP, but users must log on to their account on an HTTPS page. In this example, the site would have two bindings: one for the HTTP portion and one for the HTTPS portion.

Out of the box, you cannot add bindings for protocols other than HTTP and HTTPS by using IIS Manager. If you want to add a binding for a different protocol, such as a protocol supported by Windows Communication Foundation (WCF), use one of the other administration tools. However, if you install the IIS File Transfer Protocol (FTP) server, you can add FTP bindings by using IIS Manager. There might also be other modules or third-party functionality available for download that extend the UI.

Require SSL for Your Site

Secure Sockets Layer (SSL) encryption protects confidential or personal information sent between a client and a server. When SSL is enabled, remote clients access your site by using URLs that start with https://.

First configure a server certificate and create an HTTPS binding to enable any SSL settings. Then require Secure Sockets Layer (SSL) encryption in one or more of the following circumstances:

1. When confidential or personal content on your server must be protected by an encrypted channel.
2. When you want users to be able to confirm the identity of your server before they transmit personal information.
3. When you want to use client certificates to authenticate clients that access your server.

Client Certificates

When you want clients to verify their identity before they access content on your web server, configure client certificates. By default, client certificates are ignored.

Before you can configure client certificate on your website, configure a server certificate and create an HTTPS binding to enable any Secure Sockets Layer (SSL) settings.

If you want all clients to verify their identity, specify that client certificates are required. If some clients can access content without first verifying their identity, specify that client certificates are accepted.

Configure an ASP.NET Website on IIS

To install an IIS web server and configure it for ASP.NET web applications, follow the steps listed.

- Step 1: Install IIS and ASP.NET Modules
- Step 2: Configure ASP.NET Settings
- Step 3: Configure Data Source Settings
- Step 4: Configure Application Security

For planning information to review before deployment, see Plan an ASP.NET Website on IIS. For more information, see Build an ASP.NET Website on IIS.

Step 1: Install IIS and ASP.NET Modules

The first step in building an ASP.NET website on IIS 8 is to install IIS along with the ASP.NET modules. Then add your ASP.NET application files to IIS.

The following list shows the tasks required to complete this step:

Installing IIS and ASP.NET Modules

Adding the ASP.NET Application

When you are done, make sure that IIS and the ASP.NET modules are installed, and your ASP.NET application has been added to your website. Then go on to Step 2: Configure ASP.NET Settings.

Installing IIS and ASP.NET Modules

This section describes how to install IIS and the modules that support ASP.NET applications in any of the following ways:

- Install IIS on Windows Server® 2012 using the IIS Manager UI.
- Install IIS on Windows® 8 using the IIS Manager UI.

- Install IIS using the command line.

Important
To install IIS on either Windows Server 2012 or Windows 8, you must sign on as a member of the Administrators group.

If you prefer, you could use the Web Platform Installer (Web PI) to install IIS and third-party applications that run on IIS. Because the Web PI installs the latest versions of available Web Platform offerings, with just a few simple clicks you can download and install any new tools or updates. To learn more about the Web PI, see Learn more and install the Web PI.

If you support web applications written using ASP.NET 2 to 3.5, install .NET 3.5 before installing IIS and ASP.NET modules.

To install .NET 3.5 on Windows Server 2012 or Windows 8

1. On the **Start** screen, right-click the **Command Prompt** tile, and then click **Run as administrator**.
2. At the command prompt, type the following: **dism /online /enable-feature /featurename:netfx3**
3. Wait for the command to complete. It could take several minutes.
4. Close the command prompt window.

To install IIS and ASP.NET modules on Windows Server 2012 using the UI

1. On the **Start** page, click the **Server Manager** tile, and then click **OK**.
2. In **Server Manager**, select **Dashboard**, and click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, select Role-based or feature-based installation, and click **Next**.
5. On the **Select destination server** page, select **Select a server from the server pool**, select your server, and click **Next**.
6. On the **Select server roles** page, select **Web Server (IIS)**, and click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select role services** page, note the preselected role services that are installed by default, expand the **Application Development** node, and then select **ASP.NET 4.5**. (If you installed **.NET 3.5**, select **ASP.NET 3.5** also.)
10. On the **Summary of Features to Install** page, confirm your selections, and then click **Install**.
11. In the **Add features that are required for ASP.NET 4.5?** box, click **Add Features**.

    The following additional features are added:

    - **.NET Extensibility 4.5**
    - **ISAPI Extensions**

- **ISAPI Filters**
- **.NET Extensibility 3.5** (If **ASP.NET 3.5** was selected)

12. Click **Next**.
13. On the **Confirm installation selections** page, click **Install**.
14. On the **Installation progress** page, confirm that your installation of the Web Server (IIS) role and required role services completed successfully, and then click **Close**.
15. To verify that IIS installed successfully, type the following into a web browser:

    http://localhost

    The default IIS Welcome page is displayed.

To install IIS and ASP.NET modules on Windows 8 by using the UI

1. On the **Start** page, click the **Control Panel** tile.
2. In **Control Panel**, click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, click **Internet Information Services** to install the default features.
4. Expand the **Application Development Features** node and click **ASP.NET 4.5** to add the features that support ASP.NET. (If you installed **.NET 3.5**, select **ASP.NET 3.5** also.)

   The following additional features are automatically selected:

   - **.NET Extensibility 4.5**
   - **ISAPI Extensions**
   - **ISAPI Filters**
   - **.NET Extensibility 3.5** (If **ASP.NET 3.5** was selected)

5. Click **OK** to close the **Windows Features** dialog box.
6. To verify that IIS installed successfully, type the following into a web browser:

   http://localhost

   The default IIS Welcome page is displayed.

To install IIS and ASP.NET modules by using the command line

- Type the following command at a Command Prompt or into a script:

  Start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;IIS-HttpErrors;IIS-ApplicationDevelopment;IIS-ASPNET;IIS-NetFxExtensibility;IIS-ISAPIExtensions;IIS-ISAPIFilter;IIS-HealthAndDiagnostics;IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;IIS-Security;IIS-RequestFiltering;IIS-HttpCompressionStatic;IIS-WebServerManagementTools;IIS-ManagementConsole;WAS-WindowsActivationService;WAS-ProcessModel;WAS-NetFxEnvironment;WAS-ConfigurationAPI

Adding the ASP.NET Application

This section describes how to add your ASP.NET application to your website in the following ways:

- Adding an ASP.NET Application with the IIS Manager UI
- Adding an ASP.NET application with the command line

In Plan an ASP.NET Website on IIS, you made the following decisions:

- Alias string that becomes part of the application root URL.
- Application pool to run the application in.
- Physical path to the application files on the server.

To add an ASP.NET application by using the UI

1. Open IIS Manager.
   - For Windows Server 2012, on the **Start** page click the **Server Manager** tile, and then click **OK**. In **Server Manager**, click the **Tools** menu, and then click **Internet Information Services (IIS) Manager**.
   - For Windows 8, on the **Start** page type **Control Panel**, and then click the **Control Panel** icon in the search results. On the **Control Panel** screen, click **System and Security**, click **Administrative Tools**, and then **click Internet Information Services (IIS) Manager**.
2. In the **Connections** pane, expand the **Sites** node.
3. Right-click the site for which you want to create an application, and click **Add Application**.
4. In the **Alias** text box, type a value for the application URL, such as **marketing**. This value is used to access the application in a URL.
5. Click **Select** if you want to select a different application pool than the one listed in the **Application pool** box. In the **Select Application Pool** dialog box, select an application pool from the **Application pool** list and then click **OK**.
6. In the **Physical path** text box, type the physical path of the application's folder, or click the browse button (**...**) to navigate the file system to find the folder.
7. Optionally, click **Connect as** to specify credentials that have permission to access the physical path. If you do not use specific credentials, select the **Application user (pass-through authentication)** option on the **Connect As** dialog box.
8. Optionally, click **Test Settings** to verify the settings that you specified for the application.
9. Click **OK**.

To add an ASP.NET application by using the command line

- To add an application to a site, use the following syntax:

  appcmd add app /site.name:*string* **/path:***string* **/physicalPath:***string*

  The variable **site.name** *string* is the name of the website to which you want to add the application. The variable **path** *string* is the virtual path of the application, such as

/application, and **physicalPath** *string* is the physical path of the application content in the file system.

For example, to add an application named **marketing** to a site named **contoso**, with content at **c:\application**, type the following at the command prompt, and then press ENTER:

appcmd add app /site.name:contoso /path:/marketing /physicalPath:c:\application

See Also

- Step 2: Configure ASP.NET Settings
- Configure an ASP.NET Website on IIS
- Build an ASP.NET Website on IIS
- Step 1: Plan IIS Web Server and ASP.NET Modules Installation

Step 2: Configure ASP.NET Settings

In this phase of building your website, you configure the following IIS server and website settings that support ASP.NET:

2.1. Session State Settings

2.2. Pages and Controls Settings

2.3. Application Settings

2.4. .NET Compilation Settings

2.5. .NET Globalization Settings

2.1. Session State Settings

In Plan an ASP.NET Website on IIS, you decided to use one of the following session-state-storage options:

- **In-process**: Session state is stored in the worker process where the ASP.NET application runs.
- **State Server**: Session state is stored outside the worker process where the ASP.NET application runs.
- **SQL Server**: Session state is stored in a SQL Server database.

You might have also chosen to store session data on the client computer in a cookie. You can even combine the use of cookies with one of the other storage methods. For example, you could store authentication data in a cookie while storing other session-state data in a SQL Server database.

The following sections describe how to configure session state based on the planning decisions you have made:

- Store Session State in Process
- Store Session State by using State Server
- Store Session State by using SQL Server
- Cookie Mode for Session State

Store Session State in Process

This section describes how to configure in-process session state by using either the IIS Manager UI or the command line.

To enable in-process session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **In process**.
4. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
5. In the **Time-out** field, enter a time-out value in the format hh:mm:ss. For example, enter 00:15:00 for 15 minutes.
6. (Optional) Check the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for remote connections.
7. Click **Apply** in the **Actions** pane.

To enable in-process session state by using the command line

- To enable in-process session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /mode:InProc

  The variable **mode:InProc** enables in-process session state. By default, in-process session state is enabled.

  Note
  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.
- To specify the default time that a session object is maintained after the last request associated with the object is made, use the following syntax:

  appcmd set config /section:asp /timeout:*timeSpan*

  The variable *timeSpan* represents the maximum time (hh:mm:ss) that a session object is maintained after the last request associated with the object is made. The default value is

00:20:00.

Store Session State by using State Server

This section describes how to configure a state server for session-state storage by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Define a connection string for the state server.
- Specify the number of seconds to wait before the connection time-out.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie.

To configure a state server to maintain session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **State Server.**
4. Type a connection string in the **Connection string** text box, or click **Create** to create a connection string.
5. Type a time-out value in the **Time-out (in seconds)** text box. The default time-out value is 10 seconds.
6. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
7. (Optional) Select the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for connections to the database.
8. Click **Apply** in the **Actions** pane.

To configure a state server to maintain session state by using the command line

- To configure a state server to maintain session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /mode:StateServer /stateConnectionString:*string* **/stateNetworkTimeout:** *timeSpan* **/useHostingIdentity:True|False**

  The variable **mode:StateServer** sets the session state mode to store session data in a state server. By default, the variable *string* defines the connection string that the state server uses. The default setting is tcpip=loopback:42424. The variable *timeSpan* sets the time, in seconds, that the connection to the state server is maintained. The default is 10 seconds. The variable **useHostingIdentity:True|False** enables or disables the use of Windows authentication and the host process identity (either ASP.NET or a Windows

service identity) for connections to the database. The default value is **True**.

> **Note**
> When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

Store Session State by using SQL Server

This section describes how to configure a SQL Server database for session-state storage using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Define a connection string for the database.
- Specify the number of seconds to wait before the connection time-out.
- Specify the number of seconds to wait before trying to reconnect.
- Decide whether to enable a custom database.
- Decide whether to enable compression.
- Decide whether to store any session state data in a cookie.

To configure a SQL server to maintain session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Session State Mode Settings** area, click **SQL Server**.
4. Type a connection string in the **Connection string** text box, or click **Create** to create a connection string. If the SQL server is on the server you are managing, select **LocalSqlServer** in the **Connection string** text box.
5. Type a time-out value in the **Time-out** text box.
6. (Optional) Check the **Enable custom database** check box to use a custom database for storing session state data.
7. (Optional) Configure cookie settings in the **Cookie Settings** area on the **Session State** page.

   For more information about cookie settings, see Cookie Mode for Session State.
8. (Optional) Check the **Use hosting identity for impersonation** check box to use Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for remote connections to the SQL database.
9. Click **Apply** in the **Actions** pane.

To configure a SQL server to maintain session state by using the command line

1. To configure a state server to maintain session state, use the following syntax:

appcmd set config /commit:WEBROOT /section:sessionState /mode:SqlServer /sqlConnectionString:*string* **/sqlCommandTimeout:** *timeSpan* **/useHostingIdentity:True|False**

The variable **mode:SqlServer** sets the session state mode to store session data in a SQL server database. By default, the variable *string* defines the connection string that the SQL server uses. The default setting is LocalSqlServer. The variable *timeSpan* sets the time, in seconds, that the connection to the SQL server is maintained. The default is 30 seconds. The variable **useHostingIdentity:True|False** enables or disables the use of Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for connections to the database. The default value is **True**.

Note

When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

Cookie Mode for Session State

This section describes how to configure a SQL Server database for session-state storage by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions:

- Select a cookie mode: auto detect, use cookies, use device profile, or use URI.
- Unless you selected use URI, specify the name of the cookie.
- Unless you selected use URI, specify the number of minutes before the cookie times out.
- Unless you selected use cookies, decide whether to regenerate an expired session ID.

To configure Use Cookies mode for session state by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Session State**.
3. On the **Session State** page, in the **Cookie Settings** area, select **Use Cookies** from the **Mode** drop-down list.
4. Type a cookie name in the **Name** text box, or use the default cookie name, ASP.NET_SessionId.
5. Type a time-out value in the **Time-out** text box, or use 20 minutes, the default time-out value.
6. Click **Apply** in the **Actions** pane.

To configure Use Cookies mode for session state by using the command line

- To configure use cookies mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState

/cookieless:UseCookies /cookieName:string **/timeout:**timeSpan

The variable **cookieless:UseCookies** configures IIS 8 to use cookies mode for session state. This is the default value. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. For example, to set a cookie-mode cookie for session state (called MyCookie and that expires after 40 minutes), type the following at the command prompt, and then press Enter:

appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseCookies /cookieName:MyCookie /timeout:40

Note

When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure use device profile cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseDeviceProfile /cookieName:string **/timeout:**timeSpan **/regenerateExpiredSessionId:True|False**

  The variable **cookieless:UseDeviceProfile** configures IIS 8 to use device-profile cookie mode for session state. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

  Note

  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure auto-detect cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:AutoDetect /cookieName:string **/timeout:**timeSpan **/regenerateExpiredSessionId:True|False**

  The variable **cookieless:AutoDetect** configures IIS 8 to use auto-detect cookie mode for session state. The variable *string* is the name of the cookie. The default value is **ASP.NET_SessionId**. The variable *timeSpan* sets the time, in minutes, after which the cookie times out. The default is 20 minutes. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

  Note

When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure URI cookie mode for session state, use the following syntax:

  appcmd set config /commit:WEBROOT /section:sessionState /cookieless:UseUri /regenerateExpiredSessionId:True|False

  The variable **cookieless:UseUri** configures IIS 8 to use URI cookie mode for session state. The variable **regenerateExpiredSessionId:True|False** enables or disables the regeneration of expired session IDs.

  Note
  When you use Appcmd.exe to configure the **<sessionState>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.2. Pages and Controls Settings

This section describes how to configure ASP.NET pages and control settings by using either the IIS Manager UI or the command line.

Note
These settings apply only to ASP.NET Web Forms. They do not apply to ASP.NET MVC or ASP.NET Web Pages.

Edit Pages and Controls

You can edit the settings of existing pages and controls, and the changes are reflected into the Web.config file.

To edit settings for pages and controls by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Pages and Controls**.
3. On the **Pages and Controls** page, edit settings as necessary.
4. When finished, click **Apply** in the **Actions** pane.

To edit settings for pages and controls by using the command line

- To enable or disable page output buffering, use the following syntax:

  appcmd set config /commit:WEBROOT /section:pages /buffer:True|False

  The variable **buffer:True** enables buffering of page output. The default is **True**.

Add a Custom Control

In Plan an ASP.NET Website on IIS, you made the following design decisions for each custom control you want to add:

- Specify the tag prefix of the control.
- Specify the .NET namespace of the control.
- Specify the assembly the control is in.

To add a custom control by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Pages and Controls**.
3. In the **Actions** pane, click **Register Controls**.
4. In the **Actions** pane, click **Add Custom Control**.
5. In the **Add Custom Control** dialog box, type a tag prefix in the **Tag prefix** text box.
6. In the **Namespace** text box, type the namespace to which the custom control belongs. This is the same namespace that was specified in the application code.
7. In the **Assembly** text box, type the name of the source or assembly for the custom control, and then click **OK**.

To add a custom control by using the command line

- To add a custom control, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/pages /+"controls.[assembly='*string***',namespace='***string***',src='***string***', tagName='***string***', tagPrefix='***string***']"**

  The variable **assembly** *string* is the name of the assembly that contains the COM control implementation and requires that the **namespace** attribute is also set. The variable **namespace** *string* is the namespace that is associated with the tag prefix and requires that the **assembly** attribute is also set, if the COM control is not located in the application code directory. The variable **src** *string* is the name of the file that contains the user control and requires that the **tagName** attribute is also set. The variable **tagName** *string* is the name of the control to use in the page and requires that the **src** attribute is also set. The variable **tagPrefix** *string* is the tag prefix that is being mapped to a source file or namespace and assembly. This attribute requires one of the following combinations of other attributes:

  - **namespace** to define a COM control, if the control is in the application code directory.
  - **namespace** and **assembly** to define a COM control.
  - **tagName** and **src** to define a user control.

  For example, to add a custom control that uses the specified source for a user control, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:system.web/pages

/+"controls.[tagPrefix='MyTag', tagName='MyControl',src='controls/MyControl.asx']"

Note

When you use Appcmd.exe to configure the **<controls>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.3. Application Settings

This section describes how to configure application settings by using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions for each application setting you want to configure:

1. Specify a name for the setting.
2. Specify a value for the setting.

To create an application setting by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Application Settings**.
3. In the **Actions** pane, click **Add**.
4. In the **Add Application Setting** dialog box, type a name for the application setting in the **Name** text box.
5. In the **Add Application Setting** dialog box, type a value for the application setting in the **Value** text box, and then click **OK**.

To create an application setting by using the command line

- To add an application setting, use the following syntax:

  appcmd set config /commit:MACHINE /section:appSettings /+"[key='*string*',value='*string*']"

  The variable **key** *string* is the name of the application setting and the variable **value** *string* is the value for the application setting. For example, to add an application setting for the name of your application, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:appSettings /+"[key='Application Name',value='MyApplication']"

2.4. .NET Compilation Settings

Edit .NET compilation settings in IIS 8 when you want to control how ASP.NET code is compiled. This section describes how to edit the compilation settings for you ASP.NET application.

To configure .NET compilation settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Compilation**.
3. On the **.NET Compilation** page, edit settings as necessary.
4. When finished, click **Apply** in the **Actions** pane.

To configure .NET compilation settings by using the command line

- To enable or disable batch compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /batch:True|False

  The variable **batch:True** enables batch compilation, and the variable **batch:False** disables batch compilation. The default value is **True**.
- To change the maximum file size (in KB) of the batch compilation, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /maxBatchGeneratedFileSize:*int*

  The variable *int* is the maximum file size (in KB). The default value is 1000.
- To specify the maximum number of pages per batched compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /maxBatchSize:*int*

  The variable *int* is the maximum number of pages per batched compilation.
- To define the time-out period, in seconds, for batch compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /batchTimeout:*timeSpan*

  The variable *int* is the maximum number of pages per batched compilation.
- To specify compilation of retail or debug binaries for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /debug:True|False

  The variable **debug:True** enables compilation of debug binaries and the variable **debug:False** enables compilation of retail binaries. The default value is **False**.
- To specify the number of times resources can be dynamically recompiled before an application restarts for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /numRecompilesBeforeAppRestart:*int*

  The variable *int* is the number of times resources are dynamically recompiled before the application is restarted. The default is 15.
- To enable or disable URL line pragma for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /urlLinePragmas:True|False

The variable **urlLinePragmas:True** enables URL line pragma and the variable **urlLinePragmas:False** disables URL line pragma. The default value is **False**.

- To enable or disable Visual Basic explicit compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /explicit:True|False

  The variable **explicit:True** enables Visual Basic explicit compilation and the variable **explicit:False** disables explicit compilation. The default value is **False**.

- To enable or disable Visual Basic strict compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /strict:True|False

  The variable **strict:True** enables Visual Basic strict compilation and the variable **strict:False** disables strict compilation. The default value is **False**.

- To add an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /+"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.

- To remove an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /-"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.

- To remove an ASP.NET compilation processing directive for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /-"assemblies.[assembly='*string***']"**

  The variable *string* is the ASP.NET compilation processing directive.

- To change the default programming language that is used in dynamic compilation files for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /defaultLanguage:*string*

  The variable *string* is the default programming language. The default value is vb. For example, to change the default programming language that is used in dynamic compilation files from Visual Basic to C#, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:compilation /defaultLanguage:c#

- To specify the directory to use for temporary file storage during compilation for a site, use the following syntax:

  appcmd set config /commit:WEBROOT /section:compilation /tempDirectory:*string*

  The variable *string* is the directory path.

  Note

  When you use Appcmd.exe to configure the **<compilation>** element at the

global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

2.5. .NET Globalization Settings

This section describes how to edit the globalization settings for you ASP.NET application by using either the IIS Manager UI or the command line.

To edit .NET globalization settings by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Globalization**.
3. On the **.NET Globalization** page, in the property sheet, click to select the global setting you want to edit, and select a value from the drop-down list.
4. In the **Actions** pane, click **Apply**.

To edit .NET globalization settings by using the command line

- To edit the default culture for processing Web requests, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /culture:*string*

  The variable *string* is the default culture for processing web requests. For example, to change the default culture for processing Web requests to US English, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:WEBROOT /section:globalization /culture:en-us

  Note
  When you use Appcmd.exe to configure the **<globalization>** element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To enable or disable client-based culture, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /enableClientBasedCulture:True|False

  The variable **enableClientBasedCulture:True** enables client-based culture and the variable **enableClientBasedCulture:False** disables client-based culture. The default is **False**.

- To edit the default culture for processing locale-dependent resource searches, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /uiCulture:*string*

  The variable *string* is the default culture for processing locale-dependent resource searches.

- To edit the default culture for processing locale-dependent resource searches, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /uiCulture:*string*

  The variable *string* is the default culture for processing locale-dependent resource searches.
- To change the default encoding when parsing .aspx, .asmx, and .asax files, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /fileEncoding:*string*

  The variable *string* is the default encoding to be used when parsing .aspx, .asmx, and .asax files.
- To change the header encoding for responses, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /responseHeaderEncoding:*string*

  The variable *string* is the header encoding used for responses. The default is utf-8.
- To edit the content encoding for responses, use the following syntax:

  appcmd set config /commit:WEBROOT /section:globalization /responseEncoding:*string*

  The variable *string* is the content encoding used for responses. The default is utf-8.

Step 3: Configure Data Source Settings

In this phase of building your ASP.NET website, you configure the data source settings that are available in IIS.

3.1. Data Source Connection Strings

3.2. ASP.NET Providers

3.3. .NET Profiles

3.4. .NET Roles

3.5. .NET Users

3.1. Data Source Connection Strings

This section describes how to create a database connection string in IIS using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions about the need to add a connection string:

1. Specify a name for the connection string.
2. Specify the server the database is on.

3. Specify the name of the database.
4. Provide the credentials, unless using Windows integrate security.

To create a database connection string by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Connection Strings**.
3. On the **Connection Strings** page, click **Add** in the **Actions** pane.
4. In the **Add Connection String** dialog box, type a name for the connection string, such as MyConnection, in the **Name** text box.

   Note
   The name that you enter in IIS Manager is the same name that you reference in your application code to retrieve data by using this connection string.
5. With the **SQL Server** option selected, type the name of the server that hosts the database in the **Server** text box and type the name of the database in the **Database** text box.
6. Under **Credentials**:
   - Select **Use Windows Integrated Security**.

     -Or-
   - Select **Specify credentials** and click **Set**. Type a user name and password for an account that can connect to the server and database in the **User name** and **Password** text boxes. Then type the same password in the **Confirm password** text box, and click **OK**.
7. Click **OK**.

To create a database connection string by using the command line

- To create a database connection string, use the following syntax:

  appcmd set config /commit:MACHINE /section:connectionStrings /+"[connectionString='*string***',name='***string***',providerName='***string***']"**

  The variable **connectionString** *string* is the connection string value, the variable **name** *string* is the key to use to access the connection string value, and the variable **providerName** *string* is the name of the ADO.NET provider to use to access the underlying data store. For example, to create a connection string for an application that uses the Northwind database, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:connectionStrings /+"[connectionString='Data Source=localhost;Integrated Security=SSPI;Initial Catalog=Northwind;',name='Northwind',providerName='System.Data.SqlClient ']"

3.2. ASP.NET Providers

This section describes how to add a .NET provider to IIS using either the IIS Manager UI or the command line.

In Plan an ASP.NET Website on IIS, you made the following design decisions about adding a .NET provider to IIS:

- Select the IIS feature for the provider to provide: .NET Profiles, .NET Roles, or .NET Users.
- Select the provider type.
- Enter a name for the provider.
- If the feature selected is .NET Users, select any of the following provider behaviors that you want:
  - Enable password reset
  - Enable password retrieval
  - Requires question and answer
  - Requires unique email
  - Store password in secure format
- Provide the name of the connection string to the database.
- Enter the name of the application.

To add a provider by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Providers**.
3. On the **Providers** page, under **Feature**, select one of the following features:
   - **.NET Roles**: to configure the provider to provide an interface between the ASP.NET role management service (the "role manager") and role data sources.
   - **.NET Users**: to configure the provider to provide an interface between the ASP.NET membership service and membership data sources.
   - **.NET Profile**: to configure the provider to provide an interface between the ASP.NET profile service and profile data sources.
4. In the **Actions** pane, click **Add**.
5. In the **Add Provider** dialog box, select a provider type from the **Type** drop-down list.
6. In the **Name** text box, type a name for the provider.
7. If the **.NET Users** feature was selected in step #3, in the **Profile properties** section, under **Behavior**, set the value of one or more of the following behaviors to **True** to enable the behavior:
   - **EnablePasswordReset**: indicates whether passwords can be reset by using the provider *ResetPassword* method. The default setting is **False**.
   - **EnablePasswordRetrieval**: indicates whether passwords can be retrieved by using the provider *GetPassword* method. The default setting is **False**.

     Important

Some providers, such as the Active Directory provider, do not support the retrieval of passwords. For these providers, the value of the **enablePasswordRetrieval** attribute is always **False** and cannot be changed in configuration.

- **RequiresQuestionAndAnswer**: indicates whether a password answer is supplied when the program calls the provider *GetPassword* and *ResetPassword* methods. The default setting is **False**.
- **RequiresUniqueEmail**: indicates whether each registered user must have a unique e-mail address. The default setting is **False**.
- **StorePasswordInSecureFormat**: indicates whether passwords are hashed. The default setting is **False**.

8. Under **Data**, type the name of the connection string that is used to connect to the database in the **ConnectionStringName** text box.

   Note
   This is the same connection string that was configured under the **Connection Strings** feature. For more information about how to configure connection strings, see **Configuring database connection strings**. If the *WindowsTokenRoleProvider* was selected in step #5 of this procedure, a connection string name is optional.

9. Optionally, in the **Profile properties** section, under **General**, type the virtual path of the application in the **ApplicationName** text box. If you do not specify a value in the **ApplicationName** text box, the membership API defaults to HttpContext.Current.Request.ApplicationPath. This path can be determined only at runtime.
10. Optionally, in the **Profile properties** section, under **General**, type a description of the provider in the **Description** text box.
11. Click **OK**.

To add a provider by using the command line

1. To add a .NET role provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:roleManager /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. For example, to configure an ASP.NET application to use the **SqlRoleProvider** class to store and retrieve role information, type the following at the command prompt, and then press Enter:

   appcmd set config /commit:MACHINE /section:roleManager /+"providers.[name='SqlProvider',type='System.Web.Security.SqlRoleProvider',con

nectionStringName='SqlServices',applicationName='SampleApplication']"

2. To add a .NET user provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:membership /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***',enablePasswordReset='True|False',enablePasswordRetrieval='True|False',requiresQuestionAndAnswer='True|False',requiresUniqueEmail='True|False',storePasswordInSecureFormat='True|False']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. The variable **enablePasswordReset True** enables password reset by using the provider *ResetPassword* method. The default setting is **False**. The variable **enablePasswordRetrieval True** enables password retrieval by using the provider *GetPassword* method. The default setting is **False**. The variable **requiresQuestionAndAnswer True** requires that a password answer must be supplied when the program calls the provider *GetPassword* and *ResetPassword* methods. The default setting is **False**. The variable **requiresUniqueEmail True** requires that each registered user has a unique e-mail address. The default setting is **False.** The variable **storePasswordInSecureFormat True** requires that all passwords must be hashed. The default setting is **False**. For example, to configure an ASP.NET application to use the **SqlMembershipProvider** class to store and retrieve user information, type the following at the command prompt, and then press Enter:

   appcmd set config /commit:MACHINE /section:membership /+"providers.[name='SqlProvider',type='System.Web.Security.SqlMembershipProvider',connectionStringName='SqlServices',applicationName='/',enablePasswordRetrieval='False',enablePasswordReset='True',requiresQuestionAndAnswer='True']"

3. To add a .NET profile provider, use the following syntax:

   appcmd set config /commit:MACHINE /section:profile /+"providers.[name='*string***',type='***string***',connectionStringName='***string***',applicationName='***string***']"**

   The variable **name** *string* is the name of the provider. The variable **type** *string* is the provider type. The variable **connectionStringName** *string* is the name of the connection string that is used to connect to the database. The variable **applicationName** *string* is optional and is the virtual path of the application. For example, to configure an ASP.NET application to use the **SqlProfileProvider** class to store and retrieve profile information, type the following at the command prompt, and then press Enter:

   appcmd set config /commit:MACHINE /section:profile /+"providers.[name='SqlProvider',type='System.Web.Profile.SqlProfileProvider',connectionStringName='SqlServices',applicationName='SampleApplication']"

3.3. .NET Profiles

This section describes how to add a profile property and a profile group.

In Plan an ASP.NET Website on IIS, you made the following design decisions about profile properties and groups to IIS:

- For each profile property, provide the property name, data type (such as string or Boolean), a default value, a serialization option (string, XML, binary, or provider specific), whether it is read-only, and whether it is available for anonymous users.
- For each profile group, provide the group name.

To add a .NET profile property by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Profile**.
3. In the **Actions** pane, click **Add Property** to add a .NET profile property or, to add a .NET profile property to a group, select the group to which you want to add the .NET profile property, and then click **Add Property To Group**.
4. In the **Add .NET Profile Property** dialog box, type a name for the .NET profile property in the **Name** text box.
5. Under **Data Type**, select one of the following data types:
   - **System.Boolean**: to configure the .NET profile property to have a value of either **True** or **False**.
   - **System.Char**: to configure the .NET profile property to contain Unicode Characters.
   - **System.DateTime**: To configure the .NET profile property to contain dates and times with values that range from 12:00:00 midnight, January 1, 0001 Anno Domini (A.D.) or Common Era (C.E.) through 11:59:59 P.M., December 31, 9999 A.D. (C.E.)
   - **System.Decimal**: to configure .NET profile property to contain decimal numbers that range from positive 79,228,162,514,264,337,593,543,950,335 to negative 79,228,162,514,264,337,593,543,950,335. The **Decimal** value type is appropriate for financial calculations that require many significant integral and fractional digits and no round-off errors.
   - **System.Double**: To configure the .NET profile property to contain a double-precision 64-bit number with values that range from negative 1.79769313486232e308 to positive 1.79769313486232e308, as well as positive or negative zero, **PositiveInfinity**, **NegativeInfinity**, and Not-a-Number (NaN). For more information about the **Double** value type, see Double Structure.
   - **System.Int32**: to configure the .NET profile property to contain a signed integer with a value that ranges from negative 2,147,483,648 through positive 2,147,483,647.
   - **System.Int64**: to configure the .NET profile property to contain an integer with a value that ranges from negative 9,223,372,036,854,775,808 through positive 9,223,372,036,854,775,807.
   - **System.Single**: To configure the .NET profile property to contain a single-precision 32-bit number with values that range from negative 3.402823e38 to positive

3.402823e38, as well as positive or negative zero, **PositiveInfinity**, **NegativeInfinity**, and Not-a-Number (NaN). For more information about the **Single** value type, see Single Structure.

- **System.String**: To configure the .NET profile property to be a sequential collection of **Char** structures that represents a string of text. This setting is the default.

6. In the **Default value** text box, type a value that the property will be initialized with.
7. Under **Serialization option**, select one of the following serialization formatters:
   - **String**: Select this option when the settings property is serialized as plaintext. This setting is the default.
   - **XML**: Select this option to serialize only public properties and fields. XML serialization does not preserve type fidelity. This option is useful when you want to provide or consume data without restricting the application that uses the data. Because XML is an open standard, it is an attractive choice for sharing data on the Web.
   - **Binary**: Select this option to preserve type fidelity. Binary serialization is useful for preserving the state of an object between different invocations of an application. For example, you can share an object between different applications by serializing it to the Clipboard. You can serialize an object to a stream, to a disk, to memory, over the network, and so on.
   - **Provider Specific**: Select this option when the settings provider has implicit knowledge of the property or its type and can pick an appropriate serialization mechanism. This option is often used for custom serialization.
8. Check the **Read only** box to configure the .NET profile property so that it cannot be modified.
9. Check the **Available for anonymous users** box to make the .NET profile property available for unauthenticated users.
10. Click **OK**.

To add a .NET profile property by using the command line

- To add a .NET Profile property, use the following syntax:

  appcmd set config /commit:MACHINE /section:profile /+"properties.[name='*string***',type='***string***',defaultValue='***string***',serializeAs='String|Xml|Binary|ProviderSpecific',readOnly='True|False',allowAnonymous='True|False']"**

  The variable **name** *string* is the name of the profile property. The variable **type** *string* is the data type. The variable **defaultValue** *string* is the value with which the profile property is initialized. The variable **serializeAs 'String|Xml|Binary|ProviderSpecific** sets the serialization formatter. The variable **readOnly True** configures the .NET profile property so that it cannot be modified. The default value is **False**. The variable **allowAnonymous True** makes the .NET profile property available for unauthenticated users. The default value is **False**. For example, to specify a profile property that holds a recent search list collection, type the following at the command prompt, and then press Enter:

  appcmd set config /commit:MACHINE /section:profile /+"properties.

[name='RecentSearchList',type='System.Collections.Specialized.StringCollection', serializeAs='Xml',allowAnonymous='true']"

To add a .NET profile group by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Profile**.
3. In the **Actions** pane, click **Add Group**.
4. In the **Add Group** dialog box, type a name for the .NET profile group in the **Name** text box, and then click **OK**.

3.4. .NET Roles

This section describes how to add a .NET role by using the IIS Manager UI.

To add a .NET role by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Roles**.
3. In the **Actions** pane, click **Add**.
4. In the **Add .NET Role** dialog box, type the name of the role in the **Name** text box, and then click **OK**.

3.5. .NET Users

This section describes how to configure .NET users by using the IIS Manager UI.

In Plan an ASP.NET Website on IIS, you made the following design decisions for each user you want to add to IIS:

- User Name
- E-mail
- Password
- Question
- Answer
- Role or roles

To add a .NET user by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Users**.
3. In the **Actions** pane, click **Add**.
4. In the **.NET User Account Details** dialog box, enter the following information:
   - User Name (must be unique).

- E-mail (must use standard format: name@domain.com).
- Password (must be a strong password).
- Confirm Password (must match the password).
- Question (enter a custom question or select from the list).
- Answer to the question.

5. Click **Next** to select roles for this user. If you have not enabled **.NET Roles**, click **Finish**.
6. In the **.NET User Roles** dialog box, optionally select available roles from the **Roles** box, and then click **Finish**.

Step 4: Configure Application Security

In this phase of building your ASP.NET website, you configure the security settings that are available in IIS. The following sections discuss common security settings for ASP.NET applications:

4.1. Isolate Web Applications

4.1. .NET Trust Levels

4.2. .NET Authentication

4.3. Machine Key Settings

4.4. TLS/SSL Communication

4.1. Isolate Web Applications

Implement the following recommendations to isolate websites and web applications on your server.

- Use one application pool per website or web application.
- Limit access to site folders and files to the application pool identity.
- Set up a separate ASP.NET temp folder per site and only give access to the application pool identity.
- Make sure to set an ACL (access control list) on each site root to allow only access to the application pool identity.

If you have more than one application per application pool, consider creating enough application pools and moving some of the applications to the new pools.

To create an application pool

1. Open IIS Manager.
2. In the **Connections** pane, click **Application Pools**.
3. In the **Actions** pane, click **Add Application Pool**.

4. In the **Name** box, type a unique name for the application pool.
5. Select the **.NET Framework version** and **Managed pipeline mode**.
6. Click **OK**.

To move an application to another application pool

1. Open IIS Manager.
2. In the **Connections** page, select the website or web application you want to move.
3. In the **Actions** pane, click **Basic Settings**.
4. On the **Edit Site** dialog, click **Select** to open the **Select Application Pool** dialog, and then select the application pool from the **Application pool** menu.
5. Click **OK** to close the **Select Application Pool** dialog, and click **OK** to close the **Edit Site** menu.

To add an application pool identity to a folder or file ACL

1. Open Windows Explorer and navigate to the folder or file.
2. Right click the folder or file, and then click **Properties**.
3. Select the **Security** tab, and then click **Edit**.
4. Click **Add**, click **Locations**, and select your server as the location to search.
5. In the **Enter the object names to select** box, type **IIS APPPOOL\applicationPoolName**, where *applicationPoolName* is the application pool identity.
6. Click **OK**, click **OK**, and click **OK** again to close the dialogs.

4.1. .NET Trust Levels

This section describes how to set application trust level by using either IIS Manager UI or the command line.

To set a trust level by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **.NET Trust Levels**.
3. On the **.NET Trust Levels** page, select a trust level from the **Trust level** drop-down list, and then click **Apply** in the **Actions** pane.

To set a trust level by using the command line

1. To set a trust level, use the following syntax:

   appcmd set config /commit:WEBROOT /section:trust /level: Full | High | Medium | Low | Minimal

   The level attribute uses one of five values that correspond to preconfigured CAS policy

files. For example, to set a trust level of **Full**, type the following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:trust /level:Full

Note
When you use Appcmd.exe to configure the trust element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

4.2. .NET Authentication

In Plan an ASP.NET Website on IIS, you made design decisions about what authentication mode was right for your application. The following sections decide how to configure authentication for your ASP.NET application:

1. ASP.NET Forms Authentication
2. ASP.NET Impersonation Authentication

ASP.NET Forms Authentication

This section describes how to configure ASP.NET forms authentication by using either the IIS Manager UI or the command line.

To configure forms authentication by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Authentication**.
3. On the **Authentication** page, select **Forms Authentication**.
4. In the **Actions** pane, click **Enable** to use Forms authentication with the default settings.
5. In the **Actions** pane, click **Edit**.
6. In the **Edit Forms Authentication Settings** dialog box, in the **Login URL** text box, type the name of the page where clients log in.
7. In the **Authentication cookie time-out (in minutes)** text box, type the number of minutes you want to use for the time-out value.
8. From the **Mode** list, select the cookie mode you want to use.
9. In the **Name** text box, type the name of the cookie.
10. From the **Protection mode** list, select the protection mode you want to use.
11. Select the **Requires SSL** check box.
12. Select the **Extend cookie expiration on every request** check box, and then click **OK**.

To configure forms authentication by using the command line

- To enable forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /mode:

None | Windows | Windows Live ID | Forms

By default, IIS 8 sets the **mode** attribute to Windows, which disables Forms authentication. If you set the attribute to Forms, you enable Forms authentication. For example, to enable Forms authentication, type the following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:system.web/authentication /mode:Forms

Note

When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To specify the login URL for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.loginURL:*string*

  The variable **forms.loginURL** *string* is the name of the page where clients login. The default value is Login.aspx. For example, to specify the login URL for Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.loginURL:login.aspx

  Note

  When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To specify the authentication time-out for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.timeout:*TimeSpan*

  The variable **forms.timeout** *TimeSpan* is the time in minutes when the cookie used for authentication expires. The default value is 30 minutes. For example, to specify the authentication time-out for Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.timeout:30

  Note

  When you use Appcmd.exe to configure the authentication element at the global level in IIS 8, you must specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

- To configure the cookie name for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.name:*string*

  The variable **forms.name** *name* is the name of the cookie used for Forms authentication. The default value is .ASPXAUTH. For example, to configure the cookie name for Forms authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.name:.ASPXUTH
- To configure the cookie mode for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.cookieless: UseUri | UseCookies | AutoDetect | UseDeviceProfile

  The default value for **forms.cookieless** is **UseDeviceProfile**. For example, to configure the cookie mode for Forms authentication to use the setting Use Device Profile, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.cookieless:UseDeviceProfile
- To configure the cookie protection mode for Forms authentication, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.protection: All | None | Encryption | Validation

  The default value for **forms.protection** is **All**. For example, to configure the cookie protection mode for Forms authentication to use the setting Encryption and Validation, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.protection:All
- To require SSL for an authentication cookie, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.requireSSL: True | False

  The default value for **forms.requireSSL** is False. If you set this attribute to **True**, you require SSL. For example, to require SSL for an authentication cookie, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.requireSSL:True
- To cache frequently requested content, use the following syntax:

  appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.slidingExpiration: True | False

The default value for **forms.slidingExpiration** is **True**. For example, to cache frequently requested content, type the following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:system.web/authentication /forms.slidingExpiration:True

ASP.NET Impersonation Authentication

To configure Impersonation Authentication by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Authentication**.
3. On the **Authentication** page, select **ASP.NET Impersonation**.
4. In the **Actions** pane, click **Enable** to use ASP.NET Impersonation authentication with the default settings.
5. Optionally, in the **Actions** pane, click **Edit** to set the security principal.
6. In the **Edit ASP.NET Impersonation Settings** dialog box, select either **Specific user** or **Authenticated user**. Whichever you decide, IIS uses this identity for the security context of the ASP.NET application. By default, IIS 8 is set to impersonate the authenticated user.
7. Click **OK** to finish or proceed to the next optional steps to change the identity to impersonate.
8. Optionally, click **Set** to change the **Specific user** identity.
9. In the **Set Credentials** dialog box, enter the name of an existing user account in **User name,** the password associated with that user account in **Password**, and then the exact same value in **Confirm password** for a new account IIS should use for anonymous access.
10. Click **OK** to close the **Set Credentials** dialog box.
11. Click **OK** to close the **Edit ASP.NET Impersonation Settings** dialog box.

To configure Impersonation Authentication by using the command line

- To enable or disable ASP.NET Impersonation, use the following syntax:

  appcmd set config /commit:WEBROOT /section:identity /impersonate:true | **false**

  By default, IIS sets the **impersonate** attribute to false, which disables ASP.NET Impersonation authentication. If you set the attribute to true, you enable ASP.NET Impersonation authentication. For example, to enable ASP.NET Impersonation authentication, type the following at the command prompt, and then press ENTER:

  appcmd set config /commit:WEBROOT /section:identity /impersonate:true
- Optionally, you can set the account for IIS to impersonate, using the following syntax:

appcmd set config /commit:WEBROOT /section:identity /userName:*string* **/password:***string*

The variable **userName** *string* is the account IIS uses to impersonate and the variable **password** *string* is the password. For example, to use an account named **Moe** for IIS to impersonate, type the following at the command prompt, and then press ENTER:

appcmd set config /commit:WEBROOT /section:identity /userName:Moe /password:pass@word1

Note

When you use Appcmd.exe to configure the identity element at the global level in IIS 8, specify **/commit:WEBROOT** in the command so that configuration changes are made to the root Web.config file instead of ApplicationHost.config.

4.3. Machine Key Settings

This section describes how to generate machine keys for your ASP.NET application by using the IIS Manager UI.

To generate machine keys by using the UI

1. Open IIS Manager and navigate to the level you want to manage.
2. In **Features View**, double-click **Machine Key**.
3. On the **Machine Key** page, select a validation method from the **Validation method** list. The default validation method is **SHA1**.
4. Choose an encryption method from the **Encryption method** list. The default encryption method is **Auto**.
5. Optionally, configure settings for validation and decryption keys.
6. In the **Actions** pane, click **Generate Keys**, and then click **Apply**.

4.4. TLS/SSL Communication

This section describes how to configure TLS/SSL security for your application.

After obtaining a server certificate from a certification authority (CA), work through the procedures in the following sections:

1. SSL Binding
2. Require SSL for Your Site
3. Client Certificates

SSL Binding

This section describes how to add an SSL binding to your site by using either the IIS Manager UI or the command line.

To add an SSL binding to a site

1. Open IIS Manager.
2. In the **Connections** pane, expand the **Sites** node in the tree, and then click to select the site for which you want to add a binding.
3. In the **Actions** pane, click **Bindings**.
4. In the **Site Bindings** dialog box, click **Add**.
5. In the **Add Site Binding** dialog box, in the **Type** list, select **https**.
6. From the **IP address** list, select **All Unassigned** (unless there is a specific IP address you want to use).
7. In the **Port** box, type the number of the port (the default is 443).
8. In the **Host name** box, type the name of the host computer.
9. If you want multiple secure websites to be served using the same IP address, select the **Require Server Name Indication** check box.
10. From the **SSL certificate** list, select the certificate for your website. If your certificate doesn't appear in the list, click **Select** and search for the certificate using the **Select Certificate** dialog box.
11. Click **OK**.

1. To add a binding to a site, use the following syntax:

   appcmd set site /site.name:*string*
   /+bindings.[protocol='*string***',bindingInformation='***string***']**

   The variable **site.name** *string* is name of the site to which you want to add a binding. The variable **protocol** *string* is the protocol that you want to use, and the variable **bindingInformation** *string* is the combination of IP address, port, and host header.

   For example, to configure a site named **contoso** to have an HTTPS binding for all IP addresses, on port 443, without a host header, type the following at the command prompt, and then press ENTER:

   appcmd set site /site.name:contoso
   /+bindings.[protocol='https',bindingInformation='*:443:']

Require SSL for Your Site

This section describes how to require SSL for your website by using the IIS Manager UI or the command line.

To require SSL using the UI

1. Open IIS Manager and navigate to the level you want to manage.

   Make sure that you are at the site, application, or directory level; **SSL Settings** are not

available at the server level.

Note
If you want to configure SSL at the file level, navigate to the file in **Content View** and then click **Switch to Features View** in the **Actions** pane.

2. In **Features View**, double-click **SSL Settings**.
3. On the **SSL Settings** page, select **Require SSL**.
4. In the **Actions** pane, click **Apply**.

To require SSL using the command line

- To require SSL, use the following syntax:

  appcmd set config "*site* | *URL***" /section:access /sslFlags:Ssl /commit:APPHOST**

  The variable *site* | *URL* is the site, application, virtual directory, or file where you want IIS 8 to require SSL. For example, to require SSL for the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "Default Web Site" /section:access /sslFlags:Ssl /commit:APPHOST
- To require SSL for the file iisstart.htm on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "http://localhost/iisstart.htm" /section:access /sslFlags:Ssl /commit:APPHOST

Client Certificates

This section describes how to

To specify client certificates by using the UI

1. Open IIS Manager and navigate to the level you want to manage.

   Make sure that you are at the site, application, or directory level; **SSL Settings** are not available at the server level.

   Note
   If you want to configure SSL at the file level, navigate to the file in **Content View** and then click **Switch to Features View** in the **Actions** pane.
2. In **Features View**, double-click **SSL Settings**.
3. On the **SSL Settings** page, optionally select **Require SSL**. You do not need SSL to **Ignore** or **Accept** client certificates.
4. On the **SSL Settings** page, in the **Client certificates** area, use one of the following procedures:
   - Select **Ignore** if you do not want to accept a client certificate even if a client presents one.

- Select **Accept** to accept client certificates.
- Select **Require** to require client certificates. To use **Require Client Certificates**, you must enable **Require SSL**.

5. In the **Actions** pane, click **Apply**.

To specify client certificates by using the command line

- To specify whether to use client certificates, use the following syntax:

  appcmd set config "*site* | *URL***"/section:access /sslFlags: Ssl | SslNegotiateCert | SslRequireCert /commit:APPHOST**

  The variable *site* | *URL* is the site, application, virtual directory, or file where you want IIS to enable client certificates. For example, to accept client certificates for the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "Default Web Site"/section:access /sslFlags:SslNegotiateCert /commit:APPHOST

- To accept client certificates for the file iisstart.htm on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "http://localhost/iisstart.htm"/section:access /sslFlags:SslNegotiateCert /commit:APPHOST

  You can specify one or more of the values for the sslFlags attribute. If you want more than one value, separate each value with a comma (,). For example, to specify a requirement for both SSL and client certificates on the Default Web Site, type the following at the command prompt, and then press ENTER:

  appcmd set config "Default Web Site"/section:access /sslFlags:Ssl,SslRequireCert /commit:APPHOST

Build an FTP Site on IIS

File Transfer Protocol (FTP) simple is a protocol for transferring files between computer systems. IIS 8 includes an FTP server that is easy to configure.

The document shows how to install and configure the FTP server on an existing IIS 8 web server. The first two steps are required. All other steps are optional but recommended.

In this document

- Prerequisites
- Step 1: Install FTP on an Existing IIS Web Server
- Step 2: Add an FTP Site
- Step 3: Configure FTP Site Defaults
- Step 4: Configure Firewall Support

- Step 5: Configure User Isolation
- Step 6: Configure Directory Browsing Options
- Step 7: Configure Logon Attempt Restrictions
- Step 8: Configure Request Filtering
- Step 9: Configure FTP Logging
- Step 10: Configure FTP Messages

Prerequisites

To get the most from this tutorial, you must have access to a computer that runs one of the following operating systems:

- Windows Server® 2012
- Windows® 8

Step 1: Install FTP on an Existing IIS Web Server

This step shows you how to install the FTP service on an existing IIS web server that runs on either Windows Server 2012 or Windows 8.

To install FTP on Windows Server 2012

1. On the **Start** screen, click the **Server Manager** tile, then click **Yes**.
2. In the **Server Manager Dashboard**, click **Add roles and features**.
3. If the **Before you begin** page of **Add Roles and Features Wizard** is displayed, click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation**, and click **Next**.
5. On the **Select destination server** page, select **Select a server from the server pool**, select your server from the **Server Pool** list, and then click **Next**.
6. On the **Select server roles** page, expand the **Web Server (IIS)** node, and then expand the **FTP Server** node.
7. Select the **FTP Server** check box and the **FTP Service** check box, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.

To install FTP on Windows 8

1. On the **Start** screen, type **Control Panel**, and then click the **Control Panel** icon in the search results.
2. Click **Programs**, and then click **Turn Windows features on or off**.
3. In the **Windows Features** dialog box, expand the **Internet Information Services** node.
4. Expand the **FTP Server** node.

5. Select the **FTP Server** check box and the **FTP Service** check box, and then click **OK**.

Step 2: Add an FTP Site

Once the FTP service is installed on your IIS web server, you can add one or more FTP sites. Add an FTP site when you want to enable clients to transfer files to and from a site by using the File Transfer Protocol (FTP).

Note

Because FTP settings are contained in the sites section, changing any FTP setting also forces website application recycling. If you want to avoid this side effect, add a site that is configured exclusively for FTP, instead of for both HTTP and FTP.

To add an FTP site

1. Open IIS Manager.
2. In the **Connections** pane, expand the server node and click the **Sites** node.
3. In the **Actions** pane, click **Add FTP Site** to open the **Add FTP Site** wizard.
4. On the **Site Information** page, in the **FTP site name** box, type a unique friendly name for the FTP site.
5. In the **Physical path** box, type the physical path or click the browse button (**...**) to locate the physical path of the content directory.
6. Click **Next** to open the **Binding and SSL Settings** page.
7. Under **Binding**, in the **IP Address** list, select or type an IP address if you do not want the IP address to remain **All Unassigned**.
8. In the **Port** box, type the port number.
9. Optionally, in the **Virtual Host** box, type a host name if you want to host multiple FTP sites on a single IP address. For example, type **www.contoso.com**.
10. Clear the **Start FTP site automatically** box if you want to start the site manually.
11. Under **SSL**, from the **SSL Certificate** list, select a certificate. Optionally, click **View** to open the **Certificates** dialog box and verify information about the selected certificate.
12. Select one of the following options:
    - **Allow SSL**: Allows the FTP server to support both non-SSL and SSL connections with a client.
    - **Require SSL**: Requires SSL encryption for communication between the FTP server and a client.
13. Click **Next** to open the **Authentication and Authorization Information** page.
14. Under **Authentication**, select the authentication method or methods that you want to use:
    - **Anonymous**: Allows any user to access content providing only the user name **anonymous** or **ftp**. (Most, but not all, FTP clients enter user name for you automatically.)

- **Basic**: Requires users to provide a valid user name and password to access content. Because Basic authentication transmits unencrypted passwords across the network, use this authentication method only when you know that the connection between the client and FTP server is secure, such as by using Secure Sockets Layer (SSL).

15. Under **Authorization**, from the **Allow access to** list, select one of the following options:
    - **All Users**: All users, whether they are anonymous or identified, can access the content.
    - **Anonymous Users**: Anonymous users can access the content.
    - **Specified Roles or User Groups**: Only members of certain roles or user groups can access the content. Type the role or user group in the corresponding box.
    - **Specified Users**: Only specified users can access the content. Type the user name in the corresponding box.
16. If you selected an option from the **Allow access to** list, select one or both of the following permissions:
    - **Read**: Permits authorized users to read content from the directory.
    - **Write**: Permits authorized users to write to the directory.
17. Click **Finish**.

Step 3: Configure FTP Site Defaults

Change an FTP site default value when you want new FTP sites to use a different default value.

Note
When you change a default value, existing sites are not overridden with the new value. Change the value for any existing sites manually.

To configure FTP site defaults

1. Open IIS Manager.
2. In the **Connections** pane, click the server node and the **Sites** node.
3. In the **Actions** pane, click **FTP Site Defaults**.
4. Edit the settings, and then click **OK**.

The following table lists the settings available for configuration in the **FTP Site Defaults** dialog box.

Setting Category	Setting Name	Description
General	Allow UTF-8	Specifies whether to use UTF8 encoding. Default is true.
General	Start Automatically	If true, the FTP site is started upon creation or when the FTP service is started. Default is

Setting Category	Setting Name	Description
		true.
Connections	Control Channel Time-out	Specifies the time-out (in seconds) when a connection times out due to inactivity.
Connections	Data Channel Time-out	Specifies the time-out (in seconds) when the data channel times out due to inactivity.
Connections	Disable Socket Pooling	Specifies whether socket pooling is used for sites distinguished by IP address rather than port number or host name.
Connections	Max Connections	Specifies the maximum simultaneous connections to a server.
Connections	Reset On Max Connections	Specifies whether to disconnect FTP session when sending max connections response.
Connections	Server Listen Backlog	Specifies the number of outstanding sockets that can be queued.
Connections	Unauthenticated Time-out	Specifies the timeout (in seconds) between when a new connection is made and authentication succeeds.
Credential Caching	Enabled	Specifies whether credential caching is enabled for the FTP service.
Credential Caching	Flush Interval	Specifies the cache lifetime, in seconds, for the credentials that are stored in the cache.
File Handling	Allow Reading Files While Uploading	Specifies whether files can be read while being transferred to the server.

Setting Category	Setting Name	Description
File Handling	Allow Replace on Rename	Specifies whether files can overwrite other files when renamed.
File Handling	Keep Partial Uploads	Specifies whether to keep files that have been partially uploaded.

Step 4: Configure Firewall Support

Use the **FTP Firewall Support** feature to configure the following settings that enable the FTP server to accept passive data connections from a firewall:

- **Data Channel Port Range**: Specify a range of ports for passive data connections. You must also open that same range of ports on your firewall. This option can only be configured at the server-level, and the port range of “0-0” means to use the ephemeral port range of the server. An ephemeral port is a short-lived transport protocol port that TCP/IP allocates from a predefined range.
- **External IP Address of Firewall**: Specify the external IP address of your firewall so that clients know which IP address to use when they communicate with the FTP server through the firewall.

To configure firewall support

1. Open IIS Manager.
2. In the **Connections** pane, select the server node.
3. In **Features View**, double-click **FTP Firewall Support**.
4. In the **Data Channel Port Range** box, type a range of port numbers (separated by a hyphen). For example, type **5000-6000**. Or type **0-0** to use the default port range specified in Windows TCP/IP settings.

   Note
   Do not use ports 0-1024 because these ports are reserved ports.
5. In the **External IP Address of Firewall** box, type the IP address of your firewall.
6. In the **Actions** pane, click **Apply**.

Step 5: Configure User Isolation

The user isolation feature allows you to configure your FTP server to isolate users, which prevents users from accessing the directories of other users on the same FTP site. If you choose not to isolate users, they share a common directory structure.

For example, you can choose not to isolate users on a site that offers only download capabilities for shared content or for a site that does not require the protection of data between users.

If you want to isolate users on your site, you can choose one of the following isolation options:

1. **User name directory**: Isolates user sessions to the physical or virtual directory with the same name of the FTP user account. The user sees only their FTP root location and is, therefore, restricted from navigating higher up the physical or virtual directory tree. Any global virtual directories that are created are ignored.
2. **User name physical directory**: Isolates user sessions to the physical directory with the same name of the FTP user account. The user sees only their FTP root location and is, therefore, restricted from navigating higher up the physical directory tree. Any global virtual directories that are created apply to all users.
3. **FTP home directory configured in Active Directory**: Isolates user sessions to the home directory that is configured in the Active Directory account settings for each FTP user.

To configure user isolation

1. Open IIS Manager.
2. In **Features View**, double-click **FTP User Isolation**.
3. If you don't want to isolate users, under **Do not isolate users. Start users in**, select one of the following options:
   - **FTP root directory**: specifies that all FTP sessions start in the root directory for the FTP site. This option disables all user isolation and starting folder logic.
   - **User name directory**: specifies that all FTP sessions start in the physical or virtual directory with the same name of the currently logged-on user if the folder exists; otherwise, the FTP session starts in the root directory for the FTP site.
4. If you want to isolate users, under **Isolate users. Restrict users to the following directory**, select one of the following options:

   User name directory (disable global virtual directories): isolates user sessions to the physical or virtual directory with the same name of the FTP user account.

   User name physical directory (enable global virtual directories): isolates user sessions to the physical directory with the same name of the FTP user account.

   FTP home directory configured in Active Directory: isolates user sessions to the home directory that is configured in the Active Directory account settings for each FTP user.

   Custom: This option is an advanced feature, and enables developers to create custom providers that provide home directory lookups based on their unique business needs.
5. If you selected **FTP home directory configured in Active Directory** in the previous step, click the **Set** button, and then type a user name and password in the **User name** and **Password** boxes of the **Set Credentials** dialog box that has access to your Active Directory server. Enter the password again in the **Confirm Password** box, then click **OK**.
6. In the **Actions** pane, click **Apply**.

Step 6: Configure Directory Browsing Options

The directory browsing feature gives you control over what is displayed when users browse FTP directories.

To configure directory browsing options

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Features View**, double-click **FTP Directory Browsing**.
4. On the **FTP Directory Browsing** page, under **Directory Listing Style**, select one of the following options:
   - **MS-DOS**: Displays directory content in a manner consistent with MS-DOS.
   - **UNIX**: Displays directory content in a manner consistent with UNIX.
5. Under **Directory Listing Options**, select the information you want to display in directory listings. You can select any of the following options:
   a. **Virtual directories**: Shows virtual directories.
   b. **Available bytes**: Displays file size in bytes.
   c. **Four-digit years**: Displays years by using four digits rather than two.
6. In the **Actions** pane, click **Apply**.

Step 7: Configure Logon Attempt Restrictions

This feature enables you to configure the maximum number of failed sign-on attempts that you allow within a specified time before the IP address is denied.

To configure logon attempt restrictions

1. Open IIS Manager.
2. In the **Connections** pane, select the server level.
3. In **Features View**, double click **FTP Logon Attempt Restrictions**.
4. On the **FTP Logon Attempt Restrictions** page, select the **Enable FTP Logon Attempt Restrictions** check box.
5. In the **Maximum number of failed login attempts** box, type a positive integer. The default value is 4.
6. In the **Time period (in seconds)** box, type a positive integer. The default is 30.
7. Select either **Deny IP addresses based on the number of failed login attempts** or **Write to the log only**. If you choose **Write to the log only**, IIS will not restrict clients even if the maximum number of failed attempts is exceeded.
8. In the **Actions** page, click **Apply**.

Step 8: Configure Request Filtering

Use the **FTP Request Filtering** feature page to define the request filtering settings for your FTP site. FTP request filtering is a security feature that allows internet service providers (ISPs) and application service providers to restrict protocol and content behavior.

To configure request filtering

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Feature View**, double click the **FTP Request Filtering** icon.
4. To change the general request filtering settings, click **Edit Feature Settings** in the **Actions** pane. Then edit the general settings as required.
5. To add a filter based on file name extensions, select the **File Name Extension** tab. In the **Actions** pane, click either **Allow File Name Extension** or **Deny File Name Extensions**. Then type the extension in the **File name extension** box, and click **OK**.
6. To add a filter based on a URL segment (for example, a folder name), select the **Hidden Segments** tab. In the Actions pane, click **Add Hidden Segment**. Then type the segment in the **Hidden segment** box, and click **OK**.
7. To add a filter based on a URL sequence, select the **Denied URL Sequences** tab. In the **Actions** pane, click **Add URL Sequence**. Then type the sequence in the **URL sequence** box, and click **OK**.
8. To add a filter based on a command, select the **Commands** tab. In the **Actions** pane, click either **Allow Command** or **Deny Command**. Then type the command in the **Command** box, and click **OK**.

   Warning
   The ability to add a command-based filter is an advanced feature. If you use this feature incorrectly, you could deny access to all FTP clients on your server.

Step 9: Configure FTP Logging

You can use the **FTP Logging** feature to configure logging features at the server or site level, and to configure logging settings.

To configure FTP logging

1. Open IIS Manager.
2. In the **Connections** pane, select either the server level or the site level.
3. In **Feature View**, double click the **FTP Logging** icon.
4. In the **One log file per** menu, select either **Site** or **Server**.
5. Under **Log File**, click **Select W3C Fields**, and then select the information you want go log.
6. Under **Directory**, either type the path to the base folder where you want the FTP log files

stored, or click **Browse** to browse to the base folder.

7. Under **Encoding**, select either **UTF8** (single-byte and multi-byte characters) or **ANSI** (single-byte characters only).
8. Under **Log File Rollover**, select how you want FTP to create new log files from the following list:
   a. **Schedule**: select **Hourly**, **Daily**, **Weekly**, or **Monthly** to create new log files on a fixed time interval.
   b. **Maximum file size (in bytes)**: enter a positive integer to create new log files when the file size exceeds that number of bytes.
   c. **Do not create new log files**
9. Select the **Use local time for file naming and rollover** check box when you want log file naming and rollover to be based on the local time zone instead of Coordinated Universal Time (UTC).
10. In the **Actions** pane, click **Apply**.

Step 10: Configure FTP Messages

Use the **FTP Messages** feature to modify the settings for messages sent when a user connects to your FTP site.

To configure FTP messages

1. In the **Connections** pane, select either the server level or the site level.
2. In **Features View**, double-click **FTP Messages**.
3. On the **FTP Messages** page, under **Message Behavior**, select how you want your FTP messages to behave. You can select any of the following options:
   a. **Suppress default banner**: Specifies whether to display the default identification banner for the FTP server.
   b. **Support user variables in messages**: Specifies whether to display a specific set of user variables in FTP messages. The following user variables are supported:
      - %BytesReceived% - The number of bytes sent from the server to the client for the current session.
      - %BytesSent% - The number of bytes sent from the client to the server for the current session.
      - %SessionID% - The unique identifier for the current session.
      - %SiteName% - The name of the FTP site that is hosting the current session.
      - %UserName% - The account name of the currently logged-on user.
   c. **Show detailed messages for local requests**: Specifies whether to display detailed error messages when the FTP client is connecting to the FTP server on the server itself (local host).
4. Under Message Text, type messages in the following boxes:
   a. **Banner**: Specifies the message the FTP server displays when FTP clients first

connect to the FTP server.

b. **Welcome**: Specifies the message the FTP server displays when FTP clients have logged in to the FTP server.

c. **Exit**: Specifies the message the FTP server displays when FTP clients log off the FTP server.

d. **Maximum Connections**: Specifies the message the FTP server displays when clients try to connect and cannot because the FTP service has reached the maximum number of client connections allowed.

5. In the **Actions** pane, click **Apply**.

See Also

- Hosting-Friendly Web Server Platform (IIS): Scenario Overview
- Build a Static Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

IIS Manager Feature Overview

IIS Manager is the IIS 8 UI that helps you perform the following tasks:

- Manage IIS and ASP.NET in one tool.
- View health and diagnostic information that includes the ability to view currently running requests in real time.
- Configure user and role authorization for websites and web applications.
- Connect to a web server, site, or application remotely using HTTP.

This document provides an overview of the online help for the features included in the IIS Manager UI.

IIS manager features

The following table shows the features of IIS Manager and their descriptions.

Feature	Description
.NET Authorization Rules	Configure rules for authorizing users to access your websites and applications.
.NET Compilation	Manage ASP.NET application code configuration settings.

Feature	Description
.NET Error Pages	Configure HTTP error responses to return when errors occur.
.NET Globalization	Configure .NET Framework globalization settings.
.NET Profile	Manage a list of profile properties that are used to track any custom information your application requires.
.NET Roles	Manage a list of user groups. User groups offer the ability to categorize a set of users and perform security-related operations, such as authorization, on a defined set of users.
.NET Trust Levels	Configure the level of code access security (CAS) that is applied to an application.
.NET Users	Manage a list of user identities for an application.
Application Pools	Manage a list of application pools on a web server.
Application Settings	Manage a list of key/value pairs that are stored in the Web.config file of your application.
Applications	Manage the list of applications in a website.
ASP	Manage a list of classic ASP configuration settings.
Authentication	Configure the authentication methods that clients can use to gain access to your content.
Authorization Rules	Manage the list of Allow or Deny rules that control access to content.
Centralized Certificates	Configure and manage a central certificate store for a server farm.
CGI	Configure Common Gateway Interface (CGI) applications on the Web server.
Compression	Provide faster transmission times between IIS and compression-enabled browsers.
Configuration Editor	Manage your configuration files by pulling sections into IIS Manager.

Feature	Description
Connection Strings	Manage a list of connection strings generally used by managed code applications.
Default Document	Configure the list of default documents (such as Default.htm).
Directory Browsing	Modify the content settings for browsing a directory on the web server.
Error Pages	Manage a list of custom HTTP error messages.
Failed Request Tracing Rules	Manage a list of tracing rules for failed requests.
FastCGI Settings	Use the Fast Common Gateway Interface (FastCGI) protocol to communicate with external applications.
Feature Delegation	Configure the delegation state of IIS Manager features for sites and applications on your web server.
FTP Authentication	Configure the authentication methods that FTP clients can use to gain access to your content.
FTP Authorization Rules	Manage the list of Allow or Deny rules that control access to content.
FTP Directory Browsing	Modify the content settings for browsing a directory on the FTP server.
FTP Firewall Support	Modify the settings for passive connections when FTP clients are connecting to an FTP server that is located behind a firewall server.
FTP IP Address and Domain Restrictions	Define and manage rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a domain name or names.
FTP Logging	Configure logging features at the server or site level, and to configure logging settings.
FTP Logon Attempt Restrictions	Modify the number of times FTP allows a user to attempt unsuccessfully to log in within a specified time period before denying access to the IP address.

Feature	Description
FTP Messages	Modify the settings for messages sent when a user connects to your FTP site.
FTP Network Security	Configure login security parameters and protect your FTP server from network-based attacks.
FTP Request Filtering	Manage the request filtering settings that allow Internet service providers (ISPs) and application service providers to restrict protocol and content behavior.
FTP Sites	Represents the UI components that are available on the IIS Manager **Sites** page when FTP is installed.
FTP SSL Settings	Manage encryption for both control channel and data channel transmissions between your FTP server and clients.
FTP User Isolation	Manage the user isolation mode and restrict users to their own directories.
Handler Mappings	Manage a list of handlers that process requests for specific file types.
HTTP Redirect	Configure how incoming requests are redirected to a new destination.
HTTP Response Headers	Manage a list of name and value pairs that contain information about a requested page, and to configure common HTTP headers.
IIS Manager	Represents the UI components available at the top level of IIS Manager.
IIS Manager Permissions	Manage IIS Manager users, Windows users, and members of Windows groups that are allowed to connect to a site or an application.
IIS Manager Users	Manage user accounts that can be permitted to connect to sites or applications on the web server.
IP Address and Domain Restrictions	Define and manage rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a domain name or names.

Feature	Description
ISAPI and CGI Restrictions	Manage a list of restrictions to ISAPI and CGI programs.
ISAPI Filters	Manage a list of .dll files that change or enhance the functionality provided by IIS.
Logging	Configure how IIS logs requests made to the web server and when new log files are created.
Machine Key	Configure hashing and encryption settings used for application services, such as view state, Forms authentication, membership and roles, and anonymous identification.
Management Service	Configure the management service that enables computer and domain administrators to remotely manage a Web server that uses IIS Manager.
MIME Types	Manage a list of Multipurpose Internet Mail Extensions (MIME) types so you can identify the types of content that can be served from a Web server to a browser or a mail client.
Modules	Manage a list of native and managed code modules that perform specific tasks in the request-processing pipeline, such as authentication or compression.
Output Caching	Configure output caching settings and to configure rules that control caching of served content.
Pages and Controls	Configure the settings of ASP.NET pages and controls.
Providers	Manage a list of providers.
Request Filtering	Configure filtering rules that let you restrict protocol and content behavior.
Server Certificates	View the names of certificates, the fully qualified domain names (FQDNs) of hosts to which certificates have been issued, and the FQDNs of the servers that issued the certificates.

Feature	Description
Session State	Configure the behavior of information maintained across browser sessions.
Shared Configuration	Share IIS configuration files and encryption keys between one or more IIS servers.
Sites	Manage a list of sites on a web server.
SMTP E-mail	Configure delivery of e-mail from applications that use the System.Net.Mail API.
SSL Settings	Manage data encryption of transmissions between your server and clients.
Virtual Directories	Manage a list of virtual directories in an application.
WebDAV Authoring Rules	Manage the list of authoring rules that control access to content.
Worker Processes	Manage a list of worker processes running in application pools on a Web server.

.NET Authorization Rules

Use the **.NET Authorization Rules** page to configure rules for authorizing users to access your websites and applications. You can add Allow or Deny rules, and specify users, roles, or user groups to whom to allow access to your web content.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- **UI Elements for .NET Authorization Rules**
- Allow or Deny Authorization Rule Dialog Boxes

UI Elements for .NET Authorization Rules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Mode	Lists the type of action enabled by the .NET authorization rule. The value can be either Allow or Deny.
Users	Lists the users who are allowed or denied access to web content as specified by the .NET authorization rule.
Roles	Lists the roles that are allowed or denied access to web content as specified by the .NET authorization rule.
Verbs	Lists the verbs to which the .NET authorization rules are applied.
Entry Type	Lists the entry type for the .NET authorization rule.

Actions Pane Elements

Element Name	Description
Add Allow Rule	Opens the **Add Allow Authorization Rule** dialog box that lets you configure a rule allowing access to your web content to certain roles, users, or user groups.
Add Deny Rule	Opens the **Add Deny Authorization Rule** dialog box that lets you configure a rule denying access to your web content to certain roles, users, or user groups.
Edit	Opens the **Edit Allow Authorization Rule** dialog box or the **Edit Deny Authorization Rule** dialog box that lets you change the selected .NET authorization rule.
Remove	Removes the selected .NET authorization rule.

Allow or Deny Authorization Rule Dialog Boxes

Use the **Add Allow Authorization Rule**, **Edit Allow Authorization Rule**, **Add Deny Authorization Rule**, and **Edit Deny Authorization Rule** dialog boxes to add or edit the properties for .NET authorization rules.

For more information, see the **.NET Authorization Rules Page**.

Element Name	Description
All users	Select this option to manage access to content for both anonymous and authenticated users. **Note** Be sure to position this rule underneath any rules that grant access to content. If this rule is at the top of the rules list, all users are denied access to content.
All anonymous users	Select this option to manage access to content for users who are not authenticated. **Note** If you use this rule, all users must have a valid user account and password to be able to authenticate.
Specified roles or user groups	Select this option to manage access to content for specific roles or user groups. **Note** If you use this rule, all members of the specified roles and groups must have a valid Basic or Custom user account and password to be able to authenticate.
Specified users	Select this option to manage access to content for a specific user account. **Note** If you use this option, all users must have a valid Basic or Custom user account and password to be able to authenticate.
Apply this rule to specific verbs	Select this option to apply the authorization rule to specific verbs, such as GET or POST.

.NET Compilation

Use the **.NET Compilation** properties page to manage ASP.NET application code configuration settings.

Use the **Display** list to select one of the following options that specifies how the settings are displayed: **Friendly Names**, **Configuration Names**, or **Both Names**.

Related scenarios

- Build an ASP.NET Website on IIS

UI Elements for .NET Compilation

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Batch Compilations	Indicates whether batch processing is supported. The default is **true**.
Maximum File Size	Specifies the maximum size, in kilobytes, of the generated source files per batched compilation. The default is 1000 KB.
Maximum Size of Batch	Specifies the maximum number of pages per batched compilation. The default is 1000.
Time-out (hh:mm:ss)	Indicates the time-out period for batch compilation. If compilation cannot be completed in the time-out period, the compiler reverts to single compilation mode for the current page. The default is 00:15:00 (15 minutes).
Debug	Specifies whether release or debug binaries are used for compilation. When **true**, debug binaries are used for compilation. The default value is **false**.
Enable Prefetch	Enables ASP.NET to take advantage of Windows Prefetch functionality. The default value is **false**.

Element Name	Description
	Note The Windows Prefetch service must be configured and running for this setting to take effect.
Enable Profiling Optimizations	Specifies profiling optimizations enabled for managed assemblies using idle CPU cores. Allowed values are **All** or **None**. The default value is **All**.
Number of Recompiles	Indicates the number of dynamic recompiles of resources that can occur before the application restarts. The default is 15.
Optimize Compilations	Specifies whether dynamic compilation will recompile an entire site if a top-level file is changed. The default value is **false**.
Url Line Pragmas	Specifies whether instructions to the compiler use physical paths or URLs. When **true**, URLs are used for the compiler instructions instead of physical paths. The default value is **false**.
Explicit Compile Option	Specifies whether to set the Microsoft Visual Basic **explicit** compile option. If **True**, all variables must be declared by using a **Dim**, **Private**, **Public**, or **ReDim** statement. The default is **True**.
Strict Compile Option	Specifies whether to enable the Visual Basic **strict** compile option. When **True**, compilation explicitly disallows any data-type conversions in which data loss would occur, and any conversion between numeric types and strings. The default value is **False**.
Assemblies	Defines the set of assemblies, from the \bin folder or from the Global Assembly Cache (GAC), which is referenced during compilation. The GAC stores .NET assemblies that are designated to be shared by several applications on the computer. Use an asterisk ("*") to include every assembly in the directory.
Default Language	Specifies the default programming language,

Element Name	Description
	such as C# or Visual Basic, to use in dynamic compilation files. The default programming language is **vb**, which specifies Visual Basic as the default programming language.
Code Sub Directories	Defines subdirectories that contain files that are compiled at run time. You can have any number of subdirectories.
Temporary Directory	Specifies the directory to use for temporary file storage during compilation. The default is an empty string that puts temporary files in a location under the Temporary ASP.NET Files directory.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

.NET Error Pages

Use the **.NET Error Pages** feature to configure HTTP error responses to return when errors occur. You can create custom error pages or detailed error messages that contain troubleshooting information.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for .NET Error Pages
- Add or Edit Custom Error Page Dialog Box
- Edit ASP.NET Error Pages Settings Dialog Box

UI Elements for .NET Error Pages

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Group by	Use this drop-down list to sort your error pages by **No Grouping** or **Entry Type**. This helps you find certain error pages easily, such as the Local entry type.
Status Code	Lists the status code for the .NET error page.
Redirect	Lists the URL address for the .NET error page to which the user is redirected when an error occurs.
Entry Type	Lists the entry type for the .NET error page.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Custom Error Page** dialog box that lets you configure an error page. Enter the status code and URL to which the user is redirected when an error occurs.
Edit	Opens the **Edit Custom Error Page** dialog box that lets you change the URL for the selected error page. You can change the status code by using the **Change Status Code** option in the **Actions** pane.
Edit Feature Settings	Opens the **Edit Error Pages Settings** dialog box that lets you configure the mode and default page properties for all error pages. For more information, see **Edit ASP.NET Error Pages Settings Dialog Box**.
Change Status Code	Allows you to change the status code for the selected error page.

Element Name	Description
Remove	Removes the selected error page.

Add or Edit Custom Error Page Dialog Box

Use the **Add Custom Error Page** and **Edit Custom Error Page** dialog boxes to add or edit the status code and URL for your error pages.

For more information, see the **.NET Error Pages** topic.

Element Name	Description
Status Code	Enter the status code for your error page. You cannot edit the status code by using the **Edit Custom Error Page** dialog box. Use the **Change Status Code** option in the **Actions** pane instead.
Absolute URL	Enter the URL address to which you want to redirect users when an error occurs.

Edit ASP.NET Error Pages Settings Dialog Box

Use the **Edit Error Pages Settings** dialog box to configure the mode and default page properties for your error pages. These settings affect all error pages. For more information, see **.NET Error Pages**.

Element Name	Description
Mode	Select a mode for the error pages: **On**, **Off**, or **Remote Only**.
Redirect Mode	Select a mode for response redirection: **ResponseRedirect** or **ResponseRewrite**. **ResponseRedirect** causes the browser to redirect the client to a different URL. **ResponseRewrite** modifies the request URL before it is returned to the client. The default is **ResponseRedirect**.
Default Page	**Absolute URL** specifies the default URL address to which to direct a browser, if an error occurs. If this attribute is not specified, you

Element Name	Description
	receive a generic error message instead. The URL can be absolute or relative. An example of an absolute URL is `www.contoso.com/ErrorPage.htm`. A relative URL, such as `/ErrorPage.htm`, is relative to the Web.config file that specified the URL for this attribute, not to the webpage in which the error occurred. A URL starting with a tilde (~), such as `~/ErrorPage.htm`, indicates that the specified URL is relative to the root path of the application.
Miscellaneous	Select **Allow Nested Errors** to specify whether additional error information is displayed in browsers when a server-side error occurs during processing of an HTTP redirect for a custom error. The default is **False**.

.NET Globalization

Use the **NET Globalization** feature page to configure .NET Framework globalization settings. Globalization is the process of internationalizing application code, then localizing the application to other languages and cultures. The internationalization process makes it possible to translate, store, retrieve, and present application content for any locale, using the same application code base whenever possible. Locale is the combination of language and cultural environment. It includes date formats, times, currencies, telephone numbers, and so on. Localization means adapting your application to other locales by translating and formatting content according to culture, preferably without touching the code itself.

You can configure the following .NET Framework globalization settings in IIS:

- Culture settings such as the UI culture
- Encoding settings such as encoding for response headers

Related scenarios

- Build an ASP.NET Website on IIS

UI Elements for .NET Globalization

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Culture	Specifies the default culture for processing incoming web requests.
Enable Client Based Culture	Gets or sets a value indicating whether the **Culture** and **UICulture** properties should be based on the **AcceptLanguage** header field value that is sent by the client browser.
UI Culture	Specifies the default culture for processing locale-dependent resource searches.
File	Specifies the file encoding for .aspx, .asmx, and .asax file parsing. Unicode and UTF-8 files that are saved with the byte-order mark prefix are automatically recognized, regardless of the value for this attribute.
Requests	Specifies the assumed encoding of each incoming request, including posted data and the query string. If the request includes a request header that contains an **Accept-Charset** attribute, it overrides this attribute in the configuration. The default encoding is UTF-8, which is specified in the globalization section in the Machine.config file that was created when the .NET Framework was installed. If request encoding is not specified in a Machine.config or Web.config file, encoding defaults to the Regional and Language Options locale setting for the computer. In single-server applications, this attribute and the **responseEncoding** attribute should be the same. For the less common case (multiple-server applications where the default server encodings are different for each server), you can vary the request and response encoding by using local Web.config files.
Response Headers	Specifies the content encoding of response headers.

Element Name	Description
Responses	Specifies the content encoding of responses. The default encoding is UTF-8, which is specified in the globalization section in the Machine.config file that was created when the .NET Framework was installed. If response encoding is not specified in a Machine.config or Web.config file, encoding defaults to the Regional and Language Options locale setting for the computer. In single-server applications, this attribute and the **responseEncoding** attribute should be the same. For the less common case (multiple-server applications where the default server encodings are different for each server), you can vary the request and response encoding by using local Web.config files.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

.NET Profile

Use the **.NET Profile** feature page to manage a list of profile properties that are used to track any custom information your application requires.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for .NET Profile

- Add or Edit .NET Profile Property Dialog Box
- Add or Edit Group Dialog Box
- Edit Profile Settings Dialog Box

UI Elements for .NET Profile

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the profile property.
Data Type	Displays the data type for the property. This will be either a common data type such as **String**, **Int32**, **DateTime**, **StringCollection**, and so on, or a custom data type that contains custom information that is passed to the profile provider.
Default Value	Displays the value to which the property is initialized.
Allow Anonymous	Displays whether the property is managed for anonymous users. The value is either **True** or **False**.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Property	Opens the **Add .NET Profile Property** dialog box from which you can add a .NET profile property.
Add Group	Opens the Add Group dialog box from which you can add a group to group .NET profile

Element Name	Description
	properties.
Add Property to Group	Opens the **Add .NET Profile Property to** *Name* **Group** dialog box from which you can add .NET profile properties to the selected group.
Edit	Opens the **Edit .NET Profile Property** or **Edit .NET Profile Property in** *Name* **Group** dialog box from which you can edit the selected .NET profile property.
Rename	Enables the **Name** field of the selected .NET profile property so that you can rename the property.
Remove	Removes the item that is selected from the list on the feature page.
Disable	Disables the .NET Profile feature.
Set Default Provider	Opens the **Edit Profile Settings** dialog box from which you can select a default provider for the feature.
Providers	Opens the **Providers** feature page from which you can add additional providers for the .NET Profile feature.

Add or Edit .NET Profile Property Dialog Box

Use the **.NET Profile Property** dialog box to set or edit profile properties that track any custom information that your application requires, including:

- User information, such as an address or city.
- Preferences, such as a color scheme or list of stocks to follow.
- Custom information about the current session, such as a shopping cart.

Element Name	Description
Name	Defines the name of the profile property.
Data Type	Specifies the type for the property. The default is **String**. You can specify any .NET class as the type (**Int32**, **DateTime**, **StringCollection**,

Element Name	Description
	and so on).
Default value	Specifies the value with which the property is initialized.
Serialization option	Specifies the serialization formatter (**String**, **Binary**, **XML**, or **Provider Specific** serialization). The default is **String**.
Read only	Specifies whether the property can be modified.
Available for anonymous users	Specifies whether the property is managed for anonymous users. By default, this check box is clear.

Add or Edit Group Dialog Box

Use the **Add Group** dialog box to add a profile group. Groups are helpful for organizing profile properties.

Element Name	Description
Name	Defines the name of the profile group.

Edit Profile Settings Dialog Box

Use the **Edit Profile Settings** dialog box to define the default interface between the ASP.NET membership service and membership data sources.

Element Name	Description
Default Provider	Defines the default interface between the ASP.NET membership services and membership data sources. **Note** The default ASP.NET computer configuration contains a default **SqlProfileProvider** instance named **AspNetSqlProfileProvider** that connects to SQL Server on the local computer. By default, the ASP.NET profile feature uses this instance of the

Element Name	Description
	provider. Alternatively, you can specify a different default provider. For more information about ASP.NET profile providers, see ASP.NET Profile Providers.

.NET Roles

Use the **.NET Roles** feature page to manage a list of user groups. User groups offer the ability to categorize a set of users and perform security-related operations, such as authorization, on a defined set of users.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for .NET Roles
- Add .NET Role Dialog Box
- Edit .NET Roles Settings Dialog Box

UI Elements for .NET Roles

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Lists the name of the role.
Users	Shows the number of users assigned to a role.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add .NET Role** dialog box from which you can add role names to this provider.
View Users	Opens the **.NET Users** feature page, filtered by the selected role.
Rename	Enables you to edit the selected role.
Remove	Removes the item that is selected from the list on the feature page.
Disable	Disables the **.NET Roles** feature page.
Set Default Provider	Opens the **Edit .NET Roles Settings** dialog box, so you can set the default provider.
Providers	Opens the **Providers** feature page.

Add .NET Role Dialog Box

Use the **Add .NET Role** dialog box to add role names to the database referenced by this provider. Roles enable you to categorize users and then perform a security related task, such as authorization, more efficiently.

Element Name	Description
Name	Requests the name of the role.

Edit .NET Roles Settings Dialog Box

Use the **Edit .NET Roles Settings** dialog box to set the default provider for membership roles.

Element Name	Description
Default Provider	Sets the default provider for .NET Roles. IIS 8 includes the following providers: • **AspNetSqlRoleProvider** - Stores role information in a SQL Server database. • **AspNetWindowsTokenRoleProvider** - Uses role information based on Windows domain accounts. The Windows provider is useful only if your application runs on a network where all users have domain

Element Name	Description
	accounts.

.NET Trust Levels

Use the **.NET Trust Levels** feature page to set the trust element in the Web.config file. The trust element enables you to configure the level of code access security (CAS) that is applied to an application.

Related scenarios

- Build an ASP.NET Website on IIS

UI Elements for .NET Trust Levels

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description

Trust Level	Displays the name of the trust level. Available options include:
	• **Full (internal) -** Specifies unrestricted permissions. Grants the ASP.NET application permissions to access any resource that is subject to operating system security. All privileged operations are supported.
	• **High (web_hightrust.config) -** Specifies a high level of code access security, which means that the application cannot do any one of the following things by default:
	  • Call unmanaged code.
	  • Call serviced components.
	  • Write to the event log.
	  • Access Message Queuing service queues.
	  • Access ODBC, OleDb, or Oracle data sources.
	• **Medium (web_mediumtrust.config) -**

Element Name	Description
	Specifies a medium level of code access security, which means that, in addition to High Trust Level restrictions, the ASP.NET application cannot do any of the following things by default: • Access files outside the application directory. • Access the registry. • Make network or web service calls. • **Low (web_lowtrust.config) -** Specifies a low level of code access security, which means that, in addition to Medium Trust Level restrictions, the application cannot do any of the following things by default: • Write to the file system. • Call the Assert method. • **Minimal (web_minimaltrust.config) -** Specifies a minimal level of code access security, which means that the application has only execute permissions.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

.NET Users

Use the **.NET Users** feature page to manage a list of user identities for an application.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for .NET Users
- Edit .NET User Dialog Box
- Add .NET User Roles
- Add .NET User Account Details
- Edit .NET Users Settings Dialog Box

UI Elements for .NET Users

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description.
Name	Displays a unique user identity.
E-mail Address	Displays the e-mail address for a user.
Created	Displays the date the user identity was created.
Last Login	Displays the date on which the user last logged in to the application.

Actions Pane Elements

Element Name	Description
Add	Opens the **.NET User Account Details** wizard from which you can add a new user.
Set Default Provider	Opens the **Edit .NET Users Settings** dialog box, so you can set the default provider.
Reset Password	Resets the selected user's password.
Edit	Opens the **Edit .NET User** dialog box from which you can edit a user's role or e-mail address. This action is available only when an item is selected from the list on the feature page.
Remove	Removes the item that is selected from the list

Element Name	Description
	on the feature page.
Providers	Opens the **Providers** feature page.
Roles	Opens the **.NET Roles** feature page.

Edit .NET User Dialog Box

Use the **Edit .NET User** dialog box to modify a user's role or e-mail address.

Element Name	Description
User Name	Defines a unique user identity.
E-mail	Sets a default e-mail address for a user.
Roles	Defines the role or roles to which a user belongs.

Add .NET User Roles

Use **.NET User Roles** Wizard to specify roles for a new user. Select the roles for the user, and then click **Finish**. If you want to edit user information that is on the previous page of the wizard, click **Previous**.

Element Name	Description
Users	Specifies user names available to assign roles.
Roles	Specifies a list of available roles from the default provider. You can associate a single user name with multiple roles.

Add .NET User Account Details

Use the **.NET User Account Details** wizard to add a new user. If the Roles feature is not enabled, the **.NET User Account Details** page is the only page in the wizard. Enter the information for the new user and then click **Next** to add roles, or **Finish** if the Roles feature is not enabled.

Element Name	Description
User Name	Defines a new user identity. This value must be unique.
E-mail	Sets a default e-mail address for a user.
Password	Configures a password for the user.
Confirm Password	Confirms that the password matches what was entered in the **Password** box.
Question	Specifies a security question that users can use to reset their passwords. Select one of the default questions from the drop-down list, or enter a custom question.
Answer	Specifies the answer to the security question.

Edit .NET Users Settings Dialog Box

Use the **Edit .NET Users Settings** dialog box to set the default provider.

Element Name	Description
Default Provider	Sets the default provider for .NET Users. IIS 8 includes the following provider: • **AspNetSqlMembershipProvider** - Use this provider for storing user information in a SQL Server database.

Application Pools

Use the **Application Pools** feature page to manage a list of application pools on a Web server.

Sort the list by clicking one of the column headings on the feature page or select a value from the **Group by** drop-down list to group similar items. Search the list by typing search criteria in the **Filter** box, expand the **Go** list to select the field on which to search, and then click **Go** or press ENTER. Click **Show All** to remove the filter criteria.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS

- Build a PHP Website on IIS

In this document

- UI Elements for Application Pools
- Add or Edit Application Pool Dialog Box
- Application Pool Identity Dialog Box
- Recycling Conditions Wizard
- Recycling Events to Log Wizard

UI Elements for Application Pools

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of each application pool.
Status	Displays whether an application pool is started.
.NET Framework Version	Displays the version of the .NET Framework that an application pool uses. If the application pool was configured to not use the .NET Framework, the value is **No Managed Code**.
Managed Pipeline Mode	Displays the request-processing mode that an application pool uses. Values can be **Integrated** or **Classic**. When the value is **Integrated**, IIS uses the integrated IIS and ASP.NET request-processing pipeline to process requests in the application pool. When the value is **Classic**, IIS uses separate IIS and ASP.NET request-processing pipelines to process requests in the application pool.
Identity	Displays the identity used by an application pool. Values can be **Network Service**, **Local Service**, **Local System**, or the name of the user account for a custom identity.
Applications	Displays the number of applications that an application pool contains.

Actions Pane Elements

Element Name	Description
Add Application Pool	Opens the **Add Application Pool** dialog box from which you can add an application pool to the Web server.
Set Application Pool Defaults	Opens the **Application Pool Defaults** dialog box from which you can set default values that apply to all application pools that you add to the web server. When you select a setting on the dialog box, a description of the setting appears at the bottom.
Start	Starts the selected application pool.
Stop	Stops the selected application pool. And causes the Windows Process Activation Service (WAS) to shut down all running worker processes serving that application pool. An administrator must restart a stopped application pool or else requests made to applications in that application pool will receive HTTP 503-Service Unavailable errors.
Recycle	Stops and restarts the selected application pool. Restarting an application pool causes the application pool to be temporarily unavailable until the restart is complete.
Basic Settings	Opens the **Edit Application Pool** dialog box from which you can edit the settings that were specified when the application pool was created. This action is available only when an item is selected from the list on the feature page.
Recycling	Opens the **Edit Application Pool Recycling Settings** wizard from which you can specify conditions under which to recycle an application pool and configure how recycling events are logged.
Advanced Settings	Opens the **Advanced Settings** dialog box from which you can configure advanced settings for the selected application pool.

Element Name	Description
Rename	Enables the **Name** field of the selected application pool so that you can rename the application pool.
Remove	Removes the item that is selected from the list on the feature page. **Note** You cannot remove an application pool that contains applications. You must move any applications out of the application pool before you can remove it.
View Applications	Opens the **Applications** feature page from which you can view the applications that belong to the selected application pool.

Add or Edit Application Pool Dialog Box

Use the **Add Application Pool** and **Edit Application Pool** dialog boxes to create or edit an application pool. An application pool is a group of one or more URLs that are served by a worker process or a set of worker processes. Application pools set boundaries for the applications that they contain; these boundaries prevent applications in one application pool from affecting applications in another application pool.

Element Name	Description
Name	Type a unique, friendly name for your application pool. You can use this name to identify the application pool later, and to simplify the assignment of applications to the application pool. **Note** You cannot rename an application pool from the **Edit Application Pool** dialog box. To rename an application pool, use the **Rename** action in the **Actions** pane on the **Application Pools** feature page.
.NET Framework version	Select the version of the .NET Framework that

Element Name	Description
	is loaded by this application pool. If the applications that you want to assign to this application pool do not contain managed code, select the **No Managed Code** option from the list. An application pool can load only one version of the .NET Framework; all applications in the application pool must use the same version. **Note** If you select .NET Framework version 1.1, the **Managed pipeline mode** list will be disabled. Additionally, if you are running IIS 8 on 64-bit Windows, and you have installed .NET Framework version 1.1, the value will not appear in the **.NET Framework Version** list when you add or edit an application pool. This is because .NET Framework version 1.1 installs under the Framework directory, but IIS reads values from the Framework64 directory on 64-bit Windows. You must set the **managedRuntimeVersion** attribute to **v1.1** for the application pool in the configuration file directly or by using a tool such as Appcmd.exe.
Managed pipeline mode	Select one of the following values from the list to specify how IIS processes requests for managed content: • **Integrated** - IIS processes requests for managed content with the integrated IIS and ASP.NET request-processing pipeline. • **Classic** - IIS processes requests for managed content with the separate IIS and ASP.NET request-processing pipelines. Use this mode only when the applications in the application pool cannot run in **Integrated** mode.
Start application pool immediately	Select this option to enable the application pool to start whenever the Windows Process

Element Name	Description
	Activation Service (WAS) is started. If you disable automatic startup, you must start the application pool manually or else requests made to applications in the application pool will return an HTTP 503-Service Unavailable error.

Application Pool Identity Dialog Box

Use the **Application Pool Identity** dialog box to configure the account under which an application pool's worker process runs. You can select one of the predefined security accounts or configure a custom account.

Note

If you use a custom identity, make sure that the user account you specify is a member of the IIS_IUSRS group on the Web server so that the account has proper access to resources. Additionally, when you use Windows and Kerberos authentication in your environment, you might need to register a Service Principle Name (SPN) with the domain controller (DC).

To configure an application pool identity

1. On the **Application Pools** feature page, select an application pool from the list.
2. In the **Actions** pane, click **Advanced Settings**.
3. Under **Process Model**, locate the **Identity** field and click **…** (the Properties button) to open the **Application Pool Identity** dialog box.
4. Select an identity option:
   - Select **Built-in account** to use a predefined security account, and then select one of the accounts from the list.
   - Select **Custom account** to configure a custom account, and then click **Set** to open the **Set Credentials** dialog box from which you can specify the user name and password for the account.

Element Name	Description
Built-in account	Select this option to use one of the predefined security accounts. Then select one of the following accounts: • **LocalService** - The Local Service account is a member of the Users group and has the same user rights as the Network Service account, but the Local Service

Element Name	Description
	account is limited to the local computer. Use this account when the worker process in your application pool does not require access outside the Web server on which it runs. • **LocalSystem** - The Local System account has all user rights, and it is part of the Administrators group on the Web server. Whenever possible, avoid using the Local System account because it presents a more serious security risk for your Web server. • **NetworkService** - By default, the Network Service account is selected. It is a member of the Users group and has user rights that are required to run applications. It can interact throughout an Active Directory-based network by using the credentials of the computer account. This account provides the most security against an attack that might try to take over the Web server.
Custom account	Select this option to configure a custom account. Then click the corresponding **Set** button to configure the user name and password for the account.
Set	Opens the **Set Credentials** dialog box from which you can specify credentials for the custom account.
User name	Type the alias of the user account under which you want the worker process to run.
Password	Type the password for the user account that you specified in the **User name** box.
Confirm password	Retype the password for the user account that you specified in the **User name** box.

Recycling Conditions Wizard

Use the **Recycling Conditions** page of the **Edit Application Pool Recycling Settings** Wizard to configure IIS to periodically restart worker processes in an application pool. Restarting worker

processes can help you to recover valuable system resources and to manage faulty worker processes better.

Element Name	Description
Regular time intervals (in minutes)	Select this option to specify a time interval, in minutes, at which you want IIS to recycle the worker process. You might choose this option if you have an application that causes problems when it runs for an extended time. Based on what you know about the application, you should set the value to be less than the length of time elapsed before application failure.
Fixed number of requests	Select this option to specify the number of requests after which you want IIS to recycle the worker process. You might choose this option if you have an application that causes problems after reaching a certain number of requests. Based on what you know about the application, you should configure the value to be less than the number of requests processed before application failure.
Specific time(s)	Select this option to specify a time or times at which you want IIS to recycle the worker process in a 24-hour period. For example, to recycle a worker process at 4:30 A.M. and 4:30 P.M., enter **4:30 AM, 4:30 PM**. The time that you specify uses the local time on the Web server. You might choose this option if you have an application that causes problems when it runs for an extended time and you want to recycle the application pool at a specific time, such as a time that is late at night or early in the morning, to avoid a negative impact on users. Based on what you know about the application, you should set the interval to be frequent enough to prevent application failure.
Virtual memory usage (in KB)	Select this option to specify the maximum number of kilobytes of your system's common virtual memory that can be used by a worker process before that process is recycled. You might choose this option when you notice a

Element Name	Description
	steady increase in the virtual memory used on your server. This might indicate that an application reserves memory multiple times, which fragments the memory heap. Entering too high a value can severely decrease system performance. At first, you should set the virtual memory threshold to be less than 70 percent of available virtual memory, and then adjust the setting if you have to.
Private memory usage (in KB)	Select this option to specify the maximum privately allocated system physical memory that a worker process can use before the process is recycled. You might choose this option when you have an application that leaks memory. Entering too high a value can severely decrease system performance. At first, you should set this value to be less than 60 percent of the available physical memory on the server, and then adjust this setting if you have to.

Recycling Events to Log Wizard

Use the **Recycling Events to Log** page of the **Edit Application Pool Recycling Settings** Wizard to configure IIS to log an event when a worker process is recycled.

You can configure IIS to log recycling events that you set up, such as at a fixed interval. You can also configure IIS to log recycling events that occur at runtime, such as when an ISAPI declares itself unhealthy.

Element Name	Description
Regular time intervals	Select this option to log an event when a worker process is recycled at a specified time interval. This option is available only when the **Regular time intervals (in minutes)** option is selected and a time interval is specified on the previous wizard page.
Virtual memory usage	Select this option to log an event when a worker process is recycled after using a specified amount of virtual memory. This option is available only when the **Virtual memory**

Element Name	Description
	usage (in KB) option is selected and a number of kilobytes is specified on the previous wizard page.
Number of requests	Select this option to log an event when a worker process is recycled after reaching a specified number of requests. This option is available only when the **Fixed number of requests** option is selected and a number of requests is specified on the previous wizard page.
Scheduled time(s)	Select this option to log an event when a worker process is recycled at a specified time. This option is available only when the **Specific time(s)** option is selected and a time is specified on the previous wizard page.
Private memory usage	Select this option to log an event when a worker process is recycled after using a specified amount of physical memory. This option is available only when the **Private memory usage (in KB)** option is selected and a number of kilobytes is specified on the previous wizard page.
On-demand	Select this option to log an event when you recycle a worker process by using IIS Manager or Appcmd.exe to correct a problem.
Configuration changes	Select this option to log an event when a change to configuration causes the application pool to recycle.
Unhealthy ISAPI	Select this option to log an event when an ISAPI extension reports to the worker process that it is unhealthy.

Application Settings

Use the **Application Settings** feature page to manage a list of key/value pairs that are stored in the Web.config file of your application. Application settings can be accessed from anywhere within the application.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Application Settings
- Add or Edit Application Setting Dialog Box

UI Elements for Application Settings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the application setting variable.
Value	Displays the value assigned to the application setting variable.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Application Setting** dialog box in which you can add an application setting variable and value.
Edit	Opens the **Edit Application Setting** dialog box in which you can edit the selected application setting.

Element Name	Description
Remove	Removes the item that is selected from the list on the feature page.

Add or Edit Application Setting Dialog Box

Use the **Add Application Setting** and **Edit Application Setting** dialog boxes to add or edit global variables that are stored in a central location. These name/value pairs can be accessed from anywhere in an application.

Element Name	Description
Name	Defines a unique name for the application setting.
Value	Defines a value for the application setting.

Applications

Use the **Applications** feature page to manage the list of applications in a website.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items. Search the list by typing search criteria in the **Filter** box, expand the **Go** list to select the field on which to search, and then click **Go** or press ENTER. Click **Show All** to remove the filter criteria.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Applications
- Add or Edit Application Dialog Box

UI Elements for Applications

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Virtual Path	Displays the virtual path used in a URL to access the application.
Physical Path	Displays the physical path of the directory that holds application content.
Site	Displays the site to which the application belongs.
Application Pool	Displays the application pool in which the application runs.

Actions Pane Elements

The following table describes the actions available when connected to a server.

When you are connected to a site and select an application on the **Applications** page, the available actions are **Browse** and **View Virtual Directories**.

Element Name	Description
Add Application	Opens the **Add Application** dialog box in which you can add an application.
Set Application Defaults	Opens the **Application Defaults** dialog box in which you can configure default settings that apply to all applications.
Change Application Pool	Opens the **Select Application Pool** dialog box in which you can change the application pool to which the selected application belongs. This action is the only available action on the **Applications** page that displays when you view applications in an application pool. This action is not available on the **Applications** page that displays when you view applications in a site.
Explore	Opens the physical directory (in Windows Explorer) that is mapped to the root virtual directory for the selected application.
Edit Permissions	Opens the Windows Properties dialog box for

Element Name	Description
	the physical directory that is mapped to the root virtual directory for the selected application.
Remove	Removes the item that is selected from the list on the feature page.
Browse	Opens the selected application in an internet browser.
Basic Settings	Opens the **Edit Application** dialog box in which you can edit the settings that were specified when the selected application was created.
Advanced Settings	Opens the **Advanced Settings** dialog box in which you can configure advanced settings for the selected application.
View Virtual Directories	Opens the **Virtual Directories** feature page in which you can view the virtual directories that belong to the selected application.

Add or Edit Application Dialog Box

Use the **Add Application** and the **Edit Application** dialog boxes to add an application to a site, or to edit an existing application. An application is a grouping of files that delivers content or provides services over protocols, such as HTTP. In addition to belonging to a site, an application belongs to an application pool, which isolates the application from applications in other application pools on the server. When you create an application in IIS, the application's path becomes part of the site's URL. You designate a directory as the application root, or starting point, for the application and then specify properties specific to that application.

Element Name	Description
Site name	Displays the name of the site that will contain the application.
Path	Displays the application's path. If you create an application at the site level, this box will display */*. If you create an application within another application, this box will display the path of the existing application that will contain the new application, such as /myApp.

Element Name	Description
Alias	Type a name that clients can use to access content from a Web browser. For example, if your site address is http://www.contoso.com/ and you create an application for this site called /marketing, users can access the application from their Web browser by typing http://www.contoso.com/marketing.
Application pool	Displays the selected application pool for the application.
Select	Opens the **Select Application Pool** dialog box from which you can select the application pool in which the application will run.
Physical path	Type or navigate to the physical path where the application content is stored. The content can reside on the local computer or come from a remote share. If your content is stored on the local computer, enter the physical path, such as C:\Content. If your content is stored in a remote share, enter the UNC path, such as \\Server\Share. Optionally, click **Connect as** to provide credentials for an account that is authorized to access content in the physical path.
…	Opens the **Browse for Folder** dialog box from which you can select the physical location of your application content.
Connect as	Opens the **Connect As** dialog box from which you can select how to connect to the path that you typed in the **Physical path** box. By default, **Application user (pass-through authentication)** is selected.
Test Settings	Opens the **Test Settings** dialog box from which you can view a list of test results to assess whether the path settings are valid.

ASP

Use the **ASP** properties page to manage a list of ASP (Classic) configuration settings.

Use the **Display** list to select one of the following options that specifies how the settings are displayed: **Friendly Names**, **Configuration Names**, or **Both Names**.

Related scenarios

- Build a Classic ASP Website on IIS

UI Elements for ASP

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Code Page	Specifies the default code page for an ASP application.
Enable Buffering	Specifies whether output from an ASP application is buffered. The default is **True**.
Enable Chunked Encoding	Specifies whether HTTP 1.1 chunked transfer encoding is enabled for the World Wide Web publishing service (WWW service). The default is **True**.
Enable HTML Fallback	Controls the behavior of ASP when the request queue is full. When set to **True**, ASP searches for an HTML file that has _asp added to the file name and return the file if found. For example, if the name of the requested .asp file is Hello.asp, the name of the .htm file that is returned is Hello_asp.htm. If the file does not exist, a 500.13 (Server too busy) HTTP error is returned to the client. This error isl also returned if the value is set to **False**. The default value for this attribute is **True**.
Enable Parent Paths	Specifies whether an ASP page allows paths relative to the current directory (using the "...\" notation) or above the current directory. The default is **False**.

Element Name	Description
Client Connection Test Interval	Specifies the amount of time that a request should be queued. If the request is queued longer than the time specified, ASP checks to determine whether the client is still connected before running the request. If the client is no longer connected, the request is not processed and is deleted from the queue. The default is 00:00:03 seconds.
Maximum Requesting Entity Body Limit	Specifies the maximum number of bytes allowed in the entity-body of an ASP request. The default is 200000 bytes.
Queue Length	Specifies the maximum number of concurrent ASP requests that are permitted in the queue. The default is 3000.
Request Queue Time-out	Specifies the amount of time that an ASP script request is allowed to wait in the queue. The default is 00:00:00 (infinite).
Response Buffering Limit	Sets the maximum size of the ASP buffer. If response buffering is turned on, this property controls the maximum number of bytes that an ASP page can write to the response buffer before a flush occurs. The default is 4194304 bytes.
Script Time-out	Specifies the default length of time that ASP pages allow a script to run before attempting to stop the script and writing an event to the Windows event log. The default is 00:01:30 minutes.
Threads Per Processor Limit	Specifies the maximum number of worker threads per processor that IIS can create to run ASP requests. The default is 25.
Locale ID	Sets the default locale identifier (LCID) for an application. LCIDs define how dates, times, and currencies are formatted. The default is 1033, which represents US English (en-us).
Restart On Config Change	Determines whether an ASP application can be automatically restarted when a change is made to critical configuration properties of that

Element Name	Description
	application. The default is **True**.
Calculate Line Numbers	Specifies whether ASP should calculate and store the line number of each executed line of code in order to provide that number in an error report. The default is **True**.
Catch COM Component Exceptions	Specifies whether ASP pages trap exceptions thrown by components. If set to **False**, the ASP page does not catch exceptions thrown by the component. This can lead to the exception being handled by some other layer of the process, such as VB Script or the IIS worker process (which typically ends the worker process). The default is **True**.
Enable Client-side Debugging	Specifies whether ASP client-side debugging is enabled. The default is **False**.
Enable Log Error Requests	Controls whether the Web server writes ASP errors to the application section of the Windows event log. By default, ASP errors are written to the client browser and to the IIS log files. The default is **True**.
Enable Server-side Debugging	Specifies whether ASP debugging is enabled on the server. The default is **False**.
Log Errors to NT Log	Specifies which ASP errors are written to the Windows event log. By default, ASP errors are written to the client browser and to the IIS log files. The default is **False**.
Run On End Functions Anonymously	Specifies whether the **SessionOnEnd** and **ApplicationOnEnd** global ASP functions should be run as the anonymous user. If set to **False**, the functions are not run. The default is **True**.
Script Error Message	Specifies the error message to send to the browser if specific debugging errors are not sent to the client. The default message sent is "An error occurred on the server when processing the URL. Please contact the system administrator."

Element Name	Description
Send Errors To Browser	Specifies whether the Web server writes debugging specifics (file name, error, line number, description) to the client browser in addition to logging them to the IIS request log. The default is **True**.
Script Language	Specifies the default script language for all ASP applications that are running on the Web server. The default is VBScript.
Cache Directory Path	Contains the name of the directory that ASP uses to store compiled ASP templates to disk after overflow of the in-memory cache. The default is *systemroot*\inetpub\temp\ASP Compiled Templates.
Enable Type Library Caching	Specifies whether type libraries are cached on the server. The default is **True.**
Maximum Disk Cached Files	Specifies the maximum number of compiled ASP templates that can be stored. The default is 2000.
Maximum Memory Cached Files	Specifies the number of precompiled script files to cache. If set to 0, no script files are cached. If set to 4291967295, all the requested script files are cached. This property is used to tune performance based on the amount of available memory and the amount of script file traffic. The default is 500.
Maximum Script Engines Cached	Specifies the maximum number of scripting engines that ASP pages keep cached in memory. The default is 250.
Enable Side by Side Component	Enables COM+ side-by-side assemblies, which allow ASP applications to specify which version of a system DLL or COM component to use. The default is **False**.
Enable Tracker	Enables COM+ tracker, which then lets administrators or developers debug ASP applications. The default is **False**.
Execute In MTA	Specifies whether ASP can run scripts in a multithreaded compartment. The default is

Element Name	Description
	False.
Honor Component Threading Model	Specifies whether IIS examines the threading model of any components that your application creates. The default is **False**.
Partition ID	Set this property to the GUID of the COM+ partition. When this property is defined, set the **Use Partition** element to **True**.
Side By Side Component	Set this property to the name of the COM+ application. When this property is defined, set the **Enable Side by Side Component** element to **True**.
Use Partition	Isolates applications into their own COM+ partitions. When this property is set to **True**, you must define a value for the **Partition ID** element. The default is **False**.
Enable Session State	Enables session state persistence for the ASP application. The default is **True**.
Maximum Sessions	Specifies the maximum number of concurrent sessions that IIS allows. The default is 2147483647.
New ID On Secure Connection	Generates a new cookie when a transition from a non-secure to a secure connection is made. The default is **True**.
Time-out	Specifies how long a session object is maintained after the last request associated with the object is made. The default is 00:20:00.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Authentication

Use the **Authentication** feature page to configure the authentication methods that clients can use to gain access to your content.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Authentication
- Active Directory Client Certificate Authentication
- Anonymous Authentication
- Edit Anonymous Authentication Credentials Dialog Box
- Set Credentials Dialog Box (Anonymous)
- Basic Authentication
- Edit Basic Authentication Settings Dialog Box
- Digest Authentication
- Edit Digest Authentication Settings Dialog Box
- Forms Authentication
- Edit Forms Authentication Settings Dialog Box
- Windows Authentication
- Advanced Settings Dialog Box (Windows)
- ASP.NET Impersonation
- ASP.NET Impersonation Settings Dialog Box
- Set Credentials Dialog Box (Impersonation)

UI Elements for Authentication

The following tables describe the UI elements that are available on the **Authentication** feature page.

Element Name	Description
AD Client Certificate Authentication	AD Client Certificate authentication allows you to use Active Directory directory service features to map users to client certificates for

Element Name	Description
	authentication. Mapping users to client certificates lets you automatically authenticate users, without other authentication methods such as Basic, Digest, or Integrated Windows authentication. This kind of authentication is not available on Windows Vista Home Premium.
Anonymous Authentication	Anonymous authentication allows any user to access any public content without providing a user name and password. By default, Anonymous authentication is enabled in IIS 8. **Note** Use Anonymous authentication when you want all clients who visit your site to be able to view its content.
ASP.NET Impersonation	ASP.NET impersonation allows you to run ASP.NET applications under a context other than the default ASPNET account. Use impersonation with other IIS authentication methods or set up an arbitrary user account.
Basic Authentication	Basic authentication requires that users provide a valid user name and password to gain access to content. **Note** Basic authentication transmits passwords across the network with weak encryption. Use Basic authentication only when you know that the connection between the client and the server is secure.
Digest Authentication	Digest authentication uses a Windows domain controller to authenticate users who request access to content on your server. Consider using Digest authentication when you need improved security over Basic authentication. This kind of authentication is not available on Windows Vista Home Premium.

Element Name	Description
	Note Any browser that does not support the HTTP 1.1 protocol cannot support Digest authentication.
Forms Authentication	Forms authentication uses client-side redirection to forward unauthenticated users to an HTML form where they can enter their credentials, which are usually a user name and password. After the credentials are validated, users are redirected to the page they originally requested. **Important** Because Forms authentication sends the user name and password to the web server as plain text, use Secure Sockets Layer (SSL) encryption for the logon page and all other pages in your application.
Windows Authentication	Windows authentication uses NTLM or Kerberos protocols to authenticate clients. Windows authentication is best suited for an intranet environment. Windows authentication is not suited for use on the Internet because that environment does not require or encrypt user credentials. This kind of authentication is not available on Windows Vista Home Premium. **Important** The default setting for Windows authentication is Negotiate. This setting means that the client can select the appropriate security support provider.

Active Directory Client Certificate Authentication

Use **Active Directory (AD) Client Certificate** authentication to map client certificate information across many servers.

Note
Wildcard mapping with AD Client Certificate authentication is less efficient than the native IIS wildcard mapping.

If you select AD Client Certificate authentication, you cannot use IIS certificate mapping for any other sites hosted on the selected server.

Feature Page Elements

Element Name	Description
Active Directory Client Certificate Authentication	Select to manage **Active Directory Client Certificate** authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on Active Directory Client Certificate authentication.
Disable	Turns off Active Directory Client Certificate authentication.

Anonymous Authentication

Use Anonymous authentication to allow any user to access any public content without providing a user name and password challenge to the client browser. By default, Anonymous authentication is enabled in IIS 7.0.

If only selected users should view some content and you intend to use Anonymous authentication, configure the appropriate NTFS file system permissions to prevent anonymous users from accessing that content. If you want only registered users to view selected content, configure an authentication method for that content that requires a user name and password, for example, Basic or Digest authentication.

Feature Page Elements

Element Name	Description
Anonymous Authentication	Select to manage Anonymous authentication.

Actions Pane Elements

Element Name	Description
Disable	Turns off Anonymous authentication.
Enable	Turns on Anonymous authentication.
Edit	Displays the **Edit Anonymous Authentication Credentials** dialog box where you can set the security principle that anonymous users will use to connect to the site. You can perform this action only when you select **Anonymous Authentication** from the list on the feature page.

Edit Anonymous Authentication Credentials Dialog Box

Use the **Edit Anonymous Authentication Credentials** dialog box to change the account that IIS uses to access your sites and applications.

Element Name	Description
Specific user	Type the specific account name that you want IIS to use to access your site or application. By default, IIS 8 uses IUSR as the user name for anonymous access. This user name is created when you install IIS 8.
Set	Opens the **Set Credentials** dialog box, where you can specify an account name and password for IIS to use for Anonymous authentication.
Application pool identity	Select **Application pool identity** to enable IIS processes to run using the account that is currently specified on the property page for the application pool. By default, this is the Network Service account. **Important** If you use the Network Service account, you grant anonymous users all the internal network access

Element Name	Description
	associated with that account.

Set Credentials Dialog Box (Anonymous)

Use the **Set Credentials** dialog box to change the account that IIS uses for anonymous access to your sites and applications. By default, IIS 8 uses IUSR as the account for anonymous access. This account is created when you install IIS 8.

Element Name	Description
User name	Type the name of the account that you want IIS to use to access your sites and applications. **Important** Changing the anonymous user account from IUSR to another network account may affect security because you are giving anonymous users the administrative credentials associated with that account.
Password	Type the password associated with the account that you want IIS to use for anonymous access to your sites and applications. **Important** If the password associated with this account is changed, you must also change the password in the **Set Credentials** dialog box. If the password on the account and the password in the **Set Credentials** dialog box do not match, the associated site or application is disabled.
Confirm password	Retype the password associated with the account that you want IIS to use for anonymous access to your sites and applications.

Basic Authentication

Use Basic authentication to require that users provide a valid user name and password to access content. All major browsers support this authentication method and it works across firewalls and proxy servers. The disadvantage of Basic authentication is that it transmits passwords across the network by using weak encryption. Use Basic authentication only when you know that the connection between the client and the server is secure.

Disable Anonymous authentication if you use Basic authentication. The first request that all browsers send to a server is for anonymous access to server content. If you do not disable Anonymous authentication, users can anonymously access all the content on your server, including restricted content.

Feature Page Elements

Element Name	Description
Basic Authentication	Select to manage Basic authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on Basic authentication.
Disable	Turns off Basic authentication.
Edit	Displays the **Edit Basic Authentication Settings** dialog box where you can set the default domain and realm. You can perform this action only when you select **Basic Authentication** the list on the feature page.

Edit Basic Authentication Settings Dialog Box

Use the **Edit Basic Authentication Settings** dialog box to configure Basic authentication for a site.

Element Name	Description
Default domain	Type the name of a domain against which you want users to be authenticated by default. Any users who do not provide a domain name when

Element Name	Description
	they log on to your site are authenticated against this domain.
Realm	Type the DNS domain name or address that uses the credentials that have been authenticated against the default domain. Providing a **Realm** is optional for Basic authentication. **Important** If you enter the default domain name in the **Realm** box, your internal Microsoft Windows domain name may be exposed to external users during the user name and password challenge.

Digest Authentication

Use Digest authentication to offer authentication that is significantly more secure than Basic authentication. Digest authentication is also supported by all modern browsers, and works through proxy and firewall servers.

To use Digest authentication successfully, disable Anonymous authentication. The first request that all browsers send to a server is for anonymous access to server content. If you do not disable Anonymous authentication, users can anonymously access all the content on your server, including restricted content.

Feature Page Elements

Element Name	Description
Digest Authentication	Select to manage Digest authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on Digest authentication.
Disable	Turns off Digest authentication.
Edit	Opens the **Edit Digest Authentication**

Element Name	Description
	Settings dialog box where you can set the realm against which the credentials should be authenticated. You can perform this action only when you select **Digest Authentication** from the list on the feature page.

Edit Digest Authentication Settings Dialog Box

Use the **Edit Digest Authentication Settings** dialog box to define the realm name.

Element Name	Description
• **Realm**	Type the DNS domain name or address that uses the credentials that have been authenticated against your internal Windows domain. Providing a **Realm** is mandatory for Digest authentication.

Forms Authentication

Use Forms authentication to provide authentication for high-traffic sites or applications on public servers. This authentication mode lets you manage client registration and authentication at the application level, instead of relying on the authentication mechanisms provided by the operating system.

Feature Page Elements

Element Name	Description
Forms Authentication	Select to manage Forms authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on Forms authentication.
Disable	Turns off Forms authentication.
Edit	Opens the **Edit Forms Authentication**

Element Name	Description
	Settings dialog box where you can set various parameters for Forms authentication, such as cookie settings and the login URL that will redirect unauthenticated clients. You can perform this action only when you select **Forms Authentication** from the list on the feature page.

Edit Forms Authentication Settings Dialog Box

Use Forms authentication to manage client registration and authentication at the application level, instead of relying on the authentication mechanisms provided by the operating system.

Important

Because Forms authentication sends the user name and password to the server as plaintext, use Secure Sockets Layer (SSL) encryption for the logon page and for all other pages in your application except the home page.

Element Name	Description
Login URL	Specifies the URL to which the request is redirected for logon if no valid authentication cookie is found. The default value is login.aspx.
Authentication cookie time-out (in minutes)	Specifies the time, in integer minutes, after which the cookie expires. The default value is 30. If the **SlidingExpiration** attribute is true, the **time-out** attribute is a sliding value, expiring at the specified number of minutes after the time the last request was received. To prevent compromised performance, and to avoid multiple browser warnings for users who have cookie warnings turned on, the cookie is updated when more than half the specified time has elapsed.
Mode	Specifies where to store the Forms authentication ticket. The options are: • **Do not use cookies** - Cookies are not used. • **Use cookies** - Cookies are always used, regardless of device.

Element Name	Description
	• **Auto Detect** - Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET checks to determine whether cookies are enabled. • **Use device profile** - Cookies are used if the device profile supports cookies. Otherwise, no cookies are used. ASP.NET does not check to determine whether cookies are enabled on devices that support cookies. This setting is the default.
Name	Sets the name of the Forms authentication cookie. The default is .ASPXAUTH.
Protection mode	Specifies the type of encryption, if any, to use for cookies. The options are: • **Encryption and validation** - Specifies that both data validation and encryption are used to help protect the cookie. This option uses the configured data validation algorithm (based on the **<machineKey>** element). Triple-DES (3DES) is used for encryption, if available and if the key is long enough (48 bytes or more). **Encryption and validation** is the default, and recommended, value. • **None** - Specifies that both encryption and validation are disabled for sites that use cookies only for personalization and that have weaker security requirements. Microsoft does not recommend that you use this setting; however, it is the least resource-intensive way to enable personalization by using the .NET Framework. • **Encryption** - Specifies that the cookie is encrypted using Triple-DES or DES, but data validation is not performed on the cookie. Cookies used in this manner might be subject to plaintext attacks. • **Validation** - Specifies that a validation scheme verifies that the contents of an

Element Name	Description
	encrypted cookie have not been changed in transit. The cookie is created using cookie validation by concatenating a validation key with the cookie data, computing a message authentication code (MAC), and appending the MAC to the outgoing cookie.
Requires SSL	Specifies whether an SSL connection is required to transmit the authentication cookie. By default, this setting is disabled.
Extend cookie expiration on every request	Specifies whether sliding expiration is enabled. Sliding expiration resets an active authentication cookie's time to expire upon each request during a single session. By default, this setting is enabled.

Windows Authentication

Use Windows authentication only in an intranet environment. This authentication enables you to use authentication on your Windows domain to authenticate client connections.

Feature Page Elements

Element Name	Description
Windows Authentication	Select to manage Windows authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on Windows authentication.
Disable	Turns off Windows authentication.
Advanced Settings	Opens the **Advanced Settings** dialog box where you can enable or disable kernel mode authentication. You can perform this action only when you enable **Windows Authentication** from the list on the feature page.

Advanced Settings Dialog Box (Windows)

Use the **Advanced Settings** dialog box to specify whether Windows authentication is performed in kernel mode. By default, IIS enables kernel-mode authentication, which may improve authentication performance and prevent authentication problems with application pools configured to use a custom identity.

As a best practice, do not disable this setting if you use Kerberos authentication and have a custom identity on the application pool.

ASP.NET Impersonation

Use ASP.NET Impersonation authentication when you want to run your ASP.NET application under a security context different from the default.

If you enable impersonation for an ASP.NET application, that application can run in one of two different contexts: either as the user authenticated by IIS or as an arbitrary account that you set up. For example, if you were using Anonymous authentication and chose to run the ASP.NET application as the authenticated user, the application would run under an account setup for anonymous users (typically, IUSR). Likewise, if you chose to run the application under an arbitrary account, it would run under whatever security context was set up for that account.

Feature Page Elements

Element Name	Description
ASP.NET Impersonation	Select to manage ASP.NET Impersonation authentication.

Actions Pane Elements

Element Name	Description
Enable	Turns on ASP.NET Impersonation authentication.
Disable	Turns off ASP.NET Impersonation authentication.
Edit	Opens the **Edit ASP.NET Impersonation Settings** dialog box where you can specify the identity under which the ASP.NET application should run. You can perform this action only when you select **ASP.NET Impersonation** from the list on the feature page.

ASP.NET Impersonation Settings Dialog Box

Use the **ASP.NET Impersonation Settings** dialog box to specify the account your ASP.NET application should impersonate.

Element Name	Description
Specific user	Type the name of the specific account that you want your ASP.NET application to impersonate.
Set	Opens the **Set Credentials** dialog box, where you can specify an account name and password for IIS to use for ASP.NET impersonation.
Authenticated user	Select **Authenticated user** to enable the ASP.NET application to run under the security context of the user account authenticated by IIS.

Set Credentials Dialog Box (Impersonation)

Use the **Set Credentials** dialog box to specify the account that ASP.NET should impersonate.

Element Name	Description
User name	Type the name of the account that you want ASP.NET to impersonate.
Password	Type the password associated with the account that you want ASP.NET to impersonate.
Confirm password	Retype the password associated with the account that you want ASP.NET to impersonate.

Authorization Rules

Use the **Authorization Rules** page to manage the list of Allow or Deny rules that control access to content.

Sort the list by clicking one of the feature page column headings.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Authorization Rules
- Allow Authorization Rule Dialog Boxes

UI Elements for Authorization Rules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Mode	Displays the type of rule. Values can be either **Allow** or **Deny**. **Mode** values indicate whether the rule is designed to allow or deny access to content. **Note** If a role, user, or group has been specifically denied access by one rule, it cannot be granted access by another.
Users	Displays the user type, user names, or user groups to whom the rule applies.
Roles	Displays the Microsoft Windows role or roles to which the rule applies, for example, the Administrator role**.**
Verbs	Displays the HTTP verb or verbs to which this rule applies, for example, GET or POST.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Allow Rule	Opens the **Add Allow Authorization Rule** dialog box from which you can define rules for access to content.
Add Deny Rule	Opens the **Add Deny Authorization Rule** dialog box from which you can define rules for access to content.
Edit	Opens the **Edit Allow Authorization** or **Edit Deny Authorization** dialog box from which you can edit authorization rules. This action is available only when an item is selected from the list on the feature page.
Remove	Removes the item that is selected from the list on the feature page.

Allow Authorization Rule Dialog Boxes

Use the **Add Allow Authorization Rule**, the **Edit Allow Authorization Rule,** the **Add Deny Authorization Rule**, or the **Edit Deny Authorization Rule** dialog box to define rules for access to content.

Element Name	Description
All users	Select this option to manage access to content for both anonymous and authenticated users. By default, there is an allow rule configured for **All users**.
All anonymous users	Select this option to manage access to content for users that are not authenticated.
Specified roles or user groups	Select this option to manage access to content for specific Microsoft Windows roles or user groups.
Specified users	Select this option to manage access to content for a specific user account.
Apply this rule to specific verbs	Specifies that the rule applies to specific HTTP verbs such as GET or POST.

Centralized Certificates

Use the **Centralized Certificates** feature page to configure and manage a central certificate store for your server farm. Rather than replicating all SSL certificates across all servers in the web farm. The certificates are stored on a single server that hosts the central certificate store. This not only saves memory on all servers in the farm, but it changes the limit of SSL sites on a single server from 500 to more than 10,000.

Related scenarios

- Build a Web Farm with IIS Servers

In this document

- UI Elements for Centralized Certificates
- Edit Centralized Certificates Settings Dialog Box

UI Elements for Centralized Certificates

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
File name	Displays the file names of the certificates.
Name	Displays the names of certificates that have been issued to clients that are running on either Internet or intranet hosts.
Issued To	Displays the fully qualified domain names (FQDNs) of either the Internet or intranet hosts to which certificates have been issued.
Issued By	Displays the FQDNs of servers that have issued certificates to clients that are running on either Internet or intranet hosts.
Expiration Date	Displays the date that the certificate expires.
Certificate Hash	Displays binary data produced by using a hashing algorithm. Although this data uniquely identifies a certificate, the hash data cannot be used to trace a certificate because hashing is a one-way process.

Actions Pane Elements

Element Name	Description
Explore certificates location	Opens Windows Explorer in the directory where the certificates are stored on the Central Certificate Store server.
Edit Feature Settings	Opens the **Edit Centralized Certificates Settings** dialog box.

Edit Centralized Certificates Settings Dialog Box

Element Name	Description
Enable Centralized Certificates	Select the **Enable Centralized Certificates** check box if you want to create a central certificate store for your web farm. Otherwise, clear the check box.
Physical path	Type the physical path to the directory on the central certificate store server where you want the certificates stored.
User name	Type the name of the user account to use when accessing the central certificate store.
Password	Type the user account password.
Confirm password	Type the user account password again to confirm the password.
Private Key Password (Optional)	If you want to define a private-key password for accessing certificates, type the password, and then type it again to confirm it. You can define only one private-key password for the central certificate store.

CGI

Use the **CGI** properties page to configure Common Gateway Interface (CGI) applications on the Web server. You can set the CGI time-out value for a CGI script, isolate a CGI application in its own console, or configure a CGI application to run at the system level or at the user level.

Use the **Display** list to select one of the following options that specifies how the settings are displayed: **Friendly Names**, **Configuration Names**, or **Both Names**.

Related scenarios

- Build a Static Website on IIS
- Build a PHP Website on IIS

UI Elements for CGI

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Time-out (hh:mm:ss) [timeout]	Specifies the time-out (in seconds) for CGI applications. The default is 15 minutes (00:15:00).
Use New Console For Each Invocation [createCGIWithNewConsole]	Specifies whether a CGI application runs in its own console. If the value is set to **True**, each CGI application creates a new console when started. The default value is **False**.
Impersonate User [createProcessAsUser]	Specifies whether a CGI process is created in the system context or in the context of the requesting user. The default value is **True**.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Compression

Use the **Compression** feature page to provide faster transmission times between IIS and compression-enabled browsers. Compression is helpful if your site uses lots of bandwidth, or if you want to use bandwidth more effectively. Compression can improve performance if your network bandwidth is restricted as it is, for example, with mobile phone clients. Compression also helps improve performance in a data center environment.

IIS provides the following compression options:

- Static files only
- Dynamic application responses only
- Both static files and dynamic application responses

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

UI Elements for Compression

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Enable static content compression	Configures IIS to compress static content. **Note** Unlike dynamic responses, compressed static responses can be cached on disk across multiple requests without degrading CPU resources. On the next request, a compressed file can be retrieved from disk, which improves performance because the CPU does not have to compress the file again.
Enable dynamic content compression	Configures IIS to compress dynamic content.

Element Name	Description
	Note Compression of dynamic application responses can affect CPU resources because IIS does not cache compressed versions of dynamic output. If compression is enabled for dynamic responses and IIS receives a request for a resource that contains dynamic content, the response that IIS sends is newly compressed every time it is requested. Because dynamic compression consumes significant CPU time and memory resources, use it only on servers that have clients with slow network connections but that have CPU time to spare.
Only compress files larger than (in bytes)	Defines the minimum file size that you want IIS to compress. The default size is 256 bytes. **Note** This option is available at only the server level.
Cache directory	Defines the path of a local directory where a static file is cached after it is compressed, either until it expires or until the content changes. For security reasons, this temporary directory must be on a local drive on an NTFS-formatted partition. The directory cannot be compressed, and should not be shared. **Note** This box is available at only the server level.
Per application pool disk space limit (in MB)	Sets the maximum amount of space, in megabytes, you want IIS to use when compressing static content. When this setting is defined, IIS automatically empties the temporary directory when the set limit is reached. The default limit is 100 MB per application pool.

Element Name	Description
	Note This option is available at only the server level.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Configuration Editor

Configuration Editor lets you access and manage your configuration files by pulling sections into IIS Manager. The available sections are determined by where you are in your server, site, or application. Within each section, you can edit elements, attributes, and collections. This feature is available only to administrators.

Use the **Configuration Editor** page to manage sections of your configuration files. You can also use **Configuration Editor** to search through your configuration files and to generate scripts.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build a Web Farm with IIS Servers

In this document

- UI Elements for Configuration Editor
- Configuration Search Dialog Box
- Collection Editor Dialog Box

UI Elements for Configuration Editor

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Section	Use the **Section** drop-down list to select the section that you want to configure. If you want to type the section path, use the full path to access the data section, for example, **system.data** instead of **data**.
Core content section	The center pane displays the information for the section that you want to configure. The section path appears above the properties and their values. You can edit these property values by clicking the value.
Type and Validation Information	The text box that appears at the bottom of the Configuration Editor page displays the data type and validation information, such as the value range, for the currently selected property.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made to the element, attribute, or collection.
Cancel	Cancels the current changes that you have made to the element, attribute, or collection. Once you click **Apply**, you cannot use **Cancel** to revert your changes.
Generate Script	Creates a script of the actions that you have recently performed. Opens the **Script Dialog** window that displays the script for your action in three programming languages: C#, JavaScript, and AppCmd. The generate script functionality is only enabled after you take action, such as changing the value of a property. You must generate a script that includes your action before you click **Apply**. Note that the script will not include immediate actions, such as locking a section, editing a

Element Name	Description
	collection, or reverting to parent.
Search Configuration	Opens the **Configuration Search** dialog box that lets you search for specific sections in your configuration files. For more information, see the **Configuration Search Dialog Box** topic.
Lock Section	Locks the selected section so that additional changes cannot be made.
Unlock Section	Unlocks the selected section to let you make changes. Note that you can only unlock a section at the location where you originally locked it. For example, if you locked a section by using ApplicationHost.config, you cannot unlock it by using Web.config.
Lock Attribute	Locks the selected attribute.
Unlock Attribute	Unlocks the selected attribute to let you make changes. Note that you can only unlock an attribute at the location where you originally locked it.
Lock Element	Locks the selected element.
Unlock Element	Unlocks the selected element to let you make changes. Note that you can only unlock an element at the location where you originally locked it.
Remove	Removes the selected element and its settings.
Edit Items	Opens the **Collection Editor** dialog box for you to add, edit, or remove properties in a collection. For more information, see the **Collection Editor Dialog Box** topic.
Revert to Parent	Changes the values in the current section so that it inherits the values from the parent configuration. You can only use this functionality at the section level.

Configuration Search Dialog Box

Use the **Configuration Search** dialog box to search for specific sections in your configuration files.

Element Name	Description
Find Section	Type the path of the section that you want to search for in your configuration file. As you type each character, the search results filter to include only the sections that match.
Left pane	Displays the search results for your section search in either Hierarchy or Flat View.
Right pane	Displays the configuration value from the section that you searched for in your configuration file.

Collection Editor Dialog Box

Use the **Collection Editor** to add, edit, or remove properties in a collection. The Collection Editor contains two panes. The **Items** pane lists all items in the collection for the section that you are configuring. In the **Properties** pane, you can change the values of the property items in the collection by clicking the row that you want to change. The text box that appears at the bottom of the **Collection Editor** dialog box displays the data type and validation information, such as the value range, for the currently selected property.

Element Name	Description
Add	Adds a new row at the bottom of the list of items. The new row uses the schema default values.
Clear All	Removes all items from the collection. A dialog box asks you to confirm this action.
Lock Item	Locks the selected item in the collection.
Remove	Deletes the selected item from the configuration file section.

Connection Strings

Use the **Connection Strings** feature page to manage a list of database connection strings that are used by managed code applications.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for Connection Strings
- Add or Edit Connection String Dialog Box
- Set Credentials Dialog Box

UI Elements for Connection Strings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the connection string.
Connection String	Displays the connection string details.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Connection String** dialog box from which you can configure a SQL or custom connection string.
Edit	Opens the **Edit Connection String** dialog box from which you can edit the selected connection string.

Element Name	Description
Remove	Removes the item that is selected from the list on the feature page.

Add or Edit Connection String Dialog Box

Use the **Add Connection String** and **Edit Connection String** dialog boxes to create or edit either a SQL or custom connection string for use in an application that connects to a database.

Element Name	Description
Name	Type a name for the connection string, such as **StagingSqlServer**. The name that you enter should be the same name that you reference in your application code to retrieve data that uses this connection string.
Server	Type the name of the SQL server that hosts the database.
Database	Type the name of the SQL server database.
Credentials	Defines the security credentials that are used to connect to the database.
Use Windows Integrated Security	Select this option to configure the connection string so that the application uses the current Windows identity established on the operating system thread to access the SQL Server database. Integrated security requires that SQL Server run on the same computer as IIS, unless the administrator has configured delegation between computers. Integrated security also requires that all application users have to be on the same domain so that their credentials are available to the Web server.
Specify credentials	Select this option to configure the connection string to use a SQL Server user name and password. This setting is helpful when the SQL Server database is on a computer other than the Web server.
Set	Opens the **Set Credentials** dialog box from which you can configure a user name and

Element Name	Description
	password to connect to the SQL Server database.
Custom	Type a custom database connection string. Use this setting instead of a SQL connection string when you want to connect to a different data store or configure additional connection string settings.

Set Credentials Dialog Box

Use the **Set Credentials** dialog box to configure a connection string to use a SQL Server user name and password. Set these credentials when the SQL Server database is on a computer other than the Web server.

Element Name	Description
Name	Type the name of the account that you want IIS to use to access your SQL Server database.
Password	Type the password associated with your SQL Server database.
Confirm password	Retype the password associated with your SQL Server database.

Default Document

Use the **Default Document** feature page to configure the list of default documents. If a user accesses your site or application without specifying a document name (for example, by requesting http://www.contoso.com/ instead of http://www.contoso.com/Default.htm), you can configure IIS to serve a default document such as Default.htm. IIS returns the first default document in the list that matches a file name in the directory.

Note

To improve performance, make sure that the first default document in the list exists on disk and that it is the first file name in the list.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS

- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

1. UI Elements Default Documents
2. Add Default Document Dialog Box

UI Elements Default Documents

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the file name that is added as a default document.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Default Document** dialog box in which you can add a file name to the default document list.
Remove	Removes the item that is selected from the list on the feature page.
Move Up	Moves up a selected item in the list. **Note** When items in the list are reordered at a child level, the child will no longer inherit settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you

Element Name	Description
	want to inherit settings from a parent level, revert all changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. **Note** When items in the list are reordered at a child level, the child will no longer inherit settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Disable	Disables the **Default Documents** feature. When default documents are disabled, and a user accesses your site or application without specifying a document name, client browsers receive a 403—Forbidden error. This error is returned because the web server cannot determine which file to serve, and the client cannot view the contents of the directory. **Note** If the **Default Documents** feature is disabled but the **Directory Browsing** feature is enabled, client browsers receive a directory listing instead of a 403—Forbidden error.
Enable	Enables the **Default Documents** feature.
Revert to Inherited	Reverts the feature to inherit settings from the parent configuration. This action deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.

Add Default Document Dialog Box

Use the **Add Default Document** feature page to add a file name to the list of default documents.

Element Name	Description
Name	Type the file name or relative path that you want to add to the list of default documents. This value is added to the beginning of this list.

Directory Browsing

Use the **Directory Browsing** feature page to modify the content settings for browsing a directory on the Web server. When you configure directory browsing, all subdirectories use the same settings unless you override them at a lower level.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS

In this document

UI Elements for Directory Browsing

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Time	Select this option to display the last modified time for each file in a directory listing.
Size	Select this option to display the size of each file in a directory listing.
Extension	Select this option to display the file name extension of each file in a directory listing.
Date	Select this option to display the last modified

Element Name	Description
	date for each file in a directory listing.
Long date	Select this option to display, in extended format, the last modified date for each file in a directory listing.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Enable	Enables the **Directory Browsing** feature. When the **Default Document** feature is disabled but the **Directory Browsing** feature is enabled, client browsers receive a directory listing instead of a 403—Forbidden error. When both the **Default Document** feature and the **Directory Browsing** feature are enabled, IIS returns a default document if found, or a directory listing if no default document is found.
Disable	Disables the **Directory Browsing** feature. When the **Default Document** feature and the **Directory Browsing** feature are disabled, client browsers receive a 403—Forbidden error.

Error Pages

Use the **Error Pages** feature page to manage a list of custom HTTP error messages.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS

- Build a PHP Website on IIS

In this document

- UI Elements for Error Pages
- Add or Edit Custom Error Page Dialog Box
- Edit Error Pages Settings Dialog Box
- Set Localized Custom Error Path Dialog Box

UI Elements for Error Pages

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Status Code	Displays the HTTP status code to which the error page applies.
Path	Displays the path of the error page if you choose the **File** path type, or the URL of the custom error page if you use either the **Execute URL** or the **Redirect** path type.
Type	Displays the path type (file, execute URL, or redirect URL) of the error page.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Custom Error Page** dialog box in which you can add a custom error to the list of HTTP errors.
Edit	Opens the **Edit Custom Error Page** dialog box in which you can edit the selected custom error.
Change Status Code	Enables the **Status Code** field of the selected

Element Name	Description
	list item so you can change the status code to which the error page applies.
Remove	Removes the item that is selected from the list on the feature page.
Edit Feature Settings	Opens the **Edit Error Pages Settings** dialog box in which you can configure settings that apply to the entire **Error Pages** feature.

Add or Edit Custom Error Page Dialog Box

Use the **Add Custom Error Page** and **Edit Custom Error Page** dialog boxes to configure IIS to return a custom error message whenever a specific HTTP error occurs.

Element Name	Description
Status code	Type the HTTP status code (and optional sub-status code) to which the custom error page applies.
Insert content from static file into the error response	Select this option when the error page should include static content, such as from an HTML file. In the corresponding **File path** box, type or browse for the physical path of the static file. **Caution** If the file that you configure in the **File path** box is not static content, such as a .aspx file, the code will be displayed instead of running when the response is returned.
Browse	Opens the file menu from which you can browse for the static file.
Try to return the error file in the client language	Select this option when the error should be returned in the client language (if possible). When you select this check box, click the **Set** button to configure the path.
Set	Opens the **Set Localized Custom Error Path** dialog box from which you can configure the path of localized error files.

Element Name	Description
	Note This button is enabled only when you select **Try to return the error file in the client language**.
Execute a URL on this site	Select this option when the error page should include dynamic content, such as from an ASP file. In the corresponding **URL (relative to site root)** box, type the relative path to the custom error.
Respond with a 302 redirect	Select this option when you want the client to be redirected to a different URL. In the corresponding **Absolute URL** box, type the full URL to which the client is redirected.

Edit Error Pages Settings Dialog Box

Use the **Edit Error Pages Settings** dialog box to return a default error message any time an HTTP error occurs for which there is no predefined error page.

Element Name	Description
Custom error pages	Under **Error Responses**, select this option to respond with only custom error pages for both local requests and for requests that are made from other computers.
Detailed errors	Under **Error Responses**, select this option to respond only with detailed errors for both local requests and for requests that are made from other computers. **Caution** Detailed error messages might display sensitive server information to the client.
Detailed errors for local requests and custom error pages for remote requests	Under **Error Responses**, select this option to display detailed errors for local requests and custom errors for requests that are made from other computers.

Element Name	Description
Path	Under **Default Page**, type or browse for the path of the default error page. Depending on the option you select from the **Path type** drop-down list, the value that you type in the **Path** field should be one of the following: • If you select **File**, specify a physical path. • If you select **Execute URL**, specify a URL that is relative to the site root URL. • If you select **Redirect**, type a URL.
Path type	Under **Default Page**, select one of the following options from the **Path type** drop-down list: • **File**: The default error is static content, such as from an HTML file. • **Execute URL**: The default error is dynamic content, such as from an ASP file. • **Redirect**: The default error is at a redirection URL.

Set Localized Custom Error Path Dialog Box

Use the **Set Localized Custom Error Path** dialog box to configure the path of localized error files.

Element Name	Description
Root directory path	Type the physical path of the root directory of the localized error files. The root directory path is the part of the path before the locale directory, such as EN-US. For example, if the full path of the error files for the EN-US locale is d:\errors\EN-US\messages\404.htm, you type **d:\errors**.
Browse	Opens the file menu from which you can browse for the path.
Relative file path	Type the relative physical path of the localized error file. The relative file path is the part of the path after the locale directory, such as EN-US. For example, if the full path of the error files for the EN-US locale is d:\errors\EN-

Element Name	Description
	US\messages\404.htm, you type **\messages\404.htm**.

Failed Request Tracing Rules

Use the **Failed Request Tracing Rules** feature page to manage a list of tracing rules for failed requests.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Failed Request Tracing Rules
- Specify Content to Trace Wizard
- Define Trace Conditions Wizard
- Select Trace Providers Wizard

UI Elements for Failed Request Tracing Rules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Path	Displays the URL that was identified by a tracing rule for failed requests. The URL can contain, at most, one wildcard and must be local to the directory where the failed request definition is set. For example, the wildcard patterns, such as *.aspx and ind* match the URL www.contoso.com/index.aspx.
Associated Providers	Displays the trace providers associated with a

Element Name	Description
	tracing rule for failed requests. A provider defines the functional area for tracing. One example is the ASP.NET trace provider, which traces transitions into and out of managed code, including *.aspx requests.
Status Codes	Displays the HTTP status codes and sub-status code combinations identified by a tracing rule for failed requests. There can be multiple status codes in this list, with commas that separate each code.
Time Taken	Displays the maximum time, in seconds, that the request should take.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Failed Request Tracing Rule** wizard in which you can add a rule to trace for failed requests.
Edit	Opens the **Edit Failed Request Tracing Rule** wizard in which you can edit the selected failed request tracing rule.
Remove	Removes the item that is selected from the list on the feature page.
Edit Site Tracing	Opens the **Edit Web Site Failed Request Tracing Settings** dialog box in which you enable and configure failed request tracing settings for the website. This action is available only when you are connected to a server and you have selected a specific site.
View Trace Logs	Opens the trace log file directory. This action is available only when you are connected to a

Element Name	Description
	server.
Revert to Inherited	Resets the feature to inherit settings from the parent configuration. This action deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other actions in the **Actions** pane do not appear until you select the unordered list format.
Move Up	Moves up a selected item in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the

Element Name	Description
	Actions pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.

Specify Content to Trace Wizard

Use the **Specify Content to Trace** page of the **Add Failed Request Tracing Rule** Wizard to configure the content you want to trace with a tracing rule for failed requests.

Element Name	Description
All content (*)	Select this option to trace all files in a directory.
ASP.NET (*.aspx)	Select this option to trace all .aspx files in a directory.
ASP (*.asp)	Select this option to trace all .asp files in a directory.
Custom	Select this option and type a value, such as "xyz.exe" or "*.jpg", to trace a failure for a custom set of content. This value can contain, at most, one wildcard and must be local to the directory where the failed request definition is set. For example, the wildcard patterns, such as *.aspx and ind* match the URL www.contoso.com/index.aspx.

Define Trace Conditions Wizard

Use the **Define Trace Conditions** page of the **Add Failed Request Tracing Rule** Wizard or **Edit Failed Request Tracing Rule** Wizard to define a trace condition for a failed request tracing rule. This trace condition allows you to capture trace events for a failed request and log these trace events, as they occur, without having to reproduce the error.

If more than one trace condition is specified, the first condition that is met generates the failed request trace log file. For example, if you configure a trace rule to log requests that generate status code 404 and requests that take 4 seconds or longer, the trace log generates an event when the first of the following conditions is met:

- A request generates status code 404.

- A request takes longer than 4 seconds.

Element Name	Description
Status code(s)	Select this option when you want a trace event to be logged for responses that generate the specified HTTP status code (and optionally a sub-status code). In the corresponding box, type the status code or codes you want to trace. Separate multiple status codes by using commas (,) between each code, or type a status code range by using a hyphen (-) between two codes. You can also refine your status codes by using sub-status codes, such as 404.2 or 404.3.
Time taken (in seconds)	Select this option when you want to trace events that exceed the time interval allocated for the request to complete processing. In the corresponding box, type the maximum time, in seconds, that a request should take.
Event severity	Select this option when you want to trace events based on severity. In the corresponding list, select one of the following event severities on which to trace: • **Error**: Provides information about components that experience an error and cannot continue to process requests. These errors usually indicate a server-side problem. • **Critical Error**: Provides information about actions that can cause a process to exit. • **Warning**: Provides information about components that experience an error but that can continue to process the request.

Select Trace Providers Wizard

Use the **Select Trace Providers** page of the **Add Failed Request Tracing Rules** Wizard or **Edit Failed Request Tracing Rules** Wizard to define what to trace for a URL. This includes the trace providers, verbosity level, and areas for the provider to trace.

Element Name	Description
Providers	Defines the functional area for tracing. You can select one or more of the following trace providers: • **ASP** - When you want to trace the start and completion of the execution of an ASP request. • **ASP.NET** - When you want to see transitions into and out of managed code. This setting includes requests for .aspx files and any requests that are processed by managed modules, such as static files that are configured to use forms-based authentication. • **ISAPI Extension** - When you want to trace the transition of a request into and out of an ISAPI extension process. • **WWW Server** - When you want to trace requests through the IIS worker process.
Verbosity	Specifies how much data each provider reports to the trace log. Verbosity levels are inclusive. For example, if you select **Errors**, then errors, critical errors, and general events will be included in the log file. You can select one or more of the following verbosity levels: • **General** - Provides context information for the request activity, for example, IIS 8, IISGeneralEvents::GENERAL_REQUEST_START logs the URL and the verb for the request. • **Critical Errors** - Provides information about actions that can cause a process to end or that are about to cause a process to end. • **Errors** - Provides information about components that try to run on the server, encounter an error, and cannot continue to process requests. These errors are conditions that generally result in the termination of processing a request. An example would be an authorization event. • **Warnings** - Provides information about components that try to run on the server, encounter an error, but can continue to process requests.

Element Name	Description
	• **Information** - Provides general information about requests. • **Verbose** - Provides detailed information about requests.
Areas	Defines the areas for the provider to trace. If the provider selected is the ASP.NET trace provider, select one or more of the following areas to trace: • **Infrastructure** - When you want to trace events that are primarily related to entering and leaving various parts of the ASP.NET infrastructure. • **Module** - When you want to trace events that are logged when a request enters and leaves various HTTP pipeline modules. This area is also used to capture trace events for managed modules. • **Page** - When you want to generate trace events that correspond to the execution of specific ASP.NET page-related events, such as **Page_Load**. This area is also used to capture ASP.net Page **Trace.Write** and **Trace.Warn** events. • **AppServices** - When you want to trace those events that are logged as part of the new application services functionality. If the provider selected is the WWW Server trace provider, select one or more of the following areas to trace: • **Authentication** - When you want to trace authentication attempts, including the name of the authenticated user, the authentication scheme (Anonymous, Basic, or others), and the results of the authentication attempt (successful, failed, error, and so forth). • **Security** - When you want to generate trace events where requests are rejected by the server running IIS. For example, if a client request was denied access to a resource. • **Filter** - When you want to determine how long it takes an ISAPI filter to process requests.

Element Name	Description
	• **StaticFile** - When you want to trace how long it takes requests for static files to be completed or to see how filters might be changing the requests. • **CGI** - When a request is made for a CGI file and you want to generate trace events. • **Compression** - When a response is compressed and you want to generate trace events. • **Cache** - When you want to generate trace events for cache operations associated with the request. • **RequestNotifications** - When you want to capture all request notifications, both on entrance and on exit. • **Module** - When you want to trace events that are logged when a request enters and leaves various HTTP pipeline modules. This area is also used to capture trace events for managed modules.

FastCGI Settings

Web servers use the FastCGI (Fast Common Gateway Interface) protocol to communicate with external applications. The FastCGI protocol defines how information is passed between the web server and an external FastCGI application. When a request is sent, the web server starts a FastCGI application to process the request, and generate dynamic responses to send back to the client.

FastCGI applications handle HTTP requests and each one consists of a collection of process pools. Within the pools, each process handles requests one at a time and, upon completion, returns to the pool until it receives another request.

Use the **FastCGI Settings** page to configure process pool settings for the Fast CGI applications on your Web server. You can add FastCGI applications for use by sites on your web server, and specify configuration properties such as the processing variables and maximum requests for these applications.

Related scenarios

- Build a PHP Website on IIS

In this document

- UI Elements for FastCGI Settings

- Add or Edit FastCGI Application Dialog Box
- Environment Variables Collection Editor

UI Elements for FastCGI Settings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Full Path	Lists the physical path of the process executable to use in the process pool.
Arguments	Lists the arguments to pass to the selected FastCGI application at start time. If several FastCGI process pools exist for a single process executable, these arguments can be used to identify a FastCGI process pool.
Max. Instances	Lists the maximum number of FastCGI processes to allow in the application process pool for the selected FastCGI application. This number also represents the maximum number of concurrent requests that the FastCGI application can handle. The default value is **4**. For the Windows® 7 and Windows Server® 2008 R2 operating systems only, setting the **MaxInstances** property to 0 lets IIS automatically select the optimal number of FastCGI processes for the current execution environment.
Instance Max. Requests	Lists the maximum number of requests that the selected FastCGI application processes before recycling. The default value is **200**.

Actions Pane Elements

Element Name	Description
Add Application	Opens the **Add FastCGI Application** dialog

Element Name	Description
	box that lets you configure general and processing properties.
Edit	Opens the **Edit FastCGI Application** dialog box that lets you change the general and processing properties for the selected application.
Remove	Removes the selected FastCGI application.

Add or Edit FastCGI Application Dialog Box

Use the **Add FastCGI Application** and **Edit FastCGI Application** dialog boxes to add or edit general and processing properties for an application.

For more information about FastCGI applications, see the **FastCGI Settings Page**.

Element Name	Description
Full Path	Specifies the physical path of the process executable to use in the FastCGI application pool.
…	Click to find the physical location of your FastCGI process executable file.
Arguments	Specifies the command-line arguments that are passed to the FastCGI process. These arguments can be used to uniquely identify a FastCGI process pool in cases where several FastCGI process pools exist for a single process executable.
Environment Variables	Specifies the collection of environment variables that are set for the process executable program. This property is optional. Click the Browse (**…**) button to open the **EnvironmentVariables Collection Editor**, from which you can add and remove members and define property values. For more information, see the **Environment Variables Editor Dialog Box** UI Help page.
Instance MaxRequests	Controls the FastCGI process-recycling behavior. Specifies the maximum number of

Element Name	Description
	requests that a FastCGI application is allowed to handle before the process is recycled. The default value is **200**.
Max Instances	Specifies the maximum number of FastCGI processes to allow in the application process pool for the selected FastCGI application. This number also represents the maximum number of concurrent requests that the FastCGI application can handle. The default value is **4**. For the Windows® 7 and Windows Server® 2008 R2 operating systems only, setting the MaxInstances property to 0 lets IIS automatically select the optimal number of FastCGI processes for the current execution environment.
Monitor changes to file	This property applies to the Windows 7 and Windows Server® 2008 R2 operating systems only. Specifies the path of a file. Changes to this file trigger recycling of FastCGI processes. You can specify the absolute path and also the path that is relative to the location of the FastCGI executable.
Standard error mode	This property applies to the Windows 7 and Windows Server® 2008 R2 operating systems only. Specifies how to handle errors that are reported by a FastCGI process. The following options exist: • **ReturnStderrIn500** — This module sets the status code to 500 and sends what is received on STDERR as a response. This is the default value and has the same behavior as IIS 8. • **ReturnGeneric500** — This module sets the status code to 500, but will return a generic 500. This configuration is useful if you want to enable detailed error logging for a server but do not want to return these errors to users.

Element Name	Description
	• **IgnoreAndReturn200** — This module sends what is received on STDOUT as a response with the status code 200. The text on the STDERR stream is ignored. This property is useful if you want to use debug statements for tracking purposes. • **TerminateProcess** — This module terminates the FastCGI process and returns the generic 500 error message.
ActivityTimeout	Specifies the maximum length of time that a FastCGI process for the application is allowed to run without communicating with IIS before the process times out. This timeout can be used to detect and shut down hung processes. The default value is **30** (seconds).
FlushNamedpipe	Specifies whether the named pipe is flushed at the end of each request and before shutting down the application. If the named pipe is flushed, the value is set to **true**. The default value is **false**. This property applies only when the named-pipe protocol is used.
Protocol	Specifies the protocol to use to communicate with the FastCGI process. The possible values for the **Protocol** property are **0 (NamedPipe)** or **1 (TCP)**. The default value is **0 (NamedPipe)**.
Idle Timeout	Specifies the time interval that a FastCGI process for the application can remain idle, after which an idle process ends. The default value is **300** (seconds).
Queue Length	Specifies the maximum number of requests that are queued for the FastCGI application pool. When the queue is full, subsequent requests return the HTTP error code 503 (Service Unavailable) to clients. This error code indicates that the application is too busy. The default value is **1000**.
Rapid Fails PerMinute	Specifies the number of FastCGI process failures allowed in a minute before the FastCGI

Element Name	Description
	handler takes the process offline. The default value is **10**.
Request Timeout	Specifies the maximum time allowed for a request to the application. If a FastCGI process takes longer than the specified time on a single request, it is terminated. The default value is **90** seconds.

Environment Variables Collection Editor

Environment variables are strings that contain information such as file name or path. These variables control the behavior of various programs on a computer.

When you add or edit a FastCGI application, you can define values for the application properties. For example, you can define values for environment variables. These environment variables are set by the FastCGI module during the start time of a FastCGI process. The environment variables are set for the FastCGI process only and are not seen by any other processes that are running on that computer. Click the Browse (**…**) button when you set the **EnvironmentVariables** property to open the **EnvironmentVariables Collection Editor** dialog box. This lets you specify optional environment variables that are set in the FastCGI executable. If any of the supplied variable names collide with existing server variables, the ones specified in this property are used.

Element Name	Description
Members	Contains a list of the members in the environment variables collection. Each member consists of a name-value pair. To add a member, type a name and value in the **Name** **Properties** pane, and then click **Add**. To remove a member, highlight the member in the list, and then click **Remove**. To change the order in which these members are stored in the ApplicationHost.config file, use the **Up** and **Down** arrows. Because all name-value pairs are unique, the order will not affect functionality.
Name	Type a unique name for the member.
Value	Type a value for the member. This information is stored with the corresponding name in the ApplicationHost.config file.
Sort icons	Use these icons to sort the name-value

Element Name	Description
	collection pairs by category, property pages, or alphabetically.

Feature Delegation

Use the **Feature Delegation** feature page to configure the delegation state of IIS Manager features for sites and applications on your Web server.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- Feature Delegation Overview
- UI Elements for Feature Delegation
- Custom Site or Application Delegation
- Copy Delegation Dialog Box

Feature Delegation Overview

When you configure the delegation state of a feature from IIS Manager, you specify whether the feature's related configuration section is locked or unlocked in the server-level configuration files (ApplicationHost.config and root Web.config) for IIS 8. When you lock a feature, configuration can only be read from and written to the server-level configuration file for that feature. However, you can unlock a feature when you want to read from or write to lower-level configuration files, such as a Web.config file in your site or application.

Note

When you configure delegation state for features in IIS Manager, you control only the configuration sections and the corresponding levels at which settings can be set in the configuration files.

You can use Web.config files to read and save configuration settings for any or all features in your sites and applications. You can then easily copy configuration from one computer to another. For example, you might develop an application on a development or test computer and specify application features to write to Web.config. This lets you package all the application's content easily, including its configuration file, and transfer it to another computer.

Important

If you develop a site or an application for a production environment, ask someone who is an administrator on the production computer whether the features that you configure in the site's or application's Web.config can be unlocked in the production environment. If the feature you configure in Web.config is locked in the production environment, it will cause run-time errors because your chosen configuration is invalid.

Feature Delegation pairs with the **IIS Manager Users** and **IIS Manager Permissions** features to enable non-administrative users to view and configure delegated features in their sites or applications by using IIS Manager. The delegation state of a feature determines whether users can configure the feature in their sites and applications. For example, if you want to let users configure data connection strings, select **Connection Strings** from the list on the **Feature Delegation** page, and then click **Read/Write** from the **Actions** pane or the right-click menu.

The **Feature Delegation** page affects all sites and applications on the Web server. The **Feature Delegation** settings that you configure at the server level are for all sites on the server; the settings that you configure at the site level are for all applications in that site. If you want to configure delegation states for features in a specific site or application, use the **Custom Site Delegation** and **Custom Application Delegation** pages.

Caution

If you have configured delegation states for features and you want to change those states later, review how the features have been delegated at lower levels in the configuration. Make sure that you understand how the changes affect those features at the lower levels. For example, suppose that you configure a feature to be Read Only for sites, and a user at the site level configured a more restrictive setting by removing delegation for that feature. If you reconfigure the feature to be Read/Write at the parent level, you might inadvertently change the state that the user set. As a best practice, use a test computer to test how changes to the delegation states affect your environment before you change the production environment.

UI Elements for Feature Delegation

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the UI feature in IIS Manager. Select a feature and then click the delegation state that you want from the **Actions** pane or from the right-click menu.
Delegation	Displays the delegation state of the UI feature

Element Name	Description
	at lower levels.

Actions Pane Elements

The following table describes the delegation options available for IIS 8 features in IIS Manager and how the delegation options affect the configuration files and user interface (UI). These features are configuration-based features. If you extend IIS Manager to include third-party features, your delegation states might differ from this table if they are user-specified delegation states. Additionally, individual features that you configure on the **Feature Delegation** page might be protocol-specific and not applicable to all sites and applications.

Delegation Option	Description for Windows Vista operating systems	Description for Windows Server 2008 operating systems
Read/Write	When you select **Read/Write** for a feature, you unlock the feature's related configuration sections in the server-level configuration file. Configuration changes for that feature will then be read from and written to Web.config files in sites or applications.	When you select **Read/Write** for a feature, you unlock the feature's related configuration sections in the server-level configuration file. Configuration changes for that feature will then be read from and written to Web.config files in sites or applications. Additionally, configuring a feature to be Read/Write enables nonadministrative users to see and configure the feature in IIS Manager for sites or applications to which they are allowed to connect.
Read Only	When you select **Read Only** for a feature, you lock the feature's related configuration sections in the server-level configuration file. Configuration cannot be written to Web.config files in sites or applications. **Note** If you set configuration settings for a feature in	When you select **Read Only** for a feature, you lock the feature's related configuration sections in the server-level configuration file. Configuration cannot be written to Web.config files in sites or applications. Additionally, nonadministrative users cannot configure the feature in IIS Manager for their

Delegation Option	Description for Windows Vista operating systems	Description for Windows Server 2008 operating systems
	a Web.config file, but the feature is Read Only, you will receive an error if you try to configure the feature at the site or application level, and you will receive runtime errors in your sites and applications.	sites or applications, but they can see the feature as Read Only in IIS Manager to see how the feature is configured in the server-level configuration file. **Note** If you set configuration settings for a feature in a Web.config file, but the feature is Read Only, you will receive an error if you try to configure the feature at the site or application level, and you will receive run-time errors in your sites and applications.
Not Delegated	When you select **Not Delegated** for a feature, you lock the feature's related configuration sections in the server-level configuration file. Configuration cannot be read from and written to Web.config files in sites or applications.	When you select **Not Delegated** for a feature, you lock the feature's related configuration sections in the server-level configuration file. Configuration cannot be read from and written to Web.config files in sites or applications. Additionally, nonadministrative users will not see the feature in IIS Manager and they will not be able to configure the feature at the site and application levels.
Reset to Inherited	When you select **Reset to Inherited** for a feature, the feature inherits the delegation state that is set at the parent level.	When you select **Reset to Inherited** for a feature, the feature inherits the delegation state that is set at the parent level.
Configuration Read/Write	When you select **Configuration Read/Write** for a feature, you	When you select **Configuration Read/Write** for a feature, you

Delegation Option	Description for Windows Vista operating systems	Description for Windows Server 2008 operating systems
	unlock the feature's configuration sections in the server-level configuration file. Configuration changes for that feature will then be read from and written to Web.config files in sites or applications. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.	unlock the feature's configuration sections in the server-level configuration file. Configuration changes for that feature will then be read from and written to Web.config files in sites or applications. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.
Configuration Read Only	When you select **Configuration Read Only** for a feature, you lock the feature's configuration sections in the server-level configuration file. Configuration cannot be written to Web.config files in sites or applications; however, the application can continue to write to the database. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.	When you select **Configuration Read Only** for a feature, you lock the feature's configuration sections in the server-level configuration file. Configuration cannot be written to Web.config files in sites or applications; however, the application can continue to write to the database. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.
Reset All Delegation	Resets the delegation states of all the features to states set at the parent level. At the server	Resets the delegation states of all the features to states set at the parent level. At the server

Delegation Option	Description for Windows Vista operating systems	Description for Windows Server 2008 operating systems
	level, this sets the delegation states to the states as specified in the **overrideModeDefault** settings in the server-level configuration file.	level, this sets the delegation states to the states as specified in the **overrideModeDefault** settings in the server-level configuration file.
Custom Site Delegation or **Custom Application Delegation**	Not available.	Opens the Custom Site Delegation feature page or the Custom Application Delegation feature page in which you can configure custom delegation states for an individual site or application.

Custom Site or Application Delegation

Use the **Custom Site Delegation** and **Custom Application Delegation** feature pages to configure the default delegation state of IIS Manager features for a specific site or application on your web server. The delegation state of a feature determines whether nonadministrative users who have been granted permission to a site or application can configure the feature in that site or application. For example, you can configure the delegation state of a feature, such as **Connection Strings**, as Read/Write when you want to allow users to configure that feature in their site or application.

The features that you configure on this page affect only the site or application that you select from the **Sites** or **Applications** drop-down lists. You can use **Copy Delegation** to open the dialog box and select additional sites or applications, or both, if you want to apply the same delegation states of features to other sites and applications. You can also configure the default delegation state of features for all sites and applications on the Web server by using the **Feature Delegation** page.

Caution

If you have configured delegation states for features and you want to change those states later, review how the features have been delegated at lower levels in your configuration to make sure that you understand how the changes affect those features at the lower levels. For example, if you originally configured a feature to be Read Only for sites, and a user at the site level configured a more restrictive setting by removing delegation for that feature, you might inadvertently change the state that was set by the user if you configure the feature to be Read/Write at the parent level. As a best practice, use a test computer to test how changes to the delegation states affect your environment before you make changes to a production environment.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Feature Page Elements

Element Name	Description
Sites or **Applications**	Select the site or application for which you want to configure the delegation state of UI features. **Note** To configure custom delegation for a site, you must be connected to a server in IIS Manager. To configure custom delegation for an application, you must be connected to a site in IIS Manager.
Copy Delegation	Click **Copy Delegation** to open the **Copy Delegation** dialog box and copy the custom delegation states to other sites or applications.
Name	Displays the name of the UI feature in IIS Manager. Select a feature and then click the delegation state that you want from the **Actions** pane or from the right-click menu.
Delegation	Displays the delegation state of the UI feature at lower levels.

Actions Pane Elements

The following table describes the delegation options available for IIS 8 features in IIS Manager and how the delegation options affect the configuration files and user interface (UI). These features are configuration-based features. If you extend IIS Manager to include third-party features, your delegation states might differ from this table if they are user-specified delegation states.

Element Name	Description
Read/Write	Enables nonadministrative users for a site or application to see and configure the selected feature at the site and application levels by using IIS Manager.
Read Only	Enables nonadministrative users for a site or

Element Name	Description
	application to see, but not configure, the selected feature at the site and application levels by using IIS Manager.
Configuration Read/Write	Enables nonadministrative users for a site or application to see and configure the selected feature at the site and application levels by using IIS Manager. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.
Configuration Read Only	Enables nonadministrative users for a site or application to see, but not configure, the selected feature at the site and application levels by using IIS Manager. This option is available only for features, such as **.NET Users** and **.NET Roles**, that have configuration in both a configuration file and a database. **Note** This setting does not affect the database permissions.
Not Delegated	Removes the selected feature from being delegated. This feature is not displayed in IIS Manager and cannot be configured by nonadministrative users at the site and application levels.
Reset to Inherited	Resets the selected feature to the delegation state specified at the parent level.
Reset All Delegation	Resets the delegation states of all the features to the delegation states specified at the parent level.
Default Delegation	Opens the **Feature Delegation** page in which you can configure default delegation states for sites or applications.

Copy Delegation Dialog Box

Use the **Copy Delegation** dialog box to select sites and applications to which you want to copy custom delegation states for features in IIS Manager. You can copy custom delegation states from one site to another when you are connected to a server in IIS Manager, and copy custom delegation states from one application to another when you are connected to a site in IIS Manager.

Element Name	Description
Site_name or *application_name*	Select the check box next to the sites or applications to which you want to apply the custom delegation state that you have configured on the **Custom Site Delegation** or **Custom Application Delegation** page.

FTP Authentication

Use the **FTP Authentication** feature page to configure the authentication methods that FTP clients can use to gain access to your content. You can sort this list by name, status, or type by clicking the appropriate column heading. By using the Group by drop-down list, you can also group authentication features by type or status.

By default, no authentication methods are enabled. Enable an authentication method if you want to allow FTP users to access your content. There are two types of authentication methods: Built-In and Custom.

- Built-In authentication methods are integral parts of the FTP server. These authentication methods can be enabled or disabled, but cannot be removed from the FTP server.
- Custom authentication methods are implemented through an installable component. These authentication methods can be enabled or disabled, and they can be added to or removed from the FTP server.

Note

Clicking **Custom Providers** in the task list displays the FTP Custom Providers Dialog Box dialog box.

Note

Basic authentication works with Active Directory (AD) user isolation. However, if you enable Custom authentication or any other form of authentication when AD user isolation is already enabled, that other form of authentication will not work. For more information on FTP user isolation, including AD user isolation, see Configuring FTP 7 User Isolation.

Note
"FTP" and "Anonymous" are reserved words. You cannot create IIS Manager user accounts that contain these names.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for FTP Authentication
- Custom Providers Dialog Box
- Add or Edit Custom Authentication Provider Dialog Box
- Edit Anonymous Authentication Credentials Dialog Box
- Edit Basic Authentication Settings

UI Elements for FTP Authentication

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Anonymous Authentication	Anonymous authentication is a built-in authentication method that allows any user to access any public content by providing an anonymous user name and password. By default, Anonymous authentication is disabled. **Note** Use Anonymous authentication when you want all clients who visit your FTP site to be able to view its content.
ASP.NET Authentication	ASP.NET authentication is a custom authentication method that requires users to provide a valid .NET user name and password to gain access to content. The .NET account can be from an ASP.NET user database that is shared with your web content, or from a separate ASP.NET user database. **Note** ASP.NET authentication requires that a

Element Name	Description
	provider and possibly a connection string be configured for access to an ASP.NET user database.
Basic Authentication	Basic authentication is a built-in authentication method that requires users to provide a valid Windows user name and password to gain access to content. The user account can be local to the FTP server, or a domain account. **Note** Basic authentication transmits unencrypted passwords across the network. Use Basic authentication only when you know that the connection between the client and the server is secured using SSL.
IIS Manager Authentication	IIS Manager authentication is a custom authentication method that requires users to provide a valid IIS Manager user name and password to gain access to content. IIS Manager authentication requires that the IIS Management Service is installed and configured to use both Windows credentials and IIS Manager credentials. (The IIS Management Service does not have to be running when you use IIS Manager authentication.) **Note** IIS Manager authentication transmits unencrypted passwords across the network. Use IIS Manager authentication only when you know that the connection between the client and the server is secured using SSL.

Actions Pane Elements

Element Name	Description
Custom Providers	Opens the **Custom Providers** dialog box so that you can enable the default custom providers or your own custom providers.
Enable	Enables the selected provider. This option is only available if the selected provider is disabled.
Disable	Disables the selected provider. This option is only available if the selected provider is enabled.
Edit	Enables you to edit settings for Built-In providers. **Note** To edit settings for Custom providers, access the **Custom Providers** dialog box from the global level and click **Edit**.
Remove	Removes the selected provider. This option is not available for Built-in providers.

Note

When you configure your FTP authentication settings, also configure your FTP authorization settings.

Custom Providers Dialog Box

Use the **Custom Providers** dialog box to add or remove FTP custom authentication providers for your FTP server.

Element Name	Description
Select one or more registered providers	Displays the list of FTP custom authentication providers that have been registered on the FTP server. Selecting an FTP custom authentication provider from the list will enable that provider for the whole FTP server or for a specific FTP site depending on the node that is highlighted in the IIS Manager tree. This element appears at both global and site

Element Name	Description
	levels. **Note** Custom authentication providers can be either a managed-code class or a native-code COM module.
Register	Displays the **Add Custom Authentication Provider** dialog box. You can use this dialog box to add new FTP custom authentication providers to your FTP server.
Edit	Displays the **Edit Custom Authentication Provider** dialog box. You can use this dialog box to edit the settings for an FTP custom authentication provider that you have already registered. This element only appears at the global level.
Remove	Removes an FTP custom authentication provider that you have already registered. This element only appears at the global level.

Add or Edit Custom Authentication Provider Dialog Box

Use the **Add Custom Authentication Provider** or **Edit Custom Authentication Provider** dialog boxes to add new FTP custom authentication providers to your FTP server, or to edit the settings for an FTP custom authentication provider that you have already added to your FTP server.

Element Name	Description
Name	Specifies the name that will be displayed in the **Select one or more registered custom providers** list for the **FTP Custom Providers Dialog Box** dialog box.
Class ID	Specifies the class name or GUID for a COM class when the **Native Provider (COM)** option is used.
Type	Specifies the managed-type for a .NET class when the **Managed Provider (.NET)** option is

Element Name	Description
	used.
Application Domain	Specifies the application domain for a .NET class when the **Managed Provider (.NET)** option is used.

Edit Anonymous Authentication Credentials Dialog Box

Use the **Edit Anonymous Authentication Credentials** dialog box to specify the credentials to use for anonymous user authentication.

Element Name	Description
User name	Specifies the user account for anonymous user authentication.
Password	Specifies the password for the user account.
Confirm password	Confirms the password for the user account.

Edit Basic Authentication Settings

Use the **Edit Basic Authentication Settings** dialog box to set the default realm for Basic authentication.

Element Name	Description
Default domain	Specifies the default domain for Basic authentication.

FTP Authorization Rules

You can use the **FTP Authorization Rules** page to manage the list of Allow or Deny rules that control access to content. The rules appear on a list, and you can change the order of the rules to grant access to some users and deny access to others at the same time. Additionally, you can use the **FTP Authorization Rules** page to view information about other rules such as Mode, Users, Roles, or Permissions.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for FTP Authorization Rules
- Allow and Deny Authorization Rule Dialog Boxes

UI Elements for FTP Authorization Rules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Mode	Displays the type of rule. Values can be either **Allow** or **Deny**. Mode values indicate whether the rule is designed to allow or deny access to content. **Note** If a role, user, or group has been specifically denied access by one rule, it cannot be granted access by another. To grant access, create the rules with the Allow rule at the top of the list.
Users	Displays the user names or user groups to whom the rule applies.
Roles	Displays the Microsoft Windows role or .NET roles to which the rule applies, for example, the Administrator role.
Permissions	Displays the FTP permissions to which this rule applies, for example, Read or Write.

Actions Pane Elements

Element Name	Description
Add Allow Rule	Opens the **Add Allow Authorization Rule** dialog box so that you can create an Allow rule.
Add Deny Rule	Opens the **Add Deny Authorization Rule**

Element Name	Description
	dialog box so that you can create a Deny rule.
Edit	Enables you to edit the selected rule.
Remove	Removes the selected rule.

Note

When you configure your FTP authorization settings, also configure your FTP authentication settings.

Allow and Deny Authorization Rule Dialog Boxes

Use the **Add Allow Authorization Rule**, the **Edit Allow Authorization Rule,** the **Add Deny Authorization Rule**, or the **Edit Deny Authorization Rule** dialog box to define rules for access to content.

Important
Microsoft Windows authorization rules specify that if a user has been denied access, the Deny rule takes precedence over all other rules.

Element Name	Description
All Users	Select this option to manage access to content for both anonymous and authenticated users. **Note** Be sure to position this rule underneath any rules that grant access to content. If this rule is at the top of the rules list, all users are denied access to content.
All Anonymous Users	Select this option to manage access to content for users who are not authenticated. **Note** If you use this rule, all users must have a valid basic or custom authentication user account and password to be able to authenticate.
Specified roles or user groups	Select this option to manage access to content for specific Microsoft Windows roles or user groups.

Element Name	Description
	Note If you use this rule, all members of the specified roles and groups must have a valid basic or custom authentication user account and password to be able to authenticate.
Specified users	Select this option to manage access to content for a specific user account. **Note** If you use this option, all users must have a valid basic or custom authentication user account and password to be able to authenticate.
Read	Specifies whether the specified users have read permissions.
Write	Specifies whether the specified users have write permissions.

FTP Directory Browsing

Use the **FTP Directory Browsing** feature page to modify the content settings for browsing a directory on the FTP server. When you configure directory browsing, all directories use the same settings.

Related scenarios

- Build an FTP Site on IIS

UI Elements for FTP Directory Browsing

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Directory Listing Style	Specifies the format to use when listing the

Element Name	Description
	contents of a directory. The options are as follows: • **MS-DOS** - Displays files and folder using the MS-DOS format. • **UNIX** - Displays files and folder using the UNIX format.
Virtual directories	Specifies whether to display virtual directories when listing the contents of a directory. If enabled, displays virtual directories; otherwise, virtual directories are hidden.
Available bytes	Specifies whether to display the available bytes when listing the contents of a directory. **Note** The available bytes listed reflects the size left on disk, or, if a disk-level or folder-level quota is enabled, the available bytes remaining.
Four-digit years	Specifies the year format to use when displaying the last modified date for each file. If enabled, displays dates with four-digit years; otherwise, displays dates with two-digit years.

Actions Pane Elements

Element Name	Description
Apply	Applies changes to the **FTP Directory Browsing** feature page.
Cancel	Cancels changes to the **FTP Directory Browsing** feature page.

Note
In addition to configuring the settings for directory browsing on an FTP server, you can use the FTP Messages feature to customize the messages that are displayed when users access your FTP server.

FTP Firewall Support

Use the **FTP Firewall Support** feature page to modify the settings for passive connections when FTP clients are connecting to an FTP server that is located behind a firewall server.

Note
If you configure these options, you only control the behavior of the FTP server. You still have to configure your firewall to accept FTP connections and route them appropriately to your FTP server.

Related scenarios

- Build an FTP Site on IIS

UI Elements for FTP Firewall Support

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Data Channel Port Range	Specifies the port range for passive connections used for data channel connections. The valid range for ports is 1025 through 65535. (Ports from 1 through 1024 are reserved for use by system services.) **Note** You can enter a special port range of "0-0" to configure the FTP server to use the Windows TCP/IP ephemeral port range, which is often set to use ports 1025 through 5000 by default.
External IP address of Firewall	Specifies the IPv4 address for the external IP address of the firewall server. **Note** This setting is required when your FTP server is located behind a firewall server and you are using SSL, or your firewall does not modify the IP addresses in the FTP packets that are

Element Name	Description
	exchanged by the client and server.

Actions Pane Elements

Element Name	Description
Apply	Applies changes to the **FTP Firewall Support** feature page.
Cancel	Cancels changes to the **FTP Firewall Support** feature page.

FTP IP Address and Domain Restrictions

Use the **IP Address and Domain Restrictions** feature page to define and manage rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a domain name or names.

To configure restrictions based on domain names, enable domain name restrictions by clicking **Edit Feature Settings** in the task list, and then set the **Enable domain name restrictions** option.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for IP Address and Domain Restrictions
- Add Allow or Deny Restriction Rule Dialog Box
- Edit IP and Domain Restrictions Dialog Box

UI Elements for IP Address and Domain Restrictions

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Mode	Displays the type of rule. Values are either **Allow** or **Deny**. The **Mode** value indicates whether the rule is designed to allow or deny access to content.
Requester	Displays a specific IP address, range of IP addresses, or domain name that is defined in the **Add Allow Restriction Rule** and **Add Deny Restriction Rule** dialog boxes. You can specifically allow or deny a requester access to content.

Actions Pane Elements

Element Name	Description
Add Allow Entry	Opens the **Add Allow Restriction Rule** dialog box from which you can define rules that allow access to content for a specific IP address, a range of IP addresses, or a DNS domain name.
Add Deny Entry	Opens the **Add Deny Restriction Rule** dialog box from which you can define rules that allow access to content for a specific IP address, a range of IP addresses, or a DNS domain name.
Remove	Removes the item that is selected from the list on the feature page.
Edit Feature Settings	Opens the **Edit IP and Domain Restrictions Settings** dialog box from which you can configure settings that apply to the entire IP and domain name restrictions feature.
Revert to Parent	Reverts the feature to inherit settings from the parent configuration. This action deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other

Element Name	Description
	actions in the **Actions** pane do not appear until you select the unordered list format.
Move Up	Moves up a selected item in the list. This action is available only when you view items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when you view items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.

Add Allow or Deny Restriction Rule Dialog Box

Use either the **Add Allow Restriction Rule** or the **Add Deny Restriction Rule** dialog box to define rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a DNS domain name.

Element Name	Description
Specific IP Address	**Add Allow Restriction Rule** - Type an IP address in the **Specific IP Address** box in the **Add Allow Restriction Rule** dialog box when you want to allow access to content for a specific IP address. **Add Deny Restriction Rule** - Type an IP Address in the **Specific IP Address** box in the **Add Deny Restriction Rule** dialog box when you want to deny access to content for a specific IP address.
IP address range	**Add Allow Restriction Rule** - Type the lowest value of the range of IP addresses that you have chosen to use in the **IP Address range** box in the **Add Allow Restriction Rule** dialog box. Do this action when you want to allow access to content for a range of IP addresses. Next, enter the subnet mask. **Add Deny Restriction Rule** - Type the lowest value of the range of IP addresses that you have chosen to use in the **IP address range** box in the **Add Deny Restriction Rule** dialog box. Do this action when you want to deny access to content for a range of IP address. Next, enter the subnet mask.
Mask	**Add Allow Restriction Rule** - Type a subnet mask in the **Mask** box in the **Add Allow Restriction Rule** dialog box. Do this action when you want to allow access to content for a range of IP address. When IIS evaluates this subnet mask with the IP address entered in the **IP address range** box, the upper and lower boundaries of an IP address space are defined.

Element Name	Description
	Add Deny Restriction Rule - Type the subnet mask associated with the range of IP addresses in the **Mask** box in the **Add Deny Restriction Rule** dialog box. Do this action when you want to deny access to content for a range of IP address. When IIS evaluates this subnet mask with the IP address entered in the **IP address range** box, the upper and lower boundaries of an IP address space are defined.
Domain name	**Add Allow Restriction Rule** - Type a fully qualified DNS domain name in the **Domain name** box in the **Add Allow Restriction Rule** dialog box when you want to allow access to content for a DNS domain. **Add Deny Restriction Rule** - Type a fully qualified DNS domain name in the **Domain name** box in the **Add Deny Restriction Rule** dialog box when you want to deny access to content for a DNS domain. To see the **Domain name** option, you must first enable domain name restrictions by using **Edit Feature Settings**.

Edit IP and Domain Restrictions Dialog Box

Use the **Edit IP and Domain Restrictions** dialog box to define access restrictions for unspecified clients or to enable domain name restrictions for all rules.

Element Name	Description
Access for unspecified clients	Defines access restrictions for unspecified clients. This setting defines whether to allow or deny access to clients not specified by any other rule.
Enable domain name restrictions	Enables rules that restrict access by domain name. This rule significantly affects server performance because it requires a DNS lookup

Element Name	Description
	for every request.

FTP Logging

Use the **FTP Logging** feature page to configure logging features at the server or site level, and to configure logging settings.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for FTP Logging
- Information To Log Dialog Box

UI Elements for FTP Logging

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
One log file per	Specifies whether to use a single log file for the entire FTP server or to use separate log files on a site-by-site basis. The options are: **Server** - Specifies that you use a single log file for the entire FTP server. **Site** - Specifies that you use separate log files for each FTP site. **Note** Using a single log file for your server helps when you are using log parsing utilities. But if it becomes too large, a single log file could affect the overall performance of the FTP server.
Select W3C Fields	Displays the **Information To Log** dialog box. **Note**

Element Name	Description
	This version of the FTP server only supports W3C Extended Log File format.
Directory	Specifies the base folder to store log files. **Note** When you are using separate log files on a site-by-site basis, the base folder is appended with the service name and unique ID for the site. For example, these identifiers might be appended: FTPSVC1 or FTPSVC2. **Warning** When you specify a UNC path for the logging folder, make sure that you configure the permissions for the FTP service to create log files on the destination share.
Encoding	Specifies the file encoding for log files. The options are: **UTF8** - Enables both single-byte and multibyte characters in one string. **ANSI** - Enables only single-byte characters in a string. **Note** UTF8-encoded log files enable you to read text-based entries in a language other than English. Additionally, if your FTP server serves URLs in a language other than the one supported by the server's default code page, enable UTF-8 encoding for log content.
Schedule	Specifies creating new log files that are based on a fixed time interval. The options are: • **Hourly**: Creates a log file every hour. • **Daily**: Creates a log file daily. • **Weekly**: Creates a log file each week. • **Monthly**: Creates a log file once a month.

Element Name	Description
Maximum file size (in bytes)	Specifies whether to create new log files when file size exceeds the maximum file size entered in the **Maximum File Size** box.
Do not create new log files	Specifies that you use the current log file without creating a new one. Selecting this option sets the size of the current log file to unlimited. **Note** Using a single log file for your server helps when you are using log parsing utilities. But if it becomes too large, a single log file could affect the overall performance of the FTP server.
Use local time for file naming and rollover	Specifies rotating the log files that are based on the local time zone (if using the **Schedule** option), instead of Coordinated Universal Time (UTC). By default, log files roll over at midnight UTC instead of at midnight according to the local time zone.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Disable	Disables the feature.
View Logs	Opens the log file directory.

Information To Log Dialog Box

Use the **Information to Log** dialog box to configure the logging options or other content you want to include in your W3C Extended Log Files.

Element Name	Description
Date (date)	Logs the date on which the request occurred. Selected by default.
Time (time)	Logs the time, in Coordinated Universal Time (UTC), at which the request occurred. Selected by default.
Client IP Address (c-ip)	Logs the IP address of the client that made the request. Selected by default.
User Name (cs-username)	Logs the name of the authenticated user who accessed the server. Anonymous users are indicated by a hyphen. Selected by default.
Service Name (s-sitename)	Logs the internet service name and instance number that was running on the client at the time that the incident was logged.
Server Name (s-computername)	Logs the name of the server on which the log file entry was generated.
Server IP Address (s-ip)	Logs the IP address of the server on which the log file entry was generated. Selected by default.
Server Port (s-port)	Logs the server port number that is configured for the service. Selected by default.
Method (cs-method)	Logs the HTTP method, such as `GET`, that is used in the request. Selected by default.
URI Stem (cs-uri-stem)	Logs the Uniform Resource Identifier (URI) that is the target of the action. Selected by default. **Note** The URI for FTP requests usually applies only for a file or folder and does not log the full path of the file or folder. This makes reading log files difficult when you are working with a large volume of content. Because of this limitation, you can enable the **Full Path** field to list the detailed path information for the request.
Protocol Status (sc-status)	Logs the HTTP status code. Selected by default.

Element Name	Description
Win32 Status (sc-win32-status)	Logs the Windows status code. Selected by default.
Bytes Sent (sc-bytes)	Logs the number of bytes that the server sent.
Bytes Received (cs-bytes)	Logs the number of bytes that the server received.
Time Taken (time-taken)	Logs the length of time that the action took in milliseconds.
Server Port (s-port)	Logs the server port number that is configured for the service. Selected by default.
Host (cs-host)	Logs the host header name, if any. **Note** The host name that you configure for your website might appear differently in the log files because HTTP.sys logs host names by using the Punycode-encoded format. (Punycode encodes Unicode strings into the limited character set permitted in host names. For more information, see RFC 3492.)
Protocol Substatus (cs-substatus)	Logs the protocol version, HTTP or FTP, that the client used. Selected by default.
Session ID (x-session)	Logs the FTP session identifier for the client's session. Selected by default.
Full Path (x-fullpath)	Logs the full relative path from the FTP root directory for the target of the action. Selected by default.
Additional Information (x-debug)	Logs descriptive information for the sc-status code. For example, if the sc-status code is 552, the additional information logged may be "User disk quota exceeded."
Client Port (c-port)	Logs the port of the client that made the request.

FTP Logon Attempt Restrictions

Use the **FTP Logon Attempt Restrictions** feature page to modify the number of times FTP will allow user to attempt unsuccessfully to log in within a specified time period before denying access to the IP address.

Related scenarios

- Build an FTP Site on IIS

UI Elements for FTP Logon Attempt Restrictions

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Enable FTP Logon Attempt Restrictions	Enables and disables the FTP Logon Attempt Restrictions feature.
Maximum number of failed login attempts	Sets the number of times a user can attempt to log in unsuccessfully.
Time period (in seconds)	Sets the time period within which the failed attempts are accumulated.
Deny IP addresses based on the number of failed login attempts	Denies the IP address if the user attempts to log in unsuccessfully the maximum number of times within the given time period.
Write to the log only	Writes to the log if the user attempts to log in unsuccessfully the maximum number of times within the given time period. The IP address is still allowed access.

Actions Pane Elements

Element Name	Description
Apply	Applies changes to the **FTP Logon Attempt Restrictions** feature page.
Cancel	Cancels changes to the **FTP Logon Attempt**

Element Name	Description
	Restrictions feature page.

FTP Messages

Use the **FTP Messages** feature page to modify the settings for messages sent when a user connects to your FTP site.

Related scenarios

- Build an FTP Site on IIS

UI Elements for FTP Messages

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Suppress default banner	Specifies whether to display the default identification banner for the FTP server. If enabled, displays the default banner; otherwise, the default banner is not displayed. **Note** If **Suppress default banner** is enabled and a banner message has not been specified in **Banner**, the FTP server will display a blank banner when an FTP client connects to your server.
Support user variables in messages	Specifies whether to display a specific set of user variables in FTP messages. If enabled, displays user variables in FTP messages; otherwise, all message text will be displayed as entered. The supported user variables are: • **%BytesReceived%** - The number of bytes sent from the server to the client for the current session. • **%BytesSent%** - The number of bytes

Element Name	Description
	sent from the client to the server for the current session. • **%SessionID%** - The unique identifier for the current session. • **%SiteName%** - The name of the FTP site that is hosting the current session. • **%UserName%** - The account name of the currently logged-on user.
Show detailed messages for local requests	Specifies whether to display detailed error messages when the FTP client is connecting to the FTP server on the server itself. If enabled, displays detailed error messages only to the local host; otherwise, detailed error messages are not displayed. **Note** Detailed error messages are shown only to the local host.
Banner	Specifies the message the FTP server displays when FTP clients first connect to the FTP server. **Note** By default, this message is blank. If **Suppress default banner** is enabled and a banner message has not been specified in **Banner**, the FTP server will display a blank banner when an FTP client connects to your server.
Welcome	Specifies the message the FTP server displays when FTP clients have logged in to the FTP server. **Note** By default, this message is blank.
Exit	Specifies the message the FTP server displays when FTP clients log off the FTP server. **Note** By default, this message is blank.
Maximum connections	Specifies the message the FTP server displays

Element Name	Description
	when clients try to connect and cannot because the FTP service has reached the maximum number of client connections allowed. **Note** By default, this message is blank.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Note

In addition to configuring the settings for messages the FTP server displays when a user connects to your FTP site, you can use the FTP Directory Browsing Page feature to customize the content settings for browsing a directory on the FTP server.

FTP Network Security

Use **FTP Network Security** to configure login security parameters and protect your FTP server from network-based attacks. This feature enables you to deny access to an IP address based on a specified number of failed login attempts. However, you can, if desired, configure your site to log the failed login attempts without denying access.

Related scenarios

- Build an FTP Site on IIS

UI Elements for FTP Network Security

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Deny IP addresses based on the number of failed login attempts	Restricts clients from access the FTP server after a specified number of failed login attempts.
Maximum number of failed login attempts	Sets the maximum failed login attempts after which the client is denied access. The default is 4 attempts.
Time period (in seconds)	Sets the number of seconds to wait before returning from a failed login attempt. The default is 30 seconds.
Write to the log only	Logs the failed attempts but allows access to the FTP server. This setting will not restrict clients even if the maximum failed attempts are exceeded.

Action Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

FTP Request Filtering

Use the **FTP Request Filtering** feature page to define the request filtering settings for your FTP site. FTP request filtering is a security feature that allows internet service providers (ISPs) and application service providers to restrict protocol and content behavior. For example, using the **File Name Extensions** tab you can specify a list of file name extensions that are allowed or denied.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Element List
- File Name Extensions
- Allow or Deny File Name Extension Dialog Box

- Hidden Segments
- Add Hidden Segment Dialog Box
- Denied URL Sequences
- Add Deny Sequence Dialog Box
- Edit FTP Request Filtering Settings Dialog Box
- FTP Commands
- Allow or Deny Command Dialog Box

UI Element List

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
File Name Extensions	Specifies a list of file name extensions for which the FTP service allows or denies access.
Hidden Segments	Specifies a list of hidden segments for which the FTP service denies access and does not display in directory listings.
Denied URL Sequences	Specifies a list of URL sequences for which the FTP service denies access.
Commands	Specifies a list of FTP commands for which the FTP service allows or denies access.

Actions Pane Elements

Element Name	Description
Edit Feature Settings	Opens the **Edit FTP Request Filtering Settings** dialog box for you to configure general properties and FTP request limits.
Allow File Name Extension	Opens the **Allow File Name Extension** dialog box for you to add a file name extension to the list of allowed file name extensions.
Deny File Name Extension	Opens the **Deny File Name Extension** dialog box for you to add a file name extension to the

Element Name	Description
	list of denied file name extensions.
Add Hidden Segment	Opens the **Add Hidden Segment** dialog box for you to add a hidden segment to the list of hidden segments.
Add URL Sequence	Opens the **Add Deny URL Sequence** dialog box for you to add a URL sequence to the list of denied URL sequences.
Allow Command	Opens the **Allow Command** dialog box for you to add an FTP command to the list of allowed FTP commands.
Deny Command	Opens the **Deny Command** dialog box for you to add an FTP command to the list of denied FTP commands. **Warning** Using this feature incorrectly can prevent access to your server. For example, if you deny access to the USER and PASS commands, users are not able to log in to your FTP server.
Remove	Removes a file name extension, hidden segment, URL sequence, or command from the list.

File Name Extensions

Use FTP **File Name Extensions** to define the list of file name extensions for which the FTP service will either allow or deny access. Configuring specific file name extensions allows web server administrators to customize which file name extensions that the FTP service allows or denies, which you can use to tighten security on your server. For example, if you deny access to *.exe and *.com files, you can prevent internet clients from uploading executable files to your web server.

Note

The **Edit FTP Request Filtering Settings** dialog box allows you to specify whether unlisted file name extensions are allowed. You can fine-tune the file access restrictions for your server by specifying that unlisted file name extensions are not allowed, and using the FTP File Name Extensions feature to manually specify the list of file name extensions that are allowed on your server.

Feature Page Elements

Element Name	Description
File Name Extension	Displays the file name extension for which the FTP service either allows or denies access.
Allowed	Displays the status of the file name extension, which will be either **True** if the file name extension is allowed, or **False** if the file name extension is denied.

Actions Pane Elements

Element Name	Description
Allow File Name Extension	Opens the **Allow File Name Extension** dialog box for you to add a file name extension to the list of allowed file name extensions.
Deny File Name Extension	Opens the **Deny File Name Extension** dialog box for you to add a file name extension to the list of denied file name extensions.
Remove	Removes a file name extension, hidden segment, URL sequence, or command from the list.
Edit Feature Settings	Opens the **Edit FTP Request Filtering Settings** dialog box for you to configure general properties and FTP request limits.

Allow or Deny File Name Extension Dialog Box

Use the **Allow File Name Extension** or **Deny File Name Extension** dialog box to add a file name extension to the list of file name extensions for which the FTP service will either allow or deny access.

Element Name	Description
File name extension	Specifies the file name extension for which the FTP service either allows or denies access.

Hidden Segments

Use FTP **Hidden Segments** to define the list of URL segments for which the FTP service will deny access and will not display in directory listings. For example, to prevent access to the Bin directory of your web application, you could add the Bin directory as a hidden segment for your FTP site. When an FTP client logs in to your FTP site, the Bin folder is not displayed in directory listings. If the FTP client attempts to change to the Bin folder, the FTP service returns an access denied error message to the FTP client.

Feature Page Elements

Element Name	Description
Segment	Displays the URL segment for which the FTP service denies access and does not display in directory listings.

Actions Pane Elements

Element Name	Description
Add Hidden Segment	Opens the **Add Hidden Segment** dialog box for you to add a hidden segment to the list of hidden segments.
Remove	Removes a hidden segment from the list.
Edit Feature Settings	Opens the **Edit FTP Request Filtering Settings** dialog box for you to configure general properties and FTP request limits.

Add Hidden Segment Dialog Box

Use the **Add Hidden Segment** dialog box to add a URL segment to the list of URL segments for which the FTP service will deny access.

Element Name	Description
Hidden segment	Specifies the URL segment for which the FTP service denies access. **Note** The FTP service does not show hidden

Element Name	Description
	segments in directory listings.

Denied URL Sequences

Use FTP **Denied URL Sequences** to define the list of URL sequences for which the FTP service will deny access. For example, to prevent access to the Bin directory of your web application, you could add the Bin directory as a denied URL sequences for your FTP site. When an FTP client logs in to your FTP site, the Bin folder is displayed in directory listings; however, if the FTP client attempts to change to the Bin folder, the FTP service returns an access denied error message to the FTP client.

Feature Page Elements

Element Name	Description
URL Sequence	Displays the FTP URL sequence for which the FTP service denies access.

Actions Pane Elements

Element Name	Description
Add URL Sequence	Opens the **Add Deny Sequence** dialog box for you to add a URL sequence to the list of denied URL sequences.
Remove	Removes a file name extension, hidden segment, URL sequence, or command from the list.
Edit Feature Settings	Opens the **Edit FTP Request Filtering Settings** dialog box for you to configure general properties and FTP request limits.

Add Deny Sequence Dialog Box

Use the FTP **Add Deny Sequence** dialog box to add a URL sequence to the list of URL sequences for which the FTP service will deny access.

Element Name	Description
URL sequence	Specifies the URL sequence for which the FTP service denies access.

Edit FTP Request Filtering Settings Dialog Box

Use the Edit FTP Request Filtering Settings dialog box to specify general FTP request filtering settings, request limits, and command filtering options.

Element Name	Description
Allow unlisted file name extensions	Select this option to allow unlisted file name extensions.
Allow high-bit characters	Select this option to allow high-bit characters in FTP operations.
Maximum content length (Bytes)	Specifies the maximum length, in bytes, for FTP content. **Note** Specifying zero (0) means unlimited.
Maximum URL length (Bytes)	Specifies the maximum length, in bytes, for an FTP URL.
Maximum command length (Bytes)	Specifies the maximum length, in bytes, for an FTP command.
Allow unlisted commands	Select this option to allow FTP commands that are not in the FTP **Commands** list.

FTP Commands

Use FTP **Commands** to define the list of commands for which the FTP service will either allow or deny access. Configuring specific commands allows Web server administrators to customize the list of FTP commands that the FTP service allows, which you can use to tighten security on your server. For example, if you deny access to the FTP SYST command, you can prevent internet clients from determining your server's operating system.

Note

The **Edit FTP Request Filtering Settings** dialog box allows you to specify whether unlisted commands are allowed, which you can use to manually specify which commands are allowed.

Warning

Using the FTP **Commands** feature incorrectly can prevent access to your FTP server. For example, if you deny access to the USER or PASS commands, users are not able to log in to your FTP server. In addition, even though you could specify that unlisted commands are not allowed in the **Edit FTP Request Filtering Settings** dialog box and use the **Commands** list to specify only the FTP commands that are allowed, this configuration is not recommended due to the complexity involved in specifying the correct list of FTP commands that are required for normal FTP operation.

Feature Page Elements

Element Name	Description
Command	Displays the FTP command for which the FTP service either allows or denies access.
Allowed	Displays the status of the FTP command, which will be either **True** if the FTP command is allowed, or **False** if the FTP command is denied.

Actions Pane Elements

Element Name	Description
Allow Command	Opens the **Allow Command** dialog box for you to add an FTP command to the list of allowed FTP commands.
Deny Command	Opens the **Deny Command** dialog box for you to add an FTP command to the list of denied FTP commands. **Warning** Using this feature incorrectly can prevent access to your server. For example, if you deny access to the USER and PASS commands, users are not able to log in to your FTP server.
Remove	Removes an FTP command from the list.
Edit Feature Settings	Opens the **Edit FTP Request Filtering Settings** dialog box for you to configure

Element Name	Description
	general properties and FTP request limits.

Allow or Deny Command Dialog Box

Use the **Allow Command** or **Deny Command** dialog box to add an FTP command to the list of commands for which the FTP service will either allow or deny access.

Warning
Using this feature incorrectly can prevent access to your server. For example, if you deny access to the USER and PASS commands, users are not able to log in to your FTP server.

Element Name	Description
Command	Specifies the FTP command for which the FTP service either allows or denies access.

FTP Sites

The FTP Sites feature represents the UI components that are available on the IIS Manager **Sites** page when FTP is installed. Those components include additions to settings pages and dialog boxes.

Related scenarios

- Build an FTP Site on IIS

In this document

- Add FTP Site Wizard
- FTP Site Defaults Dialog Box
- FTP Site Advanced Settings Dialog Box
- Add FTP Site Publishing Wizard
- FTP Current Sessions Page

Add FTP Site Wizard

The Add FTP Site wizard guides you through the steps of creating an FTP site. The wizard displays the following pages in the order shown.

- Site Information Page
- Binding and SSL Settings Page

- Authentication and Authorization Information Page

Site Information Page

Use the **Site Information** page of the FTP Site Publishing Wizard to set site settings.

Element Name	Description
FTP site name	Specifies the friendly name for the FTP site.
Physical path	Specifies the physical path for the FTP site.

Binding and SSL Settings Page

Use the **Binding and SSL Settings** page of the FTP Site Publishing Wizard to set bindings and SSL.

Element Name	Description
IP Address	Specifies the IP address for the FTP site. The default value is **All Unassigned**.
Port	Specifies the port on which the FTP site listens for incoming requests. The default value for FTP is 21.
Enable Virtual Host Names	Specifies whether FTP virtual host names are enabled. **Note** Some FTP clients do not support FTP virtual host names.
Virtual Host	Specifies a virtual host name. This setting enables you to have multiple FTP sites on a single IP address by using the default port.
Start FTP site automatically	Specifies whether to start the FTP site automatically when the FTP service starts.
No SSL	Specifies that SSL is not used.
Allow SSL	Specifies whether SSL is allowed.
Require SSL	Specifies whether SSL is required. This setting is the default setting.
SSL Certificate	Specifies the SSL Certificate to use.

Element Name	Description
	Note This drop-down menu is blank if no SSL Certificates are installed.

Authentication and Authorization Information Page

Use the **Authentication and Authorization Information** page of the FTP Site Publishing Wizard to set authentication and authorization settings.

Element Name	Description
Anonymous	Specifies whether Anonymous authentication is enabled.
Basic	Specifies whether Basic authentication is enabled.
Allow access to	Specifies who is authorized to log on to the FTP site. Available options: • **All users** • **Anonymous users** • **Specified roles or user groups** • **Specified users** If you use **Specified roles or user groups** or **Specified users**, use the additional box to specify access for specific users, roles, or groups.
Read	Specifies whether the rule allows read permission.
Write	Specifies whether the rule allows write permission.

FTP Site Defaults Dialog Box

Use the **FTP Site Defaults** dialog box to set site defaults for FTP sites.

Element Name	Description
Allow UTF8	Specifies whether UTF8 activity should be

Element Name	Description
	allowed (this setting includes UTF8 file names). The default is **True**.
Start Automatically	Specifies whether to start the site upon creation or when you start IIS. Starting a site sets this property to true; stopping a site sets it to false. The default is **True**.
Control Channel Timeout	Specifies the time-out (in seconds) when an FTP session is automatically disconnected because of inactivity. The default is 120.
Data Channel Timeout	Specifies the time-out (in seconds) when the data channel is automatically disconnected because of inactivity. If you use the default time-outs for both the control and data channel, the data channel is torn down first, and then the session is disconnected if there is still no activity after 120 seconds. The default is 30.
Disable Socket Pooling	Specifies whether socket pooling is used for sites distinguished by IP address instead of by port number or host name. The default is **False**.
Max Connections	Specifies the maximum number of simultaneous connections to a server. The default is 4294967295.
Reset On Max Connections	Specifies whether to disconnect FTP session when sending max connections response. The default is **False**.
Server Listen Backlog	Specifies the number of outstanding sockets that can be queued. The default is 60.
Unauthenticated Timeout	Specifies the time-out (in seconds) between when a new connection is made and authentication is attempted.

Element Name	Description
	The default is 30.
Allow Reading Files While Uploading	Specifies whether files can be read while being transferred to the server. The default is **False**My.
Allow Replace on Rename	Specifies whether the files can overwrite other files when renamed. The default is **False**.
Keep Partial Uploads	Specifies whether to keep files that have been partially uploaded. The default is **False**.

FTP Site Advanced Settings Dialog Box

Use the **Advanced Settings** dialog box to manage or change site defaults for individual FTP sites.

Element Name	Description
Allow UTF8	Specifies whether UTF8 activity should be allowed (this setting includes UTF8 file names). The default is **True**.
Bindings	Specifies the combination of IP address, virtual host, and port where FTP should listen for incoming requests.
ID	Specifies the unique ID number for the site. This number is used instead of directory names for log and trace files.
Name	Specifies a unique name for the site.
Physical Path	Specifies the physical path of the content for the virtual directory.
Start Automatically	Specifies whether to start the site on creation or when you start IIS. Starting a site sets this property to **True**; stopping a site sets it to **False**. The default is **True**.
Control Channel Timeout	Specifies the time-out (in seconds) when an FTP session is automatically disconnected because of inactivity.

Element Name	Description
	The default is 120.
Data Channel Timeout	Specifies the time-out (in seconds) when the data channel is automatically disconnected because of inactivity. If you use the default time-outs for both the control and data channel, the data channel is torn down first, and then the session is disconnected if there is still no activity after 120 seconds. The default is 30.
Disable Socket Pooling	Specifies whether socket pooling is used for sites that are distinguished by IP address instead of port number or host name. The default is **False**.
Max Connections	Specifies the maximum number of simultaneous connections to a server. The default is 4294967295.
Reset On Max Connections	Specifies whether to disconnect the FTP session when sending the maximum connections response. The default is **False**.
Server Listen Backlog	Specifies the number of outstanding sockets that can be queued. The default is 60.
Unauthenticated Timeout	Specifies the time-out (in seconds) between when a new connection is made and when authentication is attempted. The default is 30.
Allow Reading Files While Uploading	Specifies whether files can be read as they are being transferred to the server. The default is **False**.
Allow Replace on Rename	Specifies whether files can overwrite other files when they are renamed. The default is **False**.
Keep Partial Uploads	Specifies whether to keep files that have been

Element Name	Description
	partially uploaded. (For example, when a file transfer ends prematurely.) The default is **False**. **Note** If the client closes the connection gracefully because of a failure, partially uploaded data might be deleted.

Add FTP Site Publishing Wizard

The Add FTP Site Publishing wizard guides you through the steps of adding FTP publishing to an existing website. The wizard displays the following pages in the order shown.

1. Binding and SSL Settings Page
2. Authentication and Authorization Information Page

Binding and SSL Settings Page

Use the **Binding and SSL Settings** page of the FTP Site Publishing Wizard to set bindings and SSL.

Element Name	Description
IP Address	Specifies the IP address for the FTP site. The default value is **All Unassigned**.
Port	Specifies the port on which the FTP site listens for incoming requests. The default value for FTP is 21.
Enable Virtual Host Names	Specifies whether FTP virtual host names are enabled. **Note** Some FTP clients do not support FTP virtual host names.
Virtual Host	Specifies a virtual host name. This setting enables you to have multiple FTP sites on a single IP address by using the default port.
Start FTP site automatically	Specifies whether to start the FTP site automatically when the FTP service starts.

Element Name	Description
No SSL	Specifies that SSL is not used.
Allow SSL	Specifies whether SSL is allowed.
Require SSL	Specifies whether SSL is required. This setting is the default.
SSL Certificate	Specifies the SSL Certificate to use. **Note** This drop-down menu is blank if no SSL Certificates are installed.

Authentication and Authorization Information Page

Use the **Authentication and Authorization Information** page of the FTP Site Publishing Wizard to set authentication and authorization settings.

Element Name	Description
Anonymous	Specifies whether Anonymous authentication is enabled.
Basic	Specifies whether Basic authentication is enabled.
Allow access to	Specifies who is authorized to log on to the FTP site. Available options: • **All users** • **Anonymous users** • **Specified roles or user groups** • **Specified users** If you use **Specified roles or user groups** or **Specified users**, use the additional box to specify access for specific users, roles, or groups.
Read	Specifies whether the rule allows read permission.
Write	Specifies whether the rule allows write permission.

FTP Current Sessions Page

Use the **Current Sessions** feature page to monitor current sessions for the FTP site.

Feature Page Elements

Element Name	Description
User name	Specifies the user name for this session.
Session start time	Specifies the start time for the current session.
Current command	Specifies the current command.
Previous command	Specifies the previous command.
Command start time	Specifies the start time for the command.
Bytes sent	Specifies the bytes sent.
Bytes received	Specifies the bytes received.
Session ID	Specifies the session ID.
Client IP	Specifies the client IP address.

Actions Pane Elements

Element Name	Description
Disconnect Session	Enables you to disconnect the selected session.

FTP SSL Settings

Use the **FTP SSL Settings** feature page to manage encryption for both control channel and data channel transmissions between your FTP server and clients.

Note
These settings work in both intranet and internet environments.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for FTP SSL Settings

- Advanced SSL Policy Dialog Box

UI Elements for FTP SSL Settings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
SSL Certificate	Specifies the SSL certificate that you want to use for SSL. **Note** Selecting an SSL certificate enables SSL for a specific FTP site or FTP server. To disable SSL, select **Not Selected** from the drop-down list.
View	Opens the **Certificate** dialog box for the selected certificate.
Allow SSL Connections	Select this setting to allow data encryption for both the control channel and data channel when an SSL certificate is selected.
Require SSL Connections	Select this setting to require data encryption for both the control channel and data channel when an SSL certificate is selected. **Note** SSL requires additional processor time and slows down data transmissions because of the additional encryption data that must be transmitted.
Custom	Select this setting to customize the individual data encryption requirements for the control channel and data channel when an SSL certificate is selected. **Note** Selecting the **Custom** option enables the **Advanced** button; clicking **Advanced** will display the **Advanced**

Element Name	Description
	SSL Policy dialog box.
Use 128-bit encryption for SSL connections	Select this setting to require stronger encryption. You can use 128–bit SSL to help secure transmissions between your FTP server and clients in either an intranet or internet environment. **Note** 128–bit SSL requires more processor time and slows down data transmissions more than the 40–bit version does. This setting is because 128–bit SSL transmits additional encryption data to provide stronger encryption. Therefore, before requiring 128–bit SSL, make sure that your data transmissions need stronger encryption.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Advanced SSL Policy Dialog Box

Use the **Advanced SSL Policy** dialog box to customize the individual data encryption requirements for the control channel and data channel between your FTP server and clients.

Element Name	Description
Control Channel: Allow	Specifies that data encryption is allowed for the control channel. This option allows the client to choose whether to use SSL for the control channel.

Element Name	Description
Control Channel: Require	Specifies that data encryption is required for the control channel. This option causes all clients to use SSL for all activity on the control channel.
Control Channel: Require only for credentials	Specifies that data encryption is required for the control channel only when user credentials are transmitted. **Note** Selecting this option causes all clients to use SSL for the control channel when user credentials are transmitted, and then allows the client to choose whether to continue to use SSL for the control channel.
Data Channel: Allow	Specifies that data encryption is allowed for the data channel. This option allows the client to choose whether to use SSL for the data channel.
Data Channel: Require	Specifies that data encryption is required for the data channel. This option causes all clients to use SSL for all activity on the data channel.
Data Channel: Deny	Specifies that data encryption is denied for the data channel. This option explicitly denies the use of SSL for all activity on the data channel.

FTP User Isolation

Use the **FTP User Isolation** feature page to define the user isolation mode for your FTP site. FTP user isolation is a solution for internet service providers (ISPs) who want to offer their customers individual FTP directories for uploading content. FTP user isolation prevents users from viewing or overwriting other users' content by restricting users to their own directories. Users cannot navigate higher up the directory tree because their top-level directory appears as the root of the FTP service. Within their specific site, users can create, modify, or delete files and folders.

Related scenarios

- Build an FTP Site on IIS

In this document

- UI Elements for FTP User Isolation

- Set Credentials Dialog Box

UI Elements for FTP User Isolation

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Do not isolate users. Start users in: FTP root directory	Select this option to specify that you do not want to isolate users. All FTP sessions starts in the root directory of the FTP site. **Caution** If they have sufficient permissions, any FTP user can potentially access the content of any other FTP user.
Do not isolate users. Start users in: User name directory	Select this option to specify that you do not want to isolate users. All FTP sessions starts in the physical or virtual directory with the same name of the currently logged-on user if the folder exists; otherwise, the FTP session starts in the root directory of the FTP site. **Note** To specify the starting directory for anonymous access, create a physical or virtual directory folder named **default** in the root directory of the FTP site. **Caution** If they have sufficient permissions, any FTP user can potentially access the content of any other FTP user.
Isolate users. Restrict users to the following directory: User name directory (disable global virtual directories)	Select this option to specify that you want to isolate FTP user sessions to the physical or virtual directory with the same name of the FTP user account. The user sees only their FTP root location and is restricted from navigating higher up the directory tree. **Note** To create home directories for each user, you first must create a physical or virtual directory under your FTP server's root folder that is named after your domain or

<table>
<tr><th>Element Name</th><th>Description</th></tr>
<tr><td></td><td>named LocalUser for local user accounts. Next, create a physical or virtual directory for each user account that accesses your FTP site. The following table lists the home directory syntax for the authentication providers that are included with the FTP service:

<table>
<tr><th>User Account Types</th><th>Home Directory Syntax</th></tr>
<tr><td>Anonymous users</td><td>%FtpRoot%\LocalUser\Public</td></tr>
<tr><td>Local Windows user accounts
(Requires Basic authentication.)</td><td>%FtpRoot%\LocalUser\%UserName%</td></tr>
<tr><td>Windows domain accounts
(Requires Basic authentication.)</td><td>%FtpRoot%\%UserDomain%\%UserName%</td></tr>
<tr><td>IIS Manager or ASP.NET custom authentication user accounts</td><td>%FtpRoot%\LocalUser\%UserName%</td></tr>
</table>

Note
%FtpRoot% is the root directory for your FTP site: for example, C:\Inetpub\Ftproot.

Important
Global virtual directories are ignored. No FTP users can access virtual directories that are configured at the root-level of your FTP site. All virtual directories must be defined explicitly under a user’s physical or virtual home directory path.</td></tr>
<tr><td>Isolate users. Restrict users</td><td>Select this option to specify that you want to isolate FTP user</td></tr>
</table>

<table>
<tr><th>Element Name</th><th>Description</th></tr>
<tr><td>to the following directory: User name physical directory (enable global virtual directories)</td><td>sessions to the physical directory with the same name of the FTP user account. The user sees only their FTP root location and is restricted from navigating higher up the directory tree.

Note
To create home directories for each user, you first must create a physical directory under your FTP server's root folder that is named after your domain or named LocalUser for local user accounts. Next, create a physical directory for each user account that accesses your FTP site. The following table lists the home directory syntax for the authentication providers that are included with the FTP service:

<table>
<tr><th>User Account Types</th><th>Home Directory Syntax</th></tr>
<tr><td>Anonymous users</td><td>%FtpRoot%\LocalUser\Public</td></tr>
<tr><td>Local Windows user accounts
(Requires Basic authentication.)</td><td>%FtpRoot%\LocalUser\%UserName%</td></tr>
<tr><td>Windows domain accounts
(Requires Basic authentication)</td><td>%FtpRoot%\%UserDomain%\%UserName%</td></tr>
<tr><td>IIS Manager or ASP.NET custom authentication user accounts</td><td>%FtpRoot%\LocalUser\%UserName%</td></tr>
</table>

Note
%FtpRoot% is the root directory for your FTP site; for example, C:\Inetpub\Ftproot.</td></tr>
</table>

Element Name	Description
	Important Global virtual directories are enabled. All virtual directories that are configured at the root-level of your FTP site can be accessed by all FTP users, if those users have sufficient permissions. **Caution** When global virtual directories are enabled, all FTP users can potentially access the content of other FTP users, if those users have sufficient permissions.
Isolate users. Restrict users to the following directory: FTP home directory configured in Active Directory	Select this option to specify that you want to isolate FTP user sessions to the home directory that is configured in the Active Directory account settings for each FTP user. When a user's object is located in the Active Directory container, the *FTPRoot* and *FTPDir* properties are extracted to provide the full path of the user's home directory. If the FTP service can successfully access the path, the user is positioned within their home directory, which represents their FTP root location. The user sees only their FTP root location and is restricted from navigating higher up the directory tree. The user is denied access if either the *FTPRoot* or *FTPDir* property do not exist, or, if these two together do not form a valid and accessible path. **Note** This mode requires an Active Directory server that runs using the Windows Server 2003 operating system or a later operating system. A Windows 2000 Active Directory can also be used but requires manual extension of the User Object schema.
Custom	This option specifies that you want to isolate FTP user sessions by using a custom provider. **Important** This option is an advanced feature that can be selected only by modifying the FTP configuration settings in your ApplicationHost.config file.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Set Credentials Dialog Box

Use the **Set Credentials** dialog box to specify the Active Directory credentials for your FTP server to use when it contacts your Active Directory server to retrieve FTP home directory settings.

Element Name	Description
User name	Specifies the user account that the FTP server uses to contact your Active Directory server.
Password	Specifies the password for the user account.
Confirm password	Confirms the password for the user account.

Handler Mappings

Use the **Handler Mappings** feature page to manage a list of handlers that process requests for specific file types.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

1. UI Elements for Handler Mappings
2. Add or Edit Managed Handler Dialog Box
3. Add or Edit Script Map Dialog Box
4. Add or Edit Wildcard Script Map Dialog Box
5. Add or Edit Module Mapping Dialog Box

6. Request Restrictions Dialog Box
7. Edit Feature Permissions Dialog Box

UI Elements for Handler Mappings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the handler mapping.
Path	Displays the file name or file name extension for which the handler processes a response. The value can be either a specific file name, such as WebResource.axd, or a wildcard (*) with a specific extension, such as *.exe.
State	Displays whether the handler is enabled or disabled based on the access policy set in the **Edit Feature Permissions** dialog box and the required access setting that is specified by the handler mapping. If a handler requires a type of access that is not enabled in the access policy at that level, the handler is disabled.
Path Type	Displays one of the following values for the type of path to which the handler is mapped: • **Unspecified**, if the mapping does not apply to a specific path type. • **File**, if the mapping applies to a file. • **Directory**, if the mapping applies to a directory. • **File or Directory**, if the mapping is for either a file or a directory.
Handler	Displays the module, such as IsapiModule, or managed type, such as System.Web.Handlers.TraceHandler, that responds to the request as specified in the mapping.
Entry Type	Displays whether the item is local or inherited.

Element Name	Description
	Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Managed Handler	Opens the **Add Managed Handler** dialog box in which you can add a mapping to a managed code handler.
Add Script Map	Opens the **Add Script Map** dialog box in which you can add a mapping to a script processor.
Add Wildcard Script Map	Opens the **Add Wildcard Script Map** dialog box in which you can add a handler mapping to handle every file name extension.
Add Module Mapping	Opens the **Add Module Mapping** dialog box in which you can add a mapping to a module.
Edit	Opens an edit dialog box from which you can edit the selected handler mapping.
Rename	Enables the **Name** field so you can change the friendly name of the selected handler mapping. **Note** This action is available only at the server level.
Lock	Locks the selected handler mapping so it cannot be overridden at lower levels in configuration. **Note** This action is available only at the server level.
Unlock	Unlocks the selected handler mapping so it can be overridden at lower levels in configuration. **Note** This action is available only at the

Element Name	Description
	server level.
Remove	Removes the item that is selected from the list on the feature page.
Edit Feature Permissions	Opens the **Edit Feature Permissions** dialog box in which you can configure permissions that apply to the entire **Handler Mappings** feature.
Revert to Inherited	Reverts the feature to inherit settings from the parent configuration. This option deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other actions in the **Actions** pane do not appear until you select the unordered list format.
Move Up	Moves up a selected item in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits

Element Name	Description
	settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.

Add or Edit Managed Handler Dialog Box

Use the **Add Managed Handler** and **Edit Managed Handler** dialog boxes to add new or edit existing managed code handlers on the Web server. *Managed handlers* are written in managed code and respond to specific requests on the web server. For example, the **PageHandlerFactory-Integrated** handler mapping specifies that the `System.Web.UI.PageHandlerFactory` handler process requests for .aspx files.

Note

By default, when you add a managed handler, the handler runs only in application pools configured to use Integrated mode. However, if you want to run the handler in application pools configured to use ISAPI mode, add the handler to the `<httphandlers>` section in `System.web` from a tool such as Appcmd.exe.

Element Name	Description
Request path	Type a file name extension or file name with an extension for which you want the handler to process requests. For example, if you want the handler to process all requests made for files with the extension .abc, enter *.abc; or if you want the handler to run all requests made for a specific file, type the file name and its extension, such as Myfile.abc.
Type	Type the type name of the managed handler, such as System.Web.Handlers. When you add a managed handler, make sure

Element Name	Description
	that the assembly to which the class belongs is registered in the system.web/compilation section of your Web.config file. Otherwise type the assembly qualified class name of the managed handler. If the class is defined in the App_Code or bin directory you can specify just the class name.
Name	Type a descriptive friendly name for the handler mapping.
Request Restrictions	Opens the **Request Restrictions** dialog box from which you can specify additional, optional restrictions for the mapping.

Add or Edit Script Map Dialog Box

Use the **Add Script Map** and **Edit Script Map** dialog boxes to add new or edit existing native handlers on the Web server. *Native handlers*, also known as script engines, are native code .exe or .dll files that respond to specific requests.

Note
Script maps provide backward compatibility with earlier versions of IIS. Executables must be written to the CGI specification, and dynamic-link libraries must support the ISAPI extension interfaces.

If you map a type of request to an .exe file (the script engine), CgiModule loads the associated executable when a request enters the server and it matches the handler mapping. If you map a type of request to a .dll file (the script engine), IsapiModule loads the DLL when a request enters the server and it matches the handler mapping. For example, IIS 8 includes a handler mapping for Active Server Pages (ASP). All requests for .asp files are processed by the Asp.dll script engine, which IsapiModule loads because Asp.dll is an ISAPI extension.

Note
If you configure script maps for executables, make sure that CgiModule is registered and enabled on the **Modules** page. Likewise, if you configure script maps for DLLs, make sure that IsapiModule is registered and enabled on the **Modules** page.

After you click **OK** to add or save the handler mapping, the executable must be added to the **ISAPI and CGI Restrictions** list to enable it to run.

Element Name	Description
Request path	Type a file name extension or file name with an

Element Name	Description
	extension for which you want the handler to process requests. For example, if you want the handler to run all requests made for ASP files, enter *.asp; or if you want the handler to run all requests made for a specific file, enter the file name and its extension, such as Myfile.asp. **Note** If you want to configure a wildcard script map, use the **Add Wildcard Script Map** action on the **Handler Mappings** page.
Executable	Type the path to the CGI program or to the ISAPI extension in the physical file system. **Note** If you specify the path of a .exe file that has spaces in the path, enclose the path in double quote marks ("). You can also specify command-line parameters with .exe files. These instructions do not apply to .dll files.
...	Locate the .exe or .dll file in the physical file system.
Name	Type a descriptive name for the handler mapping.
Request Restrictions	Opens the **Request Restrictions** dialog box from which you can specify additional, optional restrictions for the mapping.

Add or Edit Wildcard Script Map Dialog Box

Use the **Add Wildcard Script Map** and **Edit Wildcard Script Map** dialog boxes to add or edit a wildcard mapping. The ISAPI extension that you configure for the mapping handles every file name extension, and intercepts every request before the requested page is sent to its mapped handler.

After you click **OK** to add or save the handler mapping, the executable must be added to the **ISAPI and CGI Restrictions** list to enable it to run.

Element Name	Description
Request path	Displays an asterisk (*). **Note** If you want to configure a script map for a specific file name extension or file name with extension, use the **Add Script Map** action on the **Handler Mappings** page.
Executable	Type the path of the ISAPI extension in the physical file system.
…	Locate the .exe or .dll file in the physical file system.
Name	Type a descriptive name for the handler mapping.

Add or Edit Module Mapping Dialog Box

Use the **Add Module Mapping** and **Edit Module Mapping** dialog boxes to add new or edit existing module mappings on the Web server. You can map a specific file or file name extension to a native module on the Web server, so that when a user requests the file or a file that has the specified extension, the module processes the request. For example, IIS 8 maps all requests for .asp files to IsapiModule, so that module is called when a user requests a file that has the .asp file name extension.

Element Name	Description
Request path	Type a file name extension or file name with an extension for which you want the handler to process requests. For example, if you want the handler to process all requests made for files with the extension .abc, enter *.abc; or if you want the handler to run all requests made for a specific file, enter the file name and its extension, such as Myfile.abc.
Module	Select the module that you want to process requests for resources specified in the **Request path** box.
Executable (optional)	Optionally, type a path to the .dll or .exe file when the module you selected in the **Module**

Element Name	Description
	box requires an executable to run, such as the FastCGI module. **Note** After you click **OK** to add or save the handler mapping, the executable must be added to the **ISAPI and CGI Restrictions** list to enable it to run. If you specify a .exe file, you can also type optional command-line parameters after the path. Command-line parameters are not supported for .dll files. **Note** If you specify a path for a .exe file, and the path contains spaces, enclose the path with double quotation marks ("). For example, type "c:\files\handler files\handler.exe". Quotation marks are not required for .dll files.
Name	Type a descriptive friendly name for the handler mapping.
Request Restrictions	Opens the **Request Restrictions** dialog box from which you can specify additional, optional restrictions for the mapping.

Request Restrictions Dialog Box

Use the **Request Restrictions** dialog box to specify optional restrictions, such as the requested resource type or HTTP verb, for which you want a handler to process requests. Configure restrictions only if you have a specific need to limit the conditions under which a handler processes requests. When you do not configure restrictions, handlers process requests regardless of the resource type requested and the HTTP verbs specified.

Element Name	Description
Mapping	Configure options on this tab if you want to restrict the handler mapping to requests for specific resource types, such as for a file or a folder. If you do not configure options on this

Element Name	Description
	tab, the handler processes requests regardless of the resource type specified in the request.
Invoke handler only if request is mapped to	Select this check box on the **Mapping** tab if you want the handler mapping to apply only to requests for a specific resource type, such as a request for a file or a folder. If you select this check box, you must also specify the resource type by selecting an option under this check box.
File	Select this option if you want the handler to process requests for a specific file name or file name extension. For example, the ASPClassic handler processes requests for requests made for .asp files. This option is available only if you select the **Invoke handler only if request is mapped to** check box.
Folder	Select this option if you want the handler to process requests for a folder. Use this restriction if you have a folder with content that you want a specific handler to process. This option is available only if you select the **Invoke handler only if request is mapped to** check box.
File or folder	Select this option if you want the handler to process requests for files and folders. This option is available only if you select the **Invoke handler only if request is mapped to** check box.
Verbs	Configure options on this tab if you want to restrict a handler mapping to specific HTTP verbs that clients provide in HTTP headers when they make requests.
All verbs	Select this option if you want the handler to process requests regardless of the verbs that the client specifies. This option is selected by default.
One of the following verbs	Select this option if you want the handler to process requests only when a client provides a

Element Name	Description
	specific HTTP verb or specific HTTP verbs in a request. In the corresponding text box, type the HTTP verb or verbs to which you want the handler mapping to apply. For example, type GET if you want the handler to process requests when a client uses the GET verb.
Access	Select one of the options on this tab to configure the access that this handler requires to run in a virtual directory. This setting (together with the feature's access policy that is set on the **Edit Feature Permissions** dialog box) determines whether a handler can run. **Important** Make sure that the required access setting for a handler is correct or else the handler might run unintentionally. For example, if you change the handler's required access from **Execute** to **Read** for the ISAPI dll handler, ISAPI extensions will be able to run even if only **Read** is enabled in the feature's access policy. Select one of the following options: • **None**: The handler runs even when no access policy options are enabled. • **Read**: The handler will run when **Read** is enabled in the access policy. • **Write**: The handler will run when **Write** is enabled in the access policy. • **Script**: The handler will run when **Scripts** is enabled in the access policy. This selection is the default. • **Execute**: The handler will run when **Execute** is enabled in the access policy.

Edit Feature Permissions Dialog Box

Use the **Edit Feature Permissions** dialog box to configure the access policy that specifies the type of feature rights allowed for handlers at the Web server, site, application, directory, or file

levels in IIS. The feature rights that you can enable or disable in the access policy are: read, scripts, and execute. The access policy, together with a handler's required access setting, determines whether a handler can run. If a handler requires a type of feature right that is not enabled in the access policy, the handler is disabled and all requests that are processed by the handler (based on the handler mapping) fail, unless there is another handler that can process the request.

When you select an option in the **Edit Feature Permissions** dialog box, the **State** column on the **Handler Mappings** page displays **Enabled** for the handlers that the selection enables. Similarly, when you clear a selection in the **Edit Feature Permissions** dialog box, the **State** column on the **Handler Mappings** page displays **Disabled** for the handlers that the selection enables. You can preview the handlers that are enabled or disabled by viewing the **Handler Mappings** page, and then clicking **OK** to dismiss the **Edit Feature Permissions** dialog box. If you click **Cancel** instead of **OK**, the changes that you make in the dialog box will not be saved.

Note

To configure the required access setting for a handler, you can edit the handler mapping and click **Request Restrictions** to configure the setting on the **Access** tab.

For example, you might enable **Read** and **Scripts** at the web server level, but decide to disable **Scripts** for a specific site that serves only static content. This setting prevents the server from running scripts for that site if a user adds a handler mapping for a script or executable at the site level.

Element Name	Description
Read	Select the **Read** check box to enable handlers that require read access to a virtual directory or clear the **Read** check box to disable handlers that require read access to a virtual directory. If you want to serve static content, or want to configure default documents and directory browsing, enable read in an access policy. By default, read permissions are enabled.
Scripts	Select the **Scripts** check box to enable handlers that require script rights to a virtual directory or clear the **Scripts** check box to disable handlers that require script rights in a virtual directory.
Execute	Select the **Execute** check box to enable handlers that require execute rights in a virtual directory or clear the **Execute** check box to disable handlers that require execute rights in a virtual directory. The **Execute** check box is

Element Name	Description
	enabled only when the **Scripts** check box is selected. If you want to enable executable files, such as .exe, .dll, and .com files, to run in addition to scripts, enable execute in an access policy only. **Important** For security and performance reasons, enable executable rights only for programs that you have tested and that your applications require.

HTTP Redirect

Use the **HTTP Redirect** feature page to enable redirection and to configure how incoming requests are redirected to a new destination.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

UI Elements for HTTP Redirect

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Redirect requests to this destination	Select this option to enable redirection and to specify the URL to which to redirect requests. Requests can be redirected to a virtual path, such as http://www.contoso.com/sales, or to a file, such as http://www.contoso.com/sales/default.aspx.
Redirect all requests to exact destination	Select this option when you want to redirect

Element Name	Description
(instead of relative to destination)	clients to the exact URL as specified in the **Redirect requests to this destination** box. When this option is not selected, the destination is relative to the value specified in **Redirect requests to this destination**. For example, if you configure the redirection destination to be http://www.contoso.com/sales, and an incoming request is http://www.contoso.com/marketing/default.aspx, IIS will redirect the request to http://www.contoso.com/sales/default.aspx.
Only redirect requests to content in this directory (not subdirectories)	Select this option when you want to redirect clients to content in only the directory that is specified in the **Redirect requests to this destination box**. When this option is not selected, requests are redirected to the location in the **Redirect requests to this destination** box, and to any subdirectories under that location. For example, if you configure the redirection destination to be http://www.contoso.com/sales, requests can be served from directories under /sales when the **Redirect all requests to exact destination (instead of relative to destination)** option is not also selected.
Status code	Select one of the following options to specify the redirect status code that is sent to clients: • **Found (302)**: Tells the web client to issue a new request to the location. This option is the default. • **Permanent (301)**: Tells the web client that the location for the requested resource has permanently changed. • **Temporary (307)**: Prevents a web browser from losing data when the browser issues an HTTP POST request.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

HTTP Response Headers

Use the **HTTP Response Headers** feature page to manage a list of name and value pairs that contain information about a requested page, and to configure common HTTP headers.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for HTTP Response Headers
- Add or Edit Custom HTTP Response Header Dialog Box
- Set Common HTTP Response Headers Dialog Box

UI Elements for HTTP Response Headers

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name part of an HTTP response header name-value pair.
Value	Displays the value part of an HTTP response header name-value pair.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current

Element Name	Description
	configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Custom HTTP Response Header** dialog box in which you can add a header.
Set Common Headers	Opens the **Set Common HTTP Response Headers** dialog box in which you can configure the HTTP keep-alive and Web content expiration headers.
Edit	Opens the **Edit Custom HTTP Response Header** dialog box in which you can edit a header. This action is available only when an item is selected from the list on the feature page.
Remove	Removes the item that is selected from the list on the feature page.

Add or Edit Custom HTTP Response Header Dialog Box

Use the **Add Custom HTTP Response Header** and **Edit Custom HTTP Response Header** dialog boxes to add or edit name and value pairs that contain information about a requested page. For example, you can create a custom header named “author” that contains the name of the content author for a site.

Element Name	Description
Name	Type the name part of an HTTP response header name-value pair.
Value	Type the value part of an HTTP response header name-value pair.

Set Common HTTP Response Headers Dialog Box

Use the **Set Common HTTP Response Headers** dialog box to configure headers that enable HTTP keep-alive and specify when to expire Web content.

Element Name	Description
Enable HTTP keep-alive	Select this option to keep a client/server connection open across multiple requests to the server. An open connection improves performance when a client makes multiple requests for Web page content because the server can more quickly return the content for each request. **Note** You must use HTTP keep-alive when you use integrated security or connection-based authentication services such as Integrated Windows authentication.
Expire Web content	Select this option to configure the expiration of Web content. A Web server returns this header value to a client browser in an HTML file header. The browser compares that value with the current date to determine whether to display a page that is cached on the client computer or to request an updated page from the server. Then select one of the following options to specify when to expire the content: • **Immediately** - expires the content immediately after it is delivered. This setting is best for content that contains sensitive information that you do not want to be cached or that is updated very frequently. • **After -** sets the amount of time after which the content will be expired. This setting is best for content that is updated regularly, such as on a daily or weekly basis. Type a value in the corresponding box and select one of the following values from the list: **Second(s)**, **Minute(s)**, **Hour(s)**, **Day(s)**. • **On (in Coordinated Universal Time**

Element Name	Description
	(UTC)) - sets an exact date and time when the content will expire. This setting is best for content that is not expected to change frequently.

IIS Manager

The IIS Manager feature represents the UI components available at the top level of IIS Manager. These components include the **Start** page, the **Home** page, settings pages and dialog boxes.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS

In this document

- Start Page
- Home Page
- Content View Page
- Defaults Dialog Boxes
- Advanced Settings Dialog Boxes
- New Feature Available Dialog Boxes
- Change .NET Framework Version Dialog Box
- Create New Connection
  - Specify Server Connection Details Wizard Page
  - Specify Site Connection Details Wizard Page
  - Specify Application Connection Details Wizard Page
  - Provider Credentials Wizard Page
  - Specify a Connection Name Wizard Page
  - Certificate Name Mismatch Help Identifier Dialog Box
  - Certificate Verification Help Identifier Dialog Box

Start Page

Use the **Start** page to reopen recent connections and create new connections to web servers, sites, and applications by using IIS Manager. You can also click links to online resources and review news about IIS.

Element Name	Description
Recent connections	Displays a list of your recent connections to web servers, sites, and applications. Double-click a connection in the list to reopen the connection. You might need to retype your account credentials to connect.
Connection tasks	Displays the following connection tasks: • **Connect to localhost**: Connects to the local computer by using the credentials of the Windows user who is currently logged on. • **Connect to a server**: Opens a wizard from which you can connect to a server by using a Windows administrator account. • **Connect to a site**: Opens a wizard from which you can connect to a site by using a Windows or IIS Manager user account that has permission to connect to the site. • **Connect to an application**: Opens a wizard from which you can connect to an application by using a Windows or IIS Manager user account that has permission to connect to the application.
Online resources	Displays a list of links to online resources about IIS.
IIS News	Displays recent IIS news. Click **Enable IIS News** in the **IIS News** bar to enable this feature.

Home Page

Use the **Home** page to view and access the list of features that are available for configuration in IIS Manager.

Use the **Group by** drop-down list to group similar features, and use the **View** drop-down list to display the features in different formats, such as icons or details. If you select **Details** from the **View** list, you can sort the features by clicking one of the columns on the **Home** page.

The **Home** page is displayed when both of the following are true:

- You have selected **Features View**.

- You have clicked a Web server, site, application, virtual directory, or physical directory in the tree in the **Connections** pane, or you have selected a file in **Content View** and then clicked **Switch to Features View**.

Content View Page

Use the **Content View** page to view the contents of an item selected in the tree in the **Connections** pane. For example, select a site and click **Content View** to see the list of any applications, virtual directories, physical directories, and files in that site.

You can right-click an object in the **Content** list, and click **Switch to Features View** to go to the object's home page. From the **Home** page, you can configure features for the object, such as configuring authentication settings for an application.

Note

IIS Manager does not display files in the tree in the **Connections** pane. To configure features for individual files, right-click the file in **Content View** and then click **Switch to Features View**.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items. Search the list by typing search criteria in the **Filter** box, expand the **Go** list to select the field on which to search, and then click **Go** or press ENTER. Click **Show All** to remove the filter criteria.

Feature Page Elements

Element Name	Description
Name	Displays the name of each object.
Type	Displays the object type, such as application or virtual directory.

Actions Pane Elements

The actions available in the **Actions** pane depend on the type of object that you select on the **Content** page. Additionally, when you are connected to a site or an application, the only actions available are **Switch to Features View**, **Browse**, and **Refresh**.

Element Name	Description
Advanced Settings	Opens the **Advanced Settings** dialog box from which you can configure advanced settings for the selected object.
Add Application	Opens the **Add Application** dialog box from

Element Name	Description
	which you can add an application.
Add Application Pool	Opens the **Add Application Pool** dialog box from which you can add an application pool to the Web server.
Add Virtual Directory	Opens the **Add Virtual Directory** dialog box from which you can add a virtual directory.
Add Web Site	Opens the **Add Web Site** dialog box from which you can add a website that uses the HTTP or HTTPS protocols for communication.
Browse	Opens the selected object in an internet browser.
Change .NET Framework Version	Opens the **Change .NET Framework Version** dialog box, which enables you to select the version of the .NET Framework to use by default.
Convert to Application	Converts the directory into an application.
Edit Bindings	Opens the **Site Bindings** dialog box from which you can add, edit, and remove site bindings for the selected site.
Explore	Opens the physical directory (in Windows Explorer) that is mapped to the selected object.
Edit Permissions	Opens the Windows Properties dialog box for the physical directory that is mapped to the selected object.
Refresh	Refreshes the **Content** page.
Remove	Removes the item that is selected from the list on the page.
Remove Connection	Removes the selected connection.
Rename	Enables the **Name** field of the selected object so that you can rename it.
Restart	Stops and restarts the selected site. Restarting a site causes the site to be temporarily unavailable until the restart is complete.
Start	Starts the selected site.

Element Name	Description
Stop	Stops the selected site. Stopping a site causes the site to be unavailable until it is started.
Switch to Features View	Opens the **Home** page in **Features View** for the selected object.
View Applications	Opens the **Applications** feature page from which you can view the applications that belong to the selected site.
View Virtual Directories	Opens the **Virtual Directories** feature page from which you can view the virtual directories that belong to the selected site's root application.

Defaults Dialog Boxes

Use the **Defaults** dialog boxes to edit the default values of properties that have default values, such as recycling settings for application pools, limits for sites, and the application pool in which applications run. The **Defaults** dialog boxes do not contain properties that are frequently unique, such as site bindings or the physical path of site, application, and virtual directory content.

Each of the following objects in IIS has a **Default** dialog box from which you can edit default values for all objects of that type:

- Application pool
- Site
- Application
- Virtual directories

When you change the default value of a property, all objects of that type receive the new value when:

- You create an object of that type.
- An existing object does not already have a specific value set that differs from the default value for the property.

> **Note**
> If you must change the value for a specific site or application, use the **Advanced Settings** dialog box.

For information about a specific property, select the property in the **Defaults** dialog box and view the information that is displayed in the property description pane underneath the property grid. You can make the property description pane larger by clicking the bar between the property grid and the property description pane, and then dragging upward. You can also resize the entire **Defaults** dialog box to see more properties at the same time.

Advanced Settings Dialog Boxes

Use the **Advanced Settings** dialog boxes to edit the properties for each of the following objects in IIS:

- Application pool
- Site
- Application
- Virtual directory

The **Advanced Settings** dialog box is useful when you want to edit a specific property or group of properties for an object without using an edit wizard or a dialog box, or when a property is not displayed in a wizard or dialog box.

Note
If you want to change a default value for all application pools, sites, applications, or virtual directories, you can do so on the appropriate **Defaults** dialog box for any of these objects.

For information about a specific property, select the property in the **Advanced Settings** dialog box and view the information that displays in the property description pane underneath the property grid. You can make the property description pane larger by clicking the bar between the property grid and the property description pane, and then dragging upward. You can also resize the entire **Advanced Settings** dialog box to see more properties at the same time.

New Feature Available Dialog Boxes

Use the **New Feature Available** dialog box to download and install code to support any new IIS Manager features that are available on the server to which you are connecting. After you install a feature, you can use it in IIS Manager on subsequent connections to the same server or to any servers that have that same feature unless a server administrator removes delegation for the feature at a later time.

Element Name	Description
Assembly	Displays the name of each .dll file that is available to download and install on your computer. Select the check box next to each assembly that you want to download and install, and then click **OK**. **Note** When the publisher of the code cannot be verified as having a valid digital signature, you are prompted with a security warning before the code is

Element Name	Description
	downloaded to your computer. Click **Run** if you trust the publisher and want to continue with the download, or click **Don't Run** if you want to cancel the download.
Features	Displays the feature or features that are available with each .dll file.
Size	Displays the size of the .dll file.
Version	Displays the version of the .dll file.

Change .NET Framework Version Dialog Box

Use **Change .NET Framework Version** to set the default .NET Framework version for the Web server.

Element Name	Description
.NET Framework Version	Lists the .NET Framework versions available on the server. Select the version and click **OK**.

Create New Connection

Use the **Create New Connection** icon to select between creating a connection to a server, a site, or a web application.

Element Name	Description
Connect to a Server	Launches the **Connect to Server** wizard, which displays the Specify Server Connection Details Wizard Page page and guides you through creating a connection to a Web server.
Connect to a Site	Launches the **Connect to Site** wizard, which displays the Specify Site Connection Details Wizard Page page and guides you through creating a connection to a Web server.
Connect to an Application	Launches the **Connect to Application** wizard, which displays the Specify Application Connection Details Wizard Page page and

Element Name	Description
	guides you through creating a connection to a Web server.

Specify Server Connection Details Wizard Page

Use the **Specify Server Connection Details** wizard page to specify the name of the server to which you want to connect in IIS Manager.

If you receive an error when you try to connect to the server, check for the following possible issues:

- You mistyped a value, such as the server name or user credentials.
- The Web server to which you are connecting is not running IIS 8. You cannot use IIS Manager in IIS 8 to connect to a Web server that runs a previous version of IIS.
- The management service is not started on the Web server.
- A Firewall rule is preventing the connection.

Element Name	Description
Server name	Type the name of the server to which you want to connect. If the server is using a port other than 8172 (the default port configured for the management service), type the server name and the port number separated by a colon (:). For example, type **contoso:8080** to connect to a server named **contoso** with the management service configured to accept connections on port 8080. Alternately, you can type the IPv4 or IPv6 address in block braces (**[]**).

Specify Site Connection Details Wizard Page

Use the **Specify Site Connection Details** wizard page to specify the site to which you want to connect in IIS Manager. The server administrator should give you the information that you must have to connect to the site.

If you receive an error when you try to connect to the server, check for the following possible issues:

- You mistyped a value, such as the server name or user credentials.
- The Web server to which you are connecting is not running IIS 8. You cannot use IIS Manager in IIS 8 to connect to a Web server that runs a previous version of IIS.
- The management service is not started on the Web server.

- A Firewall rule is preventing the connection.

Element Name	Description
Server name	Type the name of the server that hosts the site to which you want to connect. If the server is using a port other than 8172 (the default port configured for the management service), type the server name and the port number separated by a colon (:). For example, type **contoso:8080** to connect to a server named contoso with the management service configured to accept connections on port 8080. Alternately, you can type the IPv4 or IPv6 address in block braces (**[]**).
Site name	Type the friendly name of the site to which you want to connect.

Specify Application Connection Details Wizard Page

Use the **Specify Site Connection Details** wizard page to specify the web application to which you want to connect in IIS Manager. The server administrator should give you the information that you must have to connect to the application.

If you receive an error when you try to connect to the server, check for the following possible issues:

- You mistyped a value, such as the server name or user credentials.
- The Web server to which you are connecting is not running IIS 8. You cannot use IIS Manager in IIS 8 to connect to a Web server that runs a previous version of IIS.
- The management service is not started on the Web server.
- A Firewall rule is preventing the connection.

Element Name	Description
Server name	Type the name of the server that hosts the site to which you want to connect. If the server is using a port other than 8172 (the default port configured for the management service), type the server name and the port number separated by a colon (:). For example, type **contoso:8080** to connect to a server named contoso with the management service configured to accept connections on port 8080.

	Alternately, you can type the IPv4 or IPv6 address in block braces (**[]**).
Site name	Type the friendly name of the site to which you want to connect.

Provider Credentials Wizard Page

Use the **Provide Credentials** wizard page to specify credentials to connect to a server, site, or application by using IIS Manager. The credentials specified on this page can be either Windows user accounts or IIS Manager user accounts.

If you receive an error when you try to connect to the server, check for the following possible issues:

- You mistyped a value, such as the server name or user credentials.
- The Web server to which you are connecting is not running IIS 8. You cannot use IIS Manager in IIS 8 to connect to a Web server that runs a previous version of IIS.
- The management service is not started on the Web server.
- A Firewall rule is preventing the connection.

Element Name	Description
User name	Type the user name of an account that is authorized to connect to the server, site, or application. **Note** Only server administrators can connect to a server by using IIS Manager.
Password	Type the password that is associated with the user name.

Specify a Connection Name Wizard Page

Use the **Specify a Connection Name** wizard page to specify a friendly name for your connection in IIS Manager.

Element Name	Description
Connection name	Type a friendly name for the connection. This connection name appears in the tree in IIS Manager.

Certificate Name Mismatch Help Identifier Dialog Box

This error message indicates that the server name that you entered does not match the name that appears on the server's certificate.

You might have caused this error by entering the secure server's IP address when the DNS or NETBIOS name appeared on the server's certificate. If this error occurs, it can appear as if you are being maliciously redirected to a server other than the one to which you had intended to connect.

To avoid seeing this message in the future, connect to the server by entering the address that appears on the secure server's certificate.

Certificate Verification Help Identifier Dialog Box

This message indicates that there is a problem with the certificate issued to the server to which you are connecting. The certificate might have expired or been revoked. If the certificate has expire or is revoked, resolve the problem before you can connect to the server.

This message can also appear if you are connecting to the server by using a name that does not appear on the server's certificate. If the name is not on the server's certificate, it can seem that you are being maliciously redirected to a server other than the one to which you had intended to connect. You are given the option to connect, but you should view the certificate to determine whether to connect to the server.

This message can also appear because of some combination of the errors listed above. If this happens, use the **View details** link to investigate the details of the problem and determine the best course of action.

IIS Manager Permissions

Use the **IIS Manager Permissions** feature page to manage IIS Manager users, Windows users, and members of Windows groups that are allowed to connect to a site or an application. Users can configure delegated features in any sites and applications to which you grant them permission from the **IIS Manager Permissions** page.

This feature is available only for server connections. When you open the feature at the server level in IIS Manager, you can view the users who are granted permission to all sites and applications on the Web server, and you can select a user to remove the user's permission to a site or an application. To grant a user permission to a site or an application, select the site or application and then open the **IIS Manager Permissions** feature to configure the users who are allowed to connect to that site or application.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS

- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS

In this document

- UI Elements for IIS Manager Permissions
- Allow User Dialog Box
- Users Dialog Box

UI Elements for IIS Manager Permissions

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the user name that is used to connect to sites or applications in IIS Manager.
Path	Displays the path of the object that the administrator can manage. For example, the path might be a site name, such as **Default Web Site**, or the path of an application in the default website, such as **Default Web Site/calendar**.
Level	Displays the level of the administrator. The possible values are **Site** or **Application**.
Type	Displays the type of credentials that the administrator uses. The value can be: • **Windows Group**: if the user is a member of a Windows group on the server. • **Windows User**: if the user is a Windows user on the server. • **IIS Manager User**: if the user is not a Windows user or a member of any Windows groups on the server. IIS Manager credentials are specific to IIS, and Windows does not recognize nor do other applications on the server.

Actions Pane Elements

Element Name	Description
Allow User	Opens the **Allow User** dialog box in which you can select a user to grant them permission to connect to a site or an application.
Deny User	Removes the selected user from the **IIS Manager Permissions** list.
Show all users	Displays users at the current level and lower levels.
Show only site users or **Show only application users**	Displays users at only the current level.

Allow User Dialog Box

Use the **Allow User** dialog box to allow IIS Manager users, Windows users, or members of Windows groups to connect to a site or an application. Users can configure delegated features in any sites and applications to which you grant them permission.

Element Name	Description
Windows	Select this option when you want to allow a Windows user or members of a Windows group to connect to the site or application. When you select this option, the Windows account must already be created on the Web server or the domain to which the server is joined.
IIS Manager	Select this option when you want to allow an IIS Manager user to connect to the site or application. When you select this option, the IIS Manager user account must already be created at the server level in IIS Manager. **Note** This option is enabled only when you configure the management service to accept connections from both Windows users and IIS Manager users.
Select	If you selected the **Windows** option, click **Select** to open the **Select User or Group**

Element Name	Description
	dialog box and search for a Windows user or group. **Note** To include Windows groups in a search, on the **Select User or Group** dialog box click **Object Types**, select **Groups**, and then click **OK**. Additionally, if your computer is joined to a domain and you want to search for users or groups on the domain, click **Locations**, select a location, and then click **OK**. If you selected the **IIS Manager** option, click **Select** to open the **Users** dialog box and select one or more IIS Manager users.

Users Dialog Box

Use the **Users** dialog box to select an IIS Manager user who you want to allow to connect to a site or an application. Users can configure delegated features in any sites and applications to which you grant them permission.

Element Name	Description
Select user	Select the IIS Manager user name that you want to allow to connect to the site or application.

IIS Manager Users

Use the **IIS Manager Users** feature page to manage user accounts that can be permitted to connect to sites or applications on the Web server. IIS Manager credentials are specific to IIS, and they are not recognized by Windows or by any other applications on the server.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS

- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS

In this document

- UI Elements for IIS Manager Users
- Add User and Change Password Dialog Boxes

UI Elements for IIS Manager Users

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the user name that is used to connect to sites or applications in IIS Manager.
Status	Displays whether the user account is enabled or disabled. Only enabled users can use IIS Manager to connect to the sites and applications to which they are granted permission.

Actions Pane Elements

Element Name	Description
Add User	Opens the **Add User** dialog box in which you can add an IIS Manager user account.
Change Password	Opens the **Change Password** dialog box in which you can change the password of an IIS Manager user.
Enable	Enables the selected user. This action is available only when a disabled item is selected from the list on the feature page.
Disable	Disables the selected user. This action is available only when an enabled item is selected

Element Name	Description
	from the list on the feature page.
Remove	Removes the selected item from the list. This also removes any permissions that were granted to the selected user in sites and applications.

Add User and Change Password Dialog Boxes

Use the **Add User** dialog box to add a user account to IIS Manager. IIS Manager user credentials are valid only in IIS Manager, and these credentials are used to authenticate the user when the user connects to a site or an application by using IIS Manager. Users can configure delegated features in any sites and applications to which a server administrator grants them permission.

Use the **Change Password** dialog box to change the password of the selected IIS Manager user account.

Element Name	Description
User name	Type a user name of your choice in the box. **Note** This box is not enabled when you are changing a user's password.
Password	Type a password in the box. To enhance security, you should type a strong password. Strong passwords have a minimum of six characters and contain a combination of uppercase and lowercase letters, numbers, and symbols.
Confirm password	Retype the password that you typed in the **Password** box.

IP Address and Domain Restrictions

Use the **IP Address and Domain Restrictions** feature page to define and manage rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a domain name or names.

Sort the list by clicking one of the column headings on the feature page, or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for IP Address and Domain Restrictions
- Add Allow or Add Deny Restriction Rule Dialog Boxes
- Edit IP and Domain Restrictions Dialog Box
- Dynamic IP Restriction Settings Dialog Box

UI Elements for IP Address and Domain Restrictions

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Mode	Displays the type of rule. Values are either **Allow** or **Deny**. The **Mode** value indicates whether the rule is designed to allow or deny access to content.
Requester	Displays a specific IP address, range of IP addresses, or domain name that is defined in the **Add Allow Restriction Rule** and **Add Deny Restriction Rule** dialog boxes. You can specifically allow or deny a requester access to content.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Allow Entry	Opens the **Add Allow Restriction Rule** dialog box from which you can define rules that allow access to content for a specific IP address, a range of IP addresses, or a DNS domain name.
Add Deny Entry	Opens the **Add Deny Restriction Rule** dialog box from which you can define rules that allow access to content for a specific IP address, a range of IP addresses, or a DNS domain name.
Remove	Removes the item that is selected from the list on the feature page.
Edit Feature Settings	Opens the **Edit IP and Domain Restrictions Settings** dialog box from which you can configure settings that apply to the entire IP and domain name restrictions feature.
Revert to Inherited	Reverts the feature to inherit settings from the parent configuration. This action deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other actions in the **Actions** pane do not appear until you select the unordered list format.
Move Up	Moves up a selected item in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using

Element Name	Description
	the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.
Edit Dynamic Restriction Settings	Displays the **Dynamic IP Restriction Setting** dialog box from which you can restrict IP addresses that have too many concurrent requests or too many requests for a given time period.

Add Allow or Add Deny Restriction Rule Dialog Boxes

Use either the **Add Allow Restriction Rule** or the **Add Deny Restriction Rule** dialog box to define rules that allow or deny access to content for a specific IP address, a range of IP addresses, or a DNS domain name.

Element Name	Description
Specific IP Address	**Add Allow Restriction Rule** - Type an IP address in the **Specific IP Address** box in the **Add Allow Restriction Rule** dialog box when

Element Name	Description
	you want to allow access to content for a specific IP address. **Add Deny Restriction Rule** - Type an IP Address in the **Specific IP Address** box in the **Add Deny Restriction Rule** dialog box when you want to deny access to content for a specific IP address.
IP address range	**Add Allow Restriction Rule** - Type the lowest value of the range of IP addresses that you have chosen to use in the **IP Address range** box in the **Add Allow Restriction Rule** dialog box. Do this action when you want to allow access to content for a range of IP addresses. Next, enter the subnet mask. **Add Deny Restriction Rule** - Type the lowest value of the range of IP addresses that you have chosen to use in the **IP address range** box in the **Add Deny Restriction Rule** dialog box. Do this action when you want to deny access to content for a range of IP address. Next, enter the subnet mask.
Mask	**Add Allow Restriction Rule** - Type a subnet mask in the **Mask** box in the **Add Allow Restriction Rule** dialog box. Do this action when you want to allow access to content for a range of IP address. When IIS evaluates this subnet mask with the IP address entered in the **IP address range** box, the upper and lower boundaries of an IP address space are defined. **Add Deny Restriction Rule** - Type the subnet mask associated with the range of IP addresses in the **Mask** box in the **Add Deny Restriction Rule** dialog box. Do this action when you want to deny access to content for a range of IP address. When IIS evaluates this subnet mask with the IP address entered in the

Element Name	Description
	IP address range box, the upper and lower boundaries of an IP address space are defined.
Domain name	**Add Allow Restriction Rule** - Type a fully qualified DNS domain name in the **Domain name** box in the **Add Allow Restriction Rule** dialog box when you want to allow access to content for a DNS domain. **Add Deny Restriction Rule** - Type a fully qualified DNS domain name in the **Domain name** box in the **Add Deny Restriction Rule** dialog box when you want to deny access to content for a DNS domain. To see the **Domain name** option, first enable domain name restrictions, using **Edit Feature Settings**.

Edit IP and Domain Restrictions Dialog Box

Use the **Edit IP and Domain Restrictions** dialog box to define access restrictions for unspecified clients or to enable domain name restrictions for all rules.

Element Name	Description
Access for unspecified clients	Defines access restrictions for unspecified clients. This setting defines whether to allow or deny access to clients not specified by any other rule.
Enable domain name restrictions	Enables rules that restrict access by domain name. This rule significantly affects server performance because it requires a DNS lookup for every request.
Enable Proxy Mode	Enables requests to come through a proxy server.
Deny Action Type	Selects the type of action to be taken when a request is denied. The following list shows the available actions: • Unauthorized

Element Name	Description
	• Forbidden • Not Found • Abort

Dynamic IP Restriction Settings Dialog Box

Use the Dynamic IP Restriction Settings dialog box to restrict IP addresses that have too many concurrent requests or too many requests for a given time period.

Element Name	Description
Deny IP Address based on the number of concurrent requests	Denies requests from an IP address when the number of concurrent requests exceeds the specified **Maximum number of concurrent requests**.
Deny IP based on the number of requests over a period of time	Denies requests from an IP address when the number of requests exceeds the specified **Maximum number of requests** for a given **Time Period (in milliseconds)**.
Enable Logging Only Mode	Specifies that if one of the previous rules is exceeded the event is logged and the request is allowed rather than denied.

ISAPI and CGI Restrictions

Use the **ISAPI and CGI Restrictions** feature page to manage a list of restrictions to ISAPI and CGI programs.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for ISAPI and CGI Restrictions
- Add or Edit ISAPI or CGI Restriction Dialog Box

- Edit ISAPI and GCI Restrictions Settings Dialog Box

UI Elements for ISAPI and CGI Restrictions

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Description	Lists the description of the ISAPI or CGI program.
Restriction	Shows whether the specific program is allowed to run.
Path	Shows the path of the .dll or .exe file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add ISAPI or CGI Restriction** dialog box from which you can manage restrictions to CGI (.exe) or ISAPI (.dll) programs.
Deny	Specifies the restriction setting for a CGI (.exe) or ISAPI (.dll) program. This action is available only when an item is selected from the list on the feature page.
Allow	Specifies the restriction setting for a CGI (.exe) or ISAPI (.dll) program. This action is available only when an item is selected from the list on the feature page.
Edit	Opens the **Edit ISAPI or CGI Restriction** dialog box from which you can manage restrictions to CGI (.exe) or ISAPI (.dll) programs. This action is available only when an item is selected from the list on the feature page.

Element Name	Description
Remove	Removes the item that is selected from the list on the feature page.
Edit Feature Settings	Opens the **Edit ISAPI and CGI Restrictions Settings** dialog box from which you can configure settings that apply to the entire ISAPI and CGI restrictions feature.

Add or Edit ISAPI or CGI Restriction Dialog Box

Use the **Add ISAPI or CGI Restriction** and **Edit ISAPI or CGI Restriction** dialog boxes to manage restrictions for CGI (.exe) and ISAPI (.dll) programs.

Element Name	Description
ISAPI or CGI path	Type the physical path to the ISAPI extension or CGI program.
Description	Provide a brief description of the program.
Allow extension path to execute	Select this option to enable this program to run.

Edit ISAPI and GCI Restrictions Settings Dialog Box

Use the **Edit ISAPI and CGI Restrictions Settings** dialog box to enable unspecified CGI or ISAPI modules to run. If you leave these settings unchanged, only the CGI or ISAPI modules specifically listed can run on this server.

Caution

Enabling either **Allow unspecified CGI modules** or **Allow unspecified ISAPI modules** is a security risk. Your server could become susceptible to computer viruses or worms that exploit CGI or ISAPI technologies. To reduce this risk, you should create specific restrictions for all ISAPI extensions or CGI files that must run on your server.

Element Name	Description
Allow unspecified CGI modules	Enables unspecified CGI modules to run on this server.
Allow unspecified ISAPI modules	Enables unspecified ISAPI modules to run on this server.

ISAPI Filters

Use the **ISAPI Filters** feature page to manage a list of .dll files that change or enhance the functionality provided by IIS. For example, an ISAPI filter can control which files are mapped to a URL, modify the response sent by the server, or perform other actions to modify the behavior of the server.

Note
ISAPI filters and extensions are replaced by the new IIS 8 modules API. For more information about the IIS 8 modules API, see Native-Code HTTP Modules.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for ISAPI Filters
- Add or Edit ISAPI Filter Dialog Box

UI Elements for ISAPI Filters

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the ISAPI filter.
Executable	Displays the path of the ISAPI filter .dll file, for example, %SystemDrive%\Windows\Microsoft.NET\Framework\v2.0.50727\aspnet_filter.dll.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add ISAPI Filter** dialog box in which you can add an ISAPI filter.
Edit	Opens the **Edit ISAPI Filter** dialog box in which you can edit the selected ISAPI filter.
Rename	Enables the **Name** field of the selected ISAPI Filter so you can rename it.
Remove	Removes the item that is selected from the list on the feature page. **Note** If you are removing an ISAPI filter at the server level, restart WWW Service (W3SVC) for the changes to take effect.
Revert to Inherited	Reverts the feature to inherit settings from the parent configuration. This setting deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other actions in the **Actions** pane do not appear until you select the unordered list format.
Move Up	Moves up a selected item in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using

Element Name	Description
	the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.

Add or Edit ISAPI Filter Dialog Box

Use the **Add ISAPI Filter** and **Edit ISAPI Filter** dialog boxes to add or edit a .dll file that changes or enhances the functionality of IIS.

Element Name	Description
Filter name	Type the name of the ISAPI filter.
Executable	Type the path of the ISAPI filter .dll file. For example, %*SystemDrive*%\Windows\Microsoft.NET\Framework\v2.0.50727\aspnet_filter.dll.

Logging

Use the **Logging** feature page to configure how IIS logs requests made to the web server and when new log files are created.

Related scenarios

- Configure Logging in IIS

In this document

- UI Elements for Logging
- W3C Logging Fields Dialog Box

UI Elements for Logging

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
One log file per	Specifies whether the web server maintains one log file for the whole server or one log file for each site on the server. When you select **Site**, the settings that you specify on the **Logging** page at the server level are default values for all sites on your server. You can then open the **Logging** page at the site level to configure specific settings for a site.
Format	Specifies the format of the log file that is created when a request is logged. When you configure the web server to have one log file, you can select one of the following log file formats: • **Binary** Configures IIS to use the binary centralized logging format. By using this format, IIS creates one log file for all sites on the web server. Every site writes request hit log information as binary unformatted data to this log file. Because this kind of logging conserves valuable memory and CPU resources, it is appropriate for use in an ISP environment, where web servers can host many sites, or in any high-traffic situation.

Element Name	Description
	Note To extract data from this log file format, you must use a tool, such as LogParser 2.2. • **W3C** Configures IIS to use the centralized W3C log file format to log information about all sites on the server. This format is handled by HTTP.sys, and is a *customizable* ASCII text-based format, which means that you specify the fields that are logged. Specify the fields that are logged on the **W3C Logging Fields** dialog box by clicking **Select Fields** on the **Logging** page. Fields are separated by spaces, and time is recorded in Coordinated Universal Time (UTC). When you configure the web server to have one log file per site, you can select one of the following log file formats: • **IIS** Configures IIS to use the Microsoft IIS log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by commas, and time is recorded as local time. The following list is a list of fields that are logged when you use the IIS log file format: • **Client IP address** • **User name** • **Date** • **Time** • **Service and instance** • **Server name** • **Server IP address** • **Time taken**

Element Name	Description
	• **Client bytes sent** • **Server bytes sent** • **Service status code** (A value of 200 indicates that the request was fulfilled successfully.) • **Windows status code** (A value of 0 indicates that the request was fulfilled successfully.) • **Request type** • **Target of operation** • **Parameters** (the parameters that are passed to a script.) Not all fields will contain data. When a field does not contain data, a hyphen (-) appears as a placeholder. When a field contains a non-printable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format. • **NCSA** Configures IIS to use the National Center for Supercomputing Applications (NCSA) Common log file format to log information about a site. This format is handled by HTTP.sys, and is a *fixed* ASCII text-based format, which means that you cannot customize the fields that are logged. Fields are separated by spaces, and time is recorded as local time with the Coordinated Universal Time (UTC) offset. The following list is a list of fields that are logged when you use the NCSA Common log file format: • **Remote host address** • **Remote log name** (This value is always a hyphen.) • **User name** • **Date, time, and UTC offset** • **Request and protocol version** • **Service status code** (A value of 200 indicates that the request was fulfilled

Element Name	Description
	successfully.) • **Bytes sent** Not all fields will contain data. When a field does not contain data, a hyphen (-) appears as a placeholder. When a field contains a non-printable character, HTTP.sys replaces it with a plus sign (+) to preserve the log file format. • **W3C** Use the W3C Extended logging format to log information about a site. This format is handled by HTTP.sys, which writes a *customizable* format, which means that you specify the fields that are logged. Specify the fields that are logged on the **W3C Logging Fields** dialog box by clicking **Select Fields** on the **Logging** page. Fields are separated by spaces, and time is recorded in Coordinated Universal Time (UTC). In Windows Server 2012 and earlier, HTTP.sys writes to an ASCII text-based file at the location specified by the **Directory** option. In Windows Server 2012 R2 and later, you have the additional option of writing the log information as an Event Tracing for Windows (ETW) event. For more information on ETW, see Event Tracing. • **Custom** Configures IIS to use a custom format for a custom logging module. When you select this option, the **Logging** page becomes disabled, because custom logging cannot be configured in IIS Manager.
Select Fields	Opens the **W3C Logging Fields** dialog box in which you can select which fields to log. This button is available only when the **W3C** log file format is selected.

Element Name	Description
Directory	Specifies the physical path where the log file or files are stored. The default value is %*SystemDrive*%\inetpub\logs\LogFiles.
Encoding	Specifies how the log file or log files are encoded: UTF-8 or ANSI. Select UTF-8 when you want to enable both single-byte and multibyte characters in one string. This encoding enables you to read text-based logs, for example, W3C Extended, IIS, and National Center for Supercomputing Applications (NCSA) Common formats, in a language other than English. Additionally, if your web server serves URLs in a language or dialect other than the one supported by the server's default code page, you should enable UTF-8 encoding for log content. By default, IIS tries to serve URLs in code pages other than the server's default code page. For security reasons, you might want to enable the UTF-8 format to reduce the effect of an attack that could cause UTF-8 URLs not to translate correctly to the default code page.
Log Event Destination	In Windows Server 2012 R2 and later, you can specify where you want the log events written: • **Log file only**: The log events are written to a log file in the location specified by the **Directory** option. • **ETW event only**: The log events are written as ETW events. • **Both log file and ETW event**: The log events are written both to the log file and as ETW events.
Schedule	Configures IIS to create new log files based one of the following values: • **Hourly**: a new log file is created each hour. • **Daily**: a new log file is created each day. • **Weekly**: a new log file is created each week. • **Monthly**: a new log file is created each

Element Name	Description
	month.
Maximum file size (in bytes)	Configures IIS to create a new log file when the file reaches a certain size (in bytes). The minimum file size is 1048576 bytes. If this attribute is set to a value less than 1048576 bytes, the default value is implicitly assumed as 1048576 bytes.
Do not create new log files	Configures IIS to not create new log files. This means that there is a single log file that will continue to grow as information is logged.
Use local time for file naming and rollover	Specifies that log file naming and time for log file rollover uses the local server time. When this is not selected, Coordinated Universal Time (UTC) is used. **Note** Regardless of this setting, timestamps in the actual log file will use the time format for the log format that you select from the **Format** list. For example, NCSA and W3C log file formats use UTC time format for timestamps.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Disable	Disables the feature.
View Log Files	Opens the log file directory.

W3C Logging Fields Dialog Box

Use the **W3C Logging Fields** dialog box to select the information that you want to include in log files. This dialog box is available only when you select **W3C** from the **Format** drop-down list on the **Logging** feature page.

Element Name	Description
Date (date)	Logs the date on which the request occurred. Selected by default.
Time (time)	Logs the time, in Coordinated Universal Time (UTC), at which the request occurred. Selected by default.
Client IP Address (c-ip)	Logs the IP address of the client that made the request. Selected by default.
User Name (cs-username)	Logs the name of the authenticated user who accessed the server. Anonymous users are indicated by a hyphen. Selected by default.
Service Name (s-sitename)	Logs the Internet service name and instance number that was running on the client at the time that the incident was logged.
Server Name (s-computername)	Logs the name of the server on which the log file entry was generated.
Server IP Address (s-ip)	Logs the IP address of the server on which the log file entry was generated. Selected by default.
Server Port (s-port)	Logs the server port number that is configured for the service. Selected by default.
Method (cs-method)	Logs the HTTP method, such as **GET**, that is used in the request. Selected by default.
URI Stem (cs-uri-stem)	Logs the Uniform Resource Identifier (URI) that is the target of the action, such as Default.htm. Selected by default.
URI Query (cs-uri-query)	Logs the query, if any, that the client was trying to perform. A URI query is necessary only for dynamic pages. Selected by default.
Protocol Status (sc-status)	Logs the HTTP status code. Selected by default.

Element Name	Description
Protocol Substatus (sc-substatus)	Logs the HTTP substatus code. Selected by default.
Win32 Status (sc-win32-status)	Logs the Windows status code. Selected by default.
Bytes Sent (sc-bytes)	Logs the number of bytes that the server sent.
Bytes Received (cs-bytes)	Logs the number of bytes that the server received.
Time Taken (time-taken)	Logs the length of time that the action took in milliseconds. Selected by default.
Protocol Version (cs-version)	Logs the protocol version, HTTP or FTP, that the client used.
Host (cs-host)	Logs the host header name, if any. **Note** The host name that you configure for your website might appear differently in the log files because HTTP.sys logs host names by using the Punycode-encoded format.
User Agent (cs(User-Agent))	Logs the browser from which the request came. Selected by default.
Cookie (cs(Cookie))	Logs the content of the cookie sent or received, if any.
Referer (cs(Referer))	Logs the site that the user last visited. This site provided a link to the current site.

Enhanced Logging with Custom Fields

In Windows Server 2012 R2 and later, you can perform enhanced logging by adding custom fields to the event log in addition to the standard fields listed above. This dialog box is available only when you select **W3C** from the **Format** drop-down list on the **Logging** feature page. Using the **Log Event Destination** option, you can specify that you want the log events to be sent to a log file, to Event Tracing for Windows (ETW), or both.

To add a custom field, in the **W3C Logging Fields** dialog box, click **Add Field...**, specify a **Field Name**, then select the **Source Type** - you can select **Request Header**, **Response Header**, or **Server Variable**. Depending on the Source Type you select, the **Source** field is updated to list the standard fields available. You can select one of the fields given, or type a custom source

string. For more information on request and response headers, see HTTP protocol reference information. For more information on server variables, see IIS Server Variables.

Note
Enhanced logging cannot log a server variable with a name that contains lower-case characters. To include a server variable in the event log, ensure that its name consists of all upper-case characters.

Note that if you add custom fields to the event log, the total size of the logged fields is limited to 64K bytes. If the total size of the custom fields that you define is greater, then the logged content is truncated to 64K bytes.

Note
When enhanced logging is enabled, Kernel caching for request handling is disabled by IIS.

Machine Key

Use the **Machine Key** feature page to configure hashing and encryption settings used for application services, such as view state, Forms authentication, membership and roles, and anonymous identification. Machine keys are also used to verify out-of-process session state identification.

Note
If you deploy your application in a web farm, make sure that the configuration files on each server in the web farm have the same value for the validation key and decryption keys, which are used for hashing and decryption respectively. Otherwise, you cannot guarantee which server handles successive requests.

Related scenarios

- Build an ASP.NET Website on IIS
- Build a Web Farm with IIS Servers

UI Elements for Machine Key

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Validation method	Select one of the following options to specify the validation method the machine key uses:

Element Name	Description
	• **AES** - Advanced Encryption Standard (AES) is relatively easy to implement and requires little memory. AES has a key size of 128, 192, or 256 bits. This method uses the same private key to encrypt and decrypt data, whereas a public-key method must use a pair of keys. • **MD5** - Message Digest 5 (MD5) is used for digital signing of applications, for example, mail messages. This method produces a 128-bit message digest, which is a compressed form of the original data. MD5 can provide some protection against computer viruses and programs that mimic harmless applications but are destructive. • **SHA1** – This method is the default setting. SHA1 is considered to be more secure than MD5 because it produces a 160-bit message digest. Use SHA1 encryption whenever possible. • **TripleDES** - Triple Data Encryption Standard (TripleDES) is a minor variation of Data Encryption Standard (DES). It is three times slower than regular DES but can be more secure because it has a key size of 192 bits. If performance is not your primary consideration, consider using TripleDES. • **HMACSHA256** - Hash-based Message Authentication Code (HMAC) mixes a secret key with the message data, hashes the result with the hash function, mixes that hash value with the secret key again, and then applies the hash function a second time. The output hash is 256 bits in length. • **HMACSHA384** - Hash-based Message Authentication Code (HMAC) with an output has that is 384 bits long. • **HMACSHA512** - Hash-based Message Authentication Code (HMAC) with an output has that is 512 bits long.
Encryption method	Select one of the following options to specify the encryption method the machine key uses: • **Auto** – This method is the default setting.

Element Name	Description
	Auto works with whichever encryption method you specified. • **AES** - Advanced Encryption Standard (AES) is relatively easy to implement and requires little memory. AES has a key size of 128, 192, or 256 bits. This method uses the same private key to encrypt and decrypt data, whereas a public-key method must use a pair of keys. • **TripleDES** - Triple Data Encryption Standard (TripleDES) is a minor variation of DES. It is three times slower than regular DES but can be more secure because it has a key size of 192 bits. If performance is not your primary consideration, consider using TripleDES. • **DES** - Data Encryption Standard (DES) uses a 56-bit key to both encrypt and decrypt data. If your server, site, or application does not require the strongest security, consider using DES.
Validation key	Computes a Message Authentication Code (MAC) to confirm the integrity of the data. This key is appended to either the Forms authentication cookie or the view state for a specific page. Select one of the following options to specify how the validation key is generated: • **Automatically generate at runtime**: Instructs ASP.NET to generate a random key at runtime. • **Generate a unique key for each application**: Isolates applications from one another by generating a unique key for each application based on the application ID of each application. If your application is deployed in a web farm, duplicate your application's key across all servers in the farm.
Decryption key	Used to encrypt and decrypt Forms authentication tickets and view state.

Element Name	Description
	Select one of the following options to specify how the decryption key is generated: • **Automatically generate at runtime**: Instructs ASP.NET to generate a random key at runtime. • **Generate a unique key for each application**: Isolates applications from one another by generating a unique key for each application based on the application ID of each application. If your application is deployed in a web farm, duplicate your application's key across all servers in the farm.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Generate Keys	Generates a validation key and a decryption key in the corresponding boxes on the feature page.

Management Service

Use the **Management Service** feature page to configure the management service for IIS Manager. The management service enables computer and domain administrators to remotely manage a web server that uses IIS Manager. The service also enables nonadministrative users to locally and remotely manage delegated features of sites and applications on a web server that uses IIS Manager.

Notes

Only administrators who use IIS Manager on the local computer can configure the management service.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Management Service
- Add Allow or Deny Connection Rule Dialog Box

UI Elements for Management Service

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Enable remote connections	Select this option when you want to enable users to connect remotely to the Web server by using IIS Manager.
Windows credentials only	Select this option when you want to allow only users who have Windows accounts to connect to the Web server.
Windows credentials or IIS Manager credentials	Select this option when you want to allow users who have either Windows accounts or IIS Manager accounts to connect to the Web server.
IP address	Type or select an IP address if you have a specific IP address that you want users to use when they connect to the web server. When you do not want to specify an IP address, select **All Unassigned**.
Port	Type the port that you want the connection to use when users connect to the Web server. The default is port 8172.
SSL certificate	Select the certificate that you want to use for the Web server. You can use the **Server Certificates** feature if you want to add a certificate.

Element Name	Description
Log requests to	Select this option if you want the Web server to log information when requestors try to connect to the Web server. Then configure the physical directory where the Web server logs files.
IP Address Restrictions	Displays connection restrictions. Use the **Allow** and **Deny** buttons to add restrictions to the list, and the Delete button to delete a selected entry from the **IP Address Restrictions** list. **Note** These options are available only when you enable remote connections.
Access for unspecified clients	Select **Allow** or **Deny** to specify whether to allow or deny requestors who do not match entries in the **IP Address Restrictions** list. If you configure the management service to deny access for unspecified clients, but do not add any allow rules to the **IP Address Restrictions list**, clients will be unable to connect to the server by using IIS Manager.
Allow	Opens the **Add Allow Connection Rule** dialog box from which you can add a rule to allow requestors based on an IP address or a range of IP addresses.
Deny	Opens the **Add Deny Connection Rule** dialog box from which you can add a rule to deny requestors based on an IP address or a range of IP addresses.
Delete	Deletes the selected entry in the **IP Address Restrictions** list.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.

Element Name	Description
Cancel	Cancels the changes that you have made on the feature page.
Restart	Restarts the management service.
Start	Starts the management service. Start the management service to connect locally to sites and applications. If you also want to allow remote connections to sites and applications on the Web server, additionally enable remote connections.
Stop	Stops the management service. When the management service is stopped, you can configure settings for the management service. However, the service must be started before it can accept connections.

Add Allow or Deny Connection Rule Dialog Box

Use the **Add Allow Connection Rule** or **Add Deny Connection Rule** dialog box to add entries to the **IP Address Restrictions** list that allow or deny remote connections to the Web server. You can allow or deny remote connections based on a specific IP address or a range of IP addresses.

Note

Only administrators who are using IIS Manager on the local computer can configure the Management Service.

Element Name	Description
Specific IP address	Select this option when you want to allow or deny remote connections from a specific IP address. After you select this option, type the IP address in the corresponding box.
IP address range	Select this option when you want to allow or deny remote connections from a range of IP addresses. After you select this option, type a range of IP addresses in the corresponding box.
Subnet mask	Type a subnet mask for the IP address that you specified in the **IP address range** box.

MIME Types

Use the **MIME Types** feature page to manage a list of Multipurpose Internet Mail Extensions (MIME) types so you can identify the types of content that can be served from a Web server to a browser or a mail client.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for MIME Types
- Add or Edit MIME Type Dialog Box

UI Elements for MIME Types

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Extension	Displays the file name extension, for example, .chm.
MIME Type	Displays the MIME type, for example, application/octet-stream.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add MIME Type** dialog box from which you can add a MIME type.

Element Name	Description
Edit	Opens the **Edit MIME Type** dialog box from which you can edit the selected MIME type.
Remove	Removes the item that is selected from the list on the feature page.

Add or Edit MIME Type Dialog Box

Use the **Add MIME Type** and **Edit MIME Type** dialog boxes to create and edit Multipurpose Internet Mail Extensions (MIME) types. MIME types help clients handle new file name extensions appropriately. If IIS does not recognize the file name extension requested by the client, IIS sends the content as the default MIME type, which is Application. This MIME type signifies that the file contains application data, and it usually means that clients cannot process the file.

Element Name	Description
Extension	Type a file name extension. This is a static file name extension, for example, .chm.
MIME Type	Type a MIME type, for example, application/octet-stream.

Modules

Use the **Modules** feature page to manage a list of native and managed code modules that perform specific tasks in the request-processing pipeline, such as authentication or compression.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Modules
- Add or Edit Managed Module Dialog Box
- Configure Native Modules Dialog Box
- **Register Native Module or Edit Native Module Registration Dialog Box**

UI Elements for Modules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the name of the module.
Code	Displays the physical path of the native code file (.dll) or of the type in the managed library.
Module Type	Displays whether the module is implemented with native code or managed code. The value is either Native or Managed.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Managed Module	Opens the **Add Managed Module** dialog box from which you can add a managed module.
Configure Native Module	Opens the **Configure Native Module** dialog box from which you can enable native modules by adding them to the list of modules. At the server level, you can also register, edit registration, and remove registration for native modules.
Edit	Opens the **Edit Managed Module** or **Edit Native Module Registration** dialog box from which you can edit the selected module. **Note** The ability to edit the registration of a native module is available only at the server level.

Element Name	Description
Lock	Locks the module so it cannot be overridden at lower levels in configuration. **Note** This action is available only at the server level.
Unlock	Unlocks the module so it can be overridden at lower levels in configuration. **Note** This action is available only at the server level.
Remove	Removes the item that is selected from the list on the feature page. When a managed module is selected, the module is removed from the Web server. When a native module is selected, the module is removed from the list of modules, which disables the module. However, it is still registered on the Web server. To remove the registration of a native module, click **Configure Native Modules**, select a module, and then click **Remove**. **Note** The ability to remove the registration of a native module is available only at the server level. You can remove more than one module at a time by pressing the CTRL key when selecting each module that you want to remove.
Revert to Inherited	Reverts the feature to inherit settings from the parent configuration. This action deletes local configuration settings, including items from the list, for this feature. This action is not available at the server level.
View Ordered List	Displays the list in order of configuration. When you select the ordered list format, you can only move items up and down in the list. Other actions in the **Actions** pane do not appear until you select the unordered list format.

Element Name	Description
Move Up	Moves up a selected item in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
Move Down	Moves a selected item down in the list. This action is available only when viewing items in the ordered list format. **Note** When items in the list are reordered at a child level, the child no longer inherits settings from the parent level. This loss of inheritance includes any items that are added to or removed from the list at the parent level. If you want to inherit settings from a parent level, revert all of the changes at the child level by using the **Revert to Inherited** action in the **Actions** pane.
View Unordered List	Displays the list in an unordered format. When you select the unordered list format, you can sort and group items in the list, and perform actions in the **Actions** pane.

Add or Edit Managed Module Dialog Box

Use the **Add Managed Module** and **Edit Managed Module** dialog boxes to add or edit managed modules that process requests made to the Web server. IIS includes several managed modules

to process parts of requests, such as the FormsAuthentication module that authenticates users in applications that use Forms authentication.

Element Name	Description
Name	Type a friendly name for the managed module.
Type	Select or type the managed type for the module, such as System.Web.Security.FormsAuthentication. When you add a managed module, make sure that the assembly to which the class belongs is registered in the system.web/compilation section of your Web.config file. Otherwise type the assembly qualified class name of the managed module. If the class is defined in the App_Code or bin directory you can specify just the class name.
Invoke only for requests to ASP.NET applications or managed handlers	Select this check box if you want the managed module to process only requests made to ASP.NET applications or managed handlers. When you select this option, the module does not invoke for requests to other types of content, such as HTML files.

Configure Native Modules Dialog Box

Use the **Configure Native Modules** dialog box to enable registered modules by adding them to the **Modules** list, add or edit registration for native modules, or remove native modules from the Web server. After you register a module, also add it to the **Modules** list in order for the native module to start processing requests.

Element Name	Description
Select one or more registered modules to enable	Select the registered module or modules that you want to add to the **Modules** list on the Web server. Adding a module to the **Modules** list will enable the module to process requests that are made to the server.
Register	Opens the **Register Native Module** dialog box from which you can register a new native module.

Element Name	Description
	Note This button is available only at the server level.
Edit	Opens the **Edit Native Module Registration** dialog box from which you can edit a registered native module. This button is available only when a registered module is selected from the **Select one or more registered modules to enable** list. **Note** This button is available only at the server level.
Remove	Removes the selected module's registration from the Web server. This button is available only at the server level, and only when a registered module is selected from the **Select one or more registered modules to enable** list. **Note** When you remove a module from the list on the **Modules** feature page, the module is disabled but it is still registered on the Web server.

Register Native Module or Edit Native Module Registration Dialog Box

Use the **Register Native Module** and **Edit Native Module Registration** dialog boxes to add or edit native modules that process requests made to the Web server. IIS includes many native modules to process parts of requests, such as the `BasicAuthenticationModule` module authenticates users in applications that use Basic authentication.

Note

After you register a native module from this dialog box, add it to the **Modules** list on the web server before the module can process requests. You can perform this task from the **Add Native Module** dialog box as soon as either the **Add Native Module** or the **Edit Native Module Registration** dialog box is closed.

Element Name	Description
Name	Type a friendly name for the native module.
Path	Type the path of the native module file on the physical file system, such as C:\Modules\MyModule.dll, or click the browse button (**...**) to browse for the file.

Output Caching

Use the **Output Caching** feature page to configure output cache settings and to configure rules that control caching of served content.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Output Caching
- Add or Edit Output Cache Rule Dialog Box
- Edit Output Cache Settings Dialog Box
- Advanced Output Caching Settings Dialog Box

UI Elements for Output Caching

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Extension	Displays the file name extension for which the cache rule applies.
User-Mode Policy	Displays the user-mode cache behavior for the rule:

Element Name	Description
	• **Cache until change**: content is cached until IIS receives a file or configuration change notification. • **Cache for time period**: content is cached for a specific duration unless a file or configuration change notification clears the cache before the duration elapses. • **Prevent all caching**: user-mode caching has been disabled. • **Do not cache**: user-mode caching has not been configured.
Kernel-Mode Policy	Displays the kernel-mode caching behavior for the rule: • **Cache until change**: content is cached until IIS receives a file or configuration change notification. • **Cache for time period**: content is cached for a specific duration unless a file or configuration change notification clears the cache before the duration elapses. • **Prevent all caching**: kernel-mode caching has been disabled. • **Do not cache**: kernel-mode caching has not been configured.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Cache Rule** dialog box in which you can add a cache rule.
Edit	Opens the **Edit Cache Rule** dialog box in which you can edit the selected cache rule.
Remove	Removes the selected caching rule.

Element Name	Description
Edit Feature Settings	Opens the **Edit Output Cache Settings** dialog box in which you can configure settings that apply to the entire output caching feature.

Add or Edit Output Cache Rule Dialog Box

Use the **Add Cache Rule** or **Edit Cache Rule** dialog box to configure how IIS caches content for a specific file name extension. You can configure settings for both user-mode and kernel-mode caching.

Element Name	Description
File name extension	Specifies the file name extension (such as .aspx) to which the caching rule applies. To apply the rule to all file name extensions, type an asterisk (*) in the box.
User-mode caching	Configures the rule to store content in the output cache, which is in user mode. Select an option under **File Cache Monitoring** to configure how or when the output cache is updated for any files that are cached because of this rule.
Using file change notifications (under the **User-mode caching option**)	Configures the files (that are cached because of this rule) to be removed from the output cache when IIS receives a file or configuration change notification.
At time intervals (hh:mm:ss) (under the **User-mode caching option**)	Configures the files (that are cached because of this rule) to be removed from the output cache after a specific duration unless a file or configuration change notification clears the cache before the duration elapses. The value in the corresponding box specifies the length of time that content is stored in the output cache. The time is formatted as hh:mm:ss (where hh is hours, mm is minutes, and ss is seconds). The default duration is 30 seconds. Setting the value too low can affect server performance by forcing items to be updated more frequently than necessary.

Element Name	Description
	Note When you select both user-mode caching and kernel-mode caching, and configure each mode to use the **At time intervals (hh:mm:ss)** option, specify the same duration value for both modes. If the values are not the same, the duration that is specified for kernel-mode caching is used.
Prevent all caching	Configures the files (that match the file name extension) not to be stored in the output cache. This setting corresponds to the **DisableCache** setting in configuration, and differs from the **DontCache** setting that appears in configuration when you do not select the **User-mode caching** option in IIS Manager. The **DontCache** setting means that the output cache module does not set the user-mode caching policy, but if another module sets the caching policy, the output cache module caches the response. The **DisableCache** setting means that the output cache does not cache the response even when another module tries to enable user-mode caching by setting the user-mode caching policy.
Advanced	Opens the **Advanced Output Caching Settings** dialog box in which you can configure the cache rule to cache different versions of a response based on query strings and headers. These options are available only for user-mode caching.
Kernel-mode caching	Configures the rule to store content in the kernel cache, which is in kernel mode. Select an option under **File Cache Monitoring** to configure how or when the kernel cache is updated for any files that are cached because of this rule.
Using file change notifications (under the **Kernel-mode caching option**)	Configures the files (that are cached because of this rule) to be removed from the kernel

Element Name	Description
	cache when IIS receives a file or configuration change notification.
At time intervals (hh:mm:ss) notifications (under the **Kernel-mode caching option**)	Configures the files (that are cached because of this rule) to be removed from the kernel cache after a specified time elapses. The value in the corresponding box specifies the length of time that content is stored in the kernel cache. The time is formatted as hh:mm:ss (where hh is hours, mm is minutes, and ss is seconds). The default duration is 30 seconds. Setting the value too low can affect server performance by forcing items to be updated more frequently than necessary. **Note** When you select both user-mode caching and kernel-mode caching, and configure each mode to use the **At time intervals (hh:mm:ss)** option, specify the same duration value for both modes. If the values are not the same, the duration that is specified for kernel-mode caching is used.
Prevent all caching notifications (under the **Kernel-mode caching option**)	Configures the files (that match the file name extension) not to be stored in the kernel cache. This setting corresponds to the **DisableCache** setting in configuration, and differs from the **DontCache** setting that appears in configuration when you do not select the **Kernel-mode caching** option in IIS Manager. The **DontCache** setting means that the output cache module does not set the kernel-mode caching policy, but if another module sets the caching policy, the output cache module caches the response. The **DisableCache** setting means that the kernel cache does not cache the response even when another module tries to enable kernel-mode caching by setting the kernel-mode caching policy.

Edit Output Cache Settings Dialog Box

Use the **Edit Output Cache Settings** dialog box to configure settings that affect the entire **Output Caching** feature. These settings affect all the output caching rules that you configure on the **Output Caching** feature page.

Element Name	Description
Enable cache	Enables the IIS output cache, which stores cached responses in user mode. **Note** The IIS output cache is similar to the ASP.NET output cache. However, this IIS cache is a native output cache that offers increased performance over the managed output cache in ASP.NET.
Enable kernel cache	Enables the kernel cache, which stores cached responses in kernel mode. Performance is improved when responses are returned from the kernel cache without transitioning to user mode.
Maximum cached response size (in bytes)	Specifies the maximum size of a cached response for both the user-mode and kernel-mode caches. The default value is 262144 bytes. This field is enabled at the server level only; it is read-only at all other levels.
Cache size limit (in MB)	Configures the size limit of both the user-mode and kernel-mode caches. You can type a size (in MB) or type 0. If you type 0, IIS uses half of the available physical memory or virtual memory, whichever is less. This field is enabled at the server level only; it is read-only at all other levels.

Advanced Output Caching Settings Dialog Box

Use the **Advanced Output Caching Settings** dialog box to configure caching of different file versions based on query strings and headers. These settings are available only for the user-mode cache.

Element Name	Description
Query string	Type a query string or a comma-separated list of query strings to define the values that the Web server uses to vary cached responses to the URL. You can also type an asterisk (*) to configure the rule to vary cached responses for all query strings.
Headers	Type a header or comma-separated list of headers to define the values that the Web server uses to vary cached responses to the URL. You can also type an asterisk (*) to configure the rule to vary cached responses for all headers.

Pages and Controls

Use the **Pages and Controls** feature page to configure the settings of ASP.NET pages and controls. These settings control the behavior of ASP.NET pages, set how ASP.NET pages and controls are compiled on the server, define the namespaces that are included for each page, and define the session state and validation services settings.

Use the **Display** list to select one of the following options that specifies how the settings are displayed: **Friendly Names**, **Configuration Names**, or **Both Names**.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Pages and Controls
- Controls Page
- Add or Edit Custom Control and Add or Edit User Control Dialog Boxes

UI Elements for Pages and Controls

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Buffer	Sets a value (true or false) that specifies whether .aspx pages and .ascx controls use response buffering. The default setting is **True**.
Master Page File	Specifies the master page path relative to the local configuration file.
Style Sheet Theme	Specifies the name of the "named theme" folder that is used to apply the theme before control declarations. This option differs from the theme attribute, which defines the theme to apply after the control declarations.
Theme	Specifies the name of the theme that is used for the pages that are in the scope of the configuration file. The specified theme must exist as either an application or a global theme. If the theme does not exist, an **HttpException** exception is thrown.
Enable Authenticated View State	Sets a value (true or false) that specifies whether ASP.NET should run a message authentication code (MAC) on the page's view state when the page is posted back from the client. The default setting is **True**.
Enable View State	Sets a value (true or false) that indicates whether view state is enabled, disabled, or read only. The default setting is **True**.
Maximum Page State Field Length	Sets the maximum number of characters that a single view-state field can contain.
Base Type for Pages	Specifies the base type to use when the pages are stand-alone only. This option is overridden by the **inherits** attribute in a stand-alone file.
Base Type for User Controls	Specifies a code-behind class that user controls inherit by default.
Compilation Mode	Specifies whether an ASP.NET page or control should be compiled at run time.
Namespaces	Defines a collection of import directives to use

Element Name	Description
	during assembly precompilation.
Enable Session State	Sets a value (true or false) that specifies whether the session state is enabled, disabled, or read-only. The default setting is **True**.
Validate Request	Sets a value (true or false) that determines whether ASP.NET examines input from the browser for dangerous values. The default setting is **True**.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Register Controls	Opens the **Controls** feature page from which you can add and edit custom controls.

Controls Page

Use the **Controls** feature page to manage a list of user controls and custom controls.

Note

User controls are not available at the server level in IIS Manager.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Feature Page Elements

Element Name	Description
Tag	For custom controls, displays the tag prefix assigned to the control. For user controls, displays the tag prefix and tag name for the user control.

Element Name	Description
Source or Assembly	Displays the user control path or assembly name in which the custom control type is defined.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add Custom Control	Opens the **Add Custom Control** dialog box in which you can add a custom control.
Add User Control	Opens the **Add User Control** dialog box in which you can add a user control. This action is not available at the server level.
Edit	When a custom control is selected, this action opens the **Edit Custom Control** dialog box in which you can edit the custom control. When a user control is selected, this action opens the **Edit User Control** dialog box in which you can edit the user control. You cannot edit a user control at the server level.
Remove	Removes the item that is selected from the list on the feature page.

Add or Edit Custom Control and Add or Edit User Control Dialog Boxes

Use the **Add Custom Control** and **Edit Custom Control** dialog boxes to specify or edit the tag prefix/namespace mapping for a custom control that will be used in multiple pages in an application. Use the **Add User Control** and **Edit User Control** dialog boxes to configure user controls, which are containers into which you can put markup and Web server controls.

For more information about user and custom controls, see ASP.NET User Controls Overview.

Element Name	Description
Tag prefix	Sets the tag prefix assigned to the control, for example, <aspSample: …>
Namespace	Defines the namespace to which the custom control belongs. This setting is the same namespace that was specified in the application code. This box is available only for custom controls.
Assembly	Sets the assembly name in which the custom control type is defined. This box is available only for custom controls.
Tag Name	Specifies the tag name of the user control. This box is available only for user controls.
User Control (.ascx) Path	Specifies the file path of the user control. This box is available only for user controls.

Providers

Use the **Providers** feature page to manage a list of ASP.NET data providers.

Sort the list by clicking one of the feature page column headings or select a value from the **Group by** drop-down list to group similar items.

Related scenarios

- Build an ASP.NET Website on IIS

In this document

- UI Elements for Providers
- Add or Edit Provider Dialog Box

UI Elements for Providers

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Feature	Select a feature from the list to display providers for that feature: **.NET Roles**, **.NET Users**, and **.NET Profile**.
Name	Displays the name of the provider. This is a friendly name that is assigned to the provider for easy identification and recognition.
Type	Displays the managed code type for the provider. All providers are implemented in managed code.
Entry Type	Displays whether the item is local or inherited. Local items are read from the current configuration file, and inherited items are read from a parent configuration file.

Actions Pane Elements

Element Name	Description
Add	Opens the **Add Provider** dialog box in which you can add a provider for the selected feature.
Edit	Opens the **Edit Provider** dialog box in which you can edit the selected provider.
Rename	Enables the **Name** field of the selected provider so that you can rename the provider.
Remove	Removes the item that is selected from the list on the feature page.
Connection Strings	Opens the **Connection Strings** feature page from which you can add a connection string that can be used by a provider.

Add or Edit Provider Dialog Box

Use the **Add Provider** and **Edit Provider** dialog boxes to add or modify a provider in IIS when you have an application that uses a provider-based service to store data in a database or other data store. The properties that are displayed in the dialog box depend on which type of provider you choose to add or edit.

The 3 different types of providers you can add or edit are:

- **Roles** - provides an interface between the ASP.NET role management service (the "role manager") and role data sources.
- **Users** - provides an interface between the ASP.NET membership service and membership data sources.
- **Profile** - provides an interface between the ASP.NET profile service and profile data sources.

Note

The **Add Provider** and **Edit Provider** dialog boxes enable you to add or edit most of the properties for a provider, but not all of them. You can configure additional properties for each provider in the Web.config file of your web application. For more information about providers see ASP.NET 2.0 Provider Model: Introduction to the Provider Model.

Element Name	Description
Type	Sets the provider type.
Name	Defines the name of the provider.
Behavior	Sets some of the behaviors for Membership User providers. Setting the value of one or more of the following behaviors to **True** enables that behavior: • **EnablePasswordReset** - indicates whether passwords can be reset using the provider's **ResetPassword** method. The default setting is **False**. • **EnablePasswordRetrieval** - indicates whether passwords can be retrieved using the provider's **GetPassword** method. The default setting is **False**. • **RequiresQuestionAndAnswer** - indicates whether a password answer must be supplied when calling the provider's **GetPassword** and **ResetPassword** methods. The default setting is **False**. • **RequiresUniqueEmail** - indicates whether each registered user must have a unique e-mail address. The default setting is **False**. • **StorePasswordInSecureFormat** - indicates whether passwords should be hashed. The default setting is **False**.
Data	Configures the name of the connection string used to connect to the database. Enter the name

Element Name	Description
	of the connection string in the **ConnectionStringName** check box.
General	Defines general settings for the provider. These settings are optional and include: • **ApplicationName** - specifies the virtual path of the application in the **ApplicationName** box. If you do not specify a value in the **ApplicationName** box, the membership API defaults to **HttpContext.Current.Request.ApplicationPath**. This value can be determined only at runtime. • **Description** - gives a description of the provider.
Misc	Enables you to configure name and value pairs that are specific to the provider. To add and remove name and value pairs, click the **…** button next to the **Provider Specific Settings** field in the **Misc** section.

Request Filtering

Use the **Request Filtering** feature page to configure filtering rules for your website. Web request filtering, similarly to FTP request filtering, lets you restrict protocol and content behavior.

Related scenarios

- Configure Request Filtering in IIS
- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Request Filtering
- Edit Request Filtering Settings Dialog Box
- File Name Extensions Tab
- Allow or Deny File Name Extension Dialog Boxes
- Rules Tab
- Add Filtering Rule Dialog Box

- Hidden Segments Tab
- Add Hidden Segment Dialog Box
- URL Tab
- Allow URL Dialog Box
- Add Deny Sequence Dialog Box
- HTTP Verbs Tab
- Allow or Deny Verb Dialog Boxes
- Headers Tab
- Add Header Dialog Box
- Query Strings Tab
- Allow or Deny Query String Dialog Boxes

UI Elements for Request Filtering

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
File Name Extensions	Specifies a list of file name extensions for which the request-filtering service allows or denies access.
Rules	Lists the filtering rules and the specific parameters for which the request-filtering service should scan. These parameters include headers, file name extensions, and deny strings.
Hidden Segments	Specifies a list of hidden segments for which the request-filtering service denies access. These segments do not display in directory listings.
URL	Specifies a list of URL sequences for which the request-filtering service allows or denies access.
HTTP Verbs	Specifies a list of HTTP verbs for which the request-filtering service allows or denies access.

Element Name	Description
Headers	Specifies the headers and their size limits for which the request-filtering service denies access.
Query Strings	Specifies the query strings for which the request-filtering service denies access.

Actions Pane Elements

Element Name	Description
Edit Feature Settings	Opens the **Edit Request Filtering Settings** dialog box that lets you configure general properties and request limits.
Remove	Removes a file name extension, rule, hidden segment, URL sequence, HTTP verb, header, or query string from the list.

Edit Request Filtering Settings Dialog Box

Use the **Edit Request Filtering Settings** dialog box to configure general settings and request limits on your Web server.

Element Name	Description
Allow unlisted file name extensions	Select this option to allow unlisted file name extensions.
Allow unlisted verbs	Select this option to allow unlisted verbs.
Allow high-bit characters	Select this option to allow high-bit characters in request filter operations. Examples of high-bit characters are: Ж, Ы, and Я.
Allow double escaping	Select this option to allow double escaping.
Maximum allowed content length (Bytes)	Specifies the maximum length, in bytes, for content. **Note** Specifying zero (0) for this value means that the length is unlimited.

Element Name	Description
Maximum URL length (Bytes)	Specifies the maximum length, in bytes, for a URL string.
Maximum query string (Bytes)	Specifies the maximum length, in bytes, for a query string.

File Name Extensions Tab

Use the File Name Extensions tab to create a list of file name extensions for which the request-filtering module allows or denies access. This helps restrict certain file types and increase security on your web server.

Feature Page Elements

Element Name	Description
File Name Extension	Displays the file name extension for which the request-filtering module either allows or denies access.
Allowed	Displays the status of the file name extension. The status is True if the file name extension is allowed, or False if the file name extension is denied.

Actions Pane Elements

Element Name	Description
Allow File Name Extension	Opens the **Allow File Name Extension** dialog box that lets you add a file name extension to the list of allowed file name extensions.
Deny File Name Extension	Opens the **Deny File Name Extension** dialog box that lets you add a file name extension to the list of denied file name extensions.

For common elements, see the **Request Filtering Page**.

Allow or Deny File Name Extension Dialog Boxes

Use the **Allow File Name Extension** or **Deny File Name Extension** dialog box to add a file name extension to the list of file name extensions for which the request-filtering module will allow or deny access.

Element Name	Description
File name extension	Specifies the file name extension for which the request-filtering service either allows or denies access.

Rules Tab

Use the Rules tab to configure request-filtering rules for allowing or denying access to your Web server based on several parameters, such as headers and deny strings. For examples of filtering rules, see Common URLScan Scenarios.

Feature Page Elements

Element Name	Description
Name	Displays the name of the request-filtering rule.
Scan	Displays the value that the request-filtering rule scans for the URL, query string, or any combination thereof.
Applies To	Displays the parameters to which the request-filtering rule applies.
Deny Strings	Displays the strings to which the request-filtering rule denies access.

Actions Pane Elements

Element Name	Description
Add Filtering Rule	Opens the **Add Filtering Rule** dialog box that lets you add a rule with specific parameters for which to scan.

For common elements, see the **Request Filtering Page**.

Add Filtering Rule Dialog Box

Use the **Add Filtering Rule** dialog box to add a rule with specific parameters to the request-filtering module's rule list. This action enables the module to deny requests that fit the criteria specified by the rule.

Element Name	Description
Name	Specifies the name of the filtering rule that is added to the list.
Scan url	Select this option and the request-filtering module scans a request URL.
Scan query string	Select this option and the request-filtering module scans a request query string.
Scan Headers	Adds a user-specified header to the list of headers to search for in the request. Use the asterisk (*) to add a row to the table.
Applies To	Adds a file name extension to the list of file name extensions to which this request-filtering rule applies. Use the asterisk (*) to add a row to the table.
Deny Strings	Adds a string to the list of strings to search for in the request. If the specified string is found in the request URL, query string, or header, then the request is denied. Use the asterisk (*) to add a row to the table.

Hidden Segments Tab

Use the **Hidden Segments** tab to define the list of URL segments for which the request-filtering module will deny access and will exclude from directory listings. A URL segment is the part of the URL path that lies between the slash (/) marks.

Feature Page Elements

Element Name	Description
Segment	Displays the URL segment for which the request-filtering service denies access and which it does not display in directory listings.

Actions Pane Elements

Element Name	Description
Add Hidden Segment	Opens the **Add Hidden Segment** dialog box that lets you add a hidden segment to the list of hidden segments.

For common elements, see the **Request Filtering Page**.

Add Hidden Segment Dialog Box

Use the **Add Hidden Segment** dialog box to add a URL segment to the list of URL segments for which the request-filtering module will deny access. A URL segment is a part of the URL path that lies between the slash (/) marks.

Element Name	Description
Hidden segment	Specifies the URL segment for which the request-filtering module denies access. **Note** URL segments that are in the hidden segments list does not appear in directory listings.

URL Tab

Use the **Deny URL Sequences** tab to create a list of URL sequences for which the request-filtering module will deny access. For example, you can specify “admin/config.xml” as a URL sequence, which denies requests to http://contoso.com/application/admin/config.xml.

Feature Page Elements

Element Name	Description
Deny URL Sequence	Displays the URL sequence for which the request-filtering module will deny access.

Actions Pane Elements

Element Name	Description
Allow URL	Opens the **Add Always Allowed URL** dialog box that lets you add a URL to the list of allowed URLs.
Deny Sequence	Opens the **Add Deny Sequence** dialog box that lets you add a sequence to the list of denied sequences.

For common elements, see the **Request Filtering Page**.

Allow URL Dialog Box

Use the **Add Always Allowed URL** dialog box to add a URL sequence to the list of URL sequences for which the request-filtering module will always grant access.

Element Name	Description
URL	Specifies the URL for which the request-filtering module always grants access.

Add Deny Sequence Dialog Box

Use the **Add Deny Sequence** dialog box to add a URL sequence to the list of URL sequences for which the request-filtering module will deny access.

Element Name	Description
URL sequence	Specifies the URL sequence for which the request-filtering module denies access.

HTTP Verbs Tab

Use the **HTTP Verbs** tab to create a list of verbs for which the request-filtering module will allow or deny access. Several examples of HTTP verbs are: GET, POST, and HEAD.

Feature Page Elements

Element Name	Description
Verb	Displays the verb for which the request-filtering module either allows or denies access.
Allowed	Displays the status of the verb, which will be either **True** if the verb is allowed, or **False** if the verb is denied. **Note** Only verbs that appear in the list and have **Allowed** set to **False** will be blocked. All other verbs are allowed.

Actions Pane Elements

Element Name	Description
Allow Verb	Opens the **Allow Verb** dialog box that lets you add a verb to the list of allowed verbs.
Deny Verb	Opens the **Deny Verb** dialog box that lets you add a verb to the list of denied verbs.

For common elements, see the **Request Filtering Page**.

Allow or Deny Verb Dialog Boxes

Use the **Allow Verb** dialog box to add a verb to the list of HTTP verbs for which the request-filtering module will allow access. Use the **Deny Verb** dialog box to add a verb to the list of HTTP verbs for which the request-filtering module will deny access.

Element Name	Description
Verb	Specifies the HTTP verb for which the request-filtering module either allows or denies access.

Headers Tab

Use the **Headers** tab to create a list of headers for which the request-filtering module will deny access if the value of a header is larger than the specified size.

Feature Page Elements

Element Name	Description
Header	Displays the header for which the request-filtering module denies access if the size is larger than the specified size.
Size Limit	Displays the maximum size allowed for the header. For example, specifying a value of 100 would limit the length of a content-type header to 100 bytes.

Actions Pane Elements

Element Name	Description
Add Header	Opens the **Add Header** dialog box that lets you add a header to the list of denied headers.

For common elements, see the **Request Filtering Page**.

Add Header Dialog Box

Use the **Add Header** dialog box to specify a size limit for a particular HTTP header. If a header's length exceeds the specified value, then the HTTP request that contains this header is denied.

Element Name	Description
Header	Specifies the header for which the request-filtering module checks the size limit.
Size limit	Specifies the maximum character length for the header's value.

Query Strings Tab

Use the **Query Strings** tab to create a list of query strings for which the request-filtering module will either always allow or deny access. An example of a query string is "%3b", which is used to catch SQL injection attempts.

Feature Page Elements

Element Name	Description
Query String	Displays the query string for which the request-filtering module either allows or denies access.
Action	Displays the status of the query string, which will be either **Always allow** or **Deny**.

Actions Pane Elements

Element Name	Description
Allow Query String	Opens the **Allow Query String** dialog box that lets you add a query string that the request-filtering service will always allow.
Deny Query String	Opens the **Deny Query String** dialog box that lets you add a query string that the request-filtering service will deny.

For common elements, see the **Request Filtering Page**.

Allow or Deny Query String Dialog Boxes

Use either the **Allow Query String** or **Deny Query String** dialog box to add a query string to the list of query strings for which the request-filtering module will either always allow or deny access.

Element Name	Description
Query string	Specifies the query string for which the request-filtering service either always allows or denies access.

Server Certificates

Use the **Server Certificates** feature page to view the names of certificates, the fully qualified domain names (FQDNs) of hosts to which certificates have been issued, and the FQDNs of the servers that issued the certificates.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS

- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build a Web Farm with IIS Servers

In this document

- UI Elements for Server Certificates
- Import Certificate Dialog Box
- Request Certificate Wizard
- Complete Certificate Request Dialog Box
- **Create Certificate Wizard**
- Create Self-Signed Certificate Dialog Box
- Export Certificate Dialog Box
- Renew an Existing Certificate Wizard

UI Elements for Server Certificates

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the names of certificates that have been issued to clients that are running on either Internet or intranet hosts. **Note** Certificates are not required to have names. You might have to view other columns to obtain information about certificates.
Issued To	Displays the FQDNs of either the Internet or intranet hosts to which certificates have been issued.
Issued By	Displays the FQDNs of servers that have issued certificates to clients that are running on either Internet or intranet hosts.
Expiration Date	Displays the date that the certificate expires.
Certificate Hash	Displays binary data produced by using a hashing algorithm. Although this data uniquely

Element Name	Description
	identifies a certificate, the hash data cannot be used to trace a certificate because hashing is a one-way process.
Certificate Store	Displays the name of the provider that stores the certificate.

Actions Pane Elements

Element Name	Description
Import	Opens the **Import Certificate** dialog box to restore a lost or damaged certificate that you previously backed up, or to install a certificate sent to you by another user or certification authority (CA).
Create Certificate Request	Opens the **Request Certificate** wizard to provide information about your organization to an external certification authority.
Complete Certificate Request	Opens the **Complete Certificate Request** dialog box to install the certificates that you receive from your certification authority.
Create Domain Certificate	Opens the **Create Certificate** wizard to provide information about your organization to an internal certification authority.
Create Self-Signed Certificate	Opens the **Create Self-Signed Certificate** dialog box to create certificates to use in server testing environments and for troubleshooting third-party certificates.
View	Opens the **Certificate** dialog box so that you can view details about a certificate. Select a certificate to see this option.
Export	Opens the **Export Certificate** dialog box to export certificates from a source server when you want to apply the same certificate to a target server, or when you want to back up a certificate and its associated private key. Select a certificate to see this option.

Element Name	Description
Remove	Removes the item that is selected from the list on the feature page. Select a certificate to see this option.

Import Certificate Dialog Box

Use the **Import Certificate** dialog box to restore a lost or damaged certificate that you previously backed up, or to install a certificate sent to you by another user or certification authority (CA).

Element Name	Description
Certificate file (.pfx)	Type a file name in the **Certificate file (.pfx)** box or click **Browse** to navigate to the name of a file where the exported certificate is stored.
Password	Type the password in the **Password** field, if the certificate was exported with a password.
Select Certificate Store	Displays the name of the provider that stores the certificate.
Allow this certificate to be exported	Select **Allow this certificate to be exported** if you want to be able to export the certificate, or clear **Allow this certificate to be exported** if you do not want to allow additional exports of this certificate.

Request Certificate Wizard

Use the **Request Certificate** wizard to request a certificate from a certification authority (CA).

Distinguished Name Properties Wizard Page

Use the **Distinguished Name Properties** dialog box to provide information about your organization to an internal or external certification authority.

Element Name	Description
Common name	Type a name for the certificate.
Organization	Type the name of the organization for which the certificate is used.

Element Name	Description
Organizational unit	Type the name of the department or division in the organization in which the certificate is used.
City/locality	Type the unabbreviated name of the city or locality where your organization or organizational unit is located.
State/province	Type the unabbreviated name of the state or province where your organization or organizational unit is located.
Country/region	Type the name of the country or region where your organization or organizational unit is located.

Cryptographic Service Provider Properties Wizard Page

Use the **Cryptographic Service Provider** Wizard page to select either **Microsoft RSA SChannel Cryptographic Provider** or **Microsoft DH SChannel Cryptographic Provider** to provide certificates that can encrypt transmissions between your server and clients. Additionally, you can adjust the level of security for your transmission by changing the bit length associated with the cryptographic service provider.

Element Name	Description
Cryptographic service provider	Select either **Microsoft RSA SChannel Cryptographic Provider** or **Microsoft DH SChannel Cryptographic Provider**. The default cryptographic service provider is **Microsoft RSA SChannel Cryptographic Provider**. **Note** Select **Microsoft DH SChannel Cryptographic Provider** when you must exchange a secret key over a network that is not secure and you have had no prior communication with the other party.
Bit length	Select a bit length that the provider you selected uses. By default, the RSA SChannel provider uses a bit length of 1024, and the DH

Element Name	Description
	SChannel provider uses a bit length of 512. **Note** A longer bit length increases the level of encryption. However, it can decrease performance because it requires the transmission of additional bits.

File Name Wizard Page

Use the **File Name** dialog box to name and then save your certificates to the appropriate storage location.

Element Name	Description
Specify a file name for certificate request	Type a file name in the **Specify a file name for the certificate request** field.
…	Navigate to a file name under which to store the certificate.

Complete Certificate Request Dialog Box

Use the **Complete Certificate Request** dialog box to install the certificates that you receive from your certification authority (CA). Additionally, provide a **Friendly name** for the certificate that you want to install to complete the certificate installation process.

Element Name	Description
File name containing certification authority's response	Type the path of the file that contains the response from the certification authority in the **File name containing certification authority's response** box, or click **Browse** to navigate to the location in which the file from the certification authority is stored. **Important** Complete this process to install a certificate on your server.
Friendly name	Type a name in the **Friendly name** box to complete the certificate installation process.

Element Name	Description
Select a certificate store for the new certificate	Select from a list of available certificate providers.

Create Certificate Wizard

Use the **Create Certificate** wizard to create a domain certificate. A domain certificate is an internal certificate that is not issued by an external certification authority (CA).

Distinguished Name Properties Wizard Page

Use the **Distinguished Name Properties** dialog box to provide information about your organization to an internal or external certification authority.

Element Name	Description
Common name	Type a name for the certificate.
Organization	Type the name of the organization for which the certificate is used.
Organizational unit	Type the name of the department or division in the organization in which the certificate is used.
City/locality	Type the unabbreviated name of the city or locality where your organization or organizational unit is located.
State/province	Type the unabbreviated name of the state or province where your organization or organizational unit is located.
Country/region	Type the name of the country or region where your organization or organizational unit is located.

Online Certification Authority Wizard Page

Use the **Online Certification Authority** Wizard page to identify an online certification authority (CA) server in your Windows domain. Additionally, supply the CA server that you want to use with a **Friendly name** to complete the **Create Domain Certificate** Wizard.

Element Name	Description
Specify Online Certification Authority	Type the path of a CA server that is in your Windows domain, or click **Select** to search for a CA server that is in your domain and display the **Select Certification Authority** dialog box. **Note** Domain certificates are not appropriate for use with external clients that are not members of your internal Windows domain.
Friendly name	Type a name for the CA server that you want to use in the **Friendly name** box. Type a name in the **Friendly name** box to complete the **Create Domain Certificate** Wizard.

Select Certification Authority Dialog Box

Use the **Select Certification Authority** dialog box to select the internal certification authority (CA) that you want to use.

Element Name	Description
Select a certification authority you want to use	Lists the friendly names of CA and the fully qualified domain name (FQDN) of the computer that hosts the CA. Select the CA that you want to use.

Create Self-Signed Certificate Dialog Box

Use the **Create Self-Signed Certificate** dialog box to create certificates to use in server testing environments and for troubleshooting third-party certificates.

You can view the properties of your self-signed certificate on the **Server Certificates Page**.

Element Name	Description
Specify a friendly name for the certificate	Type a friendly name in the **Name** box to create a self-signed certificate. **Note** The certificates you create with this

Element Name	Description
	feature are not from a trusted certification authority (CA) source. Therefore, use self-signed certificates only to help secure data transmissions between your server and clients inside a test environment.

Export Certificate Dialog Box

Use the **Export Certificate** dialog box to export certificates from a source server when you want to apply the same certificate to a target server, or when you want to back up a certificate and its associated private key.

Note
If you associate a password with the certificate, whoever imports the certificate must know the password before the certificate can be applied to the target server.

Element Name	Description
Export to	Type a file name in the **Export to** box or click **Browse** to navigate to the name of a file in which to store the certificate for exporting.
Password	Type a password in the **Password** box if you want to associate a password with the exported certificate.
Confirm password	Retype the password in the **Confirm password** box and then click **OK**.

Renew an Existing Certificate Wizard

Use the **Renew an Existing Certificate** wizard to renew a certificate that is about to expire.

Important
You cannot renew a certificate that has already expired. If you try to renew a certificate that has expired, the certification authority (CA) rejects the request, and you will see an error message similar to "Error Verifying Request Signature or Signing Certificate. A required certificate is not within its validity period when verifying against the current system clock or the timestamp in the signed file." This message is also displayed in the Failed Requests node of the issuing CA. If your certificate has already expired, request a new certificate instead of renewing the existing certificate.

Element Name	Description
Renew an existing certificate	Select this option to renew an existing certificate from an internal certification authority (CA) on your domain.
Create a renewal certificate request	Select this option to package your renewal information for later submission to a CA.
Complete certificate renewal request	Select this option to complete the certificate renewal request with the certificate you received from a CA.

Session State

Use the **Session State** feature page to configure the behavior of information maintained across browser sessions.

Related scenarios

- Build an ASP.NET Website on IIS

UI Elements for Session State

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Not enabled	Disables session state.
In process	Stores session state data for a managed-code application in the worker process where the application runs. This setting is the default.
Custom	Configures IIS to use a custom provider to handle session state for ASP.NET applications.
State Server	Enables the Windows Aspnet_state.exe state service and stores session state outside the worker process where the application runs. The advantage of this configuration is that session state is preserved when the application's worker process recycles. Using a state server is

Element Name	Description
	recommended for medium-sized web applications. The settings to configure are: • **Connection string** - sets the connection string that is used to connect to the state server. • **Time-out (in seconds)** - sets the time, in seconds, that the connection will be maintained. The default is 10 seconds. **Important** The Windows state service (Aspnet_state.exe) must be running for out-of-process session state to take effect. By default, this service is installed when ASP.NET is installed and is configured for manual start. Change the start behavior to Automatic.
SQL Server	Configures IIS to use a SQL Server database to store session state data instead of storing it in the worker process where the application runs. The advantage of this configuration is that session state is preserved when the application's worker process recycles. In addition, session state is preserved if either the Windows state service or the web server goes down. The settings to configure are: • **Connection string** - sets the connection string that is used to connect to the state server. • **Time-out (in seconds)** - sets the time, in seconds, that the connection will be maintained. The default is 10 seconds. **Important** Before you configure a SQL Server for session state, run the InstallSqlState.sql script on the server. By default, this script is stored in *systemroot*\Microsoft.NET\Framework\V2.0.50727.
Enable custom database	Enables a custom SQL Server database for storing session state data.
Mode	Defines how cookies are used to store session state data. The options are: • **Auto Detect** - uses cookies if the browser supports cookies; otherwise, no cookies are used. For desktop browsers that are known to support cookies, ASP.NET tries to use cookies when cookie support is enabled in

Element Name	Description
	the browser. When you use the Auto-Detect cookie mode, you should require that expired session IDs be regenerated. Doing this task lets a web server expire and regenerate tokens, which give potential attackers less time to capture a cookie and gain access to Web server content. Also consider changing the time-out value to less than the 20-minute default. • **Use Cookies** - associates session information with client information for the duration of a user's connection to a website. Cookies are passed on together with all requests between a client and a Web server in an HTTP header. Using cookies is a more efficient way to track session state than any of the methods that do not use cookies, because cookies do not require any redirection. Additionally, cookies let users bookmark web pages, and they retain state if a user leaves one site to visit another and then returns to the original site. **Note** Consider changing the time-out value to less than the 20-minute default so a potential attacker has less time to capture a cookie and gain access to website content. • **Use Device Profile** - uses cookies if the device profile supports cookies, otherwise, no cookies are used. If the device profile indicates support for cookies, they are used regardless of whether the user has disabled cookie support. When you use the Use Device Profile cookie mode, you should require that expired session IDs be regenerated. Doing this task lets a web server expire and regenerate tokens, which give potential attackers less time to capture a cookie and gain access to Web server content. Also consider changing the time-out value to less than the 20-minute default. • **Use URI** - embeds the session ID as a query string in the Uniform Resource Identifier (URI) request, and then the URI is redirected to the originally requested URL. The changed URI request is used for the duration of the session, so no cookie is necessary. When you use a URI, require that expired session IDs be regenerated. Doing this task lets a web server expire and regenerate tokens, which give potential attackers less time to capture a cookie and gain

Element Name	Description
	access to Web server content.
Name	Sets a name for the cookie. The default is ASP.NET_SessionID.
Time-out (in minutes)	Sets the time, in minutes, that a cookie persists. The default is 20 minutes.
Regenerate expired session ID	Tells IIS to reject and reissue session IDs that do not have corresponding sessions active in the database. By default, this feature is supported only for cookieless session IDs, but this feature can be extended in order to manage cookie/arbitrary session IDs by implementing a custom session ID manager.
Use hosting identity for impersonation	Enables Windows authentication and the host process identity (either ASP.NET or a Windows service identity) for remote connections.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Shared Configuration

Use the **Shared Configuration** feature page to share IIS configuration files and encryption keys between one or more IIS servers. You can also use this page to remove a computer from a group that was set up to share configuration settings.

Related scenarios

1. Build a Web Farm with IIS Servers

In this document

- UI Elements for Shared Configuration
- Export Configuration Dialog Box
- Set Credentials Dialog Box

- Encryption Keys Password Dialog Box

UI Elements for Shared Configuration

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Enable shared configuration	Configures IIS to use shared configuration files. **Note** You must exit and restart IIS Manager after enabling or disabling this option.
Physical path	Describes the location from which to read configuration files and encryption keys. This can be a network path or a folder on the local computer.
User name	Specifies the user name IIS should use to access the location for configuration file storage.
Password	Specifies the password needed to authenticate the user name that IIS uses to access the location for configuration file storage.
Confirm password	Specifies that the user must retype the password.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.
Export Configuration	Opens the **Export Configuration** dialog box to export IIS configuration files and encryption

Element Name	Description
	keys between one or more IIS servers.

Export Configuration Dialog Box

Use the **Export Configuration** dialog box to export IIS configuration files and encryption keys to a specific location.

Element Name	Description
Physical path	Describes the location to export configuration files and encryption keys. This can be a network path or a folder on the local computer.
Connect As	Enables access to the **Set Credentials** dialog box. By default, you access the **Physical path** using whatever user account you are currently logged in with. To connect using a different user account, use the **Set Credentials** dialog box.
Encryption keys password	Specifies the password to encrypt the encryption keys.
Confirm password	Specifies that the user must retype the password.

Set Credentials Dialog Box

Use the **Set Credentials** dialog box to connect to a physical path using a user account different from the user account that you are currently logged in with.

Element Name	Description
User name	Specifies the user name for the user account you want to use to access a physical path for shared configuration.
Password	Specifies the password needed to authenticate the **User name**.
Confirm password	Specifies that the user must retype the password.

Encryption Keys Password Dialog Box

Use the **Encryption Keys Password** dialog box to enter the password for the encryption keys archive.

Element Name	Description
Enter encryption key password	Specifies the password needed for IIS to decrypt the encrypted encryption keys. This password is set when you use IIS Manager to export configuration files.

Sites

Use the **Sites** feature page to manage a list of sites on a Web server.

Sort the list by clicking one of the column headings of the feature page or select a value from the **Group by** drop-down list to group similar items. Search the list by typing search criteria in the **Filter** box, expand the **Go** list to select the field on which to search, and then click **Go** or press ENTER. Click **Show All** to remove the filter criteria.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

In this document

- UI Elements for Sites
- Add Website Dialog Box
- Edit Site Dialog Box
- Select Application Pool Dialog Box
- Connect As Dialog Box
- Set Credentials Dialog Box
- Site Bindings Dialog Box
- Add or Edit Site Binding Dialog Box
- Advanced Settings
- Edit Website Limits Dialog Box
- Edit Website Failed Request Tracing Settings Dialog Box

UI Elements for Sites

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Name	Displays the unique name of the site.
ID	Displays the unique ID of the site.
Status	Displays whether the site is started, or displays error information if the site has a problem and cannot start. When a site uses more than one protocol, the status is displayed separately for each protocol, with the status followed by the protocol in parentheses. For example, a site that uses the HTTP or HTTPS protocol will display **Started (http)** when it is started.
Binding	Displays the binding or bindings for the site.
Path	Displays the physical path of the directory that holds the site's content.

Actions Pane Elements

Some of the actions described in the following table are available for a site only when the site has an HTTP or HTTPS binding. These actions appear in the **Actions** pane under **Manage Web Site**.

Element Name	Description
Add Website	Opens the **Add Website** dialog box from which you can add a website that uses the HTTP or HTTPS protocols for communication.
Set Website Defaults	Opens the **Web Site Defaults** dialog box from which you can set default values that apply to all websites that you add to the Web server.
Bindings	Opens the **Site Bindings** dialog box from which you can add, edit, and remove site bindings.

Element Name	Description
Basic Settings	Opens the **Edit Site** dialog box from which you can edit the settings that were specified when the selected site was created. **Note** You cannot edit bindings from the **Edit Site** dialog box. Open the bindings dialog box to edit the site's bindings.
Explore	Opens the physical directory (in Windows Explorer) for the site.
Edit Permissions	Opens the Windows Properties dialog box for the physical directory that is mapped to the root virtual directory for the selected site's root application.
Remove	Removes the item that is selected from the list on the feature page.
Rename	Enables the **Name** field of the selected site so that you can rename the site.
View Applications	Opens the **Applications** feature page from which you can view the applications that belong to the site.
View Virtual Directories	Opens the **Virtual Directories** feature page from which you can view the virtual directories that belong to the site's root application.
Restart	Stops and restarts the selected site. Restarting a site causes the site to be temporarily unavailable until the restart is complete.
Start	Starts the selected site.
Stop	Stops the selected site. Stopping a site causes the site to be unavailable until it is started.
Browse	Opens the selected site in an internet browser. When the site has more than one binding, multiple browse links display.
Advanced Settings	Opens the **Advanced Settings** dialog box from which you can configure advanced settings for the selected site.

Element Name	Description
Failed Request Tracing	Opens the **Edit Web Site Failed Request Tracing** dialog box from which you can enable failed request tracing and configure related settings for the selected site.
Limits	Opens the **Edit Web Site Limits** dialog box from which you can configure bandwidth and connection limits for the selected site.

Add Website Dialog Box

Use the **Add Website** dialog to add a website to your Web server. Websites have bindings that consist of a port number, an IP address, and an optional host name or names.

Element Name	Description
Site name	Type a friendly name for your site. The friendly name appears in the tree of IIS Manager and can be changed later if you want to rename the site. **Note** The site name must be a unique name, because you cannot have sites that have the same name on an IIS 8 Web server.
Application pool	Displays the selected application pool for the site. When you type a website name, a new application pool is created with the same name as the site (up to 64 characters maximum). The application pool is configured to use .NET Framework version 2.0 and the integrated pipeline mode by default. You can edit the settings for this application pool later or select a different application pool for your site.
Select	Opens the **Select Application Pool** dialog box from which you can select the application pool in which the site will run.
Physical path	Type the physical path where the site content is stored. The content can reside on the local

Element Name	Description
	computer or come from a remote directory or share. If your content is stored on the local computer, type the physical path, such as C:\Content. If your content is stored in a remote share, type the Universal Naming Convention (UNC) path, such as \\Server\Share. Optionally click **Connect as** to specify credentials to connect to the physical path. If you do not provide credentials, the web server uses pass-through authentication. This means that content is accessed by using the application user's identity, and configuration files are accessed by using the application pool's identity.
…	Opens the **Browse for Folder** dialog box from which you can select the physical location of your site content.
Connect as	Opens the **Connect As** dialog box from which you can select how to connect to the path that you typed in the **Physical path** box. By default, **Application user (pass-through authentication)** is selected.
Test Settings	Opens the **Test Settings** dialog box from which you can view a list of test results to assess whether the path settings are valid.
Type	Select HTTP if you want the website to have an HTTP binding, or select HTTPS if you want the website to have a Secure Sockets Layer (SSL) binding.
IP address	Select an IP address from the **IP address** list or type an IP address that users can use to access this site. **Note** The percent symbol (%) is not a valid character for IP addresses. If you type an IPv6 address with this character, Windows Process Activation Service (WAS) is unable to start the site

Element Name	Description
	because the binding is invalid. For example, the link-local IPv6 address ends with "%8". If you copy this address from an **ipconfig** query in a command window and paste it into the **IP address** box in IIS Manager, make sure to remove the "%8" portion of the IP address. If you select **All Unassigned**, this site will respond to requests for all IP addresses on the port and optional host name that you specify for this site, unless another site on the server has a binding on the same port but with a specific IP address. For example, the Default website binding specifies **All Unassigned** for **IP address**, 80 for **Port**, and no host name. If the server has a second site named Contoso with a binding that specifies 172.30.189.179 for IP address on port 80 and no host name, Contoso receives all HTTP requests to port 80 on IP address 172.30.189.179, and Default Website continues to receive HTTP requests to port 80 on any IP address other than 172.30.189.179.
Port	Type the port on which HTTP.sys must listen for requests made to this website. If you select HTTP from the **Type** drop-down list, the default port is 80; if you select HTTPS from the drop-down list, the default port is 443. If you specify a port different from the default ports, clients must specify the port number in requests to the server or they will not connect to the website.
Host name	Type a host name if you want to assign one or more host names, also known as domain names, to one computer that uses a single IP address. If you specify a host name, clients must use the host name instead of the IP address to access the website. If this website is available on the Internet, type the domain name of the website as users type it

Element Name	Description
	in a browser, for example, www.contoso.com. If your website has more than one domain name, such as www.contoso.com and contoso.com, create a separate binding for each host name. If your website is available on an intranet, you do not have to specify a host name if users type the server name in a browser, for example, http://server_name. However, if the DNS server in your environment is configured to store other names for this Web server, you can create a separate binding for each host name so that users can use the other names stored by the DNS server. **Note** The host name that you configure for your website might appear differently in the log files because HTTP.sys logs host names by using the Punycode-encoded format.
Require Server Name Indicator	Determines whether the website requires Server Name Indication (SNI). SNI extends the SSL and TLS protocols to indicate what hostname the client is attempting to connect to. It allows multiple secure websites with different certificates to use the same IP address. This check box is displayed when the binding type is **https**.
Use Centralized Certificate Store	Determines whether website certificates are retrieved from the centralized certificate store. Select the check box to look up certificates in the centralized certificate store. Clear the check box to look up certificates locally. This check box is displayed when the binding type is **https** and centralized certificates are enabled.
SSL certificate	Select the certificate that you want the website to use for SSL. The **SSL certificate** list appears only when you select **HTTPS** from the **Type** list.

Element Name	Description
	Note To add or configure server certificates, use the **Server Certificates** feature in IIS Manager.
View	Opens the **Certificate** dialog box from which you can view information about the selected certificate in the **SSL certificate** list. The **View** button appears only when you select **HTTPS** from the **Type** list and select a certificate from the **SSL certificate** list.
Start Web site immediately	Select this option to start the website immediately after it is created. When the website is started, it listens for requests. You might want to delay starting the website if you must make other configuration changes, such as changes to authentication or authorization, after the site is created. **Note** If you specify an invalid binding, the website is created but is not able to start until the binding is fixed.

Edit Site Dialog Box

Use the **Edit Site** dialog box to edit the properties of an existing site.

Note

To edit binding information for the site, use the **Site Bindings** dialog box.

Select Application Pool Dialog Box

Use the **Select Application Pool** dialog box to select an application pool for your site or application. An application pool is configured to use one version of the .NET Framework, and one of two modes for processing managed code. Select an application pool that is configured to use the version of the .NET Framework and the managed-processing mode that the site or application requires. If your application does not require the .NET Framework, you can select an application pool that is configured for no managed code.

Note

To add an application pool, use the **Add Application Pool** dialog box on the **Application Pools** page.

Element Name	Description
Application pool	Select the application pool that you want your site or application to use.
Properties	Displays the .NET Framework version and the pipeline mode that the application pool is configured to use. **Note** The pipeline mode is not displayed when the application pool is configured to use .NET Framework version 1.1.

Connect As Dialog Box

Use the **Connect As** dialog box to select the method that IIS uses to access content from the physical path that was specified for the site, application, or virtual directory. The physical path can be a path to a directory on the local computer or to a directory or share on a remote computer.

Element Name	Description
Specific user	Select this option to provide credentials for a specific user account that has access to the physical path.
Set	Opens the **Set Credentials** dialog box from which you can type the user name and password for the user account that has access to the physical path. This user is used to access all of the content in the specified physical path. This button is available only when you select the **Specific user** option in the **Connect As** dialog box.
Application user (pass-through authentication)	Select this option if you want to use pass-through authentication. With this option, IIS uses the credentials of the requesting user to access the physical path. For anonymous requests, IIS uses the identity configured for the anonymous authentication to

Element Name	Description
	access the physical path. By default, this identity is the built-in IUSR account. For authenticated requests, IIS uses the authenticated set of credentials from the requesting user to access the physical path. Make sure that the application pool identity that serves this application has Read access to the physical path, so that the authenticated user can access the content in the physical path.

Set Credentials Dialog Box

Use the **Set Credentials** dialog box to provide user credentials that can access content from the physical path that was specified for the site, application, or virtual directory. The physical path can be a path to a directory on the local computer or to a directory or share on a remote computer.

Element Name	Description
User name	Type the user name of a Windows user account that has access to the physical path.
Password	Type the password for the user account that you specified in the **User name** box.
Confirm password	Retype the password for the user account that you specified in the **Password** box.

Site Bindings Dialog Box

Use the **Site Bindings** dialog box to configure bindings for a site.

Element Name	Description
Type	Displays the protocol for each site binding.
Host Name	Displays the host name, if any, for each site binding.
Port	Displays the port number for each site binding.
IP Address	Displays the IP address for each site binding. If an IP address is not specified, the IP Address field displays an asterisk (*), which means that

Element Name	Description
	the binding is for all IP addresses.
Binding Information	Displays binding information for site bindings that use protocols other than: • HTTP • HTTPS • FTP
Add	Opens the **Add Site Binding** dialog box from which you add a binding to the site.
Edit	Opens the **Edit Site Binding** dialog box from which you can edit the selected binding.
Remove	Removes the selected binding from the site. This button is available only when a site has more than one binding.
Browse	Opens the site's default page in a web browser by using the selected binding. **Note** This action is available only for HTTP and HTTPS bindings.
Close	Closes the **Site Bindings** dialog box.

Add or Edit Site Binding Dialog Box

Use the **Add Site Binding** and the **Edit Site Binding** dialog boxes to add bindings to a site or to edit an existing site binding. For example, your site might need an HTTPS binding on port 443 in order to use Secure Sockets Layer (SSL) in addition to an HTTP binding on port 80. You can access the **Add Site Binding** dialog box by clicking **Add** from the **Site Bindings** dialog box.

Note
If you add a duplicate binding to the Web server, only one site with that binding can run at a time. Additionally, if your duplicate binding is an HTTPS binding, any changes that are made to the certificate on one binding will affect the certificate on the other binding.

Element Name	Description
Type	Select a value from the list to specify the protocol for the site binding.
IP address	Select an IP address from the **IP address** list or

Element Name	Description
	type an IP address that users can use to access this site. **Note** The percent symbol (%) is not a valid character for IP addresses. If you type an IPv6 address with this character, Windows Process Activation Service (WAS) is unable to start the site because the invalid binding is invalid. For example, the link-local IPv6 address ends with "%8". If you copy this address from an **ipconfig** query in a command window and paste it into the **IP address** box in IIS Manager, make sure to remove the "%8" portion of the IP address. If you select **All Unassigned**, this site will respond to requests for all IP addresses on the port and optional host name that you specify for this site, unless another site on the server has a binding on the same port but with a specific IP address. For example, the Default website binding specifies **All Unassigned** for **IP address**, 80 for **Port**, and no host name. If the server has a second site named Contoso with a binding that specifies 172.30.189.179 for IP address on port 80 and no host name, Contoso receives all HTTP requests to port 80 on IP address 172.30.189.179, and Default Website continues to receive HTTP requests to port 80 on any IP address other than 172.30.189.179.
Port	Type the port on which HTTP.sys must listen for requests made to this site. If you select HTTP from the **Type** drop-down list, the default port is 80; if you select HTTPS from the drop-down list, the default port is 443. If you specify a port different from the default ports, clients must specify the port number in requests to the server or they do not connect to the site.

Element Name	Description
Host name	Type a host name if you want to assign one or more host names, also known as domain names, to one computer that uses a single IP address. If you specify a host name, clients must use the host name instead of the IP address to access the site. If this site is available on the Internet, type the domain name of the site as users type it in a browser, for example, www.contoso.com. If your site has more than one domain name, such as www.contoso.com and contoso.com, create a separate binding for each host name. If your site is available on an intranet, you do not have to specify a host name if users type the server name in a browser, for example, http://server_name. However, if the DNS server in your environment is configured to store other names for this Web server, you can create a separate binding for each host name so that users can use the other names stored by the DNS server. **Note** The host name that you configure for your Web site might appear differently in the log files because HTTP.sys logs host names by using the Punycode-encoded format.
SSL certificate	Select the certificate that you want the site to use for SSL. The **SSL certificate** list appears only when you select **HTTPS** from the **Type** list. **Note** To add or configure server certificates, use the **Server Certificates** feature in IIS Manager.
View	Opens the **Certificate** dialog box from which you can view information about the selected certificate in the **SSL certificate** list. The **View** button appears only when you select **HTTPS** from the **Type** list and select a certificate from

Element Name	Description
	the **SSL certificate** list.

Advanced Settings

Use the **Advanced Settings** dialog box to change the properties of the following objects in IIS:

- Application pool
- Website
- FTP site
- Web application
- Virtual directory

The **Advanced Settings** dialog box is useful when you want to modify a specific property or group of properties for an object without using an edit wizard or a dialog box, or when a property does not display in a wizard or dialog box.

Note
If you want to change the default value for all application pools, websites, web applications, or virtual directories, you can use the appropriate **Defaults** dialog box for any one of these objects.

For information about a specific property, select the property in the **Advanced Settings** dialog box and view the information that displays in the property description pane underneath the property grid. You can make the property description pane larger by clicking the bar between the property grid and the property description pane, and then dragging upward. You can also resize the entire **Defaults** dialog box to see more properties at the same time.

Edit Website Limits Dialog Box

Use the **Edit Web Site Limits** dialog box to configure performance settings for your website based on bandwidth usage and connection limits. For example, by restricting either bandwidth or the number of connections, or both, on a low-priority website, you enable other, higher-priority sites to handle larger traffic loads. You can adjust these settings as network traffic and usage changes.

Element Name	Description
Limit bandwidth usage (in bytes)	Select this option to limit the amount of traffic allowed to a website based on bandwidth usage. In the corresponding box, enter a value (in bytes) at which you want to limit the website traffic. The value must be an integer between 1024 and 4294967295 (unlimited).

Element Name	Description
Connection time-out (in seconds)	Type a number in the box to set the length of time (in seconds) before the Web server disconnects an inactive user. This setting guarantees that all connections are closed if the HTTP protocol cannot close a connection.
Limit number of connections	Select this option to limit the number of connections allowed to a website. In the corresponding box, enter the number of connections to which you want to limit the website. The value must be an integer between 0 and 4294967295 (unlimited). Setting the number to be unlimited circumvents constant administration if your connections tend to fluctuate. However, system performance can be negatively affected if the number of connections exceeds your system resources. Restricting a website to a specified number of connections can keep performance stable.

Edit Website Failed Request Tracing Settings Dialog Box

Use the **Edit Website Failed Request Tracing Settings** dialog box to change previously configured tracing rules for failed requests.

Element Name	Description
Enable	Select this option to enable failed request trace logging for a specific site.
Directory	Type the path where you want to store the failed request trace log files. **Note** As a best practice, log files, such as log files for failed request tracing, should be stored in a directory other than system root.
Maximum number of trace files	Type the maximum number of failed request trace log files that you want to keep for the site.

Element Name	Description
	If no value is entered, the default maximum number of log files that is stored is 50. When this limit is reached, new log files are created and older ones are removed.

SMTP E-mail

Use the **SMTP E-mail** feature page to configure delivery of e-mail from applications that use the System.Net.Mail API. You can configure the sender e-mail address on outgoing e-mail, and decide to either deliver e-mail immediately to an online SMTP server or store e-mail on disk where it can be retrieved for later delivery.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for SMTP E-mail
- Set Credentials Dialog Box

UI Elements for SMTP E-mail

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
E-mail address	Sets the e-mail address from which e-mail messages are sent.
Deliver e-mail to SMTP server	Configures System.Net.Mail to deliver e-mail messages immediately to an SMTP server.
SMTP Server	Defines the unique name of the SMTP server.
Use localhost	Sets the name of the SMTP server to LocalHost. Setting the name to LocalHost means that System.Net.Mail sends e-mail to

Element Name	Description
	the SMTP server on the local computer.
Port	Sets the TCP port that System.Net.Mail uses in trying to connect to the target SMTP server. Port 25 is the SMTP standard TCP port and is the default setting and the recommended setting.
Authentication Settings	Specifies the authentication mode and credentials if your SMTP server requires these.
Not required	Requires no user name or password from the user in order to connect to the SMTP server.
Windows	Uses the application identity. Use this option if you are setting up a smart host to relay messages in your network or if you are setting up a server to receive e-mail from internal sites, such as a company intranet.
Specify credentials	Requires a valid user name and password. Credentials are sent as clear text across the network. To specify credentials, click **Set**. Enter the valid user name in the **User name** text box and the valid password in the **Password** text box. Re-enter the password in the **Confirm Password** text box and then click **OK**.
Store e-mail in pickup directory	Configures System.Net.Mail to store e-mail messages in a file location on disk for later delivery by an application such as an ASP.NET application, or by a user, such as an administrator.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Set Credentials Dialog Box

Use the **Set Credentials** dialog box to configure IIS to use a Simple Mail Transfer Protocol (SMTP) server user name and password.

Element Name	Description
User name	Type the name of the account that you want IIS to use to access your SMTP server.
Password	Type the password associated with your SMTP server.
Confirm password	Retype the password associated with your SMTP server.

SSL Settings

Use the **SSL Settings** page to manage data encryption of transmissions between your server and clients. Additionally, by selecting **Ignore**, **Accept**, or **Require** certificates you can require a client to be identified before gaining access to content.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build an FTP Site on IIS
- Build a Web Farm with IIS Servers

UI Elements for SSL Settings

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Require SSL	Select **Require SSL** (Secure Sockets Layer) to enable a 40–bit data encryption method that you can use to help secure transmissions

Element Name	Description
	between your server and clients. This option setting works in both intranet and internet environments.
Require 128 Bit SSL	Select **Require 128 bit SSL** to provide stronger encryption than the 40–bit version. You can use 128–bit SSL to help secure transmissions between your server and clients in either an intranet or internet environment.
Ignore	This option is the default. This setting does not accept client certificates if they are provided. **Note** This option does not require clients to verify their identity before gaining access to your content. Therefore, this option is the least secure of these settings.
Accept	Select this setting if you want to accept client certificates (if they are provided), and to verify client identity before allowing the client to gain access to content.
Require	Select this option to require that certificates verify client identity before allowing the client to gain access to content.

Actions Pane Elements

Element Name	Description
Apply	Saves the changes that you have made on the feature page.
Cancel	Cancels the changes that you have made on the feature page.

Virtual Directories

Use the **Virtual Directories** feature page to manage a list of virtual directories in an application.

Sort the list by clicking one of the feature page column headings.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Virtual Directories
- Add or Edit Virtual Directory Dialog Box

UI Elements for Virtual Directories

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Application Path	Displays the application that contains each virtual directory.
Virtual Path	Displays the virtual path that is used in a URL to access the virtual directory.
Physical Path	Displays the physical path of the directory that holds the virtual directory's content.
Identity	Displays the user name of the custom identity (if configured) that is used to access content from the physical directory that is mapped to the virtual directory. If the identity is left blank, pass-through authentication is used to access the content.

Actions Pane Elements

The following table describes the actions available when connected to a server. When you are connected to a site or an application and select a virtual directory on the **Virtual Directories** page, the only available action is **Browse**.

Element Name	Description
Add Virtual Directory	Opens the **Add Virtual Directory** dialog box in which you can add a virtual directory.
Set Virtual Directory Defaults	Opens the **Virtual Directory Defaults** dialog box in which you can configure default settings that apply to all virtual directories at the current level and below.
Explore	Opens the physical directory (in Windows Explorer) that is mapped to the selected virtual directory.
Edit Permissions	Opens the Windows Properties dialog box for the physical directory that is mapped to the selected virtual directory.
Remove	Removes the item that is selected from the list on the feature page.
Browse	Opens the selected virtual directory in an internet browser.
Basic Settings	Opens the **Edit Virtual Directory** dialog box in which you can edit the settings that were specified when the selected virtual directory was created.
Advanced Settings	Opens the **Advanced Settings** dialog box in which you can configure advanced settings for the selected virtual directory.

Add or Edit Virtual Directory Dialog Box

Use the **Add Virtual Directory** and **Edit Virtual Directory** dialog boxes to add and edit virtual directories in your sites and applications. *Virtual directories* are pointers to physical content stored on your local computer or in shares on remote computers. If you want to include content that does not physically reside in your site's or application's directory, you can create a virtual directory that includes content from elsewhere on your Web server or from another computer in your network.

Element Name	Description
Site name	Displays the name of the site that contains the virtual directory.
Path	Displays the application that contains the virtual directory. If you create a virtual directory at the site level, this box will display **/**. If you create a virtual directory at the application level, this box displaya the name of the application, such as /myApp.
Alias	Type a name for the virtual directory that clients can use to access content from a web browser. For example, if your site address is http://www.contoso.com/ and you create a virtual directory for this site called /marketing, users can access the virtual directory from their web browser by typing http://www.contoso.com/marketing/.
Physical path	Type or navigate to the physical path where the virtual directory content is stored. The content can reside on the local computer or come from a remote share. If your content is stored on the local computer, enter the physical path, such as C:\Content. If your content is stored in a remote share, enter the UNC path, such as \\Server\Share. The path that you specify must exist, or you might receive a configuration error. Optionally, click **Connect as** to provide credentials for an account that is authorized to access content in the physical path.
…	Opens the **Browse for Folder** dialog box from which you can select the physical location of your virtual directory content.
Connect as	Opens the **Connect As** dialog box from which you can select how to connect to the path that you typed in the **Physical path** box. By default, **Application user (pass-through authentication)** is selected.
Test Settings	Opens the **Test Settings** dialog box from which you can view a list of test results to assess

Element Name	Description
	whether the path settings are valid.

WebDAV Authoring Rules

You can use the **WebDAV Authoring Rules** page to manage the list of authoring rules that control access to content. The rules appear on a list, and you can change the order of the rules to grant access to some users and deny access to others at the same time. Additionally, you can use the **WebDAV Authoring Rules** page to view information about other rules such as Path, Users, Roles, or Access.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS
- Build a Web Farm with IIS Servers

In this document

- UI Elements for WebDAV Authoring Rules
- Add Authoring Rule Dialog Box
- WebDAV Settings Page

UI Elements for WebDAV Authoring Rules

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Path	Displays the file name or file name extension for which the handler processes a response. The value can be either a specific file name, such as WebResource.axd, or a wildcard (*) with a specific extension, such as *.exe.
Users	Displays the user type, user names, or user groups to whom the rule applies.
Roles	Displays the Microsoft Windows role or .NET

Element Name	Description
	roles to which the rule applies, for example, the Administrator role.
Access	Displays the access permissions to which this rule applies, for example, Read or Write.
Entry Type	Displays whether the item is local or inherited.

Actions Pane Elements

Element Name	Description
Enable WebDAV	Enables the WebDAV feature.
Disable WebDAV	Disables the WebDAV feature.
Add Authoring Rule	Opens the **Add Authoring Rule** dialog box so that you can create an authoring rule.
Edit	Enables you to modify the selected rule.
Remove	Removes the selected rule.
Move Up	Moves up a selected item in the list.
Move Down	Moves a selected item down in the list.
WebDAV Settings	Opens the **WebDAV Settings** page so that you can change the settings for the WebDAV feature.

Add Authoring Rule Dialog Box

Use the **Add Authoring Rule** or **Edit Authoring Rule** dialog box to define rules for access to content.

Element Name	Description
All content	Select this option to apply the authoring rule to all content.
Specified content	Select this option to apply the authoring rule only to content that matches the rule. Type a file name extension or file name with an extension for which you want the rule to apply

Element Name	Description
	for WebDAV requests. For example, if you want the rule to apply for all WebDAV requests made for ASP files, enter *.asp. If you want the rule to apply to all requests made for a specific file, enter the file name and its extension, such as Myfile.asp.
All users	Select this option to manage access to content for all users.
Specified roles or user groups	Select this option to manage access to content for specific Microsoft Windows roles or user groups. **Note** If you use this rule, all members of the specified roles and groups must have valid user accounts and passwords before they are able to authenticate.
Specified users	Select this option to manage access to content for a specific user account. **Note** If you use this option, all specified users must have valid user accounts and passwords before they are able to authenticate.
Read	Specifies whether the users, groups, or roles have read access.
Source	Specifies whether the users, groups, or roles have access to the source for files that are mapped to handlers.
Write	Specifies whether the users, groups, or roles have write access.

WebDAV Settings Page

Use the **WebDAV Settings** page to modify the settings for the WebDAV feature.

Feature Page Elements

Element Name	Description
Allow Unknown MIME Types	Specifies whether access is allowed to content types that do not have explicit MIME types defined.
Default MIME Type	Specifies the default MIME type to use for content types that do not have explicit MIME types defined.
Allow Locks	Specifies whether WebDAV file locking is enabled. The default is **True**.
Lock Store	Specifies the name of the provider to use for WebDAV locking. The default is **webdav_simple_lock**.
Require Lock for Writing	Specifies whether locking is required before clients can save changes to files on the server. The default is **False**.
Allow Anonymous Property Queries	Specifies whether internet clients can request properties anonymously.
Allow Custom Properties	Specifies whether custom properties are supported. **Note** File system properties are always supported (for example, file size and modified date)
Allow Property Queries with Infinite Depth	Specifies whether internet clients can request properties for content that is deeper than the current parent and child collection. **Note** Property queries with infinite depth can consume a great deal of system processing time before they complete.
Property Store	Specifies the provider to use when internet clients request custom properties.
Allow File Name Extension Filtering	Specifies whether to block WebDAV requests based on file name extensions. The default is

Element Name	Description
	True.
Allow Hidden Segment Filtering	Specifies whether to block WebDAV requests for files in the `hiddenSegments` collection. The default is **True**.
Allow Verb Filtering	Specifies whether to block WebDAV requests based on the HTTP verb. The default is **True**.
Allow Hidden Files to be Listed	Specifies whether content that is marked as hidden on the server are returned in file listings.
Allow Multiple Property Sections	Specifies whether WebDAV properties can be specified in multiple XML <prop></prop> sections. **Note** Some internet clients use this syntax.
Return Compact XML	Specifies whether the WebDAV extension should suppress sending a carriage-return/line-feed sequence after each line of XML.
Return IsCollection Property	Specifies whether the WebDAV extension should return the IsCollection property for content. **Note** Some internet clients use this property.
Return IsHidden Property	Specifies that the WebDAV extension should return the IsHidden property for content. **Note** Some internet clients use this property.
Return MS-Author-Via Header	Specifies whether the WebDAV extension should return the MS-Author-Via HTTP header. **Note** Some Microsoft WebDAV clients use this header.
Require SSL Access	Specifies that all WebDAV requests must use SSL.

Actions Pane Elements

Element Name	Description
WebDAV Authoring Rules	Opens the **WebDAV Authoring Rules** feature page so that you can modify the list of WebDAV authoring rules.

Worker Processes

Use the **Worker Processes** feature page to manage a list of worker processes running in application pools on a Web server. Sort the list by clicking one of the feature page column headings.

Related scenarios

- Build a Static Website on IIS
- Build a Classic ASP Website on IIS
- Build an ASP.NET Website on IIS
- Build a PHP Website on IIS

In this document

- UI Elements for Worker Processes
- Requests Page

UI Elements for Worker Processes

The following tables describe the UI elements that are available on the feature page and in the **Actions** pane.

Feature Page Elements

Element Name	Description
Application Pool Name	Displays the name of the application pool. In Web gardens or if you have a worker process that is shutting down at the same time another worker process is starting, the same application pool can be listed more than once in the grid to account for different worker processes running in that application pool.
Process ID	Displays the worker process identifier (ID) for

Element Name	Description
	that worker process.
State	Displays the state of a process, such as starting, running, or stopping.
CPU %	Displays the percentage of time that the worker process has used the CPU since the data was last updated. This option corresponds to **CPU Usage** in Windows Task Manager.
Private Bytes (KB)	Displays the current size of memory this is committed by a worker process. This memory cannot be shared with other processes. This option corresponds to **Virtual Memory Size** in Windows Task Manager.
Virtual Bytes (KB)	Displays the current size of the virtual address space for a worker process.

Actions Pane Elements

Element Name	Description
View Current Requests	Opens the **Requests** feature page in which you can view the requests currently processing in the selected application pool.

Requests Page

Use the **Requests** feature page to manage a list of currently running requests in a worker process. This information can help you determine where a problem occurs in a specific area of a site or application.

Sort the list by clicking one of the feature page column headings. Search the list by typing search criteria in the **Filter** box, expand the **Go** list to select the field on which to search, and then click **Go** or press ENTER. Click **Show All** to remove the filter criteria.

Element Name	Description
Web Site ID	Displays the site identifier (ID) for the specific request.
Url	Displays the requested URL.

Element Name	Description
Verb	Displays the HTTP verb used in the request.
Client IP	Displays the IP address of the client that made the request.
State	Displays the current pipeline module state that the request is in.
Module Name	Displays the current module that the request is in.
Time Elapsed	Displays the amount of time that the request has been in process.

IIS Manager UI (IIS 8)

IIS 8 provides a UI, IIS Manager, which lets you:

- Manage IIS and ASP.NET in one tool.
- View health and diagnostic information that includes the ability to view currently running requests in real time.
- Configure user and role authorization for sites and applications.
- Delegate site and application configuration to non-administrators.
- Connect to a web server, site, or application remotely by using HTTP.

In this document

- Navigation Toolbar
- Connection Manager and Tree
- Workspace
- Actions Pane

Navigation Toolbar

IIS Manager contains a toolbar that you can use to navigate in the UI. Depending on your location or selection in the UI, the following buttons are available:

- **Back**: goes back one page view.
- **Forward**: goes forward one page view.
- **Refresh Page**: updates the view of the current selected object in the UI.
- **Stop**: stops the current action in the UI from completing.
- **Home**: navigates to the home page of the current selection, such as the web server home page or a site home page.

- **Help**: opens list of links to Help documentation.

Connection Manager and Tree

IIS 8 has a **Connections** pane and toolbar that lets you connect to web servers, sites, and applications. When you connect to a web server, site, or application from the **Connections** pane, the tree loads the connection along with any children of the parent connection. For example, if you connect to a site, the tree displays the site connection and the applications within that site.

The **Connections** toolbar contains the following buttons that let you interact with the tree:

- **Create New Connection**: starts the connection wizard to connect to a web server, a site, or an application.
- **Save Current Connections**: saves the connection information for the current list of connections in the tree.
- **Up**: goes up one level in the hierarchy.
- **Delete Connection**: removes the selected connection in the tree. If you click this button when you are on a node other than the connection node, such as a physical folder, the parent connection is removed from the list.

Workspace

The workspace is the area between the **Connections** pane and the **Actions** pane. The workspace has two views: **Features View** and **Content View**.

When you select **Features View**, you can configure features for an object, such as a site or an application. For example, when you select a site in the tree and then double-click the **Modules** feature on the site's Home page, you can configure the modules for that site.

When you select **Content View**, you can view the actual content of the object that you select in the tree. For example, when you select a site in the tree, the workspace displays the contents of that site, such as its physical directories, virtual directories, and files. If you select an object in the list, such as a file, you can click **Switch to Features View** in the **Actions** pane or in the right-click menu to configure features for that object. You can also view the Windows properties of content from the **Actions** pane when you are in **Content View**.

Actions Pane

You can use the **Actions** pane to configure IIS, ASP.NET, and IIS Manager settings. For example, you can use the **Actions** pane to open dialog boxes and wizards that let you complete tasks in IIS Manager, such as creating a site, adding a connection string for an application, or configuring recycling settings for an application pool. Additionally, you can complete other tasks from the **Actions** pane, such as starting and stopping a site or the web server, or delegating a feature.

Note

Items in the **Actions** pane are also available from the shortcut menu when you right-click an object in IIS Manager.

Procedures

This topic includes the following procedures:

Open IIS Manager (IIS 8)

Manage Connections in IIS Manager (IIS 8)

Navigation in IIS Manager (IIS 8)

Start or Stop the Web Server (IIS 8)

View the Contents of a Site, Application, or Directory (IIS 8)

Open IIS Manager (IIS 8)

You can open IIS Manager if you want to configure IIS 8 settings by using the user interface (UI). This procedure assumes that you are viewing the **Control Panel** in the default **Category** view.

Note

IIS is not enabled be default when you install a Windows operating system. To use the IIS Manager, you must install or enable IIS on your computer. For information about installing IIS, see **IIS 8 Installation and Deployment**.

To open IIS Manager from the Start screen

1. On the Start screen, click **Control Panel**.
2. Click **System and Security**, and then click **Administrative Tools**.
3. In the **Administrative Tools** window, double-click **Internet Information Services (IIS) Manager**.

To open IIS Manager from the Desktop.

1. On the desktop, hover the mouse cursor in the upper right corner of the screen, and then click **Settings**.
2. Click **Control Panel**.
3. Click **System and Security**, and then click **Administrative Tools**.
4. In the **Administrative Tools** window, double-click **Internet Information Services (IIS) Manager**.

To open IIS Manager from the Search box

1. On the desktop, hover the mouse cursor in the upper right corner of the screen, and then click **Search**.

2. In the **Search** box, type **inetmgr** and press ENTER.

See Also

IIS Manager UI (IIS 8)

Manage Connections in IIS Manager (IIS 8)

You can connect remotely to web servers, sites, and applications by using IIS Manager. Remote connections let server administrators configure IIS and ASP.NET settings without having to log on to the physical server or by using Terminal Server. By using delegation in IIS 8, a server administrator can delegate specific features to a site or an application. Permitted users can then connect to that site or application, and they can configure the features that are delegated to that site or application.

This task includes the procedures described in this topic.

Connect to a web server by specifying connection details manually

By using IIS Manager, you can connect to a web server remotely to configure IIS settings. You can use this method if you cannot log on to the web server directly or the server is not configured to use Terminal Server for remote administration.

Important
Only web server administrators on a web server can remotely connect to the server by using IIS Manager. Additionally, the server administrator must enable remote connections and must specify whether to accept only Windows credentials or both Windows and IIS Manager user credentials.

To connect to a web server by using IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).
2. In the **Connections** pane, expand **Create New Connection** in the toolbar.
3. Click **Connect to a Server** to open the **Connect to Server** Wizard.
4. On the **Specify Server Connection Details** page, in the **Server name** box, type the name of the server to which you want to connect. If you must specify a port number over which to connect, type the port number with a colon after the server name. For example, type **WESTSRV01:8080**. To connect to the local computer, type **localhost** instead of the server name.
5. Click **Next**. On the **Provide Credentials** page type a user name and password in the **User name** and **Password** boxes.

6. Click **Next**.
7. On the **Specify a Connection Name** page, in the **Connection Name** box, type a friendly name for your connection and then click **Finish**. This friendly name appears in the tree after the connection is made.

Connect to a Site by Using IIS Manager

By using IIS Manager, you can connect remotely to a site to configure IIS settings. Connecting remotely is helpful if you are an IIS Manager user or a Windows user who wants to configure your site through the UI. Remote connection is also convenient if you are a server administrator who wants to save changes to the site in the site's Web.config file instead of in the server-level configuration files, ApplicationHost.config and the root Web.config.

Note

To log on to a site as an IIS Manager user or a Windows user, you must have been given the appropriate permissions and provided with logon information by a server administrator. Additionally, the server administrator must enable remote connections and must specify whether to accept only Windows credentials or both Windows and IIS Manager user credentials.

To connect to a site by using IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).
2. In the **Connections** pane, expand **Create New Connection** in the toolbar.
3. Click **Connect to a Site** to open the **Connect to Site** Wizard.
4. On the **Specify Site Connection Details** page:
   - In the **Server name** box, type the name of the server that hosts the site to which you want to connect. If you must specify a port number over which to connect, type the port number with a colon after the server name. For example, type **WESTSRV01:8080**. If you connect to a site that is hosted on the local computer, type **localhost** instead of the server name.
   - In the **Site name** box, type the friendly name of the site to which you want to connect. For example, type **Default Web Site**.
5. Click **Next**.
6. On the **Provide Credentials** page, type a user name and password in the **User name** and **Password** boxes.
7. Click **Next**.
8. On the **Specify a Connection Name** page, in the **Connection Name** box, type a friendly name for your connection and then click **Finish**. This friendly name appears in the tree after the connection is made.

Connect to an Application by Using IIS Manager

By using IIS Manager, you can connect remotely to an application to configure IIS settings. Connecting remotely is helpful if you are an IIS Manager user or a Windows user who wants to configure your web application through the UI, or if you are a server administrator who wants to save changes to the application in the application's Web.config file instead of in the server-level configuration files, ApplicationHost.config and the root Web.config.

Note

To log on as an IIS Manager user or a Windows user, you must have been given the appropriate permissions and provided with logon information by a server administrator. Additionally, the server administrator must enable remote connections and must specify whether to accept only Windows credentials or both Windows and IIS Manager user credentials.

To connect to an application by using IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).
2. In the **Connections** pane, expand **Create New Connection** in the toolbar.
3. Click **Connect to an Application** to open the **Connect to Application** Wizard.
4. On the **Specify Application Connection Details** page:
   - In the **Server name** box, type the name of the server that hosts the site to which you want to connect. If you must specify a port number over which to connect, type the port number with a colon after the server name. For example, type **WESTSRV01:8080**. If you connect to an application that is hosted on the local computer, type **localhost** instead of the server name.
   - In the **Site name** box, type the friendly name of the site that contains the application to which you want to connect.
   - In the **Application name** box, type the friendly name of the application to which you want to connect.
5. Click **Next**.
6. On the **Provide Credentials** page, type a user name and password in the **User name** and **Password** boxes. These credentials can be for either a Windows account or an IIS Manager user account.
7. Click **Next**.
8. On the **Specify a Connection Name** page, in the **Connection Name** box, type a friendly name for your connection and then click **Finish**. This friendly name appears in the tree after the connection is made.

Save a List of Current Connections in IIS Manager

When you have several connections open in IIS Manager, and you want to save those connections so that they return when you reopen the tool, you can save that current list of

connections. When you reopen the tool, the saved connections display in the tree in the **Connections** pane, but you must expand a connection so that it actually opens the connection.

Note
You must connect to a web server, site, or application before it is added to the tree in the **Connections** pane.

To save a list of current connections in IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).
2. Connect to any of the web servers, sites, or applications for which you want to save connections.
3. In the **Connections** pane, click **Save Current Connections** in the toolbar.

See Also

IIS Manager UI (IIS 8)

Navigation in IIS Manager (IIS 8)

In IIS 8, server administrators and non-administrators (IIS Manager users and Windows users) can connect remotely to a Web server by using IIS Manager. The procedures in this topic explain how to navigate to different levels in IIS Manager after you have connected remotely to a Web server, site, or application.

Server administrators can configure all settings and navigate to all areas of the user interface (UI). Non-administrators can configure any features delegated to the site or application to which they are allowed to connect, and can navigate only to areas of the UI that are related to those delegated features.

This task includes the procedures described in this topic.

Navigate to a web server in IIS Manager

When you want to configure settings at the web server, site, and application levels in IIS 8, you must be a server administrator. To be a server administrator, the following must be true:

- When you are logged on to the local server, your Windows account must be part of the **Administrators** group.
- When you are connecting remotely through IIS Manager, provide your Windows administrative credentials and the server name.

To navigate to a web server in IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager

(IIS 8).

2. If you are not already connected to the server you want to configure, connect by using the tools in the **Connections** pane.
3. In the **Connections** pane, click the server node in the tree.
4. On the *Server* **Home** page, double-click a feature to configure the feature for the server.

Navigate to a site in IIS Manager

Server administrators and non-administrators can configure settings for sites in IIS Manager. Non-administrators must be granted permission to a site and provided with logon details by a server administrator. Then the user can provide credentials to connect to the site, and configure settings for that site and any applications within that site.

Note

Server administrators have access to all sites and applications on the web server.

To navigate to a site in IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).
2. In the **Connections** pane:
   - If you are a server administrator, expand the **Sites** node and click a site in the tree.
   - If you are non-administrator who has permission to connect to a site, use the tools in the **Connections** pane to connect to the site that you want to manage. Then click the site in the tree.
3. In **Features View**, on the *Site* **Home** page, double-click a feature to configure that feature for the site.

Navigate to an Application in IIS Manager

Server administrators and non-administrators can configure settings for applications in IIS Manager. Non-administrators must be granted permission to an application and provided with logon details by a server administrator. Then the user can provide credentials to connect to that application and then can configure settings for the application.

Note

Server administrators have access to all sites and applications on the web server. Non-administrators who connect to a site have access to delegated features in that site and any applications within that site.

To navigate to an application in IIS Manager

1. Open IIS Manager. For information about opening IIS Manager, see Open IIS Manager (IIS 8).

2. In the **Connections** pane:
   - If you are a server administrator, expand the **Sites** node. Expand the site that contains the application that you want to manage, and then click the application in the tree.
   - If you are a non-administrator who has permission to connect to a site, use the tools in the **Connections** pane to connect to the site that contains the application that you want to manage. Expand the site in the tree, and then click the application.
   - If you are a non-administrator who has permission to connect to an application, use the tools in the **Connections** pane to connect to the application that you want to manage. Then click the application in the tree.
3. On the application **Home** page, double-click a feature to configure the feature for the application.

See Also

IIS Manager UI (IIS 8)

Manage Connections in IIS Manager (IIS 8)

Start or Stop the Web Server (IIS 8)

You can stop IIS when you want to stop the web server from serving content, such as when you upgrade applications. To stop the web server, you must stop the Windows Process Activation Service (WAS) and the World Wide Web Publishing Service (W3SVC). You can then start the web server when you are ready for the web server to serve content again.

When you stop WAS and W3SVC, all sessions connected to your web server are dropped. Any in-memory session state is lost. All sites are unavailable until these services are restarted. Therefore, avoid stopping and restarting your web server if you can.

Note

You must be a server administrator to perform this procedure.

To start or stop a web server

You can perform this procedure by using the user interface (UI) or a command line. For information about opening IIS Manager, see Open IIS Manager (IIS 8).

Using the UI

1. Open IIS Manager and navigate to the web server node in the tree.
2. In the **Actions** pane, click **Start** if you want to start the web server, **Stop** if you want to stop the web server, or **Restart** if you want to first stop IIS, and then start it again.

Using a command line

1. Open an elevated command-line window.
2. At the command prompt, type **net stop WAS** and press ENTER; type **Y** and then press ENTER to stop W3SVC as well.
3. To start the web server, type **net start W3SVC** and press ENTER to start both WAS and W3SVC.

Note

If you want to stop IIS first, and then start it again, type **iisreset /restart** and press ENTER.

See Also

IIS Manager UI (IIS 8)

View the Contents of a Site, Application, or Directory (IIS 8)

In IIS 8, you can use **Content View** to view the contents of sites, applications, virtual directories, and physical directories. For example, you can view a list of all the applications, physical directories, virtual directories, and files in a site.

Additionally, you can use **Content View** to locate a file, such as an HTML file, to configure settings for the file from the **Features View**. For example, you might have a web page that occasionally does not load correctly. You can locate the specific file in **Content View**, and then change to **Features View** to configure logging and tracing settings for that particular page.

Note

Content View is the only way to locate and configure settings for files in IIS Manager. Sites, applications, and physical and virtual directories are all available in the tree in the **Connections** pane.

To view the contents of a site, an application, or a directory

1. Open IIS Manager and navigate to the level you want to manage. For information about opening IIS Manager, see Open IIS Manager (IIS 8). For information about navigating to locations in the UI, see Navigation in IIS Manager (IIS 8).
2. Click **Content View** at the bottom of the workspace area.
3. Select an item in the list to view additional information, or right-click an item and click **Switch to Features View** to change to **Features View** and configure IIS settings for the item.

See Also

IIS Manager UI (IIS 8)

Security Best Practices for IIS 8

This document contains a list of recommendations for improving the security of your IIS 8 web server. While following these recommendations does not guarantee freedom from security issues, these recommendations can significantly reduce your risk.

The recommendations are grouped I the following categories:

- Installation and Configuration
- Web Application Isolation
- Authentication
- Request Filtering
- Application Pool Identities
- More Security Practices

Installation and Configuration

- Do not run IIS on a domain controller or a backup domain controller.

  First, there are no local accounts on a domain controller. Local accounts are important to the security of many IIS server installations. Placing an IIS web server and domain controller on the same computer seriously limits your security account options. Second, any new exploit that compromises your web server could also compromise your entire network when the web server and the domain controller are on the same computer.

- Install only the IIS modules you need.

  IIS 8 is composed of more than 40 modules, which allow you to add modules you need and remove any modules you don't need. If you install only the modules you need, you reduce the surface area that is exposed to potential attacks.

- Periodically remove unused or unwanted modules and handlers.

  Look for modules and handlers that you no longer use and remove them from your IIS installation. Strive to keep your IIS surface area as small as possible.

- For high volume installations of IIS, run other resource-intensive products like SQL Server or Exchange on separate computers.

- Keep you antivirus software up to date.

  Install and run the latest version of antivirus software on the server.

- Move the `Inetpub` folder from your system drive to a different drive.

  By default IIS 8 sets up the `Inetpub` folder on your system drive (usually the C drive). If you move the folder to a different partition, you can save space on your system drive and improve security. For information about how to the `Inetpub` folder, see the following blog post: Moving the INETPUB directory to a different drive.

Web Application Isolation

- Isolate web applications.

  Separate different applications into different sites with different application pools.

- Implement the principle of least privilege.
  Run your worker process as a low privileged identity (virtual application pool identity) that is unique per site.

- Isolate ASP.NET temp folders.

  Set up a separate ASP.NET temp folder per site and only give access to appropriate process identity.

- Isolate content.

  Make sure to set an ACL (access control list) on each site root to allow only access to the appropriate process identity.

Authentication

- If you use Windows authentication, turn on extended protection.

  Extended protection protects against credential relaying and phishing attacks when using Windows authentication. For more information about extended protection and how to turn it on in IIS, see Configure Extended Protection in IIS 7.5.

- Be aware that configuring Anonymous authentication along with another authentication type for the same website can cause authentication problems.

If you configure Anonymous authentication and another authentication type, the result is determined by the order in which the modules run. For example, if Anonymous authentication and Windows authentication are both configured and Anonymous authentication runs first, Windows authentication never runs.

- Disable anonymous access to server directories and resources.

  When you want to grant a user the access to server directories and resources, use an authentication method that is not anonymous.

- Do not allow anonymous writes to the server.

  Authenticate the user with a method that is not anonymous before allowing the user to upload anything to you website or FTP site.

Request Filtering

- Ensure that request filtering rules are enabled.

  Request filters restrict the types of HTTP requests that IIS 8 processes. By blocking specific HTTP requests, request filters help prevent potentially harmful requests from reaching the server. The request filter module scans incoming requests and rejects requests that are unwanted based upon the rules that you set up. Both websites and FTP sites should have the protection that request filter rules provide. For more information about request filtering, see Configure Request Filtering in IIS.

- Ensure that request limits are set to reasonable values.

  Think carefully about the values you assign to configuration parameters. For example, make sure that an upper limit value is higher than a lower limit value. Otherwise, the filter may never trigger.

Application Pool Identities

- Don’t use the built-in service identities (such as Network Service, Local Service, or Local System).

  For maximum security, application pools should run under the application pool identity that is generated when the application pool is created. The accounts that are built in to IIS are ApplicationPoolIdentity, NetworkService, LocalService, and LocalSystem. The default (recommended) and most secure is ApplicationPoolIdentity.

- Using a custom identity account is acceptable, but be sure to use a different account for each application pool.

More Security Practices

- Make periodic backups of the IIS server.

  Do a complete system-state backup every day or two. Also do it before major software upgrades or configuration changes.

- Limit permissions granted to non-administrators.

  Look for folders that non-administrators have write permissions and script execution permissions to, and remove the permissions.

- Turn on SSL and maintain SSL certificates.

  Renew the certificate or choose a new certificate for the site. An expired certificate becomes invalid and can prevent users from accessing your site.

- Use SSL when you use Basic authentication.

  Use Basic authentication with an SSL binding, and make sure that the site or application is set to require SSL. Alternatively, use a different method of authentication. If you use Basic authentication without SSL, credentials are sent in plaintext that might be intercepted by malicious code. If you want to continue using Basic authentication, you need to check the site bindings to make sure that an HTTPS binding is available for the site, and then configure the site to require SSL.

- When you set feature delegation rules, don’t make rules that are more permissive than the defaults.

- For a classic ASP application, turn off debug mode.

Appcmd.exe (IIS 8)

IIS 8 provides a command-line tool, Appcmd.exe, to configure and query objects on your web server, and to return output in text or XML. The following examples are tasks that you can complete with Appcmd.exe:

- Create and configure sites, applications, application pools, and virtual directories.
- Start and stop sites.
- Start, stop, and recycle application pools.

- View information about worker processes and requests that are running on the web server.

Appcmd.exe provides consistent commands for common queries and configuration tasks, which reduces the complexity of learning syntax. For example, you can use the **list** command to search for information about an object, such as an application, or you can use the **add** command to create an object, such as a site.

Additionally, you can combine commands together to return more complex data about objects on your web server or to perform tasks that are more complex. For example, you can complete complex tasks like stopping all sites that have requests that have been running longer than 60 seconds.

In this document

- Start Appcmd.exe
- Get help about Appcmd.exe
- Working with objects in Appcmd.exe
- Working with commands in Appcmd.exe
- Working with attributes and values in Appcmd.exe
- Errors and the AppHostConfig attribute
- Configure server-level settings by using Appcmd.exe
- Configure settings for a site, application, virtual directory, or URL by using Appcmd.exe
- Use Appcmd.exe to change configuration in a Web.config file
- Use Appcmd.exe to change configuration in a parent-level configuration file

Start Appcmd.exe

Open Appcmd.exe if you want to configure IIS 8 settings by using a command-shell window. Appcmd.exe resides in the %windir%\system32\inetsrv directory, which is available only to the Administrator account or to users who are members of the Administrators group on the computer. In addition, members of the Administrators group must start Appcmd.exe with elevated permissions to view and change settings in the server-level configuration files: Machine.config, the root Web.config file (in the .NET Framework directory), and in ApplicationHost.config. Users who view or change Web.config files in site or application directories must have access to read and write to files in those directories.

Note

If you are running 64-bit Windows, use Appcmd.exe from the %windir%\system32\inetsrv directory, not the %windir%\syswow64\inetsrv directory.

To start Appcmd.exe

1. Click **Start**, and then click **All Programs**.
2. Click **Accessories**, and then click **Command Prompt**.
3. At the Command Prompt, type **cd %windir%\system32\inetsrv**, and then press ENTER.

Note
To run the Command Prompt as an administrator, right-click the icon in the **Start** menu, and then click **Run as administrator**.

Get help about Appcmd.exe

Appcmd.exe contains built-in help information. To view help information about Appcmd.exe, type a command and then type **/?**.

For example, if you want to view introductory help information about Appcmd.exe, type the following command at the command prompt, and then press ENTER:

appcmd /?

To view help information about a specific Appcmd.exe object, such as the commands that you can use with the **site** object, type the following command at the command prompt, and then press ENTER:

appcmd site /?

To view help information about the settings that you can configure by using an Appcmd.exe object and command, such as the **site** object and **add** command, type the following at the command prompt, and then press ENTER:

appcmd add site /?

Working with objects in Appcmd.exe

Use objects in Appcmd.exe to specify what you want to view or change in configuration. To view a list of Appcmd.exe objects, type the following command at the command prompt, and then press ENTER:

appcmd /?

You can then configure the object by using commands. For more information about commands, see Working with commands in Appcmd.exe.

Working with commands in Appcmd.exe

Use Appcmd.exe commands to view or change configuration settings for an Appcmd.exe object. Each of these objects has commands that enable you to work with that object. To view the commands for an object, use the following syntax:

appcmd *object* **/?**

The variable *object* is the name of the object for which you want to view commands. For example, to view a list of commands for the **site** object, type the following command at the command prompt, and then press ENTER:

appcmd site /?

For more information about Appcmd.exe objects, see Working with objects in Appcmd.exe.

Working with attributes and values in Appcmd.exe

Specify values for attributes when you want to limit the results of a list command, or when you want to change the value of an attribute for an Appcmd.exe object.

For example, to list application pools configured to use Network Service, type the following command at the command prompt, and then press ENTER:

appcmd list apppool /processModel.identityType:NetworkService

Alternatively, to configure an application pool named Marketing to use Network Service, type the following command at the command prompt, and then press ENTER:

appcmd set apppool /apppool.name:Marketing /processModel.identityType:NetworkService

The **processModel.identityType** attribute is used in both of the previous examples. In the first example, the attribute limits the query results by listing only application pools that are configured to use Network Service. In the second example, the attribute specifies that the Marketing application pool use the Network Service identity.

Errors and the AppHostConfig attribute

In Windows 8 and Windows Server 2012, the `AppHostConfig` attribute was added. The `AppHostConfig` attribute enables you to edit an AppHost configuration file that is not the default AppHost configuration file. Although this attribute is useful in many situations, invalid values are ignored under the following verb/object combinations:

Verb	Object
Start or Stop	Site
Start, Stop, or Recycle	AppPool
All verbs	Backup
All verbs	Request
All verbs	WP

Configure server-level settings by using Appcmd.exe

In IIS 8, you can configure settings in server-level configuration files or in Web.config files that are in a folder for a site, application, or virtual directory. Server-level configuration files are Machine.config, the root Web.config file, and ApplicationHost.config. These three files are available only to the Administrator account or to users who are members of the Administrators group on the web server. In addition, members of the Administrators group must open a command-shell window that uses their elevated administrator permissions to read and change

settings in the server-level configuration files. For more information about starting Appcmd.exe in a command-shell window that uses elevated administrator permissions, see Start Appcmd.exe.

By default, ApplicationHost.config is the file from which Appcmd.exe reads configuration and to which Appcmd.exe writes changes. When you want to read changes from or write changes to Machine.config or the root Web.config file, include a *commit* parameter in your command. Specify MACHINE for the *commit* parameter when you want to affect the Machine.config file or specify WEBROOT when you want to affect the root Web.config file.

For example, to list the machineKey section in the root Web.config file, type the following command at the command prompt, and then press ENTER:

appcmd list config /section:machineKey /commit:WEBROOT

Configure settings for a site, application, virtual directory, or URL by using Appcmd.exe

In IIS 8, you can configure settings in server-level configuration files or in Web.config files for a site, application, or virtual directory on your web server. Users who view or change settings in Web.config files in site or application directories must have read and write access to files in those directories. Additionally, you must delegate features that you want to configure at the site, application, or virtual directory level in IIS 8. For more information about configuration files and delegating features, see **Working With Configuration Files in IIS 7.0** and **Delegate Web Site and Web Application Management**.

Use Appcmd.exe to change configuration in a Web.config file

When you use Appcmd.exe to configure settings in a Web.config file for a specific site, application, or virtual directory, specify the path of the site, application, or virtual directory for which you want the configuration to apply. For example, to disable the default document feature for a web site named **Contoso**, type the following command at the command prompt, and then press ENTER:

appcmd set config "Contoso" /section:defaultDocument /enabled:false

In the previous example, the configuration changes are made in the Web.config file in the **Contoso** directory, and all the **Contoso** applications and virtual directories inherit the configuration from that Web.config file. However, if you want to enable default documents for a specific application in the **Contoso** site, you can use Appcmd.exe to provide the path of the application. For example, to enable default documents for an application named **Gallery**, type the following command at the command prompt, and then press ENTER:

appcmd set config "Contoso/Gallery" /section:defaultDocument /enabled:true

The command in the previous example enables default documents for the **Gallery** application in the **Contoso** site, and writes the configuration changes to the Web.config file in **Gallery's** directory.

Use Appcmd.exe to change configuration in a parent-level configuration file

The previous section explained how to use Appcmd.exe to configure features for a specific site, application, or virtual directory by writing changes to a Web.config file. You can also configure features for a specific site, an application, a virtual directory, or even a URL by using location tags in a parent configuration file.

When you want to configure settings for a site, application, or virtual directory but you want to store the settings in a parent configuration file, such as in ApplicationHost.config instead of in a Web.config file, use the Appcmd.exe **commit** parameter with your command. By using the examples in the previous section, you can now write the configuration changes to ApplicationHost.config instead of to the site and application Web.config files.

First, to disable default documents for the web site named **Contoso**, type the following command at the command prompt, and then press ENTER:

appcmd set config "Contoso" /section:defaultDocument /enabled:false /commit:APPHOST

The previous command disables default documents for the **Contoso** site, and writes the configuration settings with a location tag to ApplicationHost.config. The location tag specifies that the configuration applies to the **Contoso** site. The configuration section resembles the following example:

```
<location path="Contoso">

      <system.webServer>

            <defaultDocument enabled="false" />

      </system.webServer>

</location>
```

Second, to enable default documents for the application named **Gallery**, type the following command at the command prompt, and then press ENTER:

appcmd set config "Contoso/Gallery" /section:defaultDocument /enabled:true /commit:APPHOST

The previous command enables default documents for the **Gallery** application, and writes the configuration settings with a location tag to ApplicationHost.config. The location tag specifies that the configuration applies to the **Gallery** application. The configuration section looks like the following example:

```
<location path="Contoso/Gallery">

      <system.webServer>

            <defaultDocument enabled="true" />

      </system.webServer>

</location>
```

Configuration Store

IIS 8 uses a new XML-based configuration store that is modeled after the ASP.NET configuration. IIS configuration is stored in the ApplicationHost.config file and can also be distributed among Web.config files for sites, applications, and directories. Settings configured at one level are inherited automatically by lower levels, unless they have been locked to prevent changes. By default, the server administrator is the only user who has permission to view and edit the ApplicationHost.config file.

In this document

- Working with configuration files in IIS 8.0
- Copy configuration files to a remote server

Working with configuration files in IIS 8.0

In IIS 8, the configuration system enables you to customize web server administration by delegating configuration sections from server-level configuration files to Web.config files in your sites, applications, and virtual or physical directories.

Delegating configuration in IIS 8.0

In IIS 8, the locking and unlocking of sections in configuration is known as *feature delegation*, and these actions enable you to define at what level IIS reads and writes configuration. By default, all sections in configuration are locked and can be configured only in server-level configuration files except for the following sections (also referred to as *features*):

- Default document
- Directory browsing
- HTTP protocol
- HTTP redirect

You can unlock more configuration sections to let non-administrators view and change configuration for the unlocked sections in Web.config files for specific sites, applications, directories, or even URLs. For example, you might unlock the connectionStrings section to let non-administrators configure settings to connect their applications to databases. Alternatively, if you want only administrators to view and change configuration, you can lock all configuration sections and store settings only in server-level configuration files.

Configuration levels

In IIS 8, you can configure settings at the following levels:

- Web server
- Site
- Application
- Virtual or physical directory

- URL (also known as file-level configuration)

Note
To configure settings at a child level, a configuration section must be unlocked (also known as delegated) at the parent level or levels. For example, to configure a feature at the application level, the related configuration section or sections must be delegated at both the server and the site levels.

Configuration files

Configuration exists in a physical directory in either server-level configuration files or in Web.config files. Every configuration file maps to a specific site, application, or virtual directory.

Server-level configuration is stored in the following configuration files:

- Machine.config. This file is located in %windir%\Microsoft.NET\Framework*framework_version*\CONFIG.
- Root Web.config for the .NET Framework. This file is located in %windir%\Microsoft.NET\Framework*framework_version*\CONFIG.
- ApplicationHost.config. This file is located in %windir%\system32\inetsrv\config.

Site, application, and virtual and physical directory configuration can be stored in one of the following locations:

- A server-level configuration file. When configuration for a site, application, directory, or URL is stored in a server-level configuration file, use a location tag to specify the site, application, directory, or URL to which the configuration applies.
- A parent-level Web.config file. When configuration for an application, directory, or URL is stored in a parent-level configuration file, use a location tag to specify the child at which the configuration applies.
- The Web.config file for the site, the application, or the directory. When you configure settings for an application, directory, or URL, the configuration is stored in the same directory as the site, application, or directory. You do not need to use location tags.

Storing configuration settings in a parent configuration file is helpful when:

- You want to store configuration settings in a configuration file that is accessible by only certain users or groups. For example, the ApplicationHost.config file is available only to the Administrator account and to the members of the Administrators group on a specific computer, as well as to domain administrators when a computer is part of a domain.
- You want to configure a feature at the URL-level (also known as file-level).

Inheritance

Sites, applications, virtual directories, physical directories, and files inherit configuration settings from the server-level configuration files and from any parent configuration files. For example, an application inherits configuration settings from the server-level configuration files and from the site's root application Web.config file. The following table illustrates the configuration files in which you can store an object's configuration.

Note
Configuration sections must be delegated if configuration files other than server-level configuration files are used.

Configuration Object	Server-Level Configuration Files	Parent-Level Configuration Files	Other Configuration Files	Comments
Web server	Machine.config Root Web.config ApplicationHost.config	Not applicable	Not applicable	No location tags are needed. Server-level configuration can only exist in server-level configuration files.
Site	Machine.config Root Web.config ApplicationHost.config	Not applicable	Site's root application Web.config	Use location tags to store site configuration in server-level configuration files.
Application	Machine.config Root Web.config ApplicationHost.config	Site's root application Web.config	Application's Web.config	Use location tags to store application configuration in server-level configuration files or a parent configuration file.
Virtual or Physical Directory	Machine.config Root Web.config ApplicationHost.config	Parent site's root application Web.config Parent application's Web.config	Directory's Web.config file	Use location tags to store directory configuration in server-level configuration files, or in a parent configuration

Configuration Object	Server-Level Configuration Files	Parent-Level Configuration Files	Other Configuration Files	Comments
				file.
URL (or file)	Machine.config Root Web.config ApplicationHost.config	Parent site's root application Web.config Parent application's Web.config Parent directory's Web.config	Not applicable	Use location tags to store URL-specific configuration in server-level and parent-level configuration files.

Copy configuration files to a remote server

IIS 8 uses XML files to specify web server, site, and application configuration settings. The primary configuration file for IIS 8 is ApplicationHost.config. IIS 8 also uses Web.config files for applications or directories. These files can be copied from one web server or site to another to apply the same settings to multiple objects.

Requirements

Levels

This procedure can be performed at the following levels in IIS:

- Web server
- Site
- Application
- Physical and virtual directories

Required Permissions

To perform this procedure, you must be a server administrator or you must have write access to the target location.

To copy and deploy a configuration file

You can perform this procedure by running **xcopy** commands in a command-line window.

Copy configuration to another location on the same computer

To locate help information about **xcopy**, type the following at the command prompt, and then press ENTER:

xcopy /?

To copy configuration folders and files to another location on the same computer, use the following syntax:

xcopy *source destination* **/E**

The variable *source* is the location from which you are copying the folders and files, and the variable *destination* is the location to which you are copying the folders and files. The parameter **/E** specifies that all the contents of the source location be copied to the destination location.

For example, to copy the IIS configuration folder and all its contents from the system root to another path on your local computer, type the following at the command prompt, and then press ENTER:

xcopy %windir%\system32\inetsrv\config c:\config /E

Copy configuration from one computer to another computer

To copy a source from your local computer to another computer, use the **net use** command to connect to the computer, and then use the **xcopy** command to copy the folder and files. To locate help about **net use**, type the following at the command prompt, and then press ENTER:

net use /?

For example, to copy the IIS configuration folder and all its contents from the system root to another computer, use the following syntax:

net use *computername**drive* *** /USER:***domain**username*

The variable *computername* is the name of the computer to which you are connecting, and the variable *drive* is the drive to which you want to connect. The parameter * prompts you for a password to connect to the computer and drive after you run the command. The variable *domain* is the domain to which the user belongs, and the variable *username* is the user name of the user who is connecting to the computer.

Once you connect to the computer to which you want to copy files, you can use the **xcopy** syntax as specified in the previous section. For example, to connect to another computer and copy all the IIS configuration folders and files from the system root on your local computer to the other computer, type each of the following commands at the command prompt, and then press ENTER:

net use \\contoso\c$ * /USER:SALES\joseph

xcopy %windir%\system32\inetsrv\config \\contoso\c$\config /E

After you are finished copying folders and files to the remote computer, close the connection. To do this type the following at the command prompt, and then press ENTER:
net use \\contoso\c$ /delete

See Also

IIS Manager UI (IIS 8)

Windows Management Instrumentation (WMI)

IIS 8 uses Windows Management Instrumentation (WMI) to build scripts for Web administration. The IIS 8 WMI provider namespace, WebAdministration, contains classes and methods that enable you to script administration of web sites, web applications, and their associated objects and properties.

With appropriate delegation settings, you can use WMI scripts in IIS 8 to access, read, and modify key configuration files such as ApplicationHost.config and Web.confg. You can add or remove modules at the global or section level, and configure commonplace features such as custom HTTP errors, Mime Maps, SSL bindings, and ASP settings by using the IIS 8 WMI provider.

WMI scripting continues to be important in IIS 8 for automating frequently repeated tasks such as creating web sites, but it can also take advantage of the new features exposed by the Runtime Status and Control API (RSCA) to monitor and diagnose sites. For example, you can monitor HTTP requests, worker processes, application pools, and application domains by using WMI scripting. You can also access the new diagnostics and troubleshooting features in IIS 8, such as failed request tracing, by scripting the IIS 8 WMI provider.

See Also

IIS WMI Provider Reference

IIS Manager UI (IIS 8)

Windows Deployment Services Overview

Windows Deployment Services (WDS) enables you to deploy Windows operating systems over the network, which means that you do not have to install each operating system directly from a CD or DVD.

For information about what is new or included in your version of Windows Deployment Services, see What's New in Windows Deployment Services in Windows Server 2012.

For step-by-step guidance on how to use the Windows Deployment Services, see Windows Deployment Services Getting Started Guide for Windows Server 2012.

Benefits of Windows Deployment Services

1. Allows network-based installation of Windows operating systems, which reduces the complexity and cost when compared to manual installations.
2. Supports deploying images for mixed environments including Windows 7and Windows Server 2008 R2 through Windows 8.1and Windows Server 2012 R2.

3. Uses standard Windows Setup technologies including Windows Preinstallation Environment (Windows PE), .wim files, and image-based setup.
4. Transmits data and images by using multicast functionality.
5. Allows you to create images of a reference computer using the Image Capture Wizard, which is an alternative to the ImageX tool.
6. Allows you to add driver packages to the server and configure them to be deployed to client computers along with the install image.

Prerequisites for installing Windows Deployment Services

The following are requirements for installing this role, depending on whether you choose the default installation (both Deployment Server and Transport Server), or only the Transport Server role service.

Deployment Server and Transport Server	Transport Server
• **Active Directory Domain Services (AD DS)**. A Windows Deployment Services server must be a member of an Active Directory Domain Services (AD DS) domain or a domain controller for an AD DS domain. The AD DS domain and forest versions are irrelevant; all domain and forest configurations support Windows Deployment Services. **Note** AD DS is not required if the WDS server is configured in Standalone mode. • **DHCP**. You must have a working Dynamic Host Configuration Protocol (DHCP) server with an active scope on the network because Windows Deployment Services uses PXE, which relies on DHCP for IP addressing. • **DNS**. You must have a working Domain Name System (DNS) server on the network before you can run Windows Deployment	• For Windows Server 2008: The only prerequisite is that you must be a member of the Local Administrators group on the server to install Transport Server. Also note that a PXE provider is not installed with Transport Server, so you must create a custom PXE provider to network boot. • For Windows Server 2008 R2: You must be a member of the Local Administrators group to install Transport Server. In addition, if you are using Transport Server to network boot, your environment must contain DHCP (Windows Server 2008 R2 contains a PXE provider, which allows you to network boot). • For Windows Server 2012 and Windows Server® 2012 R2, there are no additional prerequisites.

Deployment Server and Transport Server	Transport Server
Services. • **NTFS volume**. The server running Windows Deployment Services requires an NTFS file system volume for the image store. • **Credentials**. To install the role, you must be a member of the Local Administrators group on the server. To initialize the server, you must be a member of the Domain Users group. **Note** To initialize the WDS server in Standalone mode, you need not be a member of the Domain Users group. For more information, see Required Permissions.	

Tools for managing Windows Deployment Services

The following are tools that you can use to manage your server:

7. 0.1. **Windows Deployment Services MMC snap-in** - Provides an easy way to manage images, computers, and common server settings. You can perform almost all tasks from the MMC snap-in (for example, you cannot prestage client computers). Note that the snap-in is not available when you install only the Transport Server role service.
2. **WDSUTIL command-line tool** - enables you to manage the full functionality of the server. WDSUTIL also enables you to script common tasks.
3. **Windows PowerShell cmdlets for WDS** – enables you to manage and monitor the WDS server using Windows PowerShell scripts.

Deploying and configuring WDS in a multi-computer environment

It is not common to split a WDS server's functionality among different computers.

Installing the WDS role on virtual machines

You can run Windows Deployment Services on Hyper-V virtual machines. Depending on the capacity of your server running Hyper-V, you can install, configure and use Windows Deployment Services. You can also create a virtual machine to act as a client machine when building and testing your deployment images. This is especially useful for building your custom images. Note, however, that the performance will often be degraded, particularly during the Trivial File Transfer Protocol (TFTP) download phase. This phase is very resource-intensive and may fail if insufficient resources are available on your server running Hyper-V.

Warning
WDS cannot be run on a virtual machine in Windows Azure.

Support for the WDS role in a clustered environment

The Windows Deployment Services server service is not cluster-aware. However, you can run multiple WDS servers in a network to provide fault-tolerance and load balancing.

Considerations for managing this role remotely

You can manage Windows Deployment Services remotely by installing the Windows Deployment Services server role on a local server and then using that server to connect to the remote server in the data center or you can use the Windows PowerShell cmdlets for Windows Deployment Services to manage a remote server.

Note
Windows Deployment Service cannot be managed using the Remote Server Administration Tools.

Managing the WDS role using the Server Core installation option

You cannot use Windows Deployment Services with the Server Core installation option.

Configuring the WDS role for high availability

WDS is designed to automatically fail over to other WDS servers. We recommend that you deploy two WDS servers at a minimum and then additional servers to increase redundancy for higher availability.

Note
Installing images from WDS over a wide area network is not recommended.

Known issues

When using the Dynamic Driver Provisioning (DDP) feature on a WDS server running Windows Server 2012 if you attempt to add a driver to a Windows 8.1 or Windows Server 2012

R2 boot.wim you may receive the following error when running the "Add driver packages to image" wizard:

An error occurred while trying to execute command. Error Code: 0xc0000135

This is a known issue related to DISM version included with Windows Server 2012 that is used by WDS to service images. This version of DISM supports servicing Windows Server 2012 and earlier boot images but does not support later versions of boot images. You can resolved this issue by upgrading your WDS server to Windows Server 2012 R2 which contains a DISM version that supports servicing Windows 8.1 and Windows Server 2012 R2 boot.wim's or you can install the Windows 8.1 ADK and manually add the driver to the boot.wim. For instructions for how to add drivers, see Add and Remove Drivers Offline Using DISM

Common usage scenarios

The following are common scenarios for Windows Deployment Services.

Scenario one: The small business

Fabrikam, Inc. is a manufacturer of towels with custom designs. It is a small business with a single office. Monica Brink, Fabrikam's resident IT professional, is responsible for maintaining the IT infrastructure for the company, which consists of 25 client computers running Windows XP SP2 Professional and a single server running Windows Server 2003 with SP2. The server functions as a file print server, Web server, Exchange server, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP) server, and domain controller. The computers are linked by a 100-MBps Ethernet connection.

Monica is given the task of moving all of the client computers to the Windows 7 (or Windows 8.1) operating system and upgrading the single server to Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2. It is important that she makes the switch as quickly as possible so that office workers can continue their work.

In the past, she deployed a new operating system one computer at a time. This took her around 45 minutes per computer (almost 19 hours to set up the operating system on all the client computers). For almost three days, Monica was unavailable to work on anything else. Then she would spend almost as much time installing the applications on each computer.

Monica is the only IT professional at Fabrikam, which means that she also must help teach users about the new operating system. Therefore, it is important that she minimizes the amount of time she spends on deployment. To accomplish this, Monica chooses to use Windows Deployment Services because she can:

- Save time by running several installations simultaneously.
- Use a custom install image with preinstalled applications.
- Create an image by using the Windows Deployment Services Image Capture Wizard.

To begin, Monica does the following:

1. Upgrades her server to Windows Server 2008 R2, or Windows Server 2012 or Windows Server 2012 R2.

2. Installs the Windows Deployment Services server role.
3. Adds the Boot.wim from the product DVD (which contains a Windows PE image, Setup.exe, and supporting files) by using the Windows Deployment Services MMC snap-in.
4. Adds the Install.wim from the product DVD to the Windows Deployment Services server by using the MMC snap-in.
5. Uses the MMC snap-in to create a capture image from the boot image she added in step 3. This image contains Windows PE and a wizard that will capture her custom image into a .wim file.

All users at Fabrikam have the same desktop hardware, which was purchased from a single vendor. To deploy a standard image to the computers, Monica does the following:

1. Boots a reference computer from the network and installs the Install.wim onto it, which contains the standard version of Windows.
2. Installs Microsoft Office, the company's towel-design application, and the latest drivers from the manufacturer's site.
3. Uses Sysprep to generalize the operating system.
4. Reboots the computer into the capture image.
5. Uses the Image Capture Wizard to recapture the operating system and upload it directly to the Windows Deployment Services server.

Now, Monica is ready to install the new operating systems. She does not need to migrate any user data, because all of the employees store their user data on a server (rather than on their hard disks). She reboots a client computer and then presses F12 to perform a network boot. This boots her into the Boot.wim file, which guides her through the installation process. She selects the disk partition and image she wants, and then the installation begins. While waiting for the image to be applied to the first computer, Monica boots another computer and starts the same process on that one.

Scenario two: The medium-sized business

Northwind Traders is a shipping firm with three offices: a central office in Tooth City, and branch offices in the towns of Brushville and Flosston. Ron Gable is one of six IT staff members at Northwind Traders. His responsibility is maintaining the 250 client computers used by the company's employees. These are mostly desktop computers, but the sales force uses laptops for customer presentations. There are 200 computers in the central office in Tooth City, and 25 each in the Brushville and Flosston offices. Each site has an internal network running at 100 MB per second (MBps), and the branch sites are connected to the Tooth City office by a T1 line. Ron has three Windows Deployment Services servers at the Tooth City office and one in each of the branch offices, which are administered remotely.

Ron's supervisor has tasked him with deploying Windows 7 (or Windows 8.1) to the whole company. Previously, this would have involved many expensive trips to Brushville and Flosston, and it would have taken Ron several weeks to complete. He wants to use Windows Deployment Services to deploy Windows remotely; however, company policy dictates that there can be only one DHCP server on the corporate network, and this server is located at the Tooth City office.

Remotely deploying images to the 50 computers at the branch offices would cause immense congestion on the connection.

Ron chooses to use Windows Deployment Services because with unattended setup, he can:

- Deploy Windows to computers at the branch sites without being physically present there.
- Use his existing replication solution to deliver images to the branch site servers.
- Use the network boot referral system to minimize network traffic between the branch sites and the central office.

Ron prestages each client computer and assigns the appropriate branch office server as the referral server for each. This ensures that the client downloads files and images from the local server, which minimizes traffic on the line between the offices.

Ron has two standard operating system configurations — one for the desktop computers and one for laptops that contains the sales presentations and drivers for projectors. Therefore, he builds two images: one with the desktop configuration, and one with the laptop configuration (with no applications). He stores all the user data on one of the servers, so he can deploy Windows without preserving any existing data on the client computers.

Ron uses Windows System Image Manager (Windows SIM) to author two image unattend files — one for the desktop computers and one for the laptops. These files automate the installation, so Ron does not need to be present at each computer during the installation. They also automatically install Microsoft Office and the line-of-business application that the company uses for package tracking. He uses the Windows Deployment Services management tools to associate the unattend files with the images.

Next, Ron configures Windows Deployment Services so when a computer is restarted, it will boot from the network automatically and deploy the appropriate image (without requiring the users to press F12). After the image is applied to each computer, the computer is automatically joined to the corporate domain and restarted. He also configures the deployment so that when the computer restarts, the computer will not network boot unless F12 is pressed. This way, the computer then boots to the hard disk drive and finishes the installation process. This prevents a boot loop, in which the computer would continue booting into Setup. When the installation is completed, the computer is ready for the user to log on.

Scenario three: The large enterprise

Shu Ito is the network architect for Wide World Importers, a large enterprise with 5,000 employees in offices all over the world. The major employee centers are in the United States and Germany, and there are 13 branch offices in other countries. Shu has five servers available to him in the U.S. hub, two in the German hub, and one in each of the branch offices. The servers at the hubs are connected to the corporate Ethernet on 1 Gbps network interface cards (NICs); the other computers are on 100-MBps NICs. The hubs are connected by T3 lines, and the other sites are connected by T1 lines. All of the servers are hired on two-year leases.

Wide World Importers is replacing the accounting department’s 200 computers with computers running Windows 7 (or Windows 8.1). Shu would also like to deploy a Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2 image to any newly leased servers in the U.S.

office. The servers in the German office and the branch sites are the responsibility of the local administrators. Currently, deployments at Wide World Importers are done by using RIS, and Shu wants to ensure that the existing computer building processes are preserved with the move to Windows Deployment Services. In addition, it is important that each computer is deployed with an operating system in a language that is appropriate for the users in that country or region.

Shu chooses to use Windows Deployment Services because it enables him to do the following:

- Use appropriate language packs to reduce the required number of images.
- Manage all of his Windows Deployment Services servers from a single computer.
- Use multicast deployments to preserve bandwidth while deploying images to many computers concurrently.
- Write scripts to automate common management tasks.

Shu upgrades his servers to Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 and creates his images. The vast majority of his deployments will be in English or German, so he creates Windows 7 (or Windows 8.1) image in each language. Other languages will be installed by using external language packs, and applications will be downloaded by using Systems Management Server (SMS). After creating the images, Shu uploads the images and language packs to the Windows Deployment Services server.

Shu authors unattend files with Windows SIM. He then uses File Replication Service (FRS) to copy the images, language packs, and unattend files to the Windows Deployment Services servers around the world. Of the accounting computers used by Wide World Importers, 150 are in the U.S. office, 30 are in the German office, and the remaining 20 are scattered around the world. Shu uses multicasting to deploy to the 150 computers in the U.S. office simultaneously. To do this, he creates a multicast transmission for the relevant image on his Windows Deployment Services server.

To preserve the state and data on the previous computers, Shu uses the User State Migration Tool (USMT) to save all of the data and user configurations to a shared folder on the primary Windows Deployment Services server. Then he sets up each computer to boot from its local Windows Deployment Services server and to start automated setup by using the unattend files. The computers in the U.S. office will automatically join the multicast transmission, while the computers in the other offices will deploy using unicasting. When the installation is completed, Shu runs a task with USMT to migrate the user data to each computer.

When the lease on a server expires and the server is replaced, Shu can use Windows Deployment Services to deploy his Windows Server images in the same way that he performed the RIS deployment.

Scenario Four: A Custom Deployment Using Transport Server

John Woods is the server maintenance engineer at the A. Datum Corporation data center. He is responsible for maintaining the 300 servers used by A. Datum Corporation's major customers. One of these customers is Adventure Works.

Adventure Works uses 40 servers to run a career Web site (which is backed by a database) for circus performers. After the release of a popular film about circus life, Adventure Works expects

an increase in the use of their Web site. They order 10 additional servers to handle the anticipated traffic.

John wants to deploy operating systems to these servers by using Windows Deployment Services. He does not have AD DS running in this environment; he wants to store the configuration information for his computers in a SQL Server database. As a result, he cannot use the standard Windows Deployment Services solution. In addition, he wants to partition the disks in a standard configuration and also copy data (some for database servers, some for Web servers) before the unattended setup begins. John chooses to use Windows Deployment Services because he can:

- Write a plug-in that reads configuration data for the computers from a data store other than AD DS (the data store is typically a database or a flat file).
- Write scripts (to run in Windows PE) that perform preinstallation tasks and then call Setup to install the operating system.

John creates 10 computer accounts in his database for his 10 new servers, and he populates them with the required information. He installs the Windows Deployment Services server role on his server (choosing to install only the Transport Server role service). He then writes a PXE provider (a plug-in that reads information from the database and passes it to Windows Deployment Services) and registers it with the server. He creates a custom boot image that contains Windows PE along with startup scripts to partition the disks and copy the data. Then he uses ImageX to capture one of his existing servers as an install image.

After performing these initial tasks, John connects his servers to the network and boots them. They boot into Windows PE by using the configuration stored in the database. His scripts run to prepare each computer for deployment, and the scripts end by running ImageX to apply the operating system image on each computer.

Additional references

For information about getting started with Windows Deployment Services, see Windows Deployment Services Getting Started Guide for Windows Server 2012.

For more information about Windows PowerShell cmdlets for WDS, see Windows Deployment Services Cmdlets in Windows PowerShell.

What's New in Windows Deployment Services in Windows Server 2012 R2

This document describes the updated Windows Deployment Services (WDS) features in Windows Server® 2012 R2.

For more information about new features for Windows Deployment Services introduced with Windows Server 2012, see What's New for Windows Deployment Services in Windows Server 2012.

For more information about Windows Deployment Services, see the Windows Deployment Services Overview.

Role description

Windows Deployment Services (WDS) is a server role that enables you to remotely deploy Windows operating systems. You can use it to set up new computers by using a network-based installation. This means that you do not have to install each operating system directly from a CD, USB drive or DVD. To use Windows Deployment Services, you should have a working knowledge of common desktop deployment technologies and networking components, including Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Active Directory Domain Services (AD DS). It is also helpful to understand the Preboot eXecution Environment (also known as Pre-Execution Environment).

New and changed functionality

The following table compares the features and functionality included in each version of Windows Deployment Services:

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
Operating systems deployed	Legacy mode: Windows 2000, Windows XP, Windows Server 2003 Mixed mode: Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008 Native mode: Windows 2000 Professional, Windows XP, Windows Server 2003, Windows Vista,	Windows 2000 Windows XP Windows Server 2003 Windows Vista Windows Server 2008 Windows 7 Windows Server 2008 R2	Windows XP Windows Server 2003 Windows Vista SP1 Windows Server 2008 Windows 7 Windows Server 2008 R2 Windows Server 2012 Windows 8 Windows 8.1 Windows Server 2012 R2	Windows XP Windows Server 2003 Windows Vista SP1 Windows Server 2008 Windows 7 Windows Server 2008 R2 Windows Server 2012 Windows 8 Windows 8.1 Windows Server 2012 R2

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
	Windows Server 2008, Windows 7, Windows Server 2008 R2			
Image types deployed	Legacy mode: RISETUP and RIPREP Mixed mode: RISETUP, RIPREP, .wim Native mode: .wim	For a new installation, only .wim images are supported. If you upgrade from Windows Server 2003, you can convert any RIPREP images to .wim format after the upgrade. However, RISETUP images are not supported.	.wim .vhd images of Windows Server 2008 R2 are supported as part of an unattended installation (managed by using the command line only). **Note** .vhd images are only used for booting from VHD scenarios, and they are not applied directly to volumes.	.wim and .vhd images are both supported similar to Windows Server 2008 R2. .vhd images can now be managed via the WDS management snap-in in addition to the command-line. In addition, .vhdx files (a new file format available in Windows Server 2012) are supported directly and over multicast. The .vhdx format supports a sparse, dynamic representation of a disk, which produces a relatively small disk image file (compared to directly storing the full contents of the disk). The .vhdx format is

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
				supported for VDH boot scenarios in addition to direct application to system volumes (similar to .wim deployments).
Boot environment	Legacy mode: OSChooser Mixed mode: OSChooser and Windows Preinstallation Environment (Windows PE) Native mode: Windows PE	Windows PE	Windows PE	Windows PE
Administration experience	Legacy mode: Remote Installation Services (RIS) toolset Mixed mode: RIS toolset to manage RISETUP and RIPREP images. Microsoft Management Console (MMC) snap-in and WDSUTIL to manage .wim images. Native mode: MMC snap-in and WDSUTIL	MMC snap-in and WDSUTIL	MMC snap-in and WDSUTIL	MMC snap-in and WDSUTIL. Windows PowerShell cmdlets added in Windows Server 2012 R2.

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
Multicasting	Not supported.	Supported for install images. (The Boot.wim file must be from the DVD for Windows Server 2003, Windows Vista with SP1, Windows Server 2008, Windows 7, or Windows Server 2008 R2.)	Supported for install images. (The Boot.wim file must be from the DVD for Windows Server 2008, Windows Vista with SP1, Windows 7, or Windows Server 2008 R2.) Supported for boot images for computers with EFI (managed by using the command line only). Provides the ability to automatically disconnect slow clients and divide transmissions into multiple streams based on client speeds (the boot image must be Windows 7 or Windows Server 2008 R2). Provides support for multicasting in environments that use IPv6 (the boot image must be from Windows Server 2008, Windows Vista with SP1, Windows 7, or Windows Server 2008 R2).	Provides support for TFTP and multicasting over IPv6 and DHCPv6. Improved multicast deployment by eliminating the need for making a local copy of the install.wim file. You can apply the install.wim file while it is being downloaded without significant impact to the application process. Applying the install.wim file concurrently with the multicast download decreases overall deployment time, offers higher disk write throughput, and requires less space on the client system. Reduced default block size fits within an Ethernet MTU and works on customer

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
				hardware that does not support IP fragmentation.
Driver provisioning	Not supported.	Not supported.	Provides the ability to deploy driver packages to client computers based on their hardware as part of an installation (the install image must be Windows Vista with SP1, Windows 7, or Windows Server 2008). Provides the ability to add driver packages to boot images (Windows 7 or Windows Server 2008 R2 images only).	Extends filters for driver groups to support manufacturer Model Number and **Device Groups**. Driver import automatically detects and prevents duplicate driver packages from being added to the driver store.
Extensibility	Supports custom PXE providers.	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). However, you must develop a way to boot clients. Supports custom multicast providers.	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). Includes a PXE provider to boot clients. Supports custom multicast providers.	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). Includes a PXE provider to boot clients. Supports custom multicast providers.

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012and Windows Server 2012 R2
Extensible Firmware Interface (EFI)	Supports network booting of Itanium-based computers with EFI.	Supports network booting of Itanium-based and x64-based computers with EFI/UEFI.	Supports network booting of x64-based computers with UEFI. Includes Auto-add functionality. Includes DHCP referral to direct clients to a specific PXE server. Provides the ability to deploy boot images by using multicasting.	Support x86 clients (32-bit processors) with Unified Extensible Firmware Interface (UEFI) to network boot and complete an end-to-end deployment by using WDS.
PowerShell cmdlet scripting	Not supported	Not supported	Not supported	Supported on Windows Server 2012 R2 only

Windows PowerShell cmdlets for WDS

A WDS server running Windows Server 2012 R2 can be managed using the Windows PowerShell cmdlets for WDS. Using Windows PowerShell cmdlets you can add driver packages, add client images, enable and disable boot and install images, and do many other common WDS tasks. For a full reference see Windows PowerShell cmdlets for WDS.

Additional references

For information about getting started with Windows Deployment Services, see Windows Deployment Services Getting Started Guide for Windows Server 2012.

For information about the Windows Assessment and Deployment Kit (Windows ADK), see Windows Deployment with the Windows ADK.

What's New in Windows Deployment Services in Windows Server 2012

This document describes the updated Windows Deployment Services (WDS) features in Windows Server 2012.

For more information about Windows Deployment Services, see the Windows Deployment Services Overview.

Role description

Windows Deployment Services (WDS) is a server role that enables you to remotely deploy Windows operating systems. You can use it to set up new computers by using a network-based installation. This means that you do not have to install each operating system directly from a CD, USB drive or DVD. To use Windows Deployment Services, you should have a working knowledge of common desktop deployment technologies and networking components, including Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Active Directory Domain Services (AD DS). It is also helpful to understand the Preboot eXecution Environment (also known as Pre-Execution Environment).

New and changed functionality

The following table compares the features and functionality included in each version of Windows Deployment Services:

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
Operating systems deployed	Legacy mode: Windows 2000, Windows XP, Windows Server 2003 Mixed mode: Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008 Native mode: Windows 2000	Windows 2000 Windows XP Windows Server 2003 Windows Vista Windows Server 2008 Windows 7 Windows Server 2008 R2	Windows XP Windows Server 2003 Windows Vista SP1 Windows Server 2008 Windows 7 Windows Server 2008 R2	Windows XP Windows Server 2003 Windows Vista SP1 Windows Server 2008 Windows 7 Windows Server 2008 R2 Windows Server 2012 Windows 8

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
	Professional, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2			
Image types deployed	Legacy mode: RISETUP and RIPREP Mixed mode: RISETUP, RIPREP, .wim Native mode: .wim	For a new installation, only .wim images are supported. If you upgrade from Windows Server 2003, you can convert any RIPREP images to .wim format after the upgrade. However, RISETUP images are not supported.	.wim .vhd images of Windows Server 2008 R2 are supported as part of an unattended installation (managed by using the command line only). **Note** .vhd images are only used for booting from VHD scenarios, and they are not applied directly to machine volumes.	.wim and .vhd images are both supported similar to Windows Server 2008 R2. .vhd images can now be managed via the WDS management snap-in in addition to the command-line. In addition, .vhdx files (a new file format available in Windows Server 2012) are supported directly and over multicast. The .vhdx format supports a sparse, dynamic representation of a disk, which produces a relatively small disk image file (compared to

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
				directly storing the full contents of the disk). The .vhdx format is supported for VDH boot scenarios in addition to direct application to system volumes (similar to .wim deployments).
Boot environment	Legacy mode: OSChooser Mixed mode: OSChooser and Windows Preinstallation Environment (Windows PE) Native mode: Windows PE	Windows PE	Windows PE	Windows PE
Administration experience	Legacy mode: Remote Installation Services (RIS) toolset Mixed mode: RIS toolset to manage RISETUP and RIPREP images. Microsoft Management Console (MMC) snap-in and WDSUTIL to manage .wim images.	MMC snap-in and WDSUTIL	MMC snap-in and WDSUTIL	MMC snap-in and WDSUTIL

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
	Native mode: MMC snap-in and WDSUTIL			
Multicasting	Not supported.	Supported for install images. (The Boot.wim file must be from the DVD for Windows Server 2003, Windows Vista with SP1, Windows Server 2008, Windows 7, or Windows Server 2008 R2.)	Supported for install images. (The Boot.wim file must be from the DVD for Windows Server 2008, Windows Vista with SP1, Windows 7, or Windows Server 2008 R2.) Supported for boot images for computers with EFI (managed by using the command line only). Provides the ability to automatically disconnect slow clients and divide transmissions into multiple streams based on client speeds (the boot image must be Windows 7 or Windows Server 2008 R2). Provides support for multicasting in environments that use IPv6 (the boot image must be from Windows Server 2008, Windows Vista with	Provides support for TFTP and multicasting over IPv6 and DHCPv6. Improved multicast deployment by eliminating the need for making a local copy of the install.wim file. You can apply the install.wim file while it is being downloaded without significant impact to the application process. Applying the install.wim file concurrently with the multicast download decreases overall deployment time, offers higher disk write throughput, and requires less space on the client system. Reduced default block size fits within an

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
			SP1, Windows 7, or Windows Server 2008 R2).	Ethernet MTU and works on customer hardware that does not support IP fragmentation.
Driver provisioning	Not supported.	Not supported.	Provides the ability to deploy driver packages to client computers based on their hardware as part of an installation (the install image must be Windows Vista with SP1, Windows 7, or Windows Server 2008). Provides the ability to add driver packages to boot images (Windows 7 or Windows Server 2008 R2 images only).	Extends filters for driver groups to support manufacturer Model Number and **Device Groups**. Driver import automatically detects and prevents duplicate driver packages from being added to the driver store.
Extensibility	Supports custom PXE providers.	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). However, you must develop a way to boot clients. Supports custom	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). Includes a PXE provider to boot clients. Supports custom	Supports custom PXE providers. Supports transmitting data and images by using multicasting on a stand-alone server (Transport Server). Includes a PXE provider to boot clients. Supports custom multicast

Feature	Windows Server 2003	Windows Server 2008	Windows Server 2008 R2	Windows Server 2012
		multicast providers.	multicast providers.	providers.
Extensible Firmware Interface (EFI)	Supports network booting of Itanium-based computers with EFI.	Supports network booting of Itanium-based and x64-based computers with EFI/UEFI.	Supports network booting of x64-based computers with UEFI. Includes Auto-add functionality. Includes DHCP referral to direct clients to a specific PXE server. Provides the ability to deploy boot images by using multicasting.	Support x86 clients (32-bit processors) with Unified Extensible Firmware Interface (UEFI) to network boot and complete an end-to-end deployment by using WDS.

ARM Architecture and support

WDS can now deploy images to ARM clients, which is a CPU architecture that is specially engineered for low-cost, low-power consumption devices such as tablets, cell phones, GPS units, portable game consoles, network routers, and media players.

What value does this change add?

The ability to deploy images to ARM clients in addition to x86 and x64 architectures means a broader range of devices will be capable of running Windows.

What works differently?

Provides support for Windows on ARM machines as deployment clients.

Provides support for all existing deployment features on Windows on ARM clients, most notably multicast.

Note
PXE boot is not currently supported by the network drivers on ARM clients.

WDS infrastructure for custom deployments

What value does this change add?

New features that offer the ability to control all aspects of the deployment process.

What works differently?

The improvements include a variety of configuration options that allow administrators to more tightly control the deployment payload (such as images and driver packages) that is sent to client computers. These improvements include the following:

- Install image filters, which are similar to the set of driver group filters.
- Support for boot and install image priority to influence the ordering of these images as they appear in Boot Manager and WDS client image selection menus.
- The Expected Deployment Results Wizard, which allows administrators to view deployment information such as the set of matching driver groups that would be sent to a prestaged device.
- Ability to control which clients are able to boot from the PXE server.
- Control over the boot parameters of PXE clients including boot program, prompt policy, and boot.wim instance.
- Ability to control the WIM and VHD images that are deployed to the client.
- More control over the drivers that are deployed to the client.
- Control over the unattend file(s) that are used to customize the setup experience for the client.
- Lower-level WDSUTIL commands that allow administrators to set custom metadata tags and values on deployment payload and prestaged devices that get matched to directly influence the deployment process.

WDS Management Console

What value does this change add?

Extensions to the management console provide users with more convenient features to configure deployment options for their client computers.

What works differently?

The following set of features were added:

- **Create Client Unattend** dialog to make it easy for administrators to author client unattend files.
- Approve Pending Device Wizard, which simplifies and extends the functionality of the previous approve option.
- Built-in support for prestaged devices that supersede the **Remote Install** tab of the **Active Directory Users & Computers** snap-in.

WDSclient.exe

What value does this change add?

WDSclient.exe is a new standalone client that can perform Dynamic Driver Provisioning (DDP) queries, direct VHD application, and metadata queries.

What works differently?

The WDSclient tool supports the above functionality for use in custom deployment scripts.

Standalone server mode

What value does this change add?

Standalone server mode removes the dependency on Active Directory.

What works differently?

Starting in Windows Server 2012, Windows Deployment Services can be installed in a Standalone server mode. This removes the dependency on Active Directory. You still require DHCP, DNS and sufficient permissions to install and configure Windows Deployment Services. In this scenario, a local store is used to retain information about pre-staged devices.

DDP enhancements

What value does this change add?

DDP enhancements provide more control over provisioning drivers to client computers.

What works differently?

DDP enhancements include:

- Support for hardware model.
- Support for image IDs and custom device groups.
- Duplicate driver package avoidance to prevent the same package from being added multiple times to the driver store.

Expected Deployment Results Wizard

What value does this change add?

The Expected Deployment Results Wizard enables administrators to model the process of deploying a computer and seeing what boot images, install images, driver groups the computer will be offered from the server.

What works differently?

Expected Deployment Results Wizard is designed to help administrators efficiently test configuration changes to their servers and see how they impact deployments to all or specific (prestaged) computers. The tool is especially valuable when used to test some of the advanced configuration options afforded by the WDS infrastructure for custom deployments.

TFTP enhancements

What value does this change add?

Trivial File Transfer Protocol (TFTP) enhancements result in improved performance.

What works differently?

TFTP (Trivial File Transfer Protocol) has been enhanced and delivers improved results in performance.

You use the Windows Deployment Services Trivial File Transfer Protocol (TFTP) server to download the files that are needed to do a network boot using the Pre-Boot Execution Environment (PXE). PXE technology is a standard created by Intel that establishes a common and consistent set of pre-boot services within the boot firmware. The end goal is to enable a client to do a network boot and receive a network boot program (NBP) from a network boot server.

TFTP enhancements include:

- **Scalable buffer management** Provides support for a shared client buffer; allows buffering an entire file instead of a fixed size buffer for each client. Scalable TFTP buffer feature allows maintaining a single buffer per file in the server. When the server is buffering a file in shared mode, different sessions can read from the same shared buffer.
- **Scalable port management** Ability to use a dynamic or a fixed range of UDP ports to service clients with shared UDP port allocation. Sharing the same server port among different TFTP sessions improves scalability because there are sufficient ports when more clients are actively using the server.
- **Variable-size transmission window** Allows the client and server to determine the largest workable window size, resulting in improved TFTP performance. Provides the ability to dynamically determine the optimal window size.
- **Maximum TFTP block size** Previously implemented as a registry setting, this is now exposed to users through WDSUTIL and the WDS MMC snap-in.

Troubleshooting enhancements

What value does this change add?

Expanded tracing for troubleshooting issues with minimal impact to system performance.

What works differently?

Tracing has moved from plain text files to ETW logging, which enables efficient logging that works well even when diagnosing potential race conditions.

Boot Image and Install Image Priorities

What value does this change add?

When you have multiple boot or install images available to client computers, clients will be presented with a boot and an install menu that displays the selection of images to choose from.

What works differently?

Windows Deployment Services now allows you to set priorities to control the order that both boot and install image listings are presented to clients. This ability is integrated directly into the Windows Deployment Services user interface.

Additional references

For information about getting started with Windows Deployment Services, see Windows Deployment Services Getting Started Guide for Windows Server 2012.

For information about the Windows Assessment and Deployment Kit (Windows ADK), see **Windows Deployment with the Windows ADK**.

Windows Deployment Services Getting Started Guide for Windows Server 2012

This guide contains step-by-step guidance for how to use the Windows Deployment Services role. This guide focuses on the functionality of the complete installation of Windows Deployment Services.

In this guide

Contents

- What is Windows Deployment Services?
- What is Windows Deployment Services?
- What's new in Windows Deployment Services?
- Who should use this guide?
- Benefits of Windows Deployment Services
- Installing Windows Deployment Services
- Configuring Windows Deployment Services
- Installing Windows Deployment Services integrated with Active Directory
- Steps for adding images
- Creating custom install images
- Creating discover images
- Performing an unattended installation
- Creating a multicast transmission
- Additional References

Quick start checklist

To get up and running quickly, perform the following steps. Then you can use the instructions in the rest of this guide to perform more advanced tasks such as creating your own install images, creating discover images, or configuring an unattended installation.

Task	Reference
Install Windows Deployment Services.	Installing Windows Deployment Services
Configure the server and add the default images (Install.wim and Boot.wim) that are	Configuring Windows Deployment Services

Task	Reference
included on the installation media (in the \Sources folder).	
Install an operating system.	Installing an install image

What is Windows Deployment Services?

Windows Deployment Services enables you to deploy Windows operating systems. You can use it to set up new computers by using a network-based installation. This means that you do not have to install each operating system directly from installation media, for example a DVD or USB drive.

What's new in Windows Deployment Services?

Windows Deployment Services has evolved as each version is released. You can read about the new features and improvements in Windows Deployment Services in What's New in Windows Deployment Services.

Who should use this guide?

Windows Deployment Services is intended for deployment specialists who are responsible for the deployment of Windows operating systems. This guide assumes that you have a working knowledge of common desktop deployment technologies, as well as networking components such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Active Directory Domain Services (AD DS). The target audiences are:

- IT planners or analysts evaluating Windows 7, Windows 8, Windows® 8.1, Windows Server 2008 R2, Windows Server 2012, and Windows Server® 2012 R2.
- Enterprise IT planners or designers.
- Deployment specialists interested in deploying images to computers.

Benefits of Windows Deployment Services

Windows Deployment Services provides the following benefits:

- Allows network-based installation of Windows operating systems, including Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2. It reduces the complexity and cost when compared to manual installations.
- Deploys Windows images to computers without operating systems.
- Supports mixed environments that include Windows Vista, Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows XP, and Windows Server 2003.

- Uses standard Windows Server setup technologies, including Windows PE, .wim files, and image-based setup.

Installing Windows Deployment Services

Prerequisites for installing Windows Deployment Services

The following are requirements for installing this role, depending on whether you choose the default installation (both Deployment Server and Transport Server), or only the Transport Server role service.

Deployment Server and Transport Server (integrated with Active Directory)	Transport Server
• **AD DS**. A Windows Deployment Services server must be either a member of an AD DS domain or a domain controller for an AD DS domain. The AD DS domain and forest versions are irrelevant; all domain and forest configurations support Windows Deployment Services. • **DHCP**. You must have a working DHCP server with an active scope on the network because Windows Deployment Services uses PXE, which relies on DHCP for IP addressing. • **DNS**. You must have a working DNS server on the network before you can run Windows Deployment Services. • **NTFS volume**. The server running Windows Deployment Services requires an NTFS file system volume for the image store. • **Credentials**. To install the role, you must be a member of the Local Administrators group on the server. To initialize the server, you must be a member of the Domain Users group.	The only prerequisite is that you must be a member of the Local Administrators group on the server to install Transport Server.

Deployment Server and Transport Server (Standalone mode, no dependency on Active Directory)	Transport Server
• **NTFS volume**. The server running Windows Deployment Services requires an NTFS file system volume for the image store. • **DNS**. You must have a working DNS server on the network before you can run Windows Deployment Services. • **DHCP**. You must have a working DHCP server with an active scope on the network because Windows Deployment Services uses PXE, which relies on DHCP for IP addressing. • **Credentials**. To install the role, you must be a member of the Local Administrators group on the server.	The only prerequisite is that you must be a member of the Local Administrators group on the server to install Transport Server.

Installation Methods

You can install Windows Deployment Services by using Server Manager, or using Windows PowerShell.

To install using Server Manager

1. In Server Manager, click **Manage**.
2. Click **Add roles and features**.
3. Select Role-based or feature-based installation and choose the server to deploy WDS.
4. On the **Select server roles** page select the **Windows Deployment Services** check box.
5. Click **Next** and follow the wizard to completion.

During installation, on the **Select role services** page, the wizard presents the option to select role services to be installed in Windows Deployment Services. You can choose to install the Deployment Server and Transport Server, or leave both roles selected.

- **Transport Server**. This option provides a subset of the functionality of Windows Deployment Services. It contains only the core networking parts. You can use Transport Server to create multicast namespaces that transmit data (including operating system images) from a standalone server. You can also use it if you want to have a PXE server that allows clients to PXE boot and download your own custom setup application. You should use this option if you want to use either of these scenarios , but you do not want to incorporate all of Windows Deployment Services.

- **Deployment Server**. This option provides the full functionality of Windows Deployment Services, which you can use to configure and remotely install Windows operating systems. Note that Deployment Server is dependent on the core parts of Transport Server.

  Notes

  When you install Windows Deployment Services, it will automatically provide the option (selected by default) to install the Remote Server Administration Tools.

  You can also add this feature by using the Add Roles and Features Wizard and checking **Remote Server Administration Tools** on the **Select feature** page.

Windows PowerShell equivalent commands

The following Windows PowerShell cmdlet or cmdlets perform the same function as the preceding procedure. Enter each cmdlet on a single line, even though they may appear word-wrapped across several lines here because of formatting constraints.

In Windows PowerShell, unlike in the Add Roles and Features Wizard, management tools and snap-ins for a role are not included by default. To include management tools as part of a role installation, add the `-IncludeManagementTools` parameter to the cmdlet. If you are installing roles and features on a server that is running a Server Core Installation and you add a role's management tools to an installation, you are prompted to change the installation option to a minimal-shell option that allows the management tools to run. Otherwise, management tools and snap-ins cannot be installed on servers that are running the Server Core Installation option of Windows Server.

To install using Windows PowerShell

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** page, type any part of the name **Windows PowerShell**. Right-click the shortcut for Windows PowerShell when it is displayed on the **Start** page in the **Apps** results, click **Advanced**, and then click **Run as Administrator**. To pin the Windows PowerShell shortcut to the **Start** page, right-click the shortcut, and then click **Pin to Start**.
2. Type the following, and then press **Enter**, where *computer_name* represents a remote computer on which you want to install Hyper-V. To install WDS directly from a console session, do not include `-ComputerName <computer_name>` in the command.

   ```
   Install-WindowsFeature –Name WDS -ComputerName
   <computer_name> -IncludeManagementTools
   ```

3. To view a list of available and installed roles and features on the local server, type **Get-WindowsFeature** and then press **Enter**. The results of the cmdlet contain the command names of roles and features that have been added to this computer.

   Note

   In Windows PowerShell 4.0, there is no need to import the Server Manager

cmdlet module into the Windows PowerShell session before running cmdlets that are part of the module. A module is automatically imported the first time you run a cmdlet that is part of the module. Also, Windows PowerShell cmdlets are not case-sensitive.

4. When the installation is finished, verify installation by running `Get-WindowsFeature`. If you installed WDS remotely, include the `ComputerName` parameter (`Get-WindowsFeature -ComputerName` <*computer_name*>) to view a list of roles and features that are installed on the server.

Note

This will install both the Deployment Server and the Transport Server roles. The previous command line tool, Servermanagercmd.exe, was deprecated in Windows Server 2012. It is recommended that you use the Windows PowerShell cmdlets.

Install Windows Deployment Services

To install the Windows Deployment Services role

1. Sign in to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type servermanager.exe, and then click **Server Manager**.
3. Click **Manage**.
4. Click **Add roles and features**, click **Next**.
5. On the **Before you begin** page of the **Add Roles and Features Wizard**, click **Next**.
6. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
7. On the **Select destination server** page, select the appropriate server, and then click **Next**. The local server is selected by default.
8. On the **Select server roles** page, scroll down and then select the **Windows Deployment Services** check box. Click **Next**.
9. Remote Server Administration Tools are required to manage this feature. Select **Include management tools (if applicable)**. Click **Add Features**. Click **Next**.
10. On the **Select features** page, click **Next**.
11. On the **Select role services** page, select the role services to install for Windows Deployment Services. If you wish to install both the Deployment Server and Transport Server, leave these role services selected. Click **Next**.
12. On the **Confirm installation selections** page, click **Install**.
13. Windows Deployment Services will now be added to the server. Installation progress will now be illustrated in the Add Roles and Features Wizard.

Note

For an Active Directory integrated installation, the Deployment Server requires that Active Directory Domain Services, DHCP, and DNS services are available on your network. Transport Server does not require any additional roles or services. Both of these services require an NTFS partition for the file store.

Note
Before you begin, you need to configure Windows Deployment Services by running either the Windows Deployment Services Configuration Wizard or WDSUtil.exe. You will also need to add at least one boot image and one install image to the image store.

Note
To install Windows operating systems from a Windows Deployment Services server, either the client computers must be PXE-enabled, or you must use the latest version of the Windows Pre-Installation Environment (Windows PE). See Windows Assessment and Deployment Kit (ADK) for Windows 8.1in Additional References.

Configuring Windows Deployment Services

After you install the Windows Deployment Service role, you must configure the server. Once you have used the instructions in this section to configure the server, add a boot image, and an install image, you will be ready to deploy images. Then, you can use the instructions in the rest of this guide to perform more advanced tasks like creating your own install images, creating discover images, or configuring an unattended installation.

Prerequisites for configuring Windows Deployment Services

1. If Windows Deployment Services is integrated with Active Directory, the server must be a member of an Active Directory Domain Services (AD DS) domain, or a domain controller for an AD DS domain. If Windows Deployment Services is installed in Standalone mode, it can be configured without having a dependency on Active Directory.
2. There is an active DHCP server on the on Windows Deployment Services uses Pre-Execution Environment (PXE), which relies on DHCP for IP addressing.
3. There is an active DNS server on the network.
4. The server has an NTFS file system partition on which to store images.

Steps for configuring Windows Deployment Services for Standalone Server

There are two configuration options for Windows Deployment Services. It can be configured to integrate with Active Directory or be configured as a standalone server.

To configure the server role in Standalone mode, use the following procedure. Then see the following section to add images to the server.

To configure Windows Deployment Services in Standalone mode

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Tools**, and then click **Windows Deployment Services** to launch the Windows Deployment Services MMC-snap (or console).
4. In the left pane of the Windows Deployment Services MMC snap-in, expand the list of servers.
5. Right-click the desired server, click **Configure Server**.
6. On the **Before you begin** page, click **Next**.
7. On the **Install options** page, choose **Standalone server**.
8. On the **Remote Installation Folder Locations** page, choose the default path or enter your own path to the remote installation folder. Click **Next**.
9. On the **PXE Server Initial Settings** page, choose desired option to define which client computers this server will respond to. Click **Next**. This will complete the configuration of Windows Deployment Services.
10. When the configuration is completed clear the **Add images to server now** check box, and then click **Finish**.
11. If you want to modify any of the settings of the server, right-click the server in the MMC-snap in, and click **Properties**.

Now that you have configured the server, you will need to add images. These images include a boot image (which is the bootable environment that you initially boot the computer into, and the install images (which are the actual images that you deploy). For instructions, see Steps for adding images.

Installing Windows Deployment Services integrated with Active Directory

This section describes installing Windows Deployment Services integrated with Active Directory. You may skip this if you have chosen to install Standalone mode, and proceed to Steps for adding images.

Prerequisites for installing Windows Deployment Services

The following are requirements for installing this role, depending on whether you choose the default installation (both Deployment Server and Transport Server), or only the Transport Server role service.

To install Windows Deployment Services integrated with Active Directory, the following prerequisites are required.

Deployment Server and Transport Server	Transport Server
• **AD DS**. A Windows Deployment Services	The only prerequisite is that you must be a

server must be either a member of an AD DS domain or a domain controller for an AD DS domain. The AD DS domain and forest versions are irrelevant; all domain and forest configurations support Windows Deployment Services. • **DHCP**. You must have a working DHCP server with an active scope on the network because Windows Deployment Services uses PXE, which relies on DHCP for IP addressing. • **DNS**. You must have a working DNS server on the network before you can run Windows Deployment Services. • **NTFS volume**. The server running Windows Deployment Services requires an NTFS file system volume for the image store. • **Credentials**. To install the role, you must be a member of the Local Administrators group on the server. To initialize the server, you must be a member of the Domain Users group.	member of the Local Administrators group on the server to install Transport Server.

Steps for configuring Windows Deployment Services integrated with Active Directory

To configure the Windows Deployment Services server role integrated with Active Directory, use the following procedure.

To configure Windows Deployment Services integrated with Active Directory

1. Log on to the server as a member of the local Administrators group.
2. Server Manager will start automatically. If it does not automatically start, click **Start**, type **servermanager.exe**, and then click **Server Manager**.
3. Click **Tools**, and then click **Windows Deployment Services** to launch the Windows Deployment Services MMC-snap (or console).
4. In the left pane of the Windows Deployment Services MMC snap-in, expand the list of servers.
5. Right-click the desired server, click **Configure Server**.
6. On the **Before you begin** page, click **Next**.
7. On the **Install options** page, choose **Integrated with Active Directory**.
8. On the **Remote Installation Folder Locations** page, choose the default path or enter

your own path to the remote installation folder. Click **Next**.

9. If your server is running as a DHCP Server, you will see the **Proxy DHCP Server** page. Select **Do not listen on DHCP and DHCPv6 ports** and **Configure DHCP options for Proxy DHCP**. Click **Next**.
10. On the **PXE Server Initial Settings** page, choose desired option to define which client computers this server will respond to. Click **Next**. This will complete the configuration of Windows Deployment Services.
11. When the configuration is completed, click **Finish**.
12. If you want to modify any of the settings of the server, right-click the server in the MMC-snap in, and click **Properties**.

Now that you have configured the server integrated with Active Directory, you will need to add your images. These images include a boot image (which is the bootable environment that you initially boot the computer into, and the install images (which are the actual images that you deploy). For instructions, see the next section, Steps for adding images.

Steps for adding images

You must add at least one boot image and one install image before you will be able to boot to the Windows Deployment Services server and install an image.

- **Boot images**. Boot images are Windows PE images that you boot a client computer into to perform an operating system installation. In most scenarios, you should use the Boot.wim file from the installation media (in the \Sources folder). The Boot.wim file contains Windows PE and the Windows Deployment Services client.
- **Install images**. Install images are the operating system images that you deploy to the client computer. You can also use the Install.wim file from the installation media (in the \Sources folder), or you can create your own install image by using the steps in Creating custom install images later in this guide.

The following instructions use Windows 7 installation media as an example for creating an install image that can be deployed with Windows Deployment Services. You can also use the same instructions for server installation media.

To add the default images, use the following procedures.

To add the default boot image included on the product installation media

1. In the left pane of the Windows Deployment Services MMC snap-in, right-click the **Boot Images** node, and then click **Add Boot Image**.
2. Browse to choose the default boot image (Boot.wim) on the Windows installation media, located in the **\Sources** folder.
3. Click **Open**, and then click **Next**.
4. On the **Image Metadata** page, enter the desired image name and description. Click **Next**.
5. On the **Summary** page, click **Next**.

Repeat this procedure to add any additional boot images.

Installing an install image

After you have at least one boot and one install image on the server, you can deploy an install image.

Prerequisites for installing an install image

- The client computer must be capable of performing a PXE boot.
- The client computer must have at least 512 MB of RAM, which is the minimum amount of RAM for using Windows PE.
- The client must meet the system requirements for the operating system of the install image.
- A local user account must be a created on the Windows Deployment Services server.

Steps for installing an install image

To perform a PXE boot on a computer to install an image, use the following procedure.

To install an operating system

1. Configure the BIOS of the computer to enable PXE booting, and set the boot order so that it is booting from the network first.
2. Restart the computer, and when prompted, press F12 to start the network boot.
3. Select the appropriate boot image from the boot menu. (This boot image selection menu will be available only if you have two or more boot images on the server.
4. At the **Install Windows** page, choose your **Locale and Keyboard or Input Method**. Click **Next**.
5. At the **Connect to Your WDS Server** authentication dialog, enter your local user account and password. Click **OK**.
6. At the **Install Windows** dialog box, select the operating system you wish to install. If you only have one operating system to install, choose that and click **Next**.
7. If prompted at the **Where do you want to install Windows** page, select the partition you wish to install the operating system on. Click **Next**.
8. When the installation is completed, the computer will restart and Setup will continue.

Boot Image and Install Image Priorities

When you have multiple boot or install images available to client computers, clients will be presented with a boot and an install menu that displays the selection of images to choose from.

Windows Deployment Services allows you to set priorities to control the order that both boot and install image listings are presented to clients. This ability is integrated directly into Windows Deployment Services.

Steps for configuring the boot menu

To configure menu order for boot images

1. Open the Windows Deployment Services MMC snap-in.
2. Click the **Boot Images** node. You will see your boot images appear in the right hand side of your Windows Deployment Services MMC snap-in.
3. Right-click your desired boot image from the right-hand side of your Windows Deployment Services MMC snap-in. Click **Properties**.
4. In the **Image Properties** dialog, on the **General** tab, enter in your desired priority into the Priority text box. The items that appear first on your install image menu are the ones with the lowest value.
5. Click **OK**.

To configure menu order for install images

1. Open the Windows Deployment Services MMC snap-in.
2. Double-click the **Install Images** node. You will see your image group (or image groups) appear as a sub menu item. They will also appear in the right hand side of your Windows Deployment Services MMC snap-in.
3. Click your desired **Image Group**.
4. Right-click your desired image within your image group from the right-hand side of your Windows Deployment Services MMC snap-in. Click **Properties**.
5. On the **Image Properties** dialog, in the **General** tab, enter in your desired priority into the **Priority** text box. The items that appear first on your install image menu are the ones with the lowest value.
6. Click **OK**.

When you have completed this procedure and you perform a PXE boot on a client computer, a boot or install menu with the menu order you set using priorities will appear. (if those images apply to that computer).

Priorities are pre-populated with a default value that lets you place images higher or lower on the list. The items that appear first on the list are the ones with the lowest value.

Creating custom install images

Optionally, you can create custom install images for Windows 8.1, Windows 8, and Windows 7. To do this, use the instructions in this section to create a capture boot image, prepare a reference computer using Sysprep, and then capture the operating system using the Image Capture Wizard.

Prerequisites for creating custom install images

- You must ensure that there is enough disk space to create and store the new images.

- You must be a member of the Local Administrators group on the Windows Deployment Services server.

Known issues when creating custom install images

When you boot into the capture image, the Image Capture Wizard will start. Note the following points about this wizard:

- You will only see drives that contain operating systems prepared with Sysprep.
- If you do not run Sysprep on the computer before you boot into the image, you will not see any drives to capture.

You must enter a local location to save the new install image; otherwise, you will not be able to capture the image. This location can be a mapped network drive but we recommend that you use a local location to avoid image corruption in the event of a network malfunction.

Steps for creating a capture image

To create an install image, you must first create a capture image. Capture images are boot images that you boot a client computer into to capture the operating system into a .wim file. You can also create media (a CD, DVD, USB drive, or other type of media) that contains a capture image, and then boot a computer from the media. These images provide an alternative to the command-line utility, ImageX.exe. Except in advanced scenarios, you can create a capture image by using the Boot.wim file from the Windows installation media.

To create a capture image

1. In the Windows Deployment Services MMC snap-in, expand the **Boot Images** node.
2. Right-click the image to use it as a capture image. In most cases, you can just use the Boot.wim file from the installation media.
3. Click **Create Capture Image**.
4. Type in your Image Name, Image Description, and the location and file name where you want to save a local copy of the file. You must specify a location in case there is a problem with the network when you deploy the capture image. Click Next.
5. Allow the Create Capture Image Wizard to complete.
6. Tick **Add Image to the Windows Deployment Server now**. Click **Next**
7. Enter the location of the Windows Image file that contains the images. Click **Next**.
8. Enter your **Image Name** and **Image Description**. Click **Next**.
9. On the **Summary** page, click **Next**.
10. Click **Finish**.

After you have created the capture image, follow the instructions in the next section to boot a computer into the capture image and capture the operating system.

Steps for creating an install image

Now that you have a capture image, you need to prepare a reference computer and then create the install image. The reference computer can be a computer with a standard Windows installation or a Windows installation that has been configured for your environment. First, you boot a computer (which has been prepared with Sysprep) into the capture image. Then a wizard creates an install image of the reference computer and saves it as a .wim file. After that, you can deploy the .wim file to a computer.

To create a custom install image

1. Create a reference computer (install the operating system, applications, and make any other changes that you want).
2. Ensure that you have the correct version of Sysprep.exe on the computer.
3. At a command prompt on the reference computer, change folders to \Windows\System32\Sysprep or the folder that contains Sysprep.exe and Setupcl.exe.
4. Type one of the following:
   - On computers running Windows 8.1, Windows 8, Windows 7, Windows Server 2012 R2, Windows Server 2012, and Windows Server 2008 R2 run the command `sysprep /oobe /generalize /reboot`. If you prefer, you can also use the Sysprep graphical user interface by double-clicking Sysprep.exe.
5. When the computer restarts, perform a network boot on the computer by pressing F12.
6. In the boot menu, select the capture boot image that you created in the preceding procedure, and then press ENTER.
7. You will be presented with the Windows Deployment Services Image Capture Wizard. Click **Next**.
8. On the **Directory to Capture** page, select **Volume to capture**, enter your **Image name** and **Image description**. Click **Next**.

   Important
   Referring to **Volume to capture**, you will see only drives that contain operating systems prepared with Sysprep. If you do not run the command in step 4, you will not see any drives to capture.
9. Click **Browse** next to Name and location and browse to a local folder where you want to store the captured install image. Type a name for the image, using the .wim file name extension, and then click **Save**. Note that this location can be a mapped network drive.

   Important
   You must enter a local location to save the new image; otherwise, you will not be able to capture the image. This requirement is enforced to avoid image corruption if there is a network malfunction.
10. Select the **Upload Image to a Windows Deployment Services** check box.
11. Click **Connect**. If prompted for credentials, provide a user name and password for an account with permissions to connect to the Windows Deployment Services server.

12. Select your **Image Group Name**. Click **Next**. The wizard will now complete and create a custom installation image and store it in the Windows Deployment Services store.
13. Click **Finish**.

When this process is complete, you can PXE boot a client computer to install this image. The image will be listed in the installation option.

Creating discover images

If you have a computer that is not PXE enabled, you can create a discover image and use it to install an operating system on that computer. Otherwise, you can skip this section. When you create a discover image and save it to media (a CD, DVD, USB flash drive, or other type of media), you can then boot a computer to the media.

The discover image on the media locates a Windows Deployment Services server, and the server deploys the install image to the computer. You can configure discover images to target a specific Windows Deployment Services server. This means that if you have multiple servers in your environment, you can create a discover image for each one, and then name them based on the name of the server.

Prerequisites for creating discover images

- Have CD or DVD or flash drive to store the image.
- Have a disk-burning utility if you are burning the image to CD or DVD.

Steps for creating discover images

After you create the discover image, you will need to create media that contains the image. You must create discover images by using the Boot.wim file from your Windows installation media.

To create a discover image

1. In the Windows Deployment Services MMC snap-in, expand the **Boot images** node.
2. Right-click the image you want to use as a discover image.
3. Click **Create Discover Image**.
4. In the **Metadata and Locations** page, enter your **Image Name** and **Image Description**. In the location and file name, click **Browse** and select the location for storing your discover boot image. Enter the name of your discover image. Optionally, enter the name of your Windows Deployment Services server that you want to respond when you boot a computer into this image. Click **Next**.
5. Click **Finish**.

Use the following procedure to create media that contains the image.

This following example details how to create media that contains a discovery image. However, it is an example only; you can substitute any values you require. For example, different architecture and folder names and locations.

To create media that contains the discover image

1. Download and install the Windows Assessment and Deployment Kit for Windows 8.
2. Log on to the server as a member of the local Administrators group.
3. Click **Start**, click **Deployment and Imaging Tools**. Otherwise, type **Deployment and Imaging Tools**, and then click **Deployment and Imaging Tools** from the search results pane. You will be presented with a Command Prompt window.
4. To create a Windows PE build environment, type:

   ```
   CopyPE  x86   c:\MYWinPEFolder.
   ```

5. To copy the discover image that you created in the preceding procedure, type:

   ```
   copy /y c:\discover.wim C:\MYWinPEFolder\media\sources
   ```

6. To change back to the Windows PE tools folder, type:

   ```
   cd C:\Program Files (x86)\Windows Kits\8.0\Assessment and
   Deployment Kit\Deployment Tools\x86\Oscdimg
   ```

7. To create the bootable .iso image, type:

   ```
   oscdimg -n"C:\Program Files (x86)\Windows Kits\8.0\Assessment
   and Deployment Kit\Deployment Tools\x86\Oscdimg\eftsboot.com"
   C:\MYWinPEFolder\media\sources c:\discover.iso
   ```

8. Use a utility that can create a CD or DVD to transfer the .iso image to the appropriate media.

   Note
   Using the Windows copy command to copy the image to media will not create a bootable disk.

9. Boot the client computer to the media. You may need to change the boot order in the client BIOS to do this. When the client boots to the media, your experience will be exactly the same as if you had booted to the network.

Performing an unattended installation

Optionally, you can automate the entire installation. To do this, you use two different unattend files: one for the Windows Deployment Services UI screens, and one for the latter phases of Setup.

- **Windows Deployment Services client unattend file**. This file uses the Unattend.xml format and is stored on the Windows Deployment Services server in the C:\RemoteInstall\WDSClientUnattend folder. It is used to automate the Windows Deployment Services client user interface screens (such as entering credentials, choosing an install image, and configuring the disk).
- **Image unattend file**. This file can use either the Unattend.xml or Sysprep.inf format. It is stored in a subfolder (either OEM structure or \Unattend) in the per-image folder. It is used

to automate the remaining phases of setup (for example, offline servicing, Sysprep specialize, and mini-setup).

Prerequisites for performing an unattended installation

- Ensure that you are a member of the Local Administrators group on the Windows Deployment Services server.
- Ensure that you have sufficient credentials to add a computer to a domain.

Steps for performing an unattended installation

To automate the installation, create the appropriate unattend file depending on whether you are configuring the Windows Deployment Services screens or Windows Setup. We recommend that you use Windows System Image Manager (included as part of the Windows Assessment and Deployment Kit) to author the unattend files. The Windows System Image Manager (Windows SIM) creates and manages unattended Windows Setup answer files in a graphical user interface (GUI).

You will then need to copy the unattend file to the appropriate location, and assign it for use. You can assign it at the server level or the client level. The server level assignment can further be broken down by architecture, enabling you to have different settings for x86-based and x64-based clients. Assignment at the client level overrides the server-level settings.

To associate a client unattend file by architecture

1. Create an Unattend.xml file with settings applicable to Windows Deployment Services.
2. Copy the client unattend file to a folder in the **RemoteInstall** folder. For example: **RemoteInstall\WDSClientUnattend**.
3. Open the Windows Deployment Services MMC snap-in, right-click the server that contains the image that you want to associate the unattend file with, and then click **Properties**.
4. On the **Client** tab, select **Enable unattended installation**, browse to the appropriate unattend file, and then click **Open**.
5. Click **OK** to close the **Properties** page.

To associate an image unattend file

1. Create an Unattend.xml file or Sysprep.inf file.
2. The next steps depend on which file you create:
   - **Save Sysprep.inf files.** Save these files to the OEM structure of the image (for example, D:\RemoteInstall\Images\Windows XP\winxpsp2\$OEM$\$1\sysprep\sysprep.inf). Now when you deploy the image, Setup will automatically locate and use the Sysprep.inf file.
   - **Save Unattend.xml files.** Save Unattend.xml files to any location, and then associate the file with an image. To do this, open the WDS MMC snap-in, right-click the image that you want to associate with the unattend file, and then click

Properties. On the **General** tab, click **Allow image to install in unattend mode**, click **Select File**, browse to select the unattend file, and then click **OK** twice. The Unattend.xml file will be saved to the following location:

\RemoteInstall\Images\<imagegroup>\<imagename>\Unattend\ImageUnattend.xml.

Creating a multicast transmission

Multicast transmissions enable you to deploy an image to a large number of client computers without overburdening the network. When you create a multicast transmission, the image is sent over the network only once, which can drastically reduce the amount of network bandwidth that is used. Note that data is transferred only if clients request data.

If no clients are connected (that is, if the transmission is idle), data will not be sent over the network. For information about using multicast functionality with Transport Server, see Using Transport Server.

Prerequisites for creating a multicast transmission

- Routers that support multicasting. In particular, Internet Group Membership Protocol (IGMP) snooping should be enabled on all devices. This will cause your network hardware to forward multicast packets only to those devices that are requesting data. If IGMP snooping is turned off, multicast packets are treated as broadcast packets, and will be sent to every device in the subnet.
- At least one install image that you want to transmit on the server.
- The Boot.wim file provided with Windows Server 2012 R2or Windows Server 2012.

Steps for creating a multicast transmission

To enable multicasting for an image, you must create a multicast transmission for it. You have two options for creating a transmission:

- Right-click the **Multicast Transmission** node, and then click **Create Multicast Transmission**.
- Right-click an image, and then click **Create Multicast Transmission**.

There are two types of multicast transmissions:

- **Auto-Cast**. This option indicates that as soon as an applicable client requests an install image, a multicast transmission of the selected image begins. Then, as other clients request the same image, they too are joined to the transmission that is already started.
- **Scheduled-Cast**. This option sets the start criteria for the transmission based on the number of clients that are requesting an image and/or a specific day and time. If you do not select either of these check boxes, the transmission will not start until you manually start it.

Steps for configuring transmissions

After a transmission has been created, you can perform the following actions:

- **Start the transmission**. If the transmission is the Scheduled-Cast type, there is at least one client, and the transmission has not started yet, you can right-click the transmission and then click **Start**.
- **Delete the transmission**. If you right-click the transmission and click **Delete**, the multicast transmission stops and each client installation will fall back to using unicast transmission. That is, the client installations will not be deleted or stopped, but they will not use the multicast transmission to complete the installation.
- **Deactivate the transmission**. If you right-click and then click **Deactivate**, each client that is currently installing will continue, but no new clients will be joined to the transmission. After each current client installation is completed, the transmission will be deleted. If there are no clients when you click this option, the transmission will be deleted instantly.
- **View the transmission's properties**. To view the properties, right-click the transmission, and then click **Properties**. Note that you cannot edit the properties of a transmission after it is created. To make a change after you have created a transmission, you need to delete it and then recreate it.
- **Refresh the transmissions and data**. To do this, right-click a transmission, and then click **Refresh**. You can also refresh the data by pressing F5.

Steps for configuring clients in a transmission

After you have created the transmission, client computers can join it by selecting the image in the Windows Deployment Services client installation screens. Client computers can also join a transmission by using Wdsmcast.exe, a command-line tool included in the Windows Assessment and Deployment Kit (ADK). You will find command line tool from the C:\Program Files (x86)\Windows Kits\8.0\Assessment and Deployment Kit\Deployment Tools\x86\wdsmcast folder.

When there are clients in a transmission, you can do any of the following:

- **View clients and see progress**. To view any connected clients, expand the **Multicast Transmissions** node, and click the image.
- **Stop a client installation**. To stop the installation completely, right-click a client, and then click **Disconnect**. You should use this option with caution because the installation will fail and the computer could be left in an unusable state.
- **Disconnect a client from a multicast transmission**. To discontinue the transmission for a particular client but continue to transfer the image through unicasting, right-click the client, and then click **Bypass multicast**.

Additional References

Understanding the Windows Assessment and Deployment Kit (ADK) for Windows 8.1

Download the Windows Assessment and Deployment Kit for Windows 8.1.

Features available in the Windows Assessment and Deployment Kit (ADK) include:

Application Compatibility Toolkit (ACT)

The Application Compatibility Toolkit (ACT) helps IT Professionals understand potential application compatibility issues by identifying which applications are or are not compatible with the new versions of the Windows operating system. ACT helps to lower costs for application compatibility evaluation by providing an accurate inventory of the applications in your organization. ACT helps you to deploy Windows more quickly by helping to prioritize, test, and detect compatibility issues with your apps. By using ACT, you can become involved in the ACT Community and share your risk assessment with other ACT users. You can also test your web applications and web sites for compatibility with new releases of Internet Explorer. For more information, see Application Compatibility Toolkit.

Deployment Tools

Deployment tools enable you to customize, manage, and deploy Windows images. Deployment tools can be used to automate Windows deployments, removing the need for user interaction during Windows setup. Tools included with this feature are Deployment Imaging Servicing and Management (DISM) command line tool, DISM PowerShell cmdlets, DISM API, Windows System Image Manager (Windows SIM), and OSCDIMG. For more information, see Deployment and Imaging Tools Technical Reference.

User State Migration Tool (USMT)

USMT is a scriptable command line tool that IT Professionals can use to migrate user data from a previous Windows installation to a new Windows installation. By using USMT, you can create a customized migration framework that copies the user data you select and excludes any data that does not need to be migrated. Tools included with the feature are ScanState, Loadstate, and USMTUtils command line tools. For more information, see User State Migration Tool Technical Reference.

Volume Activation Management Tool (VAMT)

The Volume Activation Management Tool (VAMT) enables IT professionals to automate and centrally manage the activation of Windows, Windows Server, Windows ThinPC, Windows POSReady 7, select add-on product keys, and Office for computers in their organization. VAMT can manage volume activation using retail keys (or single activation keys), multiple activation keys (MAKs), or Windows Key Management Service (KMS) keys. For more information, see Volume Activation Management Tool Technical Reference.

Windows Performance Toolkit (WPT)

Windows Performance Toolkit includes tools to record system events and analyze performance data in a graphical user interface. Tools available in this toolkit include Windows Performance Recorder, Windows Performance Analyzer, and Xperf.

Windows Assessment Toolkit

Tools to discover and run assessments on a single computer. Assessments are tasks that simulate user activity and examine the state of the computer. Assessments produce metrics for various aspects of the system, and provide recommendations for making improvements.

Windows Assessment Services

Tools to remotely manage settings, computers, images, and assessments in a lab environment where Windows Assessment Services is installed. This application can run on any computer with access to the server that is running Windows Assessment Services.

Windows Preinstallation Environment (Windows PE)

Minimal operating system designed to prepare a computer for installation and servicing of Windows.

Windows System Resource Manager Overview

You can use Windows System Resource Manager to allocate processor and memory resources to applications, users, Remote Desktop Services sessions, and Internet Information Services (IIS) application pools.

With Windows System Resource Manager for the Windows Server® 2012 operating system, you can manage server processor and memory usage with standard or custom resource policies. Managing your resources can help ensure that all the services provided by a single server are available on an equal basis or that your resources will always be available to high-priority applications, services, or users.

Windows System Resource Manager only manages processor resources when the combined processor load is greater than 70 percent. This means that it does not actively limit the resources that can be used by each consumer when processor load is low. When there is contention for processor resources, resource allocation policies help ensure minimum resource availability based on the management profile that you define.

Role/Feature description

You can use Windows System Resource Manager to:

- Manage system resources (processor and memory) with preconfigured policies, or create custom policies that allocate resources per process, per user, per Remote Desktop Services session, or per Internet Information Services (IIS) application pool.
- Use calendar rules to apply different policies at different times without manual intervention or reconfiguration.
- Automatically select resource policies that are based on server properties and events (such as cluster events or conditions) or changes to installed physical memory or number of processors.
- Collect resource usage data locally or in a custom SQL database. Resource usage data from multiple servers can be consolidated on a single computer running Windows System Resource Manager.
- Create a computer group to help organize Remote Desktop Session Host servers that you want to manage. Policies can easily be exported or modified for an entire computer group.

Practical applications

Because Windows Server 2008 R2 is designed to give as many resources as possible to non-operating system tasks, a server running a single role usually does not require resource management. However, when multiple applications and services are installed on a single server, they are not aware of competing processes. An unmanaged application or service will typically use all available resources to complete a task. Thus, it is important to use a tool such as Windows System Resource Manager to manage system resources on multipurpose servers. Using Windows System Resource Manager provides two key benefits:

- More services can run on a single server because service availability can be improved through dynamically managed resources.
- High-priority users or system administrators can access the system even during times of maximum resource load.

Methods of resource management

Windows System Resource Manager includes five built-in resource management policies that you can use to quickly implement management. In addition, you can create custom resource management policies to meet your specific needs.

Built-in resource management policies

You can enable built-in resource management policies by selecting the type of policy to use. No further configuration is required.

Policy	Description
Equal per process	When the **Equal_Per_Process** resource allocation policy is managing the system, each running process is given equal treatment. For example, if a server that is running ten processes reaches 70 percent processor utilization, Windows System Resource Manager will limit each process to using 10 percent of the processor resources while they are in contention. Note that resources not used by low utilization processes will be allocated to other processes.
Equal per user	When the **Equal_Per_User** resource allocation policy is managing the system, processes are grouped according to the user account that is running them, and each of these process groups is given equal treatment. For example, if four users are running processes on the server,

Policy	Description
	each user will be allocated 25 percent of the system resources to complete those processes. A user running a single application is allocated the same resources as a user running several applications. This policy is especially useful for application servers.
Equal per session	When the **Equal_Per_Session** resource allocation policy is managing the system, resources are allocated on an equal basis for each session connected to the system. This policy is for use with RD Session Host servers.
Equal per IIS application pool	When the **Equal_Per_IISAppPool** resource allocation policy is managing the system, each running IIS application pool is given equal treatment, and applications that are not in an IIS application pool can only use resources that are not being consumed by IIS application pools.
Weighted Remote Sessions	When the **Weighted_Remote_Sessions** resource allocation policy is managing the system, the processes are grouped according to the priority assigned with the user account. For example, if three users are remotely connected, the user assigned Premium priority will receive highest priority access to the CPU, the user assigned Standard priority will receive second priority to the CPU, and the user assigned Basic priority will receive lowest priority to the CPU. This policy is for use with RD Session Host servers. **Note** When **Weighted_Remote_Sessions** is set as the managing policy, system management is delegated to the Windows Server 2012 scheduler, and Windows System Resource Manager only profiles the system. Setting or removing **Weighted_Remote_Sessions** as the

Policy	Description
	managing policy requires a restart of the computer imposed by the kernel.

Custom resource management

You can use custom resource management methods to identify resource users and allocate resources to them based on your own criteria.

Feature	Description
Process matching criteria	Enable you to select services or applications to be managed by resource allocation policy rules. You can choose by file name or command, or you can specify users or groups. For example, you could create a process matching criterion that applies management to the application **iexplore.exe** when it is run by the user Administrator.
Resource allocation policies	Allocate processor and memory resources to processes that are specified by the process matching criteria that you create.
Exclusion lists	Exclude applications, services, users, or groups from management by Windows System Resource Manager. **Note** You can also use command-line path matching in a resource allocation policy to exclude an application from management by only that policy.
Scheduling	Use a calendar interface to control one-time events or recurring changes to resource allocation. Different resource allocation policies can be active at different times of day, on different days of the week, or according to other scheduling paradigms.
Conditional policy application	Automatically switch resource allocation policies in response to certain system events (such as installing new memory or additional processors, starting or stopping a node, or

Feature	Description
	changing the availability of a resource group in a cluster).

Removed or deprecated functionality

Windows System Resource Manager (WSRM) is deprecated beginning with Windows Server® 2012. You should begin planning now to use alternate methods for any applications, code, or scenarios that depend on this feature. Windows System Resource Manager is an administrative tool that controls how CPU and memory resources are allocated. For more information, see. Deprecated Features for Windows 7 and Windows Server 2008 R2

Windows Server Backup Feature Overview

This topic provides information about the Windows Server Backup feature.

Did you mean…

- Windows Azure Backup
- Data Protection Manager

Feature description

Windows Server Backup is a feature that provides a set of wizards and other tools for you to perform basic backup and recovery tasks for the server it is installed on.

Practical applications

You can use Windows Server Backup to back up a full server (all volumes), selected volumes, the system state, or specific files or folders—and to create a backup that you can use for bare metal recovery. You can recover volumes, folders, files, certain applications, and the system state. And, in case of disasters like hard disk failures, you can perform a bare metal recovery. (To do this, you will need a backup of the full server or just the volumes that contain operating system files, and the Windows Recovery Environment—this will restore your complete system onto your old system or a new hard disk.)You can use Windows Server Backup to create and manage backups for the local computer or a remote computer. And, you can schedule backups to run automatically. Windows Server Backup is intended for use by everyone who needs a basic backup solution—from small business to large enterprises—but is even suited for smaller organizations or individuals who are not IT professionals.

New and changed functionality

The following functionality has been changed or added to Windows Server Backup in Windows Server 2012:

Feature/functionality	Previous operating system	New operating system
Backup and restore of individual virtual machines from a Hyper-V host server.	Virtual machines were backed up as part of a volume backup. They could not be backed up or restored separately.	When you select items to backup or recover you can now select individual virtual machines to include in the backup and to restore from a recovery point.
Improvements to managing backup versions and backup retention	Space and version management was not available. Windows Server Backup used space as needed.	If you are backing up to a hard disk or volume you can specify a deletion policy to determine whether backups should be deleted only when necessary to make space for additional backups or if they should be deleted after a certain number of backups have passed.
Ability to backup volumes greater than	Backups of volumes greater	Backup can now support larger volumes and hard drive sectors.

Feature/functionality	Previous operating system	New operating system
2 terabytes (TB) and with 4 K sector sizes.	than 2 TB was not supported and 512 byte sectors were required.	
Support for backup of Cluster Shared Volumes (CSVs)	Backups of CSVs was not supported.	Limited support for backing up CSVs as follows: 1. Virtual machines hosted on CSV's cannot be added as part of backup configuration 2. Windows Server Backup has to be configured on all nodes to ensure that backup and recovery will be available in the event of a failure on one of the nodes in the cluster. 3. Volumes recovery not supported 4. Security access control lists are not applicable on CSV file service root. Therefore, file recovery to the root of CSV volume is not supported.
Reporting system state components.	Win32 service files were reported by the System Writer as part of the System Component and were included in the system state backup	Win32 service files are not reported by the System Writer as part of the System Component for Windows Server 2012. Win32 service files are reported by the System Writer as part of the System Component for Windows 8. To include or exclude the Win32 service files in the system state of your computers, set the value of the registry key **HKLM\System\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState**. • To include the Win32 service files set the value to 0. • To exclude the Win32 service files set the value to 1 On the server operating system, deleting this registry key is equivalent to setting it to 1 and causes Win32 Service files to be excluded from System Writer's system component specifications. On the client operating system, deleting this registry key is equivalent to setting it to 0 and causes Win32 Service files to be included in System Writer's system component specifications.

Feature/functionality	Previous operating system	New operating system
	for both client and server operating systems.	

Virtual Machine Support

Windows Server Backup can be used within a virtual machine environment in the same manner that it is used on a physical server to provide backup and recovery services. It can also be used from within a Windows Azure virtual machine for Bare Metal Backup & Recovery, System State Backup & Recovery, File Folder Backup and recovery.

More information

For more information about backing up or recovering computers using Storage Spaces, see Windows Server Backup and Storage Pools

For more information about backing up or recovering computers running Active Directory Domain Services, see Administering Active Directory Backup and Recovery.

For more information about backing up computers running Hyper-V, see Planning for Backup.

Windows Server Backup and Storage Pools

This document is intended to be used as a starting point for understanding the limitations and implications of using the Windows Server Backup feature on a server running Windows Server 2012 that is using Storage Spaces.

Terminology:

- Online mode: When computer is booted into Windows.
- Offline mode: when computer is booted into Windows RE
- Bare-Metal Recovery (BMR): the process of reformatting a computer from scratch after a catastrophic failure which involves reinstalling the operating system and software applications and then, if possible, restoring data and settings. Windows RE System Recovery includes the ability to perform BMR.

- Automated System Recovery (ASR): the portion of Windows RE System Recovery that captures the configuration of disks and volumes on a system and restores them in bare metal recovery scenarios.
- LUN (Logical Unit Number): are used to identify devices, such as hard drives & partitions, connected to a computer. Each device is assigned a LUN, from 0 to 7, which serves as the device's unique address.
- Storage pool: is a collection of physical disks. Pools permit storage aggregation, flexible capacity expansion, and delegated administration
- Storage Spaces: the name of a storage subsystem that lets you group disks (drives) into one or more storage pools. You can then create virtual disks storage space with associated attributes such as a desired level of resiliency, thin or fixed provisioning, automatic or controlled allocation on heterogeneous storage media, and precise administrative control from available capacity (free space) in the pool.
- Virtual disks: A collection of one or more physical disks from a previously created storage pool.

This document includes the following sections:

- What is a storage pool?
- Backup and recovery process and storage pools
- Windows Server Backup Considerations
- Bare-metal recovery considerations
- Example 1: Same Server Recovery Process with disk recreation
- Example 2: Same server recovery without disk recreation
- Example 3: Alternate server recovery
- Critical volumes and storage spaces
- Special Considerations when using RAID and Storage Spaces
- Sample Scripts

What is a storage pool?

A storage pool enables a system administrator to group physical disks together to form a pool of resources from which they can designate spaces to provision disks, resulting in a form of storage virtualization.

A storage pool is created using physical disks. The physical disks should be free and not host any volume/partition. It is important to note that the physical disks assigned to a storage pool are hidden from Windows. This means that Windows Server Backup (Windows Server Backup) and Automated System Recovery (ASR) are not able to enumerate the physical disks. After the pool is created a storage administrator can use the pool as a basis on which to create a virtual disk. These virtual disks can be created with different RAID settings – Mirror, Spanned, and Parity. These virtual disks will be visible in Windows (DiskMgmt/Diskpart) as normal disks and thus are visible to Windows Server Backup and ASR.

Backup and recovery process and storage pools

Windows Server Backup does not differentiate a storage pool from other storage. Therefore the presence of a storage pool (both in the online mode and offline Windows RE mode) and its storage space are handled by Windows Server Backup in the same way as normal disks and volumes. Windows Server Backup can be used for any backup or recovery operations either to or from storage pool disks.

Windows Server Backup Considerations

The following list details what you should be aware of when using Windows Server Backup with Storage Spaces:

- Data from Storage Spaces can be backed up and recovered like normal disks and volumes.
- The system volume cannot be hosted on a storage space.
- Windows Server Backup and ASR do not have the ability to recreate storage pools and storage space in Windows RE.
- When using Storage Spaces it's recommended that administrators perform only a "Critical Volume BMR" if the storage pool encounters a complete failure for any reason. This recovery options restores only the critical volumes and does not attempt to recover data volumes. Once the storage pool has been restored to operation you can then perform data recovery from within the operating system.
- In case of complete server failure, you should recover the critical disks using BMR, boot into the operating system, recreate the storage space and then perform data recovery. It's not recommended to recover data on Storage Spaces using BMR.
- Windows Server Backup requires that all disks hosting critical volumes should be available in Windows RE for a BMR. If any critical volume is on a volume hosted on a storage space, then that storage space should be made online and mounted before starting BMR (assuming the physical disk is still intact). This can be done by using a DISKPART command. In a complete disk or server failure scenario, the physical disks host no storage spaces. In such a scenario, one of the physical disks should be made online and made available for BMR to proceed. All disks that host data or meant to be used for storage spaces should be excluded while performing BMR to avoid data loss scenarios.

  Note
  Currently it is recommended that a storage space virtual disk be used to host only one volume. This recommendation makes the BMR recovery scenarios easier.

Bare-metal recovery considerations

The following list details what you should be aware of when using bare-metal recovery (BMR) with Storage Spaces:

- BMR is started from Windows RE. By default, the Storage Spaces LUN would be in an attached and mounted state in Windows RE. Before you proceed with recovery, you should

either make the Storage Spaces disks offline or specifically exclude them from the recovery process.

- You cannot recreate the storage pool configuration in Windows RE Environment.
- Storage pool and storage space can be either attached or detached using WMI scripts. For more information, see Sample Scripts .
- The underlying locally attached disks assigned to storage pool will not be visible from Diskpart.exe.

  Warning
  As part of BMR backup, Windows Server Backup does not record storage pool or Storage Spaces LUN configuration information.
- Windows Server Backup and ASR are not aware of the underlying physical disks that host storage pools.
- Windows Server Backup and ASR are not aware of the storage pool configuration, the resiliency settings and other details specific to storage pool.
- Windows Server Backup and ASR will treat storage space virtual disks the same as any other locally attached disks.
- ASR will not store the storage pool configuration info during BMR backup and it cannot recreate storage pools and Storage Spaces during recovery.
- It is recommended that you store the storage pool and Storage Space configuration details as part of your disaster recovery plan. Later this information can be used to recreate the storage pool.
- ASR can recreate and recover volumes on disks. Since ASR treats the Storage Spaces as normal disks, it can recreate volumes on Storage Spaces.
- The storage space virtual disks have a different Bus Type than Physical disks (SCSI/IDE/SATA). The Bus Type is used to during partition layout recreation in BMR. BMR will attempt to recreate volumes on disks with same bus type.

Example 1: Same Server Recovery Process with disk recreation

This example describes performing a bare-metal recovery with disk recreation when some or all the storage pool volumes are part of backup. The sample server configuration has the following details:

- Storage pool and storage space virtual disks are configured on the server.
- Data volumes have been created on the storage space virtual disks.
- Windows Server Backup has been run and performed either a Full Server Backup or a backup for bare-metal recovery.
- The backup includes some or all of the volumes on the storage space virtual disks.

Recovery is being performed for the following reasons:

- Hardware failure of storage disks
- Disk replacement

In this situation the administrator needs to recover only the system volumes and can perform the rest of the data recover using Windows Server Backup. To do this the Storage Spaces must be excluded from recovery. You can choose to use the diskpart command-line tool or a WMI script to detach the storage space virtual disk before starting the repair process or if the disks are online you can use the wizard to exclude the storage space from the recovery.

Recovery steps – combination

1. Boot the server into Windows Recovery Environment. You can do this either by booting the server from the Windows Setup disc that has the same architecture of the system that you are trying to recover or select the **Settings** charm > **Power**, and then hold the Shift key while clicking **Restart**.
2. After either of these actions is performed, the Boot Options menu will be displayed. Select **Start recovery, troubleshooting, and diagnostic tools**.
3. Click **Re-image your computer**.
4. Setup searches the hard disk drives for an existing Windows installation and then displays the results in **System Recovery Options**. Click **Next**.
5. On the **System Recovery Options** page, click **System Image Recovery**. This opens the **Re-image your computer** page.
6. Select **Only restore system drives**
7. On the **Choose additional restore options** page, do the following optional tasks, and then click **Next**:
   - Select the **Format and repartition disks** check box to delete existing partitions and reformat the destination disks to be the same as the backup. This enables the **Exclude disks** button. Click this button and then select the check boxes associated with the storage space .The disk that contains the backup that you are using is automatically excluded.

     Notes
   - Unless a disk is excluded, data on it can be lost—regardless of whether it was part of the backup or whether it has volumes that are being restored.
   - You should not exclude the system disk—the first disk in the BIOS boot order. (This disk is usually referred to as Disk 0, but in some conditions, Diskmgmt.msc and Diskpart.exe may label it as something else, for example, Disk 1/2.) If the system disk (Disk 0) is excluded then Windows will try to do recovery on BIOS Disk 1. But after the recovery, the system will not start and it may fail with an error that **Bootmgr** is missing. The BIOS will always use the first disk in the boot order to search for this file and if it is missing the computer will not start.
   - In **Exclude disks**, if you do not see all the disks that are attached to the computer, you might need to install the associated drivers for the storage device.
   - Select the **Only restore system drives** check box to perform an operating system–only recovery.
   - Click **Install drivers** to install device drivers for the hardware that you are recovering to.

- Click **Advanced** to specify whether the computer is automatically restarted and the disks are checked for errors immediately after the recovery.

8. Confirm the details for the restoration, and then click **Finish**. The recovery will succeed as long as all the critical volumes (volumes containing operating system components) are recovered.
9. Re-configure the storage pool either from Server Manager or using the Windows PowerShell cmdlets. For detailed steps, see To create a storage pool
10. Create virtual disks for the storage space in the storage pool. For detailed steps, see To create a virtual disk.
11. Create volumes in the storage space. For detailed steps, see To create a volume

    Important
    Make sure that the volumes you create are either of the same or greater size as the volumes included in the backup you are restoring
12. Use Windows Server Backup to recover the volumes to the storage spaces.
13. Use Windows Server Backup to recover the volumes for other physical volumes.

Example 2: Same server recovery without disk recreation

This example describes performing a bare-metal recovery without disk recreation when some or all the storage pool volumes are part of backup. The sample server configuration has the following details:

- Storage pool and storage space virtual disks are configured on the server.
- Data volumes have been created on the storage space virtual disks.
- Windows Server Backup has been run and performed either a Full Server Backup or a backup for bare-metal recovery.
- The backup includes some or all of the volumes on the storage space virtual disks.

Recovery is being performed for the following reasons:

- Operating system instability
- System file corruption
- Persistent malware

To recover the operating system to a previous known good state the administrator wants to recover the server using a Full Server Backup that includes System volumes and data volumes in a storage pool. There is no need for partition layout and disk configuration changes between backup and recovery. The storage space volumes are in a healthy state and have an attached/online status when viewed from Windows Recovery Environment.

Recovery steps – full server

1. Boot the server into Windows Recovery Environment. You can do this either by booting the server from the Windows Setup disc that has the same architecture of the system that

you are trying to recover or select the **Settings** charm > **Power**, and then hold the Shift key while clicking **Restart**.

2. After either of these actions is performed, the Boot Options menu will be displayed. Select **Start recovery, troubleshooting, and diagnostic tools**.
3. Click **Re-image your computer**.
4. Setup searches the hard disk drives for an existing Windows installation and then displays the results in **System Recovery Options**. Click **Next**.
5. On the **System Recovery Options** page, click **System Image Recovery**. This opens the **Re-image your computer** page.
6. Select **Use the latest available system image (recommended).**
7. On the **Choose additional restore options** lear the **Format and repartition disks** check box page and then click **Next**.
8. Confirm the details for the restoration, and then click **Finish**. The recovery will succeed as long as all the critical volumes (volumes containing operating system components) are recovered.
9. Log on to the server and verify that all data and storage pools are recovered.

Example 3: Alternate server recovery

This example describes performing a bare-metal recovery onto a new server when some or all the storage pool volumes are part of backup. The sample server configuration has the following details:

- Storage pool and storage space virtual disks are configured on the server.
- Data volumes have been created on the storage space virtual disks.
- Windows Server Backup has been run and performed either a Full Server Backup or a backup for bare-metal recovery.
- The backup includes some or all of the volumes on the storage space virtual disks.

Recovery is being performed for the following reasons:

- Hardware replacement
- Operating system migration

In this situation the storage space no longer exists and the physical storage must be used to create both the operating system and the storage pool.

Recovery steps – alternate server

1. Boot the server into Windows Recovery Environment. You can do this either by booting the server from the Windows Setup disc that has the same architecture of the system that you are trying to recover or select the **Settings** charm > **Power**, and then hold the Shift key while clicking **Restart**.
2. After either of these actions is performed, the Boot Options menu will be displayed. Select **Start recovery, troubleshooting, and diagnostic tools**.
3. Click **Re-image your computer**.

4. Setup searches the hard disk drives for an existing Windows installation and then displays the results in **System Recovery Options**. Click **Next**.
5. On the **System Recovery Options** page, click **System Image Recovery**. This opens the **Re-image your computer** page.
6. Click **Restore a different backup**, and then do one of the following:
   - On the **Select the location of the system image page**, click the computer that contains the backup that you want to use, and then click **Next**.

     Important
     If the storage location contains backups of multiple computers, make sure that you click the row for the backups for the computer that you want to use.

     Then, on the **Select the system image to restore** page, click the backup that you want to use, and then click **Next**.
   - Click **Advanced** to browse for a backup in a remote shared folder on the network and provide the Universal Naming Convention (UNC) path to the backup, or if your backup is on a device, to install a device driver. (To install a drive, the driver needs to be present in local system. You cannot install a driver from the network and, instead, need to provide a local path to the .inf file to install a driver.) Click **Next**.

     Important
     If you use a domain in your environment, and if the backup storage location is on a computer that is a member of that domain, then computer containing the storage location should be on the IPsec boundary to be accessible by non-domain computers. When a computer boots into Windows Recovery Environment, it becomes a non-domain computer and cannot access the usual network shares. Only those computers that allow non-domain computers to access the share can be used as a backup storage location in this way.

     You can also address this issue by adding the computer that is your backup storage location to a workgroup and placing the backup in a shared folder. A computer running Windows Recovery Environment behaves as if it is in a workgroup, enabling you to access the shared folder with the backup.
7. Select **Only restore system drives**
8. On the **Choose additional restore options** page, do the following optional tasks, and then click **Next**:
   - Select the **Format and repartition disks** check box to delete existing partitions and reformat the destination disks to be the same as the backup. This enables the **Exclude disks** button. Click this button and then select the check boxes associated with the physical disks you are going to use for the storage pool .The disk that contains the backup that you are using is automatically excluded.

     Notes
   - Unless a disk is excluded, data on it can be lost—regardless of whether it was part of the backup or whether it has volumes that are being restored.

- You should not exclude the system disk—the first disk in the BIOS boot order. (This disk is usually referred to as Disk 0, but in some conditions, Diskmgmt.msc and Diskpart.exe may label it as something else, for example, Disk 1/2.) If the system disk (Disk 0) is excluded then Windows will try to do recovery on BIOS Disk 1. But after the recovery, the system will not start and it may fail with an error that **Bootmgr** is missing. The BIOS will always use the first disk in the boot order to search for this file and if it is missing the computer will not start.
- In **Exclude disks**, if you do not see all the disks that are attached to the computer, you might need to install the associated drivers for the storage device.
- Select the **Only restore system drives** check box to perform an operating system–only recovery.
- Click **Install drivers** to install device drivers for the hardware that you are recovering to.
- Click **Advanced** to specify whether the computer is automatically restarted and the disks are checked for errors immediately after the recovery.

9. Confirm the details for the restoration, and then click **Finish**. The recovery will succeed as long as all the critical volumes (volumes containing operating system components) are recovered.
10. Log on to the server and verify that the operating system image is as you expected.
11. Re-configure the storage pool either from Server Manager or using the Windows PowerShell cmdlets. For detailed steps, see To create a storage pool
12. Create virtual disks for the storage space in the storage pool. For detailed steps, see To create a virtual disk.
13. Create volumes in the storage space. For detailed steps, see To create a volume

    Important
    Make sure that the volumes you create are either of the same or greater size as the volumes included in the backup you are restoring
14. Use Windows Server Backup to recover the volumes to the storage spaces.
15. Use Windows Server Backup to recover the volumes for other physical volumes.

Critical volumes and storage spaces

Bare-metal recovery treats Boot, System, and System State Writer reported volumes as Critical. All of these volumes must be recoverable from within Windows Recovery Environment for a BMR restore to succeed. If a System State Writer marks a storage space volume as critical then you should make sure that a suitable virtual disk or physical disk of equal or greater size is in place to recover the critical volume.

In the alternate hardware example because the storage space can't be created in Windows RE; BMR would require a physical disk to use instead of a storage space in order to recover the critical volume. You should make sure that enough physical disks with right sizes are attached to the server to recover critical volumes and complete BMR. Data volumes should be recovered

from within the re-imaged operating system and storage pool. The following recommendations will help you obtain the necessary information to recover successfully:

- Windows Server Backup does not display the disk sizes required for BMR to succeed. You should record the disk sizes requirements along with your backup data.
- If a single volume is hosted on the disk then volume size would be similar to disk size and user should attach a disk of equal or higher size.

For example, consider a server backup that includes a critical volume hosted on storage pool space (virtual disk) and BMR is performed to an alternate server where the storage pool does not exist. When BMR starts it queries the server to find suitable disks on which to recover the critical volumes. If enough disks with required sizes are not found then recovery would fail with a generic message - "**Not all disks required to recover the System are present. Please attach the disks and retry recovery**".

At this point you can use the following procedure to query the volume size. For this example, the following parameters are used:

Version	07/27/2013-22:10
Backup Target	D:

1. Open a command prompt with administrator privileges and run the following command to see a list of Critical volumes:

   ```
   wbadmin start sysrecovery –version:07/27/2013-22:10 -
   backuptarget:D:
   ```

2. The critical volumes will be listed after the phrase **You have chosen to recover volume(s)**. Record the volumes listed then press N.

   Warning
   Do not continue recovery.

3. Then, use the following command to get the size of each volume included in the backup target:

   ```
   wbadmin get items –version:07/27/2013-22:10 -backuptarget:D:
   ```

   Match up the critical volumes from your list with the sizes listed for that volume.

Warning
This method works when there is only a single volume per disk. If a storage space has multiple volumes, the amount of space required cannot be easily estimated. It is best to make sure that data is stored separately from critical volumes so that data can be recovered in a separate process.

Special Considerations when using RAID and Storage Spaces

The Storage Spaces can be configured with different RAID/Resiliency protection level and other configuration settings. As mentioned above, ASR (which is used to recreate partition layout during BMR) is not aware of RAID configuration. Hence, it will treat all virtual disks as equal without considering the underlying resiliency settings protection. ASR will use the best fit algorithm and recover the volumes.

In certain cases, it might recover a volume hosted on a storage space with resiliency setting Simple to another storage space with resiliency setting Mirrored or vice versa if the storage space virtual disks are recreated between backup and recovery.

If you chose to recover all volumes during recovery then the volumes might be recovered to different disks than you desire. To avoid unwanted disk assignment, make sure that all the non-critical disks including physical disks and Storage Spaces are excluded from BMR recovery.

Sample Scripts

The following sample scripts are provided for you to modify as necessary for your own use.

WMIC Script to enumerate storage pools which are Read/Only, and change them to Read/Write:

The following WMIC script can be used to modify the storage pool's Read/Write state from within Windows RE, where PowerShell is unavailable. This script will enumerate all storage pools with **IsReadonly == $True**, and then set them to **$False**.**SetAttributes for StoragePool** takes two input arguments, **ClearOnDeallocate** and **IsReadOnly**. 0 is the default value for **ClearOnDeallocate**.

```
echo off

set _isReadOnlyValue="TRUE"

for /f "usebackq skip=1" %%i in (`wmic /namespace:\\root\microsoft\windows\storage  path
MSFT_StoragePool get IsReadOnly`) do (

   echo I am here...

         if /i "%%i"==%_isReadOnlyValue%  (

        echo Found the pool that IsReadOnly == %_isReadOnlyValue%

               wmic /namespace:\\root\microsoft\windows\storage  path  MSFT_StoragePool
Where IsReadOnly=%_isReadOnlyValue%  call SetAttributes  0,0

   )

)

echo on
```

WMIC Script to Enumerate ManualAttach Spaces, and change them to automatically attach

Note

The storage pool exposing the storage spacemust be in Read/Write mode in order to perform this operation.

The following script differs from the previous on by looping over VirtualDisks that **IsManualAttach** is TRUE, and then sets the **IsManualAttach** to false.

SetAttributes for VirtualDisk takes two input arguments, **Access** and **IsManualAttach**. 0 is the default value for **Access**.

```
echo off

set _isManualAttachValue="TRUE"

for /f "usebackq skip=1" %%i in (`wmic /namespace:\\root\microsoft\windows\storage  path
MSFT_VirtualDisk get IsManualAttach`) do (

   echo I am here...

     if /i "%%i"==%_isManualAttachValue%  (

        echo Found the pool that IsReadOnly == %_isManualAttachValue%

         wmic /namespace:\\root\microsoft\windows\storage  path  MSFT_VirtualDIsk Where
IsManualAttach=%_isManualAttachValue%  call SetAttributes  0,0
```

Windows Server Essentials Experience Overview

This topic applies to a server running Windows Server 2012 R2 Essentials, or to a server running Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter with the Windows Server Essentials Experience role installed.

The Windows Server Essentials Experience role is available in Windows Server 2012 R2 Standard and Windows Server 2012 R2 Datacenter. When you install the Windows Server Essentials Experience role, you can take advantage of all the features that are available to you in Windows Server 2012 R2 Essentials without the locks and limits enforced in it. The Windows Server Essentials Experience role enables you to:

- Protect your server and client data by backing up your server and all the client computers within your network.
- Manage your users and groups through the simplified server Dashboard. In addition, integration with Windows Azure Active Directory enables easy data access for Microsoft online services users (for example, Office 365, Exchange Online, and SharePoint Online users) through their domain credentials.
- Store your company's data in a centralized location.
- Integrate your server with Microsoft online services, such as Office 365, Exchange Online, SharePoint Online, and Windows Intune.

    - Integration with Office 365 enables you to synchronize and manage Office 365 user accounts and access through the Dashboard.
    - Integration with SharePoint Online enables you to create and manage SharePoint libraries through the Dashboard.
- Use the Anywhere Access functionalities on the server (such as Remote Web Access and virtual private networks) to access your server, network computers, and data from remote locations in a highly secure manner.
- Access data from any location and on any device by using the company's customized web portal (through Remote Web Access).
- Manage mobile devices that access company emails by using Office 365 through Active Sync from the Dashboard.
- Monitor network health and obtain customized health reports. The reports can be created on demand, customized, and sent by email to specified recipients.

For more information about Windows Server Essentials Experience, see the following topics:

- **Install and Configure Windows Server 2012 R2 Essentials**
- **Overview of the Dashboard in Windows Server Essentials**
- **Manage Windows Server Essentials**
- **Use Windows Server Essentials**

See also

- **Windows Server 2012 R2 Essentials and Windows Server 2012 Essentials**
- Server Roles and Technologies in Windows Server 2012 R2 and Windows Server 2012

Windows Server Update Services Overview

The Windows Server Update Service (WSUS) enables information technology administrators to deploy the latest Microsoft product updates. By using WSUS, administrators can fully manage the distribution of updates that are released through Microsoft Update to computers in their network. This topic provides an overview of this server role and more information about how to deploy and maintain WSUS.

Did you mean…

- Windows Update
- Software Deployment

WSUS server role description

The WSUS server provides the features that administrators need to manage and distribute updates through a management console. In addition, a WSUS server can be the update source for other WSUS servers within the organization. The WSUS server that acts as an update source is called an *upstream server*. In a WSUS implementation, at least one WSUS server in the

network must connect to Microsoft Update to get available update information. The administrator can determine, based on network security and configuration, how many other servers connect directly to Microsoft Update.

Practical applications

Update management is the process of controlling the deployment and maintenance of interim software releases into production environments. It helps you maintain operational efficiency, overcome security vulnerabilities, and maintain the stability of your production environment. If your organization cannot determine and maintain a known level of trust within its operating systems and application software, it might have a number of security vulnerabilities that, if exploited, could lead to a loss of revenue and intellectual property. Minimizing this threat requires you to have properly configured systems, use the latest software, and install the recommended software updates.

The core scenarios where WSUS adds value to your business are:

- Centralized update management
- Update management automation

New and changed functionality

Warning

Upgrade from any version of Windows Server that supports WSUS 3.2 to Windows Server® 2012 R2 requires that you first uninstall WSUS 3.2.

In Windows Server 2012, upgrading from any version of Windows Server with WSUS 3.2 installed is blocked during the installation process if WSUS 3.2 is detected, and you are prompted to first uninstall Windows Server Update Services prior to upgrading Windows Server 2012.

However, because of changes in Windows Server 2012 R2, when upgrading from any version of Windows Server and WSUS 3.2 to Windows Server 2012 R2, the installation is not blocked. Failure to uninstall WSUS 3.2 prior to performing a Windows Server 2012 R2 upgrade will cause the post installation tasks for WSUS in Windows Server 2012 R2 to fail. In this case, the only known corrective measure is to format the hard drive and reinstall Windows Server 2012 R2.

Windows Server Update Services is a built-in server role that includes the following enhancements:

- Can be added and removed by using the Server Manager
- Includes Windows PowerShell cmdlets to manage the ten most important administrative tasks in WSUS
- Adds SHA256 hash capability for additional security
- Provides client and server separation: Versions of the Windows Update Agent (WUA) can ship independently of WSUS

Feature and functionality	Windows Server 2008 R2	Windows Server 2012 and Windows Server 2012 R2
Inclusion of Windows PowerShell cmdlets to manage the ten most important administrative tasks in WSUS		X
Security enhancements with SHA256 hash capability		X
Client and server separation: Versions of the Windows Update Agent (WUA) can ship independently of WSUS		X

Using Windows PowerShell to manage WSUS

For system administrators to automate their operations, they need coverage through command-line automation. The main goal is to facilitate WSUS administration by allowing system administrators to automate their day-to-day operations.

What value does this change add?

By exposing core WSUS operations through Windows PowerShell, system administrators can increase productivity, reduce the learning curve for new tools, and reduce errors due to failed expectations resulting from a lack of consistency across similar operations.

What works differently?

In earlier versions of the Windows Server operating system, there were no Windows PowerShell cmdlets, and update management automation was challenging. The Windows PowerShell cmdlets for WSUS operations add flexibility and agility for the system administrator.

Removed or deprecated functionality

In this release of WSUS, the API is updated to the .NET Framework 4.5.

Hardware requirements

The minimum hardware configuration for WSUS is:

- Processor: 1.4 gigahertz (GHz) x64 processor (2 Ghz or faster is recommended)
- Memory: WSUS requires an additional 1.5 GB of RAM - above and beyond what is required by Windows Server 2012 R2 or Windows Server 2012.
- Available disk space: 10 GB (40 GB or greater is recommended)
- Network adapter: 100 megabits per second (Mbps) or greater

See also

The following table lists references for more information about WSUS:

Content type	References
Product evaluation	Windows Server Update Services TechNet Page
Planning and Deployment	Deploy Windows Server Update Services in Your Organization
Migration	Migrate Windows Server Update Services to Windows Server 2012
Community resources	WSUS Forum

Migrate Windows Server Update Services to Windows Server 2012

This document describes the process to migrate an existing Windows Server Update Services (WSUS) 3.0 SP2 server role to a destination server that is running Windows Server 2012 or Windows Server 2012 R2. This document includes instructions for moving the updates, settings, target groups, and computers to the new server. By using the tools that are described in this document, you can simplify the migration process, reduce migration time, increase the accuracy of the migration process, and help eliminate possible conflicts that might otherwise occur during the migration process.

- Step 1: Plan for WSUS Migration
- Step 2: Prepare to Migrate WSUS
- Step 3: Migrate WSUS
- Step 4: Verify the WSUS Migration

Step 1: Plan for WSUS Migration

The first step in the migration of your Windows Server Update Services (WSUS) to Windows Server 2012 or Windows Server 2012 R2 is to understand the supported and unsupported scenarios and the supported operating systems for this migration. The following checklist describes the steps involved in planning for your WSUS migration.

Task	Description
1.1. Know supported operating systems	Review the list of supported source operating systems and WSUS versions.
1.2. Review supported migration scenarios	Review the list of supported migration scenarios.
1.3. Review migration scenarios that are not supported	Review the list of unsupported migration scenarios.

1.1. Know supported operating systems

Migration from the following operating systems is supported on Windows Server 2012 and Windows Server 2012 R2:

- Windows Server 2008 R2 running WSUS 3.0 SP2
- Windows Server 2008 (full installation option) running WSUS 3.0 SP2
- Windows Server 2003 SP2 running WSUS 3.0 SP2

1.2. Review supported migration scenarios

The following WSUS migration scenarios are supported:

- Windows Server 2012 Standard and Datacenter editions, and servers running Windows Server 2012 R2 can be used as source or destination servers.
- Windows Server 2012 Enterprise edition can be used as a source server.
- Migration between physical operating systems and virtual operating systems.
- Migration from a source server that is running SQL Server 2005 to a destination server that is running SQL Server 2008 R2 SP1.
- Migration from a source server that is running Windows Internal Database to a destination server that is running SQL Server 2008 R2 SP1.
- Migration from a domain to a workgroup or from a workgroup to a domain. However, if the source server is running SQL Server from a remote location, migration from the domain to a workgroup is not supported.
- The destination server must meet the Windows Server 2012 or Windows Server 2012 R2 WSUS role minimum system requirements for hardware and software.

Important

For more information about minimum system requirements and hardware capacity requirements for the WSUS server, see the Deploy Windows Server Update Services in Your Organization.

1.3. Review migration scenarios that are not supported

The following WSUS migration scenarios are not supported:

- Migration from an unsupported version of WSUS (prior to WSUS 3.0 SP2). Upgrade the existing WSUS server to a supported version before you migrate the WSUS server role to Windows Server 2012 or Windows Server 2012 R2.
- Migration from a Server Core installation option (WSUS 3.0 SP2 does not support a Server Core installation).
- Migration from a domain that is using SQL Server from a remote location to a workgroup.
- Migration from a source server that is running SQL Server to a destination server that is running Windows Internal Database.
- Migration from a source server that stores updates on Microsoft Update to a destination server that stores updates on a local WSUS server, and vice versa. Changing the configuration during the migration process is not supported.
- Migration from a source server to a destination server that is running an operating system in a different system UI language. The system UI language is the language of the localized installation package that was used to set up the Windows operating system. For example, you cannot use Windows Server Migration Tools to migrate roles, operating system settings, data, or shares from a computer that is running Windows Server 2008 in the French system UI language to a computer that is running Windows Server 2008 R2 in the German system UI language.

See also

- Step 2: Prepare to Migrate WSUS
- Migrate Windows Server Update Services to Windows Server 2012
- WSUS server role description

Step 2: Prepare to Migrate WSUS

The second step in the migration of your Windows Server Update Services (WSUS) server role involves preparing the destination and source servers. The following checklist describes the steps involved to prepare for your WSUS migration.

Task	Description
2.1. Prepare before you start the migration	Review the recommended guidelines before starting the migration process.
2.2. Prepare the destination server	Understand the steps that must be completed on the WSUS destination server before the

Task	Description
	migration.
2.3. Prepare the source server	Understand the steps that must be completed on the WSUS source server before the migration.

2.1. Prepare before you start the migration

This migration procedure assumes a working knowledge of deployment basics for Windows Server Update Services (WSUS) 3.0 SP2. For more information about WSUS deployment, see Deploy Windows Server Update Services in Your Organization. We recommend that you make the following decisions and preparations before you start the migration process:

Warning

Upgrade from any version of Windows Server that supports WSUS 3.2 to Windows Server® 2012 R2 requires that you first uninstall WSUS 3.2.

In Windows Server 2012, upgrading from any version of Windows Server with WSUS 3.2 installed is blocked during the installation process if WSUS 3.2 is detected, and you are prompted to first uninstall Windows Server Update Services prior to upgrading Windows Server 2012.

However, because of changes in Windows Server 2012 R2, when upgrading from any version of Windows Server and WSUS 3.2 to Windows Server 2012 R2, the installation is not blocked. Failure to uninstall WSUS 3.2 prior to performing a Windows Server 2012 R2 upgrade will cause the post installation tasks for WSUS in Windows Server 2012 R2 to fail. In this case, the only known corrective measure is to format the hard drive and reinstall Windows Server 2012 R2.

- Configure a location to store updates on the source server. Changing the content storage configuration as part of the migration process is not supported. You can store updates on the local WSUS server or on Microsoft Update. If you want the destination server to store updates in a different location than the source server, the new location must be configured on the source server before migration.
- Confirm that the destination server meets the minimum WSUS hardware requirements and database requirements. For more information about those requirements see Deploy Windows Server Update Services in Your Organization on Microsoft TechNet.

2.2. Prepare the destination server

Before migrating WSUS, set up a new Windows Server 2012 in your organization as the WSUS destination server and install WSUS server role on the destination server. After you have successfully installed the WSUS server role, the Configuration Wizard starts automatically. Close the Configuration Wizard. Do not try to sync the updates at this point, because you will copy the

update binary files later in the migration process. The WSUS installation procedure assumes that updates for the new server come from Windows Update.

After this is complete, follow these guidelines:

- If you have decided to use the full installation of SQL Server as the WSUS database, install SQL Server 2008 R2 Standard or SQL Server 2008 R2 Enterprise.
- Download a graphical tool to manage your database on the destination server from Microsoft SQL Server Management Studio Express or Microsoft SQL Server 2008 R2 Management Studio Express.
- Open TCP port 7000 and make sure that it is not being used by other applications. This port is used by Send-SmigServerData and Receive-SmigServerData to perform the data transfer.
- If the destination server is not joined to the source server's domain, visually verify that the time, date, and time zone on the destination server are synchronized with the source server. Use the Windows Control Panel to update the date and time if it is necessary.

Important

For more information about minimum system requirements and hardware capacity requirements for the WSUS server, see the Deploy Windows Server Update Services in Your Organization.

2.3. Prepare the source server

Review and take action based on the following guidelines:

- Refer to Appendix A: Migration Data Collection Worksheet to collect data about the source server.
- Open TCP port 7000 and make sure that it is not being used by other applications. This port is used by Send-SmigServerData and Receive-SmigServerData to perform the data transfer.
- If you have changed the default behavior of Windows Firewall (or another firewall program) to block outgoing traffic on computers that are running Windows Server 2012, you must enable outgoing traffic on UDP port 7000.
- Download a graphical tool for managing your database on the source server at Microsoft SQL Server Management Studio Express.

See also

- Step 3: Migrate WSUS
- Step 1: Plan for WSUS Migration
- WSUS server role description

Step 3: Migrate WSUS

During the third step in the migration of your Windows Server Update Services (WSUS) server role, you will migrate binaries and security groups, back up the database, change the server identity, and apply security settings The following checklist describes the steps involved.

Task	Description
3.1. Migrate WSUS update binaries	Move WSUS update binaries from the source server to the destination server.
3.2. Migrate WSUS security groups	Migrate local users and groups manually or by using Windows Server Migration Tools.
3.3. Back up the WSUS database	Use SQL Server Management Studio to back up and restore the WSUS database, computer groups, update approvals, and WSUS settings.
3.4. Change the WSUS server identity	Change the WSUS server identity on the destination server. Performing this step ensures that there is no effect on clients that are managed by WSUS during the migration process.
3.5. Apply security settings	Configure security settings on the new server. This includes configuring security settings on the destination server that you were using on the source server.
3.6. Review additional considerations	Review some additional actions that you should take after the migration is complete.

3.1. Migrate WSUS update binaries

Use your preferred method to copy WSUS update binaries in the WSUS folder from the source server to the destination server (for example, Windows Server Migration Tools, Windows Explorer, Xcopy, or Robocopy). If you decide to use Windows Server Migration Tools to migrate WSUS update binaries to a destination server that is running Windows Server 2012, see Migrate WSUS Update Binaries from the Source Server to the Destination Server Using Windows Server Migration Tools.

Important
Migrating WSUS update binaries is unnecessary if update files are stored on Microsoft Update.

3.2. Migrate WSUS security groups

You have the option of manually migrate only the WSUS Administrators and WSUS Reporters local security groups. Or, you can use Windows Server Migration Tools to migrate all local users and groups (including the WSUS Administrators and WSUS Reporters local security groups) from the source server to the destination server.

Warning
The WSUS Server Migration Tools can be installed on the server using the Server Manager Add Features option.

Before you perform this procedure, verify that the destination server can resolve the names of domain users who are members of the local group during the import operation. If the source and destination servers are in different domains, the destination server must be able to contact a global catalog server for the forest in which the source domain user accounts are located. Use the following guidelines:

- **If the source server is a member of the domain and the destination server is a domain controller**: Imported local users are elevated to domain users, and imported local groups become Domain Local groups on the destination server.
- **If the source server is a domain controller, and the destination server is not**: Domain Local groups are migrated as local groups, and domain users are migrated as local users.

Use the following procedure to manually migrate users to the WSUS Administrators and WSUS Reporters local security groups.

To manually migrate local users and groups

1. Right-click in the Taskbar, click **Properties**, highlight **Toolbars**, and then click **Address**.
2. Type **lusrmgr.msc**, and then press ENTER.
3. In the console tree of the Local Users and Groups MMC snap-in, double-click **Users**.
4. Manually create a list of the local users.
5. In the console tree of the Local Users and Groups MMC snap-in, double-click **Groups**.
6. Manually add the users from the source server to the WSUS Administrators and WSUS Reporters groups.

Use the following procedure to use Windows Server Migration Tools to migrate users to the WSUS Administrators and WSUS Reporters local security groups.

To use Windows Server Migration Tools to migrate users

1. Open a Windows PowerShell session on the source server and on the destination server.
2. Type the command below and press ENTER:

   ```
   Add-PSSnapin Microsoft.Windows.ServerManager.Migration
   ```

3. In the Windows PowerShell session on the source server, type the following command to export local users and groups to a migration store:

   ```
   Export-SmigServerSetting -User <Enabled | Disabled | All> -Group -Path
   ```

```
<MigrationStorePath> -Verbose
```

MigrationStorePath represents the path of the location where you want to store migrated data. You can also use one of the following values with the **-User** parameter:

- **Enabled**: Export only enabled local users
- **Disabled**: Export only disabled local users
- **All**: Export enabled and disabled local users

4. Press ENTER.

Important

You are prompted to provide a password to encrypt the migration store. Remember this password, because you must provide the same password to import data from the migration store on the destination server. If the path is not a shared location to which the destination server has access, you must copy the migration store to the destination server manually, or to a location that this destination server can access as it runs the **Import-SmigServerSetting** cmdlet.

5. In the Windows PowerShell session on the destination server, type the following command to import local users and groups from the migration store that you created in Step 2:

```
Import-SmigServerSetting -User <Enabled | Disabled | All> -Group -Path
<MigrationStorePath> -Verbose
```

MigrationStorePath represents the path of the location from which you want to import migrated data. You can also use one of the following values with the **-User** parameter:

- **Enabled**: Export only enabled local users
- **Disabled**: Export only disabled local users
- **All**: Export enabled and disabled local users

6. Press ENTER.

Warning

After you enter the **Import-SmigServerSetting** cmdlet, you are prompted to provide the same password to decrypt the migration store that you created during the export process.

3.3. Back up the WSUS database

WSUS servers can be configured to use Windows Internal Database, the database software that is included with WSUS, or the full version of SQL Server. Regardless of which database option the source server is running, perform the following procedures to back up the WSUS database on the source server and restore the database to the destination server.

For an overview of backup and command-line syntax, see the following topics in SQL Server TechCenter:

- Backup Overview

- BACKUP (Transact-SQL)

For an overview of restore and command-line syntax, see the following topics in SQL Server TechCenter:

- Restore and Recovery Overview
- RESTORE (Transact-SQL)

Important
SQL Server Management Studio must be run with elevated administrator permissions throughout this procedure.

To back up the WSUS database on the source server

1. After you connect to the appropriate instance of the database in Object Explorer, click the server name to expand the server tree.

   Note
   If the source server is using Windows Internal Database, the query changes depending on which version of WSUS you are currently running. For WSUS 3.2, the query is: **\\.\pipe\mssql$microsoft##ssee\sql\query**, and for WSUS on Windows Server 2012, the query is: **\\.\pipe\Microsoft##WID\tsql\query**.
2. Expand **Databases**, and select the **SUSDB** database.
3. Right-click the database, point to **Tasks**, and then click **Back Up**. The **Back Up Database** dialog box appears.
4. In the **Database** list, verify the database name.
5. In the **Backup type** list, select **Full**.
6. Select **Copy Only Backup**. A copy-only backup is a SQL Server backup that is independent of the sequence of conventional SQL Server backups.
7. For **Backup component**, click **Database**.
8. Accept the default backup set name that is suggested in the **Name** text box, or enter a different name for the backup set.
9. Optionally, in the **Description** text box, enter a description of the backup set.
10. Specify when the backup set will expire and can be overwritten without explicitly skipping verification of the expiration data.
11. Choose the backup destination by clicking Disk.

    Important
    To remove a backup destination, select it and then click **Remove**. To view the contents of a backup destination, select it and then click **Contents**.
12. In the **Select a page** pane, click **Options** to view or select the advanced options.
13. On the **Overwrite Media** option, click one of the following:
    - **Back up to the existing media set** For this option, click **Append to the existing backup set** or **Overwrite all existing backup sets**. Optionally:
      - Click **Check media set name and backup set expiration** to cause the backup

operation to verify the date and time at which the media set and backup set expire.

- Enter a name in the **Media set name** text box. If no name is specified, a media set with a blank name is created. If you specify a media set name, the media (tape or disk) is checked to see whether the name matches the name that you enter here.

- **Back up to a new media set, and erase all existing backup sets** For this option, enter a name in the **New media set name** text box, and, optionally, describe the media set in the **New media set description** text box.

14. In the **Reliability** section, optionally select:
    - **Verify backup when finished.**
    - **Perform checksum before writing to media.**
    - **Continue on checksum error.**
15. SQL Server 2008 Enterprise support backup compression. By default, whether a backup is compressed depends on the value of the **Backup-compression default** server configuration option. Regardless of the current server-level default, you can compress the backup or prevent compression at this time. To compress the backup, select **Compress backup**, or to prevent compression, select **Do not compress backup**.

To restore the WSUS database backup on the destination server by using SQL Server Management Studio

1. After you connect to the appropriate instance of the database in Object Explorer, click the server name to expand the server tree.

   Important
   If the source server is using Windows Internal Database, the database name is: **\\.\pipe\Microsoft##WID\tsql\query**.
2. Click **New Query** and copy the following SQL command to drop the WSUS database

   USE masterGOALTER DATABASE SUSDB SET SINGLE_USER WITH ROLLBACK IMMEDIATEGODROP DATABASE SUSDBGO
3. Click **Execute**, to run the query
4. Run the following query:

   RESTORE DATABASE [SUSDB] FROM DISK = N'C:\SUSDB.bak' WITH FILE = 1, MOVE N'SUSDB' TO N'c:\Windows\WID\Data\susdb.mdf', MOVE N'SUSDB_log' TO N'c:\Windows\WID\Data\SUSDB_log.ldf', NOUNLOAD, STATS = 10

   Important
   Drive C: that is used in this example will vary according to the actual storage location for the files.
5. In the **Backup type** list, select **Full**.

   Warning
   Running the previous query will result in the following error message:

Msg 3605, Level 16, State 1, Line 5 Schema verification failed for database 'SUSDB'.Msg 3013, Level 16, State 1, Line 5RESTORE DATABASE is terminating abnormally

Disregard the error message and continue.

6. Open an elevated command prompt in Windows Server 2012, and run the following command:

   %programfiles%\update services\tools\wsusutil postinstall [sql parameter] [content parameter]

   Important
   For WID, do not specify the SQL parameter.

To restore the WSUS database backup on the destination server by using SQL Server Management Studio

1. After you connect to the appropriate instance of the database in **Object Explorer**, click the server name to expand the server tree.

   Important
   If the source server is using Windows Internal Database, the database name is *\\.\pipe\Microsoft##WID\tsql\query*.

2. Click **New Query** and copy the following SQL command to drop the WSUS database (SUSDB):

```
USE master

GO

ALTER DATABASE SUSDB SET SINGLE_USER WITH ROLLBACK IMMEDIATE

GO

DROP DATABASE SUSDB

GO
```

3. Click **Execute** to run the query.
4. Run the following query:

```
RESTORE DATABASE [SUSDB] FROM  DISK = N'C:\SUSDB.bak' WITH  FILE = 1,  MOVE
N'SUSDB' TO N'c:\Windows\WID\Data\susdb.mdf',  MOVE N'SUSDB_log' TO
N'c:\Windows\WID\Data\SUSDB_log.ldf',  NOUNLOAD,  STATS = 10
```

   Important
   Drive C: that is mentioned in this example will vary according to the actual storage location for the files.

5. This will result in the following error message:

```
Msg 3605, Level 16, State 1, Line 5

Schema verification failed for database 'SUSDB'.

Msg 3013, Level 16, State 1, Line 5
```

```
RESTORE DATABASE is terminating abnormally
```

Disregard the error message and continue.

6. Open an elevated command prompt in Windows Server 2012, and run the following command:

```
%programfiles%\update services\tools\wsusutil postinstall [sql parameter] [content parameter]
```

Important
For WID, do not specify the SQL parameter.

When a database is restored to a different server, it contains a set of users and permissions, although there may be no corresponding user log on information, or the log on information may not be associated with the same users. This condition is known as having "orphaned users." See article 168644 in the Microsoft Knowledge Base for instructions about how to resolve orphaned users.

After you restore a SQL Server 2005 database to SQL Server 2008, the database becomes available immediately, and it is then automatically upgraded.

3.4. Change the WSUS server identity

The WSUS server identity on the destination server must be changed. Performing this step guarantees that WSUS-managed clients are not affected during the migration process. If the source server and the destination server run with the same identity, and a change is made to one of the servers, the communication between the client and server will fail.

To change the WSUS server identity

1. On the destination server, open a Windows PowerShell session with elevated user rights and run the following script:

```
$updateServer = get-wsusserver

$config = $updateServer.GetConfiguration()

$config.ServerId = [System.Guid]::NewGuid()

$config.Save()
```

2. As soon as the server identity is changed, run the following command to generate a new encryption key:

```
%ProgramFiles%\Update Services\Tools\wsusutil.exe postinstall
```

3.5. Apply security settings

Refer to the settings that you recorded in the Migration Data Collection Worksheet, and then complete the following tasks to apply the security settings that you were using on the source server to the destination server.

- **SMTP server settings**: If you are using an authenticating proxy or the email notification feature to an SMTP server that requires a password (or both), you must manually configure the proxy and email notification to the SMTP server, and enter the SMTP password on the new destination server if you are using email notification.
- **Code signing certificate**: If you are using an advanced management tool that exposes local update publishing (such as Microsoft System Center Essentials 2007 or Microsoft System Center Configuration Manager 2007), copy the code-signing certificate.

To initialize a trust relationship between the update server and its clients, use the following procedures to point the downstream servers to the new WSUS server, and point the WSUS clients to the new WSUS server.

Point the downstream servers to the new WSUS server

If you have downstream servers in your WSUS configuration, and if the server identity on the destination server was changed, perform the following procedure to point them to the new WSUS server.

To connect a downstream server to an upstream server

1. In the navigation pane of the downstream WSUS Server Administration console, click **Options**.
2. Click **Update Source and Proxy Server**, and then click the **Update Source** tab.
3. Select the **Synchronize from another Windows Server Update Services server** check box, and then type the server name and port number in the corresponding text boxes.
4. Repeat step 1 through step 3 for additional downstream servers, if applicable. The synchronization can take several minutes to several hours to finish.
5. To confirm that the downstream servers are synchronizing with the upstream server, in the WSUS Administration Console on the upstream WSUS server, click **Downstream Servers**. In the **Status** pane, confirm that the server's **Last Synchronization** date is after the date that the previous steps were completed.

Point the WSUS clients to the new WSUS server

If the server identity on the destination server was changed, use the following procedure to point the WSUS clients to the new WSUS destination server.

To point the WSUS clients to the new destination server

1. Open the Local Group Policy Editor, and in **Specify intranet Microsoft update service policy**, change the URL to reflect the new WSUS server.
2. Update the Group Policy settings that are used to point WSUS clients to the WSUS server by entering the FQDN of the new WSUS server. After you have updated the Group Policy settings, WSUS clients will synchronize with the new WSUS server.
3. To force the clients to detect the new destination server, open a command prompt, and run `wuauclt.exe /resetauthorization /detectnow`.

Note

To make sure that WSUS clients point to the new WSUS server immediately, you must force detection, which causes WSUS to update computer group membership. If you do not force a detection, it can take up to four hours for clients to point to the new WSUS server.

4. Depending on the number of clients, the initial synchronization can take several minutes to several hours to finish. To confirm that the synchronization is complete, in the WSUS Administration Console, expand **Computers**, and then click **All Computers**. In the **Status** pane, click **Any**, and then click **Refresh**. Confirm that the computers that you expect to see synchronizing to this WSUS server are listed. The **Last Contact Date** has to be refreshed with a post-migration time stamp.

   Tip

   To force a report that the **Last Contact Date** was updated, run `wuauclt.exe /resetauthorization /detectnow`, and then run `wuauclt.exe /reportnow`.

5. After the clients have synchronized, confirm that clients are installing approved updates based on your WSUS configuration settings. In the WSUS Administration Console, click **Reports**, and then click **Computer Tabular Status**. Select the **Report Options** that are applicable to the clients, and then click **Run Report**.
6. To make sure that no WSUS clients are still pointing at the old WSUS server, wait a week and then open the WSUS Administration Console on the old WSUS server. Expand **Computers**, and then click **All Computers**. In the **Status** pane, click **Any**, and then click **Refresh**. Sort on **Last Status Report**. There should be no clients that have a **Last Status Report** date after the date that the synchronization completed.

3.6. Review additional considerations

After the migration is complete, consider the following:

- It is important to have a backup plan for restoring the WSUS server role if there is a migration failure. You do not need to roll back the migration on the source server because the migration process makes no changes to it. You do not need to roll back the migration on the destination server because it is a new server.
- After you have confirmed that no WSUS clients are contacting the old WSUS server, you can uninstall WSUS from the source server. To perform this operation, see the section titled **Retire the WSUS role on the source server (optional)** in the Windows Server Update Services 3.0 SP2 Migration Guide topic: Post-migration Tasks for WSUS.

See also

- Step 4: Verify the WSUS Migration
- Step 2: Prepare to Migrate WSUS
- WSUS server role description

Step 4: Verify the WSUS Migration

The final step in the migration of your Windows Server Update Services (WSUS) server role is to verify that the migration was performed correctly and if the clients can obtain updates from the new WSUS server.

Task	Description
4.1. Verify the destination server configuration	Verify if the destination server is synchronized.
4.2. Verify client computer functionality	Verify if the clients are correctly obtaining updates from the new WSUS server.

4.1. Verify the destination server configuration

Perform the following procedure on the new WSUS destination server to verify that it is configured properly and functioning correctly before you point the WSUS clients and any downstream servers to the new WSUS server.

1. In Server Manager, click **Tools**, and then click **Windows Server Update Services**.
2. In the WSUS Administration Console, expand **Computers**, and verify that all the Computer Groups that existed on the source server are displayed.
3. Expand **Synchronizations**. In the **Actions** pane, click **Synchronize now**. After the synchronization is complete, (this may take several minutes), confirm that **Succeeded** is displayed in the **Results** column.

If the synchronization fails, click **Options**. Confirm that the **Update Source and Proxy Server** settings and password are correct. Confirm that the firewall access is configured correctly for the new server's environment. Make the necessary changes, and then run the synchronization again.

4.2. Verify client computer functionality

Select a client computer so that you can force a detection to verify that the client and server communication is functioning correctly. Use the ollowing procedure to perform this verification:

1. Open a command prompt and type `wuauclt.exe /detectnow` to force the detection.
2. After the detection is finished, open Windows Explorer and check the **%WinDir%\WindowsUpdate.log** to verify that the forced detection was successful.

See also

- Step 3: Migrate WSUS
- WSUS server role description

Deploy Windows Server Update Services in Your Organization

Windows Server Update Services (WSUS) enables information technology administrators to deploy the latest Microsoft product updates. In Windows Server 2012, WSUS is a server role that can be installed to manage and distribute updates. A WSUS server can be the update source for other WSUS servers within the organization. The WSUS server that acts as an update source is called an upstream server. In a WSUS implementation, at least one WSUS server in the network must connect to Microsoft Update to get available update information. The administrator can determine, based on network security and configuration, how many other servers connect directly to Microsoft Update.

This document provides conceptual information for planning the deployment of Windows Server Update Service in Windows Server 2012. It also provides step-by-step procedures for installing WSUS in the organization.

- Step 1: Prepare for Your WSUS Deployment
- Step 2: Install the WSUS Server Role
- Step 3: Configure WSUS
- Step 4: Approve and Deploy WSUS Updates

Step 1: Prepare for Your WSUS Deployment

The first step in the deployment of Windows Server Update Services (WSUS) is to make important decisions, such as deciding the WSUS deployment scenario, choosing a network topology, and understanding the system requirements. The following checklist describes the steps that are involved in preparing for your deployment.

Task	Description
1.1. Review considerations and system requirements	Review the list of considerations and system requirements to ensure that you have all the necessary hardware and software to deploy WSUS.
1.2. Choose a WSUS deployment scenario	Decide which WSUS deployment scenario will be used.
1.3. Choose a WSUS storage strategy	Decide which WSUS storage strategy best fits your deployment.
1.4. Choose WSUS update languages	Decide which WSUS update languages will be installed.

Task	Description
1.5. Plan WSUS computer groups	Plan the WSUS computer group approach that you will use for your deployment.
1.6. Plan WSUS performance considerations	Plan a WSUS design for optimized performance.
1.7. Plan Automatic Updates settings	Plan how you will configure the automatic updates settings for your scenario.

1.1. Review considerations and system requirements

System Requirements:

Before you enable the WSUS server role, confirm that the server meets the system requirements and confirm that you have the necessary permissions to complete the installation by adhering with the following guidelines:

- Server hardware requirements to enable WSUS role are bound to hardware requirements. The hardware requirements are listed in the topic: Windows Server Update Services Overview.
- If you install roles or software updates that require you to restart the server when installation is complete, restart the server before you enable the WSUS server role.
- Microsoft .NET Framework 4.0 must be installed on the server where the WSUS server role will be installed.
- The NT Authority\Network Service account must have Full Control permissions for the following folders so that the WSUS Administration snap-in displays correctly:
  - %windir%\Microsoft.NET\Framework\v4.0.30319\Temporary ASP.NET Files

    Important
    This path might not exist prior to install Web Server Role that contains Internet Information Services (IIS).
  - %windir%\Temp
- Confirm that the account you plan to use to install WSUS is a member of the Local Administrators group.

Installation Considerations:

During the installation process, WSUS will install the following by default:

- .NET API and Windows PowerShell cmdlets
- Windows Internal Database (WID), which is used by WSUS
- Services used by WSUS, which are:
  - Update Service

- Reporting Web Service
- Client Web Service
- Simple Web Authentication Web Service
- Server Synchronization Service
- DSS Authentication Web Service

Features on Demand Considerations

Be aware that configuring client computers (including servers) to update by using WSUS will result in the following limitations:

1. Server roles that have had their payloads removed using Features on Demand cannot be installed on demand from Microsoft Update. You will must either provide an installation source at the time you try to install such server roles, or configure a source for Features on Demand in Group Policy.
2. Windows client editions will not be able to install .NET 3.5 on demand from the web. The same considerations as server roles apply to .NET 3.5.

Configuring a Features on Demand installation source does not involve WSUS 3.x or WSUS on Windows Server® 2012 R2. For information on how to configure Features on Demand on Windows 8 and later operating systems, please see Configure Features on Demand in Windows Server.

WSUS database requirements

WSUS requires one of the following databases:

- Windows Internal Database (WID)
- Microsoft SQL Server 2012 Standard Edition
- Microsoft SQL Server 2012 Enterprise Edition
- Microsoft SQL Server 2012 Express Edition
- Microsoft SQL Server 2008 R2 SP1 Standard Edition
- Microsoft SQL Server 2008 R2 SP1 Enterprise Edition
- Microsoft SQL Server 2008 R2 SP1 Express Edition

Important
SQL Server Express 2008 R2 has a database size limitation of 10 GB. This database size is likely to be sufficient for WSUS, although there is no appreciable benefit to using this database instead of WID. WID database has a minimum RAM memory requirement of 2 GB, beyond the standard Windows Server system requirements.

You can install the WSUS role on a computer that is separate from the database server computer. In this case, the following additional criteria apply:

1. The database server cannot be configured as a domain controller.
2. The WSUS server cannot run Remote Desktop Services.
3. The database server must be in the same Active Directory domain as the WSUS server, or it must have a trust relationship with the Active Directory domain of the WSUS server.

4. The WSUS server and the database server must be in the same time zone or be synchronized to the same Coordinated Universal Time (Greenwich Mean Time) source.

1.2. Choose a WSUS deployment scenario

This section describes the basic features of all WSUS deployments. Use this section to familiarize yourself with a simple deployment with a single WSUS server, in addition to more complex scenarios, such as a WSUS server hierarchy or a WSUS server on an isolated network segment.

Simple WSUS deployment

The most basic WSUS deployment consists of a server inside the corporate firewall that serves client computers on a private intranet, as shown in Figure 1. The WSUS server connects to Microsoft Update to download updates. This is known as *synchronization*. During synchronization, WSUS determines if any new updates have been made available since the last time you synchronized. If it is your first time synchronizing WSUS, all updates are made available for download.

Note
Initial synchronization can take over an hour. All synchronizations after that should be significantly quicker.

By default, the WSUS server uses port 8530 for HTTP protocol and port 8531 for HTTPS protocol to obtain updates from Microsoft. If there is a corporate firewall between your network and the Internet, you will have to open these ports on the server that communicates directly to Microsoft Update. If you are planning to use custom ports for this communication, you open those ports instead. You can configure multiple WSUS servers to synchronize with a parent WSUS server.

The following figure shows a simple WSUS server scenario in which an administrator can set up a server running WSUS inside the corporate firewall, which synchronizes content directly with Microsoft Update and distributes updates to client computers.

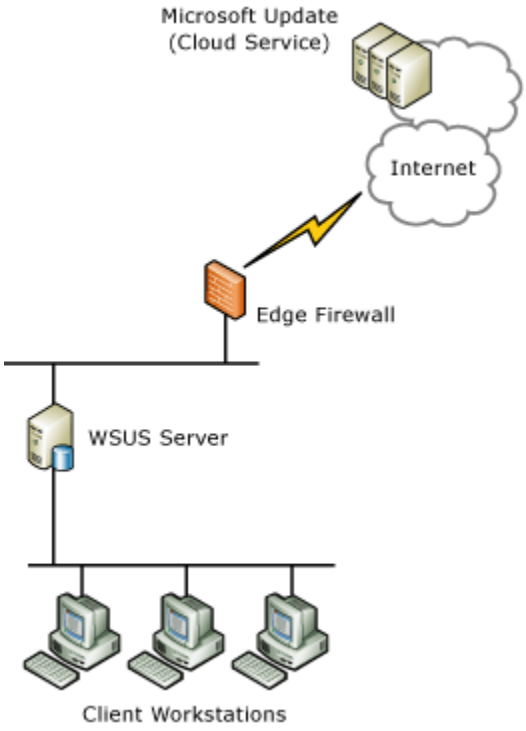

Figure 1 Basic WSUS deployment

Multiple WSUS servers

Administrators can deploy multiple servers running WSUS that synchronize all content within their organization's intranet. In Figure 2, only one server is exposed to the Internet. In this configuration, this is the only server that downloads updates from Microsoft Update. This server is set up as the upstream server—the source to which the downstream servers synchronize. When applicable, servers can be located throughout a geographically dispersed network to provide the best connectivity to all client computers. The following figure shows an example of multiple internal synchronized WSUS servers.

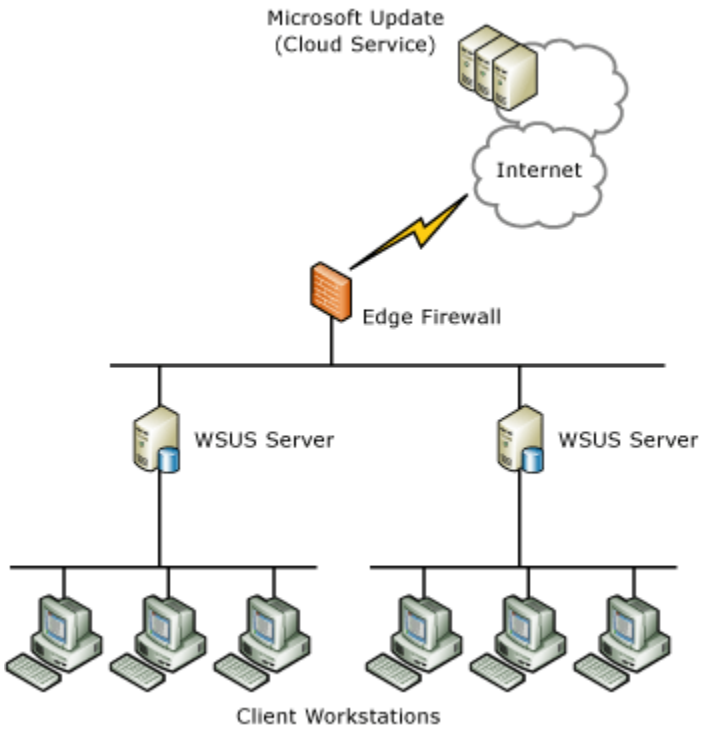

Figure 2 Multiple WSUS servers

Disconnected WSUS server

If corporate policy or other conditions limit computer access to the Internet, administrators can set up an internal server to run WSUS. An example of this is a server that is connected to the intranet but is isolated from the Internet. After downloading, testing, and approving the updates on this server, an administrator would export the update metadata and content to a DVD. The update metadata and content is imported from the DVD to servers running WSUS within the intranet.

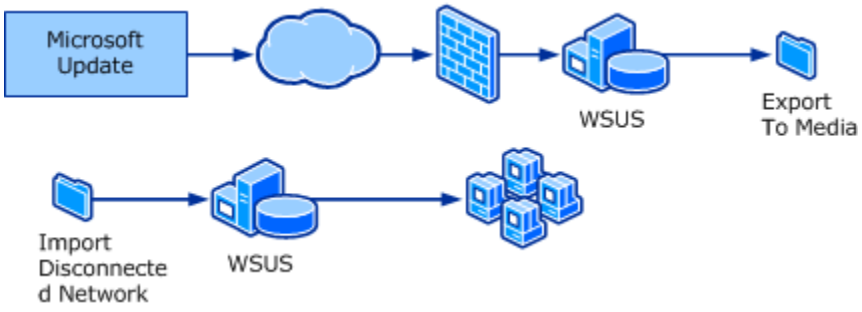

Figure 3 Importing updates with external media

WSUS server hierarchies

You can create complex hierarchies of WSUS servers. Because you can synchronize one WSUS server with another WSUS server instead of with Microsoft Update, you need to have only a single WSUS server that is connected to Microsoft Update. When you link WSUS servers together, there is an upstream WSUS server and a downstream WSUS server. A WSUS server hierarchy deployment offers the following benefits:

- You can download updates one time from the Internet and then distribute the updates to client computers by using downstream servers. This method saves bandwidth on the corporate Internet connection.
- You can download updates to a WSUS server that is physically closer to the client computers, for example, in branch offices.
- You can set up separate WSUS servers to serve client computers that use different languages of Microsoft products.
- You can scale WSUS for a large organization that has more client computers than one WSUS server can effectively manage.

Note
We recommend that you do not create a WSUS server hierarchy that is more than three levels deep. Each level adds time to propagate updates throughout the connected servers. Although there is no theoretical limit to a hierarchy, only deployments with a hierarchy of five levels deep have been tested by Microsoft Corporation.

You can connect WSUS servers in Autonomous mode (to achieve distributed administration) or in Replica mode (to achieve centralized administration). You do not have to deploy a server hierarchy that uses only one mode: you can deploy a WSUS solution that uses both autonomous and replica WSUS servers.

Autonomous mode

The Autonomous mode, also called distributed administration, is the default installation option for WSUS. In Autonomous mode, an upstream WSUS server shares updates with downstream servers during synchronization. Downstream WSUS servers are administered separately, and they do not receive update approval status or computer group information from the upstream server. By using the distributed management model, each WSUS server administrator selects update languages, creates computer groups, assigns computers to groups, tests and approves updates, and makes sure that the correct updates are installed to the appropriate computer groups. The following image shows how you might deploy autonomous WSUS servers in a branch office environment:

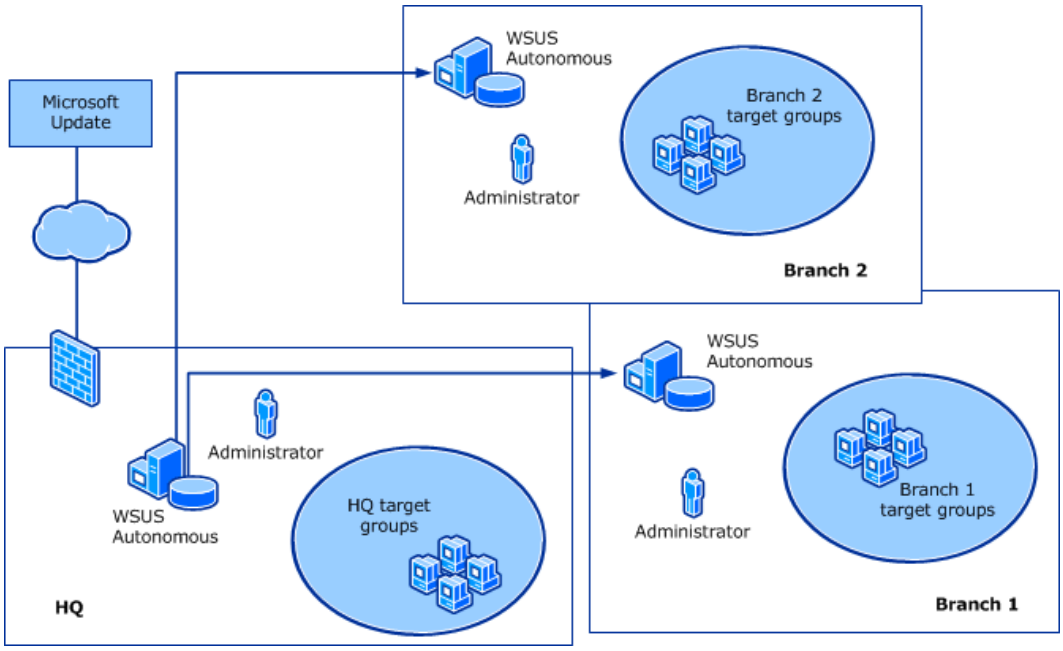

Figure 4 Autonomous servers

Replica mode

The Replica mode, also called centralized administration, works by having an upstream WSUS server that shares updates, approval status, and computer groups with downstream servers. Replica servers inherit update approvals and are not administered separately from the upstream WSUS server. The following image shows how you might deploy replica WSUS servers in a branch office environment.

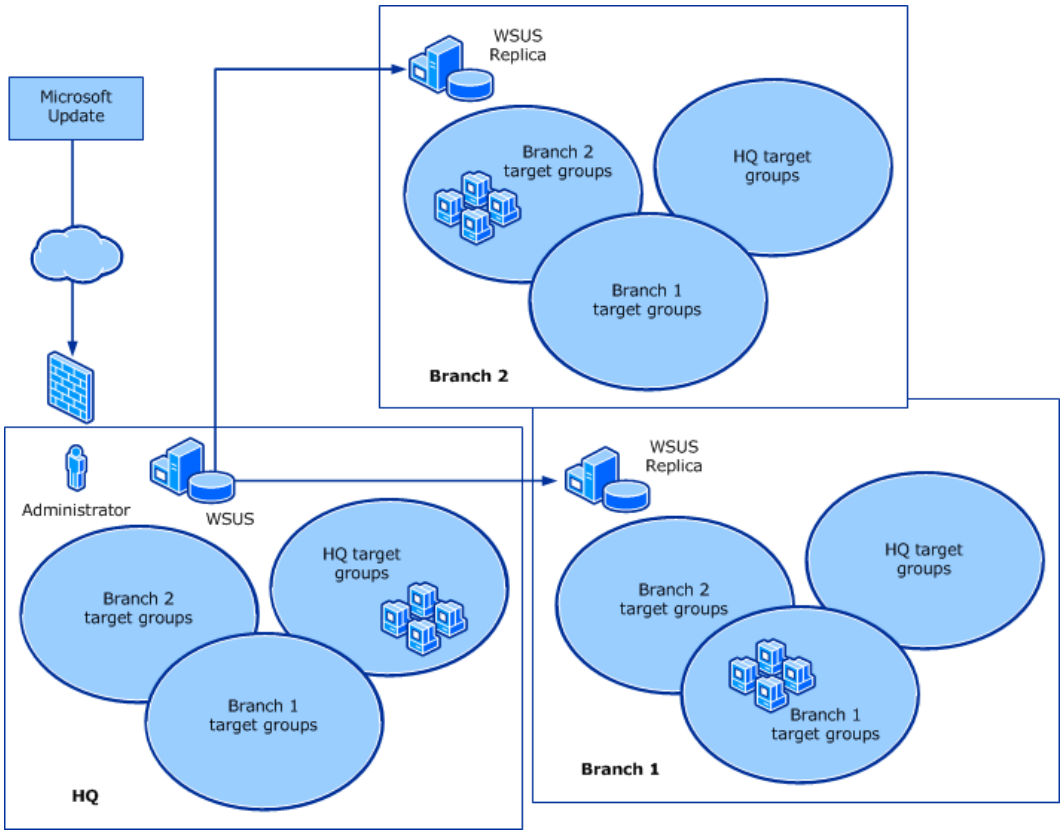

Figure 5 Replica servers

Note

If you set up several replica servers to connect to a single upstream WSUS server, do not schedule synchronization to run at the same time on each replica server. This practice will avoid sudden surges in bandwidth usage.

Branch offices

You can leverage the Branch Office feature in Windows to optimize WSUS deployment. This type of deployment offers the following advantages:

1. Helps reduce WAN link utilization and improves application responsiveness. To enable BranchCache acceleration of content that is served by the WSUS server, install the BranchCache feature on the server and the clients, and ensure that the BranchCache service has started. No other steps are necessary.
2. In branch offices that have low-bandwidth connections to the central office but high-bandwidth connections to the Internet, the Branch Office feature can also be used. In this case you may want to configure downstream WSUS servers to get information about which updates to install from the central WSUS server, but download the updates from Microsoft Update.

Network Load Balancing

Network Load Balancing (NLB) increases the reliability and performance of your WSUS network. You can set up multiple WSUS servers that share a single failover cluster running SQL Server 2008 R2 SP1. In this configuration you must use a full SQL Server installation, not the Windows Internal Database installation that is provided by WSUS, and the database role must be installed on all WSUS front-end servers. You can also have all the WSUS servers use a distributed file system (DFS) to store their content.

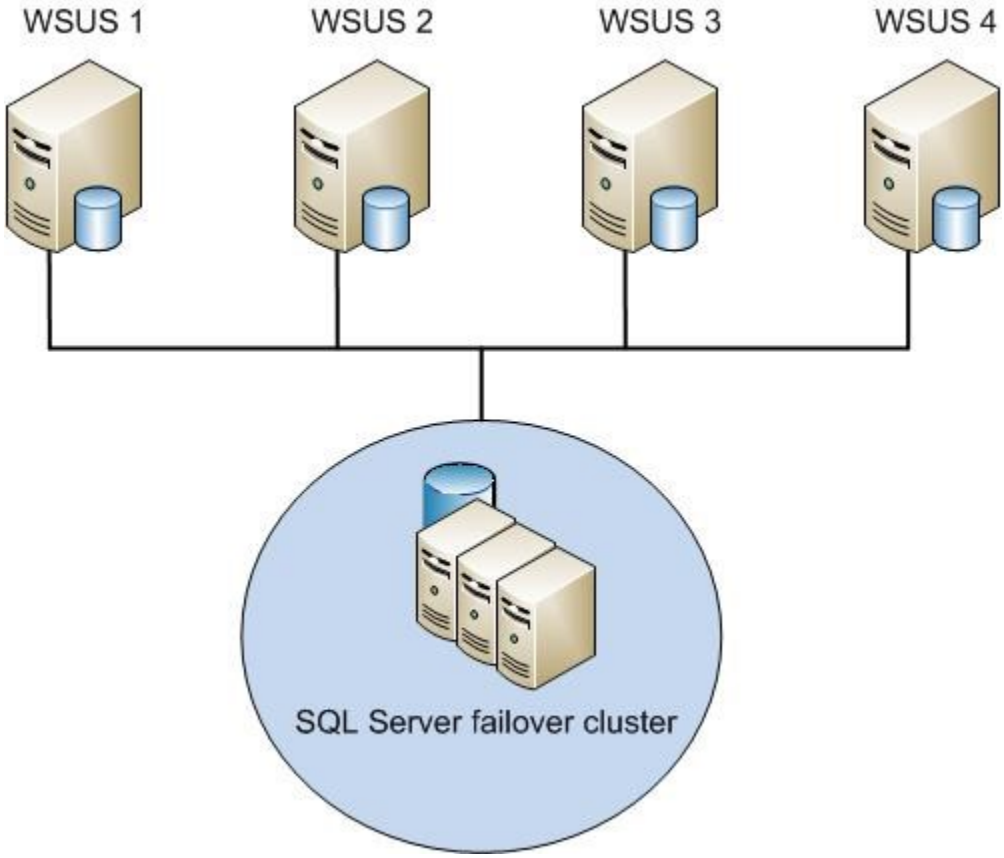

Figure 6 Failover cluster

WSUS deployment with roaming client computers

If the network includes mobile users who log on to the network from different locations, you can configure WSUS to let roaming users update their client computers from the WSUS server that is closest to them geographically. Figure 7 shows one WSUS server deployed in each region, and each region is a DNS subnet. All client computers are directed to the same WSUS server, which resolves in each subnet to the nearest physical WSUS server.

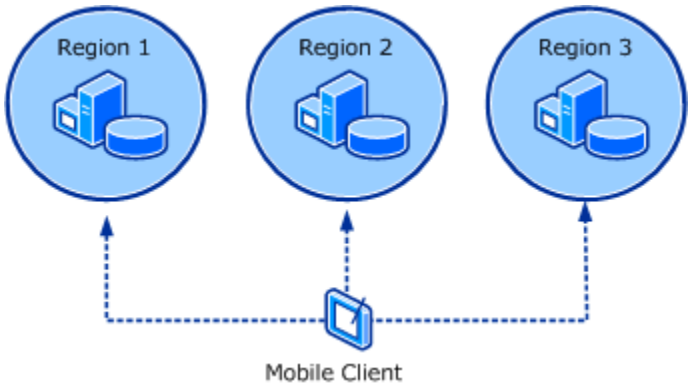

Figure 7 Servers in regions

1.3. Choose a WSUS storage strategy

Windows Server Update Services (WSUS) uses two types of storage systems: a database to store WSUS configuration and update metadata, and an optional local file system to store update files. Before you install WSUS, you should decide how you want to implement storage.

Updates are composed of two parts: metadata that describes the update, and the files that are required to install the update. Update metadata is typically much smaller than the actual update, and it is stored in the WSUS database. Update files are stored on a local WSUS server or on a Microsoft Update Web server.

WSUS database

WSUS requires a database for each WSUS server. WSUS supports the use of a database that resides on a different computer than the WSUS server, with some restrictions. For a list of supported databases and remote database limitations, see the 1.1. Review considerations and system requirements section in this document.

The WSUS database stores the following information:

- WSUS server configuration information
- Metadata that describes each update
- Information about client computers, updates, and interactions

If you install multiple WSUS servers, you must maintain a separate database for each WSUS server, whether it is an autonomous or a replica server. You cannot store multiple WSUS databases on a single instance of SQL Server, except in Network Load Balancing (NLB) clusters that use SQL Server failover.

SQL Server, SQL Server Express, and Windows Internal Database provide the same performance characteristics for a single-server configuration, where the database and the WSUS service are located on the same computer. A single-server configuration can support several thousand WSUS client computers.

Caution

Do not attempt to manage WSUS by accessing the database directly. Directly manipulating the database can cause database corruption. The corruption might not be immediately obvious, but it can prevent upgrades to the next version of the product. You can manage WSUS by using the WSUS console or WSUS application programming interfaces (APIs).

WSUS with Windows Internal Database

By default, the installation wizard creates and uses a Windows Internal Database that is named SUSDB.mdf. This database is located in the %windir%\wid\data\ folder, where %windir% is the local drive on which the WSUS server software is installed. Windows Internal Database is included with Windows Server 2012.

WSUS supports Windows authentication only for the database. You cannot use SQL Server authentication with WSUS. If you use Windows Internal Database for the WSUS database, WSUS Setup creates an instance of SQL Server that is named server\Microsoft##WID, where server is the name of the computer. With either database option, WSUS Setup creates a database named SUSDB. The name of this database is not configurable.

We recommend that you use Windows Internal Database in the following cases:

- The organization has not already purchased and does not require a SQL Server product for any other application.
- The organization does not require an NLB WSUS solution.
- You intend to deploy multiple WSUS servers (for example, in branch offices). In this case, you should consider using Windows Internal Database on the secondary servers, even if you will use SQL Server for the root WSUS server. Because each WSUS server requires a separate instance of SQL Server, you will quickly experience database performance issues if only one instance of SQL Server handles multiple WSUS servers.

Windows Internal Database does not provide a user interface or any database management tools. If you select this database for WSUS, you must use external tools to manage the database. For more information, see:

- Backup and Restore WSUS Data and Backing Up Your Server
- Reindex the WSUS database

WSUS with SQL Server

We recommend that you use SQL Server with WSUS in the following cases:

1. You require an NLB WSUS solution.
2. You already have at least one instance of SQL Server installed.
3. You cannot run the SQL Server service under a local non-system account or by using SQL Server authentication. WSUS supports Windows authentication only.

WSUS update storage

When updates are synchronized to your WSUS server, the metadata and update files are stored in two separate locations. Metadata is stored in the WSUS database. Update files can be stored on your WSUS server or on Microsoft Update servers, depending on how you have configured your synchronization options. If you choose to store update files on your WSUS server, client computers will download approved updates from the local WSUS server. If not, client computers will download approved updates directly from Microsoft Update. The option that makes the most sense for your organization will depend on network bandwidth to the Internet, network bandwidth on the intranet, and local storage availability.

You can select a different update storage solution for each WSUS server that you deploy.

Local WSUS server storage

Local storage of update files is the default option when you install and configure WSUS. This option can save bandwidth on the corporate connection to the Internet because client computers download updates directly from the local WSUS server.

This option requires that the server have sufficient disk space to store all needed updates. At a minimum, WSUS requires 20 GB to store updates locally; however, we recommend 30 GB based on tested variables.

Remote storage on Microsoft Update servers

You can store updates remotely on Microsoft Update servers. This option is useful if most client computers connect to the WSUS server over a slow WAN connection, but they connect to the Internet over a high-bandwidth connection.

In this case, the root WSUS server synchronizes with Microsoft Update and receives the update metadata. After you approve the updates, the client computers download the approved updates from Microsoft Update servers.

Figure 8 shows a WSUS configuration in which updates are stored on Microsoft Update. In this case, branch offices retrieve approved updates directly from Microsoft Update servers. This saves disk space and network bandwidth in the organization. This storage option can offer faster downloads for geographically distributed client computers.

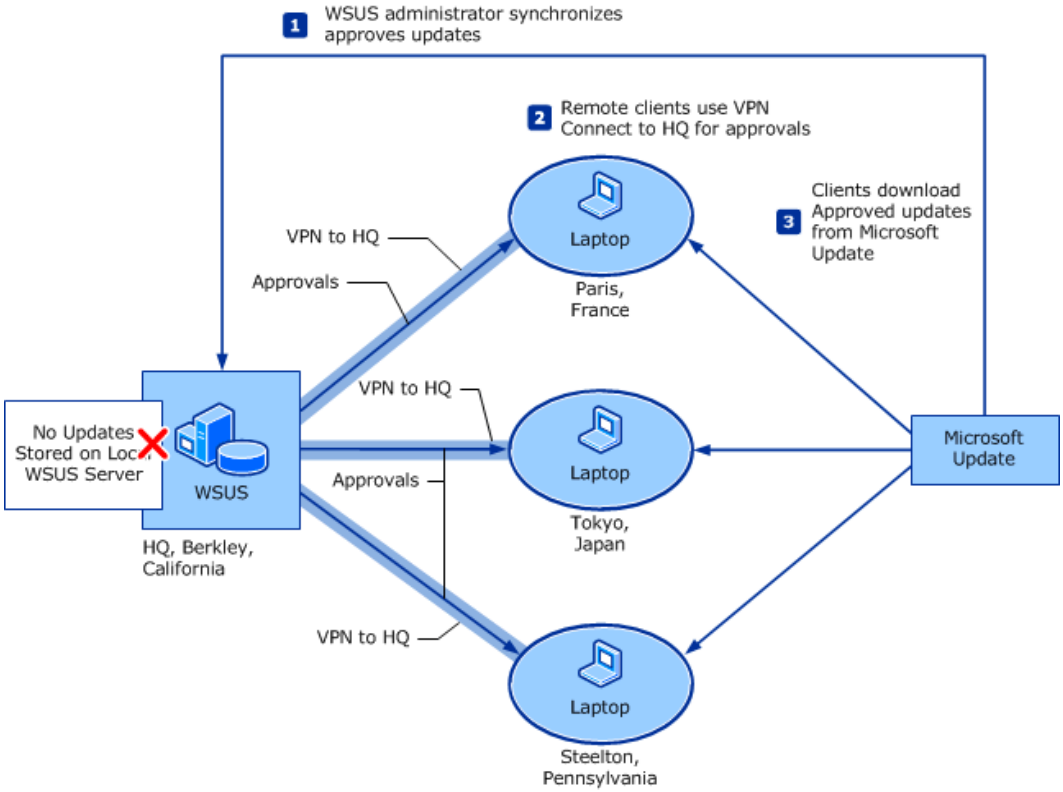

Figure 8 WSUS configuration

1.4. Choose WSUS update languages

When you deploy a WSUS server hierarchy, you should determine which language updates are required throughout the organization. You should configure the root WSUS server to download updates in all languages that are used throughout the entire organization.

For example, the main office might require English and French language updates, but one branch office requires English, French, and German language updates, and another branch office requires English and Spanish language updates. In this situation, you would configure the root WSUS server to download updates in English, French, German, and Spanish. You would then configure the first branch office WSUS server to download updates in English, French, and German only, and configure the second branch office to download updates in English and Spanish only.

The **Choose Languages** page of the WSUS Configuration Wizard allows you to get updates from all languages or from a subset of languages. Selecting a subset of languages saves disk space, but it is important to choose all the languages that are needed by all the downstream servers and client computers of a WSUS server.

Following are some important notes about the update language that you should keep in mind before configure this option:

- Always include English in addition to any other languages that are required throughout your organization. All updates are based on English language packs.
- Downstream servers and client computers will not receive all the updates they need if you have not selected all the necessary languages for the upstream server. Make sure you select all the languages that will be needed by all the client computers that are associated with all the downstream servers.
- You should generally download updates in all languages on the root WSUS server that synchronizes to Microsoft Update. This selection guarantees that all downstream servers and client computers will receive updates in the languages that they require.

If you are storing updates locally, and you have set up a WSUS server to download updates in a limited number of languages, you may notice that there are updates in languages other than the ones you specified. Many update files are bundles of several different languages, which include at least one of the languages specified on the server.

Upstream servers

Important
Configure upstream servers to synchronize updates in all languages that are required by downstream replica servers. You will not be notified of needed updates in the unsynchronized languages.

Updates will appear as **Not Applicable** on client computers that require the language. To avoid this, make sure all operating system languages are included in your WSUS server's synchronization options. You can see all the operating system languages by going to the **Computers** view of the WSUS Administration Console and sorting the computers by operating system language. However, you may want to include more languages if there are Microsoft applications in more than one language (for example, if the French version of Microsoft Word is installed on some computers that use the English version of Windows Server® 2008 R2).

Choosing languages for an upstream server is not the same as choosing languages for a downstream server. The following procedures explain the differences.

To choose update languages for a server synchronizing from Microsoft Update

1. In the WSUS Configuration Wizard:
   - To get updates in all languages, click **Download updates in all languages, including new languages**.
   - To get updates only for specific languages, click **Download updates only in these languages**, and then select the languages for which you want updates.

To choose update languages for a downstream server

1.
   a. If the upstream server has been configured to download update files in a subset of languages: In the WSUS Configuration Wizard, click **Download updates only in these languages (only languages marked with an asterisk are supported by the upstream server)**, and then select the languages for which you want updates.

Note
You should do this even though you want the downstream server to download the same languages as the upstream server.

b. If the upstream server has been configured to download update files in all languages: In the WSUS Configuration Wizard, click **Download updates in all languages supported by the upstream server**.

Notes
You should do this even though you want the downstream server to download the same languages as the upstream server. This setting causes the upstream server to download updates in all languages, including languages that were not originally configured for the upstream server. If you add languages to the upstream server, you should copy the new updates to its replica servers.

Changing language options on the upstream server alone might cause a mismatch between the number of updates that are approved on the central server and the number of updates approved on the replica servers.

1.5. Plan WSUS computer groups

WSUS allows you to target updates to groups of client computers, so you can ensure that specific computers always get the right updates at the most convenient times. For example, if all the computers in one department (such as the Accounting team) have a specific configuration, you can set up a group for that team, decide which updates their computers need and what time they should be installed, and then use WSUS reports to evaluate the updates for the team.

Important
If a WSUS server is running in replica mode, computer groups cannot be created on that server. All the computer groups that are needed for client computers of the replica server must be created on the WSUS server that is the root of the WSUS server hierarchy. For more information about replica mode, see Manage WSUS Replica Servers in the WSUS 3.0 SP2 Operations Guide.

Computers are always assigned to the **All Computers** group, and they remain assigned to the **Unassigned Computers** group until you assign them to another group. Computers can belong to more than one group.

Computer groups can be set up in hierarchies (for example, the Payroll group and the Accounts Payable group below the Accounting group). Updates that are approved for a higher group will automatically be deployed to lower groups, in addition to the higher group. In this example, if you approve Update1 for the Accounting group, the update will be deployed to all the computers in the Accounting group, all the computers in the Payroll group, and all the computers in the Accounts Payable group.

Because computers can be assigned to multiple groups, it is possible for a single update to be approved more than once for the same computer. However, the update will be deployed only once, and any conflicts will be resolved by the WSUS server. To continue with the previous example, if ComputerA is assigned to the Payroll group and the Accounts Payable group, and Update1 is approved for both groups, it will be deployed only once.

You can assign computers to computer groups by using one of two methods, server-side targeting or client-side targeting. Following are the definitions for each method:

- **Server-side targeting**: You manually assign one or more client computers to multiple groups simultaneously.
- **Client-side targeting**: You use Group Policy or edit the registry settings on client computers to enable those computers to automatically add themselves into the previously created computer groups.

Conflict Resolution

The server applies the following rules to resolve conflicts and determine the resultant action on clients:

1. Priority
2. Priority of Install and Uninstall
3. Priority of Deadlines

Priority

The actions associated with the group of the highest priority override the actions of other groups. The deeper a group appears within the hierarchy of groups, the higher its priority. Priority is assigned only based on depth; all branches have equal priority. For example, a group two levels beneath the Desktops branch has a higher priority than a group one level beneath the Server branch.

In the following text example of the Update Services console hierarchy pane, for a WSUS server named WSUS-01, computer groups named Desktop Computers and Server have been added to the default **All Computers** group. Both the Desktop Computers and Server groups are at the same hierarchical level.

Update Services

WSUS-01

Updates

Computers

All Computers

Unassigned Computers

Desktop Computers

Desktops-L1

Desktops-L2

Servers

Servers-L1

Downstream Servers

Synchronizations

Reports

Options

In this example, the group two levels beneath the Desktop Computers branch (Desktops L2) has a higher priority than the group one level beneath the Server branch (Servers L1). Accordingly, for a computer that has membership in both the Desktops-L2 and the Servers-L1 groups, all actions for the Desktops-L2 group take priority over actions specified for the Servers-L1 group.

Priority of Install and Uninstall

Install actions override uninstall actions. Required installs override optional installs (optional installs are only available through the API and changing an approval for an update using the WSUS Administration Console will clear all optional approval.)

Priority of Deadlines

Actions that have a deadline override those with no deadline. Actions with earlier deadlines override those with later deadlines.

1.6. Plan WSUS performance considerations

There are some areas that you should carefully plan before deploy WSUS so that you can have optimized performance. The key areas are:

- Network setup
- Deferred download
- Filters
- Installation
- Large update deployment
- Background Intelligent Transfer Service

Network setup

To optimize performance in WSUS networks, consider the following suggestions:

1. Set up WSUS networks in a hub-and-spoke topology rather than in a hierarchical topology.
2. Use DNS netmask ordering for roaming client computers, and configure roaming client computers to obtain updates from the local WSUS server.

Deferred download

You can approve updates, and download the update metadata before you download the update files, this method is called *deferred downloads*. When you defer downloads, an update is downloaded only after it is approved. We recommend that you defer downloads because it optimizes network bandwidth and disk space.

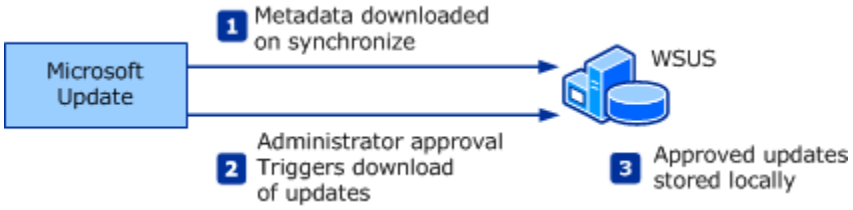

Figure 9 Deferred download

In a hierarchy of WSUS servers, WSUS automatically sets all downstream servers to use the deferred download setting of the root WSUS server. You can change this default setting. For example, you can configure an upstream server to perform full, immediate synchronizations, and then configure a downstream server to defer the downloads.

If you deploy a hierarchy of connected WSUS servers, we recommend that you do not deeply nest the servers. If you enable deferred downloads and a downstream server requests an update that is not approved on the upstream server, the downstream server's request forces a download on the upstream server. The downstream server then downloads the update on a subsequent synchronization. In a deep hierarchy of WSUS servers, delays can occur as updates are requested, downloaded, and then passed through the server hierarchy. By default, deferred downloads are enabled when you store updates locally. You can change this option manually.

Filters

WSUS lets you filter update synchronizations by language, product, and classification. In a hierarchy of WSUS servers, WSUS automatically sets all downstream servers to use the update filtering options that are selected on the root WSUS server. You can reconfigure download servers to receive only a subset of the languages.

By default, the products to be updated are Windows and Office, and the default classifications are Critical updates, Security updates, and Definition updates. To conserve bandwidth and disk space, we recommend that you limit languages to those that you actually use.

Installation

Updates typically consist of new versions of files that already exist on the computer that is being updated. On a binary level, these existing files might not differ very much from updated versions. The express installation files feature identifies the exact bytes between versions, creates and distributes updates of only those differences, and then merges the existing file together with the updated bytes.

Sometimes this feature is called *delta delivery* because it downloads only the delta (difference) between two versions of a file. Express installation files are larger than the updates that are distributed to client computers because the express installation file contains all possible versions of each file that is to be updated.

You can use express installation files to limit the bandwidth that is consumed on the local network, because WSUS transmits only the delta applicable to a particular version of an updated component. However, this comes at the cost of additional bandwidth between your WSUS server, any upstream WSUS servers, and Microsoft Update, and requires additional local disk space. By default, WSUS does not use express installation files.

Figure 10 shows an update that is first distributed by using express installation files, and then distributed without using express installation files. In this example, the use of express installation files incurs an initial download size that is three times the size of the actual update. However, this cost is lessened by the reduced amount of bandwidth that is required to update client computers on the corporate network. When the express installation files feature is disabled, the initial download is smaller, but the full size of the download must then be distributed to each client computer on the corporate network.

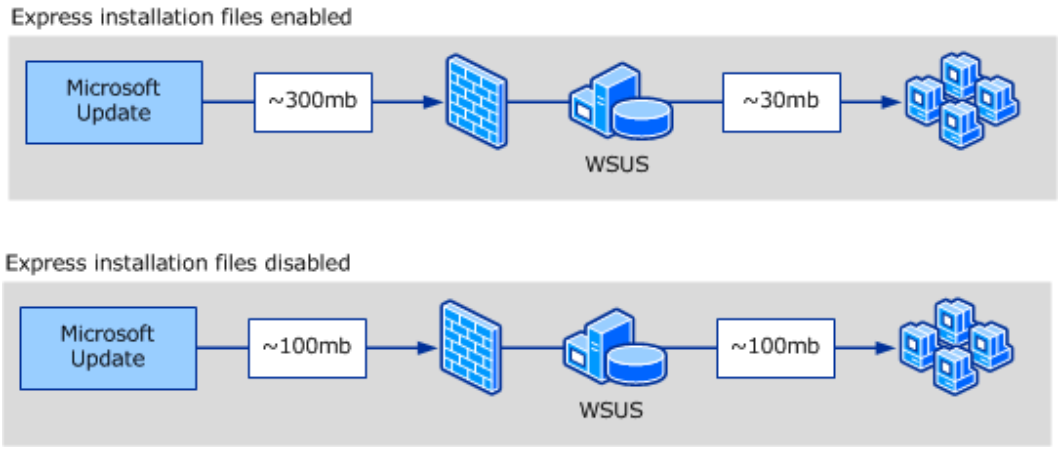

Figure 10 Express installation files

Note

The file sizes in this image are provided for illustration only. Each update and express installation file size varies. The size of each update that is distributed to client computers depends on the state of the computer that is being updated.

Not all updates are good candidates for distribution by using express installation files. If you select this option, you obtain express installation files for all updates. If you do not store updates locally, the Windows Update Agent will decide whether to download the express installation files or the full-file update distributions.

Large update deployment

When you deploy large updates (such as service packs), you can avoid saturating the network by using the following practices:

1. Use Background Intelligent Transfer Service (BITS) throttling. BITS bandwidth limitations can be controlled by time-of-day, but they apply to all applications that are using BITS. To learn how to control BITS throttling, please see Group Policies

2. Use Internet Information Services (IIS) throttling to limit throttling to one or more web services.
3. Use computer groups to control the rollout. A client computer identifies itself as a member of a particular computer group when it sends information to the WSUS server. The WSUS server uses this information to determine which updates should be deployed to this computer. You can set up multiple computer groups and sequentially approve large service pack downloads for a subset of these groups.

Background Intelligent Transfer Service

WSUS uses the Background Intelligent Transfer Service (BITS) protocol for all its file transfer tasks. This includes downloads to client computers and server synchronizations. BITS enables programs to download files by using spare bandwidth. BITS maintains file transfers through network disconnections and computer restarts. For more information, see: Background Intelligent Transfer Service.

1.7. Plan Automatic Updates settings

You can specify a deadline to approve updates on the WSUS server. The deadline causes client computers to install the update at a specific time, but there are a number of different situations, depending on whether the deadline has expired, whether there are other updates in the queue for the computer to install, and whether the update (or another update in the queue) requires a restart.

By default, Automatic Updates polls the WSUS server for approved updates every 22 hours minus a random offset. If new updates need to be installed, they are downloaded. The time between each detection cycle can be manipulated from 1 to 22 hours.

You can manipulate the notification options as follows:

1. If Automatic Updates is configured to notify the user of updates that are ready to be installed, the notification is sent to the System log and to the notification area of the client computer.
2. When a user with appropriate credentials clicks the notification area icon, Automatic Updates displays the available updates to install. The user must click **Install** to start the installation. A message appears if the update requires the computer to be restarted to complete the update. If a restart is requested, Automatic Updates cannot detect additional updates until the computer is restarted.

If Automatic Updates is configured to install updates on a set schedule, applicable updates are downloaded and marked as ready to install. Automatic Updates notifies users who have

appropriate credentials by using a notification area icon, and an event is logged in the System log.

At the scheduled day and time, Automatic Updates installs the update and restarts the computer (if necessary), even if no local administrator is logged on. If a local administrator is logged on and the computer requires a restart, Automatic Updates displays a warning and a countdown for the restart. Otherwise, the installation occurs in the background.

If the computer must be restarted, and any user is logged on, a similar countdown dialog box is displayed, which warns the user about the impending restart. You can manipulate computer restarts with Group Policy.

After the new updates are downloaded, Automatic Updates polls the WSUS server for the list of approved packages to confirm that the packages it downloaded are still valid and approved. This means that, if a WSUS administrator removes updates from the list of approved updates while Automatic Updates is downloading updates, only the updates that are still approved are actually installed.

See also

- Step 2: Install the WSUS Server Role
- Deploy Windows Server Update Services in Your Organization
- Windows Server Update Services Overview

Step 2: Install the WSUS Server Role

The next step in the deployment of your WSUS server is to install the WSUS server role. The following procedure describes how to install the WSUS server role by using Server Manager.

Important
This installation procedure only covers how to install WSUS using Windows Internal Database (WID). The procedures to install WSUS using Microsoft SQL Server are documented in this article.

To install the WSUS server role

1. Log on to the server on which you plan to install the WSUS server role by using an account that is a member of the Local Administrators group.
2. In **Server Manager**, click **Manage**, and then click **Add Roles and Features**.
3. On the **Before you begin** page, click **Next**.
4. In the **Select installation type** page, confirm that **Role-based or feature-based installation** option is selected and click **Next**.
5. On the **Select destination server** page, choose where the server is located (from a server pool or from a virtual hard disk). After you select the location, choose the server that you want to install the WSUS server role, and then click **Next**.

6. On the **Select server roles** page, select **Windows Server Update Services**. **Add features that are required for Windows Server Update Services** opens. Click **Add Features**, and then click **Next**.
7. On the **Select features** page. Retain the default selections, and then click **Next**.

   Important
   WSUS only requires the default Web Server role configuration. If you are prompted for additional Web Server role configuration while setting up WSUS you can safely accept the default values and continue setting up WSUS.
8. On the **Windows Server Update Services** page, click **Next**.
9. On the **Select Role Services** page, leave the default selections, and then click **Next**.

   Tip
   You must select one Database type. If the database options are all cleared (not selected), post installation tasks will fail.
10. On the **Content location selection** page, type a valid location to store the updates. For example, you can create a folder named WSUS_database at the root of drive K specifically for this purpose, and type **k:\WSUS_database** as the valid location.
11. Click **Next**. The **Web Server Role (IIS)** page opens. Review the information, and then click **Next**. In **Select the role services to install for Web Server (IIS)**, retain the defaults, and then click **Next**.
12. On the **Confirm installation selections** page, review the selected options, and then click **Install**. The WSUS installation wizard runs. This might take several minutes to complete.
13. Once WSUS installation is complete, in the summary window on the **Installation progress** page, click **Launch Post-Installation tasks**. The text changes, requesting: **Please wait while your server is configured**. When the task has finished, the text changes to: **Configuration successfully completed**. Click **Close**.
14. In **Server Manager**, verify if a notification appears to inform you that a restart is required. This can vary according to the installed server role. If it requires a restart make sure to restart the server to complete the installation.

Important

At this point the installation process is finished, however for WSUS to be functional you need to proceed to Step 3: Configure WSUS.

See also

- Step 3: Configure WSUS
- Step 1: Prepare for Your WSUS Deployment
- Deploy Windows Server Update Services in Your Organization
- Role description

Step 3: Configure WSUS

After installing the WSUS server role on your server, you need to properly configure it. The following checklist describes the steps involved in performing the initial configuration for your WSUS server.

Task	Description
3.1. Configure network connections	Configure the cluster network by using the Network Configuration Wizard.
3.2. Configure WSUS by using the WSUS Configuration Wizard	Use the WSUS Configuration wizard to perform the base WSUS configuration.
3.3. Configure computer groups	Create computer groups in the WSUS administration console to manage updates in your organization.
3.4. Configure client updates	Specify how and when automatic updates are applied to client computers.
3.5. Secure WSUS with the Secure Sockets Layer Protocol	Configure Secure Sockets Layer (SSL) protocol to help protect Windows Server Update Services (WSUS).

3.1. Configure network connections

Before you start the configuration process, be sure that you know the answers to the following questions:

1. Is the server's firewall configured to allow clients to access the server?
2. Can this computer connect to the upstream server (such as the server that is designated to download updates from Microsoft Update)?
3. Do you have the name of the proxy server and the user credentials for the proxy server, if you need them?

By default, WSUS is configured to use Microsoft Update as the location from which to obtain updates. If you have a proxy server on the network, you can configure WSUS to use the proxy server. If there is a corporate firewall between WSUS and the Internet, you might have to configure the firewall to ensure that WSUS can obtain updates.

Tip
Although Internet connectivity is required to download updates from Microsoft Update, WSUS offers you the ability to import updates onto networks that are not connected to the Internet.

When you have the answers for these questions, you can start configuring the following WSUS network settings:

- **Updates** Specify the way this server will obtain updates (from Microsoft Update or from another WSUS server).
- **Proxy** If you identified that WSUS needs to use a proxy server to have Internet access, you need to configure proxy settings in the WSUS server.
- **Firewall** If you identified that WSUS is behind a corporate firewall, there are some additional steps that must be done at the edge device to properly allow WSUS traffic.

3.1.1. Connection from the WSUS server to the Internet

If there is a corporate firewall between WSUS and the Internet, you might have to configure that firewall to ensure WSUS can obtain updates. To obtain updates from Microsoft Update, the WSUS server uses port 443 for HTTPS protocol. Although most of corporate firewalls allow this type of traffic, there are some companies that restrict Internet access from the servers due the company's security policies. If your company restricts access, you need to obtain authorization to allow Internet access from WSUS to the following list of URLs:

- http://windowsupdate.microsoft.com
- http://*.windowsupdate.microsoft.com
- https://*.windowsupdate.microsoft.com
- http://*.update.microsoft.com
- https://*.update.microsoft.com
- http://*.windowsupdate.com
- http://download.windowsupdate.com
- http://download.microsoft.com
- http://*.download.windowsupdate.com
- http://wustat.windows.com
- http://ntservicepack.microsoft.com

Important
For a scenario about where WSUS is failing to obtain updates due firewall configurations, see article 885819 in the Microsoft Knowledge Base.

The following section describes how to configure a corporate firewall that is positioned between WSUS and the Internet. Because WSUS initiates all the network traffic, you it is not necessary to configure Windows Firewall on the WSUS server. Although the connection between Microsoft Update and WSUS requires ports 80 and 443 to be open, you can configure multiple WSUS servers to synchronize with a custom port.

3.1.2. Connection between WSUS servers

WSUS upstream and downstream servers will synchronize on the port configured by the WSUS Administrator. By default, these ports are configured as follows:

- On WSUS 3.2 and earlier, port 80 for HTTP and 443 for HTTPS
- On WSUS 6.2 and later (at least Windows Server 2012), port 8530 for HTTP and 8531 for HTTPS

The firewall on the WSUS server must be configured to allow inbound traffic on these ports.

3.1.3. Connection between clients (Windows Update Agent) and WSUS servers

The listening interfaces and ports are configured in the IIS site(s) for WSUS and in any Group Policy settings used to configure client PCs. The default ports are the same as those specified in the preceding section **Connection between WSUS servers**, and the firewall on the WSUS server must also be configured to allow inbound traffic on these ports.

Configure the proxy server

If the corporate network uses proxy servers, the proxy servers must support HTTP and SSL protocols and use basic authentication or Windows authentication. These requirements can be met by using one of the following configurations:

1. A single proxy server that supports two protocol channels. In this case, set one channel to use HTTP and the other channel to use HTTPS.

   Note
   You can set up one proxy server that handles both protocols for WSUS during the WSUS server software installation.

2. Two proxy servers, each of which supports a single protocol. In this case, one proxy server is configured to use HTTP, and the other proxy server is configured to use HTTPS.

To set up two proxy servers, each of which will handle one protocol for WSUS, use the following procedure:

To set up WSUS to use two proxy servers

1. Log on to the computer that is to be the WSUS server by using an account that is a member of the local Administrators group.
2. Install the WSUS server role. During the WSUS Configuration Wizard (discussed in the next section) do not specify a proxy server.
3. Open a command prompt (Cmd.exe) as an administrator. To open a command prompt as an administrator, go to **Start**. In **Start Search**, type **Command Prompt**. At the top of the Start menu, right-click **Command Prompt**, and then click **Run as administrator**. If the **User Account Control** dialog box appears, enter the appropriate credentials (if requested), confirm that the action it displays is what you want, and then click **Continue**.
4. In the Command Prompt window, go to the C:\Program Files\Update Services\Tools folder. Type the following command:

   wsusutil ConfigureSSLProxy [< proxy_server proxy_port>] –enable, where:

a. proxy_server is the name of the proxy server that supports HTTPS.
b. proxy_port is the proxy server port number.

5. Close the Command Prompt window.

To add the proxy server that uses the HTTP protocol to the WSUS configuration, use the following procedure:

To add a proxy server that uses the HTTP protocol

1. Open the WSUS Administration Console.
2. In the left pane, expand the server name, and then click **Options**.
3. In the **Options** pane, click **Update Source and Update Server**, and then click the **Proxy Server** tab.
4. Use the following options to modify the existing proxy server configuration:

To change or add a proxy server to the WSUS configuration

a. Select the check box for **Use a proxy server when synchronizing**.
b. In the **Proxy server name** text box, type the name of the proxy server.
c. In the **Proxy port number** text box, type the port number of the proxy server. The default port number is 80.
d. If the proxy server requires that you use a specific user account, select the **Use user credentials to connect to the proxy server** check box. Type the required user name, domain, and password into the corresponding text boxes.
e. If the proxy server supports basic authentication, select the **Allow basic authentication (password is sent in cleartext)** check box.
f. Click **OK**.

To remove a proxy server from the WSUS configuration

a. To remove a proxy server from the WSUS configuration, clear the check box for **Use a proxy server when synchronizing**.
b. Click **OK**.

3.2. Configure WSUS by using the WSUS Configuration Wizard

This procedure assumes that you are using the WSUS Configuration Wizard, which appears the first time you launch the WSUS Management Console. Later in this topic, you will learn how to perform these configurations by using the **Options** page:

To configure WSUS

1. In the Server Manager navigation pane, click **Dashboard**, click **Tools**, and then click **Windows Server Update Services**.

   Note
   If the **Complete WSUS Installation** dialog box appears, click **Run**. In the **Complete WSUS Installation** dialog box, click **Close** when the installation successfully finishes.

2. The Windows Server Update Services Wizard opens. On the **Before you Begin** page, review the information, and then click **Next**.
3. Read the instructions on the **Join the Microsoft Update Improvement Program** page and evaluate if you want to participate. If you want to participate in the program. Retain the default selection, or clear the checkbox, and then click **Next**.
4. On the **Choose Upstream Server** page, there are two options:
   a. Synchronize the updates with Microsoft Update
   b. Synchronize from another Windows Server Update Services server
      - If you choose to synchronize from another WSUS server, specify the server name and the port on which this server will communicate with the upstream server.
      - To use SSL, select the **Use SSL when synchronizing update information** check box. The servers will use port 443 for synchronization. (Make sure that this server and the upstream server support SSL).
      - If this is a replica server, select the **This is a replica of the upstream server** check box.
5. After selecting the proper options for your deployment, click **Next** to proceed.
6. On the **Specify Proxy Server** page, select the **Use a proxy server when synchronizing** check box, and then type the proxy server name and port number (port 80 by default) in the corresponding boxes.

   Important
   You must complete this step if you identified that WSUS needs a proxy server to have Internet access.

7. If you want to connect to the proxy server by using specific user credentials, select the **Use user credentials to connect to the proxy server** check box, and then type the user name, domain, and password of the user in the corresponding boxes. If you want to enable basic authentication for the user who is connecting to the proxy server, select the **Allow basic authentication (password is sent in cleartext)** check box.
8. Click **Next**. On the **Connect to Upstream Server** page, click **Start Connecting**.
9. When it connects, click **Next** to proceed.
10. On the **Choose Languages** page, you have the option to select the languages from which WSUS will receive updates—all languages or a subset of languages. Selecting a subset of languages will save disk space, but it is important to choose all of the

languages that are needed by all the clients of this WSUS server. If you choose to get updates only for specific languages, select **Download updates only in these languages**, and then select the languages for which you want updates; otherwise, leave the default selection.

> **Warning**
> If you select the option **Download updates only in these languages**, and this server has a downstream WSUS server connected to it, this option will force the downstream server to also use only the selected languages.

11. After selecting the appropriate language options for your deployment, click **Next** to continue.
12. The **Choose Products** page allows you specify the products for which you want updates. Select product categories, such as Windows, or specific products, such as Windows Server 2008. Selecting a product category selects all the products in that category.
13. Select the appropriate product options for your deployment, and then click **Next**.
14. On the **Choose Classifications** page, select the update classifications that you want to obtain. Choose all the classifications or a subset of them, and then click **Next**.
15. The **Set Sync Schedule** page enables you to select whether to perform synchronization manually or automatically.
    - If you choose **Synchronize manually**, you must start the synchronization process from the WSUS Administration Console.
    - If you choose **Synchronize automatically**, the WSUS server will synchronize at set intervals.

    Set the time for the **First synchronization**, and then specify the number of **Synchronizations per day** that you want this server to perform. For example, if you specify that there should be four synchronizations per day, starting at 3:00 A.M., synchronizations will occur at 3:00 A.M., 9:00 A.M., 3:00 P.M., and 9:00 P.M.
16. After selecting the appropriate synchronization options for your deployment, click **Next** to continue.
17. On the **Finished** page, you have the option to start the synchronization now by selecting the **Begin initial synchronization** check box. If you do not select this option, you need to use WSUS Management Console to perform the initial synchronization. Click **Next** if you want to read more about additional settings, or you can click **Finish** to conclude this wizard and finish the initial WSUS setup.
18. After you click **Finish**, the WSUS Management Console appears.

Now that you have performed the basic WSUS configuration, read the next sections for more details about changing the settings by using WSUS Management Console.

3.3. Configure computer groups

Computer groups are an important part of Windows Server Update Services (WSUS) deployments. Computer groups permit you to test and target updates to specific computers. There are two default computer groups: All Computers and Unassigned Computers. By default,

when each client computer first contacts the WSUS server, the server adds that client computer to both of these groups.

You can create as many custom computer groups as you need to manage updates in your organization. As a best practice, create at least one computer group to test updates before you deploy them to other computers in your organization.

Use the following procedure to create a new group and assign a computer to this group:

To create a computer group

1. In the WSUS Administration Console, under **Update Services**, expand the WSUS server, expand **Computers**, right-click **All Computers**, and then click **Add Computer Group**.
2. In the **Add Computer Group** dialog box, in **Name**, specify the name of the new group, and click then **Add**.
3. Click **Computers**, and then select the computers that you want to assign to this new group.
4. Right-click the computer names that you selected in the previous step, and then click **Change Membership**.
5. In the **Set Computer Group Membership** dialog box, select the test group that you created, and then click **OK**.

3.4. Configure client updates

WSUS Setup automatically configures IIS to distribute the latest version of Automatic Updates to each client computer that contacts the WSUS server. The best way to configure Automatic Updates depends on the network environment.

- In an environment that uses Active Directory directory service, you can use an existing domain–based Group Policy Object (GPO) or create a new GPO.
- In an environment without Active Directory, use the Local Group Policy Editor to configure Automatic Updates, and then point the client computers to the WSUS server.

Important
The following procedures assume that your network runs Active Directory. These procedures also assume that you are familiar with Group Policy and you use it to manage the network.

Use the following procedures to configure Automatic Updates for client computers:

- Configure Automatic Updates in Group Policy
- **Configure client computer groups to use the WSUS server**

Perform the first two procedures on the domain–based GPO of your choice, and perform the last procedure at a command prompt on the client computer.

Configure Automatic Updates in Group Policy

If you have set up Active Directory in your network, you can configure one or multiple computers simultaneously by including them in a Group Policy Object (GPO), and then configuring that GPO with WSUS settings. We recommend that you create a new GPO that contains only WSUS settings.

Link this WSUS GPO to an Active Directory container that is appropriate for your environment. In a simple environment, you might link a single WSUS GPO to the domain. In a more complex environment, you might link multiple WSUS GPOs to several organizational units (OUs), which will enable you to apply different WSUS policy settings to different types of computers.

To enable WSUS through a domain GPO

1. In the Group Policy Management Console (GPMC), browse to the GPO on which you want to configure WSUS, and then click **Edit**.
2. In the GPMC, expand **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Windows Components**, and then click **Windows Update**.
3. In the details pane, double-click **Configure Automatic Updates**. The **Configure Automatic Updates** policy opens.
4. Click **Enabled**, and then select one of the following options under the **Configure automatic updating** setting:
   - **Notify for download and notify for install**. This option notifies a logged-on administrative user before you download and install the updates.
   - **Auto download and notify for install**. This option automatically begins downloading updates and then notifies a logged-on administrative user before installing the updates. By default, this option is selected.
   - **Auto download and schedule the install**. This option automatically begins downloading updates and then installs the updates on the day and time that you specify.
   - **Allow local admin to choose setting**. This option lets local administrators to use Automatic Updates in Control Panel to select a configuration option. For example, they can choose a scheduled installation time. Local administrators cannot disable Automatic Updates.
5. Select **Enable client-side targeting**, select **Enabled**, and then type the name of the WSUS computer group to which you want to add this computer in the **Target group name for this computer** box.

   Note
   Enable client-side targeting enables client computers to add themselves to target computer groups on the WSUS server, when Automatic Updates is redirected to a WSUS server. If the status is set to Enabled, this computer will identify itself as a member of a particular computer group when it sends information to the WSUS server, which uses it to determine which updates are

deployed to this computer. This setting indicates to the WSUS server which group the client computer will use. You must create the group on the WSUS server, and add domain-member computers to that group.

6. Click **OK** to close the **Enable client-side targeting** policy and return to the Windows Update details pane.
7. Click **OK** to close the **Configure Automatic Updates** policy and return to the Windows Update details pane.
8. In the **Windows Update** details pane, double-click **Specify intranet Microsoft update service location**.
9. Click **Enabled**, and then, server in the **Set the intranet update service for detecting updates** and **Set the intranet statistics server** text boxes, type the same URL of the WSUS server. For example, type *http://servername* in both boxes (where *servername* is the name of the WSUS server).

   Warning
   When you type the intranet address of your WSUS server make sure to specify which port is going to be used. By default WSUS will use port 8530 for HTTP and 8531 for HTTPS. For example, if you are using HTTP, you should type **http://servername:8530**.
10. Click **OK**.

After you set up a client computer, it will take several minutes before the computer appears on the **Computers** page in the WSUS Administration Console. For client computers that are configured with a domain-based Group Policy Object, it can take about 20 minutes for Group Policy to apply the new policy settings to the client computer. By default, Group Policy updates in the background every 90 minutes, with a random offset of 0–30 minutes. If you want to update Group Policy sooner, you can open a Command Prompt window on the client computer and type **gpupdate /force**.

For client computers that are configured by using the Local Group Policy Editor, the GPO is applied immediately, and the update takes about 20 minutes. If you begin detection manually, you do not have to wait 20 minutes for the client computer to contact WSUS.

Because waiting for detection to start can be a time-consuming process, you can use the following procedure to initiate detection immediately.

To initiate WSUS detection

1. On the client computer, open a Command Prompt window with elevated privileges.
2. Type **wuauclt.exe /detectnow**, and then press ENTER.

3.5. Secure WSUS with the Secure Sockets Layer Protocol

You can use the Secure Sockets Layer (SSL) protocol to help secure the WSUS deployment. WSUS uses SSL to authenticate client computers and downstream WSUS servers to the WSUS server. WSUS also uses SSL to encrypt update metadata.

Important
Clients and downstream servers that are configured to use Transport Layer Security (TLS) or HTTPS must also be configured to use a fully qualified domain name (FQDN) for their upstream WSUS server.

WSUS uses SSL for metadata only, not for update files. This is the same way that Microsoft Update distributes updates. Microsoft reduces the risk of sending update files over an unencrypted channel by signing each update. In addition, a hash is computed and sent together with the metadata for each update. When an update is downloaded, WSUS checks the digital signature and hash. If the update has been changed, it is not installed.

Limitations of WSUS SSL deployments

You must consider the following limitations when you use SSL to secure a WSUS deployment:

1. Using SSL increases the server workload. You should expect a 10 percent loss of performance because of the cost of encrypting all the metadata that is sent over the network.
2. If you use WSUS with a remote SQL Server database, the connection between the WSUS server and the database server is not secured by SSL. If the database connection must be secured, consider the following recommendations:

- Move the WSUS database to the WSUS server.
- Move the remote database server and the WSUS server to a private network.
- Deploy Internet Protocol security (IPsec) to help secure network traffic. For more information about IPsec, see Creating and Using IPsec Policies.

Configure SSL on the WSUS server

WSUS requires two ports for SSL: one port that uses HTTPS to send encrypted metadata, and one port that uses HTTP to send updates. When you configure WSUS to use SSL, consider the following:

- You cannot configure the whole WSUS website to require SSL because all traffic to the WSUS site would have to be encrypted. WSUS encrypts update metadata only. If a computer attempts to retrieve update files on the HTTPS port, the transfer will fail.

  You should require SSL for the following virtual roots only:

  - **SimpleAuthWebService**
  - **DSSAuthWebService**
  - **ServerSyncWebService**
  - **APIRemoting30**

- **ClientWebService**

  You should not require SSL for the following virtual roots:

  - **Content**
  - **Inventory**
  - **ReportingWebService**
  - **SelfUpdate**

- The certificate of the certification authority (CA) must be imported into the local computer Trusted Root CA store, or the Windows Server Update Service Trusted Root CA store on downstream WSUS servers. If the certificate is only imported to the Local User Trusted Root CA store, the downstream WSUS server will not be authenticated on the upstream server.

  For more information about how to use SSL certificates in IIS, see Require Secure Sockets Layer (IIS 7).

- You must import the certificate to all computers that will communicate with the WSUS server. This includes all client computers, downstream servers, and computers that run the WSUS Administration Console. The certificate should be imported into the local computer Trusted Root CA store or into the Windows Server Update Service Trusted Root CA store.
- You can use any port for SSL. However, the port that you set up for SSL also determines the port that WSUS uses to send clear HTTP traffic. Consider the following examples:
  - If you use the industry standard port of 443 for HTTPS traffic, WSUS uses the industry standard port 80 for clear HTTP traffic.
  - If you use any port other than 443 for HTTPS traffic, WSUS will send clear HTTP traffic over the port that numerically comes before the port for HTTPS. For example, if you use port 8531 for HTTPS, WSUS will use port 8530 for HTTP.
- You must re-initialize *ClientServicingProxy* if the server name, SSL configuration, or port number are changed.

To configure SSL on the WSUS root server

1. Log on to the WSUS server by using an account that is a member of the WSUS Administrators group or the local Administrators group.
2. Go to **Start**, type **CMD**, right-click **Command Prompt**, and then click **Run as administrator**.
3. Navigate to the *%ProgramFiles%***\Update Services\Tools** folder.
4. In the Command Prompt window, type the following command:

   Wsusutil configuressl *certificateName*

   where:

   certificateName is the DNS name of the WSUS server.

Configure SSL on client computers

When you configure SSL on client computers, you should consider the following issues:

- You must include a URL for a secure port on the WSUS server. Because you cannot require SSL on the server, the only way to make sure that client computers can use a security channel is by using a URL that specifies HTTPS. If you use any port other than 443 for SSL, you must include that port in the URL also.
- The certificate on a client computer must be imported into the Local Computer Trusted Root CA store or Automatic Update Service Trusted Root CA store. If the certificate is imported to the Local User's Trusted Root CA store only, Automatic Updates will fail server authentication.
- The client computers must trust the certificate that you bind to the WSUS server. Depending on the type of certificate that is used, you might have to set up a service to enable the client computers to trust the certificate that is bound to the WSUS server.

Configure SSL for downstream WSUS servers

The following instructions configure a downstream server to synchronize to an upstream server that uses SSL.

To synchronize a downstream server to an upstream server that uses SSL

1. Log on to the computer by using a user account that is a member of the local Administrators group or the WSUS Administrators group.
2. Click **Start**, click **All Programs**, click **Administrative Tools**, and then click **Windows Server Update Service**.
3. In the right pane, expand the server name.
4. Click **Options**, and then click **Update Source and Proxy Server**.
5. On the **Update Source** page, select **Synchronize from another Windows Server Update Services server**.
6. Type the name of the upstream server into the **Server name** text box. Type the port number that the server uses for SSL connections into the **Port number** text box.
7. Select the **Use SSL when synchronizing update information** check box, and then click **OK**.

Additional SSL resources

The steps that are required to set up a certification authority, bind the certificate to the WSUS website, and establish a trust between the client computers and the certificate are beyond the scope of this guide. For more information and for instructions about how to install certificates and set up this environment, see the following topics:

- Suite B PKI Step-by-Step Guide
- Implementing and Administering Certificate Templates
- Active Directory Certificate Services Upgrade and Migration Guide
- Configure Certificate Autoenrollment

3.6. Complete IIS Configuration

By default, anonymous read access is enabled for the default and all new IIS websites. Some applications, notably Windows SharePoint Services, may remove anonymous access. If this has occurred, you must re-enable the anonymous read access before you can successfully install and operate WSUS.

To enable anonymous read access, follow the steps for the applicable version of IIS:

1. Enable Anonymous Authentication (IIS 7), as documented in the IIS 7 Operations Guide.
2. Enabling Anonymous Authentication (IIS 6.0), as documented in the IIS 6.0 Operations Guide.

3.6. Configure a Signing Certificate

WSUS has the ability to publish custom update packages to update Microsoft and non-Microsoft products. WSUS can automatically sign these custom update packages for you with an Authenticode certificate. To enable custom update signing, you must install a package signing certificate on your WSUS server. There are several considerations associated with custom update signing.

1. **Certificate Distribution**. The private key must be installed on the WSUS server, and the public key must be explicitly installed in the trusted certificate store on all client PCs and servers which are to receive custom-signed updates.
2. **Expiration**. When the self-signed certificate expires or nears expiration, WSUS will log events in the event log.
3. **Certificate Updates/Revocation**. If you wanted to update or revoke a certificate (i.e. after discovering that it expired), WSUS offered no functionality to enable this. Accomplishing this turned into a manual task that was very hard to either do by hand or automate successfully.

See also

- Step 4: Approve and Deploy WSUS Updates
- Step 2: Install the WSUS Server Role
- Deploy Windows Server Update Services in Your Organization

Step 4: Approve and Deploy WSUS Updates

Computers in a computer group automatically contact the WSUS server over the next 24 hours to obtain updates. You can use the WSUS reporting feature to determine whether those updates were deployed to the test computers. When the tests are successfully completed, you can approve the updates for the applicable computer groups in your organization. The following checklist describes the steps to approve and deploy updates by using WSUS management console.

Task	Description
4.1. Approve and deploy WSUS updates	Approve and deploy WSUS updates by using the WSUS Management Console.
4.2. Configure auto-approval rules	Configure WSUS to automatically approve installation of updates for selected groups, and how to approve revisions to existing updates.
4.3. Review installed updates with WSUS Reports	Review the updates that were installed, the computers that received those updates and other details by using the WSUS Reporting feature.

4.1. Approve and deploy WSUS updates

Use the following procedure to approve and deploy updates.

To approve and deploy WSUS updates

1. On the WSUS Administration Console, click **Updates**. In the right pane, an update status summary is displayed for **All Updates**, **Critical Updates**, **Security Updates**, and **WSUS Updates**.
2. In the **All Updates** section, click **Updates needed by computers**.
3. In the list of updates, select the updates that you want to approve for installation in your test computer group. Information about a selected update is available in the bottom pane of the **Updates** panel. To select multiple contiguous updates, hold down the **SHIFT** key while clicking the update names. To select multiple noncontiguous updates, press down the **CTRL** key while clicking the update names.
4. Right-click the selection, and then click **Approve**.
5. In the **Approve Updates** dialog box, select your test group, and then click the down arrow.
6. Click **Approved for Install**, and then click **OK**.
7. The **Approval Progress** window appears, which shows the progress of the tasks that affect update approval. When the approval process is complete, click **Close**.

4.2. Configure auto-approval rules

Automatic Approvals enables you to specify how to automatically approve installation of updates for selected groups, and how to approve revisions to existing updates.

To configure Automatic Approvals

1. In the WSUS Administration Console, under **Update Services**, expand the WSUS

server, and then click **Options**. The **Options** window opens.

2. In **Options**, click **Automatic Approvals**. The Automatic Approvals dialog opens.
3. In **Update Rules**, click **New Rule**. The **Add Rule** dialog opens.
4. In **Add Rule**, in **Step 1: Select Properties**, select any single option, or combination of options from the following:
   - **When an update is in a specific classification**
   - **When an update is in a specific product**
   - **Set a deadline for the approval**
5. In **Step 2: Edit the properties**, click each of the options listed, and then select the appropriate options for each.
6. In **Step 3: Specify a name**, type a name for your rule, and then click **OK**.
7. Click **OK** to close the Automatic Approvals dialog.

4.3. Review installed updates with WSUS Reports

24 hours after you approve the updates, you can use the WSUS Reports feature to determine whether the updates were deployed to the test group computers. To check the status of an update, you can use the WSUS Reports feature as follows.

To review updates

1. In the navigation pane of the WSUS Administration Console, click **Reports**.
2. On the **Reports** page, click the **Update Status Summary** report. The **Updates Report** window appears.
3. If you want to filter the list of updates, select the criteria that you want to use, for example, **Include updates in these classifications**, and then click **Run Report**.
4. You will see the **Updates Report** pane. You can check the status of individual updates by selecting the update in the left section of the pane. The last section of the report pane shows the status summary of the update.
5. You can save or print this report by clicking the applicable icon on the toolbar.
6. After you test the updates, you can approve the updates for installation on the applicable computer groups in your organization.

See also

- Step 3: Configure WSUS
- Deploy Windows Server Update Services in Your Organization

Management and Tools for Windows Server 2012 R2 and Windows Server 2012

This section contains tool and reference information for IT pros using Windows Server 2012.

- **Windows PowerShell Module Reference for Windows Server and Client**
  This topic lists the Windows PowerShell modules included with Windows Server 2012 R2, Windows 8.1, Windows Server 2012, and Windows 8. The Windows PowerShell modules in the list support automating the features of those operating systems and provide links to the cmdlet references for each module. These modules enable you to use Windows PowerShell to administer, maintain, configure, and develop new features for Windows Server 2012 R2, Windows 8.1, Windows Server 2012, and Windows 8.
- Manage Multiple, Remote Servers with Server Manager
  The content in this section describes how to use Server Manager in Windows Server 2012 R2 and Windows Server 2012 to manage both local and remote Windows-based servers from desktop computers.
- Deploy Remote Server Administration Tools
  To ease remote server management, you can download and install Remote Server Administration Tools for Windows 8.1. Remote Server Administration Tools for Windows 8.1 includes Server Manager, Microsoft Management Console (MMC) snap-ins, consoles, Windows PowerShell cmdlets and providers, and some command-line tools for managing roles and features that run on Windows Server 2012 R2 and Windows Server 2012.
- Install and Use Windows PowerShell Web Access
  The content in this section describes how to configure and use Windows PowerShell Web Access, a feature that allows a server that is running Web Server (IIS) to act as a Windows PowerShell gateway and provide a web-based Windows PowerShell console to users in your organization.
- Getting Started with Windows PowerShell Workflow
  The content in this section describes Windows PowerShell Workflow, a feature that lets IT pros and developers leverage the benefits of workflows with the automation capabilities and ease of Windows PowerShell.
- **Command-Line Reference**
  The content in this section describes the command-line tools that enable IT pros to create batch files or scripting tools that automate common management tasks.
- Datacenter Abstraction Layer (DAL) Overview
  The content in this section describes DAL, an extensible framework that provides a set of management infrastructure components and schemas that can emulate device-specific details, such as those of network switches.
- Background Intelligent Transfer Service (BITS) Overview
  The content in this section describes BITS, an optional feature that improves the network bandwidth use and efficiency of file transfers.

Windows PowerShell Support for Windows Server

Most of the features in Windows Server provide support for Windows PowerShell® by including modules that enable you to install, configure, and use the feature at a Windows PowerShell prompt in addition to the traditional graphical user interface.

For more information, see the following:

1. Windows PowerShell
2. Windows PowerShell Core Module Reference
3. Windows Server and Client Module Reference

Manage Multiple, Remote Servers with Server Manager

Server Manager is a management console in Windows Server® 2012 R2 and Windows Server® 2012 that helps IT professionals provision and manage both local and remote Windows-based servers from their desktops, without requiring either physical access to servers, or the need to enable Remote Desktop protocol (RDP) connections to each server. Although Server Manager is available in Windows Server 2008 R2 and Windows Server 2008, Server Manager was updated in Windows Server 2012, to support remote, multi-server management, and help increase the number of servers an administrator can manage.

In our tests, Server Manager in Windows Server 2012 R2 and Windows Server 2012 can be used to manage up to 100 servers that are configured with a typical workload. The number of servers that you can manage by using a single Server Manager console can vary depending on the amount of data that you request from managed servers, and hardware and network resources available to the computer running Server Manager. As the amount of data you want to display approaches that computer's resource capacity, you can experience slow responses from Server Manager, and delays in the completion of refreshes. To help increase the number of servers that you can manage by using Server Manager, we recommend limiting the event data that Server Manager gets from your managed servers, by using settings in the **Configure Event Data** dialog box. Configure Event Data can be opened from the **Tasks** menu in the **Events** tile. If you need to manage an enterprise-level number of servers in your organization, we recommend evaluating products in the Microsoft® System Center suite.

This topic and its subtopics provide information about how to use features in the Server Manager console. This topic contains the following sections.

- Review initial considerations and system requirements
- Tasks that you can perform in Server Manager
- Start Server Manager
- Restart remote servers
- Export Server Manager settings to other computers

Review initial considerations and system requirements

The following sections list some initial considerations that you need to review, as well as hardware and software requirements for Server Manager.

Hardware requirements

Server Manager is installed by default with all editions of Windows Server 2012 R2 and Windows Server 2012. No additional hardware requirements exist for Server Manager.

Software and configuration requirements

Server Manager is installed by default with all editions of Windows Server 2012. Although you can use Server Manager to manage Server Core installation options of Windows Server 2012 and Windows Server 2008 R2 that are running on remote computers, Server Manager does not run directly on Server Core installation options.

To fully manage remote servers that are running Windows Server 2008 or Windows Server 2008 R2, install the following updates on the servers you want to manage, in the order shown.

To manage servers that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 by using Server Manager in Windows Server 2012 R2, apply the following updates to the older operating systems.

- .NET Framework 4.5
- Windows Management Framework 4.0. The Windows Management Framework 4.0 download package updates Windows Management Instrumentation (WMI) providers on Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. The updated WMI providers let Server Manager collect information about roles and features that are installed on the managed servers. Until the update is applied, servers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 have a manageability status of **Not accessible**.
- The performance update associated with Knowledge Base article 2682011 allows Server Manager to collect performance data from Windows Server 2008 and Windows Server 2008 R2. This performance update is not necessary on servers that are running Windows Server 2012.

To manage servers that are running Windows Server 2008 R2 or Windows Server 2008, apply the following updates to the older operating systems.

- .NET Framework 4
- Windows Management Framework 3.0 The Windows Management Framework 3.0 download package updates Windows Management Instrumentation (WMI) providers on Windows Server 2008 and Windows Server 2008 R2. The updated WMI providers let Server Manager collect information about roles and features that are installed on the managed servers. Until the update is applied, servers that are running Windows Server 2008 or Windows Server 2008 R2 have a manageability status of **Not accessible – Verify earlier versions run Windows Management Framework 3.0**.
- The performance update associated with Knowledge Base article 2682011 allows Server Manager to collect performance data from Windows Server 2008 and Windows Server 2008 R2.

Server Manager runs in the Minimal Server Graphical Interface; that is, when the Server Graphical Shell feature has been uninstalled. The Server Graphical Shell feature is installed by default on Windows Server 2012 R2 and Windows Server 2012. If you uninstall Server Graphical Shell, the Server Manager console runs, but some applications or tools available from the console are not available. Internet browsers cannot run without Server Graphical Shell, so webpages and applications such as HTML Help (The MMC F1 Help, for example) cannot be opened. You cannot open dialog boxes for configuring Windows automatic updating and feedback when Server Graphical Shell is not installed; commands that open these dialog boxes in the Server Manager console are redirected to run **sconfig.cmd**.

Manage remote computers from a client computer

The Server Manager console is included with Remote Server Administration Tools for Windows® 8.1 and Remote Server Administration Tools for Windows® 8. Note that when Remote Server Administration Tools is installed on a client computer, you cannot manage the local computer by using Server Manager; Server Manager cannot be used to manage computers or devices that are running a Windows client operating system. You can use Server Manager to manage only Windows-based servers.

Server Manager Source Operating System	Targeted at Windows Server 2012 R2	Targeted at Windows Server 2012	Targeted at Windows Server 2008 R2 or Windows Server 2008	Targeted at Windows Server 2003
Windows 8 or Windows Server 2012	Not supported	Full support	After Software and configuration requirements are satisfied, can perform most management tasks, but no role or feature	Limited support; online and offline status only

Server Manager Source Operating System	Targeted at Windows Server 2012 R2	Targeted at Windows Server 2012	Targeted at Windows Server 2008 R2 or Windows Server 2008	Targeted at Windows Server 2003
			installation or uninstallation	
Windows 8.1 or Windows Server 2012 R2	Full support	Full support	After Software and configuration requirements are satisfied, can perform most management tasks, but no role or feature installation or uninstallation	Limited support; online and offline status only

To start Server Manager on a client computer

1. Follow instructions in Deploy Remote Server Administration Tools to install Remote Server Administration Tools for Windows 8.1 or Remote Server Administration Tools for Windows 8.
2. On the **Start** screen, click **Server Manager**. The **Server Manager** tile is available after you install Remote Server Administration Tools.
3. If neither the **Administrative Tools** nor the **Server Manager** tiles are displayed on the **Start** screen after installing Remote Server Administration Tools, and searching for Server Manager on the **Start** screen does not display results, verify that the **Show administrative tools** setting is turned on. To view this setting, hover the mouse cursor over the upper right corner of the **Start** screen, and then click **Settings**. If **Show administrative tools** is turned off, turn the setting on to display tools that you have installed as part of Remote Server Administration Tools.

For more information about running Remote Server Administration Tools for Windows 8 to manage remote servers, see Remote Server Administration Tools on the TechNet Wiki.

Configure remote management on servers that you want to manage

Security

By default, Server Manager and Windows PowerShell remote management is enabled in Windows Server 2012 R2 and Windows Server 2012.

To perform management tasks on remote servers by using Server Manager, remote servers that you want to manage must be configured to allow remote management by using Server Manager and Windows PowerShell. If remote management has been disabled on Windows Server 2012 R2 or Windows Server 2012, and you want to enable it again, perform the following steps.

To configure Server Manager remote management on Windows Server 2012 R2 or Windows Server 2012 by using the Windows interface

1.

Note

The settings that are controlled by the **Configure Remote Management** dialog box do not affect parts of Server Manager that use DCOM for remote communications.

Do one of the following to open Server Manager if it is not already open.

- On the Windows taskbar, click the Server Manager button.
- On the **Start** screen, click **Server Manager**.

2. In the **Properties** area of the **Local Servers** page, click the hyperlinked value for the **Remote management** property.
3. Do one of the following, and then click **OK**.
   - To prevent this computer from being managed remotely by using Server Manager (or Windows PowerShell if it is installed), clear the **Enable remote management of this server from other computers** check box.
   - To let this computer be managed remotely by using Server Manager or Windows PowerShell, select **Enable remote management of this server from other computers**.

To enable Server Manager remote management on Windows Server 2012 R2 or Windows Server 2012 by using Windows PowerShell

1. Do one of the following.
   - To run Windows PowerShell as an administrator from the **Start** screen, right-click the **Windows PowerShell** tile, and then click **Run as Administrator**.
   - To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.
2. Type the following, and then press **Enter** to enable all required firewall rule exceptions.

   Configure-SMRemoting.exe -Enable

   Note

   This command also works in a command prompt that has been opened with elevated user rights (Run as Administrator).

   If enabling remote management fails, see about_Remote_Troubleshooting on Microsoft TechNet for troubleshooting tips and best practices.

To enable Server Manager and Windows PowerShell remote management on older operating systems

- Do one of the following.
  - To enable remote management on servers that are running Windows Server 2008 R2, see Remote Management with Server Manager in the Windows Server 2008 R2 Help.
  - To enable remote management on servers that are running Windows Server 2008, see Enable and Use Remote Commands in Windows PowerShell.
  - To enable remote management on servers that are running Windows Server 2003, enable WMI DCOM exceptions in Windows Firewall. For more information about how to do this on servers that are running Windows Server 2003, see Connecting Through Windows Firewall on MSDN.

Tasks that you can perform in Server Manager

Server Manager makes server administration more efficient by allowing administrators to do tasks in the following table by using a single tool. In Windows Server 2012 R2 and Windows Server 2012, both standard users of a server and members of the Administrators group can perform management tasks in Server Manager, but by default, standard users are prevented from performing some tasks, as shown in the following table.

Administrators can use two Windows PowerShell cmdlets in the Server Manager cmdlet module, Enable-ServerManagerStandardUserRemoting and Disable-ServerManagerStandardUserRemoting, to further control standard user access to some additional data. The **Enable-ServerManagerStandardUserRemoting** cmdlet can provide one or more standard, non-Administrator users access to event, service, performance counter, and role and feature inventory data.

Important
Server Manager cannot be used to manage a newer release of the Windows Server operating system. Server Manager running on Windows Server 2012 or Windows 8 cannot be used to manage servers that are running Windows Server 2012 R2.

Task Description	Administrators (including the built-in Administrator account)	Standard Server Users
Add remote servers to a pool of servers that Server Manager can be used to manage.	Yes	No
Create and edit custom groups of servers, such as servers that are in a specific geographic location or serve a specific	Yes	Yes

Task Description	Administrators (including the built-in Administrator account)	Standard Server Users
purpose.		
Install or uninstall roles, role services, and features on the local or on remote servers that are running Windows Server 2012 R2 or Windows Server 2012. For definitions of roles, role services, and features, see Roles, Role Services, and Features.	Yes	No
View and make changes to server roles and features that are installed on either local or remote servers. **Note** In Server Manager, role and feature data is displayed in the base language of the system, also called the system default GUI language, or the language selected during installation of the operating system.	Yes	Standard users can view and manage roles and features, and perform tasks such as viewing role events, but cannot add or remove role services.
Start management tools such as Windows PowerShell or MMC snap-ins. You can start a Windows PowerShell session targeted at a remote server by right-clicking the server in the **Servers** tile, and then clicking **Windows PowerShell**. You can start MMC snap-ins from the **Tools** menu of the Server Manager console, and then point the MMC toward a remote computer after the snap-in is open.	Yes	Yes

Task Description	Administrators (including the built-in Administrator account)	Standard Server Users
Manage remote servers with different credentials by right-clicking a server in the **Servers** tile, and then clicking **Manage As**. You can use **Manage As** for general server and File and Storage Services management tasks.	Yes	No
Perform management tasks associated with the operational lifecycle of servers, such as starting or stopping services; and start other tools that allow you to configure a server's network settings, users and groups, and Remote Desktop connections.	Yes	Standard users cannot start or stop services. They can change the local server's name, workgroup, or domain membership and Remote Desktop settings, but are prompted by User Account Control to provide Administrator credentials before they can complete these tasks. They cannot change remote management settings.
Perform management tasks associated with the operational lifecycle of roles that are installed on servers, including scanning roles for compliance with best practices.	Yes	Standard users cannot run Best Practices Analyzer scans.
Determine server status, identify critical events, and analyze and troubleshoot configuration issues or failures.	Yes	Yes
Customize the events, performance data, services, and Best Practices Analyzer results about which you want to be alerted on the Server Manager dashboard.	Yes	Yes
Restart servers.	Yes	No

Task Description	Administrators (including the built-in Administrator account)	Standard Server Users
Refresh data that is displayed in the Server Manager console about managed servers.	Yes	No

Note
Server Manager can receive only online or offline status from servers that are running Windows Server 2003. Server Manager cannot be used to add roles and features to servers that are running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003.

Start Server Manager

Server Manager starts automatically by default on servers that are running Windows Server 2012 when a member of the Administrators group logs on to a server. If you close Server Manager, restart it in one of the following ways. This section also contains steps for changing the default behavior, and preventing Server Manager from starting automatically.

To start Server Manager from the Start screen

- On the Windows **Start** screen, click the **Server Manager** tile.

To start Server Manager from the Windows desktop

- On the Windows taskbar, click **Server Manager**.

To prevent Server Manager from starting automatically

1. In the Server Manager console, on the **Manage** menu, click **Server Manager Properties**.
2. In the **Server Manager Properties** dialog box, fill the check box for **Do not start Server Manager automatically at logon**. Click **OK**.
3. Alternatively, you can prevent Server Manager from starting automatically by enabling the Group Policy setting, **Do not start Server Manager automatically at logon**. The path to this policy setting, in the Local Group Policy Editor console, is Computer Configuration\Administrative Templates\System\Server Manager.

Restart remote servers

You can restart a remote server from the **Servers** tile of a role or group page in Server Manager.

Important
Restarting a remote server forces the server to restart, even if users are still logged on to the remote server, and even if programs with unsaved data are still open. This behavior

is different from shutting down or restarting the local computer, on which you would be prompted to save unsaved program data, and verify that you wanted to force logged-on users to log off. Be sure that you can force other users to log off of remote servers, and that you can discard unsaved data in programs that are running on the remote servers.

If an automatic refresh occurs in Server Manager while a managed server is shutting down and restarting, refresh and manageability status errors can occur for the managed server, because Server Manager cannot connect to the remote server until it is finished restarting.

To restart remote servers in Server Manager

1. Open a role or server group home page in Server Manager.
2. Select one or more remote servers that you have added to Server Manager. Press and hold **Ctrl** as you click to select multiple servers at one time. For more information about how to add servers to the Server Manager server pool, see Add Servers to Server Manager.
3. Right-click selected servers, and then click **Restart Server**.

Export Server Manager settings to other computers

In Server Manager, your list of managed servers, changes to Server Manager console settings, and custom groups that you have created are stored in the following two files. You can reuse these settings on other computers that are running the same release of Server Manager (not computers that are running the Server Core installation option) or Windows 8. Remote Server Administration Tools must be running on Windows client-based computers to export Server Manager settings to those computers.

- %*appdata*%\Microsoft\Windows\ServerManager\ServerList.xml
- %*appdata*%\Local\Microsoft_Corporation\ServerManager.exe_StrongName_*GUID*\6.2.0.0\user.config

Notes

- Manage As (or alternate) credentials for servers in your server pool are not stored in the roaming profile. Server Manager users must add them on each computer from which they want to manage.
- The network share roaming profile is not created until a user logs on to the network, and then logs off for the first time. The **Serverlist.xml** file is created at this time.

You can export Server Manager settings, make Server Manager settings portable, or use them on other computers in one of the following two ways.

- To export settings to another domain-joined computer, configure the Server Manager user to have a roaming profile in Active Directory Users and Computers. You must be a Domain Administrator to change user properties in Active Directory Users and Computers.

- To export settings to another computer in a workgroup, copy the preceding two files to the same location on the computer from which you want to manage by using Server Manager.

To export Server Manager settings to other domain-joined computers

1. In Active Directory Users and Computers, open the **Properties** dialog box for a Server Manager user.
2. On the **Profile** tab, add a path to a network share to store the user's profile.
3. Do one of the following.
   - On U.S. English (en-us) builds, changes to the **ServerList.xml** file are automatically saved to the profile. Go on to the next step.
   - On other builds, copy the following two files from the computer that is running Server Manager to the network share that is part of the user's roaming profile.
     - %*appdata*%\Microsoft\Windows\ServerManager\ServerList.xml
     - %*localappdata*%\Microsoft_Corporation\ServerManager.exe_StrongName_*GUID*\6.2.0.0\user.config
4. Click **OK** to save your changes and close the **Properties** dialog box.

To export Server Manager settings to computers in workgroups

- On a computer from which you want to manage remote servers, overwrite the following two files with the same files from another computer that is running Server Manager, and that has the settings you want.
  - %*appdata*%\Microsoft\Windows\ServerManager\ServerList.xml
  - %*localappdata*%\Microsoft_Corporation\ServerManager.exe_StrongName_*GUID*\6.2.0.0\user.config

Manage the Local Server and the Server Manager Console

In Windows Server® 2012 R2 and Windows Server® 2012, Server Manager lets you manage both the local server (if you are running Server Manager on Windows Server, and not on Windows 8.1 or Windows 8) and remote servers that are running Windows Server 2003 and newer releases of the Windows Server operating system.

The **Local Server** page in Server Manager displays server properties, events, service and performance counter data, and Best Practices Analyzer (BPA) results for the local server. Event, service, BPA, and performance tiles function as they do on role and server group pages. For more information about configuring the data that is displayed in these tiles, see View and

Configure Performance, Event, and Service Data and Run Best Practices Analyzer Scans and Manage Scan Results.

Menu commands and settings in the Server Manager console heading bars apply globally to all servers in your server pool, and let you use Server Manager to manage the entire server pool.

This topic contains the following sections.

- Shut down the local server
- Configure Server Manager properties
- Manage the Server Manager console
- Customize tools that are displayed in the Tools menu
- Manage roles on role home pages

Shut down the local server

The **Tasks** menu in the local server **Properties** tile lets you start a Windows PowerShell session on the local server, open the **Computer Management** MMC snap-in, or open MMC snap-ins for roles or features that are installed on the local server. You can also shut down the local server by using the **Shut Down Local Server** command in this **Tasks** menu. The **Shut Down Local Server** command is also available for the local server in the **Servers** tile on the **All Servers** page, or on any role or group page in which the local server is represented.

Shutting down the local server by using this method, unlike shutting down Windows Server 2012 R2 or Windows Server 2012 from the **Start** screen, opens the **Shut Down Windows** dialog box, which lets you specify reasons for shutdown in the **Shutdown Event Tracker** area.

Note

Only members of the Administrators group can shut down or restart a server. Standard users cannot shut down or restart a server. Clicking the **Shut Down Local Server** command logs standard users off server sessions. This matches the experience of a standard user running the **Alt+F4** shutdown command from the server desktop.

Configure Server Manager properties

You can view or change the following settings in the **Properties** tile on the **Local Server** page. To change a setting's value, click the hypertext value of the setting.

Notes

Typically, the properties displayed in the Local Server **Properties** tile can only be changed on the local server. You cannot change the local server properties from a remote computer by using Server Manager because the **Properties** tile can only get information about the local computer, not remote computers.

Because many properties displayed in the **Properties** tile are controlled by tools that are not part of Server Manager (Control Panel, for example), changes to **Properties** settings are not always displayed in the **Properties** tile immediately. By default, data in the

Properties tile is refreshed every two minutes. To refresh **Properties** tile data immediately, click **Refresh** in the Server Manager address bar.

Setting	Description
Computer name	Displays the computer friendly name, and opens the **System Properties** dialog box, which lets you change the server's name, domain membership, and other system settings such as user profiles.
Domain (or Workgroup, if the server is not joined to a domain)	Displays the domain or workgroup of which the server is a member. Opens the **System Properties** dialog box, which lets you change the server's name, domain membership, and other system settings such as user profiles.
Windows Firewall	Displays Windows Firewall status for the local server. Opens **Control Panel\System and Security\Windows Firewall**. For more information about configuring Windows Firewall, see Windows Firewall with Advanced Security and IPsec.
Remote management	Displays Server Manager and Windows PowerShell remote management status. Opens the **Configure Remote Management** dialog box. For more information about remote management, see Configure Remote Management in Server Manager.
Remote Desktop	Shows whether users can connect to the server remotely by using Remote Desktop sessions. Opens the **Remote** tab of the **System Properties** dialog box.
NIC Teaming	Shows whether the local server is participating in NIC teaming. Opens the **NIC Teaming** dialog box, and lets you join the local server to a NIC team if desired. For more information about NIC Teaming, see the NIC Teaming white paper.
Ethernet	Displays the networking status of the server. Opens **Control Panel\Network and Internet\Network Connections**.
Operating system version	This read-only field displays the version number

Setting	Description
	of the Windows operating system that the local server is running.
Hardware information	This read-only field displays the manufacturer and model name and number of the server hardware.
Last installed updates	Displays the day and time that Windows updates were last installed. Opens **Control Panel\System and Security\Windows Update**.
Windows Update	Displays Windows Update settings for the local server. Opens **Control Panel\System and Security\Windows Update**.
Last checked for updates	Displays the day and time that the server last checked for available Windows updates. Opens **Control Panel\System and Security\Windows Update**.
Windows Error Reporting	Displays Windows Error Reporting opt-in status. Opens the **Windows Error Reporting Configuration** dialog box. For more information about Windows Error Reporting, its benefits, privacy statements, and opt-in settings, see Windows Error Reporting.
Customer Experience Improvement Program	Displays Windows Customer Experience Improvement Program opt-in status. Opens the **Customer Experience Improvement Program Configuration** dialog box. For more information about Windows Customer Experience Improvement Program, its benefits, and opt-in settings, see Windows Customer Experience Improvement Program.
Internet Explorer (IE) Enhanced Security Configuration	Shows whether IE Enhanced Security Configuration (also known as IE hardening or IE ESC) is turned on or off. Opens the **Internet Explorer Enhanced Security Configuration** dialog box. IE Enhanced Security Configuration is a security measure for servers that prevents web pages from opening in Internet Explorer. For more information about IE Enhanced

Setting	Description
	Security Configuration, its benefits, and settings, see Internet Explorer: Enhanced Security Configuration.
Time zone	Displays the local server's time zone. Opens the **Date and Time** dialog box.
Product ID	Displays the Windows activation status and product ID number (if Windows has been activated) of the Windows Server 2012 R2 or Windows Server 2012 operating system. This is not the same number as the Windows product key. Opens the **Windows Activation** dialog box.
Processors	This read-only field displays manufacturer, model name, and speed information about the local server's processors.
Installed memory (RAM)	This read-only field displays the amount of available RAM, in gigabytes.
Total disk space	This read-only field displays the amount of available disk space, in gigabytes.

Manage the Server Manager console

Global settings that apply to the entire Server Manager console, and to all remote servers that have been added to the Server Manager server pool, are found in the heading bars at the top of the Server Manager console window.

Add servers to Server Manager

The command that opens the **Add Servers** dialog box, and lets you add remote physical or virtual servers to the Server Manager server pool, is in the **Manage** menu of the Server Manager console. For detailed information about how to add servers, see Add Servers to Server Manager.

Refresh data that is displayed in Server Manager

You can configure the refresh interval for data that is displayed in Server Manager on the **Server Manager Properties** dialog box, which you open from the **Manage** menu.

To configure the refresh interval in Server Manager

1. On the **Manage** menu in the Server Manager console, click **Server Manager Properties**.
2. In the **Server Manager Properties** dialog box, specify a time period, in minutes, for the amount of elapsed time you want between refreshes of the data that is displayed in Server Manager. The default is 10 minutes. Click OK when you are finished.

Refresh limitations

Refresh applies globally to data from all servers that you have added to the Server Manager server pool. You cannot refresh data or configure different refresh intervals for individual servers, roles, or groups.

When servers that are in a cluster are added to Server Manager, whether they are physical computers or virtual machines, the first refresh of data can fail, or display data only for the host server for clustered objects. Subsequent refreshes show accurate data for physical or virtual servers in a server cluster.

Data that is displayed in role home pages in Server Manager for Remote Desktop Services, IP Address Management, and File and Storage Services does not refresh automatically. Refresh data that is displayed in these pages manually, by pressing **F5** or clicking **Refresh** in the Server Manager console heading while you are on those pages.

Add or remove roles or features

The commands that open the Add Roles and Features Wizard and Remove Roles and Features Wizard, and let you add or remove roles, role services, and features to servers in your server pool, are in the **Manage** menu of the Server Manager console, and the **Tasks** menu of the **Roles and Features** tile on role or group pages. For detailed information about how to add or remove roles or features, see Install or Uninstall Roles, Role Services, or Features.

In Server Manager, role and feature data is displayed in the base language of the system, also called the system default GUI language, or the language selected during installation of the operating system.

Create server groups

The command that opens the **Create Server Group** dialog box, and lets you create custom groups of servers, is in the **Manage** menu of the Server Manager console. For detailed information about how to create server groups, see Create and Manage Server Groups.

Prevent Server Manager from opening automatically at logon

The **Do not start Server Manager automatically at logon** check box in the **Server Manager Properties** dialog box controls whether Server Manager opens automatically at logon for members of the Administrators group on a local server. This setting does not affect Server Manager behavior when it is running on Windows 8.1 or Windows 8 as part of Remote Server Administration Tools. For more information about configuring this setting, see Manage Multiple, Remote Servers with Server Manager.

Zoom in or out

To zoom in or out on your view of the Server Manager console, you can either use the **Zoom** commands on the **View** menu, or press **Ctrl+Plus (+)** to zoom in and **Ctrl+Minus (-)** to zoom out.

Customize tools that are displayed in the Tools menu

The **Tools** menu in Server Manager includes soft links to shortcuts in the **Administrative Tools** folder in **Control Panel/System and Security**. The **Administrative Tools** folder contains a list of shortcuts or LNK files to available management tools, such as MMC snap-ins. Server Manager populates the **Tools** menu with links to those shortcuts, and copies the folder structure of the **Administrative Tools** folder to the **Tools** menu. By default, tools in the Administrative Tools folder are arranged in a flat list, sorted by type and by name. In the Server Manager **Tools** menu, items are sorted only by name, not by type.

To customize the **Tools** menu, copy tool or script shortcuts that you want to use to the **Administrative Tools** folder. You can also organize your shortcuts in folders, which create cascading menus in the **Tools** menu. Additionally, if you want to restrict access to the custom tools on the **Tools** menu, you can set user access rights on both your custom tool folders in Administrative Tools, or directly on the original tool or script files.

We recommend against reorganizing system and administrative tools, and any management tools associated with roles and features that are installed on the local server. Moving role and feature management tools can prevent successful uninstallation of those management tools, when necessary. After uninstallation of a role or feature, a nonfunctional link to a tool whose shortcut has been moved might remain in the **Tools** menu. If you reinstall the role, a duplicate link to the same tool is created in the **Tools** menu, but one of the links will not work.

Role and feature tools that are installed as part of Remote Server Administration Tools on a Windows client-based computer can be organized into custom folders, however. Uninstalling the parent role or feature has no effect on the tool shortcuts that are available on a remote computer that is running Windows 8.1 or Windows 8.

The following procedure describes how to create an example folder called *MyTools*, and move shortcuts for two Windows PowerShell scripts into the folder that are then accessible from the Server Manager Tools menu.

To customize the Tools menu by adding shortcuts in Administrative Tools

1. Create a new folder called *MyTools* in a convenient location.

   Note
   Because of restrictive access rights on the **Administrative Tools** folder, you are not allowed to create a new folder directly in the **Administrative Tools** folder; you must create a new folder elsewhere (such as on the Desktop), and then copy the new folder to the **Administrative Tools** folder.

2. Move or copy *MyTools* to **Control Panel/System and Security/Administrative Tools**. By default, you must be a member of the Administrators group on the computer to make changes to the **Administrative Tools** folder.
3. If you do not need to restrict user access rights to your custom tool shortcuts, go on to step 6. Otherwise, right-click either the tool file (or the *MyTools* folder), and then click **Properties**.
4. On the **Security** tab of the file's **Properties** dialog box, click **Edit**.
5. For users for whom you want to restrict tool access, clear check boxes for **Read & execute**, **Read**, and **Write** permissions. These permissions are inherited by the tool shortcut n the **Administrative Tools** folder.

   If you edit access rights for a user while the user is using Server Manager (or while Server Manager is open), then your changes are not shown in the **Tools** menu until the user restarts Server Manager.

   Notes

   If you restrict access to an entire folder that you have copied to Administrative Tools, restricted users can see neither the folder nor its contents in the Server Manager **Tools** menu.

   Edit permissions for the folder in the **Administrative Tools** folder. Because hidden files and folders in Administrative Tools are always displayed in the Server Manager **Tools** menu, do not use the **Hidden** setting on a file or folder's **Properties** dialog box to restrict user access to your custom tool shortcuts.

   Deny permissions always overwrite **Allow** permissions.
6. Right-click the original tool, script, or executable file for which you want to add entries on the **Tools** menu, and then click **Create shortcut**.
7. Move the shortcut to the *MyTools* folder in Administrative Tools.
8. Refresh or restart Server Manager, if necessary, to see your custom tool shortcut in the **Tools** menu.

Manage roles on role home pages

After you add servers to the Server Manager server pool, and Server Manager collects inventory data about servers in your pool, Server Manager adds pages to the navigation pane for roles that are discovered on managed servers. The **Servers** tile on role pages lists managed servers that are running the role. By default, **Events**, **Best Practices Analyzer**, **Services**, and **Performance** tiles display data for all servers that are running the role; selecting specific servers in the **Servers** tile limits the scope of events, services, performance counters, and BPA results to selected servers only. Management tools are typically available in the Server Manager console **Tools** menu, after a role or feature has been installed or discovered on a managed server. You can also right-click server entries in the **Servers** tile for a role or group, and then start the management tool that you want to use.

In Windows Server 2012, the following roles and feature have management tools that are integrated into Server Manager console as pages.

- **File and Storage Services.** File and Storage Services pages include custom tiles and commands for managing volumes, shares, iSCSI virtual disks, and storage pools. When you open the File and Storage Services role home page in Server Manager, a retracting pane opens that displays custom management pages for File and Storage Services. For more information about deploying and managing File and Storage Services, see File and Storage Services.
- **Remote Desktop Services.** Remote Desktop Services pages include custom tiles and commands for managing sessions, licenses, gateways, and virtual desktops. For more information about deploying and managing Remote Desktop Services, see Remote Desktop Services (RDS).
- **IP Address Management (IPAM).** The IPAM role page includes a custom **Welcome** tile containing links to common IPAM configuration and management tasks, including a wizard for provisioning an IPAM server. The IPAM home page also includes tiles for viewing the managed network, configuration summary, and scheduled tasks.

  There are some limitations to IPAM management in Server Manager. Unlike typical role and group pages, IPAM has no **Servers**, **Events**, **Performance**, **Best Practices Analyzer**, or **Services** tiles. There is no Best Practices Analyzer model available for IPAM; Best Practices Analyzer scans on IPAM are not supported. To access servers in your server pool that are running IPAM, create a custom group of those servers that are running IPAM, and access the server list from the **Servers** tile on the custom group page. Alternatively, access IPAM servers from the **Servers** tile on the **All Servers** group page.

  Dashboard thumbnails also display limited rows for IPAM, compared to thumbnails for other roles and groups. By clicking the IPAM thumbnail rows, you can view events, performance data, and manageability status alerts for servers that are running IPAM. IPAM-related services can be managed from pages for server groups that contain IPAM servers, such as the page for the **All Servers** group.

  For more information about deploying and managing IPAM, see IP Address Management (IPAM).

See Also

Manage Multiple, Remote Servers with Server Manager

Add Servers to Server Manager

Create and Manage Server Groups

View and Configure Performance, Event, and Service Data

File and Storage Services

Remote Desktop Services (RDS)

IP Address Management (IPAM)

Configure Remote Management in Server Manager

In Windows Server® 2012 R2 and Windows Server® 2012, you can use Server Manager to perform management tasks on remote servers. Remote management is enabled by default on servers that are running Windows Server 2012. To manage a server remotely by using Server Manager, you add the server to the Server Manager server pool.

You can use Server Manager to manage remote servers that are running older releases of Windows Server, but the following updates are required to fully manage these older operating systems.

To manage servers that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 by using Server Manager in Windows Server 2012 R2, apply the following updates to the older operating systems.

- .NET Framework 4.5
- Windows Management Framework 4.0. The Windows Management Framework 4.0 download package updates Windows Management Instrumentation (WMI) providers on Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. The updated WMI providers let Server Manager collect information about roles and features that are installed on the managed servers. Until the update is applied, servers that are running Windows Server 2012, Windows Server 2008 R2 or Windows Server 2008 have a manageability status of **Not accessible**.
- The performance update associated with Knowledge Base article 2682011 allows Server Manager to collect performance data from Windows Server 2008 and Windows Server 2008 R2. This performance update is not necessary on servers that are running Windows Server 2012.

To manage servers that are running Windows Server 2008 R2 or Windows Server 2008, apply the following updates to the older operating systems.

- .NET Framework 4
- Windows Management Framework 3.0 The Windows Management Framework 3.0 download package updates Windows Management Instrumentation (WMI) providers on Windows Server 2008 and Windows Server 2008 R2. The updated WMI providers let Server Manager collect information about roles and features that are installed on the managed servers. Until the update is applied, servers that are running Windows Server 2008 or Windows Server 2008 R2 have a manageability status of **Not accessible – Verify earlier versions run Windows Management Framework 3.0**.
- The performance update associated with Knowledge Base article 2682011 allows Server Manager to collect performance data from Windows Server 2008 and Windows Server 2008 R2.

For detailed information about how to add servers that are in workgroups to manage, or manage remote servers from a workgroup computer that is running Server Manager, see Add Servers.

Enabling or disabling remote management

In Windows Server 2012 R2 and Windows Server 2012, remote management is enabled by default. Before administrators can connect to a computer that is running Windows Server 2012 R2 or Windows Server 2012 remotely by using Server Manager, Server Manager remote management must be enabled on the destination computer if it has been disabled. The procedures in this section describe how to disable remote management, and how to re-enable remote management if it has been disabled. In the Server Manager console, the remote management status for the local server is displayed in the **Properties** area of the **Local Server** page.

Local administrator accounts other than the built-in Administrator account may not have rights to manage a server remotely, even if remote management is enabled. The Remote User Account Control (UAC) **LocalAccountTokenFilterPolicy** registry setting must be configured to allow local accounts of the Administrators group other than the built-in administrator account to remotely manage the server.

In Windows Server 2012 R2 and Windows Server 2012, Server Manager relies on Windows Remote Management (WinRM) and the Distributed Component Object Model (DCOM) for remote communications. The settings that are controlled by the **Configure Remote Management** dialog box only affect parts of Server Manager and Windows PowerShell that use WinRM for remote communications. They do not affect parts of Server Manager that use DCOM for remote communications. For example, Server Manager uses WinRM to communicate with remote servers that are running Windows Server 2012 R2 or Windows Server 2012, but uses DCOM to communicate with servers that are running Windows Server 2008 and Windows Server 2008 R2, but do not have the Windows Management Framework 4.0 or the Windows Management Framework 3.0 update applied. Microsoft Management Console (MMC) and other legacy management tools use DCOM. For more information about how to change these settings, see To configure MMC or other tool remote management over DCOM in this topic.

Note

Procedures in this section can be completed only on computers that are running Windows Server. You cannot enable or disable remote management on a computer that is running Windows 8.1 or Windows 8 by using these procedures, because the client operating system cannot be managed by using Server Manager.

To enable WinRM remote management, select one of the following procedures.

- To enable Server Manager remote management by using the Windows interface
- To enable Server Manager remote management by using Windows PowerShell
- To enable Server Manager remote management by using the command line
- To enable Server Manager and Windows PowerShell remote management on earlier releases of Windows Server

To disable WinRM and Server Manager remote management, select one of the following procedures.

- To disable remote management by using Group Policy

- To disable remote management by using an answer file during unattended installation

To configure DCOM remote management, see To configure MMC or other tool remote management over DCOM.

To enable Server Manager remote management by using the Windows interface

1.

   Note
   The settings that are controlled by the **Configure Remote Management** dialog box do not affect parts of Server Manager that use DCOM for remote communications.

   On the computer that you want to manage remotely, open Server Manager, if it is not already open. On the Windows taskbar, click **Server Manager**. On the **Start** screen, click the **Server Manager** tile.
2. In the **Properties** area of the **Local Servers** page, click the hyperlinked value for the **Remote management** property.
3. Do one of the following, and then click **OK**.
   - To prevent this computer from being managed remotely by using Server Manager (or Windows PowerShell if it is installed), clear the **Enable remote management of this server from other computers** check box.
   - To let this computer be managed remotely by using Server Manager or Windows PowerShell, select **Enable remote management of this server from other computers**.

To enable Server Manager remote management by using Windows PowerShell

1. On the computer that you want to manage remotely, do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then on the app bar, click **Run as Administrator**.
2. Type the following, and then press **Enter** to enable all required firewall rule exceptions.

   Configure-SMRemoting.exe -enable

To enable Server Manager remote management by using the command line

1. On the computer that you want to manage remotely, open a command prompt session with elevated user rights. To do this, on the **Start** screen, type **cmd**, right-click the **Command Prompt** tile when it is displayed in the **Apps** results, and then on the app bar, click **Run as Administrator**.
2. Run the following executable file.

   %windir%\system32\Configure-SMRemoting.exe

3. Do one of the following:
   - To disable remote management, type **Configure-SMRemoting.exe -disable**, and then press **Enter**.
   - To enable remote management, type **Configure-SMRemoting.exe -enable**, and then press **Enter**.
   - To view the current remote management setting, type **Configure-SMRemoting.exe -get**, and then press ENTER.

To enable Server Manager and Windows PowerShell remote management on earlier releases of Windows Server

- Do one of the following:
  - To enable remote management on servers that are running Windows Server 2012, see To enable Server Manager remote management by using the Windows interface in this topic.
  - To enable remote management on servers that are running Windows Server 2008 R2, see Remote Management with Server Manager in the Windows Server 2008 R2 Help.
  - To enable remote management on servers that are running Windows Server 2008, see Enable and Use Remote Commands in Windows PowerShell.
  - To enable remote management on servers that are running Windows Server 2003, see How to Install Windows PowerShell on Windows Server 2003 and Enable Remote Management to enable remote management by using Windows PowerShell, or Installation and Configuration for Windows Remote Management to enable remote management by using **winrm**.

To configure MMC or other tool remote management over DCOM

1. Do one of the following to open the Windows Firewall with Advanced Security snap-in.
   - In the **Properties** area of the **Local Server** page in Server Manager, click the hypertext value for the **Windows Firewall** property, and then click **Advanced settings**.
   - On the **Start** screen, type **WF.msc**, and then click the snap-in tile when it is displayed in the **Apps** results.
2. In the tree pane, select **Inbound Rules**.
3. Verify that exceptions to the following firewall rules are enabled, and have not been disabled by Group Policy settings. If any are not enabled, go on to the next step.
   - COM+ Network Access (DCOM-In)
   - Remote Event Log Management (NP-In)
   - Remote Event Log Management (RPC)
   - Remote Event Log Management (RPC-EPMAP)
4. Right-click the rules that are not enabled, and then click **Enable Rule** on the context menu.

5. Close the Windows Firewall with Advanced Security snap-in.

To disable remote management by using Group Policy

1. Do one of the following to open Local Group Policy Editor.
   - On a server that is running Windows Server 2012 R2 or Windows Server 2012, on the **Start** screen, type **gpedit.msc**, and then click the **gpedit** tile when it is displayed.
   - On a server that is running Windows Server 2008 R2 or Windows Server 2008, in the **Run** dialog box, type **gpedit.msc**, and then press **Enter**.
2. Open **Computer Configuration\Administrative Templates\Windows Components\Windows Remote Management (WinRM)\WinRM Service**.
3. In the content pane, double-click **Allow remote server management through WinRM**.
4. In the dialog box for the **Allow remote server management through WinRM** policy setting, select **Disabled** to disable remote management. Click **OK** to save your changes and close the policy setting dialog box.

To disable remote management by using an answer file during unattended installation

1. Create an unattended installation answer file for Windows Server 2012 installations by using Windows System Image Manager (Windows SIM). For more information about how to create an answer file and use Windows SIM, see What is Windows System Image Manager? and Step-by-Step: Basic Windows Deployment for IT Professionals.
2. In your answer file, locate the setting **Microsoft-Windows-Web-Services-for-Management-Core\EnableServerRemoteManagement**.
3. To disable Server Manager remote management by default on all servers to which you want to apply the answer file, set **Microsoft-Windows-Web-Services-for-Management-Core \EnableServerRemoteManagement** to **False**.

   Notes

   This setting disables remote management as part of the operating system setup process. Configuring this setting does not prevent an administrator from enabling Server Manager remote management on a server after operating system setup is complete. Administrators can enable Server Manager remote management again by using steps in To enable Server Manager remote management by using the Windows interface or To enable Server Manager remote management by using Windows PowerShell in this topic.

   If you disable remote management by default as part of an unattended installation, and do not enable remote management on the server again after installation, servers to which this answer file is applied cannot be fully managed by using Server Manager. Servers that are running Windows Server 2012 R2 or Windows Server 2012—and that have remote management disabled by default—generate manageability status errors in the Server Manager console after they are added to the Server Manager server pool.

Windows Remote Management (WinRM) listener settings

Server Manager relies on default WinRM listener settings on the remote servers that you want to manage. If the default authentication mechanism or the WinRM listener port number on a remote server has been changed from default settings, Server Manager cannot communicate with the remote server.

The following list shows default WinRM listener settings for managing by using Server Manager.

- The WinRM service is running.
- A WinRM listener is created to accept HTTP requests through port number 5985.
- Port number 5985 is enabled in Windows Firewall settings to allow requests through WinRM.
- Both **Kerberos** and **Negotiate** authentication types are enabled.

The default port number is 5985 for WinRM to communicate with a remote computer.

For more information about how to configure WinRM listener settings, at a command prompt, type **winrm help config**, and then press ENTER.

See Also

Add Servers to Server Manager

Windows PowerShell: about_Remote_Troubleshooting on the Windows Server TechCenter

Description of User Account Control

Install or Uninstall Roles, Role Services, or Features

In Windows Server® 2012 R2 and Windows Server® 2012, the Server Manager console and Windows PowerShell® cmdlets for Server Manager allow installation of roles and features to local or remote servers, or offline virtual hard disks (VHDs). You can install multiple roles and features on a single remote server or offline VHD in a single Add Roles and Features Wizard or Windows PowerShell session.

Important

Server Manager cannot be used to manage a newer release of the Windows Server operating system. Server Manager running on Windows Server 2012 or Windows 8 cannot be used to install roles, role services, and features on servers that are running Windows Server 2012 R2.

You must be logged on to a server as an administrator to install or uninstall roles, role services, and features. If you are logged on to the local computer with an account that does not have administrator rights on your target server, right-click the target server in the **Servers** tile, and then click **Manage As** to provide an account that has administrator rights. The server on which you

want to mount an offline VHD must be added to Server Manager, and you must have Administrator rights on that server.

For more information about what roles, role services, and features are, see Roles, Role Services, and Features.

This topic contains the following sections.

- Install roles, role services, and features by using the Add Roles and Features Wizard
- Install roles, role services, and features by using Windows PowerShell cmdlets
- Remove roles, role services, and features by using the Remove Roles and Features Wizard
- Remove roles, role services, and features by using Windows PowerShell cmdlets
- Install roles and features on multiple servers by running a Windows PowerShell script
- Install .NET Framework 3.5 and other features on-demand

Install roles, role services, and features by using the Add Roles and Features Wizard

In a single session in the Add Roles and Features Wizard, you can install roles, role services, and features on the local server, a remote server that has been added to Server Manager, or an offline VHD. For more information about how to add a server to Server Manager to manage, see Add Servers to Server Manager.

Note

If you are running Server Manager on Windows Server 2012 or Windows 8, you can use the Add Roles and Features Wizard to install roles and features only on servers and offline VHDs that are running Windows Server 2012. If you are running Server Manager on Windows Server 2012 R2 or Windows 8.1, you can use the Add Roles and Features Wizard to install roles and features only on servers and offline VHDs that are running Windows Server 2012 R2.

To install roles and features by using the Add Roles and Features Wizard

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** screen, click the **Server Manager** tile.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Before you begin** page, verify that your destination server and network environment are prepared for the role and feature you want to install. Click **Next**.
4. On the **Select installation type** page, select **Role-based or feature-based installation** to install all parts of roles or features on a single server, or **Remote Desktop Services installation** to install either a virtual machine-based desktop infrastructure or a session-based desktop infrastructure for Remote Desktop Services. The **Remote Desktop**

Services installation option distributes logical parts of the Remote Desktop Services role across different servers as needed by administrators. Click **Next**.

5. On the **Select destination server** page, select a server from the server pool, or select an offline VHD. To select an offline VHD as your destination server, first select the server on which to mount the VHD, and then select the VHD file. For information about how to add servers to your server pool, see Add Servers to Server Manager. After you have selected the destination server, click **Next**.

   Notes

   To install roles and features on offline VHDs, target VHDs must meet the following requirements.

   - VHDs must be running the release of Windows Server that matches the version of Server Manager you are running. See the note at the start of Install roles, role services, and features by using the Add Roles and Features Wizard.
   - VHDs cannot have more than one system volume or partition.
   - The network shared folder in which the VHD file is stored must grant the following access rights to the computer (or local system) account of server that you have selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.
     - **Read/Write** access on the **File Sharing** dialog box.
     - **Full Control** access on the **Security** tab , file or folder **Properties** dialog box.

6. Select roles, select role services for the role if applicable, and then click **Next** to select features.

   As you proceed, the Add Roles and Features Wizard automatically informs you if conflicts were found on the destination server that can prevent selected roles or features from installation or normal operation. You are also prompted to add any roles, role services, or features that are required by the roles or features that you have selected.

   Additionally, if you plan to manage the role remotely, either from another server, or from a Windows client-based computer that is running Remote Server Administration Tools, you can opt not to install management tools and snap-ins for roles on the destination server. By default, in the Add Roles and Features Wizard, management tools are selected for installation.

7. On the **Confirm installation selections** page, review your role, feature, and server selections. If you are ready to install, click **Install**.

   You can also export your selections to an XML-based configuration file that you can use for unattended installations with Windows PowerShell. To export the configuration you specified in this Add Roles and Features Wizard session, click **Export configuration settings**, and then save the XML file to a convenient location.

   The **Specify an alternate source path** command on the **Confirm installation selections** page lets you specify an alternate source path for the files that are required to install roles and features on the selected server. In Windows Server 2012, Features on Demand lets you reduce the amount of disk space used by the operating system, by

removing role and feature files from servers that are exclusively managed remotely. If you have removed role and feature files from a server by using the `Uninstall-WindowsFeature -Remove` cmdlet, you can install roles and features on the server in the future by specifying an alternate source path, or a share on which required role and feature files are stored. The source path or file share must grant **Read** permissions either to the **Everyone** group (not recommended for security reasons), or to the computer account (*DOMAIN**SERVERNAME*$) of the destination server; granting user account access is not sufficient. For more information about Features on Demand, see Windows Server Installation Options.

You can specify a WIM file as an alternate feature file source when you are installing roles, role services, and features on a running, physical server. The source path for a WIM file should be in the following format, with **WIM** as a prefix, and the index in which the feature files are located as a suffix: **WIM:e:\sources\install.wim:4**. However, you cannot use a WIM file directly as a source for installing roles, role services, and features to an offline VHD; you must either mount the offline VHD and point to its mount path for source files, or you must point to a folder that contains a copy of the contents of the WIM file.

8. After you click **Install**, the **Installation progress** page displays installation progress, results, and messages such as warnings, failures, or post-installation configuration steps that are required for the roles or features that you installed. In Windows Server 2012 R2 and Windows Server 2012, you can close the Add Roles and Features Wizard while installation is still in progress, and view installation results or other messages in the **Notifications** area at the top of the Server Manager console. Click the **Notifications** flag icon to see more details about installations or other tasks that you are performing in Server Manager.

Install roles, role services, and features by using Windows PowerShell cmdlets

The Server Manager deployment cmdlets for Windows PowerShell function similarly to the GUI-based Add Roles and Features Wizard and Remove Roles and Features Wizard, with an important difference. In Windows PowerShell, unlike in the Add Roles and Features Wizard, management tools and snap-ins for a role are not included by default. To include management tools as part of a role installation, add the `IncludeManagementTools` parameter to the cmdlet. If you are installing roles and features on a server that is running the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012, you can add a role's management tools to an installation, but GUI-based management tools and snap-ins cannot be installed on servers that are running the Server Core installation option of Windows Server. Only command-line and Windows PowerShell management tools can be installed on the Server Core installation option.

To install roles and features by using the Install-WindowsFeature cmdlet

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

Note

If you are installing roles and features on a remote server, you do not need to run Windows PowerShell with elevated user rights.

- On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
- On the Windows **Start** screen, right-click the tile for Windows PowerShell, and then on the app bar, click **Run as Administrator**.

2. Type **Get-WindowsFeature** and then press **Enter** to view a list of available and installed roles and features on the local server. If the local computer is not a server, or if you want information about a remote server, run **Get-WindowsFeature -ComputerName <***computer_name***>**, in which *computer_name* represents the name of a remote computer that is running Windows Server 2012 R2 or Windows Server 2012.The results of the cmdlet contain the command names of roles and features that you add to your cmdlet in step 4.

Note

In Windows PowerShell 3.0 and later releases of Windows PowerShell, there is no need to import the Server Manager cmdlet module into the Windows PowerShell session before running cmdlets that are part of the module. A module is automatically imported the first time you run a cmdlet that is part of the module. Also, neither Windows PowerShell cmdlets nor the feature names used with the cmdlets are case-sensitive.

3. Type **Get-Help Install-WindowsFeature**, and then press **Enter** to view the syntax and accepted parameters for the `Install-WindowsFeature` cmdlet.
4. Type the following, and then press **Enter**, where *feature_name* represents the command name of a role or feature that you want to install (obtained in step 2), and *computer_name* represents a remote computer on which you want to install roles and features. Separate multiple values for *feature_name* by using commas. The `Restart` parameter automatically restarts the destination server if required by the role or feature installation.

```
Install-WindowsFeature –Name <feature_name> -ComputerName
<computer_name> -Restart
```

To install roles and features on an offline VHD, add both the `ComputerName` parameter and the `VHD` parameter. If you do not add the `ComputerName` parameter, the cmdlet assumes that the local computer is mounted to access the VHD. The `ComputerName` parameter contains the name of the server on which to mount the VHD, and the `VHD` parameter contains the path to the VHD file on the specified server.

Notes

You must add the `ComputerName` parameter if you are running the cmdlet from a computer that is running Windows 8.1 or Windows 8.

To install roles and features on offline VHDs, target VHDs must meet the following requirements.

- VHDs must be running the release of Windows Server that matches the version of Server Manager you are running. See the note at the start of Install roles, role services, and features by using the Add Roles and Features Wizard.
- VHDs cannot have more than one system volume or partition.
- The network shared folder in which the VHD file is stored must grant the following access rights to the computer (or local system) account of server that you have selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.
  - **Read/Write** access on the **File Sharing** dialog box.
  - **Full Control** access on the **Security** tab , file or folder **Properties** dialog box.

```
Install-WindowsFeature -Name <feature_name> –VHD <path> -
ComputerName <computer_name> -Restart
```

Example: The following cmdlet installs the Active Directory Domain Services role and the Group Policy Management feature on a remote server, ContosoDC1. Management tools and snap-ins are added by using the `IncludeManagementTools` parameter, and the destination server is to be restarted automatically, if installation requires that the servers be restarted.

```
Install-WindowsFeature –Name AD-Domain-Services,GPMC –
ComputerName ContosoDC1 –IncludeManagementTools -Restart
```

5. When installation is finished, verify installation by opening the **All Servers** page in Server Manager, selecting a server on which you installed roles and features, and viewing the **Roles and Features** tile on the page for the selected server. You can also run the `Get-WindowsFeature` cmdlet targeted at the selected server (Get-WindowsFeature -ComputerName <*computer_name*>) to view a list of roles and features that are installed on the server.

Remove roles, role services, and features by using the Remove Roles and Features Wizard

You must be logged on to a server as an administrator to uninstall roles, role services, and features. If you are logged on to the local computer with an account that does not have administrator rights on your uninstallation target server, right-click the target server in the **Servers** tile, and then click **Manage As** to provide an account that has administrator rights. The server on which you want to mount an offline VHD must be added to Server Manager, and you must have Administrator rights on that server.

To remove roles and features by using the Remove Roles and Features Wizard

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the

Windows taskbar.

- On the Windows **Start** screen, click the **Server Manager** tile.

2. On the **Manage** menu, click **Remove Roles and Features**.
3. On the **Before you begin** page, verify that you have prepared for removing roles or features from a server. Click **Next**.
4. On the **Select destination server** page, select a server from the server pool, or select an offline VHD. To select an offline VHD, first select the server on which to mount the VHD, and then select the VHD file.

   Notes

   The network shared folder in which the VHD file is stored must grant the following access rights to the computer (or local system) account of server that you have selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.

   - **Read/Write** access on the **File Sharing** dialog box.
   - **Full Control** access on the **Security** tab , file or folder **Properties** dialog box.

   For information about how to add servers to your server pool, see Add Servers to Server Manager. After you have selected the destination server, click **Next**.

   Note

   You can use the Remove Roles and Features Wizard to remove roles and features from servers that are running the same release of Windows Server that supports the version of Server Manager that you are using. You cannot remove roles, role services, or features from servers that are running Windows Server 2012 R2, if you are running Server Manager on Windows Server 2012 or Windows 8. You cannot use the Remove Roles and Features Wizard to remove roles and features from servers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2003.
5. Select roles, select role services for the role if applicable, and then click **Next** to select features.

   As you proceed, the Remove Roles and Features Wizard automatically prompts you to remove any roles, role services, or features that cannot run without the roles or features that you are removing.

   Additionally, you can opt to remove management tools and snap-ins for roles on the destination server. By default, in the Remove Roles and Features Wizard, management tools are selected for removal. You can leave management tools and snap-ins if you plan to use the selected server to manage the role on other remote servers.
6. On the **Confirm removal selections** page, review your role, feature, and server selections. If you are ready to remove the roles or features, click **Remove**.
7. After you click **Remove**, the **Removal progress** page displays removal progress, results, and messages such as warnings, failures, or post-removal configuration steps that are required, such as restarting the destination server. In Windows Server 2012 R2 and

Windows Server 2012, you can close the Remove Roles and Features Wizard while removal is still in progress, and view removal results or other messages in the **Notifications** area at the top of the Server Manager console. Click the **Notifications** flag to see more details about removals or other tasks that you are performing in Server Manager.

Remove roles, role services, and features by using Windows PowerShell cmdlets

The Server Manager deployment cmdlets for Windows PowerShell function similarly to the GUI-based Remove Roles and Features Wizard, with an important difference. In Windows PowerShell, unlike in the Remove Roles and Features Wizard, management tools and snap-ins for a role are not removed by default. To remove management tools as part of a role removal, add the `IncludeManagementTools` parameter to the cmdlet. If you are uninstalling roles and features from a server that is running the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012, this parameter removes command-line and Windows PowerShell management tools for the specified roles and features.

To remove roles and features by using the Uninstall-WindowsFeature cmdlet

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note
   If you are uninstalling roles and features from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
2. Type **Get-WindowsFeature** and then press **Enter** to view a list of available and installed roles and features on the local server. If the local computer is not a server, or if you want information about a remote server, run **Get-WindowsFeature -ComputerName <***computer_name***>**, in which *computer_name* represents the name of a remote computer that is running Windows Server 2012 R2 or Windows Server 2012.The results of the cmdlet contain the command names of roles and features that you add to your cmdlet in step 4.

   Note
   In Windows PowerShell 3.0 and later releases of Windows PowerShell, there is no need to import the Server Manager cmdlet module into the Windows PowerShell session before running cmdlets that are part of the module. A module is automatically imported the first time you run a cmdlet that is part of the module. Also, neither Windows PowerShell cmdlets nor the feature names used with the cmdlets are case-sensitive.

3. Type **Get-Help Uninstall-WindowsFeature**, and then press **Enter** to view the syntax and accepted parameters for the `Uninstall-WindowsFeature` cmdlet.
4. Type the following, and then press **Enter**, where *feature_name* represents the command name of a role or feature that you want to remove (obtained in step 2), and *computer_name* represents a remote computer from which you want to remove roles and features. Separate multiple values for *feature_name* by using commas. The `Restart` parameter automatically restarts destination servers if required by the role or feature removal.

```
Uninstall-WindowsFeature –Name <feature_name> -ComputerName
<computer_name> -Restart
```

To remove roles and features from an offline VHD, add both the `ComputerName` parameter and the `VHD` parameter. If you do not add the `ComputerName` parameter, the cmdlet assumes that the local computer is mounted to access the VHD. The `ComputerName` parameter contains the name of the server on which to mount the VHD, and the `VHD` parameter contains the path to the VHD file on the specified server.

Notes

You must add the `ComputerName` parameter if you are running the cmdlet from a computer that is running Windows 8.1 or Windows 8.

The network shared folder in which the VHD file is stored must grant the following access rights to the computer (or local system) account of server that you have selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.

- **Read/Write** access on the **File Sharing** dialog box.
- **Full Control** access on the **Security** tab , file or folder **Properties** dialog box.

```
Uninstall-WindowsFeature –Name <feature_name> –VHD <path> -
ComputerName <computer_name> -Restart
```

Example: The following cmdlet removes the Active Directory Domain Services role and the Group Policy Management feature from a remote server, ContosoDC1. Management tools and snap-ins are also removed, and the destination server is to be restarted automatically, if removal requires that the servers be restarted.

```
Uninstall-WindowsFeature –Name AD-Domain-Services,GPMC –
ComputerName ContosoDC1 –IncludeManagementTools -Restart
```

5. When removal is finished, verify that the roles and features are removed by opening the **All Servers** page in Server Manager, selecting the server from which you removed roles and features, and viewing the **Roles and Features** tile on the page for the selected server. You can also run the `Get-WindowsFeature` cmdlet targeted at the selected server (Get-WindowsFeature -ComputerName <*computer_name*>) to view a list of roles and features that are installed on the server.

Install roles and features on multiple servers by running a Windows PowerShell script

Although you cannot use the Add Roles and Features Wizard to install roles, role services, and features on more than one target server in a single wizard session, you can use a Windows PowerShell script to install roles, role services, and features on multiple target servers that you are managing by using Server Manager. The script that you use to perform batch deployment, as this process is called, points to an XML configuration file that you can create easily by using the Add Roles and Features Wizard, and clicking **Export configuration settings** after advancing through the wizard to the **Confirm installation selections** page of the Add Roles and Features Wizard.

Important

All target servers that are specified in your script must be running the release of Windows Server that matches the version of Server Manager you are running on the local computer. For example, if you are running Server Manager on Windows® 8, you can install roles, role services, and features on servers that are running Windows Server 2012. If GUI-based management tools are added to the installation, the installation process automatically converts target servers that are running the Server Core installation option of Windows Server to the full installation option (server with a full GUI).

The script provided in this section is an example of how batch deployment can be performed by using the `Install-WindowsFeature` cmdlet and a Windows PowerShell script. There are other possible scripts and methods of performing batch deployment to multiple servers. To search for or provide other scripts for deploying roles and features, search the Script Center Repository.

To install roles and features on multiple servers

1. If you have not already done so, create an XML configuration file that contains the roles, role services, and features that you want installed on multiple servers. You can create this configuration file by running the Add Roles and Features Wizard, selecting roles, role services, and features that you want, and clicking **Export configuration settings** after advancing through the wizard to the **Confirm installation selections** page. Save the configuration file to a convenient location. You do not need to click **Install** or complete the wizard if you are running it only to create a configuration file.
2. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
3. Copy and paste the following script into your Windows PowerShell session.

```
function Invoke-WindowsFeatureBatchDeployment {

    param (
```

```
        [parameter(mandatory)]
        [string[]] $ComputerNames,
        [parameter(mandatory)]
        [string] $ConfigurationFilePath
    )

    # Deploy the features on multiple computers
simultaneously.
    $jobs = @()
    foreach($ComputerName in $ComputerNames) {
        $jobs += Start-Job -Command {
            Install-WindowsFeature -ConfigurationFilePath
$using:ConfigurationFilePath -ComputerName
$using:ComputerName -Restart
        }
    }

    Receive-Job -Job $jobs -Wait | Select-Object Success,
RestartNeeded, ExitCode, FeatureResult
}
```

Target servers are automatically restarted if required by the roles and features that you select.

4. Run the function by doing the following.
   a. Create a variable in which to store the names of your target computers, separated by commas. In the following example, the variable `$ServerNames` stores the names of target servers *Contoso_01* and *Contoso_02*. Press **Enter**.

      ```
      # Sample Invocation
      $ServerNames = 'Contoso_01', 'Contoso_02'
      Invoke-WindowsFeatureBatchDeployment -ComputerNames
      $ServerNames -ConfigurationFilePath
      C:\Users\sampleuser\Desktop\DeploymentConfigTemplate.xml
      ```

   b. To run the function, type the following, and then press **Enter**, where `$ServerNames` is an example of the variable that you created in the preceding step, and *C:\Users\Sampleuser\Desktop\DeploymentConfigTemplate.xml* is an example of a path to the configuration file that you created in step 1.

      Invoke-WindowsFeatureBatchDeployment –ComputerNames $ServerNames –ConfigurationFilePath

C:\Users\Sampleuser\Desktop\DeploymentConfigTemplate.xml

5. When installation is finished, verify installation by opening the **All Servers** page in Server Manager, selecting a server on which you installed roles and features, and viewing the **Roles and Features** tile on the page for the selected server. You can also run the `Get-WindowsFeature` cmdlet targeted at a specific server (`Get-WindowsFeature -ComputerName` <*computer_name*>) to view a list of roles and features that are installed on the server.

Install .NET Framework 3.5 and other features on-demand

Starting with Windows Server 2012 and Windows 8, the feature files for .NET Framework 3.5 (which includes .NET Framework 2.0 and .NET Framework 3.0) are not available on the local computer by default. The files have been removed. Files for features that have been removed in a Features on Demand configuration, along with feature files for .NET Framework 3.5, are available through Windows Update. By default, if feature files are not available on the destination server that is running Windows Server 2012 R2 or Windows Server 2012, the installation process searches for the missing files by connecting to Windows Update. You can override the default behavior by configuring a Group Policy setting or specifying an alternate source path during installation, whether you are installing by using the Add Roles and Features Wizard GUI or a command line.

You can install .NET Framework 3.5 by doing one of the following.

- Use To install .NET Framework 3.5 by running the Install-WindowsFeature cmdlet to add the `Source` parameter, and specify a source from which to get .NET Framework 3.5 feature files. If you do not add the `Source` parameter, the installation process first determines if a path to feature files has been specified by Group Policy settings, and if no such path is found, searches for missing feature files by using Windows Update.
- Use To install .NET Framework 3.5 by using the Add Roles and Features Wizard to specify an alternate source file location on the **Confirm installation options** page of the Add Roles and Features Wizard.
- Use To install .NET Framework 3.5 by using DISM to get files from Windows Update by default, or by specifying a source path to installation media.

Configure alternate sources for feature files in Group Policy for .NET Framework 3.5 or other features, if feature files are not found on the local computer.

Important

When you are installing feature files from a remote source, the source path or file share must grant **Read** permissions either to the **Everyone** group (not recommended for security reasons), or to the computer (local system) account of the destination server; granting user account access is not sufficient.

Servers that are in workgroups cannot access external file shares, even if the computer account for the workgroup server has **Read** permissions on the external share. Alternate

source locations that work for workgroup servers include installation media, Windows Update, and VHD or WIM files that are stored on the local workgroup server.

You can specify a WIM file as an alternate feature file source when you are installing roles, role services, and features on a running, physical server. The source path for a WIM file should be in the following format, with **WIM** as a prefix, and the index in which the feature files are located as a suffix: **WIM:e:\sources\install.wim:4**. However, you cannot use a WIM file directly as a source for installing roles, role services, and features to an offline VHD; you must either mount the offline VHD and point to its mount path for source files, or you must point to a folder that contains a copy of the contents of the WIM file.

To install .NET Framework 3.5 by running the Install-WindowsFeature cmdlet

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note
   If you are installing roles and features from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
   - On a server that is running the Server Core installation option of Windows Server 2012 R2 or Windows Server 2012, type **powershell** into a command prompt, and then press **Enter**.

2. Type the following command, and then press **Enter**. In the following example, the source files are located in a side-by-side store (abbreviated to as **SxS**) in installation media on drive D.

   ```
   Install-WindowsFeature NET-Framework-Core –Source
   D:\Sources\SxS
   ```

   If you want the command to use Windows Update as a source for missing feature files, or if a default source has already been configured by using Group Policy, you do not need to add the `Source` parameter unless you want to specify a different source.

To install .NET Framework 3.5 by using the Add Roles and Features Wizard

1. On the **Manage** menu in Server Manager, click **Add Roles and Features**.
2. Select a destination server that is running Windows Server 2012 R2 orWindows Server 2012.
3. On the **Select features** page of the Add Roles and Features Wizard, select **.NET Framework 3.5**.
4. If the local computer is allowed to do so by Group Policy settings, the installation process attempts to get missing feature files by using Windows Update. Click **Install**; you do not

need to go on to the next step.

If Group Policy settings do not allow this, or you want to use another source for the .NET Framework 3.5 feature files, on the **Confirm installation selections** page of the wizard, click **Specify an alternate source path**.

5. Provide a path to a side-by-side store (referred to as **SxS**) in installation media, or to a WIM file. In the following example, installation media is located on drive D.

   D:\Sources\SxS

   To specify a WIM file, add a **WIM:** prefix, and add the index of the image to use in the WIM file as a suffix, as shown in the following example.

   WIM:*server_name***\share\install.wim:3**

6. Click **OK**, and then click **Install**.

To install .NET Framework 3.5 by using DISM

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note
   If you are installing roles and features from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
   - On a server that is running the Server Core installation option, type **powershell** into a command prompt, and then press **Enter**.

2. Run one of the following DISM commands.

   - If the computer has access to Windows Update, or a default source file location has already been configured in Group Policy, run the following command.

     ```
     DISM /Online /Enable-Feature /FeatureName:NetFx3 /All
     ```

   - If the computer has access to installation media, run a command similar to the following. In the following example, the operating system installation media is located on drive D. The `LimitAccess` parameter prevents the command from attempting to contact Windows Update or a server that is running WSUS.

     ```
     DISM /Online /Enable-Feature /FeatureName:NetFx3 /All
     /LimitAccess /Source:d:\sources\sxs
     ```

     Note
     The DISM command is case-sensitive.

Configure alternate sources for feature files in Group Policy

The Group Policy setting described in this section specifies authorized source locations for .NET Framework 3.5 files, and other feature files that have been removed as part of a Features on

Demand configuration. The policy setting **Specify settings for optional component installation and component repair** is located in the **Computer Configuration\Administrative Templates\System** folder in the Group Policy Management Console or Local Group Policy Editor.

Note
You must be a member of the Administrators group to change Group Policy settings on the local computer. If Group Policy settings for the computer you want to manage are controlled at the domain level, you must be a member of the Domain Administrators group to change Group Policy settings.

To configure a default alternate source path in Group Policy

1. In Local Group Policy Editor or Group Policy Management Console, open the following policy setting.

   Computer Configuration\Administrative Templates\System\Specify settings for optional component installation and component repair
2. Select **Enabled** to enable the policy setting, if it is not already enabled.
3. In the **Alternate source file path** text box in the **Options** area, specify a fully qualified path to a shared folder or a WIM file. To specify a WIM file as an alternate source file location, add the prefix **WIM:** to the path, and add the index of the image to use in the WIM file as a suffix. The following are examples of values that you can specify.
   - Path to a shared folder: *****server_name***\share***folder_name*
   - Path to a WIM file, in which **3** represents the index of the image in which the feature files are found: **WIM:***server_name***\share\install.wim:3**
4. If you do not want computers that are controlled by this policy setting to search for missing feature files in Windows Update, select **Never attempt to download payload from Windows Update**.
5. If the computers that are controlled by this policy setting typically receive updates through WSUS, but you prefer to go through Windows Update and not WSUS to find missing feature files, select **Contact Windows Update directly to download repair content instead of Windows Server Update Services (WSUS)**.
6. Click **OK** when you are finished changing this policy setting, and then close the Group Policy editor.

See Also

Windows Server Installation Options

Microsoft .NET Framework 3.5 Deployment Considerations

How to Enable or Disable Windows Features

Configure Features on Demand in Windows Server

This topic describes how to remove feature files in a Features on Demand configuration by using the Uninstall-WindowsFeature cmdlet.

Features on Demand is a feature, introduced in Windows® 8 and Windows Server 2012, that allows you to remove role and feature files (sometimes called feature *payload*) from the operating system to conserve disk space, and install roles and features from remote locations or installation media instead of from local computers. You can remove feature files from running physical or virtual computers. You can also add feature files to or remove feature files from Windows image (WIM) files or offline virtual hard disks (VHDs) to create a reproducible copy of Features on Demand configurations.

In a Features on Demand configuration, when feature files are not available on a computer, if an installation requires those feature files, Windows Server 2012 R2 or Windows Server 2012 can be directed to get the files from a side-by-side feature store (a shared folder that contains feature files, and is available to the computer on the network), from Windows Update, or from installation media. By default, when feature files are not available on the target server, Features on Demand searches for missing feature files by performing the following tasks, in the order shown.

1. Searching in a location that has been specified by users of the Add Roles and Features Wizard or DISM installation commands
2. Evaluating the configuration of the Group Policy setting, **Computer Configuration\Administrative Templates\System\Specify settings for optional component installation and component repair**
3. Searching Windows Update

You can override default Features on Demand behavior by doing any of the following.

- Specifying an alternate source path as part of the `Install-WindowsFeature` cmdlet, by adding the `Source` parameter
- Specifying an alternate source path on the **Confirm installation options** page while you are installing features by using the Add Roles and Features Wizard
- Configuring the Group Policy setting, **Specify settings for optional component installation and component repair**

This topic contains the following sections.

- Create a feature file or side-by-side store
- Methods of removing feature files
- Remove feature files by using Uninstall-WindowsFeature

Create a feature file or side-by-side store

This section describes how to set up a remote feature file shared folder (also called a side-by-side store) that stores the files required to install roles, role services, and features on servers that

run Windows Server 2012 R2 or Windows Server 2012. After you have set up a feature store, you can install roles, role services, and features on servers that are running those operating systems, and specify the feature store as the location of installation source files.

To create a feature file store

1. Create a shared folder on a server on your network. For example, *network**share**sxs*.
2. Verify that you have the correct permissions assigned to the feature store. The source path or file share must grant **Read** permissions either to the **Everyone** group (not recommended for security reasons), or to the computer accounts (*DOMAIN**SERVERNAME*$) of servers on which you plan to install features by using this feature store; granting user account access is not sufficient.

   You can access file sharing and permissions settings by doing either of the following on the Windows desktop.

   - Right-click the shared folder, click **Properties**, and then change allowed users and their access rights to the folder on the **Security** tab.
   - Right-click the shared folder, point to **Share with**, and then click **Specific people**.

   Note
   Servers that are in workgroups cannot access external file shares, even if the computer account for the workgroup server has **Read** permissions on the external share. Alternate source locations that work for workgroup servers include installation media, Windows Update, and VHD or WIM files that are stored on the local workgroup server.
3. Copy the **Sources\SxS** folder from your Windows Server installation media to the shared folder that you created in step 1.

Methods of removing feature files

Two methods are available for removing feature files from Windows Server in a Features on Demand configuration.

- The `Remove` parameter of the `Uninstall-WindowsFeature` cmdlet lets you delete feature files from a server or offline virtual hard disk (VHD) that is running Windows Server 2012 R2 or Windows Server 2012. Valid values for the `Remove` parameter are the names of roles, role services, and features.
- Deployment Image Servicing and Management (DISM) commands let you create custom WIM files that conserve disk space by omitting feature files that are either not needed, or can be obtained from other, remote sources. For more information about using DISM to prepare custom images, see How to Enable or Disable Windows Features.

Remove feature files by using Uninstall-WindowsFeature

You can use the Uninstall-WindowsFeature cmdlet both to uninstall roles, role services, and features from servers and offline VHDs that are running Windows Server 2012 R2 or Windows Server 2012, and to delete feature files. You can both uninstall and delete the same roles, role services, and features in the same command if desired.

Important

When you delete feature files for a role, role service, or feature, roles, role services, and features that depend upon the files you are removing are also deleted. If you are deleting feature files for a role service or subfeature, and no other role services or subfeatures for the parent role or feature remain installed, then files for the entire parent role or feature are deleted. To view all feature files that would be deleted by the `Uninstall-WindowsFeature -Remove` command, add the `whatIf` parameter to the command to run it and view results without actually deleting feature files.

To remove role and feature files by using Uninstall-WindowsFeature

1. Do one of the following to open a Windows PowerShell session with elevated user rights.

   Note

   If you are uninstalling roles and features from a remote server, you do not need to run Windows PowerShell with elevated user rights.

   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click the Windows PowerShell tile, and then on the app bar, click **Run as Administrator**.
   - On a server that is running the Server Core installation option, type **powershell** into a command prompt, and then press **Enter**.

2. Type the following, and then press **Enter**.

   ```
   Uninstall-WindowsFeature -Name <feature_name> -ComputerName
   <computer_name> -Remove
   ```

   Example: Remote Desktop Licensing is the last remaining role service of Remote Desktop Services that is installed. The command uninstalls Remote Desktop Licensing, and then deletes feature files for the entire Remote Desktop Services role from the specified server, *contoso_1*.

   ```
   Uninstall-WindowsFeature -Name RDS-Licensing -ComputerName
   contoso_1 -Remove
   ```

   Example: In the following example, the command removes Active Directory Domain Services and Group Policy Management from an offline VHD. The role and feature are first uninstalled, then their feature files removed entirely from the offline VHD, *Contoso.vhd*.

Notes

You must add the `ComputerName` parameter if you are running the cmdlet from a computer that is running Windows 8.1 or Windows 8.

If you enter the name of a VHD file from a network share, that share must grant **Read** and **Write** permissions to the computer account of the server that you selected to mount the VHD. User-only account access is not sufficient. The share can grant **Read** and **Write** permissions to the **Everyone** group to allow access to the VHD, but for security reasons, this is not recommended.

```
Uninstall-WindowsFeature –Name AD-Domain-Services,GPMC –VHD
C:\WS2012VHDs\Contoso.vhd –ComputerName ContosoDC1
```

See Also

Install or Uninstall Roles, Role Services, or Features

Windows Server Installation Options

How to Enable or Disable Windows Features

Deployment Image Servicing and Management (DISM) Overview

Add Servers to Server Manager

In Windows Server® 2012 R2 Windows Server® 2012, you can manage multiple remote servers by using a single Server Manager console. Servers that you want to manage by using Server Manager can be running Windows Server® 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003. Note that you cannot manage a newer release of Windows Server with an older release of Server Manager.

This topic describes how to add servers to the Server Manager server pool.

Notes

In our tests, Server Manager in Windows Server 2012 R2 and Windows Server 2012 can be used to manage up to 100 servers that are configured with a typical workload. The number of servers that you can manage by using a single Server Manager console can vary depending on the amount of data that you request from managed servers, and hardware and network resources available to the computer running Server Manager. As the amount of data you want to display approaches that computer's resource capacity, you can experience slow responses from Server Manager, and delays in the completion of refreshes. To help increase the number of servers that you can manage by using Server Manager, we recommend limiting the event data that Server Manager gets from your managed servers, by using settings in the **Configure Event Data** dialog box. Configure Event Data can be opened from the **Tasks** menu in the **Events** tile. If you

need to manage an enterprise-level number of servers in your organization, we recommend evaluating products in the Microsoft® System Center suite.

Server Manager can receive only online or offline status from servers that are running Windows Server 2003. Although you can use Server Manager to perform management tasks on servers that are running Windows Server 2008 R2 or Windows Server 2008, you cannot add roles and features to servers that are running Windows Server 2008 R2, Windows Server 2008 or Windows Server 2003.

Server Manager cannot be used to manage a newer release of the Windows Server operating system. Server Manager running on Windows Server 2012 or Windows 8 cannot be used to manage servers that are running Windows Server 2012 R2.

This topic contains the following sections.

- Add servers to manage
- Provide credentials with the Manage As command

Provide credentials with the Manage As command

As you add remote servers to Server Manager, some of the servers that you add might require different user account credentials to access or manage them. To specify credentials for a managed server that are different from those you use to log on to the computer on which you are running Server Manager, use the **Manage As** command after you add a server to Server Manager, which is accessible by right-clicking the entry for a managed server in the **Servers** tile of a role or group home page. Clicking **Manage As** opens the **Windows Security** dialog box, in which you can provide a user name that has access rights on the managed server, in one of the following formats.

- *User name*
- *User name*@example.domain.com
- *Domain**User name*

The **Windows Security** dialog box that is opened by the **Manage As** command cannot accept smart card credentials; providing smart card credentials through Server Manager is not supported. Credentials that you provide for a managed server by using the **Manage As** command are cached, and persist as long as you are managing the server by using the same computer on which you are currently running Server Manager, or as long as you do not overwrite them by specifying blank or different credentials for the same server. If you export your Server Manager settings to other computers, or configure your domain profile to be roaming to allow Server Manager settings to be used on other computers, **Manage As** credentials for servers in your server pool are not stored in the roaming profile. Server Manager users must add them on each computer from which they want to manage.

After you add servers to manage by following procedures in this topic, but before you use the **Manage As** command to specify alternate credentials that might be required to manage a server that you have added, the following manageability status errors can be displayed for the server:

- Kerberos target resolution error

- Kerberos authentication error
- Online - Access denied

Note
Roles and features that do not support the **Manage As** command include Remote Desktop Services (RDS) and IP Address Management (IPAM) Server. If you cannot manage the remote RDS or IPAM server by using the same credentials you are using on the computer on which you are running Server Manager, try adding the account you typically use to manage these remote servers to the Administrators group on the computer that is running Server Manager. Then, log on to the computer that is running Server Manager with the account you use to manage the remote server that is running RDS or IPAM.

Add servers to manage

You can add servers to Server Manager to manage by using any of three methods in the **Add Servers** dialog box.

- **Active Directory® Domain Services** Add servers to manage that Active Directory finds in the same domain as the local computer.
- **Domain Name System (DNS) entry** Search for servers to manage by computer name or IP address.
- **Import multiple servers** Specify multiple servers to import in a file that contains servers listed by computer name or IP address.

To add servers to the server pool

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** screen, click the Server Manager tile.
2. On the **Manage** menu, click **Add Servers**.
3. Do one of the following.
   - On the **Active Directory** tab, select servers that are in the current domain. Press **Ctrl** while selecting to select multiple servers. Click the right-arrow button to move selected servers to the **Selected** list.
   - On the **DNS** tab, type the first few characters of a computer name or IP address, and then press **Enter** or click **Search**. Select servers that you want to add, and then click the right-arrow button.
   - On the **Import** tab, browse for a text file that contains the DNS names or IP addresses of computers that you want to add, one name or IP address per line.
4. When you are finished adding servers, click **OK**.

Add and manage servers in workgroups

Although adding servers that are in workgroups to Server Manager might be successful, after they are added, the **Manageability** column of the **Servers** tile—on a role or group page that includes a workgroup server—can display **Credentials not valid** errors that occur while trying to connect to or collect data from the remote, workgroup server.

These or similar errors can occur in the following conditions.

- The managed server is in the same workgroup as the computer that is running Server Manager.
- The managed server is in a different workgroup from the computer that is running Server Manager.
- One of the computers is in a workgroup, while the other is in a domain.
- The computer that is running Server Manager is in a workgroup, and remote, managed servers are on a different subnet.
- Both computers are in domains, but there is no trust relationship between the two domains.
- Both computers are in domains, but there is only a one-way trust relationship between the two domains.
- The server you want to manage has been added by using its IP address.

To add remote workgroup servers to Server Manager

1. On the computer that is running Server Manager, add the workgroup server name to the **TrustedHosts** list. This is a requirement of NTLM authentication. To add a computer name to an existing list of trusted hosts, add the `Concatenate` parameter to the command. For example, to add the `Server01` computer to an existing list of trusted hosts, use the following command.

   ```
   Set-Item wsman:\localhost\Client\TrustedHosts Server01 -
   Concatenate -Force
   ```

2. Determine whether the workgroup server that you want to manage is in the same subnet as the computer on which you are running Server Manager.

   If the two computers are in the same subnet, or if the workgroup server's network profile is set to **Private** in the **Network and Sharing Center**, go on to the next step.

   If they are not in the same subnet, or if the workgroup server's network profile is not set to **Private**, on the workgroup server, change the inbound **Windows Remote Management (HTTP-In)** setting in Windows Firewall to explicitly allow connections from remote computers by adding the computer names on the **Computers** tab of the setting's **Properties** dialog box.

3.

   Security

   Running the cmdlet in this step overrides User Account Control (UAC) measures that prevent elevated processes from running on workgroup computers unless the built-in Administrator or the System account is running the processes. The

cmdlet lets members of the Administrators group manage the workgroup server without logging on as the built-in Administrator. Allowing additional users to manage the workgroup server can reduce its security; however, this is more secure than providing built-in Administrator account credentials to what might be multiple people who are managing the workgroup server.

To override UAC restrictions on running elevated processes on workgroup computers, create a registry entry called **LocalAccountTokenFilterPolicy** on the workgroup server by running the following cmdlet.

```
New-ItemProperty -Name LocalAccountTokenFilterPolicy -path
HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Syst
em -propertyType DWord -value 1
```

4. On the computer on which you are running Server Manager, open the **All Servers** page.
5. If the computer that is running Server Manager and the target workgroup server are in the same workgroup, skip to the last step. If the two computers are not in the same workgroup, right-click the target workgroup server in the **Servers** tile, and then click **Manage as**.
6. Log on to the workgroup server by using the built-in Administrator account for the workgroup server.
7. Verify that Server Manager is able to connect to and collect data from the workgroup server by refreshing the **All Servers** page, and then viewing the manageability status for the workgroup server.

To add remote servers when Server Manager is running on a workgroup computer

1. On the computer that is running Server Manager, add remote servers to the local computer's **TrustedHosts** list in a Windows PowerShell session. To add a computer name to an existing list of trusted hosts, add the `Concatenate` parameter to the command. For example, to add the `Server01` computer to an existing list of trusted hosts, use the following command.

   ```
   Set-Item wsman:\localhost\Client\TrustedHosts Server01 -
   Concatenate -Force
   ```

2. Determine whether the server that you want to manage is in the same subnet as the workgroup computer on which you are running Server Manager.

   If the two computers are in the same subnet, or if the workgroup computer's network profile is set to

   Private in the **Network and Sharing Center**, go on to the next step.

   If they are not in the same subnet, or if the workgroup computer's network profile is not set to **Private**, on the workgroup computer that is running Server Manager, change the inbound **Windows Remote Management (HTTP-In)** setting in Windows Firewall to explicitly allow connections from remote computers by adding the computer names on the **Computers** tab of the setting's **Properties** dialog box.
3. On the computer on which you are running Server Manager, open the **All Servers** page.

4. Verify that Server Manager is able to connect to and collect data from the remote server by refreshing the **All Servers** page, and then viewing the manageability status for the remote server. If the **Servers** tile still displays a manageability error for the remote server, go on to the next step.
5. Log off of the computer on which you are running Server Manager, and then log on again by using the built-in Administrator account. Repeat the preceding step, to verify that Server Manager is able to connect to and collect data from the remote server.

If you have followed the procedures in this section, and you continue to have problems managing workgroup computers, or managing other computers from workgroup computers, see about_Remote_Troubleshooting on the Microsoft website.

Add and manage servers in clusters

You can use Server Manager to manage servers that are in failover clusters (also called server clusters or MSCS). Servers that are in failover clusters—whether the cluster nodes are physical or virtual—have some unique behaviors and management limitations in Server Manager.

- Both physical and virtual servers in clusters are automatically added to Server Manager when one server in the cluster is added to Server Manager. Similarly, when you remove a clustered server from Server Manager, you are prompted to remove other servers in the cluster.
- Server Manager does not display data for clustered virtual servers, because the data is dynamic, and is identical to data for the server on which the virtual clustered node is hosted. You can select the server that is hosting the virtual server to view its data.
- If you add a server to Server Manager by using the server's virtual cluster object name, the virtual object name is displayed in Server Manager instead of the physical server name (expected).
- You cannot install roles and features on a clustered virtual server.

See Also

Manage Multiple, Remote Servers with Server Manager

Create and Manage Server Groups

View and Configure Performance, Event, and Service Data

This topic describes how to view and configure the event log entries, performance counters, and service alerts that are displayed for local and remote servers in Server Manager.

Event, service, and performance log data is displayed in two places in the Server Manager console in Windows Server 2012 R2 and Windows Server® 2012.

- On the dashboard, you can click the **Events**, **Performance**, and **Services** rows of thumbnails to configure event, performance, and service log data that you want to see for

roles, the entire Server Manager server pool, user-created groups of servers, and the local server. Clicking the hypertext rows opens **Detail View** dialog boxes that let you specify the data about which you want to be alerted in the dashboard. After you configure event, service, and performance log data that you want to be highlighted in the dashboard thumbnails, log entries that match the criteria you have specified are listed at the bottom of the **Detail View** dialog boxes.

- **Events**, **Services**, and **Performance** tiles are part of role and group home pages. Commands on the **Tasks** menu of these tiles let you specify the data that you want collected from managed servers. The tiles include filters and queries to further limit the log entries that are displayed in the tile, if desired.

This topic contains the following sections.

- What are thumbnails?
- View and configure events
- View and configure performance log data
- Manage services and configure service alerts
- View and copy event, service, or performance entries

What are thumbnails?

Thumbnails are displayed on the Server Manager dashboard for each role (a role's thumbnail reflects data collected about all of the servers in the Server Manager pool that are running the role), for each server group, for the **All Servers** group (all of the servers in the Server Manager pool), and for the local server. After Server Manager gets data from managed servers, thumbnails are automatically created for roles that are running on servers in the server pool.

If the Server Manager console is running on a client computer as part of Remote Server Administration Tools, there is no **Local Server** thumbnail.

The thumbnail displays a quick view of the status and manageability of roles, servers, and server groups. The thumbnail heading row changes color—and highlighted numbers are displayed in the left margin—when events, performance counters, Best Practices Analyzer results, services, or general manageability issues meet criteria that you configure in the **Detail View** dialog boxes opened by clicking thumbnail rows. The following table describes the data displayed in thumbnails.

Thumbnail Row	Description
Manageability	The manageability of a server includes several measures: whether the server is online or offline, whether it is accessible and reporting data to Server Manager, whether the user who is logged on to the local computer has adequate user rights to access or manage the remote server, whether the remote server is running all of the software that is required to

Thumbnail Row	Description
	manage it remotely, or whether the server is configured in a way that allows it to be queried and managed by using Server Manager. The only manageability data that Server Manager can collect from a server that is running Windows Server 2003 is whether the server is online or offline. For detailed information about manageability status errors and how to resolve them, see the Server Manager Troubleshooting Guide.
Events	You can configure the **Events** row of a thumbnail to display alerts when events are logged that match severity levels, sources, time periods, servers, or event IDs that you specify. View details about events, and change the alerts you want to see by clicking the **Events** row, and opening the **Events Detail View** dialog box for the role or server group.
Services	You can configure the **Services** row to display alerts when services are found in a role or server group that match startup types, service status, service names, and servers that you specify in the **Services Detail View** dialog box. After a server has been added to the Server Manager server pool, service alerts about the Shell Hardware Detection service can be displayed if there are no users logged on to the managed server. This occurs because the Shell Hardware Detection service runs only when users are logged on to the managed server, or connected to a Remote Desktop session on the managed server. To avoid seeing Shell Hardware Detection service alerts for this case, click **Services** in the thumbnails for server groups, including the **All Servers** group. In the **Services Detail View** dialog box, on the **Services** drop-down list, clear the check box for **Shell Hardware Detection**, and then click **OK**.
Performance	You can configure the **Performance** row to

Thumbnail Row	Description
	display alerts for a role or server group when performance alerts occur that match resource types, servers, or time periods that you specify in the **Performance Detail View** dialog box. By default, performance counters are turned off. Managed servers that are running operating systems newer than Windows Server 2003, and for which performance counters have not been started, typically show manageability status errors of **Online – Performance counters not started** in the **Servers** tile of role or group pages. To turn performance counters on for managed servers, on the **All Servers** page, right-click entries in the **Performance** tile that show a **Counter Status** value of **Off**, and then click **Start Performance Counters**. You can also start performance counters by right-clicking entries for servers in the **Servers** tile of role or group pages, and then clicking **Start Performance Counters**.
BPA results	You can configure the **BPA results** row to display alerts for a role or server group when BPA scan results are found that match severity levels, servers, or BPA categories that you specify in the **BPA Results Detail View** dialog box.

View and configure events

In this section, learn how to configure what event log data is collected from servers in the Server Manager server pool, and which events you want highlighted in thumbnails.

Note

The events about which you are alerted in thumbnails are a subset of the total events that you instruct Server Manager to collect from managed servers. Although changing event criteria in the **Configure Event Data** dialog box in **Events** tiles can change the numbers of alerts you see on the Server Manager dashboard, changing the event alert criteria in thumbnails has no effect on the event log data that is collected from managed servers.

To configure the events collected from managed servers

1. In the Server Manager console, open any page except the dashboard. You can configure the events that you want collected from managed servers in the **Events** tile on role, server group, or local server pages.
2. On the **Tasks** menu of the **Events** tile, click **Configure Event Data**.
3. Select the event severity levels that you want to be collected from the servers in the selected group. By default, **Critical**, **Error**, and **Warning** severity levels are selected.
4. Specify a time period during which the events occur. The default age for events is 24 hours.
5. Select the event log files from which you want events to be collected. The defaults are **Application**, **Setup**, and **System**.
6. To save your changes, click **OK** to close the **Configure Event Data** dialog box. Event data automatically refreshes when your changes are saved.

To configure the events highlighted in thumbnails

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** screen, click the Server Manager tile.
2. On the dashboard page, in a thumbnail in the **Roles and Server Groups** tile, click the **Events** row.
3. In the **Events Detail View** dialog box, add a severity level to the events that you want displayed. By default, thumbnails only display alert highlighting for critical events. Note that the number of events displayed in the **Detail View** dialog box increases when you add a severity level about which you want to be alerted.
4. In the **Event sources** field, select the event sources about which you want to be alerted. The default is **All**.
5. If this thumbnail is for a role that is installed on multiple servers, or a group of multiple servers, you can select the servers for which you want event alerts in the **Servers** drop-down list.
6. In the **Time period** field, specify a time period up to 1440 minutes, 24 hours, or 1 day.
7. In the **Event IDs** field, type the event ID numbers of specific events about which you want to be alerted. You can type a range of event IDs separated by a dash (**-**), and exclude event IDs from the range by typing the dash before the event ID or range of event IDs that you want to exclude. For example, the value **1,3,5-99,-76** means that alerts are raised for event IDs 1 and 3, and all events with IDs between 5 and 99 except for event ID 76.
8. As you change the criteria for which alerts are displayed, the number of event alerts that are displayed in the results pane at the bottom of the dialog box might change. Select entries in the list and click **Hide Alerts** to prevent them from affecting the alert count that is displayed in the source thumbnail. Press and hold **Ctrl** as you select alerts to select multiple alerts at one time. You can do this for alerts that match your event alerting

criteria, but that you do not need to see.

9. Click **Show All** to return hidden alerts to the list.
10. Click **OK** to save your changes, close the **Detail View** dialog box, and view the event alert changes in the source thumbnail.

View and configure performance log data

In this section, learn how to configure what performance log data is collected from servers in the Server Manager server pool, and which performance counter alerts you want highlighted in thumbnails.

By default, performance counters are turned off. Managed servers that are running operating systems newer than Windows Server 2003, and for which performance counters have not been started, typically show manageability status errors of **Online – Performance counters not started** in the **Servers** tile of role or group pages. To turn performance counters on for managed servers, on the **All Servers** page, right-click entries in the **Performance** tile that show a **Counter Status** value of **Off**, and then click **Start Performance Counters**. You can also start performance counters by right-clicking entries for servers in the **Servers** tile of role or group pages, and then clicking **Start Performance Counters**.

Notes

The performance alerts you view in thumbnails are a subset of the total performance counter data that you instruct Server Manager to collect from managed servers. Although changing performance alert criteria in the **Configure Performance Alerts** dialog box in **Performance** tiles can change the numbers of alerts you see on the Server Manager dashboard, changing the performance alert criteria in thumbnails has no effect on the performance log data that is collected from managed servers.

Because of this, the maximum age of performance data that you can display in thumbnails cannot be greater than the maximum graph display period that is configured in the **Configure Performance Alerts** dialog box. For example, if the **Graph display period** value in **Configure Performance Alerts** is **1 day**, the maximum value for the **Time period** field in a **Performance Detail View** dialog box that you have opened from the Server Manager dashboard can be **1 day**, **24 hours**, or **1,440 minutes**.

To configure the performance log data collected from managed servers

1. In the Server Manager console, open any page except the dashboard. You can configure the performance data that you want collected from managed servers in the **Performance** tile on role, server group, or local server pages.
2. To collect performance log data from managed servers, performance counters must be turned on. If performance counters are turned off, right-click an entry in the **Performance** tile list, and then click **Start Performance Counters**. Performance counter data collection can require some time, depending on the number of servers from which data is collected, and available network bandwidth. View the status in the **Counter Status** column.

3. On the **Tasks** menu of the **Performance** tile, click **Configure Performance Alerts**.
4. For the servers in the selected group, or that are running the selected role, specify at what percent of CPU usage you want performance counter alerts collected by Server Manager. The default is 85%.
5. Specify the remaining available memory, in megabytes, that servers must have before you want performance counter alerts collected. The default is 2 MB.
6. Specify a time period that is displayed by the graphs for the resources **CPU Usage** and **Available Memory** in the **Performance** tile on the selected page. The default period is one day. Click **Save**.

   Note that the number of performance alerts in the **Performance** tile, and the mapping of the alerts over time as displayed by the graph, can change after you click **Save**.

   Note
   For virtual machines that have Dynamic Memory turned on, increasing the performance alerts threshold can result in false positive alerts.
7. To refresh the list of performance alerts that are collected from your servers, on the **Tasks** menu, click **Refresh**.

To configure the performance alerts highlighted in thumbnails

1. On the dashboard page, in a thumbnail in the **Roles and Server Groups** tile, click the **Performance** row.
2. In the **Performance Detail View** dialog box, select or clear check boxes for resource performance thresholds about which you want to be alerted in the **Resource type** field. Note that the number of performance alerts displayed in the **Detail View** dialog box can increase when you add a resource performance threshold about which you want to be alerted.
3. If this thumbnail is for a role that is installed on multiple servers, or a group of multiple servers, you can select the servers for which you want performance alerts in the **Servers** drop-down list.
4. In the **Time period** field, specify a time period up to 1440 minutes, 24 hours, or 1 day.
5. As you change the criteria for which alerts are displayed, the number of alerts that are displayed in the results pane at the bottom of the dialog box might change. Click **Hide Alerts** to hide all alerts older than the current time, and prevent them from affecting the alert count that is displayed in the source thumbnail.
6. Click **Show All** to return hidden alerts to the list.
7. Click **OK** to save your changes, close the **Detail View** dialog box, and view the performance alert changes in the source thumbnail.

To view the properties of performance alerts

1. Do one of the following.
   - On the dashboard page, in a thumbnail in the **Roles and Server Groups** tile, click the **Performance** row.

- Open a role or group home page, and locate the **Performance** tile for the role or group.

2. Double-click a performance alert in the list to view its properties. Alternatively, you can right-click a performance alert, and then click **View Properties**.
3. In the **Performance Alert Properties** dialog box, select log entries to view information about the processes that are associated with the entry in the **Processes** area.
4. When you are finished viewing performance alert properties, close the dialog box.

Analyze performance data and solve problems

For more information about analyzing performance counter data that you view in Server Manager, and solving performance problems on managed servers, see the following resources.

- Analyzing performance data
- Solving performance problems

For more information about advanced performance monitoring and analysis tools that are available for Windows Server 2012 R2 and Windows Server 2012, including Server Performance Advisor 3.0, see Performance Tools for Windows 8.

Manage services and configure service alerts

In this section, learn how to start, stop, restart, pause, or resume services that are displayed in the **Services** tile on role and server group pages in Server Manager. You can also configure the services about which you are alerted in thumbnails on the Server Manager dashboard.

Note

You cannot change the start type for services, service dependencies, recovery options, or other service properties in the Services tile in Server Manager. To change service properties other than the service status, open the **Services** snap-in. A shortcut to open the **Services** snap-in is available on the **Tools** menu in Server Manager.

To start, stop, restart, pause, or resume a service

1. In the Server Manager console, open any page except the dashboard (in other words, any role or group home page).
2. In the **Services** tile for the role or group, right-click a service.
3. In the context menu, click the action that you want to perform on the service. If the service is stopped, the only action you can perform is to start the service. Similarly, if the service is paused, the only action you can perform is to resume the service.
4. Note that when you start, stop, restart, pause, or resume a service, the value of the **Status** column changes for the service in the **Services** tile.

To configure service alerts highlighted in thumbnails

1. On the dashboard page, in a thumbnail in the **Roles and Server Groups** tile, click the

Services row.

2. In the **Services Detail View** dialog box, select the startup types for services about which you want to be alerted. By default, **Automatic (Delayed Start)** and **Automatic** are selected.
3. Select the service statuses about which you want to be alerted. By default, **All** is selected.
4. Select services about which you want to be alerted. By default, **All** is selected.
5. Select the servers associated with the role or group for which you want to receive alerts about services. By default, **All** is selected.
6. As you change the criteria for which alerts are displayed, the number of alerts that are displayed in the results pane at the bottom of the dialog box might change. Click **Hide Alerts** to hide all alerts older than the current time, and prevent them from affecting the alert count that is displayed in the source thumbnail.
7. Click **Show All** to return hidden alerts to the list.
8. Click **OK** to save your changes, close the **Detail View** dialog box, and view the service alert changes in the source thumbnail.

View and copy event, service, or performance entries

You can copy event, service, or performance entry properties in both the **Detail View** dialog boxes and the **Events** and **Performance** tiles for a role or group. Right-click an event or performance entry, and then click **Copy**.

The **Events** tile also lets you preview event properties in the bottom half of the tile by selecting an event in the list. To copy the properties shown in the preview, right click the preview pane, and then click **Copy**.

See Also

Manage Multiple, Remote Servers with Server Manager

Filter, Sort, and Query Data in Server Manager Tiles

View Task Details and Notifications

In Server Manager in Windows Server® 2012 R2 or Windows Server® 2012, when you perform management tasks such as adding roles and features, starting services, refreshing data that is displayed in the Server Manager console, or creating a custom group of servers, a notification is displayed in the **Notifications** area of the Server Manager console header. Notifications, and the **Task Details** dialog box that you can open from the **Notifications** menu by clicking the flag icon, display the status of user tasks or requests, show you if a task failed, and help you troubleshoot problems by pointing to detailed error messages about task failures.

The Notifications area

The **Notifications** area in the Server Manager menu bar, marked by a flag icon, displays the results of tasks that you start in Server Manager. Notifications inform you whether tasks that you started in Server Manager were successful or failed. When notifications are available for you to view, the number of available notifications is displayed next to the flag icon. If a task failed, could only be partially completed (for example, if it could not be completed on all of the remote servers on which you wanted to perform the task), or completed with warnings, the **Notifications** flag becomes red. Tasks for which notifications are displayed include the following.

- Manually refreshing the data displayed in Server Manager (notifications are displayed for automatic refreshes only if the refreshes fail)
- Starting or stopping services
- Installing or uninstalling roles, role services, and features
- Starting a Best Practices Analyzer (BPA) scan
- Adding remote servers to manage (notifications are displayed for failures to contact or refresh the data displayed for remote servers)

Items in the **Notifications** menu show a progress bar, a brief description of the task, the name of the target server for the task (or one of the target servers, if multiple target servers were selected), a link to a related control or dialog box if applicable, and a **Tasks** menu. The **Tasks** menu shows commands that apply to the active notification (the notification over which your mouse cursor is hovered). For example, if you stop a service, you can click **Restart** on the **Tasks** menu in the notification to restart the service.

Notifications are particularly useful for installing or uninstalling roles, role services, and features. For example, if you start a feature installation on a remote server, you can close the Add Roles and Features Wizard while the installation is still in progress, but the active task remains in the **Notifications** list. The **Notifications** item shows a progress bar for the installation, and lets you reopen the Add Roles and Features Wizard, if needed, by clicking **Add Roles and Features Wizard**. The items in this list notify you if an installation failed, or if additional configuration steps are required to finish deploying the feature.

Notifications also play an important part in troubleshooting problems with tasks or processes in Server Manager. For more information about how to use messages in the **Notifications** area and **Task Details** dialog box to troubleshoot unsuccessful tasks or processes, see the Server Manager Troubleshooting Guide.

To delete a notification that you no longer want to see from the **Notifications** list, hover your mouse cursor over the notification, and then click **Remove Task** (**X**).

Viewing and troubleshooting tasks by using Task Details

The **Task Details** command at the bottom of the **Notifications** menu opens the **Task Details** dialog box, which provides full descriptions of task events (starting, stopping, warnings, successes, or failures). Like other list controls in Server Manager, such as **Events**, **Services**, and

Best Practices Analyzer tiles, you can filter and create queries to run on the tasks that are displayed in the **Task Details** dialog box. (For more information about filtering and creating queries on list controls, see Filter, Sort, and Query Data in Server Manager Tiles.) In the top pane, you can review notifications as they have been displayed in the **Notifications** menu, and see how many notifications have been generated about the same task. Selecting a notification in the top pane displays full details about the notification in the bottom pane.

The bottom pane is particularly useful for troubleshooting failed tasks. If Server Manager cannot connect to or get data for a server that is a member of the server pool, entries in this pane often contain detailed messages, including the full text of underlying Windows Remote Management (WinRM), networking, or security problems that prevent Server Manager from communicating with a target server.

See Also

Filter, Sort, and Query Data in Server Manager Tiles

Server Manager Troubleshooting Guide

Run Best Practices Analyzer Scans and Manage Scan Results

In Windows management, *best practices* are guidelines that are considered the ideal way, under typical circumstances, to configure a server as defined by experts. For example, it is considered a best practice for most server applications to keep open only those ports required for the applications to communicate with other networked computers, and block unused ports. Although best practice violations, even crucial ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Best Practices Analyzer (BPA) is a server management tool that is available in Windows Server® 2012 R2, Windows Server 2012, and Windows Server 2008 R2. BPA can help administrators reduce best practice violations by scanning roles that are installed on managed servers that are running Windows Server 2012 or Windows Server 2008 R2, and reporting best practice violations to the administrator.

You can run Best Practices Analyzer (BPA) scans either from Server Manager, by using the BPA GUI, or by using cmdlets in Windows PowerShell®. Starting with Windows Server 2012, you can scan one role or multiple roles at one time, on multiple servers, whether you use the Best Practices Analyzer tile in the Server Manager console or Windows PowerShell cmdlets to run scans. You can also instruct BPA to exclude or ignore scan results that you do not want to see.

This topic contains the following sections.

- Find BPA
- How BPA works

- Performing Best Practices Analyzer scans on roles
- Manage scan results

Find BPA

You can find the Best Practices Analyzer tile on role and server group pages of Server Manager in Windows Server 2012 R2 and Windows Server 2012, or you can open a Windows PowerShell session with elevated user rights to run Best Practices Analyzer cmdlets.

How BPA works

BPA works by measuring a role's compliance with best practice rules in eight different categories of effectiveness, trustworthiness, and reliability. Results of measurements can be any of the three severity levels described in the following table.

Severity level	Description
Error	Error results are returned when a role does not satisfy the conditions of a best practice rule, and functionality problems can be expected.
Information	Information results are returned when a role satisfies the conditions of a best practice rule.
Warning	Warning results are returned if the results of noncompliance can cause problems if changes are not made. The application might be compliant as operating currently, but may not satisfy the conditions of a rule if changes are not made to its configuration or policy settings. For example, a scan of Remote Desktop Services might show a warning result if a license server is unavailable to the role, because even if no remote connections are active at the time of the scan, not having the license server prevents new remote connections from obtaining valid client access licenses.

Rule categories

The following table describes the best practice rules categories against which roles are measured during a Best Practices Analyzer scan.

Category Name	Description
Security	Security rules are applied to measure a role's relative risk for exposure to threats such as unauthorized or malicious users, or loss or theft of confidential or proprietary data.
Performance	Performance rules are applied to measure a role's ability to process requests and perform its prescribed duties in the enterprise within expected periods of time given the role's workload.
Configuration	Configuration rules are applied to identify role settings that might require modification for the role to perform optimally. Configuration rules can help prevent conflicts in settings that can result in error messages or prevent the role from performing its prescribed duties in an enterprise.
Policy	Policy rules are applied to identify Group Policy or Windows Registry settings that might require modification for a role to operate optimally and securely.
Operation	Operation rules are applied to identify possible failures of a role to perform prescribed tasks in the enterprise.
Predeployment	Predeployment rules are applied before an installed role is deployed in the enterprise. They let administrators evaluate, before the role is used in production, whether best practices were satisfied.
Postdeployment	Postdeployment rules are applied after all required services have started for a role, and after the role is running in the enterprise.
Prerequisites	Prerequisite rules explain configuration settings, policy settings, and features that are required for a role before BPA can apply specific rules from other categories. A prerequisite in scan results indicates that an incorrect setting, a missing program, an incorrectly enabled or disabled policy, a registry

Category Name	Description
	key setting, or other configuration has prevented BPA from applying one or more rules during a scan. A prerequisite result does not imply compliance or noncompliance. It means that a rule could not be applied, and is not therefore part of the scan results.

Performing Best Practices Analyzer scans on roles

You can perform BPA scans on roles by using either the BPA GUI in Server Manager, or by using Windows PowerShell cmdlets.

In Windows Server 2012 R2 and Windows Server 2012, some roles prompt you to specify additional parameters, such as the names of specific servers or shares that are running parts of the role, or the IDs of submodels, before starting a BPA scan. For BPA scans on models that require you to specify additional parameters, use the BPA cmdlets; the BPA GUI cannot accept additional parameters such as submodel IDs. For example, the submodel ID **FSRM** represents the File Services BPA submodel for File Server Resource Manager, a role service of File and Storage Services. To run a scan only on the File Server Resource Manager role service, run a BPA scan by using Windows PowerShell cmdlets, and add the parameter `SubModelId` to your cmdlet.

Although you cannot pass additional parameters to a scan that you start in the BPA GUI, the BPA tile in Server Manager displays results for the most recent BPA scan, regardless of how the scan was started.

- Scanning roles by using the BPA GUI
- Scanning roles by using Windows PowerShell cmdlets

Scanning roles by using the BPA GUI

Follow these steps to scan one or more roles in the BPA GUI.

To scan roles by using the BPA GUI

1. Do one of the following to open Server Manager if it is not already open.
   - On the Windows taskbar, click the Server Manager button.
   - On the **Start** screen, click the Server Manager tile.
2. In the navigation pane, open a role or group page.

   Running BPA scans from a role or group page scans all roles that are installed on servers in that group.

3. On the **Tasks** menu of the **Best Practices Analyzer** tile, click **Start BPA Scan**.
4. Depending on the number of rules that are evaluated for the role or group you selected, the BPA scan can require a few minutes to finish.

Scanning roles by using Windows PowerShell cmdlets

Use the following procedures to scan one or more roles by using Windows PowerShell cmdlets.

Note
The procedures in this section do not show all BPA cmdlets and parameters. For more information about BPA operations in Windows PowerShell, in your Windows PowerShell session, enter **Get-Help** *BPACmdlet* **-full**, where *BPACmdlet* can be one of the following values. You can also find BPA cmdlet Help topics on the Windows Server TechCenter.

- **Get-BPAModel**
- **Invoke-BPAModel**
- **Get-BPAResult**
- **Set-BPAResult**

To scan a single role by using Windows PowerShell cmdlets

1. Do one of the following to run Windows PowerShell with elevated user rights.
   - To run Windows PowerShell as an administrator from the **Start** screen, right-click the **Windows PowerShell** tile in the **Apps** results, and then on the app bar, click **Run as administrator**.
   - To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.
2. Starting in Windows PowerShell 3.0, cmdlet modules are automatically imported into your Windows PowerShell session the first time a cmdlet from the module is used. There is no need to import or load the BPA cmdlet module.
3. Find the model IDs of all roles for which BPA scans can be performed by entering the **Get-BPAModel** cmdlet, as shown in the following example.

   Get-BPAModel
4. In the results of step 3, locate the model IDs of the roles on which you want to run a BPA scan.
5. Enter either of the following commands to start the BPA scan for a specific role. For multiple roles, separate model IDs with commas.

   Invoke-BPAModel -ModelId *ModelID_from_Step3*

   Invoke-BPAModel *ModelID_from_Step3*

   Example: `Invoke-BPAModel -ModelId Microsoft/Windows/DNSServer,Microsoft/Windows/FileServices`

   Note
   The model ID includes the entire path that is shown in the **Id** column; for

example, **Microsoft/Windows/Hyper-V**.

You can also start a scan on a specific role from the results of step 3 by piping the results of the `Get-BPAModel` cmdlet into the `Invoke-BPAModel` cmdlet as shown in the following example.

Get-BPAModel *Model_ID* **| Invoke-BPAModel**

Running this cmdlet without specifying a model ID pipes all models that are returned by the `Get-BPAModel` cmdlet into the `Invoke-BPAModel` cmdlet, starting scans on all models that are available on servers that have been added to the Server Manager server pool.

To scan all roles by using Windows PowerShell cmdlets

1. Open a Windows PowerShell session with elevated user rights, if one is not already open. See the preceding procedure for instructions.
2. Pipe all roles for which BPA scans can be performed into the `Invoke-BPAModel` cmdlet to start scans.

   Get-BPAModel | Invoke-BPAModel
3. When the scan is complete, Windows PowerShell returns results similar to the following, for each role that was scanned.

```
ModelId                 :  Microsoft/Windows/FileServices
SubmodelId              :
Success                 :  True
ScanTime                :  1/01/2012  12:18:40 PM
ScanTimeUtcOffset       :  -08:00:00
Detail                  :  {server_name1, server_name2}
```

Manage scan results

After a BPA scan is completed in the GUI, you can view scan results in the BPA tile. When you select a result in the tile, a preview pane in the tile displays result properties, including an indication of whether the role is compliant with the associated best practice. If a result is not compliant, and you want to know how to resolve the problems described in the result properties, hyperlinks in error and warning result properties open detailed resolution Help topics on the Windows Server TechCenter.

Note

BPA scan results are not automatically saved or archived. Running a new scan on a model or submodel overwrites the results of the last scan. To save BPA scan results for archiving, printing, or sending to others, see To view or save BPA results from Windows PowerShell sessions in different formats in this section.

Exclude and include BPA results

If you do not need to see some BPA results, such as results that occur frequently in your BPA scans but require no resolution, you can exclude the results, by using either the BPA GUI or BPA cmdlets in Windows PowerShell. Results can be included again at any time.

Note

When you exclude results, they are also excluded from view on managed servers. Other administrators cannot see excluded results on managed servers. To exclude results from view in a local Server Manager console only, create a custom query instead of using the **Exclude Result** command.

Exclude scan results

The **Exclude** setting is persistent; results that you exclude remain excluded in future scans of the same model on the same computer, unless they are included again.

You can exclude scan results by using the `Set-BPAResult` cmdlet with the `Exclude` parameter. As in the Best Practices Analyzer tile in Server Manager, you can exclude individual result objects, or you can also exclude a set of results whose fields (category, title, and severity, for example) are equal to or contain specified values. For example, you can exclude all **Performance** results from a set of scan results for a model.

Note

You must run at least one BPA scan on a model before you can use procedures in this section.

To exclude scan results by using the GUI

1. Open a role or server group page in Server Manager.
2. In the Best Practices Analyzer tile for the role or server group, right-click a result in the list, and then click **Exclude Result**.

   The result is no longer displayed in the list of results.
3. To view excluded results in the GUI, run the built-in **Excluded results** query. Click **Saved Search Queries**, and then click **Excluded results**.

   After running the **Excluded results** query, note that the tile subheading text, a description of the results that are displayed in the list, changes to **Excluded results**. Only excluded results are displayed in the list.

To exclude scan results by using Windows PowerShell cmdlets

1. Open a Windows PowerShell session with elevated user rights.
2. Exclude specific results from a model scan by running the following command.

   Get-BPAResult -ModelId *Model ID* | **Where { $_.**<*Field Name*> **-eq "***Value***"} | Set-BPAResult -Exclude $true**

   The preceding command retrieves BPA scan result items for the model ID that is

represented by *Model ID*.

The second section of the command filters the results of the `Get-BPAResult` cmdlet to retrieve only those scan results for which the value for a result field, represented by *Field Name*, matches the text in quotation marks.

The final section of the command, following the second pipe character, excludes the results that are filtered by the previous section of the cmdlet.

Example: `Get-BPAResult -Microsoft/Windows/FileServices | Where { $_.Severity -eq "Information"} | Set-BPAResult -Exclude $true`

Include scan results

When you want to view scan results that were excluded, you can include those scan results. The **Include** setting is persistent; included results remain included in future scans of the same model on the same computer.

To include scan results by using the GUI

1. Open a role or server group page in Server Manager.
2. In the Best Practices Analyzer tile for the role or server group, right-click an excluded result in the **Excluded results** query list, and then click **Include Result**.

   The result is no longer displayed in the list of excluded results. Clear the query by clicking **Clear All** to view the included result in the list of all included results.

To include scan results by using Windows PowerShell cmdlets

1. Open a Windows PowerShell session with elevated user rights.
2. Include specific results from a model scan by typing the following command, and then pressing **Enter**.

   Get-BPAResult -ModelId *Model Id* **| Where { $_.**<*Field Name*> **-eq "***Value***" } | Set-BPAResult** -**Exclude $false**

   The preceding command retrieves BPA scan result items for the model represented by *Model Id*.

   The second part of the command, after the first pipe character (**|**) filters the results of the **Get-BPAResult** cmdlet to retrieve only those scan results for which the value of the result field, represented by *Field Name*, matches the text in quotation marks.

   The final part of the command, after the second pipe character, includes results that are filtered by the second part of the cmdlet, by setting the value of the **-Exclude** parameter to **false**.

   Example: `Get-BPAResult -Microsoft/Windows/FileServices | Where { $_.Severity -eq "Information"} | Set-BPAResult -Exclude $false`

View and export BPA scan results in Windows PowerShell

To view and manage scan results by using Windows PowerShell cmdlets, see the following procedures. Before you can use any of the following procedures, run at least one BPA scan on at least one model or submodel.

To view results of the most recent scan of a role by using Windows PowerShell

1. Open a Windows PowerShell session with elevated user rights.
2. Get the results of the most recent scan for a specified model ID. Type the following, in which the model is represented by *Model ID*, and then press **Enter**. You can get results for multiple model IDs by separating the model IDs with commas.

   Get-BPAResult *Model ID*

   Example: `Get-BPAResult Microsoft/Windows/DNSServer,Microsoft/Windows/FileServices`

   If you scanned a submodel of a model, such as a role service, get the results for only that submodel by including the submodel ID in the cmdlet.

   Example: `Get-BPAResult Microsoft/Windows/FileServices -SubmodelID FSRM`

To view or save BPA results from Windows PowerShell sessions in different formats

- In Windows PowerShell, each BPA result resembles the following.

```
ResultNumber     :  14
ResultId         :  1557706192
ModelId          :  Microsoft/Windows/FileServices
SubModelId       :  FSRM
RuleId           :  16
ComputerName     :  server_name1, server_name2
Context          :  FileServices
Source           :  server_name1
Severity         :  Information
Category         :  Configuration
Title            :  Access Denied Remediation requires remote
management be enabled on this server
Problem          :
Impact           :
Resolution       :
Compliance       :  The File Server Best Practices Analyzer
scan has determined that you are in compliance with this best
practice.
```

```
Help             :
Excluded         :  False
```

Do one of the following.

- To format BPA results in a table, run the following cmdlet, adding the result properties that you want to see from the preceding example.

  Get-BPAResult *Model ID* **| Format-Table -Property** *property1,property2,property3…*

  Example: `Get-BPAResult Microsoft/Windows/FileServices | Format-Table -Property ModelId,SubModelId,ComputerName,Source,Severity,Category,Title,Problem,Impact,Resolution,Compliance,Help`

- To format BPA results in a GUI-based grid viewer, with a text-string filter, and column headings that can be clicked to sort results, run the following cmdlet.

  Get-BPAResult *Model ID* **| OGV**

- To export BPA results to an HTML file that can be archived or sent to email recipients, run the following cmdlet, where *Path* represents the path and file name to which you want to save the HTML results.

  Get-BPAResult *Model ID* **| ConvertTo-Html | Set-Content** *Path*

  Example: `Get-BPAResult Microsoft/Windows/FileServices | ConvertTo-Html | Set-Content C:\BPAResults\FileServices.htm`

- To export BPA results to a comma-separated values (CSV) text file, run the following cmdlet, where *Path* represents the path and text file name to which you want to save the CSV results. CSV results can be imported into Microsoft® Excel, or other programs that display data in spreadsheets or grids.

  Get-BPAResult *Model ID* **| Export-CSV** *Path*

  Example: `Get-BPAResult Microsoft/Windows/FileServices | Export-CSV C:\BPAResults\FileServices.txt`

See Also

Best Practices Analyzer resolution content on the Windows Server TechCenter

Filter, Sort, and Query Data in Server Manager Tiles

Manage Multiple, Remote Servers with Server Manager

Create and Manage Server Groups

This topic describes how to create custom, user-defined groups of servers in Server Manager in Windows Server® 2012 R2 and Windows Server® 2012.

Server groups

Servers that you add to the server pool are displayed on the **All Servers** page in Server Manager. You can create custom groups of servers that you have added. Server groups allow you to view and manage a smaller subset of your server pool as a logical unit; for example, you can create a group called **Accounting Servers** for all servers in your organization's Accounting department, or a group called **Chicago** for all servers that are geographically located in Chicago. After you create a server group, the group's home page in Server Manager displays information about events, services, performance counters, Best Practices Analyzer results, and installed roles and features for the group as a whole.

Servers can be a member of more than one group.

To create a new server group

1. On the **Manage** menu, click **Create Server Group**.
2. In the **Server group name** text box, type a friendly name for your server group, such as **Accounting Servers**.
3. Add servers to the **Selected** list from the server pool, or add other servers to the group by using the **Active Directory**, **DNS**, or **Import** tabs. For more information about how to use these tabs, see Add Servers to Server Manager in this topic.
4. When you have finished adding servers to the group, click **OK**. The new group is displayed in the Server Manager navigation pane under the **All Servers** group.

To edit an existing server group

1. Do one of the following.
   - In the Server Manager navigation pane, right-click a server group, and then click **Edit Server Group**.
   - On the home page for the server group, open the **Tasks** menu on the **Servers** tile, and then click **Edit Server Group**.
2. Change the group name, or add or remove servers from the group.

   Note
   Removing servers from a server group does not remove servers from Server Manager. Servers that you remove from a group remain in the **All Servers** group, in the server pool.
3. When you are finished with changes to the group, click **OK**.

To delete an existing server group

1. Do one of the following.
   - In the Server Manager navigation pane, right-click a server group, and then click **Delete Server Group**.
   - On the home page for the server group, open the **Tasks** menu on the **Servers** tile,

and then click **Delete Server Group**.

2. When you are prompted to make sure you want to delete the server group, click **Yes**.

   Note
   Deleting a server group does not remove servers from Server Manager. Servers that were in a deleted group remain in the **All Servers** group, in the server pool.

3. When you are finished with changes to the group, click **OK**.

See Also

Add Servers to Server Manager

Manage Multiple, Remote Servers with Server Manager

Filter, Sort, and Query Data in Server Manager Tiles

In Windows Server 2012 R2 and Windows Server® 2012, tiles in Server Manager let you filter and sort data, and create and save custom queries. You can sort, use keyword filters, and run queries on list entries in the Events, Performance, Best Practices Analyzer, Services, and Roles and Features tiles on server role or group pages in Server Manager.

This topic contains the following sections.

- Filter list entries in tiles
- Sort list entries in tiles
- Create and run custom queries on tile data

Filter list entries in tiles

The **Filter** text box is a quick way of reducing the list of entries that are displayed in a tile to only those entries that contain a specified text string.

To apply a filter to the list of entries in a tile

1. Open a role or server group page in Server Manager.
2. In the **Filter** text box of an Events, Performance, Best Practices Analyzer, Services, or Roles and Features tile, type a string on which you want to filter.

   For example, if you want to see only events with an event ID of 1014, type **1014** in the **Filter** text box. All collected events that contain the string **1014** in at least one field are returned as results.

3. Note that the filter changes the description text under the tile title. Instead of **All** results, it says **Filtered results**.

4. To clear the filter, delete the string in the filter box, or click **X**.

Sort list entries in tiles

Sort list entries in Server Manager tiles by clicking column headings. Clicking a column heading the first time sorts column values in ascending alphanumeric order (arrow pointing up); clicking again sorts column values in descending alphanumeric order (arrow pointing down).

Create and run custom queries on tile data

You can create custom queries in the Events, Performance, Best Practices Analyzer, Services, or Roles and Features tiles in Server Manager. By default, the area of the tile toolbar in which you select criteria to build a custom query is hidden; click **Expand** (chevron button at right edge of tile toolbar) to display query criteria.

To create a custom query for tile data

1. Open a role or server group page in Server Manager.
2. In an Events, Performance, Best Practices Analyzer, Services, or Roles and Features tile, expand the query-building area by clicking **Expand**.
3. Click **Add criteria** to open a list of attributes (or fields) that apply to the entries in the tile.
4. Select criteria to add. When you are finished, click **Add**. Criteria that you selected are added to the query-building area.
5. Click the hypertext operator to select an operator. For numerical or date and time criteria, for example, the default is **less than or equal to**.
6. Specify acceptable values for the criteria. For example, if you selected **Date and Time**, provide a date in the format *m/d/yyyy*.
7. Repeat these steps from step 3 forward to add more criteria to your query.

   You can add duplicates of criteria that are already in your query, but the duplicates are added to the query with the **or** operator.

   For example, to query for event IDs 1003 or 1014, you would first add the ID criteria to your query, make the value of ID equal to **1003**, and then add a second ID criteria to your query, making the value of the second ID equal to **1014**. The resulting query is **and ID equals 1003 or ID equals 1014**.
8. When you are finished adding criteria and specifying operators and values, click **Save** to save the query.
9. Enter a friendly name for the query. For example, the query created in the preceding step can be named **Licensing events**.
10. When you are finished viewing query results, click **Clear All** to clear all filters and queries, and display all entries in the list.
11. To run a saved query, click **Saved Search Queries**, and click the name of the saved query that you want to run.

12. To delete a saved query, click **Saved Search Queries**, and then click **X** by the name of the saved query that you want to delete.

See Also

Manage Multiple, Remote Servers with Server Manager

View and Configure Performance, Event, and Service Data

Keyboard Shortcuts for Server Manager

Because Server Manager was fully redesigned starting in Windows Server® 2012, keyboard shortcuts that worked in the Server Manager console in Windows Server 2008 R2 or Windows Server 2008 are not necessarily the same commands. This topic describes the new keyboard shortcuts and access keys for Server Manager in Windows Server 2012 R2 and Windows Server 2012.

Commands that do not have their own keyboard shortcuts or access keys are accessible by pressing the **Tab** key, and tabbing through their control group when it is in focus.

Access keys

Active Control Area in Server Manager	Control Group	Access Key
Welcome tile	Welcome tile - Quick Start tab	**Alt+Q**
	Welcome tile - What's New tab	**Alt+W**
	Welcome tile - Learn More tab	**Alt+L**
	Welcome tile Hide command	**Alt+D**
Role and group thumbnails	Roles and Server Groups tile	**Alt+R**
Console header controls	Back button in the address bar	**Alt+Left arrow** or **Backspace**
	Forward button in the address bar	**Alt+right arrow**
	Refresh	**F5**
	Notifications area, open Task Details dialog box	**Alt+N**
	Manage menu	**Alt+M**
	View menu	**Alt+V**

Active Control Area in Server Manager	Control Group	Access Key
	Help menu	**Alt+H**
	Open Server Manager Help	**F1**
	Zoom in	**Ctrl+Plus (+)**
	Zoom out	**Ctrl+Minus (-)**
	Display console at 100%	**Ctrl+0**
Tiles on role, group, or Local Server pages	Local Server page Properties tile	**Alt+P**
	Role, group, or local server page page Events tile	**Alt+E**
	Role, group, or local server page Services tile	**Alt+R**
	Role, group, or local server page Best Practices Analyzer (BPA) tile	**Alt+B**
	Role, group, or local server page Performance tile	**Alt+O**
	Role, group, or local server page Roles and Features tile	**Alt+A**
	All Servers page Servers tile	**Alt+A**
Navigating within Local Server Properties tile	Computer name	**Alt+C**
	Last installed updates	**Alt+L**
	Domain or Workgroup	**Alt+D**
	Windows Update	**Alt+W**
	Last checked for updates	**Alt+S**
	Remote management	**Alt+R**
	Windows Firewall	**Alt+F**
	Remote Desktop	**Alt+K**
	Windows Error Reporting	**Alt+G**
	NIC Teaming	**Alt+T**
	Customer Experience	**Alt+X**

Active Control Area in Server Manager	Control Group	Access Key
	Improvement Program	
	Wired Ethernet connection	**Alt+O**
	IE Enhanced Security Configuration	**Alt+Y**
	Time zone	**Alt+Z**
	Product ID	**Alt+U**
Navigating within Events, Services, BPA, Performance, and Roles and Features tiles	Tasks menu	**Alt+T**
	Filter control	**Alt+F**
	Query control	**Alt+Q**
	Save queries	**Alt+S**

Deploy Remote Server Administration Tools

This topic supports both Remote Server Administration Tools for Windows® 8.1 and Remote Server Administration Tools for Windows® 8.

- Remote Server Administration Tools for Windows 8.1
- Remote Server Administration Tools for Windows 8

To help ease remote server management, you can download and install Remote Server Administration Tools. Remote Server Administration Tools includes Server Manager, Microsoft Management Console (MMC) snap-ins, consoles, Windows PowerShell® cmdlets and providers, and some command-line tools for managing roles and features that run on Windows Server.

Remote Server Administration Tools includes Windows PowerShell cmdlet modules that can be used to manage roles and features that are running on remote servers. Although Windows PowerShell remote management is enabled by default on Windows Server 2012 R2 and Windows Server 2012, it is not enabled by default on Windows 8.1 or Windows 8. To run cmdlets that are part of Remote Server Administration Tools against a remote server, run **Enable-PSRemoting** in a Windows PowerShell session that has been opened with elevated user rights (that is, Run as Administrator) on your Windows client computer after installing Remote Server Administration Tools.

Remote Server Administration Tools for Windows 8.1

Use Remote Server Administration Tools for Windows 8.1 to manage specific technologies on computers that are running Windows Server® 2012 R2, Windows Server® 2012, and in limited cases, Windows Server® 2008 R2, or Windows Server® 2008. A few of the tools might work for managing roles and features on Windows Server 2003.

Remote Server Administration Tools for Windows 8.1 includes support for remote management of computers that are running the Server Core installation option or the Minimal Server Interface configuration of Windows Server 2012 R2, Windows Server 2012, and in limited cases, the Server Core installation options of Windows Server 2008 R2, or Windows Server 2008. However, Remote Server Administration Tools for Windows 8.1 cannot be installed on any versions of the Windows Server operating system.

Tools available in this release

For a list of the tools available in Remote Server Administration Tools for Windows 8.1 and Remote Server Administration Tools for Windows 8, see the table in Remote Server Administration Tools (RSAT) for Windows Vista, Windows 7, Windows 8, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.

System requirements

Remote Server Administration Tools for Windows 8.1 can be installed only on computers that are running Windows 8.1. Remote Server Administration Tools cannot be installed on computers that are running Windows RT 8.1 or other system-on-chip devices.

Remote Server Administration Tools for Windows 8.1 runs on both x86- and x64-based editions of Windows 8.1.

Important
Remote Server Administration Tools for Windows 8.1 should not be installed on a computer that is running administration tools packs for Windows 8, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 or Windows 2000 Server®. Remove all older versions of Administration Tools Pack or Remote Server Administration Tools—including earlier prerelease versions, and releases of the tools for different languages or locales—from the computer before you install Remote Server Administration Tools for Windows 8.1.

To use this release of Server Manager to access and manage remote servers that are running Windows Server 2012, Windows Server 2008 or Windows Server 2008 R2, you must install several updates to make the older Windows Server operating systems manageable by using Server Manager. For detailed information about how to prepare Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008 for management by using Server Manager in Remote

Server Administration Tools for Windows 8.1, see Manage Multiple, Remote Servers with Server Manager.

Windows PowerShell and Server Manager remote management must be enabled on remote servers to manage them by using tools that are part of Remote Server Administration Tools for Windows 8.1. Remote management is enabled by default on servers that are running Windows Server 2012 R2 and Windows Server 2012. For more information about how to enable remote management if it has been disabled, see Manage multiple, remote servers with Server Manager.

Install, uninstall, or disable Remote Server Administration Tools for Windows 8.1

Remote Server Administration Tools for Windows 8.1 has the same one-step installation process as Remote Server Administration Tools for Windows 8. Before the release of Remote Server Administration Tools for Windows 8, after running the MSU installer program, users were required to turn on specific tools that they wanted to use in the **Turn Windows features on or off** dialog box. This step has been eliminated; after you run the MSU installation program, all tools are enabled by default.

To install supported language packs when the system UI language is not supported by the download package

1.

Note

This procedure applies only if your system UI language does not match one of the available Remote Server Administration Tools languages. If your computer's system UI language matches an available Remote Server Administration Tools language, go on to the next procedure, To install Remote Server Administration Tools for Windows 8.1.

Follow instructions in KB 2839636 to install supported Windows 8.1 language packs from Windows Update. Be sure to install one of the language packs that is supported by Remote Server Administration Tools.

Remote Server Administration Tools for Windows 8.1 is available in cs-CZ, de-DE, en-US, es-ES, fr-FR, hu-HU, it-IT, ja-JP, ko-KR, nl-NL, pl-PL, pt-BR, pt-PT, ru-RU, sv-SE, tr-TR, zh-CN, and zh-TW.

a. On the **Start** screen, type **language** to search for language settings.
b. In the search results, click or tap **Add a language to this device**.
c. On the **Time and Language** page, under **Languages**, click or tap **Add a language**.
d. Select and install a language that is supported by RSAT.

2. When the language pack installation is complete, verify that the language pack is installed by running `get-windowspackage -online` in a Windows PowerShell session that has been started with elevated user rights (Run as Administrator).

The language pack entry starts with the following:

Microsoft-Windows-Client-LanguagePack-Package~(GUID)~(architecture)~(supported language)~6.3.9600.XXXXX

3. After you have installed a supported language pack, go on to the next procedure, To install Remote Server Administration Tools for Windows 8.1.

To install Remote Server Administration Tools for Windows 8.1

1. Download the Remote Server Administration Tools for Windows 8.1 package from the Microsoft Download Center. You can either run the installer from the Download Center website, or save the download package to a local computer or share.
2. If you save the download package to a local computer or share, double-click the installer program, **Windows6.3-KB2693643-x64.msu** or **Windows6.3-KB2693643-x86.msu**, depending on the architecture of the computer on which you want to install the tools.
3. When you are prompted by the **Windows Update Standalone Installer** dialog box to install the update, click **Yes**.
4. Read and accept the license terms. Click **I accept**.
5. Installation requires a few minutes to finish.

To uninstall Remote Server Administration Tools for Windows 8

1. On the desktop, hover in the upper right corner of the screen, and then click **Settings**.
2. On the **Desktop** menu, click **Control Panel**.
3. Under **Programs**, click **Uninstall a program**.
4. Click **View installed updates**.
5. Right-click **Update for Microsoft Windows (KB2693643)**, and then click **Uninstall**.
6. When you are asked if you are sure you want to uninstall the update, click **Yes**.

To turn off specific tools

1. On the desktop, hover in the upper right corner of the screen, and then click **Settings**.
2. On the **Desktop** menu, click **Control Panel**.
3. Click **Programs**, and then in **Programs and Features** click **Turn Windows features on or off**.
4. In the **Windows Features** dialog box, expand **Remote Server Administration Tools**, and then expand either **Role Administration Tools** or **Feature Administration Tools**.
5. Clear the check boxes for any tools that you want to turn off.

   Note
   If you turn off Server Manager, the computer must be restarted, and tools that were accessible from the **Tools** menu of Server Manager must be opened from the **Administrative Tools** folder.
6. When you are finished turning off tools that you do not want to use, click **OK**.

Run Remote Server Administration Tools

Notes

After installing Remote Server Administration Tools for Windows 8.1, the **Administrative Tools** folder is not displayed on the **Start** screen. This is a known issue in this release of Remote Server Administration Tools for Windows 8.1. You can access the tools from the following locations.

- The **Tools** menu in the Server Manager console.
- **Control Panel\System and Security\Administrative Tools**.
- A shortcut saved to the desktop from the **Administrative Tools** folder (to do this, right click the **Control Panel\System and Security\Administrative Tools** link, and then click **Create Shortcut**).

The tools installed as part of Remote Server Administration Tools for Windows 8.1 cannot be used to manage the local computer. Regardless of the tool you run, you must specify a remote server or multiple remote servers on which to run the tool. Because most tools are integrated with Server Manager, you add remote servers that you want to manage to the Server Manager server pool before managing the server by using the tools in the **Tools** menu. For more information about how to add servers to your server pool, and create custom groups of servers, see Add servers to Server Manager and Create and manage server groups.

In Remote Server Administration Tools for Windows 8.1, all GUI-based server management tools, such as MMC snap-ins and dialog boxes, are accessed from the **Tools** menu of the Server Manager console. Although the computer that runs Remote Server Administration Tools for Windows 8.1 runs a client-based operating system, after installing the tools, Server Manager, included with Remote Server Administration Tools for Windows 8.1, opens automatically by default on the client computer. Note that there is no **Local Server** page in the Server Manager console that runs on a client computer.

To start Server Manager on a client computer

1. On the **Start** screen, click **Administrative Tools**.
2. In the **Administrative Tools** folder, double-click **Server Manager**.
3. If neither the **Administrative Tools** nor the **Server Manager** tiles are displayed on the **Start** screen after installing Remote Server Administration Tools, and searching for Server Manager on the **Start** screen does not display results, verify that the **Show administrative tools** setting is turned on. To view this setting, hover the mouse cursor over the upper right corner of the **Start** screen, and then click **Settings**. If **Show administrative tools** is turned off, turn the setting on to display tools that you have installed as part of Remote Server Administration Tools for Windows 8.

Although they are not listed in the Server Manager console **Tools** menu, Windows PowerShell cmdlets and Command Prompt management tools are also installed for roles and features as part of Remote Server Administration Tools. For example, if you open a Windows PowerShell session with elevated user rights (Run as Administrator), and run the cmdlet `Get-Command -Module RDManagement`, the results include a list of Remote Desktop Services cmdlets that are now

available to run on the local computer after installing Remote Server Administration Tools, as long as the cmdlets are targeted at a remote server that is running all or part of the Remote Desktop Services role.

To start Windows PowerShell with elevated user rights (Run as administrator)

1. On the **Start** screen, right-click the **Windows PowerShell** tile, and then click **Run as administrator**.
2. To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.

Note

You can also start a Windows PowerShell session that is targeted at a specific server by right-clicking a managed server in a role or group page in Server Manager, and then clicking **Windows PowerShell**.

Remote Server Administration Tools for Windows 8

Use Remote Server Administration Tools for Windows 8 to manage specific technologies on computers that are running Windows Server 2012, and in limited cases, Windows Server 2008 R2, or Windows Server 2008. Some of the tools work for managing roles and features on Windows Server 2003.

Remote Server Administration Tools for Windows 8 includes support for remote management of computers that are running the Server Core installation option or the Minimal Server Interface configuration of Windows Server 2012, and in limited cases, the Server Core installation options of Windows Server 2008 R2, or Windows Server 2008. However, Remote Server Administration Tools for Windows 8 cannot be installed on any versions of the Windows Server operating system.

Tools available in this release

For a list of the tools available in Remote Server Administration Tools for Windows 8 and Remote Server Administration Tools for Windows 8, see the table in Remote Server Administration Tools (RSAT) for Windows Vista, Windows 7, Windows 8, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.

System requirements

Remote Server Administration Tools for Windows 8 can be installed only on computers that are running Windows 8. Remote Server Administration Tools cannot be installed on computers that are running Windows RT or other system-on-chip devices.

Remote Server Administration Tools for Windows 8 runs on both x86- and x64-based editions of Windows 8.

Important

Remote Server Administration Tools for Windows 8 should not be installed on a computer that is running administration tools packs for Windows Server 2008 R2, Windows Server 2008, Windows Server 2003 or Windows 2000 Server®. Remove all older versions of Administration Tools Pack or Remote Server Administration Tools—including earlier prerelease versions, and releases of the tools for different languages or locales—from the computer before you install Remote Server Administration Tools for Windows 8.

To use this release of Server Manager to access and manage remote servers that are running Windows Server 2008 or Windows Server 2008 R2, you must install several updates to make the older Windows Server operating systems manageable by using Server Manager. For detailed information about how to prepare Windows Server 2008 R2 and Windows Server 2008 for management by using Server Manager in Remote Server Administration Tools for Windows 8, see Manage Multiple, Remote Servers with Server Manager.

Windows PowerShell and Server Manager remote management must be enabled on remote servers to manage them by using tools that are part of Remote Server Administration Tools for Windows 8. Remote management is enabled by default on servers that are running Windows Server 2012. For more information about how to enable remote management, see Manage multiple, remote servers with Server Manager.

Install, uninstall, or disable Remote Server Administration Tools for Windows 8

Remote Server Administration Tools for Windows 8 has a faster, one-step installation process. In earlier releases of Remote Server Administration Tools, after running the MSU installer program, users were required to turn on specific tools that they wanted to use in the **Turn Windows features on or off** dialog box. This step has been eliminated; after you run the MSU installation program, all tools are enabled by default.

To install supported language packs when the system UI language is not supported by the download package

1.

Note

This procedure applies only if your system UI language does not match one of the available Remote Server Administration Tools languages. If your computer's system UI language matches an available Remote Server Administration Tools language, go on to the next procedure, To install Remote Server Administration Tools for Windows 8.

Follow instructions in KB 2607607 to install supported Windows 8 language packs from Windows Update. Be sure to install one of the language packs that is supported by Remote Server Administration Tools for Windows 8.

Remote Server Administration Tools for Windows 8 is available in the following

languages: cs-CZ, de-DE, en-US, es-ES, fr-FR, hu-HU, it-IT, ja-JP, ko-KR, nl-NL, pl-PL, pt-BR, pt-PT, ru-RU, sv-SE, tr-TR, zh-CN, zh-HK, and zh-TW.

2. When the language pack installation is complete, verify that the language pack is installed by running `get-windowspackage -online` in a Windows PowerShell session that has been started with elevated user rights (Run as Administrator).

   The language pack entry starts with the following:

   Microsoft-Windows-Client-LanguagePack-Package~(GUID)~(architecture)~(supported language)~6.2.9200.16384

3. After you have installed a supported language pack, go on to the next procedure, To install Remote Server Administration Tools for Windows 8.

To install Remote Server Administration Tools for Windows 8

1. Download the Remote Server Administration Tools for Windows 8 package from the Microsoft Download Center. You can either run the installer from the Download Center website, or save the download package to a local computer or share.
2. If you save the download package to a local computer or share, double-click the installer program, **Windows6.2-KB2693643-x64.msu** or **Windows6.2-KB2693643-x86.msu**, depending on the architecture of the computer on which you want to install the tools.
3. When you are prompted by the **Windows Update Standalone Installer** dialog box to install the update, click **Yes**.
4. Read and accept the license terms. Click **I accept**.
5. Installation requires a few minutes to finish.

To uninstall Remote Server Administration Tools for Windows 8

1. On the desktop, hover in the upper right corner of the screen, and then click **Settings**.
2. On the **Desktop** menu, click **Control Panel**.
3. Under **Programs**, click **Uninstall a program**.
4. Click **View installed updates**.
5. Right-click **Update for Microsoft Windows (KB2693643)**, and then click **Uninstall**.
6. When you are asked if you are sure you want to uninstall the update, click **Yes**.

To turn off specific tools

1. On the desktop, hover in the upper right corner of the screen, and then click **Settings**.
2. On the **Desktop** menu, click **Control Panel**.
3. Click **Programs**, and then in **Programs and Features** click **Turn Windows features on or off**.
4. In the **Windows Features** dialog box, expand **Remote Server Administration Tools**, and then expand either **Role Administration Tools** or **Feature Administration Tools**.
5. Clear the check boxes for any tools that you want to turn off.

Note

If you turn off Server Manager, the computer must be restarted, and tools that were accessible from the **Tools** menu of Server Manager must be opened from the **Administrative Tools** folder.

6. When you are finished turning off tools that you do not want to use, click **OK**.

Run Remote Server Administration Tools

The tools installed as part of Remote Server Administration Tools for Windows 8 cannot be used to manage the local computer. Regardless of the tool you run, you must specify a remote server or multiple remote servers on which to run the tool. Because most tools are integrated with Server Manager, you add remote servers that you want to manage to the Server Manager server pool before managing the server by using the tools in the **Tools** menu. For more information about how to add servers to your server pool, and create custom groups of servers, see Add servers to Server Manager and Create and manage server groups.

In Remote Server Administration Tools for Windows 8, all GUI-based server management tools, such as MMC snap-ins and dialog boxes, are accessed from the **Tools** menu of the Server Manager console. Although the computer that runs Remote Server Administration Tools for Windows 8 runs a client-based operating system, after installing the tools, Server Manager, included with Remote Server Administration Tools for Windows 8, opens automatically by default on the client computer. Note that there is no **Local Server** page in the Server Manager console that runs on a client computer.

To start Server Manager on a client computer

1. On the **Start** screen, click **Administrative Tools**.
2. In the **Administrative Tools** folder, double-click **Server Manager**.
3. If neither the **Administrative Tools** nor the **Server Manager** tiles are displayed on the **Start** screen after installing Remote Server Administration Tools, and searching for Server Manager on the **Start** screen does not display results, verify that the **Show administrative tools** setting is turned on. To view this setting, hover the mouse cursor over the upper right corner of the **Start** screen, and then click **Settings**. If **Show administrative tools** is turned off, turn the setting on to display tools that you have installed as part of Remote Server Administration Tools for Windows 8.

Although they are not listed in the Server Manager console **Tools** menu, Windows PowerShell cmdlets and Command Prompt management tools are also installed for roles and features as part of Remote Server Administration Tools. For example, if you open a Windows PowerShell session with elevated user rights (Run as Administrator), and run the cmdlet `Get-Command -Module RDManagement`, the results include a list of Remote Desktop Services cmdlets that are now available to run on the local computer after installing Remote Server Administration Tools, as long as the cmdlets are targeted at a remote server that is running all or part of the Remote Desktop Services role.

To start Windows PowerShell with elevated user rights (Run as administrator)

1. On the **Start** screen, right-click the **Windows PowerShell** tile, and then click **Run as administrator**.
2. To run Windows PowerShell as an administrator from the desktop, right-click the **Windows PowerShell** shortcut in the taskbar, and then click **Run as Administrator**.

Note

You can also start a Windows PowerShell session that is targeted at a specific server by right-clicking a managed server in a role or group page in Server Manager, and then clicking **Windows PowerShell**.

See Also

Remote Server Administration Tools for Windows 8.1

Remote Server Administration Tools for Windows 8

Remote Server Administration Tools (RSAT) for Windows Vista, Windows 7, Windows 8, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2

Manage multiple, remote servers with Server Manager

Windows PowerShell Support for Windows Server 2012 R2 and Windows Server 2012

Install and Use Windows PowerShell Web Access

Windows PowerShell® Web Access, first introduced in Windows Server® 2012, acts as a Windows PowerShell gateway, providing a web-based Windows PowerShell console that is targeted at a remote computer. It enables IT Pros to run Windows PowerShell commands and scripts from a Windows PowerShell console in a web browser, with no Windows PowerShell, remote management software, or browser plug-in installation necessary on the client device. All that is required to run the web-based Windows PowerShell console is a properly-configured Windows PowerShell Web Access gateway, and a client device browser that supports JavaScript® and accepts cookies.

Examples of client devices include laptops, non-work personal computers, borrowed computers, tablet computers, web kiosks, computers that are not running a Windows-based operating system, and cell phone browsers. IT Pros can perform critical management tasks on remote Windows-based servers from devices that have access to an Internet connection and a web browser.

After successful gateway setup and configuration, users can access a Windows PowerShell console by using a web browser. When users open the secured Windows PowerShell Web

Access website, they can run a web-based Windows PowerShell console after successful authentication.

Windows PowerShell Web Access setup and configuration is a three-step process:

1. Installing Windows PowerShell Web Access
2. Configuring the gateway
3. Configuring authorization rules that allow users access to the web-based Windows PowerShell console

Before you install and configure Windows PowerShell Web Access, we recommend that you read this entire guide, which includes instructions about how to install, secure, and uninstall Windows PowerShell Web Access. The Use the Web-based Windows PowerShell Console topic describes how users sign in to the web-based console, and covers limitations and differences between the web-based Windows PowerShell console and the **powershell.exe** console. End users of the web-based console should read Use the Web-based Windows PowerShell Console, but do not need to read the rest of this guide.

This topic does not provide in-depth Web Server (IIS) operations guidance; only those steps required to configure the Windows PowerShell Web Access gateway are described in this topic. For more information about configuring and securing websites in IIS, see the IIS documentation resources in the See Also section.

The following diagram shows how Windows PowerShell Web Access works.

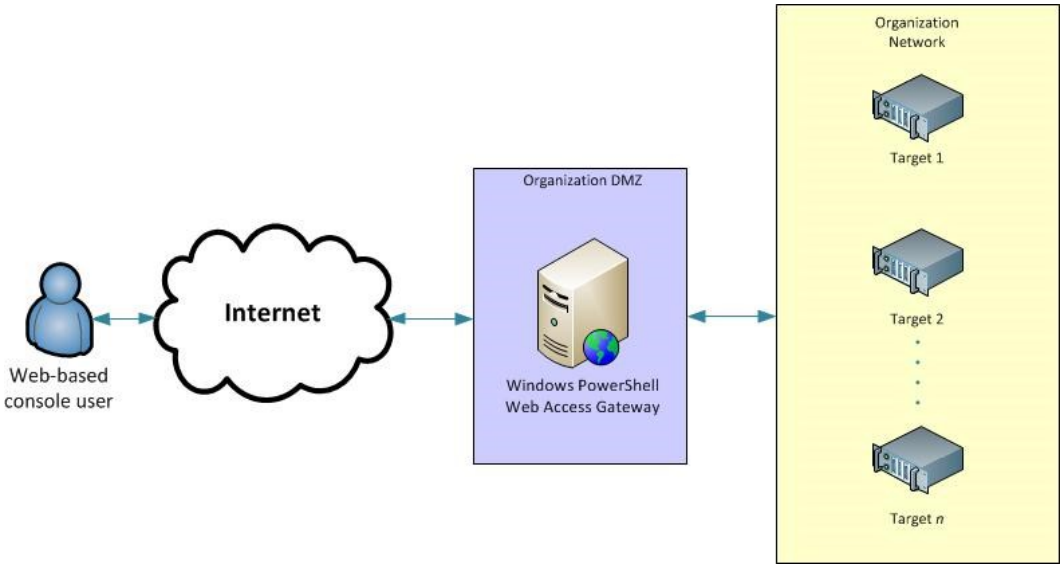

In this topic:

- Requirements for running Windows PowerShell Web Access
- Browser and client device support
- Recommended (quick) deployment
- Custom deployment
- Configure a genuine certificate

Requirements for running Windows PowerShell Web Access

Windows PowerShell Web Access requires Web Server (IIS), .NET Framework 4.5, and Windows PowerShell 3.0 or Windows PowerShell 4.0 to be running on the server on which you want to run the gateway. You can install Windows PowerShell Web Access on a server that is running Windows Server 2012 R2 or Windows Server 2012 by using either the Add Roles and Features Wizard in Server Manager, or Windows PowerShell deployment cmdlets for Server Manager. When you install Windows PowerShell Web Access by using Server Manager or its deployment cmdlets, required roles and features are automatically added as part of the installation process.

Windows PowerShell Web Access allows remote users to access computers in your organization by using Windows PowerShell in a web browser. Although Windows PowerShell Web Access is a convenient and powerful management tool, the web-based access poses security risks, and should be configured as securely as possible. We recommend that administrators who configure the Windows PowerShell Web Access gateway use available security layers, both the cmdlet-based authorization rules included with Windows PowerShell Web Access, and security layers that are available in Web Server (IIS) and third-party applications. This documentation includes both unsecure examples that are only recommended for test environments, as well as examples that are recommended for secure deployments.

Browser and client device support

Windows PowerShell Web Access supports the following Internet browsers. Although mobile browsers are not officially supported, many may be able to run the web-based Windows PowerShell console. Other browsers that accept cookies, run JavaScript, and run HTTPS websites are expected to work, but are not officially tested.

Supported desktop computer browsers

- Windows® Internet Explorer® for Microsoft Windows® 8.0, 9.0, 10.0, and 11.0
- Mozilla Firefox® 10.0.2
- Google Chrome™ 17.0.963.56m for Windows
- Apple Safari® 5.1.2 for Windows
- Apple Safari 5.1.2 for Mac OS®

Minimally-tested mobile devices or browsers

- Windows Phone 7 and 7.5
- Google Android WebKit 3.1 Browser Android 2.2.1 (Kernel 2.6)
- Apple Safari for iPhone operating system 5.0.1
- Apple Safari for iPad 2 operating system 5.0.1

Browser requirements

To use the Windows PowerShell Web Access web-based console, browsers must do the following.

- Allow cookies from the Windows PowerShell Web Access gateway website.
- Be able to open and read HTTPS pages.
- Open and run websites that use JavaScript.

Recommended (quick) deployment

You can install the Windows PowerShell Web Access gateway on a server that is running Windows Server 2012 R2 or Windows Server 2012 by using either Windows PowerShell cmdlets, or by using the Add Roles and Features Wizard that is opened from within Server Manager. For quick installation and configuration, use Windows PowerShell cmdlets, as described in this section.

- Step 1: Install Windows PowerShell Web Access
- Step 2: Configure the gateway
- Step 3: Configure a restrictive authorization rule

Step 1: Install Windows PowerShell Web Access

To install Windows PowerShell Web Access by using Windows PowerShell cmdlets

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then click **Run as Administrator**.

   Note
   In Windows PowerShell 3.0 and 4.0, there is no need to import the Server Manager cmdlet module into the Windows PowerShell session before running cmdlets that are part of the module. A module is automatically imported the first time you run a cmdlet that is part of the module. Also, Windows PowerShell cmdlets are not case-sensitive.
2. Type the following, and then press **Enter**, where *computer_name* represents a remote computer on which you want to install Windows PowerShell Web Access, if applicable. The `Restart` parameter automatically restarts destination servers if required.

   ```
   Install-WindowsFeature –Name WindowsPowerShellWebAccess -
   ComputerName <computer_name> -IncludeManagementTools -Restart
   ```

   Note
   Installing Windows PowerShell Web Access by using Windows PowerShell

cmdlets does not add Web Server (IIS) management tools by default. If you want to install the management tools on the same server as the Windows PowerShell Web Access gateway, add the `IncludeManagementTools` parameter to the installation command (as provided in this step). If you are managing the Windows PowerShell Web Access website from a remote computer, install the IIS Manager snap-in by installing Remote Server Administration Tools for Windows 8.1 or Remote Server Administration Tools for Windows 8 on the computer from which you want to manage the gateway.

To install roles and features on an offline VHD, you must add both the `ComputerName` parameter and the `VHD` parameter. The `ComputerName` parameter contains the name of the server on which to mount the VHD, and the `VHD` parameter contains the path to the VHD file on the specified server.

```
Install-WindowsFeature –Name WindowsPowerShellWebAccess –VHD
<path> -ComputerName <computer_name> -IncludeManagementTools
-Restart
```

3. When installation is complete, verify that Windows PowerShell Web Access was installed on destination servers by running the **Get-WindowsFeature** cmdlet on a destination server, in a Windows PowerShell console that has been opened with elevated user rights. You can also verify that Windows PowerShell Web Access was installed in the Server Manager console, by selecting a destination server on the **All Servers** page, and then viewing the **Roles and Features** tile for the selected server. You can also view the readme file for Windows PowerShell Web Access.
4. After Windows PowerShell Web Access is installed, you are prompted to review the readme file, which contains basic, required setup instructions for the gateway. These setup instructions are also in the following section, Step 2: Configure the gateway. The path to the readme file is `C:\Windows\Web\PowerShellWebAccess\wwwroot\README.txt`.

Step 2: Configure the gateway

The **Install-PswaWebApplication** cmdlet is a quick way to get Windows PowerShell Web Access configured. Although you can add the `UseTestCertificate` parameter to the `Install-PswaWebApplication` cmdlet to install a self-signed SSL certificate for test purposes, this is not secure; for a secure production environment, always use a valid SSL certificate that has been signed by a certification authority (CA). Administrators can replace the test certificate with a signed certificate of their choice by using the IIS Manager console.

You can complete Windows PowerShell Web Access web application configuration either by running the `Install-PswaWebApplication` cmdlet or by performing GUI-based configuration steps in IIS Manager. By default, the cmdlet installs the web application, **pswa** (and an application pool for it, **pswa_pool**), in the **Default Web Site** container, as shown in IIS Manager; if desired, you can instruct the cmdlet to change the default site container of the web application. IIS Manager offers configuration options that are available for web applications, such as changing the port number or the Secure Sockets Layer (SSL) certificate.

Security

We strongly recommend that administrators configure the gateway to use a valid certificate that has been signed by a CA.

- To configure the Windows PowerShell Web Access gateway with a test certificate by using Install-PswaWebApplication
- To configure the Windows PowerShell Web Access gateway with a genuine certificate by using Install-PswaWebApplication and IIS Manager

To configure the Windows PowerShell Web Access gateway with a test certificate by using Install-PswaWebApplication

1. Do one of the following to open a Windows PowerShell session.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar.
   - On the Windows **Start** screen, click **Windows PowerShell**.
2. Type the following, and then press **Enter**.

   Install-PswaWebApplication -UseTestCertificate

   Security

   The `UseTestCertificate` parameter should only be used in a private test environment. For a secure production environment, we recommend using a valid certificate that has been signed by a CA.

   Running the cmdlet installs the Windows PowerShell Web Access web application within the IIS Default Web Site container. The cmdlet creates the infrastructure required to run Windows PowerShell Web Access on the default website, https://<server_name>/pswa. To install the web application in a different website, provide the website name by adding the `WebSiteName` parameter. To change the name of the web application (the default is `pswa`), add the `WebApplicationName` parameter.

   The following settings are configured by running the cmdlet. You can change these manually in the IIS Manager console, if desired.

   - Path: /pswa
   - ApplicationPool: pswa_pool
   - EnabledProtocols: http
   - PhysicalPath: %*windir*%/Web/PowerShellWebAccess/wwwroot

   Example: `Install-PswaWebApplication -webApplicationName myWebApp -useTestCertificate`

   In this example, the resulting website for Windows PowerShell Web Access is https://<*server_name*>/myWebApp.

   Note

   You cannot sign in until users have been granted access to the website by adding authorization rules. For more information, see Step 3: Configure a restrictive authorization rule and Authorization Rules and Security Features of

Windows PowerShell Web Access.

To configure the Windows PowerShell Web Access gateway with a genuine certificate by using Install-PswaWebApplication and IIS Manager

1. Do one of the following to open a Windows PowerShell session.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar.
   - On the Windows **Start** screen, click **Windows PowerShell**.
2. Type the following, and then press **Enter**.

   Install-PswaWebApplication

   The following gateway settings are configured by running the cmdlet. You can change these manually in the IIS Manager console, if desired. You can also specify values for the `WebsiteName` and `WebApplicationName` parameters of the `Install-PswaWebApplication` cmdlet.
   - Path: /pswa
   - ApplicationPool: pswa_pool
   - EnabledProtocols: http
   - PhysicalPath: %*windir*%/Web/PowerShellWebAccess/wwwroot
3. Open the IIS Manager console by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar. On the **Tools** menu in Server Manager, click **Internet Information Services (IIS) Manager**.
   - On the Windows **Start** screen, click **Server Manager**.
4. In the IIS Manager tree pane, expand the node for the server on which Windows PowerShell Web Access is installed until the **Sites** folder is visible. Expand the **Sites** folder.
5. Select the website in which you have installed the Windows PowerShell Web Access web application. In the **Actions** pane, click **Bindings**.
6. In the **Site Binding** dialog box, click **Add**.
7. In the **Add Site Binding** dialog box, in the **Type** field, select **https**.
8. In the **SSL certificate** field, select your signed certificate from the drop-down menu. Click **OK**. See To configure an SSL certificate in IIS Manager in this topic for more information about how to obtain a certificate.

   The Windows PowerShell Web Access web application is now configured to use your signed SSL certificate. You can access Windows PowerShell Web Access by opening https://<server_name>/pswa in a browser window.

   Note
   You cannot sign in until users have been granted access to the website by adding authorization rules. For more information, see Step 3: Configure a restrictive authorization rule and Authorization Rules and Security Features of Windows PowerShell Web Access.

Step 3: Configure a restrictive authorization rule

After Windows PowerShell Web Access is installed and the gateway is configured, users can open the sign-in page in a browser, but they cannot sign in until the Windows PowerShell Web Access administrator grants users access explicitly. Windows PowerShell Web Access access control is managed by using the set of Windows PowerShell cmdlets described in the following table. There is no comparable GUI for adding or managing authorization rules. For more detailed information about Windows PowerShell Web Access cmdlets, see the cmdlet reference topics, Windows PowerShell Web Access Cmdlets.

For more detail about Windows PowerShell Web Access authorization rules and security, see Authorization Rules and Security Features of Windows PowerShell Web Access.

To add a restrictive authorization rule

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then click **Run as Administrator**.
2. **Optional step for restricting user access by using session configurations:** Verify that session configurations that you want to use in your rules already exist. If they have not yet been created, use instructions for creating session configurations in about_Session_Configuration_Files on MSDN.
3. Type the following, and then press **Enter**.

   ```
   Add-PswaAuthorizationRule –UserName <domain\user |
   computer\user> -ComputerName <computer_name> -
   ConfigurationName <session_configuration_name>
   ```

   This authorization rule allows a specific user access to one computer on the network to which they typically have access, with access to a specific session configuration that is scoped to the user's typical scripting and cmdlet needs. In the following example, a user named `JSmith` in the `Contoso` domain is granted access to manage the computer `Contoso_214`, and use a session configuration named `NewAdminsOnly`.

   ```
   Add-PswaAuthorizationRule –UserName Contoso\JSmith -
   ComputerName Contoso_214 -ConfigurationName NewAdminsOnly
   ```

4. Verify that the rule has been created by running either the **Get-PswaAuthorizationRule** cmdlet, or **Test-PswaAuthorizationRule -UserName <domain\user | computer\user> -ComputerName** <computer_name>. For example, **Test-PswaAuthorizationRule –UserName Contoso\JSmith –ComputerName Contoso_214**.

After you have configured an authorization rule, you are ready for authorized users to sign in to the web-based console and begin using Windows PowerShell Web Access.

Custom deployment

You can install the Windows PowerShell Web Access gateway on a server that is running Windows Server 2012 R2 or Windows Server 2012 by using the Add Roles and Features Wizard in Server Manager. After Windows PowerShell Web Access is installed, you can customize the configuration of the gateway in IIS Manager.

Step 1: Install Windows PowerShell Web Access

To install Windows PowerShell Web Access by using the Add Roles and Features Wizard

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** screen, click **Server Manager**.
2. On the **Manage** menu, click **Add Roles and Features**.
3. On the **Select installation type** page, select **Role-based or feature-based installation**. Click **Next**.
4. On the **Select destination server** page, select a server from the server pool, or select an offline VHD. To select an offline VHD as your destination server, first select the server on which to mount the VHD, and then select the VHD file. For information about how to add servers to your server pool, see the Server Manager Help. After you have selected the destination server, click **Next**.
5. On the **Select features** page of the wizard, expand **Windows PowerShell**, and then select **Windows PowerShell Web Access**.
6. Note that you are prompted to add required features, such as .NET Framework 4.5, and role services of Web Server (IIS). Add required features and continue.

   Note
   Installing Windows PowerShell Web Access by using the Add Roles and Features Wizard also installs Web Server (IIS), including the IIS Manager snap-in. The snap-in and other IIS management tools are installed by default if you use Add Roles and Features Wizard. If you install Windows PowerShell Web Access by using Windows PowerShell cmdlets as described in the following procedure, management tools are not added by default.
7. On the **Confirm installation selections** page, if the feature files for Windows PowerShell Web Access are not stored on the destination server that you selected in step 4, click **Specify an alternate source path**, and provide the path to the feature files. Otherwise, click **Install**.
8. After you click **Install**, the **Installation progress** page displays installation progress, results, and messages such as warnings, failures, or post-installation configuration steps that are required for Windows PowerShell Web Access. After Windows PowerShell Web

Access is installed, you are prompted to review the readme file, which contains basic, required setup instructions for the gateway. These instructions are also included in this topic. The path to the readme file is `C:\Windows\Web\PowerShellWebAccess\wwwroot\README.txt`.

Step 2: Configure the gateway

Instructions in this section are for installing the Windows PowerShell Web Access web application in a subdirectory—and not in the root directory—of your website. This procedure is the GUI-based equivalent of the actions performed by the `Install-PswaWebApplication` cmdlet. This section also includes instructions for how to use IIS Manager to configure the Windows PowerShell Web Access gateway as a root website.

- To use IIS Manager to configure the gateway in an existing website
- To use IIS Manager to configure the gateway as a root website with a test certificate
-

To use IIS Manager to configure the gateway in an existing website

1. Open the IIS Manager console by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar. On the **Tools** menu in Server Manager, click **Internet Information Services (IIS) Manager**.
   - On the Windows **Start** screen, type any part of the name **Internet Information Services (IIS) Manager**. Click the shortcut when it is displayed in the **Apps** results.
2. Create a new application pool for Windows PowerShell Web Access. Expand the node of the gateway server in the IIS Manager tree pane, select **Application Pools**, and click **Add Application Pool** in the **Actions** pane.
3. Add a new application pool with the name **pswa_pool**, or provide another name. Click **OK**.
4. In the IIS Manager tree pane, expand the node for the server on which Windows PowerShell Web Access is installed until the **Sites** folder is visible. Select the **Sites** folder.
5. Right-click the website (for example, **Default Web Site**) to which you would like to add the Windows PowerShell Web Access website, and then click **Add Application**.
6. In the **Alias** field, type pswa, or provide another alias. The alias becomes the virtual directory name. For example, **pswa** in the following URL represents the alias specified in this step: https://<server_name>/pswa.
7. In the **Application pool** field, select the application pool that you created in step 3.
8. In the **Physical path** field, browse for the location of the application. You can use the default location, %windir%/Web/PowerShellWebAccess/wwwroot. Click **OK**.
9. Follow the steps in the procedure To configure an SSL certificate in IIS Manager in this topic.
10. **Optional security step:** With the website selected in the tree pane, double-click **SSL**

Settings in the content pane. Select **Require SSL**, and then in the **Actions** pane, click **Apply**. Optionally, in the **SSL Settings** pane, you can require that users connecting to the Windows PowerShell Web Access website have client certificates. Client certificates help to verify the identity of a client device user. For more information about how requiring client certificates can increase the security of Windows PowerShell Web Access, see Authorization Rules and Security Features of Windows PowerShell Web Access in this guide.

11. Open a browser session on a client device. For more information about supported browsers and devices, see Browser and client device support in this topic.
12. Open the new Windows PowerShell Web Access website, https://<*gateway_server_name*>/pswa.

    The browser should display the Windows PowerShell Web Access console sign-in page.

    Note
    You cannot sign in until users have been granted access to the website by adding authorization rules.
13. In a Windows PowerShell session that has been opened with elevated user rights (Run as Administrator), run the following script, in which *application_pool_name* represents the name of the application pool that you created in step 3, to give the application pool access rights to the authorization file.

```
$applicationPoolName = "<application_pool_name>"
$authorizationFile =
"C:\windows\web\powershellwebaccess\data\AuthorizationRules.x
ml"
c:\windows\system32\icacls.exe $authorizationFile /grant ('"'
+ "IIS AppPool\$applicationPoolName" + '":R') > $null
```

    To view existing access rights on the authorization file, run the following command:

```
c:\windows\system32\icacls.exe $authorizationFile
```

To use IIS Manager to configure the gateway as a root website with a test certificate

1. Open the IIS Manager console by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar. On the **Tools** menu in Server Manager, click **Internet Information Services (IIS) Manager**.
   - On the Windows **Start** screen, type any part of the name **Internet Information Services (IIS) Manager**. Click the shortcut when it is displayed in the **Apps** results.
2. In the IIS Manager tree pane, expand the node for the server on which Windows PowerShell Web Access is installed until the **Sites** folder is visible. Select the **Sites** folder.
3. In the **Actions** pane, click **Add Website**.
4. Type a name for the website, such as **Windows PowerShell Web Access**.

5. An application pool is automatically created for the new website. To use a different application pool, click **Select** to select an application pool to associate with the new website. Select the alternate application pool in the **Select Application Pool** dialog box, and then click **OK**.
6. In the **Physical path** text box, navigate to %*windir*%/Web/PowerShellWebAccess/wwwroot.
7. In the **Type** field of the **Binding** area, select **https**.
8. Assign a port number to the website that is not already in use by another site or application. To locate open ports, you can run the **netstat** command in a Command Prompt window. The default port number is 443.

   Change the default port if another website is already using 443, or if you have other security reasons for changing the port number. If another website that is running on your gateway server is using your selected port, a warning is displayed when you click **OK** in the **Add Website** dialog box. You must use an unused port to run Windows PowerShell Web Access.
9. Optionally, if needed for your organization, specify a host name that makes sense to your organization and users, such as **www.contoso.com**. Click **OK**.
10. For a more secure production environment, we strongly recommend providing a valid certificate that has been signed by a CA. You must provide an SSL certificate, because users can only connect to Windows PowerShell Web Access through an HTTPS website. See To configure an SSL certificate in IIS Manager in this topic for more information about how to obtain a certificate.
11. Click **OK** to close the **Add Website** dialog box.
12. In a Windows PowerShell session that has been opened with elevated user rights (Run as Administrator), run the following script, in which *application_pool_name* represents the name of the application pool that you created in step 4, to give the application pool access rights to the authorization file.

    ```
    $applicationPoolName = "<application_pool_name>"
    $authorizationFile =
    "C:\windows\web\powershellwebaccess\data\AuthorizationRules.x
    ml"
    c:\windows\system32\icacls.exe $authorizationFile /grant ('"'
    + "IIS AppPool\$applicationPoolName" + '":R') > $null
    ```

    To view existing access rights on the authorization file, run the following command:

    ```
    c:\windows\system32\icacls.exe $authorizationFile
    ```
13. With the new website selected in the IIS Manager tree pane, click **Start** in the **Actions** pane to start the website.
14. Open a browser session on a client device. For more information about supported browsers and devices, see Browser and client device support in this document.
15. Open the new Windows PowerShell Web Access website.

    Because the root website points to the Windows PowerShell Web Access folder, the

browser should display the Windows PowerShell Web Access sign-in page when you open https://< *gateway_server_name*>. You should not need to add **/pswa** to the URL.

Note
You cannot sign in until users have been granted access to the website by adding authorization rules.

Step 3: Configure a restrictive authorization rule

After Windows PowerShell Web Access is installed and the gateway is configured, users can open the sign-in page in a browser, but they cannot sign in until the Windows PowerShell Web Access administrator grants users access explicitly. Windows PowerShell Web Access access control is managed by using the set of Windows PowerShell cmdlets described in the following table. There is no comparable GUI for adding or managing authorization rules. For more detailed information about Windows PowerShell Web Access cmdlets, see the cmdlet reference topics, Windows PowerShell Web Access Cmdlets.

For more detail about Windows PowerShell Web Access authorization rules and security, see Authorization Rules and Security Features of Windows PowerShell Web Access.

To add a restrictive authorization rule

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then click **Run as Administrator**.
2. **Optional step for restricting user access by using session configurations:** Verify that session configurations that you want to use in your rules already exist. If they have not yet been created, use instructions for creating session configurations in about_Session_Configuration_Files on MSDN.
3. Type the following, and then press **Enter**.

   ```
   Add-PswaAuthorizationRule –UserName <domain\user |
   computer\user> -ComputerName <computer_name> -
   ConfigurationName <session_configuration_name>
   ```

   This authorization rule allows a specific user access to one computer on the network to which they typically have access, with access to a specific session configuration that is scoped to the user's typical scripting and cmdlet needs. In the following example, a user named `JSmith` in the `Contoso` domain is granted access to manage the computer `Contoso_214`, and use a session configuration named `NewAdminsOnly`.

   ```
   Add-PswaAuthorizationRule –UserName Contoso\JSmith -
   ComputerName Contoso_214 -ConfigurationName NewAdminsOnly
   ```

4. Verify that the rule has been created by running either the **Get-PswaAuthorizationRule** cmdlet, or **Test-PswaAuthorizationRule -UserName <domain\user | computer\user> -**

ComputerName <computer_name>. For example, **Test-PswaAuthorizationRule –UserName Contoso\JSmith –ComputerName Contoso_214**.

After you have configured an authorization rule, you are ready for authorized users to sign in to the web-based console and begin using Windows PowerShell Web Access.

Configure a genuine certificate

For a secure production environment, always use a valid SSL certificate that has been signed by a certification authority (CA). The procedure in this section describes how to obtain and apply a valid SSL certificate from a CA.

To configure an SSL certificate in IIS Manager

1. In the IIS Manager tree pane, select the server on which Windows PowerShell Web Access is installed.
2. In the content pane, double click **Server Certificates**.
3. In the **Actions** pane, do one of the following. For more information about configuring server certificates in IIS, see Configuring Server Certificates in IIS 7.
   - Click **Import** to import an existing, valid certificate from a location on your network.
   - Click **Create Certificate Request** to request a certificate from a CA such as VeriSign™, Thawte, or GeoTrust®. The certificate's common name must match the host header in the request. For example, if the client browser requests http://www.contoso.com/, then the common name must also be http://www.contoso.com/. This is the most secure and recommended option for providing the Windows PowerShell Web Access gateway with a certificate.
   - Click **Create a Self-Signed Certificate** to create a certificate that you can use immediately, and have signed later by a CA if desired. Specify a friendly name for the self-signed certificate, such as **Windows PowerShell Web Access**. This option is not considered secure, and is recommended only for a private test environment.
4. After creating or obtaining a certificate, select the website to which the certificate is applied (for example, **Default Web Site**) in the IIS Manager tree pane, and then click **Bindings** in the **Actions** pane.
5. In the **Add Site Binding** dialog box, add an **https** binding for the site, if one is not already displayed. If you are not using a self-signed certificate, specify the host name from step 3 of this procedure. If you are using a self-signed certificate, this step is not required.
6. Select the certificate that you obtained or created in step 3 of this procedure, and then click **OK**.

Using the web-based Windows PowerShell console

After Windows PowerShell Web Access is installed and the gateway configuration is finished as described in this topic, the Windows PowerShell web-based console is ready to use. For more information about getting started in the web-based console, see Use the Web-based Windows PowerShell Console.

See Also

Internet Information Services (IIS) 7.0 Documentation

IIS Manager 7.0 Help

Configure Web Server Security (IIS 7)

IPsec Deployment Resources

Authorization Rules and Security Features of Windows PowerShell Web Access

Windows PowerShell® Web Access in Windows Server® 2012 R2 and Windows Server® 2012 has a restrictive security model. Users must explicitly be granted access before they can sign in to the Windows PowerShell Web Access gateway and use the web-based Windows PowerShell console.

- Configuring authorization rules and site security
- Session management

Configuring authorization rules and site security

After Windows PowerShell Web Access is installed and the gateway is configured, users can open the sign-in page in a browser, but they cannot sign in until the Windows PowerShell Web Access administrator grants users access explicitly. Windows PowerShell Web Access access control is managed by using the set of Windows PowerShell cmdlets described in the following table. There is no comparable GUI for adding or managing authorization rules. For more detailed information about Windows PowerShell Web Access cmdlets, see the cmdlet reference topics, Windows PowerShell Web Access Cmdlets.

Administrators can define 0-*n* authentication rules for Windows PowerShell Web Access. The default security is restrictive rather than permissive; zero authentication rules means no users have access to anything.

Add-PswaAuthorizationRule and Test-PswaAuthorizationRule in Windows Server 2012 R2 include a Credential parameter that allows you to add and test Windows PowerShell Web Access authorization rules from a remote computer, or from within an active Windows PowerShell Web

Access session. As with other Windows PowerShell cmdlets that have a Credential parameter, you can specify a PSCredential object as the value of the parameter. To create a PSCredential object that contains credentials you want to pass to a remote computer, run the Get-Credential cmdlet.

Windows PowerShell Web Access authentication rules are whitelist rules. Each rule is a definition of an allowed connection between users, target computers, and particular Windows PowerShell session configurations (also referred to as endpoints or runspaces) on specified target computers.

Security

A user needs only one rule to be true to get access. If a user is given access to one computer with either full language access or access only to Windows PowerShell remote management cmdlets, from the web-based console, the user can log on (or hop) to other computers that are connected to the first target computer. The most secure way to configure Windows PowerShell Web Access is to allow users access only to constrained session configurations (also called endpoints or runspaces) that allow them to accomplish specific tasks that they normally need to perform remotely.

Name	Description	Parameters
Add-PswaAuthorizationRule	Adds a new authorization rule to the Windows PowerShell Web Access authorization rule set.	• ComputerGroupName • ComputerName • ConfigurationName • RuleName • UserGroupName • UserName • Credential (Windows Server 2012 R2 and later)
Remove-PswaAuthorizationRule	Removes a specified authorization rule from Windows PowerShell Web Access.	• Id • RuleName
Get-PswaAuthorizationRule	Returns a set of Windows PowerShell Web Access authorization rules. When it is used without parameters, the cmdlet returns all rules.	• Id • RuleName
Test-PswaAuthorizationRule	Evaluates authorization rules to determine if a specific user, computer, or	• ComputerName • ConfigurationName • RuleName

Name	Description	Parameters
	session configuration access request is authorized. By default, if no parameters are added, the cmdlet evaluates all authorization rules. By adding parameters, administrators can specify an authorization rule or a subset of rules to test.	• UserName • Credential (Windows Server 2012 R2 and later)

The preceding cmdlets create a set of access rules which are used to authorize a user on the Windows PowerShell Web Access gateway. The rules are different from the access control lists (ACLs) on the destination computer, and provide an additional layer of security for web access. More details about security are described in the following section.

If users cannot pass any of the preceding security layers, they receive a generic “access denied” message in their browser windows. Although security details are logged on the gateway server, end users are not shown information about how many security layers they passed, or at which layer the sign-in or authentication failure occurred.

For more information about configuring authorization rules, see Configuring authorization rules in this topic.

Security

The Windows PowerShell Web Access security model has four layers between an end user of the web-based console, and a target computer. Windows PowerShell Web Access administrators can add security layers through additional configuration in the IIS Manager console. For more information about securing websites in the IIS Manager console, see Configure Web Server Security (IIS 7). For more information about IIS best practices and preventing denial-of-service attacks, see Best Practices for Preventing DoS/Denial of Service Attacks. An administrator can also buy and install additional, retail authentication software.

The following table describes the four layers of security between end users and target computers.

Order	Layer	Description
1	Web Server (IIS) security features, such as client certificate authentication	Windows PowerShell Web Access users must always provide a user name and password to authenticate their accounts on the gateway. However, Windows PowerShell

Order	Layer	Description
		Web Access administrators can also turn optional client certificate authentication on or off (see step 10 of “To use IIS Manager to configure the gateway in an existing website” in Install and Use Windows PowerShell Web Access). The optional client certificate feature requires end users to have a valid client certificate, in addition to their user names and passwords, and is part of Web Server (IIS) configuration. When the client certificate layer is enabled, the Windows PowerShell Web Access sign-in page prompts users to provide valid certificates before their sign-in credentials are evaluated. Client certificate authentication automatically checks for the client certificate. If a valid certificate is not found, Windows PowerShell Web Access informs users, so they can provide the certificate. If a valid client certificate is found, Windows PowerShell Web Access opens the sign-in page for users to provide their user names and passwords. This is one example of additional security settings that are offered by Web Server (IIS). For more information about other IIS security features, see Configure Web Server Security (IIS 7).
2	Windows PowerShell Web Access forms-based gateway	The Windows PowerShell Web Access sign-in page requires a

Order	Layer	Description
	authentication	set of credentials (user name and password) and offers users the option of providing different credentials for the target computer. If the user does not provide alternate credentials, the primary user name and password that are used to connect to the gateway are also used to connect to the target computer. The required credentials are authenticated on the Windows PowerShell Web Access gateway. These credentials must be valid user accounts on either the local Windows PowerShell Web Access gateway server, or in Active Directory®. After a user is authenticated at the gateway, Windows PowerShell Web Access checks authorization rules to verify if the user has access to the requested target computer. After successful authorization, the user's credentials are passed along to the target computer.
3	Windows PowerShell Web Access authorization rules	After a user is authenticated at the gateway, Windows PowerShell Web Access checks authorization rules to verify if the user has access to the requested target computer. After successful authorization, the user's credentials are passed along to the target computer. These rules are evaluated only after a user has been

Order	Layer	Description
		authenticated by the gateway, and before a user can be authenticated on a target computer.
4	Target authentication and authorization rules	The final layer of security for Windows PowerShell Web Access is the target computer's own security configuration. Users must have the appropriate access rights configured on the target computer, and also in the Windows PowerShell Web Access authorization rules, to run a Windows PowerShell web-based console that affects a target computer through Windows PowerShell Web Access. This layer offers the same security mechanisms that would evaluate connection attempts if users tried to create a remote Windows PowerShell session to a target computer from within Windows PowerShell by running the **Enter-PSSession** or **New-PSSession** cmdlets. By default, Windows PowerShell Web Access uses the primary user name and password for authentication on both the gateway and the target computer. The web-based sign-in page, in a section titled **Optional connection settings**, offers users the option of providing different credentials for the target computer, if they are required. If the user does not provide alternate

Order	Layer	Description
		credentials, the primary user name and password that are used to connect to the gateway are also used to connect to the target computer. Authorization rules can be used to allow users access to a particular session configuration. You can create restricted runspaces or session configurations for Windows PowerShell Web Access, and allow specific users to connect only to specific session configurations when they sign in to Windows PowerShell Web Access. You can use access control lists (ACLs) to determine which users have access to specific endpoints, and further restrict access to the endpoint for a specific set of users by using authorization rules described in this section. For more information about restricted runspaces, see Constrained Runspaces on MSDN.

Configuring authorization rules

Administrators likely want the same authorization rule for Windows PowerShell Web Access users that is already defined in their environment for Windows PowerShell remote management. The first procedure in this section describes how to add a secure authorization rule that grants access to one user, signing in to manage one computer, and within a single session configuration. The second procedure describes how to remove an authorization rule that is no longer needed.

If you plan to use custom session configurations to allow specific users to work only within restricted runspaces in Windows PowerShell Web Access, create your custom session configurations before you add authorization rules that refer to them. You cannot use the Windows

PowerShell Web Access cmdlets to create custom session configurations. For more information about creating custom session configurations, see about_Session_Configuration_Files on MSDN.

Windows PowerShell Web Access cmdlets support one wildcard character, an asterisk (*). Wildcard characters within strings are not supported; use a single asterisk per property (users, computers, or session configurations).

Note
For more ways you can use authorization rules to grant access to users and help secure the Windows PowerShell Web Access environment, see Other authorization rule scenario examples in this topic.

To add a restrictive authorization rule

1. Do one of the following to open a Windows PowerShell session with elevated user rights.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then click **Run as Administrator**.
2. **Optional step for restricting user access by using session configurations:** Verify that session configurations that you want to use in your rules already exist. If they have not yet been created, use instructions for creating session configurations in about_Session_Configuration_Files on MSDN.
3. Type the following, and then press **Enter**.

   ```
   Add-PswaAuthorizationRule –UserName <domain\user |
   computer\user> -ComputerName <computer_name> -
   ConfigurationName <session_configuration_name>
   ```

   This authorization rule allows a specific user access to one computer on the network to which they typically have access, with access to a specific session configuration that is scoped to the user's typical scripting and cmdlet needs. In the following example, a user named `JSmith` in the `Contoso` domain is granted access to manage the computer `Contoso_214`, and use a session configuration named `NewAdminsOnly`.

   ```
   Add-PswaAuthorizationRule –UserName Contoso\JSmith -
   ComputerName Contoso_214 -ConfigurationName NewAdminsOnly
   ```

4. Verify that the rule has been created by running either the **Get-PswaAuthorizationRule** cmdlet, or **Test-PswaAuthorizationRule -UserName <domain\user | computer\user> -ComputerName** <computer_name>. For example, **Test-PswaAuthorizationRule –UserName Contoso\JSmith –ComputerName Contoso_214**.

To remove an authorization rule

1. If a Windows PowerShell session is not already open, see step 1 of To add a restrictive authorization rule in this section.
2. Type the following, and then press **Enter**, where *rule ID* represents the unique ID number

of the rule that you want to remove.

```
Remove-PswaAuthorizationRule -ID <rule ID>
```

Alternatively, if you do not know the ID number, but know the friendly name of the rule you want to remove, you can get the name of the rule, and pipe it to the `Remove-PswaAuthorizationRule` cmdlet to remove the rule, as shown in the following example: Get-PswaAuthorizationRule -RuleName <*rule name*> | Remove-PswaAuthorizationRule.

Note

You are not prompted to confirm whether you want to delete the specified authorization rule; the rule is deleted when you press **Enter**. Be sure that you want to remove the authorization rule before running the **Remove-PswaAuthorizationRule** cmdlet.

Other authorization rule scenario examples

Every Windows PowerShell session uses a session configuration; if one is not specified for a session, Windows PowerShell uses the default, built-in Windows PowerShell session configuration, called Microsoft.PowerShell. The default session configuration includes all cmdlets that are available on a computer. Administrators can restrict access to all computers by defining a session configuration with a restricted runspace (a limited range of cmdlets and tasks that their end users could perform). A user who is granted access to one computer with either full language access or only the Windows PowerShell remote management cmdlets can connect to other computers that are connected to the first computer. Defining a restricted runspace can prevent users from accessing other computers from their allowed Windows PowerShell runspace, and improves the security of your Windows PowerShell Web Access environment. The session configuration can be distributed (by using Group Policy) to all computers that administrators want to make accessible through Windows PowerShell Web Access. For more information about session configurations, see about_Session_Configurations. The following are some examples of this scenario.

- An administrator creates an endpoint, called **PswaEndpoint**, with a restricted runspace. Then, the administrator creates a rule, ***,*,PswaEndpoint**, and distributes the endpoint to other computers. The rule allows all users to access all computers with the endpoint **PswaEndpoint**. If this is the only authorization rule defined in the rule set, computers without that endpoint would not be accessible.
- The administrator created an endpoint with a restricted runspace called **PswaEndpoint**,and wants to restrict access to specific users. The administrator creates a group of users called **Level1Support**, and defines the following rule: **Level1Support,*,PswaEndpoint**. The rule grants any users in the group **Level1Support** access to all computers with the **PswaEndpoint** configuration. Similarly, access can be restricted to a specific set of computers.
- Some administrators provide certain users more access than others. For example, an administrator creates two user groups, **Admins** and **BasicSupport**. The administrator also creates an endpoint with a restricted runspace called **PswaEndpoint**, and defines the following two rules: **Admins,*,*** and **BasicSupport,*,PswaEndpoint**. The first rule provides

all users in the **Admin** group access to all computers, and the second rule provides all users in the **BasicSupport** group access only to those computers with **PswaEndpoint**.

- An administrator has set up a private test environment, and wants to allow all authorized network users access to all computers on the network to which they typically have access, with access to all session configurations to which they typically have access. Because this is a private test environment, the administrator creates an authorization rule that is not secure. The administrator runs the cmdlet `Add-PswaAuthorizationRule * * *`, which uses the wildcard character ***** to represent all users, all computers, and all configurations. This rule is the equivalent of the following: `Add-PswaAuthorizationRule –UserName * -ComputerName * -ConfigurationName *`.

  Security

  This rule is not recommended in a secure environment, and bypasses the authorization rule layer of security provided by Windows PowerShell Web Access.

- An administrator must allow users to connect to target computers in an environment that includes both workgroups and domains, where workgroup computers are occasionally used to connect to target computers in domains, and computers in domains are occasionally used to connect to target computers in workgroups. The administrator has a gateway server, *PswaServer*, in a workgroup; and target computer *srv1.contoso.com* is in a domain. User *Chris* is an authorized local user on both the workgroup gateway server and the target computer. His user name on the workgroup server is *chrisLocal*; and his user name on the target computer is *contoso\chris*. To authorize access to srv1.contoso.com for Chris, the administrator adds the following rule.

  ```
  Add-PswaAuthorizationRule –userName PswaServer\chrisLocal –
  computerName srv1.contoso.com –configurationName
  Microsoft.PowerShell
  ```

  The preceding rule example authenticates Chris on the gateway server, and then authorizes his access to *srv1*. On the sign-in page, Chris must provide a second set of credentials in the **Optional connection settings** area (*contoso\chris*). The gateway server uses the additional set of credentials to authenticate him on the target computer, *srv1.contoso.com*.

  In the preceding scenario, Windows PowerShell Web Access establishes a successful connection to the target computer only after the following have been successful, and allowed by at least one authorization rule.

  a. Authentication on the workgroup gateway server by adding a user name in the format *server_name\user_name* to the authorization rule
  b. Authentication on the target computer by using alternate credentials provided on the sign-in page, in the **Optional connection settings** area

  Note

  If gateway and target computers are in different workgroups or domains, a trust relationship must be established between the two workgroup computers, the two domains, or between the workgroup and the domain. This relationship cannot be configured by using Windows PowerShell Web Access authorization rule cmdlets. Authorization rules do not define a trust relationship between computers; they can

only authorize users to connect to specific target computers and session configurations. For more information about how to configure a trust relationship between different domains, see Creating Domain and Forest Trusts. For more information about how to add workgroup computers to a trusted hosts list, see Remote Management with Server Manager.

Using a single set of authorization rules for multiple sites

Authorization rules are stored in an XML file. By default, the path name of the XML file is %windir%\Web\PowershellWebAccess\data\AuthorizationRules.xml.

The path to the authorization rules XML file is stored in the **powwa.config** file, which is found in %windir%\Web\PowershellWebAccess\data. The administrator has the flexibility to change the reference to the default path in **powwa.config** to suit preferences or requirements. Allowing the administrator to change the location of the file lets multiple Windows PowerShell Web Access gateways use the same authorization rules, if such a configuration is desired.

Session management

By default, Windows PowerShell Web Access limits a user to three sessions at one time. You can edit the web application's **web.config** file in IIS Manager to support a different number of sessions per user. The path to the **web.config** file is $Env:Windir\Web\PowerShellWebAccess\wwwroot\Web.config.

By default, Web Server (IIS) is configured to restart the application pool if any settings are edited. For example, the application pool is restarted if changes are made to the **web.config** file. Because Windows PowerShell Web Access uses in-memory session states, users signed in to Windows PowerShell Web Access sessions lose their sessions when the application pool is restarted.

Setting default parameters on the sign-in page

If your Windows PowerShell Web Access gateway is running on Windows Server 2012 R2, you can configure default values for the settings that are displayed on the Windows PowerShell Web Access sign-in page. You can configure values in the **web.config** file that is described in the preceding paragraph. Default values for the sign-in page settings are found in the **appSettings** section of the web.config file; the following is an example of the **appSettings** section. Valid values for many of these settings are the same as those for the corresponding parameters of the New-PSSession cmdlet in Windows PowerShell. For example, the `defaultApplicationName` key, as shown in the following code block, is the value of the **$PSSessionApplicationName** preference variable on the target computer.

```
<appSettings>

        <add key="maxSessionsAllowedPerUser" value="3"/>

        <add key="defaultPortNumber" value="5985"/>
```

```
        <add key="defaultSSLPortNumber" value="5986"/>
        <add key="defaultApplicationName" value="WSMAN"/>
        <add key="defaultUseSslSelection" value="0"/>
        <add key="defaultAuthenticationType" value="0"/>
        <add key="defaultAllowRedirection" value="0"/>
        <add key="defaultConfigurationName" value="Microsoft.PowerShell"/>
</appSettings>
```

Time-outs and unplanned disconnections

Windows PowerShell Web Access sessions time out. In Windows PowerShell Web Access running on Windows Server 2012, A time-out message is displayed to signed-in users after 15 minutes of session inactivity. If the user does not respond within five minutes after the time-out message is displayed, the session is ended, and the user is signed out. You can change time-out periods for sessions in the website settings in IIS Manager.

In Windows PowerShell Web Access running on Windows Server 2012 R2, sessions time out, by default, after 20 minutes of inactivity. If users are disconnected from sessions in the web-based console because of network errors or other unplanned shutdowns or failures, and not because they have closed the sessions themselves, the Windows PowerShell Web Access sessions continue to run, connected to target computers, until the time-out period on the client side lapses. The session is disconnected after either the default 20 minutes, or after the time-out period specified by the gateway administrator, whichever is shorter.

If the gateway server is running Windows Server 2012 R2, Windows PowerShell Web Access lets users reconnect to saved sessions at a later time, but when network errors, unplanned shutdowns, or other failures disconnect sessions, users cannot see or reconnect to saved sessions until after the time-out period specified by the gateway administrator has lapsed.

See Also

Install and Use Windows PowerShell Web Access

about_Session_Configurations

Windows PowerShell Web Access Cmdlets

Use the Web-based Windows PowerShell Console

Windows PowerShell® Web Access lets Windows PowerShell® users sign in to a Secure Sockets Layer (SSL)-secured website to use Windows PowerShell sessions, cmdlets, and scripts to manage a remote computer. Because the Windows PowerShell console runs in a web

browser, it can be opened from a variety of client devices, including cell phones, tablet computers, public computing kiosks, laptop computers, or shared or borrowed computers. The web-based Windows PowerShell console is targeted at a remote computer that is specified by users as part of the sign-in process. This topic describes how to sign in to and start using the Windows PowerShell Web Access web-based console.

This topic does not describe how to use Windows PowerShell or run cmdlets or scripts. For information about how to use Windows PowerShell, and scripting resources, see the See Also section at the end of this topic.

In this topic:

- **Supported browsers and client devices**
- Signing in to Windows PowerShell Web Access
- Signing out and timing out
- Differences in the web-based Windows PowerShell console
- See Also

Supported browsers and client devices

Windows PowerShell Web Access supports the following Internet browsers. Although mobile browsers are not officially supported, many may be able to run the web-based Windows PowerShell console. Other browsers that accept cookies, run JavaScript, and run HTTPS websites are expected to work, but are not officially tested.

Supported desktop computer browsers

- Windows® Internet Explorer® for Microsoft Windows® 8.0, 9.0, 10.0, and 11.0
- Mozilla Firefox® 10.0.2
- Google Chrome™ 17.0.963.56m for Windows
- Apple Safari® 5.1.2 for Windows
- Apple Safari 5.1.2 for Mac OS®

Minimally-tested mobile devices or browsers

- Windows Phone 7 and 7.5
- Google Android WebKit 3.1 Browser Android 2.2.1 (Kernel 2.6)
- Apple Safari for iPhone operating system 5.0.1
- Apple Safari for iPad 2 operating system 5.0.1

Browser requirements

To use the Windows PowerShell Web Access web-based console, browsers must do the following.

- Allow cookies from the Windows PowerShell Web Access gateway website.

- Be able to open and read HTTPS pages.
- Open and run websites that use JavaScript.

Signing in to Windows PowerShell Web Access

Your Windows PowerShell Web Access administrator should provide you with a URL that is the address of your organization's Windows PowerShell Web Access gateway website. By default, this website address is https://<server_name>/pswa. Before you sign in to Windows PowerShell Web Access, be sure that you have the name or IP address of the remote computer that you want to manage. You must be an authorized user on the remote computer, and it must be configured to allow remote management. For more information about configuring your computer to allow remote management, see Enable and Use Remote Commands in Windows PowerShell. The simplest method of configuring your computer to allow remote management is to run the **Enable-PSRemoting -force** cmdlet on the computer, in a Windows PowerShell session that has been opened with elevated user rights (**Run as Administrator**).

To sign in to Windows PowerShell Web Access

1. Open the Windows PowerShell Web Access website in an Internet browser window or tab.
2. On the Windows PowerShell Web Access sign-in page, provide your network user name, password, and the name of the computer that you want to manage (and on which you are an authorized user). If the Windows PowerShell Web Access administrator has instructed you to use a URI to a custom site or proxy server instead of a computer name, select **Connection URI** in the **Connection type** field, and then provide the URI.

   Notes
   - If the destination computer is in a workgroup, use the following syntax to provide your user name and sign in to the computer:<*workgroup_name*>\<*user_name*>.
   - If the destination computer is the gateway server, you can specify **localhost** in the **Computer name** field.
   - If the destination computer is the gateway server, and the gateway server is in a workgroup, you can use **localhost** in the **Computer name** field, but do not use localhost\<*user_name*> in the **User name** field. You must use <*workgroup name*>\<*user_name*>.
3. The **Optional Connection Settings** section relates to the authorization requirements of the remote computer that you want to manage. For more information about the parameters that are equivalent to optional connection settings, see the Enter-PSSession cmdlet Help.

   Typically, the credentials you use to pass through the Windows PowerShell Web Access gateway are the same that are recognized by the remote computer that you want to manage. However, if you want to use different credentials to manage the remote computer that you specified in step 2, expand the **Optional Connection Settings** section, and provide the alternate credentials. Otherwise, skip to step 6.

4. If the Windows PowerShell Web Access administrator has created a custom session configuration for Windows PowerShell Web Access users, type the name of the session configuration name in the **Configuration name** field. For more information about session configurations, see about_Session_Configurations on the Microsoft website.
5. Keep the **Authentication type** set to **Default** unless you have been instructed to do otherwise by the Windows PowerShell Web Access administrator.
6. Click **Sign in**.

Signing out and timing out

Any of the following signs you out of a web-based Windows PowerShell session.

- Clicking **Sign out** in the lower right corner of the console. (Windows Server 2012 only)
- Clicking **Save** or **Exit** in the lower right corner of the console (Windows Server 2012 R2 only). Clicking **Save** saves and closes your Windows PowerShell Web Access session; you can reconnect to the session later. When you sign in to Windows PowerShell Web Access again, Windows PowerShell Web Access displays a list of your saved sessions; you can either select and reconnect to a saved session, or start a new session. The maximum number of open sessions that users are allowed, both saved and active, is configured by the gateway administrator.

  Clicking **Exit** signs you out of the Windows PowerShell Web Access session without saving it.
- Attempting to sign in to manage a different remote computer in the same browser session, or in a new tab of the same browser session. (This does not apply if the gateway server is running Windows Server 2012 R2; Windows PowerShell Web Access running on Windows Server 2012 R2 does allow multiple user sessions in new tabs in the same browser session.) For more information about how to use more than one active session on the same computer, see “Connecting to multiple target computers simultaneously” in the Limitations of the web-based console section of this topic.
- 20 minutes of inactivity in the session. The gateway administrator can customize the inactivity time-out period; for more information, see Session management.
  - If you are disconnected from a session in the web-based console because of a network error or other unplanned shutdown or failure, and not because you have closed the session yourself, the Windows PowerShell Web Access session continues to run, connected to the target computer, until the time-out period on the client side lapses. By default, this time-out period is 20 minutes, and is configured by the gateway administrator. The session is disconnected after either the default 20 minutes, or after the time-out period specified by the gateway administrator, whichever is shorter.

    If the gateway server is running Windows Server 2012 R2, Windows PowerShell Web Access lets users reconnect to saved sessions at a later time, but you cannot see or reconnect to saved sessions until after the time-out period specified by the gateway administrator has lapsed.
- Closing the browser window or tab.

- Turning off the client device on which the browser is running, or disconnecting it from the network.
- Running the **Exit** command in the web console. This command does not work if the session configuration to which you are connected to is configured to support NoLanguage mode, or is in a restricted runspace.

If you want to sign in again, open the Windows PowerShell Web Access web page again, and sign in by following the steps in To sign in to Windows PowerShell Web Access in this topic.

Differences in the web-based Windows PowerShell console

After signing in to Windows PowerShell Web Access, a web-based Windows PowerShell console opens in your browser window or tab. Because the console is connected to the remote computer that you specified during the sign-in process, only those Windows PowerShell cmdlets or scripts that are available on the remote computer can be used in the console. This section describes other limitations of Windows PowerShell Web Access consoles, and differences between Windows PowerShell Web Access consoles and the installed **PowerShell.exe** console.

Functional disparity with PowerShell.exe

The majority of Windows PowerShell host functionality is available in the Windows PowerShell Web Access web-based console, but there are some features that are not available.

- **Nested progress displays.** Windows PowerShell Web Access displays a progress GUI for cmdlets that report progress, but only top-level progress information is displayed.
- **Input color modification.** The input color (both foreground and background) cannot be changed. The style of output, warning, verbose, and error messages can all be changed by running a script.
- **PSHostRawUserInterface.** Windows PowerShell Web Access is implemented over Windows PowerShell remote management, and uses a remote runspace. Windows PowerShell Web Access does not implement some methods in this interface; for example, any command that writes to the Windows console. Commands such as **PowerTab** do not work in Windows PowerShell Web Access.
- **Function keys.** Windows PowerShell Web Access does not support some function keys, in many cases because the commands are reserved by the browser.

Unsupported Function Key	Action
Ctrl+C	In Windows PowerShell Web Access, **Ctrl+C** is used by the browser to copy content. The console offers a **Cancel** button, and users can also use **Ctrl+Q** to cancel commands.
Alt-space, e, l	Scroll through the screen buffer

Unsupported Function Key	Action
Alt+Space, e, f	Search for text in the screen buffer
Alt+Space, e, k	Select text to be copied from the screen buffer
Alt+Space, e, p	Paste clipboard contents into the Windows PowerShell console
Alt+Space, c	Close the Windows PowerShell console
Ctrl+Break	Force the Windows PowerShell window to close
Ctrl+Home	Deletes from the beginning of the current command line
Ctrl+End	Deletes to end of the command line
F1	Move cursor one character to the right on your command line
F2	Creates a new command by copying your last command up to the character that you type
F3	Complete the command line with content from your last command line
F4	Deletes characters from cursor position
F5	Scan backward through your command history. To access commands in the command history in Windows PowerShell Web Access, click the **History** scroll buttons in the web-based console.
F7	Interactively select a command from your command history
F8	Scan history displaying commands that match current text
F9	Run a specific numbered command from history
Page Up	Run the first command in the history
Page Down	Run the last command in the history
Alt+F7	Clear the command history list

Limitations of the web-based console

- **Double-hop.** You can encounter the double-hop (or connecting to a second computer from the first connection) limitation if you try to create or work on a new session by using Windows PowerShell Web Access. Windows PowerShell Web Access uses a remote runspace, and currently, **PowerShell.exe** does not support establishing a remote connection to a second computer from a remote runspace. If you attempt to connect to a second remote computer from an existing connection by using the **Enter-PSSession** cmdlet, for example, you can get various errors, such as "Cannot get network resources."

  To avoid double-hop errors, your administrator should configure CredSSP authentication in your organization's network environment. For more information about configuring CredSSP authentication, see CredSSP for second-hop remoting on the Microsoft website. You can also provide explicit credentials when you want to manage a second remote computer; implicit credentials are unlikely to allow the second hop.
- Windows PowerShell Web Access uses and has the same limitations as a remote Windows PowerShell session. Commands that directly call Windows console APIs, such as those for console-based editors or text-based menu programs, do not work because the commands do not read or write to standard input, output, and error pipes. Therefore, commands that launch an executable file, such as **notepad.exe**, or display a GUI, such as `OpenGridView` or `ogv`, do not work. Your experience is affected by this behavior; to you, it appears that Windows PowerShell Web Access is not responding to your command.
- Tab completion does not work in a session configuration with a restricted runspace or one that is in **NoLanguage** mode. Although administrators can configure a session to support tab completion, it is discouraged for security reasons, because it can expose the following information to unauthorized users.
  - Internal file system paths
  - Shared folders on internal computers
  - Variables in the runspace
  - Loaded types or.NET Framework namespaces
  - Environment variables
- Users who are signed in to a **NoLanguage** session configuration or a restricted runspace in Windows PowerShell Web Access cannot run the **Exit** command to end the session. To sign out, users should click **Sign Out** on the console page.
- **Connecting to multiple target computers simultaneously.** If the gateway server is running Windows Server 2012, Windows PowerShell Web Access allows only one remote computer connection per browser session; it does not allow users to sign in once, and connect to multiple remote computers by using separate browser tabs. When you open a new tab or new browser window, Windows PowerShell Web Access prompts you to disconnect your current session and start a new session, so that you can connect to the new (or the same) remote computer. If two or more separate sessions to different remote computers are desired, however, a feature in Internet Explorer lets you create a new session. To start a new browser session in Internet Explorer, press **ALT**, open the **File** menu, and then select **New Session**. Then, open the Windows PowerShell Web Access website in the new session, and sign in to access another remote computer.

When the Windows PowerShell Web Access gateway is running on Windows Server 2012 R2, users can open multiple connections to remote computers in different browser tabs. If you want to open more than one connection to a remote computer by using the web-based Windows PowerShell console, check with your Windows PowerShell Web Access gateway administrator to see if this feature is supported by the gateway server.

- **Persistent Windows PowerShell sessions (Reconnection).** After you time out of the Windows PowerShell Web Access gateway, the remote connection between the gateway and the target computer is closed. This stops any cmdlets or scripts that are currently in process. You are encouraged to use the Windows PowerShell **-Job** infrastructure when you are performing long-running tasks, so that you can start jobs, disconnect from the computer, reconnect later, and have jobs persist. Another benefit of using **-Job** cmdlets is that you can start them by using Windows PowerShell Web Access, sign out, and then reconnect later, either by running Windows PowerShell Web Access or another host (such as Windows PowerShell® Integrated Scripting Environment (ISE)).
- **Console resizing.** The **PowerShell.exe** console window can be resized in the following three ways.
  - Drag and adjust the console window size with a mouse
  - Change the height and width properties by using a GUI for console properties
  - Changing the height and width of console windows with a cmdlet

    The console window for Windows PowerShell Web Access can be configured by using the cmdlets as follows. In the following example, a user changes the width of Windows PowerShell Web Access console to **20**.

```
$newSize = $Host.UI.RawUI.WindowSize
$newSize.Width = $newSize.Width - 20


$oldSize = $Host.UI.RawUI.WindowSize


$Host.UI.RawUI.WindowSize = $newSize
```

    You can change the height of the console in a similar manner.

    Additional examples for customizing the console view are available in the Windows PowerShell Team Blog.

See Also

Windows PowerShell Cmdlet Reference

Windows PowerShell on Microsoft TechNet

TechNet Script Center Repository

Script Center - Hey, Scripting Guy!

Windows PowerShell Team Blog

Troubleshooting Access Problems in Windows PowerShell Web Access

Troubleshooting access problems

The following table identifies some common problems that users might experience when they are attempting to connect to a remote computer by using Windows PowerShell Web Access, and includes suggestions for resolving the problems.

Problem	Possible cause and solution
Sign-in failure	Failure could occur because of any of the following. • An authorization rule that allows the user access to the computer, or a specific session configuration on the remote computer, does not exist. Windows PowerShell Web Access security is restrictive; users must be granted explicit access to remote computers by using authorization rules. For more information about creating authorization rules, see Authorization Rules and Security Features of Windows PowerShell Web Access in this guide. • The user does not have authorized access to the destination computer. This is determined by access control lists (ACLs). For more information, see "Signing in to Windows PowerShell Web Access" in Use the Web-based Windows PowerShell Console, or the Windows PowerShell Team Blog. • Windows PowerShell remote management might not be enabled on the destination computer. Verify that it is enabled on the computer to which the user is trying to connect. For more information, see "How to Configure Your Computer for Remoting" in about_Remote_Requirements in the Windows PowerShell About Help Topics.

When users try to sign in to Windows PowerShell Web Access in an Internet Explorer window, they are shown an **Internal Server Error** page, or Internet Explorer stops	This can occur for users who have signed in with a domain name that contains Chinese characters, or if one or more Chinese characters are part of the gateway server name. To work around this issue, the user should install and run Internet Explorer 10, and then perform the following steps.
	1. Change the Internet Explorer **Document Mode** setting to **IE10 standards**.
	  a. Press **F12** to open the Developer Tools console.
	  b. In Internet Explorer 10, click **Browser Mode**, and then select **Internet Explorer 10**.
	  c. Click **Document Mode**, and then click **IE10 standards**.
	  d. Press **F12** again to close the Developer Tools console.

<table>
<tr><th>Problem</th><th>Possible cause and solution</th></tr>
<tr><td>responding. This issue is specific to Internet Explorer.</td><td>2. Disable automatic proxy configuration.
a. In Internet Explorer 10, click Tools, and then click Internet Options.
b. In the Internet Options dialog box, on the Connections tab, click LAN settings.
c. Clear the Automatically detect settings check box. Click OK, and then click OK again to close the Internet Options dialog box.</td></tr>
<tr><td>Cannot connect to a remote workgroup computer</td><td>If the destination computer is a member of a workgroup, use the following syntax to provide your user name and sign in to the computer: <workgroup_name>\<user_name></td></tr>
<tr><td>Cannot find Web Server (IIS) management tools, even though the role was installed</td><td>If you installed Windows PowerShell Web Access by using the <code>Install-WindowsFeature</code> cmdlet, management tools are not installed unless the <code>IncludeManagementTools</code> parameter is added to the cmdlet. For an example, see “To install Windows PowerShell Web Access by using Windows PowerShell cmdlets” in Install and Use Windows PowerShell Web Access. You can add the IIS Manager console and other IIS management tools that you need by selecting the tools in an Add Roles and Features Wizard session that is targeted at the gateway server. The Add Roles and Features Wizard is opened from within Server Manager.</td></tr>
<tr><td>The Windows PowerShell Web Access website is not accessible</td><td>If Enhanced Security Configuration is enabled in Internet Explorer (IE ESC), you can add the Windows PowerShell Web Access website to the list of trusted sites, or disable IE ESC. You can disable IE ESC in the Properties tile on the Local Server page in Server Manager.</td></tr>
<tr><td>The following error message is displayed while trying to connect when the gateway server is the destination computer, and is also in a workgroup: An authorization failure occurred. Verify that you are authorized to connect to the</td><td>When the gateway server is also the destination server, and it is in a workgroup, specify the user name, computer name, and user group name as shown in the following table. Do not use a dot (.) by itself to represent the computer name.
<table>
<tr><th>Scenario</th><th>UserName Parameter</th><th>UserGroup Parameter</th><th>ComputerName Parameter</th><th>ComputerGroup Parameter</th></tr>
<tr><td>Gateway server is in a domain</td><td>Server_name\user_name, Localhost\user_name, or .\user_name</td><td>Server_name\user_group, Localhost\user_group, or .\user_group</td><td>Fully qualified name of gateway server, or Localhos</td><td>Server_name\computer_group, Localhost\computer_group, or .\computer_group</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Problem</th><th>Possible cause and solution</th></tr>
<tr><td>destination computer.</td><td>
<table>
<tr><td></td><td></td><td></td><td>t</td><td></td></tr>
<tr><td>Gateway server is in a workgroup</td><td>Server_name\user_name, Localhost\user_name, or .\user_name</td><td>Server_name\user_group, Localhost\user_group or .\user_group</td><td>Server name</td><td>Server_name\computer_group, Localhost\computer_group or .\computer_group</td></tr>
</table>

Sign in to a gateway server as target computer by using credentials formatted as one of the following.
• Server_name\user_name
• Localhost\user_name
• .\user_name
</td></tr>
<tr><td>A security identifier (SID) is displayed in an authorization rule instead of the syntax user_name/computer_name</td><td>Either the rule is no longer valid, or the Active Directory Domain Services query failed. An authorization rule is usually not valid in scenarios where the gateway server was at one time in a workgroup, but was later joined to a domain.</td></tr>
<tr><td>Cannot sign in to a target computer that has been specified in authorization rules as an IPv6 address with a domain.</td><td>Authorization rules do not support an IPv6 address in form of a domain name. To specify a destination computer by using an IPv6 address, use the original IPv6 address (that contains colons) in the authorization rule. Both domain and numerical (with colons) IPv6 addresses are supported as the target computer name on the Windows PowerShell Web Access sign-in page, but not in authorization rules. For more information about IPv6 addresses, see How IPv6 Works.</td></tr>
</table>

See Also

Authorization Rules and Security Features of Windows PowerShell Web Access

Use the Web-based Windows PowerShell Console

about_Remote_Requirements

Uninstall Windows PowerShell Web Access

Follow steps in this topic to delete the Windows PowerShell Web Access website and application from the gateway server that is running either Windows Server 2012 R2 or Windows Server 2012. Before you begin, notify users of the web-based console that you are removing the website.

Uninstalling Windows PowerShell Web Access

Before you uninstall Windows PowerShell Web Access from the gateway server, either run the `Uninstall-PswaWebApplication` cmdlet to remove the website and Windows PowerShell Web Access web applications, or use the IIS Manager procedure, To delete the Windows PowerShell Web Access website and web applications by using IIS Manager.

Uninstalling Windows PowerShell Web Access does not uninstall IIS or any other features that were installed automatically because Windows PowerShell Web Access requires them to run. The uninstallation process leaves features installed upon which Windows PowerShell Web Access was dependent; you can uninstall those features separately if needed.

Recommended (quick) uninstallation

Procedures in this section help you uninstall both the Windows PowerShell Web Access web application and the Windows PowerShell Web Access feature by using Windows PowerShell cmdlets.

Step 1: Delete the web application

If you have specified your own, custom website name, add the `WebsiteName` parameter to your command, and specify the website name. If you have used a custom web application (not the default application, **pswa**, add the `WebApplicationName` parameter to your command, and specify the name of the web application.

To delete the website and web applications by using the Uninstall-PswaWebApplication cmdlet

1. Do one of the following to open a Windows PowerShell session.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar.
   - On the Windows **Start** screen, click **Windows PowerShell**.
2. Type **Uninstall-PswaWebApplication**, and then press **Enter**.
3. If you are using a test certificate, add the `DeleteTestCertificate` parameter to the cmdlet, as shown in the following example.

```
Uninstall-PswaWebApplication -DeleteTestCertificate
```

Step 2: Uninstall Windows PowerShell Web Access

To uninstall Windows PowerShell Web Access by using Windows PowerShell cmdlets

1. Do one of the following to open a Windows PowerShell session with elevated user rights. If a session is already open, go on to the next step.
   - On the Windows desktop, right-click **Windows PowerShell** on the taskbar, and then click **Run as Administrator**.
   - On the Windows **Start** screen, right-click **Windows PowerShell**, and then click **Run as Administrator**.
2. Type the following, and then press **Enter**, where *computer_name* represents a remote server from which you want to remove Windows PowerShell Web Access. The `-Restart` parameter automatically restarts destination servers if required by the removal.

   ```
   Uninstall-WindowsFeature –Name WindowsPowerShellWebAccess -
   ComputerName <computer_name> -Restart
   ```

   To remove roles and features from an offline VHD, you must add both the `-ComputerName` parameter and the `-VHD` parameter. The `-ComputerName` parameter contains the name of the server on which to mount the VHD, and the `-VHD` parameter contains the path to the VHD file on the specified server.

   ```
   Uninstall-WindowsFeature –Name WindowsPowerShellWebAccess –
   VHD <path> -ComputerName <computer_name> -Restart
   ```
3. When removal is finished, verify that you removed Windows PowerShell Web Access by opening the **All Servers** page in Server Manager, selecting a server from which you removed the feature, and viewing the **Roles and Features** tile on the page for the selected server. You can also run the `Get-WindowsFeature` cmdlet targeted at the selected server (Get-WindowsFeature -ComputerName <*computer_name*>) to view a list of roles and features that are installed on the server.

Custom uninstallation

Procedures in this section help you uninstall both the Windows PowerShell Web Access web application and the Windows PowerShell Web Access feature by using the Remove Roles and Features Wizard in Server Manager, and the IIS Manager console.

Step 1: Delete the web application

To delete the Windows PowerShell Web Access website and web applications by using IIS Manager

1. Open the IIS Manager console by doing one of the following. If it is already open, go on to the next step.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar. On the **Tools** menu in Server Manager, click **Internet Information**

Services (IIS) Manager.

- On the Windows **Start** screen, type any part of the name **Internet Information Services (IIS) Manager**. Click the shortcut when it is displayed in the **Apps** results.

2. In the IIS Manager tree pane, select the website that is running the Windows PowerShell Web Access web application.
3. In the **Actions** pane, under **Manage Website**, click **Stop**.
4. In the tree pane, right-click the web application in the website that is running the Windows PowerShell Web Access web application, and then click **Remove**.
5. In the tree pane, select **Application Pools**, select the Windows PowerShell Web Access application pool folder, click **Stop** in the **Actions** pane, and then click **Remove** in the content pane.
6. Close IIS Manager.

Note

The certificate is not deleted during uninstallation. If you created a self-signed certificate or used a test certificate and want to remove it, delete the certificate in IIS Manager.

Step 2: Uninstall Windows PowerShell Web Access

To uninstall Windows PowerShell Web Access by using the Remove Roles and Features Wizard

1. If Server Manager is already open, go on to the next step. If Server Manager is not already open, open it by doing one of the following.
   - On the Windows desktop, start Server Manager by clicking **Server Manager** in the Windows taskbar.
   - On the Windows **Start** screen, click **Server Manager**.
2. On the **Manage** menu, click **Remove Roles and Features**.
3. On the **Select destination server** page, select the server or offline VHD from which you want to remove the feature. To select an offline VHD, first select the server on which to mount the VHD, and then select the VHD file. After you have selected the destination server, click **Next**.
4. Click **Next** again to skip to the **Remove features** page.
5. Clear the check box for **Windows PowerShell Web Access**, and then click **Next**.
6. On the **Confirm removal selections** page, click **Remove**.

See Also

Install and Use Windows PowerShell Web Access

IIS Manager 7.0 Help

Getting Started with Windows PowerShell Workflow

IT Pros often automate the management of their multi-device environments by running sequences of long-running tasks, or *workflows*, that can affect multiple managed computers or devices at the same time. Windows PowerShell Workflow lets IT Pros and developers leverage the benefits of Windows Workflow Foundation with the automation capabilities and ease of Windows PowerShell. Windows PowerShell Workflow functionality was introduced in Windows Server® 2012 and Windows 8, and is part of Windows PowerShell 3.0 and newer releases of Windows PowerShell. Windows PowerShell Workflow helps automate the distribution, orchestration, and completion of multi-device tasks, freeing users and administrators to focus on higher-level tasks.

Windows PowerShell, first introduced with the release of Windows Vista and Windows Server 2008, is a distributed automation engine, command-line shell, and scripting language that is constructed on the Microsoft® .NET Framework. It is designed especially for automating Windows management.

Windows PowerShell Workflow is a key feature of Windows PowerShell 3.0 and Windows PowerShell 4.0. For more information about Windows PowerShell, see What's New in Windows PowerShell.

This topic includes the following sections.

- Overview of Windows PowerShell Workflow
- How Windows PowerShell Workflow and Windows PowerShell scripts differ
- Creating and importing workflows by using the Visual Studio Workflow Designer

Overview of Windows PowerShell Workflow

A workflow is a sequence of programmed, connected steps that perform long-running tasks or require the coordination of multiple steps across multiple devices or managed nodes. Windows PowerShell Workflow lets IT pros and developers author sequences of multi-device management *activities*, or single tasks within a workflow, as workflows. By design, workflows can be long-running, repeatable, frequent, parallelizable, interruptible, stoppable, and restartable. They can be suspended and resumed; they can also continue after an unexpected interruption, such as a network outage or computer restart.

Windows PowerShell workflows can be authored or defined by using Windows PowerShell syntax, or XAML files. For more information about authoring a Windows PowerShell script-based workflow, see Writing a Script Workflow. For more information about XAML-based workflows, see Creating and importing workflows by using the Visual Studio Workflow Designer in this topic.

By using the `RunAs` functionality of Windows PowerShell, custom session configurations let delegated or subordinate IT Pros run workflows or activities within a workflow.

Activities

An *activity* is a specific task that you want a workflow to perform. Just as a script is composed of one or more commands, a workflow is composed of one or more activities that are carried out in a sequence. A script can also be used as a single command in another script, and a workflow can be used as an activity within another workflow.

Benefits of Windows PowerShell Workflow

The following list describes many of the benefits of Windows PowerShell Workflow.

- **Use Windows PowerShell scripting syntax**

  IT Pros can use Windows PowerShell scripting skills to author script-based workflows by using the extensible Windows PowerShell language. Windows PowerShell script-based workflows are easy to write, and can be shared easily by pasting them into an email, or publishing them on webpages.

- **Multi-device management**

  You can simultaneously apply workflow tasks to hundreds of managed nodes. Windows PowerShell Workflow adds common parameters to workflows automatically, such as `PSComputerName`, to enable multi-device management scenarios. For more information about parameters that are reserved for workflows, see “Reserved Words in Workflows” in Workflow Authoring Reference Topics.

- **Running a single task to manage complex, end-to-end processes**

  You can combine related scripts or commands that act on an entire scenario into a single workflow. Status and progress of activities within the workflow are visible at any time.

- **Automated failure recovery**

  Workflows survive both planned and unplanned interruptions, such as computer restarts. You can suspend workflow operation, then restart or resume the workflow from the point at which it was suspended. You can author checkpoints as part of your workflow, so that you can resume the workflow from the last persisted task (or checkpoint), instead of restarting the workflow from the beginning.

- **Connection and activity retries**

  By using workflow common parameters, workflow users can retry connections to managed nodes if network connection failures occur. Workflow authors can also specify activities that must run again if the activity cannot be completed on one or more managed nodes (for example, if a target computer was offline while the activity was running).

- **Connect and disconnect**

  Users can connect and disconnect from the computer that is running the workflow, but the workflow remains running. For example, if you are running the workflow and managing the workflow on two different computers, you can log off of or restart the computer from which you are managing the workflow, and monitor workflow operations from another computer (such as a home computer) without interrupting the workflow.

- **Task scheduling**

Workflow tasks can be scheduled, and started when specific conditions are met, as with any other Windows PowerShell cmdlet or script.

How Windows PowerShell Workflow and Windows PowerShell scripts differ

In a script, the entire script runs in the same *runspace*, the operating environment that defines which commands, variables, and other elements are available. In a workflow, each activity within the workflow can run in a different runspace. Variables that are created at the top level of the entire workflow are available throughout the workflow; if they are created at the script or command level, they are available for the command or script, but not available to the entire workflow.

In general, you should consider using a workflow instead of a cmdlet or script when you must meet any of the following requirements.

- You need to perform a long-running task that combines multiple steps in a sequence.
- You need to perform a task that runs on multiple devices.
- You need to perform a task that requires checkpointing or persistence.
- You need to perform a long-running task that is asynchronous, restartable, parallelizable, or interruptible.
- You need to run a task on a large scale, or in high availability environments, potentially requiring throttling and connection pooling.

Creating and importing workflows by using the Visual Studio Workflow Designer

In addition to authoring workflows that use Windows PowerShell scripts, you can author workflows by using the Visual Studio Workflow Designer. Workflows that you author in Visual Studio Workflow Designer have a XAML file name extension. You can invoke and import workflows in Windows PowerShell Workflow that are authored in Workflow Designer, and expose Windows PowerShell cmdlets as activities within the Workflow Designer.

For more information about authoring workflows in the Visual Studio Workflow Designer, see Using the Workflow Designer in the Visual Studio documentation on MSDN.

Windows PowerShell Workflow activities in the Visual Studio Workflow Designer

Windows PowerShell includes built-in activities for most Windows PowerShell cmdlets. These activities are packaged in assemblies with names that match the name of the Windows PowerShell module that contains the cmdlets. The following is the list of Windows PowerShell activity assemblies that are installed in the Global Assembly Cache (GAC).

- Microsoft.PowerShell.Activities

- Microsoft.PowerShell.Core.Activities
- Microsoft.PowerShell.Diagnostics.Activities
- Microsoft.PowerShell.Management.Activities
- Microsoft.PowerShell.Security.Activities
- Microsoft.PowerShell.Utility.Activities
- Microsoft.WSMan.Management.Activities

For more information about adding an activity to the Toolbox from an assembly, see How to: Add Activities to the Toolbox.

See Also

Configuring Your Workflow Environment

Running a Windows PowerShell Workflow

Writing a Script Workflow

What's New in Windows Workflow Foundation 4.0

about_Jobs

about_Workflow_Common_Parameters

about_Workflows

Writing a Script Workflow

The topics in this section explain how to write a workflow in the Windows PowerShell® scripting language; the same language you use to write functions and scripts in Windows PowerShell. You can run the workflows that you write at the command line, include them in scripts, and include them in script modules.

For information about writing workflows in XAML, the native language of Windows Workflow Foundation, see Writing a Windows PowerShell Workflow in the MSDN Library.

Planning a Workflow

Windows PowerShell Workflow is a powerful solution that can gather data from and make changes to hundreds of computers every time it runs. Plan the workflow design for simplicity and efficiency. Keep in mind that the workflow can suspend itself or be suspended by users, that it can be controlled from numerous PSSessions that connect to a single server-side process, that it can survive restarts, include commands and functions that run in parallel, and that each command runs its own session.

Use the following planning guidelines

- If you are new to workflows, review the reference topics in this section, including Syntactic Differences Between Script Workflows and Scripts.

- Begin by enumerating the tasks that the workflow will perform. Mark the sections that can run concurrently or do not need to run in a predetermined order. Within those sections, mark any that must run in sequential order.
- Organize tasks into functions within the workflow. For each task, use an existing Windows PowerShell cmdlet or create a new command and add a checkpoint after each critical step.
- Workflows are designed to run on multiple target computers. You do not need to create remote sessions or use remote commands, such as those that use the Invoke-Command cmdlet, in the workflow.
- As with any script that runs on multiple computers, consider variations among the computers that might affect the operation of the script, including varying operating systems, devices, file system organization, environment variables, and Windows PowerShell versions.
- Plan the help topics for the workflow in advance. Keep a record of information that you need to communicate to users, including a preferred workflow architecture, preferred session configuration characteristics, and permissions.

To write the workflow, use a script editor, such as Windows PowerShell Integrated Scripting Environment (ISE), that enforces workflow syntax and highlights syntax errors. The syntactic differences between scripts and workflows are significant, so a tool that knows workflows, as well as scripts, will save you significant coding and testing time.

For more information about tasks that are well-suited to workflows, see "How Workflows and Windows PowerShell Scripts Differ" in Getting Started with Windows PowerShell Workflow.

The Workflow Keyword

Begin with the **workflow** keyword, which identifies a workflow command to Windows PowerShell. The **workflow** keyword is required in a script workflow. The name of the workflow follows the **workflow** keyword. The body of the workflow is enclosed in braces.

The following is the syntax diagram for the **Workflow** keyword.

```
workflow Test-Workflow

{

    ...

}
```

For more information about keywords in the Windows PowerShell scripting language, see about_Language_Keywords.

Naming Workflows and Workflow Elements

A workflow is a Windows PowerShell command type. Select a name with a verb-noun format. Use the Get-Verb cmdlet and the Approved Verbs for Windows PowerShell Commands help topic to select an approved verb for the name. Windows PowerShell generates a warning when commands that do not include approved verbs are imported from a module. Use a descriptive

noun and a noun prefix, if necessary, to prevent command name conflicts when the workflow is imported into a session.

Parameter names and variable names in workflows can include only letters, numbers, and the hyphen (-) and underscore (_) characters. Avoid using the hyphen character in parameter names, because each use of a hyphenated parameter name in the workflow and in calls to the workflow must be enclosed in braces. Do not use the names of the workflow common parameters, workflow runtime variables, or other reserved words, such as "workflow" and "parallel".

For more information about workflow runtime parameters, see the PSWorkflowRuntimeVariable Enumeration help topic in the MSDN Library.

For more information about reserved words, see "Reserved Words in Workflows" in Syntactic Differences Between Script Workflows and Scripts and about_Reserved_Words.

How to Add Parameters to a Workflow

To add parameters to a workflow, use the **Param** keyword with the optional **Parameter** attribute. These are the same techniques that you use to add parameters to a function.

The following code blocks show the syntax of the **Param** keyword and the **Parameter** attribute in a script workflow.

```
workflow Test-Workflow

{

param ([Type]$<ParameterName>)

}

workflow Test-Workflow

{

    Param

    (

        [Parameter(Mandatory=<$True | $False>)]

        [<Type>]

        $<ParameterName>

    )

}
```

You can also use the **CmdletBinding** attribute to specify the **ConfirmImpact**, **DefaultParameterSetName**, **HelpUri**, and **SupportsShouldProcess** attributes.

The following code block shows the syntax of a sample workflow that uses the **CmdletBinding** attribute.

```
workflow Test-Workflow

{
```

```
    [CmdletBinding(ConfirmImpact=<String>,
        DefaultParameterSetName=<String>,
        HelpURI=<URI>,
        PositionalBinding=<Boolean>)]

    Param
    (
        [parameter(Mandatory=$true)]
        [String[]]
        $<ParameterName>
    )
}
```

You do not need to add workflow common parameters to the workflow. Windows PowerShell Workflow adds the common parameters and workflow common parameters to all workflows, including simple workflows that do not have the **CmdletBinding** attribute or the **Parameter** attribute.

For more information about parameters in functions and workflows, see the following help topics.

- about_Functions_CmdletBindingAttribute
- about_Functions_Advanced_Parameters
- about_CommonParameters
- about_WorkflowCommonParameters

How to Add Activities to a Workflow

An *activity* is the basic unit of work in a workflow. Every command and expression that you use in a workflow runs as an activity. In Windows PowerShell Workflow, an activity is very much like a cmdlet. You run an activity by typing the activity name followed by its parameters. Windows PowerShell Workflow converts many of the Windows PowerShell included cmdlets to activities. There are only a small set of excluded cmdlets. For a list of excluded cmdlets, see "Excluded Cmdlets" in Using Activities in Script Workflows.

Most of the Windows PowerShell Core cmdlets have been implemented as activities. To use them in workflow, just type the cmdlet name or the cmdlet alias.

For example, the following workflow includes the Get-Process activity.

```
workflow Test-Workflow
{
    Get-Process -Name PowerShell
    gps -Name Winword
```

```
}
```

When a cmdlet does not have an matching activity and it is not explicitly excluded, Windows PowerShell Workflow automatically runs the cmdlet in the **inlineScript** activity and returns the output to the workflow.

For example, the following workflow uses the Get-WindowsFeature cmdlet, which runs implicitly in an **inlineScript** activity.

```
workflow Test-Workflow

{

    Get-Process -Name PowerShell

    Get-WindowsFeature -Name PowerShell, PowerShell-v2

}
```

You can even run cmdlets that are excluded because they are not implemented as activities, but you must run them in an **inlineScript** activity. For example, the following workflow uses an **inlineScript** activity to run the Get-Variable cmdlet, which is excluded.

```
workflow Test-Workflow

{

    Get-Process -Name PowerShell

    Get-WindowsFeature -Name PowerShell, PowerShell-v2

    InlineScript { Get-Variable -Name PSHome }

}
```

You can also use expressions in a workflow, including arithmetic expressions and comparison statements. Windows PowerShell Workflow automatically runs expressions in a special activity that is designed for the purpose.

For example, the following workflow includes an expression. Like all expressions, it is run in an activity.

```
workflow Test-Workflow

{

    Get-Process -Name PowerShell

    Get-WindowsFeature -Name PowerShell, PowerShell-v2

    InlineScript { Get-Variable PSHome }


    320GB / 3MB

}
```

For more information about activities, see Using Activities in Script Workflows.

How to Use the Parameters of Activities

Activities are commands that have parameters, just like cmdlets. When a cmdlet is converted to an activity, the syntax and parameters remain the same. There are a few exceptions, which are listed in the "Activities in Script Workflows" help topic.

However, positional parameters are not valid in activities, so all parameter names are required. Parameter aliases and abbreviations are permitted. For example, in the following sample workflow, the **Name** parameters of the Get-Process, Get-WindowsFeature, and Get-Variable cmdlets appears in the commands.

```
workflow Test-Workflow

{

    Get-Process -Name PowerShell

    Get-WindowsFeature -Name PowerShell, PowerShell-v2

    InlineScript { Get-Variable -Name PSHome }

}
```

Dynamic parameters are also not valid in activities. To call a dynamic parameter, enclose the command in an **InlineScript** activity.For example, the following workflow includes an **inlineScript** activity that uses the **CodeSigningCert** dynamic parameter of the Get-ChildItem cmdlet. The **CodeSigningCert** parameter is added by the Certificate provider and works only in the Cert: drive.

```
workflow Test-Workflow

{

   InlineScript {Get-ChildItem -Path Cert:\CurrentUser -CodeSigningCert}

}
```

For more information about the **inlineScript** activity, see Running Windows PowerShell Commands in a Workflow.

How to Use Activity Common Parameters

Windows PowerShell Workflow adds a set of activity common parameters to activities. These parameters allow you to set options that are critical in a multi-computer environment.

Some of the workflow common parameters are also activity common parameters. This feature lets you create activity-specific exceptions to the values of the workflow common parameter values. For example, you can use the **PSComputerName** parameter of an activity to run an activity only on selected computers, or the **PSCredential** parameter to run an activity with alternate credentials.

The activity common parameters are valid on most, but not all, activities. For example, the activity common parameters are not valid on the **Suspend-Workflow** and **Checkpoint-Workflow**

activities. Also, the activity common parameters are not available on cmdlets or expressions that Windows PowerShell Workflow automatically runs in an activity. The activity common parameters are available on the **InlineScript** activity, but not on commands in the **InlineScript** script block.

The activities in the following workflow use the activity common parameters where they are valid.

```
workflow Test-Workflow

{

    Get-Process -Name PowerShell -PSComputerName Server01, Server 12

    InlineScript { Get-Variable -Name PSHome } -PSRunningTimeoutSec 3600


    #No activity common parameters.

    Get-WindowsFeature -Name PowerShell, PowerShell-v2


    #No activity common parameters.

    320GB / 3MB

}
```

For more information about the activity common parameters, see Using Activities in Script Workflows.

How to Get Common Parameter Values

Windows PowerShell Workflow adds the workflow runtime variables to all workflows. The workflow runtime variables include the values of the common parameters, workflow common parameters, and other values that are critical to the workflows.

The variable values are extremely useful. Take the time to review them and use them in your workflows. For example, when you run a workflow, the **PSComputerName** parameter takes the names of the target computers, but when you access the $PSComputerName variable in a workflow, it contains the name of the computer on which the workflow is currently running.

For information about the workflow runtime variables, see PSWorkflowRuntimeVariable Enumeration.

To access the parameter values from within the workflow, use the variable name. You do not need to define the variable.

For example, the following workflow uses the value of the **PSConnectionRetryCount** common parameter by referring to it as the **$PSConnectionRetryCount** variable.

```
workflow Test-Workflow

{

    "Retry count is $PSConnectionRetryCount."

}
```

The following example shows that when you use the **PSConnectionRetryCount** parameter when running the workflow, the parameter value is visible in the workflow.

```
PS C:\>Test-Workflow -PSConnectionRetryCount 4

Retry count is 4.
```

To change the values of the runtime and common parameters in a workflow, use the **Set-PSWorkflowData** activity.

For example, the following workflow includes a command that changes the **PSConnectionRetryCount** value.

```
workflow Test-Workflow

{

    "Retry count is $PSConnectionRetryCount."

    Set-PSWorkflowData -PSConnectionRetryCount 5

    "Retry count is $PSConnectionRetryCount."

}


PS C:\>Test-Workflow -PSConnectionRetryCount 4

Retry count is 4.

Retry count is 5.
```

For more information about variables in Windows PowerShell Workflow, see the following help topics.

- Using Variables in Script Workflows
- PSWorkflowRuntimeVariable Enumeration
- about_Get-PSWorkflowData
- about_Set-PSWorkflowData

How to Get Runtime Variables and Common Parameter Values

Windows PowerShell Workflow adds the workflow runtime variables to all workflows. The workflow runtime variables include the values of the common parameters, workflow common parameters, and other values that are critical to the workflows.

The variable values are extremely useful. Take the time to review them and use them in your workflows. For example, when you run a workflow, the **PSComputerName** parameter takes the names of the target computers, but when you access the $PSComputerName variable in a workflow, it contains the name of the computer on which the workflow is currently running.

For information about the workflow runtime variables, see PSWorkflowRuntimeVariable Enumeration.

To access the parameter values from within the workflow, use the variable name. You do not need to define the variable.

For example, the following workflow uses the value of the **PSConnectionRetryCount** common parameter by referring to it as the **$PSConnectionRetryCount** variable.

```
workflow Test-Workflow

{

    "Retry count is $PSConnectionRetryCount."

}
```

The following example shows that when you use the **PSConnectionRetryCount** parameter when running the workflow, the parameter value is visible in the workflow.

```
PS C:\>Test-Workflow -PSConnectionRetryCount 4

Retry count is 4.
```

To change the values of the runtime and common parameters in a workflow, use the **Set-PSWorkflowData** activity.

For example, the following workflow includes a command that changes the **PSConnectionRetryCount** value.

```
workflow Test-Workflow

{

    "Retry count is $PSConnectionRetryCount."

    Set-PSWorkflowData -PSConnectionRetryCount 5

    "Retry count is $PSConnectionRetryCount."

}


PS C:\>Test-Workflow -PSConnectionRetryCount 4

Retry count is 4.

Retry count is 5.
```

For more information about variables in Windows PowerShell Workflow, see the following help topics.

- Using Variables in Script Workflows
- PSWorkflowRuntimeVariable Enumeration

How to Run a Script in a Workflow

To run a script (a .ps1 file) in a workflow, enclose the call to the script in an **InlineScript** activity. Like other commands and activities, the script runs on all workflow target computers and the output is returned to the workflow

For example, the following workflow includes an InlineScript activity that runs a script. The script is stored in a shared directory. The workflow uses the PSPersist activity parameter to take a checkpoint when the script completes.

```
workflow Test-Workflow

{

   $AssetData =InlineScript {\\Server01\Share01\Get-AssetData.ps1 -All} -PSPersist

   ...

}
```

Because workflows are designed to run on multiple computers, use an absolute path to the script file, rather than a relative path.

How to Run Commands in Parallel

The **Parallel** and **ForEach -Parallel** keywords optimize your workflow by running commands concurrently and in indeterminate order. The **Sequence** keyword runs selected commands sequentially within a **Parallel** script block. These keywords are valid only in a workflow. They are valid in nested workflows, but not valid in nested functions in a workflow. By default, commands run in sequence in which each command completes before the next command begins.

Because parallel execution can provide significant performance improvements, especially when the workflow targets numerous computers, use it whenever possible.

The following types of commands and activities are good candidates for parallel execution.

- Commands that do not share data, such as Get-Process and a Get-Service commands.
- Commands that run on a collection of like items. Most **ForEach** statements are good candidates for **ForEach -Parallel** execution.

Parallel

The commands in a **Parallel** script block can run concurrently. The order in which they run is not determined.

The following diagram shows the syntax of a workflow with the **Parallel** keyword and script block.

```
workflow Test-Workflow

{

    Parallel

    {

         <Activity>

         <Activity>

...

    }
```

```
}
```

For example, the following workflow includes a Parallel script block that runs activities that get processes and services on the computer. Because the calls to Get-Process and Get-Service are independent, they can run concurrently and in any order.

```
workflow Test-Workflow

{

    Parallel

    {

         Get-Process

         Get-Service

    }

}
```

ForEach -Parallel

The **Parallel** parameter of the **ForEach** keyword runs the commands in a **ForEach** script block once for each item in the specified collection. The items in the collection are processed in parallel. The commands in the script block run sequentially.

Like the **ForEach** statement in Windows PowerShell, the variable that contains collection ($<Collection>) must be defined before the **ForEach -Parallel** statement, but the variable that represents the current item ($<item>) is defined in the **ForEach -Parallel** statement.

The following diagram shows the command syntax.

```
workflow Test-Workflow

{

    ForEach -Parallel ($<item> in $<collection>)

    {

       <Activity1>

       <Activity2>

 ...

    }

}
```

For example, the following workflow contains a **Foreach -Parallel** statement that processes the disks that the **Get-Disk** activity gets. The commands in the script block run sequentially, but they are run on the disks in parallel.

```
workflow Test-Workflow

{

    $Disks = Get-Disk
```

```
    ForEach -Parallel ($Disk in $Disks)
    {
        $DiskPath = $Disk.Path
        $Disk | Initialize-Disk
        Set-Disk -Path $DiskPath
    }
}
```

Sequence

The **Sequence** keyword runs commands in sequence within a **Parallel** script block. The **Sequence** script block runs in parallel with other commands, but the commands within the sequence script block run sequentially and in the specified order.

The following code block show the syntax of a **Sequence** script block. Activity 3 might run before or at the same time as Activity 1 and Activity 2, but Activity 4 does not run until Activity 3 is complete.

```
workflow Test-Workflow
{
    parallel
    {
       <Activity1>
       <Activity2>

       sequence
       {
           <Activity3>
           <Activity4>
           ...
       }
        ...
    }
}
```

Adding Nested Functions and Nested Workflows

In Windows PowerShell® Workflow, you can include functions and nested workflows to reuse and organize the commands in a script workflow. This topic explains the rules for creating and calling nested functions and workflows.

Adding Nested Functions and Nested Workflows

You can include functions and nested workflows to organize the commands in a script workflow. There are no syntactic limits to nesting depth. For example, the following workflow includes nested workflows and nested functions.

```
workflow One
{
    "One"
    Two
    workflow Two
    {
        "Two"
        Three
        Function Three
        {
            "Three"
            function Five {"Five"}
            Workflow Four {"Four"; Five}
            Four
        }
    }
}
```

Using Parameters and Variables in Nested Functions and Workflows

In general, child scopes inherit from the parent scope. Nested functions and workflows in a workflow can get the values of workflow parameters and the values of variables created in the top-level workflow scope. However, the values of variables created and changed in the child scope are not available in the parent scope.

Note

TROUBLESHOOTING NOTE: The values of variables created in the workflows scope might not be available to nested functions. The values of workflow parameters and of variables created in the workflow scope might not be available to nested workflows.

To assure that nested workflows and functions have access to the values of workflow parameters and the values of variables created in workflow scope, define the parameter on the nested function or workflow. Then, when calling the nested function or workflow, pass the parameter value.

For example, the following sample workflow includes the Get-ModuleAlias nested workflow. To assure that the nested workflow has access to the value of the ModuleName parameter of the workflow, define a similar parameter on the nested workflow and pass the ModuleName parameter value when calling the nested workflow.

```
Workflow Get-ModuleProperties

{

    Param

    (

         [parameter(Mandatory=$true)]

         [String]

         $ModuleName

    )


    workflow Get-ModuleAlias

    {

      Param

      (

         [parameter(Mandatory=$true)]

         [String]

         $ModuleNameForAlias,


       )

       ((Get-Module -Name $ModuleNameForAlias -ListAvailable).ExportedAliases).Keys

    }

    Get-ModuleAlias -ModuleName $ModuleNameForAlias

}
```

Calling Nested Workflows and Functions

This section explains the rules for calling nested functions and workflows in a workflow that differ from the rules for nested elements in standard functions and scripts.

- The order in which the nested workflows and functions are defined in the workflow is not important. You can call a function or nested workflow before it is defined in the workflow, as shown in the following example.

```
Workflow Test-Workflow
{
    param ($ModuleName)

    function Initialize-Files
    { ... }

    workflow Get-Metadta
    {
        ...
        Initialize-Files
        Get-Configuration
        Get-DataSource
        Get-Module -ModuleName $ModuleName
        function Get-DataSource
        {
           ...
        }
    }

    function Get-Configuration
    {
       ...
    }
}
```

- Nested workflows can call workflows and functions in the current scope and any parent scope, as shown in the following example.

```
# In Test-Nested.ps1
function Zero {"Zero"}
```

```
workflow One
{
    "One"
    Two
    Three
    Four
    workflow Four
    {
        "Four"; Two; Three; Five;
        workflow Five {"Five"; Two; Three; Zero}
    }
}
function Two {"Two"}
workflow Three {"Three"}
One
```

The output of the Test-Nested.ps1 script is as follows.

```
PS C:\Test-Nested.ps1
One
Two
Three
Four
Two
Three
Five
Two
Three
Zero
```

- Recursive calling, as shown in the following example, is not permitted in workflows.

```
workflow one
{
    "One"
    Two
}
```

```
workflow two
{
    "Two"
    One
}
```

Recursive calls are permitted in scripts and functions.

- Workflow and function names can be used only once in each workflow. Name conflicts generate errors, as shown in the following example.

```
workflow Test-Workflow
{
    workflow Test1 {"Test1"}
    function Test1 {"Hello, Test1"}
}


At line:4 char:5
+     function Test1 {"Hello, Test1"}
+     ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
The function or workflow 'Test1' cannot be redefined.
    + CategoryInfo          : ParserError: (:) [],
ParseException
    + FullyQualifiedErrorId : FunctionRedefinitionNotAllowed
```

- If the name of a nested workflow or function is the same as the name of a workflow or function in a parent scope, calls to that name will call the workflow or function in the current (workflow) scope.

  In the following example, the call to "Test-Nested" calls the Test-Nested workflow. If the Test-Nested workflow had a different name, the call to "Test-Nested" would call the Test-Nested function.

```
function Test-Nested {"Test-Nested Function"}
workflow Test-Workflow
{
    Workflow Test-Nested {"Test-Nested Workflow"}
    Test-Nested
}
PS C:\>Test-Workflow
```

```
Test-Nested Workflow
```

In a standard Windows PowerShell script or function, names of nested functions can be reused. Calls invoke the last function defined with the specific name before the call.

Running Windows PowerShell Commands in a Workflow

To run commands or expressions in a workflow that are valid in Windows PowerShell, but not valid in workflows, run the commands in an **inlineScript** activity. You can use also an **inlineScript** activity to run Windows PowerShell scripts (.ps1 files) in a workflow.

The InlineScript activity

The **inlineScript** activity runs commands in a standard, non-workflow Windows PowerShell session and then returns the output to the workflow. It is valid only in workflows.

The commands in an **inlineScript** script block run in a single session and can share data, such as the values of variables. By default, the InlineScript session runs "out-of-process," that is, it runs in its own process, not in the workflow process. However, you can change this default by changing the value of the **OutOfProcessActivity** property of the session configuration.

The following diagram shows the syntax of the **inlineScript** activity.

```
inlineScript {<script block>} <WorkflowCommonParameters>
```

Because it is an activity, but not a cmdlet, the **inlineScript** activity has the activity common parameters, but it does not have the Windows PowerShell common parameters.

When creating **inlineScript** commands, remember that the commands in an InlineScript script block do not have the features of a workflow, such as automatic multi-computer management. To run the commands in **inlinescript** activity on remote computers, you need to establish and manage connections to the remote computer from within the **inlinescript** block and then run the commands remotely.

Rules for Using the InlineScript Activity

An **inlineScript** script block can include all valid Windows PowerShell commands and expressions, including commands, expressions, and syntax that are not otherwise valid in a workflow. The commands and expressions in the script block share all state and data, including the values of variables.

You can place an **inlineScript** activity anywhere in a workflow or nested workflow, including inside a loop or control statement or a **Parallel** or **Sequence** script block.

The **inlineScript** activity has the activity common parameters, including PSPersist. However, the commands and expressions in an **inlineScript** script block do not have the features of workflows, such as checkpointing ("persistence") or the activity common parameters.

Variables in InlineScript

By default, the variables that are defined in a workflow are not visible to the commands in the InlineScript script block. To make workflow variables visible to the InlineScript, use the **$Using** scope modifier.

The following example shows how to make the values of top-level workflow variables available to the commands in an **inlineScript** script block.

```
workflow Test-Workflow

{

    $a = 3


    # Without $Using, the $a workflow variable is not visible

    # in inline script.

    InlineScript {"Inline A0 = $a"}


    # $Using imports the variable and its current value.

    InlineScript {"Inline A1 = $Using:a"}

}


PS C:\> Test-Workflow

Inline A0 =

Inline A1 = 3
```

InlineScript commands can change the value of the variable that was imported from workflow scope, but the changes are not visible in workflow scope. To make them visible, return the changed value to the workflow scope, as shown in the following example.

```
workflow Test-Workflow

{

    $a = 3


    # Changes to an InlineScript variable do not affect

    # the workflow variable.

    InlineScript {$a = $using:a+1; "Inline A = $a"}
```

```
    "Workflow A = $a"


    # To change the value in workflow scope, return the new value.
     $a = InlineScript {$a = $Using:a+1; $a}
     "Workflow New A = $a"
}


PS C:\> Test-Workflow
Inline A = 4
Workflow A = 3
Workflow New A = 4
```

Security

Troubleshooting Note: A statement with the $Using scope modifier should appear before any use of the variable in the InlineScript script block. In Windows PowerShell 3.0, the $Using scope modifier can appear anywhere in the InlineScript script block, but this behavior might change in future versions of Windows PowerShell.

Running InlineScript in the Workflow Process

To improve reliability, the commands in an **InlineScript** script block run in their own process, outside of the process in which the workflow runs, and then return their output to the workflow process. To direct Windows PowerShell to run the **InlineScript** activity in the workflow process, remove the **InlineScript** value from the **OutOfProcessActivity** property of the session configuration, such as by using the **OutOfProcessActivity** parameter of the New-PSWorkflowExecutionOption cmdlet.

The following command sequence shows how to remove the InlineScript value from the **OutOfProcessActivity** property and then to restore the value to the Microsoft.PowerShell.Workflow session configuration by deleting the session configuration, and then running the Enable-PSRemoting cmdlet to recreate it.

```
#To run InlineScript in the workflow process
[ADMIN]: PS C:\> (Get-PSSessionConfiguration
Microsoft.PowerShell.Workflow).OutOfProcessActivity
InlineScript
[ADMIN]: PS C:\> $o = New-PSWorkflowExecutionOption -OutOfProcessActivity ""
[ADMIN]: PS C:\> Set-PSSessionConfiguration Microsoft.PowerShell.Workflow -
SessionTypeOption $o -Force
```

```
[ADMIN]: PS C:\> (Get-PSSessionConfiguration
Microsoft.PowerShell.Workflow).OutOfProcessActivity
[ADMIN]: PS C:\>

#To restore the default setting and run InlineScript out of process
[ADMIN]: PS C:\> (Get-PSSessionConfiguration Microsoft.PowerShell.Workflow) | Unregister-
PSSessionConfiguration -Force
[ADMIN]: PS C:\> Enable-PSRemoting -Force
WinRM is already set up to receive requests on this computer.
WinRM is already set up for remote management on this computer.
[ADMIN]: PS C:\> (Get-PSSessionConfiguration
Microsoft.PowerShell.Workflow).OutOfProcessActivity
InlineScript
```

Examples

The following workflow uses an InlineScript activity to run a Windows PowerShell script (.ps1 file) in a workflow.

```
Workflow Test-Workflow
{
    InlineScript
    {
        \\Server01\Share01\Get-AssetData.ps1 -All
    }
}
```

The InlineScript in the following workflow includes commands that are valid in Windows PowerShell, but not in workflows, including the use of the New-Object cmdlet with the ComObject parameter. The results are returned to the workflow so they can be used in workflow activities.

```
workflow Test-Workflow
{
    $ie = InlineScript
    {
New-Object -ComObject InternetExplorer.Application
-property @{navigate2="www.microsoft.com"}
    }
```

```
    $ie.LocationName
}

PS C:>Test-Workflow
Microsoft Corporation: Software, Smartphones, Online, Games, Cloud Computing, IT Business
Technology, Downloads
```

Making a Workflow Suspend Itself

One of the many benefits of running Windows PowerShell® workflows is that you can suspend and resume them as needed. This is crucial to some scenarios, in which parts of a workflow might be managed by different administrators or require manual intervention.

To suspend a workflow that you are running, use the Suspend-Job cmdlet. But you can also use **Suspend-Workflow** activity to suspend a workflow from within the workflow.

The Suspend-Workflow Activity

The **Suspend-Workflow** activity stops workflow processing temporarily. Before suspending, Windows PowerShell Workflow takes a checkpoint so the state and data are preserved when the workflow is resumed.

To resume the workflow, the user running the workflow uses the Resume-Job cmdlet. You cannot resume a workflow from within the workflow.

After the workflow is resumed, use the Windows PowerShell Job cmdlets to manage the workflow job and get its output, even if you were running the workflow interactively before it was suspended.

Use the Get-Job cmdlet to get the workflow job. To get the output of the workflow job, use the Receive-Job cmdlet.

Suspending a Workflow

The **Suspend-Workflow** activity returns a job object that represents the suspended job, even if you did not explicitly run the workflow as a job, such as by using the **AsJob** parameter.

The job object that **Suspend-Workflow** returns includes the job ID and job state. When the **Suspend-Workflow** activity is successful, the resulting job state is **Suspended**, as shown in the following example.

```
Workflow Test-Suspend
{
    $a = Get-Date
    Suspend-Workflow
```

```
    (Get-Date)- $a
}

PS C:\>Test-Suspend

Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
8      Job8            PSWorkflowJob   Suspended     True            localhost
Test-Suspend
```

Resuming a Workflow Job

To resume the workflow job, use the Resume-Job cmdlet. The **Resume-Job** cmdlet returns the workflow job object immediately, even though it might not yet be resumed. To get the current state of the workflow job, use the Get-Job cmdlet, as shown in the following example.

```
PS C:\>Resume-Job -Name Job8

Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
8      Job8            PSWorkflowJob   Suspended     True            localhost
Test-Suspend

PS C:\>Get-Job -Name Job8

Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
8      Job8            PSWorkflowJob   Completed     True            localhost
Test-Suspend
```

Getting the Output of a Workflow Job

To get the output of a workflow job, use the Receive-Job cmdlet. The output shows that the value of the $a variable, which was defined before the workflow was suspended, is available when the workflow is resumed.

```
PS C:\>Get-Job -Name Job8 | Receive-Job


Days              : 0
Hours             : 0
Minutes           : 0
Seconds           : 19
Milliseconds      : 823
Ticks             : 198230041
TotalDays         : 0.000229432917824074
TotalHours        : 0.00550639002777778
TotalMinutes      : 0.330383401666667
TotalSeconds      : 19.8230041
TotalMilliseconds : 19823.0041
PSComputerName    : localhost
```

Restarting the Computer in a Workflow

Windows PowerShell® Workflow is designed to be robust and recoverable. The features that are built into Windows PowerShell allow you to restart the target computers and, if the workflow is not yet complete, to reconnect to the target computers automatically and continue with the workflow commands. You can even restart the computer that is running workflow, which allows you to run the workflow on the local computer or include the local computer in the list of target computers.

Restarting closes the connection to the target computers, but when you use the **Wait** parameter of the Restart-Computer cmdlet to wait for the restart to complete, Windows PowerShell Workflow automatically reconnects to the target computer and handles all disconnection and reconnection details for you.

If the connection to a target computer cannot be re-established, the workflow throws a non-terminating error and continues running on the remaining target computers.

Restart-Computer

To restart the target computers in a workflow, use the Restart-Computer cmdlet, which is implemented as an activity.

To continue running workflow commands after restarting the target computers, use the **Wait** parameter of the **Restart-Computer** activity. The **Wait** parameter waits for the target computer to restart and then reconnects to the target computer before running the next workflow command. You can use a command like this one anywhere in a workflow without disrupting workflow processing. In workflows, the **Wait** parameter is valid on the local computer, as well as on remote computers.

For example, the following workflow renames the computer and then restarts the computer to make the name change effective before configuring the disks. The command includes the **Wait** parameter of **Restart-Computer**, which waits for Windows PowerShell to restart on the target computer before running the next (Get-Disk) command.

Note that the workflow does not include any Disconnect-PSSession or Connect-PSSession commands, and it does not include any explicit checkpoints.

```
workflow New-ComputerSetup

{

   ...

   $cn =  Get-CimInstance -ClassName Win32_ComputerSystem

   if ($cn.SystemType -like "64*")

   {

       $NewCnName = $PSComputerName + "64"

       Rename-Computer -ComputerName $PSComputerName -NewName $NewCnName

   }

   Restart-Computer -Wait

   $disks = Get-Disk

   ...

}
```

Restarting the Workflow Computer

The commands in a workflow can restart the computer that is running the workflow, such as when the workflow is running on the local computer or the workflow computer is one of the target computers. If you use the **Wait** parameter of the Restart-Computer activity to continue the workflow after the restart, Windows PowerShell Workflow automatically takes a checkpoint and suspends the workflow before restarting the workflow computer.

This section describes two procedures that resume a workflow after the workflow computer is restarted. The first procedure is manual. The second uses a scheduled job.

Manually Resuming the Workflow

To resume a workflow manually after restarting the workflow computer, use the following procedure.

- Use the Import-Module command to import the **PSWorkflow** module.

  Without this module, the Job cmdlets do not recognize workflow jobs.
- Use the Get-Job cmdlet to get the suspended workflow job and the Resume-Job cmdlet to resume the workflow job. To wait for the workflow job to resume, use the **Wait** parameter of **Resume-Job**.
- You can use the job cmdlets to manage a suspended workflow job even if you didn't start the workflow with the **AsJob** parameter.

For more information about cmdlets that behave differently as activities, see the "Activities that Differ from Cmdlets" section of Using Activities in Script Workflows.

For example, the following commands run a workflow that restarts the local computer. When the workflow is suspended to perform the restart, the commands in the example resume it.

```
PS C:\>New-ComputerSetup


Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
8      Job8            PSWorkflowJob   Suspended     True            localhost
New-ComputerSetup


# Resume the suspended workflow job
PS C:\>Import-Module PSWorkflow
PS C:\>Get-Job -Name Job8 | Resume-Job -Wait


Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
8      Job8            PSWorkflowJob   Running       True            localhost
New-ComputerSetup
```

Automate Resuming the Workflow

You can create a scheduled job that resumes a workflow after the workflow computer restarts. Be sure to create the scheduled job on the computer that is running the workflow. When the workflow completes, disable or unregister the scheduled job so that it does not run every time the workflow computer starts.

Use the following procedure to resume a suspended workflow when the workflow computer starts.

1. Use the New-JobTrigger cmdlet to create an **AtStartup** trigger for the scheduled job.

   An **AtStartup** trigger runs a scheduled job every time the computer starts.
2. Use the Register-ScheduledJob cmdlet to create a scheduled job on the computer that is running the workflow. The commands in the scheduled job resume the suspended workflow job.
3. When running the workflow, use the **JobName** workflow common parameter. The scheduled job uses the job name to identify the workflow job.

For example, the following commands create the "ResumeWorkflow" scheduled job, which resumes the "ComputerSetup" workflow job as soon as the computer restarts. When you run the New-ComputerSetup workflow on the local computer in a job named "ComputerSetup," the workflow is suspended just before the computer is restarted. When the computer restarts, the ResumeWorkflow scheduled job resumes the workflow job.

The Get-Job command verifies that the workflow completed and the final Unregister-ScheduledJob deletes the scheduled job.

```
PS C:\>$AtStartup = New-JobTrigger -AtStartup


PS C:\>Register-ScheduledJob -Name ResumeWorkflow -Trigger $AtStartup -ScriptBlock
{Import-Module PSWorkflow; Get-Job ComputerSetup -State Suspended | Resume-Job}


Id         Name               JobTriggers     Command
Enabled
--         ----               -----------     -------                                -
------
1          ResumeWorkflowJob  1               Import-Module PSWorkflow; Get-Job CnS...
True


# Run the workflow. It is suspended when the computer restarts.
PS C:\>New-ComputerSetup -JobName ComputerSetup


Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
```

```
--     ----            -------------   -----         -----------     --------
-------
9      ComputerSetup   PSWorkflowJob   Suspended     True            localhost
New-ComputerSetup


# Verify that the workflow was resumed and completed successfully
PS C:\>Import-Module PSWorkflow
PS C:\>Get-Job ComputerSetup

Id     Name            PSJobTypeName   State         HasMoreData     Location
Command
--     ----            -------------   -----         -----------     --------
-------
9      ComputerSetup   PSWorkflowJob   Completed     True            localhost
New-ComputerSetup


# Delete the scheduled job
PS C:\>Unregister-ScheduledJob -Name ComputerSetup
```

See Also

Using Activities in Script Workflows

Adding Custom Activities to a Script Workflow

You can use built-in activities and custom activities in a Windows PowerShell® workflow. The library of built-in activities, including cmdlets that are implemented as activities, becomes available when you use the **workflow** keyword. In addition, developers can write custom activities in C# and include them in XAML workflows and script workflows.

To add a custom activity to a script workflow, use the **Assembly** parameter of the **#Requires** statement. The **Assembly** parameter identifies the assembly in which the custom activities are defined. When Windows PowerShell Workflow encounters the **Assembly** parameter, it loads the specified assembly and then processes the custom activities defined in the assembly so that they are available for use in the workflow. The **Assembly** parameter is valid only in workflows.

Valid value formats for the **Assembly** parameter of the **#Requires** statement are as follows:

- Fully qualified assembly description

  ```
  #requires -Assembly MyCommand.Activities.dll Version=1.1.0.0,
  Culture=neutral, PublicKeyToken=11cd3656ab363bc4"
  ```

- Fully qualified path to an assembly

  ```
  #requires -Assembly C:\Users\Admin01\MyCommand.Activities.dll
  ```

- Relative path to an assembly in the local directory or in the global assembly cache

  ```
  #requires -Assembly MyCommand.Activities.dll
  ```

To use custom activities from multiple assemblies, enter multiple #Requires statements. The following sample workflow gets custom activities from two assemblies and uses them in the workflow.

```
workflow Test-Workflow

{

    #requires -Assembly C:\Users\Admin01\MyCommand.Activities.dll.dll

    #requires -Assembly MyFunction.Activities.dll.dll

    Get-CustomData -All

    Set-CustomDataLegacy -PSComputerName ...

}
```

For information about writing a custom activity, see "Creating a Workflow Activity from a Windows PowerShell Cmdlet" in the MSDN Library.

Adding Checkpoints to a Script Workflow

This topic explains how to add commands to a Windows PowerShell workflow that "take a checkpoint," that is, they save workflow state and data to disk while the workflow runs. If the workflow is interrupted, it does not have be restarted; only resumed from the last checkpoint. This saving process is also known as *persistence* or *checkpointing*.

Windows PowerShell Workflow is designed for the real world, where network interruptions, reboots, delays, and power outages are part of the job. Any long-running task that affects multiple, heterogeneous machines in a distributed architecture must anticipate, respond, recover, continue, and complete after intentional and unintentional interruptions. Checkpointing is a critical part of this strategy.

For more information about support for robust and recoverable workflows, see Getting Started with Windows PowerShell Workflow.

What is a checkpoint?

A *checkpoint* is a snapshot of the current state of the workflow, including the current values of variables, and any output generated up to that point, and it saves it to disk. You can add multiple

checkpoints to a workflow by using different checkpoint techniques. Windows PowerShell automatically uses the data in newest checkpoint for the workflow to recover and resume the workflow if the workflow is interrupted, intentionally or unintentionally.

The checkpoint data (or *persisted data*) for a workflow is saved in the user profile on the hard drive of the computer that is hosting the workflow session.

When you run the workflow as a job, such as by using the **AsJob** workflow common parameter, the workflow checkpoints are retained until you delete the job, such as by using the Remove-Job cmdlet. Otherwise, the workflow checkpoints are deleted when the workflow completes.

Where to place checkpoints

You can place checkpoints anywhere in a workflow, including before and after each command or expression, or after each activity in a workflow. Windows PowerShell Workflow allows you to do so.

However, you probably will not add checkpoints wherever you are permitted, because the advantages of checkpoints are balanced by a performance cost. Every time you take a checkpoint, data is collected and written to disk on the workflow server computer. In a long workflow that runs on multiple computers, this delay is perceptible and affects the time it takes the workflow to complete.

When planning checkpoints, consider the process of resuming a workflow. Be sure that the time it takes to rerun a section of the workflow if it is interrupted is greater than the time it takes to write the checkpoint state and data to disk. Remember that the checkpoint is taken every time the workflow runs on each target computer; as though the entire workflow was enclosed in a **ForEach** statement.

Add a checkpoint after each significant part of the workflow completes; a part that you do not want to repeat if the workflow is interrupted. This could be a part of the workflow that takes a long time to complete, such as gathering input for a complex calculation or a complex element of output. It could be a section that requires resources that are not always available, or requires the coordination of multiple computers, connections, and resources.

When the workflow definitions is complete, use the testing process to refine your checkpoints. Add and remove (comment out) checkpoints to achieve both a reasonable running time and a reasonable recovery time.

How to add a checkpoint

The following techniques checkpoint (save state and data to disk) a workflow so that it can recover and resume without being restarted. All of the techniques described below save state and data to the same place on disk and enable the same resuming options.

PSPersist workflow common parameter

This **PSPersist** workflow common parameter adds checkpoints to the workflow. It is used when running the workflow command. **PSPersist** cannot cancel, ignore, or suppress explicit checkpoints in the workflow. It is valid on all workflows.

By default, Windows PowerShell Workflow adds a checkpoint at the beginning and end of the workflow, in addition to any checkpoints that are specified in the workflow.

$True

Adds a checkpoint at the beginning and end of the workflow and after each activity, in addition to checkpoints that are specified in the workflow.

$False

No checkpoints at the beginning or end of the workflow. Takes a checkpoint only when a checkpoint is specified in the workflow.

The following examples run the Test-Workflow workflow with the PSPersist workflow common parameter.

```
PS C:\>Test-Workflow -PSPersist $True

PS C:\>Test-Workflow -PSPersist $False
```

PSPersist activity common parameter

The **PSPersist** activity common parameter with a value of $True takes a checkpoint after the activity completes. It is valid on all activities in a workflow. It is not valid on expressions or on the commands in an InlineScript script block.

This parameter takes a Boolean value. The following are valid values.

$True

Takes a checkpoint after the activity completes.

$False

Does not add any checkpoints. This value has no effect.

The following workflow uses the **PSPersist** activity common parameter to take a checkpoint after an activity completes.

```
Workflow Test-Workflow

{

    Get-Process -Name wsmprovhost -PSPersist $True

}
```

Checkpoint-Workflow Activity

The **Checkpoint-Workflow** activity takes a checkpoint immediately. You can use a **Checkpoint-Workflow** activity multiple times in a workflow and place it in after any command or expression in a workflow, except that it cannot appear in an **InlineScript** script block. The Checkpoint-Workflow activity does not take any parameters, including common parameters and workflow common parameters.

```
Workflow Test-Workflow

{

    $p = Get-Process

    $s = Get-Service

    Checkpoint-Workflow

}
```

$PSPersistPreference Preference Variable

The **$PSPersistPreference** preference variable with a value of $True takes a checkpoint after every activity that follows the variable definition. It is equivalent to setting the value of the **PSPersist** activity parameter to $true for all subsequent activities. To stop adding checkpoints, set the value of the **$PSPersistPreference** preference variable to $False.

Unlike other preference variables, this variable is effective only in a workflow; not in a Windows PowerShell session. You can define it in a session or in a Windows PowerShell profile, but it has no effect on workflows run in the session.

This parameter takes a Boolean value. The following are valid values.

$True

Takes a checkpoint after each activity, in addition to checkpoints that are specified in the workflow. This value affects all subsequent activities in the workflow.

$False

Stops adding additional checkpoints. Takes a checkpoint only when the checkpoint is specified in the workflow.

The following workflow uses the $PSPersistPreference variable to add checkpoints to a workflow. It adds checkpoints to all activities between `$PSPersistPreference=$True` and `$PSPersistPreference=$False`.

```
Workflow Test-Workflow

{
```

```
    $WinRm = Get-Service Winrm


    $PSPersistPreference = $true
    $Assets = InlineScript {\\Server\Share\Get-AssetData.ps1}
    ...
    $PSPersistPreference = $false


    Parallel
    {
        ...
}
```

Checkpoints in Pipeline and Parallel Script Blocks

When checkpoints are set on activities in a pipeline, the checkpoint is not taken until the pipeline completes.

When checkpoints are set on activities in a **Parallel** script block, the checkpoints is not taken until the **Parallel** script block has run on all target computers. However, when checkpoints are set on activities in a **Sequence** script block, checkpoints are taken after each activity completes on each target computer.

Writing Help for a Script Workflow

This topic explains how to provide help topics for the script workflows that you author. For general information about writing help for commands in a module, see How to Write Help for Windows PowerShell Modules in the MSDN Library.

Supporting the Get-Help Cmdlet

The Get-Help cmdlet gets help for commands in modules, including the script workflows in modules. To write help topics for the script workflows that you author, use the following procedure.

1. Include the script workflow in a script module (.psm1).
2. Use the XML-based cmdlet help topic format to write a help topic for the script workflow.
3. Add an **.ExternalHelp** comment that specifies the name of the XML-based cmdlet help topic for the workflow.

Comment-based help is not supported in workflows, except for the **.ExternalHelp** comment that identifies the help file for the workflow.

There are no requirements for the file names of help topics for script modules. The recommended name format is as follows:

```
<Script-module-base-name>.psm1-help.xml
```

For example, the recommended help file for the DeploymentWorkflow.psm1 script module is as follows.

```
DeploymentWorkflow.psm1-help.xml
```

To enable the **Get-Help** cmdlet to locate the help topic for the script workflow, use an **.ExternalHelp** comment. The follow example shows an .ExternalHelp comment for in a script workflow in the DeploymentWorkflow.psm1 script module.

```
#.ExternalHelp DeploymentWorkflow.psm1-help.xml

Workflow New-Package

{

     ...

}

Workflow Set-Package

{

     ...

}
```

For more information about using the .ExternalHelp comment, see about_Comment_Based_Help.

Supporting Online Help

The **Online** parameter of the Get-Help cmdlet opens the online version of the help topic for a command in the user's default Internet browser. The **Online** parameter gets the Uniform Resource Identifier (URI) for the online help topic from the **HelpUri** property of the command.

To supply a value for the **HelpUri** property of a script workflow, add the **HelpUri** attribute of the **CmdletBinding** attribute, as shown in the following example. This is the same technique that is used to add the HelpUri attribute to an advanced function.

```
 #.ExternalHelp DeploymentWorkflow.psm1-help.xml

Workflow New-Package

{

    [CmdletBinding(HelpURI="http://www.fabrikam.com/?TopicID=NewPackage")]


    Param ([Parameter]$ServerName)

     ...

}
```

To verify that the HelpUri property value is effective, use the Get-Command cmdlet to get the HelpUri property of the script workflow.

```
Workflow New-Package

{

    [CmdletBinding(HelpURI="http://www.fabrikam.com/?LinkID=242132")]


    Param ([Parameter]$ServerName)

}


PS C:\> (Get-Command Deploy-Update).HelpUri

http://www.fabrikam.com/?TopicID=NewPackage
```

For more information about supporting online help for the commands in a module, see Supporting Online Help in the MSDN Library.

Supporting Updatable Help

The Update-Help and Save-Help cmdlets download the newest help files for a module from the Internet or a file share and install them on the local computer. You can provide updated help topics for the workflows in a script module by using the same procedure you use for types of command in the module.

For more information about supporting Updatable Help for the commands in a module, see Supporting Online Help in the MSDN Library.

See Also

about_Comment_Based_Help

Get-Help

Supporting Online Help

Supporting Online Help

Update-Help

Save-Help

Saving Your Workflow in a Module

After authoring a Windows PowerShell script workflow as described in Writing a Script Workflow, or authoring a XAML workflow, you can save the workflow as a module. Storing workflows in a module makes it easier for you to package, distribute, describe, and install them easily.

Prepare your workflow by saving it in a module

The example workflow in the following procedure, `Install-Role.ps1`, is a Windows PowerShell script workflow.

To save a workflow in a module

1. Create a module directory. The following example command creates a module directory called **ManagedRoles** in the user's **Modules** directory.

   mkdir C:\Users\CoolDev\WindowsPowerShell\Modules\ManagedRoles

   Important
   Do not name a module the same name as any other modules that are located in your module path.

2. Copy at least one Windows PowerShell script-based or XAML workflow file into the module directory, as shown in the following example command.

   copy-item .\Install-Role.ps1 -dest C:\Users\CoolDev\WindowsPowerShell\Modules\ManagedRoles

See Also

Getting Started with Windows PowerShell Workflow

Writing a Script Workflow

Configuring Your Workflow Environment

Running a Windows PowerShell Workflow

about_Modules

Workflow Authoring Reference Topics

Use the topics in this section for reference while learning to write Windows PowerShell script workflows.

The workflow authoring reference topics are as follows.

- Syntactic Differences Between Script Workflows and Scripts
- Using Activities in Script Workflows

- Using Variables in Script Workflows

Syntactic Differences Between Script Workflows and Scripts

This topic lists and explains the primary syntactic differences between Windows PowerShell scripts and script workflows, including commands and syntax that are valid in a workflow.

Review this topic in advance and use it as a reference. However, the best way to learn the syntax differences in workflows is to write workflows in a script editor that enforces valid workflow syntax, such as Windows PowerShell® Integrated Scripting Environment (ISE). If you type a syntax error, Windows PowerShell ISE recognizes and explains the error.

Scripts and Script Workflows

A *script workflow* is a workflow written in the Windows PowerShell scripting language. You can save and share your workflows in Windows PowerShell scripts (.ps1 files) and script modules (.psm1 files).

For more information about scripts and modules, see about_Scripts and Writing a Windows PowerShell Module in the MSDN Library.

The Workflow Keyword

Workflows begin with the **workflow** keyword, followed by a workflow name in verb-noun format. The commands and expressions in the workflow are enclosed in the workflow script block.

```
workflow <Verb-Noun>
{
    ...
}
```

The Get-Command cmdlet recognizes all workflows in a session. The command type is Workflow.

```
workflow Test-Workflow
{
    "Hello, World"
}


PS C:\>Get-Command Test-Workflow
```

```
CommandType     Name                ModuleName

-----------     ----                ----------

Workflow        Test-Workflow
```

Workflow Names

Workflows have names with the verb-noun format. Use approved verbs for the verb in the workflow name.

For a list of approved verbs, use the Get-Verb cmdlet. For more information, see Approved Verbs for Windows PowerShell Commands in the MSDN library.

The noun names in a workflow can include only letters, numbers, and the hyphen (-) and underscore (_) characters. Avoid names with hyphen and underscore characters, because they might be confusing to users.

CmdletBinding Attribute

The **CmdletBinding** attribute is optional and valid in a workflow. Use it to add attributes to the workflow command. In workflows, the **CmdletBinding** attribute supports the following parameters.

- ConfirmImpact
- DefaultParameterSetName
- HelpUri
- PositionalBinding

For more information about the **CmdletBinding** attribute and its parameters, see about_Functions_CmdletBindingAttribute.

Workflow Parameters

The parameters in a workflow are almost exactly like the parameters in a function or script. However, parameter names in a workflow can include only letters, numbers, and the hyphen (-) and underscore (_) characters. Avoid names with hyphen characters, because they must be enclosed in braces to be interpreted correctly.

You can have one or many parameter sets. You can use a simple **Param** statement or the optional **Parameter** attribute. Parameter attributes and parameter validation are supported in workflows, but not in nested workflows.

Workflow Common Parameters

Windows PowerShell Workflow adds the Windows PowerShell common parameters and workflow common parameters to all workflows, including simple workflows that do not use the **CmdletBinding** attribute or the **Parameter** attribute.

Windows PowerShell Workflow also adds the some of the workflow common parameters to activities in a workflow. This feature lets you create exceptions to the workflow common parameters for specific activities.

The values of common parameters and workflow common parameters are available within a workflow as variables. For more information, see Using Variables in Script Workflows.

For more information, see about_CommonParameters and about_WorkflowCommonParameters.

Note

Troubleshooting Note: Because all workflow common parameters are available on workflows, there is no need to define these parameters on any workflow. If you redefine the **PSComputerName** parameter on a workflow, the parameter name is case-sensitive. Incorrect capitalization results in unintended fan-out behavior.

Activity Parameters

When calling a cmdlet or activity in a workflow, explicit parameter names are required, even for positional parameters. You can use a parameter alias or the first few characters of a parameter name, just so they are sufficient to distinguish the parameter name from the names of other parameters in the parameter set. However, you may not omit the parameter name.

Splatting is not permitted in workflow activities or in calls to workflows.

Dynamic parameters, which are parameters that are added to a command by a provider, module or another cmdlet, are not valid in workflows. To use a dynamic parameter, run the command in an InlineScript activity, as shown in the example below.

```
workflow Test-Workflow

{

    InlineScript

    {

        $Signing = dir cert:\CurrentUser -CodeSigningCert

    }

}
```

Activity Common Parameters

Windows PowerShell Workflow adds a set of activity common parameters to the InlineScript activity and to cmdlets that are implemented as activities. For a list of the activity common parameters, see Using Activities in Script Workflows.

Some of the workflow common parameters are also activity common parameters. This feature lets you create activity-specific exceptions to the values of the workflow common parameter values.

The activity common parameters are not valid on the **Suspend-Workflow** and **Checkpoint-Workflow** activities. Also, the activity common parameters are not valid on cmdlets or expressions that Windows PowerShell Workflow automatically runs in an activity. The activity common parameters are available on the **InlineScript** activity, but not on commands in the **InlineScript** script block.

Commands Run in Independent Sessions

Each command in a workflow runs in its own session or "runspace" in the workflow process. Windows PowerShell Workflow handles the processing; you do not need to manage these sessions.

Typically, commands that run in different sessions cannot share data, such as the values of variables created in the session. However, Windows PowerShell Workflow adds all workflow variables (variables in the top-level workflow scope) to every command session. For information about getting other variables values, see Using Variables in Script Workflows.

However, if you run commands in a workflow that change the current session, such as a command that imports a module into the current session, that command does not affect the sessions in which subsequent workflow commands run. To import a module that a command requires (and one that is not imported automatically), use the **PSRequiredModules** activity parameter.

The commands in an **InlineScript** script block are an exception to the independent session rule. By default, these commands run in a single session in their own Windows PowerShell process. For more information about the **InlineScript** activity, see Running Windows PowerShell Commands in a Workflow.

No Method Calls

You cannot invoke methods in a workflow. The activities in workflows are converted to XAML and they return serialized (XML-formatted) representations of objects. These objects have properties and property values, but the methods are not available.

To call a method, place the command in an **InlineScript** activity.

```
workflow Get-PSModulePath

{

    InlineScript

    {

        [System.Environment]::GetEnvironmentVariable("PSModulePath")

    }
```

```
}
```

Statements in a Workflow

Most statements that are valid in scripts and functions are also valid in workflows. The following list describes statements that work differently in workflows than they work in scripts or functions.

- **Begin**, **Process**, and **End** statements are not valid in workflows.

- Dot-sourcing is not valid in workflows.

- The Invocation (or Call) operator (&) is not valid in workflows.

- **Switch** statements must use the **CaseSensitive** parameter. Switch clauses must be case-sensitive. Switch clauses must include only constants. They cannot include comparison statements, regular expressions, file references, or script blocks. All Switch clauses must be of the same type or be strongly typed.

- Trap statements are not valid in workflows. Instead, use **Try-catch-finally** statements, which are valid.

- The **#Requires** statement has an **Assembly** parameter that adds the activities defined in an assembly to the workflow. The **Assembly** parameter is valid only in workflows.

  To specify an assembly, enter a fully qualified module description or an absolute or relative path. If the assembly is in the global assembly cache (GAC), enter only the assembly name. The following formats are valid.

  ```
  #requires -Assembly MyWorkflow.dll
  #requires -Assembly C:\Users\User01\MyWorkflow.dll
  #requires -Assembly
  C:\Windows\Microsoft.NET\assembly\GAC_MSIL\Accessibility\v4.0_4.
  0.0.0__12a3b4e5adc7\MyWorkflow.dll
  #requires -Assembly "MyWorkflow.dll, Version=1.1.0.0,
  Culture=neutral, PublicKeyToken=123a324b243c2"
  ```

Control Statements in a Workflow

Use the following rules to write valid control statements in a workflow.

- You cannot change variable values in loop conditions. To change a variable value, place the change statement in the loop body.

  The following statement is not valid in a workflow.

```
while($i++ -lt 100){if ($i -gt 10) Write-Host $i}
```

Instead, change the variable value in the process block of the loop, as shown in the following example.

```
while($i -lt  99){if ($i++ -gt 10) Write-Host $i}
```

- **Break** and **Continue** statements are not valid in workflows. Instead, use an **If** statement to control loop execution.
- Loop labels are not valid in workflows.

Reserved Words in Workflows

Reserved words for workflows are words that cannot be used as the names of workflows or of workflow variables or parameters, because they are used by Windows PowerShell or Windows PowerShell Workflow.

In addition to the reserved words for Windows PowerShell functions and scripts, the following categories of words are reserved in workflows.

1. Workflow keywords: **workflow**, **parallel**, and **sequence**.
2. The names of workflow runtime variables. For a complete list, see PSWorkflowRuntimeVariable Enumeration.

   The names of workflow common parameters. For a complete list, see about_WorkflowCommonParameters.
3. The names of activity common parameters. For a complete list, see "Activity Common Parameters" in Using Activities in Script Workflows.
4. The names of Windows PowerShell common parameters. For a complete list, see about_CommonParameters.

The names of commands, including cmdlets, functions, CIM commands, workflows, and activities, are not reserved. However, workflow and function names can be used only once in each workflow.

If two commands in the session have the same name but different command types, the command precedence rules determine which command runs when the command name is called. If two commands in the session have the same name and the same command type, the command added to the session last runs when the command name is called. To distinguish between commands with the same name, or to run a command that is hidden by a command with higher precedence, use the module-qualified name of the command (<ModuleName\CommandName>, such as PSScheduledJob\New-JobTrigger.

For more information about command precedence rules, see about_Command_Precedence.

Using Provider Drives in Workflows

All drives added by Windows PowerShell included providers are valid in workflows. The core providers are Alias, Certificate, Environment, FileSystem, Function, Registry, Variable, and WS-Management.

To use a drive added by provider in an installed module, use the **PSRequiredModules** activity parameter or use an Import-Module command in an **InlineScript** activity.

For example, the following commands use both techniques to access the IIS: drive in a workflow. The IIS: drive is added by the IIS provider in the WebAdministration module.

```
workflow Get-IIS { dir -Path IIS: -PSRequiredModules WebAdministration }


workflow Get-IIS { InlineScript { Import-Module WebAdministration; dir IIS: } }
```

Using Activities in Script Workflows

This topic lists the Windows PowerShell Workflow activities and workflow activity parameters. It lists the Windows PowerShell cmdlets that are implemented as activities for Windows PowerShell Workflow and those that are excluded.

Script Workflow Activities

The following table lists activities that are add to Windows PowerShell for Windows PowerShell Workflow. They are valid only in workflows.

Activity	Description	Topic Reference
CheckPoint-Workflow (alias = PSPersist)	Takes a checkpoint. Saves the state and data of a workflow in progress. If the workflow is interrupted or rerun, it can restart from any checkpoint. Use the Checkpoint-Workflow activity along with the PSPersist workflow common parameter and the PSPersistPreference variable to make your workflow robust and recoverable.	about_Checkpoint-Workflow, Adding Checkpoints to a Script Workflow
ForEach -Parallel	Runs the statements in the script block once for each item in a collection. The items are processed in parallel. The statements in the script block run sequentially.	about_Foreach-Parallel

Parallel	All statements in the script block can run at the same time. The order of execution is undefined.	about_Parallel
Sequence	Creates a block of sequential statements within a parallel script block. The **Sequence** script block runs in parallel with other activities in the **Parallel** script block. However, the statements in the **Sequence** script block run in the order in which they appear. Valid only within a **Parallel** script block.	about_Sequence
Suspend-Workflow	Stops a workflow temporarily. To resume the workflow, use the Resume-Job cmdlet.	about_Suspend-Workflow

Cmdlet Activities

All cmdlets in the following Windows PowerShell included modules are available as activities, except for those listed in the Excluded Cmdlets section. All other cmdlets and functions in **Microsoft.PowerShell.*** modules are implemented as activities for use in a workflow.

- Microsoft.PowerShell.Core
- Microsoft.PowerShell.Host
- Microsoft.PowerShell.Diagnostics
- Microsoft.PowerShell.Management
- Microsoft.PowerShell.Security
- Microsoft.PowerShell.Utility
- Microsoft.Wsman.Management

If a cmdlet is not in a Windows PowerShell Core module, but it is not specifically excluded, Windows PowerShell Workflow automatically runs the cmdlet in an **InlineScript** activity and returns the output to the workflow.

Activities that Differ from Cmdlets

The following Windows PowerShell Core cmdlets are implemented as activities for use in Windows PowerShell Workflow. However, the behavior of the activity differs somewhat from the behavior of the cmdlet.

- **Invoke-Expression**

You can use the Invoke-Expression cmdlet to include inline XAML in a script workflow. In workflows, the **Invoke-Expression** cmdlet has a **Language** parameter that takes a value of "XAML". Use the **Command** parameter to specify the path to the XAML file, as shown in the following example.

```
Workflow Test-Workflow
{
    Invoke-Expression -Language XAML -Command C:\ps-test\Get-
AssetData.xaml
}
```

For more information about the **Invoke-Expression** cmdlet, see Invoke-Expression.

- **New-Object**

  In workflows, only the **Typename** and **Argument** parameters are valid.

  For more information about the New-Object cmdlet, see New-Object.

- **Restart-Computer**

  In workflows, the **Wait** parameter is valid when restarting the local computer, as well as remote computers.

Excluded Cmdlets

The following Windows PowerShell cmdlets have not been implemented as workflow activities. Because they are excluded, Windows PowerShell Workflow does not automatically run these cmdlets in an **InlineScript** activity automatically. Typically, these cmdlets are not well suited to workflow tasks.

To run these cmdlets, include them in an **inlineScript** activity. The following example shows a call to the Get-Variable cmdlet in an **inlineScript** activity.

```
Workflow Test-Workflow
{
     InlineScript { Get-Variable}
}
```

Excluded Cmdlets

Add-History	Invoke-History

Add-PSSnapin	New-Alias
Clear-History	New-Variable
Clear-Variable	Out-GridView
Complete-Transaction	Remove-PSBreakpoint
Debug-Process	Remove-PSSnapin
Disable-PSBreakpoint	Remove-Variable
Enable-PSBreakpoint	Set-Alias
Enter-PSSession	Set-PSBreakpoint
Exit-PSSession	Set-PSDebug
Export-Alias	Set-StrictMode
Export-Console	Set-TraceMode
Get-Alias	Set-Variable
Get-History	Start-Transaction
Get-PSBreakpoint	Start-Transcript
Get-PSCallStack	Stop-Transcript
Get-PSSnapin	Trace-Command
Get-Transaction	Undo-Transaction
Get-Variable	Use-Transaction
Import-Alias	

Activity Common Parameters

Windows PowerShell Workflow adds the activity common parameters to activities that are derived from the **PSActivity** base class. This category includes the **InlineScript** activity and Windows PowerShell cmdlets that are implemented as activities, such as Get-Process and Get-WinEvent.

The activity common parameters are not valid on the **Suspend-Workflow** and **Checkpoint-Workflow** activities and they are not added to cmdlets or expressions that Windows PowerShell Workflow automatically runs in an **InlineScript** or similar activity. The activity common parameters are available on the **InlineScript** activity, but not on commands in the **InlineScript** script block.

The following table lists the activity common parameters, showing the parameters that are shared with the workflow common parameters and with the Windows PowerShell common parameters, as well as the parameters that are unique to activities.

For information about the workflow common parameters, see about_WorkflowCommonParameters. For information about the Windows PowerShell common parameters, see about_CommonParameters. The parameters that are unique to activities are described in this topic.

The activity common parameters are extremely useful. For example, you can use the **PSComputerName** parameter to run particular activities on only a subset of the target computers. Or, you might use the **PSConnectionRetryCount** and **PSConnectionRetryIntervalSec** parameters to adjust the retry values for particular activities.

The following example shows how to use the workflow common parameters as variables in the workflow and as parameters at the activity level. This command runs the **Get-Eventlog** cmdlet only on computers in the User01 domain.

```
Workflow Test-Workflow

{

     $User01Domain = Get-Content -Path .\User01Computers.txt

     $Log = (Get-EventLog -LogName "Windows PowerShell" -PSComputerName $UserDomain}

     if ($Log) ...

}
```

Workflow Common Parameters
PSAllowRedirection
PSApplicationName
PSAuthentication
PSCertificateThumbprint
PSComputerName
PSConfigurationName
PSConnectionRetryCount
PSConnectionRetryIntervalSec
PSConnectionURI
PSCredential
PSPersist
PSPort
PSSessionOption
PSUseSSL

Windows PowerShell Common Parameters
Debug
ErrorAction
Verbose
WarningAction

Activity-Specific Common Parameters

Name	Definition
AppendOutput <SwitchParameter>	Adds the output of the activity to the value of the variable. By default, assigning a value to a variable replaces the variable value. For example, the following commands add a process object to the service object in the $x variable. `workflow Test-Workflow { $x = Get-Service; $x = Get-Process -AppendOutput }`
DisplayName <String>	Specifies a friendly name for the activity. The **DisplayName** value appears in the progress bar while the workflow runs and in the value of the Progress property of the workflow job. When the **PSProgressMessage** parameter is also included in the command, the progress bar content appears in `<DisplayName>:<PSProgressMessage>` format.
Input <Object[]>	Submits a collection of objects to an activity. This is an alternative to piping objects to the activity one at a time.
MergeErrorToOutput <SwitchParameter>	Adds errors to the output stream. Use this parameter with **Parallel** and **ForEach -Parallel** keywords to collect errors and output from multiple parallel commands in a single collection.
PSActionRetryCount <Int32>	Tries repeatedly to run the activity if the first attempt fails. The default value, 0, does not retry.
PSActionRetryIntervalSec <Int32>	Determines the interval between action retries in seconds. The default value, 0, retries the action immediately. This parameter is valid only when the **PSActionRetryCount** parameter is also used in the command.
PSActionRunningTimeoutSec <Int32>	Determines how long the activity can run on each target computer. If the activity does not complete before the timeout expires, Windows PowerShell Workflow generates a terminating error and

	stops processing the workflow on the affected target computer.

PSDebug <PSDataCollection[DebugRecord]>	Adds debug messages from the activity to the specified debug record collection, instead of writing the debug messages to the console or to the value of the **Debug** property of the workflow job. You can add debug messages from multiple activities to the same debug record collection object.
	To use this activity common parameter, use the New-Object cmdlet to create a **PSDataCollection** object with a type of **DebugRecord** and save the object in a variable. Then, use the variable as the value of the **PSDebug** parameter of one or more activities, as shown in the following example.
	`workflow Test-Workflow`
	`{`
	`$debugCollection = New-Object -Type System.Management.Automation.PSDataCollection[System.Management.Automation.DebugRecord]`
	`InlineScript {\\Server01\Share01\Get-AssetData.ps1} -PSDebug $debugCollection -Debug $True`
	`InlineScript {\\Server01\Share01\Set-AssetData.ps1} -PSDebug $debugCollection -Debug $True`
	`if ($debugCollection -like "*Missing*") { ...}`
	`}`

PSDisableSerialization <Boolean>	Directs the activity to return "live" (not serialized) objects to the workflow. The resulting objects have methods, as well as properties, but they cannot be saved when a checkpoint is taken.

PSError <PSDataCollection[ErrorRecord]>	Adds error messages from the activity to the specified error record collection, instead of writing the error messages to the console or to the value of the **Error** property of the workflow job. You can add error messages from multiple activities to the same error record collection object.
	To use this activity common parameter, use the New-Object cmdlet to create a **PSDataCollection** object with a type of **ErrorRecord** and save the object in a variable. Then, use the variable as the value of the **PSError** parameter of one or more activities, as shown in the following example.
	`workflow Test-Workflow`
	`{`
	`$errorCollection = New-Object -Type`

	System.Management.Automation.PSDataCollection[System.Management.Automation.ErrorRecord]
	InlineScript {\\Server01\Share01\Get-AssetData.ps1} -PSError $errorCollection
	InlineScript {\\Server01\Share01\Set-AssetData.ps1} -PSError $errorCollection
	if ($errorCollection.Count -gt 2) { ...}
	}

PSProgress <PSDataCollection[ProgressRecord]>	Adds progress messages from the activity to the specified progress record collection, instead of writing the progress messages to the console or to the value of the **Progress** property of the workflow job. You can add progress messages from multiple activities to the same progress record collection object.
PSProgressMessage <String>	Specifies a friendly description of the activity. The **PSProgressMessage** value appears in the progress bar while the workflow runs. When the **DisplayName** is also included in the command, the progress bar content appears in <DisplayName>:<PSProgressMessage> format. This parameter is particularly useful for identifying activities in a **ForEach -Parallel** script block. Without this message, activities in all parallel branches are identified by the same name.
PSRemotingBehavior <RemotingBehavior>	Specifies how remoting is managed when the activity is run on target computers. **PowerShell** is the default value. Valid values are: **None**: The activity is not run on remote computers. **PowerShell**: Windows PowerShell remoting is used to run the activity on target computers. **Custom**: The activity supports its own type of remoting. This value is valid when the cmdlet that is being implemented as an activity sets the value of the **RemotingCapability** attribute to **SupportedByCommand** and the command includes the **ComputerName** parameter.
PSRequiredModules <String[]>	Imports the specified modules before running the command. Enter the module names. The modules must be installed on the computer that is running the workflow. Modules that are installed in a path specified in the **PSModulePath** environment variable are automatically imported on first use of any command in the module. Use this parameter to import modules that

	are not in a **PSModulePath** location. Because each activity in a workflow runs in its own session, an Import-Module command imports a module only into the session in which it runs. It does not import the module into sessions in which other activities run.
PSVerbose <PSDataCollection[VerboseRecord]>	Adds verbose messages from the activity to the specified verbose record collection, instead of writing the verbose messages to the console or to the value of the **Verbose** property of the workflow job. You can add verbose messages from multiple activities to the same verbose record collection object.
PSWarning <PSDataCollection[WarningRecord]>	Adds warning messages from the activity to the specified warning record collection, instead of writing the warning messages to the console or to the value of the **Warning** property of the workflow job. You can add warning messages from multiple activities to the same warning record collection object.
Result	This parameter is valid only in XAML workflows.

UseDefaultInput <Boolean>	Accepts all workflow input as input to the activity by value.
	For example, the Get-Process activity in the following sample workflow uses the UseDefaultInput activity common parameter to get input that is passed to the workflow. When you run the workflow with input, that input is used by the activity.
	`workflow Test-Workflow`
	`{`
	`  Get-Service -UseDefaultInput $True`
	`}`
	
	`PS C:\>Test-Workflow -InputObject WinRm`
	
	`Status  Name  DisplayName  PSComputerName`
	`------  ----  -----------  --------------`
	`Running  winrm  Windows Remote Management (WS-Manag... localhost`

Using Variables in Script Workflows

This topic explains the rules for using variables in Windows PowerShell Workflow.

Automatic Variables in Workflows

All Windows PowerShell automatic variables are valid in workflows, except for the automatic variables in the following list. For a complete list of automatic variables, see about_Automatic_Variables.

- $Args
- $Error
- $MyInvocation
- $PID
- $PSBoundParameters
- $PsCmdlet
- $PSCommandPath
- $PSScriptRoot
- $StackTrace

In addition, Windows PowerShell Workflow adds the Workflow Runtime Variables to every workflow. For more information, see the PSWorkflowRunTimeVariable Enumeration help topic in the MSDN Library.

Windows PowerShell Preference Variables

All Windows PowerShell preference variables are valid in workflows. For a complete list of preference variables, see about_Preference_Variables.

Workflow Preference Variables

Workflow preference variables are valid only in a workflow. Unlike Windows PowerShell preference variables that are set in a session and affect all commands that are run in the session, workflow preference variables are set in a workflow and affect only that workflow.

The following sections list and explain the workflow preference variables.

$PSParentActivityId <string>

Establishes an identifier for activities in the current scope and all child scopes. The identifier appears in the progress records for activities in the affected scopes.

For example, in the following workflow, the Get-Data nested workflow includes the PSParentActivityID preference variable with a value of "GetData_Section". That ID will appear on

all progress records of all activities run in the Get-Data nested workflow and in any nested workflows or functions in that workflow, such as the Get-SourceData workflow.

```
Workflow Test-Workflow

{

    Workflow Get-Data

    {

        $PSParentActivityId = "GetData_Section"

        ...


        Workflow Get-SourceData

        {

          ...

        }

    }


    Get-Data

}
```

$PSPersistPreference <Boolean>

Adds a checkpoint to a group of activities. A value of $True takes a checkpoint after every activity that follows the variable definition. This value is equivalent to setting the value of the PSPersist activity parameter to $true for all affected activities. To turn off the behavior, set the value of the variable to $False.

In the following workflow, the $PSPersistPreference variable takes a checkpoint after Activity 2, 3, and 4.

```
Workflow Test-Workflow

{

    Activity1

    $PSPersistPreference = $true

    Activity2

    Activity3

    Activity4

    $PSPersistPreference = $false

    Activity5
```

```
}
```

This variable takes a Boolean value. Valid values are:

1. **$True**: Adds a checkpoint after each activity, in addition to checkpoints that are specified in the workflow. It affects only activities that follow the variable definition.
2. **$False**: Stops adding checkpoints. Checkpoints are taken when specified in the workflow

$PSRunInProcessPreference

The workflow variable that controls whether activities are run in the process. If this variable is specified, all activities in the enclosing scope are run in the workflow process.

Workflow Runtime Variables and Common Parameter Variables

Windows PowerShell Workflow adds the workflow runtime variables to every workflow. The runtime variables include the workflow common parameters and their values, so you do not have to redefine those variables in your workflow. For more information about the Workflow Runtime Variables, see PSWorkflowRuntimeVariable Enumeration. For more information about the workflow common parameters, see about_WorkflowCommonParameters.

For example, the following workflow uses the $PSComputerName, $ParentJobName, and $JobName variables. The values of these variables are passed to the workflow and are available for use in the workflow.

```
Workflow Test-Workflow

{

    $PSComputerName

    $ParentJobName

    $JobName

}


PS C:\>Test-Workflow -PSComputerName Server01, Server12 -JobName TestWFJob -AsJob

Id     Name            PSJobTypeName   State         HasMoreData     Location

--     ----            -------------   -----         -----------     --------

83     TestWFJob       PSWorkflowJob   Completed     True            Server01,.



PS C:\> Receive-Job -Keep -ID 83

Server01
```

```
TestWFJob
Job84
Server12
TestWFJob
Job85
```

The workflow runtime variables include other very useful variables. For example, the following workflow uses a few of the runtime variables to get information about the values of session options and the version of Windows PowerShell running on the target computer.

```
Workflow Get-SessionOptions
{
    "Computer: $PSComputerName"
    "Parent Command Name: $ParentCommandName"
    $p = $PSVersionTable.PSVersion.Major
    "PSVersion: $p"
    "Session Options:"
    $PSSessionOption
}

PS C:\ps-test> Get-SessionOptions -PSComputerName Server01
Computer: Server01
Parent Command Name: Test-Workflow
PSVersion: 4
Session Options:
PSComputerName                    : localhost
PSSourceJobInstanceId             : b8ecbf21-a49e-4326-a0f1-4829fcdff4a8
MaximumConnectionRedirectionCount : 5
NoCompression                     : False
NoMachineProfile                  : False
ProxyAccessType                   : None
ProxyAuthentication               : Negotiate
ProxyCredential                   :
SkipCACheck                       : False
```

```
SkipCNCheck                         : False
SkipRevocationCheck                 : False
OperationTimeout                    : 00:03:00
NoEncryption                        : False
UseUTF16                            : False
IncludePortInSPN                    : False
OutputBufferingMode                 : None
Culture                             :
UICulture                           :
MaximumReceivedDataSizePerCommand :
MaximumReceivedObjectSize           : 209715200
ApplicationArguments                :
OpenTimeout                         : 00:03:00
CancelTimeout                       : 00:01:00
IdleTimeout                         : -00:00:00.0010000
```

To change the values of workflow runtime variables within a workflow, use the Set-PSWorkflowData activity.

For example, the following workflow includes Set-PSWorkflowData commands that increment the value of the PSConnectionRetryCount workflow common parameter and reverse the value of the Verbose common parameter from within the workflow.

```
workflow Test-Workflow
{
    "Before: PSConnectionRetryCount is $PSConnectionRetryCount."
     Set-PSWorkflowData -PSConnectionRetryCount ($PSConnectionRetryCount + 1)
    "After: PSConnectionRetryCount is $PSConnectionRetryCount."


    "Before: Verbose is $verbose."
    if ($Verbose) {Set-PSWorkflowData -OtherVariableName Verbose -Value $False}
    else {Set-PSWorkflowData -OtherVariableName Verbose -Value $True}
    "After: Verbose is $Verbose."
}
PS C:\> Test-Workflow -PSConnectionRetryCount 4 -Verbose
Before: PSConnectionRetryCount is 4.
After: PSConnectionRetryCount is 5.
```

```
Before: Verbose is True.

After: Verbose is False.
```

Rules for Variables in Workflows

This section explains the general rules for using variables in workflows. The sections that follow explain the rules that are specific to activities, loops, and other control statements.

- You cannot store the results of sub-expressions in a variable, because sub-expressions cannot be persisted.

```
#Incorrect:
$a = $(Get-Process | foreach {$_.Name)


#Correct:
$a = (Get-Process | foreach {$_.Name)
```

- You cannot assign a variable value to another variable in the same statement.

```
#Incorrect:
$a = $b = 23


#Correct:
$b = 23
$a = $b
```

- You cannot assign a value to an environment variable.

```
#Incorrect:
$env:Temp = "C:\Tmp"


#Correct:
InlineScript
{[System.Environment]::SetEnvironmentVariable("C:\Tmp")}
```

- You cannot use variable references to change property values.

```
#Incorrect:
$o = New-PSSessionOption
$o.SkipCACheck = $true


#Correct:
```

```
$o = New-PSSessionOption -SkipCACheck
```

Rules for Variables in Loops, Inline Scripts, and Parallel statements

Variables that you create in the top-level scope of a Windows PowerShell workflow are called *workflow variables*. Like the variables in a script or function, the commands and expressions in the top-level scope of a workflow can view and change the value of variables that other commands and expressions create, as shown in the following example.

```
Workflow Test-Workflow

{

$a = 23

$a++

$a

}


PS C:\>Test-Workflow

24
```

However, nested scopes, such as loops, inline scripts, and Parallel and Sequence statements, have slightly different rules.

Note

Each command in a workflow runs in its own session ("runspace"). Therefore, without intervention, the command would not be able to view or change the variables created by other commands. However, Windows PowerShell Workflow adds all workflow variables to all of the sessions in which top-level commands run. This feature allows the commands to share variables and view and change their values.

The following table shows the variable rules for different types of nested statements in workflows.

Rules for Nested Statements in Workflows	**View workflow variable values**	**Change workflow variable values**	**Internal variables are visible in workflow scope**
Loops/Control Statements	Yes	Yes	Yes
Parallel / Sequence	Yes	No. Use $Workflow.	No, return the variable.

InlineScript	No. Use $Using.	No. Return the variable. $Workflow is invalid.	No, return the variable.

Variables in Loops and Control Statements

Loop statements, such as those that use the **For**, **While**, and **Do** keywords, and control statements, such as those that use the **If** and **ForEach** keywords, share workflow variables. These statements can view and change workflow variables and the top-level commands and expressions in a workflow can view and change the variables that are created in the statements.

The following example shows that the **ForEach** statement can access the $a workflow variable and the top-level commands in the workflow can access the $i variable that is defined in the **ForEach** statement.

```
Workflow Test-Workflow

{

$a = 1,2,3


foreach ($number in $a)

       {

         "A = $number"

         $i++

       }

$i

}


PS C:\>Test-Workflow

A = 1

A = 2

A = 3

3
```

Variables in Parallel and Sequence Statements

The statements in **Parallel** and **Sequence** script blocks can view the values of workflow variables. However, by default, the statements in **Parallel** and **Sequence** script blocks cannot change the values of workflow variables and cannot create variables that have the same name as workflow variables.

To change the value of a workflow variable in a **Parallel** or **Sequence** block, use the **$Workflow** scope modifier. When you do, the value that is assigned in the **Parallel** or **Sequence** block is available in the workflow scope.

The following example shows how to use the **$Workflow** scope modifier to change the value of a workflow variable and to make the new variable value available in workflow scope.

```
Workflow Test-Workflow
{
    #Workflow (top-level) variable
    $w = 4

    Parallel
    {
        Sequence
        {
            # You can get the values of workflow variables.
            "InSequenceW = $w"

            # To change the value of a workflow variable,
            # use the $Workflow scope modifier.
            $Workflow:w = 5
            "InSequenceWNew = $w"
        }
    }
    # The change is visible in workflow scope.
    "WorkflowW = $w"
}

PS C:\>Test-Workflow
InSequenceW = 4
InSequenceWNew = 5
WorkflowW = 5
```

Variables that are defined in a child scope, such as a Parallel or Sequence block, are not visible in any parent scope. To make them visible, return the value to the parent scope, as shown in the following example.

```
Workflow Test-Workflow
{
    Parallel
    {
        Sequence
        {
            #Variables defined in parallel/sequence ...
            $s = 3
        }
    }
    # ... are not available in workflow scope.
    "S = $s"



    $TFromSequence = Parallel
    {
        Sequence
        {
            #To export variable values from
            # parallel/sequence, return the value.
            $t = 3
            $t
        }
    }
    "T = $TFromSequence"
}


PS C:\>Test-Workflow
S =
T = 3
```

Variables in an InlineScript Activity

The **InlineScript** activity runs commands and expressions in a standard Windows PowerShell session and returns the results to the workflow.

By default, the variables that are defined in a workflow are not visible to the commands in an **InlineScript** script block. To make workflow variables visible in **InlineScript**, use the **$Using** scope modifier. The **$Using** scope modifier is required only once for each variable in the **InlineScript** script block, as shown in the following example.

```
workflow Test-Workflow
{
    $a = 3

    # Without $Using, the $a workflow variable is not visible
    # in inline script.
    InlineScript {"Inline A0 = $a"}

    # Using imports the variable and its current value.
    InlineScript {"Inline A1 = $Using:a"}
}

PS C:\> Test-Workflow
Inline A0 =
Inline A1 = 3
```

Commands in an InlineScript script block can change the value of variables that were imported from workflow scope, but the changes are not visible in workflow scope. To make them visible, return the changed value to the workflow scope, as shown in the following example.

```
workflow Test-Workflow
{
    $a = 3

    # Changes to variable value in an InlineScript do not affect
    # the value of the workflow variable.
    InlineScript {$a = $Using:a + 1; "Inline A = $a"}
    "Workflow A = $a"

    # To change the value in workflow scope, return the new value.
     $a = InlineScript {$a = $Using:a+1; $a}
     "Workflow New A = $a"
```

```
}

PS C:\> Test-Workflow

Inline A = 4

Workflow A = 3

Workflow New A = 4
```

Note

Troubleshooting Note: A statement with the **$Using** scope modifier and the imported variable should appear before any use of the variable in the **InlineScript** script block. In Windows PowerShell 3.0 and newer releases of Windows PowerShell, the **$Using** scope modifier can appear anywhere in the **InlineScript** script block.

Configuring Your Workflow Environment

Now that you understand the definition and benefits of workflows in Windows PowerShell, as described in Getting Started with Windows PowerShell Workflow, you are ready to set up your environment to run workflows.

This topic contains the following sections.

- Supported operating systems and architectures
- Workflow configurations
- Prepare computers to run workflows

Supported operating systems and architectures

Windows PowerShell Workflow runs on any system on which Windows PowerShell 3.0 or Windows PowerShell 4.0 is available. For a complete list of Windows PowerShell requirements, see Windows PowerShell System Requirements.

Download and install Windows Management Framework 3.0 or Windows Management Framework 4.0 on managed nodes if your workflows contain Common Information model (CIM) activities. CIM is an extensible, object-oriented data model that contains information about different parts of an enterprise. For more information about CIM, see Common Information Model on MSDN. You can also install Windows Management Framework (WMF) on managed nodes that are running Windows Server 2008 R2 with Service Pack 1, Windows 7 with Service Pack 1, or Windows Management Framework 3.0 on Windows Server 2008 with Service Pack 2 to run workflows on those computers, but this is not required unless your workflows contain CIM activities.

Workflow configurations

Workflow operations involve the following three nodes. The nodes can all be located on the same computer, or distributed across multiple computers.

- The node on which the workflow is running, sometimes called the workflow server.
- The node on which an administrator is managing the workflow or viewing workflow status. This is sometimes called the *client*.
- Roles, features, settings, or physical or virtual devices on which the workflow actions are carried out. These are called managed nodes.

Computers that are running any of the supported operating systems and architectures can function as any part of a workflow configuration. The following diagram shows the possible setup configurations for workflows.

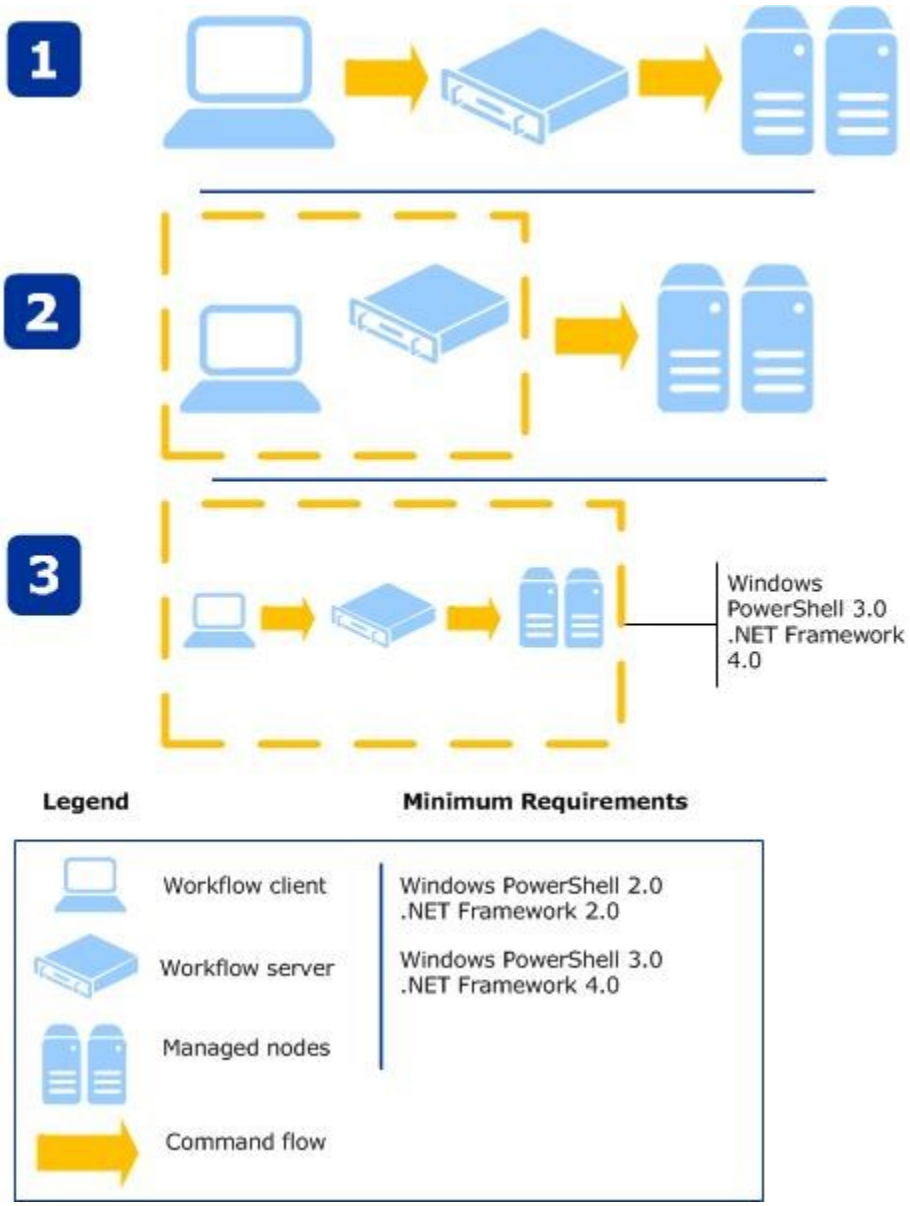

Configuration 1: The server, client, and managed nodes are running on separate computers. Note that the command flow is from client, through server, to managed nodes. In this scenario, the client node must be running Windows PowerShell 2.0 to communicate with the computer that is running the workflow, but Windows PowerShell 3.0 or Windows PowerShell 4.0 is not required on the client.

Configuration 2: The workflow server and client nodes are both running on the same computer. The versions of Windows PowerShell and .NET Framework that are running on that computer satisfy the highest requirements, those for the computer that is running the workflow. Managed nodes are running on one or more remote devices.

Configuration 3: The workflow server, client, and managed nodes are running on the same computer or device. The versions of Windows PowerShell and .NET Framework that are running on that computer satisfy the highest requirements, those for the workflow server.

Download locations for .NET Framework and Windows PowerShell

For information about download locations for required versions of .NET Framework and Windows PowerShell, see Windows PowerShell System Requirements.

Prepare computers to run workflows

Procedures for setting up your computers to run workflows are available in this section.

To prepare for remote management

1. On the computer from which you are running and managing the workflow, start Windows PowerShell with the **Run as administrator** option as described in Starting Windows PowerShell. Alternatively, you can run a new Windows PowerShell session as administrator by running the following command in a session that is already open.

   Start-Process PowerShell -Verb RunAs

2. Do one of the following to enable Windows PowerShell remote management on participating computers, such as the workflow server and managed nodes.

   Note
   While Windows PowerShell remote management must be enabled on the server node, you do not need to enable it on managed nodes unless your workflow requires it. Neither Windows PowerShell nor WMF is required on managed nodes, unless your workflow contains CIM activities.

   - On a workflow server that is running Windows Server 2012 R2 or Windows Server 2012, Windows PowerShell remote management is enabled by default. Go on to the next step.
   - On a workflow server that is running Windows 8.1 or Windows 8, right-click the **Windows PowerShell** tile, and then in the app bar, click **Run as administrator**. Type **Enable-PSRemoting**, and then press **Enter**. Press **Y** to confirm that you want to enable remote management.
   - On a workflow server that is running Windows Server 2008 R2 or Windows 7, Windows Management Framework 4.0 or Windows Management Framework 3.0 must be installed first. After WMF is installed, click **Start**, point to **All Programs**, click **Accessories**, click **Windows PowerShell**, right-click the **Windows PowerShell** shortcut, and then click **Run as Administrator**. Type **Enable-PSRemoting**, and then press **Enter**. Press **Y** to confirm that you want to enable remote management.
   - On a workflow server that is running Windows Server 2008, Windows Management Framework 3.0 must be installed first. After WMF 3.0 is installed, open a Windows PowerShell session with elevated user rights. To do this, click **Start**, click **All**

Programs, open the **Windows PowerShell** folder, right-click the **Windows PowerShell** shortcut, and then click **Run as administrator**. Type **Enable-PSRemoting**, and then press **Enter**. Press **Y** to confirm that you want to enable remote management.

Workflow session configurations

A session configuration is a group of settings on the local computer that define the access rights and commands that are available for the Windows PowerShell sessions that are created when remote or local users connect to Windows PowerShell on a computer.

The **Enable-PSRemoting** command that you run on the workflow server in the preceding section creates default session configurations, including the recommended and default session configuration for workflows. We recommend using the default session configuration because it provides the following benefits to your workflows automatically.

- **Performance optimization.** Default settings of the New-PSWorkflowExecutionOption cmdlet, such as activity time-out periods and maximum allowed processes per activity, are designed to help constrain resource usage and improve workflow performance.
- **SharedHost mode.** SharedHost mode allows users to start a workflow on one computer, log off, log on to another computer, and reconnect to the workflow. The workflow session configuration uses the same process for a user across different computers, as long as the user is connecting to the same session configuration, and is logged on to the new computer with the same access rights as the user had on the first computer.
- **Automatic multi-device support.** The workflow session configuration is optimized for running workflows on multiple devices, offering the following benefits.
  - Workflow jobs are automatically created for each target device.
  - Windows PowerShell Workflow automatically throttles the number of simultaneous workflows and remote operations that can run on a target device.
  - Manages lifecycle of remote connections – creation, reuse, and deletion

You can create a custom session configuration for your organization's needs, which might differ from the workflow session configuration. You can add your workflows to your custom session configuration, so they are available automatically when a user connects to the session configuration.

For more information about session configurations, see about_Session_Configurations.

The following procedure shows how to create a custom session configuration named **Microsoft.PowerShell.MySession**.

To create custom session configurations

1. Start a Windows PowerShell session with elevated user rights (Run as Administrator). The procedure To prepare for remote management in this topic describes how to do this.
2. We recommend increasing the default Windows Remote Management (WinRM) memory quota to 1 GB in sessions that are using the session configuration for workflows. To do this, run the following command to create a session configuration option object, and then

save the object in the **$MoreMemory** variable. The variable is used in the next step.

$MoreMemory = New-PSSessionConfigurationOption -MaxMemoryPerShellMB 1000

> **Note**
> The value shown in the preceding command is an example. You can use lower values; the minimum requirement is 400 MB.

3. To add the new session configuration, run the **Register-PSSessionConfiguration** cmdlet as shown in the following script sample.
   a. Set the value of the **SessionConfigurationType** parameter to **Workflow**.
   b. Set the value of the **ModulePath** parameter to the location of the modules that contain your workflows.
   c. Set the value of the **SessionConfigurationOption** parameter to the **$MoreMemory** variable that you created in the preceding step.

```
Register-PSSessionConfiguration -Name
Microsoft.PowerShell.MySession `
-SessionConfigurationType Workflow `
-ModulePath $pshome\modules\MyModule`
-SessionConfigurationOption $MoreMemory
```

See Also

Getting Started with Windows PowerShell Workflow

Running a Windows PowerShell Workflow

Writing a Script Workflow

Running a Windows PowerShell Workflow

After setting up your Windows PowerShell Workflow environment as described in Configuring Your Workflow Environment, and authoring or finding a workflow, you are ready to start running your workflow.

Running a workflow is like running any Windows PowerShell command. You can run the `Get-Command` cmdlet to discover workflows, and run `Get-Help` to get available help about workflows. Running a workflow is as simple as entering the *verb-noun* name of the workflow (or *verb-noun* `-PsComputerName` *managednode1*, *managednode2*).

This topic contains the following sections.

- Create a session to the computer that runs your workflow
- Connect or reconnect to a running workflow
- Run the workflow

- Run workflows as jobs
- View data about workflows
- Add checkpoints to a workflow
- Get help for your workflow
- Find error logs for workflows

Create a session to the computer that runs your workflow

You create a new session to connect to the computer that is running your workflow if you are starting the workflow from the beginning, or if you have closed (and not merely disconnected from) an earlier session that is connected to the workflow computer. If you are reconnecting to a workflow that is already in progress, you do not need to create a new session to connect to the computer that is running the workflow; you can connect to the existing session. The following procedure is an example of how to create a new session on a local or remote computer.

To create a new workflow session

1. Start Windows PowerShell with the **Run as administrator** option as described in Starting Windows PowerShell. You do not need to run Windows PowerShell with elevated user rights (Run as Administrator) if you are connecting from a remote computer.
2. Start a remote Windows PowerShell session that is connected to the computer that you want to run your workflow, and save the session in a variable. This session can be on your local computer (**localhost**) or on a remote computer. For example, the following commands start a remote session on the local computer, and save the session in the `$WFServer` variable. The session is configured to use the default workflow session configuration that is described in Configuring Your Workflow Environment.

   In the following example, the command starts a workflow session on the local computer, and saves the session in the **$WFServer** variable. By using the `New-PSWorkflowSession` cmdlet, the session is configured to use the default workflow session configuration.

   $WFServer = New-PSWorkflowSession

   In the following example, the command starts the workflow session on a different computer, *Server01*, and saves the session in the **$WFServer** variable.

   $WFServer = New-PSWorkflowSession -ComputerName Server01 -credential *DomainName\UserName*

   Because the workflow server session runs on a remote computer in the preceding example, the user must provide explicit credentials. Explicit credentials are not required if the session configuration includes **RunAs** functionality, where all users share a set of credentials that are defined in the session configuration.
3. Alternatively, you can create a new Windows PowerShell session by running the `New-PSSession` cmdlet, as described in New-PSSession. When you run `New-PSSession` instead of `New-PSWorkflowSession`, you can add the `ConfigurationName` parameter to specify a

session configuration other than the default Windows PowerShell Workflow session configuration. To specify the default Windows PowerShell Workflow session configuration, add `-ConfigurationName Microsoft.PowerShell.Workflow` to the command.

You can now run workflows in this session, either from a module, or by starting a standalone XAML or Windows PowerShell script-based workflow.

Connect or reconnect to a running workflow

If you are reconnecting to a workflow that you have already started, you can specify the job ID or name of the workflow to reconnect. You can also run the Connect-PSSession cmdlet to reconnect to a disconnected session that was connected to the computer that is running your workflow.

To connect or reconnect to a running workflow

1. Start a new session that is connected to the computer that is running your workflow. Save the session connection in a variable, as shown in the following example.

   $WFServer=New-PSWorkflowSession -ComputerName Server01

2. Get the list of jobs that are running on the computer that you specified in step 1.

   Invoke-Command $WFServer {Get-Job}

3. In the results, locate the workflow to which you want to reconnect.
4. If the workflow is not suspended, go on to the next step. To resume a suspended workflow, run the following command, where *$WFServer* represents the variable into which your session connection is saved.

   Invoke-Command $WFServer {Resume-Job <Job name or ID>}

5. To get information about the progress and status of a workflow, run `Receive-Job` as shown in the following example.

   Invoke-Command $WFServer {Receive-Job <Job name or ID>}

Run the workflow

The following procedure shows how to run a workflow that is packaged in a module. This step is required only if you did not specify a root module or a nested module key in the module manifest.

To import and run a workflow in a module

1. If you do not have an existing active session to the computer on which you want to run the workflow, create a new session as described in To create a new workflow session.
2. Copy the module directory that you created in Saving Your Workflow in a Module to the computer on which you are running the workflow. For convenience, copy it to the **Modules** directory (for example, `$env:C:\Users\<user_name>\Documents\WindowsPowerShell\Modules`) on that computer.
3. Optionally, import the module into the session that is connected to the computer running the workflow by using either of the following two commands. In Windows PowerShell 3.0,

modules are imported into a session automatically the first time a user runs a cmdlet that is part of the module. Add **Invoke-Command** to run an **Import-Module** command in the workflow server session. The optional **Verbose** parameter displays the functions that the module adds to the session.

Invoke-Command -Session $<name of session variable> -ScriptBlock {Import-Module <ModuleName> -Verbose}

For example, the following command imports the **RoleManager** module into the workflow server session.

Invoke-Command $WFServer {Import-Module RoleManager -Verbose}

The **Import-Module** cmdlet imports workflows into your session as functions. You can now run the workflows in the session that is connected to the computer on which you want to run your workflow.

4. Type the workflow name, and press **Enter** to start the workflow. For example, to start the workflow that was saved in a module in Saving Your Workflow in a Module, type **Install-Role -PSComputerName <managed_node_names>**, and then press **Enter**.

The following procedure explains how to run a XAML-based workflow that is not in a module in the session that is connected to the computer on which you want to run your workflow. After you import the workflow XAML files, you can use the workflows in your workflow server session. For more information about XAML-based workflows that you can use with Windows PowerShell Workflow, see “Creating and importing workflows using the Visual Studio Workflow Designer” in Getting Started with Windows PowerShell Workflow.

To import and run a workflow from a XAML workflow file

1. Start a workflow server session, as described in To create a new workflow session.
2. Run the following command to import the workflow XAML file into your workflow server session.

   Invoke-Command -Session $<name of session variable> -Scriptblock {Import-Module <Path_To_Xaml_File>}

   For example, the following command imports the **Install-Role** and **Set-Role** workflows into the workflow server session.

   Invoke-Command $WFServer {Import-Module -Path D:\Temp\Install-Role.xaml, D:\Temp\Set-Role.xaml}

Run workflows as jobs

You can use the **AsJob** workflow common parameter that is added to all workflows in Windows PowerShell to run a workflow as a Windows PowerShell job. For more information about jobs, see about_Jobs. You can use `Job` cmdlets to start, stop, suspend, or resume workflows. To run a workflow as a job, enter a command in the following format.

Verb-Noun -PSComputerName *ManagedNode01*, *ManagedNode02* -AsJob

Note

`Invoke-Command` and the variable in which your workflow session is stored (in this example, **$WFServer**) are necessary in the following commands only if you have not already created a session that is connected to the computer running your workflow. If you are already in that session, you need only to run the commands that are inside braces.

To run workflows as jobs

- The following are examples of running and managing workflows by using `Job` cmdlets.
  - Use the **AsJob** parameter to run the **Install-Role** workflow as a job. The **PSComputerName** parameter specifies the computers on which the roles are installed. In the following example, a workflow is run as a job.

    Invoke-Command $WFServer {Install-Role -PSComputerName Server01, Server02 -AsJob}
  - Get all the jobs that are running in the workflow server session. Each workflow instance is run as a separate job. The jobs are shown as running on the local computer (**localhost**), because the job object is created locally in the workflow server session, even though the job is running on two different remote computers.

    Invoke-Command $WFServer {Get-Job}

    The results are similar to the following.

```
Id   Name   State      HasMoreData   Location    Command
--   ----   -----      -----------   --------    -------
1    Job1   Running    True          Server01    install-role
2    Job2   Running    True          Server02    install-role
```

  - Get the workflow jobs that affect a specific managed node, by adding the `Filter` parameter.

    Invoke-Command $WFServer {Get-Job -Filter {PSComputerName = "Server01"}
  - Get the status of a job.

    Invoke-Command $WFServer { (get-job -name Job1).state}
  - Wait for a job to finish.

    Invoke-Command $WFServer {wait-job -name Job1}
  - Get the results of a job.

    Invoke-Command $WFServer {receive-job -name Job1}
  - Delete or remove a job.

    Invoke-Command $WFServer {remove-job -name Job1}

To suspend and resume workflows as jobs

- The following are examples of suspending and resuming workflows by using `Job` cmdlets.

- Suspend a workflow as a job.
  Invoke-Command $WFServer {suspend-job -name Job1}
- Resume a suspended workflow as a job.
  Invoke-Command $WFServer {resume-job -name Job1}

View data about workflows

Workflows are added to a session as functions. You can use the same commands that you would use to examine functions to examine your workflows.

Note
If you are already in the session that is connected to the computer on which your workflow is running, you do not need to add `Invoke-Command $variable`. Run only the commands that are inside braces.

To view information about workflows in your session

- The following commands are examples of information that you can get about your workflow.
  - Find functions that a workflow module added to your session.
    Invoke-Command $WFServer {get-command -module RoleManager}
  - Examine a workflow or function.
    Invoke-Command $WFServer {ls function:\Install-Role}
  - Use the `Get-Command` cmdlet to display the properties of a workflow function.
    Invoke-Command $WFServer {get-command Install-Role }
  - Use the `Get-Command` cmdlet to get the syntax of a workflow function.
    Invoke-Command $WFServer {get-command install-role -syntax}
  - Examine the objects that a workflow function returns.
    Invoke-Command $WFServer {install-role | get-member}

Add checkpoints to a workflow

You can add checkpoints to the workflows that you run and suppress its default checkpoints. You cannot remove or suppress explicit checkpoints in a workflow.

A *checkpoint* is a snapshot of the workflow state and data that is taken while the workflow runs. The checkpoint is saved to disk on the computer on which the workflow runs or to a SQL database. If the workflow is suspended or interrupted, it can be resumed from the most recent checkpoint, rather than restarting.

Most workflows have explicit checkpoints. Additional checkpoints are useful when the network is not stable, the workflow is frequently interrupted, or the workflow is running on a very large number of devices, However, adding too many checkpoints can make the workflow run very

slowly. Be sure that the time it takes to write the checkpoint data to disk is less than the time it would take to rerun the workflow if it is interrupted.

To adjust the checkpoints in a workflow, use the **PSPersist** workflow common parameter. The following are the valid values.

$True

Adds a checkpoint after each activity. These checkpoint are additions to the default checkpoints taken at the beginning and end of the workflow and any checkpoints that are explicitly specified in the workflow.

$False

Suppresses the default checkpoints that are taken at the beginning and end of the workflow. This value does not affect explicit checkpoints.

The following examples run the Set-WorkflowState workflow with the PSPersist parameter set to $True and $False.

```
PS C:\>Set-WorkflowState -PSPersist $True

PS C:\>Set-WorkflowState -PSPersist $False
```

Get help for your workflow

If Help is available for a workflow that you are running, you can get help about what the workflow does, and how to run it, by typing the following, and then pressing **Enter**.

Get-Help <*WorkflowVerb-WorkflowNoun*>

For more information about how to write help for a workflow, see Writing Help for a Script Workflow.

Find error logs for workflows

Error logs for workflows are stored on the computer that is running the workflow (sometimes called the workflow server). You can find them in the following Event Tracing for Windows (ETW) channel, or path in the Event Viewer snap-in. Event ID 45079 shows each activity that is run.

Applications and Services Logs\Microsoft\Windows\PowerShell

See Also

Getting Started with Windows PowerShell Workflow

Configuring Your Workflow Environment

Writing a Script Workflow

about_Jobs

Command-Line Reference for Windows Server 2012

The content in this collection will describe the command-line tools that enable IT pros to create batch files or scripting tools that automate common management tasks.

The following sections will contain information about the syntax and usage of the command-line tools.

- *A-Z List* contains an alphabetical list of all of the commands.
- *Command-Line Syntax Key* explains the formatting used in this reference.
- *Technology Lists* contains lists of commands by their associated technology.

For the complete collection, see **Command-line Reference**.

Windows Remote Management (WinRM) Overview

This topic provides introductory overview information about Windows Remote Management (WinRM). Also contained within this topic, are links to additional information pertaining to WinRM.

Did you mean…

- Windows Remote Management
- WS-Management Protocol on MSDN
- Remote Access (DirectAccess, Routing and Remote Access) Overview

WinRM description

Windows Remote Management (WinRM) is the Microsoft implementation of WS-Management Protocol, a standard Simple Object Access Protocol (SOAP)-based, firewall-friendly protocol that allows hardware and operating systems, from different vendors, to interoperate.

The WS-Management protocol was developed by a group of hardware and software manufacturers as a public standard for remotely exchanging management data with any computer device that implements the protocol.

Practical applications

WinRM is designed to improve hardware management in a network environment with various devices running a variety of operating systems. The entire design of the service is focused on monitoring and managing remote computers by implementing an interoperable standard protocol.

You can use WinRM scripting objects, the WinRM command-line tool, or the Windows Remote Shell command line tool WinRS to obtain management data from local and remote computers

that may have baseboard management controllers (BMCs). If the computer runs a Windows-based operating system version that includes WinRM, the management data is supplied by Windows Management Instrumentation (WMI).

You can also obtain hardware and system data from WS-Management protocol implementations running on operating systems other than Windows in your enterprise. WinRM establishes a session with a remote computer through the SOAP-based WS-Management protocol rather than a connection through DCOM, as WMI does. Data returned to WS-Management protocol are formatted in XML rather than in objects.

The Intelligent Platform Management Interface (IPMI) WMI provider is a standard WMI provider with classes that obtain baseboard management controller (BMC) sensor data from computers with appropriate hardware. IPMI data can be accessed using the WinRM scripting API, the WMI Scripting, or COM APIs. For more information about the WMI classes for which the IPMI provider supplies data, see Intelligent Platform Management Interface (IPMI) Classes.

See also

Following are links to other topics about WinRM.

Content type	References
Product evaluation	**WinRM IIS Extension Overview**
Planning	Content not available
Deployment	Content not available
Operations	Windows Remote Management Command-Line Tool (Winrm.cmd)
Troubleshooting	Content not available
Security	Content not available
Tools and settings	Content not available
Community resources	PowerShell Server vs. PowerShell Remoting Enable and Use Remote Commands in Windows PowerShell Understanding PowerShell Remote Management
Related technologies	WS-Management Protocol Understanding SOAP

Windows Management Instrumentation (WMI) Overview

This topic provides overview information about Windows Management Instrumentation (WMI). This topic also contains links to additional information that is related to WMI.

Did you mean…

- Windows Remote Management
- Windows Management Instrumentation (WMI) on the Microsoft Developer Network (MSDN)

WMI description

Windows Management Instrumentation (WMI) is the Microsoft implementation of Web-Based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components. CIM is developed and maintained by the Distributed Management Task Force (DMTF).

WBEM is an industry initiative, overseen by the Distributed Management Task Force (DMTF), to develop a standardized, nonproprietary means for accessing and sharing management information in an enterprise network. It is intended to alleviate the problems usually associated with collecting management and diagnostic data that can include different types of hardware, protocols, operating systems, and distributed applications.

CIM is an extensible, object-oriented data model that contains information about different parts of an enterprise. Through WMI, a developer can use the CIM to create classes that represent hard disk drives, applications, network routers, or even user-defined technologies, such as a networked air conditioner.

The ability to obtain management data from remote computers is what makes WMI useful. Remote WMI connections are made through DCOM. An alternative is to use Windows Remote Management (WinRM), which obtains remote WMI management data using the WS-Management SOAP-based protocol.

WMI offers a simple interface to management data that includes:

1. A rich and consistent model of the Windows operating system's behavior, configuration, and status.
2. A COM API that supplies a single point of access to all management information.
3. Interoperability with other Windows 2000 management services.
4. A flexible architecture that enables vendors to extend the information model to cover new devices, applications, and services.
5. A rich query language that enables SQL-based queries of the information model.
6. A scriptable API that allows seamless local and remote system administration.

The WMI behavior can be modified in **Server Manager** / **Tools** / **Computer Management** / **Services and Applications** / **WMI Control**, and in **Server Manager** / **Tools** / **Group Policy Management** / **Forest: (ForestName)** / **domains** / **DomainName** / **WMI Filters**.

Practical applications

By viewing and making changes to a CIM class, a manager can control different aspects of the enterprise. For example, a manager could query a CIM class instance representing a desktop workstation. The manager could then run a script to modify the CIM workstation instance. WMI would translate any change to the workstation CIM class instance into a change to the actual workstation.

See also

Following is a table of references related to WMI.

Content type	References
Product evaluation	Content not available
Planning	Content not available
Deployment	Content not available
Operations	WMI Scripting Primer: Part 1 WMI Scripting Primer: Part 2 WMI Scripting Primer: Part 3
Troubleshooting	Troubleshooting and Tips
Security	Content not available
Tools and settings	Content not available
Community resources	Probing the Depths of WMI, a Hey, Scripting Guy article.
Related technologies	Common Information Model (CIM) WMI

Datacenter Abstraction Layer (DAL) Overview

This topic provides overview information about the Datacenter Abstraction Layer (DAL). Within this topic you will find introductory information about the DAL: what the DAL is, and what it consists of. This topic also includes introductory information about how the DAL can be used to monitor and manage systems on your network, such as baseboard management controllers (BMCs), top-of-rack (TOR) switches, power distribution units (PDUs), and storage devices.

Did you mean…

- Hardware Abstraction Layer (HAL)
- Windows Remote Management (WinRM)
- Microsoft.WSMan.Management Module
- Configure Remote Management in Server Manager

DAL description

Datacenter abstraction layer (DAL) is an extensible framework that provides a set of management infrastructure components and schemas that can emulate device-specific details. DAL management infrastructure components provide management tools with direct access to the device management resources. The tools are standard and they can be automated, and Admins can use them without dealing with the complexity and specificity of the devices being managed.

The following image shows the relationship between the applications and services that use the DAL (shown in red), the DAL (shown in Blue), and managed devices (shown in green):

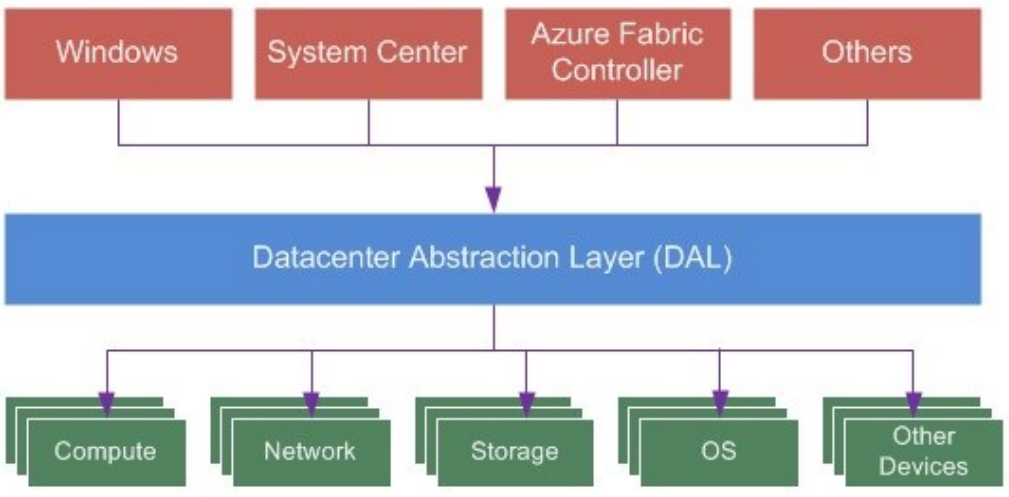

The DAL concept consists of a management infrastructure layer which implements the Common Information Model (CIM) as the modeling language, and Web Services for Management (WS-Man) as the preferred protocol.

Windows Server 2012, includes a lightweight, open-source, cross-platform, and scalable CIM object manager (CIMOM) architecture called Open Management Infrastructure (OMI), which implements the CIM plus WS-Man as the basis for managing devices. The goal of OMI is to help device vendors expedite the adoption of CIM and WS-Man as the basic technologies for data center device management.

Common Information Model

To implement DAL, the industry must agree on a set of standards and technologies to abstract this layer. The Common Information Model (CIM) as defined by the Distributed Management Task Force (DMTF) provides the means to actively control and manage hardware and software resources. By using a common model of information, management software can work with many implementations of the common model, without complex and costly conversion operations or loss of information.

CIM is a common definition of management information for a huge range of systems, networks, applications, and services. CIM defines the following:

- **Managed Object Format** (MOF): A language that is based on the interface description language (IDL), which is used to describe management objects
- **CIM Schema**: A standardized model of management objects that is written in MOF syntax (for example, Processes, Computers, and Printers)
- **CIM Profiles**: Collections of CIM models and associated behaviors for a particular management domain (for example, Power, Virtualization, and Storage)

WS-Man

WS-Man is designed to work with CIM objects to help perform all management operations securely and remotely. WS-Man is a Simple Object Access Protocol (SOAP) that is standardized by the DMTF based on:

- WS-Addressing
- WS-Transfer
- WS-Enumerating
- WS-Eventing
- WS-Management CIM Binding Specification
- WS-CIM Mapping Specification

Windows Remote Management (WinRM) is the Microsoft implementation of WS-Man, which provides the following for clients and services:

- Windows PowerShell provider for configuration
- Windows PowerShell cmdlets to enable access to any remote management endpoint that implements WS-Man plus CIM
- Windows PowerShell remoting

Practical applications

As IT organizations consolidate, evolve, and operate as service providers for their customers, they must deploy service-oriented infrastructures and management processes. Combined with the expansion of server and storage virtualization, the trend is increasingly placing pressure on organizations to provide greater levels of service while simultaneously reducing costs.

Data centers today are made up of different classes of devices from a variety of device manufacturers. The rapid expansion of data centers and cloud computing increases the importance of consistently managing these devices. Prior to DAL, most of the devices in data centers were managed by using protocols and schemas, and for some, by proprietary solutions. The lack of a consistent management platform makes the task of management complicated, inefficient, and time consuming.

For example, to meet its growing system and maintenance needs, Windows Azure acquires a variety of devices from a variety of manufacturers. The Windows Azure Fabric Controller is responsible for managing a variety of data center hardware such as blades, load balancers, PDUs, switches, and routers. Many of those devices are managed by using different protocols, schemas, and proprietary solutions in some instances. To simplify device management, Windows Azure includes a plug-in system to abstract those complexities. However, Windows Azure also requires vendor-specific solutions that are built on their plug-in model, which results in an added layer of management complexity and reliability issues.

Meeting the challenges that face complex, growing data center environments requires a management approach that is easy to use and facilitates end-to-end management operations. To address these issues, DAL provides a consistent developer abstraction layer to interact with the wide array of DAL-compliant devices from a variety of vendors. This is similar to the hardware abstraction layer (HAL) that Windows provides for hardware application developers.

For system administrators:

DAL provides a consistent data center infrastructure management platform by using Windows PowerShell.

For an example, let’s look at the job of an IT pro who is working for a cloud services vendor (CSV). This job involves collecting inventory information for the hardware in the data center and occasionally remotely troubleshooting or maintaining the hosts.

These tasks present two challenges:

1. Some of the hardware supports Systems Management Architecture for Server Hardware (SMASH), some supports Intelligent Platform Management Interface (IPMI), and some hardware can use only the vendor proprietary protocols.
2. SMASH and WS-Man can be difficult to understand. The IT pro needs a simple task-oriented way of doing things. The hardware management module in the data center cannot be used due to restrictions on scripts that are downloaded from Internet.

Solution: Windows Server 2012 R2provides one set of Windows PowerShell cmdlets that can be used to manage all of the IPMI and SMASH-based devices.

For developers:

DAL enables developers to invoke a CIM provider that provides a high-level, task-oriented abstraction by using the Windows Management Infrastructure with managed and native client APIs.

For example, a developer on the Cloud Infrastructure Solution team is working on a cloud infrastructure solution deployment tool, which will be used to manage hardware. As per the reference architecture, the hardware can be managed by using IPMI or SMASH for. A list of the typical tasks for managing hardware includes:

- Setting the power state (O or Off)
- Changing the boot order (persistent and one-time)
- Updating firmware
- Fetching event log entries from BMCs
- Providing a hardware inventory
- Creating firmware versions

The developer can use the set of DAL APIs to provide a consistent experience across protocols. These APIs provide a layer of abstraction that completely hides the complexity and details of IPMI- or SMASH-based implementation.

Summary of benefits provided by the DAL:

- Enables device management by using a common abstraction layer that provides standard protocol and schemas
- Helps data centers transition into a defined standards-based management environment
- Provides a ready-to-use solution for device management in Windows, reducing the need for fabric controllers like the Microsoft System Center Virtual Machine Manager (VMM) and CIS to create custom solutions

Note

For information about how properties that are exposed by the DAL are mapped to the values found on hardware devices and how to determine what properties are expected to have empty values when a warning is presented about absent profiles in Windows PowerShell, see Physical Computer System View (PCSV) Property Mappings.

Important functionality

This section provides brief summary information about DAL functionality.

Base management controllers:

The Physical Computer System Profile (PCSP) uses the CIM to implement the basic functionality of the data center management interface . The PCSV is designed to provide a task-based abstraction layer for compute node management.

From the DAL perspective, the PCSP is capable of abstracting different BMC implementation, regardless of what the underlying implementation consists.

TOR

To managethe TOR switch, there is a new set of client APIs and Windows PowerShell cmdlets. It is possible for device vendors to implement the TOR CIM schema in their switches, based on CIM plus WS-Man.

Microsoft has collaborated with CISCO and Arista to implement a newly designed switch profile CIM schema into their switches by using the open-source version of OMI.

See also

This section provides links to information that is relevant to the DAL.

Content type	References
Product evaluation	Datacenter Abstraction Layer Center
Operations	PowerShell Out-of-Band Hardware Management Module In-Band Management using IPMI and PowerShell For details about using the PCSV cmdlets in Windows PowerShell, see Physical Computer System View (PCSV) Cmdlets in Windows PowerShell
Technical Reference	Physical Computer System View (PCSV) Property Mappings provides details about how properties that are exposed by the DAL are mapped to the values found on the hardware devices. This document serves as a reference for hardware manufacturers and developers of management firmware to ensure that their implementation is consistent with the DAL management model. IT pros can use the tables to determine what properties are expected to have empty values when they see a warning about absent profiles in Windows PowerShell.
Tools and settings	Content not available
Community resources	Content not available
Related technologies	Open Management Infrastructure Physical Computer System View Profile (PCSVP)

Content type	References
	Systems Management Architecture for Server Hardware Web Services for Management (WS-Management Specification Software Defined Networking, Enabled in Windows Server 2012 and System Center 2012 SP1, Virtual Machine Manager Windows Management Infrastructure Blog

Physical Computer System View (PCSV) Property Mappings

This topic provides information about how properties that are exposed by the Datacenter Abstraction Layer (DAL) are mapped to the values found on the hardware device itself. The DAL exposes device information by using a Physical Computer System View (PCSV) instance. This document serves as a reference for hardware manufacturers and developers of management firmware to ensure that their implementation is consistent with the DAL management model. Additionally, IT pros can use the following tables to determine what properties are expected to have empty values when they see a warning about absent profiles in Windows PowerShell.

SMASH/DASH property mappings

The following table lists the Systems Management Architecture for Server Hardware (SMASH) and Desktop and Mobile Architecture for System Hardware (DASH) profiles and properties that are used by the DAL to display information about the hardware device. You can use this table to determine which Distributed Management Task Force (DMTF) profiles and specific properties are required in the management firmware implementation.

For devices that conform to the SMASH and DASH standards, the management firmware implementation is required to have a registered Base Server or Base Desktop and Mobile profile to be managed by the DAL. In addition, the following optional profiles must be registered for the DAL to expose the corresponding property values (refer to the following table):

- Boot Control
- Physical Asset
- Software Inventory

If these profile registrations are absent, a client using the DAL to manage your device will not receive an error message that indicates the absence of these profiles, and the corresponding

PCSV properties will have empty values. In certain client implementations (for example, Windows PowerShell) the user will see a warning message that indicates the profiles are not present.

PCSV Property	DMTF Profile Name	Property Mapping	Comment
Dedicated	Base Server/Base Desktop and Mobile	ComputerSystem.Dedicated	No comment
EnabledState	Base Server/Base Desktop and Mobile	ComputerSystem.EnabledState	No comment
OneTimeBootSource	Boot Control	BootSourceSetting.StructuredBootString (where ComputerSystem -> ElementSettingData.IsNext={1|3} -> BootConfigSetting -> OrderedComponent.AssignedSequence=1 -> BSS)	No comment
StructuredBootString	Boot Control	BootSourceSetting.StructuredBootString (where ComputerSystem -> ElementSettingData.IsNext={1|3} -> BootConfigSetting -> OrderedComponent.AssignedSequence!=0 -> BSS)	No comment
HealthState	Computer System	ComputerSystem.HealthState	No comment
IdentifyingDescriptions	Computer System	ComputerSystem.IdentifyingDescriptions	No comment
OperationalStatus	Computer System	ComputerSystem.OperationalStatus	No comment
OtherDedicatedDescriptions	Compute	ComputerSystem.OtherDedicated	No comment

PCSV Property	DMTF Profile Name	Property Mapping	Comment
	r System	Desriptions	
OtherIdentifyingInfo	Compute r System	ComputerSystem.OtherIdentifyingInfo	No comment
RequestedState	Compute r System	ComputerSystem.RequestedState	No comment
InstanceID	N/A	Computed value	Populated by PcsvDevice provider as 0-indexed integer
TargetAddress	N/A	Computed value	The device address given to the provider
FRUInfoSupported	Physical Asset	Chassis -> PhysicalAssetCapabilities.FRUInfoSupported	No comment
Manufacturer	Physical Asset	Chassis.Manufacturer	No comment
Model	Physical Asset	Chassis.Model	No comment
PartNumber	Physical Asset	Chassis.PartNumber	No comment
SerialNumber	Physical Asset	Chassis.SerialNumber	No comment
SKU	Physical Asset	Chassis.SKU	No comment
Tag	Physical Asset	Chassis.Tag	No comment
Version	Physical Asset	Chassis.Version	No comment
CurrentBIOSBuildNumber	Software Inventory	SoftwareIdentity.BuildNumber Classification = 11 "BIOS"	No comment

PCSV Property	DMTF Profile Name	Property Mapping	Comment
CurrentBIOSMajorVersion	Software Inventory	SoftwareIdentity.MajorVersion Classification = 11 "BIOS"	No comment
CurrentBIOSMinorVersion	Software Inventory	SoftwareIdentity.MinorVersion Classification = 11 "BIOS"	No comment
CurrentBIOSRevisionNumber	Software Inventory	SoftwareIdentity.RevisionNumber Classification = 11 "BIOS"	No comment
CurrentBIOSVersionString	Software Inventory	SoftwareIdentity.VersionString Classification = 11 "BIOS"	No comment
CurrentManagementFirmwareBuildNumber	Software Inventory	SoftwareIdentity.BuildNumber Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement
CurrentManagementFirmwareElementName	Software Inventory	SoftwareIdentity.ElementName Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement
CurrentManagementFirmwareMajorVersion	Software Inventory	SoftwareIdentity.MajorVersion Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement
CurrentManagementFirmwareMinorVersion	Software Inventory	SoftwareIdentity.MinorVersion Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement

PCSV Property	DMTF Profile Name	Property Mapping	Comment
CurrentManagementFirmwareRevisionNumber	Software Inventory	SoftwareIdentity.RevisionNumber Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement
CurrentManagementFirmwareVersionString	Software Inventory	SoftwareIdentity.VersionString Classification = 10 "Firmware" (Dedicated = 28 "Management Controller")	The Dedicated value refers to the associated ManagedElement

IPMI property mappings

The following table lists the Intelligent Platform Management Interface (IPMI) commands and their associated return values, which are used by the DAL to populate the corresponding PCSV properties.

Property Name	IPMI Mapping
InstanceID	Computed value
CurrentBIOSVersionString	Get System Info Parameters - Parameter 1 (System Firmware Version)
CurrentManagementFirmwareElementName	OOB MAP Firmware (hardcoded)
CurrentManagementFirmwareMajorVersion	Get Device Id - Firmware Revision 1
CurrentManagementFirmwareMinorVersion	Get Device Id - Firmware Revision 2
CurrentManagementFirmwareVersionString	[MajV]-[MinV]-[Get Device Id - Auxilliary Firmware Revision Info]
Dedicated	28 (hardcoded)
EnabledState	Get Chassis Status - Current Power State
FRUInfoSupported	(Read Fru Data.BoardInfo present | Read Fru

Property Name	IPMI Mapping
	Data.ProductInfo present)
HealthState	0 (hardcoded)
IdentifyingDescriptions	Manufacturer ID-Product ID (hardcoded)
Manufacturer	Read Fru Data - Board Info - Manufacturer
Model	Read Fru Data - Board Info - ProductName
OperationalStatus	0 (hardcoded)
OtherIdentifyingInfo	[Get Device Id - Manufacturer Id]-[Get Device Id - Product Id]
PartNumber	Read Fru Data - BoardInfo - Part Number
RequestedState	12 (hardcoded)
SerialNumber	Read Fru Data - BoardInfo - Serial Number
StructuredBootString	CIM:Network:1, CIM:Hard-Disk:1, CIM:CD/DVD:1, CIM:USB:1, CIM:Floppy:1 (hardcoded)
Version	Read Fru Data - Product Info - Version
TargetAddress	Computed value

Unimplemented properties

A number of PCSV properties are not implemented in the first version of the DAL, which was released withWindows Server 2012 R2. These properties are listed in the following table.

Null Values
Caption
Description
ElementName
LogCurrentNumberOfRecords
LogInstanceID
LogMaxNumberOfRecords
LogOverwritePolicy
LogState

Null Values
MemoryBlockSize
MemoryConsumableBlocks
MemoryNumberOfBlocks
NumberOfProcessorCores
NumberOfProcessors
NumberOfProcessorThreads
NumericSensorBaseUnits
NumericSensorContext
NumericSensorCurrentReading
NumericSensorCurrentState
NumericSensorElementName
NumericSensorEnabledState
NumericSensorHealthState
NumericSensorLowerThresholdCritical
NumericSensorLowerThresholdFatal
NumericSensorLowerThresholdNonCritical
NumericSensorOtherSensorTypeDescription
NumericSensorPrimaryStatus
NumericSensorRateUnits
NumericSensorSensorType
NumericSensorUnitModifier
NumericSensorUpperThresholdCritical
NumericSensorUpperThresholdFatal
NumericSensorUpperThresholdNonCritical
OSEnabledState
OSType
OSVersion
PersistentBootConfigOrder

Null Values
PowerAllocationLimit
PowerUtilizationMode
PowerUtilizationModesSupported
ProcessorCurrentClockSpeed
ProcessorFamily
ProcessorMaxClockSpeed

See also

- Datacenter Abstraction Layer (DAL) Overview
- Physical Computer System View (PCSV) Cmdlets in Windows PowerShell

Background Intelligent Transfer Service (BITS) Overview

This topic provides overview information about BITS, as well as links to other resources pertaining to BITS.

Did you mean…

- Background Intelligent Transfer Service in the Microsoft Dev Center.
- Managed Reference for BITS PowerShell Commands in the Microsoft Dev Center.

BITS description

Background Intelligent Transfer Service (BITS) asynchronously transfers files in the foreground or background, controls the flow of the transfers to preserve the responsiveness of other network applications, and automatically resumes file transfers if a transfer session is interrupted, such as from disconnecting from the network or restarting the computer.

Background transfers are optimal in that BITS uses idle network bandwidth to transfer the files and will increase or decrease the rate at which files are transferred based on the amount of idle network bandwidth available. If a network application begins to consume more bandwidth, BITS decreases its transfer rate to preserve the user's interactive experience.

BITS continues to transfer files after an application exits if the user who initiated the transfer remains logged on and a network connection is maintained. If a connection is lost or if the user logs off, BITS suspends the transfer. BITS persists transfer information while the user is logged

off, across network disconnects, and during computer restarts. However, BITS will not force a connection. When the user logs on again, BITS resumes the user's transfer job.

BITS provides one foreground and three background priority levels that you use to prioritize transfer jobs. Higher priority jobs preempt lower priority jobs. Jobs at the same priority level share transfer time, which prevents a large job from blocking small jobs in the transfer queue. Lower priority jobs do not receive transfer time until all higher priority jobs are complete or in an error state.

BITS uses the Windows BranchCache for peer caching.

BranchCache in Windows 7 and Windows Server 2008 R2 Overview

Practical applications

Use BITS for applications that need to transfer files asynchronously between a client and a server. There are three types of transfer jobs:

1. **Download job** downloads files to the clientAsynchronously transfer files in the foreground or background
2. **Upload job** uploads a file to the server
3. **Upload-reply** job uploads a file to the server and receives a reply file from the server application

Server Manager information

When you install BITS using the Add Roles and Features Wizard, you have the option to install the BITS server Extensions Tools. The tools are required to manage BITS. However the tools to not need to be installed on the same server on which you are installing BITS.

See also

The following table provides links to information that is relevant to BITS.

Content type	References
Operations	Managed Reference for BITS PowerShell Commands Using Windows PowerShell to Create BITS Transfer Jobs Manage BITS (Background Intelligent Transfer Service) with Windows PowerShell Transferring (Large) Files Using BITs
Troubleshooting	Content not available
Getting started	About BITS

Content type	References
Tools and settings	Content not available
Community resources	
Related technologies	BranchCache in Windows 7 and Windows Server 2008 R2 Overview